

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 10, 2013

S. Perreault
Viagenie
W. George
Time Warner Cable
T. Tsou
Huawei Technologies (USA)
July 9, 2012

Turning off IPv4 Using DHCPv6
draft-perreault-sunset4-noipv4-00

Abstract

This memo defines a new DHCPv6 option for indicating to a dual-stack host or router that IPv4 is to be turned off.

Status of this Memo

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

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

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

This Internet-Draft will expire on January 10, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
2. Terminology	3
3. The No-IPv4 Option	3
3.1. Wire Format	3
3.2. Semantics	4
3.3. Example	6
4. Open Issues	7
5. Security Considerations	7
6. IANA Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Authors' Addresses	8

1. Introduction

When a dual-stack host makes a DHCPv4 request, it typically interprets the absence of a response as a failure condition. This makes it difficult to deploy such nodes in an IPv6-only network.

Take for example a home router that is dual-stack capable but provisioned with an IPv6-only WAN connection. When the router boots, it typically assigns an IPv4 address to its LAN interface, starts services on that interface, and starts handing out IPv4 addresses to clients on the LAN by answering DHCPv4 requests. This is done unconditionally, without taking the status of the IPv4 connectivity on the WAN interface into account. Hosts on the LAN, in turn, install a default route pointing to the router and start behaving as if IPv4 connectivity was available. IPv4 packets destined to the Internet get dropped at the router and timeouts happen. The end result is that IPv4 remains fully active on the LAN and on the router itself even when it is desired that it be turned off.

A new mechanism is needed to indicate the absence of IPv4 connectivity. Given that the goal is turning off IPv4, this new signaling mechanism shall be transported over IPv6. Therefore, we introduce a new DHCPv6 [RFC3315] option for the purpose of explicitly indicating to the DHCPv6 client that IPv4 connectivity is unavailable.

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].

The following terms are also used in this document:

Upstream Interface: An interface on which the No-IPv4 DHCPv6 option is received by a DHCPv6 client.

3. The No-IPv4 Option

3.1. Wire Format

The No-IPv4 DHCPv6 option is used to signal the unavailability of IPv4 connectivity. The format of the No-IPv4 option is:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               OPTION_NO_IPV4               | option-len |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   v4-level   |
+---+---+---+---+---+

```

option-code OPTION_NO_IPV4 (TBD).

option-len 1.

v4-level Level of IPv4 functionality.

The DHCPv6 client MUST place the OPTION_NO_IPV4 option code in the Option Request Option ([RFC3315] section 22.7). Servers MAY include the option in responses (if they have been so configured). Servers MAY also place the OPTION_NO_IPV4 option code in an Option Request Option contained in a Reconfigure message.

3.2. Semantics

The option applies to the link on which it is received by the DHCPv6 client. It is used to indicate to the client that it should disable some or all of its IPv4 functionality. What should be disabled depends on the value of v4-level.

v4-level can take the following values:

- 0 - IPv4 fully enabled: This is equivalent to the absence of the No-IPv4 option. It is included here so that a DHCPv6 server can explicitly re-enable IPv4 access by including it in a Reply message following a Reconfigure.
- 1 - No IPv4 upstream, local IPv4 permitted: Any kind of IPv4 connectivity is unavailable on the link on which the option is received. Therefore, any attempts to provision IPv4 by the host or to use IPv4 in any fashion, on that link, will be useless. IPv4 MAY be dropped, blocked, or otherwise ignored on that link.

Upon reception of the No-IPv4 option with value 1, the following IPv4 functionality MUST be disabled on the Upstream Interface:

- A. IPv4 addresses MUST NOT be assigned.
- B. Currently-assigned IPv4 addresses MUST be unassigned.

- C. Dynamic configuration of link-local IPv4 addresses [RFC3927] MUST be disabled.
- D. IPv4, ICMPv4, or ARP packets MUST NOT be sent.
- E. IPv4, ICMPv4, or ARP packets received MUST be ignored.
- F. DNS A queries MUST NOT be sent, even transported over IPv6.

If all DHCPv6-configured interfaces receive the No-IPv4 option with value 1 or 2, and no other interface provides IPv4 connectivity to the Internet, IPv4 is partially shut down, leaving only local connectivity active. On the Upstream Interface, IPv4 MUST be shut down as listed above. On other interfaces, IPv4 addresses MUST NOT be assigned except for the following:

- * Loopback (127.0.0.0/8)
- * Link Local (169.254.0.0/16) [RFC3927]
- * Private-Use (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16) [RFC1918]

- 2 - No IPv4 at all: This is intended to be a stricter version of the above.

The host or router running the DHCPv6 client that receives this option MUST disable IPv4 functionality on the Upstream Interface in the same way as for value 1.

If all DHCPv6-configured interfaces received the No-IPv4 option with exclusively value 2, and no other interface provides IPv4 connectivity to the Internet, IPv4 is completely shut down. In particular:

- A. IPv4 address MUST NOT be assigned to any interface.
- B. Currently-assigned IPv4 addresses MUST be unassigned.
- C. Dynamic configuration of link-local IPv4 addresses [RFC3927] MUST be disabled.
- D. IPv4, ICMPv4, or ARP packets MUST NOT be sent on any interface.
- E. IPv4, ICMPv4, or ARP packets received on any interface MUST be ignored.

- F. In the above, "any interface" includes loopback interfaces. In particular, the 127.0.0.1 special address **MUST** be removed.
- G. Server programs listening on IPv4 addresses (e.g., a DHCPv4 server) **MAY** be shut down.
- H. DNS A queries **MUST NOT** be sent, even transported over IPv6.
- I. If the host or router also runs a DHCPv6 server, it **SHOULD** include the No-IPv4 option with value 2 in DHCPv6 responses it sends to clients that request it, unless prohibited by local policy. If it currently has active clients, it **SHOULD** send a Reconfigure to each of them with the `OPTION_NO_IPV4` included in the Option Request Option.

The intent is to remove all traces of IPv4 activity. Once the No-IPv4 option with value 2 is activated, the network stack should behave as if IPv4 functionality had never been present. For example, a modular kernel implementation could accomplish the above by unloading the IPv4 kernel module at run time.

3.3. Example

A dual-stack home gateway is set up with a single WAN uplink and is configured to use DHCPv4 and DHCPv6 to automatically obtain IPv4 and IPv6 connectivity. On the LAN side, it has one link with multiple hosts.

When it boots, the router assigns 192.168.1.1/24 to its LAN interfaces and starts a DHCPv4 server listening on it. It hands out addresses 191.168.1.100-199 to clients. It also starts an IPv6 Router Advertisement daemon as well as a stateless DHCPv6 server, also listening on the LAN interfaces.

On the WAN side, it starts two provisioning procedures in parallel: one for IPv4 and one for IPv6.

At this point, the ISP does not know if the router supports IPv6-only operation. Therefore, by default, the ISP responds to DHCPv4 requests as usual.

As part of the IPv6 provisioning procedure, the router sends a DHCPv6 request containing `OPTION_NO_IPV4` in an Option Request Option. The ISP's DHCPv6 server's reply includes the No-IPv4 option with value 2. When this procedure finishes, the ISP has determined that this customer will run in IPv6-only mode and starts dropping all IPv4 packets at the first hop. If an IPv4 address was assigned, it is reclaimed, and possibly reassigned to another subscriber.

The home router aborts the IPv4 provisioning procedure (if it is still running) and deactivates all IPv4 functionality. It shuts down its DHCPv4 server. It also configures its own stateless DHCPv6 server to send the No-IPv4 option to clients that request it.

As an optimization, the router could delay setting up IPv4 by a few seconds (10 seconds seems reasonable). If the IPv6 procedure completes with the No-IPv4 option during that time, IPv4 will never have been set up and the router will operate in pure IPv6-only mode from the start.

4. Open Issues

- o A legacy IPv4-only device connected to a network running in mode 2 (no IPv4 at all) will presumably keep retrying forever, e.g. sending DHCPDISCOVER messages endlessly. Do we want a way to signal to that host that IPv4 will never be available? But since that device was not updated for IPv6, it is doubtful that it would be updated to understand this new signaling. Could we reuse/overload some existing signaling that would have the same effect?

5. Security Considerations

One security concern is that an attacker could use the No-IPv4 option to deny IPv4 access to a victim. However, unprotected vanilla DHCP can already be exploited to cause such a denial of service ([RFC2131] section 7).

TO BE COMPLETED

6. IANA Considerations

IANA is requested to assign value TBD with description OPTION_NO_IPV4 in the "DHCP Option Codes" table which is part of the dhcpv6-parameters registry [1].

7. References

7.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

- [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.
- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, May 2005.

7.2. Informative References

- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

URIs

- [1] <<http://www.iana.org/assignments/dhcpv6-parameters>>

Authors' Addresses

Simon Perreault
Viagenie
246 Aberdeen
Quebec, QC G1R 2E1
Canada

Phone: +1 418 656 9254
Email: simon.perreault@viagenie.ca
URI: <http://viagenie.ca>

Wes George
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
USA

Email: wesley.george@twcable.com

Tina Tsou
Huawei Technologies (USA)
2330 Central Expressway
Santa Clara, CA 95050
USA

Phone: +1 408 330 4424
Email: tina.tsou.zouting@huawei.com

