

DHC Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 8, 2014

P. Patil
Cisco
M. Boucadair
France Telecom
D. Wing
T. Reddy
Cisco
February 4, 2014

IP Connectivity Status Notifications for DHCPv6
draft-ietf-dhc-conn-status-00

Abstract

This specification extends DHCPv6 so that a DHCPv6 Relay Agent can dynamically inform the DHCPv6 server about the IP connectivity status of a host. The IP connectivity status information is also triggered by any change in the connectivity as provided to the host. The DHCPv6 server uses this information as an input to its decision-making about configuration parameters to be conveyed to that host.

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 8, 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.	Terminology	3
3.	Problem Statement: Focus on DNS Reconfiguration	3
4.	Host Connectivity Status Option	4
5.	DHCPv6 Relay Agent Behavior	5
5.1.	Relay Forward	6
5.2.	Reconfigure Request	6
6.	DHCPv6 Server Behavior	6
6.1.	Relay Forward	7
6.2.	Reconfigure Request	7
7.	Host Tracking	7
8.	Security Considerations	7
9.	IANA Considerations	8
10.	References	8
10.1.	Normative References	8
10.2.	Informative References	8
	Authors' Addresses	9

[1.](#) Introduction

Some networks are expected to support IPv4-only, dual-stack, and IPv6-only hosts at the same time. Due to devices capabilities and available connectivity types, providing generic configuration from a DHCP server to connected hosts is sub-optimal in most cases, and may even break functionality in some cases. The network infrastructure is usually well equipped to be aware of the connectivity delivered to connected hosts. The network can also track and detect transitions from single to dual-stack or vice-versa.

This document specifies a DHCPv6 extension for relay agents to indicate status of hosts connectivity to remote DHCPv6 servers. The information passed by a relay is generic and a DHCPv6 server can interpret and process this information to make a more informed decision on the configuration parameters that a client is to receive.

The DHCPv6 server can either be configured or have built-in logic to use this information as desired, which is outside the scope of this document.

[Section 3](#) describes a typical problem that can be solved owing to the mechanism described in this specification. A DHCPv6 server prioritizes the DNS servers to be sent back to a requesting client based on host connectivity characteristics provided by the DHCPv6 relay agent.

While the host stack can be upgraded to send this information to the DHCPv6 server on its own, a generalized upgrade of all DHCPv6 client implementations on all operating systems is extremely difficult.

[DISCUSSION NOTE: A companion solution could be to define a container that can be used to return per-AF specific configuration parameters to the client. In such a scheme, the server blindly returns all pieces of configuration and it is up to the client to make use of the appropriate set of parameters according to its available connectivity. This alternative assumes an update at the dhcp client's side. This approach can be seen as complimentary to the one defined in this specification. The document will be updated to reflect consensus of the WG on whether the additional option is to be specified.]

2. Terminology

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

Dual-Stack host: Denotes a host that is configured with both an IPv4 address and IPv6 prefix and is reachable using both IPv4 and IPv6 connectivity.

3. Problem Statement: Focus on DNS Reconfiguration

Default address selection rules specified in [[RFC6724](#)] prefers IPv6 over IPv4. If a dual-stack host is configured to use a DNS64 server [[RFC6147](#)], it will send its DNS queries to that DNS64 server which will synthesize a AAAA response if no A records are found. Thus, a dual-stack host will always use IPv6 if a DNS lookup was involved, even if IPv4 could have been used more optimally.

In some deployments, if NAT44 [[RFC3022](#)] and NAT64 [[RFC6146](#)] are deployed within the same network, it is preferable to use NAT44 over NAT64 because of scale, performance and application incompatibility issues (e.g., FTP) [[RFC6384](#)]. At the same time, native IPv6 can still be preferred over IPv4.

A DHCPv6 relay agent can observe host characteristics on a network to determine if a host is IPv4-only, dual-stack, or IPv6-only and also

detect transitions from single to dual-stack or vice-versa. This information can be used by the DHCPv6 relay agent to influence the DHCPv6 server to send appropriately prioritized DNS Servers to the client. The DHCPv6 server can implement the following based on connectivity information received from the DHCPv6 relay agent.

- o IPv6-only transition to Dual-Stack: In case a host is IPv6-only, it is provided with a DNS64 server. When transitioning to dual-stack, an IPv4 DNS server is assigned as a consequence of obtaining an IPv4 Address. The DHCPv6 relay agent can detect this and send a RECONFIGURE_REQUEST message [[RFC6977](#)] to the DHCPv6 server indicating that the host needs to be provided with a regular DNS server. In lieu of this mechanism, the host would continue to use the DNS64 server until the host stack reinitializes.
- o Dual-Stack to IPv6-only: In case a host is dual-stack, it is provided with a regular DNS server followed by DNS64 server. When transitioning to IPv6-only, the DHCPv6 relay agent can detect this change and send a RECONFIGURE_REQUEST message to the DHCPv6 server indicating that the host needs to be assigned a DNS64 server only. In lieu of this mechanism, the host would continue to use the regular DNS Server which is inaccessible and eventually time out to fail over to the DNS64 Server. The host will take additional time to fully initialize causing delays in connection.

4. Host Connectivity Status Option

The option (Figure 1) includes an 8-bit status code that indicates specific host connectivity characteristics. The option can be included by a DHCPv6 relay agent in RELAY-FORW and RECONFIGURE-REQUEST.

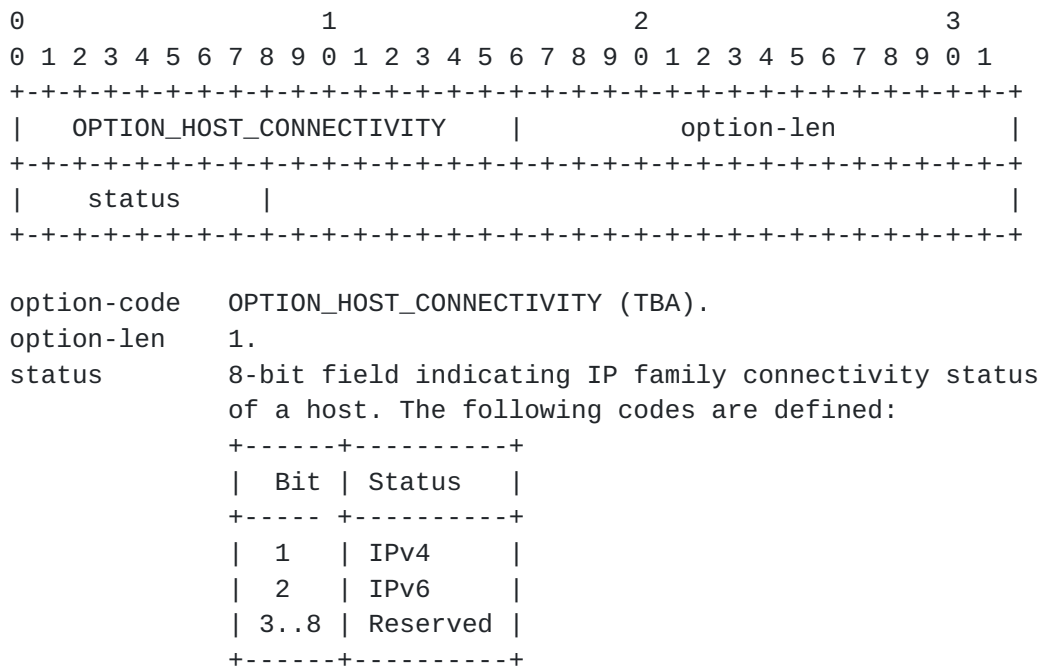


Figure 1: Relay Agent Host Connectivity Option message format

- o IPv4 : The host that is configured with an IPv4 address and is reachable using IPv4.
- o IPv6 : The host that is configured with an IPv6 prefix and is reachable using IPv6.
- o If both bits are enabled, the host is dual-stack i.e the host that is configured with both an IPv4 address and IPv6 prefix and is reachable using both IPv4 and IPv6 connectivity.

5. DHCPv6 Relay Agent Behavior

DHCPv6 relay agents that implement this specification MUST be configurable for tracking host connectivity and inserting the `OPTION_HOST_CONNECTIVITY` option in `RELAY-FORW` and `RECONFIGURE-REQUEST` messages.

To be able to notify details of hosts' connectivity, a DHCPv6 relay agent must be able to track host connectivity. A DHCPv6 relay agent can detect host connectivity type using mechanisms discussed in [Section 7](#). The DHCPv6 relay agent then includes this information in the appropriate DHCPv6 message.

DHCPv6 relay agents need to maintain connectivity state of each host it can track. This ensures that notifications to the DHCPv6 server, especially DHCPv6 `RECONFIGURE-REQUEST`, are accurately sent when there

is a change in status. If a DHCPv6 relay agent loses state due to some reason (e.g., during restart events), it will build state again using the mechanisms described in [Section 7](#) and then send appropriate notifications to the server. Such notifications are redundant and a DHCPv6 server can choose to ignore such redundant notifications from the DHCPv6 relay agent. Redundant notifications are also possible when DHCPv6 relay agents are deployed in fault tolerant mode.

[5.1.](#) Relay Forward

DHCPv6 relay agents that implement this specification MAY include the option `OPTION_HOST_CONNECTIVITY` in the `RELAY_FORW` to indicate status of host connectivity.

[5.2.](#) Reconfigure Request

DHCPv6 relay agents that implement this specification MUST be configurable for sending the `RECONFIGURE_REQUEST` message. The DHCPv6 relay agent generates a Reconfigure-Request [[RFC6977](#)] anytime status of host connectivity changes by including `OPTION_HOST_CONNECTIVITY` in the request.

[6.](#) DHCPv6 Server Behavior

A DHCPv6 server that supports `OPTION_HOST_CONNECTIVITY` may either have specific configuration or built-in logic to process information available in the option and send configuration parameters in DHCPv6 responses. How the server consumes and acts on the information obtained in the option is outside the scope of this document.

The DHCPv6 server may use this connectivity information, if available, in addition to other DHCPv6 relay agent option data, other options included in the DHCPv6 client messages, server configuration, and physical network topology information in order to assign appropriate configuration to the client.

The server MUST ignore the option if it doesn't recognize the status in the `OPTION_HOST_CONNECTIVITY` option. The server SHOULD maintain the latest status received from the DHCPv6 relay agent. The server can use this state to match against subsequent notifications and only further process if there is change in status. A DHCPv6 relay agent could, for reasons such as restart, fault-tolerant mode etc, send redundant notifications and matching of status at the server will avoid unnecessary processing and message exchanges.

6.1. Relay Forward

Upon receiving a RELAY-FORW message containing OPTION_HOST_CONNECTIVITY, the server can send appropriate configuration in the RELAY-REPLY response. The server MUST NOT return this option in a RELAY-REPLY message.

6.2. Reconfigure Request

Upon receiving a RECONFIFURE-REQUEST message containing an OPTION_HOST_CONNECTIVITY option, the server MUST follow the mechanism described in [[RFC6977](#)] to create and send Reconfigure message. The server MUST NOT return this option in a RECONFIGURE-REPLY message.

7. Host Tracking

DHCPv6 relay agents can actively keep track of all IPv4/IPv6 addresses and associated lease times assigned to hosts via the respective DHCP servers. DHCPv6 relay agents can therefore detect transitions from single to dual-stack and vice-versa efficiently. In addition to this technique, DHCPv6 relay agents closest to the client can detect transitions using snooping mechanisms. Network devices today use mechanisms such as ARP and NDP snooping (bindings learnt by snooping all NDP traffic, NS, NA, RS, RA) to determine host characteristics such as IPv4/IPv6 - MAC - DUID bindings. IPv4/IPv6 and MAC counters are also used to determine host liveliness.

First hop devices that implement first hop security features can also track IP address bindings and determine binding updates such as temporary addresses, deprecated addresses, etc. Existing work such as [[I-D.ietf-savi-dhcp](#)] and [[I-D.levy-abegnoli-savi-plbt](#)] also aim to active current host bindings, all of which can be leveraged to track host addresses.

These mechanisms help determine if a particular IP address family is inactive, has reverted to using a single stack even though it initially had dual-stack capabilities and detect active dual-stack usage after long periods of single-stack activity.

Other techniques to track host connectivity can be envisaged. It is out of scope of this document to provide an exhaustive list of host tracking techniques.

8. Security Considerations

This document describes an application of the mechanism specified in [[RFC6977](#)].

Host tracking mechanisms MUST be reliable. If a DHCPv6 relay agent is compromised, it may be used to force an uncompromised DHCPv6 server abuse DHCPv6 clients by triggering repetitive reconfigurations. Security considerations described in [RFC6977] are applicable to this specification.

9. IANA Considerations

IANA is requested to assign the following new DHCPv6 Option Code in the registry maintained in <http://www.iana.org/assignments/dhcpv6-parameters>:

- o OPTION_HOST_CONNECTIVITY

10. References

10.1. Normative References

- [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.
- [RFC6977] Boucadair, M. and X. Pournard, "Triggering DHCPv6 Reconfiguration from Relay Agents", [RFC 6977](#), July 2013.

10.2. Informative References

- [I-D.ietf-savi-dhcp] Bi, J., Wu, J., Yao, G., and F. Baker, "SAVI Solution for DHCP", [draft-ietf-savi-dhcp-18](#) (work in progress), June 2013.
- [I-D.levy-abegnoli-savi-plbt] Levy-Abegnoli, E., "Preference Level based Binding Table", [draft-levy-abegnoli-savi-plbt-02](#) (work in progress), March 2010.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3646](#), December 2003.

- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), April 2011.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [RFC 6147](#), April 2011.
- [RFC6384] van Beijnum, I., "An FTP Application Layer Gateway (ALG) for IPv6-to-IPv4 Translation", [RFC 6384](#), October 2011.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", [RFC 6724](#), September 2012.

Authors' Addresses

Prashanth Patil
Cisco Systems, Inc.
Bangalore
India

Email: praspati@cisco.com

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: dwing@cisco.com

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tiredy@cisco.com