Peering into Peering

Modern Peering Network Telemetry

Phil Bedard – TME @Cisco

Peering Days EU

3/7/2018

Old Methods vs. Modern Telemetry

- Limited insight
- Shallower depth
- Inefficient and limited delivery speed

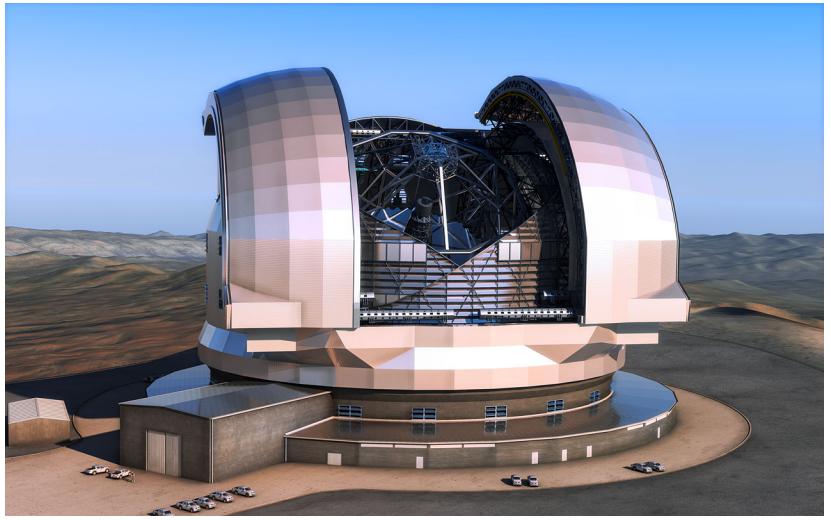

- Broader network view
- Increased resolution

What is telemetry?

- **Telemetry can mean different things to different people.**

- The dictionary says “Telemetry” is the process of using the onboard measurement capabilities of a a device to capture system data and communicating the data to a remote monitoring station.

- Telemetry in the context of networking: Collecting operational state data from a router, switch, or network of devices and sending it to a collection entity which may do meaningful things with it.

- “Sensors” are created to collect data, additional resources package the data and send it to collection entity for further processing and analysis.

- Telemetry inherently follows a “push” model once sensors and collection entities are configured on the device.

Types of modern network Telemetry

Periodic Streaming Telemetry

- Data is collected on node, “pushed” to collection entity at periodic intervals
- Cisco calls this model-driven telemetry (MDT)
- Best suited for time-series data, EG: interface statistics, router CPU
- Can also apply to network topology, EG: delay measurement between nodes
- Optimized data collection and optimized transport
- NETCONF/RESTCONF subscriptions can also be considered “streaming telemetry”

Event Driven Telemetry

- Data is pushed asynchronously from node based on state change or monitored event
- SNMP Traps, Syslog, Cisco EEM, Junos event scripts, and RMON are examples of existing event driven telemetry
- Modern approaches use YANG models and same structured encoding as periodic streaming telemetry
- BGP Monitoring Protocol (BMP) can also be thought of as event-driven telemetry

Flow Data

- Okay this isn’t modern
- Flow data is not tied to a specific state or attribute of a device, but to a network of interconnected devices is a critical measurement

What is not modern telemetry?

SNMP / NETCONF

- On demand and follows a “pull” paradigm, requiring periodic “GET” from collector
- Difficult to scale, especially using higher frequency collection
- Inefficient encoding and transport

Screen scraping

- We all may have been there, but also not “telemetry”
- Often employed when standard MIBs are not supported
- Common use case is collecting BGP RIB information

Bulk statistics gathering via ftp/scp/etc.

- Used when SNMP collection is not efficient for large data sets
- Typically large collection intervals, not suitable for near real-time applications

Peering telemetry at a glance

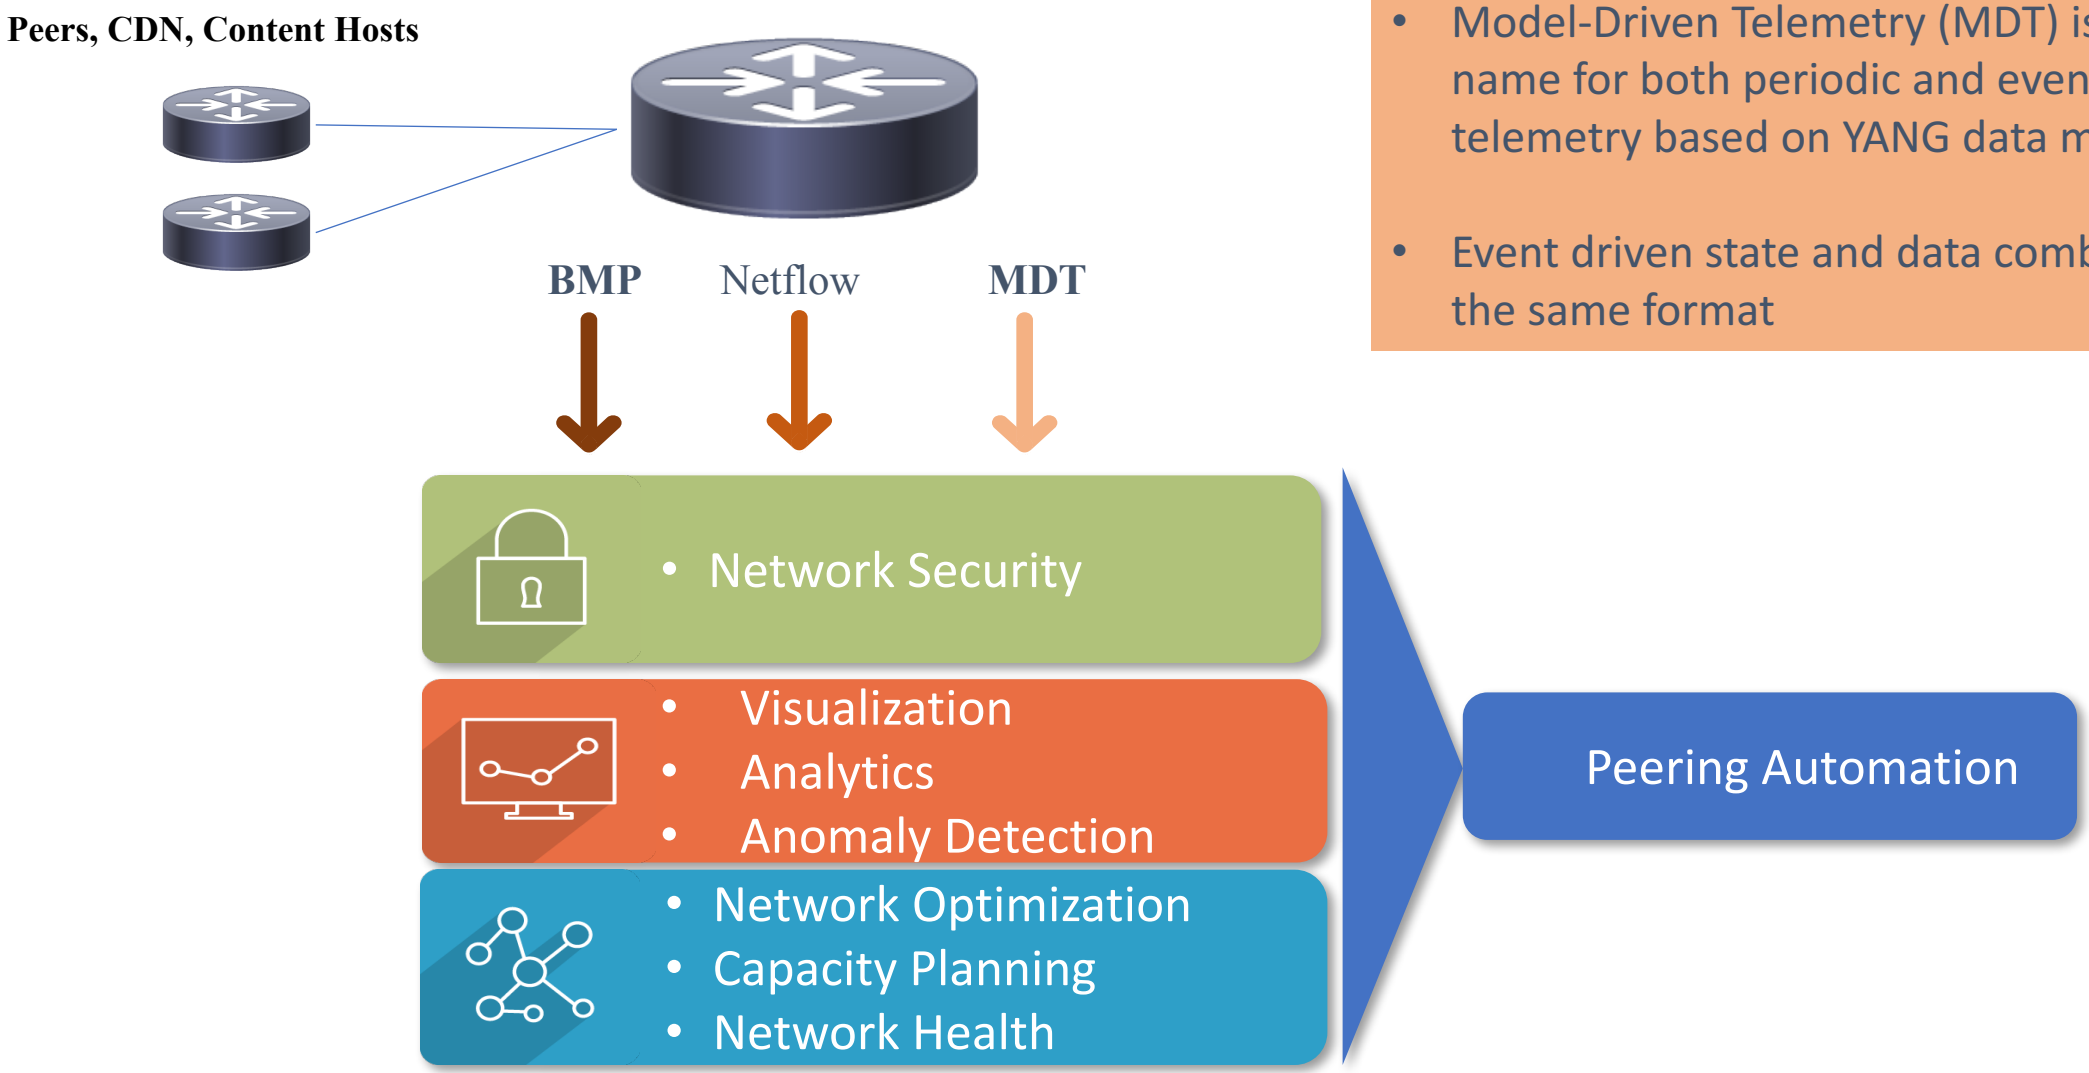

- Model-Driven Telemetry (MDT) is Cisco's name for both periodic and event based telemetry based on YANG data models
- Event driven state and data combined into the same format

Foundations of Model-Driven Telemetry

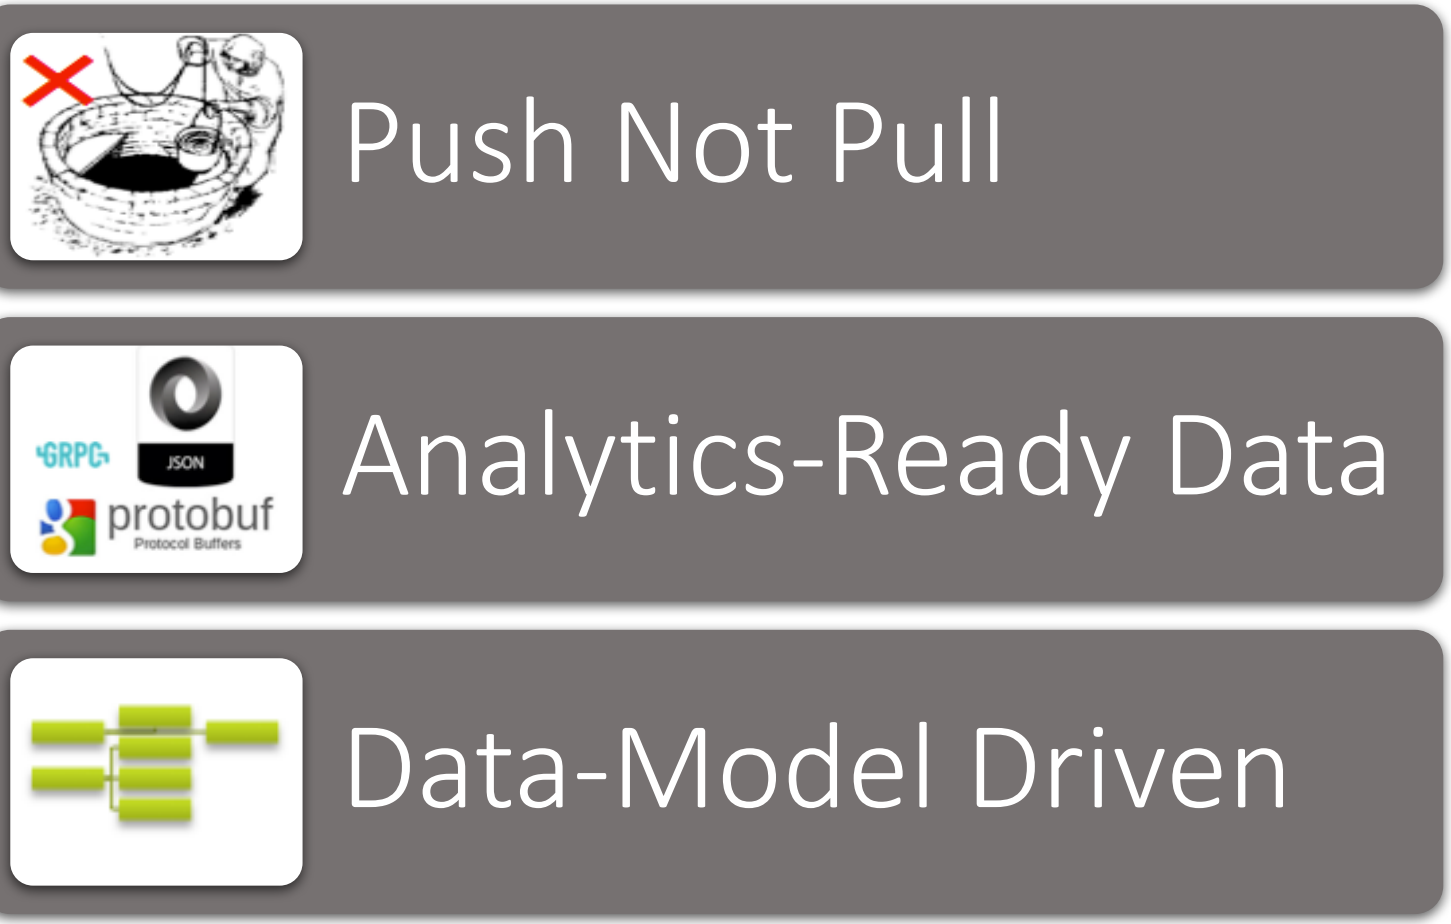

Performance

Automation

Model-Driven Telemetry for Peering

Higher Resolution Interface Statistics

- Quickly detect anomalies when coupled with thresholds or machine learning
- Increased visibility into traffic patterns
- Expose hidden oscillations
- See instant impact of network changes or maintenance events

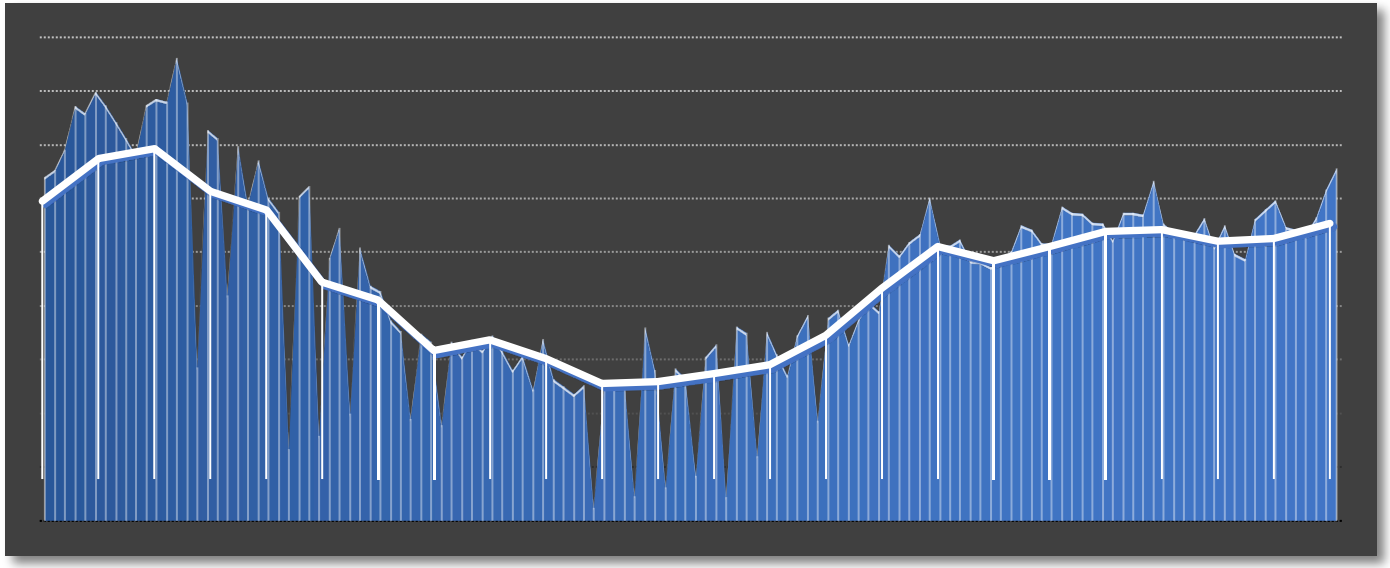

Network and Device Health Monitoring

- Monitoring queuing resources, can be important across peering or fabric where ingress/egress interfaces are the same speed. Similar in concept to datacenter microburst detection
- Monitor hardware FIB capacity and RIB memory

Fast Peering Anomaly Detection

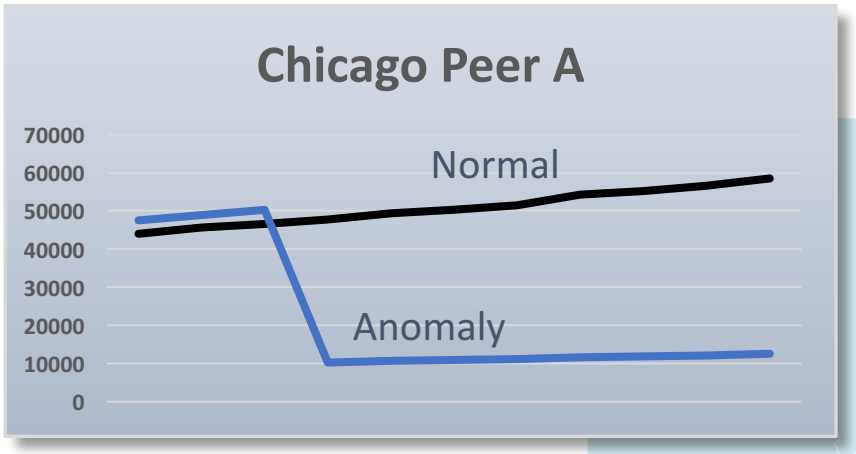

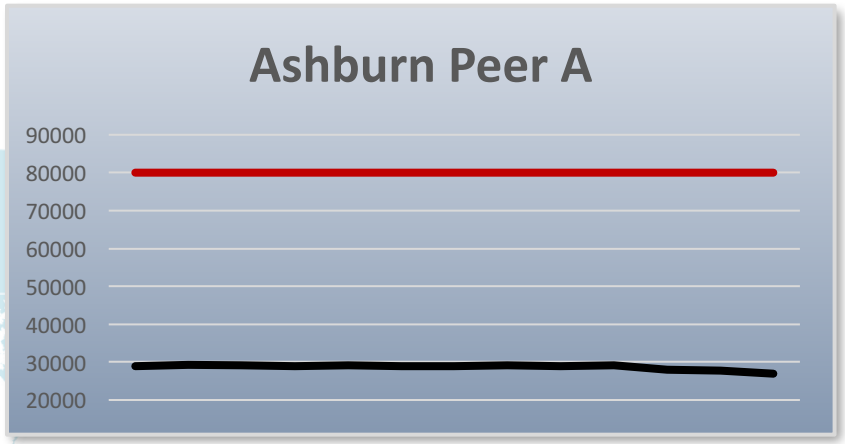

Unexpected shift from Chicago to Atlanta results in capacity exhaustion while Ashburn has unused capacity

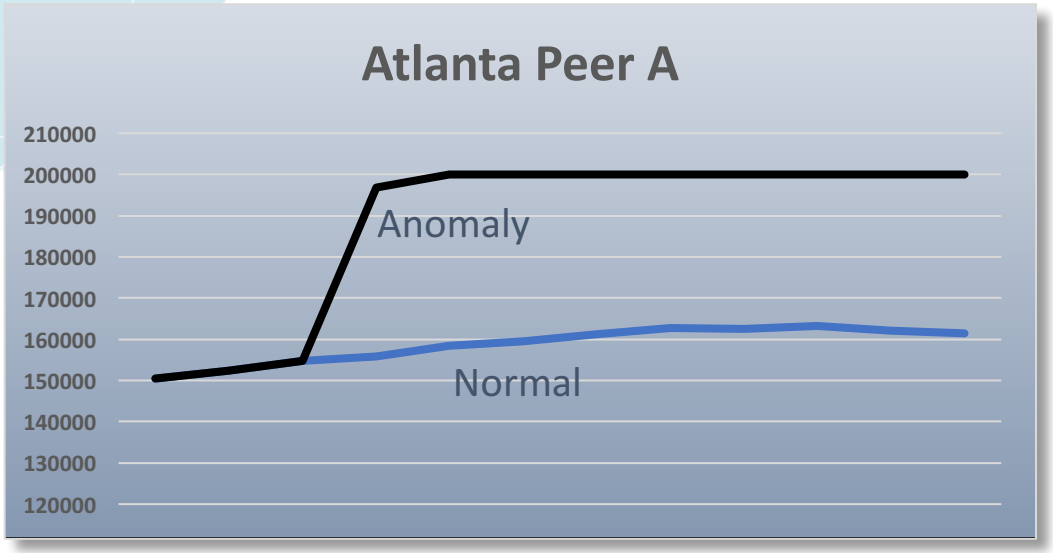

BGP Monitoring Protocol

Initial version defined by RFC 7854

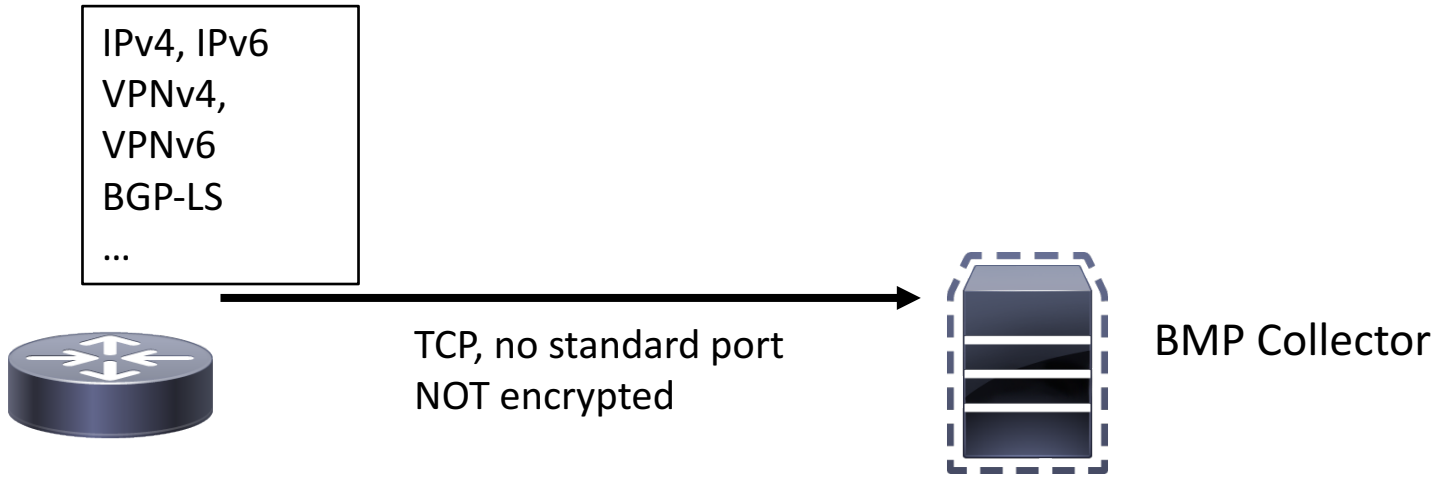

BMP Message Type	Data
Route Monitoring	Per-peer NLRI and ongoing NLRI updates
Statistics Report	14 periodic stats values, EG: denied prefixes, RIB counts
Peer Down Notification	Peer down, includes local/remote notification msg
Peer Up Notification	Peer in Established state, includes open msg
Initiation Message	sysName, sysDescr, additional info
Termination Message	Termination reason, additional info
Route Mirroring	Exact copy of BGP message and context

BMP Route Monitoring Points

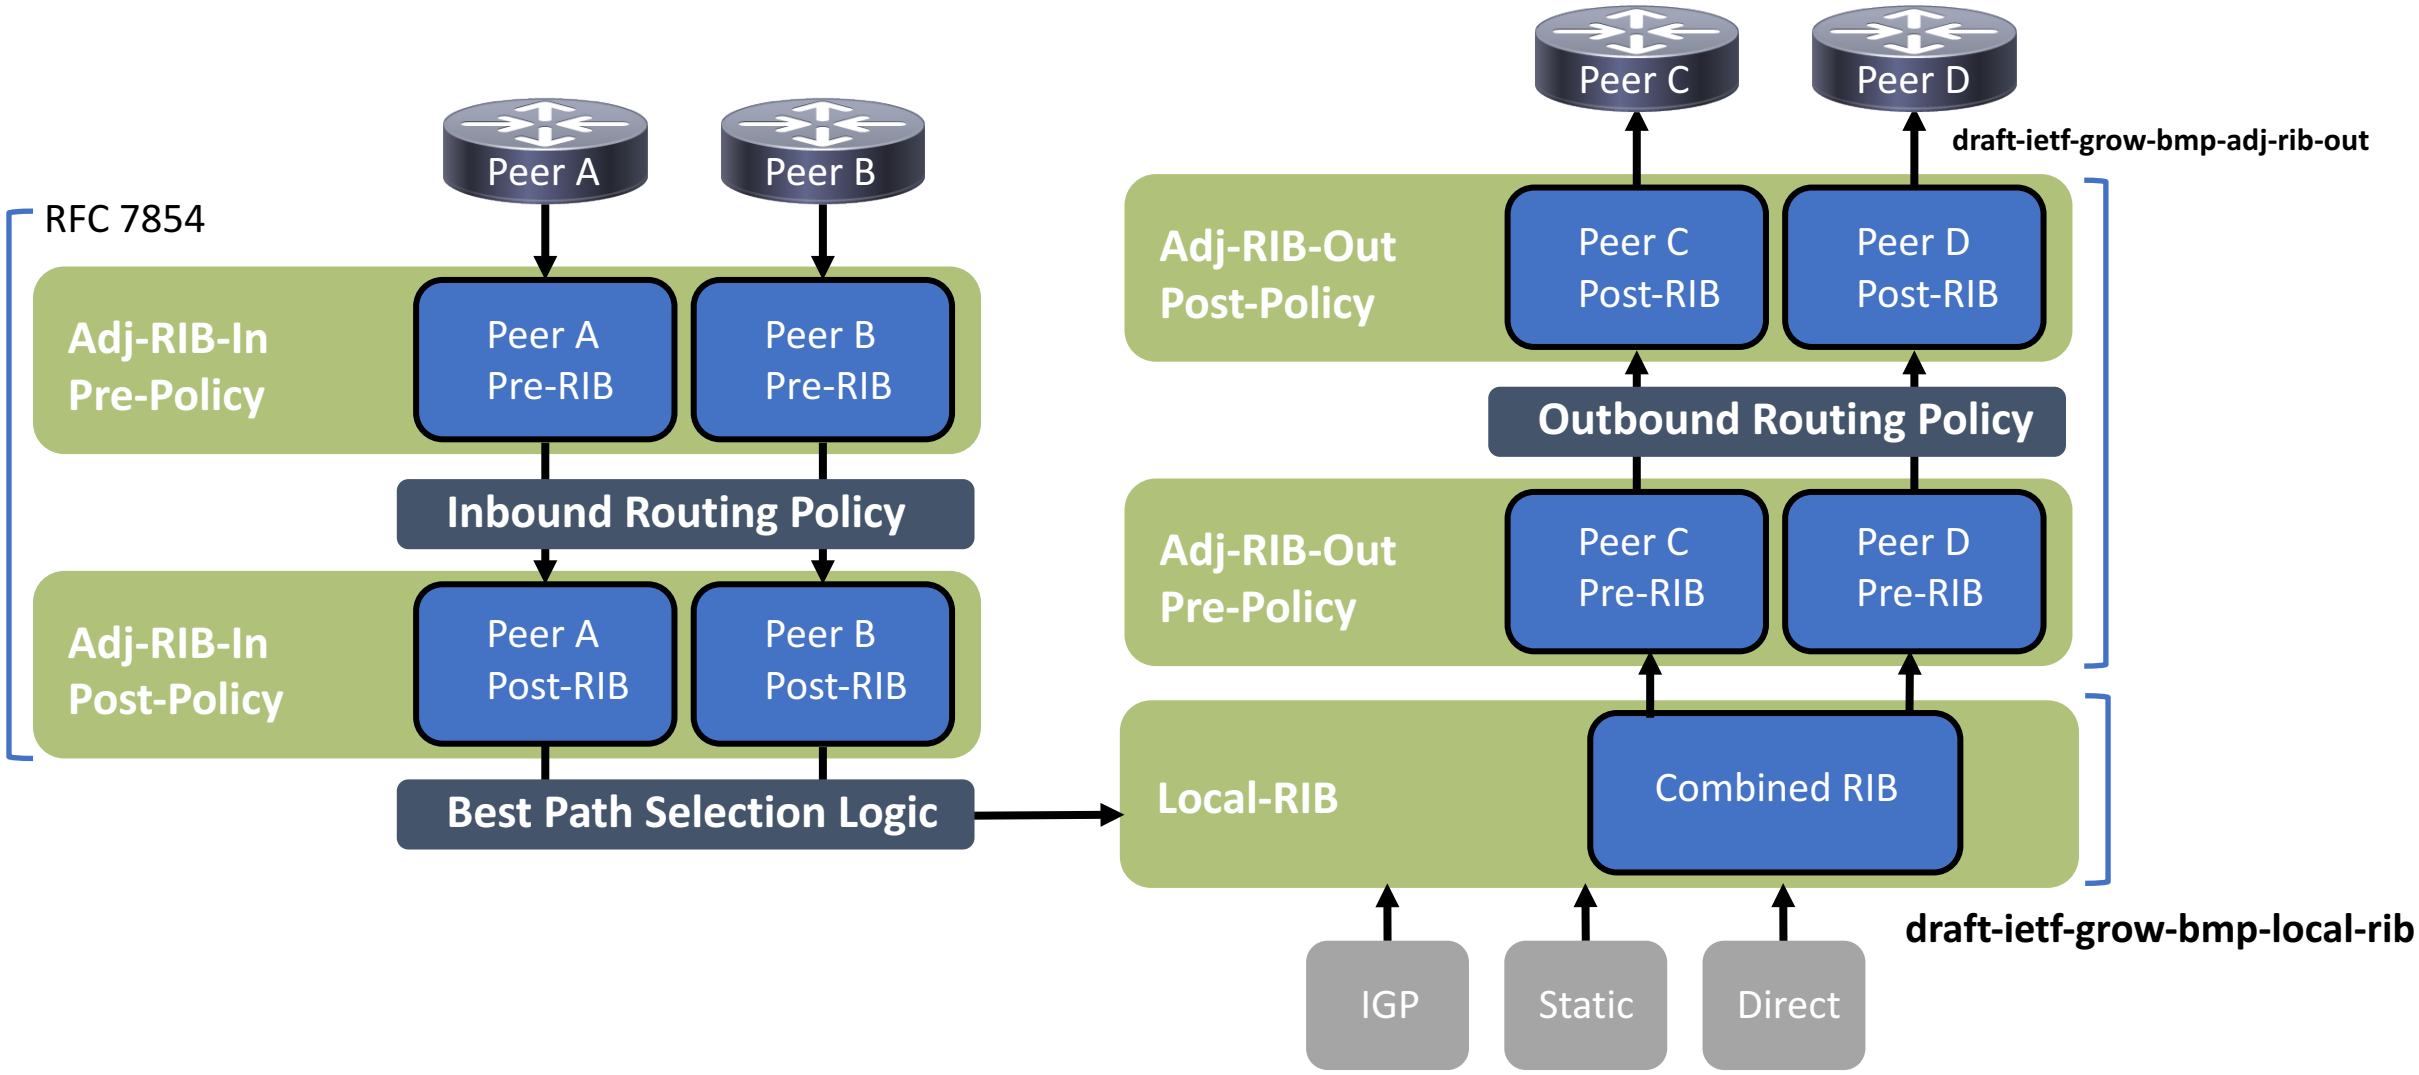

BMP RFC7854 Peering Use Cases

General Use Cases

- Monitor peers and prefixes for instability
- Statistic report data such as invalidated updates and rejected prefixes per-peer can be used to quickly identify issues
- Routing convergence analytics
- Monitor IGP topology using BGP-LS data

Pre-policy use cases

- Pre-policy sees what each peer is advertising to you prior to routing policy application
- Detect route leaks affecting your prefixes
- Simulate inbound routing policy changes
- Monitor for ”bad” attributes such as invalid/private ASNs, long ASN lengths, internal communities, bogon prefixes etc.

Post-policy use cases

- Monitor what the router installs in RIB from each peer
- Use diff from pre-policy to easily detect specific rejected prefixes
- Looking glass applications without router interaction

Local-RIB and Adj-RIB-Out Use Cases

Local-RIB use cases

- “Best” path in routing table as seen from each router. Collector does not need to know SPF or routing-policy to determine best path
- Enhances looking-glass capability without device interaction
- Tracking historical changes and performing root-cause analysis
- Validating inbound routing policy changes by comparing Adj-RIB-In and Local-RIB

Adj-RIB-Out use cases

- Catch policy change errors
- Along with remote route-server data can verify your prefixes are being accepted and re-advertised correctly without screen scraping and manual analysis

Advertisement Verification

Automated without human interaction

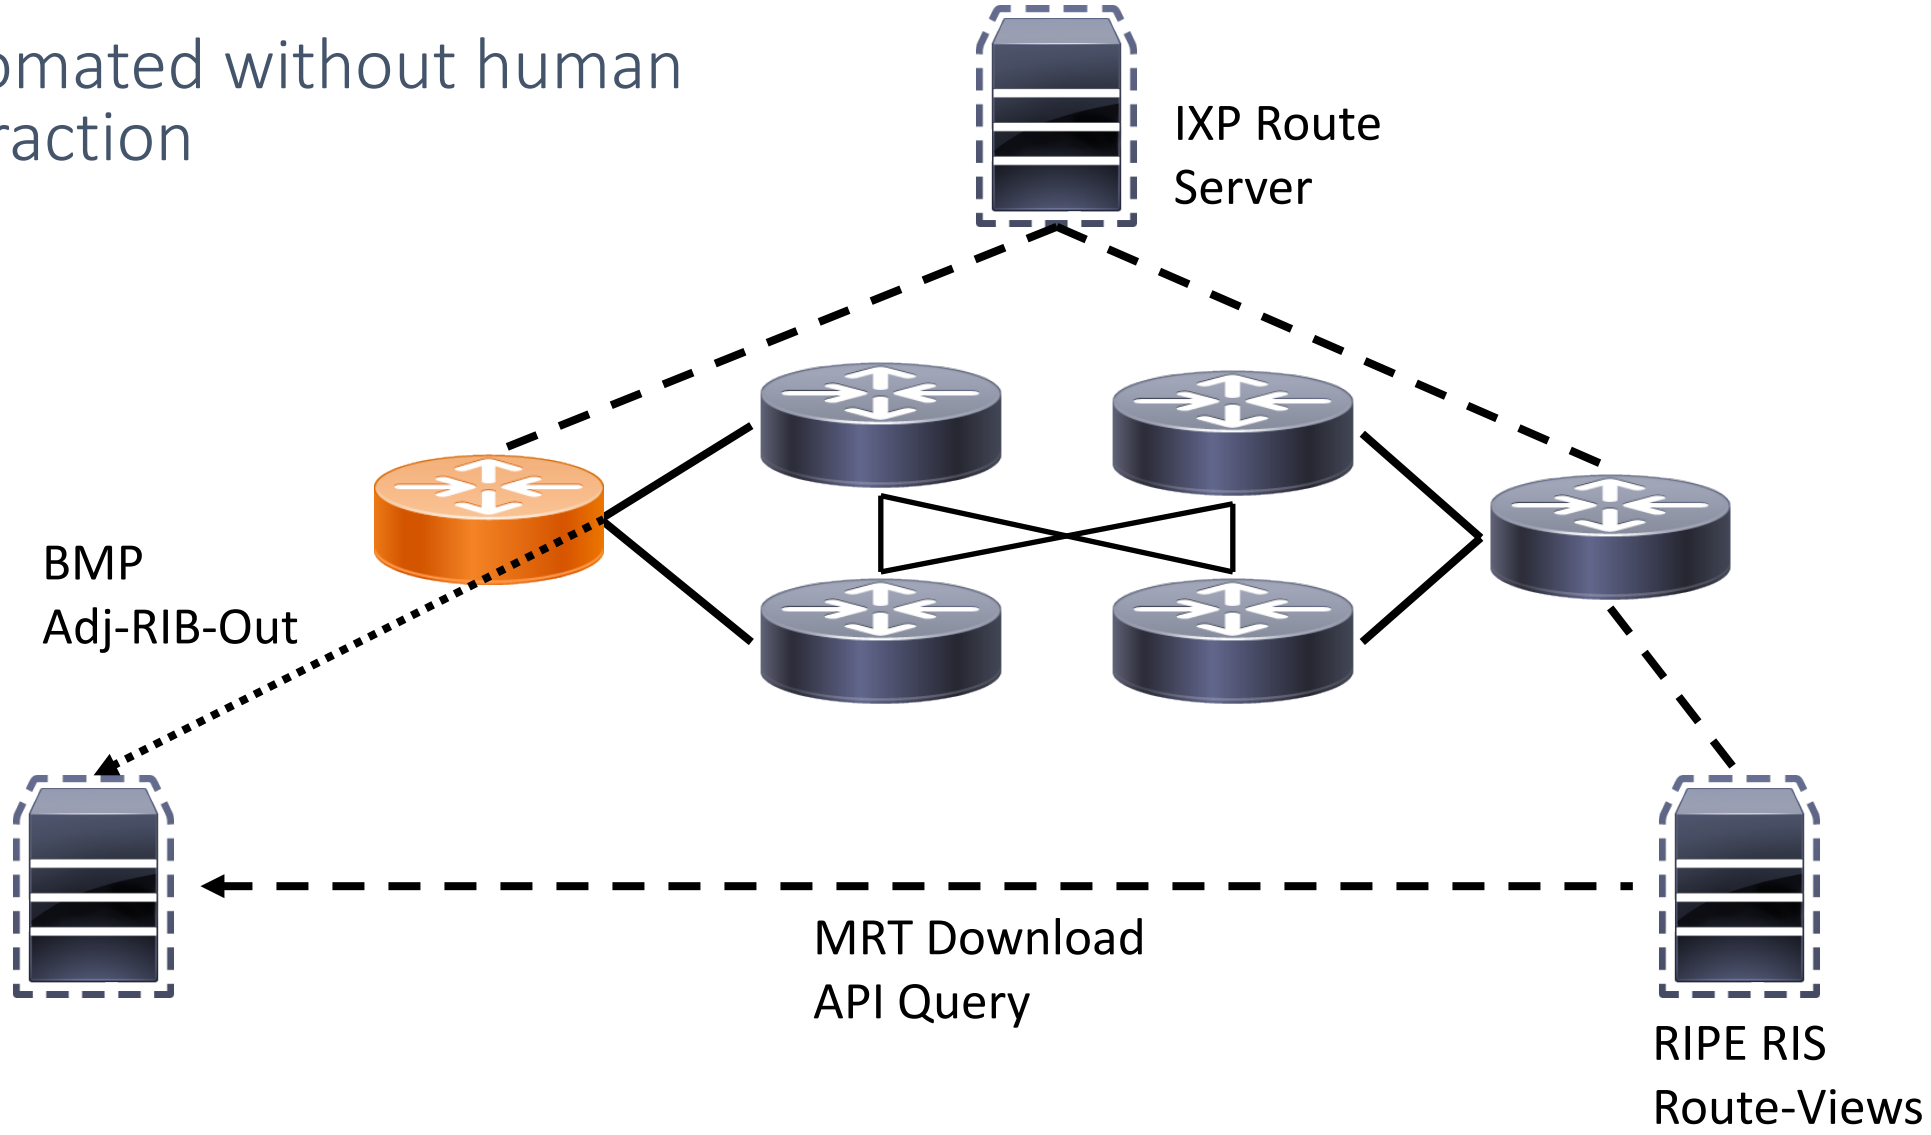

Getting Started

Model-Driven Telemetry

- IOS-XRv virtual routers support MDT
- **Pipeline** is an open-source MDT collector written in Go available at: https://github.com/cisco/bigmuddy-network-telemetry-pipeline
- Pipeline accepts Input via Cisco gRPC or JSON telemetry, Kafka. Output to Kafka, Prometheus, InfluxDB, text
- Many good telemetry blogs on https://xrdocs.com

BMP

- XR 5.3.4 for Adj-RIB-In post-policy, 6.2.2 for Adj-RIB-In pre-policy
- Local-RIB and Adj-RIB-Out under development
- SNAS, formerly OpenBMP, available at https://snas.io
- SNAS can output to Kafka or store in local database for use with SNAS UI and apps

SNAS


CISCO