

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

J. Arkko
A. Lindem
Ericsson
B. Paterson
Cisco Systems
July 9, 2012

Prefix Assignment in a Home Network
draft-arkko-homenet-prefix-assignment-02

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 January 10, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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	7
6. Prefix Assignment in OSPFv3	9
6.1. Usable Prefix TLV	9
6.2. Assigned Prefix TLV	10
6.3. OSPFv3 Prefix Assignment	11
6.3.1. Making a New Assignment	14
6.3.2. Checking for Conflicts Across the Entire Network	15
6.3.3. Deprecating an Assigned Prefix	15
6.3.4. Verifying and Making a Local Assignment	16
7. ULA Generation	16
8. Hysteresis	18
9. Manageability Considerations	18
10. Security Considerations	19
11. IANA Considerations	19
12. Timer Constants	19
13. References	19
13.1. Normative References	19
13.2. Informative References	20
Appendix A. Acknowledgments	20
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.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 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 a usable 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 a usable prefix. Parts of the usable 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 usable 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 usable prefix exists 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 assignment to the list of assigned prefixes on that interface.

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

- o An assignment from a usable prefix MUST be checked against possible other assignments from the same usable prefix on the same link by neighboring routers, to avoid unnecessary assignments. Assignments MUST also be examined against all existing assignments from the same usable 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 usable 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.

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.

5. Design Choices

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.

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.acee-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.acee-ospf-ospfv3-autoconfig] with suitable TLVs. For prefix assignment, two TLVs are used. The Usable Prefix TLV (Section 6.1) 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 6.2) 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 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.acee-ospf-ospfv3-autoconfig].

6.1. Usable Prefix TLV

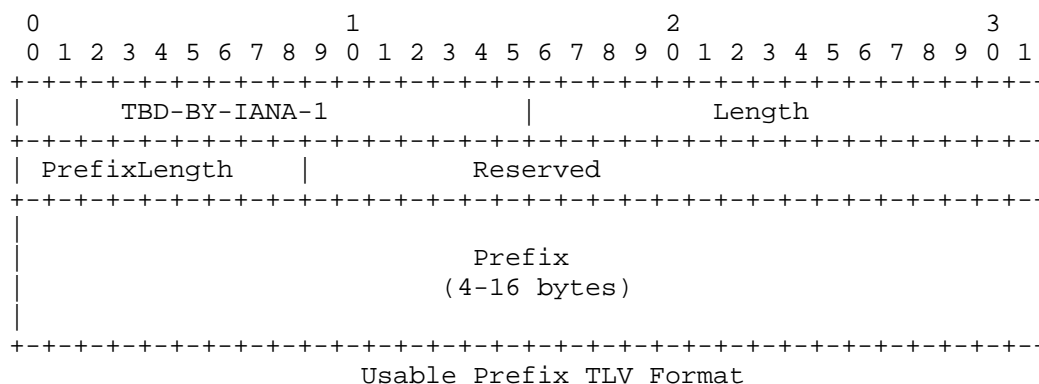
The Usable Prefix TLV is defined for the OSPFv3 Auto-Configuration (AC) LSA [I-D.acee-ospf-ospfv3-autoconfig]. It will have type TBD-BY-IANA-1 and MUST be advertised in the LSID OSPFv3 AC LSA with an LSID of 0. It MAY occur once or multiple times and the information from all TLV instances is retained. 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 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.acee-ospf-ospfv3-autoconfig] Section 5.2.2 in the same LSA.



6.2. Assigned Prefix TLV

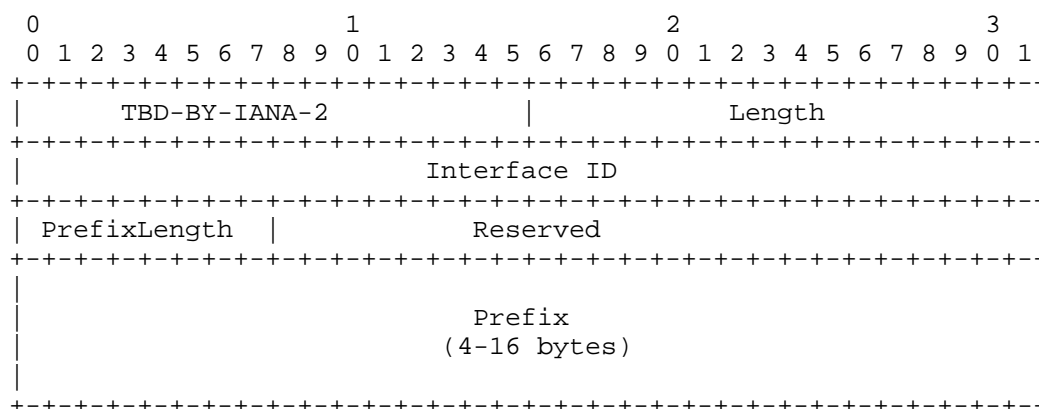
The Assigned Prefix TLV is defined for the OSPFv3 Auto-Configuration (AC) LSA [I-D.acee-ospf-ospfv3-autoconfig]. It will have type TBD-BY-IANA-2 and MUST be advertised in the LSID OSPFv3 AC LSA with an LSID of 0. It MAY occur once or multiple times and the information from all TLV instances is retained. 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 [RFC5340].

This TLV MUST be advertised by the 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

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.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 SHOULD 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.

Discussion: Note that while having multiple routers advertise the same usable address space (or address space that covers another router's usable 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 Usable 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 Usable 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 Usable 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 Usable 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 Usable 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 Usable Prefix in the LSA database and directly connected OSPFv3 interface. For the purposes of this discussion, the Usable Prefix will be referred to as the Current Usable Prefix, and the interface will be referred to as the Current Interface. The following steps will be performed for each tuple (Usable Prefix, OSPFv3 interface):

1. The OSPFv3 Router will search all AC LSAs for a Usable Prefix TLV describing a prefix which contains but is not equal to the Current Usable Prefix. If such a prefix is found, the algorithm is skipped for the Current Usable 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 Usable 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 Usable 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 Usable 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 Usable 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 Usable 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 Usable 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 Usable 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 Usable 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.

A Usable 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 usable prefix. This is done by scheduling the re-origination of an AC LSA that does not include the Usable 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 a usable 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.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 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.acee-ospf-ospfv3-autoconfig]:

- o The Usable 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.acee-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration", draft-acee-ospf-ospfv3-autoconfig-00 (work in progress), October 2011.

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.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.
- [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. 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, 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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 17, 2013

T. Chown, Ed.
University of Southampton
J. Arkko
Ericsson
A. Brandt
Sigma Designs
O. Troan
Cisco Systems, Inc.
J. Weil
Time Warner Cable
July 16, 2012

Home Networking Architecture for IPv6
draft-ietf-homenet-arch-04

Abstract

This text describes evolving networking technology within increasingly large residential home networks. The goal of this document is to define an architecture for IPv6-based home networking, while describing the associated principles, considerations and requirements. The text briefly highlights the specific implications of the introduction of IPv6 for home networking, discusses the elements of the architecture, and suggests how standard IPv6 mechanisms and addressing can be employed in home networking. The architecture describes the need for specific protocol extensions for certain additional functionality. It is assumed that the IPv6 home network is not actively managed, and runs as an IPv6-only or dual-stack network. There are no recommendations in this text for the IPv4 part of the network.

Status of this Memo

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

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

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

This Internet-Draft will expire on January 17, 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	4
1.1. Terminology and Abbreviations	5
2. Effects of IPv6 on Home Networking	6
2.1. Multiple subnets and routers	6
2.2. Global addressability and elimination of NAT	7
2.3. Multi-Addressing of devices	7
2.4. Unique Local Addresses (ULAs)	8
2.5. Naming, and manual configuration of IP addresses	9
2.6. IPv6-only operation	9
3. Homenet Architecture	10
3.1. General Principles	10
3.1.1. Reuse existing protocols	11
3.1.2. Minimise changes to hosts and routers	11
3.2. Homenet Topology	11
3.2.1. Supporting arbitrary topologies	11
3.2.2. Network topology models	11
3.2.3. Dual-stack topologies	16
3.2.4. Multihoming	17
3.3. A Self-Organising Network	18
3.3.1. Homenet realms and borders	19
3.3.2. Largest possible subnets	19
3.3.3. Handling multiple homenets	20
3.3.4. Coordination of configuration information	20
3.4. Homenet Addressing	20
3.4.1. Use of ISP-delegated IPv6 prefixes	20
3.4.2. Stable internal IP addresses	22
3.4.3. Internal prefix delegation	22
3.4.4. Privacy	24
3.5. Routing functionality	24
3.6. Security	25

3.6.1.	Addressability vs reachability	26
3.6.2.	Filtering at borders	27
3.6.3.	Device capabilities	27
3.6.4.	ULAs as a hint of connection origin	27
3.7.	Naming and Service Discovery	27
3.8.	Other Considerations	30
3.8.1.	Proxy or Extend?	30
3.8.2.	Quality of Service	30
3.8.3.	Operations and Management	31
3.9.	Implementing the Architecture on IPv6	31
4.	Conclusions	32
5.	References	32
5.1.	Normative References	32
5.2.	Informative References	33
Appendix A.	Acknowledgments	36
Appendix B.	Changes	36
B.1.	Version 04	36
B.2.	Version 03	36
B.3.	Version 02	38
Authors' Addresses	38

1. Introduction

This document focuses on evolving networking technology within increasingly large residential home networks and the associated challenges with their deployment and operation. There is a growing trend in home networking for the proliferation of networking technology in an increasingly broad range of devices and media. This evolution in scale and diversity sets requirements on IETF protocols. Some of these requirements relate to the introduction of IPv6, others to the introduction of specialised networks for home automation and sensors.

While at the time of writing some complex home network topologies exist, most operate based on IPv4, employ solutions that we would like to avoid such as (cascaded) network address translation (NAT), or require expert assistance to set up. In IPv6 home networks, there are likely to be scenarios where internal routing is required, for example to support private and guest networks, in which case such networks may use increasing numbers of subnets, and require methods for IPv6 prefixes to be delegated to those subnets. The assumption of this document is that the homenet is as far as possible self-organising and self-configuring, and is thus not pro-actively managed by the residential user.

The architectural constructs in this document are focused on the problems to be solved when introducing IPv6 with an eye towards a better result than what we have today with IPv4, as well as a better result than if the IETF had not given this specific guidance. The document aims to provide the basis and guiding principles for how standard IPv6 mechanisms and addressing [RFC2460] [RFC4291] can be employed in home networking, while coexisting with existing IPv4 mechanisms. In emerging dual-stack home networks it is vital that introducing IPv6 does not adversely affect IPv4 operation. We assume that the IPv4 network architecture in home networks is what it is, and can not be affected by new recommendations. Future deployments, or specific subnets within an otherwise dual-stack home network, may be IPv6-only, in which case considerations for IPv4 impact would not apply.

This architecture document proposes a baseline homenet architecture, based on protocols and implementations that are as far as possible proven and robust. The scope of the document is primarily the network layer technologies that provide the basic functionality to enable addressing, connectivity, routing, naming and service discovery. While it may, for example, state that homenet components must be simple to deploy and use, it does not discuss specific user interfaces, nor does it discuss specific physical, wireless or data-link layer considerations.

[RFC6204] defines basic requirements for customer edge routers (CERs). The scope of this text is the internal homenet, and thus specific features on the CER are out of scope for this text. While the network may be dual-stack or IPv6-only, the definition of specific transition tools on the CER, as introduced in RFC 6204-bis [I-D.ietf-v6ops-6204bis] with DS-Lite [RFC6333] and 6rd [RFC5969], are considered issues for that RFC, and are thus also out of scope of this text.

1.1. Terminology and Abbreviations

In this section we define terminology and abbreviations used throughout the text.

- o "Advanced Security". Describes advanced security functions for a CER, as defined in [I-D.vyncke-advanced-ipv6-security], where the default inbound connection policy is generally "default allow".
- o CER: Customer Edge Router. A border router at the edge of the homenet.
- o LLN: Low-power and lossy network.
- o NAT: Network Address Translation. Typically referring to IPv4 Network Address and Port Translation (NAPT) [RFC3022].
- o NPTv6: Network Prefix Translation for IPv6 [RFC6296].
- o PCP: Port Control Protocol [I-D.ietf-pcp-base].
- o "Simple Security". Defined in [RFC4864] and expanded further in [RFC6092]; describes recommended perimeter security capabilities for IPv6 networks.
- o ULA: IPv6 Unique Local Addresses [RFC4193].
- o UPnP: Universal Plug and Play. Includes the Internet Gateway Device (IGD) function, which for IPv6 is UPnP IGD Version 2 [IGD-2].
- o VM: Virtual machine.
- o WPA2: Wi-Fi Protected Access, as defined by the Wi-Fi Alliance.

2. Effects of IPv6 on Home Networking

Service providers are deploying IPv6, content is becoming available on IPv6 (accelerated recently by the World IPv6 Launch event) and support for IPv6 is increasingly available in devices and software used in the home. While IPv6 resembles IPv4 in many ways, it changes address allocation principles, making multi-addressing the norm, and allowing direct IP addressability of home networking devices from the Internet. This section presents an overview of some of the key implications of the introduction of IPv6 for home networking, that are simultaneously both promising and problematic.

2.1. Multiple subnets and routers

The introduction of IPv6 for home networking enables the potential for every home network to be delegated enough address space to provision globally unique prefixes for each subnet in the home. Such subnetting is not common practice in existing IPv4 homenets, but is very likely to become increasingly standard in future IPv6 homenets.

While simple layer 3 topologies involving as few subnets as possible are preferred in home networks, the incorporation of dedicated (routed) subnets remains necessary for a variety of reasons. For instance, an increasingly common feature in modern home routers is the ability to support both guest and private network subnets. Likewise, there may be a need to separate building control or corporate extensions from the main Internet access network, or different subnets may in general be associated with parts of the homenet that have different routing and security policies. Further, link layer networking technology is poised to become more heterogeneous, as networks begin to employ both traditional Ethernet technology and link layers designed for low-power and lossy networks (LLNs), such as those used for certain types of sensor devices. Constraining the flow of certain traffic from Ethernet links to much lower capacity links thus becomes an important topic.

Documents that provide some more specific background and depth on this topic include: [I-D.herbst-v6ops-cpeenhancements], [I-D.baker-fun-multi-router], and [I-D.baker-fun-routing-class].

The addition of routing between subnets raises the issue of how to extend mechanisms such as service discovery which currently rely on link-local addressing to limit scope. There are two broad choices; extend existing protocols to work across the scope of the homenet, or introduce proxies for existing link-layer protocols. This topic is discussed later in the document.

There will also be the need to discover which routers in the homenet

are the border router(s) by an appropriate mechanism. Here, there are a number of choices. These include an appropriate service discovery protocol, or the use of a well-known name, resolved by some local name service. Both might have to deal with handling more than one router responding in multihomed environments.

2.2. Global addressability and elimination of NAT

Current IPv4 home networks typically receive a single global IPv4 address from their ISP and use NAT with private [RFC1918] addresses for devices within the network. An IPv6 home network removes the need to use NAT given the ISP offers a sufficiently large globally unique IPv6 prefix to the homenet, allowing every device on every link to be assigned a globally unique IPv6 address.

The end-to-end communication that is potentially enabled with IPv6 is on the one hand an incredible opportunity for innovation and simpler network operation, but it is also a concern as it exposes nodes in the internal networks to receipt of otherwise unwanted traffic from the Internet. There may thus be an expectation of improved host security to compensate for this, at least in general networked devices, but it must be noted that many devices may also (for example) ship with default settings that make them readily vulnerable to compromise by external attackers if globally accessible, or may simply not have robustness designed-in because it was either assumed such devices would only be used on private networks or the device itself doesn't have the computing power to apply the necessary security methods.

IPv6 networks may or may not have filters applied at their borders, i.e. at the homenet CER. [RFC4864], [RFC6092] and [I-D.vyncke-advanced-ipv6-security] discuss such filtering, and the merits of "default allow" against "default deny" policies for external traffic initiated into a homenet. It is important to distinguish between addressability and reachability. While IPv6 offers global addressability through use of globally unique addresses in the home, whether they are globally reachable or not would depend on the firewall or filtering configuration, and not, as is commonly the case with IPv4, the presence or use of NAT.

2.3. Multi-Addressing of devices

In an IPv6 network, devices may acquire multiple addresses, typically at least a link-local address and a globally unique address. They may also have an IPv4 address if the network is dual-stack, a Unique Local Address (ULA) [RFC4193] (see below), and one or more IPv6 Privacy Addresses [RFC4941].

Thus it should be considered the norm for devices on IPv6 home networks to be multi-addressed, and to need to make appropriate address selection decisions for the candidate source and destination address pairs. Default Address Selection for IPv6 [I-D.ietf-6man-rfc3484bis] provides a solution for this, though it may face problems in the event of multihoming, where nodes will be configured with one address from each upstream ISP prefix. In such cases the presence of upstream ingress filtering requires multi-addressed nodes to select the correct source address to be used for the corresponding uplink, to avoid ISP BCP 38 ingress filtering, but the node may not have the information it needs to make that decision based on addresses alone. We discuss such challenges in the multihoming section later in this document.

2.4. Unique Local Addresses (ULAs)

[RFC4193] defines Unique Local Addresses (ULAs) for IPv6 that may be used to address devices within the scope of a single site. Support for ULAs for IPv6 CERNs is described in [RFC6204]. A home network running IPv6 may deploy ULAs for stable communication between devices (on different subnets) within the network where the externally allocated global prefix changes over time (e.g. due to renumbering within the subscriber's ISP) or where external connectivity is temporarily unavailable.

A counter-argument to using ULAs is that it is undesirable to aggressively deprecate global prefixes for temporary loss of connectivity, so for a host to lose its global address there would have to be a connection breakage longer than the lease period, and even then, deprecating prefixes when there is no connectivity may not be advisable. It should also be noted that there may be timers on the prefix lease to the homenet, on the internal prefix delegations, and on the Router Advertisements to the hosts. Despite this counter-argument, while setting a network up there may be a period with no connectivity, in which case ULAs would be required for inter-subnet communication. In the case where LLNs are being set up in a new home/deployment, individual LLNs may, at least initially, each use their own /48 ULA prefix.

Default address selection mechanisms should ensure a ULA source address is used to communicate with ULA destination addresses when appropriate, in particular when the ULA destination lies within a /48 ULA prefix known to be used within the same homenet. Note that unlike the IPv4 private RFC 1918 space, the use of ULAs does not imply use of host-based IPv6 NAT, or NPTv6 prefix-based NAT [RFC6296], rather that external communications should use a node's additional globally unique IPv6 source address.

2.5. Naming, and manual configuration of IP addresses

Some IPv4 home networking devices expose IPv4 addresses to users, e.g. the IPv4 address of a home IPv4 CER that may be configured via a web interface. Users should not be expected to enter IPv6 literal addresses in homenet devices or applications, given their much greater length and apparent randomness to a typical home user. While shorter addresses, perhaps ones registered with IANA from ULA-C space [I-D.hain-ipv6-ulac], could be used for specific devices/services, in general it is better to not expose users to real IPv6 addresses. Thus, even for the simplest of functions, simple naming and the associated (ideally zero configuration) discovery of services is imperative for the easy deployment and use of homenet devices and applications.

In a multi-subnet homenet, naming and service discovery should be expected to be capable of operating across the scope of the entire home network, and thus be able to cross subnet boundaries. It should be noted that in IPv4, such services do not generally function across home router NAT boundaries, so this is one area where there is room for improvement in IPv6.

2.6. IPv6-only operation

It is likely that IPv6-only networking will be deployed first in "greenfield" homenet scenarios, or perhaps as one element of an otherwise dual-stack network. Running IPv6-only adds additional requirements, e.g. for devices to get configuration information via IPv6 transport (not relying on an IPv4 protocol such as IPv4 DHCP), and for devices to be able to initiate communications to external devices that are IPv4-only. Thus, for example, the following requirements are amongst those that should be considered in IPv6-only environments:

- o Ensuring there is a way to access content in the IPv4 Internet. This can be arranged through incorporating NAT64 [RFC6144] and DNS64 [RFC6145] functionality in the home gateway router, for instance. Such features are outside the scope of this document however, being CER functions.
- o DNS discovery mechanisms are enabled for IPv6. Both stateless DHCPv6 [RFC3736] [RFC3646] and Router Advertisement options [RFC6106] may have to be supported and turned on by default to ensure maximum compatibility with all types of hosts in the network. This requires, however, that a working DNS server is known and addressable via IPv6, and that such discovery options can operate through multiple routers in the homenet.

- o All nodes in the home network support operations in IPv6-only mode. Some current devices work well with dual-stack but fail to recognise connectivity when IPv4 DHCP fails, for instance.

The widespread availability of robust solutions to these types of requirements will help accelerate the uptake of IPv6-only homenets.

3. Homenet Architecture

The aim of this architecture text is to outline how to construct advanced IPv6-based home networks involving multiple routers and subnets using standard IPv6 protocols and addressing [RFC2460] [RFC4291]. In this section, we present the elements of such a home networking architecture, with discussion of the associated design principles.

Existing IETF work [RFC6204] defines the "basic" requirements for Customer Edge Routers, while [I-D.ietf-v6ops-6204bis] extends RFC 6204 to describe additional features. The homenet architecture is focused on the internal homenet, rather than the CER(s). In general, home network equipment needs to be able to operate in networks with a range of different properties and topologies, where home users may plug components together in arbitrary ways and expect the resulting network to operate. Significant manual configuration is rarely, if at all, possible, given the knowledge level of typical home users. Thus the network should, as far as possible, be self-configuring.

The equipment also needs to be prepared to handle at least

- o Routing
- o Prefix configuration for routers
- o Name resolution
- o Service discovery
- o Network security

The remainder of this document describes the principles by which a homenet architecture may deliver these properties.

3.1. General Principles

There is little that the Internet standards community can do about the physical topologies or the need for some networks to be separated at the network layer for policy or link layer compatibility reasons.

However, there is a lot of flexibility in using IP addressing and inter-networking mechanisms. This architecture text discusses how this flexibility should be used to provide the best user experience and ensure that the network can evolve with new applications in the future. The principles described in this text should be followed when designing homenet solutions.

3.1.1. Reuse existing protocols

It is desirable to reuse existing protocols where possible, but at the same time to avoid consciously precluding the introduction of new or emerging protocols. A generally conservative approach, giving weight to running code, is preferable. Where new protocols are required, evidence of commitment to implementation by appropriate vendors or development communities is highly desirable. Protocols used should be backwardly compatible, and forward compatible where changes are made.

3.1.2. Minimise changes to hosts and routers

Where possible, any requirement for changes to hosts and routers should be minimised, though solutions which, for example, incrementally improve with host changes may be acceptable.

3.2. Homenet Topology

In this section we consider homenet topologies, and the principles we may apply in designing an architecture to support as wide a range as possible of such topologies.

3.2.1. Supporting arbitrary topologies

There should ideally be no built-in assumptions about the topology in home networks, as users are capable of connecting their devices in "ingenious" ways. Thus arbitrary topologies and arbitrary routing will need to be supported, or at least the failure mode for when the user makes a mistake should be as robust as possible, e.g. de-activating a certain part of the infrastructure to allow the rest to operate. In such cases, the user should ideally have some useful indication of the failure mode encountered.

3.2.2. Network topology models

Most IPv4 home network models at the time of writing tend to be relatively simple, typically a single NAT router to the ISP and a single internal subnet but, as discussed earlier, evolution in network architectures is driving more complex topologies, such as the separation of visitor and private networks.

In general, the models described in [RFC6204] and its successor RFC 6204-bis [I-D.ietf-v6ops-6204bis] should be supported by the IPv6 home networking architecture. The functions resident on the CER itself are, as stated previously, out of scope of this text.

There are a number of properties or attributes of a home network that we can use to describe its topology and operation. The following properties apply to any IPv6 home network:

- o Presence of internal routers. The homenet may have one or more internal routers, or may only provide subnetting from interfaces on the CER.
- o Presence of isolated internal subnets. There may be isolated internal subnets, with no direct connectivity between them within the homenet. Isolation may be physical, or implemented via IEEE 802.1q VLANs.
- o Demarcation of the CER. The CER(s) may or may not be managed by the ISP. If the demarcation point is such that the customer can provide or manage the CER, its configuration must be simple. Both models must be supported.

Various forms of multihoming are likely to be more prevalent with IPv6 home networks, as discussed further below. Thus the following properties should also be considered for such networks:

- o Number of upstream providers. A typical homenet might just have a single upstream ISP, but it may become more common for there to be multiple ISPs, whether for resilience or provision of additional services. Each would offer its own prefix. Some may or may not be walled gardens.
- o Number of CERs. The homenet may have a single CER, which might be used for one or more providers, or multiple CERs. The presence of multiple CERs adds additional complexity for multihoming scenarios, and protocols like PCP that need to manage connection-oriented state mappings.

A separate discussion of physical infrastructures for homenets is included in and [I-D.arkko-homenet-physical-standard].

In the following sections we give some examples of the types of homenet topologies we may see in the future. This is not intended to be an exhaustive or complete list, rather an indicative one to facilitate the discussion in this text.

3.2.2.1. A: Single ISP, Single CER, Internal routers

Figure 1 shows a network with multiple local area networks. These may be needed for reasons relating to different link layer technologies in use or for policy reasons, e.g. classic Ethernet in one subnet and a LLN link layer technology in another. In this example there is no single router that a priori understands the entire topology. The topology itself may also be complex, and it may not be possible to assume a pure tree form, for instance (home users may plug routers together to form arbitrary topologies including loops).

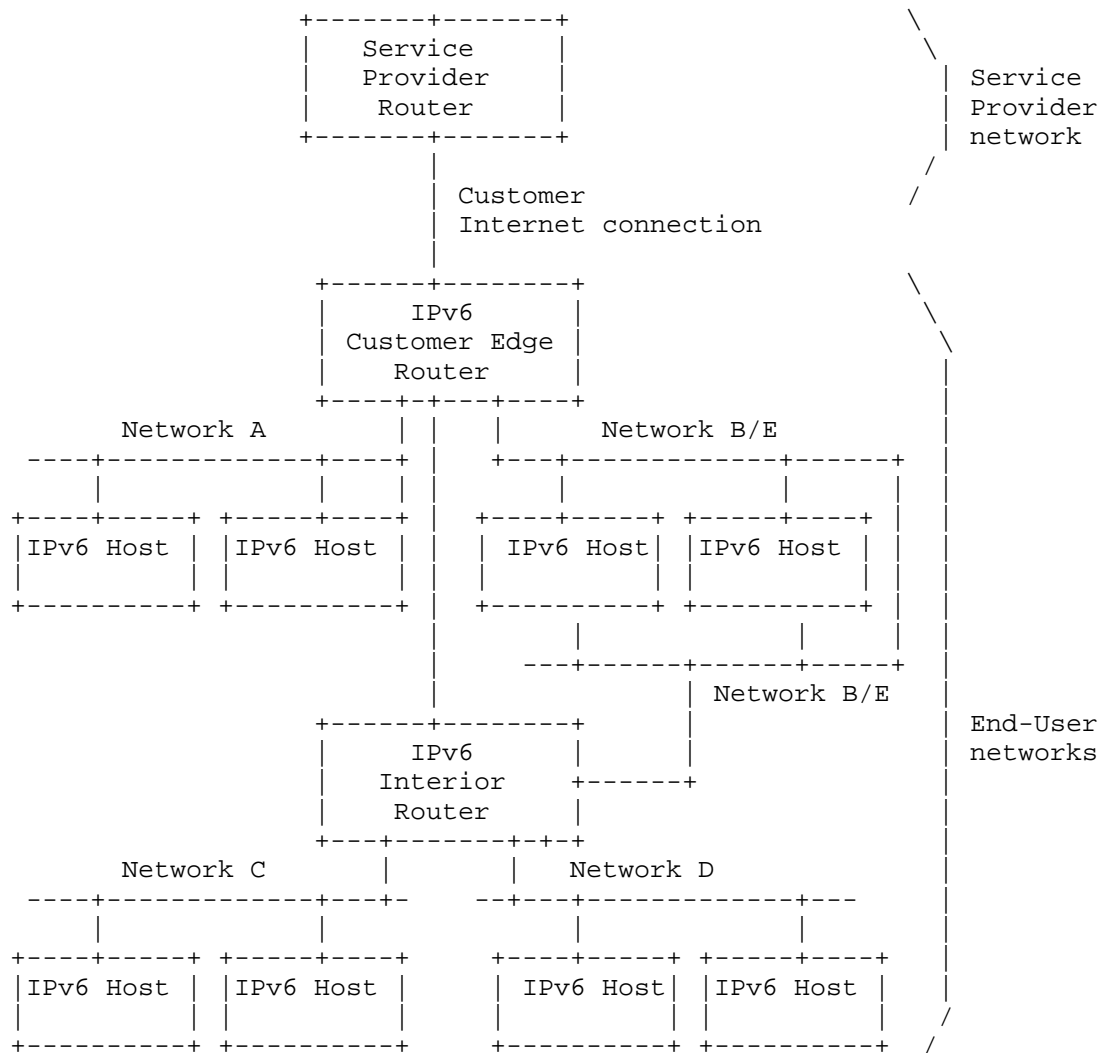


Figure 1

3.2.2.2. B: Two ISPs, Two CERs, Shared subnet

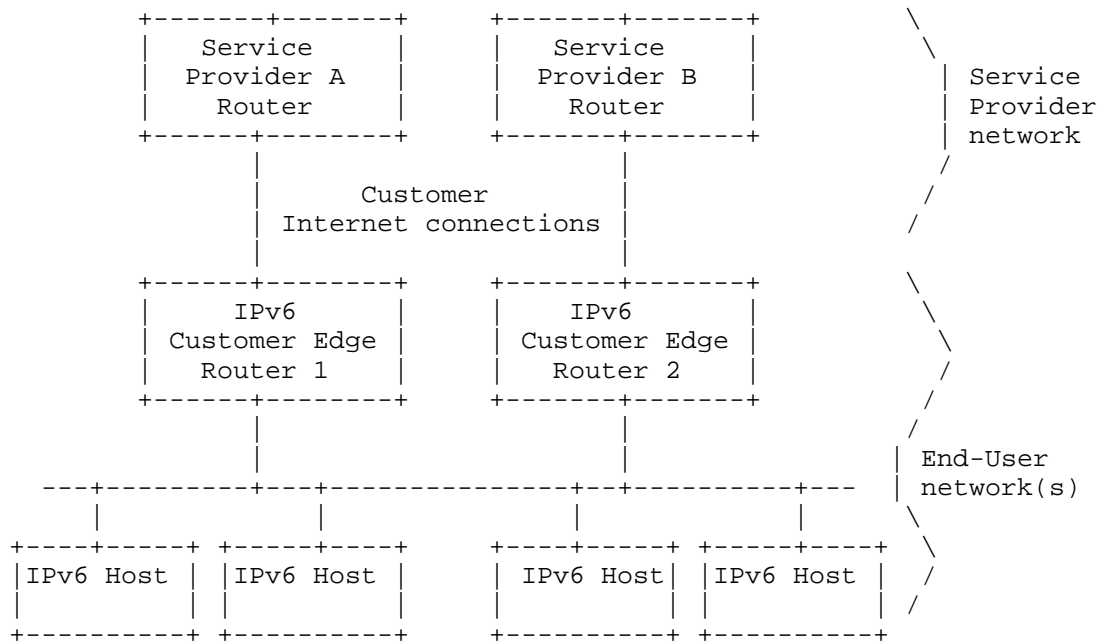


Figure 2

Figure 2 illustrates a multihomed homenet model, where the customer has connectivity via CER1 to ISP A and via CER2 to ISP B. This example shows one shared subnet where IPv6 nodes would potentially be multihomed and receive multiple IPv6 global addresses, one per ISP. This model may also be combined with that shown in Figure 1 to create a more complex scenario with multiple internal routers. Or the above shared subnet may be split in two, such that each CER serves a separate isolated subnet, which is a scenario seen with some IPv4 networks today.

3.2.2.3. C: Two ISPs, One CER, Shared subnet

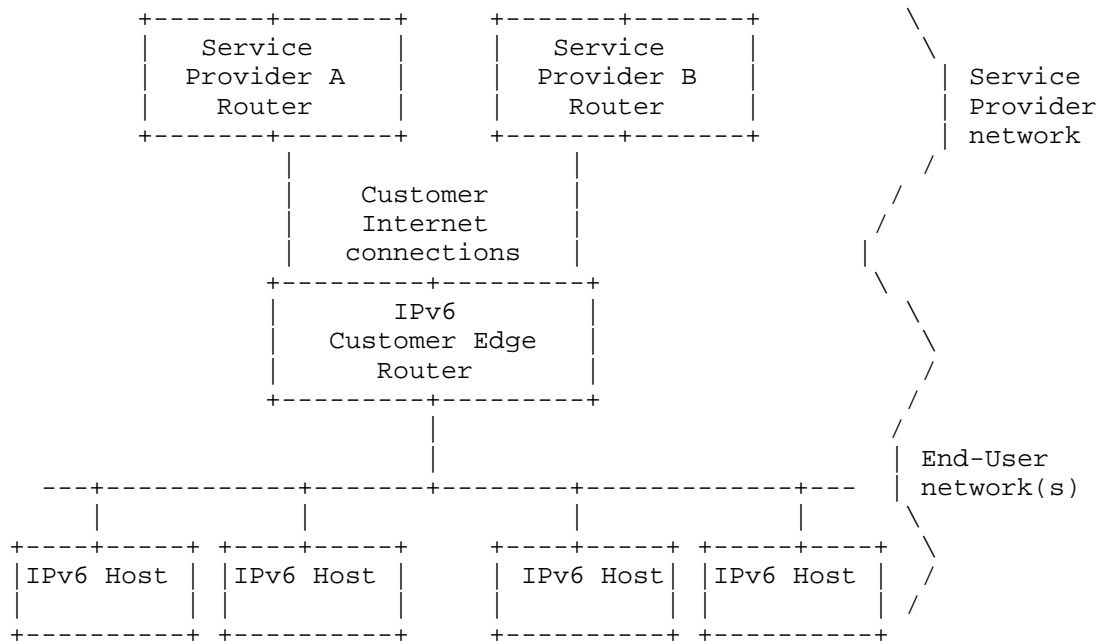


Figure 3

Figure 3 illustrates a model where a home network may have multiple connections to multiple providers or multiple logical connections to the same provider, with shared internal subnets.

In general, while the architecture may focus on likely common topologies, it should not preclude any arbitrary topology from being constructed.

3.2.3. Dual-stack topologies

It is expected that most homenet deployments will for the immediate future be dual-stack IPv4/IPv6. In such networks it is important not to introduce new IPv6 capabilities that would cause a failure if used alongside IPv4+NAT, given that such dual-stack homenets will be commonplace for some time. That said, it is desirable that IPv6 works better than IPv4 in as many scenarios as possible. Further, the homenet architecture must operate in the absence of IPv4.

A general recommendation is to follow the same topology for IPv6 as is used for IPv4, but not to use NAT. Thus there should be routed

IPv6 where an IPv4 NAT is used, and where there is no NAT there should be bridging if the link layer allows this.

In some cases IPv4 NAT home networks may feature cascaded NATs, which may include cases where NAT routers are included within VMs, or where Internet connection sharing services are used. IPv6 routed versions of such cases will be required. We should thus note that routers in the homenet may not be separate physical devices; they may be embedded within other devices.

3.2.4. Multihoming

A homenet may be multihomed to multiple providers, as the network models above illustrate. This may either take a form where there are multiple isolated networks within the home or a more integrated network where the connectivity selection needs to be dynamic. Current practice is typically of the former kind, but the latter is expected to become more commonplace.

The general multihoming problem is broad, and solutions suggested to date within the IETF may include complex architectures for monitoring connectivity, traffic engineering, identifier-locator separation, connection survivability across multihoming events, and so on. It is thus important that the homenet architecture should as far as possible minimise the complexity of any multihoming support. So we should limit the support to the smallest subset of the overall problem to meet the requirements of the topologies described above. This means that the homenet architecture should not try to make another attempt at solving complex multihoming, and we should prefer to support scenarios for which solutions exist today.

In the general homenet architecture, hosts should be multi-addressed with globally unique prefixes from each ISP they may communicate with or through. An alternative for a homenet would be to deploy NPTv6 [RFC6296] at the CER, with ULAs then typically used internally, but this mode is not considered by this text. If NPTv6 is used, the internal part of the homenet (which is the scope of this text) simply sees only the one (ULA) prefix in use. It should be noted that running NPTv6 has an architectural cost, due to the prefix translation used.

When multi-addressing is in use, hosts need some way to pick source and destination address pairs for connections. A host may choose a source address to use by various methods, which would typically include [I-D.ietf-6man-rfc3484bis]. Applications may of course do different things, and this should not be precluded.

For the single CER Network Model C, multihoming may be offered by

source routing at the CER. With multiple exit routers, the complexity rises. Given a packet with a source address on the network, the packet must be routed to the proper egress to avoid BCP 38 filtering at an ISP that did not delegate the prefix the address is chosen from. While the packet might not take an optimal path to the correct exit CER, the minimum requirement is that the packet is not dropped. It is of course highly desirable that the packet is routed in the most efficient manner to the correct exit.

There are various potential approaches to this problem, one example being described in [I-D.v6ops-multihoming-without-ipv6nat]. Another is discussed in [I-D.baker-fun-multi-router], which explores support for source routing throughout the homenet. This approach would however likely require relatively significant routing changes to route the packet to the correct exit given the source address. Such changes should preferably be minimised.

There are some other multihoming considerations for homenet scenarios. First, it may be the case that multihoming applies due to an ISP migration from a transition method to a native deployment, e.g. a 6rd [RFC5969] sunset scenario, as discussed in [I-D.townsley-troan-ipv6-ce-transitioning]. Second, one upstream may be a "walled garden", and thus only appropriate to be used for connectivity to the services of that provider; an example may be a VPN service that only routes back to the enterprise business network of a user in the homenet. While we should not specifically target walled garden multihoming as a principal goal, it should not be precluded.

Host-based methods such as Shim6 [RFC5533] have been defined, but of course require support in the hosts. There are also application-oriented approaches such as Happy Eyeballs [RFC6555]; simplified versions of this are for example already implemented in some commonly-used web browsers. The homenet architecture should not preclude use of such tools should hosts include their support.

3.3. A Self-Organising Network

A home network architecture should be naturally self-organising and self-configuring under different circumstances relating to the connectivity status to the Internet, number of devices, and physical topology. While the homenet should be self-organising, it should be possible to manually adjust (override) the current configuration.

While a goal of the homenet architecture is for the network to be as self-organising as possible, there may be instances where some manual configuration is required, e.g. the entry of a WPA2 key to apply wireless security, or to configure a shared routing secret. The

latter may be relevant when considering how to bootstrap a routing configuration. It is highly desirable that only one such key is needed for any set of functions, to increase usability for the homenet user.

3.3.1. Homenet realms and borders

The homenet will need to be aware of the extent of its own "site", which will define the borders for ULAs, site scope multicast, service discovery and security policies. The homenet will have one or more borders with external connectivity providers and potentially also have borders within the internal network (e.g. for policy-based reasons). It should be possible to automatically perform border discovery for the different borders. Such borders determine for example the scope of where prefixes, routing information, network traffic, service discovery and naming may be shared. The default internally should be to share everything.

A simple homenet model may just consider three types of realm and the borders between them. For example if the realms are the homenet, the ISP and the visitor network, then the borders will include that from the homenet to the ISP, and that from the homenet to a guest network. Regardless, it should be possible for additional types of realms and borders to be defined, e.g. for some specific Grid or LLN-based network, and for these to be detected automatically, and for an appropriate default policy to be applied as to what type of traffic/data can flow across such borders.

It is desirable to classify the external border of the home network as a unique logical interface separating the home network from service provider network/s. This border interface may be a single physical interface to a single service provider, multiple layer 2 sub-interfaces to a single service provider, or multiple connections to a single or multiple providers. This border makes it possible to describe edge operations and interface requirements across multiple functional areas including security, routing, service discovery, and router discovery.

It should be possible for the homenet user to override any automatically determined borders and the default policies applied between them.

3.3.2. Largest possible subnets

Today's IPv4 home networks generally have a single subnet, and early dual-stack deployments have a single congruent IPv6 subnet, possibly with some bridging functionality. More recently, some vendors have started to introduce "home" and "guest" functions, which in IPv6

would be implemented as two subnets.

Future home networks are highly likely to have one or more internal routers and thus need multiple subnets, for the reasons described earlier. As part of the self-organisation of the network, the homenet should subdivide itself to the largest possible subnets that can be constructed within the constraints of link layer mechanisms, bridging, physical connectivity, and policy.

While it may be desirable to maximise the chance of link-local protocols operating across a homenet by maximising the size of a subnet, multi-subnet home networks are inevitable, so their support must be included.

3.3.3. Handling multiple homenets

It is important that self-configuration with "unintended" devices is avoided. Methods are needed for devices to know whether they are intended to be part of the same homenet site or not. Thus methods to ensure separation between neighbouring homenets are required. This may require use of some unique "secret" for devices/protocols in each homenet. Some existing mechanisms exist to assist home users to associate devices as simply as possible, e.g. "connect" button support.

3.3.4. Coordination of configuration information

The network elements will need to be integrated in a way that takes account of the various lifetimes on timers that are used on different elements, e.g. DHCPv6 PD, router, valid prefix and preferred prefix timers.

3.4. Homenet Addressing

The IPv6 addressing scheme used within a homenet must conform to the IPv6 addressing architecture [RFC4291]. The homenet will need to adapt to the prefixes made available to it through the prefix delegation method used by its upstream ISP.

3.4.1. Use of ISP-delegated IPv6 prefixes

A homenet may receive an arbitrary length IPv6 prefix from its provider, e.g. /60, /56 or /48. The offered prefix may be stable or change from time to time. Some ISPs may offer relatively stable prefixes, while others may change the prefix whenever the CER is reset. Some discussion of IPv6 prefix allocation policies is included in [RFC6177] which discusses why, for example, a one-size-fits-all /48 allocation is not desirable. The home network needs to

be adaptable to such ISP policies, and thus make no assumptions about the stability of the prefix received from an ISP, or the length of the prefix that may be offered. However, if only a /64 is offered by the ISP, the homenet may be severely constrained, or even unable to function.

The internal operation of the home network should also not depend on the availability of the ISP network at any given time, other than for connectivity to services or systems off the home network. This implies the use of ULAs for stable internal communication, as described in the next section.

In practice, it is expected that ISPs will deliver a relatively stable home prefix to customers. The norm for residential customers of large ISPs may be similar to their single IPv4 address provision; by default it is likely to remain persistent for some time, but changes in the ISP's own provisioning systems may lead to the customer's IP (and in the IPv6 case their prefix pool) changing. It is not expected that ISPs will support Provider Independent (PI) addressing for general residential homenets.

When an ISP needs to restructure and in doing so renumber its customer homenets, "flash" renumbering is likely to be imposed. This implies a need for the homenet to be able to handle a sudden renumbering event which, unlike the process described in [RFC4192], would be a "flag day" event, which means that a graceful renumbering process moving through a state with two active prefixes in use would not be possible. While renumbering is an extended version of an initial numbering process, the difference between flash renumbering and an initial "cold start" is the need to provide service continuity.

There may be cases where local law means some ISPs are required to change IPv6 prefixes (current IPv4 addresses) for privacy reasons for their customers. In such cases it may be possible to avoid an instant "flash" renumbering and plan a non-flag day renumbering as per RFC 4192.

The customer may of course also choose to move to a new ISP, and thus begin using a new prefix. In such cases the customer should expect a discontinuity, and not only may the prefix change, but potentially also the prefix length, if the new ISP offers a different default size prefix, e.g. a /60 rather than a /56. Regardless, it's desirable that homenet protocols support rapid renumbering and that operational processes don't add unnecessary complexity for the renumbering process.

The 6renum WG is studying IPv6 renumbering for enterprise networks.

It is not currently targetting homenet, but may produce outputs that are relevant. The introduction of any new homenet protocols should not make any form of renumbering any more complex than it already is.

3.4.2. Stable internal IP addresses

The network should by default attempt to provide IP-layer connectivity between all internal parts of the homenet as well as to and from the external Internet, subject to the filtering policies or other policy constraints discussed later in the security section.

ULAs should be used within the scope of a homenet to support routing between subnets regardless of whether a globally unique ISP-provided prefix is available. It would be expected that ULAs would be used alongside one or more such global prefixes in a homenet, such that hosts become multi-addressed with both globally unique and ULA prefixes. Default address selection would then enable ULAs to be preferred for internal communications between devices that are using ULA prefixes generated within the same homenet.

ULA addresses will allow constrained LLN devices to create permanent relationships between IPv6 addresses, e.g. from a wall controller to a lamp. Symbolic host names would require additional non-volatile memory. Updating global prefixes in sleeping LLN devices might also be problematic.

ULAs may be used for all devices, not just those intended to only have internal connectivity. ULAs used in this way provide stable internal communications should the ISP-provided prefix (suddenly) change, or external connectivity be temporarily lost. The use of ULAs should be restricted to the homenet scope through filtering at the border(s) of the homenet, as described in RFC 6092.

3.4.3. Internal prefix delegation

As mentioned above, there are various sources of prefixes, e.g. they may be globally unique prefixes originating from ISP(s), they may be globally unique or ULA prefixes allocated by "master" router(s) in the homenet, or they may be ULAs allocated by LLN gateways. There may also be a prefix associated with NAT64, if in use in the homenet.

From the homenet perspective, a single prefix from each ISP should be received on the border CER [RFC3633]. Then each subnet in the homenet should receive a prefix from within the ISP-provided prefix(es). The ISP should only see the aggregate from the homenet, and not single /64 prefixes allocated within the homenet.

Delegation should be autonomous, and not assume a flat or

hierarchical model. This text makes no assumption about whether the delegation of prefixes is distributed or centralised. The assignment mechanism should provide reasonable efficiency, so that typical home network prefix allocation sizes can accommodate all the necessary /64 allocations in most cases, and not waste prefixes. A currently typical /60 allocation gives 16 /64 subnets. Duplicate assignment of multiple /64s to the same network should be avoided. The network should behave as gracefully as possible in the event of prefix exhaustion, though the options in such cases may be limited.

Where multiple CERS exist with multiple ISP prefix pools, it is expected that routers within the homenet would assign themselves prefixes from each ISP they communicate with/through.

Where ULAs are used, most likely but not necessarily in parallel with global prefixes, one router should be elected to offer ULA prefixes for the homenet. The router should generate a /48 ULA for the site, and then delegate /64's from that ULA prefix to subnets. In the normal state, a single /48 ULA should be used within the homenet. In cases where two /48 ULAs are generated within a homenet, the network should still continue to function.

Delegation within the homenet should give each link a prefix that is persistent across reboots, power outages and similar short-term outages. Addition of a new routing device should not affect existing persistent prefixes, but persistence may not be expected in the face of significant "replumbing" of the homenet. Persistent prefixes should not depend on router boot order. Such persistent prefixes may imply the need for stable storage on routing devices, and also a method for a home user to "reset" the stored prefix should a significant reconfiguration be required (though ideally the home user should not be involved at all).

The delegation method should support renumbering, which would typically be "flash" renumbering in that the homenet would not have advance notice of the event or thus be able to apply the types of approach described in [RFC4192]. As a minimum, delegated ULA prefixes within the homenet should remain persistent through an ISP-driven renumbering event.

Several proposals have been made for prefix delegation within a homenet. One group of proposals is based on DHCPv6 PD, as described in [I-D.baker-homenet-prefix-assignment], [I-D.chakrabarti-homenet-prefix-alloc], [RFC3315] and [RFC3633]. The other uses OSPFv3, as described in [I-D.arkko-homenet-prefix-assignment]. More detailed analysis of these approaches needs to be made against the requirements/principles described above.

3.4.4. Privacy

There are no specific privacy concerns discussed in this text. It should be noted as above that many ISPs are expected to offer relatively stable IPv6 prefixes to customers, and thus the network prefix associated with the host addresses they use may not change over a reasonably long period of time. This exposure is similar to IPv4 networks that expose the same IPv4 global address via use of NAT, where the IPv4 address received from the ISP may change over time, but not necessarily that frequently.

Hosts inside an IPv6 homenet may get new IPv6 addresses over time regardless, e.g. through Privacy Addresses [RFC4941].

3.5. Routing functionality

Routing functionality is required when there are multiple routers deployed within the internal home network. This functionality could be as simple as the current "default route is up" model of IPv4 NAT, or, more likely, it would involve running an appropriate routing protocol.

The homenet routing protocol should preferably be an existing deployed protocol that has been shown to be reliable and robust, and it is preferable that the protocol is "lightweight". It is desirable that the routing protocol has knowledge of the homenet topology, which implies a link-state protocol is preferable. If so, it is also desirable that the announcements and use of LSAs and RAs are appropriately coordinated. This would mean the routing protocol gives a consistent view of the network, and that it can pass around more than just routing information.

Multiple interface PHYs must be accounted for in the homenet routed topology. Technologies such as Ethernet, WiFi, MoCA, etc must be capable of coexisting in the same environment and should be treated as part of any routed deployment. The inclusion of the PHY layer characteristics including bandwidth, loss, and latency in path computation should be considered for optimising communication in the homenet. Multiple upstreams should be supported, as described in the multihoming section earlier. This should include load-balancing to multiple providers, and failover from a primary to a backup link when available. The protocol however should not require upstream ISP connectivity to be established to continue routing within the homenet.

To support multihoming within a homenet, a routing protocol that can make routing decisions based on source and destination addresses is desirable, to avoid upstream ISP ingress filtering problems. In

general the routing protocol should support multiple ISP uplinks and delegated prefixes in concurrent use.

The routing environment should be self-configuring, as discussed previously. An example of how OSPFv3 can be self-configuring in a homenet is described in [I-D.acee-ospf-ospfv3-autoconfig]. Minimising convergence time should be a goal in any routed environment, but as a guideline a maximum convergence time of around 30 seconds should be the target.

Any routed solution will require a means for determining the boundaries of the homenet. Borders may include but are not limited to the interface to the upstream ISP, or a gateway device to a separate home network such as a SmartGrid or similar LLN network. In some cases there may be no border such as occurs before an upstream connection has been established. The border discovery functionality may be integrated into the routing protocol itself, but may also be imported via a separate discovery mechanism.

In general, LLN or other networks should be able to attach and participate the same way as the main homenet, or alternatively map/be gatewayed to the main homenet. Current home deployments use largely different mechanisms in sensor and basic Internet connectivity networks. IPv6 VM solutions may also add additional routing requirements.

[I-D.howard-homenet-routing-comparison] contains evaluations of common routing protocols made against the type of requirements described above.

3.6. Security

The security of an IPv6 homenet is an important consideration. The most notable difference to the IPv4 operational model is the removal of NAT, the introduction of global addressability of devices, and thus a need to consider whether devices should have global reachability. However, there are other challenges introduced, e.g. default filtering policies at the borders between other homenet realms.

There is no defined "threat model" as such for the type of IPv6 homenet described in this text. Such a document may be very useful. It may include a variety of perspectives, from probing for specific types of home appliance being present, to potential denial of service attacks. Hosts need to be able to operate securely, end-to-end where required, but also be robust against malicious traffic direct towards them. We simply note at this point that software on home devices will have an increase in security if it allows its software to be

updated regularly.

3.6.1. Addressability vs reachability

An IPv6-based home network architecture should embrace and naturally offer a transparent end-to-end communications model as described in [RFC2775]. Each device should be addressable by a globally unique address, and those addresses must not be altered in transit. Security perimeters can (via policy) restrict end-to-end communications, and thus while a host may be globally addressable it may not be globally reachable.

In IPv4 NAT networks, the NAT provides an implicit firewall function. [RFC4864] describes a "Simple Security" model for IPv6 networks, whereby stateful perimeter filtering can be applied instead where global addresses are used. RFC 4864 implies an IPv6 "default deny" policy for inbound connections be used for similar functionality to IPv4 NAT. It should be noted that such a "default deny" approach would effectively replace the need for IPv4 NAT traversal protocols with a need to use a signalling protocol to request a firewall hole be opened. Thus to support applications wanting to accept connections initiated into home networks where a "default deny" policy is in place support for a signalling protocol such as UPnP or PCP [I-D.ietf-pcp-base] is required. In networks with multiple CERs, the signalling would need to handle the cases of flows that may use one or more exit routers. CERs would need to be able to advertise their existence for such protocols.

[RFC6092] expands on RFC 4864, giving a more detailed discussion of IPv6 perimeter security recommendations, without mandating a "default deny" approach. Indeed, RFC 6092 does not proscribe a particular mode of operation, instead stating that CERs must provide an easily selected configuration option that permits a "transparent" mode of operation, thus ensuring a "default allow" model is available. The homenet architecture text makes no recommendation on the default setting, and refers the reader to RFC 6092, which in turn simply states that a CER should provide functionality sufficient to support the recommendations in that RFC.

Advanced Security for IPv6 CPEs [I-D.vyncke-advanced-ipv6-security] takes the approach that in order to provide the greatest end-to-end transparency as well as security, security policies must be updated by a trusted party which can provide intrusion signatures and other "active" information on security threats. This might for example allow different malware detection profiles to be configured on a CER. Such methods should be able to be automatically updating.

3.6.2. Filtering at borders

It is desirable that there are mechanisms to detect different types of borders within the homenet, as discussed previously, and then the means to apply different types of filtering policies at those borders, e.g. whether naming and service discovery should pass a given border. Any such policies should be able to be easily applied by typical home users, e.g. to give a visitor in a "guest" network access to media services in the home, or access to a printer in the residence. Simple mechanisms to apply policy changes, or associations between devices, will be required.

There are cases where full internal connectivity may not be desirable, e.g. in certain utility networking scenarios, or where filtering is required for policy reasons against guest network subnet(s). Some scenarios/models may as a result involve running isolated subnet(s) with their own CERs. In such cases connectivity would only be expected within each isolated network (though traffic may potentially pass between them via external providers).

LLNs provide an another example of where there may be secure perimeters inside the homenet. Constrained LLN nodes may implement WPA2-style network key security but may depend on access policies enforced by the LLN border router.

3.6.3. Device capabilities

In terms of the devices, homenet hosts should implement their own security policies in accordance to their computing capabilities. They should have the means to request transparent communications to be initiated to them, either for all ports or for specific services. Users should have simple methods to associate devices to services that they wish to operate transparently through (CER) borders.

3.6.4. ULAs as a hint of connection origin

It has been suggested that using ULAs would provide an indication to applications that received traffic is locally sourced. This could then be used with security settings to designate where a particular application is allowed to connect to or receive traffic from.

3.7. Naming and Service Discovery

Naming and service discovery must be supported in the homenet. The service(s) providing this function must support unmanaged operation.

The most natural way to think about such naming and service discovery is to enable it to work across the entire homenet residence (site),

disregarding technical borders such as subnets but potentially respecting policy borders such as those between visitor and internal network realms.

Users will want simple ways to name devices, or be provided with appropriate ways for devices to generate unique names within the homenet. Users may typically perform device (re)naming and discovery through GUI interfaces that hide the local domain name element from them. Users may also wish to associated named devices to Internet domains, so that devices in their homenet can be accessed remotely. Thus from the user's perspective a device is given a name; the user may expect that same unqualified name to be valid within the local name service or through an Internet name service. This implies relative name resolution should be supported, i.e. there is some naming convention that allows name resolution while mitigating the need for the user to know an absolute location in the Internet name space. Or that there is some means to discover the domain transparently to the user.

Homenet devices may thus appear in one or more local homenet name spaces and also in one or more Internet name spaces. While typically there would be only one local name space, there may be scenarios where segmentation of that name space may be desirable. The naming system will be required to work internally or externally, be the user within the homenet or outside it, and there may be multiple naming domains used for any given device, e.g. Internet, home or guest domains. It is likely that a home user will want access to many of the devices and services in their home while "roaming" elsewhere. However, it may be the case that not all devices in the homenet are made available by name via an Internet name space, and that a "split view" is preferred for certain devices.

The homenet name service must therefore at the very least co-exist with Internet name services. There are naming protocols that are designed to be configured and operate Internet-wide, like unicast-based DNS, but also protocols that are designed for zero-configuration local environments, like mDNS. Consideration should be made for how these interact with each other in a homenet scenario.

The homenet name service should support both lookups and discovery. A lookup would operate via a direct query to a known service, while discovery may use multicast messages (as per mDNS and DNS-SD) or a service where applications register in order to be found.

Name resolution and service discovery for reachable devices must continue to function if the local network is disconnected from the global Internet, e.g. a local media server should still be available even if the Internet link is down for an extended period. This

implies the local network should also be able to perform a complete restart in the absence of external connectivity, and have local naming and discovery operate correctly. This might be achieved via a local cache and an authoritative local name service. Also, a change in ISP should also not affect local naming and service discovery.

There should be consideration of the security of any local name space. A typical problem here may be that many homenets may use a common "well-known" local domain suffix, e.g. .local, and this may be ambiguous to a device that could attach to multiple homenets that use that name, but this is also part of the "avoid joining unintended networks" problem. A method to utilise a local trust anchor is desirable.

With the introduction of new "dotless" top level domains, there is potential for ambiguity between for example a local host called "computer" and (if it is registered) a .computer gTLD. This suggests some implicit local name space is probably required. Such a name space should also be configurable to something else by the user. Discovery of a name service for access to external Internet resources is also a fundamental requirement in a multi-subnet homenet; the problem is not just name and service discovery within the homenet itself.

In some parts of the homenet, e.g. LLNs, devices may be sleeping, in which case a proxy for such nodes may be required, that can respond for example to multicast service discovery requests. Those same parts of the network may have less capacity for multicast traffic that may be flooded from other parts of the network. In general, message utilisation should be efficient considering the network technologies the service may need to operate over.

A desirable target may be a fully functional, self-configuring secure local name service so that all devices can be referred to by name, and these FQDNs are resolved locally. This could make clean use of ULAs and multiple ISP-provided prefixes much easier. Such a local name service should be (by default) authoritative for the local name space in both IPv4 and IPv6. A dual-stack residential gateway should include a dual-stack DNS server.

Current service discovery protocols are generally aimed at single subnets. If service discovery is to operate across the an entire homenet, by adopting an approach like that proposed as Extended mDNS (xmDNS) [I-D.lynn-homenet-site-mdns], then support may be required for IPv6 multicast across the scope of the whole homenet.

3.8. Other Considerations

This section discusses some other considerations for home networking that may affect the architecture.

3.8.1. Proxy or Extend?

There are two broad choices for allowing services that would otherwise be link-local to work across a homenet site. In the example of service discovery, one is to take protocols like mDNS and have them run over site multicast within the homenet. This is fine if all hosts support the extension, and the scope within any internal borders is well-understood. But it's not backwards-compatible with existing link-local protocols. The alternative is to proxy service discovery across each link, to propagate it. This is more complex, but is backwards-compatible. It would need to work with IPv6, and dual-stack.

The homenet architecture proposes that any existing protocols that are designed to only work within a subnet should be extended to work across subnets, rather than defining proxy capabilities for each of those functions. However, while it is desirable to extend protocols to site scope operation rather than providing proxy functions on subnet boundaries, the reality is that until all hosts can use site-scope discovery protocols, existing link-local protocols would need to be proxied anyway.

Some protocols already have proxy functions defined and in use, e.g. DHCPv6 relays, in which case those protocols would be expected to continue to operate that way.

3.8.2. Quality of Service

Support for QoS in a multi-service homenet may be a requirement, e.g. for a critical system (perhaps healthcare related), or for differentiation between different types of traffic (file sharing, cloud storage, live streaming, VoIP, etc). Different media types may have different such properties or capabilities.

However, homenet scenarios should require no new QoS protocols. A DiffServ [RFC2475] approach with a small number of predefined traffic classes should generally be sufficient, though at present there is little experience of QoS deployment in home networks. It is likely that QoS, or traffic prioritisation, methods will be required at the CER, and potentially around boundaries between different media types (where for example some traffic may simply not be appropriate for some media, and need to be dropped to avoid drowning the constrained media).

There may also be complementary mechanisms that could be beneficial to application performance and behaviour in the homenet domain, such as ensuring proper buffering algorithms are used as described in [Gettys11].

3.8.3. Operations and Management

The homenet should be self-organising and configuring as far as possible, and thus not be pro-actively managed by the home user. Thus protocols to manage the network are not discussed in this architecture text.

However, users may be interested in the status of their networks and devices on the network, in which case simplified monitoring mechanisms may be desirable. It may also be the case that an ISP, or a third party, might offer management of the homenet on behalf of a user, in which case management protocols would be required. How such management is done is out of scope of this document; many solutions exist.

3.9. Implementing the Architecture on IPv6

This architecture text encourages re-use of existing protocols. Thus the necessary mechanisms are largely already part of the IPv6 protocol set and common implementations. There are though some exceptions. For automatic routing, it is expected that existing routing protocols can be used as is. However, a new mechanism may be needed in order to turn a selected protocol on by default.

Some functionality, if required by the architecture, would add significant changes or require development of new protocols, e.g. support for multihoming with multiple exit routers would likely require extensions to support source and destination address based routing within the homenet.

Some protocol changes are however required in the architecture, e.g. for name resolution and service discovery, extensions to existing multicast-based name resolution protocols are needed to enable them to work across subnets, within the scope of the home network site.

Some of the hardest problems in developing solutions for home networking IPv6 architectures include discovering the right borders where the domain "home" ends and the service provider domain begins, deciding whether some of the necessary discovery mechanism extensions should affect only the network infrastructure or also hosts, and the ability to turn on routing, prefix delegation and other functions in a backwards compatible manner.

4. Conclusions

This text defines principles and requirements for a homenet architecture. The principles and requirements documented here should be observed by any future texts describing homenet protocols for routing, prefix management, security, naming or service discovery.

5. References

5.1. Normative References

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, 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.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.
- [RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.
- [RFC6204] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, April 2011.

[I-D.ietf-v6ops-6204bis]

Singh, H., Beebe, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", draft-ietf-v6ops-6204bis-09 (work in progress), May 2012.

5.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, February 2000.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.
- [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address

Assignment to End Sites", BCP 157, RFC 6177, March 2011.

- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, April 2012.
- [I-D.baker-fun-multi-router]
Baker, F., "Exploring the multi-router SOHO network",
draft-baker-fun-multi-router-00 (work in progress),
July 2011.
- [I-D.lynn-homenet-site-mdns]
Lynn, K. and D. Sturek, "Extended Multicast DNS",
draft-lynn-homenet-site-mdns-00 (work in progress),
March 2012.
- [I-D.townsley-troan-ipv6-ce-transitioning]
Townesley, M. and O. Troan, "Basic Requirements for
Customer Edge Routers - multihoming and transition",
draft-townsley-troan-ipv6-ce-transitioning-02 (work in
progress), December 2011.
- [I-D.baker-fun-routing-class]
Baker, F., "Routing a Traffic Class",
draft-baker-fun-routing-class-00 (work in progress),
July 2011.
- [I-D.howard-homenet-routing-comparison]
Howard, L., "Evaluation of Proposed Homenet Routing
Solutions", draft-howard-homenet-routing-comparison-00
(work in progress), December 2011.
- [I-D.herbst-v6ops-cpeenancements]
Herbst, T. and D. Sturek, "CPE Considerations in IPv6
Deployments", draft-herbst-v6ops-cpeenancements-00 (work
in progress), October 2010.
- [I-D.vyncke-advanced-ipv6-security]
Vyncke, E., Yourtchenko, A., and M. Townesley, "Advanced
Security for IPv6 CPE",
draft-vyncke-advanced-ipv6-security-03 (work in progress),
October 2011.

- [I-D.ietf-6man-rfc3484bis]
Thaler, D., Draves, R., Matsumoto, A., and T. Chown,
"Default Address Selection for Internet Protocol version 6
(IPv6)", draft-ietf-6man-rfc3484bis-06 (work in progress),
June 2012.
- [I-D.v6ops-multihoming-without-ipv6nat]
Troan, O., Miles, D., Matsushima, S., Okimoto, T., and D.
Wing, "IPv6 Multihoming without Network Address
Translation", draft-v6ops-multihoming-without-ipv6nat-00
(work in progress), March 2011.
- [I-D.baker-homenet-prefix-assignment]
Baker, F. and R. Droms, "IPv6 Prefix Assignment in Small
Networks", draft-baker-homenet-prefix-assignment-01 (work
in progress), March 2012.
- [I-D.arkko-homenet-prefix-assignment]
Arkko, J., Lindem, A., and B. Paterson, "Prefix Assignment
in a Home Network",
draft-arkko-homenet-prefix-assignment-02 (work in
progress), July 2012.
- [I-D.acee-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration",
draft-acee-ospf-ospfv3-autoconfig-03 (work in progress),
July 2012.
- [I-D.ietf-pcp-base]
Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P.
Selkirk, "Port Control Protocol (PCP)",
draft-ietf-pcp-base-26 (work in progress), June 2012.
- [I-D.hain-ipv6-ulac]
Hain, T., Hinden, R., and G. Huston, "Centrally Assigned
IPv6 Unicast Unique Local Address Prefixes",
draft-hain-ipv6-ulac-02 (work in progress), July 2010.
- [I-D.chakrabarti-homenet-prefix-alloc]
Nordmark, E., Chakrabarti, S., Krishnan, S., and W.
Haddad, "Simple Approach to Prefix Distribution in Basic
Home Networks", draft-chakrabarti-homenet-prefix-alloc-01
(work in progress), October 2011.
- [I-D.arkko-homenet-physical-standard]
Arkko, J. and A. Keranen, "Minimum Requirements for
Physical Layout of Home Networks",
draft-arkko-homenet-physical-standard-00 (work in

progress), March 2012.

[Gettys11]

Gettys, J., "Bufferbloat: Dark Buffers in the Internet", March 2011, <<http://www.ietf.org/proceedings/80/slides/tsvarea-1.pdf>>.

[IGD-2]

UPnP Gateway Committee, "Internet Gateway Device (IGD) V 2.0", September 2010, <<http://upnp.org/specs/gw/UPnP-gw-WANIPConnection-v2-Service.pdf>>.

Appendix A. Acknowledgments

The authors would like to thank Aamer Akhter, Mark Andrews, Dmitry Anipko, Fred Baker, Ray Bellis, Cameron Byrne, Brian Carpenter, Stuart Cheshire, Lorenzo Colitti, Robert Cragie, Ralph Droms, Lars Eggert, Jim Gettys, olafur Gudmundsson, Wassim Haddad, Joel M. Halpern, David Harrington, Lee Howard, Ray Hunter, Joel Jaeggli, Heather Kirksey, Ted Lemon, Kerry Lynn, Erik Nordmark, Michael Richardson, Barbara Stark, Sander Steffann, Dave Taht, Dave Thaler, Mark Townsley, JP Vasseur, Curtis Villamizar, Dan Wing, Russ White, and James Woodyatt for their contributions within homenet WG meetings and on the WG mailing list.

Appendix B. Changes

This section will be removed in the final version of the text.

B.1. Version 04

Changes made include:

- o Moved border section from IPv6 differences to principles section.
- o Restructured principles into areas.
- o Added summary of naming and service discovery discussion from WG list.

B.2. Version 03

Changes made include:

- o Various improvements to the readability.

- o Removed bullet lists of requirements, as requested by chair.
- o Noted 6204bis has replaced advanced-cpe draft.
- o Clarified the topology examples are just that.
- o Emphasised we are not targetting walled gardens, but they should not be precluded.
- o Also changed text about requiring support for walled gardens.
- o Noted that avoiding falling foul of ingress filtering when multihomed is desirable.
- o Improved text about realms, detecting borders and policies at borders.
- o Stated this text makes no recommendation about default security model.
- o Added some text about failure modes for users plugging things arbitrarily.
- o Expanded naming and service discovery text.
- o Added more text about ULAs.
- o Removed reference to version 1 on chair feedback.
- o Stated that NPTv6 adds architectural cost but is not a homenet matter if deployed at the CER. This text only considers the internal homenet.
- o Noted multihoming is supported.
- o Noted routers may not be separate devices, they may be embedded in devices.
- o Clarified simple and advanced security some more, and RFC 4864 and 6092.
- o Stated that there should be just one secret key, if any are used at all.
- o For multihoming, support multiple CERs but note that routing to the correct CER to avoid ISP filtering may not be optimal within the homenet.

- o Added some ISPs renumber due to privacy laws.
- o Removed extra repeated references to Simple Security.
- o Removed some solution creep on RIOS/RAs.
- o Load-balancing scenario added as to be supported.

B.3. Version 02

Changes made include:

- o Made the IPv6 implications section briefer.
- o Changed Network Models section to describe properties of the homenet with illustrative examples, rather than implying the number of models was fixed to the six shown in 01.
- o Text to state multihoming support focused on single CER model. Multiple CER support is desirable, but not required.
- o Stated that NPTv6 not supported.
- o Added considerations section for operations and management.
- o Added bullet point principles/requirements to Section 3.4.
- o Changed IPv6 solutions must not adversely affect IPv4 to should not.
- o End-to-end section expanded to talk about "Simple Security" and borders.
- o Extended text on naming and service discovery.
- o Added reference to RFC 2775, RFC 6177.
- o Added reference to the new xmDNS draft.
- o Added naming/SD requirements from Ralph Droms.

Authors' Addresses

Tim Chown (editor)
University of Southampton
Highfield
Southampton, Hampshire SO17 1BJ
United Kingdom

Email: tjc@ecs.soton.ac.uk

Jari Arkko
Ericsson
Jorvas 02420
Finland

Email: jari.arkko@piuha.net

Anders Brandt
Sigma Designs
Emdrupvej 26A, 1
Copenhagen DK-2100
Denmark

Email: abr@sdesigns.dk

Ole Troan
Cisco Systems, Inc.
Drammensveien 145A
Oslo N-0212
Norway

Email: ot@cisco.com

Jason Weil
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
USA

Email: jason.weil@twcable.com

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 30, 2013

E. Kline
Google Inc.
July 29, 2012

Default Perimeter Identification
draft-kline-default-perimeter-00

Abstract

Automatic, simple identification of when traffic is crossing a perimeter is highly desirable for a variety of home network uses. This document proposes a set of default tests to be applied to traffic scheduled for forwarding, which can be used collectively to identify this perimeter in some (but not all) environments.

Status of this Memo

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

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

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

This Internet-Draft will expire on January 30, 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
1.1. Requirements Language	3
2. Terminology	3
3. Fundamental recommendations	4
3.1. Network node default security practices	4
3.2. State changes and logging	5
4. Useful perimeter signals	5
4.1. Product-defined interface purposes	6
4.2. Routing adjacency	6
4.3. Links requiring subscriber information	6
4.4. Links requiring existing network-layer connectivity	7
4.5. Links that are fundamentally point-to-point in nature	7
5. IP over Ethernet	7
5.1. DHCPv6 PD, if and only if...	8
5.2. Other tricks?	8
6. Additional considerations	8
6.1. Physical vs virtual interfaces	8
6.2. Mixed zone next-hops on the same interface	9
6.3. Perimeter and protocol version	9
7. Acknowledgements	9
8. IANA Considerations	9
9. Security Considerations	10
10. References	10
10.1. Normative References	10
10.2. Informative References	10
Appendix A. Additional Stuff	10
Author's Address	10

1. Introduction

Automatic, simple identification of when traffic is crossing a perimeter is highly desirable for a variety of home network uses. This document proposes a set of default tests to be applied to traffic scheduled for forwarding, which can be used collectively to identify this perimeter.

Of note are limitations introduced by the ubiquitous use of IP over Ethernet (IPoE) Internet access methods. By design these architectures make it difficult (at best) to distinguish any difference between a LAN port in an enterprise and a home Internet connection.

Nevertheless, where practical, an automated mechanism of perimeter discovery permits home devices to define default definitions of the "interior", i.e. the home, and the "exterior", usually the greater Internet. Once identified, a device could apply default security policies to traffic transiting the perimeter.

Specifying the default policy that should be applied to traffic crossing this perimeter is out of scope of this document. Implementors should remain mindful of recommended practices, e.g. RFC 4864 [RFC4864], et cetera.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Terminology

In order to address perimeter identification at a manageable scale the scope has been limited to discussing concepts of "interior", "exterior", and "perimeter". Working definitions in use by this document are as follows.

Interior

The interior is broadly defined to include the collection of interfaces (physical or virtual), nodes, and forwarding next-hops collectively under the control of a single logical administrative domain.

Exterior

The exterior is broadly defined to include all interfaces (physical or virtual), nodes, and forwarding next-hops collectively NOT under the control of any single logical administrative domain and specifically NOT under the control of the administrative domain which defines the interior.

Perimeter

The perimeter is taken to be the collection of all ephemeral demarcations within the collection of interior nodes which forward traffic such that any IP packet transiting that demarcation can be said to be crossing either from the interior toward the exterior or from the exterior toward the interior. This is independent of whether or not such traffic is permitted by policy to complete its transiting from one zone to the other.

Expressly not discussed herein are architectures having multiple interior networks, nor the relationships between them as separate from their relationship to their common exterior or any common perimeter. The architectures under discussion have a single interior, single exterior, and a single logical perimeter between them.

The definition of perimeter is such that, for example, an IP packet arriving on a NAT device's interior interface that is "hairpinned" and retransmitted out the interior interface is not considered to have touched nor crossed the perimeter.

The definition of perimeter as written technically permits traffic being forwarded over an interface to be classified as transiting a perimeter or not based on the classification of the next-hop. The implications of this are discussed in a later section.

3. Fundamental recommendations

The application of security policies at the perimeter and possible relaxation of security policies within the interior are apt to have administrative consequences. Some fundamental recommendations for nodes operating in this environment follow.

3.1. Network node default security practices

By default all network nodes SHOULD follow best current security practices. Any node may at times find itself in a hostile environment. This is obviously true of mobile nodes when, for instance, connecting to a variety of public "Wi-Fi" networks. In

such environments mobile nodes cannot be sure that there is any network device acting in the mobile node's own best security interests with respect to others on the local network.

It is equally true of more traditionally "fixed" nodes: any compromised neighbor nodes ("fixed" or mobile) may be used as a conduit for further compromise. Indeed, one study of a captured "botnet" ([TORPIG], section 5.2.4) found that roughly 78.9% of infected hosts had RFC 1918 [RFC1918] addresses, commonly used in IPv4 NAT deployments.

3.2. State changes and logging

Devices conforming to this and other homenet documents MUST continuously evaluate the interior, exterior, and perimeter classifications of interfaces and traffic. These may change at any time, for example if new devices are added into the network or a power event causes all equipment to reset, and specific ordering of device bring-up within a homenet network MAY NOT be assumed.

Nodes compliant with this and other homenet documents SHOULD administratively log the perimeter classification of interfaces (both physical and virtual), the reason(s) for such classification, and times at which such classifications are made or changed.

4. Useful perimeter signals

This is a non-trivial task as it is tantamount to automated discovery of administrative boundaries.

Many architectures fundamentally make it difficult or impossible to detect administrative boundaries and rely on various mechanisms of user or administrator invention to create or modify such boundaries. Other hints about administrative boundaries may be insecure, unreliable, operationally impractical, or may place arbitrary requirements upon the architectural where previously no such requirement existed.

Nevertheless there are some signals that may be useful. Which signals are available and useful vary with the access architecture, and in some cases there may be virtually no reliable information to securely determine a perimeter. An physically or cryptographically authenticated routing protocol may be the highest fidelity signal for determining the interior, and thereby the exterior and perimeter.

4.1. Product-defined interface purposes

Many products come with interfaces labeled with their intended use. Examples include home routers with RJ-45 ports labeled "WAN" and "LAN", or perhaps with symbols indicating "The Internet" and "inside the home". Other examples include wireless routers with defined separate WLAN and "Guest" ESSIDs. In such cases where interior and exterior are clearly delineated a homenet device SHOULD by default consider traffic forwarded between interfaces of differing regions as traversing a boundary.

4.2. Routing adjacency

Some networks may employ a physically or cryptographically secure routing protocol. Within such networks, traffic received from and scheduled to be forwarded to next-hops with whom an adjacency has been formed SHOULD by default be classified as interior and not considered to be transiting a perimeter.

Similarly, traffic forwarded to or received from next-hops with whom no adjacency has been formed SHOULD by default be classified as exterior. A next-hop with whom no adjacency has been formed but which nevertheless constitutes the next-hop for a learned or configured route SHOULD by default be considered exterior.

If (and only if?) an interface has only interior next-hops then traffic originating from nodes on links on that interface SHOULD by default be considered to be interior... Discuss: [two routers each sharing a hub with two upstreams on it]. HELP: need more thought to clarify forwarding traffic to on-link destinations versus to next-hops for further forwarding. Don't want to accidentally classify interior on-link nodes as exterior because no adjacency is formed.

HELP: much more thought is required here. What about bring-up order? What about an interior node attempting and failing to start an authenticated adjacency: it's a problem if the interface flips into exterior classification.

HELP: find language that doesn't, for example, define interior to be the entire Internet when RPKI is in use.

4.3. Links requiring subscriber information

One obvious administrative boundary is a link that requires subscriber credentials in order for that link to successfully forward and receive general traffic. Examples include authenticated PPPoE sessions, 3G/LTE PDP contexts (requiring a SIM card associated with a customer account), and authenticated VPN links. By default, all

traffic traversing such a link SHOULD be considered to be traversing a perimeter.

4.4. Links requiring existing network-layer connectivity

By default, all traffic traversing any interface that encapsulates (decapsulates) its payload in a layer higher than or equal to the network (IP) layer in order to forward (receive) traffic SHOULD be considered to be traversing a perimeter. Examples of such interfaces include: Teredo, 6to4, 6rd, 4rd, PPTP and L2TP tunnels, et cetera.

In cases where the exact layer of encapsulation is not necessarily clearly defined or agreed upon, e.g. MPLS interfaces, traffic traversing such interfaces SHOULD also by default be considered to be traversing a perimeter.

In the absence of default perimeter classification, such links would provide a mechanism to breach an otherwise existing perimeter and generally complicate the definition and discovery of the interior. In cases where such interfaces are desired to be classified as part of the interior, and traffic traversing them also classified as interior traffic, another means MAY use to re-classify accordingly.

4.5. Links that are fundamentally point-to-point in nature

Most home networking technology supports more than two nodes on the same logical link communicating directly. By default, traffic traversing from such a "shared access" link which is classified as interior to one which is fundamentally point-to-point in nature (e.g. PPPoE, PPPoA, or some other future link type) SHOULD be considered as transiting a perimeter.

Additionally, traffic transiting a homenet device from such a "shared access" link which is classified as interior to one on which no on-link neighbor discovery and/or communication is permitted by configuration of the node itself, e.g. an 802.11 SSID on which the node acting as an infrastructural access point forbids direct neighbor communications, SHOULD be considered as crossing a perimeter.

HELP: what does this mean for 6lowpan networks inside the home?

5. IP over Ethernet

The ubiquity of IPoE undoubtedly greatly simplifies network architectures and node requirements for connecting to such networks. However, it can be difficult at best for a homenet device to

determine if it is fully in the interior of a network or part of the perimeter.

5.1. DHCPv6 PD, if and only if...

DHCPv6 PD is at this time the most common method for supplying "SoHo" networks with a routable prefix block. If (and only if) a means of distributing prefixes among interior routers is devised that does NOT use DHCPv6 Prefix Delegation, then a link on which DHCPv6 PD succeeded SHOULD be considered an administrative boundary and traffic traversing this interface SHOULD be considered to be traversing a perimeter.

If DHCPv6-PD is to be used within the interior then this signal is not useful.

5.2. Other tricks?

(TBD) If you're DHCPv6-setting-up-the-reverse-DNS then that interface SHOULD be considered part of the perimeter.

(TBD) If DHCPv4'ing an RFC 1918,6598,... address? What other prefixes and updates will be required as we run out of IPv4?

6. Additional considerations

Everything herein needs more thought and work.

6.1. Physical vs virtual interfaces

In certain configurations it may be desirable that the perimeter defined on a virtual interface also be extended to include the physical interface(s) over which such traffic is forwarded/received. For example, consider a router configured with a PPPoE virtual interface on a physical 802.3 interface. In such a configuration, the security policy applied to traffic transiting the PPPoE interface should most likely also be applied to non-PPPoE traffic transitting the physical interface. If not, the "interior" region would otherwise be logically extended to include the upstream access link. As there is no guaranteed of administrative boundary XXX, a default configuration SHOULD consider the physical interface a perimeter.

By way of contrast, consider an entirely interior router which also has a VPN interface, through which traffic may be passed to, say, a resident's company network. While a VPN virtual interface SHOULD be considered a logical demarcation point, the physical interface through which VPN-encapsulated traffic is transmitted need not

necessarily be classified as such. Instead, what may be desired is that traffic to/from interfaces that are interior to this VPN-enabled router may pass through either the VPN interface or any "upstream" interfaces, but traffic originating from "upstream" interfaces may be default DENIED transit through the VPN interface. While this situation may be far from the norm for networks, it nevertheless affords the maintenance of a simple mental model of a hierarchical network.

To address these uses, by default the physical interface(s) through which a virtual interface's traffic is forward should also be considered a perimeter, unless other means determine that it is in fact an interior interface. Note that in all cases such a virtual interface is considered a perimeter.

Discuss: applying the policy to the entire interface for some "infrastructural" connections (e.g. PPPoE). Virtual link versus physical link, virtual interface versus physical interface.

6.2. Mixed zone next-hops on the same interface

By default, if one then all. TBD: explain.

6.3. Perimeter and protocol version

In cases where one IP version's perimeter might be determined to differ in some way from another IP version's identified perimeter, the potential for confusion and misconfiguration, and therefore security risk, increases. In the interests of simplicity, and in keeping with the principle of least surprise, traffic transiting links or forwarding to (received from) next-hops which would transit a perimeter in one protocol version SHOULD be considered as transiting a perimeter for traffic of all protocol versions.

7. Acknowledgements

The author gratefully acknowledges the constructive input and criticism of Lorenzo Colitti, Mark Townsley, and Ole Troan.

Thanks also must go to pleasant, peaceful and productive trips on the Japan Rail (JR) Shinkansen ("bullet train").

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

All drafts are required to have a security considerations section. See RFC 3552 [RFC3552] for a guide.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2. Informative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.

[TORPIG] Stone-Gross, B., "Your Botnet is My Botnet: Analysis of a Botnet Takeover", 2009, <<http://www.cs.ucsb.edu/~seclab/projects/torpig/torpig.pdf>>.

Appendix A. Additional Stuff

This becomes an Appendix.

Author's Address

Erik Kline
Google Inc.
Roppongi 6-10-1, 26th Floor
Minato, Tokyo 106-6126
JP

Phone: +81 03 6384 9000
Email: ek@google.com

HOMENET
Internet-Draft
Intended status: Standards Track
Expires: December 31, 2012

W. Cloetens
SoftAtHome
P. Lemordant
D. Migault (Ed)
Francetelecom - Orange
July 2012

IPv6 Home Network Front End Naming Delegation
draft-mglt-homenet-front-end-naming-delegation-00.txt

Abstract

This document proposes a Naming Delegation Architecture that makes possible End Users to reach the hosts or services of their Home Network using Names instead of IP addresses.

This document shows how the Naming Delegation between the CPE and the ISP can be set so the CPE is not exposed on the Internet. This document describes an Naming Architecture where ISPs provide Front End Delegating DNS Servers whereas the CPEs constitute a Back End Network of Delegated DNS Servers. All DNS queries for any Home Network are addressed to the Delegating Front End Server. The response is expected to be stored on a CPE, and the Front End Delegating DNS Server sends a DNS Query to that CPE before answering to the initial DNS query.

The negotiation between the CPE and the ISP is using DHCP Options. This document provides options so Front End Delegating and the Delegated DNS Servers configure their respective Zone files and so that CPEs restrict access and protect themselves from unauthorized DNS Queries.

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 December 31, 2012.

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. Requirements notation	2
2. Introduction	2
3. Terminology	4
4. Front End Naming Delegation Architecture Overview	4
4.1. Home Network Naming Architecture Requirements	4
4.2. Front End Naming Delegation Architecture Description	6
4.3. Front End Naming Delegation Configuration	6
4.4. Difference between the Front End Delegating DNS Server and traditional DNS Recursive DNS Server	8
4.5. How the Front End Configuration impacts the CPE	9
5. Protocol Exchange	10
5.1. CPE Request Creation and Transmission for the Front End Naming Delegation Architecture	10
5.2. ISP DHCP Server Responding to the CPE Request for the Front End Naming Delegation Architecture	10
5.3. CPE Receiving the ISP DHCP Response for the Front End Naming Delegation Architecture	11
6. DHCP Options	11
6.1. Delegated DNS Architecture Option	11
6.2. Front End Delegating Information Option	12
6.3. Delegating Authorized Resolvers Option	12
7. IANA Considerations	13
8. Security Considerations	13
9. Acknowledgment	13
10. References	13
10.1. Normative References	14
10.2. Informational References	14
Authors' Addresses	14

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 [RFC2119].

2. Introduction

[I-D.mglt-homenet-naming-delegation] describes the Naming Delegation Architecture that makes possible Services and Objects of a Home Network to be globally reachable with Names on the Internet. For that purpose, the Customer Premise Equipment (CPE) hosts the authoritative DNS Server of the Home Network. The zone associated to the Home Network ("my-homenet") is a subzone of a zone managed by the ISP ("example."). This zone is attached to the global DNS Architecture. Because the ISP delegates the Naming service to the CPE, we call the DNS server responsible for "example." the Delegating DNS Server, and the DNS server responsible for "my-homenet.example." the Delegated DNS Server. The Delegated DNS Server runs on the CPE, and [I-D.mglt-homenet-naming-delegation] describes how the CPE can automatically set the Naming Delegation between the Delegated and the Delegating DNS Server. Necessary pieces of information to configure the respective DNS Zones are exchanged between the DHCP client of the CPE and the ISP DHCP Server through DHCP Options.

The resulting Naming Delegation Architecture [I-D.mglt-homenet-naming-delegation] results in a CPE hosting a Service on the Internet. CPEs have not been designed for heavy load, and, as a result, the Delegating exposes the Home Network to potential Deny of Service attacks. The Front End Naming Delegation Architecture proposed in this document is an alternative to the Naming Delegation Architecture [I-D.mglt-homenet-naming-delegation] where the ISP provides Front End Delegating Servers that handles the whole DNS traffic. The CPE remains responsible for the zone "my-homenet.example.", but only responds to DNS queries sent by the Front End Delegating Servers. For this reason we call the CPE Delegated DNS Server the Back End Delegated DNS Server.

The Front End Naming Delegation Architecture can be seen as providing a Authoritative DNS Server for all the Home Networks: the Front End Delegating DNS Server. However this Authoritative Server distributes the Zone between multiple nodes (the CPE). The CPE constitutes the Back End Network. The Front End Delegating DNS Server receives DNS query from the Internet, and to respond requires to retrieve this information on the CPE hosting this information. In this document, the Front End Delegating DNS Server uses the DNS protocol to retrieve this information from the CPE. Other protocols may have been chosen.

The Front End Naming Delegation Architecture is based on the Naming Delegation Architecture [I-D.mglt-homenet-naming-delegation] and addresses the same requirements. It addresses the Deny of Service Security issue. On the other hand, it requires the ISP to provide an adapted infrastructure, and that all DNS traffic is (partly) handled by the ISP. The document shows how the CPE can be configured

automatically and be part of the Front End Naming Delegation Architecture.

In this document we only considered IPv6 and DHCP. As such DHCP MUST be understood as DHCPv6. We also assume the reader has read [I-D.mglt-homenet-naming-delegation]

3. Terminology

This document uses the terminology defined in [I-D.mglt-homenet-naming-delegation], and introduces the following terminology:

- Front End Delegating DNS Server or Delegating DNS Server: The DNS Server of the ISP that handles with the DNS queries addressed to the Home Network.
- Back End Delegated DNS Server or Delegated DNS Server: CPE are hosting a DNS Service
- Front End Delegating Information: Information like FQDNs and IP addresses of the Front End Delegating DNS Servers. These pieces of information are provided from the ISP DHCP Server to the CPE so it can properly configure its DNS zone file.
- Delegating Authorized Resolvers: The hosts that are authorized to send DNS queries to the CPE. These Resolvers can be the Front End Delegating DNS Servers, but we keep these functions independent since some ISP may use dedicated Interfaces for the Front End Delegating DNS Server and for the Delegating Authorized Resolvers.

4. Front End Naming Delegation Architecture Overview

4.1. Home Network Naming Architecture Requirements

The Home Network Naming Requirements for the Naming Delegation listed in [I-D.mglt-homenet-naming-delegation] are:

- 1: Centralized Naming Configuration: The CPE is responsible to bind Names and IP addresses for the whole Home Network.
- 2: Automatic Configuration: The CPE MUST be able to set the Naming architecture when plugged, with minimum configuration from the End User.
- 3: Advanced Configuration enable: The CPE enables advanced specific configurations.
- 4: Privacy Protection By Design: The Names and the Home Network IP address plan is administrated by the CPE and are not communicated to the ISP. This prevents the ISP to be aware of the hosts, Services and Objects that compose the Home Network.
- 5: Make the Home Network Naming Architecture Scalable: The Naming Architecture MUST be scalable and designed to handle a large increase of Objects, Services and hosts in each Home Networks.

The Naming Delegation Architecture fulfills these requirements, and we consider this architecture as the base architecture. However, this architecture major drawback is that the CPE hosts the Delegated DNS Server. CPE are usually not designed to handle heavy traffic, and thus are sensitive to DoS attacks. The Front End Naming Delegation Architecture adds one requirement to the currently designed Naming Delegation Architecture [I-D.mgmt-homenet-naming-delegation]:

- 6: ISP Infrastructure MUST protect the Naming Architecture: The CPE MUST NOT expose the Home Network Naming service to DoS attacks. The ISP MUST be able to provide the necessary infrastructure that handle DoS attacks, or heavy loads.

In order to match Requirement 6, the Front End Naming delegation Architecture introduces Front End DNS Delegating Server that handles with all DNS traffic. This means that all DNS queries that concern the Home Network are addressed to the Front End DNS Delegating Server of the ISP and are not addressed to the CPE. CPEs belong to the Back End DNS Network.

The Front End DNS Naming Delegation Architecture fulfills all the above Requirements. However, Requirement 4 needs to be balanced against Requirement 6. Requirement 6 requires the ISP to handle all DNS queries that concern the Home Network. This makes the ISP aware of all queried Services, Objects and hosts in the Home Network. This may, in that sense, reduces the Privacy of the Home Network compared to the Naming Delegation Architecture. In fact with the Naming Delegation Architecture, the DNS query is directly sent to the CPE when the DNS client has the IP address of the CPE in its cache. In that case, the ISP is not aware of the existence of the queried FQDN. However, if the DNS client does not have the IP address of the CPE, then the DNS query is sent first to the ISP Delegating Server. In this latter case, the Front End DNS Naming Delegation Architecture does not provide less privacy.

4.2. Front End Naming Delegation Architecture Description

Figure 1 shows how the Resolution is performed. In [1], the Resolver sends a DNS query to the Front End Delegated Server for the host "hots1.my-homenet.example.". The Front End Delegated Server does not have the response in its cache or in its zone file. The Front End Delegating DNS Server MUST send a query to the Back End Delegated DNS Server. The IP address of the Back End Delegated DNS Server MUST NOT be revealed to the Resolver, for example by setting the NS field in the DNS Zone File. In Figure 1, we mentioned the Delegated Server Information Database where this IP address is stored. The Front End Delegating Server sends the DNS(SEC) query to the Back End Delegated Server hosting the zone "my-homenet.example.". The source IP address used is one of the Delegating Authorized Resolvers IP addresses. This query is represented in [2]. The Back End Delegated Server responds in [3] with the DNS(SEC) Response. Note that the "AUTHORITY" and "ADDITIONAL SECTION" of the DNS response MUST indicate the FQDN and the IP addresses of the Front End Delegated DNS Server. These pieces of information have been provided by the ISP DHCP Server with the Front End Delegating Information DHCP Option. The CPE can also be configured to respond without these fields. Finally in [4], the Front End Delegating Server forwards the DNS(SEC) response to the Resolver. "AUTHORITY" and "ADDITIONAL SECTION" fields MUST be filled in appropriately.

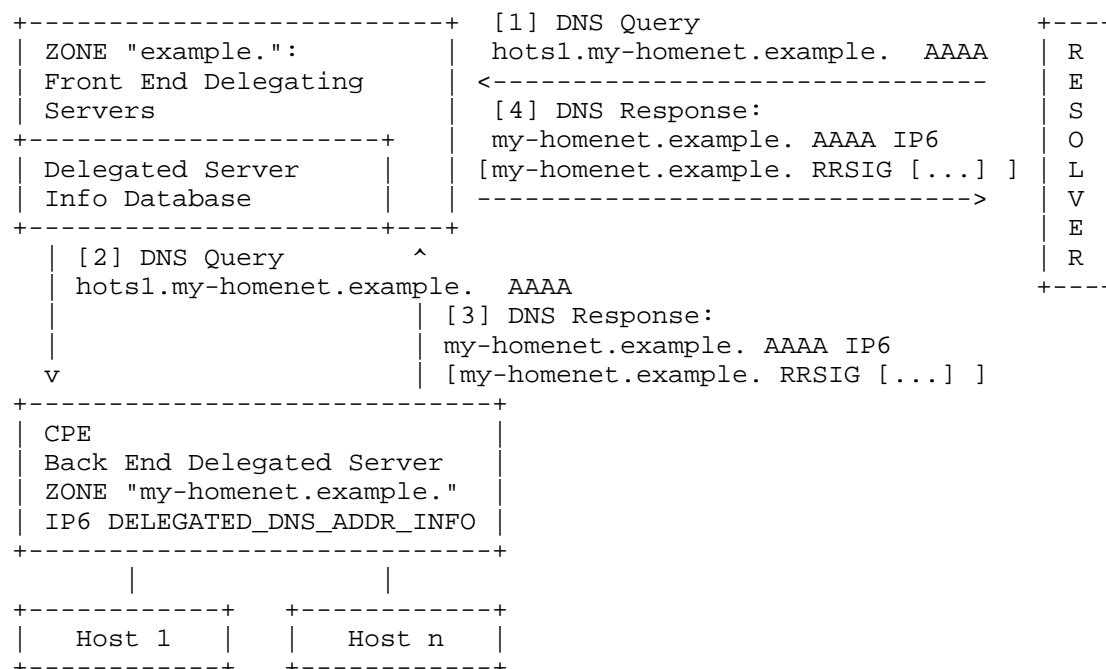


Figure 1: DNS Resolution with the Home Network Front End Naming Delegation Architecture

4.3. Front End Naming Delegation Configuration

Figure 2 describes the Interactions between the CPE and the ISP DHCP Server.

Similarly to [I-D.mglt-homenet-naming-delegation], the CPE hosts a DHCP Server (DHCP_SRV) that is used to assign IP addresses and FQDNs to the Hosts of the Home Network. In this document we considered DHCP, but other protocols can also be used in combination with DHCP or instead of DHCP. The CPE also has a DHCP Client (DHCP_CLT) that is used to exchange information with the ISP DHCP Server. This document describes how these exchanges properly configure the Front End Naming Delegation Architecture. The CPE also hosts a Authoritative DNS Server (DNS_SRV) that is responsible of the subzone associated to the Home Network. This Authoritative DNS Server is called the Back End Delegated Server. At last the CPE also has a Firewall (FIREWALL), that can be configured with security Policies. In this document, the CPE is not expected to received DNS queries from any other peer but the Front End Delegation DNS Servers, that are in the ISP Network.

In Figure 2. the CPE sends a DHCP Request for a Front End Naming Delegation Architecture (DELEGATED_DNS_ARCHITECTURE). Similarly to the Naming Delegation Architecture, the CPE provides the necessary information so the ISP can derive the IP address of the Back End Delegated DNS Server (DELEGATED_DNS_ADDR_INFO). If the CPE wants a DNSSEC Delegation to be set it also provides the Delegation of Signing Information (DS). In our case, the CPE also sends a request for a Prefix Delegation (IA_PD).

To the difference with [I-D.mglt-homenet-naming-delegation], the IP address of the Back End Delegated DNS Server is not mentioned in the Zone file of the Front End Delegating DNS Server. In this document, the Back End Delegated DNS Server is not expected to receive any DNS query from anyone but the Front End Delegating DNS Server. For DNS Resolvers, the only Authoritative DNS Server they are aware of is the Front End Delegating DNS Server.

Similarly to [I-D.mglt-homenet-naming-delegation], the ISP DHCP Server provides the CPE the IP Prefix so the CPE can configure its Prefix Delegation. To set the DNS(SEC) Naming Delegation the ISP DHCP Server indicates the type of Naming Delegation Architecture agreed between the CPE and the ISP DHCP Server (DELEGATED_DNS_ARCHITECTURE). In addition, the ISP DHCP Server, provides the Delegated Domain (DELEGATED_DOMAIN) as well as the IP addresses and FQDNs of the Front End Delegating DNS Servers (FRONT_END_DELEGATING_INFO). These pieces of information are necessary to configure the zone file of the Home Network. In fact the zone file MUST be configured with the Front End Delegating Server as the authoritative servers. In addition, the ISP DHCP Server may also provide the IP addresses or subnet prefix of the Delegating Authorized Resolvers (DELEGATING_AUTH_RESOLVERS). These Resolvers are the only hosts supposed to send DNS queries to the CPE. DNS queries from any other IP address MUST be discarded.

Upon receiving these pieces of information, the Front End Delegating Server and the Back End Delegated Server configure their Zones. In addition the CPE also configures its Firewall, so to discard any DNS queries but those emitted from the Delegating Authorized Resolvers.

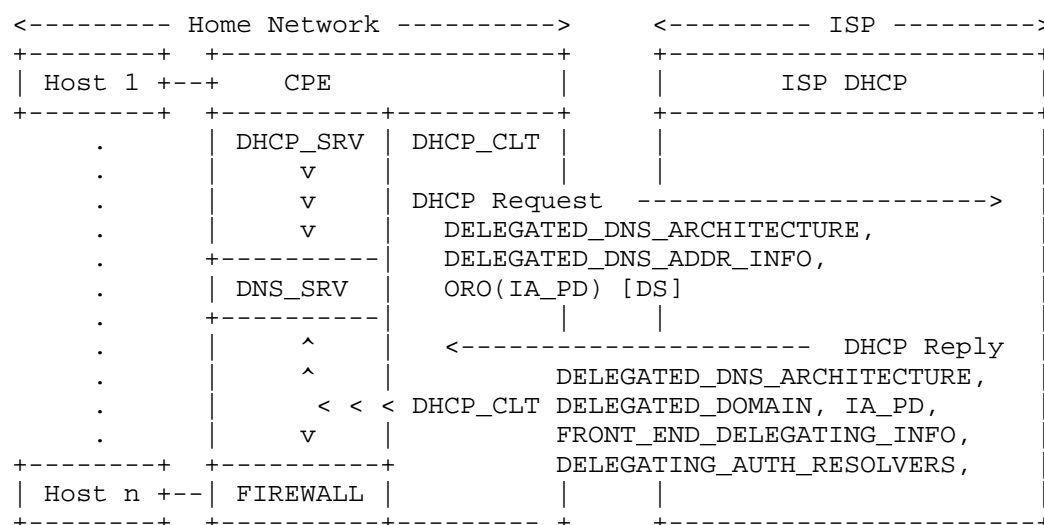


Figure 2: Front End Naming Delegation Architecture

4.4. Difference between the Front End Delegating DNS Server and traditional DNS Recursive DNS Server

From Figure 1, one may assimilate the Front End Delegating DNS Server to a Recursive DNS Resolver. The main differences are:

- 1. The Front End Delegating DNS Server only proceeds to Resolution for the FQDNs that are hosted in one of the Back End Delegated DNS Servers.

- 2. The Back End Delegated DNS Servers are not Public DNS. More especially, the Delegated DNS Server may have a public IP address, but the DNS Service is not provided for any Resolver but the authorized Resolvers.

As a result, the Front End Delegating DNS Server is a mixed mode between Authoritative and Recursive DNS Server. As an Authoritative Server, the Response [4] in figure 1 MUST have a Authoritative Answer (AA) bit set, which indicates the Response is from an Authoritative Server. Then the Resolution [2] and [3] in figure 1 MUST be processed even if the Recursion Desired (RD) bit is not set in the DNS query [1].

It is also recommended that the Front End Delegating DNS Server provides the Authoritative and Additional Section of the Response in [4], without considering the sections of [3]. In other word, it is recommended not to forward these section from [3], and the CPE should be configured not to provide these sections in [3].

4.5. How the Front End Configuration impacts the CPE

Figure 2 shows that the ISP DHCP Server provides the IP addresses of the Front End Delegating DNS Server as well as the Name of the Front End Delegating DNS Server. These are the information the Back End Delegated DNS Server MUST put in its Zone file. More especially in the NS fields.

Figure 2 also shows that the ISP DHCP Server provides the CPE the IP addresses or subnet prefix of the Authorized Delegating Resolvers. These are the IP addresses authorized to send DNS queries that should not be discarded on the WAN Interface. Any other DNS query on the WAN should be discarded. These rules are set by the Firewall as represented in Figure 2.

The Firewall rules does not prevent the CPE to be a DNS forwarder or a DNS Resolver for the hosts of the Home Network. In fact the CPE can still receive DNS queries from the LAN Interface. The issue is that the CPE may provide Multiple DNS Services. In this document, we consider the CPE provides at least a DNS Authoritative servers on its WAN Interface for the Authorized Delegating Resolvers. For the LAN Interface, the CPE may be configured in various ways, depending on the ISP DNS Infrastructure. A first configuration consists in configuring the CPE LAN DNS Service into a DNS forwarder. In that case, the CPE DHCP server of the Home Network provides an IP address of the CPE for the DNS Resolver. DNS queries for the Home Network are answered by the CPE, others are forwarded to the Resolver of the ISP. This resolver is provided via DHCP. Another alternative consists in configuring the CPE as a Recursive DNS Server. Without any specific configurations, DNS queries for the Home Network are sent to the Front End Delegating DNS Server. Optimization may be done to bypass the Front End Delegating DNS Server for the Home Network Zone and are CPE or software implementation specific.

5. Protocol Exchange

5.1. CPE Request Creation and Transmission for the Front End Naming Delegation Architecture

When the CPE wants to set a Front End Naming Delegation Architecture, it requests this set up to the ISP DHCP Server. For that purpose, we consider two new naming-delegation-action:

SET_FRONT_END_NAMING_DELEGATION_WITH_DNS when the delegation is only performed with DNS or SET_FRONT_END_NAMING_DELEGATION_WITH_DNSSEC if the CPE wants a DNSSEC delegation. These naming-delegation-actions are proposed in the Delegated DNS Architecture DHCP Option (OPTION_DELEGATED_DNS_ARCHITECTURE). Then, the CPE proceeds as described in [I-D.mglt-homenet-naming-delegation].

5.2. ISP DHCP Server Responding to the CPE Request for the Front End Naming Delegation Architecture

When the DHCP Server receives a Delegated DNS Architecture DHCP Option (OPTION_DELEGATED_DNS_ARCHITECTURE), Delegated DNS Address Information DHCP Option (OPTION_DELEGATED_DNS_ADDR_INFO) or a Delegation of Signing DHCP Option (OPTION_DS), the DHCP Server proceeds as described in [I-D.mglt-homenet-naming-delegation].

In addition, when the naming-delegation-action is set to SET_FRONT_END_NAMING_DELEGATION_WITH_DNS or SET_FRONT_END_NAMING_DELEGATION_WITH_DNSSEC, the DHCP Server MUST include in the Response the two additional DHCP Options. The Front End Delegating Information DHCP Option (OPTION_FRONT_END_DELEGATING_INFO) which indicates the FQDNs of the Front End Delegating Servers and their associated IP addresses. Then, it also MUST include the Delegating Authorized Resolvers DHCP Option (OPTION_DELEGATING_AUTH_RESOLVERS) which indicates the IP addresses or subnet prefixes of the Authorized Delegating Resolvers.

Note that Naming Delegation is set differently for the Front End Naming Delegation Architecture and for the Naming Delegation Architecture. More specifically, in the Front End Naming Delegation,

the ISP DHCP Server MUST NOT make the IP address of the Delegated DNS Server public in its zone file.

5.3. CPE Receiving the ISP DHCP Response for the Front End Naming Delegation Architecture

Similarly to [I-D.mglt-homenet-naming-delegation], if the CPE has not received all expected DHCP Options, or cannot proceed to the configuration of the Naming Delegation Architecture, it MUST either clear the Naming Delegation settings or proceed to the appropriated settings.

When the CPE receives the Delegating Authorized Resolvers DHCP Option (OPTION_DELEGATING_AUTH_RESOLVERS), the CPE may update its Firewall rules. The Front End Delegating Information DHCP Option (OPTION_FRONT_END_DELEGATING_INFO) is used to configure the DNS zone of the Home Network.

The CPE may receive the Delegating Authorized Resolvers or the Front End Delegating Information DHCP Option from the ISP DHCP Server that are not the response to a Delegated DNS Architecture DHCP Option. This may happen if the ISP DHCP Server is updating or modifying its Front End Delegating DNS Server or the associated Delegating Authorized Resolvers. In that case, the CPE MUST make sure the message provides from the ISP DHCP Server and updates its Firewall rules as well as its DNS zone file.

6. DHCP Options

The options detailed in this section are

- Delegated DNS Architecture: (OPTION_DELEGATED_DNS_ARCHITECTURE) is used by the DHCP Client on the CPE to inform how the Naming Delegation Architecture should be configured. In return, it is used by the ISP DHCP Server to report the Status Code.
- Front End Delegating Information DHCP Option: (OPTION_FRONT_END_DELEGATING_INFO) is used by the ISP DHCP Server to provide the CPE the FQDN and IP addresses of the Authoritative DNS Server of the Home Network Zone file. These Authoritative DNS Servers are the Front End DNS Server.
- Delegating Authorized Resolvers DHCP Option: (OPTION_DELEGATING_AUTH_RESOLVERS) is used by the DHCP Server to provide the CPE the IP addresses or subnet prefixes of the Delegating Authorized Resolvers. These are the resolvers authorized to send DNS(SEC) queries.

6.1. Delegated DNS Architecture Option

The Delegated DNS Architecture DHCP Option is defined in [I-D.mglt-homenet-naming-delegation]. This document adds two new naming-delegation-actions defined below:

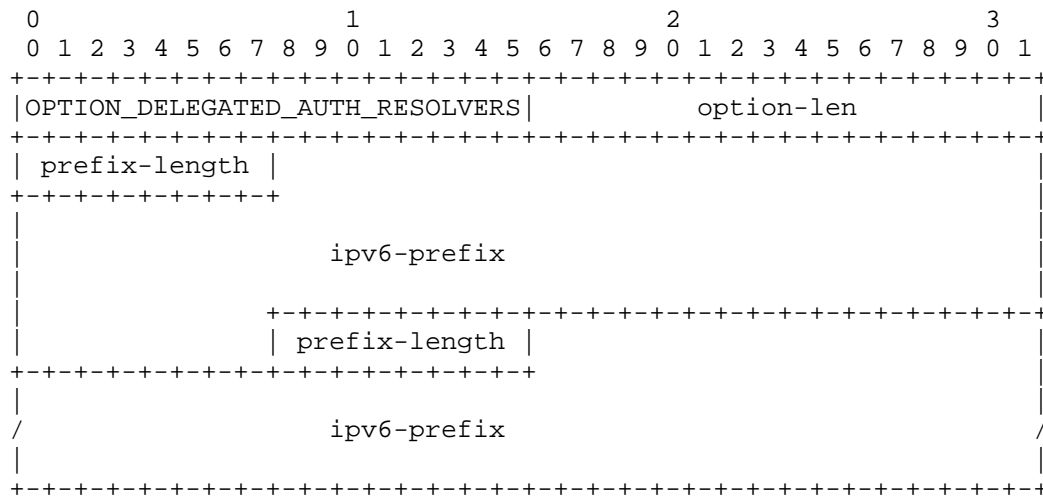
- SET_FRONT_END_NAMING_DELEGATION_WITH_DNS - 2 - : Indicates that the DHCP Server MUST set the Front End Naming Delegation Architecture with only DNS, and MUST NOT consider DNSSEC Delegation.
- SET_FRONT_END_NAMING_DELEGATION_WITH_DNSSEC - 3 - : Indicates that the DHCP Server MUST set the Front End Naming Delegation Architecture with DNSSEC.

6.2. Front End Delegating Information Option



- option-code: OPT_FRONT_END_DELEGATING_INFO (16 bits)
- option-len: Length (16 bits) of the Front End Delegating Information Option in octets.
- front-end-length: Length (16 bits) of the Front End Delegating Server.
- front-end-fqdn-length: Length (16 bits) of the Front End Delegating Server FQDN.
- ipv6-address: IPv6 Address (128 bits).

6.3. Delegating Authorized Resolvers Option



- option-code: OPTION_DELEGATED_AUTH_RESOLVERS (16 bits)
- option-len: Length (16 bits) of the Delegating Authorized Resolvers Option in octets.
- prefix-length: Length (8 bits) for this prefix in bits.
- ipv6-prefix: IPv6 address or IPv6 prefix used by the authoritative DNS server to send DNS queries to the delegated domain name.

7. IANA Considerations

This document adds two new DHCP Options:

- OPTION_FRONT_END_DELEGATING_INFO: TBD
- OPTION_DELEGATING_AUTH_RESOLVERS: TBD

8. Security Considerations

This document addresses the DoS security issue of [I-D.mglt-homenet-naming-delegation]. Other security considerations remains as described in [I-D.mglt-homenet-naming-delegation].

9. Acknowledgment

The authors wish to thank Ole Troan for pointing out issues with the IPv6 routed home concept and placing the scope of this document in a wider picture, Mark Townsley for encouragement and injecting a healthy debate on the merits of the idea, Ulrik de Bie for providing alternative solutions, Paul Mockapetris for pointing out issues of the trustworthiness of a reverse lookup, and Christian Jacquenet for seeing the value from a Service Provider point of view.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

10.2. Informational References

[I-D.mglt-homenet-naming-delegation]
Cloetens, W., Lemordant, P. and D. Migault, "IPv6 Home Network Naming Delegation Architecture", Internet-Draft draft-mglt-homenet-naming-delegation-00, July 2012.

Authors' Addresses

Wouter Cloetens
SoftAtHome
vaartdijk 3 701
3018 Wijkmaal
Belgium

Email: wouter.cloetens@softathome.com

Philippe Lemordant
Francetelecom - Orange
2 avenue Pierre Marzin
22300 Lannion
France

Phone: +33 2 96 05 35 11
Email: philippe.lemordant@orange.com

Daniel Migault
Francetelecom - Orange
38 rue du General Leclerc
92794 Issy-les-Moulineaux Cedex 9
France

Phone: +33 1 45 29 60 52
Email: mglt.ietf@gmail.com

HOMENET
Internet-Draft
Intended status: Standards Track
Expires: January 31, 2013

W. Cloetens
SoftAtHome
P. Lemordant
D. Migault (Ed)
Francetelecom - Orange
July 30, 2012

IPv6 Home Network Naming Delegation Architecture
draft-mglt-homenet-naming-delegation-00.txt

Abstract

This document describes the Naming Delegation Architecture that makes IPv6 Home Network globally reachable with Names or Fully Qualified Domain Names (FQDN). In this architecture, the Customer Premise Equipment (CPE) acts as the DNS Authoritative Server of the Home Network also called the Delegated DNS Server. The Naming Delegation is configured between the Delegated DNS Server and the Delegating DNS Server managed by the ISP.

The use case considered in this document is an End User that subscribes its ISP a specific Delegated Domain for its Home Network. This document describes how the CPE automatically sets the Naming Delegation between the Delegating and Delegated DNS Server.

The Naming Delegation is requested by the CPE. The CPE DHCP Client and the ISP DHCP Server exchange DHCP Options to properly set the Naming Delegation. More specifically, the CPE DHCP Client (resp. the ISP DHCP Server) configures the DNS(SEC) Zones of the Delegated DNS Server (resp. Delegating DNS Server). For the Delegating DNS Server, the necessary pieces of information required to set the Naming Delegation are the IP address of the Delegated DNS Server, and if DNSSEC is used, the Delegation of Signing Information. For the Delegated DNS Server, the necessary information is the Delegated Domain associated to the Home Network.

Status of this Memo

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

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

Internet-Drafts are draft documents valid for a maximum of six months

and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 31, 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. Requirements notation	4
2. Introduction	4
3. Terminology	6
4. Home Network Naming Architecture Requirements	7
5. Home Network Delegating Architecture Overview	8
5.1. Fulfilling Home Network Naming Architecture Requirements	8
5.2. Naming Delegation Architecture Description	9
5.3. Naming Delegation Configuration Environment Description	11
5.4. Naming Delegation DHCP Configuration Description	13
6. Protocol Exchange	15
6.1. CPE Request Creation and Transmission for Naming Delegation Architecture	15
6.2. ISP DHCP Server Responding to the CPE Request for Naming Delegation Architecture	16
6.2.1. Case 1: No Delegated DNS Architecture DHCP Option in conjunction with Delegated Address Information or Delegated Domain DHCP Option	16
6.2.2. Case 2: No Delegated DNS Architecture DHCP Option in conjunction with Option Request DHCP Option for a Delegated Domain DHCP Option	16
6.2.3. Case 3: Delegated DNS Architecture DHCP Option	16
6.2.4. Processing the Delegated DNS Address Information DHCP Option	19
6.2.5. Processing the Delegation of Signing DHCP Option	19
6.3. CPE Receiving the ISP DHCP Response for the Naming Delegation Architecture	19
7. DHCP Options	19
7.1. Delegated DNS Architecture Option	20
7.2. Delegated Domain Option	22
7.3. Delegated DNS Address Information Option	23
7.4. Delegated Delegation of Signing Option	23
8. IANA Considerations	24
9. Security Considerations	24
9.1. Names are less secured than IP addresses	24
9.2. Names are less volatile than IP address	25
9.3. DNSSEC is recommended to authenticate DNS hosted data	25
9.4. Channel between the CPE and ISP DHCP Server MUST be secured	26
9.5. CPEs are sensitive to DoS	26
10. Acknowledgment	26
11. References	27
11.1. Normative References	27
11.2. Informational References	27
Authors' Addresses	28

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 [RFC2119].

2. Introduction

Home Networks used to be composed of a single or a set of PCs connected to a CPE to access the Internet. Now they have evolved to a large set of applications and objects or devices managed by the CPE. Among these applications are Media applications like Video, Music and Photos Stations, Backup applications, File sharing applications with FTP and Web Stations, Access applications with VPN Stations, and others like Surveillance Station, Printing Stations. With the Internet of Things (IoT) the number of objects attached to the CPE is expected to increase in the coming years.

Then, services and objects in the Home Networks should be made reachable from anywhere on the Internet. IPv6 removes the need for NAT and makes this possible with a global reachability. But IPv6 addresses remain inconvenient. In fact, most End Users prefer using Names to access these services. Furthermore Names make communications independent from IP renumbering, or changes of IP addresses. Then, if IP addresses plan remains opaque for End Users, on the other hand, they easily understand the Naming hierarchical model. More specifically, if "my-homenet" is the Delegated Domain associated to my Home Network, it makes sense that "my-service.my-homenet" is the "my-service" in "my-homenet".

To assign Names to objects and services of the Home Network, the Home Network should be provided a Naming Architecture. For most End Users, the CPE manages the Home Network, that is to say, it provides access to the Internet, discovers the devices, and interconnects them between each other. As a result, the CPE is the natural device to centralize the Naming service of the Home Network.

Home Networks should be operational with the least configuration. End Users, expect to subscribe to an ISP, plug with minimum configuration the CPE and access to the Internet and to their services from anywhere on the Internet. The CPE interconnects the Home Network to the ISP's Network, and the CPE gets from the ISP all the necessary pieces of information to set up the connectivity. In some cases, the CPE is even provided by the ISP. In order to make services and objects of the Home Network reachable with Names, the ISP is likely to provide the CPE the Delegated Domain associated to the Home Network, and set up the necessary delegation to make the

Home Network DNS Zone reachable from the Internet. More specifically, the End User subscribes its ISP an Internet connectivity, and registered its Home Network Delegated Domain "my-homenet". When the CPE is plugged, as it requests an IP prefix, it also requests the Delegated Domain - like "my-homenet.example.". From then, all devices requesting IP addresses via DHCP or using alternative protocols are registered by the CPE in the zone "my-homenet.example.". When a communication is initiated with "a-device.my-homenet.example.", a DNS query is sent to the ISP authoritative DNS server of the zone "example.". This server is called the Delegating DNS Server and delegates the query to the CPE which acts as the authoritative server of "my-homenet.example." and sends back the response.

This architecture is called the "Home Network Naming Delegation Architecture" because, the ISP is not hosting the DNS zone of the Home Network but is delegating the Home Network zone to the CPE. There are multiple motivations for this delegation architecture. First delegation preserves the Home Network privacy, by avoiding ISPs to know the Home Network hosts. Furthermore, ISP are unlikely to be able to scale their Naming infrastructure for all services and devices of the Home Networks. As a result, ISPs are looking to distribute the Naming service between the CPEs, and delegate to each CPE their associated Home Network zone.

The purpose of this document is to describe an architecture that automatically configures the Naming architecture of the Home Network. More specifically, when the End User plugs its CPE, the CPE is being assigned by the ISP a Delegated Domain that has been pre-registered by the End User to the ISP. This Delegated Domain designates the Home Network, and the CPE is expected to act as an authoritative DNS server of this Zone. When a node of the Home Network is requesting using DHCP an IP address, the CPE can provide the node the IP address and updates the zone file of the Home Network.

This document assumes that the communication between the CPE and the ISP DHCP Server is protected. This document does not specify which mechanism should be used. [RFC3315] proposes a DHCP authentication and message exchange protection, [RFC4301], [RFC5996] proposes to secure the channel at the IP layer.

This document does not provide any mechanism that protects the CPE from being exposed on the Internet. In fact, CPE are low power devices, and the Naming Delegation described in this document exposes the CPE on the Internet by publishing its IP address and making the DNS Service hosted on the CPE. This issue is addressed in [I-D.mglt-homenet-front-end-naming-delegation] which describes the Front End Naming Delegation Architecture. In this architecture, the

ISP's infrastructure protects the CPE from heavy load.

This document only deals with IPv6 IP addresses and DHCPv6 [RFC3315]. When we mention DHCP, it MUST be understood as DHCPv6.

3. Terminology

This sections defines terminology specific to IPv6 and DHCP used in this document.

- Home Network: Designates the objects and Services that are hosted in the Home Network of the End User.
- Home Network Naming Architecture: Designates the Architecture that makes possible to reach a device, an object or a service in the Home Network by using Names like Fully Qualified Domain Names.
- Home Network Naming Delegation Architecture or Naming Delegation Architecture: Designates the Naming Architecture Described in this document. The ISP delegates the Naming management of the Home Network to the Delegated DNS Servers. Consistency with the Global Naming Architecture is provided by the ISP. The Delegation occurs between Delegating DNS Servers hosted by the ISP and Delegated DNS Servers hosted in the Home Network.
- Internet Service Provider (ISP): The End User has subscribed to the ISP. The ISP is aware of End User credential and the Delegated Domain of the Home Network. The ISP is expected to provide the CPE the required information to properly configure the DNS Zone.
- Delegating DNS Server: Designates the Authoritative DNS Server of the ISP. The Home Network is a subzone of the Delegating DNS Server. This subzone is handled by the Delegated DNS Server.
- Customer Premise Equipment (CPE): Designates the device that hosts the DNS and DHCP Service in the Home Network. This device sets the IP and Naming interconnection between the ISP Network and Home Network.
- Delegated DNS Server: Designates the DNS Authoritative Server that handles the Hosts of the Home Network.

- Delegated Delegation of Signing Option: Designates the DHCP Option that makes possible the DNSSEC Delegation between the Delegated DNS Server and the Delegating DNS Server.
- Delegated DNS Addressing Information Option: Designates the DHCP Option that makes possible the Delegation between the Delegated DNS Server and the Delegating DNS Server for both DNS and DNSSEC. With this option, the Delegating DNS Server is informed of the IP addressing information - the interface and the subnet identifier - used by the Delegated DNS Server.
- Delegated Domain: Designates the domain Name associated to the Home Network. In this document, the Delegated Domain is reserved by the End User to the ISP at the subscription of the Internet Access. It is then communicated to the CPE by the ISP, so the CPE configures properly its Delegated DNS Server.
- Fully Qualified Domain Name (FQDN): Name that fits the general DNS requirements.

4. Home Network Naming Architecture Requirements

The Home Network Naming Architecture is defined by two parties the End User and the ISP. Both of them have specific requirements.

The End User requirements we are considering are the following:

- 1: Centralized Naming Configuration: Configuring a Network, is most of the time more convenient when done in a centralized way. Home Networks now may have only a few nodes, which makes a per-node configuration possible, for example using DynDNS like service, to assign a FQDN to each node. However, the number of nodes is expected to grow in the next future, and we recommend now to specify a centralized way for configuring the Home Network Naming Architecture.
- 2: Automatic Configuration: Most End User do not want to configure, their Home Network, and configuration MUST be minimal. The procedure should consider those 90% of End Users
- 3: Advanced Configuration enable: Some End Users have various specific requirements, and they SHOULD be able to match these requirements. This means that the Automatic Configuration may be disable.

- 4: Privacy Protection By Design: Most End User does not want to provide anyone, including their ISP, the content of their zone, like network topology, or the devices and services hosted in the Home Network. On the other hand the content of the zone should be publicly published. DNS makes this possible for two reasons. First, DNS makes the content of the zone public, without publishing the whole zone - at least AXFR queries must be disabled. Then, DNS is a distributed databases with delegation mechanisms, that preserves the privacy of subzones toward upper zones. Note that as explained in Section 9 the Naming Delegation Architecture described in this document protects the End User's privacy by not providing the complete DNS zone. However, one MUST be aware that using Names exposes their Home Networks to the Internet since names are expected to provide less randomness than the standard IPv6 numbering. Then Names are more associated to an identity than IP addresses are. Thus, allowing PTR DNS queries may also affect the End User's privacy.

The ISP requirements, other than fulfilling the End Users' requirements are the following:

- 1: Make the Home Network Naming Architecture Scalable: ISPs can hardly foresee the evolution of Home Networks, that is to say the number of devices that will belong to them, or the number of requests, updates associated to each FQDN. Architectures that would make the ISP deal with all FQDNs is definitively out of scope. Delegation management of the Zone to CPE makes local management handled locally, and Delegating the zone makes CPE dealing with their zone traffic.

5. Home Network Delegating Architecture Overview

5.1. Fulfilling Home Network Naming Architecture Requirements

The CPE is designed to provide connectivity to the Home Network, to discover and connect all Hosts of the Home Network. As such, it is a good candidate to bind FQDNs and IP addresses. In this document, we consider the CPE as the device that centralizes the configuration of the Delegation Home Network Naming Architecture. This fulfills the End User Requirement 1.

The CPE should not be configured, and should get the necessary information to properly configure the Delegation Home Network Naming Architecture. These pieces of information, like the Delegated Domain assigned to the Home Network are provided by the ISP. On the other hand, the CPE may also be able to provide information to the ISP.

For example, the CPE may provide the ISP the Delegated DNS IP Address Information, that is to say the Interface and Subnet Identifier of the Home Network Authoritative DNS, or the Delegated Delegation of Signing which is the hash of public key of the Home Network Authoritative DNS server. In this document, we call the Home Network Authoritative DNS server the Delegated DNS Server. These pieces of information are device related and local information. They are not related to the configuration of the Delegation Home Network Naming Architecture. This fulfills the End User Requirement 2.

The CPE should set the Naming Delegation Architecture by requesting for it. The CPE can be configured to not request these pieces of information so the Home Network can have a specific Naming configuration. A specific Naming configuration could be for example, that the FQDN assigned to the Home Network is different from the one attributed by the ISP. This fulfills the End User Requirement 3.

The CPE acts as an authoritative DNS server for the Home Network. This prevents communication of the DNS zone to any third party. As a result, this makes the DNS zone publicly available, while protecting the privacy of the Home Network. This fulfills the End User Requirement 4.

The CPE provides the Home Network Authoritative DNS server or Delegated DNS Server. This function is an added function to the service/device discovery, routing service, DHCP service, Naming resolution service, provided by the CPE. The CPE seems to be the most adapted device, for most End Users cases, to host the Delegated DNS Server. This service includes handling with the DNS queries concerning the Home Network and updating the zone for the various devices. The load generated by the Delegated DNS Server is expected to be handled by the CPE, and CPE may be designed to handle such traffic. On the other hand, it is hardly possible ISPs can handle with this traffic for all Home Networks. The Delegation Home Network Naming Architecture is adopted for its scalability. This fulfills the ISP Requirement 1.

5.2. Naming Delegation Architecture Description

Figure 1 describes a DNS resolution with the Naming Delegation Architecture. The resolution can be done using DNS or DNSSEC. In the Architecture described in figure 1, the IPv6 address MUST be global.

In the example below, the Zone of the ISP is called "example.". The End User of the CPE has registered to the ISP the Delegated Domain "my-homenet", and the Home Network can be globally reachable under the name "my-homenet.example.". A host in the Home Network "host1"

has been assigned an IPv6, and has been registered in the Home Network with the name "host1.my-homenet.example.". Note that the architecture makes host1 globally reachable under the name "host1.my-homenet.example.".

The End User is likely to use alternate names which will require the use of DNAME [RFC6672] and CNAME [RFC2118] . In other words, the Naming Delegation Architecture described in this document does not prevent the End User to register a service or a host under an alternative name such as "host1-alternative-name.example.net". For that purpose, the End User may redirect manually "host1-alternative-name.example.net" to "host1.my-homenet.example." using CNAME [RFC2118]. Similarly, the Home Network can also be registered under an alternate domain name such as "my-alternate-homenet.net". Redirecting the zone requires to use DNAME. In both case, the configuration is performed by the End User, and is independent to the configuration between the ISP and the End User.

In figure 1, the Resolver is getting the IP address of "host1.my-homenet.example.". A DNS(SEC) Query is sent to the Delegating DNS Server responsible of "example.". Then "example." responds with the delegating information, so the resolver can send the DNS Query to the Delegated DNS Server responsible of "my-homenet.example.". The delegating pieces of information are, the Name and IP address of the Delegated DNS Server, and if DNSSEC is available and requested the Delegation of Signing. These pieces of information may have been provided by the Delegated DNS Address Information and Delegated Delegation of Signing DHCP Options.

Then, the Resolver sends the DNS(SEC) Query to the Home Network Delegated DNS Server which responds with the requested DNS(SEC) information.

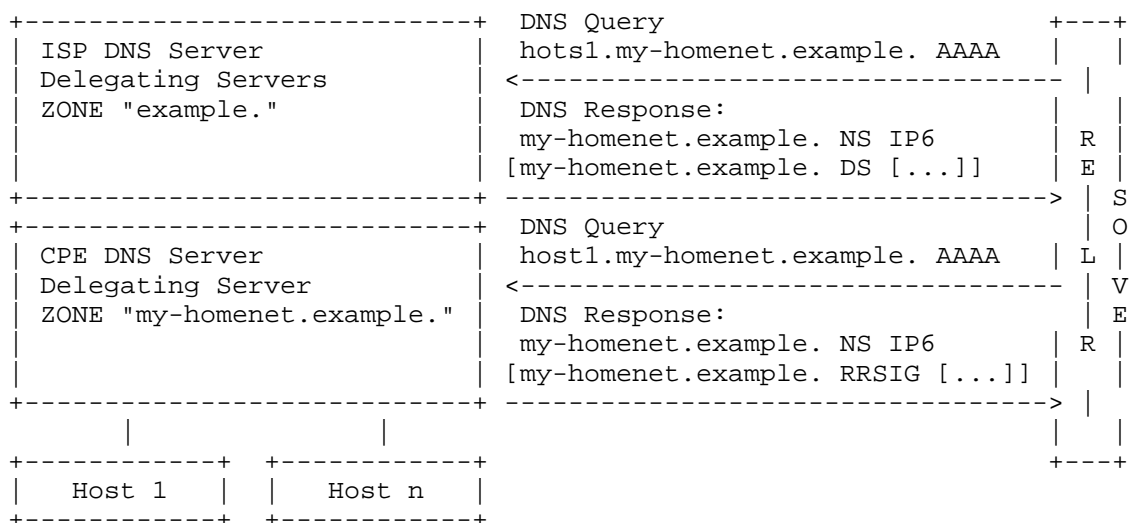


Figure 1: DNS Resolution with the Home Network Delegating Architecture

5.3. Naming Delegation Configuration Environment Description

Figure 2 shows the DHCP exchange between the CPE and the ISP DHCP Server. This exchange sets the Home Network Naming Delegation Architecture.

As mentioned in figure 2, the CPE is in the Home Network and implements three functions: the DHCP Client (DHCP_CLT), the DHCP Server (DHCP_SRV) and the Delegated DNS Server (DNS_SRV).

- CPE DHCP Client (DHCP_CLT): is responsible for getting parameters from the ISP. In figure 2, the CPE DHCP Client requests the ISP an IPv6 Prefix Delegation (IA_PD) [RFC3633]. The CPE DHCP Client also requests to set a Naming Delegation Architecture (DELEGATED_DNS_ARCHITECTURE), and provides the necessary pieces of information to set up the Naming Delegation Architecture (DELEGATED_DNS_ADDR_INFO, DELEGATED_DNSSEC_DS). In return, the CPE DHCP Client (DHCP_CLT) is expected to receive from the ISP DHCP Server, the Delegated Domain Name (DELEGATED_DOMAIN) and the IPv6 Prefix Delegation (IA_PD). These pieces of information are useful to configure the Home Network DNS Zone file, of the CPE Delegated DNS Server (DNS_SRV).
- CPE DHCP Server (DHCP_SRV): The CPE DHCP server hosted by the CPE is not mandatory for the Naming Delegation Architecture. We mentioned it in Figure 2 as most of the CPEs are responsible for assigning IPv6 Addresses to the Hosts of the Home Network.

Figure 2 considers that the IPv6 Address of the Hosts are assigned via DHCP, and that while assigning the IPv6 prefixes, the DHCP Server populates the Home Network DNS Zone file of the CPE Delegated DNS Server (DNS_SRV).

- CPE Delegated DNS Server (DNS_SRV): The CPE Delegated DNS Server hosts the Naming Service of the Home Network. The DNS Server can implement DNS or DNSSEC. This function interacts with the CPE DHCP Client (DHCP_CLT) so the Naming Delegation is properly set with the ISP, and the CPE DHCP Server (DHCP_SRV) which manages names for the hosts of the Home Network.

The ISP DHCP Server is in the ISP Network and is the counter part of the CPE DHCP Client (DHCP_CLT). As the CPE DHCP Client (DHCP_CLT) interacts with the Delegated DNS Server, the ISP DHCP Server also interact with the ISP Delegating DNS Server. In fact the ISP DHCP Server is in charge of setting the Naming Delegation upon request of the CPE DHCP Client (DHCP_CLT). Furthermore, when the Home Network Prefix Delegation is not any more active, the ISP DHCP Server MUST remove the Naming Delegation settings.

Hosts are the devices of the Home Network. Figure 2, illustrates the case, where these hosts have been assigned an IPv6 prefix from the DHCP Server of the CPE. We use the "stateful address autoconfiguration protocol", as defined in [RFC3315] but other protocols like "IPv6 Stateless Address Autoconfiguration" [RFC4862] may also be used. This will not affect the Naming Delegation Architecture.

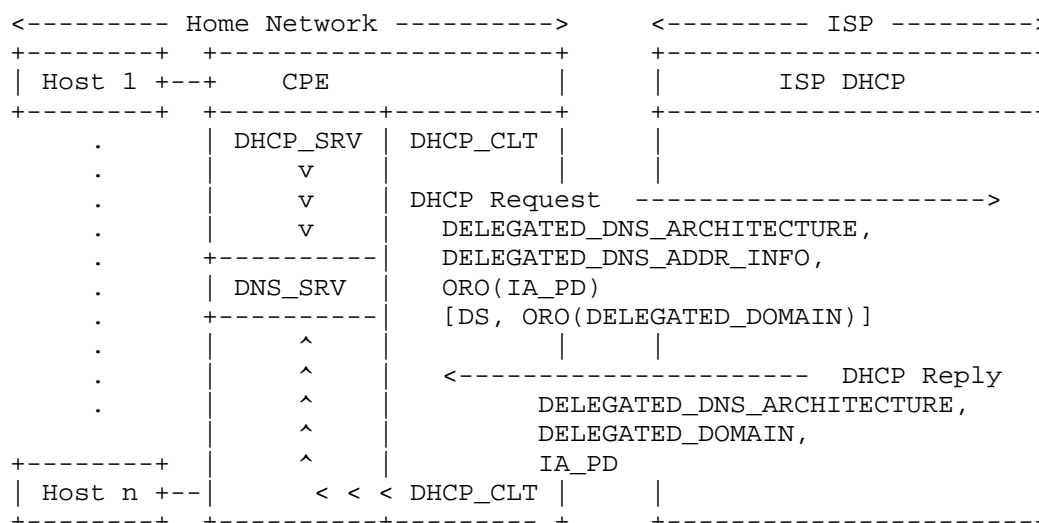


Figure 2: Naming Delegation Architecture

5.4. Naming Delegation DHCP Configuration Description

Figure 2 illustrates how the CPE provides and get the necessary information to set the Naming Delegation. In this document, all parameters are provided and received using DHCP Options.

First of all, in order to set the Home Network Naming Delegation, the CPE MUST have a Delegated Prefix. In our case, the CPE is requesting the Delegated Prefix to the ISP DHCP Server with the Identity Association Prefix Delegation DHCP Option (IA_PD), as defined in [RFC3633], [RFC3769]. To Request the Option from the ISP DHCP Server, the CPE uses the Option Request DHCP Option (ORO) [RFC3315].

The CPE uses the Delegated DNS Architecture DHCP Option (OPTION_DELEGATED_DNS_ARCHITECTURE) to specify the naming-delegation-action to perform. The CPE provides a ordered list of alternative naming-delegation-actions. One of these actions will be chosen by the ISP DHCP Server. The naming-delegation-actions considered in this document are Clear the Naming Delegation Settings, Set it with DNS or Set it with DNSSEC. Figure 2 illustrates the case where the CPE Sets the Naming Delegation Architecture with DNS or with DNSSEC.

In order to set the Naming Delegation Architecture between the Delegating DNS Server and the Delegated DNS Server, the CPE MUST provide some pieces of information. First the Delegating DNS Server MUST be aware of the IP address used for the Delegated DNS Server.

Since the CPE is requesting a Prefix Delegation, it is not aware of the IP address. That is why, the CPE MUST provide pieces of information that enables the ISP DHCP Server to derive the IP address. In fact the CPE provides the Subnet Identifier and the Interface Identifier using the Delegated Address Information DHCP Option (OPTION_DELEGATED_DNS_ADDR_INFO). The ISP DHCP Server is aware of the assigned prefix, and thus can derive the IP address of the Delegated DNS Server.

The calculation of the CPE IPv6 address used for the delegated DNS server is done as follows:

```

0                                     63|64                                     127
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| IPv6 prefix | subnet-ID | interface-ID |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
subnet-ID length = 64 - IPv6 prefix length

```

Figure 3: CPE IP address Format

If DNSSEC is used, the CPE MUST also provide the Delegation of Signing (DS) Information [RFC4034]. This is done using the Delegation of Signing DHCP Option (OPTION_DS)

In figure 2, we mentioned the Delegated Domain DHCP Option that can optionally be requested. In fact, with Delegated DNS Architecture DHCP Option requesting the ISP to Set the Naming Delegation Architecture, the ISP is expected to send back the Delegated Domain. However, in some cases, for example if the CPE wants to check the ISP has provisioned a Delegated Domain, the CPE may request the Delegated Domain without setting the Naming Delegation Architecture. In that case, the CPE, MUST request the Delegated Domain DHCP Option (OPTION_DELEGATED_DOMAIN).

The ISP DHCP Server processes the various DHCP Options, and provides the Prefix Delegation, the Delegated DNS Architecture, and the Delegated Domain DHCP Options. The Prefix Delegation Option provides the IPv6 Prefix assigned to the Home Network. The Delegated DNS Architecture DHCP Option indicates the Naming Delegation set by the ISP, as well as Status Code. The Delegated Domain DHCP Option provides the Domain the owner of the CPE has registered.

The ISP DHCP Server MUST keep the Naming Delegation Architecture coherent with the Prefix Delegation. If the Prefix Delegation is using DHCP, then, the ISP DHCP Server MUST unset the Naming Delegation Architecture when the Prefix Delegation expires. How the DHCP Server should proceed is out of scope of this document.

6. Protocol Exchange

In this document, we do not consider the CPE and the ISP have pre-agreed on some parameters. In other words, all necessary information for configuring the Home Network Naming Delegation Architecture are sent via DHCP Options. The ISP is in charge of identifying the CPE owner - that is to say the End User - and is aware of the Delegated Domain the End User has subscribed for.

For clarity, we designated the CPE DHCP Client by the CPE.

6.1. CPE Request Creation and Transmission for Naming Delegation Architecture

The CPE provides the ISP DHCP Server an ordered list of naming-delegation-actions which starts with the most preferred action. The ISP DHCP Server can choose one of these actions and process it. These naming-delegation-actions are carried by the Delegated DNS Architecture DHCP Option (OPTION_DELEGATED_DNS_ARCHITECTURE). If the CPE wants to remove the Naming Delegation Architecture, it sets the action to CLEAR. Otherwise, it sets the action to SET_NAMING_DELEGATION_WITH_DNS or SET_NAMING_DELEGATION_WITH_DNSSEC.

The Naming Delegation cannot be set if the CPE has not been provided a Prefix Delegation. So, if the CPE has not been assigned a Prefix, it MUST either get first a prefix before setting the Naming Delegation Architecture. If the Prefix Delegation is provided via the ISP DHCP Server, then the CPE can simultaneously send a DHCP Request for a Prefix Delegation with the Identity Association Prefix Delegation DHCP Option and for setting the Naming Delegation Architecture.

If SET_NAMING_DELEGATION_WITH_DNS or SET_NAMING_DELEGATION_WITH_DNSSEC is one of the naming-delegation-action carried by the Delegated DNS Architecture DHCP Option, then the CPE MUST provide the Delegated Address Information DHCP Option (OPTION_DELEGATED_DNS_ADDR_INFO).

If SET_NAMING_DELEGATION_WITH_DNSSEC is one of the naming-delegation-action carried by the Delegated DNS Architecture DHCP Option, then the CPE MUST provide the Delegation of Signing DHCP Option (OPTION_DS).

If the CPE does not want to set the Naming Delegation Architecture, but wants to know the Delegated Domain, then, the CPE MUST send a Delegated Domain DHCP Option (OPTION_DELEGATED_DOMAIN) with no Delegated DNS Architecture DHCP Option (OPTION_DELEGATED_DNS_ARCHITECTURE).

6.2. ISP DHCP Server Responding to the CPE Request for Naming Delegation Architecture

6.2.1. Case 1: No Delegated DNS Architecture DHCP Option in conjunction with Delegated Address Information or Delegated Domain DHCP Option

When the DHCP Server receives a Delegated Address Information DHCP Option or a Delegated Domain DHCP Option it MUST check if there is a Delegated DNS Architecture DHCP Option. If not, these DHCP Options MUST be discarded.

6.2.2. Case 2: No Delegated DNS Architecture DHCP Option in conjunction with Option Request DHCP Option for a Delegated Domain DHCP Option

If the DHCP Server receives an Option Request DHCP Option for a Delegated Domain DHCP Option, but no Delegated DNS Architecture DHCP Option. The DHCP Server MUST NOT proceed to any configuration settings. The ISP DHCP Server returns the Delegated Domain DHCP Option. Otherwise, it MUST return a Delegated DNS Architecture DHCP Option with a single action set to NONE and the Status Code indicating the reason of failure.

Possible failure reasons are: If the DHCP Server understands the Delegated Domain DHCP Option but does not provide the Naming Delegation Service, the DHCP Server MUST return a Status Code set to NamingDelegationUnavailable. Then, if the Naming Delegation Service is Available, the DHCP MUST check if the CPE has been identified or authenticated according to local policies. If that is not the case, the DHCP Server MUST return a Status Code set to UnauthorizedRequester. If the CPE is authorized to request a Delegated Domain DHCP Option, the DHCP Server MUST check the Delegated Domain has been provisioned, and if that is not the case, it MUST send a Status Code set to UnprovisionedDelegatedDomain. For any other failure, the DHCP Server MUST send a Status Code UnspecFail.

In case of success the DHCP Server does not return Delegated DNS Architecture DHCP Option or Status Code.

6.2.3. Case 3: Delegated DNS Architecture DHCP Option

When a Delegated DNS Architecture DHCP Option is received, the DHCP Server MUST check an Option Request for Identity Association Prefix Delegation (IA_PD) has not been provided. If that is the case, the DHCP Server MUST proceed first to this Option. Then, the Delegated DNS Architecture DHCP Option should only be processed, if the

Identity Association Prefix Delegation has been processed successfully. If no Identity Association Prefix Delegation has been requested the DHCP Server may consider the CPE has no Prefix and send a Delegated DNS Architecture DHCP Option with the status code MissingPrefixDelegationRequest. On the other hand, the DHCP Server may also assume the CPE got a Prefix from another way and proceeds to the Delegated DNS Architecture DHCP Option.

When a Delegated DNS Architecture DHCP Option is received and the Naming Delegation is already set. If the naming-delegation-action is set to NONE, the packet do not proceed to any change. For all other naming-delegation-action, the ISP DHCP Server MUST process the DHCP Option. In case of success, the Naming Delegation MUST be updated. In any other case, the ISP DHCP Server MUST clear the Naming Delegation settings.

From now, the DHCP processes the Delegated DNS Architecture DHCP Option. Preliminary checks are performed in case of failure, the DHCP Server sends a Delegated DNS Architecture DHCP Option with a single naming-delegation-action set to NONE and the Status Code indicating the reason of failure. If the DHCP Server understands this Option, but does not provide the Naming Delegation Service, the DHCP Server MUST return a Status Code set to NamingDelegationUnavailable. Then the DHCP MUST check the CPE is authorized for this Option. If not, the DHCP Server sends a Status Code set to UnauthorizedRequester. At last, it MUST check if Delegated Domain has been provisioned otherwise the DHCP Server MUST send a Status Code set to UnprovisionedDelegatedDomain. For any other reasons, a Status Code set to UnspecFail MUST be sent.

The DHCP Server then looks at the naming-delegation-actions mentioned by the CPE. The CPE has ordered these actions according to their preference, and the most preferred naming-delegation-action is put first. Naming-delegation-actions are proposed by the CPE, thus the DHCP Server MUST skip any naming-delegation-action it does not understand or its local policies prevent to apply for the CPE. Note that the ordered list is only used to chose a naming-delegation-action to be applied. If the chosen naming-delegation-action fails, the DHCP Server does not have to try other naming-delegation-action with lower preference.

To prevent long proposition lists of naming-delegation-actions, the DHCP Server may send a Status Code TooManyNamingDelegationActions. If the naming-delegation-actions list is void, the DHCP MUST send a Status Code set to VoidNamindDelegationActionList. If none of the naming-delegation-action is acceptable, the DHCP Server MUST send a Status Code of NoApplicableNamingDelegationAction. These Status Code are reported in a Delegated DNS Architecture DHCP Option with naming-

delegation-action set to NONE.

In this document, the naming-delegation-action considered can be CLEAR, SET_NAMING_DELEGATION_WITH_DNS, SET_NAMING_DELEGATION_WITH_DNSSEC. Any other proposition is skipped by the DHCP Server.

If CLEAR is the chosen naming-delegation-action, there not reason the DHCP Server cannot remove the configurations settings. In response, the DHCP Server MUST send a Delegated DNS Architecture with a single naming-delegation-action set CLEAR. In case of success, the Status Code MUST be set to Success, otherwise, it MUST be set to UnspecFail.

For both SET_NAMING_DELEGATION_WITH_DNS and SET_NAMING_DELEGATION_WITH_DNSSEC naming-delegation-actions, the DHCP MUST have an IP address for the Delegated DNS Server. This IP address can be pre-agreed. In this document we consider that this IP address can be derived from the parameters provided by the Delegated DNS Address Information DHCP Option. It is up to the DHCP Server to define how to proceed between the pre-agreed IP address and the one derived from the Delegated DNS Address Information DHCP Option. There may be multiple Delegated DNS Address Information DHCP Options, and the DHCP Server may chose to consider all of these IP Addresses. On the other hand, the DHCP Server may also chose to send a Status Code set to DelegatedIPAddressConflict. This Status Code is sent in a Delegated DNS Architecture DHCP Option with naming-delegation-action set to the corresponding naming-delegation-action.

The DHCP Server accepts the Delegated DNS Address Information DHCP Options it should first proceed to it. If there are multiple Delegated DNS Address Information DHCP Options, the DHCP Server may process to all of them. It may proceed to the Naming Delegation Architecture Configuration if at least one IP address is valid or if all IP addresses are valid.

For the SET_NAMING_DELEGATION_WITH_DNSSEC naming-delegation-action, the DHCP Server MUST check a Delegation of Signing DHCP Option has been provided. If not a Status Code set to MissingDNSSECDelegationOfSigning.

If the Delegated DNS Address Information and the Delegation of Signing DHCP Options have been processed successfully, the DHCP Server MUST configure the Delegating Server, with the IP address(es) and DS record in its zone. Values for the TTL are defined according to the DHCP Timer. The TTL value MUST NOT be greater than the valid-lifetime of the Prefix [RFC3633]. Then, the DHCP Server sends back the Delegated DNS Architecture DHCP Option with a Status Code set to Success.

6.2.4. Processing the Delegated DNS Address Information DHCP Option

Global Unicast IPv6 Addresses are composed of the ISP assigned prefix, that is usually composed of 56 bits, followed by the subnet-ID, typically composed of 8 bits and the interface-ID composed of 64 bits.

In order to set properly the Naming delegation, one MUST make sure the DHCP Server and the CPE agree on the IP address of the Delegated DNS Server. The CPE may not be aware of its ISP assigned prefix and has requested an Identity Association Prefix Delegation DHCP Option for it. The CPE may also have pre-agreed a ISP assigned prefix. In both cases, the CPE and the DHCP Server MUST make sure they agree on the same subnet-ID, that is to say with the same length. The subnet-ID is defined by setting all unknown bits of the ISP assigned prefix to zero. If the number of zeros does not match the size of the ISP assigned prefix, the DHCP Server MUST send a Delegated DNS Architecture DHCP Option with a Status Code set to SubnetIDNonMatchingISPDelegatedPrefixLength Status Code.

For clarification on the agreed IP address of the Delegated DNS Server, the DHCP Server may send in the DHCP Reply the Delegated DNS Address Information DHCP Option with the complete information. In that case, the DHCP Server MUST add a Status Code set to Success.

6.2.5. Processing the Delegation of Signing DHCP Option

The Format of the DS RDATA is defined in [RFC4034].

6.3. CPE Receiving the ISP DHCP Response for the Naming Delegation Architecture

The Delegated DNS Architecture DHCP Option (OPTION_DELEGATED_DNS_ARCHITECTURE) informs the CPE whether the Naming Delegation Architecture has been set as well as the configuration used by the ISP.

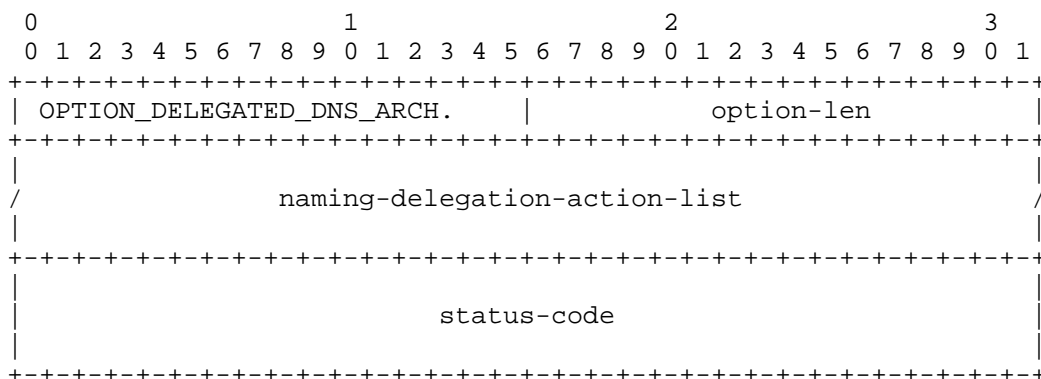
7. DHCP Options

The options detailed in this section are

- Delegated DNS Architecture (OPTION_DELEGATED_DNS_ARCHITECTURE): is used by the DHCP Client on the CPE to inform how the Naming Delegation Architecture should be configured. In return, it is used by the ISP DHCP Server to report the Status Code.

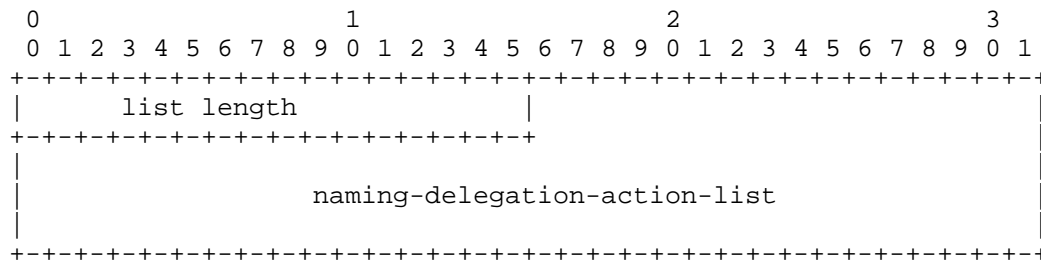
- Delegated Domain (OPTION_DELEGATED_DOMAIN): is used by the DHCP Server to advertise the CPE the Delegated Domain of the Home Network. This Delegated Domain has been reserved and assigned by the End User during the subscription. This option is used to configure properly the DNS zone file of the CPE.
- Delegated DNS Address Information (OPTION_DELEGATED_DNS_ADDR_INFO): is used by the CPE to advertise the DHCP Server which interface and subnet identifier is used by the CPE to build the IPv6 address using the delegated IPv6 prefix to host the DNS Server. This option is used so the DELEGATING_SERVERS can properly fix the delegation.
- Delegated Delegation of Signing (OPTION_DELEGATED_DNSSEC_DS): is used by the CPE so the DELEGATING_SERVERS can properly fix the DNSSEC Naming Delegation.

7.1. Delegated DNS Architecture Option



- option-code: OPTION_DELEGATED_DNS_ARCHITECTURE.
- option-len: Length of the delegated-naming-action-list field, the status-code and the status-message in octets.
- naming-delegation-action-list: The list of the actions the CPE is ready to accept.
- status-code: The Status Code of the operation as specified in [RFC3315]. This option may be absent if operation is successful.

The naming-delegation-action-list is encoded as follows:



- list length: Length of the 'naming-delegation-action-list' field in octets
- naming-delegation-action-list: List of proposed actions by the CPE to the ISP DHCP Server.

The naming-delegation-actions are 1 octet length, and the following values are considered in this document:

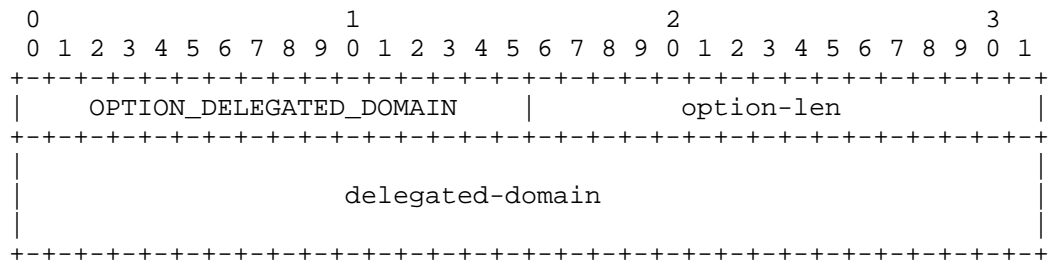
- NONE - 0 - : Indicates that the DHCP Server MUST remove the Naming Delegation Architecture Configuration settings on the Delegating DNS Server.
- CLEAR - 1 - : Indicates that the DHCP Server MUST remove the Naming Delegation Architecture Configuration settings on the Delegating DNS Server.
- SET_NAMING_DELEGATION_WITH_DNS - 2 - : Indicates that the DHCP Server MUST set the Naming Delegation Architecture with only DNS, and MUST NOT consider DNSSEC Delegation.
- SET_NAMING_DELEGATION_WITH_DNSSEC - 3 - : Indicates that the DHCP Server MUST set the Naming Delegation Architecture with DNSSEC.

The Status code 1 octet length and this section considers the following values:

- Success - 0 - :
- UnspecFail - 1 - :
- MissingPrefixDelegationRequest - TBD - :
- NamingDelegationUnavailable - TBD - :


- UnauthorizedRequester - TBD - :
- UnprovisionedDelegatedDomain - TBD - :
- TooManyNamingDelegationActions - TBD - :
- VoidNamindDelegationActionList - TBD - :
- NoApplicableNamingDelegationAction - TBD - :
- SubnetIDNonMatchingISPDelegatedPrefixLength - TBD - :
- DelegatedIPAddressConflict - TBD - :
- MissingDNSSECDelegationOfSigning - TBD - :

7.2. Delegated Domain Option



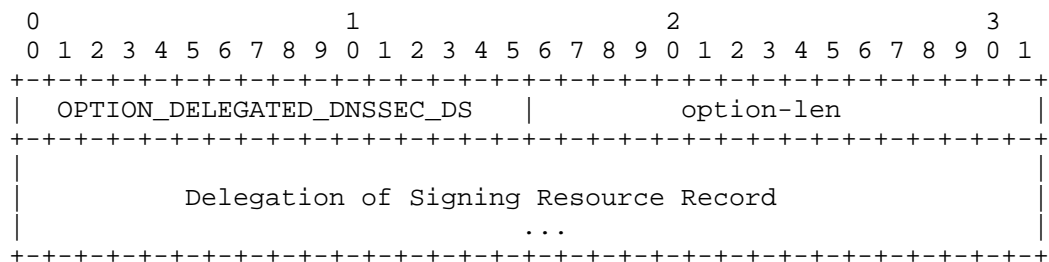
- option-code: OPTION_DELEGATED_DOMAIN
- option-len: Length of the 'Delegated Domain' field in octets.
- delegated-domain: The Delegated Domain encoded as specified in [RFC1035]

7.3. Delegated DNS Address Information Option



- option-code: OPTION_DELEGATED_DNS_ADDR_INFO
- option-len: Length (16) of the Delegated DNS addressing information.
- subnet-ID: The identifier of a subnet used by the authoritative DNS server for the delegated domain name. Only the last 'm' bits are significant. The 'm' value is equal to (64 - 'n') where 'n' is the delegated prefix length. The subnet-ID may be dynamically truncated by the DHCP server and client to match the 'm' size (depending on the delegated prefix length).
- interface-ID: The interface-ID of the IPv6 address used by the authoritative DNS server for the delegated domain name.

7.4. Delegated Delegation of Signing Option



- option-code: OPTION_DELEGATED_DNSSEC_DS

- option-len: Length of the 'Delegated Domain' field in octets.
- DS Resource Record: The DS Resource Record as defined in [RFC4034], Section 5.

8. IANA Considerations

This document introduces Status Code that are carried in the DHCP Options defined in this document. The Status Code detailed in this document are:

- NamingDelegationServiceNotProvided TBD
- UnauthorizedForNamingDelegationService TBD
- NoDelegatedDomainProvisionned TBD
- NoDelegatedDnsAddrInfo TBD
- DelegationSetWithDns TBD
- DelegationSetWithDnssec TBD
- AcceptingOnlyDnssecNamingDelegation TBD
- UnableToSetNamingDelegation TBD
- SubnetIDNonMatchingISPDelegatedPrefixLength TBD

The DHCP options detailed in this document are:

- OPTION_DELEGATED_DNS_ARCHITECTURE: TBD
- OPTION_DELEGATED_DOMAIN: TBD
- OPTION_DELEGATED_DNS_ADDR_INFO: TBD
- OPTION_DELEGATED_DNSSEC_DS: TBD

9. Security Considerations

9.1. Names are less secured than IP addresses

This document describes how an End User can make its services and devices from its Home Network reachable on the Internet with Names rather than IP addresses. This exposes the Home Network to attacker

since names are expected to provide less randomness than IP addresses. The naming delegation protects the End User's privacy by not providing the complete zone of the Home Network to the ISP. However, using the DNS with names for the Home Network exposes the Home Network and its components to dictionary attacks. In fact, with IP addresses, the Interface Identifier is 64 bit length leading to 2^{64} possibilities for a given subnetwork. This is not to mention that the subnet prefix is also of 64 bit length, thus providing another 2^{64} possibilities. On the other hand, names use either for the Home Network domain or for the devices presents less randomness (livebox, router, printer, nicolas, jennifer, ...) and thus exposes the devices to dictionary attacks.

9.2. Names are less volatile than IP address

IP addresses may be used to locate a device, a host or a Service. However, Home Network are not expected to be assigned the same Prefix over time. As a result observing IP addresses provides some ephemeral information about who is accessing the service. On the other hand, Names are not expected to be as volatile as IP addresses. As a result, logging Names, over time, may be more valuable than logging IP addresses, especially to profile End User's characteristics.

PTR provides a way to bind an IP address to a Name. In that sense responding to PTR DNS Queries may affect the End User's Privacy. For that reason we recommend that End Users may choose to respond or not to PTR DNS queries

9.3. DNSSEC is recommended to authenticate DNS hosted data

The document describes how the Secure Delegation can be set between the Delegating DNS Server and the Delegated DNS Server.

Deploying DNSSEC is recommended since in some cases the information stored in the DNS is used by the ISP or an IT department to grant access. For example some Servers may performed a PTR DNS query to grant access based on host names. With the described Delegating Naming Architecture, the ISP or the IT department MUST take into consideration that the CPE is outside its area of control. As such, with DNS, DNS responses may be forged, resulting in isolating a Service, or not enabling a host to access a service. ISPs or IT department may not base their access policies on PTR or any DNS information. DNSSEC fulfills the DNS lack of trust, and we recommend to deploy DNSSEC on CPEs.

9.4. Channel between the CPE and ISP DHCP Server MUST be secured

In the document we consider that the channel between the CPE and the ISP DHCP Server is trusted. More specifically, we suppose the CPE is authenticated and the exchanged messages are protected. The current document does not specify how to secure the channel. [RFC3315] proposes a DHCP authentication and message exchange protection, [RFC4301], [RFC5996] propose to secure the channel at the IP layer.

In fact, the channel MUST be secured because the CPE provides necessary information for the configuration of the Naming Delegation. Unsecure channel may result in setting the Naming Delegation with an non legitimate CPE. The non legitimate CPE would then be redirected the DNS traffic that is intended for the legitimate CPE. This makes the CPE sensitive to three types of attacks. The first one is the Deny Of Service Attack, if for example DNS traffic for a lot of CPEs are redirected to a single CPE. CPE are even more sensitive to this attack since they have been designed for low traffic. The other type of traffic is the DNS traffic hijacking. A malicious CPE may redirect the DNS traffic of the legitimate CPE to one of its server. In return, the DNS Servers would be able to provide DNS Responses and redirect the End Users on malicious Servers. This is particularly used in Pharming Attacks. A third attack may consists in isolating a Home Network by misconfiguring the Naming Delegation for example to a non-existing DNS Server, or with a bad DS value.

9.5. CPEs are sensitive to DoS

The Naming Delegation Architecture involves the CPE that hosts a DNS Server for the Home Network. CPE have not been designed for handling heavy load. The CPE are exposed on the Internet, and their IP address is publicly published on the Internet via the DNS. This makes the Home Network sensitive to Deny of Service Attacks. The Naming Delegation Architecture described in this document does not address this issue. The issue is addressed in the Front End Naming Delegation Architecture described in [I-D.mglt-homenet-front-end-naming-delegation].

10. Acknowledgment

The authors wish to thank Ole Troan for pointing out issues with the IPv6 routed home concept and placing the scope of this document in a wider picture, Mark Townsley for encouragement and injecting a healthy debate on the merits of the idea, Ulrik de Bie for providing alternative solutions, Paul Mockapetris for pointing out issues of the trustworthiness of a reverse lookup, and Christian Jacquenet for seeing the value from a Service Provider point of view.

11. References

11.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.
- [RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", RFC 6672, June 2012.

11.2. Informational References

- [I-D.mglt-homenet-front-end-naming-delegation] Cloetens, C., Lemordant, P., and D. Migault (Ed), "IPv6 Home Network Front End Naming Delegation", draft-mglt-homenet-front-end-naming-delegation-00 (work in progress), July 2012.
- [RFC2118] Pall, G., "Microsoft Point-To-Point Compression (MPPC) Protocol", RFC 2118, March 1997.
- [RFC3769] Miyakawa, S. and R. Droms, "Requirements for IPv6 Prefix Delegation", RFC 3769, June 2004.

Authors' Addresses

Wouter Cloetens
SoftAtHome
vaartdijk 3 701
3018 Wijnmaal
Belgium

Phone:
Email: wouter.cloetens@softathome.com

Philippe Lemordant
Francetelecom - Orange
2, avenue Pierre Marzin
22300 Lannion
France

Phone: +33 2 96 05 35 11
Email: philippe.lemordant@orange.com

Daniel Migault
Francetelecom - Orange
38, rue du General Leclerc
92794 Issy-les-Moulineaux Cedex 9
France

Phone: +33 1 45 29 60 52
Email: mglt.ietf@gmail.com

