

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 26, 2013

J. Arkko
A. Lindem
Ericsson
B. Paterson
Cisco Systems
October 23, 2012

Prefix Assignment in a Home Network
draft-arkko-homenet-prefix-assignment-03

Abstract

This memo describes a prefix assignment mechanism for home networks. It is expected that home gateway routers are allocated an IPv6 prefix through DHCPv6 Prefix Delegation (PD) or that a prefix is made available through other means. This prefix needs to be divided among the multiple subnets in a home network. This memo describes a mechanism for such division, or assignment, via OSPFv3. This is an alternative design to also using DHCPv6 PD for the 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, 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	4
3. Role of Prefix Assignment	4
4. Router Behavior	5
4.1. Sending Router Advertisements	7
4.2. DNS Discovery	7
5. Design Choices	8
5.1. DNS Discovery	8
5.2. Prefix Assignment	8
6. Prefix Assignment in OSPFv3	9
6.1. Aggregated Prefix TLV	10
6.2. Assigned Prefix TLV	11
6.3. OSPFv3 Prefix Assignment	12
6.3.1. Making a New Assignment	15
6.3.2. Checking for Conflicts Across the Entire Network	15
6.3.3. Deprecating an Assigned Prefix	16
6.3.4. Verifying and Making a Local Assignment	16
7. ULA Generation	16
8. Hysteresis	18
9. Manageability Considerations	19
10. Security Considerations	19
11. IANA Considerations	19
12. Timer Constants	19
13. References	20
13.1. Normative References	20
13.2. Informative References	20
Appendix A. Changes in Version -02	21
Appendix B. Changes in Version -03	21
Appendix C. Acknowledgments	21
Authors' Addresses	21

1. Introduction

This memo describes a prefix assignment mechanism for home networks. It is expected that home gateway routers are allocated an IPv6 prefix through DHCPv6 Prefix Delegation (PD) [RFC3633], or that a prefix is made available by some other means. Manual configuration may be needed in some networks, for instance when the ISP does not support DHCPv6-based prefix delegation. In other cases, such as networks that have do not yet have an Internet connection, Unique Local Address (ULA) [RFC4193] prefixes can be automatically generated. For the purposes of this document, we refer to the prefix reserved for a home network as a prefix allocation.

A prefix allocation needs to be divided among the multiple subnets in a home network. For the purposes of this document, we refer to this process as prefix assignment. This memo describes a mechanism for prefix assignment via OSPFv3 [RFC5340].

The OSPv3-based mechanism is an alternative design to also using DHCPv6 PD 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 use of address space than hierarchical assignment through DHCPv PD. This may be important for home networks that only get a /60 prefix 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 and Section 6 describe how the assignment can be performed in a distributed manner via OSPFv3 in the entire home network. Finally, Section 7 specifies the procedures for automatic generation of ULA prefixes, Section 8 explains the hysteresis principles applied to prefix assignment and de-assignment, Section 9 explains what administrative interfaces are useful for advanced users that wish to manually interact with the mechanisms, Section 10 discusses the security aspects of the design, Section 11 explains the necessary IANA actions, and Section 12 defines the necessary timer constants.

An analysis of a mechanism reminiscent of the one described in this specification has been published in the SIGCOMM IPv6 Workshop [SIGCOMM.IPV6]. Further analysis is encouraged.

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.ietf-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 are required.

The key requirements for this distributed mechanism are as follows.

- o A prefix allocated to a home gateway router within the home network is used to assign /64 prefixes on each subnet that requires one.

Note that several methods may be used to allocate such an aggregated prefix.

- o The assignment mechanism should provide reasonable efficiency. As a practical benchmark, some ISPs may employ /60 allocations to individual subscribers. As a result, the assignment mechanism should avoid wasting too many prefixes so that this set of 16 /64 prefixes is not exhausted 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 separate assignments 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. Simple network changes are in this case defined as those that could be resolved through either deletion or addition of a prefix assignment. For instance, the addition of a new router without changing connections between existing routers requires just the assignment of new prefixes for the new networks that the router introduces. There are no stability requirements across more complex types of network reconfiguration events. For instance, if a network is separated into two networks connected by a newly inserted router, this may lead to renumbering all networks within the home.

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. Multiple assignments can be made for the same interface. For example, this can be useful in a multi-homing setting.

Similarly, it is also possible that it is necessary to assign either a global prefix delegated from the ISP or a local, Unique Local Address (ULA) prefix [RFC4193]. The mechanisms in this memo are applicable to both types of prefixes. The details of the generation of ULA-based prefixes is covered in Section 7.

The mechanisms in this memo can also be used in standalone or ad hoc networks where no global prefixes or Internet connectivity are available, by distributing ULA prefixes within 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

allocation(s) that it can use for this assignment. Each such prefix allocation is called an aggregated prefix. Parts of the aggregated prefix may already be assigned for some purpose; a coordinated assignment 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 6, when they are used. Each prefix is associated with the Router ID of the router that assigned it.
- o Whenever the router finds a combination of an interface and aggregated prefix that is not used on the interface, it SHOULD make a new prefix assignment. That is, the router checks to see if an interface and aggregated prefix exists such that there are no assigned prefixes within that interface that are more specific than the aggregated prefix. In this situation the router makes an allocation from the aggregated prefix (if possible) and adds the assignment to the list of assigned prefixes on that interface.

Note: The above implies that when there are multiple aggregated prefixes, each network will be assigned multiple prefixes.

- o An assignment from an aggregated prefix MUST be checked against possible other assignments from the same aggregated prefix on the same link by neighboring routers, to avoid unnecessary assignments. Assignments MUST also be examined against all existing assignments from the same aggregated prefix across the network, to avoid collisions. Assignments are made for individual /64 prefixes. The choice of a /64 prefix 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 assignments are made within a network.
- o In order to provide a stable assignment, the router MUST store assignments affecting directly connected interfaces and automatically generated ULA prefixes 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 a aggregated prefix exists that is less specific than the prefix in the assignment (or it is the prefix itself in the assignment), and the prefix is currently unassigned.

4.1. Sending Router Advertisements

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.

4.2. DNS Discovery

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 in routers.

The above requires, however, that a working DNS server is known and addressable via IPv6. The mechanism in [RFC3736] and [RFC3646] can be used for this. It is RECOMMENDED that each router attempts to discover an existing DNS server. Typically, such a server will be provided by an ISP. However, in some cases no such server can be found. For instance, an ISP may provide only IPv4 DNS server addresses, or a router deep within the home network is unaware of the IPv6 DNS servers that a home gateway router has discovered. In these situations it is RECOMMENDED that each router turns on a local DNS relay that fetches information from the IPv4 Internet (if a working IPv4 DNS server is available) or a full DNS server that fetches information from the DNS root.

As a result of these recommendations, as long as there is reachability to at least the Internet, every router within the home network will either know the IPv6 address of a DNS server or it itself runs a server that can fetch information from the Internet. As a result, the router can provide information about the server address to hosts in directly connected networks.

5. Design Choices

5.1. DNS Discovery

The DNS discovery recommendations in Section 4.2 ensure that an IPv6-only home network can resolve names. However, these recommendations are suboptimal in the sense that different routers in the home may provide different DNS servers, or multiple local DNS servers have to be run where it would have been possible to point to one, or even point to the one provided by the ISP. However, better coordination for the DNS server selection would require some form of additional communication between the routers in the home network. The authors solicit opinions from the Working Group on whether this is something that should be specified. However, the current design is easy to deploy even when not all routers within the network support Homenet specifications yet; the mechanism provides an incremental improvement to IPv6 DNS reachability even when the first Homenet router is deployed.

5.2. Prefix Assignment

The OSPFv3-based prefix assignment protocol needs to detect two types of conflicts:

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

Several design decisions were needed to construct the protocol:

1. How to determine the winner in case of a conflict?

The algorithm in Section 6 ensures that the OSPFv3 Router with the numerically lower OSPFv3 Router ID removes its assignment and schedules an advertisement of LSAs that no longer describe such an assignment. That is, the router with the highest Router ID wins in a conflict situation.

2. How to ensure that a network-wide conflict can be detected?

We chose to define new LSA extensions -- TLVs within the new Autoconfiguration LSA -- that are flooded throughout the network. Another possible design would have been to re-use existing OSPFv3 LSAs, and by assuming that if a router advertises a prefix then it has made an assignment. The advantage of the design that we chose is that we get to specify what information is needed in the

new TLVs. This is particularly important, as not all existing OSPFv3 LSAs are extensible. A downside is that assignments will not be visible, unless the router using an assignment implements this specification and advertises the new LSAs. Had we reused existing LSAs, a manual assignment for a legacy router could have been handled, as the legacy router would have advertised the prefix assigned to it.

3. How to ensure that both local and network-wide conflicts can be detected?

We chose to employ the same new Autoconfiguration LSA TLVs for this purpose, and correlate neighbors through the Router IDs and Interface IDs that they advertise in these TLVs. The OSPFv3 Router with a numerically lower OSPFv3 Router ID should accept the global IPv6 prefix from the neighbor with the highest OSPFv3 Router ID.

6. 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.ietf-ospf-ospfv3-autoconfig].

An overview of OSPFv3-based prefix assignment is as follows. OSPFv3 routers that are capable of auto-configuration advertise an OSPFv3 Auto-Configuration (AC) LSA [I-D.ietf-ospf-ospfv3-autoconfig] with suitable TLVs. For prefix assignment, two TLVs are used. The Aggregated Prefix TLV (Section 6.1) advertises an aggregated prefix, usually the prefix that has been delegated to the home gateway router from the ISP through DHCPv6 PD. These aggregated 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 6.2) is used to communicate assignments that routers make out of the aggregated 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 makes an OSPFv3 advertisement with the 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 are available. This situation is detected through an advertisement where a different router claims the assignment of the same prefix. In this

situation the router with the numerically lower OSPFv3 Router ID has to select another prefix and immediately withdraw any assignments and advertisements that may have been advertised in OSPFv3. See also Section 5.2 in [I-D.ietf-ospf-ospfv3-autoconfig].

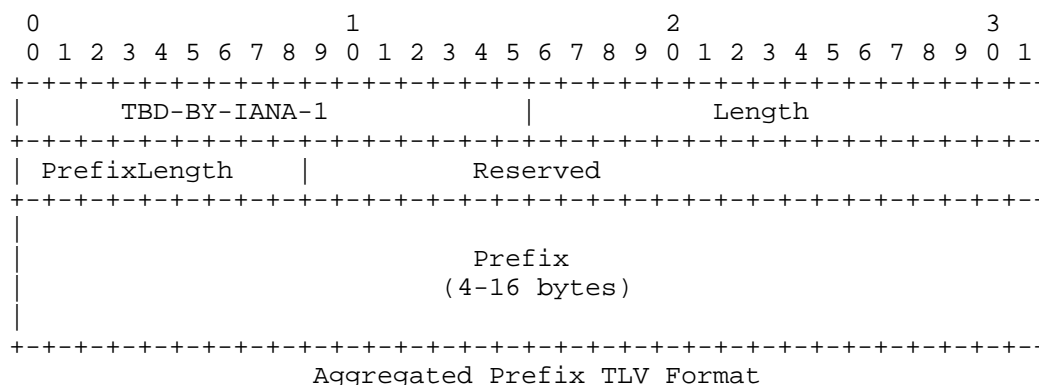
6.1. Aggregated Prefix TLV

The Aggregated Prefix TLV is defined for the OSPFv3 Auto-Configuration (AC) LSA [I-D.ietf-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 an aggregated 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 the procedures defined in this memo.

This TLV SHOULD be advertised by every home gateway router that has either a manual, DHCPv6 PD-based, or generated ULA 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.ietf-ospf-ospfv3-autoconfig] Section 5.2.2 in the same LSA.



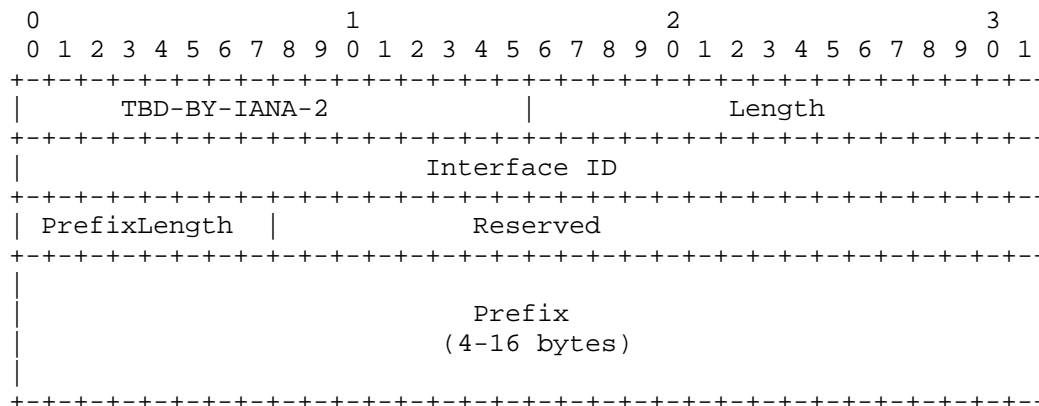
6.2. Assigned Prefix TLV

The Assigned Prefix TLV is defined for the OSPFv3 Auto-Configuration (AC) LSA [I-D.ietf-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 [RFC5340].

This TLV MUST be advertised by the router that has made assignment from an aggregated 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.ietf-ospf-ospfv3-autoconfig] Section 5.2.2 in the same LSA.



Assigned Prefix TLV Format

6.3. 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.ietf-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 SHOULD advertise this prefix by including the Aggregated 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.

Discussion: Note that while having multiple routers advertise the same aggregated address space (or address space that covers another router's aggregated address space) is a configuration error, it should not result in any adverse effects, as long as assignments from such space are still checked for collisions against all other assignments from the same address space.

When an OSPFv3 Router detects a change in the set of AC LSAs in its LSA database, it will run the prefix assignment algorithm. The purpose of this algorithm is to determine, for each Aggregated Prefix in the database, whether or not a new prefix needs to be assigned for each of its attached IPv6 interfaces and whether or not existing assignments need to be deprecated. The algorithm also detects and removes assignments for which there is no longer a corresponding Aggregated Prefix. Before the algorithm is run, all existing assignments in assigned prefix lists for directly connected interfaces must be marked as "invalid" and will be deleted at the end of the algorithm if they are still in this state. An assigned prefix is considered to be "valid" if all the following conditions are met:

- o A containing Aggregated Prefix TLV exists in reachable AC LSA(s).
- o An Assigned Prefix TLV that matches this assignment exactly (same prefix, same router and interface ID associated with the assignment) exist in reachable AC LSA(s).
- o Any router advertising an assignment for the same link and Aggregated Prefix has a lower Router ID than the source of this assignment.
- o If this router is the source of the assignment, any router in the network that has assigned the same prefix on a different link has a lower Router ID than this router.

Note that this definition of a "valid assignment" depends on the router running the algorithm: in particular, a router is not expected to detect prefix collisions or duplicate prefix assignments that do not concern assignments for which it is the responsible router. It is the role of the responsible router to detect these cases and update its AC LSAs accordingly. A router is, however, expected to react to these updates from other routers which translate into additions or removals of Aggregated Prefix or Assigned Prefix TLVs.

The router is expected to have made a snapshot of the LSA database before running this algorithm. The prefix assignment algorithm consists of the following steps run once per combination of Aggregated Prefix in the LSA database and directly connected OSPFv3 interface. For the purposes of this discussion, the Aggregated Prefix will be referred to as the Current Aggregated Prefix, and the interface will be referred to as the Current Interface. The following steps will be performed for each tuple (Aggregated Prefix, OSPFv3 interface):

1. The OSPFv3 Router will search all AC LSAs for an Aggregated Prefix TLV describing a prefix which contains but is not equal to the Current Aggregated Prefix. If such a prefix is found, the algorithm is skipped for the Current Aggregated Prefix as it either has or will be run for the shorter prefix.
2. The OSPFv3 router will examine its list of neighbors to find all neighbors in state greater than Init (these neighbors will be referred to as active neighbors).
3. The following three steps will serve to determine whether an assignment needs to be made on the link:

i.

The OSPFv3 router will determine whether or not it has the highest Router ID of all active OSPFv3 routers on the link.

ii.

If OSPFv3 active neighbors are present on the link, the router will determine whether any of them have already assigned an IPv6 prefix. This is done by examining the AC LSAs of all the active neighbors on the link and looking for any that include an Assigned Prefix TLV with the same OSPFv3 Router ID and Interface ID as the neighbor has. If one is found and it is a subnet of the IPv6 prefix advertised in the Aggregated Prefix TLV, the router stores this prefix and the Router ID of the router advertising it for reference in the next step. If

several such prefixes are found, only the prefix and Router ID with the numerically highest Router ID are stored.

iii.

The router will compare its Router ID with the highest Router ID among neighbors which have made an assignment on the link. If it is higher (or if no assignments have been made by any neighbors), it will determine whether or not it is already the source of an assignment for the Current Interface from the Current Aggregated Prefix.

4. There are four possibilities at this stage:

- * The router has already made an assignment on the link and has a higher Router ID than all eventual neighbors which have also made an assignment. In this case, the router's existing assignment takes precedence over all other eventual existing assignments on the link, but the router must determine whether its assignment is still valid throughout the whole network. This is described in Section 6.3.2.
- * An assignment has been made by a neighbor on the link, and the router either has not made an assignment on the link, or has a lower Router ID than the neighbor. In this case, the neighbor's assignment takes precedence over all eventual existing assignments on the link (including assignments made by the router), and the router must update the assigned prefix list of the Current Interface as well as check assignments on other interfaces for potential collisions. This is described in Section 6.3.4.
- * No assignment has been made by anyone on the link, and the router has the highest Router ID on the link. In this case, it must make an assignment from the Current Aggregated Prefix. This is described in Section 6.3.1.
- * No assignment has been made by anyone on the link, and the router does not have the highest Router ID on the link. In this case, the algorithm exits as the router is not responsible for prefix assignment on the link.

Once the algorithm has been run for each Aggregated Prefix and each interface, the router must delete all assignments that are not marked as valid on all assigned prefix lists and deprecate the corresponding addresses. If this leads to deleting an assignment that this router was responsible for, or if AC LSA origination was scheduled during the algorithm, it must originate a new AC LSA advertising the

changes. The router MUST also deprecate deleted prefixes as specified in Section 6.3.3.

6.3.1. Making a New Assignment

This procedure is executed when no assignment exists on the link and the router is responsible for making an assignment. The router MUST:

1. Examine all the AC LSAs not advertised by this router that include Assigned Prefix TLVs that are subnets of the Current Aggregated Prefix, as well as all assignments made by this router, to determine which prefixes are already assigned.
2. Examine former prefix assignments stored in non-volatile storage for the interface. Starting with the most recent assignment, if the prefix is both a subnet of the Current Aggregated Prefix and is currently unassigned, reuse the assignment for the interface.
3. If no unused former prefix assignment is found, and an unassigned /64 subnet of the Current Aggregated Prefix exists, assign that prefix to the interface.
4. If no OSPFv3 neighbors have been discovered and previous prefix assignments exist, the router can make the assignments immediately. Otherwise, the hysteresis periods specified in Section 8 are applied before making an assignment.
5. In the event that no assignment could be made to the interface, a warning must be raised. However, the router MUST remain in a state where it continues to assign prefixes through OSPFv3, as prefixes may later become available.
6. Once a global IPv6 prefix is assigned, the router will mark it as valid and schedule re-origination of the AC LSA including the Assigned Prefix TLV once all Aggregated Prefixes and interfaces have been examined.

6.3.2. Checking for Conflicts Across the Entire Network

This procedure is executed for every assignment that the router intends to make or retain as the router responsible for an assignment.

The router MUST verify that this assignment is still valid across the whole network. This assigned prefix will be referred to as the Current Assigned Prefix. The router will search for a reachable AC LSA in the LSA database that is advertised by a router with a higher Router ID and contains an Assigned Prefix equal to the Current

Assigned Prefix. If such an LSA is found, it needs to be deprecated as described in Section 6.3.3. Otherwise, the router will mark its assignment as valid.

6.3.3. Deprecating an Assigned Prefix

This procedure is executed when the router's earlier assignment of a prefix can no longer be used. The following steps MUST be followed:

1. If the the prefix was in an interface's assigned prefix list, it is removed.
2. If this router was the source of the prefix assignment, schedule re-origination of the modified AC LSA once the algorithm has finished.
3. The router MUST also deprecate the prefix, if it had been advertised in Router Advertisements on an interface. The prefix is deprecated by sending Router Advertisements with the lifetime set to 0 [RFC4861] for the prefix in question.

6.3.4. Verifying and Making a Local Assignment

This procedure is executed when an assignment by a neighbor already exists, and takes precedence over all other assignments on the link. The router must check whether or not it is already aware of this assignment. It will search for the assigned prefix matching the neighbor's assignment and Router ID in the Current Interface's assigned prefix list. If it is already present, the router will mark it as valid. Otherwise, the router will check that no assignment on any directly connected interface collides with the neighbor's assignment. If a collision is found and the colliding prefix takes priority over the neighbor's assignment (higher Router ID), the router will silently ignore the neighbor's assignment. If a collision is found but the neighbor's assignment takes priority, the old assignment is removed as described in Section 6.3.3. If the neighbor's assignment takes priority, or if no collision was found, the router will provision the interface with the prefix, add it to the list of assigned prefixes using the neighbor's Router ID and mark it as valid.

7. ULA Generation

For ULA-based prefixes, it is necessary to elect a router as the generator of such prefixes, have it perform the generation, and then employ the prefixes for local interfaces and the entire router network. This section specifies these procedures, and recommends the

generation of ULAs when no connectivity can be established otherwise. However, the use of ULAs in parallel with global IPv6 prefixes is outside the scope of this memo. The mechanisms in this memo could be used for that as well.

When an OSPFv3 Router detects a change in the set of AC LSAs in its LSA database, it will run the ULA generation algorithm. The purpose of this algorithm is to determine whether a new ULA prefix needs to be generated. There is no need for this router to generate a new ULA prefix when any of the following conditions are met:

i.

An Aggregated Prefix TLV exists in an AC LSA advertised by a reachable router in the LSA database, with either global or ULA address space.

ii.

A reachable router in the OSPFv3 topology with a higher Router ID than this OSPFv3 router exists.

iii.

This router has assignments from either IPv4 or IPv6 global address space on any interface, or there is connectivity to the global Internet.

Discussion: This rule is necessary in order to prevent autoconfiguration-capable routers from unnecessarily creating ULA address space in networks where autoconfiguration is not in use. Similarly, from an IPv6 "happy eyeballs" perspective it is desirable to not create local islands of IPv6 connectivity when there is IPv4 connectivity (even through a NAT).

If none of the above conditions are met after applying the hysteresis principles from Section 8, the router SHOULD perform the following actions:

1. Generate a new 48-bit ULA prefix as specified in [RFC4193], Section 3.2.
2. Record the new prefix in stable storage, per rules in Section 4.
3. Advertise the new prefix allocation in OSPFv3 as specified in Section 6.3.

4. Assign /64 prefixes from the new prefix for its own use, as a part of the general algorithm for making prefix assignments (also in Section 6.3).

If the router has made such an allocation, it SHOULD continue to advertise the prefix in OSPFv3 for as long as conditions i) through iii) do not apply, with the exception of the generated ULA prefix that this router is advertising.

If the router has made such an allocation, and any of the conditions become true (except for the case of the ULA prefix that the router is advertising) even after applying the hysteresis principles from Section 8, then the OSPFv3 router SHOULD withdraw the advertisement for the aggregated prefix. This is done by scheduling the re-origination of an AC LSA that does not include the Aggregated Prefix TLV with the ULA. Note that as a result of the general algorithm for making prefix assignments, any /64 prefix assignments from the ULA prefix will eventually be deprecated.

8. Hysteresis

A network may experience temporary connectivity problems, routing protocol convergence may take time, and a set of devices may be coming up at the same time due to power being turned on in a synchronous manner. Due to these reasons it is important that the prefix allocation and assignment mechanisms do not react before the situation is allowed to stabilize. To allow for this, a hysteresis principle is applied to new or withdrawn automatically generated prefixes and prefix assignments.

A new automatically generated ULA prefix SHOULD NOT be allocated before the router has waited NEW_ULA_PREFIX seconds for another prefix or reachable OSPFv3 router to appear. See Section 12 for the specific time value.

A previously automatically generated ULA prefix SHOULD NOT be taken out of use before the router has waited TERMINATE_ULA_PREFIX seconds.

A new prefix assignment within an aggregated prefix SHOULD NOT be committed before the router has waited NEW_PREFIX_ASSIGNMENT seconds for another prefix or reachable OSPFv3 router to appear. Note the exceptions to this rule in Section 6.3.1, item 4.

A previously assigned prefix SHOULD NOT be taken out of use before the router has waited TERMINATE_PREFIX_ASSIGNMENT seconds.

9. 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 allocation and assignment on a given interface. This setting can be a part of disabling the entire routing auto-configuration [I-D.ietf-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 the prefixes assigned on each interface, and where they came from (manual configuration, DHCPv6 PD, OSPFv3).

10. Security Considerations

Security can be always added later.

11. IANA Considerations

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

- o The Aggregated Prefix TLV in Section 6.1 takes the value TBD-BY- IANA-1 (suggested value is 2).
- o The Assigned Prefix TLV in Section 6.2 takes the value TBD-BY- IANA-2 (suggested value is 3).

12. Timer Constants

NEW_ULA_PREFIX	20 seconds
TERMINATE_ULA_PREFIX	120 seconds
NEW_PREFIX_ASSIGNMENT	20 seconds
TERMINATE_PREFIX_ASSIGNMENT	240 seconds

13. References

13.1. 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.ietf-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration", draft-ietf-ospf-ospfv3-autoconfig-00 (work in progress), October 2012.

13.2. 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.ietf-homenet-arch]
Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "Home Networking Architecture for IPv6", draft-ietf-homenet-arch-06 (work in progress), October 2012.
- [I-D.chelius-router-autoconf]
Chelius, G., Fleury, E., and L. Toutain, "Using OSPFv3 for IPv6 router autoconfiguration", draft-chelius-router-autoconf-00 (work in progress), June 2002.

[I-D.dimitri-zospf]

Dimitrelis, A. and A. Williams, "Autoconfiguration of routers using a link state routing protocol", draft-dimitri-zospf-00 (work in progress), October 2002.

[SIGCOMM.IPV6]

Chelius, G., Fleury, E., Sericola, B., Toutain, L., and D. Binet, "An evaluation of the NAP protocol for IPv6 router auto-configuration", ACM SIGCOMM IPv6 Workshop, Kyoto, Japan, 2007.

Appendix A. Changes in Version -02

These changes were extensive, including the definition of a new algorithm for making allocations, adding support for DNS server discovery, adding support for ULA-based address space generation, and adding specifications for a hysteresis mechanism.

Appendix B. Changes in Version -03

This version updated references to the most current ones, and changed the "usable prefix" terminology to "aggregated prefix". The requirements for turning on DNS relays or servers were also clarified.

Appendix C. Acknowledgments

The authors would like to thank to Tim Chown, Fred Baker, Mark Townsley, Lorenzo Colitti, Ole Troan, Ray Bellis, Markus Stenberg, Wassim Haddad, Joel Halpern, Samita Chakrabarti, Michael Richardson, Anders Brandt, Erik Nordmark, Laurent Toutain, and Ralph Droms for interesting discussions in this problem space. The authors would also like to point out some past work in this space, such as those in [I-D.chelius-router-autoconf] or [I-D.dimitri-zospf].

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

Benjamin Paterson
Cisco Systems
Paris
France

Email: benjamin@paterson.fr

