Frame Relay + MPLS

And how they fit together

F R A M E

P

R E L A Y

S

Acknowledgements

The following have made substantial and gratefully received contributions to the improvement of this white paper: Bernard da Costa (Bell Canada), Michael Walsh (Nortel Networks).

February 2001

Prepared for The Frame Relay Forum

by Harman H Hopkins

Accent on Networks

Table of Contents

Synopsis

Examining the opportunities for Frame Relay and Multiprotocol Label Switching (MPLS).

The role of IP in today's networks continues to expand and grow in importance. However, Frame Relay also continues to generate significant revenue in service provider networks. As Frame Relay networks are often used to carry IP traffic, there is tremendous discussion about how both IP and Frame Relay networks will evolve to new technologies and architectures. One current hot topic is the role of MPLS in the deployment of next generation networks. This paper looks at the similarities and differences between Frame Relay and MPLS, and how they can operate together. We examine the reasons behind the development of MPLS and discuss likely scenarios in which both Frame Relay and MPLS have important roles to play.

Overview

Frame Relay and MPLS. Do they fit together? In this paper we examine the complementary aspects of these two approaches. We show how Frame Relay switches could become MPLS switches and, importantly, how we can position Frame Relay to be an attractive means of access to an MPLS core network. In following this trail, we shall expose the need for interworking between Frame Relay and MPLS.

Frame Relay

Frame Relay is a high-speed communications technology used throughout the world to connect LAN, SNA, Internet and even voice applications. From the beginning, users embraced Frame Relay enthusiastically because it was developed in response to a clear market need, namely the need for a cheaper alternative to leased lines. Developed by and for data communications users, Frame Relay provided and continues to provide the right technology at the right time.

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 DLCI becomes a simple reference to a virtual connection.

Frame Relay suits the delivery of traffic with defined service quality. This is because network resource allocations can be applied to each connection. On the other hand, where specific resource reservation is not required, Frame Relay connections can be established with a Committed Information Rate (CIR) = 0. Traffic management is also supportable and desirable. This includes the ability to steer traffic along explicit routes

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.

Internets and the Internet Protocol

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. In a connectionless network, a packet travels from one router to the next, each router deciding how to forward that packet. Routers analyse 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 and traffic engineering capabilities are less than ideal because IP routing gives only coarse control over explicit routes.

MPLS

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 labelled 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.

The Development of MPLS

Before we investigate how MPLS works, let us see where it came from and why the IETF developed it. MPLS has its roots in a number of proprietary approaches including:

- Cisco's Tag Switching
- IBM's Aggregate Route-based IP Switching (ARIS)
- IP Switching (Ipsilon, now part of Nokia)
- IP Navigator (Lucent/Ascend)

All of these approaches had the goal of producing efficient and scalable IP networks and although each approach differed in detail, each contributed in its own way to the formulation of MPLS.

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.

Let us look briefly at each of these.

Simplified forwarding

In a traditional IP network, a router switches (forwards) packets from an input interface to an output interface. However, that is not its only function. A router updates routing information as well. To forward packets, it has to examine the IP packet header of every packet. Furthermore, routers often support multiple protocols and interface types. The two jobs, forwarding and routing, are different from each other. Forwarding is how the router transfers data. Routing is a control function defining where the router transfers the data.

Switches, as distinct from routers, perform fewer functions. Optimised for the task of forwarding, switches reduce processing and as a result, are faster. Examples are Frame Relay or ATM switches that forward data based on simple label lookup procedures. On the other hand, Frame Relay and ATM switches use their own proprietary or standardised routing and control processes that duplicate IP routing and control functions. As an example, ATM switches commonly use the ATM Forum's PNNI (Private Network to Network Interface) protocols to determine routes through an ATM network. Or, to use another example, some Frame Relay switches use standard Internet routing protocols such as OSPF.

How does MPLS help? It helps by standardising 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 realise the high speed of switching and eliminate unnecessary duplication of control and routing.

Scalable networks

There are many aspects to the design of scalable networks. Here, we look at the N^2 problem and routing adjacency.

Frame Relay networks enjoy enormous success in supporting corporate networks. A typical network might connect several branch offices to a headquarters site using permanent virtual circuits, resulting in a star-based topology. This fits the typical organisation structure since most communication is between branches and HQ. In fact, the vast majority of networks deployed today are based on this model, with little hard evidence to show that there is a real need for any-to-any connectivity.

However, of late, propelled by the ubiquity of the Internet and by new e-business models, network designers are turning their attention to the future, with the expectation that networks may have to support any-to-any connectivity. While a Frame Relay network can provide this, it results in a full mesh of permanent connections. This full mesh requires $n(n–1)/2$ connections. This is the N^2 problem, i.e. the connections and their management grow as N^2.

The other issue is routing adjacency. In the mesh network described above, we directly connect all the access routers and create routing adjacencies between them. Therefore routing traffic can be very large and can grow at a rate exceeding N^2.

The Interior Gateway Protocol Problem

All the edge routers are peers. This results in a non-scalable network.

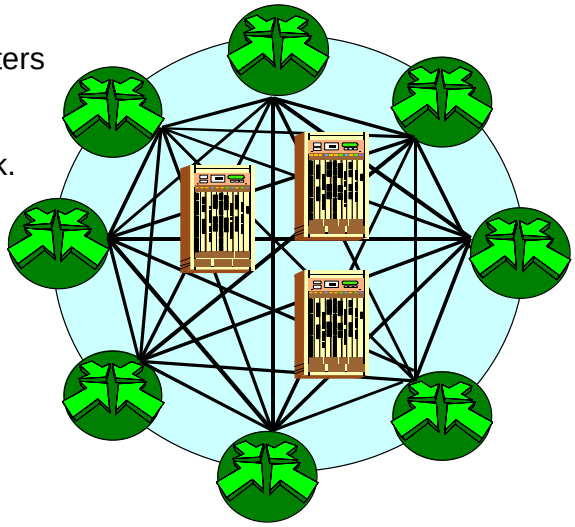

What is the solution? Well, we can eliminate the mesh of routing adjacencies if our switches support IP routing.

An "IP Routed" solution

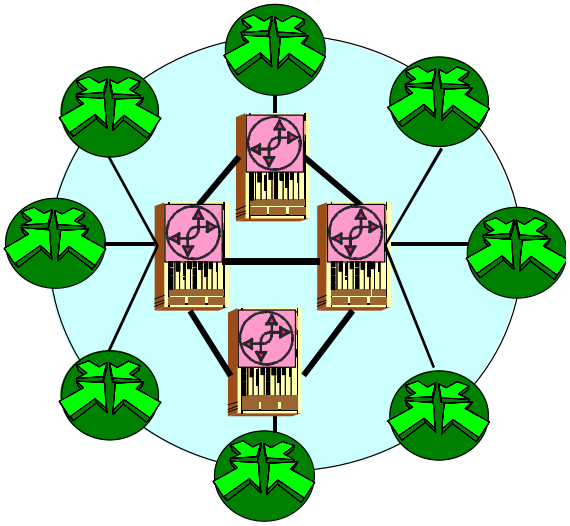

In this case the layer 2 switches participate in IP routing, which relieves the edge routers of N^2 routing adjacencies. Data forwarding makes use of the underlying switch.

Quality of service

To achieve a defined quality of service we need to allocate network resources. This ensures that when we forward data, it is done in a way that meets a QoS objective. Frame Relay and ATM services already provide such features at layer 2; MPLS promises to do the same for IP Internetworks. We also must recognise that different applications require different treatment of their data. Voice and video require low delay and delay variation, but may tolerate occasional packet loss. Other data is more tolerant of delay but requires reliable end-to-end transport.

MPLS helps us in meeting both requirements. We shall see later that MPLS, like Frame Relay or ATM, adopts a connection-oriented approach. In fact, if a Frame Relay or ATM network uses a standardised IP routing protocol (e.g. OSPF) and a standardised 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 standardisation 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.).

Traffic engineering

The efficient use of network resources requires control. This control is about how we ensure that the bandwidth available in a network is well used. This, in turn, means that we need information about available resources, e.g. links and the means to direct traffic over those links. In MPLS we achieve this by setting up explicit routes. This is in contrast to normal IP routing based on the shortest path. We can establish explicit routes manually or by a routing algorithm that takes account of a set of constraints. These constraints normally include the bandwidth required for the path. A path calculation selects only paths that can satisfy the constraints. MPLS paths are called Label Switched Paths or LSPs.

MPLS: how it works

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 standardised the control functions so that we can establish a path or connection across whatever forwarding engines we like.

MPLS forwarding

An example will help us to grasp the essentials of MPLS forwarding. Let us trace how an IP packet arriving at the ingress of an MPLS network is transported to the egress of the network. The sequence followed is:

1. The IP packet enters at the ingress to the MPLS network.
2. The packet is assigned to a path and a label attached. This process first classifies the packet and then adds the label. In fact all of the packets that fall into the same classification get the same label. More formally we say a packet is assigned to a Forwarding Equivalence Class or FEC. We will see an explanation of FEC shortly.
3. The labelled packet is sent to the next MPLS node.
4. This node looks at the label - **the IP header is not examined.**
5. The next hop is chosen by reference to a label forwarding table. This table has entries for the incoming interface and label value and corresponding entries for the output interface and the outgoing label value. Thus, the table entries may determine that a packet arriving on (say) interface 1 with label value (say) x will be switched to an output interface (say) 7 with a label value of (say) u.
6. The new label is written and the packet sent on its way to the next MPLS node.
7. This process continues until the packet reaches the last MPLS node (egress).
8. The label is stripped (popped). This may expose another label or an IP header. In the latter case, the packet is delivered to the final destination using standard IP procedures.

Readers familiar with Frame Relay will see that this is analogous to how Frame Relay works. To see this consider the following:

1. An IP packet arrives at the FR CPE (router)
2. The router encapsulates the IP packet in a frame and adds the DLCI (label)
3. The frame is sent to the Frame Relay network (ingress node)
4. This node looks at the DLCI (label) – the IP header is not examined
5. The next hop is chosen by reference to a DLCI (label) forwarding table
6. A new DLCI (label) is written and the frame sent on its way to the next FR node
7. This process continues until the frame reaches the remote CPE (router)
8. The DLCI (label) is stripped (decapsulated) and the IP packet delivered to the final destination

Returning to MPLS, the example may give rise to a number of questions:

In (2), how was the packet assigned to a particular path?

> *We define a Forwarding Equivalence Class or FEC. This could be based on the packet's IP destination address but may also take other information (e.g. the interface that the packet arrived on, Type Of Service (TOS) or Diff-Serv marking). If we use the IP address, then, we assign all packets with the same destination address to the same path. Since the label is a reference to the path, all these packets have the same label value.*

The fifth step (5) assumes that the node has formed a label forwarding table. How do we achieve this?

> *The population of the label values in the label forwarding table is handled by a label distribution protocol. There are two main choices, RSVP-TE (the TE stands for traffic engineering) or LDP (the Label Distribution Protocol) and its traffic engineering enhancement CR-LDP (Constraint-based Label Distribution Protocol).*

The virtual path through the network is defined by entries in label forwarding tables in each node. The nodes are called Label Switched Routers or LSRs. But what determines the physical route for the path. In other words, which LSRs are transited?

> *The physical path can be determined by standard IP routing procedures. So when we set up a path using a label distribution protocol, it will follow the same route that an unlabelled IP packet would take. Remember we are discussing a signalling or control process that sets up the LSP. Data can only be forwarded after the path is established. Alternatively, the physical path may be explicitly defined for traffic engineering purposes. In this case the route may be defined manually or by constraint-based routing.*

MPLS Labels

We have talked a lot about labels without defining what they look like. Because MPLS is designed to use different link layers, the label format will reflect the characteristics of the link layer used. For example, where the link layer protocol has no convenient field for an MPLS label, the MPLS label is inserted (or shimmed) between the link and network

Generic MPLS Shim Label

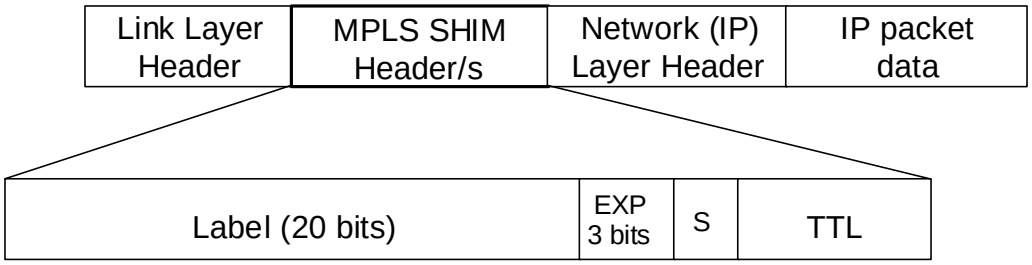

layer headers. This gives rise to the term "shim header". A shim header would be used where the link layer is PPP or Ethernet and it is illustrated in the figure above. Apart from a field for the label value, the shim header contains a Time To Live (TTL) field. This field is decremented by Label Switched Routers along the Label Switched Path (LSP), so that if a packet enters a routing loop its TTL will eventually expire and the packet will be discarded. The MPLS architecture allows a hierarchy of labels. This is similar to the use of VPI and VCI labels in an ATM network. However, ATM only has a two-level label structure, but in MPLS the number of shim headers is unlimited. Even if multiple shim headers exist, only the top header (the one next to the link layer header) is used for forwarding. If the current top-level label is removed (popped), it may expose another shim header, which then becomes the new top-level label. Because multiple shim headers may exist, the field marked S (one bit long) is used to indicate if this shim header is the last of a stack of headers. (There are several situations where using more than one shim header is useful). Finally, the EXP (or Experimental) bits can be used to mark packets for different forwarding treatment, perhaps based on the incoming IP packet's Diff-Serv marking.

If our link layer uses Frame Relay, then MPLS makes use of the existing Frame Relay label field (DLCI). In other words, the MPLS label value is written to the FR DLCI field.

Frame Relay Encapsulation

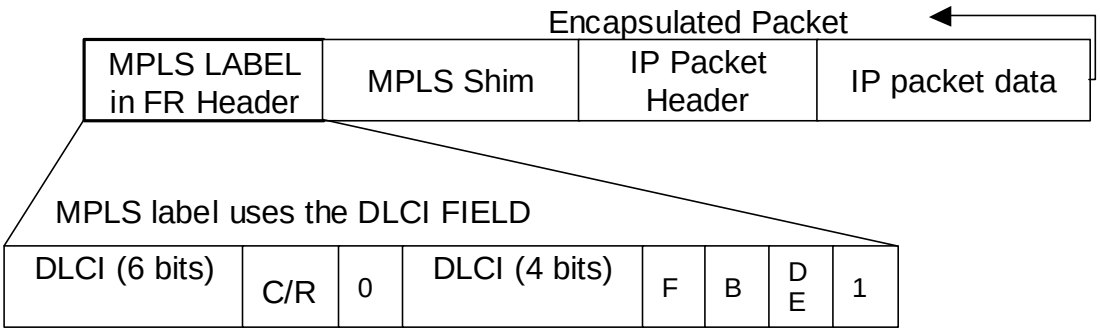

- The active or top level label is carried in the DLCI field
- Two or four octet addresses can be used
- A shim header is used to support TTL
- Shim is invisible to FR (part of FR data field)

Notice that although the MPLS label uses the DLCI field, we also have inserted a shim header after the FR header and before the IP header. This enables the processing of Time To Live (TTL) as we shall see later.

Label Distribution

There are a number of different ways in which we can distribute MPLS labels. One common method is called ‘Downstream on Demand’ and it is shown in the figure below.

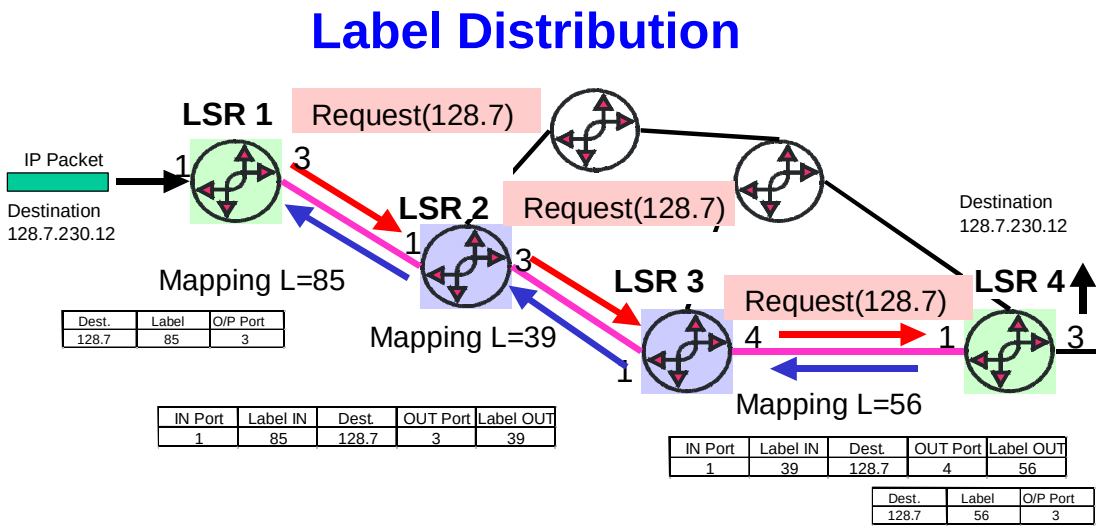

In the diagram, an IP packet arrives at LSR1 (Ingress LSR). This packet must then be mapped to an LSP. To do this, each LSP has a Forwarding Equivalence Class (FEC) associated with it and the FEC identifies the set of IP packets to be mapped to the particular LSP. If there is an LSP with a Host address FEC element identical to the packet’s destination address, the packet is mapped to that LSP. (If there are multiple matching LSPs, one LSP is selected).

Each LSR requests a label binding; this message flows downstream from the ingress (LSR1) via LSR2 and LSR3, to the egress (LSR4). The path taken by the request message will be obtained from IP routing (e.g. OSPF) as the next hop for the destination address or from an explicit list of LSRs. The actual label values to be used are decided by each LSR. Label allocations start at the egress (LSR4) and work back towards the ingress. For example, the figure shows that LSR4 decided to allocate a label value of 56; this is communicated upstream to LSR 3. LSR3 binds this label (56) to the output port (4), thus filling in the label forwarding tables as shown. This process is repeated until all the forwarding tables are completed. We have now set up a unidirectional path from ingress to egress. A full-duplex connection requires a second setup in the opposite direction.

Frame Relay +MPLS

We now turn to look at how Frame Relay and MPLS fit together. The more general, but perhaps less likely scenario is to use Frame Relay switches as Label Switched Routers (LSRs).

Frame Relay-based MPLS

Since Frame Relay is a label switching system, using Frame Relay as one of the underlying layer 2 technologies for supporting MPLS is reasonable. Then, Frame Relay switches become Label Switched Routers (LSRs).

When a Frame Relay switch is an LSR, we carry the top-level MPLS label in the Frame Relay DLCI field. As we have seen, by adding a label stack or shim header we can include other information such as TTL. Frame Relay switches have no knowledge of this since they act only on the DLCI field and the shim header is part of the data carried inside the Frame Relay information field. Furthermore, the particular characteristics of Frame Relay addresses have to be taken into account. Since the DLCI field can be 10 or 23 bits long, there must be a way of identifying this so that label values can be correctly assigned.

Another factor to be considered is the issue of decrementing the Time-To-Live (TTL). The MPLS shim header contains a TTL field that can be decremented by each LSR. However, Frame Relay switches do not have access to this field and thus cannot process the TTL value. Instead, the TTL value is decremented at the ingress by the number of hops in the LSP before the packet is encapsulated in Frame Relay. The label distribution protocol is used to communicate the actual hop count value to the ingress.

This scenario is entirely feasible and is illustrated below. This shows what happens when an IP packet enters an MPLS network that uses FR-based LSRs. The second diagram illustrates the situation at the egress.

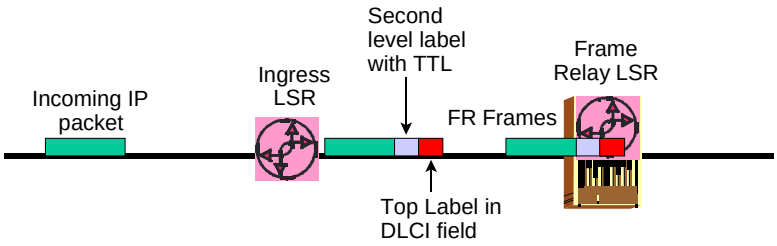

- An IP packet arrives at the ingress LSR
- The packet is classified to a Forwarding Equivalence Class (FEC)
- A label stack with at least one entry is added
- The TTL field in the top label is filled with the IP TTL value
- This TTL value is reduced by the number of hops in the MPLS segment
- The label value for the FEC and therefore the Label Switched Path is written to the frame relay DLCI field
- The labelled packet is passed to the frame relay forwarding layer.
- The frame relay frame is forwarded to the next FR-LSR by consulting the label forwarding table

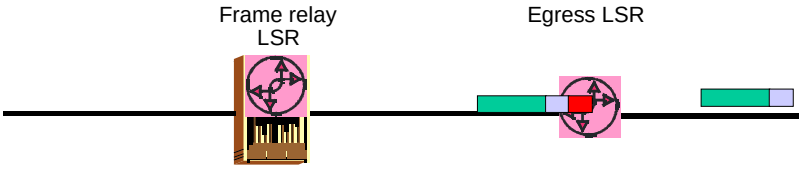

- Frame arrives at the Egress LSR
- The current (top) label is removed (popped)
- If it is the last label, the network layer is inferred from the label value
- If it is not the end of the LSP the packet is forwarded according to the type of link for the next hop (e.g. ATM, PPP)
- The TTL is decremented
- The packet is sent to the destination (or next hop)

Perhaps now is a good time to pause to take stock of the situation. We could ask, how likely is it that Frame Relay switches will be upgraded to perform LSR functions as illustrated? Looking at current networks, it is a fact that most service providers' networks support Frame Relay services using ATM in the core. Currently, MPLS deployment is aimed at core or backbone networks. This makes the large-scale deployment of Frame Relay-LSRs less likely. Instead, we are much more likely to see ATM or router-based MPLS core networks with Frame Relay used as a highly appropriate access to these MPLS core networks.

Frame Relay access to MPLS

As we have already said, Frame Relay has enjoyed great success and is used by many organisations. On the other hand, service providers faced with the problem of building scalable internetworks, are migrating to MPLS in their backbone networks. These MPLS backbones may be built using router technology or ATM-based MPLS. What we need is a strategy for the evolution of Frame Relay services that takes MPLS into account.

As it happens, Frame Relay has already tackled a similar situation, i.e. for interworking between Frame Relay and ATM. The need for Frame Relay to ATM interworking became inescapable as service providers migrated to ATM core backbones. Two scenarios emerged, network interworking and service interworking. Similarly, as these core networks embrace MPLS, Frame Relay will again have to adapt to the new situation. Used as an access method, Frame Relay will enjoy the benefits of the new backbone. This in turn will require new definitions for network and service interworking, matters to which we now turn our attention.

Layer 2-based solutions

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.
- Easily supports Virtual Private Networking by supporting multiprotocol traffic and by hiding private IP addressing schemes from the core transport network
- 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.

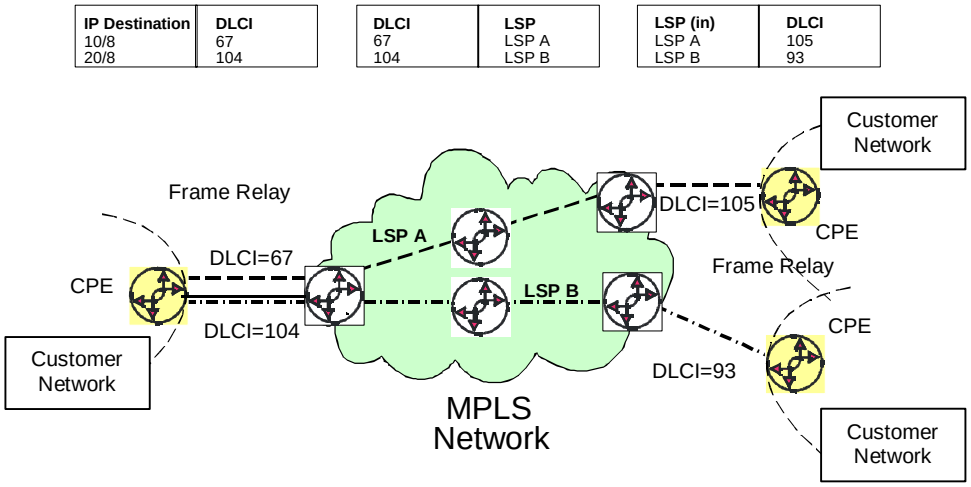

In the figure above, we see an example of tunnelling Frame Relay across an MPLS network. Assume that the Frame Relay CPE has a routing table that relates the destination IP address to a DLCI. An IP packet arriving at the CPE will be encapsulated in a Frame Relay frame, with a DLCI required for the PVC to the IP destination. (In the example above, DLCI = 67 is associated with a destination in the network at the top right). The frame is transported across to the ingress LSR, which then maps the incoming DLCI value (67) to Label Switched Path (LSP A). Frames are tunnelled across the MPLS network to the Egress LSR. The Egress LSR then maps them to another DLCI value for delivery to the destination.

This scenario could also be used to link Frame Relay service providers over an MPLS core network (See Figure below). In this case, the MPLS network maintains a mesh of Label Switched Paths connecting each ingress and egress Label Switched Router.

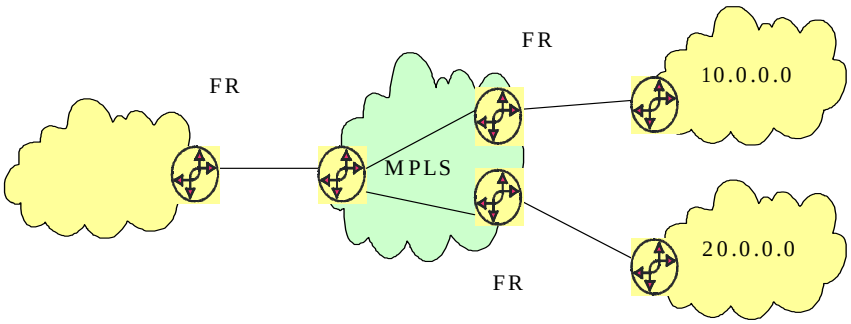

This is, however, somewhat of a halfway house approach since no layer 3 lookup is done at the ingress to the MPLS network. That leaves us with a question. Is it possible to construct layer 3 VPNs using MPLS? The answer is decidedly yes. We shall return to this issue later.

Before we move on, we should look at the pros and cons of this solution. On the benefit side, it is a simple way of supporting existing Frame Relay services. It provides traffic aggregation onto LSP trunks since each trunk carries all the transit traffic to a given destination network. On the other hand, there is an implied level of manual configuration. This includes the configuration of routing tables in the CPE and DLCI to Label lookup tables in the ingress/egress LSRs. Finally, we do not have the integration with IP routing that a layer 3 solution would provide. The resilience of this solution also needs consideration. Although both the Frame Relay and MPLS networks may support alternative backup paths for use if the primary path fails, there is no obvious mechanism for coordination between the networks.

Layer 2-based approach: Advantages and Disadvantages

To summarise, the Layer 2-based approach provides the following advantages:

- Support for VPNs with CPE control over routing.
- Retains existing customer service and CPE.
- Supports existing Frame Relay customer and service provider investments.
- Provides Multiprotocol support.
- Enables the interconnection of VPNs that use private IP addresses
- Supports the inherent security associated with Frame Relay networking.
- Enables the service provider to migrate to an MPLS backbone network.
- The MPLS backbone can use a range of underlying technologies including router-based LSRs, ATM-LSRs, FR-LSRs and in future, MPLS-controlled optical or TDM switches.
- MPLS Label Switched Path (LSPs) establishment uses dynamic signalling (LDP or RSVP-TE).
- QoS can be supported on the Frame Relay connections with service classes and priorities. The MPLS network can map to similar features using resource reservation and/or Diff-Serv and constraint-based routing.
- The MPLS network can optimise network utilisation with in-built traffic engineering features.

Some drawbacks to this approach are:

- Scalability: May require manual configuration and, for full mesh connectivity, suffers from the problems of N^2 router adjacency and number of connections (PVCs and LSPs) of the overlay approach.

So perhaps we should look at solutions that take account of the network layer.

Layer 3-based scenarios

A layer 3-based solution requires that the ingress and egress to the MPLS network do IP lookups. Therefore, if we are transporting IP packets across a Frame Relay access connection, the ingress LSR will unpack the IP from the Frame Relay frame, do an IP lookup and assign the packet to a Forwarding Equivalence Class or FEC. The ingress LSR will contain a mapping from FEC to a label value and will add this label to the IP packet. Using MPLS forwarding, the labelled packet transits the MPLS network and arrives at the egress LSR where the label is removed. An IP lookup is performed by the egress LSR and the packet encapsulated in a Frame Relay frame for delivery to the destination.

Enabling layer 3 lookup procedures at the ingress provides several advantages. For a start, the number of routing adjacencies is reduced since CPE are not peers but are peered with the ingress LSR. This moves the meshing requirement away from the CPE. Instead, we need to provide mesh connectivity between the ingress/egress LSRs. By operating at the network layer we can also examine quality of service markings in the incoming IP packet and ensure that the packets are classified to a suitable FEC and hence a Label Switched Path that supports the required QOS. The QOS marking could be in an IP Type Of Service (TOS) field or a Diff-Serve code point.

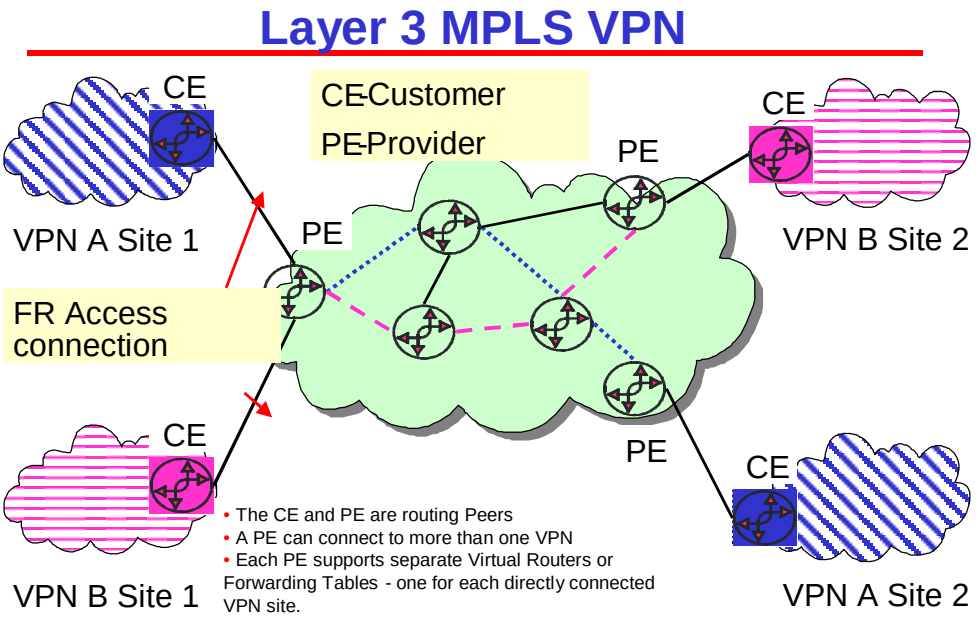

A layer 3-based approach to support Virtual Private Networking is illustrated below. Notice that we are still using Frame Relay as the access network. Clearly, this solution is appropriate when the VPNs are using IP as the network layer protocol. However, VPNs often use other protocols; indications are that as much as 40% of Frame Relay network traffic may be non-IP. What are we to do? Well, we can encapsulate “the other” protocols in IP. But keep in mind this may not work too well, especially for protocols like SNA.

Layer 3-based approach: Advantages and Disadvantages

Some advantages of the Layer 3 approach include:

- Scalability for any-to-any connections to VPN sites.
- Meshing in the core network is the responsibility of the service provider (CPE not involved).
- Provides a fully routed IP network solution.
- Only the Provider Edge (PE) devices know about VPN routing. The core devices (LSRs) do not see any VPN specific routing.

Some disadvantages are:

- It's an IP-centric solution and may not apply to all multiprotocol VPN situations.
- As yet, no IETF standards for MPLS VPNs exist. The documents describing MPLS VPN scenarios are only informational. The result is that current solutions tend to be vendor specific.
- Some solutions require modification to standard routing protocols in order to carry VPN specific information.

Enhancing FR to MPLS Interworking

As we have said, a number of network providers are either planning to or changing their backbone networks to MPLS. Perhaps now is a good time to think about how to ensure that the two technologies can interwork in a smooth manner. In the first instance we should address the way in which Frame Relay PVCs are used in conjunction with an MPLS core network. In particular, are any new Frame Relay Implementation Agreements or Standards required? For example:

A definition for FR to MPLS Network interworking might include a number of requirements such as:

- Mapping fault conditions from the MPLS side to the FR side.
- Automatic configuration of FR PVCs to Label Switched Paths.
- Multiprotocol encapsulation procedures
- Mapping of FR Service Class/Priority to MPLS equivalents.

Although of less importance due to the low demand for SVCs, we could also define how MPLS and Frame Relay SVCs interwork. This requires signalling conversion between FR signalling and the label distribution protocols used to establish MPLS Label Switched Paths.

Conclusion

In this paper, we have looked at the basics of MPLS. In so doing we have discovered that MPLS and Frame Relay have certain affinities. Both are examples of label switching and indeed, Frame Relay switches can be upgraded to become MPLS Label Switched Routers (LSRs).

In this paper we have touched on some likely scenarios in which both Frame Relay and MPLS have important roles to play. Frame Relay services continue to generate significant and growing revenues while service providers are increasingly looking to MPLS to provide the basis of their next-generation core networks. These facts demonstrate the importance of providing continued support for the customers of Frame Relay services while migrating the core networks to MPLS. The provision of this support reinforces the need for Frame Relay access to MPLS core networks. The general utility of Frame Relay, together with its large installed customer base, is a clear indicator of the need to define and optimise interworking between the two technologies.