RouteViews

BGP Global Monitoring Infrastructure and Services

RouteViews

A collaborative router looking glass to share BGP views among network operators and researchers.

RouteViews was founded at the University of Oregon's Advanced Network Technology Center (ANTC) in 1995. Data archives began in 1997 and amount to 50TBs (compressed) today.

The group is currently led by the Network Startup Resource Center (NSRC) group engineering team at the University of Oregon.

RouteViews collaborates with the Center for Applied Internet Data Analysis (CAIDA) working with NSF grants that support *Designing a Global Measurement Infrastructure to Improve Internet Security*, GMI3S (OAC-2131987), and an *Integrated Library for Advancing Network Data Science,* ILANDS (CNS-2120399).

RouteViews is supported with financial and in-kind donations by multiple additional organizations:

https://www.routeviews.org/routeviews/index.php/supporters/

NSRC

NSRC supports the growth of global Internet infrastructure by providing engineering assistance, collaborative technical workshops, training, and other resources to university, research & education networks worldwide. NSRC is partially funded by the IRNC program of the NSF (OAC-2029309) and Google with other contributions from public and private organizations.

UNIVERSITY OF OREGON

The University of Oregon is a public research institution in Eugene, Oregon, USA founded in 1876. UO is renowned for its research prowess and commitment to teaching. Both NSRC and RouteViews are based at the UO.

Why RouteViews?

It's YOUR Internet

- Originally conceived in 1995 as a tool for Internet operators to look at the BGP table from different backbones and locations around the world to troubleshoot and to assess:
  - Reachability, hijacks, peer visibility, mass withdrawals, and RPKI status
- Operators who find it a valuable tool also peer to contribute to the value
- The 26-year data-set of BGP information archived by RouteViews since 1997 has become an invaluable research resource
  - RouteViews data has been used in over 1000 research papers.
  - http://www.routeviews.org/routeviews/index.php/papers/

RouteViews Collector Map

http://www.routeviews.org/routeviews/index.php/map/

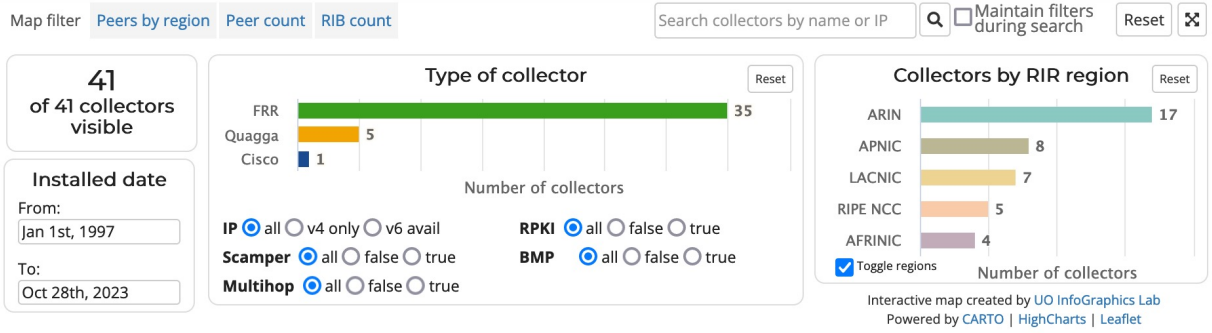

Collector Deployment

Locations

- At Internet Exchange Points
  - Or DataCentres hosting major interconnections
- Goal is to maximise opportunities for Autonomous Networks to peer with the RouteViews collector
- There are also multi-hop collectors hosted at the University of Oregon

Collector

- One interface on the IXP peering LAN
  - Peering with operators who wish to supply a BGP feed to the RouteViews collector
- One interface for management and backhaul
  - Managing the collector
  - Sending the collected BGP table & BGP updates for archival at the University of Oregon

Collectors

Hardware

Commodity

- 1RU
- 8-16 Cores
- 32GB RAM
- 400GB-1TB SSD
- 2x 10 GB ethernet

Shipped to the IX

- Last resort only

Hosted

Virtual Machine

- 4 vCPU
- 32GB RAM
- 200GB Storage
- Peering Interface
- Transit Interface

Preferred option

- Quicker to deploy
- Easier to maintain
- Easier to update

Software

OpenSource

- Ubuntu LTS with FRR

- Legacy: CentOS with Quagga

Collector Deployment

Multi-hop

- Advantages
  - If you can reach the collector, you can peer.
- Disadvantages
  - Multi-hop peerings are subject to the routing anomalies RouteViews seeks to observe and archive.

IX hosted/collocated

- Advantages
  - Better positioned to address multi-hop issues.
  - Geographic diversity.
  - Peering diversity.
  - Scalable.
- Disadvantages
  - More infrastructure to manage.

Collector Data

Multi-Threaded Routing Toolkit (MRT)

- https://tools.ietf.org/html/rfc6396
- MRT provides a standard for parsing or dumping routing information to a binary file.
- RouteViews Dumps consist of BGP RIBs and UPDATES
  - RIBs are archived every 2 hours
  - UPDATEs are archived every 15 minutes

Data Access

- MRT files are bzipped and rsynced back to http://archive.routeviews.org/ on above schedule
- They can be accessed via http, ftp and rsync
- Map view tool is interactive

Direct BGP Monitoring Protocol Feed (BMP)

- New Model, in testing right now
- BMP upstream from collectors, not FROM peers

Collector Data

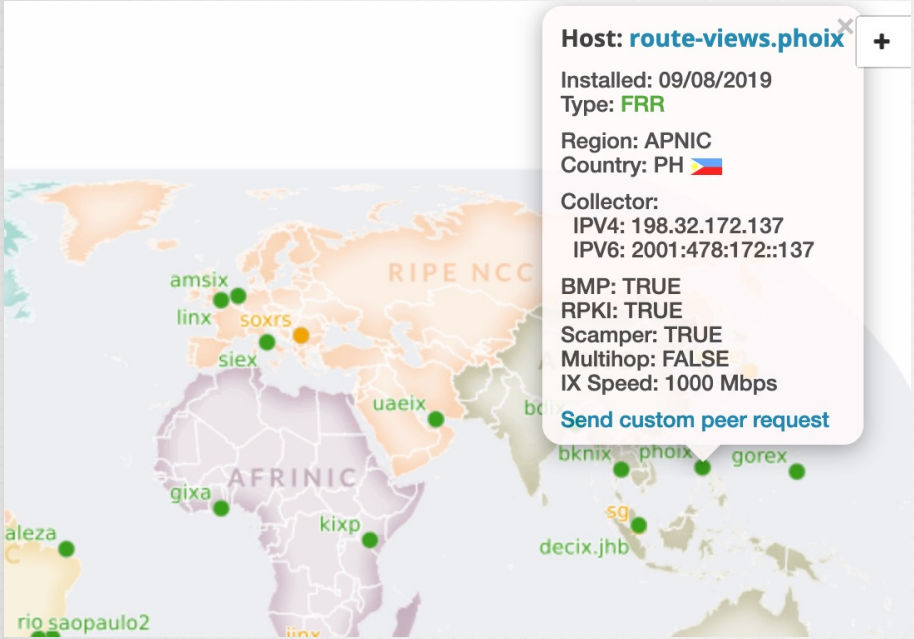

Index of /route-views.phoix/bgpdata

Name	Last modified	Size	Description
Parent Directory		-	
2019.08/	2019-08-08 22:58	-	
2019.09/	2019-08-28 23:01	-	
2019.10/	2019-09-28 23:01	-	
2019.11/	2019-10-28 23:01	-	
2019.12/	2019-11-28 23:01	-	
2020.01/	2019-12-28 23:01	-	
2020.02/	2020-01-28 23:01	-	
2020.03/	2020-02-28 23:01	-	
2020.04/	2020-03-28 23:01	-	
2020.05/	2020-04-28 23:01	-	

http://archive.routeviews.org/route-views.phoix/bgpdata/

Peering HowTo

BGP Configuration

- Please send full-table if you are able to
- Please do not send:
  - Default route
  - Private address space
- We do not accept/want ADD-PATH TX/RX
- We do not send any routes to you
- When peering with multi-hop collectors, remember to set ebgp-multihop

MRT Tools

BGPKIT, RIPE LIBBGPDUMP, NTT BGPDUMP2, etc

- https://bgpkit.com/ (BGP toolkit written in Rust)
  - https://bgpkit.com/parser
- https://github.com/bgpkit
  - https://github.com/bgpkit/bgpkit-parser
  - https://github.com/bgpkit/peer-stats
  - https://github.com/bgpkit/pybgpkit

- https://github.com/RIPE-NCC/bgpdump (Last updated 2020)
- https://github.com/cawka/bgpparser (Last updated 2015 ☹)
- https://github.com/yasuhiro-ohara-ntt/bgpdump2
- https://github.com/t2mune/mrtparse (python)
- https://github.com/rfc1036/zebra-dump-parser (perl) (Last updated 2014 ☹)

Route Collection: Global Community Cooperation

RouteViews Collector Map

Extends coverage
Provides Redundancy
Complimentary Services
Common files/tools

RIS Collector Map

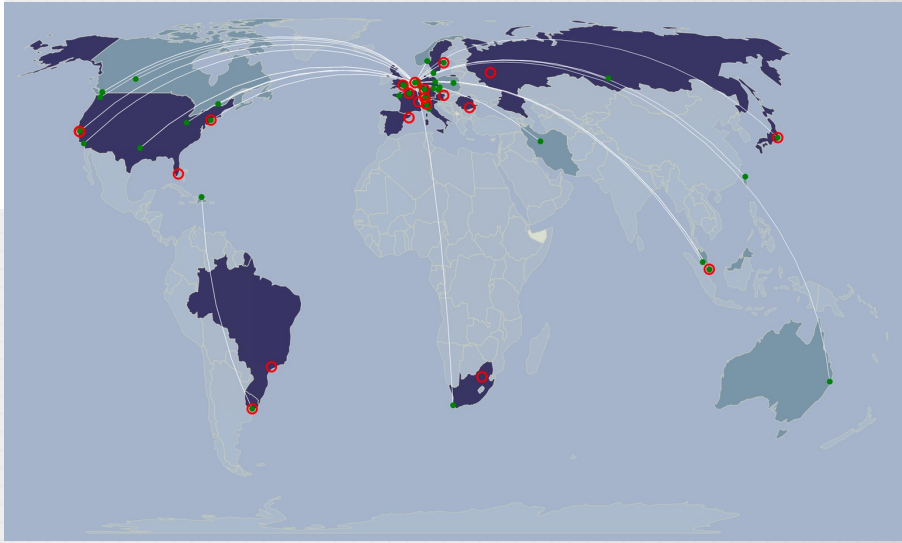

Operations

Access

- BGP is the backbone of the Global Routing Infrastructure.
- To ensure its stability, it needs to be constantly monitored.
- RouteViews provides:
  - Command-Line/Looking Glass
  - Prefix Visibility, Verify Convergence, Path Stability
  - Comparing Local/Regional/Global Views
  - Troubleshooting Reachability
  - Access to historical BGP data, i.e. “When did this happen??”

Operations

Accessing a Collector

- `telnet://route-views*.routeviews.org`
- No username necessary.
- Users are able to run show commands, e.g. `show ip bgp x.x.x.x`
- Telnet access is rate-limited to prevent automation overuse
- PLEASE don't script. There is an API for that (http://api.routeviews.org).

Gotchas

- Why not SSH?!
  - RouteViews data are publicly available. We've got nothing to hide.
- `show ip route x.x.x.x` next-hop is incorrect!
  - Remember, this is a collector
  - There's no data-plane, thus no true FIB
  - Kernel default-route points to transit provider next-hop

Operations

CLI Access

- Worldwide CLI access – how to access a collector
- telnet://route-views.routeviews.org
  - route-views, route-views{2,3,4,6} are all housed at University of Oregon in the United States, and each collector has eBGP Multihop sessions with peers from around the world

- Legacy Naming Scheme
  - telnet://route-views.phoix.routeviews.org
  - Other collector locations accessible via these 3rd level domains (replace “phoix”): amsix, eqix, telxatl, bknix, soxrs, chicago, bdix, uaeix, fortaleza, gixa, gorex, mwix, jinx, jhb, napafrica, peru, linx, flix, kixp, ny, isc, perth, nwax, rio, siex, sfmix, chile, saopaulo, sg, sydney & wide

- New Naming Scheme
  - (exchange).(closest airport).routeviews.org
  - ie: decix.jhb.routeviews.org – DE-CIX Asia, Johor Bahru

Operations

Looking Glass

- Looking Glass front-end as an alternative access is being tested
- This will likely replace the telnet access in the near future

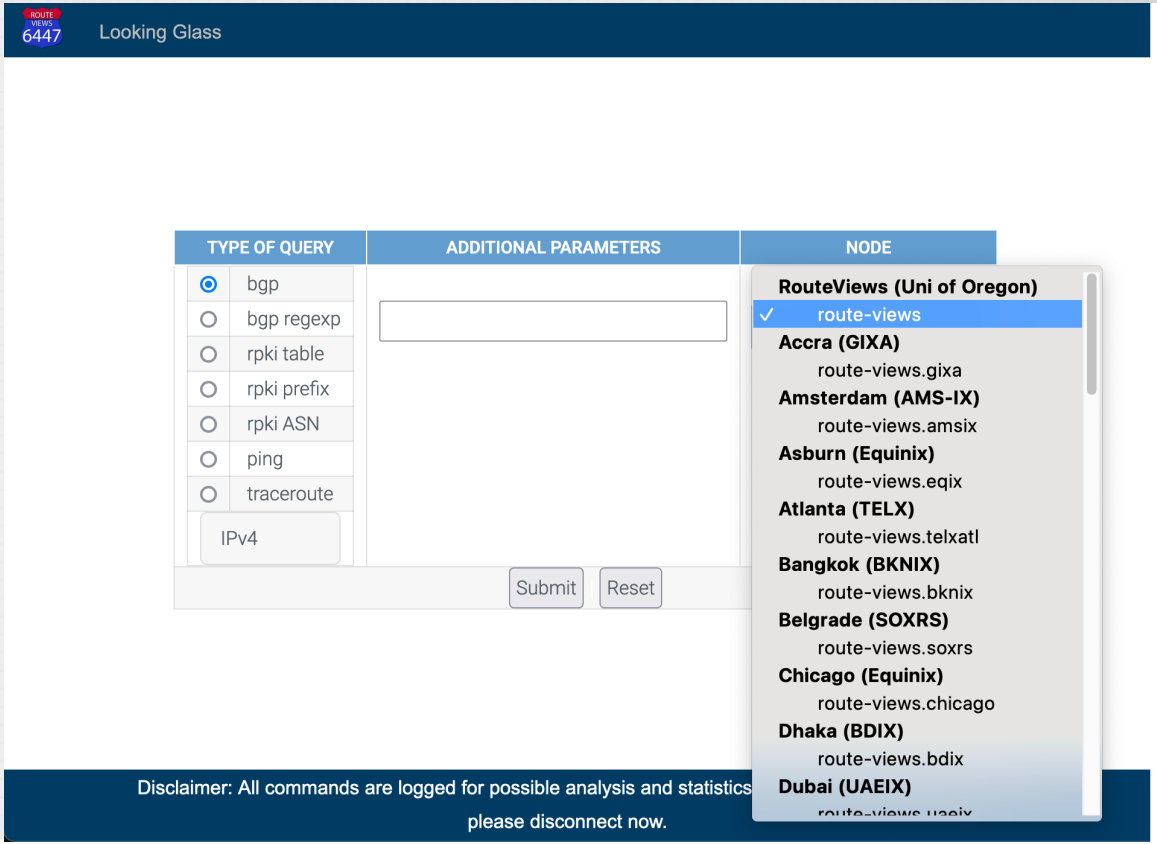

Use Cases

New(er) Collector Features

BMP

- BMP data will feed tools like BGPStream and ARTEMIS.
- Or write your own Kafka consumer for raw BMP data.
- Limited access during trial period.
- Wider availability to follow on completion of trials.

RPKI

- RPKI data is accessible directly from the collectors.
- We also are establishing an archive of RPKI ROA data.
- Working on back-filling that data set from the RIR/CAs.

Use Cases

BGP Data Distribution

- “Message-based” data distribution (Kafka)
  - Real-time streaming telemetry
  - Per-message timestamps, with meta-data
  - Middle-layer abstraction, multi-client access (facilitates analysis and services)
  - Automated consolidating and sequencing
- RPKI validation and archival

BMP & Open/GoBMP

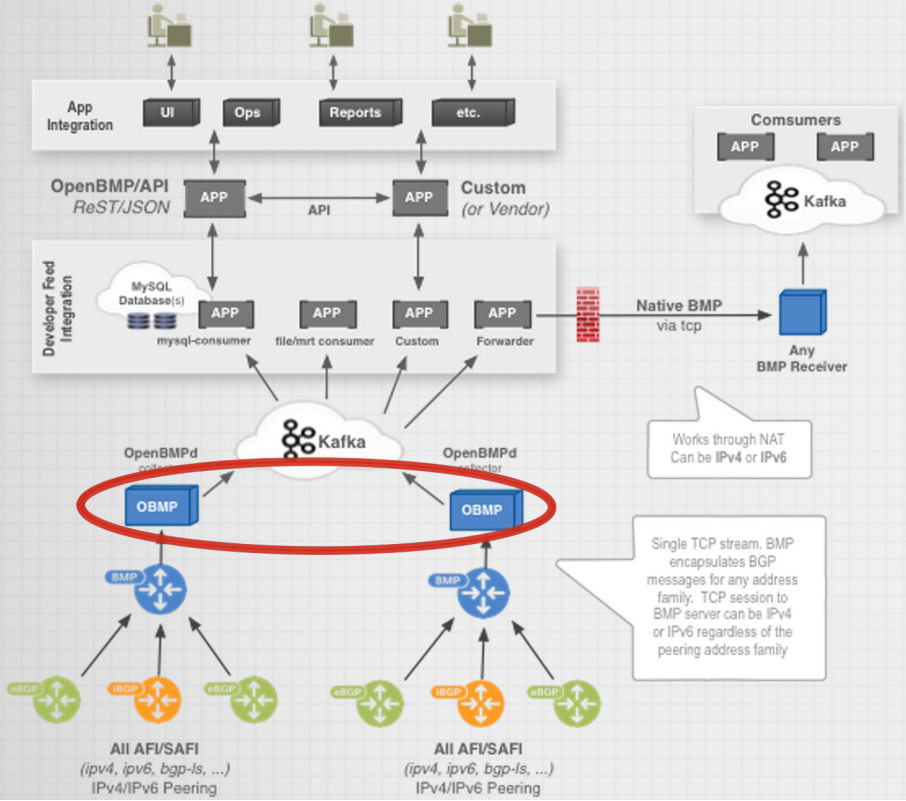

- BGP Monitoring Protocol (BMP) is an IETF standard: RFC7854
  - Available now – (Cisco, Juniper, Arista & FRR)
  - Consolidates peers/collectors
  - Splits collector, peer and update messages into separate streams
- OpenBMP is OpenSource (under the Linux Foundation)
  - Latest update: 2018 😕
- GoBMP
  - Latest update: September 2022 😀

Apache Kafka

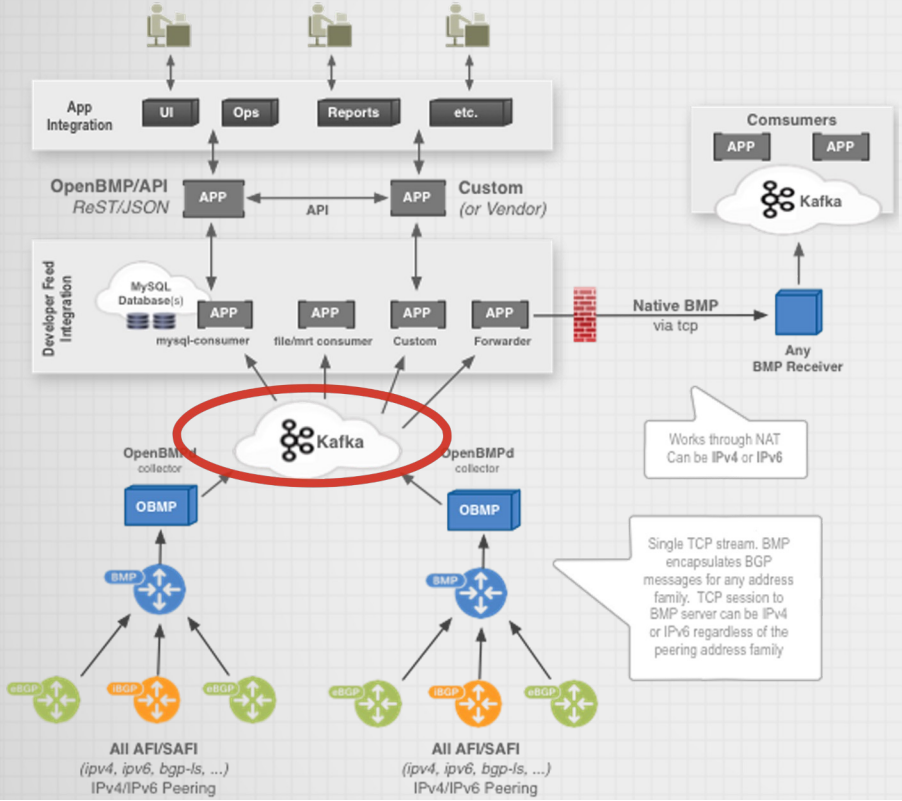

Apache Kafka comprises the message bus for OpenBMP

- Proven to scale
- Mature client API
  - Clients in 16 different programming languages

- Send email to:
  - help@routeviews.org

Use Cases

BGPStream is a project of CAIDA group at UC San Diego:
https://bgpstream.caida.org

BGPSTREAM

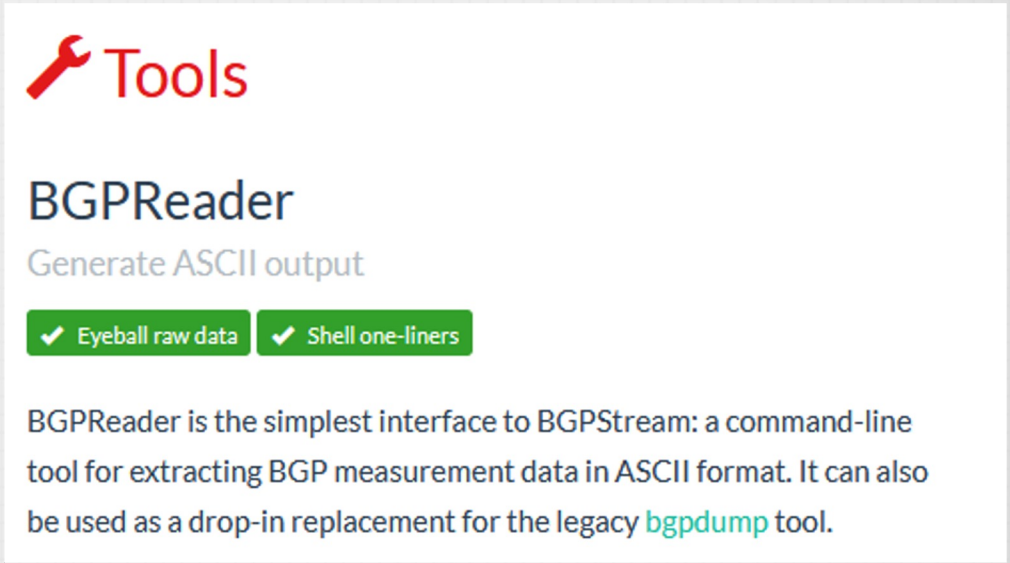

Use Cases

BGPStream is a project of CAIDA group at UC San Diego:
https://bgpstream.caida.org

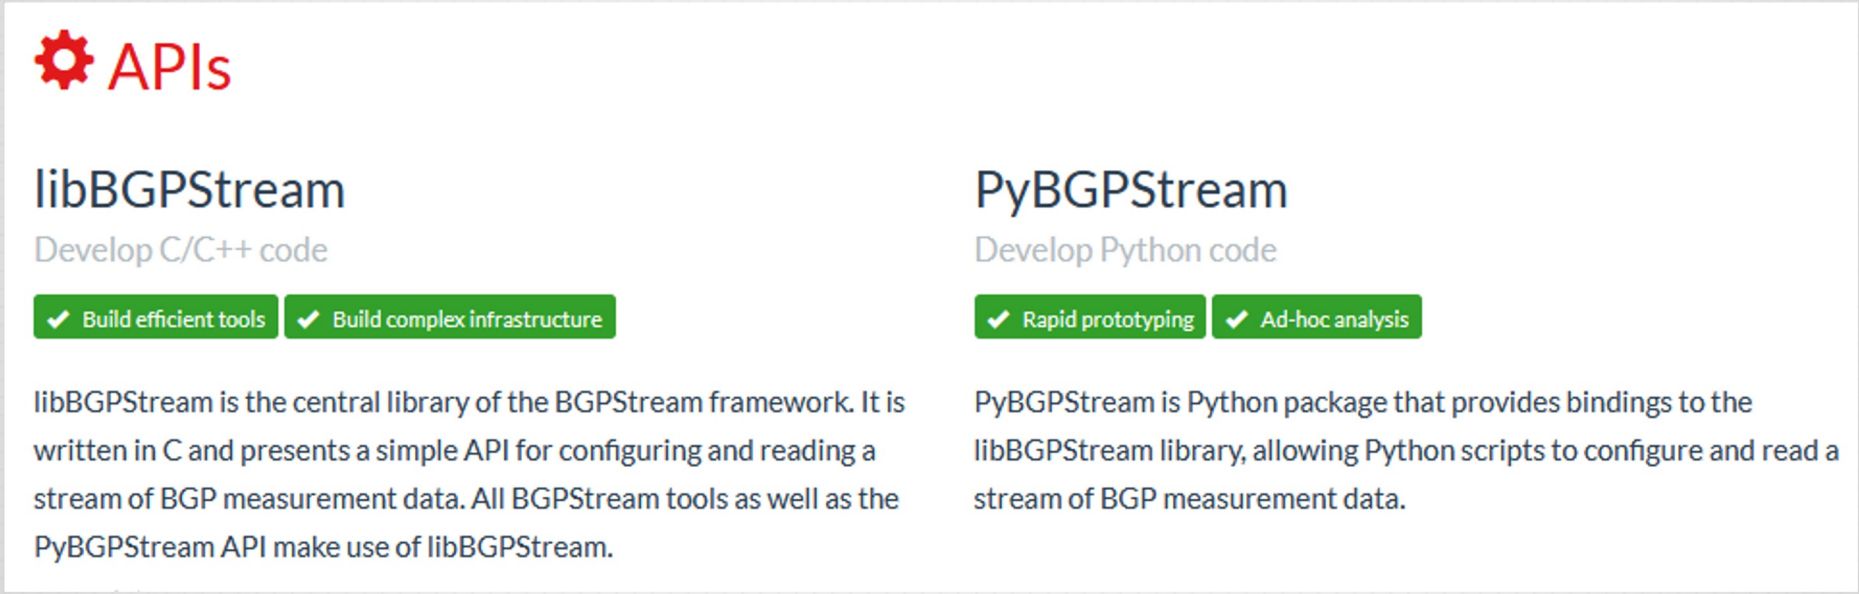

Use Cases

Other Open Source Tools: Artemis

- An open-source tool to monitor, detect, and mitigate BGP hijacks
- Real-time detection and notifications of BGP prefix hijacking attacks/events
- https://github.com/FORTH-ICS-INSPIRE/artemis

Use Cases

Other Open Source Tools: BGPKIT

- BGPKIT Parser
  - Rust-based MRT/BGP Data Parser
- BGPKIT Broker
  - REST API for searching archive files across public data collection projects. Data updated in real-time.
- BPPKIT Monocle
  - A commandline application to search, parse, and process BGP information in public sources

- https://bgpkit.com/

Use Cases

Not so Open Source Tools…

TECH

Cisco acquires ThousandEyes for around $1 billion to make deeper push into software

PUBLISHED THU, MAY 28 2020•4:38 PM EDT | UPDATED THU, MAY 28 2020•5:51 PM EDT

Ari Levy
@LEVYNEWS

Jordan Novet
@JORDANNOVET

SHARE

Use Cases

Operations

77 peers, multi-hop

Not all peers are up.. ☹

```
route-views2.routeviews.org> sh ip bgp sum

IPv4 Unicast Summary (VRF default):
BGP router identifier 128.223.51.102, local AS number 6447 vrf-id 0
BGP table version 53814890
RIB entries 1867939, using 342 MiB of memory
Peers 77, using 54 MiB of memory

Neighbor        V         AS   MsgRcvd   MsgSent   TblVer  InQ OutQ  Up/Down State/PfxRcd   PfxSnt Desc
4.68.4.46       4       3356         0     36081        0    0    0    never       Active        0 Level3
5.101.110.2     4      14061         0         0        0    0    0    never      Connect        0 DIGITALOCEAN
12.0.1.63       4       7018  22762421     72706        0    0    0 2d13h50m       918220        0 ATT
37.139.139.17   4      57866  15583523    146075        0    0    0 4d03h24m       920755        0 Fusix
43.226.4.1      4      63927         0         0        0    0    0    never      Connect        0 Rise
45.61.0.85      4      22652  21203029    146091        0    0    0 4d03h25m       921946        0 FIBRENOIRE
62.115.128.137  4       1299  63541287     73041        0    0    0 4d03h23m       903897        0 Telia
64.71.137.241   4       6939  20081396     73037        0    0    0 13:06:54       943729        0 Hurricane Electric
64.71.255.61    4        812         0         0        0    0    0    never      Connect        0 Sprint
66.185.128.1    4       1668         0         0        0    0    0    never      Connect        0 AOL
67.219.192.18   4      19653         0         0        0    0    0    never       Active        0 CTSTelecom
68.67.63.245    4      22652         0         0        0    0    0    never       Active        0 FIBRENOIRE
80.241.176.31   4      20771         0         0        0    0    0    never      Connect        0 CAUCASUS
85.114.0.217    4       8492  25929838    146099        0    0    0 01w1d10h       932986        0 OBITRU
87.121.64.4     4      57463   5292260     72989        0    0    0 4d03h16m       443778        0 NETIXLTD
89.149.178.10   4       3257  21296299     73041        0    0    0 4d03h22m       918375        0 Tiscali
91.209.102.1    4      39756         0         0        0    0    0    never      Connect        0 HOSTWAY-RO
91.218.184.60   4      49788  28486819     73045        0    0    0 01w1d03h       923119        0 NEXTHOPNO
```

Lots of full tables

Use Cases

Operations

```
BGP routing table entry for 41.77.223.0/24, version 53271948
Paths: (29 available, best #2, table default)
  Not advertised to any peer
  23673 6939 33763 37447
    203.189.128.233 from 203.189.128.233 (203.189.128.233)
      Origin IGP, valid, external, rpki validation-state: not found
      Community: 23673:50 23673:4636 23673:65401
      Last update: Sat Oct 28 00:55:33 2023
…
  3257 37468 37273 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447
 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447
 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447
 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447
 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447 37447
    89.149.178.10 from 89.149.178.10 (213.200.83.26)
      Origin incomplete, metric 10, valid, external, rpki validation-state: not found
      Community: 3257:2993 3257:4000 3257:8027 3257:50001 3257:50110 3257:54400 3257:54401 65000:174 65000:1299 65000:3356 65000:4134
 65000:10429
      Last update: Fri Oct 27 01:38:19 2023
…
  1299 6939 6939 33763 37447
    62.115.128.137 from 62.115.128.137 (2.255.254.88)
      Origin IGP, valid, external, rpki validation-state: not found
      Last update: Wed Oct 25 14:02:01 2023
```

What is AS37447 trying to achieve by prepending 97 times??

Use Cases

Invalid ROAs

AS64503 – this is RFC5398 Documentation ASN!

```
I*  36.90.220.0/23    203.189.128.233                0 23673 55329 9304 7713 64503 i
I*                    212.66.96.126                  0 20912 49367 6762 7713 64503 i
I*                    94.156.252.18          0       0 34224 6762 7713 64503 i
I*>                   85.114.0.217                   0 8492 3216 7713 64503 i
```

AS64021 & AS137451 are origins?

```
I*  43.255.37.0/24    203.189.128.233                0 23673 55329 6939 7473 9381 64021 i
I*                    64.71.137.241                  0 6939 7473 9381 64021 i
I*                    12.0.1.63                      0 7018 2914 64050 137451 i
I*                    105.16.0.247                   0 37100 2914 64050 137451 i
I*                    212.66.96.126                  0 20912 3257 2914 64050 137451 i
I*                    94.156.252.18          0       0 34224 6453 9381 64021 ?
I*>                   129.250.1.71           0       0 2914 64050 137451 i
...
```

Use Cases

Invalid ROAs

```
Whois says:

% whois -h jwhois.apnic.net 43.255.37.0/24
…
inetnum:        43.0.0.0 - 43.255.255.255
netname:        APNIC-AP-ERX
descr:          Asia Pacific Network Information Center, Pty. Ltd.
descr:          Regional Internet Registry for the Asia-Pacific Region
…
remarks:        This /8 was transfered from ARIN to APNIC on 30 Oct 2007
remarks:        The original message for this range is as below.
remarks:        Japan Inet has advised IANA to delegate the management of this /8 to APNIC.
remarks:        OrgName: Japan Inet
remarks:        OrgID: JAPANI
remarks:        Address: Kokusai-Kougyou-Kanda Bldg.
remarks:        Address: 2-3-4 Uchikanda
remarks:        Address: Chiyoda-ku
remarks:        City: Tokyo
remarks:        Country: JP
…
remarks:        RTechHandle: ZI92-ARIN
remarks:        RTechName: IPv6 Promotion Council of Japan
remarks:        RTechPhone: +81-3-3548-2601
remarks:        RTechEmail: admin@v6nic.net
```

43/8 was used by IPv6 Promotion Council of Japan for IPv6 transition experimental purposes

No specific entry for this /24 in any RIR database

Uses Cases

Invalid ROAs

```
route-views2.routeviews.org> sh rpki prefix-table | include ^43.255
…
43.255.30.0                                24 -  24   133199
43.255.30.0                                24 -  24   133861
43.255.31.0                                24 -  24   133199
43.255.31.0                                24 -  24   133861
43.255.24.0                                22 -  22   18119
43.255.36.0                                22 -  22   0
43.255.53.0                                24 -  24   59214
43.255.52.0                                24 -  24   59214
43.255.54.0                                24 -  24   59214
…
```

This /22 is in APNIC's AS0 TAL – meaning it is unassigned (classed as a bogon)

RouteViews validators include LACNIC's and APNIC's AS0 TALs – so RouteViews `sh bgp` output can indicate prefixes in the BGP table that are flagged by the RIRs as unassigned.

Hence the invalid ROA displayed for 43.255.37.0/24

AS64021 & AS137451 are both announcing part of an IP address block that is not theirs to use.

Other Bits

RouteViews email list

- https://lists.nsrc.org/listinfo/routeviews-users
- Also available on the Contact page at routeviews.org
- A place to ask questions and receive updates on RouteViews activities.
- Hosted by the wonderful folks at NSRC.

RouteViews in Asia

Looking for new Views

- The RouteViews team is continually on the lookout for new locations
- More views means better troubleshooting and diagnosis of the global BGP routing system

- Preferences:
  - At an IX
  - Rich variety of domestic, regional, and international peers
  - IPv4 & IPv6 (of course!)
  - Host willing to **sponsor** either a VM or the hosting of a 1RU server

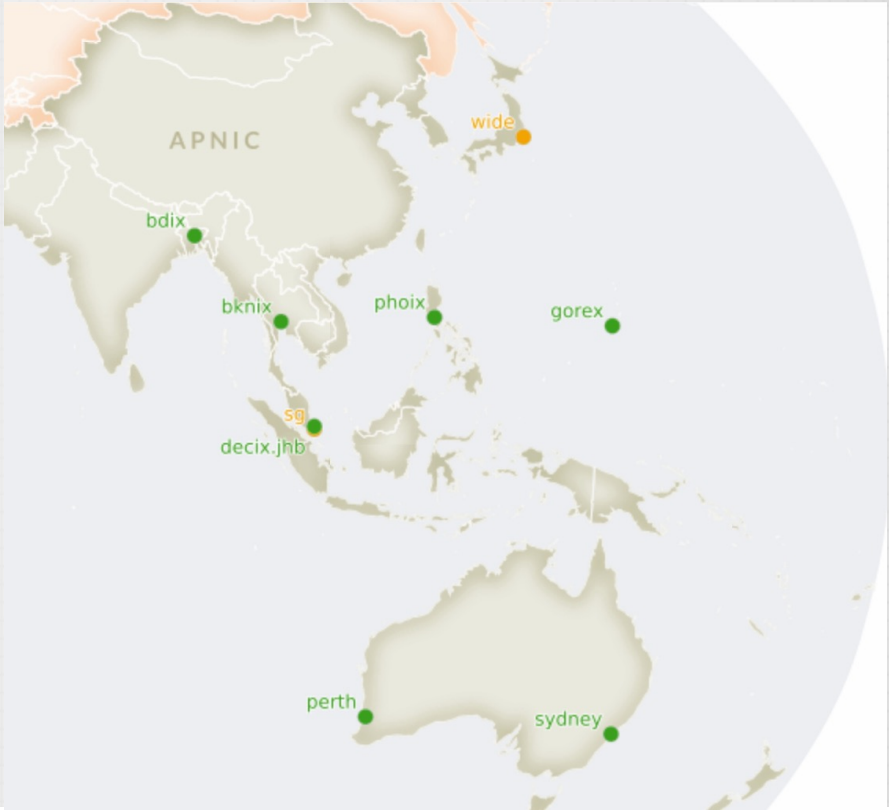

RouteViews Impact

Geoff Huston wrote in his report, "BGP in 2022 – the routing table":

"I should take a moment to mention the RouteViews Project. It was originally intended to offer a multi-perspective real-time view of the inter-domain routing system, allowing network operators to examine the current visibility of route objects from various points in the inter-domain topology.

What makes RouteViews so unique is that it archives these routing tables every two hours and has done so for more than two decades. It also archives every BGP update message. ***This vast collection of data is a valuable research data source in its own right****, and here we are just taking a tiny slice of this data set to look at longer-term routing growth trends.*

The folk at the RouteViews Project, with support from the University of Oregon and the US National Science Foundation, should be commended for their efforts here. This is a very unique data set if you are interested in the evolution of the Internet over the years."

RouteViews Impact

Aftab Siddiqui

"The MANRS Observatory relies heavily on BGPStream and GRIP for the detection of BGP related incidents such as BGP Leaks and BGP mis-origination. It is also very critical to verify that any incident highlighted by these services can be verified independently and to do that ***we require raw BGP data which is made available by 2 sources: RIPE RIS and RouteViews.*** *Diversity of data sources is once again very important to verify any such incidents. NSRC, which manages the RouteViews project, ensures that the routing data they provide is accurate and they have promptly addressed any issues or concerns raised by the MANRS team, whether it is related to changes in the MRT format causing problems in data parsing or helping with BMP data.* ***Actively maintaining RouteViews provides community service by NSRC.***

MANRS has gained a lot of good reputation in the community due to the support and expertise provided by its partners such as NSRC. NSRC included MANRS Action explanation and implementation guidelines in their training courses for network operators and R&E networks, in their technical video content, and has been promoting various MANRS programs to respective communities specifically in Asia Pacific and Africa where the MANRS participation is low as compared to other regions."

RouteViews Impact

Barrett Lyon: https://www.youtube.com/watch?v=vo5glK9czIE

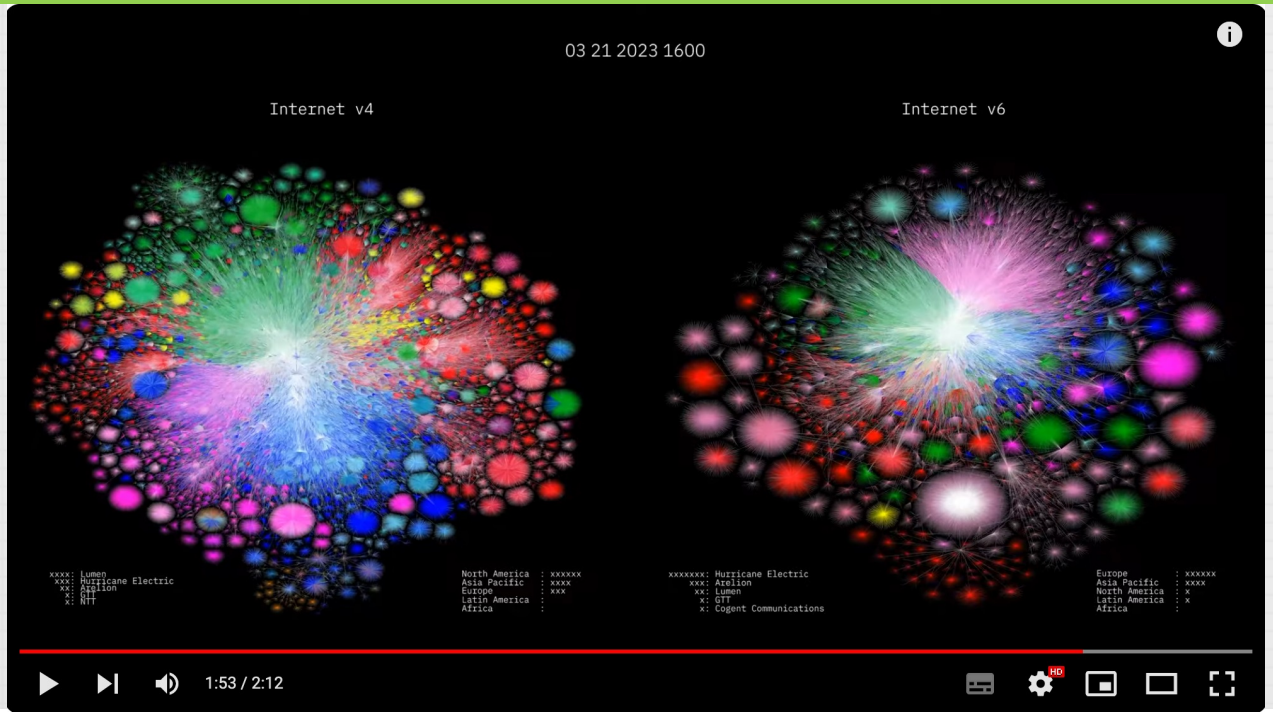

THANK YOU

Questions?