

Homenet Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 6, 2016

M. Stenberg
S. Barth
P. Pfister
Cisco Systems
July 5, 2015

Home Networking Control Protocol
draft-ietf-homenet-hncp-07

Abstract

This document describes the Home Networking Control Protocol (HNCP), an extensible configuration protocol and a set of requirements for home network devices on top of the Distributed Node Consensus Protocol (DNCP). It enables automated configuration of addresses, naming, network borders and the seamless use of a routing protocol.

Status of This Memo

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

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

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

This Internet-Draft will expire on January 6, 2016.

Copyright Notice

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

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

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

Table of Contents

1. Introduction	3
2. Requirements language	3
3. DNCP Profile	3
4. Common Links	4
5. Border Discovery	5
6. Autonomic Address Configuration	7
6.1. External Connections	7
6.1.1. External Connection TLV	7
6.1.2. Delegated Prefix TLV	8
6.1.3. Prefix Domain TLV	9
6.1.4. DHCP Data TLVs	10
6.2. Prefix Assignment	11
6.2.1. Prefix Assignment Algorithm Parameters	11
6.2.2. Assigned Prefix TLV	12
6.2.3. Making New Assignments	13
6.2.4. Applying Assignments	14
6.2.5. DHCPv6-PD Excluded Prefix Support	14
6.2.6. Downstream Prefix Delegation Support	15
6.3. Node Address Assignment	15
6.4. Local IPv4 and ULA Prefixes	16
6.5. Special Purpose Prefixes	17
7. Configuration of Hosts and non-HNCP Routers	18
7.1. DHCPv6 for Addressing or Configuration	18
7.2. Sending Router Advertisements	19
7.3. DHCPv6 for Prefix Delegation	19
7.4. DHCPv4 for Addressing and Configuration	20
7.5. Multicast DNS Proxy	20
8. Naming and Service Discovery	20
8.1. DNS Delegated Zone TLV	21
8.2. Domain Name TLV	22
8.3. Node Name TLV	23
9. Securing Third-Party Protocols	23
10. HNCP Versioning and Capabilities	24
11. Requirements for HNCP Routers	25
12. Security Considerations	27
12.1. Border Determination	27
12.2. Security of Unicast Traffic	28
12.3. Other Protocols in the Home	28
13. IANA Considerations	29
14. References	29
14.1. Normative references	29
14.2. Informative references	30

Appendix A. Changelog [RFC Editor: please remove]	31
Appendix B. Draft source [RFC Editor: please remove]	32
Appendix C. Implementation [RFC Editor: please remove]	32
Appendix D. Acknowledgements	32
Authors' Addresses	33

1. Introduction

HNCP synchronizes state across a small site in order to allow automated network configuration. The protocol enables use of border discovery, address prefix distribution [I-D.ietf-homenet-prefix-assignment], naming and other services across multiple links.

HNCP provides enough information for a routing protocol to operate without homenet-specific extensions. In homenet environments where multiple IPv6 source-prefixes can be present, routing based on source and destination address is necessary [RFC7368].

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 RFC 2119 [RFC2119].

3. DNCP Profile

HNCP is defined as a profile of DNCP [I-D.ietf-homenet-dncp] with the following parameters:

- o HNCP uses UDP datagrams on port HNCP-UDP-PORT as a transport over link-local scoped IPv6, using unicast and multicast (All-Homenet-Routers is the HNCP group address). Received datagrams with an IPv6 source or destination address which is not link-local scoped MUST be ignored. Unicast replies to multicast and unicast messages MUST be sent to the IPv6 source address and port of the original message. Each node MUST be able to receive (and potentially reassemble) UDP datagrams with a payload of at least 4000 bytes.
- o HNCP operates on multicast-capable interfaces only. HNCP routers MUST assign a unique 32-bit endpoint identifier to each interface for which HNCP is enabled. The value zero is reserved for internal purposes. Implementations MAY use a value equivalent to the `sin6_scope_id` for the given interface.

- o HNCP unicast traffic SHOULD be secured using DTLS [RFC6347] as described in DNCP if exchanged over unsecured links. UDP on port HNCP-DTLS-PORT is used for this purpose. A node implementing HNCP security MUST support the DNCP Pre-Shared Key method, SHOULD support the DNCP Certificate Based Trust Consensus and MAY support the PKI-based trust method.
- o HNCP uses opaque 32-bit node identifiers (DNCP_NODE_IDENTIFIER_LENGTH = 32). A node implementing HNCP SHOULD generate and use a random node identifier. If using a random node identifier and there is a node identifier collision, the node MUST immediately generate and use a new random node identifier which is not used by any other node.
- o HNCP nodes MUST ignore all Node State TLVs received via multicast on a link which has DNCP security enabled in order to prevent spoofing of node state changes.
- o HNCP nodes use the following Trickle parameters:
 - * k SHOULD be 1, as the timer reset when data is updated and further retransmissions should handle packet loss.
 - * Imin SHOULD be 200 milliseconds but MUST NOT be lower. Note: Earliest transmissions may occur at Imin / 2.
 - * Imax SHOULD be 7 doublings of Imin (i.e. 25.6 seconds) but MUST NOT be lower.
- o HNCP nodes MUST use the leading 64 bits of MD5 [RFC1321] as DNCP non-cryptographic hash function H(x).
- o HNCP nodes MUST use DNCP's keep-alive extension on all endpoints. The following parameters are suggested:
 - * Default keep-alive interval (DNCP_KEEPA_LIVE_INTERVAL): 20 seconds.
 - * Multiplier (DNCP_KEEPA_LIVE_MULTIPLIER): 2.1.

4. Common Links

HNCP uses the concept of Common Links for some of its applications. A Common Link usually refers to a link layer broadcast domain with certain properties and is used, e.g., to determine where prefixes should be assigned or which neighboring nodes participate in the election of a DHCP(v6) server. The Common Link is computed separately for each local interface, and it always contains the local

interface. Additionally, if the local interface is not in ad-hoc mode, it also contains the set of interfaces that are bidirectionally reachable from the given local interface, i.e. every remote interface of a remote node meeting all of the following requirements:

- o The local node publishes a Neighbor TLV with:
 - * Neighbor Node Identifier = remote node's node identifier
 - * Neighbor Endpoint Identifier = remote interface's endpoint identifier
 - * Endpoint Identifier = local interface's endpoint identifier
- o The remote node publishes a Neighbor TLV with:
 - * Neighbor Node Identifier = local node's node identifier
 - * Neighbor Endpoint Identifier = local interface's endpoint identifier
 - * Endpoint Identifier = remote interface's endpoint identifier

A node MUST be able to detect whether two of its local interfaces are connected, e.g. by detecting an identical remote interface being part of the Common Links of both local interfaces.

5. Border Discovery

HNCP router's interfaces are either internal, external or of a different category derived from the internal one. This section defines the border discovery algorithm. It is suitable for both IPv4 and IPv6 (single or dual-stack) and determines whether an HNCP interface is internal, external, or uses another fixed category. The algorithm is derived from the edge router interactions described in the Basic Requirements for IPv6 Customer Edge Routers [RFC7084]. This algorithm MUST be implemented by any router implementing HNCP.

The border discovery auto-detection algorithm works as follows, with evaluation stopping at first match:

1. If a fixed category is configured for the interface, it MUST be used.
2. If a delegated prefix could be acquired by running a DHCPv6 client on the interface, it MUST be considered external.

3. If an IPv4 address could be acquired by running a DHCPv4 client on the interface it MUST be considered external.
4. Otherwise the interface MUST be considered internal.

In order to avoid conflicts between border discovery and HNCP routers running DHCPv4 [RFC2131] or DHCPv6-PD [RFC3633] servers, each router MUST implement the following mechanism based on The User Class Option for DHCPv4 [RFC3004] and its DHCPv6 counterpart [RFC3315]:

- o An HNCP router running a DHCP client on an HNCP interface MUST include a DHCP User-Class consisting of the ASCII-String "HOMENET".
- o An HNCP router running a DHCP server on an HNCP interface MUST ignore or reject DHCP-Requests containing a DHCP User-Class consisting of the ASCII-String "HOMENET".

A router MUST allow setting a category of either auto-detected, internal or external for each interface which is suitable for both internal and external connections. In addition the following specializations of the internal category are defined to modify the local router behavior:

Leaf category: This declares an interface used by client devices only. Such an interface acts as an internal interface with the exception that HNCP or routing protocol traffic MUST NOT be sent on the interface, and all such traffic received on the interface MUST be ignored. This category SHOULD be supported.

Guest category: This declares an interface used by untrusted client devices only. In addition to the restrictions of the Leaf category, HNCP routers MUST enable firewalling rules such that connected devices are unable to reach other devices inside the HNCP network or query services advertised by them unless explicitly allowed. This category SHOULD be supported.

Ad-hoc category: This configures an interface to be ad-hoc (Section 4). Ad-hoc interfaces are considered internal but no assumption is made on the the link transitivity properties. Support for this category is OPTIONAL.

Hybrid category: This declares an interface to be internal while still running DHCPv4 and DHCPv6-PD clients on it. It is assumed that the link is under control of a legacy, trustworthy non-HNCP router, still within the same network. Detection of this category automatically in addition to manual configuration is out of scope of this document. Support for this category is OPTIONAL.

Each router MUST continuously scan each active interface that does not have a fixed category in order to dynamically reclassify it if necessary. The router therefore runs an appropriately configured DHCPv4 and DHCPv6 client as long as the interface is active including states where it considers the interface to be internal. The router SHOULD wait for a reasonable time period (5 seconds as a default), during which the DHCP clients can acquire a lease, before treating a newly activated or previously external interface as internal. Once it treats a certain interface as internal it MUST start forwarding traffic with appropriate source addresses between its internal interfaces and allow internal traffic to reach external networks according to the routes it publishes. Once a router detects an interface transitioning to external it MUST stop any previously enabled internal forwarding. In addition it SHOULD announce the acquired information for use in the network as described in later sections of this draft if the interface appears to be connected to an external network.

6. Autonomic Address Configuration

This section specifies how HNCP routers configure host and router addresses. At first border routers share information obtained from service providers or local configuration by publishing one or more External Connection TLVs. These contain other TLVs such as Delegated Prefix TLVs which are then used for prefix assignment. Finally, HNCP routers obtain addresses either statelessly or using a specific stateful mechanism and hosts and legacy routers are configured using SLAAC or DHCP.

In all TLVs specified in this section which include a prefix, IPv4 prefixes are encoded using the IPv4-mapped IPv6 addresses format [RFC4291]. The prefix length of such IPv4 prefix is set to 96 plus the IPv4 prefix length.

6.1. External Connections

Each HNCP router MAY obtain external connection information from one or more sources, e.g., DHCPv6-PD [RFC3633], NETCONF [RFC6241] or static configuration. This section specifies how such information is encoded and advertised.

6.1.1. External Connection TLV

An External Connection TLV is a container-TLV used to gather network configuration information associated with a single external connection. A node MAY publish an arbitrary number of instances of this TLV.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type: EXTERNAL-CONNECTION (33) | Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Nested TLVs                               |

```

The External Connection TLV is a container which:

- o MAY contain an arbitrary number of Delegated Prefix TLVs.
- o MUST NOT contain multiple Delegated Prefix TLVs with identical or overlapping prefixes. In such a situation, the External Connection TLV MUST be ignored.
- o MAY contain at most one DHCPv6 Data TLV and at most one DHCPv4 Data TLV encoding options associated with the External Connection but MUST NOT contain more than one of each otherwise the External Connection TLV MUST be ignored.
- o MAY contain other TLVs for future use. Such additional TLVs MUST be ignored.

6.1.1.2. Delegated Prefix TLV

The Delegated Prefix TLV is used by HNCP routers to advertise prefixes which are allocated to the whole network and will be used for prefix assignment. Any Delegated Prefix TLV MUST be nested in an External Connection TLV.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type: DELEGATED-PREFIX (34) | Length: >= 9 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Valid Lifetime                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Preferred Lifetime                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Prefix Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Prefix [+ nested TLVs]                       +
|

```

Valid Lifetime: The time in seconds the delegated prefix is valid. The value is relative to the point in time the Node-Data TLV was last published. It MUST be updated whenever the node republishes its Node-Data TLV.

Preferred Lifetime: The time in seconds the delegated prefix is preferred. The value is relative to the point in time the Node-Data TLV was last published. It MUST be updated whenever the node republishes its Node-Data TLV.

Prefix Length: The number of significant bits in the Prefix.

Prefix: Significant bits of the prefix padded with zeroes up to the next byte boundary.

Nested TLVs: Other TLVs included in the Delegated Prefix TLV and starting at the next 32-bit boundary following the end of the encoded prefix:

- * Zero or more Prefix Domain TLVs. In absence of any such TLV the prefix is assumed to be generated by an HNCP-router and for internal use only.
- * If the encoded prefix represents an IPv6 prefix, at most one DHCPv6 Data TLV MAY be included, and any included DHCPv4 Data TLV MUST be ignored.
- * If the prefix represents an IPv4 prefix (encoded as an IPv4-mapped IPv6 prefix), at most one DHCPv4 Data TLV MAY be included, and any included DHCPv6 Data TLV MUST be ignored.
- * It MAY contain other TLVs for future use. Such additional TLVs MUST be ignored.

6.1.3. Prefix Domain TLV

The Prefix Domain TLV contains information about the origin and applicability of a delegated prefix. This information can be used to determine whether prefixes for a certain domain (e.g. local reachability, internet connectivity) do exist or should be acquired and to make decisions about assigning prefixes to certain links or to fine-tune border firewalls. See Section 6.5 for a more in-depth discussion.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type: PREFIX-DOMAIN (43)   |   Length: >= 1   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Domain Type   |                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Value                                     |
|                                     |

```

Domain Type: The type of the domain identifier.

0 : Internet connectivity (no Value).

1-128 : Explicit destination prefix with the Domain Type being the actual length of the prefix (Value contains significant bits of the destination prefix padded with zeroes up to the next byte boundary).

129 : DNS Zone (Value contains an RFC 1035 [RFC1035] encoded DNS label sequence).

130 : Opaque UTF-8 string (e.g. for administrative purposes).

131-255: Reserved for future additions.

Value: A variable length identifier of the given type.

6.1.4. DHCP Data TLVs

Auxiliary connectivity information is encoded as a stream of DHCP options. Such TLVs MUST only be present in an External Connection TLV or a Delegated Prefix TLV. When included in an External Connection TLV, they MUST contain DHCP options which are relevant to the whole External Connection. When included in a Delegated Prefix, they MUST contain DHCP options which are specific to the Delegated Prefix.

The DHCPv6 Data TLV uses the following format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type: DHCPV6-DATA (37)   |           Length: > 0           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               DHCPv6 option stream            |

```

DHCPv6 option stream: DHCPv6 options encoded as specified in [RFC3315].

The DHCPv4 Data TLV uses the following format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type: DHCPV4-DATA (38)   |           Length: > 0           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               DHCPv4 option stream            |

```

DHCPv4 option stream: DHCPv4 options encoded as specified in [RFC2131].

6.2. Prefix Assignment

HNCP uses the Distributed Prefix Assignment Algorithm specified in [I-D.ietf-homenet-prefix-assignment] in order to assign prefixes to HNCP internal links and uses the terminology defined there.

6.2.1. Prefix Assignment Algorithm Parameters

All HNCP nodes running the prefix assignment algorithm MUST use the following parameters:

Node IDs: HNCP node identifiers are used. The comparison operation is defined as bit-wise comparison.

Set of Delegated Prefixes: The set of prefixes encoded in Delegated Prefix TLVs which are not strictly included in prefixes encoded in other Delegated Prefix TLVs. Note that Delegated Prefix TLVs included in ignored External Connection TLVs are not considered. It is dynamically updated as Delegated Prefix TLVs are added or removed.

Set of Shared Links: The set of Common Links associated with internal, leaf, guest or ad-hoc interfaces. It is dynamically updated as HNCP interfaces are added, removed, or switch from one category to another. When multiple interfaces are detected as belonging to the same Common Link, prefix assignment is disabled on all of these interfaces except one.

Set of Private Links: This document defines Private Links representing DHCPv6-PD clients or as a mean to advertise prefixes included in the DHCPv6 Exclude Prefix option. Other implementation-specific Private Links may be defined whenever a prefix needs to be assigned for a purpose that does not require a consensus with other HNCP routers.

Set of Advertised Prefixes: The set of prefixes included in Assigned Prefix TLVs advertised by other HNCP routers (Prefixes advertised by the local node are not in this set). The associated Advertised Prefix Priority is the priority specified in the TLV. The associated Shared Link is determined as follows:

- * If the Link Identifier is zero, the Advertised Prefix is not assigned on a Shared Link.

- * If the other node's interface identified by the Link Identifier is included in one of the Common Links used for prefix assignment, it is considered as assigned on the given Common Link.
- * Otherwise, the Advertised Prefix is not assigned on a Shared Link.

Advertised Prefixes as well as their associated priorities and associated Shared Links MUST be updated as Assigned Prefix TLVs are added, updated or removed, and as Common Links are modified.

ADOPT_MAX_DELAY: The default value is 0 seconds (i.e. prefix adoption MAY be done instantly).

BACKOFF_MAX_DELAY: The default value is 4 seconds.

RANDOM_SET_SIZE: The default value is 64.

Flooding Delay: The default value is 5 seconds.

Default Advertised Prefix Priority: When a new assignment is created or an assignment is adopted - as specified in the prefix assignment algorithm routine - the default Advertised Prefix Priority to be used is 2.

6.2.2. Assigned Prefix TLV

Published Assigned Prefixes MUST be advertised using the Assigned Prefix TLV:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Type: ASSIGNED-PREFIX (35)  |          Length: >= 6          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Endpoint Identifier              |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Rsv. | Prty. | Prefix Length |                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Prefix                           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Endpoint Identifier: The endpoint identifier of the local interface that belongs to the Common Link the prefix is assigned to, or 0 if the Common Link is a Private Link (e.g., when the prefix is assigned for downstream prefix delegation).

Rsv.: Bits are reserved for future use. They MUST be set to zero when creating this TLV, and their value MUST be ignored when processing the TLV.

Prt: The Advertised Prefix Priority from 0 to 15.

0-1 : Low priorities.

2 : Default priority.

3-7 : High priorities.

8-11 : Administrative priorities. MUST NOT be used unless configured otherwise.

12-14: Reserved for future use.

15 : Provider priorities. MAY only be used by the router advertising the corresponding delegated prefix and based on static or dynamic configuration (e.g., for excluding a prefix based on DHCPv6-PD Prefix Exclude Option [RFC6603]).

Prefix Length: The number of significant bits in the Prefix field.

Prefix: The significant bits of the prefix padded with zeroes up to the next byte boundary.

6.2.3. Making New Assignments

Whenever the Prefix Assignment Algorithm subroutine is run on a Common Link and whenever a new prefix may be assigned (case 1 of the subroutine), the decision of whether the assignment of a new prefix is desired MUST follow these rules:

If the Delegated Prefix TLV contains a DHCPv4 or DHCPv6 Data TLV, and the meaning of one of the DHCP options is not understood by the HNCP router, the creation of a new prefix is not desired. This rule applies to TLVs inside Delegated Prefix TLVs but not to those inside External Connection TLVs.

If the remaining preferred lifetime of the prefix is 0 and there is another delegated prefix of the same IP version used for prefix assignment with a non-null preferred lifetime, the creation of a new prefix is not desired.

Otherwise, the creation of a new prefix is desired, if the Delegated Prefix is either locally generated (does not have any Prefix Domain TLVs) or intended for internet access (has a Prefix

Domain TLV of type 0). Local desirability policies MAY override or provide additional desirability rules for delegated prefixes, e.g., by matching different Prefix Domain TLV values.

If the considered delegated prefix is an IPv6 prefix, and whenever there is at least one available prefix of length 64, a prefix of length 64 MUST be selected unless configured otherwise. In case no prefix of length 64 would be available, a longer prefix MAY be selected even without configuration.

If the considered delegated prefix is an IPv4 prefix (Section 6.4 details how IPv4 delegated prefixes are generated), a prefix of length 24 SHOULD be preferred.

In any case, a router MUST support a mechanism suitable to distribute addresses from the considered prefix if the link is intended to be used by clients. In this case a router assigning an IPv4 prefix MUST support the L-capability and a router assigning an IPv6 prefix MUST support serving router advertisements. In addition if an assigned IPv6 prefix is not suitable for Stateless Address Autoconfiguration the router MUST also support the H-capability as defined in Section 10.

6.2.4. Applying Assignments

The prefix assignment algorithm indicates when a prefix is applied to the respective Common Link. When that happens each router connected to said link:

- MUST create an appropriate route for said prefix, indicating it is directly reachable on the respective link and advertise said route using the chosen routing protocol.

- MUST participate in the client configuration election as described in Section 7, if the link is intended to be used by clients.

- MAY add an address from said prefix to the respective network interface as described in Section 6.3, e.g., if it is to be used as source for locally originating traffic.

6.2.5. DHCPv6-PD Excluded Prefix Support

Whenever a DHCPv6 Prefix Exclude option [RFC6603] is received with a delegated prefix, the excluded prefix MUST be advertised as assigned to a Private Link with the maximum priority (i.e. 15).

The same procedure MAY be applied in order to exclude prefixes obtained by other means of configuration.

6.2.6. Downstream Prefix Delegation Support

When an HNCP router receives a request for prefix delegation, it SHOULD assign one prefix per delegated prefix in the network. This set of assigned prefix is then delegated to the client, after it has been applied as described in the Prefix Assignment Algorithm. Each client MUST be considered as an independent Private Link and delegation MUST be based on the same set of Delegated Prefixes as the one used for Common Link prefix assignments.

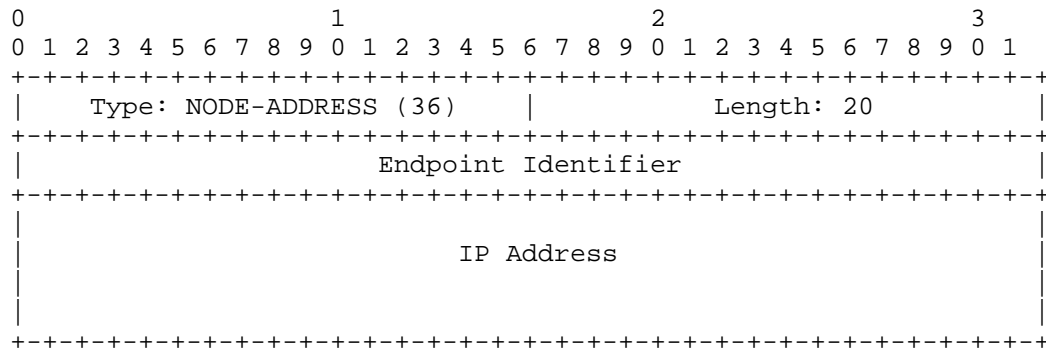
The assigned prefixes MUST NOT be given to clients before they are applied, and MUST be withdrawn whenever they are destroyed. As an exception to this rule, in order to shorten delays of processed requests, a router MAY prematurely give out a prefix which is advertised but not yet applied if it does so with a valid lifetime of not more than 30 seconds and ensures removal or correction of lifetimes as soon as possible.

6.3. Node Address Assignment

This section specifies how HNCP nodes reserve addresses for their own use. Nodes MAY, at any time, try to reserve a new address from any applied Assigned Prefix. Each HNCP router MUST announce at least one IPv6 address and - if it supports IPv4 - at least one IPv4 address, whenever matching prefixes are assigned to at least one of its Common Links. These addresses are published using Node Address TLVs and used to locally reach HNCP nodes for other services. Nodes SHOULD NOT create and announce more than one assignment per IP version to avoid cluttering the node data with redundant information unless a special use case requires it.

Stateless assignment based on Modified EUI64 interface identifiers [RFC4291] SHOULD be used for address assignment whenever possible, otherwise (e.g., for IPv4) the following method MUST be used instead: For any assigned prefix for which SLAAC cannot be used, the first quarter of the addresses are reserved for routers HNCP based address assignments, whereas the last three quarters are left to the DHCPv6 (resp. DHCPv4) elected router (Section 10 specifies the DHCP server election process). For instance, if the prefix 192.0.2.0/24 is assigned and applied to a Common Link, addresses included in 192.0.2.0/26 are reserved for HNCP nodes and the remaining addresses are reserved for the elected DHCPv4 server.

HNCP routers assign themselves addresses using the Node Address TLV:



Endpoint Identifier: The endpoint identifier of the local interface that belongs to the Common Link the prefix is assigned to, or 0 if it is not assigned on an HNCP enabled link.

IP Address: The globally scoped IPv6 address, or the IPv4 address encoded as an IPv4-mapped IPv6 address [RFC4291].

The process of obtaining addresses is specified as follows:

- o A router **MUST NOT** start advertising an address if it is already advertised by another router.
- o An assigned address **MUST** be in the first quarter of an assigned prefix currently applied on a Common Link which includes the interface specified by the endpoint identifier.
- o An address **MUST NOT** be used unless it has been advertised for at least ADDRESS_APPLY_DELAY consecutive seconds, and is still currently being advertised. The default value for ADDRESS_APPLY_DELAY is 3 seconds.
- o Whenever the same address is advertised by more than one node, all but the one advertised by the node with the highest node identifier **MUST** be removed.

6.4. Local IPv4 and ULA Prefixes

HNCP routers can create an ULA or private IPv4 prefix to enable connectivity between local devices. These prefixes are inserted in HNCP as if they were delegated prefixes. The following rules apply:

An HNCP router **SHOULD** create a ULA prefix if there is no other IPv6 prefix with a preferred time greater than 0 in the network. It **MAY** also do so, if there are other delegated IPv6 prefixes, but none of which is locally generated (i.e., without any Prefix

Domain TLV) and has a preferred time greater than 0. However, it MUST NOT do so otherwise. In case multiple locally generated ULA prefixes are present, only the one published by the node with the highest node identifier is kept among those with a preferred time greater than 0 - if there is any.

An HNCP router MUST create a private IPv4 prefix [RFC1918] whenever it wishes to provide IPv4 internet connectivity to the network and no other private IPv4 prefix with internet connectivity currently exists. It MAY also enable local IPv4 connectivity by creating a private IPv4 prefix if no IPv4 prefix exists but MUST NOT do so otherwise. In case multiple IPv4 prefixes are announced, only the one published by the node with the highest node identifier is kept among those with a Prefix Domain of type 0 - if there is any. The router publishing a prefix with internet connectivity MUST announce an IPv4 default route using the routing protocol and perform NAT on behalf of the network as long as it publishes the prefix, other routers in the network MAY choose not to.

Creation of such ULA and IPv4 prefixes MUST be delayed by a random timespan between 0 and 10 seconds in which the router MUST scan for other nodes trying to do the same.

When a new ULA prefix is created, the prefix is selected based on the configuration, using the last non-deprecated ULA prefix, or generated based on [RFC4193].

6.5. Special Purpose Prefixes

Some prefixes may have a special meaning and are not regularly used for internal or internet connectivity, instead they may provide access to special services like VPNs, sensor networks, VoIP, IPTV, etc. Care must be taken that these prefixes are properly integrated and dealt with in the network, in order to avoid breaking connectivity for devices who are not aware of their special characteristics.

Special purpose prefixes are distinguished using Prefix Domain TLVs (Section 6.1.3). Their contents MAY be partly opaque to HNCP nodes, and their identification and usage depends on local policy. However the following general rules MUST be adhered to:

Special rules apply when making address assignments for prefixes with Prefix Domain TLVs other than type 0, as described in Section 6.2.3

In presence of any type 1 to 128 Prefix Domain TLV the prefix is specialized to reach destinations denoted by any such Prefix Domain TLV, i.e. in absence of a type 0 Prefix Domain TLV it is not usable for general internet connectivity. An HNCP router MAY enforce this restriction with appropriate packet filtering rules to provide increased security.

The presence of a type 129 (DNS zone) Prefix Domain TLV indicates that the delegated prefix or its associated external connection is specialized to reach destinations within the given DNS zone. An HNCP router providing name resolving services SHOULD prefer DNS servers listed in the associated external connection's DHCPv4 or DHCPv6 Data TLVs when resolving domains from that zone.

7. Configuration of Hosts and non-HNCP Routers

HNCP routers need to ensure that hosts and non-HNCP downstream routers on internal links are configured with addresses and routes. Since DHCP-clients can usually only bind to one server at a time, a per-link and per-service election takes place.

HNCP routers may have different capabilities for configuring downstream devices and providing naming services. Each router MUST therefore indicate its capabilities as specified in Section 10 in order to participate as a candidate in the election.

7.1. DHCPv6 for Addressing or Configuration

In general Stateless Address Autoconfiguration is used for client configuration for its low overhead and fast renumbering capabilities, however stateful DHCPv6 can be used in addition by administrative choice, to e.g. collect hostnames and use them to provide naming services or whenever stateless configuration is not applicable.

The designated stateful DHCPv6 server for a Common Link (Section 4) is elected based on the capabilities described in Section 10. The winner is the router (connected to the Common Link) advertising the greatest H-capability. In case of a tie, Capability Values and node identifiers are considered (greatest value is elected). The elected router MUST serve stateful DHCPv6 and MUST provide naming services for acquired hostnames as outlined in Section 8. Stateful addresses SHOULD be assigned in a way not hindering fast renumbering even if the DHCPv6 server or client do not support the DHCPv6 reconfigure mechanism. In case no router was elected, stateful DHCPv6 is not provided and each router assigning IPv6-prefixes on said link MUST provide stateless DHCPv6 service.

7.2. Sending Router Advertisements

All HNCP routers MUST send Router Advertisements periodically via multicast and via unicast in response to Router Solicitations.

- o The "Managed address configuration" flag MUST be set whenever a router connected to the link is advertising a non-null H-capability and MUST NOT be set otherwise. The "Other configuration" flag MUST always be set.
- o The default Router Lifetime MUST be set to an appropriate non-null value whenever an IPv6 default route is known in the HNCP network and MUST be set to zero otherwise.
- o A Prefix Information Option MUST be added for each assigned and applied IPv6 prefix on the given link. The autonomous address-configuration flag MUST be set whenever the prefix is suitable for stateless configuration. The preferred and valid lifetimes MUST be smaller than the preferred and valid lifetimes of the delegated prefix the prefix is from. When a prefix is removed, it MUST be deprecated as specified in [RFC7084].
- o A Route Information Option [RFC4191] MUST be added for each delegated IPv6 prefix known in the HNCP network. Additional ones SHOULD be added for each non-default IPv6 route with an external destination prefix advertised by the routing protocol.
- o A Recursive DNS Server Option and a DNS Search List Option MUST be included with appropriate contents.
- o To allow for optimized routing decisions for clients on the local link routers SHOULD adjust their Default Router Preference and Route Preferences [RFC4191] so that the priority is set to low if the next hop of the default or more specific route is on the same interface as the Route Advertisement being sent on. Similarly the router MAY use the high priority if it is certain it has the best metric of all routers on the link for all routes known in the network with the respective destination.

Every router sending Router Advertisements MUST immediately send an updated Router Advertisement via multicast as soon as it notices a condition resulting in a change of any advertised information.

7.3. DHCPv6 for Prefix Delegation

The designated DHCPv6 server for prefix-delegation on a Common Link is elected based on the capabilities described in Section 10. The winner is the router (connected to the Common Link) advertising the

greatest P-capability. In case of a tie, Capability Values are compared, and router with the greatest value is elected. In case of another tie, the router with the highest node identifier is elected among the routers with tied Capability Values. The elected router MUST provide prefix-delegation services [RFC3633] on the given link and follow the rules in Section 6.2.6.

7.4. DHCPv4 for Addressing and Configuration

The designated DHCPv4 server on a Common Link (Section 4) is elected based on the capabilities described in Section 10. The winner is the router (connected to the Common Link) advertising the greatest L-capability. In case of a tie, Capability Values are compared, and router with the greatest value is elected. In case of another tie, the router with the highest node identifier is elected among the routers with tied Capability Values. The elected router MUST provide DHCPv4 services on the given link.

The DHCPv4 serving router MUST announce itself as router [RFC2132] to clients if and only if there is an IPv4 default route known in the network. In addition, the router SHOULD announce a Classless Static Route Option [RFC3442] for each non-default IPv4 route advertised in the routing protocol with an external destination.

DHCPv4 lease times SHOULD be short (i.e. not longer than 5 minutes) in order to provide reasonable response times to changes.

7.5. Multicast DNS Proxy

The designated MDNS [RFC6762] proxy on a Common Link is elected based on the capabilities described in Section 10. The winner is the router (connected to the Common Link) advertising the greatest M-capability. In case of a tie, Capability Values are compared, and router with the greatest value is elected. In case of another tie, the router with the highest node identifier is elected among the routers with tied Capability Values. The elected router MUST provide an MDNS-proxy on the given link and announce it as described in Section 8.

8. Naming and Service Discovery

Network-wide naming and service discovery can greatly improve the user-friendliness of a network. The following mechanism provides means to setup and delegate naming and service discovery across multiple HNCP routers.

Each HNCP router SHOULD provide and announce an auto-generated or user-configured name for each internal Common Link (Section 4) for

which it is the designated DHCPv4, stateful DHCPv6 server, MDNS proxy, or for which it provides forward or reverse DNS services on behalf of connected devices. HNCP routers providing name resolving services MUST use the included DNS server address to resolve names belonging to the zone.

Each HNCP router SHOULD announce a node name for itself to be easily reachable and MAY do so on behalf of other devices. HNCP routers providing name resolving services MUST resolve these names to their respective IP addresses.

The following TLVs are defined and MUST be supported by all nodes implementing naming and service discovery:

8.1. DNS Delegated Zone TLV

This TLV is used to announce a forward or reverse DNS zone delegation in the HNCP network. Its meaning is roughly equivalent to specifying an NS and A/AAAA record for said zone. There MUST NOT be more than one delegation for the same zone in the whole DNCP network. In case of a conflict the announcement of the node with the highest node identifier takes precedence and all other nodes MUST cease to announce the conflicting TLV.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type: DNS-DELEGATED-ZONE (39) | Length: >= 17 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IP Address                               |
|                               |                                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Reserved |L|B|S|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Zone (DNS label sequence - variable length)
|

```

IP Address : The IPv6 address of the authoritative DNS server for the zone; IPv4 addresses are represented as IPv4-mapped addresses [RFC4291]. The special value of :: (all-zero) means the delegation is available in the global DNS-hierarchy.

Reserved : Those bits MUST be set to zero when creating the TLV and ignored when parsing it unless defined in a later specification.

L-bit : DNS-SD [RFC6763] Legacy-Browse, indicates that this delegated zone should be included in the network's DNS-SD legacy

browse list of domains at `lb._dns-sd._udp.(DOMAIN-NAME)`. Local forward zones SHOULD have this bit set, reverse zones SHOULD NOT.

B-bit : (DNS-SD [RFC6763] Browse) indicates that this delegated zone should be included in the network's DNS-SD browse list of domains at `b._dns-sd._udp. (DOMAIN-NAME)`. Local forward zones SHOULD have this bit set, reverse zones SHOULD NOT.

S-bit : (fully-qualified DNS-SD [RFC6763] domain) indicates that this delegated zone consists of a fully-qualified DNS-SD domain, which should be used as base for DNS-SD domain enumeration, i.e. `_dns-sd._udp.(Zone)` exists. Forward zones MAY have this bit set, reverse zones MUST NOT. This can be used to provision DNS search path to hosts for non-local services (such as those provided by an ISP, or other manually configured service providers). Zones with this flag SHOULD be added to the search domains advertised to clients.

Zone : The label sequence of the zone, encoded as the domain names are encoded DNS messages as specified in [RFC1035]. The last label in the zone MUST be empty.

8.2. Domain Name TLV

This TLV is used to indicate the base domain name for the network. It is the zone used as a base for all non fully-qualified delegated zones and node names. In case of conflicts the announced domain of the node with the greatest node identifier takes precedence. By default, i.e., if no node advertises such a TLV., ".home" is used. This TLV MUST NOT be announced unless the domain name was explicitly configured by an administrator.

```

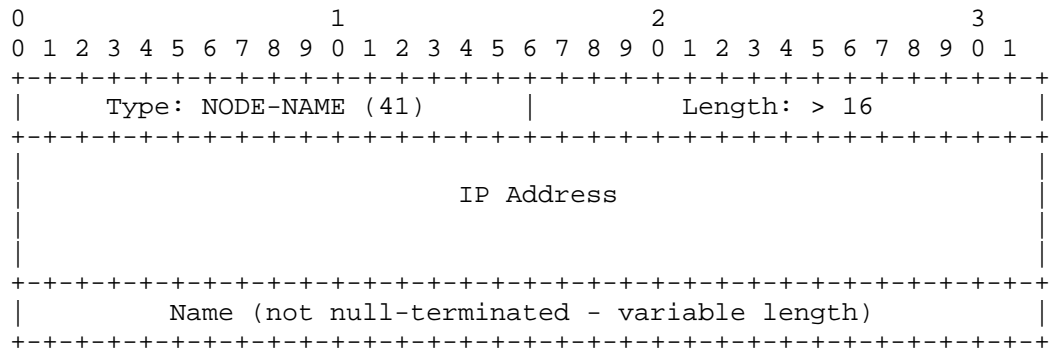
0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Type: DOMAIN-NAME (40)  |          Length: > 0          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Domain (DNS label sequence - variable length)          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Domain: The label sequence encoded according to [RFC1035]. Compression MUST NOT be used. The zone MUST end with an empty label.

8.3. Node Name TLV

This TLV is used to assign the name of a node in the network to a certain IP address. In case of conflicts the announcement of the node with the greatest node identifier for a name takes precedence and all other nodes MUST cease to announce the conflicting TLV.




IP Address: The IP address associated with the name. IPv4 addresses are encoded using IPv4-mapped IPv6 addresses.

Name: The name of the node as a single DNS label (up to 63 characters, no leading length byte).

9. Securing Third-Party Protocols

Pre-shared keys (PSKs) are often required to secure IGPs and other protocols which lack support for asymmetric security. The following mechanism manages PSKs using HNCP to enable bootstrapping of such third-party protocols and SHOULD therefore be used if such a need arises. The following rules define how such a PSK is managed and used:

- o If no Managed-PSK-TLV is currently being announced, an HNCP router MUST create one after a random delay of 0 to 10 seconds with a 32 bytes long random key and add it to its node data.
- o In case multiple routers announce such a TLV at the same time, all but the one with the greatest node identifier stop advertising it and adopt the remaining one.
- o The router currently advertising the Managed-PSK-TLV must generate and advertise a new random one whenever an unreachable node is purged as described in DNCP.




PSKs for individual protocols are derived from the random PSK through the use of HMAC-SHA256 [RFC6234] with a pre-defined per-protocol HMAC-key in ASCII-format. The following HMAC-keys are currently defined to derive PSKs for the respective protocols:

"ROUTING": to be used for IGPs

10. HNCP Versioning and Capabilities

Multiple versions of HNCP based on compatible DNCP profiles may be present in the same network when transitioning between HNCP versions and HNCP routers may have different capabilities to support clients. The following mechanism describes a way to announce the currently active version and User-agent of a node. Each node MUST include an HNCP-Version-TLV in its Node Data and MUST ignore (except for DNCP synchronization purposes) any TLVs with a type greater than 32 published by nodes not also publishing an HNCP-Version TLV or publishing such a TLV with a different Version number.

Capabilities are indicated by setting M, P, H and L fields in the TLV. The "capability value" is a metric indicated by interpreting the bits as an integer, i.e. $(M \ll 12 \mid P \ll 8 \mid H \ll 4 \mid L)$.



Version: Indicates which version of HNCP is currently in use by this particular node. It MUST be set to 1. Nodes with different versions are considered incompatible.

Reserved: Bits are reserved for future use. They MUST be set to zero when creating this TLV, and their value MUST be ignored when processing the TLV.

M-capability: Priority value used for electing the on-link MDNS [RFC6762] proxy. It MUST be set to some value between 1 and 7 included (4 is the default) if the router is capable of proxying MDNS and 0 otherwise. The values 8-15 are reserved for future use.

P-capability: Priority value used for electing the on-link DHCPv6-PD server. It MUST be set to some value between 1 and 7 included (4 is the default) if the router is capable of providing prefixes through DHCPv6-PD (Section 6.2.6) and 0 otherwise. The values 8-15 are reserved for future use.

H-capability: Priority value used for electing the on-link DHCPv6 server offering non-temporary addresses. It MUST be set to some value between 1 and 7 included (4 is the default) if the router is capable of providing such addresses and 0 otherwise. The values 8-15 are reserved for future use.

L-capability: Priority value used for electing the on-link DHCPv4 server. It MUST be set to some value between 1 and 7 included (4 is the default) if the router is capable of running a legacy DHCPv4 server offering IPv4 addresses to clients and 0 otherwise. The values 8-15 are reserved for future use.

User-Agent: The user-agent is a human-readable UTF-8 string that describes the name and version of the current HNCP implementation.

11. Requirements for HNCP Routers

Each router implementing HNCP is subject to the following requirements:

- o It MUST implement HNCP-Versioning, Border Discovery, Prefix Assignment and Configuration of hosts and non-HNCP routers as defined in this document.
- o It MUST implement and run the method for securing third-party protocols whenever it uses the security mechanism of HNCP.

- o It SHOULD implement support for the Service Discovery and Naming TLVs as defined in this document.
- o It MUST implement and run a routing protocol appropriate for the given link type on all of the interfaces it sends and receives HNCP traffic on. The protocol MUST support source-specific routes and MUST correctly propagate those also for the external destinations that may have only implicit source-specific information, such as a combination of a DHCPv6 PD-derived prefix and a non-source-specific default route.
- o It MUST use adequate security mechanisms for the routing protocol on any interface where it also uses the security mechanisms of HNCP. If the security mechanism is based on a PSK it MUST use a PSK derived from the Managed-PSK to secure the IGP.
- o It MAY be able to provide connectivity to IPv4-devices using DHCPv4.
- o It SHOULD be able to delegate prefixes to legacy IPv6 routers using DHCPv6-PD.
- o In addition, normative language of Basic Requirements for IPv6 Customer Edge Routers [RFC7084] applies with the following adjustments:
 - * The section "WAN-Side Configuration" applies to HNCP interfaces classified as external.
 - * If the CE sends a size-hint as indicated in WPD-2, the hint MUST NOT be determined by the number of LAN-interfaces of the CE, but SHOULD instead be large enough to at least accommodate prefix assignments announced for existing delegated or ULA-prefixes, if such prefixes exist and unless explicitly configured otherwise.
 - * The dropping of packets with a destination address belonging to a delegated prefix mandated in WPD-5 MUST NOT be applied to destinations that are part of any prefix announced using an ASSIGNED-PREFIX TLV by any HNCP router in the network.
 - * The section "LAN-Side Configuration" applies to HNCP interfaces classified as internal.
 - * The requirement L-2 to assign a separate /64 to each LAN interface is replaced by the participation in the prefix assignment mechanism (Section 6.2) for each such interface.

- * The requirement L-12 to make DHCPv6 options available is adapted, in that a CER SHOULD publish the subset of options using the DHCPv6 Data TLV in an External Connection TLV. Similarly it SHOULD do the same for DHCPv4 options in a DHCPv4 Data TLV. DHCPv6 options received inside an OPTION_IAPREFIX [RFC3633] MUST be published using a DHCPv6 Data TLV inside the respective Delegated Prefix TLV. HNCP routers SHOULD make relevant DHCPv6 and DHCPv4 options available to clients, i.e. options contained in External Connection TLVs that also include delegated prefixes from which a subset is assigned to the respective link.

12. Security Considerations

HNCP enables self-configuring networks, requiring as little user intervention as possible. However this zero-configuration goal usually conflicts with security goals and introduces a number of threats.

General security issues for existing home networks are discussed in [RFC7368]. The protocols used to set up addresses and routes in such networks to this day rarely have security enabled within the configuration protocol itself. However these issues are out of scope for the security of HNCP itself.

HNCP is a DNCP-based state synchronization mechanism carrying information with varying threat potential. For this consideration the payloads defined in DNCP and this document are reviewed:

- o Network topology information such as HNCP nodes and their common links.
- o Address assignment information such as delegated and assigned prefixes for individual links.
- o Naming and service discovery information such as auto-generated or customized names for individual links and routers.

12.1. Border Determination

As described in Section 5, an HNCP router determines the internal or external state on a per-link basis. A firewall perimeter is set up for the external links, and for internal links, HNCP and IGP traffic is allowed.

Threats concerning automatic border discovery cannot be mitigated by encrypting or authenticating HNCP traffic itself since external routers do not participate in the protocol and often cannot be

authenticated by other means. These threats include propagation of forged uplinks in the homenet in order to e.g. redirect traffic destined to external locations and forged internal status by external routers to e.g. circumvent the perimeter firewall.

It is therefore imperative to either secure individual links on the physical or link-layer or preconfigure the adjacent interfaces of HNCP routers to an adequate fixed-category in order to secure the homenet border. Depending on the security of the external link eavesdropping, man-in-the-middle and similar attacks on external traffic can still happen between a homenet border router and the ISP, however these cannot be mitigated from inside the homenet. For example, DHCPv4 has defined [RFC3118] to authenticate DHCPv4 messages, but this is very rarely implemented in large or small networks. Further, while PPP can provide secure authentication of both sides of a point to point link, it is most often deployed with one-way authentication of the subscriber to the ISP, not the ISP to the subscriber.

12.2. Security of Unicast Traffic

Once the homenet border has been established there are several ways to secure HNCP against internal threats like manipulation or eavesdropping by compromised devices on a link which is enabled for HNCP traffic. If left unsecured, attackers may perform arbitrary eavesdropping, spoofing or denial of service attacks on HNCP services such as address assignment or service discovery.

Detailed interface categories like "leaf" or "guest" can be used to integrate not fully trusted devices to various degrees into the homenet by not exposing them to HNCP and IGP traffic or by using firewall rules to prevent them from reaching homenet-internal resources.

On links where this is not practical and lower layers do not provide adequate protection from attackers, DNCP secure mode MUST be used to secure traffic.

12.3. Other Protocols in the Home

IGPs and other protocols are usually run alongside HNCP therefore the individual security aspects of the respective protocols must be considered. It can however be summarized that many protocols to be run in the home (like IGPs) provide - to a certain extent - similar security mechanisms. Most of these protocols do not support encryption and only support authentication based on pre-shared keys natively. This influences the effectiveness of any encryption-based

security mechanism deployed by HNCP as homenet routing information is thus usually not encrypted.

13. IANA Considerations

IANA is requested to maintain a registry for HNCP TLV-Types. This registry inherits TLV-Types and allocation policy defined in DNCP [I-D.ietf-homenet-dncp], but is independent with regard to all TLV-Types not specified or reserved by DNCP. Particularly, other DNCP profile may have there own registries, using same TLV numbers.

The following TLV-Types are defined in this document:

- 32: HNCP-Version
- 33: External-Connection
- 34: Delegated-Prefix
- 35: Assigned-Prefix
- 36: Node-Address
- 37: DHCPv4-Data
- 38: DHCPv6-Data
- 39: DNS-Delegated-Zone
- 40: Domain-Name
- 41: Node-Name
- 42: Managed-PSK

HNCP requires allocation of UDP port numbers HNCP-UDP-PORT and HNCP-DTLS-PORT, as well as an IPv6 link-local multicast address All-Homenet-Routers.

14. References

14.1. Normative references

- [I-D.ietf-homenet-dncp]
Stenberg, M. and S. Barth, "Distributed Node Consensus Protocol", draft-ietf-homenet-dncp-07 (work in progress), July 2015.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.
- [RFC6603] Korhonen, J., Savolainen, T., Krishnan, S., and O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", RFC 6603, May 2012.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.
- [I-D.ietf-homenet-prefix-assignment] Pfister, P., Paterson, B., and J. Arkko, "Distributed Prefix Assignment Algorithm", draft-ietf-homenet-prefix-assignment-07 (work in progress), June 2015.

14.2. Informative references

- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, November 2013.
- [RFC3004] Stump, G., Droms, R., Gu, Y., Vyaghrapuri, R., Demirtjis, A., Beser, B., and J. Privat, "The User Class Option for DHCP", RFC 3004, November 2000.
- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

- [RFC7368] Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, October 2014.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC6234] Eastlake, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, May 2011.
- [RFC1321] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, April 1992.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, February 2013.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, February 2013.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.
- [RFC3442] Lemon, T., Cheshire, S., and B. Volz, "The Classless Static Route Option for Dynamic Host Configuration Protocol (DHCP) version 4", RFC 3442, December 2002.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

14.3. URIs

[3] <http://www.openwrt.org>

[4] <http://www.homewrt.org/doku.php?id=run-conf>

Appendix A. Changelog [RFC Editor: please remove]

draft-ietf-homenet-hncp-07: Using version 1 instead of version 0, as existing implementations already use it.

draft-ietf-homenet-hncp-06: Various edits based on feedback, hopefully without functional delta.

draft-ietf-homenet-hncp-05: Renamed "Adjacent Link" to "Common Link". Changed single IPv4 uplink election from MUST to MAY. Added explicit indication to distinguish (IPv4)-PDs for local connectivity and ones with uplink connectivity allowing e.g. better local-only IPv4-connectivity.

draft-ietf-homenet-hncp-04: Change the responsibility for sending RAs to the router assigning the prefix.

draft-ietf-homenet-hncp-03: Split to DNCP (generic protocol) and HNCP (homenet profile).

draft-ietf-homenet-hncp-02: Removed any built-in security. Relying on IPsec. Reorganized interface categories, added requirements languages, made manual border configuration a MUST-support. Redesigned routing protocol election to consider non-router devices.

draft-ietf-homenet-hncp-01: Added (MAY) guest, ad-hoc, hybrid categories for interfaces. Removed old hnetv2 reference, and now pointing just to OpenWrt + github. Fixed synchronization algorithm to spread also same update number, but different data hash case. Made purge step require bidirectional connectivity between nodes when traversing the graph. Edited few other things to be hopefully slightly clearer without changing their meaning.

draft-ietf-homenet-hncp-00: Added version TLV to allow for TLV content changes pre-RFC without changing IDs. Added link id to assigned address TLV.

Appendix B. Draft source [RFC Editor: please remove]

This draft is available at <https://github.com/fingon/ietf-drafts/> in source format. Issues and pull requests are welcome.

Appendix C. Implementation [RFC Editor: please remove]

A GPLv2-licensed implementation of HNCP is currently under development at <https://github.com/sbyx/hnetd/> and binaries are available in the OpenWrt [3] package repositories. See [4] for more information. Feedback and contributions are welcome.

Appendix D. Acknowledgements

Thanks to Ole Troan, Mark Baugher, Mark Townsley and Juliusz Chroboczek for their contributions to the draft.

Thanks to Eric Kline for the original border discovery work.

Authors' Addresses

Markus Stenberg
Helsinki 00930
Finland

Email: markus.stenberg@iki.fi

Steven Barth
Halle 06114
Germany

Email: cyrus@openwrt.org

Pierre Pfister
Cisco Systems
Paris
France

Email: pierre.pfister@darou.fr