

Homenet
Internet-Draft
Intended status: Standards Track
Expires: November 14, 2021

D. Migault
Ericsson
R. Weber
Nominum
M. Richardson
Sandelman Software Works
R. Hunter
Globis Consulting BV
May 13, 2021

Simple Provisioning of Public Names for Residential Networks
draft-ietf-homenet-front-end-naming-delegation-15

Abstract

Home owners often have IPv6 devices that they wish to access over the Internet using names.

Outsourcing the DNS servers to a DNS infrastructure protects against possible DDoS attacks as well as sudden renumbering of the home network. It has been possible to register and populate a DNS Zone with names since DNS became a thing, but it has been an activity typically reserved for experts. This document automates the process through creation of a Homenet Naming Authority (HNA), whose responsibility is to select, sign and publish names to a set of publicly visible servers.

This document describes the mechanism that enables the HNA to outsource the naming service to the DNS Outsourcing Infrastructure (DOI) via a Distribution Manager (DM).

In addition, this document deals with publication of a corresponding reverse zone.

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 <https://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 November 14, 2021.

Copyright Notice

Copyright (c) 2021 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Selecting Names to Publish	5
1.2. Alternative solutions	6
2. Terminology	7
3. Architecture Description	8
3.1. Architecture Overview	8
3.2. Distribution Manager Communication Channels	11
4. Control Channel	12
4.1. Information to Build the Public Homenet Zone	13
4.2. Information to build the DNSSEC chain of trust	13
4.3. Information to set the Synchronization Channel	14
4.4. Deleting the delegation	14
4.5. Messages Exchange Description	14
4.5.1. Retrieving information for the Public Homenet Zone.	14
4.5.2. Providing information for the DNSSEC chain of trust	15
4.5.3. Providing information for the Synchronization Channel	16
4.5.4. HNA instructing deleting the delegation	17
4.6. Securing the Control Channel	17
4.7. Implementation Concerns	18
5. Synchronization Channel	19
5.1. Securing the Synchronization Channel	20
6. DM Distribution Channel	20
7. HNA Security Policies	20
8. DNSSEC compliant Homenet Architecture	21
9. Homenet Reverse Zone Channels Configuration	21
10. Homenet Public Zone Channel Configurations	22
11. Renumbering	24
11.1. Hidden Primary	24
12. Privacy Considerations	25

13. Security Considerations	26
13.1. HNA DM channels	26
13.2. Names are less secure than IP addresses	26
13.3. Names are less volatile than IP addresses	27
14. Information Model for Outsourced information	27
14.1. Outsourced Information Model	28
15. IANA Considerations	30
16. Acknowledgment	30
17. Contributors	30
18. References	31
18.1. Normative References	31
18.2. Informative References	34
Appendix A. Envisioned deployment scenarios	36
A.1. CPE Vendor	36
A.2. Agnostic CPE	36
Appendix B. Example: A manufacturer provisioned HNA product flow	37
Authors' Addresses	38

1. Introduction

The Homenet Naming Authority (HNA) is responsible for making devices within the home network accessible by a public name within the home network as well as from outside the home network (e.g. the Internet). IPv6 connectivity provides the possibility of global end to end IP connectivity.

The use of a DNS zone for each home network is a reasonable and scalable way to make the set of public names visible. There are a number of ways to populate such a zone. This specification proposes a way based on a number of assumptions about typical home networks.

1. The names of the devices accessible from the Internet are stored in the Public Homenet Zone, served by a DNS authoritative server.
2. It is unlikely that home networks will contain sufficiently robust platforms designed to host and expose to the Internet a service such as the DNS and as such would expose the home network to DDoS attacks.
3. [RFC7368] emphasizes that the home network is subject to connectivity disruptions with the ISP. But, names used within the home MUST be resilient against such disruption.

This specification makes the public names resolvable within both the home network and on the Internet, even when there are disruptions.

This is achieved by having a function inside the home network that builds, signs, publishes, and manages a Public Homenet Zone, thus

providing bindings between public names, IP addresses, and other RR types.

The management of the names can be under the responsibility the Customer Premises Equipment (CPE) does. Other devices within the home network could fulfill this role e.g. a Network Attached Storage server, but for simplicity, this document assumes the function is located on one of the CPE devices.

The homenet architecture [RFC7368] makes it clear that a home network may have multiple CPEs. The management of the Public Homenet Zone involves DNS specific mechanisms that cannot be distributed over multiple servers (primary server), when multiple nodes can potentially manage the Public Homenet Zone, a single node needs to be selected per outsourced zone. This selected node is designated as providing the HNA function.

The process by which a single HNA is selected per zone is not in scope for this document. It is envisioned that a future document will describe an HNCP mechanism to elect the single HNA.

CPEs, which may host the HNA function are usually low powered devices not designed for supporting a heavy traffic. As a result, hosting an authoritative DNS service visible to the Internet may expose the home network to resource exhaustion and other attacks. On the other hand, if the only copy of the public zone is on the Internet, then Internet connectivity disruptions would make the names unavailable inside the homenet.

In order to avoid resource exhaustion and other attacks, this document describes in Section 3.1 an architecture that outsources the authoritative naming service of the home network. More specifically, the HNA builds the Public Homenet Zone and outsources it to a DNS Outsourcing Infrastructure (DOI) via a Distribution Manager (DM). The DOI is in charge of publishing the corresponding Public Homenet Zone on the Internet. The transfer of DNS zone information is achieved using standard DNS mechanisms involving primary and secondary DNS servers, with the HNA being a stealth primary, and the DM a secondary.

In order to keep the Public Homenet Zone up-to-date Section 5 describes how the HNA and the DOI synchronize the Public Homenet Zone.

The architecture is explicitly designed to enable fully functional DNSSEC, and the Public Homenet Zone is expected to be signed with a secure delegation. DNSSEC key management and zone signing are handled by the HNA.

Section 10 discusses management and configuration of the Public Homenet Zone. It shows that the HNA configuration of the DOI can involve no or little interaction with the end user. More specifically, it shows that the existence of an account in the DOI is sufficient for the DOI to push the necessary configuration. In addition, when the DOI and CPE are both managed by an ISP, the configuration can be entirely automated - see Section 9.

Section 9 discusses management of one or more reverse zones. It shows that management of the reverse zones can be entirely automated and benefit from a pre-established relation between the ISP and the home network. Note that such scenarios may also be met for the Public Homenet Zone.

Section 11 discusses how renumbering should be handled. Finally, Section 12 and Section 13 respectively discuss privacy and security considerations when outsourcing the Public Homenet Zone.

The Public Homenet Zone is expected to contain public information only in a single universal view. This document does not define how the information required to construct this view is derived.

It is also not in the scope of this document to define names for exclusive use within the boundaries of the local home network. Instead, local scope information is expected to be provided to the home network using local scope naming services. mDNS [RFC6762] and DNS-SD [RFC6763] are two examples of these services. Currently mDNS is limited to a single link network. However, future protocols and architectures [I-D.ietf-homenet-simple-naming] are expected to leverage this constraint as pointed out in [RFC7558].

1.1. Selecting Names to Publish

While this document does not create any normative mechanism by which the selection of names to publish, this document anticipates that the home network administrator (a human), will be presented with a list of current names and addresses present on the inside of the home network.

The administrator would mark which devices (by name), are to be published. The HNA would then collect the IPv6 address(es) associated with that device, and put the name into the Public Homenet Zone. The address of the device can be collected from a number of places: mDNS [RFC6762], DHCP [RFC6644], UPnP, PCP [RFC6887], or manual configuration.

A device may have a Global Unicast Address (GUA), a Unique Local IPv6 Address (ULA), as well as IPv6-Link-Local addresses, IPv4-Link-Local

Addresses, and RFC1918 addresses. Of these the link-local are never useful for the Public Zone, and should be omitted. The IPv6 ULA and the RFC1918 addresses may be useful to publish, if the home network environment features a VPN that would allow the home owner to reach the network.

The IPv6 ULA addressees are significantly safer to publish, as the RFC1918 addressees are likely to be confusing to any other entity.

In general, one expects the GUA to be the default address to be published. However, during periods when the home network has connectivity problems, the ULA and RFC1918 addressees can be used inside the home, and the mapping from public name to locally useful location address would permit many services secured with HTTPS to continue to operate.

1.2. Alternative solutions

An alternative existing solution in IPv4 is to have a single zone, where a host uses a RESTful HTTP service to register a single name into a common public zone. This is often called "Dynamic DNS", and there are a number of commercial providers, including Dyn, Gandi etc. These solutions were typically used by a host behind the CPE to make its CPE IPv4 address visible, usually in order to enable incoming connections. This is the most common scenario considered in this section, while some variant may also consider the client being hosted in the CPE.

For a very few number (one to three) of hosts, the use of such a system provides an alternative to the architecture described in this document. The alternative - even adapted to IPv6 and ignoring those associated to an IPv4 development - does suffer from some severe limitations:

- o the CPE/HNA router is unaware of the process, and cannot respond to queries for these names when there are disruptions in connectivity. This makes the home user or application dependent on having to resolve different names in the event of outages or disruptions.
- o the CPE/HNA router cannot control the process. Any host can do this regardless of whether or not the home network administrator wants the name published or not. There is therefore no possible audit trail.
- o the credentials for the dynamic DNS server need to be securely transferred to all hosts that wish to use it. This is not a problem for a technical user to do with one or two hosts, but it

does not scale to multiple hosts and becomes a problem for non-technical users.

- o "all the good names are taken" - current services put everyone's names into some small set of zones, and there are often conflicts. Distinguishing similar names by delegation of zones was among the primary design goals of the DNS system.
- o The RESTful services do not always support all RR types. The homenet user is dependent on the service provider supporting new types. By providing full DNS delegation, this document enables all RR types and also future extensions.

There is no technical reason why a RESTful service could not provide solutions to many of these problems, but this document describes a DNS-based solution.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Customer Premises Equipment: (CPE) is a router providing connectivity to the home network.

Homenet Zone: is the DNS zone for use within the boundaries of the home network: 'home.arpa' (see [RFC8375]). This zone is not considered public and is out of scope for this document.

Registered Homenet Domain: is the domain name that is associated with the home network.

Public Homenet Zone: contains the names in the home network that are expected to be publicly resolvable on the Internet.

Homenet Naming Authority(HNA): is a function responsible for managing the Public Homenet Zone. This includes populating the Public Homenet Zone, signing the zone for DNSSEC, as well as managing the distribution of that Homenet Zone to the DNS Outsourcing Infrastructure (DOI).

DNS Outsourcing Infrastructure (DOI): is the infrastructure responsible for receiving the Public Homenet Zone and publishing it on the Internet. It is mainly composed of a Distribution Manager and Public Authoritative Servers.

Public Authoritative Servers: are the authoritative name servers for the Public Homenet Zone. Name resolution requests for the Homenet Domain are sent to these servers. For resiliency the Public Homenet Zone SHOULD be hosted on multiple servers.

Homenet Authoritative Servers: are authoritative name servers within the Homenet network.

Distribution Manager (DM): is the (set of) server(s) to which the HNA synchronizes the Public Homenet Zone, and which then distributes the relevant information to the Public Authoritative Servers.

Homenet Reverse Zone: The reverse zone file associated with the Public Homenet Zone.

Reverse Public Authoritative Servers: equivalent to Public Authoritative Servers specifically for reverse resolution.

Reverse Distribution Manager: equivalent to Distribution Manager specifically for reverse resolution.

Homenet DNSSEC Resolver: a resolver that performs a DNSSEC resolution on the home network for the Public Homenet Zone. The resolution is performed requesting the Homenet Authoritative Servers.

DNSSEC Resolver: a resolver that performs a DNSSEC resolution on the Internet for the Public Homenet Zone. The resolution is performed requesting the Public Authoritative Servers.

3. Architecture Description

This section provides an overview of the architecture for outsourcing the authoritative naming service from the HNA to the DOI. Note that Section 14 defines necessary parameter to configure the HNA.

3.1. Architecture Overview

Figure 1 illustrates the architecture where the HNA outsources the publication of the Public Homenet Zone to the DOI.

The Public Homenet Zone is identified by the Registered Homenet Domain Name - myhome.example. The ".local" as well as ".home.arpa" are explicitly not considered as Public Homenet zones and represented as Homenet Zone in Figure 1.

The HNA SHOULD build the Public Homenet Zone in a single view populated with all resource records that are expected to be published on the Internet. The HNA also signs the Public Homenet Zone. The HNA handles all operations and keying material required for DNSSEC, so there is no provision made in this architecture for transferring private DNSSEC related keying material between the HNA and the DM.

Once the Public Homenet Zone has been built, the HNA outsources it to the DOI as described in Figure 1. The HNA acts as a hidden primary [RFC8499] while the DM behaves as a secondary responsible to distribute the Public Homenet Zone to the multiple Public Authoritative Servers that DOI is responsible for. The DM has three communication channels:

- o a DM Control Channel (Section 4) to configure the HNA and the DOI,
- o a DM Synchronization Channel (Section 5) to synchronize the Public Homenet Zone on the HNA and on the DM,
- o one or more Distribution Channels (Section 6) that distribute the Public Homenet Zone from the DM to the Public Authoritative Server serving the Public Homenet Zone on the Internet.

There might be multiple DM's, and multiple servers per DM. This document assumes a single DM server for simplicity, but there is no reason why each channel needs to be implemented on the same server or use the same code base.

It is important to note that while the HNA is configured as an authoritative server, it is not expected to answer to DNS requests from the public Internet for the Public Homenet Zone. More specifically, the addresses associated with the HNA SHOULD NOT be mentioned in the NS records of the Public Homenet zone, unless additional security provisions necessary to protect the HNA from external attack have been taken.

The DOI is also responsible for ensuring the DS record has been updated in the parent zone.

Resolution is performed by the DNSSEC resolvers. When the resolution is performed outside the home network, the DNSSEC Resolver resolves the DS record on the Global DNS and the name associated to the Public Homenet Zone (myhome.example) on the Public Authoritative Servers.

When the resolution is performed from within the home network, the Homenet DNSSEC Resolver MAY proceed similarly. On the other hand, to provide resilience to the Public Homenet Zone in case of WAN connectivity disruption, the Homenet DNSSEC Resolver SHOULD be able

to perform the resolution on the Homenet Authoritative Servers. These servers are not expected to be mentioned in the Public Homenet Zone, nor to be accessible from the Internet. As such their information as well as the corresponding signed DS record MAY be provided by the HNA to the Homenet DNSSEC Resolvers, e.g., using HNCP [RFC7788]. Such configuration is outside the scope of this document. Since the scope of the Homenet Authoritative Servers is limited to the home network, these servers are expected to serve the Homenet Zone as represented in Figure 1.

How the Homenet Authoritative Servers are provisioned is also out of scope of this specification. It could be implemented using primary and secondary servers, or via rsync. In some cases, the HNA and Homenet Authoritative Servers may be combined together which would result in a common instantiation of an authoritative server on the WAN and inner homenet interface. Note that [RFC6092] REC-8 states this must not be the default configuration. Other mechanisms may also be used.

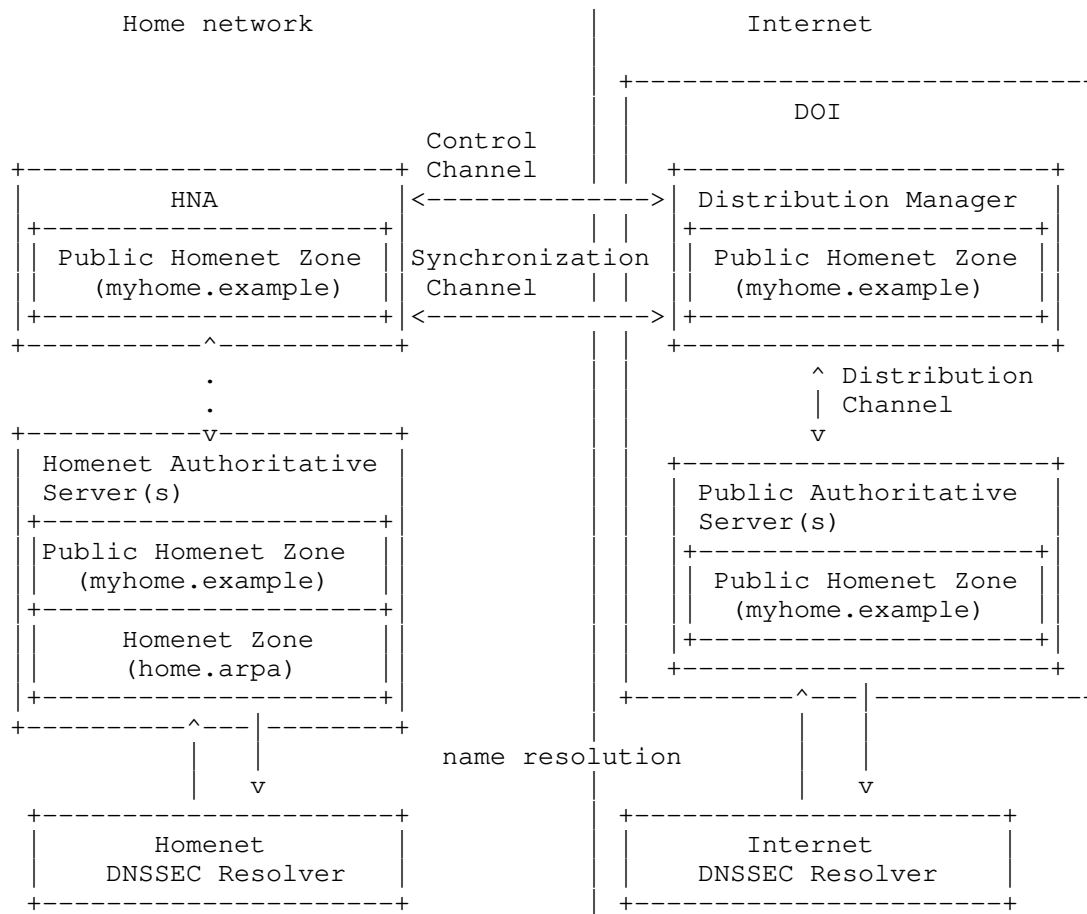


Figure 1: Homenet Naming Architecture

3.2. Distribution Manager Communication Channels

This section details the DM channels, that is the Control Channel, the Synchronization Channel and the Distribution Channel.

The Control Channel and the Synchronization Channel are the interfaces used between the HNA and the DOI. The entity within the DOI responsible to handle these communications is the DM and communications between the HNA and the DM MUST be protected and mutually authenticated. While Section 4.6 discusses in more depth the different security protocols that could be used to secure, it is RECOMMENDED to use TLS with mutually authentication based on certificates to secure the channel between the HNA and the DM.

The Control Channel is used to set up the Synchronization Channel. We assume that the HNA initiates the Control Channel connection with the DM and as such has a prior knowledge of the DM identity (X509 certificate), the IP address and port number to use and protocol to set secure session. We also assume the DM has knowledge of the identity of the HNA (X509 certificate) as well as the Registered Homenet Domain. For more detail to see how this can be achieved, please see Section 10.

The information exchanged between the HNA and the DM uses DNS messages protected by DNS over TLS (DoT) [RFC7858]. Other specifications may consider protecting DNS messages with other transport layers, among others, DNS over DTLS [RFC8094], or DNS over HTTPs (DoH) [RFC8484] or DNS over QUIC [I-D.ietf-dprive-dnsquic]. There was consideration to using a standard TSIG [RFC2845] or SIG(0) [RFC2931] to perform a dynamic DNS update to the DM. There are a number of issues with this. The first one is that TSIG or SIG(0) make scenarios where the end user needs to interact via its web browser more complex. More precisely, authorization and access control granted via OAUTH would be unnecessarily complex with TSIG or SIG(0).

The main issue is that the Dynamic DNS update would also update the parent zone's (NS, DS and associated A or AAAA records) while the goal is to update the DM configuration files. The visible NS records SHOULD remain pointing at the cloud provider's anycast addresses. Revealing the address of the HNA in the DNS is not desirable. Refer to Section 4.2 for more details.

This specification assumes:

- o the DM serves both the Control Channel and Synchronization Channel on a single IP address, single port and using a single transport protocol.
- o By default, the HNA uses a single IP address for both the Control and Synchronization channel. However, the HNA MAY use distinct IP addresses for the Control Channel and the Synchronization Channel - see Section 5 and Section 4.3 for more details.

The Distribution Channel is internal to the DOI and as such is not the primary concern of this specification.

4. Control Channel

The DM Control Channel is used by the HNA and the DOI to exchange information related to the configuration of the delegation which includes information to build the Public Homenet Zone (Section 4.1),

information to build the DNSSEC chain of trust (Section 4.2) and information to set the Synchronization Channel (Section 4.3).

4.1. Information to Build the Public Homenet Zone

When the HNA builds the Public Homenet Zone, it must include information that it retrieves from the DM relating to how the zone is to be published.

The information includes at least names and IP addresses of the Public Authoritative Name Servers. In term of RRset information this includes:

- o the MNAME of the SOA,
- o the NS and associated A and AAA RRsets of the name servers.

The DOI MAY also provide operational parameters such as other fields of SOA (SERIAL, RNAME, REFRESH, RETRY, EXPIRE and MINIMUM). As the information is necessary for the HNA to proceed and the information is associated to the DOI, this information exchange is mandatory.

4.2. Information to build the DNSSEC chain of trust

The HNA SHOULD provide the hash of the KSK (DS RRset), so the that DOI provides this value to the parent zone. A common deployment use case is that the DOI is the registrar of the Registered Homenet Domain, and as such, its relationship with the registry of the parent zone enables it to update the parent zone. When such relation exists, the HNA should be able to request the DOI to update the DS RRset in the parent zone. A direct update is especially necessary to initialize the chain of trust.

Though the HNA may also later directly update the values of the DS via the Control Channel, it is RECOMMENDED to use other mechanisms such as CDS and CDNSKEY [RFC7344] for transparent updates during key roll overs.

As some deployments may not provide a DOI that will be able to update the DS in the parent zone, this information exchange is OPTIONAL.

By accepting the DS RR, the DM commits in taking care of advertising the DS to the parent zone. Upon refusal, the DM clearly indicates it does not have the capacity to proceed to the update.

4.3. Information to set the Synchronization Channel

The HNA works as a primary authoritative DNS server, while the DM works like a secondary. As a result, the HNA MUST provide the IP address the DM is using to reach the HNA. The synchronization Channel will be set between that IP address and the IP address of the DM. By default, the IP address used by the HNA in the Control Channel is considered by the DM and the specification of the IP by the HNA is only OPTIONAL. The transport channel (including port number) is the same as the one used between the HNA and the DM for the Control Channel.

4.4. Deleting the delegation

The purpose of the previous sections were to exchange information in order to set a delegation. The HNA MUST also be able to delete a delegation with a specific DM. Upon an instruction of deleting the delegation, the DM MUST stop serving the Public Homenet Zone.

4.5. Messages Exchange Description

There are multiple ways this information could be exchanged between the HNA and the DM. This specification defines a mechanism that re-use the DNS exchanges format. The intention is to reuse standard libraries especially to check the format of the exchanged fields as well as to minimize the additional libraries needed for the HNA. The re-use of DNS exchanges achieves these goals. Note that while information is provided using DNS exchanges, the exchanged information is not expected to be set in any zone file, instead this information is uses as commands between the HNA and the DM.

The Control Channel is not expected to be a long term session. After a predefined timer - similar to those used for TCP - the Control Channel is expected to be terminated - by closing the transport channel. The Control Channel MAY be re-opened at any time later.

The provisioning process SHOULD provide a method of securing the Control Channel, so that the content of messages can be authenticated. This authentication MAY be based on certificates for both the DM and each HNA. The DM may also create the initial configuration for the delegation zone in the parent zone during the provisioning process.

4.5.1. Retrieving information for the Public Homenet Zone.

The information provided by the DM to the HNA is retrieved by the HNA with an AXFR exchange [RFC1034]. AXFR enables the response to contain any type of RRsets. The response might be extended in the

future if additional information will be needed. Alternatively, the information provided by the HNA to the DM is pushed by the HNA via a DNS update exchange [RFC2136].

To retrieve the necessary information to build the Public Homenet Zone, the HNA MUST send a DNS request of type AXFR associated to the Registered Homenet Domain. The DM MUST respond with a zone template. The zone template MUST contain a RRset of type SOA, one or multiple RRset of type NS and zero or more RRset of type A or AAAA.

- o The SOA RR indicates to the HNA the value of the MNAME of the Public Homenet Zone.
- o The NAME of the SOA RR MUST be the Registered Homenet Domain.
- o The MNAME value of the SOA RDATA is the value provided by the DOI to the HNA.
- o Other RDATA values (RNAME, REFRESH, RETRY, EXPIRE and MINIMUM) are provided by the DOI as suggestions.

The NS RRsets carry the Public Authoritative Servers of the DOI. Their associated NAME MUST be the Registered Homenet Domain.

The TTL and RDATA are those expected to be published on the Public Homenet Zone. The RRsets of Type A and AAAA MUST have their NAME matching the NSDNAME of one of the NS RRsets.

Upon receiving the response, the HNA MUST validate format and properties of the SOA, NS and A or AAAA RRsets. If an error occurs, the HNA MUST stop proceeding and MUST log an error. Otherwise, the HNA builds the Public Homenet Zone by setting the MNAME value of the SOA as indicated by the SOA provided by the AXFR response. The HNA SHOULD set the value of NAME, REFRESH, RETRY, EXPIRE and MINIMUM of the SOA to those provided by the AXFR response. The HNA MUST insert the NS and corresponding A or AAAA RRset in its Public Homenet Zone. The HNA MUST ignore other RRsets. If an error message is returned by the DM, the HNA MUST proceed as a regular DNS resolution. Error messages SHOULD be logged for further analysis. If the resolution does not succeed, the outsourcing operation is aborted and the HNA MUST close the Control Channel.

4.5.2. Providing information for the DNSSEC chain of trust

To provide the DS RRset to initialize the DNSSEC chain of trust the HNA MAY send a DNS update [RFC2136] message.

The DNS update message is composed of a Header section, a Zone section, a Pre-requisite section, and Update section and an additional section. The Zone section MUST set the ZNAME to the parent zone of the Registered Homenet Domain - that is where the DS records should be inserted. As described [RFC2136], ZTYPE is set to SOA and ZCLASS is set to the zone's class. The Pre-requisite section MUST be empty. The Update section is a DS RRset with its NAME set to the Registered Homenet Domain and the associated RDATA corresponds to the value of the DS. The Additional Data section MUST be empty.

Though the pre-requisite section MAY be ignored by the DM, this value is fixed to remain coherent with a standard DNS update.

Upon receiving the DNS update request, the DM reads the DS RRset in the Update section. The DM checks ZNAME corresponds to the parent zone. The DM SHOULD ignore non empty the Pre-requisite and Additional Data section. The DM MAY update the TTL value before updating the DS RRset in the parent zone. Upon a successful update, the DM should return a NOERROR response as a commitment to update the parent zone with the provided DS. An error indicates the MD does not update the DS, and other method should be used by the HNA.

The regular DNS error message SHOULD be returned to the HNA when an error occurs. In particular a FORMERR is returned when a format error is found, this includes when unexpected RRsets are added or when RRsets are missing. A SERVFAIL error is returned when a internal error is encountered. A NOTZONE error is returned when update and Zone sections are not coherent, a NOTAUTH error is returned when the DM is not authoritative for the Zone section. A REFUSED error is returned when the DM refuses to proceed to the configuration and the requested action.

4.5.3. Providing information for the Synchronization Channel

The default IP address used by the HNA for the Synchronization Channel is the IP address of the Control Channel. To provide a different IP address, the HNA MAY send a DNS Update message.

Similarly to the Section 4.5.2, the HNA MAY specify the IP address using a DNS update message. The Zone section sets its ZNAME to the parent zone of the Registered Homenet Domain, ZTYPE is set to SOA and ZCLASS is set to the zone's type. Pre-requisite is empty. The Update section is a RRset of type NS. The Additional Data section contains the RRsets of type A or AAAA that designates the IP addresses associated to the primary (or the HNA).

The reason to provide these IP addresses is to keep them unpublished and prevent them to be resolved.

Upon receiving the DNS update request, the DM reads the IP addresses and checks the ZNAME corresponds to the parent zone. The DM SHOULD ignore a non empty Pre-requisite section. The DM configures the secondary with the IP addresses and returns a NOERROR response to indicate it is committed to serve as a secondary.

Similarly to Section 4.5.2, DNS errors are used and an error indicates the DM is not configured as a secondary.

4.5.4. HNA instructing deleting the delegation

To instruct to delete the delegation the HNA SHOULD send a DNS UPDATE Delete message.

The Zone section sets its ZNAME to the Registered Homenet Domain, the ZTYPE to SOA and the ZCLASS to zone's type. The Pre-requisite section is empty. The Update section is a RRset of type NS with the NAME set to the Registered Domain Name. As indicated by [RFC2136] Section 2.5.2 the delete instruction is set by setting the TTL to 0, the Class to ANY, the RDLENGTH to 0 and the RDATA MUST be empty. The Additional Data section is empty.

Upon receiving the DNS update request, the DM checks the request and removes the delegation. The DM returns a NOERROR response to indicate the delegation has been deleted. Similarly to Section 4.5.2, DNS errors are used and an error indicates the delegation has not been deleted.

4.6. Securing the Control Channel

The control channel between the HNA and the DM MUST be secured at both the HNA and the DM.

Secure protocols (like TLS [RFC8446] SHOULD be used to secure the transactions between the DM and the HNA.

The advantage of TLS is that this technology is widely deployed, and most of the devices already embed TLS libraries, possibly also taking advantage of hardware acceleration. Further, TLS provides authentication facilities and can use certificates to mutually authenticate the DM and HNA at the application layer, including available API. On the other hand, using TLS requires implementing DNS exchanges over TLS, as well as a new service port.

The HNA SHOULD authenticate inbound connections from the DM using standard mechanisms, such as a public certificate with baked-in root certificates on the HNA, or via DANE [RFC6698]. The HNA is expected

to be provisioned with a connection to the DM by the manufacturer, or during some user-initiated onboarding process, see Section 10.

The DM SHOULD authenticate the HNA and check that inbound messages are from the appropriate client. The DM MAY use a self-signed CA certificate mechanism per HNA, or public certificates for this purpose.

IPsec [RFC4301] and IKEv2 [RFC7296] were considered. They would need to operate in transport mode, and the authenticated end points would need to be visible to the applications, and this is not commonly available at the time of this writing.

A pure DNS solution using TSIG and/or SIG(0) to authenticate message was also considered. Section 10 envisions one mechanism would involve the end user, with a browser, signing up to a service provider, with a resulting OAuth2 token to be provided to the HNA. A way to translate this OAuth2 token from HTTPS web space to DNS SIG(0) space seems overly problematic, and so the enrollment protocol using web APIs was determined to be easier to implement at scale.

Note also that authentication of message exchanges between the HNA and the DM SHOULD NOT use the external IP address of the HNA to index the appropriate keys. As detailed in Section 11, the IP addresses of the DM and the hidden primary are subject to change, for example while the network is being renumbered. This means that the necessary keys to authenticate transaction SHOULD NOT be indexed using the IP address, and SHOULD be resilient to IP address changes.

4.7. Implementation Concerns

The Hidden Primary Server on the HNA differs from a regular authoritative server for the home network due to:

Interface Binding: the Hidden Primary Server will almost certainly listen on the WAN Interface, whereas a regular Homenet Authoritative Servers would listen on the internal home network interface.

Limited exchanges: the purpose of the Hidden Primary Server is to synchronize with the DM, not to serve any zones to end users, or the public Internet.

As a result, exchanges are performed with specific nodes (the DM). Further, exchange types are limited. The only legitimate exchanges are: NOTIFY initiated by the Hidden Primary and IXFR or AXFR exchanges initiated by the DM.

On the other hand, regular authoritative servers would respond to any hosts, and any DNS query would be processed. The HNA SHOULD filter IXFR/AXFR traffic and drop traffic not initiated by the DM. The HNA MUST at least allow SOA lookups of the Homenet Zone.

5. Synchronization Channel

The DM Synchronization Channel is used for communication between the HNA and the DM for synchronizing the Public Homenet Zone. Note that the Control Channel and the Synchronization Channel are by construction different channels even though there they may use the same IP address. In fact the Control Channel is set between the HNA working as a client using port number YYYY (a high range port) toward a service provided by the DM at port number XX (well known port).

On the other hand, the Synchronization Channel is set between the DM working as a client using port ZZZZ (a high range port) toward a service a service provided by the HNA at port XX.

As a result, even though the same couple of IP addresses may be involved the Control Channel and the Synchronization Channel are always distinct channels.

Uploading and dynamically updating the zone file on the DM can be seen as zone provisioning between the HNA (Hidden Primary) and the DM (Secondary Server). This can be handled via AXFR + DNS Update.

The use of a primary / secondary mechanism is RECOMMENDED instead of the use of DNS Update. The primary / secondary mechanism is RECOMMENDED as it scales better and avoids DoS attacks. Note that even when UPDATE messages are used, these messages are using a distinct channel as those used to set the configuration.

Note that there is no standard way to distribute a DNS primary between multiple devices. As a result, if multiple devices are candidate for hosting the Hidden Primary, some specific mechanisms should be designed so the home network only selects a single HNA for the Hidden Primary. Selection mechanisms based on HNCP [RFC7788] are good candidates.

The HNA acts as a Hidden Primary Server, which is a regular authoritative DNS Server listening on the WAN interface.

The DM is configured as a secondary for the Registered Homenet Domain Name. This secondary configuration has been previously agreed between the end user and the provider of the DOI as part of either the provisioning or due to receipt of DNS Update messages on the DM Control Channel.

The Homenet Reverse Zone MAY also be updated either with DNS UPDATE [RFC2136] or using a primary / secondary synchronization.

5.1. Securing the Synchronization Channel

The Synchronization Channel uses standard DNS requests.

First the primary notifies the secondary that the zone must be updated and leaves the secondary to proceed with the update when possible/convenient.

Then, a NOTIFY message is sent by the primary, which is a small packet that is less likely to load the secondary.

Finally, the AXFR [RFC1034] or IXFR [RFC1995] query performed by the secondary is a small packet sent over TCP (Section 4.2 [RFC5936]), which mitigates reflection attacks using a forged NOTIFY.

The AXFR request from the DM to the HNA SHOULD be secured and the use of TLS is RECOMMENDED [I-D.ietf-dprive-xfr-over-tls]

When using TLS, the HNA MAY authenticate inbound connections from the DM using standard mechanisms, such as a public certificate with baked-in root certificates on the HNA, or via DANE [RFC6698]. In addition, to guarantee the DM remains the same across multiple TLS session, the HNA and DM MAY implement [RFC8672].

The HNA SHOULD apply an ACL on inbound AXFR requests to ensure they only arrive from the DM Synchronization Channel. In this case, the HNA SHOULD regularly check (via DNS resolution) that the address of the DM in the filter is still valid.

6. DM Distribution Channel

The DM Distribution Channel is used for communication between the DM and the Public Authoritative Servers. The architecture and communication used for the DM Distribution Channels is outside the scope of this document, and there are many existing solutions available e.g. rsynch, DNS AXFR, REST, DB copy.

7. HNA Security Policies

This section details security policies related to the Hidden Primary / Secondary synchronization.

The HNA, as Hidden Primary SHOULD drop any queries from the home network. This could be implemented via port binding and/or firewall rules. The precise mechanism deployed is out of scope of this

document. The Hidden Primary SHOULD drop any DNS queries arriving on the WAN interface that are not issued from the DM. The Hidden Primary SHOULD drop any outgoing packets other than DNS NOTIFY query, SOA response, IXFR response or AXFR responses. The Hidden Primary SHOULD drop any incoming packets other than DNS NOTIFY response, SOA query, IXFR query or AXFR query. The Hidden Primary SHOULD drop any non protected IXFR or AXFR exchange, depending on how the synchronization is secured.

8. DNSSEC compliant Homenet Architecture

[RFC7368] in Section 3.7.3 recommends DNSSEC to be deployed on both the authoritative server and the resolver. The resolver side is out of scope of this document, and only the authoritative part of the server is considered.

This document assumes the HNA signs the Public Homenet Zone.

Secure delegation is achieved only if the DS RRset is properly set in the parent zone. Secure delegation is performed by the HNA or the DOIs.

The DS RRset can be updated manually with nsupdate for example. This requires the HNA or the DOI to be authenticated by the DNS server hosting the parent of the Public Homenet Zone. Such a trust channel between the HNA and the parent DNS server may be hard to maintain with HNAs, and thus may be easier to establish with the DOI. In fact, the Public Authoritative Server(s) may use Automating DNSSEC Delegation Trust Maintenance [RFC7344].

9. Homenet Reverse Zone Channels Configuration

The Public Homenet Zone is associated to a Registered Homenet Domain and the ownership of that domain requires a specific registration from the end user as well as the HNA being provisioned with some authentication credentials. Such steps are mandatory unless the DOI has some other means to authenticate the HNA. Such situation may occur, for example, when the ISP provides the Homenet Domain as well as the DOI.

In this case, the HNA may be authenticated by the physical link layer, in which case the authentication of the HNA may be performed without additional provisioning of the HNA. While this may not be so common for the Public Homenet Zone, this situation is expected to be quite common for the Reverse Homenet Zone.

More specifically, a common case is that the upstream ISP provides the IPv6 prefix to the Homenet with a IA_PD [RFC8415] option and manages the DOI of the associated reverse zone.

This leave place for setting up automatically the relation between HNA and the DNS Outsourcing infrastructure as described in [I-D.ietf-homenet-naming-architecture-dhc-options].

In the case of the reverse zone, the DOI authenticates the source of the updates by IPv6 Access Control Lists. In the case of the reverse zone, the ISP knows exactly what addresses have been delegated. The HNA SHOULD therefore always originate Synchronization Channel updates from an IP address within the zone that is being updated.

For example, if the ISP has assigned 2001:db8:f00d::/64 to the WAN interface (by DHCPv6, or PPP/RA), then the HNA should originate Synchronization Channel updates from, for example, 2001:db8:f00d::2.

An ISP that has delegated 2001:db8:babe::/56 to the HNA via DHCPv6-PD, then HNA should originate Synchronization Channel updates an IP within that subnet, such as 2001:db8:babe:0001::2.

With this relation automatically configured, the synchronization between the Home network and the DOI happens similarly as for the Public Homenet Zone described earlier in this document.

Note that for home networks connected to by multiple ISPs, each ISP provides only the DOI of the reverse zones associated to the delegated prefix. It is also likely that the DNS exchanges will need to be performed on dedicated interfaces as to be accepted by the ISP. More specifically, the reverse zone associated to prefix 1 will not be possible to be performs by the HNA using an IP address that belongs to prefix 2. Such constraints does not raise major concerns either for hot standby or load sharing configuration.

With IPv6, the domain space for IP addresses is so large that reverse zone may be confronted with scalability issues. How the reverse zone is generated is out of scope of this document. [RFC8501] provides guidance on how to address scalability issues.

10. Homenet Public Zone Channel Configurations

This document does not deal with how the HNA is provisioned with a trusted relationship to the Distribution Manager for the forward zone.

This section details what needs to be provisioned into the HNA and serves as a requirements statement for mechanisms.

The HNA needs to be provisioned with:

- o the Registered Domain (e.g., myhome.isp.example)
- o the contact info for the Distribution Manager (DM), including the DNS name (FQDN), possibly including the IP literal, and a certificate (or anchor) to be used to authenticate the service
- o the DM transport protocol and port (the default is DNS over TLS, on port 853)
- o the HNA credentials used by the DM for its authentication.

The HNA will need to select an IP address for communication for the Synchronization Channel. This is typically the WAN address of the RG router, but could be an IPv6 LAN address in the case of a home with multiple ISPs (and multiple border routers). This is detailed in Section 4.5.3 when the NS and A or AAAA RRsets are communicated.

The above parameters MUST be be provisioned for ISP-specific reverse zones, as per [I-D.ietf-homenet-naming-architecture-dhc-options]. ISP-specific forward zones MAY also be provisioned using [I-D.ietf-homenet-naming-architecture-dhc-options], but zones which are not related to a specific ISP zone (such as with a DNS provider) must be provisioned through other means.

Similarly, if the HNA is provided by a registrar, the HNA may be handed pre-configured to end user.

In the absence of specific pre-established relation, these pieces of information may be entered manually by the end user. In order to ease the configuration from the end user the following scheme may be implemented.

The HNA may present the end user a web interface where it provides the end user the ability to indicate the Registered Homenet Domain or the registrar for example a preselected list. Once the registrar has been selected, the HNA redirects the end user to that registrar in order to receive a access token. The access token will enable the HNA to retrieve the DM parameters associated to the Registered Domain. These parameters will include the credentials used by the HNA to establish the Control and Synchronization Channels.

Such architecture limits the necessary steps to configure the HNA from the end user.

11. Renumbering

This section details how renumbering is handled by the Hidden Primary server or the DM. Both types of renumbering are discussed i.e. "make-before-break" and "break-before-make" (aka flash renumbering).

In the make-before-break renumbering scenario, the new prefix is advertised, the network is configured to prepare the transition to the new prefix. During a period of time, the two prefixes old and new coexist, before the old prefix is completely removed.

In the break-before-make renumbering scenario, the new prefix is advertised making the old prefix obsolete.

Renumbering has been extensively described in [RFC4192] and analyzed in [RFC7010] and the reader is expected to be familiar with them before reading this section.

11.1. Hidden Primary

In a renumbering scenario, the HNA or Hidden Primary is informed it is being renumbered. In most cases, this occurs because the whole home network is being renumbered. As a result, the Public Homenet Zone will also be updated. Although the new and old IP addresses may be stored in the Public Homenet Zone, we recommend that only the newly reachable IP addresses be published.

To avoid reachability disruption, IP connectivity information provided by the DNS SHOULD be coherent with the IP in use. In our case, this means the old IP address SHOULD NOT be provided via the DNS when it is not reachable anymore. Let for example TTL be the TTL associated with a RRset of the Public Homenet Zone, it may be cached for TTL seconds. Let T_NEW be the time the new IP address replaces the old IP address in the Homenet Zone, and T_OLD_UNREACHABLE the time the old IP is not reachable anymore.

In the case of the make-before-break, seamless reachability is provided as long as $T_OLD_UNREACHABLE - T_NEW > 2 * TTL$. If this is not satisfied, then devices associated with the old IP address in the home network may become unreachable for $2 * TTL - (T_OLD_UNREACHABLE - T_NEW)$. In the case of a break-before-make, $T_OLD_UNREACHABLE = T_NEW$, and the device may become unreachable up to $2 * TTL$. Of course if $T_NEW \geq T_OLD_UNREACHABLE$, the disruption is increased.

Once the Public Homenet Zone file has been updated on the Hidden Primary, the Hidden Primary needs to inform the DOI that the Public Homenet Zone has been updated and that the IP address to use to retrieve the updated zone has also been updated. Both notifications

are performed using regular DNS exchanges. Mechanisms to update an IP address provided by lower layers with protocols like SCTP [RFC4960], MOBIKE [RFC4555] are not considered in this document. Instead the IP address of the HNA is updated using the Synchronization Channel as described in Section 4.3.

12. Privacy Considerations

Outsourcing the DNS Authoritative service from the HNA to a third party raises a few privacy related concerns.

The Public Homenet Zone lists the names of services hosted in the home network. Combined with blocking of AXFR queries, the use of NSEC3 [RFC5155] (vs NSEC [RFC4034]) prevents an attacker from being able to walk the zone, to discover all the names. However, the attacker may be able to walk the reverse DNS zone, or use other reconnaissance techniques to learn this information as described in [RFC7707].

In general a home network owner is expected to publish only names for which there is some need to be able to reference externally. Publication of the name does not imply that the service is necessarily reachable from any or all parts of the Internet. [RFC7084] mandates that the outgoing-only policy [RFC6092] be available, and in many cases it is configured by default. A well designed User Interface would combine a policy for making a service public by a name with a policy on who may access it.

In many cases, the home network owner wishes to publish names for services that only they will be able to access. The access control may consist of an IP source address range, or access may be restricted via some VPN functionality. The purpose of publishing the name is so that the service may be access by the same name both within the home, and outside the home. Sending traffic to the relevant IPv6 address causes the relevant VPN policy to be enacted upon.

While the problem of getting access to internal names has been solved in Enterprise configurations with a split-DNS, and such a thing could be done in the home, many recent improvements to VPN user interfaces make it more likely that an individual might have multiple connections configured. For instance, an adult child checking on the state of a home automation system for a parent.

In addition to the Public Homenet Zone, pervasive DNS monitoring can also monitor the traffic associated with the Public Homenet Zone. This traffic may provide an indication of the services an end user accesses, plus how and when they use these services. Although,

caching may obfuscate this information inside the home network, it is likely that outside your home network this information will not be cached.

13. Security Considerations

This document exposes a mechanism that prevents the HNA from being exposed to the Internet and served DNS request from the Internet. These requests are instead served by the DOI. While this limits the level of exposure of the HNA, the HNA remains exposed to the Internet with communications with the DOI. This section analyses the attack surface associated to these communications. In addition, the DOI exposes data that are related to the home network. This section also analyses the implication of such exposure.

13.1. HNA DM channels

The channels between HNA and DM are mutually authenticated and encrypted with TLS [RFC8446] and its associated security considerations apply. To ensure the multiple TLS session are continuously authenticating the same entity, TLS may take advantage of second factor authentication as described in [RFC8672].

At the time of writing TLS 1.2 or TLS 1.3 can be used and TLS 1.3 (or newer) SHOULD be supported.

The DNS protocol is subject to reflection attacks, however, these attacks are largely applicable when DNS is carried over UDP. The interfaces between the HNA and DM are using TLS over TCP, which prevents such reflection attacks. Note that Public Authoritative servers hosted by the DOI are subject to such attacks, but that is out of scope of our document.

Note that in the case of the Reverse Homenet Zone, the data is less subject to attacks than in the Public Homenet Zone. In addition, the DM and RDM may be provided by the ISP - as described in [I-D.ietf-homenet-naming-architecture-dhc-options], in which case DM and RDM might be less exposed to attacks - as communications within a network.

13.2. Names are less secure than IP addresses

This document describes how an end user can make their services and devices from his home network reachable on the Internet by using names rather than IP addresses. This exposes the home network to attackers, since names are expected to include less entropy than IP addresses. In fact, with IP addresses, the Interface Identifier is 64 bits long leading to up to 2^{64} possibilities for a given

subnetwork. This is not to mention that the subnet prefix is also of 64 bits long, thus providing up to 2^{64} possibilities. On the other hand, names used either for the home network domain or for the devices present less entropy (livebox, router, printer, nicolas, jennifer, ...) and thus potentially exposes the devices to dictionary attacks.

13.3. 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 a time invariant prefix by ISPs. As a result, observing IP addresses only 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 an 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 end users may choose not to respond to PTR DNS queries and MAY instead return a NXDOMAIN response.

14. Information Model for Outsourced information

This section is non-normative for the front-end protocol. It specifies an optional format for the set of parameters required by the HNA to configure the naming architecture of this document.

In cases where a home router has not been provisioned by the manufacturer (when forward zones are provided by the manufacturer), or by the ISP (when the ISP provides this service), then a home user/owner will need to configure these settings via an administrative interface.

By defining a standard format (in JSON) for this configuration information, the user/owner may be able to just copy and paste a configuration blob from the service provider into the administrative interface of the HNA.

This format may also provide the basis for a future OAUTH2 [RFC6749] flow that could do the setup automatically.

The HNA needs to be configured with the following parameters as described by this CDDL [RFC8610]. These are the parameters necessary to establish a secure channel between the HNA and the DM as well as to specify the DNS zone that is in the scope of the communication.

```

hna-configuration = {
  "registered_domain" : tstr,
  "dm"                : tstr,
  ? "dm_transport"    : "DoT"
  ? "dm_port"         : uint,
  ? "dm_acl"          : hna-acl / [ +hna-acl ]
  ? "hna_auth_method" : hna-auth-method
  ? "hna_certificate" : tstr
}

```

```

hna-acl          = tstr
hna-auth-method /= "certificate"

```

For example:

```

{
  "registered_domain" : "n8d234f.r.example.net",
  "dm"                : "2001:db8:1234:111:222::2",
  "dm_transport"      : "DoT",
  "dm_port"           : 4433,
  "dm_acl"            : "2001:db8:1f15:62e:21c::/64"
                        or [ "2001:db8:1f15:62e:21c::/64", ... ]
  "hna_auth_method"   : "certificate",
  "hna_certificate"   : "-----BEGIN CERTIFICATE-----\nMIIDTjCCFGy....",
}

```

14.1. Outsourced Information Model

Registered Homenet Domain (zone) The Domain Name of the zone.
Multiple Registered Homenet Domains may be provided. This will generate the creation of multiple Public Homenet Zones. This parameter is MANDATORY.

Distribution Manager notification address (dm) The associated FQDNs or IP addresses of the DM to which DNS notifies should be sent. This parameter is MANDATORY. IP addresses are optional and the FQDN is sufficient and preferred. If there are concerns about the security of the name to IP translation, then DNSSEC should be employed.

As the session between the HNA and the DM is authenticated with TLS, the use of names is easier.

As certificates are more commonly emitted for FQDN than for IP addresses, it is preferred to use names and authenticate the name of the DM during the TLS session establishment.

Supported Transport (dm_transport) The transport that carries the DNS exchanges between the HNA and the DM. Typical value is "DoT" but it may be extended in the future with "DoH", "DoQ" for example. This parameter is OPTIONAL and by default the HNA uses DoT.

Distribution Manager Port (dm_port) Indicates the port used by the DM. This parameter is OPTIONAL and the default value is provided by the Supported Transport. In the future, additional transport may not have default port, in which case either a default port needs to be defined or this parameter become MANDATORY.

Note that HNA does not defines ports for the Synchronization Channel. In any case, this is not expected to part of the configuration, but instead negotiated through the Configuration Channel. Currently the Configuration Channel does not provide this, and limits its agility to a dedicated IP address. If such agility is needed in the future, additional exchanges will need to be defined.

Authentication Method ("hna_auth_method"): How the HNA authenticates itself to the DM within the TLS connection(s). The authentication meth of can typically be "certificate", "psk" or "none". This Parameter is OPTIONAL and by default the Authentication Method is "certificate".

Authentication data ("hna_certificate", "hna_key"): : The certificate chain used to authenticate the HNA. This parameter is OPTIONAL and when not specified, a self-signed certificate is used.

Distribution Manager AXFR permission netmask (dm_acl): The subnet from which the CPE should accept SOA queries and AXFR requests. A subnet is used in the case where the DNS Outsourced Infrastructure consists of a number of different systems. An array of addresses is permitted. This parameter is OPTIONAL and if unspecified, the CPE the IP addresses specified in the dm_notify parameters or the IP addresses that result from the DNS(SEC) resolution when dm_notify specifies a FQDN.

For forward zones, the relationship between the HNA and the forward zone provider may be the result of a number of transactions:

1. The forward zone outsourcing may be provided by the maker of the Homenet router. In this case, the identity and authorization could be built in the device at manufacturer provisioning time. The device would need to be provisioned with a device-unique credential, and it is likely that the Registered Homenet Domain would be derived from a public attribute of the device, such as a

serial number (see Appendix B or [I-D.richardson-homerouter-provisioning] for more details).

2. The forward zone outsourcing may be provided by the Internet Service Provider. In this case, the use of [I-D.ietf-homenet-naming-architecture-dhc-options] to provide the credentials is appropriate.
3. The forward zone may be outsourced to a third party, such as a domain registrar. In this case, the use of the JSON-serialized YANG data model described in this section is appropriate, as it can easily be copy and pasted by the user, or downloaded as part of a web transaction.

For reverse zones, the relationship is always with the upstream ISP (although there may be more than one), and so [I-D.ietf-homenet-naming-architecture-dhc-options] is always the appropriate interface.

The following is an abridged example of a set of data that represents the needed configuration parameters for outsourcing.

15. IANA Considerations

This document has no actions for IANA.

16. 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 HNA and low power devices; Olafur Gudmundsson for clarifying DNSSEC capabilities of small devices; Simon Kelley for its feedback as dnsmasq implementer; Andrew Sullivan, Mark Andrew, Ted Lemon, Mikael Abrahamson, and Ray Bellis for their feedback on handling different views as well as clarifying the impact of outsourcing the zone signing operation outside the HNA; Mark Andrew and Peter Koch for clarifying the renumbering.

17. Contributors

The co-authors would like to thank Chris Griffiths and Wouter Cloetens that provided a significant contribution in the early versions of the document.

18. References

18.1. Normative References

- [I-D.ietf-dprive-xfr-over-tls]
Toorop, W., Dickinson, S., Sahib, S., Aras, P., and A. Mankin, "DNS Zone Transfer-over-TLS", draft-ietf-dprive-xfr-over-tls-11 (work in progress), April 2021.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, DOI 10.17487/RFC1995, August 1996, <<https://www.rfc-editor.org/info/rfc1995>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2136] Vixie, P., Ed., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, DOI 10.17487/RFC2136, April 1997, <<https://www.rfc-editor.org/info/rfc2136>>.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, DOI 10.17487/RFC4034, March 2005, <<https://www.rfc-editor.org/info/rfc4034>>.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, DOI 10.17487/RFC4192, September 2005, <<https://www.rfc-editor.org/info/rfc4192>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4555] Eronen, P., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", RFC 4555, DOI 10.17487/RFC4555, June 2006, <<https://www.rfc-editor.org/info/rfc4555>>.
- [RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, DOI 10.17487/RFC4960, September 2007, <<https://www.rfc-editor.org/info/rfc4960>>.

- [RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, DOI 10.17487/RFC5155, March 2008, <<https://www.rfc-editor.org/info/rfc5155>>.
- [RFC5936] Lewis, E. and A. Hoenes, Ed., "DNS Zone Transfer Protocol (AXFR)", RFC 5936, DOI 10.17487/RFC5936, June 2010, <<https://www.rfc-editor.org/info/rfc5936>>.
- [RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, DOI 10.17487/RFC6092, January 2011, <<https://www.rfc-editor.org/info/rfc6092>>.
- [RFC6644] Evans, D., Droms, R., and S. Jiang, "Rebind Capability in DHCPv6 Reconfigure Messages", RFC 6644, DOI 10.17487/RFC6644, July 2012, <<https://www.rfc-editor.org/info/rfc6644>>.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012, <<https://www.rfc-editor.org/info/rfc6698>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, DOI 10.17487/RFC6887, April 2013, <<https://www.rfc-editor.org/info/rfc6887>>.
- [RFC7010] Liu, B., Jiang, S., Carpenter, B., Venaas, S., and W. George, "IPv6 Site Renumbering Gap Analysis", RFC 7010, DOI 10.17487/RFC7010, September 2013, <<https://www.rfc-editor.org/info/rfc7010>>.
- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.

- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC7344] Kumari, W., Gudmundsson, O., and G. Barwood, "Automating DNSSEC Delegation Trust Maintenance", RFC 7344, DOI 10.17487/RFC7344, September 2014, <<https://www.rfc-editor.org/info/rfc7344>>.
- [RFC7368] Chown, T., Ed., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, DOI 10.17487/RFC7368, October 2014, <<https://www.rfc-editor.org/info/rfc7368>>.
- [RFC7558] Lynn, K., Cheshire, S., Blanchet, M., and D. Migault, "Requirements for Scalable DNS-Based Service Discovery (DNS-SD) / Multicast DNS (mDNS) Extensions", RFC 7558, DOI 10.17487/RFC7558, July 2015, <<https://www.rfc-editor.org/info/rfc7558>>.
- [RFC7707] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", RFC 7707, DOI 10.17487/RFC7707, March 2016, <<https://www.rfc-editor.org/info/rfc7707>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8375] Pfister, P. and T. Lemon, "Special-Use Domain 'home.arpa.'", RFC 8375, DOI 10.17487/RFC8375, May 2018, <<https://www.rfc-editor.org/info/rfc8375>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8555] Barnes, R., Hoffman-Andrews, J., McCarney, D., and J. Kasten, "Automatic Certificate Management Environment (ACME)", RFC 8555, DOI 10.17487/RFC8555, March 2019, <<https://www.rfc-editor.org/info/rfc8555>>.

18.2. Informative References

- [I-D.ietf-dprive-dnsoquic]
Huitema, C., Mankin, A., and S. Dickinson, "Specification of DNS over Dedicated QUIC Connections", draft-ietf-dprive-dnsoquic-02 (work in progress), February 2021.
- [I-D.ietf-homenet-naming-architecture-dhc-options]
Migault, D., Weber, R., and T. Mrugalski, "DHCPv6 Options for Home Network Naming Authority", draft-ietf-homenet-naming-architecture-dhc-options-12 (work in progress), April 2021.
- [I-D.ietf-homenet-simple-naming]
Lemon, T., Migault, D., and S. Cheshire, "Homenet Naming and Service Discovery Architecture", draft-ietf-homenet-simple-naming-03 (work in progress), October 2018.
- [I-D.richardson-homerouter-provisioning]
Richardson, M., "Provisioning Initial Device Identifiers into Home Routers", draft-richardson-homerouter-provisioning-00 (work in progress), November 2020.
- [RFC2845] Vixie, P., Gudmundsson, O., Eastlake 3rd, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, DOI 10.17487/RFC2845, May 2000, <<https://www.rfc-editor.org/info/rfc2845>>.
- [RFC2931] Eastlake 3rd, D., "DNS Request and Transaction Signatures (SIG(0)s)", RFC 2931, DOI 10.17487/RFC2931, September 2000, <<https://www.rfc-editor.org/info/rfc2931>>.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, DOI 10.17487/RFC6749, October 2012, <<https://www.rfc-editor.org/info/rfc6749>>.

- [RFC8094] Reddy, T., Wing, D., and P. Patil, "DNS over Datagram Transport Layer Security (DTLS)", RFC 8094, DOI 10.17487/RFC8094, February 2017, <<https://www.rfc-editor.org/info/rfc8094>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.
- [RFC8501] Howard, L., "Reverse DNS in IPv6 for Internet Service Providers", RFC 8501, DOI 10.17487/RFC8501, November 2018, <<https://www.rfc-editor.org/info/rfc8501>>.
- [RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <<https://www.rfc-editor.org/info/rfc8610>>.
- [RFC8672] Sheffer, Y. and D. Migault, "TLS Server Identity Pinning with Tickets", RFC 8672, DOI 10.17487/RFC8672, October 2019, <<https://www.rfc-editor.org/info/rfc8672>>.

Appendix A. Envisioned deployment scenarios

A number of deployment have been envisioned, this section aims at providing a brief description. The use cases are not limitations and this section is not normative.

A.1. CPE Vendor

A specific vendor with specific relations with a registrar or a registry may sell a CPE that is provisioned with provisioned domain name. Such domain name does not need to be necessary human readable.

One possible way is that the vendor also provisions the HNA with a private and public keys as well as a certificate. Note that these keys are not expected to be used for DNSSEC signing. Instead these keys are solely used by the HNA to proceed to the authentication. Normally the keys should be necessary and sufficient to proceed to the authentication. The reason to combine the domain name and the key is that DOI are likely handle names better than keys and that domain names might be used as a login which enables the key to be regenerated.

When the home network owner plugs the CPE at home, the relation between HNA and DM is expected to work out-of-the-box.

A.2. Agnostic CPE

An CPE that is not preconfigured may also take advantage to the protocol defined in this document but some configuration steps will be needed.

1. The owner of the home network buys a domain name to a registrar, and as such creates an account on that registrar
2. Either the registrar is also providing the outsourcing infrastructure or the home network needs to create a specific account on the outsourcing infrastructure. * If the DOI is the registrar, it has by design a proof of ownership of the domain name by the homenet owner. In this case, it is expected the DOI provides the necessary parameters to the home network owner to configure the HNA. A good way to provide the parameters would be the home network be able to copy/paste a JSON object - see Section 14. What matters at that point is the DOI being able to generate authentication credentials for the HNA to authenticate itself to the DOI. This obviously requires the home network to provide the public key generated by the HNA in a CSR.

- o If the DOI is not the registrar, then the proof of ownership needs to be established using protocols like ACME [RFC8555] for example that will end in the generation of a certificate. ACME is used here to the purpose of automating the generation of the certificate, the CA may be a specific CA or the DOI. With that being done, the DOI has a roof of ownership and can proceed as above.

Appendix B. Example: A manufacturer provisioned HNA product flow

This scenario is one where a homenet router device manufacturer decides to offer DNS hosting as a value add.

[I-D.richardson-homerouter-provisioning] describes a process for a home router credential provisioning system. The outline of it is that near the end of the manufacturing process, as part of the firmware loading, the manufacturer provisions a private key and certificate into the device.

In addition to having a assymmetric credential known to the manufacturer, the device also has been provisioned with an agreed upon name. In the example in the above document, the name "n8d234f.r.example.net" has already been allocated and confirmed with the manufacturer.

The HNA can use the above domain for itself. It is not very pretty or personal, but if the owner wishes a better name, they can arrange for it.

The configuration would look like:

```
{
  "dm_notify" : "2001:db8:1234:111:222::2",
  "dm_acl"    : "2001:db8:1234:111:222::/64",
  "dm_ctrl"   : "manufacturer.example.net",
  "dm_port"   : "4433",
  "ns_list"   : [ "ns1.publicdns.example", "ns2.publicdns.example"],
  "zone"      : "n8d234f.r.example.net",
  "auth_method" : "certificate",
  "hna_certificate": "-----BEGIN CERTIFICATE-----\nMIIDTjCCFGy....",
}
```

The dm_ctrl and dm_port values would be built into the firmware.

Authors' Addresses

Daniel Migault
Ericsson
8275 Trans Canada Route
Saint Laurent, QC 4S 0B6
Canada

EMail: daniel.migault@ericsson.com

Ralf Weber
Nominum
2000 Seaport Blvd
Redwood City 94063
US

EMail: ralf.weber@nominum.com

Michael Richardson
Sandelman Software Works
470 Dawson Avenue
Ottawa, ON K1Z 5V7
Canada

EMail: mcr+iETF@sandelman.ca

Ray Hunter
Globis Consulting BV
Weegschaalstraat 3
Eindhoven 5632CW
NL

EMail: v6ops@globis.net

Homenet
Internet-Draft
Intended status: Standards Track
Expires: November 14, 2021

D. Migault
Ericsson
R. Weber
Akamai
T. Mrugalski
Internet Systems Consortium, Inc.
May 13, 2021

DHCPv6 Options for Home Network Naming Authority
draft-ietf-homenet-naming-architecture-dhc-options-14

Abstract

This document defines DHCPv6 options so an agnostic Homenet Naming Authority (HNA) can automatically proceed to the appropriate configuration and outsource the authoritative naming service for the home network. In most cases, the outsourcing mechanism is transparent for the end user.

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 <https://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 November 14, 2021.

Copyright Notice

Copyright (c) 2021 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 (<https://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. Terminology	2
2. Introduction	2
3. Procedure Overview	3
4. DHCPv6 Option	4
4.1. Registered Homenet Domain Option	4
4.2. Distribution Manager Option	5
4.2.1. Supported Transport	6
4.3. Reverse Distribution Manager Server Option	6
5. DHCPv6 Behavior	7
5.1. DHCPv6 Server Behavior	7
5.2. DHCPv6 Client Behavior	7
5.3. DHCPv6 Relay Agent Behavior	7
6. IANA Considerations	7
7. Security Considerations	8
8. Acknowledgments	8
9. Contributors	8
10. References	8
10.1. Normative References	8
10.2. Informative References	9
Appendix A. Scenarios and impact on the End User	11
Appendix B. Base Scenario	11
B.1. Third Party Registered Homenet Domain	11
B.2. Third Party DNS Infrastructure	12
B.3. Multiple ISPs	12
Authors' Addresses	13

1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The reader should be familiar with
[I-D.ietf-homenet-front-end-naming-delegation].

2. Introduction

[I-D.ietf-homenet-front-end-naming-delegation] specifies how an entity designated as the Homenet Naming Authority (HNA) outsources a Public Homenet Zone to an Outsourcing DNS Infrastructure (DOI).

This document describes how a network can provision the HNA with a specific DOI. This could be particularly useful for a DOI partly managed by an ISP, or to make home networks resilient to HNA replacement. The ISP delegates an IP prefix to the home network as well as the associated reverse zone. The ISP is thus aware of the owner of that IP prefix, and as such becomes a natural candidate for hosting the Homenet Reverse Zone - that is the Reverse Distribution Manager (RDM) and potentially the Reverse Public Authoritative Servers.

In addition, ISPs often identify the line of the home network with a name. Such name is used for their internal network management operations and is not a name the home network owner has registered to. ISPs may leverage such infrastructure and provide the homenet with a specific domain name designated as per [I-D.ietf-homenet-front-end-naming-delegation] a Homenet Registered Domain. Similarly to the reverse zone, ISPs are aware of who owns that domain name and may become a natural candidate for hosting the Homenet Zone - that is the Distribution Manager (DM) and the Public Authoritative Servers.

This document describes DHCPv6 options that enable an ISP to provide the necessary parameters to the HNA, to proceed. More specifically, the ISP provides the Registered Homenet Domain, necessary information on the DM and the RDM so the HNA can manage and upload the Public Homenet Zone and the Reverse Public Homenet Zone as described in [I-D.ietf-homenet-front-end-naming-delegation].

The use of DHCPv6 options may make the configuration completely transparent to the end user and provides a similar level of trust as the one used to provide the IP prefix - when provisioned via DHCP.

3. Procedure Overview

This section illustrates how a HNA receives the necessary information via DHCPv6 options to outsource its authoritative naming service to the DOI. For the sake of simplicity, and similarly to [I-D.ietf-homenet-front-end-naming-delegation], this section assumes that the HNA and the home network DHCPv6 client are collocated on the Customer Edge (CE) router [RFC7368]. Note also that this is not mandatory and the DHCPv6 client may instruct remotely the HNA and the DHCPv6 either with a proprietary protocol or a protocol that will be defined in the future. In addition, this section assumes the responsible entity for the DHCPv6 server is configured with the DM and RDM. This means a Registered Homenet Domain can be associated to the DHCPv6 client.

This scenario is believed to be the most popular scenario. This document does not ignore scenarios where the DHCPv6 server does not have privileged relations with the DM or RDM. These cases are discussed in Appendix A. Such scenarios do not necessarily require configuration for the end user and can also be zero-config.

The scenario considered in this section is as follows:

1. The HNA is willing to outsource the Public Homenet Zone or Homenet Reverse Zone. The DHCPv6 client is configured to include in its Option Request Option (ORO) the Registered Homenet Domain Option (OPTION_REGISTERED_DOMAIN), the Distribution Manager Option (OPTION_DIST_MANAGER) and the Reverse Distribution Manager Option (OPTION_REVERSE_DIST_MANAGER) option codes.
 2. The DHCPv6 server responds to the HNA with the requested DHCPv6 options based on the identified homenet. The DHCPv6 client passes the information to the HNA.
 3. The HNA is authenticated (eventually by a self signed certificate) by the DM and the RDM. The HNA builds the Homenet Zone (or the Homenet Reverse Zone) and proceed as described in [I-D.ietf-homenet-front-end-naming-delegation]. The DHCPv6 options provide the necessary non optional parameters described in Section 14 of [I-D.ietf-homenet-front-end-naming-delegation]. The HNA may complement the configurations with additional parameters via means not yet defined. Section 14 of [I-D.ietf-homenet-front-end-naming-delegation] describes such parameters that MAY take some specific non default value.
4. DHCPv6 Option

This section details the payload of the DHCPv6 options.

4.1. Registered Homenet Domain Option

The Registered Domain Option (OPTION_REGISTERED_DOMAIN) indicates the FQDN associated to the home network.




Figure 1: Registered Domain Option

- o option-code (16 bits): `OPTION_REGISTERED_DOMAIN`, the option code for the Registered Homenet Domain (TBD1).
- o option-len (16 bits): length in octets of the Registered Homenet Domain field as described in [RFC8415].
- o Registered Homenet Domain (variable): the FQDN registered for the homenet encoded as described in Section 10 of [RFC8415].

4.2. Distribution Manager Option

The Distributed Manager Option (`OPTION_DIST_MANAGER`) provides the HNA with the FQDN of the DM as well as the transport protocols for the communication between the HNA and the DM.




Figure 2: Distribution Manager Option

- o option-code (16 bits): `OPTION_DIST_MANAGER`, the option code for the Distribution Manager Option (TBD2).
- o option-len (16 bits): length in octets of the enclosed data as described in [RFC8415].

- o Supported Transport (16 bits): defines the supported transport by the DM (see Section 4.2.1). Each bit represents a supported transport, and a DM MAY indicate the support of multiple modes. The bit for DNS over TLS [RFC7858] MUST be set.
- o Distribution Manager FQDN (variable): the FQDN of the DM encoded as described in Section 10 of [RFC8415].

4.2.1. Supported Transport

The Supported Transport field of the DHCPv6 option indicates the supported transport protocols. Each bit represents a specific transport mechanism. A bit sets to 1 indicates the associated transport protocol is supported. The corresponding bits are assigned as described in Figure 3 and Section 6.

Bit Position	Transport Protocol	Reference
0	DNS over TLS	This-RFC
1-15	unallocated	

Figure 3: Supported Transport

DNS over TLS: indicates the support of DNS over TLS as described in [RFC7858].

4.3. Reverse Distribution Manager Server Option

The Reverse Distribution Manager Option (OPTION_REVERSE_DIST_MANAGER) provides the HNA with the FQDN of the DM as well as the transport protocols for the communication between the HNA and the DM.




Figure 4: Reverse Distribution Manager Option

- o option-code (16 bits): OPTION_REVERSE_DIST_MANAGER, the option code for the Reverse Distribution Manager Option (TBD3).
- o option-len (16 bits): length in octets of the option-data field as described in [RFC8415].
- o Supported Transport (16 bits): defines the supported transport by the RDM (see Section 4.2.1). Each bit represents a supported transport, and a RDM MAY indicate the support of multiple modes. The bit for DNS over TLS [RFC7858] MUST be set.
- o Reverse Distribution Manager FQDN (variable): the FQDN of the RDM encoded as described in section 10 of [RFC8415].

5. DHCPv6 Behavior

5.1. DHCPv6 Server Behavior

Sections 17.2.2 and 18.2 of [RFC8415] govern server operation in regards to option assignment. As a convenience to the reader, we mention here that the server will send option foo only if configured with specific values for foo and if the client requested it. In particular, when configured the DHCPv6 server sends the Registered Homenet Domain Option, Distribution Manager Option, the Reverse Distribution Manager Option when requested by the DHCPv6 client by including necessary option codes in its ORO.

5.2. DHCPv6 Client Behavior

The DHCPv6 client includes Registered Homenet Domain Option, Distribution Manager Option, the Reverse Distribution Manager Option in an ORO as specified in Sections 18.2.1, 18.2.2, 18.2.4, 18.2.5, 18.2.6, and 21.7 of [RFC8415].

Upon receiving a DHCPv6 option described in this document in the Reply message, the HNA SHOULD proceed as described in [I-D.ietf-homenet-front-end-naming-delegation].

5.3. DHCPv6 Relay Agent Behavior

There are no additional requirements for the DHCPv6 Relay agents.

6. IANA Considerations

IANA is requested to assign the following new DHCPv6 Option Codes in the registry maintained in: <https://www.iana.org/assignments/dhcpv6-parameters/dhcpv6-parameters.xhtml#dhcpv6-parameters-2>.

Value	Description	Client ORO	Singleton Option
TBD1	OPTION_REGISTERED_DOMAIN	Yes	No
TBD2	OPTION_DIST_MANAGER	Yes	Yes
TBD3	OPTION_REVERSE_DIST_MANAGER	Yes	Yes

IANA is requested to maintain a new number space of Supported Transport parameter in the Distributed Manager Option (OPTION_DIST_MANAGER) or the Reverse Distribution Manager Option (OPTION_REVERSE_DIST_MANAGER). The different parameters are defined in Figure 3 in Section 4.2.1. Future code points are assigned under Specification Required as per [RFC8126].

7. Security Considerations

The security considerations in [RFC8415] are to be considered. The use of DHCPv6 options provides a similar level of trust as the one used to provide the IP prefix. The link between the HNA and the DHCPv6 server may benefit from additional security for example by using [I-D.ietf-dhc-sedhcpv6].

8. Acknowledgments

We would like to thank Marcin Siodelski, Bernie Volz and Ted Lemon for their comments on the design of the DHCPv6 options. We would also like to thank Mark Andrews, Andrew Sullivan and Lorenzo Colliti for their remarks on the architecture design. The designed solution has been largely been inspired by Mark Andrews's document [I-D.andrews-dnsop-pd-reverse] as well as discussions with Mark. We also thank Ray Hunter for its reviews, its comments and for suggesting an appropriated terminology.

9. Contributors

The co-authors would like to thank Chris Griffiths and Wouter Cloetens that provided a significant contribution in the early versions of the document.

10. References

10.1. Normative References

[I-D.ietf-homenet-front-end-naming-delegation]
Migault, D., Weber, R., Richardson, M., and R. Hunter,
"Simple Provisioning of Public Names for Residential
Networks", draft-ietf-homenet-front-end-naming-
delegation-14 (work in progress), April 2021.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

10.2. Informative References

- [I-D.andrews-dnsop-pd-reverse]
Andrews, M., "Automated Delegation of IP6.ARPA reverse zones with Prefix Delegation", draft-andrews-dnsop-pd-reverse-02 (work in progress), November 2013.
- [I-D.ietf-dhc-sedhcpv6]
Li, L., Jiang, S., Cui, Y., Jinmei, T., Lemon, T., and D. Zhang, "Secure DHCPv6", draft-ietf-dhc-sedhcpv6-21 (work in progress), February 2017.
- [I-D.sury-dnsextn-cname-dname]
Sury, O., "CNAME+DNAME Name Redirection", draft-sury-dnsextn-cname-dname-00 (work in progress), April 2010.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, DOI 10.17487/RFC2181, July 1997, <<https://www.rfc-editor.org/info/rfc2181>>.

- [RFC6672] Rose, S. and W. Wijngaards, "DNAME Redirection in the DNS", RFC 6672, DOI 10.17487/RFC6672, June 2012, <<https://www.rfc-editor.org/info/rfc6672>>.
- [RFC7368] Chown, T., Ed., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, DOI 10.17487/RFC7368, October 2014, <<https://www.rfc-editor.org/info/rfc7368>>.

Appendix A. Scenarios and impact on the End User

This section details various scenarios and discuss their impact on the end user. This section is not normative and limits the description of a limited scope of scenarios that are assumed to be representative. Many other scenarios may be derived from these.

Appendix B. Base Scenario

The base scenario is the one described in Section 3 in which an ISP manages the DHCPv6 server, the DM and RDM.

The end user subscribes to the ISP (foo), and at subscription time registers for example.foo as its Registered Homenet Domain example.foo.

In this scenario, the DHCPv6 server, DM and RDM are managed by the ISP so the DHCPv6 server and as such can provide authentication credentials of the HNA to enable secure authenticated transaction with the DM and the Reverse DM.

The main advantage of this scenario is that the naming architecture is configured automatically and transparently for the end user. The drawbacks are that the end user uses a Registered Homenet Domain managed by the ISP and that it relies on the ISP naming infrastructure.

B.1. Third Party Registered Homenet Domain

This section considers the case when the end user wants its home network to use example.com not managed by her ISP (foo) as a Registered Homenet Domain. This section still consider the ISP manages the home network and still provides example.foo as a Registered Homenet Domain.

When the end user buys the domain name example.com, it may request to redirect the name example.com to example.foo using static redirection with CNAME [RFC2181], [RFC1034], DNAME [RFC6672] or CNAME+DNAME [I-D.sury-dnsextn-cname-dname].

This configuration is performed once when the domain name example.com is registered. The only information the end user needs to know is the domain name assigned by the ISP. Once this configuration is done no additional configuration is needed anymore. More specifically, the HNA may be changed, the zone can be updated as in Appendix B without any additional configuration from the end user.

The main advantage of this scenario is that the end user benefits from the Zero Configuration of the Base Scenario Appendix B. Then, the end user is able to register for its home network an unlimited number of domain names provided by an unlimited number of different third party providers. The drawback of this scenario may be that the end user still rely on the ISP naming infrastructure. Note that the only case this may be inconvenient is when the DNS servers provided by the ISPs results in high latency.

B.2. Third Party DNS Infrastructure

This scenario considers that the end user uses example.com as a Registered Homenet Domain, and does not want to rely on the authoritative servers provided by the ISP.

In this section we limit the outsourcing to the DM and Public Authoritative Server(s) to a third party. The Reverse Public Authoritative Server(s) and the RDM remain managed by the ISP as the IP prefix is managed by the ISP.

Outsourcing to a third party DM can be performed in the following ways:

1. Updating the DHCPv6 server Information. One can imagine a GUI interface that enables the end user to modify its profile parameters. Again, this configuration update is done once-for-ever.
2. Upload the configuration of the DM to the HNA. In some cases, the provider of the CE router hosting the HNA may be the registrar and provide the CE router already configured. In other cases, the CE router may request the end user to log into the registrar to validate the ownership of the Registered Homenet Domain and agree on the necessary credentials to secure the communication between the HNA and the DM. As described in [I-D.ietf-homenet-front-end-naming-delegation], such settings could be performed in an almost automatic way as to limit the necessary interactions with the end user.

B.3. Multiple ISPs

This scenario considers a HNA connected to multiple ISPs.

Suppose the HNA has been configured each of its interfaces independently with each ISPS as described in Appendix B. Each ISP provides a different Registered Homenet Domain.

The protocol and DHCPv6 options described in this document are fully compatible with a HNA connected to multiple ISPs with multiple Registered Homenet Domains. However, the HNA should be able to handle different Registered Homenet Domains. This is an implementation issue which is outside the scope of the current document.

If a HNA is not able to handle multiple Registered Homenet Domains, the HNA may remain connected to multiple ISP with a single Registered Homenet Domain. In this case, one entity is chosen to host the Registered Homenet Domain. This entity may be one of the ISP or a third party. Note that having multiple ISPs can be motivated for bandwidth aggregation, or connectivity fail-over. In the case of connectivity fail-over, the fail-over concerns the access network and a failure of the access network may not impact the core network where the DM and Public Authoritative Primaries are hosted. In that sense, choosing one of the ISP even in a scenario of multiple ISPs may make sense. However, for sake of simplicity, this scenario assumes that a third party has been chosen to host the Registered Homenet Domain. Configuration is performed as described in Appendix B.1 and Appendix B.2.

With the configuration described in Appendix B.1, the HNA is expect to be able to handle multiple Homenet Registered Domain, as the third party redirect to one of the ISPs servers. With the configuration described in Appendix B.2, DNS zone are hosted and maintained by the third party. A single DNS(SEC) Homenet Zone is built and maintained by the HNA. This latter configuration is likely to match most HNA implementations.

The protocol and DHCPv6 options described in this document are fully compatible with a HNA connected to multiple ISPs. To configure or not and how to configure the HNA depends on the HNA facilities. Appendix B and Appendix B.1 require the HNA to handle multiple Registered Homenet Domain, whereas Appendix B.2 does not have such requirement.

Authors' Addresses

Daniel Migault
Ericsson
8275 Trans Canada Route
Saint Laurent, QC 4S 0B6
Canada

EMail: daniel.migault@ericsson.com

Ralf Weber
Akamai

EMail: ralf.weber@akamai.com

Tomek Mrugalski
Internet Systems Consortium, Inc.
950 Charter Street
Redwood City 94063
US

EMail: tomasz.mrugalski@gmail.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 26, 2019

T. Lemon
Nibbhaya Consulting
D. Migault
Ericsson
S. Cheshire
Apple Inc.
October 23, 2018

Homenet Naming and Service Discovery Architecture
draft-ietf-homenet-simple-naming-03

Abstract

This document describes how names are published and resolved on homenets, and how hosts are configured to use these names to discover services on homenets. It presents the complete architecture, and describes a simple subset of that architecture that can be used in low-cost homenet routers.

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 <https://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on April 26, 2019.

Copyright Notice

Copyright (c) 2018 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

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

Table of Contents

1. Introduction	3
2. Requirements	3
2.1. Managed LAN versus Homenet	4
2.1.1. Multiple Provisioning Domains	5
2.2. Homenet-specific considerations	5
3. Terminology	6
4. Name	6
5. Authority	8
5.1. Reachability	8
5.2. Link Names	8
5.3. Authoritative name service for the homenet domain	9
5.4. Authoritative name service for per-link subdomains of the homenet domain	10
5.5. Authoritative name service for the ULA reverse mapping domain	10
5.6. Authoritative name service for the RFC1918 reverse mapping domains	10
6. Resolution	11
6.1. Round Robining	13
6.2. Retransmission	13
6.3. DNS Stateful Operations and DNS Push	13
6.4. Multicast DNS	14
6.5. Host behavior	14
7. Publication	14
7.1. DNSSD Service Registration Protocol	14
7.2. Homenet Reverse Mapping Update Protocol	15
7.2.1. Adding ULA reverse mappings	15
7.2.2. Adding RFC1918 reverse mappings	16
8. Host Configuration	16
9. Globally Unique Names	16
10. DNSSEC Validation	17
10.1. How trust is established	17
11. Homenet Delegation Registration Protocol	18
12. Using the Local Namespace While Away From Home	19
13. Expected Host Behavior	19
14. Management Considerations	19
15. Privacy Considerations	20
16. Security Considerations	20
17. IANA considerations	20
17.1. Homenet Reverse Registration Protocol	20
17.2. Homenet Delegation Registration Protocol	20
17.3. Unique Local Address Reserved Documentation Prefix	21

18. References	21
18.1. Normative References	21
18.2. Informative References	23
Authors' Addresses	23

1. Introduction

This document is a homenet architecture document. The term 'homenet' refers to a set of technologies that allow home network users to have a local-area network (LAN) with more than one physical link and, optionally, more than one internet service provider. Home network users are assumed not to be knowledgeable in network operations, so homenets automatically configure themselves, providing connectivity and service discovery within the home with no operator intervention. This document describes the aspect of homenet automatic configuration that has to do with service discovery and name resolution.

This architecture provides a minimal set of features required to enable seamless service discovery on a multi-link home network, but does not attempt to provide feature parity with a managed LAN.

This document begins by presenting a motivational list of requirements and considerations, which should give the reader a clear idea of the scope of the problem being solved. It then explains how each requirement is addressed, and provides references for relevant standards documents describing the details of the implementation. Not all requirements are addressed by this architecture document, but the basic requirements are satisfied, and this document serves as a foundation upon which solutions to the remaining problems can be built.

2. Requirements

Name service on a local area network (LAN) requires the following:

- o Name: a forward domain under which information about local services will be published
- o Authority: a name server that is authoritative for at least one forward domain and one or two reverse domains that are applicable to that network and is capable of signing and publishing the zones using DNSSEC
- o Resolution: a full-service caching DNS resolver that fully supports EDNS(0) and queries with the DO bit set
- o Publication: a mechanism that

- * allows services on the LAN to publish information about the services they provide
- * allows services to publish information on how to reach them
- * manages the lifetime of such information, so that it persists long enough to prevent spoofing, but protects end users from seeing stale information
- o Host configuration: one or more automatic mechanisms (e.g. DHCP or RA) that provide:
 - * caching resolver information to hosts on the LAN
 - * information about how services on the LAN can publish information
- o Trust: some basis for trusting the information that is provided by the service discovery system

2.1. Managed LAN versus Homenet

A managed network is one that has a (human) manager, or operator. The operator has authority over the network, and the authority to publish names in a forward DNS tree, and reverse names in the reverse tree. The operator has the authority to sign the respective trees with DNSSEC, and acquire TLS certificates for hosts/servers within the network.

On a managed LAN, many of these services can be provided by operators. When a new printer is added to the network, it can be added to the service discovery system (the authoritative server) manually. When a printer is taken out of service, it can be removed. In this scenario, the role of "publisher" is filled by the network operator.

In many managed LANs, establishment of trust for service discovery is simply on the basis of a belief that the local resolver will give a correct answer. Once the service has been discovered and chosen, there may be some security (e.g., TLS) that protects the connection to the service, but the trust model is often just "you're connected to a network you trust, so you can trust the printer that you discovered on this network."

A homenet does not have an operator, so functions that would normally be performed by the operator have to happen automatically. This has implications for trust establishment--since there is no operator

controlling what services are published locally, some other mechanism is required for basic trust establishment.

2.1.1. Multiple Provisioning Domains

Additionally, whereas in a managed LAN with multiple links to the Internet, the network operator can configure the network so that multihoming is handled seamlessly, in a homenet, multihoming must be handled using multiple provisioning domains [RFC7556].

When a host on a homenet connects to a host outside the homenet, and the homenet is multihomed, the source address that the host uses for connecting determines which upstream ISP connection is used. In principle, this is not a problem, because the Internet is a fully connected network, so any host that is on the Internet can be reached by any host on the Internet, regardless of how that host connects to the Internet.

Unfortunately in practice this is not always the case. Some ISPs provide special services to their end users that are only accessible when connected through the ISP. When such a service is discovered using that ISP's name server, a response will be provided that will only work if the host connects using a prefix provided by that ISP. If another ISP's prefix is used, the connection will fail.

In the case of content delivery networks (CDNs), using the name service of one ISP and then connecting through a second ISP may seem to work, but may provide very poor service.

In order to address this problem, the homenet naming architecture takes two approaches. First, for hosts that do not support provisioning domain separation, we make sure that all ISP name servers are consulted in such a way that Happy Eyeballs will tend to work. Second, for hosts that do support provisioning domain separation, we provide information to the hosts to identify provisioning domains, and we provide a mechanism that hosts can use to indicate which provisioning domain to use for a particular DNS query.

2.2. Homenet-specific considerations

A naming architecture for homenets therefore adds the following considerations:

- o All of the operations mentioned here must reliably function automatically, without any user intervention or debugging.

- o Because user intervention cannot be required, naming conflicts must be resolved automatically, and, to the extent possible, transparently.
- o Devices that provide services must be able to publish those services on the homenet, and those services must be available from any part of the homenet, not just the link to which the device is attached.
- o Homenets must address the problem of multiple provisioning domains, in the sense that the DNS may give a different answer depending on whether caching resolvers at one ISP or another are queried.

An additional requirement from the Homenet Architecture [RFC7556] is that hosts are not required to implement any homenet-specific capabilities in order to discover and access services on the homenet. This architecture may define optional homenet-specific features, but hosts that do not implement these features must work on homenets.

3. Terminology

This document uses the following terms and abbreviations:

HNR Homenet Router

SHNR Homenet Router implementing simple homenet naming architecture

AHNR Homenet Router implementing advanced homenet naming architecture

ISP Internet Service Provider

Forward Mapping A mapping between a host name or service name and some information about that host or service.

Reverse Mapping A mapping between an IP address and the host that has that IP address.

Homenet Domain A domain name that is used for publishing the names of devices and services that are present on the homenet. By default, 'home.arpa.'

4. Name

In order for names to be published on a homenet, it is necessary that there be a set of domain names under which such names are published. These domain names, together, are referred to as the "local domains."

By default, homenets publish names for forward lookups under the reserved domain 'home.arpa.' [RFC8375] publishing names.

So a host called 'example' that published its name on the homenet would publish its records on the domain name 'example.home.arpa.'. Because 'home.arpa.' is used by all homenets, it has no global meaning, and names published under the domain 'home.arpa' cannot be used outside of the homenet on which they are published.

How to publish names outside of the homenet is out of scope for this document. However, in order to address the problem of validating names published on the homenet using DNSSEC, it is necessary that the homenet have a globally valid delegation from the root. This allows hosts on the homenet to validate names published on the homenet using the DNS root trust anchor ([RFC4033] section 3.1).

It is not necessary that this delegation work for hosts off the homenet. HNRs implementing this specification do not answer queries from outside the homenet; however, when a validating resolver inside the homenet attempts to validate the chain of trust up to the root zone, the chain of trust will validate correctly, because the answer given for internally-available zones will be signed by a DS record that is present in the delegation externally.

If there is a valid delegation from the root, the homenet domain will be the name of the delegated domain. By default, there will be no delegation from the root; in this case, the homenet domainname will be 'home.arpa.'

In addition to the homenet domain, names are needed for reverse lookups. These names are dependent on the IP addressing used on the homenet. If the homenet is addressed with IPv4, a reverse domain corresponding to the IPv4 subnet [RFC1034] section 5.2.1 should be constructed. For example, if the homenet is allocating local IP addresses out of net 10 [RFC1918], a domain, '10.in-addr.arpa' would be required. Like 'home.arpa.', '10.in-addr.arpa' is a locally-served zone, and has no validity outside of the homenet.

If the homenet is addressed with IPv6, it is expected to have a unique local address prefix. The reverse mapping domain for hosts on any link in the subnet will be a subdomain of the reverse zone for the subset of the ULA prefix that is being advertised on that link. Every service on the homenet that supports IPv6 is expected to be reachable at an address that is configured using the ULA prefix. Therefore there is no need for any IPv6 reverse zone to be populated other than the ULA zone. So for example if the homenet's ULA prefix is fc00:2001:db8::/48, then the reverse domain name for the homenet would end in '8.b.d.0.1.0.0.2.0.0.d.f.ip6.arpa'.

5. Authority

There are two types of authoritative name service on the homenet. Every link on the homenet has a zone that is a subdomain of the homenet's primary domain. Authority for these zones is local to the HNR that is currently authoritative for that zone. The contents of these zones are served using DNSSD Discovery Proxy [I-D.ietf-dnssd-hybrid]. Consequently, there is no need for database replication in the case that a new HNR is elected; that HNR simply takes over the Discovery Relay function.

Name service for the homenet domain itself may be stateless or stateful. HNRs are not required to implement stateful service. If one or more HNRs on the homenet are capable of providing this service, then one of those HNRs is elected to act as the primary nameserver for the homenet domain; one or more HNRs may also act as secondary servers.

Name service for reverse mapping subdomains is only provided if one or more HNRs can provide stateful service. If no such server is present, the reverse mapping subdomains are not served. If stateful servers are present, the primary and secondary servers for these subdomains will be the same as for the homenet domain.

5.1. Reachability

Whether the homenet domain is a global domain name or not, HNRs answering queries for domains on the homenet do not respond to queries from off the homenet unless configured to do so. Exposing services on the homenet for browsing off the homenet creates many opportunities for security issues; as such, even an HNR configured to answer queries from prefixes off the homenet do not provide answers for names of devices on the homenet unless configured to do so. How reachability of names published on the homenet is managed is out of scope for this document: an HNR implementing only this document checks the source address of every query to see if it is within a prefix belonging to the homenet; if not, the HNR does not answer the query.

5.2. Link Names

Each link must have a name. These names are determined using HNCP. Each router will have zero or more wired links, each of which must be labeled. In addition, each router will have zero or more wireless links. Each of these links will be named by the frequency band the link supports, 2.4ghz or 5ghz.

The HNR is named using its manufacturer name. If, as is likely, two or more HNRs from the same manufacturer are present on a homenet, then the HNR name is made up of the manufacturer name plus as many hexadecimal digits as are required from the HNRs link layer addresses so as to disambiguate them.

When shipping multiple HNRs as a kit, manufacturers are advised to arrange that each HNR has a different number in the lowest four bits of the link-layer address. Manufacturers are also advised to print that link layer address, in full, somewhere on the outside of the HNR where it can be seen by the user. Since most HNRs will have more than one interface, the manufacturer should be consistent in choosing which link-layer address is printed on the outside and used to identify the router.

The name given to a link is the name of the HNR, plus a hyphen ('-'), plus name of the interface of that HNR that is attached to the link. In the event that this name must be displayed to the user, this should give the user enough information to figure out which link is being referenced. In the event that the HNR that is providing authoritative service for that link changes, the link name changes. This should only happen if the network topology changes.

If the appearance of a new HNR requires that the name of an existing HNR change, then the names of all the links managed by that existing HNR change to reflect the new name.

5.3. Authoritative name service for the homenet domain

All HNRs must be capable of providing authoritative name service for the homenet domain. HNRs that provide only stateless authoritative service publish the information that is required for hosts to do DNS Service Discovery over DNS, using the local resolver as a DNSSD Discovery Broker.

Some contents are required for the homenet domain, whether it is stateful or stateless.

- o Every link on the homenet has a name that is a subdomain of the homenet domain. The zone associated with the homenet domain contains a delegation for each of these subdomains.
- o In order for DNSSD service discovery to work, a default browsing domain must be published. The default browsing domain is simply the homenet domain.
- o If DNSSD SRP is supported (that is, if stateful authoritative service is present), then an SRV record must be published, along

with a list of available registration zones containing exactly one entry, for the homenet domain ([I-D.sctl-service-registration] section 2).

- o Also if DNSSD SRP is supported, then one or more A and/or AAAA records must be published under the name that the SRV record points to, which should be a single label subdomain of the homenet domain.

Both stateful and stateless authoritative servers provided by HNRs must support DNS Stateful Operations [I-D.ietf-dnsop-session-signal] and DNS Push [I-D.ietf-dnssd-push] for the names for which they are authoritative.

5.4. Authoritative name service for per-link subdomains of the homenet domain

Per-link subdomains of the homenet domain are served by DNSSD Discovery Proxies. Although these proxies generally do caching, no long-lived state is kept by them. DNSSD Discovery Proxies running on HNRs must support DNS Stateful Operations and DNS Push.

5.5. Authoritative name service for the ULA reverse mapping domain

The ULA reverse mapping domain itself is only published if stateful name service is available. It is represented as a single zone, which contains no delegations: every reverse mapping for an address in the ULA prefix is simply published in the ULA zone.

In order to permit registration of reverse mappings in this domain, it must contain an SRV record for the label `_homenet-rrp._tcp` at the top level, pointing to the primary server for the domain.

5.6. Authoritative name service for the RFC1918 reverse mapping domains

If IPv4 service is being provided on the homenet, and if stateful name service is being provided on the homenet, then either one or sixteen reverse mapping zones for the RFC1918 prefix in use must be provided. If more than one RFC1918 prefix is in use, reverse mapping zones for all such prefixes must be provided.

Like the ULA reverse mapping zone, the RFC1918 reverse mapping zones must each contain an SRV record on the label `_homenet-rrp._tcp` at the top level, pointing to the name of the primary server for the zone.

The RFC1918 reverse mapping zone contains the entire address space of the RFC1918 prefix that is in use on the homenet. Section 3 of RFC1918 defines three prefixes that may be used. The homenet will

use all of one of these three prefixes. Of these, the 172.16.0.0 prefix is subdivided on a 12-bit boundary, and therefore must be represented as 16 separate zones. The 10.0.0.0/8 and 192.168.0.0/16 prefixes are each represented as a single zone.

The zone to be updated is therefore the 10.in-addr.arpa zone for all addresses in 10.0.0.0/8, and the 168.192.in-addr.arpa zone for all addresses in 192.168.0.0/16. For addresses in the 172.16.0.0/12 prefix, the zone to be updated is the subdomain of 172.in-addr.arpa that corresponds to bits 8-11 of the prefix: a number between 16 and 31, inclusive.

Also like the ULA zone, the RFC1918 reverse mapping zones contain no delegations: if there is a single zone, then every reverse mapping published for an address in the RFC1918 prefix in use on the homenet is published directly under this zone. If there are sixteen zones, each address is published in its respective zone. Because the zone 172.in-addr.arpa is not available to be served locally, its locally served subdomains are simply served individually with no delegation.

6. Resolution

Name resolution on the homenet must accomplish two tasks: resolving names that are published on the homenet, and resolving names that are published elsewhere. This is accomplished by providing several functional layers.

1. The set of caching nameservers provided by the ISP or ISPs through which the homenet gains access to the global internet, if any (homenets can operate standalone as well).
2. The set of stateful name servers on the homenet that are authoritative for the homenet domain as a whole, and for any reverse mapping zones that are provided by the homenet. This layer is optional, and may or may not be present. If present, it is provided by one or more HNRs on the homenet that support stateful service.
3. The set of stateless name servers on the homenet that are authoritative for the homenet domain as a whole. These are not used if one or more stateful servers are present.
4. The set of stateless DNSSD Discovery Proxies that are authoritative for each of the links in the homenet.
5. A DNS routing proxy. Hereafter we refer to this as the DNS proxy.

The reason that these are described as layers is that it's quite possible that all of the DNS services on the homenet might be provided by a single service listening on port 53; how the request is routed then depends on the question being asked. So the services described as running on HNRs are treated as functional blocks which may be connected internally, if the question being asked can be answered directly by the HNR that received it, or they may be separate name servers running on different HNRs, if the question can be answered within the homenet, or it may be that the HNR receiving the query forwards it to an ISP caching name server.

The routing works as follows. When a request is received (opcode=0, Q/R=0), the DNS proxy looks at the owner name in the question part of the message.

- o If the name is a subdomain of the homenet domain, the query is local.
- o If the name is a subdomain of a locally-valid ULA reverse mapping domain, the query is local.
- o If the name is a subdomain of a locally valid RFC1918 reverse mapping zone, the query is local.
- o If the name is a subdomain of any locally-served zone, as defined in Locally Served DNS Zones [localzones], but is not otherwise identified as local, the response is NXDOMAIN.
- o Otherwise, the query is not local.

Local queries are further divided. If the query is for a link subdomain, the DNS proxy consults the table that maps per-link subdomains to the HNRs that serve them. Either the HNR that serves this link subdomain is the HNR that received the question, or not. If it is, then the DNS proxy passes the query directly to the local DNSSD Discovery Proxy. Otherwise, it forwards the query to the DNSSD Discovery Proxy on the HNR that is providing Discovery Proxy service for that link.

If the query is for the homenet subdomain, and stateful authoritative service for the homenet subdomain is present on the homenet, then either the HNR receiving the query provides stateful authoritative service, or not. If it does, then the query is passed directly to the local authoritative server. If not, then the HNR looks in the table of authoritative servers generated by HNCP and forwards the request to one of these servers. Queries for the reverse mapping zones are handled the same way.

Otherwise, the query is examined to see if it contains an EDNS(0) Provisioning Domain option. If not, it round-robins across the resolvers provided by each ISP in such a way that each ISP is tried in succession, and the same ISP is not asked the same question repeatedly. If the query does contain the EDNS(0) Provisioning Domain option, then that option is used to select which ISP's resolvers are used for the round robin.

6.1. Round Robining

There are several cases above where there may be a choice of servers to which to forward queries. It's assumed that when the query can be satisfied by the HNR that received it, round robing is not required. If there is a specific HNR that is responsible for a particular link, then round robing is likewise not required. However, if the query is for a stateful authoritative server, and the HNR that received it does not provide this service, and there is more than one HNR on the homenet that does provide the service, the HNR that received the query round robs it across the available set of HNRs that could answer it.

Similarly, if the query is to be sent to an ISP's resolver, and the ISP has provided more than one resolver, round robing is done across the set of resolvers provided by that ISP. If the query is to be attempted at every ISP, then that is accomplished by round-robing in such a way that each ISP is tried in succession, rather than all the resolvers at one ISP, and then all the resolvers at the next ISP, and so on.

6.2. Retransmission

For queries that can't be resolved locally by the HNR that received them, retransmission as described in [RFC1035] is performed.

6.3. DNS Stateful Operations and DNS Push

DNS proxies on HNRs are required to support DNS Stateful Operations and DNS Push. When a DNS Push operation is requested on a name that can be satisfied by the HNR that received it, it is handled locally. When such an operation is requested on a name that is local to the homenet, but can't be satisfied by the HNR that received it, a DNS Stateful operation is started with the HNR that is responsible for it.

6.4. Multicast DNS

In addition to consulting the local resolver, hosts on the homenet may attempt to discover services directly using Multicast DNS. HNRs may filter out incoming Multicast DNS queries, forcing the client to do service discovery using the DNS protocol. If such filtering is not done, the client will be able to discover services on the link to which it is attached, but will not be able to discover services elsewhere.

It is believed that all currently-available hosts support DNSSD using the DNS protocol. Support for mDNS on the local link is therefore not required. However, if an mDNS query returns the same answer as the DNS protocol query, this is not expected to be a problem.

6.5. Host behavior

Hosts that support the RA Provisioning Domain option direct queries to the name server(s) of the provisioning domain they will use for communication using the EDNS(0) provisioning domain option. In practice this means that a host that supports PvDs will keep a set of provisioning information for each prefix that it received from the router, and will either choose a prefix to use based on its own criteria, or will attempt to connect using every PvD at once or in sequence. Answers to queries sent for a particular provisioning domain will only be used with source addresses for prefixes that are in that provisioning domain.

7. Publication

Names are published either using Multicast DNS Service Discovery [RFC6762] or DNSSD Service Registration Protocol ([I-D.sctl-service-registration]). Reverse mappings are published using Homenet Reverse Mapping Update Protocol Section 7.2.

7.1. DNSSD Service Registration Protocol

HNRs that provide stateful authoritative service also publish information acquired using DNSSD Service Registration Protocol [I-D.sctl-service-registration]. DNSSD SRP does not explicitly support population of the reverse zone; hosts that wish to provide reverse mapping information must first establish their hostname using DNSSD SRP; once established, the key used to authenticate the DNSSD SRP update is also used to update the reverse name.

Support for SRP provides several advantages over DNSSD Discovery Proxy. First, DNSSD SRP provides a secure way of claiming service names. Second, a claimed name is valid for the entire network

covered by the SRP service, not just an individual link, as is the case with mDNS. Third, SRP does not use multicast, and is therefore more reliable on links with constrained multicast support [I-D.ietf-mboned-ieee802-mcast-problems].

Support for the DNSSD SRP service is not sufficient to achieve full deployment of DNSSD SRP: it is also necessary that services advertise using DNSSD SRP. Requiring such support is out of scope for this document; our goal is simply to specify a way in which DNSSD SRP can be supported on homenets, so that that as adoption of SRP increases on devices providing service, it can actually be used.

7.2. Homenet Reverse Mapping Update Protocol

This is an extension to the DNSSD Service Registration protocol. The purpose is to allow for updates of reverse mappings. Hosts wishing to publish reverse mappings first publish their hostname using DNSSD SRP. When this process has successfully completed, the host can add reverse mappings to the ULA reverse mapping domain and to the RFC1918 reverse mapping domain, if present.

7.2.1. Adding ULA reverse mappings

The host first determines the ULA prefix. If there is more than one ULA prefix active, the ULA prefix with the longest preferred lifetime is used. A ULA prefix can be identified because it matches the prefix `fc00::/7` ([RFC4193] section 3.1). The actual prefix is then the first 48 bits of the advertised prefix or the IP address in that prefix.

Because the ULA reverse mapping zone contains no delegations, all updates go to that one zone. To determine where to send the updates, the host first queries the SRV record under the label `_homenet-rrp._tcp` at the top of the ULA reverse mapping zone. It then uses the name contained in the SRV record to look up A and/or AAAA records to which to send the update.

The update is then signed using SIG(0) with the key that was used for the DNSSD SRP registration. The update is then sent using DNS Update [RFC2136] to one of the IP addresses received during the A/AAAA resolution step. The update is sent using TCP; if a TCP connection to one of the addresses fails, each subsequent address is tried in succession; if none of the addresses is reachable, the update fails, and may be retried after a reasonable period (on the order of an hour) has elapsed.

7.2.2. Adding RFC1918 reverse mappings

RFC1918 reverse mapping updates use the same mechanism as ULA reverse mapping updates. The host must first determine which zone to update, as described earlier in Section 5.6. Once the zone has been determined, the reverse mapping is updated as described in Section 7.2.1.

8. Host Configuration

Each HNR provides a Homenet DNS Proxy. When an HNR provides the DNS resolver IP address to hosts on the link using RA, DHCPv4 or DHCPv6, it provides its own address. The IPv4 address that it provides is a valid IPv4 address on the link to which the host is attached. The IPv6 address it provides is an address in the homenet's ULA prefix that is valid on the link to which the host is attached.

When sending router advertisements, the homenet includes the PvD ID RA option [I-D.ietf-intarea-provisioning-domains] in each RA. Because the PvD ID RA option can only be sent once per RA message, if the homenet is connected to more than one ISP, the prefixes for each ISP must be advertised in different RA options. In this case, the prefix for the ULA should also be sent in a separate RA.

If the configuration received from the ISP includes a Domain Name (DHCPv4) or Domain Search List (DHCPv4 or DHCPv6) option, the domain name provided is used to identify the PvD. In the case of Domain Search List options, if there is more than one, the first one is used. For the ULA prefix, the homenet domain is used to identify the PvD.

In order to facilitate DNSSD bootstrapping, any DHCPv4, DHCPv6 or RA Domain Search List options contain only a single domain name, the homenet domain. This allows hosts to quickly bootstrap DNS Service Discovery using the local domain name, as described in [RFC6763] section 11.

9. Globally Unique Names

Homenets are not required to have globally unique names. Homenets operating according to this specification do not publish names in such a way that they can be resolved by hosts that aren't on the homenet. However, such names are useful for DNSSEC validation.

There are three ways that homenets can get global names:

1. They can be manually configured by the user. How this is done is out of scope for this document.

2. They can publish a delegation with the ISP, using a Homenet Delegation Registration Protocol Section 11.
3. They can publish a delegation with some other provider, using Homenet Delegation Registration Protocol Section 11. How this is configured is out of scope for this document.

Homenets are also not required to support global delegations for reverse mapping of global IPv4 and IPv6 addresses. How this would be done is out of scope for this document.

10. DNSSEC Validation

DNSSEC validation for 'home.arpa' requires installing a per-homenet trust anchor on the host, and is therefore not practical. Validation for locally-served reverse zones for the ULA and RFC1918 addresses would likewise require a trust anchor to be installed on the host, and likewise are not practical.

If DNSSEC validation is to be done for the homenet, the homenet must acquire a global name, and must be provided with a secure delegation. Secure delegations must also be provided from the homenet domain to each of the per-link subdomains.

Each HNR on a homenet generates its own private/public key pair that can serve as a trust anchor. These keys are shared using HNCP [RFC7788]. HNRs MUST NOT use pre-installed keys: each HNR MUST generate its own key. The HNR responsible for authoritative Discovery Proxy service on a particular link signs the zone for that link; delegations from the homenet domain zone to each per-link subdomain zone include a DS record signed by the ZSK of the homenet zone.

10.1. How trust is established

Every HNR has its own public/private key pair. A DS record for each such public key is published in the delegation for the homenet domain. If stateless authoritative service for the homenet zone is being provided, then each HNR signs its own homenet zone. The signed zone should be very stable, although the delegations may change when the network topology changes. The HNR can therefore sign the zone using its private key whenever it changes. Each HNR will have a copy of the zone signed with a different key, but since all of the ZSKs are present in the DS RRset at the delegation point, validation will succeed.

If stateful authoritative service is being provided, the HNR that is acting as primary signs the zone, and all the secondaries serve

copies acquired using zone transfers. If the HNR that is primary goes away, then a secondary becomes primary and signs the zone before beginning to provide service. Again, since all of the HNR's public keys exist in the DS RRset at the delegation, the zone can be validated.

11. Homenet Delegation Registration Protocol

Homenet Delegation Registration Protocol (HDeRP) operates similarly to DNSSD Service Registration Protocol. When a homenet is not connected to an ISP that supports HDeRP, and then an ISP connection becomes available, the HNR that is connected to the ISP determines whether HDeRP is available. This is done by first determining the ISP domain.

If the connection to the ISP is IPv4-only, this will be either the DHCPv4 Domain Name option or, if not present, the only domain name in the DHCPv4 Domain Name Search List option. If the Domain Name Search List option contains more than one name, HDeRP is not supported by the ISP.

If the connection to the ISP is dual-stack or IPv6-only, then the DHCPv6 Domain Search List option obtained through DHCPv6 Prefix Delegation is used. If it is not present, or if it contains more than one domain name, HDeRP is not supported by the ISP.

Once the ISP domain has been discovered, the HNR looks for an SRV record owned by the name `_homenet-derp._tcp` under the ISP domain. If this is not present, HDeRP is not supported. If the SRV record is present, then the HNR looks for A and AAAA records on the hostname provided in the HNR. If present, these are used when attempting the update.

The HNR then constructs a DNS update. The DNS update creates a delegation for the zone `home.arpa`, with a DS record for each HNR on the homenet, containing that HNR's public key. The HNR doing the update lists its key as the first key in the DS RRset. The update is signed using SIG(0) with the private key of the HNR that is constructing it. As with DNSSD SRP, the update includes an Update Lease EDNS(0) option, specifying a key lifetime of a week.

The HNR then attempts to connect to the hostname provided in the SRV record, in a round-robin fashion across the set of IP addresses discovered during the A/AAAA lookup phase. When it has successfully connected, it sends the DNS update.

The HDeRP server validates the update by checking the SIG(0) signature of the update against the first key in the DS RRset. If

the update is successfully validated, then the server generates a domain name and sends a reply back to the HNR on the same TCP connection, including the NOERROR (0) RCODE, and including in the query section the actual domain name that was generated.

This domain name then becomes the homenet name. Subsequent updates use the homenet name rather than 'home.arpa'. It is not necessary that the same HNR do the update; if a different HNR does the update, it lists its public key first in the DS RRset, and signs the update using its private key.

The HDeRP is responsible for removing the delegation if it is not refreshed for the length of its lifetime. HNRs should attempt to refresh the delegation when half the lifetime has experienced, then again at 5/8ths, and again at 7/8ths of the lifetime. If the ISP becomes unavailable, and a different ISP becomes available that supports HDeRP, the homenet should migrate to the new ISP.

12. Using the Local Namespace While Away From Home

This document does not specify a way for service discovery to be performed on the homenet by devices that are not directly connected to a link that is part of the homenet.

13. Expected Host Behavior

It is expected that hosts will fall into one of two categories: hosts that are able to discover DNS-SD browsing domains, and hosts that are not. Hosts that can discover DNS-SD browsing domains can be expected to successfully use service discovery across the entire homenet. Hosts that do not will only be able to discover services on the particular local subnet of the homenet to which they happen to be attached at any given time.

This is not considered to be a problem, since it is understood by the authors that the vast majority of hosts that are capable of doing mDNS discovery are also capable of doing DNS-SD discovery as described in [RFC6763].

14. Management Considerations

This architecture is intended to be self-healing, and should not require management. That said, a great deal of debugging and management can be done simply using the DNS Service Discovery protocol.

15. Privacy Considerations

Privacy is somewhat protected in the sense that names published on the homenet are only visible to devices connected to the homenet. This may be insufficient privacy in some cases.

The privacy of host information on the homenet is left to hosts. Various mechanisms are available to hosts to ensure that tracking does not occur if it is not desired. However, devices that need to have special permission to manage the homenet will inevitably reveal something about themselves when doing so.

16. Security Considerations

There are some clear issues with the security model described in this document, which will be documented in a future version of this section. A full analysis of the avenues of attack for the security model presented here have not yet been done, and must be done before the document is published.

17. IANA considerations

17.1. Homenet Reverse Registration Protocol

IANA is requested to add a new entry to the Service Names and Port Numbers registry for homenet-rrp with a transport type of tcp. No port number is to be assigned. The reference should be to this document, and the Assignee and Contact information should reference the authors of this document. The Description should be as follows:

Availability of Homenet Reverse Registration Protocol service for a given domain is advertised using an SRV record with an owner name of "_homenet-rrp._tcp.<domain>." in that domain, which gives the target host and port where Homenet Reverse Registration service is provided for the named domain.

17.2. Homenet Delegation Registration Protocol

IANA is requested to add a new entry to the Service Names and Port Numbers registry for homenet-derp with a transport type of tcp. No port number is to be assigned. The reference should be to this document, and the Assignee and Contact information should reference the authors of this document. The Description should be as follows:

Availability of Homenet Delegation Registration Protocol service for a given domain is advertised using an SRV record with an owner name of "_homenet-derp._tcp.<domain>." in that domain, which gives the

target host and port where Homenet Delegation Registration service is provided for the named domain.

17.3. Unique Local Address Reserved Documentation Prefix

IANA is requested to add an entry to the IPv6 Special-Purpose Address Registry for the prefix fc00:2001:db8::/48. The Name shall be "Unique Local Address Documentation Prefix." The reference RFC will be this document, once published. The date will be the date the entry was added. All other fields will be the same as for the Documentation prefix, 2001:db8::/32.

18. References

18.1. Normative References

[I-D.ietf-dnsop-session-signal]

Bellis, R., Cheshire, S., Dickinson, J., Dickinson, S., Lemon, T., and T. Pusateri, "DNS Stateful Operations", draft-ietf-dnsop-session-signal-16 (work in progress), September 2018.

[I-D.ietf-dnssd-hybrid]

Cheshire, S., "Discovery Proxy for Multicast DNS-Based Service Discovery", draft-ietf-dnssd-hybrid-08 (work in progress), March 2018.

[I-D.ietf-dnssd-push]

Pusateri, T. and S. Cheshire, "DNS Push Notifications", draft-ietf-dnssd-push-15 (work in progress), September 2018.

[I-D.ietf-intarea-provisioning-domains]

Pfister, P., Vyncke, E., Pauly, T., Schinazi, D., and W. Shao, "Discovering Provisioning Domain Names and Data", draft-ietf-intarea-provisioning-domains-03 (work in progress), October 2018.

[I-D.sctl-service-registration]

Cheshire, S. and T. Lemon, "Service Registration Protocol for DNS-Based Service Discovery", draft-sctl-service-registration-02 (work in progress), July 2018.

[localzones]

Internet Assigned Numbers Authority, "Locally-Served DNS Zones", n.d., <<https://www.iana.org/assignments/locally-served-dns-zones/locally-served-dns-zones.xhtml>>.

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC2136] Vixie, P., Ed., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, DOI 10.17487/RFC2136, April 1997, <<https://www.rfc-editor.org/info/rfc2136>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, Ed., "Framework for Content Distribution Network Interconnection (CDNI)", RFC 7336, DOI 10.17487/RFC7336, August 2014, <<https://www.rfc-editor.org/info/rfc7336>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <<https://www.rfc-editor.org/info/rfc7556>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.

[RFC8375] Pfister, P. and T. Lemon, "Special-Use Domain 'home.arpa.'", RFC 8375, DOI 10.17487/RFC8375, May 2018, <<https://www.rfc-editor.org/info/rfc8375>>.

18.2. Informative References

[I-D.ietf-mboned-ieee802-mcast-problems]
Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-02 (work in progress), August 2018.

Authors' Addresses

Ted Lemon
Nibbhaya Consulting
P.O. Box 958
Brattleboro, Vermont 05301
United States of America

Email: mellon@fugue.com

Daniel Migault
Ericsson
8400 boulevard Decarie
Montreal, QC H4P 2N2
Canada

Email: daniel.migault@ericsson.com

Stuart Cheshire
Apple Inc.
1 Infinite Loop
Cupertino, California 95014
USA

Phone: +1 408 974 3207
Email: cheshire@apple.com

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: May 12, 2019

S. Cheshire
Apple Inc.
T. Lemon
Nibbhaya Consulting
November 8, 2018

Unicast Service Discovery Autoconfiguration
draft-sctl-dnssd-unicast-autoconfig-00

Abstract

This document considers the requirements for adding a Thread mesh to an existing home network, where the infrastructure of that existing home network was designed with no prior knowledge of Thread, and provides no special or unusual facilities designed to support this.

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 <https://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 12, 2019.

Copyright Notice

Copyright (c) 2018 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

[Authors' note: As an initial draft, in places this document presents several alternatives that are being considered. We invite feedback and comments to help evolve this document.]

Because multicast can be inefficient and unreliable [Mcast], work is taking place to enable DNS-Based Service Discovery [RFC6763] to operate with less reliance on multicast [Roadmap]. One current target use case for this work is Thread [Thread] wireless mesh networking.

Thread wireless mesh networking uses IEEE 802.15.4 radios, which use little power, and are suitable for battery-powered devices. The Thread protocol organizes the network nodes into a mesh, typically with a Thread border router that connects the mesh to the home network. For the purposes of this document we will refer to the home network, be it Ethernet, or Wi-Fi, or both, or other similar technologies, simply as the home network. The home network forms a backbone to which one or more Thread mesh networks connect via Thread border routers.

Existing work describes how DNS-Based Service Discovery can be performed using unicast on such a network. Devices on the Thread mesh offering services use Service Registration Protocol [RegProt] to register their services at a Service Registration Server. Devices seeking to discover these services send unicast queries to the Service Registration Server using unicast DNS [RFC1034] [RFC1035] for single individual queries, and using DNS Push Notifications [Push] where ongoing change notification is required.

Certain configuration information is required for this to work. Devices on the Thread mesh offering services need to know what names to use when registering those services, and to what address they should send their service registrations. Devices seeking to discover these services need to know what names to use when constructing their queries, and to what address they should send those queries. In addition, IPv6 address prefixes need to be chosen and configured for both the home network and the Thread mesh network(s), and communicated, in order to facilitate unicast communication between clients and the services they have discovered.

For proof-of-concept experiments, the necessary information can be configured manually, and this has been done successfully. For deployment, we need to determine how the necessary information will be learned and configured automatically in real-world scenarios.

The Thread wireless mesh protocol includes mechanisms to perform configuration tasks on the mesh, like electing a lead router, and communicating this information to devices on the Thread mesh. This existing mechanism can be extended fairly simply to facilitate the necessary Service Registration Protocol configuration tasks. The Service Registration Protocol [RegProt] specification document advocates that if a device offering a service has no information regarding the domain in which to register that service, it should use the special use domain name [RFC6761] "services.arpa" to indicate that the Service Registration Server should substitute a domain of its choice, and that same mechanism is recommended in this case.

On the home network side of the Thread border router, there are several possibilities. The necessary configuration tasks could be handled by the home network's main gateway, by a collection of Homenet routers using HNCP, or independently by the Thread border router.

1.1. Configuration by Main Gateway

The home network's main gateway could handle the necessary configuration tasks.

The main gateway could be responsible for selecting IPv6 address prefixes for each of the links in the network, and communicating that information to the relevant routers, perhaps using DHCPv6 prefix delegation.

The information about what domain name to use for service discovery can be communicated to client devices on the home network using DHCP or IPv6 router advertisement options. Currently this is done using the respective "DNS search list" options, though new options for this specific purpose could be defined in the future. If the user has a registered globally unique domain name for this purpose and the main gateway is configured with this information, then that domain name can be communicated to client devices. In the absence of a registered globally unique domain name the special-use domain name [RFC6761] "home.arpa" [RFC8375] should be used as a reasonable out-of-the-box default.

Similarly, the information about what DNS recursive resolver to use can be communicated to client devices on the home network using DHCP or IPv6 router advertisement options. If the main gateway configures its own address as the DNS recursive resolver for clients to use, it can ensure that operations using "home.arpa" are handled appropriately. Sending queries for names within "home.arpa" to public recursive resolvers on the Internet will not yield useful results, because names within "home.arpa" are not globally unique.

They are unique only within the local network, and hence queries for those names need to be handled within the local network.

1.2. Configuration using HNCP

A complex home network with multiple links and multiple routers could be managed using HNCP. However, at this time, this remains a future possibility, since it is likely to be some time before HNCP is widely used.

1.3. Self-Configuration by Thread Border Routers

The previous two scenarios assume that the home network's main gateway, or its HNCP mesh, has specific capabilities to configure and support the use of unicast DNS service discovery.

An alternative scenario is to consider the case where a Thread border router is added to an existing home network, which has no special mechanisms in place to support this operation.

The remainder of this document explores this scenario.

One possibility to keep in mind is that in this scenario, adding one or more Thread border routers to an existing home network that doesn't itself use HNCP, the Thread border router(s) themselves could use HNCP as the protocol to communicate between each other to coordinate their operation on the network.

2. Adding Thread Mesh to Existing Home Network

This section explores the requirements for connecting a Thread mesh, via a Thread border router, to a typical home network. For the purposes of this document, it is assumed that the existing network infrastructure is fixed and cannot be changed. Changes or new functionality may be implemented as required in the Thread devices on the Thread mesh, in the Thread border router, or in the devices on the home network that will be communicating with the Thread devices. Since this document assumes no changes to the existing network infrastructure, it is necessary to state the assumptions about that existing network infrastructure.

We consider a typical home network to be a single multicast/broadcast domain. If there are multiple Ethernet switches or Wi-Fi access points, they are configured so that together they provide a single logical link. If there is a NAT gateway, it is at the network egress point. (A NAT gateway on the path between two devices on a home network makes communication between those two devices considerably more complicated, and this document does not address that case.)

In order to add a Thread mesh usefully to an existing home network, several things need to be accomplished. The goal is to accomplish these objectives without requiring changes to the existing infrastructure on that home network.

1. Delivery of unicast traffic in both directions, from home network to Thread mesh, and from Thread mesh to home network.
2. Enabling services offered by devices on the Thread mesh to be discovered by clients seeking those services.
3. Enabling services offered by devices on the home network to be discovered by clients on the Thread mesh seeking those services.

2.1. Unicast Delivery

If HNCP were in use on the network, then Thread border routers could participate and use HNCP to manage their configuration.

In the absence of HNCP, Thread border routers need a way to self-configure, without assistance from the home network's existing infrastructure.

What is proposed is that Thread border routers select a 32-bit random number, and use that to construct an IPv6 ULA prefix for their connected mesh, which is very likely to be unique in that home. The Thread border router then advertises reachability to that IPv6 ULA prefix onto the home network using IPv6 Router Advertisements. In principle, this can be done independently of whatever other IPv6 prefixes, if any, are being advertised on the home network by the home network's existing main gateway. [It has been reported, however, that there are at least some client devices that do not properly handle receiving multiple independent IPv6 Router Advertisements like this, so some investigation and bug fixing may be required to make this work.]

In the case where there are multiple independent Thread border routers connected to the home network, serving separate Thread meshes, we want to avoid the situation where two different Thread border routers choose the same randomly-selected IPv6 ULA prefix. This can be achieved by having the Thread border routers listen for IPv6 Router Advertisements before selecting their IPv6 ULA prefix. If a Thread border router receives IPv6 Router Advertisements offering reachability to its IPv6 ULA prefix via a different path, then this indicates that an inadvertent duplication may have occurred, and the Thread border router should select a different IPv6 ULA prefix for its mesh.

2.2. Discovery of Services on the Thread Mesh

To facilitate unicast discovery of services on the Thread mesh, four things need to be determined:

1. How a device on the Thread mesh, offering services, knows what parent domain name to use when registering services.
2. How that device knows to what address its service registrations should be sent (if the name does not fall under a registered globally unique domain name).
3. How a client device, on the Thread mesh or the home network, seeking services, knows what parent domain name to use querying to discover services.
4. How that device knows to what address its unicast service discovery queries should be sent (if the name does not fall under a registered globally unique domain name).

Devices on the Thread mesh should register services using the parent domain "services.arpa". This indicates that the Service Registration Server should automatically substitute an appropriate domain.

The Thread mesh management protocol can be used to configure devices on the Thread mesh with the address to which they should send their service registrations.

The Thread border router needs to communicate, to devices on the home network, how they can discover services on the Thread mesh.

This involves communicating the service discovery domain. In principle, this could be a registered globally unique domain name, in which case the normal DNS delegation mechanism using NS records allows the client to discover what server is authoritative for those names. In many cases though, the Thread border router will not have a registered globally unique domain name allocated. To provide out-of-the-box automatic operation, the Thread border router needs to be able to generate its own locally unique name to use. The special use domain name "local" is not suitable, because of its implied semantics that these names are resolved using link-local multicast DNS [RFC6762]. The special use domain name "home.arpa" is not suitable, because of its implied coordination via HNCP, and the home network's main gateway may not support HNCP [RFC8375]. To provide out-of-the-box automatic operation, this document proposes a new special use domain name "adhoc.arpa" for this purpose. By default a Thread border router will use the name "thread.adhoc.arpa". If this name is

already in use on the home network, then a new unique name will be selected, such as "thread-2.adhoc.arpa".

The Thread border router needs to communicate the service discovery domain to peers on the home network. In the case that the service discovery domain falls under the "adhoc.arpa" name, the Thread border router also needs to communicate that queries for these names need to be sent to the Thread border router directly, not to the client's default DNS recursive resolver.

Three alternatives are being considered

1. Use link-local Multicast DNS queries and records to convey the service discovery domain, and optionally the address to which queries should be sent.
2. Define a new IPv6 router advertisement option to communicate the service discovery domain, and optionally the address to which queries should be sent.
3. Add this information to the Multiple Provisioning Domain Router Advertisement option [RFC7556] [MPvD].

One question to answer is whether the Multicast DNS records or IPv6 router advertisement options would directly convey the domain name to use for service discovery, or a base name used to derive domain enumeration queries of the form lb._dns-sd._udp.<domain> [RFC6763].

Another question is whether to use a single Multicast DNS record or IPv6 router advertisement option that communicates both the domain name and the address to use for queries, or a pair of records/options, one carrying the name to use for service discovery, and a second, if necessary, associating an "adhoc.arpa" name with the address to use for those queries.

With the appropriate configuration methods defined, and implemented on client devices, client devices on the home network would discover additional domains to use for service discovery, and send appropriate service discovery queries to Thread border routers on the home network.

The same discovery domain, and optionally the address to which queries should be sent, is communicated to client devices on the Thread mesh using the Thread mesh management protocol.

2.3. Discovery of Services on the Home Network

To facilitate devices on the Thread mesh discovering services offered on the home network, advertised using Multicast DNS, a Discovery Proxy [DisProx] is implemented in the Thread border router.

As above in Section 2.2 the Thread mesh management protocol is used to communicate a discovery domain, and the address to which queries should be sent for that discovery domain, to client devices on the Thread mesh.

The address in this case is the address of the Thread border router. The discovery domain could be some generated unique name under "adhoc.arpa", or it could be some fixed special use domain name. The fixed name could be a simple fixed string like "lan.arpa", or it could be a special reserved name under "adhoc.arpa" such as "services.adhoc.arpa". The latter is probably preferred because it avoids having to request multiple special use domain names [RFC6761]. Alternatively, we could organize all the required special names such that they fall under a single reserved special use domain name "services.arpa."

When the Thread border router receives a query for a name under this discovery domain, it uses the Discovery Proxy mechanism [DisProx] to perform Multicast DNS queries on behalf of the client, returning the results to the client.

3. Security Considerations

As an informational document, this document introduces no new Security Considerations of its own. The various referenced documents each describe their own relevant Security Considerations as appropriate.

4. Domain Name Reservation Considerations

As currently envisaged, this document may end up requesting a special use domain name [RFC6761]. If so, once the special properties are fully determined, this section will be populated with the appropriate text.

5. Informative References

- [DisProx] Cheshire, S., "Discovery Proxy for Multicast DNS-Based Service Discovery", draft-ietf-dnssd-hybrid-08 (work in progress), March 2018.
- [Mcast] Perkins, C., McBride, M., Stanley, D., Kumari, W., and J. Zuniga, "Multicast Considerations over IEEE 802 Wireless Media", draft-ietf-mboned-ieee802-mcast-problems-03 (work in progress), October 2018.
- [MPvD] Korhonen, J., Krishnan, S., and S. Gundavelli, "Support for multiple provisioning domains in IPv6 Neighbor Discovery Protocol", draft-ietf-mif-mpvd-ndp-support-03 (work in progress), February 2016.
- [Push] Pusateri, T. and S. Cheshire, "DNS Push Notifications", draft-ietf-dnssd-push-14 (work in progress), March 2018.
- [RegProt] Cheshire, S. and T. Lemon, "Service Registration Protocol for DNS-Based Service Discovery", draft-sctl-service-registration-00 (work in progress), July 2017.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <<https://www.rfc-editor.org/info/rfc6761>>.

- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC7556] Anipko, D., Ed., "Multiple Provisioning Domain Architecture", RFC 7556, DOI 10.17487/RFC7556, June 2015, <<https://www.rfc-editor.org/info/rfc7556>>.
- [RFC8375] Pfister, P. and T. Lemon, "Special-Use Domain 'home.arpa.'", RFC 8375, DOI 10.17487/RFC8375, May 2018, <<https://www.rfc-editor.org/info/rfc8375>>.
- [Roadmap] Cheshire, S., "Service Discovery Road Map", draft-cheshire-dnssd-roadmap-03 (work in progress), October 2018.
- [Thread] "Thread Wireless Mesh Networking", <<https://www.threadgroup.org/>>.

Authors' Addresses

Stuart Cheshire
Apple Inc.
One Apple Park Way
Cupertino, California 95014
USA

Phone: +1 (408) 996-1010
Email: cheshire@apple.com

Ted Lemon
Nibbhaya Consulting
P.O. Box 958
Brattleboro, Vermont 05301
United States of America

Email: mellon@fugue.com