



## **NISCC Technical Note 02/2006**

**Issued 24 April 2006**

### **Security considerations for IPv6**

IPv6 is an extremely complicated protocol suite, with a large number of Internet RFCs linked to it, and as such there are many different areas which have security implications. This document attempts to summarise current thoughts on the security of IPv6 and linked protocols and technologies.

The reader of this document should be reasonably familiar with IPv6 (its general form and addressing) and have a technical background. The majority of security observations within this document identify points where IPv6 imposes new or increased risks; a few topics include points where IPv6 reduces risk, however this is by no means complete.

Recommendations are offered when best practice is available; however this document is not a comprehensive lockdown document.

Reference to any specific commercial product, process or service by trade name, trademark manufacturer, or otherwise, does not constitute or imply its endorsement, recommendation, or favouring by NISCC. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes.

NISCC shall also accept no responsibility for any errors or omissions contained within this document. In particular, NISCC shall not be liable for any loss or damage whatsoever, arising from the usage of information contained in this document.

**National Infrastructure  
Security Co-ordination Centre**  
PO Box 832  
London, SW1P 1BG

Tel: 0870 487 0748  
Fax: 0870 487 0749  
Email: [enquiries@niscc.gov.uk](mailto:enquiries@niscc.gov.uk)  
Web: [www.niscc.gov.uk](http://www.niscc.gov.uk)

## Introduction

1. This document attempts to summarise current thoughts on the security of IPv6 and linked protocols and technologies. There has been some research into this, with a large number of papers and Internet RFC<sup>1</sup>s produced. However, no cohesive document yet exists which pulls these separate findings together; it is the purpose of this paper to be that cohesive document. This paper would not be possible without the large amount of work already performed in academia and elsewhere; the references. All the acronyms used in this document are defined, from the common to the obscure, and all external documents referred to are indicated and linked to.
2. IPv6 is an extremely complicated protocol suite, with a large number of Internet RFCs linked to it, and as such there are many different areas which have security implications. These different areas are covered below, followed by some general conclusions. Recommendations are offered when best practice is available; however this document is not a comprehensive lockdown document, and should not be used as such. Some comments are also made on possible mitigations which are not currently best practice; these should not be taken as policy or any form of formal recommendation. If there is sufficient feedback to NISCC we will produce a document that does contain formal recommendations.
3. The reader of this document should be reasonably familiar with IPv6 (its general form and addressing) and have a technical background; no description will be offered for core aspects of IPv6, although brief précis are given. The majority of security observations within this document identify points where IPv6 imposes new or increased risks; a few topics include points where IPv6 reduces risk, however this is by no means complete.

## IPv6 Addressing

4. The primary, or certainly the most commonly hailed, reason for the upgrade to IPv6 is the improvement to the address space. IPv4 has a 32-bit address field, which has proved insufficient for modern usage - with worldwide unique addressing of a plethora of common devices being desirable in the near future. The 128-bit address space of IPv6 is  $2^{96}$  times that of IPv4 - thus theoretically solving this problem. Additionally, there have been changes to how this address space is split up, partly to ease address aggregation - see Router Renumbering, below, for more detail.

---

<sup>1</sup> Request for Comment. For the definitive list of RFCs see [http://www.ietf.org/iesg/1rfc\\_index.txt](http://www.ietf.org/iesg/1rfc_index.txt)

5. RFC3513<sup>2</sup> defines the Addressing Architecture for IPv6. It describes three different types of address: unicast, multicast, and anycast, with explicit broadcast being dropped in favour of a specific case of multicast. In most cases the address is made up of a top 64 bits representing the network (and subnet), and the low 64 bits being the host. The host part of the address can be any hard-coded value, or by default (during autoconfiguration) an encoded form of the MAC<sup>3</sup> address. A privacy extension is available, to instead use a pseudo-random number for the host part of the address. The network part is broken down by address type, and scoping - global, link-local, or (the now deprecated - see below) site-local.
6. These changes will affect security in numerous ways. Firstly, the massively increased address space will naturally lead to a lower density of hosts, and hence make IP-targeted (as compared to user targeted - e.g. email) worm propagation vastly more difficult. However the use of multicast addresses may, once a single host is compromised by other means, lead to the extremely rapid compromise of all vulnerable hosts on that network - potentially speeding propagation on networks with large subnets.
7. Due to the numerous address types, and expanded use of multicast, interfaces will always have a plethora of addresses associated with them; in the minimal usable case there will be a link-local unicast, a global unicast, and the host will be listening to the link-local all-hosts multicast and the solicited-nodes multicast address for each unicast and anycast address owned. Moreover, if privacy extensions are used, the low 64 bits of the locally bound addresses will be periodically changing. This can cause difficulties in selecting source IP address for outbound packets, lead to enlarged local routing tables, and cause difficulties with both host, and traditional network, firewalls.
8. There are currently two camps regarding the use of Network Address Translation (NAT) - one insists that there is no need for NAT as under IPv6 there are easily sufficient addresses, whereas the others point to the benefits of NAT in network encapsulation, and pseudo-security. Each of these arguments has merit, and this paper does not take a position on them. No matter what the answer, it is certain that the network model for IPv6 networks will be very much end-to-end. This is a shift from current models, where NAT and similar lead to communication channels being set up unidirectionally, from client to server.

---

<sup>2</sup> See <http://www.ietf.org/rfc/rfc3513.txt?number=3513>

<sup>3</sup> Media Access Control address. "A unique code assigned to most forms of networking hardware. The address is permanently assigned to the hardware, so limiting a wireless network's access to hardware -- such as wireless cards -- is a security feature employed by closed wireless networks. But an experienced hacker -- armed with the proper tools -- can still figure out an authorized MAC address, masquerade as a legitimate address and access a closed network." See [edition.cnn.com/2004/TECH/internet/10/25/glossary/](http://edition.cnn.com/2004/TECH/internet/10/25/glossary/)

9. Multicast addressing was available in IPv4 but was not (until recently) especially utilised on the Internet. IPv6 makes a great deal of use of multicast in the local network, and may do so in the wider network. Multicast addresses could be used as 'smurf'-style<sup>4</sup> amplifiers for Distributed Denial of Service (DDoS) attacks, or network enumeration. In a worst-case scenario, multicast could be used for extremely rapid, lightweight, worm propagation. If multicast is taken up more within the WAN<sup>5</sup> and Internet, there could be issues from larger numbers of requests to join multicast groups, leading to massively increased network traffic as reports ripple through networks, and potentially to Denial of Services (DoSs) of routers due to their multicast tables being filled; this can be mitigated in many ways, including rate-limiting of acceptance for multicast joins, setting sensible limits on router multicast tables, and border filtering by ISP<sup>6</sup>s and similar.
10. Anycast is a technique where multiple hosts on a network listen for the same IP address, and rely on routing to differentiate between them. This has the benefits of geographical (from a network perspective) load balancing, and improved redundancy against technical problems. Anycast is already used on the Internet, with the F DNS<sup>7</sup> root nameserver having been using anycast since November 2002, which has helped defend it against numerous sizable DDoS attacks. Improved support for anycast thus has the potential to improve the stability of large services, although with a possible increase in management complexity - most of which should be handled automatically by routing protocols.
11. Finally, scoping is a technique for giving information within the IPv6 address as to the expected scope of the packet - i.e. how far on the network it should travel. There are two main scopings, global and link-local, with a third (site-local) - discussed below - being deprecated. Fundamentally, link-local scoped addresses are guaranteed not to be routed, providing some level of protection from attacks against Neighbour Discovery (considered below) and others, whilst providing some level of assurance that link-local scoped packets will not accidentally be forwarded onto other networks which would lead to network confusion and potential information confidentiality loss.

---

<sup>4</sup> See <http://www.cert.org/advisories/CA-1998-01.html>

<sup>5</sup> A wide area network is a geographically dispersed telecommunications network. The term distinguishes a broader telecommunication structure from a local area network (LAN).

<sup>6</sup> Internet Service provider.

<sup>7</sup> Domain Name System.

## Deprecation of Site-Local Addressing

12. In the original IPv6 addressing architecture RFC1884<sup>8</sup> there were two local-use address spaces - Link-Local and Site-Local. It was envisaged the Site-Local addresses could be used internally within a LAN<sup>9</sup> or WAN, with border routers guaranteeing that these would not be routed into other networks, similar to the private IPv4 addresses specified within RFC1918<sup>10</sup>. Issues with this included the substantial confusion over what a 'site' was, and technical problems with source address selection, DNS, and routing. It was confirmed in April 2003 that this feature is being deprecated, and formally deprecated with RFC3879<sup>11</sup> in September 2004. However it is noted that it is still present in the most recent addressing architecture RFC, RFC3513. It is further noted that different operating systems appear to handle site-local addressing differently.
13. With these points in mind there are several security implications for the deprecation of site-local addressing. Whilst there is almost unanimous agreement that this deprecation is required, there may be some implications in how different legacy implementations of IPv6 handle site-local addresses - partly this is due to the delays in the deprecation process and the fact that in order to be compliant with the current addressing architecture document RFC3513, site-local addressing should be supported.
14. Possible issues include loss of confidentiality if, for example, a number of hosts are using site-local addressing, and yet border routers no longer support it and hence route site-local packets as global packets.
15. The IPv6 IETF<sup>12</sup> working group confirming the deprecation of site-local in a timely manner and producing a revised RFC updating RFC3513 accordingly would help to reduce problems related to this. Furthermore, organisations may be able to determine, and mitigate, the risk by firstly checking that devices used for enforcement of routing decisions (firewalls and routers) behave correctly if site-local addresses are received, and secondly configuring external IDS<sup>13</sup>s (as a minimum) to detect site-local addresses - thus detecting attempted use of site-local addressing.
16. Even though site-local scoping is deprecated, there is still a great deal of interest in some form of limited-scope addressing and it looks likely that

---

<sup>8</sup> See <http://www.ietf.org/rfc/rfc1884.txt?number=1884>

<sup>9</sup> Local Area Network.

<sup>10</sup> See <http://www.ietf.org/rfc/rfc1918.txt?number=1918>

<sup>11</sup> See <http://www.ietf.org/rfc/rfc3879.txt?number=3879>

<sup>12</sup> Internet Engineering Task Force. "A large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet." See <http://www.ietf.org/>

<sup>13</sup> Intrusion Detection System. See [http://en.wikipedia.org/wiki/Intrusion-detection\\_system](http://en.wikipedia.org/wiki/Intrusion-detection_system)

some form of this will be accepted in the future. Unique-local unicast addressing has been proposed<sup>14</sup>, which uses a /7 portion of the IPv6 address space. The specific address space is globally scoped, but is well known as unique-local and hence will not be routed or aggregated in the wider Internet. Additionally, by having non-ambiguous addresses debugging of leaks is facilitated, as is the joining of multiple scoped networks as occurs in larger corporate and similar networks. Additionally, RFC4007<sup>15</sup> discusses scoping within IPv6 in more depth.

## Next Header Field

17. Another major change to the base IPv6 header is that it has been simplified, with the majority of fields from within the IPv4 header now being represented as options which follow in a linked-list from the IPv6 header. This linked-list uses the Next Header (NH) field to point to the next header along.
18. This vastly simplifies the header, reducing the cost to process the packet by intermediary routers (assuming all they look at is the destination, and possibly source, IPv6 address), but at the cost of increased complexity for endpoints. This complexity introduces several possible problems.
19. Unfortunately, the assumption that intermediary routers only examine the IPv6 address fields is fallacious. The hop-by-hop options header, which if present should be immediately follow the IPv6 header, must be examined by intermediaries. Furthermore, source-routing is accomplished by a Routing Header. The destination field in the IPv6 header is only the first hop - the routing header contains source-routed hops and the ultimate destination. It will often not be possible to restrict use of source-routing as the optimisations within MobileIPv6 rely on it. All of these could lead to increases in processing load on routers, with certain packets possibly becoming more processor-intensive in their IPv6 forms than their equivalent would have been within IPv4.
20. The purposes behind the redesign of option handling between IPv4 and IPv6 were twofold: firstly to simplify the main header, and secondly to ease support of future options as required. However, as there is no length field required in any header, then all preceding headers must be parsed and understood in order to access a later header.
21. RFC2460<sup>16</sup> specifies that if a header is unrecognised, the packet should be dropped and an ICMP<sup>17</sup> Parameter Problem message be returned.

---

<sup>14</sup> RFC 4193. See <http://www.ietf.org/rfc/rfc4193.txt?number=4193>

<sup>15</sup> See <http://www.ietf.org/rfc/rfc4007.txt?number=4007>

<sup>16</sup> See <http://www.ietf.org/rfc/rfc2460.txt?number=2460>

<sup>17</sup> Internet Control Message Protocol. See <http://www.ietf.org/rfc/rfc2463.txt?number=2463>

Whilst this is satisfactory for the concept of unfiltered end-to-end communication, there are obvious issues as soon as filtering is required based on higher-level protocols than IPv6. In these cases all headers must be parsed - so what should be done if, for example, a firewall ruleset specifies TCP<sup>18</sup> ports, and yet an intermediary header is unrecognised? In this instance it would be logical to return a forged Parameter Problem packet, but there are problems here as well. For example, the header may be a required but proprietary extension, or IPSec<sup>19</sup> may be being used with AH, and so the packet couldn't be forged. With this in mind it seems conceivable that no further usable option headers will be defined, due to obvious legacy issues. It would therefore be wise for vendors and government not to create proprietary headers, and for filtering and IDS rules to drop and flag on any headers other than those specified within RFC2460. There is additionally no real requirement for use of new headers by organisations concerned solely with network communication endpoints; the Destination Options header contains TLV<sup>20</sup> encoded options, in which any proprietary extensions could be included without impact.

22. Another negative side-effect of the complexity is the additional processing time required at intermediary filters (e.g. firewalls) where rules refer to higher-level protocols, as is normally the case. The impact of this additional processing is small but measurable, and hence could increase load on key devices - normal load-balancing techniques must be used to mitigate this. It is possible that this additional processing could be used to increase the effect of DDoS attacks - however the magnitude of the additional impact will not be the conclusive factor in any DDoS attack.
23. RFC2460 also specifies the order in which option headers should appear, although then goes on to say that "IPv6 nodes must accept and attempt to process extension headers occurring in any order and occurring any number of times in the same packet.". This has a number of negative security implications.
24. Firstly, as mentioned above, the impact of delays due to processing time could be multiplied by using multiple headers. This is especially the case if

---

<sup>18</sup> Transmission Control Protocol.

<sup>19</sup> Internet Protocol Security. "A framework of open standards developed by the IETF that provides security for transmission of sensitive information over unprotected networks such as the Internet. It acts at the network level and implements the following standards: IPSec, Internet Key Exchange (IKE), DES, MD5, SHA, Authentication Header (AH), and Encapsulating Security Payload (ESP)." See [www.vovida.org/document/openosp/DCLv2-TestResults.html](http://www.vovida.org/document/openosp/DCLv2-TestResults.html)

<sup>20</sup> "Type, Length, Value". A way of encoding information that formed the basis of ASN.1 BER. BER stands for Basic Encoding Rules - "Standard rules for encoding data units described in ASN.1 (Abstract Syntax Notation One). Sometimes incorrectly lumped under the term ASN.1, which properly refers only to the abstract syntax description language, not the encoding technique. See [www.ee.oulu.fi/research/ouspg/sage/glossary/](http://www.ee.oulu.fi/research/ouspg/sage/glossary/)

we can 'stack' IPSec headers, e.g. Encapsulating Security Payload (ESP) header (and payload) within ESP header, within ESP header etc.

25. Secondly, different devices may handle headers in different orders. For example, some implementations may use the first fragmentation header encountered, others may use the last, and others may use both. This (whether using the fragmentation header, or others) could lead to the success of methods for IDS and firewall evasion techniques seen in the early 1990s, and discussed within RFC1858<sup>21</sup>.
26. With these two points in mind, the best-practice recommendations made within RFC2460 should be adhered to, specifically that "Each extension header should occur at most once, except for the Destinations Options header which should occur at most twice" and that headers must occur in the order specified therein. Furthermore, it is advisable that these be enforced at network boundaries (and optionally internally, dependent on security model), and IDSs used to trigger when this is not the case (non-trivial with signature-based NIDS<sup>22</sup>).

## Neighbour Discovery

27. Neighbour Discovery (ND) is defined in RFC2461<sup>23</sup>. Simply put, this refers to the IPv6 version of IPv4 ARP<sup>24</sup>, together with some autoconfiguration, neighbour unreachability detection (NUD), redirection, and duplicate address detection (DAD). The RFC is, upon initial consideration, deceptively simple. On deeper investigation it is extremely complex, with numerous issues. Many of these issues are exact IPv6 ND versions of IPv4 ARP, whereas some are wholly new to IPv6.
28. An excellent RFC has been produced looking exclusively at the trust models and threats within IPv6 ND (RFC3756<sup>25</sup>), and the majority of this section is based on the findings within that document. Fundamentally, it was found that if IPSec is not used, then any host on a subnet could perform man-in-the-middle and spoofing attacks on any other host (and therefore DoS attacks as a special case of these), through a wide range of attacks. There are several defences against this.

---

<sup>21</sup> See <http://www.ietf.org/rfc/rfc1858.txt?number=1858>

<sup>22</sup> Network[-based] Intrusion Detection System. "A network intrusion detection system (NIDS) is a system that tries to detect malicious activity such as denial of service attacks, port-scans or even attempts to crack into computers by monitoring network traffic." See <http://en.wikipedia.org/wiki/Nids>

<sup>23</sup> See <http://www.ietf.org/rfc/rfc2461.txt?number=2461>

<sup>24</sup> Address Resolution Protocol. See RFC 826 <http://www.ietf.org/rfc/rfc0826.txt?number=826>

<sup>25</sup> See <http://www.ietf.org/rfc/rfc3756.txt?number=3756>

29. Firstly, RFC2461 states that ND messages must be received with a TTL<sup>26</sup> of 255, thus ensuring that no routed packets can be used for these attacks, as any routing will decrement the TTL from its value, which is at most 255. However, one of the key transition techniques is tunnelling - and it is not clear from the specifications how a host will handle a tunnelled packet with IPv6 TTL of 255.
30. Secondly, RFC2461 recommends the use of IPSec (namely the AH header) for the majority of ND packets. Whilst this would provide ample protection from the majority, and probably all, feasible attacks, there are issues with using IPSec here. For more details the reader is referred to the IPSec section below.
31. Thirdly, a large number of attacks can be defended against through turning off large proportions of autoconfiguration, and manually specifying key link-layer to IPv6 mappings. Specifically, by disabling handling of all received autoconfiguration messages, we can defend against attacks which cause us to form, and use, arbitrary IP addresses and routers. By manually (e.g. using the 'arp' command) setting the MAC address for a trusted router, man-in-the-middle attacks cannot be performed from the host to the router (although the reverse can still happen). Unfortunately, the management overhead of this is extremely high - in almost all production environments this particular defence would impose an insurmountably onerous network management overhead. Consider, for example, the overhead involved in manually reconfiguring every host on a class B network whenever the router used as default gateway is unavailable. Furthermore, the use of load-balancing or hot-standby routers would be negated, as only a single link-layer address can normally be associated with a single IP address.
32. Fourthly, NIDS can be set to look for certain very specific 'incorrect' packets - for example multiple replies to Address Solicitation requests (the equivalent of ARP requests). Whilst not directly defending against attack, this will flag up in a timely manner the majority of attacks.
33. Ultimately however, without use of IPSec, there is no guaranteed method for defending against attacks between hosts on the local subnet; the situation that would exist for some 'insider' attacks. IPSec is in most cases not a viable choice, partly due to bootstrapping issues. To solve some of these issues, 'Securing Neighbour Discovery' (SEND) has been designed; however this is still only in draft stage, and no production-quality implementations currently exist.

---

<sup>26</sup> Time To Live. A value carried in the header of each IP packet that is decremented by each router through which it passes. When TTL is decremented to 0, routers will not forward it, and it no longer 'lives'.

## ICMPv6

34. Control of IPv6 is performed by ICMPv6, which plays much the same role as ICMPv4 for IPv4 with the same strengths and weaknesses. However, there is additional functionality within ICMPv6 which deserves special consideration; it is used for fragmentation control and Path MTU<sup>27</sup> discovery, the same as ICMPv4. Unlike IPv4, IPv6 fragmentation takes place wholly at the endpoints - i.e. broadly equivalent to IPv4 with the Don't Fragment bit set. ICMPv6 messages related to fragmentation control, specifically the Packet-Too-Big message, must be allowed through all network filters in order for correct operation to occur. Additionally, firewall rules must allow proper operation of ND, although obviously only on the local network - however note that most firewalls are currently configured not to allow IP traffic to be sent from the firewall itself - a ruleset that will require modification<sup>28</sup>.
35. Ultimately, conformance with the IPv6 specifications will force a sizeable change in network boundary filtering; control traffic that is currently dropped at network boundaries will instead be passed. In particular, ICMPv6 messages related to multicast, neighbour discovery, and path MTU discovery will have to pass in many cases. These protocols may expose channels for probing, illicit communication, and attack. One likely implication is that monitoring and anomaly detection, joined with active response, will have to play a greater role in network defence. Additionally, there may be increased use of proxying technologies, as these allow the higher-level protocols to still be used whilst mitigating risks from lower-level protocols.

## Router Renumbering

36. Over the last few years, BGP<sup>29</sup> and route tables for the core routers of the Internet have been getting exponentially larger, leading to performance issues when routing. IPv6 has attempted to solve this in two ways.
37. Firstly, by having a larger address space, there will be less of a requirement for efficient use of address space, and hence address space fragmentation may not occur to the same degree as with IPv4.
38. Secondly, the concept of router renumbering (RR). This is defined in RFC2894<sup>30</sup>, and provides a mechanism for renumbering entire networks by sending a single ICMPv6 message. The concept is that an RR packet

---

<sup>27</sup> Maximum Transmission Unit, the largest packet size [in bytes] that can be supported by a particular network implementation.

<sup>28</sup> For a more complete list of these, and other, recommendations, the reader is referred to <http://seanconvery.com/v6-v4-threats.pdf>

<sup>29</sup> Border Gateway Protocol. See RFC 1771 <http://www.ietf.org/rfc/rfc1771.txt?number=1771>

<sup>30</sup> See <http://www.ietf.org/rfc/rfc2894.txt?number=2894>

will be sent to the core router for that network, which will then through ND packets inform all the hosts on its network of the new address. Obviously, only hosts using address autoconfiguration will be updated through this technique - however this is the expected case for clients. When used, this will in theory allow entire networks to easily be moved in address space, allowing optimal aggregation.

39. This RFC is one of few that specify that IPSec AH must be used, for obvious reasons. Due to this, the primary security implications are dependent on IPSec security, and hence out of scope of this document. There is the possibility of a race condition for firewalls etc during the transition between old and new address spaces - this would need to be closely managed at the time.
40. It is the author's belief that Router Renumbering will not be an especially common occurrence for numerous technical, psychological, and managerial reasons. However, the functionality is present and so protection must be in place against both accidental and hostile use of it.
41. The A6 DNS record is a major facilitator of router renumbering. However, as discussed under DNS, below, there is currently a large amount of discussion which may result in the A6 record being deprecated, hence rendering router renumbering even less attractive.

## Mobile IPv6

42. The phrase Mobile IPv6 covers a suite of RFCs to allow more optimal mobile networking than is currently the case for IPv4. Currently, and in the less optimal version of Mobile IPv6, packets are sent via a 'home agent', which forwards them to the care-of address of the mobile node. This is extremely inefficient, especially when both the correspondent node and mobile node may be on the same network, or indeed within the same cell for 3G telephony. Therefore with Mobile IPv6, route optimisation was designed in. Simply put, the correspondent node (the node communicating with the mobile node) acquires from the Home Agent the current care-of address of the mobile node, and addresses it's packets to it. A number of ICMPv6 types, an additional IPv6 destination option, and a new IPv6 protocol, have all been defined to allow this.
43. A number of threat scenarios are detailed in an IETF Internet Draft<sup>31</sup>. Included in these scenarios are broad requirements for methods to mitigate these. Vendors and network operators implementing Mobile IPv6 networks are recommended to read these.
44. Obviously, these additional items must be allowed or blocked as required by intermediary firewalls. Due to the complexity of the protocol, and number of holes that must be punched through firewalls, it is

---

<sup>31</sup> See <http://www3.ietf.org/proceedings/02mar/I-D/draft-ietf-mobileip-mipv6-scrty-reqts-02.txt>

recommended that home agents be situated on a network of their own, e.g. a DMZ<sup>32</sup>, with appropriately configured firewalls.

## IPSec

45. Whilst it is not within the scope of this document to examine encryption algorithms or technologies per se., IPSec plays such an important role within IPv6 that it must be considered - indeed, there is some justification in the claim by IPv6 detractors that the default answer to any question regarding the security of IPv6 is 'just use IPSec'.
46. IPSec within IPv6 is represented through two headers, both of which are mentioned in the main IPv6 RFC - RFC2460. These are the Authentication Header (AH) specified in RFC2402<sup>33</sup>, and the Encapsulating Security Header (ESP) in RFC2406<sup>34</sup>. Simply, AH provides authentication and cryptographic integrity checking of the packet payload, and ESP encrypts the packet payload together with a providing a signature, hence providing cryptographic confidentiality and integrity.
47. There are three core problems with securing IPv6 control and application protocols through IPSec; the requirement for some form of PKI<sup>35</sup>, the computationally intensive nature of cryptographic algorithms, and the latency inherent in key-exchange protocols and cryptographic calculations.
48. First, for successful IPSec deployment that provides trusted authentication and non-repudiation there must be some form of pre-shared key material (normally X.509<sup>36</sup> certificates), and optionally a PKI. This is a requirement for even the simplest unidirectional host authentication - the client requires the public-key of the server (either manually distributed or automatically, but ultimately signed through a chain of trust by a host whose public-key the client has, i.e. ultimately there must be some manual distribution). For other currently existent cryptographic protocols, namely SSL<sup>37</sup>

---

<sup>32</sup> De-Militarised Zone (or De-Militarized Zone). See [http://en.wikipedia.org/wiki/Demilitarized\\_zone\\_%28computing%29](http://en.wikipedia.org/wiki/Demilitarized_zone_%28computing%29)

<sup>33</sup> See <http://www.ietf.org/rfc/rfc2402.txt?number=2402>

<sup>34</sup> See <http://www.ietf.org/rfc/rfc2406.txt?number=2406>

<sup>35</sup> Public Key Encryption. "An arrangement which provides for third-party vetting of, and vouching for, user identities. It also allows binding of public keys to users. This is usually carried by software at a central location together with other coordinated software at distributed locations. The public keys are typically in certificates. See [http://en.wikipedia.org/wiki/Public\\_key\\_infrastructure](http://en.wikipedia.org/wiki/Public_key_infrastructure)

<sup>36</sup> X.509 is an ITU-T standard for public key infrastructure (PKI). X.509 specifies, amongst other things, standard formats for public key certificates and a certification path validation algorithm. See <http://en.wikipedia.org/wiki/X.509>

<sup>37</sup> Secure Sockets Layer. "SSL provides endpoint authentication and communications privacy over the Internet using cryptography. In typical use, only the server is authenticated (i.e. its identity is ensured) while the client remains unauthenticated; mutual authentication requires public key infrastructure (or PKI) deployment to clients. The protocols allow client/server

- authentication of websites on the Internet, this is accomplished through the host operating system being provided with the certificates of several trusted certificate authorities at install-time.
49. When a security association (SA) is required to be formed between hosts for even the most trivial of transactions (e.g. a ping of a host on the local subnet requires AH be used to ensure no spoofing occurs, requiring both sender and receiver to have public keys signed by a common trusted authority), the overhead of such transactions could be immense. Additionally, management of certificate revocation lists, for, for example, a dial-up ISP whose clients are continually acquiring new IPv6 addresses, or for users using the privacy extensions to IPv6 addressing and hence changing their IPv6 addresses often, quickly becomes a massive task.
  50. Bootstrapping becomes problematic for initial host autoconfiguration, requiring the host to already have the certificate for the router used for autoconfiguration, or an alternate layer-3 protocol must be used for verification of chain of trust if a certificate is proffered at autoconfiguration time.
  51. All of these key-material-related problems, combined with the difficulty of end-user technical support by ISPs etc due to lack of user training, would lead to a nightmare for ISPs, with the probable outcome that, in general, ISPs (and others) will not bother to use IPsec.
  52. The second problem with widespread use of IPsec is the computational expense of cryptographic operations. Whilst this is not a major concern for desktop and server machines (although still not necessarily trivial for servers due to the number of parallel communications), one of the key areas of prospective early uptake of IPv6 is mobile devices. Miniature devices currently in fact appear to represent one of the key long-term markets and drivers for IPv6. However, these devices do not currently, in general, have the processing power to perform the number and magnitude of cryptographic operations required. This is a particular problem for 3G mobile telephony, as IPsec is key to secure use of Mobile IPv6, as mentioned previously. It is possible therefore that some device vendors may not support the full cipher-suite expected.
  53. Thirdly, for time-sensitive or short-duration transactions, such as the 3G and mobile-IP networks, the latency inherent in cryptographically secure key-exchange protocols may be far in excess of that allowed for handovers to be transparent to users. This is due to both the amount of time required by low-power devices to perform cryptographic operations, and also the network latency in the challenge-response transactions required by cryptographically secure key-exchanges. Elliptic curve algorithms may ease both this and the previous issue, however this is a

---

applications to communicate in a way designed to prevent eavesdropping, tampering, and message forgery.” See [http://en.wikipedia.org/wiki/Secure\\_Sockets\\_Layer](http://en.wikipedia.org/wiki/Secure_Sockets_Layer)

relatively new technology, saddled with both licensing issues and lack of standardisation.

## Transition Mechanisms (IPv4 to IPv6)

54. In the most generic case, two transition mechanisms have been postulated, and it is certain that both will be used: tunnelling is the process of encapsulating IPv6 packets within the payload of IPv4 (or higher level protocols - UDP in the case of Teredo<sup>38</sup>), and dual-stack refers to running both IPv4 and IPv6 stacks on a host at the same time. When tunnelling is being used, by definition the hosts at the endpoints of the tunnel must be dual-stack.
55. Both dual-stack and tunnel transition mechanisms incur a substantial risk of being used to bypass network security restrictions. Both misconfiguration and malicious attack can create situations where traffic can flow into and out of a network unrestricted by filters or firewall rulesets. Incident response and forensic analyses have shown that Internet hijackers frequently use IPv6 traffic and tunnels as a means to hide their illicit communications.
56. For both techniques, the primary security implication is the probability of misconfiguration allowing the circumvention, whether accidental or otherwise, of network security restrictions - e.g. firewall rulesets.
57. It is a standard premise of IT security that whitelists are more secure than blacklists, and this must be applied to firewall rulesets. In the dual-stack configuration, assuming the IPv4 firewalling is well configured and that IPv6 is actually required, it is recommended that the IPv6 firewalling be configured to be functionally identical to the IPv4 rules for global unicast addresses, unless local policy states otherwise. Multicast should only be allowed through if there is a business case for it, and should be restricted to only those expected. Link-local scoped addresses should not be routed, however routers (and non-routing firewalls) will need to be configured to allow link-local addresses to be sent outbound, originating from the router, for DAD and similar purposes. Standard rules about Martian filtering<sup>39</sup> etc still apply. For other recommendations on transition from IPv4 to IPv6, see the documents by Savola<sup>40</sup>, by DuPont and Castelluccia<sup>41</sup>, by Renard<sup>42</sup>

---

<sup>38</sup> The Teredo protocol defines a method to access the IPv6 Internet from behind a NAT device. It consists of encapsulating IPv6 packets through UDP over IPv4 between Teredo clients and Teredo relays, with the help of Teredo servers. See <http://www.ietf.org/internet-drafts/draft-huitema-v6ops-teredo-05.txt>

<sup>39</sup> "A packet that contains an invalid source or destination address is considered to be Martian and discarded." See [www.freesoft.org/CIE/RFC/1812/197.htm](http://www.freesoft.org/CIE/RFC/1812/197.htm)

<sup>40</sup> See <http://www.6net.org/publications/standards/draft-savola-v6ops-firewalling-01.txt>

<sup>41</sup> See <http://www.watersprings.org/pub/id/draft-dupont-ipv6-ingress-filtering-00.txt>

<sup>42</sup> See <http://www.usipv6.com/6sense/2004/nov/november05.htm>

and by the US Government Accountability Office<sup>43</sup>. It should be noted that when a service binds to an IPv6 port, it by default also binds to the corresponding IPv4 port - this must be borne in mind when crafting firewall rules.

58. IPv6 within IPv4 tunnels presents an unrestricted ingress path for network traffic. This is similar to the current risk from VPNs, however due to cryptographic checks VPNs cannot be injected into from arbitrary hosts; this is not the case for IPv6 tunnels unless IPsec is used. As such, it is imperative that they be controlled, and standard best practice for VPNs provides a good start for this. Tunnel endpoints should occur within a DMZ, and firewalls should be used to enforce that no other endpoints are allowed. Of specific concern here are some of the 'automatic' tunnelling techniques, such as v4 compatible, and 6to4 addressing - where hosts may be configured and listening as tunnel endpoints without system administrator or host owner knowledge.
59. One postulated attack utilising tunnels is to interfere with IPv6 neighbour discovery at the tunnel endpoint - forging tunnelled IPv6 packets to have a TTL of 255, and hoping that the tunnel endpoint doesn't decrement the TTL before handling the tunnelled packet. It may therefore be advisable for firewalls to be configured not to allow tunnelled traffic with an IPv6 TTL of 255 to pass, NIDSs should have signatures to detect this, and operating system vendors should verify that their code decrements the TTL of received tunnelled IPv6 packets immediately upon decapsulation, and before the decapsulated packet is processed.

## Covert Channels

60. There is no evidence that covert channels have been considered during the development of the IPv6 suite of protocols, with the end result that the protocols are 'leaky' at best. Renard<sup>42</sup> mentions one particular field - others are certainly present. There is furthermore very little evidence that much, if any, research has been performed into IPv6 covert channels, which implies there is currently a dearth of IDS signatures and the like to detect these. Due to transition mechanisms the entire protocol can itself be seen as a covert channel, as it is not IPv4-based and hence many IDS technologies and some misconfigured firewalls will just ignore it<sup>44</sup>.

## Network Security Technologies

61. The state of the marketplace for network security technologies is moving too rapidly for this document to make other than sweeping generalisations. The current status of network security products for IPv6 is rather poor, but

---

<sup>43</sup> See <http://www.gao.gov/cgi-bin/getrpt?GAO-05-471>

<sup>44</sup> See <http://seclists.org/lists/focus-ids/2002/Dec/0065.html>

improving. Both NIDS and firewalls are slightly behind the curve compared to operating system development and deployment of IPv6, although there are several commercial offerings that now support it. The quality of this support has not, to the best of the author's knowledge, been independently tested in sufficient depth yet. Based on the experience of IPv4 firewall and NIDS technologies, together with the increase complexity of IPv6 versus IPv4, it seems inconceivable that problems will not be found in these relatively immature products.

62. The increase in use of tunnelling during transition, use of IPSec, multiple and changing host addresses, and general focus on end-to-end networking will all require a modification of the traditional security architecture of border protection. IDSs and firewalls will still, of course, play a major role at the border; however it is also extremely likely that there will be an increased need for host-based IDSs<sup>45</sup> and firewalls. The ability of these to handle IPv6 traffic has not yet been proven

## Programming Interface

63. Obviously, the majority of networking programs are going to require modification in order to operate correctly on both IPv4 and IPv6. Considerable work has gone into techniques for easing this, with three separate iterations of RFCs for basic operation (ending with RFC3493<sup>46</sup>) and two for advanced operations (ending with RFC3542<sup>47</sup>).
64. The basic API provides definitions for AF\_INET6, sockaddr\_in6, etc, which are effectively a 1:1 mapping of IPv4 techniques. The largest change is the method for performing name to address translation - getaddrinfo and related functions are introduced. Alternately, use can be made of the more generic sockaddr\_storage, etc, techniques, also defined in this RFC. There are several places for programmers to make mistakes during porting.
65. Firstly, new functions such as getaddrinfo malloc data structures, which require extended freeing through functions such as freeaddrinfo. Otherwise, memory leaks will occur. This is more important for servers than clients.
66. It is important that programs are written to handle three items correctly: -
  - Allow enough space for text representations of IPv6 addresses (47 bytes minimum) where required - rather than the 16 bytes for IPv4,

---

<sup>45</sup> See <http://www.alchemistowl.org/arrigo/Papers/SPI2003-IDS-and-IPv6.pdf>

<sup>46</sup> See <http://www.ietf.org/rfc/rfc3493.txt?number=3493>

<sup>47</sup> See <http://www.ietf.org/rfc/rfc3542.txt?number=3542>

- Allow enough space for binary representations of IPv6 addresses (16 bytes) where required - rather than the 4 bytes for IPv4,
  - Handle the parsing of addresses correctly - use getaddrinfo instead of manually parsing,
67. If all three of these are correct, porting a simple program from IPv4 to IPv6/4 should not introduce any new issues. Programs should, if able to use both IPv4 and IPv6 at the network level, allow the user to be able to configure both, and display both, as applicable. This will lead to a requirement for applications to handle additional configuration options.
68. Whilst strictly related to transition mechanisms, both IPv4 compatible and 6to4 addressing (i.e. techniques for mapping IPv4 addresses onto IPv6 addresses) have some issues that impact on applications. Hagino describes these issues<sup>48</sup>, namely that end applications have no way to differentiate between IPv6 and IPv4 addresses - hence applications may use inappropriate addresses. For example IPv6 addresses which, when mapped into IPv4 address space, are inappropriate IPv4 addresses (such as broadcast addresses) may be used by an attacker, allowing them to be used in 'smurf' style DoS attacks.

## DNS

69. Whilst DNS as such is not part of the IPv6 specification, it plays an extremely large role in any non-trivial IPv6 network, and as such some discussion of DNS must be entered into. Two competing standards have been proposed in order to provide DNS resolution between IPv6 addresses and host names, including modification of the MX, NS and SRV record types. These are the AAAA (commonly termed quad-A), and A6 record types for forward mapping, with corollaries for reverse mapping.
70. The AAAA record was first defined in RFC1886<sup>49</sup>, which was obsoleted by the now current RFC3596<sup>50</sup>. This is an extremely simple technique, which is effectively identical to the IPv4 A record type. Reverse mapping is performed by use of an IP6.ARPA domain made up of the IPv6 address in reverse order represented as nibbles, dot-separated, similar to the IPv4 INADDR.ARPA. Due to this being the first proposed standard, it has received widespread acceptance and is supported in BIND<sup>51</sup> versions 4, 8 and 9.

---

<sup>48</sup> See "Possible abuse against IPv6 transition technologies" at <http://playground.iijlab.net/i-d/draft-itojun-ipv6-transition-abuse-01.txt>

<sup>49</sup> See <http://www.ietf.org/rfc/rfc1886.txt?number=1886>

<sup>50</sup> See <http://www.ietf.org/rfc/rfc3596.txt?number=3596>

<sup>51</sup> Berkeley Internet Name Domain. See <http://en.wikipedia.org/wiki/BIND>

71. The AAAA record was perceived as not being flexible enough for IPv6 in certain areas. With this in mind, the A6 record was proposed in RFC2874<sup>52</sup>, and updated in RFC3152<sup>53</sup>, RFC3226<sup>54</sup>, RFC3363<sup>55</sup>, and RFC3364<sup>56</sup>. The A6 record addresses this perceived lack of flexibility by using bit-string labels, DNAME records, and increasing usage of glue records, leading to significantly more choice and power, but at the cost of increased complexity.
72. Bit-string labels (RFC2673<sup>57</sup>) are a new way of treating arbitrary strings of bits as a hierarchical sequence of one-bit domain labels. Their use allows close linkage between the hierarchical addressing of IPv6 and that of DNS, together with allowing address delegation to be made on bit boundaries, rather than nibble boundaries as is the case for AAAA.
73. DNAME records provide similar functionality to that of CNAME records, but rename entire trees rather than individual hosts. Namely, they are a technique for aliasing or substituting addresses, and in some instances may significantly simplify management of DNS by allowing both forward and reverse mappings to be managed together.
74. It is inefficient and inelegant in the extreme to have two competing methods for using DNS with IPv6, and yet discussion of which was the better option appears to have been one of the more bitter and drawn-out of the exchanges on IPv6. This has certainly lead to delays in the widespread roll-out of IPv6 as the root DNS servers have understandably proven loath to roll out and support both versions, and hence have until recently rolled out neither. This is slowly changing, with 5 of the 13 root servers now supporting IPv6, and many of these supporting both A6 and AAAA.
75. As an attempt to aid the decision on whether to use A6 or AAAA, RFC3364 was written. This RFC concisely describes the pros and cons of each technique, concluding that if it is believed that rapid renumbering and GSE<sup>58</sup>-like routing will occur then A6 should be used, otherwise AAAA is preferable. It also recommends that bit-labelling should not be used, and use of DNAMEs should be minimised.

---

<sup>52</sup> See <http://www.ietf.org/rfc/rfc2874.txt?number=2874>

<sup>53</sup> See <http://www.ietf.org/rfc/rfc3152.txt?number=3152>

<sup>54</sup> See <http://www.ietf.org/rfc/rfc3226.txt?number=3226>

<sup>55</sup> See <http://www.ietf.org/rfc/rfc3363.txt?number=3363>

<sup>56</sup> See <http://www.ietf.org/rfc/rfc3364.txt?number=3364>

<sup>57</sup> See <http://www.ietf.org/rfc/rfc2673.txt?number=2673>

<sup>58</sup> "Global, Site, End system". An early proposed routing architecture. See <http://www.ietf.org/rfc/rfc3364.txt?number=3364>

76. From a security perspective, the impact of this confusion is obvious. Complexity and ambiguity are the enemy of security, and by having multiple DNS protocols for IPv6, one of which is very complex, there is bound to be negative security effects. Firstly, tried-and-trusted software (DNS server and client code) has had to be modified multiple times, possibly leading to the accidental insertion of vulnerabilities. By having two DNS protocols, the complexity of IDS and firewall technologies has also been increased, leading to an increased likelihood of vulnerabilities involving false-negatives and rule bypassing, respectively. It is hoped that a final firm decision will be made to deprecate one or the other, and that the relevant RFCs will be updated to show this, in a timely manner. It will then be possible for IDS rules and similar to be put in place, enforcing that only the lucky protocol be used at network boundaries. Additionally, the offending functionality can be removed from DNS clients and servers, reducing the risk imposed by inclusion of legacy code.

## Conclusions

77. The phrase IPv6 actually implies a significant number of interwoven protocols, only some of which have been discussed within this document. There are significant similarities between IPv6 and IPv4, at least as they have been defined within the relevant RFCs - although many of these areas of similarity are actually features within IPv4 that were never widely utilised. There are also areas of significant difference.
78. Development of IPv6 and its associated protocols has been a slow process, with RFCs being continually produced, updated, deprecated, obsoleted, and standardised. This has produced a large number of legacy versions of code and implementations, which in the short term may have negative security implications. It is hoped however that such a successive improvement of protocol specifications will in the longer term lead to a better security picture than would otherwise have been the case.
79. Security has been specifically considered many times during the evolution of the protocol, although often the result of these considerations has been that in order to improve security, IPSec should be used, sometimes without sufficient consideration of the difficulties involved in widespread use of IPSec, and the limitations of those devices most likely to benefit from Mobile IPv6. This will lead to these devices possibly not being as secure as would have been the case if non-cryptographic, or certainly non-IPSec, solutions had been more strongly sought. That being said however, with the increase in processing power in all devices it is probable that these problems will be overcome.
80. The increase in complexity of IPv6 and other related protocols will have a major effect on border defences of networks. The proliferation of addresses, increased focus on end-to-end protocols with associated changes in protocols, increased complexity of protocols such as the IPv6 Next Header chain, and use of transition technologies, will all make

traditional firewalling vastly more complex and less effective. This will therefore require an increased utilisation of Intrusion Detection and related technologies, and make well-designed internal network architectures more important. If networking is more end-to-end focussed, it is logical that defences against threats from network protocols must also be strengthened at the endpoints of communication, rather than at network boundaries - hence host security (whether HIDS<sup>59</sup>, host-based-firewalls, mandatory access controls, or otherwise) is very likely to be of greater import.

81. Work has progressed well in implementing IPv6 on host operating systems, such that all major suppliers now support IPv6. However, the lack of proliferation of IPv6 in the corporate and government marketplaces has led to firewall and IDS vendors ignoring the protocol until recently. That is now changing, perhaps due to increased interest from the world at large, or perhaps due to the perception of access to a large market in the form of the US Department of Defense, with their 2008 IPv6 deadline. In either case, it can be argued that IDS and firewall vendors are currently playing catch-up, and that there has been little by way of independent verification of the quality of many products compared to the situation with IPv4. The immaturity of this field currently affects security negatively. There is no doubt that the vendors will respond quickly to corporate interest though, and so this problem is likely to become less of an issue in the future, although the complexity of the protocols means that it is still likely that problems will be evident even in the longer term.
82. Fundamentally, due to its increased complexity, possible over-reliance on cryptographic safeguards, and relative immaturity, IPv6 has the potential to negatively impact network security. However, this is the case for any new protocol or technology, and the question is whether proper management and mitigation of these issues will reduce the risk to below that of the current protocols. Without knowledge of these risks, management and mitigation cannot be performed, and so it is imperative that awareness of these be spread, and that research into both current and future risks and mitigations be performed.

## Acknowledgments

83. This document is a summary of a large number of other documents. It has been attempted to give credit through references where due, but the large number of sources means that this is by no means complete, for which the author apologises, and hopes that they accept his thanks here.

---

<sup>59</sup> Host-based Intrusion detection System. "A Host-based Intrusion Detection System (HIDS), as a special category of an Intrusion-Detection System, focuses its monitoring and analysis on the internals of a computing system rather than on its external interfaces (as a Network Intrusion Detection System (NIDS) would do)."

See [http://en.wikipedia.org/wiki/Host-based\\_intrusion\\_detection\\_system](http://en.wikipedia.org/wiki/Host-based_intrusion_detection_system)