Arjan Koopen

BCP38 violation hunting as an IX member

NLNOG Day 2023

BCP38 recap (1)

What is BCP38?

BCP38 = Anti IP-spoofing

http://www.bcp38.info/ : “BCP38 is **RFC2827**: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing.”

Part of MANRS: https://www.manrs.org/netops/guide/antispoofing/

CAIDA spoofer project: https://www.caida.org/projects/spoofer/

BCP38 recap (2)

Why is IP Source Address Spoofing a problem? What can it be abused for?

- Trigger large DDoS attacks (for example DNS/NTP/STUN/etc amplification attacks) → “reflection”
- Trigger UDP flooding/replay attacks with “untraceable” random source IP’s → “hiding”

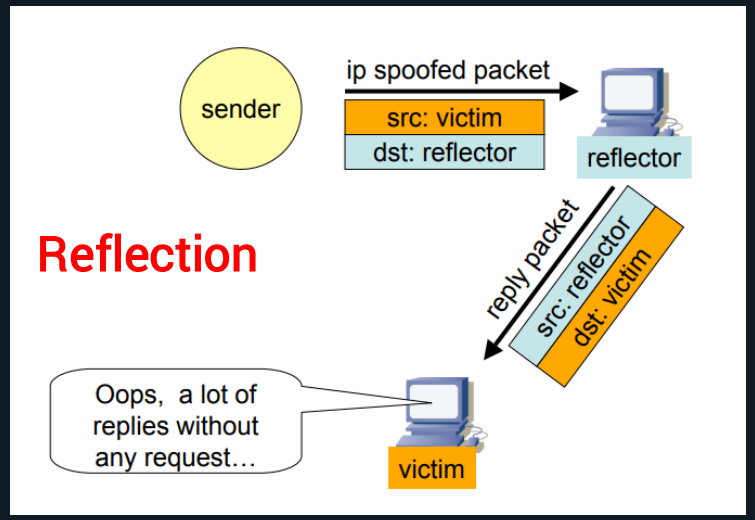

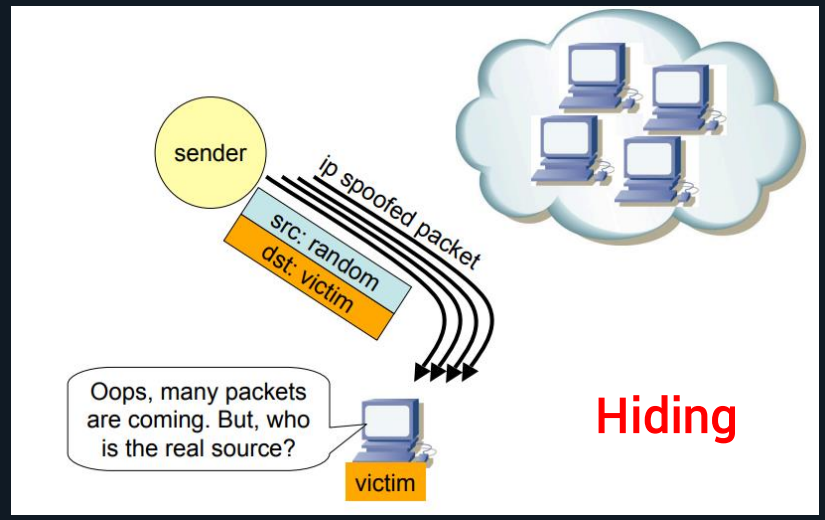

Where should you apply BCP38?

Ideally on interfaces where “endpoints” are terminated at Layer-3. Add rpf-check/uRPF or static ACL here.

AS49544 IP spoofed traffic received

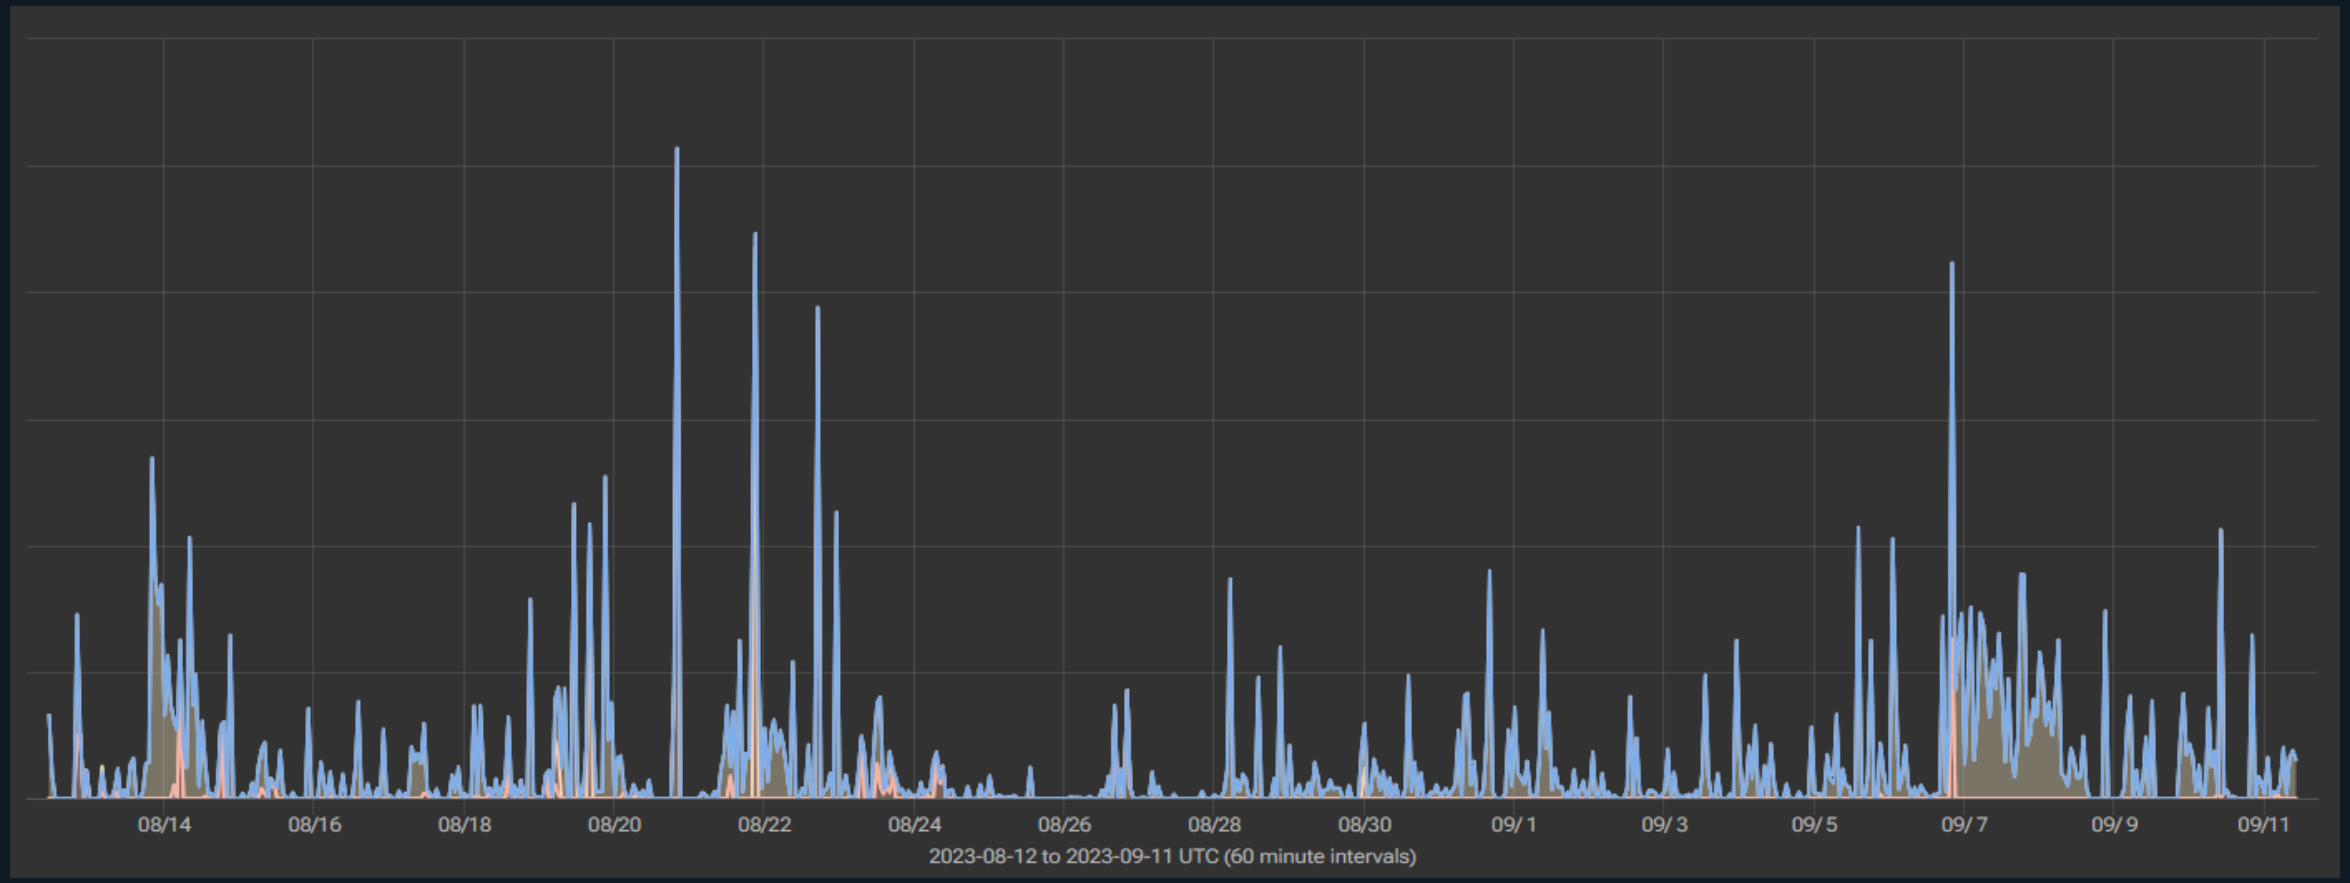

Ongoing stream of spoofed traffic (spoofing AS49544 IP’s); about 83% via transit; 17% IX/PNI; triggering amplification attacks (“reflection”)

UDP flooding traffic with BOGON spoofed IP's

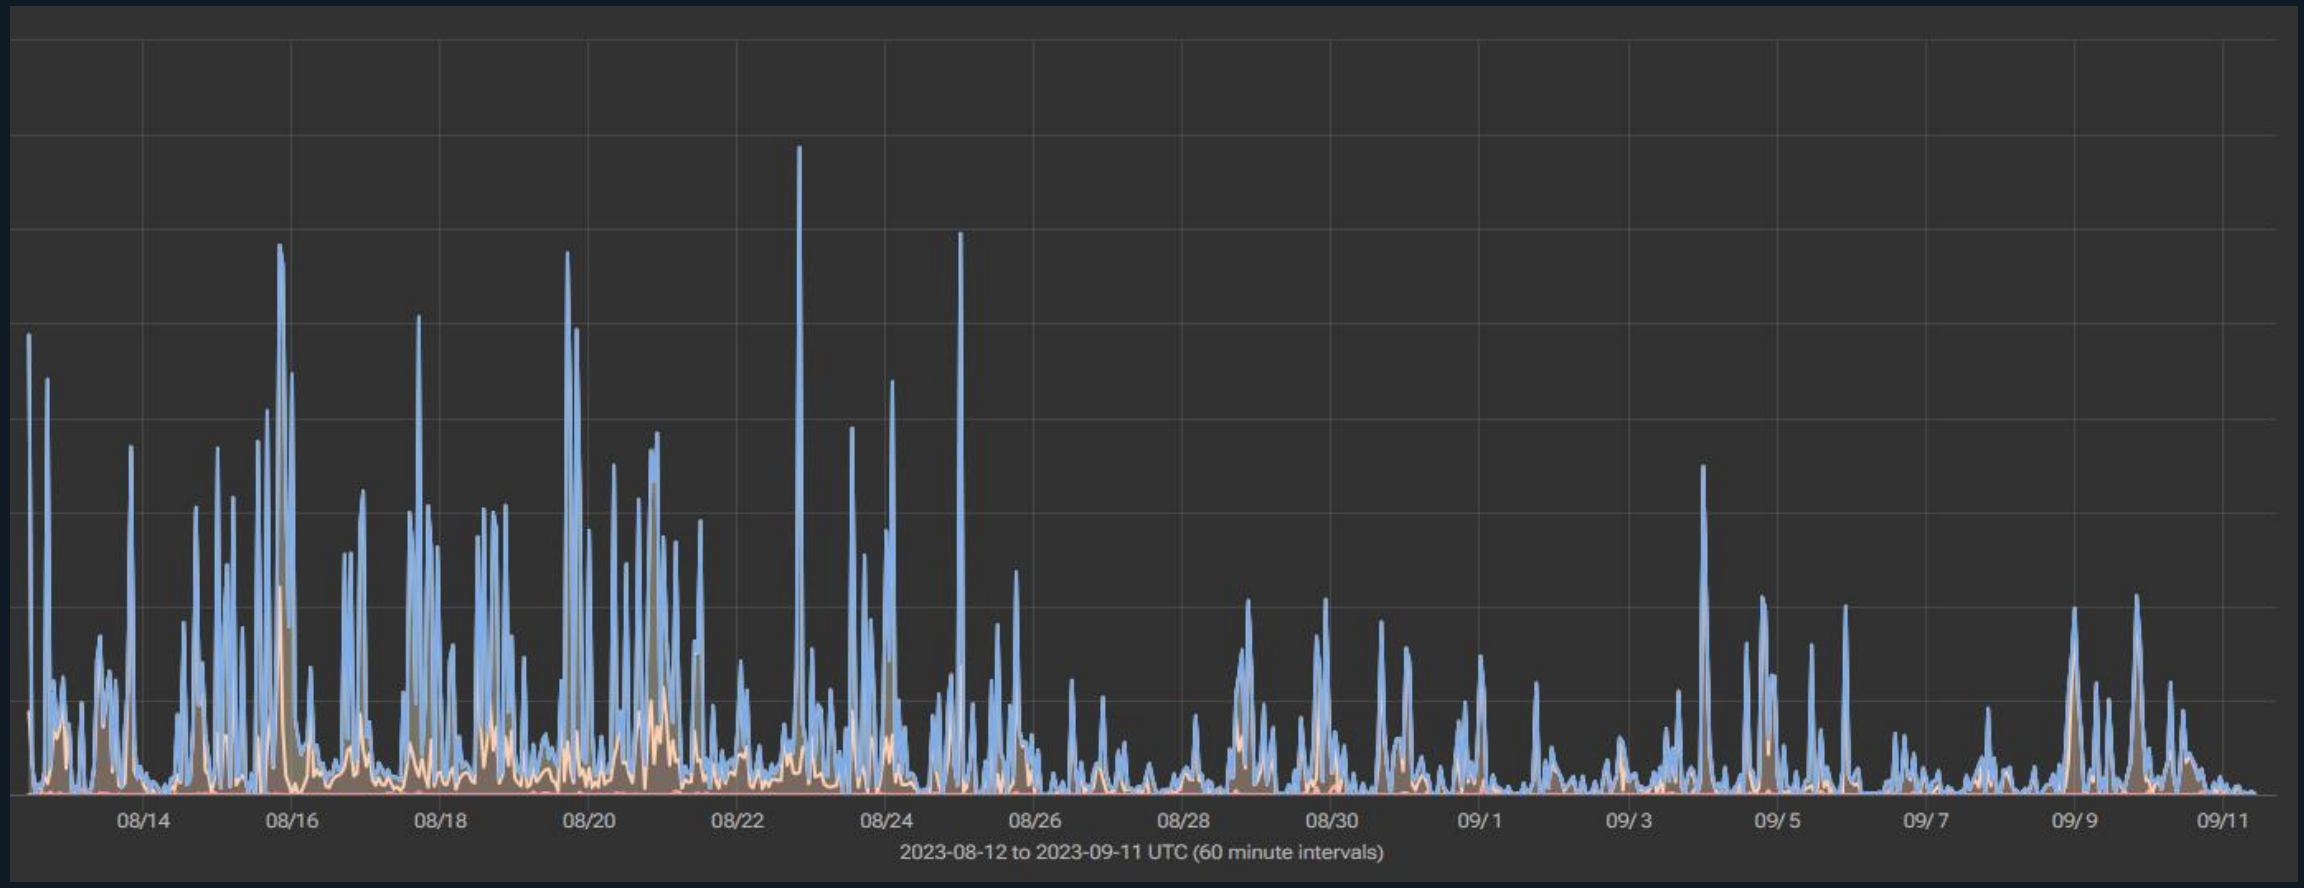

Ongoing stream of spoofed traffic ("hiding"); spoofing BOGON ranges; about 55% via IX/PNI

Most of these attacks spoofing the whole IPv4 internet

Traffic sourced from IP's *not* in visible DFZ, no BOGONs

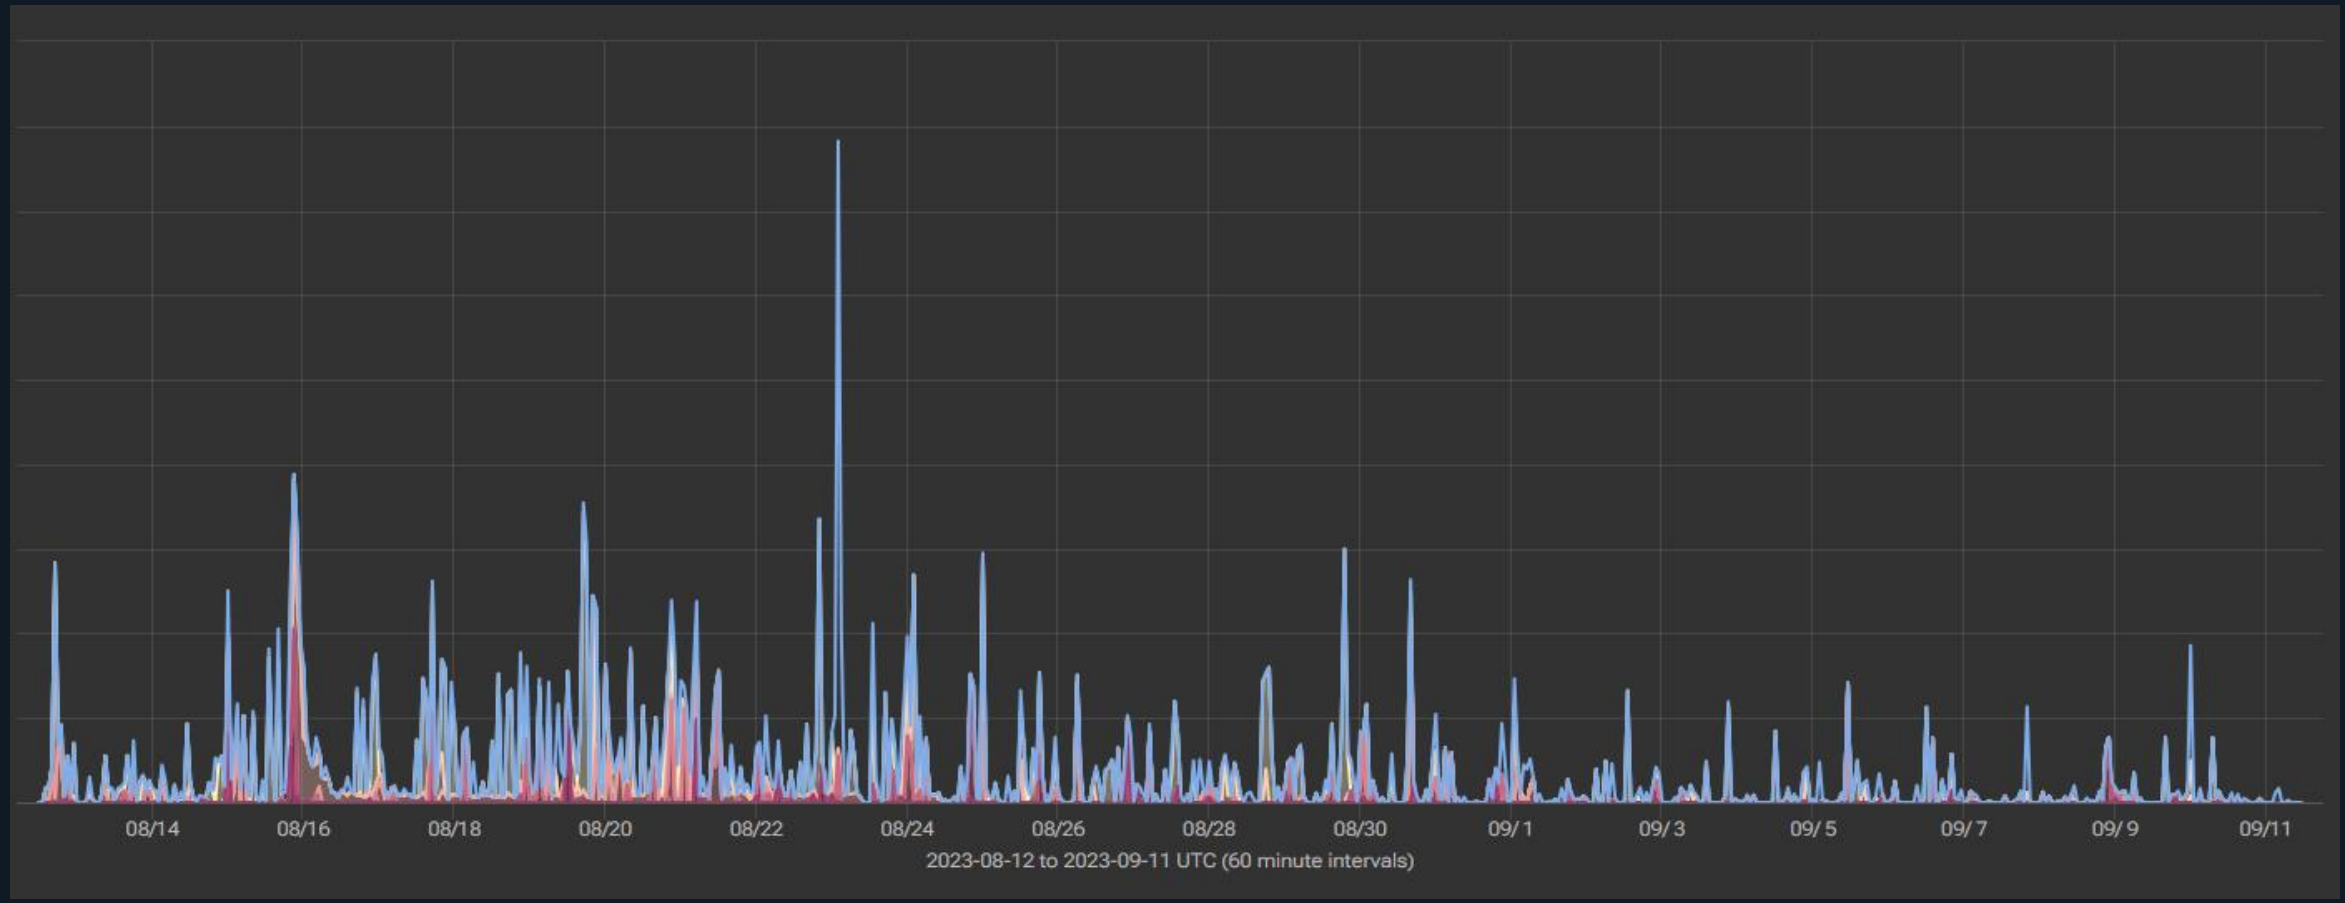

Hiding attacks

Destination ASN = 49544, Source ASN = 0, no BOGONs

“dweilen met de kraan open”

How to detect BCP38 violations in your network?

Flow monitoring (for example IPFIX)

- Look for flows on external facing interfaces with source addresses in BOGON ranges and/or your own prefixes
- Works for well for private-peering and transit interfaces (assuming your transit partner is willing to help you with your investigations)
- Works poorly on Internet Exchange interfaces when the Layer-2 information is missing in the flow data (source MAC address). Of course there are many ISP's on the same LAN, need Layer-2 information to figure out who is sending the spoofed traffic.

Firewall filter/ACL + syslog

- Not great, not terrible
- Works great on Internet Exchange interfaces, assuming your router will send Layer-2 information along and assuming the TCAM/performance penalty is not too high
- Filter on BOGON ranges and your own prefixes

Example BOGON IPv4: 0.0.0.0/8, 127.0.0.0/8, 192.0.2.0/24, 198.18.0.0/15, 198.51.100.0/24, 203.0.113.0/24, 224.0.0.0/4, 240.0.0.0/4. Beter to exclude RFC1918/6598/APIPA due to NAT misconfig/malfunctions.

Syslog example Juniper

Juniper firewall syslog example

Applied on `family inet` interface:

Meaning both Layer-2/3/4 information is included!

Grok pattern

```
%{DATA:device} %{DATA:fpc} PFE_FW_SYSLOG_ETH_IP: FW: %{DATA:interface}
%{DATA:filter_action} %{DATA:eth_proto} %{COMMONMAC:src_mac_address} ->
%{COMMONMAC:dst_mac_address}%{SPACE}%{DATA:ip_proto} %{IP:src_ip_address}
%{IP:dst_ip_address}%{SPACE}%{NUMBER:src_port}%{SPACE}%{NUMBER:dst_post}
```

Connecting the dots

1. Router sends violation syslog data to Graylog.
2. Graylog uses “extractor” using Grok pattern and will aggregate syslog data based on source MAC address, device and interface. Script will pull list of violating MAC addresses.
3. Figure out which Peering LAN IP address is bound to the violating MAC address (using LibreNMS API).
4. (optional) Who is being targeted in our network? Look up tenant/customer using Netbox API.
5. Using the Peering LAN IP, look up which ASN using this IP using the PeeringDB API. Also lookup contact details.
6. Send e-mail using gathered information and save to database for historic reasons.

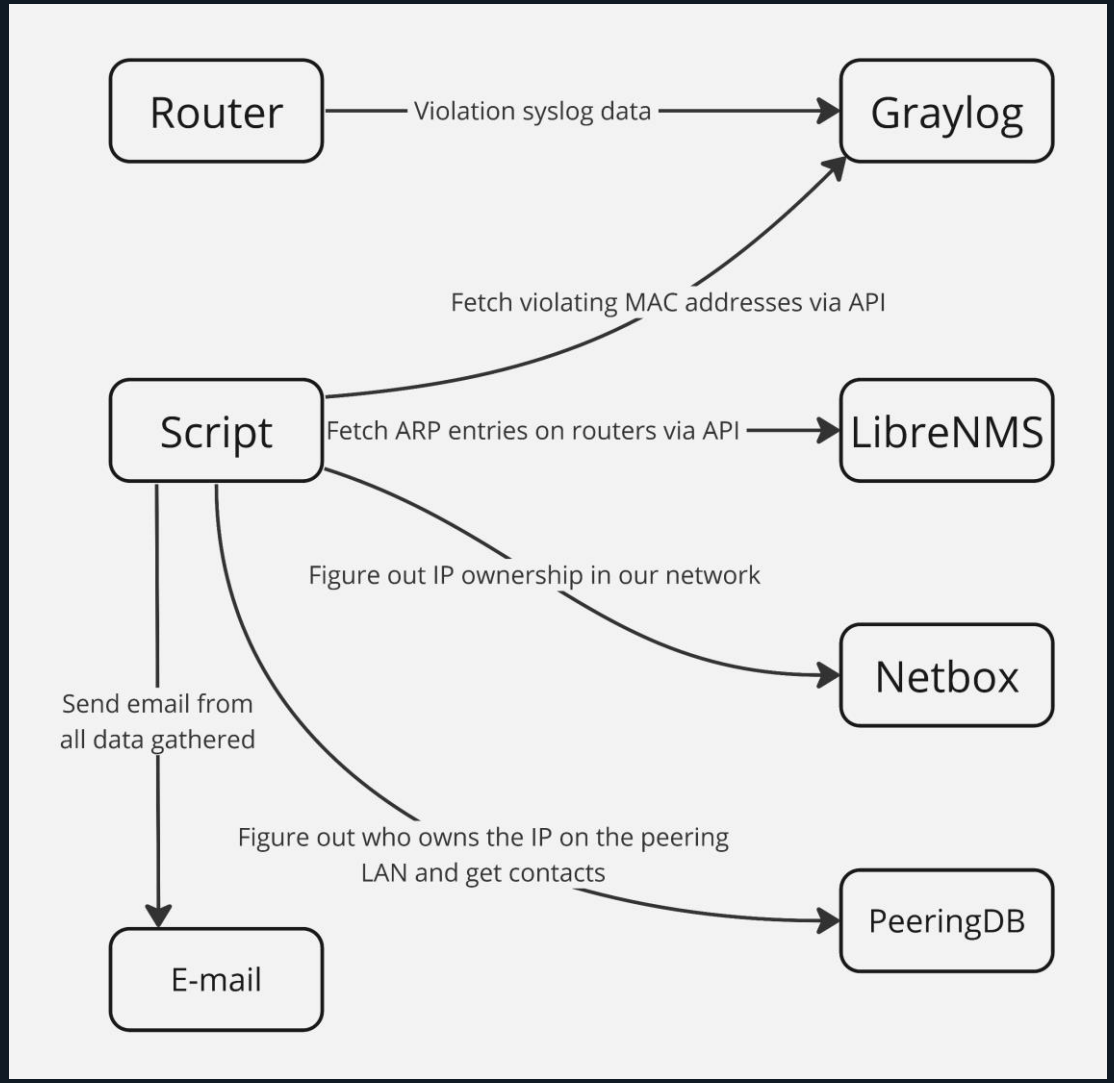

Testing from AS64404/EventInfra – Craft packets with scapy

```
from scapy.all import *

A = '255.255.13.37' # spoofed source IP address
B = '5.200.0.0' # destination IP address
C = 1337 # source port
D = 1337 # destination port
payload = "yada yada yada" # packet payload

spoofed_packet = IP(src=A, dst=B) / UDP(sport=C, dport=D) / payload
send(spoofed_packet)
```

```
e4:1d:2d:2f:6e:f1 > 3c:8c:93:54:79:19, ethertype IPv4 (0x0800), length 56:
255.255.13.37.1337 > 5.200.0.0.1337: UDP, length 14
```

Detection/response on AS49544 side

64404 [Mail][Edit][Delete]	EventInfra	(demo entry)		

IX name	IX IPv4	IX MAC	Latest timestamp	Latest src IPv4	Latest dst IPv4	Hits
Frys-IX [logs]	185.1.203.141	e4:1d:2d:2f:6e:f1	2023-09-16T12:58:20.000Z	255.255.13.37	5.200.0.0	1K

```
Hello EventInfra / AS64404,

We are i3D.net / AS49544.

We have received IP-spoofed traffic from your network via the Internet Exchange(s) mentioned below. See below for details and example logging. The source IP's are either part
of BOGON IP ranges or are part of prefixes NOT to be originated by you.

===

IX name: Frys-IX
Your IP on the IX: 185.1.203.141
Your router's MAC address on the IX: e4:1d:2d:2f:6e:f1
Latest target IP in our network: 5.200.0.0

Latest logs, timestamps in UTC, format includes [source MAC address (you)] --> [destination MAC address (us)] [proto] [source IP address (spoofed address coming from you)]
[destination IP address (in our network)] [source port] [destination port]:

2023-09-16T12:58:20.000Z nlams1-rt001i-2 fpc0 PFE_FW_SYSLOG_ETH_IP: FW: ae14.0 A 0800 e4:1d:2d:2f:6e:f1 -> 3c:8c:93:54:79:19 udp 255.255.13.37 5.200.0.0 1337 1337 (1 packets)
[...]
2023-09-16T12:55:41.000Z nlams1-rt001i-2 fpc0 PFE_FW_SYSLOG_ETH_IP: FW: ae14.0 A 0800 e4:1d:2d:2f:6e:f1 -> 3c:8c:93:54:79:19 udp 255.255.13.37 5.200.0.0 1337 1337 (2 packets)

===

This IP-spoofing can be abused for triggering DDoS attacks.

We are wondering if EventInfra does implement BCP38? For resources see:

- http://www.bcp38.info/
- https://www.manrs.org/netops/guide/antispoofing/
- https://tools.ietf.org/html/bcp38

If a network implements BCP38, this spoofed traffic would never leave that network.

Could you please look into this issue? Thank you.
```

Experiences gathered

- Response rate on emails is not great.
- But when they respond, most ISP's we contact are willing to help. Usually there are misconfigs, router bugs, etc
- In a number of cases, a downstream BGP customer of the offending ISP is generating spoofed traffic; which means that customer should implement BCP38.
- In some rare cases (of no response/action) we can escalate via the relevant Internet Exchange. But not all Internet Exchanges have a good process for dealing with these issues.
- Large chunk of IP spoofed traffic is received via transit, but not all transit providers are willing to help.

Some spoofing incident stats from the past 2 months in the i3D.net/AS49544 network (only on Internet Exchanges!):

- 126 different ASN's with BCP38 violations detected.
- 203 unique MAC addresses with BCP38 violations detected.

What’s next?

- Get more operators to take a similar approach: detect spoofing incidents coming into your own network and contact offending parties
- Work closer with transit providers to find offending parties
- Talk to Internet Exchanges about improving incident handling
- Convert syslog approach towards IPFIX/flow-data (new router’s Network OS feature?)
- Better spoofing detection for our IP-transit customers (using their dynamic route prefix-list, etc)

Thanks!
Questions?