

Tactical Use of OSPFv3 over IPv4: Analysis and Implementation

I. Chen
MITRE Corporation
McLean, VA, USA
ingwherchen@mitre.org

Dr. Feng Xie
MITRE Corporation
McLean, VA, USA
fxie@mitre.org

Dr. Shukri Abdallah
MITRE Corporation
McLean, VA, USA
sabdallah@mitre.org

Abstract—The Open Shortest Path First (OSPF) interior routing protocol is an open-standard that is widely deployed in enterprise and tactical networks. OSPFv2 only supports IPv4, while OSPFv3 originally only supported IPv6. More recently, OSPFv3 standards have been enhanced so that OSPFv3 can carry both IPv4 and IPv6 topology information and so that OSPFv3 can be deployed over IPv4. By implementing RFC7949, OSPFv3 becomes the routing protocol of choice over both IPv4 and IPv6 transports. This paper shows that in low-speed tactical networks, replacing OSPFv2 with OSPFv3 over IPv4 transport not only facilitates the eventual migration to IPv6, but also enables reduction in overhead bandwidth when compared to OSPFv3 over IPv6 transport. The needed enhancements were implemented based on an open-source routing software called FRRouting.

Keywords—OSPFv2, OSPFv3, IPv4 transport, IPv6 Transport, RFC 7949, Address Families, FRRouting

I. INTRODUCTION

Due to the sunset of IPv4, the *Internet Engineering Task Force (IETF)*, that is the official standards body for the Internet, has reduced standardization efforts for IPv4-only protocol extensions. As a result of this policy, many extensions to the Open Shortest Path First (OSPF) interior routing protocol are standardized only in OSPF version 3 (OSPFv3) [1], and are not available in IPv4-only OSPF version 2 (OSPFv2) [2]. With the standardization of OSPFv3 over IPv4 transport in RFC7949 [3], it is now possible to deploy OSPFv3 using IPv4 transport, thereby allowing for the use of these IPv6-only OSPFv3 extensions in IPv4-only or in mixed IPv4+IPv6 network deployments, without any standardization delays and without requiring proprietary extensions. This presents an opportunity to immediately take advantage of OSPFv3-unique features useful for tactical environments, such as those designed to support Mobile Ad-hoc Networks (MANETs) [4][5]. At the same time, retaining the existing IPv4-only deployment allows for continued use of established addressing schemes, eliminates the additional overhead caused by the larger IPv6 header, reduces potential attack surface, and simultaneously better positions the network for future deployment of IPv6.

In the rest of the paper, Section II describes the evolution of OSPFv3 as the feature-rich routing protocol of choice for IPv4 and IPv6 networks. Section III describes the features essential and useful to tactical applications and how IPv4-only OSPFv2 and IPv6-only OSPFv3 affect their performance. Section IV presents a qualitative packet analysis. Section V identifies some

lessons learned from implementation of OSPFv3 over IPv4. Section VI provides conclusions to date and planned future work.

II. EVOLUTION OF OSPFv3

Originally, the OSPFv3 standard specified that OSPFv3 be carried only within IPv6 packets and specified that OSPFv3 only carry IPv6 routing information [6]. More recently, the IETF has standardized several extensions that enable OSPFv3 to fully support both IPv4 networks and IPv6 networks.

The OSPFv3 Address Families (AF) extension enables OSPFv3 to carry both IPv4 and IPv6 routing information [7]. The 8-bit instance id in the OSPFv3 header is partitioned into five ranges for four different address families, IPv6 unicast, IPv6 multicast, IPv4 unicast, and IPv4 multicast. Prefixes of different address families are encoded in various OSPFv3 Link State Advertisements (LSAs) in a Type-Length-Value format. This extension is widely supported in commercial IP routers, e.g., Arista, Cisco, Juniper, and Nokia.

More recently, the IETF standardized a method for carrying OSPFv3 over IPv4 in RFC7949. However, none of the previously mentioned commercial router vendors support this extension, possibly because most Internet Service Providers (ISPs) use IS-IS rather than OSPF as their interior routing protocol. In the open-source world, neither FRRouting (FRR) [8] nor Quagga [9] supports this extension at present.

By implementing these extensions, an IP router can use OSPFv3 to replace the older OSPFv2 routing protocol both in IPv4-only deployments and in mixed IPv4-IPv6 deployments, while also deploying new capabilities that are only available in OSPFv3, such as integrated support for MANET routing algorithms.

III. TACTICAL REQUIREMENTS AND MOTIVATIONS

A. Addressing Scheme and Migration

Due to the popularity of IP, the original 32-bit size of IPv4 addresses was determined in the early 1990s to be inadequate. IPv6 was standardized with 128-bit addresses in the mid-1990s and is intended to eventually replace IPv4. However, actual transition from IPv4 to IPv6 has been neither straight forward nor rapid.

Green-field networks and applications, such as 5G IoT networks [10], might more easily adapt to IPv6 than existing

IPv4 network deployments because it is unnecessary to renumber the networks, a labor-intensive and error-prone process. Many existing IPv4 networks deploy IPv4-IPv6 network address translation devices at their exterior edges as that can be simpler and lower-cost than either deploying IPv6 in parallel with IPv4 or transitioning to an IPv6-only network deployment.

Like many other IPv4 network deployments, tactical networks often have an existing IPv4 core that is difficult to migrate to IPv6 due to the sizable task of renumbering an existing IPv4 network. Furthermore, US DoD tactical networks do not suffer from a shortage of IPv4 address space, so they lack a common commercial motivation to migrate to dual IPv4+IPv6 or to IPv6-only network deployments. While many commercial networks are primarily built over high-speed optical links, tactical networks often rely upon relatively low-speed SATCOM and line-of-sight radio links. We also observe that many IPv4-only networks have deployed MAC-layer protocol security filters that block all IPv6 packets, even link-local IPv6 packets, to reduce the potential attack surface. By allowing OSPFv3 to run over IPv4 transport, the feature rich OSPFv3 replaces OSPFv2 in the network instead of the network migrating or converting to IPv6 addressing and IPv6 transport.

B. New Features

Tactical network migration to IPv6 might be driven by new standards-track features that specifically support MANETs by optimizing protocol synchronization and reducing protocol bandwidth requirement. However, the advantages of these new features to reduce bandwidth usage might be lost due to the much larger IPv6 header. Encapsulating OSPFv3 packets with a smaller IPv4 header is an advantage.

IV. PACKET ANALYSIS

We performed protocol data unit (PDU) analysis by comparing the size, in terms of byte count, of each protocol PDU that is expected on the wire, with OSPFv2 using IPv4 transport, OSPFv3 using IPv6 transport, and OSPFv3 using IPv4 transport enabled by RFC7949. To match OSPFv2's authentication mechanism that is embedded in OSPFv2 header, we assumed that OSPFv3 would use IPsec transport mode and Authentication Header (AH) [11].

Table 1 provides a summary of PDU byte count comparison. Each cell lists the byte count of a header or a packet. For some packets, the byte count is a function of the number of a sub-field within the packet. For example, the size of Database Description Packet (DBP) is the size of OSPF Packet Header (LPH) plus 8 bytes plus the byte count of n_1 LSA Header (LH) sub-field. As expected, OSPFv2 with IPv4 transport generates the smallest PDUs, and OSPFv3 with IPv6 transport generates the largest PDUs. OSPFv3 with IPv4 transport falls between OSPFv2/IPv4 and OSPFv3/IPv6. In general, OSPFv3 has slightly larger PDUs, but OSPFv3 with IPv4 transport benefits from the smaller IPv4 addresses that ultimately reduce the PDU size.

In Table 2, we took the packet captures from an OSPFv2 router within a simulation run of a network of six OSPFv2 routers and compared the transmitted OSPFv2 PDUs to what we expected using OSPFv3 over IPv4. For this comparison, we assumed the same sequence of events in both OSPFv2 and OSPFv3 over IPv4. In this simulated router, we would expect a 2% increase in protocol PDU total bytes for the duration of the simulation run of OSPFv2.

Table 1: PDU byte count comparison.

Type	OSPFv2/IPv4	OSPFv3/IPv6	OSPFv3/IPv4
IP Header	20	40	20
OSPF Packet Header (OPH)	24	16	16
Authentication	Included in OPH	24	24
OSPF Hello Packet (OHP)	20	20	20
LSA Header (LH)	20	20	20
Database Description Packet (DBP)	$OPH+8+n_1*LH$	$OPH+12+n_1*LH$	$OPH+12+n_1*LH$
Link State Request Packet (LSRP)	$OPH+n_2*12$	$OPH+n_2*12$	$OPH+n_2*12$
Link State acknowledgement Packet (LSAP)	$OPH+n_3*LH$	$OPH+n_3*LH$	$OPH+n_3*LH$
Link State Update Packet (LSUP)	$OPH+n_4*LSA$	$OPH+n_4*LSA$	$OPH+n_4*LSA$
Type-5 External LSAs	16	48	36
Type-4 Summary LSAs	8	12	12
Type-3 Summary LSAs	8	24	12
Type-2 Network LSAs	4	4	4
Type-1 Router LSAs	$4+n_5*8$	$4+n_5*16$	$4+n_5*16$
OSPFv3 Link LSA	Not applicable	44	20
OSPFv3 Intra-Area-Prefix LSA	In Router LSA	52	12+8

Table 2: Comparison of OSPFv2 packet captures from simulation with expected OSPFv3 over IPv4 PDU.

Type	OSPFv2/IPv4	OSPFv3/IPv4	Change over OSPFv2
Hello Packet (count)	349	349	Not applicable
Hello Packet (byte)	30634	36218	+2%
Database Description Packet (count)	11	11	Not applicable
Database Description Packet (byte)	1170	1390	+2%
Link State Request Packet (count)	3	3	Not applicable
Link State Request Packet (byte)	246	294	+2%
Link State Acknowledgement Packet (count)	8	8	Not applicable
Link State Acknowledgement Packet (byte)	716	844	+2%
Link State Update (count)	39	39	Not applicable
Link State Update (byte)	7086	9258	+3%
Total (byte)	39852	48004	+2%

V. PROTOTYPE IMPLEMENTATION

We have built a prototype implementation of RFC7949, based on FRR Release 7.3, an open-source IP routing stack. FRR, originally a fork of Quagga, was chosen over Quagga for multiple reasons: (1) FRR is fully open-source software that runs on a variety of UNIX/POSIX operating systems. (2) FRR provides an IETF standards-compliant implementation of OSPFv2/v3 with deployment experience both in commercial and tactical environments. (3) FRR is supported by a broad and active community of developers, vendors, and operators without being dependent upon any single company.

A. Configuration Design

The most challenging part of this prototype is to provide operators with the full control of the choice to carry OSPFv3 over IPv6 or over IPv4 transport. Control could be either at the OSPFv3 protocol/instance level or at the OSPFv3 interface level. After careful consideration, we decided that (a) the default OSPFv3 behavior shall use IPv6 as the transport to maintain interoperability with existing OSPFv3 deployments, (b) the operator shall be able to use IPv4 as the transport for an arbitrary set of OSPFv3 interfaces while using IPv6 for the compliment set of OSPFv3 interfaces, and (c) given the potential for a large number of OSPFv3 interfaces, the operator shall be able to change the OSPFv3 transport for multiple interfaces with a single command. This approach is consistent with the Principle of Least Surprise, while also supporting easy deployment of OSPFv3 in IPv4-only tactical environments.

The implementation supports commands to configure OSPFv3 transport both at the protocol/instance level and at the interface level. If there is no configuration command for the transport, then OSPFv3 is always carried over IPv6. If the only OSPFv3 transport configuration is at the protocol/instance level, then all interfaces inherit that setting. If the OSPFv3 transport control is configured at both levels, then the interface level configuration takes precedence over the protocol/instance configuration.

B. Operational Considerations

The operational aspects of changing OSPFv3 transport shall avoid or minimize the potential for either inconsistent or unexpected behavior. This complexity comes from the requirement to cope with various types of interface events, such as addition or removal of IPv4/v6 addresses associated with an interface, which are events that occur outside of OSPFv3. Accordingly, our prototype implements the following rules:

1. If the transport control is explicitly configured, either at the protocol/instance level or at the OSPFv3 interface level, the specified transport mechanism **MUST** be used.
2. If an interface is not assigned an IP address required for the transport control configuration (except for unnumbered IPv4 interfaces), then the OSPFv3 interface becomes inactive.
3. Removal of an IP address may result in de-activation of the corresponding OSPFv3 interface if the required transport configuration cannot be satisfied.
4. Addition of an IPv4/IPv6 address may result in re-activation of the corresponding OSPFv3 interface if the required transport configuration can be satisfied.
5. Changing the OSPFv3 transport at the process/instance level may result in activation/de-activation of the associated OSPFv3 interfaces depending on whether the new transport requirement on each interface can be satisfied.
6. Changing the OSPFv3 transport at the OSPFv3 interface level may result in activation/de-activation of the OSPFv3 interface depending on whether the new transport requirement on the interface can be satisfied.

C. Implementation Details

The modular design of OSPFv3 in FRR Release 7.3 makes the implementation of RFC7949 relatively straight forward. An indicator for the underlying IPv4 transport address was added to both the OSPFv3 interface and to the OSPFv3 neighbor structure. An AF_INET socket is added to handle OSPFv3 packets encapsulated in an IPv4 packet. When IPv4 is used as

the transport mechanism, a different pseudo-header is used to calculate the OSPFv3 checksum and the corresponding AF_INET socket is used for transmitting and receiving OSPFv3 packets. Additional configuration logic and operational rules are implemented as described previously.

The implementation improved FRR's sending of OSPFv3 packets by replacing the IPv6 link local address with an interface identifier. Replacing the address with the interface identifier increases modularity and reduces the potential for operational issues should an interface address change for any reason.

VI. CONCLUSIONS AND FUTURE WORK

Migration from IPv4 to IPv6 did not happen as rapidly as originally thought. In many deployments, both commercial and tactical, it is clear both IPv4 and IPv6 technologies will have to coexist for an extended period. Using IPv4 as the transport for OSPFv3 not only facilitates the eventual migration, but also enables reduction in overhead bandwidth when compared to OSPFv3 over IPv6 transport, which is extremely important for many mission-critical tactical deployments. Our effort is motivated by the absence of this feature in either COTS or open-source IP routers. We are the first to implement this capability and the first to publish analysis of the savings on the protocol PDUs. Future work may include experiments of the feature in an emulated mission critical environment, such as EMANE [12].

ACKNOWLEDGEMENTS

The view, opinions, and/or findings contained in this report are those of The MITRE Corporation and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

Approved for Public Release; Distribution Unlimited. Public Release Case Number 21-2244.

© 2021 The MITRE Corporation.

REFERENCES

- [1] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, IETF, Reston, VA, USA, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [2] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, IETF, Reston, VA, USA, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [3] Chen, I., Lindem, A., and R. Atkinson, "OSPFv3 over IPv4 for IPv6 Transition", RFC 7949, DOI 10.17487/RFC7949, IETF, Reston, VA, USA, August 2016, <<https://www.rfc-editor.org/info/rfc7949>>.
- [4] Ogier, R. and P. Spagnolo, "Mobile Ad Hoc Network (MANET) Extension of OSPF Using Connected Dominating Set (CDS) Flooding", RFC 5614, DOI 10.17487/RFC5614, IETF, Reston, VA, USA, August 2009, <<https://www.rfc-editor.org/info/rfc5614>>.
- [5] Ogier, R., "Use of OSPF-MDR in Single-Hop Broadcast Networks", RFC 7038, DOI 10.17487/RFC7038, IETF, Reston, VA, USA, October 2013, <<https://www.rfc-editor.org/info/rfc7038>>.
- [6] Crawford, M., Narten, T., and S. Thomas, "Transmission of IPv6 Packets over Token Ring Networks", RFC 2470, DOI 10.17487/RFC2470, IETF, Reston, VA, USA, December 1998, <<https://www.rfc-editor.org/info/rfc2470>>.
- [7] Lindem, A., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, DOI 10.17487/RFC5838, IETF, Reston, VA, USA, April 2010, <<https://www.rfc-editor.org/info/rfc5838>>.
- [8] "FRRouting open source router source code." [Online]. Available: <https://www.frrouting.org>
- [9] "Quagga open source router source code." [Online]. Available: <https://www.quagga.net>
- [10] ETSI GR IP6 011 v1.1.1 (2018-10) Group Report, "IPv6-Based 5G Mobile Wireless Internet; Deployment of IPv6-Based 5G Mobile Wireless Internet", Sophia Antipolis Cedex, France, October 2018.
- [11] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, IETF, Reston, VA, USA, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.
- [12] "The Extendable Mobile Ad-Hoc Network Emulator (EMANE)." [Online]. Available: <http://cs.itd.nrl.navy.mil/work/emane/index.php>