

DHC WG
Internet-Draft
Intended status: Standards Track
Expires: August 18, 2014

S. Jiang
B. Liu
Huawei Technologies Co., Ltd
February 14, 2014

Stateless Reconfiguration in Dynamic Host Configuration Protocol for
IPv6 (DHCPv6)
draft-jiang-dhc-stateless-reconfiguration-01

Abstract

This document defines a mechanism to propagate reconfigure messages towards stateless configured DHCPv6 clients.

Status of This Memo

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

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

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

This Internet-Draft will expire on August 18, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Requirements Language	4
3. Stateless Reconfiguration Use Cases	4
4. New DHCPv6 Specifications	5
4.1. Multicast Address	5
4.2. Stateless Reconfigure Message	5
5. Stateless Reconfiguration Procedure	5
5.1. Server Behavior	6
5.2. Relay Agent Behavior	7
5.3. Client Behavior	8
6. Security Considerations	8
7. IANA Considerations	8
8. Acknowledgements	9
9. References	9
9.1. Normative References	9
9.2. Informative References	9
Authors' Addresses	10

1. Introduction

[RFC3736] defines a stateless configuration procedure using DHCPv6. With it, the network configure information, such as the addresses of DNS recursive name servers, can be propagated to nodes, which have obtained their IPv6 addresses through some other mechanism. The basic scenario is that a newly online client initiates an information request to DHCPv6 server, then the server responses with requested configuration information. This mechanism is called the Stateless DHCPv6 services, because DHCPv6 servers do not maintain any dynamic state for individual clients, including the unicast addresses of clients.

However, the specification of stateless DHCPv6 service lacks a mechanism to inform these configured clients if some configuration information is changed. Transplanting Reconfigure message of [RFC3315] into stateless DHCPv6 services does not solve the issue, because in stateful DHCPv6, servers send Reconfigure messages to clients using their unicast addresses.

The lifetime option for DHCPv6 [RFC4242] assigns a lifetime to configuration information obtained through DHCPv6. At the expiration of the lifetime, the host contacts the DHCPv6 server to obtain updated configuration information. This lifetime gives the network administrator another mechanism to configure hosts with new configuration by controlling the time at which the host refreshes the list. However, such mechanism is not flexible enough: one aspect is the minimum of refresh time is 10 minutes, which is so long that it

might not be suitable for unplanned configuration changes; the other aspect is, in order to update the configuration quickly, the short time setting would cause un-necessary refresh all the time.

This document defines a mechanism to propagate a newly defined Stateless-Reconfigure message towards stateless configured DHCPv6 clients. It requests a mechanism for the DHCPv6 server to be aware of all relay agent destinations.

{Question to WG No.1} There are three potential mechanisms to create relay agent destinations on the DHCPv6 server:

a) Static configuration

Network administrators manually configure static unicast addresses of all relay agents on the DHCPv6 server.

Pros: no need to update any protocol/function implementation in relays; allows fast deployment in current network.

Cons: cost significant human management burden; error-prone, mistakenly configuring the relay addresses or leaving out some relays are expected.

b) Define a new ALL_RELAY_AGENT multicast address

The DHCPv6 server could send the stateless reconfiguration messages directly to the new multicast address.

Pros: a solid coverage of all relays.

Cons: network administrators need to maintain an all-relay-agent multicast group; all relays and DHCPv6 servers need to be updated to know the new multicast address.

c) DHCPv6 server dynamic learning

the DHCPv6 server dynamically records unicast addresses of all relay agents from client Information-request messages and maintains the relay addresses list. A keepalive mechanism is needed between relay agents and servers to track the availability of the list entries.

Pros: automatic processing without human intervene.

Cons: requires more function update to the DHCPv6 server; the keepalive mechanism requires more function/protocol burden to the whole DHCP system.

[Editor Notes] the current form of this document is based on only the first mechanism of above three. If the WG decided to change or include other mechanism, the document would be updated accordingly.

The document newly defines a new link-scope well-known all-client multicast address and a new DHCPv6 message type for stateless reconfiguration. Correspondent server behavior, agent behavior and client behavior are specified in details.

The design of new DHCPv6 elements and procedures obey the recommendations and guidance of [I-D.ietf-dhc-option-guidelines].

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] when they appear in ALL CAPS. When these words are not in ALL CAPS (such as "should" or "Should"), they have their usual English meanings, and are not to be interpreted as [RFC2119] key words.

3. Stateless Reconfiguration Use Cases

This section described scenarios where stateless reconfiguration are expected.

- Configuration error

Configuration errors, either caused by human or program, are hard to be immune in networks. Especially, human errors is identified as one of the top reasons of network failure. In stateless DHCPv6, if the administrators/program accidentally mis-configure the parameters (e.g. DNS), then significant network failure would be expected. Current protocols just lack the ability to eliminate the configuration errors when such accident happens. The hosts configured with wrong parameters can only wait until the wrong parameters lifetime expired then to refresh them. This would not be acceptable especially when the lifetime was long. The stateless reconfiguration mechanism could be highly expected in this scenario.

- Emergent event

The network needs to initially update the already configured parameters within a short period due to some emergent events; and waiting the clients to refresh the parameters according to the lifetime is just un-acceptable. These scenarios would also require stateless reconfiguration.

4. New DHCPv6 Specifications

This section define new DHCPv6 elements requested by the stateless configuration mechanism, including multicast address constant, and message type.

4.1. Multicast Address

`ALL_CLIENT_MULTICAST_ADDRESS` (FF02::xxxx, TBD1) A link-scoped multicast address used by a DHCPv6 server or relay agent to communicate with neighboring (i.e., on-link) clients. All clients are members of this multicast group

4.2. Stateless Reconfigure Message

A new Stateless-Reconfigure message, which is mainly based on server to clients advertise model, is defined in order to distinguish from the existing Reconfigure message, which is mainly based on server/client one-to-one model.

[Editor Notes] According to the results of Qestion 2 and Question 4 (in Section 5 & 7 below), there might be two new messages needed. Current document uses the alternative of one new message.

`STATELESS-RECONFIGURE-TRIGGER` Message type value is TBD2. It follows the message format specification, defined in Section 6 of [RFC3315]. A server sends a Stateless-Reconfigure message to a client to inform the client that the server has new or updated configuration parameters, and that the client is to initiate an Information-Request transaction with the server in order to receive the updated information.

5. Stateless Reconfiguration Procedure

{Question to WG No.2} There could be two kind of stateless DHCPv6 reconfiguration modes as the following, which one is proper? Or we should support both?

- Trigger mode

The server sends out a multicast Stateless-Reconfiguration message to `ALL_CLIENT_MULTICAST_ADDRESS`. As response, every client is requested to initiate an Information-Request message back to the server. The server can then inform the changed configuration information to clients.

This mode is similar with stateful DHCPv6 reconfiguration and also provide the potential possibility that the server response to information-request differently according to various user policies.

- Push mode

The server sends out a multicast Stateless-Reconfiguration message to ALL_CLIENT_MULTICAST_ADDRESS to directly advertise new configuration to the clients. The clients then update the parameters accordingly.

Trigger mode requires every client to initial individual request to servers. This is reasonable for the stateful information that need to be maintained and tracked in the servers, e.g. clients' IP addresses. But for the stateless information shared among clients (such as DNS), it might not necessary. Some resource constrained networks (e.g. a 802.15.4e/g based mesh network) might have efficiency problem with the trigger mode. These scenarios might significantly benefit from the push mode stateless reconfiguration mechanism. However, push mode seems breaking the traditional behavior model of DHCP. Whether it is a good break needs further discussion.

[Editor Notes] the current form of this document is based on triggering client information-request model, which complies the traditional behavior model of DHCPv6. If the WG chooses to directly advertise new configuration, the document would be updated accordingly.

5.1. Server Behavior

When the network configuration information on a stateless DHCPv6 server changes, the server creates and transmit a new Stateless-Reconfigure message towards all clients following the below steps:

- o The server sets the "msg-type" field to STATELESS-RECONFIGURE. The server sets the transaction-id field to 0. The server MUST include a Server Identifier option containing its DUID in the Reconfigure message.
- o The server MAY include an Option Request option to inform the client of what information has been changed or new information that has been added.
- o The server MUST NOT include a Reconfigure Message option (defined in section 22.19 of [RFC3315]), which is mandated in Reconfigure message to indicate the client to respond a Renew or an Information-Request message. It is because there is only one possible response on the client follow a Stateless-Reconfigure message - an Information-request message.

- o The server MUST NOT include any other options in the Reconfigure except as specifically allowed in the definition of individual options.
- o The server sends Stateless-Reconfigure message to its direct local link using ALL_CLIENT_MULTICAST_ADDRESS.
- o Simultaneously, the server uses a Relay-Reply message (as described in Section 20.3 of [RFC3315]) to send the Stateless-Reconfigure message to all relay agents in its static relay-agent-destination record using the unicast address of these relay agents. The peer-address of the Relay-Reply message MUST be filled by Relay-Reply message ALL_CLIENT_MULTICAST_ADDRESS.

Notes: since there is no previous Relay-Forward message that went through multiple relay agents and the server has to send the Relay-Reply message through the return same path, the server should be able to send the Relay-Reply message to the relay agent that direct connects with clients. Consequently, the Relay-Reply message SHOULD NOT contain another Relay-Reply message.

The below is an example of a typical Relay-Reply message that contains a Stateless-Reconfigure message:

```
msg-type: RELAY-REPLY
hop-count: 0
link-address: 0
peer-address: ALL_CLIENT_MULTICAST_ADDRESS
Relay Message option: <Stateless-Reconfigure message>
```

Servers MUST discard any received Stateless-Reconfigure messages.

5.2. Relay Agent Behavior

The relay agent extracts the Stateless-Reconfigure message from the Relay Message option and relays it to all clients. If the relay agent is attached to multiple links, it MUST broadcast the Stateless-Reconfigure message on every links. This behavior is compliance with normal behavior of relaying a Relay-reply message, defined in Section 20.2 of [RFC3315].

Relay agents MUST discard any received Stateless-Reconfigure messages. By design, relay agents do not process any directly received Stateless-Reconfigure messages.

The result is that the relay agent sends out a Stateless-Reconfigure message towards all client on the local link using ALL_CLIENT_MULTICAST_ADDRESS.

5.3. Client Behavior

Clients MUST discard any Stateless-Reconfigure messages that meets any of the following conditions:

- o the message was not multicast to the client using ALL_CLIENT_MULTICAST_ADDRESS.
- o the message does not include a Server Identifier option.
- o the message contains a Reconfigure Message Option, defined in Section 22.19 of [RFC3315].

Upon receipt of a valid Stateless-Reconfigure message, after a random delay time, the client responds with an Information-request message. The random delay time is designed to avoid congested Information-request on the server. While the transaction is in progress, the client silently discards any Stateless-Reconfigure messages it receives.

{Question to WG No.3} Should we define a maximum time of random delay time? If yes, should it come from server by a new option?

6. Security Considerations

Malicious server sends Stateless Reconfigure message to cause all clients response. There is the risk of denial of service attacks against DHCP clients and server. {Current auth mechanism cannot work in this broadcast model, server public key model maybe work.}

Since the clients response to Information-Request using the standard mechanism, defined in [RFC3315], the chance that receive wrong configuration information from malicious attackers does not raise.

7. IANA Considerations

Per this document, IANA has assigned one new well-known Multicast Address in the "IPv6 Multicast Address Space Registry" registry (currently located at <http://www.iana.org/assignments/ipv6-multicast-addresses>) for the following attribute:

ALL_CLIENT_MULTICAST_ADDRESS: (FF02::xxxx, TBD1).

Per this document, IANA has assigned one new DHCPv6 message type in the "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)" registry (currently located at <http://www.iana.org/assignments/dhcpv6-parameters>) for the following attribute:

STATELESS-RECONFIGURE Message Type, TBD2.

{Question to WG No.4} As raised in Question 2, if we support both Trigger mode and Push mode, then there should be two kind of corresponding messages. We could use two message types to distinguish them; or use a flag in one message type. Which is better?

8. Acknowledgements

The authors would like to thanks the valuable comments made by Suresh Krishnan, Ted Lemon, Bernie Voltz and other members of DHC WG.

This document was produced using the xml2rfc tool [RFC2629].

9. References

9.1. Normative References

- [I-D.ietf-dhc-option-guidelines]
Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", draft-ietf-dhc-option-guidelines-17 (work in progress), January 2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.

9.2. Informative References

- [RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.
- [RFC4242] Venaas, S., Chown, T., and B. Volz, "Information Refresh Time Option for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 4242, November 2005.

Authors' Addresses

Sheng Jiang
Huawei Technologies Co., Ltd
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: jiangsheng@huawei.com

Bing Liu
Huawei Technologies Co., Ltd
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: leo.liubing@huawei.com