AmLight Express and Protect (AmLight-ExP), #OAC-1451018

2019 SAACC Meeting: AmLight-ExP updates since last meeting

and Prace

Los Ang

Julio Ibarra, PI Heidi Morgan, Co-PI Chip Cox, Co-PI Jeronimo Bezerra, Chief Network Architect Florida International University

Outline

- AmLight-ExP Vision
- Current Status
 - Network infrastructure
 - Monitoring and Measurement
- Demonstration
- Next Steps
 - Adding bandwidth
 - Increasing resilience
 - New Control and Data Planes
 - Improving Network Telemetry



AmLight Express & Protect Vision



- Community-operated network infrastructure
- Leased capacity on two submarine cable systems, evolving to a hybrid model that includes spectrum from Boca Raton to Sao Paulo
- Express (spectrum) capacity will provide up to 6 optical channels, which will be lit with 100G transponders today
- Protect (leased) capacity ring will back up the Express capacity



Current Status: Network Infrastructure

- 100G ring Miami-Fortaleza, Fortaleza-Sao Paulo, Sao Paulo-Santiago, Santiago-Panama, and Panama-Miami (solid yellow)
- 10G ring from Miami-Sao Paulo-Miami for protection (red dashed)
- 10G Miami-Santiago for protection (yellow dashed)
- 100G and 10G rings are diverse, operating on multiple submarine cables
- Total upstream capacity presently at 230Gbps!





Current Status: Network Infrastructure [2]





Current Status: Monitoring and Measurement

- Each AmLight PoP has a perfSonar node with two 10G NICs
 - Maddash portal available
 - Nodes being installed in Panama City and Fortaleza
- Three 100G network taps installed to support the IRNC AMIS project
 - Two in Miami
 - One in La Serena
- Two 100G DTNs for high bandwidth experiments and measurements
 - La Serena and Miami





What can we do with all this bandwidth and equipment?



Evaluating the network: SC18 demo!

ICT	
Services Documents Suf	ev l
LSST 100 Gbps Ne	atwork Demonstration at Supercomputing Conference
2018	in en penienen an en perieen paning contenere
2010	
Home Bogs Sister bog USST H	to rebit sweetow concompation of preparational coversions 2008
LSST BUILDERS LIST	LSST 100 Gbos Network Demonstration at Supercomputing Conference
	2018
To see paraword protected	
content log in using your LSS7	CI AN INLESS SIZE . BY MADO
account credientials	November 20. 2018 - The LSST Network Engineering Team (NET) had a strong presence at
Urarlania	the Supercomputing 2018 Conference-ISCIRI in Datas, TX, last week, including a
Usertogin	successful demonstration of the data transfer capabilities of the fiber optic networks that
Quemama *	will be used during LSST operations. Digital data were transferred from the Base Site in La
	Serena, Ohle, to the LSST Data Facility at the National Center for Supercomputing
	Applications INCSA) in Champaign IL. During the data transfer demonstration, a peak rate
	of soo Gipabits/Necond IGb/to was achieved for short periods, and a sustained rate of 60
Password 1	which an an an and the set of the
Password 1	GD/5 was achieved over a three hour period, exceeding the sent target, sins test was run
Password '	over links provisioned by several networking organizations: REUNA from La Serena to Santiago. FIU/Amlight from Santiago
Password	over links provisioned over a three nour period, exceeding the sent target, this test was nan over links provisioned by several networking organizations: REUNA from La Serena to Santiago. PIU/AmBight from Santiago to Mami, SCinet trom Hiami to Chicago IStarlight, and NCSA from Chicago to Champaign. SCinet links provided by



Evaluating the network: SC18 demo!





Evaluating the network: SC18 demo!

National's Utilization - DTN in La Baraina

SuperComputing 2018 +

Network Utilization

in this too, two graphs are plusen.

The left graph shows the network unbuston of a disk to disk data transfer, from La Serena, Ohle to Ohenpatgo/Rinols, The disk to disk data transfer is set to use to more than 400bps and it is performed, using the freehold/MOTM toxil.

The right graph shows the remerk utilization of a memory to memory data transfer, also from La bernas/Dole to Divergarge/Binelo. The memory-tomemory data transfer is set to use no more than IdDlays. Traffic is generated using NUTTOP







Next Steps: New Links

Spectrum activation in motion:

- Two 100G channels via Fortaleza (RNP and ANSP)
- Two 100G channels from Miami to Sao Paulo (ANSP)
- Total upstream capacity by May 2019:
 - 630Gbps
- Total network capacity by May 2019:
 - 1,230 Gbps / 1.23 Tbps





Next Steps: Going North!



Americas Lightpaths **Express & Protect**

Next Steps: New Data Plane

AmLight-ExP data plane refreshment:

- With the increasing number of 100G ports, a new switching solution has to be provided
- NoviFlow 100G Tofino switch seems to be the best solution so far:
 - 32 x 100G ports, 64 x 100G ports, and 32 x 400G ports (July)
 - Supports SDN, programmable data planes, and network telemetry

Solution cost:

- Extreme Networks: \$19,900/per 100G LR4 interface
- NoviFlow: \$2,525/per 100G LR4 interface



Next Steps: New Control Plane

AmLight-ExP control plane refreshment:

- Current SDN controller will be replaced by a new solution developed between AmLight and Kytos Development Team
- OpenFlow 1.0 will be replaced
- New solution will provide better interface to users and network devices, telemetry reports in real time, and microservices will be used for future network functions.





Kytos UI











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, without requiring intervention or work by the control 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
 - Streamed network telemetry for useful microburst detection and queue utilization at a sub-second interval
 - Complete view of network state in the flow's path



Network Telemetry





Questions addressed by INT

How did this packet get here?

- The sequence of network devices a packet visited along its path.
 - LAG? No problem. ECMP? No problem. Layer 2 network? No problem!

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.



THANK YOU!













