130
TECHNOLOGY WHITE PAPER

Frame Relay & Ethernet Coexistence:
Supporting Legacy and New Services with a Single MPLS Infrastructure

Timothy Wu, Riverstone Networks

ABSTRACT *Why deploy Multiprotocol Label Switching (MPLS) in metropolitan area networks?* This paper introduces an important reason: the integrated support of existing Frame Relay deployments with new Ethernet services.

To date, much attention has been focused on MPLS' role in traffic engineering and creating private business networks, or VPNs. The third role — integrating Frame Relay and Ethernet services — may prove equally significant over the long run. The goal of this paper is to explain how MPLS supports Ethernet/Frame Relay integration, and why this is an important advantage for service providers.

The significance of Frame Relay/Ethernet integration must be understood in economic, not technological, terms. Frame Relay (and by this we mean Frame services backed by ATM aggregation) represents the mainstream of modern business data networking, with nearly 2 million deployed ports and approximately $13 billion in global revenue. Metro Ethernet, on the other hand, is the consensus choice for the next generation of business networking, based primarily on its revolutionary bandwidth and easy compatibility with LAN Ethernet. This places a heavy focus on the technology that can bridge the gap between today's Frame Relay and tomorrow's Ethernet.

The apparent answer is an increasingly popular MPLS tunneling technology commonly referred to as "Martini tunneling." Over the last year, the Martini Draft has emerged as the industry standard for tunneling various Layer 2 protocols over an MPLS network. Riverstone Networks has implemented the Martini approach in the RS router family.

In this paper, we focus on the use of Martini tunneling to support both Frame Relay and Ethernet services in metropolitan networks. Part I of the paper discusses the state of Frame Relay deployments worldwide. Part II explains how Ethernet and MPLS have become part of the metro landscape. Part III details how MPLS can be used to carry both Frame Relay and Ethernet traffic over a combined Ethernet/MPLS infrastructure. Part IV provides greater detail on L2 MPLS VPN technology, while Part V provides a practical guide to deploying a multi-service MPLS metro network.

River STONE NETWORKS

PART I: THE CONTINUING SIGNIFICANCE OF FRAME RELAY ACCESS AND ATM AGGREGATION

Today, Frame Relay remains the "workhorse" of private business networking. A mature yet still-growing technology, it may be justifiably hailed as a success, and is laudable for its focus on end-user needs. In this section we conclude that Frame Relay's continuing importance in data networking means that service providers designing new Metro network infrastructures should take Frame Relay access services into account.

Some have asked: "How long will Frame Relay continue to be relevant?" To answer this question, we need to focus on two facts. First, Frame Relay is primarily an access technology today. In public networks, sheer scale makes the access side of the network much more expensive to change, as compared with aggregation portions. The endurance of the local loop is testament to this point.

Second, people like Frame Relay. Its popularity with end-users comes from its wide availability, relatively low cost, simplicity, and easy familiarity. For carriers and service providers, Frame Relay remains an important source of revenue and represents an investment in both equipment and operational expertise.

For that reason, we should expect Frame Relay access networks to be a factor in metro networks for years to come. The Yankee Group and others have projected that Frame Relay will not only remain relevant, but also continue to grow through the mid-2000s.

So exactly how big is Frame Relay? Since Frame Relay services were launched in 1991, market revenue worldwide has grown from $1.7 billion to an estimated $12.7 billion in 2001; of that, approximately $7 billion is in the United States, according to Vertical Systems Group. This year, customer installations of carrier-based Frame Relay services will reach 1.78 million ports worldwide, with more than a million service ports in the U.S.

Frame Relay is now used principally for LAN interconnection and Internet access. However, Frame Relay's multiprotocol features also continue to be important to organizations that have yet to transition to an all-IP environment, — the technology is used for carrying IBM SNA, bisync, and other traffic types. The Frame Relay Forum estimates that up to 40 percent of the traffic on current Frame Relay networks is non-IP.

Frame Relay's future appears secure, at least as a relatively low-bandwidth (45 Meg and below) access technology. But we must return to the point made earlier. Even senior members of the Frame Relay Forum acknowledge that Frame Relay's future is as an access technology, and not for aggregation. What about the aggregation side?

This is where ATM enters the picture. ATM/Frame Relay interoperability has been key to scaling Frame Relay deployments, and today ATM is typically used to aggregate Frame Relay traffic. Most users of public ATM services are actually public Frame Relay end-users who use ATM in locations where higher speeds are needed than Frame Relay can accommodate.

It is ATM's aggregation role in metro networks that is under competitive pressure from the technologies discussed here. While carriers and service providers need to sustain their current revenue streams, they must also position themselves to exploit emerging and service trends within the metro to remain competitive. As this paper details, MPLS/Ethernet networks offer a far less expensive and more scalable means of carrying Frame Relay access traffic — and even ATM end user traffic — while also supporting Ethernet services. With this is mind, we can turn to the growth of Metro Ethernet.

PART II: THE RISE OF ETHERNET IN THE METRO

While Frame Relay is established as the incumbent data networking service in the Metro, Ethernet services are clearly the emerging choice. Our story begins in the United States, with newer network operators such as Intellispace, and Telseon. These companies have demonstrated not only that Ethernet is viable for metro networking but that considerable demand for Ethernet-based services exists. These operators find the capital and operational cost of Ethernet equipment low relative to traditional MAN gear, while their customers enjoy pricing as much as 50 percent lower than equivalent TDM connections.

The success of these start-ups has led, in turn, to a worldwide trend of incumbent and national service providers offering Ethernet-based services. In the United States, AT&T announced in September 2001 the launch of both high-speed Ethernet MAN and Internet services. Their metro service is being offered in seven cities initially and will be expanded to AT&T's 75 local network service areas. The Internet service is being offered as part of AT&T's Managed Internet Service and is initially being offered in New York and San Francisco, with plans to offer services in nine cities by early 2002. In turn, regional Bell operating companies such as SBC, BellSouth, and Verizon have indicated that they will begin to offer Ethernet services over the next few years.

Internationally, the Ethernet model has also been adopted by both national carriers and new service providers. Among those offering Ethernet services are National Telcos like Korea Telecom, China Telecom, and Dubai Telecom, in addition to newer service providers like Storm Communications across Europe, Hutchison Global Crossing in Hong Kong, and Crosswave Communications in Japan. Often the overseas model is focused on residential Ethernet services, though private business networking is also popular.

As for any networking technology, what drives service provider interest in Ethernet is the ability to support new services and deliver old services in less expensive ways. Ethernet is a mature, well standardized technology; its strengths include high speed, low cost, and operational simplicity. With an Ethernet infrastructure, metro service providers can offer competitively priced data access in the high megabit — and even gigabit — range.

In short, we can see from these deployments that metro Ethernet is emerging as a consensus choice for future deployments in metro area access and aggregation networking, particularly where fiber is available. But we would be remiss to conclude without mentioning the contribution made by MPLS to the long-term viability of these networks.

MPLS makes Ethernet a serious private networking technology. By bringing low-cost connection-oriented capabilities to the Ethernet environment, it creates the features demanded by businesses: bandwidth guarantees, fast recovery times, customized rerouting, the ability to easily connect a number of branch offices to a central headquarters over a "private" network and other features usually associated with ATM or Frame. In a combined Ethernet/MPLS architecture, MPLS expands Ethernet beyond high-speed connectivity to make it a platform that can create SLA-backed services.

PART III: MPLS: BRINGING FRAME RELAY AND ETHERNET TOGETHER

For private networking, Frame Relay/ATM and Ethernet are the most important network technologies in today's metro networks. Let's now look at the role MPLS can play in integrating Ethernet and Frame Relay traffic over a single shared infrastructure, to allow service providers to offer new Ethernet services and support Frame Relay services at the same time.

It is MPLS' ability to tunnel various types of traffic that is key. In the recent Martini Drafts, the IETF has focused on specifying MPLS Layer-2 transport tunnels to accommodate Frame Relay and other forms of legacy traffic. Specifically, the Martini drafts define encapsulation and label distribution mechanisms that can be used to transport Frame Relay, ATM, Ethernet, High-level Data Link Control (HDLC) and Point-to-Point Protocol (PPP) traffic within Label Switched Paths (LSPs) across an MPLS network.

One result is a network like that pictured below: A multi-purpose MPLS/Ethernet cloud that aggregates both Ethernet and Frame Relay service traffic.

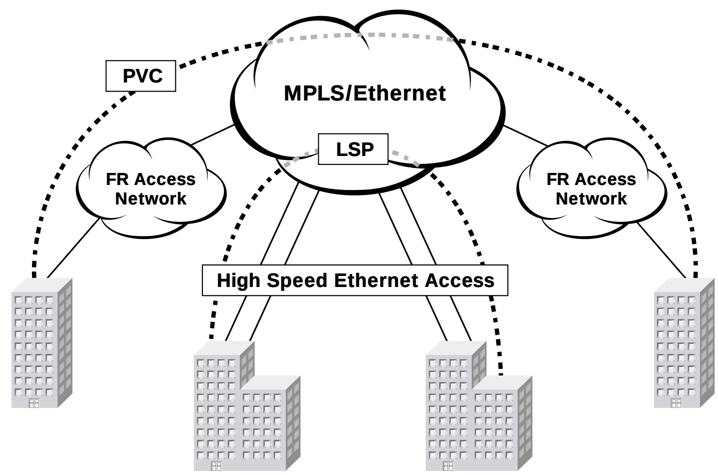

Figure 1: Multi-purpose MPLS/Ethernet Cloud Supporting Frame Relay and Ethernet sites.

New Directions in Layer-2 VPNs

To truly understand how MPLS can be used to integrate Ethernet and Frame Relay, we need to dig more deeply into the topic of Layer 2 VPNs. Aside from simple Internet access, Layer 2 VPNs are the most important business application for both Frame Relay and Ethernet metro networks. The striking innovation of a metro MPLS network is that it treats Ethernet and Frame Relay Layer 2 VPNs the same way: Both are supported and aggregated using Martini tunneling technology.

In the broadest sense, a private data network or VPN is simply a means of providing private communications over a public network infrastructure. In an L2 VPN service, customer premise equipment (CPE) receives Layer 2 (data link layer) service from the service provider. In this paper, we use the term "L2 VPN" to refer to a basic technology approach that service providers can use to create a variety of services, such as VPNs, Transparent LAN Service (TLS), and virtual leased lines.

With L2 VPNs, the customer and provider edge devices are adjacent to each other at the data link layer, not at the IP layer, on the access network. Consequently, the provider's equipment forwards customer data packets based on information in the packets' data link layer headers, such as a Frame Relay DLCI or the Ethernet 802.1q virtual LAN (VLAN) tag. In essence, the CPE sees the provider's equipment as a Layer 2 device such as a Frame Relay switch or an Ethernet VLAN switch.

Layer 2 VPNs have been around for years in the form of Frame Relay permanent virtual connection (PVC)-based services. For example, with Frame Relay, network operators create "tunnels" by establishing the data link connection identifiers (DLCIs) for a pair of communicating end points.

It is important to understand that L2 VPNs represent a fundamentally different approach than what are called "Layer 3" or "private routed" MPLS VPNs. A full treatment of this topic is beyond the scope of this paper (for more information, refer to Wu, "MPLS VPNs: Layer 2 or Layer 3? Understanding the Choice." Available at http://www.riverstonenet.com/technology/mpls_vpn.shtml). In short, we believe that Layer 2 MPLS VPNs have clear and incontrovertible advantages for service providers who want to support existing Frame Relay customers. but also have a scalable, future-proof solution as part of their business plan.

The key advantage is the flexibility provided by the Martini tunneling technology. Not all frame traffic is IP. Because L2 VPNs are multi-protocol in nature, they can support both IP and non-IP traffic. And by operating at Layer 2, L2 VPNs eliminate the need for service providers to participate in a customer's Layer 3 routing, rendering things like a customer's private IP addressing scheme irrelevant. This is significant because service providers can exploit this feature to interconnect the great majority of enterprise customer networks that use private IP addresses.

Most important, there's no need to modify or upgrade existing Frame Relay customer equipment, while customers are assured of the same or better data security as associated with Frame Relay, with the option of layering additional encryption as needed. With L2 VPNs, the customer's access to the service provider network remains the same and the associated service appears the same. For example, when using MPLS L2 tunneling technology such as Riverstone Networks provides with its Martini draft implementation, Frame Relay customers continue to access the metro network using Frame Relay. These Frame Relay frames are then transported via an MPLS-based tunnel across the provider's Ethernet/MPLS metro backbone and delivered as frames to the customer's receiving Frame Relay site.

With a combined Ethernet/MPLS architecture that supports Martini tunneling, metro service providers can handle customer traffic as well as link to other Frame Relay service providers.

We've now explained the basics of how an MPLS/Ethernet network can support both Frame Relay and Ethernet services over the same infrastructure. This section explains in greater depth exactly how Martini tunneling works.

MPLS creates tunnels using label switching rather than network-layer encapsulation. The difference with, say, IP tunneling, is that it is possible to use control protocols, such as the MPLS Label Distribution Protocol (LDP), to set up emulated virtual circuits (VCs) that carry the Protocol Data Units (PDUs) of Layer 2 protocols across a network. A number of such emulated VCs can be carried in a single tunnel so long as the PDUs are encapsulated.

The IETF Martini Internet Drafts leverage these MPLS capabilities to create tunnels for Layer 2 traffic. Specifically, the two Martini drafts (discussed in more depth below) define a technique for encapsulating Frame Relay, ATM, Ethernet, and other traffic types within MPLS frames as well as the necessary label distribution procedures for setting up the LSPs used for tunneling.

The Martini drafts have become firmly established in the industry as the standard for creating MPLS-based Layer 2 VPNs. Implementations on these drafts have become available on metro routers; in fact, a robust implementation of Martini tunneling is available across Riverstone's router product line. At the time of this writing, over 2,000 ports have been deployed.

In addition, the Martini drafts are serving as the starting point for further extensions to the basic MPLS L2 tunnel model. For example, in the case of Ethernet encapsulation, we have proposed extensions to the Martini model to support multipoint-to-multipoint connectivity, allowing for broadcasting and multicasting across VPN and TLS connections. These extensions are codified in the Lasserre draft, draft-lasserre-tls-mpls-00.txt. Another important extension is the "decoupled Transparent LAN Service" model, which is described in greater detail in draft-kompella-ppvpn-dtls-00.txt

Inside the Martini Drafts

While it is common to see references to "the Martini Draft," there are actually two Martini Drafts. The first, formally "draft-martini-l2circuit-encap-mpls-03.txt," specifies the emulated VC encapsulation for several L2 protocols. It also specifies a demultiplexer field, which is used in an MPLS environment to distinguish individual emulated VCs within a single tunnel.

The first Martini draft makes tunneling as simple as possible using the concept of the "Control Word." In most cases, it's not necessary to transport the Layer 2 encapsulation across the network. Rather, the Layer 2 header can be stripped by the ingress LSR (also known as a Label Edge Router or LER) and reproduced by the egress LSR using information carried in special control bits known as the Control Word.

The second Martini draft (draft-martini-l2circuit-trans-mpls-07.txt) specifies the label distribution procedures for transporting encapsulated L2 PDUs across an MPLS network. Currently, this Martini draft specifies the Label Distribution Protocol (LDP) for LSP set up. Members of the IETF are also exploring the use of Border Gateway Protocol (BGP) to distribute label blocks and mappings to Frame Relay DLCIs, ATM Virtual Channel Identifiers (VCIs), and other identifiers. We are actively tracking such work, which will undoubtedly evolve over time.

An important aspect of the Martini approach is the idea of using separate "Virtual Circuit" (VC) and "Tunnel" labels to create a two-level hierarchy of MPLS labels. This is a simple idea but rather ingenious in application. The tunnel label determines the path a packet takes through the MPLS "core," but knows nothing about the content of the packet. For its part, the VC label identifies what the egress LSR should do with the packet, based on its VPN, or VLAN, identifier. All services look like a VC to the MPLS network, and network operators provision a service by associating each end point with a common VC identifier (VCID).

Each supported Layer 2 technology is encapsulated slightly differently. For Frame Relay, the encapsulation consists of the tunnel label, VC label, Control Word, and Frame Relay PDU. An encapsulated Frame Relay frame is transported without the Frame Relay header (Q.922 address) or the Frame Check Sequence (FCS). The Control Word, which is required for Frame Relay encapsulation, carries the FECN (forward explicit congestion notification), BECN (backward explicit congestion notification), DE (discard eligible), C/R (Command Response) bits, and the PDU length; sequence number is optional. Figure 2 illustrates the Frame Relay Layer 2 encapsulation.

While the Control Word is required, its use is optional. For example, if an ingress LSR chooses not to use the Control Word, it must set the flags in the Control Word to 0. Authors of the Martini draft, however, indicate that use of the Control Word is "desirable."

For ATM tunneling, the Martini draft specifies encapsulations for both ATM AAL5 traffic and for entire ATM cells; Riverstone supports AAL5 encapsulation. Per the Martini draft, an ingress router performing AAL5 encapsulation is required to reassemble AAL5 Common Part Convergence Sublayer (CPCS) PDUs from the incoming ATM VC and transport each CPCS PDU as a single packet; both the ATM header and AAL5 trailer are dropped.

As with Frame Relay encapsulation, the Control Word is required in ATM AAL5 encapsulation, though its use is optional. In the AAL5 case, the Control Word contains a transport byte, which indicates whether the packet contains an ATM cell or an AAL5 CPCS-PDU, the EFCI (explicit forward congestion indication) bit, the CLP (cell loss priority) bit, the command/response field bit, length, and sequence number.

Should a service provider want to offer encapsulated, rather than native, Ethernet services, the Martini encapsulation draft specifies that, for simple port-to-port transport, the entire Ethernet frame without the preamble and FCS is transported. Use of the Control Word is optional.

Riverstone supports the Martini Layer 2 tunneling technology on a variety of its metro Ethernet devices. To date, Riverstone has implemented Frame Relay, ATM AAL5, and Ethernet encapsulations. Riverstone provides Martini encapsulation on a per-interface basis, giving service providers the flexibility to support Frame Relay, ATM, Ethernet, and other encapsulations on a single piece of equipment. As we noted earlier, we have also implemented extensions to the Martini Ethernet specification to allow for multipoint transparent LAN services, following the approach documented in the previously cited IETF Draft draft-lasserre-tls-mpls-00.txt.

PART IV: IN PRACTICE: DEPLOYING MARTINI TUNNELS

Now that we've gone over how this Layer 2 tunneling technology works, let's look at how it can be deployed. This section suggests practical strategies for deploying multi-service MPLS networks, stopping just short of providing actual configurations. For readers interested in the full details of configuration commands, please visit the Riverstone's advanced MPLS documentation, available from the page http://www.riverstonenet.com/support/mpls/index.shtml.

For service providers and carriers looking to support both legacy customer traffic and high-speed native Ethernet connections, the challenge is to move from a traditional Frame Relay/ATM environment, as depicted in Figure 2, to a parallel or hybrid environment based on a Gigabit Ethernet/MPLS metro backbone.

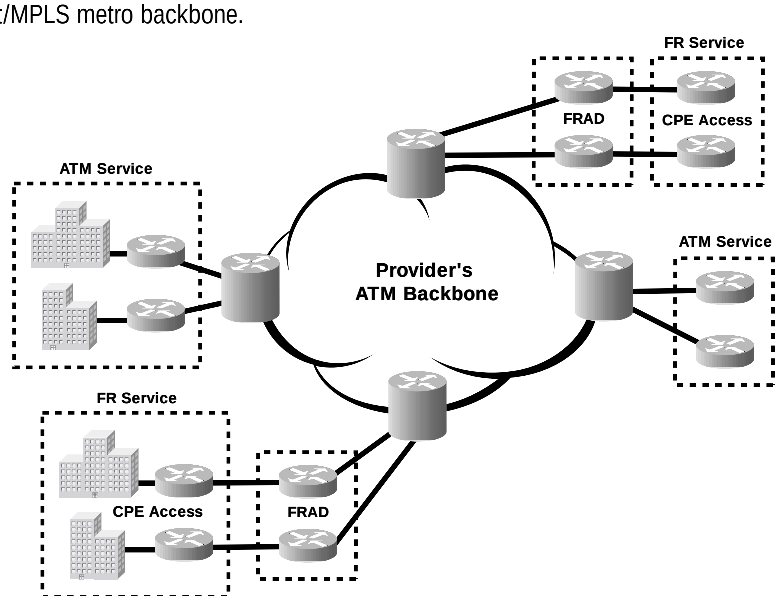

Figure 2: Existing FR/ATM metro network

Similarly, greenfield Gig E metro operators who have deployed a pure Ethernet environment, as shown in Figure 3, may want to accommodate Frame Relay or even ATM customer access.

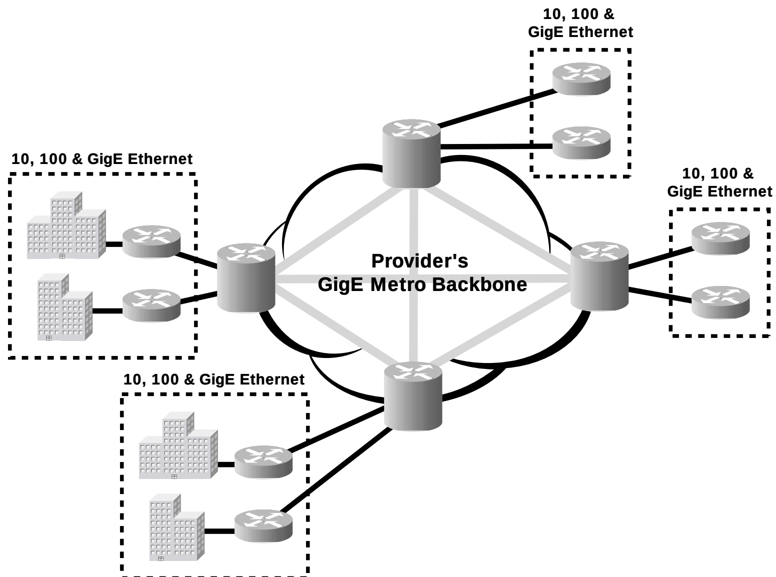

Figure 3: "Pure" Gig E metro network

Network operators have the flexibility to deploy Ethernet/MPLS routers in tandem with Frame Relay and ATM gear in existing POPs or to install Ethernet/MPLS equipment in new POPs. Within a given POP, service providers may initially want to run separate Ethernet/MPLS and Frame Relay/ATM networks, and then migrate the Frame Relay and ATM traffic to the Ethernet/MPLS infrastructure over time.

Greenfield metro Ethernet providers who want to tunnel legacy customer traffic can add Frame Relay and ATM interfaces to their metro core routers if they deploy equipment that supports the Martini drafts. By deploying Riverstone routers at the edge of their metro networks, service providers and carriers can position themselves to support high-speed native Ethernet services as well as to tunnel legacy traffic.

Providers with existing Frame Relay/ATM metro architectures may opt to deploy Riverstone routers in separate racks within the same POP as the Frame Relay/ATM gear. When deployed in an existing POP, the Riverstone router could become part of the provider edge network, operating directly downstream of an ATM switch or FRAD. In this scenario, Frame Relay or ATM traffic from a customer site would come into an ATM switch, for example, and be passed to the MPLS router.

Operating as an MPLS LSR, the router would perform the appropriate Martini encapsulation. In the case of ATM, the Riverstone router would perform the reassembly function to create a CPCS PDU, along with other encapsulation functions. The router would then map each incoming Frame Relay DLCI or ATM VCI/VPI to an individual VC LSP. These VC LSPs would then be transported via a tunnel LSP across the Ethernet metro backbone to the appropriate egress LSR. At the opposite edge of the backbone, the egress router strips off the tunnel label and does a reverse table mapping for each VC LSP to the correct DLCI or VCI/VPI. The Frame Relay or ATM traffic is then forwarded to the customer destination. This process is illustrated in Figure 3.

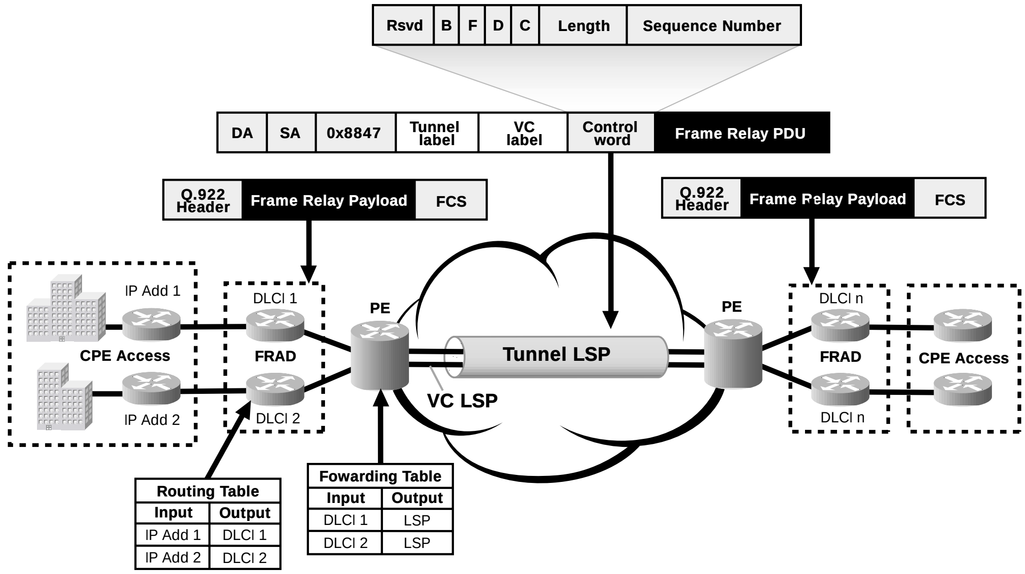

Figure 4: Frame Relay tunneled across MPLS metro

In the case of new Ethernet metro POPs, the Riverstone router at the metro edge can directly accept customer Ethernet as well as legacy Frame Relay traffic. As in the previous scenario, the router would encapsulate the Frame Relay traffic, map each incoming DLCI to a VC LSP, transport the encapsulated traffic across a tunnel LSP, perform the decapsulation and reverse table mapping, and forward the Frame Relay frames to the customer destination.

For service providers who want to offer any-to-any Layer 2 connectivity, Riverstone provides bridging functionality. A detailed discussion of this bridging technology is beyond the scope of this paper; consult Riverstone advanced technical documentation at http://www.riverstonenet.com/technology. At a basic level, Riverstone's bridging allows for any L2 technology, be it Ethernet, Frame Relay, or ATM to be terminated and then bridged across a Gigabit Ethernet metro backbone, and then delivered in any supported L2 frame format.

When we take all of these technologies and put them together, the resulting network looks something like that pictured below:

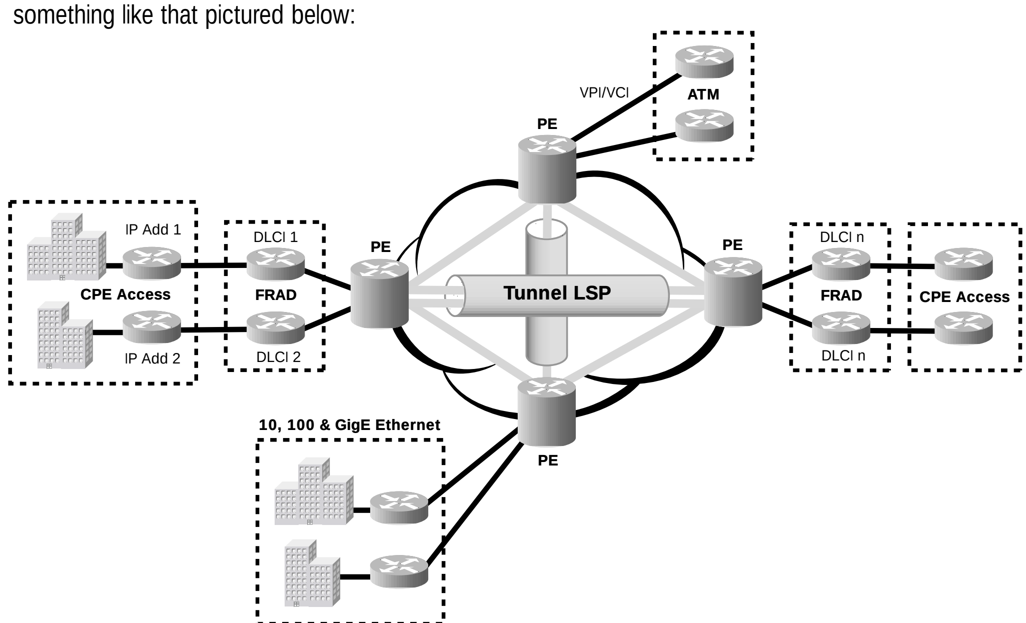

Figure 5: Mixed Legacy/Next Generation Network

CONCLUSION

A combined Ethernet/MPLS metro architecture can increase a provider's competitiveness by providing infrastructure that protects revenue from existing services while also supporting new services. Martini-based tunneling technology has emerged as the standard for accommodating existing Frame Relay (and ATM) traffic across such a combined infrastructure.

The bottom line for customers is that their Frame Relay service appears unchanged, while the service provider is able to sell new Ethernet services — including both premium high-speed access and LAN extension services. For traditional carriers and providers, such an architecture allows them to migrate their metro core network to Ethernet/MPLS in a flexible fashion. Providers can, for example, cut over existing Frame Relay and ATM traffic to the Ethernet/MPLS backbone on an as-needed basis.

Some providers may find it more cost-effective to support higher speed and/or higher volumes of Frame Relay traffic over a new Ethernet/MPLS infrastructure than to expand their existing ATM infrastructure. For their part, greenfield metro Ethernet operators can leverage a combined Ethernet/MPLS architecture to accommodate broader customer traffic types beyond native Ethernet. The bottom line for carriers and service providers is that they can continue to provide profitable Frame Relay and ATM services and expand them as needed, supported by a high-bandwidth Ethernet/MPLS infrastructure.

SOME MAJOR RELATED IETF DRAFTS

The Martini Drafts
"*Transport of Layer 2 Frames Over MPLS*", draft-martini-l2circuit-trans-mpls-07.txt

"*Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*", draft-martini-l2circuit-encap-mpls-03.txt

The Lasserre TLS Draft
Transparent VLAN Services over MPLS, draft-lasserre-tls-mpls-00.txt

Decoupled TLS
Decoupled Transparent LAN Services, draft-kompella-ppvpn-dtls-00.txt

Others
Virtual Private Switched Network Services over an MPLS Network, draft-vkompella-ppvpn-vpsn-mpls-00.txt

MPLS-based Layer 2 VPNs, draft-kompella-ppvpn-l2vpn-00.txt

Extensions to MPLS-based Layer 2 VPNs, draft-shah-mpls-l2vpn-ext-00.txt

END NOTES

1. "*Frame Relay Service Update: The Latest in a Legacy Technology,*" **Yankee Group Report**, Communications Services for the New E-conomy, Vol. 2 No. 5 (2001); see also BCR 9/19 (2001).

2. See *Frame Relay Survey, Vertical Systems Group 2001;* see also Comments of Rosemary Cochran in "*A Bright Future for ATM,*" **Business Communications Review**, September 2001, at 27.

3. *For more on this topic, see "Making The Most of Ethernet in the Metro,"* available at http://www.riverstonenet.com/technology/mpls_ethernet.shtml.

4. "*MPLS VPNs: Layer 2 or Layer 3? Understanding the Choice,*" is available at http://www.riverstonenet.com/technology.

Acronyms

ACL	*Access Control List*
ANSI	*American National Standards Institute*
ASIC	*Application-Specific Integrated Circuit*
ASP	*Application Service Provider*
ATM	*Asynchronous Transfer Mode*
CBR	*Constant Bit Rate*
CWDM	*Coarse Wave Division Multiplexing*
DS1/DS3	*Digital Signal, Level 1 (1.54 Mbps) or 3 (44.7 Mbps)*
DSL	*Digital Subscriber Line*
DWDM	*Dense Wave Division Multiplexing*
DVMRP	*Distance Vector Multicast Protocol*
E1/E2	*European Trunk 1/2 (2 Mbps/34.3 Mbps)*
ERP	*Enterprise Resource Planning*
HSSI	*High Speed Serial Interface*
ISP	*Internet Service Provider*
ITU	*International Telecommunications Union*
LAN	*Local Area Network*
LEC	*Local Exchange Carrier*
MAC	*Media Access Control*
MAN	*Metropolitan Area Network*
MDU	*Multiple Dwelling Unit*
MLPPP	*Multi Layer Point-to-Point Protocol*
MPLS	*Multiple Protocol Label Switching.*
	See "MPLS in Metro IP Networks,"
	http://www.riverstonenet.com/technology/mpls.shtml
MTU	*Multiple Tenant Unit*
OC-3/OC-12	*Optical Carrier 3/12 (155 Mbps/622 Mbps)*
PDH	*Plesiochronous Digital Hierarchy*
PIM	*Protocol Independent Multicast*
POS	*Packet over SONET*
PPP	*Point-to-Point Protocol*
PVC	*Private Virtual Circuit*
QoS	*Quality of Service*
RED	*Random Early Discard*
SONET	*Synchronous Optical NETwork*
	See http://www.techguide.com/comm/sec_html/sonet.shtml
SLA	*Service Level Agreement*
SPE	*Synchronous Payload Envelope*
SRP	*Spatial Reuse Protocol*
	See RFC 2892
T1	*Trunk 1 (1.544 Mbps)*
TCP/IP	*Transport Control Protocol/Internet Protocol*
TDM	*Time Division Multiplexing*
UBR	*Undefined Bit Rate*
VBR	*Variable Bit Rate*
VLAN	*Virtual LAN*
VoD	*Video on Demand*
WAN	*Wide Area Network*
WDM	*Wave Division Multiplexing*
WRED	*Weighted Random Early Discard*

Riverstone Networks, Inc.
5200 Great America Parkway, Santa Clara, CA 95054 USA

877 / 778-9595 or **408 / 878-6500** or **www.riverstonenet.com**

Printed in the USA

v 1.0 11/01