

IPv6 Security Hotspots

SWITCH

Frank Herberg, Head SWITCH-
CERT (Commercial Sectors)
frank.herberg@switch.ch

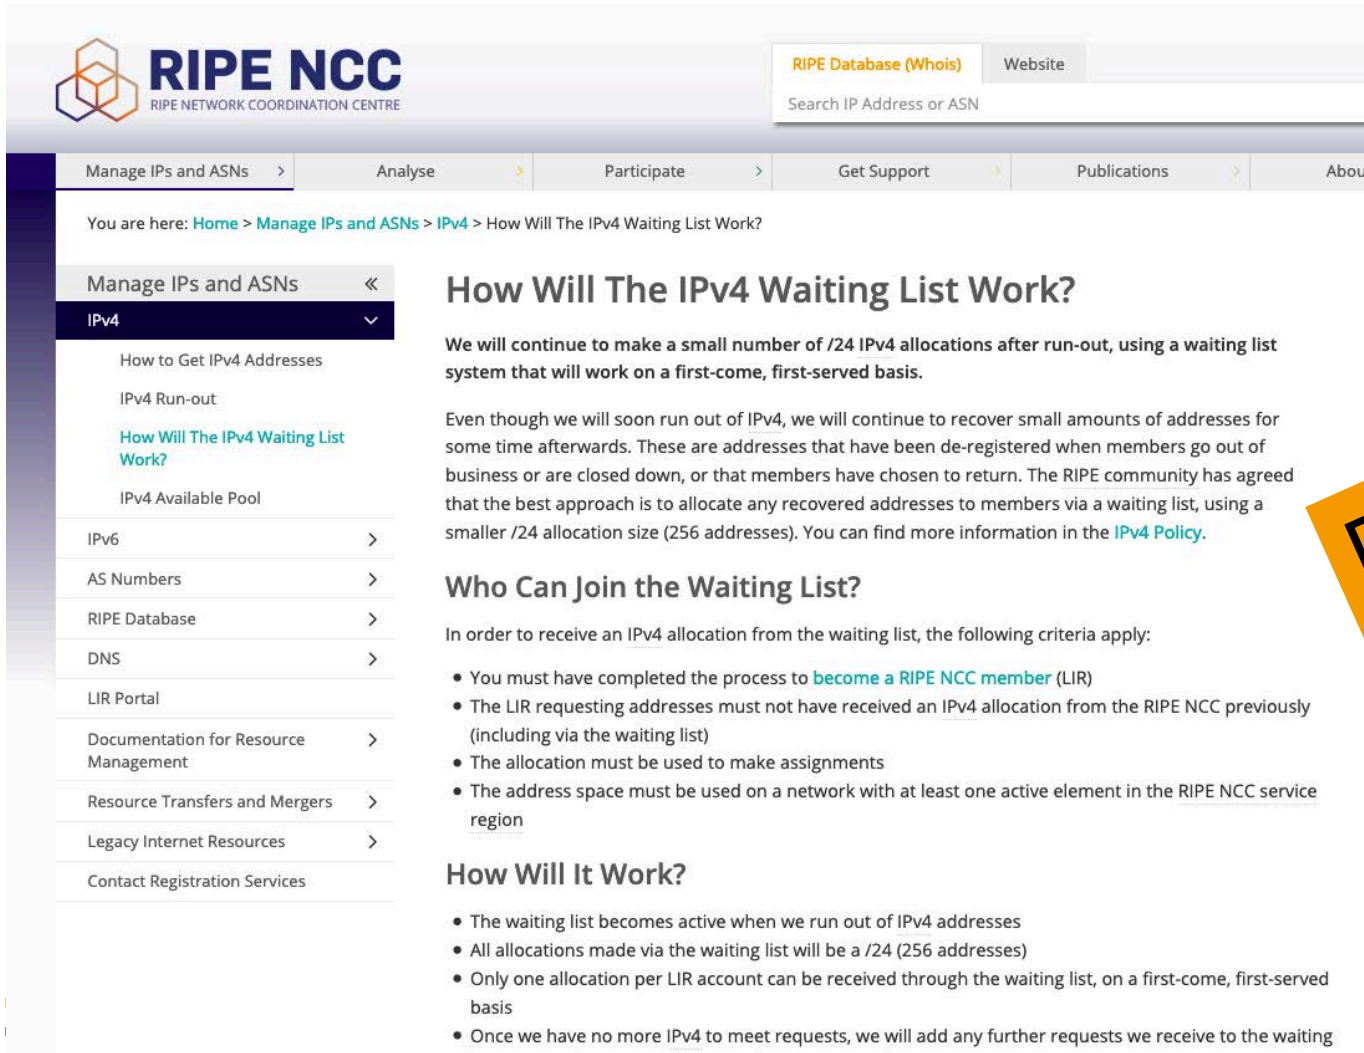
SwissIPv6Council, 11.11.2019

Agenda

- SWITCH-CERT
- IPv4 – aktuelle Situation
- **IPv6 Security Hotspots**
- 8-Punkte-Plan
- Etwa 300 weitere Slides
- Diskussion
- Ca. 19 Uhr → Aperero

My $::/32$ is
bigger than
your $/32$

Welcome to the IPv4 Address Waiting List...



The screenshot shows the RIPE NCC website interface. At the top left is the RIPE NCC logo (RIPE NETWORK COORDINATION CENTRE). To the right is a search bar with a dropdown menu showing 'RIPE Database (Whois)' and 'Website'. Below the search bar is a navigation menu with items: 'Manage IPs and ASNs', 'Analyse', 'Participate', 'Get Support', 'Publications', and 'About'. The main content area is titled 'How Will The IPv4 Waiting List Work?' and contains the following text:

You are here: [Home](#) > [Manage IPs and ASNs](#) > [IPv4](#) > How Will The IPv4 Waiting List Work?

How Will The IPv4 Waiting List Work?

We will continue to make a small number of /24 IPv4 allocations after run-out, using a waiting list system that will work on a first-come, first-served basis.

Even though we will soon run out of IPv4, we will continue to recover small amounts of addresses for some time afterwards. These are addresses that have been de-registered when members go out of business or are closed down, or that members have chosen to return. The RIPE community has agreed that the best approach is to allocate any recovered addresses to members via a waiting list, using a smaller /24 allocation size (256 addresses). You can find more information in the [IPv4 Policy](#).

Who Can Join the Waiting List?

In order to receive an IPv4 allocation from the waiting list, the following criteria apply:

- You must have completed the process to [become a RIPE NCC member](#) (LIR)
- The LIR requesting addresses must not have received an IPv4 allocation from the RIPE NCC previously (including via the waiting list)
- The allocation must be used to make assignments
- The address space must be used on a network with at least one active element in the [RIPE NCC service region](#)

How Will It Work?

- The waiting list becomes active when we run out of IPv4 addresses
- All allocations made via the waiting list will be a /24 (256 addresses)
- Only one allocation per LIR account can be received through the waiting list, on a first-come, first-served basis
- Once we have no more IPv4 to meet requests, we will add any further requests we receive to the waiting

The left sidebar contains a navigation menu with the following items:

- Manage IPs and ASNs <<
- IPv4 >>
 - How to Get IPv4 Addresses
 - IPv4 Run-out
 - [How Will The IPv4 Waiting List Work?](#)
 - IPv4 Available Pool
- IPv6 >
- AS Numbers >
- RIPE Database >
- DNS >
- LIR Portal
- Documentation for Resource Management >
- Resource Transfers and Mergers >
- Legacy Internet Resources >
- Contact Registration Services

Digital
Transformation
ahead?

...or buy them from former car dealers



Easy to Buy IPv4 Any Time

The team at IPv4 Connect has been in the data center and network technology space for over a decade. In that time, we have developed long-lasting relationships with industry-leading brands and key influencers. Our network of partners consists of data centers, network operators, hosting providers, ISPs, managed

IPv4 Brokers
Easily Buy and Sell IPv4 Addresses in Any Region Around the World

IPv4 – Buy, Sell and Lease with Ease in ARIN | RIPE | APNIC

IPv4.GLOBAL
Powered by Hilco Streambank

BROWSE AUCTIONS PRIOR SALES SALES PROCESS SELL IPv4

BUY NOW	BUY NOW	BUY NOW	BUY NOW
<p>/24 Block registered in ARIN Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$5,760.00 \$/ADDRESS: \$22.50</p> <p>ENDS IN: 4d 13h 11m</p>	<p>/24 Block registered in ARIN Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$5,888.00 \$/ADDRESS: \$23.00</p> <p>ENDS IN: 4d 13h 14m</p>	<p>/24 Block registered in ARIN Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$6,144.00 \$/ADDRESS: \$24.00</p> <p>ENDS IN: 4d 13h 17m</p>	<p>/23 Block registered in ARIN Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$11,000.00 \$/ADDRESS: \$24.00</p> <p>ENDS IN: 4d 13h 21m</p>
<p>/22 Block registered in ARIN Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$23,040.00 \$/ADDRESS: \$22.50</p> <p>ENDS IN: 4d 13h 31m</p>	<p>/19 Block registered in RIPE Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$163,840.00 \$/ADDRESS: \$20.00</p> <p>ENDS IN: 4d 14h 2m</p>	<p>/18 Block registered in ARIN Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$303,104.00 \$/ADDRESS: \$18.50</p> <p>ENDS IN: 4d 14h 2m</p>	<p>/21 Block registered in ARIN Transferable to: ARIN, APNIC, RIPE</p> <p>SALE PRICE: \$49,152.00 \$/ADDRESS: \$24.00</p> <p>ENDS IN: 4d 14h 2m</p>



IP Address Research, Recovery and Sales
100,000+ IPv4 IPs in inventory and ready for immediate transfer
sales-at-legacyresources.net
[\(877\) 822-8238](tel:(877)822-8238) or [\(707\) 520-4477](tel:(707)520-4477)
Contact Us Today To Sell or Buy IPv4 address blocks/subnets
Over 450,000 IPs sold in 2016



It is important that the possible security implications of IPv6 are well understood and considered during the design and deployment of IPv6 networks, rather than as an afterthought.

– Internet Society

Basic IT Security concept: “Complexity is the enemy of security”

- less clear / transparent
- bigger attack surface
- higher probability of (admin.) errors
- higher probability of bugs



More IP addresses to monitor and correlate

- Multiple IPv6 addresses per interface - plus the IPv4 address in a Dual Stack env.
- “Happy eyeballs” leads to unpredictable source address choice (RFC 6555,8305)
- Certain Mobile devices configure new IPv6 address each time they wake up
- IPv6 address notation isn't unique



The screenshot shows a webpage from NetworkWorld (FROM IDG) with a red header. The main content is an article titled "Using Dual Protocol for SIEMs Evasion" by Scott Hogg, dated February 24, 2013. The article discusses how attackers use IPv4 and IPv6 to evade detection by IPS, SIEMs, and reputation filtering. It mentions that attackers use a specific methodology involving malware propagation and command-and-control networks. A "RELATED" section on the right lists other articles about IPv6 deployment and government progress. A video thumbnail is also visible at the bottom right of the article content.

NETWORKWORLD
FROM IDG

INSIDER Sign I

CISCO SUBNET An independent Cisco community [View more](#)

Home > Cisco Subnet

CORE NETWORKING AND SECURITY
By Scott Hogg, Network World | FEB 24, 2013 12:51 PM PT

About | 
Scott Hogg is the CTO for Global Technology Res Inc. (GTRI). Scott provides network engineering, security consulting, and training services to his c

Using Dual Protocol for SIEMs Evasion

Attackers using IPv4 and IPv6 can avoid detection by IPS, SIEMs, reputation filtering, more



It is just a fact of life that attackers and defenders are now operating in a dual-protocol world. With the addition of IPv6, attackers are learning new tricks and defenders will need to anticipate and protect against those new attacks. Attackers will try to use IPv4 and IPv6, each alone or in combination, for their exploits. We can predict that attacks will use a combination of IPv4 and IPv6 in a way that could allow an attacker to avoid detection by today's protection mechanisms.

Attackers commonly use a specific methodology when using [malware propagation](#) and command-and-control networks for exploitation. However, attackers use a different standard methodology when performing a targeted attack. Attackers start with reconnaissance, exploring and scanning, exploitation, maintaining access, covering up tracks, and leveraging access to expand to other systems.

RELATED

IPv6 Is Not an All-or-Nothing Proposit

U.S. Government Progress on IPv6 Deployment

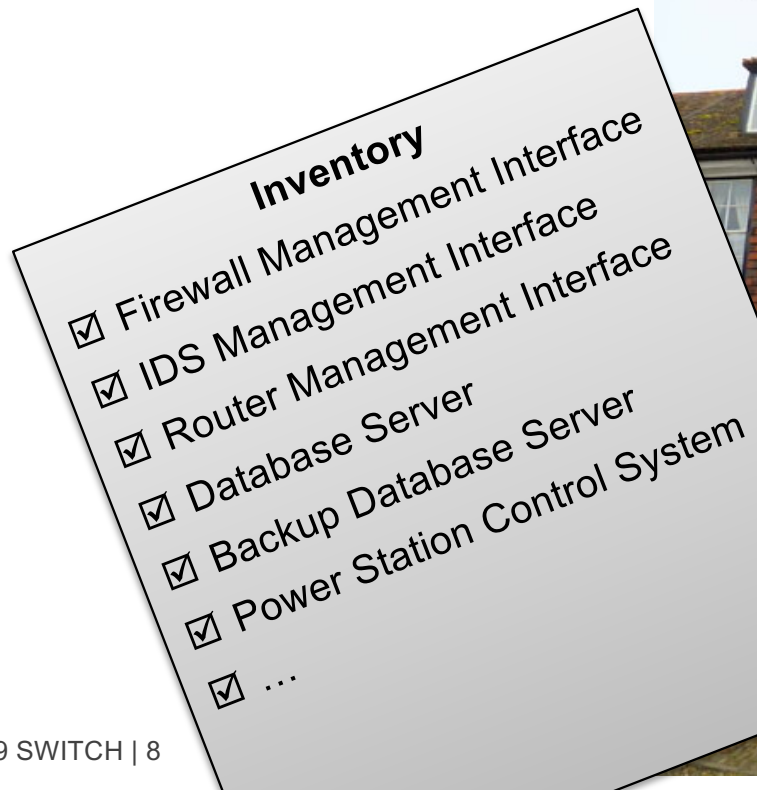
IPv6 deployment starts at the network

VIDEO
Four devices to get a b night's sleep

<http://www.networkworld.com/article/2224154/cisco-subnet/using-dual-protocol-for-siems-evasion.html>

Dual Stack and ACLs

- IPv4 based Access Control Lists (ACLs) only protect access via IPv4
- Enable IPv6? → Review all your ACLs! → Inventory??
- Maintain ACLs x2




IPv6 Blacklists

- IP reputation based Spam block lists for IPv6 are tricky:
 - difficult for vast IPv6 address space
 - Sender can utilize ‘nearly unlimited’ source addresses
 - Blacklisting of address ranges can lead to overblocking



Chained Extension Headers can be complex



- The number of EHs is **not limited**
- The number of options within an (Hop-by-Hop or Destination) Options Header is **not limited**
- There is **no defined order** of EHs (only a recommendation)
- EH have **different formats**

Possible Threats

- High Number of EHs / Manipulation of the EHs (fuzzing)
 - evade FW / IPS / RA-Guard
 - might crash or DOS the destination system
- Combine EH & Fragmentation to make it worse
 - by putting the attack into many small fragments
 - by combination of multiple extension headers and fragmentation so that layer 4 header is in 2nd fragment
- Use EH as Covert Channel

Mitigation Options: inspect EH / sanity checks / drop unknown

Homework: Resources on Extension Headers & Fragmentation Issues

- Excellent Paper:

Antonios Atlasis “Evasion of High-End IDPS Devices at the IPv6 Era”

<https://www.blackhat.com/docs/eu-14/materials/eu-14-Atlasis-Evasion-Of-High-End-IDPS-Devices-At-The-IPv6-Era-wp.pdf>

- Tool: chiron

<https://github.com/aatlasia/Chiron>

- RFC 6980, RFC 7112

ICMPv6 is much more complex than ICMP

Error-Messages (1-127)

1:Destination Unreachable 2:Packet too big (PMTUD)
3:Time Exceeded (Hop Limit) 4:Parameter Problem

Info-Messages (Ping)

128:Echo Request 129:Echo Reply

Multicast Listener Discovery (MLD, MLD2)

130:Multicast Listener Query 131/143:Multicast Listener Report/2
132:Multicast Listener Done

Neighbor Discovery (NDP), Stateless Autoconfiguration (SLAAC)

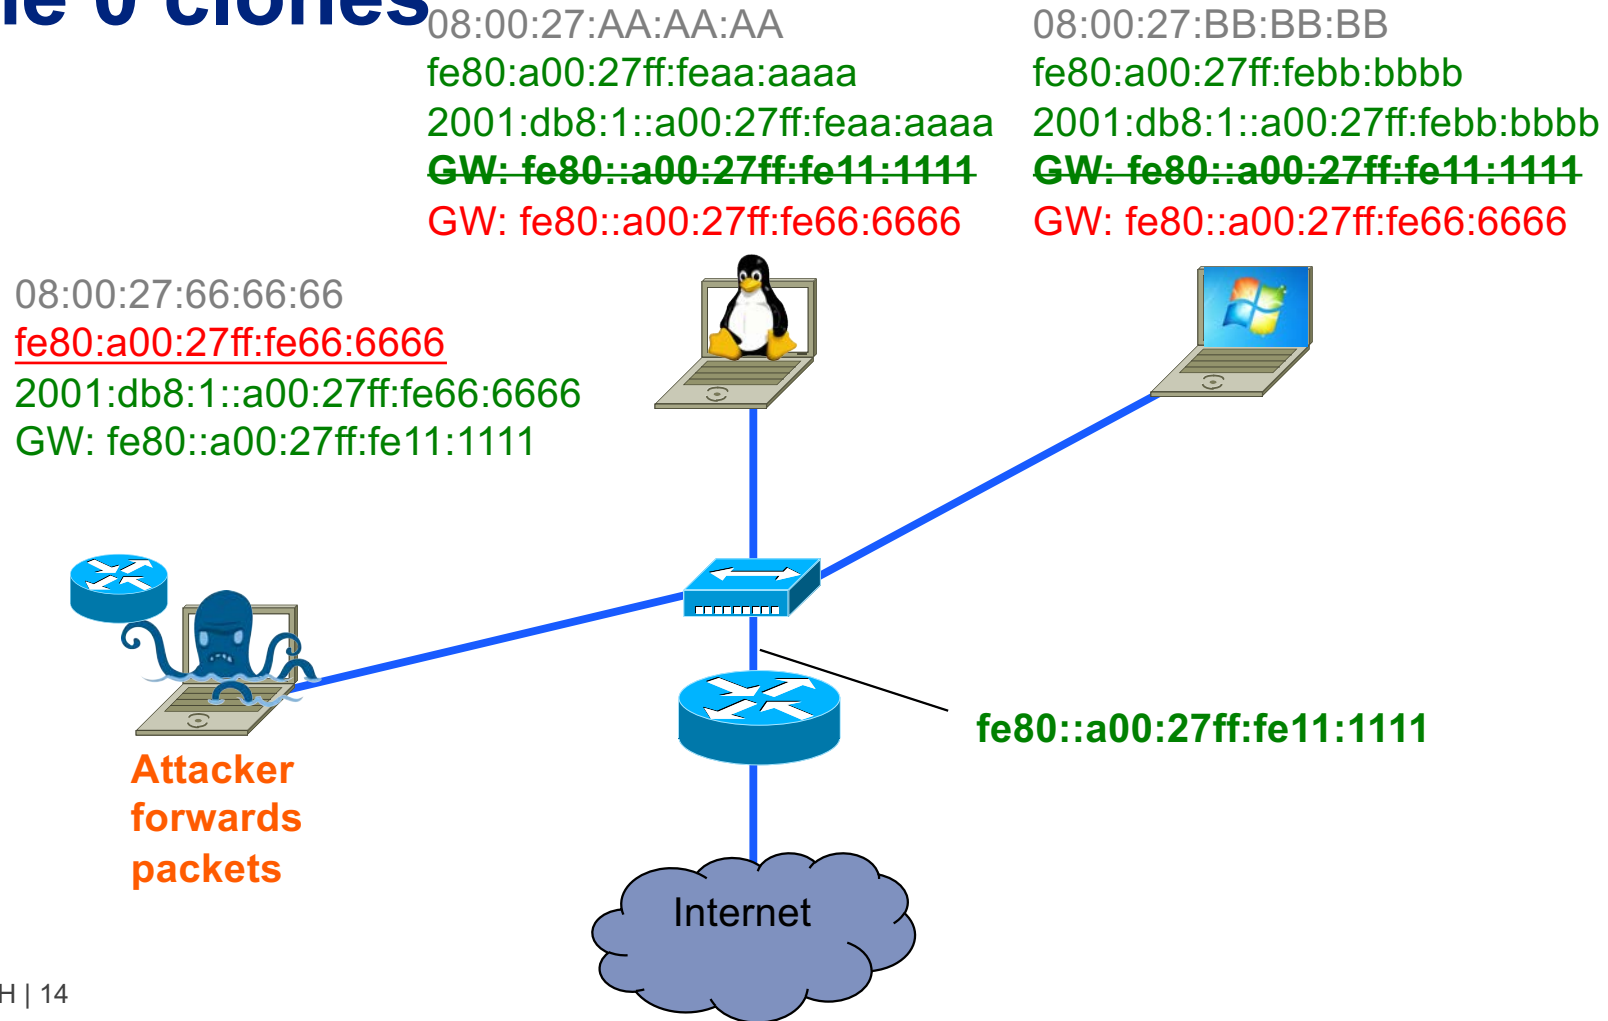
133:Router Solicitation 134:Router Advertisement
135:Neighbor Solicitation (DAD) 136:Neighbor Advertisement
(DAD) 137:Redirect Message

Other (Router Renumbering, Mobile IPv6, Inverse
NS/NA,...) 138-153

Filtering ICMPv6 is
more complex
**see RFC 4890 (38
pages)**

Several new attack
vectors (local,
remote)

Example: MITM-Attack with rogue RA plus lifetime 0 clones



Example: DOS-Attack with rogue RA flooding

ipconfig

taskmgr: CPU load

The image shows a Windows desktop during a DOS attack. On the left, a command prompt window displays the output of the 'ipconfig' command, listing numerous IPv6 addresses. On the right, the Windows Task Manager is open, showing that the CPU usage is at 100%. A blue arrow points from the 'ipconfig' text to the command prompt window, and another blue arrow points from the 'taskmgr: CPU load' text to the Task Manager window.

RFC 6104: Different Mitigation Approaches

- Disable RA processing (but it's needed for DHCPv6)
- Filter on Switch: RA-Guard (can be bypassed using EH)
- Host based filters configured to accept RAs only from valid Router addresses (works only in managed environment)
- Deprecation Daemon: Detect incorrect RAs and then in turn send a deprecating RA with a router lifetime of zero (not for flooding)
- Partitioning, Microsegmentation or Host Isolation
- DHCPv6-only? No: RA informs about use of DHCPv6

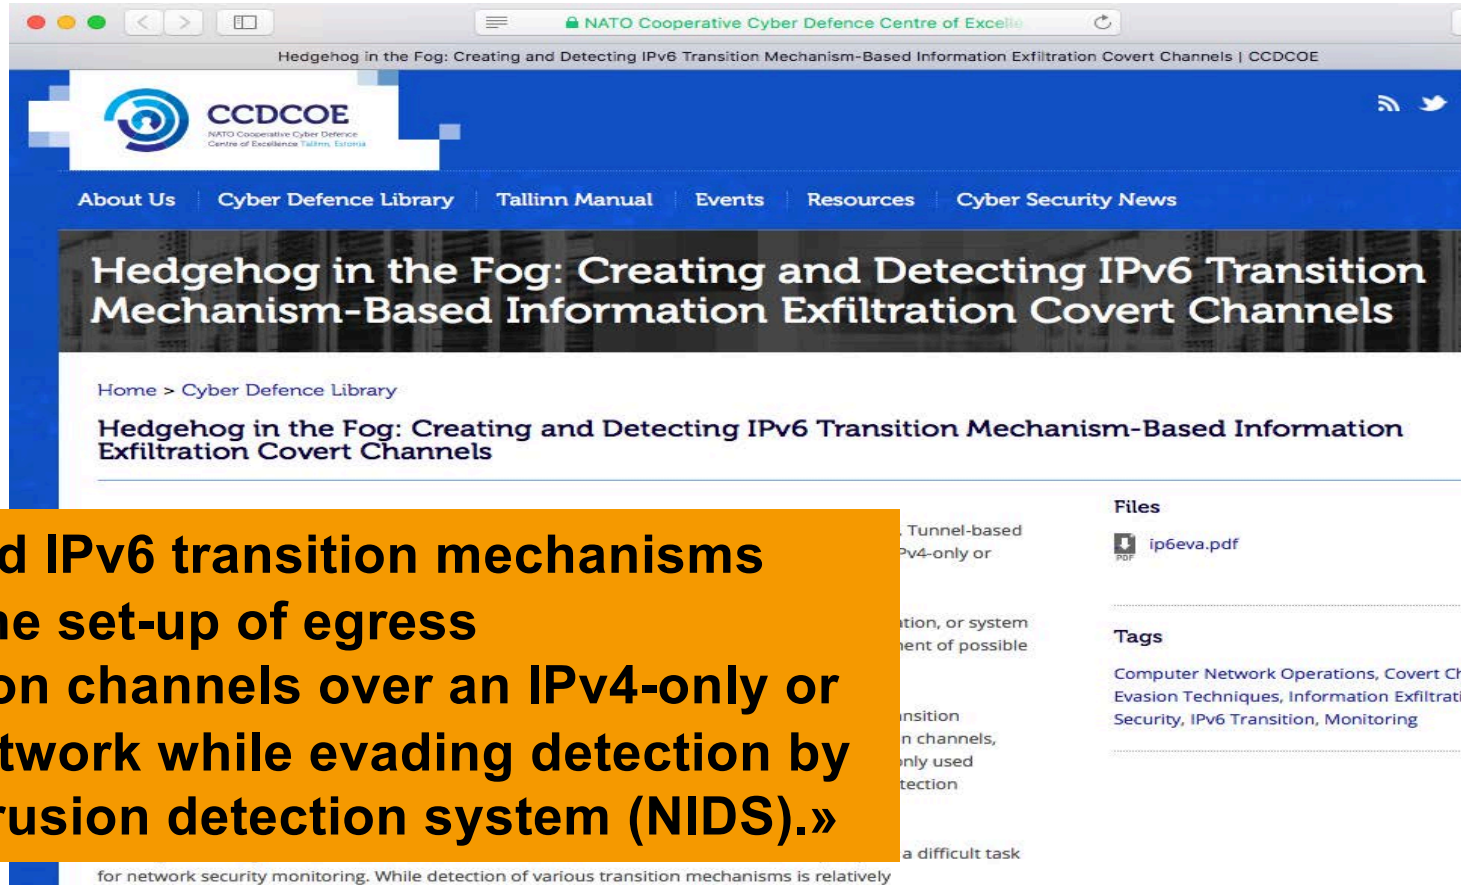
Recommended Tools:

- <https://github.com/vanhauser-thc/thc-ipv6>
- <https://www.si6networks.com/tools/ipv6toolkit/>

Homework for you: which Rogue RA Mitigation measures are wise?

Zone	Rogue RA Mitigation Measure	cost (+ o -)	feasibility	effect (+ o -)
Internal Network	Router-Preference=high / Monitor NDP Managed Switch (RAGuard, PACLs)	+/-	+	0/+
Internal Server-Zone	Router-Preference=high / Monitor NDP Disable RA processing	+	+	+
DMZ	Router-Preference=high / Monitor NDP Disable RA processing	+	+	+
Guestnet Wired	Router-Preference=high Managed Switch with RA Guard or Port ACLs	-	+	+
Guestnet Wireless	Router-Preference=high Partitioning	+/o	+	+

Data Exfiltration using IPv6 tunnels



The screenshot shows a web browser window displaying a page from the NATO Cooperative Cyber Defence Centre of Excellence (CCDCOE). The page title is "Hedgehog in the Fog: Creating and Detecting IPv6 Transition Mechanism-Based Information Exfiltration Covert Channels". The page content includes a navigation menu with links for "About Us", "Cyber Defence Library", "Tallinn Manual", "Events", "Resources", and "Cyber Security News". Below the navigation menu, there is a large heading for the article. A yellow callout box is overlaid on the page, containing text about tunnel-based IPv6 transition mechanisms. To the right of the main content, there is a "Files" section with a PDF file named "ip6eva.pdf" and a "Tags" section with various keywords.

«Tunnel-based IPv6 transition mechanisms could allow the set-up of egress communication channels over an IPv4-only or dual-stack network while evading detection by a network intrusion detection system (NIDS).»

for network security monitoring. While detection of various transition mechanisms is relatively

NATO Whitepaper on data exfiltration over IPv6 transition mechanisms
<https://ccdcoe.org/multimedia/hedgehog-fog-creating-and-detecting-ipv6-transition-mechanism-based-information.html>

Detect IPv6 tunnels in network logs

Look inside logs / NetFlow records:

- IPv4 Protocol type 41 (ISATAP, 6to4 traffic)
- IPv4 to UDP 3544 (Teredo traffic)
- Traffic to 192.88.99.1 (6to4 anycast server)
- DNS server log: resolution of "ISATAP"

➔ Better: deploy native IPv6 to avoid tunnels

low data rate exfiltration using only Layer 3 protocol headers (/64 IID dest address 64 bits)



Source <https://youtu.be/WWTtl8ebfg8>

Using flow label field as covert channel

README.md

IPv6teal

IPv6teal is a Python 3 tool to stealthily exfiltrate data from an internal network using a [covert channel](#) built on top of the IPv6 header `Flow label` field.

It is made of 2 components:

- [exfiltrate.py](#): Client-side component, used to exfiltrate data from an internal machine
- [receive.py](#): Server-side component, used to received the exfiltrated data

Jump to: [Background](#) | [Usage](#) | [F.A.Q.](#)

Background

IPv6 packets have a [header](#) containing a 20-bit field, `Flow label`.

Offsets	Octet	0				1				2				3																			
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Version				Traffic Class				Flow Label																							
4	32	Payload Length								Next Header				Hop Limit																			
8	64																																

Reconnaissance / Scanning the network / Cyber Kill Chain

- Sequentially scanning IPv6 address space is not feasible anymore
- DNS bruteforcing: common hostnames
 - with 1900 words get 90% of systems in DNS
- Alive bruteforcing: typical addresses
 - with 2000 addresses get 66% of the systems
- Combined (and use of brain):
 - ca. 90-95% of servers are found

➔ Target Discovery is still possible

Shodan: Participate in pool.ntp.org as IPv6 endpoints; if NTP clients connect for time sync => scan them

[Pool] shodan.io actively infiltrating ntp.org IPv6 pools for scanning purposes

Luca BRUNO [lucab at debian.org](mailto:lucab@debian.org)

Wed Jan 27 11:24:06 UTC 2016

- Previous message (by thread): [\[Pool\] Question about score for 89.101.218.6](#)
- Next message (by thread): [\[Pool\] shodan.io actively infiltrating ntp.org IPv6 pools for scanning purposes](#)
- Messages sorted by: [\[date \]](#) [\[thread \]](#) [\[subject \]](#) [\[author \]](#)

[cross-posted to pool-ntp and oss-sec]

Hi,
while reviewing network logs this morning I spotted some anomalies related to scan probes, ntp.org pools and IPv6.

It looks like Brad already observed and blogged about this some days ago, but I haven't seen this discussed in the usual ntp-pools, Debian and oss-sec ML, so I'm reposting this here:
<http://netpatterns.blogspot.de/2016/01/the-rising-sophistication-of-network.html>

In summary, some machines (which seem related to the shodan.io scanning project) are actively participating in pool.ntp.org as IPv6 endpoints. However, clients connecting to them for NTP timesync, are subsequently scanned by probes originating from *.scan6.shodan.io hosts.

Confirming original report from Brad, I can add that those scanners seem to implement some kind of rate-limiting: they will timeout NTP and won't re-scan recent clients when doing multiple/subsequent NTP requests. Moreover, this is not targeted/restricted to the Debian pool only, but plague the whole IPv6 pool, as seen on a sample query to the RedHat pool:

```

~ ~ ~
$ dig +short -t AAAA 2.rhel.pool.ntp.org | grep -E ':[[:xdigit:]]00[[:xdigit:]]$'
2a03:b0c0:3:d0::18:b001
$ dig +short -x 2a03:b0c0:3:d0::18:b001
analog.data.shodan.io.
~ ~ ~

```



Acht-Punkte-Plan

1. ACLs für IPv6 überprüfen / ebenso Blacklists (im IPv4-only-Netz IPv6 deny)
2. ICMPv6-Angriffe verstehen und geeignete Massnahmen ergreifen (selber ausprobieren mit den bestehenden Angriffstools)
3. ICMPv6 nicht komplett am Router blocken (und existierende Filter regelmässig reviewen, siehe RFC 4890)
4. Security-Monitoring/IDS/IPS/SIEM für Dual-Stack überdenken (feature parity, Nachvollziehbarkeit bei IPv6-Adressen / Korrelation von Multiprotokollangriffen)
5. Data exfiltration über IPv6 entdecken (vor allem Tunnel)
6. Extension Header Angriffe verstehen und Gegenmassnahmen prüfen (und regelmässig reviewen)
7. Security-Tools IPv6-fähig machen
8. Wissenslücken schliessen

SWITCH IPv6 Security Training now freely available at first.org

<https://www.first.org/education/trainings>

Available Trainings

- FIRST CSIRT Basic Course
- FIRST Threat intel Pipelines Course
- DDoS Mitigation Fundamentals
- Mastering CVSSv3
- PSIRT Training
- Incident Handling for Policy makers
- Conducting Exercises to improve Incident Response
- IPv6 Security
- Third party training material

NEW

Attribution-NonCommercial-ShareAlike
4.0 International (CC BY-NC-SA 4.0)

SWITCH

Working for a better digital world

@frankherberg
frank.herberg@switch.ch

