AmLight ExP
Americas Lightpaths Express & Protect

Internet2 Technology Exchange - Dec 06th, 2022

Evaluating INT, JTI, and sFlow @ AmLight

Renata Frez – Senior Network Engineer - RNP/AmLight

Overview

- Introduction to AmLight
- Tools/Frameworks in use at AmLight
- Juniper Telemetry Interface (JTI)
- In-band Network Telemetry (INT)
- How does In-band Network Telemetry (INT) work?
- Identifying Bursts: SNMP x JTI x INT (Tests)
- But… What is within the bursts? Using sFlow
- New: INT Collector 2.0 – Detecting Microbursts
- Conclusion / Future Work

Introduction to AmLight

- AmLight Express and Protect (AmLight-ExP) (NSF International Research Network Connections (IRNC) program)
- 600Gbps of upstream capacity between the U.S. and Latin America, and 100Gbps to Africa
- NAPs: Florida(3), Brazil(2), Chile, Puerto Rico, Panama, and South Africa
- Routers: Juniper and RARE/Freerouter
- Switches: Brocade, Dell, Corsa, Noviflow, and P4 Whiteboxes
- Production SDN Infrastructure since 2014:
  - Orchestrators: OESS and Kytos-NG
  - OpenFlow 1.0 and 1.3 as southbound interfaces
- Programmable Data Plane:
  - In production since 2021. Enables INT (In-band Network Telemetry) reporting
- **Next step: Autonomic network architecture!**
  More information: https://www.youtube.com/watch?v=CRnKKuP9I3Y

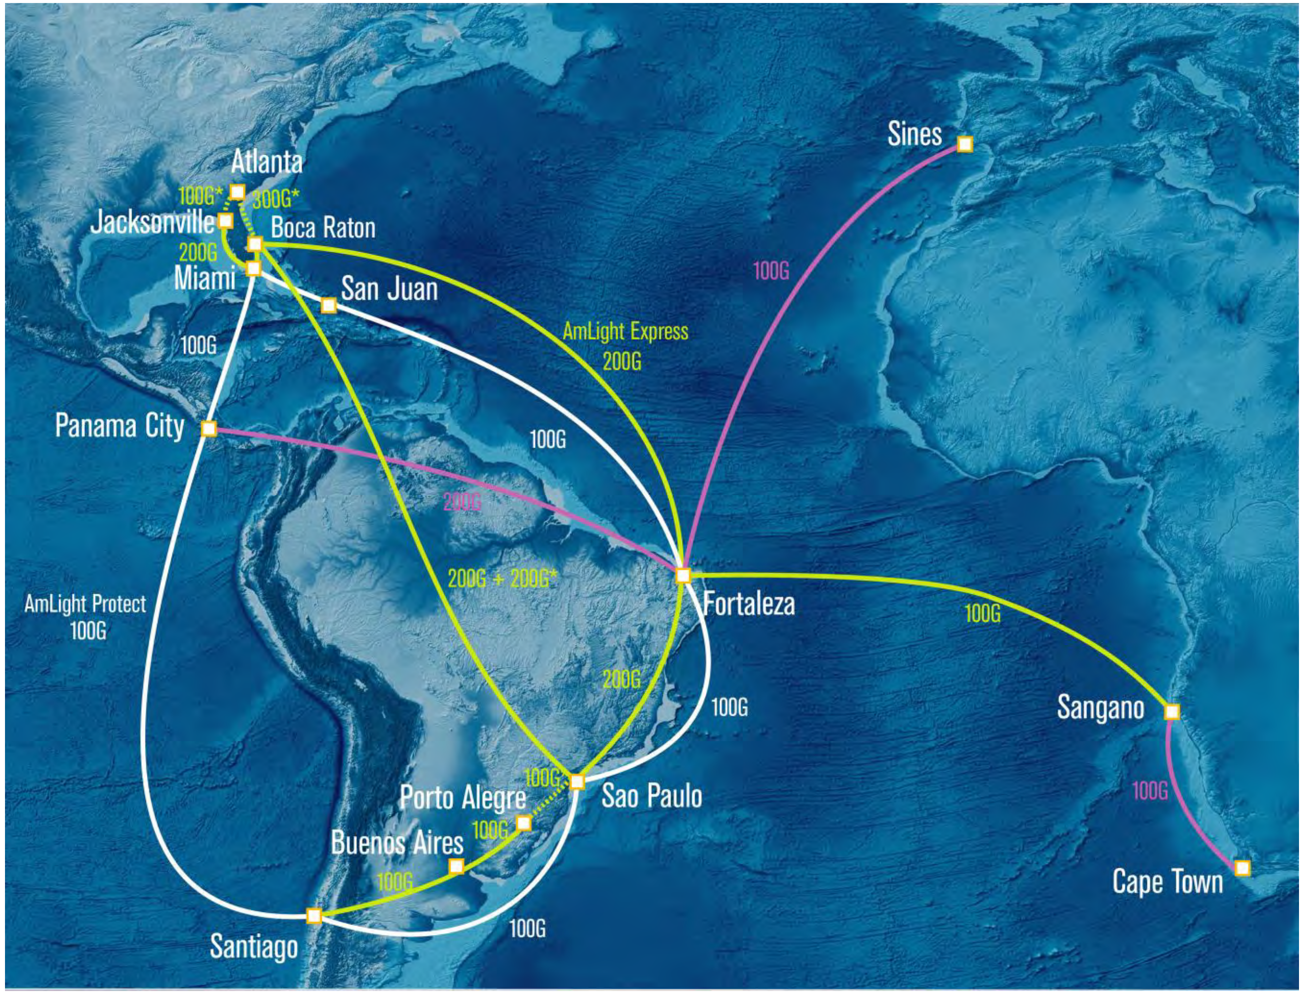

AmLight ExP
Americas Lightpaths Express & Protect

Tools/Frameworks in use at AmLight

Tool/Framework	Accuracy depends on:	Challenges:	Used for:
SNMP	➢ Data Plane counters collection interval. ➢ SNMP collector polling.	➢ Low interval ➔ higher CPU utilization. ➢ High interval ➔ lower accuracy.	➢ General monitoring.
sFlow	➢ Sampling rate.	➢ Low sampling rate ➔ more storage required ➔ higher CPU utilization. ➢ High sampling rate ➔ lower accuracy.	➢ Troubleshooting unusual events. ➢ TOP N reports.
Juniper Telemetry Interface (JTI)	➢ Data sending interval.	➢ Low interval ➔ more storage required. ➢ High interval ➔ lower accuracy.	➢ Environments that require more granular information.
In-band Network Telemetry (INT)	➢ Real time. Complete visibility.	➢ Processing all data collected in real time.	➢ Troubleshooting short-time events.

Juniper Telemetry Interface (JTI)

- As the number of devices and metrics generated by them has grown, the need for a non-impacting CPU tool has become critical.

- JTI is the Juniper telemetry solution that enables periodic data streaming as Protocol Buffers. In our environment, each device streams data every 2 seconds (lowest value for Packet Forwarding Engine Sensors).

- Examples of telemetry information streamed:
  - Interface counters, Optical counters, Routing information, Line Card information, and many others

In-band Network Telemetry (INT)

- INT is a P4 application that records network telemetry data in the packet while the packet traverses a path between two points in the network

- Since telemetry is exported directly from the Data Plane, the Control Plane is not affected:
  - Translation: you can track/monitor/evaluate EVERY single packet at line rate and in real time.

- Examples of telemetry information added:
  - Timestamp, ingress port, egress port, queue buffer utilization, sequence #, and many others

How does INT work?

1 – User sends a TCP or UDP packet unaware of INT

2 – First switch (INT Source Switch) pushes an INT header + metadata

3 – Every INT switch pushes its metadata. Non-INT switches just ignore INT content

4 – Last switch (INT Sink Switch) extracts the telemetry and forwards original packet to destination

5 – Last switch (INT Sink Switch) forwards the 1:1 telemetry report to the Telemetry Collector

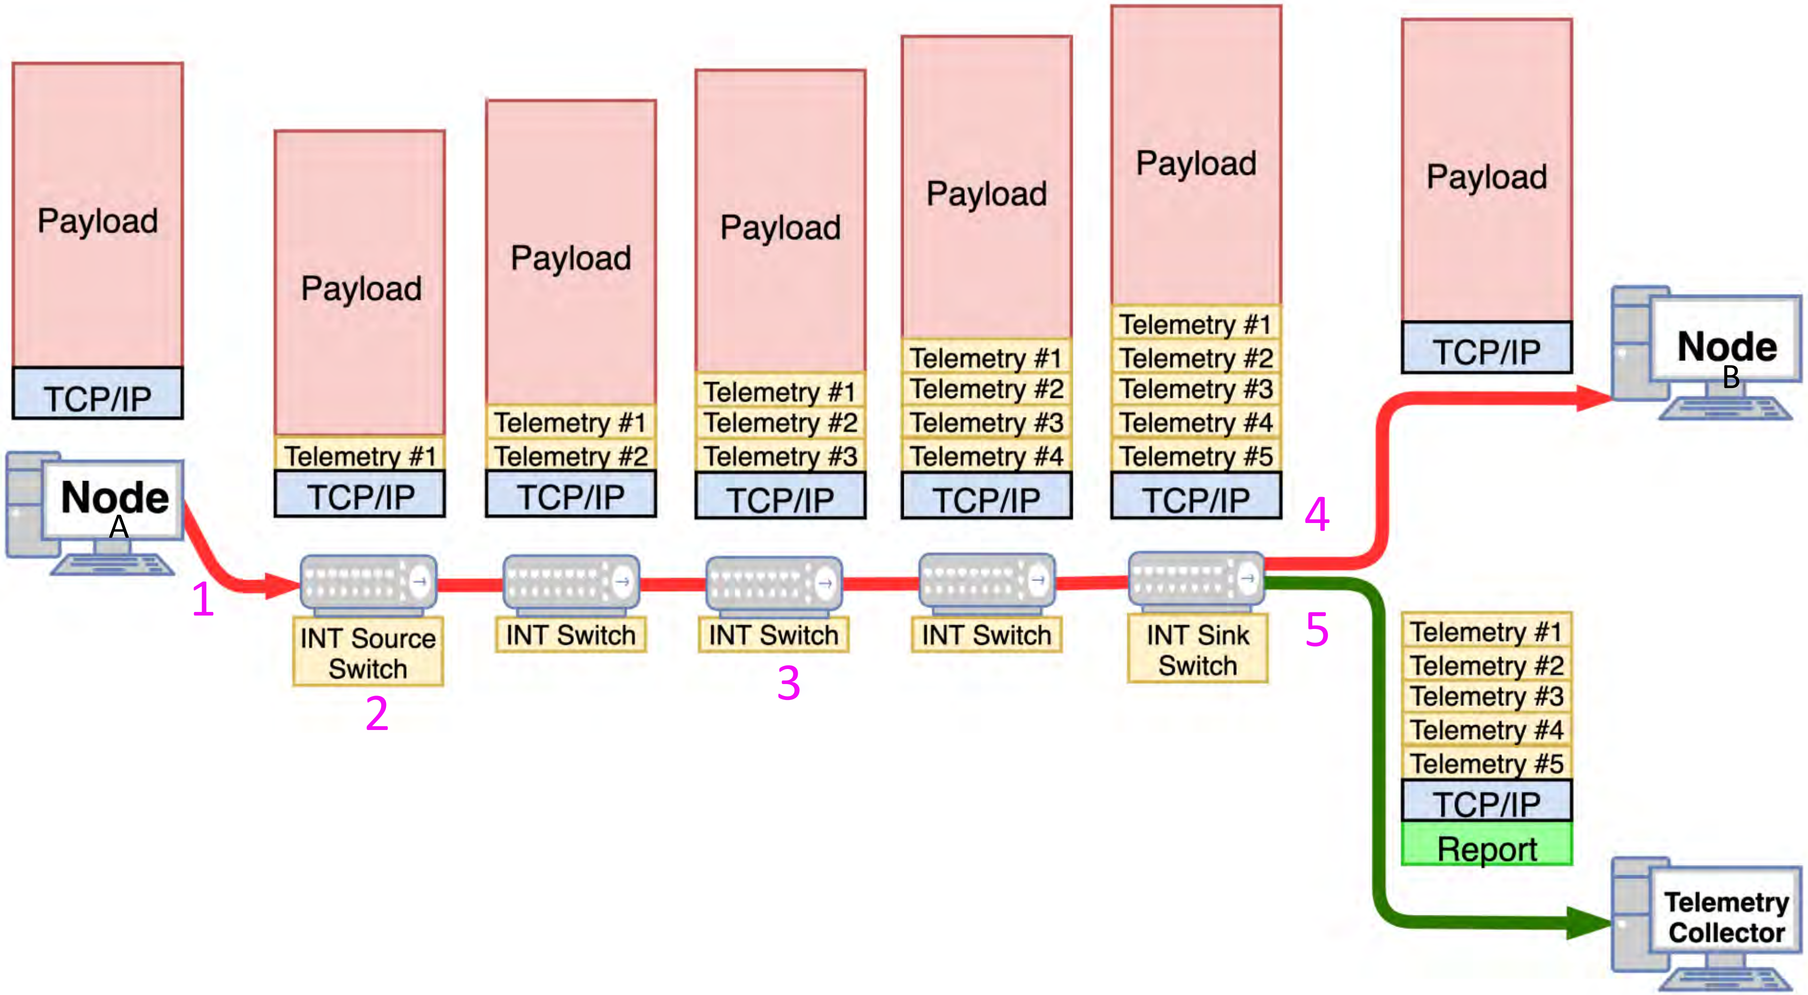

Simulations…

Demo Setup – Tools Comparison

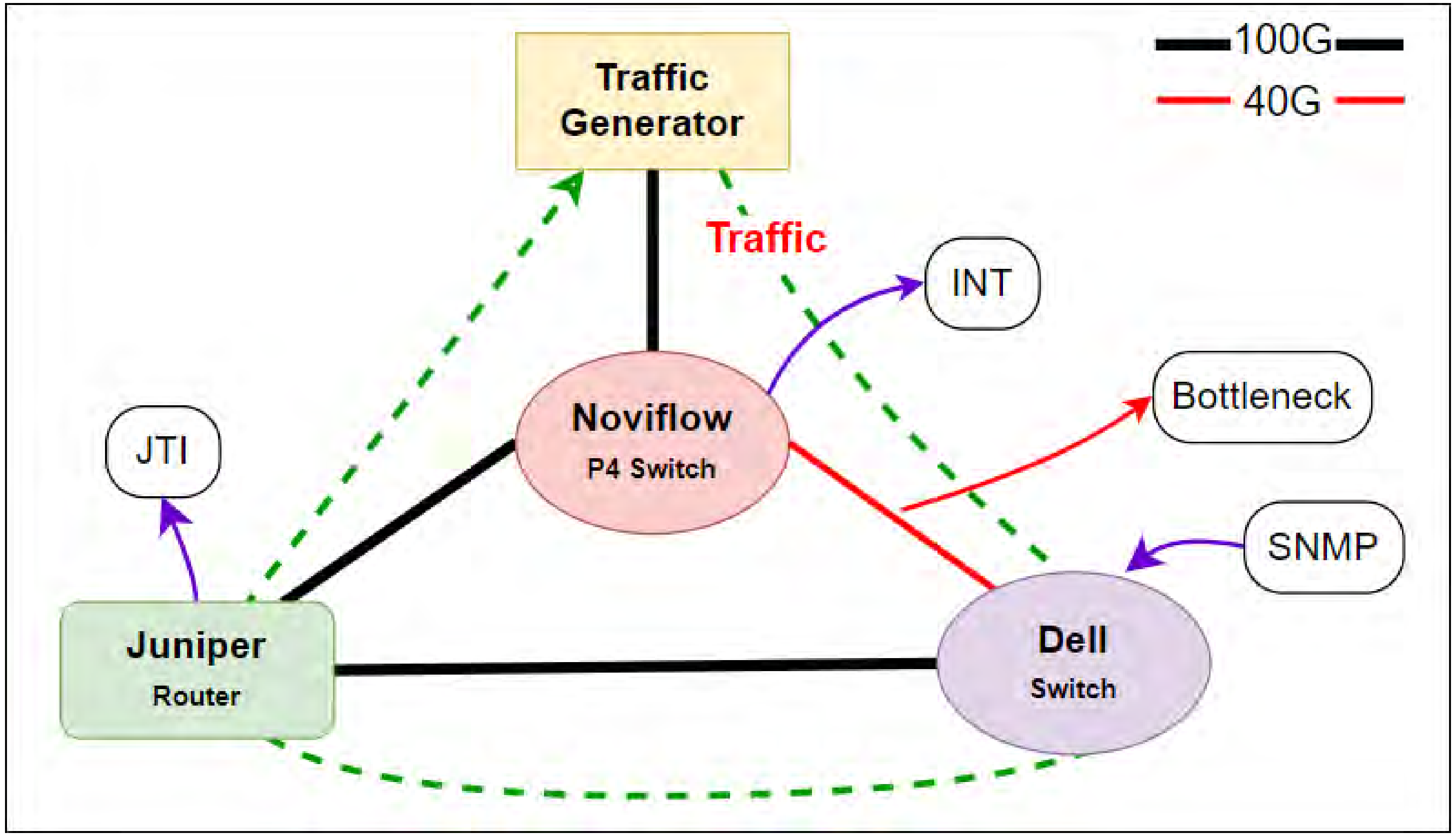

- EXFO Traffic Generator
- Dell (Switch OpenFlow) = SNMP polling every **14s** (lowest possible value).
- Juniper (Router) = JTI enabled sending telemetry every **2s** (lowest possible value).
- Noviflow (P4 Programmable Switch) = INT enabled for all packets, i.e., **real-time**. Database stores information every **100ms**.
- All graphs were taken from Grafana.

Identifying Bursts: SNMP x JTI x INT [Test 1]

- Interval: 30s.
- 2 Streams: Continuous and Burst.
- Continuous Traffic: 20G.
- Burst: 10x 30G.
- Burst duration: 2.5s.
- Interval between bursts: 2.5s.

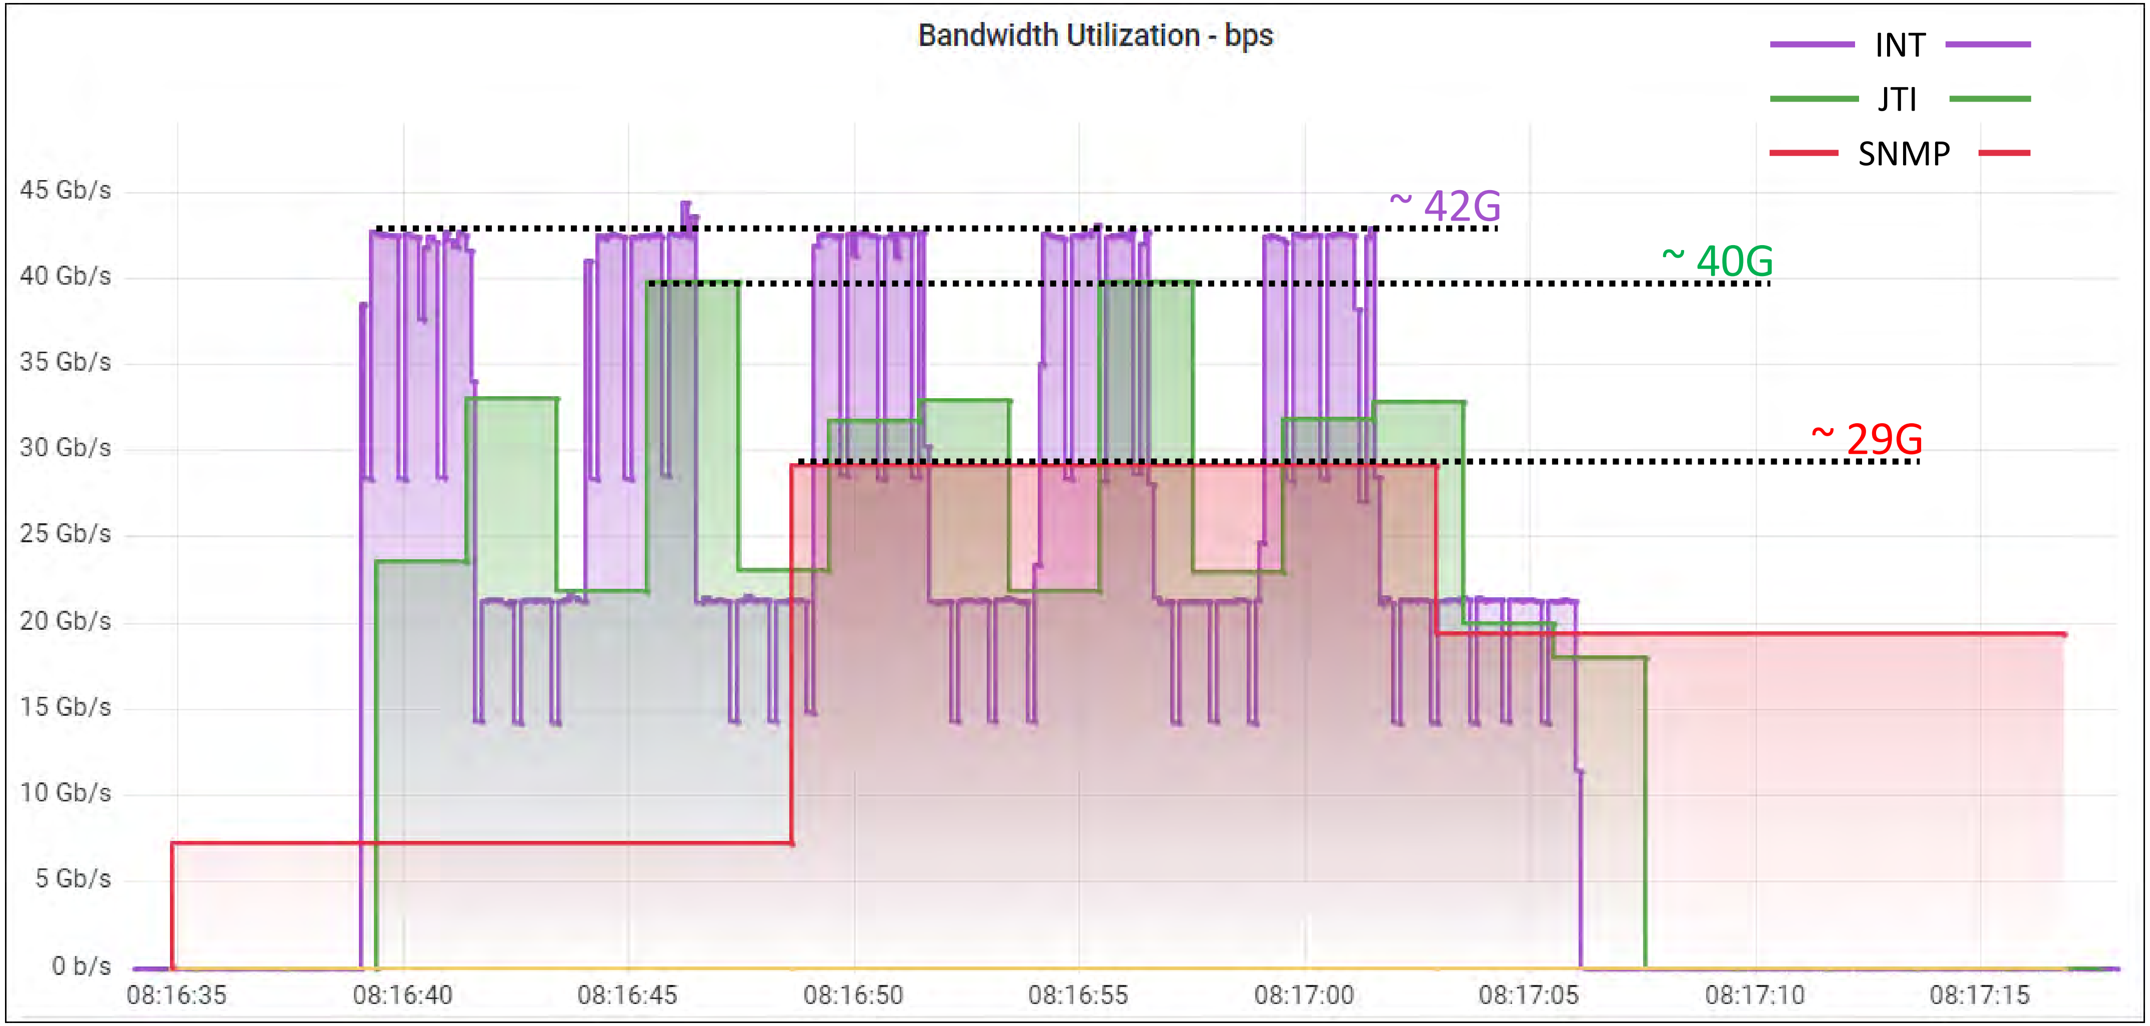

Identifying Bursts: SNMP x JTI x INT [Test 2]

- Interval: 30s.
- 2 Streams: Continuous and Burst.
- Continuous Traffic: 20G.
- Burst: **15x** 30G.
- Burst duration: **200ms**.
- Interval between bursts: **1.8s**.

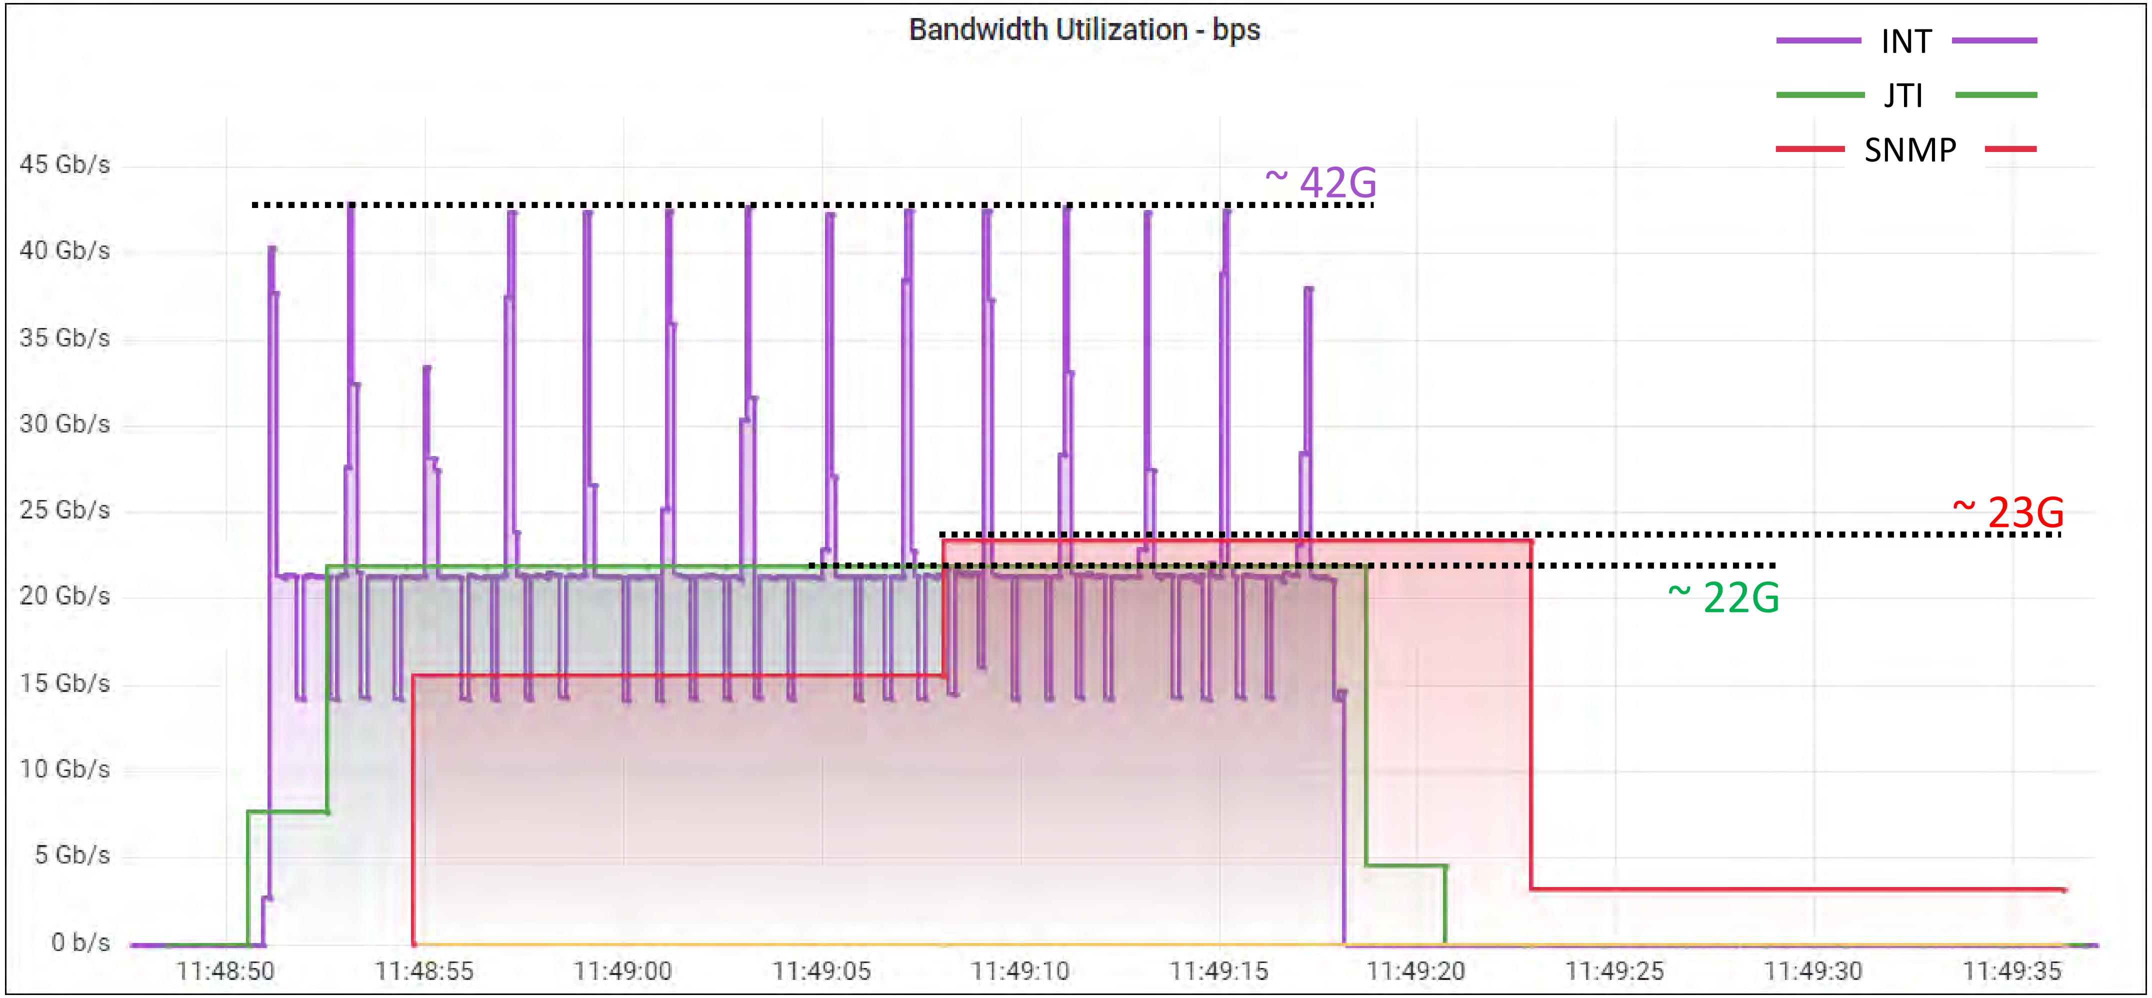

Identifying Bursts: SNMP x JTI x INT [Test 3]

- Interval: 30s.
- 2 Streams: Continuous and Burst.
- Continuous Traffic: 20G.
- Burst: **30x** 30G.
- Burst duration: **50ms**.
- Interval between bursts: **0.95s**.

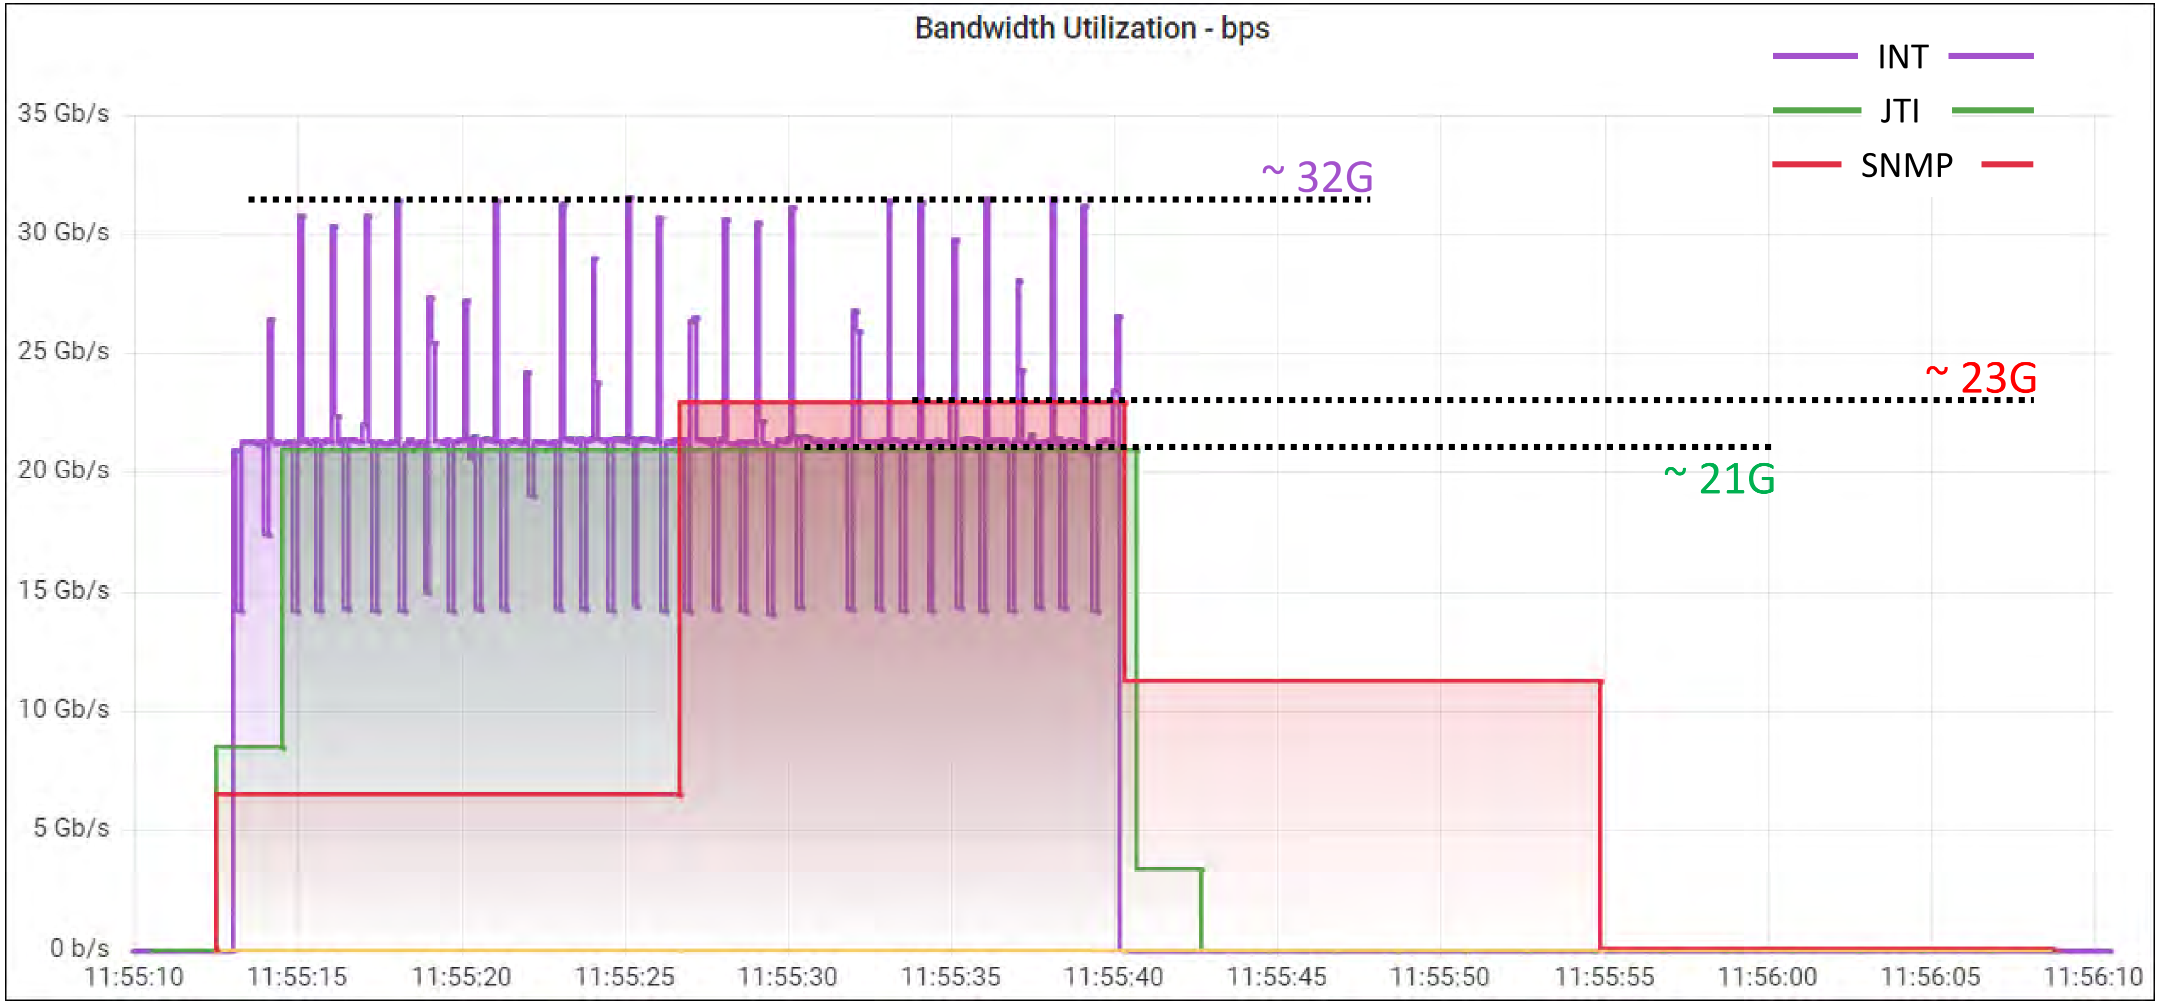

AmLight ExP
Americas Lightpaths Express & Protect

Identifying Bursts: SNMP x JTI x INT [Test 3]

- Interval: 30s.
- 2 Streams: Continuous and Burst.
- Continuous Traffic: 20G.
- Burst: **30x** 30G.
- Burst duration: **50ms**.
- Interval between bursts: **0.95s**.

Traffic Generator Results

Stream 1

	Average	Minimum	Maximum
Throughput (Gbit/s)	19.8177	19.7942	19.8473
Jitter (ms)	0.00015	< 0.00001	0.01276
Latency (ms)	0.03349	0.01748	0.40493
	Seconds	**Count**	**Rate**
Frame Loss	27	68360	9.1E-03
Out-of-Sequence	0	0	0.0E00

Stream 2

	Average	Minimum	Maximum
Throughput (Gbit/s)	1.1895	1.1599	1.2128
Jitter (ms)	0.00113	< 0.00001	0.38676
Latency (ms)	0.39435	0.01770	0.40517
	Seconds	**Count**	**Rate**
Frame Loss	27	115983	2.0E-01
Out-of-Sequence	0	0	0.0E00

But… What is within the bursts? Using sFlow

V4 DDoS - TCP ACK and ACK/PSH Floods

In an ACK flood attack or ACK-PUSH Flood, attackers send spoofed ACK (or ACK-PUSH) packets at very high packet rates that fail to belong to any current session within the firewall's state-table and/or server's connection list. The ACK (or A… **Show More**

Dest IP/CIDR 13

Baseline 9.76 Kpackets/s

147.46 Kpackets/s
1410% above baseline
57% above threshold

1.24 Gbits/s

114 Unique Src IPs

Alert | Ingress Interfaces | Traffic Patterns | Source Countries | Source Services | Packet Size Distribution

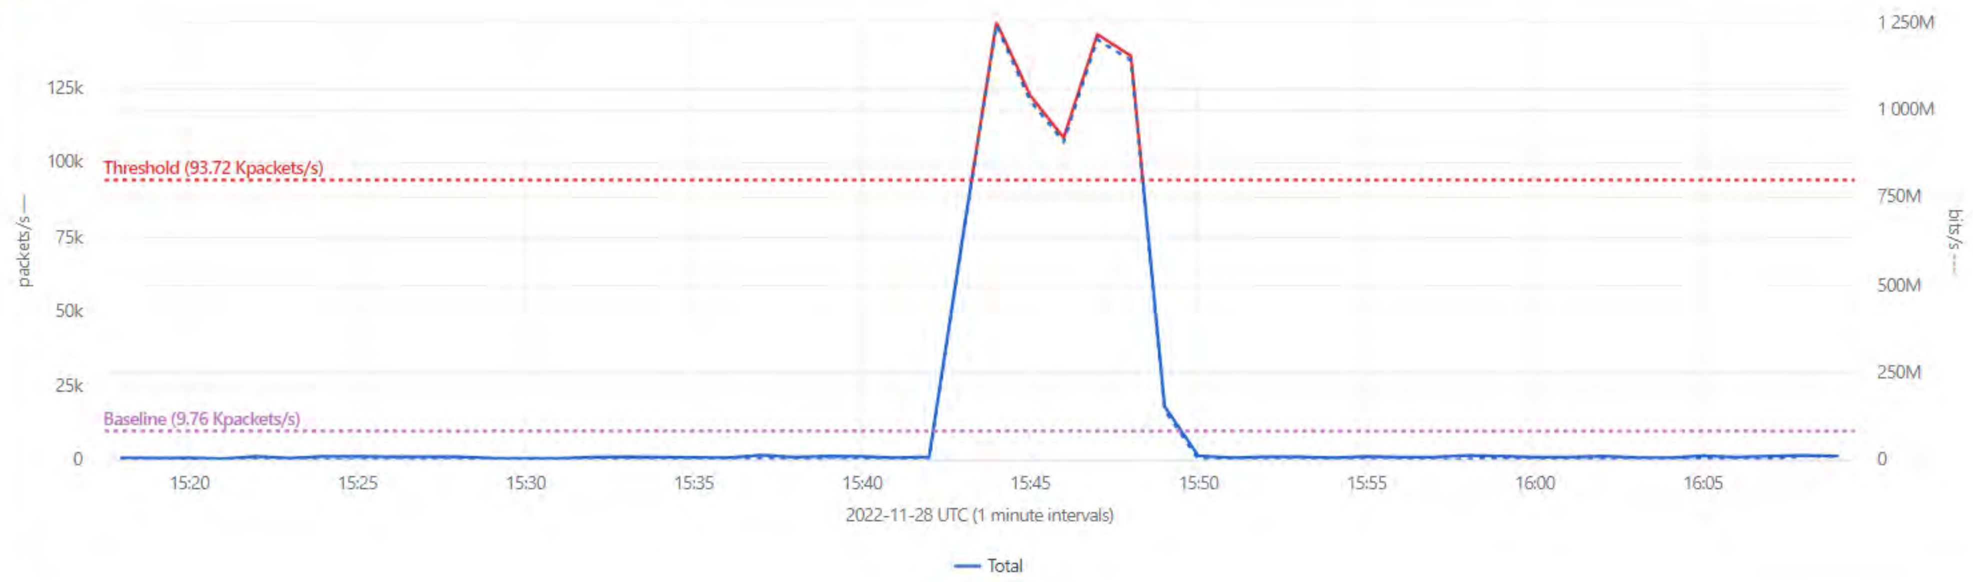

Severity
Major

Alert Start Time
2022-11-28 15:48

Event End Time
2022-11-28 15:50

Alert End Time
2022-11-28 16:10

Alert ID
200452091

Policy
V4 DDoS - TCP ACK and ACK/PSH Floods

Frequency
This alert has happened roughly 2x per week in the last 30 days
Show all Occurrences

Dest IP/CIDR
.13 has been found in 6 other alerts in the last 7 days.

AmLight ExP
Americas Lightpaths **Express & Protect**

Improvements for INT Collector

New: INT Collector 2.0 – Detecting Microbursts

- The AmLight INT Collector 2.0 will support detecting microbursts as short as 10ms.
- The figure shows 10 microbursts, each lasting 20ms, using up to 19Gbps Microbursts.

Start Time (UTC)	Duration (s)	Max BW (Gbps)
2022-10-09T13:10:37.304385768	0.02	16.35
2022-10-09T13:10:37.400937960	0.02	17.44
2022-10-09T13:10:37.499991784	0.02	18.88
2022-10-09T13:10:37.598316288	0.02	19.01
2022-10-09T13:10:37.696891136	0.02	18.97
2022-10-09T13:10:37.795097088	0.02	18.91
2022-10-09T13:10:37.893028608	0.02	19.09
2022-10-09T13:10:37.992322792	0.02	18.66
2022-10-09T13:11:58.794430952	0.06	53.41
2022-10-09T13:12:01.507265768	0.04	41.48
2022-10-09T13:13:21.666561768	0.04	20.83

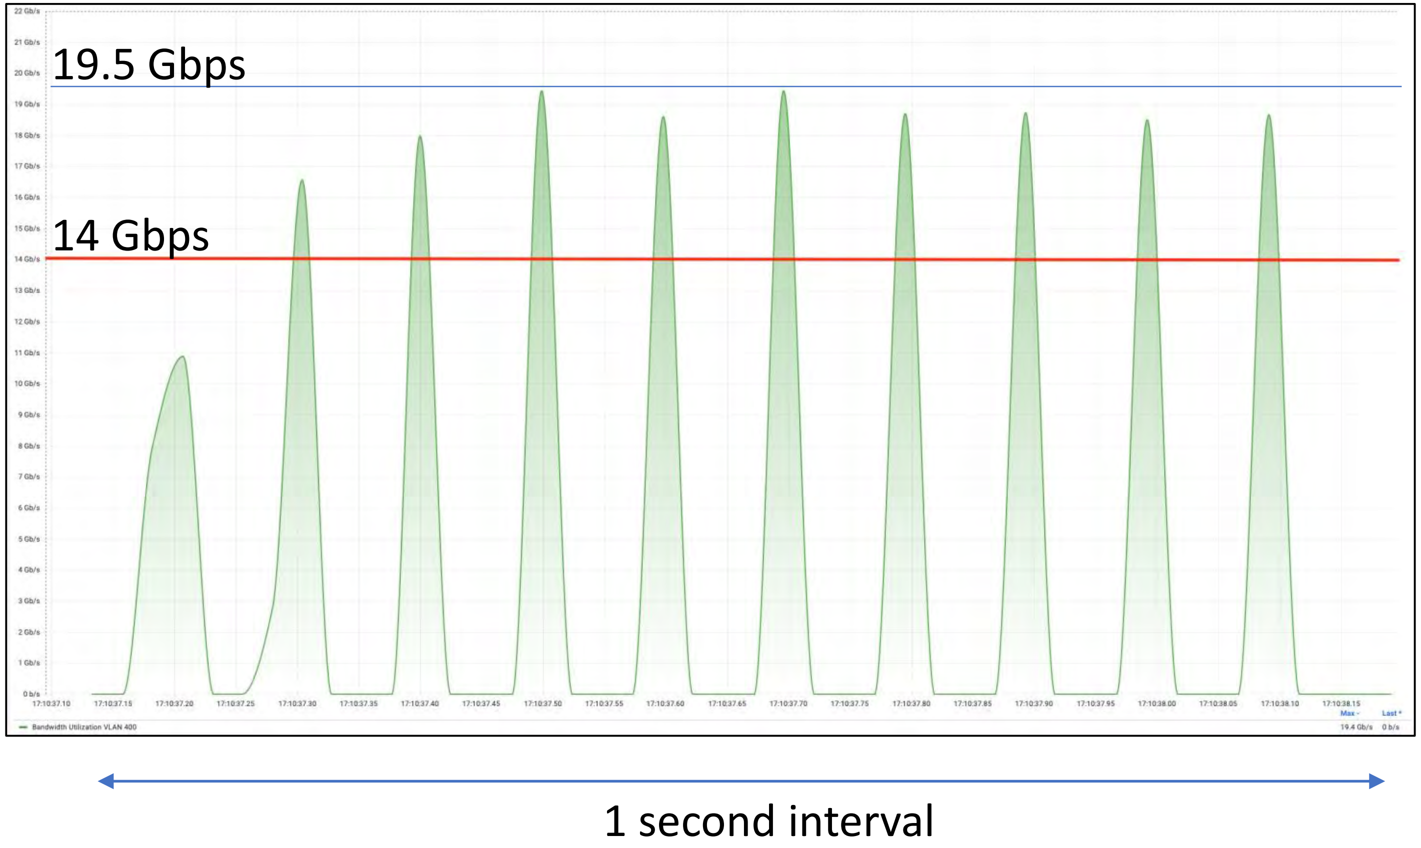

Conclusion / Future Work

- Monitoring every and any packet is possible with In-band network telemetry!
- JTI and INT have increased the network visibility beyond our expectations.
- Combining INT and legacy monitoring tools enables AmLight to track any performance issues and user complaints.
- New telemetry solutions will help achieve the Vera Rubin Observatory's Service Level Agreement (SLA).
- More tests are needed using sFlow to monitor interfaces' counters and compare the accuracy to other tools.
- Combining INT with learning tools will enable AmLight to move towards a closed-loop orchestration SDN network.
    - AmLight towards Autonomic Networking Architecture (ANA):
        - Self-configuration
        - Self-healing
        - Self-optimizing
        - Self-protection

AmLight ExP
Americas Lightpaths Express & Protect

Thanks! / Questions? / Comments?

Evaluating INT, JTI, and sFlow @ AmLight

Renata Frez – RNP/AmLight <renata@amlight.net>