

IPv6 Security Brief

Last Updated: October 2011

Introduction

This brief paper summarizes the known threats and mitigation techniques for IPv6. Little explanation is given as it is assumed that the reader is familiar with IPv6 and with network security. A couple of "security myths" are debunked: reconnaissance is impossible and IPv6 is more secure thanks to IP Security (IPSec). IPv6 is shown as being roughly as secure as IPv4 (some aspects being more secure, some less secure) with a short-term temporary issue: the lack of IPv6 knowledge by network architects and security officers and lack of experience of running multiple protocols (IPv6 and IPv4) in parallel in the same network.

This document has three parts:

- Why IPv6 security is similar to IPv4 security
- What the security differences are between IPv6 and IPv4
- How to operate an IPv6 network in a secure way

IPv6 Is Similar to IPv4

IPSec Is Not the Savior

IPSec is considered an integral part of IPv6. Therefore, a common myth of IPv6 is that IPv6 is secure because IPSec is mandated. However, this is not really true:

- Although the IETF mandates that all IPv6 nodes have IPSec available, the actual use of IPSec is optional.
- If all communications between two IPv6 nodes are encrypted then the network (which is usually trusted because it is centrally managed) becomes blind and cannot inspect the traffic or enforce a security policy.

In short, IPSec on IPv6 should be reserved for the same cases as in IPv4: remote access virtual private networks (VPNs) or site-to-site VPNs.

Layer 2 Issues

Neighbor Discovery (ND), including stateless address autoconfiguration, suffers from the same lack of authentication as Address Resolution Protocol (ARP) and Dynamic Host Configuration Protocol (DHCP) on IPv4 networks. On IPv6 networks, ND spoofing attacks are real. The latter is often seen in IPv6 networks due to a misconfigured IPv6 host.

The mitigation techniques include:

- Using the network switches to block invalid ND packets (as done in IPv4 with DHCP snooping and dynamic ARP inspection). The solution from Cisco is called First Hop Security and is delivered in several phases, the first phase (RA-guard to protect against rogue Router Advertisement messages) is available since Summer 2010 and the other phases (including NDP inspection) is expected end of 2011 early 2012 (depending on the platform).

-
- Using a cryptographically improved ND protocol (secure neighbor discovery = SEND), which is implemented in Cisco IOS® routers but is not implemented in any popular host operating systems. Therefore, SEND is only useful to secure interrouter traffic. Cisco is working on a solution for switches using SEND even for non-SEND hosts.

Layer 3 Issues

Spoofing

Spoofing IPv6 addresses is as easy as spoofing IPv4 addresses (assuming that IPsec Authentication Header [AH] is not used in either of the protocol families). The mitigation technique is also identical: bogon filters (dropping obviously wrong source/destination addresses) and unicast reverse path forwarding checks. As the address space of IPv6 is larger, a purely random source address has a small probability of having a recognized prefix and hence a high probability of being dropped by the first router.

Source Routing

Since December 2007, source routing with routing header type 0 (RH0) is disabled by default in IPv6, it is therefore identical to IPv4. Source routing cannot be used by an attacker to bypass some security policies or to mount an amplification attack.

Routing header type 2 (RH2) is used by mobile IPv6 and cannot be used by attackers for any malevolent purpose.

Layer 4 and Above

Routing Protocols

Routing protocols used by IPv6 are either identical to the IPv4 ones or an improvement of those. This includes also the authentication mechanism that should be configured to prevent route hijacking. Open Shortest Path First Version 3 (OSPFv3) uses IPsec AH in transport mode rather than a Hash-based Message Authentication Code (HMAC) in the link-state advertisement (LSA), but this is equivalent.

Applications

All vulnerabilities in the application layer (such as SQL injection or cross-site scripting) are of course independent of the network layer. The attacks and mitigations techniques are strictly identical. The same applies for social engineering attacks such as phishing.

Network Denial of Service

IPv6 networks can be targets of Denial of Service (DoS) attacks:

- Sheer flooding of data plane packets: Attack on the available link utilization.
- Flooding the control or management planes: Attack on the IPv6 node resources such as CPU or memory.

Again, the detection and mitigation techniques are identical to the IPv4 world: using control plane policing and detecting flooding attacks with NetFlow or other techniques.

IPv6 Is Different from IPv4

Addresses

Much has been touted about the larger address space of IPv6, including some myths.

Larger Address Space: Reconnaissance

Another IPv6 security myth is that, thanks to the huge address space, an attacker cannot find targets by generating a random IPv6 address. While a random address will lead to nowhere, there are numerous ways for a malicious person to find targets: trying all Domain Name Service (DNS) names, assuming hexadecimal addresses with a textual meaning such as BAD:F00D or ABBA:BABE. The attacker can also use a previously broken into host by looking in its log files or connections tables.

Email or application worms and viruses also do not need to know anything about the network layer; so, they will happily propagate to all email addresses in your address book. The only difference with IPv4 is for the worms propagating at the network layer (such as Blaster): those worms will need to be modified in order to replace the IPv4 address scanning by other techniques as described above.

Again, there is no significant difference between IPv4 and IPv6.

Larger Address Space: Reputation

In the IPv4 world there are several databases providing a reputation for each IPv4 address; this is mainly used for tagging some email messages as spam.

In IPv6, there is no such database yet:

- The amount of IPv6 traffic in mid-2011 is mostly negligible and cannot be used to build such a database.
- The huge address space and different allocation sizes per subscribers (is it a /48 or a /56?) present technical challenges for the implementation.

No doubt that in a couple of years those two problems will be fixed.

Link-Local Address

The IPv6 link-local addresses (LLAs) have a very interesting security property because they cannot be reached from outside of the link, so there is no malicious traffic that can be sent to an LLA from a remote link.

A network operator can use this property on all links between routers by configuring only LLAs; routing protocols work with LLAs, and no remote attackers can launch an attack on these routers.

Privacy Extensions Address

In order to prevent the potential tracking of a host by identifying the host with its static EUI-64 Interface Identifier (IID) part of his IPv6 address, a host can use the privacy extension, which uses a random number as the IID. With privacy extension addresses, a host cannot be tracked any more as the address changes over time.

This is fine for a residential user but not acceptable for hosts inside a managed organization: the security/network operators must be able to track a malicious or misconfigured host within their network.

It is recommended to disable privacy extensions within a managed network.

Need for Topology Hiding

IPv4 Network Address and Port Translation (NAPT) is assumed to offer some security benefits:

- Only allows connections initiated from the inside;
- Hides internal hosts and topology from the outside world.

While the first benefit has little value in 2011 because most of the malware is downloaded through an inside-initiated web connection, there is some very minimal value in the topology hiding—minimal because it brings little useful value to the attacker (except perhaps for social engineering) and can also be derived from other information (DNS name in email headers, for instance, or observed values of the IP Hop-Limit field).

It is highly probable that firewall vendors will implement a function for an IPv6 equivalent of the security policy offered by NAT even though it breaks one important IPv6 benefit: end-to-end transparency. IPv6 Privacy Extension Addresses provide a similar function as NAT to hide the internal topology, and the policy to require flows to be initiated only from the inside is something that is already implemented in IPv6 firewalls.

Extension Headers

Parsing

The optional presence of variable-length extension headers (EHs) between the IPv6 headers and the Layer 4 header requires the parsing and understanding of all those headers in order to get access to Layer 4 information (for example to apply Quality of Service [QoS] or Access-Control List [ACL] filtering). Depending on the device implementation this can represent a performance impact.

Some early IPv6 firewalls were unable to parse this chain of extension headers. This is no more the case.

Fragmentation

Fragmentation in IPv6 has properties:

- Routers never fragment: this means that Path MTU discovery must work (free flow of Internet Control Message Protocol Version 6 [ICMPv6] unreachable packet-too-big is required) or there must be a careful setting of MTU or TCP MSS.
- Fragmentation can happen in the middle of an EH chain. If this happens, then the first fragment does not have any Layer 4 information. This case was forbidden in IPv4. So, policies will have to be extended to cover this very rare case.

Boundary Conditions

While some EHs can appear only once or twice, there is also a semidefined order within the EH chain. This means too many options and some IPv6 implementations suffered or still suffer from bugs to handle boundary conditions (such as an EH appearing twice while it should occur only once) leading to buffer overflow attacks or similar. The ASA firewall can be configured to drop packets with malformed EH chains.

Transition Mechanisms

Dual Stack

A dual-stack network is as secure as its weakest protocol family. This is called fate-sharing; for example, if the IPv6 access is not protected while the IPv4 is controlled, then the malicious user will use IPv6 for the attacks. It is really important to have congruent security policies for IPv4 and IPv6.

Tunnels

Tunnels can be convenient to transport IPv6 over an IPv4-only network. They can also be misused by the attacker to inject or to sniff IPv6 packets, to gain unauthorized access to an IPv6 network, and even to launch an amplification attack by looping between two tunnels. When tunnels are used to send sensitive traffic over a public network, they should be secured by adding IPsec authentication and confidentiality that can prevent both the

injection/sniffing attacks and unauthorized access. This is a specific case where IPsec is useful. The looping attack can only be mitigated by careful configuration of tunnel headends.

Latent Threat

The fact that modern hosts can be attacked over IPv6 even when connected to an IPv4-only network is called the IPv6 latent threat. Most recent host OSs have IPv6 enabled by default and some of them even try very hard to establish tunnels when there is no native IPv6 connectivity. If the host has not secured its IPv6 access (for example it has only an IPv4 firewall configured), then a link-local attacker can launch an attack on this host by sending a Router Advertisement [RA] message to trigger IPv6 stateless autoconfiguration on the target, or an off-link attacker can attack its victim over an automatic tunnel.

Actually, it is not linked to IPv6 but to the lack of knowledge of the security officers. Enabling native IPv6 on the network is the best way to counter this threat.

Operating an IPv6 Network

Enforcing a Security Policy

In 2011, all the security techniques (from firewalls and VPN to deep packet inspection) also have an IPv6 implementation (except for the IP address reputation) for a couple of years. Therefore, a security officer has no excuse for not securing the IPv6 part of the network.

Training

While malicious people already know about IPv6 and actively use it (for example for a botnet command and control channel or to escape an IPv4-only lawful interception), the good people are falling behind.

Training of security architects and operators to IPv6 networking (how to trace back an IPv6 address, how to understand log files, and so on) is really an urgent matter in 2011 if not already done.

Conclusion

IPv6 is mostly IPv4 with larger addresses and there is no significant difference between IPv4 and IPv6 with respect to security. In some cases (link-local addresses) IPv6 is slightly more secure, and in other cases (difficulties to parse the extension headers) IPv6 is slightly less secure. Some IPv6 security myths simply do not stand. Security techniques and devices do exist to enforce a security policy for the IPv6 traffic and should be used.

The lack of IPv6 training for network and security staff is probably the biggest threat for operation in 2011–2012.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)