

Frame Relay + MPLS

By Robert A. Downing, II
Senior Marketing Manager, Larscom Inc.
on behalf of the Frame Relay Forum

What it Was Used For

Frame Relay is a high-speed communications technology that is used today throughout the world to connect LAN, SNA, Internet and even voice applications. In the beginning, however, Frame Relay was developed in response to a clear market demand for a cheaper alternative to leased lines. As networks grew in size, and increasing types of platforms incorporated networking support, it became necessary to not only reduce the cost of networking, but also the sheer numbers of physical circuits required (and envisioned for the future). Developed by and for data communications users, Frame Relay provided and continues to provide the optimal transport technology for data. Simply put, Frame Relay is a way of sending information over a wide area network (WAN) that divides the information into frames or packets. Each frame has a label that the network uses to decide the destination of the frame. Frame Relay can carry multiple network layer protocols (including IP). Because Frame Relay uses a connection-oriented approach, the Frame Relay label or Data Link Connection Identifier (DLCI) becomes a simple reference to a virtual connection.

Frame Relay suits the delivery of traffic with defined service quality, and having longer duration to the streams. This is because network resource allocations can be applied to each connection on the basis of communications between participating switches. On the other hand, where specific network resource reservation is not required, Frame Relay connections can be established with a Committed Information Rate (CIR) = 0. Traffic management is also supported and desirable. This includes the ability to steer traffic along explicit routes in the shared network. Data transfer in Frame Relay operates at layer 2 (or link layer) of the OSI seven-layer model. This has considerable benefits since it is transparent to network layer protocols and simply provides pipes or virtual circuits across a network. Things start to look rather different when we consider Internetworking and layer 3, the network layer.

QoS Solutions Used in Frame Relay Networks

In a shared network, it is necessary to "meter" the traffic entering into and transiting the elements that comprise the network infrastructure. When a connection is made to a Frame Relay network, characteristics are established that provide for a Committed Information Rate, a Maximum Burst Rate, and Burst values ('commit' and 'exceed') that govern the proportion of Frame Relay resources consumed by traffic over that connection. The Committed Information Rate (CIR) specifies the projected demand on the network for transport resources. It reflects a moving average value, which is approached over time using the Burst-commit (Bc) and Burst-exceed (Be) rates. The combination of values allows for the efficient use of the network from both Provider and a Customer perspective. This provided a mechanism for efficient allocation of resources, but as Frame networks enjoyed increasing subscription, an enforcement mechanism was required in the event of saturation.

Carried within the Frame Relay header is a mechanism that allows for adaptive treatment of traffic by both CPE and switching equipment in the event of congestion. Designated the 'Discard Eligibility Indicator,' or 'DE-bit,' it is used to designate individual frames as candidates for elimination to reduce congestion in the network. Implemented at ingress to the network by the CPE, the bit allows prioritization to be applied at the end-network by its administration, giving preference to mission-critical traffic. Applied at the switch, the provider has the ability to provision on a demand-basis across multiple end-networks according to priorities set by the subscriber.

In support of this quality mechanism, additional field mappings were implemented in the Frame Relay header defining switch-to-switch signaling mechanisms in support of over-all network congestion control. To better coordinate the discard mechanism at ingress switches, the switches communicate using Backward- and Forward-Explicit Congestion Control.

While the implementation of Discard logic is initially made using information having local significance (local congestion and traffic marked DE), the switch has the option of indicating a congested state on frames entering the transport network proper. As this congestion information is accumulated within the infrastructure, the network has the ability to indicate in a backward manner that congested state to the ingress switch points, providing network state that can be used in conjunction with the local information to further impose congestion control. It should be noted that absence of these indicators will serve as relaxation mechanisms within the network as well.

When the Nature of Traffic Changed...

The Internet Protocol (IP) and networks based on IP have shown enormous growth. This is as true for corporate internets/extranets as for the public Internet. IP is, however, a connectionless network layer protocol, as opposed to the connection-oriented Frame network. In a connectionless network, a packet travels from one router to the next, each router deciding how to forward that packet. Routers analyze each packet and route it independently of other packets. In this case routing is performed many times. This form of networking has consequences for quality of service and traffic engineering. Quality of service is impacted because no pre-established path exists on which to allocate resources. Traffic-engineering capabilities are less than ideal because IP routing gives only coarse control over explicit routes. Congestion control imposed at the route level between participating network switches begins to lose effectiveness.

The addition of Voice and Video Constant-Bit-Rate (or real time Variable-Bit-Rate) traffic to the network served to further complicate architected mechanisms to provide preferential treatment to traffic streams transiting the Frame network. Voice and Video connections have characteristics in common that minimize the established mechanisms of provisioned Bit Rate and Discard Eligible metering imposed by network elements. More transient in nature, and generally of shorter duration, rtVBR traffic can only be minimally metered using mechanisms dependent upon a sliding-window logic as has been implemented in Frame Relay. It is possible for the (necessary) accumulation and then sharing of congestion information between participating switches to outlast a significant proportion of this newer type of traffic stream.

What this Involved on the Needed Level of QoS, and How MPLS can Satisfy this Need?

What was required, then, was a means of imposing preferential treatments on these short-lived, non-deterministic connections that would occur outside of the Layer 2 network of Frame Relay yet contribute to the overall metering of the infrastructure. As visibility to the route determined for the traffic streams became clouded from the perspective of the Frame network elements, responsibility for the Quality of Service for those connections migrated beyond the Frame network to the Layer 3 switching elements. Given the real time requirements of these types of traffic, connection setup and provisioning at the higher layer is being done by the Layer 3 switch immediately adjacent to the originating (and terminating) host. It is only logical, then, that the QoS determination made by these switching elements be communicated to, and coordinated with, the switches responsible for the Layer 2 topology of the network.

Multiprotocol Label Switching (MPLS) is a development by the Internet community. It seeks to combine the flexibility of the IP network layer with the benefits conferred by a connection-oriented approach to networking. MPLS, like Frame Relay, is a label-switched system that can carry multiple network layer protocols. Similar to Frame Relay, MPLS sends information over a wide area network (WAN) in frames or packets. Each frame/packet is labeled and the network uses the label to decide the destination of the frame. In an MPLS network we can define explicit paths or let IP routing decide the path. MPLS networks can use Frame Relay, ATM and PPP as the link layer. A key feature is to separate network control and data forwarding. This makes MPLS extensible to many environments including SDH and Optical networks.

In essence, MPLS sets out to address requirements for:

- efficient and simplified high speed forwarding of IP packets.
- provision of scalable networks.
- control over quality of service (QOS).
- Traffic Engineering and the control of traffic routing.

Description of the Transport of Frames on a MPLS Network

MPLS is a generic, label-switched network solution. It uses the principle first established in ISDN, of separating control from data transfer. This architecture enables MPLS to use a range of different forwarding methods. For example, forwarding engines (switches) can operate using Frame Relay or ATM protocols. In this respect, we can use existing hardware for the forwarding function in an MPLS network. What is new in MPLS is that the IETF has standardized the control functions so that we can establish a path or connection across whatever forwarding engines we like.

MPLS, like Frame Relay or ATM, adopts a connection-oriented approach. In fact, if a Frame Relay or ATM network uses a standardized IP routing protocol (e.g. OSPF) and a standardized protocol to communicate DLCI (label) values, we will have built something that looks like an MPLS network. So, there is nothing that new in MPLS.

To provide different values for QoS or Class of Service (CoS) we set up paths across an MPLS network and allocate resources to these paths. This may be done today in Frame Relay networks to support QoS on PVCs. In fact, the network features (e.g. different queues and queue scheduling) are the same for Frame Relay and MPLS. So, MPLS does not of itself provide QoS but by using a connection-oriented approach it facilitates QoS support. From a standardization viewpoint, MPLS can use the model developed for Integrated Services (Intserv) where RSVP (Resource Reservation Protocol) makes reservations. MPLS can also support the IETF's model of Differentiated Services (Diff-Serv) to ensure that each type of traffic receives an appropriate class of forwarding treatment (Gold, Silver, Bronze etc.).

It is evident, then, that the most effective implementation between two label-switching transports would be that allowing for the interchange of label information between transports, for inclusion in the native network header.

When using a layer 2-based approach, an MPLS core network transports Frame Relay frames from ingress to egress. That is, we take frames from a Frame Relay access network, deliver them to the ingress of an MPLS network, transport them across the MPLS network and deliver frames to another Frame Relay network. Simply stated, this is "frames in and frames out". This is likely to be a popular and powerful solution since it addresses a number of issues of current concern. For example, it allows service providers to migrate their core network to MPLS while protecting Frame Relay revenue streams; there is no requirement to modify, interfere with or upgrade the customer's equipment; it easily supports Virtual Private Networking by supporting multiprotocol traffic and by hiding private IP addressing schemes from the core transport network; it provides the data security associated with Frame Relay services; the VPN implementation is controlled from the CPE (the customer or service provider may manage this).

How does MPLS help? It helps by standardizing routing and control across multiple layer 2 technologies. In turn this lets us use layer 2 switches for forwarding while integrating IP routing and control with these switches. This way we can realize the high speed of switching and eliminate unnecessary duplication of control and routing.

Conversely, when the Frame Relay network is the core architecture, the "tag" information from the MPLS label at ingress to the Frame network can serve as the Frame Relay header value for the DLCI. It would not be sufficient, however, to only modify the DLCI in the header. There would also be a 'shim' header inserted to preserve the Layer 3 information like Time-To-Live and other LSP values.

Enhancements Needed?

Essentially, there are no enhancements or modifications required beyond an instance of the Frame Relay/MPLS Interworking Implementation Agreement (issued jointly by the Frame Relay <http://www.frforum.com/> and MPLS Industry Standards Forums <http://www.mplsforum.org/>). The development on the part of the equipment vendors is analogous to that done in the effort to implement Frame Relay to ATM Interworking following the Implementation Agreements from those respective

Forums. The platform would be required to support MPLS Label-Switched Paths, and Frame Relay (Private or Switched) Virtual Circuits at the same time. This is, however, fairly commonplace in today's multiprotocol routers.

While it is a given that networks, the traffic they carry, and the standards governing that transport will change, the current work between the respective Forums has gone a long way toward integrating the two complimentary transports. It is evident that a mechanism that allows for the transfer of label information between switching processes, avoiding processing intensive and thus latency generating inspection of information deeper in the frame, allows for the exchange of Quality of Service determinations while increasing throughput of network platforms.

Acknowledgements go to Harman Hopkins, authors of a Frame Relay/MPLS White Paper for the Frame Relay Forum, Kendra Gross of Porchivina Associates, and the Frame Relay Forum. Thanks also to Larscom, Incorporated, for providing me the opportunity to participate in the Forum, and indulge in such literary exercise.

In his 20 years of IT industry experience, Robert A. Downing has served in a number of roles including systems engineering, business development, and product management. Rob is currently a senior marketing manager at Larscom Incorporated where he has overall marketing responsibility for the Larscom (<http://www.larscom.com>) 6000 family of service delivery platform products. He also serves as participating architect for future product development. Rob has also been elected to the Frame Relay Forum (<http://www.frforum.com/>) Board of Directors as co-Vice President of Marketing.