IPv6 Security

Attacks and Countermeasures in a Nutshell

Johanna Ullrich, Katharina Krombholz, Heidelinde Hobel, Adrian Dabrowski, Edgar Weippl

The history of computers is full of underestimation: 640 kilobyte, 2-digit years, and 32-bit Internet addresses. IPv6 was invented to overcome the latter as well as to revise other drawbacks and security vulnerabilities of its predecessor IPv4. Initially considered the savior in terms of security because of its mandatory IPsec support, it turned out not to be the panacea it was thought to be. Outsourcing security to IPsec but eventually removing it as well as other design decisions led to a number of vulnerabilities. They range from the already known spoofing of answers to link-layer address requests to novel possibilities regarding node tracking. In an effort to fix them, a vast amount of updates have been introduced.

In this paper, we discuss security and privacy vulnerabilities with regard to IPv6 and their current countermeasures. In a second step, vulnerabilities and countermeasures are systematized by the appliance of an extendible common language for computer security incidents. Our evaluation shows that a large part of vulnerabilities can be mitigated but several security challenges remain. We deduce three main research challenges for IPv6 security, namely address assignment and structure, securing local network discovery, and address selection for reconnaissance.

This is a reprint of the authors' article published in the 8th USENIX Workshop on Offensive Technologies (WOOT), 2014.

1 Introduction

The Internet Protocol (IP) is the principal communication protocol of the Internet. Its fast expansion led to a shortage of IPv4 addresses and triggered the current transformation process to the revised version IPv6 with an address range of 2^{128}. Even though the new version was updated multiple times, the basic security and privacy design was made in 1998. However, a full deployment in the 2010s means distinct security vulnerabilities. In 2011, the Internet Assigning Number Authority (IANA) distributed its last IPv4 addresses to the Regional Internet Registries *IPv4 Address Report* (no date), and some of them have already run out of addresses. This way, the prolonged transformation to IPv6 gains momentum.

In the narrower sense, IPv6 is only a new transport layer header. However, this is accompanied by a long list of upgrades and revisions of related technologies, which were closely tied to IPv4. This includes new entry types for the Domain Name System (DNS), the Internet Control Message Protocol (ICMP) version 6 or a redefinition of the pseudo-header for checksum calculation. As a result, some known IPv4 vulnerabilities are not relevant for IPv6, while other flaws still remain. Certainly, the enhancement of functionalities implies new security vulnerabilities.

For the successful worldwide adoption of IPv6, security and privacy aspects in the protocol suite have been examined thoroughly in recent years. The results have been published in various scientific papers, *Requests for Comments* (RFCs), videos and blogs. It is, therefore, a time-consuming and tedious task to collect all the findings and to obtain a comprehensive understanding of this topic. In addition to scientific work, we included non-scientific contributions from hacker blogs to complete our systematization with security challenges that were detected in the wild. The overall goal of this paper is to summarize and systematize the IPv6 vulnerabilities as well as the associated countermeasures in a nutshell. In the following, we assemble IPv6 vulnerabilities and evaluate appropriate countermeasures to provide a complete and comprehensive checklist for researchers, developers and administrators. Furthermore, we deduce major future research challenges, namely address assignment and structure, securing local network discovery, and address selection for reconnaissance.

The remainder of this paper is structured as follows: Section 2 introduces IPv6 and related technologies. Section 3 summarizes currently known security vulnerabilities, while Section 4 considers privacy in relation to IPv6. Section 5 presents excerpts of the systematization, providing two tables describing vulnerabilities/countermeasures according to a common language to describe computer security incidents and a matrix showing their adequacy. Finally, Section 6 discusses the main research challenge related to IPv6,

Table 1: IPv6 Header Format (Deering and Hinden 1998)

Size in Bits	Field Name	Comment
4	Version	set to *6*
8	Traffic Class	replaces *Type of Services*
20	Flow Label	for packet flow marking
16	Payload Length	incl. *IPv6 Extension Headers*
8	Next Header	
8	Hop Limit	replaces *Time to Live*
128	Source Address	
128	Destination Address	

Section 7 compares IPv4 to IPv6 in a number of aspects, and Section 8 concludes this work.

2 Background on IPv6

In comparison to IPv4, its successor IPv6 encompasses four major modifications: (1) The address length has been quadrupled to 128 bit, providing $3.4 \cdot 10^{38}$ unique addresses. These contain a subnet prefix and an interface identifier, and are represented by 8 quadruples of hexadecimal values separated by colons (Hinden and Deering 2006). (2) Regarding the amount of receivers, three types of addresses are distinguished: *unicast*, *anycast* and *multicast* addresses. There are no *broadcast* addresses in IPv6. (3) The header format has been simplified and fixed to 40 byte, as shown in Table 1. Fragmentation and other optional functionality has been shifted to optional *extension headers*, which are inserted between the IP and the upper-layer protocol header. (4) Fragmentation has further been limited to end nodes with the objective of router offloading. (5) Formerly mandatory IPsec (Kent 2005a,b; Kent and Seo 2005) is seen as its fifth major modification before being released as optional (Jankiewicz et al. 2011).

With IP being the Internet's main protocol, many constitutive Internet technologies are heavily tied to it and the change to version 6 resulted in updates of related protocols. One of them is the *Internet Control Message Protocol* (ICMPv6) (Conta, Deering and Gupta 2006). In spite of a reduced number of message types, its scope has increased beyond error and diagnostic messages. Performing now also address resolution by means of the *Neighbor Discovery Protocol* (NDP) (Narten, Nordmark et al. 2007), it is also the successor of the *Address Resolution Protocol* (ARP) and responsible for router discovery.

IPv6 addresses are either configured manually, statefully (such as by *Dynamic Host Configuration Protocol* (DHCPv6) (Droms et al. 2003)[1]), or by the newly introduced *Stateless Autoconfiguration* (SLAAC) (Thomson and Narten 1998; Thomson, Narten and Jinmei 2007), providing plug-and-play connectivity. With SLAAC,

---

1   This is a reprint of the authors' article published in the 8th USENIX Workshop on Offensive Technologies (WOOT), 2014.

1   The stateless DHCP approach is technically speaking not a means of address assignment because it does not maintain a client state (Droms 2004).

the host first creates a link-local address on its own. After receiving a *router advertisement*, the node generates global addresses with the announced network prefixes. Recommended network prefix sizes for end sites are between /48 and /64 (IAB and IESG 2001; Narten, Huston et al. 2011).

Due to the increasing number of mobile nodes, mobility support (Perkins et al. 2011) has gained importance. It allows nodes to remain transparently reachable via the same address while wandering through the network. In case the mobile node is in a foreign network, it provides its actual address to its router by means of a binding update. This provides two possibilities for correspondent nodes to communicate with the mobile node: The communication can be passed on to the home agent, which tunnels the traffic on to the mobile node. Alternatively, route optimization allows direct communication without the home agent by using a certain routing header.

The transformation from version 4 to 6 takes time and is accompanied by a phase of co-existence. Some nodes are capable of both protocols, while others are limited to one or the other. Therefore, transition technologies that bridge this gap have been developed, which can be divided into two main types: (1) Tunneling delivers a packet as another packet's payload. (Nordmark and Gilligan 2005) provides a general description on tunneling IPv6 over IPv4, while (Haas and Hares 2006) is a specification for tunneling other protocols over IPv6. Currently, there are a high number of different technologies tunneling IPv6 over IPv4: *6to4* (Carpenter and K. Moore 2001; Huitema 2001), *IPv6 rapid deployment* (Despres 2010; Townsley and Troan 2010), *6over4* (Carpenter and Jung 1999; Wu et al. 2010), *ISATAP* (Templin et al. 2008) and Teredo (Huitema 2006; D. Thaler 2011). (2) Alternatively, protocol translation, i. e., the translation of IPv4 into IPv6 headers and vice versa, can be used. Due to being tightly connected, IP translation also includes ICMP translation. The first specification *Network Address Translation - Protocol Translation (NAT-PT)* has been criticized by (Aoun and Davies 2007; Tsirtsis and Srisuresh 2000) for numerous reasons, e.g. lacking support of DNSSEC. Its successor is standardized in (Bagnulo, Matthews et al. 2011; Bagnulo, Sullivan et al. 2011; Baker et al. 2011; Bao et al. 2010; Li et al. 2011). However, tunneling is currently preferred.

3 Security Vulnerabilities

In the course of the development of the new Internet Protocol version, changes in and supplements to functionality were made. These enhancements, however, yield different behavior and therefore often result in novel security vulnerabilities. In this section, we summarize fundamental security vulnerabilities in IPv6 and present feasible countermeasures. We organize them by intended functionality, starting with extension headers, fragmentation and other native header fields. Subsequently, *Neighbor* and *Multicast Listener*

Discovery are discussed, followed by tunneling and mobility support.

3.1 Extension Headers

extension headers provide optional functionality and are inserted before the next-layer protocol header. Two of them are of further interest for security: (1) The *routing header type 0* holds a list of addresses that have to be visited en route to the receiver. By alternating the two addresses, the packet cycles between two nodes, causing traffic amplification on a remote path and possibly resulting in denial of service (Abley et al. 2007). This extension header was more harmful than beneficial and was finally deprecated (Abley et al. 2007)[2].

Offloading routers was a major focus during development. *IPv6 extension headers* are, therefore, only allowed to be processed at the end nodes. The only exception is the *Hop-by-Hop header* and its *Router Alert option*, which may be used for updating in the future. However, this option may also cause a decrease in router performance when many packets are sent (Partridge and Jackson 1999).

Initially, extension headers and options did not have to follow a certain format, therefore, middleboxes are not necessarily able to process new extension headers. Later, a uniform format for extension headers was standardized (Krishnan et al. 2012).

3.2 Fragmentation

IPv6 did not explicitly prohibit the reassembly of overlapping fragments initially despite this being a well-known security threat that can be used, e. g., to evade firewalls (Davies, Krishnan et al. 2007). The best-known way of doing so is overwriting the TCP SYN flag. The countermeasure in IPv4 was dropping fragments with an offset of one byte (Ziemba et al. 1995). But this is no appropriate mitigation for IPv6 because an arbitrary number of *extension headers* can be inserted prior to the next-layer protocol header and cause any offset.

Such insertions are also able to shift flags or port numbers to succeeding fragments. Common firewalls collect incoming packet fragments and reassemble them in any case, but reassembly implementations differ, making IPv6 vulnerable to the same attack scenarios as IPv4 (Miller 2001; Ziemba et al. 1995). These differences in reassembly can also be used to fingerprint operating systems (Atlasis 2012).

As a consequence, overlapping fragments are now explicitly forbidden because benign nodes do not have any need of sending overlaps (Krishnan 2009). Further, deep packet inspection should treat initial fragments without flags or port numbers with suspicion as there is a guaranteed MTU in IPv6. Finally, fragmentation is still a stateful process within a stateless

---

2   *Routing header type 0* differs from the benign *type 2* (Johnson et al. 2004) used for mobile applications.

protocol with the risk of memory overflow.

Specific to IPv6 are *atomic fragments*. These packets consist of only one fragment and are used in protocol translation to deliver an identifier for fragmentation in IPv4 (F. Gont 2013). Unfortunately, these fragments can cause dropping of benign fragments that have the same identifier. Thus, the two types of fragments should be handled in isolation from each other.

3.3 Mandatory IPv6 Header Fields

Similar to the *Router Alert option*, a high number of different *flow labels* is able to decrease router performance because the latter has to store a state for every label value. A malicious attacker can also gain access someone else's quality of service by using the same *flow label* (Amante et al. 2011).

3.4 Neighbor Discovery

Neighbor discovery has many security implications due to its philosophy of trusting everybody on the local network. Assuming an attacker has managed to reach the local network, they can perform a variety of malicious actions.

Address Resolution   Spoofing attacks that provide wrong link-layer addresses are still possible (Figure 1a). Attackers are further able to prevent victims from address assignment by answering to *duplicate neighbor detection*. One applied countermeasure is *Optimistic Duplicate Address Detection*. Here, the node assumes that its address is unique in any case (N. Moore 2006).

Router Advertisement Spoofing   Any node on the local network is able to announce itself as a router (see Figure 1b), or spoof a router's announcement. A number of variations of this attack are known: (1) Setting the router's lifetime to zero kicks the reminder from the client's configuration. (2) Announcing an arbitrary prefix lets the clients assume this prefix is local (Nikander et al. 2004; Yang et al. 2007).

(3) Flooding the network with *router advertisements* with various prefixes causes clients to configure one address per announcement and may lead to denial of service. These problems are not fully solved by using DHCP, as the attacker can force the node to abandon DHCP. As a countermeasure, the router advertisement guard – a middlebox filtering illegitimate announcements – is proposed (Chown and Venaas 2011; Levy-Abegnoli et al. 2011).

Advertisements may also be sent unintentionally due to misconfiguration. Preferences of benign announcements should therefore be high to guarantee service even in such a case (Draves and D. Thaler 2005).

Redirects   An attacker may redirect traffic by sending *redirects* and change the sender's configuration this way.

Smurf Attacks   An attacker sends a request to a multicast address, spoofing the victim's source address. Responses are returned to the victim, causing a denial of service. Adequate request types are *echo requests* or IP packets with an unknown extension header option of type 10. *echo requests* to multicast addresses must not be answered, but some implementations do. In contrast, the alternative containing an unknown option has to be answered (F. Gont 2011). Considering the latter, non-answering has been proposed (Fernando Gont and Liue 2013), but even in case of becoming a standard, an exception remains for *Packet too Big* messages for path MTU discovery.

General security mechanisms tackling all vulnerabilities together have been targeted. With IPsec being initially mandatory, neighbor discovery seemed adequately secure, but it suffered from bootstrapping problems. Securing it would require manual key exchange, and therefore, unacceptable effort. As a consequence, *Secure Neighbor Discovery* (SeND) was introduced (Arkko, Kempf et al. 2005). With this technology, cryptographically generated addresses enable the association of addresses to a public key (Aura 2005), and signing messages with the private key prevents spoofing. However, RSA is calculation intensive and the overhead makes the systems more prone to denial-of-service attacks. Even more limiting is the low support. For example, there is only one proof-of-concept implementation for Microsoft operating systems (Rafiee et al. 2011). Therefore, the only option remains to prevent attackers from joining the local network through physical protection or link-layer access control.

3.5 Multicast Listener Discovery

Multicast Listener Discovery (MLD) is a protocol maintaining information on nodes listening to multicast addresses. This allows the forwarding of packets destined for these addresses. A query router in charge of maintaining this information regularly sends general query messages asking for listening nodes. The latter answer with report messages. A malicious node can abort this forwarding of multicast-destined packets by sending a spoofed done message. The effect, however, would last only until the next general query message that is answered by the victim, initializing forwarding again.

Thus, the attacker has to attempt to itself become the query router. The query router is determined by having the lowest address. Although routers are frequently assigned ascending addresses, the lowest IPv6 interface identifier `::` (all zeros) is typically unused and addressing starts with `::1` (Heuse no date) – possibly an IPv4 legacy.

After becoming the query router, it stops sending query requests, causing an MLD denial of service. However, the old query router will start querying again if it does not see MLD requests. However, if it sends such queries only to the all-router multicast address, the other routers are satisfied while the nodes

(a) Neighbor Discovery Spoofing   (b) Router Advertisement Misuse   (c) Redirects

Figure 1: Basic Attacks with Neighbor Discovery
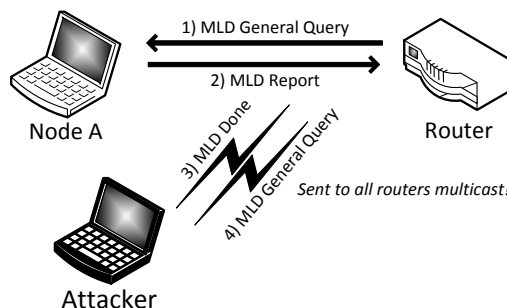

Figure 2: Multicast Listener

face deteriorated service (see Figure 2). Assigning the lowest address :: to the legitimate router is an adequate countermeasure, as explained above.

3.6 Tunneling

At the beginning of IPv6 deployment, tunneling illegitimate content over IPv6 was easy because many firewalls let any IPv6 traffic pass. While this has changed drastically, special threats arise from transition technologies due to the combination of the two IP versions.

Routing loops are an issue of automatic tunneling mechanisms, e. g. *Teredo* or *ISATAP* (G. Nakibly and Templin 2011; Gabi Nakibly and Arov 2009). Starting with a native IPv6 packet with a spoofed source address, this packet is forwarded to a tunnel ingress point. There it is encapsulated into an IPv4 packet and forwarded. At the egress point, the packet is decapsulated and equals the first, which is forwarded again to the ingress point. This causes traffic amplification because the *hop count* is only reduced on native IPv6 routers. Mitigation methods may include the general avoidance of multiple tunnels and border routers, a list of other tunnel routers' addresses to drop their packets, and checking IPv4 and IPv6 addresses for consistency (G. Nakibly and Templin 2011; Savola and Patel 2004).

Special attacks are known for *Teredo*: (1) Cycling is possible between an end node and a cone NAT supporting hair-pin routing. (2) Even endless looping is possible with a bubble request. Originally intended to open another NAT via the server, the request to open the server address causes the server to send bubbles endlessly.

Nested encapsulation means the encapsulation of

tunnel packets in packets of another tunnel, causing additional overhead through another packet header or even fragmentation. To counter this, a *Tunnel Encapsulation Limit option* limiting the number of nested tunnels has been introduced (Conta and Deering 1998).

3.7 Address Space Size

The massive expansion of address space returns vulnerabilities known from the Internet's early days. Simplistic implementations of neighbor discovery may hold too many still unanswered neighbor address requests caused by network scanning. To mitigate this denial of service, filtering unused address space and minimal subnet sizing is proposed (Gashinsky et al. 2012). There is even discussion of minimizing subnets down to e. g. a /124, but then it is likely that implementations fail due to assuming minimum subnetworks of /64.

Point-to-point links encounter the threat of ping-pong packets in case a router forwards a packet back over the incoming interface and causes packet cycling. As above, taking smaller subnets, e. g. /127 would mitigate the risk (Kohno et al. 2011). Alternatively, the latest ICMPv6 specification (Conta, Deering and Gupta 2006) mitigates this by returning an ICMPv6 *Destination Unreachable* message.

3.8 Mobile IPv6

Binding updates inform the home agent of a mobile node's current address and enable it to stay reachable via its home address. Spoofing binding updates may inform the agent of a wrong address and can be used for man-in-the-middle, hijacking, passive wire

tapping or denial-of-service attacks. In order to prevent these attacks in mobile IPv6 networks, the use of IPsec is recommended (Arkko, Devarapalli et al. 2004).

4 Privacy Issues

Since Internet-based technologies are becoming increasingly pervasive and exhibit a tendency to neglect users' privacy, addressing privacy violations is of utmost importance. In this section, we highlight privacy-related challenges along with state-of-the-art countermeasures.

4.1 Addressing

As stated above, an 128-bit IPv6 address consists of a network prefix and an interface identifier. While the first is given by the network on which the host resides, the interface identifier is independently generated by the host. Initially, the modified EUI-format containing the MAC address was proposed for generation of the interface identifier (Thomson and Narten 1998). Since using a hardware address results in unique identifiers even across different subnets, it is easy to track a node's movement through the network. A draft now even proposes their deprecation (Fernando Gont, Dave Thaler et al. 2013).

Numerous address formats have been proposed as an alternative: (1) The *Privacy Extension* generates an MD5 hash at a regular time interval – typically 24 hours – and uses this as the identifier (Narten, Draves et al. 2007). While this impedes long-term tracking, short-term tracking is still possible as the identifier does not change simultaneously with the prefix. (2) Another alternative frequently proposed is DHCPv6. However, it relies on the static *DHCP Unique Identifier* (DUID). By sniffing DUIDs locally or requesting the respective DHCP servers directly, an attacker is still able to correlate a node with its current address (Groat et al. 2011).

With *Mobile IPv6*, there is a trade-off between keeping track of all sessions during network switching and the privacy breach allowing to be traceable across different networks. By including the home address and the temporary care-of address in one packet, a potential adversary is able to eavesdrop on the communication channel and infer the device's location. This may be prevented by encryption, e. g. IPsec. However, nodes communicating with the mobile device can still track the latter. To prevent such privacy breaches, the care-of address and the home address must also be changed simultaneously (Perkins et al. 2011).

4.2 Reconnaissance

The discovery of unknown nodes is typically the first step in an attack or penetration test, but the sheer size of the address range makes brute-forcing impossible. Thus, more sophisticated methods are necessary: (1)

In 2007, an analysis of IPv6 addresses in the wild showed frequent address structures for the first time (Malone 2008). While servers and routers tend to follow the modified EUI-Format and »low« addresses, clients have a significant portion of addresses generated by the privacy extension. Further analyses are feasible by *address6* (Fernando Gont no date). Results of such analyses have resulted in *scan6* of the same toolkit. This tool searches for low-byte, IPv4-based, port-based or modified EUI addresses.

(2) Another source for addresses is DNS, which will be becoming more popular with IPv6 due to the address length. First, it is possible to query known domains. Second, reverse entries can be exploited at BIND or NDS implementations (van Dijk no date). As the response for an empty non-terminal differs from other error messages, it is possible to infer whether addresses starting with this prefix are known to this server. (3) Beyond DNS, all other sources of addresses are of interest as well, e. g. *Node Information Queries* (Crawford and Haberman 2006), *Inverse Discovery* (Conta 2001) or *whois.net* (Chown 2008).

(4) A modified version of the smurf attack is also capable of reconnaissance. Instead of spoofing the source address, the attacker inserts its own address and receives responses with previously unknown source addresses. However, one has to be aware that a high number of responses may cause a denial of service to oneself (Heuse no date). To prevent revealing individual addresses, servers listening to anycast addresses should also use this anycast address as a source address in the response (Davies, Krishnan et al. 2007).

But inherent features of IPv6 also make reconnaissance easier: (1) The assignment of more than one address to an interface is legitimate, but for reconnaissance it is sufficient to discover one. (2) Addresses expire after a preferred lifetime, but are still used for an existing connection for some time (Thomson, Narten and Jinmei 2007). (3) Clients using the privacy extension further own a stable address that can be assigned randomly or following the modified EUI format (Fernando Gont, Dave Thaler et al. 2013). (4) ICMP must not be totally filtered with IPv6. Even further, filtering *echo requests* and *responses* is said to be less important due to the alleged possible risk from scans (Davies and Mohacsi 2007). An overview on this topic is also given by (Chown 2008).

4.3 Covert Channels

Covert channels are communication channels violating system policies. In total, 22 possible covert channels have been found in the IPv6 header and its extensions (Lucena et al. 2006). The most well known covert channels are the *flow label* with 20 bit (Amante et al. 2011) and the *traffic class* with 8 bit, as their use is still vaguely defined. While the latter is allowed to be changed en route, the modification of the *flow label* was previously prohibited (Rajahalme et al. 2004). This, however, has changed: resetting is al-

lowed in case a covert channel imposes a serious risk (Amante et al. 2011). Another covert channel of 64 bit is provided by the interface identifiers. As the privacy extension causes frequently changing random addresses, it is highly unlikely that these secret messages are detected (Lindqvist 2006).

5 Systematization of Knowledge

Systematization means arranging something so as to present the content more clearly. Section 3 and Section 4 explained security and privacy vulnerabilities as well as countermeasures for IPv6 verbally. This section presents them so that they can be taken at a glance and serve as a checklist for researchers and practitioners alike. With the more in-depth verbal description in the previous sections and this systematic overview, this paper presents the subject in multiple ways, allowing it to be used as a reference guide.

The methodology has to fulfil two goals: (1) a clear arrangement and (2) a brief description of the attacks. In Section 5.1, an appropriate approach is presented. Section 5.2 contains the systematization for vulnerabilities, Section 5.3 for countermeasures and Section 5.4 shows the adequacy of countermeasures to vulnerabilities.

5.1 Methodology

(Howard 1990) developed an extendible common language for describing computer security incidents. According to this work, "an attack is a series of steps taken by an attacker to achieve an unauthorized result". It consists of a tool for exploitation, a vulnerability describing a system weakness, an event – a directed action intended to change the state of a system – and an unauthorized result. The event consists of an action performed by the attacker on a certain target. We adapted this common language to the purpose of describing IPv6 security and privacy vulnerabilities and the respective countermeasures. The original common language did not offer a description for countermeasures, but we believe describing them as a sequence of steps as well is adequate.

5.2 Systematization of Vulnerabilities

The vulnerabilities have been systematized by means of six attributes: (1) action, (2) object, (3) target, (4) unauthorized result, (5) origin, and (6) type.

The action describes the activity of the attacker and is further specified by the object and the target. The object describes the entity the action is performed on. The target defines the victim node. If the latter attribute is left free, all types of nodes are likely to be attacked. While object and target are not enumerated, a limited number of values exist for action. The following list defines them in accordance with place holders for object and target in brackets:
- *assign:* set the address for [target] to [object]

- *flood:* emit a high number of [object] to [target]
- *insert:* include [object] into [target]
- *listen:* eavesdrop on the traffic for [object]
- *scan:* iterate through the addresses of [target]
- *send:* emit a packet including [object] to [target]
- *spoof:* emit [object] to [target] pretending to be another node

The unauthorized result describes the aftermath of the malicious action. Further, the origin of a vulnerability and a threat type is defined. The attribute vulnerability indicates whether the vulnerability results from a design, implementation or configuration flaw according to the following definitions by Howard (1990):
- *configuration:* »a vulnerability resulting from an error in the configuration of a system«
- *design:* »a vulnerability inherent in the design or specification of hardware or software whereby even a perfect implementation will result in a vulnerability«
- *implementation:* »a vulnerability resulting from an error made in the software or hardware implementation of a satisfactory design«

The threat type is also limited to three values following the definitions by (C. P. Pfleeger and S. L. Pfleeger 2003):
- *interception:* »some unauthorized party has gained access to an asset«
- *interruption:* »an asset of the system becomes lost, unavailable, or unusable«
- *modification:* »an unauthorized party not only accesses but tampers with an asset«

The resulting systematization for the above described vulnerabilities is found in Table 2 and Table 3.

5.3 Systematization of Countermeasures

Countermeasures are described by the two attributes action and object, which have the same purpose as for vulnerabilities. However, the list of actions changes to the following:
- *assign:* set [object]
- *disable:* deactivate [object]
- *encrypt:* encode [object] to be secured against reading and/or tampering
- *filter[3]:* remove [object] when passing
- *isolate:* process [object] separately
- *limit:* define maximal value for [object]
- *log:* write message about [object]
- *minimize:* reduce number of [object] as much as possible
- *prohibit:* ban [object]
- *respond:* return with [object]

Object is not enumerated. The countermeasures are further classified into three groups of activity levels: (1) *detective* countermeasures discover a present attack, (2) *preventative* countermeasures are taken before an attack takes place, and (3) *reactive* countermeasures

---

3   Discarding has been included in filtering as it can also be understood as removing messages.

Table 2: Classification of Security Vulnerabilities

	ID	Vulnerability	Action	Object	Target	Unauthorized Result	Origin	Type
Security	v01	Fragmentation Header I	send	overlapping fragments		modified header fields	design	modification
	v02	Fragmentation Header II	send	port number in second fragment		middlebox evasion	design	interception
	v03	Fragmentation Header III	flood	fragments		memory shortage	design	interruption
	v04	Fragmentation Header IV	flood	atomic fragments		packet loss	design	interruption
	v05	Routing Header Type 0 I	send	routing header		traffic amplification	design	interruption
	v06	Routing Header Type 0 II	send	routing header		middlebox evasion	design	interception
	v07	Extension Header Options I	send	router alert option		increased workload	design	interruption
	v08	Extension Header Options II	spoof	invalid 10xxxx option	multicast address	multiple responses	design	interruption
	v09	Hop-by-Hop Header	send	hop-by-hop header		increased workload	design	interruption
	v10	New Extension Header	send	unknown extension header		middlebox evasion	design	interception
	v11	New Extension Header	send	unknown extension header		increased workload	design	interruption
	v12	Flow Label I	send	different flow labels		memory shortage	design	interruption
	v13	Flow Label II	send	existing flow label		quality-of-service theft	design	interruption
	v14	Neighbor Advertisement I	spoof	neighbor advertisement		wrongly resolved address	design	interruption
	v15	Neighbor Advertisement II	spoof	neighbor advertisement		traffic redirection	design	modification
	v16	Neighbor Advertisement III	spoof	neighbor advertisement		address assignment prevention	design	interruption
	v17	Router Advertisement I	spoof	router advertisement		new default router	design	modification
	v18	Router Advertisement II	spoof	router advertisement		removed default router	design	modification
	v19	Router Advertisement III	spoof	router advertisement		wrong locally-announced prefix	design	modification
	v20	Router Advertisement IV	flood	router advertisement		multiple address assignment	implementation	interruption
	v21	Router Advertisement V	spoof	router advertisement		prevention of DHCP assignment	design	interruption
	v22	Router Advertisement VI	send	router advertisement		IPv6 activation	implementation	modification
	v23	Redirect I	spoof	redirect		redirected traffic	design	modification
	v24	Redirect II	spoof	redirect		wrong locally-announced node	design	modification
	v25	Echo Request I	spoof	echo request	multicast address	multiple responses	implementation	interruption
	v26	SeND	send	authenticated messages		increased workload	design	interruption
	v27	Tunneling I	send	IPv6 packet as IPv4 payload		middlebox evasion	implementation	interception
	v28	Tunneling II	send	tunnel packet	relay router	cycling packet	implementation	interruption
	v29	Tunneling III	send	tunnel packet		cycling packet	configuration	interruption
	v30	Teredo	send	Teredo bubble	server	cycling packet	design	interruption
	v31	Nesting	insert	packet into packet		packet overhead	configuration	interruption
	v32	Fragmentation Header V	send	packet too big		inclusion of atomic fragments	design	interception
	v33	Neighbor Discovery	scan		subnetwork	memory shortage	implementation	interruption
	v34	Forwarding	send	returning packet		traffic amplification	design	interruption
	v35	Mobile IPv6 I	spoof	binding update	home agent	traffic redirection	design	modification
	v36	Multicast Listener	assign	lowest address	itself	new MDL query router	design	modification

Table 3: Classification of Privacy Vulnerabilities

	ID	Vulnerability	Action	Object	Target	Unauthorized Result	Origin	Type
Privacy	c01	Fragmentation Header VI	send	overlapping fragments		identification	implementation	interception
	c02	Modified EUI Format	scan	interface identifier	networks	tracking	design	interception
	c03	Echo Request II	send	echo request	invalid multicast address	identification of sniffing nodes	implementation	interception
	c04	Mobile IPv6 II	listen	binding update		tracking	design	interception
	c05	DHCP I	listen	DHCP traffic		tracking	design	interception
	c06	DHCP II	send	DHCP information request	DHCP server	tracking	design	interception
	c07	DNS	send	DNS request	DNS server	reconnaissance	design	interception
	c08	Reverse DNS	send	Reverse DNS query		reconnaissance	implementation	interception
	c09	Echo Request III	send	echo request	multicast address	multiple responses	implementation	interception
	c10	Extension Header Options III	send	packet with invalid option	multicast address	multiple responses	design	interception
	c11	Anycast	send		anycast address	response with unicast address	implementation	interception
	c12	Traffic Class	insert	secret information	traffic class field	leaked information	design	interception
	c13	Flow Label	insert	secret information	flow label field	leaked information	design	interception
	c14	Privacy Extension I	insert	secret information	interface identifier	leaked information	design	interception

are triggered by the attack. The resulting systematization is found in Table 4 and 5.

5.4 Vulnerabilities and Appropriate Countermeasures

Table 6 and 7 show the adequacy of countermeasures to vulnerabilities. We created a matrix where each row represents a vulnerability and each column a countermeasure. A checkmark indicates that a countermeasure is adequate. There is no distinction between various levels of mitigation, e. g. total mitigation vs. some improvement of status quo.

The introduction of a certain countermeasure may lead to new vulnerabilities. For example, the use of SeND to prevent RA attacks creates a vulnerability to denial-of-service attacks due to increased calculation efforts. Likewise, the use of the privacy extension prohibits tracking, but makes it possible for the interface identifier to be used as a covert channel. Thus, a method may be a vulnerability and a solution to another vulnerability at the same time. Further, there are vulnerabilities that cannot be mitigated easily by means of the mechanisms presented here, e.g. memory shortage due to fragment flooding.

6 Future Challenges in Research

Large-scale IPv6 deployment is unquestionably a practitioners' task. However, in this case, practice and research live in mutual symbiosis. The practical experience gained from large-scale deployments typically reveals previously unknown security issues that are not easily solved. As such, they are bounced back to research, where in-depth investigation takes place. In this paper, we described IPv6's status quo with the objective of identifying such back-bouncing topics. While many vulnerabilities have already been considered in practice, the results from our systematization suggest that there is a variety of research challenges to be investigated. In this section, we infer these main challenges regarding IPv6 and propose possible approaches for mitigation.

6.1 Addressing

Every proposed addressing solution has a serious drawback: (1) The modified EUI-format is easily traceable by benign administrators as well as attackers using out-of-the-box tools like ping. (2) The usage of DHCP does not mitigate this issue because of the unique and stable DUID, and (3) the privacy extension is highly volatile. Therefore, especially administrators fear its negative impact on logging. (4) Manual address assignment is possible for servers and routers, but not for a large amount of clients. These drawbacks highlight the lack of an adequate address assignment structure for the clients' side.

To strike a balance between full randomness and foolproof tracking, requirements for client addressing have to be defined prior to the development of another approach. From this we deduce that the ability to guess a node's address depends on a person's role: (1) Administrators must be able to correlate addresses belonging to its sub-network to physical hosts. (2) Outsiders must not be able to correlate addresses of the same physical node from different networks. This leads to the conclusion that the administrator must have an advantage in terms of knowledge, e. g., through the creation of a pseudo-random addressing scheme seeded by the administrator.

6.2 Securing the Local Network

Securing ICMPv6 with IPsec has proven to be inadequate due to a bootstrapping problem: IPv6 requires prior setup by means of *router advertisements*, *neighbor solicitations* and *neighbor advertisements*. Securing with IPsec in turn requires a previous key exchange over IP, which is not ready for use at this point.

With this insight, SeND was proposed. But even the toughest solution fails if the acceptance is low and no practical implementation is available. We therefore conclude that although high effort has been put into the development of a general security solution for ICMPv6, there are no advantages over its predecessor IPv4.

As a consequence, protection has to be provided on other layers, e. g., preventing attackers from accessing the local network (physical protection) or link-layer access control. However, physical protection is not feasible in wireless communication and the growth of cloud computing leads to shared local networks among foreign parties. Link-layer access is inappropriate with »bring-your-own-device« policies where the IT department are unable to support the various types of devices. This also applies to decentrally organized organizations like universities.

In such cases, only specific countermeasures such as *router advertisement* guards or throttles remain. The disadvantage is their limited domain and the unknown impacts of combining them. Thus, we strongly encourage researchers to pick up this topic again to develop a more practical general security solution for ICMPv6.

6.3 Reconnaissance

Even though reconnaissance in IPv6 has been considered impossible, various techniques have proven the opposite. Nevertheless, they have some drawbacks: (1) DNS querying reveals mainly servers that are intended to be found anyway. (2) Messages to multicast addresses invoking responses may result in a denial of service and their deprecation is foreseeable. (3) Eavesdropping, i. e., passive listening to network traffic, does not work for outside attackers as it is unlikely that packets originating within the victim's prefix will run into the attacker in an arbitrary location on the Internet.

Table 4: Systematization of Countermeasures

ID	Countermeasure	Action	Object
	Detective		
c01	NDP Mon	log	inconsistent NDP msg.
	Preventative		
c02	Use Anycast Address	respond	with anycast as source address
c03	DHCP	assign	addresses statefully
c04	No Forwarding	prohibit	forwarding over same interface
c05	Fragment Isolation	isolate	atomic from other fragments
c06	IPsec	encrypt	packets
c07	IPsec with Manual Keys	encrypt	packets
c08	No IPv6 Support	disable	IPv6
c09	Format Deprecation	prohibit	modified EUI format
c10	Multicast Listener Address	assign	lowest address to router
c11	No Multiple Edge Routers	disable	other edge routers
c12	No Multiple Tunnels	disable	other tunnels
c13	No Multicast Responses	prohibit	answers to multicast addresses
c14	No Overlapping Fragments	prohibit	overlapping fragments
c15	Packet Rate	limit	packet rate
c16	Physical Protection	prohibit	physical access to network
c17	Privacy Extension	assign	temporary random address
c18	RA Throttler	limit	router advertisements
c19	No RAs	disable	router advertisements
c20	No Routing Header Type 0	prohibit	routing header type 0
c21	Router Preference	assign	highest preference
c22	Segmentation	segment	network
c23	SeND	encrypt	NDP messages
c24	Subnet Size	minimize	subnet size
c25	Temporary DUID	assign	temporary DUID
c26	No Tunneling	disable	all tunnels
c27	Uniform Format	limit	number of ext. header formats

Table 5: Systematization of Countermeasures

ID	Countermeasure	Action	Object
	Reactive		
c28	Address Change	assign	new addresses simultaneously
c29	Address Checks	filter	inconsistent addresses
c30	Change Field en route	assign	default value
c31	Echo Requests	filter	echo requests
c32	Hop-by-Hop Options	filter	hop-by-hop extension header
c33	Routing Header	filter	routing headers
c34	Fragmented Packets	filter	packets with port not in 1st frag.
c35	Invalid Options	filter	options of type '10xxxx'
c36	Link Layer Access Control	filter	unauthorized clients
c37	Message Checks	filter	invalid ICMP msg.
c38	NDP Inspection	filter	inconsistent msg.
c39	RA Guard	filter	invalid router advertisements
c40	RA Filtering	filter	router alert options
c41	Router Listing	filter	msg. from other tunnel routers
c42	Tunnel Enc. Limit	limit	number of nested packets
c43	Tunnel Ingress and Exit	filter	at tunnel end points
c44	Unused Addresses	filter	unused addresses

Table 6: Evaluation of Countermeasures

	NDP Mon	Answer with Anycast Address	DHCP	No Forwarding	Fragment Isolation	IPsec	IPsec with Manual Key Configuration	IPv6 Support	Format Deprecation	Multicast Listener Address	No Multiple Edge Routers	No Multiple Tunnels	No Multicast Responses	No Overlapping Fragments	Packet Rate	Physical Protection	Privacy Extension	RA Throttler	No RAs	No Routing Header Type 0	Router Preference	Segmentation	SeND	Subnet Size	Temporary DUID	No Tunneling	Uniform Format	Address Change	Address Checks	Change Field en route	Echo Requests	Hop-by-Hop Options Header	Fragmented Packet Filtering	Invalid Options	Link Layer Access Control	Message Checks	NDP Inspection	RA Guard	RA Filtering	Router Listing	Tunnel Encapsulation Limit Option	Tunnel Ingress and Exit	Unused Addresses
Fragmentation Header I														✓																													
Fragmentation Header II																																	✓										
Fragmentation Header III																																											
Fragmentation Header IV					✓																																						
Routing Header Type 0 I																				✓																							
Routing Header Type 0 II																				✓																							
Extension Header Options I																																✓								✓			
Extension Header Options II														✓																		✓		✓									
Hop-by-Hop Header																																✓											
New Extension Header																											✓																
New Extension Header																											✓																
Flow Label I																																											
Flow Label II																																											
Neighbor Advertisement I	✓					✓							✓								✓		✓													✓	✓	✓					
Neighbor Advertisement II	✓					✓							✓								✓		✓													✓	✓	✓					
Neighbor Advertisement III	✓					✓							✓								✓		✓													✓	✓	✓					
Router Advertisement I	✓					✓							✓			✓	✓		✓		✓		✓													✓	✓		✓				
Router Advertisement II	✓					✓							✓			✓	✓				✓		✓													✓	✓		✓				
Router Advertisement III	✓					✓							✓			✓	✓				✓		✓													✓	✓		✓				
Router Advertisement IV	✓					✓							✓		✓	✓	✓				✓		✓													✓	✓		✓				
Router Advertisement V	✓					✓							✓			✓	✓				✓		✓													✓	✓		✓				
Router Advertisement VI	✓				✓	✓							✓			✓	✓				✓		✓													✓	✓		✓				
Redirect I	✓												✓								✓		✓													✓	✓						
Redirect II	✓												✓								✓		✓													✓	✓						
Echo Request I															✓				✓												✓												
SeND																			✓																	✓							
Tunneling I											✓	✓														✓																✓	
Tunneling II											✓	✓														✓									✓						✓		
Tunneling III																										✓									✓								
Teredo																										✓																	
Nesting														✓												✓																✓	

Table 7: Evaluation of Countermeasures

	NDP Mon	Answer with Anycast Address	DHCP	No Forwarding	Fragment Isolation	IPsec	IPsec with Manual Key Configuration	IPv6 Support	Format Deprecation	Multicast Listener Address	No Multiple Edge Routers	No Multiple Tunnels	No Multicast Responses	No Overlapping Fragments	Packet Rate	Physical Protection	Privacy Extension	RA Throttler	No RAs	No Routing Header Type 0	Router Preference	Segmentation	SeND	Subnet Size	Temporary DUID	No Tunneling	Uniform Format	Address Change	Address Checks	Change Field en route	Echo Requests	Hop-by-Hop Options Header	Fragmented Packet Filtering	Invalid Options	Link Layer Access Control	Message Checks	NDP Inspection	RA Guard	RA Filtering	Router Listing	Tunnel Encapsulation Limit Option	Tunnel Ingress and Exit	Unused Addresses
Fragmentation Header V																																											
Neighbor Discovery															✓									✓																			✓
Forwarding				✓																				✓																			
Mobile IPv6 I						✓	✓																																				
Multicast Listener										✓																																	
Fragmentation Header VI														✓																													
Modified EUI Format								✓									✓																										
Echo Request II																														✓													
Mobile IPv6 II						✓																								✓													
DHCP I																									✓																		
DHCP II																									✓																		
DNS																																											
Reverse DNS																																											
Echo Request III													✓																	✓													
Extension Header Options III													✓																					✓									
Anycast		✓																																									
Traffic Class																												✓															
Flow Label																												✓															
Privacy Extension I			✓																																								

Considering this, scanning is still the most promising reconnaissance type due to (1) invoking active responses from the victim, (2) revealing the stable address instead of a temporary one, (3) its local as well as global applicability, (4) its independence from certain protocols and (5) the difficulty of mitigating it due to using the inherent functionality of protocols. What seems to be legacy is brute-force scanning, i. e., iterating through all possible addresses – the method of choice in IPv4. In conclusion, research has to find new address selection algorithms for active probing to replace brute-forcing and manage the large amount of IPv6 addresses in this way. We believe that the exploitation of address structures is promising. Research, therefore, requires data sets of IPv6 addresses. Thus, we strongly encourage the collection of such data sets that make it possible to get more in-depth knowledge on assignment in various environments. Nevertheless, reconnaissance will be also dependent on the developments in addressing.

7 Generation Next - Generation Best?

Although IPv6 undoubtedly implies significant privacy and security flaws, it must be noted that neither was its ancestor fully secure, yet still contributed to today's interconnected world. Next generation IP will neither be Internet security's patron nor its tortfeasor. Thus, this chapter describes the idea behind protocol application in (1) IPv4 as we know it from today's Internet, (2) IPv6 as primarily intended before the turn of the millennium, and (3) the current state of IPv6. This allows further deliberation of the extent of security and privacy flaws in different phases of IP.

IPv4 was developed as a packet-switching protocol in 1981. At this time, Internet attacks were rare because the network was an academic network connecting universities with a high number of trusted users. This changed with the Internet's commercialization, providing targets with great financial gain and a changing user group. More central solutions came into existence to tackle corporate needs. Since then, a controversy has existed between the corporate world and academia still aiming at the end-to-end principle.

Initially, IPv6 had been planned to restore the Internet's end-to-end principle, enabling flexibility, decentrality and equality. Measures thereof were the prohibition of fragmentation or other extension headers on intermediate routers, or self-configuration of addresses by SLAAC and the restriction to basic functions in the main protocol header. Additionally, security was valued by the mandatory introduction of IPsec. However, this turned out to be a pitfall presumably caused by limited security knowledge and experience at that time. The decentralized approach was also not fully pervasive as numerous technologies were reused, e.g. DNS. Resolution of domain names even seemed to become a more vital role due to the unwieldy IPv6 addresses.

Like IPv4, the new protocol version experienced an evolution in the past decades based on gained experience as well as a changed environment, e.g. the increased number of mobile nodes. Similar to before, a trend towards centrality becomes apparent. It seems to be driven by corporate administrators who prefer to limit their users in order to achieve manageability, controllability and security. This has led to a reintroduction of various protocols, e.g. DHCP, or the wide acceptance of central middle boxes. The standardization efforts are further an action to anticipate the de-

velopment of various flavors of implementation like experienced with NAT in IPv4.

Considering all these attacks and the failed security approach with IPsec, IPv6 seems less secure and leads to the final question: Is IPv6 in general more or less secure than IPv4? Our results suggest that this protocol is less secure than it could be if the experience with its predecessor had been taken into account. Further, we conclude that IPv6 is not less secure than IPv4: (1) Fragmentation attacks are known for both versions and (2) securing the local network has always been done on lower layers. (3) SeND vulnerabilities will not play a major role due to its lacking acceptance in practice. (4) Attacks aimed at denial-of-service of routers en-route prevent the goal of router offloading, but IPv6 has at least achieved offloading from the performance-intensive task of fragmentation.

Nevertheless, one major issue remains – transition technology which causes roughly 30 percent of the presented security vulnerabilities. Originally, transition was intended as an interim phase of dual-stack nodes natively supporting both protocol versions. However, this process did not gain momentum for a long time -— also due to distrusting IPv6 security, and now the time has passed for this approach leading to tunneling and translating. In conclusion, a number of security flaws have been introduced by fearing IPv6.

8 Conclusion

In this paper, we contextualized security as well as privacy vulnerabilities of IPv6 and evaluated available countermeasures. Then, we systematized the vulnerabilities with respect to the following criteria: *action, object, target, unauthorized results, origin* and *type*. Furthermore, the countermeasures were systematized by *action, object* and *activity level*. The evaluation showed that a countermeasure could be found for the majority of vulnerabilities, which leads to the conclusion that IPv6 is a rather secure protocol. However, some countermeasures create new vulnerabilities. For example, SeND prevents *router advertisement* attacks but increases the risk of denial of service due to increased calculation effort.

Finally, we targeted imperfectly addressed vulnerabilities and identified three major research challenges left with regard to IPv6: (1) addresses providing protection against outside tracking but easy logging for administrators, (2) once more picking up the idea of a general security solution for local network discovery, (3) and the development of an address selection technique that allows reconnaissance through active probing.

About the Authors

Heidelinde Hobel is Researcher at SBA Research and PhD student at the Vienna University of Technology. She completed her master's degree in Business Informatics with distinction. Since 2013, she holds a research scholarship from the Doctoral College of Environmental Informatics (TU Vienna), aimed at utilizing synergies from the fields of informatics, geoinformation, simulation, statistics, visualization, energy research, and architecture. She has gained industry experience as Software Developer in large-scale security and risk management projects. Her research interests include, among others, risk/compliance management, location-based services and knowledge engineering, and the subsequently following privacy issues.

Katharina Krombholz is researcher at SBA Research and PhD student at the Vienna University of Technology in Vienna, Austria. She completed her master's degree in Media Informatics with distinction and received two outstanding student's awards. Her research interests include usable security and privacy and digital forensics. Katharina published more than 13 scientific papers in peer-reviewed conferences and journals. Besides her research activities at SBA Research, she is currently teaching graduate courses on digital forensics and cloud security at Vienna University of Technology and two Austrian Universities of Applied Sciences.

Adrian Dabrowski is a PhD student at TU Wien and employed at SBA Research, Vienna, Austria. His Master's Thesis on RFID locking systems was nominated for the Distinguished Young Alumnus Award of the Faculty of Informatics, the OCG Incentive Award, and won the IEEE Austria Diploma Thesis Award. He participated on the winning International Capture the Flag (iCTF) team in 2006 and 2011 and co-organized the team in 2011-2014. His work on IMSI Catchers won the Best Student Paper Award at ACSAC 2014.

Johanna Ullrich is a PhD student at SBA Research, Vienna, Austria, the national competence center for Information Security. Recipient of three outstanding student awards, she achieved a Bachelor of Electrical Engineering and Information Technology in 2010, and her Master of Automation Engineering in 2013, both with distinction from Vienna University of Technology. For her master thesis on Header Compression of IPsec in Powerline Networks, she received the diploma prize of the city of Vienna. Her research interests include network security and cloud security, and in raising awareness for security and privacy in traditional engineering.

Edgar R. Weippl is Research Director of SBA Research and associate professor at the Vienna University of Technology. His research focuses on applied concepts of IT-security; he organizes the ARES conference, and is on the editorial board of Elsevier's Computers & Security journal (COSE) and chair of SACMAT 2015, Esorics 2015 and CCS 2016.

Acknowledgment

References

Abley, J., Savola, P. & Neville-Neil, G. (2007, December). Deprecation of Type 0 Routing Headers in IPv6. RFC 5095 (Proposed Std). IETF. IETF.

Amante, S., Carpenter, B., Jiang, S. & Rajahalme, J. (2011, November). IPv6 Flow Label Specification. RFC 6437 (Proposed Std). IETF. IETF.

Aoun, C. & Davies, E. (2007, July). Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status. RFC 4966 (Informational). IETF. IETF.

Arkko, J., Devarapalli, V. & Dupont, F. (2004, June). Using IPsec to Protect Mobile IPv6 Signaling Between Mobile Nodes and Home Agents. RFC 3776 (Proposed Std). Updated by RFC 4877. IETF. IETF.

Arkko, J., Kempf, J., Zill, B. & Nikander, P. (2005, March). SEcure Neighbor Discovery (SEND). RFC 3971 (Proposed Std). Updated by RFCs 6494, 6495. IETF. IETF.

Atlasis, A. (2012). Attacking ipv6 implementation using fragmentation. *BlackHat Europe*.

Aura, T. (2005, March). Cryptographically Generated Addresses (CGA). RFC 3972 (Proposed Std). Updated by RFCs 4581, 4982. IETF. IETF.

Bagnulo, M., Matthews, P. & van Beijnum, I. (2011, April). Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers. RFC 6146 (Proposed Std). IETF. IETF.

Bagnulo, M., Sullivan, A., Matthews, P. & van Beijnum, I. (2011, April). DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers. RFC 6147 (Proposed Std). IETF. IETF.

Baker, F., Li, X., Bao, C. & Yin, K. (2011, April). Framework for IPv4/IPv6 Translation. RFC 6144 (Informational). IETF. IETF.

Bao, C., Huitema, C., Bagnulo, M., Boucadair, M. & Li, X. (2010, October). IPv6 Addressing of IPv4/IPv6 Translators. RFC 6052 (Proposed Std). IETF. IETF.

Carpenter, B. & Jung, C. (1999, March). Transmission of IPv6 over IPv4 Domains without Explicit Tunnels. RFC 2529 (Proposed Std). IETF. IETF.

Carpenter, B. & Moore, K. (2001, February). Connection of IPv6 Domains via IPv4 Clouds. RFC 3056 (Proposed Std). IETF. IETF.

Chown, T. (2008, March). IPv6 Implications for Network Scanning. RFC 5157 (Informational). IETF. IETF.

Chown, T. & Venaas, S. (2011, February). Rogue IPv6 Router Advertisement Problem Statement. RFC 6104 (Informational). IETF. IETF.

Conta, A. (2001, June). Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification. RFC 3122 (Proposed Std). IETF. IETF.

Conta, A. & Deering, S. (1998, December). Generic Packet Tunneling in IPv6 Specification. RFC 2473 (Proposed Std). IETF. IETF.

Conta, A., Deering, S. & Gupta, M. (2006, March). Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. RFC 4443 (Draft Std). Updated by RFC 4884. IETF. IETF.

Crawford, M. & Haberman, B. (2006, August). IPv6 Node Information Queries. RFC 4620 (Experimental). IETF. IETF.

Davies, E., Krishnan, S. & Savola, P. (2007, September). IPv6 Transition/Co-existence Security Considerations. RFC 4942 (Informational). IETF. IETF.

Davies, E. & Mohacsi, J. (2007, May). Recommendations for Filtering ICMPv6 Messages in Firewalls. RFC 4890 (Informational). IETF. IETF.

Deering, S. & Hinden, R. (1998, December). Internet Protocol, Version 6 (IPv6) Specification. RFC 2460 (Draft Std). Updated by RFCs 5095, 5722, 5871, 6437, 6564, 6935, 6946. IETF. IETF.

Despres, R. (2010, January). IPv6 Rapid Deployment on IPv4 Infrastructures (6rd). RFC 5569 (Informational). IETF. IETF.

Draves, R. & Thaler, D. [D.]. (2005, November). Default Router Preferences and More-Specific Routes. RFC 4191 (Proposed Std). IETF. IETF.

Droms, R. (2004, April). Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6. RFC 3736 (Proposed Std). IETF. IETF.

Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C. & Carney, M. (2003, July). Dynamic Host Configuration Protocol for IPv6 (DHCPv6). RFC 3315 (Proposed Std). Updated by RFCs 4361, 5494, 6221, 6422, 6644. IETF. IETF.

Gashinsky, I., Jaeggli, J. & Kumari, W. (2012, March). Operational Neighbor Discovery Problems. RFC 6583 (Informational). IETF. IETF.

Gont, F. [F.]. (2011, July). Security Assessment of the Internet Protocol Version 4. RFC 6274 (Informational). IETF. IETF.

Gont, F. [F.]. (2013, May). Processing of IPv6 "Atomic" Fragments. RFC 6946 (Proposed Std). IETF. IETF.

Gont, F. [Fernando]. (nodate). SI6 Networks' IPv6 Toolkit. Retrieved from http://www.si6networks.com

Gont, F. [Fernando] & Liue, W. (2013, March). Security Implications of IPv6 Options of Type 10xxxxxx. IETF. Retrieved from tools.ietf.org/html/draft-gont-6man-ipv6-smurf-amplifier-03

Gont, F. [Fernando], Thaler, D. [Dave] & Liue, W. (2013, October). Deprecating EUI-64 Based IPv6 Addresses. Internet Engineering Task Force. IETF. Retrieved from http://tools.ietf.org/html/draft-gont-6man-deprecate-eui64-based-addresses-00

Groat, S., Dunlop, M., Marchany, R. & Tront, J. (2011). What DHCPv6 says about you. In *Internet Security (WorldCIS), 2011 World Congress on* (Pages 146–151). IEEE.

Haas, J. & Hares, S. (2006, January). Definitions of Managed Objects for BGP-4. RFC 4273 (Proposed Std). IETF. IETF.

Heuse, M. (nodate). THC-IPv6-Attack-Toolkit. Retrieved from http://www.aldeid.com/wiki/THC-IPv6-Attack-Toolkit

Hinden, R. & Deering, S. (2006, February). IP Version 6 Addressing Architecture. RFC 4291 (Draft Std). Updated by RFCs 5952, 6052. IETF. IETF.

Howard, D. (1990). The Influence of verbal responses to common greetings on compliance behavior: The foot-in-the-mouth effect. *Journal of Applied Social Psychology*, *20*.

Huitema, C. (2001, June). An Anycast Prefix for 6to4 Relay Routers. RFC 3068 (Proposed Std). IETF. IETF.

Huitema, C. (2006, February). Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs). RFC 4380 (Proposed Std). Updated by RFCs 5991, 6081. IETF. IETF.

IAB & IESG. (2001, September). IAB/IESG Recommendations on IPv6 Address Allocations to Sites. RFC 3177 (Informational). Obsoleted by RFC 6177. IETF. IETF.

IPv4 Address Report. (nodate). Retrieved from http://www.potaroo.net/tools/ipv4/index.html

Jankiewicz, E., Loughney, J. & Narten, T. (2011, December). IPv6 Node Requirements. RFC 6434 (Informational). IETF. IETF.

Johnson, D., Perkins, C. & Arkko, J. (2004, June). Mobility Support in IPv6. RFC 3775 (Proposed Std). Obsoleted by RFC 6275. IETF. IETF.

Kent, S. (2005a, December). IP Authentication Header. RFC 4302 (Proposed Std). IETF. IETF.

Kent, S. (2005b, December). IP Encapsulating Security Payload (ESP). RFC 4303 (Proposed Std). IETF. IETF.

Kent, S. & Seo, K. (2005, December). Security Architecture for the Internet Protocol. RFC 4301 (Proposed Std). Updated by RFC 6040. IETF. IETF.

Kohno, M., Nitzan, B., Bush, R., Matsuzaki, Y., Colitti, L. & Narten, T. (2011, April). Using 127-Bit IPv6 Prefixes on Inter-Router Links. RFC 6164 (Proposed Std). Updated by RFC 6547. IETF. IETF.

Krishnan, S. (2009, December). Handling of Overlapping IPv6 Fragments. RFC 5722 (Proposed Std). Updated by RFC 6946. IETF. IETF.

Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J. & Bhatia, M. (2012, April). A Uniform Format for IPv6 Extension Headers. RFC 6564 (Proposed Std). IETF. IETF.

Levy-Abegnoli, E., de Velde, G. V., Popoviciu, C. & Mohacsi, J. (2011, February). IPv6 Router Advertisement Guard. RFC 6105 (Informational). IETF. IETF.

Li, X., Bao, C. & Baker, F. (2011, April). IP/ICMP Translation Algorithm. RFC 6145 (Proposed Std). Updated by RFC 6791. IETF. IETF.

Lindqvist, J. (2006). IPv6 Stateless Address Auto-configuration Considered Harmful. In *Military Communications Conference, 2006. MILCOM 2006. IEEE* (Pages 1–5). doi:10.1109/MILCOM.2006.302471

Lucena, N. B., Lewandowski, G. & Chapin, S. J. (2006). Covert Channels in IPv6. In *Proceedings of the 5th International Conference on Privacy Enhancing Technologies* (Pages 147–166). PET'05. Cavtat, Croatia: Springer. doi:10.1007/11767831_10

Malone, D. (2008). Observations of IPv6 addresses. In *Passive and Active Network Measurement* (Pages 21–30). Springer.

Miller, I. (2001, June). Protection Against a Variant of the Tiny Fragment Attack (RFC 1858). RFC 3128 (Informational). IETF. IETF.

Moore, N. (2006, April). Optimistic Duplicate Address Detection (DAD) for IPv6. RFC 4429 (Proposed Std). IETF. IETF.

Nakibly, G. [G.] & Templin, F. (2011, August). Routing Loop Attack Using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations. RFC 6324 (Informational). IETF. IETF.

Nakibly, G. [Gabi] & Arov, M. (2009). Routing Loop Attacks using IPv6 Tunnels. In *Proceedings of the 3rd USENIX conference on Offensive technologies* (Pages 7–7). USENIX Association.

Narten, T., Draves, R. & Krishnan, S. (2007, September). Privacy Extensions for Stateless Address Autoconfiguration in IPv6. RFC 4941 (Draft Std). IETF. IETF.

Narten, T., Huston, G. & Roberts, L. (2011, March). IPv6 Address Assignment to End Sites. RFC 6177 (Best Current Practice). IETF. IETF.

Narten, T., Nordmark, E., Simpson, W. & Soliman, H. (2007, September). Neighbor Discovery for IP version 6 (IPv6). RFC 4861 (Draft Std). Updated by RFC 5942. IETF. IETF.

Nikander, P., Kempf, J. & Nordmark, E. (2004, May). IPv6 Neighbor Discovery (ND) Trust Models and Threats. RFC 3756 (Informational). IETF. IETF.

Nordmark, E. & Gilligan, R. (2005, October). Basic Transition Mechanisms for IPv6 Hosts and Routers. RFC 4213 (Proposed Std). IETF. IETF.

Partridge, C. & Jackson, A. (1999, October). IPv6 Router Alert Option. RFC 2711 (Proposed Std). Updated by RFC 6398. IETF. IETF.

Perkins, C., Johnson, D. & Arkko, J. (2011, July). Mobility Support in IPv6. RFC 6275 (Proposed Std). IETF. IETF.

Pfleeger, C. P. & Pfleeger, S. L. (2003). *Security in computing*. Prentice Hall Professional.

Rafiee, H., Alsa'deh, A. & Meinel, C. (2011). WinSEND: Windows SEcure Neighbor Discovery. In *Proceedings of the 4th international conference on Security of information and networks* (Pages 243–246). ACM.

Rajahalme, J., Conta, A., Carpenter, B. & Deering, S. (2004, March). IPv6 Flow Label Specification. RFC 3697 (Proposed Std). Obsoleted by RFC 6437. IETF. IETF.

Savola, P. & Patel, C. (2004, December). Security Considerations for 6to4. RFC 3964 (Informational). IETF. IETF.

Templin, F., Gleeson, T. & Thaler, D. (2008, March). Intra-Site Automatic Tunnel Addressing Protocol (ISATAP). RFC 5214 (Informational). IETF. IETF.

Thaler, D. [D.]. (2011, January). Teredo Extensions. RFC 6081 (Proposed Std). IETF. IETF.

Thomson, S. & Narten, T. (1998, December). IPv6 Stateless Address Autoconfiguration. RFC 2462 (Draft Std). Obsoleted by RFC 4862. IETF. IETF.

Thomson, S., Narten, T. & Jinmei, T. (2007, September). IPv6 Stateless Address Autoconfiguration. RFC 4862 (Draft Std). IETF. IETF.

Townsley, W. & Troan, O. (2010, August). IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) – Protocol Specification. RFC 5969 (Proposed Std). IETF. IETF.

Tsirtsis, G. & Srisuresh, P. (2000, February). Network Address Translation - Protocol Translation (NAT-PT). RFC 2766 (Historic). Obsoleted by RFC 4966, updated by RFC 3152. IETF. IETF.

Ullrich, J., Krombholz, K., Hobel, H., Dabrowski, A. & Weippl, E. (2015). IPv6 Security: Attacks and Countermeasures in a Nutshell. *Magdeburger Journal zur Sicherheitsforschung*, *1*, 514–529. Retrieved March 30, 2015, from http://www.sicherheitsforschung-magdeburg.de/publikationen.html

van Dijk, P. (nodate). Finding v6 hosts by efficiently mapping ip6.arpa. Retrieved from http://7bits.nl/blog/posts/finding-v6-hosts-by-efficiently-mapping-ip6-arpa

Wu, J., Cui, Y., Li, X., Xu, M. & Metz, C. (2010, March). 4over6 Transit Solution Using IP Encapsulation and MP-BGP Extensions. RFC 5747 (Experimental). IETF. IETF.

Yang, X., Ma, T. & Shi, Y. (2007). Typical DoS/DDoS threats under IPv6. In *Computing in the Global Information Technology, ICCGI 2007. International Multi-Conference on* (Pages 55–55). IEEE.

Ziemba, G., Reed, D. & Traina, P. (1995, October). Security Considerations for IP Fragment Filtering. RFC 1858 (Informational). Updated by RFC 3128. IETF. IETF.