

Streaming telemetry



Vlad Hanciuta, Sergey Kolobov
iNOG::D
Dublin, Ireland
September 20, 2017

Agenda

1. Why?

- Why can't we just keep using traditional methods?

2. What?

- What else is available on the market?

3. How?

- How can we do that?

1. Why?


a. Why can't we just keep using traditional methods?

- We used to explore our networks by CLI, SNMP, syslog. Which is OK, but...
 - CLI - subject to change, difficult to script, fragile to automation
 - SNMP - not really scalable, hence not granular enough, poor feature coverage
 - syslog - not structured, not efficient

2. What?

- What else is available on the market?

	REST	gRPC	NETCONF	RESTCONF
Terminattr	✓	✓		
OpenConfig		✓	✓	✓

 The exact same open-source interface

- Terminattr is for all the raw state
- OpenConfig is for vendor-independent state
- Complementary approaches: Use OpenConfig whenever you can, Terminattr for everything else

2. What?

- What else is available on the market?

- Example: if you got tired of chasing bad fans in your switches

```
[a@b~]$ ./grpcdemo-1 -json -addr 10.10.10.209:6042 -subscribe /Sysdb/environment/cooling/status/
```

<Output omitted>

```
"notification": {  
  "path": "/Sysdb/environment/cooling/status/fan/Fan2/1",  
  "timestamp": 1474660832905749957,  
  "updates": {  
    "FanFault": false,
```

<Output omitted>

```
    "name": "Fan2/1",  
    "properties": {  
      "_ptr": "/Sysdb/environment/cooling/config/supportedFanProperties/SanAce40"},  
      "speed/value": 45,
```

<Output omitted>

2. What?

- What else is available on the market?

- Example: if you troubleshooting buffer congestions and latency

```
[a@b ~]$ curl localhost:6060/rest/LANZ/congestion or {config | error | globalBufferUsage}
{
  "congestionRecord": {
    "entryType": "UPDATE",
    "fabricPeerIntfName": "",
    "intfName": "Ethernet39",
    "portID": 141,
    "qDropCount": 0,
    "queueSize": 52166656,
    "switchID": 0,
    "timeOfMaxQLen": 0,
    "timestamp": 1499726827343980,
    "trafficClass": 1,
  }
}
```

<-Queue size in segments at time of congestion

<-Time of congestion in micro-seconds (UTC)

More info and fields description on:

<https://github.com/aristanetworks/goarista/blob/master/lanz/proto/lanz.proto>

2. What?

- **What else is available on the market?**

More examples:

- Interface counters:
 - /Sysdb/interface/counter/eth/slice/phy/ <linecard> /intfCounterDir/ <interface> /intfCounter/current
- Temperature sensors:
 - /Sysdb/environment/temperature/status/tempSensor/ <sensorname>
- show queue-monitor length:
 - /LANZ/globalBufferUsage
- show queue-monitor length drops:
 - /LANZ/congestion
- show hardware counter drop [rate]:
 - /Smash/hardware/counter/internalDrop
- show platform sand health: Combination of various paths:
 - /Sysdb/hardware/sand/{fap,fe}/status/<linecard>
 - /Sysdb/hardware/sand/system/config/cli
 - /Sysdb/hardware/sand/system/status/sand

2. What?

- What else is available on the market?

CPU/Memory/Throughput/Storage

- In a steady state, **Terminattr** uses on the order of **80MB to 400MB** of memory
- A typical switch in will generate on the order of **100/300 updates per second**
- With our HBase storage we measured **~500MB/day per device** from an average of 4kB/s on the stream

3. How?

- How can we do that?

Demo...

Thank you!
Questions?