Internet Engineering Task Force (IETF)
Request for Comments: 7949
Updates: 5838
Category: Standards Track
ISSN: 2070-1721

I. Chen
Ericsson
A. Lindem
Cisco
R. Atkinson
Consultant

August 2016

# OSPFv3 over IPv4 for IPv6 Transition

## Abstract

This document defines a mechanism to use IPv4 to transport OSPFv3 packets.  Using OSPFv3 over IPv4 with the existing OSPFv3 Address Family extension can simplify transition from an OSPFv2 IPv4-only routing domain to an OSPFv3 dual-stack routing domain.  This document updates RFC 5838 to support virtual links in the IPv4 unicast address family when using OSPFv3 over IPv4.

## Status of This Memo

## Copyright Notice

Table of Contents

## 1. Introduction

Using OSPFv3 [RFC5340] over IPv4 [RFC791] with the existing OSPFv3 address family extension can simplify transition from an IPv4-only routing domain to an IPv6 [RFC2460] or dual-stack routing domain. Dual-stack routing protocols, such as the Border Gateway Protocol [RFC4271], have an advantage during the transition, because both IPv4 and IPv6 address families can be advertised using either IPv4 or IPv6 transport. Some IPv4-specific and IPv6-specific routing protocols share enough similarities in their protocol packet formats and protocol signaling that it is trivial to deploy an initial IPv6 routing domain by transporting the routing protocol over IPv4, thereby allowing IPv6 routing domains to be deployed and tested before decommissioning IPv4 and moving to an IPv6-only network.

In the case of the Open Shortest Path First (OSPF) interior gateway routing protocol (IGP), OSPFv2 [RFC2328] is the IGP deployed over IPv4, while OSPFv3 [RFC5340] is the IGP deployed over IPv6. OSPFv3 further supports multiple address families [RFC5838], including both the IPv6 unicast address family and the IPv4 unicast address family. Consequently, it is possible to deploy OSPFv3 over IPv4 without any changes to either OSPFv3 or IPv4. During the transition to IPv6, future OSPF extensions can focus on OSPFv3, and OSPFv2 can move to maintenance mode.

This document specifies how to use IPv4 to transport OSPFv3 packets. The mechanism takes advantage of the fact that OSPFv2 and OSPFv3 share the same IP protocol number, 89. Additionally, the OSPF packet header for both OSPFv2 and OSPFv3 includes the OSPF header version (i.e., the field that distinguishes an OSPFv2 packet from an OSPFv3 packet) in the same location (i.e., the same offset from the start of the header).

If the IPv4 topology and IPv6 topology are not identical, the most likely cause is that some parts of the network deployment have not yet been upgraded to support both IPv4 and IPv6. In situations where the IPv4 deployment is a superset of the IPv6 deployment, it is expected that OSPFv3 packets would be transported over IPv4, until the rest of the network deployment is upgraded to support IPv6 in addition to IPv4. In situations where the IPv6 deployment is a superset of the IPv4 deployment, it is expected that OSPFv3 would be transported over IPv6.

Throughout this document, "OSPF" is used when the text applies to both OSPFv2 and OSPFv3. "OSPFv2" or "OSPFv3" is used when the text is specific to one version of the OSPF protocol. Similarly, "IP" is used when the text describes either version of the Internet Protocol. "IPv4" or "IPv6" is used when the text is specific to a single version of the Internet Protocol.

### 1.1. IPv4-Only Use Case

OSPFv3 only requires IPv6 link-local addresses to form adjacencies, and does not require IPv6 global-scope addresses to establish an IPv6 routing domain. However, IPv6 over Ethernet [RFC2464] uses a different EtherType (0x86dd) from IPv4 (0x0800) and the Address Resolution Protocol (ARP) (0x0806) [RFC826] used with IPv4.

Some existing deployed link-layer equipment only supports IPv4 and ARP. Such equipment contains hardware filters keyed on the EtherType field of the Ethernet frame to filter which frames will be accepted by that link-layer equipment. Because IPv6 uses a different EtherType, IPv6 framing for OSPFv3 will not work with that equipment. In other cases, Point-to-Point Protocol (PPP) might be used over a serial interface, but again only IPv4 over PPP might be supported over such an interface. It is hoped that equipment with such limitations will be eventually upgraded or replaced.

In some locations, especially locations with less communications infrastructure, satellite communications (SATCOM) are used to reduce deployment costs for data networking. SATCOM often has lower cost to deploy than running new copper or optical cables over long distances to connect remote areas. Also, in a wide range of locations including places with good communications infrastructure, Very Small Aperture Terminals (VSATs) often are used by banks and retailers to connect their branches and stores to a central location.

Some widely deployed VSAT equipment has either (A) Ethernet interfaces that only support the Ethernet Address Resolution Protocol (ARP) and IPv4, or (B) serial interfaces that only support IPv4 and PPP packets.  Such deployments and equipment still can deploy and use OSPFv3 over IPv4 today, and then later migrate to OSPFv3 over IPv6 after equipment is upgraded or replaced.  This can have lower operational costs than running OSPFv2 and then trying to make a flag-day switch to OSPFv3.  By running OSPFv3 over IPv4 now, the eventual transition to dual-stack, and then to IPv6-only, can be orchestrated.

## 2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

## 3. Encapsulation in IPv4

An OSPFv3 packet can be directly encapsulated within an IPv4 packet as the payload, without the IPv6 packet header, as illustrated in Figure 1.  For OSPFv3 transported over IPv4, the IPv4 packet has an IPv4 protocol type of 89, denoting that the payload is an OSPF packet.  The payload of the IPv4 packet consists of an OSPFv3 packet, beginning with the OSPF packet header having its OSPF version field set to 3.

An OSPFv3 packet followed by an OSPF link-local signaling (LLS) extension data block [RFC5613] encapsulated in an IPv4 packet is illustrated in Figure 2.

Since an IPv4 header without options is only 20 octets long and is shorter than an IPv6 header, an OSPFv3 packet encapsulated in a 20-octet IPv4 header is shorter than an OSPFv3 packet encapsulated in an IPv6 header.  Consequently, the link MTU for IPv6 is sufficient to transport an OSPFv3 packet encapsulated in a 20-octet IPv4 header. If the link MTU is not sufficient to transport an OSPFv3 packet in IPv4, then OSPFv3 can rely on IP fragmentation and reassembly [RFC791].

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ ---
|   4   |  IHL  |Type of Service|          Total Length         |  ^
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  |
|         Identification        |Flags|      Fragment Offset    |  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  |
|  Time to Live | Protocol (89) |         Header Checksum       | IPv4
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ Header
|                       Source Address                          |  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  |
|                    Destination Address                        |  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  |
|                    Options                    |    Padding    |  v
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ ---
|       3       |     Type      |         Packet length         |  ^
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  |
|                         Router ID                             | OSPFv3
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ Header
|                          Area ID                              |  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  |
|          Checksum             |  Instance ID  |      0        |  v
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ ---
|                        OSPFv3 Body ...                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Note: "IHL" stands for Internet Header Length.

Figure 1: An IPv4 Packet Encapsulating an OSPFv3 Packet

```
+---------------+
| IPv4 Header   |
+---------------+
| OSPFv3 Header |
|...............|
|               |
| OSPFv3 Body   |
|               |
+---------------+
|               |
| LLS Data      |
|               |
+---------------+
```

Figure 2: The IPv4 Packet Encapsulating an OSPFv3 Packet with a Trailing OSPF Link-Local Signaling Data Block

### 3.1. Source Address

For OSPFv3 over IPv4, the source address is the primary IPv4 address for the interface over which the packet is transmitted. All OSPFv3 routers on the link should share the same IPv4 subnet for IPv4 transport to function correctly.

While OSPFv2 operates on a subnet, OSPFv3 operates on a link [RFC5340]. Accordingly, an OSPFv3 router implementation MAY support adjacencies with OSPFv3 neighbors on different IPv4 subnets. If this is supported, the IPv4 data plane MUST resolve IPv4 addresses to Layer 2 addresses using ARP on multi-access networks and point-to-point over LAN [RFC5309] for direct next hops on different IPv4 subnets. When OSPFv3 adjacencies on different IPv4 subnets are supported, Bidirectional Forwarding Detection (BFD) [RFC5881] cannot be used for adjacency loss detection since BFD is restricted to a single subnet.

### 3.2. Destination Address

As defined in OSPFv2, the IPv4 destination address of an OSPF protocol packet is either an IPv4 multicast address or the IPv4 unicast address of an OSPFv2 neighbor. Two well-known link-local multicast addresses are assigned to OSPFv2, the AllSPFRouters address (224.0.0.5) and the AllDRouters address (224.0.0.6). The multicast address used depends on the OSPF packet type, the OSPF interface type, and the OSPF router’s role on multi-access networks.

Thus, for an OSPFv3-over-IPv4 packet to be sent to AllSPFRouters, the destination address field in the IPv4 packet MUST be 224.0.0.5. For an OSPFv3-over-IPv4 packet to be sent to AllDRouters, the destination address field in the IPv4 packet MUST be 224.0.0.6.

When an OSPF router sends a unicast OSPF packet over a connected interface, the destination of such an IP packet is the address assigned to the receiving interface. Thus, a unicast OSPFv3 packet transported in an IPv4 packet would specify the OSPFv3 neighbor’s IPv4 address as the destination address.

### 3.3. OSPFv3 Header Checksum

For IPv4 transport, the pseudo-header used in the checksum calculation will contain the IPv4 source and destination addresses, the OSPFv3 protocol ID, and the OSPFv3 length from the OSPFv3 header (Appendix A.3.1 of [RFC5340]). The format is similar to the UDP pseudo-header as described in [RFC768] and is illustrated in Figure 3.

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Source Address                         |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                     Destination Address                       |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |     0         | Protocol (89) |     OSPFv3 Packet Length      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Pseudo-header for OSPFv3 over IPv4

## 3.4. Operation over Virtual Links

When an OSPF router sends an OSPF packet over a virtual link, the receiving router might not be directly connected to the sending router. Thus, the destination IP address of the IP packet must be a reachable unicast IP address for the virtual link endpoint. Because IPv6 is the presumed Internet protocol and an IPv4 destination is not routable, the OSPFv3 address family extension [RFC5838] specifies that only virtual links in the IPv6 address family are supported.

As illustrated in Figure 1, this document specifies OSPFv3 transport over IPv4. As a result, OSPFv3 virtual links can be supported with IPv4 address families by simply setting the IPv4 destination address to a reachable IPv4 unicast address for the virtual link endpoint. Hence, the restriction in Section 2.8 of RFC 5838 [RFC5838] is relaxed since virtual links can now be supported for IPv4 address families as long as the transport is also IPv4. If IPv4 transport, as specified herein, is used for IPv6 address families, virtual links cannot be supported. Hence, in OSPF routing domains that require virtual links, the IP transport MUST match the address family (IPv4 or IPv6).

# 4. Management Considerations

## 4.1. Coexistence with OSPFv2

Since OSPFv2 [RFC2328] and OSPFv3 over IPv4 as described herein use exactly the same protocol and IPv4 addresses, OSPFv2 packets may be delivered to the OSPFv3 process and vice versa. When this occurs, the mismatched protocol packets will be dropped due to validation of the version in the first octet of the OSPFv2/OSPFv3 protocol header. Note that this will not prevent the packets from being delivered to the correct protocol process as standard socket implementations will deliver a copy to each socket matching the selectors.

Implementations of OSPFv3 over IPv4 transport SHOULD implement separate counters for a protocol mismatch and SHOULD provide means to suppress the ospfIfRxBadPacket and ospfVirtIfRxBadPacket SNMP notifications as described in [RFC4750] and the ospfv3IfRxBadPacket and ospv3VirtIfRxBadPacket SNMP notifications as described in [RFC5643] when an OSPFv2 packet is received by the OSPFv3 process or vice versa.

## 5. Security Considerations

OSPFv3 [RFC5340] relies on IPsec [RFC4301] for authentication and confidentiality. "Authentication/Confidentiality in OSPFv3" [RFC4552] specifies how IPsec is used with OSPFv3 over IPv6 transport. In order to use OSPFv3 with IPv4 transport as specified herein, further work such as "Authentication/Confidentiality in OSPFv2" [IPsec-OSPF] would be required.

An optional OSPFv3 Authentication Trailer [RFC7166] also has been defined as an alternative to using IPsec. The calculation of the authentication data in the Authentication Trailer includes the source IPv6 address to protect an OSPFv3 router from man-in-the-middle attacks. For IPv4 encapsulation as described herein, the IPv4 source address should be placed in the first 4 octets of Apad followed by the hexadecimal value 0x878FE1F3 repeated (L-4)/4 times, where L is the length of the hash measured in octets.

The processing of the optional Authentication Trailer is contained entirely within the OSPFv3 protocol. In other words, each OSPFv3 router instance is responsible for the authentication, without involvement from IPsec or any other IP-layer function. Consequently, except for calculation of the Apad value, transporting OSPFv3 packets using IPv4 does not change the generation or validation of the optional OSPFv3 Authentication Trailer.

## 6. References

### 6.1. Normative References

[RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <http://www.rfc-editor.org/info/rfc791>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <http://www.rfc-editor.org/info/rfc2328>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[RFC5309] Shen, N., Ed., and A. Zinin, Ed., "Point-to-Point Operation over LAN in Link State Routing Protocols", RFC 5309, DOI 10.17487/RFC5309, October 2008, <http://www.rfc-editor.org/info/rfc5309>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <http://www.rfc-editor.org/info/rfc5340>.

[RFC5838] Lindem, A., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, DOI 10.17487/RFC5838, April 2010, <http://www.rfc-editor.org/info/rfc5838>.

### 6.2. Informative References

[IPsec-OSPF] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv2", Work in Progress, draft-gupta-ospf-ospfv2-sec-01, August 2009.

[RFC768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <http://www.rfc-editor.org/info/rfc768>.

[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, DOI 10.17487/RFC0826, November 1982, <http://www.rfc-editor.org/info/rfc826>.

[RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <http://www.rfc-editor.org/info/rfc2464>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <http://www.rfc-editor.org/info/rfc4271>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <http://www.rfc-editor.org/info/rfc4301>.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <http://www.rfc-editor.org/info/rfc4552>.

[RFC4750] Joyal, D., Ed., Galecki, P., Ed., Giacalone, S., Ed., Coltun, R., and F. Baker, "OSPF Version 2 Management Information Base", RFC 4750, DOI 10.17487/RFC4750, December 2006, <http://www.rfc-editor.org/info/rfc4750>.

[RFC5613] Zinin, A., Roy, A., Nguyen, L., Friedman, B., and D. Yeung, "OSPF Link-Local Signaling", RFC 5613, DOI 10.17487/RFC5613, August 2009, <http://www.rfc-editor.org/info/rfc5613>.

[RFC5643] Joyal, D., Ed., and V. Manral, Ed., "Management Information Base for OSPFv3", RFC 5643, DOI 10.17487/RFC5643, August 2009, <http://www.rfc-editor.org/info/rfc5643>.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <http://www.rfc-editor.org/info/rfc5881>.

[RFC7166] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 7166, DOI 10.17487/RFC7166, March 2014, <http://www.rfc-editor.org/info/rfc7166>.

## Acknowledgments

The authors would like to thank Alexander Okonnikov for his thorough review and valuable feedback and Suresh Krishnan for pointing out that clear specification for the pseudo-header used in the OSPFv3 packet checksum calculation was required.  The authors would also like to thank Wenhu Lu for acting as document shepherd.

## Authors' Addresses

Ing-Wher Chen
Ericsson

Email: ichen@kuatrotech.com

Acee Lindem
Cisco

Email: acee@cisco.com

RJ Atkinson
Consultant

Email: rja.lists@gmail.com