

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 26, 2012

J. Arkko
A. Lindem
Ericsson
October 24, 2011

Prefix Assignment in a Home Network
draft-arkko-homenet-prefix-assignment-00

Abstract

This memo describes a prefix assignment mechanism for home networks. It is expected that home gateway routers are assigned an IPv6 prefix through DHCPv6 Prefix Delegation (PD). This prefix needs to be divided among the multiple subnets in a home network. This memo describes a mechanism for such division via OSPFv3. This is an alternative design to using DHCPv6 PD also for the prefix assignment. The memo is input to the working group so that it can make a decision on which type of design to pursue. It is expected that a routing-protocol based assignment uses a minimal amount of prefixes.

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 26, 2012.

Copyright Notice

Copyright (c) 2011 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

Internet-Draft

Homenet Prefixes

October 2011

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.	Role of Prefix Assignment	3
4.	Router Behavior	5
5.	Prefix Assignment in OSPFv3	6
5.1.	Usable Prefix TLV	7
5.2.	Assigned Prefix TLV	8
5.3.	OSPFv3 Prefix Assignment	9
6.	Manageability Considerations	11
7.	Security Considerations	11
8.	IANA Considerations	11
9.	References	12
9.1.	Normative References	12
9.2.	Informative References	12
Appendix A.	Acknowledgments	13
	Authors' Addresses	13

1. Introduction

This memo describes a prefix assignment mechanism for home networks. It is expected that home gateway routers are assigned an IPv6 prefix through DHCPv6 Prefix Delegation (PD) [[RFC3633](#)], or in some cases manually configured. This prefix needs to be divided among the multiple subnets in a home network. This memo describes a mechanism for such division via OSPFv3 [[RFC5340](#)].

The OSPFv3-based mechanism is an alternative design to using DHCPv6 PD also for the prefix assignment in the internal network. This memo has been written so that the working group can make a decision on which type of design to pursue. The main benefit of using a routing protocol to handle the prefix assignment is that it can provide a more efficient allocation mechanism than hierarchical assignment through DHCPv PD. This may be important for home networks that get only a /60 allocation from their ISPs.

The rest of this memo is organized as follows. [Section 2](#) defines the usual keywords, [Section 3](#) explains the main requirements for prefix assignments, [Section 4](#) describes how a home gateway router makes assignments when it itself has multiple subnets, and [Section 5](#) describes how the assignment can be performed in a distributed manner via OSPFv3 in the entire home network. Finally, [Section 6](#) explains what administrative interfaces are useful for advanced users that wish to manually interact with the mechanisms, [Section 7](#) discusses the security aspects of the design, and [Section 8](#) explains the necessary IANA actions.

2. Requirements language

In this document, the key words "MAY", "MUST", "MUST NOT", "OPTIONAL", "RECOMMENDED", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [[RFC2119](#)].

[3.](#) Role of Prefix Assignment

Given a prefix shorter than /64 for the entire home network, this prefix needs to be subdivided so that every subnet is given its own /64 prefix. In many cases there will be just one subnet, the internal network interface attached to the router. But it is also common to have two or more internal network interfaces with intentionally separate networks. For instance, "private" and "guest" SSIDs are automatically configured in many current home network routers. When all the network interfaces that require a prefix are attached to the same router, the prefix assignment problem is simple,

and procedures outlined in [Section 4](#) can be employed.

In a more complex setting there are multiple routers in the internal network. There are various possible reasons why this might be necessary [[I-D.chown-homenet-arch](#)]. For instance, networks that cannot be bridged together should be routed, speed differences between wired and wireless interfaces make the use of the same broadcast domain undesirable, or simply that router devices keep being added. In any case, it then becomes necessary to assign prefixes across the entire network, and this assignment can no longer be done on a local basis as proposed in [Section 4](#). A distributed mechanism and a protocol is required.

The key requirements for this distributed mechanism are as follows.

- o The short prefix assigned to the home gateway router must be used to assign /64 prefixes on each subnet that requires one.
- o The assignment mechanism should provide reasonable efficiency. As a practical benchmark, some ISPs may employ /60 assignments to individual subscribers. As a result, the assignment mechanism should avoid wasting too many prefixes so that this set of 16 /64 prefixes does not run out in the foreseeable future for commonly occurring network configurations.
- o In particular, the assignment of multiple prefixes to the same network from the same top-level prefix must be avoided.

Example: When a home network consists of a home gateway router

connected to another router which in turn is connected to hosts, a minimum of two /64 prefixes are required in the internal network: one between the two routers, and another one for the host-side interface on the second router. But an ineffective assignment mechanism in the two routers might have both of them asking for an assignment for this shared interface.

- o The assignments must be stable across reboots, power cycling, router software updates, and preferably, should be stable across simple network changes such as adding a new device on the stub network segments. However, stability across more complex types of network reconfiguration events is not a requirement.

In an even more complex setting there may be multiple home gateway routers and multiple connections to ISP(s). These cases are analogous to the case of a single gateway router. Each gateway will simply distribute the prefix it has, and participating routers throughout the network may assign themselves prefixes from several

gateways.

Similarly, it is also possible that it is necessary to assign both a global prefix delegated from the ISP and a local, Unique Local Address (ULA) prefix [[RFC4193](#)]. The mechanisms in this memo are applicable to both types of prefixes. For ULA-based prefixes, it is necessary to elect one or more router as the generator of such prefixes, and have it perform the generation and employ the prefixes for local interfaces and the entire router network. The generation of ULAs in this manner -- and indeed even the question of whether ULAs are needed -- is outside the scope of this memo, however. We only note that if ULA prefixes are generated, then the mechanisms in this memo can be used to subdivide that prefix for the rest of the network.

[4.](#) Router Behavior

This section describes how a router assigns prefixes to its directly connected interfaces. We assume that the router has prefix(es) that it can use for this allocation. These prefix(es) can be manually configured, acquired through DHCPv6 PD from the ISP, or learned

through the distributed prefix assignment protocols described in [Section 5](#). Each such prefix is called a usable prefix. Parts of the usable prefix may already be assigned for some purpose; a coordinated allocation from the prefix is necessary before it can actually be assigned to an interface.

Even if the assignment process is local, it still needs to follow the requirements from [Section 3](#). This is achieved through the following rules:

- o The router **MUST** maintain a list of assigned prefixes on a per-interface basis. The contents of this list consists of prefixes that the router itself has assigned to the interface, as well as prefixes assigned to the interface by other routers. The latter are learned through the mechanisms described in [Section 5](#), when they are used.
- o Whenever the router finds a combination of an interface and usable prefix that is not used on the interface, it **SHOULD** make a new assignment. That is, the router checks to see if there exists an interface and usable prefix such that there are no assigned prefixes within that interface that are more specific than the usable prefix. In this situation the router makes an allocation from the usable prefix (if possible) and adds the allocation to the list of assigned prefixes on that interface.

- o An allocation from a usable prefix **MUST** check for other allocations from the same usable prefix. Allocations are made for individual /64 prefixes. The choice of a /64 among multiple free ones **MUST** be made randomly or based on an algorithm that takes unique hardware characteristics of the router and the interface into account. This helps avoid collisions when simultaneous allocations are made within a network.
- o In order to provide a stable assignment, the router **MUST** store assignments affecting directly connected interfaces in non-volatile memory and attempt to re-use them in the future when possible. At least the 5 most recent assignments **SHOULD** be stored. Note that this applies to both its own assignments as well as assignments made by others. This ensures that the same prefix assignments are made regardless of the order that different

devices are brought up. To avoid attacks on flash memory write cycles, assignments made by others SHOULD be recorded only after 10 minutes have passed and the assignment is still valid.

- o Re-using a memorized assignment is possible when there exists a usable prefix that is less specific than the prefix in the assignment (or it is the prefix itself in the assignment), and the prefix in the assignment can be allocated for that purpose.

Once the router has assigned a prefix to an interface, it MUST act as a router as defined in [\[RFC4861\]](#) and advertise the prefix in Router Advertisements. The lifetime of the prefix SHOULD be advertised as a reasonably long period, at least 48 hours or the lifetime of the assigned prefixes, whichever is smaller. To support a variety of IPv6-only hosts in these networks, the router needs to ensure that sufficient DNS discovery mechanisms are enabled. It is RECOMMENDED that both stateless DHCPv6 [\[RFC3736\]](#) and Router Advertisement options [\[RFC6106\]](#) are supported and turned on by default. This requires, however, that a working DNS server is known and addressable via IPv6. The mechanism in [\[RFC3736\]](#) and [\[RFC3646\]](#) can be used for this.

5. Prefix Assignment in OSPFv3

This section describes how prefix assignment in a home network can be performed in a distributed manner via OSPFv3. It is expected that the router already support the auto-configuration extensions defined in [\[I-D.acee-ospf-ospfv3-autoconfig\]](#).

An overview to OSPFv3-based prefix assignment is as follows. OSPFv3 routers that are capable of auto-configuration advertise OSPFv3 Auto-Configuration (AC) LSA [\[I-D.acee-ospf-ospfv3-autoconfig\]](#) with suitable TLVs. For prefix assignment, two TLVs are used. The Usable

Prefix TLV ([Section 5.1](#)) advertises a usable prefix, usually the prefix that has been delegated to the home gateway router from the ISP through DHCPv6 PD. These usable prefixes are necessary for running the algorithm in [Section 4](#) for determining whether prefix assignments can and should be made.

The Assigned Prefix TLV ([Section 5.2](#)) is used to communicate assignments that routers make out of the usable prefixes.

An assignment can be made when the algorithm in [Section 4](#) indicates that it can be made and no other router has claimed the same assignment. The router emits an OSPFv3 advertisement with Assigned Prefix TLV included to let other devices know that the prefix is now in use. Unfortunately, collisions are still possible, when the algorithms on different routers happen to choose the same free /64 prefix or when more /64 prefixes are needed than there are available. This situation is detected through an advertisement where a different router claims the allocation of the same prefix. In this situation the router with numerically lower OSPFv3 Router ID has to select another prefix. See also [[I-D.acee-ospf-ospfv3-autoconfig](#)] [Section 5.2](#).

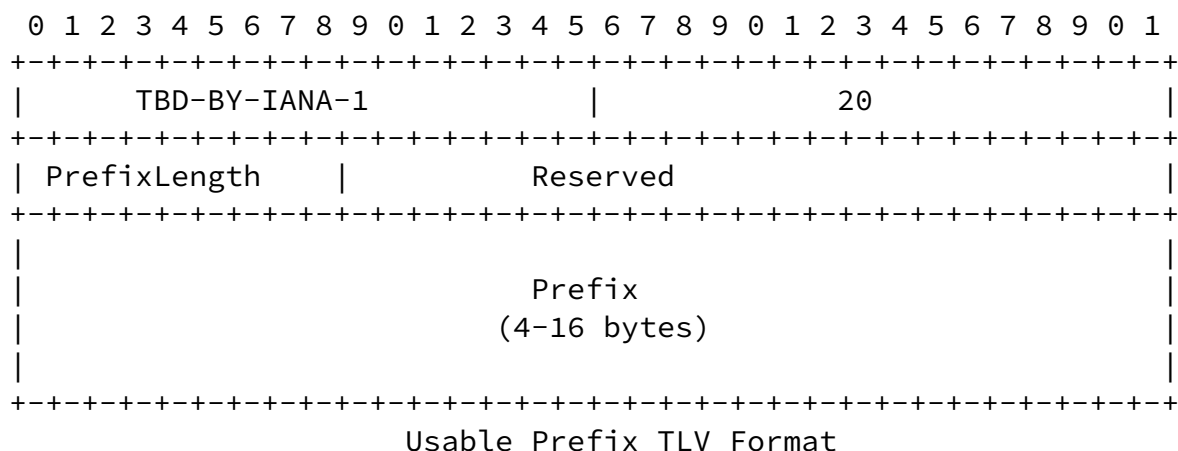
[5.1](#). Usable Prefix TLV

The Usable Prefix TLV is defined for the OSPFv3 Auto-Configuration (AC) LSA [[I-D.acee-ospf-ospfv3-autoconfig](#)]. It will have type TBD-BY-IANA-1 and MUST be advertised in the LSID OSPFv3 AC LSA with an LSID of 0. It MAY occur once or multiple times and the information from all TLV instances is retained. The length of the TLV is variable.

The contents of the TLV include a usable prefix (Prefix) and prefix length (PrefixLength). PrefixLength is the length in bits of the prefix and is an 8-bit field. The PrefixLength MUST be greater than or equal to 8 and less than or equal to 64. The prefix describes an allocation of a global or ULA prefix for the entire auto-configured home network. The Prefix is an encoding of the prefix itself as an even multiple of 32-bit words, padding with zero bits as necessary. This encoding consumes $(\text{PrefixLength} + 31) / 32$ 32-bit words and is consistent with [[RFC5340](#)]. It MUST NOT be directly assigned to any interface before following through the procedures defined above.

This TLV SHOULD be emitted by every home gateway router that has either a manual or DHCPv6 PD based prefix that is shorter than /64.

This TLV MUST appear inside an OSPFv3 Router Auto-Configuration LSA, and only in combination with the Router-Hardware-Fingerprint TLV [[I-D.acee-ospf-ospfv3-autoconfig](#)] [Section 5.2.2](#) in the same LSA.




5.2. Assigned Prefix TLV

The Assigned Prefix TLV is defined for the OSPFv3 Auto-Configuration (AC) LSA [[I-D.acee-ospf-ospfv3-autoconfig](#)]. It will have type TBD-BY-IANA-2 and MUST be advertised in the LSID OSPFv3 AC LSA with an LSID of 0. It MAY occur once or multiple times and the information from all TLV instances is retained. The length of the TLV is variable.

The contents of the TLV include an Interface ID, assigned prefix (Prefix), and prefix length (PrefixLength). The Interface ID is the same OSPFv3 Interface ID that is described in [section 4.2.1](#) or [[RFC5340](#)]. PrefixLength is the length in bits of the prefix and is an 8-bit field. The PrefixLength value MUST be 64 in this version of the specification. The prefix describes an assignment of a global or ULA prefix for a directly connected interface in the advertising router. The Prefix is an encoding of the prefix itself as an even multiple of 32-bit words, padding with zero bits as necessary. This encoding consumes $(\text{PrefixLength} + 31) / 32$ 32-bit words and is consistent with xref target="[RFC5340](#)"/>.

This TLV MUST be emitted by every home router that has made assignment from a usable prefix per [Section 4](#).

This TLV MUST appear inside an OSPFv3 Router Auto-Configuration LSA, and only in combination with the Router-Hardware-Fingerprint TLV [[I-D.acee-ospf-ospfv3-autoconfig](#)] [Section 5.2.2](#) in the same LSA.



Assigned Prefix TLV Format

5.3. OSPFv3 Prefix Assignment

OSPFv3 Routers supporting the mechanisms in the memo will learn or assign a global /64 IPv6 prefix for each IPv6 interface. Since the mechanisms described herein are based on OSPFv3, Router ID assignment as described in [[I-D.acee-ospf-ospfv3-autoconfig](#)] MUST have completed successfully.

When an OSPFv3 Router receives a global prefix through DHCPv6 prefix delegation, manual configuration, or other means, it will advertise this prefix by including the Usable Prefix TLV in its OSPFv3 AC LSA. This will trigger prefix assignment for auto-configured OSPFv3 routers within the routing domain including the originating OSPFv3 router.

When an OSPFv3 Router receives an AC LSA containing a Usable Prefix TLV, it will determine whether or not a new prefix needs to be assigned for each of its attached IPv6 interfaces. For the purposes of this discussion, the received prefix will be referred to as the Current Usable Prefix. The following steps will be performed for each IPv6 interface:

1. The OSPFv3 Router will determine whether there are any other OSPFv3 Routers connected to the same link by examining its list of neighbors.
2. If no OSPFv3 neighbors have been discovered, the router will wait TBD seconds before allocating a unique /64 IPv6 prefix for the

link as described in step 5.

3. If OSPFv3 neighbors are present on the link, the router needs to determine whether any of them have already assigned an IPv6 prefix. This is done by examining the AC LSAs for neighbors on the link and looking for any that include an Assigned Prefix TLV with the same OSPFv3 Interface ID as the neighbor. If one is found and is it a subnet of the IPv6 prefix advertised in the Usable Prefix TLV, this global IPv6 prefix has been already been assigned to the link. If more than one neighbor's Assigned Prefix TLV is found with an IPv6 prefix matching the criteria above, the Assigned Prefix advertised by the OSPFv3 router with the numerically highest OSPFv3 Router ID takes precedence.
4. If there are OSPFv3 neighbors on the link but no IPv6 Prefix is found, the task of prefix allocation is delegated to the OSPFv3 Router with the numerically highest OSPFv3 Router ID. Note that this is different from OSPFv3 Designated Router (DR) election, as described in [\[RFC5340\]](#), in that the router priority is not taken into consideration and that the election will work for networks types where no DR is elected, e.g., point-to-point links.
5. If it is determined that the OSPFv3 Router is responsible for prefix assignment on the link, it will:
 - * Examine all the AC LSA including Assigned Prefix TLVs that are subnets of the Current Usable Prefix to determine which /64s prefixes are already assigned.
 - * Examine former prefix assignments stored in non-volatile storage for interface. Starting with the most recent assignment, if the prefix is both a subnet of the Current Usable Prefix and is currently unassigned, reuse the assignment for the interface.
 - * If no unused former prefix allocation is found, allocate a new one from the subnets of the Current Usable Prefix which are unallocated.
 - * Once a global IPv6 prefix is assigned, a new instance of the

AC LSA will be re-originated including the Assigned Prefix TLV.

- * In the rare event that no global /64 IPv6 prefixes are available within the Current Usable Prefix, no IPv6 prefix is assigned and an error condition must be raised.

There are two types of conflicts that may be detected:

1. Two or more OSPFv3 routers have assigned the same IPv6 prefix for different networks.
2. Two of more OSPFv3 routers have assigned different IPv6 prefixes for the same network.

In the case of the former, the OSPFv3 Router with the numerically lower OSPFv3 Router ID must select a new prefix and advertise a new instance of its AC LSA with an updated Assign Prefix TLV for the link. In the latter case, the OSPFv3 Router with the numerically lower OSPFv3 Router ID should accept the global IPv6 prefix from the neighbor with the highest OSPFv3 Router ID and originate a new AC LSA excluding the Assigned Prefix TLV for the link.

6. Manageability Considerations

Advanced users may wish to manage their networks without automation, and there may also be situations where manual intervention may be needed. For these purposes there MUST be a configuration mechanism that allows users to turn off the automatic prefix assignment on a given interface. This setting can be a part of disabling the entire routing auto-configuration [[I-D.acee-ospf-ospfv3-autoconfig](#)].

In addition, there SHOULD be a configuration mechanism that allows users to specify the prefix that they would like the router to request for a given interface. This can be useful, for instance, when a router is replaced and there is a desire for the new router to be configured to ask for the same prefix as the old one, in order to avoid renumbering other devices on this network.

Finally, there SHOULD be mechanisms to display what prefixes the

router has been assigned, and where they came from (manual configuration, DHCPv6 PD, OSPFv3).

7. Security Considerations

Security can be always added later.

8. IANA Considerations

This memo makes two allocations out of the OSPFv3 Auto- Configuration (AC) LSA TLV namespace [[I-D.acee-ospf-ospfv3-autoconfig](#)]:

- o The Usable Prefix TLV in [Section 5.1](#) takes the value TBD-BY-IANA-1 (suggested value is 2).

Arkko & Lindem

Expires April 26, 2012

[Page 11]

Internet-Draft

Homenet Prefixes

October 2011

- o The Assigned Prefix TLV in [Section 5.2](#) takes the value TBD-BY-IANA-2 (suggested value is 3).

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3646](#), December 2003.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", [RFC 3736](#), April 2004.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.

- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), July 2008.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", [RFC 6106](#), November 2010.
- [I-D.acee-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration", [draft-acee-ospf-ospfv3-autoconfig-00](#) (work in progress), October 2011.

[9.2.](#) Informative References

- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.
- [I-D.chown-homenet-arch]
Arkko, J., Chown, T., Weil, J., and O. Troan, "Home Networking Architecture for IPv6", [draft-chown-homenet-arch-00](#) (work in progress), September 2011.

[Appendix A.](#) Acknowledgments

The authors would like to thank to Tim Chown, Fred Baker, Mark Townsley, Lorenzo Colitti, Ole Troan, Ray Bellis, Wassim Haddad, Joel Halpern, Samita Chakrabarti, Michael Richardson, and Ralph Droms for interesting discussions in this problem space.

Authors' Addresses

Jari Arkko
Ericsson
Jorvas 02420
Finland

Email: jari.arkko@piuha.net

Acee Lindem
Ericsson
Cary, NC 27519
USA

Email: acee.lindem@ericsson.com