Internet Engineering Task Force (IETF)
Request for Comments: 7439
Category: Informational
ISSN: 2070-1721

W. George, Ed.
Time Warner Cable
C. Pignataro, Ed.
Cisco

January 2015

Gap Analysis for Operating IPv6-Only MPLS Networks

Abstract

This document reviews the Multiprotocol Label Switching (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 intended to focus on gaps in the standards defining the MPLS suite, and is not intended to highlight particular vendor implementations (or lack thereof) in the context of IPv6-only MPLS functionality.

In the data plane, MPLS fully supports IPv6, and MPLS labeled packets can be carried over IPv6 packets in a variety of encapsulations. However, support for IPv6 among MPLS control-plane protocols, MPLS applications, MPLS Operations, Administration, and Maintenance (OAM), and MIB modules is mixed, with some protocols having major gaps. For most major gaps, work is in progress to upgrade the relevant protocols.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7439.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

IPv6 [RFC2460] 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 referred 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 the recommendation in RFC 6540 [RFC6540] 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 and networks that are primarily IPv6 (hereafter referred to as IPv6-primary). This document is intended to focus on gaps in the standards defining the MPLS suite, and not to highlight particular vendor implementations (or lack thereof) in the context of IPv6-only MPLS functionality.

2. Use Case

This section discusses some drivers for ensuring that MPLS completely supports IPv6-only operation. It is not intended to be a comprehensive discussion of all potential use cases, but rather a discussion of one use case to provide context and justification to undertake such a gap analysis.

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 (IoT) 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. As the number of devices to manage increases, the operator is compelled to move to IPv6. Depending on the MPLS features required, it is plausible to assume that the (existing) MPLS network will need to be extended to these IPv6-only devices.

Additionally, as the impact of IPv4 exhaustion becomes more acute, more and more aggressive IPv4 address reclamation measures will be justified. Many networks are likely to focus on preserving their remaining IPv4 addresses for revenue-generating customers so that legacy support for IPv4 can be maintained as long as necessary. As a result, it may be appropriate for some or all of the network infrastructure, including MPLS Label Switching Routers (LSRs) and Label Edge Routers (LERs), to have its IPv4 addresses reclaimed and transition toward IPv6-only operation.

3. Gap Analysis

This gap analysis aims to answer the question of what fails when one attempts to use MPLS features on a network of IPv6-only devices. The baseline assumption for this analysis is that some endpoints, as well as Label Switching Routers (Provider Edge (PE) and Provider (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 the Label Distribution Protocol (LDP) [RFC5036], RSVP - Traffic Engineering (RSVP-TE) [RFC3209], or Border Gateway Protocol (BGP) [RFC3107], and whether they are encapsulated in MPLS [RFC3032], IP [RFC4023], Generic Routing Encapsulation (GRE) [RFC4023], or Layer 2 Tunneling Protocol Version 3 (L2TPv3) [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.

A note about terminology: Gaps identified by this document are characterized as "Major" or "Minor". Major gaps refer to significant changes necessary in one or more standards to address the gap due to existing standards language having either missing functionality for IPv6-only operation or explicit language requiring the use of IPv4 with no IPv6 alternatives defined. Minor gaps refer to changes necessary primarily to clarify existing standards language. Usually

these changes are needed in order to explicitly codify IPv6 support in places where it is either implicit or omitted today, but the omission is unlikely to prevent IPv6-only operation.

3.1. MPLS Data Plane

MPLS labeled packets can be transmitted over a variety of data links [RFC3032], and MPLS labeled packets can also be encapsulated over IP. The encapsulations of MPLS in IP and GRE, as well as MPLS over L2TPv3, support IPv6. See Section 3 of RFC 4023 [RFC4023] and Section 2 of RFC 4817 [RFC4817], respectively.

Gap: None.

3.2. MPLS Control Plane

3.2.1. Label Distribution Protocol (LDP)

The Label Distribution Protocol (LDP) [RFC5036] defines a set of procedures for distribution of labels between Label Switching 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 prevent it from working in an IPv6-only network. LDP-IPv6 [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 Label Switched Path (LSP) mapping, LDP identifiers, LDP discovery, LDP session establishment, next-hop address, and LDP Time To Live (TTL) security [RFC5082] [RFC6720].

Gap: Major; update to RFC 5036 in progress via [LDP-IPv6], which should close this gap.

3.2.2. Multipoint LDP (mLDP)

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 (Section 3.2.1) with regards to control plane (discovery, transport, and session establishment) in an IPv6-only network.

2. Multipoint (MP) Forwarding Equivalence Class (FEC) Root Address: mLDP defines its own MP FECs and rules, different from LDP, to map MP LSPs. An mLDP MP FEC contains a Root Address field that is an IP address in IP networks. The current specification allows specifying the root address according to the Address Family Identifier (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 the root address to find corresponding upstream neighbor(s). This will pose a problem if an MP LSP traverses IPv4-only and IPv6-only nodes in a dual-stack network 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 a 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 a VPN environment, but a similar concept can be used/extended to solve the above issue of the IPv6-only backbone receiving an MP FEC element with an IPv4 address. The solution will require a border LSR (the one that is sitting on the border of an IPv4/IPv6 island (namely, R2 and R5 in this example)) to translate an IPv4 root address to an equivalent IPv6 address (and vice versa) through procedures similar to RFC 6512.

Gap: Major; update in progress for LDP via [LDP-IPv6], may need additional updates to RFC 6512.

3.2.3. RSVP - Traffic Engineering (RSVP-TE)

RSVP-TE [RFC3209] defines a set of procedures and enhancements to establish LSP 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.

Gap: None.

3.2.3.1. Interior Gateway Protocol (IGP)

RFC 3630 [RFC3630] specifies a method of adding traffic engineering capabilities to OSPF Version 2. New TLVs and sub-TLVs were added in RFC 5329 [RFC5329] to extend TE capabilities to IPv6 networks in OSPF Version 3.

RFC 5305 [RFC5305] specifies a method of adding traffic engineering capabilities to IS-IS. New TLVs and sub-TLVs were added in RFC 6119 [RFC6119] to extend TE capabilities to IPv6 networks.

Gap: None.

3.2.3.2. RSVP-TE Point-to-Multipoint (P2MP)

RFC 4875 [RFC4875] describes extensions to RSVP-TE for the setup of Point-to-Multipoint (P2MP) LSPs in MPLS and Generalized MPLS (GMPLS) with support for both IPv4 and IPv6.

Gap: None.

3.2.3.3. RSVP-TE Fast Reroute (FRR)

RFC 4090 [RFC4090] specifies Fast Reroute (FRR) mechanisms to establish backup LSP tunnels for local repair supporting both IPv4 and IPv6 networks. Further, [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.

Gap: None.

3.2.4. Path Computation Element (PCE)

The Path Computation Element (PCE) defined in RFC 4655 [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 RFC 5440 [RFC5440].

The PCEP specification [RFC5440] is defined for both IPv4 and IPv6 with support for PCE discovery via an IGP (OSPF [RFC5088] or IS-IS [RFC5089]) using both IPv4 and IPv6 addresses. Note that PCEP uses identical encoding of subobjects, as in RSVP-TE defined in RFC 3209 [RFC3209] that supports both IPv4 and IPv6.

The extensions to PCEP to support confidentiality [RFC5520], route exclusions [RFC5521], monitoring [RFC5886], and P2MP TE LSPs [RFC6006] have support for both IPv4 and IPv6.

Gap: None.

3.2.5. Border Gateway Protocol (BGP)

RFC 3107 [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. RFC 3107 allows BGP to distribute the label for IPv4 or IPv6 prefix in an IPv6-only network.

Gap: None.

3.2.6. Generalized Multi-Protocol Label Switching (GMPLS)

The Generalized Multi-Protocol Label Switching (GMPLS) specification includes signaling functional extensions [RFC3471] and RSVP-TE extensions [RFC3473]. The gap analysis in Section 3.2.3 applies to these.

RFC 4558 [RFC4558] specifies Node-ID Based RSVP Hello Messages with capability for both IPv4 and IPv6. RFC 4990 [RFC4990] clarifies the use of IPv6 addresses in GMPLS networks including handling in the MIB modules.

The second paragraph of Section 5.3 of RFC 6370 [RFC6370] describes the mapping from an MPLS Transport Profile (MPLS-TP) LSP_ID to RSVP-TE with an assumption that Node_IDs are derived from valid IPv4 addresses. This assumption fails in an IPv6-only network, given that there would not be any IPv4 addresses.

Gap: Minor; Section 5.3 of RFC 6370 [RFC6370] needs to be updated.

3.3. MPLS Applications

3.3.1. Layer 2 Virtual Private Network (L2VPN)

L2VPN [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 a

discovery protocol and LDP, as well as L2TPv3, as a 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 a 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 for IPv6-only operation. In its discovery procedures (Sections 3.2.2 and 6 of RFC 6074 [RFC6074]), it suggests encoding PE IP addresses in the Virtual Switching Instance ID (VSI-ID), which is encoded in Network Layer Reachability Information (NLRI) and should not exceed 12 bytes (to differentiate its AFI/SAFI (Subsequent Address Family Identifier) encoding from RFC 4761). This means that a PE IP address cannot be an IPv6 address. Also, in its signaling procedures (Section 3.2.3 of RFC 6074 [RFC6074]), it suggests encoding PE_addr in the Source Attachment Individual Identifier (SAII) and the Target Attachment Individual Identifier (TAII), which are limited to 32 bits (AII Type=1) at the moment.

RFC 6073 [RFC6073] defines the new LDP Pseudowire (PW) Switching Point PE TLV, which supports IPv4 and IPv6.

Gap: Minor; RFC 6074 needs to be updated.

3.3.1.1. Ethernet VPN (EVPN)

Ethernet VPN [EVPN] defines a method for using BGP MPLS-based Ethernet VPNs. Because it can use functions in LDP and mLDP, as well as Multicast VPLS [RFC7117], it inherits LDP gaps previously identified in Section 3.2.1. Once those gaps are resolved, it should function properly on IPv6-only networks as defined.

Gap: Major for LDP; update to RFC 5036 in progress via [LDP-IPv6] that should close this gap (see Section 3.2.1).

3.3.2. Layer 3 Virtual Private Network (L3VPN)

RFC 4364 [RFC4364] defines a method by which a Service Provider may use an IP backbone to provide IP 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 require mapping an IP packet to an IPv6-signaled LSP. RFC 4364 defines Layer 3 Virtual Private Networks (L3VPNs) for IPv4-only and has references to 32-bit BGP next-hop addresses. RFC 4659 [RFC4659] adds support for IPv6 on L3VPNs, including 128-bit BGP next-hop addresses, and discusses operation whether IPv6 is the payload or the underlying transport address family. However, RFC 4659 does not formally update RFC 4364, and thus an implementer may miss this additional set of standards unless it is explicitly identified independently of the base functionality defined in RFC 4364. Further, Section 1 of RFC 4659 explicitly identifies use case 2 as out of scope for the document.

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

Gap: Major; RFC 4659 needs to be updated to explicitly cover use case 2 (discussed in further detail below)

3.3.2.1. IPv6 Provider Edge/IPv4 Provider Edge (6PE/4PE)

RFC 4798 [RFC4798] defines IPv6 Provider Edge (6PE), which defines how to interconnect IPv6 islands over a MPLS-enabled IPv4 cloud. However, use case 2 is doing the opposite, and thus could also be referred to as IPv4 Provider Edge (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 LSPs. 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 Autonomous System (AS) to a neighboring AS.

Gap: Major; RFC 4798 covers only the "6PE" case. Use case 2 is currently not specified in an RFC.

3.3.2.2. IPv6 Virtual Private Extension/IPv4 Virtual Private Extension (6VPE/4VPE)

RFC 4659 [RFC4659] defines IPv6 Virtual Private Network Extension (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. RFC 4364 should work as defined for use case 2 above, which could also be referred to as IPv4 Virtual Private Extension (4VPE), but the RFC explicitly does not discuss this use and defines it as out of scope.

Gap: Minor; RFC 4659 needs to be updated to explicitly cover use case 2.

3.3.2.3. BGP Encapsulation Subsequent Address Family Identifier (SAFI)

RFC 5512 [RFC5512] defines the BGP Encapsulation SAFI and the BGP Tunnel Encapsulation Attribute, which can be used to signal tunneling over an IP Core that is using a single address family. 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].

Gap: None.

3.3.2.4. Multicast in MPLS/BGP IP VPN (MVPN)

RFC 6513 [RFC6513] defines the procedure to provide multicast service over an MPLS VPN backbone for downstream customers. It is sometimes referred to as Next Generation Multicast VPN (NG-MVPN) The procedure involves the below set of protocols.

3.3.2.4.1. PE-CE Multicast Routing Protocol

RFC 6513 [RFC6513] explains the use of Protocol Independent Multicast (PIM) as a Provider Edge - Customer Edge (PE-CE) protocol, while Section 11.1.2 of RFC 6514 [RFC6514] explains the use of mLDP as a PE-CE protocol.

The MCAST-VPN NLRI route-type format defined in RFC 6514 [RFC6514] is not sufficiently covering all scenarios when mLDP is used as a PE-CE protocol. The issue is explained in Section 2 of [mLDP-NLRI] along with a new route type that encodes the mLDP FEC in NLRI.

Further, [PE-CE] defines the use of BGP as a PE-CE protocol.

Gap: None.

3.3.2.4.2. P-Tunnel Instantiation

[RFC6513] explains the use of the below tunnels:

- RSVP-TE P2MP LSP
- PIM Tree
- mLDP P2MP LSP
- mLDP MP2MP LSP
- Ingress Replication

Gap: Gaps in RSVP-TE P2MP LSP (Section 3.2.3.2) and mLDP (Section 3.2.2) P2MP and MP2MP LSP are covered in previous sections. There are no MPLS-specific gaps for PIM Tree or Ingress Replication, and any protocol-specific gaps not related to MPLS are outside the scope of this document.

3.3.2.4.3. PE-PE Multicast Routing Protocol

Section 3.1 of RFC 6513 [RFC6513] explains the use of PIM as a PE-PE protocol, while RFC 6514 [RFC6514] explains the use of BGP as a PE-PE protocol.

PE-PE multicast routing is not specific to P-tunnels or to MPLS. It can be PIM or BGP with P-tunnels that are label based or PIM tree based. Enabling PIM as a PE-PE multicast protocol is equivalent to running it on a non-MPLS IPv6 network, so there are not any MPLS-specific considerations and any gaps are applicable for non-MPLS networks as well. Similarly, BGP only includes the P-Multicast Service Interface (PMSI) tunnel attribute as a part of the NLRI, which is inherited from P-tunnel instantiation and considered to be an opaque value. Any gaps in the control plane (PIM or BGP) will not be specific to MPLS.

Gap: Any gaps in PIM or BGP as a PE-PE multicast routing protocol are not unique to MPLS, and therefore are outside the scope of this document. It is included for completeness.

3.3.3. MPLS Transport Profile (MPLS-TP)

MPLS-TP does not require IP (see Section 2 of RFC 5921 [RFC5921]) and should not be affected by operation on an IPv6-only network. Therefore, this is considered out of scope for this document but is included for completeness.

Although not required, MPLS-TP can use IP. One such example is included in Section 3.2.6, where MPLS-TP identifiers can be derived from valid IPv4 addresses.

Gap: None. MPLS-TP does not require IP.

3.4. MPLS Operations, Administration, and Maintenance (MPLS OAM)

For MPLS LSPs, there are primarily three OAM mechanisms: Extended ICMP [RFC4884] [RFC4950], LSP Ping [RFC4379], and Bidirectional Forwarding Detection (BFD) for MPLS LSPs [RFC5884]. For MPLS Pseudowires, there is also Virtual Circuit Connectivity Verification (VCCV) [RFC5085] [RFC5885]. Most of these mechanisms work in pure

IPv6 environments, but there are some problems encountered in mixed environments due to address-family mismatches. The next subsections cover these gaps in detail.

Gap: Major; RFC 4379 needs to be updated to better support multipath IPv6. Additionally, there is potential for dropped messages in Extended ICMP and LSP Ping due to IP version mismatches. It is important to note that this is a more generic problem with tunneling when address-family mismatches exist and is not specific to MPLS. While MPLS will be affected, it will be difficult to fix this problem specifically for MPLS, rather than fixing the more generic problem.

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 an 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 IPv4 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 IPv4 payload, format an ICMPv4 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 IPv4-aware. It cannot parse the original IPv4 datagram, nor can it send an IPv4 message. So, no ICMP message is delivered to the source. Some specific ICMP extensions, in particular, ICMP extensions for interface and next-hop identification [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.

Gap: Major; IP version mismatches may cause dropped messages. However, as noted in the previous section, this problem is not specific to MPLS.

3.4.2. Label Switched Path Ping (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, which are later specified for IPv6 in RFC 6829 [RFC6829]. The multipath information also includes IPv6 encodings (see Section 3.3.1 of RFC 4379).

LSP Ping packets are UDP packets over either IPv4 or IPv6 (see Section 4.3 of RFC 4379). However, for IPv6 the destination IP address is a (randomly chosen) IPv6 address from the range 0:0:0:0:0:FFFF:127/104; that is, using an IPv4-mapped IPv6 address. This is a transitional mechanism that should not bleed into IPv6-only networks, as [IPv4-MAPPED] explains. The issue is that the MPLS LSP Ping mechanism needs a range of loopback IP addresses to be used as destination addresses to exercise Equal Cost Multiple Path (ECMP), but the IPv6 address architecture specifies a single address (::1/128) for loopback. A mechanism to achieve this was proposed in [LOOPBACK-PREFIX].

Additionally, RFC 4379 does not define the value to be used in the IPv6 Router Alert option (RAO). For IPv4 RAO, a value of zero is used. However, there is no equivalent value for IPv6 RAO. This gap needs to be fixed to be able to use LSP Ping in IPv6 networks. Further details on this gap are captured, along with a proposed solution, in [IPv6-RAO].

Another gap is that 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.

The mechanism described in RFC 4379 also does not sufficiently explain the behavior in certain IPv6-specific scenarios. For example, RFC 4379 defines the K value as 28 octets when the Address Family is set to IPv6 Unnumbered, but it doesn't describe how to carry a 32-bit LSR Router ID in the 128-bit Downstream IP Address field.

Gap: Major; LSP Ping uses IPv4-mapped IPv6 addresses. IP version mismatches may cause dropped messages and unclear mapping from the LSR Router ID to Downstream IP Address.

3.4.3. Bidirectional Forwarding Detection (BFD)

The BFD specification for MPLS LSPs [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).

Gap: None.

3.4.4. Pseudowire OAM

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

Additionally, for LSP Ping for pseudowires, the Pseudowire FECs are specified for IPv6 in RFC 6829 [RFC6829].

Gap: None.

3.4.5. MPLS Transport Profile (MPLS-TP) OAM

As with MPLS-TP, MPLS-TP OAM [RFC6371] does not require IP or existing MPLS OAM functions and should not be affected by operation on an IPv6-only network. Therefore, this is out of scope for this document but is included for completeness. Although not required, MPLS-TP can use IP.

Gap: None. MPLS-TP OAM does not require IP.

3.5. MIB Modules

RFC 3811 [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], the GMPLS-TE MIB specification [RFC4802] and the FRR extension [RFC6445]. "Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management" [MPLS-TC] tries to resolve this gap by marking this textual convention as obsolete.

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

Lastly, RFC 4990 [RFC4990] discusses how to handle IPv6 sources and destinations in the MPLS and GMPLS-TE MIB modules. In particular, Section 8 of RFC 4990 [RFC4990] describes a method of defining or monitoring an LSP tunnel using the MPLS-TE and GMPLS-TE MIB modules, working around some of the limitations in RFC 3811 [RFC3811].

Gap: Minor; Section 8 of RFC 4990 [RFC4990] describes a method to handle IPv6 addresses in the MPLS-TE [RFC3812] and GMPLS-TE [RFC4802] MIB modules. Work underway to update RFC 3811 via [MPLS-TC], may also need to update RFC 3812, RFC 4802, and RFC 6445, which depend on it.

4. Gap Summary

This document has reviewed a wide variety of MPLS features and protocols to determine their suitability for use on IPv6-only or IPv6-primary networks. While some parts of the MPLS suite will function properly without additional changes, gaps have been identified in others that 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. Note that because the referenced documents are works in progress and do not have consensus at the time of this document’s publication, there could be other solutions proposed at a future time, and the pointers in this document should not be considered normative in any way. Additionally, work in progress on new features that use MPLS protocols will need to ensure that those protocols support operation on IPv6-only or IPv6-primary networks, or explicitly identify any dependencies on existing gaps that, once resolved, will allow proper IPv6-only operation.

Identified Gaps in MPLS for IPv6-Only Networks

Item	Gap	Addressed in
LDP S.3.2.1	LSP mapping, LDP identifiers, LDP discovery, LDP session establishment, next-hop address, and LDP TTL security	[LDP-IPv6]
mLDP S.3.2.2	Inherits gaps from LDP, RFC 6512 [RFC6512]	Inherits [LDP-IPv6], additional fixes TBD
GMPLS S.3.2.6	RFC 6370 [RFC6370] Node ID derivation	TBD
L2VPN S.3.3.1	RFC 6074 [RFC6074] discovery, signaling	TBD
L3VPN S.3.3.2	RFC 4659 [RFC4659] does not define a method for 4PE/4VPE	TBD
OAM S.3.4	RFC 4379 [RFC4379] No IPv6 multipath support, no IPv6 RAO, possible dropped messages in IP version mismatch	[IPv6-RAO]
MIB Modules S.3.5	RFC 3811 [RFC3811] no IPv6 textual convention	[MPLS-TC]

Table 1: IPv6-Only MPLS Gaps

5. 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, follow-on work recommended by this gap analysis will need to address any effects that the use of IPv6 in their modifications may have on security.

6. References

6.1. Normative References

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001, <http://www.rfc-editor.org/info/rfc3032>.

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, May 2001, <http://www.rfc-editor.org/info/rfc3107>.

[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, <http://www.rfc-editor.org/info/rfc3209>.

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003, <http://www.rfc-editor.org/info/rfc3471>.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003, <http://www.rfc-editor.org/info/rfc3473>.

[RFC3811] Nadeau, T. and J. Cucchiara, "Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management", RFC 3811, June 2004, <http://www.rfc-editor.org/info/rfc3811>.

[RFC4023] Worster, T., Rekhter, Y., and E. Rosen, "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005, <http://www.rfc-editor.org/info/rfc4023>.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006, <http://www.rfc-editor.org/info/rfc4379>.

[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, <http://www.rfc-editor.org/info/4659>.

[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, <http://www.rfc-editor.org/info/rfc4817>.

[RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007, <http://www.rfc-editor.org/info/rfc5036>.

[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, <http://www.rfc-editor.org/info/rfc6074>.

[RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, September 2011, <http://www.rfc-editor.org/info/rfc6370>.

[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, <http://www.rfc-editor.org/info/rfc6512>.

6.2. Informative References

[EVPN] Sajassi, A., Aggarwal, R., Bitar, N., Isaac, A., and J. Uttaro, "BGP MPLS Based Ethernet VPN", Work in Progress, draft-ietf-l2vpn-evpn-11, October 2014.

[IPv4-MAPPED] Metz, C. and J. Hagino, "IPv4-Mapped Addresses on the Wire Considered Harmful", Work in Progress, draft-itojun-v6ops-v4mapped-harmful-02, October 2003.

[IPv6-RAO] Raza, K., Akiya, N., and C. Pignataro, "IPv6 Router Alert Option for MPLS OAM", Work in Progress, draft-raza-mpls-oam-ipv6-rao-02, September 2014.

[LDP-IPv6] Asati, R., Manral, V., Papneja, R., and C. Pignataro, "Updates to LDP for IPv6", Work in Progress, draft-ietf-mpls-ldp-ipv6-14, October 2014.

[LOOPBACK-PREFIX]
Smith, M., "A Larger Loopback Prefix for IPv6", Work in Progress, draft-smith-v6ops-larger-ipv6-loopback-prefix-04, February 2013.

[mLDP-NLRI]
Wijnands, I., Rosen, E., and U. Joorde, "Encoding mLDP FECs in the NLRI of BGP MCAST-VPN Routes", Work in Progress, draft-ietf-l3vpn-mvpn-mldp-nlri-10, November 2014.

[MPLS-TC] Manral, V., Tsou, T., Will, W., and F. Fondelli, "Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management", Work in Progress, draft-manral-mpls-rfc3811bis-04, September 2014.

[PE-CE] Patel, K., Rekhter, Y., and E. Rosen, "BGP as an MVPN PE-CE Protocol", Work in Progress, draft-ietf-l3vpn-mvpn-pe- ce-02, October 2014.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003, <http://www.rfc-editor.org/info/rfc3630>.

[RFC3812] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)", RFC 3812, June 2004, <http://www.rfc-editor.org/info/rfc3812>.

[RFC3813] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)", RFC 3813, June 2004, <http://www.rfc-editor.org/info/rfc3813>.

[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, <http://www.rfc-editor.org/info/rfc3815>.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005, <http://www.rfc-editor.org/info/rfc4090>.

[RFC4220] Dubuc, M., Nadeau, T., and J. Lang, "Traffic Engineering Link Management Information Base", RFC 4220, November 2005, <http://www.rfc-editor.org/info/rfc4220>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006, <http://www.rfc-editor.org/info/rfc4364>.

[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, <http://www.rfc-editor.org/info/rfc4447>.

[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, <http://www.rfc-editor.org/info/rfc4558>.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006, <http://www.rfc-editor.org/info/rfc4655>.

[RFC4664] Andersson, L. and E. Rosen, "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006, <http://www.rfc-editor.org/info/rfc4664>.

[RFC4761] Kompella, K. and Y. Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007, <http://www.rfc-editor.org/info/rfc4761>.

[RFC4762] Lasserre, M. and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007, <http://www.rfc-editor.org/info/rfc4762>.

[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, <http://www.rfc-editor.org/info/rfc4798>.

[RFC4802] Nadeau, T. and A. Farrel, "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", RFC 4802, February 2007, <http://www.rfc-editor.org/info/rfc4802>.

[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, <http://www.rfc-editor.org/info/rfc4875>.

[RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, April 2007, <http://www.rfc-editor.org/info/rfc4884>.

[RFC4950] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "ICMP Extensions for Multiprotocol Label Switching", RFC 4950, August 2007, <http://www.rfc-editor.org/info/rfc4950>.

[RFC4990] Shiomoto, K., Papneja, R., and R. Rabbat, "Use of Addresses in Generalized Multiprotocol Label Switching (GMPLS) Networks", RFC 4990, September 2007, <http://www.rfc-editor.org/info/rfc4990>.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, October 2007, <http://www.rfc-editor.org/info/rfc5082>.

[RFC5085] Nadeau, T. and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007, <http://www.rfc-editor.org/info/rfc5085>.

[RFC5088] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008, <http://www.rfc-editor.org/info/rfc5088>.

[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, <http://www.rfc-editor.org/info/rfc5089>.

[RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, September 2008, <http://www.rfc-editor.org/info/rfc5286>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008, <http://www.rfc-editor.org/info/rfc5305>.

[RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, September 2008, <http://www.rfc-editor.org/info/rfc5329>.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009, <http://www.rfc-editor.org/info/rfc5440>.

[RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", RFC 5512, April 2009, <http://www.rfc-editor.org/info/rfc5512>.

[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, <http://www.rfc-editor.org/info/rfc5520>.

[RFC5521] Oki, E., Takeda, T., and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", RFC 5521, April 2009, <http://www.rfc-editor.org/info/rfc5521>.

[RFC5640] Filsfils, C., Mohapatra, P., and C. Pignataro, "Load-Balancing for Mesh Softwires", RFC 5640, August 2009, <http://www.rfc-editor.org/info/rfc5640>.

[RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010, <http://www.rfc-editor.org/info/rfc5837>.

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010, <http://www.rfc-editor.org/info/rfc5884>.

[RFC5885] Nadeau, T. and C. Pignataro, "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010, <http://www.rfc-editor.org/info/rfc5885>.

[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, <http://www.rfc-editor.org/info/rfc5886>.

[RFC5921] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010, <http://www.rfc-editor.org/info/rfc5921>.

[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, <http://www.rfc-editor.org/info/rfc6006>.

[RFC6073] Martini, L., Metz, C., Nadeau, T., Bocci, M., and M. Aissaoui, "Segmented Pseudowire", RFC 6073, January 2011, <http://www.rfc-editor.org/info/rfc6073>.

[RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", RFC 6119, February 2011, <http://www.rfc-editor.org/info/rfc6119>.

[RFC6371] Busi, I. and D. Allan, "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", RFC 6371, September 2011, <http://www.rfc-editor.org/info/rfc6371>.

[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, <http://www.rfc-editor.org/info/rfc6388>.

[RFC6445] Nadeau, T., Koushik, A., and R. Cetin, "Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base for Fast Reroute", RFC 6445, November 2011, <http://www.rfc-editor.org/info/rfc6445>.

[RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", RFC 6513, February 2012, <http://www.rfc-editor.org/info/rfc6513>.

[RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, February 2012, <http://rfc-editor.org/info/rfc6514>.

[RFC6540] George, W., Donley, C., Liljenstolpe, C., and L. Howard, "IPv6 Support Required for All IP-Capable Nodes", BCP 177, RFC 6540, April 2012, <http://www.rfc-editor.org/info/rfc6540>.

[RFC6624] Kompella, K., Kothari, B., and R. Cherukuri, "Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling", RFC 6624, May 2012, <http://www.rfc-editor.org/info/rfc6624>.

[RFC6720] Pignataro, C. and R. Asati, "The Generalized TTL Security Mechanism (GTSM) for the Label Distribution Protocol (LDP)", RFC 6720, August 2012, <http://www.rfc-editor.org/info/rfc6720>.

[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, <http://www.rfc-editor.org/info/rfc6829>.

[RFC7117] Aggarwal, R., Kamite, Y., Fang, L., Rekhter, Y., and C. Kodeboniya, "Multicast in Virtual Private LAN Service (VPLS)", RFC 7117, February 2014, <http://www.rfc-editor.org/info/rfc7117>.

Acknowledgements

The authors wish to thank Alvaro Retana, Andrew Yourtchenko, Loa Andersson, David Allan, Mach Chen, Mustapha Aissaoui, and Mark Tinka for their detailed reviews, as well as Brian Haberman, Joel Jaeggli, Adrian Farrel, Nobo Akiya, Francis Dupont, and Tobias Gondrom for their comments.

Contributors

The following people have contributed text to this document:

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

EMail: rajiva@cisco.com

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

EMail: skraza@cisco.com

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

EMail: rbonica@juniper.net

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

EMail: rajiv.papneja@huawei.com

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
India

EMail: dhruv.ietf@gmail.com

Vishwas Manral
Ionos Networks
Sunnyvale, CA 94089
United States

EMail: vishwas@ionosnetworks.com

Nagendra Kumar
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709
United States

EMail: naikumar@cisco.com

Authors' Addresses

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

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

Carlos Pignataro (editor)
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
United States

Phone: +1-919-392-7428
EMail: cpignata@cisco.com