ROUTEVIEWS EVOLVES: Modernizing the BGP Collector for Today's Researcher

ROUTEVIEWS

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

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 22TBs (compressed) today.

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 22TBs (compressed) today.

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

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 22TBs (compressed) today.

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

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.

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 22TBs (compressed) today.

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

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.

ROUTEVIEWS

SPECIAL THANKS

JOHN KEMP
DAVE MEYER
RANDY BUSH
KIMBERLY (KC) CLAFFY
LUCY LYNCH
HANS KUHN
JOEL JAEGGLI
JOHN HEASLEY

FOOTPRINT


Finland
Iceland
Sweden
Norway
Russia
United Kingdom
Canada
Poland
Germany
Ukraine
Kazakhstan
Mongolia
France
Italy
Spain
Turkey
China
South Korea
Japan
United States
North Atlantic Ocean
Afghanistan
Iraq
Iran
Pakistan
Algeria
Libya
Egypt
Saudi Arabia
India
Thailand
Mexico
Mali
Niger
Sudan
Chad
Nigeria
Ethiopia
Venezuela
Colombia
DRC
Tanzania
Indonesia
Papua New Guinea
Brazil
Peru
Angola
Bolivia
Namibia
Madagascar
Indian Ocean
South Pacific Ocean
South Atlantic Ocean
Chile
Australia
Argentina
New Zealand

FOOTPRINT

COLLECTOR LOCATIONS

- ✓ Atlanta (digital realty)
- ✓ Chicago (equinx)
- ✓ Chile
- ✓ DC (eqix)
- ✓ Eugene (Multi-hop)
- ✓ Johannesburg (JINX, NAPAfrica)
- ✓ London (LINX)
- ✓ Miami (flix)
- ✓ Nairobi (kixp)
- ✓ Palo Alto (PAIX)
- ✓ Perth (WAIX)
- ✓ Portland (NWAX)
- ✓ Sao Paulo (IX.br x4)
- ✓ San Francisco (sfmix)
- ✓ Singapore (Equinix SG)
- ✓ Serbia (sox)
- ✓ Sydney (equinix)
- ✓ Tokyo (DIX-IE)
- ✓ Cape Town

PEERING STATS

PEERING STATS

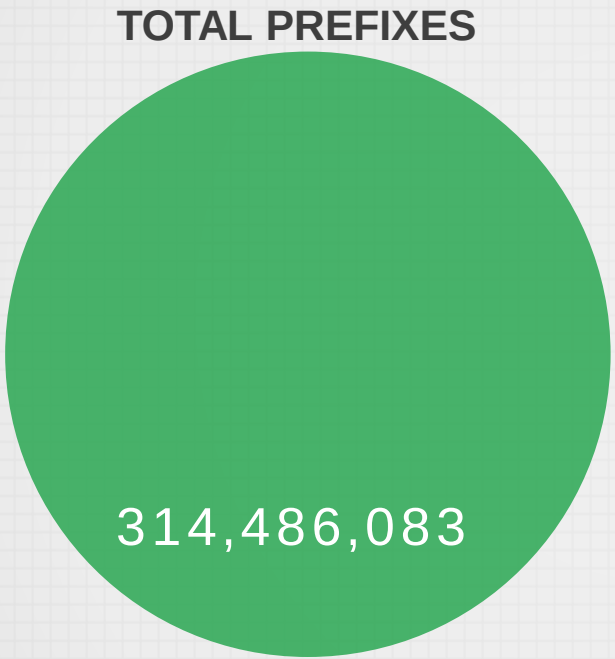

PEERING STATS

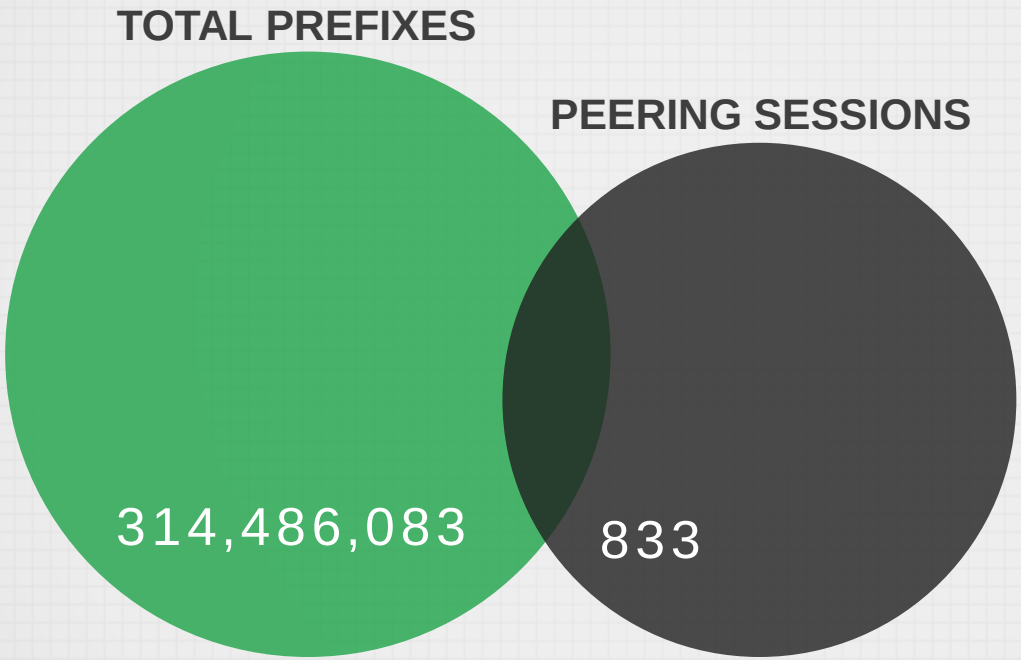

PEERING STATS

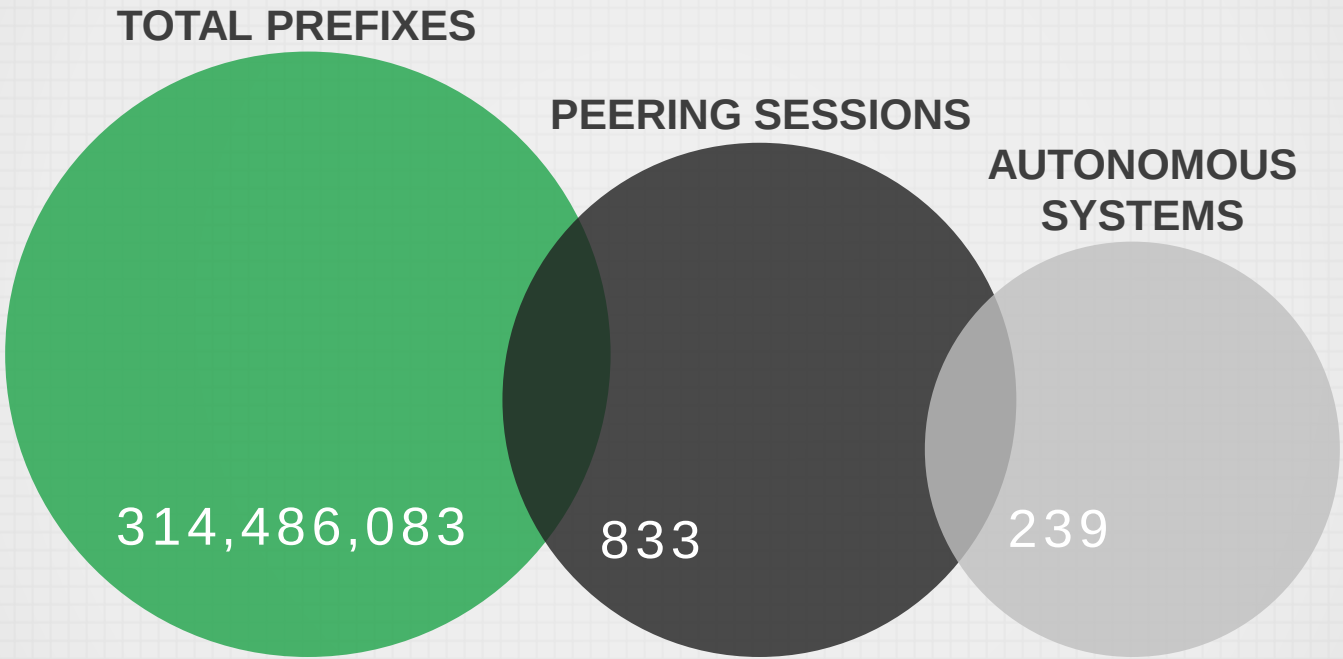

More peering information: routeviews.org/peers/peering-status.html

COLLECTORS

Commodity

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

Vendor

- ASR 1004

OpenSource

- Linux/Centos and…
- Quagga – bgpd
- FRR – bgpd

Vendor

- IOS XE

COLLECTORS OPERATIONS

CHANGE

Pros

- If you can reach the collector, you can peer

Cons

- Peerings are subject to the routing anomalies that RouteViews seeks to observe and collect

Pros

- Better positioned to address multi-hop issues
- Geographic diversity
- Peering diversity

COLLECTOR DATA

Multi-Threaded Routing Toolkit

- 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 dumped every 2 hours
  - UPDATEs are dumped every 15 minutes

DATA ACCESS

- MRT files are bzipped and rsynced back to http://archive.routeviews.org/ regularly
- They can be accessed via, http, ftp and rsync.

MRT TOOLS

RIPE libBGPdump, UCLA BGP Parser, NTT BGPdump2, etc:

- https://bitbucket.org/ripencc/bgpdump/wiki/Home
- https://github.com/cawka/bgpparser
- https://github.com/yasuhiro-ohara-ntt/bgpdump2
- https://github.com/t2mune/mrtparse (Python)
- https://github.com/rfc1036/zebra-dump-parser (Perl)

COLLECTOR ACCESSIBILITY

telnet://route-views*.routeviews.org

- No username necessary.
- Users are able to run show commands, e.g. show ip bgp x.x.x.x/x.

GOTCHAS

- Why not SSH?!
  - RouteViews data is publicly available. We've got nothing to hide.
  - We use ssh for host management.
- 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.

USE CASES

- BGP is the backbone of the Global Routing Infrastructure.
- To ensure it's 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

USE CASES

- BGP anomalies and dynamics are critical as well.
- RouteViews Provides:
  - Network Topology Monitoring
  - Route Leaks/Hi-Jacks (ex. Artemis, Cyclops)
  - Network Optimization
  - Growth, Aggregation, etc. In AS/V4/V6
  - Address Provenance
- ~500 research publications have used RouteViews data
- More info: http://www.routeviews.org/routeviews/index.php/papers/

BGP DATA DISTRIBUTION EVOLUTION

1st Generation Characteristics (current)

- File-Based storage, MRT data format

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (current)

- File-Based storage, MRT data format
- Asynchronous

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (current)

- File-Based storage, MRT data format
- Asynchronous
- Manual retrieval, sequencing, and consolidation

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (current)

- File-Based storage, MRT data format
- Asynchronous
- Manual retrieval, sequencing, and consolidation
- No post-processing

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (current)

- File-Based storage, MRT data format
- Asynchronous
- Manual retrieval, sequencing, and consolidation
- No post-processing
- Centralized model

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (future)

- “Message-based” data distribution, per-message timestamps, with meta-data

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (future)

- “Message-based” data distribution, per-message timestamps, with meta-data
- Automated consolidating and sequencing

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (future)

- “Message-based” data distribution, per-message timestamps, with meta-data
- Automated consolidating and sequencing
- Database storage and access

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (future)

- “Message-based” data distribution, per-message timestamps, with meta-data
- Automated consolidating and sequencing
- Database storage and access
- RESTful interfaces

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (future)

- “Message-based” data distribution, per-message timestamps, with meta-data
- Automated consolidating and sequencing
- Database storage and access
- RESTful interfaces
- Real-time streaming telemetry

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (future)

- “Message-based” data distribution, per-message timestamps, with meta-data
- Automated consolidating and sequencing
- Database storage and access
- RESTful interfaces
- Real-time streaming telemetry
- Middle-layer abstraction, multi-client access (facilitates analysis and services)

BGP DATA DISTRIBUTION EVOLUTION

n Characteristics (future)

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

NEXT STEPS COMMUNICATION

- Better communications for those who are interested.
  - Maintenance.
  - Outages.
  - Collector announcements.

NEXT STEPS GOVERNANCE

mittee

- Ensure RouteViews continues to meet the needs of the community.
- Comprised of research and industry members.

NEXT STEPS BMP & OpenBMP

BGP Monitoring Protocol

- https://tools.ietf.org/html/rfc7854
- Available now – Cisco, Juniper, (FRR coming soon)
- In addition to MRT attributes BMPs adds
  - Start, Stop, Peer Up, Peer Down
  - Collector Identification
  - Statistics

NEXT STEPS BMP & OpenBMP

- BMP is the IETF standard for BGP monitoring
- OpenBMPd is OpenSource (part of the Linux Foundation)
  - Consolidates peers/collectors
  - Splits collector, peer and update messages into separate streams
- Apache Kafka comprises the message bus for openbmp
  - Addresses producer/consumer problems
  - Proven to Scale
  - Mature client API
    - Clients in 16 different programming languages
  - Can be easily extended to meet future needs.

OpenBMP ARHITECTURE

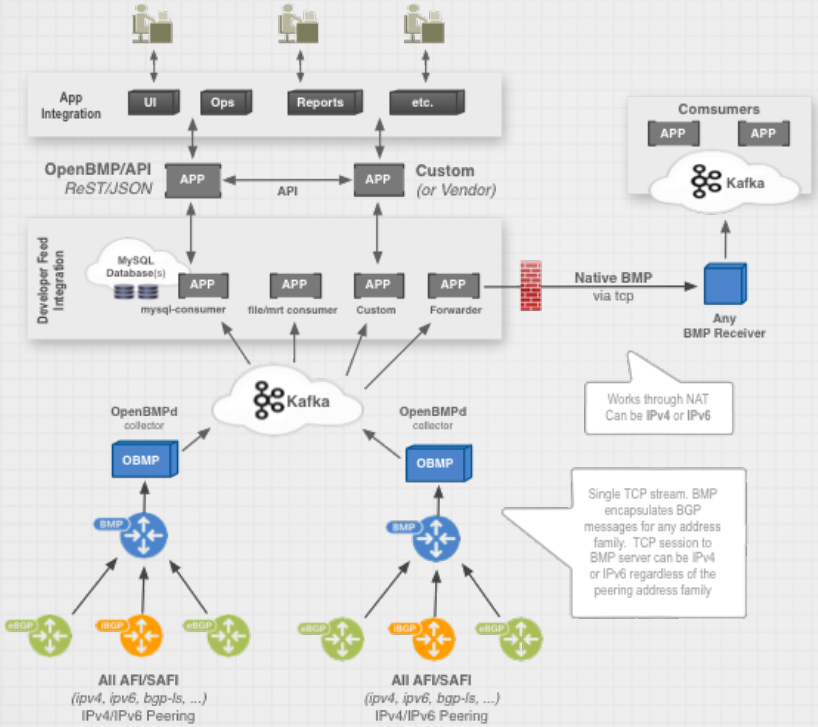

https://github.com/OpenBMP/openbmp/blob/master/docs/images/openbmp-flow.png

BMP TOOLS

- https://bgpstream.caida.org/

Languages:

- https://cwiki.apache.org/confluence/display/KAFKA/Clients

RESEACH OPPORTUNITIES

By leveraging the 2nd generation characteristics of RouteViews BGP data distribution, new and novel approaches to BGP anomaly and dynamics analysis are possible.

RESEACH OPPORTUNITIES

- Use RouteViews API data for ML supervised learning. Train models to better detect:
  - Route leaking/hijacking
  - Infrastructure/peering outages
  - Internet censorship
  - Routing policy complexity
- Validate ML models against live BMP streams

THANK YOU

Questions?