ROUTEVIEWS EVOLVES

Modernizing the BGP Collector for Today's Researcher & Operator

ROUTEVIEWS

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

ROUTEVIEWS

A collaborative routing 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 30TBs (compressed) today.

ROUTEVIEWS

A collaborative routing 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 30TBs (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, and ESnet engineers.

ROUTEVIEWS

A collaborative routing 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 30TBs (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, and ESnet engineers.

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 routing 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 30TBs (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, and ESnet engineers.

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.

FOOTPRINT

COLLECTOR LOCATIONS

- Amsterdam/Sweden (AMSIX)
- Atlanta (TELXATL on digital realty)
- Chicago (at Equinix)
- Chile
- DC/Ashburn (EQIX)
- Eugene (various multi-hop)
- Fortaleza
- Johannesburg(JINX, NAPAfrica)
- London (LINX – LON1 & 2)
- Miami (FLIX)
- Nairobi (KIXP)
- Perth (WAIX)
- Portland (NWAX)
- Rio di Janeiro (RIO)
- San Francisco (SFMIX)
- Sao Paulo (2 collectors on IX.br)
- Serbia (SOX)
- Singapore (SG on Equinix)
- Sydney (on Equinix)
- Tokyo (WIDE on DIX-IE)
- Palo Alto (PAIX on Equinix)

New collectors: MWIX, PHOIX, GIXA, BKNIX, GOREX, and DATA-IX (St. Petersburg)

PEERING STATS

PEERING STATS

PEERING STATS

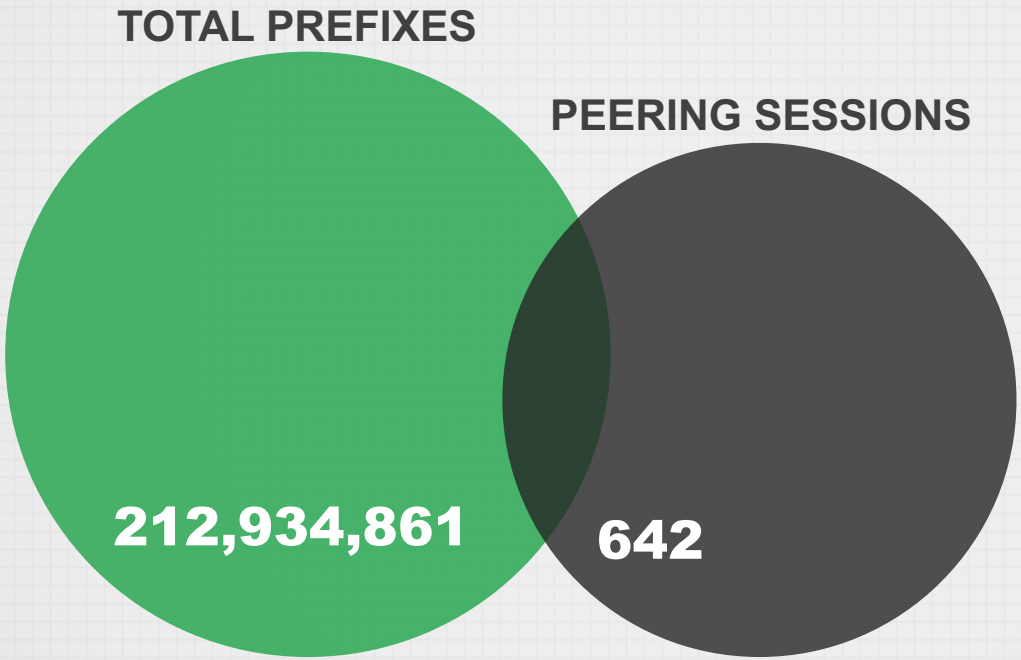

PEERING STATS

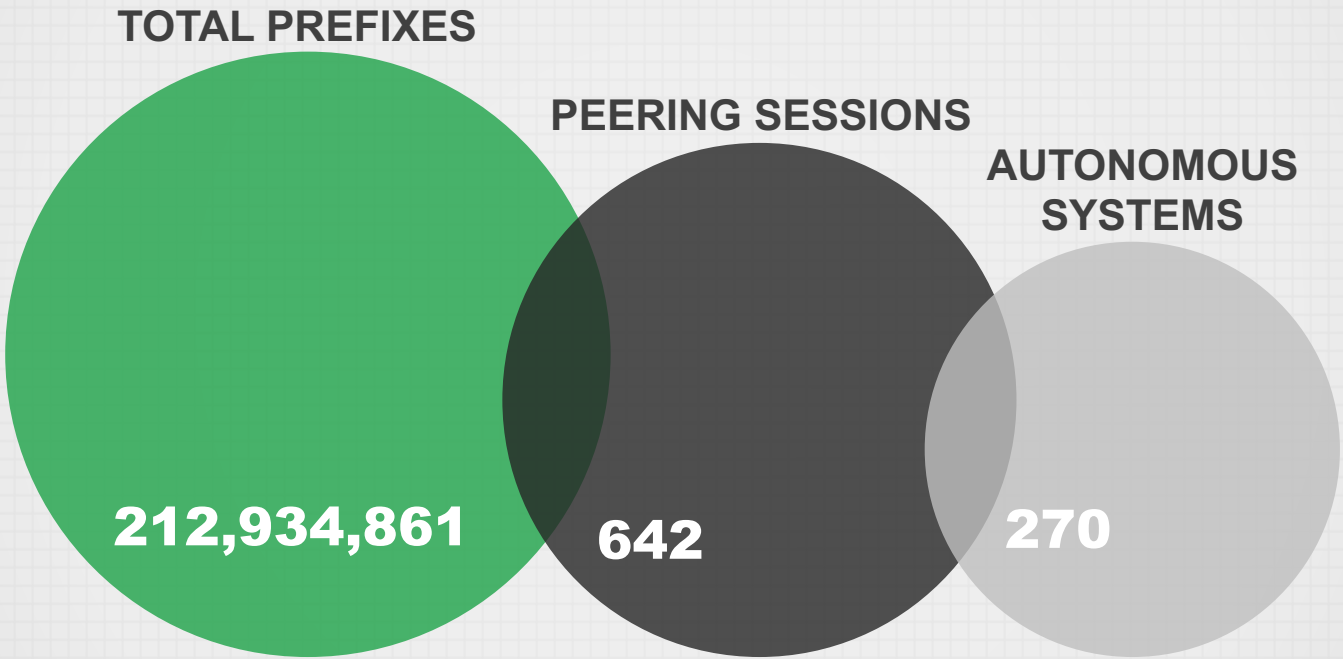

More peering information: routeviews.org/peers/peering-status.html

COLLECTORS

HARDWARE

Commodity

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

Vendor

- Cisco ASR 1004

SOFTWARE

OpenSource

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

Vendor

- IOS XE

COLLECTORS OPERATIONS

MULTI-HOP

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

INTERNET EXCHANGE

Pros

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

COLLECTOR DATA

MRT

Multi-Threaded Routing Toolkit

- https://tools.ietf.org/html/rfc6396
- MRT provides a standard for dumping routing information to a binary file.
- RouteViews MRT 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, Isolario BGPscanner:

- 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)
- https://github.com/CAIDA/libparsebgp
- https://bgpstream.caida.org/
- https://www.isolario.it/web_content/php/site_content/tools.php

COLLECTOR ACCESSIBILITY

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

- No username necessary.
- Users can 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

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

USE CASES

RESEARCH

- BGP anomalies and dynamics are critical as well.
- RouteViews Provides:
  - Network Topology Monitoring
  - Route Leaks/Hijacks (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

- File-based storage, MRT data format

BGP DATA DISTRIBUTION EVOLUTION

1st Generation Characteristics

- File-based storage, MRT data format
- Asynchronous

BGP DATA DISTRIBUTION EVOLUTION

1st Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

1st Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

1st Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

2nd Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

2nd Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

2nd Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

2nd Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

2nd Generation Characteristics

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

BGP DATA DISTRIBUTION EVOLUTION

2nd Generation Characteristics

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

ROUTEVIEWS: FAQ

- Many of the following examples can also be found on our FAQ page:
  http://www.routeviews.org/routeviews/index.php/faq/

ROUTEVIEWS: HOW TOS

- These commands can be run from all RouteViews collectors.
- For a list of collector locations, visit our interactive map.
  http://www.routeviews.org/routeviews/index.php/map/

ROUTEVIEWS: HOW TOS

- Common use cases: What routes am I advertising?

```
% telnet route-views3.routeviews.org
route-views3>show ip bgp regexp _<your ASN>$
```

ROUTEVIEWS: HOW TOS

- Common use cases: What's the best path to a prefix?

```
% telnet route-views3.routeviews.org
route-views3>show ip bgp <your prefix> bestpath
```

ROUTEVIEWS: HOW TOS

- Common use cases: How do I know which collector I'm peered with?
- We keep a list of each collectors established peers at http://www.routeviews.org/peering-status.html
- You can also curl the output and grep for your AS

```
% curl http://www.routeviews.org/peering-status.html | grep
'<your AS>'
```

RPKI OVERVIEW

Resource Public Key Infrastructure (**RPKI**) is a public key infrastructure framework designed to secure the Internet's routing infrastructure, specifically the Border Gateway Protocol. **RPKI** provides a way to connect Internet number resource information (such as IP Addresses) to a trust anchor.

RPKI OVERVIEW

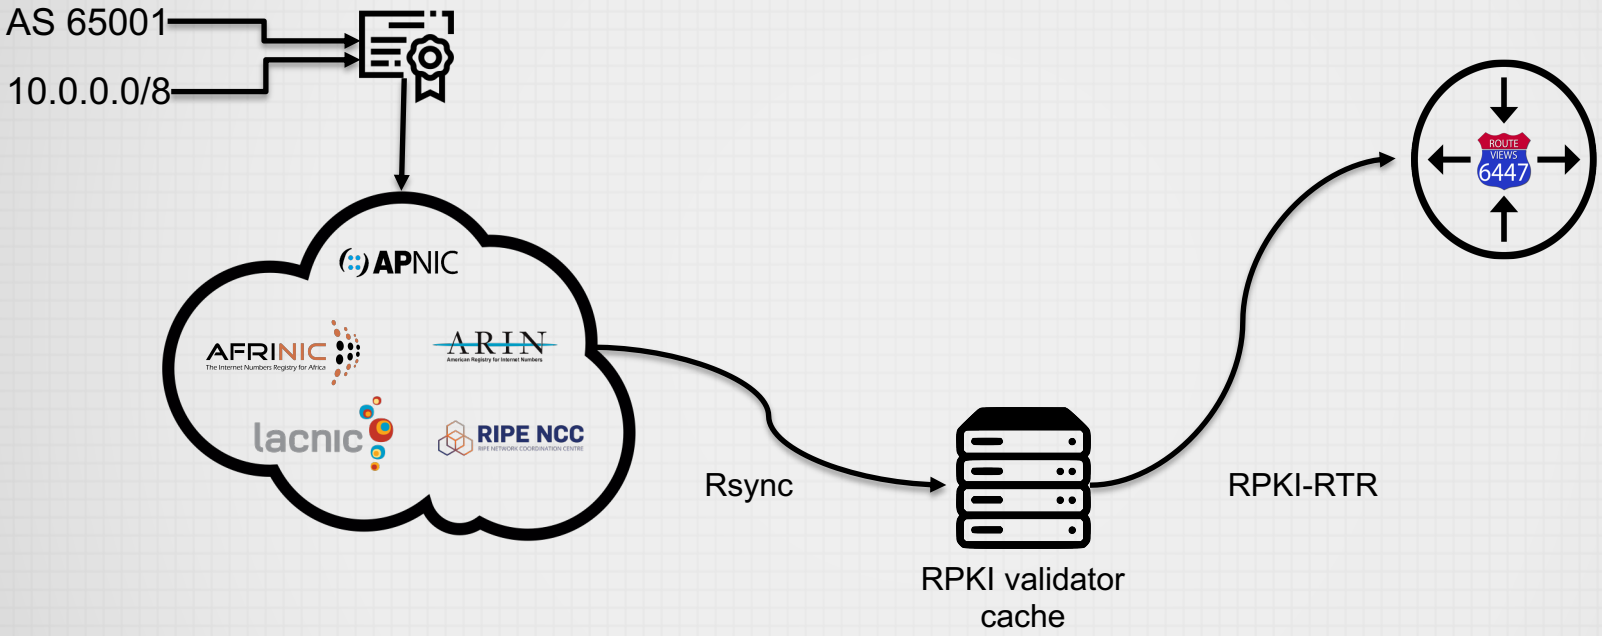

RPKI ON ROUTEVIEWS

- All the modern RouteViews collectors are connected to a RPKI validator
- We do not filter or drop any routes based on their RPKI state

ROUTEVIEWS: HOW TOS

- Common use cases: What is the RPKI state of my prefix?

```
% telnet route-views3.routeviews.org
route-views3>show rpki prefix <your prefix>
```

- You can get the whole RPKI prefix table with:

```
% telnet route-views3.routeviews.org
route-views3>show rpki prefix-table
```

- The RouteViews map has a list of RPKI enabled collectors http://www.routeviews.org/routeviews/index.php/map/

ROUTEVIEWS: HOW TOS

- Common use cases: What is the RPKI state of my prefix?

```
% curl https://api.routeviews.org/rpki?prefix=1.1.1.0/24
 {"1.1.1.0/24":{"asn":[{"13335":"valid"}],"timestamp":"2020-07-23 04:00:02"}}
```

- Or

```
% curl https://api.routeviews.org/rpki?asn=13335
 {"13335":{"prefix":[{"1.0.0.0/24":"valid"},{"1.1.1.0/24":"valid"},
{"23.227.38.0/23":"valid"},{"103.22.200.0/23":"valid"}…
```

- Entries are regenerated every 2 hours
- A full dump of valids, invalids, and unknowns is available at

```
% curl https://api.routeviews.org/rpki
```

BMP & OpenBMP

BMP

BGP Monitoring Protocol

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

BMP & OpenBMP

- 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 ARCHITECTURE

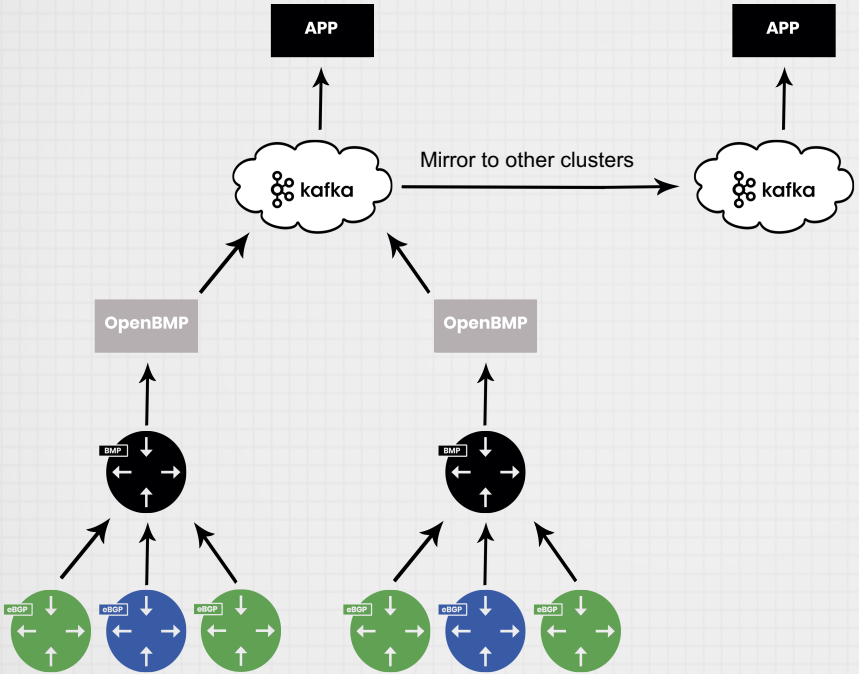

OpenBMP ARCHITECTURE

KAFKA TOPICS

A Topic is a category/feed name to which messages are stored and published.
OpenBMP uses 3 types of topics:

- Collector: Information about the openbmp collector(s).
- Router: Information about router state (up/down/name/version/etc).
- BMP: Raw bmp messages grouped by…

{{collector_group}}.{{router_group}}.{{peer_asn}}.bmp_raw

routeviews.linx.4775.bmp_raw

Kafka consumers support a subscribe pattern, which is a regex.

- Pattern: /^.*\.16509\.bmp_raw$/ - subscribe to all updates from Amazon
- Pattern: /^.*multihop.*$/ - subscribe to updates from all multihop collectors

BMP TOOLS

- https://bgpstream.caida.org/
  - https://bgpstream.caida.org/docs/install/bgpstream
- Languages:
  - https://cwiki.apache.org/confluence/display/KAFKA/Clients

ROUTEVIEWS: HOW TOS

- Common use cases: How can I see live BGP updates from a specific RouteViews peer?

```
% bgpreader -d kafka -o brokers=stream.routeviews.org:9092 -o
topic="^routeviews.*\.bmp_raw" -o data-type=bmp
U|A|1595889133.191761||is-ah-
bmp1|fortaleza|189.90.173.248|52320|45.184.144.128|194.110.144.0/22|45.184.1
44.128|6447 52320 31122 42227 39485|39485|52320:21311||
U|A|1595889133.191763||is-ah-
bmp1|fortaleza|189.90.173.248|52320|45.184.144.128|197.149.123.0/24|45.184.1
44.128|6447 52320 16637 29465 37480|37480|52320:21311||
```

ROUTEVIEWS: HOW TOS

- Bgpreader supports several filtering mechanisms out of the box.

```
-R, --router     <router>     process records from only the given router
-j, --peer-asn   <peer ASN>   return elems received by a given peer ASN
-a, --origin-asn <origin ASN> return elems originated by a given origin ASN
-k, --prefix     <prefix>     return elems associated with a given prefix
-y, --community  <community>  return elems with the specified community
-A, --aspath     <regex>      return elems that match the aspath regex
```

- Very useful for quick CLI filtering of live bgp updates.

ROUTEVIEWS: HOW TOS

- For a more programmatic approach, we can use
  - libBGPStream (C/C++ API)
  - pyBGPStream (python wrapper for libBGPStream)

ROUTEVIEWS: HOW TOS

Pybgpstream examples:
https://github.com/routeviews/tutorials/tree/master/pybgpstream

RESEACH OPPORTUNITIES

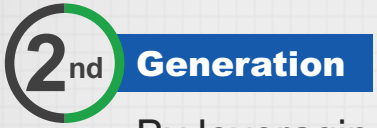

By leveraging the 2nd generation characteristics of RouteViews BGP data distribution, new and novel approaches to BGP anomaly and dynamics analysis are possible.

RESEACH OPPORTUNITIES

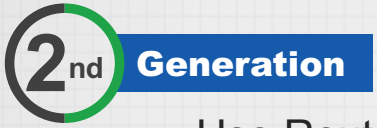

- Use RouteViews API data for ML supervised learning. Train models to better detect:
  - Route leaking/hijacking
  - Infrastructure/peering outages
  - Internet censorship
  - Routing policy complexity
- Validate ML models against live BMP streams

THANK YOU

We'd like to take this opportunity to thank everyone involved with the Routeviews project, especially the companies and consortiums that host our collectors in their datacenter or on a VM cluster. **Your contribution is invaluable to this project!**

Questions?