Internet Engineering Task Force
INTERNET DRAFT
Expires in July 2003

Francis Dupont
ENST Bretagne
Pekka Savola
CSC/FUNET
January 2003

RFC 3041 Considered Harmful

<draft-dupont-ipv6-rfc3041harmful-02.txt>

Status of this Memo

Abstract

The purpose of the privacy extensions for stateless address autoconfiguration [1] is to change the interface identifier (and the global-scope addresses generated from it) over time in order to make it more difficult for eavesdroppers and other information collectors to identify when different addresses used in different transactions actually correspond to the same node.

Current Distributed Denial of Service (DDoS) [2] attacks employ forged source addresses which can also be in the same prefixes than the real addresses of the compromised nodes used for attacks. Indeed, network ingress filtering defeats DDoS using "random" forged source addresses.

The issue developed in this document is that the behavior of a compromised node used as source in a DDoS attack with "in-prefix" spoofed source address and the behavior of nodes using temporary addresses at high rate can not be distinguished. This could make future defenses against DDoS attacks very hard.

1. Introduction

Last IPv6 addressing architecture document [3] defines the modified EUI-64 format for interface identifiers. This format is mandatory for all unicast addresses, except those that start with binary value 000 and is 64 bit long with two special bits:

- the universal/local "u" bit which indicates whether the scope of the identifier is global or local.
- the individual/group "g" bit inherited from IEEE specification.

In practice interface identifiers enter in one of these categories:

- global scoped identifiers derived from a built-in interface hardware identifier like an IEEE MAC-48 address.
- manually assigned small identifiers (::1, ::2, ...) which have, of course, a local scope.
- randomly generated identifiers (with a local scope, used when the first category of identifiers raises a privacy concern)
- identifiers derived from a key like Statically Unique and Cryptographically Verifiable identifiers [5] (also with a local scope but bound to a key with a provable ownership).

The RFC 3041 (privacy extensions) [1] defines the management of randomly generated identifiers and, in the real world, all of them.

Interface identifiers are used in the stateless address autoconfiguration [4] to create link-local addresses (in all cases) and to create global and site-local addresses (for hosts from prefix information options in router advertisements).

2. Privacy Extensions

The privacy extensions document addresses claimed privacy concerns with globally unique and/or persistent interface identifiers.

The basic issue is when a constant identifier is reused over an extended period of time and in multiple independent activities, it becomes possible for that identifier to be used to correlate seemingly unrelated activity.

Note that the interface identifier is usually only the half of the whole address, and to change the interface identifier when the prefix remains the same will not improve the privacy.

There are only two cases where privacy extensions can be justified: where the link has a very high number of nodes or where the link (and the prefix(es)) changes from time to time. In the second case a fresh interface identifier gives a whole new address which can not be tracked, but an interface identifier change between two movements should give only complexity with little benefit.

How little benefit the is to be had depends mainly on how frequently the prefix changes. For example, if ISPs assign a static prefix to a customer, changing the interface identifier never really helps. However, if ISPs assign a dynamic prefix to a customer, how often the prefix changes (compared to the rate at which new interface identifiers are generated) restricts the applicability of the privacy extensions. For example, if the prefix changes about once a month, practically the users will be trackable; on the other hand, if the prefix changes once a day it could be argued the privacy extensions could provide some extra privacy in that timeframe. In the latter case, the goal is to have a dynamic prefix out of a large dynamic prefix address block, so it is unnoticeable to the observers when a different user from the bigger address block is using the prefix under observation and when the same user has generated a new interface identifier.

Our argument is that in the second case the prefix(es) and the interface identifier should be changed at the same time, or at least that the prefix(es) should be changed often enough when compared to the interface identifier change frequency.

3. "In-Prefix" Source Addresses Spoofing

Distributed Denial of Services (DDoS) attacks are a variant of Denial of Service attacks where the attackers use a large number of compromised hosts in poorly managed domains to flood aimed targets with forged packets. In some cases, the amount of traffic is enough to overload network infrastructures near the targets.

In order to hide the real addresses of compromised hosts, to defeat easy defenses like rate limitation on detected flows, to avoid returned traffic, etc, DDoS attacks employ forged, rapidly changing source addresses. When spoofed source addresses are randomly chosen, ingress filtering [2] can check if they are topologically plausible and drop forged packets. Ingress filtering, especially based on unicast Reverse Path Forwarding Forwarding (uRPF) checking, has been enough deployed in some networks in the today Internet to encourage the attackers to also find different ways to spoof addresses to perform effective DDoS attacks.

But ingress filtering is not effective against "in-prefix" source address spoofing where forged addresses are derived from real ones by only changing the last bits so they are likely to be topologically correct. Administrators of systems under attacks have the choice between accepting some traffic from fake sources and filtering out too much traffic including legitimate traffic from close to the apparent attack source, i.e. meaning a denial of service for those legitimate sources. Of course, IPv6 gives the attackers even more bits to play with (64 bits for a link, 80 for a site); practically e.g. rate limiting by address must be changed to rate limiting by prefix.

To summarize, filtering works only when it is possible and/or easy to recognize legitimate packets from forged packets. In some cases attacks can be detected at some places (it should always be the case near the targets) but again defensive actions need a good selection criterion or they become themselves denial of service attacks.

4. Temporary vs. Forged Source Addresses

Privacy extensions create new temporary addresses with an additive rate, i.e. with 1000 nodes and a rate by node of one new temporary address per day (the default rate [1]) the resulting rate is one new address every 90 seconds. So, where changing the temporary addresses makes sense, the uses of temporary or forged addresses are very hard to distinguish.

Of course, solutions like per address network access control and outbound traffic filtering are both unlikely in poorly managed sites where the attackers find hosts to compromise, and are not very compatible with user privacy concerns.

So we recommend:

- the use of temporary addresses should be disabled by default (as in the revision of RFC 3041 [6]).
- implementations should be updated as soon as possible when their default is to use temporary addresses.
- next revisions of RFC 3041 should address the tradeoff between temporary and forged addresses.
- schemes for cryptographically generated addresses (CGAs) should take the issue described in this document into account.
- A new revision of RFC 3041 should be finished as soon as possible and released to update RFC 3041 to avoid further damage.

4. Security Considerations

This document proposed to fill the Security Considerations section of revisions [6] of RFC 3041 which is currently empty.

Cryptographically generated addresses (CGAs) are by definition verifiable but verifying a CGA can be an expensive operation and if different levels of verification are possible, levels which provide good trust are likely to be more expensive. So if a network access control should check CGAs, the design must avoid to transform it into an easy target for Denial of Service attacks.

It should also be noted that there are a lot more ways to collect privacy information than watching the addresses (e.g. User Agent logging in HTTP [7]); these may not enough to make conclusions about the node in itself, but could be used, in part, when trying to tell whether two addresses might belong to the same node.

One can argue the usage of privacy features should be unobservable [8].

5. Acknowledgments

The nature of current DDoS attacks was described by Stanislav Shaluno during an ingress filtering and home address option thread in mobile-ip and ipv6 IETF WG mailing-lists.

Thomas Narten and Richard Draves tried to explain exactly what kind of privacy temporary addresses can (not) provide. Unfortunately this answer to complaints about IEEE derived interface identifiers and privacy is not IMHO technically far more well-founded than the complaints themselves; but there was not the time for a real anonymity device (the future work section of the RFC 3041 revision [6] finishes by the same kind of considerations).

Alberto Escudero-Pascual suggested to have a look on the observability of privacy extensions [8].

6. Normative References

[1] T. Narten, R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

[2] P. Ferguson, D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", RFC 2827 / BCP 38, May 2000.

[3] R. Hinden, S. Deering, "IP Version 6 Addressing Architecture", draft-ietf-ipngwg-addr-arch-v3-11.txt (update of RFC 2373), October 2002.

[4] S. Thomson, T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

7. Informative References

[5] G. Montenegro, C. Castelluccia, "SUCV Identifiers and Addresses", draft-montenegro-sucv-03.txt, July 2002.

[6] T. Narten, R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", revision of RFC 3041, draft-ietf-ipngwg-temp-addresses-v2-00.txt (expired), July 2001.

[7] R. Fielding, et al., "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2068, January 1997.

[8] A. Escudero, "Requirements for unobservability of privacy extension in IPv6", RVK02, Stockholm, June 2002.

8. Author's Addresses

Francis Dupont
ENST Bretagne
Campus de Rennes
2, rue de la Chataigneraie
BP 78
35512 Cesson-Sevigne Cedex
FRANCE
Fax: +33 2 99 12 70 30
EMail: Francis.Dupont@enst-bretagne.fr

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