

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

F. Templin, Ed.
Boeing Research & Technology
June 24, 2019

**A Unified Stateful/Stateless Configuration Service for IPv6
draft-templin-6man-dhcpv6-ndopt-08.txt**

Abstract

IPv6 Neighbor Discovery (IPv6ND) specifies a control message set for nodes to discover neighbors, routers, prefixes and other services on the link. It also supports a manner of StateLess Address AutoConfiguration (SLAAC), while the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) specifies a separate stateful service. This document presents IPv6ND extensions for providing a unified stateful/stateless configuration service.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

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

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

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

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

Table of Contents

1.	Introduction	2
2.	DHCPv6 Options in IPv6 ND Messages	4
2.1.	The DHCPv6 Option	4
2.2.	DHCPv6 Option Usage	5
2.3.	Stateful Provisioning Requirements	6
2.4.	Implementation Considerations	7
3.	PIO Options in RS Messages	7
3.1.	The PIO Option in RS Messages	7
3.2.	PIO Option Usage	7
3.3.	Stateful Provisioning Requirements	8
3.4.	Implementation Considerations	8
4.	Embedded Prefix Assertion	9
4.1.	Embedded Prefix Assertion	9
4.2.	Embedded Prefix Usage	9
4.3.	Stateful Provisioning Requirements	9
4.4.	Implementation Considerations	10
5.	Out-of-Band Network Login Messaging	10
5.1.	Out-of-Band Network Login	10
5.2.	Out-of-Band Network Login Usage	10
5.3.	Stateful Provisioning Requirements	11
5.4.	Implementation Considerations	11
6.	Implementation Status	11
7.	IANA Considerations	11
8.	Security Considerations	11
9.	Acknowledgements	12
10.	References	12
10.1.	Normative References	12
10.2.	Informative References	12
Appendix A.	Change Log	13
	Author's Address	14

1. Introduction

IPv6 Neighbor Discovery (IPv6ND) [[RFC4861](#)] specifies a control message set for nodes to discover neighbors, routers, prefixes and other services on the link. It also supports a manner of Stateless Address AutoConfiguration (SLAAC). The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) specifies a separate service for delegation of prefixes, addresses and any other stateful information [[RFC8415](#)]. This document presents IPv6ND extensions for providing a unified stateful/stateless configuration service.

Templin

Expires December 26, 2019

[Page 2]

If the network can provide such a unified service, multi-message procedures can be condensed into a single and concise message exchange. This would ease network management as well as simplify host and router operations. It would further accommodate both stateless and stateful services in a way that combines the best aspects of both. The operating model is based on harnessing the IPv6 ND Router Solicitation (RS) / Router Advertisement (RA) functions to provide all configuration information in a single message exchange.

When a node first comes onto a link, it sends an RS to elicit an RA from one or more routers for the link. If the node also needs to acquire stateful information it then sends a DHCPv6 Solicit message to elicit a Reply message from a DHCPv6 server. This two round-trip message exchange can add delay as well as waste critical link bandwidth on low-end links (e.g., 6LoWPAN, satellite communications, aeronautical wireless, etc.). While it is possible to start both round trip exchanges at the same time, this would still result in twice as many channel access transactions as necessary. Moreover, the multicast nature of these messages could disturb other nodes on the link, e.g., resulting in an unnecessary wakeup from sleep mode.

This document proposes methods for combining all stateless and stateful configuration operations into a single, unified exchange based on IPv6ND messaging extensions. It notes that stateful exchanges should include:

- o an explicit request for stateful information
- o the identity of the requesting node
- o a transaction identification that the requesting node can use to match replies with their corresponding requests
- o any security parameters necessary for the requesting node to establish its authorization to receive stateful information

The first method is through definition of a new IPv6ND option called the "DHCPv6 Option" that combines the IPv6ND router discovery and DHCPv6 stateful processes into a single message exchange. Nodes include the DHCPv6 option in RS messages to solicit an RA message with a DHCPv6 option in return. This allows the IPv6ND and DHCPv6 functions to work together to supply the client with all needed configuration information in a minimum number of messages.

The second method proposes the inclusion of Prefix Information Options (PIOs) in RS messages for the purpose of soliciting stateful information. [[I-D.naveen-slaac-prefix-management](#)] discusses the

maintenance and management functions required for supporting the operation.

The third method entails the encoding of a prefix in the IPv6 link-local source address of the RS message. If the node is pre-configured with the prefix that it will solicit from the network, and if the network has a way of accepting the node's prefix assertion without the threat of spoofing, the network can then delegate the prefix and establish the necessary routing information.

The fourth method uses out-of-band messaging for the node to request stateful information outside of the scope of IPV6ND messaging. The out-of-band messaging could entail some sort of network login process (e.g., through Layer-2 (L2) messaging, etc.).

The following sections present considerations for nodes that employ these approaches.

2. DHCPv6 Options in IPv6 ND Messages

The first method entails the inclusion of DHCPv6 messages within IPV6ND RS and RA messages, as discussed in the following sections.

2.1. The DHCPv6 Option

The DHCPv6 option is a new IPV6ND option that simply embeds a standard DHCPv6 message per section 6 of [RFC8415], beginning with the 'msg-type' followed by the 'transaction-id' and all DHCPv6 'options'. The format of the option is as follows:

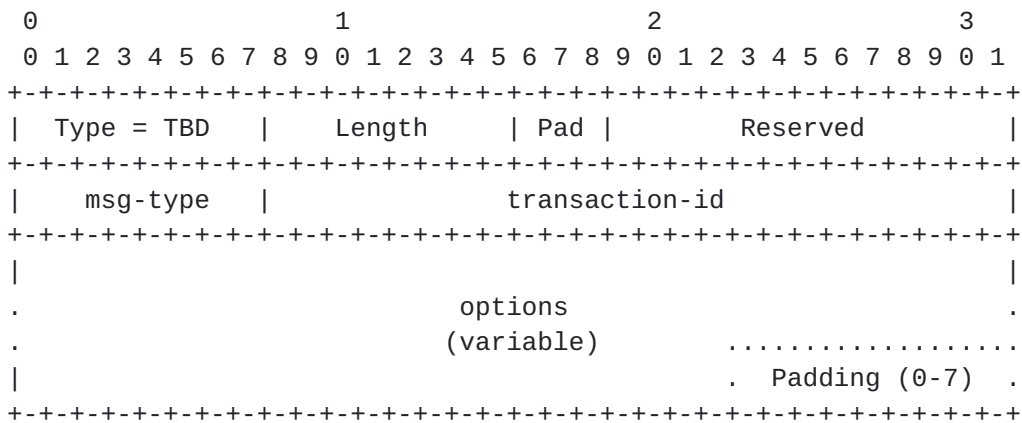


Figure 1: IPv6 ND DHCPv6 Option Format

In this format, 'Type' and 'Length' are exactly as defined in Section 4.6 of [RFC4861], 'Pad' is a 3-bit integer that encodes the padding length, 'Reserved' is included for alignment and future use,

and the rest of the option is formatted as specified in [Section 6 of \[RFC8415\]](#) except with trailing null padding added as necessary for 8 octet alignment. The length of the full DHCPv6 message is determined by $((('Length' * 8) - 4) - 'Pad')$, for a maximum message length of 2036 octets.

The 'Reserved' field MUST be set to 0 on transmission and ignored on reception. Future specifications MAY define new uses for these bits.

2.2. DHCPv6 Option Usage

When a node first comes onto the link, it creates an RS message containing a DHCPv6 option that embeds a DHCPv6 Solicit message. The Solicit may include a Rapid Commit option if a two-message exchange (i.e., instead of four) is required. The RS message may also include a Nonce option to provide an extended transaction identifier [\[RFC3971\]](#). The node then sends the RS message either to the unicast address of a specific router on the link, or to the all-routers multicast address.

When a router receives an RS message with a DHCPv6 option, if it does not recognize the option and/or does not employ a DHCPv6 relay agent or server, it returns an RA message as normal with any stateless configuration information and without including a DHCPv6 option. By receiving the RA message with no DHCPv6 option, the node can determine that the router does not recognize the option and/or does not support a DHCPv6 relay/server function. In this way, no harm will have come from the node including the DHCPv6 option in the RS, and the function is fully backwards compatible.

When a router receives an RS message with a DHCPv6 option, if it recognizes the option and employs a DHCPv6 relay agent or server, it extracts the encapsulated DHCPv6 message and forwards it to the relay agent or server. When the DHCPv6 message reaches a DHCPv6 server, the server processes the DHCPv6 Solicit message and prepares either an Advertise (four message) or Reply (two message) DHCPv6 message containing any delegated addresses, prefixes and/or any other information the server is configured to send. The server then returns the Advertise/Reply message to the router.

When the router receives the DHCPv6 Advertise/Reply message, it creates a Router Advertisement (RA) message that includes any autoconfiguration information necessary for the link and also embeds the DHCPv6 message in a DHCPv6 option within the body of the RA. (The RA also echos the Nonce value if a Nonce was included in the RS message.) The router then returns the RA as a unicast message response to the node that sent the RS.

In a two message exchange, the stateless/stateful exchange is completed when the node receives the RA. In a four message exchange, the requesting node can Decline any stateful information it does not wish to accept and/or send unicast Request options in subsequent RSeS to get RA messages with Reply options back from the router or routers of its choosing.

At any time after the initial RS/RA exchange, the node may need to issue DHCPv6 Renew, Release or Rebind messages to manage address/prefix lifetimes. In that case, the node prepares a DHCPv6 message option and inserts it in an RS message which it then sends via unicast to the router. The router in turn processes the message the same as for DHCPv6 Solicit/Reply.

At any time after the initial RS/RA exchange, the DHCPv6 server may need to issue a DHCPv6 Reconfigure message. In that case, when the router receives the DHCPv6 Reconfigure message it prepares a unicast RA message with a DHCPv6 option that encodes the Reconfigure and sends the RA as an unsolicited unicast message to the node. The node then follows the DHCPv6 client procedures for processing and responding to Reconfigure messages.

At any time after the initial RS/RA exchange, the router can initiate an unsolicited RA/Reply, e.g., to cause the node to update its configuration information quickly. In this method, the router sends a synthesized DHCPv6 Renew or Information-request message that induces the server to return a DHCPv6 Reply. The message includes the same DHCPv6 transaction-id and IPv6 ND Nonce values that the router had echoed in its initial Reply. The server then wraps the Reply message in the body of an RA message, and sends the unsolicited RA/Reply. When the node receives the unsolicited RA/Reply message, it matches the transaction-id and Nonce values with the initial RA/Reply it had received from the router. If the identification information matches, the node processes the message and initiates a new RS/RA exchange if necessary; otherwise it drops the message.

2.3. Stateful Provisioning Requirements

Using the DHCPv6 Option, the message itself includes sub-options to request stateful information. The DHCPv6 Device Unique Identifier (DUID) provides the identity of the requesting node, and the DHCPv6 transaction-id and IPv6 ND Nonce provide a unique identifier for matching RS and RA messages. Finally, the message can be protected using SEcure Neighbor Discovery (SEND) [[RFC3971](#)].

2.4. Implementation Considerations

The IPv6ND and DHCPv6 functions are typically implemented in separate router modules. In that case, the IPv6ND function extracts the DHCPv6 message from the option included in the RS message and wraps it in IP/UDP headers with the same addresses and port numbers the soliciting node would have used had it send an ordinary IP/UDP/DHCPv6 message. The IPv6ND function then acts as a Lightweight DHCPv6 Relay Agent (LDRA) [[RFC6221](#)] to forward the message to the DHCPv6 relay or server function on-board the router.

The forwarded DHCPv6 message then traverses any additional relays on the reverse path until it reaches the DHCPv6 server. When the DHCPv6 server processes the message, it delegates any necessary resources and returns a Reply via the same relay agent path as had occurred on the reverse path so that the Reply will eventually arrive back at the IPv6ND function. The IPv6ND function then prepares an RA message with any autoconfiguration information associated with the link, embeds the DHCPv6 message body in an IPv6ND DHCPv6 option, and returns the message via unicast to the node that sent the RS.

In an ideal implementation, the IPv6ND and DHCPv6 functions could be co-located in the same module on the router. In that way the two functions would be coupled as though they were in fact a single unified function without the need for any LDRA processing.

3. PIO Options in RS Messages

The second method entails the inclusion of Prefix Information Options (PIOs) in IPv6ND RS messages, as discussed in the following sections.

3.1. The PIO Option in RS Messages

This document proposes the inclusion of PIOs in RS messages to solicit and maintain prefixes that are delegated in subsequent RA messages. Prefix management is performed as discussed in [[I-D.naveen-slaac-prefix-management](#)] (an alternate prefix management proposal based on unsolicited advertisements with special flag settings is found in [[I-D.pioxfolks-6man-pio-exclusive-bit](#)]).

3.2. PIO Option Usage

When a node that wishes to request a prefix delegation first comes onto the link, it creates an RS message containing a PIO. It sets the Prefix Length to either the length of the prefix it wishes to receive or '0' (unspecified) if it will defer to the router's preference. The node then sets the Valid and Preferred Lifetimes to either its preferred values or '0' (unspecified) if it will defer to

the router's preference. The node then sets the Prefix to either the prefix it wishes to receive, or '0' (unspecified) if it will defer to the router's preference. The node then sends the RS message either to the unicast address of a specific router on the link, or to the all-routers multicast address.

When a router receives an RS message with a PIO, if it is not configured to accept PIOs in RS messages it returns an RA message as normal and without including a PIO. By receiving the RA message with no PIO, the node can determine that the router does not recognize the option and/or does not support an IPv6ND-based prefix delegation service. In this way, no harm will have come from the node including the PIO in the RS, and the function is fully backwards compatible.

When a router receives an RS message with a PIO, if it is configured to accept the option and can provide prefix delegation services it examines the fields in the message and selects a prefix to delegate to the node. If the PIO included a specific Prefix, the router delegates the node's preferred prefix if possible. Otherwise, the router selects a prefix to delegate to the node with length based on the node's Prefix Length. The router sets lifetimes matching the lifetimes requested by the node if possible, or shorter lifetimes if the node's requested lifetimes are too long. The router finally prepares a PIO containing this information and inserts it into an RA message to send back to the source of the RS.

3.3. Stateful Provisioning Requirements

Using the PIO in RS messages, the option itself requests stateful information. The RS message link-layer address can be used as the identity of the requesting node. The RS message includes a Nonce option [[RFC3971](#)] to provide a transaction identifier for matching RS and RA messages. Finally, the message can be protected using SEND the same as for the DHCPv6 option.

3.4. Implementation Considerations

Each router can implement a stateful database management service of their own choosing, but a functional alternative would be to use the standard DHCPv6 service as the back-end management service. In this way, all communications between the router's link to the requesting node are via RS/RA messaging. But, when the router receives an RS message with a PIO it can create a synthesized DHCPv6 Solicit message to send to the DHCPv6 server. This can be done in the same way as for the approach discussed in [Section 2.4](#). In this way, the node on the link over which the PIO is advertised only ever sees RS/RA messages on the front end, and the router gets to use the DHCPv6 service for stateful configuration management on the back end.

4. Embedded Prefix Assertion

The third method entails a simple RS/RA exchange with no additional options where the node asserts a prefix by embedding the prefix in the source address of the RS message. The following sections provide further details.

4.1. Embedded Prefix Assertion

In this method, the node is pre-provisioned with the prefix it will use on its downstream networks (e.g., through network management, manual configuration, etc.). To invoke this method, the node includes its pre-provisioned prefix in the link-local source address of its RS message according to the AERO address format [[I-D.templin-6man-aeroaddr](#)]. For example, if the node is pre-provisioned with the prefix 2001:db8:1000:2000::/64, it creates its IPv6 link-local source address as fe80::2001:db8:1000:2000.

4.2. Embedded Prefix Usage

When a node that wishes to assert a prefix first comes onto the link, it statelessly configures an AERO address based on its pre-provisioned prefix. The node then includes the AERO address as the source address of a standard RS message. If a router that receives the RS message has a way of verifying that the node is authorized to receive the solicited prefix, the router injects the prefix into the routing system and returns a standard RA message. When the node receives the RA message, it then has assurance that the proper routing state has been established.

The node examines the default router lifetime in the RA message as guidance for when subsequent RS/RA exchanges are necessary, i.e., the same as for normal IPv6ND. The node sends additional RS messages before the default router lifetime expires in order to keep the prefix assertion alive in the network. The RS messages may be sent either to the all-routers multicast address or to the unicast address(es) of the router(s) it received previous RAs from.

4.3. Stateful Provisioning Requirements

Using embedded prefix assertion, the network must have some way of determining the node's authority to assert its claimed prefix. This could be, e.g., through examination of the link-layer source address of the RS message. The network must also have some way of knowing the node's claimed prefix length, as the length cannot be conveyed in the RS message. If necessary, the exchange can also include some form of transaction identifier, e.g., by including a Nonce option in

the RS. Finally, the exchange can be protected using SEND the same as for the previous two methods.

4.4. Implementation Considerations

This method can be conducted using standard RS/RA messages with no extra options added to either message. It entails an administrative assignment of the node's AERO address to the upstream interface over which it will send the RS. When the router receives the standard RS message, it statelessly derives the node's prefix from the AERO address and injects the prefix into the routing system. The router then returns a standard RA message.

When the router returns the RA message, if it is configured to do so it can include a PIO option as discussed in [Section 3.1](#). The PIO option includes prefix lifetimes and the prefix length. This "hybrid" combination of methods two and three could be useful in some deployment scenarios.

As for the PIO-based service discussed in [Section 3.4](#), DHCPv6 can be used as the back-end service for stateful configuration management.

5. Out-of-Band Network Login Messaging

The fourth method entails an out-of-band messaging exchange through a "network login" procedure. During the network login, the requesting node could have an out-of-band messaging exchange with the network to set the stage for the router eventually sending an RA message as discussed in the following sections

5.1. Out-of-Band Network Login

In the out-of-band network login, the node signs into the network using, e.g., a login/password, a security certificate, etc. The node authenticates itself to the network, and can optionally have an iterative exchange to request certain aspects of the node's desired stateful configuration information. The first-hop router is then signaled to prepare an RA message to return to the node, i.e., either through some out-of-band signaling or through the node sending an RS message.

5.2. Out-of-Band Network Login Usage

When a node first comes onto the link, it engages in a network login session using some form of out-of-band messaging such as Layer-2 (L2) messaging. The session entails a security exchange where the node authenticates itself to the network and proves its authorization to receive the stateful configuration information. The network then

signals the router to send an RA message to the node, either unsolicited or in response to the node's RS message.

5.3. Stateful Provisioning Requirements

Using out-of-band messaging, the node engages in an iterative exchange where a request for stateful configuration information is conveyed. The exchange includes an identity for the requesting node and provides a unique per-message identifier so that the node can correlate its message requests with the responses it gets back from the network. Finally, the message exchange itself contains security parameters for authenticating the requesting node.

5.4. Implementation Considerations

The network login system and routers must be tightly coupled so that the network login can securely convey the requesting node's identity to the router.

As for the PIO-based service discussed in [Section 3.4](#), DHCPv6 can be used as the back-end service for managing the stateful configuration database.

6. Implementation Status

The approach discussed in [Section 2](#) has been implemented as extensions to the OpenVPN open source software distribution. The implementation is available at: <http://linkupnetworks.net/aero/AERO-OpenVPN-2.0.tgz>.

7. IANA Considerations

The IANA is instructed to assign an IPv6ND option Type value TBD for the DHCPv6 option.

The IANA is instructed to create a registry for the DHCPv6 option "Reserved" field (with no initial assignments) so that future uses of the field can be coordinated.

8. Security Considerations

Security considerations for IPv6 Neighbor Discovery [[RFC4861](#)] and DHCPv6 [[RFC8415](#)] apply to this document.

SEcure Neighbor Discovery (SEND) [[RFC3971](#)] can provide authentication for IPv6 ND messages with no need for additional securing mechanisms.

9. Acknowledgements

This work was motivated by discussions on the 6man and v6ops list. Those individuals who provided encouragement and critical review are acknowledged.

The following individuals provided useful comments that improved the document: Mikael Abrahamsson, Fred Baker, Ron Bonica, Yuçel Guven, Naveen Kottapalli, Ole Troan, Bernie Volz.

The following individuals developed IPv6ND and DHCPv6 extensions for OpenVPN: Kyle Bae, Wayne Benson, Eric Yeh.

This work is aligned with the NASA Safe Autonomous Systems Operation (SASO) program under NASA contract number NNA16BD84C.

This work is aligned with the FAA as per the SE2025 contract number DTFAWA-15-D-00030.

This work is aligned with the Boeing Information Technology (BIT) MobileNet program and the Boeing Research & Technology (BR&T) enterprise autonomy program.

10. References

10.1. Normative References

- [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, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 8415](#), DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

10.2. Informative References

[I-D.naveen-slaac-prefix-management]

Kottapalli, N., "IPv6 Stateless Prefix Management", [draft-naveen-slaac-prefix-management-00](#) (work in progress), November 2018.

[I-D.pioxfolks-6man-pio-exclusive-bit]

Kline, E. and M. Abrahamsson, "IPv6 Router Advertisement Prefix Information Option eXclusive Flag", [draft-pioxfolks-6man-pio-exclusive-bit-02](#) (work in progress), March 2017.

[I-D.templin-6man-aeroaddr]

Templin, F., "The AERO Address", [draft-templin-6man-aeroaddr-04](#) (work in progress), December 2018.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.

[RFC6221] Miles, D., Ed., Ooghe, S., Dec, W., Krishnan, S., and A. Kavanagh, "Lightweight DHCPv6 Relay Agent", [RFC 6221](#), DOI 10.17487/RFC6221, May 2011, <<https://www.rfc-editor.org/info/rfc6221>>.

Appendix A. Change Log

<< RFC Editor - remove prior to publication >>

Changes from -07 to -08:

- o Changed DHCPv6 reference to [RFC8415](#) - deprecates [RFC3315](#) and [RFC3633](#)
- o added prefix length to example in [Section 4.1](#).

Changes from -06 to -07:

- o Added "unsolicited DHCPv6 Reply" considerations
- o Added refeence to new IPv6ND-based PD proposal.
- o No longer associate the term "autoconfiguration" with the term "stateful".
- o Added URL for implementation.

Author's Address

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
USA

Email: fltemplin@acm.org