

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 10, 2013

A. Lindem
J. Arkko
Ericsson
July 9, 2012

OSPFv3 Auto-Configuration
draft-acee-ospf-ospfv3-autoconfig-03.txt

Abstract

OSPFv3 is a candidate for deployments in environments where auto-configuration is a requirement. One such environment is the IPv6 home network where users expect to simply plug in a router and have it automatically use OSPFv3 for intra-domain routing. This document describes the necessary mechanisms for OSPFv3 to be self-configuring.

Status of this Memo

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

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

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

This Internet-Draft will expire on January 10, 2013.

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.

1. Introduction

OSPFv3 [OSPFV3] is a candidate for deployments in environments where auto-configuration is a requirement. Its operation is largely unchanged from the base OSPFv3 protocol specification [OSPFV3].

The following aspects of OSPFv3 auto-configuration are described:

1. Default OSPFv3 Configuration
2. Unique OSPFv3 Router-ID generation
3. OSPFv3 Adjacency Formation
4. Duplicate OSPFv3 Router-ID Resolution

1.1. Requirements notation

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 [RFC-KEYWORDS].

1.2. Acknowledgments

This specification was inspired by the work presented in the Homenet working group meeting in October 2011 in Philadelphia, Pennsylvania. In particular, we would like to thank Fred Baker, Lorenzo Colitti, Ole Troan, Mark Townsley, and Michael Richardson.

Arthur Dimitrelis and Aidan Williams did prior work in OSPFv3 auto-configuration in the expired "Autoconfiguration of routers using a link state routing protocol" IETF Draft. There are many similarities between the concepts and techniques in this document.

Thanks for Abhay Roy and Manav Bhatia for comments regarding duplicate router-id processing.

Thanks for Alvaro Retana and Michael Barnes for comments regarding OSPFv3 Instance ID auto-configuration.

The RFC text was produced using Marshall Rose's xml2rfc tool.

3. OSPFv3 Router ID Selection

As OSPFv3 Router implementing this specification must select a unique Router-ID. A pseudo-random number SHOULD be used for the OSPFv3 Router-ID. The generation should be seeded with a variable that is likely to be unique in that environment. A good choice of seed would be some portion or hash of the Route-Hardware-Fingerprint as described in Section 5.2.2.

Since there is a possibility of a Router ID collision, duplicate Router ID detection and resolution are required as described in Section 5 and Section 5.3.

4. OSPFv3 Adjacency Formation

Since OSPFv3 uses IPv6 link-local addresses for all protocol messages other than message sent on virtual links (which are not applicable to auto-configuration), OSPFv3 adjacency formation can proceed as soon as a Router-ID has been selected and the IPv6 link-local address has completed Duplicate Address Detection (DAD) as specified in IPv6 Stateless Address Autoconfiguration [SLAAC]. Otherwise, there is no change to the OSPFv3 base specification except with respect to duplicate Router-ID detection and resolution as described in Section 5 and Section 5.3.

5. OSPFv3 Duplicate Router-ID Detection and Resolution

There are two cases of duplicate OSPFv3 Router-ID detection. One where the OSPFv3 router with the duplicate Router-ID is directly connected and one where it is not. In both cases, the resolution is for one of the routers with the duplicate OSPFv3 Router-ID to select a new one.

5.1. Duplicate Router-ID Detection for Neighbors

In this case, a duplicate Router-ID is detected if any valid OSPFv3 packet is received with the same OSPFv3 Router-ID but a different IPv6 link-local source address. Once that occurs, the OSPFv3 router with the numerically smaller IPv6 link-local address will need to select a new Router-ID as described in Section 5.3. Note that the fact that the OSPFv3 router is a neighbor on a non-virtual interface implies that the router is directly connected. An OSPFv3 router implementing this specification should assure that the inadvertent connection of multiple router interfaces to the same physical link in not misconstrued as detection of a different OSPFv3 router with a duplicate Router-ID.

5.2. Duplicate Router-ID Detection for OSPFv3 Routers that are not Neighbors

OSPFv3 Routers implementing auto-configuration, as specified herein, MUST originate an Auto-Config (AC) Link State Advertisement (LSA) including the Router-Hardware-Fingerprint Type-Length-Value (TLV). The Router-Hardware-Fingerprint TLV contains a variable length value that has a very high probability of uniquely identifying the advertising OSPFv3 router. An OSPFv3 router implementing this specification MUST compare a received self-originated Auto-Config LSA's Router-Hardware-Fingerprint TLV against its own router hardware fingerprint. If the fingerprints are not equal, there is a Router-ID conflict and the OSPFv3 Router with the numerically smaller router hardware fingerprint MUST select a new Router-ID as described in Section 5.3.

This new LSA is designated for information related to OSPFv3 Auto-configuration and, in the future, could be used other auto-configuration information, e.g., global IPv6 prefixes. However, this is beyond the scope of this document.

5.2.1. OSPFv3 Router Auto-Configuration LSA

The OSPFv3 Auto-Configuration (AC) LSA has a function code of TBD and the S1/S2 bits set to B'01' indicating Area Flooding Scope. The U bit will be set indicating that the OSPFv3 AC LSA should be flooded

neighbor adjacencies MUST be reestablished.

5.4. Change to Received Self-Originated LSA Processing

RFC 2328 [OSPFV2], Section 13.4, describes the processing of received self-originated LSAs. If the received LSA doesn't exist, the receiving router will purge it from the OSPF routing domain. If the LSA is newer than the version in the Link State Database (LSDB), the receiving router will originate a newer version by advancing the LSA sequence number and reflooding. Since it is possible for an auto-configured OSPFv3 router to choose a duplicate OSPFv3 Router-ID, OSPFv3 routers implementing this specification should detect when multiple instances of the same self-originated LSA are purged or reoriginated since this is indicative of an OSPFv3 router with a duplicate Router-ID in the OSPFv3 routing domain. When this condition is detected, the OSPFv3 Router SHOULD delay self-originated LSA processing for LSAs that have recently been purged or reflooded. This specification recommends 10 seconds as the interval defining recent self-originated LSA processing and an exponential back off of 1 to 8 seconds for the processing delay.

6. Security Considerations

A unique OSPFv3 Interface Instance ID is used for auto-configuration to prevent inadvertent OSPFv3 adjacency formation, see Section 2

The goals of security and complete OSPFv3 auto-configuration are somewhat contradictory. When no explicit security configuration takes place, auto-configuration implies that additional devices placed in the network are automatically adopted as a part of the network. However, auto-configuration can also be combined with password configuration (see below) or future extensions for automatic pairing between devices. These mechanisms can help provide an automatically configured, securely routed network.

It is RECOMMENDED that OSPFv3 routers supporting this specification also offer an option to explicitly configure a password for HMAC- SHA authentication as described in [OSPFV3-AUTH-TRAILER]. When configured, the password will be used on all auto-configured interfaces with the Security Association Identifier (SA ID) set to 1 and HMAC-SHA-256 will be used as the authentication algorithm.

7. Management Considerations

It is RECOMMENDED that OSPFv3 routers supporting this specification also allow explicit configuration of OSPFv3 parameters as specified in Appendix C of [OSPFV3]. This is in addition to the authentication key configuration recommended in Section 6. However, it is acknowledged that there may be some deployment scenarios where manual configuration is not required.

8. IANA Considerations

This specification allocates a new OSPFv3 LSA, OSPFv3 Auto-Configuration (AC) LSA, TBD, as described in Section 5.2.1.

This specification also creates a registry for OSPFv3 Auto-Configuration (AC) LSA TLVs. This registry should be placed in the existing OSPFv3 IANA registry, and new values can be allocated via IETF Consensus or IESG Approval.

Three initial values are allocated:

- o 0 is marked as reserved.
- o 1 is Router-Hardware-Fingerprint TLV (Section 5.2.2).
- o 65535 is an Auto-configuration-Experiment-TLV, a common value that can be used for experimental purposes.

9. Normative References

- [OSPFV2] Moy, J., "OSPF Version 2", RFC 2328, April 1998.
- [OSPFV3] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [OSPFV3-AF] Lindem, A., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, April 2010.
- [OSPFV3-AUTH-TRAILER] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 6506, February 2012.
- [RFC-KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.
- [SLAAC] Thomson, S., Narten, T., and J. Tatuya, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [TE] Katz, D., Yeung, D., and K. Kompella, "Traffic Engineering Extensions to OSPF", RFC 3630, September 2003.

Authors' Addresses

Acee Lindem
Ericsson
102 Carric Bend Court
Cary, NC 27519
USA

Email: acee.lindem@ericsson.com

Jari Arkko
Ericsson
Jorvas, 02420
Finland

Email: jari.arkko@piuha.net

Copyright Notice

Copyright (c) 2014 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.

Ning So
Tata Communications
2613 Fairbourne Cir.
Plano, TX 75082
USA

Email: ningso01@gmail.com

Vic Liu
China Mobile
No.32 Xuanwumen West Street, Xicheng District
Beijing, 100053
China

Email: liuzhiheng@chinamobile.com

Mehmet Toy
Comcast
1800 Bishops Gate Blvd.
Mount Laurel, NJ 08054
USA

Email: mehmet_toy@cable.comcast.com

Lei Liu
UC Davis
CA
USA

Email: liulei.kddi@gmail.com

Renwei Li
Huawei Technologies
2330 Central expressway
Santa Clara, CA
USA

Email: renwei.li@huawei.com

Gregory Cauchie
FRANCE

Email: greg.cauchie@gmail.com

Ning So
Tata Communications
2613 Fairbourne Cir.
Plano, TX 75082
USA

Email: ningso01@gmail.com

Vic Liu
China Mobile
No.32 Xuanwumen West Street, Xicheng District
Beijing, 100053
China

Email: liuzhiheng@chinamobile.com

Lei Liu
UC Davis
USA

Email: liulei.kddi@gmail.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 14, 2013

A. Retana
Hewlett-Packard Co.
A. Lindem
Ericsson
July 13, 2012

OSPF Incremental Link State Database Synchronization
draft-retana-ospf-ils-01

Abstract

The ability of OSPF to transport non-routing information to be used by other applications was defined by the Opaque LSA Option. In order to not impact the convergence of routing information, this document describes a simple process to incrementally synchronize the routing and non-routing information residing in an OSPF link-state database.

Status of this Memo

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

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

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

This Internet-Draft will expire on January 14, 2013.

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
2. Requirements Language	3
3. Incremental LSDB Synchronization Process	3
3.1. Graceful Restart	4
3.2. Flooding and Database Synchronization	5
4. Backward Compatibility	5
5. IANA Considerations	6
6. Security Considerations	6
7. Acknowledgements	6
8. References	6
8.1. Normative References	6
8.2. Informative References	6
Appendix A. Changes between the -00 and -01 versions.	7
Authors' Addresses	7

1. Introduction

Opaque LSAs [RFC5250] provide the ability for OSPFv2 [RFC2328] to transport non-routing information to be used by other applications. A similar capability exists in OSPFv3 [RFC5340] through the use of the U-bit and an appropriate LSA Function Code. Throughout this document Opaque LSAs and ones with unrecognized link-state types will be referred to simply as "opaque".

The presence of opaque information in the OSPF Link-State Database (LSDB) may result in longer database exchange times, especially in cases where the amount of data is significantly larger than the routing-specific information. In order to not impact the convergence of routing information, this document describes a simple process to incrementally synchronize the information residing in an OSPF LSDB. The process uses existing mechanisms.

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. Incremental LSDB Synchronization Process

The Incremental LSDB Synchronization (ILS) process consists of the following steps:

LSA Prioritization

The contents of the local LSDB are classified to determine which LSAs require prioritized synchronization.

In general, LSAs containing routing-specific information MUST be classified as requiring prioritized synchronization, while other LSAs MAY be classified as not requiring it.

All routers in the OSPF routing domain should use the same criteria for LSA prioritization. While this is not required for correct OSPF operation, it is required for deterministic operation of the applications using the opaque LSAs.

Prioritized LSDB Synchronization

This step corresponds to the adjacency establishment process as described in [RFC2328].

LSAs classified as not requiring prioritized synchronization MUST NOT be included in Database Description (DBD) Packets during the Database Exchange Process. The OSPF routing table structure SHOULD be calculated before moving on to the next step.

Final LSDB Synchronization

In this step, any remaining LSAs in the LSDB SHOULD be synchronized. The routers MUST use the Out-of-Band LSDB Resynchronization [RFC4811] (OOB Resync) mechanism, which provides a way to resynchronize the LSDB without affecting the advertised neighbor state.

The process is described in terms of LSAs containing (or not) routing-specific information, but it may be generalized to include any other criteria considered significant in the local network and protocol instance.

The last step in the process MAY be used recursively to achieve an incremental LSDB synchronization through different types of data, making it also applicable to environments where only non-routing information exists.

3.1. Graceful Restart

The restart of the OSPF software in a router also presents an opportunity for LSDB synchronization. Because the restarting router is still in the forwarding path, it is important for the routing information in the LSDB to be synchronized as fast as possible. ILS can be used, with minor modifications, to reduce the synchronization time and the probability of network topology changes.

Graceful OSPF Restart

Graceful OSPF Restart for OSPFv2 [RFC3623] and OSPFv3 [RFC5187] don't specify any changes to the adjacency establishment process.

ILS can be used by the Helper Neighbor during the Grace Period; if so, then the Helper Node MUST include any Grace-LSAs in the DBD Packets during the Prioritized LSDB Synchronization step.

OSPF Restart Signaling

OSPF Restart Signaling [RFC4812] defines a mechanism to inform neighbors about a local restart, in which the LSDB synchronization is achieved using OOB Resync. In other words, the Prioritized LSDB Synchronization step would use OOB Resync if the non-restarting router uses ILS. No other

changes to the process are needed.

3.2. Flooding and Database Synchronization

In order to maintain OSPF reliable flooding, separate neighbor flooding state must be maintained for the lower priority LSAs and used when determining whether a lower priority LSA has been flooded or put on a neighbor's retransmission list.

The rules for flooding lower-priority LSAs and deleting lower priority MaxAge LSAs are modified for ILS synchronization. Neighbor state MUST be maintained as to whether low-priority LSAs have been synchronized with a given neighbor. In this document, this lower priority state will be referred to as ILS-state. This ILS-state will not advance until a neighbor reaches Full state and will return to Down ILS-state should the neighbor fall from Full state.

If a given neighbor's ILS-state is Exchange or higher, then lower priority LSAs should be flooded to that neighbor. This similar to the general LSA flooding rules in Section 13.1 of [RFC2328].

Consistent with Section 14.1 in [RFC2328], a lower priority MaxAge LSA cannot be removed the Link State Database if any the router is in Exchange or Loading ILS-state.

If multiple ILS synchronizations are used for synchronization different priorities of LSAs, separate ILS-state MUST be maintained for each priority and the previously described rules MUST be applied at each priority.

4. Backward Compatibility

The operation of ILS depends on the support of OOB Resync during synchronization; no backwards compatibility issues exist there [RFC4811]. If OOB Resync is not supported by one of the routers, then the LSDB synchronization would fall back to the adjacency establishment process as described in [RFC2328].

If OOB Resync is supported, but ILS has not been implemented by all the routers involved, the operation is still backwards compatible. Note that the process (Section 3) depends on the database description by the local router. In other words, a router may decide to not fully describe the contents of its LSDB to its neighbor during the adjacency establishment process, and later use OOB Resync to incrementally describe the difference; the receiver doesn't need to be aware of ILS. The benefits of ILS may only be partially realized if not supported by all the routers involved in synchronization.

5. IANA Considerations

This memo includes no request to IANA.

6. Security Considerations

The process described in this document does not introduce any new security issues into the OSPF protocol.

7. Acknowledgements

The authors would like to thank Abhay Roy and Liem Nguyen for their comments, and Dimitri Papadimitriou for his comments and for providing the motivation for this document.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4811] Nguyen, L., Roy, A., and A. Zinin, "OSPF Out-of-Band Link State Database (LSDB) Resynchronization", RFC 4811, March 2007.

8.2. Informative References

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC3623] Moy, J., Pillay-Esnault, P., and A. Lindem, "Graceful OSPF Restart", RFC 3623, November 2003.
- [RFC4812] Nguyen, L., Roy, A., and A. Zinin, "OSPF Restart Signaling", RFC 4812, March 2007.
- [RFC5187] Pillay-Esnault, P. and A. Lindem, "OSPFv3 Graceful Restart", RFC 5187, June 2008.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

Appendix A. Changes between the -00 and -01 versions.

- o Added a new section explaining how to maintain reliable flooding for the low priority LSAS.
- o Updated authors and contact information.

Authors' Addresses

Alvaro Retana
Hewlett-Packard Co.
2610 Wycliff Road
Raleigh, NC 27607
USA

Email: alvaro.retana@hp.com

Acee Lindem
Ericsson
102 Carric Bend Court
Cary, NC 27519
USA

Email: acee.lindem@ericsson.com

