# Breaking your network with AS4/ASN32

…. in ten minutes

Andy Davidson - (NetSumo / LONAP) (*me*)
Rob Shakir (GX Networks)
Jonathan Oddy (Hostway UK)
Will Hargrave (LONAP / Imperial College London)

# Who should care about this talk

- Standards Authors - we have suggestions to fix the standard
- BGP Stack Authors - the standard is broken and we want you to know.
- Operators - there is software out there that can break your network

# Motivation

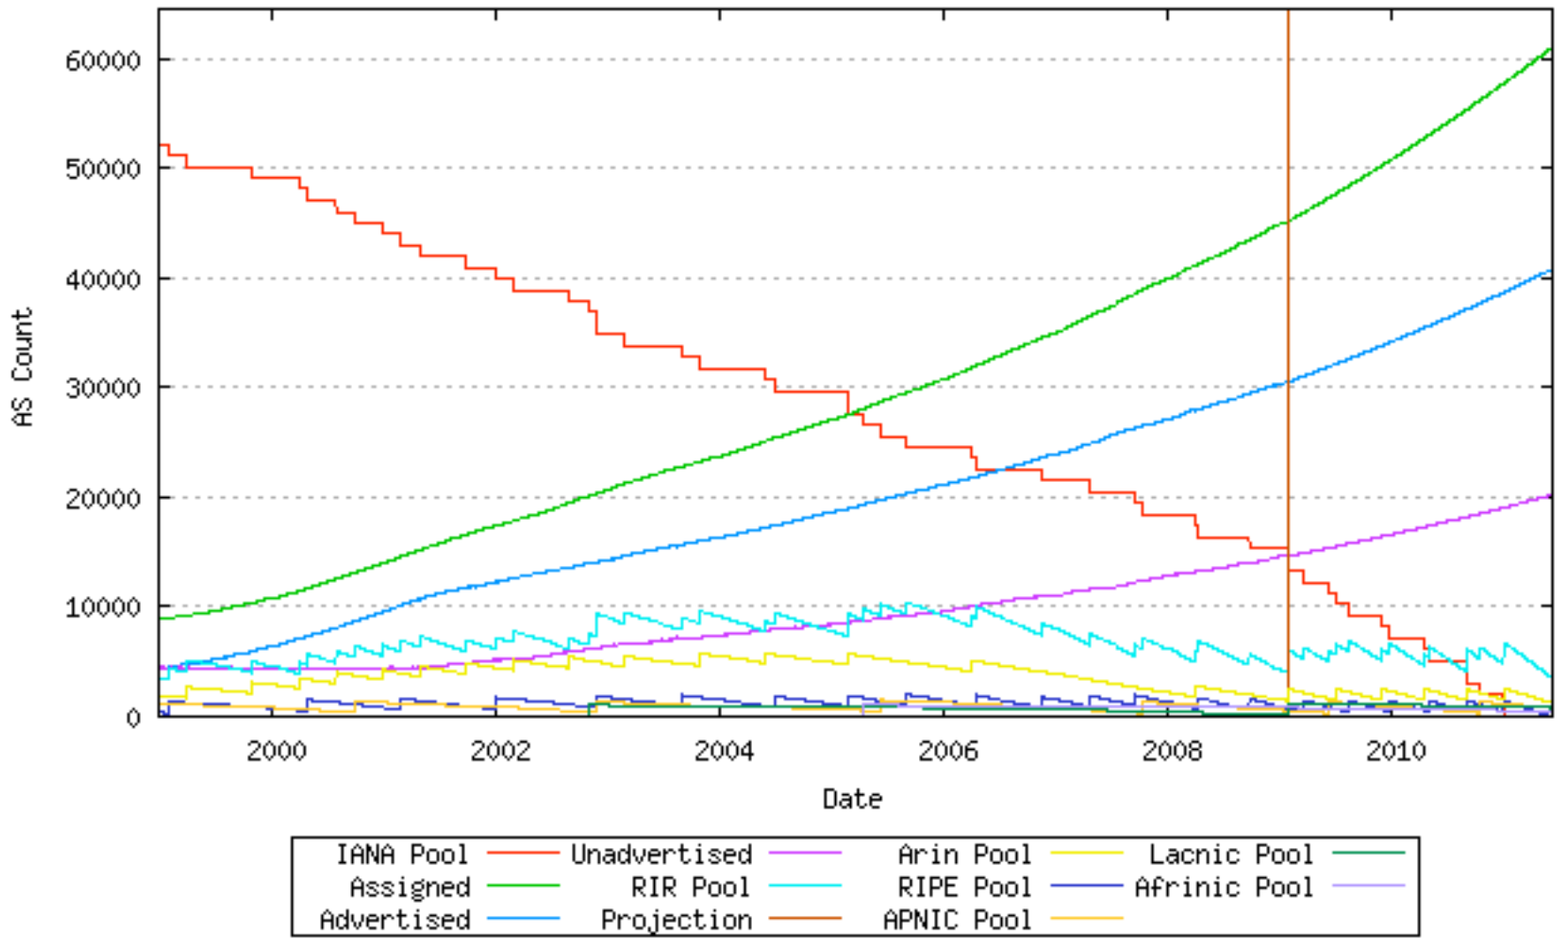


## Projection : 4 June 2011 - 16 bit ASN exhaustion

Thank you Geoff

# How do you get two whales in a car ?

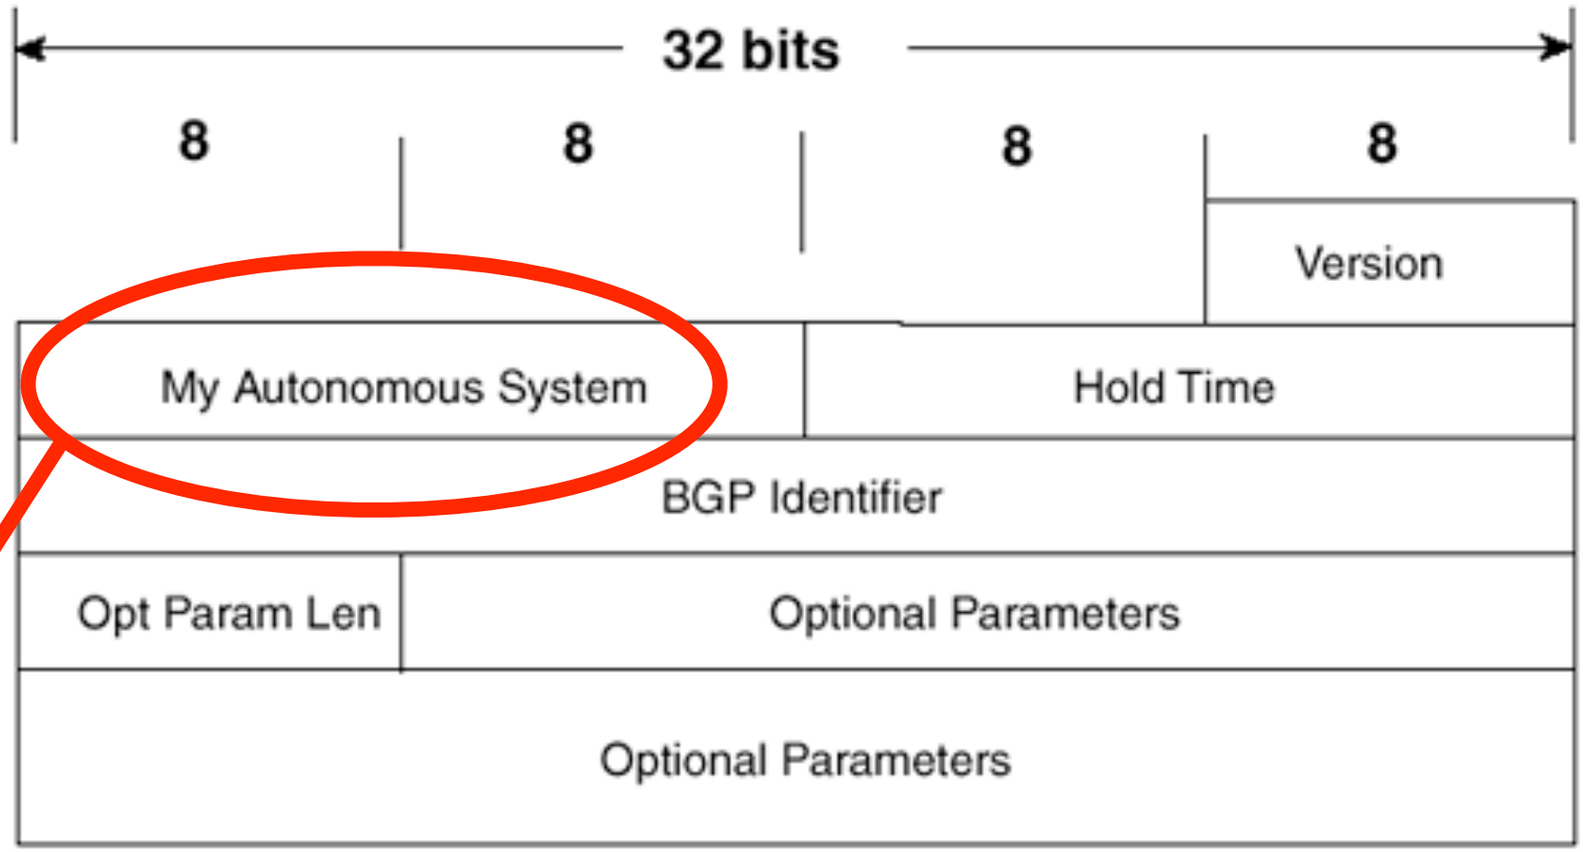


Someone's going to be AS65536 - which wont fit here

# … you get a big truck

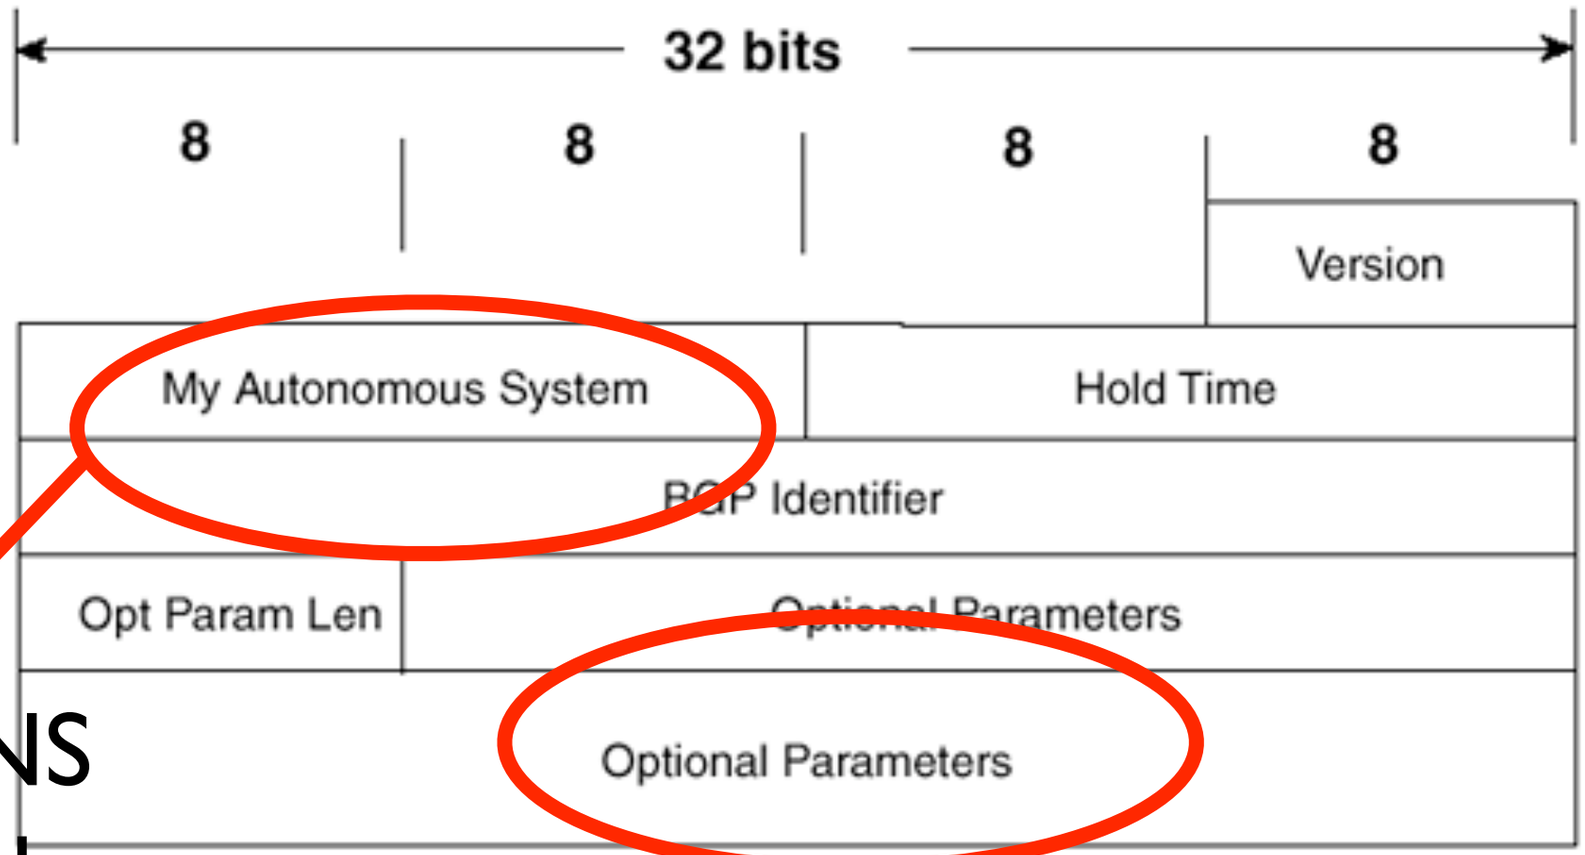


AS_TRANS
Magic Number
AS23456

My proper ASN

# ...And a bigger freeway

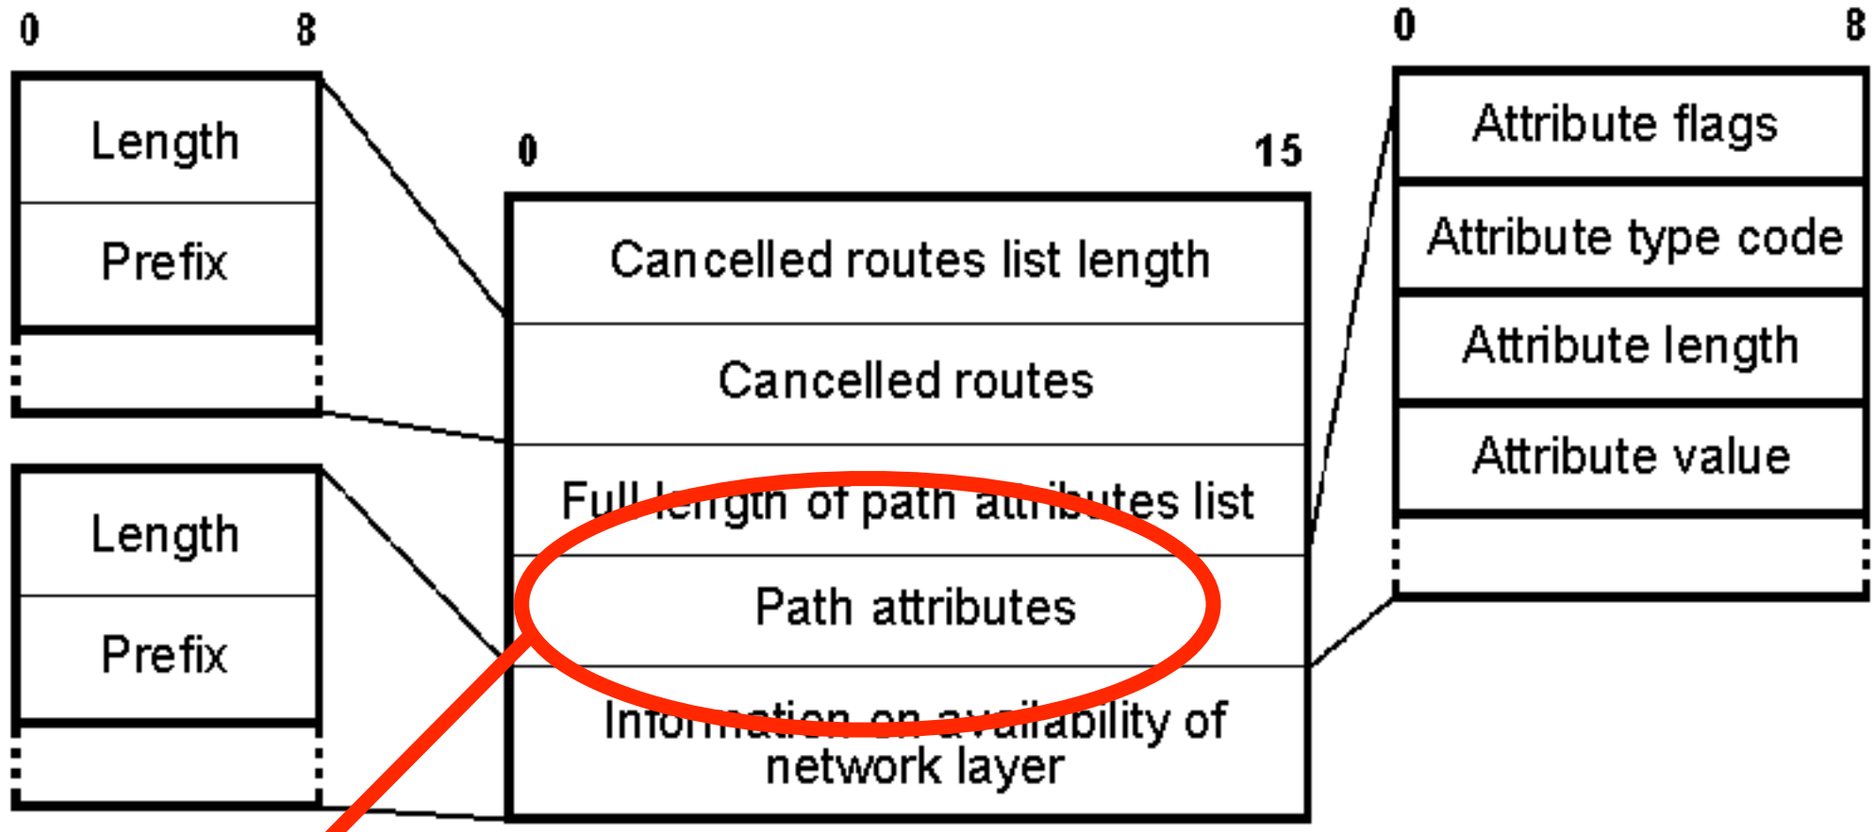


AS_PATH (loop prevention, traffic engineering)

AS_PATH gets a magic number too, and AS4_PATH contains the large numbers.

# Support is coming

| Name | Version | Notation |
|---|---|---|
| Alcatel-Lucent SR OS | >= 7.0 | asplain |
| BIRD | >= 1.0.12 | asplain |
| Brocade (Foundry) IronWare | >= 4.0.00 for the NetIron MLX and XMR, >= 2.8.00 for the BigIron RX | asdot, asdot+, asplain |
| Cisco IOS | >= 12.4(24)T, >= 12.0(32)S12 | asplain (asdot optional) |
| Cisco IOS XR | >= 3.4(1) | asdot (asplain planned for 3.9) |
| Cisco NX-OS | >= 4.0(1) | asdot (asplain planned for 4.1(3)) |
| ExtremeXOS | Need Information | Need Information |
| Juniper JUNOS | >= 9.1R1 | asplain (asdot optional) |
| Juniper JUNOSe | >= 4.1.0 | asplain |
| Force10 FTOS | >= 7.7.1.0 | asdot (asdot+, asplain optional) |
| OpenBGPD | >= 4.2, patches for 3.9 and 4.0 | asdot |
| Quagga | >= 0.99.10, patches for 0.99.6 and other versions | asplain |
| Redback SEOS | >= 2.0 | ascolon (asplain planned for end of 2009) |

http://as4.cluepon.net/index.php/Software_Support

# Whats the issue?

- December 10th 2008
- AS196629 originated 91.207.218.0/23
- AS_PATH: xx xx 35320 23456 (13 bytes)

  AS4_PATH: (65044 65057) 196629 (7 bytes)
- Confederation ASN in AS4_PATH is illegal

# What happened?

- "To prevent the possible propagation of confederation path segments outside of a confederation, the path segment types AS_CONFED_SEQUENCE and AS_CONFED_SET [RFC3065] are declared invalid for the AS4_PATH attribute."
- BGP Speakers we managed, which supported AS4, literally translated the RFC and tore down the session.
- Some networks we ran had OpenBGP edges, or OpenBGP route reflectors.
- The speakers kept flapping the sessions with their *transits* (where they were learning the route)
- **Disconnection from the internet**

# Motivation for this behaviour

- The motive is to prevent as4 upstreams transiting bad data (end user session torn down by upstream ISP)
- The reality is that networks lost their transit if bad data gets into “global bgp”

# The response?

- openbsd.misc - Claudio Jeker
- “The best thing we can do is to mark the update as ineligible so it will not propagate further and will not be used but this is a quite radical measure. On the other hand this is probably the safest way to handle this error.”
- I agree.  Drop the route, which will cause people who need to see the pfx to complain and apply pressure on the originator to fix.

# How did it leak?

- Junos introduced AS4 in 9.1R1.
- An AS with a mixed <9.1 and >9.1 network, using confederations in as4_path, updates with “dirty” transitive values can leak through egress routers running <9.1.
- If you use Junos and confeds, run >9.1R1 everywhere.

# Phew.

- We patched the affected BGP speakers.
- My pager stopped ringing.
- I went to sleep.

# Haaaaang on ….

- We suffered outages, because the BGP stack was literally interpreting the standard.
- This thread left the standard unchanged, so does this mean that other implementations are broken?

# Cisco IOS behaviour

- Installed 12.0(32)S12 on c7200vxr and singled homed it to AS15653.

```
*Jan 16 11:29:58.531: %BGP-5-ADJCHANGE: neighbor 193.239.32.2 Up
```

```
*Jan 16 11:30:02.595: %BGP-6-ASPATH: Invalid AS path (65044 65048 65062)
3.21 23456 received from 193.239.32.2: Confederation found in AS4_PATH
*Jan 16 11:30:02.595: %BGP-5-ADJCHANGE: neighbor 193.239.32.2 Down BGP
Notification sent
*Jan 16 11:30:02.595: %BGP-3-NOTIFICATION: sent to neighbor 193.239.32.2
3/1 (update malformed) 27 bytes E0111803 030000FE 140000FE 180000FE 26 FFFF
FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0050 0200 0000 3540 0101 0240 020C 0205
3D25 2114 89F8 5BA0 5BA0 4003 04C1 EF20 02E0 1118 0303 0000 FE14 0000 FE18
0000 FE26 0202 0003 0015 0000 5BA0 175B CFDA
```

**Track bug CSCsx10140**

# More behaviour

- 20 Jan 09 - Rob Shakir's post to nanog-l

- Pierfrancesco Caci was kind enough to provide me with some output from an XR box. It appears that IOS XR behaves in the same manner as Force10, and JunOS, whereby the session is not torn down, and the path is installed, albeit with a munged AS_PATH. The output below is for the prefix from 196629 which we originally analysed:

```
Path #1: Received by speaker 0
  3356 35320 3.21 23456

Given that XR box is an AS4-speaker, one would not expect to see 23456 in the
AS_PATH, the prescence of this AS seems to be a symptom of the bug (and again
occurs on Juniper/Force10).
```

For once, we're saved because vendors ignored the RFCs!

# Our recommendations

- When a path to the prefix is not already known, and an invalid AS4_PATH is received, discard the UPDATE and log.
- When a path to the pfx is already known *VIA THAT NEIGHBOUR* treat the broken UPDATE as a WITHDRAW and log.

# Note :

- AS4 / ASN32 is not inherently “bad”, in fact we need it to support the growth of the internet.
- We don’t want you to go away with the message, “do not upgrade”, we want you to take the message “**follow progress of this issue, and upgrade when safe**”.
- Safety comes when the standard is fixed