RouteViews: Internet Routing and BGP Data Collection Architecture

The University of Oregon's RouteViews platform provides detailed public views of Internet routing data. It was originally conceived in 1995 as a tool for Internet operators to obtain realtime Border Gateway Protocol (BGP) information about the global routing system from the perspectives of several different backbones and locations around the Internet. The 26-year data set of BGP information archived by RouteViews since 1997 helps network operators and researchers identify and address issues related to routing stability, security, and performance of the global Internet. Network operators utilize RouteViews data to understand reachability, hijacks, peer visibility, mass withdrawals, and RPKI status. The Network Startup Resource Center (NSRC) provides technical and operational management of the set of 40 (and growing) collectors deployed around the world.

System Architecture

Starting in 1995, the University of Oregon's RouteViews project received BGP data snapshots about the global routing system from a single collector (router) hosted at the UO Computing Center, and began archiving the raw BGP data in 1997. Over time it grew into a globally distributed network of route collectors providing route table perspectives of several different backbones and locations around the Internet. Although other tools handle related tasks, such as the various Looking Glass systems, they typically either provide a limited, constrained view of the routing system (e.g., either a single provider, or the route server) or they do not provide real-time access to routing data.

Today RouteViews encompasses a global distributed collection architecture providing raw BGP data and archiving information about the global Internet routing table with real-time updates. This growth was organic, being guided by working relationships with industry and academia in response to requests and offers to host expansion collectors and/or new peers. The infrastructure to support this growth was deployed on an as-needed basis, both for scale and to support new features, such as live interactive access BMP (BGP Monitoring Protocol).

While the RouteViews project was originally fueled by network operators to determine how the global routing system viewed *their* prefixes and/or AS space, there have been many other informative uses of RouteViews data by the research community. RouteViews is cited in more than 1000 peer-reviewed research papers. The research community has expressed interest in a greatly expanded RouteViews system, targeting a global footprint of many more new collectors and peers. This increase in scale is significant enough to justify a design process to standardize topology, components, and configuration, with the goal of a scalable solution that can be incrementally deployed and managed.

A detailed diagram of the overall topology of the RouteViews ecosystem is needed and under development, but not yet ready to add to the report at this time.

Current RouteViews Discussion for Prototyping a New RV Architecture

- Review stakeholder (researchers and network operators) requirements and expectations of the RouteViews service and dataset, to see if there are other BGP collection mechanisms beyond being present at major Internet eXchange Points worldwide.
- Explore different achitectures for scaling BGP data collection. Review the existing use of FRR (Free Range Routing) as the BGP collector daemon, and consider more scalable options (for example, virtualized versions of major vendor operating systems like Juniper or Cisco or scalable open source BGP daemons such as BIRD). Examine and test load-balancing designs which segment peers into individual containers.
- With the ever-increasing number of routes in both the global IPv4 and IPv6 BGP tables, as well as adding new collector sites and new peers, RouteViews is reviewing the existing data collection infrastructure, to increase both the use of virtualization and clustering, as well as improving the redundancy of data capture.
- While RouteViews deployments of new collectors and new peers at collectors is significantly automated, recent advances in software tools means that further improvements can be made to make the daily operational maintenance of RouteViews infrastructure more straightforward and less intensive for RV personnel. More intelligent peering models are being explored to improve efficiency. Improvements to existing automation and configuration management tools is actively ongoing.
- 26 years of historical data is of huge value to researchers and network operators alike, and RouteViews is reviewing enhancing backup options and locations so none of this data can be lost in case of issues at the main hosting site (full copies in different data centers) in Oregon. Google currently hosts a full copy of the entire RouteViews dataset, and worked with NSRC to create a pipeline to ingest all new data to provide a full backup with recovery features enabled.
- Increase community engagement and involvement via more public presentations at Network Operator Group meetings, Peering Forums, Research and Education Network meetings, and workshops for researchers who use RV data.