Why IPv6 Security Is So Hard – Structural Deficits of IPv6 & Their Implications

Enno Rey, erey@ernw.de, @enno_insinuator

#whoami

- Networking background, doing security as a full-time profession since 1997
- Taking care of LIR stuff at some enterprise LIRs
  - Including the one with probably the coolest org handle: ORG-HACK1-RIPE
- Blogging about IPv6 & other pieces at `https://insinuator.net/tag/ipv6/`

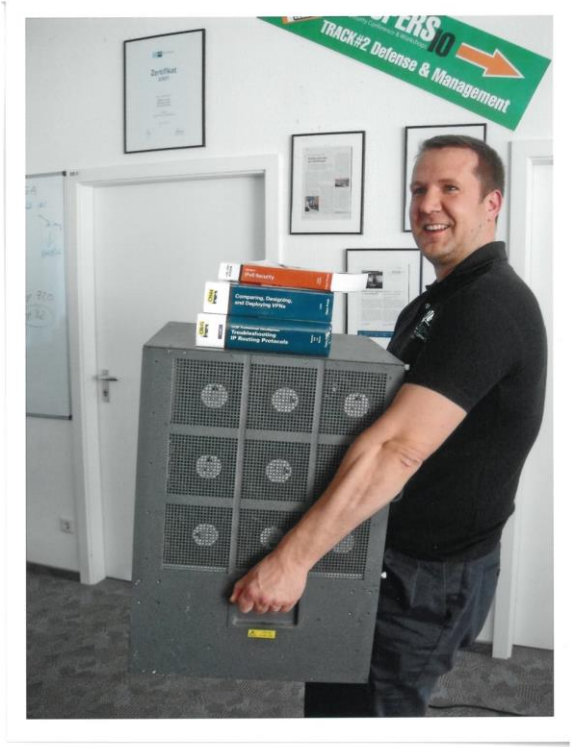

ERNW
providing security.

Agenda

- Some objectives, from a security perspective
- Properties of IPv6, and how they've developed over time
- Conclusions

Disclaimer

- This is not a rant about the IETF.
- Still many of the things I discuss in the following go back to decisions (or lack thereof) in relevant IPv6 WGs at the IETF.
- We've all heard about the creep of self-interest, (corporate) politics etc. into voluntary organizations – which effectively undermines their original purpose.

I think this point has long been reached in certain IETF circles, namely in *6man*.

Why I Think so Many Things in IETF 6man Go in the Wrong Direction

Let's look at how the actual discussion (and subsequent specification) work is done at the IETF, similar to other voluntary organizations: on mailing lists and in (f2f) meetings. As we all know, these meetings take place three times a year, each on a different continent (yes, I'm aware of remote participation, but let's be honest: at the end of the day how much impact on specification did this have this in past, in particular in heavily old boys' clubs dominated WGs like 6man?).

Further fact is: if you look at the lists of participants of the meetings, the vast majority of it is vendor personnel. This is not surprising when reflecting on the incentives different parties may have to send people to IETF meetings. How would, say, an enterprise person argue in front of her boss to attend the 51st (!) IETF meeting since the publication of RFC 2460 (especially considerung the ongoing [non-]state of deployment in large parts of that space. it's up to the reader to connect that state with the things I describe here...)?

But it's not like vendor people don't have to justify these nice trips to their bosses. Of course they have to. Here's two prevalent strategies:
- "we have that new feature. let's try to push it into an RFC, as this strengthens our market position (in general and for selling the specific thing)"
- "you know, there's this future thing called IPv6. I'm in one of the working groups where we come up with lots of creative ideas how to even make it better. my name is on one of the draft documents so I'll have to be there, at the next meeting (and we, as a vendor, demonstrated our contribution also)".

For quite some of the stakeholders (namely both the vendor in question and the respective participant[s]) these are not only legitimate but fully understandable. It's just: does this drive things in the right direction of the greater good & community? Me seems we have a classic tragedy of the commons here...

https://mailarchive.ietf.org/arch/msg/ipv6/tNR24ZN6o9APFEedk1_vTJe09Mc

Some Objectives
When It Comes to Network Security

ERNW
providing security.

Taking an Infosec Practitioner's View

- Predictability (<=> Trustworthiness)
  - “trust: the extent to which someone who relies on a system can have confidence that the system meets its specifications, i.e., that the system does what it claims to do and does not perform unwanted functions” (RFC 2828).
- Identification
  - Be able to identify actors being part of connections
    - Usually the basis for filtering
    - Helpful in the context of accountability, too.
- Ability to restrict / filter
  - To enforce security policy.

A bit more Abstract Objectives

- Keep things simple
- Avoid complexity
- Minimize state

Keep It Simple & Small

- There might be a direct relationship between (number of) lines of code and amount of vulnerabilities…
- Parsing needs CPU cycles
  - Often: more parsing → higher susceptibility to DoS
- The more protocols one uses the more attack surface might be exposed.

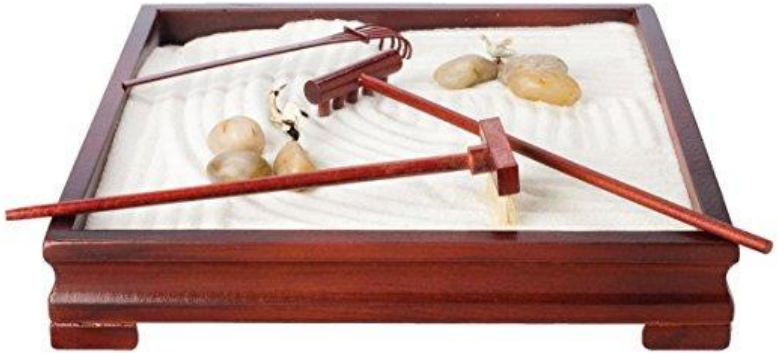

William of Ockham

Entia non sunt multiplicanda praeter necessitatem.

This translates roughly as:

More things should not be used than are necessary.

Occam‘s Razor
Phrased by a Networking Guy

- RFC 1925:

  (12) In protocol design, perfection has been reached not when there is nothing left to add, but when there is nothing left to take away.

ERNW
providing security.

Avoid Complexity

ERNW
providing security.

Complexity (I)

RIPE 74
BUDAPEST, HUNGARY
8 - 12 May 2017

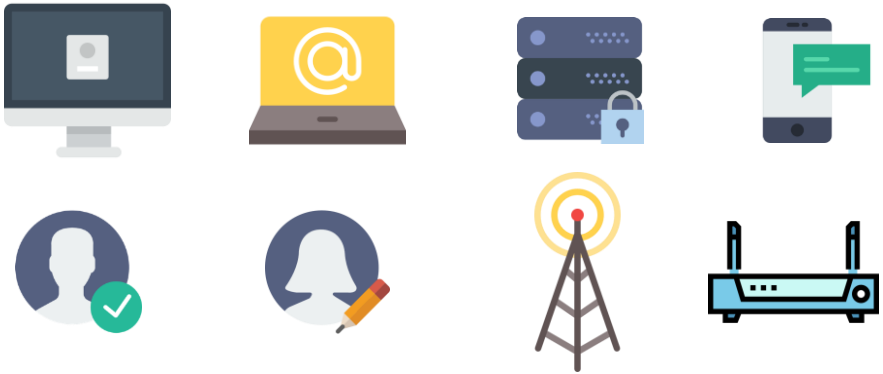

- composed of many interconnected parts

See also: https://insinuator.net/2015/05/ipv6-complexity/

Complexity (II)

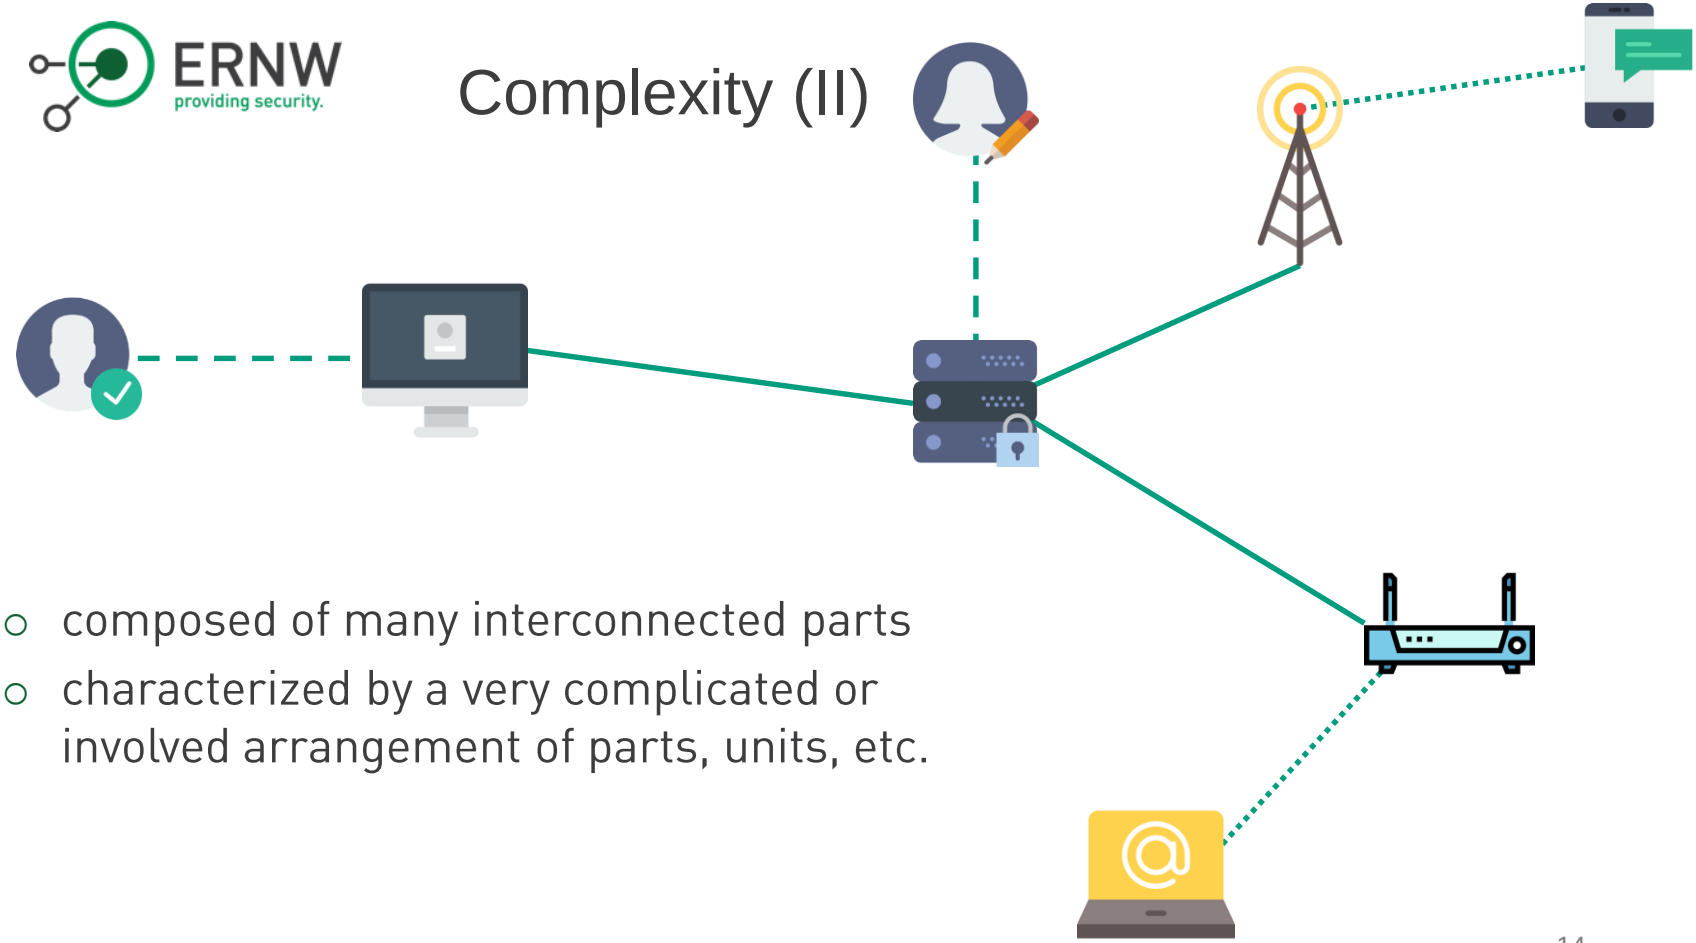

- composed of many interconnected parts
- characterized by a very complicated or involved arrangement of parts, units, etc.

Complexity (III)

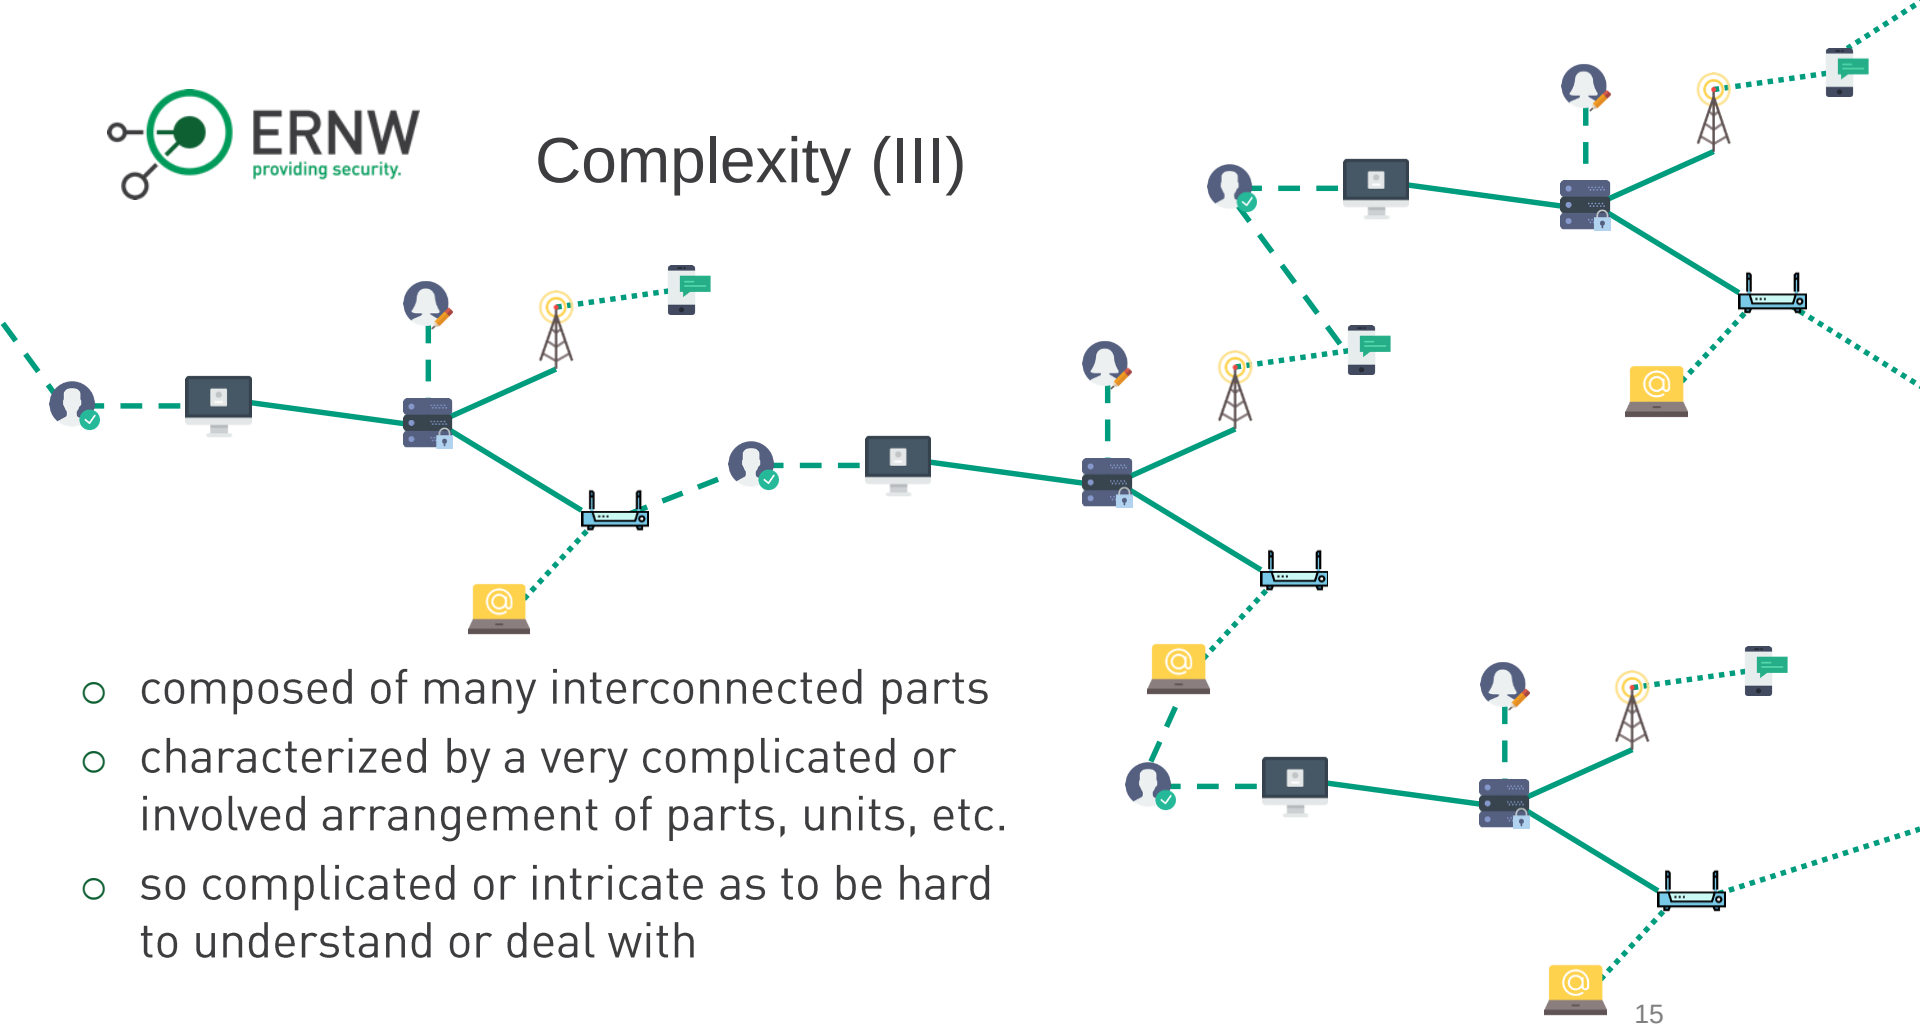

- composed of many interconnected parts
- characterized by a very complicated or involved arrangement of parts, units, etc.
- so complicated or intricate as to be hard to understand or deal with

Why the “Understanding“ Part is Crucial

- Understanding allows to
  - Develop mental model of inputs & their associated outputs
  - Predict output
- Mental model allows you to recognize when system isn‘t working correctly
  - Troubleshooting & fixing
  - Detection of security violations

IPv6 – Interactions

- Various types of relationships between SLAAC and DHCPv6
  - Unclear specs & several generations of them
  - Major vendors deviate, and still get it wrong
  - IETF WGs not aligned
    (e.g. RDNNS related momentum in v6ops vs. RFC 8106, sect. 5.3.1)
- Relationship between ND and MLD
- Relationship between RA flags, routing tables and address selection mechanisms
- Relationship between IP and other layers
  - All those lovely MTU issues come to mind.

(Minimize) State

- "State" usually encompasses several dimensions:
  - Amount of state (entries in $TABLE, RAM etc.)
  - Frequency/speed of state changes
  - Surface
    - Depth of interaction
    - Breadth of interaction

- Simple rule: the more state to be processed the higher the susceptibility to DoS

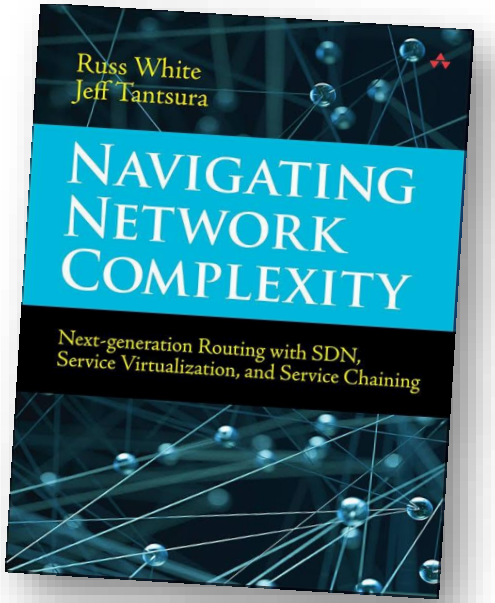

How State Can Kill a Network

"Our network switches have been observed using far more CPU than has historically been the case, we have had a variety of packet storms that appear to have been caused by forwarding loops despite the fact that we run a protocol designed to prevent such loops from taking place, and we have had a variety of unexplained switch crashes."

Massachusetts Institute of Technology

From:
http://blog.bimajority.org/2014/09/05/the-network-nightmare-that-ate-my-week/

State

Btw, some define complexity =

“anything for which there is more state than required to achieve a goal“

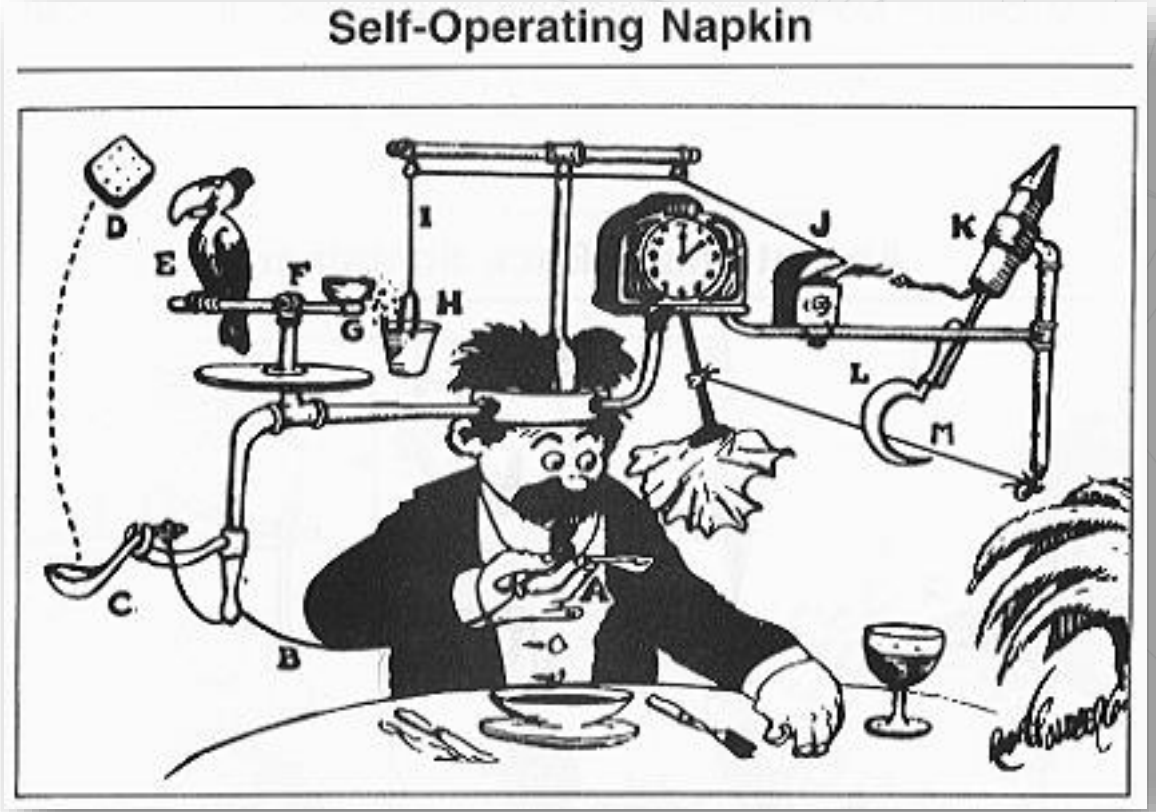

Rube Goldberg - Originally published in Collier's, September 26 1931

IPv6 Properties & Their (Non-) Evolution

Now Let's Have a Look at IPv6's Technical Properties

- 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*

Focus on Four of Them

- Multicast instead of broadcast
- Multiple address types & addresses
- Extension Headers
- Parameter provisioning

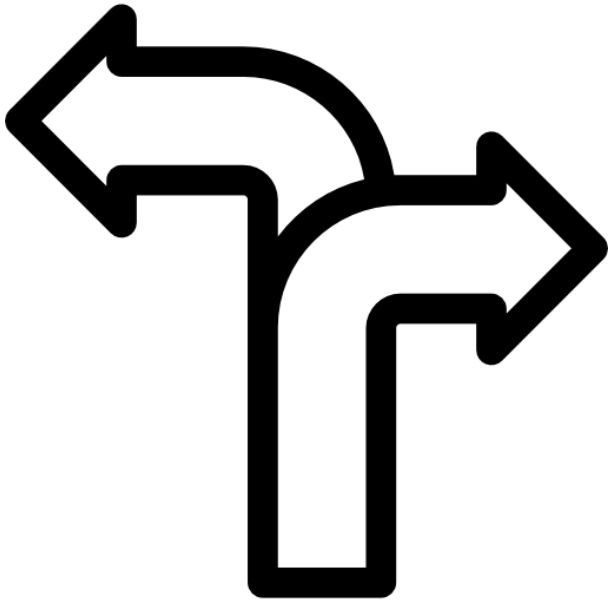

Multicast Instead of Broadcast

- Multicast based networking
  - Requires more state.
  - Usually (and in our case) requires more parsing
- One can probably write an implementation of ARP in max. 100 lines of Python code
  - Try this with ND ;-)
  - RFC 4861 has 94 pages. And has been updated by six (6) other RFCs...
- But, hey, you save some context changes/ interrupts on CPUs of local systems...

Multiple Address Types & Addresses

- IPv6 introduces the concept of a link-local address, as opposed to “global“ addresses
  - Separating the two is not a new concept
  - Still it‘s mainly associated with Ethernet networks, and doesn‘t make much sense in other types of networks, e.g. mobile/telco.

- Separating the two introduces new problems…

Multiple Address Types / Problems

- It increases (doubles?) the amount of state
  - Routing tables
  - Handling of addresses in kernel/IP stack etc.

- It creates a decision problem
  - Which address to choose for communication acts?
  - You‘re probably aware that – surprise! – there‘s several IETF documents for this.

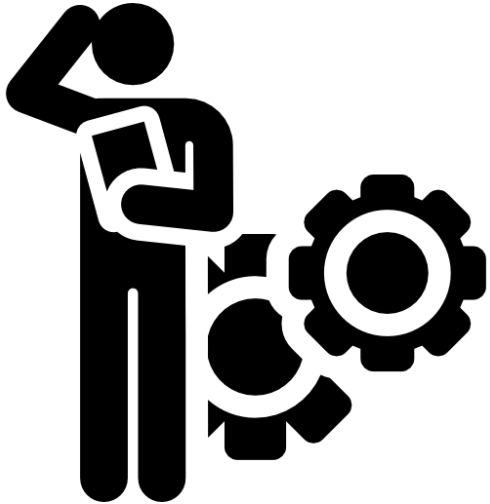

That Decision Problem

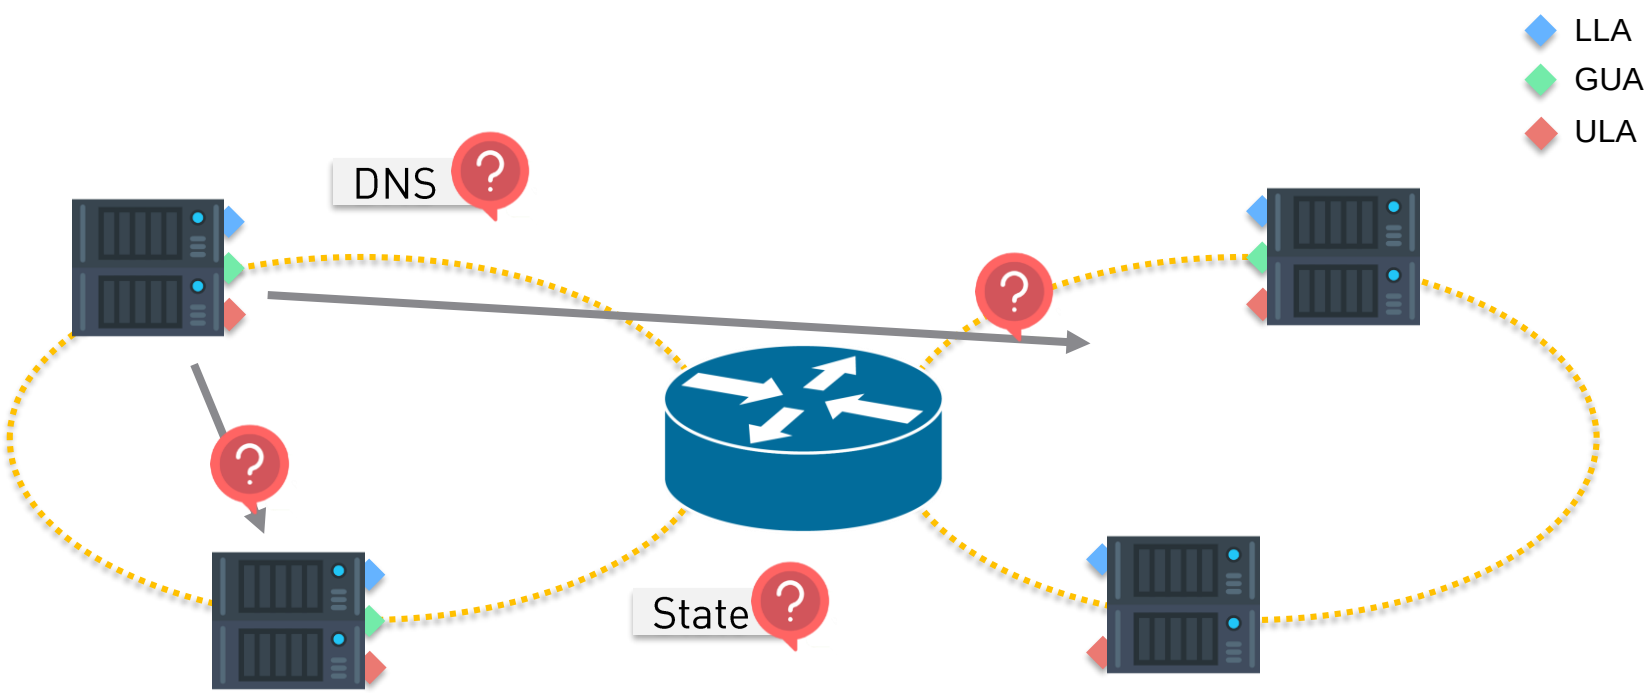

[Docs] [txt|pdf] [draft-ietf-v6ops-...] [Diff1] [Diff2] [Errata]

Errata Exist

Internet Engineering Task Force (IETF)
Request for Comments: 7934
BCP: 204
Category: Best Current Practice
ISSN: 2070-1721

L. Colitti
V. Cerf
Google
S. Cheshire
D. Schinazi
Apple Inc.
July 2016

Host Address Availability Recommendations

Extension Headers / Protocol Design

- Two main school of thoughts (re: protocol design)
  - Design a protocol that can handle many situations, and also support extensions that hadn‘t been thought of initially.
  - Design a protocol that (only) supports initial requirements.
- Looking at RFC 2460 the decision taken at the time immediately becomes clear.
- I‘m not judging this. But one must realize …

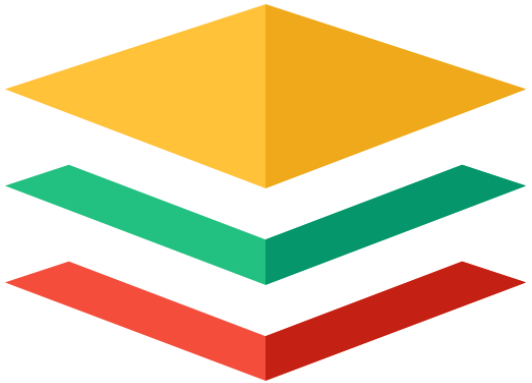

Implications of an Extensible Protocol

- Probably less predictability
- Almost certainly higher complexity
- More parsing (→ more code)
  - Also: https://youtu.be/Pru5BRrImz0
- Most probably more state needed

What an IPv6 Datagrams Looks Like…

What an IPv6 Datagrams Looks Like…

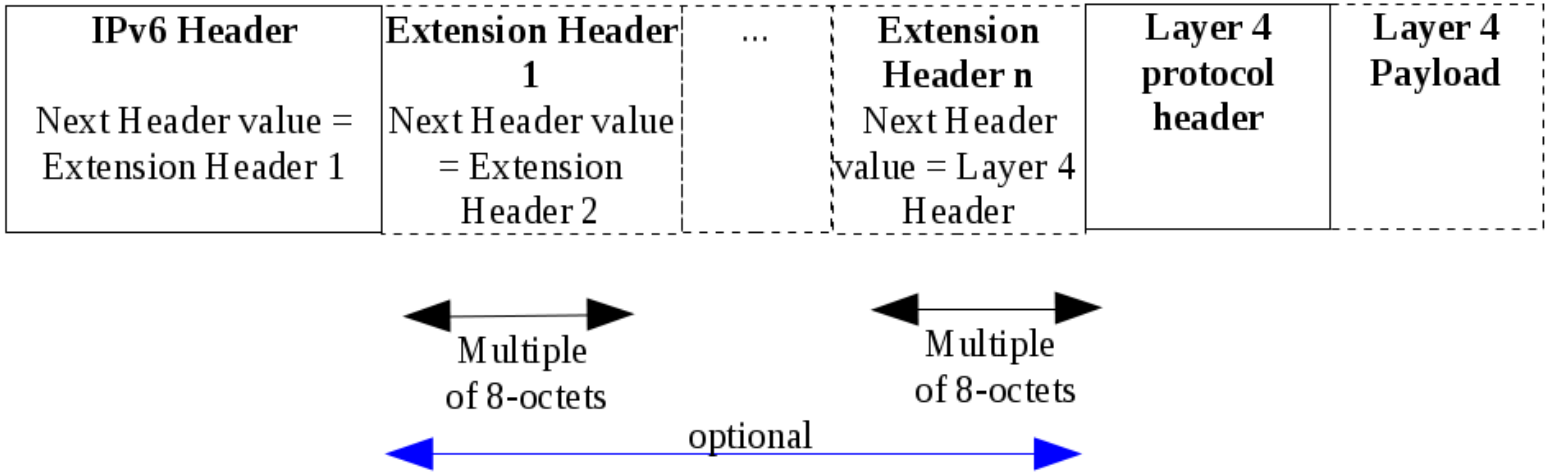

- This is the root of 3 significant problems…

Problem

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

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

Extensible Protocols Need This

- "be conservative in what you do,
  be liberal in what you accept from others"

  RFC 761

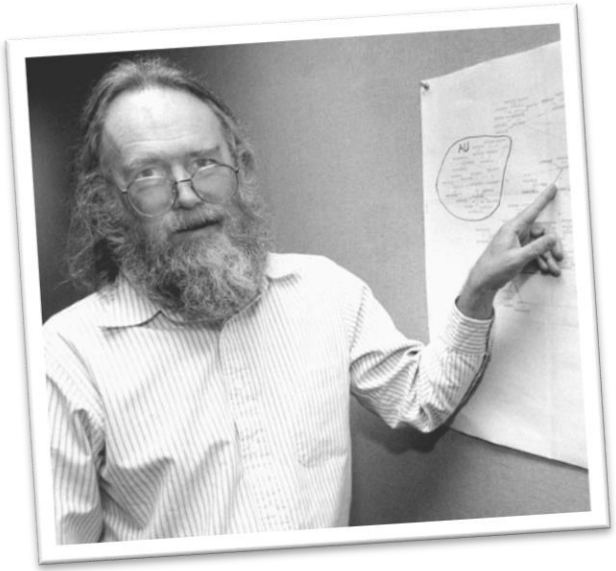

Once Upon a Time…

Postel's law was considered beneficial.

- Don't get me wrong: I'm a big fan of the *Robustness Principle*.
  - The Internet's innovation speed strongly related to it, at the time at least.
  - Imagine ITU (or IEEE for that matter) had had to specify the Internet…
- There's just one problem…

There Was a Time …

… when Postel's law was considered beneficial.

- Unfortunately, it fails once an involved party deliberately plays foul.

- Or as Eric Allman states it:
  - "The Robustness Principle was formulated in an Internet of cooperators."
    - The Robustness Principle Reconsidered, 2011, http://queue.acm.org/detail.cfm?id=1999945

Some Things Have Changed since the 80s

"Today, the motivations of some individuals using the Internet are not always entirely ethical, and, even if they are, the assumption that end nodes will always co-operate to achieve some mutually beneficial action, as implied by the end-to-end principle, is not always accurate."

[RFC 3724]

Security Problems Due to EHs

- Heavily increased parsing complexity
- Evasion of blacklist-based security controls
  - IDPS systems.
  - First Hop Security (FHS) features
  - Insufficient ACL/filtering implementations.
- For the record
  - "EHs" in the terminology of most sec ppl encompass: HBH, DestOptions, RH, FragHdr
  - AH &ESP have their (legitimate) role.
  - But nothing else…

https://www.ernw.de/download/eu-14-Atlasis-Rey-Schaefer-briefings-Evasion-of-HighEnd-IPS-Devices-wp.pdf

Parameter Provisioning

What's a *Router*?

- Wikipedia:
  - router = "a **router** is a device that forwards *data packets* between *computer networks*"
- RFC 2460:
  - router: "router - a node that forwards IPv6 packets not explicitly addressed to itself."

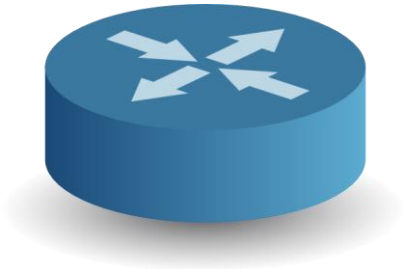

ERNW
providing security.

What's a *Router,* in IPv6?

Looking Closer

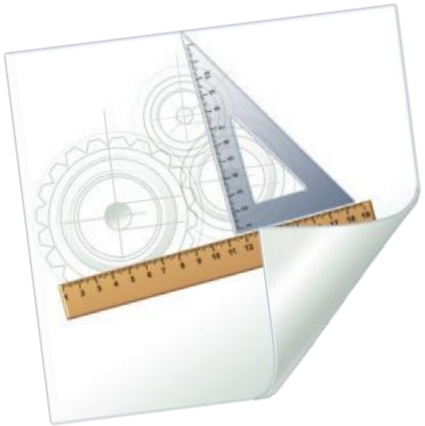

- RFC 2461: "Routers advertise their presence together with various link and Internet parameters either periodically, or in response to a Router Solicitation message".

- In the end of the day, in IPv6 a router is not just a forwarding device but a provisioning system as well.

IPv6's Trust Model

On the *local link* we're all brothers.

But Can‘t We just Filter the Bad Stuff?

There‘s RA Guard et al., right?

- Hmm… like most other *blacklist- based* security features RA Guard can be circumvented.
  - There‘s no (easy) cure for this. Choose two out of (function|speed|cost).
- Hey, we have RFC 6980 for this.
  - I for one consider this one of the most important IPv6 RFCs from the last years.
  - But it seems not easy to implement…
    - Which in turn might not be surprising…

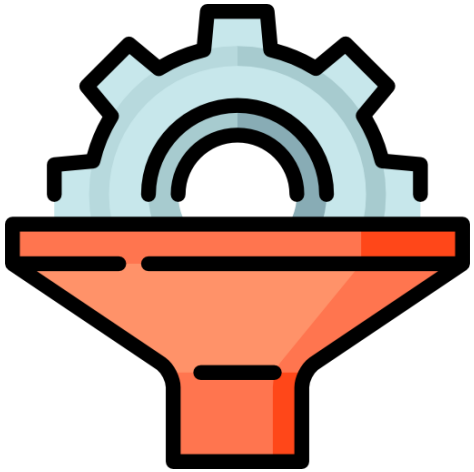

From some Recent Testing

https://insinuator.net/2017/03/testing-rfc-6980-implementations-with-chiron/

Test Case No.	Description	Chiron Options Used (in addition to baseline cmd)	Impact on Target OS' IPv6 Config (without RA Guard)	What was observed in Wireshark on Target OS? (without RA Guard)	What still got through with RA Guard enabled?	Overall Result With RA Guard Enabled
13	Two fragments, with two DestOptions in fragmentable part	-lfE 60,60 -nf 2	Added 2nd default gw, created additional address	One fragment plus RA packet which contains two DestOptions EHs	1st fragment, but *not* the RA	No impact
14	Four fragments, with two DestOptions in fragmentable part	-lfE 60,60 -nf 4	Added 2nd default gw, created additional address	Three fragments plus RA packet which contains two DestOptions	Three fragments, plus RA containing two DestOptions EHs. Nothing logged on the switch.	Successful attack
15	Two fragments, with two RoutingHdr EHs in fragmentable part	-lfE 43,43 -nf 2	Added 2nd default gw, created additional address	One fragment plus RA packet which contains two RoutingHdr EHs	Two fragments, plus RA containing EHs. "traceback" on switch console when running 15.0(2)SE2	Successful attack when switch runs 15.0(2)SE2, no impact when switch runs 15.0(2)SE10a
16	Two fragments, with two RHs and two DestOptions, in mixed order	-lfE 60,43,60,43 -nf 2	Added 2nd default gw, created additional address	One fragment plus RA packet which contains the four EHs	1st fragment, but *not* RA	No impact
17	Same as 16 but four fragments	-lfE 60,43,60,43 -nf 4	none	1st three segments only, but not RA	1st three fragments, but not RA	No impact
18	Same as 16 but three fragments	-lfE 60,43,60,43 -nf 3	Added 2nd default gw, created additional address	Two fragments, then RA containing all EHs	1st two fragments plus RA	Successful attack

Evolution. Not

Charles Darwin

Wait, Humans Learn and Standards Can Be Changed!

- Not really.
  In the IETF world standards are not withdrawn, but *deprecated*.
  - Because vendors – from their perspective fully legitimately – want to protect their investments.

Let's call this "the culture of deprecation"

withdrawn

deprecated

Culture of Deprecation & Its Consequences

- This means that in the vast majority of IPv6 stacks around there's some remnants of $SOME_PHASE_OF_IPV6_EVOLUTION.

- Which in turn heavily impedes *predictability.*

Different Generations of IPv6 Stacks

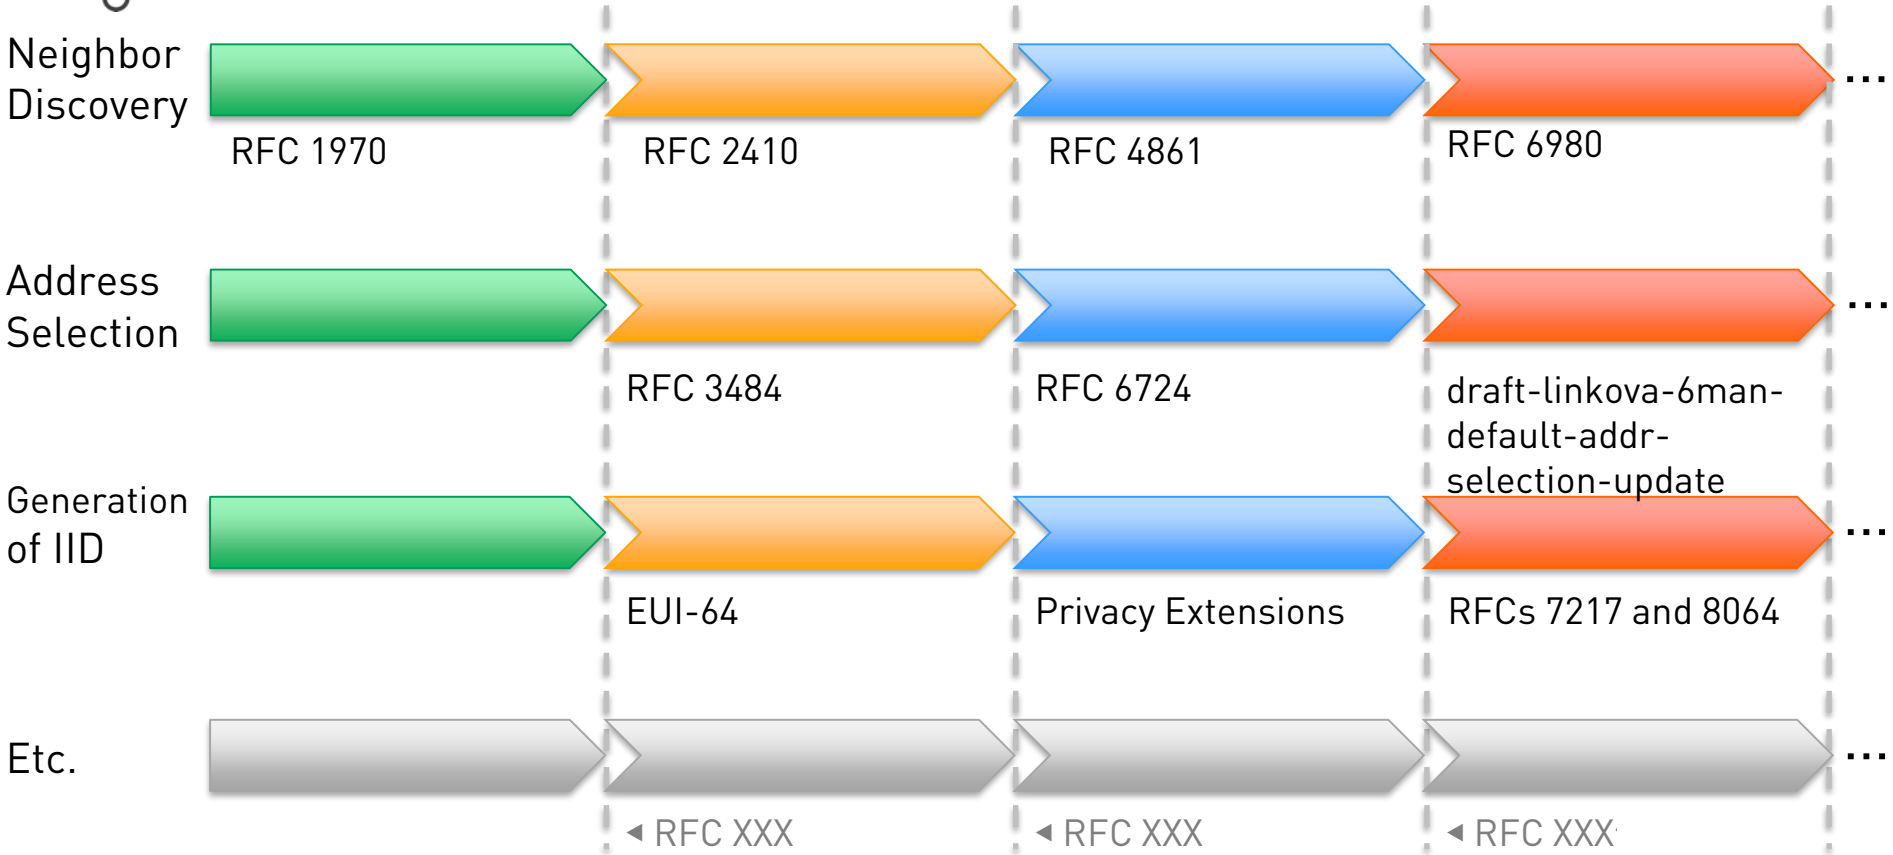

Conclusions

- IPv6 is much more complex (than IPv4)
  - On the protocol level.
  - On the operations level.

- IPv6 requires much more state
  - On L2 devices (e.g. multicast groups)
  - On L3 devices (*neighbors*)
  - On security devices

What Now?

- Try to understand
  - IPv6 *interactions* in your network.
  - where state is maintained by/for IPv6.
  - vendor agendas & incentives, namely in context of IETF stuff.
- Minimize complexity where possible
  - Drop EHs at the border of your DCs.
  - Limit interactions and/or number of protocols.
  - Keep addressing simple…
- Minimize the amount of state where possible
  - Re-think filtering approach?
  - Perform an inventory which type of state is created on different types of devices.

ERNW providing security.

There’s never enough time…

RIPE 74 BUDAPEST, HUNGARY 8 - 12 May 2017

THANK YOU…

…for yours!

@Enno_Insinuator

erey@ernw.de

ernw.de

insinuator.net

Slides available soon.

ERNW
providing security.

Sources

As indicated on slides.

Image Source:

- Icons made by Freepik from www.flaticon.com