

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

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

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

Abstract

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

Status of this Memo

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

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

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

This Internet-Draft will expire on April 26, 2013.

Copyright Notice

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

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

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

## Table of Contents

1. Introduction . . . . .	3
2. Requirements language . . . . .	4
3. Role of Prefix Assignment . . . . .	4
4. Router Behavior . . . . .	5
4.1. Sending Router Advertisements . . . . .	7
4.2. DNS Discovery . . . . .	7
5. Design Choices . . . . .	8
5.1. DNS Discovery . . . . .	8
5.2. Prefix Assignment . . . . .	8
6. Prefix Assignment in OSPFv3 . . . . .	9
6.1. Aggregated Prefix TLV . . . . .	10
6.2. Assigned Prefix TLV . . . . .	11
6.3. OSPFv3 Prefix Assignment . . . . .	12
6.3.1. Making a New Assignment . . . . .	15
6.3.2. Checking for Conflicts Across the Entire Network . . . . .	15
6.3.3. Deprecating an Assigned Prefix . . . . .	16
6.3.4. Verifying and Making a Local Assignment . . . . .	16
7. ULA Generation . . . . .	16
8. Hysteresis . . . . .	18
9. Manageability Considerations . . . . .	19
10. Security Considerations . . . . .	19
11. IANA Considerations . . . . .	19
12. Timer Constants . . . . .	19
13. References . . . . .	20
13.1. Normative References . . . . .	20
13.2. Informative References . . . . .	20
Appendix A. Changes in Version -02 . . . . .	21
Appendix B. Changes in Version -03 . . . . .	21
Appendix C. Acknowledgments . . . . .	21
Authors' Addresses . . . . .	21

## 1. Introduction

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

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

The OSPv3-based mechanism is an alternative design to also using DHCPv6 PD for the prefix assignment in the internal network. This memo has been written so that the working group can make a decision on which type of design to pursue. The main benefit of using a routing protocol to handle the prefix assignment is that it can provide a more efficient use of address space than hierarchical assignment through DHCPv PD. This may be important for home networks that only get a /60 prefix allocation from their ISPs.

The rest of this memo is organized as follows. Section 2 defines the usual keywords, Section 3 explains the main requirements for prefix assignments, Section 4 describes how a home gateway router makes assignments when it itself has multiple subnets, and Section 5 and Section 6 describe how the assignment can be performed in a distributed manner via OSPFv3 in the entire home network. Finally, Section 7 specifies the procedures for automatic generation of ULA prefixes, Section 8 explains the hysteresis principles applied to prefix assignment and de-assignment, Section 9 explains what administrative interfaces are useful for advanced users that wish to manually interact with the mechanisms, Section 10 discusses the security aspects of the design, Section 11 explains the necessary IANA actions, and Section 12 defines the necessary timer constants.

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

## 2. Requirements language

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

## 3. Role of Prefix Assignment

Given a prefix shorter than /64 for the entire home network, this prefix needs to be subdivided so that every subnet is given its own /64 prefix. In many cases there will be just one subnet, the internal network interface attached to the router. But it is also common to have two or more internal network interfaces with intentionally separate networks. For instance, "private" and "guest" SSIDs are automatically configured in many current home network routers. When all the network interfaces that require a prefix are attached to the same router, the prefix assignment problem is simple, and procedures outlined in Section 4 can be employed.

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

The key requirements for this distributed mechanism are as follows.

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

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

- o The assignment mechanism should provide reasonable efficiency. As a practical benchmark, some ISPs may employ /60 allocations to individual subscribers. As a result, the assignment mechanism should avoid wasting too many prefixes so that this set of 16 /64 prefixes is not exhausted in the foreseeable future for commonly occurring network configurations.

- o In particular, the assignment of multiple prefixes to the same network from the same top-level prefix must be avoided.

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

- o The assignments must be stable across reboots, power cycling, router software updates, and preferably, should be stable across simple network changes. Simple network changes are in this case defined as those that could be resolved through either deletion or addition of a prefix assignment. For instance, the addition of a new router without changing connections between existing routers requires just the assignment of new prefixes for the new networks that the router introduces. There are no stability requirements across more complex types of network reconfiguration events. For instance, if a network is separated into two networks connected by a newly inserted router, this may lead to renumbering all networks within the home.

In an even more complex setting there may be multiple home gateway routers and multiple connections to ISP(s). These cases are analogous to the case of a single gateway router. Each gateway will simply distribute the prefix it has, and participating routers throughout the network may assign themselves prefixes from several gateways. Multiple assignments can be made for the same interface. For example, this can be useful in a multi-homing setting.

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

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

#### 4. Router Behavior

This section describes how a router assigns prefixes to its directly connected interfaces. We assume that the router has prefix

allocation(s) that it can use for this assignment. Each such prefix allocation is called an aggregated prefix. Parts of the aggregated prefix may already be assigned for some purpose; a coordinated assignment from the prefix is necessary before it can actually be assigned to an interface.

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

- o The router MUST maintain a list of assigned prefixes on a per-interface basis. The contents of this list consists of prefixes that the router itself has assigned to the interface, as well as prefixes assigned to the interface by other routers. The latter are learned through the mechanisms described in Section 6, when they are used. Each prefix is associated with the Router ID of the router that assigned it.
- o Whenever the router finds a combination of an interface and aggregated prefix that is not used on the interface, it SHOULD make a new prefix assignment. That is, the router checks to see if an interface and aggregated prefix exists such that there are no assigned prefixes within that interface that are more specific than the aggregated prefix. In this situation the router makes an allocation from the aggregated prefix (if possible) and adds the assignment to the list of assigned prefixes on that interface.

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

- o An assignment from an aggregated prefix MUST be checked against possible other assignments from the same aggregated prefix on the same link by neighboring routers, to avoid unnecessary assignments. Assignments MUST also be examined against all existing assignments from the same aggregated prefix across the network, to avoid collisions. Assignments are made for individual /64 prefixes. The choice of a /64 prefix among multiple free ones MUST be made randomly or based on an algorithm that takes unique hardware characteristics of the router and the interface into account. This helps avoid collisions when simultaneous assignments are made within a network.
- o In order to provide a stable assignment, the router MUST store assignments affecting directly connected interfaces and automatically generated ULA prefixes in non-volatile memory and attempt to re-use them in the future when possible. At least the 5 most recent assignments SHOULD be stored. Note that this applies to both its own assignments as well as assignments made by

others. This ensures that the same prefix assignments are made regardless of the order that different devices are brought up. To avoid attacks on flash memory write cycles, assignments made by others SHOULD be recorded only after 10 minutes have passed and the assignment is still valid.

- o Re-using a memorized assignment is possible when a aggregated prefix exists that is less specific than the prefix in the assignment (or it is the prefix itself in the assignment), and the prefix is currently unassigned.

#### 4.1. Sending Router Advertisements

Once the router has assigned a prefix to an interface, it MUST act as a router as defined in [RFC4861] and advertise the prefix in Router Advertisements. The lifetime of the prefix SHOULD be advertised as a reasonably long period, at least 48 hours or the lifetime of the assigned prefixes, whichever is smaller.

#### 4.2. DNS Discovery

To support a variety of IPv6-only hosts in these networks, the router needs to ensure that sufficient DNS discovery mechanisms are enabled. It is RECOMMENDED that both stateless DHCPv6 [RFC3736] and Router Advertisement options [RFC6106] are supported and turned on by default in routers.

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

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

## 5. Design Choices

### 5.1. DNS Discovery

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

### 5.2. Prefix Assignment

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

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

Several design decisions were needed to construct the protocol:

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

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

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

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

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

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

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

## 6. Prefix Assignment in OSPFv3

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

An overview of OSPFv3-based prefix assignment is as follows. OSPFv3 routers that are capable of auto-configuration advertise an OSPFv3 Auto-Configuration (AC) LSA [I-D.ietf-ospf-ospfv3-autoconfig] with suitable TLVs. For prefix assignment, two TLVs are used. The Aggregated Prefix TLV (Section 6.1) advertises an aggregated prefix, usually the prefix that has been delegated to the home gateway router from the ISP through DHCPv6 PD. These aggregated prefixes are necessary for running the algorithm in Section 4 for determining whether prefix assignments can and should be made.

The Assigned Prefix TLV (Section 6.2) is used to communicate assignments that routers make out of the aggregated prefixes.

An assignment can be made when the algorithm in Section 4 indicates that it can be made and no other router has claimed the same assignment. The router makes an OSPFv3 advertisement with the Assigned Prefix TLV included to let other devices know that the prefix is now in use. Unfortunately, collisions are still possible, when the algorithms on different routers happen to choose the same free /64 prefix or when more /64 prefixes are needed than are available. This situation is detected through an advertisement where a different router claims the assignment of the same prefix. In this

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

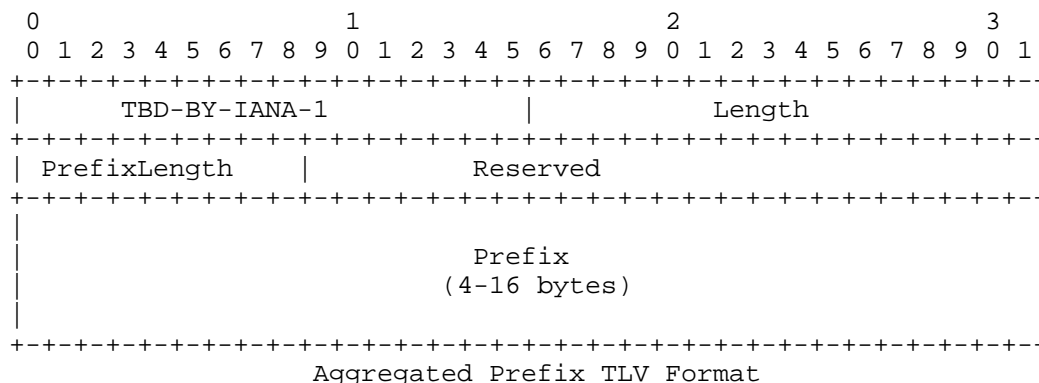
### 6.1. Aggregated Prefix TLV

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

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

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

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



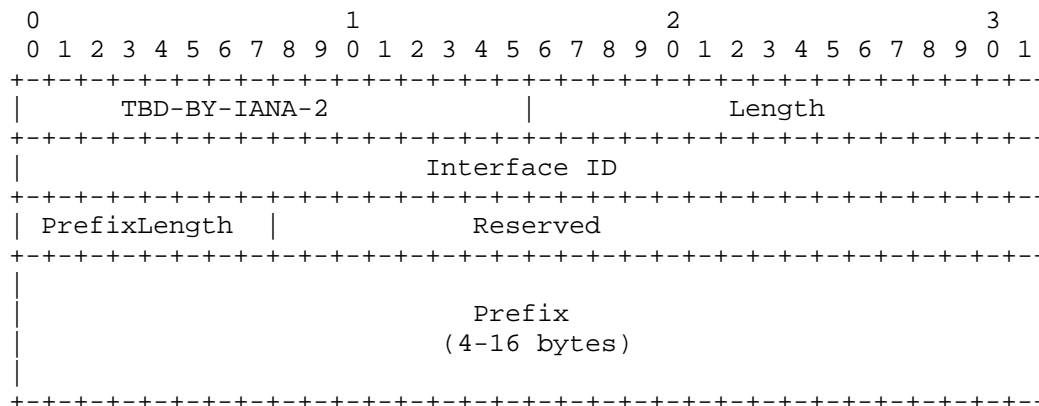
## 6.2. Assigned Prefix TLV

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

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

This TLV MUST be advertised by the router that has made assignment from an aggregated prefix per Section 4.

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



Assigned Prefix TLV Format

### 6.3. OSPFv3 Prefix Assignment

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

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

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

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

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

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

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

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

i.

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

ii.

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

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

iii.

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

4. There are four possibilities at this stage:

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

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

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

#### 6.3.1. Making a New Assignment

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

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

#### 6.3.2. Checking for Conflicts Across the Entire Network

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

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

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

#### 6.3.3. Deprecating an Assigned Prefix

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

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

#### 6.3.4. Verifying and Making a Local Assignment

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

### 7. ULA Generation

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

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

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

i.

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

ii.

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

iii.

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

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

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

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

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

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

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

## 8. Hysteresis

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

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

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

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

A previously assigned prefix SHOULD NOT be taken out of use before the router has waited TERMINATE\_PREFIX\_ASSIGNMENT seconds.

## 9. Manageability Considerations

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

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

Finally, there **SHOULD** be mechanisms to display the prefixes assigned on each interface, and where they came from (manual configuration, DHCPv6 PD, OSPFv3).

## 10. Security Considerations

Security can be always added later.

## 11. IANA Considerations

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

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

## 12. Timer Constants

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

## 13. References

## 13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3646, December 2003.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.
- [I-D.ietf-ospf-ospfv3-autoconfig]  
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration", draft-ietf-ospf-ospfv3-autoconfig-00 (work in progress), October 2012.

## 13.2. Informative References

- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [I-D.ietf-homenet-arch]  
Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "Home Networking Architecture for IPv6", draft-ietf-homenet-arch-06 (work in progress), October 2012.
- [I-D.chelius-router-autoconf]  
Chelius, G., Fleury, E., and L. Toutain, "Using OSPFv3 for IPv6 router autoconfiguration", draft-chelius-router-autoconf-00 (work in progress), June 2002.

[I-D.dimitri-zospf]

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

[SIGCOMM.IPV6]

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

#### Appendix A. Changes in Version -02

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

#### Appendix B. Changes in Version -03

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

#### Appendix C. Acknowledgments

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

#### Authors' Addresses

Jari Arkko  
Ericsson  
Jorvas 02420  
Finland

Email: jari.arkko@piuha.net

Acee Lindem  
Ericsson  
Cary, NC 27519  
USA

Email: [acee.lindem@ericsson.com](mailto:acee.lindem@ericsson.com)

Benjamin Paterson  
Cisco Systems  
Paris  
France

Email: [benjamin@paterson.fr](mailto:benjamin@paterson.fr)



Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: April 18, 2013

M. Behringer  
M. Pritikin  
S. Bjarnason  
Cisco  
October 15, 2012

Bootstrapping Trust on a Homenet  
draft-behringer-homenet-trust-bootstrap-00.txt

Abstract

A homenet must be aware of its borders, and the realms within those. This document proposes an approach to bootstrap trust in such an environment. The idea is to select one device as the trust anchor and to enroll other devices into the domain. The result is the creation of a domain of trust in the homenet, with a common trust anchor. This trust model can subsequently be used to determine boundaries, and to autonomically bootstrap network services.

Status of this Memo

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

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

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

This Internet-Draft will expire on April 18, 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.

## 1. Problem Statement

[I-D.ietf-homenet-arch] states that "It should be possible to automatically perform border discovery for the different borders." Simple approaches, such as terminating a homenet on a particular interface type do not easily allow for devices from different administrative realms to be locally connected. [I-D.ietf-homenet-arch] states further that "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."

An approach is needed that allows to establish trust inside a homenet according to a policy set by the admin of the homenet.

## 2. Approach

An autonomic device can be a router, switch, PC, smartphone, or any other device, independent of its role in the network, which has the autonomic functionality mentioned below . A homenet consists of autonomic devices and non-autonomic devices. This approach requires at least one autonomic networking device, such as a router or switch.

One autonomic device in the homenet takes on a registrar function. This could be manually enabled, for example on a smartphone autonomic app; in the absence of a registrar function, a device can also auto-select itself to take on this function, using some detection mechanism to resolve potential conflicts.

The registrar creates a trust anchor for the homenet, and subsequently acts as a registration authority, granting domain certificates to other devices.

Every autonomic device discovers neighbouring autonomic nodes through an autonomic neighbour discovery protocol. This could be implemented for example through IPv6 neighbour discovery, using a to-be-assigned well-known multicast address indicating "all autonomic nodes on this subnet".

An autonomic device signs its neighbour discovery packets. If it has a domain certificate from the domain registrar, it uses that. If not, it uses either a vendor certificate (e.g., an IEEE 802.1AR

[IDevID] credential) or a self-signed certificate.

If two autonomic homenet devices use the same trust anchor they can verify each other's certificate thus establishing that the peer is a member of the same local domain.

If one autonomic homenet device is member of a domain, and its neighbour is not, it invites the neighbour to join the domain. The device without domain credentials requests to join the first domain it is presented with. The device MUST only join a homenet domain when it is in the factory default configuration (e.g. it is not currently a member of a homenet). The domain device proxies the request to the registrar, including the device credentials of the device without domain credentials.

The registrar accepts or declines a request to join the domain, based on the credentials presented and other policy defined criteria such as proxy identity. Any authorized device currently within the domain MAY provide supplemental criteria for help making this decision. A smartphone autonomic application would be an ideal domain member to provide user interface functionality for the obtaining of supplemental criteria from end users.

If a device is accepted into the domain, it is invited to request a domain certificate through a certificate enrollment process.

The result is a common trust anchor and device certificates for all autonomic devices in a domain. These certificates can subsequently be used to determine the boundaries of the homenet, to authenticate other domain nodes, and to autonomically enable services on the homenet.

### 3. Security Considerations

The approach as outlined in this document is open to a number of attacks at bootstrap time. For example, a malicious device could pretend to be an expected device and assume its role. This is however no different to current security models in home networks.

### 4. Informative References

[I-D.ietf-homenet-arch]

Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil,  
"Home Networking Architecture for IPv6",  
draft-ietf-homenet-arch-04 (work in progress), July 2012.

[IDevID] IEEE Standard, "IEEE 802.1AR Secure Device Identifier",  
December 2009, <[http://standards.ieee.org/findstds/  
standard/802.1AR-2009.html](http://standards.ieee.org/findstds/standard/802.1AR-2009.html)>.

Authors' Addresses

Michael H. Behringer  
Cisco

Email: [mbehring@cisco.com](mailto:mbehring@cisco.com)

Max Pritikin  
Cisco

Email: [pritikin@cisco.com](mailto:pritikin@cisco.com)

Steinthor Bjarnason  
Cisco

Email: [sbjarnas@cisco.com](mailto:sbjarnas@cisco.com)



Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: April 18, 2013

T. Boot  
Infinity Networks B.V.  
October 15, 2012

BRDP for Homenet  
draft-boot-homenet-brdp-00.txt

## Abstract

This document describes the Border Router Discovery Protocol (BRDP) and all of its related components. BRDP enables multi-homing for small to medium sites, including Homenets, using Provider Aggregatable IPv6 addresses. It describes a mechanism for automated IP address configuration and renumbering, a mechanism for optimized source address selection and a new paradigm for packet forwarding, for support of multi-homed sites. BRDP prevents ingress filtering problems with multi-homed sites and supports load-balancing for multi-path transport protocols. This work also prevents routing scalability problems in the provider network and Internet Default Free Zone because small to medium multi-homed size sites would not need to request Provider Independent address blocks.

## Status of this Memo

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

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

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

This Internet-Draft will expire on April 18, 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. Detection of Homenet Perimeter interfaces on Border Routers . . . . .	5
1.2. Propagation of Border Router information . . . . .	5
1.3. Address configuration with Border Router information . . . . .	6
1.4. Address selection with Border Router information . . . . .	6
1.5. Routing based on Border Router information . . . . .	7
1.6. Deployment of the Border Router Discovery Protocol . . . . .	7
1.7. Requirements Language . . . . .	8
2. Reference Scenarios . . . . .	8
2.1. Single-homed site . . . . .	8
2.2. Small multi-homed site or DMZ . . . . .	10
2.3. Medium multi-homed site . . . . .	12
2.4. Medium multi-homed site with ULAs and DHCP server . . . . .	16
2.5. MANET site . . . . .	18
3. Border Router Discovery Protocol (BRDP) . . . . .	20
3.1. Border Router Information Option (BRIO) . . . . .	21
3.2. BRDP processing . . . . .	22
3.2.1. BRDP message generation and transmission . . . . .	23
3.2.2. BRDP message reception . . . . .	24
3.2.3. BRIO-Cache maintenance . . . . .	25
3.2.4. BRDP loop prevention . . . . .	26
3.3. Unified Path Metric (UPM) . . . . .	27
4. BRDP based Address Configuration and Prefix Delegation . . . . .	27
4.1. Border Router selection . . . . .	28
4.1.1. Border Router Selection based on UPM . . . . .	28
4.2. Address autoconfiguration . . . . .	29
4.2.1. Address and prefix configuration with SLAAC or DHCP . . . . .	29
4.2.2. Address generation and configuration for Routers . . . . .	29
4.2.3. Support for Unique Local Addresses (ULA) . . . . .	30
5. BRDP based Source Address Selection . . . . .	30
5.1. Address Selection for dynamic DNS . . . . .	30
6. BRDP based Routing . . . . .	30
6.1. Problems with default gateway routing . . . . .	31
6.2. Default gateway routing replaced with BRDP Based Routing . . . . .	31
7. BRDP and IRTF RRG goals . . . . .	32

7.1.	Scalability . . . . .	33
7.2.	Traffic engineering . . . . .	33
7.3.	Multi-homing . . . . .	33
7.4.	Loc/id separation . . . . .	33
7.5.	Mobility . . . . .	34
7.6.	Simplified renumbering . . . . .	34
7.7.	Modularity . . . . .	34
7.8.	Routing quality . . . . .	34
7.9.	Routing security . . . . .	34
7.10.	Deployability . . . . .	35
8.	Currently unaddressed issues . . . . .	35
9.	Acknowledgements . . . . .	35
10.	IANA Considerations . . . . .	35
11.	Security Considerations . . . . .	35
12.	Change log . . . . .	36
13.	References . . . . .	36
13.1.	Normative References . . . . .	36
13.2.	Informative References . . . . .	37
	Author's Address . . . . .	38

## 1. Introduction

\*\*\* Note to the reader \*\*\* This Internet Draft is submitted as an early version for a proposal for the Homenet working group. This version is a merge from earlier documents. Now that is a single document, it is to be adjusted to comply to the Homenet scenario. This is work in progress.

IPv6 provides basic functionality for multi-homing, since nodes can have multiple addresses configured on their interfaces. However, it is difficult to utilize the advantages of this, as there is a strong tendency shielding the network topology from hosts and in general routing does not support multi-homing very well. As a result, it is difficult or impossible for a host to utilize available facilities of the network, such as multi-path. Also scalability of the Internet routing system is getting a problem due to a high demand of Provider Independent (PI) addresses.

The Border Router Discovery Protocol (BRDP) enhances the IPv6 model by enabling automated renumbering in dynamically changing multi-homed environments, such that routers and hosts cooperate on address configuration and path selection. BRDP utilizes Provider Aggregatable (PA) addresses and uses them as locator. Mapping identifiers to locators is out of scope of BRDP, also because other solutions exists or are being worked on. All these solutions work fine with BRDP, as long as they don't break IPv6.

BRDP applies to edge networks. These networks can be fixed, for example enterprise networks, small offices / home offices (SOHO) or home sites (Homenets). BRDP also can be used in wireless access networks, for example wireless access networks such as 3G or 4G, wireless LANs or mobile ad hoc networks (MANETs). A nice attribute of BRDP is that it supports multi-homing in heterogeneous networks, meaning that e.g. a Homenet network can have multiple wired broadband and 3G/4G connections to the Internet simultaneously.

In a multi-homed network, nodes are connected to the Internet via multiple exit points, possibly via multiple providers. [RFC5887] argues that if a site is multi-homed, using multiple PA routing prefixes, then the interior routers need a mechanism to learn which upstream providers and corresponding PA prefixes are currently reachable and valid. Next to that, these upstream providers or PA prefixes may change over time. This requires a dynamic renumbering mechanism that can handle planned or unplanned changes in the prefixes used. BRDP proposes a mechanism for automated renumbering in larger networks that goes beyond hosts in a single subnet.

BRDP uses the following key elements:

- o Propagation of available Border Routers and corresponding prefixes, described in Section 3;
- o Address autoconfiguration and prefix delegation, using BRDP provided hints, described in Section 4;
- o Source address selection, using BRDP provided hints, described in Section 5;
- o Packet forwarding to the Border Router that corresponds with the source address prefix, in case the destination address is not found in the routing domain, described in Section 6.

#### 1.1. Detection of Homenet Perimeter interfaces on Border Routers

For fully automated deployment in Homenets, it is required that routers can discover automatically their uplink interfaces, that connect the Homenet to ISPs. Some mechanisms for automatic detection are described in [I-D.kline-default-perimeter].

After detection of an uplink interface to an ISP and reception of a prefix, the router starts acting as a border router. It starts acting as a DHCP server, with support of prefix delegation. It also configures at least an address out of the assigned prefix. This address is used as Border Router address and DHCP server address.

The BRDP protocol can also be used to assist perimeter detection. A router interface on which Border Router information is received should not be identified as an uplink interface to an ISP.

#### 1.2. Propagation of Border Router information

The propagation of available Border Routers and corresponding prefixes is implemented as an extension on the Neighbor Discovery Protocol [RFC4861]. Border Router Information Options (BRIOS) are sent with Router Advertisements, and contain information about the Border Router, such as:

- o - the Border Router address;
- o - the prefix that corresponds with that Border Router;
- o - cost indication of the path via that Border Router to the core network, i.e. the Internet Default Free Zone (DFZ).

BRIOS are disseminated downstream through the network. All nodes store the information from BRIOS they receive in a BRIO cache.

Border routers with multiple prefixes send out a BRIO for each of these prefixes. In a multi-homed network, nodes will receive multiple prefix information, from multiple upstream Border Routers or from a Border Router with multiple prefixes.

### 1.3. Address configuration with Border Router information

Routers can generate IPv6 addresses, with regular SLAAC [RFC4862]. Generation is based on Prefix Information Option from upstream routers and optionally on information in the BRIO cache, e.g. using the prefix with the lowest cost to the Internet. In addition, routers may generate /128 IPv6 address-prefixes for a management interface, based on a Border Router prefix. Routers set up reachability to these addresses automatically, by adding the generated address or prefix in the routing protocol.

With BRDP, routers automatically learn Border Routers that act as DHCP server or relay agent [RFC3633]. When routers detect an alternate path to the DFZ, with no corresponding assigned address or prefix already, new prefixes are requested for using this alternate path.

Prefixes, of which the path to the DFZ is no longer available, are put 'out of service' by routers, meaning they are not used for address assignments anymore. Optionally, if the cost to the DFZ through a Border Router is far higher than via other available paths, a router can put the corresponding prefix out of service also. Prefixes that are out of service are released.

Prefixes that are in service are configured on interfaces with a 64-bit prefix length and advertised with a Prefix Information Option in Router Advertisements. The Prefix Information lifetime is copied from lifetime information in the BRIO cache.

Hosts can use the BRDP provided information together with the Prefix Information to autoconfigure addresses, based on IPv6 Stateless Address Autoconfiguration [RFC4862]. A host may also use DHCPv6 to get addresses or "Other configuration", using multicast or with unicast to the BRDP learned DHCP server address.

### 1.4. Address selection with Border Router information

Nodes with multiple configured addresses need to select a source address for outgoing connections. Default Address Selection for IPv6 [RFC6724] defines a mechanism, used as default behavior. It is open to more advanced mechanisms or site policies. BRDP provided information can be used for a more advanced mechanism, where the hosts select automatically a source address that corresponds with a path with the lowest cost to the DFZ. When multiple Border Routers are available, automatic load distribution and multi-path transport becomes available.

### 1.5. Routing based on Border Router information

Network Ingress Filtering [RFC2827] describes the need for ingress filtering, to limit the impact of distributed denial of service attacks, by denying traffic with spoofed source addresses access. It also helps ensure that traffic is traceable to its correct source network. Ingress Filtering for Multihomed Networks [RFC3704] provides solutions for multi-homed sites. However, the proposal applicable for PA addresses requires careful planning and configuration. It suggests routing based on source address, and a path on each Border Router to all ISPs in use, either with a direct connection or with tunnels between all Border Routers. It is hard to make such mechanisms work in an automated fashion, or mechanisms are not supported on Border Routers used today. As an evolutionary approach, BRDP provided information is to be used to forward packets to their destination without ingress filtering problems. The BRIO cache contains a mapping between Border Routers and the addresses that do pass ingress filtering. So the packet forwarding heuristic can be straightforward: send packets, where the destination is not in the routing domain itself, to the Border Router that owns the prefix of the source address.

Hosts use information in the Default Router List to select a default router. For selecting the best paths, hosts may use next hop selection based on source address and path costs to the corresponding Border Router, if such information is available to the host. Such next hop determination is useful for destinations outside the edge network, i.e. the destination address does not belong to a prefix in the BRIO cache.

### 1.6. Deployment of the Border Router Discovery Protocol

Enabling BRDP in an existing network is straightforward. First, all routers have to be updated for BRDP support. At this step, Border Router information is propagated in the network enabling BRDP assisted address autoconfiguration and prefix delegation and BRDP assisted source address selection. The second step is updating all routers with the BRDP based routing mechanism. To enable this mechanism the default gateway is removed from the routing table. This second step is a flag day operation. Rolling back is easy, by just re-inserting the default gateway.

After the update of the network, additional border routers can be added to the network and will be used automatically. Also a renumbering event will take place without any manual intervention.

BRDP does not provide session continuity when paths are broken. Mobility solutions are in place, or are work in progress. Recently,

interesting developments are work in progress, such as MPTCP [I-D.ietf-mptcp-multiaddressed] and ILNP [I-D.rja-ilnp-intro]. BRDP is very useful for both of these protocols.

### 1.7. 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. Reference Scenarios

This section describes the use of BRDP in five different scenarios: a single homed Homenet, multi-homed site or DMZ, a medium multi-homed site, a medium multi-homed site with ULA with DHCP server and a MANET site.

### 2.1. Single-homed site

This scenario discusses BRDP operation for single-homed home networks. The scenario is taken from [I-D.ietf-homenet-arch].

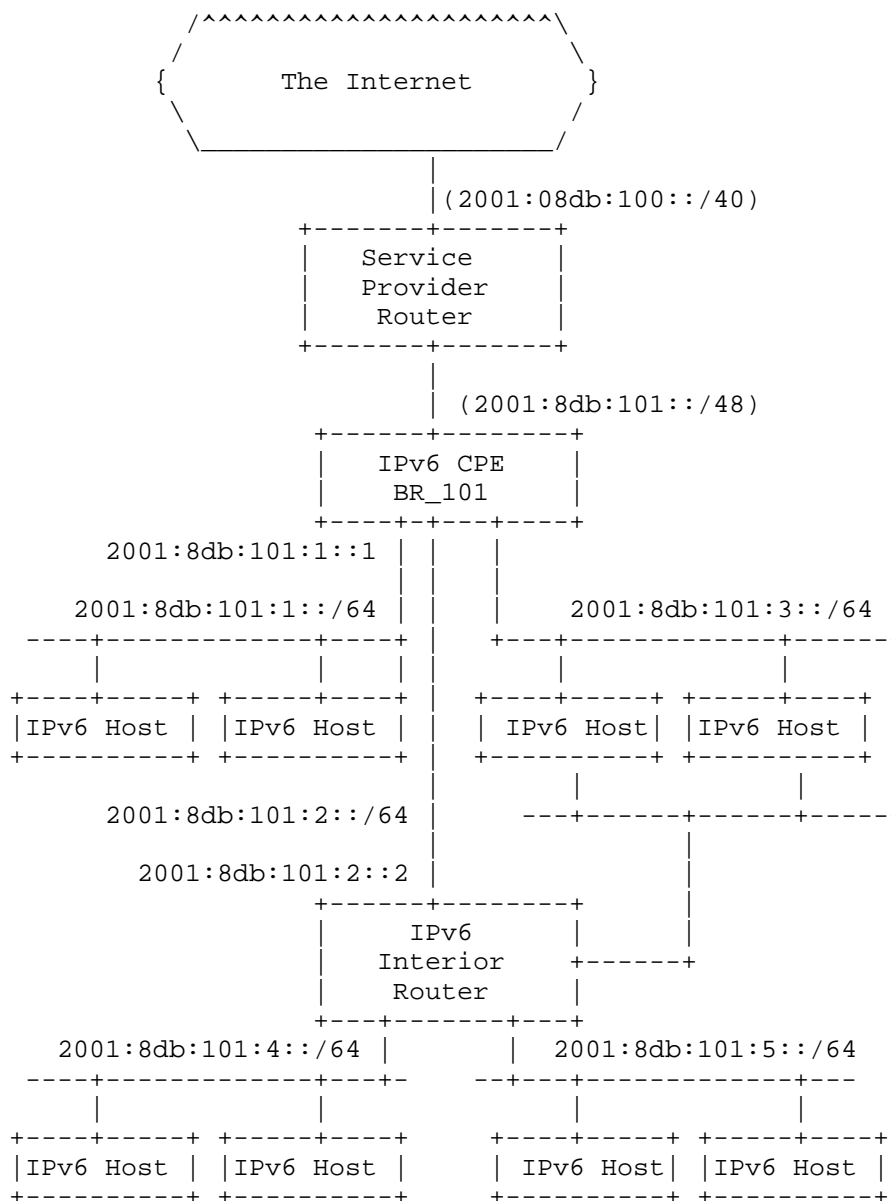


Figure 1: Scenario 1: Single-homed site

The CPE device has to discover the link to the ISP and has to get assigned an IPv6 prefix, in this scenario 2001:8db:101::/48. The CPE configures itself a global unique address prefix 2001:8db:101:1::1/64, assumed on the first interface in the homenet. It may

configure additional global unique address on other interfaces, but this is not required. This is existing functionality which is not updated by BRDP.

The CPE starts sending router advertisements. It also checks received router advertisements on already existing prefixes for the /48 prefix it has assigned by the ISP. In this scenario there is no other CPE, so no on-link prefixes exist. The CPE allocates and bind additional prefixes for all its interfaces, and send Router Advertisements with the Prefix Information Option. By then it has configured 2001:8db:101:1::/64, 2001:8db:101:2::/64 and 2001:8db:101:3::/64. The CPE router also acts as DHCP server, for the ISP provided prefix.

Now, the IPv6 hosts in the middle row learn these prefixes from Prefix Information Options sent by the CPE. They can configure IPv6 addresses, either with SLAAC or DHCP. Also, the IPv6 Interior Router can configure an IPv6 address, in this scenario on the link with prefix 2001:8db:101:2::/64.

The IPv6 Interior Router also receives the router advertisement with the onlink prefix 2001:8db:101:3::/64 on its interface on the right. It could configure an address in the prefix, but because it has already a globally unique address configured, there is no need for this. Question is if the router should echo the prefix as on-link. In this BRDP proposal, it doesn't. It is not the "delegated prefix holder".

Before the IPv6 hosts on the lower row can get their addresses, the IPv6 Interior Router has to be assigned two more prefixes. Here, BRDP starts playing its role. The CPE router advertises itself with Border Router Information Option, in its Router Advertisement. The IPv6 Interior Router learns this information, and gets the two needed prefixes from the CPE Router, using unicasted DHCP messages to the (CPE) Border Router. Two prefixes are assigned and configured, 2001:8db:101:4::/64 and 2001:8db:101:5::/64.

For full connectivity, the homenet uses an interior routing protocol. BRDP is agnostic on the routing protocol used.

## 2.2. Small multi-homed site or DMZ

This scenario discusses BRDP operation for multi-homed Small Office - Home Office (SOHO) networks and De-Militarized Zones (DMZ). The scenario is shown in Figure 2. Each provider assigns a PA /48 prefix to its customers. All addresses and prefixes are configured completely automatically. The feature of BRDP that adds value in this scenario is BRDP based Border Router selection for multi-homed

hosts. This is enabled by using BRDP based forwarding.

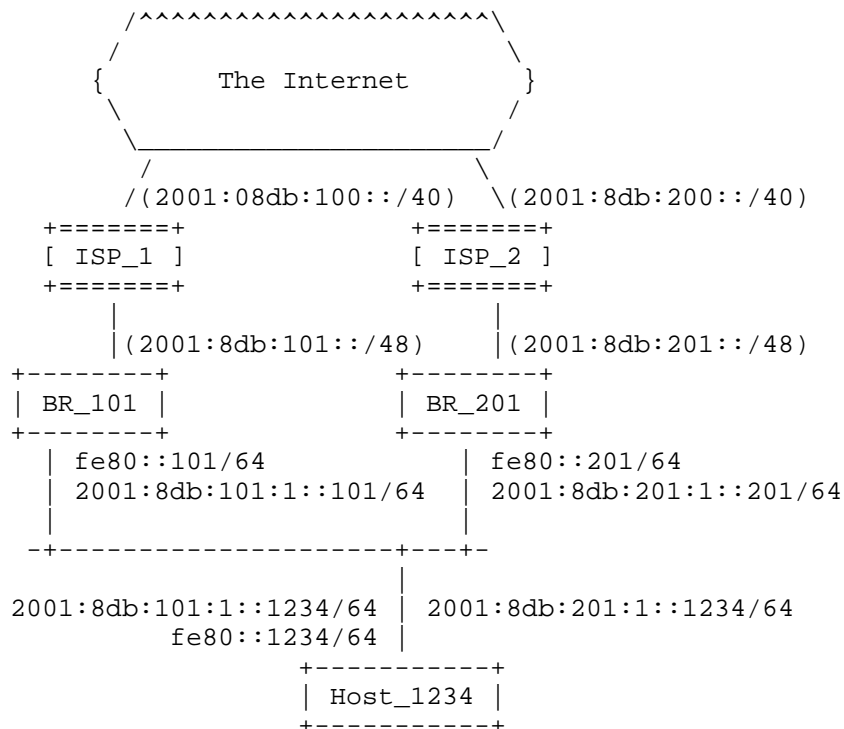


Figure 2: Scenario 2: multi-homed Small Office - Home Office (SOHO) network or DMZ

In this scenario, Host\_1234 has configured two addresses using SLAAC [RFC4862], one with prefix 2001:8db:101:1::/64 from Border Router BR\_101 and one with prefix 2001:8db:201:1::/64 from Border Router BR\_201. Host\_1234 has learned these prefixes from Prefix Information Options sent by both Border Routers according to [RFC4861]. The host has learned via BRIOs that these prefixes belong to Border Routers. The host can use optimal paths by selecting BR\_101 as default router for all packets with a source address with prefix 2001:8db:101:1::/64 and default gateway BR\_201 for all packets with a source address with prefix 2001:8db:201:1::/64. Non-optimal default router selection on hosts is handled by the routers, "misdirected" packets are forwarded to the correct Border Router.

BRDP enables routers to deliver non-optimal directed packets from attached hosts towards a Border Router that owns the prefix of the source address, if such a Border Router exists. In the above scenario, a packet sent from Host\_1234 with source address 2001:8db:

201:1::1234 to default router BR\_101 would be dropped due to on an ingress filter, when no mechanism is in place to redirect the packet. BRDP based forwarding provides such a mechanism automatically. Instead of dropping the packet, BR\_101 forwards it to BR\_201.

### 2.3. Medium multi-homed site

This scenario discusses BRDP operation for medium sized multi-homed networks. The difference with the previous scenario is that the network paths between hosts and the Border Routers have intermediate routers. The scenario is shown in Figure 3. The added value of BRDP in this scenario is the discovery of Border Routers for hosts and routers beyond the first hop as well as Border Router Selection for hosts and routers.

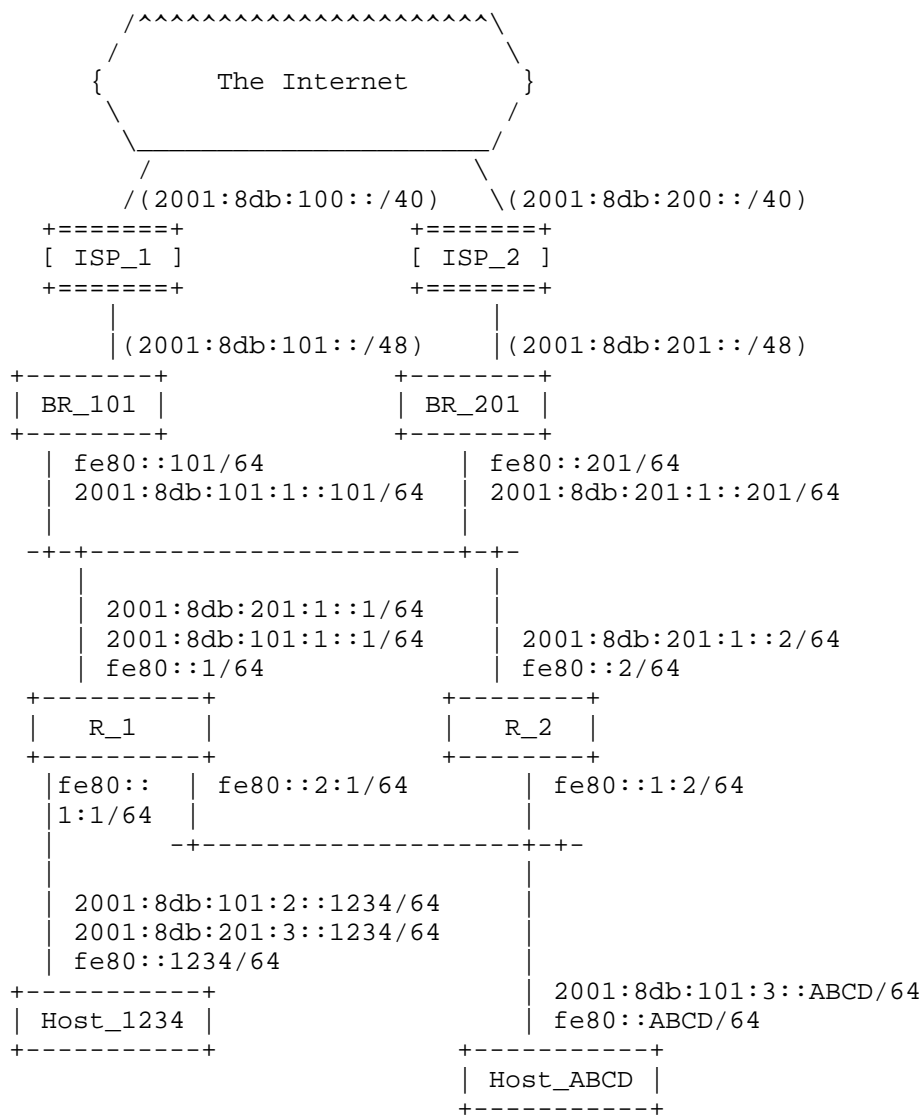


Figure 3: Scenario 3: medium sized multi-homed network

Routers can learn advertised on-link prefixes automatically via the Prefix Information Option in IPv6 ND RAs. In this scenario, routers R\_1 and R\_2 learn prefix 2001:8db:101:1::/64 from BR\_101 and prefix 2001:8db:201:1::/64 from BR\_201. Routers may autoconfigure addresses on their interfaces. In this example, R\_1 has configured addresses from both providers on its upstream interface, R\_2 only configured an address based on the prefix of BR\_201. If the routers run a routing

protocol, the learned prefixes are made reachable in the network. In the next steps of the autoconfiguration proces, the prefixes and addresses on the other links are automatically configured, but first we discuss the BRDP messages that are disseminated through the network.

Routers automatically learn Border Routers and mapping between prefixes and Border Routers using BRDP. The diagram in Figure 4 depicts BRIO message dissemination in scenario 2. The two Border Routers advertise their own address and corresponding prefix with an address prefix. Nothing prevents them from forwarding each other's BRIO message, although resending BRIO information on non-MANET interfaces is not useful. Both routers R\_1 and R\_2 forward both Border Router address prefixes, using separate BRIOs in RAs, on downstream interfaces. In this way all routers and hosts in the network are aware of all reachable Border Routers and corresponding prefixes.

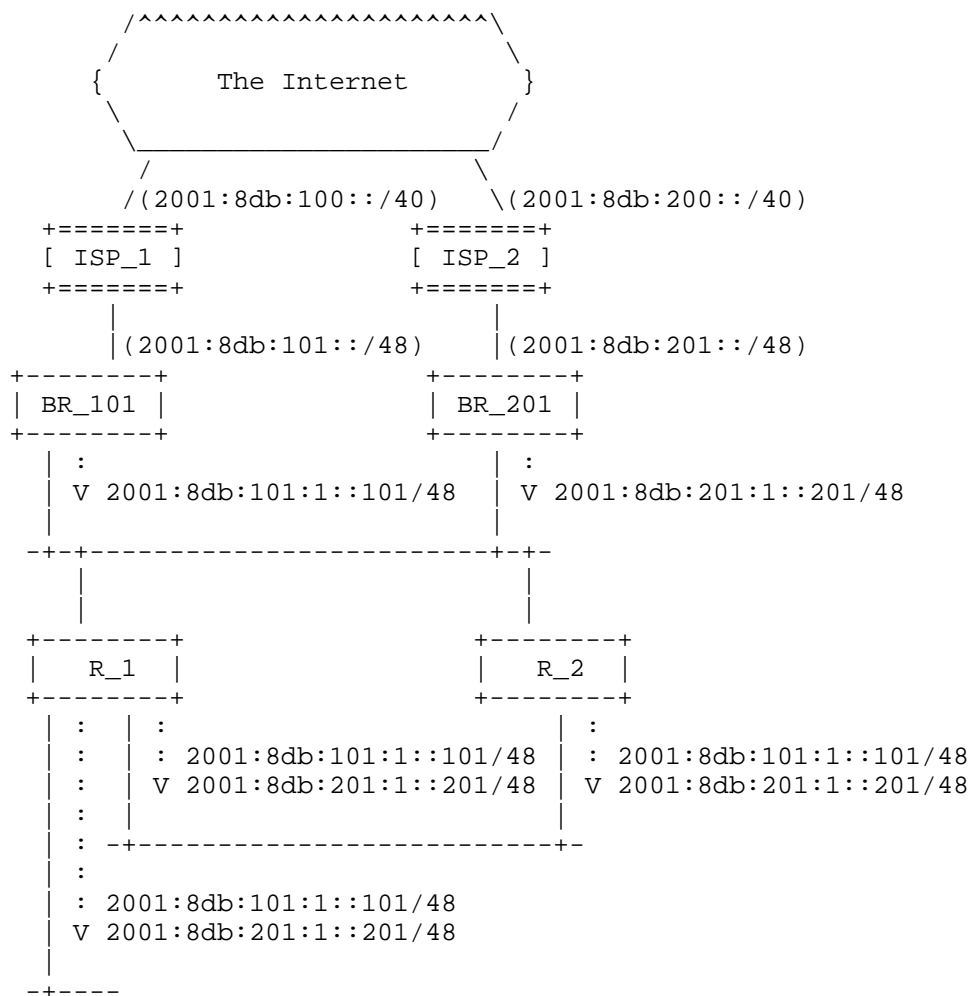


Figure 4: BRIO dissemination in Scenario 3

Routers are not required to configure global addresses on each interface. In the example, only the interface pointing to the Internet has configured global addresses. Routers may also use a (logical) management interface for global reachability.

So, the one-hop neighbours of BR\_101 and BR\_201, being R\_1 and R\_2, have learned the prefixes and configured addresses on their upstream interfaces. And all nodes in the network have learned the Border Router prefixes. The next step is to get configured addresses on the hosts in Figure 2. This is done by using DHCP Prefix Delegation. R\_1 and R\_2 request a prefix from either or both BR\_101 and BR\_201

for binding as on-link prefix on the links, and advertise those using Prefix Information Options to the hosts. This will result in a maximum of four prefixes that are advertised on the downlink of R\_1 and R\_2. Having multiple prefixes from the same ISP bound on a link is not useful. So a router requests a prefix from a Border Router only if no other prefix of that Border Router is advertised already by another router on this network segment.

In this example, R\_1 has been delegated two prefixes by DHCP PD for the link with host Host\_1234; 2001:8db:101:2::/64 and 2001:8db:201:3::/64. No other router is on this link. R\_1 or R\_2 has also been delegated a prefix on the link to host Host\_ABCD; 2001:8db:101:3::/64. It cannot be seen in Figure 2 which router has been delegated the prefix, nor if another prefix for this link has been delegated. No redundant prefix is delegated, as the routers learned with RA PIO already delegated prefixes for known Border Routers.

Now, Host\_1234 and Host\_ABCD can autoconfigure addresses for their interfaces. Host\_1234 configures two addresses, one for each Border Router. Host\_ABCD chooses not to use ISP\_2.

Nodes R\_1 and Host\_1234 can use both providers, by using two configured global addresses. Any multi-path facility can be used, either on an application layer or with a multi-path transport protocol.

Host\_ABCD may forward packets to the Internet via router R\_1 or R\_2. If R\_2 is selected as default router, R\_2 forwards the packets to BR\_101 as this Border Router corresponds to the prefix of the source address 2001:8db:101:3::ABCD. This works well, even in this case where R\_2 hasn't configured an address with a BR\_101 prefix for itself, and selected a global address from the BR\_201 prefix only.

#### 2.4. Medium multi-homed site with ULAs and DHCP server

In this example, the scenario 2 is extended by adding Unique Local Addresses (ULA) for communication within the site itself. For simplicity there is only one ISP present. The ULA IP configuration, with prefix fd00:8db::/48, is managed by DHCPv6 server DHCP\_201. The scenario is shown in Figure 5.

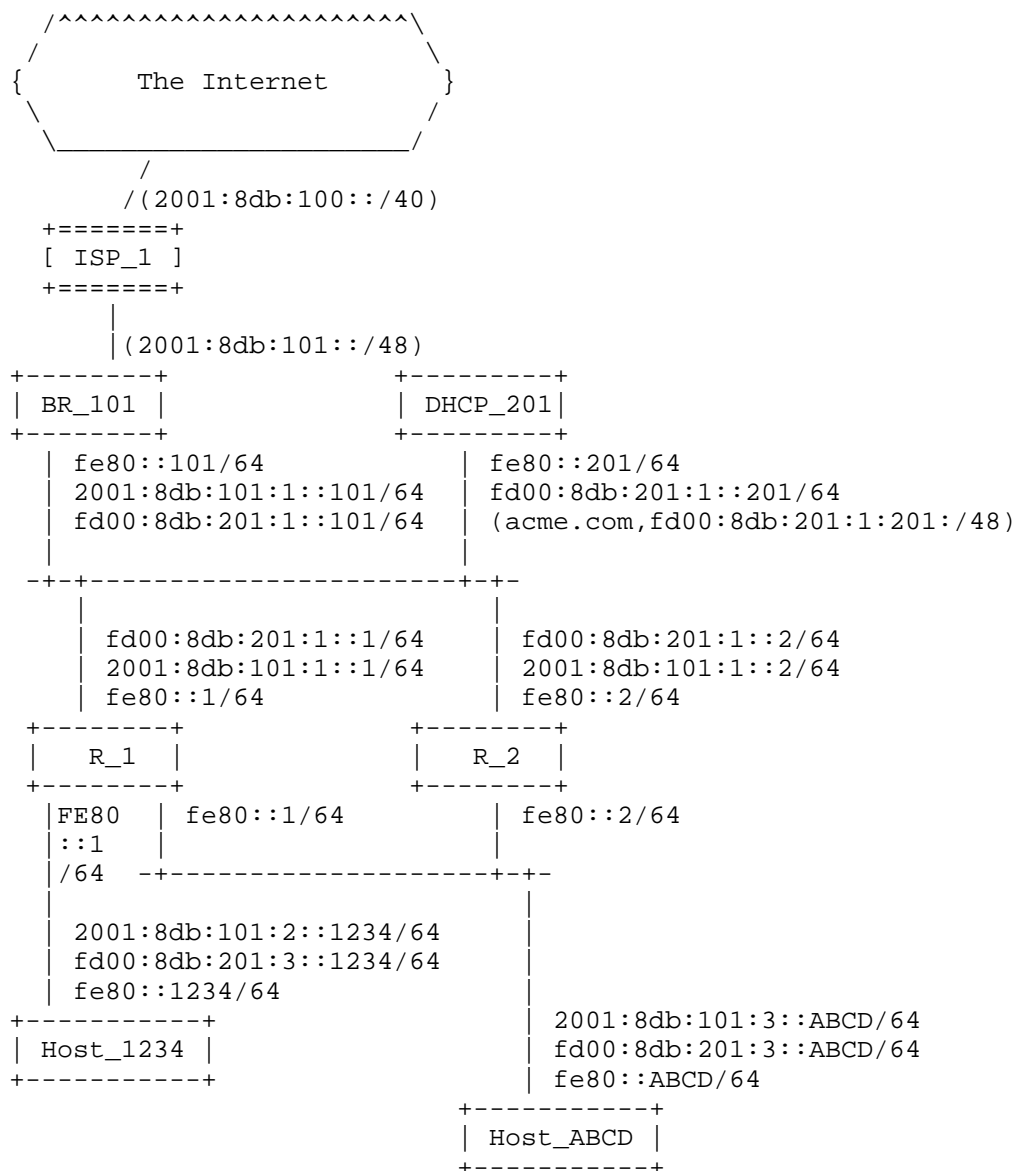


Figure 5: Scenario 4: a medium sized multi-homed site with ULAs and DHCP server.

In this scenario, all nodes have configured a ULA and a Global Unicast Address using prefix delegation in the way that was described in Section 2.2. ULA prefix delegation is automated just like PA addresses. The DHCP server is therefore implemented on a router, in

this case DHCP\_201. This router advertises the ULA prefix with BRDP, here fd00:8db:201::/48.

Although BRDP provides automatic prefix and address configuration for ULA, a network administrator is free to configure it manually, along using BRDP for global addresses.

BRDP based ULA configuration with BRDP based routing would result in routing packets with ULA destinations outside the site to the originator of the ULA prefix, in this case router DHCP\_201. DHCP\_201 is not connected to the Internet or another site owning the ULA, so packets to non-existing destinations are dropped. DHCP\_201 indicates such with the BRIO F-bit set, meaning the Border Router is floating.

This scenario, it is demonstrated that BRDP and DHCPv6 cooperate in address configuration. BRDP provides announcements of Border Routers and DHCP servers. Routers request prefixes with DHCP, and can request other parameters also. Such parameters are disseminated to other nodes, either with router advertisements or acting as DHCP server itself. Routers may also act as DHCP relay, redirecting address requests to the Border Router(s). The Router Advertisement M-bit and O-bit indicates availability of DHCPv6 services to attached nodes.

Difficulties may arise when both ULA and global addresses are used for Internet connectivity, e.g. when address translation is used. To distinguish, the Border Router not providing Internet connectivity informs nodes in the network using Service Selection suboption, similar to "Service Selection for Mobile IPv6" [RFC5149]. This procedure helps also for extranet connectivity. In this scenario, the ULA is used within the ACME Corporation, nodes are made aware by adding "acme.com" in the BRIO Service Selection Option.

It is for the reader to work out extensions for this scenario, where the ULA prefix originator is a Border Router to another site, e.g. a link from a branch office to a head quarter, or a ULA-only side connected to the Internet with NAT66.

## 2.5. MANET site

BRDP was developed for address autoconfiguration in MANETs. This scenario, see Figure 6 demonstrates the powerful multi-homing facilities provided to the MANET nodes.

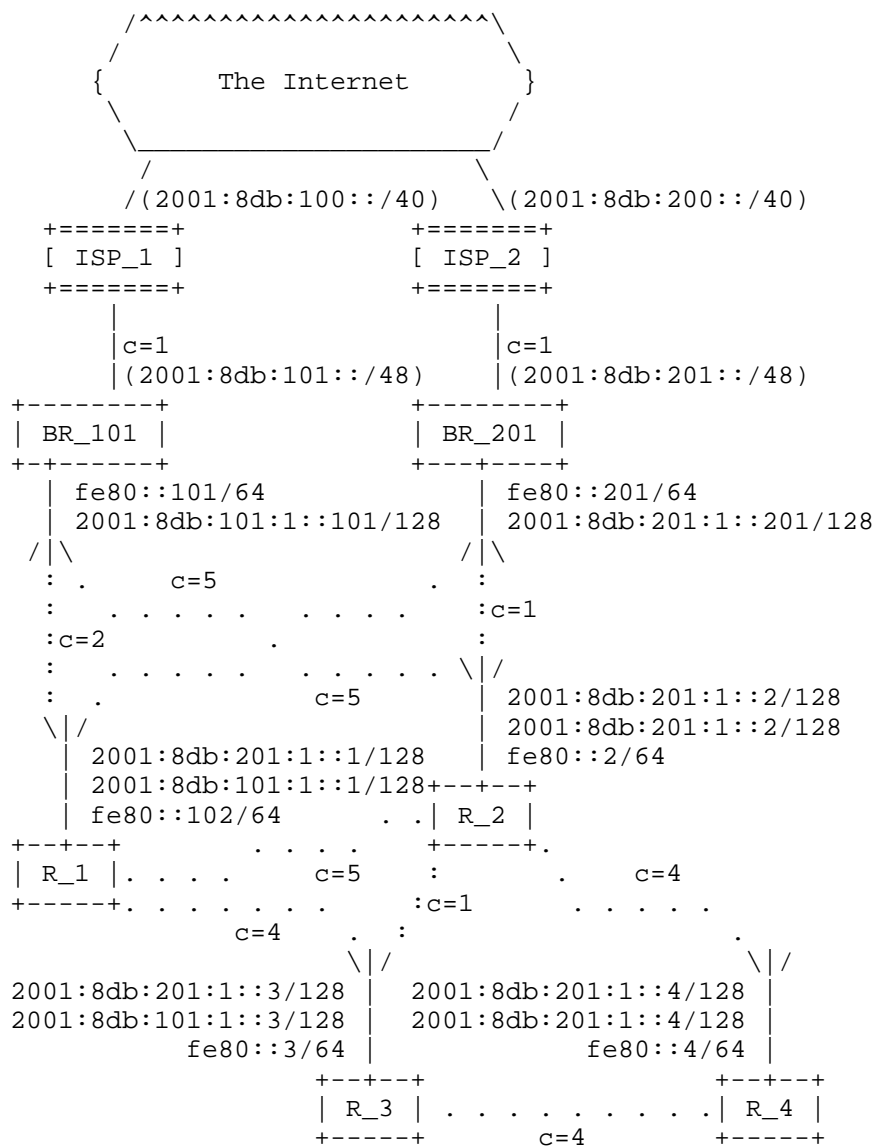


Figure 6: Scenario 5: a MANET site

On the MANET interfaces, addresses are configured using a 64-bit prefix provided by BRDP, appending it with a 64-bit Interface Identifier according to BRDP based address autoconfiguration. This creates a 128-bit prefix length as recommended in IP Addressing Model in Ad Hoc Networks [RFC5889]. Each MANET node has configured two global addresses, one for each ISP. With BRDP, the nodes are aware

of the cost of the path to the DFZ, defined as dimensionless metric for both directions of the patch. This enables optimized source address selection, and as an implicit result a Border Router and ISP selection. In the scenario, R\_1 is near to BR\_101 and the cost via this Border Router is lower than via BR\_201. The table below shows costs to the DFZ for all nodes, via both ISPs. Paths with lowest costs are marked with \*.

Costs to DFZ	Via ISP_1	Via ISP_2
BR_101	1*	7
BR_201	7	1*
R_1	3*	6
R_2	6	2*
R_3	7	3*
R_4	10	6*

The optimized source address selection facility is also of utility in the other scenarios. For example, the cost of the link to the ISP could be set depending of bandwidth and optionally on utilization. Nodes would use a near uplink to an ISP, and as a result some form of load distribution is enabled. Note that nodes still can use the alternative addresses, in fact it is recommended to use multi-path transport protocols for better load balancing and improved robustness.

For isolated MANETs, a DHCP server election mechanism can be used. Nodes may initiate to advertise a self-generated ULA. In such cases, it is recommended that a prefix is used with a 56-bit random ULA identifier (including random 16-bit Subnet ID) and 64-bit prefix length. Other nodes join this prefix, although some may wish to start or continue using their own prefix. The latter would occur in cases of a merge of previous isolated MANETs.

### 3. Border Router Discovery Protocol (BRDP)

BRDP is an extension to the IPv6 ND mechanism [RFC4861] that provides information about the reachability, availability, prefix information, quality and cost of upstream providers, and enables automated (re)numbering of possibly multi-homed routers and hosts.

BRDP adds the Border Router Information Option (BRIO) to the Router Advertisement (RA) of IPv6 ND. A BRIO contains all relevant information of an upstream Border Router and the corresponding

provider.

Border Routers initiate sending BRIO messages, other routers in the network disseminate the messages downstream through the network. Nodes store the information from received BRIOS in a BRIO cache, to be used for address generation, DHCP server discovery, address selection or packet forwarding.

A BRIO cache entry records reception of a BRIO for a single advertised prefix, received via a neighbor router. Border Routers that need to advertise multiple prefixes simply use multiple BRIOS, each with its own address prefix. For further processing of BRIO entries, only the entry with the lowest cost to a Border Router is used, for each Border Router.

When a node is multi-homed, it will receive BRIOS from multiple upstream Border Routers.

### 3.1. Border Router Information Option (BRIO)

The Border Router Information Option carries information that allows a nodes in the edge network to select and utilize a Border Router.

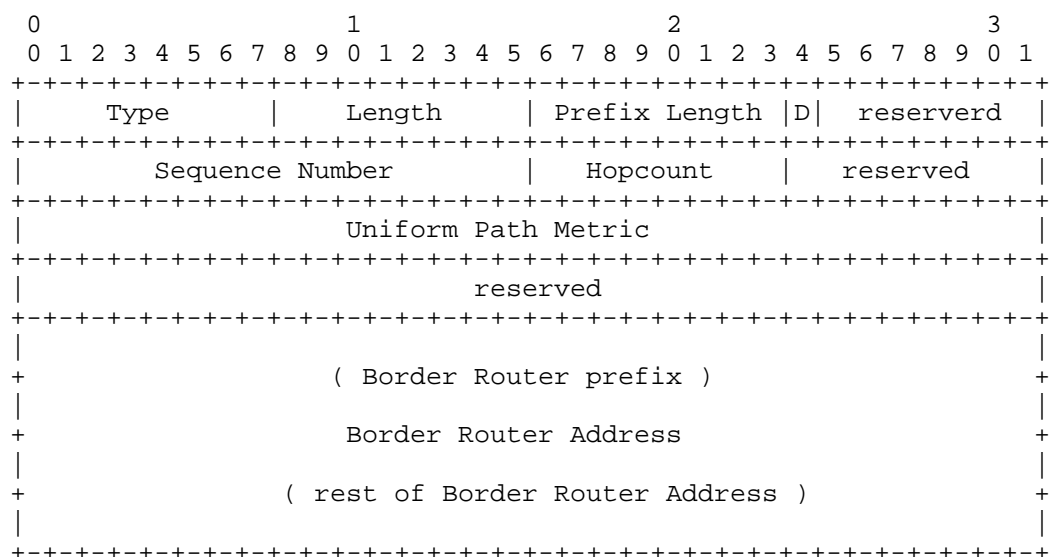


Figure 7: BRIO base option

Fields:

**Type:**

8-bit identifier of the Border Router Information Option type.  
The value of this option identifier is to be determined.

**Length:**

8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 octets. A BRIO has a length value of 4.

**Prefix Length:**

8-bit unsigned integer. The number of leading bits in the Border Router Address, that indicates the assigned prefix for that Border Router. The Prefix Length is used for BRDP Based Routing, as described in Section 6.

**DHCP (D):**

When the D-flag is set, the Border Router is acting as a DHCP server or DHCP relay agent [RFC3315].

**reserved:**

Reserved bits. Currently unused, set to 0.

**Sequence Number:**

16-bit unsigned integer. It is set by the Border Router and incremented with each new BRIO it sends on a link. When forwarding downstream, the sequence number is not changed.

**Hopcount:**

8-bit field registering the number of hops from the advertizing Router to the Border Router. Border Routers send the initial BRIO with its Hopcount set to zero. Routers increment the Hopcount by one when forwarding a BRIO.

**Uniform Path Metric (UPM):**

A measure for the cost of the bi-directional path between the upstream Router and the Default Free Zone of the Internet. Uniform Path Metric is set to some initial value by the Border Router and is incremented by each Router forwarding the BRIO.

**Border Router Address:**

128-bit address of the Border Router. For reachability, the Border Router is expected to add this own address (prefix) in the routing system.

### 3.2. BRDP processing

The main BRDP processing functions of a Router are BRDP message generation, transmission and reception and the maintenance of a BRIO-

Cache. Routers forward BRDP messages using ICMP ND Router Advertisements.

### 3.2.1. BRDP message generation and transmission

A BRDP message is part of a Router Advertisement and includes a set of BRIOs. It provides the current state of (paths to) the Border Routers listed in the set of BRIOs. BRIOs originate from a Border Router, and contain initially metric information on connectivity to the Internet. BRIOs are forwarded downstream in the edge network.

When a Router sends a ICMP ND Router Advertisement, it SHOULD include a set of BRIOs by appending them to the message. The maximum number of BRIOs in a single BRDP message is a Router configuration parameter. BRIO selection for advertisement is done based on the information stored in the BRIO-Cache. BRIOs that do not pass the loop prevention check described in Section 3.2.4 SHOULD NOT be selected.

The UPM and Hopcount fields of the advertised BRIOs are updated. An UPM-increment, based on uniformed bi-directional link metrics, is added to the UPM and the Hopcount is incremented by 1. UPM-increment MAY be governed by a hysteresis and dampening mechanism. Also forecasted information MAY be used.

Each BRIO originating from a Border Router has an increased Sequence Number. This BRIO is forwarded in the edge network and refreshes entries in BRIO-Caches of downstream Routers.

Router Advertisements are sent in response to Router Solicitation messages or unsolicited with a uniformly-distributed random interval between MinRtrAdvInterval and MaxRtrAdvInterval [RFC4861]. The MaxRtrAdvInterval falls between a minimum of 30 milliseconds, specified in [RFC6275] and a maximum of 1800 seconds, specified in [RFC4861]. In addition, the Router MAY send a Router Advertisement when an important change in a to be sent BRIO would occur.

When a Router sends Router Advertisements more frequently than an upstream Router, this Router MAY repeatedly send BRIOs with a constant Sequence Number but with an updated UPM or Hopcount.

The ICMP ND Router Advertisement MAY include the Advertisement Interval Option [RFC6275]. This option contains the interval at which the sending router sends unsolicited multicast Router Advertisements.

A Router SHOULD inform downstream Routers in case the path to a previous advertised Border Router is lost, by at least 3 times

retransmitting the previously sent BRIO with a UPM value of 4294967295.

In case a Border Router loses its connection to the infrastructure it will lose its Border Router functionality and become a normal Router. In that case it performs the same procedure as a Router that has lost the path to a previous advertised Border Router.

For each Border Router listed in the BRIO-Cache, the UPM-loop-prevention-threshold and the Hopcount-loop-prevention-threshold variables are maintained. These variables are used by the loop prevention mechanism described in Section 3.2.4. The thresholds are set or updated when sending BRDP messages. When sending a BRIO with a higher Sequence Number than the previously sent BRIO for that Border Router, the threshold variables are set to the UPM and Hopcount values in BRIO to be sent. When sending a BRIO with the same Sequence Number as the previously sent BRIO, the loop-prevention-thresholds are independently updated if either the UPM or Hopcount of the outgoing BRIO is lower than their thresholds.

A Router that detects an attractive candidate BRIO but is prohibited from using it because of the loop prevention check, MAY send a (unicast) Router Solicitation message to the Border Router. The Border Router responds to such a Router Solicitation message with a new BRIO. Sending Router Solicitations MUST be rate limited. A next version of this document would include a specification for sending the unicast Router Solicitation message.

### 3.2.2. BRDP message reception

When a BRDP message is received, the Sequence Number fields of the contained BRIOS are checked; the Sequence Number of a received BRIO MUST be equal to or higher than the Sequence Number in the cache for an existing entry in the cache, with wrap-around checking. Otherwise, the BRIO will be discarded.

BRIO messages do not need to be forwarded at fixed time intervals, because the RA intervals on different Routers are not synchronized. Therefore, large gaps in Sequence Numbers may occur. Increment values between 0 and 65000 are accepted. Increment values between 65001 and 65535 are rejected.

Information in received BRIOS is stored in a BRIO-Cache table. Other information is stored as well, such as the BRIO upstream node, a timestamp indicating when the most recent message was received and the measured or signaled RA interval.

### 3.2.3. BRIO-Cache maintenance

Each Router maintains a BRIO-Cache that stores all information on Border Routers. Unique cache entries are maintained on (Border Router Address, address of the upstream router that forwarded the BRIO) tuples. This information is obtained by receiving BRIOs, or, in case of a Border Router, by getting information from the interface that connects to the Internet. The BRIO-Cache also maintains context information for the BRIO such as the BRIO sender, link metrics and UPM-increment for this sender, history, statistics and status information. History information includes a timestamp indicating when the most recent message was received and a measured or signaled RA interval. Status information includes the BRIO selection outcome for BRIO forwarding as explained in Section 3.2.1 and the Border Router selected for address generation as explained in Section 4.

BRIO entries in the BRIO-Cache stay valid for a certain period of time. During this period, they can be used for Border Router selection by the Router, for forwarding BRIOs and for address generation. BRIO-Cache information could also be useful for source address selection [RFC6724]. The lifetime of a BRIO is determined by using the timing information sent along with the RA ([RFC6275], section 7.3) or statistics of received BRIOs.

Some values in the BRIO-Cache can be updated independent of incoming BRDP messages. A Router MAY update the UPM-increment based on link quality measurements performed in an environment with changing link metrics. A Router SHOULD indicate in its BRIO-Cache which BRIO entries are currently selected for forwarding and for address generation. Border Router Selection MAY take place after the UPM of a BRIO entry has been updated.

In case the link to the Router from which a BRIO has been received is broken, the UPM and the Hopcount of the BRIO entry in the cache are set to the maximum value, i.e. 4294967295 and 255.

A cache cleanup routine SHOULD run at regular intervals to get rid of stale entries. Stale entries are removed when the entry is not updated for 5400 seconds or all of the following conditions are met:

- o The stale entry is not used by the Router itself for address generation.
- o The stale entry was not selected for forwarding in the last three Router Advertisement.
- o The stale entry was not recently updated by a received BRIO. In this context, recently is defined as the maximum of a) three times its own unsolicited multicast Router Advertisements interval and b) three times the senders unsolicited multicast Router Advertisements interval.

Cache entries MAY also be removed, under the condition that the BRIO-Cache has reached a configured maximum number of entries and a new, to be stored BRIO is received. A removal candidate is selected based on:

- o The candidate entry is not used by the Router itself.
- o The candidate entry was not selected for forwarding in the last Router Advertisement.
- o The candidate entry is redundant; other information for the same Border Router is stored in the cache with a better UPM and / or was received more recently.
- o The candidate entry is redundant; other information for the same Service Selection Identifier is stored in the cache with a better UPM and / or was received more recently.
- o The candidate entry is less attractive; other Border Routers are stored in the cache with better UPM and / or were received more recently.

#### 3.2.4. BRDP loop prevention

A BRDP loop check mechanism prevents that a Router forwards an earlier advertized BRIO.

BRDP loop-free operation is guaranteed as long as at least one of the following conditions is true:

- o The to be sent BRIO has a higher Sequence Number than a BRIO for this Border Router that was sent before. The loop check mechanism uses wrap-around logic. Increments up to 32768 are acceptable (wrap-around logic needs checking).
- o The to be sent BRIO is generated from the same BRIO-Cache entry as the BRIO that was sent most recently.
- o The to be sent BRIO has the same Sequence Number as the BRIO for this Border Router that was sent before but the BRIO-Cache entry UPM is equal to or lower than the UPM-loop-prevention-threshold for this Border Router.
- o The to be sent BRIO has the same Sequence Number as the BRIO for this Border Router that was sent before but the BRIO-Cache entry Hopcount is equal to or lower than the Hopcount-loop-prevention-threshold for this Border Router.

In some circumstances, a Router would select a BRIO for forwarding that fails the loop prevention check. For example, the link to the upstream neighbor is lost and an alternative path is available, with a higher UPM and a higher Hopcount or with a lower Sequence Number. The Router cannot assure this candidate BRIO is not reflecting its own advertized message, therefore it should not send this BRIO. Instead, it sends a unicast Router Solicitation message to that Border Router.

### 3.3. Unified Path Metric (UPM)

Unified Path Metric (UPM) is a measure for the cost of the path between the Router and the Internet Default Free Zone. It is a united metric for both inbound and outbound paths. On each hop, the UPM is incremented with an UPM-increment, which is derived from the routing protocol and / or is obtained from lower layers.

It is on forehand not known what is more important; Border Router selection based on path metric to the Border Router or the path metric for the reverse path. In BRDP, UPM is used for optimizing Border Router selection for both the inbound and the outbound traffic. Note that actual traffic will use the path provided by the routing protocols, not by BRDP.

Since the UPM uses 32 bits, its maximum value is 4294967295. On each hop, an UPM-increment is calculated for each Router from which a BRIO has been received. UPM-increments have a value between 1 and 16777215, to support a 255 hop path, with maximum UPM increments.

Further discussion on metrics and how the UPM-increment value is determined is outside the scope of this document.

## 4. BRDP based Address Configuration and Prefix Delegation

BRDP supports stateless address autoconfiguration [RFC4862], DHCP managed IP configuration [RFC3315] and DHCP Prefix Delegation [RFC3633]. Routers can also use a variant of stateless address autoconfiguration, where BRDP provided information is used to configure Router management interfaces or used to configure off-link addresses, used in ad hoc networks [RFC5889].

BRDP adds topology awareness in address configuration. Nodes can configure multiple addresses, each to support a different facility. ULAs can be used for site internal traffic. Global addresses are mandatory for access to the Internet, assuming address translation is not used.

A node that is offered multiple prefixes for stateless address autoconfiguration or multiple addresses by DHCP chooses to configure one or more addresses. BRDP provides information for the candidate addresses. An important criterium is the costs of the path to the Internet DFZ. A node would prefer addresses with lower costs.

BRDP does not modify stateless address autoconfiguration and DHCP protocols, except that in a edge network, Routers may perform stateless address autoconfiguration from the Border Router

Information Option (BRIO), for their management or MANET interfaces. This enables edge network-wide address configuration, because BRIOs are disseminated over multiple hops in the edge network, while PIOs are link local messages only.

When a BRIO is stored in the BRIO cache table, the node checks if a corresponding address already exists for the Border Router from which this BRIO originates. If not, and a corresponding address for that Border Router is beneficial, address generation for that Border Router is triggered.

#### 4.1. Border Router selection

When a node needs to communicate to nodes on the Internet, it MUST select a (set of) Border Router(s) for address generation. A node MAY generate multiple addresses for smooth handover implementing make-before-break or distributing traffic over multiple Border Routers. A description how Border Routers can be used concurrently is out-of-scope for this document.

Information concerning available Border Routers is kept in the BRIO-Cache.

The Border Router selection mechanism MAY be triggered by received BRDP messages, changes in metrics on links to neighbors advertising BRDP messages, changes in costs to Border Routers used or on a time-driven basis.

The Border Router selection algorithm SHOULD be based on UPM. UPM is used for selecting the Border Router with the best connectivity to the Internet. The Border Router selection algorithm MAY be extended with any other information. Future defined BRIO suboptions could provide additional information, such authorization and service selection. Border Router selection MAY be based on the type of the Border Router Address, e.g. a globally unique address or a unique local address.

Border Router selection does not provide nor select a routing path to the Border Router.

##### 4.1.1. Border Router Selection based on UPM

The node uses the UPM for Border Router selection preferring the best bi-directional paths between the node and the Internet. Note that the BRIO UPM includes the initial metric set by the Border Router and is not solely a metric between the node and the Border Router. The initial metric set by Border Routers can be used for Border Router preference and for load balancing.

In order to use an up-to-date UPM in the selection procedure the UPM-increment is calculated by the node before selecting a Border Router. UPM is discussed in Section 3.3.

#### 4.2. Address autoconfiguration

Nodes should use a topologically correct address when communicating with corresponding nodes on the Internet. Topologically correct addresses should be configured for each Border Router used.

##### 4.2.1. Address and prefix configuration with SLAAC or DHCP

Nodes can use existing IPv6 address configuration protocols, such as SLAAC [RFC4862] and DHCP [RFC3315]. Nodes can use SLAAC based on prefix information, provided by the upstream router. Nodes may use DHCP multicast and neighbor routers will relay those packets to selected Border Routers with D-flag set or reply with DHCP parameters it has received from a Border Router before for itself.

Nodes using SLAAC may also query a DHCP server on a Border Router themselves for additional parameters, using the BRDP learned address of the DHCP server.

A Router should request a prefix for attached subnetworks, with DHCP-PD [RFC3633], where there is at that moment no on-link prefix for a selected Border Router.

##### 4.2.2. Address generation and configuration for Routers

A generated address for a Router management interface or a MANET has a /128 prefix. It is constructed from a 64-bit Interface Identifier and a 64-bit prefix from the Border Router Address. The generated 128-bit address SHOULD be advertised in the routing system. The generated address may be used for user traffic, either inside the edge network or traffic towards the Internet.

For the Interface Identifier used, the BRDP-based Address Generation MUST implement a mechanism for generating a highly unique Interface Identifier. Known mechanisms are:

- o Modified EUI-64 format-based Interface Identifier, [RFC4291], based on IEEE 802 48-bit MAC address or IEEE EUI-64 identifier. However, this method does not guarantee identifiers are unique as duplicate MAC addresses can occur.
- o Generation of randomized Interface Identifiers, [RFC4941].
- o Well-distributed hash function, [RFC3972].

After Address Generation, RFC4429 Optimistic Duplicate Address Detection [RFC4429] should be used. A passive Duplicate Address

Detection, based on information in the routing protocol information bases could be used as an alternative. Still, uniqueness is not fully guaranteed. Main reasons for non-uniqueness are merging of edge network segments, node movement, node misbehavior or address spoofing attacks. Details on handling a duplicate address condition are out-of-scope for this document.

A generated addresses clean-up routine SHOULD run at regular intervals to get rid of stale addresses.

#### 4.2.3. Support for Unique Local Addresses (ULA)

Address generation for globally unique addresses and unique local addresses (ULA) [RFC4193] is equivalent. If no BRIO for a unique local addresses is available, a router may start as a Border Router and DHCP server for a self generated 48-bit ULA prefix.

### 5. BRDP based Source Address Selection

As a next step, multi-homed nodes perform source address selection for new, self-initiated connections. The algorithm described in Default Address Selection for IPv6 [RFC6724] uses the concept of a "candidate set" of potential source addresses. Rule 8 of source address selection is "Uses longest match prefix". The goal of this rule is to select the address with good communications performance. If other means of choosing among source addresses for better performance is available, that should be used.

BRDP provides attributes for prefix, such as a cost metric to the Internet. This information can be used to select the "best" source address. For multi-path transport protocols, it is also important to have a mechanism to select alternative addresses. For example, rule 4 gives preference to a Home Address. Alternate addresses can be used for route optimization and to avoid overhead of the Mobile IP tunnel.

#### 5.1. Address Selection for dynamic DNS

BRDP provided information can also be utilized by address lookup protocols such as DNS. A node can register its addresses dynamically, with support of preference and load balancing if the mechanism used support such.

### 6. BRDP based Routing

BRDP introduces a new paradigm for packet forwarding for multi-homed

sites, where forwarding to a default gateway is replaced by source address based forwarding towards a corresponding Border Router. This enforces that packets will be sent via the selected upstream provider, without the need of tunneling. As such, it prevents problems with ingress filters in multi-homed edge networks [RFC3704].

The BRDP Based Routing mechanism provides basic support for load distribution over multiple Border Routers. BRDP Based Routing forwards the packets to the Border Router that corresponds with the source address. As a result, nodes can utilize multiple paths, if available. Standardization of this load balancing functionality is work in progress in the IETF MPTCP working group.

When a router forwards a packet to a next-hop node, via the interface where this packet was received, and the next-hop address was selected using BRDP based routing, then the router should not send an ICMP redirect message to that host. This is because the upstream node would cache the redirect for the destination address, while the forwarding decision was based on the source address.

#### 6.1. Problems with default gateway routing

Usually, the nexthop selection is based on the destination address. In case of default gateway routing and multiple exit routers to multiple providers, the source has no influence on what exit router is used. In case of ingress filtering and lack of a mechanism to redirect packets to exit routers that correspond to the source address, packets may be dropped.

This default gateway routing behavior blocks incremental enhancement of the Internet, e.g. through the addition of support for more dynamic networks and / or host based load distribution mechanisms. In a MANET, it also prevents the use of make-before-break [RFC3753] mechanisms.

#### 6.2. Default gateway routing replaced with BRDP Based Routing

Default gateway based routing for IPv4 is defined in [RFC1812], section 5.2.4.3:

- (5) Default Route: This is a route to all networks for which there are no explicit routes. It is by definition the route whose prefix length is zero.

With BRDP Based Routing, another type of route is introduced:

- (6) BRDP Route: This is a route to all networks for which there are no explicit routes, and a default route is not used. The nexthop IP address is found by means of a Border Router Information Cache (BRIO-Cache) lookup based on the source address and, if a matching BRIO-Cache entry is found, a subsequent FIB lookup based on the selected Border Router address.

Note that route types (3) and (4) are not defined in RFC1812.

BRDP Based Routing can be turned on and off with the existence of a default route in the IGP. This switch function might be useful in migration scenarios towards BRDP Based Routing.

The Border Router should run the IGP on the interface with the BRDP advertized Border Router address, to make sure this address is reachable in the edge network.

In the edge network, all interior routers should run BRDP and BRDP Based Routing. All interior routers will have a BRIO-Cache with information for selecting Border Routers as exit points to the Internet. A BRIO-Cache entry contains a Border Router address and a summary prefix assigned to that Border Router. BRIO-Cache lookup follows the longest-match rule.

Forwarding is solely based on FIB lookups, the nexthop IP address is found either by a FIB lookup with the destination address or by a FIB lookup with the address of the Border Router that corresponds with the source address. If the nexthop IP address lookup fails, the packet is discarded.

## 7. BRDP and IRTF RRG goals

The IRTF Routing Research Group (RRG) was chartered to explore solutions for problems on routing and addressing, when the Internet continues to evolve. It has explored a number of proposed solutions, but did not reach consensus on a single, best approach [I-D.irtf-rrg-recommendation]. In fulfillment of the routing research group's charter, the co-chairs recommend that the IETF pursue work in three areas, "Evolution" [I-D.zhang-evolution], "Identifier/Locator Network Protocol (ILNP)" [I-D.rja-ilnp-intro] and "Renumbering" [RFC5887]. BRDP fits in all three approaches.

BRDP is an evolution in IPv6 address configuration and address selection, as well as forwarding to destinations outside the routing domain. As a result, it removes a demand for Provider Independent

addresses for (small) multi-homed edge networks. BRDP enables sites to use multiple Provider Aggregatable address blocks, while being able to utilize multi-homing for improved redundancy of communications and enlarged capacity. Each site that continues to use Provider Aggregatable addresses when getting multi-homed, instead of using its own Provider Independent address space, reduces the growth of the routing tables in the Default Free Zone.

BRDP can cooperate or live next many other solutions. ILNP is a good example for cooperation, BRDP provides multi-path transport capabilities to ILNP nodes. This multi-path transport capability applies to many other approaches also, such as map&enap and nat66.

Because BRDP provides automatic address and prefix configuration, Renumbering is far less problematic. That said, legacy (IPv4) hosts, applications and network equipment is not BRDP enabled and manual address configuration will be used for many years to come.

In Design Goals for Scalable Internet Routing [I-D.irtf-rrg-design-goals], a number of design goals are defined. The role BRDP can play for these goals are briefly described in the next sections.

#### 7.1. Scalability

Because BRDP is implemented in edge networks, and not in the core, scalability of BRDP is less an issue. BRDP solves the Internet routing problem at the source, by reducing the demand for PI addresses.

#### 7.2. Traffic engineering

BRDP provides traffic engineering options to end-nodes. End-nodes can configure multiple addresses and use these for utilizing multi-path capabilities of the network. Using multi-path is being worked on by the IETF MPTCP working group.

#### 7.3. Multi-homing

The core function of BRDP is providing support for IPv6 multi-homing, without any problems caused by ingress filtering [RFC3704].

#### 7.4. Loc/id separation

BRDP does not mandate any approach for location / identification. For packet forwarding, addresses are used as locator. If addresses are used as identifiers also, for example in Mobile IP, BRDP supports route optimization where traffic uses the Home Address as identifier

and care-of addresses as locator. MPTCP provides the route optimization capability.

#### 7.5. Mobility

BRDP was defined as a solution for address autoconfiguration for ad hoc networks. With BRDP, nodes can easily configure topology correct addresses in a multi-homes ad hoc network. BRDP does not provide session continuity functions. Mobility solutions are already in place or new approaches are proposed. All approaches should work well with BRDP, as BRDP does not modify the IPv6 protocol.

#### 7.6. Simplified renumbering

BRDP makes site renumbering fully automatic. This applies to node address configuration on the IPv6 stack and prefix delegation and configuration on routers. IP addresses could be configured on many other places, either manually or using specific protocols for such purpose. Complete automatic numbering is possible if all mechanisms in use support dynamic addresses. There is definitely more work to do [RFC5887].

#### 7.7. Modularity

BRDP is a small, but important piece of the puzzle. It applies to edge networks only. It helps other mechanisms to work well in a multi-homed network using PA addresses, but also provides multi-path capabilities in multi-homed networks with PI addresses or multi-homing with connections to Extranets.

#### 7.8. Routing quality

BRDP is not a routing protocol, so it has no influence on routing quality. But the functionality of routing to a default gateway is changed. BRDP based routing supports paths to multiple Border Routers, where hosts can select which Border Router to use. In such scheme, nodes can select the route to use, based on quality of available routes. MPTCP provides this route selection functionality.

#### 7.9. Routing security

BRDP doesn't update any routing protocols. BRDBP based routing modifies the default gateway heuristic, the route to prefix `::/0` is replaced by a route to a Border Router, which corresponds with the source address of a packet. As a result, ingress filtering is distributed over all routers in the edge network and invalid packets are dropped as near to the source as possible.

The BRDP protocol runs on IPv6 NDP and inherits all security aspects. BRDP messages are disseminated in the edge network, which may enlarge the needs for protection. Implementing SeND [RFC3971] is recommended.

#### 7.10. Deployability

BRDP deployment takes place edge network by edge network. Each network that migrates to BRDP, instead of getting a PI address block, reduces the load on the Internet routing infrastructure.

For implementing BRDP on an edge network, all routers in the network must support BRDP. BRDP support for hosts is optional. Enterprise networks can migrate site by site.

#### 8. Currently unaddressed issues

BRDP based routing may have impact on multicast routing, e.g. selecting the route to a RP.

It is not fully understood how BRDP may influence host behavior on RA M and O bits, and may bypass a 1-hop router DHCP relay server for getting information for a BRDP-learned DHCP server.

Currently unaddressed issues are to be addressed in a next version of this document.

#### 9. Acknowledgements

BRDP is inspired by MANEMO technology; thanks to all who contributed to it. Thanks to Arjen Holtzer (TNO), co-author of earlier Internet drafts on BRDP. Thanks to Ran Atkinson, who guided me towards a BRDP Based Routing mechanism that does not rely on routing headers or encapsulation.

#### 10. IANA Considerations

TBD

#### 11. Security Considerations

TBD

## 12. Change log

This -00 version is gathering the material of BRDP, produced for Autoconf and RRG. It is a bit cleaned up, with removal of some details for MANET and with removal of options for emergency services, service selection and authorization.

## 13. References

### 13.1. Normative References

- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3753] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [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.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, April 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

## 13.2. Informative References

- [I-D.ietf-homenet-arch]  
Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil,  
"Home Networking Architecture for IPv6",  
draft-ietf-homenet-arch-04 (work in progress), July 2012.
- [I-D.ietf-mptcp-multiaddressed]  
Ford, A., Raiciu, C., Handley, M., and O. Bonaventure,  
"TCP Extensions for Multipath Operation with Multiple  
Addresses", draft-ietf-mptcp-multiaddressed-10 (work in  
progress), October 2012.
- [I-D.irtf-rrg-design-goals]  
Li, T., "Design Goals for Scalable Internet Routing",  
draft-irtf-rrg-design-goals-06 (work in progress),  
January 2011.
- [I-D.irtf-rrg-recommendation]  
Li, T., "Recommendation for a Routing Architecture",  
draft-irtf-rrg-recommendation-16 (work in progress),  
November 2010.
- [I-D.kline-default-perimeter]  
Kline, E., "Default Perimeter Identification",  
draft-kline-default-perimeter-00 (work in progress),  
July 2012.
- [I-D.rja-ilnp-intro]  
Atkinson, R., "ILNP Concept of Operations",  
draft-rja-ilnp-intro-11 (work in progress), July 2011.
- [I-D.zhang-evolution]  
Zhang, B. and L. Zhang, "Evolution Towards Global Routing  
Scalability", draft-zhang-evolution-02 (work in progress),  
October 2009.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering:  
Defeating Denial of Service Attacks which employ IP Source  
Address Spoofing", BCP 38, RFC 2827, May 2000.
- [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.

- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC5149] Korhonen, J., Nilsson, U., and V. Devarapalli, "Service Selection for Mobile IPv6", RFC 5149, February 2008.
- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.
- [RFC5889] Baccelli, E. and M. Townsley, "IP Addressing Model in Ad Hoc Networks", RFC 5889, September 2010.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

#### Author's Address

Teco Boot  
Infinity Networks B.V.

Email: [teco@inf-net.nl](mailto:teco@inf-net.nl)



Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: April 3, 2016

W. Haddad  
Ericsson  
D. Saucez  
INRIA Sophia Antipolis  
J. Halpern  
Ericsson  
October 1, 2015

Multihoming in Homenet  
draft-haddad-homenet-multihomed-06

Abstract

Multihoming becomes popular in residential and SOHO networks indicating the absolute necessity of fully supporting multihoming in Homenet. While the approach followed in Homenet is to delegate multihoming management to hosts, we propose to enable multihoming in Homenet by the mean of the infrastructure instead of the hosts.

Status of This Memo

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

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

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

This Internet-Draft will expire on April 3, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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

#### Table of Contents

1. Introduction . . . . .	2
2. Homenet multihoming without host involvement . . . . .	3
3. Requirements . . . . .	3
4. Homenet multihoming with MSP . . . . .	4
5. Security Considerations . . . . .	6
6. Conclusion . . . . .	6
7. Normative References . . . . .	6
Authors' Addresses . . . . .	7

#### 1. Introduction

So far, multihoming in Homenet must be supported by the hosts with solutions like Shim6 [RFC5533] or MPTCP [RFC6182] as there is no mean to use simultaneously the different ISPs of the Homenet without risking flow disruption. In this memo, we propose the creation of a new multihoming service for Homenets. The concept relies on a middlebox added between the home network and its gateways with the ISPs. On the one hand, this middlebox is in charge to redirect the home network traffic to a multihoming service provider (MSP) by selecting the most appropriate Homenet's ISPs. On the other hand, the MSP is in charge of attracting traffic normally destined to the home network and then, the MSP can eventually redirect the traffic to its final destination, the Homenet itself, such that it enters the Homenet via the most appropriate ISP.

Section 2 describes the multihoming problem in Homenet when hosts cannot support it directly. Section 3 gives the necessary requirements. Section 4 sketches a possible solution to that problem.

## 2. Homenet multihoming without host involvement

It is known that multihoming reduces costs for ISPs by allowing higher aggregated bandwidth, better quality of service, and higher robustness.

Alternatively, the access to multiple ISPs at the same time for residential and SOHO users is now a reality, e.g., ADSL + Cable + 4G, but there is currently no simple solution for home networks to exploit it. For now, the only solution is to modify end-hosts with protocols such as Shim6 or MPTCP in order for hosts to change IP addresses on elapsing communications.

We claim that multihoming for Homenets will become a reality and will provide the same benefits as those observed for the ISPs. Also, requiring every single device in the Homenet to be modified to support multihoming is not acceptable as some devices have limited resources and cannot achieve it correctly and also because it would dramatically slow down the adoption of multihoming in the Homenet. Finally, letting every device deciding of the routing strategy (e.g., shall I route my traffic via left or right ISP?) might cause management issues.

At the light of this, the question can be: How can we achieve multihoming in Homenets, without changing neither the devices connected to the Homenet, nor the protocols and operations of the Homenet's ISPs?

## 3. Requirements

In order to fix the solutions space of our problem, we have isolated four requirements.

As we are in the context of Homenet, requirement (1) is to have zero configuration need at the Homenet user level. Multihoming must be transparent for users and devices.

Also, residential and SOHO network operators (i.e., John/Jane Does) seldom have enough power to make specific settlements or negotiations with their ISP, the solution thus have to be completely independent of the network's ISPs and the ISPs cannot have any mean to forbid the solution. Requirement (2) is thus ISP independence.

Multihoming offers the possibility to implement policies, and to some extend even capabilities, at any arbitrary level. For example, the home network can determine the number of ISPs it is using simultaneously or limit flows for example to only go via one

particular ISP at a given speed. Requirement (3) is thus policies/capabilities.

Finally, and this is related to policies and capabilities, the system must be able to provide quality of service (to some extent) to ensure Quality of Experience. We call the requirement (4) Quality of Service.

#### 4. Homenet multihoming with MSP

To offer fast and efficient deployment of multihoming in residential and SOHO networks, a dedicated middlebox is added to be in charge of dealing with multihoming, on behalf of the devices. This middlebox is logically linked with a Multihoming Service Provider (MSP). The role of the MSP is to achieve the multihoming for the Homenet by using offloading: the Homenets, by the mean of the middlebox, offloads all its Internet traffic to the MSP, and the offloading is such that the traffic leverages the Homenet's multihoming capability.

The MSP can be seen as a service in the cloud (in a remote network or in devices widely deployed by the MSP in the ISPs). The service is two-fold. On the one hand, the MSP must attract the traffic sent by the Homenet to the Internet, this part is ensured by the middle-box deployed at the Homenet. On the other hand, the MSP must attract traffic sent by the Internet to the Homenet, before this last can receive it. Then, the MSP can send this traffic to the Homenet via the most relevant ISP.

The figure below gives a reference network for the multihoming service for Homenet.

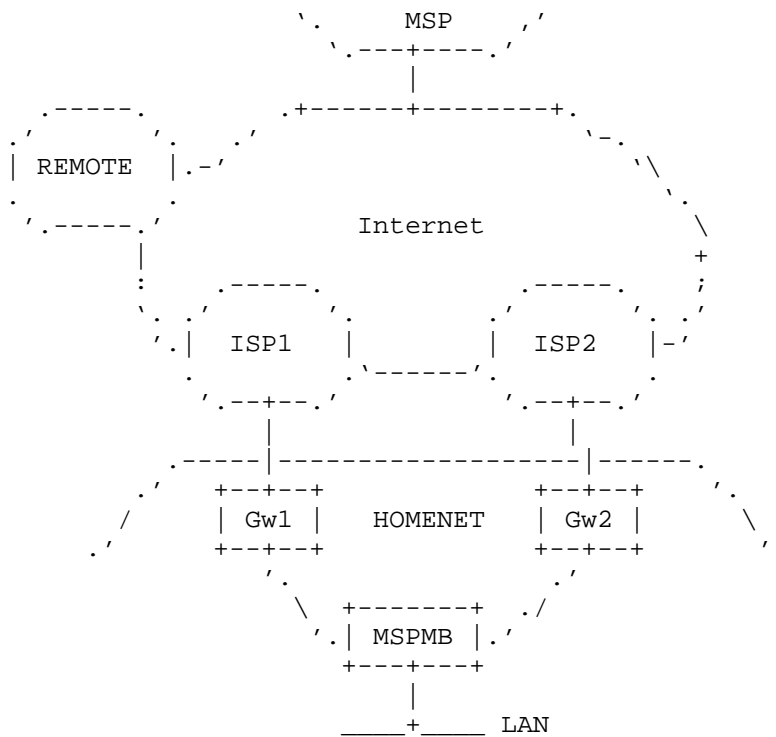


Figure 1: Reference Network

In this figure, HOMENET is the multihomed Homenet, connected to ISP1 via gateway Gw1 and to ISP2 via gateway Gw2. The remote end of communications with the Homenet is designated by REMOTE. MSPMB designates the MSP middlebox in the home network and is logically linked to the MSP multihoming service provider.

Let's imagine that the best to send traffic from the Homenet to the remote end is to go via ISP2 while for the traffic from the remote end to the Homenet it is better to go via ISP1. In this case, the traffic generated from Homenet's LAN is caught by MSPMB that divert traffic to Gw2, then crosses ISP2 and the Internet to reach MSP, then REMOTE. On the other direction, traffic sent by REMOTE goes to MSP that sends the traffic on the Internet to ISP1, then it goes to Gw1, MSPMB, and finally the LAN.

The Multihoming Service Provider (MSP) would typically be operated on an AS well connected to Homenet's ISPs. Or alternatively, a Service provider that has its own devices deployed at the Homenet's ISPs.

As Homenet is targeting IPv6 networks, communications between the Homenet and the MSP cannot rely on NAT but instead they might use encapsulation. For that purpose, LISP [RFC6830] is a perfect candidate. In this case, the MSPMB is an xTR. To ensure zero configuration at the Homenet level, the EID-to-RLOC Cache can be populated on the fly by a mapping system hosted and managed by the MSP. A major advantage of using LISP for communications between the MSP and the Homenet is that residential and SOHO networks would then have access the IPv6 Internet without the need of subscribing to IPv6 ISPs.

The service we propose answers the problem exposed in Section 3 in an elegant way. It also fulfills the four requirements stated above. Requirement (1) (zeroconf) is respected if MSPMB is given directly by the MSP, which can thus be pre-configured to access the MSP service provider. If it is not the case, the process can be simplified if a generalized name and protocol is used to configure the middlebox (e.g., msp.example.org). In addition, if Gw1 and Gw2 provide addresses by the mean of DHCPv6 or RA, addresses at the MSPMB will be configured automatically as well. Obviously, policies and capabilities need configuration either from the home network operator or the MSP directly (which is straightforward with LISP). Finally, UPnP can be used for special services provided to the Homenet by its ISPs.

## 5. Security Considerations

Traffic redirection can be used for DoS or eavesdropping.

## 6. Conclusion

Multihoming in Homenet is considered to be solved by the hosts directly. In this memo, we propose to not involving host in multihoming operations and instead rely on a Multihoming Service Provider deploying a middlebox in the Homenet network in charge of operating multihoming services.

## 7. Normative References

- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, DOI 10.17487/RFC5533, June 2009, <<http://www.rfc-editor.org/info/rfc5533>>.
- [RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, DOI 10.17487/RFC6182, March 2011, <<http://www.rfc-editor.org/info/rfc6182>>.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <<http://www.rfc-editor.org/info/rfc6830>>.

#### Authors' Addresses

Wassim Haddad  
Ericsson  
6210 Spine Road  
Boulder, CO 80301  
USA

Email: [Wassim.Haddad@ericsson.com](mailto:Wassim.Haddad@ericsson.com)

Damien Saucez  
INRIA Sophia Antipolis  
2004, Route des Lucioles BP 93  
06902 Sophia Antipolis CEDEX  
France

Email: [damien.saucez@inria.fr](mailto:damien.saucez@inria.fr)

Joel  
Ericsson  
P.O. Box 6049  
Leesburg, VA 20178  
USA

Email: [Joel.Halpern@ericsson.com](mailto:Joel.Halpern@ericsson.com)

Home Networking  
Internet-Draft  
Intended status: Informational  
Expires: May 10, 2013

E. Kline  
Google Japan  
November 6, 2012

Default Border Definition  
draft-kline-default-perimeter-01

Abstract

Automatic, simple identification of when traffic is crossing a perimeter is highly desirable for a variety of home network uses. This document describes how to use homenet routing protocol adjacencies as the primary signal of a common administrative domain (e.g. "the home"). Classification of interfaces et cetera as internal or external follow from this, as do various policy and implementation implications. One fundamental implication is that the active definition of a home network's interior is no more secure than its policy for forming homenet routing protocol adjacencies.

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 May 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
1.1. Requirements Language . . . . .	3
2. Terminology . . . . .	3
3. Dynamically determining the border . . . . .	4
3.1. Collect information about neighboring routers . . . . .	5
3.2. Classify next hops . . . . .	5
3.3. Classify interfaces . . . . .	6
3.3.1. Mixed mode . . . . .	6
3.4. State changes and logging . . . . .	6
4. Fixed-category interfaces . . . . .	6
4.1. Examples . . . . .	7
5. Security Considerations . . . . .	7
5.1. Disclamatory remarks . . . . .	7
5.2. Security of adjacency formation . . . . .	8
5.2.1. Secure links and authenticated adjacency formation . . . . .	8
5.2.2. Unsecure links . . . . .	8
5.2.3. Recommendations . . . . .	9
5.3. Example border-aware filtering policies . . . . .	9
5.3.1. Anti-spoofing on internal interfaces . . . . .	9
5.3.2. Stateful filtering on external interfaces . . . . .	10
5.3.3. Mixed-mode interface filtering . . . . .	10
6. Acknowledgements . . . . .	11
7. IANA Considerations . . . . .	11
8. References . . . . .	11
8.1. Normative References . . . . .	11
8.2. Informative References . . . . .	11
Author's Address . . . . .	12

## 1. Introduction

Automatic, simple identification of when traffic is crossing the homenet perimeter is highly desirable for a variety of home network uses. This is a non-trivial task as it is tantamount to automated discovery of administrative boundaries.

Many architectures make it difficult or impossible (by design) 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 can be insecure, unreliable, operationally impractical, or may place arbitrary requirements upon the architecture where previously no such requirement existed.

This document describes how to use homenet routing protocol adjacencies as the primary signal of a common administrative domain (e.g. "the home"). Classification of interfaces et cetera as internal or external follow from this, as do various policy and implementation implications. One fundamental implication is that the active definition of a home network's interior is no more secure than its policy for forming homenet routing protocol adjacencies.

### 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 border determination at a manageable scale the scope has been limited to discussing concepts of "interior", "exterior", and "border". Documents may reference any of the terms "internal", "interior", "external", or "exterior" as required by grammar (adjective versus noun use cases). Definitions in use by this document are as follows.

### Internal/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. This document uses the homenet routing protocol adjacencies as a indicator of membership in the same logical administrative domain.

#### External/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.

#### Border/Perimeter

The border 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.

Additionally, some implementations MAY choose to support handling questionable home network configurations which result in an interface qualifying for both interior and exterior classification simultaneously. Requirements for this are discussed further below.

Expressly not considered by this document are architectures having multiple interior networks, nor the relationships between them as separate from their relationship to their common exterior or any common border (e.g. a hierarchy of internal networks). This document is solely concerned with a single interior, a single exterior, and a single logical perimeter between the two.

### 3. Dynamically determining the border

Homenet routers may support interfaces which attempt to learn the nature of their relationship to neighboring routers, and determine where the border between the "interior" and the "exterior" lies.

Interfaces which have not yet determined their categorization may consider themselves to be in a "learning" state, where no traffic is routed but probing continues. Once nodes of any kind (e.g. either routers or hosts) are detected, classification takes place and the router exits the "learning" state.

The classification algorithm documented here is roughly:

1. Continuously collect information on all interfaces about neighboring routers (including manually configured routers).

2. Classify next hops as either "internal" or "external" primarily by homenet router protocol adjacency status.
3. Classify interfaces according to the nature of their next hops.

Manually configured next hops MUST also have their classification as either "internal" or "external" explicitly specified. As such, they can be considered to be of a fixed category, and on-going evaluation is NOT required.

Policies of various sorts can be applied and updated as appropriate based on these classifications, as and when they are determined.

### 3.1. Collect information about neighboring routers

An interface (physical or virtual) which is configured to dynamically assess its internal or external classification MUST periodically probe for routing information on the link. This includes sending Router Solicitations, DHCPv6 Prefix Delegation requests, and probing for homenet routing protocol-capable nodes.

Probing for routing information MUST be performed whenever the interface is logically up. The periodicity of probes is protocol-dependent (e.g. subject to Router Advertisement lifetimes, DHCPv6 lease timers, or homenet routing protocol timers). Wherever possible, implementations SHOULD limit the impact of probing by implementing mechanisms like exponential back-off.

Homenet routers MUST be able to identify loops, e.g. when two (or more) of the router's own interfaces are connected on the same link. Compliant routers MUST administratively log this misconfiguration, and SHOULD implement a mechanism that permits maximum continued homenet functionality if possible. For example, implementations MAY administratively disable all but one of looped interfaces.

### 3.2. Classify next hops

Routing information is used to categorize next hops as either "internal" or "external".

Routers with which a homenet routing protocol adjacency is successfully established MUST be considered "internal".

Routers of which this homenet router has knowledge but with which no homenet routing protocol adjacency is successfully establish AND from which no routing information is learned SHOULD be considered "internal". This includes "downstream" routers for which the homenet router is acting as the Delegating Router via a DHCPv6 Prefix

Delegation exchange.

All other routers with which no homenet routing protocol adjacency is successfully established MUST be considered "external".

### 3.3. Classify interfaces

An interface with no "external" next hops SHOULD be categorized as "internal". This includes interfaces serving leaf networks consisting only of hosts, an interface which has "downstream" routers for which this router is a Delegating Router, an interface with only homenet routing protocol adjacent peers, or any combination thereof.

An interface with next hops all of which are categorized as "external" MUST be categorized as "external".

#### 3.3.1. Mixed mode

Some devices MAY choose to support handling questionable home network configurations which result in an interface having both interior and exterior next hops simultaneously. This can happen if, for example, two homenet routers form an adjacency with each other over the same interface they use for communicating to "upstream" ISP routers.

All homenet routers, whether this configuration is considered supported or not, MUST administratively log and provide product-relevant notification of this configuration, preferably with recommendations for resolution.

### 3.4. State changes and logging

Home routers performing dynamic border discovery MUST continuously evaluate the interior and exterior classifications of interfaces. 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.

Homenet routers performing dynamic border discovery SHOULD administratively log the perimeter classification of all interfaces (physical and virtual), the reason(s) for such classification, and times at which such classifications are made or changed.

## 4. Fixed-category interfaces

Interfaces (physical or virtual) which have product-defined purposes or are otherwise permanently categorized by the homenet router

implementation as exclusively "internal" or exclusively "external" do not require any algorithm to determine their categorization.

Homenet routers MUST restrict relevant traffic on fixed-category interfaces according to their categorizations. Specifically, they MUST NOT originate traffic which could result in re-categorizing the interface IF the interface were dynamically attempting to learn its categorization. For example, a fixed "external" interface MUST NOT attempt to participate in the homenet routing protocol. Similarly, fixed "internal" interfaces must not issue DHCPv6 Prefix Delegation requests.

#### 4.1. Examples

Examples of product-defined interfaces include home router interfaces which are labeled for their intended use, e.g. RJ-45 ports labeled "WAN" and "LAN" or symbols indicating "The Internet" and "inside the home". Other examples include wireless routers with defined separate WLAN "home" and "guest" ESSIDs.

Another set of examples of product-defined, fixed category interfaces are those which require subscriber credentials in order for that interface to successfully pass general purpose traffic. These include authenticated PPPoE sessions and 3G/LTE PDP contexts (e.g. requiring a SIM card associated with a valid customer account). These SHOULD be classified as "exterior".

Similarly, an implementation may consider the interface a mobile device uses to provide service to "tethering" clients to be a fixed-category interface. Such interfaces SHOULD be classified as "interior".

### 5. Security Considerations

A key motivation for border identification is the application of security policies which can take into account classifications of interior, exterior, and the transition from one the other. General remarks, comments on the security of the adjacencies which form the basis of border identification, and examples of potential policies which might be applied follow.

#### 5.1. Disclamatory remarks

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, for example, when connecting to any public "Wi-Fi" network. This is, of course,

equally true of more traditionally "fixed" nodes. Any compromised neighbor node ("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.

Though it goes without saying, at all times homenet implementers MUST remain mindful of best current security practices, including but not limited to RFC 4864 [RFC4864], RFC 4890 [RFC4890], RFC 6092 [RFC6092], and others.

## 5.2. Security of adjacency formation

The security of the border definition is limited by the security applied to the formation of homenet routing protocol adjacencies: the next hops with which a homenet router forms adjacencies are the next hops the router trusts with the application of interior policies.

### 5.2.1. Secure links and authenticated adjacency formation

The trustworthiness of next hops during adjacency formation can be improved if the security of the link connecting them can be trusted. Using encrypted link technologies like 802.1x or secured "Wi-Fi" ESSIDS when forming homenet adjacencies, or authenticating homenet next hops by physical or cryptographic mechanisms limits the ability of malicious nodes to join the homenet interior.

### 5.2.2. Unsecure links

In general IP over Ethernet connections, common to residential Internet (and countless other places like some in-room hotel network) service provider deployments, create the possibility of malicious nodes attempting to join the homenet interior.

In a broad variety of circumstances users already implicitly trust unsecured links. Residential subscribers generally trust that their ISP has properly isolated their connection from any neighbors; few if any subscribers validate the ISP's DHCP server in order to thwart a malicious neighbor intervening.

In the event of a network with a single upstream where an interior next hop is formed instead of an external next hop, the homenet network as a whole would have detected no external next hops. Homenet router networks in which there are no external next hops SHOULD administratively log this configuration and SHOULD provide a means to alert the user to this condition. Note that for isolated networks of homenet routers (e.g. a lab network) this configuration

is entirely valid.

User notification alone is not sufficient protection for the homenet user, and will not correctly alert the user of a homenet with two upstream connections, one of which has mistakenly not categorized a next hop as external. To better serve the homenet user, homenet routers are **SHOULD** follow one or more of the recommendations in Section 5.2.3.

#### 5.2.3. Recommendations

Homenet router implementations that support dynamic discovery of the border (i.e. have interfaces on which the dynamic border detection described in Section 3 can be configured to operate) **SHOULD** support restricting homenet routing protocol adjacency formation to only next hops which meet some user-defined or user-verified authentication mechanism (including examples described in Section 5.2.1).

Alternatively, implementations **MAY** incorporate a mechanism (possibly physical) whereby a homenet user can disable border detection on an interface which the user wishes to force into either an interior or exterior classification (e.g. a button to force an interface to be "external" only).

#### 5.3. Example border-aware filtering policies

As a homenet router forms adjacencies and learns internal aggregate prefixes it could dynamically maintain a single logical entity encompassing all current internal prefixes in use that can be treated as a whole (e.g. an access list). Below are example filtering policies that might be applied by homenet routers with knowledge of both this prefix set and the interior or exterior classification of all interfaces.

The examples below use the string "{interior\_nets}" for refer to the grouping of all internal aggregate prefixes. The sample filtering policy rules are written in configuration pseudo-syntax that should hopefully be intuitive.

##### 5.3.1. Anti-spoofing on internal interfaces

Given knowledge of all interior network prefixes and the categorization of interfaces, all interior interfaces could apply a stateless filter designed to prevent devices in the home from originating source-address-spoofed traffic.

Using a filter configuration pseudo-syntax:

```
from !{interior_nets} to !{interior_nets} deny
... # permit or deny other kinds of traffic
```

### 5.3.2. Stateful filtering on external interfaces

Given knowledge of all interior network prefixes and the categorization of interfaces, all exterior interfaces could apply a stateful filter designed to discard traffic without matching state in the homenet router.

Using a filter configuration pseudo-syntax:

```
... # permit other kinds of good traffic first
from {interior_nets} to !{interior_nets} permit
from !{interior_nets} to {interior_nets} established permit
from any to any deny
```

### 5.3.3. Mixed-mode interface filtering

Given knowledge of all interior network prefixes and the categorization of interfaces, all mixed-mode interfaces could apply a stateful filter designed to discard exterior traffic without matching state in the homenet router and still statelessly permit internal traffic.

Using a filter configuration pseudo-syntax:

```
... # permit other kinds of good traffic first
from {interior_nets} to !{interior_nets} permit
from !{interior_nets} to {interior_nets} established permit
from {interior_nets} to {interior_nets} permit
from any to any deny
```

Because routing changes elsewhere in the home may cause traffic to shift among interior next hops which may not have state, traffic between interior routers may not be well-served by stateful filtering. One consequence for this policy on mixed-mode interfaces is that traffic from the exterior with spoofed source addresses from the "{interior\_nets}" set of prefixes may be mistakenly allowed into the home.

Filter implementations which can incorporate knowledge of the previous and next hops and their classifications can design much more precise filters. Such implementations could deny traffic with "{interior\_nets}" source addresses arriving from an exterior next hop, but permit them from an interior next hop on the same mixed-mode

interface.

## 6. Acknowledgements

Many thanks for the constructive input and criticism of Shwetha Bhandari, Lorenzo Colitti, Markus Stenberg, Mark Townsley, and Ole Troan.

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

## 7. IANA Considerations

This memo includes no request to IANA.

## 8. References

### 8.1. Normative References

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

### 8.2. 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.

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

[RFC4890] Davies, E. and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", RFC 4890, May 2007.

[RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.

[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>>.

Author's Address

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

Phone: +81 03 6384 9000

Email: [ek@google.com](mailto:ek@google.com)



HOMENET  
Internet-Draft  
Intended status: Standards Track  
Expires: May 10, 2013

D. Migault (Ed)  
Francetelecom - Orange  
W. Cloetens  
SoftAtHome  
P. Lemordant  
Francetelecom - Orange  
C. Griffiths  
Comcast Cable Communications  
November 6, 2012

IPv6 Home Network Front End Naming Delegation  
draft-mglt-homenet-front-end-naming-delegation-01.txt

Abstract

CPEs are designed to provide IP connectivity to the Home Network. Most of the CPEs are also providing the IP addresses of the nodes of the Home Network. This makes CPEs good candidates for hosting the Naming Service that would make devices reachable from the Home Network but also from the Internet.

CPEs have not been designed to host a Naming Service reachable from the Internet. This would expose the CPEs and the Home Network to resource exhaustion which would result in making the Home Network unreachable, and most probably would also affect the Home Network inner communications.

This document describes an Front End Naming Architecture where the CPEs manage the DNS(SEC) zone for its Home Network, and outsource the zone to Public Server for resolution coming from the Internet.

The goal of the document is first to describe a Naming Architecture that fulfills Home Network Naming requirements without exposing the CPE to resource exhaustion. Then we intend the CPEs to be easily configured by the End Users, and describe the necessary information the End User is expect to provide to the CPE.

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 May 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. Requirements notation . . . . .	4
2. Introduction . . . . .	4
3. Front End Naming Architecture Requirements . . . . .	5
4. Front End Naming Architecture Presentation . . . . .	6
5. Front End Naming Architecture Description . . . . .	8
5.1. Setting the Homenet Authoritative Server . . . . .	8
5.2. Setting the Homenet View . . . . .	8
5.3. Setting the Public View . . . . .	9
5.4. Synchronizing the Public View . . . . .	9
5.5. Securing the Synchronization . . . . .	10
5.6. Setting the Homenet Resolution Server . . . . .	11
5.7. Additional Views . . . . .	11
6. CPE's interface Recommendations . . . . .	11
7. Position toward Homenet Architecture . . . . .	12
8. Security Considerations . . . . .	13
8.1. Names are less secure than IP addresses . . . . .	13
8.2. Names are less volatile than IP addresses . . . . .	13
8.3. DNSSEC is recommended to authenticate DNS hosted data . . . . .	14
9. IANA Considerations . . . . .	14
10. Acknowledgment . . . . .	14
11. References . . . . .	15
11.1. Normative References . . . . .	15
11.2. Informational References . . . . .	15
Appendix A. Document Change Log . . . . .	16
Authors' Addresses . . . . .	16

## 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

IPv6 provides global IP reachability to almost all nodes of the Home Network even outside the Home Network. However End Users do not want to access to the Services hosted in the Home Network with IPv6 addresses, but prefer to use names.

CPEs are already providing IPv6 connectivity to the Home Network and generally provide IPv6 addresses or prefixes to the nodes of the Home Network. This makes the CPEs a good candidate to manage binding between names and IP addresses of the nodes. In other words, the CPE is the natural candidate for setting the DNS(SEC) zone file.

CPEs are usually low powered devices designed for the Home Network, but not for heavy traffic. CPEs can host the Naming Service for the Home Network but should not be exposed on the Internet. This would expose the CPE to resource exhaustion. As a consequence, it may isolate the Home Network from the Internet and affects the services hosted by the CPEs, thus affecting Home Network communications. As a result, CPE SHOULD NOT host the Naming Service of the Home Network for resolutions coming from the Internet.

In this document, we propose that the CPE sets the DNS(SEC) zone of the Home Network. The CPE may generate different zones one for the queries coming from the Home Network, and one for queries coming from the Internet. We respectively call these Zones Homenet View and Public View. The CPE hosts the Homenet View and responds to the associated DNS(SEC) queries coming from the Home Network. For queries coming from the Internet, the CPE outsources the Public View to Public DNS(SEC) Servers that responds to the queries.

This document describes the Front End Naming Architecture where the CPE hosts the Naming Service for the Home Network and outsources it for queries from outside the Home Network. We especially insists on the parameters the CPE requires to properly set up the Front End Naming Architecture.

Section 3 describes the Front End Naming Architecture requirements. Section 4 presents the different functional entities involved in the Front End Naming Architecture. Section 5 details configuration of the various functional entities as well as how they interact each

other. Section 6 is informative and sums up the different inputs the CPE requires from the End User to set up the Front End Naming Architecture. Section 7 positions the described architecture toward the Home Network Architecture. Finally Section 8 provides security considerations.

### 3. Front End Naming Architecture Requirements

This section lists and details goals and requirements of Front End Naming Architecture.

- REQUIREMENT 1: DNS(SEC) queries for subdomain of the Homenet Domain Name MUST be responded by the Public DNS(SEC) Servers when issued from outside the Home Network. CPE could hardly cope with heavy traffic coming from the Internet. To avoid exposing the CPE to resource exhaustion, the Naming Service is outsourced on the Public DNS(SEC) Servers for traffic coming from the Internet.
- REQUIREMENT 2: The CPE MUST NOT, by default, accept any DNS(SEC) queries from outside the Home Network. In some aspects, it rewords the previous requirement.
- REQUIREMENT 3: The IP address of the CPE SHOULD NOT be publicly published. This requirement avoids the DNS(SEC) queries incidentally ends up on the CPE.
- REQUIREMENT 5: DNS(SEC) queries for subdomain of the Homenet Domain Name MUST be responded by the CPE when issued from the Home Network. To guarantee the Home Network independence in case the Home Network has no connectivity on the Internet, the CPE MUST respond to DNS(SEC) queries for subdomain of the Homenet Domain Name coming from the Home Network.
- REQUIREMENT 6: The CPE MUST be able to update the Home Network Zone hosted on the Public DNS(SEC) Servers.
- REQUIREMENT 7: The CPE SHOULD be able to provide different views. At least the CPE should be able to handle a view for the Home Network nodes and a view for the nodes outside the Home Network. Home Network nodes that are not supposed to be reachable from outside the Home Network are not expected to be part of the latest view.

#### 4. Front End Naming Architecture Presentation

This section describes the Front End Naming Architecture and defines the notations used in this document.

- Home Network: designates all devices that are behind and managed by the CPE.
- Internet: designates the network the CPE is attached to.

The CPE connects the Home Network to the Internet. Although the different functional entities listed below MUST NOT necessarily be hosted on the CPE, we assume in this document they are hosted on the CPE:

- Homenet Resolving Server: is the DNS Resolver Server of the Home Network. Typically its IP address is announced via DHCPv6. Most of DNS(SEC) queries from nodes on the Home Network are expected to be addressed to this Homenet Resolving Server. The Homenet Resolving Server is expected to receive queries only from the Home Network.
- Homenet Authoritative Server: is the Authoritative Server of the Home Network. This server hosts bindings between FQDNs and IP addresses. Unless cached, most of the DNS(SEC) queries sent from the Home Network that concerns a node in the Home Network are expected to be forwarded to this Homenet Authoritative Server. More specifically DNS(SEC) resolutions sent from the Home Network are expected to be sent to the Homenet Resolving Server. The Homenet Resolver Server is a forwarder and forwards to the Homenet Authoritative Server queries for domain names or subdomain of the Homenet Domain Name. For other resolutions, the Homenet Resolving Server proceeds to traditional DNS(SEC) resolutions over the public DNS(SEC) infrastructure.
- Homenet View: is the DNS(SEC) zone that contains all bindings between FQDNs associated to the Homenet Domain Name and IP addresses. The Home Network may have multiple views, but for most Home Networks, a single Homenet View is expected. Information of this Homenet View is only visible from the Home Network.
- Public View: is the view that contains the bindings between FQDNs and IP addresses. Unlike the Homenet View, the Public View is expected to be publicly published. The Public View contains information visible from the Internet. It is expected that the Public View is constituted by a subset of the names of the

Homenet View. More specifically, devices that are not expected to be reachable from the Internet should not be part of the Public View. In some implementations, the Public View may be equivalent to the the Homenet View. In this latter case Public Views and Homenet Views will be represented by a single file.

- Master Public Server: is the part of the CPE that deals with the Public view of the Home Network. The Master Public Server is in charge of providing the Public View to the Public DNS(SEC) Servers. It is not necessarily a DNS(SEC) server. However, in this document we are using DNS mechanisms to synchronize the Public View in the Public DNS(SEC) Servers and the Public View on the CPE.
- WAN Interface: the CPE Interface on the Internet.
- Homenet Interfaces: the CPE Interfaces on the Home Network. There might be a single or multiple interfaces.

The other involved entities are:

- Public DNS(SEC) Servers: are the servers on the Internet hosting the Public View of the Home Network.
- Homenet Node: a Node of the Home Network
- Node: a Node located on the Internet. This Node is expected to be in most of the cases a resolving server.
- Homenet Domain Name: The domain name associated to the Home Network. There may be one or multiple domain names.

Figure 1 illustrates how a DNS(SEC) resolution is performed from a Node in the Home Network or from a node on the Internet.

The Homenet Node sends a DNS(SEC) query to the Homenet Resolving Server (1). When the Homenet Resolving Server receives the DNS(SEC) it notices that query name is a subdomain of the Homenet Domain Name (example.com), and forwards the query to the Homenet Authoritative Server that hosts the Homenet View (2). The Homenet Authoritative Server sends the response to the Homenet Resolving Server (3), which finally sends the response to the Homenet Node (4).

For a node located on the Internet, the DNS(SEC) query is requesting the Public DNS infrastructure (.com) which redirects the DNS(SEC) query to the Public DNS(SEC) Servers (a). The Public DNS(SEC) Server sends the response back to the Node (b).

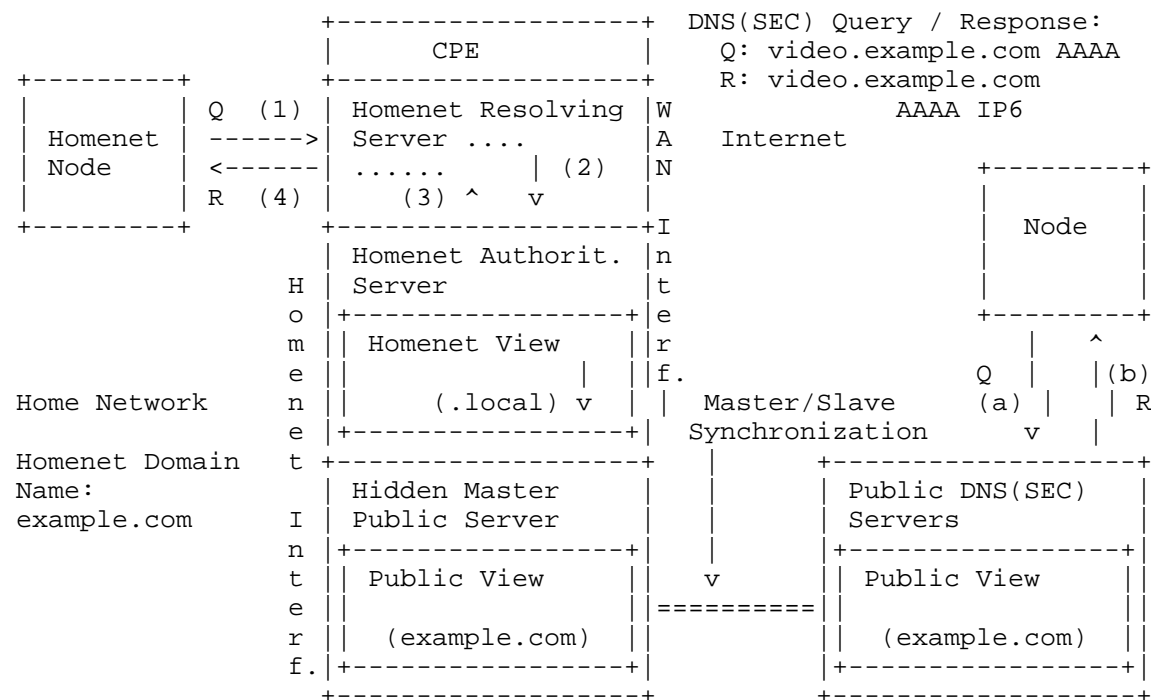


Figure 1: Front End Naming Architecture Description

5. Front End Naming Architecture Description

This section provides a more detailed description of the Front End Naming Architecture. More specifically it shows how the entities described in Section 4 are organized to fulfill the requirements of Section 3.

5.1. Setting the Homenet Authoritative Server

The Homenet Authoritative MUST be configured to reject any queries coming from outside the Home Network, i.e. not from the Homenet Interface. In other words, DNS queries related to the Homenet Domain Names MUST never be received from the WAN Interface.

5.2. Setting the Homenet View

The Homenet Authoritative Server may be authoritative for multiple Homenet Domain Names and each Homenet Domain Name may be associated with multiple views.

The CPE is expected to be provided the various Homenet Domain Names so it can properly generate the associated Homenet Zone files and the appropriate DNS(SEC) settings.

### 5.3. Setting the Public View

The Public DNS(SEC) Servers MUST handle all DNS(SEC) queries related to any Homenet Domain Names that are sent from outside the Home Network.

The CPE MUST generate the Public View. In the case of multiple Homenet Domain Names, multiple views MUST be generated, and in order to fill properly the SOA and NS field, the CPE must be provided for each Homenet Domain Name the corresponding Public DNS(SEC) name, and IP addresses.

### 5.4. Synchronizing the Public View

Uploading and dynamically updating the zone file on the Public Servers can be seen as zone provisioning between the CPE (Hidden Master) and the Public Server (Slave Server). This can be handled either in band or out of band. DNS dynamic update [RFC2136] may be used. However, in this section we detail how to take advantage of the DNS slave / master architecture to deploy updates to public zones.

The Public DNS Server is configured as a slave for the Homenet Domain Name. This slave configuration has been previously agreed between the End User and the provider of the Public DNS Servers. The CPE is hosting the Public Zone files associated to the various Homenet Domain Names and associated views. Each of these files are associated a Public Server. In order to set the master/ slave architecture, the CPE acts as a Hidden Master Public Server, which is a regular Authoritative DNS(SEC) Server listening on the WAN interface.

The Hidden Master Public Server is only expected to initiate AXFR [RFC1034], IXFR [RFC1995] transfers to configured slave DNS servers. The Hidden Master Public Server should send NOTIFY messages [RFC1996] in order to update Public DNS server zones as updates occur.

The CPE MUST be configured to send NOTIFY only when necessary. It is recommended for example that it checks first the SOA on the Public DNS Server before sending a NOTIFY. In other words, rebooting a CPE SHOULD NOT systematically trigger a NOTIFY message.

Hidden Master Public Server differs from the Homenet Authoritative Server by:

- Interface: the Homenet Authoritative Server listens on the Homenet Interface whereas the Hidden Master Public Server listen on the WAN Interface
- View: Homenet Authoritative Server hosts the names that are available on the Home Network, whereas the Hidden Master Public Server hosts the names that are publicly available. These two zones may differ since some of the nodes may not be reached from outside the Home Network.
- Traffic: Homenet Authoritative Server expects traffic from the Home Network, whereas the Hidden Master Public Server only accepts traffic from the Public Servers.
- Function: Homenet Authoritative Servers acts as an authoritative DNS Server on the Home Network, whereas the Hidden Master Public Server only synchronizes with the Public DNS Servers.

In this document, Master Public Server differs from the Homenet Authoritative Server as different functions. Both functions may be implemented by a single running instance of Authoritative Servers.

#### 5.5. Securing the Synchronization

Exchange between the Public Servers and the CPE MUST be secured, at least for integrity protection and for authentication. This is the case whatever mechanism is used between the CPE and the Public DNS(SEC) Servers.

TSIG [RFC2845] can be used to secure the DNS communications between the CPE and the Public DNS(SEC) Servers. TKEY [RFC2931] can be used for re-keying the key used for TSIG. Using TSIG and TKEY requires that this mechanism is implemented on the DNS(SEC) Server's implementation running on the CPE. One disadvantage is that TKEY does not provides authentication mechanism, and the initial shared secret must be set manually.

Protocols like TLS [RFC5246] / DTLS [RFC6347] can be used to secure the transactions between the Public Servers and the CPE. Their use would require the implementations to integrate TLS/DTLS as a security layer. TLS/DTLS can use certificates to authenticate the Public Server and the CPE. For example, the certificates can be hosted on a dongle.

IPsec [RFC4301] IKEv2 [RFC5996] can also be used to secure the transactions between the CPE and the Public Servers. IKEv2 provides multiple authentications possibilities with its EAP framework. Then, IPsec security does not require any changes of the DNS applications.

For these reasons, we recommend using IPsec.

#### 5.6. Setting the Homenet Resolution Server

The Homenet Resolving Server MUST be configured as a DNS forwarder. When a DNS(SEC) query coming from the Home Network concerns a Homenet Domain Name or a Homenet Subdomain Name, the resolution MUST be performed with the Homenet Authoritative Server. If the Home Network has multiple Homenet Domain Names, multiple forwarding rules may be applied.

To properly configure a basic configuration, the Homenet Resolving Server needs to be informed of the Homenet Domain Names and associated Homenet Authoritative Server. They may be one or multiple associated Homenet Authoritative Servers. The same Authoritative Naming Server may be used for multiple Homenet Domain Names.

#### 5.7. Additional Views

In this document, we considered the Public and Homenet View. Each of these Views may have additional views.

### 6. CPE's interface Recommendations

This section describes the various objects that are required to properly set the Front End Naming Architecture. This section is informational, and is intended to clarify the information handled by the CPE and the various settings to be done.

A Public Server is defined with the following information:

- Public Server Name: The associated FQDN of the Public Server
- IP addresses: The list of IP addresses associated to the Public Server. This list should not be provided by the End User. Instead, it should be provided by performing a DNSSEC exchange. If the Public Server Name DNS resolution cannot be performed with DNSSEC, then it is recommended to provide this field. This list of IP address is used to generate the Public View with the proper values for SOA and NS and associated AAAA fields.
- Public Server Management Name: The FQDN of the management interface. This Management interface designated who the CPE is synchronizing its Public View with.

- Public Server Management IP addresses: The list of IP addresses associated to the Public Server Management Name. This field is not expected to be filled by the End User, but to be derived from the Public Server Management Name with a DNS(SEC) query.
- Authentication Method: How the CPE authenticates itself to the Public Server.
- Authentication data: Associated Data.

To set a View one needs to have the following information:

- Homenet Domain Name: The Domain Name of the zone.
- Public Server Name (optional): The Server that are expected to host the View. It is required both to field the SOA, NS and associated AAAA as well as to define where the View has to be uploaded. If the View is the Homenet View and the Homenet Authoritative Server is hosted on the CPE, then, this information is not required.
- Rules: Defines specific rules for deriving the View. Example of rule s may be a list of FQDNs or IP addresses that MUST be included or removed...
- DNSSEC Data: DNSSEC data required to generate the DNSSEC zone. This can be the various DNSSEC Keys for example.

First the CPE MUST reject DNS queries received from the WAN Interface. Then the CPE MUST list the Homenet Views and Public Views. Homenet Views are those without the Public Server Name specified and are loaded on the Homenet Authoritative Server. The Homenet Resolving Server is configured as a forwarder for these Views. Public Views are loaded on the Master Public Server, the communications between the Master Public Server and the Public Servers are secured, with an IPsec authenticated and encrypted traffic flow for example.

## 7. Position toward Homenet Architecture

This section positions the Front End Naming Architecture toward the Naming recommendation of [I-D.chown-homenet-arch].

The Front End Naming Architecture has been designed to favor unmanaged operations. Naming configuration is automatically performed by the CPE.

The Front End Naming Architecture provides the End User a mean to assign names to their devices and associate these names to an Internet domain. With traditional naming configuration that sets an "search" field for the resolvers, the Front End Naming Architecture provides relative naming resolution. The search field is configurable on the DHCPv6 Server hosted on the CPE.

Homenet devices can be attached in multiple local and Internet name spaces. The Front End Naming Architecture works internally and externally depending where the End User is. With Views, not all devices are visible from the Internet.

The Front End Naming Architecture completely coexists with the Internet name services.

With the Homenet View hosted on the CPE, Name resolution and service discovery for reachable devices must continue to function if the local network is disconnected from the global Internet.

## 8. Security Considerations

The Front End Naming Architecture described in this document solves exposing the CPE's DNS service as a DoS attack vector.

### 8.1. Names are less secure than IP addresses

This document describes how an End User can make his services and devices from his Home Network reachable on the Internet with Names rather than IP addresses. This exposes the Home Network to attackers 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 used either for the Home Network domain or for the devices present less randomness (livebox, router, printer, nicolas, jennifer, ...) and thus exposes the devices to dictionary attacks.

### 8.2. Names are less volatile than IP addresses

IP addresses may be used to locate a device, a host or a Service. However, Home Networks 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 and may return a NXDOMAIN response.

### 8.3. 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 perform 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. IANA Considerations

This document has no actions for IANA.

## 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

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, August 1996.
- [RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", RFC 1996, August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2136] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [RFC2845] Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, May 2000.
- [RFC2931] Eastlake, D., "DNS Request and Transaction Signatures (SIG(0)s)", RFC 2931, September 2000.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.

### 11.2. Informational References

- [I-D.chown-homenet-arch]  
Arkko, J., Chown, T., Weil, J., and O. Troan, "Home Networking Architecture for IPv6",  
draft-chown-homenet-arch-01 (work in progress),  
October 2011.

#### Appendix A. Document Change Log

[RFC Editor: This section is to be removed before publication]

-01:

- \* Added C. Griffiths as co-author.

- \* Updated section 5.4 and other sections of draft to update section on Hidden Master / Slave functions with CPE as Hidden Master/Homenet Server.

- \* For next version, address functions of MDNS within Homenet Lan and publishing details northbound via Hidden Master.

-00: First version published.

#### Authors' Addresses

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

Wouter Cloetens  
SoftAtHome  
vaartdijk 3 701  
3018 Wijkmaal  
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

Chris Griffiths  
Comcast Cable Communications  
One Comcast Center  
Philadelphia, PA 19103  
US

Phone:  
Fax:  
Email: chris\_griffiths@cable.comcast.com  
URI: <http://www.comcast.com>



Internet Engineering Task Force  
Internet-Draft  
Intended status: Standards Track  
Expires: April 18, 2013

O. Troan  
Cisco  
October 15, 2012

Naming and Service Discovery in the Home Network  
draft-troan-homenet-naming-and-sd-00

Abstract

This memo describes how simple naming and service discovery is done in 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 April 18, 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 . . . . .	3
3. Terminology . . . . .	3
4. Problems and Requirements . . . . .	3
5. Security Considerations . . . . .	3
6. IANA Considerations . . . . .	3
7. Acknowledgements . . . . .	3
8. Normative References . . . . .	3
Author's Address . . . . .	3

1. Introduction

2. 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].

3. Terminology

4. Problems and Requirements

5. Security Considerations

6. IANA Considerations

7. Acknowledgements

8. Normative References

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

Author's Address

Ole Troan  
Cisco  
Oslo,  
Norway

Email: ot@cisco.com

