



BGP Bugs, Hiccups and weird stuff: Issues seen by RT-BGP Toolkit

AfPIF 2018

Martin Winter, Hurricane Electric



“Real-Time BGP Toolkit”

A quick Introduction

Traditional Looking Glass

Looking Glass

Welcome to Hurricane Electric's Network Looking Glass. The information provided by and the support of this service are on a best effort basis. These are some of our routers at core locations within our network. We also operate a public route server accessible via telnet at route-server.he.net.

Routers

Europe

- ☒ NIKHEF Amsterdam
- ☐ Itconic, Barcelona
- ☐ Telenor, Belgrade
- ☐ IPB CarrierColo Berlin
- ☐ Sitel, Bratislava
- ☐ Interxion Brussels
- ☐ NXDATA-1, Bucharest
- ☐ CE Colo Victor Hugo 1
- ☐ Interxion Copenhagen
- ☐ Interxion Dusseldorf
- ☐ Equinix Dublin (DB2)
- ☐ Interxion Frankfurt
- ☐ Equinix Geneva (GV1)
- ☐ GlobalConnect, Hamburg

- Amsterdam, NL
- Barcelona, ES
- Belgrade, RS
- Berlin, DE
- Bratislava, SK
- Zaventem, BE
- Bucharest, RO
- Budapest, HU
- Copenhagen, DK
- Düsseldorf, DE
- Dublin, IE
- Frankfurt, DE
- Geneva, CH
- Hamburg, DE

Commands

- ☐ Ping
- ☐ Traceroute
- ☐ BGP Route
- ☒ BGP Summary (IPv4)
- ☐ BGP Summary (IPv6)

Arguments

IP/Hostname:

☐ Raw output (no tables)

- ▶ **Classic Looking Glass shows view of single entity**
 - View of routing table from various location within the network of the same company

Traditional Looking Glass

core1.ams1.he.net> show ip bgp summary								
Local AS Number		6939						
Number of Neighbors Configured		972, 897 up						
Number of Routes Installed		3418476 (293988936 bytes)						
Number of Routes Advertised		120385048 (7191230 entries) (345179040 bytes)						
Number of Attribute Entries		758203 (68238270 bytes)						
Neighbor Address	ASN	State	Time	Rt:Accepted	Rt:Filtered	Rt:Sent	Rt:ToSend	
80.249.208.1	1200	ESTAB	30d18h48m	3	1	110649	0	
80.249.208.26	26496	ESTAB	25d22h39m	38	0	110649	0	
80.249.208.27	29075	ACTIV	457d 1h52m	0	0	0	110649	
80.249.208.29	8304	ESTAB	3d 0h 6m	23	0	110649	0	
80.249.208.30	8529	ESTAB	80d15h32m	302	0	110649	0	
80.249.208.32	12871	ESTAB	2d 9h29m	12	0	110649	0	
80.249.208.33	559	ESTAB	11d 0h49m	119	0	110649	0	
80.249.208.34	1103	ESTAB	3d22h59m	194	4	110649	0	
80.249.208.35	12859	ESTAB	40d11h50m	65	0	110649	0	
80.249.208.37	2686	ESTAB	38d19h12m	382	0	110649	0	
80.249.208.38	4589	ESTAB	80d15h32m	127	1	110649	0	

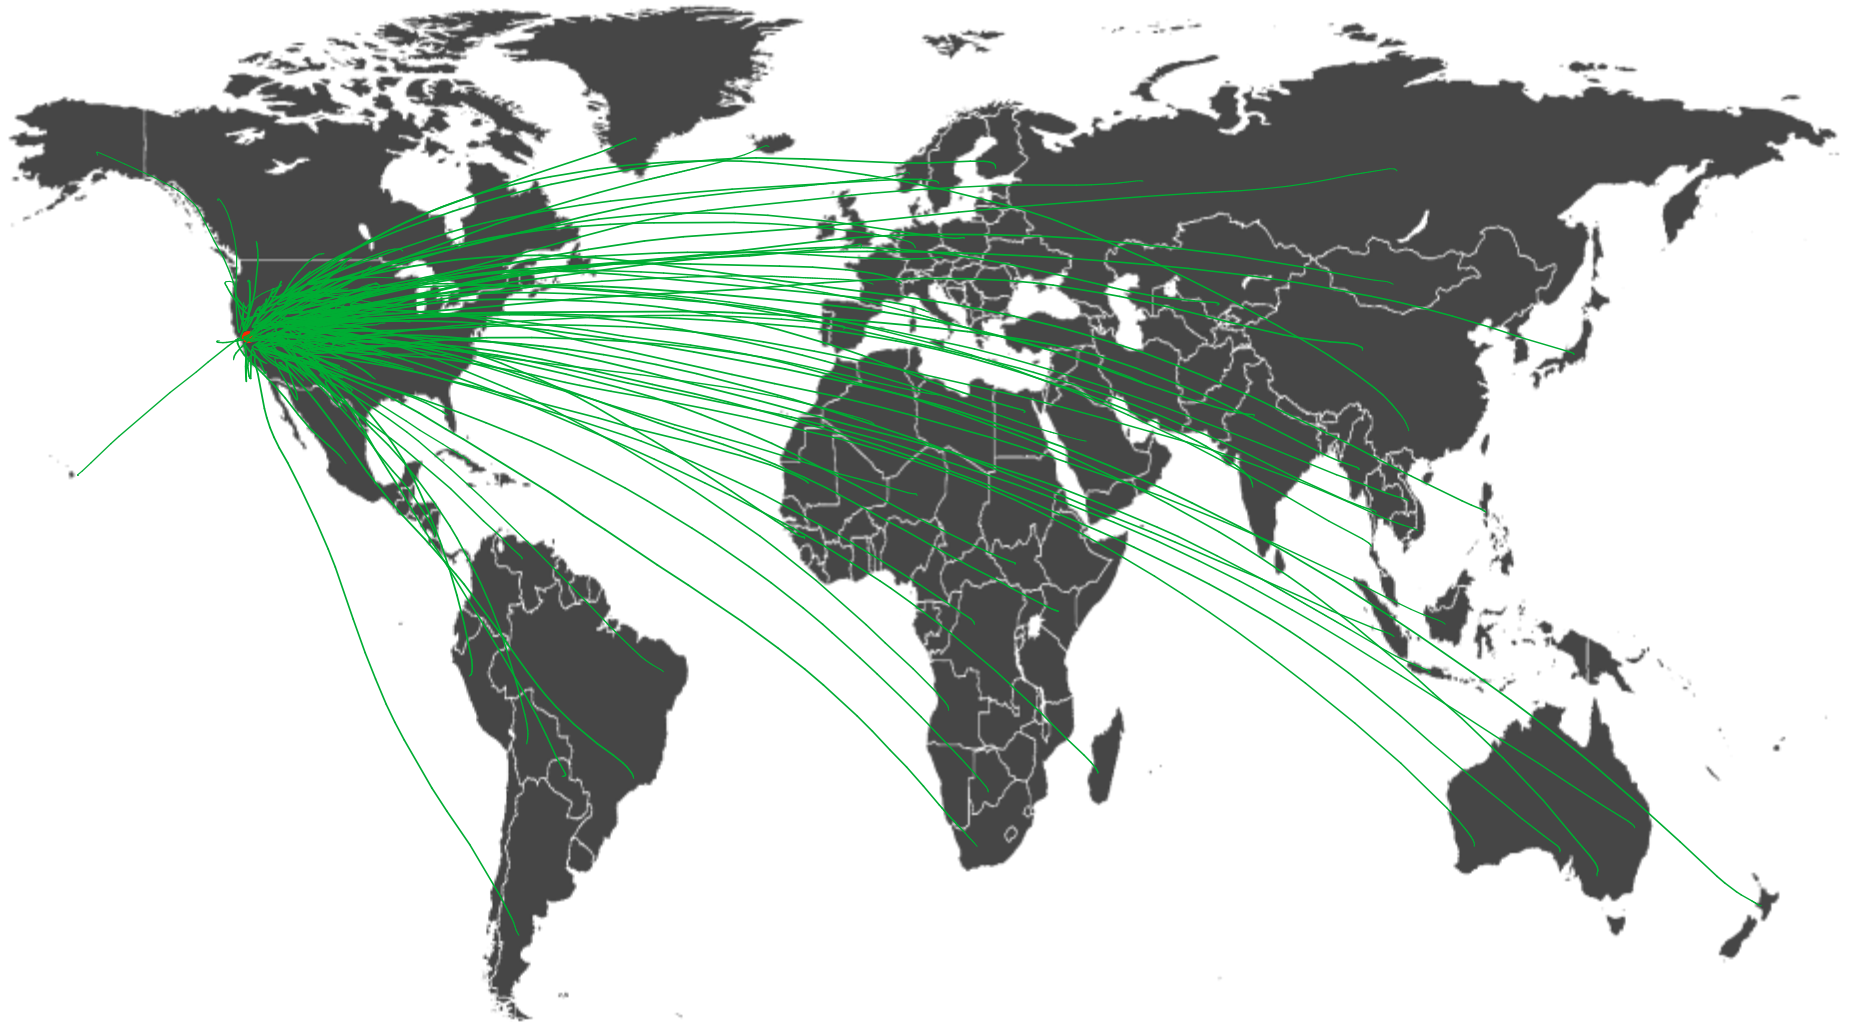
- ▶ **Classic Looking Glass mostly simple router output**
 - Showing current data from a single router at specific location.



“Real-Time BGP Toolkit”

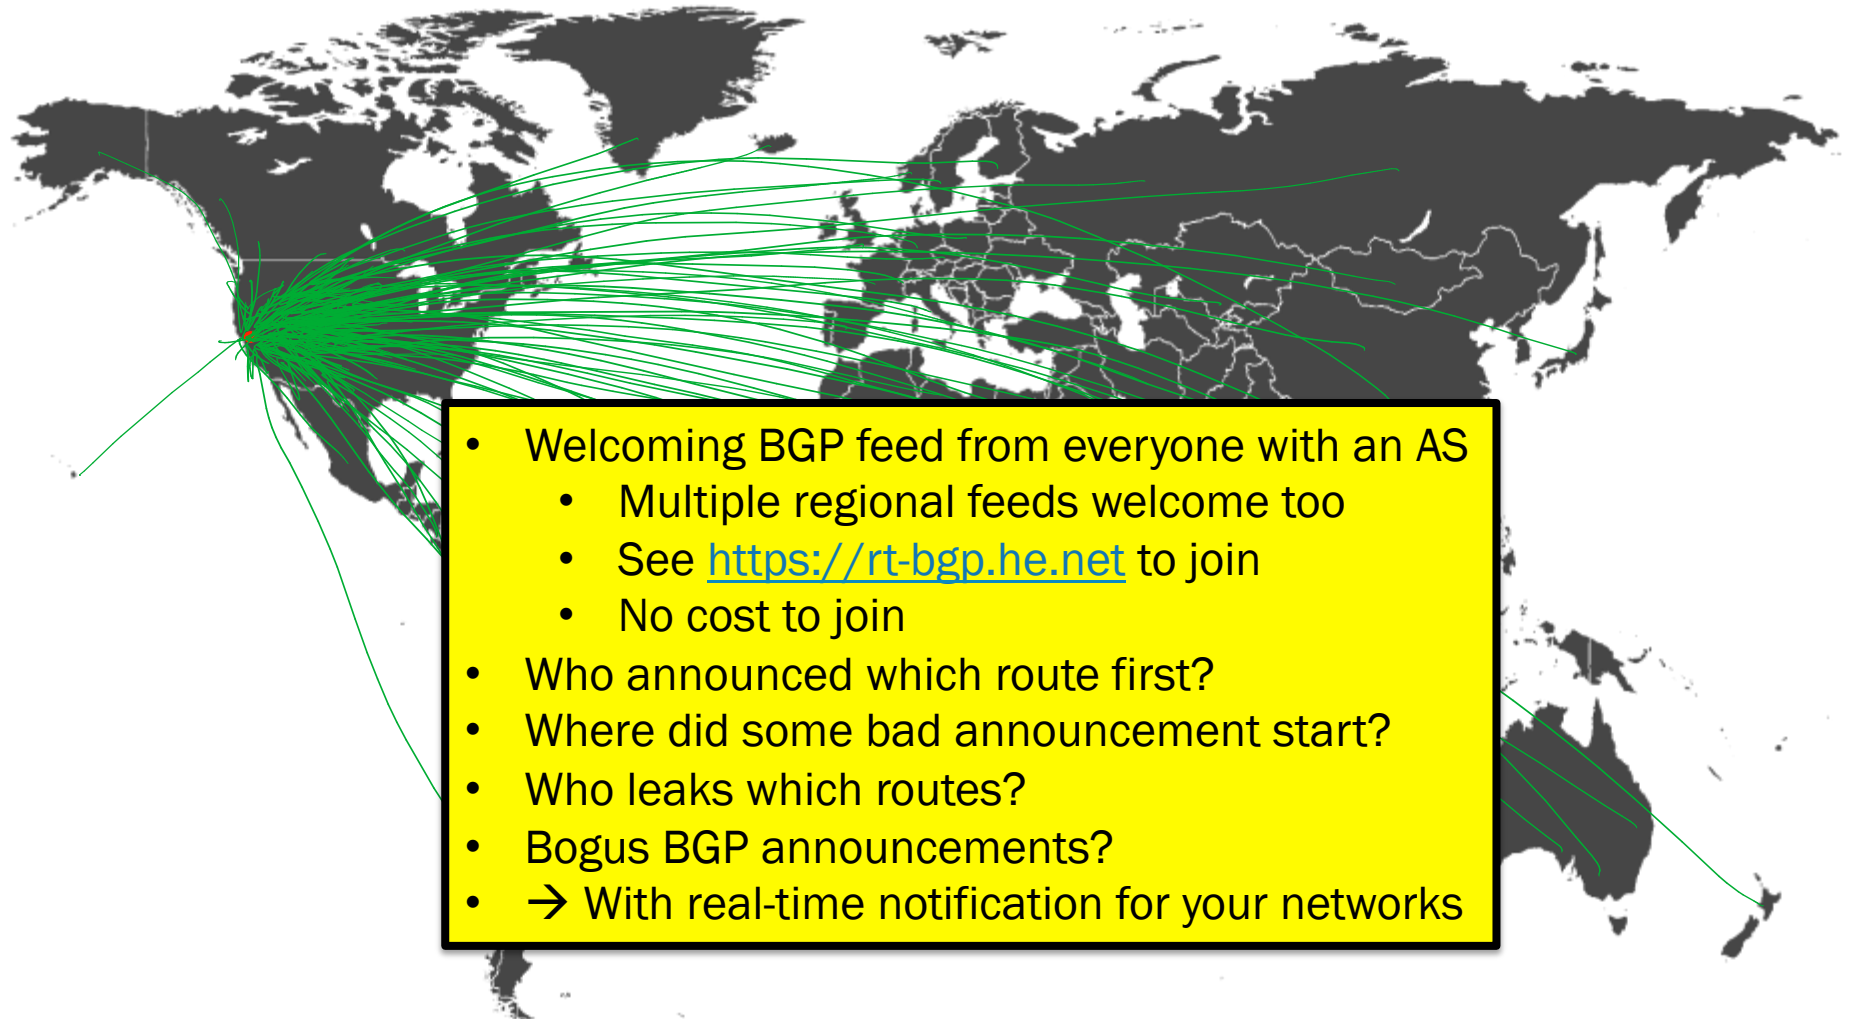
Breaking the single Entity view

Getting feeds from everywhere



Breaking the single Entity view


Getting feeds from everywhere



- Welcoming BGP feed from everyone with an AS
 - Multiple regional feeds welcome too
 - See <https://rt-bgp.he.net> to join
 - No cost to join
- Who announced which route first?
- Where did some bad announcement start?
- Who leaks which routes?
- Bogus BGP announcements?
- → With real-time notification for your networks

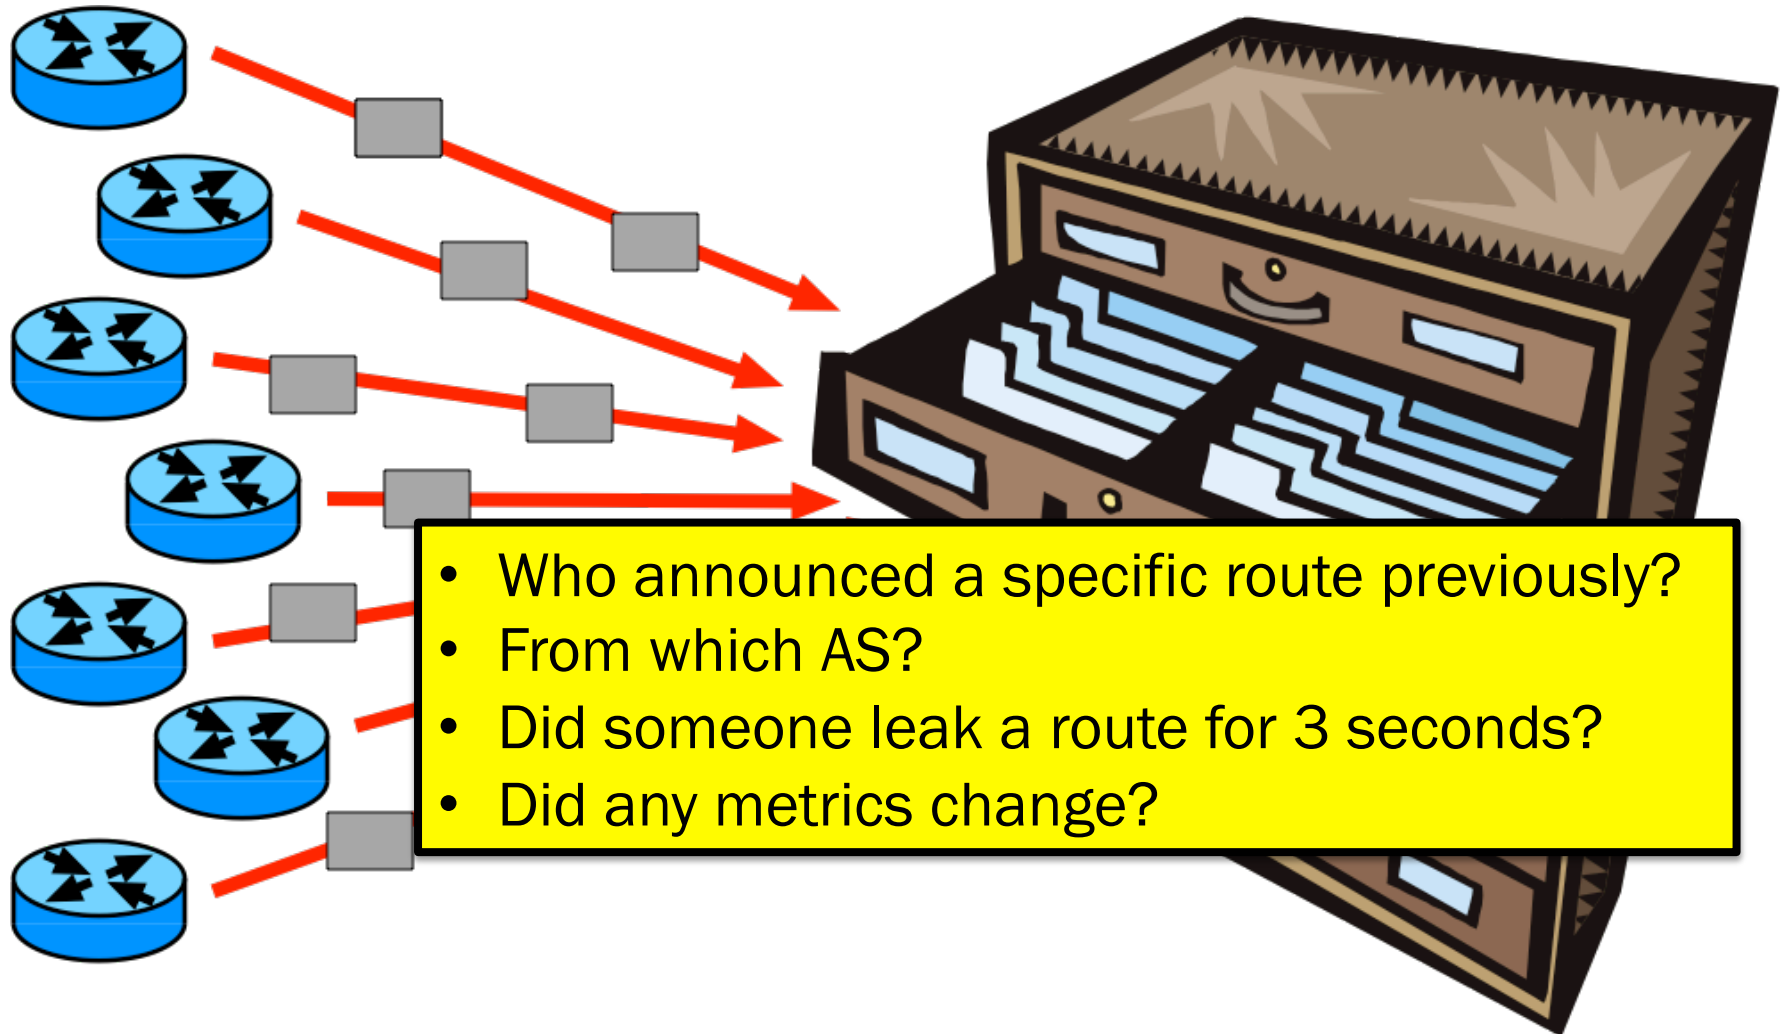
Not just Real-Time. History too

Store it all. Every single update. From every peer.



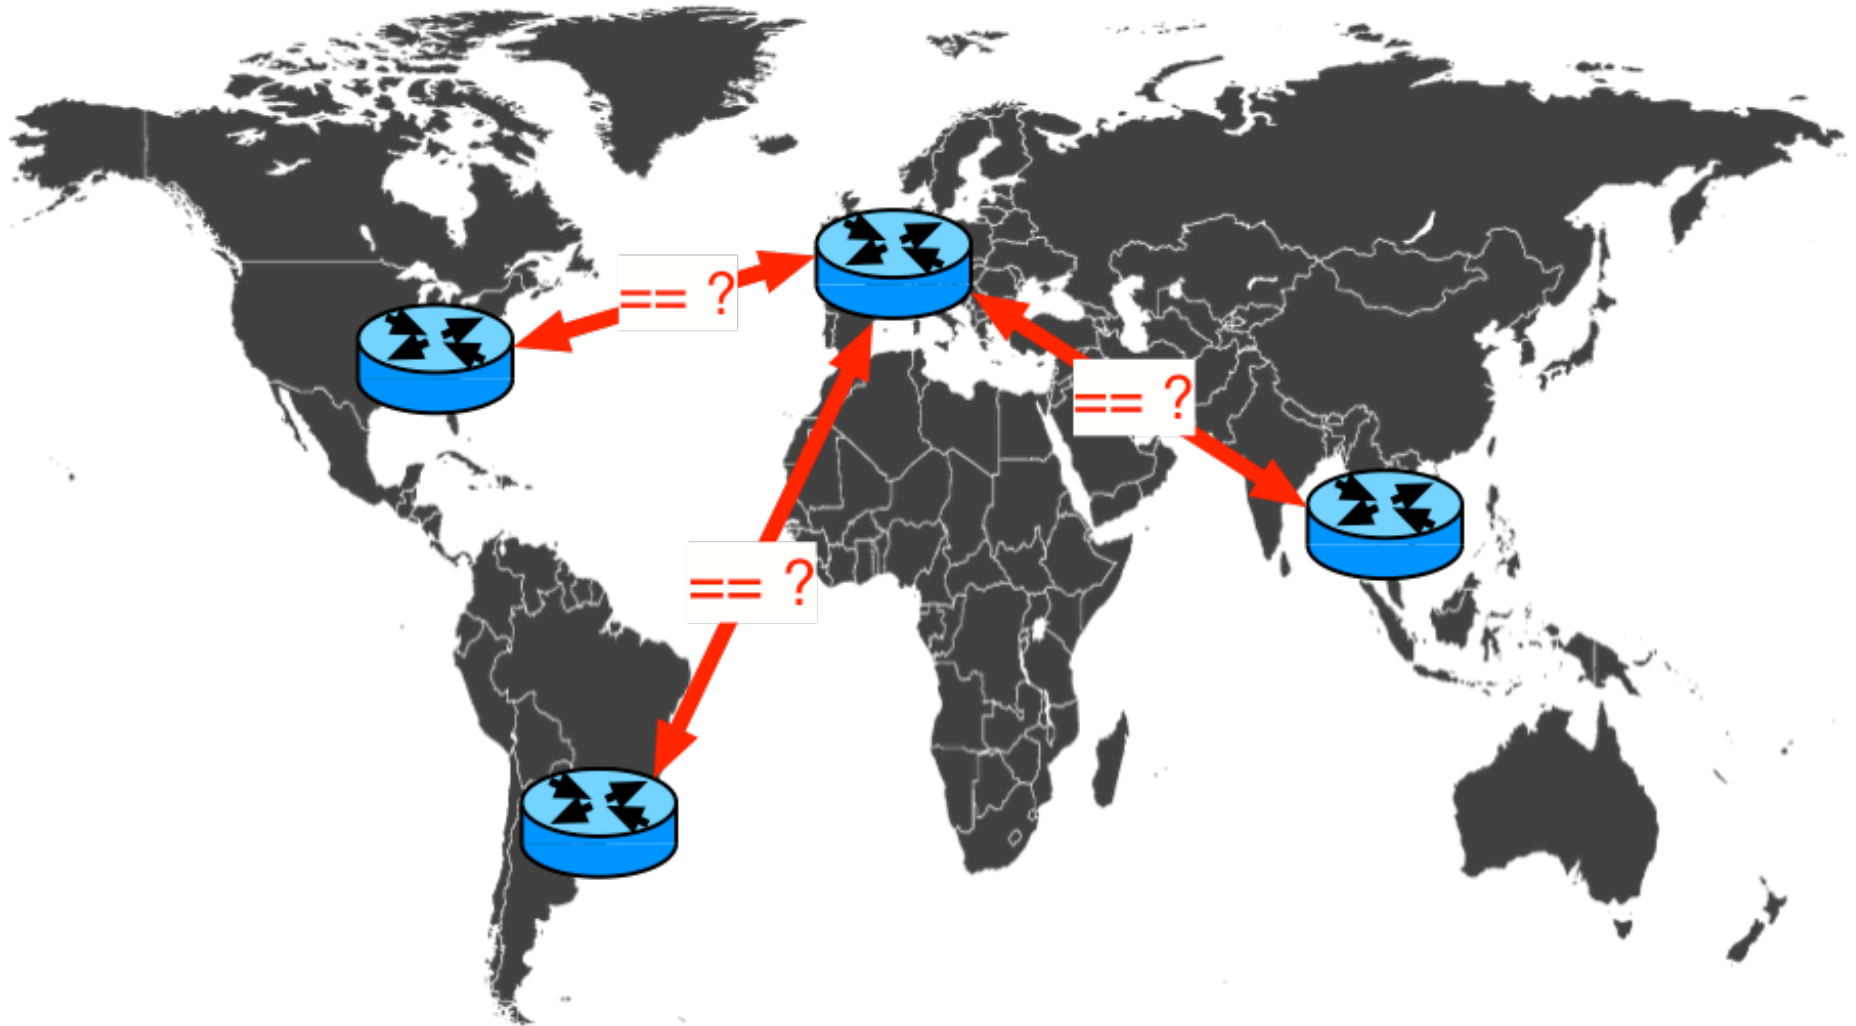
Not just Real-Time. History too

Store it all. Every single update. From every peer.



Compare the BGP feeds

Compare BGP routes between ISPs




Compare the BGP feeds

Compare BGP routes between ISPs



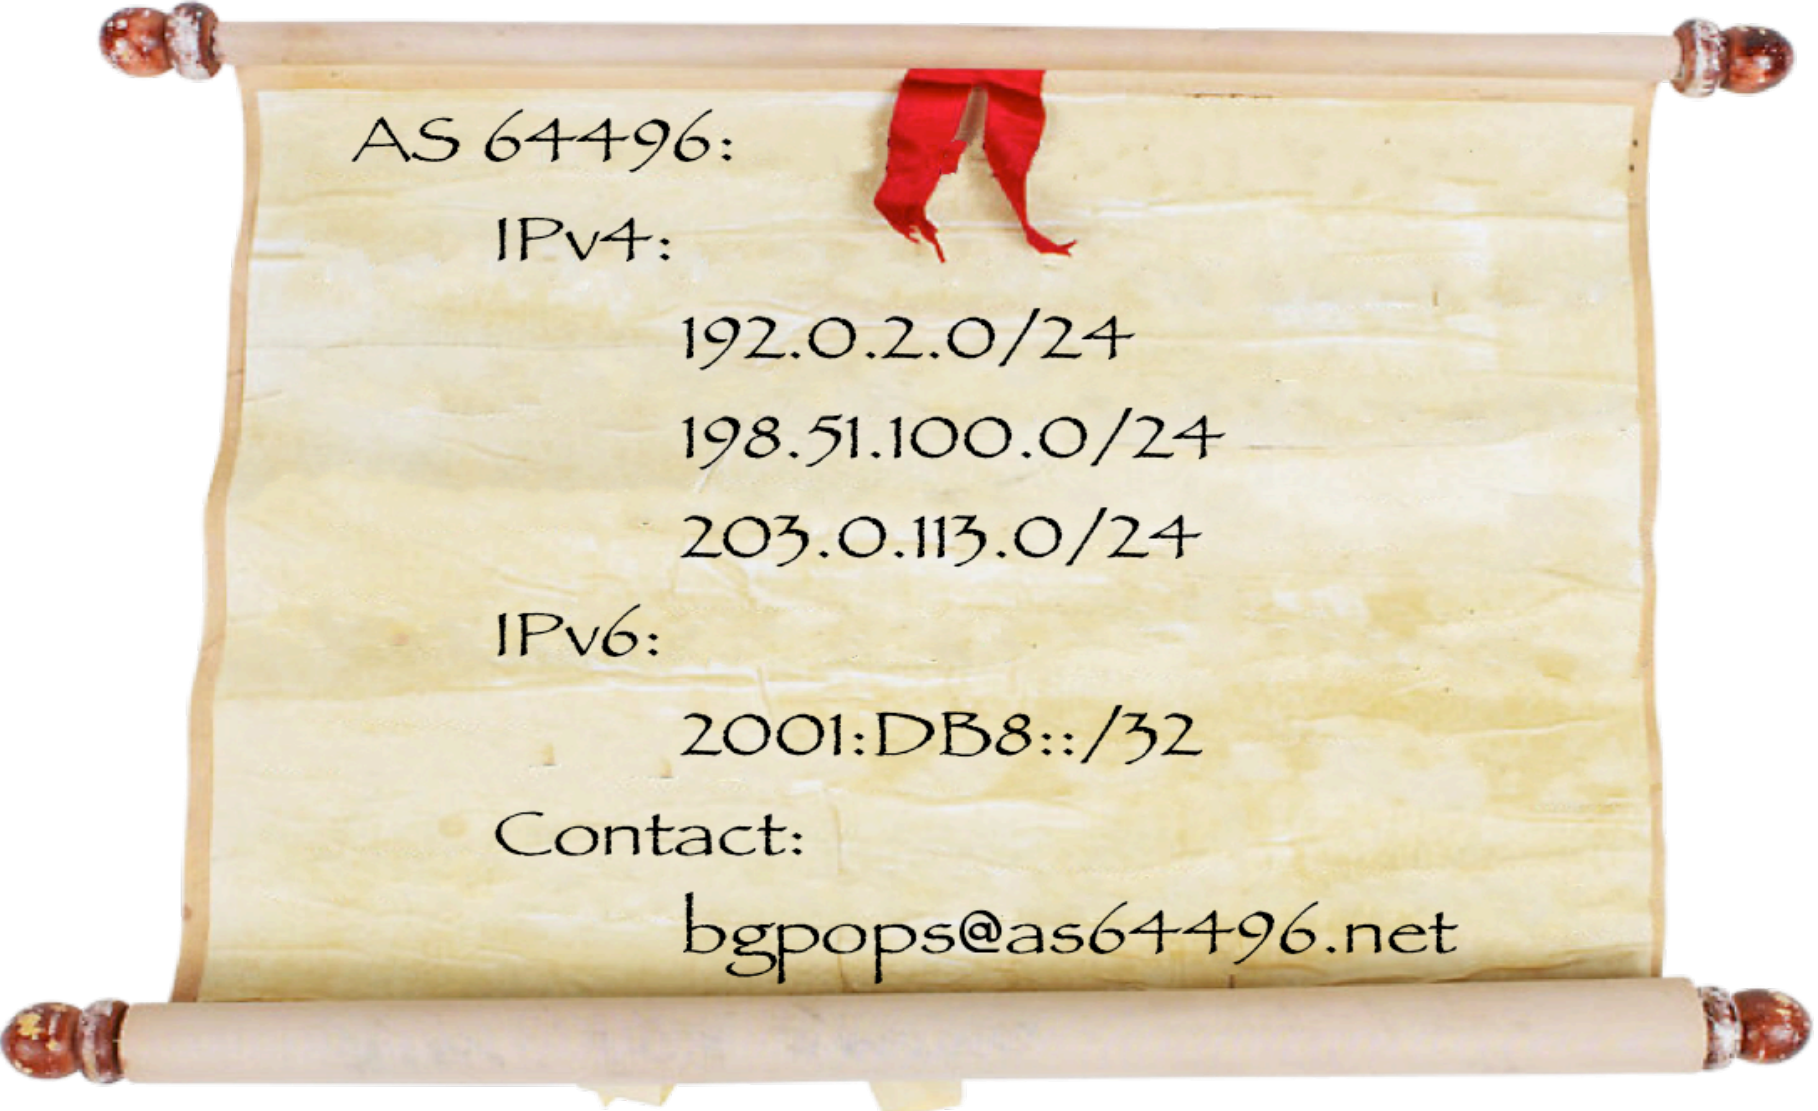
Compare the BGP feeds

Compare BGP routes between ISPs



Register routes with your AS

Get notifications on important events



AS 64496:

IPv4:

192.0.2.0/24

198.51.100.0/24

203.0.113.0/24

IPv6:

2001:DB8::/32

Contact:

bgpops@as64496.net

Register routes with your AS

Get notifications on important events

AS 64496:

IPv4:

192.0.2.0/24

Notifications for

- Routes seen announced with different source AS (Hijack?)
- More specific blocks are seen (Hijack?)
- Various bad announcements

Contact:

bgpops@as64496.net

Current (initial) features

- ▶ **Search for specific route (Current and past specified time)**
 - Show all current paths received for the route
 - Search for peers which don't have the prefix
 - Highlight different source AS for route
- ▶ **Search for specific AS number**
 - Show all routes received from the AS
- ▶ **Unassigned AS number reports**
 - Show routes sourced by unassigned AS numbers
 - Show routes with unassigned AS anywhere in AS path

Current (initial) features

- ▶ **Timeline of updates for a given prefix**
- ▶ **Hijacking detection (routes are registered with account)**
 - Detect more specific routes
- ▶ **BGPplay**

Prefix hijack report

+ additional email notifications



HURRICANE ELECTRIC
INTERNET SERVICES

QUICK LINKS

[BGPlay](#)
[BGP Toolkit Home](#)
[BGP Peer Report](#)
[BGP AS Report](#)
[BGP Prefix Report](#)
[BGP Unassigned AS Report](#)

MEMBERS

[Hijack Reports](#)
[Prefix Management](#)
[Peer Management](#)
[Profile](#)
[Activity](#)
[Log out](#)

CONTACT US

[Twitter](#)
[Facebook](#)

Hijack Report

Prefix	Infringing AS	Address	AS Path	When
204.99.140.128/25	10555	72.52.98.53	10555	2018-04-04 20:44:34.000000
204.99.140.128/25	10555	64.71.180.66	15096 10555	2018-04-04 20:44:48.000000
204.99.140.0/24	64512	64.71.180.66	15096 64512	2018-04-05 13:24:53.000000
204.99.140.0/24	10555	72.52.98.53	10555	2018-04-10 18:15:42.000000
204.99.140.0/25	10555	72.52.98.53	10555	2018-04-10 20:53:55.000000
204.99.140.0/25	10555	64.71.180.66	15096 10555	2018-04-10 20:54:09.000000

Peer comparison

Please be aware of slight update delays between peers



HURRICANE ELECTRIC
INTERNET SERVICES

QUICK LINKS

[BGPlay](#)
[BGP Toolkit Home](#)
[BGP Peer Report](#)
[BGP AS Report](#)
[BGP Prefix Report](#)
[BGP Unassigned AS Report](#)

MEMBERS

[Hijack Reports](#)
[Prefix Management](#)
[Peer Management](#)
[Profile](#)
[Activity](#)
[Log out](#)

CONTACT US

[Twitter](#)
[Facebook](#)

BGP Peer compare: 65.49.27.152 x 65.49.27.155

Summary:

☒ IPv4 ☐ IPv6

Peer 65.49.27.152 (AS6939 Hurricane Electric) has 45 different IPv4 prefixes which are missing on the other peer.

Peer 65.49.27.155 (AS6939 Hurricane Electric) has 103 different IPv4 prefixes which are missing on the other peer.

Filter

65.49.27.152 (AS6939 Hurricane Electric)

Prefix
109.161.56.0/24
113.212.125.0/24
14.188.32.0/20
14.188.64.0/20
168.205.44.0/23
170.244.211.0/24
177.107.0.0/24
177.107.12.0/24
177.107.13.0/24
177.107.14.0/24
177.107.15.0/24
177.85.193.0/24
177.85.194.0/24
184.80.188.0/23

65.49.27.155 (AS6939 Hurricane Electric)

Prefix
102.165.124.0/24
105.184.33.0/24
105.184.41.0/24
105.184.65.0/24
105.184.81.0/24
123.201.23.0/24
124.205.32.0/24
128.69.92.0/24
143.74.92.0/22
145.255.144.0/20
147.39.16.0/23
147.39.24.0/23
147.39.66.0/24
147.39.68.0/24

Unassigned AS Report

Current & past unassigned/private AS numbers announced



QUICK LINKS	
BGPlay	
BGP Toolkit Home	
BGP Peer Report	
BGP AS Report	
BGP Prefix Report	
BGP Unassigned AS Report	
MEMBERS	
Hijack Reports	
Prefix Management	
Peer Management	
Profile	
Activity	
Log out	
CONTACT US	
Twitter	
Facebook	

BGP Unassigned AS Report			
AS Path	Peer Address	Prefix	When
196624 8495 174 1273 12880 65535 205588	195.47.195.254	185.213.10.0/24	2018-05-14 14:19:38
201701 50629 174 47794 64609	185.66.192.6	94.77.232.0/24	2018-05-14 14:05:38
196624 8495 174 47794 64609	195.47.195.254	94.77.232.0/24	2018-05-14 14:05:17
15562 2914 209 721 721 721 721 27065 367 531	165.254.255.2	143.86.15.0/24	2018-05-14 13:44:21
196624 8495 174 16637 20294 64513 328292	195.47.195.254	102.164.120.0/24	2018-05-14 13:01:17
196624 8495 31025 20485 29049 199311 65456	195.47.195.254	31.171.126.0/23	2018-05-14 13:00:48
201701 50629 174 16637 20294 64513 328292	185.66.192.6	102.164.120.0/24	2018-05-14 13:00:43
201701 24482 24482 4628 9255 65010	185.66.192.6	210.24.210.0/24	2018-05-14 12:46:06
201701 24482 24482 4628 9255 65010	185.66.192.6	210.24.219.0/24	2018-05-14 12:46:06
201701 24482 24482 4628 9255 65010	185.66.192.6	210.24.209.0/24	2018-05-14 12:46:06
201701 24482 24482 4628 9255 65010	185.66.192.6	210.24.214.0/24	2018-05-14 12:46:06
201701 24482 24482 4628 9255 65010	185.66.192.6	210.24.208.0/24	2018-05-14 12:46:06
201701 49009 12731 174 16637 20294 64513 328292	185.66.192.6	102.164.120.0/24	2018-05-14 12:02:14
196624 8495 174 15699 65003 39838 39838 39838 39838	195.47.195.254	185.225.244.0/22	2018-05-14 11:45:01
201701 50629 174 15699 65003 39838 39838 39838 39838	185.66.192.6	185.225.244.0/22	2018-05-14 11:44:45
201701 50629 2914 15699 65003 39838 39838 39838	185.66.192.6	185.225.244.0/22	2018-05-14 11:44:44
201701 50629 12389 29049 199311 65456	185.66.192.6	31.171.126.0/23	2018-05-14 11:34:30
201701 50629 57463 29049 199311 65456	185.66.192.6	31.171.126.0/23	2018-05-14 11:34:29
201701 50629 174 29049 199311 65456	185.66.192.6	31.171.126.0/23	2018-05-14 11:34:16
196624 8495 174 15699 65003 39838 39838 39838 39838	195.47.195.254	185.225.244.0/22	2018-05-14 11:30:40
201701 50629 174 15699 65003 39838 39838 39838 39838	185.66.192.6	185.225.244.0/22	2018-05-14 11:30:39
201701 50629 12389 29049 199311 65456	185.66.192.6	31.171.126.0/23	2018-05-14 11:28:02
201701 50629 57463 29049 199311 65456	185.66.192.6	31.171.126.0/23	2018-05-14 11:27:59
201701 50629 174 29049 199311 65456	185.66.192.6	31.171.126.0/23	2018-05-14 11:27:49
201701 50629 174 15699 65003 39639 39639 39639 39639	185.66.192.6	185.229.216.0/22	2018-05-14 11:27:36
196624 8495 174 15699 65003 39639 39639 39639 39639	195.47.195.254	185.229.216.0/22	2018-05-14 11:27:36
196624 8495 174 15699 65003 204442 204442 204442	195.47.195.254	185.228.152.0/22	2018-05-14 11:24:00
201701 50629 174 15699 65003 39827 39827 39827 39827	185.66.192.6	185.227.92.0/22	2018-05-14 11:23:34
201701 50629 174 15699 65003 204442 204442 204442	185.66.192.6	185.228.152.0/22	2018-05-14 11:23:34

Unassigned AS Report – Prefix view

Current view of prefix



QUICK LINKS		BGP Prefix Report: 102.164.120.0/24					
BGPlay BGP Toolkit Home BGP Peer Report BGP AS Report BGP Prefix Report BGP Unassigned AS Report		Summary: Prefix announced by 20 peer(s).					
MEMBERS		When					
Hijack Reports Prefix Management Peer Management Profile Activity Log out							
CONTACT US							
Twitter Facebook							
Peer	AS Path	Origin AS					
72.52.98.53	10555 15096 6939 1299 174 16637 20294	328292					
195.47.195.254	196624 8495 174 16637 20294 64513 *	328292					
65.49.27.154	6939 1299 174 16637 20294 328292	328292					
213.144.128.184	13030 16637 20294 328292	328292					
65.49.27.157	6939 1299 174 16637 20294 328292	328292					
213.144.128.212	13030 16637 20294 328292	328292					
185.66.192.6	201701 50629 174 16637 20294 64513 *	328292					
185.142.180.67	48111 16097 16637 20294 328292	328292					
65.49.27.153	6939 1299 174 16637 20294 328292	328292					
65.49.27.156	6939 1299 174 16637 20294 328292	328292					
213.144.128.191	13030 16637 20294 328292	328292					
65.49.27.155	6939 1299 174 16637 20294 328292	328292	Not Set				
213.144.128.222	13030 16637 20294 328292	328292	Set (1)				
65.49.27.158	6939 1299 174 16637 20294 328292	328292	Not Set				
165.254.255.2	15562 2914 174 16637 20294 328292	328292	Set (0)				



BGP Bugs, Hiccups and weird stuff

Interesting things found in BGP tables

BGP Attribute 21 ??

AS_PATHLIMIT

- ▶ Anyone remember [draft-ietf-idr-as-pathlimit](#)
- ▶ Hint: Expired **11** years ago
- ▶ From the draft:

This document describes the 'AS path limit' (AS_PATHLIMIT) path attribute for BGP. This is an optional, transitive path attribute that is designed to help limit the distribution of routing information in the Internet.

By default, prefixes advertised into the BGP graph are distributed freely, and if not blocked by policy will propagate globally. This is harmful to the scalability of the routing subsystem since information that only has a local effect on routing will cause state creation throughout the default-free zone. This attribute can be attached to a particular path to limit its scope to a subset of the Internet.

BGP Attribute 21 ??

AS_PATHLIMIT

► Seen from from 3 originating AS

- 2 out of 3 answered inquiry
- Both use the same firewall vendor (Palo Alto Networks)
- Still supported in current code (as of 8.1)
 - <https://www.paloaltonetworks.com/documentation/81/pan-os/web-interface-help/network/network-virtual-routers/bgp/bgp-redist-rules-tab>

Set AS Path Limit		Enter an AS path limit for the redistributed route in the range 1-255.
-------------------	--	--

Unassigned AS-Numbers

Ignoring the leaks of private AS space



**Reserved for Documentation
and sample code [RFC5398]**

... 6461 701 703 **65817**
... 174 11845 **65610**
... 3491 11172 **65537**
... 6762 3216 **65536**
... 9829 132717 132717 132717 132717 **65542 65544**

- ▶ **Comparing the world against the IANA list of assigned AS numbers**
 - <https://www.iana.org/assignments/as-numbers/as-numbers.xhtml>

Unassigned AS-Numbers

Ignoring the leaks of private AS space



Reserved [RFC 7300]

... 45903 24176 24173 65535

... 12389 29049 12880 65535

... 32489 65535

- ▶ **Comparing the world against the IANA list of assigned AS numbers**
 - <https://www.iana.org/assignments/as-numbers/as-numbers.xhtml>

Unassigned AS-Numbers

Ignoring the leaks of private AS space

Unallocated

... 16637 36974 **327074**

... 15290 23015 23015 23015 23015 **230105**

... 9121 41933 41933 41933 41933 41933 **419333**

Typo ???

- ▶ **Comparing the world against the IANA list of assigned AS numbers**

- <https://www.iana.org/assignments/as-numbers/as-numbers.xhtml>

Unassigned AS-Numbers

Ignoring the leaks of private AS space



Unallocated

... 16637 36974 327074



MTN COTE D'IVOIRE S.A

***A lot of these announcements from unallocated AS
327074 seen behind MTN Côte d'ivoire***

Broken 4-byte AS implementation?

Is 4-byte AS support still a new thing?

- ▶ **RT-BGP uses 4-byte AS to force extended attributes**
- ▶ **One large vendor sends BGP OPEN **without** 4-byte BGP option to us (but configuration shows 4-byte AS for us correctly configured)**
- ▶ **Receiving BGP open from us (with correct 4-byte AS in BGP option) is rejected as incorrect AS**
- ▶ **Seen on ~~Foundry Brocade~~ Extreme NetIron XMR**

Broken 4-byte AS implementation

Is 4-byte AS support still a new thing?

```
isp_router# sh run | incl 64.62.153.98
neighbor 64.62.153.98 remote-as 393338
neighbor 64.62.153.98 next-hop-self
neighbor 64.62.153.98 ebgp-multihop 250
neighbor 64.62.153.98 update-source loopback 1
neighbor 64.62.153.98 remove-private-as
neighbor 64.62.153.98 filter-list 2 in
neighbor 64.62.153.98 route-map out TRANSITout
```

- ▶ **Configuration looks good...**

Broken 4-byte AS implementation

Is 4-byte AS support still a new thing?

```
isp_router# sh run | incl 64.62.153.98
neighbor 64.62.153.98 remote-as 393338
neighbor 64.62.153.98 next-hop-self
neighbor 64.62.153.98 ebgp-multihop 250
neighbor 64.62.153.98 update-source loopback 1
neighbor 64.62.153.98 remove-private-as
neighbor 64.62.153.98 filter-list 2 in
neighbor 64.62.153.98 route-map out TRANSITout
```

► Configuration looks good...

- But missing AS4 enable!

```
isp_router(config-bgp)# capability as4 enable
```

Dear Brocade (now Extreme):

Please DON'T accept config with 4-byte AS if you have the support not enabled and definitely don't try to OPEN a session to a 4-byte neighbor without having 4-byte support enabled (and then rejecting the session because of AS mismatch)

High unassigned AS number



HURRICANE ELECTRIC
INTERNET SERVICES

QUICK LINKS

[BGPlay](#)
[BGP Toolkit Home](#)
[BGP Peer Report](#)
[BGP AS Report](#)
[BGP Prefix Report](#)
[BGP Unassigned AS Report](#)
[API Information](#)

MEMBERS

[Sign up](#)
[Log in](#)

CONTACT US

[Twitter](#)
[Facebook](#)

BGP Unassigned AS Report

AS Path	Peer Address	Prefix	When
42034 9066 9145 20495 9178	2001:67c:274c:fffe::1:1	2a07:59c6:ca01::/48	2018-05-04 23:38:26
42034 9066 9145 3320 6939	2001:67c:274c:fffe::1:1	2a07:59c6:ca01::/48	2018-05-04 21:26:02
48111 16097 33891 6939 20495	185.142.180.67	2a07:59c6:ca01::/48	2018-05-04 21:25:57
15096 6939 20495 9178 23456	64.71.180.66	2a07:59c6:ca01::/48	2018-05-04 21:25:46
42034 9066 9145 20495 9178	2001:67c:274c:fffe::1:1	2a07:59c6:ca01::/48	2018-05-04 21:25:31
48111 16097 20495 9178 23456	185.142.180.67	2a07:59c6:ca01::/48	2018-05-04 21:25:26
15096 6939 20495 9178 23456	64.71.180.66	2a07:59c6:ca01::/48	2018-05-04 21:25:16
13030 126288116 16243	2001:1620:1::221	2a05:96c0::/29	2018-05-07 06:06:44
13030 126288116 16243	2001:1620:1::212	2a05:96c0::/29	2018-05-07 05:24:10
13030 1085894900 61025	2001:1620:1::221	2a03:dd40::/32	2018-05-07 06:06:43
13030 1085894900 61025	2001:1620:1::212	2a03:dd40::/32	2018-05-07 05:24:09
13030 1085894900	2001:1620:1::221	2a03:6c80::/32	2018-05-07 06:06:44
13030 1085894900	2001:1620:1::212	2a03:6c80::/32	2018-05-07 05:24:09
42034 25394 2914 174 17072	2001:67c:274c:fffe::1:1	2806:2f0:92a0::/48	2018-05-07 02:12:49
48111 16097 3257 18747 65540	185.142.180.67	2801:194::/44	2018-05-04 13:31:09
196624 8495 3257 18747 65540	195.47.195.254	2801:194::/44	2018-05-04 13:30:51
15562 2914 3257 18747 65540	2001:728:1808::2	2801:194::/44	2018-05-04 13:30:46
15562 2914 6453 18747 65540	2001:728:1808::2	2801:194::/44	2018-05-04 13:30:19
15562 2914 6453 18747 65540	2001:728:1808::2	2801:194::/44	2018-05-04 13:30:17
196624 8495 3257 18747 65540	195.47.195.254	2801:194::/44	2018-05-04 06:14:26
13030 126288116 45834	2001:1620:1::221	2407:6e00:254::/48	2018-05-07 06:06:45
13030 126288116 45834	2001:1620:1::212	2407:6e00:254::/48	2018-05-07 05:24:11

High unassigned AS number



QUICK LINKS	
BGPlay	
BGP Toolkit Home	
BGP Peer Report	
BGP AS Report	
BGP Prefix Report	
BGP Unassigned AS Report	
API Information	
MEMBERS	
Sign up	
Log in	
CONTACT US	
Twitter	
Facebook	

BGP	
AS Path	Peer
42034 9066 9145 20495 9178	2001:67c
42034 9066 9145 3320 6939	2001:67c
48111 16097 33891 6939 20495	185.1
15096 6939 20495 9178 23456	64.7
42034 9066 9145 20495 9178	2001:67c
48111 16097 20495 9178 23456	185.1
15096 6939 20495 9178 23456	64.7
13030 126288116 16243	2001:1
13030 126288116 16243	2001:1
13030 1085894900 61025	2001:1
13030 1085894900 61025	2001:1
13030 1085894900	2001:1
13030 1085894900	2001:1
42034 25394 2914 174 17072	2001:67c:274c:ffe::1:1
48111 16097 3257 18747 65540	185.142.180.67
196624 8495 3257 18747 65540	195.47.195.254
15562 2914 3257 18747 65540	2001:7
15562 2914 6453 18747 65540	2001:7
15562 2914 6453 18747 65540	2001:7
196624 8495 3257 18747 65540	195.4
13030 126288116 45834	2001:1
13030 126288116 45834	2001:1

13030 126288116 16243
13030 126288116 16243
13030 1085894900 61025
13030 1085894900 61025
13030 1085894900
13030 1085894900

13030 126288116 45834
13030 126288116 45834

High unassigned AS number

- ▶ **High number created on EBGP peer between a NetIron (Extreme) and some Juniper Router**
 - Peer is a 2-byte AS peer
 - Happened on different routers, different software versions
 - Clean up with a hard reset of the eBGP session
 - Probably bug on NetIron XMR code
- ▶ ***Still a mystery – Anyone seen this before?***
- ▶ **Check your BGP tables if you have NetIron's:**
 - `show ip bgp regex [0-9]{7}`
 - Will get routing entries with 7 or more digit AS numbers



HURRICANE ELECTRIC
INTERNET SERVICES

Try It

<https://rt-bgp.he.net>

Peer with it

AS 393338

Set up peering at <https://rt-bgp.he.net>

(create an account & login, then look for the menu option to add peering)

Contact Us

RT-BGP Toolkit
rtbgp@he.net

Martin Winter
mwinter@he.net