OPSEC   E. Vyncke, Ed.
Internet-Draft   Cisco
Intended status: Informational   K. Chittimaneni
Expires: April 14, 2020   WeWork
                                                       M. Kaeo
                                          Double Shot Security
                                                        E. Rey
                                                          ERNW
                                              October 12, 2019

           Operational Security Considerations for IPv6 Networks
                        draft-ietf-opsec-v6-20

Abstract

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

   This document analyzes the operational security issues in several
   places of a network (enterprises, service providers and residential
   users) and proposes technical and procedural mitigations techniques.
   Some very specific places of a network such as the Internet of Things
   are not discussed in this document.

Status of This Memo

Table of Contents

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 differences between IPv4 and IPv6 especially with
   respect to security. For example, all layer-2 interactions are now
   done using Neighbor Discovery Protocol [RFC4861] rather than using
   Address Resolution Protocol [RFC0826].

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

   This document complements [RFC4942] by listing all security issues
   when operating a network utilizing varying transition technologies
   and updating with ones that have been standardized since 2007. It
   also provides more recent operational deployment experiences where
   warranted.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Generic Security Considerations

2.1. Addressing Architecture

IPv6 address allocations and overall architecture are an important part of securing IPv6. Initial designs, even if intended to be temporary, tend to last much longer than expected. Although initially IPv6 was thought to make renumbering easy, in practice, it may be extremely difficult to renumber without a proper IP Addresses Management (IPAM) system.

A key task before starting an IPv6 deployment, once the sufficient knowledge has been acquired, is to prepare an addressing plan. With the abundance of address space available, an addressing plan may be structured around services along with geographic locations, which then can be a basis for more structured security policies to permit or deny services between geographic regions.

A common question is whether companies should use Provider Independent (PI) vs Provider Allocated (PA) space [RFC7381], but from a security perspective there is little difference. However, one aspect to keep in mind is who has administrative ownership of the address space and who is technically responsible if/when there is a need to enforce restrictions on routability of the space e.g. due to malicious criminal activity originating from it.

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 to limit a host to only one IPv6 address per prefix. It also recommends that the network give the host the ability to use new addresses without requiring explicit requests (for example by using SLAAC). Having multiple IPv6 addresses per interface is a major change compared to the unique IPv4 address per interface; especially for audits (see section Section 2.6.2.3).

2.1.1. Use of ULAs

Unique Local Addresses (ULAs) [RFC4193] are intended for scenarios
where interfaces are not globally reachable, despite being routed
within a domain. They formally have global scope, but [RFC4193]
specifies that they must be filtered out at domain boundaries. ULAs
are different from [RFC1918] addresses and have different use cases.
One use of ULA is described in [RFC4864].

2.1.2. Point-to-Point Links

[RFC6164] in section 5.1 documents the reasons why it can make sense
to use a /127 for inter-router point-to-point links; notably, a /127
prevents the ping-pong attack between routers not implementing
correctly [RFC4443] and also prevents a DoS attack on the neighbor
cache. The previous recommendation of [RFC3627] has been obsoleted
and marked Historic by [RFC6547]).

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

2.1.3. Loopback Addresses

Many operators reserve a /64 block for all loopback addresses in
their infrastructure and allocates a /128 out of this reserved /64
prefix for each loopback interface. This practice allows for easy to
write Access Control List to enforce a security policy about those
loopback addresses.

2.1.4. Statically Configured Addresses

When considering how to assign statically configured addresses it is
necessary to take into consideration the effectiveness of perimeter
security in a given environment. There is a trade-off between ease
of operation (where some portions of the IPv6 address could be easily
recognizable for operational debugging and troubleshooting) versus
the risk of trivial scanning used for reconnaissance. [SCANNING]
shows that there are scientifically based mechanisms that make
scanning for IPv6 reachable nodes more feasible than expected; see
also [RFC7707]. The use of well-known (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 and more to
come possible, hence, operators should not rely on the 'impossible to
find because my address is random' paradigm, even if it is common

practice to have the statically configured addresses randomly
distributed across /64 subnets and to always use DNS.

While in some environments obfuscating addresses could be considered
an added benefit; it does not preclude that perimeter rules are
actively enforced and that statically configured addresses follow
some logical allocation scheme for ease of operation (as simplicity
always helps security). Typical deployments will have a mix of
static and non static addresses.

2.1.5. Temporary Addresses - Privacy Extensions for SLAAC

Historically stateless address autoconfiguration (SLAAC) relied on an
automatically generated 64-bit interface identifier (IID) based on
the EUI-64 MAC address, which together with the /64 prefix makes up
the globally unique IPv6 address. The EUI-64 address is generated
from the 48-bit stable MAC address. [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
[RFC4941] or [RFC8064].

Randomly generating an interface ID, as described in [RFC4941], is
part of SLAAC with so-called privacy extension addresses and 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 somehow
reduce the attack exposure window.

Using [RFC4941] privacy extension addresses might prevent the
operator from building host specific access control lists (ACLs). As
[RFC4941] privacy extension addresses could also be used to obfuscate
some malevolent activities (whether on purpose or not), specific user
attribution/accountability procedures should be put in place as
described in Section 2.6.

[RFC8064] 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 address on different networks.

In some extreme use cases where user accountability is more important
than user privacy, network operators may consider to disable SLAAC
and rely only on DHCPv6; but, not all operating systems support
DHCPv6 so some hosts will not get any IPv6 connectivity. Disabling
SLAAC and privacy extensions addresses can be done for most OS and
for non-hacker users by sending RA messages with a hint to get
addresses via DHCPv6 by setting the M-bit but also disabling SLAAC by
resetting all A-bits in all prefix information options. However

attackers could still find ways to bypass this mechanism if not
enforced at the switch/ router level.

However, in scenarios where anonymity is a strong desire (protecting
user privacy is more important than user attribution), privacy
extension addresses should be used. When [RFC8064] is 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/DNS Considerations

Many environments use DHCPv6 to provision addresses and other
parameters in order to ensure audit-ability and traceability (but see
Section 2.6.1.5). A main security concern is the ability to detect
and counteract against rogue DHCP servers (Section 2.3.3). It must
be noted that as opposed to DHCPv4, DHCPv6 can lease several IPv6
addresses per client and the lease is not bound to the link-layer
address of the client but to the DHCP Unique ID (DUID) of the client
that is not always bound to the client link-layer address.

While there are no fundamental differences with IPv4 and IPv6
security concerns about DNS, there are specific consideration in
DNS64 [RFC6147] environments that need to be understood.
Specifically the interactions and the potential of interference with
DNSSEC implementation need to be understood - these are pointed out
in more detail in Section 2.7.3.2.

2.1.7. Using a /64 per host

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

2.1.8. Privacy consideration of Addresses

The reader can learn more about privacy considerations for IPv6
addresses in [RFC7721].

2.2. Extension Headers

The extension headers are an important difference between IPv4 and
IPv6. The packet structure does make a big difference. For
instance, it's trivial to find (in IPv4-based packets) the upper
layer protocol type and protocol header, while in IPv6 it actually
isn't as the extension header chain must be parsed completely. The

IANA has closed the existing empty "Next Header Types" registry to
new entries and is redirecting its users to a new "IPv6 Extension
Header Types" registry per [RFC7045].

They 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 varying extension headers is important and
this section enumerates the ones that need careful consideration.

A clarification on how intermediate nodes should handle existing
packets with extension headers and any extension headers that are
defined in the future is found in [RFC7045]. The uniform TLV format
to be used for defining future extension headers is described in
[RFC6564].

It must also be noted that there is no indication in the packet
whether the Next Protocol field points to an extension header or to a
transport header. This may confuse some filtering rules.

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

2.2.1. Order and Repetition of Extension Headers

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

A firewall or any edge device should be used to enforce the
recommended order and number of occurrences of extension headers.

2.2.2. Hop-by-Hop Options Header

The hop-by-hop options header, when present in an IPv6 packet, forces
all nodes in the path to inspect this header in the original IPv6
specification [RFC2460]. This was of course a large avenue for a
denial of service as most if not all routers cannot process this kind
of packets in hardware but have to 'punt' this packet for software
processing. Section 4.3 of the current Internet Standard for IPv6,

[RFC8200], is more sensible to this respect as the processing of hop-
by-hop options header by intermediate routers is optional.

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 to:

> 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).

Else, stateless filtering could be bypassed by a hostile party.
[RFC6980] applies a stricter rule to NDP by enforcing the drop of
fragmented NDP packets. [RFC7113] describes how RA-guard function
described in [RFC6105] should behave in presence of fragmented RA
packets.

2.2.4. IP Security Extension Header

The IPsec [RFC4301] [RFC4301] extension headers (AH [RFC4302] and ESP
[RFC4303]) are required if IPsec is to be utilized for network level
security functionality.

2.3. Link-Layer Security

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

2.3.1. ND/RA Rate Limiting

Neighbor Discovery (ND) can be vulnerable to denial of service (DoS)
attacks in which a router is forced to perform address resolution for
a large number of unassigned addresses. Possible side effects of
this attack preclude new devices from joining the network or even
worse rendering the last hop router ineffective due to high CPU
usage. Easy mitigative steps include rate limiting Neighbor

Solicitations, restricting the amount of state reserved for
unresolved solicitations, and clever cache/timer management.

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

o Ingress filtering of unused addresses by ACL. These require
   static configuration of the addresses; for example, allocating the
   addresses out of a /120 and using a specific ACL to only allow
   traffic to this /120 (of course, the actual hosts are configured
   with a /64 prefix for the link).

o Tuning of NDP process (where supported).

o Using /127 on point-to-point link per [RFC6164].

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

Additionally, IPv6 ND uses multicast extensively for signaling
messages on the local link to avoid broadcast messages for on-the-
wire efficiency. However, this has some side effects on wireless
networks, especially a negative impact on battery life of smartphones
and other battery operated devices that are connected to such
networks. The following drafts are actively discussing methods to
rate limit RAs and other ND messages on wifi networks in order to
address this issue:

o [I-D.thubert-savi-ra-throttler]

o [I-D.chakrabarti-nordmark-6man-efficient-nd]

2.3.2. RA/NA Filtering

Router Advertisement spoofing is a well-known attack vector and has
been extensively documented. The presence of rogue RAs, either
intentional or malicious, can cause partial or complete failure of
operation of hosts on an IPv6 link. For example, a host can select
an incorrect router address which can be used as a man-in-the-middle
(MITM) attack or can assume wrong prefixes to be used for stateless
address configuration (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. An
attacker can conceal the attack by fragmenting his 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 in 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 current 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 Neighbor Discovery attacks.

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

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 line of defense is fully effective when weird fragments
are dropped by routers and switches as described in Section 2.2.3.
The generated log should also be analyzed to act on violations.

A drastic technique to prevent all NDP attacks is based on isolation
of all hosts with specific configurations. 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 set-up can be established on some switches or wireless
access points. Of course, this is not always easily feasible when
hosts need to communicate with other hosts.

2.3.3. Securing DHCP

Dynamic Host Configuration Protocol for IPv6 (DHCPv6), as detailed 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 and 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. A
malicious DHCP server is established with the intent of providing
incorrect configuration information to the client to cause a denial
of service attack or to mount a-man-in-the-middle 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].

[RFC7610], DHCPv6-Shield, specifies a mechanism for protecting
connected DHCPv6 clients against rogue DHCPv6 servers. This
mechanism is based on DHCPv6 packet-filtering at the layer-2 device;
the administrator specifies the interfaces connected to DHCPv6
servers. Furthermore, 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 an end host (the mobile
hand-set) and the first-hop router (i.e., a GPRS Gateway Support Node
(GGSN) or a Packet Gateway (PGW)) on that link. The GGSN/PGW never
configures a non link-local address on the link using the advertised
/64 prefix on it. The advertised prefix must not be used for on-link
determination. There is no need for an 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 built
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
one).

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

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

2.3.5. SeND and CGA

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

SeND protects against:

o Neighbor Solicitation/Advertisement Spoofing

o Neighbor Unreachability Detection Failure

o Duplicate Address Detection DoS Attack

o Router Solicitation and Advertisement Attacks

o Replay Attacks

o Neighbor Discovery DoS Attacks

SeND does NOT:

o Protect statically configured addresses

o Protect addresses configured using fixed identifiers (i.e. EUI-
   64)

o Provide confidentiality for NDP communications

o Compensate for an unsecured link - SEND does not require that the
   addresses on the link and Neighbor Advertisements correspond

However, at this time and after many years after their
specifications, CGA and SeND do not have wide support from generic
operating systems; hence, their usefulness is limited and should not
be relied upon.

2.4. **Control Plane Security**

   [RFC6192] defines the router control plane. 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).

   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 determine the best outgoing
   interface towards the destination, and forwarding the packet out
   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
   (named router 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 OSPF or BGP adjacencies which can cause a severe
   network disruption.

   The mitigation technique is:

   o To drop non-legit control packet before they are queued to the RP
      (this can be done by a forwarding plane ACL) and

   o To rate limit the remaining packets to a rate that the RP can
      sustain. Protocol specific protection should also be done (for
      example, a spoofed OSPFv3 packet could trigger the execution of
      the Dijkstra algorithm, therefore the number of Dijsktra execution
      should be also rate limited).

   This section will consider several classes of control packets:

   o Control protocols: routing protocols: such as OSPFv3, BGP and by
      extension Neighbor Discovery and ICMP

   o Management protocols: SSH, SNMP, IPfix, etc

o Packet exceptions: which are normal data packets which requires a
   specific processing such as generating a packet-too-big ICMP
   message or having the hop-by-hop options header.

2.4.1. Control Protocols

This class includes OSPFv3, BGP, NDP, ICMP.

An ingress ACL to be applied on all the router interfaces should be
configured such as:

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

o allow BGP (identified by TCP port 179) packets from all BGP
   neighbors and drop the others

o allow all ICMP packets (transit and to the router interfaces)

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

Rate limiting of the valid packets should be done. The exact
configuration obviously depends on the available ressources of the
router (CPU, TCAM, ...).

2.4.2. Management Protocols

This class includes: SSH, SNMP, 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 such as:

o Drop packets destined to the routers except those belonging to
   protocols which are used (for example, permit TCP 22 and drop all
   when only SSH is used);

o Drop packets where the source does not match the security policy,
   for example if SSH connections should only be originated from the
   NOC, then the ACL should permit TCP port 22 packets only from the
   NOC prefix.

Rate limiting of the valid packets should be done. The exact
configuration obviously depends on the power of the Route Processor.

2.4.3. **Packet Exceptions**

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

   o generation of an ICMP packet-too-big message when a data plane
      packet cannot be forwarded because it is too large;

   o generation of an ICMP hop-limit-expired message when a data plane
      packet cannot be forwarded because its hop-limit field has reached
      0;

   o generation of an ICMP destination-unreachable message when a data
      plane packet cannot be forwarded for any reason;

   o processing of the hop-by-hop options header, new implementations
      follow section 4.3 of [RFC8200] where this processing is optional;

   o or more specific to some router implementation: an oversized
      extension header chain which cannot be processed by the hardware
      and force the packet to be punted to the generic router CPU.

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

   An ingress ACL cannot help to mitigate a control plane attack using
   those packet exceptions. The only protection for the RP is to limit
   the rate of those packet exceptions forwarded to the RP, this means
   that some data plane packets will be dropped without any ICMP
   messages back to the source which may cause Path MTU holes.

   In addition to limiting the rate of data plane packets queued to the
   RP, it is also important to limit the generation rate of ICMP
   messages both the save the RP but 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

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

   1. Authenticating neighbors/peers

   2. Securing routing updates between peers

   3. Route filtering

   [RFC7454] covers these sections specifically for BGP in detail.

   [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.

2.5.1. Authenticating Neighbors/Peers

   A basic element of routing is the process of forming adjacencies,
   neighbor, or peering relationships with other routers. From a
   security perspective, it is very important to establish such
   relationships only with routers and/or administrative domains that
   one trusts. A traditional approach has been to use MD5 HMAC, which
   allows routers to authenticate each other prior to establishing a
   routing relationship.

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

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

indeed been issued by a router that had access to the shared
authentication key.

With all authentication mechanisms, operators should confirm that
implementations can support re-keying mechanisms that do not cause
outages. There have been instances where any re-keying cause outages
and therefore the tradeoff between utilizing this functionality needs
to be weighed against the protection it provides.

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

2.5.2. Securing Routing Updates Between Peers

IPv6 initially mandated the provisioning of IPsec capability in all
nodes. However, in the updated IPv6 Nodes Requirement standard
[RFC8504] is a 'SHOULD' and no more a 'MUST' implement.
Theoretically it is possible that communication between two IPv6
nodes, especially routers exchanging routing information be encrypted
using IPsec. In practice however, deploying IPsec is not always
feasible given hardware and software limitations of various platforms
deployed, it has also an operational cost as described in the earlier
section.

2.5.3. Route Filtering

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

o Prevent IP source address spoofing when applicable by applying
   [RFC2827];

o Filter internal-use, non-globally routable IPv6 addresses at the
   perimeter;

o Discard packets from and to bogon and reserved space (see [CYMRU]
   and [RFC8190]);

o Configure ingress route filters that validate route origin, prefix
   ownership, etc. through the use of various routing databases,
   e.g., RADB. There is additional work being done in this area to
   formally validate the origin ASs of BGP announcements in [RFC8210]

Some good recommendations for filtering can be found from Team CYMRU at [CYMRU]. [RFC7454] is another valuable resource of guidance in this space.

2.6. Logging/Monitoring

In order to perform forensic research in case of any security incident or to detect abnormal behaviors, network operators should log multiple pieces of information.

This includes:

o logs of all applications when available (for example web servers);

o use of IP Flow Information Export [RFC7011] also known as IPfix;

o use of SNMP MIB [RFC4293];

o use of historical data of Neighbour Cache entries;

o use of stateful DHCPv6 [RFC8415] lease cache, especially when a
   relay agent [RFC6221] is used;

o use of Source Address Validation Improvement (SAVI) [RFC7039]
   events, especially the binding of an IPv6 address to a MAC address
   and a specific switch or router interface;

o use of RADIUS [RFC2866] for accounting records.

Please note that there are privacy issues or regulations related to how those logs are collected, kept and safely discarded. Operators are urged to check their country legislation (e.g. GDPR in the European Union).

All those pieces of information will be used for:

o forensic (Section 2.6.2.1) investigations such as who did what and
   when?

o correlation (Section 2.6.2.3): which IP addresses were used by a
   specific node (assuming the use of privacy extensions addresses
   [RFC4941])

o inventory (Section 2.6.2.2): which IPv6 nodes are on my network?

o abnormal behavior detection (Section 2.6.2.4): unusual traffic
   patterns are often the symptoms of a abnormal behavior which is in

turn a potential attack (denial of services, network scan, a node
being part of a botnet, ...)

2.6.1. Data Sources

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

2.6.1.1. Logs of Applications

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

o 2001:DB8::1 (in uppercase)

o 2001:0db8::0001 (with leading 0)

o and many other ways including the reverse DNS mapping into a FQDN
   (which should not be trusted).

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

For example, this perl script can be used:

```
   <CODE BEGINS>

   #!/usr/bin/perl -w
   use strict ;
   use warnings ;
   use Socket ;
   use Socket6 ;

   my (@words, $word, $binary_address) ;

   ## go through the file one line at a time
   while (my $line = <STDIN>) {
     chomp $line;
     foreach my $word (split /[\s+]/, $line) {
       $binary_address = inet_pton AF_INET6, $word ;
       if ($binary_address) {
         print inet_ntop AF_INET6, $binary_address ;
       } else {
         print $word ;
       }
       print " " ;
     }
     print "\n" ;
   }

   <CODE ENDS>
```

2.6.1.2. IP Flow Information Export by IPv6 Routers

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

   o in section 5.4 (IP Header fields): nextHeaderIPv6 and
      sourceIPv6Address;

   o in section 5.6 (Sub-IP fields) sourceMacAddress.

   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 by IPv6 Routers

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

   o ipIfStatsTable table which collects traffic counters per
      interface;

   o ipNetToPhysicalTable table which is the content of the Neighbor
      cache, i.e. the mapping between IPv6 and data-link layer
      addresses.

2.6.1.4. Neighbor Cache of IPv6 Routers

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

   o the SNMP MIB (Section 2.6.1.3) as explained above;

   o using streaming telemetry or NETCONF [RFC6241] to collect the
      state of the neighbor cache;

   o also by connecting over a secure management channel (such as SSH)
      and explicitly requesting a neighbor cache dump via the Command
      Line Interface or any other monitoring mechanism.

   The neighbor cache is highly dynamic as mappings are added when a new
   IPv6 address appears on the network (could be quite often with
   privacy extension addresses [RFC4941] 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 typical host such as Windows 7). This means that
   the content of the neighbor cache must periodically be fetched at an
   interval which does not exhaust the router ressources and still
   provides valuable information (suggested value is 30 seconds but to
   be checked in the actual set-up) and stored for later use.

   This is an important source of information because it is trivial (on
   a switch not using the SAVI [RFC7039] algorithm) to defeat the
   mapping between data-link layer address and IPv6 address. Let us
   rephrase the previous statement: having access to the current and
   past content of the neighbor cache has a paramount value for forensic
   and audit trail.

   Using the approach of one /64 per host (Section 2.1.7) or DHCP-PD
   replace the neighbor cache dumps by a mere caching of the allocated

/64 prefix when combined with strict enforcement rule on the router
and switches to prevent IPv6 spoofing.

2.6.1.5. Stateful DHCPv6 Lease

In some networks, IPv6 addresses/prefixes are managed by 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 it was
usually done in the IPv4 era.

It is not so easy in the IPv6 era because not all nodes will use
DHCPv6 (there are nodes which can only do stateless
autoconfiguration) but also because DHCPv6 clients are identified not
by their hardware-client address as in IPv4 but by a DHCP Unique ID
(DUID) which can have several formats: some being the data-link layer
address, some being data-link layer address prepended with time
information or even an opaque number which is useless for operation
security. Moreover, when the DUID is based on the data-link address,
this address can be of any interface of the client (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 the layer-2
switches, then the DHCP server also receives the Interface-ID
information which could be save in order to identify the interface of
the switches which received a specific leased IPv6 address. Also, if
a 'normal' (not lightweight) relay agent adds the data-link layer
address in the option for Relay Agent Remote-ID [RFC4649] or
[RFC6939], then the DHCPv6 server can keep track of the data-link and
leased IPv6 addresses.

In short, the DHCPv6 lease file is less interesting than in the IPv4
era. DHCPv6 servers that keep the relayed data-link layer address in
addition to the DUID in the lease file do not suffer from this
limitation.

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

2.6.1.6. RADIUS Accounting Log

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

2.6.1.7. Other Data Sources

   There are other data sources that must be kept exactly as in the IPv4
   network:

   o historical mapping of IPv6 addresses to users of remote access
      VPN;

   o historical mapping of MAC address to switch interface in a wired
      network.

2.6.2. Use of Collected Data

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

2.6.2.1. Forensic and User Accountability

   The forensic use case is when the network operator must locate an
   IPv6 address that was present in the network at a certain time or is
   still currently in the network.

   To locate an IPv6 address in a enterprise network where the operator
   has control over all ressources, the source of information can be, in
   decreasing order, neighbor cache, DHCP lease file. Then, the
   procedure is:

   1. based on the IPv6 prefix of the IPv6 address find the router(s)
       which is(are) used to reach this prefix (assuming that anti-
       spoofing mechanisms are used);

   2. based on this limited set of routers, on the incident time and on
       IPv6 address to retrieve the data-link address from live neighbor
       cache, from the historical data of the neighbor cache 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 on which switch
    interface was this data-link layer address. In the case of
    wireless LAN, the RADIUS log should have the mapping between user
    identification and the MAC address. If a Configuration
    Management Data Base (CMDB) is used, the mapping between the
    data-link layer address and a switch port.

At the end of the process, the interface the host originating
malicious activity or the username which was abused for malicious
activity has been determined.

To identify the subscriber of an IPv6 address is a residential
Internet Service Provider, the main source will be the DHCP-PD leased
prefix which will often be linked to a subscriber via the RADIUS log.
Alternatively, the Forwarding Information Base of the CMTS or BNG
will indicate the CPE of the subscriber and the RADIUS log can be
used to retrieve the actual subscriber.

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

2.6.2.2. Inventory

RFC 7707 [RFC7707] is about the difficulties for an attacker to scan
an IPv6 network due to the vast number of IPv6 addresses per link
(and why in some case it can still be done). While the huge
addressing space can sometime be perceived as a 'protection', it also
make 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 operation of a network.

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

The first technique is to use the IPfix information and extract the
list of all IPv6 source addresses to find all IPv6 nodes that sent
packets through a router. This is very efficient but alas will not
discover silent node that never transmitted such packets. 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 works only for local network, it consists in sending a
ICMP ECHO_REQUEST to the link-local multicast address ff02::1 which

is all IPv6 nodes on the network. All nodes should reply to this
ECHO_REQUEST per [RFC4443].

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

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

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 was 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 with canonical IPv6 addresses. Then, the neighbor
cache current (or historical) data set must be searched to find the
data-link layer address of the IPv6 address. Then, the current and
historical neighbor cache data sets must be searched for all IPv6
addresses associated to this data-link layer address: this is the
search set. The last step is to search in all log files (containing
only IPv6 address in canonical format) for any IPv6 addresses in the
search set.

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

2.6.2.4. Abnormal Behavior Detection

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

o sudden increase of traffic detected by interface counter (SNMP) or
   by aggregated traffic from IPfix records [RFC7012];

   o change of traffic pattern (number of connection per second, number
      of connection per host...) with the use of IPfix [RFC7012]

2.6.3. Summary

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

   The fact that there are multiple ways to express in a character
   string the same IPv6 address 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
   way, 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.

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 traffics are native (easier to
   observe and secure) and should have the same network processing
   (path, quality of service, ...). Dual stack allows to gradually turn
   IPv4 operations off when your 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 you
   have twice the exposure. One needs to think about protecting both
   protocols now. At a minimum, the IPv6 portion of a dual stacked
   network should maintain parity with IPv4 from a security policy point
   of view. Typically, the following methods are employed to protect
   IPv4 networks at the edge or security perimeter:

   o ACLs to permit or deny traffic;

   o Firewalls with stateful packet inspection.

   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, else the

attacker will use the protocol version having the more relax 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.

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

For many years, all host operating systems have IPv6 enabled by
default, so, it is possible even in an 'IPv4-only' network to attack
layer-2 adjacent victims over their IPv6 link-local address or over a
global IPv6 address once rogue RAs or rogue DHCPv6 addresses are
provided by an attacker.

2.7.2. Encapsulation Mechanisms

There are many tunnels used for specific use cases. Except when
protected by IPsec [RFC4301], all those tunnels have a couple of
security issues; most of them because being tunnel as described in
RFC 6169 [RFC6169];

o tunnel injection: a malevolent person knowing a few pieces of
   information (for example the tunnel endpoints and the used
   protocol) can forge a packet which looks like a legit and valid
   encapsulated packet that will gladly be accepted by the
   destination tunnel endpoint, this is a specific case of spoofing;

o traffic interception: no confidentiality is provided by the tunnel
   protocols (without the use of IPsec or alternative encryption
   methods), therefore anybody on the tunnel path can intercept the
   traffic and have access to the clear-text IPv6 packet; combined
   with the absence of authentication, a man in the middle attack can
   also be mounted;

o service theft: as there is no authorization, even a non authorized
   user can use a tunnel relay for free (this is a specific case of
   tunnel injection);

o reflection attack: another specific use case of tunnel injection
   where the attacker injects packets with an IPv4 destination
   address not matching the IPv6 address causing the first tunnel
   endpoint to re-encapsulate the packet to the destination... Hence,
   the final IPv4 destination will not see the original IPv4 address
   but only one IPv4 address of the relay router.

o bypassing security policy: if a firewall or an IPS is on the path
   of the tunnel, then it will probably neither inspect nor detect an
   malevolent IPv6 traffic contained in the tunnel.

To mitigate the bypassing of security policies, it is recommended to
block all default configuration tunnels by denying all IPv4 traffic
matching:

o IP protocol 41: this will block ISATAP (Section 2.7.2.2), 6to4
   (Section 2.7.2.7), 6rd (Section 2.7.2.3) as well as 6in4
   (Section 2.7.2.1) tunnels;

o IP protocol 47: this will block GRE (Section 2.7.2.1) tunnels;

o UDP protocol 3544: this will block the default encapsulation of
   Teredo (Section 2.7.2.8) tunnels. Teredo is now mostly never used
   and it is no more automated in most environment, so, it is less of
   a threat, however, special consideration must be taken in case of
   devices with older or non-updated operating systems may be
   present, which by default were running Teredo.

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

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

2.7.2.1. Site-to-Site Static Tunnels

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

2.7.2.2. ISATAP

ISATAP tunnels [RFC5214] are mainly used within a single
administrative domain and to connect a single IPv6 host to the IPv6
network. This often implies that those systems are usually managed
by a single entity; therefore, audit trail and strict anti-spoofing
are usually possible and this raises the overall security.

Special care must be taken to avoid looping attack by implementing
the measures of RFC 6324 [RFC6324] and of [RFC6964].

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 than 6to4 tunnels and have little security issues except
lack of confidentiality. The security considerations (Section 12) of
[RFC5969] describes how to secure the 6rd tunnels.

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

2.7.2.4. 6PE, 6VPE, and LDPv6

Organizations using MPLS in their core can also use 6PE [RFC4798] and
6VPE [RFC4659] to enable IPv6 access over MPLS. As 6PE and 6VPE are
really similar to BGP/MPLS IP VPN described in [RFC4364], the
security of these networks is also similar to the one described in
[RFC4381]. It relies on:

o Address space, routing and traffic separation with the help of
   VRFs (only applicable to 6VPE);

o Hiding the IPv4 core, hence removing all attacks against
   P-routers;

o Securing the routing protocol between CE and PE; in the case of
   6PE and 6VPE, link-local addresses (see [RFC7404]) can be used and
   as these addresses cannot be reached from outside of the link, the
   security of 6PE and 6VPE is even higher than the IPv4 BGP/MPLS IP
   VPN.

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

2.7.2.5. DS-Lite

DS-lite is more a translation mechanism and is therefore analyzed
further (Section 2.7.3.3) in this document.

2.7.2.6. Mapping of Address and Port

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

2.7.2.7. 6to4

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

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

2.7.2.8. Teredo

Teredo tunnels [RFC4380] are mainly used in a residential environment
because that can easily traverse an IPv4 NAT-PT device thanks to its
UDP encapsulation and they 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 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
accept the return UDP traffic, then Teredo actually punches a hole in
this firewall for all IPv6 traffic to the Internet and from the
Internet. While host policies can be deployed to block Teredo in an
IPv4-only network in order to avoid this firewall bypass, it would be
more efficient to block all UDP outbound traffic at the IPv4 firewall
if deemed possible (of course, at least port 53 should be left open
for DNS traffic).

Teredo is now mostly never used and it is no more automated in most
environment, so, it is less of a threat, however, special
consideration must be taken in case of devices with older or non-
updated operating systems may be present, which by default were
running Teredo.

2.7.3. Translation Mechanisms

Translation mechanisms between IPv4 and IPv6 networks are alternative
coexistence strategies while networks transition to IPv6. While a
framework is described in [RFC6144] the specific security
considerations are documented in 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.

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 prolong the use of IPv4 in a large service
provider network until the provider can deploy and effective IPv6
solution. [RFC6598] requested a specific IANA allocated /10 IPv4
address block to be used as address space shared by all access
networks using CGN. This has been allocated as 100.64.0.0/10.

Section 13 of [RFC6269] lists some specific security-related issues
caused by large scale address sharing. The Security Considerations
section of [RFC6598] also lists some specific mitigation techniques
for potential misuse of shared address space. Some Law Enforcement
Agencies have identified CGN as impeding their cyber-crime
investigations (for example Europol press release on CGN
[europol-cgn]). Many translation techniques (NAT64, DS-lite, ...)
have the same security issues as CGN when one part of the connection
is IPv4-only.

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

[RFC7422] suggests the use of deterministic address mapping in order
to reduce logging requirements for CGN. The idea is to have an
algorithm mapping back and forth the internal subscriber to public
ports.

2.7.3.2. NAT64/DNS64 and 464XLAT

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

The Security Consideration sections of [RFC6146] and [RFC6147] list
the comprehensive issues. A specific issue with the use of NAT64 is
that it will interfere with most IPsec deployments unless UDP
encapsulation is used. DNS64 has an impact on DNSSEC see section 3.1
of [RFC7050].

464XLAT [RFC6877] 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
Network Address and Port Translation (NAPT).

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

Section 11 of [RFC6333] describes important security issues
associated with this technology.

2.8. **General Device Hardening**

   There are many environments which rely too much on the network
   infrastructure to disallow malicious traffic to get access to
   critical hosts. In new IPv6 deployments it has been common to see
   IPv6 traffic enabled but none of the typical access control
   mechanisms enabled for IPv6 device access. With the possibility of
   network device configuration mistakes and the growth of IPv6 in the
   overall Internet it is important to ensure that all individual
   devices are hardened against miscreant behavior.

   The following guidelines should be used to ensure appropriate
   hardening of the host, be it an individual computer or router,
   firewall, load-balancer,server, etc device.

   o Restrict access to the device to authorized individuals

   o Monitor and audit access to the device

   o Turn off any unused services on the end node

   o Understand which IPv6 addresses are being used to source traffic
      and change defaults if necessary

   o Use cryptographically protected protocols for device management if
      possible (SCP, SNMPv3, SSH, TLS, etc)

   o Use host firewall capabilities to control traffic that gets
      processed by upper layer protocols

   o Use virus scanners to detect malicious programs

3. **Enterprises Specific Security Considerations**

   Enterprises 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 ISP, i.e., the part of the network where the ISP employees are
   connected.

   Security considerations in the enterprise can be broadly categorized
   into two sections - 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
   providers network. This is commonly achieved by enforcing a security
   policy either by implementing dedicated firewalls with stateful
   packet inspection or a router with ACLs. A common default IPv4
   policy on firewalls that could easily be ported to IPv6 is to allow
   all traffic outbound while only allowing specific traffic, such as
   established sessions, inbound (see also [RFC6092]). Here are a few
   more things that could enhance the default policy:

   o Filter internal-use IPv6 addresses at the perimeter

   o Discard packets from and to bogon and reserved space, see also
      [CYMRU] and [RFC8190]

   o Accept certain ICMPv6 messages to allow proper operation of ND and
      PMTUD, see also [RFC4890] or [REY_PF] for hosts

   o Filter specific extension headers by accepting only the required
      ones (white list approach) such as ESP, AH (not forgetting the
      required transport layers: ICMP, TCP, UDP, ...) , where possible
      at the edge and possibly inside the perimeter; see also
      [I-D.ietf-opsec-ipv6-eh-filtering]

   o Filter packets having an illegal IPv6 headers chain at the
      perimeter (and possible inside as well), see Section 2.2

   o Filter unneeded services at the perimeter

   o Implement ingress and egress anti-spoofing in the forwarding and
      control planes

   o Implement appropriate rate-limiters and control-plane policers

3.2. Internal Security Considerations:

   The internal aspect deals with providing security inside the
   perimeter of the network, including the end host. 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.

   As mentioned in Section 2.6.2, care must be taken when running
   automated IPv6-in-IP4 tunnels.

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

Hosts need to be hardened directly through security policy to protect
against security threats. The host firewall default capabilities
have to be clearly understood, especially 3rd party ones which can
have different settings for IPv4 or IPv6 default permit/deny
behavior. In some cases, 3rd party firewalls have no IPv6 support
whereas the native firewall installed by default has it. General
device hardening guidelines are provided in Section 2.8

It should also be noted that many hosts still use IPv4 for transport
for things like RADIUS, TACACS+, SYSLOG, etc. This will require some
extra level of due diligence on the part of the operator.

4. Service Providers Security Considerations

4.1. BGP

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

o Authenticating the TCP session;

o TTL security (which becomes hop-limit security in IPv6) as
   [RFC5082];

o bogon AS filtering;

o Prefix filtering.

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

4.1.1. Remote Triggered Black Hole Filtering

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

4.2. Transition Mechanism

   SP will typically use transition mechanisms such as 6rd, 6PE, MAP,
   DS-Lite which have been analyzed in the transition Section 2.7.2
   section.

4.3. Lawful Intercept

   The Lawful Intercept requirements are similar for IPv6 and IPv4
   architectures and will be subject to the laws enforced in varying
   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 what the logging retention policies will be.

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

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

   A sample architecture which was written for informational purposes is
   found in [RFC3924].

5. Residential Users Security Considerations

   The IETF Homenet working group is working on how IPv6 residential
   network should be done; this obviously includes operational security
   considerations; but, this is still work in progress.

   Residential users have usually less experience and knowledge about
   security or networking. As most of the recent hosts, smartphones,
   tablets have all 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 tunnels. Several peer-
   to-peer programs (notably Bittorrent) support IPv6 and those programs
   can initiate a Teredo tunnel through the IPv4 residential gateway,
   with the consequence of making the internal host reachable from any
   IPv6 host on the Internet. It is therefore recommended that all host
   security products (personal firewall, ...) are configured with a
   dual-stack security policy.

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

o outbound only: allowing all internally initiated connections and
   block all externally initiated ones, which is a common default
   security policy enforced by IPv4 Residential Gateway doing NAT-PT
   but it also breaks the end-to-end reachability promise of IPv6.
   [RFC6092] lists several recommendations to design such a CPE;

o open/transparent: allowing all internally and externally initiated
   connections, therefore restoring the end-to-end nature of the
   Internet for the IPv6 traffic but having a different security
   policy for IPv6 than for IPv4.

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

There is also an alternate solution which has been deployed notably
by Swisscom: open to all outbound and inbound connections at the
exception of a handful of TCP and UDP ports known as vulnerable.

6. Further Reading

There are several documents that describe in more details 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 your
convenience:

1. Guidelines for the Secure Deployment of IPv6 [NIST]

2. North American IPv6 Task Force Technology Report - IPv6 Security
    Technology Paper [NAv6TF_Security]

3. IPv6 Security [IPv6_Security_Book]

7. Acknowledgements

The authors would like to thank the following people for their useful
comments: Mikael Abrahamsson, Fred Baker, Mustafa Suha Botsali, Brian
Carpenter, Tim Chown, Lorenzo Colitti, Markus de Bruen, Tobias
Fiebig, Fernando Gont, Jeffry Handal, Lee Howard, Panos Kampanakis,
Erik Kline, Jouni Korhonen, Mark Lentczner, Jen Linkova (and her
detailed review), Gyan S. Mishra, Jordi Palet, Bob Sleigh, Donal
Smith, Tarko Tikan, Ole Troan, Bernie Volz (by alphabetical order).

8. IANA Considerations

   This memo includes no request to IANA.

9. Security Considerations

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

10. References

10.1. Normative References

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

   [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>.

10.2. Informative References

   [CYMRU]   "Packet Filter and Route Filter Recommendation for IPv6 at
              xSP routers", <http://www.team-
              cymru.org/ReadingRoom/Templates/IPv6Routers/xsp-
              recommendations.html>.

   [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>.

   [I-D.chakrabarti-nordmark-6man-efficient-nd]
             Chakrabarti, S., Nordmark, E., Thubert, P., and M.
             Wasserman, "IPv6 Neighbor Discovery Optimizations for
             Wired and Wireless Networks", draft-chakrabarti-nordmark-
             6man-efficient-nd-07 (work in progress), February 2015.

   [I-D.ietf-opsec-ipv6-eh-filtering]
             Gont, F. and W. LIU, "Recommendations on the Filtering of
             IPv6 Packets Containing IPv6 Extension Headers", draft-
             ietf-opsec-ipv6-eh-filtering-06 (work in progress), July
             2018.

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

   [I-D.thubert-savi-ra-throttler]
             Thubert, P., "Throttling RAs on constrained interfaces",
             draft-thubert-savi-ra-throttler-01 (work in progress),
             June 2012.

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

   [IPv6_Security_Book]
             Hogg, S. and E. Vyncke, "IPv6 Security",
             ISBN 1-58705-594-5, Publisher CiscoPress, December 2008.

   [NAv6TF_Security]
             Kaeo, M., Green, D., Bound, J., and Y. Pouffary, "North
             American IPv6 Task Force Technology Report - IPv6 Security
             Technology Paper", 2006,
             <http://www.ipv6forum.com/dl/white/
             NAv6TF_Security_Report.pdf>.

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

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

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

   [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G.,
              and E. Lear, "Address Allocation for Private Internets",
              BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996,
              <https://www.rfc-editor.org/info/rfc1918>.

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

   [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6
              (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460,
              December 1998, <https://www.rfc-editor.org/info/rfc2460>.

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

   [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>.

   [RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6
              Neighbor Discovery (ND) Trust Models and Threats",
              RFC 3756, DOI 10.17487/RFC3756, May 2004,
              <https://www.rfc-editor.org/info/rfc3756>.

   [RFC3924] Baker, F., Foster, B., and C. Sharp, "Cisco Architecture
              for Lawful Intercept in IP Networks", RFC 3924,
              DOI 10.17487/RFC3924, October 2004,
              <https://www.rfc-editor.org/info/rfc3924>.

   [RFC3964] Savola, P. and C. Patel, "Security Considerations for
              6to4", RFC 3964, DOI 10.17487/RFC3964, December 2004,
              <https://www.rfc-editor.org/info/rfc3964>.

   [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander,
              "SEcure Neighbor Discovery (SEND)", RFC 3971,
              DOI 10.17487/RFC3971, March 2005,
              <https://www.rfc-editor.org/info/rfc3971>.

   [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)",
              RFC 3972, DOI 10.17487/RFC3972, March 2005,
              <https://www.rfc-editor.org/info/rfc3972>.

   [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>.

   [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>.

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

   [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>.

[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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken,
              "Specification of the IP Flow Information Export (IPFIX)
              Protocol for the Exchange of Flow Information", STD 77,
              RFC 7011, DOI 10.17487/RFC7011, September 2013,
              <https://www.rfc-editor.org/info/rfc7011>.

   [RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model
              for IP Flow Information Export (IPFIX)", RFC 7012,
              DOI 10.17487/RFC7012, September 2013,
              <https://www.rfc-editor.org/info/rfc7012>.

   [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed.,
              "Source Address Validation Improvement (SAVI) Framework",
              RFC 7039, DOI 10.17487/RFC7039, October 2013,
              <https://www.rfc-editor.org/info/rfc7039>.

   [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing
              of IPv6 Extension Headers", RFC 7045,
              DOI 10.17487/RFC7045, December 2013,
              <https://www.rfc-editor.org/info/rfc7045>.

   [RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of
              the IPv6 Prefix Used for IPv6 Address Synthesis",
              RFC 7050, DOI 10.17487/RFC7050, November 2013,
              <https://www.rfc-editor.org/info/rfc7050>.

   [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic
              Requirements for IPv6 Customer Edge Routers", RFC 7084,
              DOI 10.17487/RFC7084, November 2013,
              <https://www.rfc-editor.org/info/rfc7084>.

   [RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of
              Oversized IPv6 Header Chains", RFC 7112,
              DOI 10.17487/RFC7112, January 2014,
              <https://www.rfc-editor.org/info/rfc7112>.

   [RFC7113] Gont, F., "Implementation Advice for IPv6 Router
              Advertisement Guard (RA-Guard)", RFC 7113,
              DOI 10.17487/RFC7113, February 2014,
              <https://www.rfc-editor.org/info/rfc7113>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [SCANNING]
              Barnes, R., Altmann, R., and D. Kerr, "Mapping the Great
              Void - Smarter scanning for IPv6", February 2012,
              <http://www.caida.org/workshops/isma/1202/slides/
              aims1202_rbarnes.pdf>.

   [WEBER_VPN]
              Weber, J., "Dynamic IPv6 Prefix - Problems and VPNs",
              March 2018, <https://blog.webernetz.net/wp-
              content/uploads/2018/03/TR18-Johannes-Weber-Dynamic-IPv6-
              Prefix-Problems-and-VPNs.pdf>.

Authors' Addresses

   Eric Vyncke (editor)
   Cisco
   De Kleetlaan 6a
   Diegem 1831
   Belgium

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

   Kiran K. Chittimaneni
   WeWork

   Email: kk.chittimaneni@gmail.com

   Merike Kaeo
   Double Shot Security
   3518 Fremont Ave N 363
   Seattle 98103
   USA

   Phone: +12066696394
   Email: merike@doubleshotsecurity.com

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

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