

LSR Working Group
Internet-Draft
Intended status: Informational
Expires: May 28, 2021

U. Chunduri
Futurewei USA
J. Tantsura
Apstra, Inc.
S. Hegde
Juniper Networks
November 24, 2020

IS-IS Multi Topology Deployment Considerations
draft-chunduri-lsr-isis-mt-deployment-cons-04

Abstract

This document analyzes IS-IS Multi Topology (MT) applicability in various IS-IS deployments. This document explores the nuances around the terminology and usage of various IS-IS address families, topologies with different considerations, for choosing the right combination for a specific deployment scenario.

This document also discusses various ways one can deploy IPv6 only IS-IS topology.

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 [RFC2119], RFC8174 [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 May 28, 2021.

Copyright Notice

Copyright (c) 2020 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. Need for MT in IS-IS networks	3
3. Acronyms	3
4. Topologies and Address Families	4
4.1. Single Topology Mode and Multiple Address Families	4
4.2. Multiple Topology Mode and Multiple Address Families	5
4.2.1. Transition Mode	6
4.3. IPv6 Only Topology	6
5. IS-IS MT and LFA	7
6. Acknowledgements	7
7. IANA Considerations	7
8. Security Considerations	7
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Authors' Addresses	8

1. Introduction

IS-IS originally developed for OSI [ISO.10589.1992] and extensions have been made available to support IPv4 [RFC1195]. A method for exchanging IPv6 routing information using the IS-IS routing protocol is specified in [RFC5308]. How to run a set of independent IP topologies with topology specific adjacencies, within a single IS-IS domain has been defined in IS-IS MT [RFC5120].

There are number of networks, including mobile backhaul networks seeking to use IPv6 only solutions. It is possible to conceive, various parts of the backhaul networks use IPv4 and appropriate migration strategy needed before eventually moving towards IPv6 only

network. While any IGP can be used in these networks, this document covers only IS-IS protocol aspects.

Various layer-3 DC fabric routing options (refs: openfabric, spine-leaf, controller-based) by changing or optimizing some aspects w.r.t adjacency formation, flooding optimizations, or/and mechanisms to automatically compute the location of the node in the fat tree topology are proposed recently and this document brings some of the multi topology deployment aspects relevant to these networks. Please note, part of the discussion around IS-IS MT is not specific to DC or CLOS fabrics and generally applicable to any IS-IS deployment but discussed here because of multiple proposals to use various forms of IS-IS in this context.

2. Need for MT in IS-IS networks

For mobile transport backhaul networks seeking only IPv6 network or transitioning from parts of the network with only IPv4, IS-IS MT is needed. For layer-3 DC fabric underlay, which provide reachability, only one address family (either IPv4 or IPv6) SHOULD be sufficient. However if either only IPv6 address family is needed in the underlay or deploying both IPv4 and IPv6 address families are desired discussion in Section 4 is relevant.

It is an unlikely requirement, where DC fabric to be partitioned logically to have different topologies in the underlay but this can happen in various scenarios as listed in Section 4.1. If one does the same to meet a particular requirement, it introduces a manageability complexity of these logical topologies. IS-IS MT [RFC5120] also designed to address the above need and discussion in Section 4.2 is relevant. It is worth noting, majority of the IS-IS deployments use MT primarily to have a separate logical topology for IPv6 address family.

3. Acronyms

IIH : IS-IS Hello Protocol Data Unit

LSP : Link State PDU

MT : Multi Topology

SPF : Shortest Path First

4. Topologies and Address Families

Terminology around IS-IS topologies and address families is somewhat confusing at best. Just to give an example, MT ID #2 defined in [RFC5120] says, it is "Reserved for IPv6 routing topology". While multiple MT ID's can be deployed in a network with IPv6 topologies, MT ID #2, perhaps referring to a first such topology with IPv6 only address family. This section details various topology and address family options possible with currently available IS-IS specifications with respective defined TLVs.

4.1. Single Topology Mode and Multiple Address Families

IS-IS with IPv4 address family and with wide-metrics [RFC5305] is widely deployed, with TLV 22 defined for IS Reachability and TLV 135 for IP (IPv4) reachability information. This is essentially a single topology for the entire IS-IS area/domain with a single address family (IPv4 unicast).

IS-IS can also be enabled with IPv6 unicast address family in a single topology mode along with IPv4 unicast address family. Here IPv6 uses the same underlying topology that is used for IPv4 and this can be done as specified in IS-IS IPv6 [RFC5308] which introduces TLV 236, an IPv6 reachability TLV. It is important to note same IS-IS adjacency is used for both address families and with a single SPF (decision process) both IPv4 and IPv6 reachability would be computed.

However, for the above to work effectively, both IPv4 and IPv6 address families MUST share a common network topology. That is to use IS-IS for IPv4 and IPv6 routing, any interface configured for IPv4 IS-IS MUST also be configured for IPv6 IS-IS, and vice versa. All routers within an IS-IS area (Level 1 routing) or domain (Level 2 routing) MUST also support the same set of address families: IPv4 only, IPv6 only, or both IPv4 and IPv6. Any discrepancy in the configuration w.r.t above can cause routing black holes and one such scenario is discussed below.




Figure 1: IS-IS with multiple address families

As shown, in the above diagram all routers in the network enabled with both IPv4 and IPv6 unicast address families at the IS level and single topology would be built. However, at a link level all but except one link, say if IPv6 is not configured on the link between the routers Rx and R2; due to a single IS-IS topology, the shortest path between Rx and R2 is the direct link and since IPv6 is not enabled on that link, Rx and R2 cannot exchange IPv6 data traffic even though there's an alternate path between them in the topology through Rx, R1, Ry and R2.

Hence to summarize the restrictions: all routers in the topology MUST support only IPv4, only IPv6 or both IPv4 and IPv6 address families on all links and node. In other words, network MUST be congruent. While this model is too simpler to operate, might not be flexible enough for some IS-IS deployments. Some examples where congruency is not possible as follows:

- a. When IPv6 is getting introduced in the network legacy nodes that are IPv6 incapable.
- b. Implementation issues causing IPv6 to be disabled on some nodes.
- c. Hardware scale limitations causing IPv6 to be disabled on some low-end nodes.

4.2. Multiple Topology Mode and Multiple Address Families

Multi-topology IS-IS uses multiple SPF's to compute routes and removes the restriction that all interfaces MUST support all configured address families and that all routers in an IS-IS area or domain MUST support the same set of address families. This introduces the concept of topology specific adjacency with MT IS Reachability TLV 222 and MT capable IPv4 Reachability with TLV 235 and MT capable IPv6 Reachability with TLV 237.

When MT IS-IS is enabled with IPv4 and IPv6 address families, the routers build two topologies, one for each address family (IPv4 and IPv6) and can find the optimum path for each address family even when some links in the network support only one of them. IS-IS MT [RFC5120] defines MT ID #0 for backward compatibility, as the "standard" topology and this essentially operate as IS-IS single topology mode as specified in Section 4.1 and supports both IPv4 and IPv6 address families. MT ID #2 [RFC5120] is defined for IPv6 address family in MT mode.

4.2.1. Transition Mode

Most of the vendors supported MT transition feature (though some vendors disabled to avoid confusion around this) in the IS-IS networks to facilitate MT deployments without disrupting the single topology mode. The MT transition mode allows a network operating in single topology IS-IS IPv6 [RFC5308] to continue to work while upgrading routers to include MT IS-IS IPv6 support i.e., MT ID #2 with [RFC5120]. While in transition mode, both types of TLVs (single-topology with TLVs 22/236 and MT with TLVs 222/237) are sent in LSPs for all configured IPv6 addresses, nodes can continue to process these and operate in single topology mode though being in MT mode ("standard" IS-IS topology with MT ID #0). After all routers in the area or domain have been upgraded to support MT IPv6 transition mode can be removed from the configuration. Once all routers in the area or domain are operating in MT IPv6 mode, the topological restrictions of single-topology mode can be made no longer in effect.

When transition mode is enabled, the router advertises both MT TLVs and the old style IS-IS IPv6 TLVs but the topological restrictions of the single topology mode discussed above are in effect. However, there were instances while this mode is enabled and expectations for different result in the actual deployments.

4.3. IPv6 Only Topology

Though it is theoretically possible to build IPv6 only underlay (with TLV 236 for IPv6 reachability prefixes) in single topology mode as discussed in Section 4.1, lot of legacy implementations require IPv4 address families too be configured in single topology mode (ingrained code structures for IPv4 address family). IPv6 only DC underlay network can be built with multi topology adjacencies (TLV 222) and reachability prefixes (TLV 237) with MT ID #2 as discussed above in Section 4.2. With this, any other address family can be introduced including "standard" topology MT ID #0 (Single topology mode with both address families) and there are no restrictions on which address family has to enable on which link as specified in Section 4.1.

5. IS-IS MT and LFA

IP Fast Reroute (FRR) or Loop Free Alternative (LFA) computation in MT mode are described in detail in Section 5.2 of [RFC5120].

6. Acknowledgements

Thanks to Acee Lindem, Chris Hopps, Michael Abramson and Les Ginsberg for various inputs on this work.

7. IANA Considerations

This document has no actions for IANA.

8. Security Considerations

Security concerns for IS-IS are addressed in [RFC5304] and [RFC5310]. Further security analysis for IS-IS protocol is done in [RFC7645].

This document does not introduce any change in any of the IS-IS protocol or IS-IS protocol extensions. This document also does not introduce any new security issues other than as noted in the referenced IS-IS protocol extensions.

9. References

9.1. Normative References

- [ISO.10589.1992]
International Organization for Standardization,
"Intermediate system to intermediate system intra-domain-
routing routine information exchange protocol for use in
conjunction with the protocol for providing the
connectionless-mode Network Service (ISO 8473)",
ISO Standard 10589, 1992.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and
dual environments", RFC 1195, DOI 10.17487/RFC1195,
December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.
- [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>>.
- [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

- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.
- [RFC7645] Chunduri, U., Tian, A., and W. Lu, "The Keying and Authentication for Routing Protocol (KARP) IS-IS Security Analysis", RFC 7645, DOI 10.17487/RFC7645, September 2015, <<https://www.rfc-editor.org/info/rfc7645>>.
- [RFC8518] Sarkar, P., Ed., Chunduri, U., Ed., Hegde, S., Tantsura, J., and H. Gredler, "Selection of Loop-Free Alternates for Multi-Homed Prefixes", RFC 8518, DOI 10.17487/RFC8518, March 2019, <<https://www.rfc-editor.org/info/rfc8518>>.

Authors' Addresses

Uma Chunduri
Futurewei USA
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: umac.ietf@gmail.com

Jeff Tantsura
Apstra, Inc.

Email: jefftant.ietf@gmail.com

Shraddha Hegde
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore, Karnataka 560103
USA

Email: shraddha@juniper.net

OPSEC
Internet-Draft
Intended status: Informational
Expires: November 7, 2021

E. Vyncke
Cisco
K. Chittimaneni
Square
M. Kaeo
Double Shot Security
E. Rey
ERNW
May 6, 2021

Operational Security Considerations for IPv6 Networks
draft-ietf-opsec-v6-27

Abstract

Knowledge and experience on how to operate IPv4 networks securely is available: whether it is an Internet Service Provider or an enterprise internal network. However, IPv6 presents some new security challenges. RFC 4942 describes security issues in the protocol, but network managers also need a more practical, operations-minded document to enumerate advantages and/or disadvantages of certain choices.

This document analyzes the operational security issues associated with several types of network and proposes technical and procedural mitigation techniques. This document is only applicable to managed networks, such as enterprise networks, service provider networks, or managed residential networks.

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 November 7, 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
1.1. Applicability Statement	4
2. Generic Security Considerations	4
2.1. Addressing	4
2.1.1. Use of ULAs	5
2.1.2. Point-to-Point Links	5
2.1.3. Loopback Addresses	5
2.1.4. Stable Addresses	6
2.1.5. Temporary Addresses for SLAAC	6
2.1.6. DHCP Considerations	8
2.1.7. DNS Considerations	8
2.1.8. Using a /64 per host	8
2.1.9. Privacy consideration of Addresses	8
2.2. Extension Headers	9
2.2.1. Order and Repetition of Extension Headers	9
2.2.2. Hop-by-Hop Options Header	10
2.2.3. Fragment Header	10
2.2.4. IP Security Extension Header	10
2.3. Link-Layer Security	11
2.3.1. Neighbor Solicitation Rate-Limiting	11
2.3.2. Router and Neighbor Advertisements Filtering	12
2.3.3. Securing DHCP	13
2.3.4. 3GPP Link-Layer Security	14
2.3.5. Impact of Multicast Traffic	15
2.3.6. SeND and CGA	15
2.4. Control Plane Security	16
2.4.1. Control Protocols	17
2.4.2. Management Protocols	18
2.4.3. Packet Exceptions	18
2.5. Routing Security	19
2.5.1. BGP Security	20

2.5.2.	Authenticating OSPFv3 Neighbors	20
2.5.3.	Securing Routing Updates	21
2.5.4.	Route Filtering	21
2.6.	Logging/Monitoring	21
2.6.1.	Data Sources	23
2.6.2.	Use of Collected Data	26
2.6.3.	Summary	29
2.7.	Transition/Coexistence Technologies	29
2.7.1.	Dual Stack	30
2.7.2.	Encapsulation Mechanisms	31
2.7.3.	Translation Mechanisms	35
2.8.	General Device Hardening	37
3.	Enterprises Specific Security Considerations	37
3.1.	External Security Considerations	38
3.2.	Internal Security Considerations	39
4.	Service Providers Security Considerations	40
4.1.	BGP	40
4.1.1.	Remote Triggered Black Hole Filtering (RTBH)	40
4.2.	Transition/Coexistence Mechanism	40
4.3.	Lawful Intercept	40
5.	Residential Users Security Considerations	41
6.	Further Reading	41
7.	Acknowledgements	42
8.	Security Considerations	42
9.	References	42
9.1.	Normative References	42
9.2.	Informative References	42
	Authors' Addresses	57

1. Introduction

Running an IPv6 network is new for most operators not only because they are not yet used to large-scale IPv6 networks but also because there are subtle but critical and important differences between IPv4 and IPv6, especially with respect to security. For example, all layer-2 interactions are now done using Neighbor Discovery Protocol [RFC4861] rather than using Address Resolution Protocol [RFC0826]. Also, there is no Network Address Port Translation (NAPT) defined in [RFC2663] for IPv6 even if [RFC6296] specifies a Network Prefix Translation for IPv6 (NPTv6) which is a 1-to-1 mapping of IPv6 addresses. Another important difference is that IPv6 is extensible with the use of extension headers.

IPv6 networks are deployed using a variety of techniques, each of which have their own specific security concerns.

This document complements [RFC4942] by listing security issues when operating a network (including various transition technologies). It

also provides more recent operational deployment experiences where warranted.

1.1. Applicability Statement

This document is applicable to managed networks, i.e., when the network is operated by the user organization itself. Indeed, many of the recommended mitigation techniques must be configured with detailed knowledge of the network (which are the default routers, the switch trunk ports, etc.). This covers Service Provider (SP), enterprise networks and some knowledgeable-home-user-managed residential networks. This applicability statement especially applies to Section 2.3 and Section 2.5.4.

2. Generic Security Considerations

2.1. Addressing

IPv6 address allocations and overall architecture are an important part of securing IPv6. Initial designs, even if intended to be temporary, tend to last much longer than expected. Although initially IPv6 was thought to make renumbering easy, in practice it may be extremely difficult to renumber without a proper IP Address Management (IPAM) system. [RFC7010] introduces the mechanisms that could be utilized for IPv6 site renumbering and tries to cover most of the explicit issues and requirements associated with IPv6 renumbering.

A key task for a successful IPv6 deployment is to prepare an addressing plan. Because an abundance of address space is available, structuring an address plan around both services and geographic locations allows address space to become a basis for more structured security policies to permit or deny services between geographic regions. [RFC6177] documents some operational considerations of using different prefix sizes for address assignments at end sites.

A common question is whether companies should use Provider Independent (PI) vs. Provider Allocated (PA) space [RFC7381], but from a security perspective there is little difference. However, one aspect to keep in mind is who has administrative ownership of the address space and who is technically responsible if/when there is a need to enforce restrictions on routability of the space, e.g., due to malicious criminal activity originating from it. Relying on PA address space may also increase the perceived need for address translation techniques such as NPTv6 and thereby augmenting the complexity of the operations including the security operations.

In [RFC7934], it is recommended that IPv6 network deployments provide multiple IPv6 addresses from each prefix to general-purpose hosts and it specifically does not recommend limiting a host to only one IPv6 address per prefix. It also recommends that the network give the host the ability to use new addresses without requiring explicit requests (for example by using SLAAC). Privacy Extensions as of [RFC8981] constitute one of the main scenarios where hosts are expected to generate multiple addresses from the same prefix and having multiple IPv6 addresses per interface is a major change compared to the unique IPv4 address per interface for hosts (secondary IPv4 addresses are not common); especially for audits (see section Section 2.6.2.3).

2.1.1. Use of ULAs

Unique Local Addresses (ULAs) [RFC4193] are intended for scenarios where interfaces are not globally reachable, despite being routed within a domain. They formally have global scope, but [RFC4193] specifies that they must be filtered at domain boundaries. ULAs are different from [RFC1918] addresses and have different use cases. One use of ULA is described in [RFC4864], another one is for internal communication stability in networks where external connectivity may come and go (e.g., some ISPs provide ULAs in home networks connected via a cable modem). It should further be kept in mind that ULA /48s from the fd00::/8 space (L=1) MUST be generated with a pseudo-random algorithm, per [RFC4193] section 3.2.1.

2.1.2. Point-to-Point Links

[RFC6164] in section 5.1 specifies the rationale of using /127 for inter-router point-to-point links to prevent the ping-pong issue between routers not correctly implementing [RFC4443] and also prevents a DoS attack on the neighbor cache. The previous recommendation of [RFC3627] has been obsoleted and marked Historic by [RFC6547]).

Some environments are also using link-local addressing for point-to-point links. While this practice could further reduce the attack surface of infrastructure devices, the operational disadvantages also need to be carefully considered; see also [RFC7404].

2.1.3. Loopback Addresses

Many operators reserve a /64 block for all loopback addresses in their infrastructure and allocate a /128 out of this reserved /64 prefix for each loopback interface. This practice facilitates configuration of Access Control List (ACL) rules to enforce a security policy for those loopback addresses.

2.1.4. Stable Addresses

When considering how to assign stable addresses for nodes (either by static configuration or by pre-provisioned DHCPv6 lease Section 2.1.6), it is necessary to take into consideration the effectiveness of perimeter security in a given environment.

There is a trade-off between ease of operation (where some portions of the IPv6 address could be easily recognizable for operational debugging and troubleshooting) versus the risk of trivial scanning used for reconnaissance. [SCANNING] shows that there are scientifically based mechanisms that make scanning for IPv6 reachable nodes more feasible than expected; see also [RFC7707].

Stable addresses also allow easy enforcement of a security policy at the perimeter based on IPv6 addresses. E.g., Manufacturer Usage Description (MUD) [RFC8520] is a mechanism where the perimeter defense can retrieve security policy template based on the type of internal device and apply the right security policy based on the device IPv6 address.

The use of well-known IPv6 addresses (such as ff02::1 for all link-local nodes) or the use of commonly repeated addresses could make it easy to figure out which devices are name servers, routers, or other critical devices; even a simple traceroute will expose most of the routers on a path. There are many scanning techniques possible and operators should not rely on the 'impossible to find because my address is random' paradigm (a.k.a. "security by obscurity"), even if it is common practice to have the stable addresses randomly distributed across /64 subnets and to always use DNS (as IPv6 addresses are hard for human brains to remember).

While in some environments obfuscating addresses could be considered an added benefit, it should not preclude enforcement of perimeter rules. Stable addresses following some logical allocation scheme may ease the operation (as simplicity always helps security).

Typical deployments will have a mix of stable and non-stable addresses; the stable addresses being either predictable (e.g., ::25 for a mail server) or obfuscated (i.e., appearing as a random 64-bit number).

2.1.5. Temporary Addresses for SLAAC

Historically, stateless address autoconfiguration (SLAAC) makes up the globally unique IPv6 address based on an automatically generated 64-bit interface identifier (IID) based on the EUI-64 MAC address combined with the /64 prefix (received in the Prefix Information

Option (PIO) of the Router Advertisement (RA)). The EUI-64 address is generated from the stable 48-bit MAC address and does not change even if the host moves to another network; this is of course bad for privacy as a host can be traced from network (home) to network (office or Wi-Fi in hotels). [RFC8064] recommends against the use of EUI-64 addresses; and it must be noted that most host operating systems do not use EUI-64 addresses anymore and rely on either [RFC8981] or [RFC8064].

Randomly generating an interface ID, as described in [RFC8981], is part of SLAAC with so-called privacy extension addresses and is used to address some privacy concerns. Privacy extension addresses, a.k.a., temporary addresses may help to mitigate the correlation of activities of a node within the same network and may also reduce the attack exposure window. But using [RFC8981] privacy extension addresses might prevent the operator from building host specific access control lists (ACLs). The [RFC8981] privacy extension addresses could also be used to obfuscate some malevolent activities and specific user attribution/accountability procedures should be put in place as described in Section 2.6.

[RFC8064] combined with the address generation mechanism of [RFC7217] specifies another way to generate an address while still keeping the same IID for each network prefix; this allows SLAAC nodes to always have the same stable IPv6 address on a specific network while having different IPv6 addresses on different networks.

In some specific use cases where user accountability is more important than user privacy, network operators may consider disabling SLAAC and relying only on DHCPv6; but not all operating systems support DHCPv6 so some hosts will not get any IPv6 connectivity. Disabling SLAAC and privacy extension addresses can be done for most operating systems by sending RA messages with a hint to get addresses via DHCPv6 by setting the M-bit and disabling SLAAC by resetting all A-bits in all prefix information options. However, attackers could still find ways to bypass this mechanism if not enforced at the switch/router level.

However, in scenarios where anonymity is a strong desire (protecting user privacy is more important than user attribution), privacy extension addresses should be used. When mechanisms recommended by [RFC8064] are available, the stable privacy address is probably a good balance between privacy (among different networks) and security/user attribution (within a network).

2.1.6. DHCP Considerations

Some environments use DHCPv6 to provision addresses and other parameters in order to ensure auditability and traceability (see Section 2.6.1.5 for the limitations of DHCPv6 for auditability).

A main security concern is the ability to detect and counteract rogue DHCP servers (Section 2.3.3). It must be noted that as opposed to DHCPv4, DHCPv6 can lease several IPv6 addresses per client. For DHCPv4, the lease is bound to the 'client identifier', which may contain a hardware address, or it may contain another type of identifier, such as a DNS name. For DHCPv6, the lease is bound to the client DHCP Unique ID (DUID), which may, or may not, be bound to the client link-layer address. [RFC7824] describes the privacy issues associated with the use of DHCPv6 by Internet users. The anonymity profiles [RFC7844] are designed for clients that wish to remain anonymous to the visited network. [RFC7707] recommends that DHCPv6 servers issue addresses randomly from a large pool.

2.1.7. DNS Considerations

While the security concerns of DNS are not fundamentally different between IPv4 and IPv6, there are specific considerations in DNS64 [RFC6147] environments that need to be understood. Specifically, the interactions and the potential of interference with DNSSEC ([RFC4033]) implementation need to be understood - these are pointed out in more detail in Section 2.7.3.2.

2.1.8. Using a /64 per host

An interesting approach is using a /64 per host as proposed in [RFC8273] especially in a shared environment. This allows for easier user attribution (typically based on the host MAC address) as its /64 prefix is stable even if applications within the host can change their IPv6 address within this /64 prefix.

This can also be useful for the generation of ACLs once individual systems (e.g. admin workstations) have their own prefixes.

2.1.9. Privacy consideration of Addresses

Beside the security aspects of IPv6 addresses, there are also privacy considerations: mainly because they are of global scope and visible globally. [RFC7721] goes into more detail on the privacy considerations for IPv6 addresses by comparing the manually configured IPv6 address, DHCPv6, and SLAAC.

2.2. Extension Headers

Extension headers are an important difference between IPv4 and IPv6. In IPv4-based packets, it's trivial to find the upper-layer protocol type and protocol header, while in IPv6 it is more complex since the extension header chain must be parsed completely (even if not processed) in order to find the upper-layer protocol header. IANA has closed the existing empty "Next Header Types" registry to new entries and is redirecting its users to a new "IPv6 Extension Header Types" registry per [RFC7045].

Extension headers have also become a very controversial topic since forwarding nodes that discard packets containing extension headers are known to cause connectivity failures and deployment problems [RFC7872]. Understanding the role of various extension headers is important and this section enumerates the ones that need careful consideration.

A clarification on how intermediate nodes should handle packets with existing or future extension headers is found in [RFC7045]. The uniform TLV format to be used for defining future extension headers is described in [RFC6564]. Sections 5.2 and 5.3 of [RFC8504] provide more information on the processing of extension headers by IPv6 nodes.

Vendors of filtering solutions and operations personnel responsible for implementing packet filtering rules should be aware that the 'Next Header' field in an IPv6 header can both point to an IPv6 extension header or to an upper layer protocol header. This has to be considered when designing the user interface of filtering solutions or during the creation of filtering rule sets.

There is IETF work in progress regarding filtering rules for those extension headers: [I-D.ietf-opsec-ipv6-eh-filtering] for transit routers.

2.2.1. Order and Repetition of Extension Headers

While [RFC8200] recommends the order and the maximum repetition of extension headers, there are still IPv6 implementations, at the time of writing, which support a non-recommended order of headers (such as ESP before routing) or an illegal repetition of headers (such as multiple routing headers). The same applies for options contained in the extension headers (see [I-D.kampanakis-6man-ipv6-eh-parsing]). In some cases, it has led to nodes crashing when receiving or forwarding wrongly formatted packets.

A firewall or edge device should be used to enforce the recommended order and the maximum occurrences of extension headers by dropping non-conforming packets.

2.2.2. Hop-by-Hop Options Header

In the previous IPv6 specification [RFC2460], the hop-by-hop options header, when present in an IPv6 packet, forced all nodes to inspect and possibly process this header. This enabled denial-of-service attacks as most, if not all, routers cannot process this type of packet in hardware but have to process these packets in software and hence compete with other software tasks, such as handling the control and management plane processing.

Section 4.3 of the current Internet Standard for IPv6, [RFC8200], has taken this attack vector into account and made the processing of hop-by-hop options headers by intermediate routers explicitly configurable.

2.2.3. Fragment Header

The fragment header is used by the source (and only the source) when it has to fragment packets. [RFC7112] and section 4.5 of [RFC8200] explain why it is important that:

Firewall and security devices should drop first fragments that do not contain the entire IPv6 header chain (including the transport-layer header).

Destination nodes should discard first fragments that do not contain the entire IPv6 header chain (including the transport-layer header).

If those requirements are not met, stateless filtering could be bypassed by a hostile party. [RFC6980] applies a stricter rule to Neighbor Discovery Protocol (NDP) by enforcing the drop of fragmented NDP packets (except for "Certification Path Advertisement" messages as noted in section Section 2.3.2.1). [RFC7113] describes how the RA-guard function described in [RFC6105] should behave in the presence of fragmented RA packets.

2.2.4. IP Security Extension Header

The IPsec [RFC4301] extension headers (AH [RFC4302] and ESP [RFC4303]) are required if IPsec is to be utilized for network level security. Previously, IPv6 mandated implementation of IPsec but [RFC6434] updated that recommendation by making support of the IPsec

architecture [RFC4301] a SHOULD for all IPv6 nodes which is also retained in the latest IPv6 Nodes Requirement standard [RFC8504].

2.3. Link-Layer Security

IPv6 relies heavily on NDP [RFC4861] to perform a variety of link operations such as discovering other nodes on the link, resolving their link-layer addresses, and finding routers on the link. If not secured, NDP is vulnerable to various attacks, such as router/neighbor message spoofing, redirect attacks, Duplicate Address Detection (DAD) DoS attacks, etc. Many of these security threats to NDP have been documented in IPv6 ND Trust Models and Threats [RFC3756] and in [RFC6583].

Most of the issues are only applicable when the attacker is on the same link but NDP also has security issues when the attacker is off-link, see the section below Section 2.3.1.

2.3.1. Neighbor Solicitation Rate-Limiting

NDP can be vulnerable to remote denial of service (DoS) attacks; for example, when a router is forced to perform address resolution for a large number of unassigned addresses, i.e., when a prefix is scanned by an attacker in a fast manner. This can keep new devices from joining the network or render the last-hop router ineffective due to high CPU usage. Easy mitigative steps include rate-limiting Neighbor Solicitations, restricting the amount of state reserved for unresolved solicitations, and clever cache/timer management.

[RFC6583] discusses the potential for off-link DoS in detail and suggests implementation improvements and operational mitigation techniques that may be used to mitigate or alleviate the impact of such attacks. Here are some feasible mitigation options that can be employed by network operators today:

- o Ingress filtering of unused addresses by ACL. These require stable configuration of the addresses; for example, allocating the addresses out of a /120 and using a specific ACL to only allow traffic to this /120 (of course, the actual hosts are configured with a /64 prefix for the link).
- o Tuning of NDP process (where supported), e.g., enforcing limits on data structures such as the number of neighbor cache entries in 'incomplete' state (e.g., 256 incomplete entries per interface) or the rate of NA per interface (e.g., 100 NA per second).
- o Using a /127 on a point-to-point link, per [RFC6164].

- o Using only link-local addresses on links where there are only routers, see [RFC7404]

2.3.2. Router and Neighbor Advertisements Filtering

2.3.2.1. Router Advertisement Filtering

Router Advertisement spoofing is a well-known on-link attack vector and has been extensively documented. The presence of rogue RAs, either unintentional or malicious, can cause partial or complete failure of operation of hosts on an IPv6 link. For example, a node can select an incorrect router address which can then be used for an on-path attack or the node can assume wrong prefixes to be used for SLAAC. [RFC6104] summarizes the scenarios in which rogue RAs may be observed and presents a list of possible solutions to the problem. [RFC6105] (RA-Guard) describes a solution framework for the rogue RA problem where network segments are designed around switching devices that are capable of identifying invalid RAs and blocking them before the attack packets actually reach the target nodes.

However, several evasion techniques that circumvent the protection provided by RA-Guard have surfaced. A key challenge to this mitigation technique is introduced by IPv6 fragmentation. Attackers can conceal their attack by fragmenting their packets into multiple fragments such that the switching device that is responsible for blocking invalid RAs cannot find all the necessary information to perform packet filtering of the same packet. [RFC7113] describes such evasion techniques and provides advice to RA-Guard implementers such that the aforementioned evasion vectors can be eliminated.

Given that the IPv6 Fragmentation Header can be leveraged to circumvent some implementations of RA-Guard, [RFC6980] updates [RFC4861] such that use of the IPv6 Fragmentation Header is forbidden in all Neighbor Discovery messages except "Certification Path Advertisement", thus allowing for simple and effective measures to counter fragmented NDP attacks.

2.3.2.2. Neighbor Advertisement Filtering

The Source Address Validation Improvements (SAVI) working group has worked on other ways to mitigate the effects of such attacks. [RFC7513] helps in creating bindings between a DHCPv4 [RFC2131] /DHCPv6 [RFC8415] assigned source IP address and a binding anchor [RFC7039] on a SAVI device. Also, [RFC6620] describes how to glean similar bindings when DHCP is not used. The bindings can be used to filter packets generated on the local link with forged source IP addresses.

2.3.2.3. Host Isolation

Isolating hosts for the NDP traffic can be done by using a /64 per host, refer to Section 2.1.8, as NDP is only relevant within a /64 on-link prefix; 3GPP Section 2.3.4 uses a similar mechanism.

A more drastic technique to prevent all NDP attacks is based on isolation of all hosts with specific configurations. In such a scenario, hosts (i.e., all nodes that are not routers) are unable to send data-link layer frames to other hosts, therefore, no host-to-host attacks can happen. This specific setup can be established on some switches or Wi-Fi access points. This is not always feasible when hosts need to communicate with other hosts in the same subnet, e.g., for access to file shares.

2.3.2.4. NDP Recommendations

It is still recommended that RA-Guard and SAVI be employed as a first line of defense against common attack vectors including misconfigured hosts. This recommendation also applies when DHCPv6 is used, as RA messages are used to discover the default router(s) and for on-link prefix determination. This line of defense is most effective when incomplete fragments are dropped by routers and switches as described in Section 2.2.3. The generated log should also be analyzed to identify and act on violations.

Network operators should be aware that RA-Guard and SAVI do not work as expected or could even be harmful in specific network configurations (notably when there could be multiple routers).

Enabling RA-Guard by default in managed networks (e.g., Wi-Fi networks, enterprise campus networks, etc.) should be strongly considered except for specific use cases such as the presence of homenet devices emitting router advertisements.

2.3.3. Securing DHCP

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6), as described in [RFC8415], enables DHCP servers to pass configuration parameters, such as IPv6 network addresses and other configuration information, to IPv6 nodes. DHCP plays an important role in most large networks by providing robust stateful configuration in the context of automated system provisioning.

The two most common threats to DHCP clients come from malicious (a.k.a., rogue) or unintentionally misconfigured DHCP servers. In these scenarios, a malicious DHCP server is established with the intent of providing incorrect configuration information to the

clients to cause a denial-of-service attack or to mount on-path attack. While unintentional, a misconfigured DHCP server can have the same impact. Additional threats against DHCP are discussed in the security considerations section of [RFC8415].

DHCPv6-Shield, [RFC7610], specifies a mechanism for protecting connected DHCPv6 clients against rogue DHCPv6 servers. This mechanism is based on DHCPv6 packet-filtering at the layer-2 device, i.e., the administrator specifies the interfaces connected to DHCPv6 servers. However, extension headers could be leveraged to bypass DHCPv6-Shield unless [RFC7112] is enforced.

It is recommended to use DHCPv6-Shield and to analyze the corresponding log messages.

2.3.4. 3GPP Link-Layer Security

The 3GPP link is a point-to-point like link that has no link-layer address. This implies there can only be one end host (the mobile hand-set) and the first-hop router (i.e., a GPRS Gateway Support Node (GGSN) or a Packet Gateway (PGW)) on that link. The GGSN/PGW never configures a non link-local address on the link using the advertised /64 prefix on it; see Section 2.1.8. The advertised prefix must not be used for on-link determination. There is no need for address resolution on the 3GPP link, since there are no link-layer addresses. Furthermore, the GGSN/PGW assigns a prefix that is unique within each 3GPP link that uses IPv6 stateless address autoconfiguration. This avoids the necessity to perform DAD at the network level for every address generated by the mobile host. The GGSN/PGW always provides an IID to the cellular host for the purpose of configuring the link-local address and ensures the uniqueness of the IID on the link (i.e., no collisions between its own link-local address and the mobile host's address).

The 3GPP link model itself mitigates most of the known NDP-related Denial-of-Service attacks. In practice, the GGSN/PGW only needs to route all traffic to the mobile host that falls under the prefix assigned to it. As there is also a single host on the 3GPP link, there is no need to defend that IPv6 address.

See Section 5 of [RFC6459] for a more detailed discussion on the 3GPP link model, NDP, and the address configuration details. In some mobile networks, DHCPv6 and DHCP-PD are also used.

2.3.5. Impact of Multicast Traffic

IPv6 uses multicast extensively for signaling messages on the local link to avoid broadcast messages for on-the-wire efficiency.

The use of multicast has some side effects on wireless networks, such as a negative impact on battery life of smartphones and other battery-operated devices that are connected to such networks. [RFC7772] and [RFC6775] (for specific wireless networks) discuss methods to rate-limit RAs and other ND messages on wireless networks in order to address this issue.

The use of link-layer multicast addresses (e.g., ff02::1 for the all nodes link-local multicast address) could also be misused for an amplification attack. Imagine, a hostile node sending an ICMPv6 ECHO_REQUEST to ff02::1 with a spoofed source address, then, all link-local nodes will reply with ICMPv6 ECHO_REPLY packets to the source address. This could be a DoS attack for the address owner. This attack is purely local to the layer-2 network as packets with a link-local destination are never forwarded by an IPv6 router.

This is the reason why large Wi-Fi network deployments often limit the use of link-layer multicast either from or to the uplink of the Wi-Fi access point, i.e., Wi-Fi stations are prevented to send link-local multicast to their direct neighboring Wi-Fi stations; this policy also blocks service discovery via mDNS ([RFC6762]) and LLmNR ([RFC4795]).

2.3.6. SeND and CGA

SEcure Neighbor Discovery (SeND), as described in [RFC3971], is a mechanism that was designed to secure ND messages. This approach involves the use of new NDP options to carry public key-based signatures. Cryptographically Generated Addresses (CGA), as described in [RFC3972], are used to ensure that the sender of a Neighbor Discovery message is the actual "owner" of the claimed IPv6 address. A new NDP option, the CGA option, was introduced and is used to carry the public key and associated parameters. Another NDP option, the RSA Signature option, is used to protect all messages relating to neighbor and Router discovery.

SeND protects against:

- o Neighbor Solicitation/Advertisement Spoofing
- o Neighbor Unreachability Detection Failure
- o Duplicate Address Detection DoS Attack

- o Router Solicitation and Advertisement Attacks
- o Replay Attacks
- o Neighbor Discovery DoS Attacks

SeND does NOT:

- o Protect statically configured addresses
- o Protect addresses configured using fixed identifiers (i.e., EUI-64)
- o Provide confidentiality for NDP communications
- o Compensate for an unsecured link - SeND does not require that the addresses on the link and Neighbor Advertisements correspond.

However, at this time and over a decade since their original specifications, CGA and SeND do not have support from widely deployed IPv6 devices; hence, their usefulness is limited and should not be relied upon.

2.4. Control Plane Security

[RFC6192] defines the router control plane and provides detailed guidance to secure it for IPv4 and IPv6 networks. This definition is repeated here for the reader's convenience. Please note that the definition is completely protocol-version agnostic (most of this section applies to IPv6 in the same way as to IPv4).

Preamble: IPv6 control plane security is vastly congruent with its IPv4 equivalent with the exception of OSPFv3 authentication (Section 2.4.1) and some packet exceptions (see Section 2.4.3) that are specific to IPv6.

Modern router architecture design maintains a strict separation of forwarding and router control plane hardware and software. The router control plane supports routing and management functions. It is generally described as the router architecture hardware and software components for handling packets destined to the device itself, as well as, building and sending packets originated locally on the device. The forwarding plane is typically described as the router architecture hardware and software components responsible for receiving a packet on an incoming interface, performing a lookup to identify the packet's IP next hop and best outgoing interface towards the destination, and forwarding the packet through the appropriate outgoing interface.

While the forwarding plane is usually implemented in high-speed hardware, the control plane is implemented by a generic processor (referred to as the route processor (RP)) and cannot process packets at a high rate. Hence, this processor can be attacked by flooding its input queue with more packets than it can process. The control plane processor is then unable to process valid control packets and the router can lose IGP or BGP adjacencies which can cause a severe network disruption.

[RFC6192] provides detailed guidance to protect the router control plane in IPv6 networks. The rest of this section contains simplified guidance.

The mitigation techniques are:

- o To drop non-legit or potentially harmful control packets before they are queued to the RP (this can be done by a forwarding plane ACL) and
- o To rate-limit the remaining packets to a rate that the RP can sustain. Protocol-specific protection should also be done (for example, a spoofed OSPFv3 packet could trigger the execution of the Dijkstra algorithm, therefore, the frequency of Dijkstra calculations should be also rate-limited).

This section will consider several classes of control packets:

- o Control protocols: routing protocols: such as OSPFv3, BGP, RIPng, and by extension NDP and ICMP
- o Management protocols: SSH, SNMP, NETCONF, RESTCONF, IPFIX, etc.
- o Packet exceptions: normal data packets that require a specific processing such as generating a packet-too-big ICMP message or processing the hop-by-hop options header.

2.4.1. Control Protocols

This class includes OSPFv3, BGP, NDP, ICMP.

An ingress ACL to be applied on all the router interfaces for packets to be processed by the RP should be configured to:

- o drop OSPFv3 (identified by Next-Header being 89) and RIPng (identified by UDP port 521) packets from a non link-local address (except for OSPFv3 virtual links)

- o allow BGP (identified by TCP port 179) packets from all BGP neighbors and drop the others
- o allow all ICMP packets (transit and to the router interfaces)

Note: dropping OSPFv3 packets which are authenticated by IPsec could be impossible on some routers that are unable to parse the IPsec ESP or AH extension headers during ACL classification.

Rate-limiting of the valid packets should be done, see also [RFC8541] for a side benefit for OSPv3. The exact configuration will depend on the available resources of the router (CPU, TCAM, ...).

2.4.2. Management Protocols

This class includes: SSH, SNMP, RESTCONF, NETCONF, gRPC, syslog, NTP, etc.

An ingress ACL to be applied on all the router interfaces (or at ingress interfaces of the security perimeter or by using specific features of the platform) should be configured for packets destined to the RP such as:

- o Drop packets destined to the routers except those belonging to protocols which are used (for example, permit TCP 22 and drop all others when only SSH is used);
- o Drop packets where the source does not match the security policy, for example, if SSH connections should only be originated from the Network Operation Center (NOC), then the ACL should permit TCP port 22 packets only from the NOC prefix.

Rate-limiting of valid packets should be done. The exact configuration will depend on the available router resources.

2.4.3. Packet Exceptions

This class covers multiple cases where a data plane packet is punted to the route processor because it requires specific processing:

- o generation of an ICMP packet-too-big message when a data plane packet cannot be forwarded because it is too large (required to discover the Path MTU);
- o generation of an ICMP hop-limit-expired message when a data plane packet cannot be forwarded because its hop-limit field has reached 0 (also used by the traceroute utility);

- o generation of an ICMP destination-unreachable message when a data plane packet cannot be forwarded for any reason;
- o processing of the hop-by-hop options header, new implementations follow section 4.3 of [RFC8200] where this processing is optional;
- o or more specific to some router implementation: an oversized extension header chain which cannot be processed by the hardware and force the packet to be punted to the RP.

On some routers, not everything can be done by the specialized data plane hardware which requires some packets to be 'punted' to the generic RP. This could include for example the processing of a long extension header chain in order to apply an ACL based on layer-4 information. [RFC6980] and more generally [RFC7112] highlight the security implications of oversized extension header chains on routers and updates the original IPv6 specifications, [RFC2460], such that the first fragment of a packet is required to contain the entire IPv6 header chain. Those changes are incorporated in the IPv6 standard [RFC8200]

An ingress ACL cannot mitigate a control plane attack using these packet exceptions. The only protection for the RP is to rate-limit those packet exceptions that are forwarded to the RP, this means that some data plane packets will be dropped without an ICMP message sent to the source which may delay Path MTU discovery and cause drops.

In addition to limiting the rate of data plane packets queued to the RP, it is also important to rate-limit the generation of ICMP messages. This is important both to preserve RP resources and also to prevent an amplification attack using the router as a reflector. It is worth noting that some platforms implement this rate-limiting in hardware. Of course, a consequence of not generating an ICMP message will break some IPv6 mechanisms such as Path MTU discovery or a simple traceroute.

2.5. Routing Security

Preamble: IPv6 routing security is congruent with IPv4 routing security with the exception of OSPv3 neighbor authentication (see Section 2.5.2).

Routing security in general can be broadly divided into three sections:

1. Authenticating neighbors/peers
2. Securing routing updates between peers

3. Route filtering

[RFC5082] is also applicable to IPv6 and can ensure that routing protocol packets are coming from the local network; it must also be noted that in IPv6 all interior gateway protocols use link-local addresses.

As for IPv4, it is recommended to enable a routing protocol only on interfaces where it is required.

2.5.1. BGP Security

As BGP is identical for IPv4 and IPv6 and as [RFC7454] covers all the security aspects for BGP in detail, [RFC7454] is also applicable to IPv6.

2.5.2. Authenticating OSPFv3 Neighbors

OSPFv3 can rely on IPsec to fulfill the authentication function. Operators should note that IPsec support is not standard on all routing platforms. In some cases, this requires specialized hardware that offloads crypto over to dedicated ASICs or enhanced software images (both of which often come with added financial cost) to provide such functionality. An added detail is to determine whether OSPFv3 IPsec implementations use AH or ESP-Null for integrity protection. In early implementations, all OSPFv3 IPsec configurations relied on AH since the details weren't specified in [RFC5340]. However, the document which specifically describes how IPsec should be implemented for OSPFv3 [RFC4552] specifically states that "ESP-Null MUST and AH MAY be implemented" since it follows the overall IPsec standards wording. OSPFv3 can also use normal ESP to encrypt the OSPFv3 payload to provide confidentiality for the routing information.

[RFC7166] changes OSPFv3 reliance on IPsec by appending an authentication trailer to the end of the OSPFv3 packets; it does not specifically authenticate the specific originator of an OSPFv3 packet; rather, it allows a router to confirm that the packet has been issued by a router that had access to the shared authentication key.

With all authentication mechanisms, operators should confirm that implementations can support re-keying mechanisms that do not cause outages. There have been instances where any re-keying causes outages and therefore, the tradeoff between utilizing this functionality needs to be weighed against the protection it provides. [RFC4107] documents some guidelines for crypto keys management.

2.5.3. Securing Routing Updates

IPv6 initially mandated the provisioning of IPsec capability in all nodes. However, in the updated IPv6 Nodes Requirement standard [RFC8504], IPsec is a 'SHOULD' and not a 'MUST' implement. Theoretically, it is possible that all communication between two IPv6 nodes, especially routers exchanging routing information, is encrypted using IPsec. In practice however, deploying IPsec is not always feasible given hardware and software limitations of the various platforms deployed.

Many routing protocols support the use of cryptography to protect the routing updates, the use of this protection is recommended; [RFC8177] is a YANG data model for key chains that includes re-keying functionality.

2.5.4. Route Filtering

Route filtering policies will be different depending on whether they pertain to edge route filtering vs. internal route filtering. At a minimum, IPv6 routing policy as it pertains to routing between different administrative domains should aim to maintain parity with IPv4 from a policy perspective, e.g.,

- o Filter internal-use, non-globally routable IPv6 addresses at the perimeter;
- o Discard routes for bogon [CYMRU] and reserved space (see [RFC8190]);
- o Configure ingress route filters that validate route origin, prefix ownership, etc. through the use of various routing databases, e.g., [RADB]. [RFC8210] formally validates the origin ASs of BGP announcements.

Some good guidance can be found at [RFC7454].

A valid routing table can also be used to apply network ingress filtering (see [RFC2827]).

2.6. Logging/Monitoring

In order to perform forensic research in the cases of a security incident or detecting abnormal behavior, network operators should log multiple pieces of information. In some cases, this requires a frequent poll of devices via a Network Management Station.

This logging should include, but not limited to:

- o logs of all applications using the network (including user space and kernel space) when available (for example web servers that the network operator manages);
- o data from IP Flow Information Export [RFC7011] also known as IPFIX;
- o data from various SNMP MIBs [RFC4293] or YANG data via RESTCONF [RFC8040] or NETCONF [RFC6241];
- o historical data of Neighbor Cache entries;
- o stateful DHCPv6 [RFC8415] lease cache, especially when a relay agent [RFC6221] is used;
- o Source Address Validation Improvement (SAVI) [RFC7039] events, especially the binding of an IPv6 address to a MAC address and a specific switch or router interface;
- o firewall ACL log;
- o authentication server log;
- o RADIUS [RFC2866] accounting records.

Please note that there are privacy issues or regulations related to how these logs are collected, stored, used, and safely discarded. Operators are urged to check their country legislation (e.g., General Data Protection Regulation GDPR [GDPR] in the European Union).

All those pieces of information can be used for:

- o forensic (Section 2.6.2.1) investigations such as who did what and when?
- o correlation (Section 2.6.2.3): which IP addresses were used by a specific node (assuming the use of privacy extensions addresses [RFC8981])
- o inventory (Section 2.6.2.2): which IPv6 nodes are on my network?
- o abnormal behavior detection (Section 2.6.2.4): unusual traffic patterns are often the symptoms of an abnormal behavior which is in turn a potential attack (denial-of-service, network scan, a node being part of a botnet, etc.)

2.6.1. Data Sources

This section lists the most important sources of data that are useful for operational security.

2.6.1.1. Application Logs

Those logs are usually text files where the remote IPv6 address is stored in clear text (not binary). This can complicate the processing since one IPv6 address, for example 2001:db8::1 can be written in multiple ways, such as:

- o 2001:DB8::1 (in uppercase)
- o 2001:0db8::0001 (with leading 0)
- o and many other ways including the reverse DNS mapping into a FQDN (which should not be trusted).

[RFC5952] explains this problem in detail and recommends the use of a single canonical format. This document recommends the use of canonical format [RFC5952] for IPv6 addresses in all possible cases. If the existing application cannot log using the canonical format, then it is recommended to use an external post-processing program in order to canonicalize all IPv6 addresses.

2.6.1.2. IP Flow Information Export by IPv6 Routers

IPFIX [RFC7012] defines some data elements that are useful for security:

- o nextHeaderIPv6, sourceIPv6Address, and destinationIPv6Address;
- o sourceMacAddress and destinationMacAddress.

The IP version is the ipVersion element defined in [IANA-IPFIX].

Moreover, IPFIX is very efficient in terms of data handling and transport. It can also aggregate flows by a key such as sourceMacAddress in order to have aggregated data associated with a specific sourceMacAddress. This memo recommends the use of IPFIX and aggregation on nextHeaderIPv6, sourceIPv6Address, and sourceMacAddress.

2.6.1.3. SNMP MIB and NETCONF/RESTCONF YANG Modules data by IPv6 Routers

RFC 4293 [RFC4293] defines a Management Information Base (MIB) for the two address families of IP. This memo recommends the use of:

- o ipIfStatsTable table which collects traffic counters per interface;
- o ipNetToPhysicalTable table which is the content of the Neighbor cache, i.e., the mapping between IPv6 and data-link layer addresses.

There are also YANG modules relating to the two IP addresses families and can be used with [RFC6241] and [RFC8040]. This memo recommends the use of:

- o interfaces-state/interface/statistics from ietf-interfaces@2018-02-20.yang [RFC8343] which contains counters for interfaces.
- o ipv6/neighbor from ietf-ip@2018-02-22.yang [RFC8344] which is the content of the Neighbor cache, i.e., the mapping between IPv6 and data-link layer addresses.

2.6.1.4. Neighbor Cache of IPv6 Routers

The neighbor cache of routers contains all mappings between IPv6 addresses and data-link layer addresses. There are multiple ways to collect the current entries in the Neighbor Cache, notably but not limited to:

- o the SNMP MIB (Section 2.6.1.3) as explained above;
- o using streaming telemetry or NETCONF [RFC6241] and RESTCONF [RFC8040] to collect the operational state of the neighbor cache;
- o also, by connecting over a secure management channel (such as SSH) and explicitly requesting a neighbor cache dump via the Command Line Interface (CLI) or another monitoring mechanism.

The neighbor cache is highly dynamic as mappings are added when a new IPv6 address appears on the network. This could be quite frequently with privacy extension addresses [RFC8981] or when they are removed when the state goes from UNREACH to removed (the default time for a removal per Neighbor Unreachability Detection [RFC4861] algorithm is 38 seconds for a host using Windows 7). This means that the content of the neighbor cache must periodically be fetched at an interval

which does not exhaust the router resources and still provides valuable information (suggested value is 30 seconds but this should be verified in the actual deployment) and stored for later use.

This is an important source of information because it is trivial (on a switch not using the SAVI [RFC7039] algorithm) to defeat the mapping between data-link layer address and IPv6 address. Let us rephrase the previous statement: having access to the current and past content of the neighbor cache has a paramount value for the forensic and audit trail. It should also be noted that in certain threat models this information is also deemed valuable and could itself be a target.

When using one /64 per host (Section 2.1.8) or DHCP-PD, it is sufficient to keep the history of the allocated prefixes when combined with strict source address prefix enforcement on the routers and layer-2 switches to prevent IPv6 spoofing.

2.6.1.5. Stateful DHCPv6 Lease

In some networks, IPv6 addresses/prefixes are managed by a stateful DHCPv6 server [RFC8415] that leases IPv6 addresses/prefixes to clients. It is indeed quite similar to DHCP for IPv4, so it can be tempting to use this DHCP lease file to discover the mapping between IPv6 addresses/prefixes and data-link layer addresses as is commonly used in IPv4 networking.

It is not so easy in the IPv6 networks, because not all nodes will use DHCPv6 (there are nodes which can only do stateless autoconfiguration) but also because DHCPv6 clients are identified not by their hardware-client address as in IPv4 but by a DHCP Unique ID (DUID), which can have several formats: some being the data-link layer address, some being data-link layer address prepended with time information, or even an opaque number that requires correlation with another data source to be usable for operational security. Moreover, when the DUID is based on the data-link address, this address can be of any client interface (such as the wireless interface while the client actually uses its wired interface to connect to the network).

If a lightweight DHCP relay agent [RFC6221] is used in a layer-2 switch, then the DHCP servers also receive the Interface-ID information which could be saved in order to identify the interface on which the switch received a specific leased IPv6 address. Also, if a 'normal' (not lightweight) relay agent adds the data-link layer address in the option for Relay Agent Remote-ID [RFC4649] or [RFC6939], then the DHCPv6 server can keep track of the data-link and leased IPv6 addresses.

In short, the DHCPv6 lease file is less interesting than for IPv4 networks. If possible, it is recommended to use DHCPv6 servers that keep the relayed data-link layer address in addition to the DUID in the lease file as those servers have the equivalent information to IPv4 DHCP servers.

The mapping between data-link layer address and the IPv6 address can be secured by deploying switches implementing the SAVI [RFC7513] mechanisms. Of course, this also requires that the data-link layer address is protected by using a layer-2 mechanism such as [IEEE-802.1X].

2.6.1.6. RADIUS Accounting Log

For interfaces where the user is authenticated via a RADIUS [RFC2866] server, and if RADIUS accounting is enabled, then the RADIUS server receives accounting Acct-Status-Type records at the start and at the end of the connection which include all IPv6 (and IPv4) addresses used by the user. This technique can be used notably for Wi-Fi networks with Wi-Fi Protected Address (WPA) or other IEEE 802.1X [IEEE-802.1X] wired interface on an Ethernet switch.

2.6.1.7. Other Data Sources

There are other data sources for log information that must be collected (as currently collected in IPv4 networks):

- o historical mapping of IPv6 addresses to users of remote access VPN;
- o historical mappings of MAC addresses to switch ports in a wired network.

2.6.2. Use of Collected Data

This section leverages the data collected as described before (Section 2.6.1) in order to achieve several security benefits. Section 9.1 of [RFC7934] contains more details about host tracking.

2.6.2.1. Forensic and User Accountability

The forensic use case is when the network operator must locate an IPv6 address (and the associated port, access point/switch, or VPN tunnel) that was present in the network at a certain time or is currently in the network.

To locate an IPv6 address in an enterprise network where the operator has control over all resources, the source of information can be the

neighbor cache, or, if not found, the DHCP lease file. Then, the procedure is:

1. Based on the IPv6 prefix of the IPv6 address, find the router(s) which is(are) used to reach this prefix (assuming that anti-spoofing mechanisms are used) perhaps based on an IPAM.
2. Based on this limited set of routers, on the incident time and on the IPv6 address, retrieve the data-link address from the live neighbor cache, from the historical neighbor cache data, or from SAVI events, or retrieve the data-link address from the DHCP lease file (Section 2.6.1.5).
3. Based on the data-link layer address, look-up the switch interface associated with the data-link layer address. In the case of wireless LAN with RADIUS accounting (see Section 2.6.1.6), the RADIUS log has the mapping between the user identification and the MAC address. If a Configuration Management Data Base (CMDB) is used, then it can be used to map the data-link layer address to a switch port.

At the end of the process, the interface of the host originating, or the subscriber identity associated with, the activity in question has been determined.

To identify the subscriber of an IPv6 address in a residential Internet Service Provider, the starting point is the DHCP-PD leased prefix covering the IPv6 address; this prefix can often be linked to a subscriber via the RADIUS log. Alternatively, the Forwarding Information Base (FIB) of the Cable Modem Termination System (CMTS) or Broadband Network Gateway (BNG) indicates the CPE of the subscriber and the RADIUS log can be used to retrieve the actual subscriber.

More generally, a mix of the above techniques can be used in most, if not all, networks.

2.6.2.2. Inventory

RFC 7707 [RFC7707] describes the difficulties for an attacker to scan an IPv6 network due to the vast number of IPv6 addresses per link (and why in some cases it can still be done). While the huge addressing space can sometimes be perceived as a 'protection', it also makes the inventory task difficult in an IPv6 network while it was trivial to do in an IPv4 network (a simple enumeration of all IPv4 addresses, followed by a ping and a TCP/UDP port scan). Getting an inventory of all connected devices is of prime importance for a secure network operation.

There are many ways to do an inventory of an IPv6 network.

The first technique is to use passive inspection such as IPFIX. Using exported IPFIX information and extracting the list of all IPv6 source addresses allows finding all IPv6 nodes that sent packets through a router. This is very efficient but, alas, will not discover silent nodes that never transmitted packets traversing the IPFIX target router. Also, it must be noted that link-local addresses will never be discovered by this means.

The second way is again to use the collected neighbor cache content to find all IPv6 addresses in the cache. This process will also discover all link-local addresses. See Section 2.6.1.4.

Another way that works only for a local network, consists of sending a ICMP ECHO_REQUEST to the link-local multicast address ff02::1 which addresses all IPv6 nodes on the network. All nodes should reply to this ECHO_REQUEST per [RFC4443].

Other techniques involve obtaining data from DNS, parsing log files, leveraging service discovery such as mDNS [RFC6762] and [RFC6763].

Enumerating DNS zones, especially looking at reverse DNS records and CNAMEs, is another common method employed by various tools. As already mentioned in [RFC7707], this allows an attacker to prune the IPv6 reverse DNS tree, and hence enumerate it in a feasible time. Furthermore, authoritative servers that allow zone transfers (AXFR) may be a further information source. An interesting research paper has analysed the entropy in various IPv6 addresses: see [ENTROPYIP].

2.6.2.3. Correlation

In an IPv4 network, it is easy to correlate multiple logs, for example to find events related to a specific IPv4 address. A simple Unix grep command is enough to scan through multiple text-based files and extract all lines relevant to a specific IPv4 address.

In an IPv6 network, this is slightly more difficult because different character strings can express the same IPv6 address. Therefore, the simple Unix grep command cannot be used. Moreover, an IPv6 node can have multiple IPv6 addresses.

In order to do correlation in IPv6-related logs, it is advised to have all logs in a format with only canonical IPv6 addresses [RFC5952]. Then, the neighbor cache current (or historical) data set must be searched to find the data-link layer address of the IPv6 address. Then, the current and historical neighbor cache data sets must be searched for all IPv6 addresses associated with this data-

link layer address to derive the search set. The last step is to search in all log files (containing only IPv6 addresses in canonical format) for any IPv6 addresses in the search set.

Moreover, [RFC7934] recommends using multiple IPv6 addresses per prefix, so, the correlation must also be done among those multiple IPv6 addresses, for example by discovering in the NDP cache (Section 2.6.1.4) all IPv6 addresses associated with the same MAC address and interface.

2.6.2.4. Abnormal Behavior Detection

Abnormal behavior (such as network scanning, spamming, denial-of-service) can be detected in the same way as in an IPv4 network.

- o Sudden increase of traffic detected by interface counter (SNMP) or by aggregated traffic from IPFIX records [RFC7012].
- o Rapid growth of ND cache size.
- o Change in traffic pattern (number of connections per second, number of connections per host...) observed with the use of IPFIX [RFC7012].

2.6.3. Summary

While some data sources (IPFIX, MIB, switch CAM tables, logs, ...) used in IPv4 are also used in the secure operation of an IPv6 network, the DHCPv6 lease file is less reliable and the neighbor cache is of prime importance.

The fact that there are multiple ways to express the same IPv6 address in a character string renders the use of filters mandatory when correlation must be done.

2.7. Transition/Coexistence Technologies

As it is expected that some networks will not run in a pure IPv6-only mode, the different transition mechanisms must be deployed and operated in a secure way. This section proposes operational guidelines for the most known and deployed transition techniques. [RFC4942] also contains security considerations for transition or coexistence scenarios.

2.7.1. Dual Stack

Dual stack is often the first deployment choice for network operators. Dual stacking the network offers some advantages over other transition mechanisms. Firstly, the impact on existing IPv4 operations is reduced. Secondly, in the absence of tunnels or address translation, the IPv4 and IPv6 traffic are native (easier to observe and secure) and should have the same network processing (network path, quality of service, ...). Dual stack enables a gradual termination of the IPv4 operations when the IPv6 network is ready for prime time. On the other hand, the operators have to manage two network stacks with the added complexities.

From an operational security perspective, this now means that the network operator has twice the exposure. One needs to think about protecting both protocols now. At a minimum, the IPv6 portion of a dual-stacked network should be consistent with IPv4 from a security policy point of view. Typically, the following methods are employed to protect IPv4 networks at the edge or security perimeter:

- o ACLs to permit or deny traffic;
- o Firewalls with stateful packet inspection;
- o Application firewalls inspecting the application flows.

It is recommended that these ACLs and/or firewalls be additionally configured to protect IPv6 communications. The enforced IPv6 security must be congruent with the IPv4 security policy, otherwise the attacker will use the protocol version having the more relaxed security policy. Maintaining the congruence between security policies can be challenging (especially over time); it is recommended to use a firewall or an ACL manager that is dual-stack, i.e., a system that can apply a single ACL entry to a mixed group of IPv4 and IPv6 addresses.

Application firewalls work at the application layer and are oblivious to the IP version, i.e., they work as well for IPv6 as for IPv4 and the same application security policy will work for both protocol versions.

Also, given the end-to-end connectivity that IPv6 provides, it is recommended that hosts be fortified against threats. General device hardening guidelines are provided in Section 2.8.

For many years, all host operating systems have IPv6 enabled by default, so, it is possible even in an 'IPv4-only' network to attack layer-2 adjacent victims via their IPv6 link-local address or via a

global IPv6 address when the attacker provides rogue RAs or a rogue DHCPv6 service.

[RFC7123] discusses the security implications of native IPv6 support and IPv6 transition/coexistence technologies on "IPv4-only" networks and describes possible mitigations for the aforementioned issues.

2.7.2. Encapsulation Mechanisms

There are many tunnels used for specific use cases. Except when protected by IPsec [RFC4301] or alternative tunnel encryption methods, all those tunnels have a number of security issues as described in RFC 6169 [RFC6169];

- o tunnel injection: a malevolent actor knowing a few pieces of information (for example the tunnel endpoints and the encapsulation protocol) can forge a packet which looks like a legitimate and valid encapsulated packet that will gladly be accepted by the destination tunnel endpoint. This is a specific case of spoofing;
- o traffic interception: no confidentiality is provided by the tunnel protocols (without the use of IPsec or alternative encryption methods), therefore anybody on the tunnel path can intercept the traffic and have access to the clear-text IPv6 packet; combined with the absence of authentication, an on-path attack can also be mounted;
- o service theft: as there is no authorization, even a non-authorized user can use a tunnel relay for free (this is a specific case of tunnel injection);
- o reflection attack: another specific use case of tunnel injection where the attacker injects packets with an IPv4 destination address not matching the IPv6 address causing the first tunnel endpoint to re-encapsulate the packet to the destination... Hence, the final IPv4 destination will not see the original IPv4 address but only the IPv4 address of the relay router.
- o bypassing security policy: if a firewall or an Intrusion Prevention System (IPS) is on the path of the tunnel, then it may neither inspect nor detect malevolent IPv6 traffic transmitted over the tunnel.

To mitigate the bypassing of security policies, it is often recommended to block all automatic tunnels in default OS configuration (if they are not required) by denying IPv4 packets matching:

- o IP protocol 41: this will block ISATAP (Section 2.7.2.2), 6to4 (Section 2.7.2.7), 6rd (Section 2.7.2.3), as well as, 6in4 (Section 2.7.2.1) tunnels;
- o IP protocol 47: this will block GRE (Section 2.7.2.1) tunnels;
- o UDP port 3544: this will block the default encapsulation of Teredo (Section 2.7.2.8) tunnels.

Ingress filtering [RFC2827] should also be applied on all tunnel endpoints if applicable to prevent IPv6 address spoofing.

The reflection attack cited above should also be prevented by using an IPv6 ACL preventing the hair pinning of the traffic.

As several of the tunnel techniques share the same encapsulation (i.e., IPv4 protocol 41) and embed the IPv4 address in the IPv6 address, there are a set of well-known looping attacks described in RFC 6324 [RFC6324]. This RFC also proposes mitigation techniques.

2.7.2.1. Site-to-Site Static Tunnels

Site-to-site static tunnels are described in RFC 2529 [RFC2529] and in GRE [RFC2784]. As the IPv4 endpoints are statically configured and are not dynamic, they are slightly more secure (bi-directional service theft is mostly impossible) but traffic interception and tunnel injection are still possible. Therefore, the use of IPsec [RFC4301] in transport mode to protect the encapsulated IPv4 packets is recommended for those tunnels. Alternatively, IPsec in tunnel mode can be used to transport IPv6 traffic over a non-trusted IPv4 network.

2.7.2.2. ISATAP

ISATAP tunnels [RFC5214] are mainly used within a single administrative domain and to connect a single IPv6 host to the IPv6 network. This often implies that those systems are usually managed by a single entity; therefore, audit trail and strict anti-spoofing are usually possible and this raises the overall security. Even if ISATAP is no more often used, its security issues are relevant per [KRISTOFF].

Special care must be taken to avoid a looping attack by implementing the measures of [RFC6324] and [RFC6964] (especially the section 3.6).

IPsec [RFC4301] in transport or tunnel mode can be used to secure the IPv4 ISATAP traffic to provide IPv6 traffic confidentiality and prevent service theft.

2.7.2.3. 6rd

While 6rd tunnels share the same encapsulation as 6to4 tunnels (Section 2.7.2.7), they are designed to be used within a single SP domain, in other words, they are deployed in a more constrained environment (e.g., anti-spoofing, protocol 41 filtering at the edge) than 6to4 tunnels and have few security issues other than lack of confidentiality. The security considerations (Section 12) of [RFC5969] describes how to secure 6rd tunnels.

IPsec [RFC4301] for the transported IPv6 traffic can be used if confidentiality is important.

2.7.2.4. 6PE, 6VPE, and LDPv6

Organizations using MPLS in their core can also use 6PE [RFC4798] and 6VPE [RFC4659] to enable IPv6 access over MPLS. As 6PE and 6VPE are really similar to BGP/MPLS IP VPNs described in [RFC4364], the security properties of these networks are also similar to those described in [RFC4381] (please note that this RFC may resemble a published IETF work but it is not based on an IETF review and the IETF disclaims any knowledge of the fitness of this RFC for any purpose). They rely on:

- o Address space, routing, and traffic separation with the help of VRFs (only applicable to 6VPE);
- o Hiding the IPv4 core, hence removing all attacks against P-routers;
- o Securing the routing protocol between CE and PE; in the case of 6PE and 6VPE, link-local addresses (see [RFC7404]) can be used and as these addresses cannot be reached from outside of the link, the security of 6PE and 6VPE is even higher than an IPv4 BGP/MPLS IP VPN.

LDPv6 itself does not induce new risks, see also [RFC7552].

2.7.2.5. DS-Lite

DS-lite is also a translation mechanism and is therefore analyzed further (Section 2.7.3.3) in this document as it includes IPv4 NAT.

2.7.2.6. Mapping of Address and Port

With the encapsulation and translation versions of mapping of Address and Port (MAP) (MAP-E [RFC7597] and MAP-T [RFC7599]), the access network is purely an IPv6 network and MAP protocols are used to

provide IPv4 hosts on the subscriber network access to IPv4 hosts on the Internet. The subscriber router does stateful operations in order to map all internal IPv4 addresses and layer-4 ports to the IPv4 address and the set of layer-4 ports received through the MAP configuration process. The SP equipment always does stateless operations (either decapsulation or stateless translation). Therefore, as opposed to Section 2.7.3.3, there is no state-exhaustion DoS attack against the SP equipment because there is no state and there is no operation caused by a new layer-4 connection (no logging operation).

The SP MAP equipment should implement all the security considerations of [RFC7597]; notably, ensuring that the mapping of the IPv4 address and port are consistent with the configuration. As MAP has a predictable IPv4 address and port mapping, the audit logs are easier to use as there is a clear mapping between the IPv6 address and the IPv4 address and ports.

2.7.2.7. 6to4

In [RFC3056]; 6to4 tunnels require a public routable IPv4 address in order to work correctly. They can be used to provide either single IPv6 host connectivity to the IPv6 Internet or multiple IPv6 networks connectivity to the IPv6 Internet. The 6to4 relay was historically the anycast address defined in [RFC3068] which has been deprecated by [RFC7526] and is no longer used by recent Operating Systems. Some security considerations are explained in [RFC3964].

[RFC6343] points out that if an operator provides well-managed servers and relays for 6to4, non-encapsulated IPv6 packets will pass through well-defined points (the native IPv6 interfaces of those servers and relays) at which security mechanisms may be applied. Client usage of 6to4 by default is now discouraged, and significant precautions are needed to avoid operational problems.

2.7.2.8. Teredo

Teredo tunnels [RFC4380] are mainly used in a residential environment because Teredo easily traverses an IPv4 NAT device thanks to its UDP encapsulation. Teredo tunnels connect a single host to the IPv6 Internet. Teredo shares the same issues as other tunnels: no authentication, no confidentiality, possible spoofing and reflection attacks.

IPsec [RFC4301] for the transported IPv6 traffic is recommended.

The biggest threat to Teredo is probably for an IPv4-only network as Teredo has been designed to easily traverse IPv4 NAT-PT devices which

are quite often co-located with a stateful firewall. Therefore, if the stateful IPv4 firewall allows unrestricted UDP outbound and accepts the return UDP traffic, then Teredo actually punches a hole in this firewall for all IPv6 traffic to the Internet and from the Internet. Host policies can be deployed to block Teredo in an IPv4-only network in order to avoid this firewall bypass. On the IPv4 firewall all outbound UDP should be blocked except for the commonly used services (e.g., port 53 for DNS, port 123 for NTP, port 443 for QUIC, port 500 for IKE, port 3478 for STUN, etc.).

Teredo is now hardly ever used and no longer enabled by default in most environments, so it is less of a threat, however, special consideration must be taken in cases when devices with older or non-updated operating systems may be present and by default were running Teredo.

2.7.3. Translation Mechanisms

Translation mechanisms between IPv4 and IPv6 networks are alternate coexistence strategies while networks transition to IPv6. While a framework is described in [RFC6144], the specific security considerations are documented with each individual mechanism. For the most part, they specifically mention interference with IPsec or DNSSEC deployments, how to mitigate spoofed traffic, and what some effective filtering strategies may be.

While not really a transition mechanism to IPv6, this section also includes the discussion about the use of heavy IPv4-to-IPv4 network address and port translation to prolong the life of IPv4-only networks.

2.7.3.1. Carrier-Grade NAT (CGN)

Carrier-Grade NAT (CGN), also called NAT444 CGN or Large Scale NAT (LSN) or SP NAT is described in [RFC6264] and is utilized as an interim measure to extend the use of IPv4 in a large service provider network until the provider can deploy an effective IPv6 solution. [RFC6598] requested a specific IANA allocated /10 IPv4 address block to be used as address space shared by all access networks using CGN. This has been allocated as 100.64.0.0/10.

Section 13 of [RFC6269] lists some specific security-related issues caused by large scale address sharing. The Security Considerations section of [RFC6598] also lists some specific mitigation techniques for potential misuse of shared address space. Some Law Enforcement Agencies have identified CGN as impeding their cyber-crime investigations (for example Europol press release on CGN [europol-cgn]). Many translation techniques (NAT64, DS-lite, ...)

have the same security issues as CGN when one part of the connection is IPv4-only.

[RFC6302] has recommendations for Internet-facing servers to also log the source TCP or UDP ports of incoming connections in an attempt to help identify the users behind such a CGN.

[RFC7422] suggests the use of deterministic address mapping in order to reduce logging requirements for CGN. The idea is to have a known algorithm for mapping the internal subscriber to/from public TCP and UDP ports.

[RFC6888] lists common requirements for CGNs. [RFC6967] analyzes some solutions to enforce policies on misbehaving nodes when address sharing is used. [RFC7857] also updates the NAT behavioral requirements.

2.7.3.2. NAT64/DNS64 and 464XLAT

Stateful NAT64 translation [RFC6146] allows IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP. It can be used in conjunction with DNS64 [RFC6147], a mechanism which synthesizes AAAA records from existing A records. There is also a stateless NAT64 [RFC7915], which has similar security aspects but with the added benefit of being stateless, so, less prone to a state exhaustion attack.

The Security Consideration sections of [RFC6146] and [RFC6147] list the comprehensive issues; in section 8 of [RFC6147] there are some considerations on the interaction between NAT64 and DNSSEC. A specific issue with the use of NAT64 is that it will interfere with most IPsec deployments unless UDP encapsulation is used.

Another translation mechanism relying on a combination of stateful and stateless translation, 464XLAT [RFC6877], can be used to do host local translation from IPv4 to IPv6 and a network provider translation from IPv6 to IPv4, i.e., giving IPv4-only application access to an IPv4-only server over an IPv6-only network. 464XLAT shares the same security considerations as NAT64 and DNS64, however it can be used without DNS64, avoiding the DNSSEC implications.

2.7.3.3. DS-Lite

Dual-Stack Lite (DS-Lite) [RFC6333] is a transition technique that enables a service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and IPv4 NAPT.

Security considerations with respect to DS-Lite mainly revolve around logging data, preventing DoS attacks from rogue devices (as the Address Family Translation Router (AFTR) [RFC6333] function is stateful) and restricting service offered by the AFTR only to registered customers.

Section 11 of [RFC6333] and section 2 of [RFC7785] describe important security issues associated with this technology.

2.8. General Device Hardening

With almost all devices being IPv6 enabled by default and with many end points having IPv6 connectivity to the Internet, it is critical to also harden those devices against attacks over IPv6.

The same techniques used to protect devices against attack over IPv4 should be used for IPv6 and should include, but not limited to:

- o Restrict device access to authorized individuals
- o Monitor and audit access to the device
- o Turn off any unused services on the end node
- o Understand which IPv6 addresses are being used to source traffic and change defaults if necessary
- o Use cryptographically protected protocols for device management (SCP, SNMPv3, SSH, TLS, etc.)
- o Use host firewall capabilities to control traffic that gets processed by upper-layer protocols
- o apply firmware, OS and application patches/upgrades to the devices in a timely manner
- o use multi-factor credentials to authenticate to devices
- o Use virus scanners to detect malicious programs

3. Enterprises Specific Security Considerations

Enterprises [RFC7381] generally have robust network security policies in place to protect existing IPv4 networks. These policies have been distilled from years of experiential knowledge of securing IPv4 networks. At the very least, it is recommended that enterprise networks have parity between their security policies for both protocol versions. This section also applies to the enterprise part

of all SP networks, i.e., the part of the network where the SP employees are connected.

Security considerations in the enterprise can be broadly categorized into two groups: External and Internal.

3.1. External Security Considerations

The external aspect deals with providing security at the edge or perimeter of the enterprise network where it meets the service provider's network. This is commonly achieved by enforcing a security policy either by implementing dedicated firewalls with stateful packet inspection or a router with ACLs. A common default IPv4 policy on firewalls that could easily be ported to IPv6 is to allow all traffic outbound while only allowing specific traffic, such as established sessions, inbound (see also [RFC6092]). Section 3.2 of [RFC7381] also provides similar recommendations.

Here are a few more things that could enhance the default policy:

- o Filter internal-use IPv6 addresses at the perimeter, this will also mitigate the vulnerabilities listed in [RFC7359]
- o Discard packets from and to bogon and reserved space, see also [CYMRU] and [RFC8190]
- o Accept certain ICMPv6 messages to allow proper operation of ND and PMTUD, see also [RFC4890] or [REY_PF] for hosts
- o Based on the use of the network, filter specific extension headers by accepting only the required ones (permit list approach) such as ESP, AH, and not forgetting the required transport layers: ICMP, TCP, UDP, ... This filtering should be done where applicable at the edge and possibly inside the perimeter; see also [I-D.ietf-opsec-ipv6-eh-filtering]
- o Filter packets having an illegal IPv6 headers chain at the perimeter (and if possible, inside the network as well), see Section 2.2
- o Filter unneeded services at the perimeter
- o Implement ingress and egress anti-spoofing in the forwarding and control planes, see [RFC2827] and [RFC3704]
- o Implement appropriate rate-limiters and control-plane policers based on traffic baselines

Having global IPv6 addresses on all the enterprise sites is different than in IPv4 where [RFC1918] addresses are often used internally and not routed over the Internet. [RFC7359] and [WEBER_VPN] explain that without careful design, there could be IPv6 leakages from layer-3 VPNs.

3.2. Internal Security Considerations

The internal aspect deals with providing security inside the perimeter of the network, including end hosts. Internal networks of enterprises are often different: University campus, wireless guest access, ... so there is no "one size fits all" recommendation.

The most significant concerns here are related to Neighbor Discovery. At the network level, it is recommended that all security considerations discussed in Section 2.3 be reviewed carefully and the recommendations be considered in-depth as well. Section 4.1 of [RFC7381] also provides some recommendations.

As mentioned in Section 2.7.2, care must be taken when running automated IPv6-in-IPv4 tunnels.

When site-to-site VPNs are used it should be kept in mind that, given the global scope of IPv6 global addresses as opposed to the common use of IPv4 private address space [RFC1918], sites might be able to communicate with each other over the Internet even when the VPN mechanism is not available and hence no traffic encryption is performed and traffic could be injected from the Internet into the site, see [WEBER_VPN]. It is recommended to filter at Internet connection(s) packets having a source or destination address belonging to the site internal prefix(es); this should be done for ingress and egress traffic.

Hosts need to be hardened directly through security policy to protect against security threats. The host firewall default capabilities have to be clearly understood. In some cases, 3rd party firewalls have no IPv6 support whereas the native firewall installed by default has IPv6 support. General device hardening guidelines are provided in Section 2.8.

It should also be noted that many hosts still use IPv4 for transporting logs for RADIUS, DIAMETER, TACACS+, SYSLOG, etc. Operators cannot rely on an IPv6-only security policy to secure such protocols that are still using IPv4.

4. Service Providers Security Considerations

4.1. BGP

The threats and mitigation techniques are identical between IPv4 and IPv6. Broadly speaking they are:

- o Authenticating the TCP session;
- o TTL security (which becomes hop-limit security in IPv6) as [RFC5082];
- o bogon AS filtering, see [CYMRU];
- o Prefix filtering.

These are explained in more detail in Section 2.5. Also, the recommendations of [RFC7454] should be considered.

4.1.1. Remote Triggered Black Hole Filtering (RTBH)

RTBH [RFC5635] works identically in IPv4 and IPv6. IANA has allocated the 100::/64 prefix to be used as the discard prefix [RFC6666]

4.2. Transition/Coexistence Mechanism

SPs will typically use transition mechanisms such as 6rd, 6PE, MAP, and NAT64 which have been analyzed in the transition and coexistence Section 2.7 section.

4.3. Lawful Intercept

The Lawful Intercept requirements are similar for IPv6 and IPv4 architectures and will be subject to the laws enforced in different geographic regions. The local issues with each jurisdiction can make this challenging and both corporate legal and privacy personnel should be involved in discussions pertaining to what information gets logged and with regard to the respective log retention policies for this information.

The target of interception will usually be a residential subscriber (e.g., his/her PPP session, physical line, or CPE MAC address). In the absence of IPv6 NAT on the CPE, IPv6 has the possibility to allow for intercepting the traffic from a single host (i.e., a /128 target) rather than the whole set of hosts of a subscriber (which could be a /48, /60, or /64).

In contrast, in mobile environments, since the 3GPP specifications allocate a /64 per device, it may be sufficient to intercept traffic from the /64 rather than specific /128's (since each time the device establishes a data connection it gets a new IID).

5. Residential Users Security Considerations

The IETF Homenet working group is working on standards and guidelines for IPv6 residential networks; this obviously includes operational security considerations; but this is still work in progress. [RFC8520] is an interesting approach on how firewalls could retrieve and apply specific security policies to some residential devices.

Some residential users have less experience and knowledge about security or networking than experimented operators. As most of the recent hosts (e.g., smartphones, tablets) have IPv6 enabled by default, IPv6 security is important for those users. Even with an IPv4-only ISP, those users can get IPv6 Internet access with the help of Teredo (Section 2.7.2.8) tunnels. Several peer-to-peer programs support IPv6 and those programs can initiate a Teredo tunnel through an IPv4 residential gateway, with the consequence of making the internal host reachable from any IPv6 host on the Internet. It is therefore recommended that all host security products (including personal firewalls) are configured with a dual-stack security policy.

If the residential CPE has IPv6 connectivity, [RFC7084] defines the requirements of an IPv6 CPE and does not take a position on the debate of default IPv6 security policy as defined in [RFC6092]:

- o outbound only: allowing all internally initiated connections and block all externally initiated ones, which is a common default security policy enforced by IPv4 Residential Gateway doing NAT but it also breaks the end-to-end reachability promise of IPv6. [RFC6092] lists several recommendations to design such a CPE;
- o open/transparent: allowing all internally and externally initiated connections, therefore restoring the end-to-end nature of the Internet for IPv6 traffic but having a different security policy for IPv6 than for IPv4.

[RFC6092] REC-49 states that a choice must be given to the user to select one of those two policies.

6. Further Reading

There are several documents that describe in more detail the security of an IPv6 network; these documents are not written by the IETF and

some of them are dated but are listed here for the reader's convenience:

1. Guidelines for the Secure Deployment of IPv6 [NIST]
2. North American IPv6 Task Force Technology Report - IPv6 Security Technology Paper [NAv6TF_Security]
3. IPv6 Security [IPv6_Security_Book]

7. Acknowledgements

The authors would like to thank the following people for their useful comments: Mikael Abrahamsson, Fred Baker, Mustafa Suha Botsali, Mohamed Boucadair, Brian Carpenter, Tim Chown, Lorenzo Colitti, Roman Danyliw (IESG review), Markus de Bruen, Lars Eggert (IESG review), Tobias Fiebig, Fernando Gont, Jeffry Handal, Lee Howard, Benjamin Kaduk (IESG review), Panos Kampanakis, Erik Kline, Jouni Korhonen, Warren Kumari (IESG review), Ted Lemon, Mark Lentczner, Acee Lindem (and his detailed nits), Jen Linkova (and her detailed review), Gyan S. Mishra (the document shepherd), Jordi Palet, Alvaro Retana (IESG review), Zaheduzzaman Sarker (IESG review), Bob Sleigh, Donald Smith, Tarko Tikan, Ole Troan, Bernie Volz (by alphabetical order).

8. Security Considerations

This memo attempts to give an overview of security considerations of operating an IPv6 network both for an IPv6-only network and for networks utilizing the most widely deployed IPv4/IPv6 coexistence strategies.

9. References

9.1. Normative References

- [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>>.

9.2. Informative References

- [CYMRU] Team, C., "The Bogon Reference", Existing in 2021, <<https://team-cymru.com/community-services/bogon-reference/>>.

[ENTROPYIP]

Foremski, P., Plonka, D., and A. Berger, "Entropy/IP: Uncovering Structure in IPv6 Addresses",
<<http://www.entropy-ip.com/>>.

[europol-cgn]

Europol, "ARE YOU SHARING THE SAME IP ADDRESS AS A CRIMINAL? LAW ENFORCEMENT CALL FOR THE END OF CARRIER GRADE NAT (CGN) TO INCREASE ACCOUNTABILITY ONLINE", October 2017,
<<https://www.europol.europa.eu/newsroom/news/are-you-sharing-same-ip-address-criminal-law-enforcement-call-for-end-of-carrier-grade-nat-cgn-to-increase-accountability-online>>.

[GDPR]

Union, O. J. O. T. E., "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)", April 2016,
<<https://eur-lex.europa.eu/eli/reg/2016/679/oj>>.

[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.

[IANA-IPFIX]

IANA, "IP Flow Information Export (IPFIX) Entities",
<<http://www.iana.org/assignments/ipfix>>.

[IEEE-802.1X]

IEEE, "IEEE Standard for Local and metropolitan area networks - Port-Based Network Access Control", IEEE Std 802.1X-2010, February 2010.

[IPv6_Security_Book]

Hogg, S. and E. Vyncke, "IPv6 Security", ISBN 1-58705-594-5, Publisher CiscoPress, December 2008.

[KRISTOFF]

Kristoff, J., Ghasemisharif, M., Kanich, C., and J. Polakis, "Plight at the End of the Tunnel: Legacy IPv6 Transition Mechanisms in the Wild", March 2021, <<https://dataplane.org/jtk/publications/kgkp-pam-21.pdf>>.

[NAv6TF_Security]

Kaeo, M., Green, D., Bound, J., and Y. Pouffary, "North American IPv6 Task Force Technology Report - IPv6 Security Technology Paper", 2006, <http://www.ipv6forum.com/dl/white/NAv6TF_Security_Report.pdf>.

[NIST]

Frankel, S., Graveman, R., Pearce, J., and M. Rocks, "Guidelines for the Secure Deployment of IPv6", 2010, <<http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf>>.

[RADB]

INC., M. N., "RADb The Internet Routing Registry", Existing in 2021, <<https://www.radb.net/>>.

[REY_PF]

Rey, E., "Local Packet Filtering with IPv6", July 2017, <https://labs.ripe.net/Members/enno_rey/local-packet-filtering-with-ipv6>.

[RFC0826]

Plummer, D., "An Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, DOI 10.17487/RFC0826, November 1982, <<https://www.rfc-editor.org/info/rfc826>>.

[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>>.

[RFC2131]

Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.

[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>>.

[RFC2529]

Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, DOI 10.17487/RFC2529, March 1999, <<https://www.rfc-editor.org/info/rfc2529>>.

- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, DOI 10.17487/RFC2663, August 1999, <<https://www.rfc-editor.org/info/rfc2663>>.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.
- [RFC2866] Rigney, C., "RADIUS Accounting", RFC 2866, DOI 10.17487/RFC2866, June 2000, <<https://www.rfc-editor.org/info/rfc2866>>.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, DOI 10.17487/RFC3056, February 2001, <<https://www.rfc-editor.org/info/rfc3056>>.
- [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, DOI 10.17487/RFC3068, June 2001, <<https://www.rfc-editor.org/info/rfc3068>>.
- [RFC3627] Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful", RFC 3627, DOI 10.17487/RFC3627, September 2003, <<https://www.rfc-editor.org/info/rfc3627>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.
- [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>>.
- [RFC3924] Baker, F., Foster, B., and C. Sharp, "Cisco Architecture for Lawful Intercept in IP Networks", RFC 3924, DOI 10.17487/RFC3924, October 2004, <<https://www.rfc-editor.org/info/rfc3924>>.
- [RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, DOI 10.17487/RFC3964, December 2004, <<https://www.rfc-editor.org/info/rfc3964>>.

- [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>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC4107] Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management", BCP 107, RFC 4107, DOI 10.17487/RFC4107, June 2005, <<https://www.rfc-editor.org/info/rfc4107>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.
- [RFC4293] Routhier, S., Ed., "Management Information Base for the Internet Protocol (IP)", RFC 4293, DOI 10.17487/RFC4293, April 2006, <<https://www.rfc-editor.org/info/rfc4293>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, DOI 10.17487/RFC4380, February 2006, <<https://www.rfc-editor.org/info/rfc4380>>.

- [RFC4381] Behringer, M., "Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4381, DOI 10.17487/RFC4381, February 2006, <<https://www.rfc-editor.org/info/rfc4381>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <<https://www.rfc-editor.org/info/rfc4552>>.
- [RFC4649] Volz, B., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option", RFC 4649, DOI 10.17487/RFC4649, August 2006, <<https://www.rfc-editor.org/info/rfc4649>>.
- [RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, DOI 10.17487/RFC4659, September 2006, <<https://www.rfc-editor.org/info/rfc4659>>.
- [RFC4795] Aboba, B., Thaler, D., and L. Esibov, "Link-local Multicast Name Resolution (LLMNR)", RFC 4795, DOI 10.17487/RFC4795, January 2007, <<https://www.rfc-editor.org/info/rfc4795>>.
- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, DOI 10.17487/RFC4798, February 2007, <<https://www.rfc-editor.org/info/rfc4798>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, DOI 10.17487/RFC4864, May 2007, <<https://www.rfc-editor.org/info/rfc4864>>.

- [RFC4890] Davies, E. and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", RFC 4890, DOI 10.17487/RFC4890, May 2007, <<https://www.rfc-editor.org/info/rfc4890>>.
- [RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations", RFC 4942, DOI 10.17487/RFC4942, September 2007, <<https://www.rfc-editor.org/info/rfc4942>>.
- [RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <<https://www.rfc-editor.org/info/rfc5082>>.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, DOI 10.17487/RFC5214, March 2008, <<https://www.rfc-editor.org/info/rfc5214>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [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>>.
- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, DOI 10.17487/RFC5952, August 2010, <<https://www.rfc-editor.org/info/rfc5952>>.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, DOI 10.17487/RFC5969, August 2010, <<https://www.rfc-editor.org/info/rfc5969>>.
- [RFC6092] Woodyatt, J., Ed., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, DOI 10.17487/RFC6092, January 2011, <<https://www.rfc-editor.org/info/rfc6092>>.
- [RFC6104] Chown, T. and S. Venaas, "Rogue IPv6 Router Advertisement Problem Statement", RFC 6104, DOI 10.17487/RFC6104, February 2011, <<https://www.rfc-editor.org/info/rfc6104>>.

- [RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, DOI 10.17487/RFC6105, February 2011, <<https://www.rfc-editor.org/info/rfc6105>>.
- [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, DOI 10.17487/RFC6144, April 2011, <<https://www.rfc-editor.org/info/rfc6144>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <<https://www.rfc-editor.org/info/rfc6147>>.
- [RFC6164] Kohno, M., Nitzan, B., Bush, R., Matsuzaki, Y., Colitti, L., and T. Narten, "Using 127-Bit IPv6 Prefixes on Inter-Router Links", RFC 6164, DOI 10.17487/RFC6164, April 2011, <<https://www.rfc-editor.org/info/rfc6164>>.
- [RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, DOI 10.17487/RFC6169, April 2011, <<https://www.rfc-editor.org/info/rfc6169>>.
- [RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, DOI 10.17487/RFC6177, March 2011, <<https://www.rfc-editor.org/info/rfc6177>>.
- [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>>.
- [RFC6221] Miles, D., Ed., Ooghe, S., Dec, W., Krishnan, S., and A. Kavanagh, "Lightweight DHCPv6 Relay Agent", RFC 6221, DOI 10.17487/RFC6221, May 2011, <<https://www.rfc-editor.org/info/rfc6221>>.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.

- [RFC6264] Jiang, S., Guo, D., and B. Carpenter, "An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition", RFC 6264, DOI 10.17487/RFC6264, June 2011, <<https://www.rfc-editor.org/info/rfc6264>>.
- [RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <<https://www.rfc-editor.org/info/rfc6269>>.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, DOI 10.17487/RFC6296, June 2011, <<https://www.rfc-editor.org/info/rfc6296>>.
- [RFC6302] Durand, A., Gashinsky, I., Lee, D., and S. Sheppard, "Logging Recommendations for Internet-Facing Servers", BCP 162, RFC 6302, DOI 10.17487/RFC6302, June 2011, <<https://www.rfc-editor.org/info/rfc6302>>.
- [RFC6324] Nakibly, G. and F. Templin, "Routing Loop Attack Using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations", RFC 6324, DOI 10.17487/RFC6324, August 2011, <<https://www.rfc-editor.org/info/rfc6324>>.
- [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>>.
- [RFC6343] Carpenter, B., "Advisory Guidelines for 6to4 Deployment", RFC 6343, DOI 10.17487/RFC6343, August 2011, <<https://www.rfc-editor.org/info/rfc6343>>.
- [RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <<https://www.rfc-editor.org/info/rfc6434>>.
- [RFC6459] Korhonen, J., Ed., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", RFC 6459, DOI 10.17487/RFC6459, January 2012, <<https://www.rfc-editor.org/info/rfc6459>>.
- [RFC6547] George, W., "RFC 3627 to Historic Status", RFC 6547, DOI 10.17487/RFC6547, February 2012, <<https://www.rfc-editor.org/info/rfc6547>>.

- [RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and M. Bhatia, "A Uniform Format for IPv6 Extension Headers", RFC 6564, DOI 10.17487/RFC6564, April 2012, <<https://www.rfc-editor.org/info/rfc6564>>.
- [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>>.
- [RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", BCP 153, RFC 6598, DOI 10.17487/RFC6598, April 2012, <<https://www.rfc-editor.org/info/rfc6598>>.
- [RFC6620] Nordmark, E., Bagnulo, M., and E. Levy-Abegnoli, "FCFS SAVI: First-Come, First-Served Source Address Validation Improvement for Locally Assigned IPv6 Addresses", RFC 6620, DOI 10.17487/RFC6620, May 2012, <<https://www.rfc-editor.org/info/rfc6620>>.
- [RFC6666] Hilliard, N. and D. Freedman, "A Discard Prefix for IPv6", RFC 6666, DOI 10.17487/RFC6666, August 2012, <<https://www.rfc-editor.org/info/rfc6666>>.
- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <<https://www.rfc-editor.org/info/rfc6763>>.
- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [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>>.
- [RFC6888] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, DOI 10.17487/RFC6888, April 2013, <<https://www.rfc-editor.org/info/rfc6888>>.

- [RFC6939] Halwasia, G., Bhandari, S., and W. Dec, "Client Link-Layer Address Option in DHCPv6", RFC 6939, DOI 10.17487/RFC6939, May 2013, <<https://www.rfc-editor.org/info/rfc6939>>.
- [RFC6964] Templin, F., "Operational Guidance for IPv6 Deployment in IPv4 Sites Using the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 6964, DOI 10.17487/RFC6964, May 2013, <<https://www.rfc-editor.org/info/rfc6964>>.
- [RFC6967] Boucadair, M., Touch, J., Levis, P., and R. Penno, "Analysis of Potential Solutions for Revealing a Host Identifier (HOST_ID) in Shared Address Deployments", RFC 6967, DOI 10.17487/RFC6967, June 2013, <<https://www.rfc-editor.org/info/rfc6967>>.
- [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>>.
- [RFC7010] Liu, B., Jiang, S., Carpenter, B., Venaas, S., and W. George, "IPv6 Site Renumbering Gap Analysis", RFC 7010, DOI 10.17487/RFC7010, September 2013, <<https://www.rfc-editor.org/info/rfc7010>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.
- [RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <<https://www.rfc-editor.org/info/rfc7012>>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <<https://www.rfc-editor.org/info/rfc7039>>.
- [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>>.

- [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <<https://www.rfc-editor.org/info/rfc7050>>.
- [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.
- [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>>.
- [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>>.
- [RFC7123] Gont, F. and W. Liu, "Security Implications of IPv6 on IPv4 Networks", RFC 7123, DOI 10.17487/RFC7123, February 2014, <<https://www.rfc-editor.org/info/rfc7123>>.
- [RFC7166] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 7166, DOI 10.17487/RFC7166, March 2014, <<https://www.rfc-editor.org/info/rfc7166>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.
- [RFC7359] Gont, F., "Layer 3 Virtual Private Network (VPN) Tunnel Traffic Leakages in Dual-Stack Hosts/Networks", RFC 7359, DOI 10.17487/RFC7359, August 2014, <<https://www.rfc-editor.org/info/rfc7359>>.
- [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>>.

- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", RFC 7404, DOI 10.17487/RFC7404, November 2014, <<https://www.rfc-editor.org/info/rfc7404>>.
- [RFC7422] Donley, C., Grundemann, C., Sarawat, V., Sundaresan, K., and O. Vautrin, "Deterministic Address Mapping to Reduce Logging in Carrier-Grade NAT Deployments", RFC 7422, DOI 10.17487/RFC7422, December 2014, <<https://www.rfc-editor.org/info/rfc7422>>.
- [RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <<https://www.rfc-editor.org/info/rfc7454>>.
- [RFC7513] Bi, J., Wu, J., Yao, G., and F. Baker, "Source Address Validation Improvement (SAVI) Solution for DHCP", RFC 7513, DOI 10.17487/RFC7513, May 2015, <<https://www.rfc-editor.org/info/rfc7513>>.
- [RFC7526] Troan, O. and B. Carpenter, Ed., "Deprecating the Anycast Prefix for 6to4 Relay Routers", BCP 196, RFC 7526, DOI 10.17487/RFC7526, May 2015, <<https://www.rfc-editor.org/info/rfc7526>>.
- [RFC7552] Asati, R., Pignataro, C., Raza, K., Manral, V., and R. Papneja, "Updates to LDP for IPv6", RFC 7552, DOI 10.17487/RFC7552, June 2015, <<https://www.rfc-editor.org/info/rfc7552>>.
- [RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <<https://www.rfc-editor.org/info/rfc7597>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <<https://www.rfc-editor.org/info/rfc7599>>.
- [RFC7610] Gont, F., Liu, W., and G. Van de Velde, "DHCPv6-Shield: Protecting against Rogue DHCPv6 Servers", BCP 199, RFC 7610, DOI 10.17487/RFC7610, August 2015, <<https://www.rfc-editor.org/info/rfc7610>>.

- [RFC7707] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", RFC 7707, DOI 10.17487/RFC7707, March 2016, <<https://www.rfc-editor.org/info/rfc7707>>.
- [RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <<https://www.rfc-editor.org/info/rfc7721>>.
- [RFC7772] Yourtchenko, A. and L. Colitti, "Reducing Energy Consumption of Router Advertisements", BCP 202, RFC 7772, DOI 10.17487/RFC7772, February 2016, <<https://www.rfc-editor.org/info/rfc7772>>.
- [RFC7785] Vinapamula, S. and M. Boucadair, "Recommendations for Prefix Binding in the Context of Software Dual-Stack Lite", RFC 7785, DOI 10.17487/RFC7785, February 2016, <<https://www.rfc-editor.org/info/rfc7785>>.
- [RFC7824] Krishnan, S., Mrugalski, T., and S. Jiang, "Privacy Considerations for DHCPv6", RFC 7824, DOI 10.17487/RFC7824, May 2016, <<https://www.rfc-editor.org/info/rfc7824>>.
- [RFC7844] Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity Profiles for DHCP Clients", RFC 7844, DOI 10.17487/RFC7844, May 2016, <<https://www.rfc-editor.org/info/rfc7844>>.
- [RFC7857] Penno, R., Perreault, S., Boucadair, M., Ed., Sivakumar, S., and K. Naito, "Updates to Network Address Translation (NAT) Behavioral Requirements", BCP 127, RFC 7857, DOI 10.17487/RFC7857, April 2016, <<https://www.rfc-editor.org/info/rfc7857>>.
- [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>>.
- [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont, "IP/ICMP Translation Algorithm", RFC 7915, DOI 10.17487/RFC7915, June 2016, <<https://www.rfc-editor.org/info/rfc7915>>.

- [RFC7934] Colitti, L., Cerf, V., Cheshire, S., and D. Schinazi, "Host Address Availability Recommendations", BCP 204, RFC 7934, DOI 10.17487/RFC7934, July 2016, <<https://www.rfc-editor.org/info/rfc7934>>.
- [RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017, <<https://www.rfc-editor.org/info/rfc8040>>.
- [RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", RFC 8064, DOI 10.17487/RFC8064, February 2017, <<https://www.rfc-editor.org/info/rfc8064>>.
- [RFC8177] Lindem, A., Ed., Qu, Y., Yeung, D., Chen, I., and J. Zhang, "YANG Data Model for Key Chains", RFC 8177, DOI 10.17487/RFC8177, June 2017, <<https://www.rfc-editor.org/info/rfc8177>>.
- [RFC8190] Bonica, R., Cotton, M., Haberman, B., and L. Vegoda, "Updates to the Special-Purpose IP Address Registries", BCP 153, RFC 8190, DOI 10.17487/RFC8190, June 2017, <<https://www.rfc-editor.org/info/rfc8190>>.
- [RFC8210] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1", RFC 8210, DOI 10.17487/RFC8210, September 2017, <<https://www.rfc-editor.org/info/rfc8210>>.
- [RFC8273] Brzozowski, J. and G. Van de Velde, "Unique IPv6 Prefix per Host", RFC 8273, DOI 10.17487/RFC8273, December 2017, <<https://www.rfc-editor.org/info/rfc8273>>.
- [RFC8343] Bjorklund, M., "A YANG Data Model for Interface Management", RFC 8343, DOI 10.17487/RFC8343, March 2018, <<https://www.rfc-editor.org/info/rfc8343>>.
- [RFC8344] Bjorklund, M., "A YANG Data Model for IP Management", RFC 8344, DOI 10.17487/RFC8344, March 2018, <<https://www.rfc-editor.org/info/rfc8344>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

- [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>>.
- [RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.
- [RFC8541] Litkowski, S., Decraene, B., and M. Horneffer, "Impact of Shortest Path First (SPF) Trigger and Delay Strategies on IGP Micro-loops", RFC 8541, DOI 10.17487/RFC8541, March 2019, <<https://www.rfc-editor.org/info/rfc8541>>.
- [RFC8981] Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", RFC 8981, DOI 10.17487/RFC8981, February 2021, <<https://www.rfc-editor.org/info/rfc8981>>.
- [SCANNING] Barnes, R., Altmann, R., and D. Kerr, "Mapping the Great Void - Smarter scanning for IPv6", February 2012, <http://www.caida.org/workshops/isma/1202/slides/aims1202_rbarnes.pdf>.
- [WEBER_VPN] Weber, J., "Dynamic IPv6 Prefix - Problems and VPNs", March 2018, <<https://blog.webernetz.net/wp-content/uploads/2018/03/TR18-Johannes-Weber-Dynamic-IPv6-Prefix-Problems-and-VPNs.pdf>>.

Authors' Addresses

Eric Vyncke
Cisco
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2 778 4677
Email: evyncke@cisco.com

Kiran Kumar
Square
1455 Market Street, Suite 600
San Francisco 94103
United States of America

Email: kk.chittimaneni@gmail.com

Merike Kaeo
Double Shot Security
3518 Fremont Ave N 363
Seattle 98103
United States of America

Phone: +12066696394
Email: merike@doubleshotsecurity.com

Enno Rey
ERNW
Carl-Bosch-Str. 4
Heidelberg, Baden-Wuerttemberg 69115
Germany

Phone: +49 6221 480390
Email: erey@ernw.de

v6ops
Internet-Draft
Intended status: Informational
Expires: January 6, 2020

J. Linkova
Google
July 5, 2019

Neighbor Cache Entries on First-Hop Routers: Operational Considerations
draft-linkova-v6ops-nd-cache-init-01

Abstract

Neighbor Discovery (RFC4861) is used by IPv6 nodes to determine the link-layer addresses of neighboring nodes as well as to discover and maintain reachability information. This document discusses how the neighbor discovery state machine on a first-hop router is causing user-visible connectivity issues when a new (not being seen on the network before) IPv6 address is being used.

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 January 6, 2020.

Copyright Notice

Copyright (c) 2019 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
1.1. Requirements Language	4
1.2. Terminology	4
2. Potential Solutions	5
2.1. Do Nothing	5
2.1.1. Pros	5
2.1.2. Cons	5
2.2. Change to the Registration-Based Neighbor Discovery	5
2.3. Hosts Explicitly Advertizing Their GUAs Using Existing ND Mechanisms	6
2.3.1. Host Sending Unsolicited NA	6
2.3.1.1. Pros	7
2.3.1.2. Cons	7
2.3.2. Host Sending NS to the Router Address from Its GUA	7
2.3.2.1. Pros	7
2.3.2.2. Cons	7
2.3.3. Host Sending Router Solicitation from its GUA	8
2.3.3.1. Pros	8
2.3.3.2. Cons	8
2.4. Initiating Hosts2Routers Communication	8
2.4.1. Pros	9
2.4.2. Cons	9
2.5. Tweaking Probing Algorithms	9
2.6. Routers Buffering More Packets	9
2.6.1. Pros	10
2.6.2. Cons	10
3. Recommendations	10
3.1. Avoiding Disruption	10
4. IANA Considerations	11
5. Security Considerations	11
6. Acknowledgements	12
7. References	12
7.1. Normative References	12
7.2. Informative References	13
Author's Address	13

1. Introduction

The section 7.2.5 of [RFC4861] states: "When a valid Neighbor Advertisement is received (either solicited or unsolicited), the Neighbor Cache is searched for the target's entry. If no entry exists, the advertisement SHOULD be silently discarded. There is no

need to create an entry if none exists, since the recipient has apparently not initiated any communication with the target."

This approach is perfectly suitable for host2host communications which are in most cases bi-directional and it could be expected that if a host A has an ND cache entry for the host B IPv6 address, the host B also has the corresponding ND entry for the host A address in its cache. However when a host communicates to off-link destinations via its first-hop router that logic does not apply. Here is the most typical scenario when the problem may arise:

1. When a host joins the network it receives an RA packet from the first-hop router (either a periodic unsolicited RA or a response to an RS sent by the host). The RA contains information the host needs to perform SLAAC and to configure its network stack. Among other things the host populates its ND cache with the router link-local address and potentially link-layer address (if included in the RA Source Link-Layer Address option).
2. The host starts opening connections to off-link destinations. Very common use case is a mobile device sending probes to detect the Internet connectivity and/or the captive portals presence on the network. To speed up that process many implementations are using the Optimistic Duplicate Address Detection ([RFC4429]) which allows them to send probes from their GUA before the DAD process is completed. Important point here is that at that moment the device ND cache contains all information required to send those probes (such as the default gateway LLA and the link-layer address). The router ND cache, however, might contain an entry for the device link-local address (if the device has been performing the ND process for the router LLA) but there are no entries for the device GUA.
3. Response packets for the probes (or any other traffic sent by the host) are received by the first-hop router. As the router does not have any ND cache entry for the host GUA, the router starts the neighbor discovery process by creating an INCOMPLETE cache entry and then sending an NS to the Solicited Node Multicast Address. Apparently most of the router implementations buffer only one data packet while performing the ND process for its destination. Therefore all packets for the host GUA, except for the very first one are dropped until the address resolution process is completed.
4. As many implementations send multiple probes in parallel it's very likely that all probes ex. the first one would be considered failed. If the host implements an exponential backoff for

probing it leads to user-noticeable delay in detecting network connectivity/reporting the network as usable.

The above-mentioned scenario illustrates the problem happening when the device connects to the network for the first time/after a long timeout. However the same sequence of events happen when the host starts using the new (previously unseen by the router or) GUA (e.g. a new privacy address [RFC4941]) or if the router Neighbor Cache has been flushed.

While in dual-stack networks this problem might be hidden by Happy Eyeballs ([RFC8305]) it manifests itself quite clearly in IPv6-only networks, especially wireless ones, leading to poor user experience and contributing to negative perception of IPv6-only solutions as unstable and non-deployable.

1.1. Requirements Language

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.

1.2. Terminology

ND: Neighbor Discovery, [RFC4861].

SLAAC: IPv6 Stateless Address Autoconfiguration, [RFC4862].

NS: Neighbor Solicitation, [RFC4861].

NA: Neighbor Advertisement, [RFC4861].

RS: Router Solicitation, [RFC4861].

RA: Router Advertisement, [RFC4861].

SLLA: Source link-layer Address, an option in the ND packets containing the link-layer address of the sender of the packet ([RFC4861]).

TLLA: Target link-layer Address, an option in the ND packets containing the link-layer address of the target ([RFC4861]).

GUA: Global Unicast Address ([RFC4291]).

DAD: Duplicate Address Detection, [RFC4862].

Optimistic DAD: a modification of DAD, [RFC4429].

2. Potential Solutions

The problem could be addressed from different angles. Possible approaches are:

- o Just do nothing.
- o Migrate from the "reactive" Neighbor Discovery ([RFC4861]) to the registration-based mechanisms ([RFC8505]).
- o The host explicitly advertizes its GUAs using Neighbor Discovery mechanisms.
- o The host initiates bidirectional communication to the router using the host GUA.
- o Making the probing logic on hosts more robust.
- o Increasing the buffer size on routers.

The following sections discuss those approaches in more detail.

2.1. Do Nothing

One of the possible approaches might be to declare that everything is working as intended.

2.1.1. Pros

- o No work required.

2.1.2. Cons

- o Unhappy users.
- o Many support tickets.
- o More resistance to deploy IPv6 and IPv6-Only networks.

2.2. Change to the Registration-Based Neighbor Discovery

The most radical approach would be to move away from the reactive ND as defined in [RFC4861] and expand the registration-based ND ([RFC6775], [RFC8505]) used in Low-Power Wireless Personal Area Networks (6LoWPANs) to the rest of IPv6 deployments.

This option required some investigation and discussions. More text will be added here in the following revision of the draft.

2.3. Hosts Explicitly Advertizing Their GUAs Using Existing ND Mechanisms

The Neighbor Discovery is designed to allow IPv6 nodes to discover neighboring nodes reachability and learn IPv6 to link-layer addresses mapping. Therefore ND seems to be the most appropriate tool to inform the first-hop routers about addresses the host is going to use. The following sections discuss potential approaches in more detail.

2.3.1. Host Sending Unsolicited NA

Section 4.4 of [RFC4861] says:

"A node sends Neighbor Advertisements in response to Neighbor Solicitations and sends unsolicited Neighbor Advertisements in order to (unreliably) propagate new information quickly."

Propagating information about new GUA as quickly as possible is exactly what is required to solve the problem outlined in this document. Therefore the host might send an unsolicited NA to advertize its GUA as soon as the said address enters Optimistic or Preferred state. The NA should include the target link-layer address option. To ensure that all first-hop routers receive the advertisement it could be sent to all-routers multicast address (ff02::2).

As it's been mentioned, [RFC4861] explicitly states that receiving a NA should not create a new NC entry. However the justification for that requirement ("There is no need to create an entry if none exists, since the recipient has apparently not initiated any communication with the target.") clearly does not apply for the case discussed. As per [RFC2119] "there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course.". Therefore routers creating a new NC entry upon receiving an unsolicited NA would still be compliant with [RFC4861].

It should be noted that some routing and switching platforms have implemented such behaviour already. Administrators could enable creating neighbor discovery cache entries based on unsolicited NA packets sent from the previously unknown neighbors on that interface.

2.3.1.1. Pros

- o Already implemented on some platforms.
- o In accordance with [RFC4861].

2.3.1.2. Cons

- o Allows a malicious host to execute an ND cache exhaustion attack. It's recommended that this functionality is configurable and recommendations from [RFC6583] are taken into account.
- o Requires hosts to send unsolicited NA (changes to the hosts).
- o Some wireless devices are known to fiddle with ND packets and perform various non-obvious forms of ND proxy actions. In some cases unsolicited NAs might not even reach the routers.

2.3.2. Host Sending NS to the Router Address from Its GUA

The host could force creating a STALE entry for its GUA in the router ND cache by sending the following Neighbor Solicitation message:

- o The NS source address is the host GUA.
- o The Source Link-Layer Address option contains the host link-layer address.
- o The target address is the host default gateway address (the default router address the host received in the RA).

The main disadvantage of this approach is that it would not work if the GUA the host needs to advertise is still in the Optimistic state. The section 2.2 of [RFC4429] explicitly prohibits sending Neighbor Solicitations from an Optimistic Address.

2.3.2.1. Pros

- o Router implementations which follow recommendations made in [RFC6583] might prioritize responding to NS packets to own addresses.

2.3.2.2. Cons

- o Does not work for Optimistic addresses (see section 2.2 of [RFC4429]).

- o If first-hop redundancy is deployed in the network, the NS would reach the active router only, so all backup routers (or all active routers ex. one) would not get their neighbor cache updated.
- o Some wireless devices are known to fiddle with ND packets and perform various non-obvious forms of ND proxy actions. In some cases unsolicited NAs might not even reach the routers.

2.3.3. Host Sending Router Solicitation from its GUA

The host could send a router solicitation message to 'all routers' multicast address, using its GUA as a source. If the host link-layer address is included in the Source Link-Layer Address option, the router would create a STALE entry for the host GUA (see the section 6.2.6 of [RFC4861]). However this approach can not be used if the GUA is in optimistic state: the section 2.2 of [RFC4429] explicitly prohibits using an Optimistic Address as the source address of a Router Solicitation with a SLLAO as it might disrupt the rightful owner of the address in the case of a collision. So for the optimistic addresses the host can send an RS without SLLAO included. In that case the router may respond with either a multicast or a unicast RA (only the latter would create a cache entry).

2.3.3.1. Pros

- o Unlike NS packets, RS packets would reach all routers on link, allowing all routers to update their neighbor caches and preventing packet loss in case of asymmetric routing.

2.3.3.2. Cons

- o As for the Optimistic addresses SLLAO can not be included into RS packets, the cache entry for the optimistic address would be created only if the router sends solicited RAs as unicast. In addition, there might be a random delay between receiving an RS and sending a unicast RA back (and creating a cache entry) which might undermine the idea of creating the cache entry proactively.
- o Some wireless devices are known to fiddle with ND packets and perform various non-obvious forms of ND proxy actions. In some cases RSes might not even reach the routers.

2.4. Initiating Hosts2Routers Communication

Every time the host configures a new GUA (when the address enters the Optimistic state or, if the optimistic DAD is not used, as soon as it changes the state from tentative to preferred) the host can a ping or traceroute packet to the default gateway LLA. As the RTT to the

default gateway is lower than RTT to any off-link destinations it's quite likely that the router would start the neighbor discovery process for the host GUA before the first packet of the returning traffic arrives. There are pretty good chances that the process would be completed before the actual data traffic reaches the router.

2.4.1. Pros

- o As data packets are involved, there is no potential impact caused by smart wireless infrastructure performing ND proxy.
- o Full compliance with existing standards.

2.4.2. Cons

- o Data packets to the router LLA could be blocked by security policy or control plane protection mechanism.
- o Maximum overhead for routers control plane (in addition to processing ND packets, the data packet needs to be processed as well).
- o If the first hop redundancy is implemented in the network the host ping/traceroute packet would reach the active router only. All backup routers would not receive it and therefore would not start populating the cache. So in the case of asymmetric traffic flow (packets leave the network via one router while the return flow is going via another) the backup router(s) still would not have the cache entry. (A hacky way to overcome this limitation would be sending ping/traceroute packet to 'all routers' ff02::2 multicast address).

2.5. Tweaking Probing Algorithms

While tweaking the probing logic on devices might make the problem less visible it would be still desirable to avoid packet loss everytime the new GUA is used by a host. It would be quite tricky to adjust every probing algorithm to find the right balance between prompt detection of network connectivity and false positives in IPv6-only mode.

2.6. Routers Buffering More Packets

Another way to mitigate the issue, at least partially, would be increasing the number of packets the router could buffer while performing the neighbor discovery process for the INCOMPLETE cache entry. However it would be against recommendations made in the section 7.2.2 of [RFC4861] and [RFC6583].

2.6.1. Pros

- o Does not require changes on hosts.

2.6.2. Cons

- o This approach makes the routers even more vulnerable to attack vectors described in [RFC6583]. In particular, it would amplify the impact of any scanning attack.
- o Against the recommendations from the section 7 of [RFC6583].
- o Requires router vendors support.

3. Recommendations

- o Hosts SHOULD send at least one unsolicited NA packet to all-routers multicast address (ff02::2) as soon as one of the following events happens:
 - * (if Optimistic DAD is used): a new Optimistic GUA is assigned to the host interface.
 - * (if Optimistic DAD is not used): a GUA changes the state from tentative to preferred.
- o Routers SHOULD have a configuration knob to enable creating ND cache entry upon receiving unsolicited NAs on a specific interface. This document does not change the behavior if the ND cache entry already exists when receiving an unsolicited NA.

3.1. Avoiding Disruption

If hosts following the recommendations in this document are using the DAD mechanism defined in [RFC4862], they would send unsolicited NA as soon as the address changes the state from tentative to preferred (after its uniqueness has been verified). However hosts willing to minimize network stack configuration delays might be using optimistic addresses, which means there is possibility of the address not being unique on the link. The section 2.2 of [RFC4429] discusses measures to ensure that ND packets from the optimistic address do not override any existing neighbor cache entries as it would cause traffic interruption of the rightful address owner in case of address conflict.

As hosts willing to speed up their network stack configuration are most likely to be affected by the problem outlined in this document it seems reasonable for such hosts to advertise their optimistic GUAs

by sending unsolicited NAs. The main question to consider is the potential risk of overriding the cache entry for the rightful address owner if the optimistic address happens to be duplicated.

As per section 7.2.5 of [RFC4861] if the Neighbor Cache entry for the target address already exists and is in any state other than INCOMPLETE then the only change the unsolicited NA could cause is to change the entry from REACHABLE to STALE. It would not cause any traffic interruption for the rightful address owner.

If there is no entry then it would be created/updated with the supplied LLA and its state set to STALE. In that case as soon as the entry is used for sending traffic to the host, the entry state will be changed to DELAY and the Neighbor Unreachability Detection would be started and the rightful owner LLA will be entered in the cache. So in the scenario when the rightful owner does not use the address for communication then it might be a short (a few seconds) period of time when the data packets sent from the outside could reach the host with the optimistic address. However it seems likely that hosts using Optimistic DAD would start sending/receiving traffic right away, so the first return packet would trigger the NUD process and rewrite the cache.

Another corner case is the INCOMPLETE cache entry for the address. If the host sends an unsolicited NA from the Optimistic address it would update the entry with the host LLA and set the entry to the STALE state. As the INCOMPLETE entry means that the router has started the ND process for the address and the multicast NS has been sent, the rightful owner is expected to reply with solicited NA which would recover the cache entry and set the LLA to the rightful owner's one. The risk here:

- o The data packet arrives after the unsolicited NA from the host but before the rightful owner responded with the solicited NA. Those packets would be sent to the host with the optimistic address instead of its rightful owner. However without the unsolicited NA those packets would have been dropped anyway (as the entry was in INCOMPLETE state).

4. IANA Considerations

This memo asks the IANA for no new parameters.

5. Security Considerations

One of the potential attack vector to consider is a cache spoofing when the attacker might try to install a cache entry for the victim's IPv6 address and the attacker's Link-Layer address. However it

should be noted that this document does not propose any changes for the scenario when the ND cache for the given IPv6 address already exists. Therefore it is not possible for the attacker to override any existing cache entry.

A malicious host could attempt to exhaust the neighbor cache on the router by creating a large number of STALE entries. However this attack vector is not new and this document does not increase the risk of such attack: the attacker could do it, for example, by sending a NS or RS packet with SLLAO included. All recommendations from [RFC6583] still apply.

6. Acknowledgements

Thanks to the following people (in alphabetical order) for their review and feedback: Lorenzo Colitti, Tatuya Jinmei, Erik Kline, Warren Kumari, Michael Richardson, Pascal Thubert, Loganaden Velvindron, Eric Vyncke.

7. References

7.1. 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>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, DOI 10.17487/RFC4429, April 2006, <<https://www.rfc-editor.org/info/rfc4429>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [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>>.

- [RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <<https://www.rfc-editor.org/info/rfc6775>>.
- [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>>.
- [RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.
- [RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <<https://www.rfc-editor.org/info/rfc8505>>.

7.2. Informative References

- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.
- [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>>.

Author's Address

Jen Linkova
Google
1 Darling Island Rd
Pyrmont, NSW 2009
AU

Email: furry@google.com

v6ops
Internet-Draft
Intended status: Informational
Expires: May 7, 2020

J. Palet Martinez
The IPv6 Company
A. D'Egidio
Telecentro
November 4, 2019

464XLAT Optimization
draft-palet-v6ops-464xlat-opt-cdn-caches-04

Abstract

This document proposes possible solutions to avoid certain drawbacks of IP/ICMP Translation Algorithm (SIIT) when the destinations are available with IPv6. When SIIT is used as a NAT46 and IPv4-only devices or applications initiate traffic flows to dual-stack CDNs (Content Delivery Networks), Caches or other network resources (in the operator network or Internet), those flows will be translated back to IPv4 by a NAT64. This is the case for 464XLAT and MAP-T.

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 May 7, 2020.

Copyright Notice

Copyright (c) 2019 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. Requirements Language	4
3. Problem Statement	4
4. Solution Approaches	6
4.1. Approach 1: DNS/Routing-based Solution	6
4.2. Approach 2: NAT46/CLAT/DNS-proxy-EAM-based Solution	7
4.2.1. Detection of IPv4-only devices or applications	7
4.2.2. Detection of IPv6-enabled service	8
4.2.3. Creation of EAMT entries	8
4.2.4. Forwarding path via stateful NAT for existing EAMT entries	10
4.2.5. Maintenance of the EAMT entries	10
4.2.6. Usage example	10
4.2.7. Behavior in case of multiple A/AAAA RRs	11
4.2.8. Behavior in presence/absence of DNS64	11
4.2.9. Behavior when using literal addresses or non IPv6-compliant APIs	11
4.2.10. False detection of a dual-stack host as IPv4-only	11
4.2.11. Behaviour in presence of Happy Eyeballs	12
4.2.12. Behavior in case of Foreign DNS	12
4.3. Approach 3: NAT46/CLAT-provider-EAM-based Solution	13
5. IPv6-only Services become accessible to IPv4-only devices/apps	14
6. Conclusions	14
7. Security Considerations	15
8. IANA Considerations	15
9. Acknowledgements	15
10. References	15
10.1. Normative References	15
10.2. Informative References	16
Authors' Addresses	17

1. Introduction

Different transition mechanisms, typically in the group of the so-called IPv6-only with IPv4aaS (IPv4-as-a-Service), such as 464XLAT ([RFC6877]) or MAP-T ([RFC7599]), allow IPv4-only devices or applications to connect with IPv4 services in Internet, by means of a NAT46 SIIT (IP/ICMP Translation Algorithm) as described by [RFC7915].

This is done by the implementation of SIIT at the CE (Customer Edge) Router or sometimes at the end-device, for example, the UE (User

Equipment) in cellular networks. This functionality is the CLAT (Customer Translator) in the case of 464XLAT.

The NAT46/CLAT (WAN side) is connected by IPv6-only to the operator network, which in turn, will have a reverse function, the NAT64 ([RFC6146]), known as PLAT (Provider Translator) in the case of 464XLAT. This allows to translate the IPv6-only flow back to IPv4, in order to forward it to Internet.

The translation of the packet headers is done using the IP/ICMP translation algorithm defined in [RFC7915] and algorithmically translating the IPv4 addresses to IPv6 addresses following [RFC6052].

In the case of 464XLAT, a DNS64 ([RFC6147]) optionally is in charge of the synthesis of AAAA records from the A records, so they can use a NAT64, without the need of doing a double-translation by means of the CLAT. However, the DNS64 is not useful for the IPv4-only devices or applications in the LANs, as they will not be able to use the AAAA records.

A typical 464XLAT deployment is depicted in Figure 1.




Figure 1: Typical 464XLAT Deployment

As it can be observed in the preceding picture, the situation is the same, regardless of in case of a wired network with a CE Router or a cellular network where a UE is connecting other devices (which may be IPv4-only or have IPv4-only apps), by means of a tethering functionality.

If the operator is providing direct access to Content Delivery Networks (CDNs), caches, or other resources, and they are dual-stacked, the situation can be described as shown in Figure 2.




Figure 2: Typical 464XLAT Deployment with CDNs/Caches

2. Requirements Language

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.

3. Problem Statement

If the devices or applications in the customer LAN are IPv6-capable, then the access to the CDNs, caches or other resources, will be made in an optimized way, by means of IPv6-only, not using the NAT64, as depicted in Figure 3.




Figure 3: 464XLAT access to CDNs/Caches by IPv6-capable apps

However, if the devices or applications are IPv4-only, for example, most of the SmartTVs and Set-Top-Boxes available today, a non-optimal double translation will occur (NAT46 at the CLAT and NAT64 at the PLAT), as illustrated in Figure 4.




Figure 4: 464XLAT access to CDNs/Caches by IPv4-only apps

Clearly, this is a non-optimal situation, as it means that even if there is a dual-stack service, the NAT46/CLAT translated IPv4 to IPv6 traffic flow, is unnecessarily translated back to IPv4, traversing the stateful NAT64. This has a direct impact in the need to scale the NAT64 beyond what will be actually needed if possible solutions,

in order to keep using the IPv6 path towards those services, are considered.

As shown in the Figure 4, this is also the case for many other services, not just CDNs or caches, such as VoIP access to the relevant operator infrastructure, which may be also dual-stack. This is true as well for many other dual-stack or IPv6-enabled services, which may be directly reachable from the operator infrastructure, even if they are not part of it, for example peering agreements, services in IXs, etc. In general, this will become a more frequent situation for many other services, which are not yet dual-stack.

For simplicity, across the rest of this document, references to CDNs/caches, should be understood, unless otherwise stated, as any dual-stacked resources.

This document looks into different possible solution approaches in order to optimize the IPv4-only SIIT translation providing a direct path to IPv6-capable services, as depicted in Figure 5.




Figure 5: Optimized 464XLAT access to CDNs/Caches by IPv4-only apps

4. Solution Approaches

4.1. Approach 1: DNS/Routing-based Solution

Because the IPv4-only devices will not be able to query for AAAA records, the NAT46/CLAT/CE will translate the IPv4 addresses from the A record for the CDN/cache destination, using the WKP or NSP, as configured by the operator.

If the CDN/cache provider is able to configure, in the relevant interfaces of the CDN/caches, the same IPv6 addresses that will naturally result as the translated destination addresses for the queried A records, preceded by the WKP or NSP, then having more specific routing prefixes, will result in traffic to those destinations being directly forwarded towards those interfaces, instead of needing to traverse the NAT64.

For example, let's suppose a provider using the WKP (64:ff9b::/96) and a SmartTV querying for www.example.com:

www.example.com	A	192.0.2.1
NAT46/CLAT translated to		64:ff9b::192.0.2.1
CDN IPv6 interface must be		64:ff9b::192.0.2.1
Operator must have a specific route to		64:ff9b::192.0.2.1

Note: Examples using text representation as per Section 2.3 of [RFC6052].

Because the WKP is non-routable, this solution will only be possible if the CDN/cache is in the same ASN as the provider network, or somehow interconnected without routing thru Internet.

This solution has the additional drawback of the operational complexity/issues added to the operation of the CDN/cache, and the need to synchronize any IPv4 interface address changes with the relevant IPv6 ones, and possibly with routing.

4.2. Approach 2: NAT46/CLAT/DNS-proxy-EAM-based Solution

If the NAT46/CLAT/CE, as commonly is the case, is also a DNS proxy/stub resolver, it is possible to modify the behavior and create an "internal" interaction among both of them.

This approach uses the existing IPv4 and IPv6 addresses in the A and AAAA records, respectively, so no additional complexity/issues added to the CDN/caches operations.

The following sub-sections detail this approach and provide a step-by-step example case.

4.2.1. Detection of IPv4-only devices or applications

The assumption is that, typically a dual-stack device will prefer using IPv6 as the DNS transport. So, when there is a DNS query, transported with IPv4, for an A record, and there is not a query for the AAAA record from the same IPv4 source (to the same destination), the DNS proxy/stub resolver can infer that, most probably, it is an

IPv4-only device or application.

It needs to be remarked that, if the detection of the IPv4-only device or application is done incorrectly (either not detecting it or by a false detection), no harm is caused. In the worst case, optimization will not be performed, at least, at the time being. However, optimization maybe performed later on, if a new detection succeeds (for example, another device using the same A record).

4.2.2. Detection of IPv6-enabled service

In the case of an IPv4-only detected device or application, the DNS proxy/stub resolver MUST actually perform an additional AAAA query, unless the information is already present in the Additional Section, as per Section 3 of [RFC3596]. Note that the NAT46/CLAT MUST already know the WKP or NSP being used in that network. If the response contains at least one IPv6 address not using the WKP/NSP, it means that the destination is IPv6-enabled (because at least one of the IPv6 addresses is not synthesized). This means that it is possible for the NAT46/CLAT, to create an Explicit Address Mapping ([RFC7757]).

4.2.3. Creation of EAMT entries

This way, an EAM Table (EAMT used for short, across the rest of this document) is created/maintained automatically by the DNS proxy/stub resolver in the NAT46/CLAT, and the NAT46/CLAT is responsible to prioritize any available entries in the EAMT, versus the use of any synthetic AAAA.

In order to create the EAMT entry, to determine if there is an AAAA record after an A record query, it is suggested to use the same delay value (50 milliseconds) as the "Resolution Delay" indicated by Happy Eyeballs [RFC8305]. This avoids a slight NAT64 overload and flapping between destination addresses (IPv4/IPv6), which may impact some applications, at the cost of a small extra delay for the initial communication setup, when the EAMT entry doesn't yet exist.

Each EAMT entry will contain, the fields already described in [RFC7757] and a few new ones:

1. ID: EAMT Entry Index (optional).
2. IPv4 address/prefix: By default, the prefix length is 32 bits.
3. IPv6 address/prefix: By default, the prefix length is 128 bits.
4. TTL: Because the optimization will make use of the AAAA (IPv6

address), the TTL for the EAMT entry must be the one of the AAAA RR. In normal conditions the TTL for both A and AAAA records, of a given FQDN, should be the same, so this ensures a proper behavior if there is any DNS mismatch.

5. FQDN: The one that originated the A query for this EAMT entry. Required in order to ensure a correct detection of cases such as the use of reverse-proxy with a single IPv4 address to multiple IPv6 addresses.
6. Valid/Invalid: When set to 1, means that this EAMT entry MUST NOT be used and consequently no optimization performed. It may be used also for an explicit configuration (GUI, CLI, provisioning system, etc.) to disallow optimization for any IPv4 addresses.
7. Auto/Static: When set to 1, means that this EAMT entry has been manually/statically configured, for example by means of an explicit configuration (GUI, CLI, provisioning system, etc.), so it doesn't expire with TTL.

When a new EAMT entry is first automatically created, it is marked as "Valid" and "Auto" (both bits cleared). If a subsequent A query, with a different FQDN, results in an IPv4 address that has already an EAMT entry and a different IPv6 address, it means that some reverse-proxy or similar functionality is being used by the IPv6-enabled service. In this case, the existing EAMT entry will be marked as "Invalid" (bit set). No new EAMT entry is created for that IPv4 address. Otherwise, the optimization will only allow to access the first set of IPv4/IPv6/FQDN, which may break the access to other FQDN that share the same IPv4 address and different IPv6 addresses.

In this case the EAMT entry will still expire according the TTL, which allows to re-enable optimization if a new query for the A record has changed the situation. For example, maybe the reverse-proxy has been removed, or there is now only a single device using it, so at the time being, the optimization is again possible without creating troubles to other hosts.

Note that when an EAMT entry is marked as "invalid", it will not affect the devices or applications, as they will still be able to use the regular CLAT+NAT64 flow, of course, without the optimization.

***** Open question regarding TTL and maybe FQDN and valid/auto bits. Is this always a good thing to do for EAM? Should this document update [RFC7757] to support this by default? Or it is just an "extension" as per section 3.1 of [RFC7757].

4.2.4. Forwarding path via stateful NAT for existing EAMT entries

Following this approach, if there is a valid EAMT entry, for a given IPv4-destination, the IPv6-native path pointed by the IPv6 address of that EAMT entry, will take precedence versus the NAT64 path, so the traffic will not be forwarded to the NAT64.

However, this is not sufficient to ensure that individual applications are able to keep existing connections. In many cases, audio and video streaming may use a single TCP connection lasting from minutes to hours. Instead, the CDN TTLs may be configured in the range from 10 to 300 seconds in order to allow new resolutions to switch quickly and to handle large recursive resolvers (with hundreds of thousands of clients behind them).

Consequently, the EAMT entries should not be used directly to establish a forwarding path, but instead, to create a stateful NAT entry for the 4-tuple for the duration of the session/connection.

4.2.5. Maintenance of the EAMT entries

The information in the EAMT MUST be kept timely-synchronized with the AAAA records TTL's, so the EAMT entries MUST expire on the AAAA TTL expiry and consequently be deleted.

However, EAMT entries with the Auto/Static bit set, will not be deleted.

4.2.6. Usage example

Using the same example as in the previous approach:

www.example.com	A	192.0.2.1
	AAAA	2001:db8::a:b:c:d
EAMT entry	192.0.2.1	2001:db8::a:b:c:d
NAT46/CLAT translated to		2001:db8::a:b:c:d
CDN IPv6 interface already is		2001:db8::a:b:c:d
Operator already has a specific route to		2001:db8::a:b:c:d

The following is an example of the CE behavior after the previous case has already created an EAMT entry and a reverse-proxy is detected:

1. A query for www.another-example.com A RR is received
2. www.another-example.com A 192.0.2.1
3. www.another-example.com AAAA 2001:db8::e:e:f:f

4. A conflict has been detected

5. The existing EAMT entry for 192.0.2.1 is set as invalid

4.2.7. Behavior in case of multiple A/AAAA RRs

If multiple A and/or AAAA records are available, the DNS proxy/stub resolver MUST follow existing procedures to choose each one. In other words, the chosen pair of A/AAAA records doesn't present any different result compared with a situation when this mechanism is not used.

4.2.8. Behavior in presence/absence of DNS64

This mechanism performs the same in both cases, if a DNS64 is present/used and if it is not present/used. This is explained because the mechanism is only relevant for destinations which don't have AAAA records, and in those cases DNS64 is not relevant. Furthermore, because as indicated in Section 4.2.2, the EAMT entry is not created when the service is IPv6-enabled. This is relevant because 464XLAT can be deployed/used with and without a DNS64.

4.2.9. Behavior when using literal addreses or non IPv6-compliant APIs

Because the EAMT entries are only created when the NAT46/CLAT/CE proxy/stub DNS is being used, any devices or applications that don't use DNS, will not create the relevant entries.

They will be however optimized if devices or applications using DNS, at some point, query for the same A RRs, or if EAMT entries are statically configured.

4.2.10. False detection of a dual-stack host as IPv4-only

If a dual-stack host is issuing the A query using IPv4 transport, and the AAAA query using IPv6 transport, or using different IPv4 addresses for the A and AAAA queries, the EAMT entry will be created. However, this EAMT entry may not be used by dual-stack devices or applications, because those devices or applications should prefer IPv6. If the host is preferring IPv4 for connecting to the CDN/cache or IPv6-enabled service, it will be actually using the NAT46/CLAT, including the EAMT entry and consequently IPv6, so this mechanism will be correcting an undesirable behavior. This is a special case, which actually seems to be an incoherent host or application implementation.

However, if other IPv4-only devices or applications subsequently need to connect to the same IPv6-enabled service, they will take advantage

of the already existing EAMT entry, and consequently use the IPv6-optimised path.

4.2.11. Behaviour in presence of Happy Eyeballs

Happy Eyeballs [RFC8305] is only available in dual-stack hosts. Consequently, is not affected by this mechanism because both, the A and the AAAA queries should be issued by the host as soon after one another as possible. However, if the same NAT46/CLAT/CE is serving IPv4-only hosts and dual-stack hosts and both of them are using the same destinations, an EAMT entry will be created for that destination. Consequently, a Happy Eyeballs fallback to IPv4 will actually be using the relevant EAMT entry IPv6 destination. This has the disadvantage that the IPv4-IPv6-IPv4 translation path can't be used by Happy Eyeballs-enabled applications. However, this may be actually considered as a good thing, in the sense that an operator is interested in knowing as soon as possible, if the IPv6-only network is not performing correctly, because that means also IPv4 will not be working. If the issue is related to extra IPv6 delay versus the IPv4 delay, Happy Eyeballs will not be able to offer a significative advantage here, but it looks like an acceptable trade-off.

Note that when using 464XLAT, the WAN link of the NAT46/CLAT/CE is IPv6-only. So even if Happy Eyeballs is present, the fallback to IPv4-only typically, will be slower than native IPv6 itself, because the added detail in the NAT46+NAT64 translations, when not using this optimization.

4.2.12. Behavior in case of Foreign DNS

Devices or applications may use DNS servers from other networks. For a complete description of reasons for that, refer to Section 4.4 of [I-D.ietf-v6ops-nat64-deployment]. In the case the DNS is modified, or some devices or applications use other DNS servers, the possible scenarios and the implications are:

- a. Devices configured to use a DNS proxy/resolver which is not the CE/NAT46/CLAT. In this case this optimization will not work, because the EAMT entry will not be created based on their own flows. Nevertheless, the EAMT entry may be created by other devices using the same destinations. However, the lack of EAMT entry, will not impact negatively in the user's devices/applications (the optimization is not performed). It should be noticed that users commonly, don't change the configuration of devices such as SmartTVs or STBs (if they do, some other functionalities, such as CDN/caches optimizations may not work as well), so this only happens typically if the vendor is doing it on-purpose and for good well-known reasons.

- b. DNS privacy/encryption. Hosts or applications that use mechanisms for DNS privacy/encryption, such as DoT ([RFC7858], [RFC8094]), DoH ([RFC8484]) or DoQ ([I-D.huitema-quic-dnsquic]), will not make use of the stub/proxy resolver, so the same considerations as for the previous case apply.
- c. Users that modify the DNS in their Operating Systems. This is quite frequent, however commonly Operating Systems are dual-stack, so aren't part of the problem statement described by this document and will not be adversely affected.
- d. Users that modify the DNS in the CE. This is less common. In this case, this optimization is not adversely affected, because it doesn't depend on the operator DNS, it works only based on the internal CE interaction between the NAT46/CLAT and the stub/proxy resolver. Note that it may be affected if the operator offers different "DNS views" or "split DNS", however this is not related to this optimization and will anyway impact in the other possible operator optimizations (i.e. CDN/cache features).
- e. Combinations of the above ones. No further impact, than the one already described, is observed.

4.3. Approach 3: NAT46/CLAT-provider-EAM-based Solution

Instead of using the DNS proxy/stub resolver to create the EAMT entries, the operator may push this table (or parts of it) into the CE/NAT46/CLAT, by using configuration/management mechanisms.

This solution has the advantage of not being affected by any DNS changes from the user (the EAMT is created by the operator) and ensures a complete control from the operator. However, it may impact the cases of devices with a DNS configured by the vendor.

In general, most of the considerations from the previous approach will apply.

One more advantage of this solution is that the EAMT pairs doesn't need to match the "real" IPv4/IPv6 addresses available in the A/AAAA records, as shown in the next example.

www.example.com	A	192.0.2.1
	AAAA	2001:db8::a:b:c:d
EAMT pulled/pushed entry	192.0.2.1	2001:db8::f:e:d:c
NAT46/CLAT translated to		2001:db8::f:e:d:c
CDN IPv6 interface already is		2001:db8::f:e:d:c
Operator already has a specific route to		2001:db8::f:e:d:c

EAMT may contain TTLs which probably are derived from DNS ones, or alternatively, a global TTL for the full table.

An alternative way to configure the table, is that the CE is actually pulling the table (or parts of it) from the operator infrastructure. In this case it will be mandatory that the entries have individual TTLs, again probably derived from the DNS ones.

The major drawback of this approach is that it requires a new protocol, or an extension to existing ones, in order to push or pull the EAMT, in addition to the possible impact in terms of bandwidth each time the CEs reboot, or an EAMT must be pushed to all the CEs, etc.

5. IPv6-only Services become accessible to IPv4-only devices/apps

One of the issues with the IPv6 deployment, is that those services which become IPv6-only in Internet, aren't reachable by IPv4-only devices and applications. This means that new content providers must support dual-stack even for new services, even while IPv4 public addresses aren't available.

If NAT46/CLAT/DNS-proxy-EAM approach (Section 4.2) is chosen, it can be complemented to resolve this issue, by means of making sure that IPv6-only destinations have one A resource record (even an invalid one), despite they aren't actually connected to IPv4. This will mean that those services will work fine if there is a NAT46/CLAT, and will have no impact if that one doesn't exist, not a different situation than not having an A resource record.

In fact, it may become an incentive for the IPv6 deployment in Internet services and provides the option to use an IPv4 address (maybe anycast) for the "non-valid" A resource record, that points to a "universal" web page (maybe hosted by IETF?) that displays a warning such as "Sorry, you don't IPv6 support in your operator, so this service is not available for you".

6. Conclusions

NAT46/CLAT/DNS-proxy-EAM approach (Section 4.2) seems the right solution for optimizing the access to dual-stack services, whether they are located inside or outside the ISP.

Having this type of optimization facilitates and increases the usage of IPv6, even for IPv4-only devices and applications, at the same time that decreases the use of the NAT64.

SIIT already has a SHOULD for EAM support. Should 464XLAT be updated

by this document so the CLAT has a MUST for EAM support?.

Should we recommend having A records for IPv6-only services in Internet? The A record may point to a "reserved" or "special" IPv4 address. A web page IPv4-only hosted by IETF(?) showing "sorry this web page/service is only available from IPv6 enabled operators"?.

Open question: Should we consider any other risks? If CE's implementing this optimization create troubles, it may bring the content providers to switch back to IPv4-only. So possible failure cases need to be carefully considered for every possible solution approach.

7. Security Considerations

This document does not have any new specific security considerations.

8. IANA Considerations

This document does not have any new specific IANA considerations, unless we decide to define a "special reserved IPv4 address".

9. Acknowledgements

The authors would like to acknowledge the inputs of Erik Nygren, Fred Baker, Martin Hunek and TBD ...

10. References

10.1. 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>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", STD 88, RFC 3596, DOI 10.17487/RFC3596, October 2003, <<https://www.rfc-editor.org/info/rfc3596>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.

- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <<https://www.rfc-editor.org/info/rfc6147>>.
- [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>>.
- [RFC7599] Li, X., Bao, C., Dec, W., Ed., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", RFC 7599, DOI 10.17487/RFC7599, July 2015, <<https://www.rfc-editor.org/info/rfc7599>>.
- [RFC7757] Anderson, T. and A. Leiva Popper, "Explicit Address Mappings for Stateless IP/ICMP Translation", RFC 7757, DOI 10.17487/RFC7757, February 2016, <<https://www.rfc-editor.org/info/rfc7757>>.
- [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont, "IP/ICMP Translation Algorithm", RFC 7915, DOI 10.17487/RFC7915, June 2016, <<https://www.rfc-editor.org/info/rfc7915>>.
- [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>>.
- [RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.

10.2. Informative References

- [I-D.huitema-quic-dnsquic]
Huitema, C., Shore, M., Mankin, A., Dickinson, S., and J. Iyengar, "Specification of DNS over Dedicated QUIC Connections", draft-huitema-quic-dnsquic-07 (work in progress), September 2019.

- [I-D.ietf-v6ops-nat64-deployment]
Palet, J., "Additional NAT64/464XLAT Deployment Guidelines in Operator and Enterprise Networks", draft-ietf-v6ops-nat64-deployment-08 (work in progress), July 2019.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8094] Reddy, T., Wing, D., and P. Patil, "DNS over Datagram Transport Layer Security (DTLS)", RFC 8094, DOI 10.17487/RFC8094, February 2017, <<https://www.rfc-editor.org/info/rfc8094>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.

Authors' Addresses

Jordi Palet Martinez
The IPv6 Company
Molino de la Navata, 75
La Navata - Galapagar, Madrid 28420
Spain

Email: jordi.palet@theipv6company.com
URI: <http://www.theipv6company.com/>

Alejandro D'Egidio
Telecentro
Argentina

Email: adegidio@telecentro.net.ar

v6ops
Internet-Draft
Intended status: Informational
Expires: September 10, 2020

J. Palet Martinez
The IPv6 Company
March 9, 2020

IPv6-only Terminology Definition
draft-palet-v6ops-ipv6-only-05

Abstract

This document defines the terminology regarding the usage of expressions such as "IPv6-only", in order to avoid confusions when using them in IETF and other documents. The goal is that the reference to "IPv6-only" describes the actual native functionality being used, not the actual protocol support.

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 September 10, 2020.

Copyright Notice

Copyright (c) 2020 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. Context	2
3. Definition of IPv6-only	3
4. Dual-stack	3
5. Native dual-stack	3
6. IPv6-only network	4
7. IPv6-only WAN/access	4
8. IPv6-only LAN	4
9. IPv6-only host/router	4
10. Transitional IPv6 host/router	4
11. Other cases	4
12. Security Considerations	5
13. IANA Considerations	5
14. Acknowledgements	5
Author's Address	5

1. Introduction

Due to the nature of the Internet and the different types of users, parts of a network, providers, flows, etc., there is not a single and easy way to categorically say something such as "IPv6-only".

The goal of this document is to depict this situation and agree in a common language to be used for IETF and other documents, in order to facilitate ourselves and future readers, the correct understanding of what we are talking about.

Note that all the references in this document are regarding the actual usage of IPv4/IPv6, not the support of those protocols by nodes. For example, a device or access network may support both IPv4 and IPv6, however actually is only "natively" forwarding IPv6, because the link used for that communication is only natively configured for IPv6. IPv4 may be used as well, but it is being encapsulated or translated by means of IPv6. So, from this perspective, this device is attached to an IPv6-only link.

2. Context

The transition from IPv4 to IPv6 is not something that can be done, in the large majority of the cases, overnight and in a single step. Consequently, in general, we are unable to talk about a whole network having a "single and uniform" status regarding the IPv6 support, at least not in the early deployment stages of an operator network.

Even if possible, it is not frequent to deploy new IPv6 networks which have no IPv4 connectivity at all, because at the current phase

of the universal goal of the IPv6 deployment, almost every network still need to provide some kind of "access" to IPv4 sites. It is not feasible for most of the operators to tell their customers "I can provide you IPv6 service, but you will not be able to access all Internet contents and apps, because some of them still don't support IPv6, so you will miss every content that it is IPv4-only".

Some networks may have IPv6-only support for specific purposes or services. For example, a DOCSIS provider may have decided that is worth the effort to get rid of IPv4 for the management network of the cable-modems. Or a network that provides connectivity only to IoT devices, may be IPv6-only.

However, the "end-networks", in general, need to continue supporting IPv4, as there are many devices or apps, in both corporate and end-user networks (smartTV, IP cameras, etc.), which are IPv4-only and it is not always feasible to update or replace them. Also if customer devices in a LAN are IPv6-only, they will not be able to access IPv6-only services, so this means that IPv6-only services can't be deployed unless it is done in such way that some transition mechanism solves that problem as well (example an IPv6-only Data Center, requires SIIT-DC)

In IPv6-only access networks, IPv4 support may be provided by mechanisms that allow "IPv4-as-a-service" (IPv4aaS, for example by means of encapsulation and/or translation on top of IPv6).

3. Definition of IPv6-only

Consequently, considering the context described in the section above, if we want to be precise and avoid confusing others, we can't use the terminology "IPv6-only" in a generic way, and we need to define what part of the network we are referring to.

From that perspective, we define the "IPv6-only" status in a given part(s) of a network, depending on if there is actual native forwarding of IPv4, so IPv4 is not configured neither managed.

4. Dual-stack

This can be applied to a host, router, link, network (part), etc. It means that both, IPv4 and IPv6 are reachable, without specifying how.

5. Native dual-stack

This can be applied to a host, router, link, network (part), etc. It means that both, IPv4 and IPv6 are configured/used natively (without the need of transition mechanisms).

6. IPv6-only network

IPv6-only can be used only if, a complete network, end-to-end, is actually not natively forwarding IPv4, which will mean that no-IPv4 addresses are configured, neither used for management, neither the network is providing transition/translation support, neither there is IPv4 transit/peering.

This is the end of the road of the IPv4-to-IPv6 transition, however we aren't there yet, in general at the time of writing this document, unless we are referring to special or disconnected (from IPv4) networks.

7. IPv6-only WAN/access

IPv6-only WAN or access can be used only if the WAN or access network isn't actually natively forwarding IPv4.

8. IPv6-only LAN

IPv6-only LAN(s) can be used only if the LAN(s), isn't actually natively forwarding IPv4.

9. IPv6-only host/router

IPv6-only host/router can be used only if the host/router, isn't actually using/forwarding IPv4, so IPv4 is unconfigured and/or disabled in the external facing interfaces.

Internal interfaces, such as loopback, can still be using IPv4 (internally).

10. Transitional IPv6 host/router

Transitional IPv6 host/router is a dual-stack host/router with IPv6-only WAN where IPv4 service support is provided by means of transition mechanism, IPv4aaS (IPv4-as-a-service).

11. Other cases

Similar other cases or parts of the network can be considered as IPv6-only if there is no actual native forwarding of IPv4 and in that case, after "IPv6-only" some word/short text pointing to the specific case or part of the network needs to be used. For instance, we could talk about "IPv6-only core" if a core network is only natively forwarding IPv6.

12. Security Considerations

This document does not have any specific security considerations.

13. IANA Considerations

This document does not have any IANA considerations.

14. Acknowledgements

The author would like to acknowledge the inputs from Tim Chown, Noah Maina, Lee Howard, Azael Fernandez Alcantara and Marcos Sanz Grosson.

Author's Address

Jordi Palet Martinez
The IPv6 Company
Molino de la Navata, 75
La Navata - Galapagar, Madrid 28420
Spain

Email: jordi.palet@theipv6company.com
URI: <http://www.theipv6company.com/>

v6ops
Internet-Draft
Intended status: Informational
Expires: May 7, 2020

J. Palet Martinez
The IPv6 Company
November 4, 2019

IPv6 Point-to-Point Links
draft-palet-v6ops-p2p-links-04

Abstract

This document describes different alternatives for configuring IPv6 point-to-point links, considering the prefix size, numbering choices and prefix pool to be used.

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 May 7, 2020.

Copyright Notice

Copyright (c) 2019 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. Requirements Language	3
3. The Ping-Pong Problem in Point-to-Point Links	3
4. Prefix Size Choices	3
4.1. Rationale for using /64	3
4.2. Rationale for using /127	4
4.3. Rationale for using /126 and Other Options	5
4.4. A Possible Middle-Term Choice	5
5. Numbering Choices	5
5.1. GUA (Global Unicast Addresses)	5
5.2. ULA (Unique Local Addresses)	5
5.3. Link-Local Addresses Only	6
6. Prefix Pool Choices	7
7. /64 from Customer Prefix for point-to-point links	7
7.1. Numbering Interfaces	7
7.2. Routing Aggregation of the Point-to-Point Links	8
7.3. DHCPv6 Considerations	9
7.4. Router Considerations	9
8. Security Considerations	10
9. IANA Considerations	10
10. Acknowledgements	10
11. References	10
11.1. Normative References	10
11.2. Informative References	11
Author's Address	12

1. Introduction

There are different alternatives for numbering IPv6 point-to-point links, and from an operational perspective, there may have different advantages or disadvantages that need to be taken in consideration under the scope of each specific network architecture design.

[RFC6164] describes using /127 prefixes for inter-router point-to-point links, using two different address pools, one for numbering the point-to-point links and another one for delegating the prefixes at the end of the point-to-point link. However, this doesn't exclude other choices.

This document describes alternative approaches, for the prefix size, the numbering of the link and the prefix pool.

The proposed approaches are suitable for those point-to-point links connecting ISP to customers, but not limited to those cases, and in fact, all them are being used by a relevant number of networks worldwide, in several different scenarios (service providers,

enterprise networks, etc.).

2. Requirements Language

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.

3. The Ping-Pong Problem in Point-to-Point Links

Some point-to-point links may present the ping-pong problem, (a forwarding loop). The fundamental root cause of this problem is an IPv6 implementations not performing full Neighbor Discovery (NS/NA) on addresses that the prefix says could exist on the link.

IPv6 implementations are assuming that all addresses within the prefix must exist at the other end of the point-to-point link, and send the traffic straight onto the link. If the address doesn't exist, and there is a covering route back in the other direction, the ping-pong problem occurs.

Full Neighbor Discovery is doing more than just resolving the link-layer address of an IPv6 address. Neighbor Discovery is also determining if the address exists. Even if a point-to-point link doesn't have link-layer addresses to resolve, ND determining if an address exists on the link is very beneficial because it will prevent the ping-pong problem occurring entirely regardless of the IPv6 prefix length being used on the link.

4. Prefix Size Choices

[RFC7608] already discusses about the IPv6 prefix length recommendations for forwarding, and the need for routing and forwarding implementations to ensure that longest-prefix-match works on any prefix length. So, in this document, we concentrate in the most commonly used choices, not excluding other options.

4.1. Rationale for using /64

The IPv6 Addressing Architecture ([RFC4291]) specifies that all the Interface Identifiers for all the unicast addresses (except for 000/3) are required to be 64 bits long and to be constructed in Modified EUI-64 format.

The same document also mandates the usage of the predefined subnet-router anycast address, which has cleared to zero all the bits that

do not form the subnet prefix.

Using /64 is the most common scenario and currently the best practice by the number of service providers using this approach compared to others.

Using a /64 has the advantage of being future proof and avoids renumbering, in the event that new standards take advantage of the 64 bits for other purposes, or the link becomes a point-to-multipoint, or there is a need to use more addresses in the link (e.g., monitoring equipment, managed bridges).

It has been raised also the issue of some hardware having limitations in using prefixes longer than /64, for example using extra hardware resources.

Section 5. of [RFC6164] describes possible issues when using /64 for the point-to-point links, such as the ping-pong and the neighbor cache exhaustion. However, it also states that they can be mitigated by other means, including the latest ICMPv6 [RFC4443] ND [RFC4861]. Indeed, considering the publication date of that document, those issues should not be any longer a concern. The fact is that many operators worldwide, today use /64 without any concerns, as vendors have taken the necessary code updates.

Consequently, we shall conclude that it is a valid approach to use /64 prefixes for the point-to-point links.

4.2. Rationale for using /127

[RFC6164] already do a complete review of reasons why /127 is a good approach vs other options. However, it needs to be considered that it was published a number of years ago, and most of the hardware today already incorporate mitigations.

It should be noted that, when using a /127 prefix, configuration of each of the addresses within the /127 prefix, at each respective end of the link, must be actively validated by the network operator. A missing /127 address from one end of the link, with a local route pointing out that end of the link that covers the missing /127 address, such as a default route, causes a "ping-pong" scenario to exist for the missing /127 address. The link could still be successfully carrying transit traffic, and IPv6 will not report any errors, because IPv6 doesn't require or nor check to ensure all interfaces attached to a link has addresses from all prefixes assigned to the link, excepting the Link-Local prefix per [RFC4291].

It is a valid approach to use /127 for the point-to-point links,

however is not future proof considering the comments from the previous section, and older equipment may not support it.

4.3. Rationale for using /126 and Other Options

/126 was considered by [RFC3627], and despite this document has been obsoleted, because was considering /127 as harmful, the considerations in Section 4.3 are still valid.

The same document describes options such as /112 and /120, and all those are commonly used in worldwide IPv6 deployments [IPv6-Survey], though in a lesser degree than /64 or /127.

Consequently, we shall conclude that /126, /120 and /112 are valid approaches for the point-to-point links.

4.4. A Possible Middle-Term Choice

A possible "middle-term" approach, will be to allocate a /64 for each point-to-point link, but use just one /127 out of it, making it future proof and at the same time avoiding possible issues indicated in the previous sections.

5. Numbering Choices

IPv6 provides different unicast addressing scopes which can be considered when numbering a point-to-point link.

It has been reported that certain hardware may consume resources when using numbered links. This is a very specific situation that may need to be consider on a case by case basis.

5.1. GUA (Global Unicast Addresses)

Using GUA is the most common approach. It provides full functionality for both end-points of the point-to-point link and consequently, facilitates troubleshooting.

5.2. ULA (Unique Local Addresses)

Some networks use ULAs for numbering the point-to-point links. This approach may cause numerous problems when carrying Internet traffic and therefore, is strongly discouraged. For example, if the CE needs to send an ICMPv6 message to a host outside that network (to the Internet), the packet with ULA source address will not get thru and PMTUD will break, which in turn will completely break that IPv6 connection when the MTU is not the same for all the path.

ULAs are IPv6 private addresses, not intended to be used as source or destination addresses across the Internet. This issue also exists in IPv4 when using [RFC1918] addresses on links carrying IPv4 Internet traffic. [RFC6752] discusses this issue for IPv4, with much of the discussion applying similarly to IPv6 and ULAs.

However, this approach is valid if, following Section 2.2 of [RFC4443], and despite using ULA for the point-to-point link, the router is configured with at least one GUA and the source of the ICMPv6 messages are always a GUA, per the IPv6 Default Address Selection algorithm [RFC6724].

5.3. Link-Local Addresses Only

Some networks leave the point-to-point links with only Link-Local addresses used at both ends of the link. This is sometimes improperly referred as "unnumbered", because the Link-Local addresses are also "numbers". Furthermore, [RFC4291] requires that all interfaces attached to a link have at least a Link-Local "number" or address from the Link-Local prefix.

[RFC7404] (Using Only Link-Local Addressing inside an IPv6 Network) discusses pros and cons of this alternative, which in general apply for the point-to-point links.

While this choice might work if the point-to-point link is terminated in a router, which typically will get configured with a suitable routable GUA or ULA, it will not work for devices that can't be further configured, for example if they do not support DHCPv6-PD. This is the case for hosts, when the Operating System is not expected to be a DHCPv6-PD client and are therefore left without any usable GUA to allow traffic forwarding.

In the case of a router, the route for the assigned prefix is pointed towards the link-local address on the router WAN port and the default route on the router is pointed towards the link-local address on the upstream network equipment port.

This choice seems easier to implement, compared the previous ones, but it also brings some drawbacks, such as difficulties with troubleshooting and monitoring. For example, link local addresses do not appear in traceroute, so it makes more difficult to locate the exact point of failure.

It is more useful in scenarios where it is known that only a router will be attached to the point-to-point link, and where the configured address of the router is known. Non-routers connecting to a network, which can't initiate DHCPv6-PD might experience problems and will

stay unnumbered upon connection, if a /64 prefix is not used to number the link. This may be also the case for routers, which will not be able to complete the DHCPv6-PD in unnumbered links.

The considerations indicated in the previous section, regarding not using ULA as source address of ICMPv6 messages, and instead ensuring there is at least one GUA configured for that, also apply if link-locals are used for the point-to-point link.

6. Prefix Pool Choices

The logic choice seems to use a dedicated pool of IPv6 addresses, as this is the way we are "used to" with IPv4. Actually, this is done often by means of different IPv6 pools at every PoP in a service provider network.

A possible benefit of using a dedicated IPv6 pool, is that allows applying security policies without harming the customers. This is only true if customers always have a CE at their end of the WAN link.

However, the fact that the default IPv6 link size is /64 and commonly multiple /64's are assigned to a single customer, provides an interesting alternative approach for combining "best practices" described in the precedent sections.

The following section depicts this alternative.

7. /64 from Customer Prefix for point-to-point links

Using a /64 from the customer prefix, in addition to the advantages already indicated when using /64, simplifies the addressing plan.

The use of /64 also facilitates an easier way for routing the shorter aggregated prefix into the point-to-point link. Consequently it simplifies the "view" of a more unified addressing plan, providing an easier path for following up any issue when operating IPv6 networks and typically, will have a great impact in saving expensive hardware resources (lower usage of TCAM, typically by half).

This mechanism would not work in broadcast layer two media that rely on ND, because it will try ND for all the addresses within the shorter prefix that is being routed thru the point-to-point link.

7.1. Numbering Interfaces

Often, in point-to-point links, hardware tokens are not available, or there is the need to keep certain bits (u, g) cleared, so the links can be manually numbered sequentially with most of the bits cleared

to zero. This numbering makes as well easier to remember the interfaces, which typically will become numbered as 0 (with 63 leading zero bits) for the provider side and 1 (with 63 leading zero bits) for the customer side.

Using interface identifiers as 0 and 1 is not only a very simple approach, but also a very common practice. Other different choices can as well be used as required in each case.

On the other hand, using the EUI-64, makes it more difficult to remember and handle the interfaces, but provides an additional degree of protection against port (actually address) scanning as described at [RFC7707].

7.2. Routing Aggregation of the Point-to-Point Links

Following this approach and assuming that a shorter prefix is typically delegated to a customer, for example a /48, it is possible to simplify the routing aggregation of the point-to-point links. Towards this, the point-to-point link may be numbered using the first /64 of the /48 delegated to the customer.

Let's see a practical example:

- o A service provider uses the prefix 2001:db8::/32 and is using 2001:db8:aaaa::/48 for a given customer.
- o Instead of allocating the point-to-point link from a different addressing pool, it may use 2001:db8:aaaa::/64 (which is the first /64 subnet from the 2001:db8:aaaa::/48) to number the link.
- o This means that, in the case the non-EUI-64 approach is used, the point-to-point link may be numbered as 2001:db8:aaaa::1/64 for the provider side and 2001:db8:aaaa::2/64 for the customer side.
- o Note that using the first /64 and interface identifiers 1 and 2 is a very common practice. However other values may be chosen according to each case specific needs.

In this way, as the same address pool is being used for both, the prefix and the point-to-point link, one of the advantages of this approach is to make very easy the recognition of the point-to-point link that belongs to a given customer prefix, or in the other way around, the recognition of the prefix that is linked by a given point-to-point link.

For example, making a trace-route to debug any issue to a given address in the provider network, will show a straight view, and it

becomes unnecessary one extra step to check a database that correlate an address pool for the point-to-point links and the customer prefixes, as all they are the same.

Moreover, it is possible to use the shorter prefix as the provider side numbering for the point-to-point link and keep the /64 for the customer side. In our example, it will become:

- o Point-to-point link at provider side: 2001:db8:aaaa::1/48
- o Point-to-point link at customer side: 2001:db8:aaaa::2/64

This provides one additional advantage as in some platforms the configuration may be easier saving one step for the route of the delegated prefix (no need for two routes to be configured, one for the delegated prefix, one for the point-to-point link). It is possible because the longest-prefix-match rule.

The behavior of this type of configuration has been successfully deployed in different operator and enterprise networks, using commonly available implementations with different routing protocols, including RIP, BGP, IS-IS, OSPF, along static routing, and no failures or interoperability issues have been reported.

7.3. DHCPv6 Considerations

As stated in [RFC3633], "the requesting router MUST NOT assign any delegated prefixes or subnets from the delegated prefix(es) to the link through which is received the DHCP message from the delegating router", however the approach described in this document is still useful in other DHCPv6 scenarios or non-DHCPv6 scenarios.

Furthermore, [RFC3633] was updated by Prefix Exclude Option for DHCPv6-based Prefix Delegation ([RFC6603]), precisely to define a new DHCPv6 option, which covers the case described by this document.

Moreover, [RFC3769] has no explicit requirement that avoids the approach described in this document.

7.4. Router Considerations

This approach is being used by operators in both, residential/SOHO and enterprise networks, so the routers at the customer end for those networks MUST support [RFC6603] if DHCPv6-PD is used.

In the case of Customer Edge Routers there is a specific requirement ([RFC7084]) WPD-8 (Prefix delegation Requirements), marked as SHOULD for [RFC6603]. However, in a scenario where the approach described

in this document is followed, together with DHCPv6-PD, the CE Router MUST support [RFC6603].

8. Security Considerations

This document does not have any new specific security considerations.

9. IANA Considerations

This document does not have any new specific IANA considerations.

10. Acknowledgements

The author would like to acknowledge the inputs of Mikael Abrahamsson, Brian Carpenter, Eric Vyncke, Mark Smith and TBD.

Acknowledge is also due to my co-authors of RIPE-690 (Best Current Operational Practice for Operators: IPv6 prefix assignment for end-users - persistent vs non-persistent, and what size to choose, <https://www.ripe.net/publications/docs/ripe-690>) and global community, which provided valuable inputs which have been key for this document.

Acknowledgement to co-authors, Cesar Olvera and Miguel Angel Diaz, of a previous related document (draft-palet-v6ops-point2point, 2006), as well as inputs for that version from Alain Durand, Chip Popoviciu, Daniel Roesen, Fred Baker, Gert Doering, Olaf Bonness, Ole Troan, Pekka Savola and Vincent Jardin, are also granted.

11. References

11.1. 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>>.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <<https://www.rfc-editor.org/info/rfc3633>>.
- [RFC3769] Miyakawa, S. and R. Droms, "Requirements for IPv6 Prefix Delegation", RFC 3769, DOI 10.17487/RFC3769, June 2004, <<https://www.rfc-editor.org/info/rfc3769>>.

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC6603] Korhonen, J., Ed., Savolainen, T., Krishnan, S., and O. Troan, "Prefix Exclude Option for DHCPv6-based Prefix Delegation", RFC 6603, DOI 10.17487/RFC6603, May 2012, <<https://www.rfc-editor.org/info/rfc6603>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.
- [RFC7084] Singh, H., Beebe, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", RFC 7084, DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.
- [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>>.

11.2. Informative References

- [IPv6-Survey] Palet Martinez, J., "IPv6 Deployment Survey (Residential/Household Services)", January 2018, <<https://indico.uknwf.org.uk/event/41/contribution/5/material/slides/0.pdf>>.
- [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>>.

- [RFC3627] Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful", RFC 3627, DOI 10.17487/RFC3627, September 2003, <<https://www.rfc-editor.org/info/rfc3627>>.
- [RFC6164] Kohno, M., Nitzan, B., Bush, R., Matsuzaki, Y., Colitti, L., and T. Narten, "Using 127-Bit IPv6 Prefixes on Inter-Router Links", RFC 6164, DOI 10.17487/RFC6164, April 2011, <<https://www.rfc-editor.org/info/rfc6164>>.
- [RFC6752] Kirkham, A., "Issues with Private IP Addressing in the Internet", RFC 6752, DOI 10.17487/RFC6752, September 2012, <<https://www.rfc-editor.org/info/rfc6752>>.
- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", RFC 7404, DOI 10.17487/RFC7404, November 2014, <<https://www.rfc-editor.org/info/rfc7404>>.
- [RFC7608] Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015, <<https://www.rfc-editor.org/info/rfc7608>>.
- [RFC7707] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", RFC 7707, DOI 10.17487/RFC7707, March 2016, <<https://www.rfc-editor.org/info/rfc7707>>.

Author's Address

Jordi Palet Martinez
The IPv6 Company
Molino de la Navata, 75
La Navata - Galapagar, Madrid 28420
Spain

EMail: jordi.palet@theipv6company.com
URI: <http://www.theipv6company.com/>