Network Working Group
Internet-Draft
Updates: 5340 (if approved)
Intended status: Standards Track
Expires: August 14, 2015

A. Lindem
Cisco Systems
J. Arkko
Ericsson

February 10, 2015

OSPFv3 Auto-Configuration
draft-ietf-ospf-ospfv3-autoconfig-15.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. This document updates RFC 5340 by relaxing the HelloInterval/ RouterDeadInterval checking during OSPFv3 adjacency formation and adding hysteresis to the update of self-originated Link State Advertisements (LSAs).

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 August 14, 2015.

Copyright Notice

Table of Contents

1. Introduction

OSPFv3 [OSPFV3] is a candidate for deployments in environments where auto-configuration is a requirement. This document describes extensions to OSPFv3 to enable it to operate in these environments. In this mode of operation, the protocol is largely unchanged from the base OSPFv3 protocol specification [OSPFV3]. Since the goals of auto-configuration and security can be conflicting, operators and network administrators should carefully consider their security requirements before deploying the solution described in this document. Refer to Section 8 for more information.

The following aspects of OSPFv3 auto-configuration are described in this document:

1. Default OSPFv3 Configuration

2. HelloInterval/RouterDeadInterval Flexibility

3. Unique OSPFv3 Router ID generation

4. OSPFv3 Adjacency Formation

5. Duplicate OSPFv3 Router ID Resolution

6. Self-Originated LSA Processing

OSPFv3 [OSPFV3] is updated by allowing OSPFv3 adjacencies to be formed between OSPFv3 routers with differing HelloIntervals or RouterDeadIntervals (refer to Section 3). Additionally, hysteresis has been added to the processing of stale self-originated LSAs to mitigate the flooding overhead created by an OSPFv3 Router with a duplicate OSPFv3 Router ID in the OSPFv3 routing domain (refer to Section 7.4. Both updates are fully backward compatible.

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].

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 all IPv6-capable interface on the router. An interface MAY be excluded if it is clear that running OSPFv3 on the interface is not required. For example, if manual configuration or another condition indicates that an interface is connected to an Internet Service Provider (ISP), there is typically no need to employ OSPFv3. In fact, [IPv6-CPE] specifically requires that IPv6 Customer Premise Equipment (CPE) routers do not initiate any dynamic routing protocol by default on the router’s WAN, i.e., ISP-facing, interface. In home networking environments, an interface where no OSPFv3 neighbors are found but a DHCP IPv6 prefix can be acquired may be considered an ISP-facing interface and running OSPFv3 is unnecessary.

3. OSPFv3 interfaces will be auto-configured to an interface type corresponding to their layer-2 capability. For example, Ethernet interfaces and Wi-Fi interfaces will be auto-configured as OSPFv3 broadcast networks and Point-to-Point Protocol (PPP) interfaces will be auto-configured as OSPFv3 Point-to-Point interfaces. Most extant OSPFv3 implementations do this already. Auto-configured operation over wireless networks requiring a point-to-multipoint (P2MP) topology and dynamic metrics based on wireless feedback is not within the scope of this document. However, auto-configuration is not precluded in these environments.

4. OSPFv3 interfaces MAY use an arbitrary HelloInterval and RouterDeadInterval as specified in Section 3. Of course, an identical HelloInterval and RouterDeadInterval will still be required to form an adjacency with an OSPFv3 router not supporting auto-configuration [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 IPv4 unicast OSPFv3 instance [OSPFV3-AF].

3. OSPFv3 HelloInterval/RouterDeadInterval Flexibility

Auto-configured OSPFv3 routers will not require an identical HelloInterval and RouterDeadInterval to form adjacencies. Rather, the received HelloInterval will be ignored and the received RouterDeadInterval will be used to determine OSPFv3 liveliness with the sending router. In other words, the Neighbor Inactivity Timer (Section 10 of [OSPFV2]) for each neighbor will reflect that neighbor’s advertised RouterDeadInterval and MAY be different from other OSPFv3 routers on the link without impacting adjacency formation. A similar mechanism requiring additional signaling is proposed for all OSPFv2 and OSPFv3 routers [ASYNC-HELLO].

3.1. 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).

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.

4. OSPFv3 Minimal Authentication Configuration

In many deployments, the requirement for OSPFv3 authentication overrides the goal of complete OSPFv3 autoconfiguration. Therefore, it is RECOMMENDED that OSPFv3 routers supporting this specification minimally offer an option to explicitly configure a single password for HMAC-SHA authentication as described in [OSPFV3-AUTH-TRAILER]. It is RECOMMENDED that the password entered as ASCII hexadecimal digits and that 32 or more digits to facilitate a password with a high degree of entropy. 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 used as the authentication algorithm.

5. OSPFv3 Router ID Selection

An OSPFv3 router requires a unique Router ID within the OSPFv3 routing domain for correct protocol operation. Existing Router ID selection algorithms (section C.1 in [OSPFV2] and [OSPFV3]) are not viable since they are dependent on a unique IPv4 interface address which is not likely to be available in autoconfigured deployments. An OSPFv3 router implementing this specification will select a router-id that has a high probability of uniqueness. 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 the applicable OSPFv3 router deployment. A good choice of seed would be some portion or hash of the Router-Hardware-Fingerprint as described in Section 7.2.2.

Since there is a possibility of a Router ID collision, duplicate Router ID detection and resolution are required as described in Section 7 and Section 7.3. OSPFv3 routers SHOULD maintain the last successfully chosen Router ID in non-volatile storage to avoid collisions subsequent to when an autoconfigured OSPFv3 router is first added to the OSPFv3 routing domain.

6. OSPFv3 Adjacency Formation

Since OSPFv3 uses IPv6 link-local addresses for all protocol messages other than messages 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, the only changes to the OSPFv3 base specification are supporting HelloInterval/RouterDeadInterval flexibility as described in Section 3 and duplicate Router ID detection and resolution as described in Section 7 and Section 7.3.

7. 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 duplicate resolution is for one of the routers to select a new OSPFv3 Router ID.

7.1. 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 this occurs, the OSPFv3 router with the numerically smaller IPv6 link-local address will need to select a new Router ID as described in Section 7.3. 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 an OSPFv3 neighbor with a duplicate Router ID.

7.2. Duplicate Router ID Detection for Non-Neighbors

OSPFv3 routers implementing auto-configuration, as specified herein, MUST originate an Auto-Configuration (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 detect received Auto-Configuration LSAs with its Router ID specified in the LSA header. LSAs received with the local OSPFv3 Router’s Router ID in the LSA header are perceived as self-originated (see section 4.6 of [OSPFV3]). In these received Auto-Configuration LSAs, the Router-Hardware-Fingerprint TLV is compared against the OSPFv3 Router’s own router hardware fingerprint. If the fingerprints are not equal, there is a duplicate Router ID conflict and the OSPFv3 router with the numerically smaller router hardware fingerprint MUST select a new Router ID as described in Section 7.3.

This new LSA is designated for information related to OSPFv3 Auto-configuration and, in the future, could be used for other auto-

configuration information, e.g., global IPv6 prefixes. However, this is beyond the scope of this document.

7.2.1. 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 OSPFv3 router. This specification only describes the contents of an AC LSA with a Link State ID (LSID) of 0.

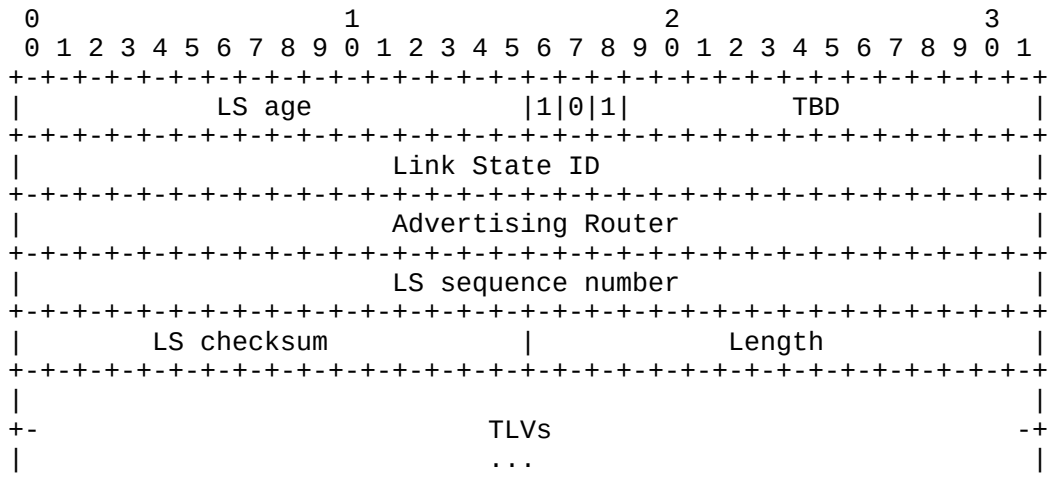

OSPFv3 Auto-Configuration (AC) LSA

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:

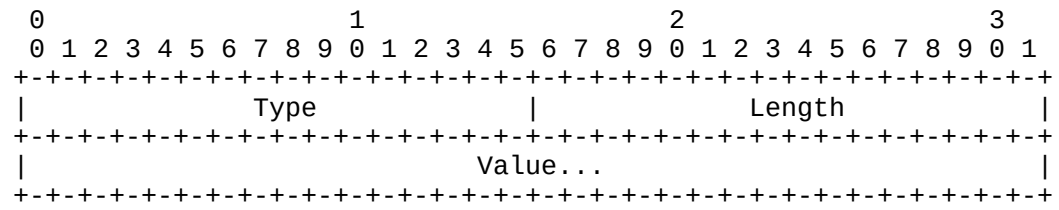

TLV Format

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.

7.2.2. 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 subsequent TLV instances. The length of the Router-Hardware-Fingerprint is variable but must be 32 octets or greater. If the Router-Hardware-Fingerprint TLV is not present as the first TLV, the AC-LSA is considered malformed and is ignored for the purposes of duplicate Router ID detection. Additionally, the event SHOULD be logged.

The contents of the hardware fingerprint MUST have an extremely high probability of uniqueness. It SHOULD be constructed from the concatenation of a number of local values that themselves have a high likelihood of uniqueness, such as MAC addresses, CPU ID, or serial numbers. It is RECOMMENDED that one or more available universal tokens (e.g., IEEE 802 48-bit MAC addresses or IEEE EUI-64 Identifiers [EUI64]) associated with the OSPFv3 router be included in the hardware fingerprint. It MUST be based on hardware attributes that will not change across hard and soft restarts.

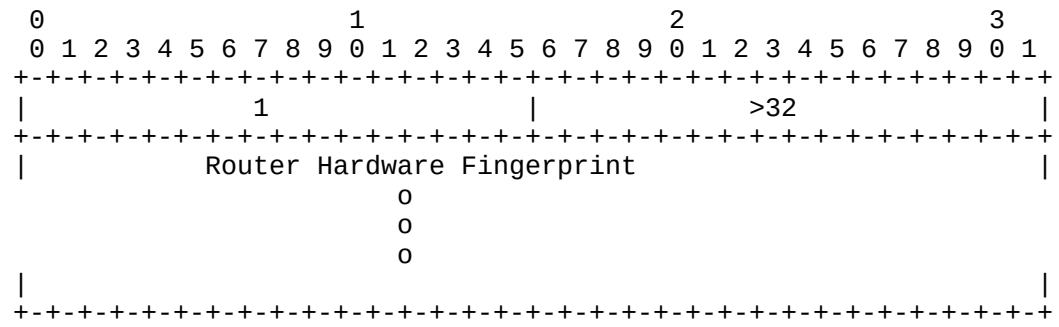

Router-Hardware-Fingerprint TLV Format

7.3. Duplicate Router ID Resolution

The OSPFv3 router selected to resolve the duplicate OSPFv3 Router ID condition must select a new OSPFv3 Router ID. The OSPFv3 router SHOULD reduce the possibility of a subsequent Router ID collision by checking the Link State Database for an OSPFv3 Auto-Configuration LSA with the newly selected Router ID and a different Router-Hardware-Fingerprint. If one is detected, a new Router ID should be selected without going through the resolution process Section 7. After selecting a new Router ID, all self-originated LSAs MUST be reoriginated, and any OSPFv3 neighbor adjacencies MUST be reestablished. The OSPFv3 router retaining the Router ID causing the conflict will reoriginate or purge stale any LSAs as described in Section 13.4 [OSPFV2].

7.4. Change to RFC 2328 Section 13.4, 'Receiving Self-Originated LSAs'

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 reoriginating. 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 reoriginated. 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.

This additional delay should allow for the mechanisms described in Section 7 to resolve the duplicate OSPFv3 Router ID conflict.

Since this mechanism is useful in mitigating the flooding overhead associated with the inadvertent or malicious introduction of an OSPFv3 router with a duplicate Router ID into an OSPFv3 routing domain, it MAY be deployed outside of autoconfigured deployments. The detection of a self-originated LSA that is being repeated reoriginated or purged SHOULD be logged.

8. 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 Section 4) or future extensions for automatic pairing between devices. These mechanisms can help provide an automatically configured, securely routed network.

In deployments where different authentication algorithm, per-interface keys, or encryption is required, OSPFv3 IPsec [OSPFV3-IPSEC] or alternate OSPFv3 Authentication trailer [OSPFV3-AUTH-TRAILER] algorithms MAY be used at the expense of additional configuration. The configuration and operational description of such deployments is beyond the scope of this document. However, a deployment could always revert to explicit configuration as described in Section 9 for features such as IPsec, per-interface keys, or alternate authentication algorithms.

The introduction, either malicious or accidental, of an OSPFv3 router with a duplicate Router ID is an attack point for OSPFv3 routing domains. This is due to the fact that OSPFv3 routers will interpret LSAs advertised by the router with the same Router ID as self-originated LSAs and attempt to purge them from the routing domain. The mechanisms in Section 7.4 will mitigate the effects of duplication.

9. Management Considerations

It is RECOMMENDED that OSPFv3 routers supporting this specification also support explicit configuration of OSPFv3 parameters as specified in Appendix C of [OSPFV3]. The would allow explicit override of autoconfigured parameters in situations where it is required (e.g.,

if the deployment requires multiple OSPFv3 areas). This is in addition to the authentication key configuration recommended in Section 4. Additionally, it is RECOMMENDED that OSPFv3 routers supporting this specification allow autoconfiguration to be completely disabled.

Since there is a small possibility of OSPFv3 Router ID collisions, manual configuration of OSPFv3 Router IDs is RECOMMENDED in OSPFv3 routing domains where route convergence due to a router ID change is intolerable.

OSPFv3 Routers supporting this specification MUST augment mechanisms for displaying or otherwise conveying OSPFv3 operational state to indicate whether or not the OSPFv3 router was autoconfigured and whether or not its OSPFv3 interfaces have been auto-configured.

10. IANA Considerations

This specification defines an OSPFv3 LSA Type for the OSPFv3 Auto-Configuration (AC) LSA, as described in Section 7.2.1. The value TBD will be allocated from the existing "OSPFv3 LSA Function Code" registry for the OSPFv3 Auto-Configuration LSA.

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 Review or, under exceptional circumstances, IESG Approval. [IANA-GUIDELINES]

Three initial values are allocated:

- 0 is marked as reserved.
- 1 is Router-Hardware-Fingerprint TLV (Section 7.2.2).
- 65535 is an Auto-configuration-Experiment-TLV, a common value that can be used for experimental purposes.

11. 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.

Thanks to Les Ginsberg for document review and recommendations on OSPFv3 hardware fingerprint content.

Thanks to Curtis Villamizar for document review and analysis of duplicate router-id resolution nuances.

Thanks to Uma Chunduri for comments during OSPF WG last call.

Thanks to Martin Vigoureux for Routing Area Directorate review and comments.

Thanks to Adam Montville for Security Area Directorate review and comments.

Thanks to Qin Wu for Operations & Management Area Directorate review and comments.

Thanks to Robert Sparks for General Area (GEN-ART) review and comments.

Thanks to Rama Darbha for review and comments.

Special thanks to Adrian Farrel for his in-depth review, copious comments, and suggested text.

Special thanks go to Markus Stenberg for his implementation of this specification in Bird.

Special thanks also go to David Lamparter for his implementation of this specification in Quagga.

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

12. References

12.1. 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 7166, 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.

12.2. Informative References

[ASYNC-HELLO] Anand, M., Grover, H., and A. Roy, "Asymmetric OSPF Hold Timer", draft-madhukar-ospf-agr-asymmetric-01.txt (work in progress), June 2013.

[EUI64] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", IEEE Tutorial http://standards.ieee.org/regauth/oui/tutorials/EUI64.html, March 1997.

[IANA-GUIDELINES] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Sectino in RFCs", RFC 5226, May 2008.

[IPv6-CPE]
Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, November 2013.

[OSPFV3-IPSEC]
Gupta, M. and S. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, June 2006.

Authors' Addresses

Acee Lindem
Cisco Systems
301 Midenhall Way
Cary, NC 27513
USA

Email: acee@cisco.com

Jari Arkko
Ericsson
Jorvas, 02420
Finland

Email: jari.arkko@piuha.net