

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 8, 2013

A. Lindem
J. Arkko
Ericsson
October 5, 2012

OSPFv3 Auto-Configuration
draft-ietf-ospf-ospfv3-autoconfig-00.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 April 8, 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
1.1.	Requirements notation	3
1.2.	Acknowledgments	3
2.	OSPFv3 Default Configuration	4
2.1.	Wait Timer Reduction	4
3.	OSPFv3 Router ID Selection	6
4.	OSPFv3 Adjacency Formation	7
5.	OSPFv3 Duplicate Router-ID Detection and Resolution	8
5.1.	Duplicate Router-ID Detection for Neighbors	8
5.2.	Duplicate Router-ID Detection for OSPFv3 Routers that are not Neighbors	8
5.2.1.	OSPFv3 Router Auto-Configuration LSA	8
5.2.2.	Router-Hardware-Fingerprint TLV	10
5.3.	Duplicate Router-ID Resolution	10
5.4.	Change to Received Self-Originated LSA Processing	11
6.	Security Considerations	12
7.	Management Considerations	13
8.	IANA Considerations	14
9.	References	15
9.1.	Normative References	15
9.2.	Informative References	15
	Authors' Addresses	16

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.

Thanks to Faraz Shamim for review and comments.

Thanks to Mark Smith for the requirement to reduce the adjacency formation delay in the back-to-back ethernet topologies that are prevalent in home networks.

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

2. OSPFv3 Default Configuration

For complete auto-configuration, OSPFv3 will need to choose suitable configuration defaults. These include:

1. Area 0 Only - All auto-configured OSPFv3 interfaces MUST be in area 0.
2. OSPFv3 SHOULD be auto-configured on for IPv6 on all interfaces intended as general IPv6-capable routers. Optionally, an interface MAY be excluded if it is clear that running OSPFv3 on the interface is not required. For example, if manual configuration or an other condition indicates that an interface is connected to an Internet Service Provider (ISP) and there is no Border Gateway Protocol (BGP) [[BGP](#)] peering, there is typically no need to employ OSPFv3. However, note that in many environments it can be useful to test whether an OSPFv3 adjacency can be established. In home networking environments, an interface where no OSPFv3 neighbors are found but a DHCP prefix can be acquired may be considered as an ISP interface.
3. OSPFv3 interfaces will be auto-configured to an interface type corresponding to their layer-2 capability. For example, Ethernet interfaces will be auto-configured as broadcast networks and Point-to-Point Protocol (PPP) interfaces will be auto-configured as Point-to-Point interfaces. Most extant OSPFv3 implementations do this already.
4. OSPFv3 interfaces MUST use the default HelloInterval, 10 seconds, and RouterDeadInterval, 40 seconds, as suggested in [Appendix C](#) of [[OSPFV3](#)].
5. All OSPFv3 interfaces SHOULD be auto-configured to use an Interface Instance ID of 0 that corresponds to the base IPv6 unicast address family instance ID as defined in [[OSPFV3-AF](#)]. Similarly, if IPv4 unicast addresses are advertised in a separate auto-configured OSPFv3 instance, the base IPv4 unicast address family instance ID value, i.e., 64, SHOULD be auto-configured as the Interface Instance ID for all interfaces corresponding to the OSPFv3 instance [[OSPFV3-AF](#)].

2.1. Wait Timer Reduction

In many situations, auto-configured OSPFv3 routers will be deployed in environments where back-to-back ethernet connections are utilized. When this is the case, an OSPFv3 broadcast interface will not come up until the other OSPFv3 router is connected and the routers will wait RouterDeadInterval seconds before forming an adjacency [[OSPFV2](#)]. In

order to reduce this delay, an auto-configured OSPFv3 router MAY reduce the wait interval to a value no less than (HelloInterval + 1), i.e., 11 seconds. Reducing the setting will slightly increase the likelihood of the Designated Router (DR) flapping but is preferable to the long adjacency formation delay. Note that this value is not included in OSPFv3 Hello packets and does not impact interoperability.

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 is 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

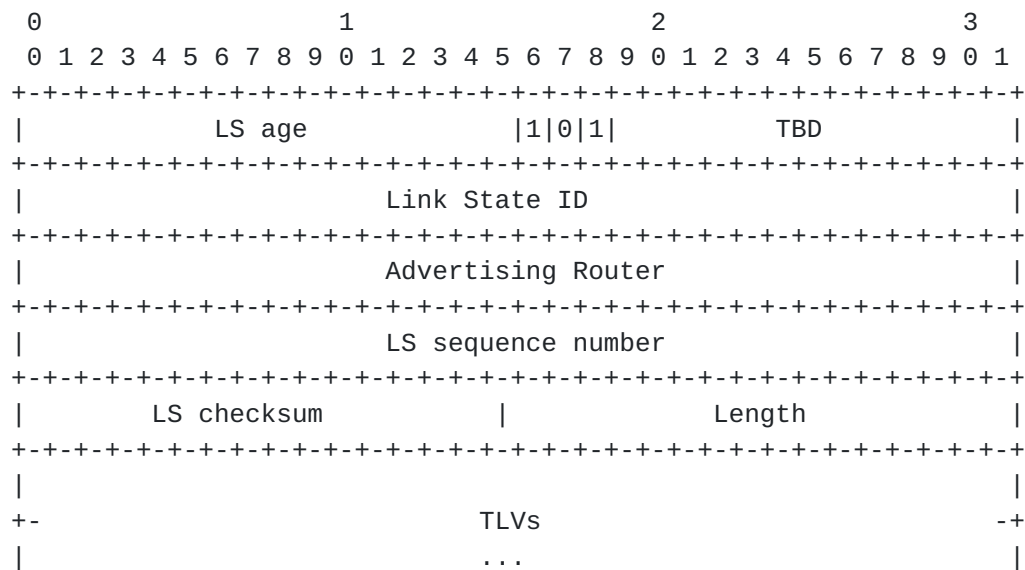
OSPFv3 Routers implementing auto-configuration, as specified herein, MUST originate an Auto-Config (AC) Link State Advertisement (LSA) including the Router-Hardware-Fingerprint Type-Length-Value (TLV). The Router-Hardware-Fingerprint TLV contains a variable length value that has a very high probability 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

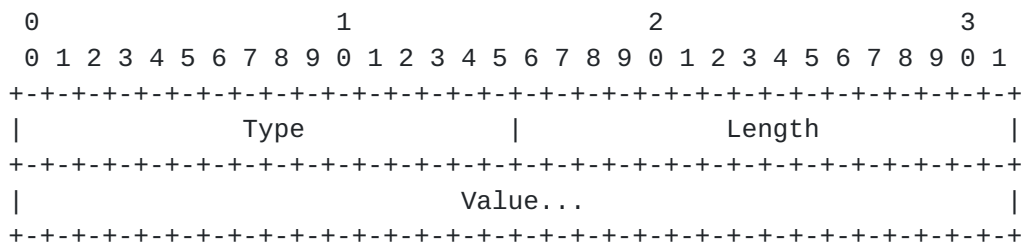
The OSPFv3 Auto-Configuration (AC) LSA has a function code of TBD and the S2/S1 bits set to 01 indicating Area Flooding Scope. The U bit will be set indicating that the OSPFv3 AC LSA should be flooded even

if it is not understood. The Link State ID (LSID) value will be a integer index used to discriminate between multiple AC LSAs originated by the same OSPF Router. This specification only describes the contents of an AC LSA with a Link State ID (LSID) of 0.



OSPFv3 Auto-Configuration (AC) LSA

The format of the TLVs within the body of an AC LSA is the same as the format used by the Traffic Engineering Extensions to OSPF [\[TE\]](#). The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets. The format of each TLV is:



TLV Format

The Length field defines the length of the value portion in octets (thus a TLV with no value portion would have a length of 0). The TLV is padded to 4-octet alignment; padding is not included in the length

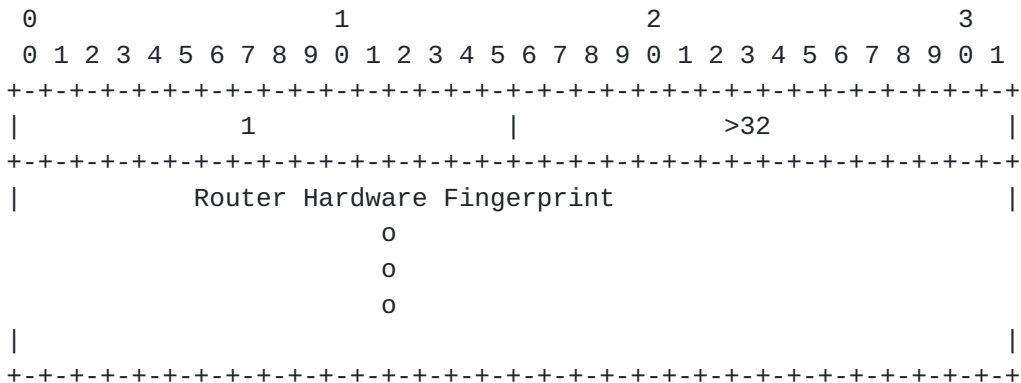
field (so a 3-octet value would have a length of 3, but the total size of the TLV would be 8 octets). Nested TLVs are also 32-bit aligned. For example, a 1-byte value would have the length field set to 1, and 3 octets of padding would be added to the end of the value portion of the TLV. Unrecognized types are ignored.

The new LSA is designated for information related to OSPFv3 Auto-configuration and, in the future, can be used other auto-configuration information, e.g., global IPv6 prefixes.

5.2.2. Router-Hardware-Fingerprint TLV

The Router-Hardware-Fingerprint TLV is the first TLV defined for the OSPFv3 Auto-Configuration (AC) LSA. It will have type 1 and MUST be advertised in the LSID OSPFv3 AC LSA with an LSID of 0. It SHOULD occur, at most, once and the first instance of the TLV will take precedence over preceding TLV instances. The length of the Router-Hardware-Fingerprint is variable but must be 32 bytes or greater.

The contents of the hardware fingerprint should be some combination of MAC addresses, CPU ID, or serial number(s) that provides an extremely high probability of uniqueness. It MUST be based on hardware attributes that will not change across hard and soft restarts.



Router-Hardware-Fingerprint TLV Format

5.3. Duplicate Router-ID Resolution

The OSPFv3 Router selected to resolve the duplicate OSPFv3 Router-ID condition must select a new OSPFv3 Router-ID. After selecting a new Router-ID, the Router-LSA with the prior duplicate Router-ID MUST be purged. all self-originated LSAs MUST be reoriginated, and any OSPFv3

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. References

9.1. 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.

9.2. Informative References

- [BGP] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.

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

