



# In-band Network Telemetry @ AmLight: Our Solution

Italo Valcy <italo@amlight.net>

Jeronimo Bezerra <jab@amlight.net>

Julio Ibarra <julio@fiu.edu>

Renata Frez <renata@amlight.net>

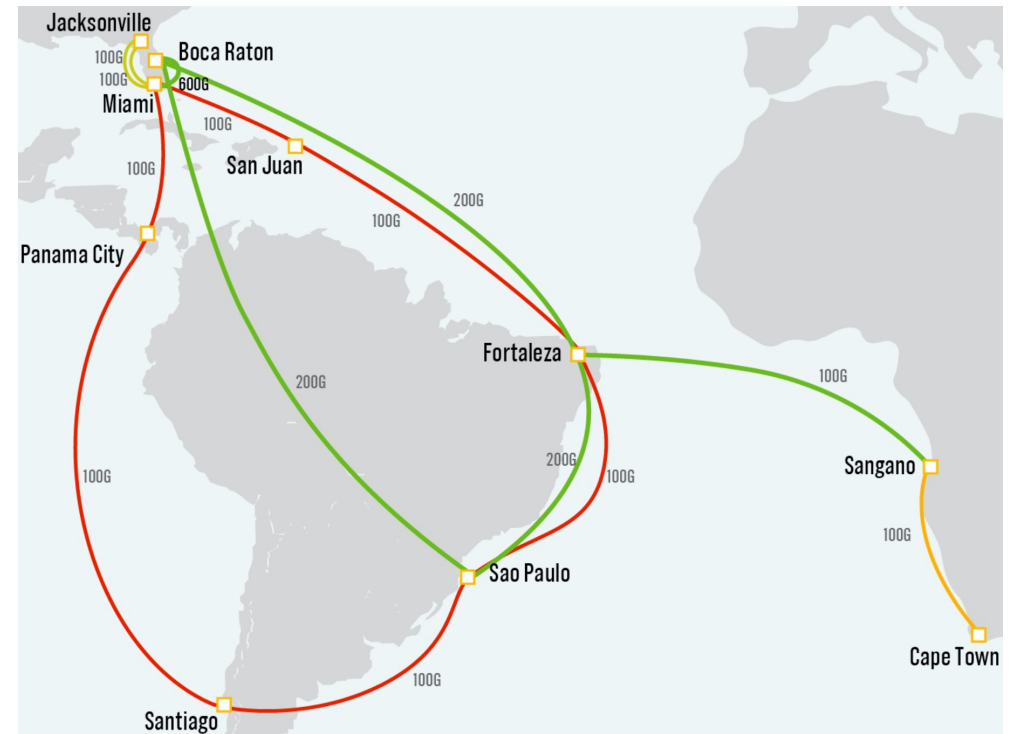
Vasilka Chergarova <vchergar@fiu.edu>

# Outline

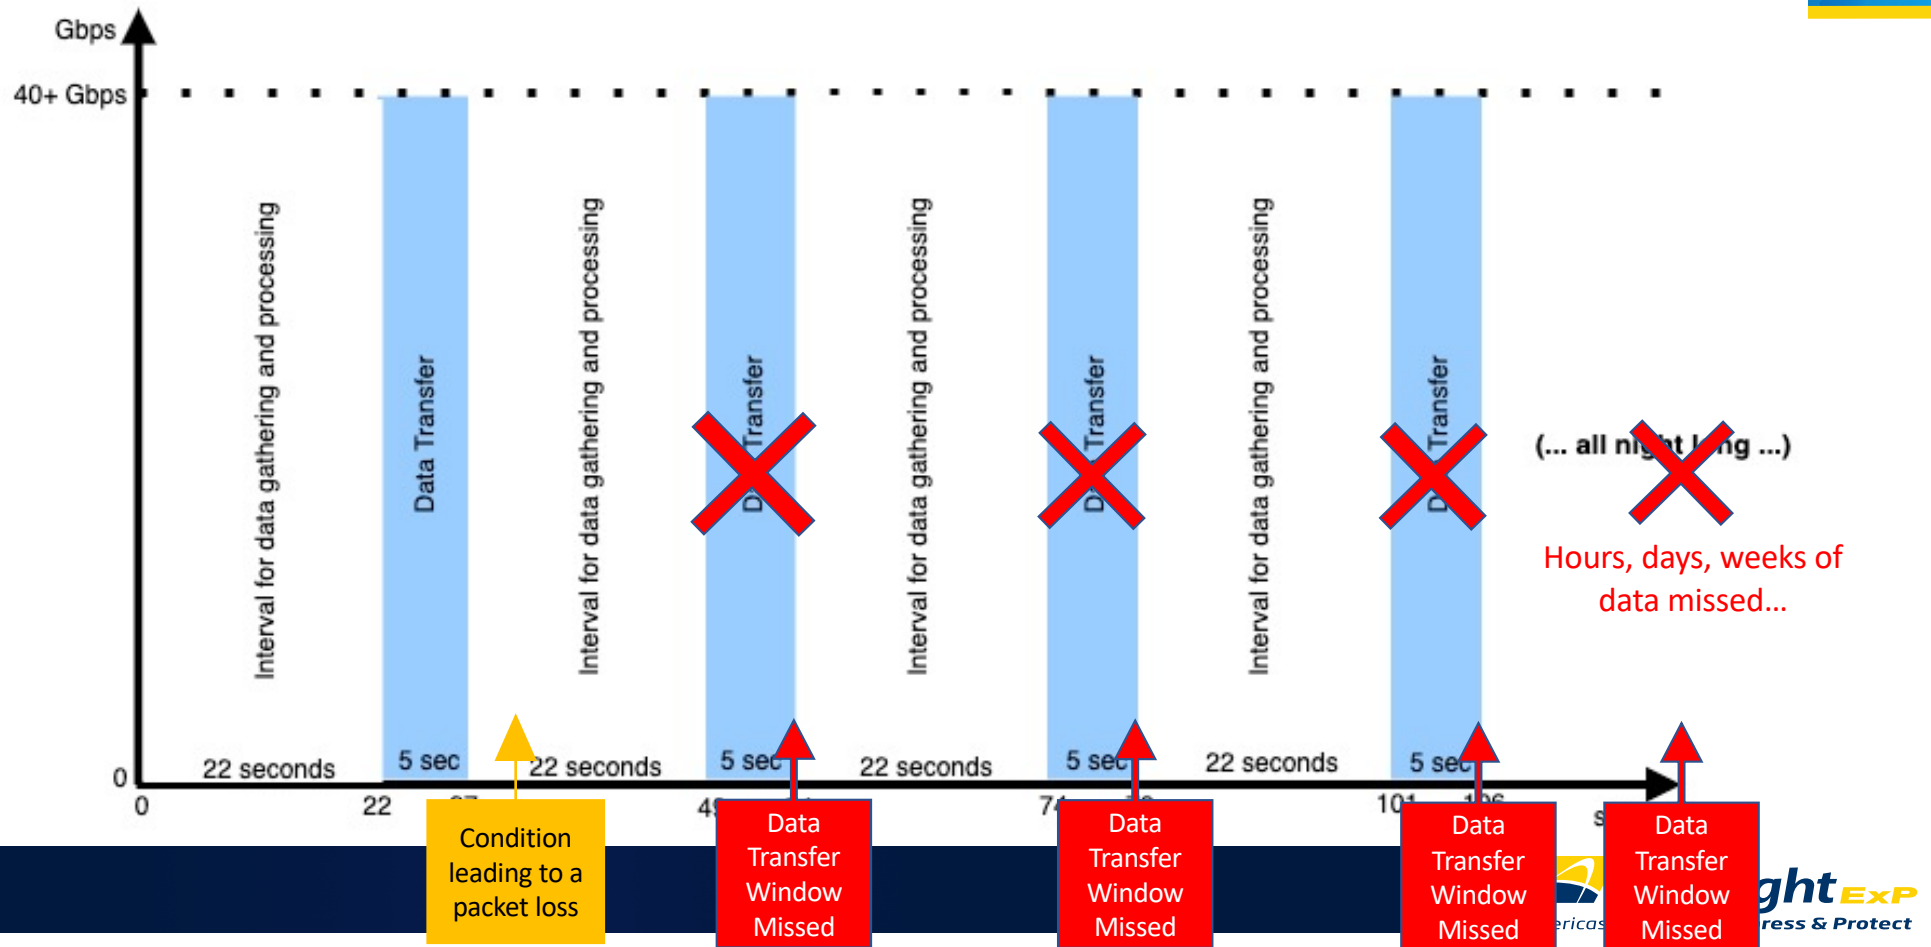
- Introducing AmLight
- The Science Driver
- What is In-band Network Telemetry (INT)?
- Challenges
- Deployment
- The Environment
- SC21 Demonstration
- Future

# 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
- Production SDN Infrastructure since 2014
- NAPs: Florida(3), Brazil(2), Chile, Puerto Rico, Panama, and South Africa
- Driver for deploying INT: **The Vera Rubin Observatory's Service Level Agreement (SLA)**



# The Use Case: Vera Rubin Obs's operation



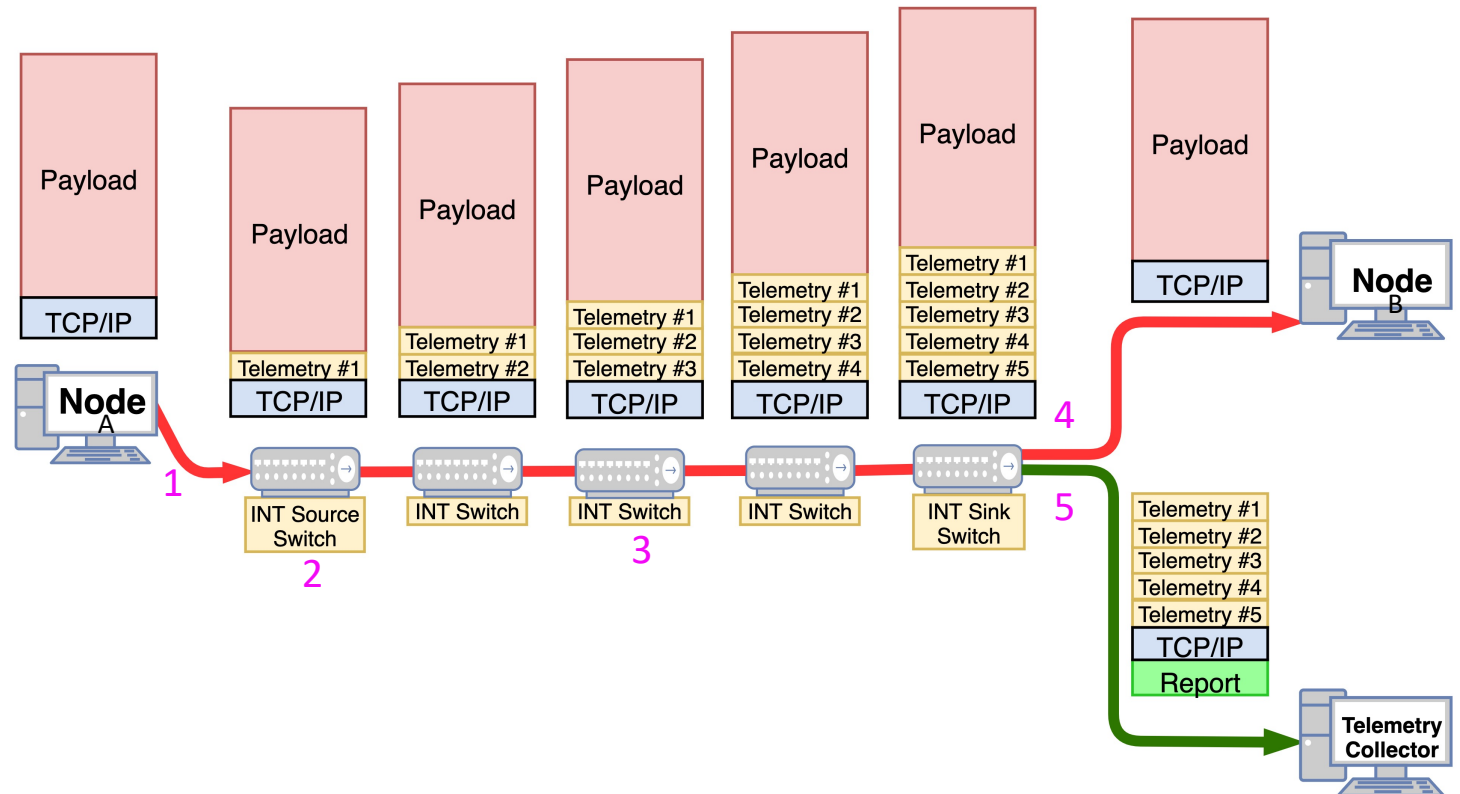
# In-band Network Telemetry (INT)

- 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.
  - The goal is to report the network state as seen by each packet.
- INT exports reports directly from the Data Plane: not impact to the Control Plane
  - Translating: *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 In-band Network Telemetry (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



# What INT metadata is being used and how?

- **Instantaneous Ingress and Egress Interface utilization**
  - Telemetry Collector monitors and reports egress interface utilization every customized interval (100-500ms)
  - Bandwidth monitored per interface & queue & VLAN
- **Instantaneous Egress Interface Queue utilization (or buffer)**
  - Useful for evaluating QoS policies
  - Useful for detecting sources of packet drops
- **Per-Node Hop Delay**
  - Useful for evaluating sources of jitter along the path
  - Useful for mitigating traffic engineering issues (under and over provisioned links)
- **L1/L2 Path Tracing**
  - EVERY packet and recording changes
  - Useful for detecting LAG or ECMP hash errors/mismatches and detecting unstable links
  - Path taken even reports the **egress queue ID**

## Challenges



# Challenge 1: Lack of commercial solution available


- AmLight-INT Project funded by NSF in 2018
- Collaboration between **FIU** and **NoviFlow** to expand AmLight SDN network towards an INT-capable domain
  - NoviFlow expanded the NoviWare OS to support INT following FIU's requirements
  - FIU developed the telemetry collector and evaluated the NoviFlow switch
- Characteristics of the NoviFlow switches @ AmLight:
  - Barefoot Tofino chip:
    - Fully programmable
    - 32 x 100G interfaces
- P4/INT 1.0 specification being followed

## Challenge 2: Receiving telemetry reports

- 100Gbps with 9000-Bytes packets → ~1.5M packets per second
- At AmLight, 4-8 switches connect Chile to the U.S.
- Telemetry reports have up to 300 bytes
- Each user packet triggers a telemetry report (1:1)
- 4.5Gbps of telemetry report for each 100Gbps flow
  - Each switch creates a single flow (No hashing possible)
- Solution in place: eBPF/XDP (eXpress Data Path)

# XDP - eXpress Data Path


- A thin layer at lowest levels of network stack for incoming packets
  - Not a bypass
- Run-time programmability via “hook”
  - No need to recompile the Kernel
- Comparable to DPDK but simpler
  - No need for dedicated CPUs, huge memory pages, hardware requirements, or licensing
- Can offload instructions to supported NICs
  - Examples: Netronome and Mellanox
- Use by service providers for DDoS mitigation
  - 20Mpps per node documented!
- Performance improvements observed:
  - From 5kpps with Python3 and C (user-space)
  - To 3 Mpps with XDP and one CPU



Source: [https://github.com/iovisor/bpf-docs/blob/master/Express\\_Data\\_Path.pdf](https://github.com/iovisor/bpf-docs/blob/master/Express_Data_Path.pdf)

# Challenge 3: Storing telemetry reports of interest

- Not feasible to save all telemetry reports (yet)
- **Solution: XDP code stores counters that report a change in the traffic behaviour:**
  - A queue that increased/decreased more than 20KBytes
  - A flow path that changed
  - A hop delay >2 microseconds
  - A total delay > 50 microseconds
  - An egress interface that is using more than 80Gbps for more than 50ms
- This data is stored in a time series db
- More granular metrics => more CPU usage
- Results:
  - Pros: Close to real time processing
  - Cons: Not so granular measurements.




# Challenge 4: Event-driven monitoring

- INT is a streaming telemetry solution:
  - Network devices “proactively” trigger notifications when events happen

- What if we don't receive a notification?

- Was it because there is no event?
- Was it because there is no traffic?
- Was it because the monitoring system is down?



- Recording events within the threshold limits is necessary

- *Every X ms, record the current state*
- *“Easy” to be done at the collector / hard to be done at the INT node (future)*

- *AmLight Telemetry Collector works as a collector (passive) but also as a requester (active)*

## Challenge 5: Storing all telemetry reports for future research (ML/AI)

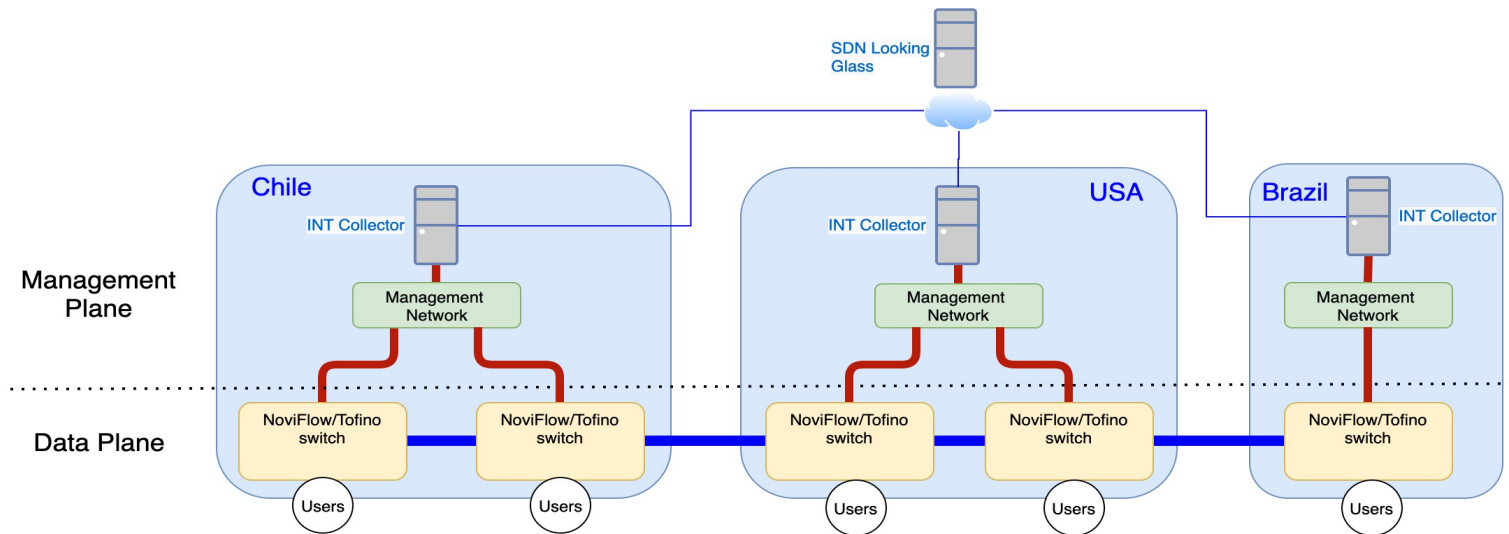
- Goal: *Store as many telemetry reports as possible for future research to enable ML/AI researchers to have grounding truth for learning algorithms*
- Each Vera Rubin Telescope (LSST) 5-second 13.6GB data transfer will generate ~337MB of telemetry data.
  - 1,334 observations/night: **450GB of telemetry data/night**
- Challenges:
  - How to save Gbps of telemetry reports without increasing OPEX (rack space, power consumption, etc.)
  - How/Where/How long to store such data?
  - How to make it available preserving privacy but without compromising research?
  - What data is really necessary from the telemetry report?
  - What has to be combined with reports to give context? Topology?
- Challenge 5 is wide-open.



Deployment

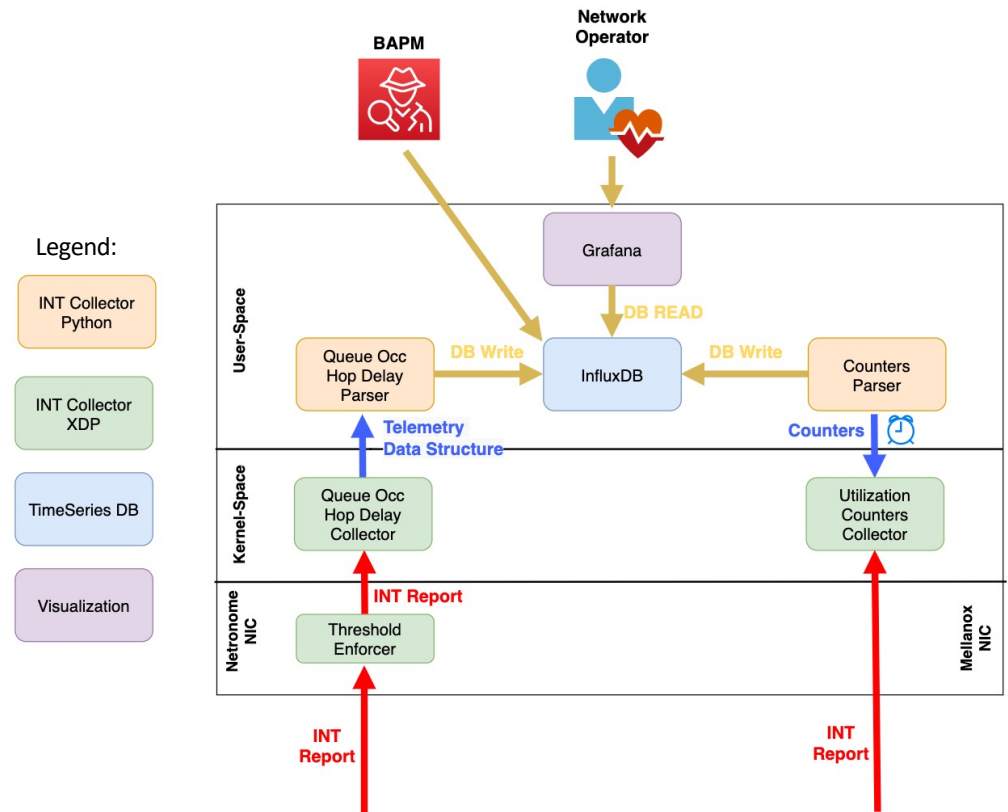
# INT Deployment at AmLight [1]

- At each AmLight site, P4 switches are replacing the current data plane
- Each pop has a Telemetry Collector parsing telemetry generated locally



# INT Deployment at AmLight [2] - AmLight Telemetry Collector

- Netronome 40G programmable NIC
- Mellanox MLX5 NIC
- A threshold enforcer application running on the Netronome card
- eBPF/XDP applications running before the Linux networking stack
- Python applications running at user space interfacing users and database
- Influxdb storing time-series data
- Grafana displaying results



Navigating through the solution:

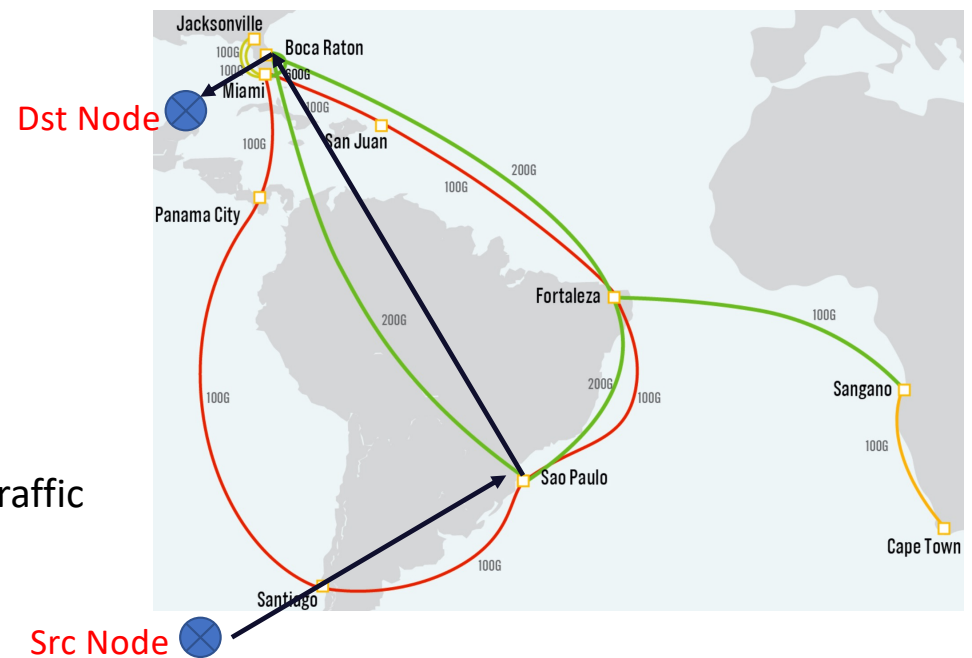
- Live: <https://int-collector.amlight.net>

## SC21 Demonstration

# Demo Setup

- The goal was to highlight the INT potential by showing INT in production.

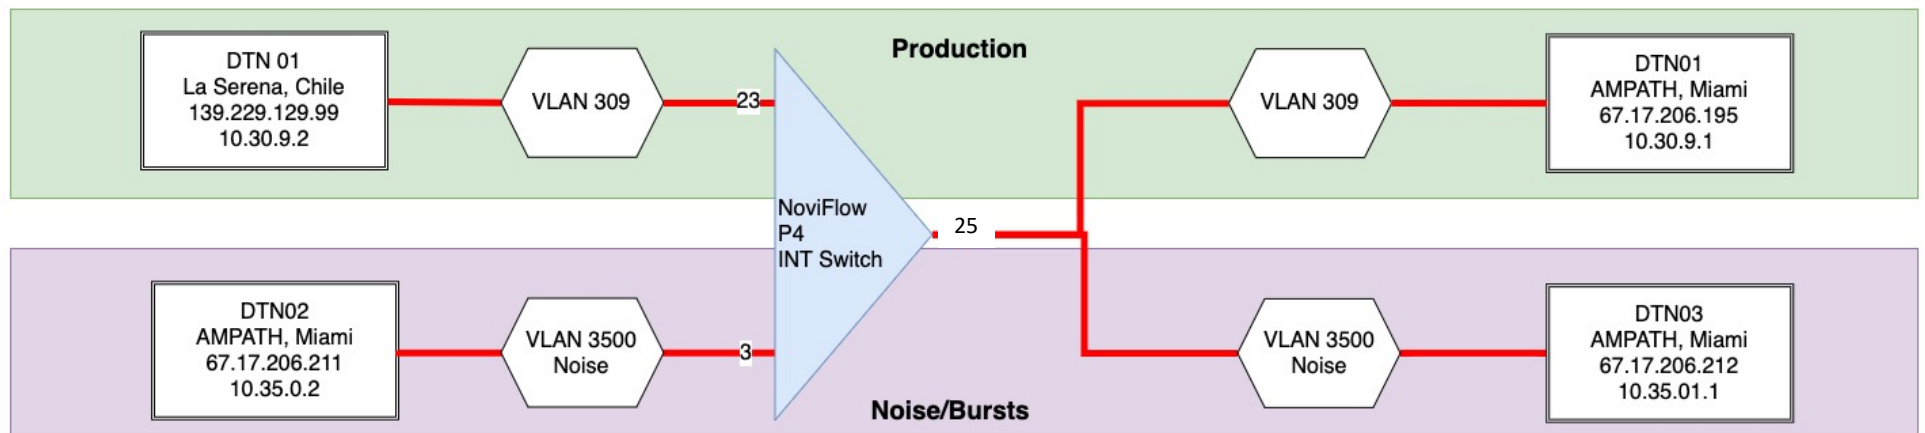
- Science traffic:
  - Source node in La Serena, Chile (@Vera Rubin DC)
  - Destination node in Miami (@AMPATH)
  - VLAN 309
- Noise traffic:
  - Congestion created at AMPATH
  - Sharing the same egress interface with the *science* traffic
  - VLAN 3500





# Demo Setup

- Science traffic comes from interface 23 and has egress interface 25
- Noise traffic comes from interface 3 and has egress interface 25 (same)
- Interface 25 shares buffer for both incoming flows



## Demo 1: Just the science traffic...

- Source in La Serena sending as much TCP data as we can generate
- Bursts and continuous traffic
- We will monitor bandwidth utilization, hop delay, buffering, and retransmissions (iperf3)

## Demo 2: The science traffic being impacted by the congestion traffic.

- Source in La Serena sending as much TCP data as we can generate
- Noise being generated in Miami, using UDP
- Bursts and continuous traffic
- We will monitor bandwidth utilization, hop delay, buffering, and retransmissions (iperf3)

### Snapshot:

<https://snapshot.raintank.io/dashboard/snapshot/WRGyNGRUbxflxN6wZEXy6295YK2tQw5r?orgId=2&from=1637001022363&to=1637006799938>

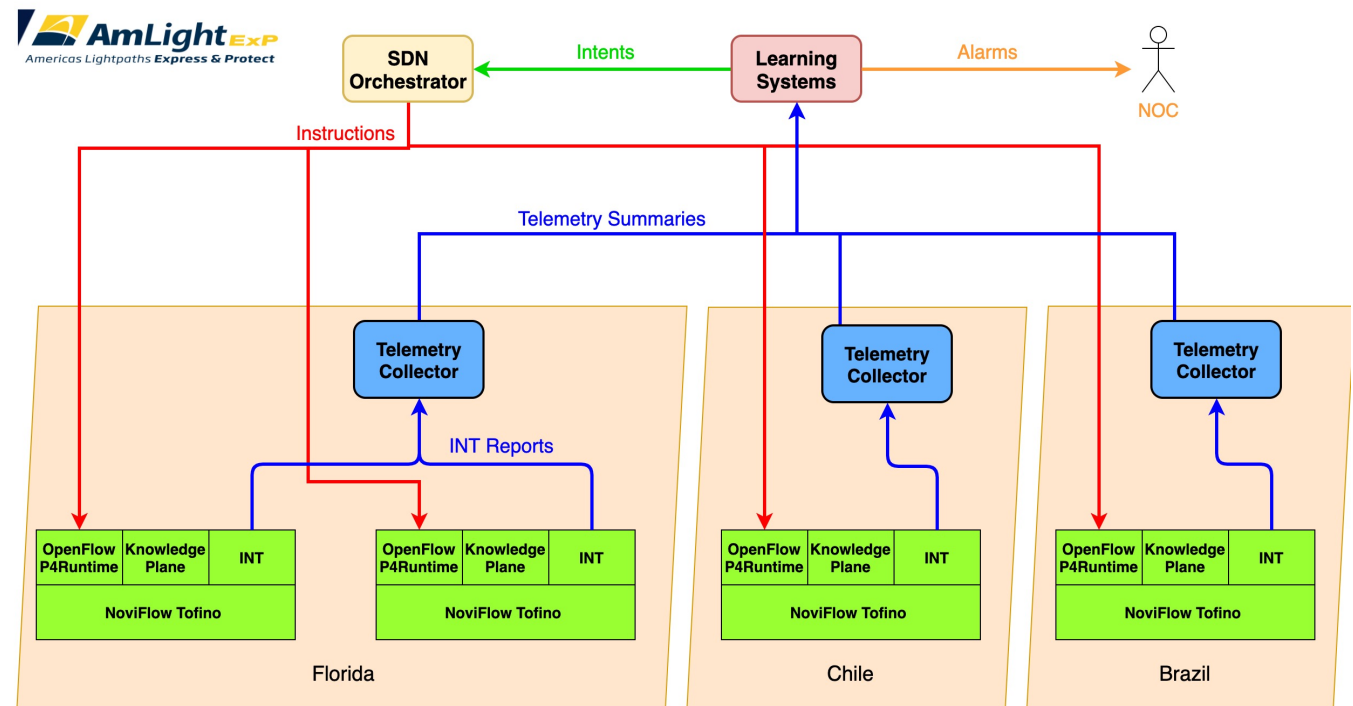
Future

# AmLight-ExP

INT-related objective: Closed-loop network orchestration by leveraging telemetry reports from the packet and optical layers, combined with Machine Learning algorithms

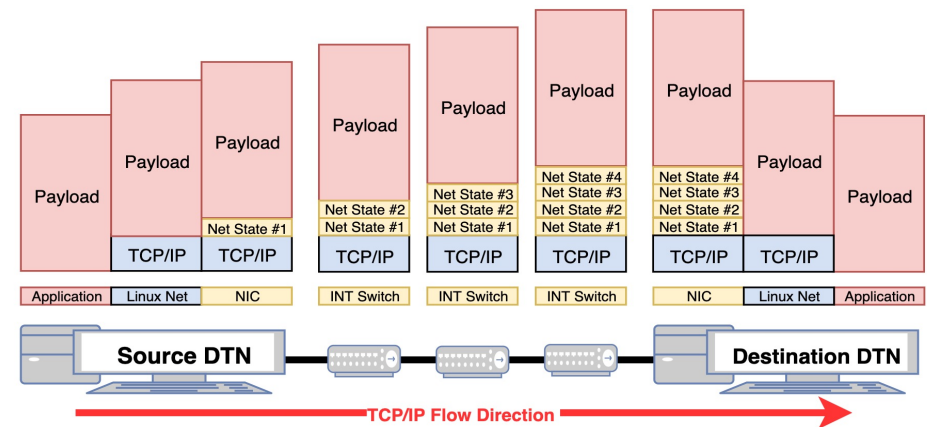
Roadmap: Self-Optimizing the network:

- Year 2: < 5 seconds
- Year 3: < 2 seconds
- Year 4: < 1 second
- Year 5: < 500 ms



# CC\* Integration: Q-Factor

- Collaboration between FIU and ESnet
- Objective: Improve data transfers over long-haul high-bandwidth programmable networks
- How: Creating an end-to-end framework where endpoints would have network state information to dynamically tune data transfer parameters in real time
  - Bandwidth and resources optimization
- Transformative:
  - Q-Factor will enable endpoints to adapt their data transfers jitter/delays, and excessive memory consumption.
- Summary of proposed activities:
  - Expanding the Management Plane to endpoints
  - Developing a Telemetry Agent to consume network state information and tune endpoints
  - Evaluating tuning at scale over multiple scenarios by leveraging AmLight and Esnet networks and testbeds







**Thank You! Questions?**

AmLight SDN Eng. Team <sdn@amlight.net>

SC21: Demo: In-band Network Telemetry @ AmLight