Converging ATM/Frame Relay over MPLS - Solving Control Plane Interworking

Crafting standards to create a scalable, manageable, converged service infrastructure utilizing an MPLS core

September 2005

Nikhil Shah

Jean Jones

Table of Contents

Objective: This is an MFA Forum white paper that provides a high level description of the "ATM and Frame Relay to MPLS Control Plane Interworking" Implementation Agreements. It also demonstrates the benefits of the approaches defined in the IAs.

1. Introduction

The emergence of new standards for Layer 2 transport, new Layer 3 applications, and multiservice interworking over Multiprotocol Label Switching (MPLS) have made MPLS a desirable framework for integrating specialized ATM, Frame Relay, and IP networks into a scalable, converged service infrastructure. Concerted efforts by the MFA Forum and other standards bodies have given MPLS the capabilities that a service provider needs to build a reliable delivery platform for premium voice, video, and data services: agnostic access, standardized encapsulation, predictable QoS, and proven traffic management.

Frame Relay and ATM networks employ dynamic signaling and routing protocols to expedite connection setup and recovery. PNNI is an example of one such protocol that is used to establish switched ATM connections across a native ATM network.

When transitioning to MPLS, service providers will expect and desire ATM and Frame Relay signaling and routing protocols (e.g. PNNI) to interoperate with IP/MPLS signaling and routing protocols for the purpose of establishing switched connections across a Packet-Switched Network (PSN). For this reason it becomes necessary to define methods that enable interworking between the client control plane functions of ATM & Frame Relay and IP/MPLS . Functions described in this paper are based upon ‘network interworking’ concepts. The objective is to create the ability to dynamically establish switched point-to-point connections from one client network end-point to another, across an IP/MPLS network.

The MFA Forum is currently developing two Implementation Agreements (IA) on the topic of ATM-and-Frame Relay to MPLS Control Plane Interworking, to increase the scalability and manageability of ATM/FR[1]-over-MPLS solutions. This white paper provides a high level overview of both IAs, which cover:

- The network reference model
- A method for implementing dynamic control plane interworking between the ATM/FR and MPLS domains

The ATM Forum defined techniques that enable switched-connection set up across an MPLS network [1]. In that approach, the client networks are ATM, and the client control plane is also ATM-based (e.g. PNNI, AINI). It defines extensions for ATM signalling (e.g. PNNI) to signal Pseudo-Wire (PW) labels between the MPLS edge devices. The client and PW control functions are collapsed into a single ATM signaling protocol.

By contrast, the MFA IAs define an approach that decouples the client control plane from the PW control plane so as to minimize the dependencies of one on the other. Each is allowed to operate and evolve separately. By decoupling label exchange for PWs from the client control plane, this approach enables different client control planes (e.g. Frame Relay and ATM) to employ the same PW control plane. Conversely, different PW control planes may be used to support inter-client network connection setups.

[1] Many aspects of the IAs apply equally to ATM and Frame Relay client networks; we use the notation “ATM/FR” to mean “ATM or Frame Relay.”

2. Establishing the basic network design

The following are some of the considerations that were made in developing the IAs:

- Minimize the impact on current ATM and Frame Relay network architectures, procedures, and mechanisms, for example, with regard to end-to-end QoS and OAM
- Utilize existing IETF mechanisms (e.g. PWE3, LDP) as much as possible
- Ensure that new standards do not adversely impact the scale and performance required in service provider networks (e.g. signaling rates, use of network resources)

Currently, the IAs define the following two cases to address ATM and Frame Relay to MPLS control plane interworking:

- A "Client network-to-Client network" Case, in which a single edge router performs both ATM and MPLS control plane functionality, and
- A "Virtual Trunk" Case, which permits the distribution of ATM and MPLS control plane functionality across physically or logically separate boxes.

The following definitions apply to both cases:

- *Client Control Plane*, which performs all ATM or Frame Relay control functions that release and maintain switched connections between client network end-points (e.g., routing, signaling, etc.)
- *Pseudo-Wire Control Plane*, which performs all the MPLS control functions required to set up, tear down, and maintain the Pseudo-Wires transporting the ATM services through the MPLS core. Procedures for the handling of PWs are as per [2].
- The *Pseudo-Wire Service Interface*, an interworking function that translates service requests to establish, release, or maintain ATM or Frame Relay connections into service requests to the Pseudo-Wire control plane to set up, tear down, or maintain MPLS Label Switched Paths (LSPs) for transport across the core

Understanding the Client Network-to-Client Network Case

Figure 1 illustrates this case. As the diagram shows, all traffic from client edge networks enter the MPLS network through Label Edge Routers (LERs) stationed at the edge of the MPLS core. These LERs map each virtual channel (VC) or virtual path (VP) into an MPLS Pseudo-Wire. The LERs then aggregate the PWs into PSN Tunnels for transport through the MPLS core.

This case assumes that both the Client Control Plane and the MPLS Pseudowire Control Plane are in one MPLS LER. This design requires only one interworking function per LER, the Pseudo-Wire Service Interface, to translate service requests between the Client and the MPLS Pseudo-Wire Control Planes.

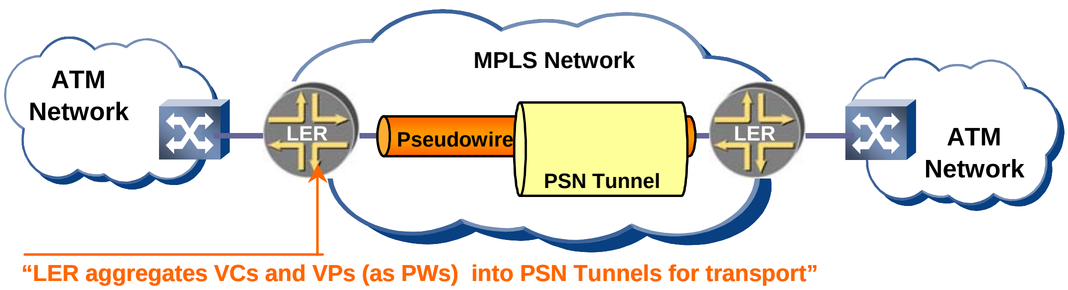

Figure 1—The Client Network-to-Client Network Case locates all ATM and MPLS control plane functions in the edge MPLS LER.

To transport ATM/FR services through the MPLS core in this scenario:

1. The ingress LER encapsulates each ATM/FR virtual connection into a single MPLS Pseudo-Wire.
2. The ingress LER aggregates those Pseudo-Wires onto a core MPLS tunnel.
3. The egress LER on the other side of the MPLS core reverses Steps 1 and 2 to complete the ATM/FR-over-MPLS connection.

For optimum scalability, the 'Client network-to-Client network' Case requires LERs that offer the full range of ATM or Frame Relay signaling and routing capabilities, robust standards-based IP/MPLS routing and traffic management, and high processing capacity.

Understanding the 'Virtual Trunking' Case

Figure 2 illustrates the Virtual Trunking Case, which makes it possible to split the ATM and MPLS routing functions, allowing the ATM switches to communicate through the MPLS core without requiring:

- Support of MPLS or Pseudo-Wire encapsulation on the ATM switches
- Support of ATM signaling and routing, or to interwork individual ATM VCs or VPs on the MPLS LERs

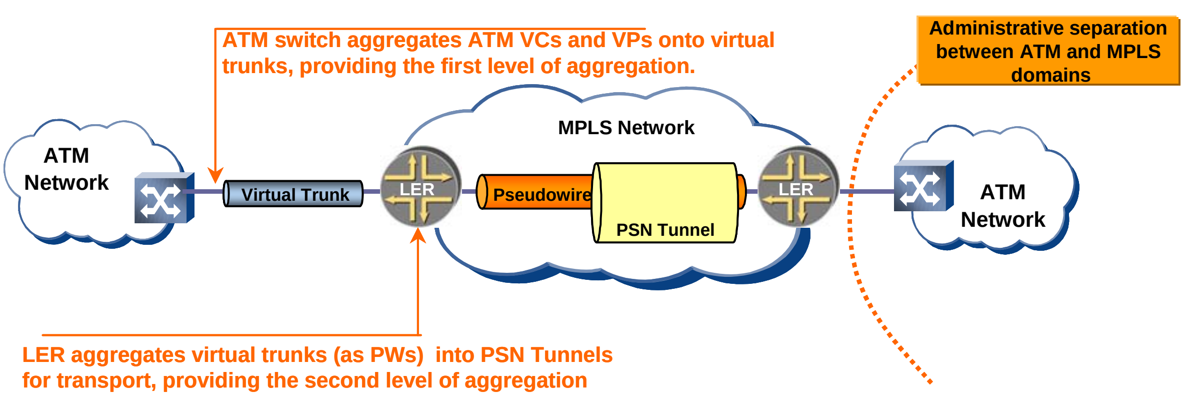

Figure 2—The Virtual Trunk Case separates the ATM and MPLS control plane functions

As in the previous ‘Client network-to-Client network’ Case, the ATM edge switches in the Virtual Trunk Case communicate with the MPLS core through LERs at the edge of that core. The Virtual Trunk Case implements the familiar:

- Client Control Plane to manage client service requests from the ATM network
- Pseudo-Wire Control Plane to maintain the Pseudo-Wires transporting those services through the MPLS core

However, the placement of these functions is quite different. Unlike the previous case, the Virtual Trunk Case locates the Client Control Plane’s ATM signaling and call control functions in a native Layer 2 switch at the edge of the ATM network—reserving the MPLS pseudowire and tunnel setup functions performed by the Pseudo-Wire Control Plane for the edge LERs.

To accommodate this functional divide, the Virtual Trunk Case requires a slightly different roster of interworking functions and interfaces. An *ATM virtual trunk (VT)* provides the ATM logical link between the two ATM switches. Adding this virtual trunk, which is simply a bundle of VPs that are encapsulated (at the LER) onto a single MPLS Pseudo-Wire, allows the ATM switch to aggregate individual ATM connections onto a higher-bandwidth trunk before transmitting them to the MPLS LER. The LER maps a virtual trunk to a Pseudo-Wire, using ATM N:1 cell mode encapsulation, establishing a one-to-one correspondence between a virtual trunk and a Pseudo-Wire. The ATM switches can use standard routing and signaling protocols to exchange routing information and to establish Virtual Path Connections (VPCs) and Virtual Channel Connections (VCCs) through the virtual trunks. To manage the virtual trunk, this case introduces a *Virtual Trunk Control Plane Interface* between the Client and Pseudo-Wire Control Planes, to establish, maintain, and release the virtual trunk connections between the ATM edge switch and the LER at the edge of the MPLS core.

To complete the end-to-end connection, the Pseudo-Wire Service Interface on the MPLS LER maps the aggregated ATM virtual trunk (rather than the individual ATM connections) onto an MPLS Pseudo-Wire. The Pseudo-Wire then emulates the ATM virtual trunk across the MPLS core.

3. Implementing dynamic interworking between ATM / Frame Relay and MPLS domains

To increase the scalability and manageability of ATM-over-MPLS solutions, an ATM-to-MPLS Control Plane Interworking standard should contain mechanisms that service providers can use to dynamically establish ATM-over-MPLS connections that preserve end-to-end QoS. Consequently, discussions are currently underway in the MFA Forum to define:

- A signaling mechanism to establish new connections dynamically
- A model for ensuring end-to-end QoS

Using dynamic signaling to establish new ATM-over-MPLS connections

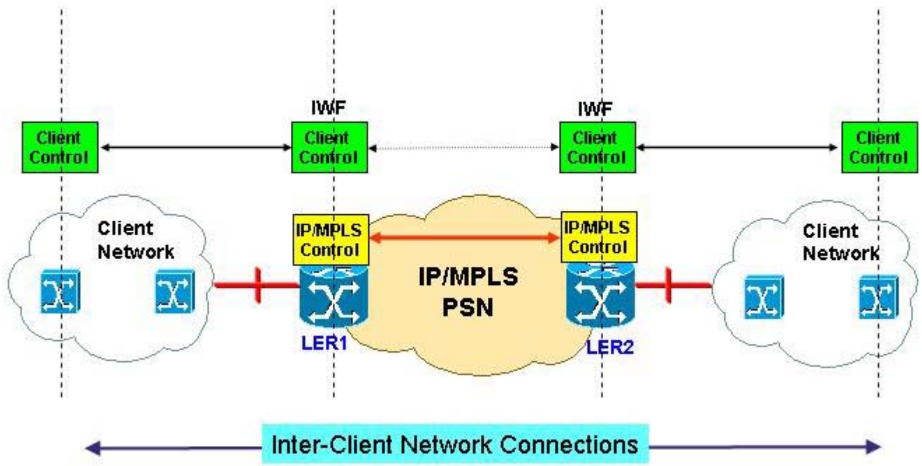

Figure 3-- High-Level Architecture.

Figure 3 depicts the high-level architecture for end-to-end call setup. Client networks operating a native client (ATM or Frame Relay) control plane attach to Label Edge Routers (LERs) that are part of an IP/MPLS PSN. The LERs operate an IP/MPLS control plane protocol (i.e. IGP/BGP, LDP, PWE3, RSVP-TE, etc.) to set up and manage label-switched paths (LSPs) across the network. The LERs either operate an instance of a co-located client control plane (for the Client network-to-Client network Case), or may interface to a client control plane running on a separate physical or logical device (for the Virtual Trunk Case).

For the Client network-to-Client network Case, client network control plane (e.g. PNNI) messages are used to dynamically route, set up, manage, and tear down inter-client network connections. Receipt of client control plane messages by an LER will dynamically trigger the corresponding IP/MPLS control plane functions (such as setup and tear-down of Pseudo-Wires). The procedures specifying the interaction between the Client Control Plane and the Pseudo-Wire Control Plane at the LERs are defined in the IA, along with the associated encoding extensions to the PWE3 control plane. In addition, the IA specifies how to establish the control PWs that carry the client control plane messages between the LERs.

As described earlier for the Client network-to-Client network Case, the virtual trunking concept permits the separation of client control and Pseudo-Wire control into different physical (or possibly logical) devices, and provides an additional level of aggregation for ATM connections. Since the ATM control plane is tunneled through virtual trunks, along with the end-to-end ATM connections, ATM connection setup is completely transparent to the LERs. The virtual trunk is initially established by operator configuration. The MFA Forum is planning to evaluate (in the future) the desirability and feasibility of dynamic setup of a virtual trunk using a control protocol.

Preserving QoS across the MPLS core

Because QoS is critical to successful ATM and Frame Relay services, particularly tariffed services backed by SLAs, operators must be able to ensure consistent QoS treatment for traffic carried over an MPLS network.

For the Client network-to-Client network Case, existing methods for ensuring QoS in MPLS networks are applicable. Since each LER sees all the client control plane messages triggering establishment of ATM and Frame Relay connections through that LER, the LER knows the QoS requirements of each client connection on it. This information can be used to manage the resources for the PSN tunnels that transport the client connections across the MPLS network, and can also be used when selecting resources for local support of Pseudo-Wires and attachment circuits.

For the Virtual Trunk Case, two scenarios are defined for providing end-to-end QoS for the virtual trunk data plane:

- Single-QOS VT - all traffic in the virtual trunk receives the same QoS treatment through the MPLS core
- Multi-QoS VT - traffic for different connections within the virtual trunk may receive different QoS treatment through the MPLS core

An operator deploying a Single-QoS VT has two options:

- Carry traffic of more than one service class in the VT. Traffic with less stringent QoS needs will receive better QoS treatment than is required.
- Create multiple VTs between a pair of LERs, associating each VT with a specific service class. The ATM edge switches bundle the VCs and VPs into virtual trunks by service class as well as by VT end-point.

In contrast, for Multi-QoS VTs, VPs within the VT can belong to different service classes. QoS treatment is defined on a per-VP basis. Support for Multi-QoS VTs is optional. One of the ways to implement Multi-QoS VTs is by provisioning the ATM edge switches with contiguous ranges of VPIs per service class. Similarly, the association between particular VPI ranges and MPLS network QoS treatment is provisioned on the MPLS Edge.

As a result, ATM edge switches can aggregate ATM VPs from multiple service classes onto one virtual trunk while maintaining the association with the QoS treatment required in the MPLS network. MPLS LERs could also bundle MPLS Pseudo-Wires from multiple classes onto the same MPLS tunnel, for example in the case of E-LSPs. Allowing one MPLS tunnel to carry traffic from multiple classes can increase scalability; instead of establishing one tunnel per class, operators can use one tunnel to carry all service classes, reducing the total number of tunnels in the core.

4. Conclusion

When complete, the ATM-and-Frame Relay-to-MPLS Control Plane Interworking IA will significantly increase the overall scalability and manageability of large MPLS core networks by:

- Enabling service providers to implement networks that follow either the Client network-to-Client network or Virtual Trunk Case, depending upon their current embedded network assets, operational preferences, and scalability requirements
- Allowing a dynamic signaling mechanism to establish ATM-over-MPLS and Frame Relay-over-MPLS connections
- Providing QoS support comparable to that of traditional ATM and Frame Relay networks

At the same time, the ATM-and-Frame Relay-to-MPLS Control Plane Interworking solution will enable service providers to ensure that they provide reliability using standard ATM, Frame Relay, and MPLS service restoration and fault management tools such as MPLS Fast Reroute and 1+1 Automatic Protection Switching (APS). These solutions also allow for slower failover using rerouting in the client control plane, in case the faster methods do not recover the connections or are not used.

Acknowledgements

The MFA Forum thanks Nikhil Shah (Lucent) and Jean Jones (Lucent) for authoring and contributing to this paper. Special appreciation to Bernard da Costa (Bell Canada), Rao Cherukuri (Juniper Networks), Mickey Spiegel (Cisco Systems), Matthew Bocci (Alcatel), Himanshu Shah (Ciena), Tom Walsh (Lucent), John Kenney (Tellabs), Peter Busschabach (Lucent), Daniel Porch (Marconi) for their contribution in reviewing the content and offering valuable suggestions.

References

[1] ATM Forum, af-aic-0178.001.pdf. ATM-MPLS Network Interworking Version 2.0. 08/2003.

[2] Martini, L., et al. Pseudowire Setup and Maintenance using LDP. IETF, draft-ietf-pwe3-control-protocol-15.txt. February, 2005

Appendix A: Acronyms

The following table lists the acronyms used in this paper.

AINI	ATM Inter-Network Interface
APS	Automatic Protection Switching
BGP	Border Gateway Protocol
IA	Implementation Agreement
IETF	Internet Engineering Task Force
E-LSP	EXP inferred LSP
LER	Label Edge Router
LSP	Label Switched Path
MP-BGP	BGP with Multi-Protocol extensions
MPLS	Multiprotocol Label Switching
PNNI	Private Network-to-Network Interface
PSN	Packet-Switched Network
PW	Pseudo-wire
PWE3	Pseudo Wire Emulation Edge to Edge
QoS	Quality of Service
RSVP	Resource Reservation Protocol

RSVP-TE	RSVP with Traffic Engineering
UNI	User to Network Interface
VCC	Virtual Channel Connection
VPC	Virtual Path Connection
VPI	Virtual Path Identifier
VT	Virtual Trunk