

Network Working Group
Internet-Draft
Expires: May 9, 2014

M. Andrews
ISC
November 5, 2013

Automated Delegation of IP6.ARPA reverse zones with Prefix Delegation
draft-andrews-dnsop-pd-reverse-02

Abstract

This document describes a method to automate the delegation of IP6.ARPA reverse zones when performing Prefix Delegations.

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 9, 2014.

Copyright Notice

Copyright (c) 2013 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. Method	3
3. Example	4
4. IANA Considerations	4
5. Security Considerations	4
6. Normative References	5
Author's Address	5

1. Introduction

This document describes a method to automate the delegation of IP6.ARPA reverse zones when performing Prefix Delegations.

This will allow home users and small businesses to have IP6.ARPA zones without manual intervention on the part of the ISP.

2. Method

1) CPE device generates a RSA key pair and stores this in non-volatile memory.

2) CPE device generates a DHCPv6 Prefix Delegation [RFC3633] request which includes a KEY-RDATA option (code point TBA), which contains a the rdata of a DNS KEY record containing a RSASHA256 key using the public components of the previously generated RSA key pair.

3) DHCP server updates DNS server based on the prefix it is delegating and the KEY-RDATA, using TSIG [RFC2845] for authentication, and responds with prefix. If this is a new prefix delegation, it will clear out all the old DNS records as part of the delegation process. If there are multiple prefixes being delegated the ISP's DNS server will be updated for all of them. If the delegated prefix is not nibble aligned then the server will update all the reverse apex names that cover the address space, i.e. 1, 2, 4 or 8 KEY records will be added all with the same rdata contents.

4) CPE device configures the nameserver built into it to serve the reverse of the delegated prefixes. Alternatively it may configure other nameservers to serve these zones, however the method to do that is out of scope for this document.

5) CPE device generates a DNS UPDATE [RFC2136] which delegates the reverse name space to itself and others if they have been configured. It uses SIG(0) [RFC2931] to sign the request, with owner name matching the reverse of the delegated prefix.

6) The ISP's DNS server is configured to accept self-signed requests (the owner name used in the SIG(0) signature matches the owner name of the data to be updated). It examines the request, looks at the KEY record added by the DHCPv6 server, and decides whether the request is valid.

3. Example

If 2001:DB8:1:4::/62 is delegated then KEY records for

```
4.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA KEY ...
5.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA KEY ...
6.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA KEY ...
7.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA KEY ...
```

will be added. The CPE device will configure the nameservers to serve all of the following zones

```
4.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
5.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
6.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
7.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
```

then will send individual UPDATE messages to delegate each of the reverse zones.

```
% nsupdate -k K4.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
update add 4.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA NS ...
send
% nsupdate -k K5.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
update add 5.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA NS ...
send
% nsupdate -k K6.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
update add 6.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA NS ...
send
% nsupdate -k K7.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA
update add 7.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA NS ...
send
```

4. IANA Considerations

Allocate a DHCPv6 code point for KEY-RDATA.

5. Security Considerations

The UPDATE requests are all signed. This is a proven method for securing UPDATE requests in the DNS.

As a RSA key is being used there is no issue with key material being sent in the clear.

Only the CPE device and the ISP itself is capable of creating,

updating or destroying the delegation.

6. Normative References

- [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., "Secret Key Transaction Authentication for
DNS (TSIG)", RFC 2931, September 2000.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic
Host Configuration Protocol (DHCP) version 6", RFC 3633,
December 2003.

Author's Address

M. Andrews
Internet Systems Consortium
950 Charter Street
Redwood City, CA 94063
US

Email: marka@isc.org

Network Working Group
Internet-Draft
Intended status: Informational
Expires: February 14, 2014

F.J. Baker
Cisco Systems
August 13, 2013

Requirements and Use Cases for Source/Destination Routing
draft-baker-rtgwg-src-dst-routing-use-cases-00

Abstract

This note attempts to capture important use cases for source/destination routing.

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 February 14, 2014.

Copyright Notice

Copyright (c) 2013 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	2
1.1. Requirements Language	3
2. Use Cases	3
2.1. Simple Egress Routing	3
2.2. General Egress Routing	4
2.3. Specialized Egress Routing	5
2.4. Intra-domain access control	7
3. Derived Requirements	8
4. IANA Considerations	8
5. Security Considerations	8
6. Privacy Considerations	8
7. Acknowledgements	8
8. References	8
8.1. Normative References	8
8.2. Informative References	8
Appendix A. Change Log	9
Author's Address	9

1. Introduction

Source/Destination routing has been proposed in the IPv6 community and specifically in homenet as a means of dealing with multihomed networks whose upstream networks give them provider-allocated addresses. An initial approach was suggested in [RFC3704], which assumed that a packet following a default route to an egress CPE Router might arrive at the wrong one, and need to be redirected to the right CPE Router. Subsequent approaches, including those listed in the bibliography, have focused on using routing protocols or routing procedures with extensions that make decisions based on both the source and the destination address.

"Source/Destination Routing" is defined as routing in which both the source and the destination address must be considered in selecting the next hop. It might be thought of as routing "to a destination with a constraint" - a router might have multiple routes to a given destination, and follow the one that also obeys the constraint, or it might have only one route to a destination but correctly fail to forward a packet that doesn't meet the constraint. From that perspective, the logic here extends to other cases in which a constraint might be placed on the route. As with all routing, a primary requirement is to follow the longest-match-first rule to the destination; following a less specific route may well take traffic to the wrong place.

As a side note, source address spoofing in this case will be limited to addresses from the indicated source prefixes, obviating the need for upstream ingress filtering. Ingress filtering within the domain in LAN switches can prevent spoofing of addresses within those prefixes.

This note attempts to capture common use cases. These will be in terms of a general statement of intent coupled with a specific example of the intent for clarity. The use cases are obviously not limited to these, but these should be a reasonably complete set.

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 [RFC2119].

2. Use Cases

The use cases proposed here are not an exhaustive set, but are representative of a set of possibilities. At least three are presently-deployed use cases; the fourth is a possible use case within an edge network.

2.1. Simple Egress Routing

One use case is as shown in Figure 1. A customer network has two or more upstream networks, and a single CPE Router. Each upstream network allocates a prefix for use in the customer network, and the customer network configures a subnet from each of those ISP prefixes on each of its LANs. The CPE Router advertises default routes into the network that are "from" each PA prefix. Apart from prefix itself, the services of the upstream ISPs are indistinguishable; they each get the customer to the Internet.

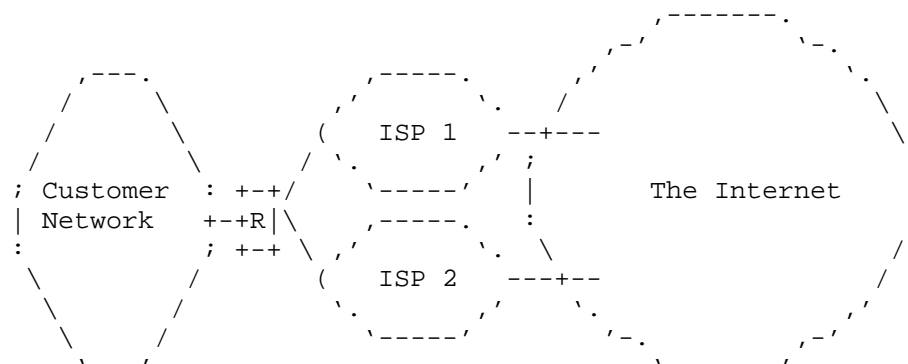


Figure 1: Egress Routing in a Multihomed Environment with One CPE Router

The big issue in this network is, of course, ingress filtering [RFC2827] by the upstream ISP. If packets intended for a remote destination pass through the wrong ISP, they will be blocked. In the ideal case, traffic following default route gets to the upstream network indicated by its source address.

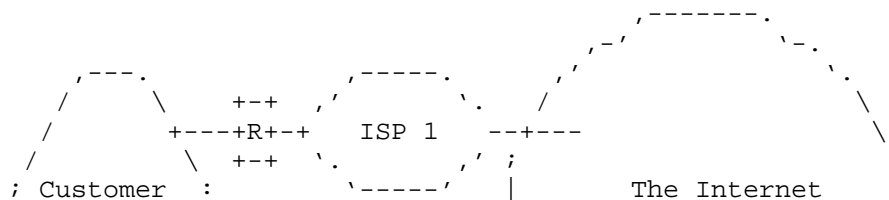
The CPE Router could, at least in concept, advertise a single default route into the network, as all traffic to an upstream ISP must pass through that CPE Router. However, should another CPE Router be added later, it would have to change its behavior to accomodate that CPE Router (as in Section 2.2). Hence, the single CPE Router must advertise two default routes into the network, one "from" each PA prefix.

In this case, the destination prefix in routing is a default route, `::/0`. The source prefix is the prefix allocated by the ISP. In this case, routing within the network is largely unchanged, as all traffic to another network goes to the CPE Router, but the CPE Router must send it to the correct ISP.

Note that in this use case, if there are other routers or internal routes in the network, there is no need for them to specify source prefixes on their routes, and if they do, the prefix specified is likely to be `::/0`. The reason is that traffic arriving from the ISPs must be delivered to destinations within the network, so routing cannot preclude them.

2.2. General Egress Routing

A more general use case is as shown in Figure 2. A customer network has two or more upstream networks, with a separate CPE Router for each one. Each upstream network allocates a prefix for use in the customer network, and the customer network configures a subnet from each of those ISP prefixes on each of its LANs. Each CPE Router advertises a default route into the customer network. Apart from prefix itself, the services of the upstream ISPs are indistinguishable; they each get the customer to the Internet.



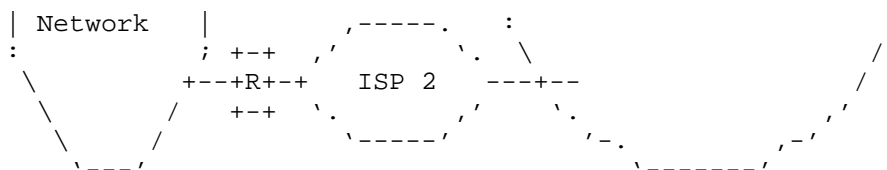


Figure 2: Egress Routing in a Multihomed Environment

The big issue in this network is again ingress filtering [RFC2827] by the upstream ISP. If packets intended for a remote destination pass through the wrong ISP, they will be blocked. Traffic following default route gets to the upstream network indicated by its source address.

In this case, the destination prefix in routing is a default route, `::/0`. The source prefix is the prefix allocated by the ISP. We want a routing algorithm that sends packets matching such a specification to the CPE Router advertising that default route.

Note that in this use case, if there are other routers or internal routes in the network, there is no need for them to specify source prefixes on their routes, and if they do, the prefix specified is likely to be `::/0`. The reason is that traffic arriving from the ISPs must be delivered to destinations within the network, so routing cannot preclude them.

2.3. Specialized Egress Routing

A more specialized use case is as shown in Figure 3. A customer network has two or more upstream networks, with one or more CPE Routers; the example shows a separate CPE Router for each one. Each upstream network allocates a prefix for use in the customer network, and the customer network configures a subnet from each of those ISP prefixes on each of its LANs. Some CPE Routers might advertise a default route into the customer network; one or more of the other CPE Routers, perhaps all of them, advertise a more-specific route. The services offered by the upstream networks differ in some important way.



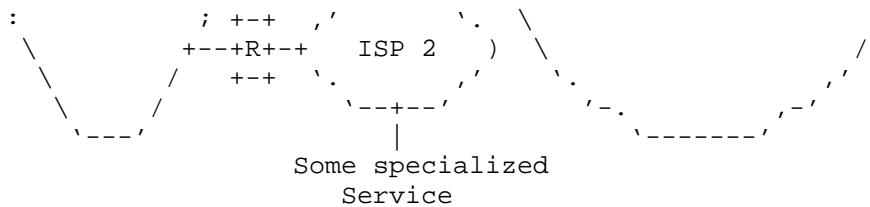


Figure 3: Egress Routing with a specialized upstream network

A specific example of such a service is the NTT B-FLETS video service in Japan; however, the use case describes any use with one or more walled gardens. In the B-FLETS case, a customer may purchase services from a number of ISPs, providing general Internet access. However, the video service requires customers accessing it to use its allocated prefix, and other ISPs (following [RFC2827]) will not accept that prefix as a source address. This is similar to the previous use cases, but

- o the only application at that "ISP" is the video service,
- o packets using the video service MUST use the video service's source and destination addresses, and
- o no other service will accept a video service address as a source address.

The big issue in this network is, once again, ingress filtering [RFC2827] by the upstream ISP, with the additional caveat that the upstream services are far from identical. If packets intended for a remote destination pass through the wrong ISP, they will be blocked. Additionally, while other ISPs advertise access to the general Internet, they may not provide service to the specialized service in question. Hence, egress routing in this case also ensures delivery to the intended destination using the bandwidth it provides. In the ideal case, traffic following default route gets to the upstream network indicated by its source address.

In this case, one or more ISPs might offer a default route as a destination prefix in routing, `::/0`. The source prefix is the prefix allocated by the ISP. In addition, the ISP offering the specialized service advertises one or more specific prefixes for those services, with appropriate source prefixes for their use. We want a routing algorithm that sends packets matching such a specification to the CPE Router advertising that indicated route, and dropping, perhaps with an ICMPv6 response, packets for which it effectively has no route.

Note that in this use case, if there are other routers or internal routes in the network, there is no need for them to specify source prefixes on their routes, and if they do, the prefix specified is likely to be `::/0`. The reason is that traffic arriving from the ISPs must be delivered to destinations within the network, so routing cannot preclude them.

2.4. Intra-domain access control

A use case within the confines of a single network is as shown in Figure 4. A network has one or more internal networks with differing access permission sets; the financial servers might only be accessible from a set of other prefixes that financial people are located in, or university grade records is only reachable from the offices of professors. This could be implemented using firewalls between the domains, or using application layer filters; in this case, the routing architecture replaces an exclusive firewall rule.

In this case, each domain advertises reachability to its prefix, listing acceptable source prefixes. Domains that are willing to be generally reached might advertise `::/0` as a source prefix, or the prefix in use in the general domain.

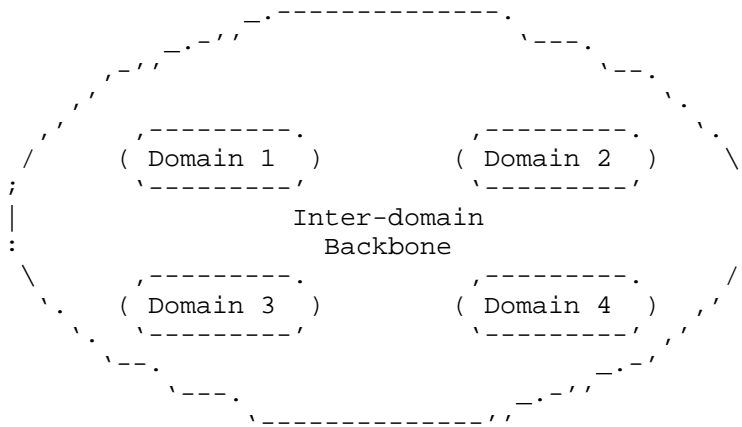


Figure 4: Intradomain Access Control

The big issue in this network is a difference in policy.

3. Derived Requirements

The use cases in can each be met if:

- o The routing protocol or mechanism includes a source prefix. It is acceptable that a default source prefix of `::/0` (all addresses) applies to routes that don't specify a prefix.
- o The routing protocol or mechanism includes a destination prefix, which may be a default route (`::/0`) or any more specific prefix up to and including a host route (`/128`).
- o The FIB lookup yields the route with the most specific (e.g. longest-match) destination prefix that also matches the source prefix constraint, or no match.

4. IANA Considerations

This memo asks the IANA for no new parameters.

5. Security Considerations

As a descriptive document, this note adds no new security risks to the network.

6. Privacy Considerations

As a descriptive document, this note adds no new privacy risks to the network.

7. Acknowledgements

This note was discussed with Acee Lindem, Jianping Wu, Juliusz Chroboczek, Les Ginsberg, Lorenzo Colitti, Mark Townsley, Markus Stenberg, Matthieu Boutier, Ole Troan, Ray Bellis, Shu Yang, and Xia Yin.

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

[I-D.baker-fun-routing-class]

Baker, F., "Routing a Traffic Class", draft-baker-fun-routing-class-00 (work in progress), July 2011.

[I-D.baker-ipv6-isis-dst-src-routing]

Baker, F., "IPv6 Source/Destination Routing using IS-IS", draft-baker-ipv6-isis-dst-src-routing-00 (work in progress), February 2013.

[I-D.baker-ipv6-ospf-dst-src-routing]

Baker, F., "IPv6 Source/Destination Routing using OSPFv3", draft-baker-ipv6-ospf-dst-src-routing-02 (work in progress), May 2013.

[I-D.boutier-homenet-source-specific-routing]

Boutier, M. and J. Chroboczek, "Source-specific Routing", draft-boutier-homenet-source-specific-routing-00 (work in progress), July 2013.

[I-D.troan-homenet-sadr]

Troan, O. and L. Colitti, "IPv6 Multihoming with Source Address Dependent Routing (SADR)", draft-troan-homenet-sadr-00 (work in progress), February 2013.

[I-D.xu-homenet-traffic-class]

Xu, M., Yang, S., Wu, J., and F. Baker, "Traffic Class Routing Protocol in Home Networks", draft-xu-homenet-traffic-class-00 (work in progress), July 2013.

[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.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

Appendix A. Change Log

Initial Version: August 2013

Author's Address

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

Email: fred@cisco.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 21, 2014

M. Behringer
M. Pritikin
S. Bjarnason
Cisco
October 18, 2013

Bootstrapping Trust on a Homenet
draft-behringer-homenet-trust-bootstrap-01.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 enrol 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 21, 2014.

Copyright Notice

Copyright (c) 2013 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. Problem Statement	2
2. Approach	2
2.1. Summary of the approach	3
2.2. Autonomic devices	3
2.3. User interface	3
2.4. The Registrar	3
2.5. Validating a device identity	4
2.6. Claiming a device	5
2.7. Services	6
2.8. Network boundaries	6
3. Security Considerations	6
4. Informative References	7
Authors' Addresses	7

1. Problem Statement

[I-D.ietf-homenet-arch] states that "A homenet will most likely also have internal borders between internal realms, e.g. a guest realm or a corporate network extension realm. It should be possible to automatically discover these 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. There should be a way for a user to administratively assert in a simple way whether or not a device belongs to a homenet."

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

2. Approach

This approach is based on making homenet devices behave in autonomic mode where devices discover each others and autonomically establish trust boundaries. See [I-D.behringer-autonomic-network-framework] for more information.

2.1. Summary of the approach

In short, the approach is:

- o The user pairs a smart phone (or similar device) with one of the devices in the homenet, for example the CPE. The smart phone acts as a user interface only.
- o The selected device becomes the trust anchor of the homenet. Technically, it acts as a certification authority (CA).
- o Devices in the homenet use a protocol to exchange identities.
- o A new device is added to the homenet by the user accepting it on the smart phone, and the CA issuing a domain certificate to the new device.
- o The boundary of the network is determined by checking the certificates of devices.

2.2. Autonomic devices

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.

2.3. User interface

The user interface can be provided by the devices themselves or through a smart phone interface. It is also possible to access the devices indirectly through the manufactures web site. Options are:

- o The user connects a PC to a physical port on network device and gets access to devices's user interface.
- o The user scans a QR code on the device using his smartphone. This will trigger the download of the manufactures autonomic app which will allow the user to connect to the device using wireless access.

2.4. The Registrar

One autonomic device in the homenet takes on a registrar function. This could be enabled using the 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 domain, and subsequently acts as a certification authority, granting domain certificates to other devices.

The user can configure a device as the registrar using one of the following options:

- o By using a smartphone app which is automatically downloaded when scanning a QR code on the device. This will then allow the user to connect to the device on an SSID which is dynamically created based on the device serial number. The device will only allow connections from smartphones using the manufactures app.
- o By connecting a PC to a physical port on the network device and gaining access to devices's user interface.

2.5. Validating a device identity

Every autonomic device discovers neighbouring autonomic nodes through an autonomic secure neighbour discovery protocol. This could be implemented for example through IPv6 secure 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 the homenet 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. This may be validated by the user. Any authorised 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 users.

The registrar can also decide to accept the device based on alternate criteria:

- o Allow any device to join within a specific time period.
- o Allow only devices with specific serial numbers to join. These can either be entered manually into the registrar or by scanning a QR code using the manufactures autonomic app on a smartphone.
- o If the device has a vendor certificate (e.g., an IEEE 802.1AR [IDevID] credential), the device can be validated using a Cloud service from the vendor.

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

A device MAY require an invitation to be signed by the manufacturer, stating that it has been claimed by the user before it decides to join the domain.

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.

2.6. Claiming a device

A device can be claimed using one of the following options:

- o Presenting a manufacturer signed "claim" over the network interface.
- o Connecting to a physical port on the network device and inserting the domain identity (public key).

Any registrar can contact the manufacturer or other trusted-by-the-device cloud resource to obtain a claim on a device. This does not require the device to be online. The claim is issued by the cloud

resource in a non-discriminatory fashion to the unauthenticated registrar. Claims can include a nonce generated by the device. The registrar may drop the nonce. The cloud service may drop the nonce. If the nonce is included in the resulting claim the device must verify this value against the current device state.

The cloud resource should offer open and non-discriminatory audit functionalities associating the privacy protected registrar public key information with the device identity and any nonce information included.

2.7. Services

As the devices have a common trust anchor, device identity can be securely established, making it possible to automatically deploy services across the domain in a secure manner.

Examples of services:

- o Device management.
- o Routing authentication.
- o Service discovery.

2.8. Network boundaries

When a device has joined the domain, it can validate the domain membership of other devices. This makes it possible to create trust boundaries where domain members have higher level of trusted than external devices. Using the autonomic User Interface, specific devices can be grouped into to sub domains and specific trust levels can be implemented between those.

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.

There are counter-measures against these attacks, with various security levels, and corresponding various ease of use. The options are (in order of increased security):

- o Only allow new devices to join in a specific time period.
- o Only allow specific devices to join by matching their serial numbers.

- o Validating the vendor certificate on new devices using the vendors Cloud portal.

In order to support a variety of use cases, devices can be claimed by a registrar without proving possession of the device in question. This would result in a nonceless, and thus always valid, claim. Future registrars are recommended to take the audit history of a device into account when deciding to join the device into their network.

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-10 (work in progress), August 2013.

[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

Max Pritikin
Cisco

Email: pritikin@cisco.com

Steinthor Bjarnason
Cisco

Email: sbjarnas@cisco.com

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 24, 2014

Z. Cao
China Mobile
A. Ding
Cambridge University / Helsinki University
October 21, 2013

Service Discovery in the Homenet Environment with Multiple Connections
draft-cao-homenet-mif-srvdis-00

Abstract

This document analyzes the problems of service discovery in a homenet multiple connection environment. A multiple connection environment consists of multiple-interfaced nodes connecting to multiple networks or multiple provisioning domains. Given a type of service a multiple-interfaced client is looking for, the discovery progress ought to return a correct pointer to the service instance that the client is able to access without trying every available channel.

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 24, 2014.

Copyright Notice

Copyright (c) 2013 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	2
2. Requirements and Terminology	3
2.1. Requirements	3
2.2. Terminology	3
3. Scenarios	3
3.1. Mif Scenario	3
3.2. Homenet Scenario	5
4. Problem Analysis	5
5. IANA Considerations	9
6. Security Considerations	9
7. References	9
7.1. Normative References	9
7.2. Informative References	9
Authors' Addresses	10

1. Introduction

A multihomed host has multiple provisioning domains via physical and/or virtual interfaces. A multihomed host receives node configuration information from each of its access networks, through various mechanisms such as DHCP, PPP and IPv6 Router Advertisements. When the received node-scoped configuration objects have different values from each administration domains, such as different DNS servers IP addresses, different default gateways or different address selection policies, the node has to decide which it will use or how it will merge them.

Issues regarding how the multi-homed host uses the configuration objects have been addressed in [RFC6418]. Current practices of how the various implementations handle these problems are introduced in [RFC6419]. [RFC6731] extends DHCPv6 to inform the host which DNS server it ought to select to send the query request, and DNS based Service Discovery (DNS-SD) has been specified in [RFC6763].

This document analyzes the problem of service discovery in a multiple connection environment. A multiple connection environment consists of multiple-interfaces nodes connecting to multiple networks or multiple provisioning domains. Given a type of service a multiple-interfaced client is looking for, the discovery progress ought to return a correct pointer to the service instance that the a client is able to access.

2. Requirements and Terminology

2.1. Requirements

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].

2.2. Terminology

Service Domain

A set of services that can be accessed by users. Besides providing services, a service domain is responsible for delivering configuration and pointers that ensure a guaranteed service access.

Service Discovery

Procedure to acquire information that is necessary to access service.

Multiple Connection Environment

Consists of multiple-interfaced nodes that connect to multiple networks or multiple provisioning domains.

3. Scenarios

We describe two scenarios in this section, one related to Multiple Interfaces, and the other one related to Home Networks (homenet).

3.1. Mif Scenario

The service discovery process can be summarized as the following five steps.

1. Service Discovery Preparation: the host determines which interface it should send a query request based on the configuration information.
2. Service Query Request: the host sends a query request to find a service. The query should include a description of the service, for example, a full-qualified domain name, a URI, or an application-specific naming of the service.
3. Service Request Handling: any entity that receives the query request should handle the request. The entity should understand

the meaning of the request, and check the semantics of the request language before giving an answer back.

4. Service Query Response: the entity that receives the query request should reply with an answer to the query. The answer should include a pointer to the service.
5. Service Access: the host accesses the service via the pointer provided in the query response. The host is supposed to be able to get the service instance via the pointer under a successful and efficient service discovery mechanism, unless the servers in such service domain encounter problems e.g. a web server is down.

Figure 1 shows a typical scenario for service discovery in a multiple connection environment. It is common in today's mobile Internet that a host is equipped with multiple network interfaces. On the service domain, different services are deployed and some services may not be accessible to a certain interface on the host due to security concern or access policy. The connectivity each interface provides may not be restricted to Internet access. For instance, WLAN and bluetooth can offer direct access to potential services e.g. printers via local ad-hoc connectivity. In such multiple connection environment, the service discovery process should return a correct point to the host and ensure that the host can access the service via this pointer. This situation makes the multiple interface service discovery different from the typical one-interface Internet access scenario. Furthermore, the growing usage of IPv6 in the homenet environment has made service discovery more challenging that requires thorough investigation.

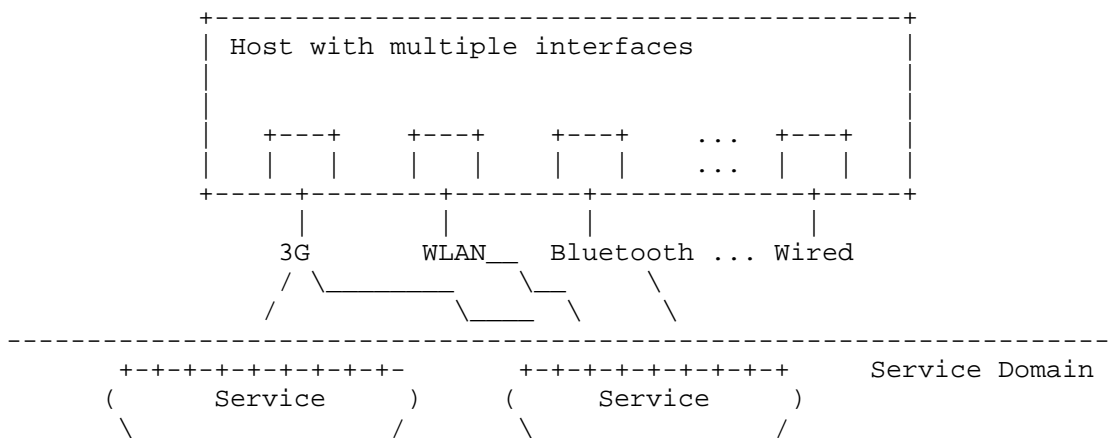


Figure 1: Multiple Interface Host with Multiple Available Services

3.2. Homenet Scenario

We also describe the issues related to the homenet architecture [I-D.ietf-homenet-arch], as depicted in Figure 2.

Suppose one MIF host is connected to three domains: homenet domain, 3gpp domain and a WiFi or enterprise domain. There is one service that is named with the private domain name, say 'temperature.ietf', which is only resolvable via the domain name service residing inside the homenet and is supported by the multicast dns service [RFC6762].

There are several problems in this scenario. First of all, since the host has two unicast dns domains configured over the 3GPP and WiFi, and as well as a multicast service discovery domain within the homenet, the host does not know which domain it should send a dns resolution request. Secondly, even if coupled with the split dns solution [RFC6731], the configuration information obtained from DHCP supports only those two unicast dns domains, but not the homenet domain which is normally considered as 'zero-configuration'. Third, the service discovery problem will become more complicated if the host is connecting to more than one home networks, i.e., multiple multicast dns domains.

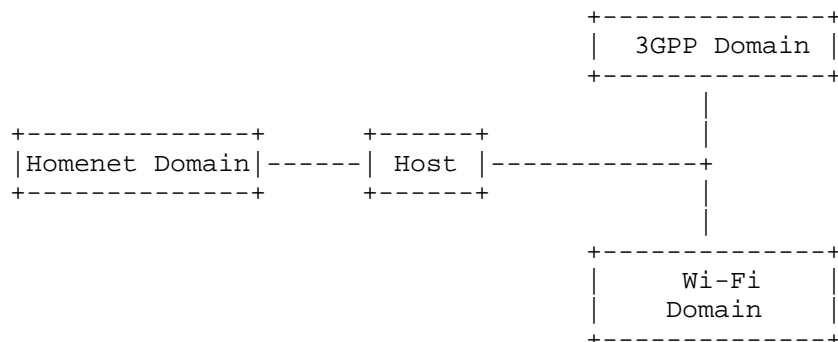


Figure 2: Homenet Scenario

4. Problem Analysis

The problems that a multiple-interfaced host may meet during the service discovery include:

1. How the query requests are sent? Because there are multiple interfaces available and multiple service rendezvous existing,

the host should decide which destination it ought to send the query to. And if there is a round-robin mechanism, the host should determine the order of the query request.

2. How to handle the query and reply? Some pointers to the service are restricted to a local scope or a certain interface, e.g., an office printer may not be accessible to the 3G-interface. The service discovery process is supposed to return a pointer that is accessible to the host.
3. How to access the service? Given the pointer to the service, the host should be able to determine from which interface it can access the service.

The existing work of [RFC6418] and [RFC6419] have covered the general problems encountered by hosts accessing multiple provisioning domains, but the focus is on connectivity and configuration. Proposal of Happy Eyeball in [I-D.ietf-mif-happy-eyeballs-extension] allows a host with multiple interfaces to pick a suitable one for access and enables automatic fallback. In a DNS based service discovery [RFC6763], the problem of domain split is analyzed in the [RFC6731]. The document defines an extension to the DHCPv4 and DHCPv6 to inform the MIF host which domain scope the Recursive DNS Server(RDNSS) is serving for, so that the "service query request" can be sent to the correct RDNSS to get an answer.

The existing proposals resolve the partial problem in the service discovery process mentioned above. To highlight the missing blocks, Figure 3 provides a 'gap' analysis. In the figure, we compare three existing solutions on service discovery, DNS-SD[RFC6763], DNS-Server-Selection [RFC6731], and MIF Happy Eyeball [I-D.ietf-mif-happy-eyeballs-extension], from three aspects as mentioned above. The DNS-Srv-Sel solution uses the defined DHCP option for the MIF host to select the corresponding DNS Server, and MIF-HE inherits this method in its most updated version. The MIF-HE can help host failover to the workable interface during service access while DNS-Srv-Sel does not handle this particular issue. The DNS-SD is not designed for a multiple interfaces environment and DNS server selection and request handling are based on standard DNS behaviors.

Aspects \ Sol	DNS-SD	DNS-Srv-Sel	MIF-HE
How to Send Query	Std. DNS behavior	DHCP Option informed	Same as DNS-Srv-Sel

How to Handle Queries	Std. DNS server behavior	selection based on option	Same as DNS-Srv-Sel
How to Access service	no guarantee of connectivity	not possible if ports rejected	Failover to the Happiest one

Figure 3: Gap Analysis of Existing Service Discovery Methods

In a complicated network as shown in Figure 4 , the host connects to the enterprise network via the wired interface, a WLAN network with the 802.11 interface, and an operator's network via the cellular interface. The three intranets have their own Firewall policies to the global Internet. On the enterprise network, many outgoing ports are restricted, and on the WLAN and operator's public network, there is more freedom. If the MIF host makes a DNS-SRV query to a service in a global domain, all the RDNS servers have the corresponding records. But say the service port number has been blocked by the enterprise network administrator, the DNS has no such information. Even if the DNS returns a pointer to the MIF host, the MIF host cannot access this service via the wired interface.

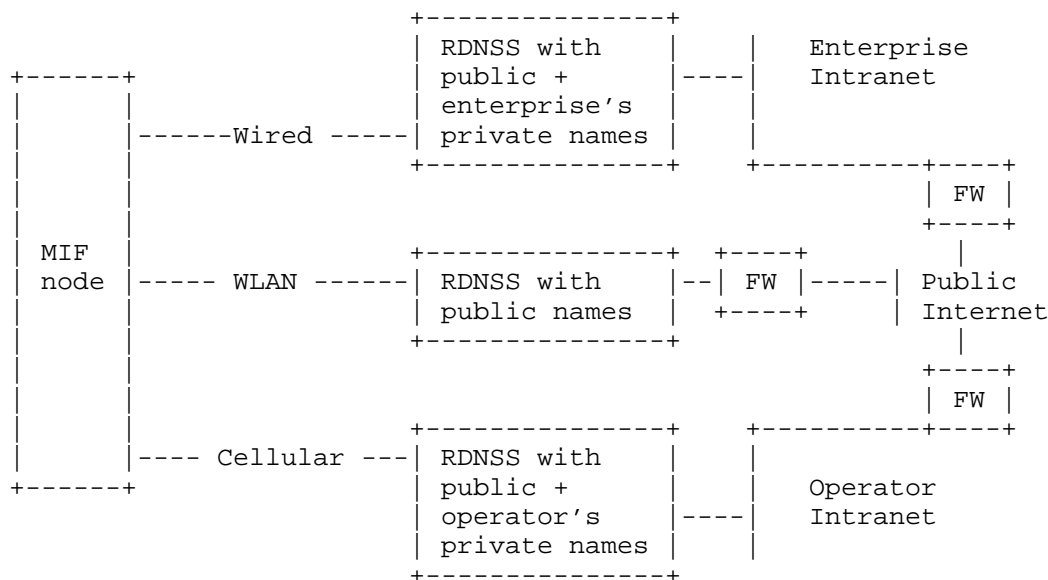


Figure 4: Scenario of Multiple Interface DNS Service Discovery

CoAP [I-D.ietf-core-coap] is an IETF designed RESTful protocol for constrained environment. CoAP defines a link-format for service discovery of the particular CoAP server, i.e., `"/.well-known/core"`. If the CoAP client has multiple access networks as shown in Figure 5, the situation turns to be more complex. For instance, if the MIF client wants to find a humidity sensing resource, but does not know which domain contains the information, it basically needs to send multiple CoAP GET requests with the well-known URL. Once it gets a response for the required resource, it can send the corresponding request to get the information. However this way is sub-optimal especially for constrained devices. MIF service discovery SHOULD consider the efficiency of the service discovery process.

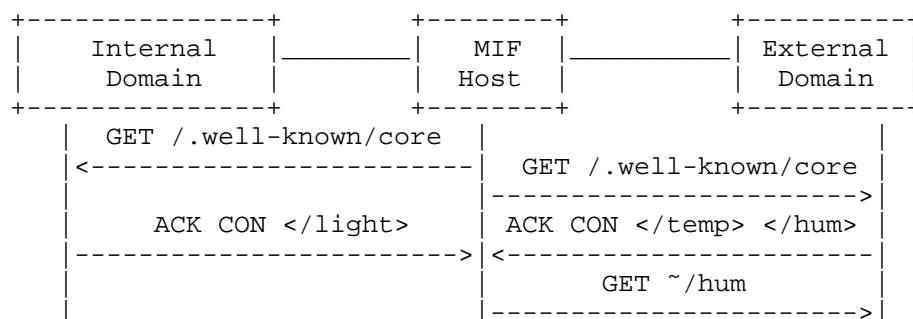


Figure 5: CoAP Service Discovery for a MIF Host

As a summary of the above analysis, the general problems and requirements of service discovery in a MIF environment can be summarized as follows:

Service Directory Service Configuration: Service directory is the entity that stores or can get the stored relationship between service names and service pointers. Different interfaces or provisioning domains have their different service directories. How to configure them on the MIF host and how the MIF host utilizes the configured information are important for the service discovery process to behave correctly.

Service Directory Selection: After the service directory information is configured on the host, the host is able to select the correct directory to send the query. The host can utilize auxiliary information available or send the query to all the directories that have been configured. The behavior of MIF host to select a correct directory is also important for a stable system.

Service Pointer/Address Resolution: The same service may have different available addresses and pointers, and some service has limited connectivity. So the resolution process should be able to return to the MIF host a record that is accessible from at least one of the interfaces. Efficiency SHOULD be taken into consideration in this phase.

Service Route Selection: With the pointers returned, the host should route the service level data to the service instance identified by the returned pointers.

5. IANA Considerations

This document has no IANA requests.

6. Security Considerations

The query response exchanges should be protected by security mechanisms. If the response contains invalid information, e.g. a pointer to a worm website, it harms. As a consequence, the service discovery should protect bogus information injected by attackers or intruders. The security consideration ought to be made by the underlining protocols, and it is out the scope of this problem statement document.

7. References

7.1. Normative References

- [I-D.ietf-mif-happy-eyeballs-extension]
Chen, G., Williams, C., Wing, D., and A. Yourtchenko,
"Happy Eyeballs Extension for Multiple Interfaces", draft-
ietf-mif-happy-eyeballs-extension-03 (work in progress),
August 2013.
- [RFC6731] Savolainen, T., Kato, J., and T. Lemon, "Improved
Recursive DNS Server Selection for Multi-Interfaced
Nodes", RFC 6731, December 2012.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762,
February 2013.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service
Discovery", RFC 6763, February 2013.

7.2. Informative References

- [I-D.ietf-core-coap]

Shelby, Z., Hartke, K., and C. Bormann, "Constrained Application Protocol (CoAP)", draft-ietf-core-coap-18 (work in progress), June 2013.

[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-10 (work in progress), August 2013.

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

[RFC6418] Blanchet, M. and P. Seite, "Multiple Interfaces and Provisioning Domains Problem Statement", RFC 6418, November 2011.

[RFC6419] Wasserman, M. and P. Seite, "Current Practices for Multiple-Interface Hosts", RFC 6419, November 2011.

Authors' Addresses

Zhen Cao
China Mobile
Xuanwumenxi Ave. No. 32
Beijing 100871
China

Phone: +86-10-52686688
Email: zehn.cao@gmail.com, caozhen@chinamobile.com

Aaron Yi Ding
Cambridge University / Helsinki University
William Gates Building, 15 JJ Thomson Ave
CB3 0FD Cambridge
United Kingdom

Phone: +44-7934034801; +358-9-19151296
Email: Aaron.Ding@cl.cam.ac.uk

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 25, 2014

T. Chown, Ed.
University of Southampton
J. Arkko
Ericsson
A. Brandt
Sigma Designs
O. Troan
Cisco Systems, Inc.
J. Weil
Time Warner Cable
October 22, 2013

IPv6 Home Networking Architecture Principles
draft-ietf-homenet-arch-11

Abstract

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

Status of this Memo

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

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

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

This Internet-Draft will expire on April 25, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Terminology and Abbreviations	5
2. Effects of IPv6 on Home Networking	6
2.1. Multiple subnets and routers	7
2.2. Global addressability and elimination of NAT	8
2.3. Multi-Addressing of devices	8
2.4. Unique Local Addresses (ULAs)	9
2.5. Avoiding manual configuration of IP addresses	10
2.6. IPv6-only operation	11
3. Homenet Architecture Principles	11
3.1. General Principles	12
3.1.1. Reuse existing protocols	12
3.1.2. Minimise changes to hosts and routers	12
3.2. Homenet Topology	13
3.2.1. Supporting arbitrary topologies	13
3.2.2. Network topology models	13
3.2.3. Dual-stack topologies	18
3.2.4. Multihoming	19
3.3. A Self-Organising Network	20
3.3.1. Differentiating neighbouring homenets	21
3.3.2. Largest practical subnets	21
3.3.3. Handling varying link technologies	22
3.3.4. Homenet realms and borders	22
3.3.5. Configuration information from the ISP	23
3.4. Homenet Addressing	23
3.4.1. Use of ISP-delegated IPv6 prefixes	23
3.4.2. Stable internal IP addresses	25
3.4.3. Internal prefix delegation	26
3.4.4. Coordination of configuration information	27
3.4.5. Privacy	28

3.5.	Routing functionality	28
3.5.1.	Multicast support	29
3.5.2.	Mobility support	30
3.6.	Security	30
3.6.1.	Addressability vs reachability	31
3.6.2.	Filtering at borders	31
3.6.3.	Partial Effectiveness of NAT and Firewalls	32
3.6.4.	Exfiltration concerns	32
3.6.5.	Device capabilities	32
3.6.6.	ULAs as a hint of connection origin	33
3.7.	Naming and Service Discovery	33
3.7.1.	Discovering services	33
3.7.2.	Assigning names to devices	34
3.7.3.	The homenet name service	35
3.7.4.	Name spaces	36
3.7.5.	Independent operation	38
3.7.6.	Considerations for LLNs	38
3.7.7.	DNS resolver discovery	38
3.7.8.	Devices roaming to/from the homenet	39
3.8.	Other Considerations	39
3.8.1.	Quality of Service	39
3.8.2.	Operations and Management	39
3.9.	Implementing the Architecture on IPv6	40
4.	Conclusions	41
5.	Security Considerations	41
6.	IANA Considerations	41
7.	References	41
7.1.	Normative References	41
7.2.	Informative References	42
Appendix A.	Acknowledgments	44
Appendix B.	Changes	45
B.1.	Version 11 (after IESG review)	45
B.2.	Version 10 (after AD review)	45
B.3.	Version 09 (after WGLC)	45
B.4.	Version 08	46
B.5.	Version 07	46
B.6.	Version 06	47
B.7.	Version 05	47
B.8.	Version 04	48
B.9.	Version 03	48
B.10.	Version 02	49
Authors' Addresses	50

1. Introduction

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

While at the time of writing some complex home network topologies exist, most are relatively simple single subnet networks, and ostensibly operate using just IPv4. While there may be IPv6 traffic within the network, e.g., for service discovery, the homenet is provisioned by the ISP as an IPv4 network. Such networks also typically employ solutions that should be avoided, such as private [RFC1918] addressing with (cascaded) network address translation (NAT) [RFC3022], or they may require expert assistance to set up.

In contrast, emerging IPv6-capable home networks are very likely to have multiple internal subnets, e.g., to facilitate private and guest networks, heterogeneous link layers, and smart grid components, and have enough address space available to allow every device to have a globally unique address. This implies that internal routing functionality is required, and that the homenet's ISP both provides a large enough prefix to allocate a prefix to each subnet, and that a method is supported for such prefixes to be delegated efficiently to those subnets.

It is not practical to expect home users to configure their networks. Thus the assumption of this document is that the homenet is as far as possible self-organising and self-configuring, i.e., it should function without pro-active management by the residential user.

The architectural constructs in this document are focused on the problems to be solved when introducing IPv6, with an eye towards a better result than what we have today with IPv4, as well as aiming at a more consistent solution that addresses as many of the identified requirements as possible. The document aims to provide the basis and guiding principles for how standard IPv6 mechanisms and addressing [RFC2460] [RFC4291] can be employed in home networking, while coexisting with existing IPv4 mechanisms. In emerging dual-stack home networks it is vital that introducing IPv6 does not adversely affect IPv4 operation. We assume that the IPv4 network architecture in home networks is what it is, and can not be modified by new

recommendations. This document does not discuss how IPv4 home networks provision or deliver support for multiple subnets. It should not be assumed that any future new functionality created with IPv6 in mind will be backward-compatible to include IPv4 support. Further, future deployments, or specific subnets within an otherwise dual-stack home network, may be IPv6-only, in which case considerations for IPv4 impact would not apply.

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

[RFC6204] defines basic requirements for customer edge routers (CERs). This document has recently been updated with the definition of requirements for specific transition tools on the CER in [I-D.ietf-v6ops-6204bis], specifically DS-Lite [RFC6333] and 6rd [RFC5969]. Such detailed specification of CER devices is considered out of scope of this architecture document, and we assume that any required update of the CER device specification as a result of adopting this architecture will be handled as separate and specific updates to these existing documents. Further, the scope of this text is the internal homenet, and thus specific features on the WAN side of the CER are out of scope for this text.

1.1. Terminology and Abbreviations

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

- o Border: a point, typically resident on a router, between two networks, e.g., between the main internal homenet and a guest network. This defines point(s) at which filtering and forwarding policies for different types of traffic may be applied.
- o CER: Customer Edge Router: A border router intended for use in a homenet, which connects the homenet to a service provider network.
- o FQDN: Fully Qualified Domain Name. A globally unique name.
- o Guest network: A part of the home network intended for use by visitors or guests to the home(net). Devices on the guest network may typically not see or be able to use all services in the

home(net).

- o Homenet: A home network, comprising host and router equipment, with one or more CERs providing connectivity to service provider network(s).
- o Internet Service Provider (ISP): an entity that provides access to the Internet. In this document, a service provider specifically offers Internet access using IPv6, and may also offer IPv4 Internet access. The service provider can provide such access over a variety of different transport methods such as DSL, cable, wireless, and others.
- o LLN: Low-power and lossy network.
- o LQDN: Locally Qualified Domain Name. A name local to the homenet.
- o NAT: Network Address Translation. Typically referring to IPv4 Network Address and Port Translation (NAPT) [RFC3022].
- o NPTv6: Network Prefix Translation for IPv6 [RFC6296].
- o PCP: Port Control Protocol [RFC6887].
- o Realm: a network delimited by a defined border. A guest network within a homenet may form one realm.
- o 'Simple Security'. Defined in [RFC4864] and expanded further in [RFC6092]; describes recommended perimeter security capabilities for IPv6 networks.
- o ULA: IPv6 Unique Local Address [RFC4193].
- o VM: Virtual machine.

2. Effects of IPv6 on Home Networking

While IPv6 resembles IPv4 in many ways, there are some notable differences in the way it may typically be deployed. It changes address allocation principles, making multi-addressing the norm, and, through the vastly increased address space, allows globally unique IP addresses to be used for all devices in a home network. This section presents an overview of some of the key implications of the introduction of IPv6 for home networking, that are simultaneously both promising and problematic.

2.1. Multiple subnets and routers

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

The introduction of IPv6 for home networking makes it possible for every home network to be delegated enough address space from its ISP to provision globally unique prefixes for each such subnet in the home. While the number of addresses in a standard /64 IPv6 prefix is practically unlimited, the number of prefixes available for assignment to the home network is not. As a result the growth inhibitor for the home network shifts from the number of addresses to the number of prefixes offered by the provider; this topic is discussed in [RFC6177] (BCP 157), which recommends that "end sites always be able to obtain a reasonable amount of address space for their actual and planned usage".

The addition of routing between subnets raises a number of issues. One is a method by which prefixes can be efficiently allocated to each subnet, without user intervention. Another is the issue of how to extend mechanisms such as zero configuration service discovery which currently only operate within a single subnet using link-local traffic. In a typical IPv4 home network, there is only one subnet, so such mechanisms would normally operate as expected. For multi-subnet IPv6 home networks there are two broad choices to enable such protocols to work across the scope of the entire homenet; extend existing protocols to work across that scope, or introduce proxies for existing link layer protocols. This topic is discussed in Section 3.7.

2.2. Global addressability and elimination of NAT

The possibility for direct end-to-end communication on the Internet to be restored by the introduction of IPv6 is on the one hand an incredible opportunity for innovation and simpler network operation, but on the other hand it is also a concern as it potentially exposes nodes in the internal networks to receipt of unwanted and possibly malicious traffic from the Internet.

With devices and applications able to talk directly to each other when they have globally unique addresses, there may be an expectation of improved host security to compensate for this. It should be noted that many devices may (for example) ship with default settings that make them readily vulnerable to compromise by external attackers if globally accessible, or may simply not have robustness designed-in because it was either assumed such devices would only be used on private networks or the device itself doesn't have the computing power to apply the necessary security methods. In addition, the upgrade cycle for devices (or their firmware) may be slow, and/or lack auto-update mechanisms.

It is thus important to distinguish between addressability and reachability. While IPv6 offers global addressability through use of globally unique addresses in the home, whether devices are globally reachable or not would depend on any firewall or filtering configuration, and not, as is commonly the case with IPv4, the presence or use of NAT. In this respect, IPv6 networks may or may not have filters applied at their borders to control such traffic, i.e., at the homenet CER. [RFC4864] and [RFC6092] discuss such filtering, and the merits of 'default allow' against 'default deny' policies for external traffic initiated into a homenet. This document takes no position on which mode is the default, but assumes the choice for the homenet to use either mode would be available.

This topic is discussed further in Section 3.6.1.

2.3. Multi-Addressing of devices

In an IPv6 network, devices will often acquire multiple addresses, typically at least a link-local address and one or more globally unique addresses. Where a homenet is multihomed, a device would typically receive a globally unique address (GUA) from within the delegated prefix from each upstream ISP. Devices may also have an IPv4 address if the network is dual-stack, an IPv6 Unique Local Address (ULA) [RFC4193] (see below), and one or more IPv6 Privacy Addresses [RFC4941].

It should thus be considered the norm for devices on IPv6 home

networks to be multi-addressed, and to need to make appropriate address selection decisions for the candidate source and destination address pairs for any given connection. Default Address Selection for IPv6 [RFC6724] provides a solution for this, though it may face problems in the event of multihoming where, as described above, nodes will be configured with one address from each upstream ISP prefix. In such cases the presence of upstream BCP 38 [RFC2827] ingress filtering requires multi-addressed nodes to select the correct source address to be used for the corresponding uplink. A challenge here is that the node may not have the information it needs to make that decision based on addresses alone. We discuss this challenge in Section 3.2.4.

2.4. Unique Local Addresses (ULAs)

[RFC4193] defines Unique Local Addresses (ULAs) for IPv6 that may be used to address devices within the scope of a single site. Support for ULAs for IPv6 CERNs is described in [RFC6204]. A home network running IPv6 should deploy ULAs alongside its globally unique prefix(es) to allow stable communication between devices (on different subnets) within the homenet where that externally allocated globally unique prefix may change over time, e.g., due to renumbering within the subscriber's ISP, or where external connectivity may be temporarily unavailable. A homenet using provider-assigned global addresses is exposed to its ISP renumbering the network to a much larger degree than before whereas, for IPv4, NAT isolated the user against ISP renumbering to some extent.

While setting up a network there may be a period where it has no external connectivity, in which case ULAs would be required for inter-subnet communication. In the case where home automation networks are being set up in a new home/deployment (as early as during construction of the home), such networks will likely need to use their own /48 ULA prefix. Depending upon circumstances beyond the control of the owner of the homenet, it may be impossible to renumber the ULA used by the home automation network so routing between ULA /48s may be required. Also, some devices, particularly constrained devices, may have only a ULA (in addition to a link-local), while others may have both a GUA and a ULA.

Note that unlike private IPv4 RFC 1918 space, the use of ULAs does not imply use of an IPv6 equivalent of a traditional IPv4 NAT [RFC3022], or of NPTv6 prefix-based NAT [RFC6296]. When an IPv6 node in a homenet has both a ULA and a globally unique IPv6 address, it should only use its ULA address internally, and use its additional globally unique IPv6 address as a source address for external communications. This should be the natural behaviour given support for Default Address Selection for IPv6 [RFC6724]. By using such

globally unique addresses between hosts and devices in remote networks, the architectural cost and complexity, particularly to applications, of NAT or NPTv6 translation is avoided. As such, neither IPv6 NAT or NPTv6 is recommended for use in the homenet architecture. Further, the homenet border router(s) should filter packets with ULA source/destination addresses as discussed in Section 3.4.2.

Devices in a homenet may be given only a ULA as a means to restrict reachability from outside the homenet. ULAs can be used by default for devices that, without additional configuration (e.g., via a web interface), would only offer services to the internal network. For example, a printer might only accept incoming connections on a ULA until configured to be globally reachable, at which point it acquires a global IPv6 address and may be advertised via a global name space.

Where both a ULA and a global prefix are in use, the ULA source address is used to communicate with ULA destination addresses when appropriate, i.e., when the ULA source and destination lie within the /48 ULA prefix(es) known to be used within the same homenet. In cases where multiple /48 ULA prefixes are in use within a single homenet (perhaps because multiple homenet routers each independently auto-generate a /48 ULA prefix and then share prefix/routing information), utilising a ULA source address and a ULA destination address from two disjoint internal ULA prefixes is preferable to using GUAs.

While a homenet should operate correctly with two or more /48 ULAs enabled, a mechanism for the creation and use of a single /48 ULA prefix is desirable for addressing consistency and policy enforcement.

A counter-argument to using ULAs is that it is undesirable to aggressively deprecate global prefixes for temporary loss of connectivity, so for a host to lose its global address there would have to be a connection breakage longer than the lease period, and even then, deprecating prefixes when there is no connectivity may not be advisable. However, it is assumed in this architecture that homenets should support and use ULAs.

2.5. Avoiding manual configuration of IP addresses

Some IPv4 home networking devices expose IPv4 addresses to users, e.g., the IPv4 address of a home IPv4 CER that may be configured via a web interface. In potentially complex future IPv6 homenets, users should not be expected to enter IPv6 literal addresses in devices or applications, given their much greater length and the apparent randomness of such addresses to a typical home user. Thus, even for

the simplest of functions, simple naming and the associated (minimal, and ideally zero configuration) discovery of services is imperative for the easy deployment and use of homenet devices and applications.

2.6. IPv6-only operation

It is likely that IPv6-only networking will be deployed first in new home network deployments, often referred to as 'greenfield' scenarios, where there is no existing IPv4 capability, or perhaps as one element of an otherwise dual-stack network. Running IPv6-only adds additional requirements, e.g., for devices to get configuration information via IPv6 transport (not relying on an IPv4 protocol such as IPv4 DHCP), and for devices to be able to initiate communications to external devices that are IPv4-only.

Some specific transition technologies which may be deployed by the homenet's ISP are discussed in [I-D.ietf-v6ops-6204bis]. In addition, certain other functions may be desirable on the CER, e.g., to access content in the IPv4 Internet, NAT64 [RFC6144] and DNS64 [RFC6145] may be applicable.

The widespread availability of robust solutions to these types of requirements will help accelerate the uptake of IPv6-only homenets. The specifics of these are however beyond the scope of this document, especially those functions that reside on the CER.

3. Homenet Architecture Principles

The aim of this text is to outline how to construct advanced IPv6-based home networks involving multiple routers and subnets using standard IPv6 addressing and protocols [RFC2460] [RFC4291] as the basis. As described in Section 3.1, solutions should as far as possible re-use existing protocols, and minimise changes to hosts and routers, but some new protocols, or extensions, are likely to be required. In this section, we present the elements of the proposed home networking architecture, with discussion of the associated design principles.

In general, home network equipment needs to be able to operate in networks with a range of different properties and topologies, where home users may plug components together in arbitrary ways and expect the resulting network to operate. Significant manual configuration is rarely, if at all, possible, or even desirable given the knowledge level of typical home users. Thus the network should, as far as possible, be self-configuring, though configuration by advanced users should not be precluded.

The homenet needs to be able to handle or provision at least

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

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

3.1. General Principles

There is little that the Internet standards community can do about the physical topologies or the need for some networks to be separated at the network layer for policy or link layer compatibility reasons. However, there is a lot of flexibility in using IP addressing and inter-networking mechanisms. This text discusses how such flexibility should be used to provide the best user experience and ensure that the network can evolve with new applications in the future. The principles described in this text should be followed when designing homenet protocol solutions.

3.1.1. Reuse existing protocols

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

3.1.2. Minimise changes to hosts and routers

In order to maximise deployability of new homenets, where possible any requirement for changes to hosts and routers should be minimised, though solutions which, for example, incrementally improve capability with host or router changes may be acceptable. There may be cases where changes are unavoidable, e.g., to allow a given homenet routing protocol to be self-configuring.

3.2. Homenet Topology

This section considers homenet topologies, and the principles that may be applied in designing an architecture to support as wide a range of such topologies as possible.

3.2.1. Supporting arbitrary topologies

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

There should be no topology scenarios which cause loss of connectivity, except when the user creates a physical island within the topology. Some potentially pathological cases that can be created include bridging ports of a router together, however this case can be detected and dealt with by the router. Loops within a routed topology are in a sense good in that they offer redundancy. Bridging loops can be dangerous but are also detectable when a switch learns the MAC of one of its interfaces on another or runs a spanning tree or link state protocol. It is only loops using simple repeaters that are truly pathological.

The topology of the homenet may change over time, due to the addition or removal of equipment, but also due to temporary failures or connectivity problems. In some cases this may lead to, for example, a multihomed homenet being split into two isolated homenets, or, after such a fault is remedied, two isolated parts reconfiguring back to a single network.

3.2.2. Network topology models

Most IPv4 home network models at the time of writing tend to be relatively simple, typically a single NAT router to the ISP and a single internal subnet but, as discussed earlier, evolution in network architectures is driving more complex topologies, such as the separation of guest and private networks. There may also be some cascaded IPv4 NAT scenarios, which we mention in the next section. For IPv6 homenets, the Network Architectures described in [RFC6204] and its successor [I-D.ietf-v6ops-6204bis] should, as a minimum, be supported.

There are a number of properties or attributes of a home network that

we can use to describe its topology and operation. The following properties apply to any IPv6 home network:

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

Various forms of multihoming are likely to become more prevalent with IPv6 home networks, where the homenet may have two or more external ISP connections, as discussed further below. Thus the following properties should also be considered for such networks:

- o Number of upstream providers. The majority of home networks today consist of a single upstream ISP, but it may become more common in the future for there to be multiple ISPs, whether for resilience or provision of additional services. Each would offer its own prefix. Some may or may not provide a default route to the public Internet.
- o Number of CERs. The homenet may have a single CER, which might be used for one or more providers, or multiple CERs. The presence of multiple CERs adds additional complexity for multihoming scenarios, and protocols like PCP that may need to manage connection-oriented state mappings on the same CER as used for subsequent traffic flows.

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

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

Figure 1 shows a home network with multiple local area networks. These may be needed for reasons relating to different link layer technologies in use or for policy reasons, e.g., classic Ethernet in

one subnet and a LLN link layer technology in another. In this example there is no single router that a priori understands the entire topology. The topology itself may also be complex, and it may not be possible to assume a pure tree form, for instance (because home users may plug routers together to form arbitrary topologies including loops).

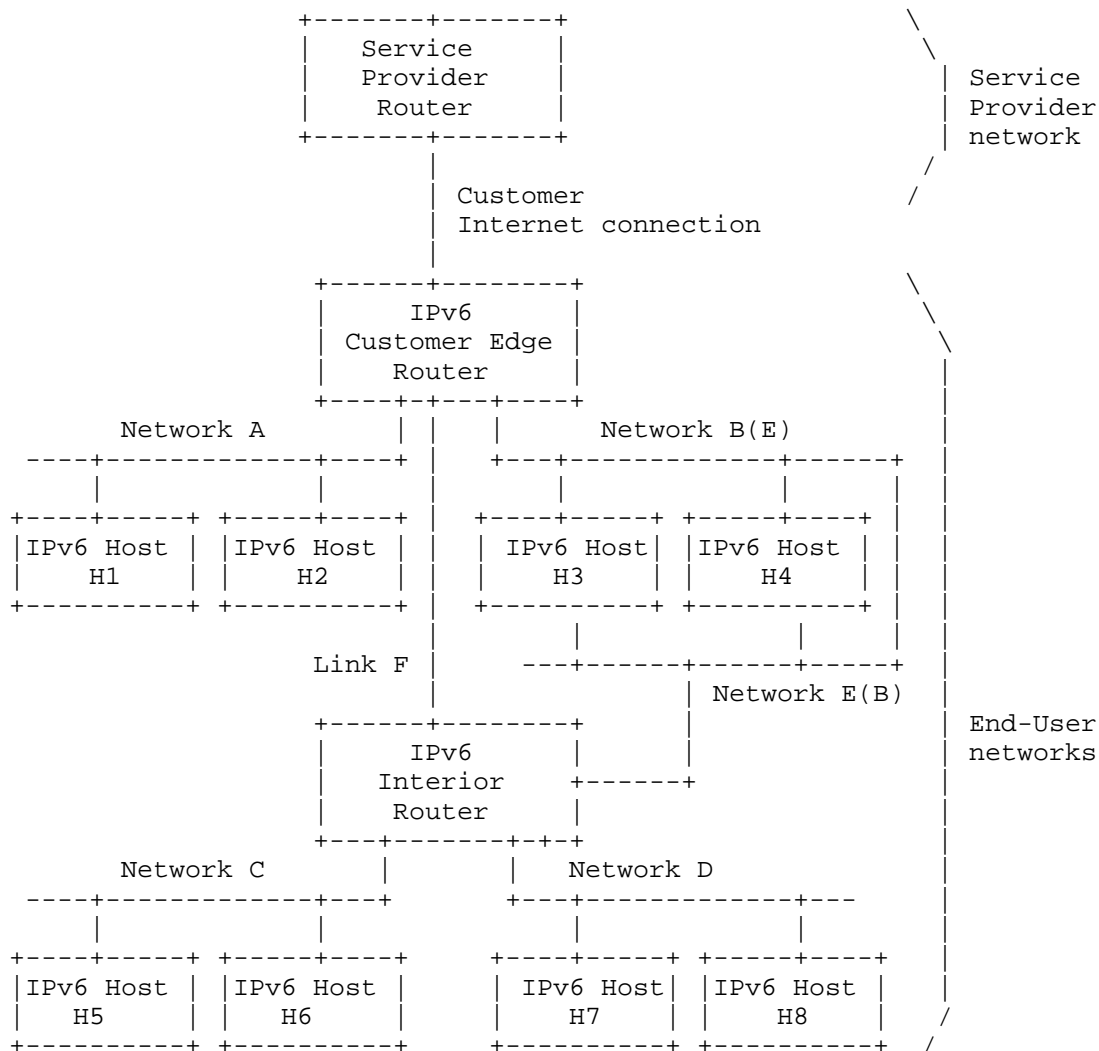


Figure 1

In this diagram there is one CER. It has a single uplink interface. It has three additional interfaces connected to Network A, Link F, and Network B. IPv6 Internal Router (IR) has four interfaces connected to Link F, Network C, Network D and Network E. Network B and Network E have been bridged, likely inadvertently. This could be as a result of connecting a wire between a switch for Network B and a switch for Network E.

Any of logical Networks A through F might be wired or wireless.

Where multiple hosts are shown, this might be through one or more physical ports on the CER or IPv6 (IR), wireless networks, or through one or more layer-2 only Ethernet switches.

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

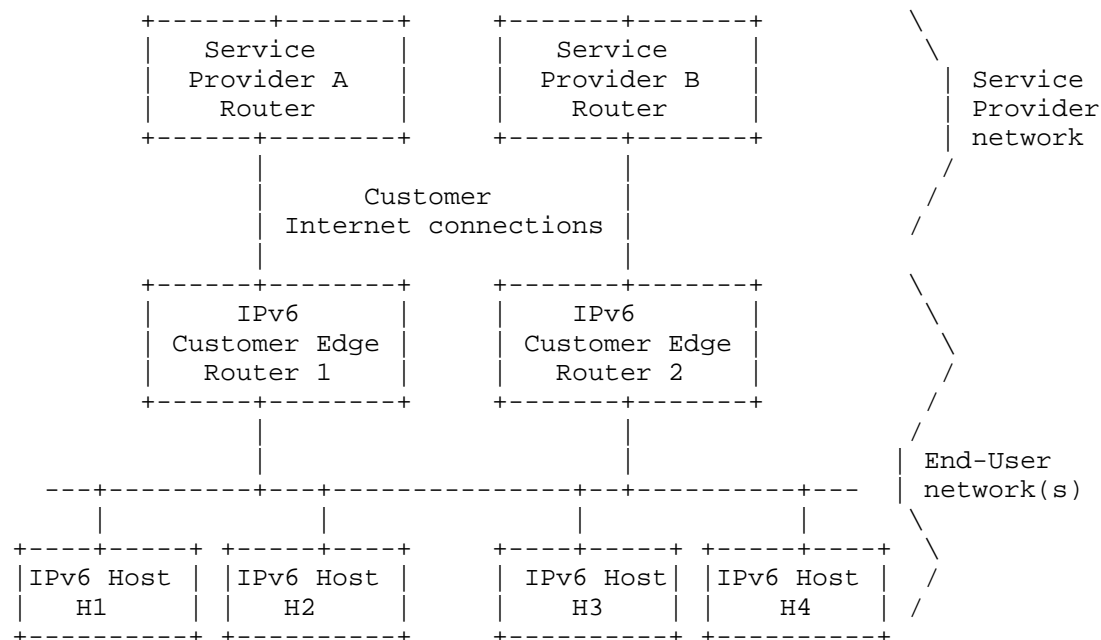


Figure 2

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

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

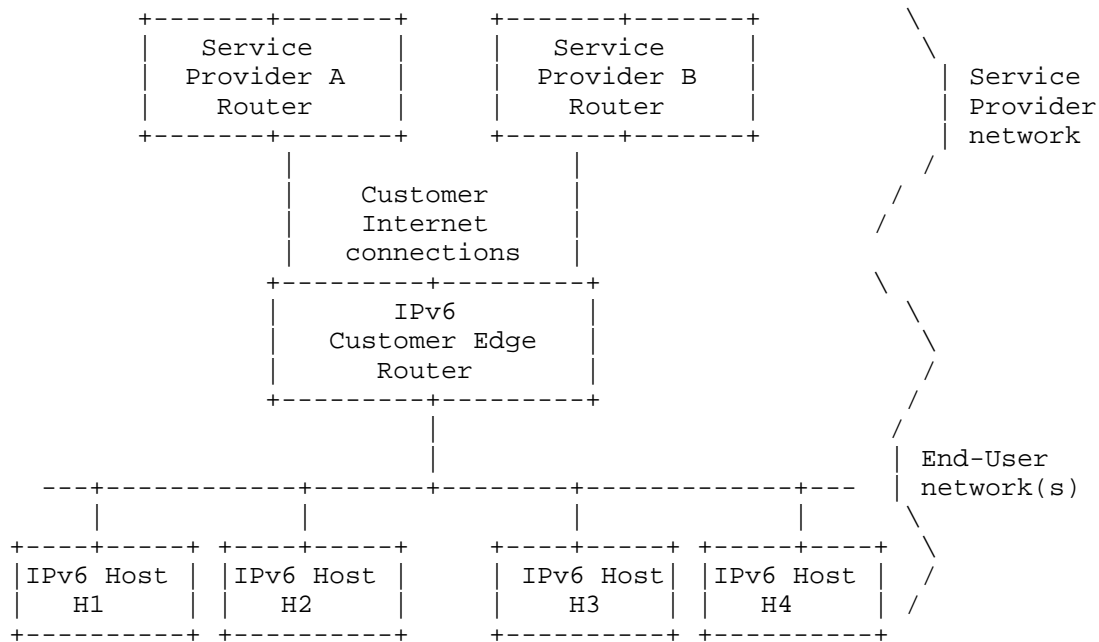


Figure 3

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

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

3.2.3. Dual-stack topologies

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

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

IPv6 where an IPv4 NAT is used and, where there is no NAT, routing or bridging may be used. Routing may have advantages when compared to bridging together high speed and lower speed shared media, and in addition bridging may not be suitable for some networks, such as ad-hoc mobile networks.

In some cases IPv4 home networks may feature cascaded NATs. End users are frequently unaware that they have created such networks as 'home routers' and 'home switches' are frequently confused. In addition, there are cases where NAT routers are included within Virtual Machine Hypervisors, or where Internet connection sharing services have been enabled. This document applies equally to such hidden NAT 'routers'. IPv6 routed versions of such cases will be required. We should thus also note that routers in the homenet may not be separate physical devices; they may be embedded within other devices.

3.2.4. Multihoming

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

In the general homenet architecture, multihomed hosts should be multi-addressed with a global IPv6 address from the global prefix delegated from each ISP they communicate with or through. When such multi-addressing is in use, hosts need some way to pick source and destination address pairs for connections. A host may choose a source address to use by various methods, most commonly [RFC6724]. Applications may of course do different things, and this should not be precluded.

For the single CER Network Model C illustrated above, multihoming may be offered by source-based routing at the CER. With multiple exit routers, as in CER Network Model B, the complexity rises. Given a packet with a source address on the home network, the packet must be routed to the proper egress to avoid BCP 38 ingress filtering if exiting through the wrong ISP. It is highly desirable that the packet is routed in the most efficient manner to the correct exit, though as a minimum requirement the packet should not be dropped.

The homenet architecture should support both the above models, i.e., one or more CERs. However, the general multihoming problem is broad, and solutions suggested to date within the IETF have included complex architectures for monitoring connectivity, traffic engineering,

identifier-locator separation, connection survivability across multihoming events, and so on. It is thus important that the homenet architecture should as far as possible minimise the complexity of any multihoming support.

An example of such a 'simpler' approach has been documented in [I-D.ietf-v6ops-ipv6-multihoming-without-ipv6nat]. Alternatively a flooding/routing protocol could potentially be used to pass information through the homenet, such that internal routers and ultimately end hosts could learn per-prefix configuration information, allowing better address selection decisions to be made. However, this would imply router and, most likely, host changes. Another avenue is to introduce support throughout the homenet for routing which is based on the source as well as the destination address of each packet. While greatly improving the 'intelligence' of routing decisions within the homenet, such an approach would require relatively significant router changes but avoid host changes.

As explained previously, while NPTv6 has been proposed for providing multi-homing support in networks, its use is not recommended in the homenet architecture.

It should be noted that some multihoming scenarios may see one upstream being a "walled garden", and thus only appropriate for connectivity to the services of that provider; an example may be a VPN service that only routes back to the enterprise business network of a user in the homenet. As per [RFC3002] (section 4.2.1) we do not specifically target walled garden multihoming as a goal of this document.

The homenet architecture should also not preclude use of host or application-oriented tools, e.g., Shim6 [RFC5533], MPTCP [RFC6824] or Happy Eyeballs [RFC6555]. In general, any incremental improvements obtained by host changes should give benefit for the hosts introducing them, but not be required.

3.3. A Self-Organising Network

The home network infrastructure should be naturally self-organising and self-configuring under different circumstances relating to the connectivity status to the Internet, number of devices, and physical topology. At the same time, it should be possible for advanced users to manually adjust (override) the current configuration.

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

The latter may be relevant when considering how to bootstrap a routing configuration. It is highly desirable that the number of such configurations is minimised.

3.3.1. Differentiating neighbouring homenets

It is important that self-configuration with 'unintended' devices is avoided. There should be a way for a user to administratively assert in a simple way whether or not a device belongs to a homenet. The goal is to allow the establishment of borders, particularly between two adjacent homenets, and to avoid unauthorised devices from participating in the homenet. Such an authorisation capability may need to operate through multiple hops in the homenet.

The homenet should thus support a way for a homenet owner to claim ownership of their devices in a reasonably secure way. This could be achieved by a pairing mechanism, by for example pressing buttons simultaneously on an authenticated and a new homenet device, or by an enrolment process as part of an autonomic networking environment.

3.3.2. Largest practical subnets

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

Future home networks are highly likely to have one or more internal routers and thus need multiple subnets, for the reasons described earlier. As part of the self-organisation of the network, the homenet should subdivide itself into the largest practical subnets that can be constructed within the constraints of link layer mechanisms, bridging, physical connectivity, and policy, and where applicable performance or other criteria. In such subdivisions the logical topology may not necessarily match the physical topology. This text does not, however, make recommendations on how such subdivision should occur. It is expected that subsequent documents will address this problem.

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

3.3.3. Handling varying link technologies

Homenets tend to grow organically over many years, and a homenet will typically be built over link-layer technologies from different generations. Current homenets typically use links ranging from 1Mbit/s up to 1Gbit/s, which is a three orders of magnitude throughput discrepancy. We expect this discrepancy to widen further as both high-speed and low-power technologies are deployed.

Homenet protocols should be designed to deal well with interconnecting links of very different throughputs. In particular, flows local to a link should not be flooded throughout the homenet, even when sent over multicast, and, whenever possible, the homenet protocols should be able to choose the faster links and avoid the slower ones.

Links (particularly wireless links) may also have limited numbers of transmit opportunities (txops), and there is a clear trend driven by both power and downward compatibility constraints toward aggregation of packets into these limited txops while increasing throughput. Transmit opportunities may be a system's scarcest resource and therefore also strongly limit actual throughput available.

3.3.4. Homenet realms and borders

The homenet will need to be aware of the extent of its own 'site', which will, for example, define the borders for ULA and site scope multicast traffic, and may require specific security policies to be applied. The homenet will have one or more such borders with external connectivity providers.

A homenet will most likely also have internal borders between internal realms, e.g., a guest realm or a corporate network extension realm. It should be possible to automatically discover these borders, which will determine, for example, the scope of where network prefixes, routing information, network traffic, service discovery and naming may be shared. The default mode internally should be to share everything.

It is expected that a realm would span at least an entire subnet, and thus the borders lie at routers which receive delegated prefixes within the homenet. It is also desirable, for a richer security model, that hosts are able to make communication decisions based on available realm and associated prefix information in the same way that routers at realm borders can.

A simple homenet model may just consider three types of realm and the borders between them, namely the internal homenet, the ISP and a

guest network. In this case the borders will include that from the homenet to the ISP, that from the guest network to the ISP, and that from the homenet to the guest network. Regardless, it should be possible for additional types of realms and borders to be defined, e.g., for some specific LLN-based network, such as Smart Grid, and for these to be detected automatically, and for an appropriate default policy to be applied as to what type of traffic/data can flow across such borders.

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

It should be possible for the homenet user to override any automatically determined borders and the default policies applied between them, the exception being that it may not be possible to override policies defined by the ISP at the external border.

3.3.5. Configuration information from the ISP

In certain cases, it may be useful for the homenet to get certain configuration information from its ISP. For example, the homenet DHCP server may request and forward some options that it gets from its upstream DHCP server, though the specific of the options may vary across deployments. There is potential complexity here of course should the homenet be multihomed.

3.4. Homenet Addressing

The IPv6 addressing scheme used within a homenet must conform to the IPv6 addressing architecture [RFC4291]. In this section we discuss how the homenet needs to adapt to the prefixes made available to it by its upstream ISP, such that internal subnets, hosts and devices can obtain the and configure the necessary addressing information to operate.

3.4.1. Use of ISP-delegated IPv6 prefixes

Discussion of IPv6 prefix allocation policies is included in [RFC6177]. In practice, a homenet may receive an arbitrary length IPv6 prefix from its provider, e.g., /60, /56 or /48. The offered prefix may be stable or change from time to time; it is generally

expected that ISPs will offer relatively stable prefixes to their residential customers. Regardless, the home network needs to be adaptable as far as possible to ISP prefix allocation policies, and thus make no assumptions about the stability of the prefix received from an ISP, or the length of the prefix that may be offered.

However, if, for example, only a /64 is offered by the ISP, the homenet may be severely constrained or even unable to function. [RFC6177] (BCP 157) states that "a key principle for address management is that end sites always be able to obtain a reasonable amount of address space for their actual and planned usage, and over time ranges specified in years rather than just months. In practice, that means at least one /64, and in most cases significantly more. One particular situation that must be avoided is having an end site feel compelled to use IPv6-to-IPv6 Network Address Translation or other burdensome address conservation techniques because it could not get sufficient address space." This architecture document assumes that the guidance in the quoted text is being followed by ISPs.

There are many problems that would arise from a homenet not being offered a sufficient prefix size for its needs. Rather than attempt to contrive a method for a homenet to operate in a constrained manner when faced with insufficient prefixes, such as the use of subnet prefixes longer than /64 (which would break stateless address autoconfiguration [RFC4862]), use of NPTv6, or falling back to bridging across potentially very different media, it is recommended that the receiving router instead enters an error state and issues appropriate warnings. Some consideration may need to be given to how such a warning or error state should best be presented to a typical home user.

Thus a homenet CER should request, for example via DHCP Prefix Delegation (DHCP PD) [RFC3633], that it would like a /48 prefix from its ISP, i.e., it asks the ISP for the maximum size prefix it might expect to be offered, even if in practice it may only be offered a /56 or /60. For a typical IPv6 homenet, it is not recommended that an ISP offer less than a /60 prefix, and it is highly preferable that the ISP offers at least a /56. It is expected that the allocated prefix to the homenet from any single ISP is a contiguous, aggregated one. While it may be possible for a homenet CER to issue multiple prefix requests to attempt to obtain multiple delegations, such behaviour is out of scope of this document.

The norm for residential customers of large ISPs may be similar to their single IPv4 address provision; by default it is likely to remain persistent for some time, but changes in the ISP's own provisioning systems may lead to the customer's IP (and in the IPv6 case their prefix pool) changing. It is not expected that ISPs will

generally support Provider Independent (PI) addressing for residential homenets.

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

There may be cases where local law means some ISPs are required to change IPv6 prefixes (current IPv4 addresses) for privacy reasons for their customers. In such cases it may be possible to avoid an instant 'flash' renumbering and plan a non-flag day renumbering as per RFC 4192. Similarly, if an ISP has a planned renumbering process, it may be able to adjust lease timers, etc appropriately.

The customer may of course also choose to move to a new ISP, and thus begin using a new prefix. In such cases the customer should expect a discontinuity, and not only may the prefix change, but potentially also the prefix length if the new ISP offers a different default size prefix. The homenet may also be forced to renumber itself if significant internal 'replumbing' is undertaken by the user. Regardless, it's desirable that homenet protocols support rapid renumbering and that operational processes don't add unnecessary complexity for the renumbering process. Further, the introduction of any new homenet protocols should not make any form of renumbering any more complex than it already is.

Finally, the internal operation of the home network should also not depend on the availability of the ISP network at any given time, other than of course for connectivity to services or systems off the home network. This reinforces the use of ULAs for stable internal communication, and the need for a naming and service discovery mechanism that can operate independently within the homenet.

3.4.2. Stable internal IP addresses

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

ULAs should be used within the scope of a homenet to support stable

routing and connectivity between subnets and hosts regardless of whether a globally unique ISP-provided prefix is available. In the case of a prolonged external connectivity outage, ULAs allow internal operations across routed subnets to continue. ULA addresses also allow constrained devices to create permanent relationships between IPv6 addresses, e.g., from a wall controller to a lamp, where symbolic host names would require additional non-volatile memory and updating global prefixes in sleeping devices might also be problematic.

As discussed previously, it would be expected that ULAs would normally be used alongside one or more global prefixes in a homenet, such that hosts become multi-addressed with both globally unique and ULA prefixes. ULAs should be used for all devices, not just those intended to only have internal connectivity. Default address selection would then enable ULAs to be preferred for internal communications between devices that are using ULA prefixes generated within the same homenet.

In cases where ULA prefixes are in use within a homenet but there is no external IPv6 connectivity (and thus no GUAs in use), recommendations ULA-5, L-3 and L-4 in RFC 6204 should be followed to ensure correct operation, in particular where the homenet may be dual-stack with IPv4 external connectivity. The use of the Route Information Option described in [RFC4191] provides a mechanism to advertise such more-specific ULA routes.

The use of ULAs should be restricted to the homenet scope through filtering at the border(s) of the homenet, as mandated by RFC 6204 requirement S-2.

Note that it is possible that in some cases multiple /48 ULA prefixes may be in use within the same homenet, e.g., when the network is being deployed, perhaps also without external connectivity. In cases where multiple ULA /48's are in use, hosts need to know that each /48 is local to the homenet, e.g., by inclusion in their local address selection policy table.

3.4.3. Internal prefix delegation

As mentioned above, there are various sources of prefixes. From the homenet perspective, a single global prefix from each ISP should be received on the border CER [RFC3633]. Where multiple CERs exist with multiple ISP prefix pools, it is expected that routers within the homenet would assign themselves prefixes from each ISP they communicate with/through. As discussed above, a ULA prefix should be provisioned for stable internal communications or for use on constrained/LLN networks.

The delegation or availability of a prefix pool to the homenet should allow subsequent internal autonomous delegation of prefixes for use within the homenet. Such internal delegation should not assume a flat or hierarchical model, nor should it make an assumption about whether the delegation of internal prefixes is distributed or centralised. The assignment mechanism should provide reasonable efficiency, so that typical home network prefix allocation sizes can accommodate all the necessary /64 allocations in most cases, and not waste prefixes. Further, duplicate assignment of multiple /64s to the same network should be avoided, and the network should behave as gracefully as possible in the event of prefix exhaustion (though the options in such cases may be limited).

Where the home network has multiple CERs and these are delegated prefix pools from their attached ISPs, the internal prefix delegation would be expected to be served by each CER for each prefix associated with it. Where ULAs are used, it is preferable that only one /48 ULA covers the whole homenet, from which /64's can be delegated to the subnets. In cases where two /48 ULAs are generated within a homenet, the network should still continue to function, meaning that hosts will need to determine that each ULA is local to the homenet.

Delegation within the homenet should result in each link being assigned a stable prefix that is persistent across reboots, power outages and similar short-term outages. The availability of persistent prefixes should not depend on the router boot order. The addition of a new routing device should not affect existing persistent prefixes, but persistence may not be expected in the face of significant 'replumbing' of the homenet. However, delegated ULA prefixes within the homenet should remain persistent through an ISP-driven renumbering event.

Provisioning such persistent prefixes may imply the need for stable storage on routing devices, and also a method for a home user to 'reset' the stored prefix should a significant reconfiguration be required (though ideally the home user should not be involved at all).

This document makes no specific recommendation towards solutions, but notes that it is very likely that all routing devices participating in a homenet must use the same internal prefix delegation method. This implies that only one delegation method should be in use.

3.4.4. Coordination of configuration information

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

timers.

3.4.5. Privacy

If ISPs offer relatively stable IPv6 prefixes to customers, the network prefix part of addresses associated with the homenet may not change over a reasonably long period of time.

The exposure of which traffic is sourced from the same homenet is thus similar to IPv4; the single IPv4 global address seen through use of IPv4 NAT gives the same hint as the global IPv6 prefix seen for IPv6 traffic.

While IPv4 NAT may obfuscate to an external observer which internal devices traffic is sourced from, IPv6, even with use of Privacy Addresses [RFC4941], adds additional exposure of which traffic is sourced from the same internal device, through use of the same IPv6 source address for a period of time.

3.5. Routing functionality

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

The homenet unicast routing protocol should be based on a previously deployed protocol that has been shown to be reliable and robust, and that allows lightweight implementations. The availability of open source implementations is an important consideration. It is desirable, but not absolutely required, that the routing protocol be able to give a complete view of the network, and that it be able to pass around more than just routing information.

Multiple types of physical interfaces must be accounted for in the homenet routed topology. Technologies such as Ethernet, WiFi, Multimedia over Coax Alliance (MoCA), etc. must be capable of coexisting in the same environment and should be treated as part of any routed deployment. The inclusion of physical layer characteristics including bandwidth, loss, and latency in path computation should be considered for optimising communication in the homenet.

The routing protocol should support the generic use of multiple customer Internet connections, and the concurrent use of multiple delegated prefixes. A routing protocol that can make routing

decisions based on source and destination addresses is thus desirable, to avoid upstream ISP BCP38 ingress filtering problems. Multihoming support should also include load-balancing to multiple providers, and failover from a primary to a backup link when available. The protocol however should not require upstream ISP connectivity to be established to continue routing within the homenet.

The routing environment should be self-configuring, as discussed previously. An example of how OSPFv3 can be self-configuring in a homenet is described in [I-D.ietf-ospf-ospfv3-autoconfig]. Minimising convergence time should be a goal in any routed environment, but as a guideline a maximum convergence time at most 30 seconds should be the target (this target is somewhat arbitrary, and was chosen based on how long a typical home user might wait before attempting another reset; ideally the routers might have some status light indicating they are converging, similar to an ADSL router light indicating it is establishing a connection to its ISP).

As per prefix delegation, it is assumed that a single routing solution is in use in the homenet architecture. If there is an identified need to support multiple solutions, these must be interoperable.

An appropriate mechanism is required to discover which router(s) in the homenet are providing the CER function. Borders may include but are not limited to the interface to the upstream ISP, a gateway device to a separate home network such as a LLN network, or a gateway to a guest or private corporate extension network. In some cases there may be no border present, which may for example occur before an upstream connection has been established. The border discovery functionality may be integrated into the routing protocol itself, but may also be imported via a separate discovery mechanism.

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

3.5.1. Multicast support

It is desirable that, subject to the capacities of devices on certain media types, multicast routing is supported across the homenet. The natural scopes for internal multicast would be link-local or site-local, with the latter constrained within the homenet, but other policy borders, e.g., to a guest subnet, or to certain media types,

may also affect where specific multicast traffic is routed.

The homenet will need to be able to automatically configure the site multicast scope and scope boundary as part of the homenet edge (border) discovery process.

There may be different drivers for multicast to be supported across the homenet, e.g., for homenet-wide service discovery should a site-scope multicast service discovery protocol be defined, or potentially for novel streaming or filesharing applications. Where multicast is routed across a homenet an appropriate multicast routing protocol is required, one that as per the unicast routing protocol should be self-configuring. It must be possible to scope or filter multicast traffic to avoid it being flooded to network media where devices cannot reasonably support it.

Multicast may also be received by or sourced from the homenet from/to external networks, e.g., where video applications use multicast to conserve the bandwidth they consume. Such multicast traffic would be greater than site scope.

The multicast environment should support the ability for applications to pick a unique multicast group to use.

3.5.2. Mobility support

Devices may be mobile within the homenet. While resident on the same subnet, their address will remain persistent, but should devices move to a different (wireless) subnet, they will acquire a new address in that subnet. It is desirable that the homenet supports internal device mobility. To do so, the homenet may either extend the reach of specific wireless subnets to enable wireless roaming across the home (availability of a specific subnet across the home), or it may support mobility protocols to facilitate such roaming where multiple subnets are used.

3.6. Security

The security of an IPv6 homenet is an important consideration. The most notable difference to the IPv4 operational model is the removal of NAT, the introduction of global addressability of devices, and thus a need to consider whether devices should have global reachability. Regardless, hosts need to be able to operate securely, end-to-end where required, and also be robust against malicious traffic directed towards them. However, there are other challenges introduced, e.g., default filtering policies at the borders between various homenet realms.

3.6.1. Addressability vs reachability

An IPv6-based home network architecture should embrace the transparent end-to-end communications model as described in [RFC2775]. Each device should be globally addressable, and those addresses must not be altered in transit. However, security perimeters can be applied to restrict end-to-end communications, and thus while a host may be globally addressable it may not be globally reachable.

[RFC4864] describes a 'Simple Security' model for IPv6 networks, whereby stateful perimeter filtering can be applied to control the reachability of devices in a homenet. RFC 4864 states in Section 4.2 that "the use of firewalls ... is recommended for those that want boundary protection in addition to host defences". It should be noted that a 'default deny' filtering approach would effectively replace the need for IPv4 NAT traversal protocols with a need to use a signalling protocol to request a firewall hole be opened, e.g., a protocol such as PCP [RFC6887]. In networks with multiple CERs, the signalling would need to handle the cases of flows that may use one or more exit routers. CERs would need to be able to advertise their existence for such protocols.

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

3.6.2. Filtering at borders

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

There are cases where full internal connectivity may not be desirable, e.g., in certain utility networking scenarios, or where filtering is required for policy reasons against guest network subnet(s). Some scenarios/models may as a result involve running

isolated subnet(s) with their own CERs. In such cases connectivity would only be expected within each isolated network (though traffic may potentially pass between them via external providers).

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

3.6.3. Partial Effectiveness of NAT and Firewalls

Security by way of obscurity (address translation) or through firewalls (filtering) is at best only partially effective. The very poor security track record of home computer, home networking and business PC computers and networking is testimony to this. A security compromise behind the firewall of any device exposes all others, making an entire network that relies on obscurity or a firewall as vulnerable as the most insecure device on the private side of the network.

However, given current evidence of home network products with very poor default device security, putting a firewall in place does provide some level of protection. The use of firewalls today, whether a good practice or not, is common practice and whatever protection afforded, even if marginally effective, should not be lost. Thus, while it is highly desirable that all hosts in a homenet be adequately protected by built-in security functions, it should also be assumed that all CERs will continue to support appropriate perimeter defence functions, as per [I-D.ietf-v6ops-6204bis].

3.6.4. Exfiltration concerns

As homenets become more complex, with more devices, and with service discovery potentially enabled across the whole home, there are potential concerns over the leakage of information should devices use discovery protocols to gather information and report it to equipment vendors or application service providers.

While it is not clear how such exfiltration could be easily avoided, the threat should be recognised, be it from a new piece of hardware or some 'app' installed on personal device.

3.6.5. Device capabilities

In terms of the devices, homenet hosts should implement their own security policies in accordance to their computing capabilities. They should have the means to request transparent communications to be able to be initiated to them through security filters in the

homenet, either for all ports or for specific services. Users should have simple methods to associate devices to services that they wish to operate transparently through (CER) borders.

3.6.6. ULAs as a hint of connection origin

As noted in Section 3.6, if appropriate filtering is in place on the CER(s), as mandated by RFC 6204 requirement S-2, a ULA source address may be taken as an indication of locally sourced traffic. This indication could then be used with security settings to designate between which nodes a particular application is allowed to communicate, provided ULA address space is filtered appropriately at the boundary of the realm.

3.7. Naming and Service Discovery

The homenet requires devices to be able to determine and use unique names by which they can be accessed on the network. Users and devices will need to be able to discover devices and services available on the network, e.g., media servers, printers, displays or specific home automation devices. Thus naming and service discovery must be supported in the homenet, and, given the nature of typical home network users, the service(s) providing this function must as far as possible support unmanaged operation.

The naming system will be required to work internally or externally, be the user within the homenet or outside it, i.e., the user should be able to refer to devices by name, and potentially connect to them, wherever they may be. The most natural way to think about such naming and service discovery is to enable it to work across the entire homenet residence (site), disregarding technical borders such as subnets but respecting policy borders such as those between guest and other internal network realms. Remote access may be desired by the homenet residents while travelling, but also potentially by manufacturers or other 'benevolent' third parties.

3.7.1. Discovering services

Users will typically perform service discovery through graphical user interfaces (GUIs) that allow them to browse services on their network in an appropriate and intuitive way. Devices may also need to discover other devices, without any user intervention or choice. Either way, such interfaces are beyond the scope of this document, but the interface should have an appropriate application programming interface (API) for the discovery to be performed.

Such interfaces may also typically hide the local domain name element from users, especially where only one name space is available.

However, as we discuss below, in some cases the ability to discover available domains may be useful.

We note that current zero-configuration service discovery protocols are generally aimed at single subnets. There is thus a choice to make for multi-subnet homenet as to whether such protocols should be proxied or extended to operate across a whole homenet. In this context, that may mean bridging a link-local method, taking care to avoid loops, or extending the scope of multicast traffic used for the purpose. It may mean that some proxy or hybrid service is utilised, perhaps co-resident on the CER. Or it may be that a new approach is preferable, e.g., flooding information around the homenet as attributes within the routing protocol (which could allow per-prefix configuration). However, we should prefer approaches that are backwardly compatible, and allow current implementations to continue to be used. Note that this document does not mandate a particular solution, rather it expresses the principles that should be used for a homenet naming and service discovery environment.

One of the primary challenges facing service discovery today is lack of interoperability due to the ever increasing number of service discovery protocols available. While it is conceivable for consumer devices to support multiple discovery protocols, this is clearly not the most efficient use of network and computational resources. One goal of the homenet architecture should be a path to service discovery protocol interoperability either through a standards based translation scheme, hooks into current protocols to allow some form of communication among discovery protocols, extensions to support a central service repository in the homenet, or simply convergence towards a unified protocol suite.

3.7.2. Assigning names to devices

Given the large number of devices that may be networked in the future, devices should have a means to generate their own unique names within a homenet, and to detect clashes should they arise, e.g., where a second device of the same type/vendor as an existing device with the same default name is deployed, or where a new subnet is added to the homenet which already has a device of the same name. It is expected that a device should have a fixed name while within the scope of the homenet.

Users will also want simple ways to (re)name devices, again most likely through an appropriate and intuitive interface that is beyond the scope of this document. Note the name a user assigns to a device may be a label that is stored on the device as an attribute of the device, and may be distinct from the name used in a name service, e.g., 'Study Laser Printer' as opposed to printer2.<somedomain>.

3.7.3. The homenet name service

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

It is highly desirable that the homenet name service must at the very least co-exist with the Internet name service. There should also be a bias towards proven, existing solutions. The strong implication is thus that the homenet service is DNS-based, or DNS-compatible. There are naming protocols that are designed to be configured and operate Internet-wide, like unicast-based DNS, but also protocols that are designed for zero-configuration local environments, like mDNS [RFC6762].

When DNS is used as the homenet name service, it typically includes both a resolving service and an authoritative service. The authoritative service hosts the homenet related zone. One approach when provisioning such a name service, which is designed to facilitate name resolution from the global Internet, is to run an authoritative name service on the CER and a secondary authoritative name service provided by the ISP or perhaps an external third party.

Where zero configuration name services are used, it is desirable that these can also coexist with the Internet name service. In particular, where the homenet is using a global name space, it is desirable that devices have the ability, where desired, to add entries to that name space. There should also be a mechanism for such entries to be removed or expired from the global name space.

To protect against attacks such as cache poisoning, where an attacker is able to insert a bogus DNS entry in the local cache, it is desirable to support appropriate name service security methods, including DNS Security Extensions (DNSSEC) [RFC4033], on both the authoritative server and the resolver sides. Where DNS is used, the homenet router or naming service must not prevent DNSSEC from operating.

While this document does not specify hardware requirements, it is worth noting briefly here that e.g., in support of DNSSEC, appropriate homenet devices should have good random number generation capability, and future homenet specifications should indicate where high quality random number generators, i.e., with decent entropy, are needed.

Finally, the impact of a change in CER must be considered. It would be desirable to retain any relevant state (configuration) that was

held in the old CER. This might imply that state information should be distributed in the homenet, to be recoverable by/to the new CER, or to the homenet's ISP or a third party externally provided service by some means.

3.7.4. Name spaces

If access to homenet devices is required remotely from anywhere on the Internet, then at least one globally unique name space is required, though the use of multiple name spaces should not be precluded. One approach is that the name space(s) used for the homenet would be served authoritatively by the homenet, most likely by a server resident on the CER. Such name spaces may be acquired by the user or provided/generated by their ISP or an alternative externally provided service. It is likely that the default case is that a homenet will use a global domain provided by the ISP, but advanced users wishing to use a name space that is independent of their provider in the longer term should be able to acquire and use their own domain name. For users wanting to use their own independent domain names, such services are already available.

Devices may also be assigned different names in different name spaces, e.g., by third parties who may manage systems or devices in the homenet on behalf of the resident(s). Remote management of the homenet is out of scope of this document.

If however a global name space is not available, the homenet will need to pick and use a local name space which would only have meaning within the local homenet (i.e., it would not be used for remote access to the homenet). The .local name space currently has a special meaning for certain existing protocols which have link-local scope, and is thus not appropriate for multi-subnet home networks. A different name space is thus required for the homenet.

One approach for picking a local name space is to use an Ambiguous Local Qualified Domain Name (ALQDN) space, such as .sitelocal (or an appropriate name reserved for the purpose). While this is a simple approach, there is the potential in principle for devices that are bookmarked somehow by name by an application in one homenet to be confused with a device with the same name in another homenet. In practice however the underlying service discovery protocols should be capable of handling moving to a network where a new device is using the same name as a device used previously in another homenet.

An alternative approach for a local name space would be to use a Unique Locally Qualified Domain Name (ULQDN) space such as .<UniqueString>.sitelocal. The <UniqueString> could be generated in a variety of ways, one potentially being based on the local /48 ULA

prefix being used across the homenet. Such a <UniqueString> should survive a cold restart, i.e., be consistent after a network power-down, or, if a value is not set on startup, the CER or device running the name service should generate a default value. It would be desirable for the homenet user to be able to override the <UniqueString> with a value of their choice, but that would increase the likelihood of a name conflict. Any generated <UniqueString> should not be predictable; thus adding a salt/hash function would be desirable.

In the (likely) event that the homenet is accessible from outside the homenet (using the global name space), it is vital that the homenet name space follow the rules and conventions of the global name space. In this mode of operation, names in the homenet (including those automatically generated by devices) must be usable as labels in the global name space. [RFC5890] describes considerations for Internationalizing Domain Names in Applications (IDNA).

Also, with the introduction of new 'dotless' top level domains, there is also potential for ambiguity between, for example, a local host called 'computer' and (if it is registered) a .computer gTLD. Thus qualified names should always be used, whether these are exposed to the user or not. The IAB has issued a statement which explains why dotless domains should be considered harmful [IABdotless].

There may be use cases where either different name spaces may be desired for different realms in the homenet, or for segmentation of a single name space within the homenet. Thus hierarchical name space management is likely to be required. There should also be nothing to prevent individual device(s) being independently registered in external name spaces.

Where a user is in a remote network wishing to access devices in their home network, there may be a requirement to consider the domain search order presented where multiple associated name spaces exist. This also implies that a domain discovery function is desirable.

It may be the case that not all devices in the homenet are made available by name via an Internet name space, and that a 'split view' is preferred for certain devices, whereby devices inside the homenet see different DNS responses to those outside.

This document makes no assumption about the presence or omission of a reverse lookup service. There is an argument that it may be useful for presenting logging information to users with meaningful device names rather than literal addresses. There are also some services, most notably email mail exchangers, where some operators have chosen to require a valid reverse lookup before accepting connections.

3.7.5. Independent operation

Name resolution and service discovery for reachable devices must continue to function if the local network is disconnected from the global Internet, e.g., a local media server should still be available even if the Internet link is down for an extended period. This implies the local network should also be able to perform a complete restart in the absence of external connectivity, and have local naming and service discovery operate correctly.

The approach described above of a local authoritative name service with a cache would allow local operation for sustained ISP outages.

Having an independent local trust anchor is desirable, to support secure exchanges should external connectivity be unavailable.

A change in ISP should not affect local naming and service discovery. However, if the homenet uses a global name space provided by the ISP, then this will obviously have an impact if the user changes their network provider.

3.7.6. Considerations for LLNs

In some parts of the homenet, in particular LLNs or any devices where battery power is used, devices may be sleeping, in which case a proxy for such nodes may be required, that could respond (for example) to multicast service discovery requests. Those same devices or parts of the network may have less capacity for multicast traffic that may be flooded from other parts of the network. In general, message utilisation should be efficient considering the network technologies and constrained devices that the service may need to operate over.

There are efforts underway to determine naming and discovery solutions for use by the Constrained Application Protocol (CoAP) [I-D.ietf-core-coap] in LLN networks. These are outside the scope of this document.

3.7.7. DNS resolver discovery

Automatic discovery of a name service to allow client devices in the homenet to resolve external domains on the Internet is required, and such discovery must support clients that may be a number of router hops away from the name service. Similarly it may be desirable to convey any DNS domain search list that may be in effect for the homenet.

3.7.8. Devices roaming to/from the homenet

It is likely that some devices which have registered names within the homenet Internet name space and that are mobile will attach to the Internet at other locations and acquire an IP address at those locations. Devices may move between different homenets. In such cases it is desirable that devices may be accessed by the same name as is used in their home network.

Solutions to this problem are not discussed in this document. They may include use of Mobile IPv6 or Dynamic DNS, either of which would put additional requirements on to the homenet, or establishment of a (VPN) tunnel to a server in the home network.

3.8. Other Considerations

This section discusses two other considerations for home networking that the architecture should not preclude, but that this text is neutral towards.

3.8.1. Quality of Service

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

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

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

3.8.2. Operations and Management

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

discussed in this document.

There may be some configuration parameters which are exposed to users, e.g., SSID name(s), or wireless security key(s). Users may also be expected to be aware of the functions of certain devices they connect, e.g., which are providing a server function, though service discovery protocols should make their selection as intuitive as possible.

As discussed in Section 3.6.1 the default setting on the homenet-ISP border for inbound traffic may be default deny, default allow, or some position inbetween. Whatever the default position, it should be possible for the user to change the setting.

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

It is expected that network management functions would be available over IPv6 transport, even where the homenet is dual-stack.

3.9. Implementing the Architecture on IPv6

This architecture text encourages re-use of existing protocols. Thus the necessary mechanisms are largely already part of the IPv6 protocol set and common implementations, though there are some exceptions.

For automatic routing, it is expected that solutions can be found based on existing protocols. Some relatively smaller updates are likely to be required, e.g., a new mechanism may be needed in order to turn a selected protocol on by default, a mechanism may be required to automatically assign prefixes to links within the homenet.

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

Some protocol changes are however required in the architecture, e.g., for name resolution and service discovery, extensions to existing zero configuration link-local name resolution protocols are needed to

enable them to work across subnets, within the scope of the home network site.

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

4. Conclusions

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

5. Security Considerations

Security considerations for the homenet architecture are discussed in Section 3.6 above.

6. IANA Considerations

This document has no actions for IANA.

7. References

7.1. Normative References

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [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.

7.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, February 2000.
- [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.
- [RFC3002] Mitzel, D., "Overview of 2000 IAB Wireless Internetworking Workshop", RFC 3002, December 2000.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.

- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, August 2010.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.
- [RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.
- [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, March 2011.
- [RFC6204] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, April 2011.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, April 2012.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, February 2013.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P.

Selkirk, "Port Control Protocol (PCP)", RFC 6887,
April 2013.

[I-D.ietf-v6ops-ipv6-multihoming-without-ipv6nat]
Troan, O., Miles, D., Matsushima, S., Okimoto, T., and D.
Wing, "IPv6 Multihoming without Network Address
Translation",
draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat-05 (work
in progress), March 2013.

[I-D.ietf-ospf-ospfv3-autoconfig]
Lindem, A. and J. Arkko, "OSPFv3 Auto-Configuration",
draft-ietf-ospf-ospfv3-autoconfig-05 (work in progress),
October 2013.

[I-D.ietf-core-coap]
Shelby, Z., Hartke, K., and C. Bormann, "Constrained
Application Protocol (CoAP)", draft-ietf-core-coap-18
(work in progress), June 2013.

[I-D.ietf-v6ops-6204bis]
Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic
Requirements for IPv6 Customer Edge Routers",
draft-ietf-v6ops-6204bis-12 (work in progress),
October 2012.

[IABdotless]
"IAB Statement: Dotless Domains Considered Harmful",
February 2013, <[http://www.iab.org/documents/
correspondence-reports-documents/2013-2/
iab-statement-dotless-domains-considered-harmful](http://www.iab.org/documents/correspondence-reports-documents/2013-2/iab-statement-dotless-domains-considered-harmful)>.

[Gettys11]
Gettys, J., "Bufferbloat: Dark Buffers in the Internet",
March 2011,
<<http://www.ietf.org/proceedings/80/slides/tsvarea-1.pdf>>.

Appendix A. Acknowledgments

The authors would like to thank Aamer Akhter, Mikael Abrahamsson, Mark Andrews, Dmitry Anipko, Ran Atkinson, Fred Baker, Ray Bellis, Teco Boot, John Brzozowski, Cameron Byrne, Brian Carpenter, Stuart Cheshire, Julius Chroboczek, Lorenzo Colitti, Robert Cragie, Elwyn Davies, Ralph Droms, Lars Eggert, Jim Gettys, Olafur Gudmundsson, Wassim Haddad, Joel M. Halpern, David Harrington, Lee Howard, Ray Hunter, Joel Jaeggli, Heather Kirksey, Ted Lemon, Acee Lindem, Kerry Lynn, Daniel Migault, Erik Nordmark, Michael Richardson, Mattia

Rossi, Barbara Stark, Markus Stenberg, Sander Steffann, Don Sturek, Andrew Sullivan, Dave Taht, Dave Thaler, Michael Thomas, Mark Townsley, JP Vasseur, Curtis Villamizar, Dan Wing, Russ White, and James Woodyatt for their comments and contributions within homenet WG meetings and on the WG mailing list. An acknowledgement generally means that person's text made it in to the document, or was helpful in clarifying or reinforcing an aspect of the document. It does not imply that each contributor agrees with every point in the document.

Appendix B. Changes

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

B.1. Version 11 (after IESG review)

Changes made include:

- o Jouni Korhonen's OPSDIR review comments addressed.
- o Elwyn Davies' gen-art review comments addressed.

B.2. Version 10 (after AD review)

Changes made include:

- o Minor changes/clarifications resulting from AD review

B.3. Version 09 (after WGLC)

Changes made include:

- o Added note about multicast into or out of site
- o Removed further personal draft references, replaced with covering text
- o Routing functionality text updated to avoid ambiguity
- o Added note that devices away from homenet may tunnel home (via VPN)
- o Added note that homenets more exposed to provider renumbering than with IPv4 and NAT
- o Added note about devices that may be ULA-only until configured to be globally addressable

- o Removed paragraph about broken CERS that do not work with prefixes other than /64
- o Noted no recommendation on methods to convey prefix information is made in this text
- o Stated that this text does not recommend how to form largest possible subnets
- o Added text about homenet evolution and handling disparate media types
- o Rephrased NAT/firewall text on marginal effectiveness
- o Emphasised that multihoming may be to any number of ISPs

B.4. Version 08

Changes made include:

- o Various clarifications made in response to list comments
- o Added note on ULAs with IPv4, where no GUAs in use
- o Added note on naming and internationalisation (IDNA)
- o Added note on trust relationships when adding devices
- o Added note for MPTCP
- o Added various naming and SD notes
- o Added various notes on delegated ISP prefixes

B.5. Version 07

Changes made include:

- o Removed reference to NPTv6 in section 3.2.4. Instead now say it has an architectural cost to use in the earlier section, and thus it is not recommended for use in the homenet architecture.
- o Removed 'proxy or extend?' section. Included shorter text in main body, without mandating either approach for service discovery.
- o Made it clearer that ULAs are expected to be used alongside globals.

- o Removed reference to 'advanced security' as described in draft-vyncke-advanced-ipv6-security.
- o Balanced the text between ULQDN and ALQDN.
- o Clarify text does not assume default deny or allow on CER, but that either mode may be enabled.
- o Removed ULA-C reference for 'simple' addresses. Instead only suggested service discovery to find such devices.
- o Reiterated that single/multiple CER models to be supported for multihoming.
- o Reordered section 3.3 to improve flow.
- o Added recommendation that homenet is not allocated less than /60, and a /56 is preferable.
- o Tidied up first few intro sections.
- o Other minor edits from list feedback.

B.6. Version 06

Changes made include:

- o Stated that unmanaged goal is 'as far as possible'.
- o Added note about multiple /48 ULAs potentially being in use.
- o Minor edits from list feedback.

B.7. Version 05

Changes made include:

- o Some significant changes to naming and SD section.
- o Removed some expired drafts.
- o Added notes about issues caused by ISP only delegating a /64.
- o Recommended against using prefixes longer than /64.
- o Suggested CER asks for /48 by DHCP PD, even if it only receives less.

- o Added note about DS-Lite but emphasised transition is out of scope.
- o Added text about multicast routing.

B.8. Version 04

Changes made include:

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

B.9. Version 03

Changes made include:

- o Various improvements to the readability.
- o Removed bullet lists of requirements, as requested by chair.
- o Noted 6204bis has replaced advanced-cpe draft.
- o Clarified the topology examples are just that.
- o Emphasised we are not targetting walled gardens, but they should not be precluded.
- o Also changed text about requiring support for walled gardens.
- o Noted that avoiding falling foul of ingress filtering when multihomed is desirable.
- o Improved text about realms, detecting borders and policies at borders.
- o Stated this text makes no recommendation about default security model.
- o Added some text about failure modes for users plugging things arbitrarily.
- o Expanded naming and service discovery text.

- o Added more text about ULAs.
- o Removed reference to version 1 on chair feedback.
- o Stated that NPTv6 adds architectural cost but is not a homenet matter if deployed at the CER. This text only considers the internal homenet.
- o Noted multihoming is supported.
- o Noted routers may not be separate devices, they may be embedded in devices.
- o Clarified simple and advanced security some more, and RFC 4864 and 6092.
- o Stated that there should be just one secret key, if any are used at all.
- o For multihoming, support multiple CERs but note that routing to the correct CER to avoid ISP filtering may not be optimal within the homenet.
- o Added some ISPs renumber due to privacy laws.
- o Removed extra repeated references to Simple Security.
- o Removed some solution creep on RIOS/RAs.
- o Load-balancing scenario added as to be supported.

B.10. Version 02

Changes made include:

- o Made the IPv6 implications section briefer.
- o Changed Network Models section to describe properties of the homenet with illustrative examples, rather than implying the number of models was fixed to the six shown in 01.
- o Text to state multihoming support focused on single CER model. Multiple CER support is desirable, but not required.
- o Stated that NPTv6 not supported.
- o Added considerations section for operations and management.

- o Added bullet point principles/requirements to Section 3.4.
- o Changed IPv6 solutions must not adversely affect IPv4 to should not.
- o End-to-end section expanded to talk about "Simple Security" and borders.
- o Extended text on naming and service discovery.
- o Added reference to RFC 2775, RFC 6177.
- o Added reference to the new xmDNS draft.
- o Added naming/SD requirements from Ralph Droms.

Authors' Addresses

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

Email: tjc@ecs.soton.ac.uk

Jari Arkko
Ericsson
Jorvas 02420
Finland

Email: jari.arkko@piuha.net

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

Email: abr@sdesigns.dk

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

Email: ot@cisco.com

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

Email: jason.weil@twcable.com

HOMENET
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

D. Migault (Ed)
Orange
October 21, 2013

DNSSEC Validators DHCP Options
draft-mglt-homenet-dnssec-validator-dhc-options-02.txt

Abstract

DNSSEC provides data integrity and authentication for DNSSEC validators. However, without valid trust anchor(s) and an acceptable value for the current time, DNSSEC validation cannot be performed. As a result, there are multiple cases where DNSSEC validation MUST NOT be performed. In addition, this list of exceptions is expected to become larger over time.

Considering an increasing number of cases where DNSSEC is disabled adds complexity to the DNSSEC validator implementations and increases the vectors that disable security.

This document assumes that DNSSEC adoption by end devices requires that end devices MUST be able to support a DNSSEC validation always set. This MUST be valid today as well as in the future.

This document describes DHCP Options to provision the DHCP Client with valid trust anchors and time so DNSSEC validation can be performed.

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 24, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Requirements notation	2
2. Introduction	2
3. Threat Model	3
3.1. Motivations for providing DNSSEC Trust Anchor	3
3.2. Motivations for providing Time	5
4. Terminology	5
5. DHCP DNSSEC Trust Anchor Options	6
5.1. DHCP DNSSEC KSK RR Trust Anchor Options	6
5.2. DHCP DNSSEC KSK CERT Trust Anchor Options	6
6. DHCP Time Option	7
7. DHCP Client Behavior	8
8. DHCP Server Behavior	9
9. DHCP Relay Agent Behavior	10
10. IANA Considerations	10
11. Security Considerations	10
12. Acknowledgment	10
13. References	10
13.1. Normative References	10
13.2. Informational References	11
Appendix A. Document Change Log	12
Author's Address	12

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

DNSSEC [RFC4033], [RFC4034], [RFC4035] adds data authentication and integrity checks to DNS [RFC1034], [RFC1035]. For signature validation, DNSSEC requires a trust anchor such as the Key Signing

Key (KSK) of the Root Zone or any other zone. Without a trust anchor, DNSSEC validation cannot be performed. In addition KSKs and signatures are valid for a given period of time. As a result, DNSSEC validation cannot be performed if time shifting is too large.

This document considers DHCP DNSSEC Trust Anchor Option and DHCP Time Option to provision a device with trusted KSKs and current time. Although our priority is to provide the Root Zone KSK, we also consider the case other trusted KSK MAY be provided, for example, if a Zone does not provide secure delegation, or to mitigate badly configured DNSSEC zones (like TLDs zones).

The main motivation for these DHCP Options is that DHCP enabled devices have DNSSEC validation always set and do not need to perform DNS resolution without DNSSEC validation. In fact, enabling DNS with no validation represents a potential way to remove security and MAY be used by attackers. Similarly, DNSSEC configuration implemented in the end users device, MAY not consider future cases and MAY introduce vulnerabilities. DHCP Options prevent this as long as the relationship between DHCP Client and DHCP Server is trusted.

This document assumes that the channel between the DHCP Client and the DHCP Server is trusted and secured with DHCP mechanisms described in [RFC3315], or IPsec [RFC4301].

3. Threat Model

This document addresses the case of a device configured with DNSSEC validation set that is plugged in, gets connectivity (using DHCP for example), but fails DNSSEC resolutions because its trust anchor KSK is not valid anymore or its local time is not valid.

This threat mainly addresses devices that can be switched off for a long period of time or devices that MAY be off-shelves for a long time before being plugged in. CPEs as well as any homenet devices are concerned by this use case.

This threat also addresses DNSSEC emergency key roll over operations. Devices that have cached the out-of-date KSK will not be able to check the signatures until the TTL has expired on all caches.

This document proposes DHCP Options that provide the necessary parameters to perform DNSSEC validation. These Options MUST be used on a trusted network over a trusted channel between the DHCP Client and the DHCP Server. These options MAY be used in conjunction of additional mechanisms.

3.1. Motivations for providing DNSSEC Trust Anchor

The first motivation for providing trusted KSKs is to provide automatic configuration of devices to enable DNSSEC validation. This avoids validator initial KSK provisioning issue as well as KSK roll over issues.

A validator MAY not be able to perform signature check with an authenticated KSK because:

- 1) It does not have a trust anchor (like the Root Zone KSK)
- 2) The KSK MAY have been authenticated, stored or cached with an expiration date valid but is not valid anymore. This MAY happen in the case of an emergency key roll over, if the device has been offline during the key roll over, or if the key roll over is not performed as described in [DPS-KSK], [RFC5011].
- 3) The chain of trust MAY have been broken. This can happen to non Root Zone KSK only and MAY not involve the responsibility of the owner of the zone. The deeper the Zone is in the hierarchy, the more likely this happens.
- 4) A DNSSEC zone MAY have been badly signed or a KSK MAY have been badly generated. The DNSSEC MAY be correct, but DNSSEC validator MAY keep for a long time the badly generated KSK, ZSK...

The goal of the DHCP DNSSEC Trust Anchor Option is to provide these validators trusted anchors like the Root Zone KSK, as well as other KSKs (TLDs...) so the validator has the proper KSKs to perform DNSSEC validation.

Most documents are currently focused on the Root Zone KSK for which recommendations and alternative mechanisms have been described. [I-D.jabley-dnsop-validator-bootstrap] provides guide lines on how to retrieve and select DNSSEC Trust Anchors. Section 5.3 and [I-D.jabley-dnssec-trust-anchor] describes mechanisms to retrieve securely the Root Zone KSK relying on TLS security. It suggests to use insecure DNS resolution to set HTTPS connections. Using HTTPS requires downloading the keyDigest id (key-label) from <https://data.iana.org/root-anchors/root-anchors.xml>, followed by an HTTPS request at <https://data.iana.org/root-anchors/key-label.crt> to get the whole certificate.

The key advantages of the DHCP DNSSEC Trust Anchor Option described in this document are that we extend the mechanism to any KSK, and validators can set DNSSEC validation for all DNS queries. However, we do not see any contradiction between recommendations provided by [I-D.jabley-dnsop-validator-bootstrap] and [I-D.jabley-dnssec-trust-

anchor] and believe the principle described in these documents SHOULD be applied by the validators. Note also that DHCP DNSSEC Trust Anchor Option only benefits to validators that are configured via DHCP.

To recover from a DNSSEC failure and remove a particular data from cache, [I-D.jabley-dnsop-dns-flush] suggests to use a NOTIFY message between Authoritative Servers and Resolvers. This mechanism is set between Recursive Server and Authoritative Servers with a specific trusted relationship. This is probably a selection of TLDs. This document, does not address the DNSSEC failure over Recursive Servers, but addresses more specifically DHCP configured devices. These are typically CPEs or End Users. We believe that configuring and restarting DNSSEC validators with DHCP Option, is an easier way to cope with this issue. First the trust relation between DHCP Server already exists, we do not need additional trusted channel between Authoritative Servers or eventually the Recursive Servers. Then basic implementations of stub resolvers, in CPE or desktops may not address NOTIFY message.

3.2. Motivations for providing Time

KSKs and signatures are always associated to an expiration time. As a result, DNSSEC validation requires that the validator knows the current time.

A number of mechanisms exists like [TSLDATE] or [RFC5905] for setting the time of the device. In addition, [RFC5908] provides a Network Time Protocol (NTP) Server Option for DHCP. The DHCP Time Option described in this document differs from [RFC5908] as it provides an estimation of the current time, instead of providing the NTP servers location information. The time value provided by the DHCP Time Option should be used only if previously mentioned mechanisms are either not implemented on the device or are unavailable. One of the reason MAY be that you MAY need valid DNS(SEC) resolution to use these protocols. The time provided by the DHCP Time Option does not have the accuracy of NTP and SHOULD be considered as a best effort value. [I-D.jabley-dnsop-validator-bootstrap] also recommends that when time has not been verified by the validator, the signature validation SHOULD be done with time off.

The key advantage of the DHCP Time Option is that it makes possible to have DNSSEC validation always set. It limits the possible DNSSEC validation variants which potentially expose the device to disable DNSSEC validations. Note also that DHCP Time Option only benefits to validators that are configured via DHCP.

4. Terminology

5. DHCP DNSSEC Trust Anchor Options

This section describes two options:

- DHCP DNSSEC KSK Trust Anchor Options: carries the KSK RRset as described in [RFC1035] with a DNSKEY RDATA as described in [RFC4033]. This data is not integrity protected, nor it can be authenticated. Such data SHOULD be trusted over a trusted DHCP channel.
- DHCP DNSSEC CERT Trust Anchor Options: Carries a certificate encoded as described in [RFC4398]. The advantage of the Certificate is that it enables authentication of the received information by a trusted party. For example, CPE providers MAY provide a trusted certification authority. Unlike DNSSEC key roll over, the CPE provider controls the key roll over of the certification authority it provides.

5.1. DHCP DNSSEC KSK RR Trust Anchor Options

The DHCP DNSSEC KSK Trust Anchor Option provides the RRset as mentioned in the DNS(SEC) Zone. In other words, it carries the RR as defined in Section 3.2. of [RFC1035] and a RDATA DNSKEY as defined in Section 2.1 of [RFC4033]. As the RR has a variable length, the DHCP DNSSEC KSK Trust Anchor Options follows the recommendation format of Section 5.9 of [I-D.ietf-dhc-option-guidelines].

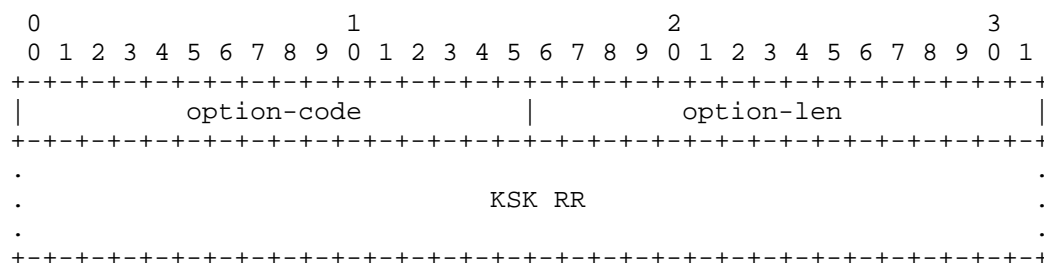


Figure 1: DHCP DNSSEC KSK Trust Anchor Options
Payload Description

- option-code: OPTION_DNSSEC_KSK_RR_TRUST_ANCHOR
- option-len: An unsigned integer giving the length of the KSK RR field in this option in octets

5.2. DHCP DNSSEC KSK CERT Trust Anchor Options

The DHCP DNSSEC CERT Trust Anchor Option provides a certificate. The CERT RR is described in [RFC4398]. Note that only the RDATA associated to the CERT is present in the DHCP Option. As the RR has a variable length, the DHCP DNSSEC KSK CERT Trust Anchor Options follows the recommendation format of Section 5.9 of [I-D.ietf-dhc-option-guidelines].

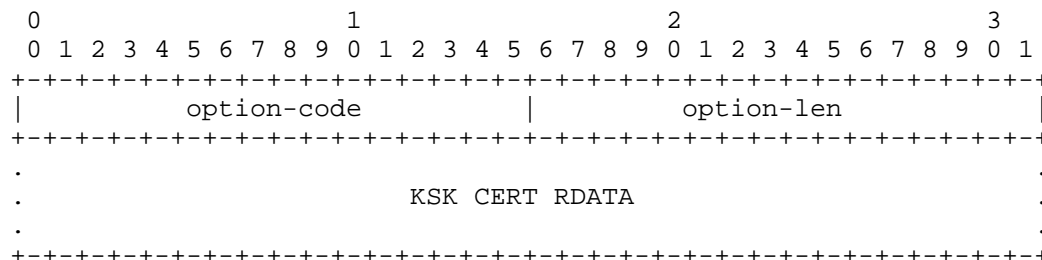


Figure 2: DHCP DNSSEC CERT Trust Anchor Options
Payload Description

- option-code: OPTION_DNSSEC_CERT_TRUST_ANCHOR
- option-len: An unsigned integer giving the length of the KSK RR field in this option in octets

The X.509 [RFC5280] certificate MUST have a keyUsage set to digitalSignature (0) and nonRepudiation (1). Subject Alternative Name DNS name indicates the name of the zone.

In order to be compliant with the certificate of the Root Zone described [I-D.jabley-dnssec-trust-anchor]. The CERT for a KSK SHOULD have a Common Name (CN) with the string "'Zone-FQDN' Zone KSK" followed by the time and date of key generation in the format specified in [RFC3339]. 'Zone-FQDN' is the name of the zone and SHOULD be the same as the one mentioned in Subject Alternative Name. The resourceRecord Attribute SHOULD be set with the DS RRset.

6. DHCP Time Option

The DHCP DNSSEC Time Option is used by the DHCP Server to indicate the Time to the DHCP Client. The Time is provided in a string format as specified in [RFC3339] and in [I-D.ietf-dhc-option-guidelines] Section 5.8.

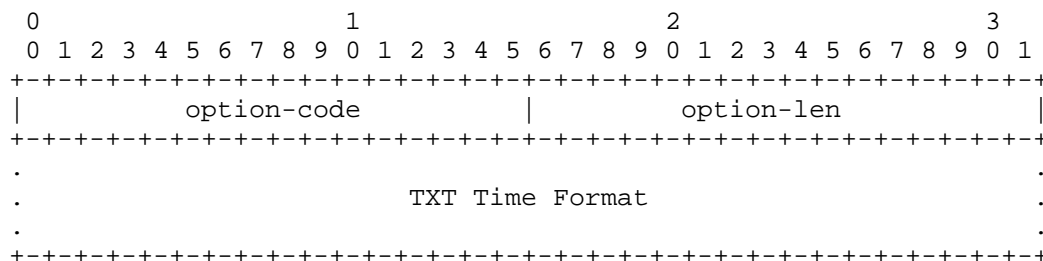


Figure 2: DHCP Time Options
Payload Description

- option-code: OPTION_TIME
- option-len: A string representing the Time

7. DHCP Client Behavior

DHCP DNSSEC KSK Trust Anchor Option, DHCP DNSSEC CERT Trust Anchor Option or DHCP Time Option described in this document are intended for DNSSEC validation. If a connected device is not performing DNSSEC validation, it MUST NOT send a DHCP an Option Request DHCP Option (ORO) [RFC3315] for any of these options, and MUST ignore all these options if provided by the DHCP Server.

The DHCP sends a DHCP ORO for one or multiple options described in the document. Motivations for sending this Option Request DHCP Option is out of scope of the document. It could be a device switched off for a long time, a device that cannot validate the DNSSEC responses.

A channel is considered trusted if 1) the DHCP Server is trusted and authenticated and 2) exchanged data between the DHCP Client and the DHCP Server is integrity protected. IPsec [RFC4301], for example, MAY be used to establish a secure channel.

Over a trusted channel, the DHCP Client that performs DNSSEC validation MAY send an ORO for any of the DHCP DNSSEC KSK Trust Anchor Option, the DHCP DNSSEC CERT Trust Anchor Option or the DHCP Time Option to a DHCP Server.

Over a trusted channel, the DHCP Client that performs DNSSEC validation SHOULD consider the DHCP DNSSEC KSK Trust Anchor Option, the DHCP DNSSEC CERT Trust Anchor Option or the DHCP Time Option sent by the DHCP Server.

Over a non trusted channel, the DHCP Client MAY only send ORO for a DHCP DNSSEC CERT Trust Anchor Option. This option is the only one that MAY be considered by the DHCP Client if sent by the DHCP Server. If the DHCP Client does not trust the signer of the certificate, the option MUST be ignored.

When a DHCP DNSSEC KSK Trust Anchor Option or a DHCP DNSSEC CERT Trust Anchor Option is accepted by the DHCP Client, it MUST remove overwrite old values for the KSK with the new one.

When a DHCP Time Option is accepted by the DHCP Client, it MUST check the difference between its clock and the time provided by the Option. It SHOULD overwrite its clock value only if the difference is too large.

In any other case, ORO requests MUST NOT be sent by the DHCP Client, and options received by the DHCP Server MUST NOT be considered by the DHCP Client. The remaining of the section details when the options MUST NOT be requested by the DHCP Client and MUST be ignored by the DHCP Client when received by the DHCP Server.

The DHCP Client MUST NOT send an ORO for a DHCP DNSSEC KSK Trust Anchor Option, a DHCP DNSSEC CERT Trust Anchor Option or a DHCP Time Option to a DHCP Server that is either not trusted or not authenticated.

All DHCP DNSSEC KSK Trust Anchor Option, a DHCP DNSSEC CERT Trust Anchor Option or a DHCP Time Option received from DHCP Server that is not authenticated or that is not trusted MUST be ignored by the DHCP Client.

The DHCP Client MUST NOT send an ORO for a DHCP DNSSEC KSK Trust Anchor Option or a DHCP Time Option to a trusted DHCP Server over an untrusted channel. A DHCP DNSSEC CERT Trust Anchor Option MAY be requested over an untrusted channel since the certificate is signed and thus can be authenticated. A DHCP DNSSEC CERT Trust Anchor Option signed by an untrusted authority MUST be ignored by the DHCP Client.

All DHCP DNSSEC KSK Trust Anchor Option or a DHCP Time Option received from DHCP Server over a channel that is not trusted MUST be ignored by the DHCP Client.

8. DHCP Server Behavior

The DHCP Server SHOULD properly answer with the requested options in the ORO, even if the DHCP Server does not consider the channel with DHCP Client as trusted.

The DHCP Server MAY also provide DHCP DNSSEC KSK Trust Anchor Option, DHCP DNSSEC CERT Trust Anchor Option or DHCP Time Option without being requested by the DHCP Client. This could for example prevent failures not detected by the DHCP Client.

9. DHCP Relay Agent Behavior

The DHCP Options described in the document do not impact the Relay Agent.

10. IANA Considerations

The DHCP options detailed in this document is:

- OPTION_DNSSEC_KSK_RR_TRUST_ANCHOR: TBD
- OPTION_DNSSEC_KSK_CERT_TRUST_ANCHOR: TBD
- OPTION_TIME: TBD

11. Security Considerations

Security has been discussed in the "DHCP Client Behavior Section". As information contained in the payloads are use to enable signature validation, these pieces of information MUST be considered only when issued by a trusted party, and when integrity protection is provided.

12. Acknowledgment

Bringing DNSSEC in Home Networks discussion has started during the IETF87 in Berlin with Ted Lemon, Ralph Weber, Normen Kowalewski, and Mikael Abrahamsson. An email discussion has also been initiated by Jim Gettys with among others, helpful remarks from Paul Wouters, Joe Abley, Michael Ridchardson.

13. References

13.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3339] Klyne, G., Ed. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, July 2002.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4398] Josefsson, S., "Storing Certificates in the Domain Name System (DNS)", RFC 4398, March 2006.
- [RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", STD 74, RFC 5011, September 2007.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.
- [RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.
- [RFC5908] Gayraud, R. and B. Lourdelet, "Network Time Protocol (NTP) Server Option for DHCPv6", RFC 5908, June 2010.

13.2. Informational References

- [DPS-KSK] Ljunggren, F., Okubo, T., Lamb, R., and J. Schlyter, "DNSSEC Practice Statement for the Root Zone KSK Operation", Root DNSSEC Design Team, URL: <http://www.root-dnssec.org/wp-content/uploads/2010/06/icann-dps-00.txt>, 2010.

[I-D.ietf-dhc-option-guidelines]

Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", draft-ietf-dhc-option-guidelines-14 (work in progress), September 2013.

[I-D.jabley-dnsop-dns-flush]

Abley, J., "A Mechanism for Remote-Triggered DNS Cache Flushes (DNS FLUSH)", draft-jabley-dnsop-dns-flush-00 (work in progress), June 2013.

[I-D.jabley-dnsop-validator-bootstrap]

Abley, J. and D. Knight, "Establishing an Appropriate Root Zone DNSSEC Trust Anchor at Startup", draft-jabley-dnsop-validator-bootstrap-00 (work in progress), January 2011.

[I-D.jabley-dnssec-trust-anchor]

Abley, J., Schlyter, J., and G. Bailey, "DNSSEC Trust Anchor Publication for the Root Zone", draft-jabley-dnssec-trust-anchor-07 (work in progress), June 2013.

[TSLDATE] error, IO., "tlsdate: secure parasitic rdate replacement", URL: <https://github.com/ioerror/tlsdate>, 2013.

Appendix A. Document Change Log

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

-00: First version published.

Author's Address

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

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

HOMENET
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2014

D. Migault (Ed)
Orange
W. Cloetens
SoftAtHome
C. Griffiths
Dyn
R. Weber
Nominum
October 20, 2013

IPv6 Home Network Naming Delegation
draft-mglt-homenet-front-end-naming-delegation-03.txt

Abstract

CPEs are designed to provide IP connectivity to home networks. Most CPEs assigns IP addresses to the nodes of the home network which makes it a good candidate for hosting the naming service. With IPv6, the naming service makes nodes reachable from the home network as well as from the Internet.

However, CPEs have not been designed to host such a naming service exposed on the Internet. This MAY expose the CPEs to resource exhaustion which would make the home network unreachable, and most probably would also affect the home network inner communications.

In addition, DNSSEC management and configuration may not be well understood or mastered by regular end users. Misconfiguration MAY also results in naming service disruption, thus these end users MAY prefer to rely on third party naming providers.

This document describes a homenet naming architecture where the CPEs manage the DNS zone associates to its home network, and outsource both DNSSEC management and naming service on the Internet to a third party designated as the Public Authoritative Servers.

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 23, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Requirements notation	2
2. Introduction	3
3. Terminology	4
4. Architecture Overview	5
5. Architecture Description	7
5.1. CPE and Public Authoritative Servers Synchronization . .	8
5.1.1. Synchronization with a Hidden Master	8
5.1.2. Securing Synchronization	9
5.2. DNS Homenet Zone configuration	10
5.3. DNSSEC outsourcing configuration	12
5.4. CPE Security Policies	13
6. Homenet Naming Configuration	13
7. Security Considerations	14
7.1. Names are less secure than IP addresses	14
7.2. Names are less volatile than IP addresses	15
8. IANA Considerations	15
9. Acknowledgment	15
10. References	16
10.1. Normative References	16
10.2. Informational References	17
Appendix A. Document Change Log	17
Authors' Addresses	18

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 end to end IP reachability from the Internet and into the Home Network. End Users to access services hosted in the Home Network with IPv6 addresses would prefer to use names instead of long and complex IPv6 addresses.

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. As a result, hosting the a DNS service on the Internet MAY expose the Home Network to resource exhaustion, which may isolate the Home Network from the Internet and affect the services hosted by the CPEs, thus affecting the overall Home Network communications. So, this document considers that the Naming Service SHOULD NOT be hosted on the CPE and SHOULD be outsourced to a third party.

In addition, the Naming Service of the Home Network is expected to be deployed with its security extension DNSSEC. DNSSEC comes with complex configurations as well as complex operation management like (DNSSEC secure delegation, DNSSEC key roll over, DNSSEC zone updates). These operations can hardly be understood by the average end user, and a misconfiguration MAY result in invalid naming resolutions that MAY make an host, or the whole home network unreachable. So, this document considers DNSSEC management operations SHOULD NOT be handled by the average end user, but SHOULD be outsourced to a third party.

This document describes an architecture where the CPE outsources the authoritative naming service and DNSSEC zone management to a third party designated as Public Authoritative Servers. It describes interactions between the CPE and the Public Authoritative Servers, that is to say the involved protocols and their respective configurations. More specifically, this document does not describe any new protocol. It provides a guide line to properly use the already existing protocols.

This document intends to efficiently deploy DNSSEC in the Home Networks in a standardized and highly flexible way. More

specifically, the described Home Network Naming architecture is expected to lead to autoconfiguration facilities for most common users, as well as enabling advanced users to have their own specific settings. In fact, some end users MAY choose to host and expose a Naming service on their CPE. Others MAY sign the zone on the CPE. Although the document does not describe these scenarios, the described architecture only requires minor modifications - such as allowing incoming DNS queries from the Internet and adding the CPE in the list of Naming servers.

The document is organized as follows. Section 4 provides an overview of the homenet naming architecture and presents the CPE and the Public Authoritative Server that handles the authoritative naming service of the home network as well as DNSSEC management operations on behalf of the CPE. Section 5 describes in details protocols and configurations to set the homenet naming architecture. Section 6 sums up the various configuration parameters that MAY be filled by the end user on the CPE for example via a GUI. Finally Section 7 provides security considerations.

3. Terminology

- Customer Premises Equipment: (CPE) is the router providing connectivity to the home network. It is configured and managed by the end user. In this document, the CPE MAY also hosts services such as DHCPv6. This device MAY be provided by the ISP.
- Registered Homenet Domain: is the Domain Name associated to the home network.
- DNS Homenet Zone: is the DNS zone associated to the home network. This zone is set by the CPE and essentially contains the bindings between names and IP addresses of the nodes of the home network. In this document, the CPE does neither perform any DNSSEC management operations such as zone signing nor provide an authoritative service for the zone. Both are delegated to the Public Authoritative Server. The CPE synchronizes the DNS Homenet Zone with the Public Authoritative Server via a hidden master / slave architecture. The Public Authoritative Server MAY use specific servers for the synchronization of the DNS Homenet Zone: the Public Authoritative Name Server Set as public available name servers for the Registered Homenet Domain.
- Public Authoritative Server: performs DNSSEC management operations as well as provides the authoritative service for the zone. In this document, the Public Authoritative Server synchronizes the

DNS Homenet Zone with the CPE via a hidden master / slave architecture. The Public Authoritative Server acts as a slave and MAY use specific servers called Public Authoritative Name Server Set. Once the Public Authoritative Server synchronizes the DNS Homenet Zone, it signs the zone and generates the DNSSEC Public Zone. Then the Public Authoritative Server hosts the zone as an authoritative server on the Public Authoritative Master(s).

- DNSSEC Public Zone: corresponds to the signed version of the DNS Homenet Zone. It is hosted by the Public Authoritative Server, which is authoritative for this zone, and is reachable on the Public Authoritative Master(s).
- Public Authoritative Master(s): are the visible name server hosting the DNSSEC Public Zone. End users' resolutions for the Homenet Domain are sent to this server, and this server is a master for the zone.
- Public Authoritative Name Server Set: is the server the CPE synchronizes the DNS Homenet Zone. It is configured as a slave and the CPE acts as master. The CPE sends information so the DNSSEC zone can be set and served.

4. Architecture Overview

Figure 1 provides an overview of the homenet naming architecture.

The CPE is in charge of building the DNS Homenet Zone that contains all FQDN bindings of the home network. The home network is associated to a FQDN, the Registered Homenet Domain (example.com). Any node in the home network is associated to a FQDN (node1.example.com) that MAY be provided via DHCP or statically configured on the CPE via a GUI for example.

The goal of the homenet naming architecture is that the CPE does not handle any DNSSEC operations and does not host the authoritative naming service while FQDNs in the Homenet Zone can be resolved with DNSSEC by any node on the Internet.

In order to achieve this goal, when a node on the Internet sends a DNS(SEC) query like for node1.example.com, this DNS(SEC) query MUST be treated by a third party designated in figure 1 as the Public Authoritative Servers.

The Public Authoritative Servers are in charge of DNS(SEC) traffic for the Registered Homenet Domain (example.com) as well as all DNSSEC management operations like zone signing, key rollover. The DNSSEC

zone hosted by the Public Authoritative Servers is called the DNSSEC Public Zone.

The purpose of our architecture is to describe how the CPE can outsource the DNS Homenet Zone hosted on the CPE to the DNSSEC Public Zone hosted on the Public Authoritative Servers. This includes description of the synchronization protocols between the CPE and the Public Authoritative Servers in Section 5.1 as well as configurations of the DNS Homenet Zone Section 5.2.

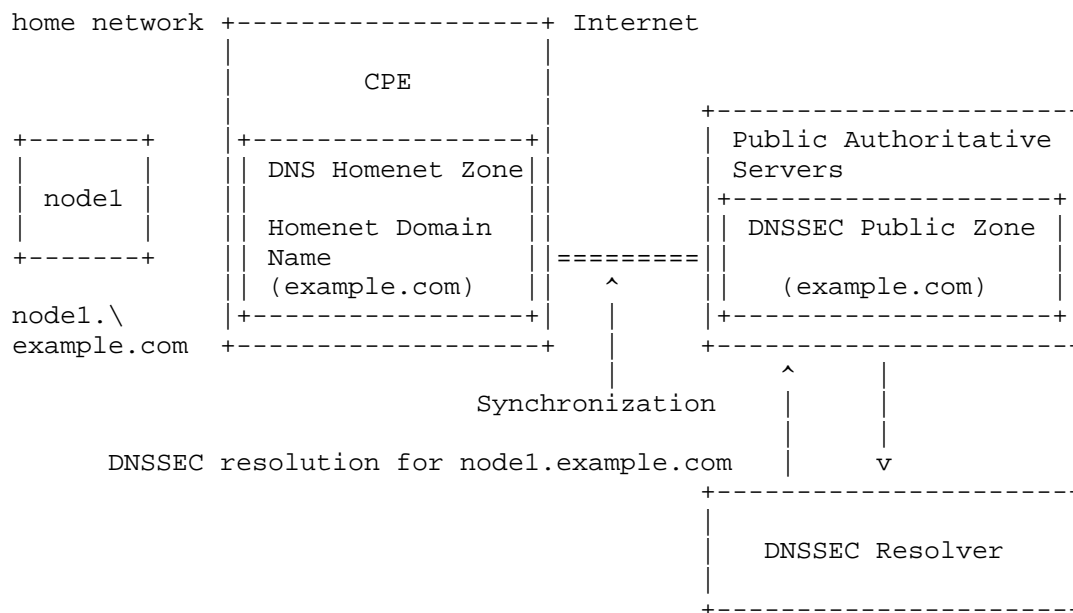


Figure 1: Homenet Naming Architecture Description

The content of the DNS Homenet Zone is out of the scope of this document. The CPE MAY host multiple services like a web GUI, DHCP [RFC6644] or mDNS [RFC6762]. These services MAY coexist and MAY be used to populate the DNS Homenet Zone. This document assumes the DNS Homenet Zone has been populated with domain names that are intended to be publicly published and that are publicly reachable. More specifically, names associated to services or devices that are not expected to be reachable from outside the home network or names bound to non globally reachable IP addresses MUST NOT be part of the DNS Homenet Zone.

Because services or devices MAY only be reached from hosts in the home network, DNS resolution MAY be handled differently from inside the network and from outside the network. This is out of scope of this document. This document is focused on outsourcing the DNS Homenet Zone to the DNS Public Authoritative Servers that are visible from outside the home network. How to deal with a homenet view and a public view is out of the scope of this document. In order to deal with different views, some CPE MAY host DNS forwarders or use DNS view mechanisms.

This document does not make any other assumption on the DNS Homenet Zone that records MUST be made public. More specifically, the DNS Homenet Zone can be a regular or a reverse zone with PTR RRsets. A CPE SHOULD consider both the normal zone as well as the reverse zone and outsource them both to the designated Public Authoritative Servers.

By outsourcing to Public Authoritative Servers, services or devices mentioned in the DNS Homenet Zone MAY be not reachable in case the home network has no internet connectivity. How to keep the naming service within the home network when it is disconnected from the public internet is out of scope of this document. CPE MAY chose for example to host an authoritative naming server for the home network or use a DNS forwarders.

Similarly, CPE MAY host a DNS(SEC) resolution service for nodes in the home network. There are multiple ways to configure the resolver service on the CPE. Detailing these various configurations is out of the scope of this document, and is considered as an implementation issue. Some implementers MAY chose to forward DNS(SEC) queries from the home network to the resolving server of its ISP or any other public resolver. In that case, the DNS(SEC) response from the Public Authoritative Servers is forwarded to the home network, which provide DNS and DNSSEC resolution for the home network. Note also that in this case, the naming service depends on the connectivity with the resolving servers. In case the home network is disconnected, the naming service MAY not be available. Alternative implementations MAY chose to take advantage of forwarders and lookup in the DNS Homenet Zone. This MAY provide only DNS responses in the home network if the CPE does not sign the DNS Homenet Zone. Other implementation MAY chose to synchronize the DNSSEC Public Zone on the CPE either using DNS master slave mechanisms, or by caching the whole zone. This latest option MAY require some additional configuration the Public Authoritative Servers.

5. Architecture Description

This section describes how the CPE and the Public Authoritative Servers SHOULD be configured to outsource authoritative naming service as well as DNSSEC management operations. Section 5.1 describes how a secure synchronization between the CPE and the Public Authoritative server is set. Section 5.2 provides guide lines for the DNS Homenet Zone set in the CPE and uploaded on the Public Authoritative Servers. Section 5.3 describes DNSSEC settings on the Public Authoritative Servers. Finally, Section 5.4 provides the security policies that SHOULD be set on the CPE.

5.1. CPE and Public Authoritative Servers Synchronization

5.1.1. Synchronization with a Hidden Master

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 Authoritative 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 Authoritative Servers. In order to set the master/ slave architecture, the CPE acts as a Hidden Master Server, which is a regular Authoritative DNS(SEC) Server listening on the WAN interface.

The Hidden Master Server is expected to accept SOA [RFC1033], AXFR [RFC1034], and IXFR [RFC1995] queries from its configured slave DNS servers. The Hidden Master Server SHOULD send NOTIFY messages [RFC1996] in order to update Public DNS server zones as updates occur. Because, DNS Homenet Zones are likely to be small, CPE MUST implement AXFR and SHOULD implement IXFR.

Hidden Master Server differs from a regular authoritative server for the home network by:

- Interface Binding: the Hidden Master Server listens on the WAN Interface, whereas a regular authoritative server for the home network would listen on the home network interface.
- Limited exchanges: the purpose of the Hidden Master Server is to synchronizes with the Public Authoritative Servers, not to serve zone. As a result, exchanges are performed with specific nodes (the Public Authoritative Servers). Then exchange types are limited. The only legitimate exchanges are: NOTIFY

initiated by the Hidden Master and IXFR or AXFR exchanges initiated by the Public Authoritative Servers. On the other hand regular authoritative servers would respond any hosts on the home network, and any DNS(SEC) query would be considered. The CPE SHOULD filter IXFR/AXFR traffic and drop traffic not initiated by the Public Authoritative Server. The CPE MUST listen for DNS on TCP and UDP and at least allow SOA lookups to the DNS Homenet Zone.

5.1.2. Securing 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 Authoritative DNS(SEC) Servers.

TSIG [RFC2845] or SIG(0) [RFC2931] can be used to secure the DNS communications between the CPE and the Public DNS(SEC) Servers. TSIG uses a symmetric key which can be managed by TKEY [RFC2930]. Management of the key involved in SIG(0) is performed through zone updates. How to roll the keys with SIG(0) is out-of-scope of this document. The advantage of these mechanisms is that they are only associated with the DNS application. Not relying on shared libraries ease testing and integration. On the other hand, using TSIG, TKEY or SIG(0) requires that these mechanisms to be implemented on the DNS(SEC) Server's implementation running on the CPE, which adds codes. Another disadvantage is that TKEY does not provides authentication mechanism.

Protocols like TLS [RFC5246] / DTLS [RFC6347] can be used to secure the transactions between the Public Authoritative Servers and the CPE. The advantage of TLS/DTLS is that this technology is widely deployed, and most of the boxes already embeds a TLS/DTLS libraries, eventually taking advantage of hardware acceleration. Then TLS/DTLS provides authentication facilities and can use certificates to authenticate the Public Authoritative Server and the CPE. On the other hand, using TLS/DTLS requires to integrate DNS exchange over TLS/DTLS, as well as a new service port. This is why we do not recommend this option.

IPsec [RFC4301] IKEv2 [RFC5996] can also be used to secure the transactions between the CPE and the Public Authoritative Servers. Similarly to TLS/DTLS, most CPE already embeds a IPsec stack, and IKEv2 provides multiple authentications possibilities with its EAP framework. In addition, IPsec can be used to protect the DNS exchanges between the CPE and the Public Authoritative Servers without any modifications of the DNS Servers or client. DNS integration over IPsec only requires an additional security policy in

the Security Policy Database. One disadvantage of IPsec is that it hardly goes through NATs and firewalls. However, in our case, the CPE is connected to the Internet, and IPsec communication between the CPE and Public Authoritative Server SHOULD NOT be impacted by middle boxes.

As mentioned above, TSIG, IPsec and TLS/DTLS may be used to secure transactions between the CPE and the Public Authentication Servers. The CPE and Public Authoritative Server SHOULD implement TSIG and IPsec.

How the PSK can be used by any of the TSIG, TLS/DTLS or IPsec protocols. Authentication based on certificates implies a mutual authentication and thus requires the CPE to manage a private key, a public key or certificates as well as Certificate Authorities. This adds complexity to the configuration especially on the CPE side. For this reason, we recommend that CPE MAY use PSK or certificate base authentication and that Public Authentication Servers MUST support PSK and certificate based authentication.

5.2. DNS Homenet Zone configuration

As depicted in figure 1, the DNSSEC Public Zone is hosted on the Public Authoritative Server, whereas the DNS Homenet Zone is hosted on the CPE. As a result, the CPE MUST configure the DNS Homenet Zone as if the DNS Homenet Zone were hosted by the Public Authoritative Servers instead of the CPE.

If one considers the case where the CPE has a single Homenet Domain Name and has an agreement with a single Public Authoritative Server. In that case, the DNS Homenet Zone SHOULD configure its Name Server RRset and Start of Authority with the ones associated to the Public Authoritative Servers. This is illustrated in figure 2. `public.autho.servers.example.net` is the domain name associated to the Public Authoritative Server, and IP1, IP2, IP3, IP4 are the IP addresses associated.

```
$ORIGIN example.com
$TTL 1h

@ IN SOA public.autho.servers.example.net
      hostmaster.example.com. (
      2013120710 ; serial number of this zone file
      1d        ; slave refresh
      2h        ; slave retry time in case of a problem
      4w        ; slave expiration time
      1h        ; maximum caching time in case of failed
                  ; lookups
      )

@ NS public.authoritative.servers.example.net

public.autho.servers.example.net  A @IP1
public.autho.servers.example.net  A @IP2
public.autho.servers.example.net  AAAA @IP3
public.autho.servers.example.net  AAAA @IP4
```

Figure 2: DNS Homenet Zone

The SOA RRset is defined in [RFC1033], [RFC1035]. This SOA is specific as it is used for the synchronization between the Hidden Master and the Public Authoritative Name Server Set and published on the DNS Public Authoritative Master.

- MNAME: indicates the primary master. In our case the zone is published on the Public Authoritative Master, and its name MUST be mentioned. If multiple Public Authoritative Masters are involved, one of them MUST be chosen. More specifically, the CPE MUST NOT place the name of the Hidden Master.
- RNAME: indicates the email address to reach the administrator. [RFC2142] recommends to use hostmaster@domain and replacing the '@' sign by '.'.
- REFRESH and RETRY: indicate respectively in seconds how often slaves need to check the master and the time between two refresh when a refresh has failed. Default value indicated by [RFC1033] are 3600 (1 hour) for refresh and 600 (10 minutes) for retry. This value MAY be long for highly dynamic content. However, Public Authoritative Masters and the CPE are expected to implement NOTIFY [RFC1996]. Then short values MAY increase the bandwidth usage for slaves hosting large number of zones. As a result, default values looks fine.

EXPIRE: is the upper limit data SHOULD be kept in absence of refresh. Default value indicated by [RFC1033] is 3600000 about 42 days. In home network architectures, the CPE provides both the DNS synchronization and the access to the home network. This device MAY be plug / unplugged by the end user without notification, thus we recommend large period.

MINIMUM: indicates the minimum TTL. Default value indicated by [RFC1033] is 86400 (1 day). For home network, this value MAY be reduced, and 3600 (1hour) seems more appropriated.

When the end user considers multiple Public Authoritative Servers for a given Registered Homenet Domain, the DNS Homenet Zone MAY contain all associated Name Servers and IP addresses.

Some additional verification can check whether the CPE IP address is mentioned in the Public Zone file, and raise a warning to the End User.

5.3. DNSSEC outsourcing configuration

In this document we assumed that the Public Authoritative Server signs the DNS Homenet Zone. Multiple variants MAY be proposed by the Public Authoritative Servers. Public Authoritative Servers MAY propose to sign the DNS Homenet Zone with keys generated by the Public Authoritative Servers and unknown to the CPE. Alternatively some MAY propose the end user to provide the private keys. Although not considered in this document some end user MAY still prefer to sign their zone with their own keys they do not communicate to the Public Authoritative Servers. All these alternatives result from a negotiation between the end user and the Public Authoritative Servers. This negotiation is performed out-of-band and is out of scope of this document.

In this document, we consider that the Public Authoritative Server has all the necessary cryptographic elements to perform zone signing and key management operations.

Note that Public Authoritative Servers described in this document accomplish different functions, and thus different entities MAY be involved.

- DNS Slave function synchronizes the DNS Homenet Zone between the CPE and the Public Authoritative Servers. The DNS Homenet Zone on the Public Authoritative Servers is not available, and the Public Authoritative Server MUST NOT address any DNS queries for that zone. As a result, the Public Authoritative Servers MAY chose a dedicated set of servers to serve the DNS Homenet Zone: the Public Authoritative Name Server Set.
- DNS Zone Signing function signs the DNS Zone Homenet Zone to generate an DNSSEC Public Zone.
- DNSSEC Authoritative Server hosts the naming service for the DNSSEC Public Zone. Any DNS(SEC) query associated to the Homenet Zone SHOULD be done using the specific servers designated as the Public Authoritative Master(s).

5.4. CPE Security Policies

This section details security policies related to the Hidden Master / Slave synchronization.

The Hidden Master, as described in this document SHOULD drop any queries from the home network. This can be performed with port binding and/or firewall rules.

The Hidden Master SHOULD drop on the WAN interface any DNS queries that is not issued from the Public Authoritative Server Name Server Set.

The Hidden Master SHOULD drop any outgoing packets other than DNS NOTIFY query, SOA response, IXFR response or AXFR responses.

The Hidden Master SHOULD drop any incoming packets other than DNS NOTIFY response, SOA query, IXFR query or AXFR query.

The Hidden Master SHOULD drop any non protected IXFR or AXFR exchange. This depends how the synchronization is secured.

6. Homenet Naming Configuration

This section specifies the various parameters required by the CPE to configure the naming architecture of this document. This section is informational, and is intended to clarify the information handled by the CPE and the various settings to be done.

Public Authoritative Servers MAY be defined with the following parameters. These parameters are necessary to establish a secure channel between the CPE and the Public Authoritative Server, and to set the appropriated DNS Homenet Zone file:

- Public Authoritative Name Server Set: The associated FQDNs or IP addresses of the Public Authoritative Server. IP addresses are optional and the FQDN is sufficient. To secure the binding name and IP addresses, a DNSSEC exchange is required. Otherwise, the IP addresses SHOULD be entered manually.
- Authentication Method: How the CPE authenticates itself to the Public Server. This MAY depend on the implementation but we should consider at least IPsec, DTLS and TSIG
- Authentication data: Associated Data. PSK only requires a single argument. If other authentication mechanisms based on certificates are used, then, files for the CPE private keys, certificates and certification authority SHOULD be specified.
- Public Authoritative Master(s): The FQDN or IP addresses of the Public Authoritative Master. It MAY correspond to the data that will be set in the NS RRsets and SOA of the DNS Homenet Zone. IP addresses are optional and the FQDN is sufficient. To secure the binding name and IP addresses, a DNSSEC exchange is required. Otherwise, the IP addresses SHOULD be entered manually.
- Registered Homenet Domain: The domain name the Public Authoritative is configured for DNS slave, DNSSEC zone signing and DNSSEC zone hosting.

Setting the DNS Homenet Zone requires the following information.

- Registered Homenet Domain: The Domain Name of the zone. Multiple Registered Homenet Domain MAY be provided. This will generate the creation of multiple DNS Homenet Zones.
- Public Authoritative Server: The Public Authoritative Servers associated to the Registered Homenet Domain. Multiple Public Authoritative Server MAY be provided.

7. Security Considerations

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

7.1. 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.

7.2. 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. IANA Considerations

This document has no actions for IANA.

9. Acknowledgment

The authors wish to thank Philippe Lemordant for its contributions on the early versions of the draft, 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, Christian Jacquenet, Francis Dupont and Ludovic Eschard for their remarks on CPE and low power devices, Olafur Gudmundsson for clarifying DNSSEC capabilities of small devices, Simon Kelley for its feedback as dnsmasq implementer.

10. References

10.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, 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.
- [RFC2142] Crocker, D., "MAILBOX NAMES FOR COMMON SERVICES, ROLES AND FUNCTIONS", RFC 2142, May 1997.
- [RFC2845] Vixie, P., Gudmundsson, O., Eastlake, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, May 2000.
- [RFC2930] Eastlake, D., "Secret Key Establishment for DNS (TKEY RR)", RFC 2930, September 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.

[RFC6644] Evans, D., Droms, R., and S. Jiang, "Rebind Capability in DHCPv6 Reconfigure Messages", RFC 6644, July 2012.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, February 2013.

10.2. Informational References

[RFC1033] Lottor, M., "Domain administrators operations guide", RFC 1033, November 1987.

Appendix A. Document Change Log

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

-03:

*Simon's comments taken into consideration

*Adding SOA, PTR considerations

*Removing DNSSEC performance paragraphs on low power devices

*Adding SIG(0) as a mechanism for authenticating the servers

*Goals clarification: the architecture described in the document 1) does not describe new protocols, and 2) can be adapted to specific cases for advance users.

-02:

*remove interfaces: "Public Authoritative Server Naming Interface" is replaced by "Public Authoritative Master(s)". "Public Authoritative Server Management Interface" is replaced by "Public Authoritative Name Server Set".

-01.3:

*remove the authoritative / resolver services of the CPE.
Implementation dependent

*remove interactions with mdns and dhcp. Implementation dependent.

*remove considerations on low powered devices

*remove position toward homenet arch

*remove problem statement section

-01.2:

- * add a CPE description to show that the architecture can fit CPEs
- * specification of the architecture for very low powered devices.
- * integrate mDNS and DHCP interactions with the Homenet Naming Architecture.
- * Restructuring the draft. 1) We start from the homenet-arch draft to derive a Naming Architecture, then 2) we show why CPE need mechanisms that do not expose them to the Internet, 3) we describe the mechanisms.
- * I remove the terminology and expose it in the figures A and B.
- * remove the Front End Homenet Naming Architecture to Homenet Naming

-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
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 Wijgmaal
Belgium

Email: wouter.cloetens@softathome.com

Chris Griffiths
Dyn
150 Dow Street
Manchester, NH 03101
US

Email: cgriffiths@dyn.com
URI: <http://dyn.com>

Ralf Weber
Nominum
2000 Seaport Blvd #400
Redwood City, CA 94063
US

Email: ralf.weber@nominum.com
URI: <http://www.nominum.com>

HOMENET
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2014

D. Migault (Ed)
Francetelecom - Orange
W. Cloetens
SoftAtHome
C. Griffiths
Dyn
R. Weber
Nominum
October 20, 2013

DHCP DNS Public Authoritative Server Options
draft-mglt-homenet-naming-architecture-dhc-options-00.txt

Abstract

The home network naming architecture as described in [I-D.mglt-homenet-front-end-naming-delegation] requires a complex naming configuration on the CPE. This configuration MAY not be handled easily by the average end user. Furthermore, such misconfiguration MAY result in making home network unreachable.

This document proposes a DHCP options that provide the CPE all necessary parameters to set up the home network naming architecture.

First, this DHCP options provide automatic configuration and avoid most end users' misconfiguration. Most average end users may not require specific configuration, and their ISP default configuration MAY fully address their needs. In that case, the naming homenet architecture configuration will be completely transparent to the end users. Then, saving naming configuration outside the CPE, makes it resilient to change of CPE or CPE upgrades. Such configuration may also be configured by the end user, via the customer area of their ISP.

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 23, 2014.

Copyright Notice

Copyright (c) 2013 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	3
2. Terminology	3
3. Introduction	4
4. Protocol Overview	6
5. Payload Description	7
5.1. DHCP Zone Public Master Option	7
5.1.1. Unpacking a DHCP Zone Public Master Option	9
5.1.2. Packing a DHCP Zone Public Master Option	9
5.1.3. DHCP Registered Domain Name Option	9
5.1.3.1. Unpacking a DHCP Registered Domain Name Option	10
5.1.3.2. Packing a DHCP Registered Domain Name Option	10
5.1.4. DHCP Master Option	10
5.1.4.1. Unpacking a DHCP Master Option	11
5.1.4.2. Packing a DHCP Master Option	12
5.1.4.3. DHCP Master FQDN Option	12
5.1.4.4. DHCP Master IP4 Option	12
5.1.4.5. DHCP Master IP6 Option	13
5.2. DHCP Public Master Upload Option	14
5.2.1. Unpacking a DHCP Public Master Upload Option	16
5.2.2. Packing a DHCP Public Master Upload Option	16
5.2.3. DHCP Master FQDN List Option	16
5.2.4. DHCP Secure Channel Options	16
5.2.4.1. Unpacking a DHCP Secure Channel Option	17
5.2.4.2. Packing a DHCP Secure Channel Option	18

5.2.4.3.	DHCP Secure Protocol Option	18
5.2.4.4.	DHCP Secure Credential Option	18
5.2.4.4.1.	DHCP PSK Credential Option	19
5.2.4.5.	DHCP Server Set Option	20
5.2.4.5.1.	DHCP Server Set IP4 Option	21
5.2.4.5.2.	DHCP Server Set IP6 Option	21
6.	DHCPv6 Server Behavior	22
7.	DHCPv6 Client Behavior	22
7.1.	Sending an ORO	23
7.2.	Receiving no DHCP Options	23
7.3.	Receiving empty DHCP Options	23
7.4.	Receiving multiple DHCP Options	24
8.	DHCPv6 Relay Behavior	24
9.	IANA Considerations	24
10.	Security Considerations	25
10.1.	DNSSEC is recommended to authenticate DNS hosted data .	25
10.2.	Channel between the CPE and ISP DHCP Server MUST be secured	26
10.3.	CPEs are sensitive to DoS	26
11.	Acknowledgment	27
12.	Document Change Log	27
13.	Pseudo Code	27
13.1.	DHCP Zone Public Master Option	27
13.1.1.	Unpacking a DHCP Zone Public Master Option	27
13.1.2.	Packing a DHCP Zone Public Master Option	28
13.1.3.	DHCP Master Option	29
13.1.3.1.	Unpacking a DHCP Master Option	29
13.1.3.2.	Packing a DHCP Master Option	30
13.2.	DHCP Public Master Upload Option	31
13.2.1.	Unpacking a DHCP Public Master Upload Option	31
13.2.2.	DHCP Secure Channel Options	32
13.2.2.1.	Unpacking a DHCP Secure Channel Option	32
14.	References	33
14.1.	Normative References	33
14.2.	Informational References	34
	Authors' Addresses	34

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. Terminology

- Customer Premises Equipment: (CPE) is the router providing connectivity to the home network. It is configured and managed by the end user. In this document, the CPE MAY also hosts

services such as DHCPv6. This device MAY be provided by the ISP.

- Registered Homenet Domain: is the Domain Name associated to the home network.
- DNS Homenet Zone: is the DNS zone associated to the home network. This zone is set by the CPE and essentially contains the bindings between names and IP addresses of the nodes of the home network. In this document, the CPE does neither perform any DNSSEC management operations such as zone signing nor provide an authoritative service for the zone. Both are delegated to the Public Authoritative Server. The CPE synchronizes the DNS Homenet Zone with the Public Authoritative Server via a hidden master / slave architecture. The Public Authoritative Server MAY use specific servers for the synchronization of the DNS Homenet Zone: the Public Authoritative Name Server Set.
- Public Authoritative Server: performs DNSSEC management operations as well as provides the authoritative service for the zone. In this document, the Public Authoritative Server synchronizes the DNS Homenet Zone with the CPE via a hidden master / slave architecture. The Public Authoritative Server acts as a slave and MAY use specific servers called Public Authoritative Name Server Set. Once the Public Authoritative Server synchronizes the DNS Homenet Zone, it signs the zone and generates the DNSSEC Public Zone. Then the Public Authoritative Server hosts the zone as an authoritative server on the Public Authoritative Master(s).
- DNSSEC Public Zone: corresponds to the signed version of the DNS Homenet Zone. It is hosted by the Public Authoritative Server, which is authoritative for this zone, and is reachable on the Public Authoritative Master(s).
- Public Authoritative Master(s): are the visible name server hosting the DNSSEC Public Zone. End users' resolutions for the Homenet Domain are sent to this server, and this server is a master for the zone.
- Public Authoritative Name Server Set: is the server the CPE synchronizes the DNS Homenet Zone. It is configured as a slave and the CPE acts as master. The CPE sends information so the DNSSEC zone can be set and served.

3. Introduction

With IPv6, nodes inside the home network are expected to be globally reachable. CPEs are already providing connectivity to the home network, and most of the time already assigns IP addresses to the nodes of the home network using for example DHCPv6.

This makes CPE good candidate for defining the DNS zone file of the home network. However, CPEs have not been designed to handle heavy traffic, nor heavy operations. As a consequence, CPE SHOULD neither host the authoritative naming service of the home network, nor handle DNSSEC operations such as zone signing. In addition, CPE are usually managed by end users, and the average end user is most likely not mastering DNSSEC to administrate its DNSSEC zone. As a result, CPE SHOULD outsource both the naming authoritative service and its DNSSEC management operations to a third party. This architecture, designated as the homenet naming architecture is described in [I-D.mglt-homenet-front-end-naming-delegation], and the third party is designated as the Public Authoritative Servers.

The home network naming architecture [I-D.mglt-homenet-front-end-naming-delegation] defines how the CPE and the Public Authoritative Servers interact together, to leverage some of the issues related to the CPE, and the DNSSEC understanding of the average end user. Even though most of the DNSSEC issues are outsourced to the Public Authoritative Servers, setting the homenet naming architecture still requires some configurations.

Configuration is fine as it provides the opportunity for advanced end users to make the naming architecture fit their specific needs. However most of the end users do not want to configure the homenet naming architecture. In most cases, the end users wants to subscribe and plug its CPE. The CPE is expected to be configured to set automatically and transparently the appropriated home network naming architecture.

Using DHCP options to provide the necessary parameters for setting the homenet naming architecture provides multiple advantages. Firstly, it makes the network configuration independent of the CPE. Any new plugged CPE configures itself according to the provided configuration parameters. Secondly, it saves the configuration outside the CPE, which prevents re-configuring the CPE when it is replaced or reset. Finally ISPs are likely to propose a default homenet naming architecture that may address most of the end users needs. For these end users, no configuration will be performed at any time. This avoids unnecessary configurations or misconfiguration that could result in isolating the home network. For more advanced end users, the configuration MAY be provided also via the web GUI of the ISP's customer area for example. This configuration MAY enable third party Public Authoritative Servers. By doing so, these end

users will also benefit from CPE-indepdent configuration and configuration backup.

This document considers the architecture described in [I-D.mglt-homenet-front-end-naming-delegation]. The DNS(SEC) zone related to the home network is configured and set by the CPE and hosted on a Public Authoritative Server. [I-D.mglt-homenet-front-end-naming-delegation] describes how the synchronization between the CPE and the Public Authoritative Server is performed. This document describes DHCP options that provide the necessary parameters to the CPE to set the architecture described in [I-D.mglt-homenet-front-end-naming-delegation].

Section 4 presents an overview of the DHCP options presented in this document and Section 5 describes the format of this option and Section 6, Section 7 and Section 8 details the behavior of respectively the DHCP Client, the DHCP Server and the DHCP Relay.

This document assumes the reader is familiar with [I-D.mglt-homenet-front-end-naming-delegation].

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

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

4. Protocol Overview

To properly configure the home network naming architecture defined in [I-D.mglt-homenet-front-end-naming-delegation], the CPE MUST:

- 1: Determine which Registered Domain are considered. Each Registered Domain is associated to a DNS Zone file. Note that a CPE MAY publish a single zone under different Registered Domain Names, or set different contents on different Registered Domain Names.

- 2: Properly generate the DNS Zone file and associate the corresponding authoritative Name Server (RRset NS) with associated IP addresses (RRsets A or AAAA). The CPE derives the NS RRset from the Registered Domain and the Public Authoritative Master. Then it MAY derive the glue A or AAAA records from the Public Authoritative Master and associated IP addresses. These pieces of information are provided by the DHCP Zone Public Master Option (OPTION_ZONE_PUBLIC_MASTER).
- 3: Upload the Zone files to the Public Authoritative Master. Uploading the DNS Homenet Zone or the DNSSEC Homenet Zone may not be done directly from the CPE to the Public Authoritative Masters. In fact, the Public Authoritative Server MAY have a dedicated server for DNS zone uploads: the Public Authoritative Name Server Set. One of the reason is that the Public Authoritative Server MAY perform some extra operations such as the DNSSEC signing before publishing the DNSSEC Homenet Zone to on the Public Authoritative Masters. As a result, for each Public Authoritative Master a secure channel MUST be established between the CPE and the Public Authoritative Name Server Set. How to set, for each Public Authoritative Master, the secure channel to the Public Authoritative Name Server Set is provided by the DHCP Public Master Upload Option (OPTION_PUBLIC_MASTER_UPLOAD).

A common way for the CPE to collect these pieces of information is to send an Option Request DHCP Option (ORO) [RFC3315] for the DHCP DNS Public Authoritative Server Option and for the DNS Public Authoritative Name Server Set Option. If the DHCP Sever understand these options, it MAY send back one or multiple instance for each option. Then, the DHCP Client sets the naming architecture.

Similarly, the DHCP Server MAY provide the DHCP DNS Public Authoritative Server Option and for the DNS Public Authoritative Name Server Set Option without any request from the DHCP Client.

Note that how the CPE manage the multiple DNS Homenet Zones is implementation dependent. It MAY synchronize all DNS Homenet Zone with the Public Authoritative Servers, or use zone redirection mechanisms like like CNAME [RFC2181], [RFC1034], DNAME [RFC6672] or CNAME+DNAME [I-D.sury-dnsexst-cname-dname]. In the first case, any update requires to update all zone, whereas redirection MAY require only updating a single DNS Homenet Zone.

5. Payload Description

5.1. DHCP Zone Public Master Option

The DHCP Zone Public Master Option (OPTION_ZONE_PUBLIC_MASTER) is used by the CPE to set the DNS Homenet Zone with the proper NS RRsets and the associated IP addresses. The DHCP Zone Public Master Option provides bindings between Registered Domain Names and Public Authoritative Master.

Following Section 9 of [I-D.ietf-dhc-option-guidelines], the DHCP Zone Master Option encapsulates the DHCP Registered Domain Name Option (OPTION_REGISTERED_DOMAIN_NAMES) that contains a list of registered domain names and the Public Authoritative DHCP Master Option (OPTION_MASTER) that contains the FQDNs and IP addresses of each Public Authoritative Masters.

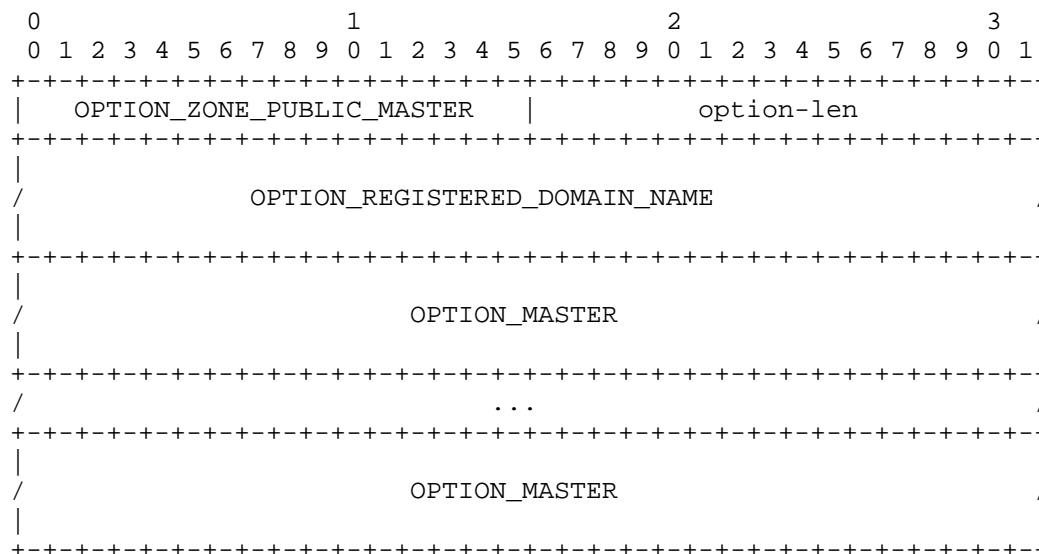


Fig 1: DHCP Zone Public Master Option

- OPTION_ZONE_PUBLIC_MASTER: the option code for the DHCP Zone Public Master Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- OPTION_REGISTERED_DOMAIN_NAME: the list of Registered Homenet Domains.
- OPTION_MASTER: the necessary information to configure properly the DNS Homenet Zone file with NS, A and AAA RRsets associated to the Public Authoritative Master(s).

5.1.1. Unpacking a DHCP Zone Public Master Option

When a DHCP Zone Public Master Option is received by the DHCP Client, if the DHCP Registered Domain Name Option does not exist or is void, the CPE ignores the DHCP Zone Public Master Option. It MAY indicate the DHCP Server supports these options but they are not properly configured. Otherwise, it selects all DNS Homenet Zone designated by the DHCP Registered Domain Name Option and adds the Public Authoritative NS, A and AAA records provided by the DHCP Master Option. If DHCP Master Option are missing, the CPE hosts the DNS Homenet Zone for the Registered Domains.

All DHCP Options are propositions. The CPE MAY chose a subset of these according to its policies.

Section 13 illustrates with pseudo code how this MAY be performed.

5.1.2. Packing a DHCP Zone Public Master Option

The DHCP Server sends a DHCP Zone Public Master Option to bind Registered Domain Names to a list of Public Authoritative Masters. How to collect these pieces of information is implementation dependent, and depends on the data structure that stores the information. However, we can reasonably assume that sending a DHCP Zone Public Master Option is composed of two phases:

- 1) First collect all Registered Domain with their associated list of Public Authoritative Masters. Basic implementation MAY build DHCP Zone Public Master Option for each Registered Domain. However, we recommend to group all Registered Domain with the same list of Public Authoritative Masters. This leads to a list of Registered Domain associated to a list of Public Authoritative Masters. Note that lists of Public Authoritative Masters are equals if they have the same Public Authoritative Masters, that is for each of them the same FQDN and the same list of IP addresses. The list is not ordered.
- 2) Then, for each binding build a new DHCP Zone Public Master Option, build DHCP Registered Domain Name Option from the list of Registered Domains, add it to the DHCP Zone Public Master Option. For each Public Authoritative Master of the list of Authoritative Masters, build a DHCP Master Option and add it to the DHCP Zone Public Master Option.

Section 13 illustrates with pseudo code how this MAY be performed.

5.1.3. DHCP Registered Domain Name Option

The DHCP Registered Domain Name Option (OPTION_REGISTERED_DOMAIN_NAME) contains a list of DNS domain names. It MAY have multiple FQDNs. This option follows the description of section 5.10 [I-D.ietf-dhc-option-guidelines].

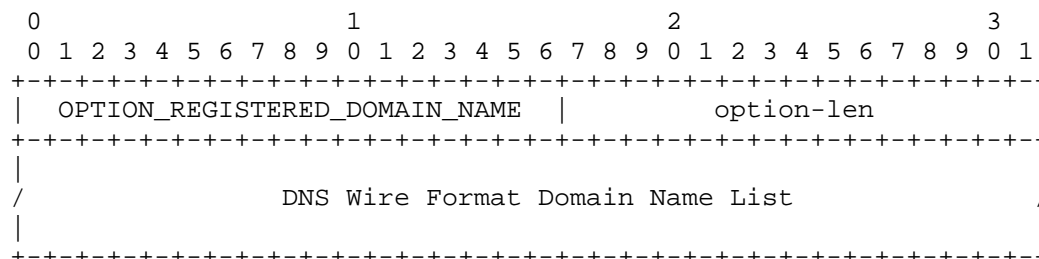


Fig 2: DHCP Registered Domain Name Option

- OPTION_REGISTERED_DOMAIN_NAME: the option code for the DHCP Registered Domain Name Option
- option-len: length in octets of the option-data field as described in [RFC3315].
- DNS Wire Format Domain Name List: The special encoding of this field supports carrying multiple instances of hosts or domain names in a single option, by terminating each instance with a byte of 0 value.

5.1.3.1. Unpacking a DHCP Registered Domain Name Option

The DHCP Registered Domain Name Option MAY return one or multiple Registered Domain Names. The DHCP Client MUST remove empty strings from the list.

5.1.3.2. Packing a DHCP Registered Domain Name Option

The DHCP Registered Domain Name Option is build from a list of non-empty strings.

5.1.4. DHCP Master Option

The DHCP Master Option provides the FQDN and associated IP addresses of the Public Authoritative Master. Following Section 9 of [I-D.ietf-dhc-option-guidelines], the DHCP Master Option encapsulates the DHCP Master FQDN Option (OPTION_MASTER_FQDN) that contains a single FQDN followed by a DHCP Master IP4 Option (OPTION_MASTER_IP4) or a DHCP Master IP6 Option (OPTION_MASTER_IP6) that contains the

associated IP addresses. Only a single DHCP Master IP6 Option or DHCP Master IP4 Option is expected.

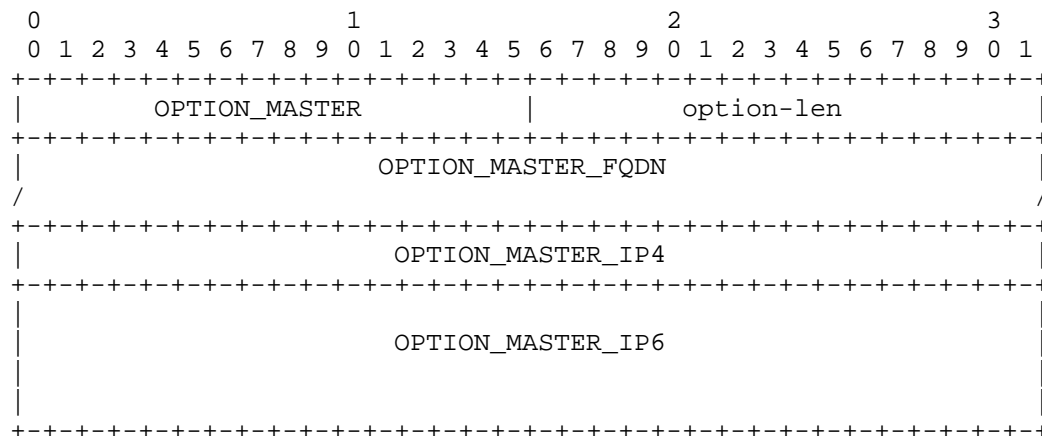


Fig 3: DHCP Master Option

- OPTION_MASTER: the option code for the DHCP Master Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- OPTION_MASTER_FQDN: the DHCP Master FQDN Option with FQDN associated to the Public Authoritative Server. This option is mandatory.
- OPTION_MASTER_IP4: the DHCP Master IP4 Option with IPv4 address associated to the Public Authoritative Server. This option is optional, however the DHCP server SHOULD provide at least one DHCP Master IP4 Option or one DHCP Master IP6 Option.
- OPTION_MASTER_IP6: the DHCP Master IP6 Option with IPv6 address associated to the Public Authoritative Server. This option is optional, however the DHCP server SHOULD provide at least one DHCP Master IP6 Option or one DHCP Master IP6 Option.

5.1.4.1. Unpacking a DHCP Master Option

The DHCP Master FQDN Option is mandatory, and a DHCP Master Option that do not encapsulate a DHCP Master FQDN Option MUST be ignored. An empty DHCP Master FQDN Option indicates the CPE and a FQDN MUST be provided by the CPE. DHCP Master IP4 Options and DHCP Master IP6 Options are optional. Following Section 8 of [I-D.ietf-dhc-option-

guidelines], providing IP addresses avoids DNS(SEC) resolutions which unnecessarily load the network and delay the configuration. As a result, it is recommended to provide the IP addresses. If at least a single non void DHCP Master IP4 Option or DHCP Master IP6 Option is provide, the DHCP Client MUST NOT perform any DNS(SEC) resolution. Otherwise, the DHCP Client SHOULD perform a DNSSEC resolution.

Section 13 illustrates with pseudo code how this MAY be performed.

5.1.4.2. Packing a DHCP Master Option

DHCP Master Options are built from the master object. If the FQDN of the Public Authoritative Master is empty, the DHCP Master Option MUST NOT be built. If no IP address has been provisioned, the DHCP Server SHOULD perform a DNS(SEC) resolution and provide the IP addresses.

Section 13 illustrates with pseudo code how this MAY be performed.

5.1.4.3. DHCP Master FQDN Option

The DHCP Master FQDN Option (OPTION_MASTER_FQDN) designates the FQDN of the Public Authoritative Server. Only one FQDN is expected. This option follows the description of section 5.10 [I-D.ietf-dhc-option-guidelines].

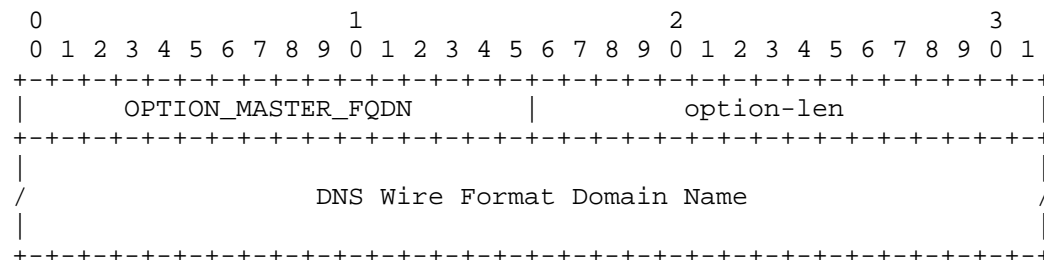


Fig 4: DHCP Master FQDN Option Format

- OPTION_MASTER_FQDN: the option code for the DHCP Master FQDN Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- DNS Wire Format Domain Name: A single FQDN.

5.1.4.4. DHCP Master IP4 Option

This section defines the IP addresses associated to the master. This option follows the recommendation of section 5.1 and 8 of [I-D.ietf-dhc-option-guidelines].

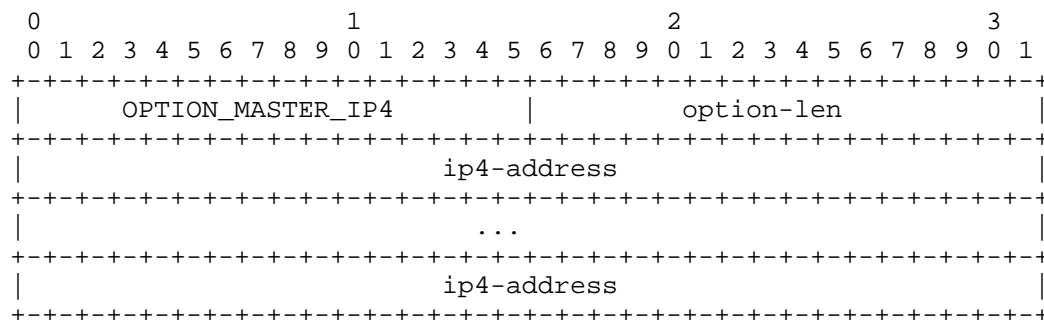


Fig 5: DHCP Master IP4 Option Format

- OPTION_MASTER_IP4: the option code for the DHCP Master IP4 Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- ip4-address: the 32 bit value of the IPv4 address.

5.1.4.5. DHCP Master IP6 Option

This section defines the IP addresses associated to the master. This option follows the recommendation of section 5.1 and 8 of [I-D.ietf-dhc-option-guidelines].

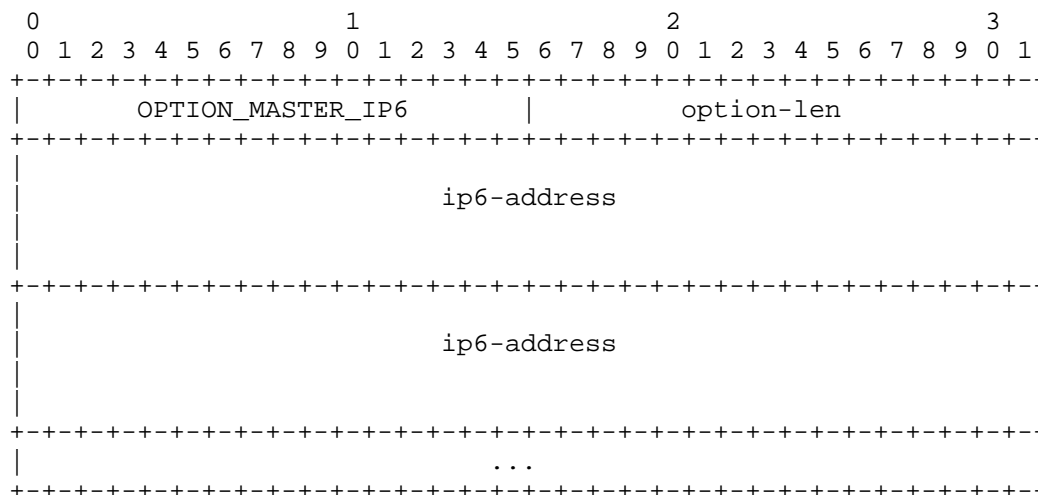


Fig 6: DHCP Master IP6 Option Format

- OPTION_MASTER_IP6: the option code for the DHCP Master IP6 Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- ip6-address: the 128 bit value of IPv6 address.

5.2. DHCP Public Master Upload Option

The DHCP Public Master Upload Option (OPTION_PUBLIC_MASTER_UPLOAD) is used to associate a secure channel for each Public Authoritative Master.

More specifically, to publish a DNS Homenet Zone on a given Public Authoritative Master, the CPE establish a secure channel with the Public Authoritative Name Server Set and upload the DNS Homenet Zone using master / slave mechanisms. It is then the responsibility of the Public Authoritative Name Server Set to publish the DNS Homenet Zone to its associated Public Authoritative Masters.

The DHCP Public Master Upload Option enables the CPE to set an address book where the key is the Public Authoritative and the value is a set of Secure Channels. For each DNS(SEC) Homenet Zone, the CPE is expect to list the Public Authoritative Master and for each of them upload the DNS(SEC) Homenet Zone file to the Public Authoritative Name Server Set via a Secure Channel.

Following Section 9 of [I-D.ietf-dhc-option-guidelines], the DHCP Public Master Upload Option encapsulates a DHCP Master FQDN List Option (OPTION_MASTER_FQDN_LIST) that contains the list of the FQDNs of the Public Authoritative Master and a list of DHCP Secure Channel Options (OPTION_SECURE_CHANNEL) that list how the CPE can upload the DNS(SEC) Homenet Zone. For each of the mentioned Public Authoritative Master, the CPE is expected to consider one of the DHCP Secure Channel Options to upload the Zone.

The DHCP Master FQDN List Option is mandatory and MUST be unique. At least one DHCP Secure Channel Options MUST be encapsulated.

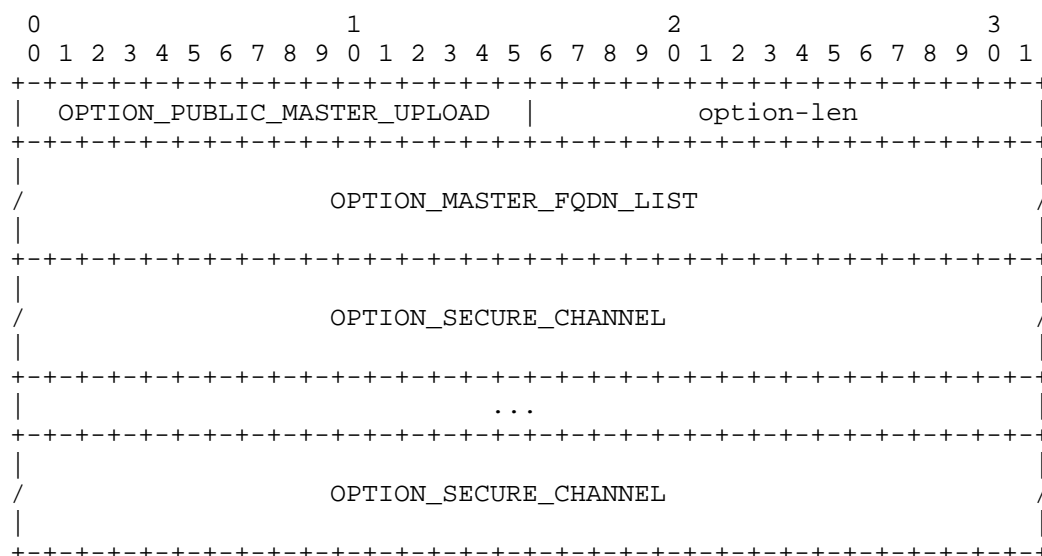


Fig 7: DHCP Public Master Upload Option Format

- OPTION_PUBLIC_MASTER_UPLOAD: the option code that corresponds to the DHCP Public Master Upload Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- OPTION_MASTER_FQDN_LIST: The DHCP Master FQDN List Option. A single DHCP Master FQDN List Option MUST be encapsulated in the DHCP Public Master Upload Option.
- OPTION_SECURE_CHANNEL: The DHCP Secure Channel Option. One or multiple DHCP Secure Channel Option MUST be encapsulated in the DHCP Public Master Upload Option.

5.2.1. Unpacking a DHCP Public Master Upload Option

When the DHCP Client receives a DHCP Public Master Upload Option, it builds a dictionary between, Public Authoritative Master FQDN and possible Secure Channels. Secure Channel that do not correspond to the CPE policy, MAY be discarded.

This binding will be used latter when the CPE uploads the DNS(SEC) Homenet Zone files. Section 13 illustrates with pseudo code how this MAY be performed.

5.2.2. Packing a DHCP Public Master Upload Option

The DHCP Server SHOULD send at least a Secure Channel for each of the Public Authoritative Master. All Public Authoritative Masters associated to the DHCP Client that are provided in a DHCP Zone Public Master Option MUST be mentioned in an DHCP Public Master Upload Option.

Packing a DHCP Public Master Upload Option is performed in a similar way as the DHCP Zone Public Master Option. A binding between the Public Authoritative Master and the Secure Channels is performed. Then, this dictionary is factorized as described in Section 5.1.2. More specifically, all keys with the same value are grouped.

5.2.3. DHCP Master FQDN List Option

The DHCP Master FQDN List Option is similar to the DHCP Master FQDN Option except that it can indicate multiple Master FQDNs.

5.2.4. DHCP Secure Channel Options

The DHCP Secure Channel Option (OPTION_SECURE_CHANNEL) indicates:

- 1: How to secure the channel, that is to say which protocols are used to secure the channel. This is indicated by the DHCP Secure Protocol Option (OPTION_SECURE_PROTOCOL).
- 2: The security credential used to set up the secure channel, that is to say the cryptographic material to authenticate the CPE and the Public Authoritative Name Server Set. This is carried by the DHCP Secure Credential Option (OPTION_SECURE_CREDENTIAL).
- 3: The Public Authoritative Name Server Set the secure channel is established with, that is to say its IP addresses. These IP addresses are carried by the DHCP Server Set Option (OPTION_SERVER_SET).

Following Section 9 of [I-D.ietf-dhc-option-guidelines], the DHCP Secure Channel Option encapsulates the DHCP Secure Protocol Option, the DHCP Secure Credential Option and the DHCP Server Set Option. Each of these options is mandatory and MUST appear only once.

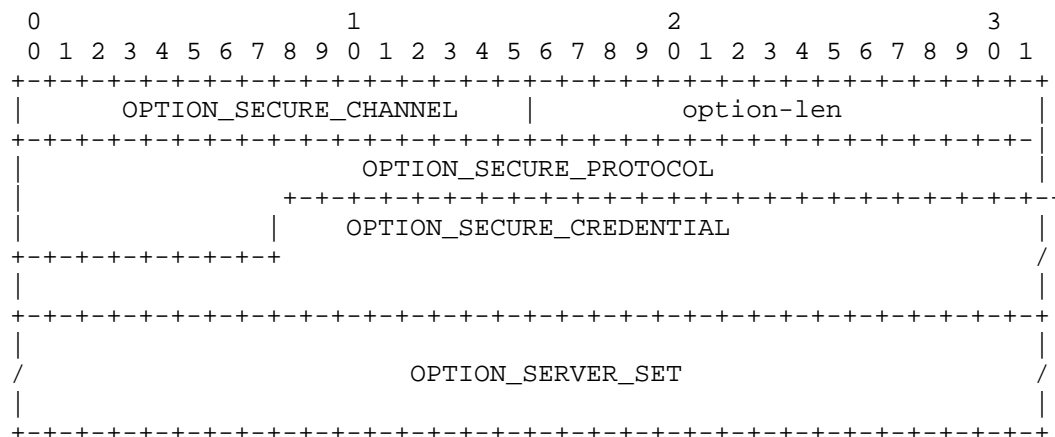


Fig 8: DHCP Secure Channel Option Format

- `OPTION_SECURE_CHANNEL`: The option code for the DHCP Secure Channel Option.
- `option-len`: length in octets of the option-data field as described in [RFC3315].
- `OPTION_SECURE_PROTOCOL`: the DHCP Secure Protocol Option. This option is mandatory and MUST appear only once.
- `OPTION_SECURE_CREDENTIAL`: the DHCP Secure Credential Option. This option is mandatory and MUST appear only once.
- `OPTION_SERVER_SET`: the DHCP Server Set Option. This option is mandatory and MUST appear only once.

5.2.4.1. Unpacking a DHCP Secure Channel Option

When the DHCP Secure Channel Option is received, the DHCP Client MUST check that DHCP Secure Protocol Option, DHCP Secure Credential Option and DHCP Server Set Option are unique. If not, the DHCP Client MUST ignore the DHCP Secure Channel Option.

The DHCP Client MUST also check proposition match its policies and Secure Credentials match the Secure Protocol.

Section 13 illustrates with pseudo code how this MAY be performed.

5.2.4.2. Packing a DHCP Secure Channel Option

Packing the DHCP Secure Channel Option from the Secure_channel object is straight forward.

5.2.4.3. DHCP Secure Protocol Option

The DHCP Secure Protocol Option (OPTION_SECURE_PROTOCOL) is a 8-bit integer option that fills recommendation of section 5.6 of [I-D.ietf-dhc-option-guidelines].

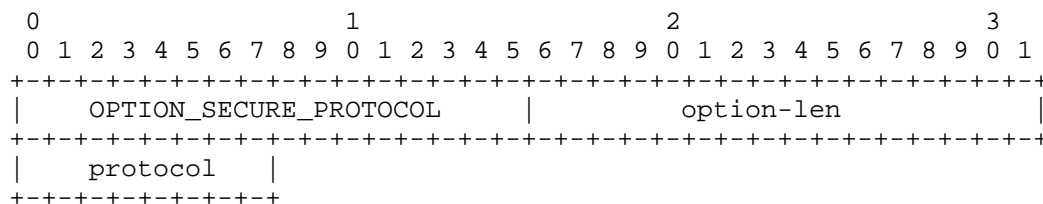


Fig 9: DHCP Secure Protocol Option Format

- OPTION_SECURE_PROTOCOL: The opcode for the DHCP Secure Protocol Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- protocol: the 8 bit value that indicates which protocol MUST be used between the CPE and the Name Server Set for this secure channel.

The protocol detailed in this document are:

- NONE: TBD
- TSIG: TBD
- IPSEC: TBD
- SIG(0): TBD

5.2.4.4. DHCP Secure Credential Option

The DHCP Secure Credential Option encapsulates the necessary element to authenticate and set up the secure channel. Currently, only the

DHCP PSK Credential Option (OPTION_PSK_CREDENTIAL) is defined in this document but the use of DHCP Secure Credential Option enables makes possible the use of different credential in the future.

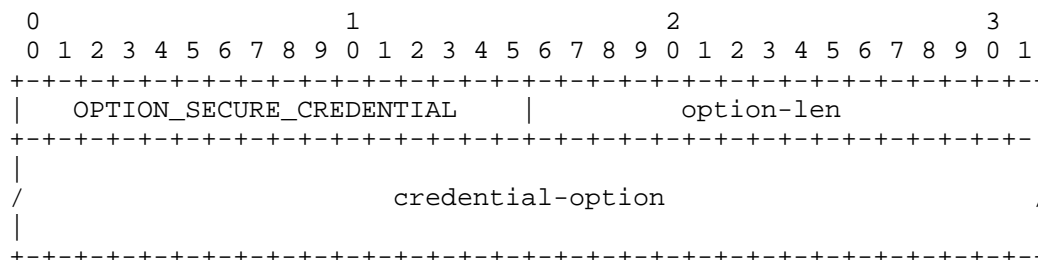


Fig 10: DHCP Secure Credential Option Format

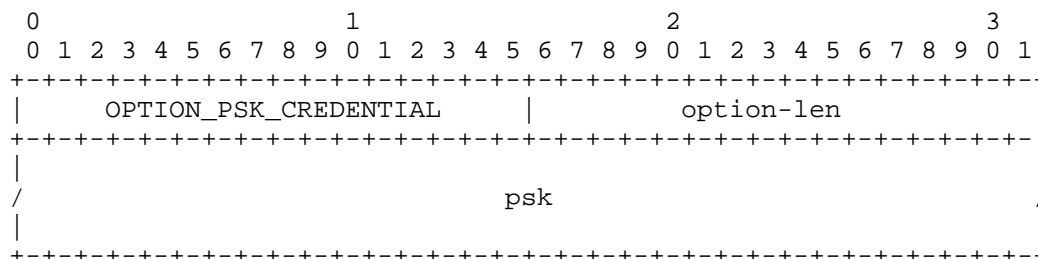
- OPTION_SECURE_CREDENTIAL: The option code for the DHCP Secure Credential Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- credential-option: A DHCP Credential Option.

5.2.4.4.1. DHCP PSK Credential Option

The DHCP PSK Credential Option (OPTION_PSK_CREDENTIAL) contains the pre-shared key. It can be used with IPsec, TSIG. If SIG(0) is selected as a protocol, the DHCP server MUST NOT provide this option, and it MUST be ignored by the DHCP Client.

Note that PSK MUST NOT be sent by the DHCP Server over an non trusted channel. In addition a DHCP Server SHOULD NOT provide the PSK unless requested explicitly by the DHCP Client. Similarly, the DHCP Client MUST NOT request the PSK if the channel between the DHCP Client and the DHCP Server is not trusted.

This option follows recommendation of section 5.9 of [I-D.ietf-dhc-option-guidelines].

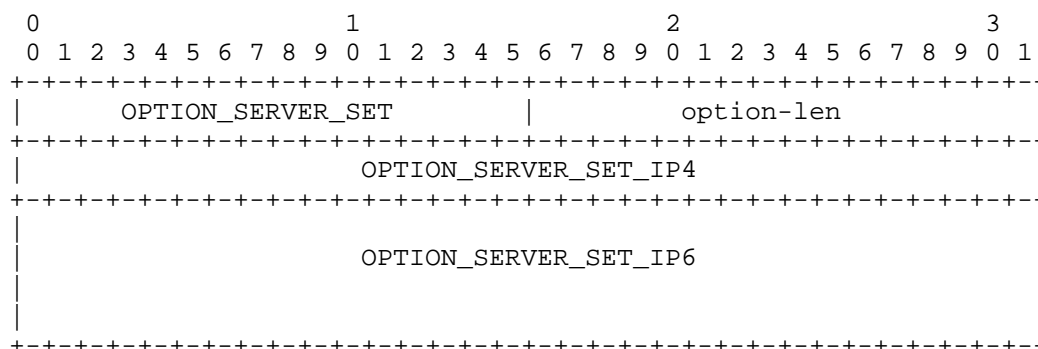


- OPTION_PSK_CREDENTIAL: The opcode for the DHCP PSK Credential Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- psk: raw preshared key.

5.2.4.5. DHCP Server Set Option

The DHCP Server Set Option (OPTION_SERVER_SET) carries the IP addresses of the Public Authoritative Name Server Set. It is recommended to have one IP address, and eventually two if the Public Authoritative Server is dual stack. In any case no more than one IP address of each family is expected. The use of IP addresses instead of FQDN follows recommendation of [I-D.ietf-dhc-option-guidelines] section 8.

The DHCP Server Set Option encapsulates the DHCP Server Set IP4 Option (OPTION_SERVER_SET_IP4) and the DHCP Server Set IP6 Option (OPTION_SERVER_SET_IP6).



- OPTION_SERVER_SET: The option code for the DHCP Server Set Option.

- option-len: length in octets of the option-data field as described in [RFC3315].
- OPTION_SERVER_SET_IP4: DHCP Server Set IP4 Option with IPv4 address associated to the Public Authoritative Name Server Set. This option is optional, however the DHCP server SHOULD provide at least one DHCP Server Set IP4 Option or one DHCP Server Set IP6 Option.
- OPTION_SERVER_SET_IP6: DHCP Server Set IP6 Option with IPv4 address associated to the Public Authoritative Name Server Set. This option is optional, however the DHCP server SHOULD provide at least one DHCP Server Set IP4 Option or one DHCP Server Set IP6 Option.

5.2.4.5.1. DHCP Server Set IP4 Option

The DHCP Server Set IP4 Option (OPTION_SERVER_SET_IP4) carries the IP address of the Public Authoritative Name Server Set. This option follows the recommendation of section 5.1 and 8 of [I-D.ietf-dhc-option-guidelines].

```

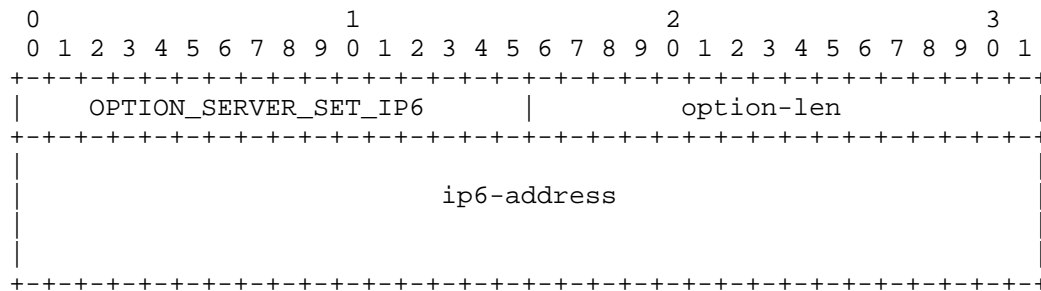
      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|  OPTION_SERVER_SET_IP4  |  option-len  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               ip4-address                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

- OPTION_SERVER_SET_IP4: The option code for the DHCP Server Set IP4 Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- ip4-address: the 32 bit value of the IPv4 address.

5.2.4.5.2. DHCP Server Set IP6 Option

The DHCP Server Set IP6 Option (OPTION_SERVER_SET_IP6) carries the IP address of the Public Authoritative Name Server Set. This option follows the recommendation of section 5.1 and 8 of [I-D.ietf-dhc-option-guidelines].



- OPTION_SERVER_SET_IP6: The option code for the DHCP Server Set IP6 Option.
- option-len: length in octets of the option-data field as described in [RFC3315].
- ip6-address: the 128 bit value of the IPv6 address.

6. DHCPv6 Server Behavior

The DHCPv6 server MAY send DHCP Zone Public Master Option and/or DHCP Public Master Upload Option upon receiving or not a DHCP Option Request Option (ORO) by the DHCP Client.

The DHCP Server MAY send zero, one or multiple DHCP Zone Public Master Option or DHCP Public Master Upload Option.

The DHCP Server MUST send these option over a trusted channel. The PSK MUST NOT be sent over a non trusted channel.

If the DHCP Server is not provisioned properly, it SHOULD send empty DHCP Zone Public Master Option or DHCP Public Master Upload Option to indicate it supports the options, but they are not provisioned properly.

Although DHCP Zone Public Master Option and DHCP Public Master Upload Option are different options, they MAY be used together by the DHCP Client. Unless there are good reasons, DHCP Servers SHOULD provide in their DHCP Public Master Upload Option a Secure Channel for all Public Authoritative Masters mentioned in the DHCP Zone Public Master Option. In other words, for all Public Authoritative Masters mentioned in the DHCP Zone Public Master Option, the DHCP Server SHOULD send a DHCP Public Master Upload Option that provides a Secure Channel.

7. DHCPv6 Client Behavior

7.1. Sending an ORO

The DHCPv6 Client MAY enable different policies to configure the Home Network Naming Architecture. More specifically, it MAY allow an end user to set manually a specific naming configuration, which may disable automatic configuration of the Home Network Naming Architecture. If automatic configuration of the Home Network is not considered, the CPE SHOULD NOT send a DHCP Option Request Option (ORO) from the CPE for a DHCP DNS Public Authoritative Server Option or for a DHCP Public Master Upload Option. This would otherwise unnecessarily load the DHCPv6 Server of the ISP and the network.

The DHCP Client SHOULD NOT send an ORO for a DHCP Zone Public Master Option or a DHCP Public Master Upload Option if the channel between the DHCP Client and the DHCP Server is not trusted.

By sending an DHCP Option Request Option (ORO) for a DHCP Zone Public Master Option or a DHCP Public Master Upload Option, the CPE indicates that the response received from the DHCP Server will define how the Naming Architecture of the CPE is configured. However, this does not necessarily mean that the CPE will automatically configure its Naming Architecture according to the elements provided by the DHCPv6 Server. In fact the CPE MAY implement various policies to configure the Naming Architecture. Some policies MAY merge configuration manually provided by the end user and those provided by the DHCPv6 Server, others MAY only accept a subset of Public Authoritative Name Servers provided by the DHCPv6 Server. Defining the selection policies of the CPE is how of scope of this document.

The remaining of the section describes how the DHCPv6 Client handles information received by the DHCPv6 Server to configure the Naming Architecture. We assume that all information are considered. Although this document restricts the description to a single use case, we believe this will be the most common and basic use case. In addition, other uses cases implementing different configuration policies only requires small modifications to the use case considered in this section.

7.2. Receiving no DHCP Options

If the DHCPv6 Client does not receive any DHCP Zone Public Master Option or DHCP Public Master Upload Option from the DHCPv6 Server. The DHCPv6 Client assumes the DHCPv6 Server does not support the option. In this case, the Naming Architecture can only be set from local settings.

7.3. Receiving empty DHCP Options

By receiving empty DHCP Zone Public Master Option or empty DHCP Public Master Upload Option. The DHCP Client assumes the DHCP Server supports these options, but they are not or badly provisioned.

7.4. Receiving multiple DHCP Options

The DHCPv6 Client MAY receive one or multiple DHCP Zone Public Master Option and/or DHCP Public Master Upload Option. From these Option the CPE is expected to :

- 1: Collect all DHCP Zone Public Master Options.
- 2: Set the DNS Homenet Zone with the Public Authoritative Servers associated to each Registered Homenet Domain. This includes setting NS and Public Authoritative Server A/AAA RRsets. The CPE is expected to proceed to some checks so these RRsets are valid.
- 3: Collect all DHCP Public Master Upload Option.
- 4: Build the `master_secure_channel_dict` dictionary that associates to each Public Authoritative Master a list of possible secure channels.
- 5: Upload the DNS Homenet Zone to the Public Authoritative Name Server Set so the zone can be published on the corresponding Public Authoritative Servers. The DNS(SEC) Homenet Zone is considered uploaded if for all Public Authoritative Masters of the DNS(SEC) Homenet Zone, upload (i.e. master/slave synchronization) succeeded at least with one Secure Channel. If for a given Public Authoritative Master, upload fails with all its Secure Channel, it MUST be removed from the DNS(SEC) Homenet Zone and upload MUST restarted. In other words synchronization MUST be performed again with all Public Authoritative Masters of the DNS(SEC) Homenet Zone (not only the remaining ones).

8. DHCPv6 Relay Behavior

DHCP Relay behavior are not modified by this document.

9. IANA Considerations

The DHCP options detailed in this document is:

- OPTION_ZONE_PUBLIC_MASTER: TBD
- OPTION_REGISTERED_DOMAIN_NAME: TBD

- OPTION_MASTER: TBD
- OPTION_MASTER_IP4: TBD
- OPTION_MASTER_IP6: TBD
- OPTION_MASTER_FQDN: TBD
- OPTION_PUBLIC_MASTER_UPLOAD: TBD
- OPTION_MASTER_FQDN_LIST: TBD
- OPTION_SECURE_CHANNEL: TBD
- OPTION_SECURE_PROTOCOL: TBD
- OPTION_SECURE_CREDENTIAL: TBD
- OPTION_SERVER_SET: TBD
- OPTION_SERVER_SET_IP4: TBD
- OPTION_SERVER_SET_IP6: TBD

The security-protocol detailed in this document are:

- NONE: TBD
- TSIG: TBD
- IPSEC: TBD
- SIG(0): TBD

The security-credential detailed in this document are:

- NONE: TBD
- PSK: TBD

10. Security Considerations

10.1. DNSSEC is recommended to authenticate DNS hosted data

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

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

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

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

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

10.3. CPEs are sensitive to DoS

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

11. Acknowledgment

We would like to thank Tomasz Mrugalski and Bernie Volz for their comments on the design of the DHCP Options.

12. Document Change Log

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

-00: version published in the homenet WG. Major modifications are:

- Reformatting of DHCP Options: Following options guide lines

- DHCPv6 Client behavior: Following options guide lines

- DHCPv6 Server behavior: Following options guide lines

-00: First version published in dhc WG.

13. Pseudo Code

This section is informational. It aims at illustrating how options MAY be handled by the DHCP Client or the DHCP Server. Not all the DHCP Options described in the document are considered. This section is not normative and implementation MAY differ.

13.1. DHCP Zone Public Master Option

13.1.1. Unpacking a DHCP Zone Public Master Option

```
## We consider the following object for the CPE:
## Class cpe
##     self.ip4_list
##     self.ip6_list
##     self.fqdn

receive_dhcp_zone_public_master_option(OPTION_ZONE_PUBLIC_MASTER):
    ## Checking existence of Registered Domains
    if OPTION_REGISTERED_DOMAIN_NAME is empty or missing:
        ignore OPTION_ZONE_PUBLIC_MASTER
    ## Checking existence of Public Authoritative Masters
    if OPTION_ZONE_PUBLIC_MASTER has no OPTION_MASTER options:
        build_option_master(cpe.fqdn, cpe.ip4_list, cpe.ip6_list)

    ## select each DNS Homenet Zone
    for zone_name in OPTION_REGISTERED_DOMAIN_NAME:
        ## adds Public Authoritative Master information to the
        ## zone_name. Typically this may consists in adding the
        ## lines:
        ## zone_name    NS master_fqdn
        ## master_fqdn A ip4
        ## master_fqdn A ip6
        for each OPTION_MASTER:
            (fqdn, ip4_list, ip6_list) = \
                get_master_option_info(OPTION_MASTER)
            add_ns(fqdn, zone_name)
            add_a(fqdn, ip4_list, zone_name)
            add_aaa(fqdn, ip6_list, zone_name)
```

Fig 11: Pseudo code for receiving a DHCP Zone
Public Master Option

13.1.2. Packing a DHCP Zone Public Master Option

The pseudo code for sending a DHCP Zone Public Master Option is presented below. It is informational and intended to illustrate text above. Implementation MAY be different.

```
## We consider the following objects for Public Authoritative Master
## Class master
##     self.ip4_list
##     self.ip6_list
##     self.fqdn

build_dhcp_zone_public_master_option():
    ## Collect all Registered Domain and associate the list of
    ## Public Authoritative Master. One way it to build the
    ## registered_domain_dict = {registered_domain:[master_list]}
    tmp_rd_dict = build_registered_domain_dict()

    ## Remove void registered domain names,
    ## Factorize registered_domain_dict and output
    rd_dict = factorize_registered_domain_dict(tmp_rd_dict)

    ## Build DHCP Zone Public Master Option
    for registered_domain_list, master_list in rd_dict.items():
        ## Builds the DHCP Zone Public Master Option Data Payload
        data = \
            build_registered_domain_name_option(registered_domain_list)
        ## Concatenation with DHCP Public Authoritative Master
        ## Options
        for master in master_list:
            data += build_master_option(master)

    ## Build the DHCP Zone Public Master Option
    return build_header_zone_public_master(data) + data
```

Fig 12: Pseudo code for sending DHCP Zone
Public Master Option

13.1.3. DHCP Master Option

13.1.3.1. Unpacking a DHCP Master Option

```
def get_master_option_info(OPTION_MASTER):
    if OPTION_MASTER is empty:
        ## Consider the CPE for the Public Authoritative Master
        ## or ignore the option
        ## build_option_master(cpe.fqdn, cpe.ip4_list, cpe.ip6_list)
        ignore OPTION_MASTER

    if OPTION_MASTER_FQDN is empty or missing:
        ignore OPTION_MASTER
    if multiple OPTION_MASTER_IP4 or\
       multiple OPTION_MASTER_IP6:
        ignore OPTION_MASTER

    fqdn = get_fqdn(OPTION_MASTER_FQDN)
    ip4_list = []
    ip6_list = []
    ## collect Public Authoritative Master's IP addresses
    if OPTION_MASTER_IP4 exists:
        append(ip4_list, get_ip4(OPTION_MASTER_IP4))
    if OPTION_MASTER_IP6 exists:
        append(ip6_list, get_ip6(OPTION_MASTER_IP6))

    ## if no IP addresses are provided perform a DNSSEC
    ## resolution
    if len(ip4_list) == 0 and len(ip6_list) == 0:
        ip4_list = dig(fqdn, A)
        ip6_list = dig(fqdn, AAAA)
```

Fig 13: Pseudo code for Unpacking DHCP Master Option

13.1.3.2. Packing a DHCP Master Option

The pseudo code illustrates how the DHCP Server builds a DHCP Master Option.

```
def build_master_option(fqdn, ip4_list, ip6_list):
    ## Do not build the data payload of the DHCP Master Option
    ## if fqdn is empty or a null string
    if fqdn is empty or fqdn has more than one fqdn:
        return error
    data = build_master_fqdn_option(fqdn)
    if ip4_list is empty:
        ip4_list = dig(fqdn, A)
    if ip6_list is empty:
        ip6_list = dig(fqdn, AAAA)
    if ip4_list is empty and ip6_list is empty:
        return error
    if ip4_list is not empty:
        data += build_master_ip4_option(ip4_list)
    if ip6_list is not empty:
        data += build_master_ip6_option(ip6_list)
    return build_header_master(data) + data
```

Fig 14: Pseudo code for Packing DHCP Master Option

13.2. DHCP Public Master Upload Option

13.2.1. Unpacking a DHCP Public Master Upload Option

The pseudo code for building the binding between Public Authoritative Master and Secure Channel MAY be as follows:

```
## We consider the master_secure_channel_dict that associates
## for each master a list of Secure Channel
## master_secure_channel_dict = {master: [secure_channel], ..., }

## This function is used to add an entry to the
## master_secure_channel_dict
def get_master_upload_option_info(OPTION_PUBLIC_MASTER_UPLOAD):
    if OPTION_MASTER_FQDN_LIST is not unique or\
       OPTION_SECURE_CHANNEL does not exist:
        ignore OPTION_PUBLIC_MASTER_UPLOAD

    secure_channel_list = []
    for all OPTION_SECURE_CHANNEL:
        sc = get_secure_channel_info(OPTION_SECURE_CHANNEL)
        secure_channel_list.append(sc)

    for each master in OPTION_MASTER_FQDN_LIST:
        master_secure_channel_dict[master] = secure_channel_list
```

Fig 15: Pseudo code for receiving a DHCP Public
Master Upload Option

13.2.2. DHCP Secure Channel Options

13.2.2.1. Unpacking a DHCP Secure Channel Option

```
## We consider the secure_channel object
## Class Secure_channel:
##     self.protocol
##     self.credential
##     self.server_set_ip4
##     self.server_set_ip6

def get_secure_channel_option_info(OPTION_SECURE_CHANNEL):

    if OPTION_SECURE_PROTOCOL is not unique or \
        OPTION_SECURE_CREDENTIAL is not unique or \
        OPTION_SERVER_SET is not unique:
        ignore OPTION_SECURE_CHANNEL

    protocol = \
        get_secure_protocol_option_info(OPTION_SECURE_PROTOCOL)
    credential = \
        get_secure_credential_option_info(OPTION_SECURE_CREDENTIAL)
    ip4, ip6 = \
        get_server_set_option_info(OPTION_SERVER_SET)

    return(Secure_channel(protocol, credential, ip4, ip6))
```

Fig 16: Pseudo code for receiving a DHCP
Secure Channel Option

14. References

14.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen,
"Internet Key Exchange Protocol Version 2 (IKEv2)", RFC
5996, September 2010.
- [RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the
DNS", RFC 6672, June 2012.

14.2. Informational References

- [I-D.ietf-dhc-option-guidelines]
Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and
S. Krishnan, "Guidelines for Creating New DHCPv6 Options",
draft-ietf-dhc-option-guidelines-14 (work in progress),
September 2013.
- [I-D.mglt-homenet-front-end-naming-delegation]
Migault, D., Cloetens, W., Griffiths, C., and R. Weber,
"IPv6 Home Network Naming Delegation", draft-mglt-homenet-
front-end-naming-delegation-02 (work in progress), July
2013.
- [I-D.sury-dnsexst-cname-dname]
Sury, O., "CNAME+DNAME Name Redirection", draft-sury-
dnsexst-cname-dname-00 (work in progress), April 2010.

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

Email: wouter.cloetens@softathome.com

Chris Griffiths
Dyn
150 Dow Street
Manchester, NH 03101
US

Email: cgriffiths@dyn.com
URI: <http://dyn.com>

Ralf Weber
Nominum
2000 Seaport Blvd #400
Redwood City, CA 94063
US

Email: ralf.weber@nominum.com
URI: <http://www.nominum.com>

Homenet
Internet-Draft University of New Hampshire - InterOperability Laboratory
Intended status: Informational
Expires: April 22, 2014

T. Winters, Ed.
October 21, 2013

Service Provider Edge Router Interaction
draft-winters-homenet-sper-interaction-00

Abstract

This document describes the interaction between a Service Provider Gateway fixed at the home edge, and the Home Networking interior routers. Assessing the interactions between existing routers implementing 6204bis and the Home Networking routers. The document will also define the interactions with the HIPnet edge router and Home Networking router.

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 22, 2014.

Copyright Notice

Copyright (c) 2013 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	2
1.1. Requirements Language	2

2.	Terminology	2
3.	6204bis Service Provider Edge Router (SPER)	3
3.1.	Addressing	3
3.2.	Routing	3
3.3.	Security	3
3.4.	Naming and Service Discovery	4
3.5.	Requirements	4
4.	HIPnet Server Provider Edge Router (SPER)	4
4.1.	Addressing	4
4.2.	Routing	4
4.3.	Security	4
4.4.	Naming and Service Discovery	4
4.5.	Requirements	5
5.	Security Considerations	5
6.	IANA Considerations	5
7.	Acknowledgements	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	6
	Author's Address	6

1. Introduction

This document defines the interactions between the future Homenet network and 6204bis Routers and HIPnet routers. Currently 6204bis and HIPnet routers are being deployed by ISPs as Service Provider Edge Routers (SPER) for IPv6 deployment. In the future the SPER will be a full Homenet routers but there will be a period of transition. This document specifies how currently deployed SPER will interact with Homenet architecture [I-D.ietf-homenet-arch]. Identifying the effects of key components of the architecture. The goal of this document is to make recommendations on issues uncovered to make the devices work with the future Homenet.

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 [RFC2119].

2. Terminology

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

- o Border: a point, typically resident on a router, between two networks, e.g. between the main internal homenet and a guest network. This defines point(s) at which filtering and forwarding policies for different types of traffic may be applied
- o SPER: Service Provider Edge Router: A border router intended for use in a homenet, which connects the homenet to a service provider network.

- o Homenet: A home network, comprising host and router equipment, with one or more SPERs providing connectivity to service provider network(s).
- o Internet Service Provider (ISP): an entity that provides access to the Internet. In this document, a service provider specifically offers Internet access using IPv6, and may also offer IPv4 Internet access. The service provider can provide such access over a variety of different transport methods such as DSL, cable, wireless, and others.
- o Home IP network HIPnet: A router intended for home or small-office use that forwards packet explicitly address to itself as defined in [I-D.grundemann-homenet-hipnet].
- o 6204bis: A router intended for home or small-office use that forwards packet explicitly address to itself as defined in [I-D.ietf-v6ops-6204bis]

3. 6204bis Service Provider Edge Router (SPER)

3.1. Addressing

A 6204bis SPER acquire LAN addressing thru DHCP Prefix Delegation [RFC3633]. A 6204bis SPER will assign a separate /64 from the delegated prefix(es) for each LAN interface. The 6202bis SPER can assign address either thru SLAAC or DHCP to LAN host. There is no mechanism for delegating any unused prefix(es) that were delegated to the 6204bis Router.

3.2. Routing

A 6204bis SPER is capable of learning default routes thru Router Advertisement on the WAN interface. All other routing information is learned when a prefix is assigned to a LAN interface by the 6204bis SPER. All other Routing information is learned from other methods such as user configuration.

The 6204bis SPER will NOT forward packets from an unrecognized source address. Any IPv6 packets routed from the Homenet would receive an ICMPv6 Destination Unreachable message. A 6204bis SPER prevents any Routing mechanisms from working.

3.3. Security

A 6204 SPER is recommended to enable a stateful [RFC6092] firewall by default. This stateful firewall will allow the homenet incoming traffic is limited to return traffic resulting from outgoing packets.

A Homenet Router with the firewall on might not allow valid traffic from devices connected to the 6204bis SPER. When a Homenet Router detects a 6204bis SPER it should allow native IPv6 traffic thru the firewall so that traffic can flow between the 6204bis SPER and Homenet.

3.4. Naming and Service Discovery

A 6204bis SPER support the ability to assign DNS servers and Domain Search List. A 6204bis SPER may support service discovery protocols such as mDNS/DNS-SD and SSDP. These protocols will only work when directly connected to the 6204 SPER LAN interfaces. There is currently no mechanism for sharing service discovery between the Homenet and IPv6 nodes connected to the LAN interfaces of the 6204 SPER.

3.5. Requirements

TBD

4. HIPnet Server Provider Edge Router (SPER)

4.1. Addressing

HIPnet SPER use DHCPv6 prefix sub-delegation to build the network [I-D.grundemann-homenet-hipnet]. The HIPnet SPER receives IPv6 prefix per [I-D.ietf-v6ops-6204bis]. If the prefix is larger than a single /64 prefix the HIPnet router will subdivide the IPv6 prefix received via DHCPv6 [RFC3315]. Using Recursive Prefix Delegation allows the Homenet to receive prefixes that can be used to address the network.

4.2. Routing

Leveraging the recursive prefix delegation method described above, a HIPnet SPER installs route to the WAN interface of the router delegated the prefixes. With this routing information the HIPnet SPER is able to properly route packets to and from the Homenet.

4.3. Security

A HIPnet SPER must enable a stateful [RFC6092] firewall by default. This stateful firewall will allow the homenet incoming traffic is limited to return traffic resulting from outgoing packets.

A Homenet Router with the firewall on might not allow valid traffic from devices connected to the HIPnet SPER. When a Homenet Router detects a HIPnet SPER it should allow native IPv6 traffic thru the firewall so that traffic can flow between the HIPnet SPER and Homenet.

4.4. Naming and Service Discovery

Both the Homenet and HIPnet have several common protocols that can be used for service discovery such as mDNS [RFC6762], DNS-SD [RFC6763], and SSDP. Both the HIPnet and Homenet Routers may have host directly connected that are using them as DNS servers. If the HIPnet SPER advertises itself as the DNS-SD server for connected host, the host could query the HIPnet SPER. The issue that arises with this configuration is the HIPnet Router currently has no method for finding the Homenet router to query when trying to resolve DNS. Future revisions of this draft should try and resolve this issue.

4.5. Requirements

TBD

5. Security Considerations

6. IANA Considerations

This document makes no request of IANA.

7. Acknowledgements

TBD

8. References

8.1. Normative References

[I-D.grundemann-homenet-hipnet]

Grundemann, C., Donley, C., Brzozowski, J., Howard, L. and V. Kuarsingh, "A Near Term Solution for Home IP Networking (HIPnet)", Internet-Draft draft-grundemann-homenet-hipnet-01, February 2013.

[I-D.ietf-homenet-arch]

Chown, T., Arkko, J., Brandt, A., Troan, O. and J. Weil, "Home Networking Architecture for IPv6", Internet-Draft draft-ietf-homenet-arch-10, August 2013.

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

Singh, H., Beebee, W., Donley, C. and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", Internet-Draft draft-ietf-v6ops-6204bis-12, October 2012.

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

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C. and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.
- [RFC6204] Singh, H., Beebee, W., Donley, C., Stark, B. and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, April 2011.

8.2. Informative References

- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, February 2013.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, February 2013.

Author's Address

Timothy Winters, editor
University of New Hampshire - InterOperability Laboratory
Durham, NH

Email: twinters@iol.unh.edu

Network Working Group
Internet-Draft
Expires: April 24, 2014

M. Xu
S. Yang
J. Wu
Tsinghua University
F. Baker
Cisco Systems
October 21, 2013

Traffic Class Routing Protocol in Home Networks
draft-xu-homenet-traffic-class-01

Abstract

Home IT staff is generally unfamiliar with network operations, making it desirable to provide a configuration-free mode of operation. Policy-based routing (in the sense of configuring one router to redirect traffic to another based on access control) and multi-topology routing both require configuration, making them undesirable.

In this document, we propose a configuration-free mechanism, in which packets will be routed towards the corresponding upstream ISPs based on both destination and source addresses.

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 24, 2014.

Copyright Notice

Copyright (c) 2013 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. Terminology	3
3. Overview	3
4. Router Behavior	5
4.1. Egress Router Behavior	5
4.2. Interior Router Behavior	5
5. TC-LSA Format	6
6. Routing Table Structure	7
7. Calculation of the Routing Table	8
8. Matching Rule	9
9. Forwarding Table Structure	9
10. Compatibility	9
11. IANA Considerations	10
12. Acknowledgments	10
13. References	10
13.1. Normative References	10
13.2. Informative References	10
Authors' Addresses	11

1. Introduction

Home networks are growing both in device count and in complexity. Today they generally contain both wired and wireless components, and may require routing to place audio/visual entertainment traffic on one path, office services on another, and wireless LANs (both IEEE-style and 4G/LTE-style) on a third. Traditionally, we have

simplified networks using a single exit router and a default route. Today, we might have multiple routers to wired upstream networks, and by separate paths LTE services, "Smart Grid" services, or health network services. Increasingly, such networks are multihomed, and multihomed using diverse access network technologies.

Traditionally, routing protocols make routing decisions solely based on destination IP addresses, packets towards the same destination will be delivered to the same next hop no matter where they come from. These protocols work well with simple home networks that have only one egress router. However, in the multi-homing scenario, packets may be dropped if forwarded only based on destination addresses [RFC3704].

Although many patch-like solutions, like policy-based routing (PBR), multi-topology routing (MTR) and layer-3 VPN can solve the problem, they complex the configurations in home networks, and are not suitable for home IT staffs. We need a configuration-free solution to help operators set up their home networks in the multi-homing scenario.

In this document, we propose a configuration-free mechanism - traffic class routing, based on OSPFv3, such that home networks can route packets towards the corresponding upstream ISPs, according to both destination and source addresses.

2. Terminology

Terminology used in this document:

- o Traffic Class (TC): Identified by (destination prefix, source prefix), all packets falling in the domain belong to the traffic class.
- o TC-Route: Identified by (destination prefix, source prefix, value), where value is the administrative value applied to the traffic class (destination prefix, source prefix).
- o TC-LSA: Link state advertisement that communicates the reachability for a traffic classes.

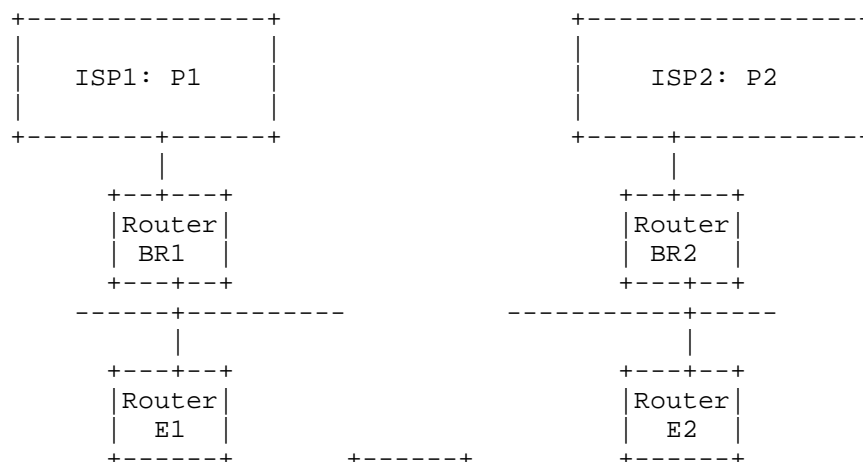
3. Overview

In a home network, traditionally, egress routers obtain delegated prefixes from upstream ISPs using DHCPv6 with prefix options [RFC3633]. The egress routers will then assign longer sub-prefixes to the other links in the home network. Each router inside the home network will act as standard OSPFv3 router, and forward packets based on their destination addresses.

With traffic class routing, after obtaining delegated prefixes and assigning sub-prefixes, egress routers will populate traffic classes (with extended LSAs), rather than destination address only, into the home network. Each router inside the home network will flood these traffic classes information. When calculating the path towards a destination address, routers will take the traffic classes into considerations. Intrinsically, in traditional routing model, the object being routed to is a destination prefix; in our routing model, the object being routed might be a destination prefix given that the packet sports a certain source prefix.

Each traffic class is associated with a cost, which is a single dimensionless metric.

For example, a site is connected to the Internet through two ISPs, ISP1 and ISP2. ISP1 delegates prefix P1 to the site, and ISP2 delegates prefix P2 to the site. After being delegated with P1, the egress router E1 of the site will advertise a traffic class - $\{::/0, P1\}$, into the site. After being delegated with P2, the egress router E2 of the site will advertise a traffic class - $\{::/0, P2\}$, into the site. Receiving these advertisements, interior router I1 will compute two paths towards $::/0$, one through router E1 for traffic from P1, the other through E2 for traffic from P2.



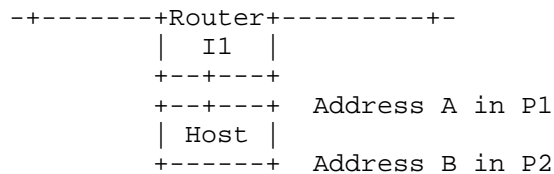


Figure 1: Multi-homing Scenario in Home Networks

4. Router Behavior

All routers behave like traditional OSPFv3 routers, however, the following behaviors are different with traditional OSPFv3 routers.

4.1. Egress Router Behavior

After obtaining delegated prefixes using DHCPv6 with prefix options, an egress router should originate TC-LSAs, i.e., extended LSAs with source prefixes appended. Egress routers then will advertise these TC-LSAs into the home network.

Note that an egress router behaves like an interior router if it receives a TC-LSA from other egress routers.

4.2. Interior Router Behavior

Receiving TC-LSAs from egress routers, an interior router should store the TC-LSAs into its LSDB, and flood it to other routers. After calculating a path to an egress router advertising reachability, i.e., a destination prefix, the interior router should decide which traffic class can follow this path towards the egress router. If a traffic class can travel through two different paths, then interior router should compare their costs, and select the path with the lowest cost.

Interior routers contains a routing table that contains all necessary information to forward an IP packet following the path of a traffic class. After computing the path towards a traffic class, interior routers should update the entry in the routing table if necessary, e.g., change the next hop towards the traffic class. The routing table structure will be described in Section 6. Calculation of routing table will be illustrated in Section 7.

At last, interior routers should update the Forwarding Information Base (FIB), which will be discussed in the next version of this document.

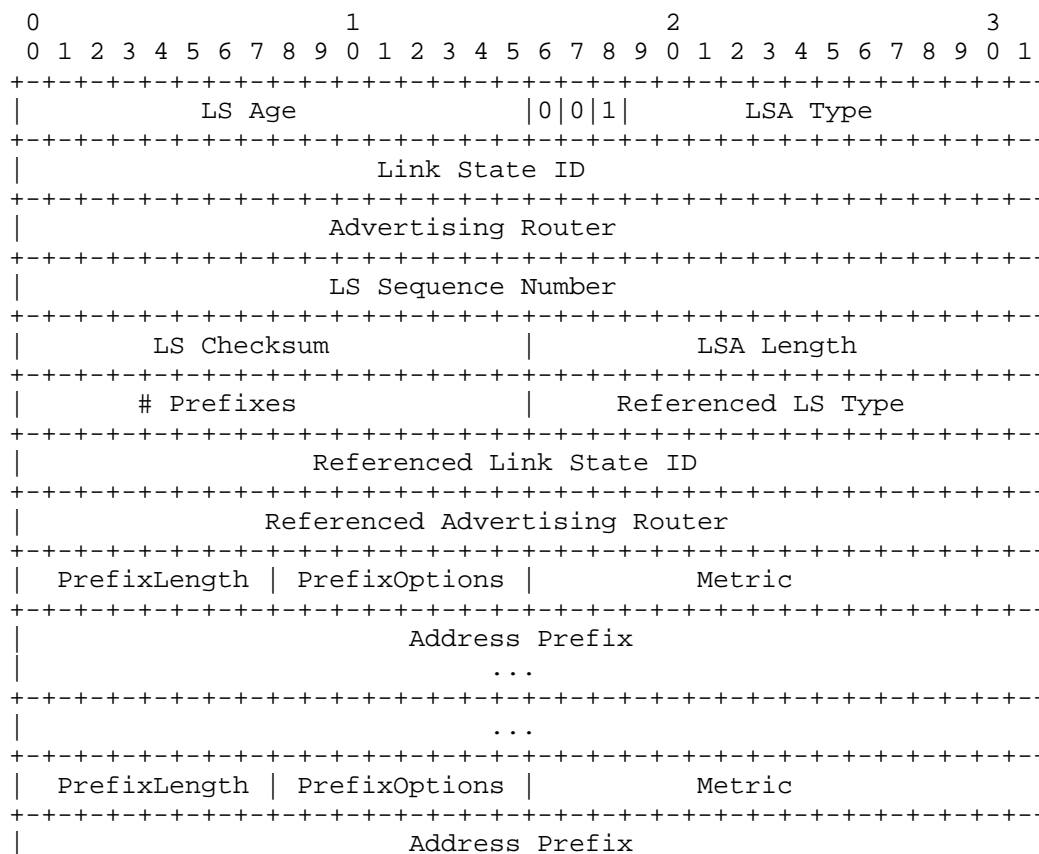
5. TC-LSA Format

TC-LSA adds TLV extensions, which contains source prefix information, based on original OSPFv3 LSA. We follow the TLV format in [I-D.baker-ipv6-ospf-dst-src-routing] and extended LSA format in [I-D.acee-ospfv3-lsa-extend].

Each extended LSA includes the traditional LSA part in [RFC5340], and one or more TLVs defined in [I-D.baker-ipv6-ospf-dst-src-routing]. But we do not need all LSAs to be extended, the LSAs need to be extended are as follows:

- o Intra-Area-Prefix-LSA: The extended LSA has type 0xA029.

The extended LSA format for Intra-Area-Prefix-LSA in multi-homing is as follows:



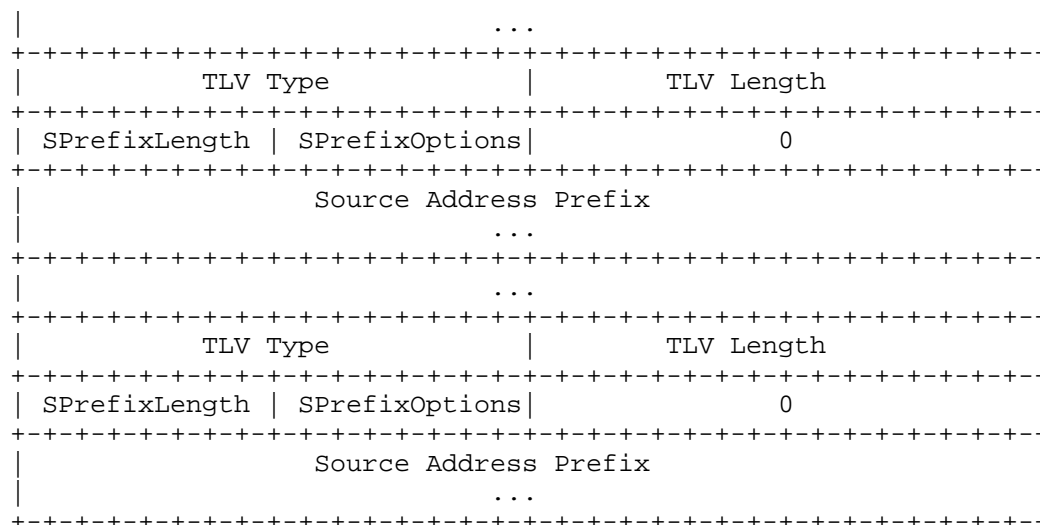


Figure 1: Multi-homing Scenario in Home Networks

All LSA header fields are the same as defined in [RFC5340], except the following:

- o LSA type: The LSA type value is 0xA029, according to [I-D.acee-ospfv3-lsa-extend];
- o LSA length: The length of the whole LSA header, including the TLVs;
- o TLV type: The type of IPv6 source prefix TLV, assigned by IANA;
- o TLV length: The value is 20 as defined in [I-D.baker-ipv6-ospf-dst-src-routing];
- o SPrefixLength, SPrefixOptions, Source Address Prefix: Representation of the IPv6 address prefix, which is delegated from the upstream ISP providers;

For simplicity, each extended LSA should only carry one source prefix, suppose there are n destination prefix d_1, d_2, \dots, d_n , and the source prefix is s , then the LSA carries n TC-route announcement, $(d_1, s, v_1), (d_1, s, v_2), \dots, (d_n, s, v_n)$, where v_i is the metric associated with destination prefix d_i .

6. Routing Table Structure

For traditional routing, the routing table structure contains all needed information to forward IP packets to the right destination. For example, destination prefixes are commonly structured into a prefix trie, where each trie nodes contain the necessary information. Routers can lookup and update the prefix trie.

With traffic classes, the routing table structure must contain all needed information to forward IP packets following the right traffic class, i.e., towards the related destination and from the related source. For each routing table entry, there are two additional fields other than the fields mentioned in [RFC5340]:

- o Source IP Address: The IP address of the source in traffic class.
- o Source Address Mask: If the source is a subnet, then it is referred to as the subnet mask.

The routing table must provide interface for update and lookup in it. For example, traffic classes can be structured into a two dimensional (or two level) trie, where each trie node in the first dimension points to a sub-trie in the second dimension. The trie nodes in the second dimension contain the necessary information to forward IP packets following the right traffic class.

7. Calculation of the Routing Table

The fundamental algorithm in OSPFv3 doesn't change. The algorithm uses the SPF approach to calculate a path to the router advertising reachability, and then uses the reachability advertisement to decide what traffic should follow that route. What we are changing is the reachability advertisement, in traditional OSPFv3, the advertisements, which is one or several kinds of LSAs, represent destination prefixes; in this document, the advertisements, which is one or several kinds of TC-LSAs, represent traffic classes.

Note that we do not have to change router-LSA and network-LSA in [RFC5340]. Thus, the first stage of Section 4.8.1 in [RFC5340] remains the same in this document. However, the second stage of Section 4.8.1 in [RFC5340] should change by a little bit. Instead of examining the list of the intra-area-prefix-LSAs, the list of extended intra-area-prefix-LSAs is examined. The cost of any advertised traffic class is the sum of the class' advertised metric plus the cost of the transit vertex (either router or transit network) identified by extended intra-area-prefix-LSAs' referenced LS type, referenced link state ID, and referenced advertising router field.

8. Matching Rule

We also adopt the LMF (longest match first) rule when a packet matches multiple routing entries. However, traffic class has two dimensions, there might exist ambiguity. For example, if there exists two routing entries, (d1, s1, nexthop1), (d2, s2, nexthop2), where d1 is longer than d2 and s2 is longer than s1, then none entry is longer than the other in both dimensions. In this situation, we must insert an additional entry into the routing table, e.g., (d1, s2, nexthop1) in the above example. The entry directs to nexthop1 rather than nexthop2, because we must guarantee reachability according to the destination prefix.

9. Forwarding Table Structure

As the format of routing table entries changes from (destination prefix, nexthop) to (destination prefix, source prefix, nexthop), The Forwarding Information Base (FIB), based on which routers forward IP packets, should be redesigned. There can be different possible FIB structures, we use a trie-based solution in this document.

In traditional routing, the destination prefixes are constructed using a prefix trie. However, in traffic class routing, the forwarding decision is based on both the destination and source prefixes. To store them, we design a new structure, based on partricia-trie to meet the needs.

In traffic class routing, the router forwards IP packets following the right traffic class based on destination and source prefixes. Routers construct a two dimensional (or two level) partricia-trie, where each trie node in the first dimension points to a sub-trie (the sub-trie can be empty). When a packet arrives, the packet should match a trie node (destination prefix) in the first dimension using the destination address. Following the sub-trie under the node, the packet should match another trie node (source prefix) in the second dimension, i.e., the sub-trie.

To implement the FIB structure, we modified Click, which is a modular software router. The FIB structure could support lookup and update. When the routing table changes, routers can update the FIB through pre-defined interfaces.

10. Compatibility

Routers can also announce the traditional destination-based LSAs at the same time. In this case, we have two choices. In the first choice, we treat the destination-based LSAs as TC-LSAs where the source prefix equal the wildcard. In the second choice, routers have

to keep two routing tables, one for destination prefix only, and the other for traffic classes. When a packet arrives, routers first lookup in the routing table storing traffic classes; If none entry matches, then routers lookup in the routing table storing destination prefixes.

11. IANA Considerations

The newly LSA types and TLVs should be assigned by IANA, please refer to [I-D.baker-ipv6-ospf-dst-src-routing] and [I-D.acee-ospfv3-lsa-extend].

12. Acknowledgments

Zheng Liu and Gautier Bayzelon provided useful input into this document.

13. References

13.1. Normative References

- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

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-07 (work in progress), February 2013.
- [I-D.baker-ipv6-ospf-dst-src-routing]
Baker, F., "IPv6 Source/Destination Routing using OSPFv3", draft-baker-ipv6-ospf-dst-src-routing-02 (work in progress), May 2013.
- [I-D.acee-ospfv3-lsa-extend]
Lindem, A., Mirtorabi, S., Roy, A., and F. Baker, "OSPFv3 LSA Extendibility", draft-acee-ospfv3-lsa-extend-00 (work in progress), May 2013.

[I-D.baker-fun-routing-class]

Baker, F., "Routing a Traffic Class", draft-baker-fun-routing-class-00 (work in progress), July 2011.

Authors' Addresses

Mingwei Xu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6278-5822
Email: xmw@csnet1.cs.tsinghua.edu.cn

Shu Yang
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6278-5822
Email: yangshu@csnet1.cs.tsinghua.edu.cn

Jianping Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6278-5983
Email: jianping@cernet.edu.cn

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

Email: fred@cisco.com