Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: January 1, 2017

M. Georgescu, Ed.
NAIST

June 30, 2016

The STRIDE towards IPv6: A Threat Model for IPv6 Transition Technologies

draft-georgescu-opsec-ipv6-trans-tech-threat-model-01

Abstract

This document provides a structured approach for analyzing the threats associated with the various IPv6 transition technologies specified by the IETF. The threat model is built around the established STRIDE threat classification and is aimed at existing IPv6 transition technologies, as well as their future developments.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 1, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

When building an IPv6 transition plan, security is arguably one of the biggest concerns for network operators, as a heterogeneous IPv4 and IPv6 environment greatly increases the attack surface. To that end, building a threat model for IPv6 transition technologies can help clarify and categorize the associated security threats. In turn, this should facilitate the search for mitigation solutions.

The security considerations of IPv6 transition technologies has generally been analyzed in each of the corresponding specifications, and some documents have discussed the general threats associated with transition technologies (see e.g. [RFC4942]).

However, more structured threat modeling has proved useful for understanding the security of intricate systems. Structured approaches allows one to discover, categorize and classify the threats according to their potential impact on the system. Considering the complicated nature of IPv6 transition technologies, threat modeling makes a good candidate for better understanding their security implications. This document follows a structured approach for analyzing the threats asociated with transition technologies, that considers the functions of a transition technology as well as the cntext in which the technology is used.

The threat model uses the established STRIDE mnemonic and threat classification. STRIDE stands for Spoofing, Tampering, Repudiation, Information Disclosure, Denial of service and Elevation of Privilege, a generic list of threats which can be used to classify various threats and provides some basic mitigation directions. Since similar transition technologies can be associated with a similar list of threats, the document considers the generic classification of IPv6 transition technologies described in [draft-bmwg-v6trans].

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

3. The Generic Categories of IPv6 Transition Technologies

Table 1 presents the generic categories described in [draft-bmwg-v6trans] and some sample IPv6 transition technologies specified by the IETF.

Table 1. IPv6 Transition Technologies Categories

	Generic category	IPv6 Transition Technology
1	Dual-stack	Dual IP Layer Operations [RFC4213]
2	Single translation	NAT64 [RFC6146], IVI [RFC6219]
3	Double translation	464XLAT [RFC6877], MAP-T [RFC7599]
4	Encapsulation	DSLite [RFC6333], MAP-E [RFC7597] Lightweight 4over6 [RFC7596] 6RD [RFC5569]

4. Building The Threat Model

To build a threat model for IPv6 transition technologies a series of steps is recommended. The steps were inspired by the threat modelling approach described in [stride-shostack]. These steps are detailed in the following subsections.

4.1. Establish the function

The function of the IPv6 transition technology needs to be clearly documented. Depending on the context, the technology can incorporate multiple services, which need to be clearly identified in order to perform an effective threat analysis.

4.2. Identify the generic category

The category should be identified considering the generic classification defined in Section 3. This step can help reuse the threat analysis data for technologies which are part of the same category.

4.3. Decompose the technology

Build a data flow diagram (DFD) and highlight the entry points, protected resources and trust boundaries. The entry points should be assigned a level of trust considering the trust boundaries.

The external entities, process, data store and data flow elements should be depicted in the same diagram. The IP protocol suite and the protocols used for the designated function should be identified as well. This can narrow down the attack surface.

Figure 1 presents the basic elements of a data flow diagram as well as general rules for their association with network elements.

```
 +-----------+
 | External  |  Represents a network node which is outside
 | Entity    |  the control of a network provider
 +-----------+
      ___
   ,'    `.     Represents a middle-box or a network node
  |  Pro   |    which processes translated or encapsulated
   \ cess/      traffic
    `---'

 ============
  Data store    Represents a node where user and provider
 ============   data is stored

  Data Flow
------------>   Data in transit exchanged between network elements

    Trust
     ()        The border which marks the part of the
     ()        network considered outside the control
     ()        of a network provider
  boundary
```

Figure 1. DFD Elements

4.4. Identify the threats

4.4.1. STRIDE-DFD Association

The STRIDE model associates the six categories of threats to each of the elements described in the DFD. Based on this association, we get an initial assessment of the threats as shown in Table 2. To clarify, a data flow, for example, is susceptible to tampering, information disclosure and denial of service threats. The initial threat assessment must be followed by a detailed analysis which should consider the protocols used in conjuncture with the transition technology.

Table2. DFD-STRIDE Associations

S	T	R	I	D	E
#		#			
O	O	O	O	O	O
	=	=	=	=	
	>		>	>	
#	External entity				
O	Process				
=	Data store				
>	Data flow				

4.4.2. Level of Trust

We associate a level of trust with each entry point. Entry points that are trusted are assumed to behave as expected. That is, if the entry point is considered trusted, we can assume the likelihood of an attack is low. Furthermore, the six categories of STRIDE attacks could be assigned a likelihood by considering their association with the DFD elements that are entry points.

For instance, let's suppose we have an untrusted entry point (High likelihood of exploitation) which is also an external entity. Spoofing and repudiation are potential threats for an external entity. By association, the two types of attacks can be considered to have a high likelihood of being exploited. Using this logic, we can assign a likelihood value to each found threat. This can represent a base for prioritizing mitigation solutions. The likelihood levels can be defined in accordance with the levels of trust assigned to the the entry points.

4.4.3. Documenting the Threats

Each discovered threat should be documented using the format presented in Table 3.

Table2. Threat Info Format

Field Name	Description
Threat-ID	A code associated with each identified threat
Description	A summarized description of the threat
STRIDE	The association with the STRIDE categories
Mitigation	Details about possible mitigation solutions
Likelihood	Likelihood of the threat being exploited
Validation	Empirical validation data

The Threat-ID is supposed to be an easy way to refer and identify the threat within the IETF. The tentative format is IETF-TDB-[associated protocol/technology]-[serial number]. IETF-TDB stands for IETF Threat Database in the hope that in the future a threat database will be maintained within the IETF. The serial number is incremented with each threat found for a particular protocol or technology.

4.4.4. Complex Threats

As the subcomponents and subprotocols interact, the threats can fuse and result in convoluted threats with a higher likelihood of exploitation. Depending on the list of discovered threats, the possibility of a fusion between threats should be analyzed.

4.5. Review, Repeat and Validate

Steps 1 and 3 have to be reviewed in the context of potential changes in the technology function and associated protocols. Step 4 should be repeated periodically, as threats may have been overlooked, or the context set by steps 1 and 3 may have changed. If the transition technologies have existing implementations, the analysis should be confirmed with empirical data.

The next sections applied the proposed threat modeling approach to the IPv6 transition technologies identified in Section 3.

5. Dual Stack Threat Model

5.1. Establish the Function

The function for dual-stack transition technologies is to ensure a safe data exchange over a dual-stack infrastructure. In other words, the data can be transfered over both IPv4 and IPv6. From a network service perspective, the main function is data forwarding. This includes interior gateway routing solutions. We start with the assumption that services such as address provision, DNS resolution or exterior gateway routing are performed by other nodes within the core network. This assumption in common for all the four generic categories of IPv6 transition technologies.

5.2. Identify the Generic Category

Since we are targeting the generic category itself, the step is unnecessary here. This stands for the other three categories as well.

5.3. Decompose the Technology

A DFD for dual-stack transition technologies is presented in Figure 2. The diagram represents a basic use case and includes a minimal set of elements.

```
  Domain A-DS                Core  ___   Domain       Domain B-DS
      +----------+      (         ,'   `.          )    ============
      | Customer |-----(------->|  DS   |------- )---> Provider
      | Device   |<----(-------- \ node/<--------)---- Data Store
      +----------+      (          `---' IPv4/IPv6)    ============
         E  P                     E  P                     E  P
_____________________________________________________________________
Legend              ___                               Trust
 +----------+  ,'   `.   ============    Data Flow  ()  E=Entry
 | External | |  Pro  |   Data store   ------------> ()    point
 | Entity   |  \ cess/   ============                ()  P=Protected
 +----------+   `---'                                     boundary
```

Figure 2. Data Flow Diagram (DFD) for Dual Stack (DS) technologies

In Domain A, which is assumed to be on the customer side we have a Customer Device which initiates the data exchange. It represents one of the entry points of the system and contains important data, which should be regarded as an asset and protected. The Customer Device is regarded as an external element because it is outside the control zone of the assumed network provider. The data request is transmitted over IPv4 or IPv6 to a Dual-stack node.

The Dual-stack node is another entry point and contains valuable topology information which should to protected as well. The Dual-stack node forwards in turn the data request to the provider data store. The Data store is the last entry point in the system and it is assumed to contain valuable data. The data reply is forwarded back to the customer device.

The only trusted entry point in the system is the Dual-stack node. The other two entry points are considered untrusted, since they are outside the control of the production network.That means they can be exploited with a higher likelihood by an attacker.

Considering the data can be transferred over both IPv4 and IPv6, we need to consider both IP protocol suites. Furthermore, the possibility of using security and routing protocols should be considered.

5.4. Identify the threats

5.4.1. STRIDE-DFD Assoctiation

By analyzing the DFD in association with the STRIDE threats per element chart, we can make the associations depicted in Table 3.

Table3. DFD-STRIDE Associations DS

S	T	R	I	D	E
#-H		#-H			
O-L	O-L	O-L	O-L	O-L	O-L
	=-H	=-H	=-H	=-H	
	>-H		>-H	>-H	
#	Customer device				
O	DS node				
=	Provider data store				
>	Data flow				

5.4.2. From Trust to Likelihood

Looking at the associations in Table 3, The Customer Device can be subject to spoofing and repudiation attacks. It being an untrusted entry point, that means there is a high likelihood of an attack. This is marked in Table 3 with H.

The Dual-stack node can be subject to all six types of attacks. However, the likelihood of that happening is low, considering it is a trusted entry point.

The Data flow is vulnerable to tampering, information disclosure and denial of service. Considering it traverses untrusted parts of the system, the level of likelihood of an attack on the data flow is high.

Lastly, the Data store could potentially be targeted by tampering, repudiation, information disclosure and denial of service attacks. The likelihood for these to happen is high as well, the data store being an untrusted entry point.

5.4.3. Documenting the Threats

The Tables 5-10 of the Appendix contain a non-exhaustive collection of existing threats, which have been collected by surveying a part of existing literature on this subject. For further documentation, each threat has been provided with a reference in the first column. For reuse purposes, the threats are organized according to the categories of protocols which would be necessary for accomplishing the function of the IPv6 transition technologies.

For dual-stack transition technologies the protocol threats associated with the IPv4 suite (Table 6), IPv6 suite(Table 7), routing (Table 10) and switching (Table 5) could potentially be exploited from the 3 entries of the system: the untrusted (High likelihood of exploitation) Customer device, the trusted (Low likelihood of exploitation) Dual-stack node (Process) and untrusted (High likelihood of exploitation) Provider Data store.

The IPv4 suite, transport layer and most of the IPv6 suite protocols are associated with all the elements of the DFD. By extrapolation, their threats have a high likelihood of occurrence. Some of the IPv6 protocol threats (Table 7), namely IETF-TDB-ND-3 to IETF-TDB-ND-6 and the Layer 2 technologies' threats (Table 5) can only be associated with routers or switches. In the context of the DFD, they could only be associated with the Dual-stack node. That means they have a low likelihood of occurrence. Similarly, the routing protocols

(Table 10) can only be associated with the Dual-stack node. By association, they also have a low likelihood of being exploited.

5.4.4. Complex Threats

By analyzing the interaction between the three elements of the DFD and the protocols used by Dual stack transition technologies, we can uncover other threats. For example, if the IETF-TDB-ARP-1(ARP cache poisoning) is used to perform a Denial of Service attack on the Dual-stack node from the Customer device, the likelihood of exploitation rises for the IETF-TDB-ND-10 (ND Replay Attacks) threats. IETF-TDB-ARP-1 could be replaced by any other DoS threat associated with the IPv4 protocol suite. This complex threat could be prevented by ensuring that the IPv4 suite DoS threats are properly mitigated. Examples of convoluted threats for the four generic IPv6 transition technologies are presented in Table 4.

Table4. Complex Threats

	ThreatID	Description	S	T	R	I	D	E	Mitigation
1 V	IETF-TDB -DS-1	IETF-TDB -ARP-1 + IETF-TDB -ND-4	H	H	H	H	H		DoS Mitigation for IPv4 suite
2 V	IETF-TDB -DS-2	IETF-TDB -ARP-1 + IETF-TDB -OSPFv3-1	H	H	H	H	H	H	Crypto authen
3 X	IETF-TDB -1transl-1	IETF-TDB IP/ICMP-3 + IETF-TDB -ICMPv6-1	H		H	H	H		No widely accepted mitigation
4 V	IETF-TDB -1transl-2	IETF-TDB -TCP-1 + IETF-TDB -ND-4	H	H	H	H	H		Block non-internal traffic
5 X	IETF-TDB	IETF-TDB -IP/ICMP-4	L	L	L	L	L		No widely accepted

	-2transl-1	+ IETF-TDB -ND-4							mitigation
6 V	IETF-TDB -2transl_2	IETF-TDB -IP/ICMP-1 + IETF-TDB -OSPFv3-1	L	L	L	L	L	L	reverse path checks
7	IETF-TDB -encaps-1	IETF-TDB -IPv6-1 + IETF-TDB -4encaps_1				H	H		IPv4 firewall before decaps
Legend									
H	associaced with High likelihood				L	associaced with Low likelihood			

Another convoluted threat can result from exploiting IPv4 or IPv6 spoofing threats to increase the likelihood of an attack on routing protocols with simple authentication, such as or IETF-TDB-OSPFv3-1, IETF-TDB-OSPFv2-1 or IETF-TDB-RIPv2-1. Since the attack could be performed from an untrusted entry point (Customer device or Data store), the likelihood of the threat being exploited rises to High. This type of attack can be mitigated by using cryptographic authentication for the routing protocols.

The list of threats can help technology implementors and network operators alike prioritize the threats and mitigate accordingly.

5.5. Review, Repeat and Validate

This step is necessary if the technology analyzed or associated protocols change. For example if the routing system were to be only OSPFv3, then the threats associated with other routing protocols could be ignored. Also, the detailed analysis of threats is far from exhaustive. In terms of convoluted new threats, only a few are presented as an example. If this was to be an updated database of threats, it would need constant update.

To further validate the presented threats, a simple penetration testbed was built. The details of the testbed are presented in Figure 3. MAP-T [RFC7599] was used as transition technology. Asamap [asamap2014], a transition implementation developed in Japan, was used as the base for MAP-T. The threats which were successfully emulated, have been marked accordingly in the first column of Table 4. In the case of the convoluted threats identified for Dual-stack transition technologies, both threats were emulated successfully by performing ARP Cache Spoofing, Neighbor Advertisement (NA) flooding and simple traffic analysis.

```
                                +----------+
      +---------+---------------+ Kali     |
      |---------|+--------------> Attacker |
      ||        ||              +---+^-----+
      ||        ||                  ||
      ||        ||           ___    ||     __
   +--v-------+ ||  (      ,'   `.  ||   ,'   `.       )      ===========
   | Win8     +-+|--(---->+  amap +-+|->+  amap +-----)----->  Ubuntu
   | Host     <--+--(-----+\ CE  /<--+--+\ BR  /<-----)-----+  Server
   +----------+     (       `+-+'  IPvY   `+-+'  IPvX )      ===========
      E  P                  E  P          E  P                    E  P
```

Figure 3. Pentestbed Setup

6. Single Translation Threat Model

To avoid redundant information, the following three subsections will only mark the differences with the threat modeling process presented for Dual-stack transition technologies.

One of the fundamental differences is that the single translation technologies would require a node to algorithmically translate the IPvX packets to IPvY, as shown in Figure 4.

6.1. Decompose the Technology

A DFD for single translation transition technologies is presented in Figure 4. The diagram represents a basic use case and includes a minimal set of elements.

```
     Domain A-IPvX                Core___   Domain       Domain B-IPvY
         +----------+ IPvX (        ,'   `.  IPvY   )    ============
         | Customer |------(------>|  Tra  |------- )--->  Provider
         | Device   |<-----(------- \ nsl /<--------)---- Data Store
         +----------+ IPvX (         `---'   IPvY   )    ============
            E  P                     E  P                    E  P
   _____________________________________________________________________
   Legend           ___                                Trust
    +----------+  ,'   `.   ============    Data Flow   ()  E=Entry
    | External | |  Pro  |   Data store   ------------> ()    point
    | Entity   |  \ cess/   ============                ()  P=Protected
    +----------+   `---'                                       boundary
```

Figure 4. DFD for 1transl technologies

6.2. Identify the threats

For both translation directions 4->6 and 6->4, the threats for the IPv4 suite (Table 6), IPv6 suite (Table 7), routing (Table 10) and switching (Table 5) should be considered. There are technologies that use stateful mapping algorithms e.g. Stateful NAT64 [RFC6146], which create dynamic correlations between IP addresses or {IP address, transport protocol, transport port number} tuples. Consequently, we need to consider the protocols used at the transport layer (Table 9) as part of the attack surface. The threats presented in Table X, associated with the IP/ICMP translation algorithm (IP/ICMP) should be considered as well.

In terms of convoluted threats, one example could be exploiting the IETF-TDB-IP/ICMP-3 threat (IPAuth does not work with IP/ICMP) which would increase the likelihood of IETF-TDB-ND-4 (Default router is killed) or IETF-TDB-ND-5 (Good router goes bad) threats being exploited. Since there is no widely-accepted mitigation for any of the three threats, this convoluted threat is laking a mitigation solution as well. Fortunately, both complex threats could not be validated empirically. An IPsec VPN connection was successfully established using UDP encapsulation between the Windows Host and the Ubuntu Server. Moreover, the IETF-TDB-ND-4 and IETF-TDB-ND-5 could not be validated empirically, as Asamap [asamap2014] does not accept RA messages when IPv6 forwarding is enabled.

If the IETF-TDB-TCP-1 threat (SYN flood) is exploited from an untrusted entry point, it increases the likelihood of a IETF-TDB-ND-10 (ND Replay attacks) threat. This threat can be mitigated by blocking packets with non-internal addresses from leaving the network. Both the SYN flood attack and the Neighbor Advertisement (NA) flooding attacks were staged successfully.

7. Double Translation Threat Model

The main difference between the Single translation case and the double translation case is the need for an extra translation device as part of the core network (Figure 5). Another important difference would be that in the untrusted zone, the Customer device and Data store would employ the same IP suite.

7.1. Decompose the Technology

A DFD for double translation transition technologies is presented in Figure 5. The diagram represents a basic use case and includes a minimal set of elements.

```
   Domain A-IPvX          ___  Core Dom ___           Domain B-IPvX
  +----------+ IPvX(    ,'   `. IPvY  ,'   `.     )      ============
  | Customer |-----(---|  Tra  |---->|  Tra  |--- )-----> Provider
  | Device   |<----(----\ nsl /<------\ nsl /<----)------ Data Store
  +----------+     (     `---'  IPvY   `---'  IPvX)      ============
     E  P               E  P          E  P                  E  P
--------------------------------------------------------------------
  Legend              ___                               Trust
  +----------+  ,'   `.   ============    Data Flow   ()  E=Entry
  | External | |  Pro  |   Data store   ------------> ()    point
  | Entity   |  \ cess/   ============                 ()  P=Protected
  +----------+   `---'                                      boundary
```

Figure 5. DFD for 2transl technologies

7.2. Identify the threats

The considered threats for the untrusted elements would be either the IPv4 suite (Table 6) or the IPv6 suite (Table 7) protocol threats. Similar to the single translation technologies, the routing (Table 10), switching (Table 5), transport layer (Table 9) and IP/ ICMP (Table 8) threats should be analyzed as well.

The use of stateful translation mechanisms can expose a double translation technology to the IETF-TDB-IP/ICMP-4 threat (DoS by exhaustion of resources). A convoluted threat can result by exploiting this threat on one of the translators and the IETF-TDB-ND-4 or IETF-TDB-ND-5 threats on the other translator. This threat would have a higher likelihood of exploitation since it is associated with an untrusted entry point. In terms of mitigation, further investigation is needed, as there are no widely accepted mitigation techniques. Although the IETF-TDB-IP/ICMP-4 threat was replicated with success, the IETF-TDB-ND-10 or IETF-TDB-ND-5 could not be emulated because of a simple built-in mitigation mechanism implemented by Asamap [asamap2014]. Router advertisement (RA) messages are not accepted while in IPv6 forwarding mode.

The IETF-TDB-IP/ICMP-4 threat can also fuse with the simple authentication threats such as IETF-TDB-OSPFv3-1 , IETF-TDB-OSPFv2-1 or IETF-TDB-RIPv2-1 to affect both translating nodes. The likelihood of the threats become higher by fusing them, since the flooding attack can be performed from an untrusted entry point, the customer network. This threat could be mitigated by using cryptographic authentication or implementing reverse path checks. The convoluted threat was validated by flooding the translation table of the first translator and forcing it to crash. OSPFv3 information disclosure was emulated with simple traffic analysis. To validate the other types of threats, a rogue router instance was created using Asamap [asamap2014].

8. Encapsulation Threat Model

Similar to double translation IPv6 transition technologies, encapsulation technologies, the core network traffic is forwarded through at least two devices, an Encapsulator and a Decapsulator (Figure 6). As the main difference, the traffic is encapsulated. This means more overhead but also more support for end-to-end security protocols. Packets are encapsulated either over IPv4 or IPv6.

8.1. Decompose the Technology

A DFD for encapsulation transition technologies is presented in Figure 6. The diagram represents a basic use case and includes a minimal set of elements.

```
  Domain A-IPvX          ___ Core Dom  ___            Domain B-IPvX
 +----------+ IPvX(    ,'   `. IPvY  ,'   `.      )      ============
 | Customer |-----(---|  Enc  |---->|  Dec  |---- )----->  Provider
 | Device   |<----(----\ aps /<------\ aps /<-----)------ Data Store
 +----------+     (     `---'  IPvY   `---'  IPvX )      ============
    E  P              E  P          E  P                    E  P
________________________________________________________________________
Legend           ___                                 Trust
 +----------+  ,'   `.    ============    Data Flow   {}  E=Entry
 | External | |  Pro  |    Data store   ------------> {}    point
 | Entity   |  \ cess/    ============                {}  P=Protected
 +----------+   `---'                                       boundary
```

Figure 6. DFD for encaps technologies

8.2. Identify the threats

For the untrusted domain devices we would consider either the IPv4 suite (Table 6) or the IPv6 suite (Table 7) threats. In addition the routing (Table 10), switching (Table 5), transport layer (Table 9) and encapsulation-related (Table 8) threats should be considered.

Convoluted threats can arise by exploiting the IETF-TDB-4encaps-1 threat (avoiding IPv4 network security measures with encapsulation). This threat can facilitate IPv6 suite DoS threats on the Decapsulator device. This convoluted threat would increase the likelihood of a successful DoS attack from the Customer Device. The threat could be mitigated by making use of an IPv4 firewall before decapsulating the packets.

9. Acknowledgments

The author would like to thank Fernando Gont for his review and useful suggestions.

This document was derrived from a template contributed by the xml2rfc project.

10. IANA Considerations

This memo includes no request to IANA.

All drafts are required to have an IANA considerations section (see Guidelines for Writing an IANA Considerations Section in RFCs [RFC5226] for a guide). If the draft does not require IANA to do anything, the section contains an explicit statement that this is the case (as above). If there are no requirements for IANA, the section will be removed during conversion into an RFC by the RFC Editor.

11. Security Considerations

This memo attempts to build a threat model for IPv6 transition technologies. The author would like to encourage the use of a similar threat modeling approach when writing the security considerations of standards developed in the IETF. To be more concrete the following steps could be reused:

R1 Identify the function

R2 Associate the technology with a generic category (if any)

R3 Decompose the technology

R4 Identify the threats

R5 Review, repeat and validate

12. References

12.1. Normative References

[draft-bmwg-v6trans]
Georgescu, M. and G. Lencse, "Benchmarking Methodology for IPv6 Transition Technologies", 2015, <https://tools.ietf.org/html/draft-ietf-bmwg-ipv6-tran-tech-benchmarking-01>.

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

[RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations", RFC 4942, DOI 10.17487/RFC4942, September 2007, <http://www.rfc-editor.org/info/rfc4942>.

12.2. Informative References

[arps] Abad, C. and R. Bonilla, "An analysis on the schemes for detecting and preventing ARP cache poisoning attacks", 2007.

[asamap2014]
Asama, M., "MAP supported Vyatta", 2014, <http://enog.jp/˜masakazu/vyatta/map/>.

[bellovin89]
Bellovin, S., "Security Problems in the TCP/IP Protocol Suite", 1989.

[harris99]
Harris, B. and R. Hunt, "TCP/IP security threats and attack methods", 1999.

[icmps] Low, C., "ICMP attacks illustrated", 2001.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <http://www.rfc-editor.org/info/rfc2328>.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, DOI 10.17487/RFC2629, June 1999, <http://www.rfc-editor.org/info/rfc2629>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <http://www.rfc-editor.org/info/rfc3552>.

[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, <http://www.rfc-editor.org/info/rfc3756>.

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

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, DOI 10.17487/RFC4213, October 2005, <http://www.rfc-editor.org/info/rfc4213>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, DOI 10.17487/RFC4443, March 2006, <http://www.rfc-editor.org/info/rfc4443>.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <http://www.rfc-editor.org/info/rfc4552>.

[RFC4822] Atkinson, R. and M. Fanto, "RIPv2 Cryptographic Authentication", RFC 4822, DOI 10.17487/RFC4822, February 2007, <http://www.rfc-editor.org/info/rfc4822>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC5569] Despres, R., "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)", RFC 5569, DOI 10.17487/RFC5569, January 2010, <http://www.rfc-editor.org/info/rfc5569>.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <http://www.rfc-editor.org/info/rfc6052>.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, DOI 10.17487/RFC6145, April 2011, <http://www.rfc-editor.org/info/rfc6145>.

[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, <http://www.rfc-editor.org/info/rfc6146>.

[RFC6219] Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The China Education and Research Network (CERNET) IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition", RFC 6219, DOI 10.17487/RFC6219, May 2011, <http://www.rfc-editor.org/info/rfc6219>.

[RFC6274] Gont, F., "Security Assessment of the Internet Protocol Version 4", RFC 6274, DOI 10.17487/RFC6274, July 2011, <http://www.rfc-editor.org/info/rfc6274>.

[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, <http://www.rfc-editor.org/info/rfc6333>.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <http://www.rfc-editor.org/info/rfc6877>.

[RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <http://www.rfc-editor.org/info/rfc7596>.

[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, <http://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, <http://www.rfc-editor.org/info/rfc7599>.

[stride-shostack] Shostack, A., "Experiences threat modeling at microsoft", 2008, <http://mail.homeport.org/˜adam/modsec08/Shostack-ModSec08-Experiences-Threat-Modeling-At-Microsoft.pdf>.

[sws] Rouiller, S., "Virtual LAN Security: weaknesses and countermeasures", 2003.

[udps] Garg, A. and A. Reddy, "Mitigation of DoS attacks through QoS regulation", 2004.

[x1037] ITU-T, "IPv6 technical security guidelines. Recommendation X.1037", 2013, <https://www.itu.int/rec/T-REC-X.1037/en>.

Appendix A. Appendix A

Table5. L2 Technologies Threats

	ThreatID	Description	S	T	R	I	D	E	Mitigation
1	IETF-TDB -VLAN-1 [x1037]	Exhaust a the FIB of an L2switch					L		IEEE 802.1x authentication
2	IETF-TDB -VLAN-2 [sws]	CAM Overflow					L		port -security features
3	IETF-TDB -VLAN-3 [sws]	Basic VLAN Hopping	L						Software update
4	IETF-TDB -VLAN-4 [sws]	Double encapsulation VLAN Hopping	L					L	Disable Auto -trunking
5	IETF-TDB -VLAN-5 [sws]	Spanning Tree Attack				L	L		Disable STP; BPDU Guard

Legend			
H	associaced with High likelihood	L	associaced with Low likelihood

Table6. IPv4 Protocol Suite Threats

	ThreatID	Description	S	T	R	I	D	E	Mitigation
1	IETF-TDB -IPv4-1 [harris99]	IP source address spoofing	H	H	H	H			Apply ACLs filter source address traffic
2	IETF-TDB	Mal		H					Version

	-IPv4-2 [RFC6274]	formed version field							checked to be 4
3	IETF-TDB -IPv4-3 [RFC6274]	forged DSCP field	H				H		Filter unrecogn ized DSCP
4	IETF-TDB -IPv4-4 [RFC6274]	Buffer overflow IP frag mentation					H		avoid illegit imate re assembly
5	IETF-TDB -ICMP-1 [harris99]	Ping o'death					H		do not accept oversized ICMP
6	IETF-TDB -ICMP-2 [bellovin89]	ICMP redirects	H	H	H	H	H		don't update routing tables with ICMP Redirects
7	IETF-TDB -ICMP-3 [icmps]	ICMP sweep for recon				H			Selective filtering of ICMP
8	IETF-TDB -ICMP-6 [icmps]	ICMP flooding					H		Selective filtering of ICMP
9	IETF-TDB -ARP-1 [arps]	ARP cache poisoning	H	H	H	H	H		Static ARP entries, arpwatch
10	IETF-TDB -ARP-2 [RFC6274]	ARP cache overrun					H		Selective drop of packets

Table7. IPv6 Protocol Suite Threats

	ThreatID	Description	S	T	R	I	D	E	Mitigation
1	IETF-TDB -IPv6-1 [RFC4942]	Routing header to evade access controls	H				H		Access controls based on dest addresses
2	IETF-TDB -IPv6-2 [RFC4942]	Site-scope multicast addresses reconnaiss ance				H			Drop packets with site-scope dest addresses
3	IETF-TDB -IPv6-3 [RFC4942]	Anycast traffic identif reconnaiss ance				H			Restrict outside anycast services
4	IETF-TDB -IPv6-4 [RFC4942]	Extension headers excessive hop-by-hop options					H		Drop packets with unknown options
5	IETF-TDB -IPv6-5 [RFC4942]	Overuse of IPv6 router alert Option					H		Filter externally generated Router Alert packets
6	IETF-TDB -IPv6-6 [RFC4942]	IPv6 fragmentation overload of reconstruct buffers					H		Mandating the size of packet fragments
7	IETF-TDB -IPv6-7 [RFC4942]	IPv4-Mapped IPv6 Addresses	H				H		Avoid IPv4 -mapped IPv6 addesses

8	IETF-TDB -ICMPv6-1 [RFC4443]	ICMPv6 spoofing	H				H		IPAuth
9	IETF-TDB -ICMPv6-2 [RFC4443]	ICMPv6 Redirects	H		H	H			IPAuth or ESP
10	IETF-TDB -ICMPv6-3 [RFC4443]	Back-to -back erroneous IP packets					H		ICMP error rate limiting
11	IETF-TDB -ICMPv6-4 [RFC4443]	Send ICMP Parameter Problem to multicast source				H	H		Secure multicast traffic
12	IETF-TDB -ICMPv6-5 [RFC4443]	ICMP passed to upper-layers					H		IPSec
14	IETF-TDB -SLAAC-1 [RFC4942]	Address Privacy Extensions Interaction with DDoS Defenses					H		Tune the change rate of the node address
15	IETF-TDB -ND-1 [RFC3756]	NS/NA Spoofing	H				H		SEND
16	IETF-TDB -ND-2 [RFC3756]	NUD failure					H		SEND
17	IETF-TDB -ND-3 [RFC3756]	Malicious Last Hop Router			L	L	L		SEND
18	IETF-TDB -ND-4 [RFC3756]	Default router is 'killed'			L	L	L		No widely accepted mitigation technique

19	IETF-TDB -ND-5 [RFC3756]	Good Router Goes Bad			L	L	L		No widely accepted mitigation technique
20	IETF-TDB -ND-6 [RFC3756]	Spoofed Redirect Message			L	L	L		SEND; Still an issue for ad-hoc cases
21	IETF-TDB -ND-7 [RFC3756]	Bogus On-Link Prefix					L		SEND
22	IETF-TDB -ND-8 [RFC3756]	Bogus Address Config Prefix					L		SEND; Still an issue for ad-hoc cases
23	IETF-TDB -ND-9 [RFC3756]	Parameter Spoofing	L		L	L			SEND; Still an issue for ad-hoc cases
24	IETF-TDB -ND-10 [RFC3756]	ND Replay attacks	H			H			SEND
25	IETF-TDB -ND-11 [RFC3756]	Neighbor Discovery DoS					H		Rate limit NS messsages
26	IETF-TDB DAD_1 [RFC3756]	DAD DoS					H		SEND
27	IETF-TDB -SEND-1 [RFC3971]	Authorization Delegation Discovery DoS					H		Cache discovered info and limit the number of discovery processes

28	IETF-TDB -MIPv6-1 [RFC4942]	Obsolete Home Address Option Mobile IPv6	H						Secure Binding Update messages

Table8. Basic Transition Technologies Threats

	ThreatID	Description	S	T	R	I	D	E	Mitigation
1	IETF-TDB-IP/ICPM-1 [RFC6052]	IPv4 spoofing with IPv4 -embedded IPv6	L						Implement reverse path checks
2	IETF-TDB -IP/ICMP-2 [RFC6145]	ESP fails with IPv6 -to- IPv4 translation				L			Use checksum -neutral addresses
3	IETF-TDB -IP/ICMP-3 [rfc6145]	Auth Headers cannot be used across IPv6- to-IPv4				L			No widely accepted mitigation
4	IETF-TDB -IP/ICMP-4 [RFC6145]	Stateful translators resources exhaustion					L		No widely accepted mitigation
5	4encaps_1 [RFC4942]	Tunneling IPv6 over IPv4 breaks IPv4 Network's security assumptions				L			route encaps traffic through IPv4 firewall before decaps

Table9. L4 Technologies Threats

	ThreatID	Description	S	T	R	I	D	E	Mitigation

1	IETF-TDB -TCP-1 [harris99]	SYN flood					H		Block non- internal addresses from leaving
2	IETF-TDB -TCP-2 [harris99]	SYN /ACK flood	H		H		H		L3/L4 Packet Filtering
3	IETF-TDB -TCP-3 [harris99]	ACK or ACK -PUSH Flood	H		H		H		L3/L4 Packet Filtering
4	IETF-TDB -TCP-4 [harris99]	Frag mented ACK Flood					H		L3/L4 Packet Filtering
5	IETF-TDB -TCP-5 [harris99]	TCP Spoofing sequence number prediction	H						Block non -internal traffic from leaving
6	IETF-TDB -TCP-6 [harris99]	TCP session hijacking sequence number prediction	H	H	H	H	H	H	Block non -internal traffic from leaving
7	IETF-TDB -TCP-7 [harris99]	RST and FIN DoS					H		L3/L4 Packet Filtering Stateful Flow Awareness
8	IETF-TDB -UDP-8 [udps]	UDP flood					H		QoS regulation ;L3/L4 Packet Filtering

9	IETF-TDB -NAT44-9 [rfc7957]	Port set exaustion					H		Address Dependent Filtering

Table10. Routing Technologies Threats

	ThreatID	Description	S	T	R	I	D	E	Mitigation
1 x	IETF-TDB -RIPv2-1 [RFC4822]	simple password authen	L	L	L	L	L	L	crypto authen
2 x	IETF-TDB -OSPFv2-1 [RFC2328]	simple password authen	L	L	L	L	L	L	crypto authen
3 x	IETF-TDB -OSPFv2-2 [RFC2328]	OSPFv2 authen sequence number prediction	L	L	L	L	L	L	crypto sequence number
4	IETF-TDB -OSPFv3-1 [RFC4552]	OSPFv3 using the same manual key	L	L	L	L	L	L	no manual keys

Legend				
H	associaced with High likelihood		L	associaced with Low likelihood

Author's Address

Marius Georgescu (editor)
NAIST
Takayama 8916-5
Nara 630-0192
Japan

Phone: +81 743 72 5216
Email: liviumarius-g@is.naist.jp
URI: http://www.ipv6net.ro

opsec F. Gont
Internet-Draft UTN-FRH / SI6 Networks
Intended status: Informational W. Liu
Expires: January 9, 2017 Huawei Technologies
R. Bonica
Juniper Networks
July 8, 2016

Recommendations on Filtering of IPv6 Packets Containing IPv6 Extension Headers
draft-ietf-opsec-ipv6-eh-filtering-01

Abstract

It is common operator practice to mitigate security risks by enforcing appropriate packet filtering. This document analyzes both the general security implications of IPv6 Extension Headers and the specific security implications of each Extension Header and Option type, and provides advice on the filtering of IPv6 packets based on the IPv6 Extension Headers and the IPv6 options they contain. Additionally, it discusses the operational and interoperability implications of discarding packets based on the IPv6 Extension Headers and IPv6 options they contain.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 9, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

Recent studies (see e.g. [RFC7872]) suggest that there is widespread filtering of IPv6 packets that contain IPv6 Extension Headers (EHs). While some operators "officially" filter packets that contain IPv6 EHs, it is possible that some of the measured packet drops be the result of improper configuration defaults, or inappropriate advice in this area.

This document analyzes both the general security implications of IPv6 EHs and the specific security implications of each EH and Option type, and provides advice on the filtering of IPv6 packets based on the IPv6 EHs and the IPv6 options they contain. Since various protocols may use IPv6 EHs (possibly with IPv6 options), discarding packets based on the IPv6 EHs or IPv6 options they contain may have implications on the proper functioning of such protocols. Thus, this document also attempts to discuss the operational and interoperability implications of such filtering policies. This document is similar in nature to [RFC7126], which addresses the same problem for the IPv4 case. However, in IPv6, the problem space is compounded by the fact that IPv6 specifies a number of IPv6 EHs, and a number of IPv6 options which may be valid only when included in specific EH types.

This document completes and complements the considerations for protecting the control plane from packets containing IP options that can be found in [RFC6192].

Section 2 of this document specifies the terminology and conventions employed throughout this document. Section 3 of this document discusses IPv6 EHs and provides advice in the area of filtering IPv6 packets that contain such IPv6 EHs. Section 4 of this document discusses IPv6 options and provides advice in the area of filtering IPv6 packets that contain such options.

2. Terminology and Conventions Used in This Document

2.1. Terminology

The terms "fast path", "slow path", and associated relative terms ("faster path" and "slower path") are loosely defined as in Section 2 of [RFC6398].

The terms "permit" (allow the traffic), "drop" (drop with no notification to sender), and "reject" (drop with appropriate notification to sender) are employed as defined in [RFC3871]. Throughout this document we also employ the term "discard" as a generic term to indicate the act of discarding a packet, irrespective of whether the sender is notified of such drops, and irrespective of whether the specific filtering action is logged.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2.2. Conventions

This document assumes that nodes comply with the requirements in [RFC7045]. Namely (from [RFC7045]),

- o If a forwarding node discards a packet containing a standard IPv6 EH, it MUST be the result of a configurable policy and not just the result of a failure to recognise such a header.

- o The discard policy for each standard type of EH MUST be individually configurable.

- o The default configuration SHOULD allow all standard IPv6 EHs.

The advice provided in this document is only meant to guide an operator in configuring forwarding devices, and is *not* to be interpreted as advice regarding default configuration settings for network devices. That is, this document provides advice with respect to operational configurations, but does not change the implementation defaults required by [RFC7045]. We note that the advice provided in this document is *not* meant to be employed by transit routers for transit traffic, since such devices should not enforce this type of filtering policy on traffic not directed to them.

We recommend that configuration options are made available to govern the processing of each IPv6 EH type and each IPv6 option type. Such configuration options may include the following possible settings:

- o Permit this IPv6 EH or IPv6 Option type

- o Discard (and log) packets containing this IPv6 EH or option type

- o Reject (and log) packets containing this IPv6 EH or option type (where the packet drop is signaled with an ICMPv6 error message)

- o Rate-limit traffic containing this IPv6 EH or option type

- o Ignore this IPv6 EH or option type (as if it was not present) and forward the packet. We note that if a packet carries forwarding information (e.g., in an IPv6 Routing Header) this might be an inappropriate or undesirable action.

We note that special care needs to be taken when devices log packet drops/rejects. Devices should count the number of packets dropped/ rejected, but the logging of drop/reject events should be limited so as to not overburden device resources.

Finally, we note that when discarding packets, it is generally desirable that the sender be signaled of the packet drop, since this is of use for trouble-shooting purposes. However, throughout this document (when recommending that packets be discarded) we generically refer to the action as "discard" without specifying whether the sender is signaled of the packet drop.

3. IPv6 Extension Headers

3.1. General Discussion

IPv6 [RFC2460] EHs allow for the extension of the IPv6 protocol. Since both IPv6 EHs and upper-layer protocols share the same namespace ("Next Header" registry/namespace), [RFC7045] identifies which of the currently assigned Internet Protocol numbers identify IPv6 EHs vs. upper-layer protocols. This document discusses the filtering of packets based on the IPv6 EHs (as specified by [RFC7045]) they contain.

> NOTE: [RFC7112] specifies that non-fragmented IPv6 datagrams and IPv6 First-Fragments MUST contain the entire IPv6 header chain [RFC7112]. Therefore, intermediate systems can enforce the filtering policies discussed in this document, or resort to simply discarding the offending packets when they fail to comply with the requirements in [RFC7112]. We note that, in order to implement filtering rules on the fast path, it may be necessary for the filtering device to limit the depth into the packet that can be inspected before giving up. In circumstances where there is such a limitation, it is recommended that implementations discard packets if, when trying to determine whether to discard or permit a packet, the aforementioned limit is encountered.

3.2. General Security Implications

Depending on the specific device architecture, IPv6 packets that contain IPv6 EHs may cause the corresponding packets to be processed on the slow path, and hence may be leveraged for the purpose of Denial of Service (DoS) attacks [I-D.gont-v6ops-ipv6-ehs-packet-drops] [Cisco-EH] [FW-Benchmark].

Operators are urged to consider IPv6 EH filtering and IPv6 options handling capabilities of different devices as they make deployment decisions in future.

3.3. Advice on the Handling of IPv6 Packets with Specific IPv6 Extension Headers

3.3.1. IPv6 Hop-by-Hop Options (Protocol Number=0)

3.3.1.1. Uses

The Hop-by-Hop Options header is used to carry optional information that should be examined by every node along a packet's delivery path.

3.3.1.2. Specification

This EH is specified in [RFC2460], and its processing rules have been updated by [RFC7045]. At the time of this writing, the following options have been specified for the Hop-by-Hop Options EH:

- Type 0x00: Pad1 [RFC2460]
- Type 0x01: PadN [RFC2460]
- Type 0x05: Router Alert [RFC2711]
- Type 0x07: CALIPSO [RFC5570]
- Type 0x08: SMF_DPD [RFC6621]
- Type 0x26: Quick-Start [RFC4782]
- Type 0x4D: (Deprecated)
- Type 0x63: RPL Option [RFC6553]
- Type 0x6D: MPL Option [RFC7731]
- Type 0x8A: Endpoint Identification (Deprecated) [draft-ietf-nimrod-eid]
- Type 0xC2: Jumbo Payload [RFC2675]
- Type 0xEE: IPv6 DFF Header [RFC6971]
- Type 0x1E: RFC3692-style Experiment [RFC4727]
- Type 0x3E: RFC3692-style Experiment [RFC4727]
- Type 0x5E: RFC3692-style Experiment [RFC4727]
- Type 0x7E: RFC3692-style Experiment [RFC4727]

- o Type 0x9E: RFC3692-style Experiment [RFC4727]
- o Type 0xBE: RFC3692-style Experiment [RFC4727]
- o Type 0xDE: RFC3692-style Experiment [RFC4727]
- o Type 0xFE: RFC3692-style Experiment [RFC4727]

3.3.1.3. Specific Security Implications

Since this EH is required to be processed by all intermediate-systems en route, it can be leveraged to perform Denial of Service attacks against the network infrastructure.

NOTE: Ongoing work essentially aims at requiring the Hop-by-Hop Option EH to be processed only in cases where the intermmediate node is making use of any functionality provided by such header (see [I-D.ietf-6man-hbh-header-handling]). However, the deployed base is likely to reflect the traditional behavior for a while, and hence the potential security problems of this EH are still of concern.

3.3.1.4. Operational and Interoperability Impact if Blocked

Discarding packets containing a Hop-by-Hop Options EH would break any of the protocols that rely on it for proper functioning. For example, it would break RSVP [RFC2205] and multicast deployments, and would cause IPv6 jumbograms to be discarded.

3.3.1.5. Advice

The recommended configuration for the processing of these packets depends on the features and capabilities of the underlying platform. On platforms that allow forwarding of packets with HBH Options on the fast path, we recommend that packets with a HBH Options EH be forwarded as normal (for instance, [RFC7045] allows for implementations to ignore the HBH Options EH when forwarding packets). Otherwise, on platforms in which processing of packets with a IPv6 HBH Options EH is carried out in the slow path, and an option is provided to rate-limit these packets, we recommend that this option be selected. Finally, when packets containing a HBH Options EH are processed in the slow-path, and the underlying platform does not have any mitigation options available for attacks based on these packets, we recommend that such platforms discard packets containing IPv6 HBH Options EHs.

Finally, we note that, for obvious reasons, RPL (Routing Protocol for Low-Power and Lossy Networks) [RFC6550] routers must not discard packets based on the presence of an IPv6 Hop-by-Hop Options EH.

3.3.2. Routing Header for IPv6 (Protocol Number=43)

3.3.2.1. Uses

The Routing header is used by an IPv6 source to list one or more intermediate nodes to be "visited" on the way to a packet's destination.

3.3.2.2. Specification

This EH is specified in [RFC2460]. [RFC2460] originally specified the Routing Header Type 0, which has been later obsoleted by [RFC5095].

At the time of this writing, the following Routing Types have been specified:

- Type 0: Source Route (DEPRECATED) [RFC2460] [RFC5095]
- Type 1: Nimrod (DEPRECATED)
- Type 2: Type 2 Routing Header [RFC6275]
- Type 3: RPL Source Route Header [RFC6554]
- Types 4-252: Unassigned
- Type 253: RFC3692-style Experiment 1 [RFC4727]
- Type 254: RFC3692-style Experiment 2 [RFC4727]
- Type 255: Reserved

3.3.2.3. Specific Security Implications

The security implications of RHT0 have been discussed in detail in [Biondi2007] and [RFC5095].

3.3.2.4. Operational and Interoperability Impact if Blocked

Blocking packets containing a RHT0 or RTH1 has no operational implications. However, blocking packets employing other routing header types will break the protocols that rely on them.

3.3.2.5. Advice

Intermediate systems should discard packets containing a RHT0 or RHT1. RHT2 and RHT3 should be permitted, as required by [RFC7045]. Other routing header types should be discarded.

3.3.3. Fragment Header for IPv6 (Protocol Number=44)

3.3.3.1. Uses

This EH provides the fragmentation functionality for IPv6.

3.3.3.2. Specification

This EH is specified in [RFC2460].

3.3.3.3. Specific Security Implications

The security implications of the Fragment Header range from Denial of Service attacks (e.g. based on flooding a target with IPv6 fragments) to information leakage attacks [RFC7739].

3.3.3.4. Operational and Interoperability Impact if Blocked

Blocking packets that contain a Fragment Header will break any protocol that may rely on fragmentation (e.g., the DNS [RFC1034]).

3.3.3.5. Advice

Intermediate systems should permit packets that contain a Fragment Header.

3.3.4. Encapsulating Security Payload (Protocol Number=50)

3.3.4.1. Uses

This EH is employed for the IPsec suite [RFC4303].

3.3.4.2. Specification

This EH is specified in [RFC4303].

3.3.4.3. Specific Security Implications

Besides the general implications of IPv6 EHs, this EH could be employed to potentially perform a DoS attack at the destination system by wasting CPU resources in validating the contents of the packet.

3.3.4.4. Operational and Interoperability Impact if Blocked

Discarding packets that employ this EH would break IPsec deployments.

3.3.4.5. Advice

Intermediate systems should permit packets containing the Encapsulating Security Payload EH.

3.3.5. Authentication Header (Number=51)

3.3.5.1. Uses

The Authentication Header can be employed for provide authentication services in IPv4 and IPv6.

3.3.5.2. Specification

This EH is specified in [RFC4302].

3.3.5.3. Specific Security Implications

Besides the general implications of IPv6 EHs, this EH could be employed to potentially perform a DoS attack at the destination system by wasting CPU resources in validating the contents of the packet.

3.3.5.4. Operational and Interoperability Impact if Blocked

Discarding packets that employ this EH would break IPsec deployments.

3.3.5.5. Advice

Intermediate systems should permit packets containing an Authentication Header.

3.3.6. Destination Options for IPv6 (Protocol Number=60)

3.3.6.1. Uses

The Destination Options header is used to carry optional information that needs be examined only by a packet's destination node(s).

3.3.6.2. Specification

This EH is specified in [RFC2460]. At the time of this writing, the following options have been specified for this EH:

- Type 0x00: Pad1 [RFC2460]
- Type 0x01: PadN [RFC2460]
- Type 0x04: Tunnel Encapsulation Limit [RFC2473]
- Type 0x4D: (Deprecated)
- Type 0xC9: Home Address [RFC6275]
- Type 0x8A: Endpoint Identification (Deprecated) [draft-ietf-nimrod-eid]
- Type 0x8B: ILNP Nonce [RFC6744]
- Type 0x8C: Line-Identification Option [RFC6788]
- Type 0x1E: RFC3692-style Experiment [RFC4727]
- Type 0x3E: RFC3692-style Experiment [RFC4727]
- Type 0x5E: RFC3692-style Experiment [RFC4727]
- Type 0x7E: RFC3692-style Experiment [RFC4727]
- Type 0x9E: RFC3692-style Experiment [RFC4727]
- Type 0xBE: RFC3692-style Experiment [RFC4727]
- Type 0xDE: RFC3692-style Experiment [RFC4727]
- Type 0xFE: RFC3692-style Experiment [RFC4727]

3.3.6.3. Specific Security Implications

No security implications are known, other than the general implications of IPv6 EHs. For a discussion of possible security implications of specific options specified for the DO header, please see the Section 4.3.

3.3.6.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain a Destination Options header would break protocols that rely on this EH type for conveying information, including protocols such as ILNP [RFC6740] and Mobile IPv6 [RFC6275], and IPv6 tunnels that employ the Tunnel Encapsulation Limit option.

3.3.6.5. Advice

Intermediate systems should permit packets that contain a Destination Options Header.

3.3.7. Mobility Header (Number=135)

3.3.7.1. Uses

The Mobility Header is an EH used by mobile nodes, correspondent nodes, and home agents in all messaging related to the creation and management of bindings in Mobile IPv6.

3.3.7.2. Specification

This EH is specified in [RFC6275].

3.3.7.3. Specific Security Implications

A thorough security assessment of the security implications of the Mobility Header and related mechanisms can be found in Section 15 of [RFC6275].

3.3.7.4. Operational and Interoperability Impact if Blocked

Discarding packets containing this EH would break Mobile IPv6.

3.3.7.5. Advice

Intermediate systems should permit packets containing this EH.

3.3.8. Host Identity Protocol (Protocol Number=139)

3.3.8.1. Uses

This EH is employed with the Host Identity Protocol (HIP), an experimental protocol that allows consenting hosts to securely establish and maintain shared IP-layer state, allowing separation of the identifier and locator roles of IP addresses, thereby enabling continuity of communications across IP address changes.

3.3.8.2. Specification

This EH is specified in [RFC5201].

3.3.8.3. Specific Security Implications

The security implications of the HIP header are discussed in detail in Section 8 of [RFC6275].

3.3.8.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain the Host Identity Protocol would break HIP deployments.

3.3.8.5. Advice

Intermediate systems should permit packets that contain a Host Identity Protocol EH.

3.3.9. Shim6 Protocol (Protocol Number=140)

3.3.9.1. Uses

This EH is employed by the Shim6 [RFC5533] Protocol.

3.3.9.2. Specification

This EH is specified in [RFC5533].

3.3.9.3. Specific Security Implications

The specific security implications are discussed in detail in Section 16 of [RFC5533].

3.3.9.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this EH will break Shim6.

3.3.9.5. Advice

Intermediate systems should permit packets containing this EH.

3.3.10. Use for experimentation and testing (Protocol Numbers=253 and 254)

3.3.10.1. Uses

These IPv6 EHs are employed for performing RFC3692-Style experiments (see [RFC3692] for details).

3.3.10.2. Specification

These EHs are specified in [RFC3692] and [RFC4727].

3.3.10.3. Specific Security Implications

The security implications of these EHs will depend on their specific use.

3.3.10.4. Operational and Interoperability Impact if Blocked

For obvious reasons, discarding packets that contain these EHs limits the ability to perform legitimate experiments across IPv6 routers.

3.3.10.5. Advice

Intermediate systems should discard packets containing these EHs. Only in specific scenarios in which RFC3692-Style experiments are to be performed should these EHs be permitted.

3.4. Advice on the Handling of Packets with Unknown IPv6 Extension Headers

We refer to IPv6 EHs that have not been assigned an Internet Protocol Number by IANA (and marked as such) in [IANA-PROTOCOLS] as "unknown IPv6 extension headers" ("unknown IPv6 EHs").

3.4.1. Uses

New IPv6 EHs may be specified as part of future extensions to the IPv6 protocol.

Since IPv6 EHs and Upper-layer protocols employ the same namespace, it is impossible to tell whether an unknown "Internet Protocol Number" is being employed for an IPv6 EH or an Upper-Layer protocol.

3.4.2. Specification

The processing of unknown IPv6 EHs is specified in [RFC2460] and [RFC7045].

3.4.3. Specific Security Implications

For obvious reasons, it is impossible to determine specific security implications of unknown IPv6 EHs. However, from security standpoint, a device should discard IPv6 extension headers for which the security implications cannot be determined. We note that this policy is allowed by [RFC7045].

3.4.4. Operational and Interoperability Impact if Blocked

As noted in [RFC7045], discarding unknown IPv6 EHs may slow down the deployment of new IPv6 EHs and transport protocols. The corresponding IANA registry ([IANA-PROTOCOLS]) should be monitored such that filtering rules are updated as new IPv6 EHs are standardized.

We note that since IPv6 EHs and upper-layer protocols share the same numbering space, discarding unknown IPv6 EHs may result in packets encapsulating unknown upper-layer protocols being discarded.

3.4.5. Advice

Intermediate systems should discard packets containing unknown IPv6 EHs.

4. IPv6 Options

4.1. General Discussion

The following subsections describe specific security implications of different IPv6 options, and provide advice regarding filtering packets that contain such options.

4.2. General Security Implications of IPv6 Options

The general security implications of IPv6 options are closely related to those discussed in Section 3.2 for IPv6 EHs. Essentially, packets that contain IPv6 options might need to be processed by an IPv6 router’s general-purpose CPU,and hence could present a DDoS risk to that router’s general-purpose CPU (and thus to the router itself). For some architectures, a possible mitigation would be to rate-limit the packets that are to be processed by the general-purpose CPU (see e.g. [Cisco-EH]).

4.3. Advice on the Handling of Packets with Specific IPv6 Options

The following subsections contain a description of each of the IPv6 options that have so far been specified, a summary of the security implications of each of such options, a discussion of possible interoperability implications if packets containing such options are discarded, and specific advice regarding whether packets containing these options should be permitted.

4.3.1. Pad1 (Type=0x00)

4.3.1.1. Uses

This option is used when necessary to align subsequent options and to pad out the containing header to a multiple of 8 octets in length.

4.3.1.2. Specification

This option is specified in [RFC2460].

4.3.1.3. Specific Security Implications

None.

4.3.1.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option would potentially break any protocol that relies on IPv6 EHs.

4.3.1.5. Advice

Intermediate systems should not discard packets based on the presence of this option.

4.3.2. PadN (Type=0x01)

4.3.2.1. Uses

This option is used when necessary to align subsequent options and to pad out the containing header to a multiple of 8 octets in length.

4.3.2.2. Specification

This option is specified in [RFC2460].

4.3.2.3. Specific Security Implications

Because of the possible size of this option, it could be leveraged as a large-bandwidth covert channel.

4.3.2.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option would potentially break any protocol that relies on IPv6 EHs.

4.3.2.5. Advice

Intermediate systems should not discard IPv6 packets based on the presence of this option.

4.3.3. Jumbo Payload (Type=0XC2)

4.3.3.1. Uses

The Jumbo payload option provides the means of specifying payloads larger than 65535 bytes.

4.3.3.2. Specification

This option is specified in [RFC2675].

4.3.3.3. Specific Security Implications

There are no specific issues arising from this option, except for improper validity checks of the option and associated packet lengths.

4.3.3.4. Operational and Interoperability Impact if Blocked

Discarding packets based on the presence of this option will cause IPv6 jumbograms to be discarded.

4.3.3.5. Advice

Intermediate systems should discard packets that contain this option. An operator should permit this option only in specific scenarios in which support for IPv6 jumbograms is desired.

4.3.4. RPL Option (Type=0x63)

4.3.4.1. Uses

The RPL Option provides a mechanism to include routing information with each datagram that an RPL router forwards.

4.3.4.2. Specification

This option is specified in [RFC6553].

4.3.4.3. Specific Security Implications

Those described in [RFC6553].

4.3.4.4. Operational and Interoperability Impact if Blocked

This option is meant to be employed within an RPL instance. As a result, discarding packets based on the presence of this option (e.g. at an ISP) will not result in interoperability implications.

4.3.4.5. Advice

Non-RPL routers should discard packets that contain an RPL option.

4.3.5. Tunnel Encapsulation Limit (Type=0x04)

4.3.5.1. Uses

The Tunnel Encapsulation Limit option can be employed to specify how many further levels of nesting the packet is permitted to undergo.

4.3.5.2. Specification

This option is specified in [RFC2473].

4.3.5.3. Specific Security Implications

Those described in [RFC2473].

4.3.5.4. Operational and Interoperability Impact if Blocked

Discarding packets based on the presence of this option could result in tunnel traffic being discarded.

4.3.5.5. Advice

Intermediate systems should not discard packets based on the presence of this option.

4.3.6. Router Alert (Type=0x05)

4.3.6.1. Uses

The Router Alert option [RFC2711] is typically employed for the RSVP protocol [RFC2205] and the MLD protocol [RFC2710].

4.3.6.2. Specification

This option is specified in [RFC2711].

4.3.6.3. Specific Security Implications

Since this option causes the contents of the packet to be inspected by the handling device, this option could be leveraged for performing DoS attacks.

4.3.6.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option would break RSVP and multicast deployments.

4.3.6.5. Advice

Intermediate systems should discard packets that contain this option. Only in specific environments where support for RSVP, multicast routing, or similar protocols is desired, should this option be permitted.

4.3.7. Quick-Start (Type=0x26)

4.3.7.1. Uses

This IP Option is used in the specification of Quick-Start for TCP and IP, which is an experimental mechanism that allows transport protocols, in cooperation with routers, to determine an allowed sending rate at the start and, at times, in the middle of a data transfer (e.g., after an idle period) [RFC4782].

4.3.7.2. Specification

This option is specified in [RFC4782], on the "Experimental" track.

4.3.7.3. Specific Security Implications

Section 9.6 of [RFC4782] notes that Quick-Start is vulnerable to two kinds of attacks:

o attacks to increase the routers' processing and state load, and,

o attacks with bogus Quick-Start Requests to temporarily tie up available Quick-Start bandwidth, preventing routers from approving Quick-Start Requests from other connections.

We note that if routers in a given environment do not implement and enable the Quick-Start mechanism, only the general security implications of IP options (discussed in Section 4.2) would apply.

4.3.7.4. Operational and Interoperability Impact if Blocked

The Quick-Start functionality would be disabled, and additional delays in TCP's connection establishment (for example) could be introduced. (Please see Section 4.7.2 of [RFC4782].) We note, however, that Quick-Start has been proposed as a mechanism that could be of use in controlled environments, and not as a mechanism that would be intended or appropriate for ubiquitous deployment in the global Internet [RFC4782].

4.3.7.5. Advice

Intermediate systems should not discard IPv6 packets based on the presence of this option.

4.3.8. CALIPSO (Type=0x07)

4.3.8.1. Uses

This option is used for encoding explicit packet Sensitivity Labels on IPv6 packets. It is intended for use only within Multi-Level Secure (MLS) networking environments that are both trusted and trustworthy.

4.3.8.2. Specification

This option is specified in [RFC5570].

4.3.8.3. Specific Security Implications

Presence of this option in a packet does not by itself create any specific new threat. Packets with this option ought not normally be seen on the global public Internet.

4.3.8.4. Operational and Interoperability Impact if Blocked

If packets with this option are discarded or if the option is stripped from the packet during transmission from source to destination, then the packet itself is likely to be discarded by the receiver because it is not properly labeled. In some cases, the receiver might receive the packet but associate an incorrect sensitivity label with the received data from the packet whose CALIPSO was stripped by an intermediate router or firewall. Associating an incorrect sensitivity label can cause the received information either to be handled as more sensitive than it really is ("upgrading") or as less sensitive than it really is ("downgrading"), either of which is problematic.

4.3.8.5. Advice

Intermediate systems that do not operate in Multi-Level Secure (MLS) networking environments should discard packets that contain this option.

4.3.9. SMF_DPD (Type=0x08)

4.3.9.1. Uses

This option is employed in the (experimental) Simplified Multicast Forwarding (SMF) for unique packet identification for IPv6 I-DPD, and as a mechanism to guarantee non-collision of hash values for different packets when H-DPD is used.

4.3.9.2. Specification

This option is specified in [RFC6621].

4.3.9.3. Specific Security Implications

TBD.

4.3.9.4. Operational and Interoperability Impact if Blocked

TBD.

4.3.9.5. Advice

TBD.

4.3.10. Home Address (Type=0xC9)

4.3.10.1. Uses

The Home Address option is used by a Mobile IPv6 node while away from home, to inform the recipient of the mobile node's home address.

4.3.10.2. Specification

This option is specified in [RFC6275].

4.3.10.3. Specific Security Implications

No (known) additional security implications than those described in [RFC6275].

4.3.10.4. Operational and Interoperability Impact if Blocked

Discarding IPv6 packets based on the presence of this option will break Mobile IPv6.

4.3.10.5. Advice

Intermediate systems should not discard IPv6 packets based on the presence of this option.

4.3.11. Endpoint Identification (Type=0x8A)

4.3.11.1. Uses

The Endpoint Identification option was meant to be used with the Nimrod routing architecture [NIMROD-DOC], but has never seen widespread deployment.

4.3.11.2. Specification

This option is specified in [NIMROD-DOC].

4.3.11.3. Specific Security Implications

TBD.

4.3.11.4. Operational and Interoperability Impact if Blocked

None.

4.3.11.5. Advice

Intermediate systems should discard packets that contain this option.

4.3.12. ILNP Nonce (Type=0x8B)

4.3.12.1. Uses

This option is employed by Identifier-Locator Network Protocol for IPv6 (ILNPv6) for providing protection against off-path attacks for packets when ILNPv6 is in use, and as a signal during initial network-layer session creation that ILNPv6 is proposed for use with this network-layer session, rather than classic IPv6.

4.3.12.2. Specification

This option is specified in [RFC6744].

4.3.12.3. Specific Security Implications

Those described in [RFC6744].

4.3.12.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option will break INLPv6 deployments.

4.3.12.5. Advice

Intermediate systems should not discard packets based on the presence of this option.

4.3.13. Line-Identification Option (Type=0x8C)

4.3.13.1. Uses

This option is used by an Edge Router to identify the subscriber premises in scenarios where several subscriber premises may be logically connected to the same interface of an Edge Router.

4.3.13.2. Specification

This option is specified in [RFC6788].

4.3.13.3. Specific Security Implications

Those described in [RFC6788].

4.3.13.4. Operational and Interoperability Impact if Blocked

Since this option is meant to be employed in Router Solicitation messages, discarding packets based on the presence of this option at intermediate systems will result in no interoperability implications.

4.3.13.5. Advice

Intermediate devices should discard packets that contain this option.

4.3.14. Deprecated (Type=0x4D)

4.3.14.1. Uses

No information has been found about this option type.

4.3.14.2. Specification

No information has been found about this option type.

4.3.14.3. Specific Security Implications

No information has been found about this option type, and hence it has been impossible to perform the corresponding security assessment.

4.3.14.4. Operational and Interoperability Impact if Blocked

Unknown.

4.3.14.5. Advice

Intermediate systems should discard packets that contain this option.

4.3.15. MPL Option (Type=0x6D)

4.3.15.1. Uses

This option is used with the Multicast Protocol for Low power and Lossy Networks (MPL), that provides IPv6 multicast forwarding in constrained networks.

4.3.15.2. Specification

This option is specified in [RFC7731], and is meant to be included only in Hop-by-Hop Option headers.

4.3.15.3. Specific Security Implications

Those described in [RFC7731].

4.3.15.4. Operational and Interoperability Impact if Blocked

TBD.

4.3.15.5. Advice

TBD.

4.3.16. IP_DFF (Type=0xEE)

4.3.16.1. Uses

This option is employed with the (Experimental) Depth-First Forwarding (DFF) in Unreliable Networks.

4.3.16.2. Specification

This option is specified in [RFC6971].

4.3.16.3. Specific Security Implications

Those specified in [RFC6971].

4.3.16.4. Operational and Interoperability Impact if Blocked

TBD.

4.3.16.5. Advice

TBD.

4.3.17. RFC3692-style Experiment (Types = 0x1E, 0x3E, 0x5E, 0x7E, 0x9E, 0xBE, 0xDE, 0xFE)

4.3.17.1. Uses

These options can be employed for performing RFC3692-style experiments. It is only appropriate to use these values in explicitly configured experiments; they must not be shipped as defaults in implementations.

4.3.17.2. Specification

Specified in RFC 4727 [RFC4727] in the context of RFC3692-style experiments.

4.3.17.3. Specific Security Implications

The specific security implications will depend on the specific use of these options.

4.3.17.4. Operational and Interoperability Impact if Blocked

For obvious reasons, discarding packets that contain these options limits the ability to perform legitimate experiments across IPv6 routers.

4.3.17.5. Advice

Intermediate systems should discard packets that contain these options. Only in specific environments where RFC3692-style experiments are meant to be performed should these options be permitted.

4.4. Advice on the handling of Packets with Unknown IPv6 Options

We refer to IPv6 options that have not been assigned an IPv6 option type in the corresponding registry ([IANA-IPV6-PARAM]) as "unknown IPv6 options".

4.4.1. Uses

New IPv6 options may be specified as part of future protocol work.

4.4.2. Specification

The processing of unknown IPv6 options is specified in [RFC2460].

4.4.3. Specific Security Implications

For obvious reasons, it is impossible to determine specific security implications of unknown IPv6 options.

4.4.4. Operational and Interoperability Impact if Blocked

Discarding unknown IPv6 options may slow down the deployment of new IPv6 options. As noted in [draft-gont-6man-ipv6-opt-transmit], the corresponding IANA registry ([IANA-IPV6-PARAM] should be monitored such that IPv6 option filtering rules are updated as new IPv6 options are standardized.

4.4.5. Advice

Enterprise intermediate systems that process the contents of IPv6 EHs should discard packets that contain unknown options. Other intermediate systems that process the contents of IPv6 EHs should permit packets that contain unknown options.

5. IANA Considerations

This document has no actions for IANA.

6. Security Considerations

This document provides advice on the filtering of IPv6 packets that contain IPv6 EHs (and possibly IPv6 options). Discarding such packets can help to mitigate the security issues that arise from the use of different IPv6 EHs and options.

7. Acknowledgements

The authors of this document would like to thank (in alphabetical order) Mikael Abrahamsson, Brian Carpenter, Mike Heard, Jen Linkova, Carlos Pignataro, Donald Smith, and Gunter Van De Velde, for providing valuable comments on earlier versions of this document.

This document borrows some text an analysis from [RFC7126], authored by Fernando Gont, Randall Atkinson, and Carlos Pignataro.

8. References

8.1. Normative References

[draft-gont-6man-ipv6-opt-transmit]
Gont, F., Liu, W., and R. Bonica, "Transmission and Processing of IPv6 Options", IETF Internet Draft, work in progress, August 2014.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <http://www.rfc-editor.org/info/rfc1034>.

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

[RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <http://www.rfc-editor.org/info/rfc2205>.

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

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <http://www.rfc-editor.org/info/rfc2473>.

[RFC2675] Borman, D., Deering, S., and R. Hinden, "IPv6 Jumbograms", RFC 2675, DOI 10.17487/RFC2675, August 1999, <http://www.rfc-editor.org/info/rfc2675>.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <http://www.rfc-editor.org/info/rfc2710>.

[RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", RFC 2711, DOI 10.17487/RFC2711, October 1999, <http://www.rfc-editor.org/info/rfc2711>.

[RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, DOI 10.17487/RFC3692, January 2004, <http://www.rfc-editor.org/info/rfc3692>.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <http://www.rfc-editor.org/info/rfc4302>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <http://www.rfc-editor.org/info/rfc4303>.

[RFC4304] Kent, S., "Extended Sequence Number (ESN) Addendum to IPsec Domain of Interpretation (DOI) for Internet Security Association and Key Management Protocol (ISAKMP)", RFC 4304, DOI 10.17487/RFC4304, December 2005, <http://www.rfc-editor.org/info/rfc4304>.

[RFC4727] Fenner, B., "Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers", RFC 4727, DOI 10.17487/RFC4727, November 2006, <http://www.rfc-editor.org/info/rfc4727>.

[RFC4782] Floyd, S., Allman, M., Jain, A., and P. Sarolahti, "Quick-Start for TCP and IP", RFC 4782, DOI 10.17487/RFC4782, January 2007, <http://www.rfc-editor.org/info/rfc4782>.

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

[RFC5201] Moskowitz, R., Nikander, P., Jokela, P., Ed., and T. Henderson, "Host Identity Protocol", RFC 5201, DOI 10.17487/RFC5201, April 2008, <http://www.rfc-editor.org/info/rfc5201>.

[RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, DOI 10.17487/RFC5533, June 2009, <http://www.rfc-editor.org/info/rfc5533>.

[RFC5570] StJohns, M., Atkinson, R., and G. Thomas, "Common Architecture Label IPv6 Security Option (CALIPSO)", RFC 5570, DOI 10.17487/RFC5570, July 2009, <http://www.rfc-editor.org/info/rfc5570>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <http://www.rfc-editor.org/info/rfc6275>.

[RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <http://www.rfc-editor.org/info/rfc6398>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <http://www.rfc-editor.org/info/rfc6550>.

[RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <http://www.rfc-editor.org/info/rfc6553>.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <http://www.rfc-editor.org/info/rfc6554>.

[RFC6621] Macker, J., Ed., "Simplified Multicast Forwarding", RFC 6621, DOI 10.17487/RFC6621, May 2012, <http://www.rfc-editor.org/info/rfc6621>.

[RFC6740] Atkinson, RJ. and SN. Bhatti, "Identifier-Locator Network Protocol (ILNP) Architectural Description", RFC 6740, DOI 10.17487/RFC6740, November 2012, <http://www.rfc-editor.org/info/rfc6740>.

[RFC6744] Atkinson, RJ. and SN. Bhatti, "IPv6 Nonce Destination Option for the Identifier-Locator Network Protocol for IPv6 (ILNPv6)", RFC 6744, DOI 10.17487/RFC6744, November 2012, <http://www.rfc-editor.org/info/rfc6744>.

[RFC6788] Krishnan, S., Kavanagh, A., Varga, B., Ooghe, S., and E. Nordmark, "The Line-Identification Option", RFC 6788, DOI 10.17487/RFC6788, November 2012, <http://www.rfc-editor.org/info/rfc6788>.

[RFC6971] Herberg, U., Ed., Cardenas, A., Iwao, T., Dow, M., and S. Cespedes, "Depth-First Forwarding (DFF) in Unreliable Networks", RFC 6971, DOI 10.17487/RFC6971, June 2013, <http://www.rfc-editor.org/info/rfc6971>.

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

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

[RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <http://www.rfc-editor.org/info/rfc7731>.

8.2. Informative References

[Biondi2007]
Biondi, P. and A. Ebalard, "IPv6 Routing Header Security", CanSecWest 2007 Security Conference, 2007, <http://www.secdev.org/conf/IPv6_RH_security-csw07.pdf>.

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

[draft-ietf-nimrod-eid]
Lynn, C., "Endpoint Identifier Destination Option", IETF Internet Draft, draft-ietf-nimrod-eid-00.txt, November 1995.

[FW-Benchmark]
Zack, E., "Firewall Security Assessment and Benchmarking IPv6 Firewall Load Tests", IPv6 Hackers Meeting #1, Berlin, Germany. June 30, 2013, <http://www.ipv6hackers.org/meetings/ipv6-hackers-1/zack-ipv6hackers1-firewall-security-assessment-and-benchmarking.pdf>.

[I-D.gont-v6ops-ipv6-ehs-packet-drops]
Gont, F., Hilliard, N., Doering, G., (Will), S., and W. Kumari, "Operational Implications of IPv6 Packets with Extension Headers", draft-gont-v6ops-ipv6-ehs-packet-drops-03 (work in progress), March 2016.

[I-D.ietf-6man-hbh-header-handling]
Baker, F. and R. Bonica, "IPv6 Hop-by-Hop Options Extension Header", draft-ietf-6man-hbh-header-handling-03 (work in progress), March 2016.

[IANA-IPV6-PARAM]
Internet Assigned Numbers Authority, "Internet Protocol Version 6 (IPv6) Parameters", December 2013, <http://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml>.

[IANA-PROTOCOLS]
Internet Assigned Numbers Authority, "Protocol Numbers", 2014, <http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml>.

[NIMROD-DOC]
Nimrod Documentation Page, , "http://ana-3.lcs.mit.edu/~jnc/nimrod/".

[RFC3871] Jones, G., Ed., "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure", RFC 3871, DOI 10.17487/RFC3871, September 2004, <http://www.rfc-editor.org/info/rfc3871>.

[RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, DOI 10.17487/RFC6192, March 2011, <http://www.rfc-editor.org/info/rfc6192>.

[RFC7126] Gont, F., Atkinson, R., and C. Pignataro, "Recommendations on Filtering of IPv4 Packets Containing IPv4 Options", BCP 186, RFC 7126, DOI 10.17487/RFC7126, February 2014, <http://www.rfc-editor.org/info/rfc7126>.

[RFC7739] Gont, F., "Security Implications of Predictable Fragment Identification Values", RFC 7739, DOI 10.17487/RFC7739, February 2016, <http://www.rfc-editor.org/info/rfc7739>.

[RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <http://www.rfc-editor.org/info/rfc7872>.

Authors' Addresses

Fernando Gont
UTN-FRH / SI6 Networks
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: http://www.si6networks.com

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

Email: liushucheng@huawei.com

Ronald P. Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20171
US

Phone: 571 250 5819
Email: rbonica@juniper.net

opsec F. Gont
Internet-Draft SI6 Networks
Intended status: Informational W. Liu
Expires: July 23, 2021 Huawei Technologies
January 19, 2021

Recommendations on the Filtering of IPv6 Packets Containing IPv6 Extension Headers at Transit Routers
draft-ietf-opsec-ipv6-eh-filtering-07

Abstract

This document analyzes the security implications of IPv6 Extension Headers and associated IPv6 options. Additionally, it discusses the operational and interoperability implications of discarding packets based on the IPv6 Extension Headers and IPv6 options they contain. Finally, it provides advice on the filtering of such IPv6 packets at transit routers for traffic *not* directed to them, for those cases where such filtering is deemed as necessary.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 23, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

IPv6 Extension Headers (EHs) allow for the extension of the IPv6 protocol, and provide support for core functionality such as IPv6 fragmentation. However, common implementation limitations suggest that EHs present a challenge for IPv6 packet routing equipment, particularly when the IPv6 header chain needs to be processed for e.g. enforcing ACLs or implementing other functions [I-D.ietf-v6ops-ipv6-ehs-packet-drops].

Recent studies (see e.g. [RFC7872]) suggest that there is widespread dropping of IPv6 packets that contain IPv6 Extension Headers (EHs). In some cases, such packet drops occur at transit routers. While some operators "officially" drop packets that contain IPv6 EHs, it is possible that some of the measured packet drops be the result of improper configuration defaults, or inappropriate advice in this area.

This document analyzes both the general security implications of IPv6 EHs, as well as the security implications of specific EH and Option types. It also provides advice on the filtering of IPv6 packets based on the IPv6 EHs and the IPv6 options they contain. Since various protocols may use IPv6 EHs (possibly with IPv6 options), discarding packets based on the IPv6 EHs or IPv6 options they contain can have implications on the proper functioning of such protocols. Thus, this document also attempts to discuss the operational and interoperability implications of such filtering policies.

The resulting packet filtering policy typically depends on where in the network such policy is enforced: when the policy is enforced in a transit network, the policy typically follows a "deny-list" approach, where only packets with clear negative implications are dropped. On the other hand, when the policy is enforced closer to the destination systems, the policy typically follows an "accept-list" approach, where only traffic that is expected to be received is allowed. The advice in this document is aimed only at transit routers that may need to enforce a filtering policy based on the EHs and IPv6 options a packet may contain, following a "deny-list" approach, and hence is likely to be much more permissive that a filtering policy to be employed at e.g. the edge of an enterprise network. The advice in this document is meant to improve the current situation of the dropping of packets with IPv6 EHs in the Internet [RFC7872] in such cases where packets are being dropped due to inappropriate or missing guidelines.

This document is similar in nature to [RFC7126], which addresses the same problem for the IPv4 case. However, in IPv6, the problem space is compounded by the fact that IPv6 specifies a number of IPv6 EHs, and a number of IPv6 options which may be valid only when included in specific EH types.

This document completes and complements the considerations for protecting the control plane from packets containing IP options that can be found in [RFC6192].

Section 2 of this document specifies the terminology and conventions employed throughout this document. Section 3 of this document discusses IPv6 EHs and provides advice in the area of filtering IPv6 packets that contain such IPv6 EHs. Section 4 of this document discusses IPv6 options and provides advice in the area of filtering IPv6 packets that contain such options.

2. Terminology and Conventions Used in This Document

2.1. Terminology

The terms "permit" (allow the traffic), "drop" (drop with no notification to sender), and "reject" (drop with appropriate notification to sender) are employed as defined in [RFC3871]. Throughout this document we also employ the term "discard" as a generic term to indicate the act of discarding a packet, irrespective of whether the sender is notified of such drops, and irrespective of whether the specific filtering action is logged.

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.

2.2. Applicability Statement

This document provides advice on the filtering of IPv6 packets with EHs at transit routers for traffic *not* explicitly destined to them, for cases in which such filtering is deemed as necessary.

2.3. Conventions

This document assumes that nodes comply with the requirements in [RFC7045]. Namely,

- o If a forwarding node discards a packet containing a standard IPv6 EH, it MUST be the result of a configurable policy and not just the result of a failure to recognise such a header.

- o The discard policy for each standard type of EH MUST be individually configurable.

- o The default configuration should allow all standard IPv6 EHs.

The advice provided in this document is only meant to guide an operator in configuring forwarding devices, and is *not* to be interpreted as advice regarding default configuration settings for network devices. That is, this document provides advice with respect to operational configurations, but does not change the implementation defaults required by [RFC7045].

We recommend that configuration options are made available to govern the processing of each IPv6 EH type and each IPv6 option type. Such configuration options should include the following possible settings:

- o Permit this IPv6 EH or IPv6 Option type.
- o Discard (and log) packets containing this IPv6 EH or option type.
- o Reject (and log) packets containing this IPv6 EH or option type (where the packet drop is signaled with an ICMPv6 error message).
- o Rate-limit traffic containing this IPv6 EH or option type.
- o Ignore this IPv6 EH or option type (as if it was not present) and forward the packet. We note that if a packet carries forwarding information (e.g., in an IPv6 Routing Header) this might be an inappropriate or undesirable action.

We note that special care needs to be taken when devices log packet drops/rejects. Devices should count the number of packets dropped/rejected, but the logging of drop/reject events should be limited so as to not overburden device resources.

Finally, we note that when discarding packets, it is generally desirable that the sender be signaled of the packet drop, since this is of use for trouble-shooting purposes. However, throughout this document (when recommending that packets be discarded) we generically refer to the action as "discard" without specifying whether the sender is signaled of the packet drop.

3. IPv6 Extension Headers

3.1. General Discussion

IPv6 [RFC8200] EHs allow for the extension of the IPv6 protocol. Since both IPv6 EHs and upper-layer protocols share the same namespace ("Next Header" registry/namespace), [RFC7045] identifies which of the currently assigned Internet Protocol numbers identify IPv6 EHs vs. upper-layer protocols. This document discusses the filtering of packets based on the IPv6 EHs (as specified by [RFC7045]) they contain.

> NOTE: [RFC8200] specifies that non-fragmented IPv6 datagrams and IPv6 First-Fragments must contain the entire IPv6 header chain [RFC7112]. Therefore, intermediate systems can enforce the filtering policies discussed in this document, or resort to simply discarding the offending packets when they fail to comply with the requirements in [RFC8200]. We note that, in order to implement

filtering rules on the fast path, it may be necessary for the filtering device to limit the depth into the packet that can be inspected before giving up. In circumstances where such a limitation exists, it is recommended that implementations provide a configuration option that specifies whether to discard packets if the aforementioned limit is encountered. Operators may then determine according to their own circumstances how such packets will be handled.

3.2. General Security Implications

In some device architectures, IPv6 packets that contain IPv6 EHs can cause the corresponding packets to be processed on the slow path, and hence may be leveraged for the purpose of Denial of Service (DoS) attacks [I-D.ietf-v6ops-ipv6-ehs-packet-drops] [Cisco-EH] [FW-Benchmark].

Operators are urged to consider the IPv6 EH and IPv6 options handling capabilities of their devices as they make deployment decisions in future.

3.3. Summary of Advice on the Handling of IPv6 Packets with Specific IPv6 Extension Headers

This section summarizes the advice provided in Section 3.4, providing references to the specific sections in which a detailed analysis can be found.

EH type	Filtering policy	Reference
IPv6 Hop-by-Hop Options (Proto=0)	Drop or Ignore	Section 3 .4.1
Routing Header for IPv6 (Proto=43)	Drop only RHT0 and RHT1. Permit other RH Types	Section 3 .4.2
Fragment Header for IPv6 (Proto=44)	Permit	Section 3 .4.3
Encapsulating Security Payload (Proto=50)	Permit	Section 3 .4.4
Authentication Header (Proto=51)	Permit	Section 3 .4.5
Destination Options for IPv6 (Proto=60)	Permit	Section 3 .4.6
Mobility Header (Proto=135)	Permit	Section 3 .4.7
Host Identity Protocol (Proto=139)	Permit	Section 3 .4.8
Shim6 Protocol (Proto=140)	Permit	Section 3 .4.9
Use for experimentation and testing (Proto=253 and 254)	Drop	Section 3 .4.10

Table 1: Summary of Advice on the Handling of IPv6 Packets with Specific IPv6 Extension Headers

3.4. Advice on the Handling of IPv6 Packets with Specific IPv6 Extension Headers

3.4.1. IPv6 Hop-by-Hop Options (Protocol Number=0)

3.4.1.1. Uses

The Hop-by-Hop Options header is used to carry optional information that may be examined by every node along a packet's delivery path.

It is expected that nodes will examine the Hop-by-Hop Options header if explicitly configured to do so.

NOTE: A previous revision of the IPv6 core specification, [RFC2460], originally required that all nodes examined and processed the Hop-by-Hop Options header. However, even before the publication of [RFC8200] a number of implementations already provided the option of ignoring this header unless explicitly configured to examine it.

3.4.1.2. Specification

This EH is specified in [RFC8200]. At the time of this writing, the following options have been specified for the Hop-by-Hop Options EH:

- Type 0x00: Pad1 [RFC8200]
- Type 0x01: PadN [RFC8200]
- Type 0x05: Router Alert [RFC2711]
- Type 0x07: CALIPSO [RFC5570]
- Type 0x08: SMF_DPD [RFC6621]
- Type 0x23: RPL Option [I-D.ietf-roll-useofrplinfo]
- Type 0x26: Quick-Start [RFC4782]
- Type 0x4D: (Deprecated)
- Type 0x63: RPL Option [RFC6553]
- Type 0x6D: MPL Option [RFC7731]
- Type 0x8A: Endpoint Identification (Deprecated) [draft-ietf-nimrod-eid]
- Type 0xC2: Jumbo Payload [RFC2675]
- Type 0xEE: IPv6 DFF Header [RFC6971]
- Type 0x1E: RFC3692-style Experiment [RFC4727]
- Type 0x3E: RFC3692-style Experiment [RFC4727]
- Type 0x5E: RFC3692-style Experiment [RFC4727]
- Type 0x7E: RFC3692-style Experiment [RFC4727]

- o Type 0x9E: RFC3692-style Experiment [RFC4727]
- o Type 0xBE: RFC3692-style Experiment [RFC4727]
- o Type 0xDE: RFC3692-style Experiment [RFC4727]
- o Type 0xFE: RFC3692-style Experiment [RFC4727]

3.4.1.3. Specific Security Implications

Legacy nodes that process this extension header might be subject to Denial of Service attacks.

NOTE: While [RFC8200] has removed this requirement, the deployed base may still reflect the traditional behavior for a while, and hence the potential security problems of this EH are still of concern.

3.4.1.4. Operational and Interoperability Impact if Blocked

Discarding packets containing a Hop-by-Hop Options EH would break any of the protocols that rely on it for proper functioning. For example, it would break RSVP [RFC2205] and multicast deployments, and would cause IPv6 jumbograms to be discarded.

3.4.1.5. Advice

Nodes implementing [RFC8200] would already ignore this extension header unless explicitly required to process it. For legacy ([RFC2460]) nodes, the recommended configuration for the processing of these packets depends on the features and capabilities of the underlying platform, the configuration of the platform, and also the deployment environment of the platform. On platforms that allow forwarding of packets with HBH Options on the fast path, we recommend that packets with a HBH Options EH be forwarded as normal. Otherwise, on platforms in which processing of packets with a IPv6 HBH Options EH is carried out in the slow path, and an option is provided to rate-limit these packets, we recommend that this option be selected. Finally, when packets containing a HBH Options EH are processed in the slow-path, and the underlying platform does not have any mitigation options available for attacks based on these packets, we recommend that such platforms discard packets containing IPv6 HBH Options EHs.

Finally, we note that, for obvious reasons, RPL (Routing Protocol for Low-Power and Lossy Networks) [RFC6550] routers must not discard packets based on the presence of an IPv6 Hop-by-Hop Options EH.

3.4.2. Routing Header for IPv6 (Protocol Number=43)

3.4.2.1. Uses

The Routing header is used by an IPv6 source to list one or more intermediate nodes to be "visited" on the way to a packet's destination.

3.4.2.2. Specification

This EH is specified in [RFC8200]. [RFC2460] had originally specified the Routing Header Type 0, which was later obsoleted by [RFC5095], and thus removed from [RFC8200].

At the time of this writing, the following Routing Types have been specified:

- Type 0: Source Route (DEPRECATED) [RFC2460] [RFC5095]
- Type 1: Nimrod (DEPRECATED)
- Type 2: Type 2 Routing Header [RFC6275]
- Type 3: RPL Source Route Header [RFC6554]
- Type 4: Segment Routing Header (SRH) [RFC8754]
- Types 5-252: Unassigned
- Type 253: RFC3692-style Experiment 1 [RFC4727]
- Type 254: RFC3692-style Experiment 2 [RFC4727]
- Type 255: Reserved

3.4.2.3. Specific Security Implications

The security implications of RHT0 have been discussed in detail in [Biondi2007] and [RFC5095]. The security implications of RHT4 (SRH) are discussed in [RFC8754].

3.4.2.4. Operational and Interoperability Impact if Blocked

Blocking packets containing a RHT0 or RHT1 has no operational implications, since both have been deprecated. Blocking packets with a RHT4 (SRH) will break Segment Routing (SR) deployments, if the filtering policy is enforced on packets being forwarded within an SR domain.

3.4.2.5. Advice

Intermediate systems should discard packets containing a RHT0 or RHT1. Other routing header types should be permitted, as required by [RFC7045].

3.4.3. Fragment Header for IPv6 (Protocol Number=44)

3.4.3.1. Uses

This EH provides the fragmentation functionality for IPv6.

3.4.3.2. Specification

This EH is specified in [RFC8200].

3.4.3.3. Specific Security Implications

The security implications of the Fragment Header range from Denial of Service attacks (e.g. based on flooding a target with IPv6 fragments) to information leakage attacks [RFC7739].

3.4.3.4. Operational and Interoperability Impact if Blocked

Blocking packets that contain a Fragment Header will break any protocol that may rely on fragmentation (e.g., the DNS [RFC1034]). However, IP fragmentation is known to introduce fragility to Internet communication [RFC8900].

3.4.3.5. Advice

Intermediate systems should permit packets that contain a Fragment Header.

3.4.4. Encapsulating Security Payload (Protocol Number=50)

3.4.4.1. Uses

This EH is employed for the IPsec suite [RFC4303].

3.4.4.2. Specification

This EH is specified in [RFC4303].

3.4.4.3. Specific Security Implications

Besides the general implications of IPv6 EHs, this EH could be employed to potentially perform a DoS attack at the destination system by wasting CPU resources in validating the contents of the packet.

3.4.4.4. Operational and Interoperability Impact if Blocked

Discarding packets that employ this EH would break IPsec deployments.

3.4.4.5. Advice

Intermediate systems should permit packets containing the Encapsulating Security Payload EH.

3.4.5. Authentication Header (Protocol Number=51)

3.4.5.1. Uses

The Authentication Header can be employed for provide authentication services in IPv4 and IPv6.

3.4.5.2. Specification

This EH is specified in [RFC4302].

3.4.5.3. Specific Security Implications

Besides the general implications of IPv6 EHs, this EH could be employed to potentially perform a DoS attack at the destination system by wasting CPU resources in validating the contents of the packet.

3.4.5.4. Operational and Interoperability Impact if Blocked

Discarding packets that employ this EH would break IPsec deployments.

3.4.5.5. Advice

Intermediate systems should permit packets containing an Authentication Header.

3.4.6. Destination Options for IPv6 (Protocol Number=60)

3.4.6.1. Uses

The Destination Options header is used to carry optional information that needs be examined only by a packet's destination node(s).

3.4.6.2. Specification

This EH is specified in [RFC8200]. At the time of this writing, the following options have been specified for this EH:

- Type 0x00: Pad1 [RFC8200]
- Type 0x01: PadN [RFC8200]
- Type 0x04: Tunnel Encapsulation Limit [RFC2473]
- Type 0x4D: (Deprecated)
- Type 0xC9: Home Address [RFC6275]
- Type 0x8A: Endpoint Identification (Deprecated) [draft-ietf-nimrod-eid]
- Type 0x8B: ILNP Nonce [RFC6744]
- Type 0x8C: Line-Identification Option [RFC6788]
- Type 0x1E: RFC3692-style Experiment [RFC4727]
- Type 0x3E: RFC3692-style Experiment [RFC4727]
- Type 0x5E: RFC3692-style Experiment [RFC4727]
- Type 0x7E: RFC3692-style Experiment [RFC4727]
- Type 0x9E: RFC3692-style Experiment [RFC4727]
- Type 0xBE: RFC3692-style Experiment [RFC4727]
- Type 0xDE: RFC3692-style Experiment [RFC4727]
- Type 0xFE: RFC3692-style Experiment [RFC4727]

3.4.6.3. Specific Security Implications

No security implications are known, other than the general implications of IPv6 EHs. For a discussion of possible security implications of specific options specified for the DO header, please see the Section 4.3.

3.4.6.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain a Destination Options header would break protocols that rely on this EH type for conveying information, including protocols such as ILNP [RFC6740] and Mobile IPv6 [RFC6275], and IPv6 tunnels that employ the Tunnel Encapsulation Limit option.

3.4.6.5. Advice

Intermediate systems should permit packets that contain a Destination Options Header.

3.4.7. Mobility Header (Protocol Number=135)

3.4.7.1. Uses

The Mobility Header is an EH used by mobile nodes, correspondent nodes, and home agents in all messaging related to the creation and management of bindings in Mobile IPv6.

3.4.7.2. Specification

This EH is specified in [RFC6275].

3.4.7.3. Specific Security Implications

A thorough security assessment of the security implications of the Mobility Header and related mechanisms can be found in Section 15 of [RFC6275].

3.4.7.4. Operational and Interoperability Impact if Blocked

Discarding packets containing this EH would break Mobile IPv6.

3.4.7.5. Advice

Intermediate systems should permit packets containing this EH.

3.4.8. Host Identity Protocol (Protocol Number=139)

3.4.8.1. Uses

This EH is employed with the Host Identity Protocol (HIP), an experimental protocol that allows consenting hosts to securely establish and maintain shared IP-layer state, allowing separation of the identifier and locator roles of IP addresses, thereby enabling continuity of communications across IP address changes.

3.4.8.2. Specification

This EH is specified in [RFC5201].

3.4.8.3. Specific Security Implications

The security implications of the HIP header are discussed in detail in Section 8 of [RFC6275].

3.4.8.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain the Host Identity Protocol would break HIP deployments.

3.4.8.5. Advice

Intermediate systems should permit packets that contain a Host Identity Protocol EH.

3.4.9. Shim6 Protocol (Protocol Number=140)

3.4.9.1. Uses

This EH is employed by the Shim6 [RFC5533] Protocol.

3.4.9.2. Specification

This EH is specified in [RFC5533].

3.4.9.3. Specific Security Implications

The specific security implications are discussed in detail in Section 16 of [RFC5533].

3.4.9.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this EH will break Shim6.

3.4.9.5. Advice

Intermediate systems should permit packets containing this EH.

3.4.10. Use for experimentation and testing (Protocol Numbers=253 and 254)

3.4.10.1. Uses

These IPv6 EHs are employed for performing RFC3692-Style experiments (see [RFC3692] for details).

3.4.10.2. Specification

These EHs are specified in [RFC3692] and [RFC4727].

3.4.10.3. Specific Security Implications

The security implications of these EHs will depend on their specific use.

3.4.10.4. Operational and Interoperability Impact if Blocked

For obvious reasons, discarding packets that contain these EHs limits the ability to perform legitimate experiments across IPv6 routers.

3.4.10.5. Advice

Operators should determine according to their own circumstances whether to discard packets containing these EHs.

3.5. Advice on the Handling of Packets with Unknown IPv6 Extension Headers

We refer to IPv6 EHs that have not been assigned an Internet Protocol Number by IANA (and marked as such) in [IANA-PROTOCOLS] as "unknown IPv6 extension headers" ("unknown IPv6 EHs").

3.5.1. Uses

New IPv6 EHs may be specified as part of future extensions to the IPv6 protocol.

Since IPv6 EHs and Upper-layer protocols employ the same namespace, it is impossible to tell whether an unknown "Internet Protocol Number" is being employed for an IPv6 EH or an Upper-Layer protocol.

3.5.2. Specification

The processing of unknown IPv6 EHs is specified in [RFC7045].

3.5.3. Specific Security Implications

For obvious reasons, it is impossible to determine specific security implications of unknown IPv6 EHs.

3.5.4. Operational and Interoperability Impact if Blocked

As noted in [RFC7045], discarding unknown IPv6 EHs may slow down the deployment of new IPv6 EHs and transport protocols. The corresponding IANA registry ([IANA-PROTOCOLS]) should be monitored such that filtering rules are updated as new IPv6 EHs are standardized.

We note that since IPv6 EHs and upper-layer protocols share the same numbering space, discarding unknown IPv6 EHs may result in packets encapsulating unknown upper-layer protocols being discarded.

3.5.5. Advice

Operators should determine according to their own circumstances whether to discard packets containing unknown IPv6 EHs.

4. IPv6 Options

4.1. General Discussion

The following subsections describe specific security implications of different IPv6 options, and provide advice regarding filtering packets that contain such options.

4.2. General Security Implications of IPv6 Options

The general security implications of IPv6 options are closely related to those discussed in Section 3.2 for IPv6 EHs. Essentially, packets that contain IPv6 options might need to be processed by an IPv6 router's general-purpose CPU,and hence could present a DDoS risk to that router's general-purpose CPU (and thus to the router itself). For some architectures, a possible mitigation would be to rate-limit the packets that are to be processed by the general-purpose CPU (see e.g. [Cisco-EH]).

4.3. Advice on the Handling of Packets with Specific IPv6 Options

The following subsections contain a description of each of the IPv6 options that have so far been specified, a summary of the security implications of each of such options, a discussion of possible interoperability implications if packets containing such options are discarded, and specific advice regarding whether packets containing these options should be permitted.

4.3.1. Pad1 (Type=0x00)

4.3.1.1. Uses

This option is used when necessary to align subsequent options and to pad out the containing header to a multiple of 8 octets in length.

4.3.1.2. Specification

This option is specified in [RFC8200].

4.3.1.3. Specific Security Implications

None.

4.3.1.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option would potentially break any protocol that relies on IPv6 options.

4.3.1.5. Advice

Intermediate systems should not discard packets based on the presence of this option.

4.3.2. PadN (Type=0x01)

4.3.2.1. Uses

This option is used when necessary to align subsequent options and to pad out the containing header to a multiple of 8 octets in length.

4.3.2.2. Specification

This option is specified in [RFC8200].

4.3.2.3. Specific Security Implications

Because of the possible size of this option, it could be leveraged as a large-bandwidth covert channel.

4.3.2.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option would potentially break any protocol that relies on IPv6 options.

4.3.2.5. Advice

Intermediate systems should not discard IPv6 packets based on the presence of this option.

4.3.3. Jumbo Payload (Type=0XC2)

4.3.3.1. Uses

The Jumbo payload option provides the means of specifying payloads larger than 65535 bytes.

4.3.3.2. Specification

This option is specified in [RFC2675].

4.3.3.3. Specific Security Implications

There are no specific issues arising from this option, except for improper validity checks of the option and associated packet lengths.

4.3.3.4. Operational and Interoperability Impact if Blocked

Discarding packets based on the presence of this option will cause IPv6 jumbograms to be discarded.

4.3.3.5. Advice

Intermediate systems should discard packets that contain this option. An operator should permit this option only in specific scenarios in which support for IPv6 jumbograms is desired.

4.3.4. RPL Option (Type=0x63)

4.3.4.1. Uses

The RPL Option provides a mechanism to include routing information with each datagram that an RPL router forwards.

4.3.4.2. Specification

This option was originally specified in [RFC6553]. It has been deprecated by [I-D.ietf-roll-useofrplinfo].

4.3.4.3. Specific Security Implications

Those described in [RFC6553].

4.3.4.4. Operational and Interoperability Impact if Blocked

This option is meant to be employed within an RPL instance. As a result, discarding packets based on the presence of this option (e.g. at an ISP) will not result in interoperability implications.

4.3.4.5. Advice

Non-RPL routers should discard packets that contain an RPL option.

4.3.5. RPL Option (Type=0x23)

4.3.5.1. Uses

The RPL Option provides a mechanism to include routing information with each datagram that an RPL router forwards.

4.3.5.2. Specification

This option is specified in [I-D.ietf-roll-useofrplinfo].

4.3.5.3. Specific Security Implications

Those described in [I-D.ietf-roll-useofrplinfo].

4.3.5.4. Operational and Interoperability Impact if Blocked

This option is meant to survive outside of an RPL instance. As a result, discarding packets based on the presence of this option would break some use cases for RPL (see [I-D.ietf-roll-useofrplinfo]).

4.3.5.5. Advice

Intermediate systems should not discard IPv6 packets based on the presence of this option.

4.3.6. Tunnel Encapsulation Limit (Type=0x04)

4.3.6.1. Uses

The Tunnel Encapsulation Limit option can be employed to specify how many further levels of nesting the packet is permitted to undergo.

4.3.6.2. Specification

This option is specified in [RFC2473].

4.3.6.3. Specific Security Implications

Those described in [RFC2473].

4.3.6.4. Operational and Interoperability Impact if Blocked

Discarding packets based on the presence of this option could result in tunnel traffic being discarded.

4.3.6.5. Advice

Intermediate systems should not discard packets based on the presence of this option.

4.3.7. Router Alert (Type=0x05)

4.3.7.1. Uses

The Router Alert option [RFC2711] is typically employed for the RSVP protocol [RFC2205] and the MLD protocol [RFC2710].

4.3.7.2. Specification

This option is specified in [RFC2711].

4.3.7.3. Specific Security Implications

Since this option causes the contents of the packet to be inspected by the handling device, this option could be leveraged for performing DoS attacks.

4.3.7.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option would break RSVP and multicast deployments.

4.3.7.5. Advice

Packets containing this option should be permitted in environments where support for RSVP, multicast routing, or similar protocols is desired.

4.3.8. Quick-Start (Type=0x26)

4.3.8.1. Uses

This IP Option is used in the specification of Quick-Start for TCP and IP, which is an experimental mechanism that allows transport protocols, in cooperation with routers, to determine an allowed sending rate at the start and, at times, in the middle of a data transfer (e.g., after an idle period) [RFC4782].

4.3.8.2. Specification

This option is specified in [RFC4782], on the "Experimental" track.

4.3.8.3. Specific Security Implications

Section 9.6 of [RFC4782] notes that Quick-Start is vulnerable to two kinds of attacks:

- attacks to increase the routers' processing and state load, and,
- attacks with bogus Quick-Start Requests to temporarily tie up available Quick-Start bandwidth, preventing routers from approving Quick-Start Requests from other connections.

We note that if routers in a given environment do not implement and enable the Quick-Start mechanism, only the general security implications of IP options (discussed in Section 4.2) would apply.

4.3.8.4. Operational and Interoperability Impact if Blocked

The Quick-Start functionality would be disabled, and additional delays in TCP's connection establishment (for example) could be introduced. (Please see Section 4.7.2 of [RFC4782].) We note, however, that Quick-Start has been proposed as a mechanism that could be of use in controlled environments, and not as a mechanism that would be intended or appropriate for ubiquitous deployment in the global Internet [RFC4782].

4.3.8.5. Advice

Intermediate systems should not discard IPv6 packets based on the presence of this option.

4.3.9. CALIPSO (Type=0x07)

4.3.9.1. Uses

This option is used for encoding explicit packet Sensitivity Labels on IPv6 packets. It is intended for use only within Multi-Level Secure (MLS) networking environments that are both trusted and trustworthy.

4.3.9.2. Specification

This option is specified in [RFC5570].

4.3.9.3. Specific Security Implications

Presence of this option in a packet does not by itself create any specific new threat. Packets with this option ought not normally be seen on the global public Internet.

4.3.9.4. Operational and Interoperability Impact if Blocked

If packets with this option are discarded or if the option is stripped from the packet during transmission from source to destination, then the packet itself is likely to be discarded by the receiver because it is not properly labeled. In some cases, the receiver might receive the packet but associate an incorrect sensitivity label with the received data from the packet whose CALIPSO was stripped by an intermediate router or firewall. Associating an incorrect sensitivity label can cause the received information either to be handled as more sensitive than it really is ("upgrading") or as less sensitive than it really is ("downgrading"), either of which is problematic.

4.3.9.5. Advice

Recommendations for handling the CALIPSO option depend on the deployment environment, rather than whether an intermediate system happens to be deployed as a transit device (e.g., IPv6 transit router).

Explicit configuration is the only method via which an intermediate system can know whether or not that particular intermediate system has been deployed within a Multi-Level Secure (MLS) environment. In many cases, ordinary commercial intermediate systems (e.g., IPv6 routers and firewalls) are the majority of the deployed intermediate systems inside an MLS network environment.

For Intermediate systems that DO NOT implement RFC-5570, there should be a configuration option to EITHER (a) drop packets containing the CALIPSO option OR (b) to ignore the presence of the CALIPSO option and forward the packets normally. In non-MLS environments, such intermediate systems should have this configuration option set to (a) above. In MLS environments, such intermediate systems should have this option set to (b) above. The default setting for this configuration option should be set to (a) above, because MLS environments are much less common than non-MLS environments.

For Intermediate systems that DO implement RFC-5570, there should be configuration options (a) and (b) from the preceding paragraph and also a third configuration option (c) to process packets containing a CALIPSO option as per RFC-5570. When deployed in non-MLS environments, such intermediate systems should have this configuration option set to (a) above. When deployed in MLS environments, such intermediate systems should have this set to (c). The default setting for this configuration option MAY be set to (a) above, because MLS environments are much less common than non-MLS environments.

4.3.10. SMF_DPD (Type=0x08)

4.3.10.1. Uses

This option is employed in the (experimental) Simplified Multicast Forwarding (SMF) for unique packet identification for IPv6 I-DPD, and as a mechanism to guarantee non-collision of hash values for different packets when H-DPD is used.

4.3.10.2. Specification

This option is specified in [RFC6621].

4.3.10.3. Specific Security Implications

None. The use of transient numeric identifiers is subject to the security and privacy considerations discussed in [I-D.irtf-pearg-numeric-ids-generation].

4.3.10.4. Operational and Interoperability Impact if Blocked

Dropping packets containing this option within a MANET domain would break SMF. However, dropping such packets at the border of such domain would have no negative impact.

4.3.10.5. Advice

Intermediate systems that are not within a MANET domain should discard packets that contain this option.

4.3.11. Home Address (Type=0xC9)

4.3.11.1. Uses

The Home Address option is used by a Mobile IPv6 node while away from home, to inform the recipient of the mobile node's home address.

4.3.11.2. Specification

This option is specified in [RFC6275].

4.3.11.3. Specific Security Implications

No (known) additional security implications than those described in [RFC6275].

4.3.11.4. Operational and Interoperability Impact if Blocked

Discarding IPv6 packets based on the presence of this option will break Mobile IPv6.

4.3.11.5. Advice

Intermediate systems should not discard IPv6 packets based on the presence of this option.

4.3.12. Endpoint Identification (Type=0x8A)

4.3.12.1. Uses

The Endpoint Identification option was meant to be used with the Nimrod routing architecture [NIMROD-DOC], but has never seen widespread deployment.

4.3.12.2. Specification

This option is specified in [NIMROD-DOC].

4.3.12.3. Specific Security Implications

Undetermined.

4.3.12.4. Operational and Interoperability Impact if Blocked

None.

4.3.12.5. Advice

Intermediate systems should discard packets that contain this option.

4.3.13. ILNP Nonce (Type=0x8B)

4.3.13.1. Uses

This option is employed by Identifier-Locator Network Protocol for IPv6 (ILNPv6) for providing protection against off-path attacks for packets when ILNPv6 is in use, and as a signal during initial network-layer session creation that ILNPv6 is proposed for use with this network-layer session, rather than classic IPv6.

4.3.13.2. Specification

This option is specified in [RFC6744].

4.3.13.3. Specific Security Implications

Those described in [RFC6744].

4.3.13.4. Operational and Interoperability Impact if Blocked

Discarding packets that contain this option will break INLPv6 deployments.

4.3.13.5. Advice

Intermediate systems should not discard packets based on the presence of this option.

4.3.14. Line-Identification Option (Type=0x8C)

4.3.14.1. Uses

This option is used by an Edge Router to identify the subscriber premises in scenarios where several subscriber premises may be logically connected to the same interface of an Edge Router.

4.3.14.2. Specification

This option is specified in [RFC6788].

4.3.14.3. Specific Security Implications

Those described in [RFC6788].

4.3.14.4. Operational and Interoperability Impact if Blocked

Since this option is meant to be employed in Router Solicitation messages, discarding packets based on the presence of this option at intermediate systems will result in no interoperability implications.

4.3.14.5. Advice

Intermediate devices should discard packets that contain this option.

4.3.15. Deprecated (Type=0x4D)

4.3.15.1. Uses

No information has been found about this option type.

4.3.15.2. Specification

No information has been found about this option type.

4.3.15.3. Specific Security Implications

No information has been found about this option type, and hence it has been impossible to perform the corresponding security assessment.

4.3.15.4. Operational and Interoperability Impact if Blocked

Unknown.

4.3.15.5. Advice

Intermediate systems should discard packets that contain this option.

4.3.16. MPL Option (Type=0x6D)

4.3.16.1. Uses

This option is used with the Multicast Protocol for Low power and Lossy Networks (MPL), that provides IPv6 multicast forwarding in constrained networks.

4.3.16.2. Specification

This option is specified in [RFC7731], and is meant to be included only in Hop-by-Hop Option headers.

4.3.16.3. Specific Security Implications

Those described in [RFC7731].

4.3.16.4. Operational and Interoperability Impact if Blocked

Dropping packets that contain an MPL option within an MPL network would break the Multicast Protocol for Low power and Lossy Networks (MPL). However, dropping such packets at the border of such networks will have no negative impact.

4.3.16.5. Advice

Intermediate systems should not discard packets based on the presence of this option. However, since this option has been specified for the Hop-by-Hop Options, such systems should consider the discussion in Section 3.4.1.

4.3.17. IP_DFF (Type=0xEE)

4.3.17.1. Uses

This option is employed with the (Experimental) Depth-First Forwarding (DFF) in Unreliable Networks.

4.3.17.2. Specification

This option is specified in [RFC6971].

4.3.17.3. Specific Security Implications

Those specified in [RFC6971].

4.3.17.4. Operational and Interoperability Impact if Blocked

Dropping packets containing this option within a routing domain that is running DFF would break DFF. However, dropping such packets at the border of such domains will have no security implications.

4.3.17.5. Advice

Intermediate systems that do not operate within a routing domain that is running DFF should discard packets containing this option.

4.3.18. RFC3692-style Experiment (Types = 0x1E, 0x3E, 0x5E, 0x7E, 0x9E, 0xBE, 0xDE, 0xFE)

4.3.18.1. Uses

These options can be employed for performing RFC3692-style experiments. It is only appropriate to use these values in explicitly configured experiments; they must not be shipped as defaults in implementations.

4.3.18.2. Specification

Specified in RFC 4727 [RFC4727] in the context of RFC3692-style experiments.

4.3.18.3. Specific Security Implications

The specific security implications will depend on the specific use of these options.

4.3.18.4. Operational and Interoperability Impact if Blocked

For obvious reasons, discarding packets that contain these options limits the ability to perform legitimate experiments across IPv6 routers.

4.3.18.5. Advice

Operators should determine according to their own circumstances whether to discard packets containing these IPv6 options.

4.4. Advice on the handling of Packets with Unknown IPv6 Options

We refer to IPv6 options that have not been assigned an IPv6 option type in the corresponding registry ([IANA-IPV6-PARAM]) as "unknown IPv6 options".

4.4.1. Uses

New IPv6 options may be specified as part of future protocol work.

4.4.2. Specification

The processing of unknown IPv6 options is specified in [RFC8200].

4.4.3. Specific Security Implications

For obvious reasons, it is impossible to determine specific security implications of unknown IPv6 options.

4.4.4. Operational and Interoperability Impact if Blocked

Discarding unknown IPv6 options may slow down the deployment of new IPv6 options. As noted in [draft-gont-6man-ipv6-opt-transmit], the corresponding IANA registry ([IANA-IPV6-PARAM] should be monitored such that IPv6 option filtering rules are updated as new IPv6 options are standardized.

4.4.5. Advice

Operators should determine according to their own circumstances whether to discard packets containing unknown IPv6 options.

5. IANA Considerations

This document has no actions for IANA.

6. Privacy Considerations

There are no privacy considerations associated with this document.

7. Security Considerations

This document provides advice on the filtering of IPv6 packets that contain IPv6 EHs (and possibly IPv6 options) at IPv6 transit routers. It is meant to improve the current situation of widespread dropping of such IPv6 packets in those cases where the drops result from improper configuration defaults, or inappropriate advice in this area.

8. Acknowledgements

The authors would like to thank Ron Bonica for his work on earlier versions of this document.

The authors of this document would like to thank (in alphabetical order) Mikael Abrahamsson, Brian Carpenter, Darren Dukes, David Farmer, Mike Heard, Bob Hinden, Christian Huitema, Jen Linkova, Carlos Pignataro, Maria Ines Robles, Donald Smith, Pascal Thubert, Ole Troan, Gunter Van De Velde, and Eric Vyncke, for providing valuable comments on earlier versions of this document.

This document borrows some text and analysis from [RFC7126], authored by Fernando Gont, Randall Atkinson, and Carlos Pignataro.

The authors would like to thank Eric Vyncke for his guidance during the publication process of this document.

Fernando would also like to thank Brian Carpenter and Ran Atkinson who, over the years, have answered many questions and provided valuable comments that have benefited his protocol-related work (including the present document).

9. References

9.1. Normative References

[I-D.ietf-roll-useofrplinfo]
Robles, I., Richardson, M., and P. Thubert, "Using RPI Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane", draft-ietf-roll-useofrplinfo-44 (work in progress), January 2021.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/info/rfc1034>.

[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>.

[RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <https://www.rfc-editor.org/info/rfc2205>.

[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>.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <https://www.rfc-editor.org/info/rfc2473>.

[RFC2675] Borman, D., Deering, S., and R. Hinden, "IPv6 Jumbograms", RFC 2675, DOI 10.17487/RFC2675, August 1999, <https://www.rfc-editor.org/info/rfc2675>.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, DOI 10.17487/RFC2710, October 1999, <https://www.rfc-editor.org/info/rfc2710>.

[RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", RFC 2711, DOI 10.17487/RFC2711, October 1999, <https://www.rfc-editor.org/info/rfc2711>.

[RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, DOI 10.17487/RFC3692, January 2004, <https://www.rfc-editor.org/info/rfc3692>.

[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>.

[RFC4727] Fenner, B., "Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers", RFC 4727, DOI 10.17487/RFC4727, November 2006, <https://www.rfc-editor.org/info/rfc4727>.

[RFC4782] Floyd, S., Allman, M., Jain, A., and P. Sarolahti, "Quick-Start for TCP and IP", RFC 4782, DOI 10.17487/RFC4782, January 2007, <https://www.rfc-editor.org/info/rfc4782>.

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

[RFC5201] Moskowitz, R., Nikander, P., Jokela, P., Ed., and T. Henderson, "Host Identity Protocol", RFC 5201, DOI 10.17487/RFC5201, April 2008, <https://www.rfc-editor.org/info/rfc5201>.

[RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, DOI 10.17487/RFC5533, June 2009, <https://www.rfc-editor.org/info/rfc5533>.

[RFC5570] StJohns, M., Atkinson, R., and G. Thomas, "Common Architecture Label IPv6 Security Option (CALIPSO)", RFC 5570, DOI 10.17487/RFC5570, July 2009, <https://www.rfc-editor.org/info/rfc5570>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <https://www.rfc-editor.org/info/rfc6275>.

[RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <https://www.rfc-editor.org/info/rfc6398>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/info/rfc6550>.

[RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <https://www.rfc-editor.org/info/rfc6553>.

[RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6554, DOI 10.17487/RFC6554, March 2012, <https://www.rfc-editor.org/info/rfc6554>.

[RFC6621] Macker, J., Ed., "Simplified Multicast Forwarding", RFC 6621, DOI 10.17487/RFC6621, May 2012, <https://www.rfc-editor.org/info/rfc6621>.

[RFC6740] Atkinson, RJ. and SN. Bhatti, "Identifier-Locator Network Protocol (ILNP) Architectural Description", RFC 6740, DOI 10.17487/RFC6740, November 2012, <https://www.rfc-editor.org/info/rfc6740>.

[RFC6744] Atkinson, RJ. and SN. Bhatti, "IPv6 Nonce Destination Option for the Identifier-Locator Network Protocol for IPv6 (ILNPv6)", RFC 6744, DOI 10.17487/RFC6744, November 2012, <https://www.rfc-editor.org/info/rfc6744>.

[RFC6788] Krishnan, S., Kavanagh, A., Varga, B., Ooghe, S., and E. Nordmark, "The Line-Identification Option", RFC 6788, DOI 10.17487/RFC6788, November 2012, <https://www.rfc-editor.org/info/rfc6788>.

[RFC6971] Herberg, U., Ed., Cardenas, A., Iwao, T., Dow, M., and S. Cespedes, "Depth-First Forwarding (DFF) in Unreliable Networks", RFC 6971, DOI 10.17487/RFC6971, June 2013, <https://www.rfc-editor.org/info/rfc6971>.

[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>.

[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>.

[RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <https://www.rfc-editor.org/info/rfc7731>.

[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>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

[RFC8900] Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O., and F. Gont, "IP Fragmentation Considered Fragile", BCP 230, RFC 8900, DOI 10.17487/RFC8900, September 2020, <https://www.rfc-editor.org/info/rfc8900>.

9.2. Informative References

[Biondi2007]
Biondi, P. and A. Ebalard, "IPv6 Routing Header Security", CanSecWest 2007 Security Conference, 2007, <http://www.secdev.org/conf/IPv6_RH_security-csw07.pdf>.

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

[draft-gont-6man-ipv6-opt-transmit]
Gont, F., Liu, W., and R. Bonica, "Transmission and Processing of IPv6 Options", IETF Internet Draft, work in progress, August 2014.

[draft-ietf-nimrod-eid]
Lynn, C., "Endpoint Identifier Destination Option", IETF Internet Draft, draft-ietf-nimrod-eid-00.txt, November 1995.

[FW-Benchmark]
Zack, E., "Firewall Security Assessment and Benchmarking IPv6 Firewall Load Tests", IPv6 Hackers Meeting #1, Berlin, Germany. June 30, 2013, <http://www.ipv6hackers.org/meetings/ipv6-hackers-1/zack-ipv6hackers1-firewall-security-assessment-and-benchmarking.pdf>.

[I-D.ietf-6man-hbh-header-handling]
Baker, F. and R. Bonica, "IPv6 Hop-by-Hop Options Extension Header", draft-ietf-6man-hbh-header-handling-03 (work in progress), March 2016.

[I-D.ietf-v6ops-ipv6-ehs-packet-drops]
Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston, G., and W. LIU, "Operational Implications of IPv6 Packets with Extension Headers", draft-ietf-v6ops-ipv6-ehs-packet-drops-03 (work in progress), January 2021.

[I-D.irtf-pearg-numeric-ids-generation]
Gont, F. and I. Arce, "On the Generation of Transient Numeric Identifiers", draft-irtf-pearg-numeric-ids-generation-06 (work in progress), January 2021.

[IANA-IPV6-PARAM]
Internet Assigned Numbers Authority, "Internet Protocol Version 6 (IPv6) Parameters", December 2013, <http://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml>.

[IANA-PROTOCOLS]
Internet Assigned Numbers Authority, "Protocol Numbers", 2014, <http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml>.

[NIMROD-DOC]
Nimrod Documentation Page, "http://ana-3.lcs.mit.edu/~jnc/nimrod/".

[RFC3871] Jones, G., Ed., "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure", RFC 3871, DOI 10.17487/RFC3871, September 2004, <https://www.rfc-editor.org/info/rfc3871>.

[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>.

[RFC7126] Gont, F., Atkinson, R., and C. Pignataro, "Recommendations on Filtering of IPv4 Packets Containing IPv4 Options", BCP 186, RFC 7126, DOI 10.17487/RFC7126, February 2014, <https://www.rfc-editor.org/info/rfc7126>.

[RFC7739] Gont, F., "Security Implications of Predictable Fragment Identification Values", RFC 7739, DOI 10.17487/RFC7739, February 2016, <https://www.rfc-editor.org/info/rfc7739>.

[RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <https://www.rfc-editor.org/info/rfc7872>.

Authors' Addresses

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

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

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

Email: liushucheng@huawei.com

OPSEC K. Chittimaneni
Internet-Draft Dropbox Inc.
Intended status: Informational M. Kaeo
Expires: January 8, 2017 Double Shot Security
E. Vyncke
Cisco
July 7, 2016

Operational Security Considerations for IPv6 Networks
draft-ietf-opsec-v6-09

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 all places of a network (enterprises, service providers and residential users) and proposes technical and procedural mitigations techniques.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 8, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

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]. Also, there are subtle differences between NAT44 [RFC2993] and NPTv6 [RFC6296] which are explicitly pointed out in the latter's security considerations section.

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", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] when they appear in ALL CAPS. These words may also appear in this document in lower case as plain English words, absent their normative meanings.

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 good IP Addresses Management (IPAM) system.

Once an address allocation has been assigned, there should be some thought given to an overall address allocation plan. With the abundance of address space available, an address allocation 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 PI vs 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 due to malicious criminal activity.

2.1.1. 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 operational deployment where some portions of the IPv6 address could be easily recognizable for operational debugging and troubleshooting versus the risk of scanning; [SCANNING] shows that there are scientifically based mechanisms that make scanning for IPv6 reachable nodes more realizable than expected; see also [RFC7707]. The use of common multicast groups which are defined for important networked devices and the use of commonly repeated addresses could make it easy to figure out which devices are name servers, routers or other critical devices.

While in some environments the security is so poor that obfuscating addresses is considered a benefit; it is a better practice to ensure that perimeter rules are actively checked and enforced and that statically configured addresses follow some logical allocation scheme for ease of operation.

2.1.2. Use of ULAs

ULAs are intended for scenarios where IP addresses will not have global scope so they should not appear in the global BGP routing table. The implicit expectation from the RFC is that all ULAs will be randomly created as /48s. Any use of ULAs that are not created as a /48 violates RFC4193 [RFC4193].

ULAs could be useful for infrastructure hiding as described in RFC4864 [RFC4864]; Alternatively Link-Local addresses RFC7404 [RFC7404] could also be used. Although ULAs are supposed to be used in conjunction with global addresses for hosts that desire external connectivity, a few operators chose to use ULAs in conjunction with some sort of address translation at the border in order to maintain a perception of parity between their IPv4 and IPv6 setup. Some operators believe that stateful IPv6 Network Address and Port Translation (NAPT) provides some security not provided by NPTv6 (the authors of this document do not share this point of view). The use of stateful IPv6 NAPT would be problematic in trying to track specific machines that may source malware although this is less of an issue if appropriate logging is done which includes utilizing accurate timestamps and logging a node's source ports RFC6302 [RFC6302]. Another typical argument in favor of ULA is that there are too many mistakes made with ACL filters at the edge and the use of ULAs could make things easier to hide internal machines.

The use of ULA does not isolate 'by magic' the part of the network using ULA from other parts of the network (including the Internet). Although section 4.1 of RFC4193 [RFC4193] explicitly states "If BGP is being used at the site border with an ISP, the default BGP configuration must filter out any Local IPv6 address prefixes, both incoming and outgoing.", the operational reality is that this guideline is not always followed. As written, RFC4193 makes no changes to default routing behavior of exterior protocols. Therefore, routers will happily forward packets whose source or destination address is ULA as long as they have a route to the destination and there is no ACL blocking those packets. This means that using ULA does not prevent route and packet filters having to be implemented and monitored. This also means that all Internet transit networks should consider ULA as source or destination as bogons packets and drop them.

It is important to carefully weigh the benefits of using ULAs versus utilizing a section of the global allocation and creating a more effective filtering strategy. It is also important to note that the IETF does not recommend the use of ULA and NPTv6.

2.1.3. Point-to-Point Links

RFC6164 [RFC6164] recommends the use of /127 for inter-router point-to-point links. A /127 prevents the ping-pong attack between routers. However, it should be noted that at the time of this writing, there are still many networks out there that follow the advice provided by RFC3627 [RFC3627] (obsoleted and marked Historic by RFC6547 [RFC6547]) and therefore continue to use /64's and/or /112's. We recommend that the guidance provided by RFC6164 be followed.

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 RFC7404 [RFC7404].

2.1.4. Temporary Addresses - Privacy Extensions for SLAAC

Normal stateless address autoconfiguration (SLAAC) relies on the automatically generated EUI-64 address, which together with the /64 prefix makes up the global unique IPv6 address. The EUI-64 address is generated from the MAC address. 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 reduce the attack exposure window.

As privacy extension addresses could also be used to obfuscate some malevolent activities (whether on purpose or not), it is advised in scenarios where user attribution is important to rely on a layer-2 authentication mechanism such as IEEE 802.1X [IEEE-802.1X] with the appropriate RADIUS accounting (Section 2.6.1.6) or to disable SLAAC and rely only on DHCPv6. However, in scenarios where anonymity is a strong desire (protecting user privacy is more important than user attribution), privacy extension addresses should be used.

Using privacy extension addresses prevents the operator from building a priori host specific access control lists (ACLs). It must be noted that recent versions of Windows do not use the MAC address anymore to build the stable address but use a mechanism similar to the one described in [RFC7217], this also means that such an ACL cannot be configured based solely on the MAC address of the nodes, diminishing the value of such ACL. On the other hand, different VLANs are often used to segregate users, in this case ACL can rely on a /64 prefix per VLAN rather than a per host ACL entry.

The decision to utilize privacy extension addresses can come down to whether the network is managed versus unmanaged. In some environments full visibility into the network is required at all times which requires that all traffic be attributable to where it is sourced or where it is destined to within a specific network. This situation is dependent on what level of logging is performed. If logging considerations include utilizing accurate timestamps and logging a node's source ports [RFC6302] then there should always exist appropriate user attribution needed to get to the source of any malware originator or source of criminal activity.

Disabling SLAAC and privacy extensions addresses can be done by sending Router Advertisement with a hint to use DHCPv6 by setting the M-bit but also disabling SLAAC by resetting all A-bits in all prefixe information options sent in the Router Advertisement message.

2.1.5. Privacy consideration of Addresses

However, there are several privacy issues still present with [RFC4941] such as host tracking, and address scanning attacks are still possible. More details are provided in Appendix A. of [RFC7217] and in [RFC7721].

2.1.6. DHCP/DNS Considerations

Many environments use DHCPv6 to allocate addresses 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 2.3.2).

DNS is often used for malware activities and while there are no fundamental differences with IPv4 and IPv6 security concerns, there are specific consideration in DNS64 RFC6147 [RFC6147] environments that need to be understood. Specifically the interactions and potential to interference with DNSsec implementation need to be understood - these are pointed out in detail in Section 2.7.3.2.

2.2. Extension Headers

TBD, a short section referring to all Fernando's I-D & RFC.

2.3. Link-Layer Security

IPv6 relies heavily on the Neighbor Discovery protocol (NDP) RFC4861 [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 [RFC3756] and in RFC6583 [RFC6583].

2.3.1. SeND and CGA

SEcure Neighbor Discovery (SeND), as described in RFC3971 [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 [RFC3972], are used to ensure that the sender of a Neighbor Discovery message is the actual "owner" of the claimed IPv6 address. A new NDP option, the CGA option, was introduced and is used to carry the public key and associated parameters. Another NDP option, the RSA Signature option, is used to protect all messages relating to neighbor and Router discovery.

SeND protects against:

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

SeND does NOT:

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

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

2.3.2. Securing DHCP

Dynamic Host Configuration Protocol for IPv6 (DHCPv6), as detailed in RFC3315 [RFC3315], 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 any large network by providing robust stateful autoconfiguration and autoregistration of DNS Host Names.

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 mount a man in the middle attack. While unintentionall, a misconfigured DHCP server can have the same impact. Additional threats against DHCP are discussed in the security considerations section of RFC3315 [RFC3315]DHCP-shield

RFC7610 [RFC7610] specifies a mechanism for protecting connected DHCPv6 clients against rogue DHCPv6 servers. This mechanism is based on DHCPv6 packet-filtering at the layer-2 device; the administrator specifies the interfaces connected to DHCPv6 servers.

It is recommended to use DHCP-shield.

2.3.3. 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 [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, route filtering, longer than /64 prefix; These require static configuration of the addresses.

- o Tuning of NDP process (where supported).

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 wifi 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:

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

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

2.3.4. ND/RA 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 [RFC6104] summarizes the scenarios in which rogue RAs may be observed and presents a list of possible solutions to the problem. RFC6105 [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 [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 [RFC6980] updates RFC4861 [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 [RFC7513] would help in creating bindings between a DHCPv4 RFC2131 [RFC2131] /DHCPv6 RFC3315 [RFC3315] assigned source IP address and a binding anchor RFC7039 [RFC7039] on a SAVI device. Also, RFC6620 [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 be be employed as a first line of defense against common attack vectors including misconfigured hosts.

2.3.5. 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 [RFC6459] for a more detailed discussion on the 3GPP link model, NDP on it and the address configuration detail.

2.4. Control Plane Security

RFC6192 [RFC6192] defines the router control plane. This definition is repeated here for the reader's convenience.

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:

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

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

This section will consider several classes of control packets:

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

- Management protocols: SSH, SNMP, IPfix, etc

- 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 extension 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 power of the Route Processor.

2.4.2. Management Protocols

This class includes: SSH, SNMP, syslog, NTP, etc

An ingress ACL to be applied on all the router interfaces 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 extension header (see also [I-D.ietf-6man-hbh-header-handling]);

- 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 [RFC6980] and more generally RFC7112 [RFC7112] highlights the security implications of oversized extension header chains on routers and updates RFC2460 [RFC2460] such that the first fragment of a packet is required to contain the entire IPv6 header chain.

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 will cause Path MTU holes. But, there is no other solution.

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.

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.

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 [RFC5340] or RFC2740 [RFC2740] that was obsoleted by the former. However, the document which specifically describes how IPsec should be implemented for OSPFv3 RFC4552 [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 [RFC7166] (which obsoletes RFC6506 [RFC6506] changes OSPFv3's reliance on IPsec by appending an authentication trailer to the end of the OSPFv3 packets. This document does not specifically provide for a mechanism that will authenticate the specific originator of a packet. Rather, it will allow 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.

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 RFC6434 [RFC6434] is now a SHOULD and not MUST implement. Theoretically it is possible, and recommended, that communication between two IPv6 nodes, including 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, as described in the earlier section. Additionally, in a protocol such as OSPFv3 where adjacencies are formed on a one-to-many basis, IPsec key management becomes difficult to maintain and is not often utilized.

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

- o Discard packets from and to bogon and reserved space

- 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 RFC6810 [RFC6810]

Some good recommendations for filtering can be found from Team CYMRU at [CYMRU].

2.6. Logging/Monitoring

In order to perform forensic research in case of any security incident or to detect abnormal behaviors, network operator 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 the Neighbor cache;

- o use of stateful DHCPv6 [RFC3315] lease cache, especially when a relay agent [RFC6221] in layer-2 switches is used;

- o use of RADIUS [RFC2866] for accounting records.

Please note that there are privacy issues related to how those logs are collected, kept and safely discarded. Operators are urged to check their country legislation.

All those pieces of information will be used for:

- o forensic (Section 2.6.2.1) research to answer questions 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, 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.

RFC 5952 [RFC5952] explains this problem in detail and recommends the use of a single canonical format (in short use lower case and suppress leading 0). This memo 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 this memo recommends the use an external program in order to canonicalize all IPv6 addresses.

For example, this perl script can be used:

```
#!/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" ;
}
```

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. It is usually available by two means:

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

- o also by connecting over a secure management channel (such as SSH or HTTPS) and explicitely requesting a neighbor cache dump.

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 every 30 seconds (to be on the safe side) 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.

2.6.1.5. Stateful DHCPv6 Lease

In some networks, IPv6 addresses are managed by stateful DHCPv6 server [RFC3315] that leases IPv6 addresses 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 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.

In short, the DHCPv6 lease file is less interesting than in the IPv4 era. DHCPv6 servers that keeps the relayed data-link layer address in addition to the DUID in the lease file do not suffer from this limitation. On a managed network where all hosts support DHCPv6, special care must be taken to prevent stateless autoconfiguration anyway (and if applicable) by sending RA with all announced prefixes without the A-bit set.

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.

2.6.2.1. Forensic

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.

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 are used to reach this prefix;

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 the DHCP lease file;

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.

At the end of the process, the interface where the malicious user was connected or the username that was used by the malicious user is found.

2.6.2.2. Inventory

RFC 7707 [RFC7707] (which obsoletes RFC 5157 [RFC5157]) is about the difficulties to scan an IPv6 network due to the vast number of IPv6 addresses per link. This has the side effect of making 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 enumerating the DNS zones, parsing log files, leveraging service discovery such as mDNS RFC6762 [RFC6762] and RFC6763 [RFC6763].

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.

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

Some text

2.7.1. Dual Stack

Dual stack has established itself as the preferred deployment choice for most network operators without an MPLS core where 6PE RFC4798 [RFC4798] is quite common. Dual stacking the network offers many advantages over other transition mechanisms. Firstly, it is easy to turn on without impacting normal IPv4 operations. Secondly, perhaps more importantly, it is easier to troubleshoot when things break. Dual stack allows you to gradually turn IPv4 operations down when your IPv6 network is ready for prime time.

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:

- ACLs to permit or deny traffic
- Firewalls with stateful packet inspection

It is recommended that these ACLs and/or firewalls be additionally configured to protect IPv6 communications. 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

2.7.2. Transition 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 being described in RFC 6169 [RFC6169]);

- 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;

- traffic interception: no confidentiality is provided by the tunnel protocols (without the use of IPsec), therefore anybody on the tunnel path can intercept the traffic and have access to the clear-text IPv6 packet;

- 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);

- 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.

- bypassing security policy: if a firewall or an IPS is on the path of the tunnel, then it will probably neither inspect not detect an malevolent IPv6 traffic contained in the tunnel.

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

- IP protocol 41: this will block ISATAP (Section 2.7.2.2), 6to4 (Section 2.7.2.4), 6rd (Section 2.7.2.5) as well as 6in4 (Section 2.7.2.1) tunnels;

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

- UDP protocol 3544: this will block the default encapsulation of Teredo (Section 2.7.2.3) tunnels.

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 ad 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 means that endpoints and and the tunnel endpoint 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 [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. 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).

2.7.2.4. 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 [RFC3068] which has been deprecated by RFC7526 [RFC7526], and is no more used by recent Operating Systems. Some security considerations are explained in RFC3694 [RFC3964].

RFC6343 [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.

2.7.2.5. 6rd

While 6rd tunnels share the same encapsulation as 6to4 tunnels (Section 2.7.2.4), 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 [RFC5969] describes how to secure the 6rd tunnels.

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

2.7.2.6. 6PE and 6VPE

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

- o Address space, routing and traffic seperation with the help of VRF (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.

2.7.2.7. DS-Lite

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

2.7.2.8. Mapping of Address and Port

With the tunnel and encapsulation 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 MUST 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 manager.

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.

RFC7422 [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

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 [RFC6145] 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 incidence on DNSSEC see section 3.1 of [RFC7050].

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 AFTR 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 agains 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.

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]
- o Accept certain ICMPv6 messages to allow proper operation of ND and PMTUD, see also [RFC4890]
- o Filter specific extension headers, where possible
- o Filter unneeded services at the perimeter
- o Implement anti-spoofing
- 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.

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:

- Authenticating the TCP session;
- TTL security (which becomes hop-limit security in IPv6);
- Prefix Filtering.

These are explained in more detail in section Section 2.5.

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 [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] (which obsoletes [RFC6204]) defines the requirements of an IPv6 CPE and does not take position on the debate of default IPv6 security policy as defined in [RFC6092]:

- 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;

- 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 ([I-D.ietf-v6ops-balanced-ipv6-security]: open to all outbound and inbound connections at the exception of an 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 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, Brian Carpenter, Tim Chown, Markus deBruen, Fernando Gont, Jeffry Handal, Panos Kampanakis, Erik Kline, Jouni Korhonen, Mark Lentczner, Bob Sleigh,Tarko Tikan (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, <http://www.rfc-editor.org/info/rfc2119>.

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

[RFC6104] Chown, T. and S. Venaas, "Rogue IPv6 Router Advertisement Problem Statement", RFC 6104, DOI 10.17487/RFC6104, February 2011, <http://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, <http://www.rfc-editor.org/info/rfc6105>.

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>.

[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-6man-hbh-header-handling]
Baker, F. and R. Bonica, "IPv6 Hop-by-Hop Options Extension Header", draft-ietf-6man-hbh-header-handling-03 (work in progress), March 2016.

[I-D.ietf-v6ops-balanced-ipv6-security]
Gysi, M., Leclanche, G., Vyncke, E., and R. Anfinsen, "Balanced Security for IPv6 Residential CPE", draft-ietf-v6ops-balanced-ipv6-security-01 (work in progress), December 2013.

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

[NAv6TF_Security]
Kaeo, , Green, , Bound, , and 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, , Graveman, , Pearce, , and Rooks, "Guidelines for the Secure Deployment of IPv6", 2010, <http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf>.

[RFC0826] Plummer, D., "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, <http://www.rfc-editor.org/info/rfc826>.

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

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

[RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, DOI 10.17487/RFC2740, December 1999, <http://www.rfc-editor.org/info/rfc2740>.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <http://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, <http://www.rfc-editor.org/info/rfc2827>.

[RFC2866] Rigney, C., "RADIUS Accounting", RFC 2866, DOI 10.17487/RFC2866, June 2000, <http://www.rfc-editor.org/info/rfc2866>.

[RFC2993] Hain, T., "Architectural Implications of NAT", RFC 2993, DOI 10.17487/RFC2993, November 2000, <http://www.rfc-editor.org/info/rfc2993>.

[RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, DOI 10.17487/RFC3056, February 2001, <http://www.rfc-editor.org/info/rfc3056>.

[RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, DOI 10.17487/RFC3068, June 2001, <http://www.rfc-editor.org/info/rfc3068>.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <http://www.rfc-editor.org/info/rfc3315>.

[RFC3627] Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful", RFC 3627, DOI 10.17487/RFC3627, September 2003, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc3924>.

[RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, DOI 10.17487/RFC3964, December 2004, <http://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, <http://www.rfc-editor.org/info/rfc3971>.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc4301>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <http://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, <http://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, <http://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", RFC 4443, DOI 10.17487/RFC4443, March 2006, <http://www.rfc-editor.org/info/rfc4443>.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <http://www.rfc-editor.org/info/rfc4552>.

[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, <http://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, <http://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, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc4942>.

[RFC5157] Chown, T., "IPv6 Implications for Network Scanning", RFC 5157, DOI 10.17487/RFC5157, March 2008, <http://www.rfc-editor.org/info/rfc5157>.

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, DOI 10.17487/RFC5214, March 2008, <http://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, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc6092>.

[RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, DOI 10.17487/RFC6144, April 2011, <http://www.rfc-editor.org/info/rfc6144>.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, DOI 10.17487/RFC6145, April 2011, <http://www.rfc-editor.org/info/rfc6145>.

[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, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc6192>.

[RFC6204] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, Ed., "Basic Requirements for IPv6 Customer Edge Routers", RFC 6204, DOI 10.17487/RFC6204, April 2011, <http://www.rfc-editor.org/info/rfc6204>.

[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, <http://www.rfc-editor.org/info/rfc6221>.

[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, <http://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, <http://www.rfc-editor.org/info/rfc6269>.

[RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, DOI 10.17487/RFC6296, June 2011, <http://www.rfc-editor.org/info/rfc6296>.

[RFC6302] Durand, A., Gashinsky, I., Lee, D., and S. Sheppard, "Logging Recommendations for Internet-Facing Servers", BCP 162, RFC 6302, DOI 10.17487/RFC6302, June 2011, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc6333>.

[RFC6343] Carpenter, B., "Advisory Guidelines for 6to4 Deployment", RFC 6343, DOI 10.17487/RFC6343, August 2011, <http://www.rfc-editor.org/info/rfc6343>.

[RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, DOI 10.17487/RFC6434, December 2011, <http://www.rfc-editor.org/info/rfc6434>.

[RFC6459] Korhonen, J., Ed., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", RFC 6459, DOI 10.17487/RFC6459, January 2012, <http://www.rfc-editor.org/info/rfc6459>.

[RFC6506] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 6506, DOI 10.17487/RFC6506, February 2012, <http://www.rfc-editor.org/info/rfc6506>.

[RFC6547] George, W., "RFC 3627 to Historic Status", RFC 6547, DOI 10.17487/RFC6547, February 2012, <http://www.rfc-editor.org/info/rfc6547>.

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, DOI 10.17487/RFC6583, March 2012, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc6666>.

[RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <http://www.rfc-editor.org/info/rfc6762>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, DOI 10.17487/RFC6763, February 2013, <http://www.rfc-editor.org/info/rfc6763>.

[RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", RFC 6810, DOI 10.17487/RFC6810, January 2013, <http://www.rfc-editor.org/info/rfc6810>.

[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, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc7039>.

[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, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc7166>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <http://www.rfc-editor.org/info/rfc7217>.

[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, <http://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, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc7526>.

[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, <http://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, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc7721>.

[SCANNING] "Mapping the Great Void - Smarter scanning for IPv6", <http://www.caida.org/workshops/isma/1202/slides/aims1202_rbarnes.pdf>.

Authors' Addresses

Kiran K. Chittimaneni
Dropbox Inc.
185 Berry Street, Suite 400
San Francisco, CA 94107
USA

Email: kk@dropbox.com

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

Phone: +12066696394
Email: merike@doubleshotsecurity.com

Eric Vyncke
Cisco
De Kleetlaan 6a
Diegem 1831
Belgium

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