Properties of IPv6 and Their Implications for Offense & Defense

Enno Rey, erey@ernw.de, @enno_insinuator

#whoami

- Some background in large scale networking, doing security as a full-time profession since '97.

- Founded (in 2001) a company performing security assessments, security research and the like
  - www.ernw.de

- Troopers ;-)
  - https://insinuator.net/2017/10/troopers-for-students/

Agenda

- Some objectives, from a security perspective
- Properties of IPv6, and their implications
- Conclusions

Common Objectives from a Network Security Perspective

- 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/interfaces 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.

Avoid Complexity

Complexity (I)

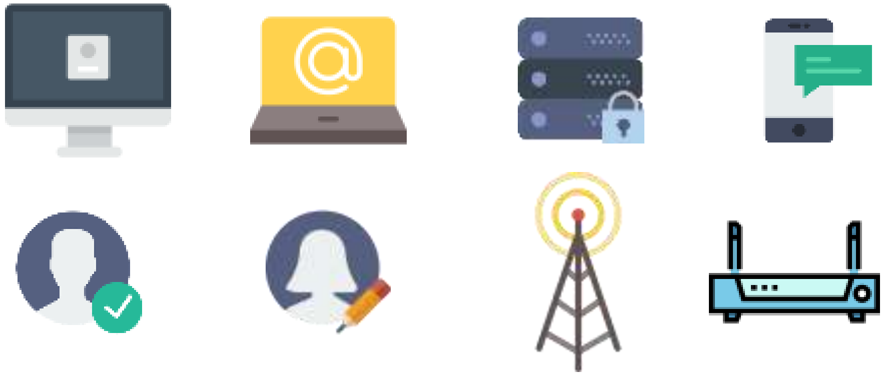

- Composed of many interconnected parts

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

Complexity (II)

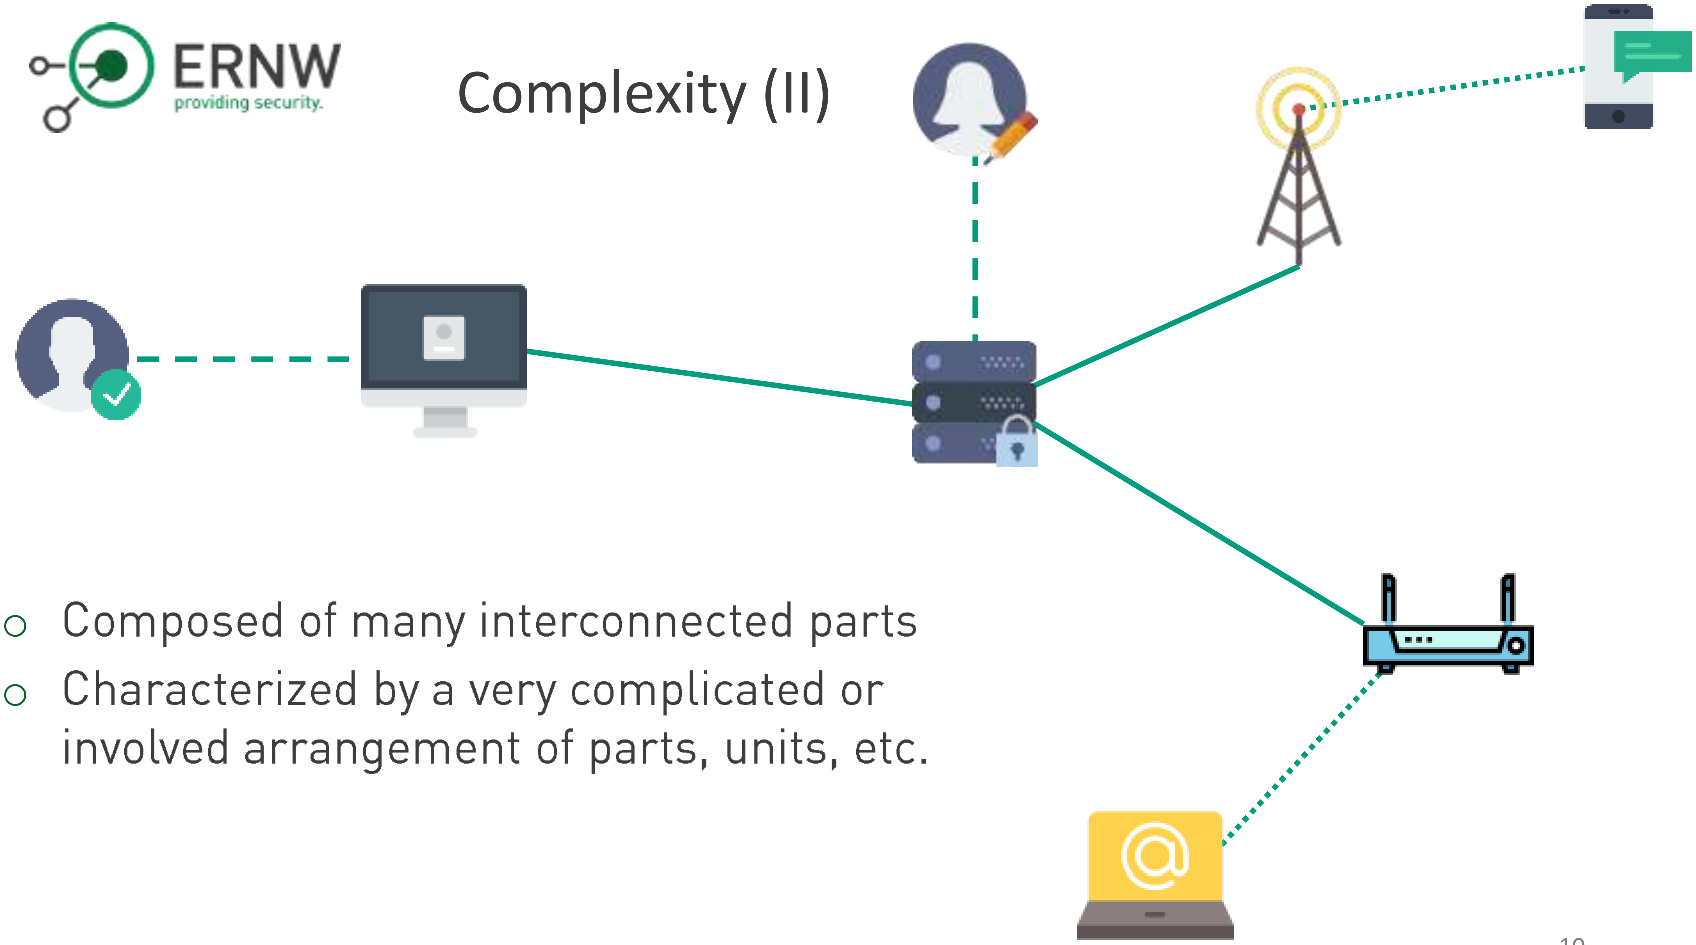

- Composed of many interconnected parts
- Characterized by a very complicated or involved arrangement of parts, units, etc.

ERNW
providing security.

Complexity (III)

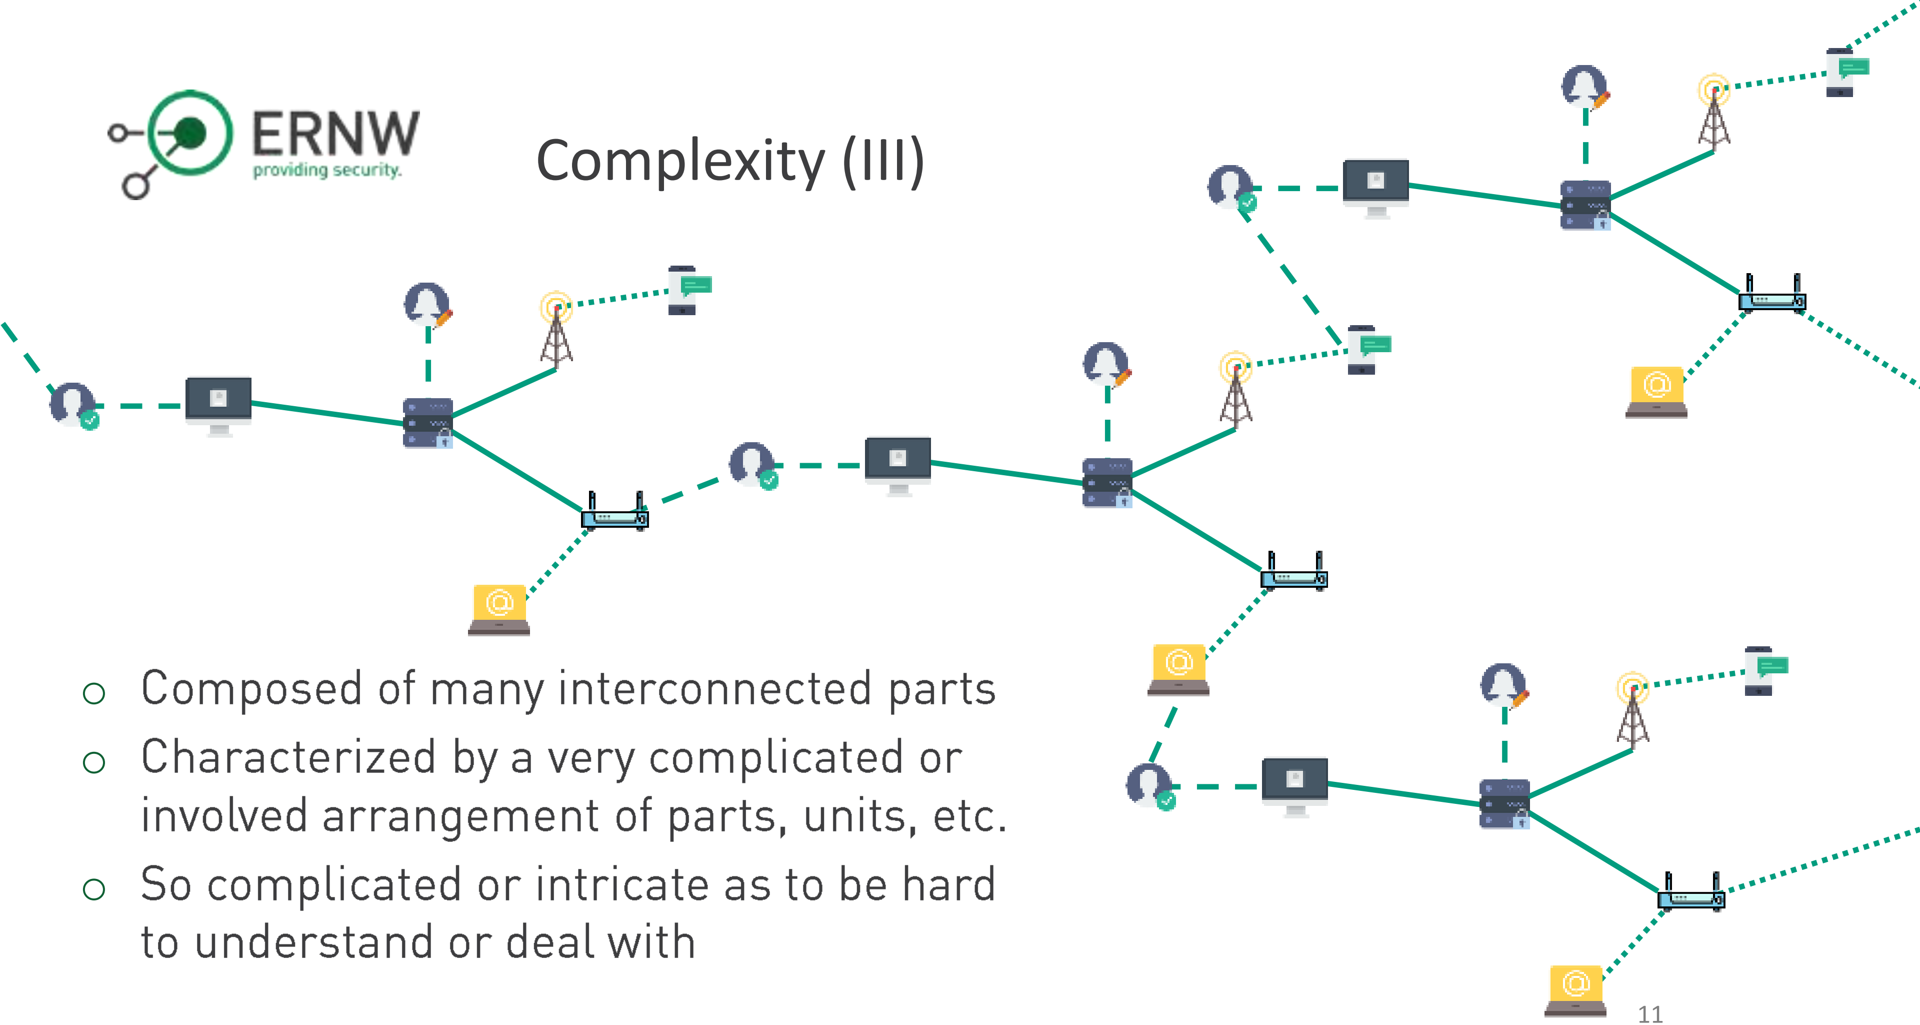

- Composed of many interconnected parts
- Characterized by a very complicated or involved arrangement of parts, units, etc.
- So complicated or intricate as to be hard to understand or deal with

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.

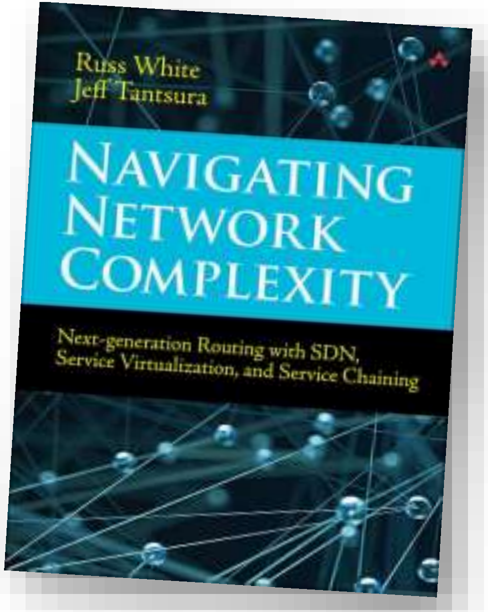
Russ White
Jeff Tantsura
Navigating Network Complexity
Next-generation Routing with SDN, Service Virtualization, and Service Chaining

IPv6 Properties

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*

Different Generations of IPv6 Stacks

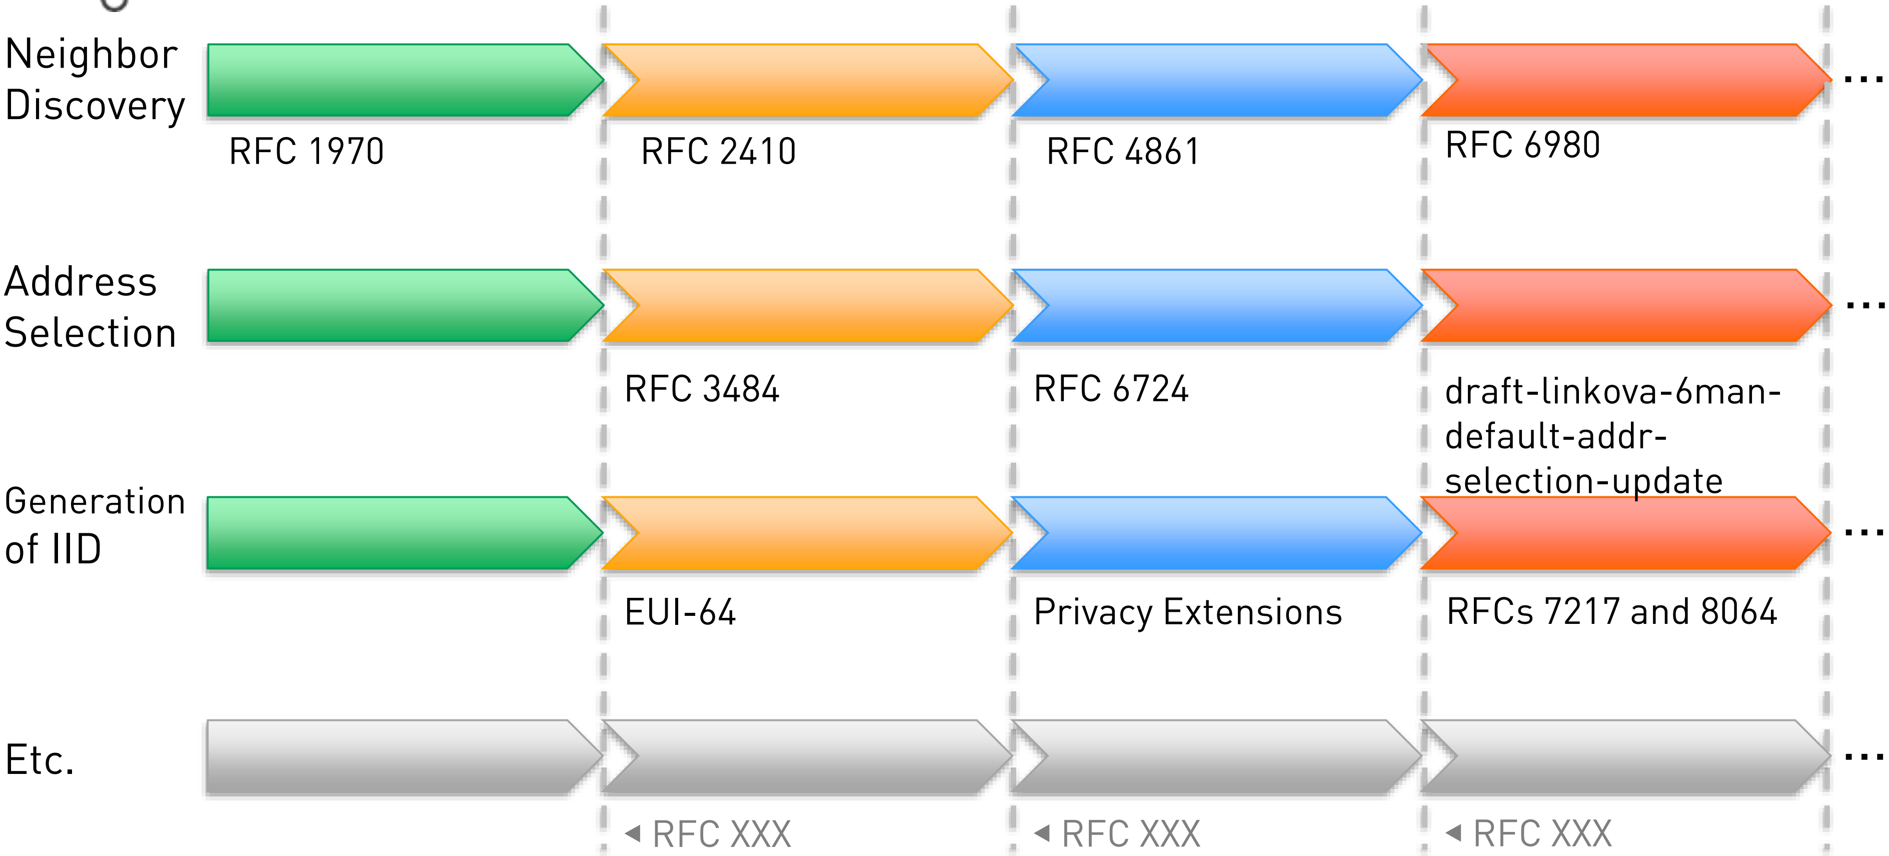

Neighbor Discovery	RFC 1970	RFC 2410	RFC 4861	RFC 6980	...
Address Selection		RFC 3484	RFC 6724	draft-linkova-6man-default-addr-selection-update	...
Generation of IID		EUI-64	Privacy Extensions	RFCs 7217 and 8064	...
Etc.		◂ RFC XXX	◂ RFC XXX	◂ RFC XXX	...

Focus on Four of Them

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

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...

Security Implications of Multicast

- In general higher complexity and more state to be maintained on devices.
- In IPv6's case there also the unfortunate relationship with MLD.
  - New vector for OS fingerprinting
  - Complex in itself => code level vulnerabilities
    - E.g. CVE-2014-0705
  - See also #TR15 talk on MLD and
    - `https://insinuator.net/tag/mld/`

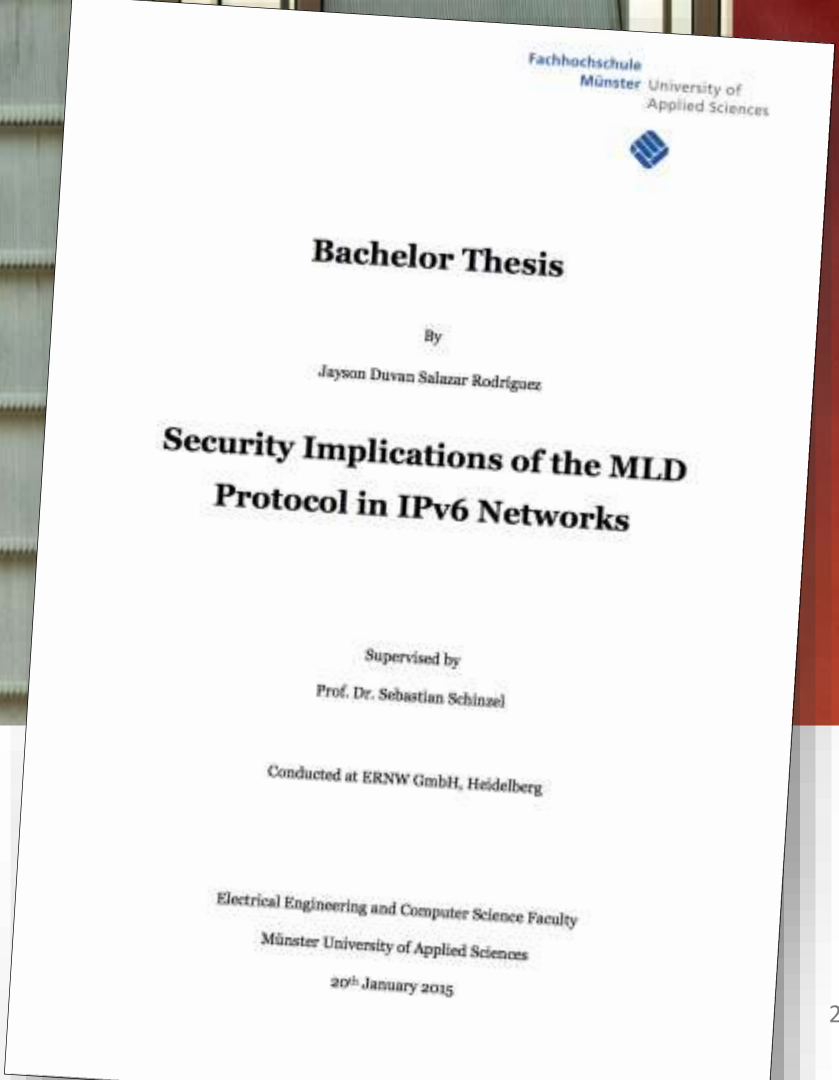

Fachhochschule Münster University of Applied Sciences

Bachelor Thesis

By

Jayson Duvan Salazar Rodriguez

Security Implications of the MLD Protocol in IPv6 Networks

Supervised by

Prof. Dr. Sebastian Schinzel

Conducted at ERNW GmbH, Heidelberg

Electrical Engineering and Computer Science Faculty

Münster University of Applied Sciences

20th January 2015

https://www.its.fh-muenster.de/doc/Security_Implications_of_MLD_in_IPv6_Networks.pdf

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…

◆ LLA

◆ GUA

◆ ULA

Multiple Address Types / Problems

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

- From an offense perspective
  - There might be several ways to interact with a device which
    - Haven‘t been thought of by an implementor.
    - Haven‘t been thought of by an operator who implemented ACLs....

While We're at It:

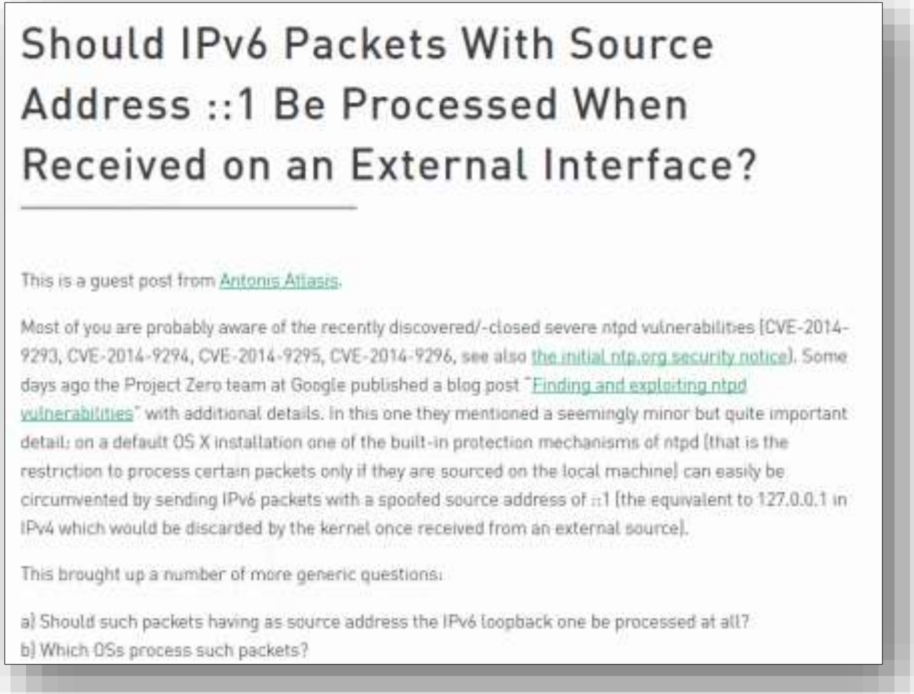

Should IPv6 Packets With Source Address ::1 Be Processed When Received on an External Interface?

This is a guest post from Antonis Atlasis.

Most of you are probably aware of the recently discovered/-closed severe ntpd vulnerabilities (CVE-2014-9293, CVE-2014-9294, CVE-2014-9295, CVE-2014-9296, see also the initial ntp.org security notice). Some days ago the Project Zero team at Google published a blog post "Finding and exploiting ntpd vulnerabilities" with additional details. In this one they mentioned a seemingly minor but quite important detail: on a default OS X installation one of the built-in protection mechanisms of ntpd (that is the restriction to process certain packets only if they are sourced on the local machine) can easily be circumvented by sending IPv6 packets with a spoofed source address of ::1 (the equivalent to 127.0.0.1 in IPv4 which would be discarded by the kernel once received from an external source).

This brought up a number of more generic questions:

a) Should such packets having as source address the IPv6 loopback one be processed at all?
b) Which OSs process such packets?

https://insinuator.net/2015/01/should-ipv6-packets-with-source-address-1-be-processed-when-received-on-an-external-interface/

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.”

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.

RA Provisioning

- M-Flag – Stateful DHCPv6 to acquire IPv6 address
- O-Flag – Stateless DHCPv6 in addition to SLAAC
- Preference Bits – Low, Med, High
- Router Lifetime – Must be >0 for Default
- Options - Prefix Information, Length, Flags
- L bit –Host installs the prefix as On Link
- A bit – Set to 0 for DHCP to work properly

```
Type: 134 (RA)
Code: 0
Checksum: 0xff78 [correct]
Cur hop limit: 64
∞ Flags: 0x84
    1... .... = Managed (M flag)
     .0.. .... = Not other (O flag)
     ..0. .... = Not Home (H flag)
     ...0 1... = Router pref: High
Router lifetime: (s)1800
Reachable time: (ms) 3600000
Retrans timer: (ms) 1000
ICMPv6 Option 3 (Prefix Info)
Prefix length: 64
∞ Flags: 0x80
    1... .... = On link (L Bit)
     .1.. .... = No Auto (A Bit)
Prefix: 2001:0db8:4646:1234::/64
```

IPv6's Trust Model

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

© 2017 WOODSTOCK.COM

Security Problems of “Elections”

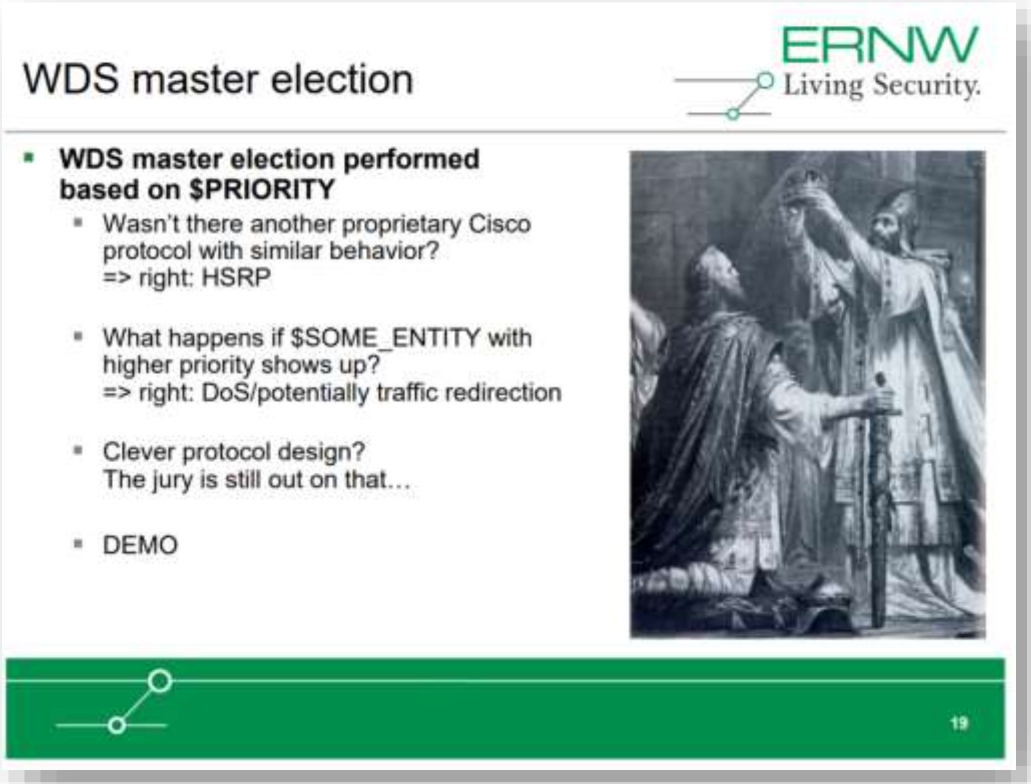

ERNW Living Security.

WDS master election

- **WDS master election performed based on $PRIORITY**
  - Wasn’t there another proprietary Cisco protocol with similar behavior?
    => right: HSRP
  - What happens if $SOME_ENTITY with higher priority shows up?
    => right: DoS/potentially traffic redirection
  - Clever protocol design?
    The jury is still out on that…
  - DEMO

19

https://media.blackhat.com/bh-eu-10/presentations/Mende_Rey/BlackHat-EU-2010-Mende-Rey-Cisco-Wlansec-slides.pdf

Do It Like Jim

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/

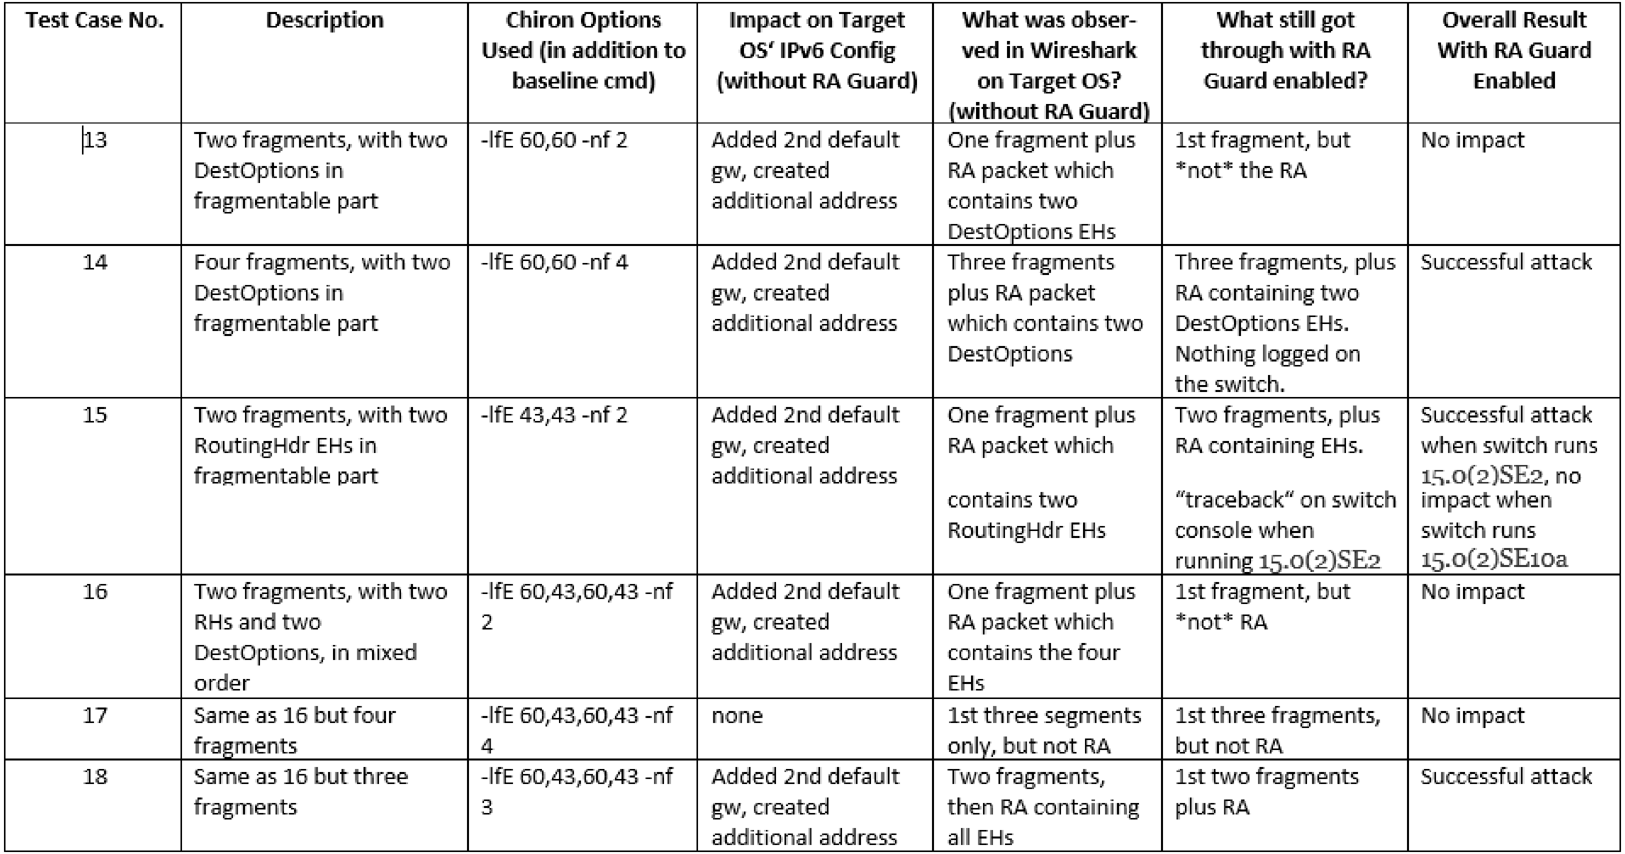

Test Case No.	Description	Chiron Options Used (in addition to baseline cmd)	Impact on Target OS‘ IPv6 Config (without RA Guard)	What was obser-ved in Wireshark on Target OS? (without RA Guard)	What still got through with RA Guard enabled?	Overall Result With RA Guard Enabled
13	Two fragments, with two 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

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/Pru5BRrlmz0
- Most probably more state needed

What an IPv6 Datagrams Looks Like…

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
Might Need This

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

RFC 761

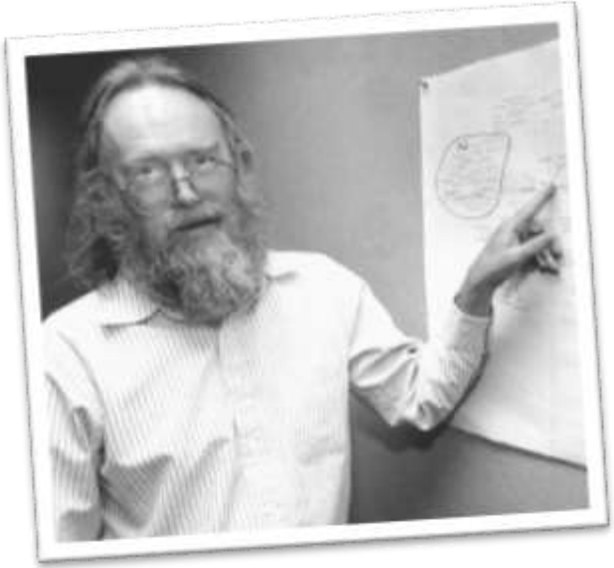

Once Upon a Time…

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…

black hat

ERNW

Evasion of High-End IDPS Devices in the Age of IPv6

Antonios Atlasis

secfu.net
aatlasis@secfu.net

Enno Rey

ERNW GmbH
erey@ernw.de

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

ERNW
providing security.

Conclusions

- From an offense perspective IPv6 is interesting
  - Different architecture => operational security has to adapt (which might not yet be the case).
  - Due its flexibility re: packet format ;-)
- From a protocol development perspective IPv6 is interesting as ... case study ;-) #fail
- From a security research perspective IPv6 is interesting
  - Many RFCs, many different implementations etc.

There’s never enough time…

THANK YOU…

…for yours!

@Enno_Insinuator

erey@ernw.de

ernw.de

insinuator.net

Slides available soon.

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

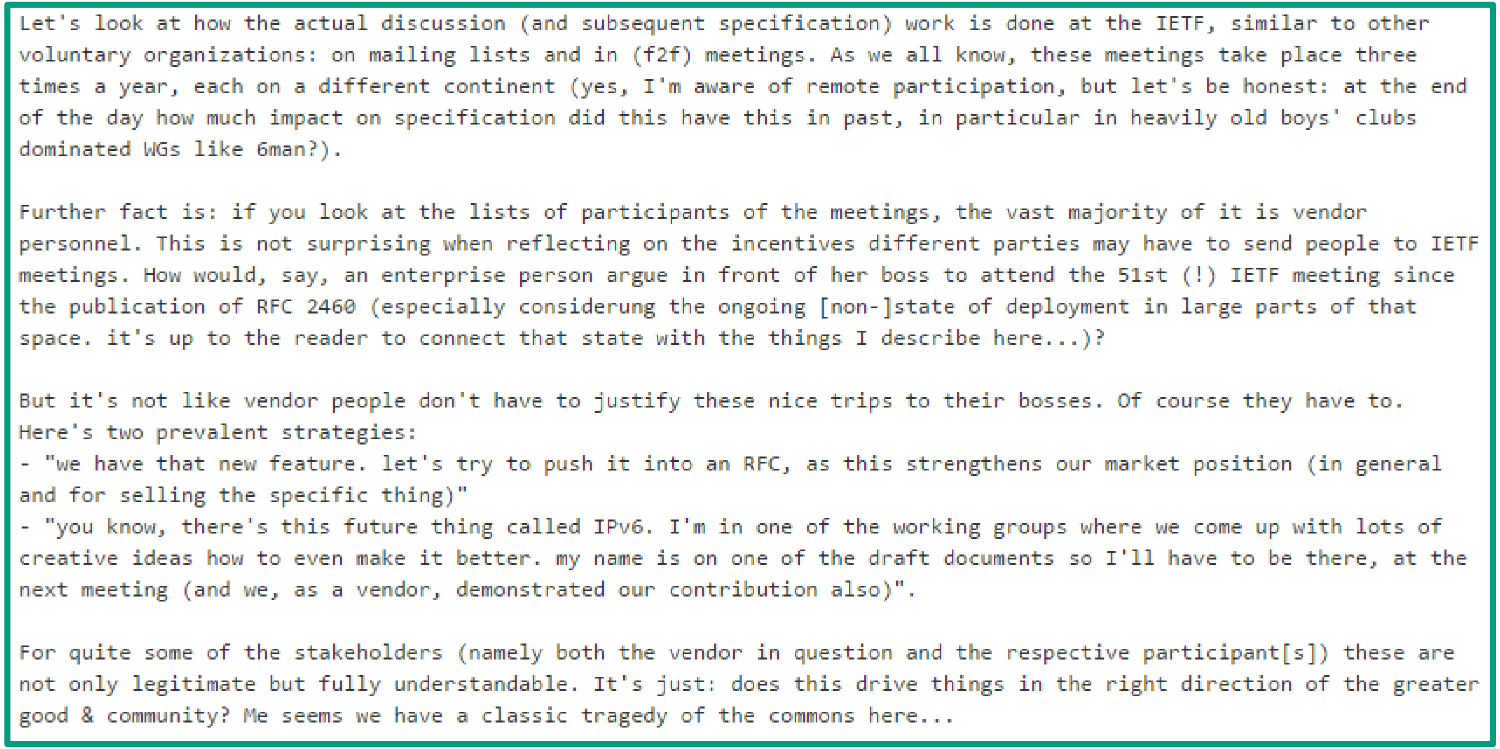

Let's look at how the actual discussion (and subsequent specification) work is done at the IETF, similar to other voluntary organizations: on mailing lists and in (f2f) meetings. As we all know, these meetings take place three times a year, each on a different continent (yes, I'm aware of remote participation, 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

Sources

As indicated on slides.

Image Source:

- Icons made by Freepik from www.flaticon.com