HURRICANE ELECTRIC
INTERNET SERVICES

# Peering Security

IX Forum 13
Sao Paulo 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 223 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

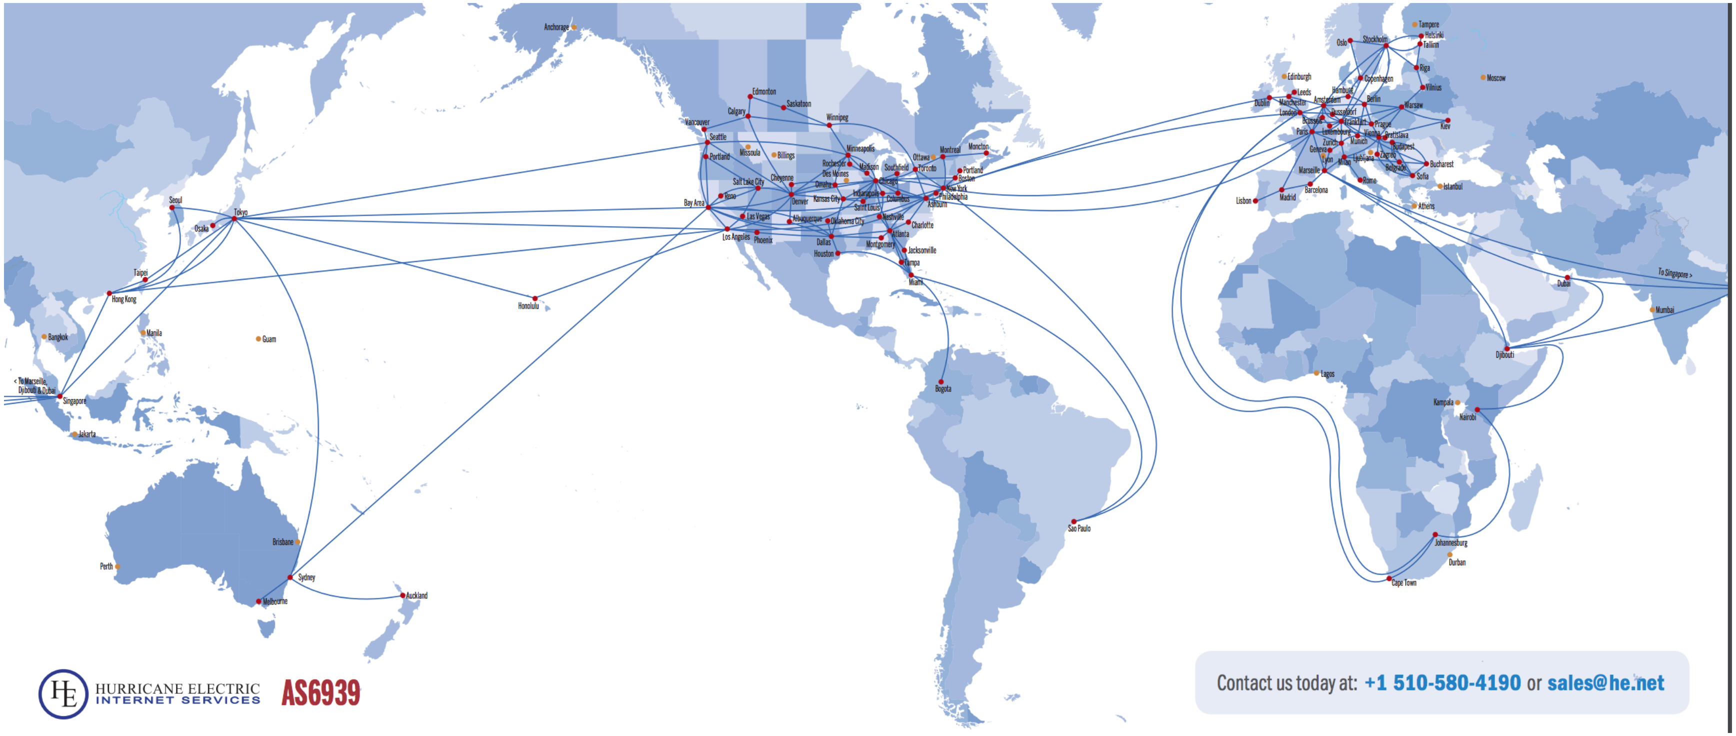

# The Most Peering Exchanges

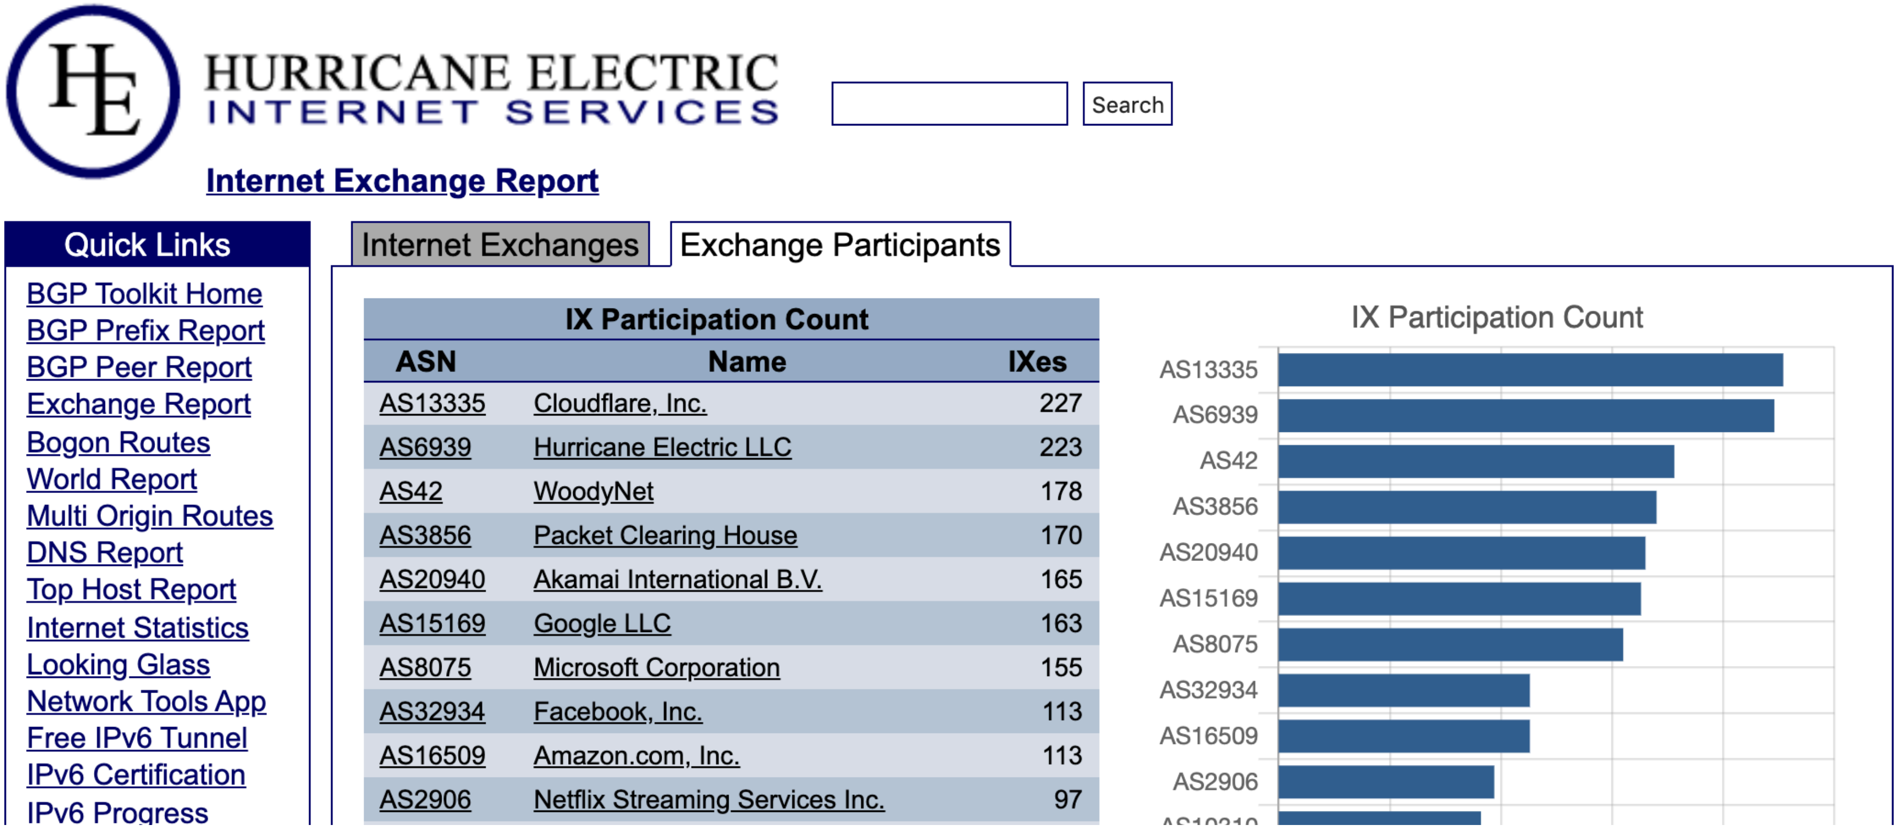

HURRICANE ELECTRIC
INTERNET SERVICES

Search

**Internet Exchange Report**

| IX Participation Count | | |
|---|---|---|
| **ASN** | **Name** | **IXes** |
| AS13335 | Cloudflare, Inc. | 227 |
| AS6939 | Hurricane Electric LLC | 223 |
| AS42 | WoodyNet | 178 |
| AS3856 | Packet Clearing House | 170 |
| AS20940 | Akamai International B.V. | 165 |
| AS15169 | Google LLC | 163 |
| AS8075 | Microsoft Corporation | 155 |
| AS32934 | Facebook, Inc. | 113 |
| AS16509 | Amazon.com, Inc. | 113 |
| AS2906 | Netflix Streaming Services Inc. | 97 |

# Why So Many Peering Exchanges?

# Why So Many Peering Exchanges?

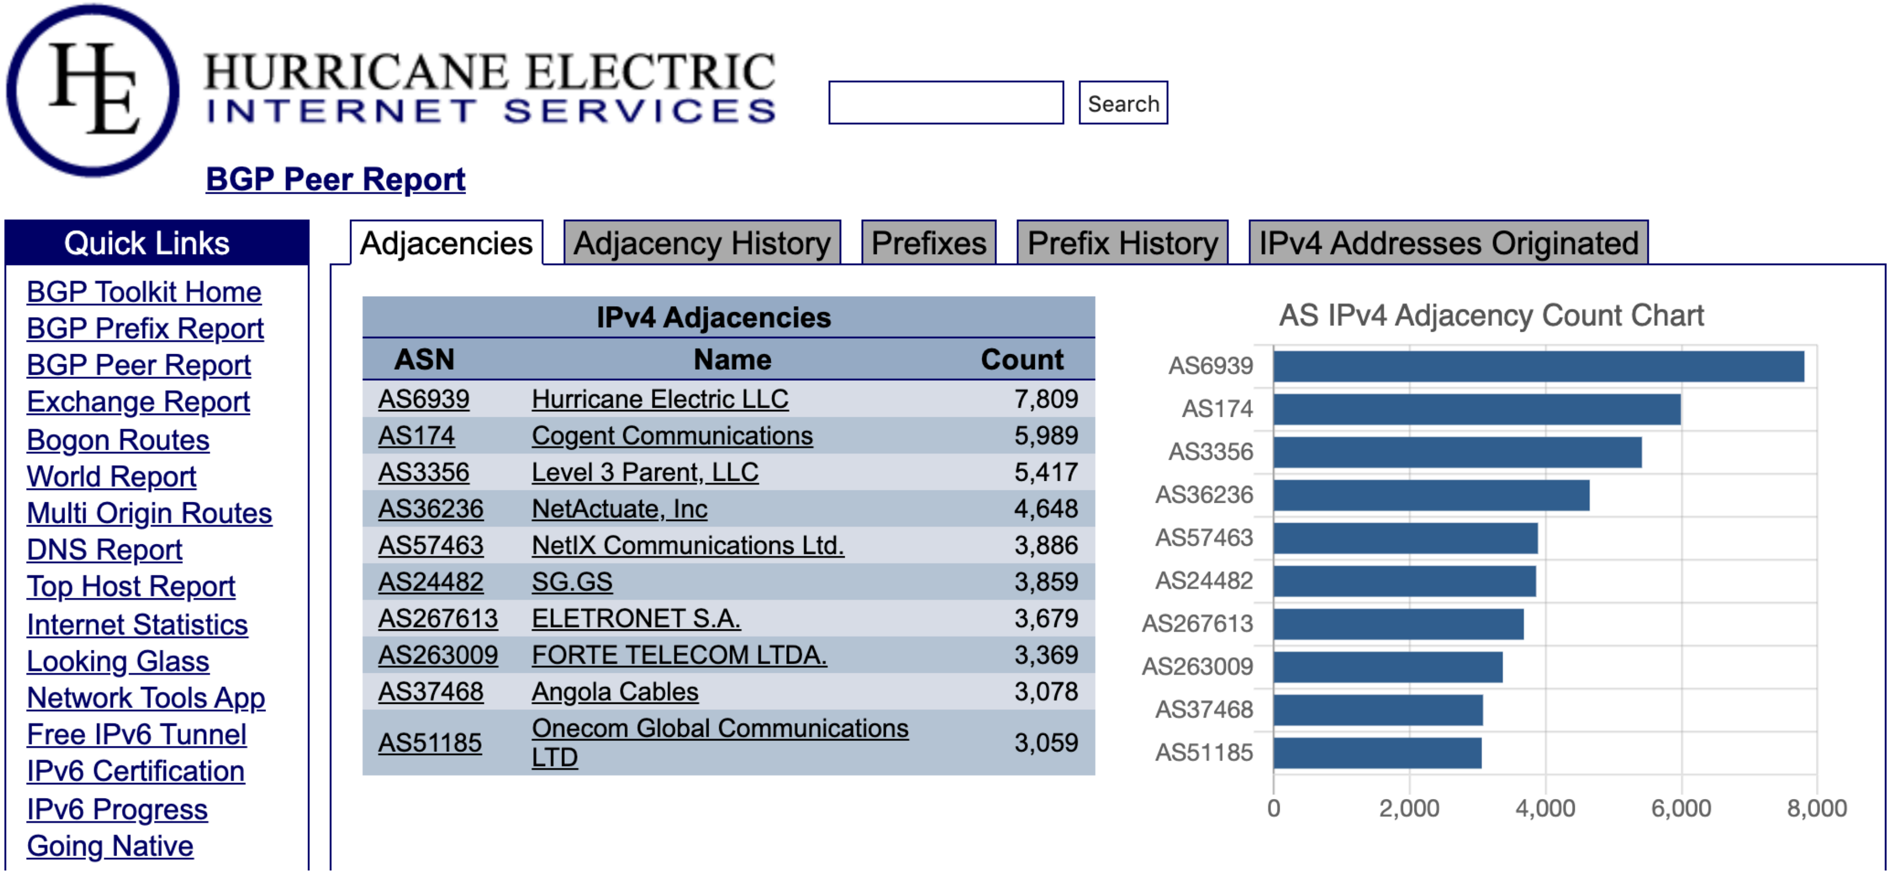

**BGP Peer Report**

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
- Going Native

**IPv4 Adjacencies**

| ASN | Name | Count |
|---|---|---|
| AS6939 | Hurricane Electric LLC | 7,809 |
| AS174 | Cogent Communications | 5,989 |
| AS3356 | Level 3 Parent, LLC | 5,417 |
| AS36236 | NetActuate, Inc | 4,648 |
| AS57463 | NetIX Communications Ltd. | 3,886 |
| AS24482 | SG.GS | 3,859 |
| AS267613 | ELETRONET S.A. | 3,679 |
| AS263009 | FORTE TELECOM LTDA. | 3,369 |
| AS37468 | Angola Cables | 3,078 |
| AS51185 | Onecom Global Communications LTD | 3,059 |

# Before we start…..

We all live in glass houses

So we shouldn't throw stones

Offer to help and drop that rock….

# 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
  - Peering tools

# 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

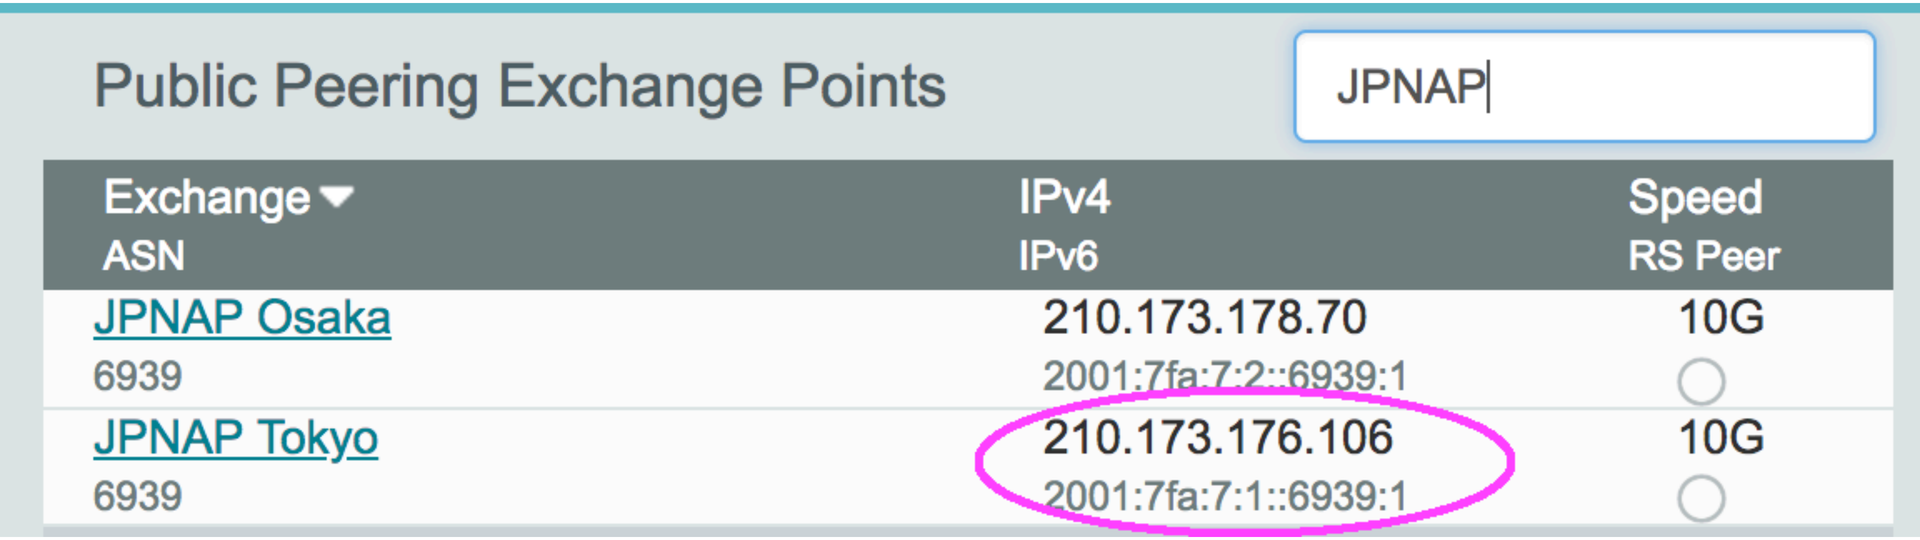

Public Peering Exchange Points

JPNAP

| Exchange<br>ASN | IPv4<br>IPv6 | Speed<br>RS Peer |
|---|---|---|
| JPNAP Osaka<br>6939 | 210.173.178.70<br>2001:7fa:7:2::6939:1 | 10G<br>○ |
| JPNAP Tokyo<br>6939 | 210.173.176.106<br>2001:7fa:7:1::6939:1 | 10G<br>○ |

# IXP Subnet

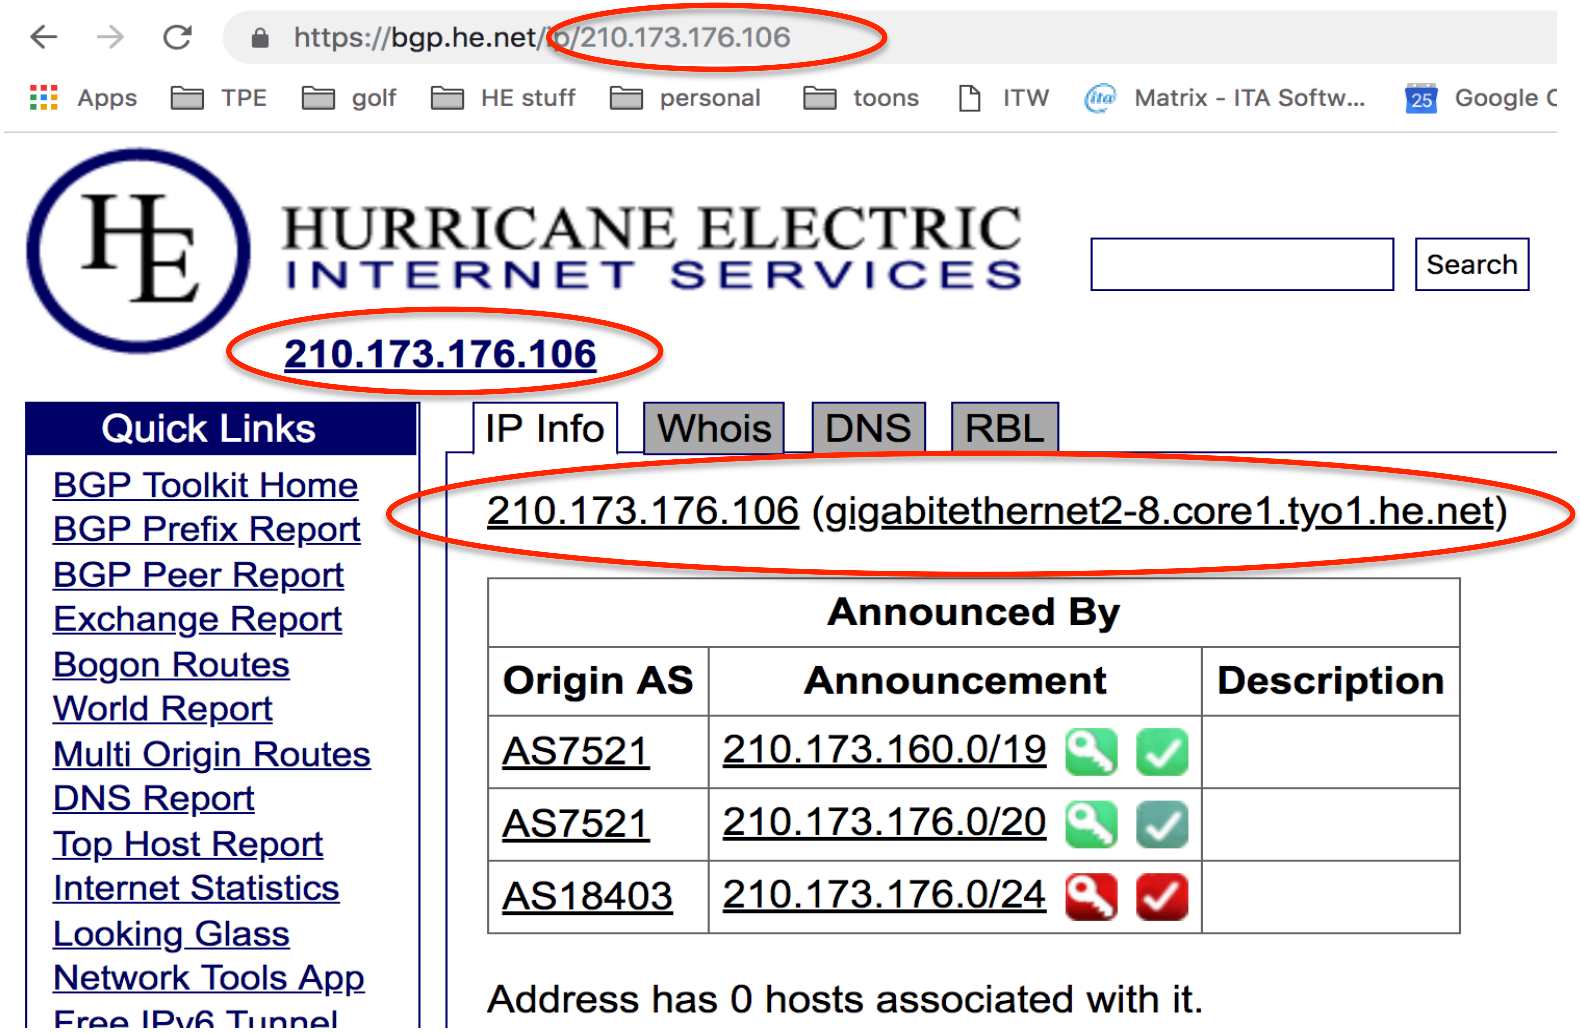

## 210.173.176.106

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

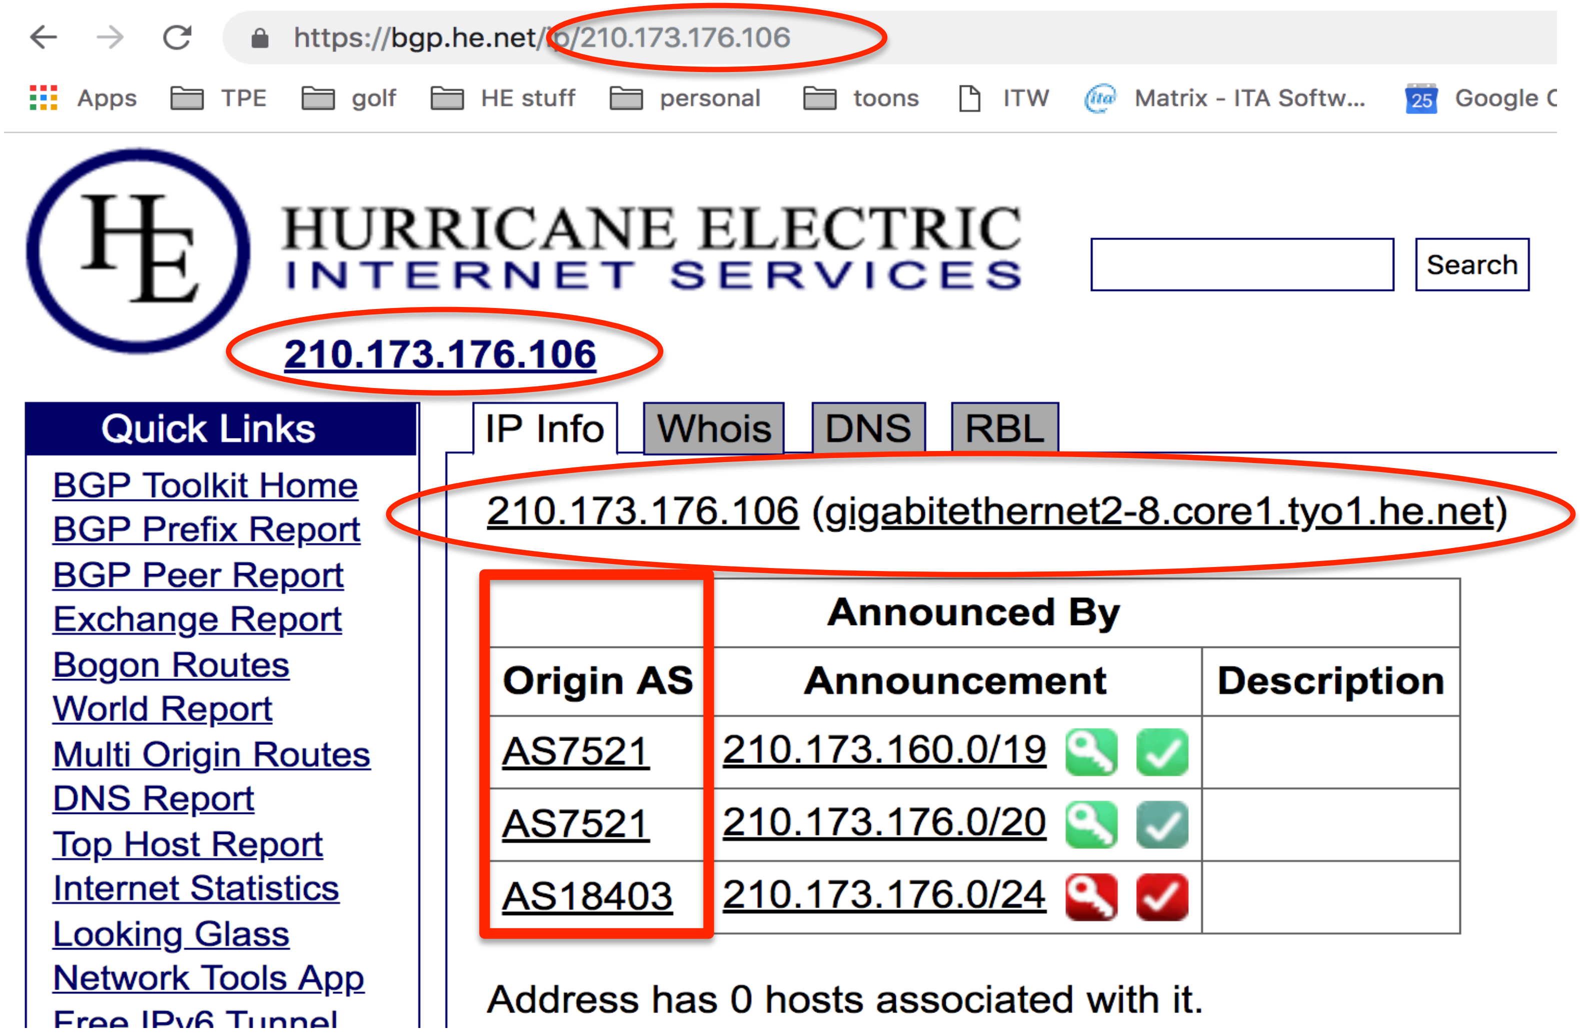

210.173.176.106

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

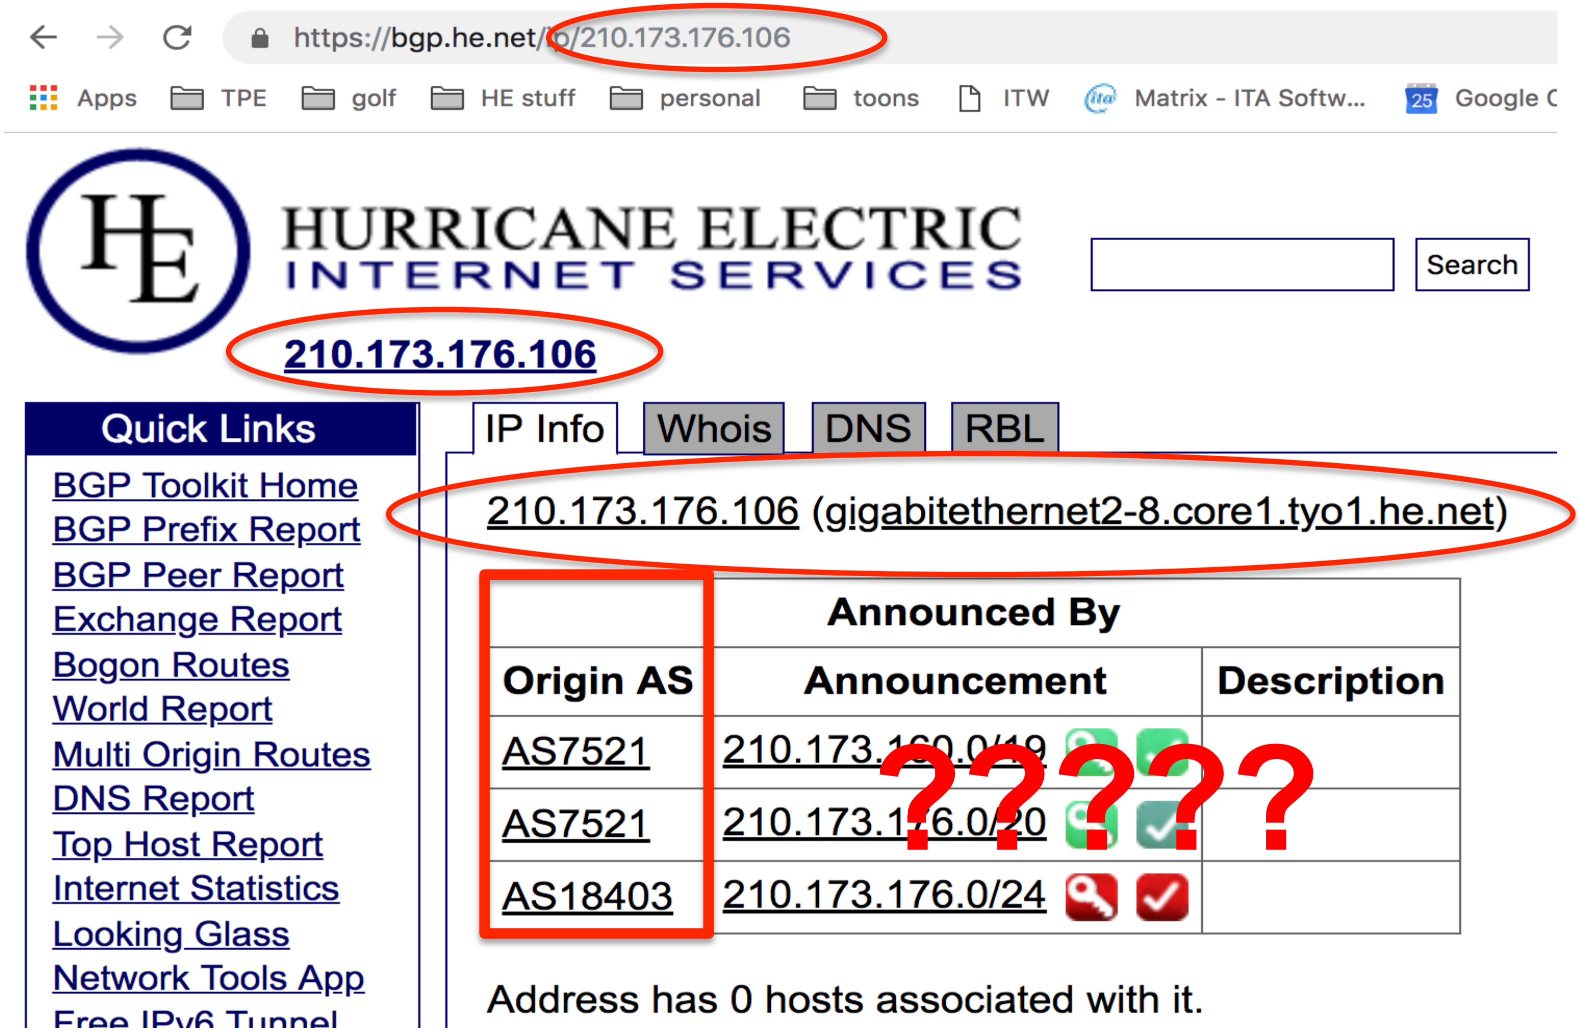

**210.173.176.106**

- 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

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

- The IX LAN is not your IP space and should not be routed.
- Some of the smaller guys

# IXP Subnet

Europe

| CC | Exchange | Speed | IPv4 | IPv6 |
|---|---|---|---|---|
| | VIX | 2x10GE | 193.203.0.185 | 2001:7f8:30:0:2:1:0:6939 |
| | BNIX | 2x10GE | 194.53.172.33 | 2001:7f8:26::a500:6939:1 |
| | B-IX Balkans | 10GE | 217.174.157.31 | 2001:7f8:8e::31 |
| | BIX.BG | 2x10GE | 193.169.198.70 | 2001:7f8:58::1b1b:0:1 |
| | NetIX | 10GE | 193.218.0.89 | 2001:67c:29f0::6939:1 |
| | MegaIX Sofia | 10GE | 91.212.235.55 | 2001:7f8:9f::a:6939:1 |
| | T-CIX Bulgaria | 10GE | 185.1.40.26 | 2001:7f8:98::26 |
| | CIXP | 10GE | 192.65.185.143 | 2001:7f8:1c:24a::1b1b:1 |

# IXP Subnet

- Some of the big ones…..

# IXP Subnet

Europe

```
CC Exchange             Speed   IPv4            IPv6
-- -------------------- ------- --------------- -----------------------

DE-CIX Frankfurt 2x100GE 80.81.192.172 2001:7f8::1b1b:0:1
France-IX Paris 2x10GE 37.49.236.10 2001:7f8:54::10
AMS-IX 2x100GE 80.249.209.150 2001:7f8:1::a500:6939:1
LINX 100GE 195.66.224.21 2001:7f8:4:0::1b1b:1
MSK-IX Moscow 2x100GE 195.208.210.40  2001:7f8:20:101::210:40
NL-IX 3x10GE 193.239.116.14 2001:7f8:13::a500:6939:1
```

# IXP Subnet

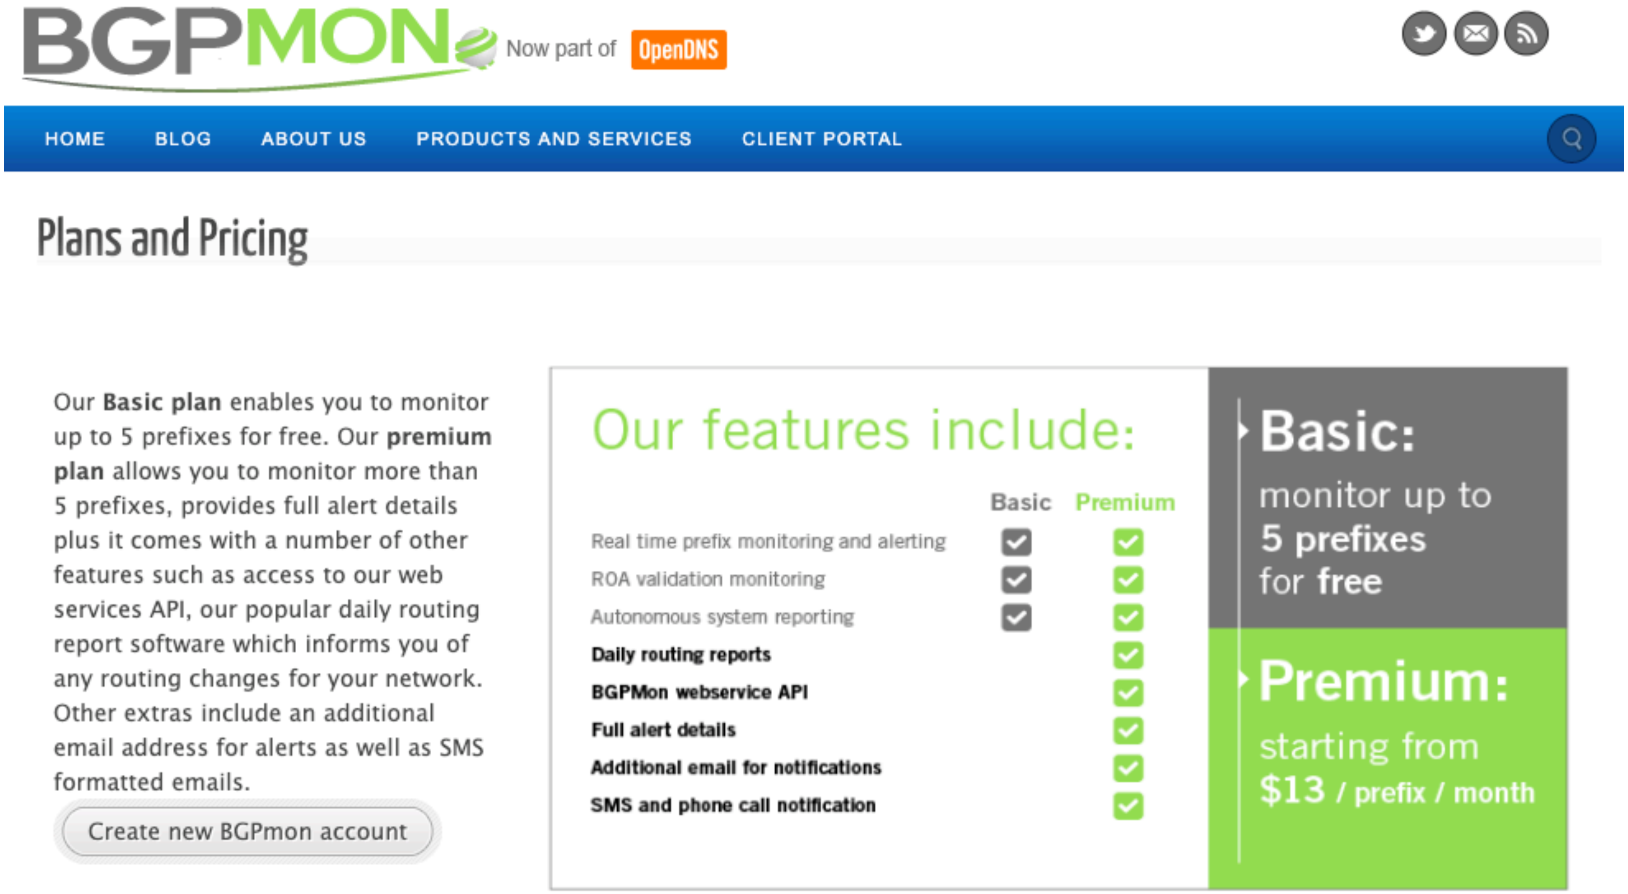

## Plans and Pricing

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

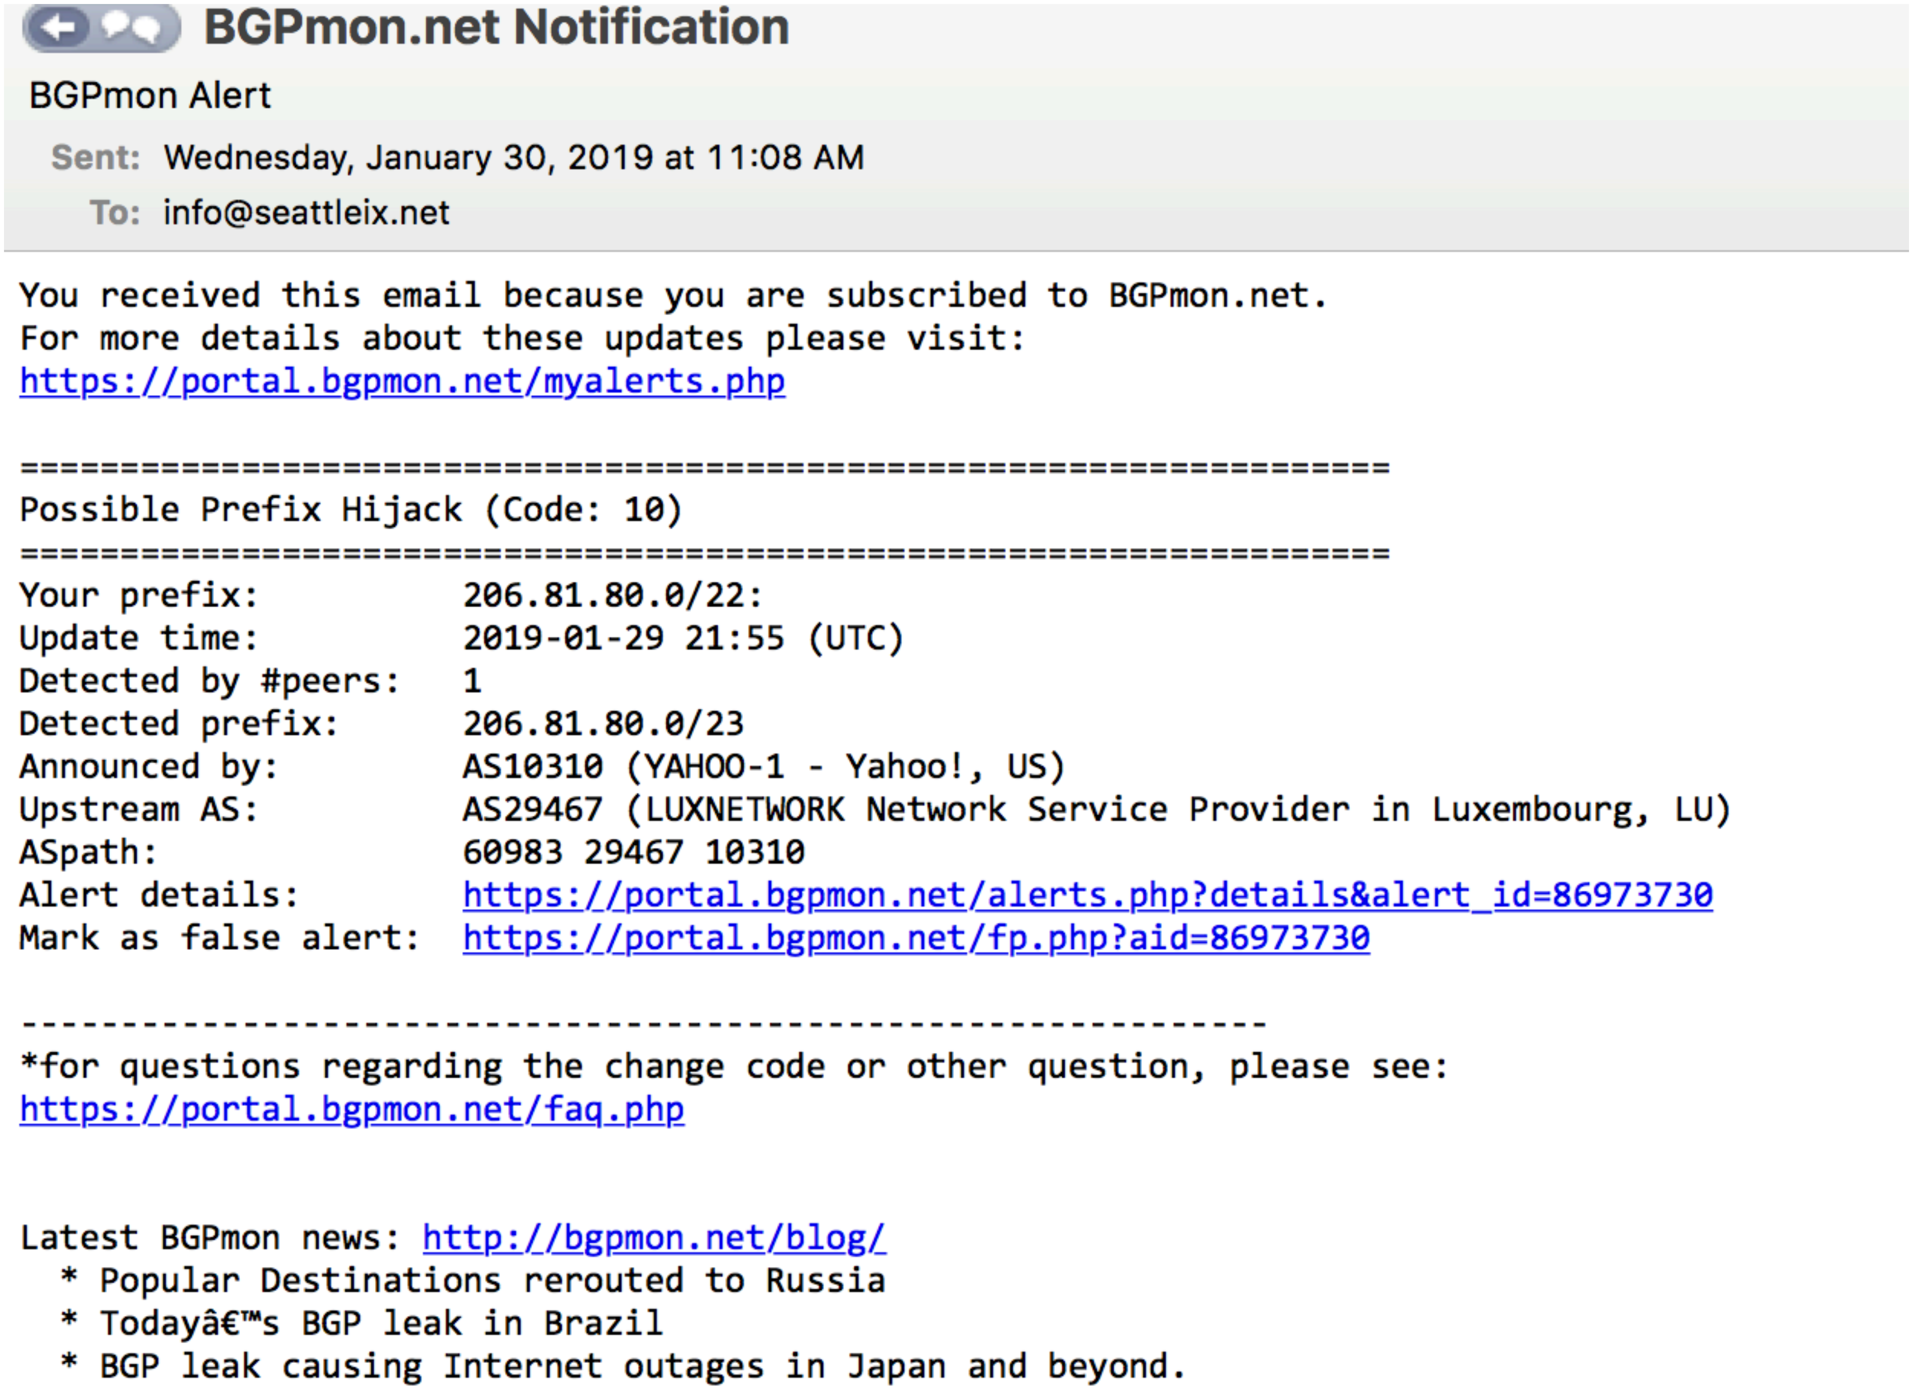

BGPmon.net Notification

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.
```

# BGPMON Replacement

https://mailman.nanog.org/pipermail/nanog/2019-August/102672.html

Thanks to Job & Massimo @NTT Ltd

# 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.

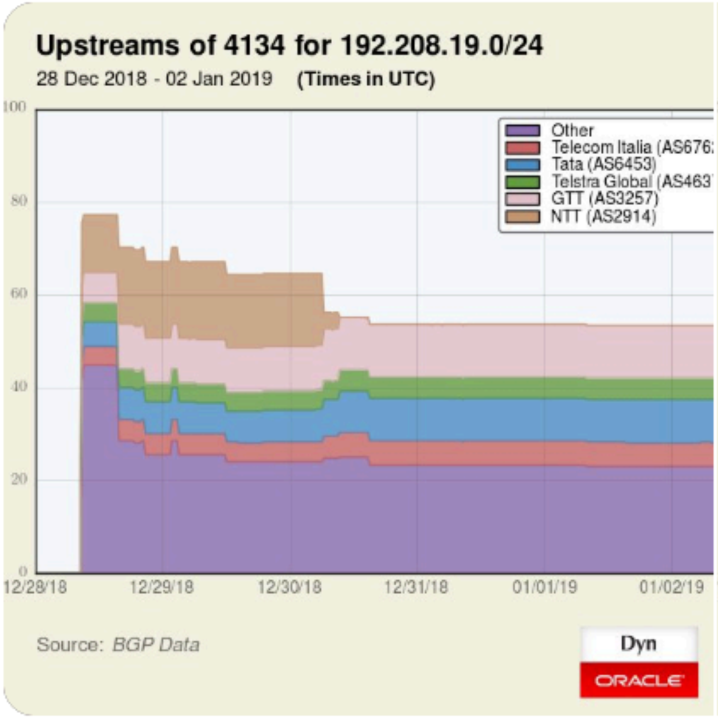

Upstreams of 4134 for 192.208.19.0/24

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

Source: *BGP Data*

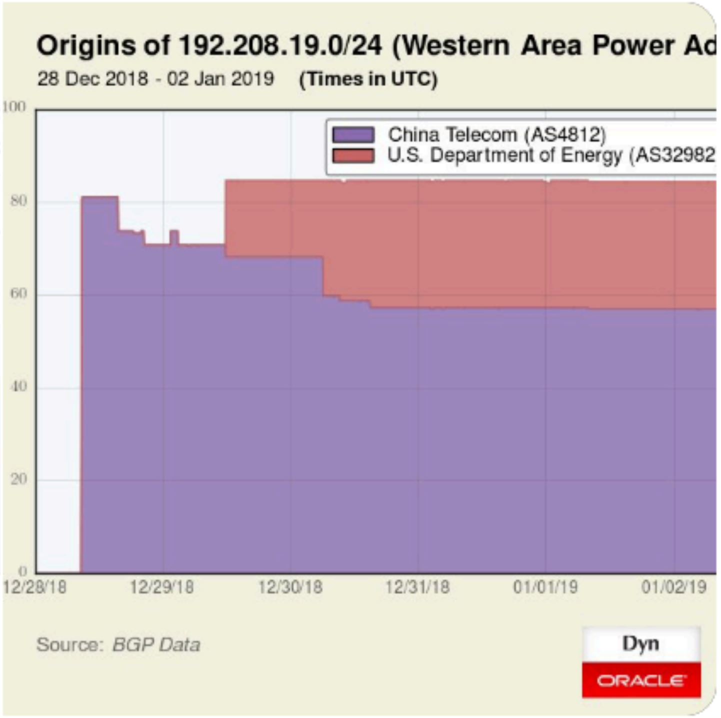

Origins of 192.208.19.0/24 (Western Area Power Ad

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

Source: *BGP Data*

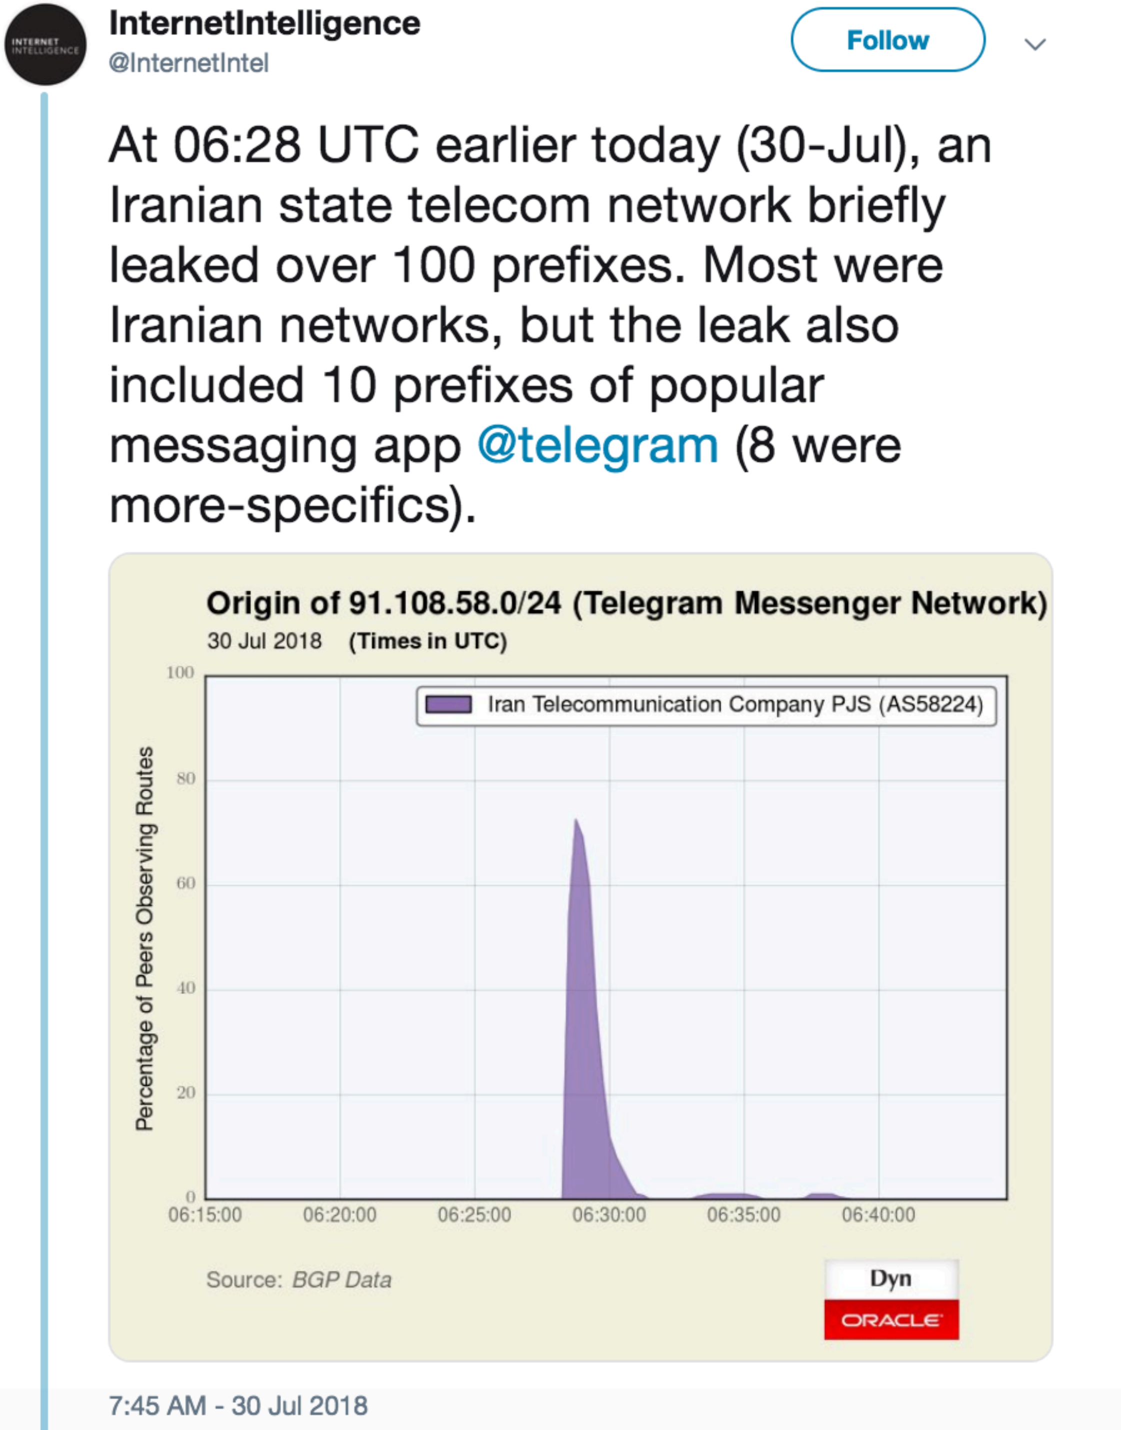

**InternetIntelligence**
@InternetIntel

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
Dyn
ORACLE

Source: *BGP Data*

7:45 AM - 30 Jul 2018

# https://bgpstream.com

- Every day there are several hijacks and leaks

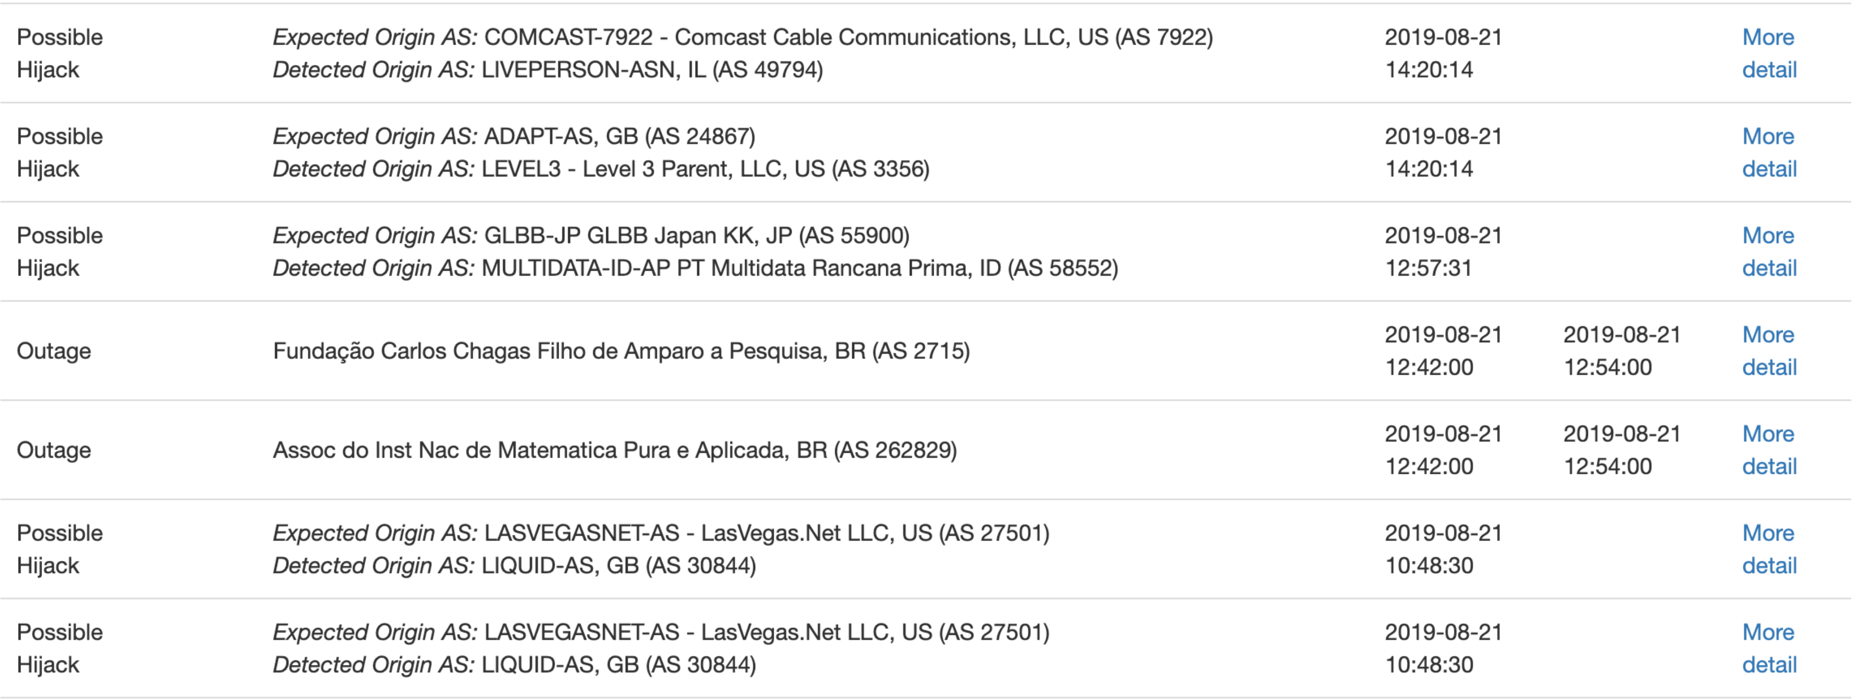

| | | | | |
|---|---|---|---|---|
| Possible Hijack | *Expected Origin AS:* COMCAST-7922 - Comcast Cable Communications, LLC, US (AS 7922)<br>*Detected Origin AS:* LIVEPERSON-ASN, IL (AS 49794) | 2019-08-21 14:20:14 | | More detail |
| Possible Hijack | *Expected Origin AS:* ADAPT-AS, GB (AS 24867)<br>*Detected Origin AS:* LEVEL3 - Level 3 Parent, LLC, US (AS 3356) | 2019-08-21 14:20:14 | | More detail |
| Possible Hijack | *Expected Origin AS:* GLBB-JP GLBB Japan KK, JP (AS 55900)<br>*Detected Origin AS:* MULTIDATA-ID-AP PT Multidata Rancana Prima, ID (AS 58552) | 2019-08-21 12:57:31 | | More detail |
| Outage | Fundação Carlos Chagas Filho de Amparo a Pesquisa, BR (AS 2715) | 2019-08-21 12:42:00 | 2019-08-21 12:54:00 | More detail |
| Outage | Assoc do Inst Nac de Matematica Pura e Aplicada, BR (AS 262829) | 2019-08-21 12:42:00 | 2019-08-21 12:54:00 | More detail |
| Possible Hijack | *Expected Origin AS:* LASVEGASNET-AS - LasVegas.Net LLC, US (AS 27501)<br>*Detected Origin AS:* LIQUID-AS, GB (AS 30844) | 2019-08-21 10:48:30 | | More detail |
| Possible Hijack | *Expected Origin AS:* LASVEGASNET-AS - LasVegas.Net LLC, US (AS 27501)<br>*Detected Origin AS:* LIQUID-AS, GB (AS 30844) | 2019-08-21 10:48:30 | | More detail |

# Basics - Routing Security

I know we can do better

# Basics - Routing Security

You must filter your peers!

# 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

- Prefix Count

Consider tightening up the limits with bgp neighbor restart/graceful

# Basics - Routing Security

## AS-Path

BBIX peer 各位 (Dear BBIX peering partners,)

さくらインターネット(AS9371)の津田です。
いつもお世話になっております。

弊社から広報しておりますAS Pathに変更が御座います。
AS Pathでのフィルタ設定が御座います場合、設定変更をお願い致します。

AS name: SAKURA-C
AS set: AS-SAKURA
AS number: 9371

▼追加するAS Path(IPv4)
^(9371_)+(2519_)+(9354_)+(10001_)+$
^(9371_)+(9370_)+(2519_)+(9354_)+(10001_)+$

# Basics - Routing Security

## Prefix list per neighbor

```
ip prefix-list AS57660 permit 37.26.208.0/20
ip prefix-list AS57660 permit 185.67.16.0/22
ip prefix-list AS57660 permit 212.67.48.0/20
```

# Basics - Routing Security

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 -J4l AS57660-IN AS57660
policy-options {
replace:
 prefix-list AS57660-IN {
    37.26.208.0/20;
    185.67.16.0/22;
    212.67.48.0/20;
 }
}
walt@staff:~$
```

# IXPs using RPKI

- IX.BR
- AMS-IX
- DE-CIX
- France-IX
- LINX
- Over 58 IXP today and more coming!

- **Downside is that not all networks peer on route servers**

- http://peering.exposed/

# http://routing.he.net

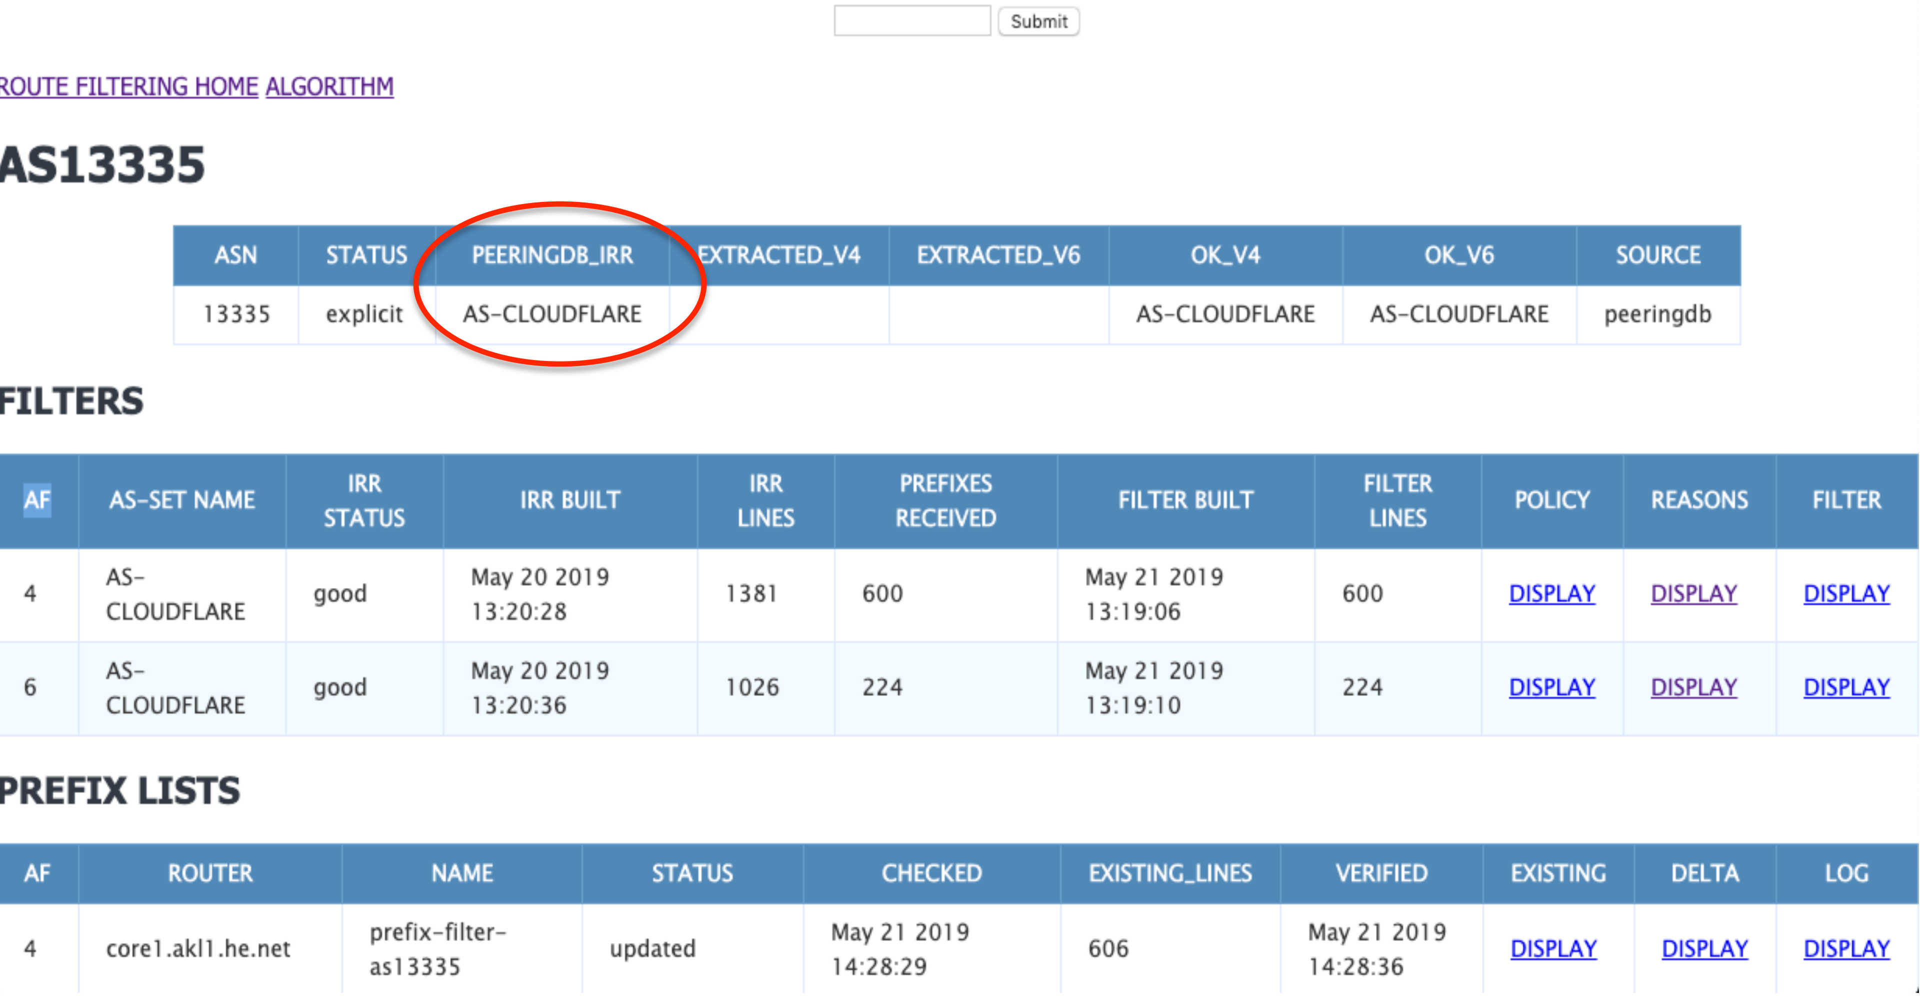

ROUTE FILTERING HOME ALGORITHM

# 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 | May 20 2019 13:20:28 | 1381 | 600 | May 21 2019 13:19:06 | 600 | DISPLAY | DISPLAY | DISPLAY |
| 6 | AS-CLOUDFLARE | good | May 20 2019 13:20:36 | 1026 | 224 | May 21 2019 13:19:10 | 224 | DISPLAY | DISPLAY | DISPLAY |

## PREFIX LISTS

| AF | ROUTER | NAME | STATUS | CHECKED | EXISTING_LINES | VERIFIED | EXISTING | DELTA | LOG |
|---|---|---|---|---|---|---|---|---|---|
| 4 | core1.akl1.he.net | prefix-filter-as13335 | updated | May 21 2019 14:28:29 | 606 | May 21 2019 14:28:36 | DISPLAY | DISPLAY | DISPLAY |

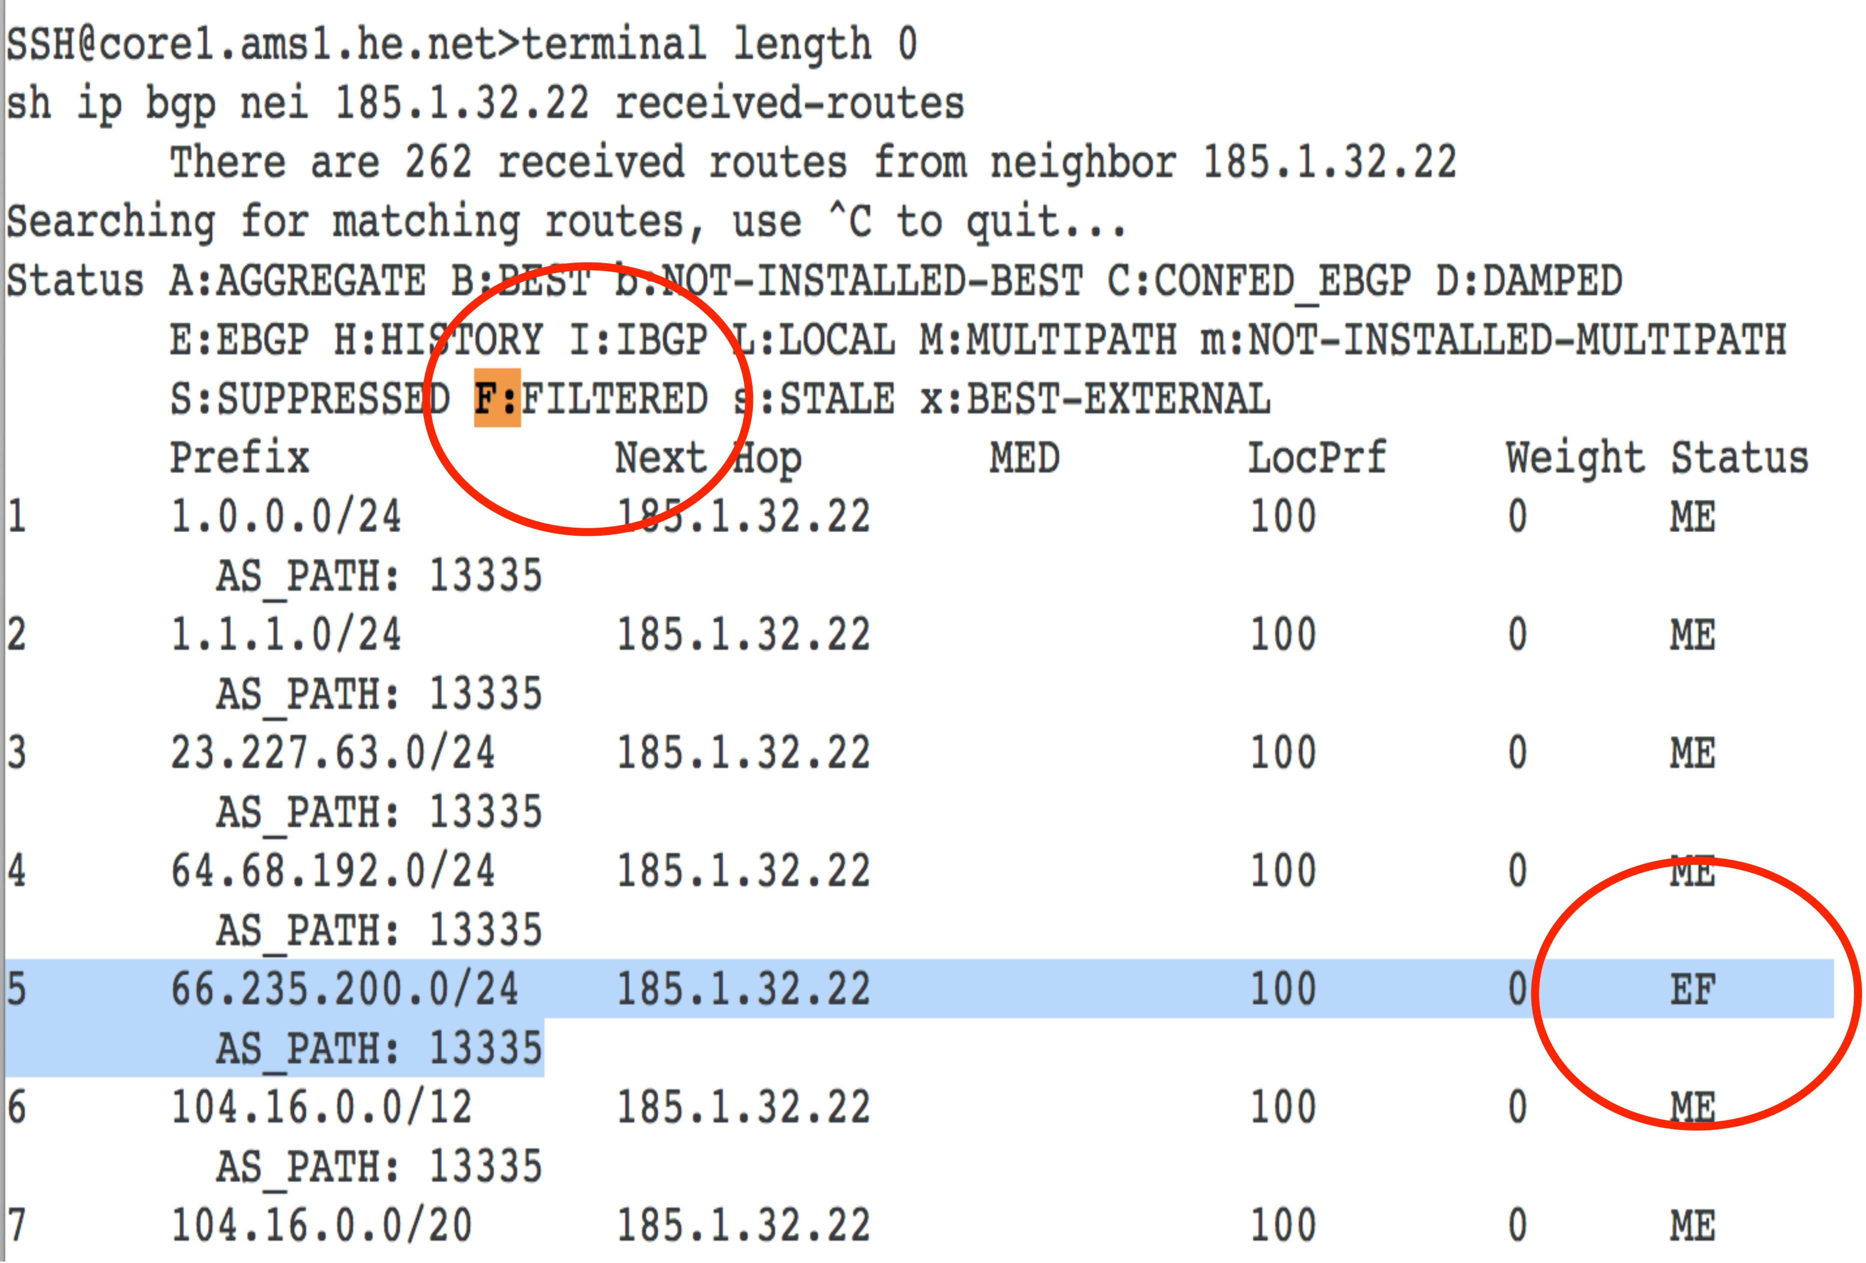

```
SSH@core1.ams1.he.net>terminal length 0
sh ip bgp nei 185.1.32.22 received-routes
       There are 262 received routes from neighbor 185.1.32.22
Searching for matching routes, use ^C to quit...
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST C:CONFED_EBGP D:DAMPED
       E:EBGP H:HISTORY I:IBGP L:LOCAL 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
```

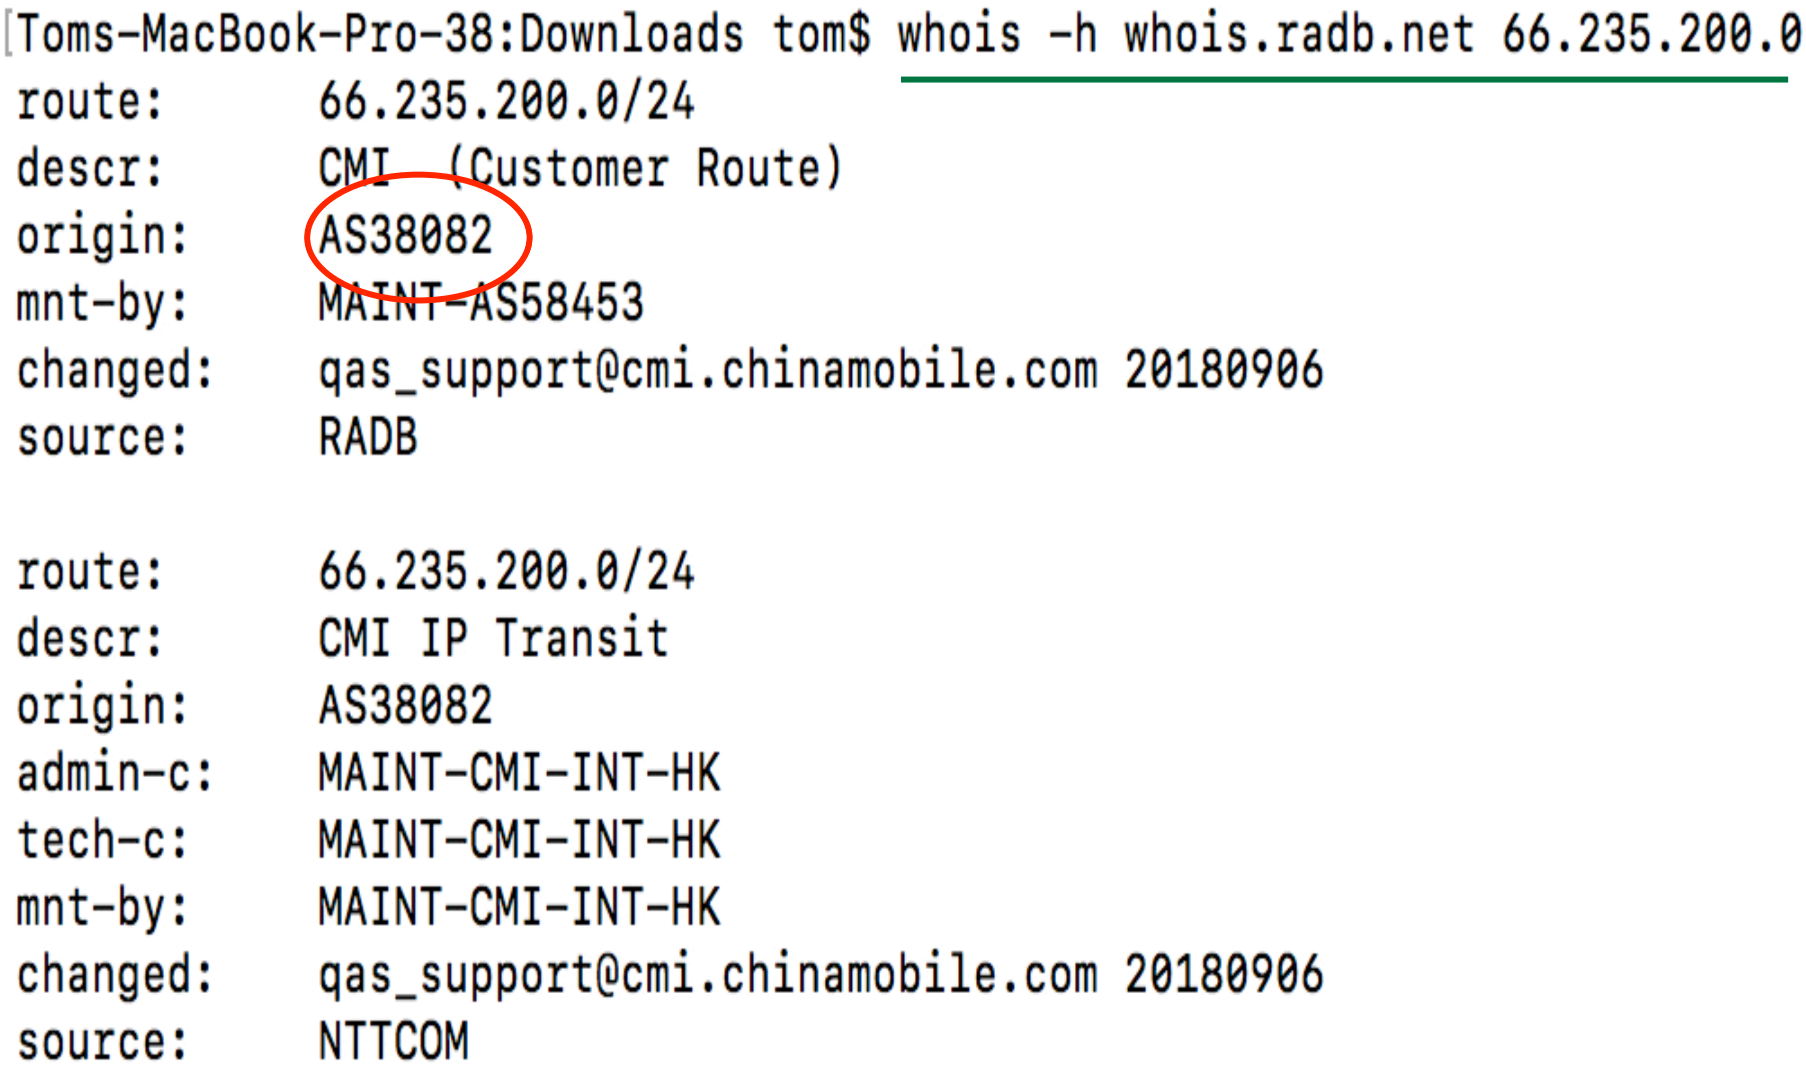

```
Toms-MacBook-Pro-38:Downloads tom$ whois -h whois.radb.net 66.235.200.0
route:          66.235.200.0/24
descr:          CMI  (Customer Route)
origin:         AS38082
mnt-by:         MAINT-AS58453
changed:        qas_support@cmi.chinamobile.com 20180906
source:         RADB

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

# Resources

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

HURRICANE ELECTRIC
INTERNET SERVICES

# Thanks!

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