



Adventures with IPv6 and Scapy **or** The Human Element in Protocol Design and Research

Enno Rey, ERNW / ScapyCon 2026





Internet Engineering Task Force (IETF)
Request for Comments: 9099
Category: Informational
ISSN: 2070-1721

É. Vyncke
Cisco
K. Chittimaneni

M. Kaeo
Double Shot Security
E. Rey
ERNW
August 2021

Operational Security Considerations for IPv6 Networks

Abstract

Knowledge and experience on how to operate IPv4 networks securely is available, whether the operator is an Internet Service Provider (ISP) or an enterprise internal network. However, IPv6 presents some new security challenges. RFC 4942 describes security issues in the protocol, but network managers also need a more practical, operations-minded document to enumerate advantages and/or disadvantages of certain choices.

This document analyzes the operational security issues associated with several types of networks and proposes technical and procedural mitigation techniques. This document is only applicable to managed networks, such as enterprise networks, service provider networks, or managed residential networks.



Black Hat US 2014

Evasion of High-End IDPS Devices at the IPv6 Era

Antonios Atlasis¹

Enno Rey²

Rafael Schaefer^{3,4}

1. secfu.net, aatlasis@secfu.net

2. ERNW GmbH, erey@ernw.de


3. ERNW GmbH, rschaefer@ernw.de

4. Bonn-Rhein-Sieg University of Applied Sciences

Protocol Politics

THE GLOBALIZATION OF INTERNET GOVERNANCE

LAURA DENARDIS



Technical Criteria for Choosing
IP The Next Generation (IPng)

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document was submitted to the IPng Area in response to RFC 1550. Publication of this document does not imply acceptance by the IPng Area of any ideas expressed within. Comments should be submitted to the big-internet@munnnari.oz.au mailing list. This RFC specifies criteria related to mobility for consideration in design and selection of the Next Generation of IP.

Table of Contents

1.	Introduction.	2
2.	Goals	3
3.	Note on Terminology	4
4.	General Principles.	4
4.1	Architectural Simplicity.	4
4.2	One Protocol to Bind Them All	4
4.3	Live Long	5
4.4	Live Long AND Prosper	5
4.5	Co-operative Anarchy.	5
5.	Criteria.	6
5.1	Scale	7
5.2	Topological Flexibility	8
5.3	Performance	9
5.4	Robust Service.	10
5.5	Transition.	12
5.6	Media Independence.	13
5.7	Unreliable Datagram Service	15
5.8	Configuration, Administration, and Operation.	16
5.9	Secure Operation.	17

Network Working Group
Request for Comments: 1883
Category: Standards Track

S. Deering, Xerox PARC
R. Hinden, Ipsilon Networks
December 1995

Internet Protocol, Version 6 (IPv6) Specification

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document specifies version 6 of the Internet Protocol (IPv6), also sometimes referred to as IP Next Generation or IPng.

Pls design an Internet Protocol similar to IPv4 but with more address space

How would you like the design delivered?

Spec document

What address width should the protocol use?

64-bit

How closely should it track IPv4's design philosophy?

Stay close to IPv4

Internet Protocol, 64-bit Addressing (IP64)

****Protocol Specification – Draft 1****

Author: Enno Rey · September 2026 · Status: Experimental

Abstract

IP64 is a network-layer protocol derived directly from the Internet Protocol version 4 (RFC 791). It retains IPv4's packet model, header semantics, fragmentation mechanism, option format, header checksum, and time-to-live handling, and changes one thing of substance: the source and destination address fields are widened from 32 bits to 64 bits, raising the address space from 2^{32} ($\approx 4.3 \times 10^9$) to 2^{64} ($\approx 1.8 \times 10^{19}$) addresses.

The design goal is that any correct IPv4 implementation can be turned into an IP64 implementation by changing the address type and the header length constants, and that a stateless IPv4↔IP64 translator can be written in a few dozen lines. Everything IPv4 got right is kept; the address shortage is fixed; nothing else is redesigned.

1. Introduction

1.1 Motivation

IPv4's 32-bit address space was exhausted at the registry level between 2011 and 2019. The deployed remedy, IPv6, solved the address problem but bundled it with a substantially different header, a new neighbor discovery protocol, new autoconfiguration semantics, and a new operational model, which is a large part of why adoption took decades. IP64 explores the

Network Working Group
Request for Comments: 1883
Category: Standards Track

S. Deering, Xerox PARC
R. Hinden, Ipsilon Networks
December 1995

Internet Protocol, Version 6 (IPv6) Specification

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract


This document specifies version 6 of the Internet Protocol (IPv6), also sometimes referred to as IP Next Generation or IPng.



1st IPv6 Security Tooling (22C3, 2005)



3. ICMPv6 Stateless Auto-Configuration



1. RS



fake_router6
Sets any IP as
default router



2. RA


1. RS:
ICMP Type = 133
Src = ::
Dst = FF02::2
query= please send RA

2. RA:
ICMP Type = 134
Src = Router Link-Local Address
Dst = FF02::1
Data= options, prefix, lifetime,
autoconfig flag

Routers send periodic as well as soliticated Router Advertisements (RA) to the all-nodes multicast address FF02::1

Clients configure their routing tables and network prefix from advertisements. => Like a DHCP-light in IPv4

But anyone can send Router Advertisements! => fake_router6



Internet Engineering Task Force (IETF)
Request for Comments: 6105
Category: Informational
ISSN: 2070-1721

E. Levy-Abegnoli
G. Van de Velde
Cisco Systems
C. Popoviciu
Technodyne
J. Mohacsi
NIIF/Hungarnet
February 2011

IPv6 Router Advertisement Guard

Abstract

Routed protocols are often susceptible to spoof attacks. The canonical solution for IPv6 is Secure Neighbor Discovery (SEND), a solution that is non-trivial to deploy. This document proposes a light-weight alternative and complement to SEND based on filtering in the layer-2 network fabric, using a variety of filtering criteria, including, for example, SEND status.



Yet another update on IPv6 security – Some notes from the IPv6-Kongress in Frankfurt

A couple of hours ago Christopher (Werny) and I gave [this presentation](#) well-organized event bringing together a number of practitioners from the field. While yesterday's talks were dominated by a certain euphoria and optimistic pioneer spirit, the second day featured some security talks which induced slight shadows to the brave new world of IPv6 ;-). I particularly enjoyed meeting Eric Vyncke from Cisco (one of the two authors of [this great book](#)) and Marc “van Hauser” Heuse who released a new version of the [THC-IPV6 tool set](#) today. We had some fruitful discussions and we took the opportunity to test some of his newly implemented attacks against “RA Guard” running on a 4948E Chris and I had brought for a demo within our talk. Unfortunately – or fortunately in terms of a [“from theory to reality”](#) approach – I have to say that Marc found a quite clever way to circumvent RA Guard by putting the actual “RA payload” into a second frame following a first one mostly containing a “long & empty” destination option (after a fragmentation header pointing to the mentioned second one). To get an idea pls see these screenshots from Wireshark.

4	0.018042	fe80::218:68ff:fe23ff02::1	ICMPv6 Router advertisement from 00:18:68:23:65:f6
5	0.031489	fe80::218:73ff:fe1eff02::1	IPv6 IPv6 fragment (nxt=IPv6 destination option (0x3c) off=0 id=0x4dccf335) [Malformed Packet]
6	0.031724	fe80::218:73ff:fe1eff02::1	ICMPv6 Router advertisement from 00:18:73:1e:de:4f
7	0.046780	fe80::218:41ff:fe54ff02::1	IPv6 IPv6 fragment (nxt=IPv6 destination option (0x3c) off=0 id=0x4dccf335) [Malformed Packet]
8	0.047018	fe80::218:41ff:fe54ff02::1	ICMPv6 Router advertisement from 00:18:41:54:56:b0
9	0.062635	fe80::218:fff:fe2e:ff02::1	IPv6 IPv6 fragment (nxt=IPv6 destination option (0x3c) off=0 id=0x4dccf335) [Malformed Packet]
10	0.062857	fe80::218:fff:fe2e:ff02::1	ICMPv6 Router advertisement from 00:18:0f:2e:f3:74

CVE-2012-2744



Public on July 10, 2012

Last modified on June 27, 2026,
1:49:39 PM UTC

Jump to

DescriptionStatementAffected PackagesCVSS Score

Description

From CVE.org

net/ipv6/netfilter/nf_conntrack_reasm.c in the Linux kernel before 2.6.34, when the nf_conntrack_ipv6 module is enabled, allows remote attackers to cause a denial of service (NULL pointer dereference and system crash) via certain types of fragmented IPv6 packets.

IPv6 Extension Headers


Problem 1: Too Many Things to Vary

- Variable types
- Variable sizes
- Variable order
- Variable number of occurrences of each one.
- Variable fields



IPv6 = f(v,w,x,y,z,)

The role of Scapy





RFC 6980

(or 'one passionate engineer's fight against the protocol priesthood')

5. Specification

Nodes MUST NOT employ IPv6 fragmentation for sending any of the following Neighbor Discovery and SEcure Neighbor Discovery messages:

- o Neighbor Solicitation
- o Neighbor Advertisement
- o Router Solicitation
- o Router Advertisement
- o Redirect
- o Certification Path Solicitation

Geistige Arbeit als Beruf

Vorträge vor dem
Freistudentischen Bund

Erster Vortrag
Prof. Max Weber
(München)

Wissenschaft als Beruf
(Referat: Seite 22)



München und Leipzig
Verlag von Duncker & Humblot
1919

A10/64
-1-

Max Weber (in Paraphrase von Hartmut Rosa):

"AI kann uns nicht sagen, wie wir leben und welche Werte wir verfolgen wollen."

Max Weber (paraphrased by Hartmut Rosa):

"AI cannot tell us how to live, nor which values are worth pursuing."



Bildnachweise/Quellen:

Folie 01 Illustration: iStockphoto/CAPToro

Folie 13 https://www.youtube.com/watch?v=ERaXv3k1_S8

Folie 14 https://media.ccc.de/v/22C3-772-en-attacking_ipv6

Folie 21 https://commons.wikimedia.org/wiki/File:Rosales_-_Sir_Isaac_Newton's_Apple_-_1.jpg

Folie 24 Zitat aus <https://www.sueddeutsche.de/kultur/wahl-sachsen-anhalt-afd-hartmut-rosa-interview-li.3541717?reduced=true>

Hindsight is 20/20

From a security perspective, which things should have been done differently in the initial IPv6 specification as of RFC 1883?

Looking at RFC 1883 (December 1995) with thirty years of hindsight, most of the security problems trace back to a few architectural choices rather than to any single field. The IETF has spent the intervening decades patching each of them individually, which is itself the strongest evidence that the original design should have been tighter.

The biggest one is the extension header model. RFC 1883 made the header chain an unbounded, arbitrarily ordered linked list of variable-length blocks, without a common TLV format for unknown headers (that only came with RFC 6564 in 2012), without any cap on chain length or total size, and without requiring that the whole chain fit into the first fragment (RFC 7112, 2014). The consequence was that a stateless device — a firewall, a RA-Guard implementation, an ACL in a router's fast path — could not reliably find the upper-layer protocol, and the spec gave it no principled basis to refuse a packet it could not parse. Practically every extension-header evasion technique published from 2011 onwards (Atlasis, Gont, your own group's work) exploits exactly this gap. A fixed order, a length ceiling, a uniform header format, and an explicit "first fragment must contain the full chain and the transport header" rule would have cost nothing in 1995.