2bc8a3431ebf781663e6b867720c64c0

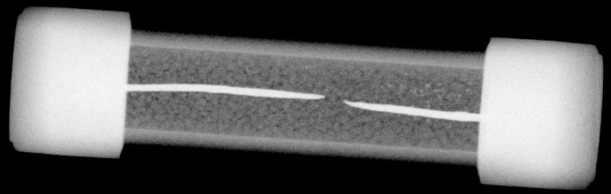

BGP Error "Handling"

Ben Cartwright-Cox
NANOG 89 (2023)

Pinky Swear

I run a business that involves being peered to many IXP route servers and other people's routers. **I have not and will not ever test for BGP bugs/exploits on customer/partner sessions without explicit consent**.

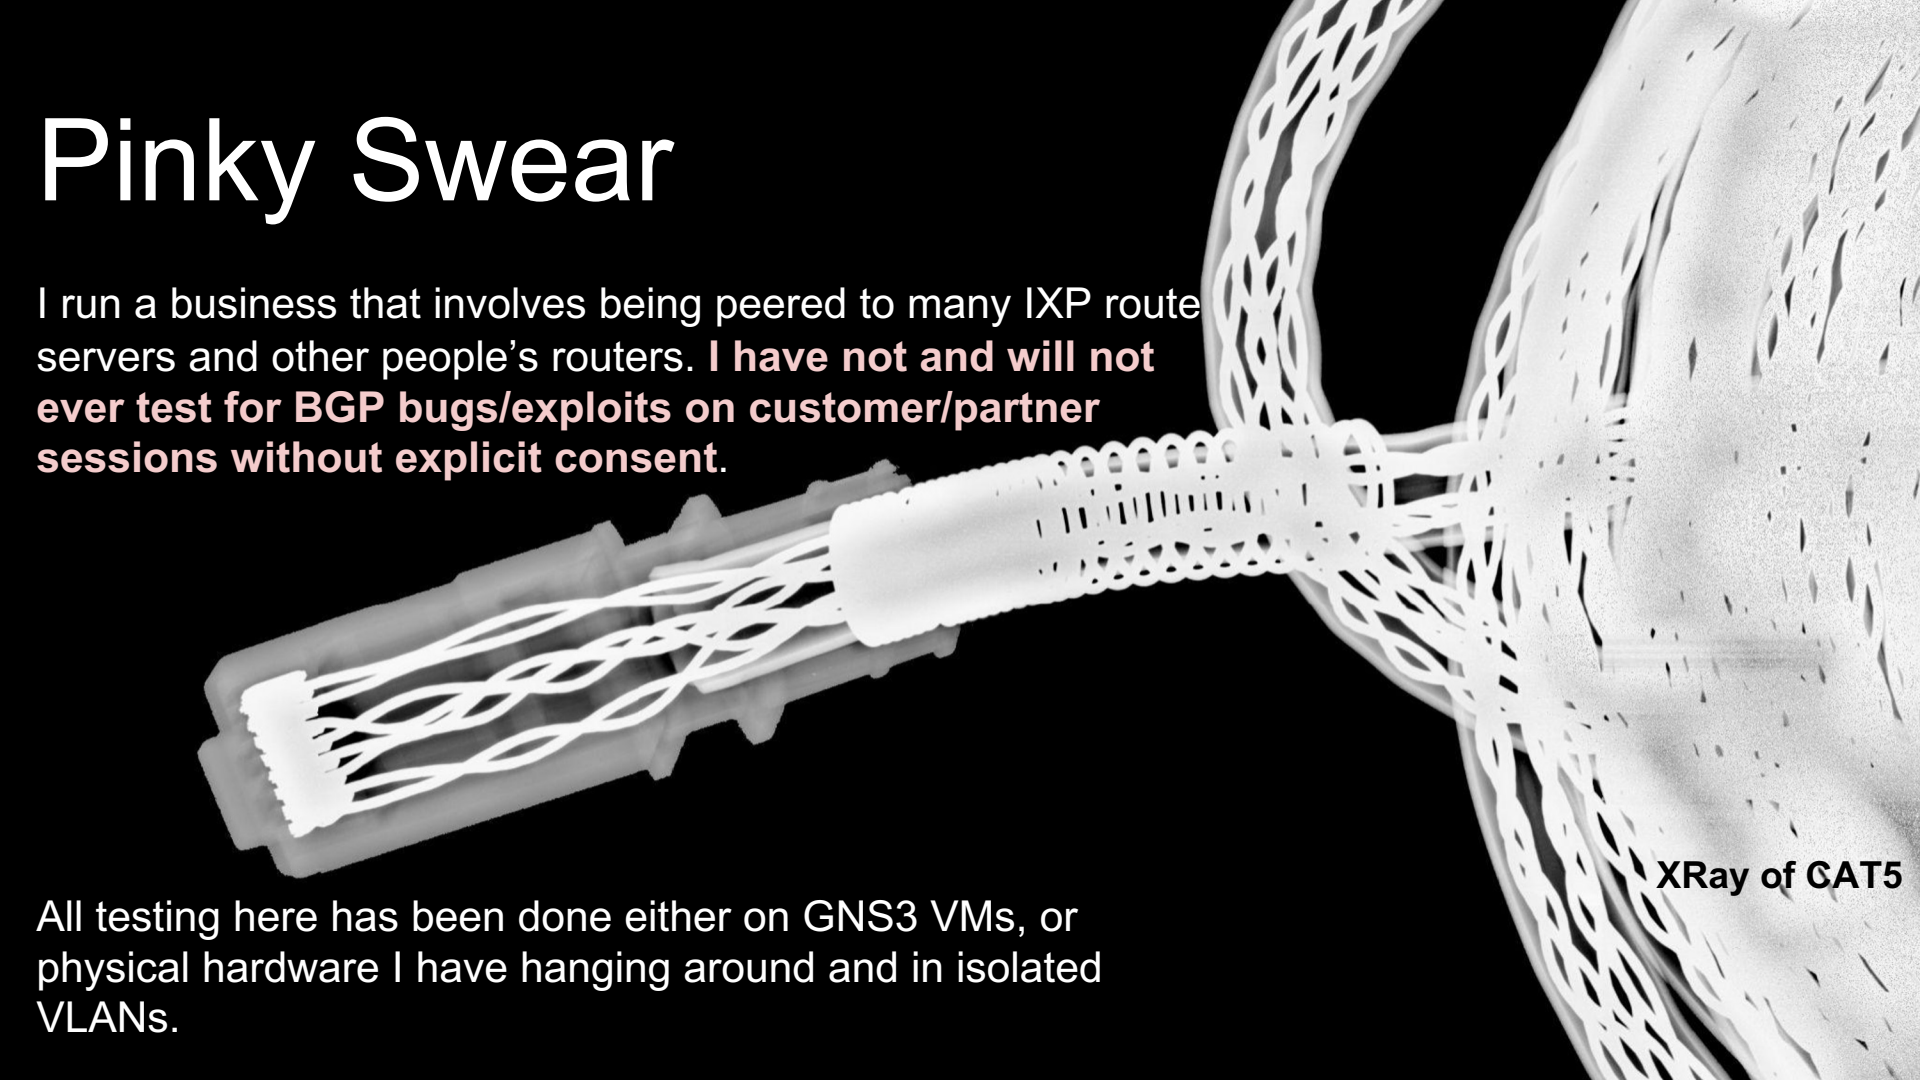

XRay of CAT5

All testing here has been done either on GNS3 VMs, or physical hardware I have hanging around and in isolated VLANs.

Recently:

- AS264366 originated a IPv6 route with a spicy BGP attribute
- This route (and spicy attribute) got carried very far
- This also seemed to cause any JunOS device that ingested it to tear down the session it received it from
- "Okay" for peering, Less okay for transit
- Colt (AS8220) got (IPv6?) de-peered from the internet
  - Other ASNs got hurt too, but Colt is the one that inconvenienced me

Recently:

- The offending payload was very boring (other than impact):
- It appears something in their network originated one of their prefixes, with BGP Attr 28 [BGP Entropy Label Capability Attribute]
  - I would assume this came from a Huawei device
- The attribute was not ***technically*** corrupted
- **This was enough to cause JunOS sessions from R17+ (lsh) to tear down the session it seems**

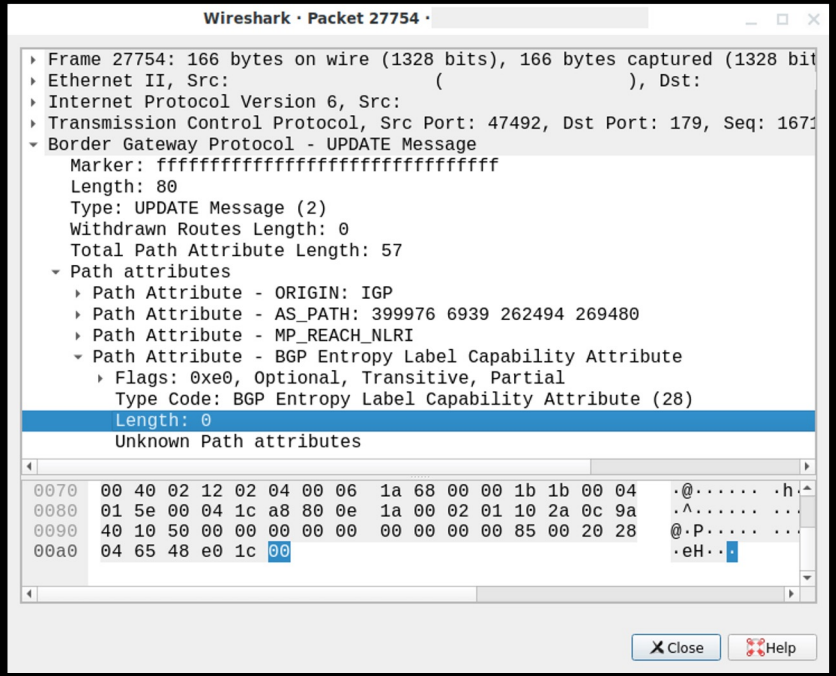

A look at BGP Attributes

- Two "sections" of a BGP UPDATE include
  - The NLRI/Withdraw data (aka prefixes)
  - The Attributes
  - * In BGP MultiProto, the NLRI/Withdraw are also in the attributes
- These attributes contain stuff like:
  - AS Path
  - Community values
  - Local Pref/MED
  - Aggregation info
  - etc

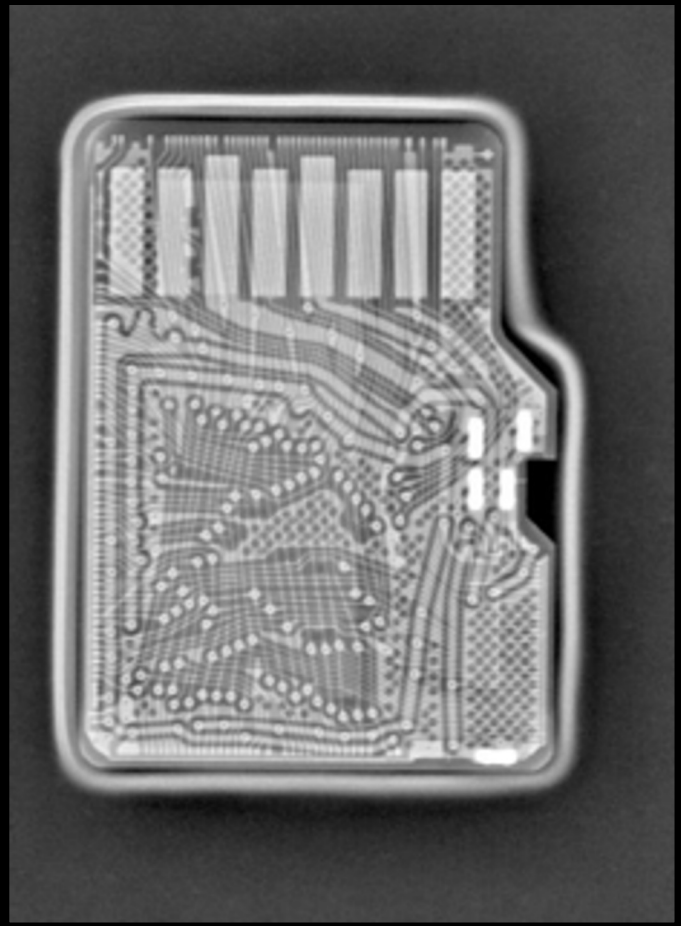

A look at BGP Attributes

- There are a lot of different BGP path attributes defined.
- Most (209) are unassigned, 14 are deprecated, and 32 are "active"
- Only a handful of these are expected on the "normal" internet routing table

- But surely they all are handled correctly, ***right????***

Value	Code
0	Reserved
1	ORIGIN
2	AS_PATH
3	NEXT_HOP
4	MULTI_EXIT_DISC
5	LOCAL_PREF
6	ATOMIC_AGGREGATE
7	AGGREGATOR
8	COMMUNITIES
9	ORIGINATOR_ID
10	CLUSTER_LIST
11	DPA (deprecated)
12	ADVERTISER (historic) (deprecated)
13	RCID_PATH / CLUSTER_ID (Historic) (deprecated)
14	MP_REACH_NLRI
15	MP_UNREACH_NLRI
16	EXTENDED COMMUNITIES
17	AS4_PATH
18	AS4_AGGREGATOR
19	SAFI Specific Attribute (SSA) (deprecated)
20	Connector Attribute (deprecated)
etc	Goes up until 255

https://www.iana.org/assignments/bgp-parameters/bgp-parameters.xhtml#bgp-parameters-2

Un-good scenario

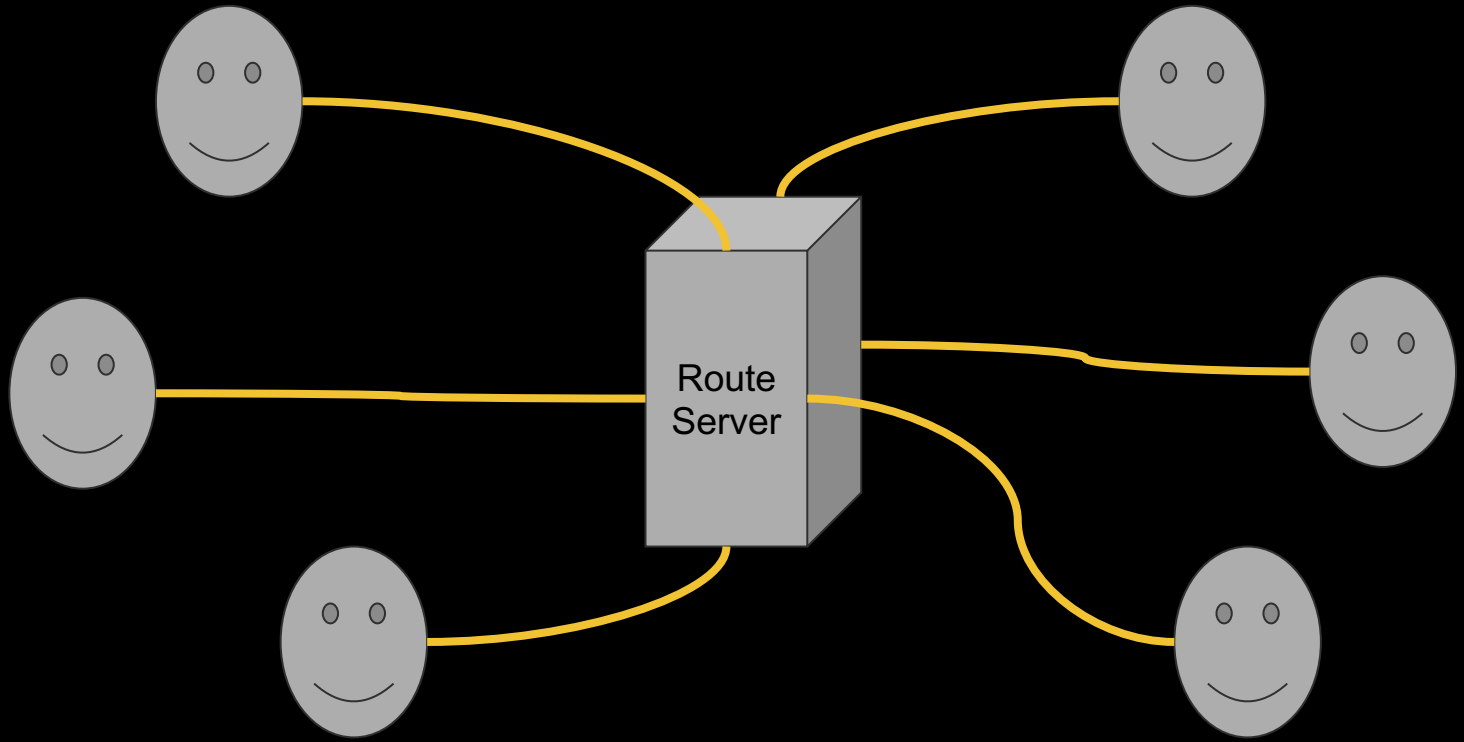

Un-good scenario

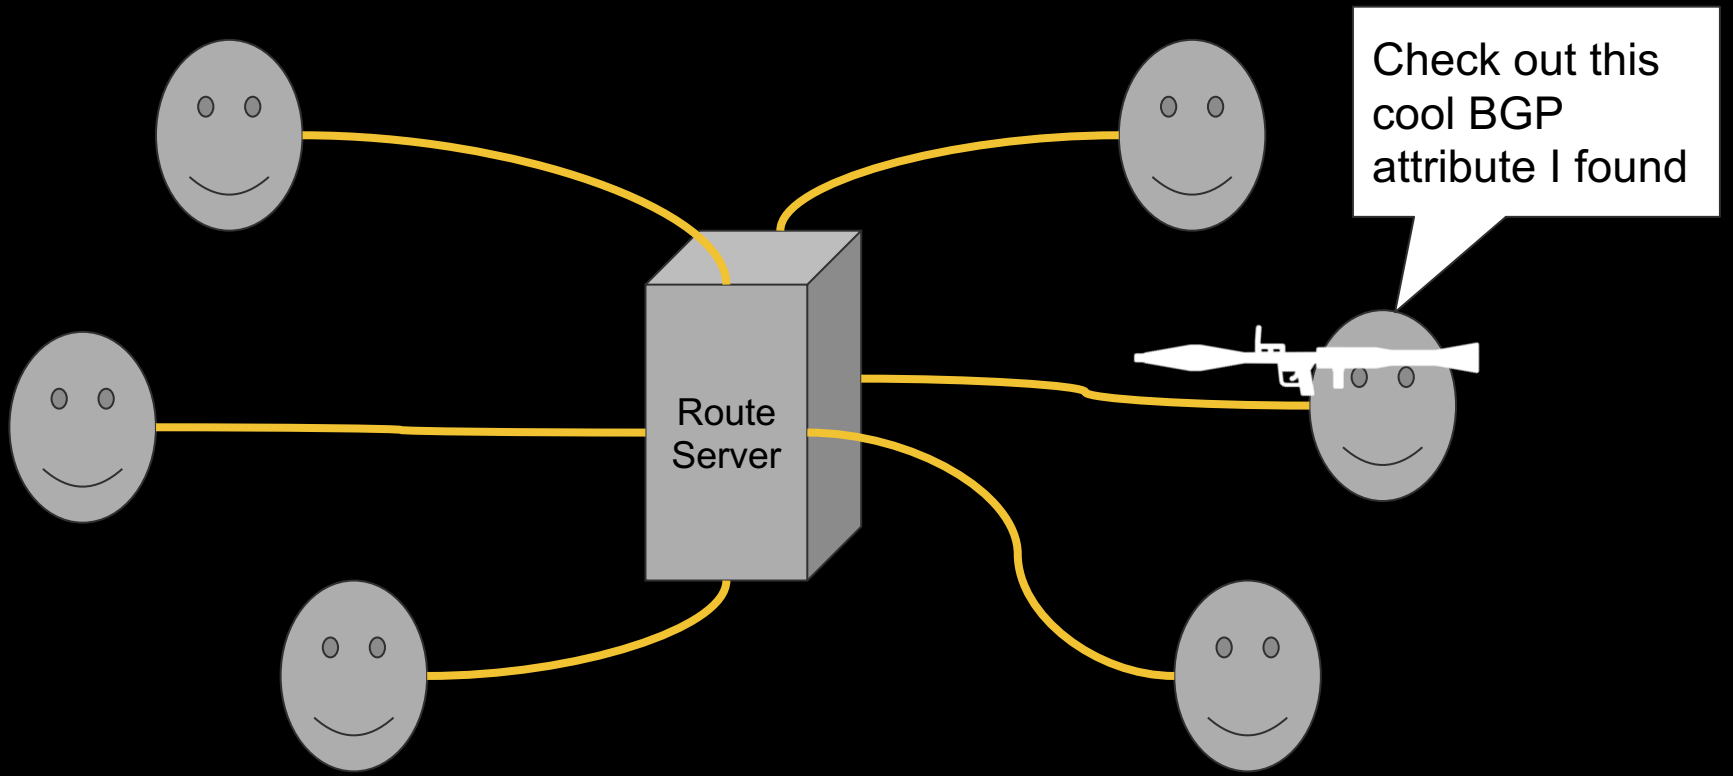

Un-good scenario

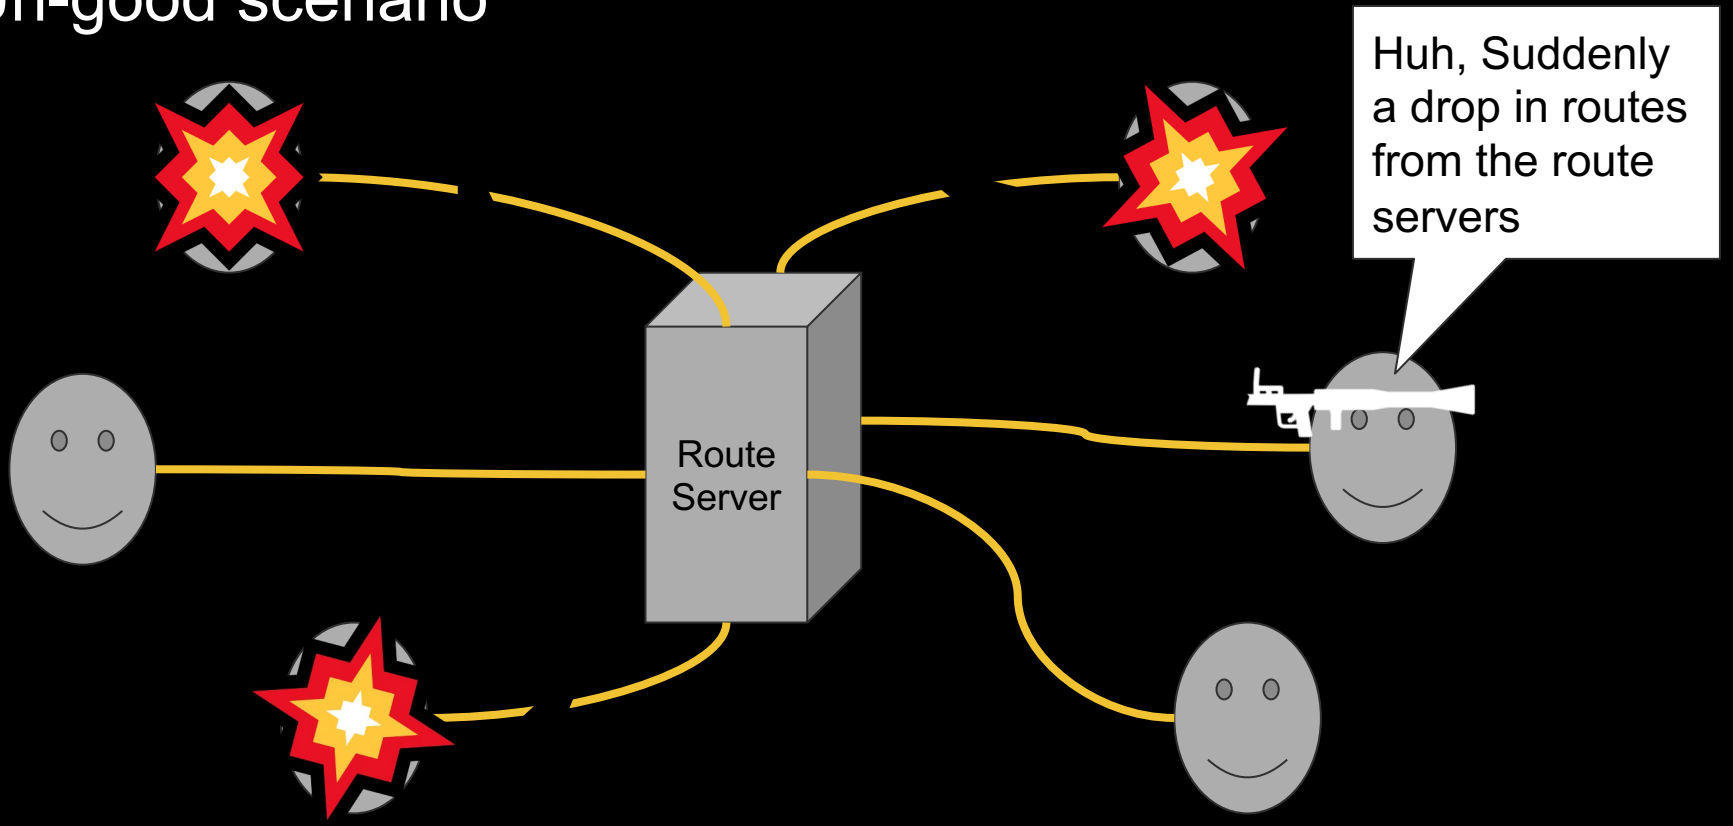

Really un-great scenario

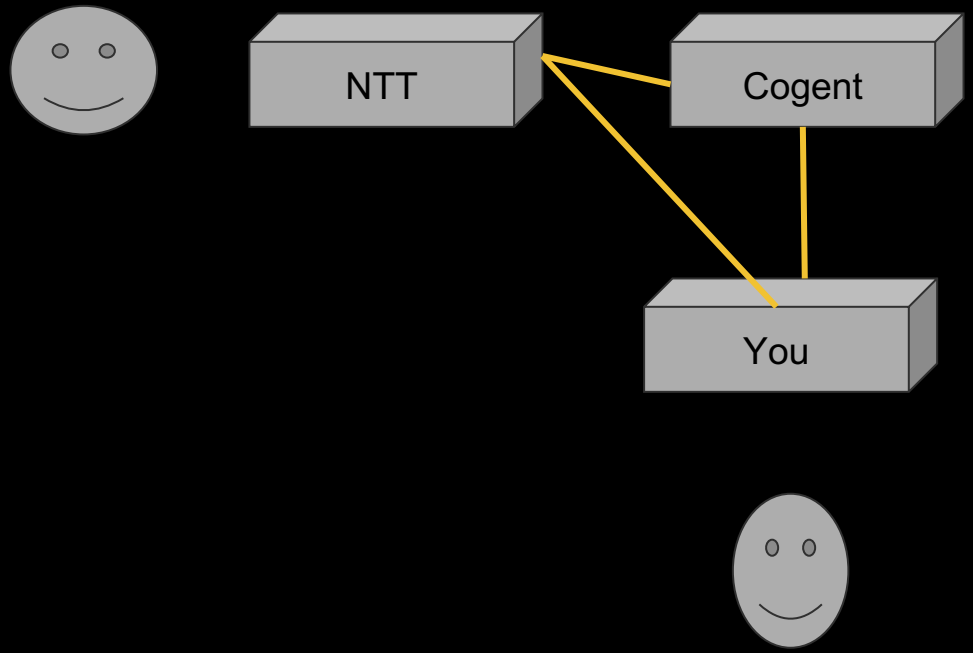

Really un-great scenario

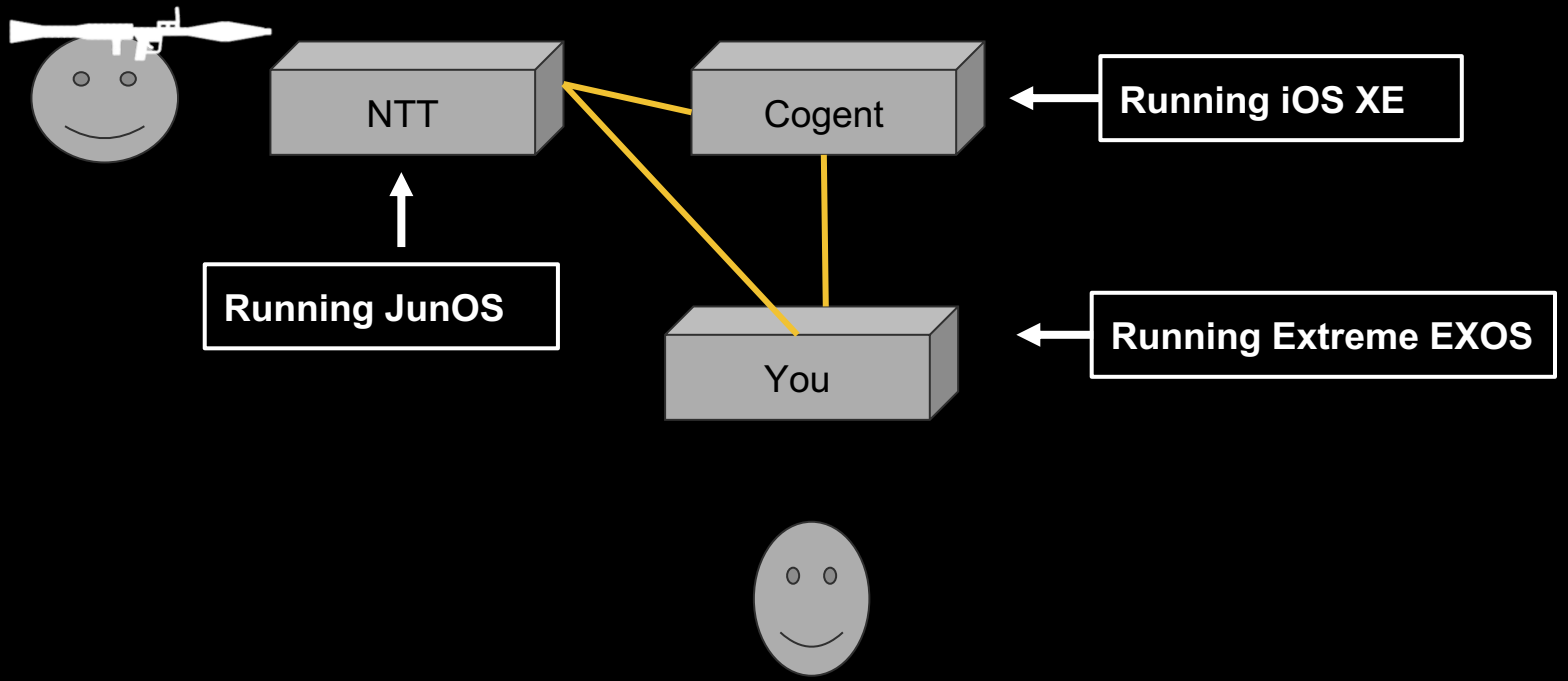

Really un-great scenario

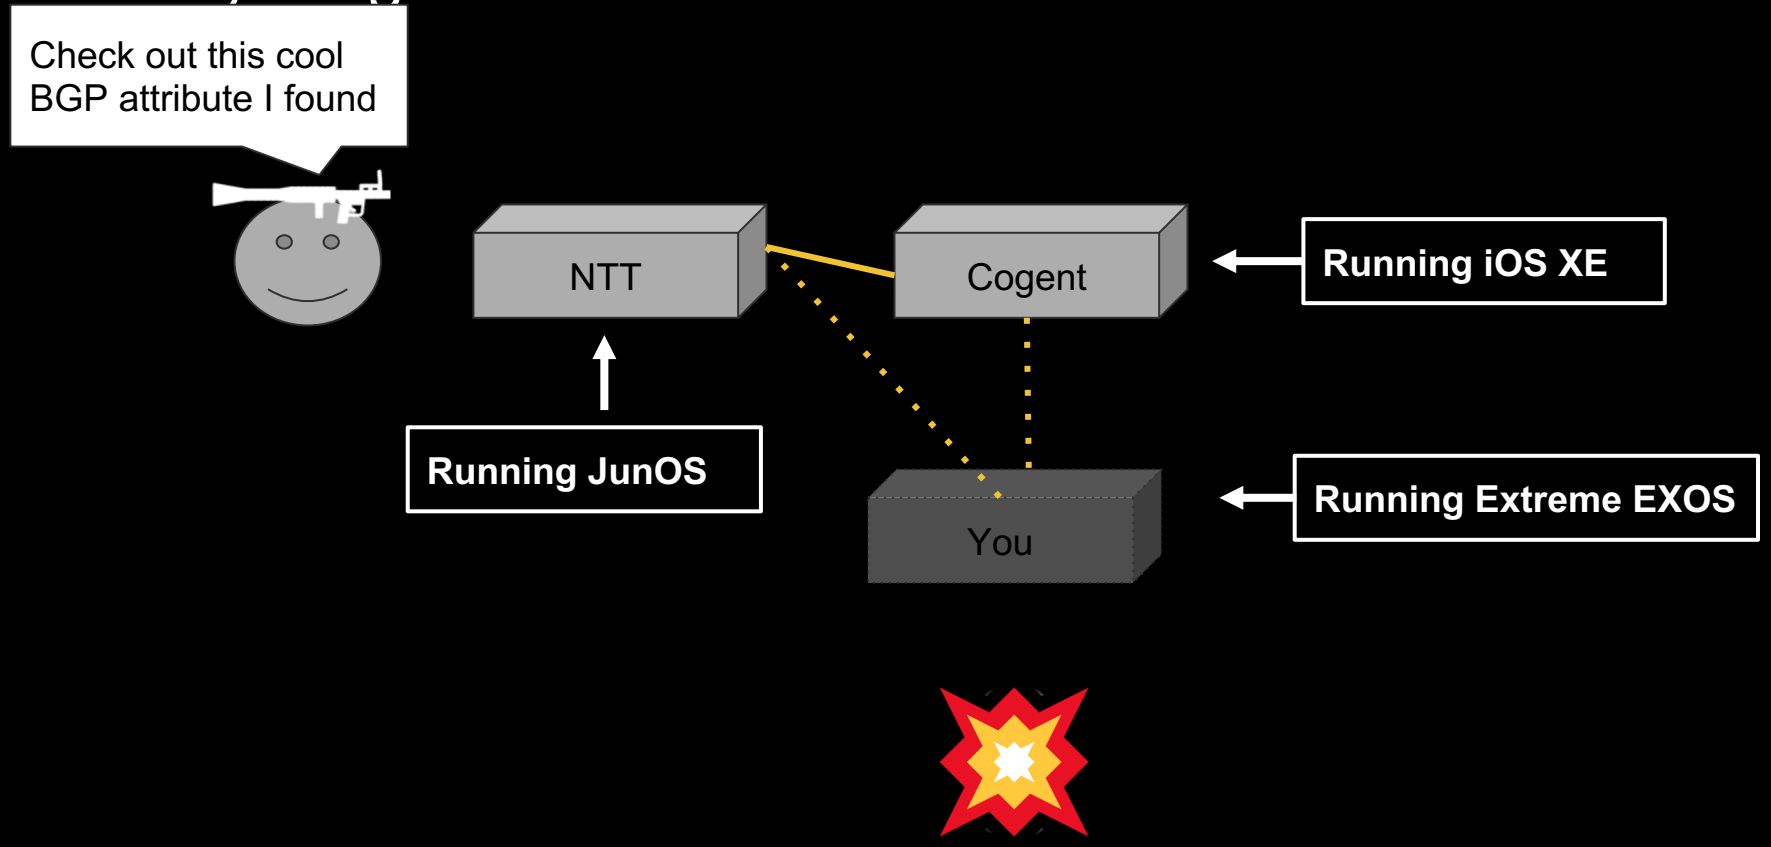

Funnily enough…

RFC 7606 / 9. Security Considerations

> This specification addresses the vulnerability of a BGP speaker to a potential attack whereby a distant attacker can generate a malformed optional transitive attribute that is not recognized by intervening routers. Since the intervening routers do not recognize the attribute, they propagate it without checking it. When the malformed attribute arrives at a router that does recognize the given attribute type, that router resets the session over which it arrived. Since significant fan-out can occur between the attacker and the routers that do recognize the attribute type, **this attack could potentially be particularly harmful.**

First time where I've seen a RFC Security Considerations be "on the money"

Fuzzing setup

- Go through all 1->255 BGP Attribute types
- Generate progressively more and more random bytes inside them
- To check for "BGP Worm" status, we will relay it though BIRD 2 to ensure it's viable that it will transmit through a Route Server

- Good fuzzers should be able to run unattended and find things
  - To check if the "Device Under Test" (DUT) router is still connected, we monitor a prefix that the DUT is originating and log a failure if the prefix is withdrawn, and wait for it to come back after session restart

Fuzzing setup

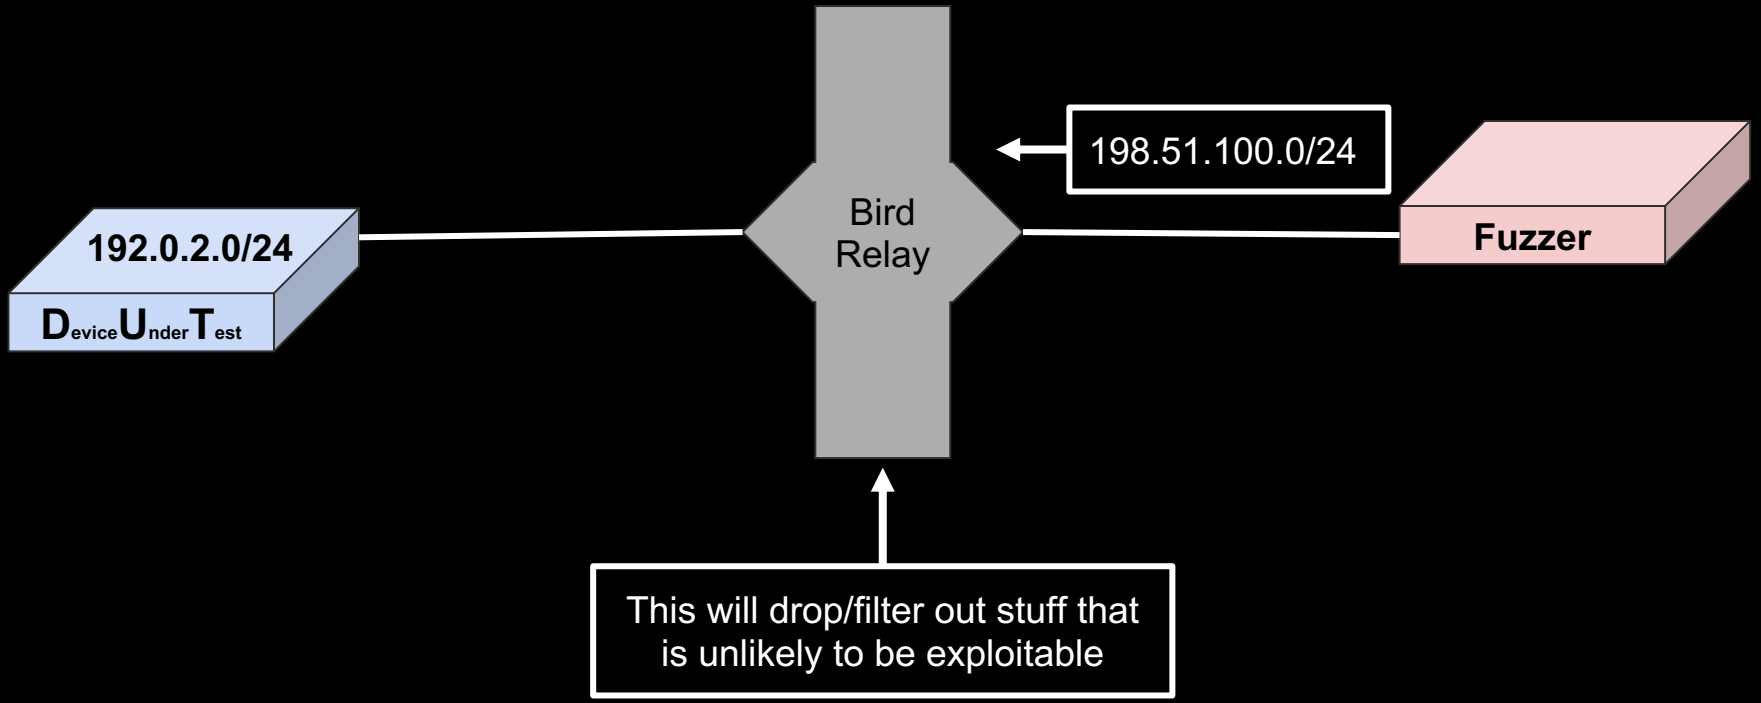

Fuzzing setup

```
# ./internet-bullets -first.hop 192.168.5.2

2023/07/05 13:50:51 Establishing Connection to first hop

2023/07/05 13:50:51 waiting for prefix to come back

2023/07/05 13:50:51 MESSAGE_OPEN

2023/07/05 13:50:51 BGP MESSAGE_KEEPALIVE sent

2023/07/05 13:50:51 MESSAGE_UPDATE

2023/07/05 13:50:51 Announce 192.0.2.0/24

2023/07/05 13:50:51 MESSAGE_UPDATE

2023/07/05 13:51:25 BGP MESSAGE_KEEPALIVE sent
```

```
$ cloc .
       1 text file.
       1 unique file.
       0 files ignored.

github.com/AlDanial/cloc v 1.82  T=0.01 s ----------------------
-------------------------------------------
Language             files         blank       comment     code
-------------------------------------------------------------------
Go                   1             45          42          217
-------------------------------------------------------------------
```

```
root@pass:/etc/bird# birdc s ro  all
BIRD 2.0.7 ready.
Table master4:
198.51.100.0/24      unicast [fuzzer 21:31:24.378] * (100)
[AS65001?]
                via 192.168.5.1 on ens5
                Type: BGP univ
                BGP.origin: Incomplete
                BGP.as_path: 65001
                BGP.next_hop: 192.168.5.1
                BGP.local_pref: 100
                BGP.community: (123,2345)
                BGP.ec [t]: 7d cc c7 30
192.0.2.0/24         unicast [nokia 21:15:11.775] * (100) [AS1i]
                via 192.0.2.1 on ens4
                Type: BGP univ
                BGP.origin: IGP
                BGP.as_path: 1
                BGP.next_hop: 192.0.2.1
                BGP.local_pref: 100
```

The fuzzer findings

MikroTik

- Zero issues, did not log a single error
- Also this was RouterOS7.7 so it's unknown if anyone actually uses this

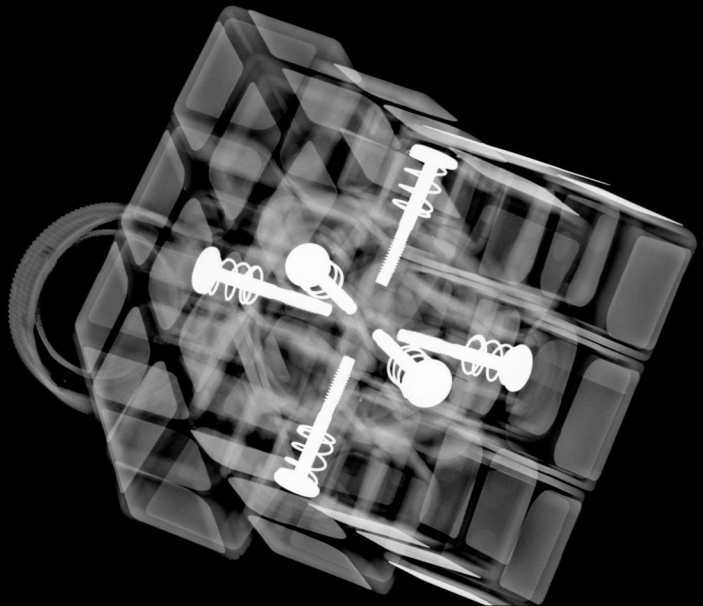

Ubiquiti

- No problem! All clear
- I suspect they forked Quagga before it grew the features that would end up problematic

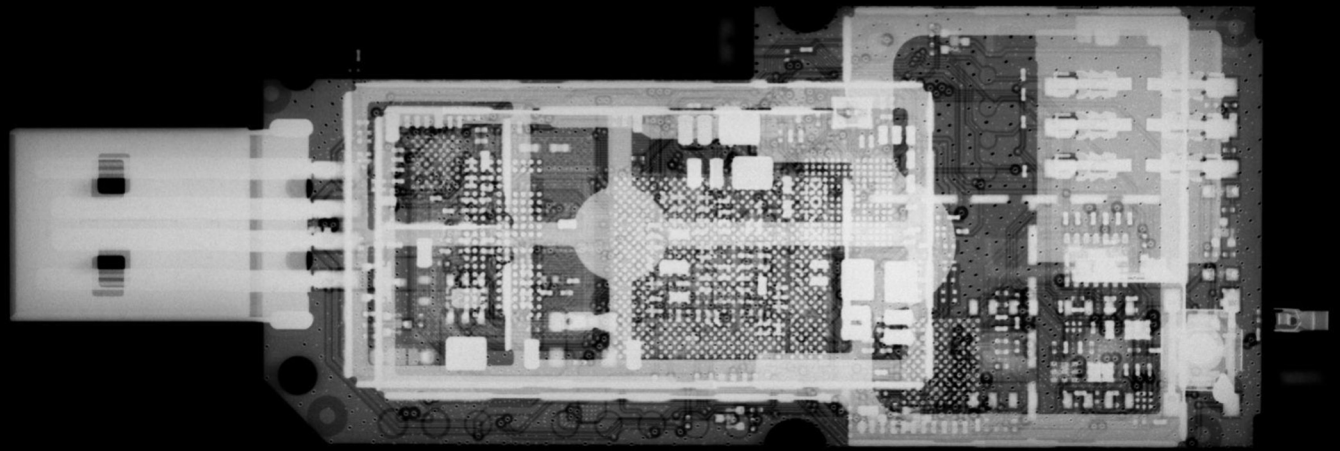

Arista EOS

- No errors, No obvious logging (though I'm sure there is some)
- You can check if withdraw behaviour has triggered for you by running:
  - show ip bgp neighbors update errors

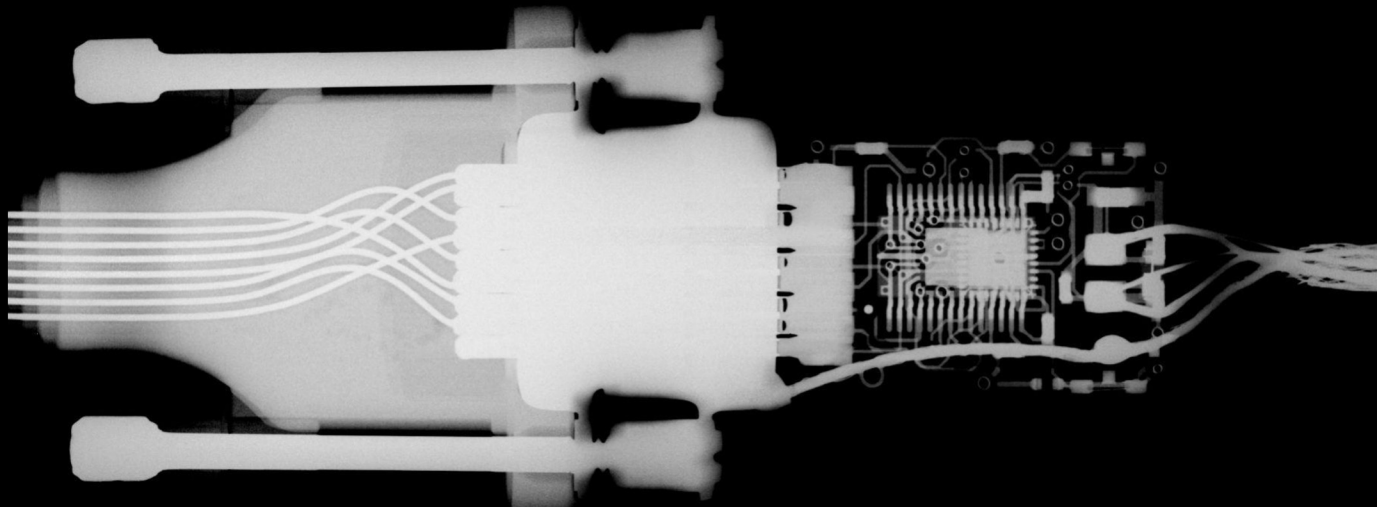

Cisco IOS-XE / IOS-XR

- No errors
- Logs the issue verbosely (maybe a little too verbose)

```
RP/0/RP0/CPU0:ios#RP/0/RP0/CPU0:Oct 10 15:22:22.596 UTC: bgp[1094]: %ROUTING-BGP-3-
MALFORM_UPDATE : Malformed UPDATE message received from neighbor 192.0.2.2 (VRF:
default) - message length 78 bytes, error flags 0x00200000, action taken
"DiscardAttr". Error details: "Error 0x00200000, Field "Attr-data", Attribute 19
(Flags 0xe0, Length 17), Data [e01311f2b6ea2c46ce1b44b82b74049c]". NLRIs: [IPv4
Unicast] 198.51.100.0/24
```

```
*Jul  5 13:51:18.582: %BGP-6-MSGDUMP_LIMIT: unsupported or mal-formatted message received from
192.0.2.2:
FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 003E 0200 0000 2340 0101 0240 020A 0202
0000 0002 0000 FDE9 4003 04C0 0002 02C0 0804 007B 0929 E01B 0164 18C6 3364
*Jul  5 13:51:18.582: %BGP-6-MALFORMEDATTR: Malformed attribute in (BGP(0) Prefixes:
198.51.100.0/24 ) received from 192.0.2.2,
*Jul  5 13:51:20.582: %BGP-6-ATTR_FLAG: BGP update error: 192.0.2.2 Wrong flag 0xE0 received for
LS attribute attribute (fixed by error handling)
```

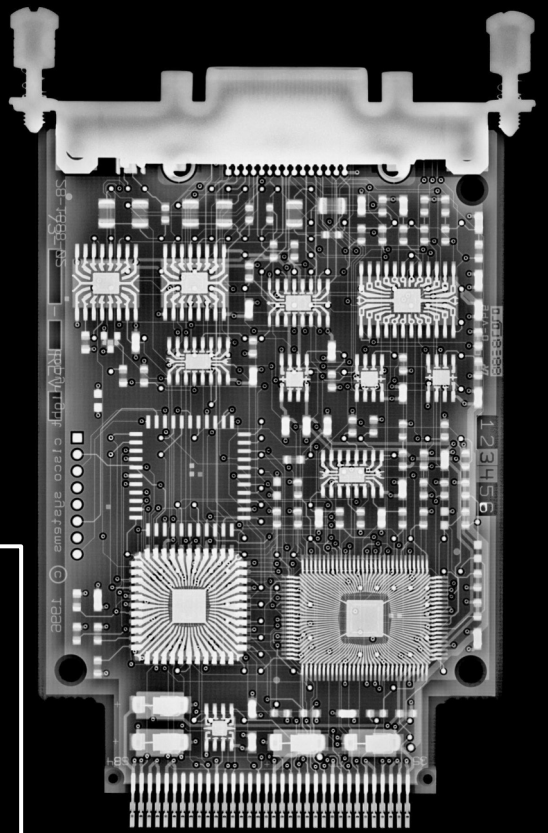

JunOS

as seen in JSA72510

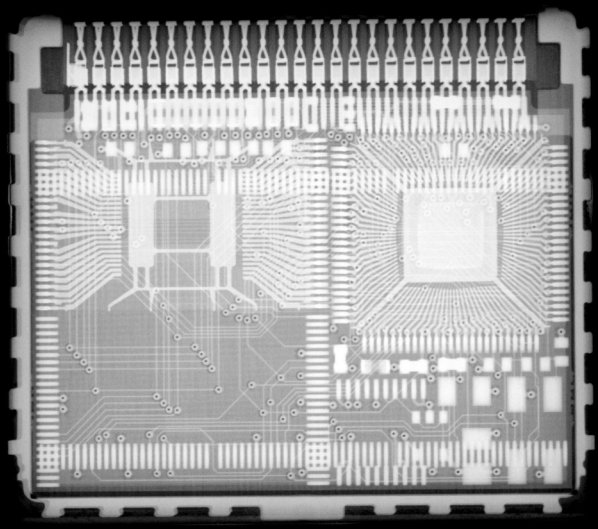

- Attr 28 [BGP Entropy Label Capability Attribute]
  - (The one that spawned this entire adventure)
- Attr 29 [BGP-LS Attribute,[RFC-ietf-idr-rfc7752bis-16]
  - The new one that this work discovered (JSA72510)
- Mitigated with:

```
[edit protocols bgp]
root# show
group FUZZ-VM {
    import yolo;
    export send-direct;
    peer-as 4200000001;
    local-as 4200000002;
    neighbor 192.0.2.2;
}
bgp-error-tolerance;
```

Lots of people already have enabled this after the previous (Attr 28) incident

Nokia SR-OS

as seen in 23-0450b

Connector Attribute (deprecated),[RFC6037]
Tunnel Encapsulation,[RFC9012]
IPv6 Extended Community
BGP-LS Attribute
BGP Prefix-SID,[RFC8669]

- **Many** ways to pop a session by default (20,23,25,29,40)
- You can mitigate it by using `update-fault-tolerance`

```
bgp
    group "eBGP"
        export "yes"
        error-handling
            update-fault-tolerance
        exit
        neighbor 192.0.2.2
            peer-as 2
        exit
    exit
    no shutdown
exit
```

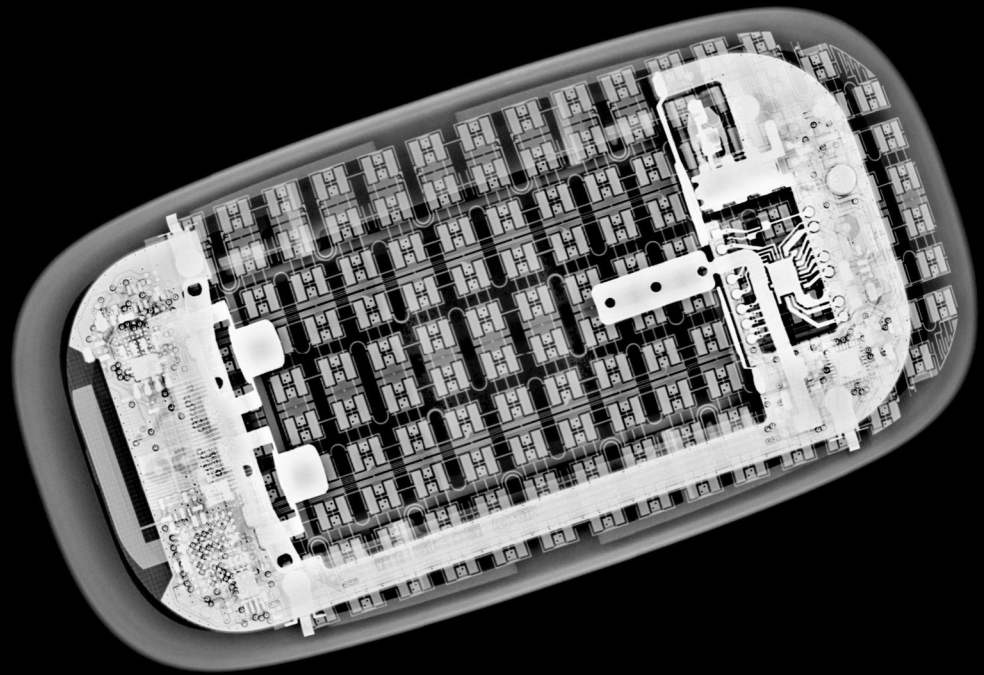

Nokia SR-OS

Connector Attribute (deprecated),[RFC6037]
Tunnel Encapsulation,[RFC9012]
IPv6 Extended Community
BGP-LS Attribute
BGP Prefix-SID,[RFC8669]

- **Many** ways to pop a session by default (20,23,25,29,40)
- You can mitigate it by using `update-fault-tolerance`

```
bgp
    group "eBGP"
        export "yes"
        error-handling
            update-fault-tolerance
        exit
        neighbor 192.0.2.2
            peer-as 2
        exit
    exit
    no shutdown
exit
```

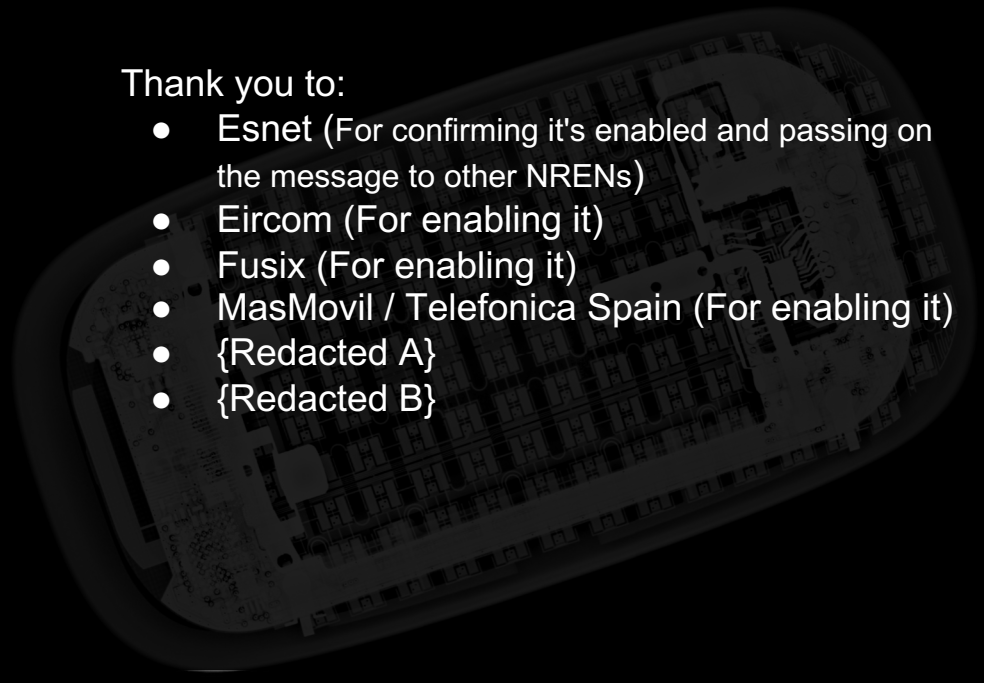

Thank you to:

- Esnet (For confirming it's enabled and passing on the message to other NRENs)
- Eircom (For enabling it)
- Fusix (For enabling it)
- MasMovil / Telefonica Spain (For enabling it)
- {Redacted A}
- {Redacted B}

Huawei NetEngine (NE40)

- No problems detected, No logs found about the errors though
  - But I may be just be unable to figure out the NetEngine CLI
- Very hard to acquire testing images for Huawei, it doesn't help that I am not allowed to import Huawei into my country

- There may be bugs in other products, I just can't test them.

FRR / Pica8 / SONIC / Loads of vendors

- Explodes on Attr 23 (Tunnel Encapsulation,[RFC9012])
- Assigned: CVE-2023-38802 / Fixed in https://github.com/FRRouting/frr/pull/14290 (Post Public disclosure)

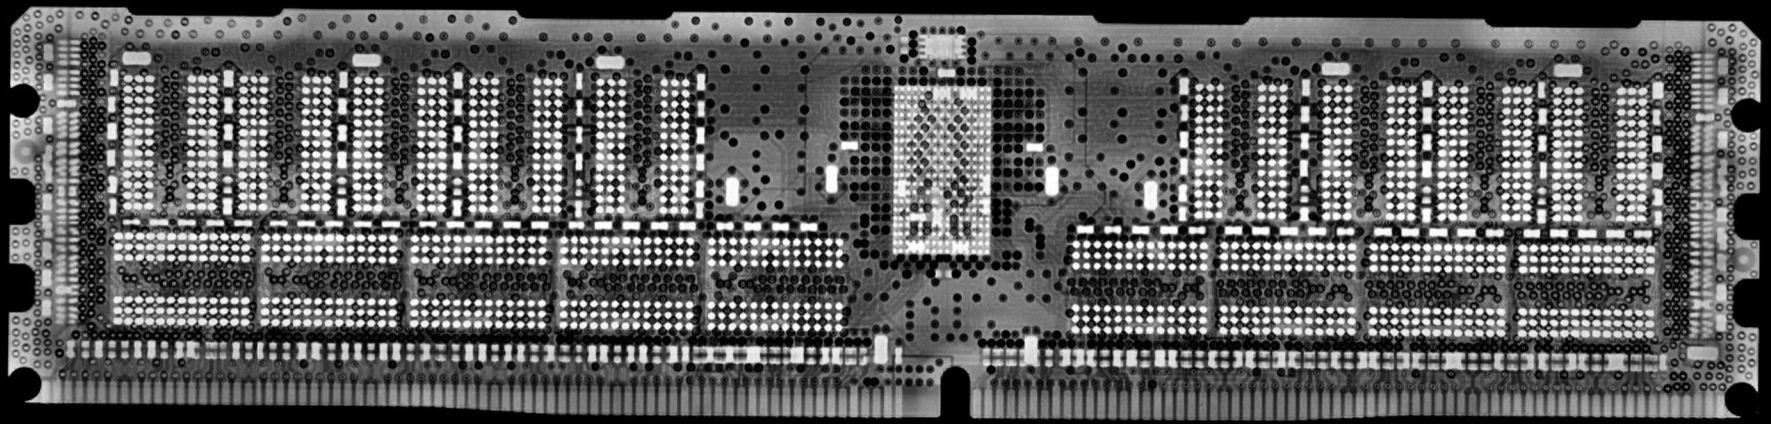

OpenBGPd (OpenBSD)

- Exploded on invalid OTC (Attr 35)
- Logs most other bad packets
- OpenBGPd is increasingly used in route servers

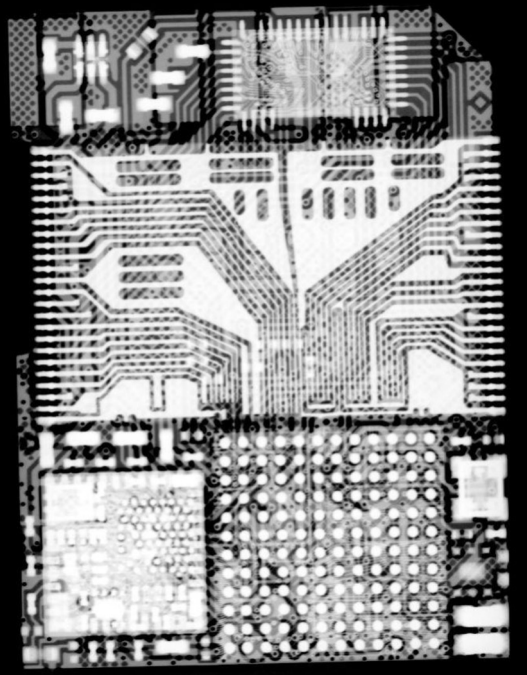

- Actually exposed more than one bug in OpenBGPd
  - Only this one was reachable to remote routers it seems
- **Fixed in OpenBSD 7.3 Errta 006**
- Assigned: CVE-2023-38283

EXOS / Extreme-ly bad

- Explodes on:
  - Attr (21)
    - AS_PATHLIMIT (deprecated),[draft-ietf-idr-as-pathlimit]
  - Attr (25)
    - IPv6 Address Specific Extended Community,[RFC5701]
- No mitigating config
- You could de-peer most of HE (and others)

- CVE: CVE-2023-40457 (Disputed by Extreme)

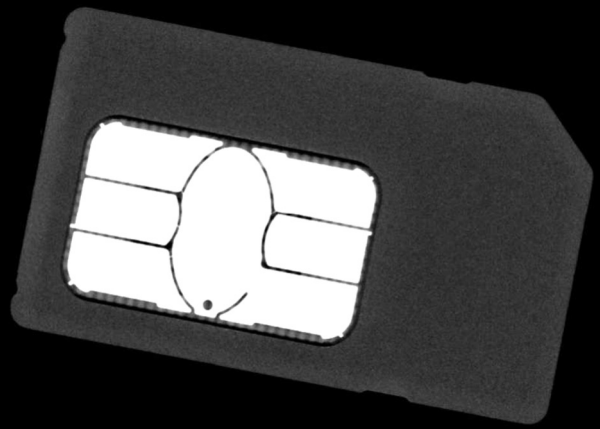

Extreme won't commit to fixing this.

After review of all the material, **we are not considering this a vulnerability due to the presence of RFC 7606**, as well as a history of documentation expressing these concerns all the way back to early 2000s, if not earlier. Malformed attributes are not a novel concept as an attack vector to BGP networks, as evidenced by RFC 7606, which is almost a decade old. As such, customers that have chosen to not require or implement RFC 7606 have done so willingly and with knowledge of what is needed to defend against these types of attacks. Thus, the expectation that we'll reset our BGP sessions based on RFC 4271 attribute handling is proper. We do abide by other RFCs, in which we claim support, that update RFC 4271. Other vendors do claim RFC 7606 support and have been sharing these controls as a mitigation to malformed attribute response. They don't appear to be producing new work product to account for these behaviors. **We are evaluating support for RFC 7606 as a future feature.** Obviously, if customers desire a different response, we'll work through our normal feature request pipelines to address. This is no different than any other RFC support request.

Full Email Exchange here: https://blog.benjojo.co.uk/asset/JgH8G5duO1

To clarify:

- Any AS can emit a BGP message with a corrupted IPv6 Address Specific Extended Community
- It will get carried around a number of global networks
- When a Extreme device running EXOS ingests this, it will reset the BGP session it came from
  - This will likely be a transit BGP session, causing that transit to flap
  - It will flap over and over because when it reconnects, it will get the same poisoned data
- Thinking about this even more, the requirement for the EXOS device to be on the edge is not even true, **a core iBGP full table device inside a network will do the same thing**

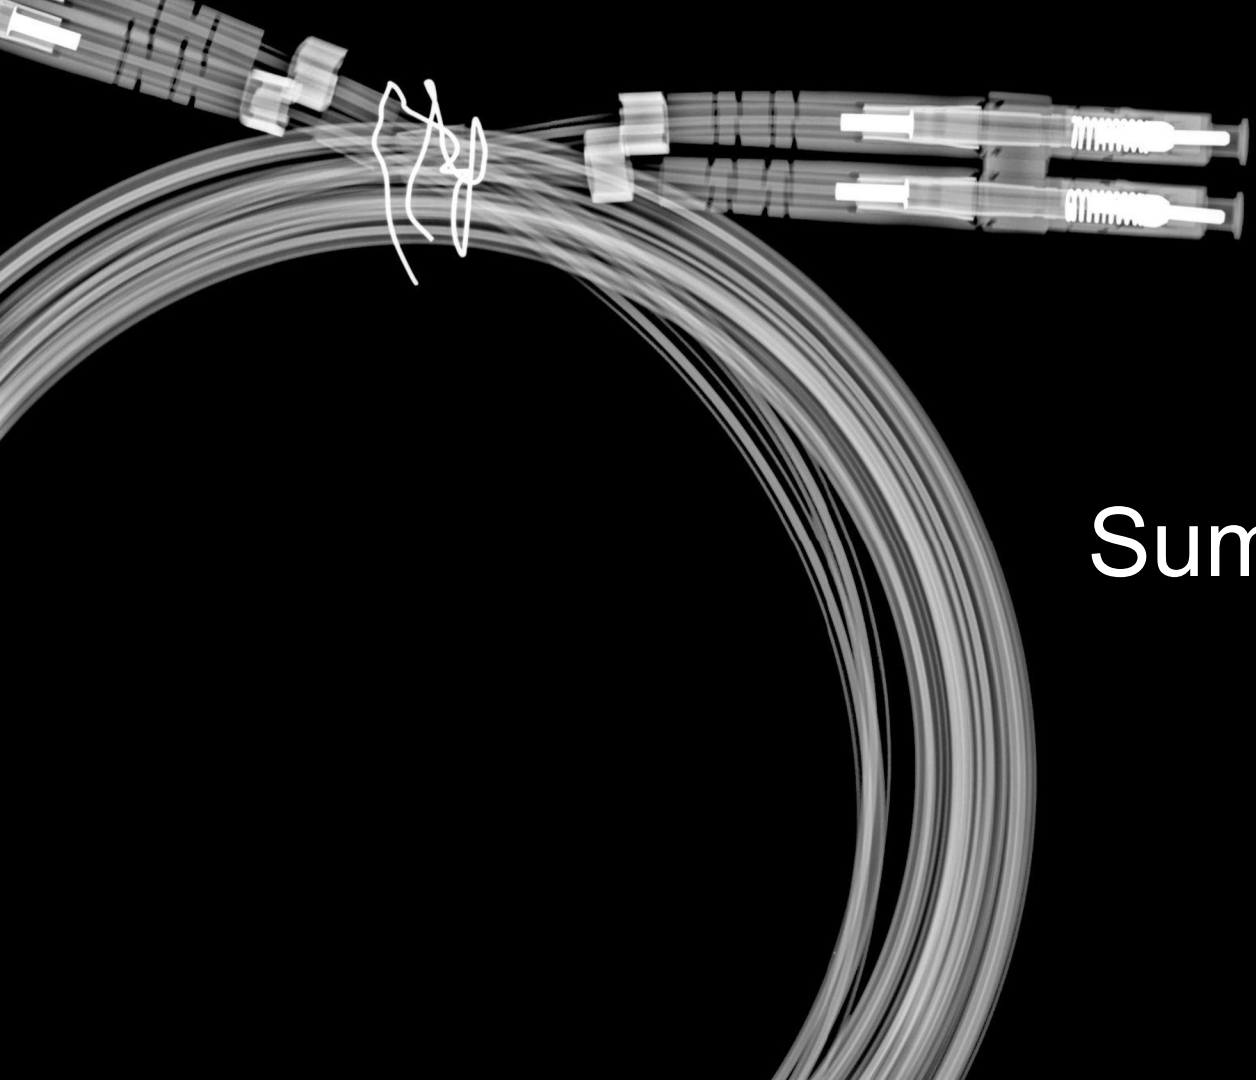

Summary

Platform Status Quo

Type	Rating	Vendors
Default RFC 7606	A+	MikroTik, Arista EOS, Cisco, Bird, GoBGP, NE40
Buggy RFC 7606	A	FRR/SONIC/VyOS, OpenBSD, NX-OS
Switchable RFC 7606	B	Juniper JunOS, Nokia SROS
No RFC 7606 At All, No Plans	F	Extreme EXOS

Remember:
This is *not* a new bug class for BGP. This type of BGP problem has been known for at least 13 years!

Security Response

Vendor	Rating	Comment
OpenBSD	A	Quick reply, only regret was telling them so early on
Juniper	B	Replied, Were polite and seemingly knowledgeable, eventually pushed out a JunOS patch, no date for default-safe behaviour.
Nokia	B	Replied, future (March 2024) SROS versions will be default-safe. Eventual customer communication.
FRR	C	Quick reply, acked issue, then disappeared afterwards! Only patched after public disclosure!
Extreme	D	Had to ask for contacts from many people, Security team ran down the clock to instead tell me they didn't think it was a issue

Security Response

None of these vendors have any bug bounty program, reporting this was a waste of my time

Vendor	Rating	Comment
OpenBSD	A	Quick reply, only regret was telling them so early on
Juniper	B	Replied, Were polite and seemingly knowledgeable, eventually pushed out a JunOS patch, no date for default-safe behaviour.
Nokia	B	Replied, future (March 2024) SROS versions will be default-safe. Eventual customer communication.
FRR	C	Quick reply, acked issue, then disappeared afterwards! Only patched after public disclosure!
Extreme	D	Had to ask for contacts from many people, Security team ran down the clock to instead tell me they didn't think it was a issue

Reporting these issues was deeply frustrating, I would argue that (in most cases) **it is not worth doing**.

Last Slide

Security Response

Vendor	Rating	Comment
Arista	A+	Correct behaviour in the first place
Cisco (XR/XE)	A+	Correct behaviour in the first place
Mikrotik	A+	Correct behaviour in the first place
Bird	A+	Correct behaviour in the first place
OpenBSD	A	Quick reply, only regret was telling them so early on
Juniper	B	Replied, Were polite and seemingly knowledgeable, eventually pushed out a JunOS patch, no date for default-safe behaviour.
Nokia	B	Replied, future (March 2024) SROS versions will be default-safe. Eventual customer communication.
FRR	C	Quick reply, acked issue, then disappeared afterwards! Only patched after public disclosure!
Extreme	D	Had to ask for contacts from many people, Security team ran down the clock to instead tell me they didn't think it was a issue

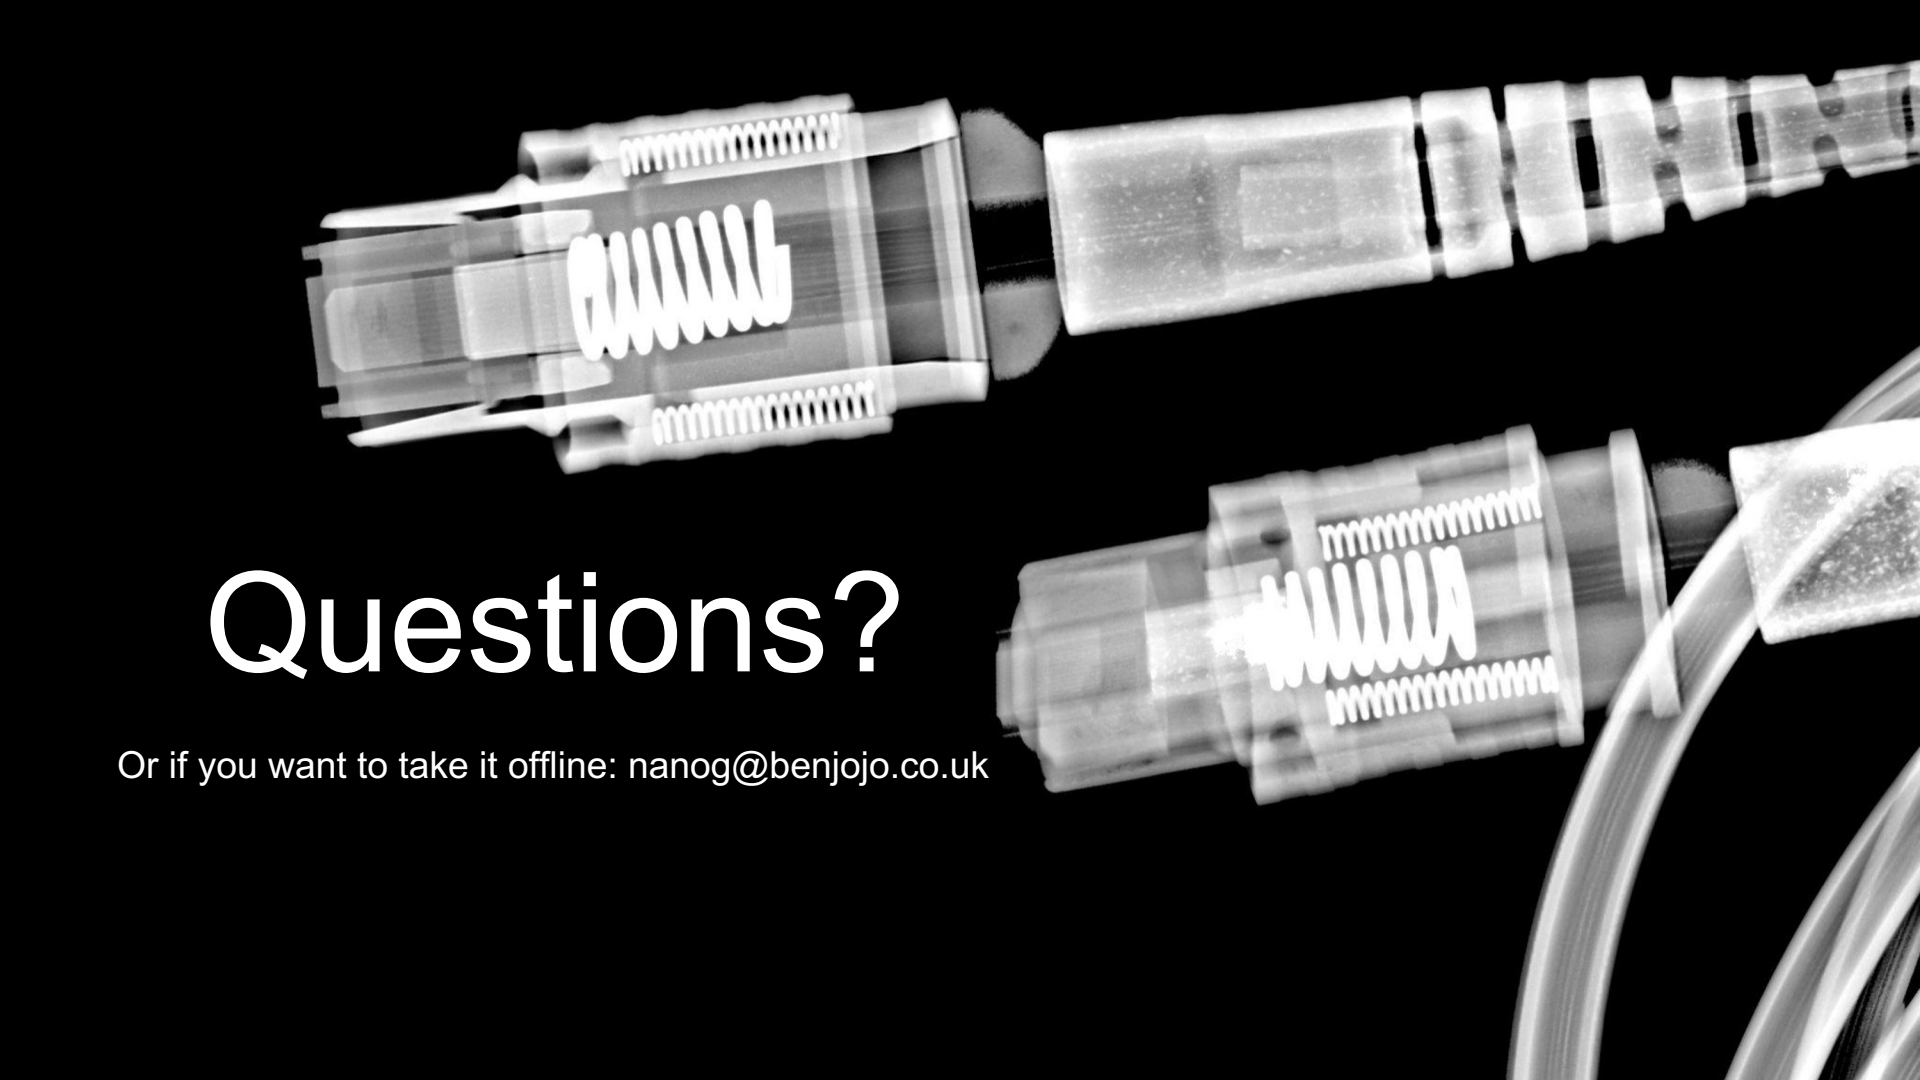

Questions?

Or if you want to take it offline: nanog@benjojo.co.uk