

Internet Engineering Task Force (IETF)
Request for Comments: 6575
Category: Standards Track
ISSN: 2070-1721

H. Shah, Ed.
Ciena
E. Rosen, Ed.
G. Heron, Ed.
Cisco
V. Kompella, Ed.
Alcatel-Lucent
June 2012

Address Resolution Protocol (ARP) Mediation for
IP Interworking of Layer 2 VPNs

Abstract

The Virtual Private Wire Service (VPWS), detailed in RFC 4664, provides point-to-point connections between pairs of Customer Edge (CE) devices. It does so by binding two Attachment Circuits (each connecting a CE device with a Provider Edge (PE) device) to a pseudowire (connecting the two PEs). In general, the Attachment Circuits must be of the same technology (e.g., both Ethernet or both ATM), and the pseudowire must carry the frames of that technology. However, if it is known that the frames' payload consists solely of IP datagrams, it is possible to provide a point-to-point connection in which the pseudowire connects Attachment Circuits of different technologies. This requires the PEs to perform a function known as "Address Resolution Protocol (ARP) Mediation". ARP Mediation refers to the process of resolving Layer 2 addresses when different resolution protocols are used on either Attachment Circuit. The methods described in this document are applicable even when the CEs run a routing protocol between them, as long as the routing protocol runs over IP.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc6575>.

Copyright Notice

Copyright (c) 2012 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	3
1.1. Conventions Used in This Document	4
2. ARP Mediation (AM) Function	5
3. IP Layer 2 Interworking Circuit	6
4. IP Address Discovery Mechanisms	6
4.1. Discovery of IP Addresses of Locally Attached IPv4 CE	7
4.1.1. Monitoring Local Traffic	7
4.1.2. CE Devices Using ARP	7
4.1.3. CE Devices Using Inverse ARP	8
4.1.4. CE Devices Using PPP	9
4.1.5. Router Discovery Method	10
4.1.6. Manual Configuration	10
4.2. How a CE Learns the IPv4 Address of a Remote CE	10
4.2.1. CE Devices Using ARP	11
4.2.2. CE Devices Using Inverse ARP	11
4.2.3. CE Devices Using PPP	11
4.3. Discovery of IP Addresses of IPv6 CE Devices	11
4.3.1. Distinguishing Factors between IPv4 and IPv6	11
4.3.2. Requirements for PEs	12
4.3.3. Processing of Neighbor Solicitations	12
4.3.4. Processing of Neighbor Advertisements	13
4.3.5. Processing Inverse Neighbor Solicitations (INs) ...	14
4.3.6. Processing of Inverse Neighbor Advertisements (INAs)	15
4.3.7. Processing of Router Solicitations	15
4.3.8. Processing of Router Advertisements	15
4.3.9. Duplicate Address Detection	16
4.3.10. CE Address Discovery for CEs Attached Using PPP ...	16
5. CE IPv4 Address Signaling between PEs	16
5.1. When to Signal an IPv4 Address of a CE	16
5.2. LDP-Based Distribution of CE IPv4 Addresses	17

6.	IPv6 Capability Advertisement	20
6.1.	PW Operational Down on Stack Capability Mismatch	21
6.2.	Stack Capability Fallback	21
7.	IANA Considerations	22
7.1.	LDP Status Messages	22
7.2.	Interface Parameters	22
8.	Security Considerations	22
8.1.	Control Plane Security	23
8.2.	Data Plane Security	24
9.	Acknowledgements	24
10.	Contributors	24
11.	References	25
11.1.	Normative References	25
11.2.	Informative References	26
Appendix A.	Use of IGPs with IP L2 Interworking L2VPNs	27
A.1.	OSPF	27
A.2.	RIP	27
A.3.	IS-IS	28

1. Introduction

Layer 2 Virtual Private Networks (L2VPNs) are constructed over a Service Provider IP/MPLS backbone but are presented to the Customer Edge (CE) devices as Layer 2 networks. In theory, L2VPNs can carry any Layer 3 protocol, but in many cases, the Layer 3 protocol is IP. Thus, it makes sense to consider procedures that are optimized for IP.

In a typical implementation, illustrated in the diagram below, the CE devices are connected to the Provider Edge (PE) devices via Attachment Circuits (ACs). The ACs are Layer 2 circuits. In a pure L2VPN, if traffic sent from CE1 via AC1 reaches CE2 via AC2, both ACs would have to be of the same type (i.e., both Ethernet, both Frame Relay, etc.). However, if it is known that only IP traffic will be carried, the ACs can be of different technologies, provided that the PEs provide the appropriate procedures to allow the proper transfer of IP packets.

2. ARP Mediation (AM) Function

The ARP Mediation (AM) function is an element of a PE node that deals with the IP address resolution for CE devices connected via a VPWS L2VPN. By placing this function in the PE node, ARP Mediation is transparent to the CE devices.

For a given point-to-point connection between a pair of CEs, the ARP Mediation procedure depends on whether the packets being forwarded are IPv4 or IPv6. A PE that is to perform ARP Mediation for IPv4 packets MUST perform the following logical steps:

1. Discover the IP address of the locally attached CE device.
2. Terminate. Do not forward ARP and Inverse ARP requests from the CE device at the local PE.
3. Distribute the IP address to the remote PE using pseudowire control signaling.
4. Notify the locally attached CE of the IP address of the remote CE.
5. Respond appropriately to ARP and Inverse ARP requests from the local CE device using the IP address of the remote CE and the hardware address of the local PE.

A PE that is to perform ARP Mediation for IPv6 packets MUST perform the following logical steps:

1. Discover the IPv6 addresses of the locally attached CE device, together with those of the remote CE device.
2. Perform the following steps:
 - a. Intercept Neighbor Discovery (ND) and Inverse Neighbor Discovery (IND) packets received from the local CE device.
 - b. From these ND and IND packets, learn the IPv6 configuration of the CE.
 - c. Forward the ND and IND packets over the pseudowire to the remote PE.

3. Intercept Neighbor Discovery and Inverse Neighbor Discovery packets received over the pseudowire from the remote PE, possibly modifying them (if required for the type of outgoing AC) before forwarding to the local CE and learning information about the IPv6 configuration of the remote CE.

Details for the procedures described above are given in the following sections.

3. IP Layer 2 Interworking Circuit

The IP Layer 2 Interworking Circuit refers to interconnection of the Attachment Circuit with the IP Layer 2 Transport pseudowire that carries IP datagrams as the payload. The ingress PE removes the data link header of its local Attachment Circuit and transmits the payload (an IP packet) over the pseudowire with or without the optional control word. If the IP packet arrives at the ingress PE with multiple data link headers (for example, in the case of bridged Ethernet PDU on an ATM Attachment Circuit), all data link headers MUST be removed from the IP packet before transmission over the pseudowire (PW). The egress PE encapsulates the IP packet with the data link header used on its local Attachment Circuit.

The encapsulation for the IP Layer 2 Transport pseudowire is described in [RFC4447]. The "IP Layer 2 Interworking Circuit" pseudowire is also referred to as "IP pseudowire" in this document.

In the case of an IPv6 L2 Interworking Circuit, the egress PE MAY modify the contents of Neighbor Discovery or Inverse Neighbor Discovery packets before encapsulating the IP packet with the data link header.

4. IP Address Discovery Mechanisms

An IP Layer 2 Interworking Circuit enters monitoring state immediately after configuration. During this state, it performs two functions:

- o Discovery of the CE IP device(s)
- o Establishment of the PW

The establishment of the PW occurs independently from local CE IP address discovery. During the period when the PW has been established but the local CE IP device has not been discovered, only broadcast/multicast IP frames are propagated between the Attachment Circuit and pseudowire; unicast IP datagrams are dropped. The IP destination address is used to classify unicast/multicast packets.

Unicast IP frames are propagated between the AC and pseudowire only when CE IP devices on both Attachment Circuits have been discovered and notified and proxy functions have completed.

The need to wait for address resolution completion before unicast IP traffic can flow is simple.

- o PEs do not perform routing operations.
- o The destination IP address in the packet is not necessarily that of the attached CE.
- o On a broadcast link, there is no way to find out the Media Access Control (MAC) address of the CE based on the destination IP address of the packet.

4.1. Discovery of IP Addresses of Locally Attached IPv4 CE

A PE MUST support manual configuration of IPv4 CE addresses. This section also describes automated mechanisms by which a PE MAY also discover an IPv4 CE address.

4.1.1. Monitoring Local Traffic

The PE devices MAY learn the IP addresses of the locally attached CEs from any IP traffic, such as link-local multicast packets (e.g., destined to 224.0.0.x), and are not restricted to the operations below.

4.1.2. CE Devices Using ARP

If a CE device uses ARP to determine the IP-address-to-MAC-address binding of its neighbor, the PE processes the ARP requests to learn the IP address of the local CE for the local Attachment Circuit.

The method described in this document only supports the case where there is a single CE per Attachment Circuit. However, customer-facing access topologies may exist whereby more than one CE appears to be connected to the PE on a single Attachment Circuit. For example, this could be the case when CEs are connected to a shared LAN that connects to the PE. In such a case, the PE MUST select one local CE. The selection could be based on manual configuration or the PE MAY optionally use the following selection criteria. In either case, manual configuration of the IP address of the local CE (and its MAC address) MUST be supported.

- o Wait to learn the IP address of the remote CE (through PW signaling) and then select the local CE that is sending the request for IP address of the remote CE.
- o Augment cross-checking with the local IP address learned through listening for link-local multicast packets (as per Section 4.1.1).
- o Augment cross-checking with the local IP address learned through the Router Discovery Protocol (as described in Section 4.1.5).
- o There is still a possibility that the local PE may not receive an IP address advertisement from the remote PE, and there may exist multiple local IP routers that attempt to 'connect' to remote CEs. In this situation, the local PE MAY use some other criteria to select one IP device from many (such as "the first ARP received"), or an operator MAY configure the IP address of the local CE. Note that the operator does not have to configure the IP address of the remote CE (as that would be learned through pseudowire signaling).

Once the local and remote CEs have been discovered for the given Attachment Circuit, the local PE responds with its own MAC address to any subsequent ARP requests from the local CE with a destination IP address matching the IP address of the remote CE.

The local PE signals the IP address of the local CE to the remote PE and MAY initiate an unsolicited ARP response to notify the IP-address-to-MAC-address binding for the remote CE to the local CE (again using its own MAC address).

Once the ARP Mediation function is completed (i.e., the PE device knows both the local and remote CE IP addresses), unicast IP frames are propagated between the AC and the established PW.

The PE MAY periodically generate ARP request messages for the IP address of the CE as a means of verifying the continued existence of the IP address and its MAC address binding. The absence of a response from the CE device for a given number of retries could be used as a trigger for withdrawal of the IP address advertisement to the remote PE. The local PE would then re-enter the address resolution phase to rediscover the IP address of the attached CE. Note that this "heartbeat" scheme is needed only where the failure of a CE device may otherwise be undetectable.

4.1.3. CE Devices Using Inverse ARP

If a CE device uses Inverse ARP to determine the IP address of its neighbor, the attached PE processes the Inverse ARP request from the Attachment Circuit and responds with an Inverse ARP reply containing

the IP address of the remote CE, if the address is known. If the PE does not yet have the IP address of the remote CE, it does not respond, but records the IP address of the local CE and the circuit information. Subsequently, when the IP address of the remote CE becomes available, the PE MAY initiate an Inverse ARP request as a means of notifying the local CE of the IP address of the remote CE.

This is the typical mode of operation for Frame Relay and ATM Attachment Circuits. If the CE does not use Inverse ARP, the PE can still discover the IP address of the local CE using the mechanisms described in Sections 4.1.1 and 4.1.5.

4.1.4. CE Devices Using PPP

The IP Control Protocol [RFC1332] describes a procedure to establish and configure IP on a point-to-point connection, including the negotiation of IP addresses. When such an Attachment Circuit is configured for IP interworking, PPP negotiation is not performed end-to-end between CE devices. Instead, PPP negotiation takes place between the CE and its local PE. The PE performs proxy PPP negotiation and informs the attached CE of the IP address of the remote CE during IP Control Protocol (IPCP) negotiation using the IP-Address option (0x03).

When a PPP link completes Link Control Protocol (LCP) negotiations, the local PE MAY perform the following IPCP actions:

- o The PE learns the IP address of the local CE from the Configure-Request received with the IP-Address option (0x03). If the IP address is non-zero, the PE records the address and responds with Configure-Ack. However, if the IP address is zero, the PE responds with Configure-Reject (as this is a request from the CE to assign it an IP address). Also, the IP-Address option is set with a zero value in the Configure-Reject response to instruct the CE not to include that option in any subsequent Configure-Request.
- o If the PE receives a Configure-Request without the IP-Address option, it responds with a Configure-Ack. In this case, the PE is unable to learn the IP address of the local CE using IPCP; hence, it MUST rely on other means as described in Sections 4.1.1 and 4.1.5. Note that in order to employ other learning mechanisms, the IPCP negotiations MUST have reached the open state.
- o If the PE does not know the IP address of the remote CE, it sends a Configure-Request without the IP-Address option.

- o If the PE knows the IP address of the remote CE, it sends a Configure-Request with the IP-Address option containing the IP address of the remote CE.

The IPCP IP-Address option MAY be negotiated between the PE and the local CE device. Configuration of other IPCP options MAY be rejected. Other Network Control Protocols (NCPs), with the exception of the Compression Control Protocol (CCP) and the Encryption Control Protocol (ECP), MUST be rejected. The PE device MAY reject configuration of the CCP and ECP.

4.1.5. Router Discovery Method

In order to learn the IP address of the CE device for a given Attachment Circuit, the PE device MAY execute the Router Discovery Protocol [RFC1256] whereby a Router Discovery Request (ICMP - Router Solicitation) message is sent using a source IP address of zero. The IP address of the CE device is extracted from the Router Discovery Response (ICMP - Router Advertisement) message from the CE. It is possible that the response contains more than one router address with the same preference level, in which case, some heuristics (such as first on the list) are necessary. The use of the Router Discovery method by the PE is optional.

4.1.6. Manual Configuration

In some cases, it may not be possible to discover the IP address of the local CE device using the mechanisms described in Sections 4.1.1 to 4.1.5. In such cases, manual configuration MAY be used. All implementations of this document MUST support manual configuration of the IPv4 address of the local CE. This is the only REQUIRED mode for a PE to support.

The support for configuration of the IP address of the remote CE is OPTIONAL.

4.2. How a CE Learns the IPv4 Address of a Remote CE

Once the local PE has received the IP address information of the remote CE from the remote PE, it will either initiate an address resolution request or respond to an outstanding request from the attached CE device.

In the event that the IPv4 address of the remote CE is manually configured, the address resolution can begin immediately as receipt of remote IP address of the CE becomes unnecessary.

4.2.1. CE Devices Using ARP

When the PE learns the IP address of the remote CE as described in Section 5.1, it may or may not already know the IP address of the local CE. If the IP address is not known, the PE **MUST** wait until it is acquired through one of the methods described in Sections 4.1.1, 4.1.2, and 4.1.5. If the IP address of the local CE is known, the PE **MAY** choose to generate an unsolicited ARP message to notify the local CE about the binding of the IP address of the remote CE with the PE's own MAC address.

When the local CE generates an ARP request, the PE **MUST** proxy the ARP response [RFC925] using its own MAC address as the source hardware address and the IP address of the remote CE as the source protocol address. The PE **MUST** respond only to those ARP requests whose destination protocol address matches the IP address of the remote CE.

4.2.2. CE Devices Using Inverse ARP

When the PE learns the IP address of the remote CE, it **SHOULD** generate an Inverse ARP request. If the Attachment Circuit requires activation (e.g., Frame Relay), the PE **SHOULD** activate it first before the Inverse ARP request. It should be noted that the PE might never receive the response to its own request, nor see any Inverse ARP request from the CE, in cases where the CE is pre-configured with the IP address of the remote CE or where the use of Inverse ARP has not been enabled. In either case, the CE has used other means to learn the IP address of its neighbor.

4.2.3. CE Devices Using PPP

When the PE learns the IP address of the remote CE, it **SHOULD** initiate a Configure-Request and set the IP-Address option to the IP address of the remote CE. This notifies the local CE of the IP address of the remote CE.

4.3. Discovery of IP Addresses of IPv6 CE Devices

4.3.1. Distinguishing Factors between IPv4 and IPv6

IPv4 uses ARP and Inverse ARP to resolve IP address and link-layer associations. Since these are dedicated address resolution protocols, and not IP packets, they cannot be carried on an IP pseudowire. They **MUST** be processed locally and the IPv4 address information they carry signaled between the PEs using the pseudowire control plane. IPv6 uses ICMPv6 extensions to resolve IP address and

link address associations. As these are IPv6 packets, they can be carried on an IP pseudowire; therefore, no IPv6 address signaling is required.

4.3.2. Requirements for PEs

A PE device that supports IPv6 MUST be capable of the following:

- o Intercepting ICMPv6 Neighbor Discovery [RFC4861] and Inverse Neighbor Discovery [RFC3122] packets received over the AC as well as over the PW,
- o Recording the IPv6 interface addresses and CE link-layer addresses present in these packets,
- o Possibly modifying these packets as dictated by the data link type of the egress AC (described in the following sections), and
- o Forwarding them towards the original destination.

The PE MUST also be capable of generating packets in order to interwork between Neighbor Discovery (ND) and Inverse Neighbor Discovery (IND). This is specified in Sections 4.3.3 to 4.3.6.

If an IP PW is used to interconnect CEs that use IPv6 Router Discovery [RFC4861], a PE device MUST also be capable of intercepting and processing those Router Discovery packets. This is required in order to translate between different link-layer addresses. If a Router Discovery message contains a link-layer address, then the PE MAY also use this message to discover the link-layer address and IPv6 interface address. This is described in more detail in Sections 4.3.7 and 4.3.8.

The PE device MUST learn a list of CE IPv6 interface addresses for its directly attached CE and another list of CE IPv6 interface addresses for the far-end CE. The PE device MUST also learn the link-layer address of the local CE and be able to use it when forwarding traffic between the local and far-end CEs. The PE MAY also wish to monitor the source link-layer address of data packets received from the CE and discard packets not matching its learned CE link-layer address.

4.3.3. Processing of Neighbor Solicitations

A Neighbor Solicitation received on an AC from a local CE SHOULD be inspected to determine and learn an IPv6 interface address (if provided, this will not be the case for Duplicate Address Detection) and any link-layer address provided. The packet MUST then be

forwarded over the pseudowire unmodified. A Neighbor Solicitation received over the pseudowire SHOULD be inspected to determine and learn an IPv6 interface address for the far-end CE. If a source link-layer address option is present, the PE MUST remove it. The PE MAY substitute an appropriate link-layer address option, specifying the link-layer address of the PE interface attached to the local AC. Note that if the local AC is Ethernet, failure to substitute a link-layer address option may mean that the CE has no valid link-layer address with which to transmit data packets.

When a PE with a local AC, which is of the type point-to-point Layer 2 circuit, e.g., FR, ATM or PPP, receives a Neighbor Solicitation from a far-end PE over the pseudowire, after learning the IP address of the far-end CE, the PE MAY use one of the following procedures:

1. Forward the Neighbor Solicitation to the local CE after replacing the source link-layer address with the link-layer address of the local AC.
2. Send an Inverse Neighbor Solicitation to the local CE, specifying the far-end CE's IP address and the link-layer address of the PE interface attached to local AC.
3. Reply to the far-end PE with a Neighbor Advertisement, using the IP address of the local CE as the source address and an appropriate link-layer address option that specifies the link-layer address of the PE interface attached to local AC. As described in Section 4.3.10, the IP address of the local CE is learned through IPv6 Control Protocol (IPv6CP) in the case of PPP and through Neighbor Solicitation in other cases.

4.3.4. Processing of Neighbor Advertisements

A Neighbor Advertisement received on an AC from a local CE SHOULD be inspected to determine and learn an IPv6 interface address and any link-layer address provided. The packet MUST then be forwarded over the IP pseudowire unmodified.

A Neighbor Advertisement received over the pseudowire SHOULD be inspected to determine and learn an IPv6 interface address for the far-end CE. If a source link-layer address option is present, the PE MUST remove it. The PE MAY substitute an appropriate link-layer address option, specifying the link-layer address of the PE interface attached to local AC. Note that if the local AC is Ethernet, failure to substitute a link-layer address option may mean that the local AC has no valid link-layer address with which to transmit data packets.

When a PE with a local AC that is of the type point-to-point Layer 2 circuit, such as ATM, FR, or PPP, receives a Neighbor Advertisement over the pseudowire, in addition to learning the remote CE's IPv6 address, it SHOULD perform the following steps:

- o If the AC supports Inverse Neighbor Discovery (IND) and the PE had already processed an Inverse Neighbor Solicitation (INS) from the local CE, it SHOULD send an Inverse Neighbor Advertisement (INA) on the local AC using source IP address information received in an ND advertisement (ND-ADV) and its own local AC link-layer information.
- o If the PE has not received any Inverse Neighbor Solicitation (INS) from the local CE and the AC supports Inverse Neighbor Discovery (IND), it SHOULD send an INS on the local AC using source IP address information received in the INA together with its own local AC link-layer information.

4.3.5. Processing Inverse Neighbor Solicitations (INSS)

An INS received on an AC from a local CE SHOULD be inspected to determine and learn the IPv6 addresses and the link-layer addresses. The packet MUST then be forwarded over the pseudowire unmodified.

An INS received over the pseudowire SHOULD be inspected to determine and learn one or more IPv6 addresses for the far-end CE. If the local AC supports IND (e.g., a switched Frame Relay AC), the packet SHOULD be forwarded to the local CE after modifying the link-layer address options to match the type of the local AC.

If the local AC does not support IND, processing of the packet depends on whether the PE has learned at least one interface address for its directly attached CE.

- o If it has learned at least one IPv6 address for the CE, the PE MUST discard the Inverse Neighbor Solicitation (INS) and generate an Inverse Neighbor Advertisement (INA) back into the pseudowire. The destination address of the INA is the source address from the INS; the source address is one of the local CE's interface addresses; and all the local CE's interface addresses that have been learned so far SHOULD be included in the Target Address List. The Source and Target link-layer addresses are copied from the INS. In addition, the PE SHOULD generate ND advertisements on the local AC using the IPv6 address of the remote CE and the link-layer address of the local PE.

- o If it has not learned at least one IPv6 and link-layer address of its directly connected CE, the INS MUST continue to be discarded until the PE learns an IPv6 and link-layer address from the local CE (through receiving, for example, a Neighbor Solicitation). After this has occurred, the PE will be able to respond to INS messages received over the pseudowire as described above.

4.3.6. Processing of Inverse Neighbor Advertisements (INAs)

An INA received on an AC from a local CE SHOULD be inspected to determine and learn one or more IPv6 addresses for the CE. It MUST then be forwarded unmodified over the pseudowire.

An INA received over the pseudowire SHOULD be inspected to determine and learn one or more IPv6 addresses for the far-end CE.

If the local AC supports IND (e.g., a Frame Relay AC), the packet MAY be forwarded to the local CE after modifying the link-layer address options to match the type of the local AC.

If the local AC does not support IND, the PE MUST discard the INA and generate a Neighbor Advertisement (NA) towards its local CE. The source IPv6 address of the NA is the source IPv6 address from the INA; the destination IPv6 address is the destination IPv6 address from the INA; and the link-layer address is that of the local AC on the PE.

4.3.7. Processing of Router Solicitations

A Router Solicitation received on an AC from a local CE SHOULD be inspected to determine and learn an IPv6 address for the CE and, if present, the link-layer address of the CE. It MUST then be forwarded unmodified over the pseudowire.

A Router Solicitation received over the pseudowire SHOULD be inspected to determine and learn an IPv6 address for the far-end CE. If a source link-layer address option is present, the PE MUST remove it. The PE MAY substitute a source link-layer address option specifying the link-layer address of its local AC. The packet is then forwarded to the local CE.

4.3.8. Processing of Router Advertisements

A Router Advertisement received on an AC from a local CE SHOULD be inspected to determine and learn an IPv6 address for the CE and, if present, the link-layer address of the CE. It MUST then be forwarded unmodified over the pseudowire.

A Router Advertisement received over the pseudowire SHOULD be inspected to determine and learn an IPv6 address for the far-end CE. If a source link-layer address option is present, the PE MUST remove it. The PE MAY substitute a source link-layer address option specifying the link-layer address of its local AC. If an MTU option is present, the PE MAY reduce the specified MTU if the MTU of the pseudowire is less than the value specified in the option. The packet is then forwarded to the local CE.

4.3.9. Duplicate Address Detection

Duplicate Address Detection [RFC4862] allows IPv6 hosts and routers to ensure that the addresses assigned to interfaces are unique on a link. As with all Neighbor Discovery packets, those used in Duplicate Address Detection will simply flow through the pseudowire, being inspected at the PEs at each end. Processing is performed as detailed in Sections 4.3.3 and 4.3.4. However, the source IPv6 address of Neighbor Solicitations used in Duplicate Address Detection is the unspecified address, so the PEs cannot learn the CE's IPv6 interface address (nor would it make sense to do so, given that at least one address is tentative at that time).

4.3.10. CE Address Discovery for CEs Attached Using PPP

The IPv6 Control Protocol (IPv6CP) [RFC5072] describes a procedure for establishing and configuring IPv6 on a point-to-point connection, including the negotiation of a link-local interface identifier. As in the case of IPv4, when such an AC is configured for IP interworking, PPP negotiation is not performed end-to-end between CE devices. Instead, PPP negotiation takes place between the CE and its local PE. The PE performs proxy PPP negotiation and informs the attached CE of the link-local identifier of its local interface using the Interface-Identifier option (0x01). This local interface identifier is used by stateless address autoconfiguration [RFC4862].

When a PPP link completes IPv6CP negotiations and the PPP link is open, a PE MAY discover the IPv6 unicast address of the CE using any of the mechanisms described above.

5. CE IPv4 Address Signaling between PEs

5.1. When to Signal an IPv4 Address of a CE

A PE device advertises the IPv4 address of the attached CE only when the encapsulation type of the pseudowire is IP Layer2 Transport (the value 0x000B, as defined in [RFC4446]). The IP Layer2 transport PW is also referred to as IP PW and is used interchangeably in this document. It is quite possible that the IPv4 address of a CE device

is not available at the time the PW labels are signaled. For example, in Frame Relay, the CE device sends an Inverse ARP request only when the Data Link Connection Identifier (DLCI) is active. If the PE signals the DLCI to be active only when it has received the IPv4 address along with the PW Forwarding Equivalence Class (FEC) from the remote PE, a deadlock situation arises. In order to avoid such problems, the PE MUST be prepared to advertise the PW FEC before the IPv4 address of the CE is known; hence, the PE uses an IPv4 address value zero. When the IPv4 address of the CE device does become available, the PE re-advertises the PW FEC along with the IPv4 address of the CE.

Similarly, if the PE detects that an IP address of a CE is no longer valid (by methods described above), the PE MUST re-advertise the PW FEC with a null IP address to denote the withdrawal of the IP address of the CE. The receiving PE then waits for notification of the remote IP address. During this period, propagation of unicast IPv4 traffic is suspended, but multicast IPv4 traffic can continue to flow between the AC and the pseudowire.

If two CE devices are locally attached to the PE on disparate AC types (for example, one CE connected to an Ethernet port and the other to a Frame Relay port), the IPv4 addresses are learned in the same manner as described above. However, since the CE devices are local, the distribution of IPv4 addresses for these CE devices is a local step.

Note that the PEs discover the IPv6 addresses of the remote CE by intercepting Neighbor Discovery and Inverse Neighbor Discovery packets that have been passed in-band through the pseudowire. Hence, there is no need to communicate the IPv6 addresses of the CEs through LDP signaling.

If the pseudowire is carrying both IPv4 and IPv6 traffic, the mechanisms used for IPv6 and IPv4 SHOULD NOT interact. In particular, just because a PE has learned a link-layer address for IPv6 traffic by intercepting a Neighbor Advertisement from its directly connected CE, it SHOULD NOT assume that it can use that link-layer address for IPv4 traffic until that fact is confirmed by reception of, for example, an IPv4 ARP message from the CE.

5.2. LDP-Based Distribution of CE IPv4 Addresses

[RFC4447] uses Label Distribution Protocol (LDP) transport to exchange PW FECs in the Label Mapping message in the Downstream Unsolicited (DU) mode. The PW FEC comes in two flavors, with some

common fields between them: Pwid and Generalized ID FEC elements. The discussions below refer to these common fields for IP L2 Interworking encapsulation.

In addition to PW FEC, this document uses an IP Address List TLV (as defined in [RFC5036]) that is to be included in the optional parameter field of the Label Mapping message when advertising the PW FEC for the IP Layer2 Transport. The use of optional parameters in the Label Mapping message to extend the attributes of the PW FEC is specified in [RFC4447].

As defined in [RFC4447], when processing a received PW FEC, the PE matches the PW ID and PW type with the locally configured PW ID and PW Type. If there is a match and if the PW Type is IP Layer2 Transport, the PE further checks for the presence of an Address List TLV [RFC5036] in the optional parameter TLVs. The processing of the Address List TLV is as follows.

- o If a PE is configured for an AC to a CE enabled for IPv4 or dual-stack IPv4/IPv6, the PE SHOULD advertise an Address List TLV with address family type of IPv4 address. The PE SHOULD process the IPv4 Address List TLV as described in this document. The PE MUST advertise and process IPv6 capability using the procedures described in Section 6.
- o If a PE does not receive any IPv4 address in the Address List TLV, it MAY assume IPv4 behavior. The address resolution for IPv4 MUST then depend on local manual configuration. In the case of mismatched configuration whereby one PE has manual configuration while the other does not, the IP address to link-layer address mapping remains unresolved, resulting in unsuccessful propagation of IPv4 traffic to the local CE.
- o If a PE is configured for an AC to a CE enabled for IPv6 only, the PE MUST advertise IPv6 capability using the procedures described in Section 6. In addition, by virtue of not setting the manual configuration for IPv4 support, IPv6-only support is realized.

We use the Address List TLV [RFC5036] to signal the IPv4 address of the local CE. This IP Address List TLV is included in the optional parameter field of the Label Mapping message.

The Address List TLV is only used for IPv4 addresses.

The fields of the IP Address List TLV are set as follows:

Length

Set to 6 to encompass 2 bytes of Address Family field and 4 bytes of Addresses field (because a single IPv4 address is used).

Address Family

Set to 1 to indicate IPv4 as defined in [RFC5036].

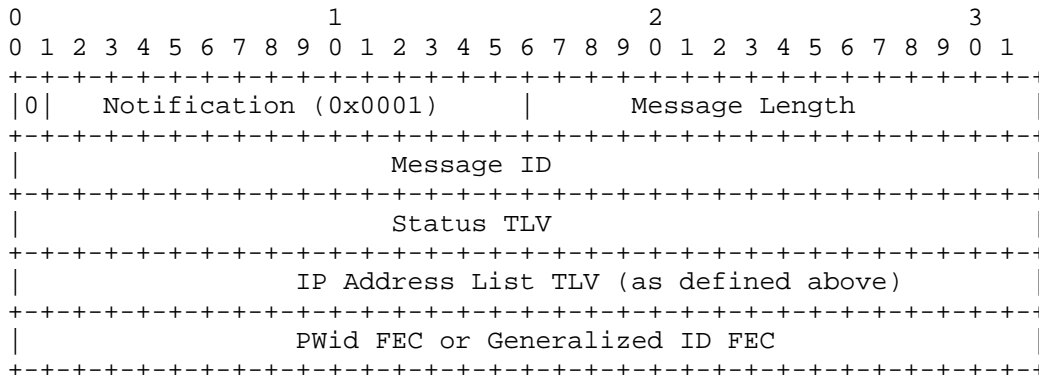
Addresses

Contains a single IPv4 address that is the address of the CE attached to the advertising PE.

The address in the Addresses field is set to all zeros to denote that the advertising PE has not learned the IPv4 address of its local CE. Any non-zero address value denotes the IPv4 address of the advertising PE's attached CE device.

The IPv4 address of the CE is also supplied in the optional parameters field of the LDP Notification message along with the PW FEC. The LDP Notification message is used to signal any change in the status of the CE's IPv4 address.

The encoding of the LDP Notification message is as follows.




The Status TLV status code is set to 0x0000002C "IP address of CE", to indicate that an IP address update follows. Since this notification does not refer to any particular message, the Message ID field is set to 0.

The PW FEC TLV SHOULD NOT include the interface parameters as they are ignored in the context of this message.

6. IPv6 Capability Advertisement

A Stack Capability Interface Parameter sub-TLV is signaled by the two PEs so that they can agree which network protocol(s) they SHOULD be using. As discussed earlier, the use of the Address List TLV signifies support for IPv4 stack, so the Stack Capability Interface Parameter sub-TLV is used to indicate whether support for IPv6 stack is required on a given IP PW.

The Stack Capability Interface Parameter sub-TLV is part of the interface parameters. The proposed format for the Stack Capability Interface Parameter sub-TLV is as follows:



Parameter ID = 0x16

Length = 4

The Stack Capability field is a bit field. Only one bit is defined in this document. When bit zero (the least significant bit with bitmask 0x0001) is set, it indicates IPv6 Stack Capability.

The presence of the Stack Capability Interface Parameter sub-TLV is relevant only when the PW type is IP PW. A PE that supports IPv6 on an IP PW MUST signal the Stack Capability Interface Parameter sub-TLV in the initial Label Mapping message for the PW. The PE nodes compare the value advertised by the remote PE with the local configuration and only use a capability that is supported by both.

The behavior of a PE that does not understand an Interface Parameter sub-TLV is specified in Section 5.5 of RFC 4447 [RFC4447].

In some deployment scenarios, it may be desirable to take a PW operationally down if there is a mismatch of the Stack Capability between the PEs. In other deployment scenarios, an operator may wish the IP version supported by both PEs to fall back to IPv4 if one of the PEs does not support IPv6. The following procedures MUST be followed for each of these cases.

6.1. PW Operational Down on Stack Capability Mismatch

If a PE that supports IPv6 and has not yet sent a Label Mapping message receives an initial Label Mapping message from the far-end PE that does not include the Stack Capability Interface Parameter sub-TLV, or one is received but it is not set to the 'IPv6 Stack Capability' value, then the PE supporting this procedure MUST NOT send a Label Mapping message for this PW.

If a PE that supports IPv6 has already sent an initial Label Mapping message for the PW and does not receive a Stack Capability Interface Parameter sub-TLV in the Label Mapping message from the far-end PE, or one is received but it is not set to 'IPv6 Stack Capability', the PE supporting this procedure MUST withdraw its PW label with the LDP status code meaning "IP Address type mismatch" (Status Code 0x0000004A). However, subsequently, if the configuration was to change at the far-end PE and a Stack Capability Interface Parameter sub-TLV in the Label Mapping message is received from the far-end PE, the local PE MUST re-advertise the Label Mapping message for the PW.

6.2. Stack Capability Fallback

If a PE that supports IPv6 and has not yet sent a Label Mapping message receives an initial Label Mapping message from the far-end PE that does not include the Stack Capability Interface Parameter sub-TLV, or one is received but it is not set to the 'IPv6 Stack Capability' value, then it MAY send a Label Mapping message for this PW but MUST NOT include the Stack Capability Interface Parameter sub-TLV.

If a PE that supports IPv6 and has already sent a Label Mapping message for the PW with the Stack Capability Interface Parameter sub-TLV but does not receive a Stack Capability Interface Parameter sub-TLV from the far-end PE in the initial Label Mapping message (or one is received but it is not set to the 'IPv6 Stack Capability' value), the PE following this procedure MUST send a Label Withdraw for its PW label with the LDP status code meaning "Wrong IP Address type" (Status Code 0x0000004B) followed by a Label Mapping message that does not include the Stack Capability Interface Parameter sub-TLV. If a Label Withdraw message with the "Wrong IP Address Type" status code is received by a PE, it SHOULD treat this as a normal Label Withdraw but MUST NOT respond with a Label Release. It MUST continue to wait for the next control message for the PW as specified in Section 6.2 of RFC 4447 [RFC4447].

7. IANA Considerations

7.1. LDP Status Messages

This document uses new LDP status codes. IANA already maintains a registry of name "Status Code Name Space" defined by [RFC5036]. The following values have been assigned:

```
0x0000002C "IP Address of CE"
0x0000004A "IP Address Type Mismatch"
0x0000004B "Wrong IP Address Type"
```

7.2. Interface Parameters

This document proposes a new Interface Parameters sub-TLV, that has been assigned from the 'Pseudowire Interface Parameters Sub-TLV type Registry'. The following value has been assigned for the Parameter ID:

```
0x16 "Stack Capability"
```

IANA has also set up a registry of "L2VPN PE stack Capabilities". This is a 16-bit field. Stack Capability bitmask 0x0001 is specified in Section 6 of this document. The remaining bitfield values (0x0002,...,0x8000) are to be assigned by IANA using the "IETF Review" policy defined in [RFC5226].

L2VPN PE Stack Capabilities:

Bit (Value)	Description
=====	=====
Bit 0 (0x0001)	- IPv6 stack capability
Bit 1 (0x0002)	- Unassigned
Bit 2 (0x0004)	- Unassigned
.	.
.	.
.	.
Bit 14 (0x4000)	- Unassigned
Bit 15 (0x8000)	- Unassigned

8. Security Considerations

The security aspect of this solution is addressed for two planes: the control plane and the data plane.

8.1. Control Plane Security

Control plane security pertains to establishing the LDP connection and to pseudowire signaling and CE IP address distribution over that LDP connection. For greater security, the LDP connection between two trusted PEs MUST be secured by each PE verifying the incoming connection against the configured address of the peer and authenticating the LDP messages, as described in Section 2.9 of [RFC5036]. Pseudowire signaling between two secure LDP peers does not pose a security issue but mis-wiring could occur due to configuration error. However, the fact that the pseudowire will only be established if the two PEs have matching configurations (e.g., PW ID, PW type, and MTU) provides some protection against mis-wiring due to configuration errors.

Learning the IP address of the appropriate CE can be a security issue. It is expected that the Attachment Circuit to the local CE will be physically secured. If this is a concern, the PE MUST be configured with the IP and MAC address of the CE when connected with Ethernet, IP and virtual circuit information (DLCI or VPI/VCI (Virtual Path Identifier / Virtual Circuit Identifier) when connected over Frame Relay or ATM, and IP address only when connected over PPP. During ARP/Inverse ARP frame processing, the PE MUST verify the received information against local configuration before forwarding the information to the remote PE to protect against hijacking of the connection.

For IPv6, the preferred means of security is Secure Neighbor Discovery (SEND) [RFC3971]. SEND provides a mechanism for securing Neighbor Discovery packets over media (such as wireless links) that may be insecure and open to packet interception and substitution. SEND is based upon cryptographic signatures of Neighbor Discovery packets. These signatures allow the receiving node to detect packet modification and confirm that a received packet originated from the claimed source node. SEND is incompatible with the Neighbor Discovery packet modifications described in this document. As such, SEND cannot be used for Neighbor Discovery across an ARP Mediation pseudowire. PEs taking part in IPv6 ARP Mediation MUST remove all SEND packet options from Neighbor Discovery packets before forwarding into the pseudowire. If the CE devices are configured to accept only SEND Neighbor Discovery packets, Neighbor Discovery will fail. Thus, the CE devices MUST be configured to accept non-SEND packets, even if they treat them with lower priority than SEND packets. Because SEND cannot be used in combination with IPv6 ARP Mediation, it is suggested that IPv6 ARP Mediation only be used with secure Attachment Circuits. An exception to this recommendation applies to an implementation that supports the SEND Proxy [RFC6496], which allows a device such as PE to act as an ND proxy as described in [RFC6496].

8.2. Data Plane Security

The data traffic between CE and PE is not encrypted, and it is possible that in an insecure environment, a malicious user may tap into the CE-to-PE connection and generate traffic using the spoofed destination MAC address on the Ethernet Attachment Circuit. In order to avoid such hijacking, the local PE may verify the source MAC address of the received frame against the MAC address of the admitted connection. The frame is forwarded to the PW only when authenticity is verified. When spoofing is detected, the PE MUST sever the connection with the local CE, tear down the PW, and start over.

9. Acknowledgements

The authors would like to thank Yetik Serbest, Prabhu Kavi, Bruce Lasley, Mark Lewis, Carlos Pignataro, and others who participated in the discussions related to this document.

10. Contributors

This document is the combined effort of many who have contributed, carefully reviewed, and provided technical clarifications. This includes the individuals listed in this section and those listed in the Editors' Addresses.

Matthew Bocci
Alcatel-Lucent
EMail: Mathew.bocci@alcatel-lucent.com

Tiberiu Grigoriu
Alcatel-Lucent
EMail: Tiberiu.Grigoriu@alcatel-lucent.com

Neil Hart
Alcatel-Lucent
EMail: Neil.Hart@alcatel-lucent.com

Andrew Dolganow
Alcatel-Lucent
EMail: Andrew.Dolganow@alcatel-lucent.com

Shane Amante
Level 3
EMail: Shane@castlepoint.net

Toby Smith
Google
EMail: tob@google.com

Andrew G. Malis
Verizon
EMail: Andy.g.Malis@verizon.com

Steven Wright
Bell South Corp
EMail: steven.wright@bellsouth.com

Waldemar Augustyn
Consultant
EMail: waldemar@wdmsys.com

Arun Vishwanathan
Juniper Networks
EMail: arunvn@juniper.net

Ashwin Moranganti
IneoQuest Technologies
EMail: Ashwin.Moranganti@Ineoquest.com

11. References

11.1. Normative References

- [RFC826] Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, November 1982.
- [RFC2390] Bradley, T., Brown, C., and A. Malis, "Inverse Address Resolution Protocol", RFC 2390, September 1998.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, April 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, October 2007.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

- [RFC3122] Conta, A., "Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification", RFC 3122, June 2001.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

11.2. Informative References

- [RFC4664] Andersson, L., Ed., and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006.
- [RFC1332] McGregor, G., "The PPP Internet Protocol Control Protocol (IPCP)", RFC 1332, May 1992.
- [RFC5072] Varada, S., Ed., Haskins, D., and E. Allen, "IP Version 6 over PPP", RFC 5072, September 2007.
- [RFC925] Postel, J., "Multi-LAN address resolution", RFC 925, October 1984.
- [RFC1256] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [RFC5309] Shen, N., Ed., and A. Zinin, Ed., "Point-to-Point Operation over LAN in Link State Routing Protocols", RFC 5309, October 2008.
- [RFC6496] Krishnan, S., Laganier, J., Bonola, M., and A. Garcia-Martinez, "Secure Proxy ND Support for SEcure Neighbor Discovery (SEND)", RFC 6496, February 2012.

Appendix A. Use of IGP with IP L2 Interworking L2VPNs

In an IP L2 interworking L2VPN, when an IGP on a CE connected to a broadcast link is cross-connected with an IGP on a CE connected to a point-to-point link, there are routing protocol related issues that MUST be addressed. The link state routing protocols are cognizant of the underlying link characteristics and behave accordingly when establishing neighbor adjacencies, representing the network topology, and passing protocol packets. The point-to-point operations of the routing protocols over a LAN are discussed in [RFC5309].

A.1. OSPF

The OSPF protocol treats a broadcast link type with a special procedure that engages in Neighbor Discovery to elect a designated router and a backup designated router (DR and BDR, respectively), with which each other router on the link forms adjacencies. However, these procedures are neither applicable nor understood by OSPF running on a point-to-point link. By cross-connecting two neighbors with disparate link types, an IP L2 interworking L2VPN may experience connectivity issues.

Additionally, the link type specified in the router Link State Advertisement (LSA) will not match for the two cross-connected routers.

Finally, each OSPF router generates network LSAs when connected to a broadcast link such as Ethernet, receipt of which by an OSPF router that believes itself to be connected to a point-to-point link further adds to the confusion.

Fortunately, the OSPF protocol provides a configuration option (`ospfIfType`) whereby OSPF will treat the underlying physical broadcast link as a point-to-point link.

It is strongly recommended that all OSPF protocols on CE devices connected to Ethernet interfaces use this configuration option when attached to a PE that is participating in an IP L2 Interworking VPN.

A.2. RIP

The RIP protocol broadcasts RIP advertisements every 30 seconds. If the multicast/broadcast traffic snooping mechanism is used as described in Section 4.1, the attached PE can learn the local CE router's IP address from the IP header of its advertisements. No special configuration is required for RIP in this type of Layer 2 IP Interworking L2VPN.

A.3. IS-IS

The IS-IS protocol does not encapsulate its PDUs in IP; hence, it cannot be supported in IP L2 Interworking L2VPNs.

Editors' Addresses

Himanshu Shah (editor)
Ciena
EMail: hshah@ciena.com

Eric Rosen (editor)
Cisco Systems
EMail: erosen@cisco.com

Giles Heron (editor)
Cisco Systems
EMail: giheron@cisco.com

Vach Kompella (editor)
Alcatel-Lucent
EMail: vach.kompella@alcatel-lucent.com