HURRICANE ELECTRIC
INTERNET SERVICES

Peering Security

LATNIC 31
Punta Cana Dominican Republic 2019

Walt Wollny, Director Interconnection Strategy
Hurricane Electric AS6939

Who is Walt Wollny?

- Hurricane Electric AS6939 – 4 years
  - Director Interconnection Strategy – supporting the network to reach to over 44 counties and over 210 Internet Exchanges. Focus on Global connectivity.
- Amazon AS16509 – 4 years
  - Developed IP Transit and Peering on five continents.
  - Primary focus on Japan, Singapore, Hong Kong, India, Taiwan, Philippines, Australia.
  - Over 62 new CDN sites.
- Microsoft AS8075 – 13 years
  - Developed IP Transit and Peering on four continents.
  - Primary focus on US, EU and South America.

Hurricane Electric Backbone

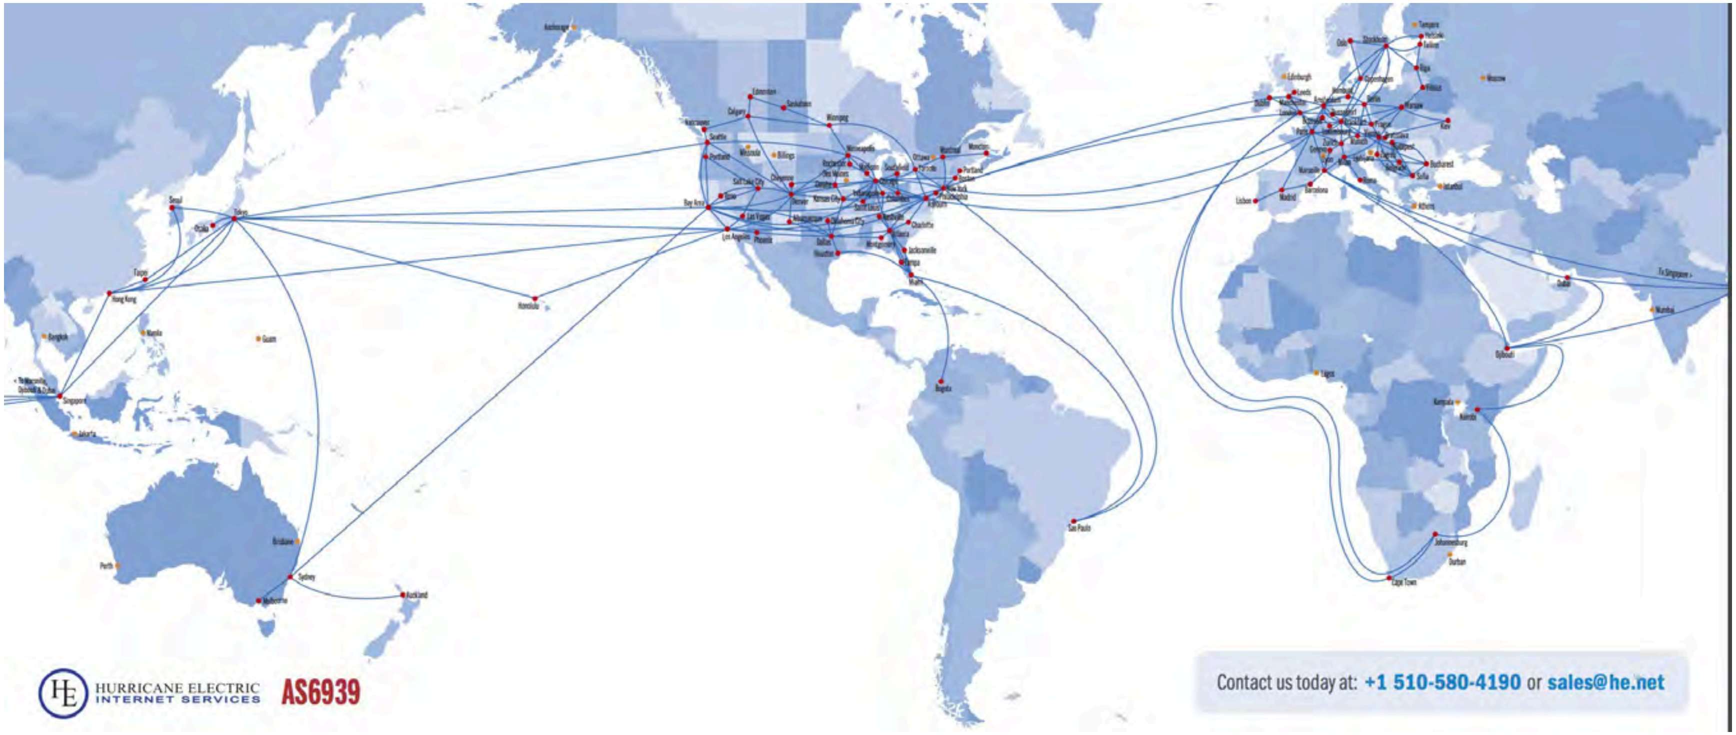

The Most Peering Exchanges

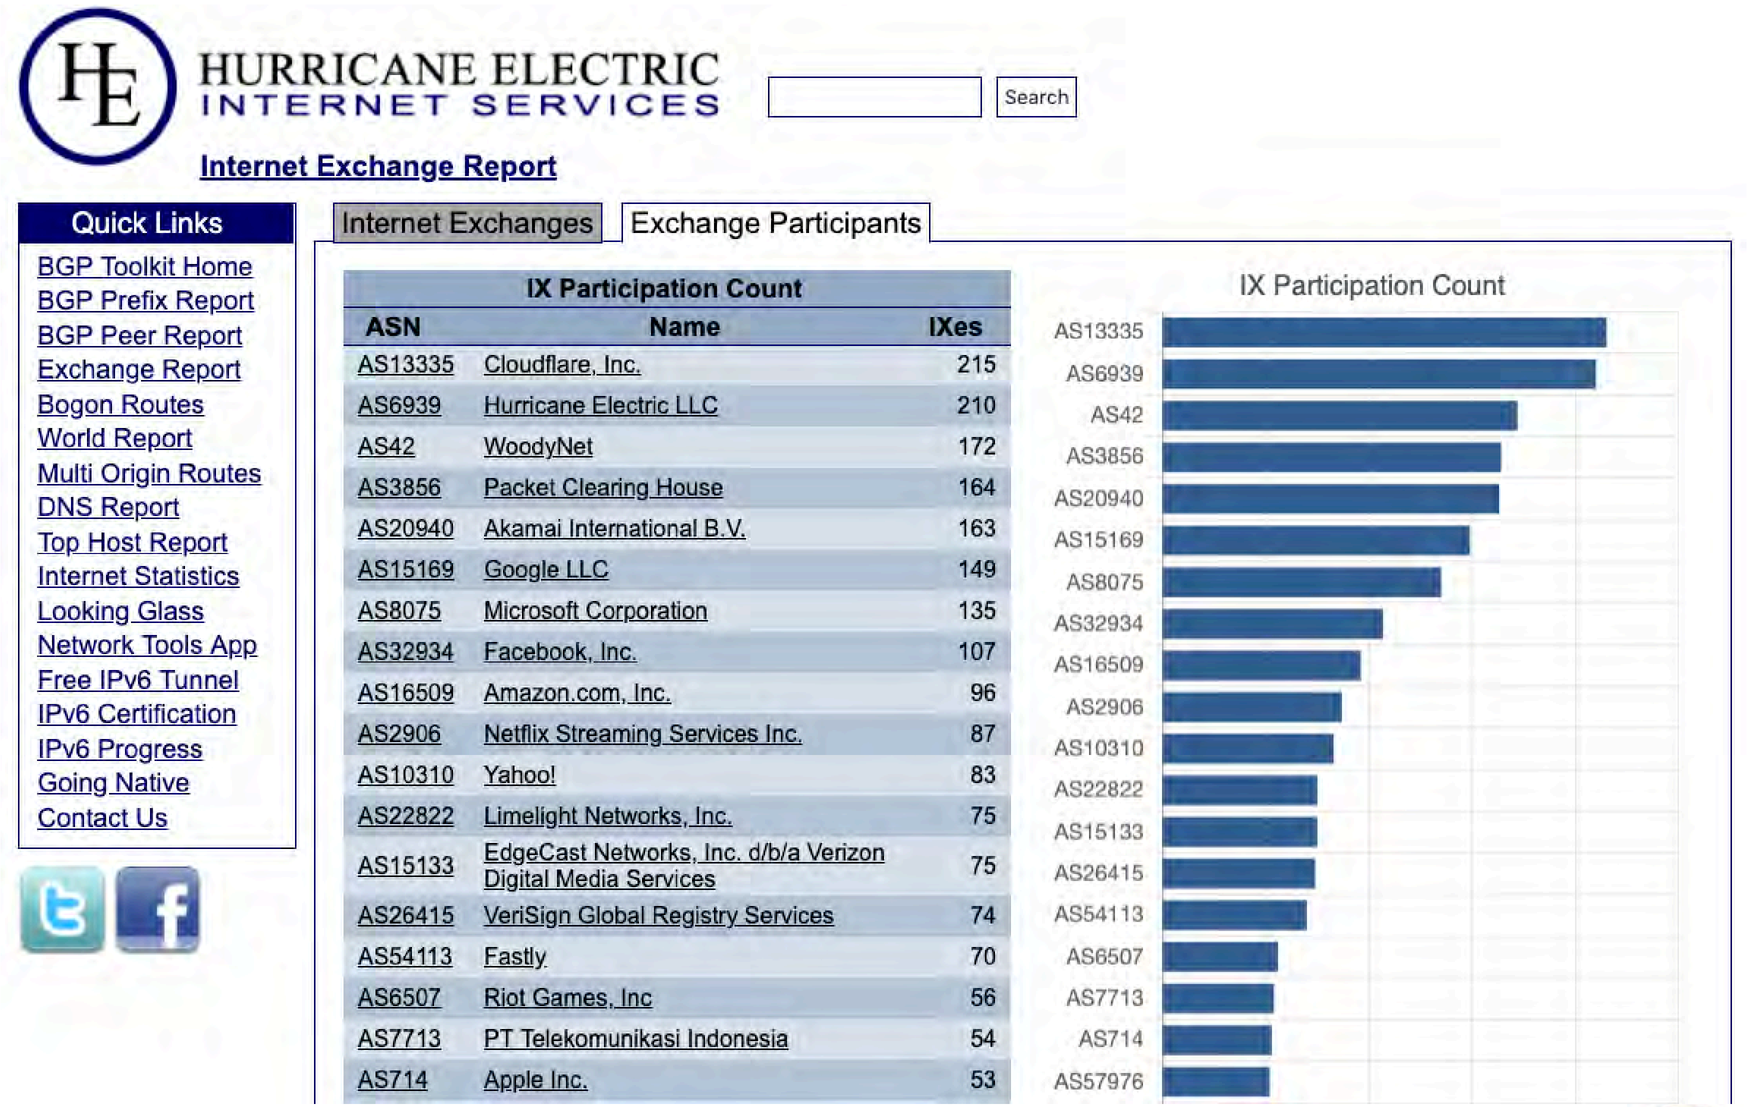

HURRICANE ELECTRIC
INTERNET SERVICES

Internet Exchange Report

Quick Links

- BGP Toolkit Home
- BGP Prefix Report
- BGP Peer Report
- Exchange Report
- Bogon Routes
- World Report
- Multi Origin Routes
- DNS Report
- Top Host Report
- Internet Statistics
- Looking Glass
- Network Tools App
- Free IPv6 Tunnel
- IPv6 Certification
- IPv6 Progress
- Going Native
- Contact Us

IX Participation Count

ASN	Name	IXes
AS13335	Cloudflare, Inc.	215
AS6939	Hurricane Electric LLC	210
AS42	WoodyNet	172
AS3856	Packet Clearing House	164
AS20940	Akamai International B.V.	163
AS15169	Google LLC	149
AS8075	Microsoft Corporation	135
AS32934	Facebook, Inc.	107
AS16509	Amazon.com, Inc.	96
AS2906	Netflix Streaming Services Inc.	87
AS10310	Yahoo!	83
AS22822	Limelight Networks, Inc.	75
AS15133	EdgeCast Networks, Inc. d/b/a Verizon Digital Media Services	75
AS26415	VeriSign Global Registry Services	74
AS54113	Fastly	70
AS6507	Riot Games, Inc	56
AS7713	PT Telekomunikasi Indonesia	54
AS714	Apple Inc.	53

Why So Many Peering Exchanges?

Why So Many Peering Exchanges?

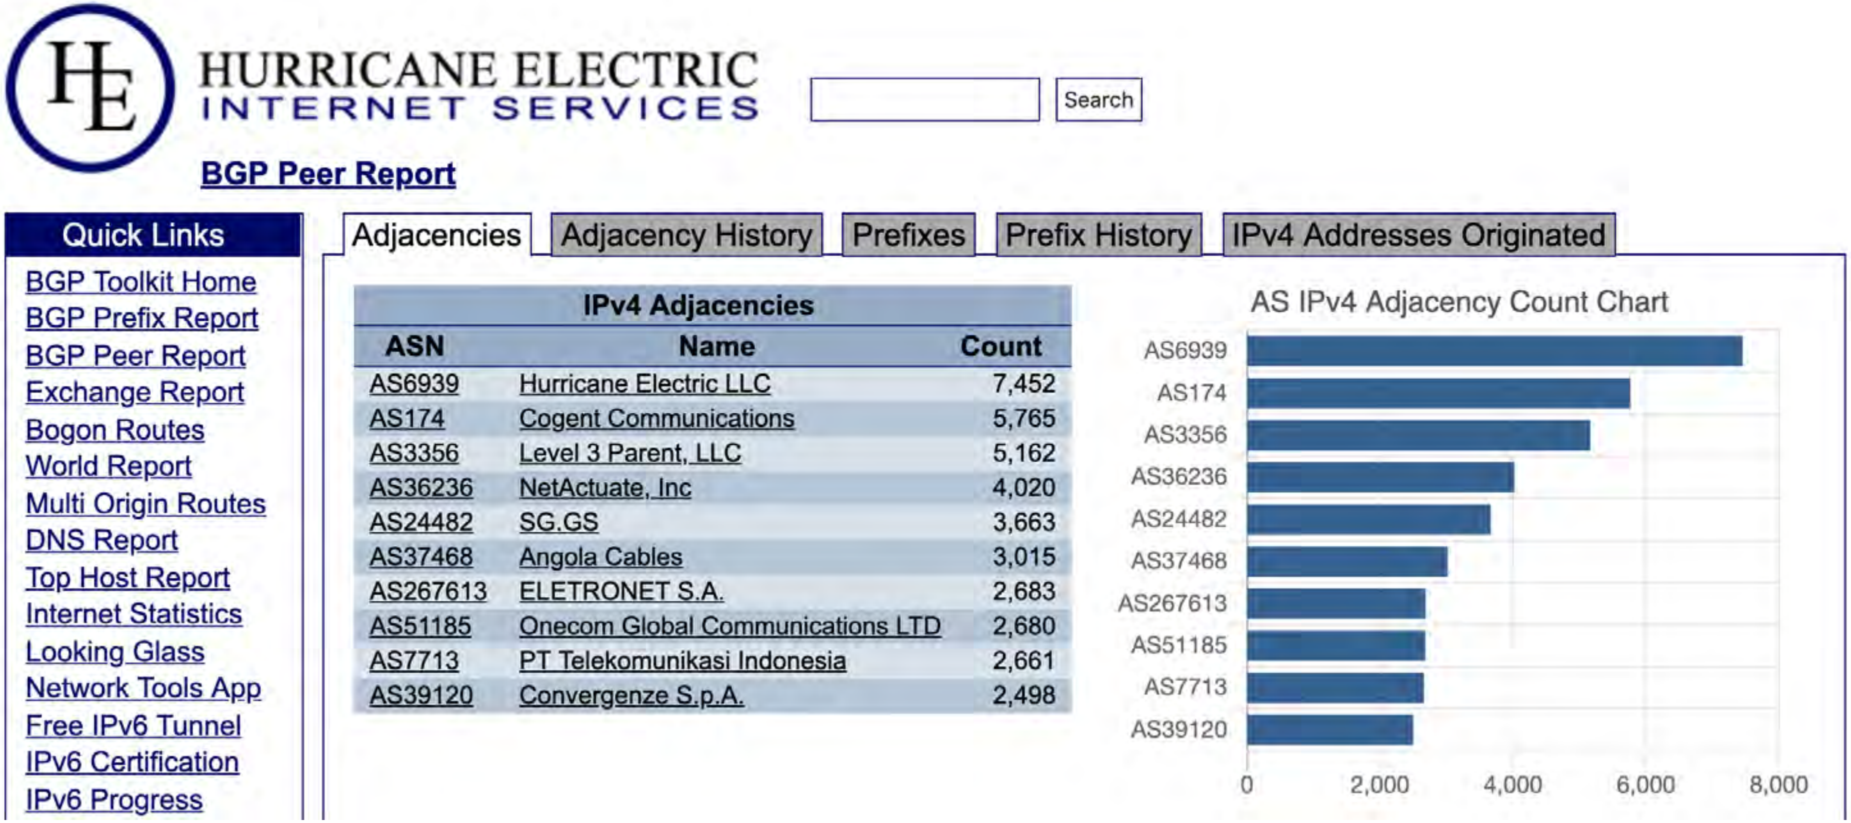

HURRICANE ELECTRIC
INTERNET SERVICES

Search

BGP Peer Report

Quick Links

- BGP Toolkit Home
- BGP Prefix Report
- BGP Peer Report
- Exchange Report
- Bogon Routes
- World Report
- Multi Origin Routes
- DNS Report
- Top Host Report
- Internet Statistics
- Looking Glass
- Network Tools App
- Free IPv6 Tunnel
- IPv6 Certification
- IPv6 Progress

Adjacencies | Adjacency History | Prefixes | Prefix History | IPv4 Addresses Originated

IPv4 Adjacencies

ASN	Name	Count
AS6939	Hurricane Electric LLC	7,452
AS174	Cogent Communications	5,765
AS3356	Level 3 Parent, LLC	5,162
AS36236	NetActuate, Inc	4,020
AS24482	SG.GS	3,663
AS37468	Angola Cables	3,015
AS267613	ELETRONET S.A.	2,683
AS51185	Onecom Global Communications LTD	2,680
AS7713	PT Telekomunikasi Indonesia	2,661
AS39120	Convergenze S.p.A.	2,498

AS IPv4 Adjacency Count Chart

What does security have to do with Peering?

A lot. Now.

Security was an afterthought, but it has become **critically** important with the increase of BGP hijacks

Some of the basics...

Basics

- Best defenses for your network?
  - Logical Port Security
  - IXP Subnet Security
  - Routing Security
  - http://routing.he.net/

Logical Port Security

- Many IXPs will post their recommended port configuration (HKIX, AMS-IX, etc).
- Don’t just connect an interface with a default configuration to an IX Port!
- Services like Proxy-ARP will disrupt the IX as well as degrade your own network.
- Most IXs allow only unicast traffic. (IPv6 multicast neighbor discovery packets are an exception.0

Logical Port Security

- Apply ACL’s to your interfaces—don’t forget to configure both IPv4 and IPv6 ACLs!
- The SIX (Seattle Internet Exchange) has a great example here.
- Your IX port is an exposed piece of your network.
- Hundreds of other networks are directly connected.
- Remove this security risk!

Logical Port Security

- Why do we care?

AMS-IX

Ticket: 341134
Subject: Instability on AMS-IX
Status: closed
Opened: 2017-06-20 16:04:56 +0200
Type: unscheduled
Scope: AMS-IX NL
Start: 2017-06-20 15:20:00 +0200
CLOSED 2017-06-21 16:54:10 +0200:

Total impact time – 1 hour 34 mins

Root cause human error

The instability was caused due to a hardware issue on the customer's NIC and due to proxy-arp being enabled after the port passed the testing phase and was moved to production.

BBIX Tokyo

Occurred time:	2018/5/16 17:28 JST
Corresponded time:	2018/5/16 17:48 JST
Recovered time:	2018/5/16 18:10 JST
Affected area:	BBIX Tokyo IX service

Total impact time – 39 mins

Root cause human error

Arp proxy response(= proxy arp) became effective when we changed the subnet mask on our monitoring router

IXP Subnet

- Your IX Port is a target for DDoS Attacks!
- Applying the best security practices will help limit the exposure.

IXP Subnet

- The IXP is responsible for protecting the infrastructure.
- The IX LAN is not your IP space and should not be routed.
- Checking this...

IXP Subnet

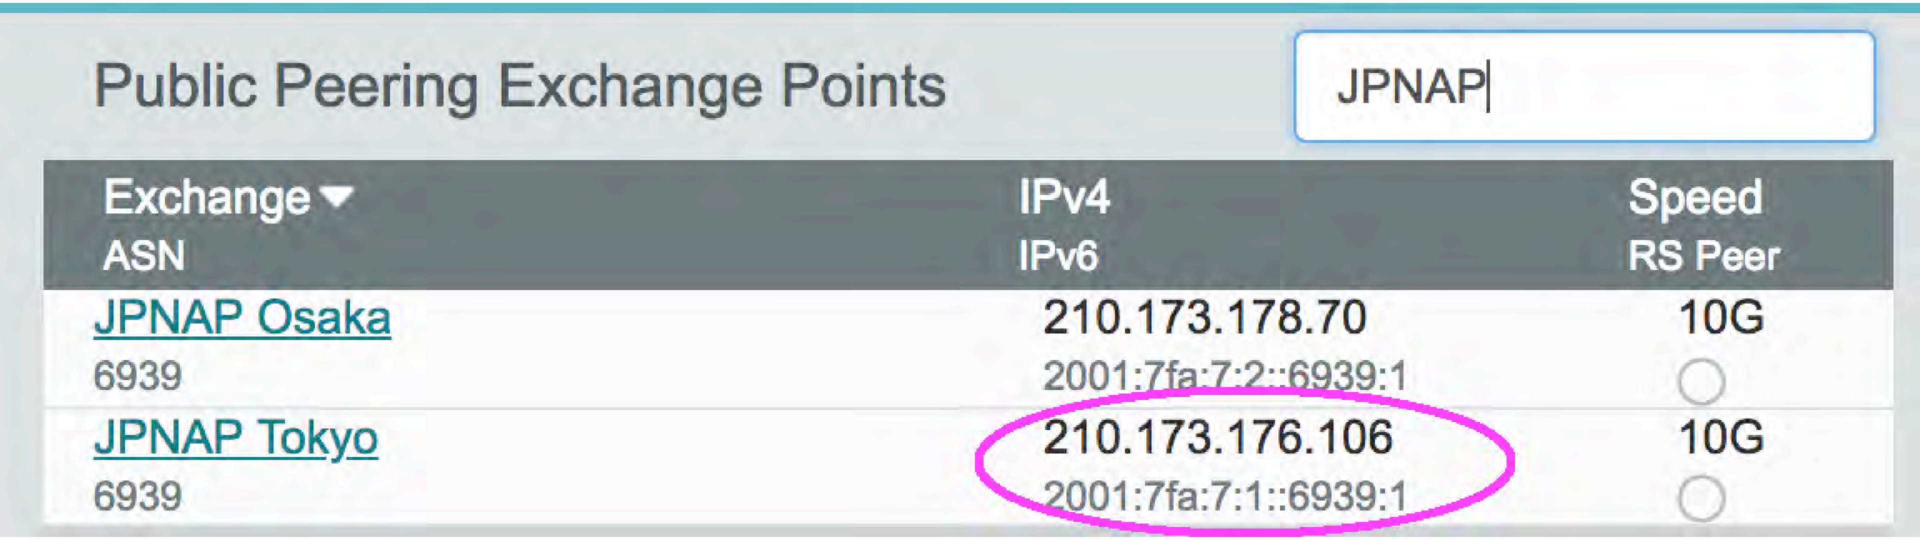

Public Peering Exchange Points

Exchange ▾ ASN	IPv4 IPv6	Speed RS Peer
JPNAP Osaka 6939	210.173.178.70 2001:7fa:7:2::6939:1	10G ○
JPNAP Tokyo 6939	210.173.176.106 2001:7fa:7:1::6939:1	10G ○

IXP Subnet

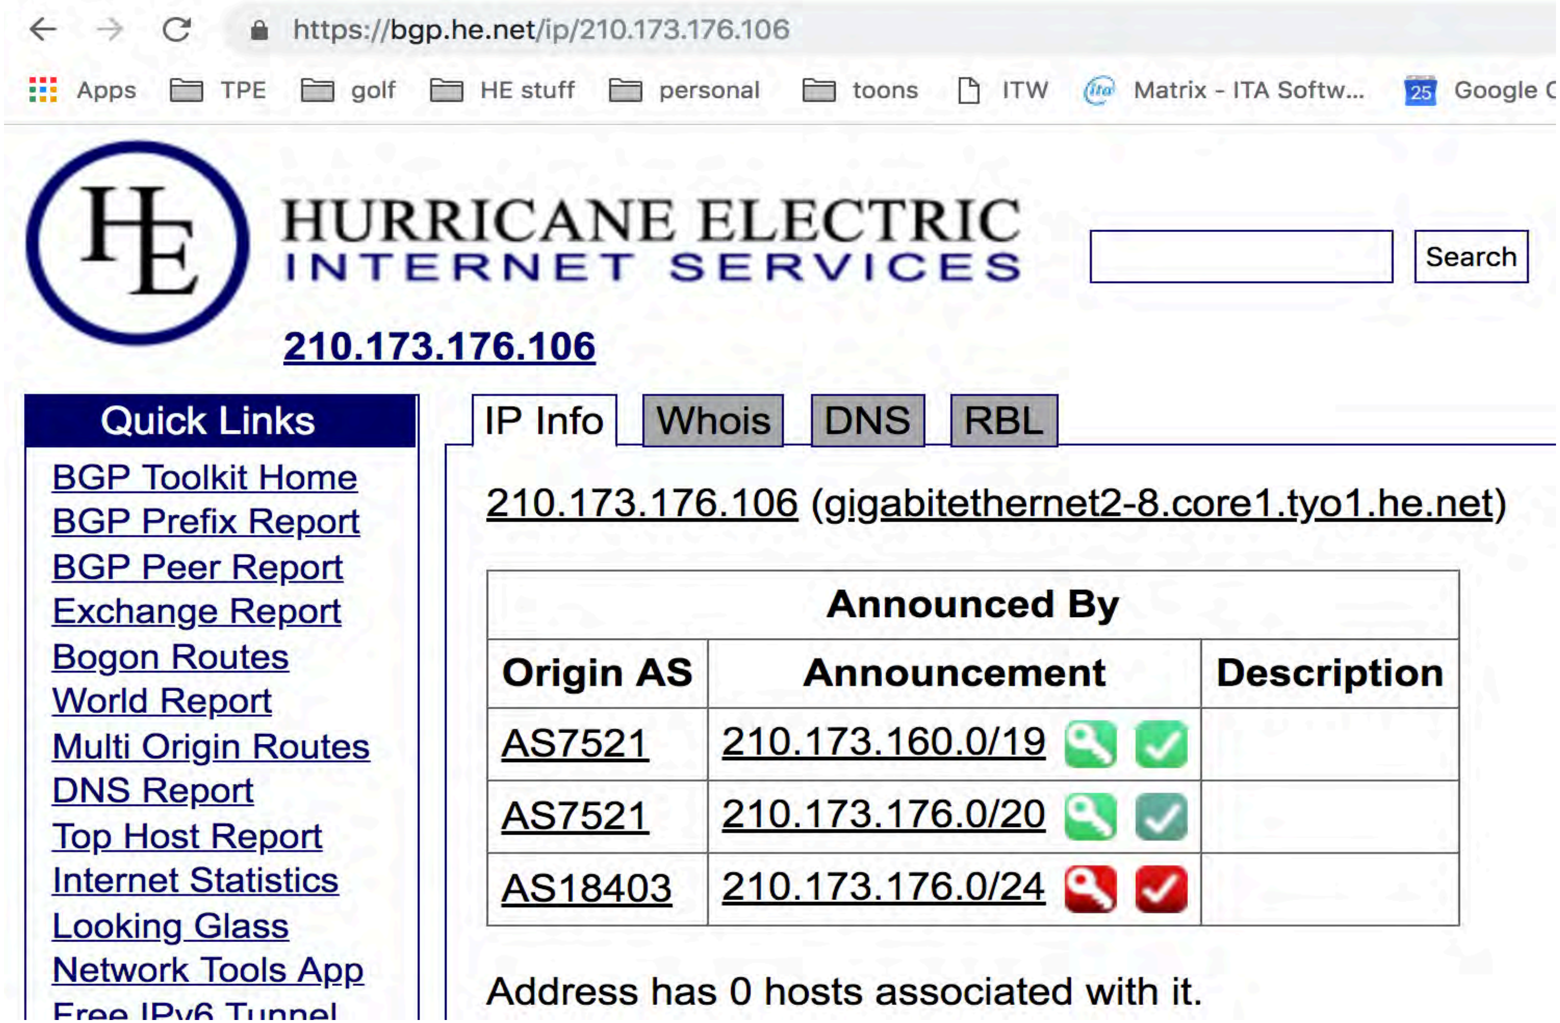

210.173.176.106

Quick Links

- BGP Toolkit Home
- BGP Prefix Report
- BGP Peer Report
- Exchange Report
- Bogon Routes
- World Report
- Multi Origin Routes
- DNS Report
- Top Host Report
- Internet Statistics
- Looking Glass
- Network Tools App
- Free IPv6 Tunnel

IP Info | Whois | DNS | RBL

210.173.176.106 (gigabitethernet2-8.core1.tyo1.he.net)

Announced By		
Origin AS	Announcement	Description
AS7521	210.173.160.0/19 🔍 ✓	
AS7521	210.173.176.0/20 🔍 ✓	
AS18403	210.173.176.0/24 🔍 ✓	

Address has 0 hosts associated with it.

IXP Subnet

Oceania

```
CC Exchange              Speed   IPv4            IPv6
-- -------------------- ------- --------------- -----------------------
AU Equinix Melbourne    10GE    183.177.61.28   2001:de8:6:1::6939:1
AU Equinix Sydney       10GE    45.127.173.24   2001:de8:6::6939:1
AU NSW-IX Sydney        10GE    218.100.52.249  2001:7fa:11:4:0:1b1b:0:1
AU VIC-IX Melbourne     10GE    218.100.78.108  2001:7fa:11:1:0:1b1b:0:1
AU MegaIX Melbourne     10GE    103.26.71.122   2001:dea:0:30::7a
AU MegaIX Sydney        10GE    103.26.68.236   2001:dea:0:10::ec
NZ APE                  10GE    192.203.154.197 2001:7fa:4:c0cb::9ac5
NZ AKL-IX               10GE    43.243.21.17    2001:7fa:11:6:0:1b1b:0:1
NZ MegaIX Auckland      10GE    43.243.22.82    2001:dea:0:40::52
```

IXP Subnet

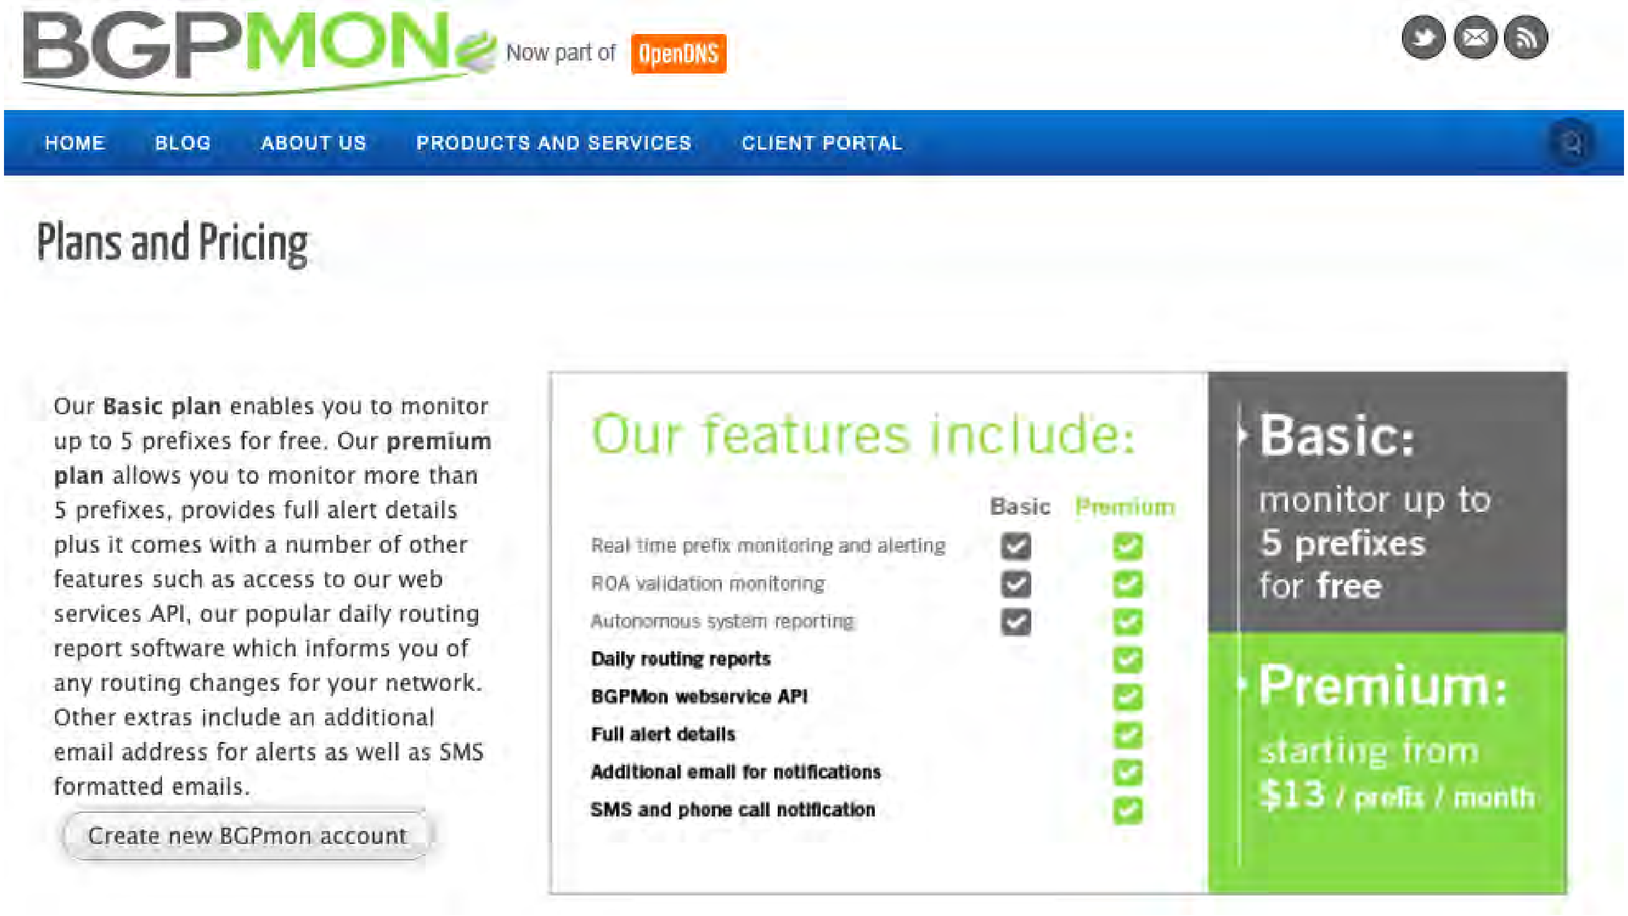

Plans and Pricing

Our **Basic plan** enables you to monitor up to 5 prefixes for free. Our **premium plan** allows you to monitor more than 5 prefixes, provides full alert details plus it comes with a number of other features such as access to our web services API, our popular daily routing report software which informs you of any routing changes for your network. Other extras include an additional email address for alerts as well as SMS formatted emails.

Our features include:

	Basic	Premium
Real time prefix monitoring and alerting	✓	✓
ROA validation monitoring	✓	✓
Autonomous system reporting	✓	✓
Daily routing reports		✓
BGPMon webservice API		✓
Full alert details		✓
Additional email for notifications		✓
SMS and phone call notification		✓

Basic: monitor up to **5 prefixes** for **free**

Premium: starting from **$13 / prefix / month**

This product is now end of life in March 2020

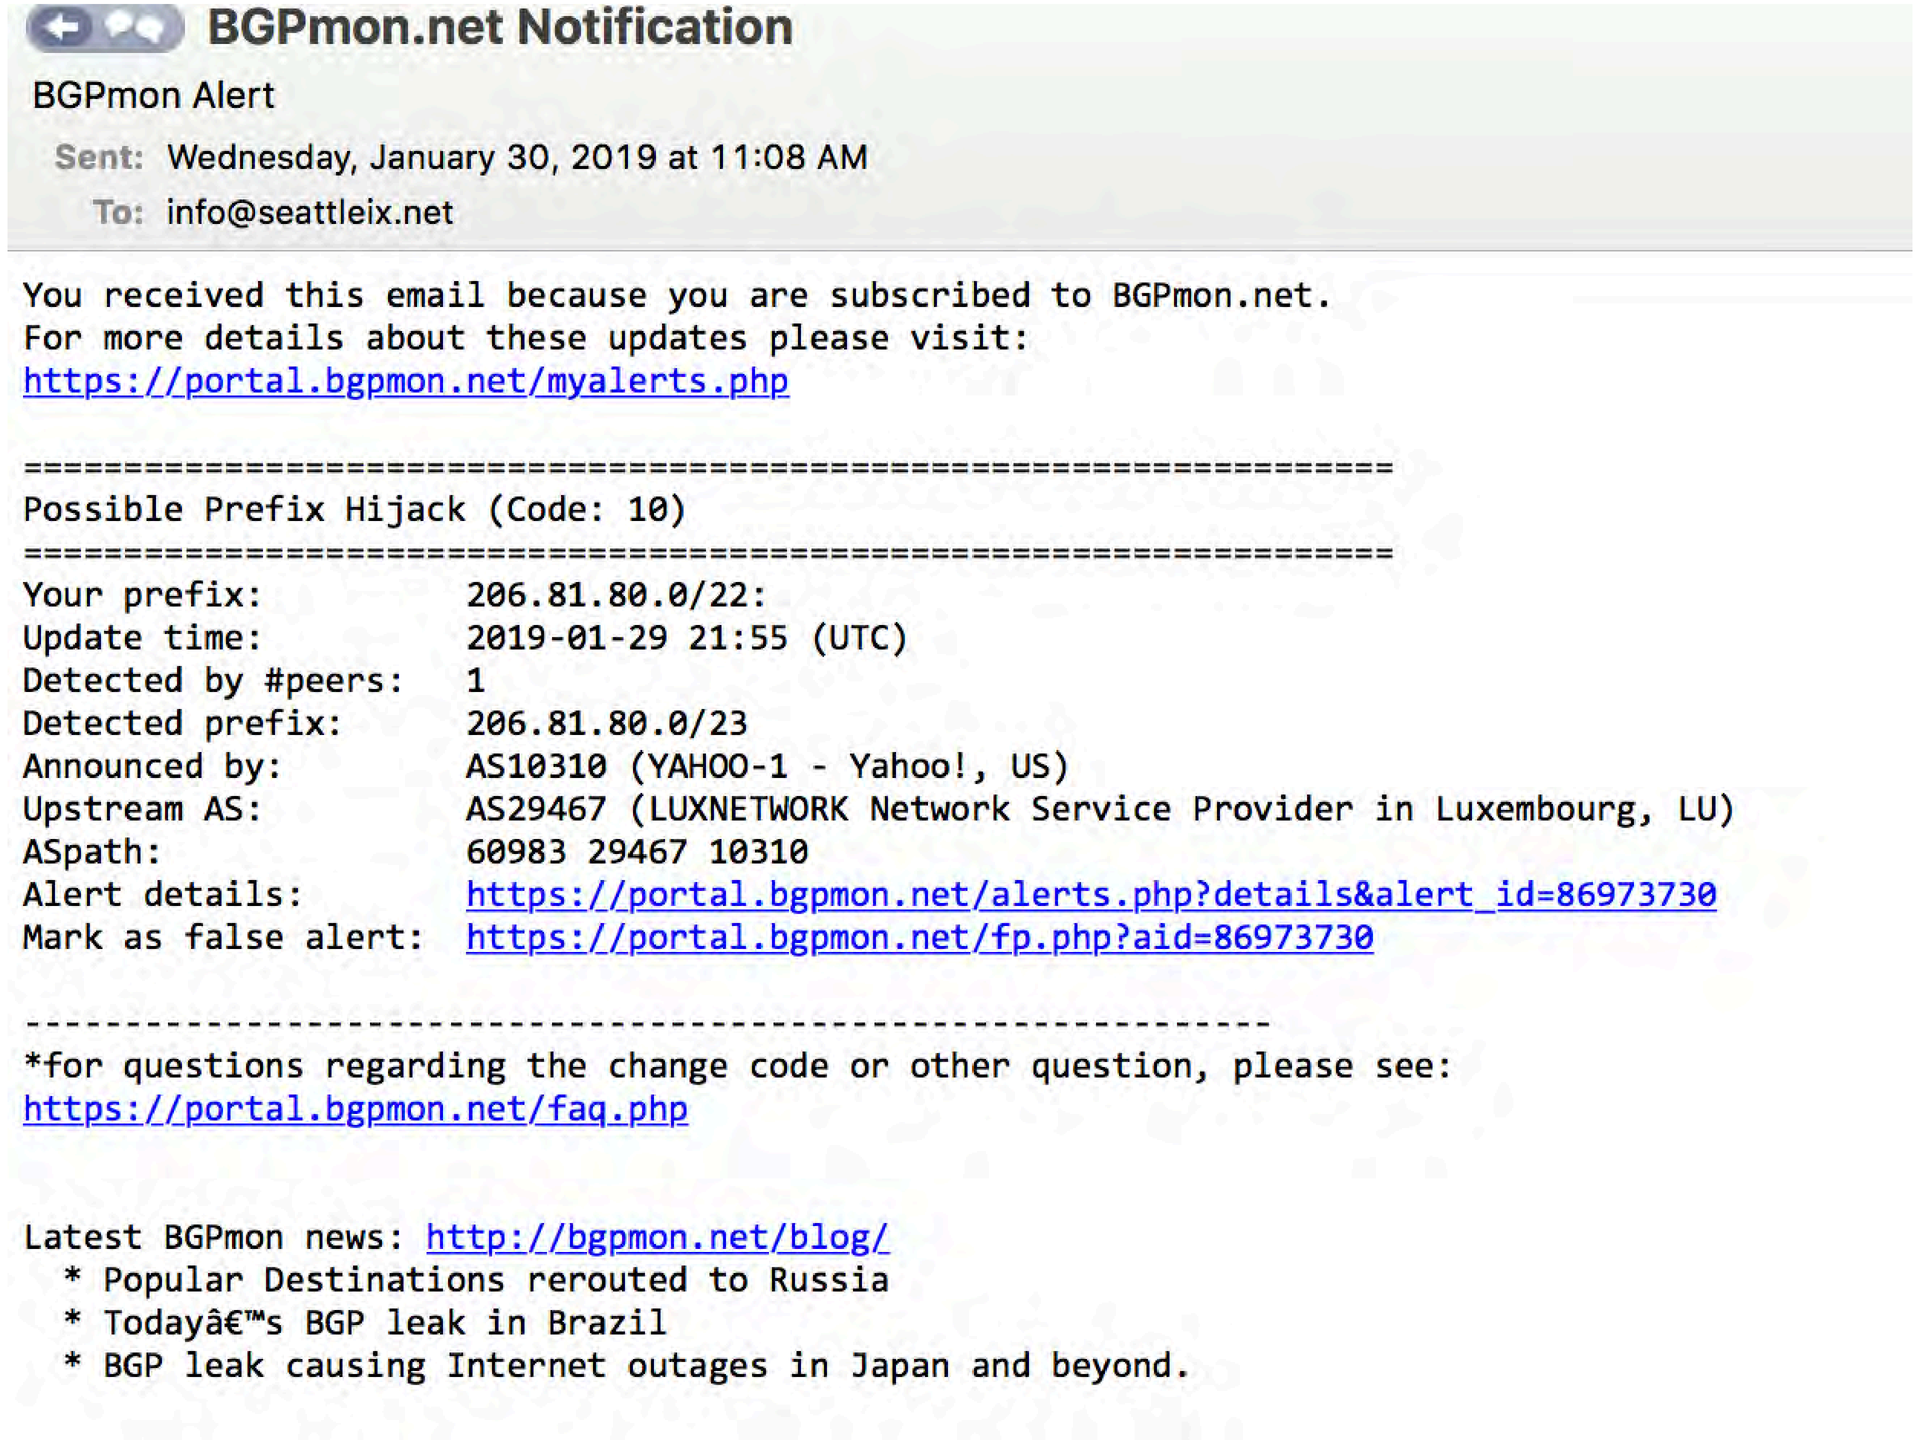

BGPmon.net Notification

BGPmon Alert

Sent: Wednesday, January 30, 2019 at 11:08 AM

To: info@seattleix.net

```
You received this email because you are subscribed to BGPmon.net.
For more details about these updates please visit:
https://portal.bgpmon.net/myalerts.php

====================================================================
Possible Prefix Hijack (Code: 10)
====================================================================
Your prefix:          206.81.80.0/22:
Update time:          2019-01-29 21:55 (UTC)
Detected by #peers:   1
Detected prefix:      206.81.80.0/23
Announced by:         AS10310 (YAHOO-1 - Yahoo!, US)
Upstream AS:          AS29467 (LUXNETWORK Network Service Provider in Luxembourg, LU)
ASpath:               60983 29467 10310
Alert details:        https://portal.bgpmon.net/alerts.php?details&alert_id=86973730
Mark as false alert:  https://portal.bgpmon.net/fp.php?aid=86973730

--------------------------------------------------------------
*for questions regarding the change code or other question, please see:
https://portal.bgpmon.net/faq.php


Latest BGPmon news: http://bgpmon.net/blog/
  * Popular Destinations rerouted to Russia
  * Todayâ€™s BGP leak in Brazil
  * BGP leak causing Internet outages in Japan and beyond.
```

IXP Subnet

Why do we care?

IXP Subnet

The DDoS That Almost Broke the Internet

Cloudflare March 2013 ~120Gbps attack on LINX

Basics - Routing Security

You must filter your peers.

- Most networks don’t filter their peers.
- This is negligent behavior.

Routing Security: Why it matters

On 28 December 2018 China Telecom hijacked a US Department of Energy prefix (192.208.19.0/24) and did not correct the problem for 6 days.

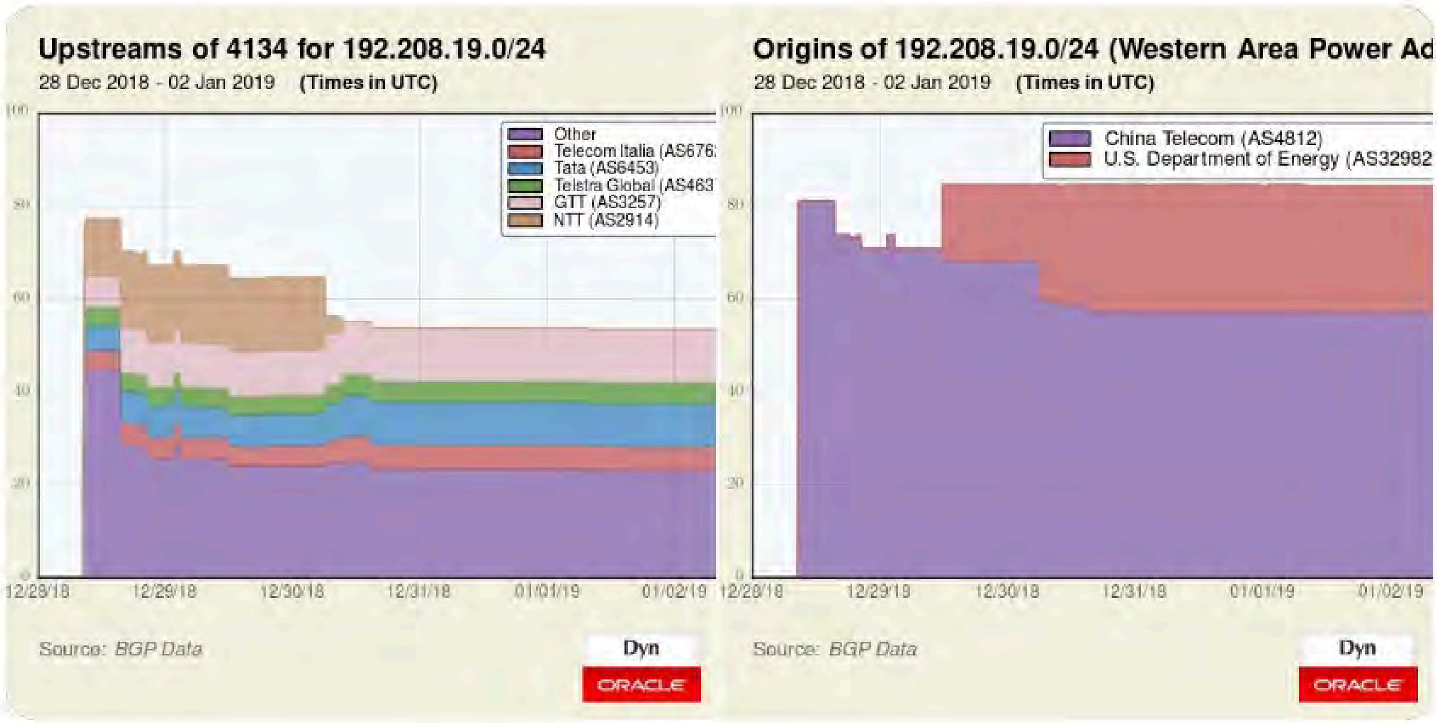

Upstreams of 4134 for 192.208.19.0/24

28 Dec 2018 - 02 Jan 2019 (Times in UTC)

Source: BGP Data

Origins of 192.208.19.0/24 (Western Area Power Ad

28 Dec 2018 - 02 Jan 2019 (Times in UTC)

Source: BGP Data

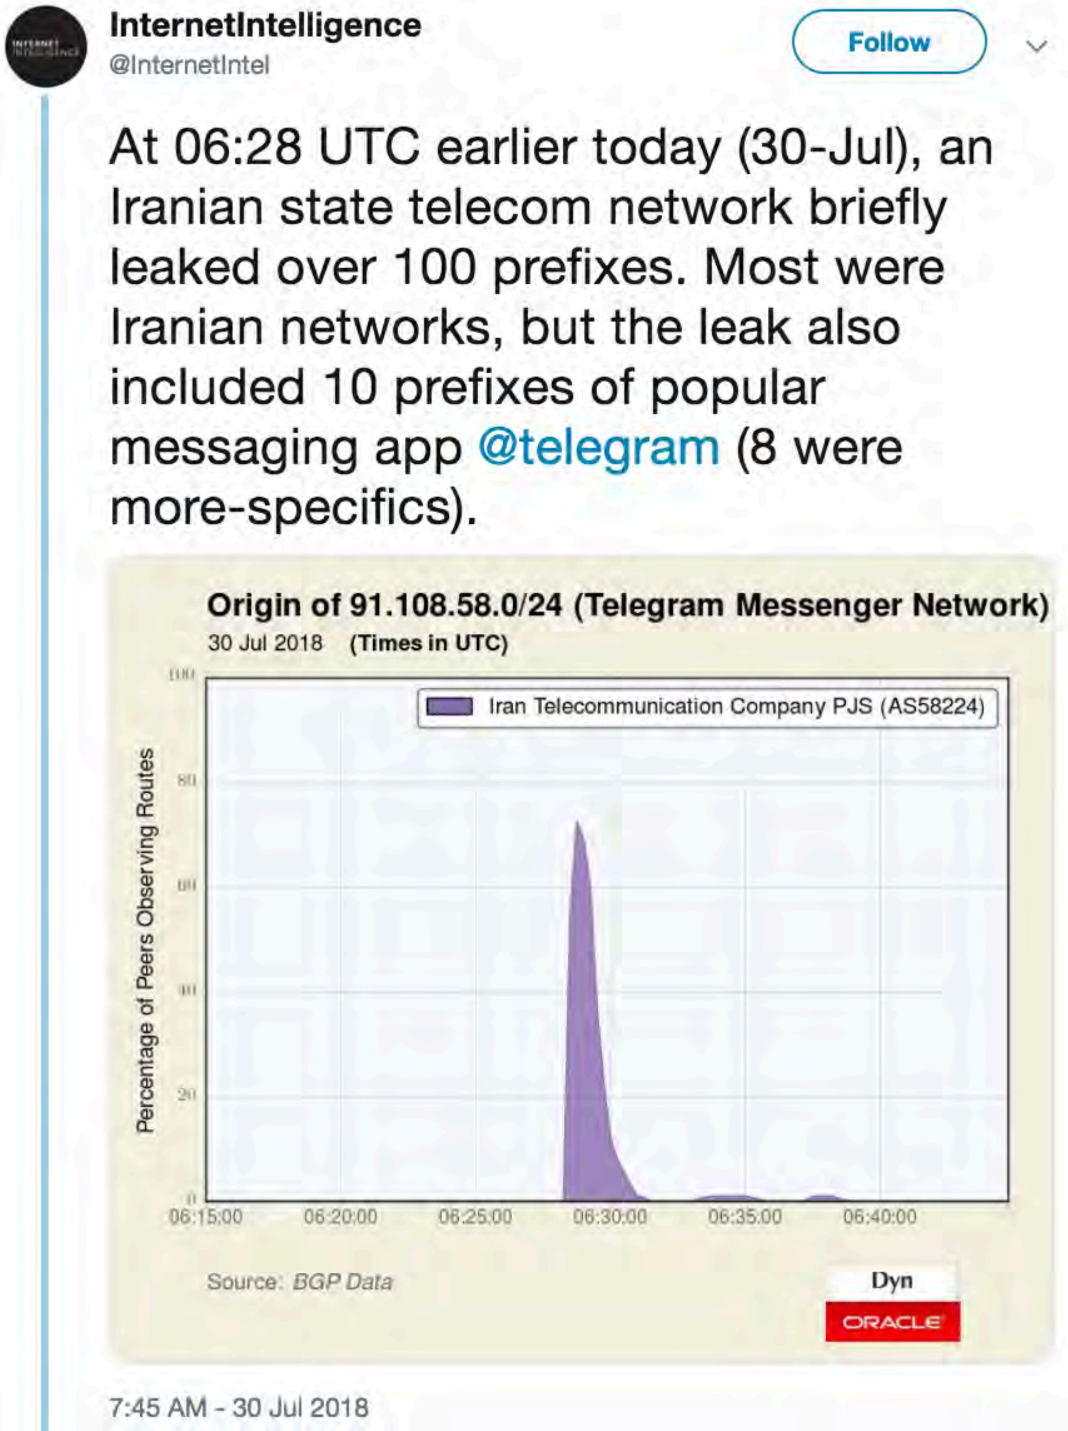

InternetIntelligence
@InternetIntel

Follow

At 06:28 UTC earlier today (30-Jul), an Iranian state telecom network briefly leaked over 100 prefixes. Most were Iranian networks, but the leak also included 10 prefixes of popular messaging app @telegram (8 were more-specifics).

Origin of 91.108.58.0/24 (Telegram Messenger Network)
30 Jul 2018 (Times in UTC)
Iran Telecommunication Company PJS (AS58224)
Percentage of Peers Observing Routes
100
80
60
40
20
0
06:15:00
06:20:00
06:25:00
06:30:00
06:35:00
06:40:00
Source: BGP Data
Dyn
ORACLE

7:45 AM - 30 Jul 2018

https://bgpstream.com

- In the last few days there have been several hijacks and leaks

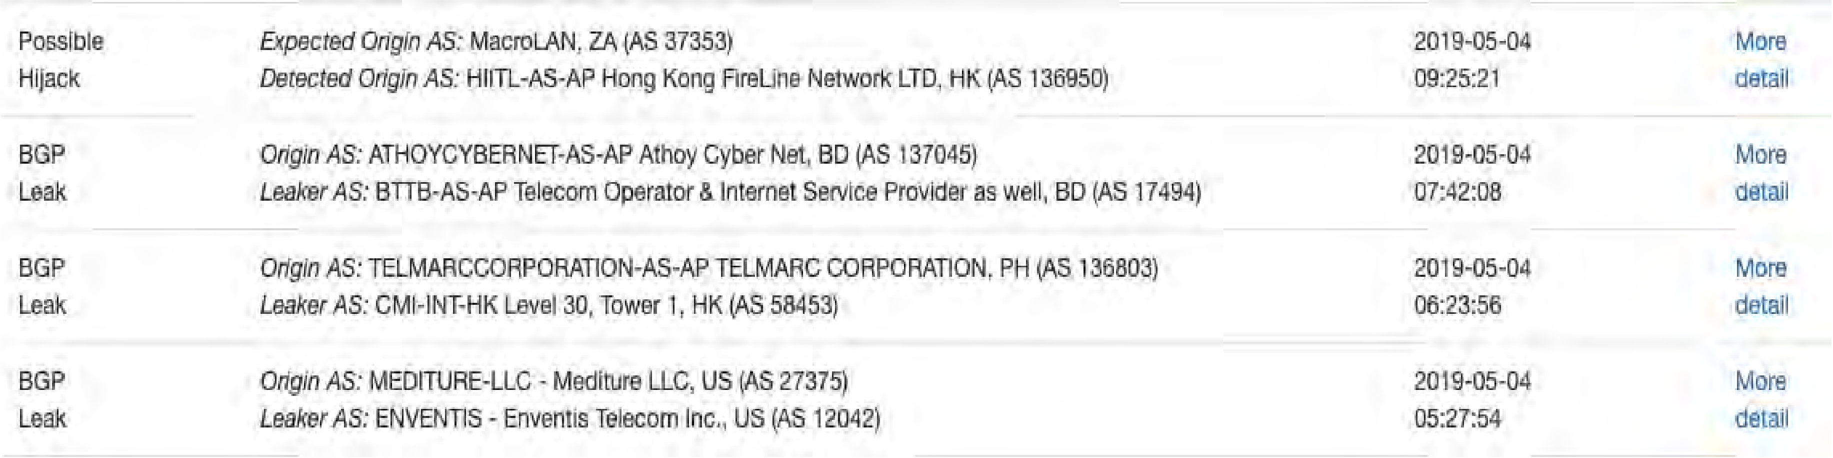

Possible Hijack	*Expected Origin AS:* MacroLAN, ZA (AS 37353) *Detected Origin AS:* HIITL-AS-AP Hong Kong FireLine Network LTD, HK (AS 136950)	2019-05-04 09:25:21	More detail
BGP Leak	*Origin AS:* ATHOYCYBERNET-AS-AP Athoy Cyber Net, BD (AS 137045) *Leaker AS:* BTTB-AS-AP Telecom Operator & Internet Service Provider as well, BD (AS 17494)	2019-05-04 07:42:08	More detail
BGP Leak	*Origin AS:* TELMARCCORPORATION-AS-AP TELMARC CORPORATION, PH (AS 136803) *Leaker AS:* CMI-INT-HK Level 30, Tower 1, HK (AS 58453)	2019-05-04 06:23:56	More detail
BGP Leak	*Origin AS:* MEDITURE-LLC - Mediture LLC, US (AS 27375) *Leaker AS:* ENVENTIS - Enventis Telecom Inc., US (AS 12042)	2019-05-04 05:27:54	More detail

Basics - Routing Security

I know we can do better

Basics - Routing Security

- Routing security is important in two directions:
  - The routes you receive
  - The routes you announce

- Starting with the routes you receive…

Basics - Routing Security

- The routes you receive can be filtered in a few ways:
  - Prefix Count
  - AS-Path
  - Prefix list
  - RPKI

Basics - Routing Security

Building filters does not have to be hard. You can script it yourself or use a tool like bgpq3. Here is an example using bgpq3 to generate a prefix list for a Juniper router:

```
walt@staff:~$ bgpq3 -J6l MyNewPrefixList AS44684
policy-options {
replace:
 prefix-list MyNewPrefixList {
    2a00:1098::/32;
    2a00:7d81:1000::/48;
    2a00:7d81:1001::/48;
    2a00:9b40::/48;
    2a06:1c80::/29;
 }
}
```

http://routing.he.net

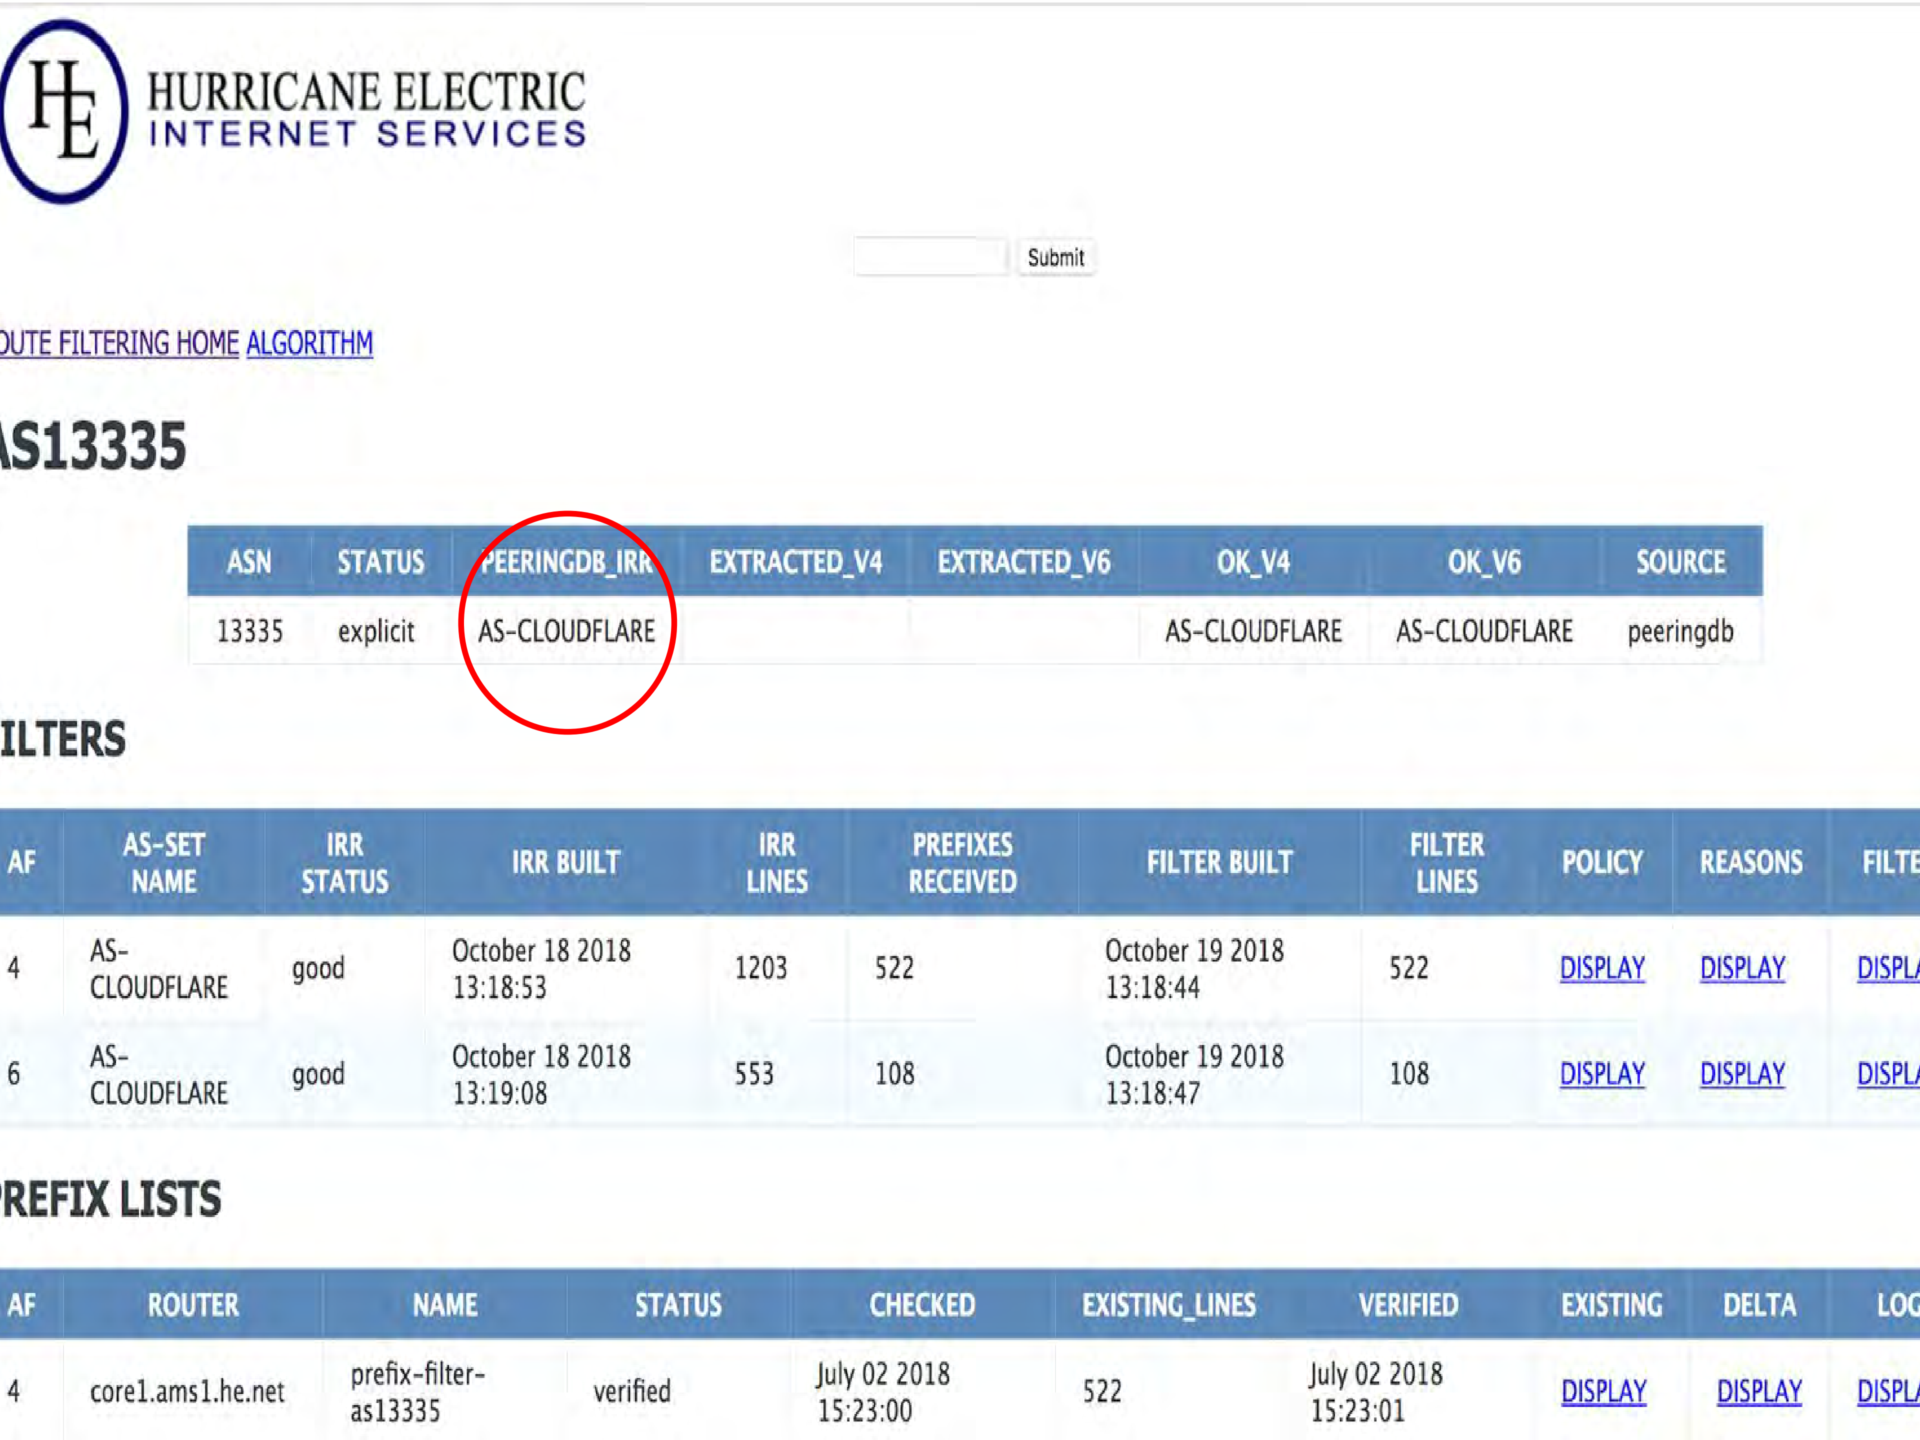

HURRICANE ELECTRIC
INTERNET SERVICES

OUTE FILTERING HOME ALGORITHM

S13335

ASN	STATUS	PEERINGDB_IRR	EXTRACTED_V4	EXTRACTED_V6	OK_V4	OK_V6	SOURCE
13335	explicit	AS-CLOUDFLARE			AS-CLOUDFLARE	AS-CLOUDFLARE	peeringdb

ILTERS

AF	AS-SET NAME	IRR STATUS	IRR BUILT	IRR LINES	PREFIXES RECEIVED	FILTER BUILT	FILTER LINES	POLICY	REASONS	FILTE
4	AS-CLOUDFLARE	good	October 18 2018 13:18:53	1203	522	October 19 2018 13:18:44	522	DISPLAY	DISPLAY	DISPL
6	AS-CLOUDFLARE	good	October 18 2018 13:19:08	553	108	October 19 2018 13:18:47	108	DISPLAY	DISPLAY	DISPL

REFIX LISTS

AF	ROUTER	NAME	STATUS	CHECKED	EXISTING_LINES	VERIFIED	EXISTING	DELTA	LOG
4	core1.ams1.he.net	prefix-filter-as13335	verified	July 02 2018 15:23:00	522	July 02 2018 15:23:01	DISPLAY	DISPLAY	DISPL

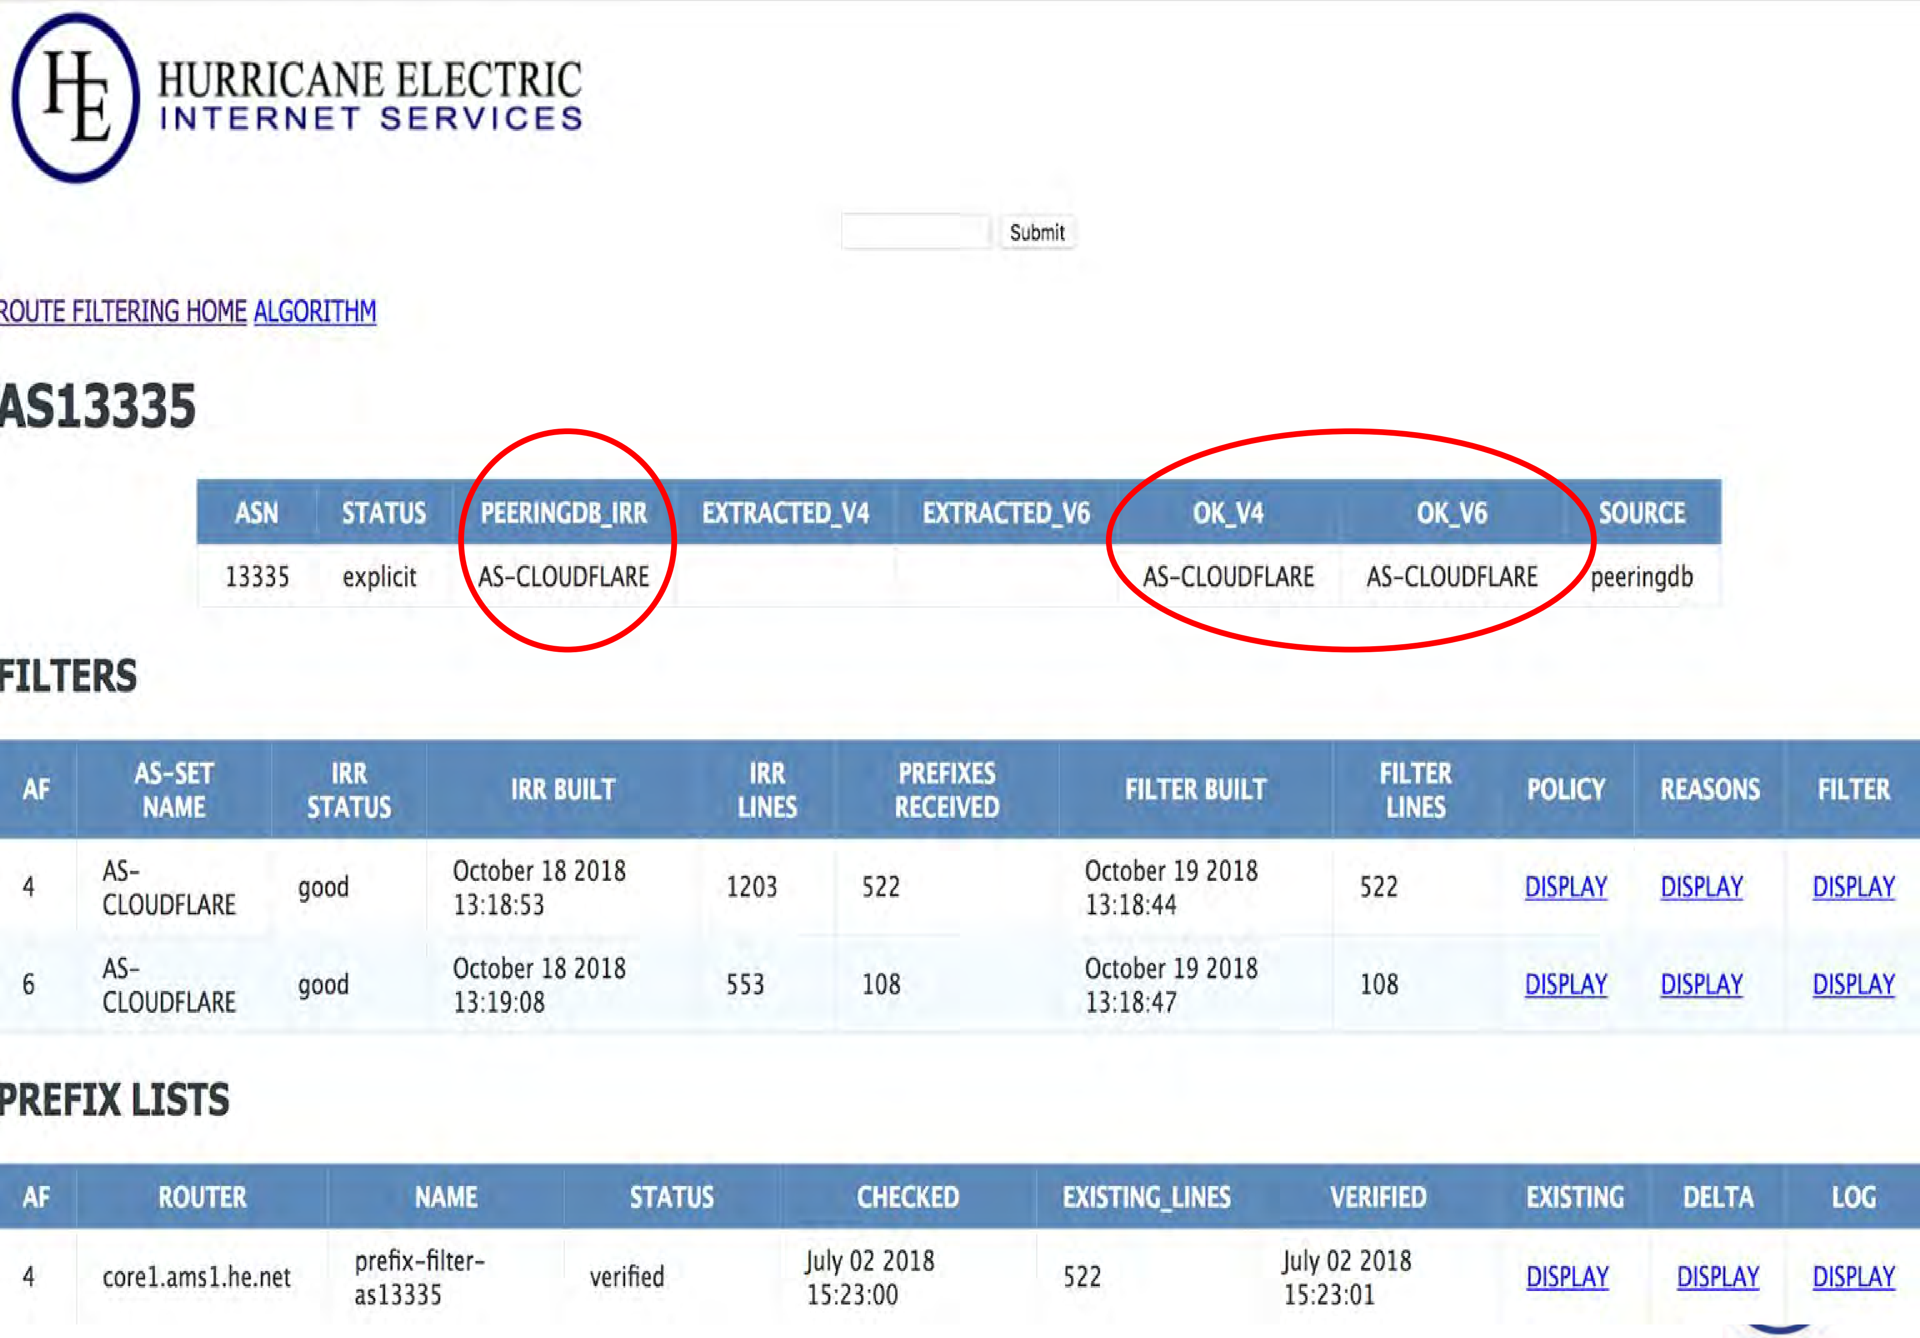

HURRICANE ELECTRIC
INTERNET SERVICES

ROUTE FILTERING HOME ALGORITHM

AS13335

ASN	STATUS	PEERINGDB_IRR	EXTRACTED_V4	EXTRACTED_V6	OK_V4	OK_V6	SOURCE
13335	explicit	AS-CLOUDFLARE			AS-CLOUDFLARE	AS-CLOUDFLARE	peeringdb

FILTERS

AF	AS-SET NAME	IRR STATUS	IRR BUILT	IRR LINES	PREFIXES RECEIVED	FILTER BUILT	FILTER LINES	POLICY	REASONS	FILTER
4	AS-CLOUDFLARE	good	October 18 2018 13:18:53	1203	522	October 19 2018 13:18:44	522	DISPLAY	DISPLAY	DISPLAY
6	AS-CLOUDFLARE	good	October 18 2018 13:19:08	553	108	October 19 2018 13:18:47	108	DISPLAY	DISPLAY	DISPLAY

PREFIX LISTS

AF	ROUTER	NAME	STATUS	CHECKED	EXISTING_LINES	VERIFIED	EXISTING	DELTA	LOG
4	core1.ams1.he.net	prefix-filter-as13335	verified	July 02 2018 15:23:00	522	July 02 2018 15:23:01	DISPLAY	DISPLAY	DISPLAY

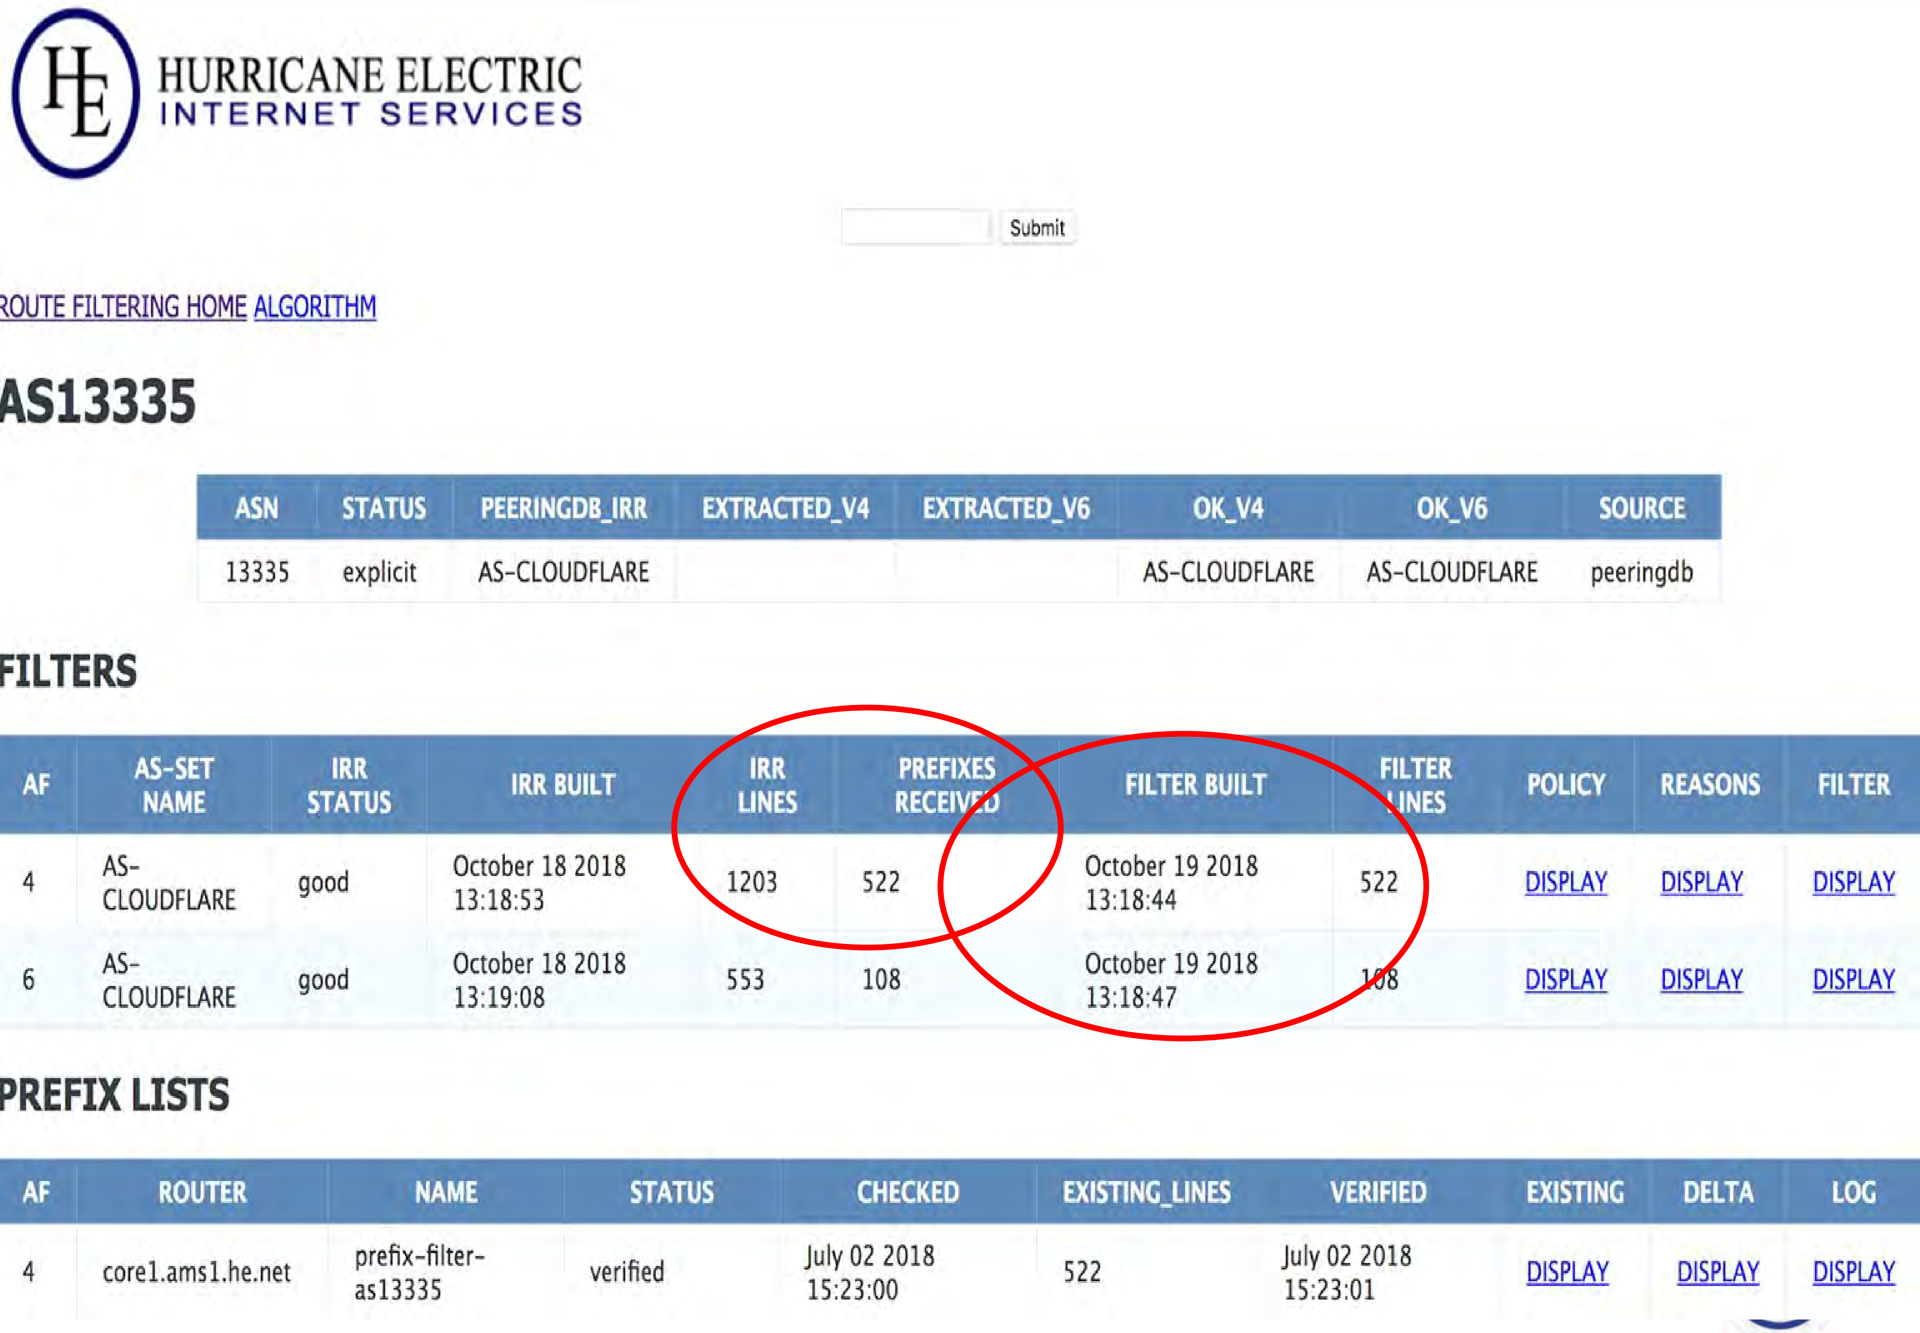

HURRICANE ELECTRIC
INTERNET SERVICES

ROUTE FILTERING HOME ALGORITHM

AS13335

ASN	STATUS	PEERINGDB_IRR	EXTRACTED_V4	EXTRACTED_V6	OK_V4	OK_V6	SOURCE
13335	explicit	AS-CLOUDFLARE			AS-CLOUDFLARE	AS-CLOUDFLARE	peeringdb

FILTERS

AF	AS-SET NAME	IRR STATUS	IRR BUILT	IRR LINES	PREFIXES RECEIVED	FILTER BUILT	FILTER LINES	POLICY	REASONS	FILTER
4	AS-CLOUDFLARE	good	October 18 2018 13:18:53	1203	522	October 19 2018 13:18:44	522	DISPLAY	DISPLAY	DISPLAY
6	AS-CLOUDFLARE	good	October 18 2018 13:19:08	553	108	October 19 2018 13:18:47	108	DISPLAY	DISPLAY	DISPLAY

PREFIX LISTS

AF	ROUTER	NAME	STATUS	CHECKED	EXISTING_LINES	VERIFIED	EXISTING	DELTA	LOG
4	core1.ams1.he.net	prefix-filter-as13335	verified	July 02 2018 15:23:00	522	July 02 2018 15:23:01	DISPLAY	DISPLAY	DISPLAY

http://routing.he.net

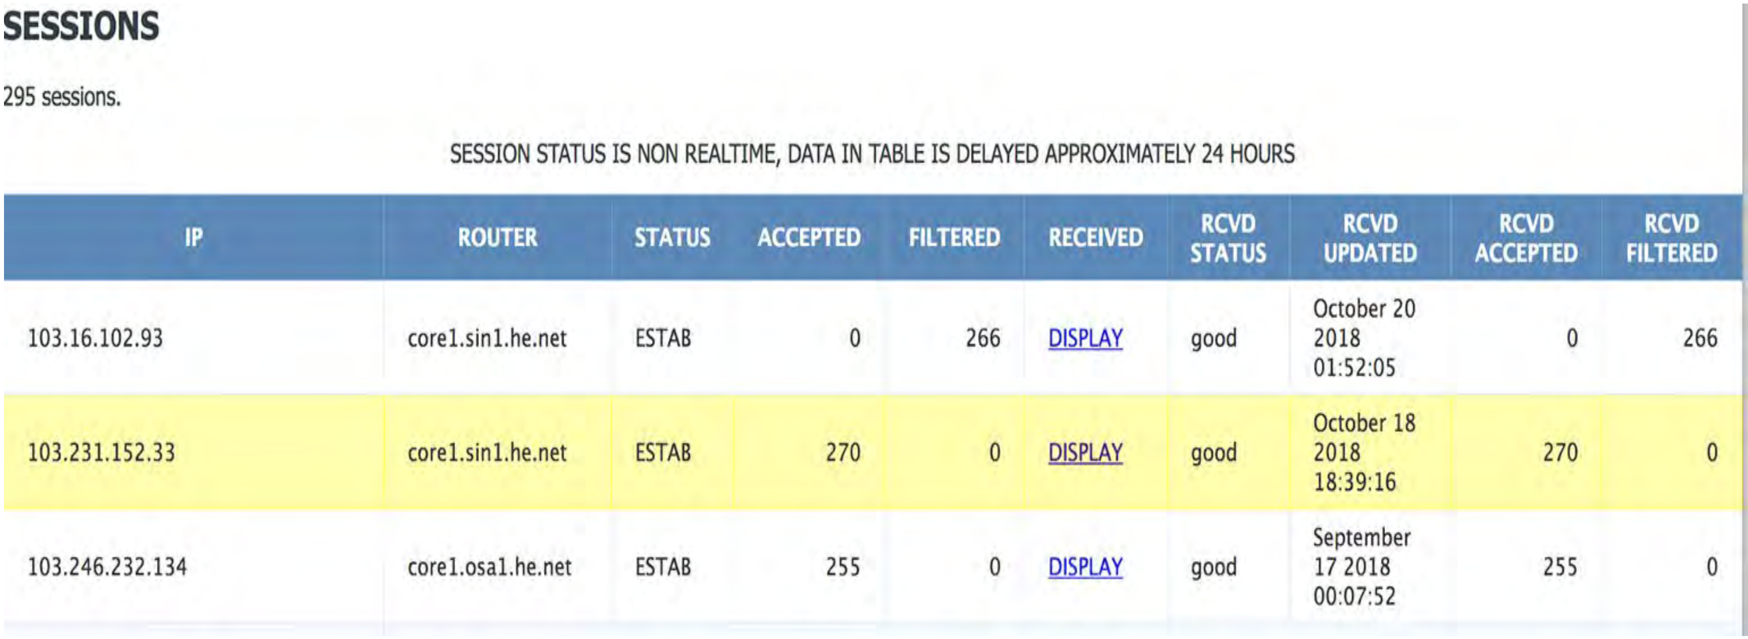

SESSIONS

295 sessions.

SESSION STATUS IS NON REALTIME, DATA IN TABLE IS DELAYED APPROXIMATELY 24 HOURS

IP	ROUTER	STATUS	ACCEPTED	FILTERED	RECEIVED	RCVD STATUS	RCVD UPDATED	RCVD ACCEPTED	RCVD FILTERED
103.16.102.93	core1.sin1.he.net	ESTAB	0	266	DISPLAY	good	October 20 2018 01:52:05	0	266
103.231.152.33	core1.sin1.he.net	ESTAB	270	0	DISPLAY	good	October 18 2018 18:39:16	270	0
103.246.232.134	core1.osa1.he.net	ESTAB	255	0	DISPLAY	good	September 17 2018 00:07:52	255	0

http://routing.he.net

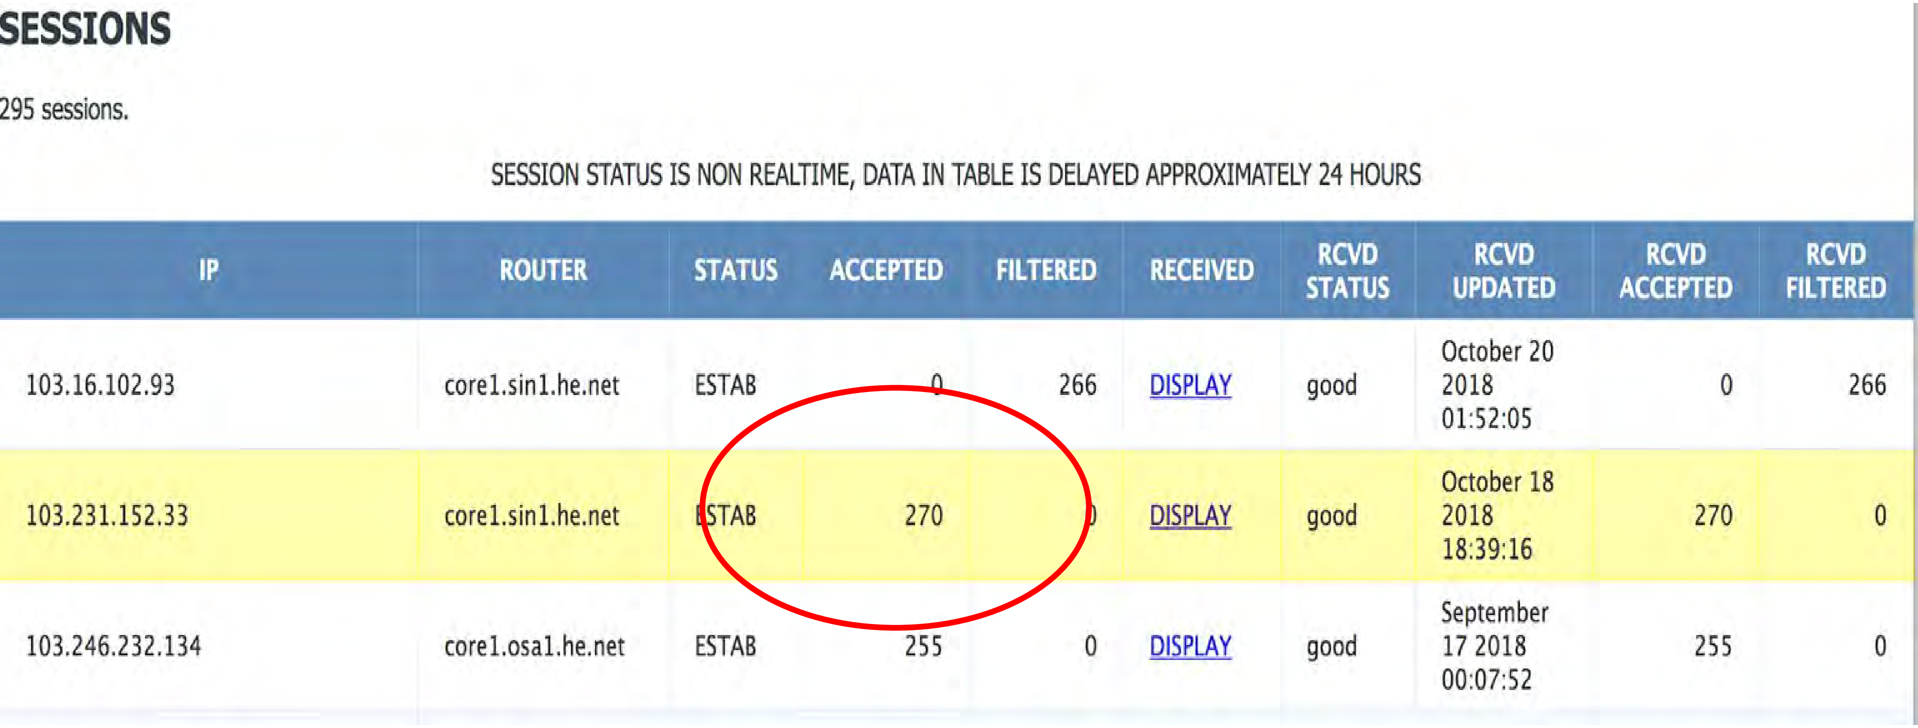

SESSIONS

295 sessions.

SESSION STATUS IS NON REALTIME, DATA IN TABLE IS DELAYED APPROXIMATELY 24 HOURS

IP	ROUTER	STATUS	ACCEPTED	FILTERED	RECEIVED	RCVD STATUS	RCVD UPDATED	RCVD ACCEPTED	RCVD FILTERED
103.16.102.93	core1.sin1.he.net	ESTAB	0	266	DISPLAY	good	October 20 2018 01:52:05	0	266
103.231.152.33	core1.sin1.he.net	ESTAB	270	0	DISPLAY	good	October 18 2018 18:39:16	270	0
103.246.232.134	core1.osa1.he.net	ESTAB	255	0	DISPLAY	good	September 17 2018 00:07:52	255	0

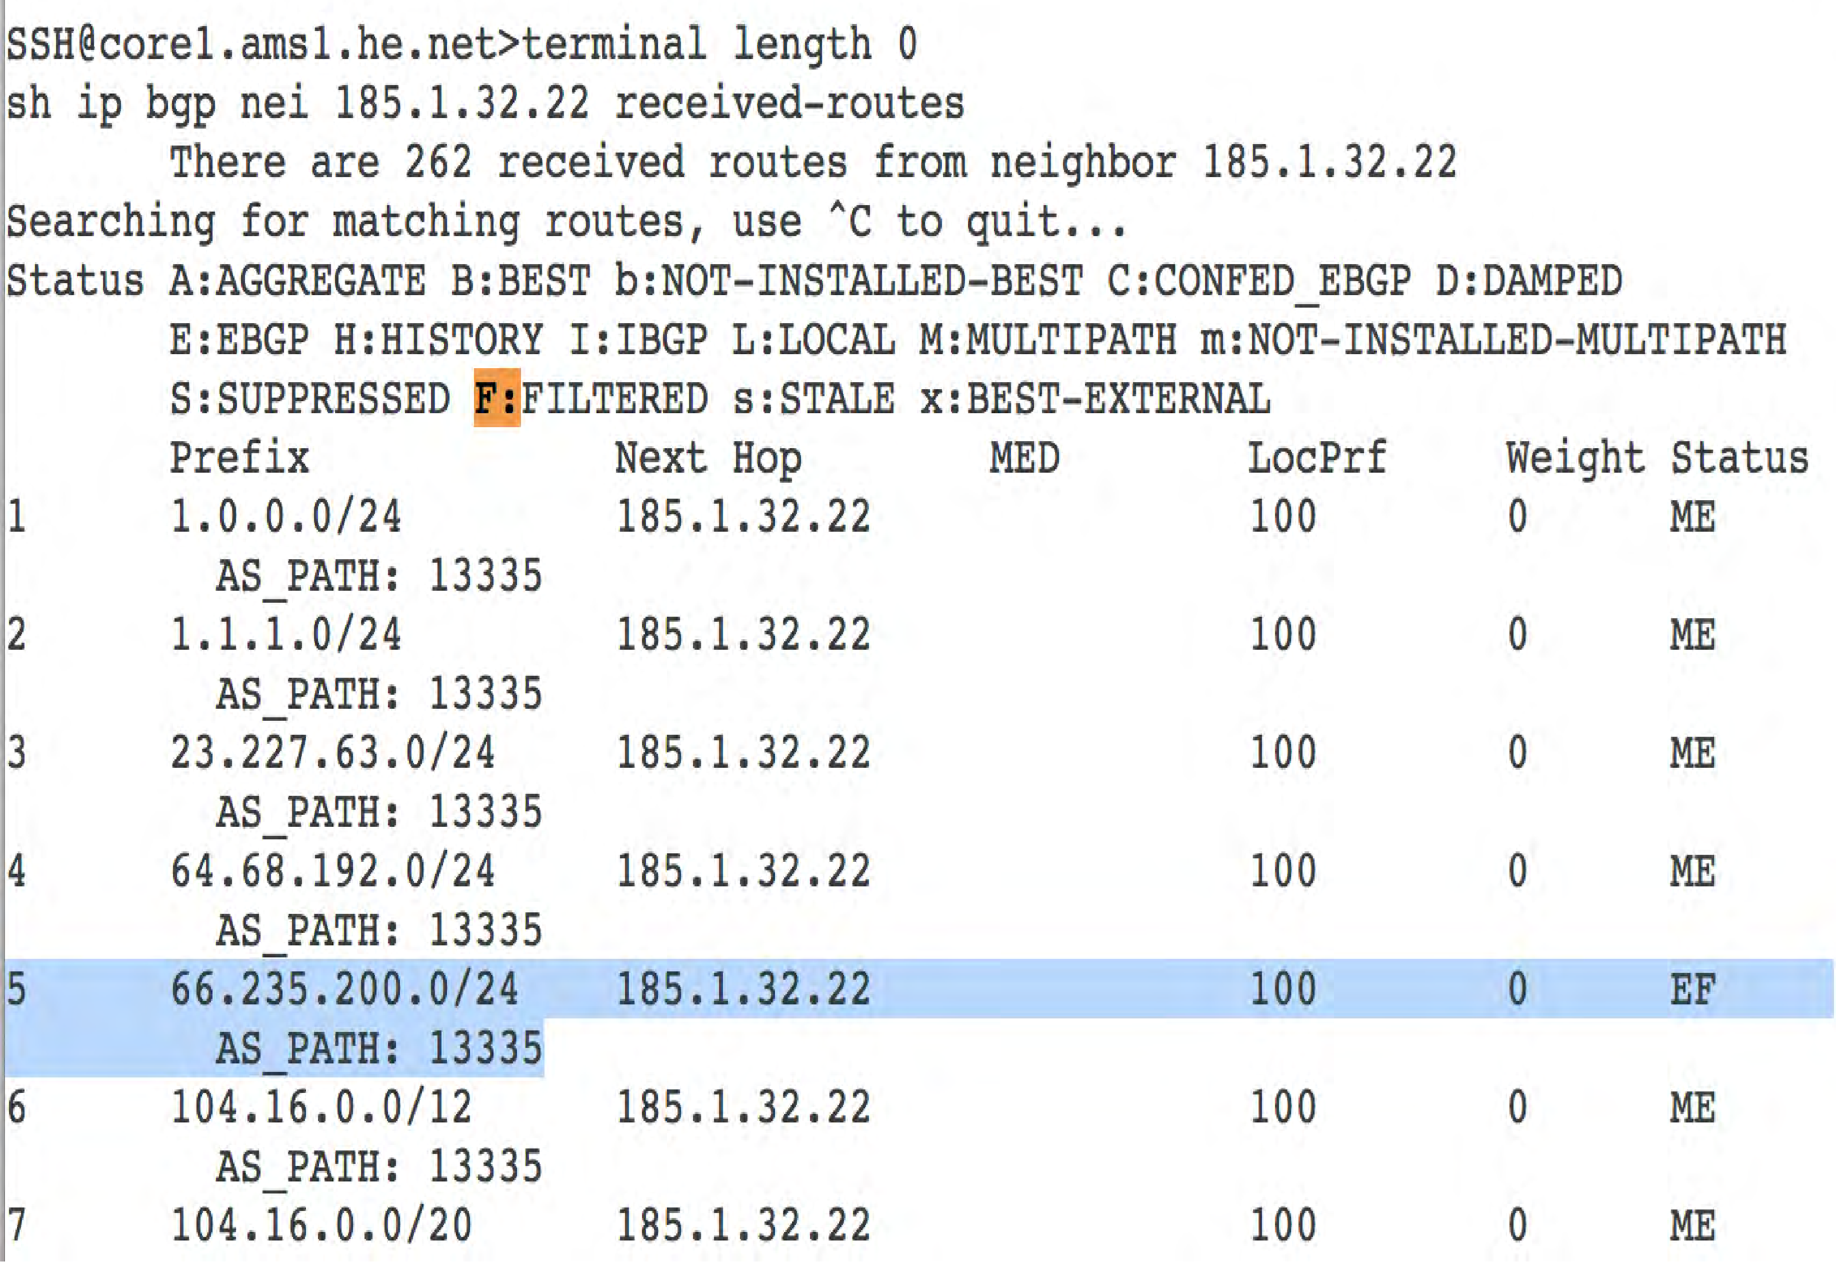

```
SSH@core1.ams1.he.net>terminal length 0
sh ip bgp nei 185.1.32.22 received-routes
       There are 262 received routes from neighbor 185.1.32.22
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
       E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH m:NOT-INSTALLED-MULTIPATH
       S:SUPPRESSED F:FILTERED s:STALE x:BEST-EXTERNAL
       Prefix             Next Hop        MED         LocPrf      Weight Status
1      1.0.0.0/24         185.1.32.22                 100         0      ME
         AS_PATH: 13335
2      1.1.1.0/24         185.1.32.22                 100         0      ME
         AS_PATH: 13335
3      23.227.63.0/24     185.1.32.22                 100         0      ME
         AS_PATH: 13335
4      64.68.192.0/24     185.1.32.22                 100         0      ME
         AS_PATH: 13335
5      66.235.200.0/24    185.1.32.22                 100         0      EF
         AS_PATH: 13335
6      104.16.0.0/12      185.1.32.22                 100         0      ME
         AS_PATH: 13335
7      104.16.0.0/20      185.1.32.22                 100         0      ME
```

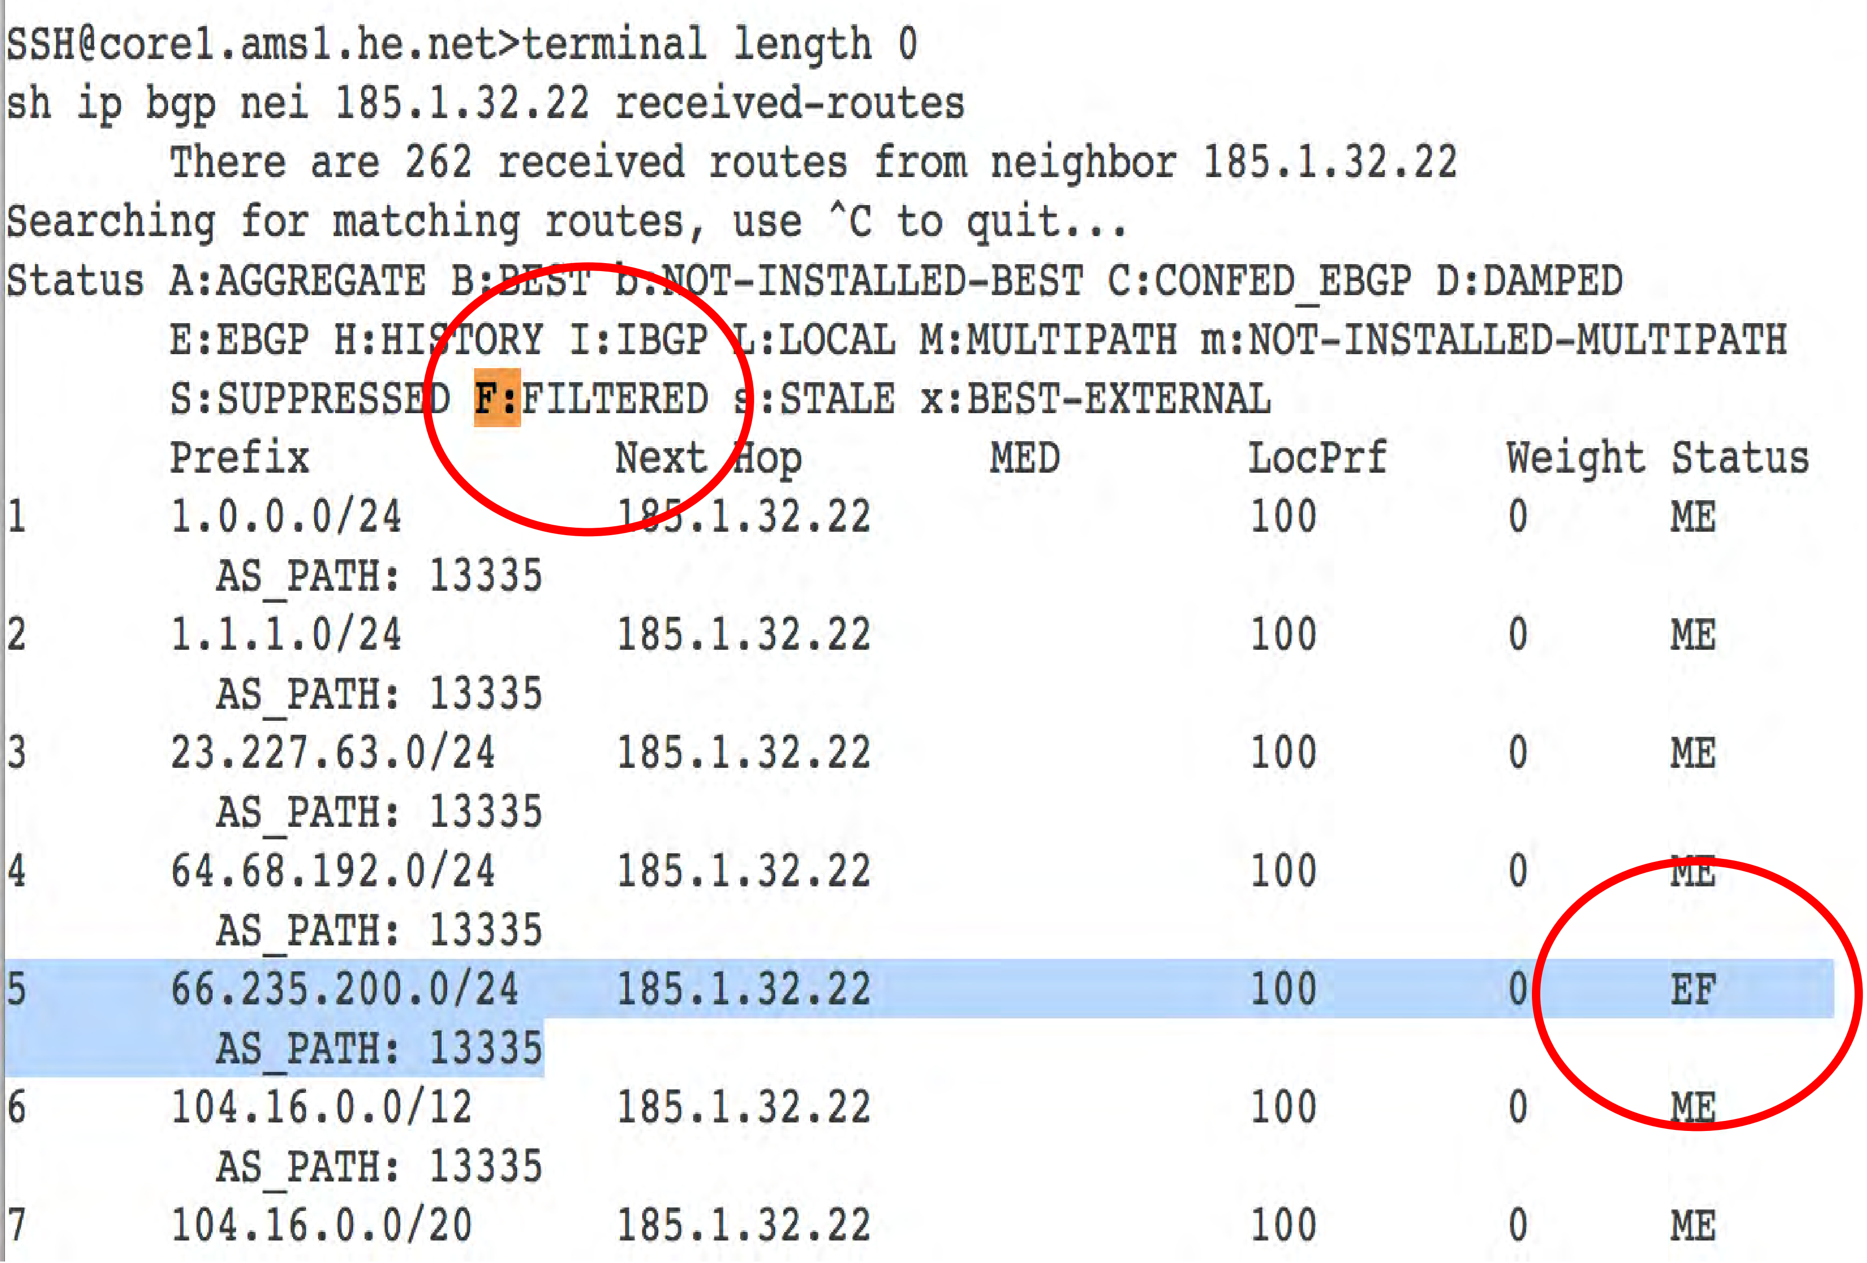

```
SSH@core1.ams1.he.net>terminal length 0
sh ip bgp nei 185.1.32.22 received-routes
       There are 262 received routes from neighbor 185.1.32.22
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
       E:EBGP H:HISTORY I:IBGP L:LOCAL M:MULTIPATH m:NOT-INSTALLED-MULTIPATH
       S:SUPPRESSED F:FILTERED s:STALE x:BEST-EXTERNAL
       Prefix             Next Hop        MED        LocPrf     Weight Status
1      1.0.0.0/24         185.1.32.22                100        0      ME
         AS_PATH: 13335
2      1.1.1.0/24         185.1.32.22                100        0      ME
         AS_PATH: 13335
3      23.227.63.0/24     185.1.32.22                100        0      ME
         AS_PATH: 13335
4      64.68.192.0/24     185.1.32.22                100        0      ME
         AS_PATH: 13335
5      66.235.200.0/24    185.1.32.22                100        0      EF
         AS_PATH: 13335
6      104.16.0.0/12      185.1.32.22                100        0      ME
         AS_PATH: 13335
7      104.16.0.0/20      185.1.32.22                100        0      ME
```

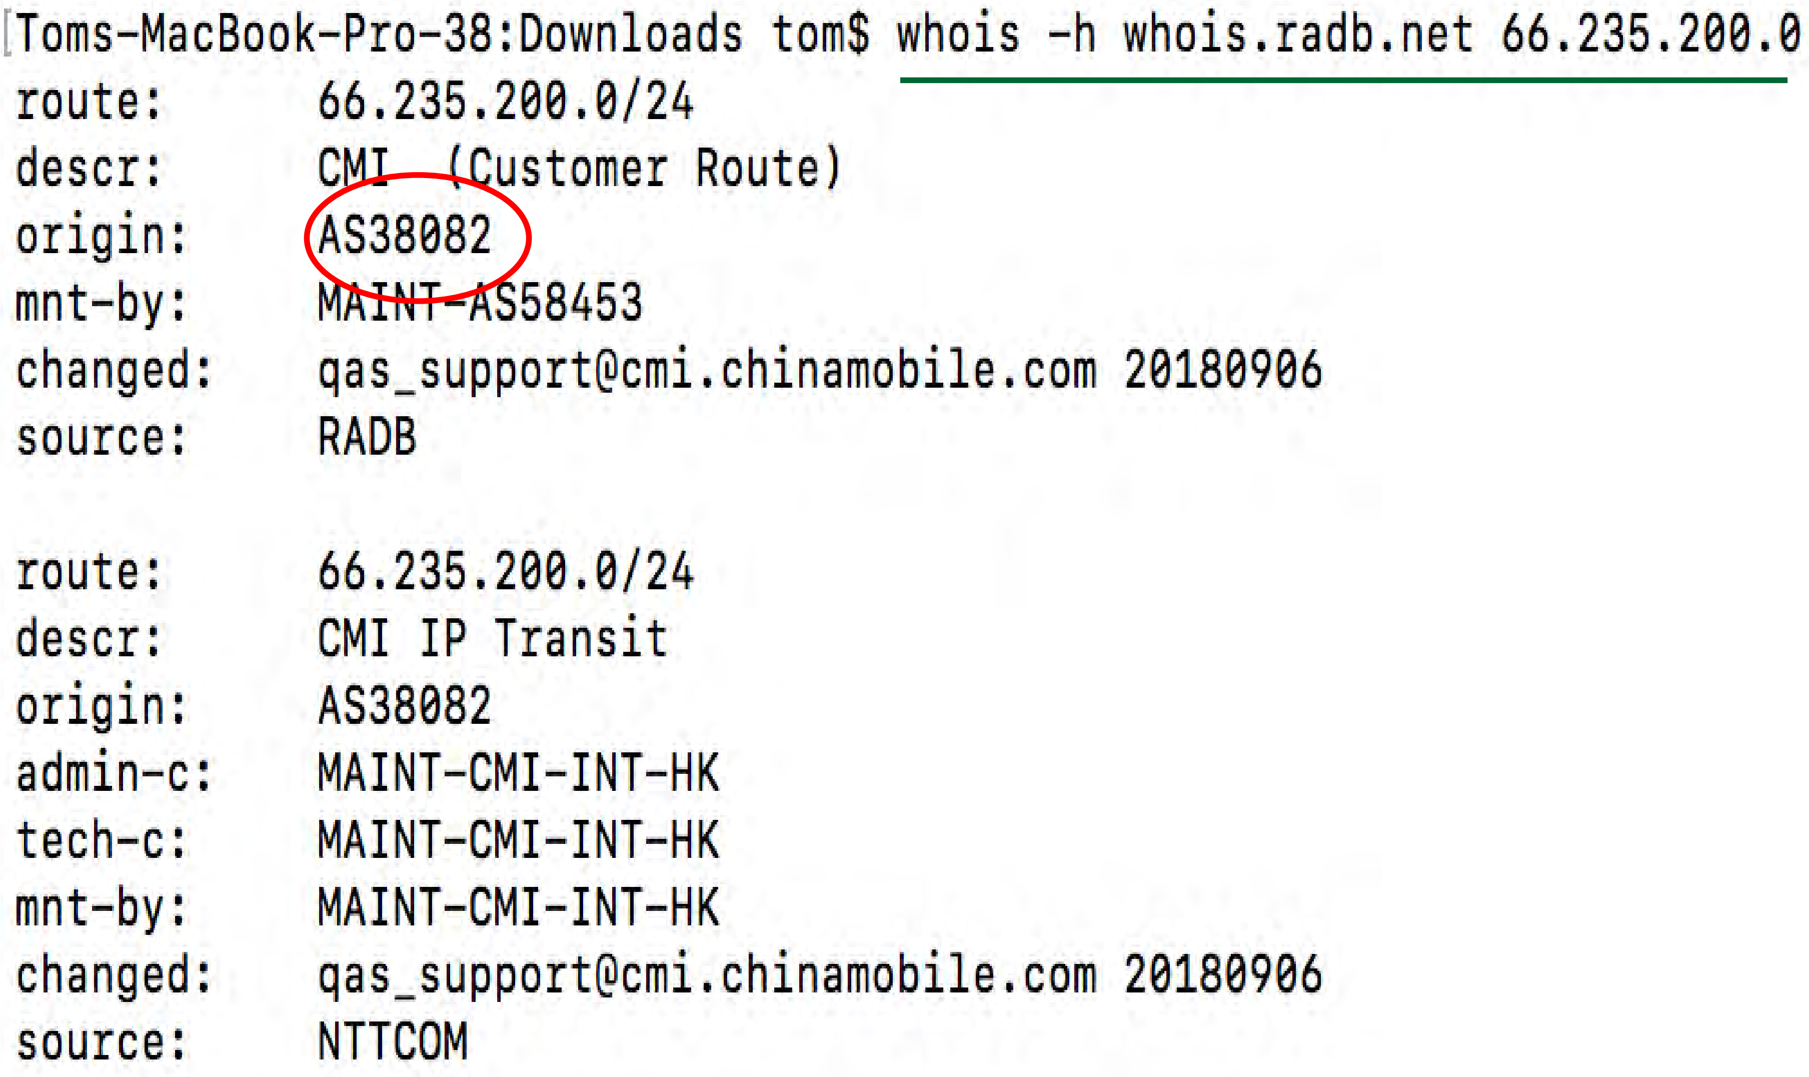

```
Toms-MacBook-Pro-38:Downloads tom$ whois -h whois.radb.net 66.235.200.0
route:      66.235.200.0/24
descr:      CMI  (Customer Route)
origin:     AS38082
mnt-by:     MAINT-AS58453
changed:    qas_support@cmi.chinamobile.com 20180906
source:     RADB

route:      66.235.200.0/24
descr:      CMI  IP Transit
origin:     AS38082
admin-c:    MAINT-CMI-INT-HK
tech-c:     MAINT-CMI-INT-HK
mnt-by:     MAINT-CMI-INT-HK
changed:    qas_support@cmi.chinamobile.com 20180906
source:     NTTCOM
```

Hurricane Electric
Route Filtering Algorithm

- Read more here

  http://routing.he.net/algorithm.html

- Example:
- xx.7.224.0/24,rejected,does not strictly match IRR policy or RIR handles
- xx.10.254.0/23,accepted,strictly matched IRR policy
- xx.17.248.0/24,accepted,strictly matched IRR policy
- xx.26.36.0/22,rejected,does not strictly match IRR policy or RIR handles
- xx.26.39.0/24,rejected,does not strictly match IRR policy or RIR handles

Hurricane Electric Route Filtering

- Please check and update your IRR or RIR handles
- Check your routing here:
  http://routing.he.net/
- We at now filtering ~90% of all our peers.
- Rolling it out slowly over the last six months

Resources

- https://www.seattleix.net/faq
- https://twitter.com/bgpstream/status/1078584924364595202?lang=en
- https://bgp.he.net
- https://github.com/snar/bgpq3
- https://bgpmon.net/
- https://bgpstream.com/
- https://bgpmon.net/

Thanks to Tom Paseka of Cloudflare.

Thanks!

Walt Wollny, Director Interconnection Strategy
Hurricane Electric AS6939
walt@he.net