draft-dionne-sunset4-v4gapanalysis-00

Jean-Philippe Dionne <jean-philippe.dionne@viagenie.ca>
Simon Perreault <simon.perreault@viagenie.ca>
Tina Tsou <tina.tsou.zouting@huawei.com>

IETF 84, Vancouver
SUNSET4 meeting

2012-07-30

Introduction

- Suppose you want to turn off IPv4.
  - That is the final stage of the migration to IPv6.
- What problems will you encounter?
- What is missing in current protocols and tools?
- We have identified 8 problems.

Remotely Disabling IPv4

- 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.

Remotely Disabling IPv4

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

Remotely Disabling IPv4

- Disabling IPv4 in the LAN
  - Problem 3: IPv4-enabled hosts inside an IPv6-only LAN can auto-configure IPv4 addresses and enable various protocols over IPv4 such as mDNS and LLMNR.

    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.

Remotely Disabling IPv4

- Disabling IPv4 in the LAN
  - 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.
  - RFC 2563: DHCPv4 option to deactivate IPv4 auto-configuration
    - IPv6 equivalent is needed.
  - Filtering at L2 level is possible.
    - Guidelines on how to do it safely are missing.

Client Connection Establishment Behaviour

- Problem 5:
  - Happy Eyeballs
  - Some implementations introduce delays which provide an advantage to IPv6, while others do not.
  - 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.

Client Connection Establishment Behaviour

- Problem 6
  - getaddrinfo() 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.

Client Connection Establishment Behaviour

- Recommendations on client connection establishment behaviour that would facilitate IPv4 sunsetting are necessary.

Disabling IPv4 in Operating Systems 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.

Disabling IPv4 in Operating Systems and Applications

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

Disabling IPv4 in Operating Systems and Applications

- It is possible to completely remove IPv4 support from an operating system as has been shown by the work of Bjoern Zeeb on FreeBSD.
  - Many dependencies on IPv4 were discovered.
- Guidelines for programmers that would be useful:
  - How to avoid creating dependencies on IPv4
  - How to discover existing dependencies
  - How to eliminate them

Conclusion

- Questions?
- Is this a good gap analysis?
- Can it be adopted as a basis of SUNSET4 work?
- Merge with draft-zhou-sunset4-scenarios?