

**SAACC Meeting**  
**April 22<sup>nd</sup>, 2020**



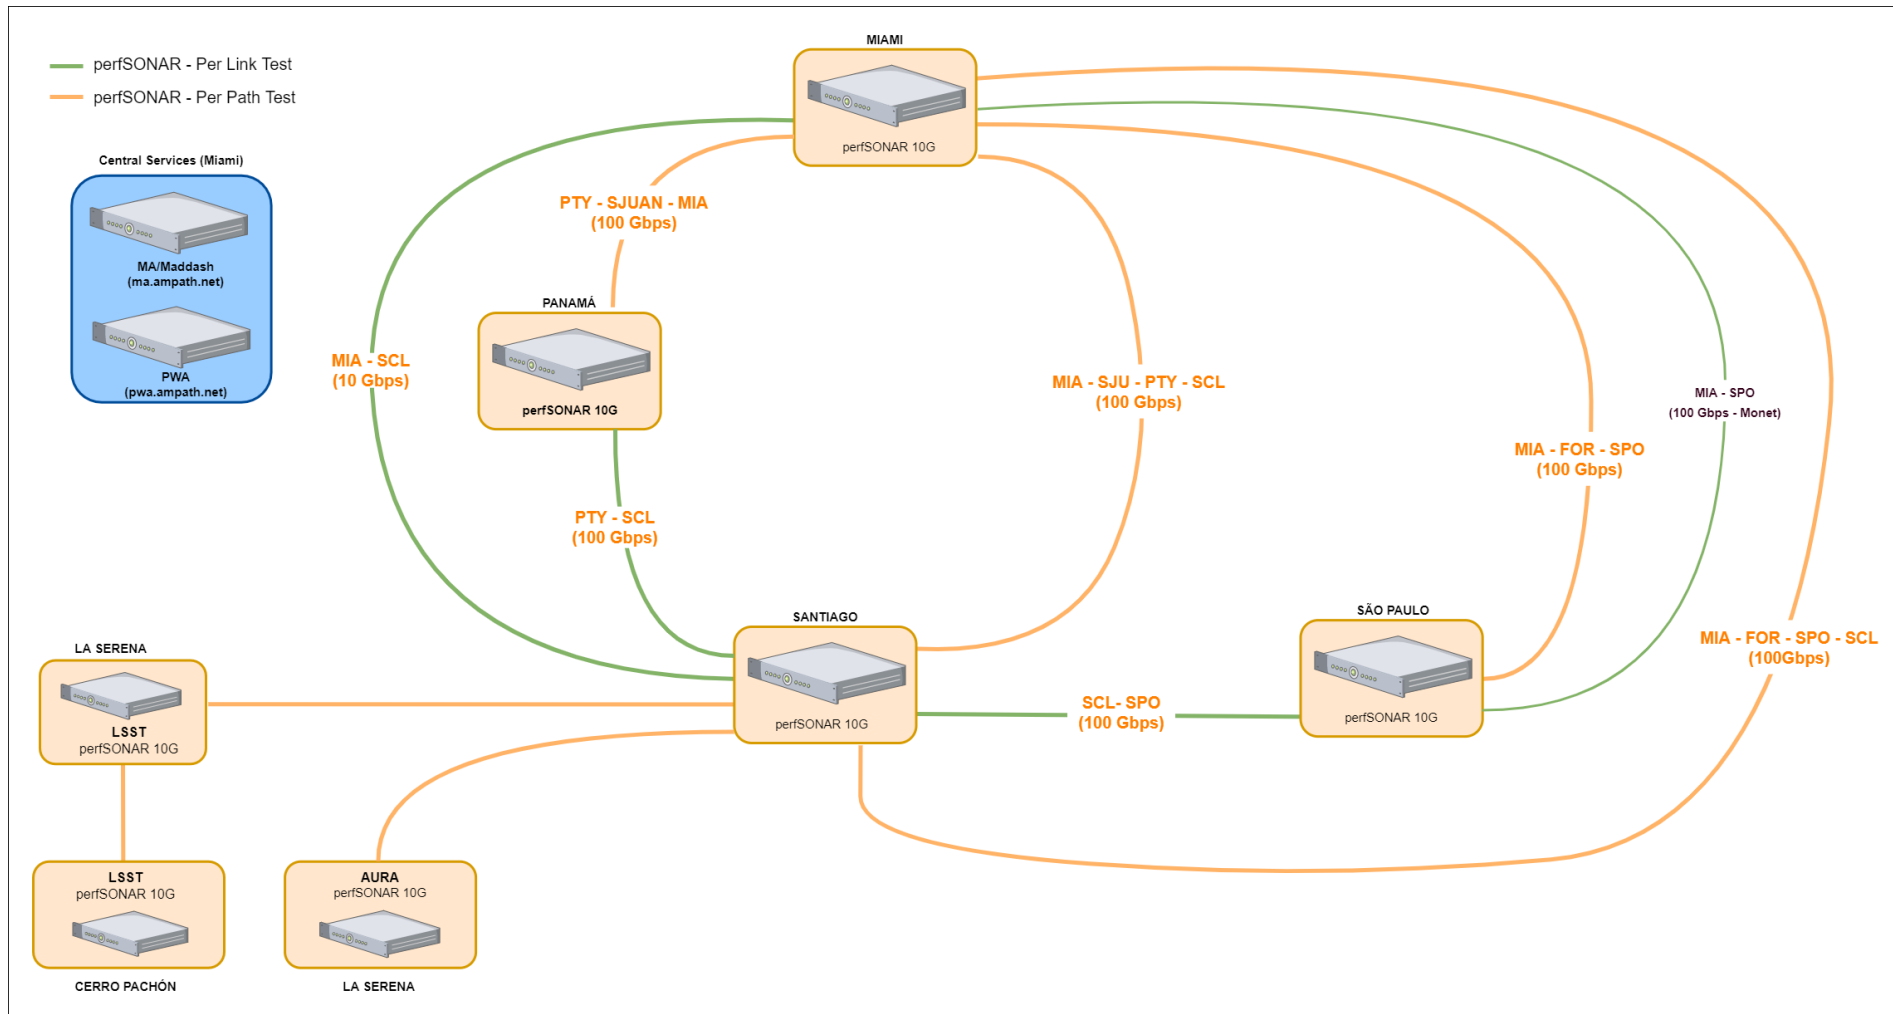
# AmLight-Exp: Performance Measurement @ AmLight

Renata Frez – Senior Network Engineer - RNP/AmLight

# Measurement Tool: perfSONAR

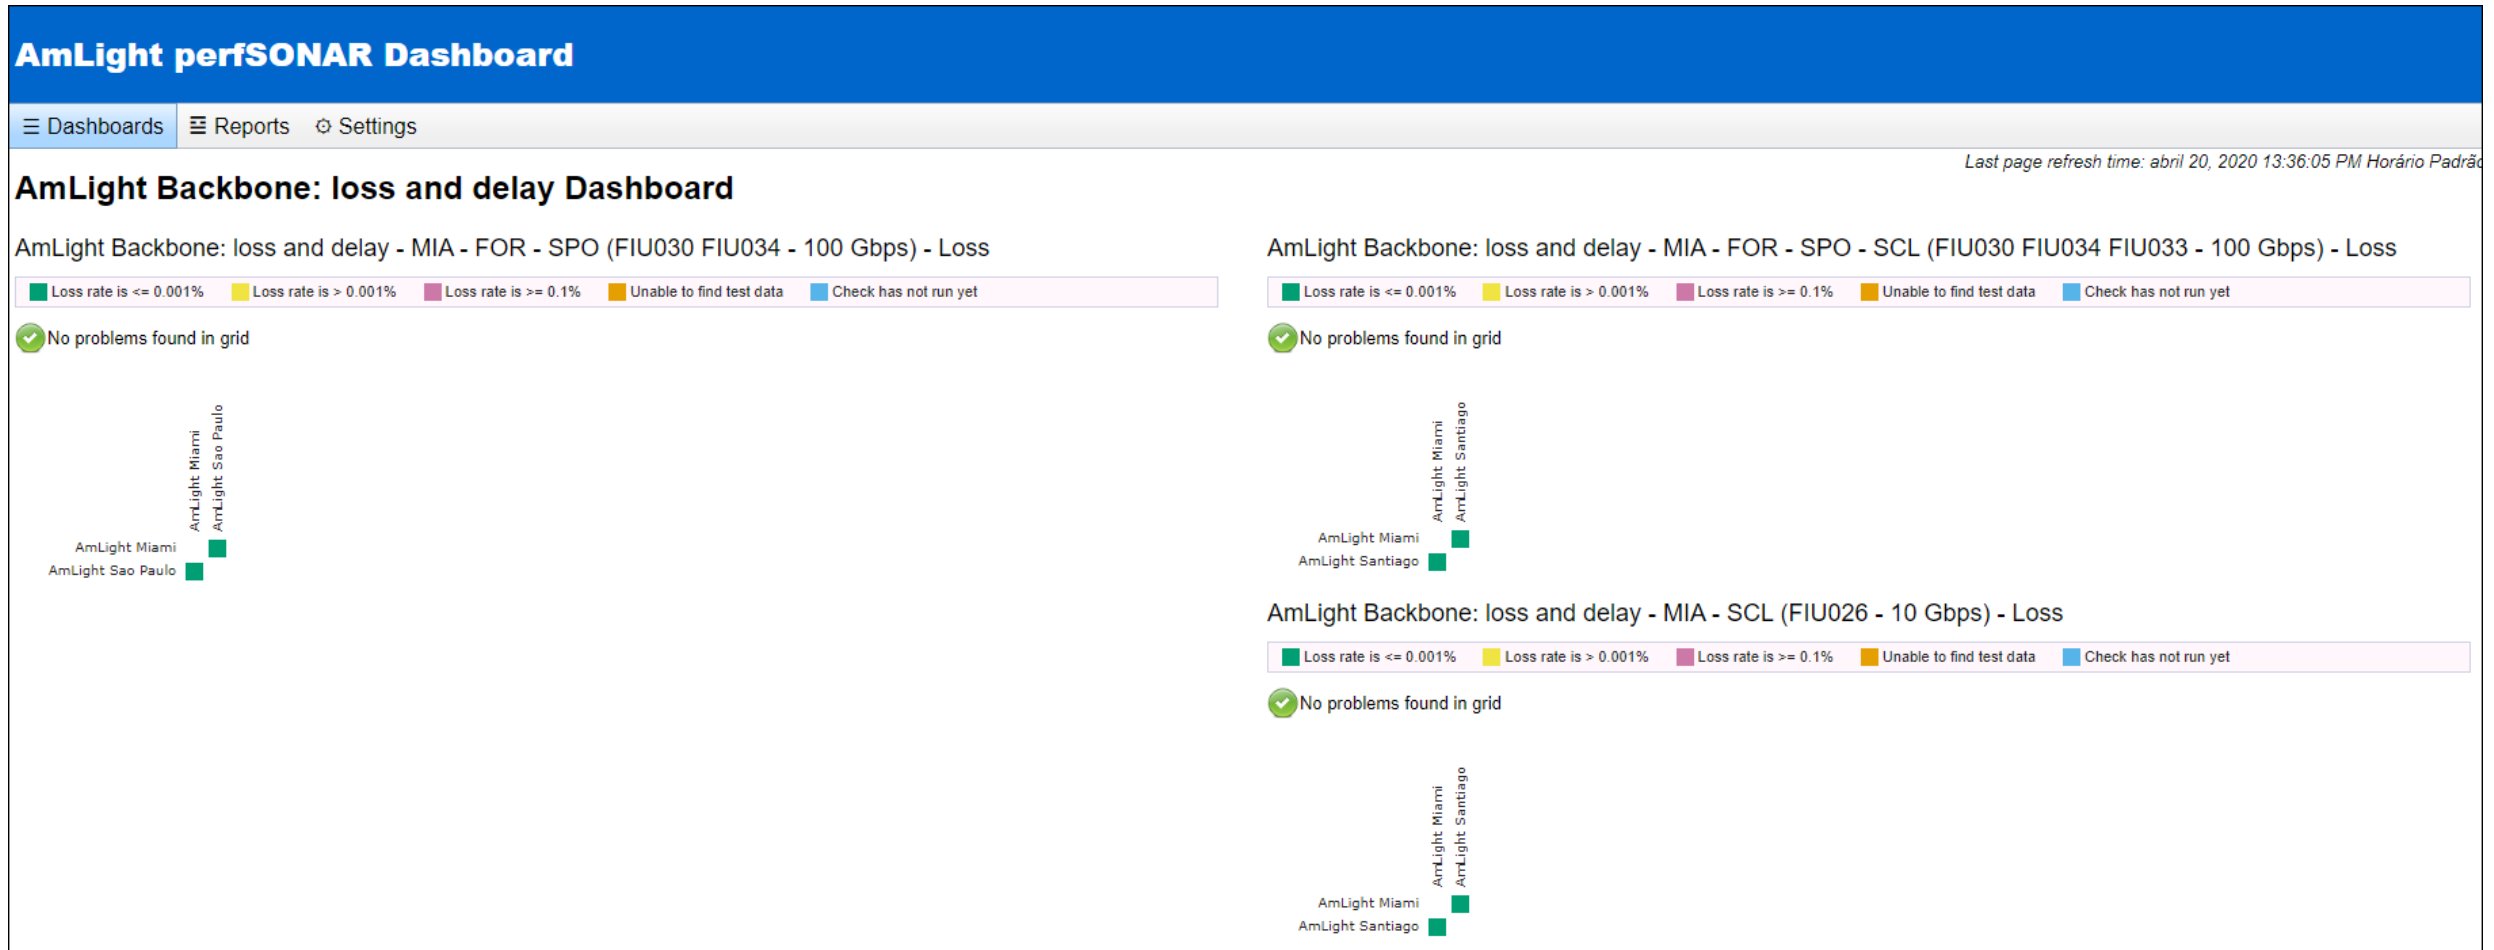
- Why use a network performance tool?
  - Save historical data
  - Availability of a large number of tests
  - Help to identify potential improvement points
- Why perfSONAR?
  - Well-known toolkit commonly used in the academic community
  - It is being developed to address the R&E network's needs
  - It is an open source tool
  - Customizable tests supported
  - More than 2000 measurement points deployed
  - Web interfaces could be used for management and data visualization
- What is AmLight testing?
  - Network throughput – Type of traffic: TCP. Interval: 4 hours. Parallel streams: 8
  - Network packet loss – Type of traffic: OWAMP. Interval: Run continuously
  - Network delay – Type of traffic: OWAMP. Runs continuously.

# AmLight – perfSONAR Measurement Topology (April 2020)



# Maddash – The perfSONAR Dashboard (1)

- Presents the data as a 2D grid enabling to visualize quickly a test with poor performance. Accessible using: <https://dashboard.ampath.net>



# Maddash – The perfSONAR Dashboard (2)

- Presents the detailed data for a chosen test.





# Live Demonstration

# Use Cases:

- Replacement of a Brocade for a Dell switch in Santiago:
  - With the perfSonar tests we were able to verify the behavior of the international links before and after the replacement.
- Integrations made by AmLight:
  - AmLight has integrated the perfSONAR environment with AmLight NMS (Zabbix).
  - Zabbix can raise alarms as soon as a perfSONAR test reports a poor performance.
  - These alarms provide a way to quickly be advised when the network performance deteriorated.

Last 50 issues						
Host	Issue	Last change	Age	Info	Ack	Actions
ps-ma-maddash	Problem: RubinObs-LS: Santiago to La Serena - Latency: Santiago - La Serena - Loss (10.7.27.2->10.7.27.1)	<a href="#">22 Apr 2020 06:09:03</a>	6h 6m 36s		No	
ps-ma-maddash	Problem: AURA: Santiago to La Serena - Latency: Santiago - La Serena - Loss (10.7.23.2->10.7.23.1)	<a href="#">22 Apr 2020 06:01:21</a>	6h 14m 18s		No	
ps-ma-maddash	Problem: RubinObs-LS: Santiago to La Serena - RTT: Santiago - La Serena - Ping Loss (10.7.27.1->10.7.27.2)	<a href="#">22 Apr 2020 05:39:06</a>	6h 36m 33s		No	
ps-ma-maddash	Problem: RubinObs-LS: Santiago to La Serena - Latency: Santiago - La Serena - Loss (10.7.27.1->10.7.27.2)	<a href="#">22 Apr 2020 05:39:04</a>	6h 36m 35s		No	
ps-ma-maddash	Problem: AURA: Santiago to La Serena - Latency: Santiago - La Serena - Loss (10.7.23.1->10.7.23.2)	<a href="#">22 Apr 2020 05:31:22</a>	6h 44m 17s		No	
ps-ma-maddash	Problem: RubinObs-LS: Santiago to La Serena - RTT: Santiago - La Serena - Ping Loss (10.7.27.2->10.7.27.1)	<a href="#">22 Apr 2020 05:09:06</a>	7h 6m 33s		No	
ps-ma-maddash	Problem: AURA: Santiago to La Serena - RTT: Santiago - La Serena - Ping Loss (10.7.23.1->10.7.23.2)	<a href="#">22 Apr 2020 05:01:25</a>	7h 14m 15s		No	
ps-ma-maddash	Problem: AURA: Santiago to La Serena - RTT: Santiago - La Serena - Ping Loss (10.7.23.2->10.7.23.1)	<a href="#">22 Apr 2020 05:01:23</a>	7h 14m 17s		No	

# Expanding the Performance Measurement: Rubin Observatory Use Case

- Ideal end-to-end South-North BW: 40+Gbps
  - 5 seconds data transfer window every 27 seconds/every night
  - 10,703 miles/17,125km from La Serena to Champaign, IL
    - AmLight is responsible for 9,551 miles/15,281km of it
  - There are 22+ paths between Santiago and Atlanta
  - 15+ data centers in the path; 20+ cross-connects in the shortest path
  - Vera Rubin LHN's SLA:
    - Acceptable packet loss: 0.0001%
    - Mean Time Between Failures: 180 days
    - Mean Time To Repair: 48h
- ➔ *We need better troubleshooting tools!*



# Expanding AmLight-ExP: AmLight-INT


- In-band Network Telemetry is a framework designed to allow for the collection and reporting of network state, by the data plane
- Network Telemetry overcomes limitations imposed by legacy technology:
  - More metrics and granularity beyond what traditional networking monitoring solutions can provide
  - Sub-second data gathering
  - Useful for microburst detection and queue utilization at a sub-second interval
  - Complete view of network state in the flow's path

# How does AmLight plan to use INT to support Rubin Obs.?

- *How did this packet get here?*
  - The sequence of network devices a packet visited along its path.
- *Why is this packet here?*
  - The set of rules a packet matched upon at every switch along the way.
- *How long was this packet delayed?*
  - The time a packet spent buffered in every switch, to the nanosecond, from end-to-end.
- *Why was this packet delayed?*
  - The flows and applications that a packet shared each queue with.

# AmLight-INT: Framework and Tools

- **INT Controller:**
  - Collects all telemetry reports
  - Saves telemetry reports after a threshold
- **QueueTop**
  - Monitors all network devices' queues
  - Pinpoint where the bottlenecks are
- **Proof Of Transit**
  - Records the path taken per packet
- Tools will be released as open source on Github soon



# Final Considerations

## Future thoughts:

- Add perfSONAR nodes to the San Juan and Fortaleza AmLight sites.
- Add TENET's perfSONAR to AmLight Maddash dashboard.
- Create more end-to-end tests to amplify our connectors' perspective.
- Study if other types of tests available with the perfSONAR toolkit could improve our daily work.



**Thank You! / Questions? / Comments?**



**AmLight-Exp: Performance Measurement @ AmLight**

**Renata Frez** - <renata@amlight.net>