Results of a Security Assessment of the Internet Protocol version 6 (IPv6)

Fernando Gont

Hack.lu 2011 Conference
Luxembourg, G. D. of Luxembourg, September 19-21, 2011

So... what is this “IPv6” thing about?

- IPv6 was developed to address the exhaustion of IPv4 addresses
- IPv6 has not yet seen broad/global deployment (current estimations are that IPv6 traffic is less than 1% of total traffic)
- However, general-purpose OSes have shipped with IPv6 support for a long time – hence part of your network is already running IPv6!
- Additionaly, ISPs and other organizations have started to take IPv6 more seriosly, partly as a result of:
  - Exhaustion of the IANA IPv4 free pool
  - Awareness activities such as the “World IPv6 Day”
  - Imminent exhaustion of the free pool of IPv4 addresses at the different RIRs
- It looks like IPv6 is finally starting to take off

Motivation for this presentation

- A lot of myths have been created around IPv6 security:
  - Security as a key component of the protocol
  - Change from network-centric to host-centric paradigm
  - Increased use of IPsec
  - etc.
- These myths have contributed to a general misunderstanding of the security properties of IPv6, thus negatively affecting the emerging (or existing) IPv6 networks.
- This presentation separates fudge from fact, and offers a more realistic view of “IPv6 security”
- Rather than delving into specific vulnerabilities, it is meant to influence the way in which you think about IPv6 security (and IPv6 in general).

Agenda

- Brief comparision of IPv6/IPv4
- Discussion of security aspects of IPv6
- Security implications of IPv6 transition/co-existence mechanisms
- Security implications of IPv6 on IPv4 networks
- Areas in which further work is needed
- Conclusions
- Questions & Answers

General considerations about IPv6 security

Some interesting aspects of IPv6 security

- There is much less experience with IPv6 than with IPv4
- IPv6 implementations are less mature than their IPv4 counterparts
- Security products (firewalls, NIDS, etc.) have less support for IPv6 than for IPv4
- The complexity of the resulting network will increase during the transition/co-existance period:
  - Two internetworking protocols (IPv4 and IPv6)
  - Increased use of NATs
  - Increased use of tunnels
  - Use of other transition/co-existance technologies
- Lack of well-trained human resources

…and even then, in many cases IPv6 will be the only option to remain in this business

Brief comparision between IPv6 and IPv4

(what changes, and what doesn't)

Brief comparision of IPv6 and IPv4

- IPv6 and IPv4 are very similar in terms of functionality (but not in terms of mechanisms)

	IPv4	IPv6
Addressing	32 bits	128 bits
Address resolution	ARP	ICMPv6 NS/NA (+ MLD)
Auto-configuration	DHCP & ICMP RS/RA	ICMPv6 RS/RA & DHCPv6 (optional) (+ MLD)
IPsec support	Optional	Mandatory (to "optional")
Fragmentation	Both in hosts and routers	Only in hosts

Security Implications of IPv6

IPv6 Addressing

Brief overview

- The main driver for IPv6 is its increased address space
- IPv6 uses 128-bit addresses
- Similarly to IPv4,
  - Addresses are aggregated into “prefixes” (for routing purposes)
  - There are different address types (unicast, anycast, and multicast)
  - There are different address scopes (link-local, global, etc.)
- It’s common for a node to be using, at any given time, several addresses, of multiple types and scopes. For example,
  - One or more unicast link-local address
  - One or more global unicast address
  - One or more link-local address

Global Unicast Addresses

- Syntax of the global unicast addresses:

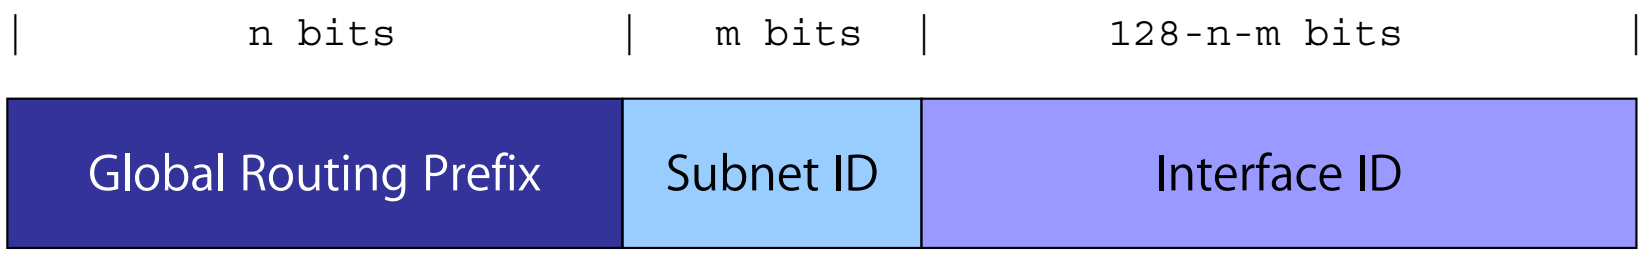

- The interface ID is typically 64-bis
- Global Unicast Addresses can be generated with multiple different criteria:
  - Use modified EUI-64 format identifierss (embed the MAC address)
  - “Privacy Addresses” (or some variant of them)
  - Manually-configured (e.g., 2001:db8::1)
  - As specified by some specific transition/co-existence technology

Implications on network reconnaissance

Myth: "The huge IPv6 address spaces makes brute-force scanning attacks impossible"

- This assumes host addresses are uniformly distributed over the subnet address space (/64)
- However, research (*) indicates that addresses do follow specific patterns:
  - SLAAC (Interface-ID based on the MAC address)
  - IPv4-based (e.g., 2001:db8::192.168.10.1)
  - "Low byte" (e.g., 2001:db8::1, 2001:db8::2, etc.)
  - Privacy Addresses (Random Interface-IDs)
  - "Wordy" (e.g., 2001:db8::dead:beef)
  - Related to specific transition-co-existence technologies (e.g., Teredo)

(*) Malone, D. 2008. *Observations of IPv6 Addresses*. Passive and Active Measurement Conference (PAM 2008, LNCS 4979), 29–30 April 2008.

Some real-world data….

- [Malone, 2008] (*) measured how IPv6 addreses are assigned to hosts and routers:

Hosts

Address Type	Percentage
SLAAC	50%
IPv4-based	20%
Teredo	10%
Low-byte	8%
Privacy	6%
Wordy	<1%
Other	<1%

Routers

Address Type	Percentage
Low-byte	70%
IPv4-based	5%
SLAAC	1%
Wordy	<1%
Privacy	<1%
Teredo	<1%
Other	<1%

(*) Malone, D. 2008. *Observations of IPv6 Addresses*. Passive and Active Measurement Conference (PAM 2008, LNCS 4979), 29–30 April 2008.

Some thoughts about network scanning

- Host scanning attacks have been found in the wild.
- As expected, they are less “brute-force” than their IPv4 counter-part.
- For servers,
  - Address predictability is irrelevant -- after all, you want them to be easily found.
- For hosts,
  - IPv6 “privacy addresses” are probably desirable.
  - Always consider the use of firewalls.

End-to-end connectivity

Brief overview

- Because of the increased IPv6 address space, it is expected that each device connected to the Internet will have a unique address.
- It is also assumed that this will "return" the "End-to-end Principle" to the Internet:
  - The network is transparent to the communication of any two nodes (e.g., intermmediate nodes do not modify the TCP port numbers, etc.)
  - Any node can establish a communication node with any other node in the network (e.g., the network does not filter "incoming connections")
  - It is usually argued that the "end-to-end principle" allows for Innovation.

Some considerations

- Even if each device receives a unique address, that does not necessarily imply “end-to-end” connectivity.
- In practice, most production networks don’t really care about innovation, but rather about getting work done.
- Users expect to use in IPv6 the same services currently available for IPv4 without “end-to-end” connectivity (web, email, social networks, etc.)
- Thus,
  - End-to-end connectivity is not necessarily a desired property in a production network (e.g., may increase host exposure unnecessarily)
  - A typical IPv6 subnet will be protected by a stateful firewall that only allows “return traffic”

Address Resolution

Brief overview

- IPv6 addresses are mapped to link-layer addresses by means of the "Neighbor Discovery" mechanism.
- ND is based on ICMPv6.
- ICMPv6 Neighbor Solicitations and Neighbor Advertisements are analogous to ARP requests and ARP replies, respectively.
- Being transported by IPv6, NS/NA messages may contain IPv6 Extension Headers, be fragmented, etc.
  - ARP is implemented directly over Ethernet, with no possibilities for Extension Headers or fragmentation.

Security considerations

- IPv4's ARP spoofing attacks can "ported" to IPv6
- Possible mitigation techniques:
  - Deploy SEND (SEcure Neighbor Discovery)
  - Monitor Neighbor Discovery traffic (e.g. with NDPMon)
  - Add static entries to the Neighbor Cache
  - Restrict access to the local network
- Unfortunately,
  - SEND is very difficult to deploy (it requires a PKI)
  - ND monitoring tools can be trivially evaded
  - Use of static Neighbor Cache entries does not scale
  - Not always is it possible to restrict access to the local network
- Conclusion: the situation is not that different from that of IPv4 (actually, it's a bit worse)

Auto-configuration

Brief overview

- There are two auto-configuration mechanisms in IPv6:
  - Stateless: SLAAC (Stateless Address Auto-Configuration), based on ICMPv6 messages (Router Solicitation y Router Advertisement)
  - Stateful: DHCPv6
- SLAAC is mandatory, while DHCPv6 is optional
- In SLAAC, "Router Advertisements" communicate configuration information such as:
  - IPv6 prefixes to use for autoconfiguration
  - IPv6 routes
  - Other configuration parameters (Hop Limit, MTU, etc.)
  - etc.

Security considerations

- By forging Router Advertisements, an attacker can perform:
  - Denial of Service (DoS) attacks
  - “Man in the Middle” (MITM) attacks
- Possible mitigation techniques:
  - Deploy SEND (SEcure Neighbor Discovery)
  - Monitor Neighbor Discovery traffic (e.g., with NDPMon)
  - Deploy Router Advertisement Guard (RA-Guard)
  - Restrict access to the local network
- Unfortunately,
  - SEND is very difficult to deploy (it requires a PKI)
  - ND monitoring tools can be trivially evaded
  - RA-Guard can be trivially evaded
  - Not always is it possible to restrict access to the local network
- Conclusion: the situation is not that different from that of IPv4 (actually, it’s a bit worse)

IPsec Support

Brief overview and considerations

Myth: "*IPv6 is more secure than IPv4 because security was incorporated in the design of the protocol, rather than as an 'add-on'*"

- This myth originated from the fact that IPsec support is mandatory for IPv6, but optional for IPv4
- In practice, this is irrelevant:
  - What is mandatory is IPsec support – not IPsec usage
  - And nevertheless, many IPv4 implementations support IPsec, while there exist IPv6 implementations that do not support IPsec
  - Virtually all the same IPsec deployment obstacles present in IPv4 are also present in IPv6
- The IETF has acknowledged this fact, and is currently changing IPsec support in IPv6 to "optional"
- Conclusion: there is no reason to expect increased use of IPsec as a result of IPv6 deployment

Security Implications of Transition/Co-existance Mechanisms

Brief overview

- The original IPv6 transition plan was dual-stack – it failed.
- Current strategy is a transition/co-existence based on a toolset:
  - Dual Stack
  - "Configured" Tunnels
  - Automatic Tunnels (ISATAP, 6to4, Teredo, etc.)
  - Translation (e.g., NAT64)
- Dual stack is usually enabled by default in most systems.
- Some automatic-tunnelling mechanisms (e.g. Teredo and ISATAP) are enabled by default in some systems (e.g., Windows Vista and Windows 7)

Security considerations

- Transition technologies increase the complexity of the network, and thus the number of potential vulnerabilities.
- Many of these technologies introduce “Single Points of Failure” in the network.
- Some of them have privacy implications:
  - Which networks/systems does your Teredo or 6to4 traffic traverse?
  - This may (or may not) be an important issue for your organization

Security considerations (II)

- Transition/co-existance traffic usually results in complex traffic (with multiple encapsulations).
- This increases the difficulty of performing Deep Packet Inspection (DPI) and (e.g. prevent the enforcement of some filtering policies or detection by NIDS).
- Example: structure of a Teredo packet.

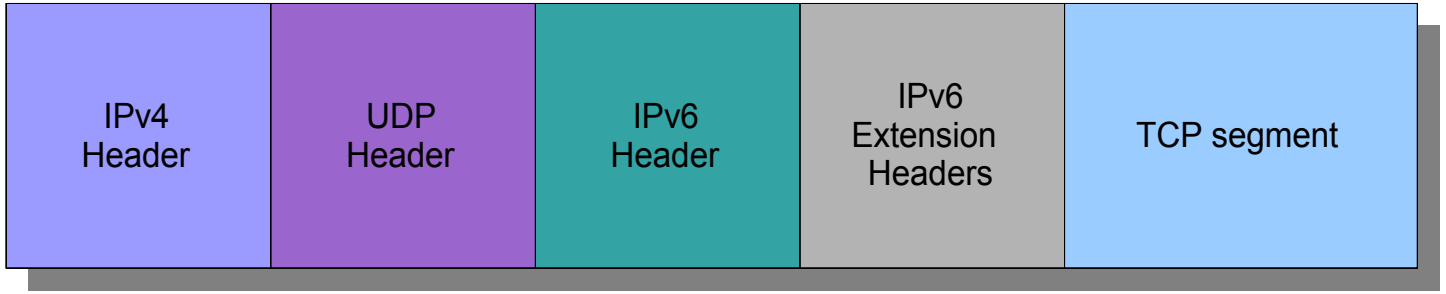

- "Exercise": write a libpcap filter to detect TCP/IPv6 packets transported over Teredo, and destined to host 2001:db8::1, TCP port 25.

Security Implications of IPv6 on IPv4 Networks

Brief overview

- Most general-purpose systems have some form of IPv6 support enabled by default.
- It may be in the form of “dual-stack”, and/or some transition/co-existence technology.
- This essentially means that an alledged “IPv4-only” network also include a partial deployment of IPv6.

Security considerations

- An attacker could readily enable the “dormant” IPv6 support at local nodes (e.g., sending ICMPv6 RAs), possibly evading network controls.
- Transition technologies such as Teredo could result in increased (and unexpected) host exposure.
- Thus,
  - Even if you don’t plan to “use” IPv6, you should consider its implications on your network.
  - If a network is meant to be IPv4-only, make sure this is actually the case.

Areas in which further work is needed

Key areas in which further work is needed

- IPv6 resiliency
  - Implementations have not really been the target of attackers, yet
  - Only a handful of publicly available attack tools
  - Lots of vulnerabilities and bugs still to be discovered.
- IPv6 support in security devices
  - IPv6 transport is not broadly supported in security devices (firewalls, IDS/IPS, etc.)
  - This is key to be able enforce security policies comparable with the IPv4 counterparts
- Education/Training
  - Pushing people to “Enable IPv6” point-and-click style is simply insane.
  - Training is needed for engineers, technicians, security personnel, etc., before the IPv6 network is running.

20 million engineers need IPv6 training, says IPv6 Forum
The IPv6 Forum - a global consortium of vendors, ISPs and national research & Education networks - has launched an IPv6 education certification programme in a bid to address what it says is an IPv6 training infrastructure that is "way too embryonic to have any critical impact.“ (**http://www.itwire.com**)

Some Conclusions

Some conclusions…

- Beware of IPv6 marketing and mythology!
- While IPv6 provides similar features than IPv4, it uses different mechanisms. – and the devil is in the small details
- The security implications of IPv6 should be considered before it is deployed (not after!)
- Most systems have IPv6 support enabled by default, and this has implications on “IPv4-only” networks!
- Even if you are not planning to deploy IPv6 in the short term, most likely you will eventually do it
- It is time to learn about and experiment with IPv6!

Questions?

Thank you!

Fernando Gont

fgont@si6networks.com

www.si6networks.com