ERNW
providing security.

OS IPv6 Behavior in Conflicting Environments

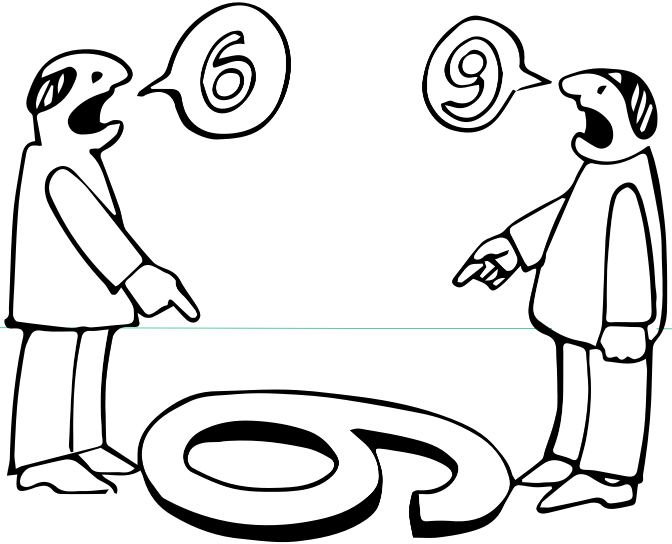

Enno Rey & Christopher Werny
{erey,cwerny}@ernw.de

Who Are We

- Old-school network guys with some background in large scale operations.
- Involved with IPv6 since a loong time and regularly blogging about IPv6 at `www.insinuator.net`.

Agenda

- Fundamentals
  - Quick Refresher of Basics & Specifications
- Results from the Lab
  - Some Surprises (?)
- Conclusions
  - What All this Means from an Operations Perspective

Related Work

https://tools.ietf.org/id/draft-ietf-v6ops-dhcpv6-slaac-problem-04.txt

- [draft-ietf-v6ops-dhcpv6-slaac-problem]
  - DHCPv6/SLAAC Interaction Problems on Address Auto-configuration. draft-ietf-v6ops-dhcpv6-slaac-problem-04
- [draft-droms-dhcpv6-issues]
  - Issues Concerning DHCP in IPv6 Specifications. draft-droms-dhcpv6-issues-00
  - Expired Apr 27 2003 (!)
  - https://tools.ietf.org/id/draft-droms-dhcpv6-issues-00.txt

Fundamentals

What the textbook tells you

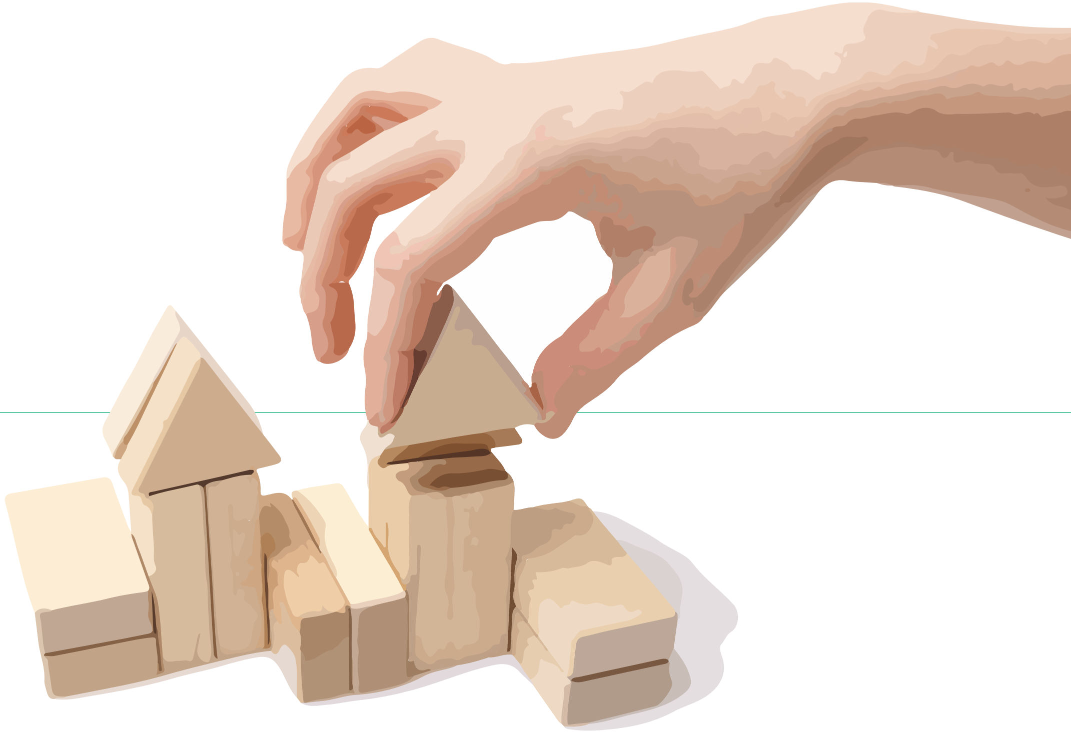

Relevant Specifications

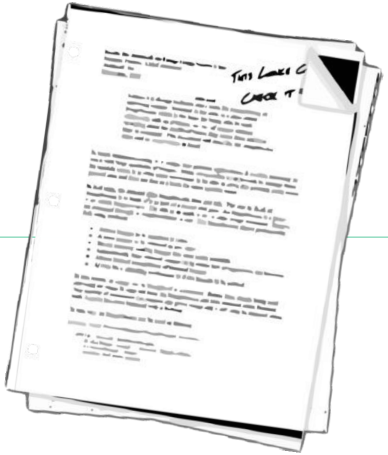

What Do the Specs Say?

- Curtain up!

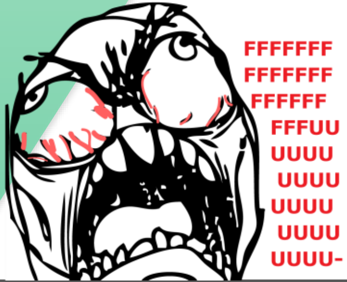

- Oh, that's an easy one. Just look at the RFCs.
- "The nice thing about standards is that you have so many to choose from."

Andrew Tanenbaum

  – This was funny, wasn't it?
  – Combine this with the *culture of deprecation* and out comes... a horrible mess.

There's Different Generations of IPv6 Stacks

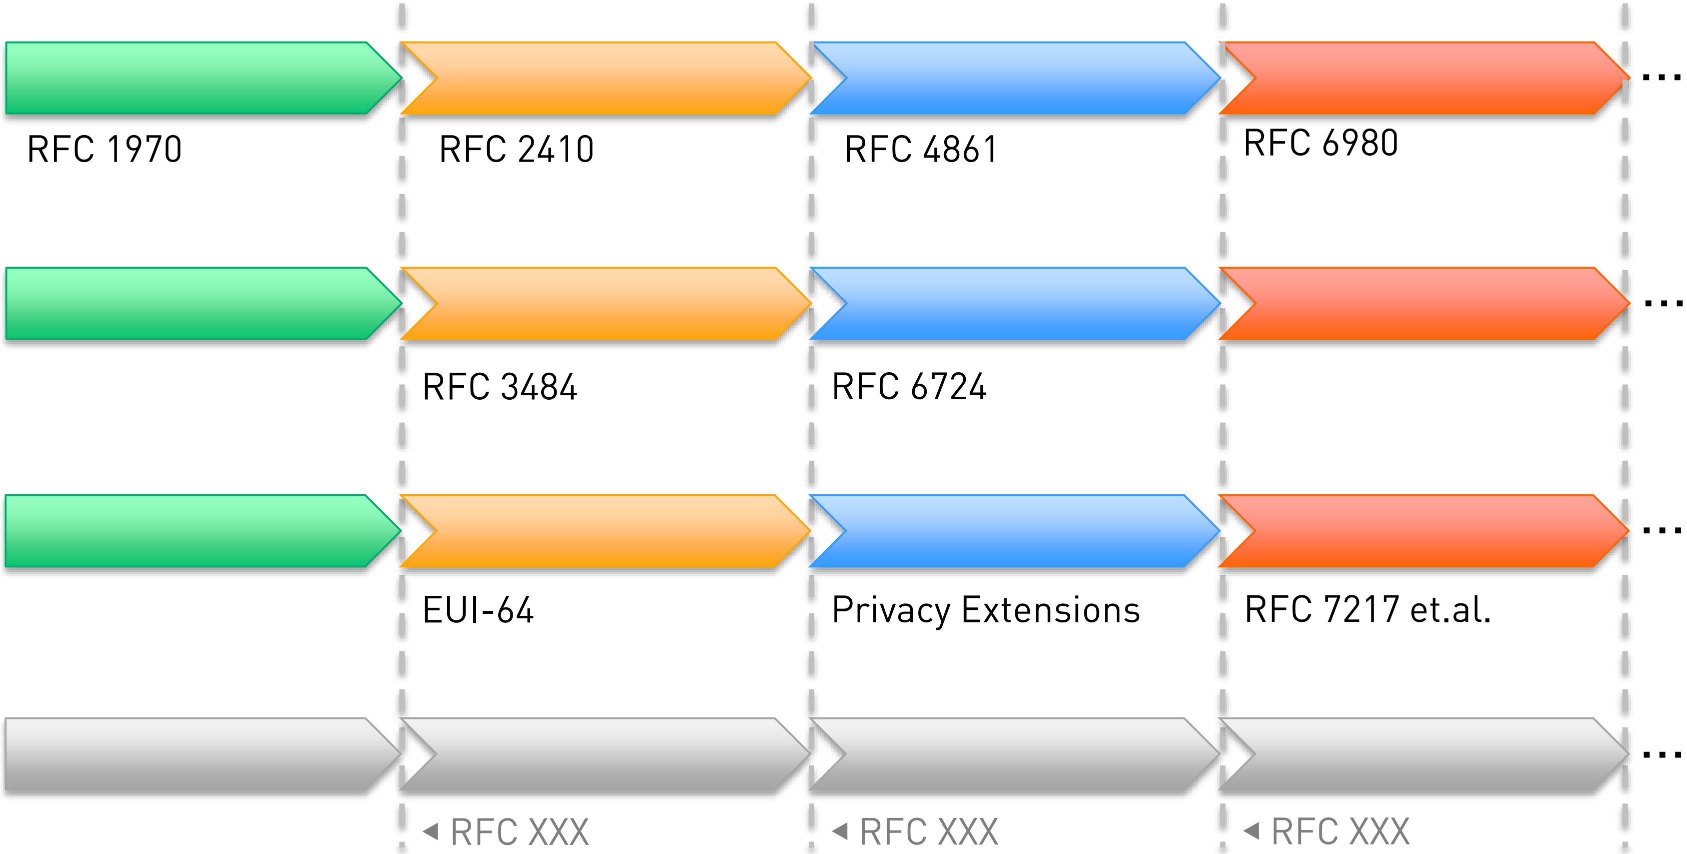

Neighbor Discovery	RFC 1970	RFC 2410	RFC 4861	RFC 6980	…
Address Selection		RFC 3484	RFC 6724		…
Generation of IID		EUI-64	Privacy Extensions	RFC 7217 et.al.	…
Etc.		◂ RFC XXX	◂ RFC XXX	◂ RFC XXX	…

RFC 2461

Note: RFC 4861, 6.2.7 on "Router Advertisement Consistency" seems to state that "inconsistencies are ok, but should be logged, by nodes".

6.3.4. Processing Received Router Advertisements

When multiple routers are present, the information advertised collectively by all routers may be a superset of the information contained in a single Router Advertisement. Moreover, information may also be obtained through other dynamic means, such as stateful autoconfiguration. Hosts accept the union of all received information; the receipt of a Router Advertisement MUST NOT invalidate all information received in a previous advertisement or from another source. However, when received information for a specific parameter (e.g., Link MTU) or option (e.g., Lifetime on a specific Prefix) differs from information received earlier, and the parameter/option can only have one value, the most recently received information is considered authoritative.

RFC 4861

- Sect. 4.2

  “If neither M nor O flags are set, this indicates that no information is available via DHCPv6.”

- If the M flag is set, the O flag is redundant and it can be ignored.

Some More Quotes

Not much RFC 2119 in there, is it?

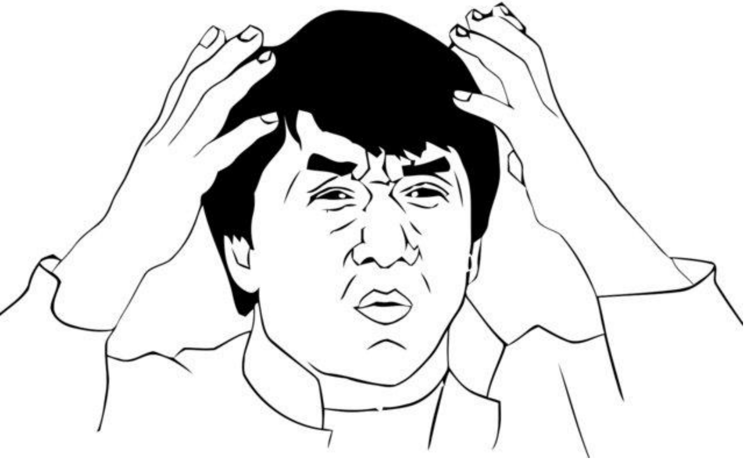

- RFC 4862, 5.5.2 *Absence of Router Advertisements*
  - "Even if a link has no routers, the DHCPv6 service to obtain addresses may still be available, and hosts may want to use the service."
- RFC 4862, 5.6 *Configuration Consistency*
  - "If the same configuration information is provided by multiple sources, the value of this information should be consistent."
  - "In any case, if there is no security difference, the most recently obtained values SHOULD have precedence over information learned earlier."

RFC 6106

“1.2 Coexistence of RA Options and DHCP Options for DNS Configuration

Two protocols exist to configure the DNS information on a host, the Router Advertisement options described in this document and the DHCPv6 options described in [RFC3646]. They can be used together.

The rules governing the decision to use stateful configuration mechanisms are specified in [RFC4861]. Hosts conforming to this specification MUST extract DNS information from Router Advertisement messages, unless static DNS configuration has been specified by the user.

If there is DNS information available from multiple Router Advertisements and/or from DHCP, the host MUST maintain an ordered list of this information as specified in Section 5.3.1.

RFC 6106

Section 5.3.1

In the case where the DNS options of RDNSS and DNSSL can be obtained from multiple sources, such as RA and DHCP, the IPv6 host SHOULD keep some DNS options from all sources.

Unless explicitly specified for the discovery mechanism, the exact number of addresses and domain names to keep is a matter of local policy and implementation choice.

However, the ability to store at least three RDNSS addresses (or DNSSL domain names) from at least two different sources is RECOMMENDED.

The DNS options from Router Advertisements and DHCP SHOULD be stored into the DNS Repository and Resolver Repository so that information from DHCP appears there first and therefore takes precedence.

Thus, the DNS information from DHCP takes precedence over that from RA for DNS queries.

In Short

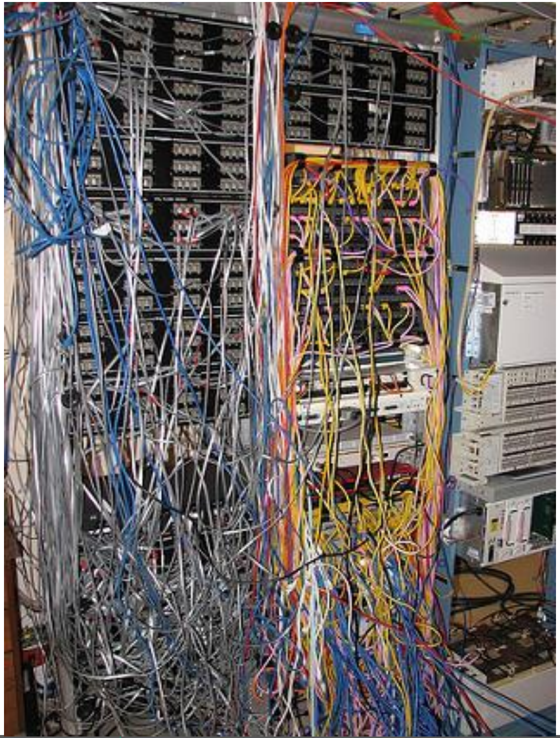

- It´s a mess!

At least on the specs level.

Problem Statement

From a High-Level Perspective

Problem Statement (I)

- IPv6 provides two mechanisms for one task, that is provisioning of IP parameters to nodes.

Problem Statement (II)

There´s two mechanisms

- They are independent.
  - Well, mostly.
- In many environments both of them are needed, in some combination.
  - In particular this applies in (wrt OSs, devices) heterogeneous environments.
    Read: probably in pretty much all of your environments.
- In some environments different groups might be responsible for operating them.
  - Why did you add this to the "problem statement"? Well…
- There´s differences as for the degree of vendor support & their strategies.

Problem Statement (III)

Let's look at the specs...

- Some properties and elements have been enhanced over time, e.g. RFC 6106.
  - Evolution is a good thing. Seriously!

- Still, there's a certain (alas, when it comes to IPv6: usual) amount of ambiguity and vagueness in the main RFCs.

Problem Statement (IV)

- The “cooperation“ of those two mechanisms has been discussed quite a bit, both in the specs and in “the relevant fora“.
- However, not so much as for scenarios where the information provided by them is conflicting.
- This is what this talk is about.

Problem Statement (IV)

Can such ("conflict scenarios") happen?

- Human error
  - Both on the *active failure* and *latent failure* level.
- Conflicting/differing vendor default settings
  - Network devices
  - CPEs!
    - Keep in mind: there might be any OS in customers' networks.
- Attacker injecting nasty packets
  - Boils down to "standard local-link sec" discussion → we will only briefly cover this.

What's a "Conflict"?

Pls define!

- #1: Both mechanisms are (maybe: somewhat) present, but only one is supposed to be used.
- #2: Both mechanisms lead to address(es) on nodes.
- #3: Both mechanisms distribute RDNSS in parallel, but different ones.
- Once you think hard enough, you'll come up with many more variants.

Additional Observations

- [draft-ietf-v6ops-dhcpv6-slaac-problem-04] explicitly discusses the role of *state transitions*.

- We can confirm that these lead to particularly interesting effects.
  - → Pay special attention in times when you perform those deliberately.
    Be prepared for surprises…

- In general the *order of events* seems to play a role, too.
  - See also test cases with two routers below.

Why the Order Might Matter – Sample

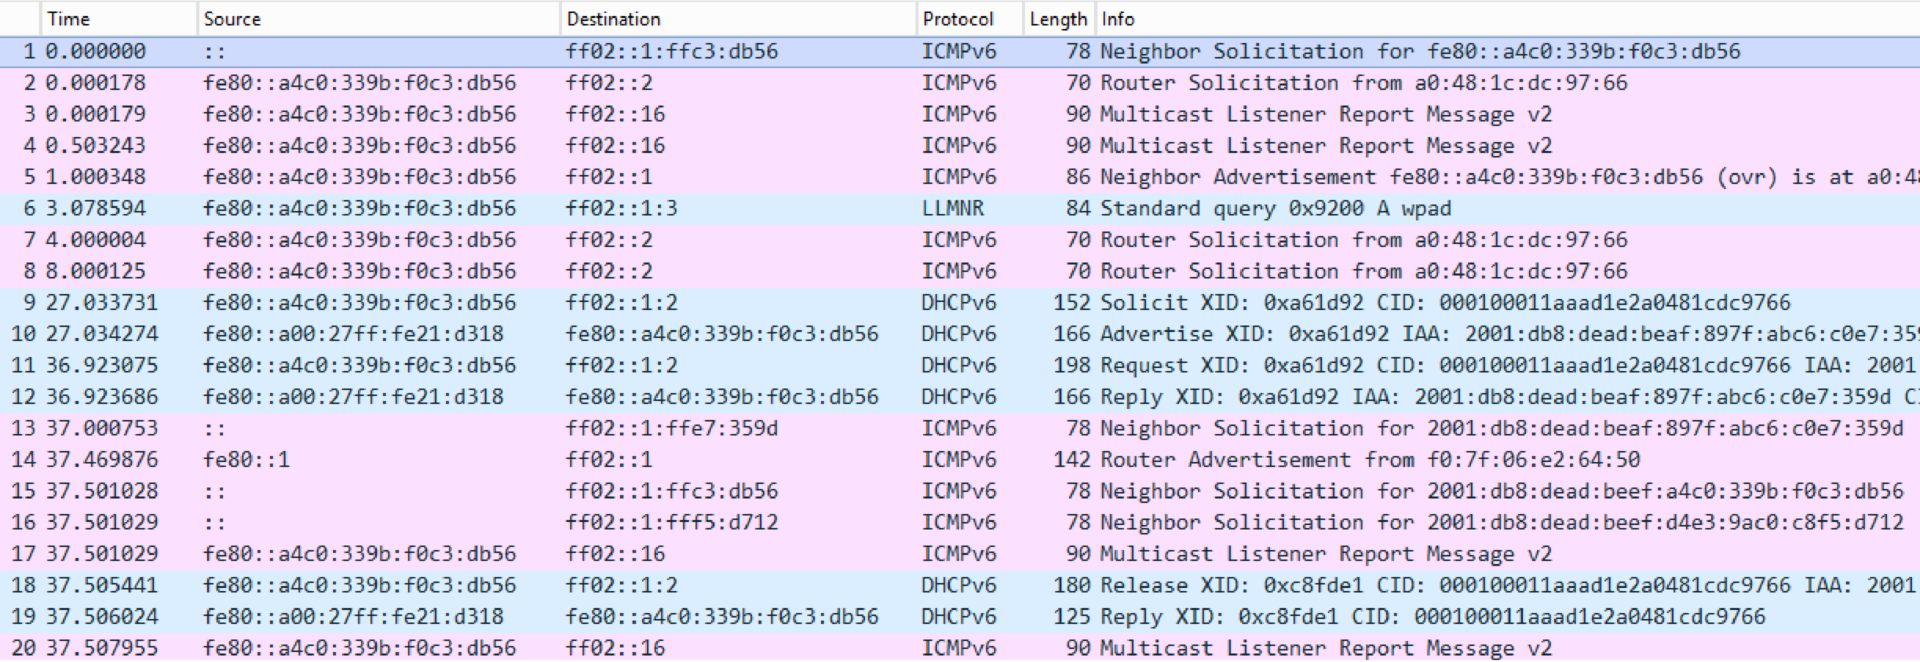

	Time	Source	Destination	Protocol	Length	Info
1	0.000000	::	ff02::1:ffc3:db56	ICMPv6	78	Neighbor Solicitation for fe80::a4c0:339b:f0c3:db56
2	0.000178	fe80::a4c0:339b:f0c3:db56	ff02::2	ICMPv6	70	Router Solicitation from a0:48:1c:dc:97:66
3	0.000179	fe80::a4c0:339b:f0c3:db56	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
4	0.503243	fe80::a4c0:339b:f0c3:db56	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
5	1.000348	fe80::a4c0:339b:f0c3:db56	ff02::1	ICMPv6	86	Neighbor Advertisement fe80::a4c0:339b:f0c3:db56 (ovr) is at a0:4
6	3.078594	fe80::a4c0:339b:f0c3:db56	ff02::1:3	LLMNR	84	Standard query 0x9200 A wpad
7	4.000004	fe80::a4c0:339b:f0c3:db56	ff02::2	ICMPv6	70	Router Solicitation from a0:48:1c:dc:97:66
8	8.000125	fe80::a4c0:339b:f0c3:db56	ff02::2	ICMPv6	70	Router Solicitation from a0:48:1c:dc:97:66
9	27.033731	fe80::a4c0:339b:f0c3:db56	ff02::1:2	DHCPv6	152	Solicit XID: 0xa61d92 CID: 000100011aaad1e2a0481cdc9766
10	27.034274	fe80::a00:27ff:fe21:d318	fe80::a4c0:339b:f0c3:db56	DHCPv6	166	Advertise XID: 0xa61d92 IAA: 2001:db8:dead:beaf:897f:abc6:c0e7:35
11	36.923075	fe80::a4c0:339b:f0c3:db56	ff02::1:2	DHCPv6	198	Request XID: 0xa61d92 CID: 000100011aaad1e2a0481cdc9766 IAA: 2001
12	36.923686	fe80::a00:27ff:fe21:d318	fe80::a4c0:339b:f0c3:db56	DHCPv6	166	Reply XID: 0xa61d92 IAA: 2001:db8:dead:beaf:897f:abc6:c0e7:359d C
13	37.000753	::	ff02::1:ffe7:359d	ICMPv6	78	Neighbor Solicitation for 2001:db8:dead:beaf:897f:abc6:c0e7:359d
14	37.469876	fe80::1	ff02::1	ICMPv6	142	Router Advertisement from f0:7f:06:e2:64:50
15	37.501028	::	ff02::1:ffc3:db56	ICMPv6	78	Neighbor Solicitation for 2001:db8:dead:beef:a4c0:339b:f0c3:db56
16	37.501029	::	ff02::1:fff5:d712	ICMPv6	78	Neighbor Solicitation for 2001:db8:dead:beef:d4e3:9ac0:c8f5:d712
17	37.501029	fe80::a4c0:339b:f0c3:db56	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
18	37.505441	fe80::a4c0:339b:f0c3:db56	ff02::1:2	DHCPv6	180	Release XID: 0xc8fde1 CID: 000100011aaad1e2a0481cdc9766 IAA: 2001
19	37.506024	fe80::a00:27ff:fe21:d318	fe80::a4c0:339b:f0c3:db56	DHCPv6	125	Reply XID: 0xc8fde1 CID: 000100011aaad1e2a0481cdc9766
20	37.507955	fe80::a4c0:339b:f0c3:db56	ff02::16	ICMPv6	90	Multicast Listener Report Message v2

From the Lab

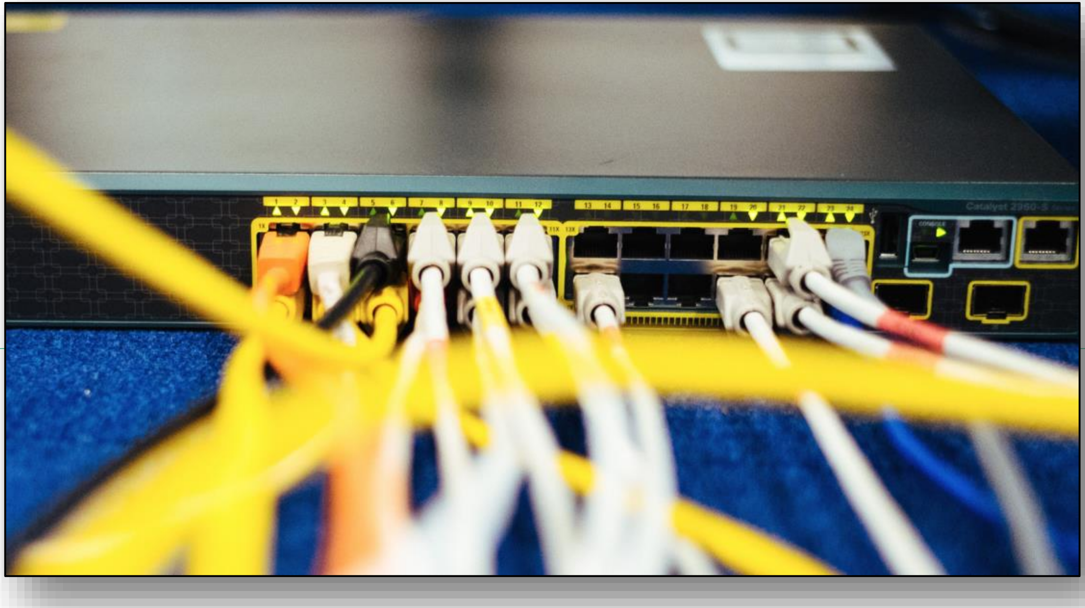

Lab Setup

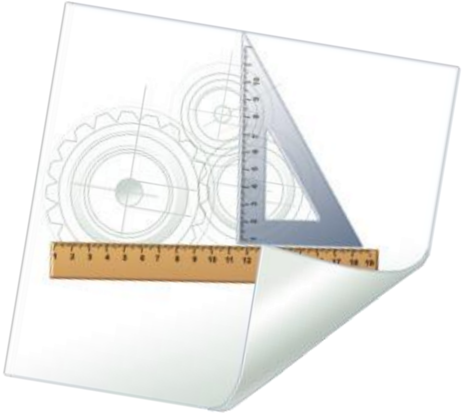

See also:
https://www.ernw.de/download/ERNW_Whitepaper_IPv6_RAs_RDNSS_DHCPv6_Conflicting_Parameters.pdf

- A DHCPv6 Server (DHCP ISC Version 4.3.1) installed on CentOs 6.6 . The DHCPv6 server is configured to provide both IPv6 addresses and RDNSS information.
- Two (2) routers Cisco 4321 using Cisco IOS Software version 15.5(1)S.
- The following OS as clients:
  - Fedora 21, kernel version 3.18.3-201 x64
  - Ubuntu 14.04.1 LTS, kernel version 3.13.0-44-generic
  - CentOS 7, kernel version 3.10.0-123.13.2.el7
  - Mac OS X 10.10.2 Yosemite
  - Windows 7, patch level Feb 2015
  - Windows 8.1, patch level Feb 2015

Case 1: One Router with the Management Flag not Set and a DHCPv6 Server

Router: M=0, A=1, O=0 and an RDNSS is advertised.

DHCPv6 server on the same link offering IPv6 addresses & RDNSS

- Fedora 21, MAC OS X, CentOS 7 and Ubuntu 14.04
  - Get an IPv6 address and an RDNSS from the IPv6 router only.
- Windows 7
  - Get an IPv6 address from the router only, but they do not get any DNS information, neither from the router nor from the DHCPv6 server. They also do not get IPv6 address from the DHCPv6 server.
- Windows 8.1
  - Get an IPv6 address from both the IPv6 router and the DHCPv6 server, despite the fact that the Management flag (M) is not set. They get RDNSS information from the DHCPv6 only.

Case 4: All Flags are Set and a DHCPv6 Server is Present

Router: M=1, A=1, O=1, and an RDNSS is advertised.

A DHCPv6 server on the same link advertising IPv6 addresses and RDNSS.

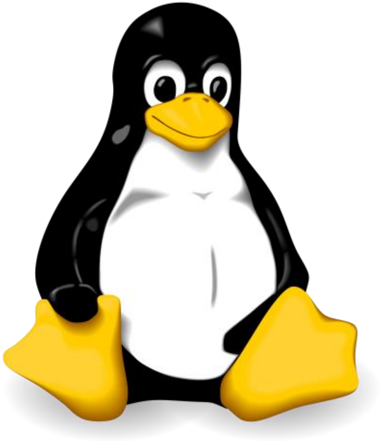

- Fedora 21 and Centos 7:
  - They get IPv6 addresses both from SLAAC and DHCPv6 server.
  - They get RDNSS both from RAs and DHCPv6 server.
  - The DNS of the RAs has higher priority.
- Ubuntu 14.04:
  - It gets IPv6 addresses both using SLAAC and from the DHCPv6 server.
  - It gets RDNSS from RAs only.
  - From the DHCPv6 server it only gets “Domain Search List” information, no RDNSS.

Case 4 Results cont‘d

- MAC OS X:
  - It gets IPv6 addresses both using SLAAC and from the DHCPv6 server.
  - It also gets RDNSS both from RAs and the DHCPv6 server.
  - The DNS server from DHCPv6 has higher priority.
- Windows 7 and Windows 8.1:
  - They get IPv6 addresses both from SLAAC and DHCPv6 server.
  - They get RDNSS only from the DHCPv6 server.

Summary

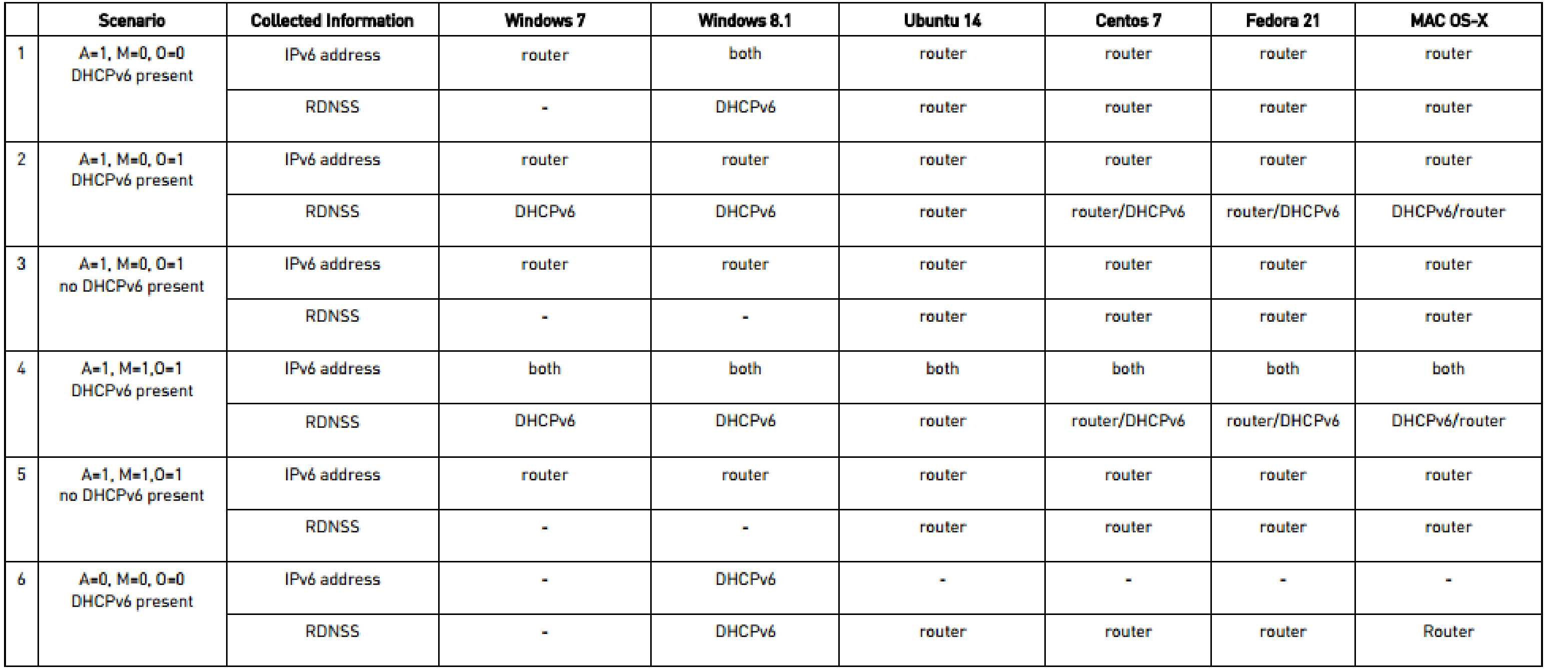

	Scenario	Collected Information	Windows 7	Windows 8.1	Ubuntu 14	Centos 7	Fedora 21	MAC OS-X
1	A=1, M=0, O=0 DHCPv6 present	IPv6 address	router	both	router	router	router	router
		RDNSS	-	DHCPv6	router	router	router	router
2	A=1, M=0, O=1 DHCPv6 present	IPv6 address	router	router	router	router	router	router
		RDNSS	DHCPv6	DHCPv6	router	router/DHCPv6	router/DHCPv6	DHCPv6/router
3	A=1, M=0, O=1 no DHCPv6 present	IPv6 address	router	router	router	router	router	router
		RDNSS	-	-	router	router	router	router
4	A=1, M=1,O=1 DHCPv6 present	IPv6 address	both	both	both	both	both	both
		RDNSS	DHCPv6	DHCPv6	router	router/DHCPv6	router/DHCPv6	DHCPv6/router
5	A=1, M=1,O=1 no DHCPv6 present	IPv6 address	router	router	router	router	router	router
		RDNSS	-	-	router	router	router	router
6	A=0, M=0, O=0 DHCPv6 present	IPv6 address	-	DHCPv6	-	-	-	-
		RDNSS	-	DHCPv6	router	router	router	Router

https://www.ernw.de/download/ERNW_Whitepaper_IPv6_RAs_RDNSS_DHCPv6_Conflicting_Parameters.pdf
https://tools.ietf.org/html/draft-ietf-v6ops-dhcpv6-slaac-problem-04

More Stuff from the Lab

Case 7: Router 1 Advertising M=0, O=0 and RDNSS, and then Router 2 advertising M=1, O=1 while DHCPv6 is Present

Initially:

One IPv6 router with the following settings:

- M=0, O=0, A=1 and RDNSS advertised and 15 seconds time interval of the RAs.

After a while (when clients are configured by the RAs of the above router) another IPv6 router with the following:

- M=1, O=1, no advertised prefix information, and 30 seconds time interval of the RAs.

MAC OS X and Ubuntu 14.04:

- Initially they get address and RDNSS from the first router.
- When they receive RAs from the second router, they never get any information (IPv6 address or RDNSS) from the DHCPv6 server.

Case 7 Results cont‘d

- Fedora 21 and Centos 7:
  - Initially they get IPv6 address and RDNSS from the RAs of the first router.
  - When they receive an RA from router 2, they also get an IPv6 address and RDNSS from the DHCPv6 server while retaining the ones (IPv6 address and RDNSS) obtained from the RAs of the first router.
  - The RDNSS obtained from the first router has a higher priority than the one obtained from the DHCPv6 server (probably because it was received first).
  - When they receive again RAs from the first router, they lose/forget the information (IPv6 address and RDNSS) obtained from the DHCPv6 server.

    → Troubleshooting nightmare…

Case 7 Results cont'd

- Windows 7:
  - Initially they get address from the first router – no RDNSS.
  - When they receive RAs from the second router, they never get any information (IPv6 address or RDNSS) from the DHCPv6 server.

Case 7 Results cont‘d

- Windows 8.1:
  - Initially, they get just an IPv6 address from the first router 1 - no RDNSS information (since they do not implement RFC 6106).
  - When they receive RAs from the second router, then they also get an IPv6 address from the DHCPv6 server, as well as RDNSS from it. They do not lose the IPv6 address obtained by the first router using SLAAC.
  - When they receive another RA from the first router, they retain all the information obtained so far (there isn't any change).

Summary

	Scenario		Collected Information	Windows 7	Windows 8.1	Ubuntu 14	Centos 7	Fedora 21	MAC OS-X
7	Initial Situation	Router 1: A=1, M=0, O=0, RDNSS, 15 sec. RA interval	IPv6 address	router	router	router	router	router	router
			RDNSS	-	-	router	router	router	router
	Later addition	Router 2: M=1, O=1, no advertised prefix, 30 sec. RA interval DHCPv6 server.	IPv6 address	router	both	router	Both	both	router
			RDNSS	-	DHCPv6	router	Router/DHCPv6	Router/DHCPv6	router
	Router 1 RAs received again		IPv6 address	router	both	router	router	router	router
			RDNSS	-	DHCPv6	router	router	router	router
8	Initial Situation	Router 2: M=1, O=1, no advertised prefix, 30 sec. RA interval DHCPv6 server	IPv6 address	DHCPv6	DHCPv6	DHCPv6	DHCPv6	DHCPv6	DHCPv6
			RDNSS	DHCPv6	DHCPv6	DHCPv6	DHCPv6	DHCPv6	DHCPv6
	Later addition	Router 1: A=1, M=0, O=0, RDNSS, 15 sec. RA interval	IPv6 address	Router 1	DHCPv6	both	Router 1	Router 1	both
			RDNSS	-	DHCPv6	Router 1	Router 1	Router 1	DHCPv6
	Router 2 RAs received again		IPv6 address	Both		both	Both	both	both
			RDNSS	DHCPv6		Router 1	Router1/DHCPv6	Router1/DHCPv6	DHCPv6

Conclusions

- Don't assume a certain OS' IPv6 behavior just because:
  - "the specs say so"
  - "another OS does it that way"
  - you have a good understanding of IPv4.
- Sorry guys ;-)
- Test, test, test!
  - Helps to gain (even more) IPv6 knowledge anyway.
  - Yes, please allocate budget for test lab.

Keep RFC 1925 in Mind

- "(4) Some things in life can never be fully appreciated nor understood unless experienced firsthand. Some things in networking can never be fully understood by someone who neither builds commercial networking equipment nor runs an operational network."

The Two Most Important RFCs Ever.

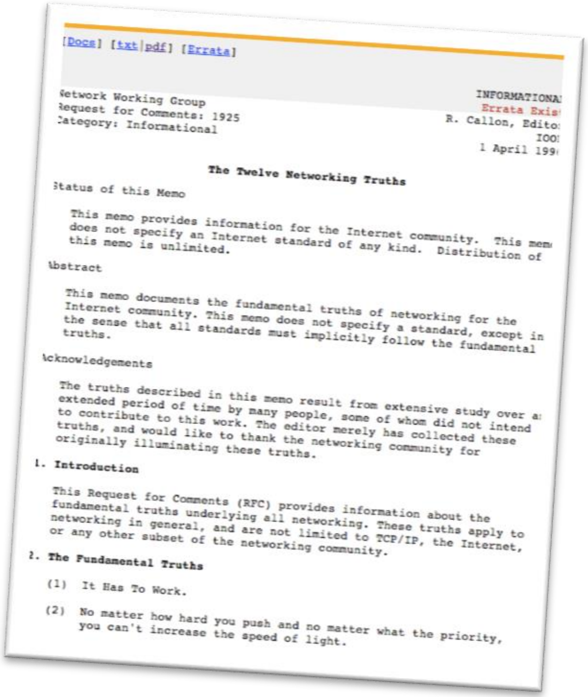

[Docs] [txt|pdf] [Errata]

etwork Working Group
equest for Comments: 1925
ategory: Informational

INFORMATIONA
Errata Exis
R. Callon, Edito
IOO
1 April 199

The Twelve Networking Truths

tatus of this Memo

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

bstract

This memo documents the fundamental truths of networking for the Internet community. This memo does not specify a standard, except in the sense that all standards must implicitly follow the fundamental truths.

cknowledgements

The truths described in this memo result from extensive study over a extended period of time by many people, some of whom did not intend to contribute to this work. The editor merely has collected these truths, and would like to thank the networking community for originally illuminating these truths.

1. Introduction

This Request for Comments (RFC) provides information about the fundamental truths underlying all networking. These truths apply to networking in general, and are not limited to TCP/IP, the Internet, or any other subset of the networking community.

2. The Fundamental Truths

(1) It Has To Work.

(2) No matter how hard you push and no matter what the priority, you can't increase the speed of light.

RFC 1925
The Twelve
Networking Truths

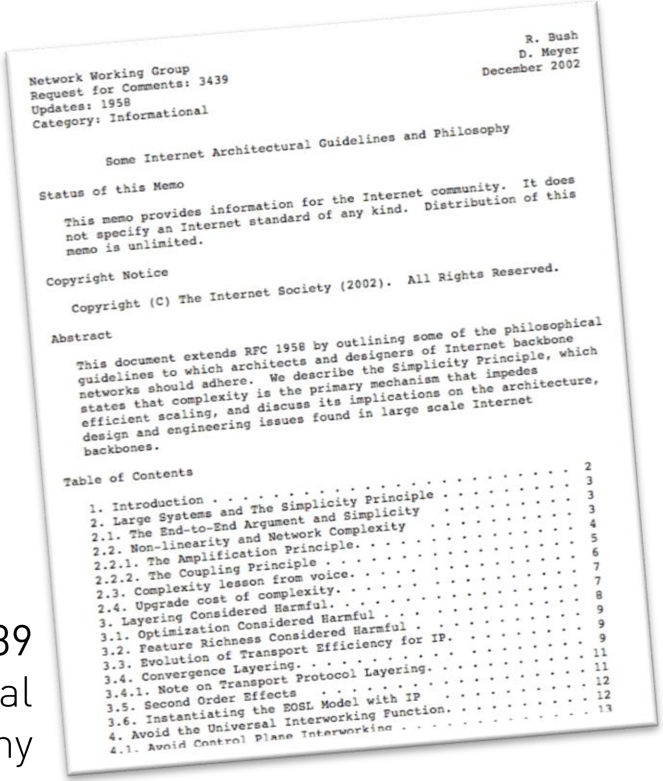

Network Working Group
Request for Comments: 3439
Updates: 1958
Category: Informational

R. Bush
D. Meyer
December 2002

Some Internet Architectural Guidelines and Philosophy

Status of this Memo

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

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

This document extends RFC 1958 by outlining some of the philosophical guidelines to which architects and designers of Internet backbone networks should adhere. We describe the Simplicity Principle, which states that complexity is the primary mechanism that impedes efficient scaling, and discuss its implications on the architecture, design and engineering issues found in large scale Internet backbones.

Table of Contents

RFC 3439
Some Internet Architectural
Guidelines and Philosophy

Operations Perspective

- Keep configs/properties related to IPv6 parameter provisioning in sync, at all times
  - IPv6 transition might be an opportunity to re-think your ops model.
  - Yes, we understand you´ll be happy to survive that one mostly unscathed, hence concentrate on one task at a time. Still #justsayin ;-)

Planning Perspective

Considerations how to set up the whole SLAAC/DHCPv6 thing

- In short: it depends ☺
- Seriously: it depends heavily on the client base you want to support. Here's some thoughts:
  - in case there's Android devices, your routers should advertise RDNSS info (RFC 6106), else the Android clients will have to rely on their IPv4 DNS or manual kludges. RFC 6106 is supported since Lollipop.
  - in case you don't have Android devices, you might go _without_ advertising RDNSS in RAs, for the simple reason of reducing potential for errors/"unexpected behavior".
  - once you go with m-flag DHCPv6 clearing the A-flag in prefix information, but leaving the L-flag set (to "circumvent RFC 5942") is usually a good idea.
    - Ofc, you can't do this once there's Android devices as those won't generate any (non LL) address then.
  - we observe that most of our customers strive to go with m-flag DHCPv6. that's just an observation...

Troubleshooting

For the poor sod responsible…

A helpful resource:
https://wikispaces.psu.edu/display/ipv6/IPv6+Rosetta+Stone

- You should know how to diagnose a node´s exact properties on the OS level
  - incl. types of addresses and order of name resolution
    - `"netsh int ipv6" commands on Win`
    - `"ip -6 add show" on Linux`
    - `btw: /etc/resolv.conf not relevant on Mac`
- The truth is in the packets…

Troubleshooting

In such scenarios

- Being familiar with the following certainly helps
  - Flags in router advertisements
  - Main DHCPv6 messages
  - IPv6 Subnet Model (RFC 5942) and its (non-) relationship with DHCPv6

Summary

- Some IPv6 RFCs merely serve as an indication & inspiration how things *could* be implemented.
- In complex & heterogeneous network you may expect surprises when it comes to the actual behavior of IPv6 nodes.
- Get your hands dirty, and (re-) read RFC 3439.

There’s never enough time...

THANK YOU...

...for yours!

Guys, we would love to see you in Heidelberg!

March, 14-18 2016
Heidelberg, Germany
Make the world a safer place.

FOLLOW THE NEWS
@WEareTROOPERS

MORE INFO, EXTENSIVE ARCHIVE & REGISTRATION @ WWW.TROOPERS.DE

Questions?

- You can reach us at: ✉
  - erey@ernw.de, www.ernw.de
  - cwerny@ernw.de, www.ernw.de
- Our blog:
  - www.insinuator.net
- Follow me at: 🐦
  - @Enno_Insinuator