

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 20, 2018

J. Chroboczek
IRIF, University of Paris-Diderot
January 16, 2018

Homenet profile of the Babel routing protocol
draft-ietf-homenet-babel-profile-05

Abstract

This document defines the subset of the Babel routing protocol and its extensions that a Homenet router must implement, as well as the interactions between HNCP and Babel.

Status of This Memo

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

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

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

This Internet-Draft will expire on July 20, 2018.

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

Table of Contents

1. Introduction	2
1.1. Requirement Language	2
1.2. Background	2
2. The Homenet profile of Babel	3
2.1. Requirements	3
2.2. Non-requirements	5
3. Interactions between HNCP and Babel	5
3.1. Requirements	6
3.2. Non-requirements	6
4. Security Considerations	7
5. IANA Considerations	7
6. Acknowledgments	7
7. References	8
7.1. Normative References	8
7.2. Informative References	8
Author's Address	9

1. Introduction

The core of the Homenet protocol suite consists of HNCP [RFC7788], a protocol used for flooding configuration information and assigning prefixes to links, combined with the Babel routing protocol [RFC6126bis]. Babel is an extensible, flexible and modular protocol: minimal implementations of Babel have been demonstrated that consist of a few hundred lines of code, while the "large" implementation includes support for a number of extensions and consists of over ten thousand lines of C code.

This document consists of two parts. The first specifies the exact subset of the Babel protocol and its extensions that is required by an implementation of the Homenet protocol suite. The second specifies how HNCP interacts with Babel.

1.1. Requirement Language

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 RFC 2119 [RFC2119].

1.2. Background

The Babel routing protocol and its extensions are defined in a number of documents:

- o RFC 6126bis [RFC6126bis] defines the Babel routing protocol. It allows Babel's control data to be carried over either link-local IPv6 or IPv4, and in either case allows announcing both IPv4 and IPv6 routes. It leaves link cost estimation, metric computation and route selection to the implementation. Distinct implementations of RFC 6126bis Babel will interoperate, in the sense that they will maintain a set of loop-free forwarding paths. However, if they implement conflicting options, they might not be able to exchange a full set of routes; in the worst case, an implementation that only implements the IPv6 subset of the protocol and an implementation that only implements the IPv4 subset of the protocol will not exchange any routes. In addition, if implementations use conflicting route selection policies, persistent oscillations might occur.
- o The informative Appendix A of RFC 6126bis suggests a simple and easy to implement algorithm for cost and metric computation that has been found to work satisfactorily in a wide range of topologies.
- o While RFC 6126bis does not provide an algorithm for route selection, its Section 3.6 suggests selecting the route with smallest metric with some hysteresis applied. An algorithm that has been found to work well in practice is described in Section III.E of [DELAY-BASED].
- o Five RFCs and Internet-Drafts define optional extensions to Babel: HMAC-based authentication [RFC7298], source-specific routing [BABEL-SS], delay-based routing [BABEL-RTT] and ToS-specific routing [ToS-SPECIFIC]. All of these extensions interoperate with the core protocol as well as with each other.

2. The Homenet profile of Babel

2.1. Requirements

REQ1: a Homenet implementation of Babel MUST encapsulate Babel control traffic in IPv6 packets sent to the IANA-assigned port 6696 and either the IANA-assigned multicast group ff02::1:6 or to a link-local unicast address.

Rationale: since Babel is able to carry both IPv4 and IPv6 routes over either IPv4 or IPv6, choosing the protocol used for carrying control traffic is a matter of preference. Since IPv6 has some features that make implementations somewhat simpler and more reliable (notably link-local addresses), we require carrying control data over IPv6.

REQ2: a Homenet implementation of Babel MUST implement the IPv6 subset of the protocol defined in the body of RFC 6126bis.

Rationale: support for IPv6 routing is an essential component of the Homenet architecture.

REQ3: a Homenet implementation of Babel SHOULD implement the IPv4 subset of the protocol defined in the body of RFC 6126bis. Use of other techniques for acquiring IPv4 connectivity (such as multiple layers of NAT) is strongly discouraged.

Rationale: support for IPv4 will likely remain necessary for years to come, and even in pure IPv6 deployments, including code for supporting IPv4 has very little cost. Since HNCP makes it easy to assign distinct IPv4 prefixes to the links in a network, it is not necessary to resort to multiple layers of NAT, with all of its problems.

REQ4: a Homenet implementation of Babel MUST implement source-specific routing for IPv6, as defined in draft-ietf-babel-source-specific [BABEL-SS].

Rationale: source-specific routing is an essential component of the Homenet architecture. Source-specific routing for IPv4 is not required, since HNCP arranges things so that a single non-specific IPv4 default route is announced (Section 6.5 of [RFC7788]).

REQ5: a Homenet implementation of Babel MUST use metrics that are of a similar magnitude to the values suggested in Appendix A of RFC 6126bis. In particular, it SHOULD assign costs that are no less than 256 to wireless links, and SHOULD assign costs between 32 and 196 to lossless wired links.

Rationale: if two implementations of Babel choose very different values for link costs, combining routers from different vendors will cause sub-optimal routing.

REQ6: a Homenet implementation of Babel SHOULD distinguish between wired and wireless links; if it is unable to determine whether a link is wired or wireless, it SHOULD make the worst-case hypothesis that the link is wireless. It SHOULD dynamically probe the quality of wireless links and derive a suitable metric from its quality estimation. The algorithm described in Appendix A of RFC 6126bis MAY be used.

Rationale: support for wireless transit links is a "killer feature" of Homenet, something that is requested by our users and easy to explain to our bosses. In the absence of dynamically

computed metrics, the routing protocol attempts to minimise the number of links crossed by a route, and therefore prefers long, lossy links to shorter, lossless ones. In wireless networks, "hop-count routing is worst-path routing".

2.2. Non-requirements

NR1: a Homenet implementation of Babel MAY perform route selection by applying hysteresis to route metrics, as suggested in Section 3.6 of RFC 6126bis and described in detail in Section III.E of [BABEL-RTT]. However, it MAY simply pick the route with the smallest metric.

Rationale: hysteresis is only useful in congested and highly dynamic networks. In a typical home network, stable and uncongested, the feedback loop that hysteresis compensates for does not occur.

NR2: a Homenet implementation of Babel MAY include support for other extensions to the protocol, as long as they are known to interoperate with both the core protocol and source-specific routing.

Rationale: a number of extensions to the Babel routing protocol have been defined over the years; however, they are useful in fairly specific situations, such as routing over global-scale overlay networks [BABEL-RTT] or multi-hop wireless networks with multiple radio frequencies [BABEL-Z]. Hence, with the exception of source-specific routing, no extensions are required for Homenet.

3. Interactions between HNCP and Babel

The Homenet architecture cleanly separates between configuration, which is done by HNCP, and routing, which is done by Babel. While the coupling between the two protocols is deliberately kept to a minimum, some interactions are unavoidable.

All the interactions between HNCP and Babel consist of HNCP causing Babel to perform an announcement on its behalf (under no circumstances does Babel cause HNCP to perform an action). How this is realised is an implementation detail that is outside the scope of this document; while it could conceivably be done using a private communication channel between HNCP and Babel, in existing implementations HNCP installs a route in the operating system's kernel which is later picked up by Babel using the existing redistribution mechanisms.

3.1. Requirements

REQ7: if an HNCP node receives a DHCPv6 prefix delegation for prefix P and publishes an External-Connection TLV containing a Delegated-Prefix TLV with prefix P and no Prefix-Policy TLV, then it MUST announce a source-specific default route with source prefix P over Babel.

Rationale: source-specific routes are the main tool that Homenet uses to enable optimal routing in the presence of multiple IPv6 prefixes. External connections with non-trivial prefix policies are explicitly excluded from this requirement, since their exact behaviour is application-specific.

REQ8: if an HNCP node receives a DHCPv4 lease with an IPv4 address and wins the election for NAT gateway, then it MUST act as a NAT gateway and MUST announce a (non-specific) IPv4 default route over Babel.

Rationale: the Homenet architecture does not use source-specific routing for IPv4; instead, HNCP elects a single NAT gateway and publishes a single default route towards that gateway ([RFC7788] Section 6.5).

REQ9: if an HNCP node assigns a prefix P to an attached link and announces P in an Assigned-Prefix TLV, then it MUST announce a route towards P over Babel.

Rationale: prefixes assigned to links must be routable within the Homenet.

3.2. Non-requirements

NR3: an HNCP node that receives a DHCPv6 prefix delegation MAY announce a non-specific IPv6 default route over Babel in addition to the source-specific default route mandated by requirement REQ7.

Rationale: since the source-specific default route is more specific than the non-specific default route, the former will override the latter if all nodes implement source-specific routing. Announcing an additional non-specific route is allowed, since doing that causes no harm and might simplify operations in some circumstances, e.g. when interoperating with a routing protocol that does not support source-specific routing.

NR4: an HNCP node that receives a DHCPv4 lease with an IPv4 address and wins the election for NAT gateway SHOULD NOT announce a source-specific IPv4 default route.

Homenet does not require support for IPv4 source-specific routing. Announcing IPv4 source-specific routes will not cause routing pathologies (blackholes or routing loops), but it might cause packets sourced in different parts of the Homenet to follow different paths, with all the confusion that this entails.

4. Security Considerations

Both HNCP and Babel carry their control data in IPv6 packets with a link-local source address, and implementations are required to drop packets sent from a global address. Hence, they are only susceptible to attacks from a directly connected link on which the HNCP and Babel implementations are listening.

The security of a Homenet network relies on having a set of "Internal", "Ad Hoc" and "Hybrid" interfaces (Section 5.1 of [RFC7788]) that are assumed to be connected to links that are secured at a lower layer. HNCP and Babel packets are only accepted when they originate on these trusted links. "External" and "Guest" interfaces are connected to links that are not trusted, and any HNCP or Babel packets that are received on such interfaces are ignored. ("Leaf" interfaces are a special case, since they are connected to trusted links but HNCP and Babel traffic received on such interfaces is ignored.) This implies that the security of a Homenet network depends on the reliability of the border discovery procedure described in Section 5.3 of [RFC7788].

If untrusted links are used for transit, which is NOT RECOMMENDED, then any HNCP and Babel traffic that is carried over such links MUST be secured using an upper-layer security protocol. While both HNCP and Babel support cryptographic authentication, at the time of writing no protocol for autonomous configuration of HNCP and Babel security has been defined.

5. IANA Considerations

This document requires no actions from IANA.

6. Acknowledgments

A number of people have helped with defining the requirements listed in this document. I am especially indebted to Barbara Stark, Markus Stenberg, and Stephen Farrell.

7. References

7.1. Normative References

- [BABEL-SS] Boutier, M. and J. Chroboczek, "Source-Specific Routing in Babel", draft-ietf-babel-source-specific-01 (work in progress), August 2017.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997.
- [RFC6126bis] Chroboczek, J. and D. Schinazi, "The Babel Routing Protocol", Internet Draft draft-ietf-babel-rfc6126bis-04, October 2017.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016.

7.2. Informative References

- [BABEL-RTT] Jonglez, B. and J. Chroboczek, "Delay-based Metric Extension for the Babel Routing Protocol", draft-jonglez-babel-rtt-extension-01 (work in progress), May 2015.
- [BABEL-Z] Chroboczek, J., "Diversity Routing for the Babel Routing Protocol", draft-chroboczek-babel-diversity-routing-01 (work in progress), February 2016.
- [DELAY-BASED] Jonglez, B. and J. Chroboczek, "A delay-based routing metric", March 2014.
- Available online from <http://arxiv.org/abs/1403.3488>
- [RFC7298] Ovsienko, D., "Babel Hashed Message Authentication Code (HMAC) Cryptographic Authentication", RFC 7298, July 2014.
- [ToS-SPECIFIC] Chouasne, G., "<https://tools.ietf.org/id/draft-chouasne-babel-tos-specific-00.xml>", draft-chouasne-babel-tos-specific-00 (work in progress), July 2017.

Author's Address

Juliusz Chroboczek
IRIF, University of Paris-Diderot
Case 7014
75205 Paris Cedex 13
France

Email: jch@irif.fr

Network Working Group
Internet-Draft
Updates: RFC7788 (if approved)
Intended status: Standards Track
Expires: March 5, 2018

P. Pfister
Cisco Systems
T. Lemon
Nominum, Inc.
September 1, 2017

Special Use Domain 'home.arpa.'
draft-ietf-homenet-dot-14

Abstract

This document specifies the behavior that is expected from the Domain Name System with regard to DNS queries for names ending with '.home.arpa.', and designates this domain as a special-use domain name. 'home.arpa.' is designated for non-unique use in residential home networks. Home Networking Control Protocol (HNCP) is updated to use the 'home.arpa.' domain instead of '.home'.

Status of This Memo

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

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

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

This Internet-Draft will expire on March 5, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Requirements Language	3
3. General Guidance	3
4. Domain Name Reservation Considerations	4
5. Updates to Home Networking Control Protocol	6
6. Security Considerations	7
6.1. Local Significance	7
6.2. Insecure Delegation	8
6.3. Bypassing Manually Configured Resolvers	8
7. Delegation of 'home.arpa.'	8
8. IANA Considerations	9
9. Acknowledgments	9
10. References	9
10.1. Normative References	9
10.2. Informative References	10
Authors' Addresses	11

1. Introduction

Users and devices within a home network (hereafter "homenet") require devices and services to be identified by names that are unique within the boundaries of the homenet [RFC7368]. The naming mechanism needs to function without configuration from the user. While it may be possible for a name to be delegated by an ISP, homenets must also function in the absence of such a delegation. This document reserves the name 'home.arpa.' to serve as the default name for this purpose, with with a scope limited to each individual homenet.

This document corrects an error in [RFC7788], replacing '.home' with 'home.arpa.' as the default domain-name for homenets. '.home' had been selected as the most user-friendly option. However, there are existing uses of '.home' that may be in conflict with this use: evidence indicates that '.home' queries frequently leak out and reach the root name servers [ICANN1] [ICANN2].

In addition, it's necessary, for compatibility with DNSSEC (Section 6), that an insecure delegation ([RFC4035] section 4.3) be present for the name. There is an existing process for allocating names under '.arpa.' [RFC3172]. No such process is available for requesting a similar delegation in the root at the request of the IETF, which does not administer that zone. As a result, all unregistered uses of '.home' (that is, all current uses at the time

of this document's publication), particularly as specified in RFC7788, are deprecated.

This document registers the domain 'home.arpa.' as a special-use domain name [RFC6761] and specifies the behavior that is expected from the Domain Name System with regard to DNS queries for names whose rightmost non-terminal labels are 'home.arpa.'. Queries for names ending with '.home.arpa.' are of local significance within the scope of a homenet, meaning that identical queries will result in different results from one homenet to another. In other words, a name ending in '.home.arpa.' is not globally unique.

Although this document makes specific reference to RFC7788, it is not intended that the use of 'home.arpa.' be restricted solely to networks where HNCP is deployed; it is rather the case that 'home.arpa.' is the correct domain for uses like the one described for '.home' in RFC7788: local name service in residential homenets.

2. Requirements Language

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.

3. General Guidance

The domain name 'home.arpa.' is to be used for naming within residential homenets. Names ending with '.home.arpa.' reference a locally-served zone, the contents of which are unique only to a particular homenet, and are not globally unique. Such names refer to nodes and/or services that are located within a homenet (e.g., a printer, or a toaster).

DNS queries for names ending with '.home.arpa.' are resolved using local resolvers on the homenet. Such queries MUST NOT be recursively forwarded to servers outside the logical boundaries of the homenet.

Some service discovery user interfaces that are expected to be used on homenets conceal information such as domain names from end users. However, it is still expected that in some cases, users will need to see, remember, and even type, names ending with '.home.arpa.'. The working group hopes that this name will in some way indicate to as many readers as possible that such domain names are referring to devices in the home, but we recognize that it is an imperfect solution.

4. Domain Name Reservation Considerations

This section specifies considerations for systems involved in domain name resolution when resolving queries for names ending with `'home.arpa.'`. Each item in this section addresses some aspect of the DNS or the process of resolving domain names that would be affected by this special use allocation. Detailed explanations of these items can be found in [RFC6761], Section 5.

1. Users can use names ending with `'home.arpa.'` just as they would use any other domain name. The `'home.arpa.'` name is chosen to be readily recognized by users as signifying that the name is addressing a service on the homenet to which the user's device is connected.
2. Application software SHOULD NOT treat names ending in `'home.arpa.'` differently than other names. In particular, there is no basis for trusting names that are subdomains of `'home.arpa.'` (see Section 6).
3. Name resolution APIs and libraries MUST NOT recognize names that end in `'home.arpa.'` as special and MUST NOT treat them as having special significance, except that it may be necessary that such APIs not bypass the locally configured recursive resolvers.

One or more IP addresses for recursive DNS servers will usually be supplied to the client through router advertisements or DHCP. For an administrative domain that uses subdomains of `'home.arpa.'`, such as a homenet, the recursive resolvers provided by that domain will be able to answer queries for subdomains of `'home.arpa.'`; other resolvers will not, or will provide answers that are not correct within that administrative domain.

A host that is configured to use a resolver other than one that has been provided by the local network may be unable to resolve, or may receive incorrect results for, subdomains of `'home.arpa.'`. In order to avoid this, it is permissible that hosts use the locally-provided resolvers for resolving `'home.arpa.'` even when they are configured to use other resolvers.

4.

- A. Recursive resolvers at sites using `'home.arpa.'` MUST transparently support DNSSEC queries: queries for DNSSEC records and queries with the DO bit set ([RFC4035] section 3.2.1). While validation is not required, it is strongly encouraged: a caching recursive resolver that does not validate answers that can be validated may cache invalid

data. This in turn will prevent validating stub resolvers from successfully validating answers.

- B. Unless configured otherwise, recursive resolvers and DNS proxies MUST behave as described in Locally Served Zones ([RFC6303] Section 3). That is, queries for 'home.arpa.' and subdomains of 'home.arpa.' MUST NOT be forwarded, with one important exception: a query for a DS record with the DO bit set MUST return the correct answer for that question, including correct information in the authority section that proves that the record is nonexistent.

So for example a query for the NS record for 'home.arpa.' MUST NOT result in that query being forwarded to an upstream cache nor to the authoritative DNS server for '.arpa.'. However, as necessary to provide accurate authority information, a query for the DS record MUST result in whatever queries are necessary being forwarded; typically, this will just be a query for the DS record, since the necessary authority information will be included in the authority section of the response if the DO bit is set.

- C. In addition to the behavior specified above, recursive resolvers that can be used in a homenet MUST be configurable to forward queries for 'home.arpa.' and subdomains of 'home.arpa.' to an authoritative server for 'home.arpa.'. This server will provide authoritative data for 'home.arpa.' within a particular homenet. The special handling for DS records for the 'home.arpa.' delegation is still required.

It is permissible to combine the recursive resolver function for general DNS lookups with an authoritative resolver for 'home.arpa.'; in this case, rather than forwarding queries for subdomains of 'home.arpa.' to an authoritative server, the resolver answers them authoritatively. The behavior with respect to forwarding queries specifically for 'home.arpa.' remains the same.

- 5. No special processing of 'home.arpa.' is required for authoritative DNS server implementations. It is possible that an authoritative DNS server might attempt to check the authoritative servers for 'home.arpa.' for a delegation beneath that name before answering authoritatively for such a delegated name. In such a case, because the name always has only local significance there will be no such delegation in the 'home.arpa.' zone, and so the server would refuse to answer authoritatively for such a zone. A server that implements this sort of check MUST be

configurable so that either it does not do this check for the 'home.arpa.' domain, or it ignores the results of the check.

6. DNS server operators MAY configure an authoritative server for 'home.arpa.' for use in homenets and other home networks. The operator for the DNS servers authoritative for 'home.arpa.' in the global DNS will configure any such servers as described in Section 7.
 7. 'home.arpa.' is a subdomain of the 'arpa' top-level domain, which is operated by IANA under the authority of the Internet Architecture Board according to the rules established in [RFC3172]. There are no other registrars for .arpa.
5. Updates to Home Networking Control Protocol

The final paragraph of Home Networking Control Protocol [RFC7788], section 8, is updated as follows:

OLD:

Names and unqualified zones are used in an HNCP network to provide naming and service discovery with local significance. A network-wide zone is appended to all single labels or unqualified zones in order to qualify them. ".home" is the default; however, an administrator MAY configure the announcement of a Domain-Name TLV (Section 10.6) for the network to use a different one. In case multiple are announced, the domain of the node with the greatest node identifier takes precedence.

NEW:

Names and unqualified zones are used in an HNCP network to provide naming and service discovery with local significance. A network-wide zone is appended to all single labels or unqualified zones in order to qualify them. 'home.arpa.' is the default; however, an administrator MAY configure the announcement of a Domain-Name TLV (Section 10.6) for the network to use a different one. In case multiple are announced, the domain of the node with the greatest node identifier takes precedence.

The 'home.arpa.' special-use name does not require a special resolution protocol. Names for which the rightmost two labels are 'home.arpa.' are resolved using the DNS protocol [RFC1035].

6. Security Considerations

6.1. Local Significance

A DNS record that is returned as a response to a query for an FQDN that is a subdomain of 'home.arpa.' is expected to have local significance. It is expected to be returned by a server involved in name resolution for the homenet the device is connected in. However, such response MUST NOT be considered more trustworthy than would be a similar response for any other DNS query.

Because 'home.arpa.' is not globally scoped and cannot be secured using DNSSEC based on the root domain's trust anchor, there is no way to tell, using a standard DNS query, in which homenet scope an answer belongs. Consequently, users may experience surprising results with such names when roaming to different homenets.

To prevent this from happening, it could be useful for the resolver on the host to securely differentiate between different homenets, and between identical names on different homenets. However, a mechanism for doing this has not yet been standardized, and doing so is out of scope for this document. It is expected that this will be explored in future work.

Locally Served Zones ([RFC6303] section 7) recommends installing trust anchors for locally served zones. However, in order for this to be effective, there must be some way of configuring the trust anchor in the host. Homenet currently specifies no mechanism for configuring such trust anchors. As a result, while this advice sounds good, it is not practicable.

Also, although in principle it might be useful to install a trust anchor for a particular instance of 'home.arpa.', it's reasonable to expect that a host with such a trust anchor might from time to time connect to more than one network with its own instance of 'home.arpa.'. Such a host would be unable to access services on any instance of 'home.arpa.' other than the one for which a trust anchor was configured.

It is in principle possible to attach an identifier to an instance of 'home.arpa.' that could be used to identify which trust anchor to rely on for validating names in that particular instance. However, the security implications of this are complicated, and such a mechanism, as well as a discussion of those implications, is out of scope for this document.

6.2. Insecure Delegation

It is not possible to install a trust anchor (a DS RR) for this zone in the '.arpa' zone. The reason for this is that in order to do so, it would be necessary to have the key-signing key for the zone ([RFC4034] Section 5). Since the zone is not globally unique, no one key would work.

An alternative would be to provide a authenticated denial of existence ([RFC4033] Section 3.2). This would be done simply by not having a delegation from the 'arpa.' zone. However, this requires the validating resolver to treat 'home.arpa.' specially. If a validating resolver that doesn't treat 'home.arpa.' specially attempts to validate a name in 'home.arpa.', an authenticated denial of existence of 'home' as a subdomain of 'arpa.' would cause the validation to fail. Therefore, the only delegation that will allow names under 'home.arpa.' to be resolved by all validating resolvers is an insecure delegation as in [RFC6303] section 7.

Consequently, unless a trust anchor for the particular instance of the 'home.arpa.' zone being validated is manually configured on the validating resolver, DNSSEC signing and validation of names within the 'home.arpa.' zone is not possible.

6.3. Bypassing Manually Configured Resolvers

In Section 4, item 3, an exception is made to the behavior of stub resolvers allowing them to query local resolvers for subdomains of 'home.arpa.' even when they have been manually configured to use other resolvers. This behavior obviously has security and privacy implications, and may not be desirable depending on the context. It may be better to simply ignore this exception and, when one or more recursive resolvers are configured manually, simply fail to provide correct answers for subdomains of 'home.arpa.'. At this time we do not have operational experience that would guide us in making this decision; implementors are encouraged to consider the context in which their software will be deployed when deciding how to resolve this question.

7. Delegation of 'home.arpa.'

In order to be fully functional, there must be a delegation of 'home.arpa.' in the '.arpa.' zone [RFC3172]. This delegation MUST NOT include a DS record, and MUST point to one or more black hole servers, for example 'blackhole-1.iana.org.' and 'blackhole-2.iana.org.'. The reason that this delegation must not be signed is that not signing the delegation breaks the DNSSEC chain of trust,

which prevents a validating stub resolver from rejecting names published under 'home.arpa.' on a homenet name server.

8. IANA Considerations

IANA is requested to record the domain name 'home.arpa.' in the Special-Use Domain Names registry [SUDN]. IANA is requested, with the approval of IAB, to implement the delegation requested in Section 7.

IANA is further requested to create a new subregistry within the "Locally-Served DNS Zones" registry [LSDZ], titled "Transport-Independent Locally-Served DNS Zones", with the same format as the other subregistries. IANA is requested to add an entry in this new registry for 'home.arpa.' with the description "Homenet Special-Use Domain", listing this document as the reference. The registration procedure for this subregistry should be the same as for the others, currently "IETF Review" ([RFC8126] Section 4.8).

9. Acknowledgments

The authors would like to thank Stuart Cheshire for his prior work on '.home', as well as the homenet chairs: Mark Townsley and Ray Bellis. We would also like to thank Paul Hoffman for providing review and comments on the IANA considerations section, Andrew Sullivan for his review and proposed text, and Suzanne Woolf and Ray Bellis for their very detailed review comments and process insights. Thanks to Mark Andrews for providing an exhaustive reference list on the topic of insecure delegations. Thanks to Dale Worley for catching a rather egregious mistake and for the Gen-Art review, and to Daniel Migault for a thorough SecDir review. Thanks to Warren Kumari for catching some additional issues, and to Adam Roach for some helpful clarifications.

10. References

10.1. Normative References

- [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>>.
- [RFC3172] Huston, G., Ed., "Management Guidelines & Operational Requirements for the Address and Routing Parameter Area Domain ("arpa")", BCP 52, RFC 3172, DOI 10.17487/RFC3172, September 2001, <<https://www.rfc-editor.org/info/rfc3172>>.

- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <<https://www.rfc-editor.org/info/rfc4035>>.
- [RFC6303] Andrews, M., "Locally Served DNS Zones", BCP 163, RFC 6303, DOI 10.17487/RFC6303, July 2011, <<https://www.rfc-editor.org/info/rfc6303>>.
- [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>>.
- [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>>.

10.2. Informative References

- [ICANN1] "New gTLD Collision Risk Mitigation", October 2013, <<https://www.icann.org/en/system/files/files/new-gtld-collision-mitigation-05aug13-en.pdf>>.
- [ICANN2] "New gTLD Collision Occurrence Management", October 2013, <<https://www.icann.org/en/system/files/files/resolutions-new-gtld-annex-1-07oct13-en.pdf>>.
- [LSDZ] "Locally-Served DNS Zones Registry", July 2011, <<https://www.iana.org/assignments/locally-served-dns-zones/locally-served-dns-zones.xhtml>>.
- [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>>.
- [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>>.
- [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>>.

- [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>>.
- [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>>.
- [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>>.
- [SUDN] "Special-Use Domain Names Registry", July 2012, <<https://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml>>.

Authors' Addresses

Pierre Pfister
Cisco Systems
Paris
France

Email: pierre.pfister@darou.fr

Ted Lemon
Nominum, Inc.
800 Bridge Parkway
Redwood City, California 94065
United States of America

Phone: +1 650 381 6000
Email: ted.lemon@nominum.com

HOMENET
Internet-Draft
Intended status: Standards Track
Expires: April 30, 2018

D. Migault (Ed)
Ericsson
R. Weber
Nominum
R. Hunter
Globis Consulting BV
C. Griffiths

W. Cloetens
SoftAtHome
October 27, 2017

Outsourcing Home Network Authoritative Naming Service
draft-ietf-homenet-front-end-naming-delegation-06

Abstract

Designation of services and devices of a home network is not user friendly, and mechanisms should enable a user to designate services and devices inside a home network using names.

In order to enable internal communications while the home network experiments Internet connectivity shortage, the naming service should be hosted on a device inside the home network. On the other hand, home networks devices have not been designed to handle heavy loads. As a result, hosting the naming service on such home network device, visible on the Internet exposes this device to resource exhaustion and other attacks, which could make the home network unreachable, and most probably would also affect the internal communications of the home network.

As result, home networks may prefer not serving the naming service for the Internet, but instead prefer outsourcing it to a third party. This document describes a mechanisms that enables the Home Network Authority (HNA) to outsource the naming service to the Outsourcing Infrastructure.

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 30, 2018.

Copyright Notice

Copyright (c) 2017 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. Requirements notation	3
2. Introduction	3
3. Terminology	4
4. Architecture Description	6
4.1. Architecture Overview	6
4.2. Example: Homenet Zone	8
4.3. Example: HNA necessary parameters for outsourcing	10
5. Synchronization between HNA and the Synchronization Server	11
5.1. Synchronization with a Hidden Primary	11
5.2. Securing Synchronization	12
5.3. HNA Security Policies	14
6. DNSSEC compliant Homenet Architecture	14
6.1. Zone Signing	14
6.2. Secure Delegation	16
7. Handling Different Views	17
7.1. Misleading Reasons for Local Scope DNS Zone	17
7.2. Consequences	18
7.3. Guidance and Recommendations	18
8. Homenet Reverse Zone	19
9. Renumbering	19
9.1. Hidden Primary	20
9.2. Synchronization Server	21
10. Privacy Considerations	22
11. Security Considerations	22

11.1.	Names are less secure than IP addresses	22
11.2.	Names are less volatile than IP addresses	23
11.3.	DNS Reflection Attacks	23
11.3.1.	Reflection Attack involving the Hidden Primary . . .	23
11.3.2.	Reflection Attacks involving the Synchronization Server	25
11.3.3.	Reflection Attacks involving the Public Authoritative Servers	26
11.4.	Flooding Attack	26
11.5.	Replay Attack	26
12.	IANA Considerations	27
13.	Acknowledgment	27
14.	References	27
14.1.	Normative References	27
14.2.	Informational References	30
Appendix A.	Document Change Log	31
Authors' Addresses	33

1. Requirements notation

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

2. Introduction

IPv6 provides global end to end IP reachability. End users prefer to use names instead of long and complex IPv6 addresses when accessing services hosted in the home network.

Customer Edge Routers and other Customer Premises Equipment (CPEs) are already providing IPv6 connectivity to the home network, and generally provide IPv6 addresses or prefixes to the nodes of the home network. In addition, [RFC7368] recommends that home networks be resilient to connectivity disruption from the ISP. This could be achieved by a dedicated device inside the home network that builds, serves or manage the Homenet Zone, thus providing bindings between names and IP addresses.

CPEs are of course good candidates to manage the binding between names and IP addresses of nodes. However, this could also be performed by another device in the home network that is not a CPE. In addition, a given home network may have multiple nodes that may implement this functionality. Since management of the Homenet Zone involves DNS specific mechanisms that cannot be distributed (primary server), when multiple nodes can potentially manage the Homenet Zone, a single node needs to be selected. This selected node is designated as the Homenet Naming Authority (HNA).

CPEs, Homenet Naming Authority, as well as home network devices are usually low powered devices not designed not for terminating heavy traffic. As a result, hosting an authoritative DNS service on the Internet may expose the home network to resource exhaustion and other attacks. This may isolate the home network from the Internet and also impact the services hosted by the such an home network device, thus affecting overall home network communication.

In order to avoid resource exhaustion and other attacks, this document describes an architecture that outsources the authoritative naming service of the home network. More specifically, the Homenet Naming Authority builds the Homenet Zone and outsources it to an Outsourcing Infrastructure. The Outsourcing Infrastructure is in charge of publishing the corresponding Public Homenet Zone on the Internet.

Section 4.1 provides an architecture description that describes the relation between the Homenet Naming Authority and the Outsourcing Architecture. In order to keep the Public Homenet Zone up-to-date Section 5 describes how the Homenet Zone and the Public Homenet Zone can be synchronized. The proposed architecture aims at deploying DNSSEC, and the Public Homenet Zone is expected to be signed with a secure delegation. The zone signing and secure delegation may be performed either by the Homenet Naming Authority or by the Outsourcing Infrastructure. Section 6 discusses these two alternatives. Section 7 discusses the consequences of publishing multiple representations of the same zone also commonly designated as views. This section provides guidance to limit the risks associated with multiple views. Section 8 discusses management of the reverse zone. Section 9 discusses how renumbering should be handled. Finally, Section 10 and Section 11 respectively discuss privacy and security considerations when outsourcing the Homenet Zone.

3. Terminology

- Customer Premises Equipment: (CPE) is a router providing connectivity to the home network.
- Homenet Naming Authority: (HNA) is a home network node responsible to manage the Homenet Zone. This includes building the Homenet Zone, as well as managing the distribution of that Homenet Zone through the Outsourcing Infrastructure.
- Registered Homenet Domain: is the Domain Name associated to the home network.
- Homenet Zone: is the DNS zone associated with the home network. It is designated by its Registered Homenet Domain. This zone

is built by the HNA and contains the bindings between names and IP addresses of the nodes in the home network. The HNA synchronizes the Homenet Zone with the Synchronization Server via a hidden primary / secondary architecture. The Outsourcing Infrastructure may process the Homenet Zone - for example providing DNSSEC signing - to generate the Public Homenet Zone. This Public Homenet Zone is then transmitted to the Public Authoritative Server(s) that publish it on the Internet.

- Public Homenet Zone: is the public version of the Homenet Zone. It is expected to be signed with DNSSEC. It is hosted by the Public Authoritative Server(s), which are authoritative for this zone. The Public Homenet Zone and the Homenet Zone might be different. For example some names might not become reachable from the Internet, and thus not be hosted in the Public Homenet Zone. Another example of difference may also occur when the Public Homenet Zone is signed whereas the Homenet Zone is not signed.
- Outsourcing Infrastructure: is the combination of the Synchronization Server and the Public Authoritative Server(s).
- Public Authoritative Servers: are the authoritative name servers hosting 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.
- Synchronization Server: is the server with which the HNA synchronizes the Homenet Zone. The Synchronization Server is configured as a secondary and the HNA acts as primary. There MAY be multiple Synchronization Servers, but the text assumes a single server. In addition, the text assumes the Synchronization Server is a separate entity. This is not a requirement, and when the HNA signs the zone, the synchronization function might also be operated by the Public Authoritative Servers.
- Homenet Reverse Zone: The reverse zone file associated with the Homenet Zone.
- Reverse Public Authoritative Servers: are the authoritative name server(s) hosting the Public Homenet Reverse Zone. Queries for reverse resolution of the Homenet Domain are sent to this server. Similarly to Public Authoritative Servers, for resiliency, the Homenet Reverse Zone SHOULD be hosted on multiple servers.

- Reverse Synchronization Server: is the server with which the HNA synchronizes the Homenet Reverse Zone. It is configured as a secondary and the HNA acts as primary. There MAY be multiple Reverse Synchronization Servers, but the text assumes a single server. In addition, the text assumes the Reverse Synchronization Server is a separate entity. This is not a requirement, and when the HNA signs the zone, the synchronization function might also be operated by the Reverse Public Authoritative Servers.
- Hidden Primary: designates the primary server of the HNA, that synchronizes the Homenet Zone with the Synchronization Server. A primary / secondary architecture is used between the HNA and the Synchronization Server. The hidden primary is not expected to serve end user queries for the Homenet Zone as a regular primary server would. The hidden primary is only known to its associated Synchronization Server.

4. Architecture Description

This section describes the architecture for outsourcing the authoritative naming service from the HNA to the Outsourcing Infrastructure. Section 4.1 describes the architecture, Section 4.2 and Section 4.3 illustrates this architecture and shows how the Homenet Zone should be built by the HNA. It also lists the necessary parameters the HNA needs to be able to outsource the authoritative naming service. These two sections are informational and non-normative.

4.1. Architecture Overview

Figure 1 provides an overview of the architecture.

The home network is designated by the Registered Homenet Domain Name -- example.com in Figure 1. The HNA builds the Homenet Zone associated with the home network. How the Homenet Zone is built is out of the scope of this document. The HNA may host or interact with multiple services to determine name-to-address mappings, such as a web GUI, DHCP [RFC6644] or mDNS [RFC6762]. These services may coexist and may be used to populate the Homenet Zone. This document assumes the Homenet Zone has been populated with domain names that are intended to be publicly published and that are publicly reachable. More specifically, names associated with services or devices that are not expected to be reachable from outside the home network or names bound to non-globally reachable IP addresses MUST NOT be part of the Homenet Zone.

Once the Homenet Zone has been built, the HNA does not host an authoritative naming service, but instead outsources it to the Outsourcing Infrastructure. The Outsourcing Infrastructure takes the Homenet Zone as an input and publishes the Public Homenet Zone. If the HNA does not sign the Homenet Zone, the Outsourcing Infrastructure may instead sign it on behalf of the HNA. Figure 1 provides a more detailed description of the Outsourcing Infrastructure, but overall, it is expected that the HNA provides the Homenet Zone. Then the Public Homenet Zone is derived from the Homenet Zone and published on the Internet.

As a result, DNS queries from the DNS resolvers on the Internet are answered by the Outsourcing Infrastructure and do not reach the HNA. Figure 1 illustrates the case of the resolution of `node1.example.com`.

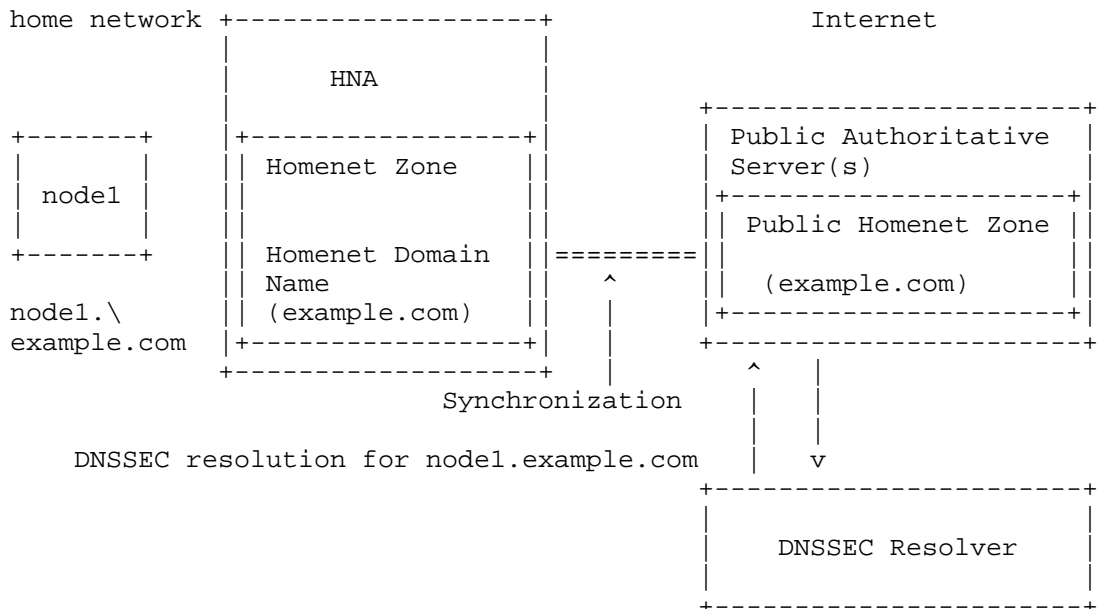


Figure 1: Homenet Naming Architecture Description

The Outsourcing Infrastructure is described in Figure 2. The Synchronization Server receives the Homenet Zone as an input. The received zone may be transformed to output the Public Homenet Zone. Various operations may be performed here, however this document only considers zone signing as a potential operation. This should occur only when the HNA outsources this operation to the Synchronization Server. On the other hand, if the HNA signs the Homenet Zone itself, the zone would be collected by the Synchronization Server and

directly transferred to the Public Authoritative Server(s). These policies are discussed and detailed in Section 6 and Section 7.




Figure 2: Outsourcing Infrastructure Description

4.2. Example: Homenet Zone

This section is not normative and intends to illustrate how the HNA builds the Homenet Zone.

As depicted in Figure 1 and Figure 2, the Public Homenet Zone is hosted on the Public Authoritative Server(s), whereas the Homenet Zone is hosted on the HNA. Motivations for keeping these two zones identical are detailed in Section 7, and this section considers that the HNA builds the zone that will be effectively published on the Public Authoritative Server(s). In other words "Homenet to Public Zone transformation" is the identity also commonly designated as "no operation" (NOP).

In that case, the Homenet Zone should configure its Name Server RRset (NS) and Start of Authority (SOA) with the values associated with the Public Authoritative Server(s). This is illustrated in Figure 3. `public.primary.example.net` is the FQDN of the Public Authoritative Server(s), and `IP1`, `IP2`, `IP3`, `IP4` are the associated IP addresses. Then the HNA should add the additional new nodes that enter the home

network, remove those that should be removed, and sign the Homenet Zone.

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

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

@ NS public.authoritative.servers.example.net

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

Figure 3: Homenet Zone

The SOA RRset is defined in [RFC1033], [RFC1035] and [RFC2308]. This SOA is specific, as it is used for the synchronization between the Hidden Primary and the Synchronization Server and published on the DNS Public Authoritative Server(s)..

- MNAME: indicates the primary. In our case the zone is published on the Public Authoritative Server(s), and its name MUST be included. If multiple Public Authoritative Server(s) are involved, one of them MUST be chosen. More specifically, the HNA MUST NOT include the name of the Hidden Primary.
- RNAME: indicates the email address to reach the administrator. [RFC2142] recommends using hostmaster@domain and replacing the '@' sign by '.'.
- REFRESH and RETRY: indicate respectively in seconds how often secondaries need to check the primary, and the time between two refresh when a refresh has failed. Default values indicated by [RFC1033] are 3600 (1 hour) for refresh and 600 (10 minutes) for retry. This value might be too long for highly dynamic content. However, the Public Authoritative Server(s) and the HNA are expected to implement NOTIFY [RFC1996]. So whilst shorter refresh timers might increase the bandwidth usage for

secondaries hosting large number of zones, it will have little practical impact on the elapsed time required to achieve synchronization between the Outsourcing Infrastructure and the Hidden Master. As a result, the default values are acceptable.

EXPIRE: is the upper limit data SHOULD be kept in absence of refresh. The default value indicated by [RFC1033] is 3600000 (approx. 42 days). In home network architectures, the HNA provides both the DNS synchronization and the access to the home network. This device may be plugged and unplugged by the end user without notification, thus we recommend a long expiry timer.

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

4.3. Example: HNA necessary parameters for outsourcing

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

Synchronization Server may be configured with the following parameters. These parameters are necessary to establish a secure channel between the HNA and the Synchronization Server as well as to specify the DNS zone that is in the scope of the communication:

- Synchronization Server: The associated FQDNs or IP addresses of the Synchronization Server. IP addresses are optional and the FQDN is sufficient. To secure the binding name and IP addresses, a DNSSEC exchange is required. Otherwise, the IP addresses should be entered manually.
- Authentication Method: How the HNA authenticates itself to the Synchronization Server. This MAY depend on the implementation but this should cover at least IPsec, DTLS and TSIG
- Authentication data: Associated Data. PSK only requires a single argument. If other authentication mechanisms based on certificates are used, then HNA private keys, certificates and certification authority should be specified.
- Public Authoritative Server(s): The FQDN or IP addresses of the Public Authoritative Server(s). It MAY correspond to the data that will be set in the NS RRsets and SOA of the Homenet Zone. IP addresses are optional and the FQDN is sufficient. To

secure the binding between name and IP addresses, a DNSSEC exchange is required. Otherwise, the IP addresses should be entered manually.

- Registered Homenet Domain: The domain name used to establish the secure channel. This name is used by the Synchronization Server and the HNA for the primary / secondary configuration as well as to index the NOTIFY queries of the HNA when the HNA has been renumbered.

Setting the Homenet Zone requires the following information.

- Registered Homenet Domain: The Domain Name of the zone. Multiple Registered Homenet Domains may be provided. This will generate the creation of multiple Public Homenet Zones.
- Public Authoritative Server(s): The Public Authoritative Server(s) associated with the Registered Homenet Domain. Multiple Public Authoritative Server(s) may be provided.

5. Synchronization between HNA and the Synchronization Server

The Homenet Reverse Zone and the Homenet Zone MAY be updated either with DNS UPDATE [RFC2136] or using a primary / secondary synchronization. The primary / secondary mechanism is preferred as it scales better and avoids DoS attacks: First the primary notifies the secondary that the zone must be updated and leaves the secondary to proceed with the update when possible. 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 query performed by the secondary is a small packet sent over TCP (section 4.2 [RFC5936]), which mitigates reflection attacks using a forged NOTIFY. On the other hand, DNS UPDATE (which can be transported over UDP), requires more processing than a NOTIFY, and does not allow the server to perform asynchronous updates.

This document RECOMMENDS use of a primary / secondary mechanism instead of the use of DNS UPDATE. This section details the primary / secondary mechanism.

5.1. Synchronization with a Hidden Primary

Uploading and dynamically updating the zone file on the Synchronization Server can be seen as zone provisioning between the HNA (Hidden Primary) and the Synchronization Server (Secondary Server). This can be handled either in band or out of band.

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 Synchronization Server is configured as a secondary for the Homenet Domain Name. This secondary configuration has been previously agreed between the end user and the provider of the Synchronization Server. In order to set the primary / secondary architecture, the HNA acts as a Hidden Primary Server, which is a regular authoritative DNS Server listening on the WAN interface.

The Hidden Primary Server SHOULD accept SOA [RFC1033], AXFR [RFC1034], and IXFR [RFC1995] queries from its configured secondary DNS server(s). The Hidden Primary Server SHOULD send NOTIFY messages [RFC1996] in order to update Public DNS server zones as updates occur. Because, the Homenet Zones are likely to be small, the HNA MUST implement AXFR and SHOULD implement IXFR.

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

- Interface Binding: the Hidden Primary Server listens on the WAN Interface, whereas a regular authoritative server for the home network would listen on the home network interface.
- Limited exchanges: the purpose of the Hidden Primary Server is to synchronize with the Synchronization Server, not to serve any zones to end users. As a result, exchanges are performed with specific nodes (the Synchronization Server). Further, exchange types are limited. The only legitimate exchanges are: NOTIFY initiated by the Hidden Primary and IXFR or AXFR exchanges initiated by the Synchronization Server. 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 Synchronization Server. The HNA MUST listen for DNS on TCP and UDP and MUST at least allow SOA lookups of the Homenet Zone.

5.2. Securing Synchronization

Exchange between the Synchronization Server and the HNA MUST be secured, at least for integrity protection and for authentication.

TSIG [RFC2845] or SIG(0) [RFC2931] MAY be used to secure the DNS communications between the HNA and the Synchronization Server. TSIG

uses a symmetric key which can be managed by TKEY [RFC2930]. Management of the key involved in SIG(0) is performed through zone updates. How keys are rolled over with SIG(0) is out-of-scope of this document. The advantage of these mechanisms is that they are only associated with the DNS application. Not relying on shared libraries eases testing and integration. On the other hand, using TSIG, TKEY or SIG(0) requires these mechanisms to be implemented on the HNA, which adds code and complexity. Another disadvantage is that TKEY does not provide authentication mechanisms.

Protocols like TLS [RFC5246] / DTLS [RFC6347] MAY be used to secure the transactions between the Synchronization Server and the HNA. The advantage of TLS/DTLS is that this technology is widely deployed, and most of the devices already embed TLS/DTLS libraries, possibly also taking advantage of hardware acceleration. Further, TLS/DTLS provides authentication facilities and can use certificates to authenticate the Synchronization Server and the HNA. On the other hand, using TLS/DTLS requires implementing DNS exchanges over TLS/DTLS, as well as a new service port. This document therefore does NOT RECOMMEND this option.

IPsec [RFC4301] IKEv2 [RFC7296] MAY also be used to secure transactions between the HNA and the Synchronization Server. Similarly to TLS/DTLS, most HNAs already embed an IPsec stack, and IKEv2 supports multiple authentication mechanisms via the EAP framework. In addition, IPsec can be used to protect DNS exchanges between the HNA and the Synchronization Server without any modifications of the DNS server or client. DNS integration over IPsec only requires an additional security policy in the Security Policy Database (SPD). One disadvantage of IPsec is that NATs and firewall traversal may be problematic. However, in our case, the HNA is connected to the Internet, and IPsec communication between the HNA and the Synchronization Server should not be impacted by middle boxes.

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

Note also that authentication of message exchanges between the HNA and the Synchronization Server SHOULD NOT use the external IP address of the HNA to index the appropriate keys. As detailed in Section 9, the IP addresses of the Synchronization Server 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.

5.3. HNA Security Policies

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

The Hidden Primary, as described in this document 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 Synchronization Server.

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.

6. 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.

Deploying DNSSEC requires signing the zone and configuring a secure delegation. As described in Section 4.1, signing can be performed either by the HNA or by the Outsourcing Infrastructure. Section 6.1 details the implications of these two alternatives. Similarly, the secure delegation can be performed by the HNA or by the Outsourcing Infrastructure. Section 6.2 discusses these two alternatives.

6.1. Zone Signing

This section discusses the pros and cons when zone signing is performed by the HNA or by the Outsourcing Infrastructure. It is RECOMMENDED that the HNA signs the zone unless there is a strong argument against this, such as a HNA that is not capable of signing

the zone. In that case zone signing MAY be performed by the Outsourcing Infrastructure on behalf of the HNA.

Reasons for signing the zone by the HNA are:

- 1: Keeping the Homenet Zone and the Public Homenet Zone equal to securely optimize DNS resolution. As the Public Zone is signed with DNSSEC, RRsets are authenticated, and thus DNS responses can be validated even though they are not provided by the authoritative server. This provides the HNA the ability to respond on behalf of the Public Authoritative Server(s). This could be useful for example if, in the future, the HNA announces to the home network that the HNA can act as a local authoritative primary or equivalent for the Homenet Zone. Currently the HNA is not expected to receive authoritative DNS queries, as its IP address is not mentioned in the Public Homenet Zone. On the other hand most HNAs host a resolving function, and could be configured to perform a local lookup to the Homenet Zone instead of initiating a DNS exchange with the Public Authoritative Server(s). Note that outsourcing the zone signing operation means that all DNSSEC queries SHOULD be cached to perform a local lookup, otherwise a resolution with the Public Authoritative Server(s) would be performed.
- 2: Keeping the Homenet Zone and the Public Homenet Zone equal to securely address the connectivity disruption independence detailed in [RFC7368] section 4.4.1 and 3.7.5. As local lookups are possible in case of network disruption, communications within the home network can still rely on the DNSSEC service. Note that outsourcing the zone signing operation does not address connectivity disruption independence with DNSSEC. Instead local lookup would provide DNS as opposed to DNSSEC responses provided by the Public Authoritative Server(s).
- 3: Keeping the Homenet Zone and the Public Homenet Zone equal to guarantee coherence between DNS responses. Using a unique zone is one way to guarantee uniqueness of the responses among servers and places. Issues generated by different views are discussed in more details in Section 7.
- 2: Privacy and Integrity of the DNSSEC Homenet Zone are better guaranteed. When the Zone is signed by the HNA, it makes modification of the DNS data -- for example for flow redirection -- impossible. As a result, signing the Homenet Zone by the HNA provides better protection for end user privacy.

Reasons for signing the zone by the Outsourcing Infrastructure are:

- 1: The HNA may not be capable of signing the zone, most likely because its firmware does not support this function. However this reason is expected to become less and less valid over time.
- 2: Outsourcing DNSSEC management operations. Management operations involve key roll-over, which can be performed automatically by the HNA and transparently for the end user. Avoiding DNSSEC management is mostly motivated by bad software implementations.
- 3: Reducing the impact of HNA replacement on the Public Homenet Zone. Unless the HNA private keys can be extracted and stored off-device, HNA hardware replacement will result in an emergency key roll-over. This can be mitigated by using relatively small TTLs.
- 4: Reducing configuration impact on the end user. Unless there are zero configuration mechanisms in place to provide credentials between the new HNA and the Synchronization Server, authentication associations between the HNA and the Synchronization Server would need to be re-configured. As HNA replacement is not expected to happen regularly, end users may not be at ease with such configuration settings. However, mechanisms as described in [I-D.ietf-homenet-naming-architecture-dhc-options] use DHCP Options to outsource the configuration and avoid this issue.
- 5: The Outsourcing Infrastructure is more likely to handle private keys more securely than the HNA. However, having all private keys in one place may also nullify that benefit.

6.2. Secure Delegation

Secure delegation is achieved only if the DS RRset is properly set in the parent zone. Secure delegation can be performed by the HNA or the Outsourcing Infrastructures (that is the Synchronization Server or the Public Authoritative Server(s)).

The DS RRset can be updated manually with nsupdate for example. This requires the HNA or the Outsourcing Infrastructure 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 Outsourcing Infrastructure. In fact, the

Public Authoritative Server(s) may use Automating DNSSEC Delegation Trust Maintenance [RFC7344].

7. Handling Different Views

The Homenet Zone provides information about the home network. Some users may be tempted to have provide responses dependent on the origin of the DNS query. More specifically, some users may be tempted to provide a different view for DNS queries originating from the home network and for DNS queries coming from the Internet. Each view could then be associated with a dedicated Homenet Zone. Note that this document does not specify how DNS queries originating from the home network are addressed to the Homenet Zone. This could be done via hosting the DNS resolver on the HNA for example.

This section is not normative. Section 7.1 details why some nodes may only be reachable from the home network and not from the global Internet. Section 7.2 briefly describes the consequences of having distinct views such as a "home network view" and an "Internet view". Finally, Section 7.3 provides guidance on how to resolve names that are only significant in the home network, without creating different views.

7.1. Misleading Reasons for Local Scope DNS Zone

The motivation for supporting different views is to provide different answers dependent on the origin of the DNS query, for reasons such as:

- 1: An end user may want to have services not published on the Internet. Services like the HNA administration interface that provides the GUI to administer your HNA might not seem advisable to publish on the Internet. Similarly, services like the mapper that registers the devices of your home network may also not be desirable to be published on the Internet. In both cases, these services should only be known or used by the network administrator. To restrict the access of such services, the home network administrator may choose to publish these pieces of information only within the home network, where it might be assumed that the users are more trusted than on the Internet. Even though this assumption may not be valid, at least this may reduce the surface of any attack.
- 2: Services within the home network may be reachable using non global IP addresses. IPv4 and NAT may be one reason. On the other hand IPv6 may favor link-local or site-local IP addresses. These IP addresses are not significant outside the boundaries of the home network. As a result, they MAY be

published in the home network view, and SHOULD NOT be published in the Public Homenet Zone.

7.2. Consequences

Enabling different views leads to a non-coherent naming system. Depending on where resolution is performed, some services will not be available. This may be especially inconvenient with devices with multiple interfaces that are attached both to the Internet via a 3G/4G interface and to the home network via a WLAN interface. Devices may also cache the results of name resolution, and these cached entries may no longer be valid if a mobile device moves between a homenet connection and an internet connection e.g. a device temporarily loses wifi signal and switches to 3G.

Regarding local-scope IP addresses, such devices may end up with poor connectivity. Suppose, for example, that DNS resolution is performed via the WLAN interface attached to the HNA, and the response provides local-scope IP addresses, but the communication is initiated on the 3G/4G interface. Communications with local-scope addresses will be unreachable on the Internet, thus aborting the communication. The same situation occurs if a device is flip / flopping between various WLAN networks.

Regarding DNSSEC, if the HNA does not sign the Homenet Zone and outsources the signing process, the two views are different, because one is protected with DNSSEC whereas the other is not. Devices with multiple interfaces will have difficulty securing the naming resolution, as responses originating from the home network may not be signed.

For devices with all its interfaces attached to a single administrative domain, that is to say the home network, or the Internet. Incoherence between DNS responses may still also occur if the device is able to perform DNS resolutions both using the DNS resolving server of the home network, or one of the ISP. DNS resolution performed via the HNA or the ISP resolver may be different than those performed over the Internet.

7.3. Guidance and Recommendations

As documented in Section 7.2, it is RECOMMENDED to avoid different views. If network administrators choose to implement multiple views, impacts on devices' resolution SHOULD be evaluated.

As a consequence, the Homenet Zone is expected to be an exact copy of the Public Homenet Zone. As a result, services that are not expected to be published on the Internet SHOULD NOT be part of the Homenet

Zone, local-scope addresses SHOULD NOT be part of the Homenet Zone, and when possible, the HNA SHOULD sign the Homenet Zone.

The Homenet Zone is expected to host public information only. It is not the scope of the DNS service to define local home network boundaries. Instead, local scope information is expected to be provided to the home network using local scope naming services. mDNS [RFC6762] DNS-SD [RFC6763] are two examples of these services. Currently mDNS is limited to a single link network. However, future protocols are expected to leverage this constraint as pointed out in [RFC7558].

8. Homenet Reverse Zone

This section is focused on the Homenet Reverse Zone.

Firstly, all considerations for the Homenet Zone apply to the Homenet Reverse Zone. The main difference between the Homenet Reverse Zone and the Homenet Zone is that the parent zone of the Homenet Reverse Zone is most likely managed by the ISP. As the ISP also provides the IP prefix to the HNA, it may be able to authenticate the HNA using mechanisms outside the scope of this document e.g. the physical attachment point to the ISP network. If the Reverse Synchronization Server is managed by the ISP, credentials to authenticate the HNA for the zone synchronization may be set automatically and transparently to the end user. [I-D.ietf-homenet-naming-architecture-dhc-options] describes how automatic configuration may be performed.

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. [I-D.howard-dnsop-ip6rdns] provides guidance on how to address scalability issues.

9. Renumbering

This section details how renumbering is handled by the Hidden Primary server or the Synchronization Server. Both types of renumbering are discussed i.e. "make-before-break" and "break-before-make".

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.

9.1. Hidden Primary

In a renumbering scenario, the 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 Homenet Zone will also be updated. Although the new and old IP addresses may be stored in the 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 plane. 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 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 $2 * TTL$.

Once the Homenet Zone file has been updated on the Hidden Primary, the Hidden Primary needs to inform the Outsourcing Infrastructure that the 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.

The Hidden Primary SHOULD inform the Synchronization Server that the Homenet Zone has been updated by sending a NOTIFY payload with the new IP address. In addition, this NOTIFY payload SHOULD be authenticated using SIG(0) or TSIG. When the Synchronization Server receives the NOTIFY payload, it MUST authenticate it. Note that the cryptographic key used for the authentication SHOULD be indexed by the Registered Homenet Domain contained in the NOTIFY payload as well as the RRSIG. In other words, the IP address SHOULD NOT be used as an index. If authentication succeeds, the Synchronization Server MUST also notice the IP address has been modified and perform a reachability check before updating its primary configuration. The routability check MAY be performed by sending a SOA request to the

Hidden Primary using the source IP address of the NOTIFY. This exchange is also secured, and if an authenticated response is received from the Hidden Primary with the new IP address, the Synchronization Server SHOULD update its configuration file and retrieve the Homenet Zone using an AXFR or a IXFR exchange.

Note that the primary reason for providing the IP address is that the Hidden Primary is not publicly announced in the DNS. If the Hidden Primary were publicly announced in the DNS, then the IP address update could have been performed using the DNS as described in Section 9.2.

9.2. Synchronization Server

Renumbering of the Synchronization Server results in the Synchronization Server changing its IP address. The Synchronization Server is a secondary, so its renumbering does not impact the Homenet Zone. In fact, exchanges to the Synchronization Server are restricted to the Homenet Zone synchronization. In our case, the Hidden Primary MUST be able to send NOTIFY payloads to the Synchronization Server.

If the Synchronization Server is configured in the Hidden Primary configuration file using a FQDN, then the update of the IP address is performed by DNS. More specifically, before sending the NOTIFY, the Hidden Primary performs a DNS resolution to retrieve the IP address of the secondary.

As described in Section 9.1, the Synchronization Server DNS information SHOULD be coherent with the IP plane. Let TTL be the TTL associated with the Synchronization Server FQDN, T_{NEW} the time the new IP address replaces the old one and $T_{\text{OLD_UNREACHABLE}}$ the time the Synchronization Server is not reachable anymore with its old IP address. Seamless reachability is provided as long as $T_{\text{OLD_UNREACHABLE}} - T_{\text{NEW}} > 2 * \text{TTL}$. If this condition is not met, the Synchronization Server may be unreachable during $2 * \text{TTL} - (T_{\text{OLD_UNREACHABLE}} - T_{\text{NEW}})$. In the case of a break-before-make, $T_{\text{OLD_UNREACHABLE}} = T_{\text{NEW}}$, and it may become unreachable up to $2 * \text{TTL}$.

Some DNS infrastructure uses the IP address to designate the secondary, in which case, other mechanisms must be found. The reason for using IP addresses instead of names is generally to reach an internal interface that is not designated by a FQDN, and to avoid potential bootstrap problems. Such scenarios are considered as out of scope in the case of home networks.

10. Privacy Considerations

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

The Homenet Zone contains a full description of the services hosted in the network. These services may not be expected to be publicly shared although their names remain accessible through the Internet. Even though DNS makes information public, the DNS does not expect to make the complete list of services public. In fact, making information public still requires the key (or FQDN) of each service to be known by the resolver in order to retrieve information about the services. More specifically, making mywebsite.example.com public in the DNS, is not sufficient to make resolvers aware of the existence web site. However, an attacker may walk the reverse DNS zone, or use other reconnaissance techniques to learn this information as described in [RFC7707].

In order to prevent the complete Homenet Zone being published on the Internet, AXFR queries SHOULD be blocked on the Public Authoritative Server(s). Similarly, to avoid zone-walking NSEC3 [RFC5155] SHOULD be preferred over NSEC [RFC4034].

When the Homenet Zone is outsourced, the end user should be aware that it provides a complete description of the services available on the home network. More specifically, names usually provides a clear indication of the service and possibly even the device type, and as the Homenet Zone contains the IP addresses associated with the service, they also limit the scope of the scan space.

In addition to the Homenet Zone, the third party can also monitor the traffic associated with the 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.

11. Security Considerations

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

11.1. 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.

11.2. Names are less volatile than IP addresses

IP addresses may be used to locate a device, a host or a service. However, home networks are not expected to be assigned 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.

11.3. DNS Reflection Attacks

An attacker performs a reflection attack when it sends traffic to one or more intermediary nodes (reflectors), that in turn send back response traffic to the victim. Motivations for using an intermediary node might be anonymity of the attacker, as well as amplification of the traffic. Typically, when the intermediary node is a DNSSEC server, the attacker sends a DNSSEC query and the victim is likely to receive a DNSSEC response. This section analyzes how the different components may be involved as a reflector in a reflection attack. Section 11.3.1 considers the Hidden Primary, Section 11.3.2 the Synchronization Server, and Section 11.3.3 the Public Authoritative Server(s).

11.3.1. Reflection Attack involving the Hidden Primary

With the specified architecture, the Hidden Primary is only expected to receive DNS queries of type SOA, AXFR or IXFR. This section analyzes how these DNS queries may be used by an attacker to perform a reflection attack.

DNS queries of type AXFR and IXFR use TCP and as such are less subject to reflection attacks. This makes SOA queries the only

remaining practical vector of attacks for reflection attacks, based on UDP.

SOA queries are not associated with a large amplification factor compared to queries of type "ANY" or to query of non existing FQDNs. This reduces the probability a DNS query of type SOA will be involved in a DDoS attack.

SOA queries are expected to follow a very specific pattern, which makes rate limiting techniques an efficient way to limit such attacks, and associated impact on the naming service of the home network.

Motivations for such a flood might be a reflection attack, but could also be a resource exhaustion attack performed against the Hidden Primary. The Hidden Primary only expects to exchange traffic with the Synchronization Server, that is its associated secondary. Even though secondary servers may be renumbered as mentioned in Section 9, the Hidden Primary is likely to perform a DNSSEC resolution and find out the associated secondary's IP addresses in use. As a result, the Hidden Primary is likely to limit the origin of its incoming traffic based on the origin IP address.

With filtering rules based on IP address, SOA flooding attacks are limited to forged packets with the IP address of the secondary server. In other words, the only victims are the Hidden Primary itself or the secondary. There is a need for the Hidden Primary to limit that flood to limit the impact of the reflection attack on the secondary, and to limit the resource needed to carry on the traffic by the HNA hosting the Hidden Primary. On the other hand, mitigation should be performed appropriately, so as to limit the impact on the legitimate SOA sent by the secondary.

The main reason for the Synchronization Server sending a SOA query is to update the SOA RRset after the TTL expires, to check the serial number upon the receipt of a NOTIFY query from the Hidden Primary, or to re-send the SOA request when the response has not been received. When a flood of SOA queries is received by the Hidden Primary, the Hidden Primary may assume it is involved in an attack.

There are few legitimate time slots when the secondary is expected to send a SOA query. Suppose T_{NOTIFY} is the time a NOTIFY is sent by the Hidden Primary, T_{SOA} the last time the SOA has been queried, TTL the TTL associated to the SOA, and T_{REFRESH} the refresh time defined in the SOA RRset. The specific time SOA queries are expected can be for example T_{NOTIFY} , $T_{\text{SOA}} + 2/3 \text{ TTL}$, $T_{\text{SOA}} + \text{TTL}$, $T_{\text{SOA}} + T_{\text{REFRESH}}$, and. Outside a few minutes following these specific time slots, the probability that the HNA discards a legitimate SOA query

is very low. Within these time slots, the probability the secondary may have its legitimate query rejected is higher. If a legitimate SOA is discarded, the secondary will re-send SOA query every "retry time" second until "expire time" seconds occurs, where "retry time" and "expire time" have been defined in the SOA.

As a result, it is RECOMMENDED to set rate limiting policies to protect HNA resources. If a flood lasts more than the expired time defined by the SOA, it is RECOMMENDED to re-initiate a synchronization between the Hidden Primary and the secondaries.

11.3.2. Reflection Attacks involving the Synchronization Server

The Synchronization Server acts as a secondary coupled with the Hidden Primary. The secondary expects to receive NOTIFY query, SOA responses, AXFR and IXFR responses from the Hidden Primary.

Sending a NOTIFY query to the secondary generates a NOTIFY response as well as initiating an SOA query exchange from the secondary to the Hidden Primary. As mentioned in [RFC1996], this is a known "benign denial of service attack". As a result, the Synchronization Server SHOULD enforce rate limiting on sending SOA queries and NOTIFY responses to the Hidden Primary. Most likely, when the secondary is flooded with valid and signed NOTIFY queries, it is under a replay attack which is discussed in Section 11.5. The key thing here is that the secondary is likely to be designed to be able to process much more traffic than the Hidden Primary hosted on a HNA.

This paragraph details how the secondary may limit the NOTIFY queries. Because the Hidden Primary may be renumbered, the secondary SHOULD NOT perform permanent IP filtering based on IP addresses. In addition, a given secondary may be shared among multiple Hidden Primaries which make filtering rules based on IP harder to set. The time at which a NOTIFY is sent by the Hidden Primary is not predictable. However, a flood of NOTIFY messages may be easily detected, as a NOTIFY originated from a given Homenet Zone is expected to have a very limited number of unique source IP addresses, even when renumbering is occurring. As a result, the secondary, MAY rate limit incoming NOTIFY queries.

On the Hidden Primary side, it is recommended that the Hidden Primary sends a NOTIFY as long as the zone has not been updated by the secondary. Multiple SOA queries may indicate the secondary is under attack.

11.3.3. Reflection Attacks involving the Public Authoritative Servers

Reflection attacks involving the Public Authoritative Server(s) are similar to attacks on any Outsourcing Infrastructure. This is not specific to the architecture described in this document, and thus are considered as out of scope.

In fact, one motivation of the architecture described in this document is to expose the Public Authoritative Server(s) to attacks instead of the HNA, as it is believed that the Public Authoritative Server(s) will be better able to defend itself.

11.4. Flooding Attack

The purpose of flooding attacks is mostly resource exhaustion, where the resource can be bandwidth, memory, or CPU for example.

One goal of the architecture described in this document is to limit the surface of attack on the HNA. This is done by outsourcing the DNS service to the Public Authoritative Server(s). By doing so, the HNA limits its DNS interactions between the Hidden Primary and the Synchronization Server. This limits the number of entities the HNA interacts with as well as the scope of DNS exchanges - NOTIFY, SOA, AXFR, IXFR.

The use of an authenticated channel with SIG(0) or TSIG between the HNA and the Synchronization Server, enables detection of illegitimate DNS queries, so appropriate action may be taken - like dropping the queries. If signatures are validated, then most likely, the HNA is under a replay attack, as detailed in Section 11.5

In order to limit the resource required for authentication, it is recommended to use TSIG that uses symmetric cryptography over SIG(0) that uses asymmetric cryptography.

11.5. Replay Attack

Replay attacks consist of an attacker either resending or delaying a legitimate message that has been sent by an authorized user or process. As the Hidden Primary and the Synchronization Server use an authenticated channel, replay attacks are mostly expected to use forged DNS queries in order to provide valid traffic.

From the perspective of an attacker, using a correctly authenticated DNS query may not be detected as an attack and thus may generate a response. Generating and sending a response consumes more resources than either dropping the query by the defender, or generating the query by the attacker, and thus could be used for resource exhaustion

attacks. In addition, as the authentication is performed at the DNS layer, the source IP address could be impersonated in order to perform a reflection attack.

Section 11.3 details how to mitigate reflection attacks and Section 11.4 details how to mitigate resource exhaustion. Both sections assume a context of DoS with a flood of DNS queries. This section suggests a way to limit the attack surface of replay attacks.

As SIG(0) and TSIG use inception and expiration time, the time frame for replay attack is limited. SIG(0) and TSIG recommends a fudge value of 5 minutes. This value has been set as a compromise between possibly loose time synchronization between devices and the valid lifetime of the message. As a result, better time synchronization policies could reduce the time window of the attack.

12. IANA Considerations

This document has no actions for IANA.

13. 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, Michael Richardson 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.

14. References

14.1. Normative References

- [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>>.

- [RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, DOI 10.17487/RFC1995, August 1996, <<https://www.rfc-editor.org/info/rfc1995>>.
- [RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", RFC 1996, DOI 10.17487/RFC1996, August 1996, <<https://www.rfc-editor.org/info/rfc1996>>.
- [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>>.
- [RFC2142] Crocker, D., "Mailbox Names for Common Services, Roles and Functions", RFC 2142, DOI 10.17487/RFC2142, May 1997, <<https://www.rfc-editor.org/info/rfc2142>>.
- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, DOI 10.17487/RFC2308, March 1998, <<https://www.rfc-editor.org/info/rfc2308>>.
- [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>>.
- [RFC2930] Eastlake 3rd, D., "Secret Key Establishment for DNS (TKEY RR)", RFC 2930, DOI 10.17487/RFC2930, September 2000, <<https://www.rfc-editor.org/info/rfc2930>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [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>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.

14.2. Informational References

- [I-D.howard-dnsop-ip6rdns]
Howard, L., "Reverse DNS in IPv6 for Internet Service Providers", draft-howard-dnsop-ip6rdns-00 (work in progress), June 2014.
- [I-D.ietf-homenet-naming-architecture-dhc-options]
Migault, D., Mrugalski, T., Griffiths, C., Weber, R., and W. Cloetens, "DHCPv6 Options for Homenet Naming Architecture", draft-ietf-homenet-naming-architecture-dhc-options-03 (work in progress), October 2015.
- [RFC1033] Lottor, M., "Domain Administrators Operations Guide", RFC 1033, DOI 10.17487/RFC1033, November 1987, <<https://www.rfc-editor.org/info/rfc1033>>.
- [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>>.
- [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>>.
- [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>>.

Appendix A. Document Change Log

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

-08

- 1: Clarification of the meaning of CPE. The architecture does not consider a single CPE. The CPE represents multiple functions.

-07:

- 1: Ray Hunter is added as a co-author.

-06:

- 2: Ray Hunter is added in acknowledgment.
- 3: Adding Renumbering section with comments from Dallas meeting
- 4: Replacing Master / Primary - Slave / Secondary

Security Consideration has been updated with Reflection attacks, flooding attacks, and replay attacks.

-05:

*Clarifying on handling different views:

- 1: How the CPE may be involved in the resolution and responds without necessarily requesting the Public Authoritative Server(s) (and eventually the Hidden Primary)
- 2: How to handle local scope resolution that is link-local, site-local and NAT IP addresses as well as Private domain names that the administrator does not want to publish outside the home network.

Adding a Privacy Considerations Section

Clarification on pro/cons outsourcing zone-signing

Documenting how to handle reverse zones

Adding reference to RFC 2308

-04:

*Clarifications on zone signing

- *Rewording

- *Adding section on different views

- *architecture clarifications

-03:

- *Simon's comments taken into consideration

- *Adding SOA, PTR considerations

- *Removing DNSSEC performance paragraphs on low power devices

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

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

-02:

- *remove interfaces: "Public Authoritative Server Naming Interface" is replaced by "Public Authoritative Server(s)y(ies)". "Public Authoritative Server Management Interface" is replaced by "Synchronization Server".

-01.3:

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

- *remove interactions with mdns and dhcp. Implementation dependent.

- *remove considerations on low powered devices

- *remove position toward homenet arch

- *remove problem statement section

-01.2:

- * add a CPE description to show that the architecture can fit CPEs

- * specification of the architecture for very low powered devices.

- * integrate mDNS and DHCP interactions with the Homenet Naming Architecture.

* Restructuring the draft. 1) We start from the homenet-arch draft to derive a Naming Architecture, then 2) we show why CPE need mechanisms that do not expose them to the Internet, 3) we describe the mechanisms.

* I remove the terminology and expose it in the figures A and B.

* remove the Front End Homenet Naming Architecture to Homenet Naming
-01:

* Added C. Griffiths as co-author.

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

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

-00: First version published.

Authors' Addresses

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

Phone: +1 (514) 452-2160
Email: daniel.migault@ericsson.com

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

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

Ray Hunter
Globis Consulting BV
Weegschaalstraat 3
5632CW Eindhoven
The Netherlands

Email: v6ops@globis.net
URI: <http://www.globis.net>

Chris Griffiths

Email: cgriffiths@gmail.com

Wouter Cloetens
SoftAtHome
vaartdijk 3 701
3018 Wijgmaal
Belgium

Email: wouter.cloetens@softathome.com

HOMENET
Internet-Draft
Intended status: Standards Track
Expires: April 30, 2018

D. Migault (Ed)
Ericsson
T. Mrugalski
ISC
C. Griffiths

R. Weber
Nominum
W. Cloetens
SoftAtHome
October 27, 2017

DHCPv6 Options for Homenet Naming Architecture
draft-ietf-homenet-naming-architecture-dhc-options-05

Abstract

The Homenet Naming Authority (HNA) is the designated device in charge of outsourcing the service to a third party, which requires setting up an architecture.

Such settings may be inappropriate for most end users. This document defines DHCPv6 options so any agnostic 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 April 30, 2018.

Copyright Notice

Copyright (c) 2017 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. Requirements notation	3
2. Terminology	3
3. Introduction	3
4. Protocol Overview	6
4.1. Architecture and DHCPv6 Options Overview	6
4.2. Mechanisms Securing DNS Transactions	9
4.3. Primary / Secondary Synchronization versus DNS Update . .	10
5. HNA Configuration	11
5.1. HNA Primary / Secondary Synchronization Configurations .	11
5.1.1. HNA / Synchronization Server	11
5.1.2. HNA / Reverse Synchronization Server	11
5.2. HNA DNS Data Handling and Update Policies	12
5.2.1. Homenet Zone Template	12
5.2.2. DNS (Reverse) Homenet Zone	12
6. Payload Description	13
6.1. Supported Authentication Methods Field	13
6.2. Update Field	14
6.3. Client Public Key Option	14
6.4. Zone Template Option	14
6.5. Synchronization Server Option	15
6.6. Reverse Synchronization Server Option	16
7. DHCP Behavior	17
7.1. DHCPv6 Server Behavior	17
7.2. DHCPv6 Client Behavior	18
7.3. DHCPv6 Relay Agent Behavior	18
8. IANA Considerations	18
9. Security Considerations	19
9.1. DNSSEC is recommended to authenticate DNS hosted data .	19
9.2. Channel between the HNA and ISP DHCP Server MUST be secured	19
9.3. HNAs are sensitive to DoS	19

10. Acknowledgments	19
11. References	19
11.1. Normative References	20
11.2. Informational References	21
Appendix A. Scenarios and impact on the End User	22
A.1. Base Scenario	22
A.2. Third Party Registered Homenet Domain	23
A.3. Third Party DNS Infrastructure	23
A.4. Multiple ISPs	25
Appendix B. Document Change Log	26
Authors' Addresses	27

1. Requirements notation

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

2. Terminology

The reader is expected to be familiar with [I-D.ietf-homenet-front-end-naming-delegation] and its terminology section. This section defines terms that have not been defined in [I-D.ietf-homenet-front-end-naming-delegation]:

- Client Public Key: designates a public key generated by the HNA.
This key is used as an authentication credential for the HNA.
- Homenet Zone Template: The template used as a basis to generate the Homenet Zone.
- DNS Template Server: The DNS server that hosts the Homenet Zone Template.
- Homenet Reverse Zone: The reverse zone file associated to the Homenet Zone.

3. Introduction

HNAs are usually constrained devices with reduced network and CPU capacities. As such, a HNA hosting on the Internet the authoritative naming service for its home network may become vulnerable to resource exhaustion attacks. Outsourcing the authoritative service to a third party avoids exposing the HNA to such attacks. This third party can be the ISP or any other independent third party.

Outsourcing the authoritative naming service to a third party requires setting up an architecture designated in this document as

the Outsourcing Infrastructure. These settings may be inappropriate for most end users that do not have the sufficient knowledge. To address this issue, this document proposes DHCPv6 options so any agnostic HNA can automatically set the Outsourcing Infrastructure. In most cases, these DHCPv6 options are sufficient and do not require any additional interaction from the end user, thus achieving a zero-config settings. In some other cases, the end user is expected to perform some limited manual configuration.

When the HNA is plugged, the DHCPv6 options described in the document enable:

- 1. To build the Homenet Zone: Building the Homenet Zone requires filling the zone with appropriated bindings such as bindings between the names and the IP addresses of the different devices of the home networks. How the HNA is aware of these binding is out of scope of the document. They may be provided, for example, by the DHCPv6 server hosted on the HNA. On the other hand, building the Homenet Zone also requires configuration parameters like the name of the Registered Domain Name associated to the home network or the Public Authoritative Server(s) the Homenet Zone is outsourced to. These configuration parameters are stored in the Homenet Zone Template. This document describes the Zone Template Option which carries the FQDN associated to the Homenet Zone Template. In order to retrieve the Homenet Zone Template, the HNA sends a query of type AXFR [RFC1034], [RFC5936].
- 2. To upload the Homenet Zone to the Synchronization Server, in charge of publishing the Homenet Zone on the Public Authoritative Server(s). This document describes the Synchronization Server Option that provides the FQDN of the appropriated server. Note that, the document does not consider whether the Homenet Zone is signed or not, and if signed, which entity is responsible to sign it. Such questions are out of the scope of the current document.
- 3. To upload the Homenet Reverse Zone to the Reverse Synchronization Server in charge of publishing the Homenet Reverse Zone on the Reverse Public Authoritative Server(s). This document describes the Reverse Synchronization Server Option that provides the FQDN of the appropriated server. Similarly to item 2., we do not consider in this document if the Homenet Reverse Zone is signed or not, and if signed who signs it.
- 4. To provide authentication credential (a public key) to the DHCP Server: Information stored in the Homenet Zone Template, the

Homenet Zone and Homenet Reverse Zone belongs to the HNA, and only the HNA should be able to update or upload these zones. To authenticate the HNA, this document defines the Client Public Key Option. This option is sent by the HNA to the DHCPv6 server and provides the Client Public Key the HNA uses to authenticate itself. This document does not describe mechanisms used to transmit the Client Public Key from the DHCPv6 server to the appropriate entities. If the DHCPv6 server is not able to provide the Client Public Key to the appropriated entities, then the end user is likely to provide manually the Client Public Key to these entities. This document illustrates two scenarios: one where the DHCPv6 server is responsible for distributing the Client Public Key to the Synchronization Servers and Reverse Synchronization Server. In the other scenarios, the Client Public Key is distributed out of band.

The DHCPv6 options described in this document make possible to configure an Outsourcing Infrastructure with no or little configurations from the end user. A zero-config setting is achieved if the the link between the HNA and the DHCPv6 server and the link between the DHCPv6 server and the various DNS servers (Homenet Zone Server, the Reverse Synchronization Server, Synchronization Server) are trusted. For example, one way to provide a trustworthy connection between the HNA and the DHCPv6 server is defined in [I-D.ietf-dhc-sedhcpv6]. When both links are trusted, the HNA is able to provide its authentication credentials (a Client Public Key) to the DHCPv6 server, that in turn forwards it to the various DNS servers. With the authentication credentials on the DNS servers, the HNA is able to securely update.

If the DHCPv6 server cannot provide the Client Public Key to one of these servers (most likely the Synchronization Server) and the HNA needs to interact with the server, then, the end user is expected to provide the HNA's Client Public Key to these servers (the Reverse Synchronization Server or the Synchronization Server) either manually or using other mechanisms. Such mechanisms are outside the scope of this document. In that case, the authentication credentials need to be provided every time the key is modified. Appendix A provides more details on how different scenarios impact the end users.

The remaining of this document is structured as follows. Section 4 provides an overview of the DHCPv6 options as well as the expected interactions between the HNA and the various involved entities. This section also provides an overview of available mechanisms to secure DNS transactions and update DNS data. Section 5 describes how the HNA may securely synchronize and update DNS data. Section 6 describes the payload of the DHCPv6 options and Section 7 details how

DHCPv6 client, server and relay agent behave. Section 8 lists the new parameters to be registered at the IANA, Section 9 provides security considerations. Finally, Appendix A describes how the HNA may behave and be configured regarding various scenarios.

4. Protocol Overview

This section provides an overview of the HNA's interactions with the Outsourcing Infrastructure in Section 4.1, and so the necessary for its setting. In this document, the configuration is provided via DHCPv6 options. Once configured, the HNA is expected to be able to update and publish DNS data on the different components of the Outsourcing Infrastructure. As a result authenticating and updating mechanisms play an important role in the specification. Section 4.2 provides an overview of the different authentication methods and Section 4.3 provides an overview of the different update mechanisms considered to update the DNS data.

4.1. Architecture and DHCPv6 Options Overview

This section illustrates how a HNA receives the necessary information via DHCPv6 options to outsource its authoritative naming service on the Outsourcing Infrastructure. For the sake of simplicity, this section assumes that the DHCPv6 server is able to communicate to the various DNS servers and to provide them the public key associated with the HNA. Once each server got the public key, the HNA can proceed to transactions in an authenticated and secure way.

This scenario has been chosen as it is believed to be the most popular scenario. This document does not ignore that scenarios where the DHCP Server does not have privileged relations with the Synchronization Server must be considered. These cases are discussed latter in Appendix A. Such scenario does not necessarily require configuration for the end user and can also be zero-config.

The scenario is represented in Figure 1.

- 1: The HNA provides its Client Public Key to the DHCP Server using a Client Public Key Option (OPTION_PUBLIC_KEY) and includes the following option codes in its Option Request Option (ORO): Zone Template Option (OPTION_DNS_ZONE_TEMPLATE), the Synchronization Server Option (OPTION_SYNC_SERVER) and the Reverse Synchronization Server Option (OPTION_REVERSE_SYNC_SERVER).
- 2: The DHCP Server makes the Client Public Key available to the DNS servers, so the HNA can secure its DNS transactions. How the Client Public Key is transmitted to the various DNS servers

is out of scope of this document. Note that the Client Public Key alone is not sufficient to perform the authentication and the key should be, for example, associated with an identifier, or the concerned domain name. How the binding is performed is out of scope of the document. It can be a centralized database or various bindings may be sent to the different servers. Figure 1 represents the specific case where the DHCP Server forwards the set (Client Public Key, Zone Template FQDN) to the DNS Template Server, the set (Client Public Key, IPv6 subnet) to the Reverse Synchronization Server and the set (Client Public Key, Registered Homenet Domain) to the Synchronization Server.

- 3: The DHCP Server responds to the HNA with the requested DHCPv6 options, i.e. the Client Public Key Option (OPTION_PUBLIC_KEY), Zone Template Option (OPTION_DNS_ZONE_TEMPLATE), Synchronization Server Option (OPTION_SYNC_SERVER), Reverse Synchronization Server Option (OPTION_REVERSE_SYNC_SERVER). Note that this step may be performed in parallel to step 2, or even before. In other words, there is no requirements that step 3 is conducted after step 2.
- 4: Upon receiving the Zone Template Option (OPTION_DNS_ZONE_TEMPLATE), the HNA performs an AXFR DNS query for the Zone Template FQDN. The exchange is authenticated according to the authentication methods defined in the Supported Authentication Methods field of the DHCP option. Once the HNA has retrieved the DNS Zone Template, the HNA can build the Homenet Zone and the Homenet Reverse Zone. Eventually the HNA signs these zones.
- 5: Once the Homenet Reverse Zone has been set, the HNA uploads the zone to the Reverse Synchronization Server. The Reverse Synchronization Server Option (OPTION_REVERSE_SYNC_SERVER) provides the Reverse Synchronization Server FQDN as well as the upload method, and the Supported Authentication Methods protocol to secure the upload.
- 6: Once the Homenet Zone has been set, the HNA uploads the zone to the Synchronization Server. The Synchronization Server Option (OPTION_SYNC_SERVER) provides the Synchronization Server FQDN as well as the upload method and the authentication method to secure the upload.




Figure 1: Protocol Overview

As described above, the HNA is likely to interact with various DNS content. More specifically, the HNA is likely to update the:

- Homenet Zone Template: if the configuration of the zone may be changed. This may include additional Public Authoritative Server(s), a different Registered Homenet Domain as the one initially proposed, or a redirection to another domain.
- Homenet Reverse Zone: every time a new device is connected or dis-connected.
- Homenet Zone: every time a new device is connected, dis-connected.

Step 2 and step 3 should be considered as independent steps and could be re-ordered. In fact, the DHCPv6 server does not have to wait for a confirmation from the DNS servers the Client Public Key has been properly received, and is operational by the DNS servers. The DHCP Server is expected to reply upon receiving the Client Public Key Option. The reply to the message with a Client Public Key Option from the DHCP Server is interpreted by the DHCPv6 client as a confirmation of the reception of the option by the DHCP Server only. It does not indicate whether the server had processed the option or

not. Debugging configurations errors or transmission error with one of the DNS servers is let to the HNA and thus is outside of the scope of the DHCPv6. First, it is unlikely a DNS server can validate that the Client Public Key will be operational for the HNA, as multiple causes of errors could occur. For example, the Client Public Key may have been changed during the transmission or by the DHCP Server, or the DNS server may be misconfigured. Second, the number of error codes would be too complex. In addition to multiple causes of errors, multiple architectures and multiple DNS servers may be involved. Third, this may cause significant DHCP Server performance degradation.

In fact, the HNA performs these updates in a secure manner. There are multiple ways to secure a DNS transaction and this document considers two mechanisms: nsupdate and primary/secondary synchronization. Section 4.2 describes the authentication method that may be use to secure the DNS transactions of the HNA. The appropriate authentication methods may, for example, be chosen according to the level of confidentiality or the level of authentication requested by the HNA transactions. Section 4.3 positions the nsupdate and primary/secondary synchronization mechanisms. The update appropriate update mechanism may depend on the for example on the update frequency or the size of the DNS data to update.

4.2. Mechanisms Securing DNS Transactions

Multiple protocols like IPsec [RFC4301] or TLS / DTLS [RFC5246] / [RFC6347] may be used to secure DNS transactions between the HNA and the DNS servers. This document limits its scope to authentication method that have been designed specifically for DNS. This includes DNSSEC [RFC4033], [RFC4034], [RFC4035] that authenticates and provides integrity protection of DNS data, TSIG [RFC2845], [RFC2930] that use a shared secret to secure a transaction between two end points and SIG(0) [RFC2931] authenticates the DNS packet exchanged.

The key issue with TSIG is that a shared secret must be negotiated between the HNA and the server. On the other hand, TSIG performs symmetric cryptography which is light in comparison with asymmetric cryptography used by SIG(0). As a result, over large zone transfer, TSIG may be preferred to SIG(0).

This document does not provide means to distribute shared secret for example using a specific DHCPv6 option. The only assumption made is that the HNA generates or is assigned a public key.

As a result, when the document specifies the transaction is secured with TSIG, it means that either the HNA and the DNS server have been

manually configured with a shared secret, or the shared secret has been negotiated using TKEY [RFC2930], and the TKEY exchanged are secured with SIG(0).

Exchanges with the DNS Template Server to retrieve the Homenet Zone Template may be protected by SIG(0), TSIG or DNSSEC. When DNSSEC is used, it means the DNS Template Server only provides integrity protection, and does not necessarily prevent someone else to query the Homenet Zone Template. In addition, DNSSEC is only a way to protect the AXFR queries transaction, in other words, DNSSEC cannot be used to secure updates. If DNSSEC is used to provide integrity protection for the AXFR response, the HNA should proceed to the DNSSEC signature checks. If signature check fails, it MUST reject the response. If the signature check succeeds, the HNA removes all DNSSEC related RRsets (DNSKEY, RRSIG, NSEC* ...) before building the Homenet Zone. In fact, these DNSSEC related fields are associated to the Homenet Zone Template and not the Homenet Zone.

Any update exchange should use SIG(0) or TSIG to authenticate the exchange.

4.3. Primary / Secondary Synchronization versus DNS Update

As updates only concern DNS zones, this document only considers DNS update mechanisms such as DNS update [RFC2136] [RFC3007] or a primary / secondary synchronization.

The Homenet Zone Template SHOULD be updated with DNS update as it contains static configuration data that is not expected to evolve over time.

The Homenet Reverse Zone and the Homenet Zone can be updated either with DNS update or using a primary / secondary synchronization. As these zones may be large, with frequent updates, we recommend to use the primary / secondary architecture as described in [I-D.ietf-homenet-front-end-naming-delegation]. The primary / secondary mechanism is preferred as it better scales and avoids DoS attacks: First the primary notifies the secondary the zone must be updated, and leaves the secondary to proceed to the update when possible. Then, the NOTIFY message sent by the primary is a small packet that is less likely to load the secondary. At last, the AXFR query performed by the secondary is a small packet sent over TCP (section 4.2 [RFC5936]) which makes unlikely the secondary to perform reflection attacks with a forged NOTIFY. On the other hand, DNS updates can use UDP, packets require more processing than a NOTIFY, and they do not provide the server the opportunity to postpone the update.

5. HNA Configuration

5.1. HNA Primary / Secondary Synchronization Configurations

The primary / secondary architecture is described in [I-D.ietf-homenet-front-end-naming-delegation]. The HNA hosts a Hidden Primary that synchronizes with a Synchronization Server or the Reverse Synchronization Server.

When the HNA is plugged its IP address may be unknown to the secondary. The section details how the HNA or primary communicates the necessary information to set up the secondary.

In order to set the primary / secondary configuration, both primary and secondaries must agree on 1) the zone to be synchronized, 2) the IP address and ports used by both primary and secondary.

5.1.1. HNA / Synchronization Server

The HNA is aware of the zone to be synchronized by reading the Registered Homenet Domain in the Homenet Zone Template provided by the Zone Template Option (OPTION_DNS_ZONE_TEMPLATE). The IP address of the secondary is provided by the Synchronization Server Option (OPTION_SYNC_SERVER).

The Synchronization Server has been configured with the Registered Homenet Domain and the Client Public Key that identifies the HNA. The only missing information is the IP address of the HNA. This IP address is provided by the HNA by sending a NOTIFY [RFC1996].

When the HNA has built its Homenet Zone, it sends a NOTIFY message to the Synchronization Servers. Upon receiving the NOTIFY message, the secondary reads the Registered Homenet Domain and checks the NOTIFY is sent by the authorized primary. This can be done using the shared secret (TSIG) or the public key (SIG(0)). Once the NOTIFY has been authenticated, the Synchronization Servers might consider the source IP address of the NOTIFY query to configure the primaries attributes.

5.1.2. HNA / Reverse Synchronization Server

The HNA is aware of the zone to be synchronized by looking at its assigned prefix. The IP address of the secondary is provided by the Reverse Synchronization Server Option (OPTION_REVERSE_SYNC_SERVER).

Configuration of the secondary is performed as illustrated in Section 5.1.1.

5.2. HNA DNS Data Handling and Update Policies

5.2.1. Homenet Zone Template

The Homenet Zone Template contains at least the related fields of the Public Authoritative Server(s) as well as the Homenet Registered Domain, that is SOA, and NS fields. This template might be generated automatically by the owner of the DHCP Server. For example, an ISP might provide a default Homenet Registered Domain as well as default Public Authoritative Server(s). This default settings should provide the HNA the necessary pieces of information to set the homenet naming architecture.

If the Homenet Zone Template is not subject to modifications or updates, the owner of the template might only use DNSSEC to enable integrity check.

On the other hand, the Homenet Zone Template might also be subject to modification by the HNA. The advantage of using the standard DNS zone format is that standard DNS update mechanism can be used to perform updates. These updates might be accepted or rejected by the owner of the Homenet Zone Template. Policies that defines what is accepted or rejected is out of scope of this document. However, this document assumes the Registered Homenet Domain is used as an index by the Synchronization Server, and SIG(0), TSIG are used to authenticate the HNA. As a result, the Registered Homenet Domain should not be modified unless the Synchronization Server can handle with it.

5.2.2. DNS (Reverse) Homenet Zone

The Homenet Zone might be generated from the Homenet Zone Template. How the Homenet Zone is generated is out of scope of this document. In some cases, the Homenet Zone might be the exact copy of the Homenet Zone Template. In other cases, it might be generated from the Homenet Zone Template with additional RRsets. In some other cases, the Homenet Zone might be generated without considering the Homenet Zone Template, but only considering specific configuration rules.

In the current document the HNA only sets a single zone that is associated with one single Homenet Registered Domain. The domain might be assigned by the owner of the Homenet Zone Template. This constraint does not prevent the HNA to use multiple domain names. How additional domains are considered is out of scope of this document. One way to handle these additional zones is to configure static redirections to the Homenet Zone using CNAME [RFC2181], [RFC1034], DNAME [RFC6672] or CNAME+DNAME [I-D.sury-dnsext-cname-dname].

6. Payload Description

This section details the payload of the DHCPv6 options. A few DHCPv6 options are used to advertise a server the HNA may be expected to interact with. Interaction may require to define update and authentication methods. Update fields are shared by multiple DHCPv6 options and are described in separate sections. Section 6.1 describes the Supported Authentication Method field, Section 6.2 describes the Update field, the remaining Section 6.3, Section 6.4, Section 6.5, Section 6.6 describe the DHCPv6 options.

6.1. Supported Authentication Methods Field

The Supported Authentication Methods field of the DHCPv6 option represented in Figure 2 indicates the authentication method supported by the DNS server. One of these mechanism MUST be chosen by the HNA in order to perform a transaction with the DNS server. See Section 4.2 for more details.




Figure 2: Supported Authentication Methods Filed

- DNS (Bit 0): indicates, when set to 1, that DNS without any security extension is supported.
- DNSSEC (Bit 1): indicates, when set to 1, that DNSSEC provides integrity protection. This can only be used for read operations like retrieving the Homenet Zone Template.
- SIG(0) (Bit 2): indicates, when set to 1, that transaction protected by SIG(0) are supported.
- TSIG (Bit 3): indicates, when set to 1, that transaction using TSIG is supported. Note that if a shared secret has not been previously negotiated between the two party, it should be negotiated using TKEY. The TKEY exchanges MUST be protected with SIG(0) even though SIG(0) is not supported.
- Remaining Bits (Bit 4-15): MUST be set to 0 by the DHCP Server and MUST be ignored by the DHCPv6 client.

A Supported Authentication Methods field with all bits set to zero indicates the operation is not permitted. The Supported

Authentication Methods field may be set to zero when updates operations are not permitted for the DNS Homenet Template. In any other case this is an error.

6.2. Update Field

The Update Field of the DHCPv6 option is represented in Figure 3. It indicates the update mechanism supported by the DNS server. See Section 4.3 for more details.




Figure 3: Update Field

- Primary / Secondary (Bit 0): indicates, when set to 1, that DNS Server supports data synchronization using a Primary / Secondary mechanism.
- DNS Update (Bit 1): indicates, when set to 1, that DNS Server supports data synchronization using DNS Updates.
- Remaining Bits (Bit 2-7): MUST be set to 0 by the DHCPv6 server and MUST be ignored by the DHCPv6 client.

6.3. Client Public Key Option

The Client Public Key Option (OPTION_PUBLIC_KEY) indicates the Client Public Key that is used to authenticate the HNA. This option is defined in [I-D.ietf-dhc-sedhcpv6].

6.4. Zone Template Option

The Zone Template Option (OPTION_DNS_ZONE_TEMPLATE) Option indicates the HNA how to retrieve the Homenet Zone Template. It provides a FQDN the HNA SHOULD query with a DNS query of type AXFR as well as the authentication methods associated to the AXFR query or the nsupdate queries. Homenet Zone Template update, if permitted MUST use the DNS Update mechanism.




Figure 4: Zone Template Option

- option-code: (16 bits): `OPTION_DNS_ZONE_TEMPLATE`, the option code for the Zone Template Option (TBD1).
- option-len (16 bits): length in octets of the option-data field as described in [RFC3315].
- Supported Authentication Methods(axfr) (16 bits): defines which authentication methods are supported by the DNS server. This field concerns the AXFR and consultation queries, not the update queries. See Section 6.1 for more details.
- Supported Authentication Methods (16 bits): defines which authentication methods are supported by the DNS server. This field concerns the update. See Section 6.1 for more details.
- Zone Template FQDN FQDN (variable): the FQDN of the DNS server hosting the Homenet Zone Template.

6.5. Synchronization Server Option

The Synchronization Server Option (`OPTION_SYNC_SERVER`) provides information necessary for the HNA to upload the Homenet Zone to the Synchronization Server. Finally, the option provides the authentication methods that are available to perform the upload. The upload is performed via a DNS primary / secondary architecture or DNS updates.




Figure 5: Synchronization Server Option

- option-code (16 bits): `OPTION_SYNC_SERVER`, the option code for the Synchronization Server Option (TBD2).
- option-len (16 bits): length in octets of the option-data field as described in [RFC3315].
- Supported Authentication Methods (16 bits): defines which authentication methods are supported by the DNS server. See Section 6.1 for more details.
- Update (8 bits): defines which update mechanisms are supported by the DNS server. See Section 4.3 for more details.
- Server Port (16 bits): defines the port the Synchronization Server is listening. When multiple transport layers may be used, a single and unique Server Port value applies to all the transport layers. In the case of DNS for example, Server Port value considers DNS exchanges using UDP and TCP.
- Synchronization Server FQDN (variable): the FQDN of the Synchronization Server.

6.6. Reverse Synchronization Server Option

The Reverse Synchronization Server Option (`OPTION_REVERSE_SYNC_SERVER`) provides information necessary for the HNA to upload the Homenet Zone to the Synchronization Server. The option provides the authentication methods that are available to perform the upload. The upload is performed via a DNS primary / secondary architecture or DNS updates.




Figure 6: Reverse Synchronization Server Option

- option-code (16 bits): `OPTION_REVERSE_SYNC_SERVER`, the option code for the Reverse Synchronization Server Option (TBD3).
- option-len (16 bits): length in octets of the option-data field as described in [RFC3315].
- Supported Authentication Methods (16 bits): defines which authentication methods are supported by the DNS server. See Section 6.1 for more details.
- Update (8 bits): defines which update mechanisms are supported by the DNS server. See Section 4.3 for more details.
- Server Port (16 bits): defines the port the Synchronization Server is listening.
- Reverse Synchronization Server FQDN (variable): The FQDN of the Reverse Synchronization Server.

7. DHCP Behavior

7.1. DHCPv6 Server Behavior

Sections 17.2.2 and 18.2 of [RFC3315] 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 DHCP Server sends the Zone Template Option, Synchronization Server Option, Reverse Synchronization Server Option when requested by the DHCPv6 client by including necessary option codes in its ORO.

The DHCP Server may receive a Client Public Key Option (OPTION_PUBLIC_KEY) from the HNA. Upon receipt of this DHCPv6 option, the DHCP Server SHOULD acknowledge the reception of the Client Public Key Option as described in Section 4.1 and communicate this credential to the available DNS Servers like the DNS Template Server, the Synchronization Server and the Reverse Synchronization Server, unless not configured to do so.

A HNA may update its Client Public Key by sending a new value in the Client Public Key Option (OPTION_PUBLIC_KEY) as this document assumes the link between the HNA and the DHCP Server is considered authenticated and trusted. The server SHOULD process received Client Public Key Option sent by the client (see step 2 in Section 4.1), unless not configured to do so.

7.2. DHCPv6 Client Behavior

The DHCPv6 client SHOULD send a Client Public Key Option (OPTION_PUBLIC_KEY) to the DHCP Server. This Client Public Key authenticates the HNA.

The DHCPv6 client sends a ORO with the necessary option codes: Zone Template Option, Synchronization Server Option and Reverse Synchronization Server Option.

Upon receiving a DHCP option described in this document in the Reply message, the HNA SHOULD retrieve or update DNS zones using the associated Supported Authentication Methods and update protocols, as described in Section 5.

7.3. DHCPv6 Relay Agent Behavior

There are no additional requirements for the DHCP Relay agents.

8. IANA Considerations

The DHCP options detailed in this document is:

- OPTION_DNS_ZONE_TEMPLATE: TBD1
- OPTION_SYNC_SERVER: TBD2
- OPTION_REVERSE_SYNC_SERVER: TBD3

9. Security Considerations

9.1. DNSSEC is recommended to authenticate DNS hosted data

It is recommended that the (Reverse) Homenet Zone is signed with DNSSEC. The zone may be signed by the HNA or by a third party. We recommend the zone to be signed by the HNA, and that the signed zone is uploaded.

9.2. Channel between the HNA and ISP DHCP Server MUST be secured

The channel MUST be secured because the HNA provides authentication credentials. Unsecured channel may result in HNA impersonation attacks.

The document considers that the channel between the HNA and the ISP DHCP Server is trusted. More specifically, the HNA is authenticated and the exchanged messages are protected. The current document does not specify how to secure the channel. [RFC3315] proposes a DHCP authentication and message exchange protection, [RFC4301], [RFC7296] propose to secure the channel at the IP layer.

9.3. HNAs are sensitive to DoS

HNA have not been designed for handling heavy load. The HNA are exposed on the Internet, and their IP address is publicly published on the Internet via the DNS. This makes the Home Network sensitive to Deny of Service Attacks. The resulting outsourcing architecture is described in [I-D.ietf-homenet-front-end-naming-delegation]. This document shows how the outsourcing architecture can be automatically set.

10. Acknowledgments

We would like to thank Marcin Siodelski and Bernie Volz 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.

11. References

11.1. Normative References

- [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>>.
- [RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", RFC 1996, DOI 10.17487/RFC1996, August 1996, <<https://www.rfc-editor.org/info/rfc1996>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC2930] Eastlake 3rd, D., "Secret Key Establishment for DNS (TKEY RR)", RFC 2930, DOI 10.17487/RFC2930, September 2000, <<https://www.rfc-editor.org/info/rfc2930>>.
- [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>>.
- [RFC3007] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, DOI 10.17487/RFC3007, November 2000, <<https://www.rfc-editor.org/info/rfc3007>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.

- [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>>.
- [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>>.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005, <<https://www.rfc-editor.org/info/rfc4035>>.
- [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>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [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>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [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>>.
- [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>>.

11.2. Informational 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.ietf-homenet-front-end-naming-delegation]

Migault, D., Weber, R., Hunter, R., Griffiths, C., and W. Cloetens, "Outsourcing Home Network Authoritative Naming Service", draft-ietf-homenet-front-end-naming-delegation-05 (work in progress), August 2016.

[I-D.sury-dnsextns-cname-dname]

Sury, O., "CNAME+DNS Name Redirection", draft-sury-dnsextns-cname-dname-00 (work in progress), April 2010.

Appendix A. Scenarios and impact on the End User

This section details various scenarios and discuss their impact on the end user.

A.1. Base Scenario

The base scenario is the one described in Section 4. It is typically the one of an ISP that manages the DHCP Server, and all DNS servers.

The end user subscribes to the ISP (foo), and at subscription time registers for example.foo as its Registered Homenet Domain example.foo. Since the ISP knows the Registered Homenet Domain and the Public Authoritative Server(s) the ISP is able to build the Homenet Zone Template.

The ISP manages the DNS Template Server, so it is able to load the Homenet Zone Template on the DNS Template Server.

When the HNA is plugged (at least the first time), it provides its Client Public Key to the DHCP Server. In this scenario, the DHCP Server and the DNS Servers are managed by the ISP so the DHCP Server can provide authentication credentials of the HNA to enable secure authenticated transaction between the HNA and these DNS servers. More specifically, credentials are provided to:

- Synchronization Server
- Reverse Synchronization Server
- DNS Template Server

The HNA can update the zone using DNS update or a primary / secondary configuration in a secure way.

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.

A.2. Third Party Registered Homenet Domain

This section considers the case when the end user wants its home network to use example.com as a Registered Homenet Domain instead of example.foo that has been assigned by the ISP. We also suppose that example.com is not managed by the ISP.

This can also be achieved without any configuration. 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-dnsex-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 A.1 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 A.1. 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.

A.3. 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 Synchronization Server and Public Authoritative Server(s) to a third party. All other DNS Servers DNS Template Server, Reverse Public Authoritative Server(s) and Reverse Synchronization Server remain managed by the ISP. The reason we consider that Reverse Public Authoritative Server(s) and Reverse Synchronization Server remains managed by the ISP are that the prefix is managed by the ISP, so outsourcing these resources requires some redirection agreement with the ISP. More specifically the ISP will need to configure the redirection on one of its Reverse DNS Servers. That said, outsourcing these resources is similar as outsourcing Synchronization Server and Public Authoritative Server(s) to a third party. Similarly, the DNS Template Server can be easily outsourced as detailed in this section

Outsourcing Synchronization Server and Public Authoritative Server(s) requires:

- 1) Updating the Homenet Zone Template: this can be easily done as detailed in Section 4.3 as the DNS Template Server is still managed by the ISP. Such modification can be performed once by any HNA. Once this modification has been performed, the HNA can be changed, the Client Public Key of the HNA may be changed, this does not need to be done another time. One can imagine a GUI on the HNA asking the end user to fill the field with Registered Homenet Domain, optionally Public Authoritative Server(s), with a button "Configure Homenet Zone Template".
- 2) Updating the DHCP Server Information. In fact the Reverse Synchronization Server returned by the ISP is modified. 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.
- 3) Upload the authentication credential of the HNA, that is the Client Public Key of the HNA, to the third party. Unless we use specific mechanisms, like communication between the DHCP Server and the third party, or a specific token that is plugged into the HNA, this operation is likely to be performed every time the HNA is changed, and every time the Client Public Key generated by the HNA is changed.

The main advantage of this scenario is that the DNS infrastructure is completely outsourced to the third party. Most likely the Client Public Key that authenticate the HNA need to be configured for every HNA. Configuration is expected to be HNA live-long.

A.4. Multiple ISPs

This scenario considers a HNA connected to multiple ISPs.

Firstly, suppose the HNA has been configured with the based scenarios exposed in Appendix A.1. The HNA has multiple interfaces, one for each ISP, and each of these interface is configured using DHCP. The HNA sends to each ISP its Client Public Key Option as well as a request for a Zone Template Option, a Synchronization Server Option and a Reverse Synchronization Server Option. Each ISP provides the requested DHCP options, with different values. Note that this scenario assumes, the home network has a different Registered Homenet Domain for each ISP as it is managed by the ISP. On the other hand, the HNA Client Public Key may be shared between the HNA and the multiple ISPs. The HNA builds the associate DNS(SEC) Homenet Zone, and proceeds to the various settings as described in Appendix A.1.

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. More specifically, multiple Registered Homenet Domains leads to multiple DNS(SEC) Homenet Zones. A basic implementation may erase the DNS(SEC) Homenet Zone that exists when it receives DHCPv6 options, and rebuild everything from scratch. This will work for an initial configuration but comes with a few drawbacks. First, updates to the DNS(SEC) Homenet Zone may only push to one of the multiple Registered Homenet Domain, the latest Registered Homenet Domain that has been set, and this is most likely expected to be almost randomly chosen as it may depend on the latency on each ISP network at the boot time. As a results, this leads to unsynchronized Registered Homenet Domains. Secondly, if the HNA handles in some ways resolution, only the latest Registered Homenet Domain set may be able to provide naming resolution in case of network disruption.

Secondly, suppose the HNA is connected to multiple ISP with a single Registered Homenet Domain. In this case, the one party 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 Synchronization Server 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 be chosen to host the Registered Homenet Domain. The DNS settings for each ISP is

described in Appendix A.2 and Appendix A.3. With the configuration described in Appendix A.2, 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 A.3, 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 A.1 and Appendix A.2 require the HNA to handle multiple Registered Homenet Domain, whereas Appendix A.3 does not have such requirement.

Appendix B. Document Change Log

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

-05: changing Master to Primary, Slave to Secondary

-04: Working Version Major modifications are:

- Re-structuring the draft: description and comparison of update and authentication methods have been integrated into the Overview section. a Configuration section has been created to describe both configuration and corresponding behavior of the HNA.
- Adding Ports parameters: Server Set can configure a port. The Port Server parameter have been added in the DHCPv6 option payloads because middle boxes may not be configured to let port 53 packets and it may also be useful to split servers among different ports, assigning each end user a different port.
- Multiple ISP scenario: In order to address comments, the multiple ISPs scenario has been described to explicitly show that the protocol and DHCPv6 options do not prevent a HNA connected to multiple independent ISPs.

-03: Working Version Major modifications are:

- Redesigning options/scope: according to feed backs received from the IETF89 presentation in the dhc WG.
- Redesigning architecture: according to feed backs received from the IETF89 presentation in the homenet WG, discussion with Mark and Lorenzo.

- 02: Working Version Major modifications are:
 - Redesigning options/scope: As suggested by Bernie Volz
- 01: Working Version Major modifications are:
 - Remove the DNS Zone file construction: As suggested by Bernie Volz
 - DHCPv6 Client behavior: Following options guide lines
 - DHCPv6 Server behavior: Following options guide lines
- 00: version published in the homenet WG. Major modifications are:
 - Reformatting of DHCPv6 options: Following options guide lines
 - DHCPv6 Client behavior: Following options guide lines
 - DHCPv6 Server behavior: Following options guide lines
- 00: First version published in dhc WG.

Authors' Addresses

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

Email: daniel.migault@ericsson.com

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

Email: tomasz.mrugalski@gmail.com
URI: <http://www.isc.org>

Chris Griffiths

Email: cgriffiths@gmail.com

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

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

Wouter Cloetens
SoftAtHome
vaartdijk 3 701
3018 Wijkmaal
Belgium

Email: wouter.cloetens@softathome.com