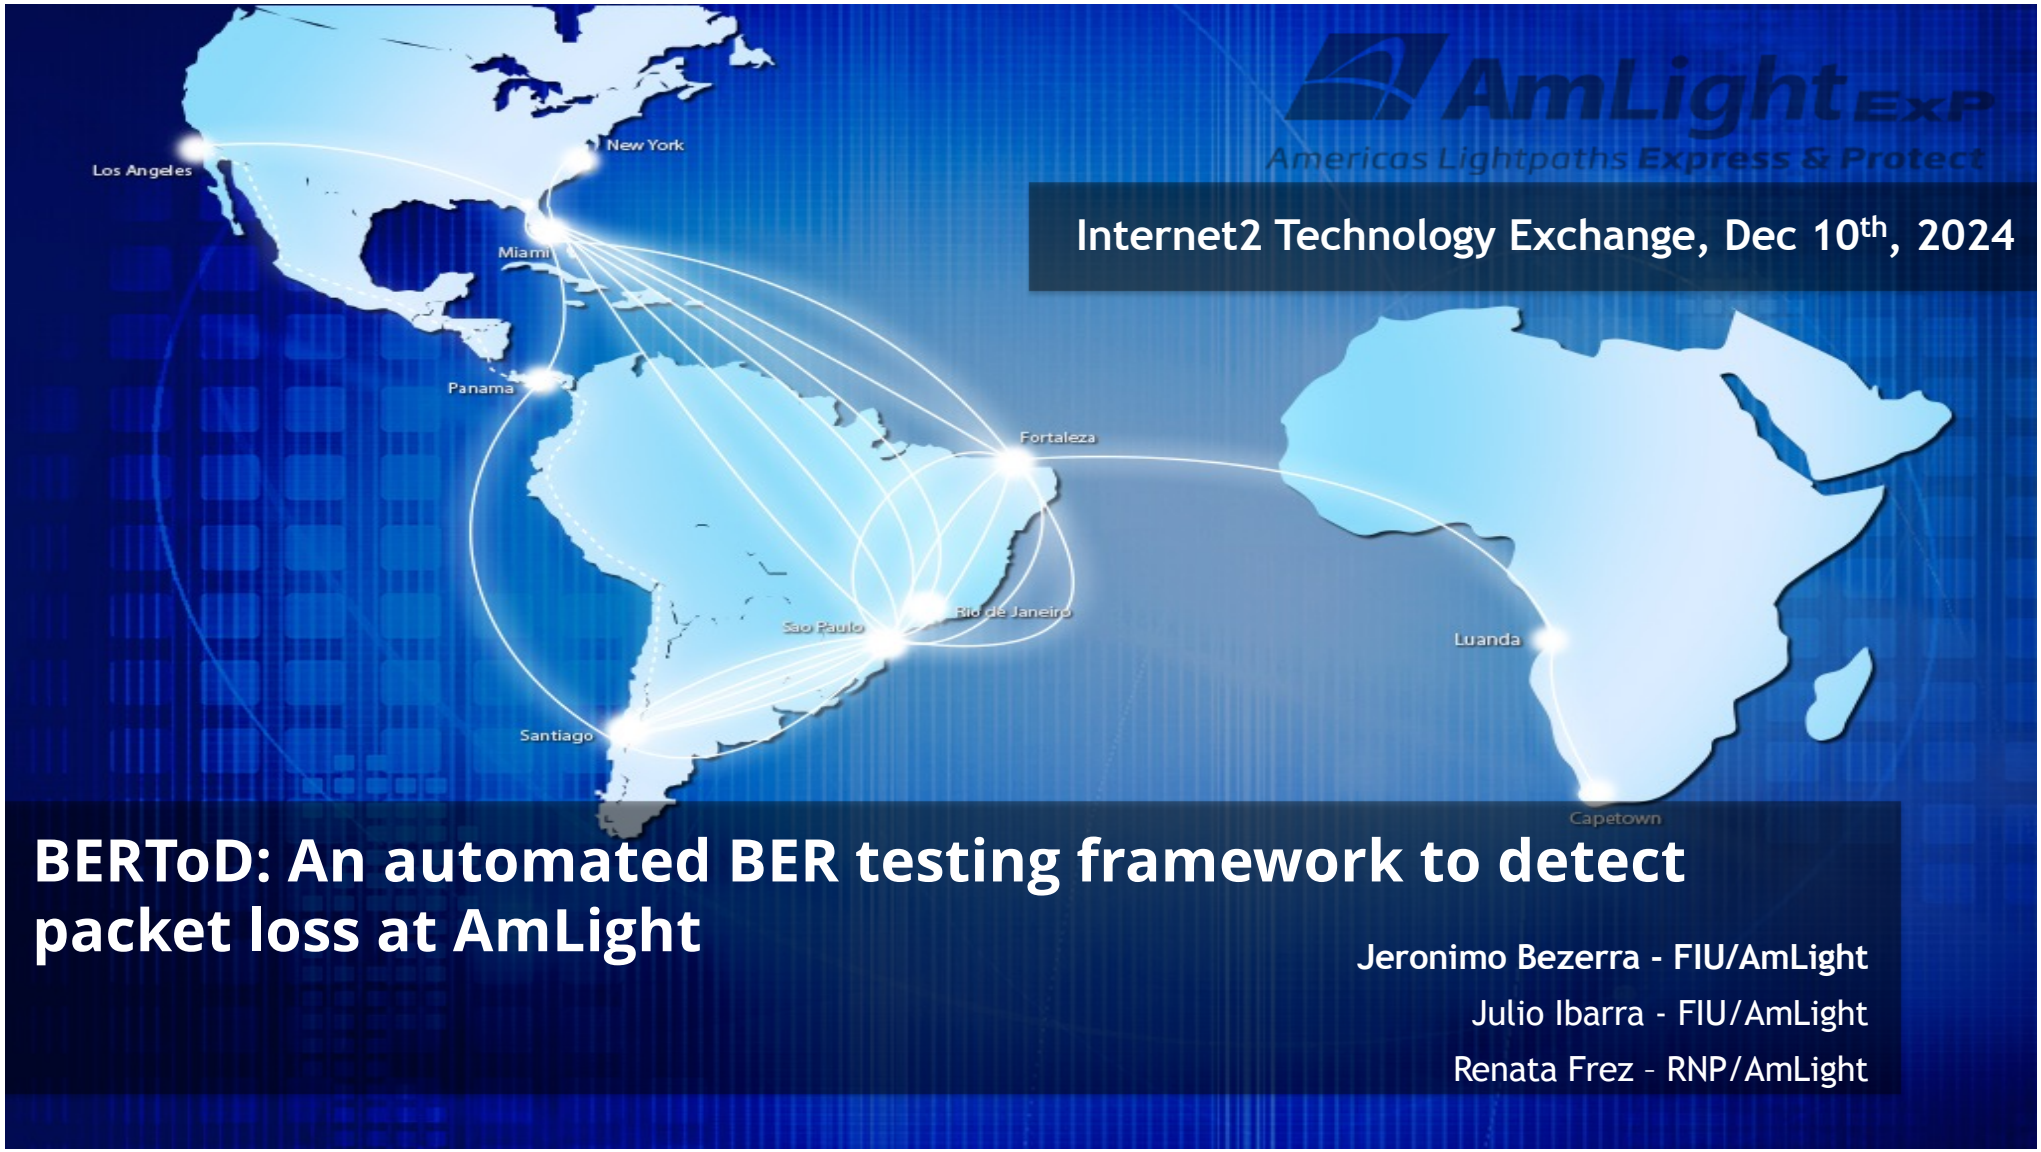

BERToD: An automated BER testing framework to detect packet loss at AmLight

Jeronimo Bezerra - FIU/AmLight
Julio Ibarra - FIU/AmLight
Renata Frez – RNP/AmLight

Outline

- Motivation
- Current approaches to detect/isolate packet loss
- How BERToD works
- Lessons Learned
- Future Work
- Conclusion

Disclaimer

- Packet vs Frame
  - Interchangeable in this presentation
- Packet Loss vs Drop:
  - Drop: We drop our packets
    - QoS, blocking topologies, traffic engineering, small buffers
  - Loss: Someone/thing loses/corrupts our packets
    - Fiber cuts, power outages, damage components
- For this talk, our focus is on packet loss!

Motivation

- Applications running over long-haul links suffer more with packet loss
  - 105 ms RTT is standard at AmLight
- AmLight has grown in complexity in the last 5 years
  - Next Slide
- Science applications are expecting better network performance
  - SLA-driven science drivers are demanding more granular measurement ($1x10^{-9}$)
- Current solutions for loss detection have fundamental, accuracy, or granularity limitations

NSF IRNC: AmLight Network – 2020-2025

- 39x 100G links:
  - 2.1+ Tbps of international connectivity
  - AmLight will reach 4.9 Tbps of total capacity[2025]
  - Dark fiber, spectrum, waves, and lit services
- 9x Sites / 19x racks:
  - Miami, Boca Raton, Jacksonville, Sao Paulo, Fortaleza, Santiago, San Juan, Panama City, Cape Town, Atlanta, and Buenos Aires
- Network and Monitoring Devices:
  - 25x programmable switches and Juniper routers
  - 10x 10G perfSonar nodes
  - 4x In-band Network Telemetry (INT) collectors
    - ~10Mpps & 96TB of telemetry data per day

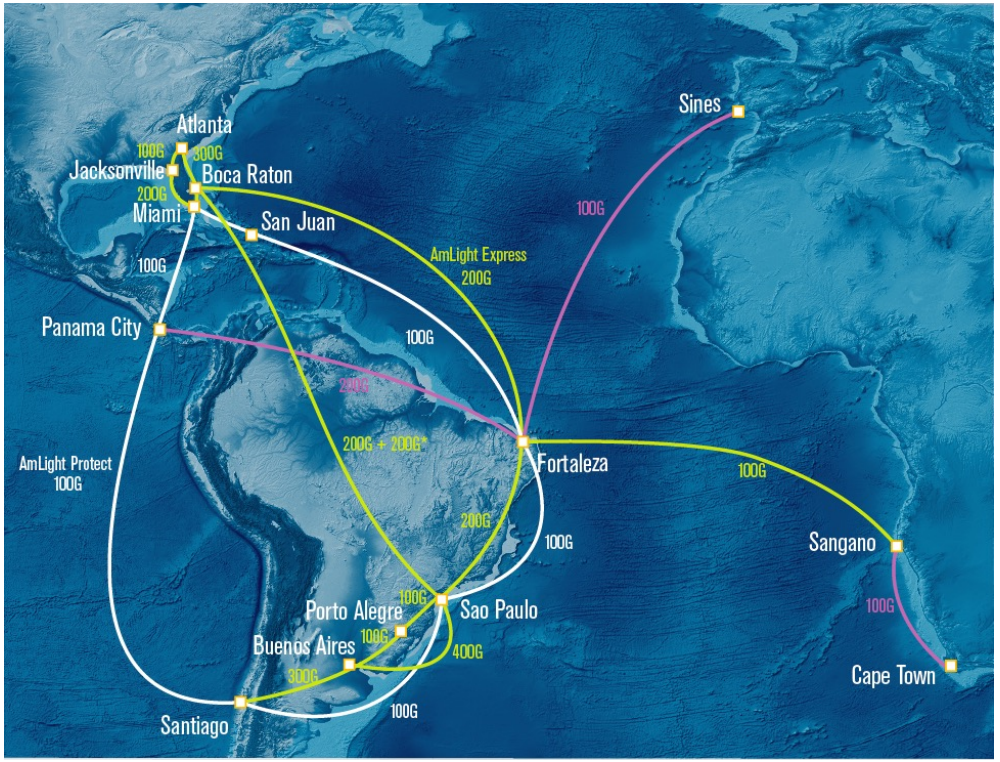

AmLight ExP
Americas Lightpaths Express & Protect

Current approaches to detect/isolate packet loss

- AmLight monitoring:
  - SDN counters, ICMP, SNMP, traceroute, optical telemetry metrics (polling-based), INT and JTI reports (streaming-based), and perfSONAR measurements and spare DTNs (software-based)

- Even with all of those, fault isolation and mitigation are still challenging with a high OPEX:
  - Evaluate results, correlate data, run extra tests, send field technicians to clean/replace suspicious components, steer traffic, and run again with different outcomes ➔ Days of work

- Existing packet generators/testers are used in an ad-hoc fashion
  - Hardware-based granularity,
  - But manual configuration and with a learning curve to configure and read results

BERToD - Bit Error Rate Test on Demand

- An automated packet loss detection framework that uses granular network telemetry (INT), SDN, and hardware-based packet generators to detect bit error rates as low as $1x10^{-9}$
- Leverages recent developments at AmLight:
  - Physical and logical loops in the SDN switches
  - Link utilization from In-band Network Telemetry (INT)
  - Topology data and service details from Kytos-ng SDN Controller
- *Near* deterministic results due to specialized network hardware being used end-to-end
  - Highly accurate with granular results

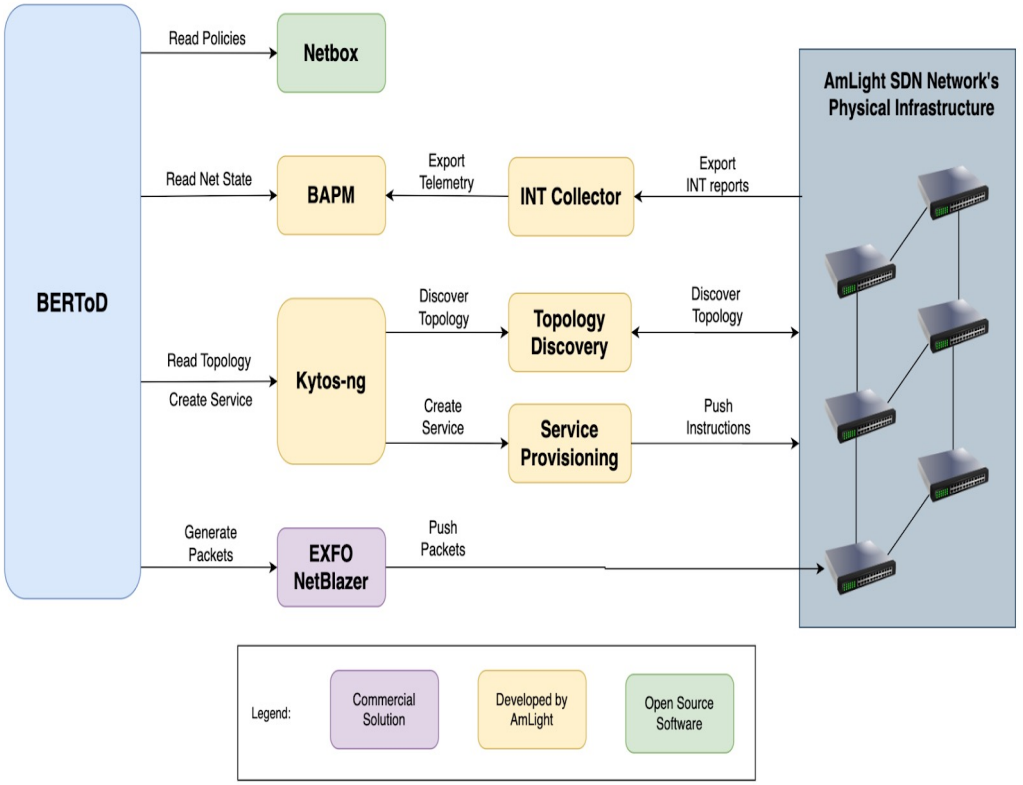

AmLight ExP
Americas Lightpaths Express & Protect

BERToD - Bit Error Rate Test on Demand [2]

- Test every possible link every 30 min:
  - Latency, jitter, frame loss, and out-of-sequence tests
  - Multiple frame sizes: 68, 256, 512, 1024, 1518, 9000 bytes
  - Each test runs for up to 10 seconds, and we send up to 1,500,000 frames
    - In case a test fails, run again with a multiplier metric (for instance, 3)
  - Choice for max bandwidth comes from the In-band Network Telemetry (INT) data
    - 50% of the available bandwidth based on the last 30 seconds (and up to 40 Gbps)

- Displaying results:
  - Last 6 hours, Last 7 days, Heatmap, and command-line outputs

- Annotations are used to document known topology events and actions to help correlate events

BERToD – Granular Individual Results

- Using Grafana to plot each test's loss, jitter, latency, and out-of-sequence
- Great way to understand the last 24 hours
- Filters available to visualize test results based on frame size and individual paths
- Not great for correlating fault events

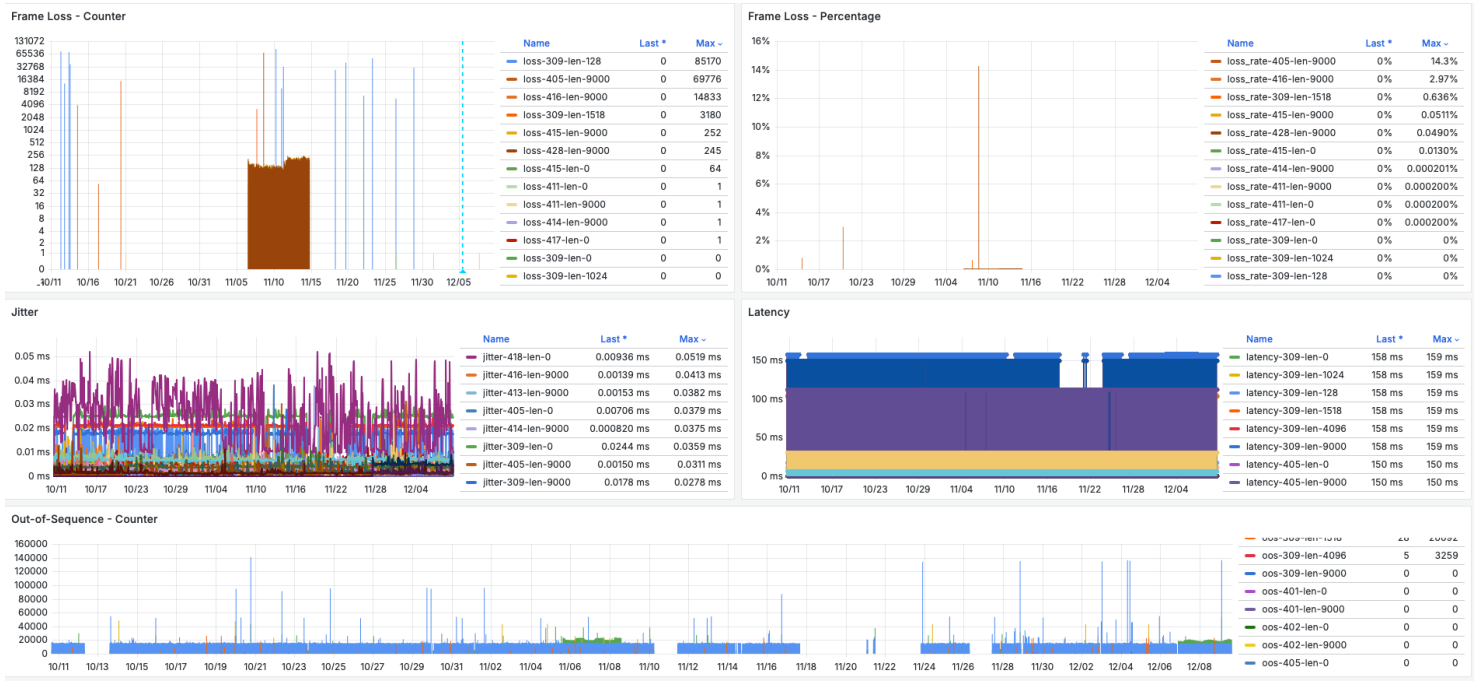

AmLight ExP
Americas Lightpaths Express & Protect

BERToD - Bit Error Rate Test on Demand [3]

- Using Grafana to plot each test's loss per day
- Great way to correlate events and identify patterns
- Filters available to visualize test results based on frame size and individual paths
- Used with annotations to add context

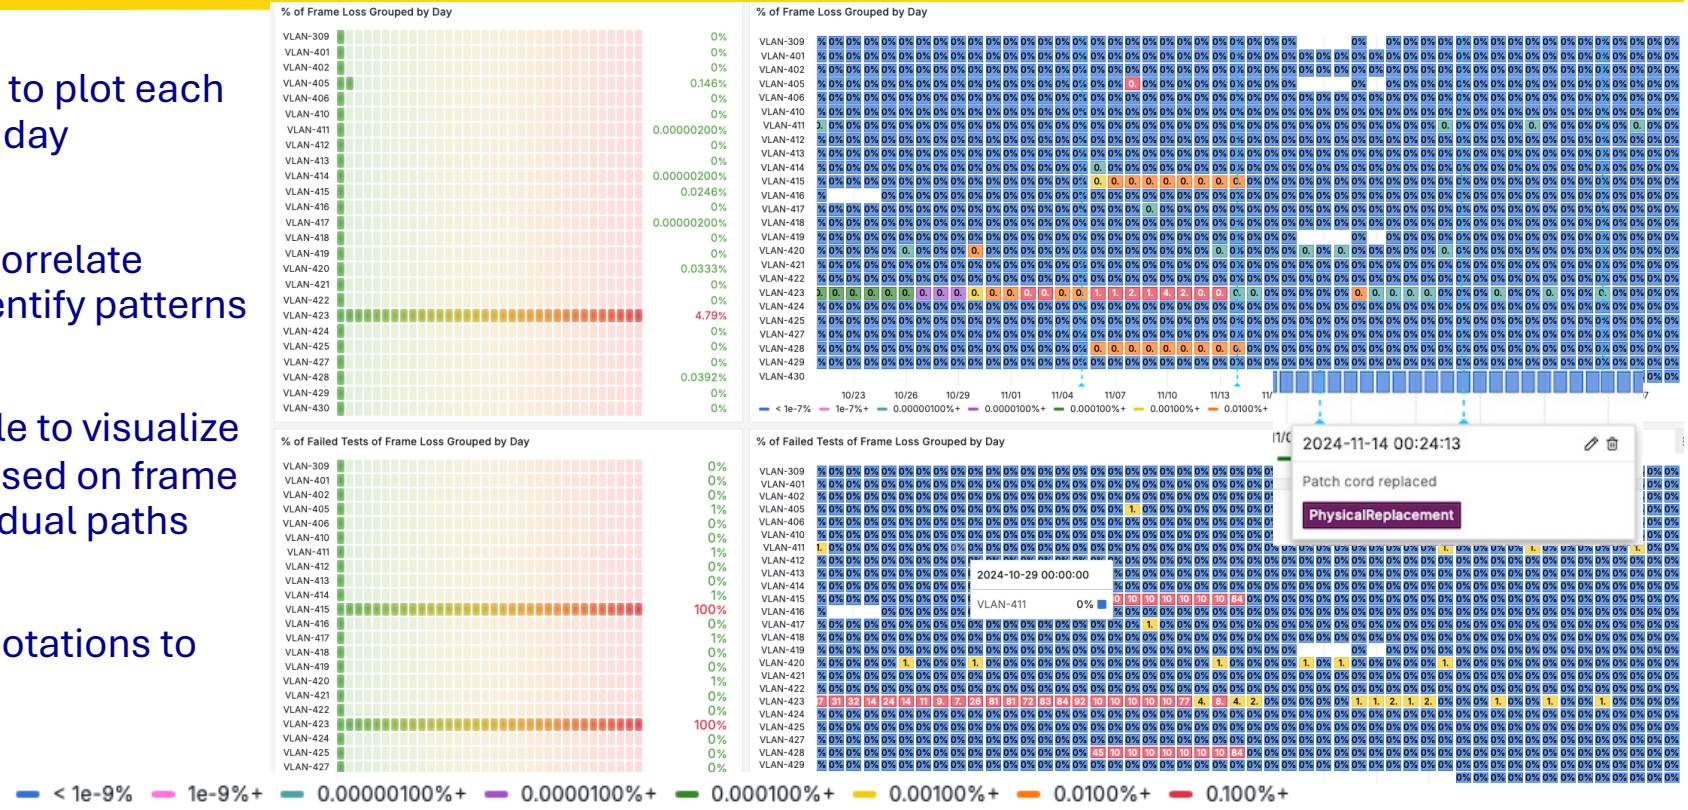

AmLight ExP
Americas Lightpaths Express & Protect

Lessons Learned

- Testing infrastructure vs testing user experience
  - To achieve deterministic results, network resources must be fully available
  - Traffic engineering at AmLight was enhanced to cover this issue ➔
- How to monitor the actual user experience?
  - Using perfSONAR and BERToD in the same queue as users
- "*There is no such innocent maintenance at the NAP*".
  - Mishandled patch cords are the main reason for sudden spikes of errors ➔
- Dirty fiber/connector is the main reasons for discreet errors (<0.0001%)
- Some vendors have weird policies, and small frames are delivered out of sequence (under investigation)

AmLight Traffic Prioritization Policy	
Queue 7	Reserved for future use.
Queue 6	Reserved for management traffic
Queue 5	Reserved for future use.
Queue 4	Reserved for "deterministic" monitoring (BERT).
Queue 3	Vera Rubin Observatory over shared links.
Queue 2	Reserved for more than best effort. Not in use.
Queue 1	(Default) Best effort traffic & BERToD for users
Queue 0	Less than Best Effort. Experiments/Microbursts

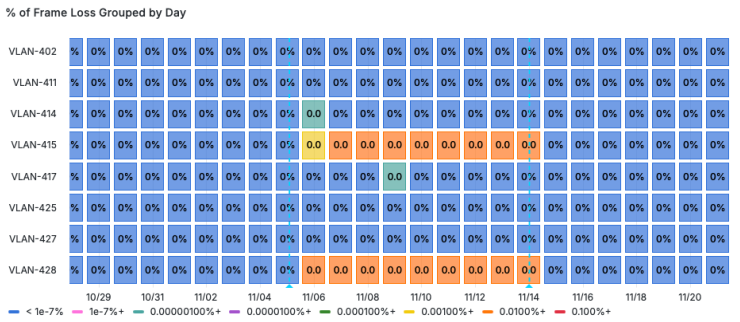

AmLight ExP
Americas Lightpaths Express & Protect

Lessons Learned – Testing Infra vs User Experience

AmLight ExP
Americas Lightpaths Express & Protect

Next steps

- Integration with the SDN Controller to test links after link flaps
  - The goal is to evaluate a link after maintenance/repair before using it again!
- Enhance the fault isolation process using all data sources available
  - SDN logs, topology changes, EVC optimizations, events/demos, and visits to the data center.
- A deep dive presentation is being scheduled for the CI Lunch and Learn!

Conclusion

- BERToD is a fantastic addition to the network monitoring portfolio thanks to the hardware-based traffic generator and enhanced network telemetry provided by the AmLight SDN solution.

- Having a hardware-based traffic generator enables quick testing with extreme accuracy
  - Helps us follow the demands of our SLA-driven science drivers

- BERToD is a great complement to perfSONAR @ AmLight.
  - While perfSONAR allows AmLight to test applications and protocols with excellent per-direction visibility, BERToD provides extreme performance visibility for applications over ultra-long paths where any packet loss causes damage.

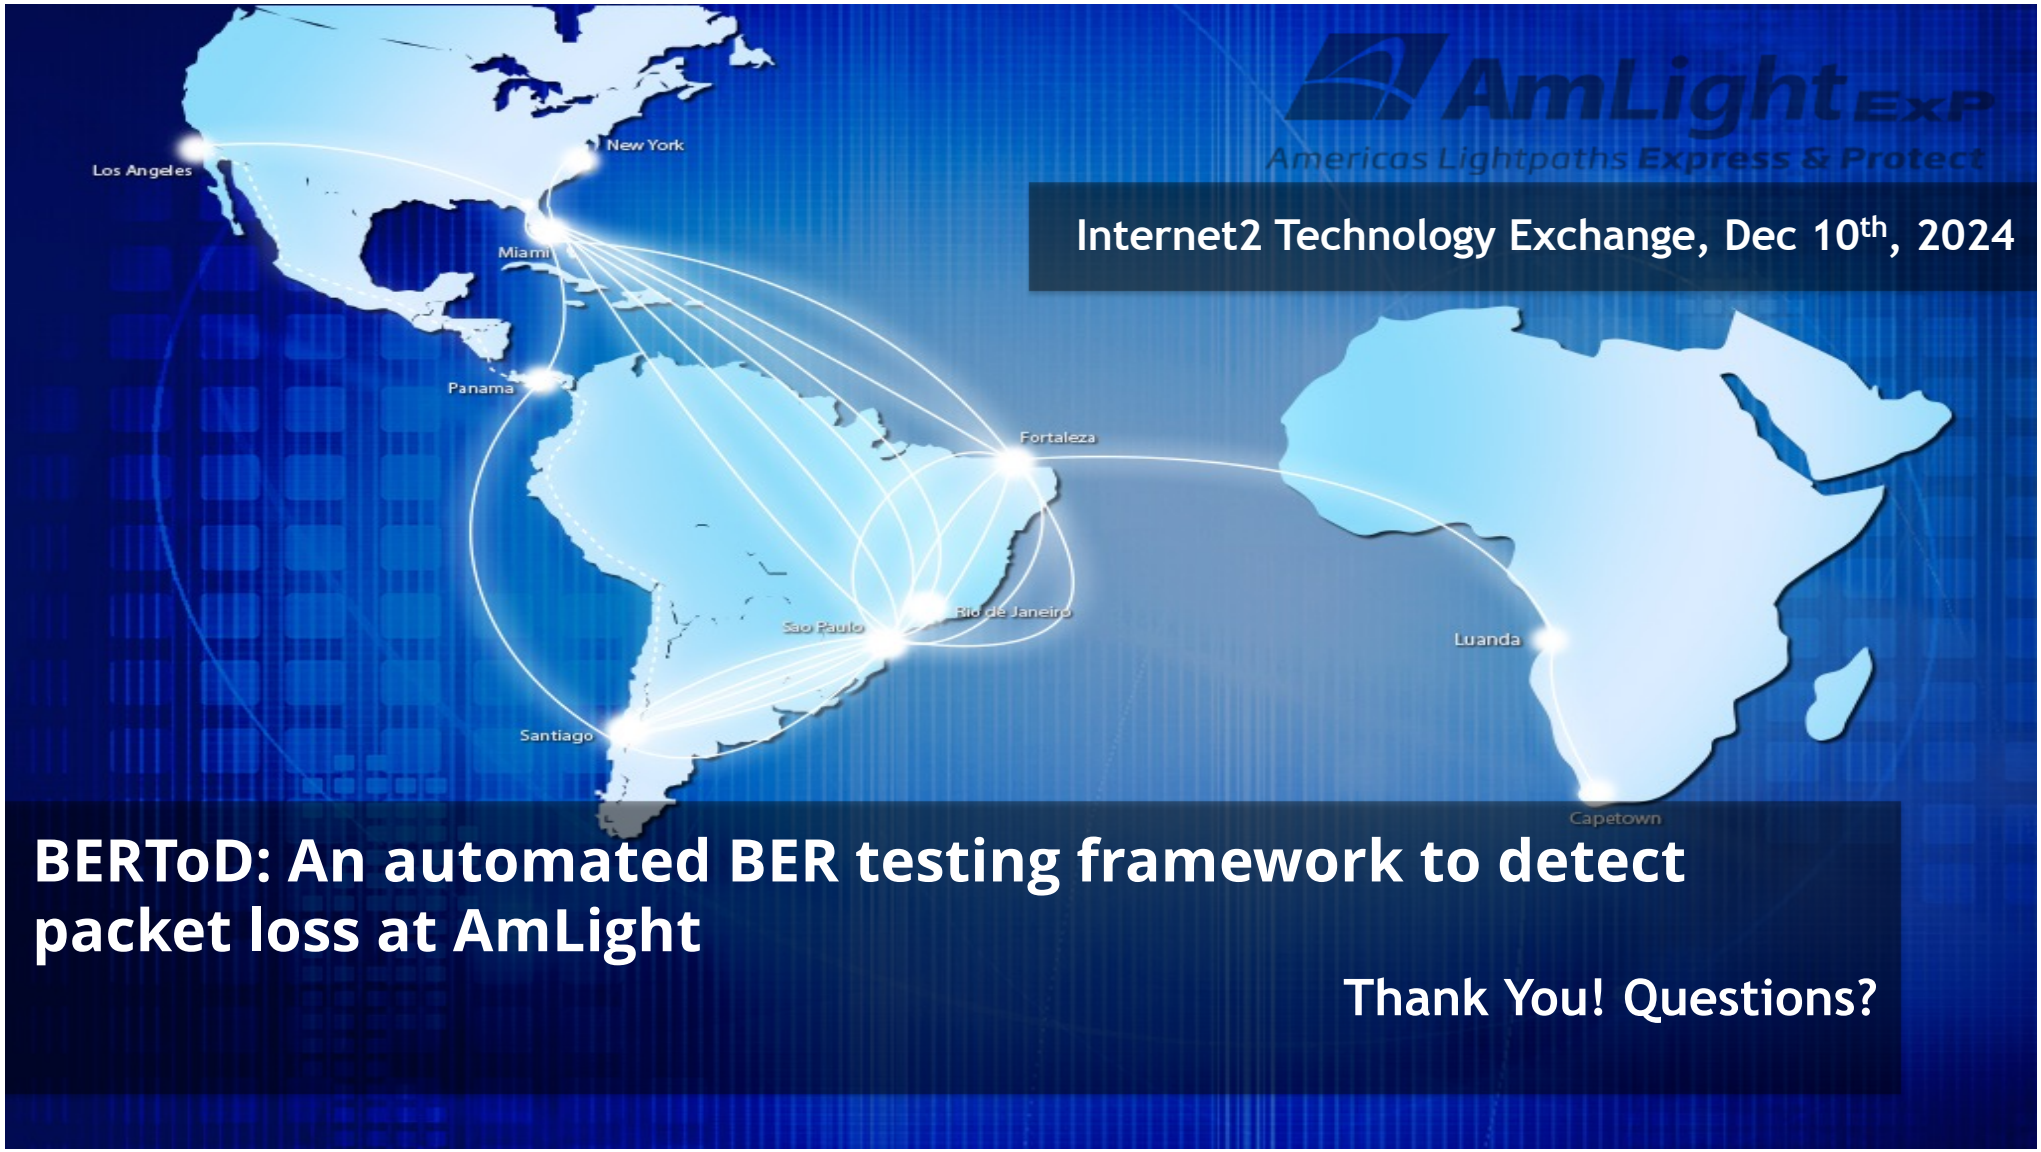
AmLight ExP
Americas Lightpaths Express & Protect
Internet2 Technology Exchange, Dec 10th, 2024
Los Angeles
New York
Miami
Panama
Fortaleza
Rio de Janeiro
Sao Paulo
Santiago
Luanda
Capetown
BERToD: An automated BER testing framework to detect packet loss at AmLight
Thank You! Questions?

BERToD - Bit Error Rate Test on Demand [5]

- Per-Hour Heatmap visualization created to help identify patterns across tests

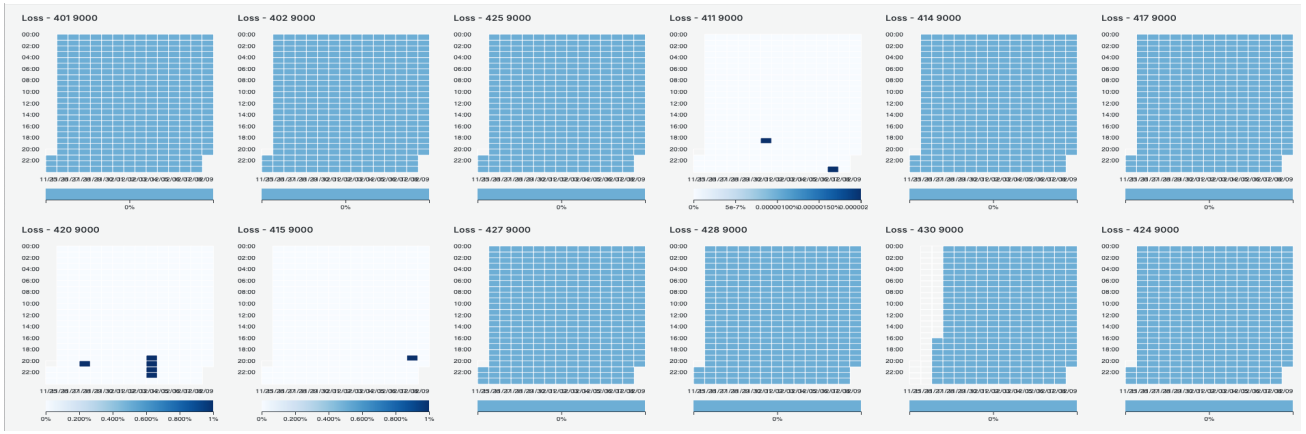

- Command-line to access full results and test configuration

2024-12-09 20:36:45	Vlan_405_BERToD_SCL-SW03	9	405	fx(9000)	10	500,000	pass	0.001/0.040	pass	150/152	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:36:58	Vlan_416_BERToD_SAO-SW03	10	416	fx(9000)	10	500,000	pass	0.001/0.040	pass	104/108	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:37:11	Vlan_410_BERToD_BCT-SW03	11	410	fx(9000)	10	491,278	pass	0.001/0.040	pass	1/1	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:37:25	Vlan_419_BERToD_SCL-SW03	12	419	fx(9000)	10	500,000	pass	0.001/0.040	pass	150/152	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:37:38	Vlan_413_BERToD_ATL-SW01	13	413	fx(9000)	10	500,000	pass	0.001/0.040	pass	12/18	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:37:52	Vlan_424_BERToD_MIA-SW14	14	424	fx(9000)	10	500,000	pass	0.001/0.040	pass	0/1	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:38:05	Vlan_418_BERToD_SAO-SW04	1	418	fx(9000)	10	500,000	pass	0.003/0.040	pass	105/108	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:38:18	Vlan_423_BERToD_SJU-SW03	2	423	fx(9000)	10	497,286	pass	0.011/0.040	pass	156/160	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:38:32	Vlan_412_BERToD_BCT-SW03	3	412	fx(9000)	10	500,000	pass	0.001/0.040	pass	1/1	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:38:46	Vlan_421_BERToD_SCL-SW04	4	421	fx(9000)	10	498,246	pass	0.001/0.040	pass	113/114	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:38:59	Vlan_422_BERToD_SJU-SW02	5	422	fx(9000)	10	500,000	pass	0.001/0.040	pass	31/33	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:39:12	Vlan_401_BERToD_MIA-SW17	6	401	fx(9000)	10	500,000	pass	0.000/0.040	pass	0/1	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:39:26	Vlan_427_BERToD_JAX-SW02	7	427	fx(9000)	10	500,000	pass	0.001/0.040	pass	7/18	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:39:39	Vlan_428_BERToD_MIA-SW18	8	428	fx(9000)	10	500,000	pass	0.001/0.040	pass	0/1	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:39:52	Vlan_429_BERToD_MIA-SW16	9	429	fx(9000)	10	485,364	pass	0.001/0.040	pass	0/1	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:40:06	Vlan_430_BERToD_MIA-SW18	10	430	fx(9000)	10	500,000	pass	0.002/0.040	pass	0/1	pass	0/0.000%	pass	0/0.000%
2024-12-09 20:40:25	Vlan_309_BERToD	4	309	emix	15	500,000	pass	0.024/0.040	pass	158/160	fail	20351/4.070%	pass	0/0.000%
2024-12-09 20:40:42	Vlan_309_BERToD-128B	5	309	fx(128)	15	285,291	fail	0.000/0.040	fail	0/160	fail	15366/5.386%	pass	0/0.000%
2024-12-09 20:41:12	Vlan_309_BERToD-1KB	6	309	fx(1024)	15	500,000	pass	0.021/0.040	pass	158/160	fail	84/0.017%	pass	0/0.000%
2024-12-09 20:41:30	Vlan_309_BERToD-1.5KB	7	309	fx(1518)	15	498,562	pass	0.020/0.040	pass	158/160	fail	18/0.004%	pass	0/0.000%
2024-12-09 20:41:47	Vlan_309_BERToD-4KB	8	309	fx(4096)	15	500,000	pass	0.021/0.040	pass	158/160	fail	4/0.001%	pass	0/0.000%
2024-12-09 20:42:05	Vlan_309_BERToD-9KB	9	309	fx(9000)	15	500,000	pass	0.017/0.040	pass	158/160	pass	0/0.000%	pass	0/0.000%

AmLight ExP
Americas Lightpaths Express & Protect

Lessons Learned #1 – Topology view vs. troubleshooting

- Sub-topologies are key to isolate faults

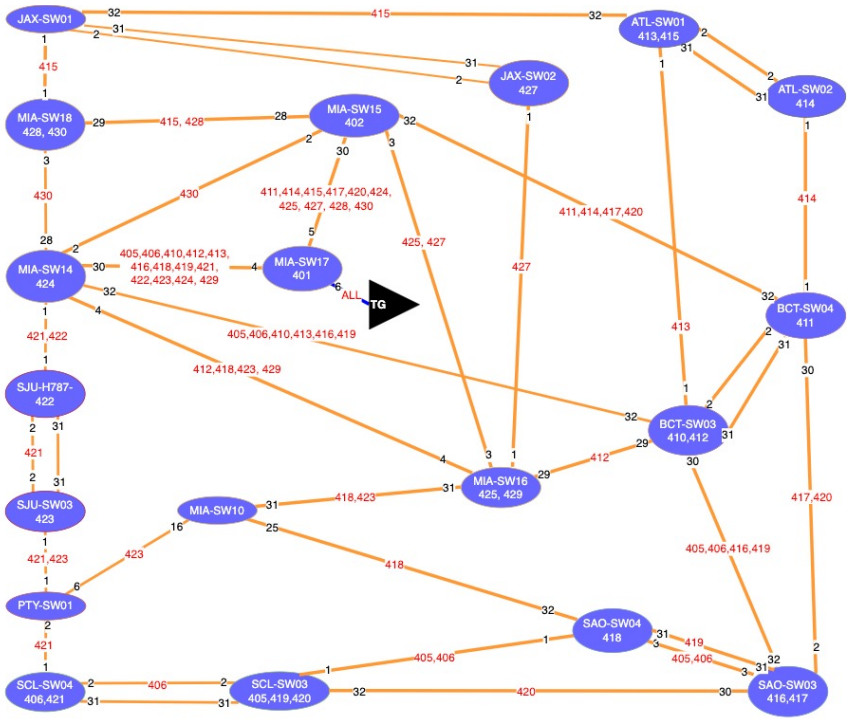

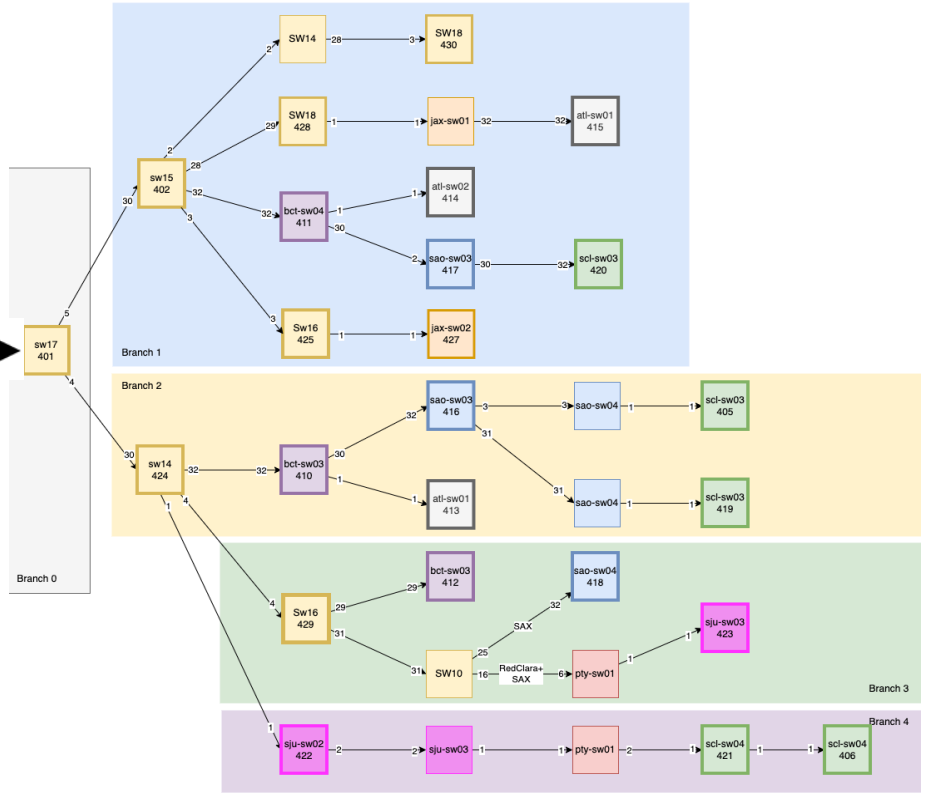

AmLight ExP
Americas Lightpaths Express & Protect