Internet Engineering Task Force
Internet Draft
Expiration Date: December 2003

P. Savola
CSC/FUNET

June 2003

IPv6 Transition/Co-existence Security Considerations

draft-savola-v6ops-security-overview-00.txt

Status of this Memo

Abstract

The transition/co-existance from IPv4 to IPv4/IPv6 causes one to consider the security considerations of such a process. In this memo, I try to give an overview of different aspects relating to IPv6: the notion of increased end-to-end transparency, implications of tunneling, the use of IPv4-mapped addresses, the considerations of IPv6 service piloting without firewalls, IPv6 protocol-specific issues, IPv6 transition/co-existence mechanism -specific issues, consequences of enabling IPv6 by default, and operational security issues when enabling IPv6 in the network infrastructure.

Internet Draft draft-savola-v6ops-security-overview-00.txt June 2003

Table of Contents

1. Introduction

The transition/co-existance from IPv4 to IPv4/IPv6 causes one to consider the security considerations of such a process. In this memo, I try to give an overview of different aspects relating to IPv6: the notion of increased end-to-end transparency, implications of tunneling, the use of IPv4-mapped addresses, the considerations of IPv6 service piloting without firewalls, IPv6 protocol-specific issues, IPv6 transition/co-existence mechanism -specific issues, consequences of enabling IPv6 by default and operational security issues when enabling IPv6 in the network infrastructure.

The document is quite raw and not very well structured. Nevertheless, feedback is sought either on already mentioned issues (whether you feel it is important or not), or on new issues which haven't been mentioned.

2. Increased End-to-End Transparency

With IPv6, increased end-to-end transparency in general can sometimes be seen as a threat. Some seem to want limited end-to-end

capabilities, e.g. in the form of private, local addressing, even when it is not necessary.

People have gotten used to the perceived, dubious security benefits of NATs and perimeter firewalls, and the bidirectionality and

transparency that IPv6 should provide may seem undesirable at times.

[XXX: I'm not sure if this section is really worth mentioning here..]

3. Tunneling May Break Operational/Security Assumptions

NATs and firewalls have been deployed extensively in the IPv4 Internet, for the good or the bad. People who deploy them typically have some security/operational requirements in mind (e.g. a desire to block inbound connection attemps), whether misguided or not.

Tunneling can change that model. IPv6-over-IPv4 tunneling is typically explicitly allowed or disallowed implicitly. Tunneling IPv6 over IPv4 UDP, however, is often an entirely different thing: as UDP must usually be allowed through, at least in part and in a possibly stateful manner, one can "punch holes" in NAT's and firewalls using UDP.

One could say that tunneling is especially questionable in home/SOHO environments where the level of network administration is not that high; in these environments the hosts may not be as managed as in others (e.g. network services might be enabled unnecessarily), leading to possible security break-ins or other vulnerabilities.

Holes can be punched both intentionally and unintentionally. In case it is a willing choice from the administrator/user, this is less of a problem (but e.g. enterprises might want to block IPv6 tunneling explicitly if some employees would do something like this willingly on their own). On the other hand, if a hole is punched transparently, without people understanding the consequences, it may be a serious threat.

When deploying tunneling solutions, especially tunneling solutions which are automatic and/or can be enabled easily by users not understanding the consequences, care should be taken not to compromise the security assumptions held by the users.

4. IPv4-mapped IPv6 Addresses

Overloaded functionality is always a double-edged sword: it may yield some deployment benefits, but often also incurs the price which comes with ambiguity.

One example of such is IPv4-mapped IPv6 addresses: a representation of IPv4 address as an IPv6 address inside an operating system. Since then, IPv4-mapped addresses have been extended to be used with a transition mechanism [SIIT], on the wire.

Therefore, it becomes difficult to unambiguously discern whether an IPv4 mapped address is really an IPv4 address represented in the IPv6 address format *or* an IPv6 address received from the wire (which may be subject to address forgery, etc.).

In addition, special cases like these, while giving deployment benefits in some arenas, require some amount of code complexity (e.g. in the implementations of bind() system calls) which we might be better off without [V4MAPPEDA] [V4MAPPEDW].

5. IPv6 Service Piloting Done Insecurely

In many cases, IPv6 service piloting is done in a manner which is considered to be less secure than as one would do with IPv4. For example, hosts and routers might not be protected by IPv6 firewalls, even if in IPv4 firewalls are being used.

The other possible alternative, in some places, is that no service piloting is done because IPv6 firewalls aren't being used -- and IPv6 deployment suffers (of course, this is also one of the nice excuses for not doing IPv6).

This problem may be partially due to a slow speed of IPv6-capable firewall development and deployment. However, it is also a problem with a lack of information: actually, there are quite a few IPv6 packet filters and firewalls already, which could be used for sufficient access controls, but network administrators may not be aware of them yet.

However, there appears to be a real lack in two areas: "personal firewalls" and enterprise firewalls; the same devices that support and are used for IPv4 today are often expected to also become IPv6-capable (even though this is is not really necessary).

Another, smaller factor may be that due to a few decisions on how IPv6 was built, it's more difficult for firewalls to be implemented and work under all the cases (e.g. when new extension headers etc. are used) [FW]: it's a bit more difficult for intermediate nodes to process the IPv6 header chains than IPv4 packets.

6. IPv6 Protocol-specific Issues

Some features of IPv6 are a bit different from IPv4, and may include some potential problems specification-wise.

Some examples come to mind right-out:

- how hosts should interact with routing headers (they must act as forwarders) [RHHOSTS]
- how routing headers may be too generic contructs to be useful for e.g. MIPv6 purposes [RHHAOSEC]
- how home address option was previously specified (fixed now) [RHHAOSEC]
- how ICMPv6 messages, in some cases, may be generated in response to multicast packets (where in IPv4 they can't) [FW]
- how the privacy IPv6 addresses may not actually provide all that much privacy (ie. the applicability is unclear) [3041HARM]

- how IPv6 has been specified wrt. middleboxes such as firewalls (e.g. when new extension headers etc. are used) [FW]

7. IPv6 Transition/Co-existence Mechanism-specific Issues

The more complicated the IPv6 transition/co-existence becomes, the more danger there is to introduce security issues in the mechanisms (which may or may not be readily apparent). Therefore it would be desirable to keep the mechanisms simple, in as small pieces as possible.

One case where such security issues have been analyzed is [6TO4SEC].

As tunneling has been proposed as a model for quite a bit more cases than are currently being used, its security properties should be analyzed in more detail. There are some generic dangers to tunneling:

- it may be easier to avoid ingress filtering checks
- it is possible to attack the tunnel interface: several IPv6 security mechanisms depend on checking that Hop Limit equals 255 on receipt and that link-local addresses are used. Sending such packets to the tunnel interface is much easier than gaining access to a physical segment and sending them there.

- automatic tunneling mechanisms are typically particularly dangerous as the other end-point is unspecified, and packets have to be accepted and decapsulated from everywhere. Therefore, special care should be observed when specifying automatic tunneling techniques.

8. Enabling IPv6 by Default Brings the Usability Down

A practical disadvantage of enabling IPv6 at the moment is that it typically brings the observed service level down a bit; that is, the usability suffers.

This is due to at least three reasons:

- global IPv6 routing is still rather unstable, leading to packet

loss, lower throughput, and higher delay [6BMESS]

- some applications can not properly handle both IPv4 and IPv6 or may have problems handling all the fallbacks and failure modes (and in some cases, e.g. if the TCP timeout kicks in, this may be very difficult)

- some DNS server implementations have flaws that severely affect DNS queries for IPv6 addresses [DNSA4]

Actually, some would be 100% ready to release IPv6 services (e.g. web) today, but that would mean trouble for many of their dual-stacked customers or users all over the world so they don't: these are often published under a separate domain or subdomain, and are practically not used that often.

9. Operational Factors when Enabling IPv6 in the Network

You have to be careful when enabling IPv6 in the network gear for multiple reasons:

IPv6-enabled router software may be unstable(r) yet; either IPv6 is unstable, or the software you have to run to be able to run IPv6 is different (from non-IPv6 parts) from the one you would run otherwise, making the software in practice more unstable -- and raising the bar for IPv6 adoption.

IPv6 processing may not happen at (near) line speed (or in the same level as IPv4). A high amount of IPv6 traffic (even legitimate, e.g. NNTP) could easily overload the software-based IPv6 processing and cause harm to IPv4 processing too.

Sometimes required features may be missing from the vendors' software releases; an example is a software enabling IPv6 telnet/SSH access, but having no ability to turn it off or limit access to it!

Sometimes the default IPv6 configuration is insecure. For example, in one vendor, if you've restricted IPv4 telnet to only a few hosts in the configuration, you need to be aware that IPv6 telnet will be automatically enabled, that the configuration commands used

previously do not block IPv6 telnet, IPv6 telnet is open to the world by default, and that you have to use a separate command to lock down the IPv6 telnet access.

Many operator networks have to run interior routing protocols for both IPv4 and IPv6. It's possible to run the both in one routing protocol, or have two separate routing protocols; either approach has its tradeoffs. If multiple routing protocols are used, one should note that this causes double the number of processing when links flap or recalculation is otherwise needed -- which might more easily overload the routers' CPU, causing slightly slower convergence time.

10. Acknowledgements

Your name might end up here :-)

11. Security Considerations

This memo tries to give an overview of security considerations of the different aspects of IPv6.

12. References

12.1. Informative

[3041HARM] Dupont, F., Savola, P., "RFC 3041 Considered Harmful", draft-dupont-ipv6-rfc3041harmful-02.txt, work-in-progress, January 2003.

[6BMESS] Savola, P., "Moving from 6bone to IPv6 Internet", draft-savola-v6ops-6bone-mess-01.txt, work-in-progress, November 2002.

[6TO4SEC] Savola, P., "Security Considerations for 6to4", draft-savola-v6ops-6to4-security-02.txt, work-in-progress, January 2003.

[DNSA4] Morishita., Y., Jinmei, T., "Common Misbehavior against DNS Queries for IPv6 Addresses", draft-morishita-dnsop-misbehavior-against-aaaa-00.txt,

work-in-progress, June 2003.

[FW] Savola, P. "Firewalling Considerations for IPv6", work-in-progress, draft-savola-v6ops-firewalling-01.txt, March 2003.

[RHHAOSEC] Savola, P. "Security of IPv6 Routing Header and Home Address Options", work-in-progress, draft-savola-ipv6-rh-ha-security-03.txt, December 2002.

[RHHAOSEC] Savola, P. "Note about Routing Header Processing on IPv6 Hosts", work-in-progress, draft-savola-ipv6-rh-hosts-00.txt, February 2002.

[SIIT] Nordmark, E., "Stateless IP/ICMP Translation Algorithm", RFC276, February 2000.

[V4MAPPEDA] Metz, C., Hagino, J., "IPv4-Mapped address API considered harmful", draft-cmetz-v6ops-v4mapped-api-harmful-00.txt, work-in-progress, Oct 2002.

[V4MAPPEDW] Metz, C., Hagino, J., "IPv4-Mapped Addresses on the Wire Considered Harmful", draft-itojun-v6ops-v4mapped-harmful-01.txt, work-in-progress, Oct 2002.

Author's Address

Pekka Savola
CSC/FUNET
Espoo, Finland
EMail: psavola@funet.fi

Intellectual Property Statement

Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.