

RouteViews

BGP Global Monitoring Infrastructure and Services



Presented by Greg Shepherd
shep@routeviews.org

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

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 and Google with other contributions from public and private organizations.

UNIVERSITY OF OREGON

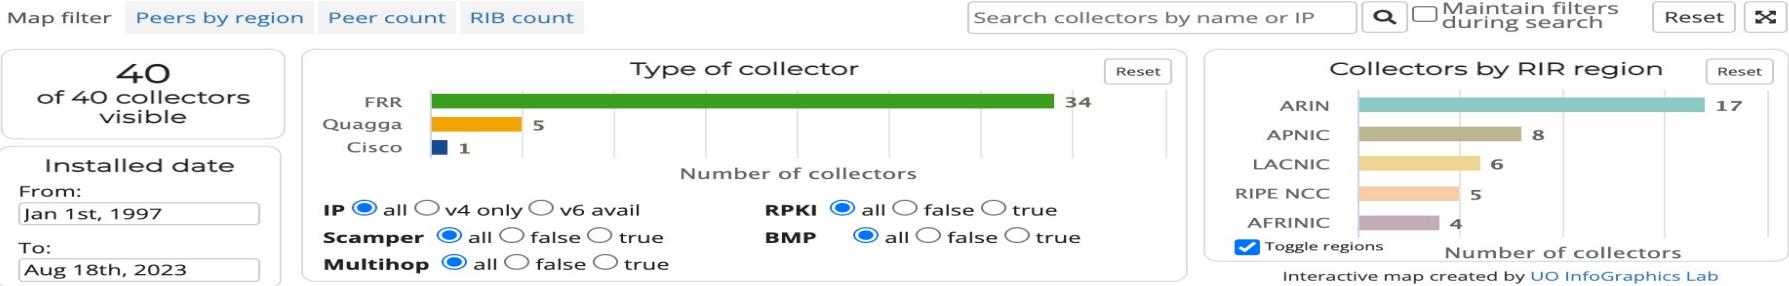
The University of Oregon is a public research institution in Eugene, Oregon, USA founded in 1876. UO is renowned for its research prowess and commitment to teaching. Both NSRC and RouteViews are based at the UO.

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/>

RouteViews



UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

COLLECTOR DEPLOYMENT

Physical

- Off the shelf hardware.
- Shipped to the IX.
- Last resort option only.

Virtual

- Preferred option.
- Much quicker deployment time.
- Easier to maintain.
- Easier to upgrade.

Collectors

Hardware

Commodity

- 8-16 Cores
- 32G-64G Ram
- 400GB-1TB SSD
- 10 GB eth

Software

OpenSource

- Linux/Centos and...
- Quagga – bgpd
- FRR – bgpd
- Gobgp
- OpenBMP
- GoBMP

Collector Deployment

Multi-hop

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

IX-Hosted/co-located

- Pros:
 - Better positioned to address multi-hop issues.
 - Geographic diversity.
 - Peering diversity.
 - Scalable.
- Cons:
 - 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

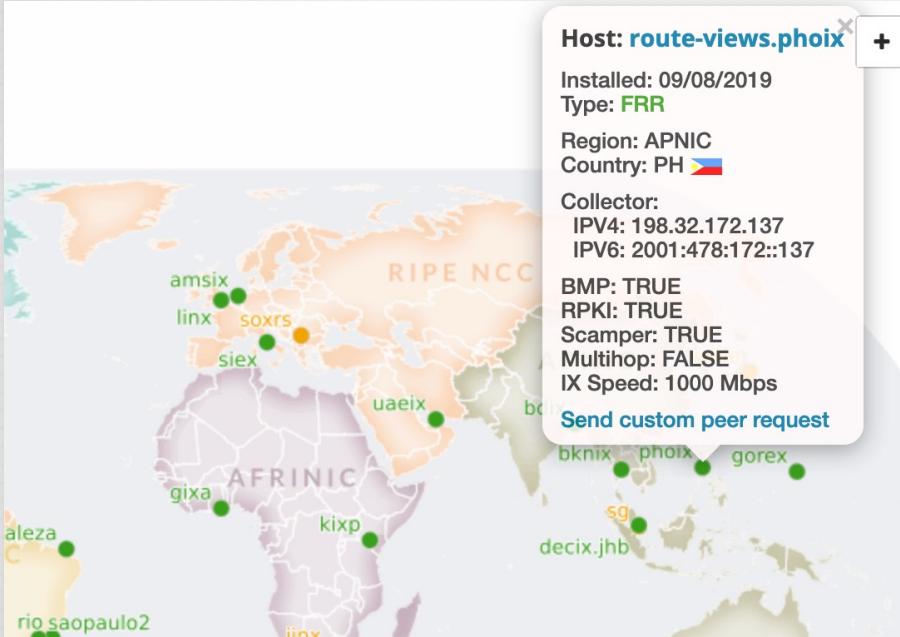
Data Access

- 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 (BMP) Feed

- 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

- Send full-table if you can
- Remove default route
- Remove NULL routes
- Remove RFC-1918 addresses
- We don't accept/want ADD-PATH TX/RX
- We don't send any routes to you
- When peering with multi-hop collectors, set ebgp-multipath

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

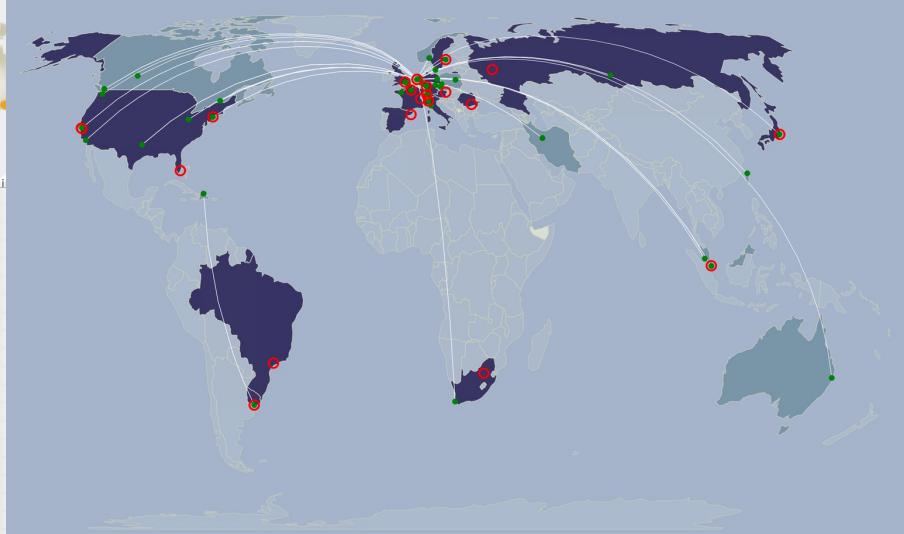
Map info and instructions



RouteViews Collector Map

Extends coverage
Provides Redundancy
Complimentary Services
Common files/tools

RIS Collector Map



RouteViews EU Collectors

```
route-views.amsix.routeviews.org> sh ip bg sum
```

IPv4 Unicast Summary (VRF default):

BGP router identifier 80.249.214.5, local AS number 6447 vrf-id 0

BGP table version 80501602

RIB entries 1823998, using 334 MiB of memory

Peers 62, using 44 MiB of memory

Two busiest collectors
cannot endure more
peers

Require additional
collectors to expand in
these locations

```
route-views.linx.routeviews.org> sh ip bg sum
```

IPv4 Unicast Summary (VRF default):

BGP router identifier 195.66.225.222, local AS number 6447 vrf-id 0

BGP table version 45222680

RIB entries 1824471, using 334 MiB of memory

Peers 76, using 54 MiB of memory

Would also like to
expand into strategic
EU locations:

Marseille anyone? :)

Use Cases

Operations

- 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??”

Use Cases

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

Use Cases

Operations

- 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 - Malaysia IX, Johor Bahru airport

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 at first.
- Wider availability to follow.

RPKI

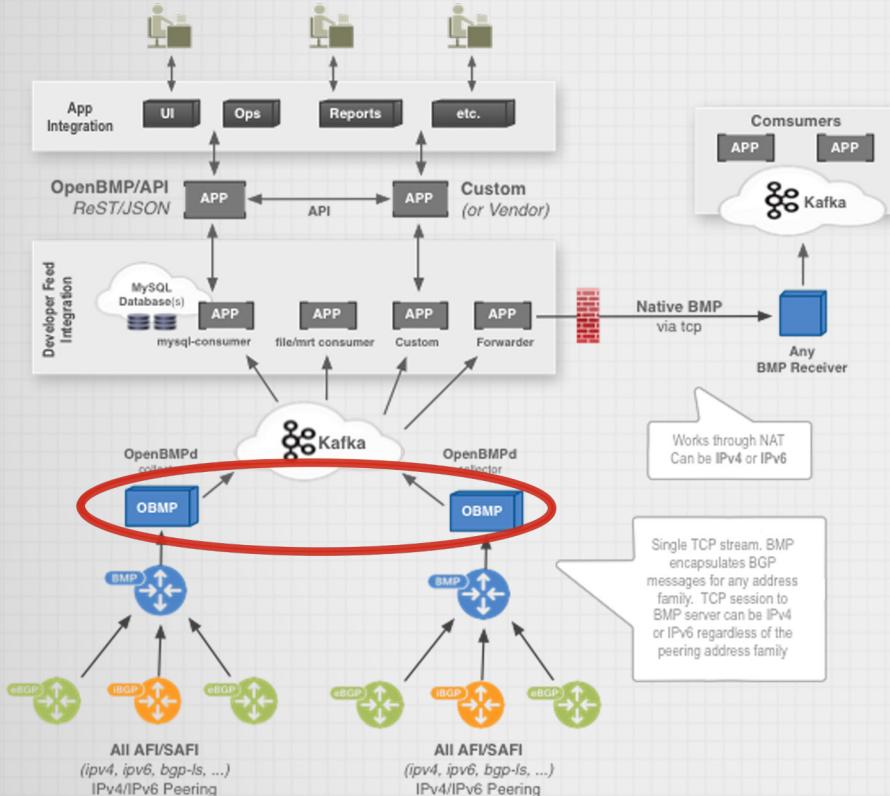
- RPKI data will be accessible directly from the collectors.
- We also are establishing an archive of RPKI ROA data.
- Working on back-filling that data set from the RIR/CAs.

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



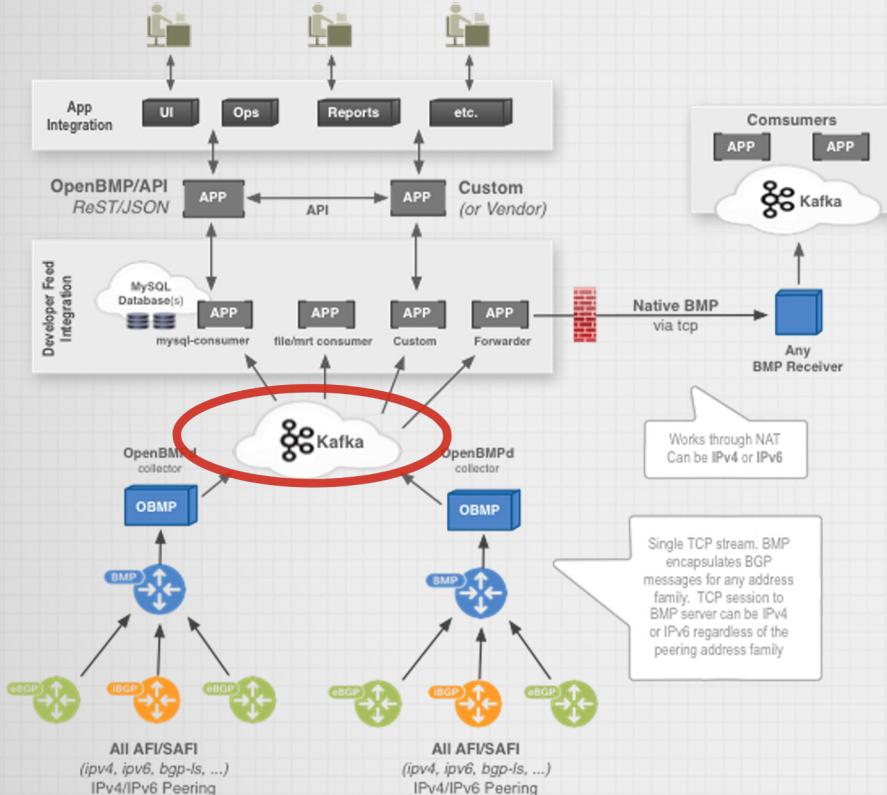
- BGP Monitoring Protocol (BMP) is an IETF standard rfc7854

- 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 😕

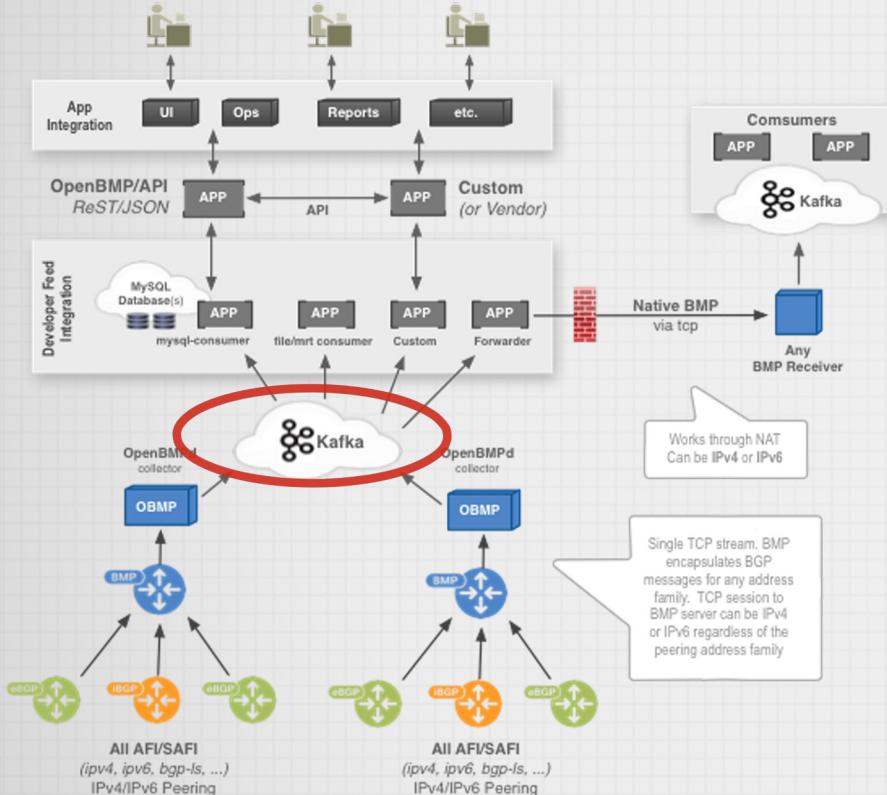
- GoBMP
 - Latest update: September 2022 😊

Apache Kafka



- Apache Kafka comprises the message bus for OpenBMP
 - Proven to scale
 - Mature client API
 - Clients in 16 different programming languages

Apache Kafka



- Send email to:
 - help@routeviews.org

Use Cases

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



BGPReader

Generate ASCII output

Eyeball raw data

Shell one-liners

BGPReader is the simplest interface to BGPStream: a command-line tool for extracting BGP measurement data in ASCII format. It can also be used as a drop-in replacement for the legacy `bgpdump` tool.

Use Cases

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

APIs

libBGPStream

Develop C/C++ code

✓ Build efficient tools

✓ Build complex infrastructure

libBGPStream is the central library of the BGPStream framework. It is written in C and presents a simple API for configuring and reading a stream of BGP measurement data. All BGPStream tools as well as the PyBGPStream API make use of libBGPStream.

PyBGPStream

Develop Python code

✓ Rapid prototyping

✓ Ad-hoc analysis

PyBGPStream is Python package that provides bindings to the libBGPStream library, allowing Python scripts to configure and read a stream of BGP measurement data.

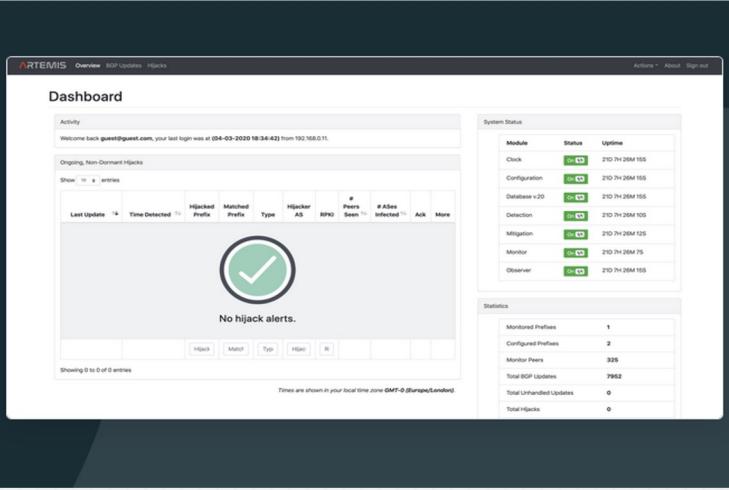
Use Cases

Other Open-Source tools: Artemis

Accurate and real-time BGP hijacking protection.

An open-source tool to monitor, detect, and mitigate BGP hijacks.

Live Demo



The screenshot shows the Artemis dashboard with a large green checkmark in the center. Below it, the text 'No hijack alerts.' is displayed. The dashboard includes a table for 'Ongoing, Non-Dormant Hijacks' with columns for Last Update, Time Detected, Hijacked Prefix, Matched Prefix, Type, Hijacker AS, ASNO, Peers Seen, # ASes Infected, and Ack. There are also sections for 'System Status' and 'Statistics'.

Last Update	Time Detected	Hijacked Prefix	Matched Prefix	Type	Hijacker AS	ASNO	Peers Seen	# ASes Infected	Ack
Showing 0 to 0 of 0 entries									

Module	Status	Uptime
Clock	OK	21D 7h 28M 155
Configuration	OK	21D 7h 28M 155
Database >20	OK	21D 7h 28M 155
Detection	OK	21D 7h 28M 105
Mitigation	OK	21D 7h 28M 125
Monitor	OK	21D 7h 28M 75
Observer	OK	21D 7h 28M 155

Statistics	
Monitored Prefixes	1
Configured Prefixes	2
Monitor Peers	325
Total BGP Updates	7962
Total Unhandled Updates	0
Total Hijacks	0

- **ARTEMIS**
 - An open-source tool to monitor, detect, and mitigate BGP hijacks
 - Real-time detection and notifications of BGP prefix hijacking attacks/events
 - <https://bgpartemis.org>

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

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

77 peers, multi-hop

```
IPv4 Unicast Summary (VRF default):  
BGP router identifier 128.223.51.102, local AS number 6447 vrf-id 0  
BGP table version 14375055  
RIB entries 1786807, using 327 MiB of memory  
Peers 77, using 54 MiB of memory
```

Not all peers are up..

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	PfxSnt	Desc
4.68.4.46	4	3356	0	15732	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	9318817	31942	0	0	0	03w1d04h	909294	0	ATT
37.139.139.17	4	57866	9727660	63869	0	0	0	03w1d04h	911576	0	Fusix
43.226.4.1	4	63927	0	0	0	0	0	never	Connect	0	Rise
45.61.0.85	4	22652	8018158	63869	0	0	0	03w1d04h	913095	0	FIBRENOIRE
62.115.128.137	4	1299	31136526	31866	0	0	0	01w3d16h	892007	0	Telia
64.71.137.241	4	6939	7260005	31935	0	0	0	03w1d04h	936399	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	13486557	63875	0	0	0	01w4d01h	924556	0	OBITRU
87.121.64.4	4	57463	13008288	32486	0	0	0	01w3d12h	413863	0	NETIXLTD
89.149.178.10	4	3257	8721640	31936	0	0	0	03w1d04h	909050	0	Tiscali
91.209.102.1	4	39756	0	0	0	0	0	never	Connect	0	HOSTWAY-RO
91.218.184.60	4	49788	17361887	31939	0	0	0	01w3d15h	915046	0	NEXTHOPNO
...											

Total number of neighbors 77

Lots of full tables

Use Cases

OPERATIONS

```
route-views2.routeviews.org> sh ip bgp 45.235.208.0/22
BGP routing table entry for 45.235.208.0/22, version 1474520
Paths: (30 available, best #25, table default)
  Not advertised to any peer
    11686 52320 22381 22381 22381 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 11432 268214
      96.4.0.55 from 96.4.0.55 (96.4.0.55)
        Origin IGP, valid, external, rPKI validation-state: not found
        Community: 11686:294
        Last update: Tue May 9 04:34:01 2023
    22652 4230 11432 268214
      45.61.0.85 from 45.61.0.85 (184.95.245.30)
        Origin IGP, valid, external, rPKI validation-state: not found
        Community: 4230:11 4230:30 4230:511 4230:5101
        Last update: Sat May 6 08:05:45 2023
    8492 31133 3356 268214 268214
      85.114.0.217 from 85.114.0.217 (85.114.0.104)
        Origin IGP, valid, external, rPKI validation-state: not found
        Community: 8492:1104 8492:1601
        Last update: Thu May 4 05:57:42 2023
...

```

What is ASN:11432 trying to achieve by prepending 23 times??

Use Cases

OPERATIONS

```
route-views2.routeviews.org> sh ip bgp rpk1 invalid
```

```
...
I* 212.193.8.0/24      168.209.255.56          0 3741 5511 6453 8551 203905 i
I*                      194.153.0.253          0 5413 8551 203905 i
I*                      45.61.0.85          0 22652 6453 8551 203905 i
I*                      91.218.184.60          0 49788 12552 8551 203905 i
I*                      198.129.33.85          710 0 293 6453 8551 203905 i
I*                      212.66.96.126          0 20912 49367 6762 61135 61135 60446 204843 204843 204843 i
I*                      202.232.0.3          0 2497 6453 8551 203905 i
I*>                     85.114.0.217          0 8492 8551 203905 i
I*                      203.189.128.233          0 23673 23764 8551 203905 i
I*                      94.156.252.18          0 0 34224 6453 8551 203905 i
```

ASN: 203905 ??

ASN: 204843 ??

Two ASN: of origin ??

Use Cases

OPERATIONS

WHOIS Lookup (212.193.8.0)

```
inetnum:          212.193.8.0 - 212.193.8.255
netname:          AjyalFiCompanyLLC
country:          PS
admin-c:          KB5060-RIPE
tech-c:           KB5060-RIPE
status:           ASSIGNED PA
abuse-c:          AR68281-RIPE
mnt-by:           lir-ae-rcstechnologies-1-MNT
mnt-by:           interlir-mnt
created:          2023-01-12T18:31:35Z
last-modified:    2023-01-12T18:31:35Z
source:           RIPE
```

Palestine

Use Cases

OPERATIONS

WHOIS Lookup (AS204843)

...

organisation: ORG-SVMY1-RIPE

org-name: STERLY VERI MERKEZI YAZILIM VE SIBER GUVENLIK HIZMETLERİ A.S.

country: TR  From Turkey ??

org-type: OTHER

address: KONAK MAH. BARIS(120) SK. OFIS ARTI BLOK NO:3 IC KAPI NO:10 NILUFER/BURSA

abuse-c: ACRO48320-RIPE

mnt-ref: ulasatakan

mnt-ref: bggroupittelecom-mnt

mnt-by: lir-tr-teknosos-1-MNT

mnt-by: Teknosos-TR

created: 2022-05-27T09:14:49Z

last-modified: 2022-12-01T17:27:16Z

source: RIPE # Filtered

Use Cases

OPERATIONS

WHOIS Lookup (AS203905)

```
...
as-block:      AS196608 - AS207419
...
admin-c:       KB5060-RIPE
tech-c:        KB5060-RIPE
status:        ASSIGNED
mnt-by:        RIPE-NCC-END-MNT
mnt-by:        DigiComm-MNT
created:       2015-10-06T13:35:54Z
last-modified: 2023-02-04T17:55:17Z
source:        RIPE # Filtered

organisation: ORG-DCL18-RIPE
org-name:      Digital Communication Company for Telecommunications and Information Technology LTD
country:       PS
org-type:      LIR
address:       Omar Al-Mokhtar st., Khaduir Building Floor #2
address:       9990300
address:       Gaza Al-Remal
address:       PALESTINE, STATE OF
```

From Palestine

Use Cases

OPERATIONS

```
route-views2.routeviews.org> sh ip bgp rpkI invalid
...
I* 212.193.8.0/24      168.209.255.56          0 3741 5511 6453 8551 203905 i
I*                      194.153.0.253          0 5413 8551 203905 i
I*                      45.61.0.85           0 22652 6453 8551 203905 i
I*                      91.218.184.60          0 49788 12552 8551 203905 i
I*                      198.129.33.85          710 0 293 6453 8551 203905 i
I*                      212.66.96.126          0 20912 49367 6762 61135 61135 60446 204843 204843 204843 i
I*                      202.232.0.3           0 2497 6453 8551 203905 i
I*>                     85.114.0.217          0 8492 8551 203905 i
I*                      203.189.128.233          0 23673 23764 8551 203905 i
I*                      94.156.252.18           0 0 34224 6453 8551 203905 i
```

ASN: 203905

ASn: 204843 Not valid origin ?

Mis-configuration? prefix Hijack??

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 question and receive updates on RouteViews activities.
- Hosted by the wonderful folks at NSRC.

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 BGPSStream 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."

= **THANK YOU**

Questions?