Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 13, 2014

W. George, Ed.
Time Warner Cable
C. Pignataro
R. Asati
K. Raza
Cisco Systems
R. Bonica
Juniper Networks
R. Papneja
D. Dhody
Huawei Technologies
V. Manral
Hewlett-Packard, Inc.
July 12, 2013

Gap Analysis for Operating IPv6-only MPLS Networks
draft-george-mpls-ipv6-only-gap-01

Abstract

This document reviews the MPLS protocol suite in the context of IPv6 and identifies gaps that must be addressed in order to allow MPLS-related protocols and applications to be used with IPv6-only networks. This document is not intended to highlight a particular vendor's implementation (or lack thereof) in the context of IPv6-only MPLS functionality, but rather to focus on gaps in the standards defining the MPLS suite.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 13, 2014.

Table of Contents

1. Introduction

IPv6 is an integral part of modern network deployments. At the time when this document was written, the majority of these IPv6 deployments were using dual-stack implementations, where IPv4 and IPv6 are supported equally on many or all of the network nodes, and single-stack primarily refers to IPv4-only devices. Dual-stack deployments provide a useful margin for protocols and features that are not currently capable of operating solely over IPv6, because they can continue using IPv4 as necessary. However, as IPv6 deployment and usage becomes more pervasive, and IPv4 exhaustion begins driving changes in address consumption behaviors, there is an increasing likelihood that many networks will need to start operating some or all of their network nodes either as primarily IPv6 (most functions use IPv6, a few legacy features use IPv4), or as IPv6-only (no IPv4 provisioned on the device). This transition toward IPv6-only operation exposes any gaps where features, protocols, or implementations are still reliant on IPv4 for proper function. To that end, and in the spirit of RFC 6540's [RFC6540] recommendation that implementations need to stop requiring IPv4 for proper and complete function, this document reviews the MPLS protocol suite in the context of IPv6 and identifies gaps that must be addressed in order to allow MPLS-related protocols and applications to be used with IPv6-only networks. This document is not intended to highlight a particular vendor's implementation (or lack thereof) in the context of IPv6-only MPLS functionality, but rather to focus on gaps in the standards defining the MPLS suite.

2. Use Case

From a purely theoretical perspective, ensuring that MPLS is fully IP version-agnostic is the right thing to do. However, it is sometimes helpful to understand the underlying drivers that make this work necessary to undertake, especially at a time when IPv6-only networking is still fairly uncommon. This section will discuss some drivers. It is not intended to be a comprehensive discussion of all potential use cases, but rather a discussion of at least one use case so that this is not seen as solving a purely theoretical problem.

IP convergence is continuing to drive new classes of devices to begin communicating via IP. Examples of such devices could include set top boxes for IP Video distribution, cell tower electronics (macro or micro cells), infrastructure Wi-Fi Access Points, and devices for machine to machine (M2M) or Internet of Things applications. In some cases, these classes of devices represent a very large deployment base, on the order of thousands or even millions of devices network-wide. The scale of these networks, coupled with the increasingly

overlapping use of RFC 1918 [RFC1918] address space within the average network, and the lack of globally-routable IPv4 space available for long-term growth begins to drive the need for many of the endpoints in this network to be managed solely via IPv6. Even if these devices are carrying some IPv4 user data, it is often encapsulated in another protocol such that the communication between the endpoint and its upstream devices can be IPv6-only without impacting support for IPv4 on user data. Depending on the MPLS features required, it is plausible to assume that the (existing) MPLS network may need to be extended to these devices.

Additionally, as the impact of IPv4 exhaustion becomes more acute, more and more aggressive IPv4 address reclamation measures will be justified. Measures that were previously seen as too complex or as netting too few addresses for the work required may become more realistic as the cost for obtaining new IPv4 addresses increases. More and more networks are likely to adopt the general stance that IPv4 addresses need to be preserved for revenue-generating customers so that legacy support for IPv4 can be maintained as long as possible. As a result, it may be appropriate for some or all of the network infrastructure, including MPLS LSRs and LERs, to have its IPv4 addresses reclaimed and transition toward IPv6-only operation.

3. Gap Analysis

This gap analysis aims to answer the question, "what breaks when one attempts to use MPLS features on a network of IPv6-only devices?" The assumption is that some endpoints as well as LSRs (PE and P routers) only have IPv6 transport available, and need to support the full suite of MPLS features defined as of the time of this document's writing at parity with the support on an IPv4 network. This is necessary whether they are enabled via LDP RFC 5036 [RFC5036], RSVP-TE RFC 5420 [RFC5420], or BGP RFC 3107 [RFC3107], and whether they are encapsulated in MPLS RFC 3032 [RFC3032], IP RFC 4023 [RFC4023], GRE RFC 4023 [RFC4023], or L2TPv3 RFC 4817 [RFC4817]. It is important when evaluating these gaps to distinguish between user data and control plane data, because while this document is focused on IPv6-only operation, it is quite likely that some amount of the user payload data being carried in the IPv6-only MPLS network will still be IPv4.

3.1. MPLS Data Plane

MPLS labeled packets can be transmitted over a variety of data links RFC 3032 [RFC3032], and MPLS labeled packets can also be encapsulated over IP. The encapsulations of MPLS in IP and Generic Routing Encapsulation (GRE) as well as MPLS over Layer 2 Tunneling Protocol

Version 3 (L2TPv3) support IPv6. See Section 3 of RFC 4023 [RFC4023] and Section 2 of RFC 4817 [RFC4817] respectively.

3.2. MPLS Control Plane

3.2.1. LDP

Label Distribution Protocol (LDP) RFC 5036 [RFC5036] defines a set of procedures for distribution of labels between label switch routers that can use the labels for forwarding traffic. While LDP was designed to use an IPv4 or dual-stack IP network, it has a number of deficiencies that prohibit it from working in an IPv6-only network. LDP-IPv6 [I-D.ietf-mpls-ldp-ipv6] highlights some of the deficiencies when LDP is enabled in IPv6 only or dual-stack networks, and specifies appropriate protocol changes. These deficiencies are related to LSP mapping, LDP identifiers, LDP discovery, LDP session establishment, next hop address and LDP TTL security RFC 5082 [RFC5082].

3.2.2. Multicast LDP

Multipoint LDP (mLDP) is a set of extensions to LDP for setting up Point to Multipoint (P2MP) and Multipoint to Multipoint (MP2MP) LSPs. These extensions are specified in RFC 6388 [RFC6388]. In terms of IPv6-only gap analysis, mLDP has two identified areas of interest:

1. LDP Control plane: Since mLDP uses the LDP control plane to discover and establish sessions with the peer, it shares the same gaps as LDP with regards to control plane (discovery, transport, and session establishment) in an IPv6-only network.

2. Multipoint (MP) FEC Root address: mLDP defines its own MP FECs and rules, different from LDP, to map MP LSPs. mLDP MP FEC contains a Root Address field which is an IP address in IP networks. The current specification allows specifying Root address according to AFI and hence covers both IPv4 or IPv6 root addresses, requiring no extension to support IPv6-only MP LSPs. The root address is used by each LSR participating in an MP LSP setup such that root address reachability is resolved by doing a table lookup against root address to find corresponding upstream neighbor(s). This will pose a problem when an MP LSP traverses islands of IPv4 and IPv4 clouds on the way to the root node.

For example, consider following setup, where R1/R6 are IPv4-only, R3/R4 are IPv6-only, and R2/R5 are dual-stack LSRs:

```
( IPv4-only )  (  IPv6-only   )  ( IPv4-only )
       R1 -- R2 -- R3 -- R4 -- R5 -- R6
       Leaf                          Root
```

Assume R1 to be a leaf node for an P2MP LSP rooted at R6 (root node). R1 uses R6's IPv4 address as the Root address in MP FEC. As the MP LSP signaling proceeds from R1 to R6, the MP LSP setup will fail on the first IPv6-only transit/branch LSRs (R3) when trying to find IPv4 root address reachability. RFC 6512 [RFC6512] defines a recursive-FEC solution and procedures for mLDP when the backbone (transit/branch) LSRs have no route to the root. The proposed solution is defined for a BGP-free core in an VPN environment, but the similar concept can be used/extended to solve the above issue of IPv6-only backbone receiving an MP FEC element with an IPv4 address. The solution will require a border LSR (the one which is sitting on border of an IPv4/IPv6 island(s) (R2 and R5) to translate an IPv4 root address to equivalent IPv6 address (and vice vera) through the procedures similar to RFC6512. The translation of root address on borders of IPv4 or IPv6 islands will also be needed for recursive FECs and procedures defined in RFC6512.

3.2.3. RSVP- TE

Resource Reservation Protocol Extensions for MPLS Traffic Engineering (RSVP-TE) RFC 3209 [RFC3209] defines a set of procedures & enhancements to establish label-switched tunnels that can be automatically routed away from network failures, congestion, and bottlenecks. RSVP-TE allows establishing an LSP for an IPv4 or IPv6 prefix, thanks to its LSP_TUNNEL_IPv6 object and subobjects.

3.2.3.1. IGP

RFC3630 [RFC3630] specifies a method of adding traffic engineering capabilities to OSPF Version 2. New TLVs and sub-TLVs were added in RFC5329 [RFC5329] to extend TE capabilities to IPv6 networks in OSPF Version 3.

RFC5305 [RFC5305] specifies a method of adding traffic engineering capabilities to IS-IS. New TLVs and sub-TLVs were added in RFC6119 [RFC6119] to extend TE capabilities to IPv6 networks.

3.2.3.2. RSVP-TE-P2MP

RFC4875 [RFC4875] describes extensions to RSVP-TE for the setup of point-to-multipoint (P2MP) LSPs in MPLS and GMPLS with support for both IPv4 and IPv6.

3.2.3.3. RSVP-TE Fast Reroute (FRR)

RFC4090 [RFC4090] specifies FRR mechanisms to establish backup LSP tunnels for local repair supporting both IPv4 and IPv6 networks. Further RFC5286 [RFC5286] describes the use of loop-free alternates to provide local protection for unicast traffic in pure IP and MPLS networks in the event of a single failure, whether link, node, or shared risk link group (SRLG) for both IPv4 and IPv6.

3.2.4. Controller, PCE

The Path Computation Element (PCE) defined in RFC4655 [RFC4655] is an entity that is capable of computing a network path or route based on a network graph, and applying computational constraints. A Path Computation Client (PCC) may make requests to a PCE for paths to be computed. The PCE communication protocol (PCEP) is designed as a communication protocol between PCCs and PCEs for path computations and is defined in RFC5440 [RFC5440].

The PCEP specification RFC5440 [RFC5440] is defined for both IPv4 and IPv6 with support for PCE discovery via an IGP (OSPF RFC5088 [RFC5088], or ISIS RFC5089 [RFC5089]) using both IPv4 and IPv6 addresses. Note that PCEP uses identical encoding of subobjects as in the Resource Reservation Protocol Traffic Engineering Extensions (RSVP-TE) defined in RFC3209 [RFC3209] which supports both IPv4 and IPv6.

The extensions of PCEP to support confidentiality RFC5520 [RFC5520], Route Exclusion RFC5521, [RFC5521] Monitoring RFC5886 [RFC5886], and P2MP RFC6006 [RFC6006] have support for both IPv4 and IPv6.

3.2.5. BGP

RFC3107 [RFC3107] specifies a set of BGP protocol procedures for distributing the labels (for prefixes corresponding to any address-family) between label switch routers so that they can use the labels for forwarding the traffic. RFC3107 allows BGP to distribute the label for IPv4 or IPv6 prefix in an IPv6 only network.

3.2.6. GMPLS

RFC4558 [RFC4558] specifies Node-ID Based RSVP Hello Messages with capability for both IPv4 and IPv6. RFC4990 [RFC4990] clarifies the use of IPv6 addresses in GMPLS networks including handling in the MIB modules.

3.3. MPLS Applications

3.3.1. L2VPN

L2VPN RFC 4664 [RFC4664] specifies two fundamentally different kinds of Layer 2 VPN services that a service provider could offer to a customer: Virtual Private Wire Service (VPWS) and Virtual Private LAN Service (VPLS). RFC 4447 [RFC4447] and RFC 4762 [RFC4762] specify the LDP protocol changes to instantiate VPWS and VPLS services respectively in an MPLS network using LDP as the signaling protocol. This is complemented by RFC 6074 [RFC6074], which specifies a set of procedures for instantiating L2VPNs (e.g. VPWS, VPLS) using BGP as discovery protocol and LDP as well as L2TPv3 as signaling protocol. RFC 4761 [RFC4761] and RFC 6624 [RFC6624] specify BGP protocol changes to instantiate VPLS and VPWS services in an MPLS network, using BGP for both discovery and signaling.

In an IPv6-only MPLS network, use of L2VPN represents connection of Layer 2 islands over an IPv6 MPLS core, and very few changes are necessary to support operation over an IPv6-only network. The L2VPN signaling protocol is either BGP or LDP in an MPLS network, and both can run directly over IPv6 core infrastructure, as well as IPv6 edge devices. RFC 6074 [RFC6074] is the only RFC that appears to have a gap wrt IPv6. In its discovery procedures (section 3.2.2 and section 6), it suggests encoding PE IP address in the VSI-ID, which is encoded in NLRI, which should not exceed 12 bytes (to differentiate its AFI/SAFI encoding from RFC4761). This means that PE IP address can NOT be an IPv6 address. Also, in its signaling procedures (section 3.2.3), it suggests encoding PE_addr in SAII and TAII, which are limited to 32-bit (AII Type=1) at the moment.

3.3.1.1. EVPN

EVPN [I-D.ietf-l2vpn-evpn] is still a work in progress. As such, it is out of scope for this gap analysis. Instead, the authors of that draft need to ensure that it supports IPv6-only operation, or if it cannot, identify dependencies on underlying gaps in MPLS protocol(s) that must be resolved before it can support IPv6-only operation.

3.3.2. L3VPN

RFC 4364 [RFC4364] defines a method by which a Service Provider may use an IP backbone to provide IP Virtual Private Networks (VPNs) for its customers. The following use cases arise in the context of this gap analysis:

1. Connecting IPv6 islands over IPv6-only MPLS network

2. Connecting IPv4 islands over IPv6-only MPLS network

Both use cases 1 and 2 require mapping an IP packet to an IPv6-signaled LSP to the remote PE, which is not explicitly defined in any RFC. RFC4364 has two MAJOR gaps. First, it is not possible to use an IPv6-only MPLS network, since RFC4364 explicitly assumes IPv4-only MPLS network i.e. BGP Next Hop is assumed to have /32 (for example, see section 5 of RFC4364]. Second, it is limited to VPN-IPv4 address-family i.e. connecting IPv4 islands over IPv4-only MPLS networks. This second gap has been fixed by 6VPE RFC 4659 [RFC4659], which defines connecting IPv6 VPN sites over an IPv4-only MPLS networks, but more work is needed to address the first gap.

The authors do not believe that there are any additional issues encountered when using L2TPv3, RSVP, or GRE (instead of LDP) as transport on an IPv6-only network.

3.3.2.1. 6PE/4PE

RFC 4798 [RFC4798] defines 6PE, which defines how to interconnect IPv6 islands over a Multiprotocol Label Switching (MPLS)-enabled IPv4 cloud. However, use case 2 is doing the opposite, and thus could also be referred to as 4PE. The method to support this use case is not defined explicitly. To support it, IPv4 edge devices need to be able to map IPv4 traffic to MPLS IPv6 core LSP's. Also, the core switches may not understand IPv4 at all, but in some cases they may need to be able to exchange Labeled IPv4 routes from one AS to a neighboring AS.

3.3.2.2. 6VPE/4VPE

RFC 4659 [RFC4659] defines 6VPE, a method by which a Service Provider may use its packet-switched backbone to provide Virtual Private Network (VPN) services for its IPv6 customers. It allows the core network to be MPLS IPv4 or MPLS IPv6, thus addressing use case 1 above. RFC4364 should work as defined for use case 2 above, which could also be referred to as 4VPE, but the RFC does not explicitly discuss this use.

3.3.2.3. BGP Encapsulation SAFI

RFC 5512 [RFC5512] defines the BGP Encapsulation SAFI and the BGP Tunnel Encapsulation Attribute, which can be used to signal tunnelling over an single-Address Family IP core. This mechanism supports transport of MPLS (and other protocols) over Tunnels in an IP core (including an IPv6-only core). In this context, load-balancing can be provided as specified in RFC 5640 [RFC5640].

3.3.2.4. NG-MVPN

TBD RFC 6513 both IPv4 and IPv6 multicast payload traffic

No IP version considerations?

3.3.3. MPLS-TP

***TBD RFC 6371 *** MPLS-TP does not require IP ("and network operation in the absence of a dynamic > control plane or IP forwarding support." RFC 5921) and thus should not be affected by operation on an IPv6-only network.

3.4. MPLS OAM

For MPLS LSPs, there are primarily three OAM mechanisms: Extended ICMP RFC 4884 [RFC4884] RFC 4950 [RFC4950], LSP Ping RFC 4379 [RFC4379], and BFD for MPLS LSPs RFC 5884 [RFC5884]. For MPLS Pseudowires, there is also Virtual Circuit Connectivity Verification (VCCV) RFC 5085 [RFC5085] RFC 5885 [RFC5885]. All of these mechanisms work in pure IPv6 environments. The next subsections cover these in detail.

3.4.1. Extended ICMP

Extended ICMP to support Multi-part messages is defined in RFC 4884 [RFC4884]. This extensibility is defined generally for both ICMPv4 and ICMPv6. The specific ICMP extensions for MPLS are defined in RFC 4950 [RFC4950]. ICMP Multi-part with MPLS extensions works for IPv4 and IPv6. However, the mechanisms described in RFC 4884 and 4950 may fail when tunneling IPv4 traffic over an LSP that is supported by IPv6-only infrastructure.

Assume the following:

- the path between two IPv4 only hosts contains an MPLS LSP
- the two routers that terminate the LSP run dual stack
- the LSP interior routers run IPv6 only
- the LSP is signaled over IPv6

Now assume that one of the hosts sends an IPv6 packet to the other. However, the packet's TTL expires on an LSP interior router. According to RFC 3032 [RFC3032], the interior router should examine the IPv6 payload, format an ICMPv6 message, and send it (over the tunnel upon which the original packet arrived) to the egress LSP. In

this case, however, the LSP interior router is not IPv6-aware. It cannot parse the original IPv6 datagram, nor can it send an IPv6 message. So, no ICMP message is delivered to the source. Some specific ICMP extensions, in particular ICMP Extensions for Interface and Next-Hop Identification RFC 5837 [RFC5837] restrict the address family of address information included in an Interface Information Object to the same one as the ICMP (see Section 4.5 of RFC 5837). While these extensions are not MPLS specific, they can be used with MPLS packets carrying IP datagrams. This has no implications for IPv6-only environments.

3.4.2. LSP Ping

The LSP Ping mechanism defined in RFC 4379 [RFC4379] is specified to work with IPv6. Specifically, the Target FEC Stacks include both IPv4 and IPv6 versions of all FECs (see Section 3.2 of RFC 4379). The only exceptions are the Pseudowire FECs later specified for IPv6 in RFC 6829 [RFC6829]. Additionally, LSP Ping packets are UDP packets over both IPv4 and IPv6 (see Section 4.3 of RFC 4379). The multipath information includes also IPv6 encodings (see Section 3.3.1 of RFC 4379). However, the mechanisms described in RFC 4379 may fail when tunneling IPv4 traffic over an LSP that is supported by IPv6-only infrastructure.

Assume the following:

- LSP Ping is operating in traceroute mode over an MPLS LSP
- the two routers that terminate the LSP run dual stack
- the LSP interior routers run IPv6 only
- the LSP is signaled over IPv6

Packets will expire at LSP interior routers. According to RFC 4379, the interior router must parse the IPv4 Echo Request, and then, send an IPv4 Echo Reply. However, the LSP interior router is not IPv4-aware. It cannot parse the IPv4 Echo Request, nor can it send an IPv4 Echo Reply. So, no reply is sent.

3.4.3. BFD

The BFD specification for MPLS LSPs RFC 5884 [RFC5884] is defined for IPv4 as well as IPv6 versions of MPLS FECs (see Section 3.1 of RFC 5884). Additionally the BFD packet is encapsulated over UDP and specified to run over both IPv4 and IPv6 (see Section 7 of RFC 5884).

3.4.4. Pseudowires

The OAM specifications for MPLS Pseudowires define usage for both IPv4 and IPv6. Specifically, VCCV RFC 5085 [RFC5085] can carry IPv4 or IPv6 OAM packets (see Section 5.1.1 and 5.2.1 of RFC 5085), and VCCV for BFD RFC 5885 [RFC5885] also defines an IPv6 encapsulation (see Section 3.2 of RFC 5885).

3.4.5. MPLS-TP OAM

*** TBD***

3.5. MIBs

RFC3811 [RFC3811] defines the textual conventions for MPLS. These lack support for IPv6 in defining MplsExtendedTunnelId and MplsLsrIdentifier. These textual conventions are used in the MPLS TE MIB specification RFC3812 [RFC3812], GMPLS TE MIB specification RFC4802 [RFC4802] and Fast ReRoute (FRR) extension RFC6445 [RFC6445]. 3811bis [I-D.manral-mpls-rfc3811bis] tries to resolve this gap by marking this textual convention as obsolete.

The other MIB specifications for LSR RFC3813 [RFC3813], LDP RFC3815 [RFC3815] and TE RFC4220 [RFC4220] have support for both IPv4 and IPv6.

4. Gap Summary

This draft has reviewed a wide variety of MPLS features and protocols to determine their suitability for use on IPv6-only networks. While some parts of the MPLS suite will function properly without additional changes, gaps have been identified in others, which will need to be addressed with follow-on work. This section will summarize those gaps, along with pointers to any work-in-progress to address them.

Identifed gaps in MPLS for IPv6-only networks

Item	Gap	Addressed in
LDP	LSP mapping, LDP identifiers, LDP discovery, LDP session establishment, next hop address and LDP TTL security	LDP-IPv6 [I-D.ietf-mpls-ldp-ipv6]
L2VPN	RFC 6074 [RFC6074] discovery, signaling	TBD
L3VPN	RFC 4364 [RFC4364] BGP next-hop, define method for 4PE/4VPE	TBD
MIBs	RFC 3811 [RFC3811] no IPv6 textual convention	3811bis [I-D.manral-mpls-rfc3811bis]

Table 1: IPv6-only MPLS Gaps

5. Acknowledgements

This draft is brought to you by the letters I, P, V, and the number 6.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

Changing the address family used for MPLS network operation does not fundamentally alter the security considerations currently extant in any of the specifics of the protocol or its features. However, the change does expose the network and protocol to some of the IPv6-specific security considerations inherent to IPv6 itself as documented in [list of RFCs?]

8. Informative References

[I-D.ietf-l2vpn-evpn]
Sajassi, A., Aggarwal, R., Henderickx, W., Balus, F., Isaac, A., and J. Uttaro, "BGP MPLS Based Ethernet VPN", draft-ietf-l2vpn-evpn-03 (work in progress), February 2013.

[I-D.ietf-mpls-ldp-ipv6]
Asati, R., Manral, V., Papneja, R., and C. Pignataro, "Updates to LDP for IPv6", draft-ietf-mpls-ldp-ipv6-08 (work in progress), February 2013.

[I-D.manral-mpls-rfc3811bis]
Manral, V., Tsou, T., Liu, W., and F. Fondelli, "Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management", draft-manral-mpls-rfc3811bis-03 (work in progress), June 2013.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, May 2001.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC3811] Nadeau, T. and J. Cucchiara, "Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management", RFC 3811, June 2004.

[RFC3812] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)", RFC 3812, June 2004.

[RFC3813] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)", RFC 3813, June 2004.

[RFC3815] Cucchiara, J., Sjostrand, H., and J. Luciani, "Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)", RFC 3815, June 2004.

[RFC4023] Worster, T., Rekhter, Y., and E. Rosen, "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[RFC4220] Dubuc, M., Nadeau, T., and J. Lang, "Traffic Engineering Link Management Information Base", RFC 4220, November 2005.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.

[RFC4558] Ali, Z., Rahman, R., Prairie, D., and D. Papadimitriou, "Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement", RFC 4558, June 2006.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

[RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, September 2006.

[RFC4664] Andersson, L. and E. Rosen, "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006.

[RFC4761] Kompella, K. and Y. Rekhter, "Virtual Private LAN Service

(VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007.

[RFC4762] Lasserre, M. and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007.

[RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, February 2007.

[RFC4802] Nadeau, T. and A. Farrel, "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", RFC 4802, February 2007.

[RFC4817] Townsley, M., Pignataro, C., Wainner, S., Seely, T., and J. Young, "Encapsulation of MPLS over Layer 2 Tunneling Protocol Version 3", RFC 4817, March 2007.

[RFC4875] Aggarwal, R., Papadimitriou, D., and S. Yasukawa, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.

[RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, April 2007.

[RFC4950] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "ICMP Extensions for Multiprotocol Label Switching", RFC 4950, August 2007.

[RFC4990] Shiomoto, K., Papneja, R., and R. Rabbat, "Use of Addresses in Generalized Multiprotocol Label Switching (GMPLS) Networks", RFC 4990, September 2007.

[RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, October 2007.

[RFC5085] Nadeau, T. and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.

[RFC5088] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang,

"OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008.

[RFC5089] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, January 2008.

[RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, September 2008.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

[RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, September 2008.

[RFC5420] Farrel, A., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", RFC 5512, April 2009.

[RFC5520] Bradford, R., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", RFC 5520, April 2009.

[RFC5521] Oki, E., Takeda, T., and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", RFC 5521, April 2009.

[RFC5640] Filsfils, C., Mohapatra, P., and C. Pignataro, "Load-Balancing for Mesh Softwires", RFC 5640, August 2009.

[RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010.

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

[RFC5885] Nadeau, T. and C. Pignataro, "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010.

[RFC5886] Vasseur, JP., Le Roux, JL., and Y. Ikejiri, "A Set of Monitoring Tools for Path Computation Element (PCE)-Based Architecture", RFC 5886, June 2010.

[RFC6006] Zhao, Q., King, D., Verhaeghe, F., Takeda, T., Ali, Z., and J. Meuric, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic Engineering Label Switched Paths", RFC 6006, September 2010.

[RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", RFC 6074, January 2011.

[RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", RFC 6119, February 2011.

[RFC6388] Wijnands, IJ., Minei, I., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, November 2011.

[RFC6445] Nadeau, T., Koushik, A., and R. Cetin, "Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base for Fast Reroute", RFC 6445, November 2011.

[RFC6512] Wijnands, IJ., Rosen, E., Napierala, M., and N. Leymann, "Using Multipoint LDP When the Backbone Has No Route to the Root", RFC 6512, February 2012.

[RFC6540] George, W., Donley, C., Liljenstolpe, C., and L. Howard, "IPv6 Support Required for All IP-Capable Nodes", BCP 177, RFC 6540, April 2012.

[RFC6624] Kompella, K., Kothari, B., and R. Cherukuri, "Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling", RFC 6624, May 2012.

[RFC6829] Chen, M., Pan, P., Pignataro, C., and R. Asati, "Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6", RFC 6829, January 2013.

Appendix A. Assignments

RFC EDITOR PLEASE REMOVE BEFORE PUBLISHING

This will track which author volunteered for which section(s):

OAM - Ron Bonica, Carlos Pignataro

LDP/mLDP (multicast) - Kamran Raza

L2VPN - Rajiv Asati, Vishwas Manral, Rajiv Papneja

L3VPN - Rajiv Asati, Vishwas Manral, Rajiv Papneja

PCE - Dhruv Dhody, Rajiv Papneja

Editors- Wes George(primary), Vishwas Manral, Rajiv Asati

Authors' Addresses

Wesley George (editor)
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20111
US

Phone: +1-703-561-2540
Email: wesley.george@twcable.com

Carlos Pignataro
Cisco Systems
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
US

Phone:
Email: cpignata@cisco.com

Rajiv Asati
Cisco Systems
7025 Kit Creek Road
Research Triangle Park, NC 27709
US

Phone:
Email: rajiva@cisco.com

Kamran Raza
Cisco Systems
2000 Innovation Drive
Ottawa, ON K2K-3E8
CA

Phone:
Email: skraza@cisco.com

Ronald Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20171
US

Phone:
Email: rbonica@juniper.net

Rajiv Papneja
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
US

Phone:
Email: rajiv.papneja@huawei.com

Dhruv Dhody
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
US

Phone:
Email: dhruv.dhody@huawei.com

Vishwas Manral
Hewlett-Packard, Inc.
19111 Pruneridge Ave.
Cupertino, CA 95014
US

Phone:
Email: vishwas.manral@hp.com