Visualizing BGP RIB Changes into Forwarding Plane by Leveraging BMP and IPFIX

Project Proposal

Swisscom

1 Description

Swisscom uses the Network Telemetry Framework [7] to correlate information and network metrics from the control plane, from the forwarding plane and from the topology (devices) to bring visibility into Swisscom MPLS and VXLAN encapsulated networks. The end goal is to understand what cause (control plane or topology event) has which impact on the forwarding plane and make it visible.

Control plane monitoring (BGP Monitoring Protocol or BMP), in particular, has recently seen important extensions, with the introduction of access to Adj-RIB-Out RIBs [4] and to the BGP instance Local-RIB [3] - both in draft status. Having access to all three BGP RIBs (Adj-In, Adj-Out and Local), together with existing forwarding plane data from IPFIX has the potential to substantially improve visibility into both private networks, such as Swisscom's and public ones.

Swisscom has access to early field trial router code, BMP collectors supporting the new BMP drafts within a Swisscom lab environment which can be made accessible externally. However, due to the scale of the network, an important pre-requisite for correlating the newly available control plane data with forwarding plane data is for the collected BMP data to be usable in a distributed programming environment, such as Apache Spark. This requires defining a new data model and serializing the BMP data (ideally using the Avro framework [1]) to enable correlation to existing IPFIX collected forwarding plane metrics, currently stored in a the real-time analytics database Druid [2].

Once this pre-requisite is completed the control plane - forwarding plane correlation is possible and the implementation of impactful use cases such as root cause analysis or real-time, deep performance monitoring are the logical next steps.

While several references exist for real-time analysis of control plane information [5, 6], the focus is on public networks, they do not include the newly available BGP RIBs, and there's no correlation to the forwarding plane.

2 Requirements

A solid understanding of large scale networks, VPNs and routing protocols, in particular BGP (Border Gateway Protocol) on a metric/schema level.

Ideally, an understanding of big data technologies such as Apache Kafka, Apache Avro, Sparks, HDFS, Druid etc. and great interest to bring visibility into networks.

Table 1: Tentative schedule for 6 months (26 weeks)

Milestone	**Estimated Effort**
Onboarding: Setting up and getting to know the lab environment	2 weeks
Baseline: Impact visualization with forwarding plane data (IPFIX/Netflow)	4 weeks
Improvement:	
1. Understanding BMP formats from field routers in the lab	4 weeks
2. Understanding Avro (input requirements from Kafka)	2 weeks
3. Putting BMP and Avro together	8 weeks
4. Impact visualization with control plane (BMP) and forwarding plane data (IPFIX/Netflow)	4 weeks
Analysis: Benefits of using BMP? (thesis write-up)	2 weeks

References

[1] Apache. Avro. `https://avro.apache.org/`.

[2] Apache. Druid. `http://druid.io/`.

[3] T. Evens, S. Bayraktar, M. Bhardwaj, and P. Lucente. Support for Local RIB in BGP Monitoring Protocol (BMP). `https://tools.ietf.org/html/draft-ietf-grow-bmp-local-rib-02`, 2018.

[4] T. Evens, S. Bayraktar, P. Lucente, and S. Zhuang. Support for Adj-RIB-Out in BGP Monitoring Protocol (BMP). `https://tools.ietf.org/html/draft-ietf-grow-bmp-adj-rib-out-03`, 2018.

[5] J. Obstfeld, X. Chen, O. Frebourg, and P. Sudheendra. Towards Near Real-Time BGP Deep Analysis: A Big-Data Approach. *arXiv:1705.08666*, May 2017.

[6] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti. BGPStream: A Software Framework for Live and Historical BGP Data Analysis. In *Internet Measurement Conference (IMC)*, Nov 2016.

[7] H. Song, T. Zhou, Z. Li, Z. Li, P. Martinez-Julia, L. Ciavaglia, and A. Wang. Network telemetry framework. `https://tools.ietf.org/html/draft-song-opsawg-ntf-02`, 2018.