

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 5, 2013

JP. Dionne
S. Perreault
Viagenie
T. Tsou
Huawei Technologies (USA)
July 4, 2012

Gap Analysis for IPv4 Sunset
draft-dionne-sunset4-v4gapanalysis-00

Abstract

Sunsetting IPv4 refers to the process of turning off IPv4 definitively. It can be seen as the final phase of the migration to IPv6. This memo analyses difficulties arising when sunseting IPv4, and identifies the gaps resulting in additional work.

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 5, 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. Related Work	3
3. Remotely Disabling IPv4	3
3.1. Indicating that IPv4 connectivity is unavailable	3
3.2. Disabling IPv4 in the LAN	3
4. Client Connection Establishment Behavior	4
5. Disabling IPv4 in Operating System and Applications	5
6. IANA Considerations	5
7. Security Considerations	5
8. Acknowledgements	5
9. Informative References	5
Authors' Addresses	7

1. Introduction

The final phase of the migration to IPv6 is the sunset of IPv4, that is turning off IPv4 definitively on the attached networks and on the upstream networks.

Some current implementations behavior make it hard to sunset IPv4. Additionally, some new features could be added to IPv4 to make its sunsetting easier. This document analyzes the current situation and proposes new work in this area.

2. Related Work

[RFC3789], [RFC3790], [RFC3791], [RFC3792], [RFC3793], [RFC3794], [RFC3795] and [RFC3796] contain surveys of IETF protocols with their IPv4 dependencies.

3. Remotely Disabling IPv4

3.1. Indicating that IPv4 connectivity is unavailable

PROBLEM 1: When an IPv4 node boots and requests an IPv4 address (e.g., using DHCP), it typically interprets the absence of a response as a failure condition even when it is not.

PROBLEM 2: Home router devices often identify themselves as default routers in DHCP responses that they send to requests coming from the LAN, even in the absence of IPv4 connectivity on the WAN.

One way to address these issues is to send a signal to an dual-stack node that IPv4 connectivity is unavailable. Given that IPv4 shall be off, the message must be delivered through IPv6.

3.2. Disabling IPv4 in the LAN

PROBLEM 3: IPv4-enabled hosts inside an IPv6-only LAN can auto-configure IPv4 addresses [RFC3927] and enable various protocols over IPv4 such as mDNS [I-D.cheshire-dnsext-multicastdns] and LLmNR [RFC4795]. This can be undesirable for operational or security reasons, since in the absence of IPv4, no monitoring or logging of IPv4 will be in place.

PROBLEM 4: IPv4 can be completely disabled on a link by filtering it on the L2 switching device. However, this may not be possible in all cases or complex to deploy. For example, an ISP is often not able to control the L2 switching device in the subscriber home network.

One way to address these issues is to send a signal to an dual-stack node that auto-configuration of IPv4 addresses is undesirable, or that direct IPv4 communications between nodes on the same link should not take place.

This problem was described in [RFC2563], which standardized a DHCP option to disable IPv4 address auto-configuration. However, using this option requires running an IPv4 DHCP server, which is contrary to the goal of IPv4 sunsetting. An equivalent way of signalling this over IPv6 is necessary,, using either Router Advertisements or DHCPv6.

Furthermore, it could be useful to have L2 switches snoop this signalling and automatically start filtering IPv4 traffic as a consequence.

Finally, it could be useful to publish guidelines on how to safely block IPv4 on an L2 switch.

4. Client Connection Establishment Behavior

PROBLEM 5: Happy Eyeballs [RFC6555] refers to the multiple approaches to dual-stack client implementations that try to reduce connection setup delays by trying both IPv4 and IPv6 paths simultaneously. Some implementations introduce delays which provide an advantage to IPv6, while others do not [Huston2012]. The latter will pick the fastest path, no matter whether it is over IPv4 or IPv6, directing more traffic over IPv4 than the other kind of implementations. This can prove problematic in the context of IPv4 sunsetting, especially for Carrier-Grade NAT phasing out.

PROBLEM 6: `getaddrinfo()` [RFC3493] sends DNS queries for both A and AAAA records regardless of the state of IPv4 or IPv6 availability. The `AI_ADDRCONFIG` flag can be used to change this behavior, but it relies on programmers using the `getaddrinfo()` function to always pass this flag to the function. The current situation is that in an IPv6-only environment, many useless A queries are made.

Recommendations on client connection establishment behavior that would facilitate IPv4 sunsetting is therefore appropriate.

5. Disabling IPv4 in Operating System and Applications

PROBLEM 7: Completely disabling IPv4 at runtime often reveals implementation bugs. Hard-coded dependencies on IPv4 abound, such as on the 127.0.0.1 address assigned to the loopback interface. It is therefore often operationally impossible to completely disable IPv4 on individual nodes.

PROBLEM 8: In an IPv6-only world, legacy IPv4 code in operating systems and applications incurs a maintenance overhead and can present security risks.

It is possible to completely remove IPv4 support from an operating system as has been shown by the work of Bjoern Zeeb on FreeBSD. [Zeeb] Removing IPv4 support in the kernel revealed many IPv4 dependencies in libraries and applications.

It would be useful for the IETF to provide guidelines to programmers on how to avoid creating dependencies on IPv4, how to discover existing dependencies, and how to eliminate them. Having programs and operating systems that behave well in an IPv6-only environment is a prerequisite for IPv4 sunsetting.

6. IANA Considerations

None.

7. Security Considerations

TODO

8. Acknowledgements

TODO

9. Informative References

[Huston2012]

Huston, G. and G. Michaelson, "RIPE 64: Analysing Dual

Stack Behaviour and IPv6 Quality", April 2012.

- [I-D.cheshire-dnsext-multicastdns]
Cheshire, S. and M. Krochmal, "Multicast DNS",
draft-cheshire-dnsext-multicastdns-15 (work in progress),
December 2011.
- [RFC2563] Troll, R., "DHCP Option to Disable Stateless Auto-
Configuration in IPv4 Clients", RFC 2563, May 1999.
- [RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W.
Stevens, "Basic Socket Interface Extensions for IPv6",
RFC 3493, February 2003.
- [RFC3789] Nesser, P. and A. Bergstrom, "Introduction to the Survey
of IPv4 Addresses in Currently Deployed IETF Standards
Track and Experimental Documents", RFC 3789, June 2004.
- [RFC3790] Mickles, C. and P. Nesser, "Survey of IPv4 Addresses in
Currently Deployed IETF Internet Area Standards Track and
Experimental Documents", RFC 3790, June 2004.
- [RFC3791] Olvera, C. and P. Nesser, "Survey of IPv4 Addresses in
Currently Deployed IETF Routing Area Standards Track and
Experimental Documents", RFC 3791, June 2004.
- [RFC3792] Nesser, P. and A. Bergstrom, "Survey of IPv4 Addresses in
Currently Deployed IETF Security Area Standards Track and
Experimental Documents", RFC 3792, June 2004.
- [RFC3793] Nesser, P. and A. Bergstrom, "Survey of IPv4 Addresses in
Currently Deployed IETF Sub-IP Area Standards Track and
Experimental Documents", RFC 3793, June 2004.
- [RFC3794] Nesser, P. and A. Bergstrom, "Survey of IPv4 Addresses in
Currently Deployed IETF Transport Area Standards Track and
Experimental Documents", RFC 3794, June 2004.
- [RFC3795] Sofia, R. and P. Nesser, "Survey of IPv4 Addresses in
Currently Deployed IETF Application Area Standards Track
and Experimental Documents", RFC 3795, June 2004.
- [RFC3796] Nesser, P. and A. Bergstrom, "Survey of IPv4 Addresses in
Currently Deployed IETF Operations & Management Area
Standards Track and Experimental Documents", RFC 3796,
June 2004.
- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic

Configuration of IPv4 Link-Local Addresses", RFC 3927,
May 2005.

[RFC4795] Aboba, B., Thaler, D., and L. Esibov, "Link-local
Multicast Name Resolution (LLMNR)", RFC 4795,
January 2007.

[RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with
Dual-Stack Hosts", RFC 6555, April 2012.

[Zeeb] "FreeBSD Snapshots without IPv4 support",
<<http://wiki.freebsd.org/IPv6Only>>.

Authors' Addresses

Jean-Philippe Dionne
Viagenie
246 Aberdeen
Quebec, QC G1R 2E1
Canada

Phone: +1 418 656 9254
Email: jean-philippe.dionne@viagenie.ca
URI: <http://viagenie.ca>

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

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

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

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

