The Layer-2 Insecurities of IPv6 *and the Mitigation Techniques*

Eric Vyncke

Cisco, Consulting Engineering

Distinguished Engineer

evyncke@cisco.com

Eric.Vyncke@ipv6council.be

No Doubt Anymore: IPv4 is Out...

BBC

News | Sport | Weather | Travel | Future

NEWS TECHNOLOGY

Home | UK | Africa | Asia | Europe | Latin America | Mid-East | US & Canada | Business | Health | Sci/Enviro

14 September 2012 Last updated at 15:08 GMT

Share

Europe hits old internet address limits

By Mark Ward
Technology correspondent, BBC News

Europe has almost exhausted its stock of old-style internet addresses.

Strict rationing of these addresses - called IPv4 - has been started by the body that hands them out in Europe.

From now on, companies can only make one more application for IPv4 addresses and, if successful, will only get 1,024 of them.

In addition, any application for more old addresses must demonstrate how an organisation is using the new, replacement, addressing scheme.

Europe's stock of old-style net addresses has effectively run dry.

Pool Reaches Final /8

pril 2011, the APNIC pool reached the Final /8 IPv4
IPv4 exhaustion in the Asia Pacific.

ion 9.10 in "Policies for IPv4 address space

vide IPv4 address space for new entrants to the

t holders will be entitled to receive a maximum
space.

y members to deploy IPv6 within their organizations.
ling IPv6 deployment, statistics, training, and related

g for quite some time," states Raúl Echeberría,
the five RIRs. "The future of the Internet is in IPv6.

... And IPv6 in In ;-)

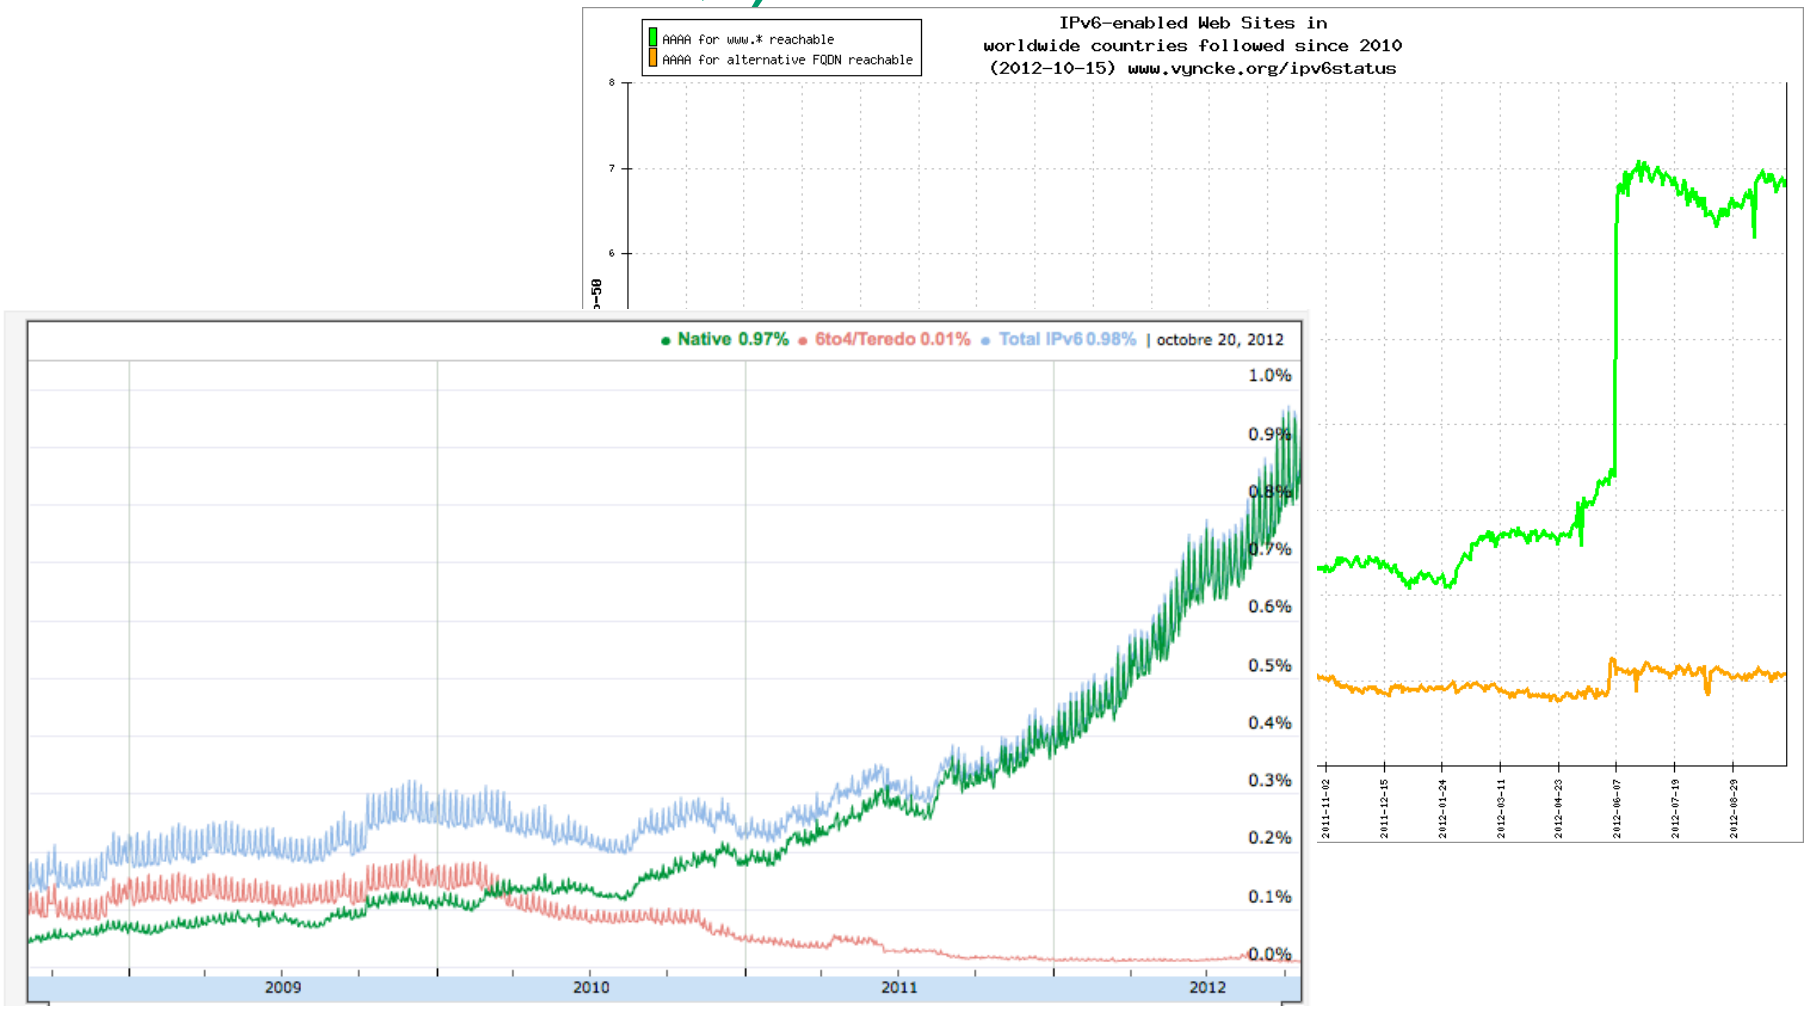

Source: http://www.google.com/intl/en/ipv6/statistics/ http://www.vyncke.org/ipv6status

IPv6 in One Slide

- IPv6 is IPv4 with larger addresses
  - 128 bits vs. 32 bits
  - **NAT no more needed => easier for applications**
    - **Simpler hence more security**
- Data-link layer unchanged: Ethernet, xDSL, …
- Transport layer unchanged: UDP, TCP, …
- Applications “unchanged”: HTTP, SSL, SMTP, …
- IPv6 is not really BETTER than IPv4 because it is ‘new’
  - IPv6 has been specified in 1995…
  - IPsec is identical in IPv4 & IPv6
  - **Only benefit is a much larger address space**

IPv6 Myths: Better, Faster, More Secure

Sometimes, newer means better and more secure

Sometimes, experience IS better and safer!

Fundamentals On Neighbor Discovery (ND)

For Your Reference

- Defined in:
  - RFC 4861 Neighbor Discovery for IP Version 6 (IPv6)
  - RFC 4862 IPv6 Stateless Address Auto-configuration
  - RFC 3971 Secure Neighbor Discovery etc.
- Used for:
  - Router discovery
  - IPv6 Stateless Address Auto Configuration (SLAAC)
  - IPv6 address resolution (replaces ARP)
  - Neighbor Unreachability Detection (NUD)
  - Duplicate Address Detection (DAD)
  - Redirection
- Operates above ICMPv6
  - Relies heavily on (link-local scope) multicast, combined with Layer 2 Multicast
- Works with ICMP messages and messages “options”

Networks are Sand Castles…

Courtesy of Curt Smith

Attacking Stateless Address Autoconfiguration with Rogue RA

Rogue Router Advertisement

Router Advertisements contains:
-Prefix to be used by hosts
-Data-link layer address of the router
-Miscellaneous options: MTU, DHCPv6 use, …

RA w/o Any Authentication Gives Exactly Same Level of Security as DHCPv4 (None)

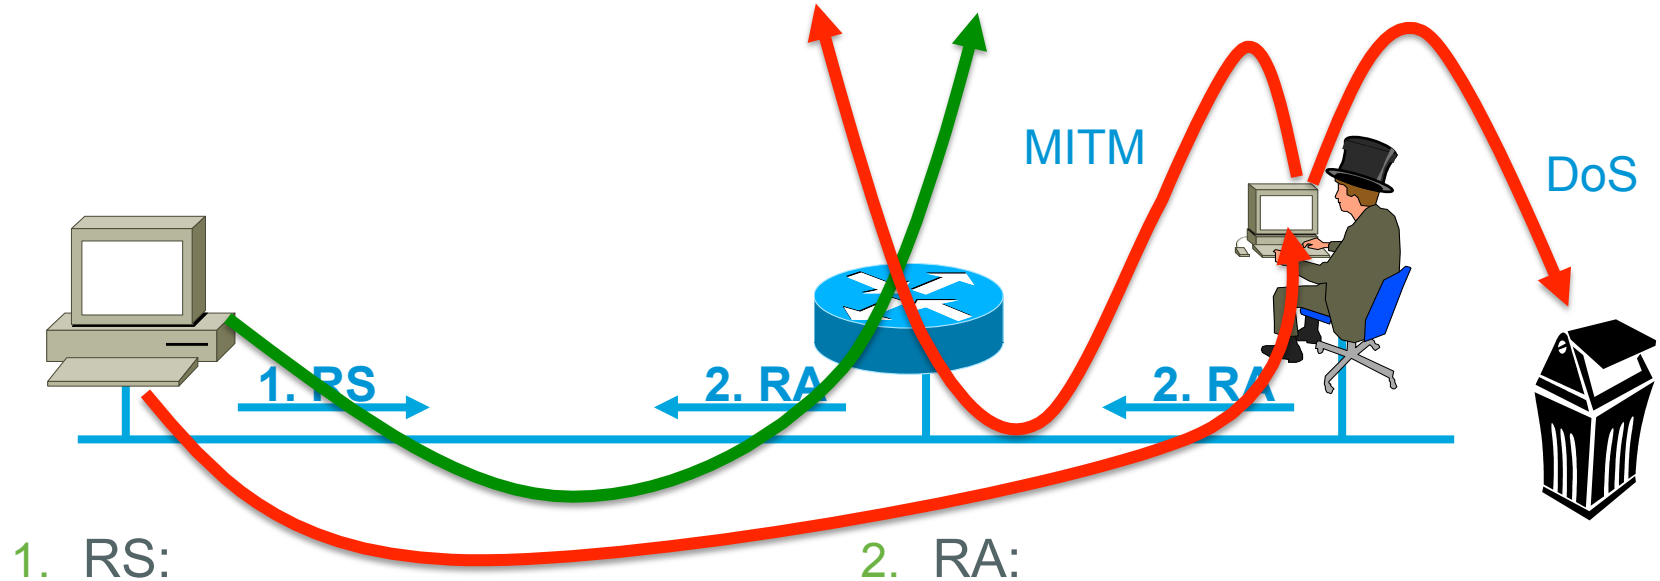

1. RS:
   Data = Query: please send RA
2. RA:
   Data= options, **prefix**, lifetime, A+M+O flags

Effect of Rogue Router Advertisements

- Devastating:
  - Denial of service: all traffic sent to a black hole
  - Man in the Middle attack: attacker can intercept, listen, modify unprotected data
- **Also affects legacy IPv4-only network** with IPv6-enabled hosts
- Most of the time from non-malicious users
- Requires layer-2 adjacency (some relief…)

- The **major blocking factor** for enterprise IPv6 deployment
- Special from THC: RA flood with different prefixes => crash Windows and a few other OS ☹ Still in 2012!

Bored at BRU Airport on Sunday at 22:30…

```
$ ifconfig en1
en1: flags=8863<UP,BROADCAST                                  T> mtu 1500
     ether 00:26:bb:xx:xx:
     inet6 fe80::226:bbff:f                                   d 0x6
     inet 10.19.19.118 netma                                  19.19.255
     media: autoselect
     status: active
```

Humm…
Is there an IPv6 Network?

```
$ ping6 -I en1 ff02::1%en1
PING6(56=40+8+8 bytes) fe80::226:bbff:fe                      ::1
16 bytes from fe80::226:bb                                    ime=0.140 ms
. . .
16 bytes from fe80::cabc                                      time=402.112 ms
^C
--- ff02::1%en1 ping6 statistics ---
4 packets transmitted, 4 packets received, +142 duplicates, 0.0% packet loss
round-trip min/avg/max/std-dev = 0.140/316.721/2791.178/412.276 ms
```

Humm…
Are there any IPv6 peers?

```
$ ndp -an
Neighbor
2001:
. . .
$ ndp -an | wc
       64
```

Let's have some fun here… Configure a tunnel, enable forwarding and transmit RA

Rogue RA – Mitigation Techniques

Where	What
Routers	Increase “legal” router preference
Hosts	Disabling Stateless Address Autoconfiguration
Routers & Hosts	SeND “Router Authorization”
Switch (First Hop)	Host isolation
Switch (First Hop)	Port Access List (PACL)
Switch (First Hop)	RA Guard

Secure Neighbor Discovery (SeND) RFC 3971

- RFC 3972 Cryptographically Generated Addresses (CGA)
  - IPv6 addresses whose interface identifiers are cryptographically generated from node public key
- SeND adds a signature option to Neighbor Discovery Protocol
  - Using node private key
  - Node public key is sent in the clear (and linked to CGA)
- Very powerful
  - If MAC spoofing is prevented
  - But, not a lot of implementations: Cisco IOS, Linux, some H3C, **third party for Windows (from Hasso-Plattner-Institut in Germany!)**

Cryptographically Generated Addresses CGA RFC 3972 (Simplified)

- Each devices has a RSA key pair (no need for cert)
- Ultra light check for validity
- Prevent spoofing a valid CGA address

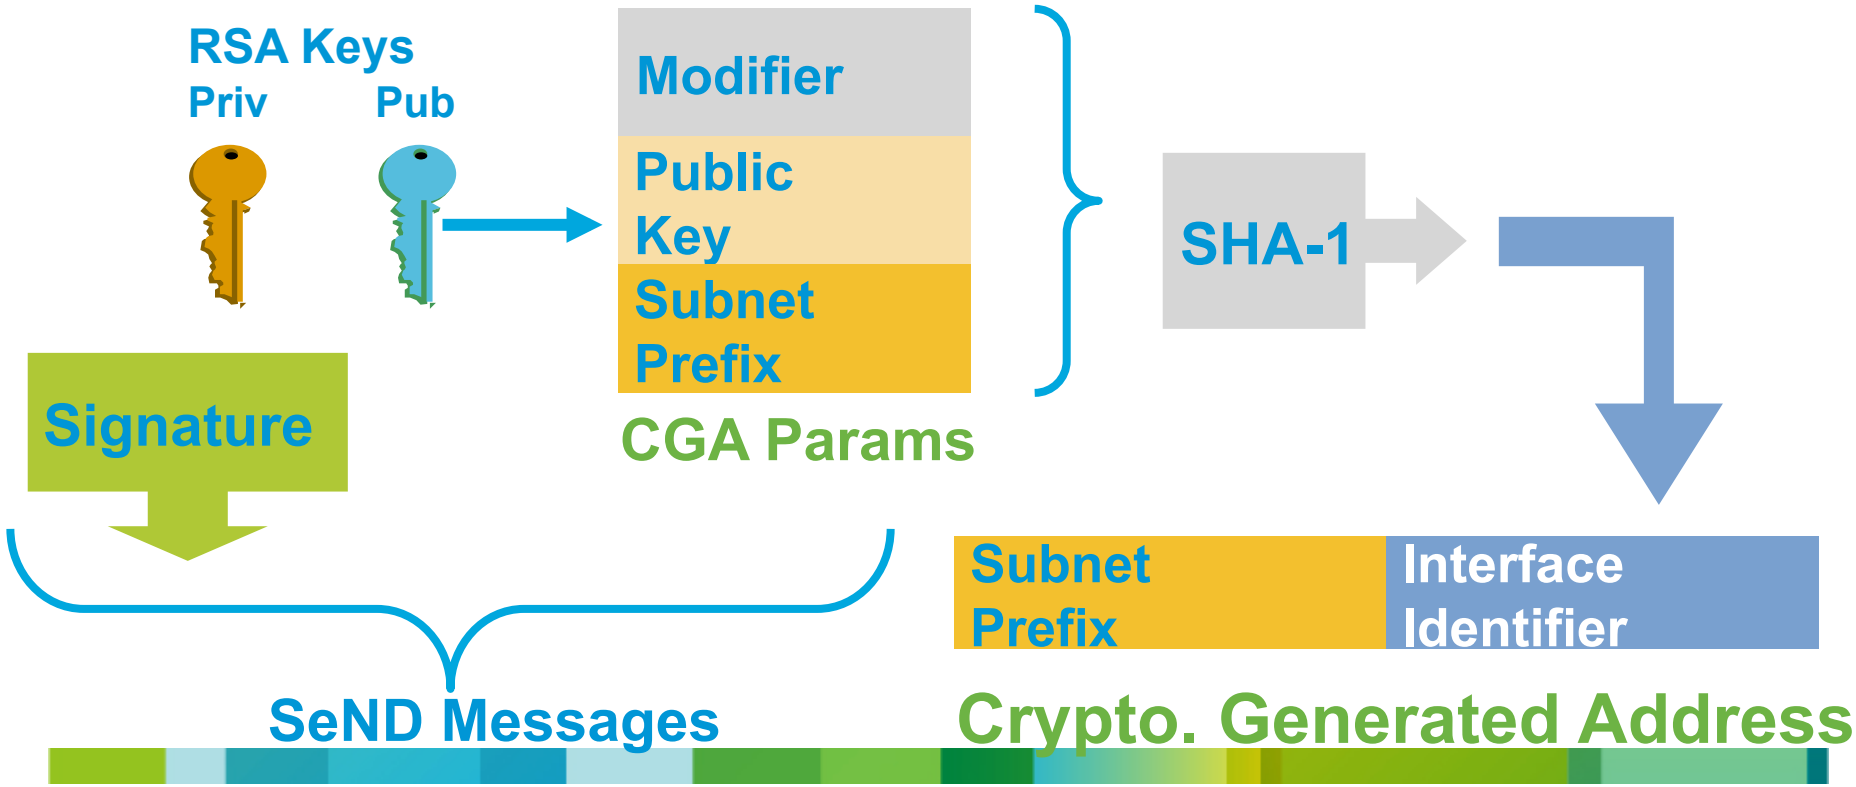

Securing Router Advertisements with SeND

- Adding a X.509 certificate to RA
- Subject Name contains the list of autho v6 prefixes

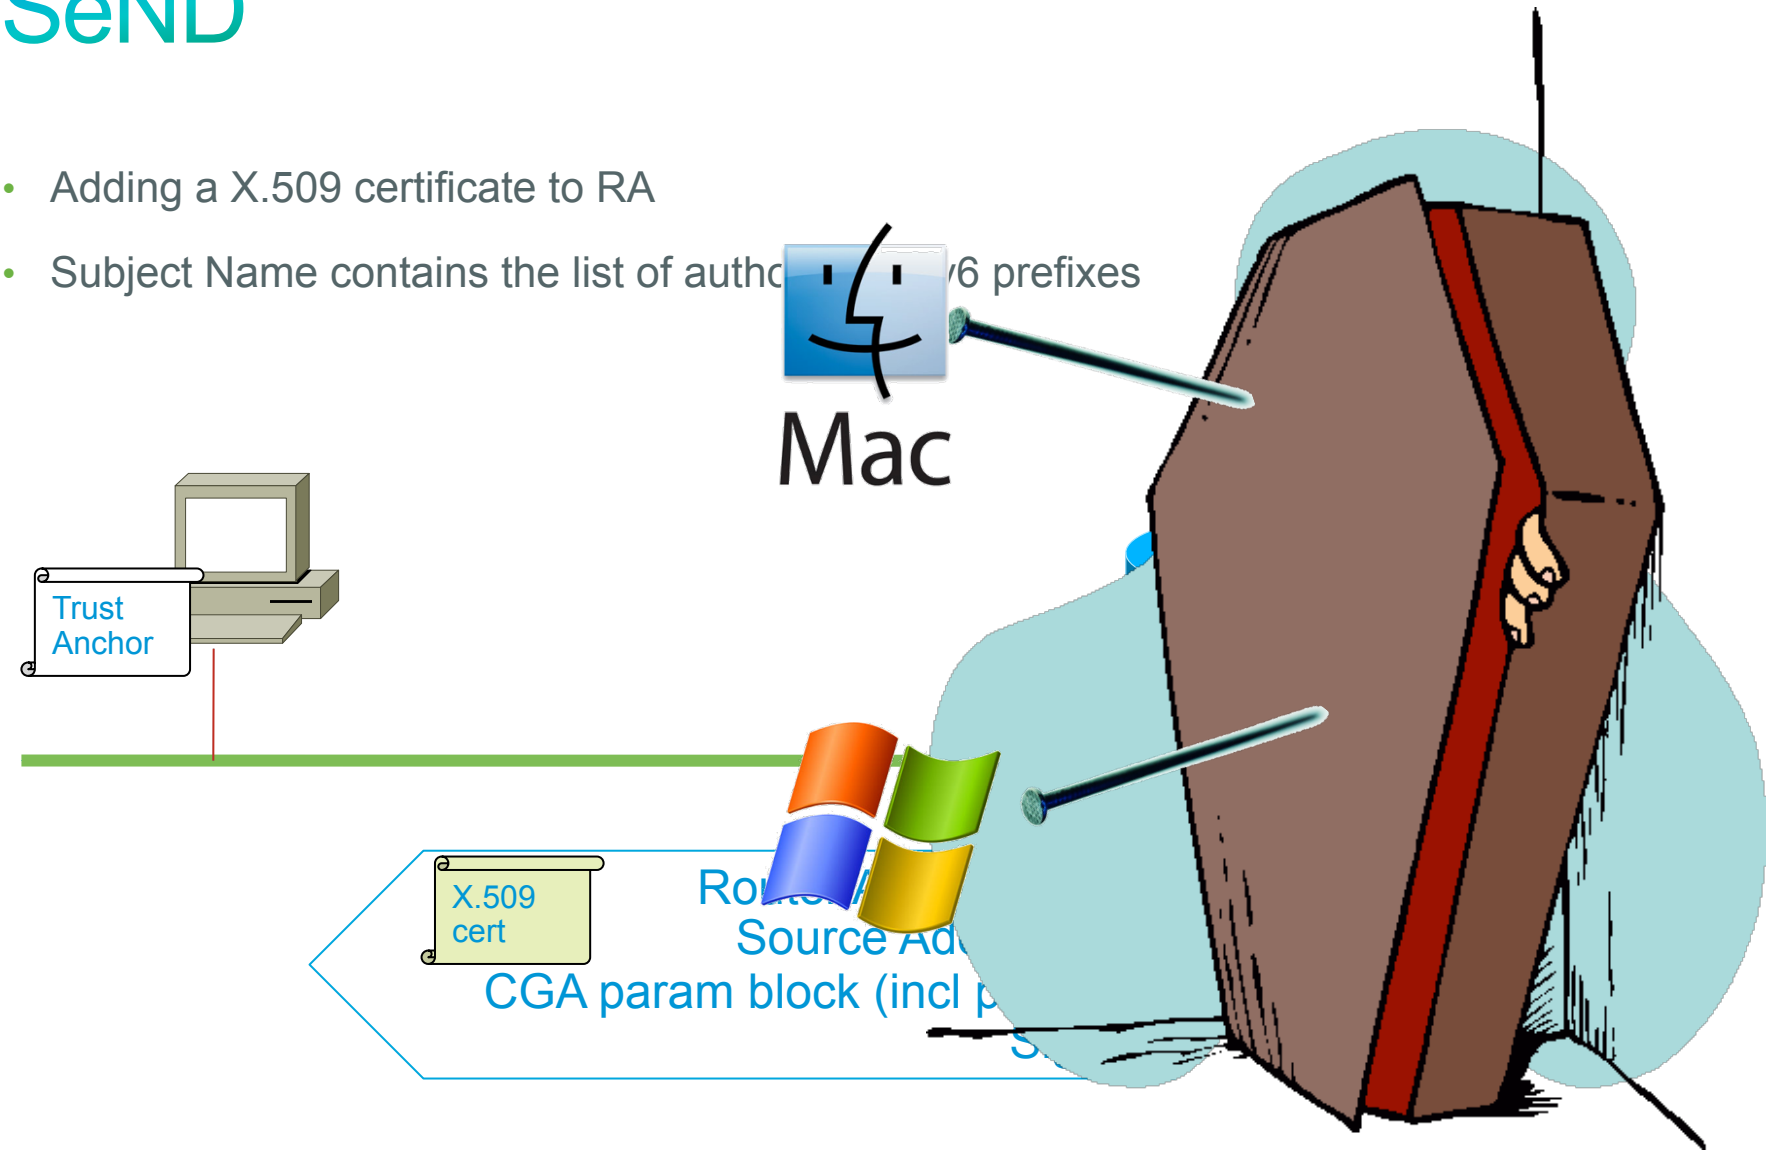

Mitigating Rogue RA: Host Isolation

- Prevent Node-Node Layer-2 communication by using:
  - Private VLANs (PVLAN) where nodes (isolated port) can only contact the official router (promiscuous port)
  - WLAN in ‘AP Isolation Mode’
  - 1 VLAN per host (SP access network with Broadband Network Gateway)
- Link-local multicast (RA, DHCP request, etc) sent only to the local official router: no harm

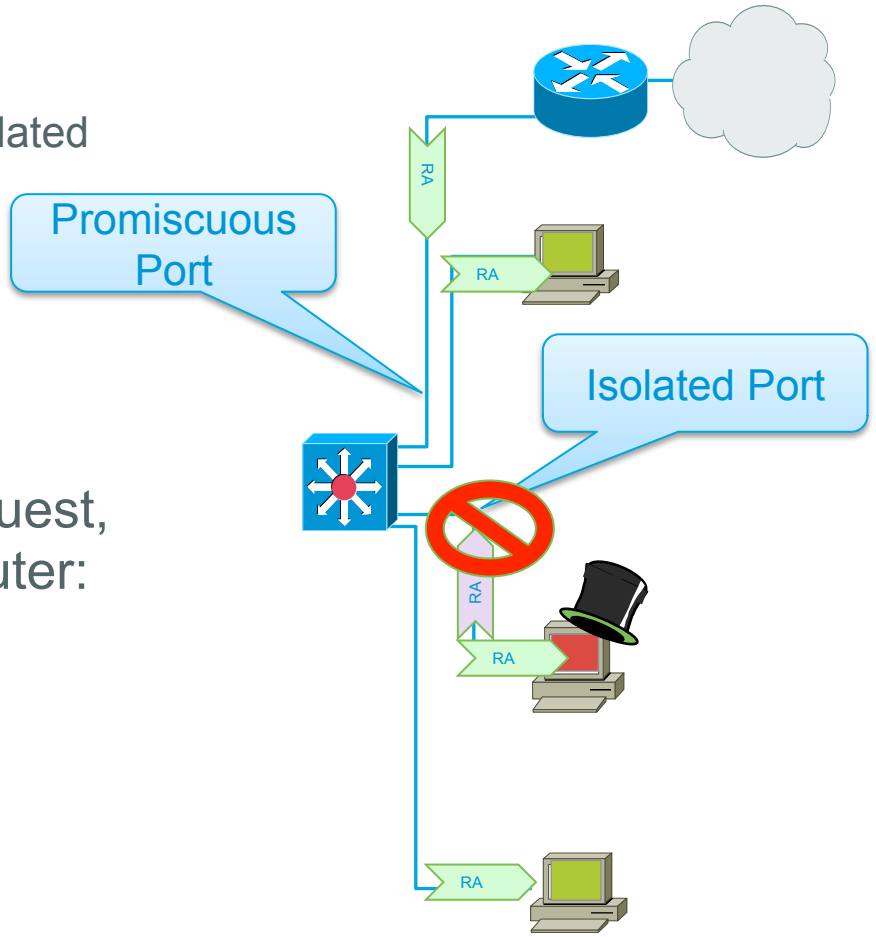

Mitigating Rogue RA: RFC 6105

- **Port ACL** blocks all ICMPv6 RA from hosts

```
interface FastEthernet0/2
  ipv6 traffic-filter ACCESS_PORT in
  access-group mode prefer port
```

- **RA-guard lite** (12.2(33)SXI4 & 12.2(54)SG): also dropping all RA received on this port

```
interface FastEthernet0/2
  ipv6 nd raguard
  access-group mode prefer port
```

- **RA-guard** (12.2(50)SY)

```
ipv6 nd raguard policy HOST device-role host
ipv6 nd raguard policy ROUTER device-role router
ipv6 nd raguard attach-policy HOST vlan 100
interface FastEthernet0/0
  ipv6 nd raguard attach-policy ROUTER
```

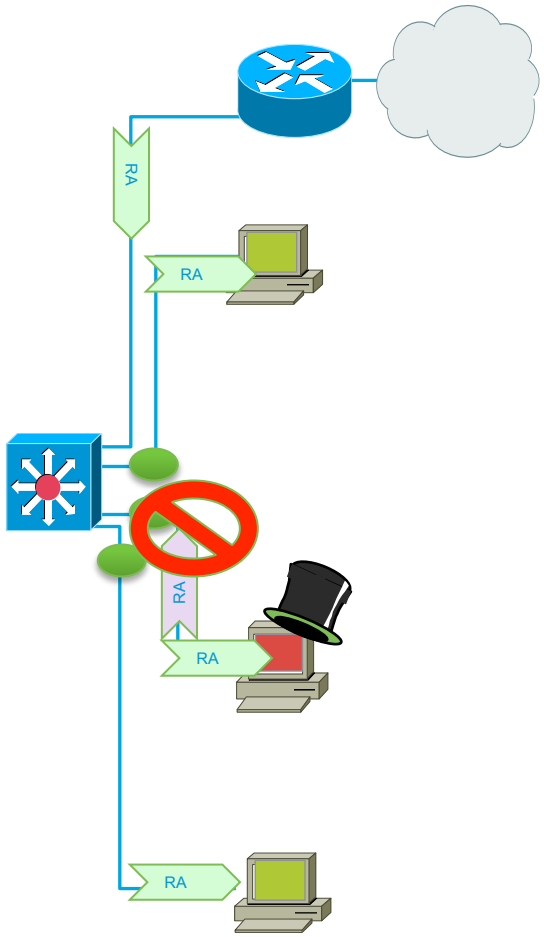

RA-Guard (RFC 6105)

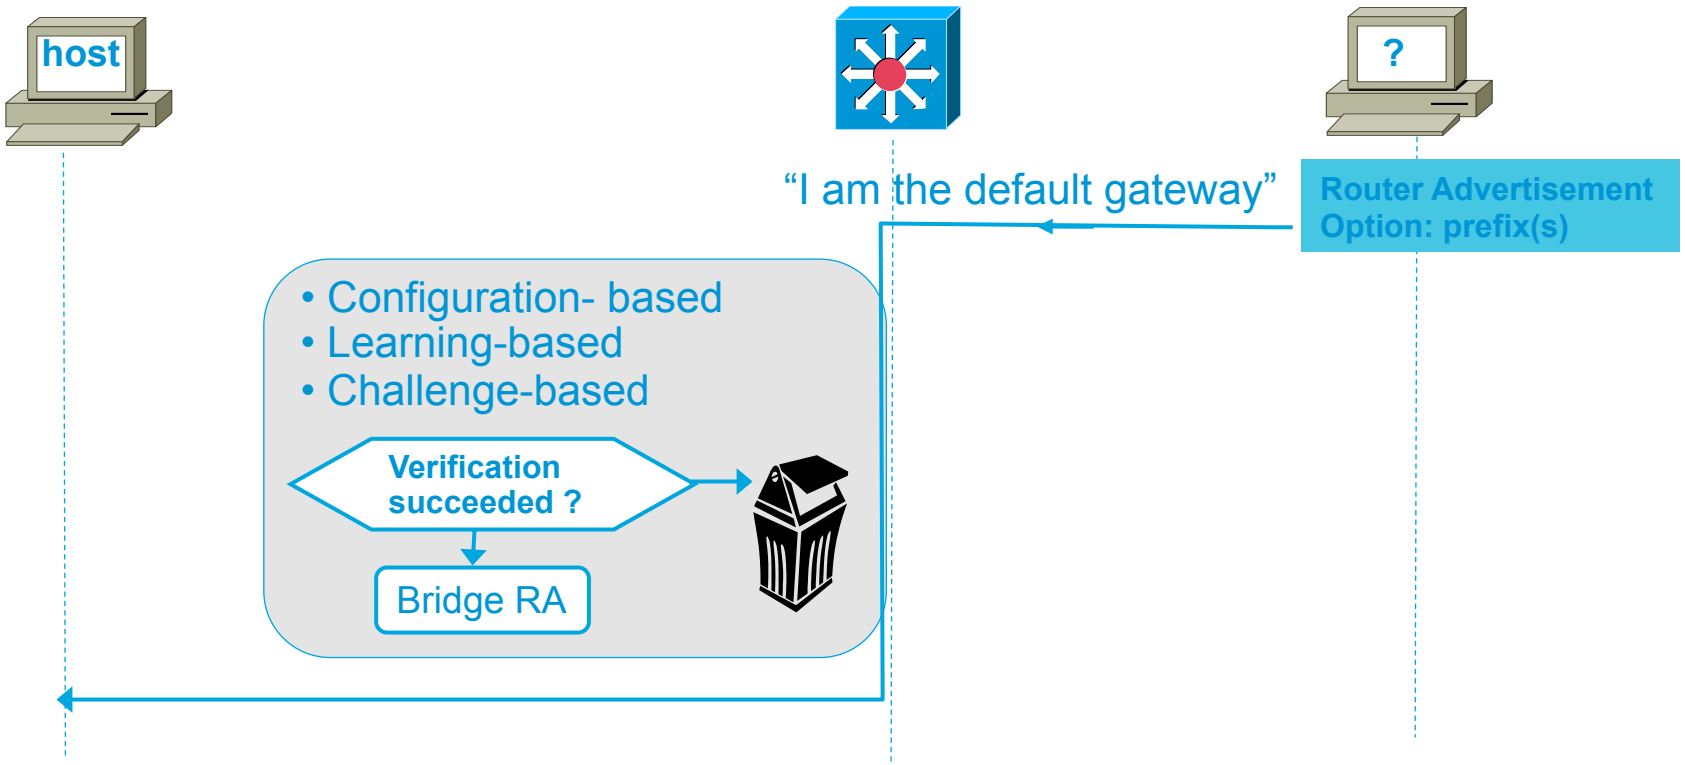

- Switch selectively accepts or rejects RAs based on various criteria's
- Can be ACL based, learning based or challenge (SeND) based.
- Hosts see only allowed RAs, and RAs with allowed content

Here comes Fragmentation…

- Extension headers chain can be so large than it is fragmented!
- RFC 3128 is not applicable to IPv6
- Layer 4 information could be in 2nd fragment

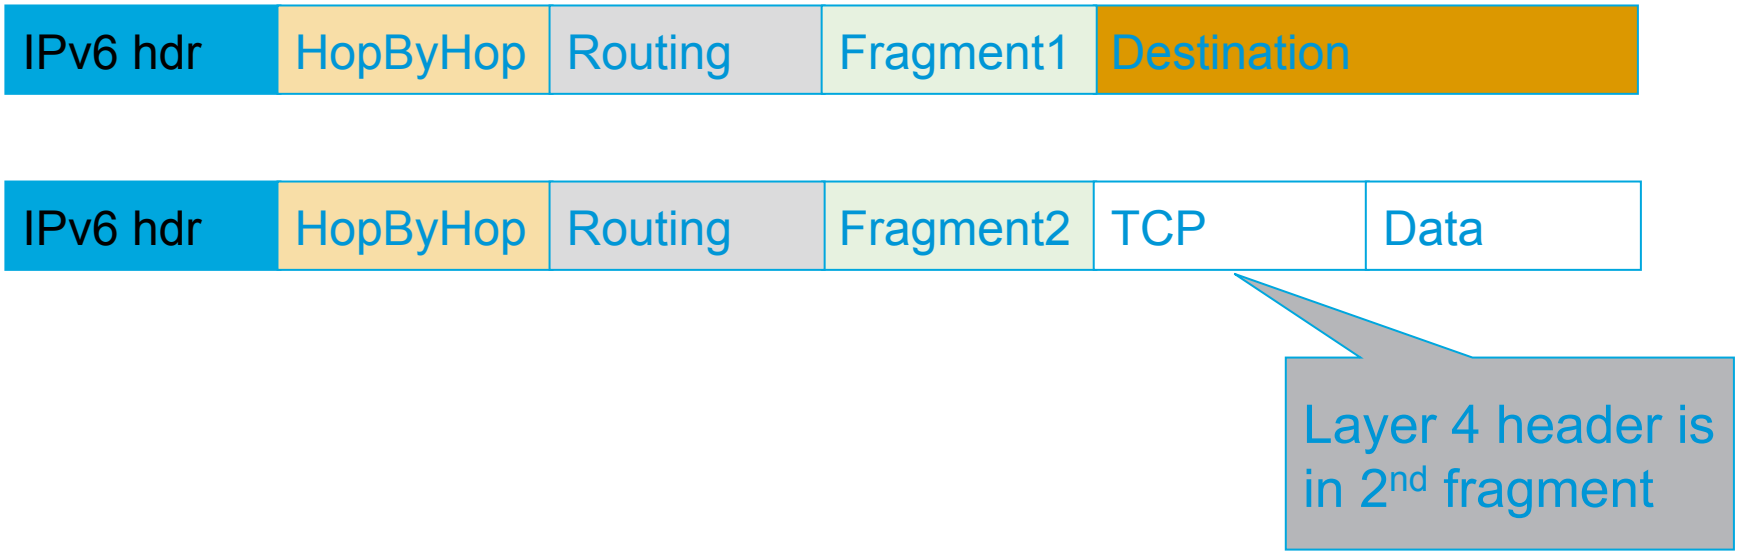

Parsing the Extension Header Chain Fragments and Stateless Filters (RA Guard)

- RFC 3128 is not applicable to IPv6, extension header can be fragmented
- ICMP header could be in 2nd fragment after a fragmented extension header
- RA Guard works like a stateless ACL filtering ICMP type 134
- THC fake_router6 –FD implements this attack which bypasses RA Guard
- ***Partial work-around: block all fragments sent to ff02::1***
  - ***'undetermined-transport' is even better***
  - ***Does not work in a SeND environment (larger packets) but then no need for RA-guard ☺***

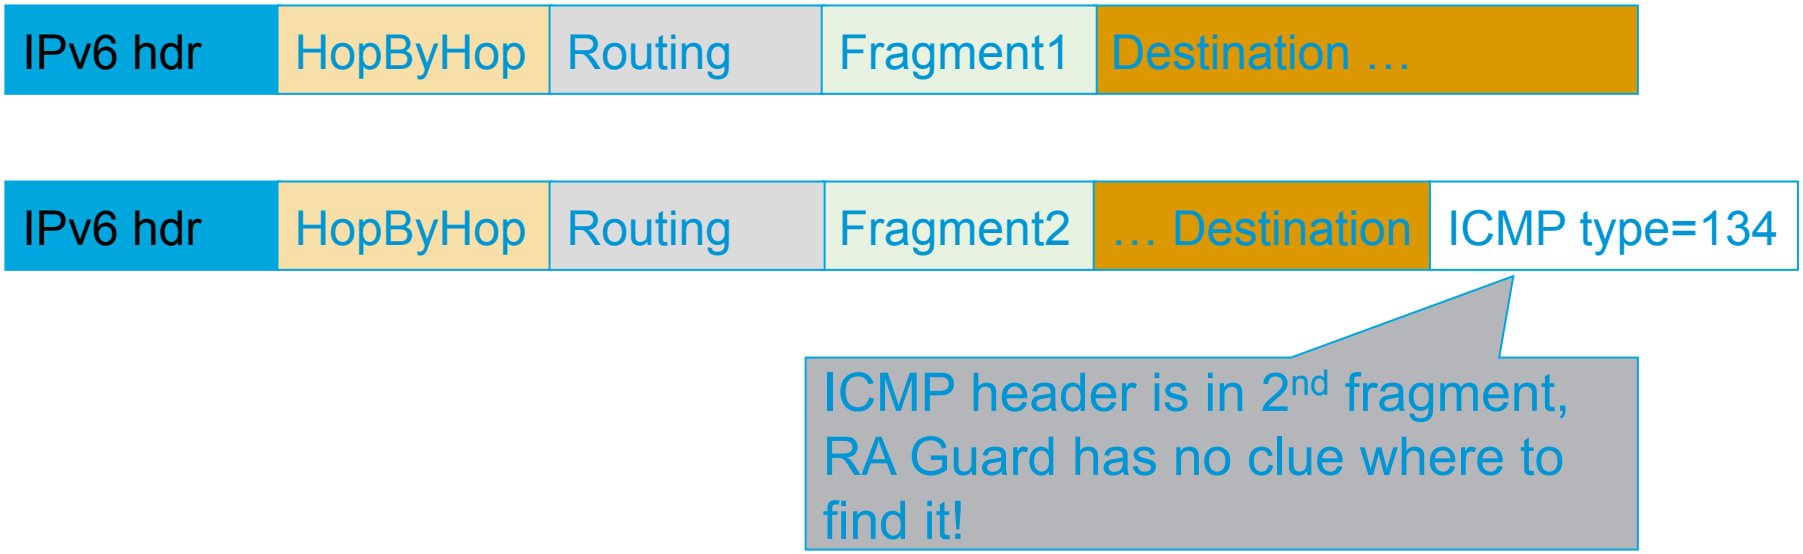

Enabling IPv6 in the IPv4 Data Center The Fool's Way

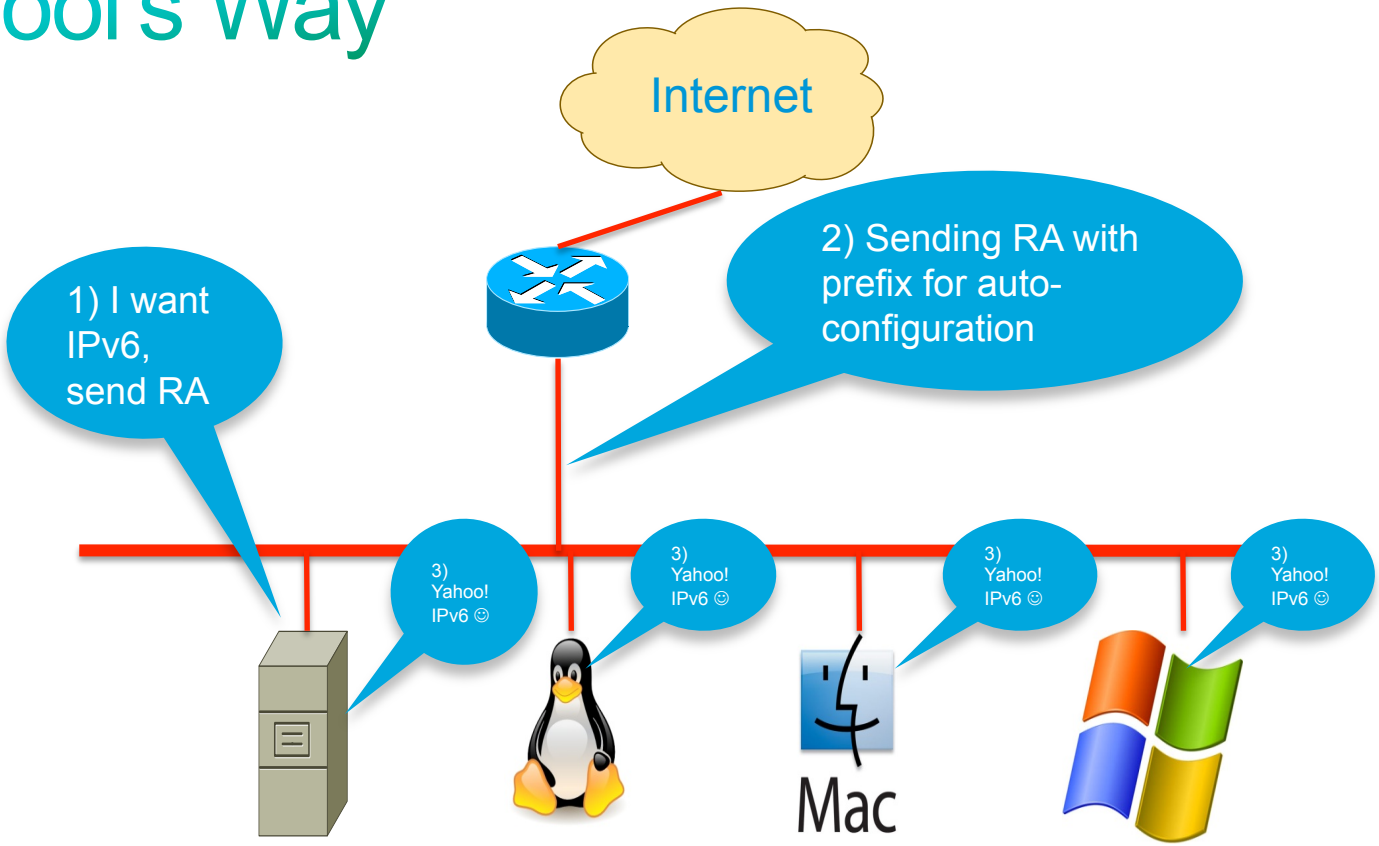

Attacking Neighbor Discovery with NDP Spoofing

Neighbor Discovery Issue#2
Neighbor Solicitation

Src = A
Dst = Solicited-node multicast of B
ICMP type = 135
Data = link-layer address of A
Query: what is your link address?

Src = B
Dst = A
ICMP type = 136
Data = link-layer address of B

A and B Can Now Exchange Packets on This Link

Security Mechanisms Built into Discovery Protocol = None

=> Very similar to ARP

Attack Tool: Parasite6 Answer to all NS, Claiming to Be All Systems in the LAN...

Neighbor Discovery Issue#3
Duplicate Address Detection

Duplicate Address Detection (DAD) Uses Neighbor Solicitation to Verify the Existence of an Address to Be Configured

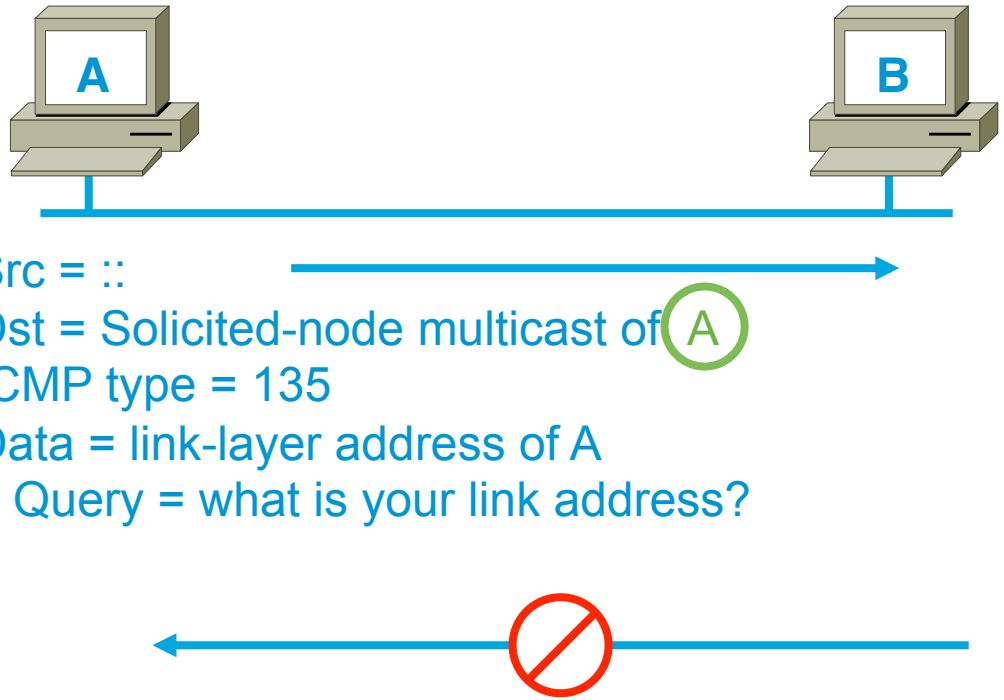

From RFC 4862 5.4:
« ***If a duplicate @ is discovered… the address cannot be assigned to the interface»***
⇔What If: Use MAC@ of the Node You Want to DoS and Claim Its IPv6 @

Attack Tool:
Dos-new-IPv6

Mitigation in IOS:
Configuring the IPv6 address as anycast disables DAD on the interface

Neighbor Advertisement can be Spoofed

- Pretty much like RA: no authentication
  - Any node can ‘steal’ the IP address of any other node
  - Impersonation leading to denial of service or MITM
- Requires layer-2 adjacency
- IETF SAVI Source Address Validation Improvements (work in progress)

NDP Spoofing Mitigations

Where	What
Routers & Hosts	configure static neighbor cache entries
Routers & Hosts	Use CryptoGraphic Addresses (SeND CGA)
Switch (First Hop)	Host isolation
Switch (First Hop)	Address watch • Glean addresses in NDP and DHCP • Establish and enforce rules for address ownership

Securing Neighbor Advertisements with SeND

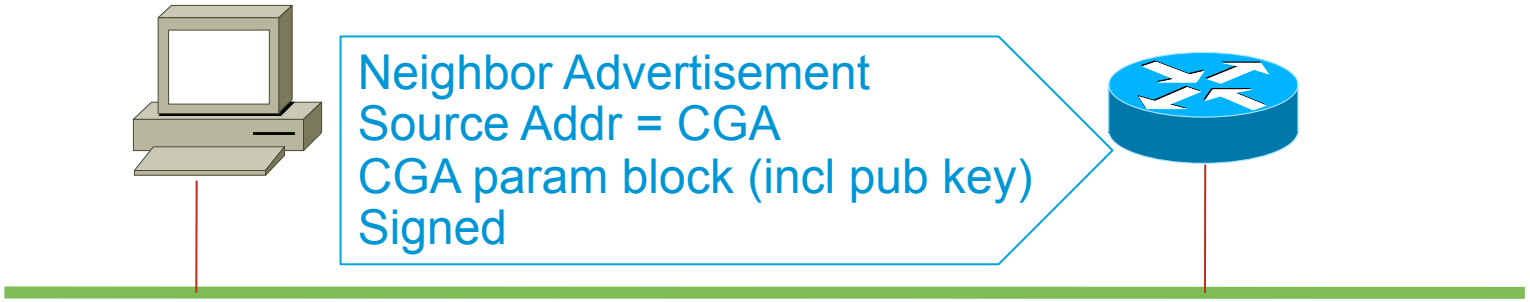

SAVI: How to Learn?

- If a switch wants to enforce the mappings < *IP address, MAC address*> how to learn them?
- Multiple source of information

  SeND: verify signature in NDP messages, then add the mapping

  DHCP: snoop all messages from DHCP server to learn mapping (same as in IPv4)

  NDP: more challenging, but '*first come, first served*'

    The first node claiming to have an address will have it

NDP Spoofing – Mitigation: Binding Integrity at the First Hop

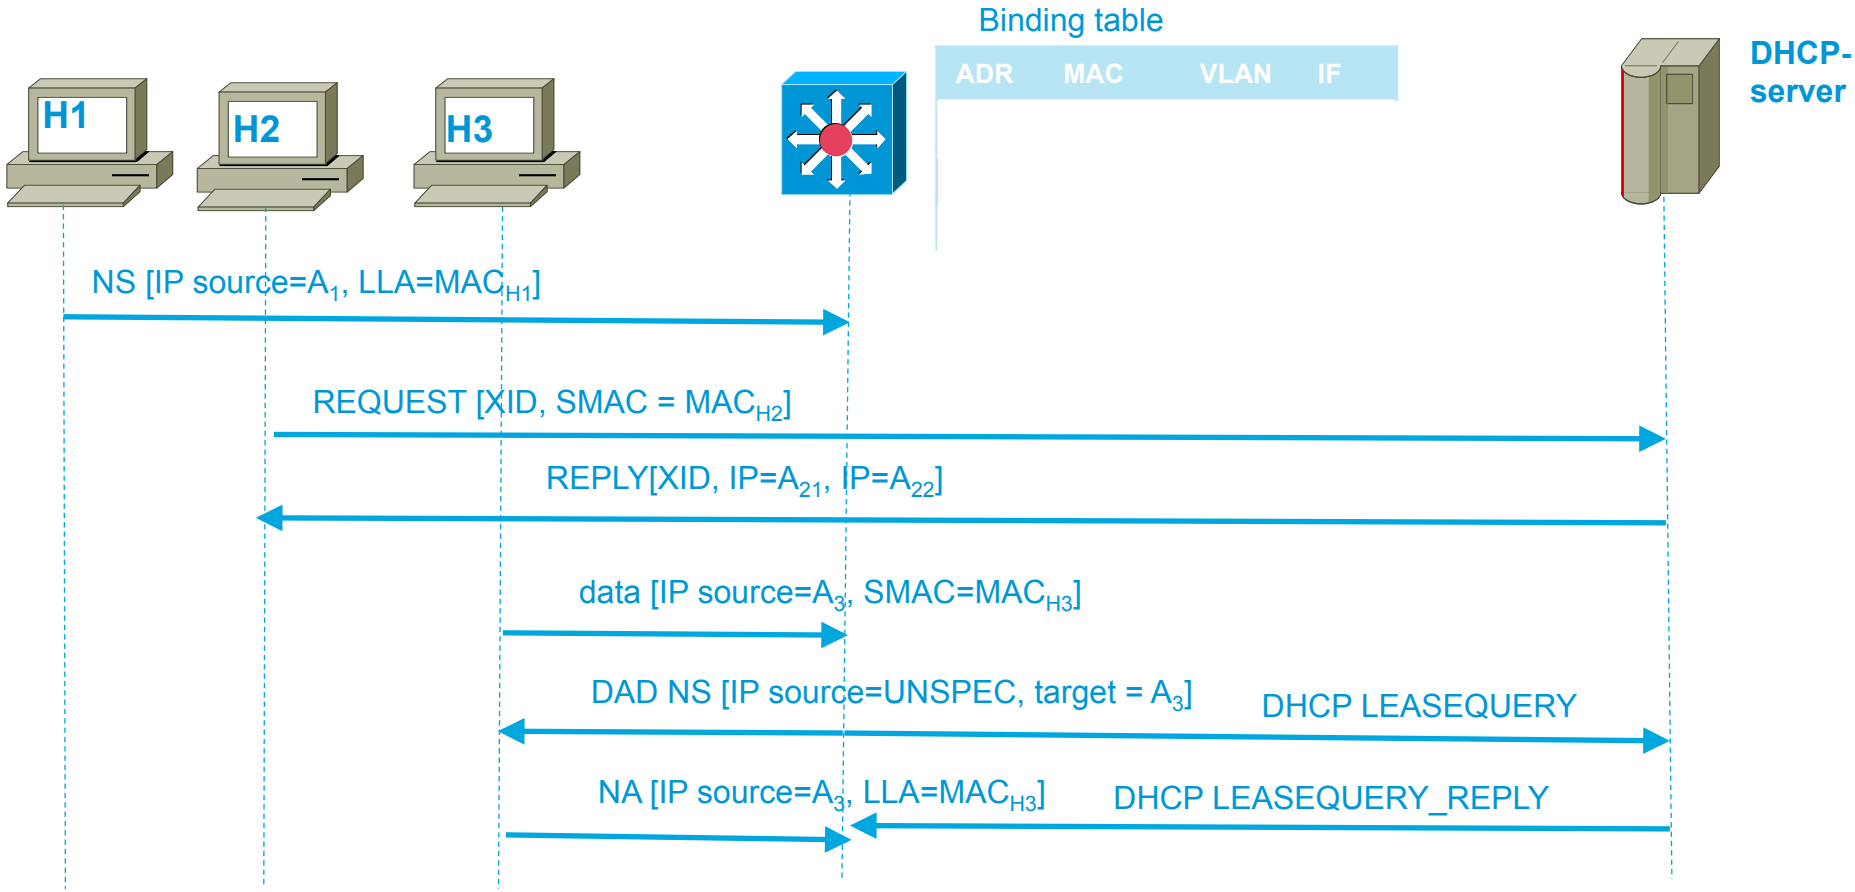

Then, drop all Neighbor Discovery packets not matching the binding...

Exhausting the Neighbor Cache

Scanning Made Bad for CPU
Remote Neighbor Cache Exhaustion

- Remote router CPU/memory DoS attack if aggressive scanning

  Router will do Neighbor Discovery... And waste CPU and memory

- Local router DoS with NS/RS/…

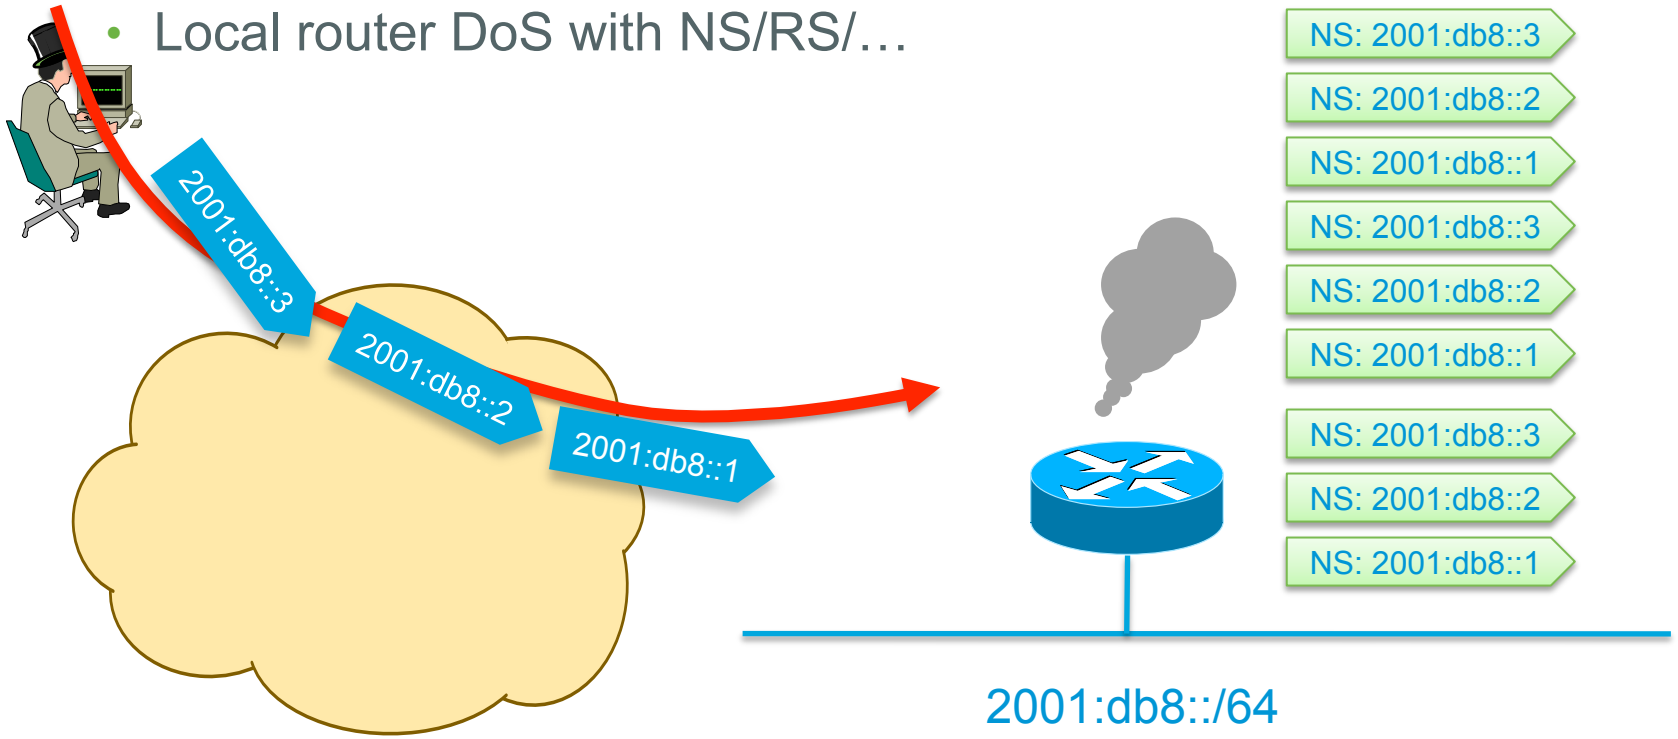

Mitigating Remote Neighbor Cache Exhaustion

- Mainly an implementation issue
  - Rate limiter on a global and per interface
  - Prioritize renewal (PROBE) rather than new resolution
  - Maximum Neighbor cache entries per interface and per MAC address
- **Internet edge/presence**: a target of choice
  - Ingress ACL permitting traffic to specific statically configured (virtual) IPv6 addresses only
  - ⇒Allocate and configure a /64 but uses addresses fitting in a /120 in order to have a simple ingress ACL

Simple Fix for Remote Neighbor Cache Exhaustion

- Ingress ACL allowing only valid destination and dropping the rest
- NDP cache & process are safe
- Requires DHCP or static configuration of hosts

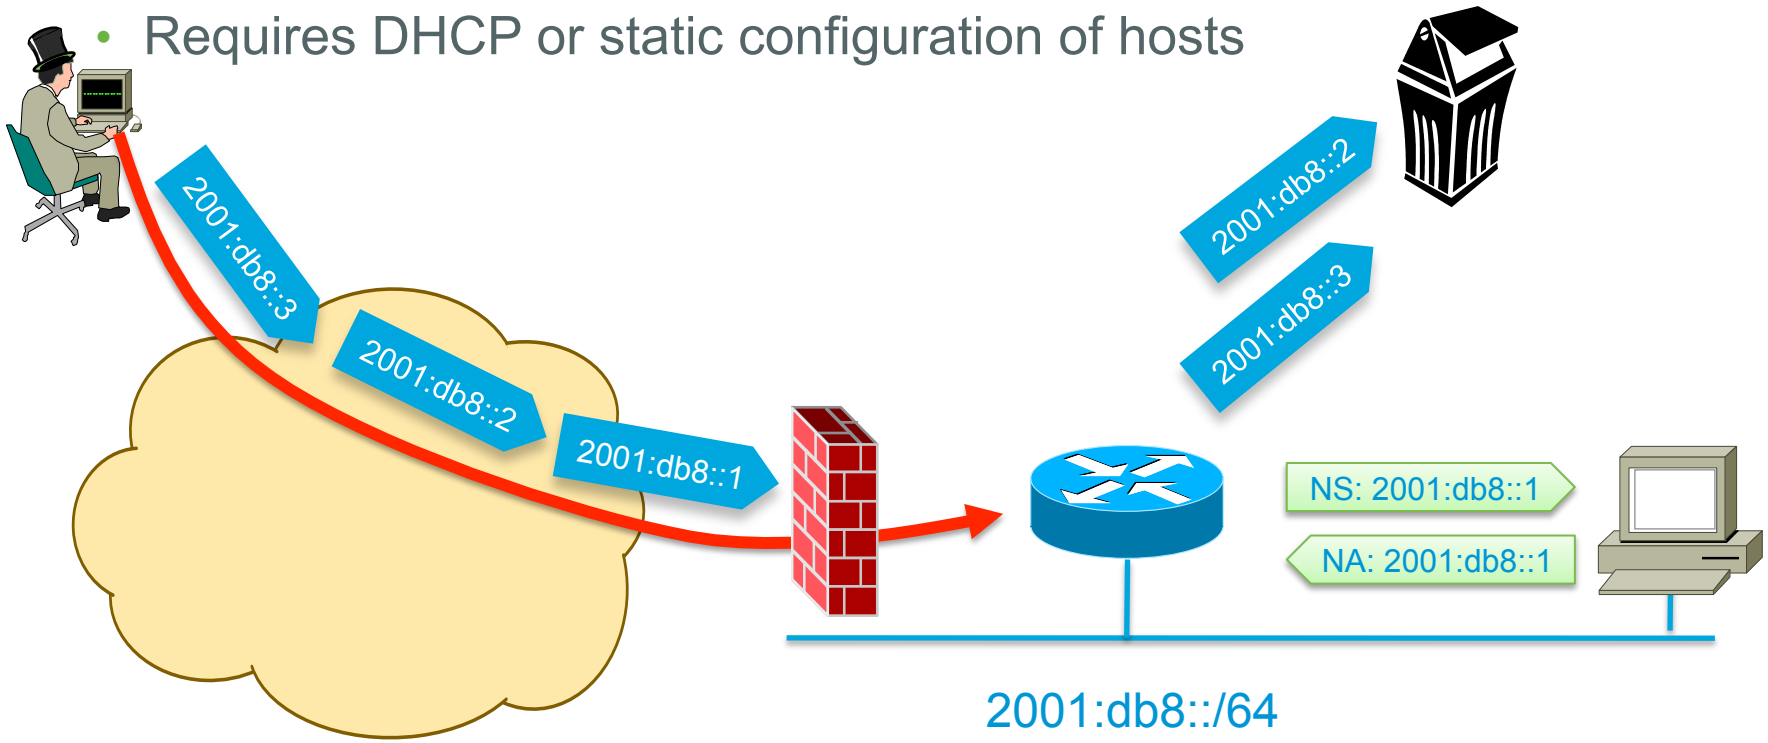

Addressing the Attendees-Exhaustion with Summary

Summary

- Without a secure layer-2, there is no upper layer security
- Rogue Router Advertisement is the most common threat
- Mitigation techniques
  - Host isolation
  - Secure Neighbor Discovery: but not a lot of implementations
  - SAVI-based techniques: discovery the ‘right’ information and dropping RA/NA with wrong information
  - Last remaining issue: (overlapped) fragments => drop all fragments…
- Neighbor cache exhaustion
  - Use good implementation
  - Expose only a small part of the addresses and block the rest via ACL
- Products are now available implementing the techniques ;-)

Any Question?

- And a shameless plug

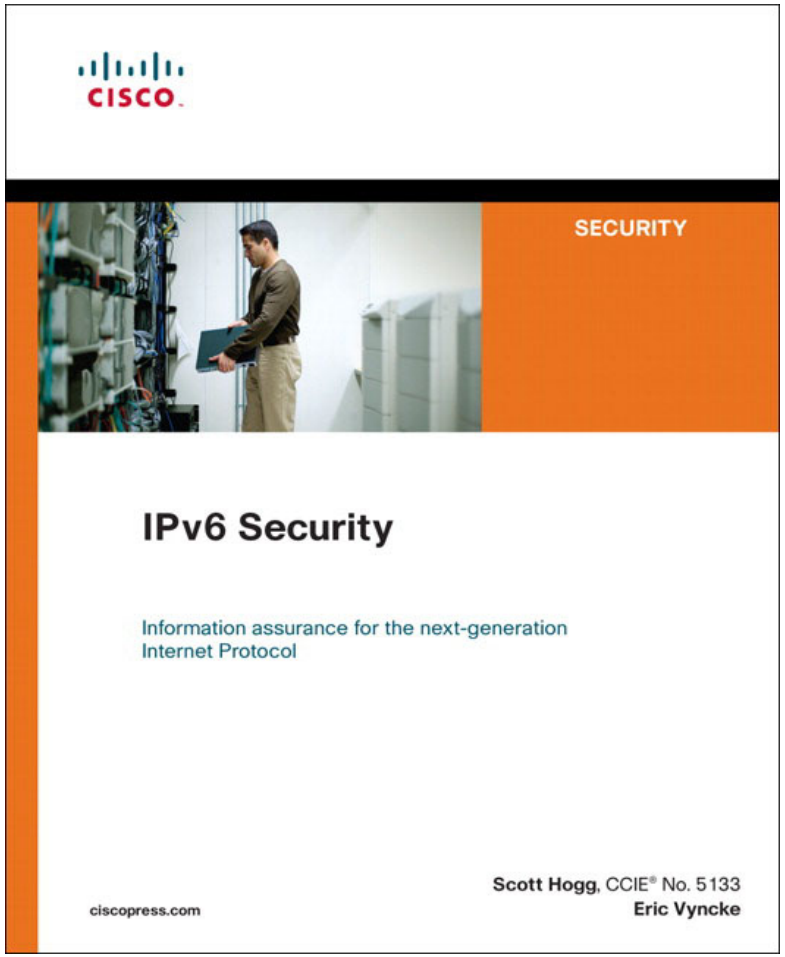

Thank you.