Internet Engineering Task Force (IETF)
Request for Comments: 7788
Category: Standards Track
ISSN: 2070-1721

M. Stenberg
S. Barth
Independent
P. Pfister
Cisco Systems

April 2016

Home Networking Control Protocol

Abstract

This document describes the Home Networking Control Protocol (HNCP), an extensible configuration protocol, and a set of requirements for home network devices. HNCP is described as a profile of and extension to the Distributed Node Consensus Protocol (DNCP). HNCP enables discovery of network borders, automated configuration of addresses, name resolution, service discovery, and the use of any routing protocol that supports routing based on both the source and destination address.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7788.

Table of Contents

1. Introduction

The Home Networking Control Protocol (HNCP) is designed to facilitate the sharing of state among home routers to fulfill the needs of the IPv6 homenet architecture [RFC7368], which assumes zero-configuration operation, multiple subnets, multiple home routers, and (potentially) multiple upstream service providers providing (potentially) multiple prefixes to the home network. While RFC 7368 sets no requirements for IPv4 support, HNCP aims to support the dual-stack mode of operation, and therefore the functionality is designed with that in mind. The state is shared as TLVs transported in the DNCP node state among the routers (and potentially advanced hosts) to enable:

- Autonomic discovery of network borders (Section 5.3) based on Distributed Node Consensus Protocol (DNCP) topology.

- Automated portioning of prefixes delegated by the service providers as well as assigned prefixes to both HNCP and non-HNCP routers (Section 6.3) using [RFC7695]. Prefixes assigned to HNCP routers are used to:

  * Provide addresses to non-HNCP aware nodes (using Stateless Address Autoconfiguration (SLAAC) and DHCP).

    * Provide space in which HNCP nodes assign their own addresses (Section 6.4).

- Internal and external name resolution, as well as multi-link service discovery (Section 8).

- Other services not defined in this document that do need to share state among homenet nodes and do not cause rapid and constant TLV changes (see the following applicability section).

HNCP is a protocol based on DNCP [RFC7787] and includes a DNCP profile that defines transport and synchronization details for sharing state across nodes defined in Section 3. The rest of the document defines behavior of the services noted above, how the required TLVs are encoded (Section 10), as well as additional requirements on how HNCP nodes should behave (Section 11).

1.1. Applicability

While HNCP does not deal with routing protocols directly (except potentially informing them about internal and external interfaces if classification specified in Section 5.3 is used), in homenet environments where multiple IPv6 source prefixes can be present, routing based on the source and destination address is necessary [RFC7368]. Ideally, the routing protocol is also zero configuration (e.g., no need to configure identifiers or metrics), although HNCP can also be used with a manually configured routing protocol.

As HNCP uses DNCP as the actual state synchronization protocol, the applicability statement of DNCP applies here as well; HNCP should not be used for any data that changes rapidly and constantly. If such data needs to be published in an HNCP network, 1) a more applicable protocol should be used for those portions, and 2) locators to a server of said protocol should be announced using HNCP instead. An example for this is naming and service discovery (Section 8) for which HNCP only transports DNS server addresses and no actual per-name or per-service data of hosts.

HNCP TLVs specified within this document, in steady state, stay constant, with one exception: as Delegated-Prefix TLVs (Section 10.2.1) do contain lifetimes, they force republishing of that data every time the valid or preferred lifetimes of prefixes are updated (significantly). Therefore, it is desirable for ISPs to provide large enough valid and preferred lifetimes to avoid unnecessary HNCP state churn in homes, but even given non-cooperating ISPs, the state churn is proportional only to the number of externally received delegated prefixes and not to the home network size, and it should therefore be relatively low.

HNCP assumes a certain level of control over host configuration servers (e.g., DHCP [RFC2131]) on links that are managed by its routers. Some HNCP functionality (such as border discovery or some aspects of naming) might be affected by existing DHCP servers that are not aware of the HNCP-managed network and thus might need to be reconfigured to not result in unexpected behavior.

While HNCP routers can provide configuration to and receive configuration from non-HNCP routers, they are not able to traverse such devices based solely on the protocol as defined in this document, i.e., HNCP routers that are connected only by different interfaces of a non-HNCP router will not be part of the same HNCP network.

While HNCP is designed to be used by (home) routers, it can also be used by advanced hosts that want to do, e.g., their own address assignment and routing.

HNCP is link-layer agnostic; if a link supports IPv6 (link-local) multicast and unicast, HNCP will work on it. Trickle retransmissions and keep-alives will handle both packet loss and non-transitive connectivity, ensuring eventual convergence.

2. Terminology

The following terms are used as they are defined in [RFC7695]:

- Advertised Prefix Priority
- Advertised Prefix
- Assigned Prefix
- Delegated Prefix
- Prefix Adoption
- Private Link
- Published Assigned Prefix
- Applied Assigned Prefix
- Shared Link

The following terms are used as they are defined in [RFC7787]:

- DNCP profile
- Node identifier
- Link
- Interface

(HNCP) node	a device implementing this specification.
(HNCP) router	a device implementing this specification, which forwards traffic on behalf of other devices.
Greatest node identifier	when comparing the DNCP node identifiers of multiple nodes, the one that has the greatest value in a bitwise comparison.
Border	separation point between administrative domains; in this case, between the home network and any other network, i.e., usually an ISP network.
Internal link	a link that does not cross borders.
Internal interface	an interface that is connected to an internal link.
External interface	an interface that is connected to a link that is not an internal link.
Interface category	a local configuration denoting the use of a particular interface. The Interface category determines how an HNCP node should treat the particular interface. The External and Internal categories mark the interface as out of or within the network border; there are also a number of subcategories of Internal that further affect local node behavior. See Section 5.1 for a list of interface categories and how they behave. The Internal or External category may also be auto-detected (Section 5.3).
Border router	a router announcing external connectivity and forwarding traffic across the network border.

Common Link	a set of nodes on a link that share a common view of it, i.e., they see each other's traffic and the same set of hosts. Unless configured otherwise, transitive connectivity is assumed.
DHCPv4	refers to the Dynamic Host Configuration Protocol [RFC2131] in this document.
DHCPv6	refers to the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [RFC3315] in this document.
DHCP	refers to cases that apply to both DHCPv4 and DHCPv6 in this document.

2.1. 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

The DNCP profile for HNCP is defined as follows:

- HNCP uses UDP datagrams on port 8231 as a transport over link-local scoped IPv6, using unicast and multicast (FF02:0:0:0:0:0:0:11 is the HNCP group address). Received datagrams where either or both of the IPv6 source or destination addresses are not link-local scoped MUST be ignored. 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.

- HNCP operates on multicast-capable interfaces only. HNCP nodes MUST assign a non-zero 32-bit endpoint identifier to each interface for which HNCP is enabled. The value 0 is not used in DNCP TLVs but has a special meaning in HNCP TLVs (see Sections 6.4 and 10.3). These identifiers MUST be locally unique within the scope of the node, and using values equivalent to the IPv6 link-local scope identifiers for the given interfaces are RECOMMENDED.

- HNCP uses opaque 32-bit node identifiers (DNCP_NODE_IDENTIFIER_LENGTH = 32). A node implementing HNCP SHOULD use a random node identifier. If there is a node identifier collision (as specified in the Node-State TLV handling of Section 4.4 of [RFC7787]), the node MUST immediately generate

and use a new random node identifier that is not used by any other node at the time, based on the current DNCP network state.

- HNCP nodes MUST use the leading 64 bits of the MD5 message digest [RFC1321] as the DNCP hash function H(x) used in building the DNCP hash tree.

- HNCP nodes MUST use DNCP's per-endpoint keep-alive extension on all endpoints. The following parameters are suggested:

  * Default keep-alive interval (DNCP_KEEPALIVE_INTERVAL): 20 seconds.

  * Multiplier (DNCP_KEEPALIVE_MULTIPLIER): 2.1 on virtually lossless links works fine, as it allows for one lost keep-alive. If used on a lossy link, a considerably higher multiplier, such as 15, should be used instead. In that case, an implementation might prefer shorter keep-alive intervals on that link as well to ensure that the timeout (equal to DNCP_KEEPALIVE_INTERVAL * DNCP_KEEPALIVE_MULTIPLIER) after which entirely lost nodes time out is low enough.

- HNCP nodes use the following Trickle parameters for the per-interface Trickle instances:

  * k SHOULD be 1, as the timer reset when data is updated, and further retransmissions should handle packet loss. Even on a non-transitive lossy link, the eventual per-endpoint keep-alives should ensure status synchronization occurs.

  * 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 [RFC6206] but MUST NOT be lower.

- HNCP unicast traffic SHOULD be secured using Datagram Transport Layer Security (DTLS) [RFC6347] as described in DNCP if exchanged over unsecured links. UDP on port 8232 is used for this purpose. A node implementing HNCP security MUST support the DNCP Pre-Shared Key (PSK) method, SHOULD support the PKI-based trust method, and MAY support the DNCP certificate-based trust consensus method. [RFC7525] provides guidance on how to securely utilize DTLS.

- HNCP nodes MUST ignore all Node-State TLVs received via multicast on a link that has DNCP security enabled in order to prevent spoofing of node state changes.

4. HNCP Versioning and Router Capabilities

Multiple versions of HNCP based on compatible DNCP profiles may be present in the same network when transitioning between HNCP versions, and for troubleshooting purposes, it might be beneficial to identify the HNCP agent version running. Therefore, each node MUST include an HNCP-Version TLV (Section 10.1) indicating the currently supported version in its node data and MUST ignore (except for DNCP synchronization purposes) any TLVs that have a type greater than 32 and that are published by nodes that didn't also publish an HNCP-Version TLV.

HNCP routers may also have different capabilities regarding interactions with hosts, e.g., for configuration or service discovery. These are indicated by M, P, H, and L values. The combined "capability value" is a metric indicated by interpreting the bits as an integer, i.e., (M << 12 | P << 8 | H << 4 | L). These values are used to elect certain servers on a Common Link, as described in Section 7. Nodes that are not routers MUST announce the value 0 for all capabilities. Any node announcing the value 0 for a capability is considered to not advertise said capability and thus does not take part in the respective election.

5. Interface Classification

5.1. Interface Categories

HNCP specifies the following categories that interfaces can be configured to be in:

Internal category: This declares an interface to be internal, i.e., within the borders of the HNCP network. The interface MUST operate as a DNCP endpoint. Routers MUST forward traffic with appropriate source addresses between their internal interfaces and allow internal traffic to reach external networks. All nodes MUST implement this category, and nodes not implementing any other category implicitly use it as a fixed default.

External category: This declares an interface to be external, i.e., not within the borders of the HNCP network. The interface MUST NOT operate as a DNCP endpoint. Accessing internal resources from external interfaces is restricted, i.e., the use of Recommended Simple Security Capabilities in Customer Premises Equipments (CPEs) [RFC6092] is RECOMMENDED. HNCP routers SHOULD announce acquired configuration information for use in the network as described in Section 6.2, if the interface appears to be connected to an external network. HNCP routers MUST implement this category.

Leaf category: This declares an interface used by client devices only. Such an interface uses the Internal category with the exception that it MUST NOT operate as a DNCP endpoint. This category SHOULD be supported by HNCP routers.

Guest category: This declares an interface used by untrusted client devices only. In addition to the restrictions of the Leaf category, HNCP routers MUST filter traffic from and to the interface 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 by HNCP routers.

Ad Hoc category: This configures an interface to use the Internal category, but no assumption is made about the link's transitivity. All other interface categories assume transitive connectivity. This affects the Common Link (Section 6.1) definition. Support for this category is OPTIONAL.

Hybrid category: This declares an interface to use the Internal category while still trying to acquire (external) configuration information on it, e.g., by running DHCP clients. This is useful, e.g., if the link is shared with a non-HNCP router under control and still within the borders of 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.

5.2. DHCP-Aided Auto-Detection

Auto-detection of interface categories is possible based on interaction with DHCPv4 [RFC2131] and DHCPv6 Prefix Delegation (DHCPv6-PD) [RFC3633] servers on connected links. HNCP defines special DHCP behavior to differentiate its internal servers from external ones in order to achieve this. Therefore, all internal devices (including HNCP nodes) running DHCP servers on links where auto-detection is used by any HNCP node MUST use the following mechanism based on "The User Class Option for DHCP" [RFC3004] and its DHCPv6 counterpart [RFC3315]:

- The device MUST ignore or reject DHCP-Requests containing a DHCP user class consisting of the ASCII string "HOMENET".

Not following this rule (e.g., running unmodified DHCP servers) might lead to false positives when auto-detection is used, i.e., HNCP nodes assume an interface to not be internal, even though it was intended to be.

5.3. Algorithm for Border Discovery

This section defines the interface classification algorithm. It is suitable for both IPv4 and IPv6 (single or dual stack) and detects the category of an interface either automatically or based on a fixed configuration. By determining the category for all interfaces, the network borders are implicitly defined, i.e., all interfaces not belonging to the External category are considered to be within the borders of the network; all others are not.

The following algorithm MUST be implemented by any node implementing HNCP. However, if the node does not implement auto-detection, only the first and last step are required. The algorithm works as follows, with evaluation stopping at first match:

1. If a fixed category is configured for an interface, it is used.

2. If a delegated prefix could be acquired by running a DHCPv6 client, it is considered external. The DHCPv6 client MUST have included a DHCPv6 user class consisting of the ASCII string "HOMENET" in all of its requests.

3. If an IPv4 address could be acquired by running a DHCPv4 client on the interface, it is considered external. The DHCPv4 client MUST have included a DHCP user class consisting of the ASCII string "HOMENET" in all of its requests.

4. The interface is considered internal.

Note that as other HNCP nodes will ignore the client due to the User Class option, any server that replies is clearly external (or a malicious internal node).

An HNCP router SHOULD allow setting the fixed category for each interface that may be connected to either an internal or external device (e.g., an Ethernet port that can be connected to a modem, another HNCP router, or a client). Note that all fixed categories except internal and external cannot be auto-detected and can only be selected using manual configuration.

An HNCP router using auto-detection on an interface MUST run the appropriately configured DHCP clients as long as the interface without a fixed category is active (including states where auto-detection considers it to be internal) and rerun the algorithm above to react to conditions resulting in a different interface category. 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.

6. Autonomous Address Configuration

This section specifies how HNCP nodes configure host and node addresses. At first, border routers share information obtained from service providers or local configuration by publishing one or more External-Connection TLVs (Section 10.2). These contain other TLVs such as Delegated-Prefix TLVs (Section 10.2.1) that are then used for prefix assignment. Finally, HNCP nodes obtain addresses either statelessly or using a specific stateful mechanism (Section 6.4). Hosts and non-HNCP routers are configured using SLAAC, DHCP, or DHCPv6-PD.

6.1. Common Link

HNCP uses the concept of Common Link both in autonomic address configuration and naming and service discovery (Section 8). A Common Link refers to the set of interfaces of nodes that see each other's traffic and presumably also the traffic of all hosts that may use the nodes to, e.g., forward traffic. Common Links are used, e.g., to determine where prefixes should be assigned or which peers participate in the election of a DHCP server. The Common Link is computed separately for each local internal interface, and it always contains the local interface. Additionally, if the local interface is not set to the Ad Hoc category (see Section 5.1), it also contains the set of interfaces that are bidirectionally reachable from the given local interface; that is, every remote interface of a remote node meeting all of the following requirements:

- The local node publishes a Peer TLV with:
  - Peer Node Identifier = remote node's node identifier
  - Peer Endpoint Identifier = remote interface's endpoint identifier
  - Endpoint Identifier = local interface's endpoint identifier
- The remote node publishes a Peer TLV with:
  - Peer Node Identifier = local node's node identifier
  - Peer 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 internal interfaces are connected, e.g., by detecting an identical remote interface being part of the Common Links of both local interfaces.

6.2. External Connections

Each HNCP router MAY obtain external connection information such as address prefixes, DNS server addresses, and DNS search paths from one or more sources, e.g., DHCPv6-PD [RFC3633], NETCONF [RFC6241], or static configuration. Each individual external connection to be shared in the network is represented by one External-Connection TLV (Section 10.2).

Announcements of individual external connections can consist of the following components:

Delegated Prefixes: Address space available for assignment to internal links announced using Delegated-Prefix TLVs (Section 10.2.1). Some address spaces might have special properties that are necessary to understand in order to handle them (e.g., information similar to [RFC6603]). This information is encoded using DHCPv6-Data TLVs (Section 10.2.2) inside the respective Delegated-Prefix TLVs.

Auxiliary Information: Information about services such as DNS or time synchronization regularly used by hosts in addition to addressing and routing information. This information is encoded using DHCPv6-Data TLVs (Section 10.2.2) and DHCPv4-Data TLVs (Section 10.2.3).

Whenever information about reserved parts (e.g., as specified in [RFC6603]) is received for a delegated prefix, the reserved parts MUST be advertised using Assigned-Prefix TLVs (Section 10.3) with the greatest priority (i.e., 15), as if they were assigned to a Private Link.

Some connections or delegated 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, Voice over IP (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 or to only selectively allow certain devices to use them. Such prefixes are distinguished using Prefix-Policy TLVs (Section 10.2.1.1). 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-Policy TLVs with type 131, as described in Section 6.3.2.

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

6.3. Prefix Assignment

HNCP uses the prefix assignment algorithm [RFC7695] in order to assign prefixes to HNCP internal links and uses some of the terminology (Section 2) defined there. HNCP furthermore defines the Assigned-Prefix TLV (Section 10.3), which MUST be used to announce Published Assigned Prefixes.

6.3.1. Prefix Assignment Algorithm Parameters

All HNCP nodes running the prefix assignment algorithm use the following values for its parameters:

Node IDs: HNCP node identifiers are used. The comparison operation is defined as bitwise comparison.

Set of Delegated Prefixes: The set of prefixes encoded in Delegated-Prefix TLVs that 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 interfaces with the Internal, Leaf, Guest, or Ad Hoc category. It is dynamically updated as interfaces are added, removed, or switched 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 as 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 nodes.

Set of Advertised Prefixes: The set of prefixes included in Assigned-Prefix TLVs advertised by other HNCP nodes (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 0, 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 is 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.

6.3.2. Making New Assignments

Whenever the prefix assignment algorithm subroutine (Section 4.1 of [RFC7695]) is run on a Common Link, and whenever a new prefix may be assigned (case 1 of the subroutine: no Best Assignment and no Current Assignment), the decision of whether the assignment of a new prefix is desired MUST follow these rules in order:

> If the Delegated-Prefix TLV contains a DHCPv6-Data TLV, and the meaning of one of the DHCP options is not understood by the HNCP node, 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-zero preferred lifetime, the creation of a new prefix is not desired.
>
> If the Delegated-Prefix TLV does not include a Prefix-Policy TLV indicating restrictive assignment (type 131) or if local policy exists to identify it based on, e.g., other Prefix-Policy TLV values and allows assignment, the creation of a new prefix is desired.
>
> Otherwise, the creation of a new prefix is not desired.

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 6.5 details how IPv4-delegated prefixes are generated), a prefix of length 24 SHOULD be preferred.

In any case, an HNCP router making an assignment 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 announce the L-capability, and a router assigning an IPv6 prefix with a length greater than 64 MUST announce the H-capability as defined in Section 4.

6.3.3. 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 forward traffic destined to said prefix to the respective link.
- 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 6.4, e.g., if it is to be used as source for locally originating traffic.

6.3.4. DHCPv6 Prefix Delegation

When an HNCP router announcing the P-Capability (Section 4) receives a DHCPv6-PD request from a client, it SHOULD assign one prefix per delegated prefix in the network. This set of assigned prefixes is then delegated to the client, after it has been applied as described in the prefix assignment algorithm. Each DHCPv6-PD 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; however, the prefix length to be delegated MAY be smaller than 64.

The assigned prefixes MUST NOT be given to DHCPv6-PD 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 that 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.

6.4. 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 node SHOULD announce an IPv6 address and -- if it supports IPv4 -- MUST announce an 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 Semantically Opaque Interface Identifiers [RFC7217] SHOULD be used for address assignment whenever possible (e.g., the prefix length is 64), otherwise (e.g., for IPv4 if supported) the following method MUST be used instead: For any assigned prefix for which stateless assignment is not used, the first quarter of the addresses are reserved for HNCP-based address assignments, whereas the last three quarters are left to the DHCP elected router (Section 4 specifies the DHCP server election process). For example, 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 nodes assign addresses to themselves and then (to ensure eventual lack of conflicting assignments) publish the assignments using the Node-Address TLV (Section 10.4).

The process of obtaining addresses is specified as follows:

- A node MUST NOT start advertising an address if it is already advertised by another node.

- An assigned address MUST be part of an assigned prefix currently applied on a Common Link that includes the interface specified by the endpoint identifier.

- 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.

- Whenever the same address is advertised by more than one node, all but the one advertised by the node with the greatest node identifier MUST be removed.

6.5. Local IPv4 and ULA Prefixes

HNCP routers can create a Unique Local Address (ULA) or private IPv4 prefix to enable connectivity between local devices. These prefixes are inserted in HNCP as if they were delegated prefixes of a (virtual) external connection (Section 6.2). 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-Policy 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 greatest 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 greatest node identifier is kept among those with a Prefix-Policy TLV of type 0 -- if there is any. The router publishing a prefix with Internet connectivity MUST forward IPv4 traffic to the Internet 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 time span between 0 and 10 seconds in which the router MUST scan for others 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].

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 4 in order to participate as a candidate in the election.

7.1. IPv6 Addressing and Configuration

In general, Stateless Address Autoconfiguration [RFC4861] is used for client configuration for its low overhead and fast renumbering capabilities. Therefore, each HNCP router sends Router Advertisements on interfaces that are intended to be used by clients and MUST at least include a Prefix Information Option for each Applied Assigned Prefix that it assigned to the respective link in every such advertisement. 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 6.1) is elected based on the capabilities described in Section 4. The winner is the router (connected to the Common Link) advertising the greatest H-capability. In case of a tie, Capability Values (Section 4) are compared, and the router with the greatest value is elected. In case of another tie, the router with the greatest node identifier is elected among the routers with tied Capability Values.

The elected router MUST serve stateful DHCPv6 and SHOULD provide naming services for acquired hostnames as outlined in Section 8; all other nodes MUST NOT. Stateful addresses SHOULD be assigned in a way that does not hinder fast renumbering even if the DHCPv6 server or client do not support the DHCPv6 reconfigure mechanism, e.g., by only handing out leases from locally generated (ULA) prefixes and prefixes with a length different from 64 and by using low renew and rebind times (i.e., not longer than 5 minutes). In case no router was elected, stateful DHCPv6 is not provided. Routers that cease to be elected DHCP servers SHOULD -- when applicable -- invalidate remaining existing bindings in order to trigger client reconfiguration.

7.2. DHCPv6 for Prefix Delegation

The designated DHCPv6 server for prefix delegation on a Common Link is elected based on the capabilities described in Section 4. The winner is the router (connected to the Common Link) advertising the greatest P-capability. In case of a tie, Capability Values (Section 4) are compared, and the router with the greatest value is elected. In case of another tie, the router with the greatest 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 6.3.4); all other nodes MUST NOT.

7.3. DHCPv4 for Addressing and Configuration

The designated DHCPv4 server on a Common Link (Section 6.1) is elected based on the capabilities described in Section 4. The winner is the router (connected to the Common Link) advertising the greatest L-capability. In case of a tie, Capability Values (Section 4) are compared, and the router with the greatest value is elected. In case of another tie, the router with the greatest node identifier is elected among the routers with tied Capability Values.

The elected router MUST provide DHCPv4 services on the given link; all other nodes MUST NOT. The elected router MUST provide IP addresses from the pool defined in Section 6.4 and MUST announce itself as router [RFC2132] to clients.

DHCPv4 lifetimes renew and rebind times (T1 and T2) SHOULD be short (i.e., not longer than 5 minutes) in order to provide reasonable response times to changes. Routers that cease to be elected DHCP servers SHOULD -- when applicable -- invalidate remaining existing bindings in order to trigger client reconfiguration.

7.4. Multicast DNS Proxy

The designated Multicast DNS (mDNS) [RFC6762] proxy on a Common Link is elected based on the capabilities described in Section 4. The winner is the router (connected to the Common Link) advertising the greatest M-capability. In case of a tie, Capability Values (Section 4) are compared, and the router with the greatest value is elected. In case of another tie, the router with the greatest 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 advertise a recursive name resolving server to clients that honor the announcements made in Delegated-Zone TLVs (Section 10.5), Domain-Name TLVs (Section 10.6), and Node-Name TLVs (Section 10.7), i.e., delegate queries to the designated name servers and hand out appropriate A, AAAA, and PTR records according to the mentioned TLVs.

Each HNCP router SHOULD provide and announce an auto-generated or user-configured name for each internal Common Link (Section 6.1) 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. This announcement is done using Delegated-Zone TLVs (Section 10.5) and MUST be unique in the whole network. In case of a conflict, the announcement of the node with the greatest node identifier takes precedence, and all other nodes MUST cease to announce the conflicting TLV. HNCP routers providing recursive name resolving services MUST use the included DNS server

address within the TLV to resolve names belonging to the zone as if there was an NS record.

Each HNCP node SHOULD announce a node name for itself to be easily reachable and MAY announce names on behalf of other devices. Announcements are made using Node-Name TLVs (Section 10.7), and the announced names MUST be unique in the whole network. In case of a conflict, the announcement of the node with the greatest node identifier takes precedence, and all other nodes MUST cease to announce the conflicting TLV. HNCP routers providing recursive name resolving services as described above MUST resolve such announced names to their respective IP addresses as if there were corresponding A/AAAA records.

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.

9. Securing Third-Party Protocols

PSKs are often required to secure (for example) IGPs and other protocols that lack support for asymmetric security. The following mechanism manages PSKs using HNCP to enable bootstrapping of such third-party protocols. The scheme SHOULD NOT be used unless it’s in conjunction with secured HNCP unicast transport (i.e., DTLS), as transferring the PSK in plaintext anywhere in the network is a potential risk, especially as the originator may not know about security (and use of DNCP security) on all links. The following rules define how such a PSK is managed and used:

- If no Managed-PSK TLV (Section 10.8) is currently being announced, an HNCP node using this mechanism 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.

- In case multiple nodes 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.

- The node currently advertising the Managed-PSK TLV MUST generate and advertise a new random one whenever an unreachable node is removed from the DNCP topology as described in Section 4.6 of [RFC7787].

PSKs for individual protocols SHOULD be derived from the random PSK using a suitable one-way hashing algorithm (e.g., by using the HMAC-based Key Derivation Function (HKDF) based on HMAC-SHA256 [RFC6234] with the particular protocol name in the info field) so that disclosure of any derived key does not impact other users of the managed PSK. Furthermore, derived PSKs MUST be updated whenever the managed PSK changes.

10. Type-Length-Value Objects

HNCP defines the following TLVs in addition to those defined by DNCP. The same general rules and defaults for encoding as noted in Section 7 of [RFC7787] apply. Note that most HNCP variable-length TLVs also support optional nested TLVs, and they are encoded after the variable-length content, followed by the zero padding of the variable-length content to the next 32-bit boundary.

TLVs defined here are only valid when appearing in their designated context, i.e., only directly within container TLVs mentioned in their definition or -- absent any mentions -- only as top-level TLVs within the node data set. TLVs appearing outside their designated context MUST be ignored.

TLVs encoding IP addresses or prefixes allow encoding both IPv6 and IPv4 addresses and prefixes. IPv6 information is encoded as is, whereas for IPv4, the IPv4-mapped IPv6 addresses format [RFC4291] is used, and prefix lengths are encoded as the original IPv4 prefix length increased by 96.

10.1. HNCP-Version TLV

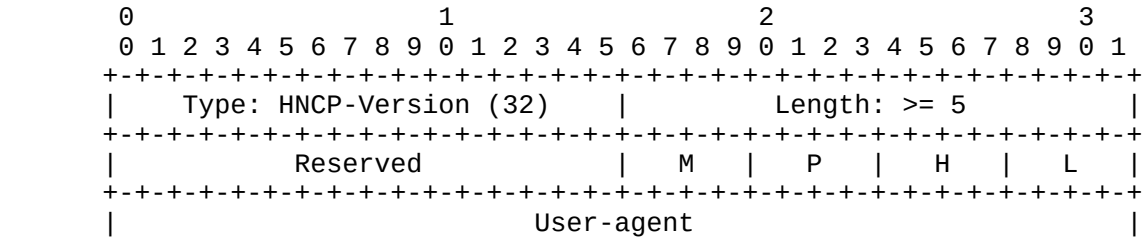

This TLV is used to indicate the supported version and router capabilities of an HNCP node as described in Section 4.

Reserved: Bits are reserved for future use. They MUST be set to 0 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 0 if the router is not capable of proxying mDNS, otherwise it SHOULD be set to 4 but MAY be set to any value from 1 to 7 to indicate a non-default priority. 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 0 if the router is not capable of providing prefixes through DHCPv6-PD (Section 6.3.4), otherwise it SHOULD be set to 4 but MAY be set to any value from 1 to 7 to indicate a non-default priority. 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 0 if the router is not capable of providing such addresses, otherwise it SHOULD be set to 4 but MAY be set to any value from 1 to 7 to indicate a non-default priority. 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 0 if the router is not capable of running a legacy DHCPv4 server offering IPv4 addresses to clients, otherwise it SHOULD be set to 4 but MAY be set to any value from 1 to 7 to indicate a non-default priority. 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.

10.2. External-Connection TLV

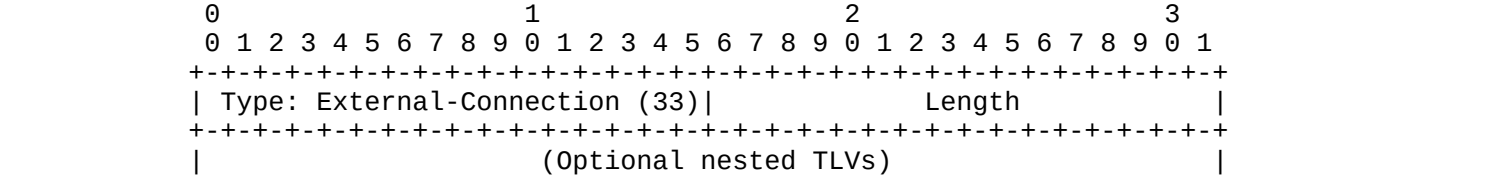

An External-Connection TLV is a container TLV used to gather network configuration information associated with a single external connection (Section 6.2) to be shared across the HNCP network. A node MAY publish an arbitrary number of instances of this TLV to share the desired number of external connections. Upon reception, the information transmitted in any nested TLVs is used for the purposes of prefix assignment (Section 6.3) and host configuration (Section 7).

10.2.1. Delegated-Prefix TLV

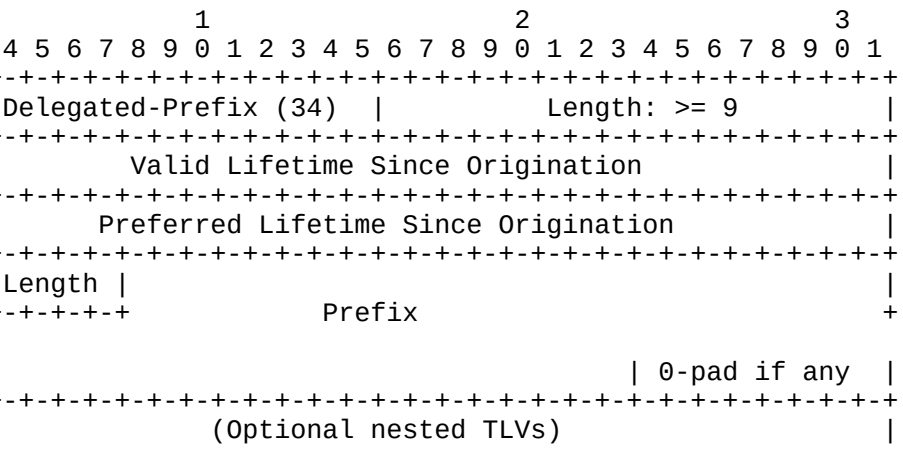

The Delegated-Prefix TLV is used by HNCP routers to advertise prefixes that are allocated to the whole network and can be used for prefix assignment. Delegated-Prefix TLVs are only valid inside External-Connection TLVs, and their prefixes MUST NOT overlap with those of other such TLVs in the same container.

Valid Lifetime Since Origination: The time in seconds the delegated prefix was valid for at the origination time of the node data containing this TLV. The value MUST be updated whenever the node republishes its Node-State TLV.

Preferred Lifetime Since Origination: The time in seconds the delegated prefix was preferred for at the origination time of the node data containing this TLV. The value MUST be updated whenever the node republishes its Node-State TLV.

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

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

10.2.1.1. Prefix-Policy TLV

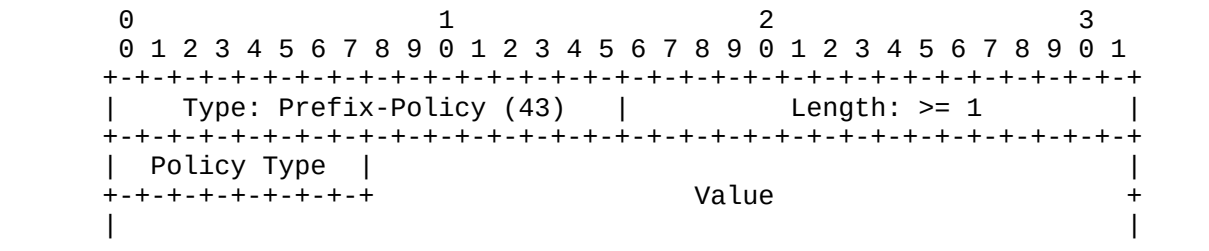

The Prefix-Policy TLV contains information about the policy or applicability of a delegated prefix. This information can be used to determine whether prefixes for a certain use case (e.g., local reachability, Internet connectivity) do exist or are to be acquired and to make decisions about assigning prefixes to certain links or to fine-tune border firewalls. See Section 6.2 for a more in-depth discussion. This TLV is only valid inside a Delegated-Prefix TLV.

Policy Type: The type of the policy identifier.

- 0: Internet connectivity (no value).
- 1-128: Explicit destination prefix with the Policy Type being the actual length of the prefix and the value containing significant bits of the destination prefix padded with zeros up to the next byte boundary.
- 129: DNS domain. The value contains a DNS label sequence encoded per [RFC1035]. Compression MUST NOT be used. The label sequence MUST end with an empty label.
- 130: Opaque UTF-8 string (e.g., for administrative purposes).
- 131: Restrictive assignment (no value).
- 132-255: Reserved for future additions.

Value: A variable-length identifier of the given type.

10.2.2. DHCPv6-Data TLV

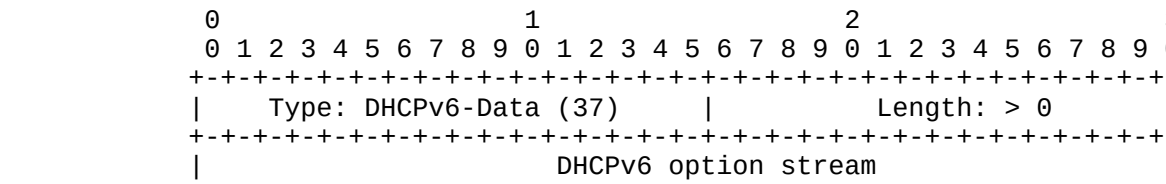

This TLV is used to encode auxiliary IPv6 configuration information (e.g., recursive DNS servers) encoded as a stream of DHCPv6 options. It is only valid in an External-Connection TLV or a Delegated-Prefix TLV encoding an IPv6 prefix and MUST NOT occur more than once in any single container. When included in an External-Connection TLV, it contains DHCPv6 options relevant to the external connection as a whole. When included in a delegated prefix, it contains options mandatory to handle said prefix.

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

10.2.3. DHCPv4-Data TLV

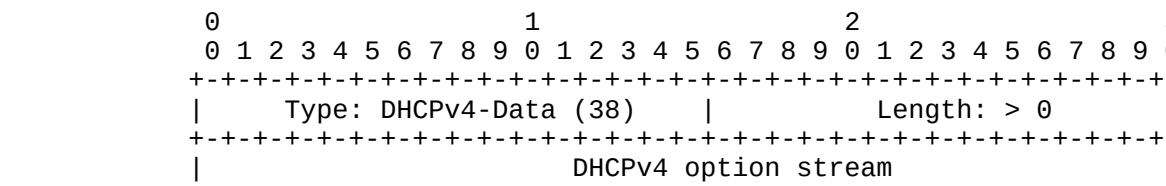

This TLV is used to encode auxiliary IPv4 configuration information (e.g., recursive DNS servers) encoded as a stream of DHCPv4 options. It is only valid in an External-Connection TLV and MUST NOT occur more than once in any single container. It contains DHCPv4 options relevant to the external connection as a whole.

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

10.3. Assigned-Prefix TLV

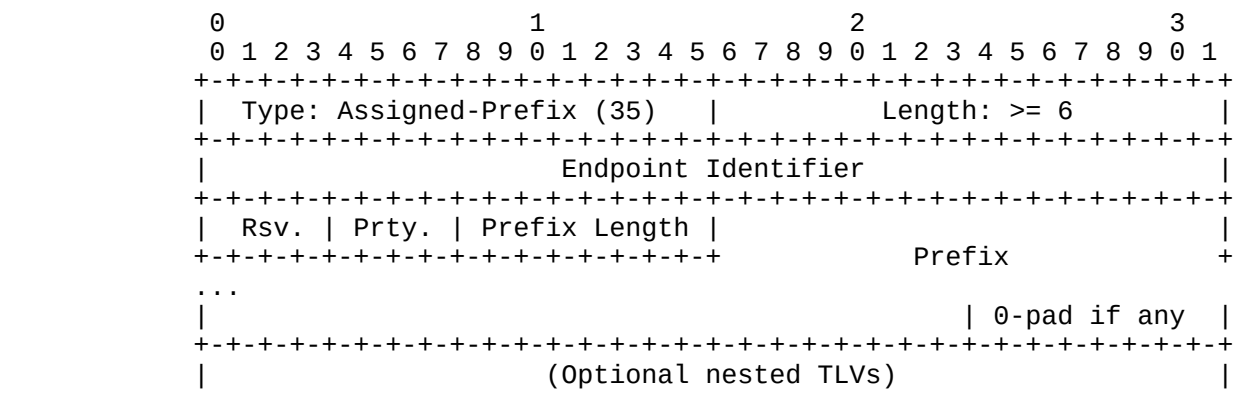

This TLV is used to announce Published Assigned Prefixes for the purposes of prefix assignment (Section 6.3).

Endpoint Identifier: The endpoint identifier of the local interface the prefix is assigned to, or 0 if it is assigned to 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 0 when creating this TLV, and their value MUST be ignored when processing the TLV.

Prty: 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 the 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 zeros up to the next byte boundary.

10.4. Node-Address TLV

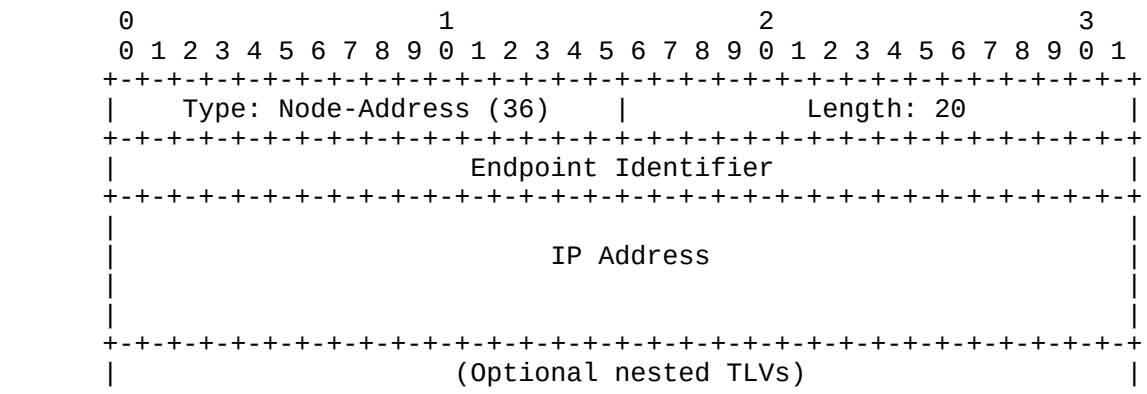

This TLV is used to announce addresses assigned to an HNCP node as described in Section 6.4.

Endpoint Identifier: The endpoint identifier of the local interface 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].

10.5. DNS-Delegated-Zone TLV

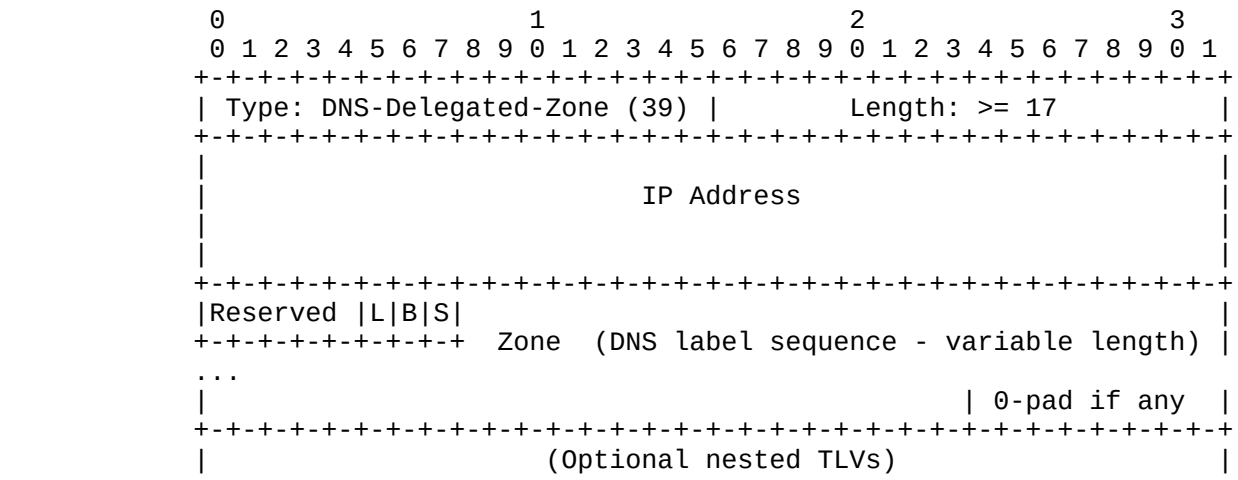

This TLV is used to announce a forward or reverse DNS zone delegation in the HNCP network. Its meaning is roughly equivalent to specifying an NS and A/AAAA record for said zone. Details are specified in Section 8.

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 zeros) means the delegation is available in the global DNS hierarchy.

Reserved: Those bits MUST be set to 0 when creating the TLV and ignored when parsing it unless defined in a later specification.

L-bit: (DNS-based Service Discovery (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 the 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 a 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 encoded according to [RFC1035]. Compression MUST NOT be used. The label sequence MUST end with an empty label.

10.6. Domain-Name TLV

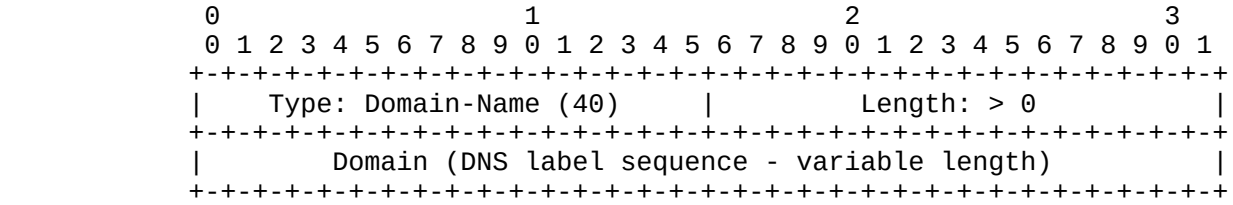

This TLV is used to indicate the base domain name for the network as specified in Section 8. This TLV MUST NOT be announced unless the domain name was explicitly configured by an administrator.

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

10.7. Node-Name TLV

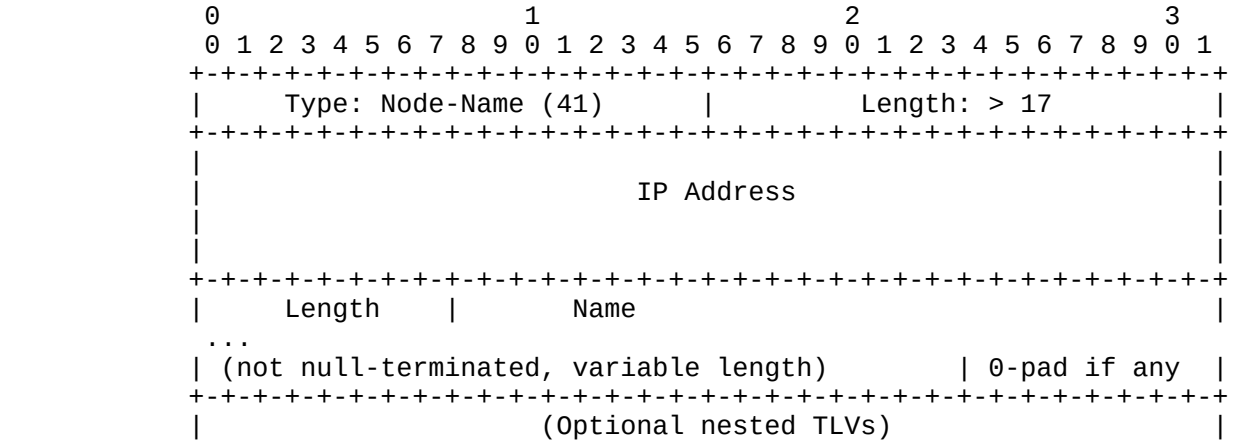

This TLV is used to assign the name of a node in the network to a certain IP address as specified in Section 8.

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

Length: The length of the name (0-63).

Name: The name of the node as a single DNS label.

10.8. Managed-PSK TLV

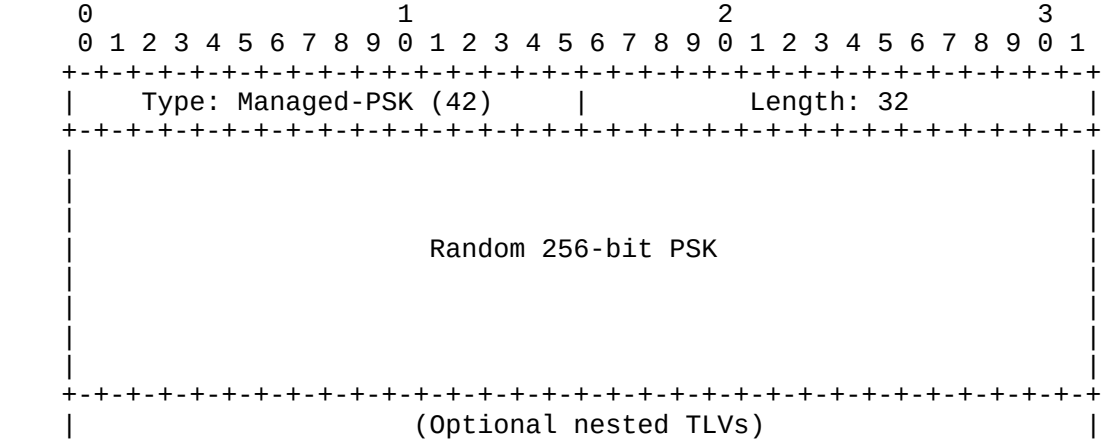

This TLV is used to announce a PSK for securing third-party protocols exclusively supporting symmetric cryptography as specified in Section 9.

11. General Requirements for HNCP Nodes

Each node implementing HNCP is subject to the following requirements:

- It MUST implement HNCP versioning (Section 4) and interface classification (Section 5).
- It MUST implement and run the method for securing third-party protocols (Section 9) whenever it uses the security mechanism of HNCP.

If the node is acting as a router, then the following requirements apply in addition:

- It MUST support Autonomous Address Configuration (Section 6) and configuration of hosts and non-HNCP routers (Section 7).
- It SHOULD implement support for naming and service discovery (Section 8) as defined in this document.

- It MAY be able to provide connectivity to IPv4 devices using DHCPv4.

- It SHOULD be able to delegate prefixes to legacy IPv6 routers using DHCPv6-PD (Section 6.3.4).

- In addition, the normative language of "Basic Requirements for IPv6 Customer Edge Routers" [RFC7084] applies with the following adjustments:

  * The generic requirements G-4 and G-5 are relaxed such that any known default router on any interface is sufficient for a router to announce itself as the default router; similarly, only the loss of all such default routers results in self-invalidation.

  * "WAN-Side Configuration" (Section 4.2) applies to interfaces classified as external.

  * If the Customer Edge (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.

  * "LAN-Side Configuration" (Section 4.3) applies to interfaces not classified as external.

  * The requirement L-2 to assign a separate /64 to each LAN interface is replaced by the participation in the prefix assignment mechanism (Section 6.3) for each such interface.

  * The requirement L-9 is modified, in that the M flag MUST be set if and only if a router connected to the respective Common Link is advertising a non-zero H-capability. The O flag SHOULD always be set.

  * The requirement L-12 to make DHCPv6 options available is adapted, in that Canonical Encoding Rules (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.

* The requirement L-13 to deprecate prefixes is applied to all delegated prefixes in the network from which assignments have been made on the respective interface. Furthermore, the Prefix Information Options indicating deprecation MUST be included in Router Advertisements for the remainder of the prefixes' respective valid lifetime but MAY be omitted after at least 2 hours have passed.

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:

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

12.1. Interface Classification

As described in Section 5.3, an HNCP node determines the internal or external state on a per-interface basis. A firewall perimeter is set up for the external interfaces, and for internal interfaces, HNCP traffic is allowed, with the exception of the Leaf and Guest subcategories.

Threats concerning automatic interface classification 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 appropriate 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 that is enabled for HNCP traffic. If left unsecured, attackers may perform arbitrary traffic redirection, eavesdropping, spoofing, or denial-of-service attacks on HNCP services such as address assignment or service discovery, and the protocols are secured using HNCP-derived keys such as routing protocols.

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 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, DTLS-based secure unicast transport 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 has set up a registry for the (decimal values within range 32-511) "HNCP TLV Types" under "Distributed Node Consensus Protocol (DNCP)". The registration procedures is ’RFC Required’ [RFC5226]. The initial contents are:

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

43: Prefix-Policy

44-511: Unassigned.

768-1023: Reserved for Private Use. This range is used by HNCP for per-implementation experimentation. How collisions are avoided is outside the scope of this document.

IANA has registered the UDP port numbers 8231 (service name: hncp-udp-port, description: HNCP) and 8232 (service name: hncp-dtls-port, description: HNCP over DTLS), as well as an IPv6 link-local multicast address FF02:0:0:0:0:0:0:11 (description: All-Homenet-Nodes).

14. References

14.1. Normative References

[RFC1321] Rivest, R., "The MD5 Message-Digest Algorithm", RFC 1321, DOI 10.17487/RFC1321, April 1992, <http://www.rfc-editor.org/info/rfc1321>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <http://www.rfc-editor.org/info/rfc2131>.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <http://www.rfc-editor.org/info/rfc2132>.

[RFC3004] Stump, G., Droms, R., Gu, Y., Vyaghrapuri, R., Demirtjis, A., Beser, B., and J. Privat, "The User Class Option for DHCP", RFC 3004, DOI 10.17487/RFC3004, November 2000, <http://www.rfc-editor.org/info/rfc3004>.

[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, <http://www.rfc-editor.org/info/rfc3315>.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <http://www.rfc-editor.org/info/rfc3633>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <http://www.rfc-editor.org/info/rfc4193>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <http://www.rfc-editor.org/info/rfc4291>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <http://www.rfc-editor.org/info/rfc4861>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, DOI 10.17487/RFC6092, January 2011, <http://www.rfc-editor.org/info/rfc6092>.

[RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <http://www.rfc-editor.org/info/rfc6206>.

[RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <http://www.rfc-editor.org/info/rfc6347>.

[RFC6603] Korhonen, J., Ed., Savolainen, T., Krishnan, S., and O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", RFC 6603, DOI 10.17487/RFC6603, May 2012, <http://www.rfc-editor.org/info/rfc6603>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <http://www.rfc-editor.org/info/rfc6763>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <http://www.rfc-editor.org/info/rfc7217>.

[RFC7695] Pfister, P., Paterson, B., and J. Arkko, "Distributed Prefix Assignment Algorithm", RFC 7695, DOI 10.17487/RFC7695, November 2015, <http://www.rfc-editor.org/info/rfc7695>.

[RFC7787] Stenberg, M. and S. Barth, "Distributed Node Consensus Protocol", RFC 7787, DOI 10.17487/RFC7787, April 2016, <http://www.rfc-editor.org/info/rfc7787>.

14.2. Informative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <http://www.rfc-editor.org/info/rfc1035>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC3118] Droms, R., Ed. and W. Arbaugh, Ed., "Authentication for DHCP Messages", RFC 3118, DOI 10.17487/RFC3118, June 2001, <http://www.rfc-editor.org/info/rfc3118>.

[RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <http://www.rfc-editor.org/info/rfc6234>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <http://www.rfc-editor.org/info/rfc6241>.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <http://www.rfc-editor.org/info/rfc6762>.

[RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <http://www.rfc-editor.org/info/rfc7084>.

[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, <http://www.rfc-editor.org/info/rfc7368>.

[RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <http://www.rfc-editor.org/info/rfc7525>.

Acknowledgments

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

Thanks to Eric Kline for the original border discovery work.

Authors' Addresses

Markus Stenberg
Independent
Helsinki 00930
Finland

Email: markus.stenberg@iki.fi

Steven Barth
Independent
Halle 06114
Germany

Email: cyrus@openwrt.org

Pierre Pfister
Cisco Systems
Paris
France

Email: pierre.pfister@darou.fr