Frame Relay over MPLS (FRoMPLS): Today, Tomorrow and the Future

Mike Walsh
President, Razor Sharp Consulting
President Emeritus, Frame Relay Forum

On behalf of the Frame Relay Forum

The author wishes to acknowledge the contributions of the FRF leadership, members and associates in the review of this document.

The Frame Relay Forum is an association of vendors, carriers, users and consultants committed to the education, promotion, and implementation of Frame Relay in accordance with international standards.

Learn more about the Frame Relay Forum at:
http://www.frforum.com/

Frame Relay Forum
www.frforum.com

Table of Contents

Introduction

Today's telecom marketplace is perhaps more interesting than ever. The fast and wild days of unbridled optimism in the 1990's have passed, but the job of providing for growing data communication needs remains. It is at this critical moment in time that we find ourselves increasingly considering the relationship of the established technology of frame relay to the emerging technology of MPLS.

Often this discussion is framed in terms of a contest, frame relay vs. MPLS. Similarities have even been drawn to the old arguments of connection-less (IP) vs. connection-oriented (frame relay, ATM) technologies with the two sides represented by the "Bell-heads" on one side and the "Net-heads" on the other. However, as this paper will make clear, frame relay and MPLS are closely related (not opposed) in terms of operation and objectives. Moreover, much like family relatives, there are more similarities between generations than differences.

Through this paper, we will explore the backgrounds of frame relay and MPLS, the advantages of each, how they can combine synergistically, current FRoMPLS implementation shortcomings, the evolution of FRoMPLS in standards work, and how this pairing of technologies will provide significant advantage.

This paper is sponsored by the Frame Relay Forum, an association of vendors, carriers, users and consultants committed to the education, promotion, and implementation of Frame Relay in accordance with international standards. Through work with the MPLS Forum and IETF, the Frame Relay Forum is one of the leaders for standardization, implementation, and rollout of effective FRoMPLS equipment and services.

Frame Relay: Over 10 years old, and still going strong

Frame relay, which represents the elder portion of the relationship, is a highly successful technology that started in 1991 with a mere 3 carriers, 30 enterprise users and a total market revenue of $1.7M, yet has grown to be a $12B+ worldwide market with a predicted total market of $21B in 2004 (source: Vertical Systems Group). Frame relay is used throughout the world for IP traffic, LAN interconnect, SNA, Internet access and even voice, among other uses.

While predictions have been made over the years that frame relay will be eclipsed by various technologies, in the end frame relay has proven to be a remarkably useful and evolving technology that has successfully fulfilled customer need. Even in the midst of the networking downturn, frame relay has survived and prospered. Why this is so deserves some discussion.

When frame relay was first developed back in 1991, it was created as a cheaper alternative to the prevalent leased lines of the day. Frame relay not only reduced cost, it also reduced the number of physical circuits and related equipment through virtualization of connections. Virtualization of connections means that all traffic will essentially follow a predetermined route through the network – looking like a private line service, except that the physical circuits are shared between multiple virtual connections. This simplicity also results in very fast and efficient switching at connection points in the network.

Because most data traffic is bursty in nature, statistical multiplexing is used to increase the efficiency of the data network, and user-specified traffic parameters are used to police traffic to insure that those network resources are not overloaded. The combination allows service providers to combine lower prices with strong service level agreements (SLAs) that the end user could count on. The end result is more efficient use of resources, lower prices, better customer guarantees and faster order fulfillment because the circuits are virtual instead of physical.

Finally, because frame relay is a layer 2 (link layer) protocol it can carry any type of layer 3 (network) protocol desired: IP, SNA, bisync, voice, etc. In addition, since frame relay leverages the high performance of today's transport networks and the resiliency power of upper layer protocols, it provides this performance with a minimum of overhead.

Frame relay has evolved over the years for increased performance, more applications, higher speeds and a number of other improvements. These basic starting point advantages provide the foundation upon which the success of frame relay was built. Today frame relay is available throughout the world providing outstanding data network connectivity at very competitive prices.

For those desiring further technical or other types of information, please refer to the excellent resources on the Frame Relay Forum website.

MPLS: The new kid on the block. Or is it?

Around 1997 interest began in finding ways of bringing some of the advantages of layer 2 technologies to IP networks. Increasingly IP was becoming the layer 3 protocol of choice for applications and, at least in theory, IP networks were thought to be cheaper from both a capital expenditure and operational expense perspective. This effort resulted in the development of a protocol now known as multiprotocol label switching (MPLS).

One of the first objectives behind MPLS was to speed up routing operations. Because routing was originally a slow and resource-intensive operation that had to be performed at each decision point in the network, an idea was advanced to determine fixed paths through a network to a specific endpoint and label the path with an identifier. Thus, a packet entering the network could have one routing path determination made, have a label put on it representing that fixed path, and then be quickly and efficiently label-switched through the network. If this sounds to you a bit like frame relay with label switching according to virtual circuit identifiers, but with IP protocols used for routing, you are not far off.

By 2000, router technology had caught up and eliminated the need for this 'fast-routing' aspect of MPLS. However, in the process of the work it was discovered that the fixed path nature of MPLS brought many of the advantages of frame relay to IP networks, most importantly traffic engineering and quality of service (QOS). One of the great challenges of pure IP networks is that its connection-less nature makes it very difficult to control network resources and ensure that traffic moves through the network without loss or delay. In fact, over half the traffic carried by frame relay networks is IP, in large part due to the need for QOS guarantees.

So, while losing its original purpose, MPLS gained a more important one: traffic engineering for QOS assurance. As the industry has shown, one can charge substantially more for a service that is predictable and controllable as opposed to one that is uncontrollable and unpredictable. An illustration of this principle would be business ATM service versus a residential broadband service. In addition, during the development of the MPLS protocol, many other tremendously important evolutionary ideas were incorporated:

- **Separation of control and data transfer planes**: Many providers today use ATM layer 2 technology inside their networks in order to provide QOS guarantees between IP routing points. While this works, it requires managing two types of infrastructure and two types of routing: IP and ATM. MPLS helps the network provider by using IP protocols for all control (routing) decisions, while also being able to use also any type of data transfer technology: frame relay, ATM, PPP, optical and others.

- **Increased integration with IP**: As layer 2 switches become MPLS-enabled, they begin working like IP devices. And as IP-routers become MPLS-enabled, they begin taking on the advantages of layer 2 technologies. With routing based on IP, and data transfer technology agnosticism, MPLS could provide an effective evolutionary path for IP.

- **Speed/size advantages**: While many MPLS networks will initially be built using existing ATM cell switching for transport, most ATM offerings will top out at OC-48, or perhaps at best OC-192. MPLS is built to go as fast as underlying technologies can go and in a world where data networking demands are going no direction but up, this is a key advantage.

- **Recovery and resiliency**: Understanding that many of the quick recovery techniques built into layer 2 networks hold advantage over the slower to converge re-routing capabilities in IP routing protocols, MPLS incorporates advanced techniques to provide resiliency and recovery.

- **Network architecture planning and provisioning**: One of the most useful aspects of MPLS has to do with how networks can be organized. While ATM provided the ability to collect a number of virtual circuits into a virtual path (2 levels of organization), MPLS provides unlimited ability to design and combine backbone tunnels at many points throughout the network infrastructure (unlimited levels of organization). Imagine this as similar to how you might make a car journey using neighborhood roads, local roads, regional roads, state roads and interstate roads. MPLS allows this exact same concept of merging traffic flows onto and off appropriately directed and sized paths. That is a major advantage for network planners.

For those desiring further technical or other types of information, please refer to the excellent resources on the MPLS Forum website.

In the end analysis, the thinking and work that has gone into MPLS has resulted in a protocol that takes advantage of the power of IP protocols and the low cost of IP equipment while delivering the traffic engineering and service level assurance of layer 2 protocols like frame relay and ATM.

So, we now return to the beginning of the paper: is it a case of MPLS vs. frame relay, or one of MPLS and frame relay? It seems clear from the analysis that it is indeed the latter: a unique opportunity to extend the outstanding success of frame relay using a well thought out evolution of backbone transport technology called MPLS.

How frame relay is being used today

To understand better how frame relay can leverage the advantages of MPLS, it is helpful to consider how frame relay is being used today.

As we previously discussed in the background section above, frame relay is used throughout the world by businesses of every size for a large number of applications including IP, SNA and voice, among other uses.

A distinction that is not always clear however is that frame relay is used both in terms of a service, and as an access method to other services.

Frame Relay as a Service

Today the most frequent use of frame relay is when a customer buys a frame relay networking service. A common example is of a customer with a large headquarters and several smaller branch locations. A frame relay virtual circuit will be purchased between each branch and the headquarters and, if traffic levels justify, occasionally links between certain branch offices. At each end of the frame relay link there is a device such as a router or a frame relay access device (FRAD) that terminates the frame relay link and connects into the office’s computing infrastructure.

This has been the classic and most popular use of frame relay, and fits the way many of the businesses of today perform their work and position their computing resources. Depending upon the information technology expertise of the business, the terminating equipment could be owned and maintained by the business, or purchased as part of a managed service from the service provider with the frame relay transport. For many business users, the use of frame relay has been an evolution from private line services configured in very much the same topology.

Frame Relay as Access

While the most popular use today for frame relay is as a service, one cannot overlook the growing role of frame relay as an access technology to other services. Frame relay's penetration, acceptance and reliability as a link layer technology has made it the obvious choice as the access layer into many of the most popular services of today.

Frame Relay Internet Access

Internet access for B2B, website, email, customer service and other applications has become critical to business customers in recent years. It should be no surprise that reliable frame relay is frequently chosen as the link layer access technology to get to Internet access services, often with the purchasing business hardly being aware that frame relay is being used.

Private IP Services

Another popular service offered by providers has been private IP services. Essentially the service is offered as an IP network service to the customer, and implemented with frame relay links to the customer's sites connecting into dedicated IP routers (physically hosted at the provider's site) that are connected to each other via a dedicated backbone. This is somewhat similar to a customer building their own IP network based on their routers on premises connected by frame relay links, but offloads the storage, maintenance and updating of the routers to the service provider.

IP-VPN Services

One of the most popular new services in recent years has been the emerging IP-VPN services. Much like the private IP services above with a single frame relay connection from each customer site into the service provider network, the new IP-VPN services leverage new techniques and technologies to 'virtualize' the routing function. In an interesting parallel to the advantages that frame relay brought with the virtualization of circuits, virtualization of the routing functions has allowed significant savings in routing and transport infrastructures, resulting in a high-performance IP networking solution at an outstanding price, all with the access link reliability brought with frame relay. As a major bonus, internal IP routing enables easy any-to-any branch connectivity, an option that is often more difficult and expensive to do with traditional frame relay service when it is required. Interestingly, one of the more popular approaches has been to use MPLS labeled paths as a means of separating different customer's traffic during passage through the IP network – yet another example of MPLS's usefulness.

Frame over DSL (FRoDSL)

DSL is known as a technology with the potential for reducing network access cost, and some providers are offering DSL as the access technology for frame relay service instead of leased lines or direct dial access. This is yet another example of how frame relay can be provided and used in a variety of ways.

Frame Relay Service over MPLS

Having provided some background on frame relay and MPLS technologies and uses, we will now turn out attention back to the current efforts to synergize the established success of frame relay service with the advantages afforded by the use of MPLS.

Why Pursue Frame Relay Service?

While frame relay has proven remarkably adaptable to a variety of uses including access to other services, today by far the most popular use of frame relay is still as a layer 2 networking service. These services represent a substantial portion of the revenues received today by service providers, and a vital means of reliable data networking for global businesses. The popularity of frame relay as a service providing circuit-level connectivity is expected to continue to be strong in the near term for the following reasons:

- **Businesses like it and are comfortable with it**: a reason that cannot be ignored. For many customers, frame relay service continues to work well for them and they have no compelling business reason to switch.
- **It matches network configurations**: many businesses today remain in the hub-and-spoke configuration for their data networking with the computing servers and other resources maintained at the headquarters.
- **Strong IT department**: customers that have the ability to install and maintain their own infrastructure are in a position to shop for only the data networking transport component.

- **Quality of service**: frame relay services have a strong technology basis and reputation for providing excellent quality and reliability of service that can be enforced through measurable SLAs.

The astute reader will notice that many of these reasons correspond to those used for selecting private lines for data networking transport, both in the past and currently. With the knowledge that today private line usage remains higher than even frame relay, it would appear clear that the use of frame relay as a layer 2 networking solution will not fade any time soon either, and will mostly likely be popular for many years to come. Based on both the current situation and the prognosis for the future, evaluation of ways to improve performance and efficiency seems prudent, if not vital.

What value does MPLS bring to FR?

Previously we have discussed the advantages of MPLS technology, but how specifically does this translate to value for frame relay providers and customers?

The following are from our list presented earlier, explained specifically for frame relay:

- **Traffic management and QOS**: The strong capabilities of frame relay in this area are similarly supported in MPLS. This allows customers to purchase the same type of high performance and reliability they have had with frame relay services in the past.

- **Separation of control and data transfer planes**: MPLS can use the existing ATM technology backbone cell switching that is probably already in place for transporting frame relay traffic. This means a service provider can avoid expensive forklift upgrades when offering FRoMPLS.

- **Increased integration with IP resources**: The ability to build and expand networks around new MPLS-enabled IP routers provides two excellent benefits: (1) lower risk of obsolescence while still maintaining current frame relay revenues and (2) the ability to converge networks onto a common MPLS-enabled backbone.

- **Speed/size advantages**: Because MPLS is built to go as fast as underlying technologies can go, as an underlying technology for frame relay services it can be used to build efficient networks as well as provide easy expansion when necessary.

- **Recovery and resiliency**: MPLS incorporates advanced techniques to provide the resiliency and recovery that are the hallmark of a good frame relay service and allows those customer expectations to be maintained with confidence.

- **Network architecture planning and provisioning**: MPLS's facility to provide unlimited ability to design and combine backbone tunnels at many points throughout the network infrastructure will allow service providers to more easily and cost effectively manage their frame relay networks.

Current Implementations of FRoMPLS: Issues and Concerns

As has been discussed, MPLS brings great advantage to the offering of frame relay services, and as one would expect, many vendors and service providers have quickly moved to implement some initial FRoMPLS services. The vast majority of these implementations are based on the IETF Martini draft specifications that described how frame relay could be encapsulated in MPLS. Already great claims are being made for the advantages provided, and excitement is high.

Recent investigations and network rollouts have revealed that while these initial implementations provide many of the advantages of frame relay, they do not provide all of the expected behaviors. Some industry experts have gone so far as to suggest that these current FRoMPLS implementations are sufficiently different from familiar frame relay service that they should have any entirely different name such as 'frame forwarding'.

The concerns reside in two specific behaviors of current FRoMPLS implementations: lack of packet order guarantees, and lack of prompt circuit loss or congestion notification. We'll discuss each of these issues in a bit more detail.

> **Lack of packet order guarantees**: Frame relay services provide what is referred to as fast packet relay. In exchange for fast packet transport and low overhead, frame relay doesn't provide extensive buffering and recovery procedures. On the rare occasion that a frame is found to be corrupted, it is quickly dropped and the next frame is evaluated. This is actually a good behavior in that many services such as voice or streaming UDP traffic can better tolerate loss of a packet than delays of many packets due to recovery time. In addition, higher-level protocols like TCP can successfully recover from a dropped packet condition. However, a major characteristic of circuit-oriented technologies is that the packets delivered successfully will be in order and many upper layer protocols count on this behavior. In particular, SNA continues to be a popular protocol transmitted on frame relay and out of order packets will cause SNA to reinitialize and create an absolute fault at the application layer. This is obviously not desired behavior, and packet misordering is a possible remnant of the connectionless basis of many MPLS implementations unless something else is done. It is estimated that 15-25% of frame relay traffic could be affected by this lack of packet order guarantee on some current FRoMPLS implementations.

Lack of prompt circuit loss or network congestion notification: Another set of objections raised is that with some current FRoMPLS services it is possible to experience circuit loss or congestion without proper notification to the end user equipment. In the case of circuit loss, it can be up to half a minute before higher layer protocols report the loss to applications (if at all), and if your networking termination equipment depends on being notified of network congestion there may be network stability problems. Again, as above, we should fully consider the impact of these issues. On the subject of circuit loss, most MPLS implementations include advanced rerouting and resiliency techniques, much like the frame relay networks in place today. Regarding congestion notification, reaction to congestion has always been a tricky subject, and as a result, many pieces of equipment are configured to not even react to the notification. Still, it seems clear that because these two are assumed behaviors from a frame service, it is best to build these into the standards and service definitions to be sure.

Some arguments can be made that in a number of situations certain upper layer protocols and other conditions can minimize the problems in the current generation of FRoMPLS implementations. However, this seems to miss the point that one of the great success factors of frame relay to date is that it is upper layer protocol agnostic and provides certain guarantees that data networks were built around. As a result, to the end user it is a simple and easy to use data transport service. Through a little bit of work, we can evolve FRoMPLS to that level.

With that in mind, let's turn to a discussion of the current initiatives to improve FRoMPLS to represent better the capabilities of frame relay combined with the potential of MPLS.

The Next Evolution for FRoMPLS: The Work of the Frame Relay Forum, MPLS Forum and IETF

The following sections will now explore the work that is being done to evolve FRoMPLS to a more advanced behavior that more closely approximates native frame relay services. This work is being pursued within three separate bodies, all working together: the IETF's PWE3 working group, the Frame Relay Forum and the MPLS Forum.

Background on Standards vs. Implementation Agreements

It may be of use to describe the differences between standards, written by groups like the IETF's PWE3, and implementation agreements, written by forums like the FRF and the MPLSF. A standard lays out the basic mechanism for how a proposed technology will work, but leaves room for interpretation via directives such as "should", "recommend", "may" and "optional" to be considered by each implementer. An implementation agreement (IA) is based on a standards document and composed by a forum of implementing member companies and users. An IA further details the technology by specifying behavior that might have been optional in the standard, clarifies areas in the standard that might have been vague, may add more requirements, and on occasion may correct errors found after a standard is issued. In some cases, an IA may even originate independently without reliance on an existing standard.

Industry experience has shown that the combination of standards work and implementation agreements has been an effective means of achieving faster technology acceptance, quicker network rollouts, decreased field problems and greater vendor interoperability. The FRF is proud to be playing a leading role in the development of a FRoMPLS implementation agreement.

The Building Blocks: How the FRoMPLS Implementation Agreement Relates to the PWE3 Draft Specification

The FRF/MPLSF FRoMPLS implementation agreement can be seen as an overall inclusive document that references and details the PWE3 draft specification, while also referencing the various existing frame relay implementation agreements and standards from FRF, ITU-T and IETF. As described above, the combination of a solid standard with a similarly solid wrapping implementation agreement has proven to be a very powerful and effective method of moving technologies through the market to users.

The FRoMPLS Reference Model

The following descriptions of the basic FRoMPLS reference model and operation are extracted from the PWE3 draft specification and the FRF/MPLSF FRoMPLS implementation agreement in order to provide a good background.

A frame relay and MPLS network interworking reference model is shown below. It consists of the following elements:

– An MPLS core network.
– Provider Edge (PE) devices providing network interworking functions between frame relay and MPLS. PEs can support frame relay UNIs to Customer Edge (CE) devices and/or NNIs.
– Frame relay devices and networks interconnecting PEs with frame relay UNIs and/or NNIs.

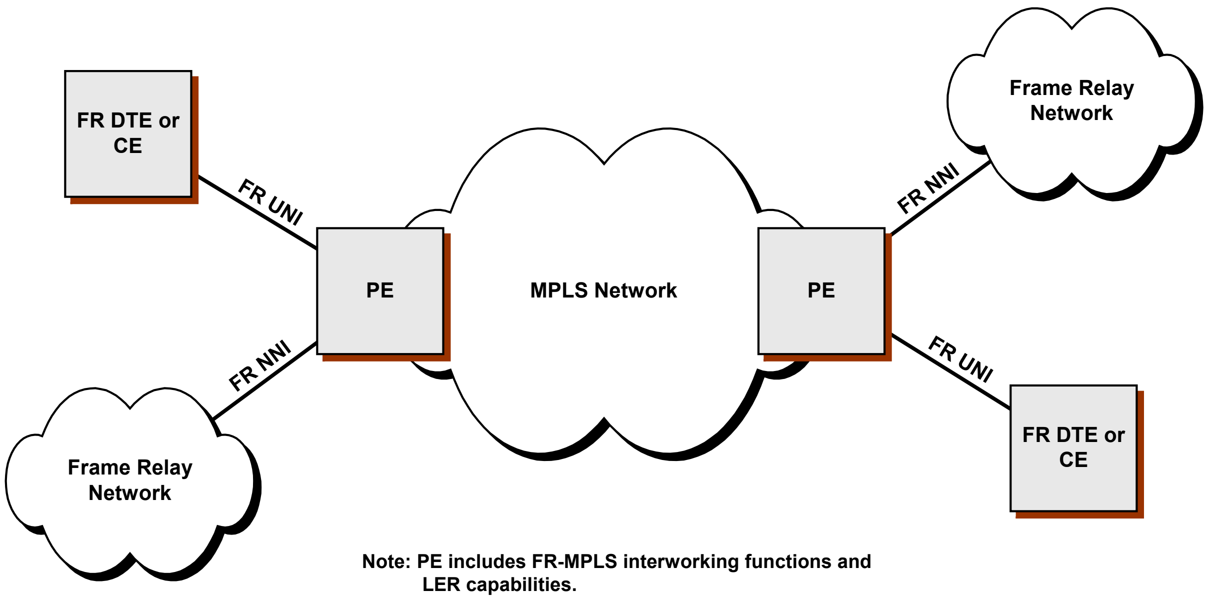

FR-MPLS network interworking connects two frame relay networks and/or devices using an MPLS network. In this architecture, frame relay networks and devices act as CE entities according to the PWE3 architecture. The frame relay service is first provisioned between each frame relay device and the corresponding PE device. A Virtual Connection Label Switched Path (VC LSP) is then set up between the two PEs to complete the frame relay Virtual Connection (VC).

The basic transport mechanism at the PE is as follows:

1. Encapsulation of frame relay specific information in a suitable frame relay over pseudo wire (FRoPW) packet,
2. Transfer of a FRoPW packet across a PSN for delivery to a peer PE,
3. Extraction of frame relay specific information from a FRoPW packet by the remote edge node,
4. Generation of native frame relay frames for forwarding across an egress port of the remote edge node,
5. Execution of any other operations required to support frame relay service.

The use of the MPLS network by two frame relay networks and/or devices is not visible to the end users. The end user protocol suites remain intact. The PE provides all mapping and encapsulation functions necessary to ensure that the service provided to the frame relay networks and/or devices is unchanged by the presence of an MPLS transport network. This is also referred to as frame relay transport over MPLS.

For further information on the technical details of this protocol, please refer to the many excellent resources available at the IETF, MPLS Forum and Frame Relay Forum websites.

What is the PWE3 Draft Specification?

The acronym PWE3 stands for pseudo-wire emulation edge-to-edge (essentially circuit emulation), and this provides a good starting point for understanding the objectives of the PWE3 draft specification work. This effort takes as its starting point the previously mentioned Martini draft work, and incorporates the contributions of many individuals and companies to provide a standard that can support as faithfully as possible frame relay services over a MPLS network.

The main frame relay pseudo-wire requirements to be met by a PE as expressed in the terms of the specification are:

MUSTS:

1. Support of bidirectional traffic via paired MPLS LSPs to match frame relay's two-way virtual circuits
2. Frame relay traffic control information must be transported (Discard Eligibility (DE), Forward Explicit Congestion Notification (FECN), Backward Explicit Congestion Notification (BECN) and Command/Response (C/R) bits)
3. Frame relay circuit status indicators must be mapped and transported.
4. Frame relay permanent virtual circuits must be supported

SHOULD:

1. Frame order should be preserved (*very strongly recommended*)
2. Variable length frame relay frames should be supported regardless of underlying MPLS technology
3. Frame relay traffic management and QOS parameters should be mapped and used to establish appropriately engineered MPLS network paths.
4. PVC link integrity check should be provided.

OPTIONAL:

1. Frame relay SVC and SPVC support is optional.

ASSUMPTIONS:

1. Transmission errors will be detected by the underlying link layer, not at the PW protocol layer.

Evaluation of the requirements imposed by the PWE3 draft specification show a very complete and well thought out set that strongly addresses many of the shortcomings of current FRoMPLS implementations. This includes requiring transport of circuit control information including congestion notification, communication of circuit status and a strong recommendation to support preservation of frame order. The PWE3 draft specification provides an outstanding foundation for an implementation agreement.

What is the Forthcoming FRF/MPLSF FRoMPLS Implementation Agreement?

The Frame Relay Forum and MPLS Forum FRoMPLS implementation agreement builds upon the work of the IETF PWE3 draft specification with implementation details and recommendations for future work.

The IA first identifies agreement with the MUST provisions of the PWE3 work, and strengthens the requirement for frame order preservation to the MUST level. The combined requirements and directives of the PWE3 draft specification and the FRoMPLS IA create a service offering that properly carries frame relay service traffic in the manner that customers are used to, and with PWE3 addresses the major shortcomings in current FRoMPLS implementations.

Any IA or standard has room for enhancements that will result in improved customer performance, and future study has been identified for work in the following areas:

- The FRoMPLS service will be required to transport variable length frame relay frames regardless of the underlying MPLS technology maximum transfer unit (MTU) size, relieving the customer end from having to perform fragmentation when necessary.
- The status of the connection will be mapped and transported as appropriate through the entire circuit, and the ability to perform a continuity check should be provided.
- Customer-specified traffic characteristics such as committed information rate, burst sizes, and maximum frame size will be mapped appropriately and used to establish suitably engineered MPLS tunnel LSPs. Similarly, the QOS values associated with a frame circuit (service classes, frame transfer/discard priorities) will be mapped and used to establish suitably engineered MPLS tunnel LSPs.

It should be noted that these items identified for future study involve areas of network equipment design and architecture that are often initially differing between vendor products. For example, there are a large number of possibilities for how different levels of frame relay QOS could be mapped to different implementations of MPLS tunnels. Areas like these are the ones that benefit the most through forum activity to specify in detail how these requirements should be implemented.

As discussed before, implementation agreements are a critical addition to the standards process as they facilitate faster technology acceptance, quicker network rollouts, decreased field problems and greater vendor interoperability. The FRoMPLS IA is an excellent addition.

Summary

This paper has described the background of frame relay and MPLS, their respective benefits, current implementation shortcomings, and the on-going standards/IA efforts of the FRF, MPLSF and IETF organizations to evolve FRoMPLS into a highly successful service offering.

It is clear that frame relay today provides enormous benefit to customers all around the world. Equally clear is that MPLS is an important technology that can reduce cost and improve network performance. The combination of these two technologies into a high performance data transport service predicts a great future for both.

The work undertaken by the Frame Relay Forum, MPLS Forum and IETF will extend the astounding frame relay service success of the past 10 years through a future enhanced with MPLS technology

For More Information on the FRF, MPLSF and IETF

The reader is encouraged to consult these organizations web sites for further information, including how to become a member and contribute to the efforts.

Frame Relay Forum: http://www.frforum.com/
MPLS Forum: http://www.mplsforum.org/
IETF: http://www.ietf.org/

Key References

1. Frame Relay and MPLS network interworking Straw Ballot text
2. draft-martini-frame-encap-mpls-01.txt

About the author:

Mike Walsh is the president and founder of Razor Sharp Consulting. As a product strategy and marketing expert, Mr. Walsh brings extensive international and domestic experience in the data communications and telecommunications industries to his clients.

Working with companies from small startups to large corporations in full-time and consulting roles, Mr. Walsh has built outstanding knowledge of the entire product cycle through activity in the areas of sales, *business development, marketing, training, product management, research and development, and manufacturing operations.*

Mr. Walsh has also worked with a number of industry forums and standards bodies including the Frame Relay Forum, ATM Forum, MPLS Forum, ITU and the IETF. He currently serves as President Emeritus of the Frame Relay Forum, where he previously worked as President and Chairman of the Board of Directors, and Chair of Market Development and Education.

Frequently consulted as an industry expert, Mr. Walsh has been interviewed by such leading magazines as Network World, Communications Week International Business Communications Review, and Telecommunications International. He has also spoken at major domestic and international trade shows including VPNcon, NetWorld+Interop, a number of private conferences, and many Web broadcast sessions.

Mr. Walsh earned advanced Master degrees in business administration and technical disciplines from North Carolina State University, and a Bachelors degree from Clarkson University. He currently serves on the Board of Directors of several non-profit organizations, and has contributed his skills to a large number of volunteer organizations.

Mr. Walsh can be reached via email at mikew@razorsharpconsulting.com

Glossary

Customer Edge: A Customer Edge (CE) is the customer device connected to a provider edge device.

Provider edge: A Provider Edge (PE) is a network edge device that provides a frame relay service over an MPLS network.

Label Switched Path: A Label Switched Path (LSP) is the path through one or more MPLS nodes at one level of the hierarchy over which packets in a particular forwarding equivalence class (FEC) are transmitted.

MPLS node: An MPLS node is a device that is aware of MPLS control protocols, will operate one or more layer three routing protocols and will be capable of forwarding packets based on LSP labels.

PseudoWire: A pseudowire (PW) is a connection between two PEs carried over an MPLS network.

PseudoWire Emulation Edge-to-Edge: pseudowire emulation edge-to-edge (PWE3) is a mechanism that emulates the essential attributes of a frame relay service (and of some other services not covered in this document) over an MPLS network.

B_c	Committed Burst size
B_e	Excess Burst size
BECN	Backward Explicit Congestion Notification

BOM	Beginning Of Message
CE	Customer Edge
CIR	Committed Information Rate
CPE	Customer Premises Equipment
C/R	Command / Response indicator
DE	Discard Eligibility
DLCI	Data Link Connection Identifier
DTE	Data Terminal Equipment
EOM	End Of Message
FEC	Forwarding Equivalence Class
FECN	Forward Explicit Congestion Notification
FR	Frame Relay
HDLC	High-level Data Link Control
IA	Implementation Agreement
IETF	Internet Engineering Task Force
ISDN	Integrated Services Digital Network
ITU-T	International Telecommunications Union-Telecommunications
LSP	Label Switched Path
LSR	Label Switch Router
MPLS	Multi Protocol Label Switching
MTU	Maximum Transfer Unit
NNI	Network-to-Network Interface
PDU	Packet Data Unit
PE	Provider Edge
PHP	Penultimate Hop Popping
POS	Packet Over SONET/SDH
PPP	Point to Point Protocol
PT	Payload Type
PVC	Permanent Virtual Connection
PW	Pseudo-Wire
PWE3	Pseudo-Wire Emulation Edge-to-Edge
QoS	Quality of Service
RFC	Request For Comments
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
SVC	Switched Virtual Connection

UNI	User-to-Network Interface
VC	Virtual Circuit / Virtual Connection
VoFR	Voice over Frame Relay