

IPv6 Operations Working Group (v6ops)
Internet-Draft
Intended status: Informational
Expires: December 13, 2021

F. Gont
SI6 Networks
N. Hilliard
INEX
G. Doering
SpaceNet AG
W. Kumari
Google
G. Huston
APNIC
W. Liu
Huawei Technologies
June 11, 2021

Operational Implications of IPv6 Packets with Extension Headers
draft-ietf-v6ops-ipv6-ehs-packet-drops-08

Abstract

This document summarizes the operational implications of IPv6 extension headers specified in the IPv6 protocol specification (RFC8200), and attempts to analyze reasons why packets with IPv6 extension headers are often dropped in the public Internet.

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 <https://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 December 13, 2021.

Copyright Notice

Copyright (c) 2021 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 (<https://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. Disclaimer	3
4. Background Information	3
5. Previous Work on IPv6 Extension Headers	5
6. Packet Forwarding Engine Constraints	7
6.1. Recirculation	8
7. Requirement to Process Layer-3/layer-4 information in Intermediate Systems	8
7.1. ECMP and Hash-based Load-Sharing	8
7.2. Enforcing infrastructure ACLs	9
7.3. DDoS Management and Customer Requests for Filtering	10
7.4. Network Intrusion Detection and Prevention	10
7.5. Firewalling	11
8. Operational and Security Implications	12
8.1. Inability to Find Layer-4 Information	12
8.2. Route-Processor Protection	12
8.3. Inability to Perform Fine-grained Filtering	12
8.4. Security Concerns Associated with IPv6 Extension Headers	12
9. IANA Considerations	14
10. Security Considerations	14
11. Acknowledgements	14
12. References	14
12.1. Normative References	14
12.2. Informative References	15
Authors' Addresses	19

1. Introduction

IPv6 Extension Headers (EHs) allow for the extension of the IPv6 protocol, and provide support for core functionality such as IPv6 fragmentation. However, common implementation limitations suggest that EHs present a challenge for IPv6 packet routing equipment and middle-boxes, and evidence exists that IPv6 packets with EHs are intentionally dropped in the public Internet in some circumstances.

This document has the following goals:

- o Raise awareness about the operational and security implications of IPv6 Extension Headers specified in [RFC8200], and present reasons why some networks resort to intentionally dropping packets containing IPv6 Extension Headers.
- o Highlight areas where current IPv6 support by networking devices maybe sub-optimal, such that the aforementioned support is improved.
- o Highlight operational issues associated with IPv6 extension headers, such that those issues are considered in IETF standardization efforts.

Section 4 provides background information about the IPv6 packet structure and associated implications. Section 5 of this document summarizes the previous work that has been carried out in the area of IPv6 extension headers. Section 6 discusses packet forwarding engine constraints in contemporary routers. Section 7 discusses why intermediate systems may need to access Layer-4 information to make a forwarding decision. Finally, Section 8 discusses the operational implications of IPv6 EHs.

2. Terminology

This document uses the term "intermediate system" to describe both routers and middle-boxes, when there is no need to distinguish between the two and where the important issue is that the device being discussed forwards packets.

3. Disclaimer

This document analyzes the operational challenges represented by packets that employ IPv6 Extension Headers, and documents some of the operational reasons why these packets are often dropped in the public Internet. This document is not a recommendation to drop such packets, but rather an analysis of why they are currently dropped.

4. Background Information

It is useful to compare the basic structure of IPv6 packets against that of IPv4 packets, and analyze the implications of the two different packet structures.

IPv4 packets have a variable-length header size, that allows for the use of IPv4 "options" -- optional information that may be of use by nodes processing IPv4 packets. The IPv4 header length is specified

in the IHL header field of the mandatory IPv4 header, and must be in the range from 20 octets (the minimum IPv4 header size) to 60 octets (accommodating at most 40 octets of options). The upper-layer protocol type is specified via the "Protocol" field of the mandatory IPv4 header.




Figure 1: IPv4 Packet Structure

IPv6 took a different approach to the IPv6 packet structure. Rather than employing a variable-length header as IPv4 does, IPv6 employs a linked-list-like packet structure, where a mandatory fixed-length IPv6 header is followed by an arbitrary number of optional extension headers, with the upper-layer header being the last header in the IPv6 header chain. Each extension header typically specifies its length (unless it is implicit from the extension header type), and the "next header" type that follows in the IPv6 header chain.

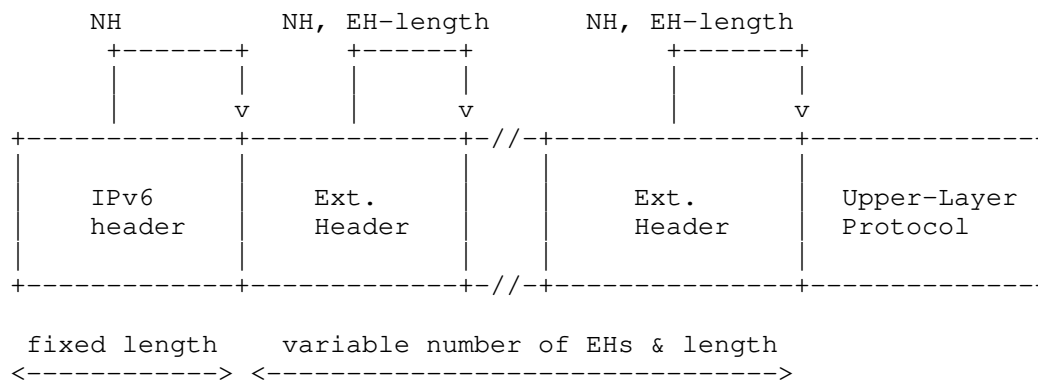


Figure 2: IPv6 Packet Structure

This packet structure has the following implications:

- o [RFC8200] requires the entire IPv6 header chain to be contained in the first fragment of a packet, therefore limiting the IPv6 extension header chain to the size of the path MTU.
- o Other than the path MTU constraints, there are no other limits to the number of IPv6 EHs that may be present in a packet. Therefore, there is no upper-limit regarding "how deep into the IPv6 packet" the upper-layer may be found.
- o The only way for a node to obtain the upper-layer protocol type or find the upper-layer protocol header is to parse and process the entire IPv6 header chain, in sequence, starting from the mandatory IPv6 header, until the last header in the IPv6 header chain is found.

5. Previous Work on IPv6 Extension Headers

Some of the operational and security implications of IPv6 Extension Headers have been discussed at the IETF:

- o [I-D.taylor-v6ops-fragdrop] discusses a rationale for which operators drop IPv6 fragments.
- o [I-D.wkumari-long-headers] discusses possible issues arising from "long" IPv6 header chains.
- o [I-D.kampanakis-6man-ipv6-eh-parsing] describes how inconsistencies in the way IPv6 packets with extension headers are parsed by different implementations could result in evasion of security controls, and presents guidelines for parsing IPv6 extension headers with the goal of providing a common and consistent parsing methodology for IPv6 implementations.
- o [I-D.ietf-opsec-ipv6-eh-filtering] analyzes the security implications of IPv6 EHs, and the operational implications of dropping packets that employ IPv6 EHs and associated options.
- o [RFC7113] discusses how some popular RA-Guard implementations are subject to evasion by means of IPv6 extension headers.
- o [RFC8900] analyzes the fragility introduced by IP fragmentation.

A number of recent RFCs have discussed issues related to IPv6 extension headers, specifying updates to a previous revision of the IPv6 standard [RFC2460], many of which have now been incorporated into the current IPv6 core standard [RFC8200] or the IPv6 Node Requirements [RFC8504]. Namely,

- o [RFC5095] discusses the security implications of Routing Header Type 0 (RTH0), and deprecates it.
- o [RFC5722] analyzes the security implications of overlapping fragments, and provides recommendations in this area.
- o [RFC7045] clarifies how intermediate nodes should deal with IPv6 extension headers.
- o [RFC7112] discusses the issues arising in a specific fragmentation case where the IPv6 header chain is fragmented into two or more fragments (and formally forbids such fragmentation).
- o [RFC6946] discusses a flawed (but common) processing of the so-called IPv6 "atomic fragments", and specified improved processing of such packets.
- o [RFC8021] deprecates the generation of IPv6 atomic fragments.
- o [RFC8504] clarifies processing rules for packets with extension headers, and also allows hosts to enforce limits on the number of options included in IPv6 EHs.
- o [RFC7739] discusses the security implications of predictable fragment Identification values, and provides recommendations for the generation of these values.
- o [RFC6980] analyzes the security implications of employing IPv6 fragmentation with Neighbor Discovery for IPv6, and formally recommends against such usage.

Additionally, [RFC8200] has relaxed the requirement that "all nodes examine and process the Hop-by-Hop Options header" from [RFC2460], by specifying that only nodes that have been explicitly configured to process the Hop-by-Hop Options header are required to do so.

A number of studies have measured the extent to which packets employing IPv6 extension headers are dropped in the public Internet:

- o [PMTUD-Blackholes] and [Linkova-Gont-IEPG90] presented some preliminary measurements regarding the extent to which packet containing IPv6 EHs are dropped in the public Internet.
- o [RFC7872] presents more comprehensive results and documents the methodology used to obtain these results.
- o [Huston-2017] and [Huston-2020] measured packet drops resulting from IPv6 fragmentation when communicating with DNS servers.

6. Packet Forwarding Engine Constraints

Most contemporary carrier-grade routers use dedicated hardware, e.g. application-specific integrated circuits (ASICs) or network processing units (NPUs), to determine how to forward packets across their internal fabrics (see [IEPG94-Scudder] and [APNIC-Scudder] for details). One of the common methods of handling next-hop lookup is to send a small portion of the ingress packet to a lookup engine with specialised hardware, e.g. ternary content-addressable memory (TCAM) or reduced latency dynamic random-access memory (RLDRAM), to determine the packet's next-hop. Technical constraints mean that there is a trade-off between the amount of data sent to the lookup engine and the overall packet forwarding rate of the lookup engine. If more data is sent, the lookup engine can inspect further into the packet, but the overall packet forwarding rate of the system will be reduced. If less data is sent, the overall packet forwarding rate of the router will be increased but the packet lookup engine may not be able to inspect far enough into a packet to determine how it should be handled.

NOTE:

For example, some contemporary high-end routers are known to inspect up to 192 bytes, while others are known to parse up to 384 bytes of header.

If a hardware forwarding engine on a contemporary router cannot make a forwarding decision about a packet because critical information is not sent to the look-up engine, then the router will normally drop the packet. Section 7 discusses some of the reasons for which a contemporary router might need to access layer-4 information to make a forwarding decision.

Historically, some packet forwarding engines punted packets of this form to the control plane for more in-depth analysis, but this is unfeasible on most contemporary router architectures as a result of the vast difference between the hardware forwarding capacity of the router and processing capacity of the control plane and the size of the management link which connects the control plane to the forwarding plane. Other platforms may have a separate software forwarding plane that is distinct both from the hardware forwarding plane and the control plane. However, the limited CPU resources of this software-based forwarding plane, as well as the limited bandwidth of the associated link results in similar throughput constraints.

If an IPv6 header chain is sufficiently long that it exceeds the packet look-up capacity of the router, the router might be unable to

determine how the packet should be handled, and thus could resort to dropping the packet.

6.1. Recirculation

Although TLV chains are amenable to iterative processing on architectures that have packet look-up engines with deep inspection capabilities, some packet forwarding engines manage IPv6 Extension Header chains using recirculation. This approach processes Extension Headers one at a time: when processing on one Extension Header is completed, the packet is looped back through the processing engine again. This recirculation process continues repeatedly until there are no more Extension Headers left to be processed.

Recirculation is typically used on packet forwarding engines with limited look-up capability, because it allows arbitrarily long header chains to be processed without the complexity and cost associated with packet forwarding engines which have deep look-up capabilities. However, recirculation can impact the forwarding capacity of hardware, as each packet will pass through the processing engine multiple times. Depending on configuration, the type of packets being processed, and the hardware capabilities of the packet forwarding engine, this could impact data-plane throughput performance on the router.

7. Requirement to Process Layer-3/layer-4 information in Intermediate Systems

The following subsections discuss some of the reasons for which intermediate systems may need to process Layer-3/layer-4 information to make a forwarding decision.

7.1. ECMP and Hash-based Load-Sharing

In the case of equal cost multi-path (ECMP) load sharing, the intermediate system needs to make a decision regarding which of its interfaces to use to forward a given packet. Since round-robin usage of the links is usually avoided to prevent packet reordering, forwarding engines need to use a mechanism that will consistently forward the same data streams down the same forwarding paths. Most forwarding engines achieve this by calculating a simple hash using an n-tuple gleaned from a combination of layer-2 through to layer-4 packet header information. This n-tuple will typically use the src/dst MAC address, src/dst IP address, and if possible further layer-4 src/dst port information.

In the IPv6 world, flows are expected to be identified by means of the IPv6 Flow Label [RFC6437]. Thus, ECMP and Hash-based Load-

Sharing should be possible without the need to process the entire IPv6 header chain to obtain upper-layer information to identify flows. [RFC7098] discusses how the IPv6 Flow Label can be used to enhance layer 3/4 load distribution and balancing for large server farms.

Historically, many IPv6 implementations failed to set the Flow Label, and hash-based ECMP/load-sharing devices also did not employ the Flow Label for performing their task. While support of [RFC6437] is currently widespread for current versions of all popular host implementations, there is still only marginal usage of the IPv6 Flow Label for ECMP and load balancing [Cunha-2020]. A contributing factor could be the issues that have been found in host implementations and middle-boxes [Jaeggli-2018].

Clearly, widespread support of [RFC6437] would relieve intermediate systems from having to process the entire IPv6 header chain, making Flow Label-based ECMP and Load-Sharing [RFC6438] feasible.

If an intermediate system cannot determine consistent n-tuples for calculating flow hashes, data streams are more likely to end up being distributed unequally across ECMP and load-shared links. This may lead to packet drops or reduced performance.

7.2. Enforcing infrastructure ACLs

Infrastructure ACLs (iACLs) drop unwanted packets destined to a network's infrastructure. Typically, iACLs are deployed because external direct access to a network's infrastructure addresses is operationally unnecessary, and can be used for attacks of different sorts against router control planes. To this end, traffic usually needs to be differentiated on the basis of layer-3 or layer-4 criteria to achieve a useful balance of protection and functionality. For example, an infrastructure may be configured with the following policy:

- o Permit some amount of ICMP echo (ping) traffic towards a router's addresses for troubleshooting.
- o Permit BGP sessions on the shared network of an exchange point (potentially differentiating between the amount of packets/seconds permitted for established sessions and connection establishment), but do not permit other traffic from the same peer IP addresses.

If a forwarding router cannot determine consistent n-tuples for calculating flow hashes, data streams are more likely to end up being distributed unequally across ECMP and load-shared links. This may lead to packet drops or reduced performance.

If a network cannot deploy infrastructure ACLs, then the security of the network may be compromised due to having more potential attack vectors open.

7.3. DDoS Management and Customer Requests for Filtering

The case of customer DDoS protection and edge-to-core customer protection filters is similar in nature to the iACL protection. Similar to iACL protection, layer-4 ACLs generally need to be applied as close to the edge of the network as possible, even though the intent is usually to protect the customer edge rather than the provider core. Application of layer-4 DDoS protection to a network edge is often automated using Flowspec [RFC8955] [RFC8956].

For example, a web site that normally only handled traffic on TCP ports 80 and 443 could be subject to a volumetric DDoS attack using NTP and DNS packets with randomised source IP address, thereby rendering traditional [RFC5635] source-based real-time black hole mechanisms useless. In this situation, DDoS protection ACLs could be configured to block all UDP traffic at the network edge without impairing the web server functionality in any way. Thus, being able to block arbitrary protocols at the network edge can avoid DDoS-related problems both in the provider network and on the customer edge link.

7.4. Network Intrusion Detection and Prevention

Network Intrusion Detection Systems (NIDS) examine network traffic and try to identify traffic patterns that can be correlated to network-based attacks. These systems generally inspect application-layer traffic (if possible), but at the bare minimum inspect layer-4 flows. When attack activity is inferred, the operator is notified of the potential intrusion attempt.

Network Intrusion Prevention Systems (IPS) operate similarly to NIDS's, but they can also prevent intrusions by reacting to detected attack attempts by e.g., triggering packet filtering policies at firewalls and other devices.

Use of extension headers can be problematic for NIDS/IPS, since:

- o Extension headers increase the complexity of resulting traffic, and the associated work and system requirements to process it.
- o Use of unknown extension headers can prevent an NIDS/IPS from processing layer-4 information.

- o Use of IPv6 fragmentation requires a stateful fragment-reassembly operation, even for decoy traffic employing forged source addresses (see e.g., [nmap]).

As a result, in order to increase the efficiency or effectiveness of these systems, packets employing IPv6 extension headers are often dropped at the network ingress point(s) of networks that deploy these systems.

7.5. Firewalling

Firewalls enforce security policies by means of packet filtering. These systems usually inspect layer-3 and layer-4 traffic, but can often also examine application-layer traffic flows.

As with NIDS/IPS (Section 7.4), use of IPv6 extension headers can represent a challenge to network firewalls, since:

- o Extension headers increase the complexity of resulting traffic, and the associated work and system requirements to process it, as outlined in [Zack-FW-Benchmark].
- o Use of unknown extension headers can prevent firewalls from processing layer-4 information.
- o Use of IPv6 fragmentation requires a stateful fragment-reassembly operation, even for decoy traffic employing forged source addresses (see e.g., [nmap]).

Additionally, a common firewall filtering policy is the so-called "default deny", where all traffic is blocked (by default), and only expected traffic is added to an "allow/accept list".

As a result, packets employing IPv6 extension headers are often dropped by network firewalls, either because of the challenges represented by extension headers or because the use of IPv6 extension headers has not been explicitly allowed.

Note that although the data presented in [Zack-FW-Benchmark] were several years old at the time of publication of this document, many contemporary firewalls use comparable hardware and software architecture, and consequently the conclusions of this benchmark are still relevant, despite its age.

8. Operational and Security Implications

8.1. Inability to Find Layer-4 Information

As discussed in Section 7, intermediate systems that need to find the layer-4 header must process the entire IPv6 extension header chain. When such devices are unable to obtain the required information, the forwarding device has the option to drop the packet unconditionally, forward the packet unconditionally, or process the packet outside the normal forwarding path. Forwarding packets unconditionally will usually allow for the circumvention of security controls (see e.g., Section 7.5), while processing packets outside of the normal forwarding path will usually open the door to DoS attacks (see e.g., Section 6). Thus, in these scenarios, devices often simply resort to dropping such packets unconditionally.

8.2. Route-Processor Protection

Most contemporary carrier-grade routers have a fast hardware-assisted forwarding plane and a loosely coupled control plane, connected together with a link that has much less capacity than the forwarding plane could handle. Traffic differentiation cannot be performed by the control plane, because this would overload the internal link connecting the forwarding plane to the control plane.

The Hop-by-Hop Options header has been particularly challenging since in most circumstances, the corresponding packet is punted to the control plane for processing. As a result, many operators drop IPv6 packets containing this extension header [RFC7872]. [RFC6192] provides advice regarding protection of a router's control plane.

8.3. Inability to Perform Fine-grained Filtering

Some intermediate systems do not have support for fine-grained filtering of IPv6 extension headers. For example, an operator that wishes to drop packets containing Routing Header Type 0 (RHT0), may only be able to filter on the extension header type (Routing Header). This could result in an operator enforcing a more coarse filtering policy (e.g., "drop all packets containing a Routing Header" vs. "only drop packets that contain a Routing Header Type 0").

8.4. Security Concerns Associated with IPv6 Extension Headers

The security implications of IPv6 Extension Headers generally fall into one or more of these categories:

- o Evasion of security controls

- o DoS due to processing requirements
- o DoS due to implementation errors
- o Extension Header-specific issues

Unlike IPv4 packets where the upper-layer protocol can be trivially found by means of the "IHL" ("Internet Header Length") IPv4 header field, the structure of IPv6 packets is more flexible and complex. This can represent a challenge for devices that need to find this information, since locating upper-layer protocol information requires that all IPv6 extension headers be examined. In turn, this presents implementation difficulties, since some packet filtering mechanisms that require upper-layer information (even if just the upper layer protocol type) can be trivially circumvented by inserting IPv6 Extension Headers between the main IPv6 header and the upper layer protocol. [RFC7113] describes this issue for the RA-Guard case, but the same techniques could be employed to circumvent other IPv6 firewall and packet filtering mechanisms. Additionally, implementation inconsistencies in packet forwarding engines can result in evasion of security controls [I-D.kampanakis-6man-ipv6-eh-parsing] [Atlasis2014] [BH-EU-2014].

Sometimes packets with IPv6 Extension Headers can impact throughput performance on intermediate systems. Unless appropriate mitigations are put in place (e.g., packet dropping and/or rate-limiting), an attacker could simply send a large amount of IPv6 traffic employing IPv6 Extension Headers with the purpose of performing a Denial of Service (DoS) attack (see Section 6.1 and Section 8 for further details).

NOTE:

In the most trivial case, a packet that includes a Hop-by-Hop Options header might go through the slow forwarding path, to be processed by the router's CPU. Alternatively, a router configured to enforce an ACL based on upper-layer information (e.g., upper layer protocol or TCP Destination Port) may need to process the entire IPv6 header chain in order to find the required information, thereby causing the packet to be processed in the slow path [Cisco-EH-Cons]. We note that, for obvious reasons, the aforementioned performance issues can affect other devices such as firewalls, Network Intrusion Detection Systems (NIDS), etc. [Zack-FW-Benchmark]. The extent to which performance is affected on these devices is implementation-dependent.

IPv6 implementations, like all other software, tend to mature with time and wide-scale deployment. While the IPv6 protocol itself has existed for over 20 years, serious bugs related to IPv6 Extension

Header processing continue to be discovered (see e.g., [Cisco-Frag], [Microsoft-SA], and [FreeBSD-SA]). Because there is currently little operational reliance on IPv6 Extension headers, the corresponding code paths are rarely exercised, and there is the potential for bugs that still remain to be discovered in some implementations.

IPv6 Fragment Headers are employed to allow fragmentation of IPv6 packets. While many of the security implications of the fragmentation / reassembly mechanism are known from the IPv4 world, several related issues have crept into IPv6 implementations. These range from denial of service attacks to information leakage, as discussed in [RFC7739], [Bonica-NANOG58] and [Atlasis2012]).

9. IANA Considerations

This document has no IANA actions.

10. Security Considerations

The security implications of IPv6 extension headers are discussed in Section 8.4. This document does not introduce any new security issues.

11. Acknowledgements

The authors would like to thank (in alphabetical order) Mikael Abrahamsson, Fred Baker, Dale W. Carder, Brian Carpenter, Tim Chown, Owen DeLong, Gorrry Fairhurst, Guillermo Gont, Tom Herbert, Lee Howard, Tom Petch, Sander Steffann, Eduard Vasilenko, Eric Vyncke, Rob Wilton, Jingrong Xie, and Andrew Yourtchenko, for providing valuable comments on earlier versions of this document.

Fernando Gont would like to thank Jan Zorz / Go6 Lab <<https://go6lab.si/>>, Jared Mauch, and Sander Steffann <<https://steffann.nl/>>, for providing access to systems and networks that were employed to perform experiments and measurements involving packets with IPv6 Extension Headers.

12. References

12.1. Normative References

- [RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", RFC 5095, DOI 10.17487/RFC5095, December 2007, <<https://www.rfc-editor.org/info/rfc5095>>.

- [RFC5722] Krishnan, S., "Handling of Overlapping IPv6 Fragments", RFC 5722, DOI 10.17487/RFC5722, December 2009, <<https://www.rfc-editor.org/info/rfc5722>>.
- [RFC6946] Gont, F., "Processing of IPv6 "Atomic" Fragments", RFC 6946, DOI 10.17487/RFC6946, May 2013, <<https://www.rfc-editor.org/info/rfc6946>>.
- [RFC6980] Gont, F., "Security Implications of IPv6 Fragmentation with IPv6 Neighbor Discovery", RFC 6980, DOI 10.17487/RFC6980, August 2013, <<https://www.rfc-editor.org/info/rfc6980>>.
- [RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of Oversized IPv6 Header Chains", RFC 7112, DOI 10.17487/RFC7112, January 2014, <<https://www.rfc-editor.org/info/rfc7112>>.
- [RFC8021] Gont, F., Liu, W., and T. Anderson, "Generation of IPv6 Atomic Fragments Considered Harmful", RFC 8021, DOI 10.17487/RFC8021, January 2017, <<https://www.rfc-editor.org/info/rfc8021>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <<https://www.rfc-editor.org/info/rfc8504>>.

12.2. Informative References

- [APNIC-Scudder]
Scudder, J., "Modern router architecture and IPv6", APNIC Blog, June 4, 2020, <<https://blog.apnic.net/2020/06/04/modern-router-architecture-and-ipv6/>>.
- [Atlasis2012]
Atlasis, A., "Attacking IPv6 Implementation Using Fragmentation", BlackHat Europe 2012. Amsterdam, Netherlands. March 14-16, 2012, <https://media.blackhat.com/bh-eu-12/Atlasis/bh-eu-12-Atlasis-Attacking_IPv6-Slides.pdf>.

[Atlasis2014]

Atlasis, A., "A Novel Way of Abusing IPv6 Extension Headers to Evade IPv6 Security Devices", May 2014, <<http://www.insinuator.net/2014/05/a-novel-way-of-abusing-ipv6-extension-headers-to-evade-ipv6-security-devices/>>.

[BH-EU-2014]

Atlasis, A., Rey, E., and R. Schaefer, "Evasion of High-End IDPS Devices at the IPv6 Era", BlackHat Europe 2014, 2014, <<https://www.ernw.de/download/eu-14-Atlasis-Rey-Schaefer-briefings-Evasion-of-HighEnd-IPS-Devices-wp.pdf>>.

[Bonica-NANOG58]

Bonica, R., "IPv6 FRAGMENTATION: The Case For Deprecation", NANOG 58. New Orleans, Louisiana, USA. June 3-5, 2013, <<https://www.nanog.org/sites/default/files/mon.general.fragmentation.bonica.pdf>>.

[Cisco-EH-Cons]

Cisco, "IPv6 Extension Headers Review and Considerations", October 2006, <http://www.cisco.com/en/US/technologies/tk648/tk872/technologies_white_paper0900aecd8054d37d.pdf>.

[Cisco-Frag]

Cisco, "Cisco IOS XR Software Crafted IPv6 Packet Denial of Service Vulnerability", June 2015, <<http://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20150611-iosxr>>.

[Cunha-2020]

Cunha, I., "IPv4 vs IPv6 load balancing in Internet routes", NPS/CAIDA 2020 Virtual IPv6 Workshop, 2020, <<https://www.cmand.org/workshops/202006-v6/slides/cunha.pdf>>.

[FreeBSD-SA]

FreeBSD, "FreeBSD Security Advisory FreeBSD-SA-20:24.ipv6: IPv6 Hop-by-Hop options use-after-free bug", September 2020, <<https://www.freebsd.org/security/advisories/FreeBSD-SA-20:24.ipv6.asc>>.

[Huston-2017]

Huston, G., "Dealing with IPv6 fragmentation in the DNS", APNIC Blog, 2017, <<https://blog.apnic.net/2017/08/22/dealing-ipv6-fragmentation-dns/>>.

[Huston-2020]

Huston, G., "Measurement of IPv6 Extension Header Support", NPS/CAIDA 2020 Virtual IPv6 Workshop, 2020, <<https://www.cmand.org/workshops/202006-v6/slides/2020-06-16-xtn-hdrs.pdf>>.

[I-D.ietf-opsec-ipv6-eh-filtering]

Gont, F. and W. Liu, "Recommendations on the Filtering of IPv6 Packets Containing IPv6 Extension Headers at Transit Routers", draft-ietf-opsec-ipv6-eh-filtering-07 (work in progress), January 2021.

[I-D.kampanakis-6man-ipv6-eh-parsing]

Kampanakis, P., "Implementation Guidelines for parsing IPv6 Extension Headers", draft-kampanakis-6man-ipv6-eh-parsing-01 (work in progress), August 2014.

[I-D.taylor-v6ops-fragdrop]

Jaeggli, J., Colitti, L., Kumari, W., Vyncke, E., Kaeo, M., and T. Taylor, "Why Operators Filter Fragments and What It Implies", draft-taylor-v6ops-fragdrop-02 (work in progress), December 2013.

[I-D.wkumari-long-headers]

Kumari, W., Jaeggli, J., Bonica, R. P., and J. Linkova, "Operational Issues Associated With Long IPv6 Header Chains", draft-wkumari-long-headers-03 (work in progress), June 2015.

[IEPG94-Scudder]

Petersen, B. and J. Scudder, "Modern Router Architecture for Protocol Designers", IEPG 94. Yokohama, Japan. November 1, 2015, <<http://www.iepg.org/2015-11-01-ietf94/IEPG-RouterArchitecture-jgs.pdf>>.

[Jaeggli-2018]

Jaeggli, J., "IPv6 flow label: misuse in hashing", APNIC Blog, 2018, <<https://blog.apnic.net/2018/01/11/ipv6-flow-label-misuse-hashing/>>.

[Linkova-Gont-IEPG90]

Linkova, J. and F. Gont, "IPv6 Extension Headers in the Real World v2.0", IEPG 90. Toronto, ON, Canada. July 20, 2014, <<http://www.iepg.org/2014-07-20-ietf90/iepg-ietf90-ipv6-ehs-in-the-real-world-v2.0.pdf>>.

- [Microsoft-SA] Microsoft, "Windows TCP/IP Remote Code Execution Vulnerability (CVE-2021-24094)", February 2021, <<https://msrc.microsoft.com/update-guide/vulnerability/CVE-2021-24094>>.
- [nmap] Fyodor, "Dealing with IPv6 fragmentation in the DNS", Firewall/IDS Evasion and Spoofing, <<https://nmap.org/book/man-bypass-firewalls-ids.html>>.
- [PMTUD-Blackholes] De Boer, M. and J. Bosma, "Discovering Path MTU black holes on the Internet using RIPE Atlas", July 2012, <<http://www.nlnetlabs.nl/downloads/publications/pmtu-black-holes-msc-thesis.pdf>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<https://www.rfc-editor.org/info/rfc2460>>.
- [RFC5635] Kumari, W. and D. McPherson, "Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF)", RFC 5635, DOI 10.17487/RFC5635, August 2009, <<https://www.rfc-editor.org/info/rfc5635>>.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, DOI 10.17487/RFC6192, March 2011, <<https://www.rfc-editor.org/info/rfc6192>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.
- [RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <<https://www.rfc-editor.org/info/rfc6438>>.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, DOI 10.17487/RFC7045, December 2013, <<https://www.rfc-editor.org/info/rfc7045>>.
- [RFC7098] Carpenter, B., Jiang, S., and W. Tarreau, "Using the IPv6 Flow Label for Load Balancing in Server Farms", RFC 7098, DOI 10.17487/RFC7098, January 2014, <<https://www.rfc-editor.org/info/rfc7098>>.

- [RFC7113] Gont, F., "Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)", RFC 7113, DOI 10.17487/RFC7113, February 2014, <<https://www.rfc-editor.org/info/rfc7113>>.
- [RFC7739] Gont, F., "Security Implications of Predictable Fragment Identification Values", RFC 7739, DOI 10.17487/RFC7739, February 2016, <<https://www.rfc-editor.org/info/rfc7739>>.
- [RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <<https://www.rfc-editor.org/info/rfc7872>>.
- [RFC8900] Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O., and F. Gont, "IP Fragmentation Considered Fragile", BCP 230, RFC 8900, DOI 10.17487/RFC8900, September 2020, <<https://www.rfc-editor.org/info/rfc8900>>.
- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <<https://www.rfc-editor.org/info/rfc8955>>.
- [RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <<https://www.rfc-editor.org/info/rfc8956>>.
- [Zack-FW-Benchmark] Zack, E., "Firewall Security Assessment and Benchmarking IPv6 Firewall Load Tests", IPv6 Hackers Meeting #1, Berlin, Germany. June 30, 2013, <<https://www.ipv6hackers.org/files/meetings/ipv6-hackers-1/zack-ipv6hackers1-firewall-security-assessment-and-benchmarking.pdf>>.

Authors' Addresses

Fernando Gont
SI6 Networks
Segurola y Habana 4310, 7mo Piso
Villa Devoto, Ciudad Autonoma de Buenos Aires
Argentina

Email: fgont@si6networks.com
URI: <https://www.si6networks.com>

Nick Hilliard
INEX
4027 Kingswood Road
Dublin 24
IE

Email: nick@inex.ie

Gert Doering
SpaceNet AG
Joseph-Dollinger-Bogen 14
Muenchen D-80807
Germany

Email: gert@space.net

Warren Kumari
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: warren@kumari.net

Geoff Huston

Email: gih@apnic.net
URI: <http://www.apnic.net>

Will (Shucheng) Liu
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: liushucheng@huawei.com

6MAN Working Group
Internet-Draft
Intended status: Informational
Expires: 1 November 2022

G. Mishra
Verizon Inc.
A. Petrescu
CEA, LIST
N. Kottapalli
Benu Networks
D. Mudric
Ciena
D. Shytyi
SFR
30 April 2022

SLAAC with prefixes of arbitrary length in PIO (Variable SLAAC) - A
Problem Statement
draft-mishra-v6ops-variable-slaac-problem-stmt-04

Abstract

In the past, various IPv6 addressing models have been proposed based on a subnet hierarchy embedding a 64-bit prefix. The last remnant of IPv6 classful addressing is a inflexible interface identifier boundary at /64. This document details the 64-bit boundary problem statement.

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 <https://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 1 November 2022.

Copyright Notice

Copyright (c) 2022 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Terminology	2
2. Introduction	3
3. Problem Statement	3
4. Variable SLAAC Use Cases	6
4.1. Permission-less Extension of the Network	6
4.2. Private Networks	6
4.3. Mobile IPv6	7
4.4. Home and SOHO	7
4.5. 3GPP V2I and V2V networking	7
4.6. Smart Traffic Lights	8
4.7. 6lo	9
4.8. Large ISP's backbone POP	9
4.9. Permission-less extension of the Network	9
5. Security Considerations	10
6. IANA Considerations	10
7. Contributors	10
8. Acknowledgements	10
9. References	10
9.1. Normative References	10
9.2. Informative References	11
Appendix A. ChangeLog	12
Authors' Addresses	12

1. Terminology

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Introduction

From the beginning, the IPv6 addressing plan was based on a 128 bit address format made up of 8 hexets which were broken down into a 64 bit four hexet prefix and 64 bit four hexet interface identifier. For example, the address 2001:db8:3:4::1 has the first 4 hexets forming the /64 prefix 2001:db8:3:4::/64, whereas the last four hexets form an interface identifier abbreviated as ::1 (a 'hexet' is a group of max 4 hex digits between two columns, e.g. "2001" and "db8" are each a hexet). A comprehensive analysis of the 64-bit boundary is provided in [RFC7421]. The history of IPv6 Classful models proposed, and the last remnant of IPv6 Classful addressing rigid network interface identifier boundary at /64 is discussed in detail as well as the removal of the fixed position of the boundary for interface addressing in draft [I-D.bourbaki-6man-classless-ipv6].

This document discusses the reasons why the interface identifier has been fixed at 64 bits, and the problems that can be addressed by changing the GUA interface identifier from fixed 64 bit size to a variable interface identifier. This change would be consistent with static and DHCPv6 stateful IPv6 address assignment. This document tries to achieve clearing the confusion related to prefix length, and provide consistency of variable length prefix across the three IPv6 addressing strategies deployed, static, DHCPv6 and Stateless Address Autoconfiguration(SLAAC), and finally update all RFCs with the new variable SLAAC standard. The solution to this problem statement is defined in draft [I-D.mishra-6man-variable-slaac]

Over the years one of the merits of increasing the prefix length, and reducing the size of the interface identifier has been incorrectly stated as the possibility of IPv6 address space exhaustion could be circumvented, or that a 64 bit interface identifier is an efficient use of address space.

3. Problem Statement

This section details the problem statement as to what is broken today with fixed length Stateless Address Autoconfiguration SLAAC [RFC4862] and why it is critical to resolve this problem.

The main problem is that SLAAC RA or PD allocates a /64 by the wireless carrier 4G, 5G, 3GPP to mobile handset or hotspot, however segmentation of the /64 via SLAAC is required so that downstream interfaces can be further sub-netted. The use case section of this draft discusses this scenario as one of the use cases for shorter interface identifier, and this use case is the only one stated here in the problem statement as this is broken today with the current SLAAC specification [RFC4862], and there is not any workaround for this use case.

There are two reasons why this was not a problem in the past, but now with increased bandwidth there are more and more devices being piled onto a single handset or mobile hotspot. In the past generations of cellular systems (e.g. 2.5G aka GPRS and some 3G) the bandwidth available to the User Equipment was not enough to accommodate several applications; bandwidth available was roughly 256Kbit/s. For that reason, users were rarely tempted to use an UE to link other devices than that UE to the Internet. However, with the arrival of 3G, 3G+ (e.g. HSDPA / HSUPA), and even more so with 4G and 4G+, the bandwidth made available to UE increased significantly; this became an average effective of 1Mbit/s and even more. With this available bandwidth, the users are more and more tempted to connect several devices to the Internet. This operation is named 'connection sharing' or 'tethering'. Another answer to this question is that IPv6 technology that is widely used to 'tether' several IP devices to a smartphone is '64share' RFC7278. This technology is used for smartphones but is not so in vehicles. One of the reasons of not being used in vehicles is the lack of scalability: a /64 prefix is shared between the UE ptp link and the subnet (typically Wi-Fi), but can not be further sub-netted to other subnets in the car.

The reason why all devices in a car cannot remain on a single /64 are as follows. These devices have different link-layer technologies, and not all WiFi could be bridged into Ethernet such as to keep all devices into one /64. They could be on links that are not bridgeable: devices on 802.11-OCB cannot be bridged, devices on Bluetooth cannot be bridged, devices on 3GPP cannot be bridged, and so on. Other than the impossibility to bridge several such link-layer technologies there is also a problem of noise: in a vehicle one wants the braking pedal signal to not be disturbed by entertainment sites such as YouTube. That physical technical requirement separation of different link layer technologies segmentation on to different smaller IPv6 subnets cannot be achieved if all devices are on a single /64, or bridged. Therefore, the only possible solution to connect these disparate devices onto a 3GPP network for internet access is to keep these separate link layer technologies segmented onto separate greater than /64 prefix subnets and breaking the /64 boundary that exists today with a Variable SLAAC solution. Thus,

when the 3GPP network gives a /64 to the car, and when there are unbridgeable technologies in the car (e.g. WiFi cant be bridged to Bluetooth), then the only possibility is to divide that /64 into two /65s. One /65 would be used on the WiFi and another /65 would be used on Bluetooth. But in order for SLAAC to work with /65 then there is a need to have the shorter interface identifier of length 63. Hence the need of lengths of PIOs other than 64 (variable plen).

There are three scenarios that require SLAAC to be able to be routed between two greater than /64 prefix segments as part of the requirement for variable length SLAAC and what is broken with the current SLAAC specification defined in [RFC4862].

The first scenario is within a car using car manufacturer single SIM for internet access and being able to bridge(Route) other link layer devices like BT via variable slaac. In this scenario the communication between downstream devices are all located within the car using the car manufacturer built in SIM card for in-vehicle communication. The in-vehicle scenario covers both the built-in car manufacturer SIM card scenario, or if the car manufacturer does not support built-in SIM card then a single mobile handset providing 3GPP internet access to all devices in the car.

The second scenario is V2V (vehicle to vehicle) between cars requiring SLAAC to subnet the >64 prefix so that the two cars have WiFi connectivity.

This third scenario is a uCPE(Universal Customer Premises Equipment) device is LTE 4G and Wi-Fi capable, and utilizes NFV (Network Function Virtualization) framework, providing SFC (Service Function Chaining), where one VNF (Virtual Network Function) is a CPE Layer 3 router and is the uCPE device which will receive a /64 prefix from 4G 3GPP Wireless provider and would like to be able to provide further segmentation. In order to provide further segmentation and subdivide the /64 into smaller longer prefix subnets variable SLAAC must be employed. In this example we would give 1st /66 to Wi-Fi users, 2nd /66 to Wired connected network device without security, 3rd /66 prefix to VNF firewall instance, and 4th /66 prefix VNF load balancer instance. The uCPE (Universal Customer Premises Equipment) defined in draft [I-D.shytyi-opsawg-vysm].

From a segmented bandwidth perspective while breaking up the /64 subnet into smaller subnets, there is not any impact to the user experience of the now shared bandwidth, as long as the cellular signal has adequate enough bars as far as signal strength to accommodate the now multiple devices sharing the single cellular signal. These scenarios described above are the problems that can only be solved with a variable SLAAC solution. There is no other solution or workaround for this problem. A possible solution to this problem has been proposed in [I-D.mishra-6man-variable-slaac].

4. Variable SLAAC Use Cases

This section describes real world use cases of variable slaac that cannot be done today and with fixed 64 bit prefix lengths.

4.1. Permission-less Extension of the Network

Permission-less extensions of the network with new links (and by implication with new routers) are not supported.

The lack of possibility to realize a permission-less extension of the network is an important problem, which appears at the edge of the network. The permission is 'granted' for end users situated at the edge of the network, and is 'granted' by advertising a prefix of length 64 inside the PIO option in a RA typically. The end user receives this prefix, forms an address, and is able to connect to the internet. However, the end user has no permission to further extend the network. Although the device is able to form subsequent prefixes of a length of, say 65, and further advertise it down in the extension of the network, no other Host in that extension of the network is able to use that advertisement; a Host cannot form an address with a prefix length 65 by using SLAAC. The Linux error text reported in the kernel log upon reception of a plen 65 is "illegal" (or similar).

4.2. Private Networks

Private networks such as Service Provider core not accessible by customers and enterprises where all hosts are trusted are the primary use case for variable SLAAC as the shorter interface identifier does not create any security issues with not having a longer 64 bit interface identifier for privacy extensions stable interface identifier [RFC8084] due to all hosts being inherently trusted. Private internal networks such as corporate intranets traditionally have always used static IPv6 addressing for infrastructure. This manual IPv6 address assignment process for network infrastructure links can take long lead times to complete deployment. By changing the behavior of SLAAC to support variable length prefix and interface

identifier allows SLAAC to be used programmatically to deploy to large scale IPv6 networks with thousands of point-to-point links. Note that network infrastructure technically does not require IPv6 addressing due to IPv6 next hop being a link local address for IGP routing protocols such as OSPF and ISIS as well as the link local address can be the peer IPv6 address for exterior gateway routing protocols such as BGP. However for hop by hop ping and traceroute capability to have IPv6 reachability at each hop for troubleshooting jitter, latency and drops it is an IPv6 recommended best practice to configure IPv6 address on all infrastructure interfaces.

4.3. Mobile IPv6

Old MIP6 (Mobile IPv6) Working Group and old Nemo Working Group's routing solution scenarios related to Mobile IPv6 ([RFC3775]) (note: nowadays most MIP-related activity is in DMM WG) where the mobile endpoint can now obtain from the home agent variable SLAAC address and not 64 bit prefix /64 address. This maybe useful in cases where a /64 can now be managed from an addressing perspective and subdivided into blocks for manageability of MIP6 endpoints instead of allocating a single /64 per endpoint.

4.4. Home and SOHO

Home and SOHO (Small Office and Home Office) environments where internet access uses a broadband service provider single or dual homed scenario. In those such Home networking Homenet environments where HNCP (Home Network Control Protocol [RFC7788] SADR (Source Address Dependent Routing) are deployed for automatic configuration for LAN Wi-Fi endpoint subnets can also now take advantage of variable length SLAAC in deployment scenarios. In cases where multiple routers are deployed in a home environment where routing prefix reachability needs to be advertised where Babel [RFC6126] routing protocol is utilized in those cases variable SLAAC can also be utilized to break up a /64 into multiple smaller subnets.

4.5. 3GPP V2I and V2V networking

In V2I networking (with 3GPP or with IEEE 802.11bd) the IP-OBUE in the vehicle receives a /64 prefix from the cellular network (or from a IP-RSU - Road-Side Unit). This /64 prefix can be used to form one address for the egress interface of the Mobile Router (MR, which is also termed 'IP-OBUE', for IP On-Board Unit, in IPWAVE WG documents such as RFC8691), but can not be used to form IP addresses for other hosts in the vehicle. In the following two paragraphs we explain this problem.

In certain 3GPP V2I networking use cases a /56 is allocated by the 3GPP infrastructure to the 4G modem of the IP-OBU in the vehicle. In such use case it is possible that the IP-OBU sub-divides the allocated /56 into multiple 'result' /64 prefixes. Such a 'result' /64 prefix could be used to form addresses for deeper subnets in the vehicle, by employing existing SLAAC and existing IPv6-over-foo specifications of Interface ID.

If in other 3GPP V2I networking use-cases the infrastructure does not allocate a /56 (or 'longer' prefix lengths such as a /57, /58../63) to the IP-OBU, i.e. a /64 is allocated to the IP-OBU, then the 'result' prefix obtained after a sub-divide operation can only be of length /65, or /66, or longer. A prefix of such length (longer than 64) can not be used with SLAAC and existing IP-over-foo Interface Identifiers, because the length of all Interface Identifiers in all IPv6-over-foo documents must always be 64, and the length of the IPv6 is always 128bit. The 64bit of an IID added to the 65bit (or more) of a prefix is larger than 128bit. It is for this reason that a SLAAC with other than 64bit Interface IDs (hence a 'Variable Prefix Length SLAAC') is needed.

The problem of /64 allocation to the vehicle is mostly present in V2I use-cases. In V2V use-cases this problem is less apparent but deserves consideration. Until now there was no clearcut design and decision about the infrastructure allocating addresses to several vehicles (just to one, in V2I, see above). In some use-cases, the prefix allocated to one vehicle could be further extended by that vehicle to delegate prefixes to other vehicles nearby which might not have 3GPP connections, but only 802.11-OCB interfaces. In such cases it is again necessary that a /64 allocated by the infrastructure to the first vehicle be further sub-divided in multiple 'result' longer-than-/64 prefixes; and one of these longer-than-64 prefixes might be used for the second vehicle (instead of being used for the internal subnets of the first vehicle); this latter vehicle will need to use a form of SLAAC and IP-over-foo that are not limited by the /64 limit.

4.6. Smart Traffic Lights

Smart traffic lights are traffic lights equipped with a communication system. Smart traffic lights are deployed at intersections of roads and serve the purpose of safely arbitrating the passage of automobiles, pedestrians and cyclists. A typical smart traffic lights setting is made of several computers, included but not limited to: a traffic lights controller, a power controller and a communication gateway. More advanced smart traffic lights are equipped with more computers for radars, detection loops, lidars, V2X wireless capabilities, Wi-Fi, Bluetooth and cellular 4G or 5G. All these computers need to use IP addresses: at least one IP address per

computer. Since smart traffic lights are deployed in areas where Internet might not be available by cable, fibre or other Wireless MAN technology the only way to connect all computers in the smart traffic lights setting is to employ a 4G (or 5G) gateway. This gateway obtains typically a /64 prefix from the network operator; there is a problem in subdividing that /64 prefix into smaller prefixes, because the obtained prefixes can not be used by SLAAC, because SLAAC uses Interface IDs of length 64 in practice. Even if the SLAAC specification is independent of the prefix length, the length of the Interface ID dictates the prefix length by side effect (128 minus IID length imposes the prefix length). SLAAC might work with a plen 65 by specification, but all IIDs in all IPv6-over-foo request that IIDs be 64; and the sum of IID len plus plen must be 128.

4.7. 6lo

6lo Working IPv6 over Network Constrained nodes working group use cases. Use cases for IoT devices where have limited network access requirements could now take advantage of variable SLAAC longer prefixes lengths /65-/128.

4.8. Large ISP's backbone POP

Large ISP backbone POPs such as IXPs where many carriers share the same backbone and ND cache exhaustion may occur due to /64 subnet size. One mitigation technique employed is the use of an ARP Sponge for IPv4 or Layer 2 multicast rate limiters for IPv6. In those particular cases a longer prefix static or variable SLAAC subnet could be utilized to reduce the maximum number of hosts on the subnet.

4.9. Permission-less extension of the Network

When one wants to extend the network, one typically wants to add new computers to it. Currently, there are two ways to achieve it: (1) ask the network administrator to provide addresses while also inserting a route towards the new subnet of devices and (2) use NAT. With IPv6, NAT is not desirable. In order to extend the network without asking for permission one needs to obtain addresses and to obtain that route inserted. In order to obtain addresses, one might take advantage of the /64 prefix typically advertised by the network to an edge of it. To do that, one needs to sub-divide the /64 prefix into /65 sub-prefixes (or longer, such as /66, /67, etc.) which could be further advertised in the extension of the network. For the action of inserting a route, the particular topic is outside the scope of this document.

5. Security Considerations

The administrator should be aware to maintain 64 bit interface identifier for privacy when connected directly to the internet so that entropy for optimal heuristics are maintained for security.

Variable length interface identifier shorter than 64 bits should be only used within corporate intranets and private networks where all hosts are trusted.

In all cases where the host is on a public network for privacy concerns to avoid traceability variable interface identifier MUST never be utilized.

6. IANA Considerations

IANA is requested to assign the new Router Advertisement flag defined in Section 5 of this document. Bit 6 is the next available bit in this registry, IANA is requested to use this bit unless there is a reason to use another bit in this registry.

IANA is also requested to register this new flag bit in the IANA IPv6 ND Router Advertisement flags Registry [IANA-RF].

7. Contributors

Contributors.

8. Acknowledgements

9. References

9.1. Normative References

[I-D.bourbaki-6man-classless-ipv6]
Bourbaki, N., "IPv6 is Classless", Work in Progress, Internet-Draft, draft-bourbaki-6man-classless-ipv6-06, 18 April 2022, <<https://www.ietf.org/internet-drafts/draft-bourbaki-6man-classless-ipv6-06.txt>>.

[I-D.mishra-6man-variable-slaac]
Mishra, G., Petrescu, A., Kottapalli, N., Mudric, D., and D. Shytyi, "SLAAC with prefixes of arbitrary length in PIO (Variable SLAAC)", Work in Progress, Internet-Draft, draft-mishra-6man-variable-slaac-06, 30 April 2022, <<https://www.ietf.org/archive/id/draft-mishra-6man-variable-slaac-06.txt>>.

- [I-D.shytyi-opsawg-vysm]
Shytyi, D., Beylier, L., and L. Iannone, "A YANG Module for uCPE management.", Work in Progress, Internet-Draft, draft-shytyi-opsawg-vysm-10, 9 September 2021, <<https://www.ietf.org/archive/id/draft-shytyi-opsawg-vysm-10.txt>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, DOI 10.17487/RFC3775, June 2004, <<https://www.rfc-editor.org/info/rfc3775>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC6126] Chroboczek, J., "The Babel Routing Protocol", RFC 6126, DOI 10.17487/RFC6126, April 2011, <<https://www.rfc-editor.org/info/rfc6126>>.
- [RFC7788] Stenberg, M., Barth, S., and P. Pfister, "Home Networking Control Protocol", RFC 7788, DOI 10.17487/RFC7788, April 2016, <<https://www.rfc-editor.org/info/rfc7788>>.
- [RFC8084] Fairhurst, G., "Network Transport Circuit Breakers", BCP 208, RFC 8084, DOI 10.17487/RFC8084, March 2017, <<https://www.rfc-editor.org/info/rfc8084>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

9.2. Informative References

- [RFC7421] Carpenter, B., Ed., Chown, T., Gont, F., Jiang, S., Petrescu, A., and A. Yourtchenko, "Analysis of the 64-bit Boundary in IPv6 Addressing", RFC 7421, DOI 10.17487/RFC7421, January 2015, <<https://www.rfc-editor.org/info/rfc7421>>.

Appendix A. ChangeLog

The changes are listed in reverse chronological order, most recent changes appearing at the top of the list.

-00: initial version.

Authors' Addresses

Gyan Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com

Alexandre Petrescu
CEA, LIST
CEA Saclay
91190 Gif-sur-Yvette
France
Phone: +33169089223
Email: Alexandre.Petrescu@cea.fr

Naveen Kottapalli
Benu Networks
300 Concord Road
Billerica, MA 01821
United States of America
Phone: +1 978 223 4700
Email: nkottapalli@benu.net

Dusan Mudric
Ciena
Canada
Phone: +1-613-670-2425
Email: dmudric@ciena.com

Dmytro Shytyi
SFR
Paris
France
Email: dmytro.shytyi@sfr.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 24 February 2022

S. Peng
Z. Li
Huawei Technologies
C. Xie
China Telecom
Z. Qin
China Unicom
G. Mishra
Verizon Inc.
23 August 2021

Processing of the Hop-by-Hop Options Header
draft-peng-v6ops-hbh-06

Abstract

This document describes the processing of the Hop-by-Hop Options Header (HBH) in today's routers in the aspects of standards specification, common implementations, and default operations. This document outlines the reasons why the Hop-by-Hop Options Header is rarely utilized in current networks. In addition, this document describes how the HBH could be used as a powerful mechanism allowing deployment and operations of new services requiring a more optimized way to leverage network resources of an infrastructure. The Hop-by-Hop Options Header is taken into consideration by several network operators as a valuable container for carrying the information facilitating the introduction of new services. The processing requirements of the HBH and the migration strategies are also suggested.

Requirements Language

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] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 <https://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 24 February 2022.

Copyright Notice

Copyright (c) 2021 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 (<https://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. Modern Router Architecture	4
3. Specification of RFC 8200	7
4. Common Implementations	8
4.1. Historical Reasons	9
4.2. Consequences	9
5. Operators' Typical Processing	9
6. New Services	10
7. Requirements	10
8. Migration Strategies	11
9. Security Considerations	12
10. IANA Considerations	12
11. Acknowledgements	12
12. References	12
12.1. Normative References	12
12.2. Informative References	13
Authors' Addresses	14

1. Introduction

Due to historical reasons, such as incapable ASICs, limited IPv6 deployments, and few service requirements, the most common Hop-by-Hop Options header (HBH) processing implementation is that the node sends the IPv6 packets with the Hop-by-Hop Options header to the control plane of the node. The option type of each option carried within the Hop-by-Hop Options header will not even be examined before the packet is sent to the control plane. Very often, such processing behavior is the default configuration or, even worse, is the only behavior of the ipv6 implementation of the node.

Such default processing behavior of the Hop-by-Hop Options header could result in various unpleasant effects such as a risk of Denial of Service (DoS) attack on the router control plane and inconsistent packet drops due to rate limiting on the interface between the router control plane and forwarding plane, which will impact the normal end-to-end IP forwarding of the network services.

This actually introduced a circular problem:

-> An implementation problem caused HBH to become a DoS vector.

-> Because HBH is a DoS vector, network operators deployed ACLs that discard packets containing HBH.

-> Because network operators deployed ACLs that discard packets containing HBH, network designers stopped defining new HBH Options.

-> Because network designers stopped defining new HBH Options, the community was not motivated to fix the implementation problem that cause HBH to become a DoS vector.

The purpose of this draft is to break the cycle described above, fixing the problem that caused HBH not actually being utilized in operators' networks so to allow a better leverage of the HBH capability.

Driven by the wide deployments of IPv6 and ever-emerging new services, the Hop-by-Hop Options Header is taken as a valuable container for carrying the information to facilitate these new services.

This document suggests the desired processing behavior and the migration strategies towards it.

2. Modern Router Architecture

Modern router architecture design maintains a strict separation of the router control plane and its forwarding plane [RFC6192], as shown in Figure 1. Either the control plane or the forwarding plane is composed of both software and hardware, but each plane is responsible for different functions.




Figure 1. Modern Router Architecture

The router control plane supports routing and management functions, handling packets destined to the device as well as building and sending packets originated locally on the device, and also drives the programming of the forwarding plane. The router control plane is generally realized in software on general-purpose processors, and its hardware is usually not optimized for high-speed packet handling. Because of the wide range of functionality, it is more susceptible to security vulnerabilities and a more likely a target for a DoS attack.

The forwarding plane is typically responsible for receiving a packet on an incoming interface, performing a lookup to identify the packet's next hop and determine the outgoing interface towards the destination, and forwarding the packet out through the appropriate outgoing interface. Typically, forwarding plane functionality is realized in high-performance Application Specific Integrated Circuits (ASICs) or Network Processors (NPs) that are capable of handling very high packet rates.

The router control plane interfaces with its forwarding plane through the Interface Z, as shown in the Figure 1, and the forwarding plane connects to other network devices via Interfaces such as X and Y. Since the router control plane is vulnerable to the DoS attack, usually a traffic filtering mechanism is implemented on Interface Z in order to block unwanted traffic. In order to protect the router control plane, a rate-limiting mechanism is always implemented on this interface. However, such rate limiting mechanism will always cause inconsistent packet drops, which will impact the normal IP forwarding.

Semiconductor chip technology has advanced significantly in the last decade, and as such the widely used network processing and forwarding process can now not only forward packets at line speed, but also easily support other feature processing such as QoS for DiffServ/MPLS, Access Control List (ACL), Firewall, and Deep Packet Inspection (DPI).

A Network Processing Unit (NPU) is a non-ASIC based Integrated Circuit (IC) that is programmable through software. It performs all packet header operations between the physical layer interface and the switching fabric such as packet parsing and forwarding, modification, and forwarding. Many equipment vendors implement these functions in fixed function ASICs rather than using "off-the-shelf" NPUs, because of proprietary algorithms.

Classification Co-processor is a specialized processor that can be used to lighten the processing load on an NPU by handling the parsing and classification of incoming packets such as IPv6 extended header HBH options processing. This advancement enables network processors to do the general process to handle simple control messages for traffic management, such as signaling for hardware programming, congestion state report, OAM, etc. Industry trend is for intelligent multi-core CPU hardware using modern NPUs for forwarding packets at line rate while still being able to perform other complex tasks such as HBH forwarding options processing without having to punt to the control plane.

Many of the packet-processing devices employed in modern switch and router designs are fixed-function ASICs to handle proprietary functions. While these devices can be very efficient for the set of functions they are designed for, they can be very inflexible. There is a tradeoff of price, performance and flexibility when vendors make a choice to use a fixed function ASIC as opposed to NPU. Due to the inflexibility of the fixed function ASIC, tasks that require additional processing such as IPv6 HBH header processing must be punted to the control plane. This problem is still a challenge today and is the reason why operators to protect against control plane DOS

attack vector must drop or ignore HBH options. As industry shifts to Merchant Silicon based NPU evolution from fixed function ASIC, the gap will continue to close increasing the viability ubiquitous HBH use cases due to now processing in the forwarding plane.

Most modern routers maintain a strict separation between forwarding plane and control plane hardware. Forwarding plane bandwidth and resources are plentiful, while control plane bandwidth and resources are constrained. In order to protect scarce control plane resources, routers enforce policies that restrict access from the forwarding plane to the control plane. Effective policies address packets containing the HBH Options Extension header, because HBH control options require access from the forwarding plane to the control plane. Many network operators perceive HBH Options to be a breach of the separation between the forwarding and control planes. In this case HBH control options would be required to be punted to control plane by fixed function ASICs as well as NPUs.

The maximum length of an HBH Options header is 2,048 bytes. A source node can encode hundreds of options in 2,048 bytes. With today's technology it would be cost prohibitive to be able to process hundreds of options with either NPU or proprietary fixed function ASIC.

While [RFC8200] required that all nodes must examine and process the Hop-by-Hop Options header, it is now expected that nodes along a packet's delivery path only examine and process the Hop-by-Hop Options header if explicitly configured to do so. This can be beneficial in cases where transit nodes are legacy hardware and the destination endpoint PE is newer NPU based hardware that can process HBH in the forwarding plane.

IPv6 Extended Header limitations that need to be addressed to make HBH processing more efficient and viable in the forwarding plane:

[RFC8504] defines the IPv6 node requirements and how to protect a node from excessive header chain and excessive header options with various limitations that can be defined on a node. [RFC8883] defines ICMPv6 Errors for discarding packets due to processing limits. Per [RFC8200] HBH options must be processed serially. However, an implementation of options processing can be made to be done with more parallelism in serial processing grouping of similar options to be processed in parallel.

The IPv6 standard does not currently limit the header chain length or number of options that can be encoded.

Each Option is encoded in a TLV and so processing of a long list of TLVs is expensive. Zero data length encoded options TLVs are a valid option. A DOS vector could be easily generated by encoding 1000 HBH options (Zero data length) in a standard 1500 MTU packet. So now imagine if you have a Christmas tree long header chain to parse each with many options.

3. Specification of RFC 8200

[RFC8200] defines several IPv6 extension header types, including the Hop-by-Hop (HBH) Options header. As specified in [RFC8200], the Hop-by-Hop (HBH) Options header is used to carry optional information that will be examined and processed by every node along a packet's delivery path, and it is identified by a Next Header value of zero in the IPv6 header.

The Hop-by-Hop (HBH) Options header contains the following fields:

-- Next Header: 8-bit selector, identifies the type of header immediately following the Hop-by-Hop Options header.

-- Hdr Ext Len: 8-bit unsigned integer, the length of the Hop-by-Hop Options header in 8-octet units, not including the first 8 octets.

-- Options: Variable-length field, of length such that the complete Hop-by-Hop Options header is an integer multiple of 8 octets long.

The Hop-by-Hop (HBH) Options header carries a variable number of "options" that are encoded in the format of type-length-value (TLV).

The highest-order two bits (i.e., the ACT bits) of the Option Type specify the action that must be taken if the processing IPv6 node does not recognize the Option Type. The third-highest-order bit (i.e., the CHG bit) of the Option Type specifies whether or not the Option Data of that option can change en route to the packet's final destination.

While [RFC2460] required that all nodes must examine and process the Hop-by-Hop Options header, with [RFC8200] it is expected that nodes along a packet's delivery path only examine and process the Hop-by-Hop Options header if explicitly configured to do so. It means that the HBH processing behavior in a node depends on its configuration.

However, in the current [RFC8200], there is no explicit specification of the possible configurations. Therefore, the nodes may be configured to ignore the Hop-by-Hop Options header, drop packets containing a Hop-by-Hop Options header, or assign packets containing a Hop-by-Hop Options header to the control plane [RFC8200]. Because of these likely uncertain processing behaviors, new hop-by-hop options are not recommended.

4. Common Implementations

In the current common implementations, once an IPv6 packet, with its Next Header field set to 0, arrives at a node, it will be directly sent to the control plane of the node. With such implementations, the value of the Next Header field in the IPv6 header is the only trigger for the default processing behavior. The option type of each option carried within the Hop-by-Hop Options header will not even be examined before the packet is sent to the control plane.

Very often, such processing behavior is the default configuration on the node, which is embedded in the implementation and cannot be changed or reconfigured.

Another critical component of IPv6 HBH processing, in some cases overlooked, is the operator core network which can be designed to use the global Internet routing table for internet traffic and in other cases use an overlay MPLS VPN to carry Internet traffic.

In the global Internet routing table scenario where only an underlay global routing table exists, and no VPN overlay carrying customer Internet traffic, the IPv6 HBH options can be used as a DOS attack vector for both the operator nodes, adjacent inter-as peer nodes as well as customer nodes along a path.

In a case where the Internet routing table is carried in a MPLS VPN overlay payload, the HBH options header does not impact the operator underlay framework and only impacts the VPN overlay payload and thus the operator underlay topmost label global table routing FEC LSP instantiation is not impacted as the operator underlay is within the operators closed domain.

However, HBH options DOS attack vector in the VPN overlay can still impact the customer CE destination end nodes as well as other adjacent inter-as operators that only use underlay global Internet routing table. In an operator closed domain where MPLS VPN overlay is utilized to carry internet traffic, the operator has full control of the underlay and IPv6 Extended header chain length as well as the number of HBH options encoded.

In the global routing table scenario for Internet traffic there is no way to control the IPv6 Extended header chain length as well as the number of HBH options encoded.

4.1. Historical Reasons

When IPv6 was first implemented on high-speed routers, HBH options were not yet well-understood and ASICs were not as capable as they are today. So, early IPv6 implementations dispatched all packets that contain HBH options to their control plane.

4.2. Consequences

Such implementation introduces a risk of a DoS attack on the control plane of the node, and a large flow of IPv6 packets could congest the control plane, causing other critical functions (including routing and network management) that are executed on the control plane to fail. Rate limiting mechanisms will cause inconsistent packet drops and impact the normal end-to-end IP forwarding of the network services.

5. Operators' Typical Processing

To mitigate this DoS vulnerability, many operators deployed Access Control Lists (ACLs) that discard all packets containing HBH Options.

[RFC6564] shows the Reports from the field indicating that some IP routers deployed within the global Internet are configured either to ignore or to drop packets having a hop-by-hop header. As stated in [RFC7872], many network operators perceive HBH Options to be a breach of the separation between the forwarding and control planes. Therefore, several network operators configured their nodes so as to discard all packets containing the HBH Options Extension Header, while others configured nodes to forward the packet but to ignore the HBH Options. [RFC7045] also states that hop-by-hop options are not handled by many high-speed routers or are processed only on a control plane.

Due to such behaviors observed and described in these specifications, new hop-by-hop options are not recommended in [RFC8200] hence the usability of HBH options is severely limited.

6. New Services

As IPv6 is being rapidly and widely deployed worldwide, more and more applications and network services are migrating to or directly adopting IPv6. More and more new services that require HBH are emerging and the HBH Options header is going to be utilized by the new services in various scenarios.

In-situ OAM (IOAM) with IPv6 encapsulation [I-D.ietf-ippm-ioam-ipv6-options] is one of the examples. IOAM in IPv6 is used to enhance diagnostics of IPv6 networks and complements other mechanisms, such as the IPv6 Performance and Diagnostic Metrics Destination Option described in [RFC8250]. The IOAM data fields are encapsulated in "option data" fields of the Hop-by-Hop Options header if Pre-allocated Tracing Option, Incremental Tracing Option, or Proof of Transit Option are carried [I-D.ietf-ippm-ioam-data], that is, the IOAM performs per hop.

Alternate Marking Method can be used as the passive performance measurement tool in an IPv6 domain. The AltMark Option is defined as a new IPv6 extension header option to encode alternate marking technique and Hop-by-Hop Options Header is considered [I-D.ietf-6man-ipv6-alt-mark].

The Minimum Path MTU Hop-by-Hop Option is defined in [I-D.ietf-6man-mtu-option] to record the minimum Path MTU along the forward path between a source host to a destination host. This Hop-by-Hop option is intended to be used in environments like Data Centers and on paths between Data Centers as well as other environments including the general Internet. It provides a useful tool for allowing to better take advantage of paths able to support a large Path MTU.

As more services start utilizing the HBH Options header, more packets containing HBH Options are going to be injected into the networks. According to the current common configuration in most network deployments, all the packets of the new services are going to be sent to the control plane of the nodes, with the possible consequence of causing a DoS on the control plane. The packets will be dropped and the normal IP forwarding may be severely impacted. The deployment of new network services involving multi-vendor interoperability will become impossible.

7. Requirements

- * The HBH options header MUST NOT become a possible DDoS Vector.

- * HBH options SHOULD be designed so that they don't reduce the probability of packet delivery. For example, an intermediate node may discard a packet because it contains more HBH options than the node can process.
- * HBH processing MUST be efficient. That is, it MUST be possible to produce implementations that perform well at a reasonable cost.
- * The Router Alert Option MUST NOT impact the processing of other HBH options that should be processed more quickly.
- * HBH Options MAY influence how a packet is forwarded. However, with the exception of the Router Alert Option, an HBH Option MUST NOT cause control plane state to be created, modified or destroyed on the processing node. As per [RFC6398], protocol developers SHOULD avoid future use of the Router Alert Option.
- * More requirements are to be added.

8. Migration Strategies

In order to achieve the desired processing behavior of the HBH options header and facilitate the ever-emerging new services to be deployed in operators' networks across multiple vendors' devices, the migration can happen in three parts as described below:

1. The source of the HBH options header encapsulation.

The information to be carried in the HBH options header needs to be first categorized and encapsulated into either control options or forwarding options, and then encapsulated in different packets.

2. The nodes within the network.

The nodes within the network are updated to the proposed behavior introduced in the previous section.

3. The edge nodes of the network.

The edge nodes should check whether the packet contains an HBH header with control or forwarding option. Packets with a control option may still be filtered and dropped while packets with forwarding option SHOULD be allowed by the ACL.

If it is certain that there is no harm that can be introduced by the HBH control options to the nodes and the services, they can also be allowed.

Note: During the migration stage, the nodes that are not yet updated will stay with their existing configurations.

9. Security Considerations

The same as the Security Considerations apply as in [RFC8200] for the part related with the HBH Options header.

10. IANA Considerations

This document does not include an IANA request.

11. Acknowledgements

The authors would like to acknowledge Ron Bonica, Fred Baker, Bob Hinden, Stefano Previdi, and Donald Eastlake for their valuable review and comments.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <<https://www.rfc-editor.org/info/rfc2460>>.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, DOI 10.17487/RFC6192, March 2011, <<https://www.rfc-editor.org/info/rfc6192>>.
- [RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <<https://www.rfc-editor.org/info/rfc6398>>.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, DOI 10.17487/RFC7045, December 2013, <<https://www.rfc-editor.org/info/rfc7045>>.

- [RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <<https://www.rfc-editor.org/info/rfc7872>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

12.2. Informative References

- [I-D.ietf-6man-ipv6-alt-mark]
Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate Marking Method", Work in Progress, Internet-Draft, draft-ietf-6man-ipv6-alt-mark-08, 26 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-6man-ipv6-alt-mark-08.txt>>.
- [I-D.ietf-6man-mtu-option]
Hinden, R. M. and G. Fairhurst, "IPv6 Minimum Path MTU Hop-by-Hop Option", Work in Progress, Internet-Draft, draft-ietf-6man-mtu-option-06, 7 August 2021, <<https://www.ietf.org/archive/id/draft-ietf-6man-mtu-option-06.txt>>.
- [I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-data-14, 24 June 2021, <<https://www.ietf.org/archive/id/draft-ietf-ippm-ioam-data-14.txt>>.
- [I-D.ietf-ippm-ioam-ipv6-options]
Bhandari, S. and F. Brockners, "In-situ OAM IPv6 Options", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-ipv6-options-06, 31 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-ippm-ioam-ipv6-options-06.txt>>.

- [RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", RFC 2711, DOI 10.17487/RFC2711, October 1999, <<https://www.rfc-editor.org/info/rfc2711>>.
- [RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", RFC 8250, DOI 10.17487/RFC8250, September 2017, <<https://www.rfc-editor.org/info/rfc8250>>.
- [RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <<https://www.rfc-editor.org/info/rfc8504>>.
- [RFC8883] Herbert, T., "ICMPv6 Errors for Discarding Packets Due to Processing Limits", RFC 8883, DOI 10.17487/RFC8883, September 2020, <<https://www.rfc-editor.org/info/rfc8883>>.

Authors' Addresses

Shuping Peng
Huawei Technologies
Beijing
China

Email: pengshuping@huawei.com

Zhenbin Li
Huawei Technologies
Beijing
China

Email: lizhenbin@huawei.com

Chongfeng Xie
China Telecom
China

Email: xiechf@chinatelecom.cn

Zhuangzhuang Qin
China Unicom
Beijing
China

Email: qinzhuangzhuang@chinaunicom.cn

Gyan Mishra
Verizon Inc.
United States of America

Email: gyan.s.mishra@verizon.com

V6OPS
Internet-Draft
Intended status: Informational
Expires: August 26, 2021

G. Fioccola
P. Volpato
Huawei Technologies
N. Elkins
Inside Products
S. Lourdez
Post Luxembourg
February 22, 2021

IPv6 Deployment Status
draft-vf-v6ops-ipv6-deployment-02

Abstract

Looking globally, IPv6 is growing faster than IPv4 and this means that the collective wisdom of the networking industry has selected IPv6 for the future. This document provides an overview of IPv6 transition deployment status and a view on how the transition to IPv6 is progressing among network operators and enterprises that are introducing IPv6 or have already adopted an IPv6-only solution. It also aims to analyze the transition challenges and therefore encourage actions and more investigations on some areas that are still under discussion. The overall IPv6 incentives are also examined.

Requirements Language

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 RFC 2119 [RFC2119].

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 <https://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 26, 2021.

Copyright Notice

Copyright (c) 2021 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 (<https://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. IPv4 Address Exhaustion	4
3. The global picture of IPv6	5
3.1. IPv6 users	5
3.2. IPv6 allocations and networks	5
4. Survey among Network Operators	7
5. Considerations for Enterprises	8
6. Observations on Content and Cloud Service Providers	8
7. Industrial Internet application	8
8. IPv6 deployments worldwide	9
8.1. IPv6 service design for Mobile, Fixed broadband and enterprises	9
8.1.1. IPv6 introduction	9
8.1.2. IPv6-only service delivery	10
9. Findings of the IPv6 Survey	11
10. IPv6 incentives	12
11. Call for action	13
11.1. Transition choices	13
11.1.1. Service providers	13
11.1.2. Enterprises	14
11.1.3. Cloud and Data Centers	16
11.1.4. Industrial Internet	16
11.1.5. Government and Regulators	17
11.2. Network Operations	17
11.3. Performance	18
11.3.1. IPv6 latency	18
11.3.2. IPv6 packet loss	18
11.3.3. Router's performance	19
11.4. IPv6 security	19
11.4.1. Protocols security issues	20
11.4.2. IPv6 Extension Headers and Fragmentation	21

11.4.3. Oversized IPv6 packets	21
12. Security Considerations	22
13. Contributors	22
14. Acknowledgements	22
15. IANA Considerations	22
16. References	22
16.1. Normative References	22
16.2. Informative References	22
Appendix A. Summary of Questionnaire and Replies	26
Authors' Addresses	30

1. Introduction

The focus of this document is to provide a survey of the deployed IPv6 transition technologies and to highlight the difficulties in the transition. This process helps to understand what is missing and how to improve the current IPv6 deployment strategies of network operators, enterprises, content and cloud service providers. The objective is to give an updated view of the practices and plans already described in [RFC6036]. The scope is to report the current IPv6 status and encourage actions and more investigations on some areas that are still under discussion as well as the main incentives for the IPv6 adoption.

[RFC6180] discussed the IPv6 deployment models and migration tools. [RFC6036] described the Service Provider Scenarios for IPv6 Deployment, [RFC7381] introduced the guidelines of the IPv6 deployment for Enterprise and [RFC6883] provided guidance and suggestions for Internet Content Providers and Application Service Providers. On the other hand, this document focuses on the end-to-end services and in particular on the device - network - content communication chain.

[ETSI-IP6-WhitePaper] reported the IPv6 Best Practices, Benefits, Transition Challenges and the Way Forward. IPv6 is becoming a priority again and a new wave of IPv6 deployment is expected, due the exhaustion of the IPv4 address space since 2010, in addition technologies like 5G, cloud, IoT require its use, governments and standard bodies (including IETF) demand it, and the device - network - content communication chain is calling for its adoption. In this regard it is possible to mention the IAB Statement on IPv6 stating that "IETF will stop requiring IPv4 compatibility in new or extended protocols".

The following sections go through the issue of IPv4 address exhaustion and give the global picture of IPv6 to show how IPv6 is growing faster than IPv4 worldwide in all measures including number of users, percentage of content, and amount of traffic. This

testifies that the key Internet industry players have decided strategically to invest and deploy IPv6 in large-scale to sustain the Internet growth.

Then it is presented the survey among network operators as well as considerations and observations for enterprises and content and cloud service providers about the IPv6 deployment and the considerations that have come out. IPv6 transition solutions for Mobile BroadBand (MBB), Fixed BroadBand (FBB) and enterprise services are ready. Dual-Stack is the most deployed solution for IPv6 introduction, while 464XLAT and Dual Stack Lite (DS-Lite) seem the most suitable for IPv6-only service delivery.

Finally, The IPv6 incentives are presented but the general IPv6 challenges are also reported in particular in relation to Architecture, Operations, Performance and Security issues. These considerations aim to start a call for action on the areas of improvement, that are often mentioned as reason for not deploying IP6.

2. IPv4 Address Exhaustion

According to [CAIR] there will be 29.3 billion networked devices by 2023, up from 18.4 billion in 2018. This poses the question on whether the IPv4 address space can sustain such a number of allocations and, consequently, if this is affecting the process of its exhaustion. The answer is not straightforward as many aspects have to be considered.

On the one hand, the RIRs are reporting scarcity of available and still reserved addresses. Table 3 of [POTAR001] shows that the available pool of the five RIRs counts a little more than 6 million IPv4 address, while the reserved pool includes another 12 million, for a total of "usable" addresses equal to 18.3 million. The same reference, in table 1, shows that the total IPv4 allocated pool equals 3.684 billion addresses. The ratio between the "usable" addresses and the total allocated brings to 0.005% of remaining space.

On the other, [POTAR001] again highlights the role of both NAT and the address transfer to counter the IPv4 exhaustion. NAT systems well fit in the current client/server model used by most of the available Internet applications, with this phenomenon amplified by the general shift to cloud. The transfer of IPv4 addresses also contributes to mitigate the the need of addresses. As an example, [IGP-GT] shows the amount of transfers to recipient organizations in the ARIN region in 2018. Cloud Service Providers (CSPs) appear to be

the most active is buying available addresses to satisfy their need of providing IPv4 connectivity to their tenants.

3. The global picture of IPv6

The utilization of IPv6 has been monitored by many agencies and institutions worldwide. Different analytics have been made available, ranging from the number of IPv6 users, its relative utilization over the Internet, to the number of carriers able to route IPv6 network prefixes. [ETSI-IP6-WhitePaper] provided several of those analytics. The scope of this section then is to summarize the status of the IPv6 adoption, so to get an indication of the relevance of IPv6 today. For the analytics listed here, the trend over the past five years is given, expressed as the Compound Annual Growth Rate (CAGR). In general, this shows how IPv6 has grown in the past few years, and that is growing faster than IPv4.

3.1. IPv6 users

[ETSI-IP6-WhitePaper] provided the main statistics about the utilization of IPv6 worldwide and references the organizations that make their measurement publicly available through their web sites. To give a rough estimation of the relative growth of IPv6, the next table shows the total number of estimated IPv6 users at December 2020 as measured by [POTAROO2], [APNIC1].

	Dec 2016	Dec 2017	Dec 2018	Dec 2019	Dec 2020	CAGR
World	300.85	473.14	543.04	990.19	1,201.09	41%

Figure 1: IPv6 users worldwide (in millions)

3.2. IPv6 allocations and networks

Regional Internet Registries (RIRs) are responsible for assigning an IPv6 address block to ISPs or enterprises. An ISP will use the assigned block to provide addresses to their end users. For example, a mobile carrier will assign one or several /64 prefixes to the end users. Several analytics are available for the RIRs. The next table shows the amount of individual allocations, per RIR, in the time period 2016-2020 [APNIC2].

Registry	Dec 2016	Dec 2017	Dec 2018	Dec 2019	Dec 2020	Cumulated	CAGR
AFRINIC	116	112	110	115	109	562	48%
APNIC	1,681	1,369	1,474	1,484	1,498	7,506	45%
ARIN	646	684	659	605	644	3,238	50%
LACNIC	1,009	1,549	1,448	1,614	1,801	7,421	65%
RIPE NCC	2,141	2,051	2,620	3,104	1,403	11,319	52%
Total	5,593	5,765	6,311	6,922	5,455	30,046	52%

Figure 2: IPv6 allocations worldwide

Note that the decline in 2020 of IPv6 allocations from the RIPE NCC could be explained with the COVID-19 measures that affect many European countries. Anyway countries all over the world have been similarly affected, but the decline in IPv6 allocation activity in 2020 is only seen in the data from the RIPE NCC.

[APNIC2] also compares the number of allocations for both address families, and the result is in favor of IPv6. The average yearly growth is 52% for IPv6 in the period 2016-2020 versus 49% for IPv4, a sign that IPv6 is growing bigger than IPv4. This is described in the next table.

Address family	Dec 2016	Dec 2017	Dec 2018	Dec 2019	Dec 2020	Cumulated	CAGR
IPv6	5,593	5,765	6,311	6,922	5,455	30,046	52%
IPv4	10,515	9,437	10,192	14,019	7,437	51,600	49%

Figure 3: Allocations per address family

The next table is based on [APNIC3], [APNIC4] and shows the percentage of ASes supporting IPv6 compared to the total ASes worldwide. The number of IPv6-capable ASes increases from 22.6% in January 2017 to 30.4% in January 2021. This equals to 14% CAGR for IPv6 enabled networks. This also shows that the number of networks supporting IPv6 is growing faster than the ones supporting IPv4, since the total (IPv6 and IPv4) networks grow at 6% CAGR.

Advertised ASN	Jan 2017	Jan 2018	Jan 2019	Jan 2020	Jan 2021	CAGR
IPv6-capable	12,700	14,500	16,470	18,600	21,400	14%
Total ASN	56,100	59,700	63,100	66,800	70,400	6%
Ratio	22.6%	24.3%	26.1%	27.8%	30.4%	

Figure 4: Percentage of IPv6-capable ASes

4. Survey among Network Operators

It was started an IPv6 poll to more than 50 network operators about the status of IPv6 deployment. This poll reveals that more than 30 operators will migrate fixed and mobile users to IPv6 in next 2 years. The IPv6 Poll has been submitted in particular to network operators considering that, as showed by the previous section, both user devices and contents seem more ready for IPv6. The answers to the questionnaire can be found in Appendix.

The main Questions asked are:

* Do you plan to move more fixed or mobile or enterprise users to IPv6 (e.g. Dual-Stack) or IPv6-only in the next 2 years? What are the reasons to do so? Which transition solution will you use, Dual-Stack, DS-Lite, 464XLAT, MAP-T/E?

* Do you need to change network devices for the above goal? Will you migrate your metro or backbone or backhaul network to support IPv6?

The result of this questionnaire highlights that major IPv6 migration will happen in next 2 years. Dual Stack is always the most adopted solution and the transition to IPv6-only is motivated in particular by business reasons like the 5G and IoT requirements. In addition it is worth mentioning that the migration of transport network (metro and backbone) is not considered a priority today for many network operators and the focus is in particular on the end to end IPv6 services.

More details about the answers received can be found in the Appendix.

5. Considerations for Enterprises

As described in [RFC7381], enterprises face different challenges than operators. The overall problem for many enterprises is to handle IPv6-based connectivity to the upstream providers, while supporting a mixed IPv4/IPv6 domain in the internal network.

The business reasons for IPv6 is unique to each enterprise especially for the internal network. But the most common drivers are on the external network due to the fact that when Internet service providers, run out of IPv4 addresses, they will provide native IPv6 and non-native IPv4. So for client networks trying to reach enterprise networks, the IPv6 experience will be better than the transitional IPv4 if the enterprise deploys IPv6 in its public-facing services. Enterprise that is or will be expanding into emerging markets or that partners with other companies who use IPv6 (larger enterprise, governments, service providers) has to deploy IPv6 or plan to do in the near term to support the long term goals. As an example it is possible to mention the emerging energy market and in particular SmartGrid where high density of IP-enabled endpoints are needed and IPv6 is a key technology.

6. Observations on Content and Cloud Service Providers

The number of addresses required to connect all of the virtual and physical elements in a Data Center and the necessity to overcome the limitation posed by [RFC1918] has been the driver to adopt IPv6 in several Content and Cloud Service Provider (CSP) networks.

Several public references discuss how most of the major players find themselves at different stages in the transition to IPv6-only in their DC infrastructure. In some cases, the transition already happened and the DC infrastructure of these hyperscalers is completely based on IPv6. This can be considered a good sign because the end-to-end connectivity between a client (e.g. an application on a smartphone) and a server (a Virtual Machine in a DC) may be based on IPv6.

7. Industrial Internet application

There are potential advantages for implementing IPv6 for IIoT (Industrial Internet of Things) applications, in particular the large IPv6 address space, the automatic IPv6 configuration and resource discovery.

However, there are still many obstacles that prevent its pervasive use. The key problems identified are the incomplete or immature tool support, the dependency on manual configuration and the poor

knowledge of the IPv6 protocols among insiders. To advance and ease the use of IPv6 for smart manufacturing systems and IIoT applications in general, a generic approach to remove these pain points is therefore highly desirable.

8. IPv6 deployments worldwide

This section reports the most deployed approaches for the IPv6 migration in MBB, FBB and enterprise.

8.1. IPv6 service design for Mobile, Fixed broadband and enterprises

The consolidated strategy, as also described in [ETSI-IP6-WhitePaper], is based on two stages, namely: (1) IPv6 introduction, and (2) IPv6-only. The first stage aims at delivering the service in a controlled manner, where the traffic volume of IPv6-based services is minimal. When the service conditions change, e.g. when the traffic grows beyond a certain threshold, then the move to the second stage may occur. In this latter case, the service is delivered solely on IPv6.

8.1.1. IPv6 introduction

In order to enable the deployment of an IPv6 service over an underlay IPv4 architecture, there are two possible approaches:

- o Enabling Dual-Stack at the CPE
- o Tunneling IPv6 traffic over IPv4, e.g. with 6rd.

So, from a technical perspective, the first stage is based on Dual-Stack [RFC4213] or tunnel-based mechanisms such as Generic Routing Encapsulation (GRE), IPv6 Rapid Deployment (6rd), Connection of IPv6 Domains via IPv4 Clouds (6to4), and others.

Dual-Stack [RFC4213] is more robust, and easier to troubleshoot and support. Based on information provided by operators with the answers to the poll (see Appendix A), it can be stated that Dual-Stack is currently the most widely deployed IPv6 solution, for MBB, FBB and enterprises, accounting for about 50% of all IPv6 deployments, see both Appendix A and the statistics reported in [ETSI-IP6-WhitePaper]. Therefore, for operators that are willing to introduce IPv6 the most common approach is to apply the Dual-Stack transition solution.

With Dual-Stack, IPv6 can be introduced together with other network upgrade and many parts of network management and IT systems can still work in IPv4. This avoids major upgrade of such systems to support IPv6, which is possibly the most difficult task in IPv6 transition.

In other words, the cost and effort on the network management and IT system upgrade are moderate. The benefits are to start to accommodate future services and save the NAT costs.

The CPE has only an IPv6 address at the WAN side and uses an IPv6 connection to the operator gateway, e.g. Broadband Network Gateway (BNG) or Packet Gateway (PGW) / User Plane Function (UPF). However, the hosts and content servers can still be IPv4 and/or IPv6. For example, NAT64 can enable IPv6 hosts to access IPv4 servers. The backbone network underlay can also be IPv4 or IPv6.

Although the Dual-Stack IPv6 transition is a good solution to be followed in the IPv6 introduction stage, it does have few disadvantages in the long run, like the duplication of the network resources and states, as well as other limitations for network operation. For this reason, when IPv6 increases to a certain limit, it would be better to switch to the IPv6-only stage.

8.1.2. IPv6-only service delivery

The second stage, named here IPv6-only, can be a complex decision that depends on several factors, such as economic factors, policy and government regulation.

[I-D.lmhp-v6ops-transition-comparison] discusses and compares the technical merits of the most common transition solutions for IPv6-only service delivery, 464XLAT, DS-lite, Lightweight 4over6 (lw4o6), MAP-E, and MAP-T, but without providing an explicit recommendation. As the poll highlights, the most widely deployed IPv6 transition solution for MBB is 464XLAT and for FBB is DS-Lite.

Based on the survey among network operators in Appendix A it is possible to analyze the IPv6 transition technologies that are already deployed or that will be deployed. The different answers to the questionnaire and in particular [ETSI-IP6-WhitePaper] reported detailed statistics on that and it can be stated that, besides Dual-Stack, the most widely deployed IPv6 transition solution for MBB is 464XLAT [RFC6877], and for FBB is DS-Lite [RFC6333], both of which are IPv6-only solutions.

Looking at the different feedback from network operators, in some cases, even when using private addresses, such as 10.0.0.0/8 space [RFC1918], the address pool is not large enough, e.g. for large mobile operators or large Data Centers (DCs), Dual-Stack is not enough, because it still requires IPv4 addresses to be assigned. Also, Dual-Stack will likely lead to duplication of several network operations both in IPv6 and IPv4 and this increases the amount of state information in the network with a waste of resources. For this

reason, in some scenarios (e.g. MBB or DCs) IPv6-only stage could be more efficient from the start since the IPv6 introduction phase with Dual-Stack may consume more resources (for example CGNAT costs).

So, in general, it is possible to state that, when the Dual-Stack disadvantages outweigh the IPv6-only complexity, it makes sense to migrate to IPv6-only. Some network operators already started this process, while others are still waiting.

9. Findings of the IPv6 Survey

Global IPv4 address depletion is reported by most network operators as the important driver for IPv6 deployment. Indeed, the main reason for IPv6 deployment given is related to the run out of private 10.0.0.0/8 space [RFC1918]. 5G and IoT service deployment is another incentive not only for business reasons but also for the need of more addresses.

The answers in Appendix shows that the IPv6 deployment strategy is based mainly on Dual Stack architecture and most of the network operators are migrating or plan to migrate in the next few years. The main motivation is related to the depletion of IPv4 addresses and to save the NAT costs.

It is interesting to see that most of the network operators have no big plans to migrate transport network (metro and backbone) soon, since they do not see business reasons. It seems that there is no pressure to migrate to native IPv6 forwarding in the short term, anyway the future benefit of IPv6 may justify in the long term a migration to native IPv6. Some network operators also said that a software upgrade can be enough to support IPv6 where it is needed for now.

This survey demonstrates that full replacement of IPv4 will take long time. Indeed the transition to IPv6 has different impacts and requirements depending on the network segment:

- o It is possible to say that almost all mobile devices are already IPv6 capable while for fixed access most of the CPEs are Dual Stack. Data Centers are also evolving and deploying IPv6 to cope with the increasing demand of cloud services.
- o While the access network seems not strongly impacted because it is mainly based on layer 2 traffic, regarding Edge and BNG, most network operators that provide IPv6 connectivity runs BNG devices in Dual Stack in order to distribute both IPv4 and IPv6.

- o For Metro and Backbone, the trend is to keep MPLS Data Plane and run IPv6/IPv4 over PE devices at the border. All MPLS services can be guaranteed in IPv6 as well through 6PE/6VPE protocols.

In this scenario it is clear that the complete deployment of a full IPv6 data plane will take more time. If we look at the long term evolution, IPv6 can bring other advantages like introducing advanced protocols developed only on IPv6 (e.g. SRv6) to implement all the controlled SLA services aimed by the 5G technology and beyond.

10. IPv6 incentives

It is possible to state that IPv6 adoption is no longer optional, indeed there are several incentives for the IPv6 deployment:

Technical incentives: all Internet technical standard bodies and network equipment vendors have endorsed IPv6 and view it as the standards-based solution to the IPv4 address shortage. The IETF, as well as other SDOs, need to ensure that their standards do not assume IPv4. The IAB expects that the IETF will stop requiring IPv4 compatibility in new or extended protocols. Future IETF protocol work will then optimize for and depend on IPv6. It is recommended that all networking standards assume the use of IPv6 and be written so they do not require IPv4 ([RFC6540]). In addition, every Internet registry worldwide strongly recommends immediate IPv6 adoption.

Business incentives: with the emergence of new digital technologies, such as 5G, IOT and Cloud, new use cases have come into being and posed more new requirements for IPv6 deployment. Over time, numerous technical and economic stop-gap measures have been developed in an attempt to extend the lifetime of IPv4, but all of these measures add cost and complexity to network infrastructure and raise significant barriers to innovation. It is widely recognized that full transition to IPv6 is the only viable option to ensure future growth and innovation in Internet technology and services. Several large networks and Data Centers have already evolved their internal infrastructures to be IPv6-only. Forward looking large corporations are also working toward migrating their enterprise networks to IPv6-only environments.

Governments incentives: governments have a huge responsibility in promoting IPv6 deployment within their countries. There are example of governments already adopting policies to encourage IPv6 utilization or enforce increased security on IPv4. So, even without funding the IPv6 transition, governments can recommend to add IPv6 compatibility for every connectivity, service or products

bid. This will encourage the network operators and vendors who don't want to miss out on government related bids to evolve their infrastructure to be IPv6 capable. Any public incentives for technical evolution will be bonded to IPv6 capabilities of the technology itself. In this regard, in the United States, the Office of Management and Budget is calling for an implementation plan to have 80% of the IP-enabled resources on Federal networks be IPv6-only by 2025. If resources cannot be converted, then the Federal agency is required to have a plan to retire them. The Call for Comment is at [US-FR] and [US-CIO].

11. Call for action

There are some areas of improvement, that are often mentioned in the literature and during the discussions on IPv6 deployment. This section lists these topics and wants to start a call for action to encourage more investigations on these aspects.

11.1. Transition choices

From an architectural perspective, a service provider or an enterprise may perceive quite a complex task the transition to IPv6, due to the many technical alternatives available and the changes required in management and operations. Moreover, the choice of the method to support the transition may depend on factors specific to the operator's or the enterprise's context, such as the IPv6 network design that fits the service requirements, the deployment strategy, and the service and network operations.

This section briefly highlights the basic approaches that service providers and enterprises may take. The scope is to raise the discussion whether actions may be taken that allow to overcome the issues highlighted and further push the adoption of IPv6.

11.1.1. Service providers

For a service provider, the IPv6 transition often refers to the service architecture (also referred to as overlay) and not to the network architecture (underlay). IPv6 is introduced at the service layer when a service requiring IPv6-based connectivity is deployed in an IPv4-based network. In this case, as already mentioned in the previous sections, a strategy is based on two stages: IPv6 introduction and IPv6-only.

For fixed operators, the massive CPE software upgrade to support Dual Stack started in most of service providers network and the traffic percentage is currently between 30% and 40% of IPv6, looking at the global statistics. This is valid for a network operator that

provides Dual Stack and gives the same opportunity for end terminal applications to choose freely the path that they want and assuming a normal internet usage. Anyway, it is interesting to see that in the latest years all major content providers have already implemented dual stack access to their services and most of them have implemented IPv6-only in their Data Centers. This aspect could affect the decision on the IPv6 adoption for an operator, but there are also other aspects like the current IPv4 addressing status, CPE costs, CGNAT costs and so on. Most operators already understood the need to adopt IPv6 in their networks and services, and also to promote the diffusion into their clients, while others are still at the edge of a massive implementation decision. Indeed, two situations are possible:

Operators that have already employed CGNAT and have introduced IPv6 in their networks, so they remain attached to a Dual Stack architecture. Although IPv6 brought them to a more technological advanced state, CGNAT, on the other end, boosts for some time their ability to supply CPE IPv4 connectivity.

Operators with a Dual Stack architecture that have introduced IPv6 both in the backbone and for the CPEs, but when reaching the limit in terms of number of IPv4 addresses available, they need to start defining and start to apply a new strategy that can be through CGNAT or with an IPv6-only approach.

For mobile operators, the situation is different since they are stretching their IPv4 address space since CGNAT translation levels have been reached and no more IPv4 public pool addresses are available. The new requirements from IoT services, 5G 3GPP release implementations, Voice over Long-Term Evolution (VoLTE) together with the constraints of national regulator lawful interception are seen as major drivers for IPv6. For these reasons, two situations are possible:

Some mobile operators choose to implement Dual-Stack as first and immediate mitigation solution.

Other mobile operators prefer to move to IPv6-only solution(e.g. 464XLAT) since Dual-Stack only mitigates and does not solve completely the IPv4 number scarcity issue.

11.1.2. Enterprises

The dual stage approach described in the previous sections can be still applicable for enterprises, even if the priorities to apply either stage are different since they have to consider both the internal and external network:

It is possible to start with Dual-Stack on hosts/OS and then in client network distribution layer. This allows the IPv6 introduction independently since both hosts/OS and client networks belong to the domain of the enterprise.

Dual-Stack can be further extended to WAN/campus core/edge routers. Also, as temporary solution, the use of NAT64 is recommended for servers/apps only capable of IPv4. Enterprise Data Center is also to be considered for the IPv6 transition. In this regard the application support needs to be taken into account, even if virtualization should make DCs simpler and more flexible.

There are additional challenges also related to the campus network and the cloud interconnection, indeed the networking may be not homogeneous. IPv6 could help to build a flat network by leveraging SD-WAN integration. The perspective of IPv6-only could also ensure better end-to-end performance.

Enterprises (private, managed networks) in US and Europe have failed to adopt IPv6, especially on internal networks. Other countries, in particular in Asia, who faced a shortage of IPv4 addresses, have moved somewhat more quickly. But, even there, the large "brick-and-mortar" enterprises find no business reason to adopt IPv6.

The enterprise engineers and technicians also don't know how IPv6 works. The technicians want to get trained yet the management does not feel that they do not want to pay for such training because they do not see a business need for adoption. This creates an unfortunate cycle where misinformation about the complexity of the IPv6 protocol and unreasonable fears about security and manageability combine with the perceived lack of urgent business needs to prevent adoption of IPv6.

In 2019 and 2020, there has been a concerted effort by some grass roots non-profits working with ARIN and APNIC to provide training [ARIN-CG] [ISIF-ASIA-G].

Having said that, some problems such as the problem of application conversion from IPv6 are quite difficult. The reliance of the economic, governmental, and military enterprise organizations on computer applications is great; the number of legacy systems, and ossification at such organizations, is also great. A number of mission-critical computer applications were written in the 1970's. While they have the source code, no one at the enterprise may be familiar with the application nor do they have the funds for external resources. So, transitioning to IPv6 is quite difficult.

The problem may be that of "First Mover Disadvantage". Understandably, corporations, having responsibility to their stockholders, have upgraded to new technologies and architectures, such as IPv6, only if it gains them revenue. Thus, legacy programs and technical debt accumulate.

11.1.3. Cloud and Data Centers

It was already highlighted how CSPs have adopted IPv6 in their internal infrastructure but are also active in gathering IPv4 addresses on the transfer market to serve the current business needs of IPv4 connectivity. This is primarily directed to serve the transition to cloud of enterprise's applications.

As noted in the previous section, most enterprises do not consider the transition to IPv6 as a priority. To this extent, the use of IPv4-based network services by the CSPs will last. Yet, CSPs are struggling to buy IPv4 addresses. If, in the next years, the scarcity of IPv4 addresses becomes more evident, it is likely that the cost of buying an IPv4 address by a CSP will be charged to an enterprise as a fee. From a financial standpoint this effect might be taken into consideration when evaluating the decision of moving to IPv6.

11.1.4. Industrial Internet

As the most promising protocol for network applications, IPv6 is frequently mentioned in relation to Internet of Things and Industry 4.0. However, its industrial adoption, in particular in smart manufacturing systems, has been much slower than expected. Indeed, it is important to provide an easy way to familiarize system architects and software developers with the IPv6 protocol and its role in the application development life cycle in order to limit the dependency on manual configuration and improve the tool support.

It is possible to differentiate types of data and access to understand how and where the IPv6 transition can happen. In the control network, determinism is required with full operational visibility and control, as well as reliability and availability. In monitoring IoT, best effort can be acceptable and low OPEX, zero-touch functions autoconfiguration, zero-configuration. For diagnostics and alerts, trust and transmissions that do not impact the control network are needed. For safety, guarantees in terms of redundancy, latency similar to the control network but with total assurance, is necessary.

For IIoT applications, it would be desirable to be able to implement a truly distributed system without dependencies to central components

like a DHCP server. In this regard the distributed IIoT applications can leverage the configuration-less characteristic of IPv6 and in this regard all the possible problems and compatibility issues with IPv6 link local addresses, SLAAC (StateLess Address Auto Configuration) needs to be investigated.

In addition, it could be interesting to have the ability to use IP based communication and standard application protocols at every point in the production process and further reduce the use of specialized communication systems like PLCs (Programmable Logic Controllers) and fieldbuses for real-time control to subsystems where this is absolutely necessary.

11.1.5. Government and Regulators

The slogan should be "stimulate if you can, regulate if you must". The global picture shows that the deployment of IPv6 worldwide is not uniform at all [G_stats], [APNIC1]. Countries where either market conditions or local regulators have stimulated the adoption of IPv6 show clear sign of growth.

As an example, zooming into the European Union area, countries such as Belgium, France and Germany are well ahead in terms of IPv6 adoption. The French National Regulator, Arcep, can be considered a good reference of National support to IPv6. [ARCEP] introduced an obligation for the operators awarded with a license to use 5G frequencies (3.4-3.8GHz) in Metropolitan France to be IPv6 compatible. As stated, "the goal is to ensure that services are interoperable and to remove obstacles to using services that are only available in IPv6, as the number of devices in use continues to soar, and because the RIPE NCC has run out of IPv4 addresses". A slow adoption of IPv6 could prevent new Internet services to widespread or create a barrier to entry for newcomers to the market. "IPv6 can help to increase competition in the telecom industry, and help to industrialize a country for specific vertical sectors".

A renewed industrial policy might be advocated in other countries and regions to stimulate IPv6 adoption. As an example, in the United States, the Office of Management and Budget is also calling for IPv6 adoption [US-FR], [US-CIO].

11.2. Network Operations

An important factor is represented by the need for training the network operations workforce. Deploying IPv6 requires it as policies and procedures have to be adjusted in order to successfully plan and complete an IPv6 migration. Staff has to be aware of the best practices for managing IPv4 and IPv6 assets. In addition to network

nodes, network management applications and equipment need to be properly configured and in some cases also replaced. This may introduce more complexity and costs for the migration.

11.3. Performance

Despite their relative differences, people tend to compare the performance of IPv6 versus IPv4, even if these differences are not so important for applications. In some cases, IPv6 behaving "worse" than IPv4 tends to re-enforce the justification of not moving towards the full adoption of IPv6. This position is supported when looking at available analytics on two critical parameters: packet loss and latency. These parameters have been constantly monitored over time, but only a few extensive researches and measurement campaigns are currently providing up-to-date information. This paragraph will look briefly at both of them, considering the available measurements. Operators are invited to bring in their experience and enrich the information reported below.

11.3.1. IPv6 latency

[APNIC5] constantly compares the latency of both address families. Currently, the worldwide average is still in favor of IPv4. Zooming at the country or even at the operator level, it is possible to get more detailed information and appreciate that cases exist where IPv6 is faster than IPv4. [APRICOT] highlights how when a difference in performance exists it is often related to asymmetric routing issues. Other possible explanations for a relative latency difference lays on the specificity of the IPv6 header which allows packet fragmentation. In turn, this means that hardware needs to spend cycles to analyze all of the header sections and when it is not capable of handling one of them it drops the packet. Even considering this, a difference in latency stands and sometimes it is perceived as a limiting factor for IPv6. A few measurement campaigns on the behavior of IPv6 in Content Delivery Networks (CDN) are also available [MAPRG-IETF99], [INFOCOM]. The TCP connect time is still higher for IPv6 in both cases, even if the gap has reduced over the analysis time window.

11.3.2. IPv6 packet loss

[APNIC5] also provides the failure rate of IPv6. Two reports, namely [RIPE1] and [APRICOT], discussed the associated trend, showing how the average worldwide failure rate of IPv6 worsened from around 1.5% in 2016 to a value exceeding 2% in 2020. Reasons for this effect may be found in endpoints with an unreachable IPv6 address, routing instability or firewall behaviours. Yet, this worsening effect may appear as disturbing for a plain transition to IPv6. Operators are

once again invited to share their experience and discuss the performance of IPv6 in their network scenarios.

11.3.3. Router's performance

It is worth mentioning the aspect of Router's performance too. IPv6 is 4 times longer than IPv4 and it is possible to do a simple calculation: the same memory on routers could permit to have 1/4 of different tables (routing, filtering, next hop). Anyway most of the routers showed a remarkably similar throughput and latency for IPv4 and IPv6. For smaller software switching platforms, some tests reported a lower throughput for IPv6 compared to IPv4 only in case of smaller packet sizes, while for larger hardware switching platforms there was no throughput variance between IPv6 and IPv4 both at larger frame sizes and at the smaller packet size.

11.4. IPv6 security

IPv6 presents a number of exciting possibilities for the expanding global Internet, however, there are also noted security challenges associated with the transition to IPv6. [I-D.ietf-opsec-v6] analyzes the operational security issues in several places of a network (enterprises, service providers and residential users).

The security aspects have to be considered to keep the same level of security as it exists nowadays in an IPv4-only network environment. The autoconfiguration features of IPv6 will require some more attention for the things going on at the network level. Router discovery and address autoconfiguration may produce unexpected results and security holes. The IPsec protocol implementation has initially been set as mandatory in every node of the network, but then relaxed to recommendation due to extremely constrained hardware deployed in some devices e.g., sensors, Internet of Things (IoT).

There are some concerns in terms of the security but, on the other hand, IPv6 offers increased efficiency. There are measurable benefits to IPv6 to notice, like more transparency, improved mobility, and also end to end security (if implemented).

As reported in [ISOC], comparing IPv6 and IPv4 at the protocol level, one may probably conclude that the increased complexity of IPv6 results in an increased number of attack vectors, that imply more possible ways to perform different types attacks. However, a more interesting and practical question is how IPv6 deployments compare to IPv4 deployments in terms of security. In that sense, there are a number of aspects to consider.

Most security vulnerabilities related to network protocols are based on implementation flaws. Typically, security researchers find vulnerabilities in protocol implementations, which eventually are "patched" to mitigate such vulnerabilities. Over time, this process of finding and patching vulnerabilities results in more robust implementations. For obvious reasons, the IPv4 protocols have benefited from the work of security researchers for much longer, and thus IPv4 implementations are generally more robust than IPv6.

Besides the intrinsic properties of the protocols, the security level of the resulting deployments is closely related to the level of expertise of network and security engineers. In that sense, there is obviously much more experience and confidence with deploying and operating IPv4 networks than with deploying and operating IPv6 networks.

Finally, implementation of IPv6 security controls obviously depends on the availability of features in security devices and tools. Whilst there have been improvements in this area, there is a lack of parity in terms of features and/or performance when considering IPv4 and IPv6 support in security devices and tools.

11.4.1. Protocols security issues

It is important to say that IPv6 is not more or less secure than IPv4 and the knowledge of the protocol is the best security measure.

In general there are security concerns related to IPv6 that can be classified as follows:

- o Basic IPv6 protocol (Basic header, Extension Headers, Addressing)
- o IPv6 associated protocols (ICMPv6, NDP, MLD, DNS, DHCPv6)
- o Internet-wide IPv6 security (Filtering, DDoS, Transition Mechanisms)

ICMPv6 is an integral part of IPv6 and performs error reporting and diagnostic functions. Since it is used in many IPv6 related protocols, ICMPv6 packet with multicast address should be filtered carefully to avoid attacks. Neighbor Discovery Protocol (NDP) is a node discovery protocol in IPv6 which replaces and enhances functions of ARP. Multicast Listener Discovery (MLD) is used by IPv6 routers for discovering multicast listeners on a directly attached link, much like Internet Group Management Protocol (IGMP) is used in IPv4.

These IPv6 associated protocols like ICMPv6, NDP and MLD are something new compared to IPv4, so they add new security threats and

the related solutions are still under discussion today. NDP has vulnerabilities [RFC3756] [RFC6583]. The specification says to use IPsec but it is impractical and not used, on the other hand, SEND (SEcure Neighbour Discovery) [RFC3971] is not widely available.

[RIPE2] describes the most important threats and solutions regarding IPv6 security.

11.4.2. IPv6 Extension Headers and Fragmentation

IPv6 Extension Headers imply some issues, in particular their flexibility also means an increased complexity, indeed security devices and software must process the full chain of headers while firewalls must be able to filter based on Extension Headers. Additionally, packets with IPv6 Extension Headers may be dropped in the public Internet.

There are some possible attacks through EHs, for example RH0 can be used for traffic amplification over a remote path and it is deprecated. Other attacks based on Extension Headers are based on IPv6 Header Chains and Fragmentation that could be used to bypass filtering, but, to mitigate this effect, Header chain should go only in the first fragment and the use of the IPv6 Fragmentation Header is forbidden in all Neighbor Discovery messages.

Fragment Header is used by IPv6 source node to send a packet bigger than path MTU and the Destination host processes fragment headers. There are several threats related to fragmentation to pay attention to e.g. overlapping fragments (not allowed) resource consumption while waiting for last fragment (to discard), atomic fragments (to be isolated).

11.4.3. Oversized IPv6 packets

A lot of additional functionality has been added to IPv6 primarily by adding Extension Headers and/or using overlay encapsulation. All of these expand the packet size and this could lead to oversized packets that would be dropped on some links.

It is better to investigate the potential problems with oversized packets in the first place. Fragmentation must not be done in transit and a better solution needs to be found, e.g. upgrade all links to bigger MTU or follow specific recommendations at the source node.

12. Security Considerations

This document has no impact on the security properties of specific IPv6 protocols or transition tools. The security considerations relating to the protocols and transition tools are described in the relevant documents.

13. Contributors

TBC

14. Acknowledgements

TBC

15. IANA Considerations

This document has no actions for IANA.

16. References

16.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

16.2. Informative References

- [APNIC1] APNIC, "IPv6 Capable Rate by country (%)", 2020, <<https://stats.labs.apnic.net/ipv6>>.
- [APNIC2] APNIC2, "Addressing 2020", 2021, <<https://labs.apnic.net/?p=1400>>.
- [APNIC3] APNIC, "BGP in 2019 - The BGP Table", 2020, <<https://blog.apnic.net/2020/01/14/bgp-in-2019-the-bgp-table/>>.
- [APNIC4] APNIC, "IPv6 in 2020", 2021, <<https://blog.apnic.net/2021/02/08/ipv6-in-2020/>>.
- [APNIC5] APNIC, "Average RTT Difference (ms) (V6 - V4) for World (XA)", 2020, <<https://stats.labs.apnic.net/v6perf/XA>>.

- [APRICOT] Huston, G., "Average RTT Difference (ms) (V6 - V4) for World (XA)", 2020, <<https://2020.apricot.net/assets/files/APAE432/ipv6-performance-measurement.pdf>>.
- [ARCEP] ARCEP, "Arcep Decision no 2019-1386, Decision on the terms and conditions for awarding licences to use frequencies in the 3.4-3.8GHz band", 2019, <https://www.arcep.fr/uploads/tx_gsavis/19-1386.pdf>.
- [ARIN-CG] ARIN, "Community Grant Program: IPv6 Security, Applications, and Training for Enterprises", 2020, <https://www.arin.net/about/community_grants/recipients/>.
- [CAIR] Cisco, "Cisco Annual Internet Report (2018-2023) White Paper", 2020, <<https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>>.
- [ETSI-IP6-WhitePaper] ETSI, "ETSI White Paper No. 35: IPv6 Best Practices, Benefits, Transition Challenges and the Way Forward", ISBN 979-10-92620-31-1, 2020.
- [G_stats] Google, "Google IPv6 Per-Country IPv6 adoption", 2021, <<https://www.google.com/intl/en/ipv6/statistics.html#tab=per-country-ipv6-adoption>>.
- [I-D.ietf-opsec-v6] Vyncke, E., Kk, C., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", draft-ietf-opsec-v6-21 (work in progress), November 2019.
- [I-D.lmhp-v6ops-transition-comparison] Lencse, G., Martinez, J., Howard, L., Patterson, R., and I. Farrer, "Pros and Cons of IPv6 Transition Technologies for IPv4aaS", draft-lmhp-v6ops-transition-comparison-06 (work in progress), January 2021.
- [IGP-GT] Internet Governance Project, Georgia Tech, "The hidden standards war: economic factors affecting IPv6 deployment", 2019, <<https://via.hypothes.is/https://www.internetgovernance.org/wp-content/uploads/IPv6-Migration-Study-final-report.pdf>>.

- [INFOCOM] Doan, T., "A Longitudinal View of Netflix: Content Delivery over IPv6 and Content Cache Deployments", 2020, <<https://dl.acm.org/doi/abs/10.1109/INFOCOM41043.2020.9155367>>.
- [ISIF-ASIA-G] ISIF Asia, "Internet Operations Research Grant: IPv6 Deployment at Enterprises. IIESoc. India", 2020, <<https://isif.asia/2020-grantees/>>.
- [ISOC] Internet Society, "IPv6 Security FAQ", 2019, <<https://www.internetsociety.org/wp-content/uploads/2019/02/Deploy360-IPv6-Security-FAQ.pdf>>.
- [MAPRG-IETF99] Bajpai, V., "Measuring YouTube Content Delivery over IPv6", 2017, <<https://www.ietf.org/proceedings/99/slides/slides-99-maprg-measuring-youtube-content-delivery-over-ipv6-00.pdf>>.
- [POTAROO1] POTAROO, "Addressing 2020", 2020, <<https://www.potaroo.net/ispcol/2021-01/addr2020.html>>.
- [POTAROO2] POTAROO, "IPv6 Resource Distribution Reports", 2021, <<https://resources.potaroo.net/iso3166/archive/>>.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <<https://www.rfc-editor.org/info/rfc3756>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <<https://www.rfc-editor.org/info/rfc3971>>.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, DOI 10.17487/RFC4213, October 2005, <<https://www.rfc-editor.org/info/rfc4213>>.

- [RFC6036] Carpenter, B. and S. Jiang, "Emerging Service Provider Scenarios for IPv6 Deployment", RFC 6036, DOI 10.17487/RFC6036, October 2010, <<https://www.rfc-editor.org/info/rfc6036>>.
- [RFC6180] Arkko, J. and F. Baker, "Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment", RFC 6180, DOI 10.17487/RFC6180, May 2011, <<https://www.rfc-editor.org/info/rfc6180>>.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <<https://www.rfc-editor.org/info/rfc6333>>.
- [RFC6540] George, W., Donley, C., Liljenstolpe, C., and L. Howard, "IPv6 Support Required for All IP-Capable Nodes", BCP 177, RFC 6540, DOI 10.17487/RFC6540, April 2012, <<https://www.rfc-editor.org/info/rfc6540>>.
- [RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, DOI 10.17487/RFC6583, March 2012, <<https://www.rfc-editor.org/info/rfc6583>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<https://www.rfc-editor.org/info/rfc6877>>.
- [RFC6883] Carpenter, B. and S. Jiang, "IPv6 Guidance for Internet Content Providers and Application Service Providers", RFC 6883, DOI 10.17487/RFC6883, March 2013, <<https://www.rfc-editor.org/info/rfc6883>>.
- [RFC7381] Chittimaneni, K., Chown, T., Howard, L., Kuarsingh, V., Pouffary, Y., and E. Vyncke, "Enterprise IPv6 Deployment Guidelines", RFC 7381, DOI 10.17487/RFC7381, October 2014, <<https://www.rfc-editor.org/info/rfc7381>>.
- [RIPE1] Huston, G., "Measuring IPv6 Performance", 2016, <<https://ripe73.ripe.net/wp-content/uploads/presentations/35-2016-10-24-v6-performance.pdf>>.
- [RIPE2] RIPE, "IPv6 Security", 2019, <<https://www.ripe.net/support/training/material/ipv6-security/ipv6security-slides.pdf>>.

- [US-CIO] The CIO Council, "Memorandum for Heads of Executive Departments and Agencies. Completing the Transition to Internet Protocol Version 6 (IPv6)", 2020, <<https://www.cio.gov/assets/resources/internet-protocol-version6-draft.pdf>>.
- [US-FR] Federal Register, "Request for Comments on Updated Guidance for Completing the Transition to the Next Generation Internet Protocol, Internet Protocol Version 6 (IPv6)", 2020, <<https://www.federalregister.gov/documents/2020/03/02/2020-04202/request-for-comments-on-updated-guidance-for-completing-the-transition-to-the-next-generation>>.

Appendix A. Summary of Questionnaire and Replies

This Appendix summarizes the questionnaire and the replies received.

1. Do you have plan to move more fixed or mobile or enterprise users to IPv6 in the next 2 years?
 - a. If yes, fixed, or mobile, or enterprise?
 - b. What're the reasons to do so?
 - c. When to start: already on going, in 12 months, after 12 months?
 - d. Which transition solution will you use, Dual-Stack, DS-Lite, 464XLAT, MAP-T/E?
2. Do you need to change network devices for the above goal?
 - a. If yes, what kind of devices: CPE, or BNG/mobile core, or NAT?
 - b. Will you migrate your metro or backbone or backhaul network to support IPv6?

Some answers below:

Answer 1: (1) Yes, IPv6 migration strategy relies upon the deployment of Dual Stack architecture. IPv4 service continuity designs is based on DS-Lite for fixed environments and 464XLAT for mobile environments. No plans to move towards MAP-E or MAP-T solutions for the time being. (2) Yes, it's a matter of upgrading CPE, routers (including BNGs), etc. Tunneling options (ISATAP, TEREDO, 6rd) will also be used for migration.

Answer 2: (1) Yes, at this moment we widely use IPv6 for mobile services while we are using DS-Lite for fixed services (FTTH and DSL). (2) We have no pressure to migrate to native IPv6 forwarding in the short term and it would represent a significant work without clear immediate benefit or business rationale. However we may see a future benefit with SRv6 which may justify in the long term a migration to native IPv6.

Answer 3: (1) Yes, fixed. The IP depletion topic is crucial, so we need to speed up the DS-Lite deployment and also Carrier Grade NAT introduction. (2) Yes, CGNAT introduction.

Answer 4: (1) No, we are rolling IPv6 users back to IPv4. DS-Lite. (2) No, it was already done. IPv6 works worse than IPv4. it is immature.

Answer 5: (1) Yes, all 3. Target is Dual-stack for fixed, mobile and enterprise. (2) Yes, we are adding specific services cards inside our FTTH equipment for dealing with CGNAT. Metro and backbone are already Dual Stack.

Answer 6: (1) Yes, Enterprises customer demand is high and the transition is on going through Dual-Stack. (2) No big plan for transport network.

Answer 7: No such requirements

Answer 8: (1) Yes, mobile. The Internet APN is not yet enabled for IPv6, this will be done soon. 464XLAT will be used to save on RFC1918 address space. (2) Yes, PGW; Metro is already IPv6 and Backbone is currently IPv4/MPLS. No native IPv6 planned as for now.

Answer 9: (1) Yes, Dual-Stack for all 3. Not all services are available on IPv6. IPv6 adoption has been stated from many years but still not finished. Dual-Stack is used. (2) No, at the moment it is 6PE solution. No plan to migrate on native IPv6.

Answer 10: (1) Yes, all 3. Ongoing transition with Dual-stack and 464XLAT. (2) No plan for Metro and Backbone.

Answer 11: No such requirements.

Answer 12: (1) Yes, mobile and fixed. To mitigate IPv4 exhaustion in 12 months, Dual-Stack is used. (2) No (hopefully). Managed by software upgrade.

Answer 13: (1) Yes, on Mobile and Fixed. Mobile: IPv4 exhaustion for the RAN transport and IPv6 roll out ongoing. Fixed: Enterprises are

requesting IPv6 and also competitors are offering it. Mobile: dual stack and 6VPE; Enterprise: Dual Stack and 6VPE. (2) No, maybe only a software upgrade.

Answer 14: (1) Yes, fixed. IPv4 address depletion, on going, Dual-Stack with NAT444. (2) No.

Answer 15: (1) Yes, Mobile. Running out of private IPv4 address space and do not want to overlap addresses. Transition on going through 464XLAT. (2) Not yet, this is not the most pressing concern at the moment but it is planned.

Answer 16: No, already on Dual-Stack for many years. Discussing IPv6-only.

Answer 17: (1) Yes, all 3, strategy on going, Dual-Stack, MAP-T. (2) Yes, CPE, BR Dual-Stack.

Answer 18: (1) Yes, Mobile, due to address deficit. It would be very likely 464XLAT. (2) It is not clear at the moment. Still under investigation. CPE, Mobile Core, NAT. For IPv6 native support no plans for today.

Answer 19: No. Difficult to do it for enterprises, and don't really care for residential customers.

Answer 20: (1) Yes, fixed, mobile. IP space depletion. Mobile and Backbone are already done, Fixed is becoming Dual-Stack. (2) Yes, ordinary CPE and small routers. Some of them needs just software upgrade. Backbone done, no plan for metro and backhaul.

Answer 21: No such requirements

Answer 22: (1) Yes, mobile, we have few enterprise requests for IPv6; fixed already Dual-Stack. We are in the exhaustion point in public IPv4 usage in mobile so we need to move to IPv6 in the terminals. Dual-Stack deployment is ongoing. (2) No, all devices already support dual-stack mode. No migration needed. We already support IPv6 forwarding in our backbone.

Answer 23: No, already Dual-Stack

Answer 24: (1) Yes, fixed. DS-Lite. (2) Yes, BNG supporting CGNAT.

Answer 25: (1) Yes, fixed. DS-Lite will be deployed. (2) Yes.

Answer 26: (1) Yes, Mobile (Fixed already Dual-Stack). IPv4 depletion and Business customers are asking for it. Dual-Stack will be deployed. (2) No.

Answer 27: (1) Yes, Mobile. Dual-Stack is on going. (2) Yes, MBH, mobile core.

Answer 28: No such requirements.

Answer 29: (1) Yes, fixed and mobile, enterprise is not certain. IPv4 addressing is not enough, fixed and mobile should be started in 12 months. (2) Telco Cloud, BNG and PEs already support IPv6.

Answer 30: (1) Yes, all 3. Government has pushed. Dual-Stack for FBB in 12 months. (2) Yes, RGs have not good readiness, but not much could be done about it. PPPoE access does not create problem in access and aggregation. BNG should only change configuration.

Answer 31: (1) Yes, mobile for 5G sites. Plan to use IPv6 soon. 6VPE in the beginning, then migrate to Dual-stack. (2) IP BH devices already support IPv6.

Answer 32: No.

Answer 33: Yes, Enterprises. We are running short of IPV4 addresses. In our Internet Core IPV4/IPV6 Dual Stack was already introduced. The rollout of IPV6 services is slow and we started with business services. From customer perspective Dual Stack is still a "must have" and this will be true for many years to come. Another thought is related to regulatory obligations. Anyway a total switch from IPv4 to IPv6 will not be possible for many more years.

Answer 34: No, we have no plans to introduce new wave of IPv6 in our network.

Answer 35: (1) Yes. Fixed, Enterprise. IPv4 addressing is not enough. Dual Stack deployment is ongoing. (2) Yes, CPE for metro and backbone.

Answer 36: (1) Yes, Fixed, Enterprise. Dual-Stack. (2) Yes, CPE for IPv6 service delivery support.

Answer 37: Yes, mobile and enterprise. 6PE is deployed on the PEs, and dual-stack. The PE supports IPv6 by modifying the live network configuration or upgrading the software.

Answer 38: Yes, both home broadband and enterprise services support IPv6. IPv6 services are basic capabilities of communication

networks. Currently 6RD, dual stack (native IPv6) in the future. The dual-stack feature does not require device changes. The home gateway is connected to the switch and the BNG. The Dual Stack can be supported through configuration changes. Both the metro and backbone networks use MPLS to provide bearer services and do not require IPv6 capabilities. IPv6 is not enabled on both the metro and backbone networks. IPv6 services are implemented through 6VPE.

Answer 39: (1) Yes, Enterprises B2B needs more IP addresses. Dual-Stack is already on going. (2) No, BNG/mobile core and NAT. Metro and Backbone already support today.

Answer 40: Not for now.

Authors' Addresses

Giuseppe Fioccola
Huawei Technologies
Riesstrasse, 25
Munich 80992
Germany

Email: giuseppe.fioccola@huawei.com

Paolo Volpato
Huawei Technologies
Via Lorenteggio, 240
Milan 20147
Italy

Email: paolo.volpato@huawei.com

Nalini Elkins
Inside Products
36A Upper Circle
Carmel Valley CA 93924
United States of America

Email: nalini.elkins@insidethestack.com

Sebastien Lourdez
Post Luxembourg

Email: sebastien.lourdez@post.lu