MPLS & Frame Relay Alliance

MPLS PVC User to Network Interface
Annex B:
MPLS Proxy Admission Control Protocol

Implementation Agreement

MPLS & FR Alliance 7.0.0

MPLS & Frame Relay Alliance Technical Committee
October 2004

Editor:

Tom Phelan
Sonus Networks

For more information contact:

The MPLS & Frame Relay Alliance
Suite 307
39555 California Street
Freemont, CA 94538 USA

Phone: +1 (510) 608-3997
Fax: +1 (510) 608-5917
E-mail: info@mplsforum.org
WWW: http://www.mplsforum.org

Table of Contents

Annex B: MPLS Proxy Admission Control Protocol

(NORMATIVE, OPTIONAL)

B.1 Introduction

The MPLS PVC User to Network Interface (UNI) [B.1] describes an interface between an MPLS network and its customers for signaling the status of PVC connections, and a way to extend that interface for other applications. This document is an annex to [B.1]; it describes extensions to the UNI protocol that are necessary to implement the MPLS Proxy Admission Control capability. It is one part of a two-document set. The operational and provisioning aspects of the MPLS Proxy Admission Control capability are described in the companion document, *MPLS Proxy Admission Control Definition* [B.2].

B.1.1 MPLS Proxy Admission Concepts

Consider the simple Traffic Engineered Multiprotocol Label Switched network (TE-MPLS, [B.3]and[B.4]) shown in Figure B-1. In this figure, a number of Provider Edge devices (PEs) are connected via unidirectional Traffic Engineered PE-to-PE LSP tunnels. The PEs provide an IP network service to the Customer Edge devices (CEs). Interconnection between the PEs and the CEs is via a layer 2 network such as Ethernet or point-to-point links. The types of CE equipment include not only routers, but also directly connected high-capacity application devices such as Public Switched Telephone Network (PSTN) gateways, telephony application servers (such as audio conferencing servers), video media servers and others.

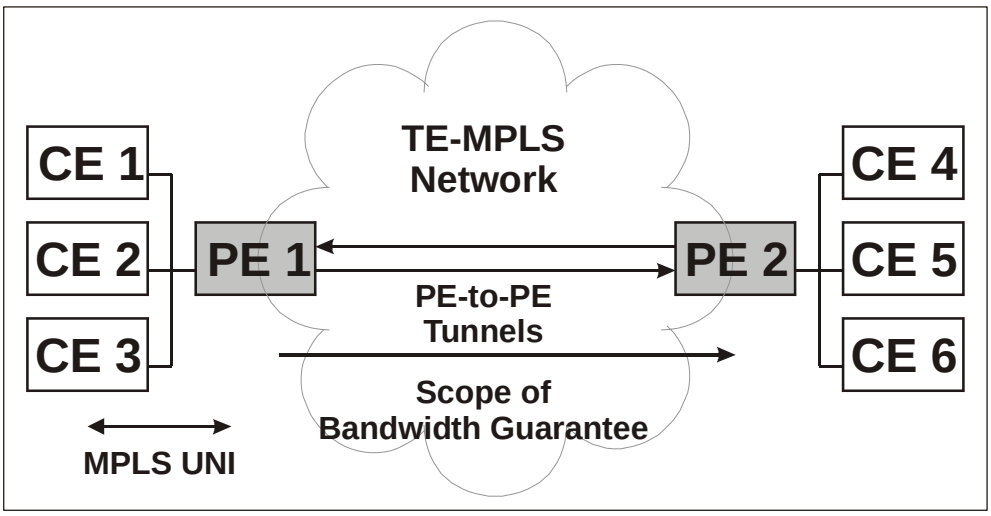

Figure B-1: MPLS QoS Proxy Admission Service reference model.

In large networks the number of PEs is often in the tens to low hundreds; the number of CEs is often in the hundreds or thousands. If a CE-to-CE LSP needs to be configured for every possible CE-to-CE combination, or each CE needs to be configured with its share of network resources, the amount of information to be configured grows with the square of the number of devices. The amount of labor required to perform the configuration, and the chance of error, are dramatically increased as the network grows. In addition, with static configuration of the LSPs or CEs, there is no opportunity to dynamically share network resources.

What each CE wants is:

- To have some portion of the capacity on the tunnels between the PEs allocated on demand for its private use.
- To be able to increase (assuming unused capacity is available) or decrease that allocation on demand.

- To be informed of network events that impact these allocations.
- To be able to unambiguously indicate which packets it transmits are to be applied to which reservation.

The MPLS User-to-Network Interface (UNI) [B.1]defines an interface between PEs and CEs over which various capabilities can be offered. One of these capabilities is LSP Connection. While this capability can supply CE-to-CE LSPs with QoS guarantees, scalability problems develop when there are hundreds or thousands of CEs to interconnect, or when pre-configured bandwidth-sharing is not adequate.

The MPLS Proxy Admission Control capability is intended to address these issues. MPLS Proxy Admission Control is a capability offered over the MPLS UNI; it allows CEs to dynamically share network resources without the need for coordinated and static configuration of the CEs. The base protocol mechanisms that support the capability are defined in the MPLS PVC UNI Implementation Agreement [B.1]; the extensions necessary to support MPLS Proxy Admission are defined in this Implementation Agreement. The functional aspects of the capability are defined in the MPLS Proxy Admission Control Definition document [B.2].

The MPLS Proxy Admission Control Protocol interactions happen between the CEs and the ingress PE. The remote PE and CEs are not directly involved. The ingress PE provides proxy admission control for the tunnel that runs from itself to the egress PE. In a realistic network each PE will have a large number of tunnels connecting it to a large number of egress PEs. The tunnel ends at the egress PE. The egress PE is *not* the penultimate hop; it is the ultimate hop. Packets are IP routed from the egress PE to the destination CE, and therefore they must have an IP header (v4 or v6). The current version of the protocol does not support other network layer protocols directly over MPLS. Support of other protocols is for future study.

B.1.2 Scope

MPLS Proxy Admission Control allocates and guarantees bandwidth inside the service provider's network, from the ingress PE to the egress PE. It does not guarantee bandwidth in the access networks from CE to ingress PE, and from egress PE to CE. Applications that desire guaranteed bandwidth in the CE-to-PE or PE-to-CE access networks must use other means to ensure the availability of bandwidth in the access networks. Possible means of achieving this include layer 2 QoS protocols such as 802.1Q/p [B.5]or session-layer application-specific admission controls.

Signaled support for other Quality of Service parameters, such as delay, jitter, drop probability, and class of service indications, is for future study.

The traffic used with this feature must be IP-addressed. There is no label on the destination PE-CE interface. Once the tunnel label is popped the traffic is IP-routed to the correct egress interface. Support for non-IP traffic is for future study.

There are PE-to-PE tunnels between PEs that participate in this feature. Bandwidth is reserved over these tunnels. This feature and the allocation of resources to CEs are transparent to intermediate or core routers.

MPLS Proxy Admission Control does not guarantee in-order delivery of packets.

B.2 Definitions and Terminology

B.2.1 Definitions

Must, Shall or Mandatory — the item is an absolute requirement of this implementation agreement.

Should — the item is desirable.

May or Optional — the item is not compulsory, and may be followed or ignored according to the needs of the implementer.

Notes — outside of Tables and Figures are informative.

B.2.2 Acronyms

CE	Customer Edge devices
FEC	Forwarding Equivalence Class
LDP	Label Distribution Protocol
LSP	Label Switched Path
MPLS	Multi Protocol Label Switching
PE	Provider Edge device
PSTN	Public Switched Telephone Network
PVC	Permanent Virtual Circuit
QoS	Quality of Service
RIL	Resource Index Label
TE-MPLS	Traffic Engineered Multi Protocol Label Switching
TLV	Type, Length, Value encoding
UNI	User-to-Network Interface

B.3 Messages and TLVs

This section describes modifications of existing MPLS UNI TLVs and messages to be used specifically for the MPLS Proxy Admission Control Protocol.

B.3.1 UNI Capability Element

The MPLS Proxy Admission Control Protocol defines a new UNI capability bit and PE and CE mode bits in the UNI Capability Element. These new bits are always zero when the MPLS Proxy Admission Control Protocol is not used.

```
UNI Capability Element value encoding:
```

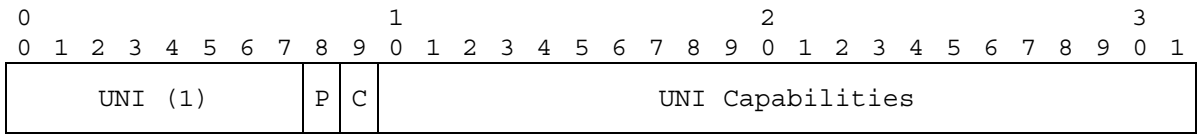


```
Bit 8:   P – PE Mode
Bit 9:   C – CE Mode
UNI Capabilities:
   Bits 10-28:  Reserved
   Bit 29:      MPLS Proxy Admission Control Protocol
   Bit 30:      SVC UNI
   Bit 31:      PVC UNI
```

If the MPLS Proxy Admission Control Protocol bit is set (=1), the SVC UNI and PVC UNI bits **must not** be set. The simultaneous operation of LSP Connection and MPLS Proxy Admission Control over the same MPLS UNI is for future study.

B.3.2 Hello Capabilities TLV

In order to distinguish among the messages of multiple capabilities operating over the same UNI connection, each MPLS Proxy Admission Control Protocol message includes a Hello Capabilities TLV. The Hello Capabilities TLV includes exactly one UNI Capability Element, and no other Capability Elements. The sender of the message sets the PE or CE bit as appropriate to its role, and sets the MPLS Proxy Admission Control Protocol UNI Capabilities bit. It **must not** set any other UNI Capabilities bits.

B.3.3 Forwarding Equivalence Classes

The MPLS Proxy Admission Control Protocol uses the Host Address (3) and Wildcard (1) FEC element types. It does not use the Prefix (2) and Opaque (4) types.

B.3.4 Label Request Message

There is always exactly one FEC element. The FEC type is always "Host Address" (3). There is always exactly one Traffic Parameters TLV. The LSPID TLV, Reverse LSPID TLV, and Reverse Traffic Parameters TLV are not used.

B.3.5 Label Mapping Message

There is always exactly one FEC element. The FEC type is always "Host Address" (3). There is always exactly one Traffic Parameters TLV. The LSPID TLV is not used. The Label Message Request ID is always included.

B.3.6 Label Release Message

There is always exactly one FEC element. The FEC type is always "Wildcard" (1). The Label TLV is required. There is always exactly one Traffic Parameters TLV.

B.3.7 Label Withdraw Message

There is always exactly one FEC element. The FEC type is always "Wildcard" (1). The Label TLV is required. There is always exactly one Traffic Parameters TLV.

B.3.8 Notification Messages

A notification message that is returned as a successful response to a Label Release message always includes exactly one Label TLV and one Traffic Parameters TLV.

B.4 Procedures

This section describes modifications to the MPLS UNI procedures to be used specifically for the MPLS Proxy Admission Control Protocol.

B.4.1 HELLO Message Procedure

A device **must** include a Hello Capabilities TLV in its HELLO Messages. The Hello Capabilities TLV **must** include one and only one UNI Capability Element. The device **must** set one and only one of the PE Mode and CE Mode bits in the UNI Capability Element.

A PE that provides the MPLS Proxy Admission Control Protocol sets the PE Mode bit in the UNI Capability Element and sets the MPLS Proxy Admission Control Protocol UNI Capabilities bit. A PE **must** ignore any received HELLO Message with the PE Mode bit set. It **should** ignore HELLO messages that do not have any supported UNI capabilities bits set.

A CE wishing to use the MPLS Proxy Admission Control Protocol sets the CE Mode bit in the UNI Capability Element and sets the MPLS Proxy Admission Control Protocol UNI Capabilities bit. A CE **must** ignore any received HELLO Message with the CE Mode bit set. It **should** ignore HELLO Messages if none of the desired UNI Capabilities bits are set.

Devices establish LDP sessions according to the rules of LDP Basic discovery, depending upon which HELLO messages have not been ignored.

B.4.2 Resource Index Label

For a granted reservation, a PE assigns a label (referred to as the Resource Index Label (RIL)) to the requesting CE. The CE encapsulates transmitted traffic to be counted against the reservation using normal MPLS encapsulation methods with the RIL as the MPLS label. On failure of the LDP session, the CE and PE **must** delete all associated RILs.

B.4.3 Making a Reservation

After LDP sessions are established, a CE that wishes to reserve capacity across the provider network towards another CE sends a Label Request message to the PE. In addition to the Hello Capabilities TLV included in all MPLS Proxy Admission Control Protocol messages, the Label Request message **must** include one (and only one) FEC TLV and one (and only one) Traffic Parameters TLV. The FEC TLV indicates the address of the remote CE towards which the reservation is to be made. The Traffic Parameters TLV indicates the desired capacity.

Conceptually, the receiving PE first searches for a network-side tunnel that is usable for the destination indicated in the FEC TLV and has sufficient available capacity to satisfy the Traffic Parameters TLV. If a suitable tunnel is not found, the PE returns a Notification message with the Status TLV set to "No Route" (0x0000000D), and the procedure ends.

If a suitable tunnel is found, the PE searches for a RIL that points to this tunnel and that has already been allocated to the CE. If no RIL is found, a new one is allocated, and the granted RIL capacity is set to zero. If no RIL can be allocated, the PE returns a Notification message with the Status TLV set to "No Label Resources" (0x0000000E) and the procedure ends.

Now that the PE has a tunnel and a RIL (new or existing), it subtracts the Traffic Parameters TLV from the available tunnel capacity and adds the Traffic Parameters TLV to the granted RIL capacity. It then returns a Label Mapping message with the contents set as follows:

- The FEC TLV is set to the value in the received Label Request message.
- The Label TLV is set to the RIL.
- The Label Request Message ID is set to the message ID of the received Label Request message.
- The Traffic Parameters TLV is set to the total capacity granted to this RIL.

When the CE receives the Label Mapping message, it checks to see if a RIL with the same value has already been allocated to it. If it has not, it creates a new allocation. In either case (new or existing RIL), the granted capacity is set to the value in the received Label Mapping Traffic Parameters TLV, and the FEC TLV is added to the set of destinations allowed for the RIL. If there is more than one destination associated with a RIL, the CE is free to distribute traffic among these destinations as it wishes, as long as it does not exceed the total granted capacity.

B.4.4 Transmitting Data

The CE indicates which reservations apply to its transmitted data packets by encapsulating those packets with RILs received from the PE, using the MPLS encapsulation method appropriate to the interface in use. The PE **may** police incoming labeled traffic to the reservation agreement, so the CE **should** ensure that its traffic is shaped properly. The PE **may** also examine the IP header of the encapsulated packet, and discard packets that would not be routed to the tunnel indicated by the incoming RIL.

If policing is performed, the traffic description parameters from the Traffic Parameters TLV in the Label Mapping message that granted the reservation are used to configure the chosen policing algorithm. For more information on choosing policing algorithms, see [B.2].

B.4.5 Adding Capacity

To add capacity to an existing reservation, the CE and PE follow the "Making a Reservation" procedure (section B.4.3). The Traffic Parameters TLV is set to the value of the additional capacity needed.

B.4.6 Releasing Capacity

To release (return) some or all of the capacity from an existing reservation, the CE sends a Label Release message. The FEC TLV is set to the Wildcard type (1). The Label TLV is required and is set to the value of the RIL from which capacity is to be released. The Traffic Parameters TLV is required and is set to the amount of the reservation capacity to be released.

If the PE accepts the message, it adds the Traffic Parameters TLV to the associated tunnel capacity, and subtracts the Traffic Parameters TLV from the capacity granted to the RIL. If the capacity granted to the RIL has gone to zero,

the PE **may** delete the RIL. Either way (remaining capacity is zero or not), the PE then returns a Notification message that includes a Traffic Parameters TLV and a Label TLV. The Status TLV is set to "Success" (0x00000000). The Label TLV is set to the RIL being released. The Traffic Parameters TLV is set to the now-current capacity granted to the RIL.

If the PE does not accept the message it returns a Notification message with the Status TLV set to indicate the error.

On receiving a Notification message that has the Status TLV set to "Success" (0x00000000), and that includes a Label TLV and a Traffic Parameters TLV, the CE sets the granted capacity for the RIL to the value in the Traffic Parameters TLV. If the granted capacity has gone to zero the CE **must** delete the RIL.

B.4.7 Withdrawing Reservations

If network events lead to the need for a PE to withdraw previously-granted reservations, it sends Label Withdraw messages for each withdrawn RIL. The FEC TLV is set to the Wildcard type (1). The Label TLV **must** be included and is set to the RIL being withdrawn. A Traffic Parameters TLV **must** be included and is set to the capacity being withdrawn from the label.

The receiving CE responds to a Label Withdraw message by releasing all of the capacity indicated in the Traffic Parameters TLV in the Label Withdraw message. To do this it follows the "Releasing Capacity" procedure (section B.4.6) with the Traffic Parameters TLV set to the full value given in the Label Withdraw message.

B.4.8 PE State Information

Conceptually, the PE needs to retain the following state information:

- For each RIL granted to each CE:
  - The associated (total) granted capacity.
- For each (network-side) tunnel:
  - The total capacity.
  - A list of RILs that are forwarded to the tunnel.
  - The currently available capacity (the total capacity minus the sum of the granted capacity for all RILs using the tunnel).

B.5 References

B.5.1 Normative References

[B.1] D. Sinicrope, A. Malis, MPLS PVC User to Network Interface, MPLS/FR Alliance 2.0.1, May 2003.

[B.2] A. Bhargava, T. Phelan, MPLS Proxy Admission Control Definition, MPLS/Frame Relay Alliance 6.0.0, January 2004.

B.5.2 Informative References

[B.3] E. Rosen, A. Viswanathan, R. Callon, Multiprotocol Label Switching Architecture, RFC3031, January 2001

[B.4] J. Boyle, V. Gill, A. Hanaan, D. Cooper, D. Awduche, B. Christian, W. S. Lai, "Applicability Statement For Traffic Engineering with MPLS", RFC3346, August 2002

[B.5] IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks – 802.1Q-1998