# AS4 / ASN32
# What you should know.

**Andy Davidson** - LONAP / NetSumo

This talk benefited from long discussions with, and extra research from **Rob Shakir** of GX Networks, **Jonathan Oddy** of Hostway UK, **Greg Hankins** of Force 10 and **Tom Scholl** of AT&T - without the input of these experts, this talk would not be as useful.

# Agenda

- What's the motivation for AS4
- How it works
  - Backwards compatible mode
  - "Native" AS4 mode
- AS_CONFED_SEQ in AS4_PATH
- Some interesting stuff that isn't about AS4
  - (But sort of relevant.)

# What does it mean

- Historically, AS numbers have been 2 byte numbers - 0 to 65,535
- 4 byte ASNs are 4 byte numbers - 0 to 4,294,967,295

4 byte isn't just the big numbers, it's a superset.

# What's the motivation?

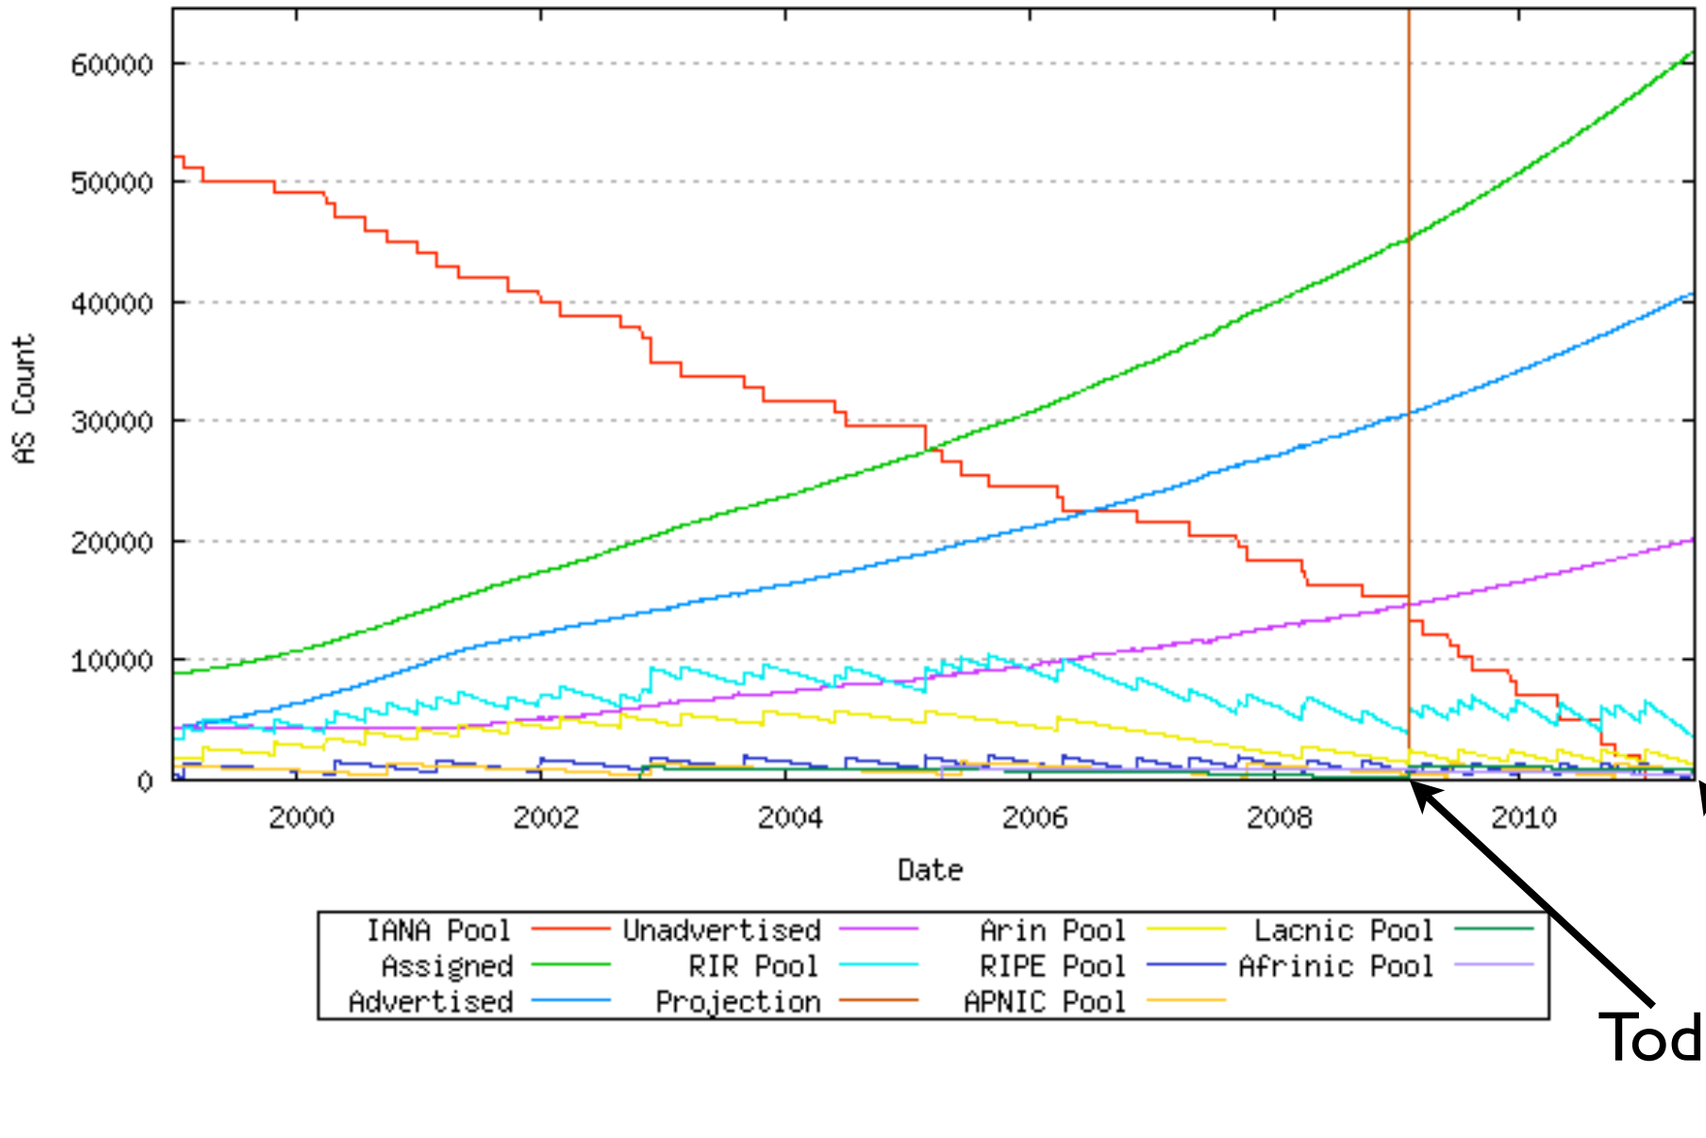


Thank you Geoff Huston
(P.S, The red lines are the scary ones)

# How to display AS4

*These ASN belong to Afrinic*

| Name | 16 bit (as2905) | 32 bit (as327676) | Notes |
|---|---|---|---|
| **asplain** | 2905 | 327676 | RIPE NCC moved to this scheme **Use this one** RFC5396 |
| asdot+ | 0.2905 | 5.1 | |
| asdot | 2905 | 5.1 | Quite popular, some, RIRs use it today |
| ascolon | 0:2905 | 5:1 | Redback |

# RIR Timescale

32 bit ASN by default

16 bit ASN on request

No distinction between 16 and 32

Unallocated 16 bit ASN LOCKED

ASN were 16 bit.

Then RFC4893 happened
- 32 bit ASN optional

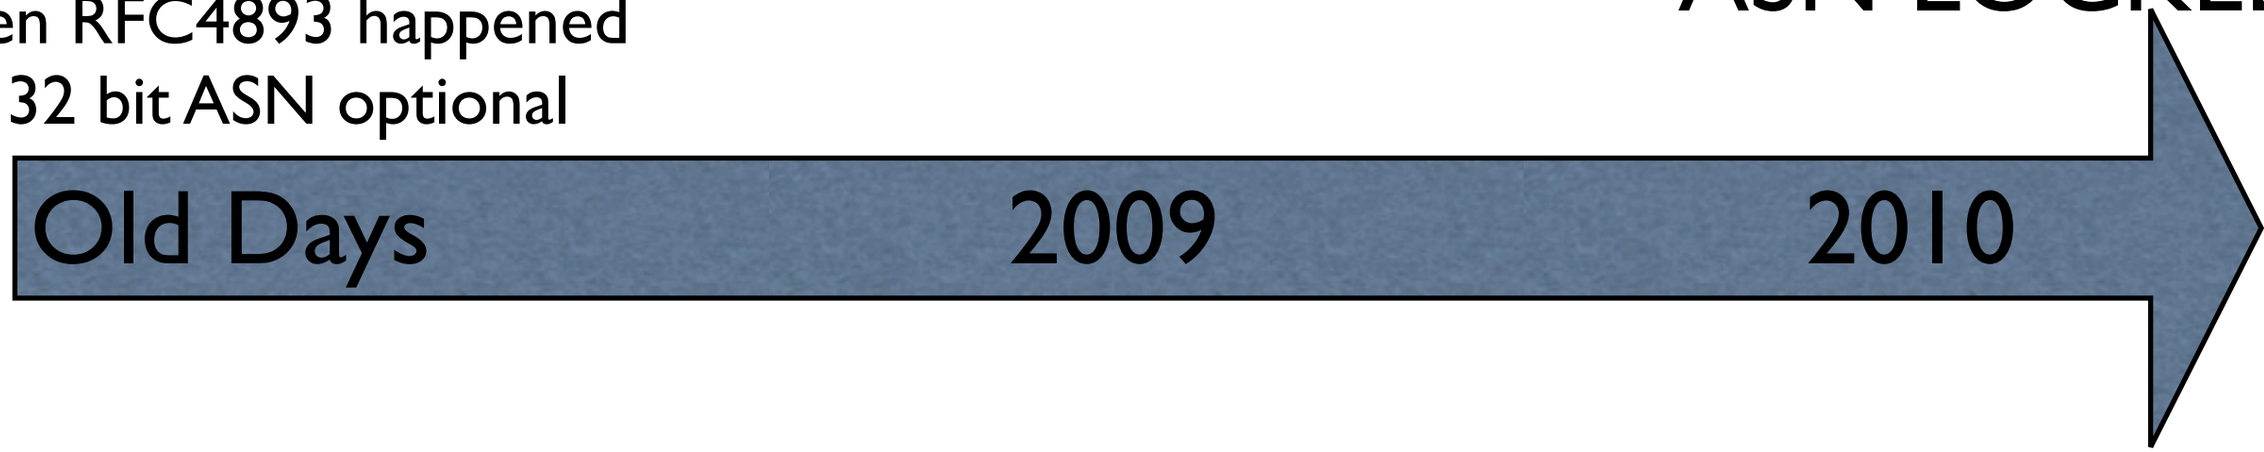

# What's the change about?

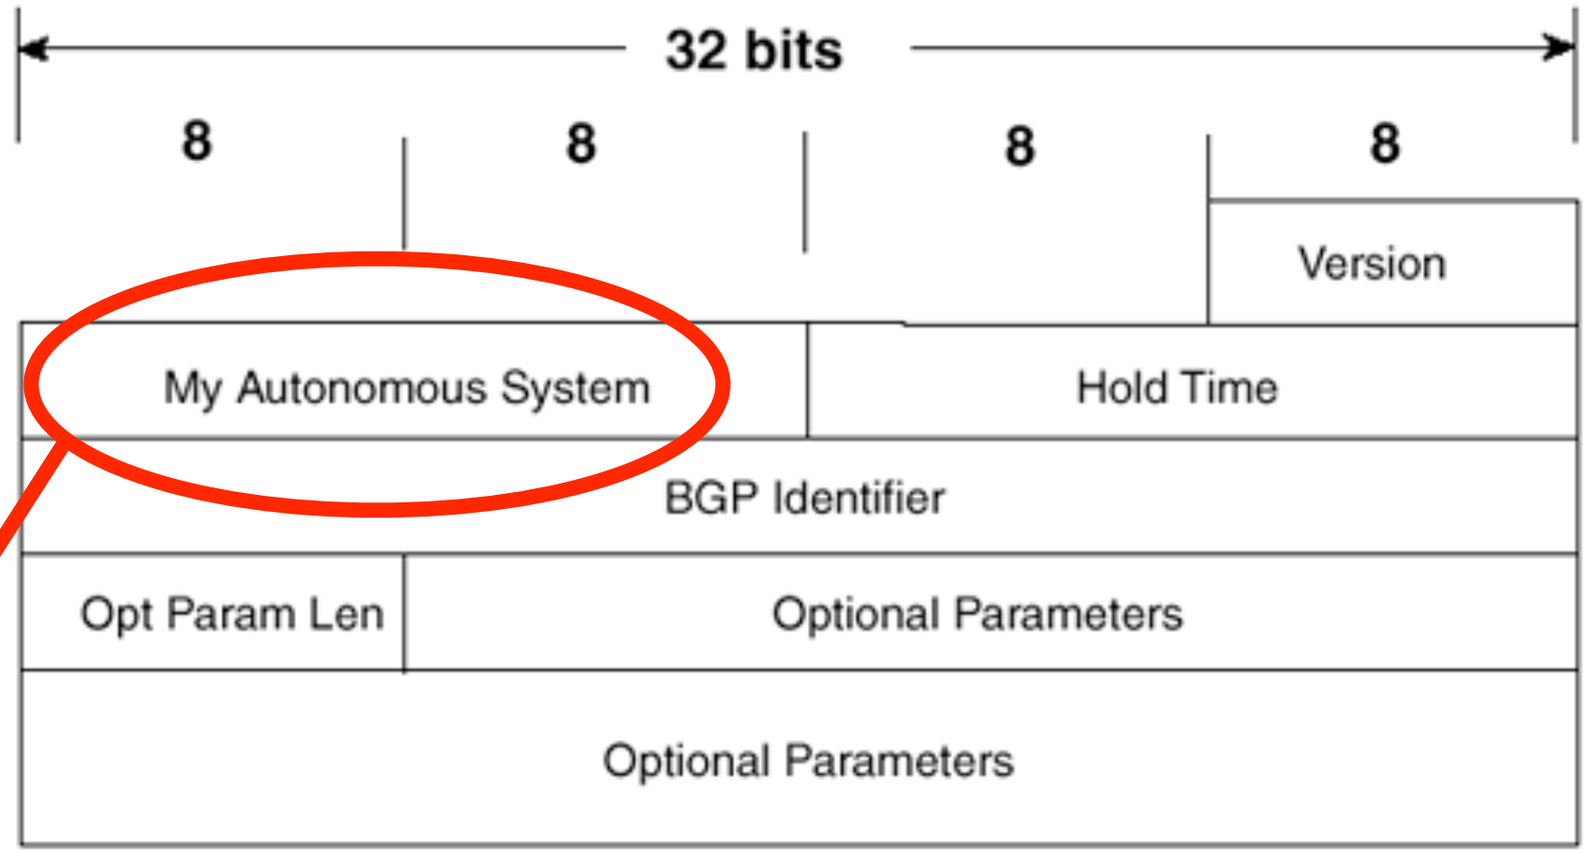


Someone's going to be AS65536 - which wont fit here

# AS4 BGP Open Packet

(It's the same)

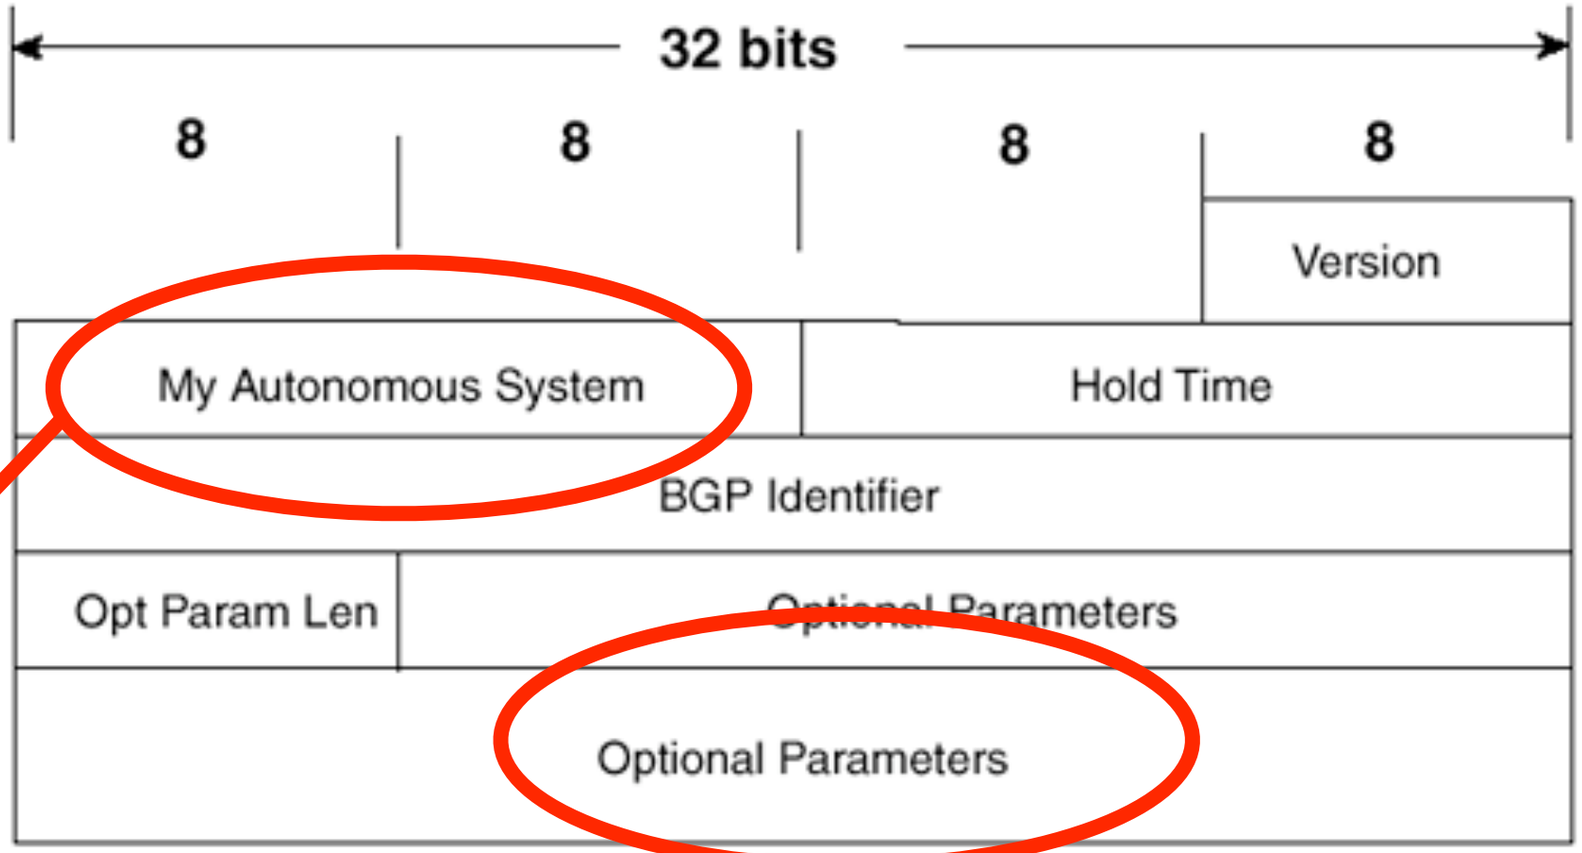


AS_TRANS
Magic Number
AS23456

My proper ASN lives in a capability value, inside here

# ASN16 to ASN32

- RFC4893 implements a new capability ("I speak ASN32") and carries the 32 bit ASN as a capability value.
- AS4_PATH is a new optional, transitive attribute in UPDATE which carries 32-bit portion of the AS_PATH.
- AS 23456 reserved for ASN32<-->ASN16 sessions (AS_TRANS)
- Both implementations co-exist, ASN32 originated prefixes in table today.

# Native ASN32

The new attributes, AS4_PATH and AS4_AGGREGATOR SHOULD NOT be carried in the UPDATE messages between NEW BGP peers. A NEW BGP speaker that receives the AS4_PATH and AS4_AGGREGATOR path attributes in an UPDATE message from a NEW BGP speaker SHOULD discard these path attributes and continue processing the UPDATE message.

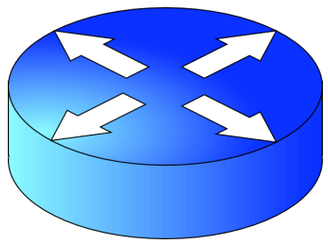

AS222222
Originates 1.0.0.0/8

AS333333
1.0.0.0/8
AS_PATH 333333 222222

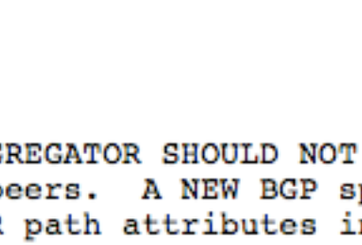

AS444444
1.0.0.0/8
AS_PATH 444444 333333 222222

# Compatibility mode

ASN16 ONLY
AS_PATH 50000 23456 23456
AS4_PATH 333333 222222 (optional, transitive)

ASN32 UPGRADED
AS_PATH 50001 50000 23456 23456
AS4_PATH 50001 50000 333333 222222
(rfc4893, 4.2.3 mandates padding)

AS50000

AS50001

AS222222

Originates 1.0.0.0/8

peers with AS23456

```
4.2.2.  Generating Updates

   When communicating with an OLD BGP speaker, a NEW speaker MUST send
   the AS path information in the AS_PATH attribute encoded with 2-octet
   AS numbers.  The NEW speaker MUST also send the AS path information
   in the AS4_PATH attribute (encoded with 4-octet AS numbers), except
   for the case where the entire AS path information is composed of 2-
   octet AS numbers only.  In this case, the NEW speaker SHOULD NOT send
   the AS4_PATH attribute.
```

AS333333

1.0.0.0/8
AS_PATH 333333 222222

AS444444

1.0.0.0/8
AS_PATH 444444 333333 222222

# Danger of comp' mode

- Everything you have needs upgrading if you make route-map/ policy based on ASN
- You could have many sessions to 23456 - all different networks
- Best Path Selection algorithm will get tricked
- NetFlow sees 32 bit world as one ASN

```
sovgw1#sh ip bgp regexp _23456$
BGP table version is 22147747, local router ID is 193.239.35.240
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 64.127.137.0/24  72.37.255.13              0    500      0 23456 i
*                   93.174.159.9                  450      0 34816 19151 18508 23456 i
*                   87.127.231.189           10             0 8468 19151 18508 23456 i
*  91.196.186.0/24  93.174.159.9                           0 34816 16150 15703 43531 23456 i
*>                  87.127.231.189           10             0 8468 15703 43531 23456 i
*> 91.207.218.0/23  72.37.255.13                500      0 23456 18508 19151 35320 23456 23456 ?
*                   93.174.159.9                            0 34816 9002 13249 13249 13249 6886 23456 ?
*                   87.127.231.189           10             0 8468 35320 23456 23456 ?
*> 91.208.44.0/24   72.37.255.13                500      0 23456 18508 25973 1299 1213 23456 i
*                   93.174.159.9                            0 34816 2914 1299 1213 23456 i
*> 93.89.236.0/22   72.37.255.13                500      0 23456 18508 25973 1299 12301 23456 i
*                   93.174.159.9                            0 34816 2914 1299 12301 23456 i
*> 94.199.136.0/21  93.174.159.9                           0 34816 2914 174 23456 i
*> 169.222.0.0/24   72.37.255.13                500      0 23456 23456 i
*                   87.127.231.189           10             0 8468 6939 7091 715 23456 i
*                   93.174.159.9                  450      0 34816 6939 7091 715 23456 i
*> 192.26.93.0      72.37.255.13                500      0 23456 18508 2914 4697 23456 i
*                   93.174.159.9                            0 34816 2914 4697 23456 i
*> 193.5.68.0/23    72.37.255.13                500      0 23456 18508 25973 6830 8758 23456 i
*                   87.127.231.189           10             0 8468 13030 8758 23456 i
*                   93.174.159.9                            0 34816 2914 13030 8758 23456 i
*> 193.31.7.0       72.37.255.13                500      0 23456 18508 19151 5539 23456 i
*                   93.174.159.9                            0 34816 2914 3549 5539 23456 i
*                   87.127.231.189           10             0 8468 5539 23456 i
*> 195.47.195.0     72.37.255.13                500      0 23456 18508 25973 6461 8495 23456 i
*                   87.127.231.189           10             0 8468 33843 8495 23456 i
*                   93.174.159.9                            0 34816 2914 1299 8495 23456 i
*> 196.1.15.0       72.37.255.13                500      0 23456 23456 i
*                   87.127.231.189           10             0 8468 3741 23456 i
*                   93.174.159.9                            0 34816 2914 174 3741 23456 i
*> 197.255.248.0/22 72.37.255.13                500      0 23456 18508 25973 3741 23456 i
*                   87.127.231.189           10             0 8468 3741 23456 i
*                   93.174.159.9                            0 34816 2914 174 3741 23456 i
*> 202.255.47.0     72.37.255.13                500      0 23456 18508 25973 2516 7667 23456 i
*                   93.174.159.9                            0 34816 2914 3356 2516 7667 23456 i
*> 205.233.128.0    72.37.255.13                500      0 23456 18508 25973 10026 7657 23754 23754 9439 23456 i
*                   93.174.159.9                            0 34816 2914 3257 10026 7657 23754 23754 9439 23456 i
```

# 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

# Communities

- Usual (new-)format is 16bits:16bits - first half usually your ASN
- New type of community proposed, "four-octet AS specific extended community" - 32bits:16bits.
- Only Quagga implements to date. (Early draft status)

# End of part one

- AS4 matters, you probably want to upgrade whether you have a 32 bit as number or not
- Upgrade your tools, train your NOC
- Lab the upgrade and track progress of AS_CONFED bugs for your software

- ASN 32 “BUG”

# AS_CONFED_SEQ in AS4_PATH - the bug

- 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.
- The speakers kept flapping the sessions with their *transits* (where they were learning the route)
- **Disconnection from the internet**
  - http://www.merit.edu/mail.archives/nanog/msg14345.html

# 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.

# Early 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**

# 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**”.

# Fix we propose

- If you see a route with illegal confed_sets in AS4, then discard (and log) it.
- If you see an UPDATE to a route with an illegal confed_set in AS4, treat it as a withdraw (and log it)
- Needs operator support on idr@ietf list.

# Cisco support

- New 12.0(32)S that fixes the issue (was due out 31st March, but not seen yet)
- First fix seen in special release 12.0(32)S11o (16th Feb 2009)
- Did what we want :

%BGP-6-MALFORMEDATTR: Malformed optional transitive attribute in (193.5.68.0/23) received from 12.122.123.3, discard confed segment found in malformed AS4_PATH;

# Some related drafts of interest

- draft-scholl-idr-advisory-00
- BGP Advisory message. “Soft Notify”
- Allows human readable communication between bgp speakers
- e.g. “Hey your update had illegal confed_set in as4 data - that sucks”
- Needs operator support.

# Another related draft

- draft-scudder-idr-optional-transitive
- Many good ideas, effectively changes the ‘malformed optional transitive attribute must lead to bgp session reset’ to ‘must NOT’ lead to session reset.
- Comments to idr@ietf

# Any Questions?

# Any Answers?

andy.davidson@netsumo.com