MPLS & Frame Relay Alliance

MPLS Proxy Admission Control Definition

Implementation Agreement

MPLS & Frame Relay Alliance 6.0.0

MPLS & Frame Relay Alliance Technical Committee
October 2004

Editors:

Anurag Bhargava
Ericsson IP Infrastructure

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

Copyright © 2004 MPLS & Frame Relay Alliance.

All rights reserved.

This document and translations of it may be copied and furnished to others, and works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the MPLS & Frame Relay Alliance, except as needed for the purpose of developing MPLS implementation agreements (in which case the procedures copyrights defined by the MPLS & Frame Relay Alliance must be followed), or as required to translate it into languages other than English

This document and the information contained herein is provided on an "AS IS" basis and THE MPLS & FRAME RELAY ALLIANCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Table of Contents

1 Introduction

Many real-time applications such as audio and video streaming, VoIP telephony and multimedia conferencing use protocols that do not take corrective action in response to network congestion. Traditionally, they rely on network over-provisioning to achieve reliable operation. There have been some attempts to improve this, but they have failed in one dimension or another. For example, classical RSVP [7] imposes an untenable burden on core routers, where the aggregation of reservations is high, while Diffserv [8] does not provide feedback to the customer equipment on the availability of network resources. Extending the MPLS network out to Customer Edge devices (CEs) is problematic when the network and the applications are owned by different organizations. The MPLS UNI LSP Connection feature improves this by creating a demarcation point that separates core label distribution from the edge. However, since the number of LSPs needed scales with the square of the number of CEs, there can be an explosion in the number of LSPs needed in large networks, burdening both the CEs and the network.

The MPLS Proxy Admission Control capability runs over the MPLS User-to-Network Interface (UNI) 2.0.1; it provides a highly scalable and lightweight method for customers of an MPLS Network service provider to dynamically reserve resources in the network for their private use. It solves the problem of scaling by removing the need to extend LSPs all the way to the CEs (in order to provide quality of service for real-time applications), and by allowing Provider Edge devices (PEs) to aggregate reservation requests. A service provider can offer this capability to its customers over the MPLS User-to-Network Interface (UNI), defined in.[4].

This document describes the operational and provisioning aspects of the MPLS Proxy Admission Control capability. It is one part of a two-document set. The companion document, *MPLS Proxy Admission Control Protocol* [5], describes the extensions to the MPLS UNI protocol necessary to implement the MPLS Proxy Admission Control capability.

1.1 Operational Environment

Consider the simple Traffic Engineered Multiprotocol Label Switched network (TE-MPLS, [1], and [2]) shown in Figure 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.

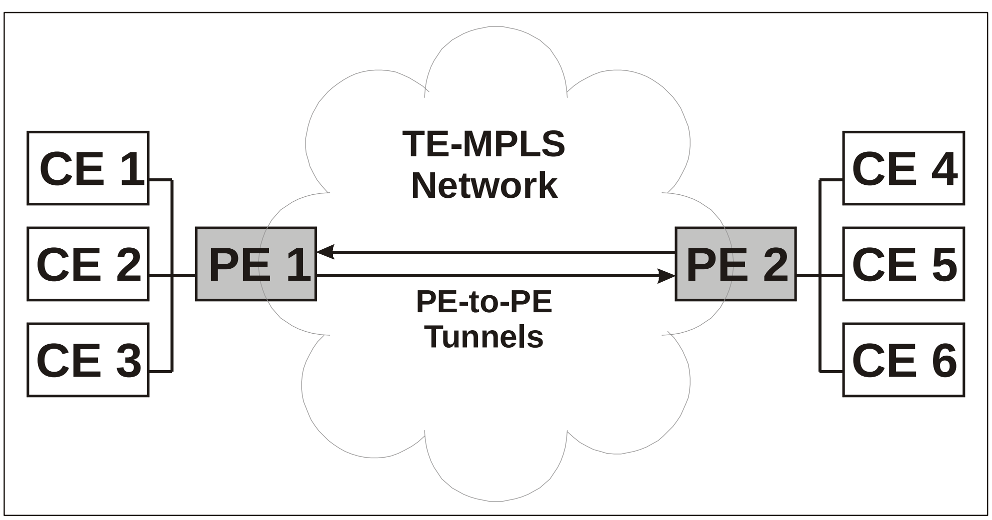

Figure 1: MPLS Proxy Admission Control environment.

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 own 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 great. 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 UNI defines an interface between PEs and CEs over which various capabilities can be offered. One of these capabilities is the LSP Connection feature. While this capability can supply CE-to-CE LSPs with QoS guarantees, problems develop when there are hundreds 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. It is offered over the MPLS UNI, allowing CEs to dynamically share network resources without the need for coordinated and static configuration of the CEs. The protocol mechanisms that support this capability are defined in the MPLS PVC UNI Implementation Agreement [4] and the MPLS Proxy Admission Control Protocol Implementation Agreement [5]. The simultaneous operation of the LSP Connection feature and the MPLS Proxy Admission Control capability over the same MPLS UNI is for further study.

1.2 Purpose

The purpose of this document is to describe the operational and provisioning aspects of the MPLS Proxy Admission Control capability.

1.3 Scope & Assumptions

The MPLS Proxy Admission Control capability provides bandwidth guarantees within the Provider Network only (from PE to PE). It does not provide any assurances for the ingress or egress networks. If an application requires end-to-end assurances, it must use additional techniques. Such additional techniques could include the use of layer 2 QoS technology such as 802.1Q/p [6], or other application-specific session layer controls.

There are PE-to-PE tunnels that participate in this feature; bandwidth is reserved over these tunnels. This feature, and the allocations of resources to CEs, is transparent to intermediate or core routers.

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.

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

2 Definitions and Terminology

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.

2.2 Acronyms

CE	Customer Edge device
FEC	Forwarding Equivalence Class
LDP	Label Distribution Protocol
LSP	Label Switched Path
LSR	Label Switching Router
MPLS	Multi Protocol Label Switching
PE	Provider Edge device
QoS	Quality of Service
RIL	Resource Index Label
SLA	Service Level Agreement
TLV	Type, Length, Value encoding
UNI	User-to-Network Interface

3 Reference Model

The reference model for the MPLS Proxy Admission Control capability is shown in Figure 2. The capability requires PEs to be connected by traffic-engineered tunnels. Source CEs interact with a local ingress PE via the MPLS UNI, and an Egress PE uses IP routing to forward packets to destination CEs. The availability of the MPLS Proxy Admission Control capability is automatically discovered by CEs, using the Hello procedures of the MPLS UNI. With the MPLS UNI, CEs make reservations for resources on PE-to-PE tunnels towards particular destination CEs. The reservations span the provider network, from ingress PE to egress PE, but do not include the access networks between the PEs and CEs.

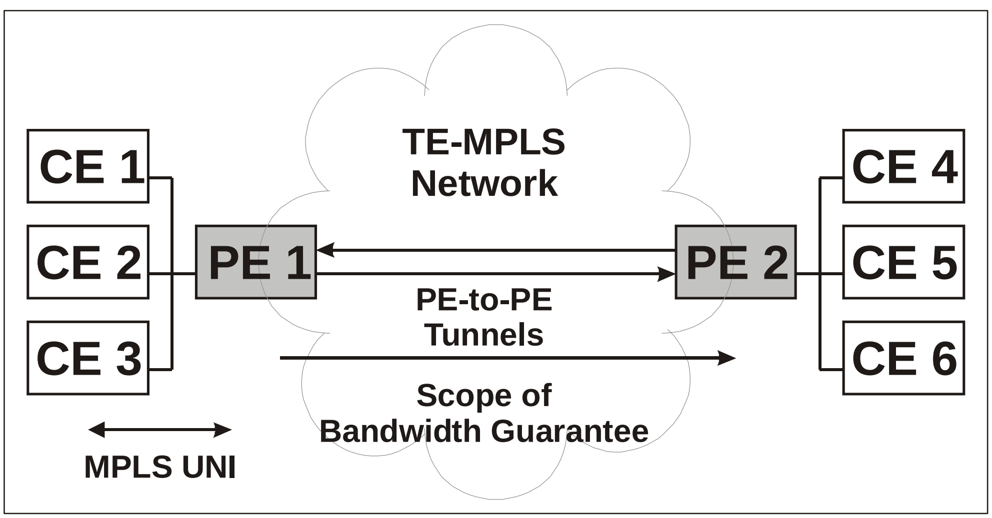

Figure 2: MPLS Proxy Admission reference model.

The MPLS Proxy Admission Control capability:

- Allows a CE to reserve network resources on demand on a PE-to-PE tunnel in the provider network towards a destination CE.
- Allows a CE to increase/decrease/release reservations on demand.
- Allows a CE to explicitly tag a packet flow to the corresponding reservation.
- Allows a PE to send notification to a CE if a request cannot be granted.
- Allows a PE to send notification to a CE if a granted reservation can no longer be supported.
- Allows a PE to aggregate requests and minimize the amount of reservation information maintained by the PE and CE.

Reservations are transparent to the core routers and egress PE in the provider network.

The use of the MPLS UNI with MPLS Proxy Admission Control capability is agreed to between a network service provider and a customer at subscription time.

3.1 Reservation Parameters

The MPLS Proxy Admission Control reservations are quantified by the following parameters:

- Peak data rate
- Peak burst size
- Committed data rate
- Committed burst size
- Excess burst size

A service provider may choose to guarantee other factors on a contractual (SLA) basis. These factors could include bounds on delay, jitter, packet loss, or commitments to certain packet treatments such as discard precedence or particular per-hop behaviors. These factors are not dynamically chosen or signaled per-reservation. The signaling of the above QoS parameters is for further study.

At subscription-time agreements may be made that limit the maximum capacity that can be reserved by the CE, and/or guarantee a minimum capacity that will be available to the CE.

Billing models are outside the scope of this document, but could include time-per-capacity models, and others.

4 Functional Operation

This describes the basic flow of operation. Protocol specifics are in the MPLS UNI Implementation Agreement [4] and the MPLS Proxy Admission Control Protocol Implementation Agreement [5].

4.1 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.

4.2 Protocol Flow

- Provisioning of this feature is done at subscription time. The willingness to provide, or desire to use, the feature is identified in the corresponding control protocol.
- A CE requests a reservation for a particular destination CE, providing a destination CE IP address and traffic parameters.

- A PE has to decide whether the request can be granted or not. If the PE accepts the request, it makes the corresponding reservation and returns a RIL to the ingress CE.
- The CE uses the RIL to mark the traffic stream so that it can be associated with the reservation. Unmarked packets to the same CE may be sent concurrently, but will be forwarded with best-effort behavior.
- A PE may return the same RIL for different reservation requests if the same PE-to-PE tunnel will serve the traffic for each reservation. The traffic parameters for the RIL will be the sum of all of the requests. The CE may distribute the aggregate traffic load among the various destinations as it wishes.
- The ingress CE may request modification of the reservation at any time. The reservation modification is subject to acceptance by the PE.
- Any topological change in the provider network that affects a granted request should result in a notification to the CE, and the withdrawal of the affected RIL. The CE must send a new request to reestablish the reservation. There is no need to notify a CE for changes in tunnel routing (such as Fast Reroute [9]) that do not affect the tunnel capacity or reachability. Any changes beyond the PE network are outside the scope of this function, and should not be reflected in any notifications to the CE.
- The PE may choose to police incoming RIL-labeled traffic using the RIL to determine the reservation contract. The algorithm used must be chosen at subscription time by mutual agreement between customer and provider. The CE should shape its transmissions to the reservation parameters (e.g. committed and peak data rates and burst sizes). The recommended policing algorithm is the two-rate three-color marker algorithm described in [11]. Other policing algorithms, such as [10] and [12], may be chosen as alternatives.
- On the egress PE-CE interface, the egress PE may provide differential treatment for MPLS Proxy Admission Control packets by taking additional actions using various mechanisms such as traffic classification.

5 References

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

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

[3] L. Anderson, P. Doolan, H. Feldman, A. Fredette, B. Thomas, LDP Specification, RFC 3036, January 2001

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

[5] Phelan, T., MPLS Proxy Admission Control Protocol, MPLS/FR Alliance 7.0.0

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

[7] Braden R., Zhang, L., Berson, S., Herzog, S., Jamin, S., Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification, RFC2205, September 1997

[8] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W., An Architecture for Differentiated Services, RFC2475, December 1998

[9] Pan, P., et al, Fast Reroute Extensions to RSVP-TE for LSP Tunnels, December 2003, draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt, work in progress.

[10] J. Heinanen, R. Guerin, A Single Rate Three Color Marker, RFC 2697, September 1999.

[11] J. Heinanen, R. Guerin, A Two Rate Three Color Marker, RFC 2698, September 1999.

[12] W. Fang, N Seddigh, B. Nandy, A Time Sliding Window Three Color Marker (TSWTCM), RFC 2859, June 2000.