



YANG-based Telemetry storage in TSDBs

IETF 118, November 2023

Kristian Larsson <kl@dev.terastrm.net>



YANG to Time Series Database Mapping

- A mapping of YANG-modelled data into label-set centric Time Series Databases
 - Convert NETCONF XML or RESTCONF JSON payload to store in TSDB

Original YANG Instance-Identifier:

`/interfaces/interface[name='eth0']/statistics/in-unicast-pkts = 5432100`

- Metric: `interfaces_interface_statistics_in_unicast_pkts`
- Value: `5432100`
- Labels:
 - `host = router-01`
 - `interfaces_interface_name = eth0`

Label-set centric TSDBs

- Rough concept shared by many TSDBs, not a strict standard
 - Metrics / time series are identified by a set of labels
 - Aka tags or dimensions – k/v
 - Sometimes time series has a “name” too, but it is similar to a label
-
- Influxdb, Prometheus, M3DB, OpenTSDB, VictoriaMetrics, Druid
 - Some variation in design

YANG + TSDB = <3

- Time Series Databases (TSDBs) are very common today
- YANG now prevalent on network devices, lots of config false
- Cisco IOS XR 7.6, 404k line of config false YANG
- Nokia SROS 20.10, 230k line of config false YANG
- Juniper 23.2, 509k line of RPC YANG (mostly for retrieving oper state)

Scale, cardinality & queries

- SP scale order of magnitude
 - Thousands PE routers * Hundreds of thousands / millions CPE
= ~Billions of metrics
 - First gen TSDB has Cardinality limit -> less of an issue today
- Use YANG model as rich index to get overview of data
 - Turn isolated islands into company wide telemetry resource
- Programmatic queries of TSDB data

To retrieve incoming unicast packets for the interface eth0:

•**InfluxQL**: `SELECT * FROM interfaces_interface_statistics_in_unicast_pkts WHERE interfaces_interface_name = 'eth0'`

•**PromQL**: `interfaces_interface_statistics_in_unicast_pkts{interfaces_interface_name="eth0"}`

Goals & next steps

- Deterministic model-driven mapping for programmatic consumption
 - Complete the specification! Talk to me if you are interested in the topic ☐
- We have some POC-level running code, make further progress
- Fit well into existing TSDBs and systems
- Self-describing data, so no strict reliance on YANG for visualization
 - All existing tooling & system that we want leverage is not YANG-aware
 - Suitable for direct human consumption
 - But extra metadata can be retrieved using YANG model
 - Dashboards and navigation can be built / organized with metadata from YANG model
 - Like show all temperature sensors on one dashboard, grouped by geography
- Encoding of 100% of YANG is NOT a goal
 - For example, anyxml / anydata is impossible / hard to present in TSDB
 - That's why we call this a “mapping” and not “encoding” as for XML / JSON
- Align with draft-lindblad-tlm-philatelist-00