

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
 - 1.1. Applicability Statement
 - 1.2. Requirements Language
2. Generic Security Considerations
 - 2.1. Addressing
 - 2.1.1. Use of ULAs
 - 2.1.2. Point-to-Point Links
 - 2.1.3. Loopback Addresses
 - 2.1.4. Stable Addresses
 - 2.1.5. Temporary Addresses for SLAAC
 - 2.1.6. DHCP Considerations
 - 2.1.7. DNS Considerations
 - 2.1.8. Using a /64 per Host
 - 2.1.9. Privacy Consideration of Addresses
 - 2.2. Extension Headers
 - 2.2.1. Order and Repetition of Extension Headers
 - 2.2.2. Hop-by-Hop Options Header
 - 2.2.3. Fragment Header
 - 2.2.4. IP Security Extension Header
 - 2.3. Link-Layer Security
 - 2.3.1. Neighbor Solicitation Rate-Limiting
 - 2.3.2. Router and Neighbor Advertisements Filtering
 - 2.3.3. Securing DHCP
 - 2.3.4. 3GPP Link-Layer Security

- 2.3.5. Impact of Multicast Traffic
- 2.3.6. SEND and CGA
- 2.4. Control Plane Security
 - 2.4.1. Control Protocols
 - 2.4.2. Management Protocols
 - 2.4.3. Packet Exceptions
- 2.5. Routing Security
 - 2.5.1. BGP Security
 - 2.5.2. Authenticating OSPFv3 Neighbors
 - 2.5.3. Securing Routing Updates
 - 2.5.4. Route Filtering
- 2.6. Logging/Monitoring
 - 2.6.1. Data Sources
 - 2.6.2. Use of Collected Data
 - 2.6.3. Summary
- 2.7. Transition/Coexistence Technologies
 - 2.7.1. Dual Stack
 - 2.7.2. Encapsulation Mechanisms
 - 2.7.3. Translation Mechanisms
- 2.8. General Device Hardening
- 3. Enterprises-Specific Security Considerations
 - 3.1. External Security Considerations
 - 3.2. Internal Security Considerations
- 4. Service Provider Security Considerations
 - 4.1. BGP
 - 4.1.1. Remote Triggered Black Hole Filtering
 - 4.2. Transition/Coexistence Mechanism
 - 4.3. Lawful Intercept
- 5. Residential Users Security Considerations

- [6. Further Reading](#)
- [7. Security Considerations](#)
- [8. IANA Considerations](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Acknowledgements](#)
- [Authors' Addresses](#)

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 (L2) interactions are now done using the Neighbor Discovery Protocol (NDP) [RFC4861] rather than the Address Resolution Protocol [RFC0826]. Also, there is no Network Address Port Translation (NAPT) defined in [RFC2663] for IPv6 even if [RFC6296] specifies an IPv6-to-IPv6 Network Prefix Translation (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 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 Providers (SPs), enterprise networks, and some knowledgeable home-user-managed residential networks. This applicability statement especially applies to Sections 2.3 and 2.5.4.

1.2. 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.

2. Generic Security Considerations

2.1. Addressing

IPv6 address allocations and overall architecture are important parts of securing IPv6. Initial designs, even if intended to be temporary, tend to last much longer than expected. Although IPv6 was initially 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) or Provider-Aggregated (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; thereby, the complexity of the operations, including the security operations, is augmented.

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 Stateless Address Autoconfiguration (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 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 the addresses described in [RFC1918] and have different use cases. One use of ULAs 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 pseudorandom algorithm, per Section 3.2.1 of [RFC4193].

2.1.2. Point-to-Point Links

Section 5.1 of [RFC6164] 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 it also prevents a denial-of-service (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 [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 [RFC7707].

Stable addresses also allow easy enforcement of a security policy at the perimeter based on IPv6 addresses. For example, [Manufacturer Usage Description \(MUD\)](#) [RFC8520] is a mechanism where the perimeter defense can retrieve the security policy template based on the type of internal device and apply the right security policy based on the device's 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 64-bit Extended Unique Identifier (EUI-64) Media Access Control (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 privacy extension addresses as described in [RFC8981] might prevent the operator from building host-specific access control lists (ACLs). These 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; however, 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 PIOs. However, attackers could still find ways to bypass this mechanism if it is 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 another type of identifier, such as a DNS name. For DHCPv6, the lease is bound to the client DHCP Unique Identifier (DUID), which may or may not be bound to the client L2 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

In addition to 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 the "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.

[IPV6-EH-FILTERING] discusses filtering rules for those extension headers at transit routers.

2.2.1. Order and Repetition of Extension Headers

While [RFC8200] recommends the order and the maximum repetition of extension headers, at the time of writing, there are still IPv6 implementations that support an order of headers that is not recommended (such as Encapsulating Security Payload (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 [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 nonconforming 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; they have to process these packets in software and, hence, this task competes with other software tasks, such as handling the control and management plane processing.

Section 4.3 of [RFC8200], the current Internet Standard for IPv6, 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 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 (Authentication Header (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 that are 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 Neighbor Discovery (ND) Trust Models and Threats" [RFC3756] and in "Operational Neighbor Discovery Problems" [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 Section 2.3.1 below.

2.3.1. Neighbor Solicitation Rate-Limiting

NDP can be vulnerable to remote 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 cleverly managing the cache/timer.

[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:

- Ingress filtering of unused addresses by ACL. These require stable configuration of the addresses, e.g., 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).
- 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).
- Using a /127 on a point-to-point link, per [RFC6164].
- 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 source IP address assigned to DHCPv4 [RFC2131] or DHCPv6 [RFC8415] 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 L2 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 in 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 an 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 L2 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 handset) and the first-hop router (i.e., a Gateway GPRS Support Node (GGSN) or a Packet Data Network 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 DoS 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 Prefix Delegation (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 L2 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 Multicast DNS (mDNS) [RFC6762] and Link-Local Multicast Name Resolution (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:

- Neighbor Solicitation/Advertisement Spoofing
- Neighbor Unreachability Detection Failure
- Duplicate Address Detection DoS Attack
- Router Solicitation and Advertisement Attacks
- Replay Attacks
- Neighbor Discovery DoS Attacks

SEND does NOT:

- protect statically configured addresses
- protect addresses configured using fixed identifiers (i.e., EUI-64)
- provide confidentiality for NDP communications
- 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 routing 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:

- to drop illegitimate or potentially harmful control packets before they are queued to the RP (this can be done by a forwarding plane ACL) and
- 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 also be rate limited).

This section will consider several classes of control packets:

Control protocols:

routing protocols, such as OSPFv3, BGP, Routing Information Protocol Next Generation (RIPng), and, by extension, NDP and ICMP

Management protocols:

Secure Shell (SSH), SNMP, Network Configuration Protocol (NETCONF), RESTCONF, IP Flow Information Export (IPFIX), etc.

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, and ICMP.

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

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

Note: Dropping OSPFv3 packets that 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 [\[RFC8541\]](#) for a side benefit for OSPv3. The exact configuration will depend on the available resources of the router (CPU, Ternary Content-Addressable Memory (TCAM), etc.).

2.4.2. Management Protocols

This class includes SSH, SNMP, RESTCONF, NETCONF, gRPC Remote Procedure Calls (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:

- drop packets destined to the routers, except those belonging to protocols that are used (for example, permit TCP 22 and drop all others when only SSH is used) and
- 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:

- 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);
- 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);

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

On some routers, not everything can be done by the specialized data plane hardware that 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 update 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 Application-Specific Integrated Circuits (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-NULLEN for integrity protection. In early implementations, all OSPFv3 IPsec configurations relied on AH since the details weren't specified in [\[RFC5340\]](#). However, the document that specifically describes how IPsec should be implemented for OSPFv3 [\[RFC4552\]](#) states that "implementations **MUST** support ESP[-NULLEN] and **MAY** support AH" 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 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 rekeying mechanisms that do not cause outages. There have been instances where any rekeying causes outages; therefore, the trade-off 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**' implementation. Theoretically, it is possible that all communication between two IPv6 nodes, especially routers exchanging routing information, is encrypted using IPsec. However, in practice, 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 rekeying functionality.

2.5.4. Route Filtering

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

- filter internal-use IPv6 addresses that are not globally routable at the perimeter;
- discard routes for bogon [CYMRU] and reserved space (see [RFC8190]); and
- 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 Autonomous Systems (ASes) 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 is not limited to:

- logs of all applications using the network (including user space and kernel space) when available (for example, web servers that the network operator manages);
- data from [IP Flow Information Export](#) [RFC7011], also known as IPFIX;
- data from various SNMP MIBs [RFC4293] or YANG data via [RESTCONF](#) [RFC8040] or [NETCONF](#) [RFC6241];
- historical data of Neighbor Cache entries;
- [stateful DHCPv6](#) [RFC8415] lease cache, especially when a [relay agent](#) [RFC6221] is used;
- [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;
- firewall ACL logs;
- authentication server logs; and
- [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] in the European Union).

All those pieces of information can be used for:

- [forensic](#) (Section 2.6.2.1) investigations: who did what and when?

- [correlation \(Section 2.6.2.3\)](#): which IP addresses were used by a specific node (assuming the use of [privacy extensions addresses \[RFC8981\]](#))?
- [inventory \(Section 2.6.2.2\)](#): which IPv6 nodes are on my network?
- [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 cleartext (not binary). This can complicate the processing since one IPv6 address, for example, 2001:db8::1, can be written in multiple ways, such as:

- 2001:DB8::1 (in uppercase),
- 2001:0db8::0001 (with leading 0), and
- many other ways, including the reverse DNS mapping into a Fully Qualified Domain Name (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:

- `nextHeaderIPv6`, `sourceIPv6Address`, and `destinationIPv6Address`
- `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

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

- `ipIfStatsTable` table, which collects traffic counters per interface, and

- `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 address families and that can be used with [RFC6241] and [RFC8040]. This memo recommends the use of:

- `interfaces-state/interface/statistics` from [ietf-interfaces@2018-02-20.yang](#) [RFC8343], which contains counters for interfaces, and
- `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:

- using the [SNMP MIB](#) (Section 2.6.1.3), as explained above;
- using streaming telemetry or [NETCONF](#) [RFC6241] and [RESTCONF](#) [RFC8040] to collect the operational state of the Neighbor Cache; and
- 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 be fetched periodically at an interval that does not exhaust the router resources and still provides valuable information (the 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 an IPv6 address. Put another way, having access to the current and past content of the Neighbor Cache has a paramount value for the forensic and audit trails. 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 L2 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 that can only do stateless autoconfiguration) and also because DHCPv6 clients are identified not by their hardware-client address, as in IPv4, but by a DHCP Unique Identifier (DUID). The DUID can have several formats: the data-link layer address, the 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 L2 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 [\[RFC4649\]](#) [\[RFC6939\]](#), 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 lease files 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 the 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 be protected by using a L2 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 Access (WPA) or other [IEEE 802.1X \[IEEE-802.1X\]](#) wired interfaces 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):

- historical mappings of IPv6 addresses to users of remote access VPN and
- 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 in [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 one or more routers that 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 Database (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 Customer Premises Equipment (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

[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 an 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, and leveraging service discovery, such as mDNS [\[RFC6762\]](#) [\[RFC6763\]](#).

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 (i.e., Authoritative Transfers (AXFRs)) may be a further information source. An interesting research paper has analyzed 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 current (or historical) Neighbor Cache data set must be searched to find the data-link layer address of the IPv6 address. Next, 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 all IPv6 addresses associated with the same MAC address and interface in the [NDP cache \(Section 2.6.1.4\)](#).

2.6.2.4. Abnormal Behavior Detection

Abnormal behavior (such as network scanning, spamming, DoS) can be detected in the same way as in an IPv4 network:

- a sudden increase of traffic detected by interface counter (SNMP) or by aggregated traffic from [IPFIX records](#) [RFC7012],
- rapid growth of ND cache size, or
- change in traffic pattern (number of connections per second, number of connections per host, etc.) observed with the use of [IPFIX](#) [RFC7012].

2.6.3. Summary

While some data sources (IPFIX, MIB, switch Content Addressable Memory (CAM) tables, logs, etc.) 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, etc.). 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:

- ACLs to permit or deny traffic,
- firewalls with stateful packet inspection, and
- 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 that has 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 L2-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 [[RFC6169](#)]:

tunnel injection:

A malevolent actor knowing a few pieces of information (for example, the tunnel endpoints and the encapsulation protocol) can forge a packet that 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.

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 cleartext IPv6 packet. Combined with the absence of authentication, an on-path attack can also be mounted.

service theft:

As there is no authorization, even an unauthorized user can use a tunnel relay for free (this is a specific case of tunnel injection).

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.

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:

IP protocol 41: This will block Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) ([Section 2.7.2.2](#)), 6to4 ([Section 2.7.2.7](#)), 6rd ([Section 2.7.2.3](#)), and 6in4 ([Section 2.7.2.1](#)) tunnels.

IP protocol 47: This will block GRE ([Section 2.7.2.1](#)) tunnels.

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 [[RFC6324](#)]. This RFC also proposes mitigation techniques.

2.7.2.1. Site-to-Site Static Tunnels

Site-to-site static tunnels are described in [[RFC2529](#)] and in [GRE](#) [[RFC2784](#)]. As the IPv4 endpoints are statically configured and are not dynamic, they are slightly more secure (bidirectional 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 an untrusted 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 in [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 in 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 IPv6 Provider Edge (6PE) [RFC4798] and IPv6 Virtual Private Extension (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:

- address space, routing, and traffic separation with the help of VRFs (only applicable to 6VPE);
- hiding the IPv4 core, hence, removing all attacks against P-routers; and
- securing the routing protocol between Customer Edge (CE) and Provider Edge (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 [RFC7552].

2.7.2.5. DS-Lite

Dual-Stack 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) -- abbreviated 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, nonencapsulated 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-PT 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 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 UDPs 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 Internet Key Exchange Protocol (IKE), port 3478 for Session Traversal Utilities for NAT (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 made in cases when devices with older or operating systems that have not been updated 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 addresses 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 cybercrime investigations (for example, see the [Europol press release on CGN \[europol-cgn\]](#)). Many translation techniques (NAT64, DS-Lite, etc.) 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 that 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 and is thereby 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 a 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 NAT.

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 endpoints 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 attacks over IPv4 should be used for IPv6 and should include but are not limited to:

- restricting device access to authorized individuals;
- monitoring and auditing access to the device;
- turning off any unused services on the end node
- understanding which IPv6 addresses are being used to source traffic and changing defaults if necessary;
- using cryptographically protected protocols for device management (Secure Copy Protocol (SCP), SNMPv3, SSH, TLS, etc.);
- using host firewall capabilities to control traffic that gets processed by upper-layer protocols;
- applying firmware, OS, and application patches/upgrades to the devices in a timely manner;
- using multifactor credentials to authenticate to devices; and
- using 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 [RFC6092]). Section 3.2 of [RFC7381] also provides similar recommendations.

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

- Filter internal-use IPv6 addresses at the perimeter; this will also mitigate the vulnerabilities listed in [RFC7359].
- Discard packets from and to bogon and reserved space; see [CYMRU] and [RFC8190].
- Accept certain ICMPv6 messages to allow proper operation of ND and Path MTU Discovery (PMTUD); see [RFC4890] or [REY_PF] for hosts.
- 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, etc. This filtering should be done where applicable at the edge and possibly inside the perimeter; see [IPV6-EH-FILTERING].
- Filter packets having an illegal IPv6 header chain at the perimeter (and, if possible, inside the network as well); see Section 2.2.
- Filter unneeded services at the perimeter.
- Implement ingress and egress anti-spoofing in the forwarding and control planes; see [RFC2827] and [RFC3704].
- 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, e.g., University campus, wireless guest access, etc., 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. 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 or prefixes; 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, third-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 Provider Security Considerations

4.1. BGP

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

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

A Remote Triggered Black Hole (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](#)).

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 /128s (since each time the device establishes a data connection, it gets a new IID).

5. Residential Users Security Considerations

The IETF Home Networking (homenet) Working Group is working on standards and guidelines for IPv6 residential networks; this obviously includes operational security considerations, but this is still a 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 and 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. Therefore, it is 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](#)]:

outbound only:

Allowing all internally initiated connections and blocking 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.

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.

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

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:

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

7. 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.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. 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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

- [RFC8200]** Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

9.2. Informative References

- [CYMRU]** Team Cymru, "The Bogon Reference", <<https://team-cymru.com/community-services/bogon-reference/>>.
- [ENTROPYIP]** Foremski, P., Plonka, D., and A. Berger, "Entropy/IP: Uncovering Structure in IPv6 Addresses", November 2016, <<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]** European Union, "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)", Official Journal of the European Union, April 2016, <<https://eur-lex.europa.eu/eli/reg/2016/679/oj>>.
- [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-2020, February 2020.
- [IPv6-EH-FILTERING]** Gont, F. and W. Liu, "Recommendations on the Filtering of IPv6 Packets Containing IPv6 Extension Headers at Transit Routers", Work in Progress, Internet-Draft, draft-ietf-opsec-ipv6-eh-filtering-08, 3 June 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsec-ipv6-eh-filtering-08>>.
- [IPv6-EH-PARSING]** Kampanakis, P., "Implementation Guidelines for parsing IPv6 Extension Headers", Work in Progress, Internet-Draft, draft-kampanakis-6man-ipv6-eh-parsing-01, 5 August 2014, <<https://datatracker.ietf.org/doc/html/draft-kampanakis-6man-ipv6-eh-parsing-01>>.
- [IPv6_Security_Book]** Hogg, S. and É. Vyncke, "IPv6 Security", CiscoPress, ISBN 1587055945, 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 (NAv6TF) Technology Report "IPv6 Security Technology Paper", July 2006, <http://www.ipv6forum.com/dl/white/NAv6TF_Security_Report.pdf>.
- [NIST]** Frankel, S., Graveman, R., Pearce, J., and M. Rooks, "Guidelines for the Secure Deployment of IPv6", December 2010, <<http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf>>.
- [RADB]** Merit Network, Inc., "RADb: The Internet Routing Registry", <<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. J., 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>>.
- [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>>.
- [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>>.

-
- [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>>.

Acknowledgements

The authors would like to thank the following people for their useful comments (in alphabetical order): 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, and Bernie Volz.

Authors' Addresses

Éric Vyncke

Cisco
De Kleetlaan 6a
1831 Diegem
Belgium
Phone: +32 2 778 4677
Email: evyncke@cisco.com

Kiran Kumar Chittimaneni

Email: kk.chittimaneni@gmail.com

Merike Kaeo

Double Shot Security
3518 Fremont Ave N 363
Seattle, 98103
United States of America
Phone: [+12066696394](tel:+12066696394)
Email: merike@doubleshotsecurity.com

Enno Rey

ERNW
Carl-Bosch-Str. 4
69115 Heidelberg Baden-Wuerttemberg
Germany
Phone: [+49 6221 480390](tel:+496221480390)
Email: erey@ernw.de