AmLight ExP

Americas Lightpaths Express & Protect

AmLight Express and Protect (AmLight-ExP):

Enabling Research and Education in the Americas via cooperation and collaboration

Julio Ibarra, PI

Heidi Morgan, Co-PI

Chip Cox, Co-PI

Jeronimo Bezerra, Chief Network Architect
Florida International University

Americas Summit of Academic Networks

January 23-24, 2019

Mexico City

Outline

- Overview
- Current Status
- Planned activities for 2019
- Science Drivers
- AmLight-INT
- AARCLight
- Conclusion

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 100G ring will back up the Express capacity

AmLight ExP
Americas Lightpaths Express & Protect

Partners and Goals

- AmLight-ExP interconnects the U.S. to key aggregation points in South and Central America (Brazil, Chile, Panama)
- 5-year Cooperative Agreement with the U.S. National Science Foundation. Year 5 starts April 1, 2019.
- Cooperative and collaborative partnerships with ANSP, RNP, CLARA, REUNA, AURA, FLR, and Internet2
- Continue evolving a rational network infrastructure using both optical spectrum and leased capacity

Cooperation and Collaboration

- Sharing physical resources
  - Network Links
  - Colocation
  - Compute resources
- Sharing operations resources
  - Network Engineers
  - Service Level Agreements (SLAs)
- Supporting science and education in the Americas
  - Astronomy community: cooperation and collaboration in SAACC
  - Large Synoptic Survey Telescope (LSST)
  - High-energy physics community in Brazil and Chile
  - Etc.

Outline

- Overview
- Current Status
- Planned activities for 2019
- Science Drivers
- AmLight-INT
- AARCLight
- Conclusion

Regional Network Infrastructure

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

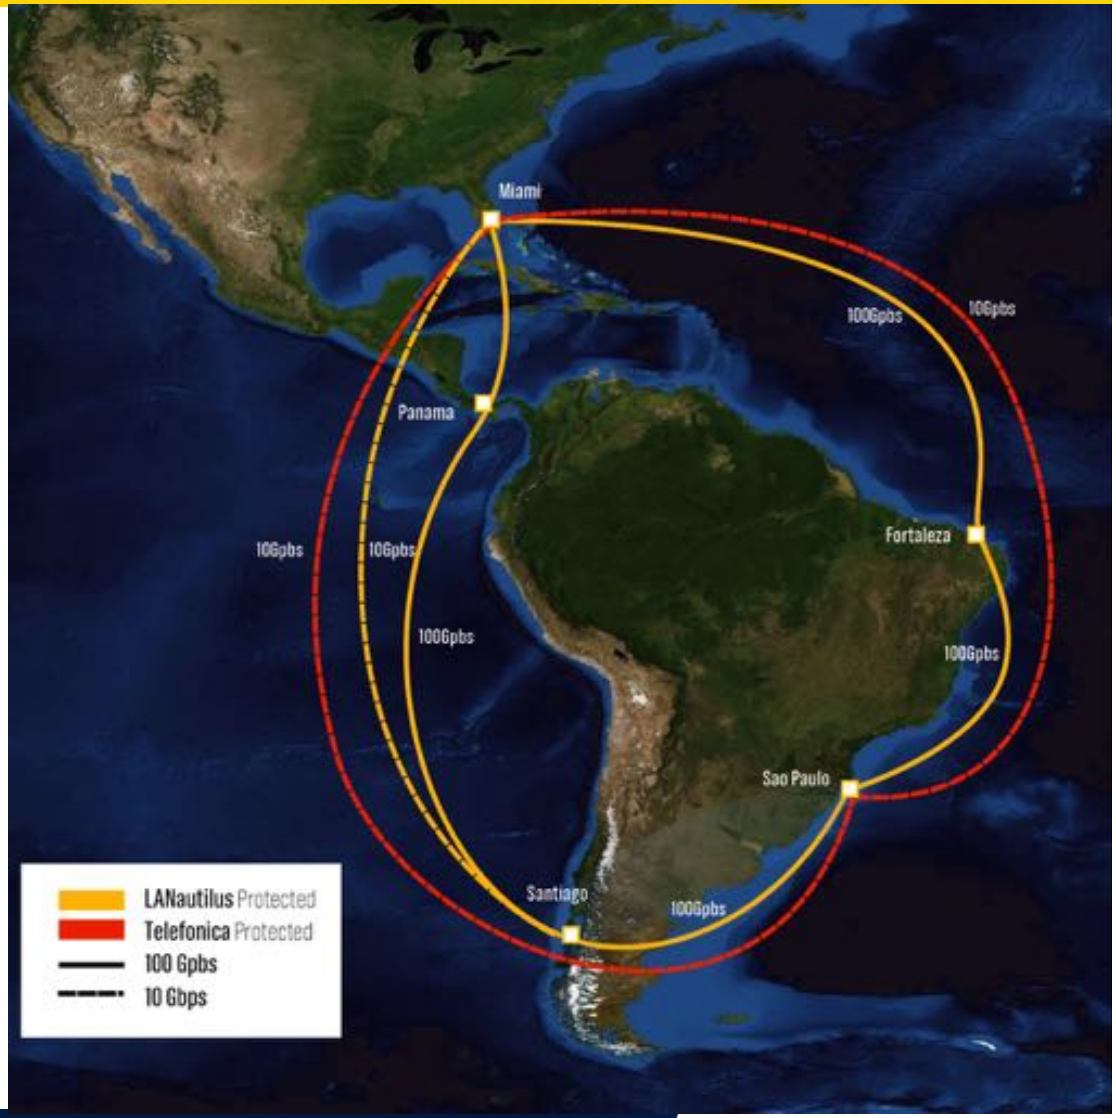

AmLight ExP
Americas Lightpaths Express & Protect

Network virtualization and programmability

- Supporting SDN in production since 2014
- Researchers use slicing/virtualization to prototype network-aware applications
  - Can implement testbeds with real network devices
  - Can validate research in a production environment, and at scale
  - E.g., Existing testbed for bandwidth prioritization and reservation to support big-data and real-time applications (LSST use case)
- SDN Looking Glass for integration of both optical and packet domains for a complete network troubleshooting and visualization

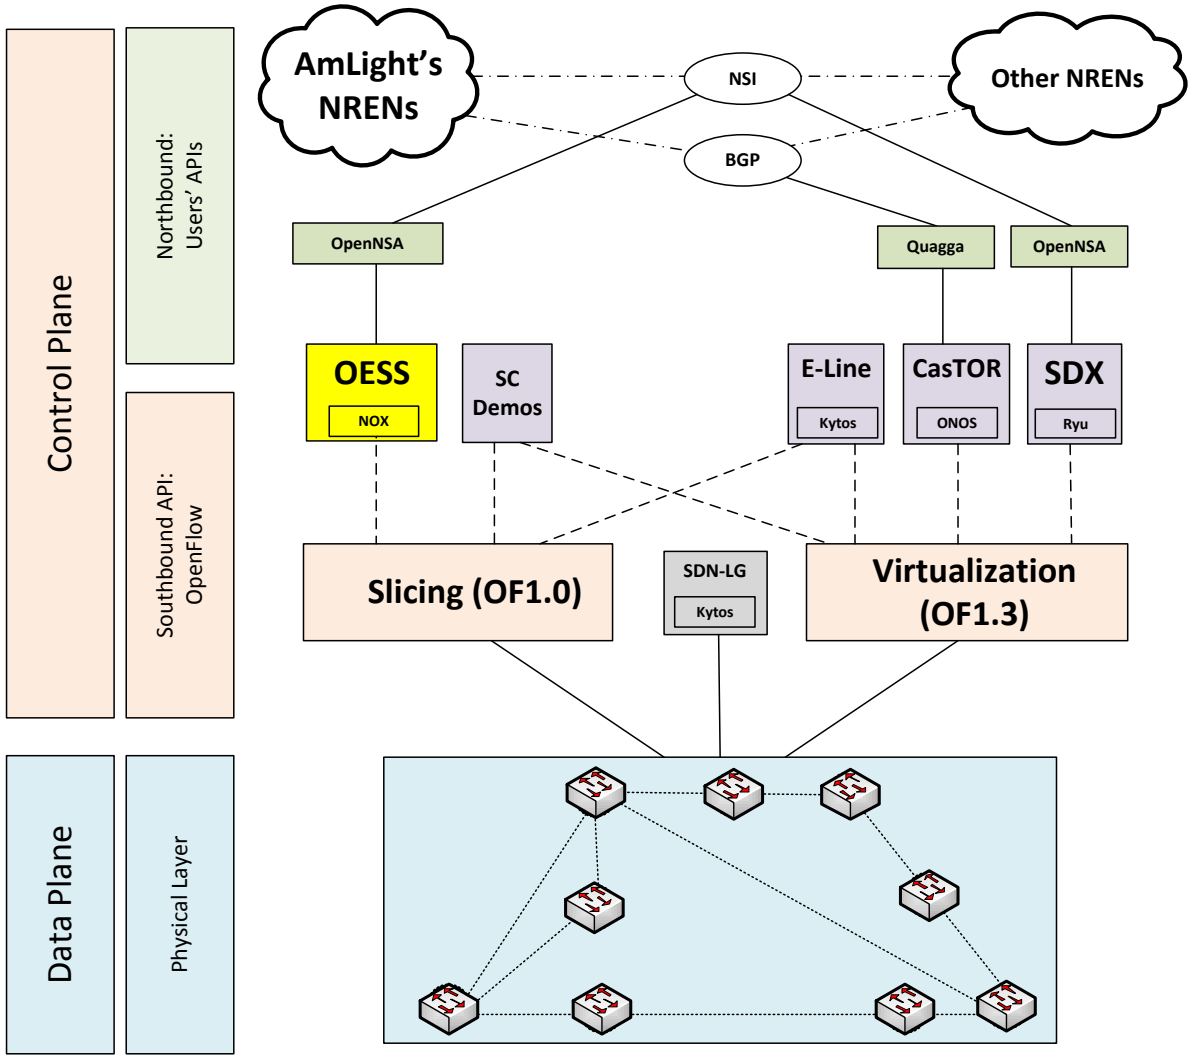

AmLight ExP
Americas Lightpaths Express & Protect

AmLight-INT project

- AmLight-INT is a new project, funded by NSF/OAC, that will include telemetry functions in the AmLight network
- Focuses on improving Operations, Administration and Maintenance (OAM) methods for the AmLight-ExP SDN production environment
- Provides a response to the LSST SLA requirements by enabling real-time network troubleshooting functions

Outline

- Overview
- Current Status
- Planned activities for 2019
- Science Drivers
- AmLight-INT
- AARCLight
- Conclusion

Build Express backbone, and enhance resilience on AmLight

- Increase AmLight network resilience and capacity
  - Activate Express:
    - Boca Raton, Fortaleza, Sao Paulo
    - 6 (green lines) x 100G links
      - 4 managed by RNP
      - 2 managed by FIU/ANSP/LSST
- Expand AmLight network
  - Add PoPs in Jacksonville and Boca Raton
- Enhance monitoring capabilities
  - Enhance SDN Looking Glass
  - Activate INT capabilities
  - Add P4+INT NICs to perfSonar nodes

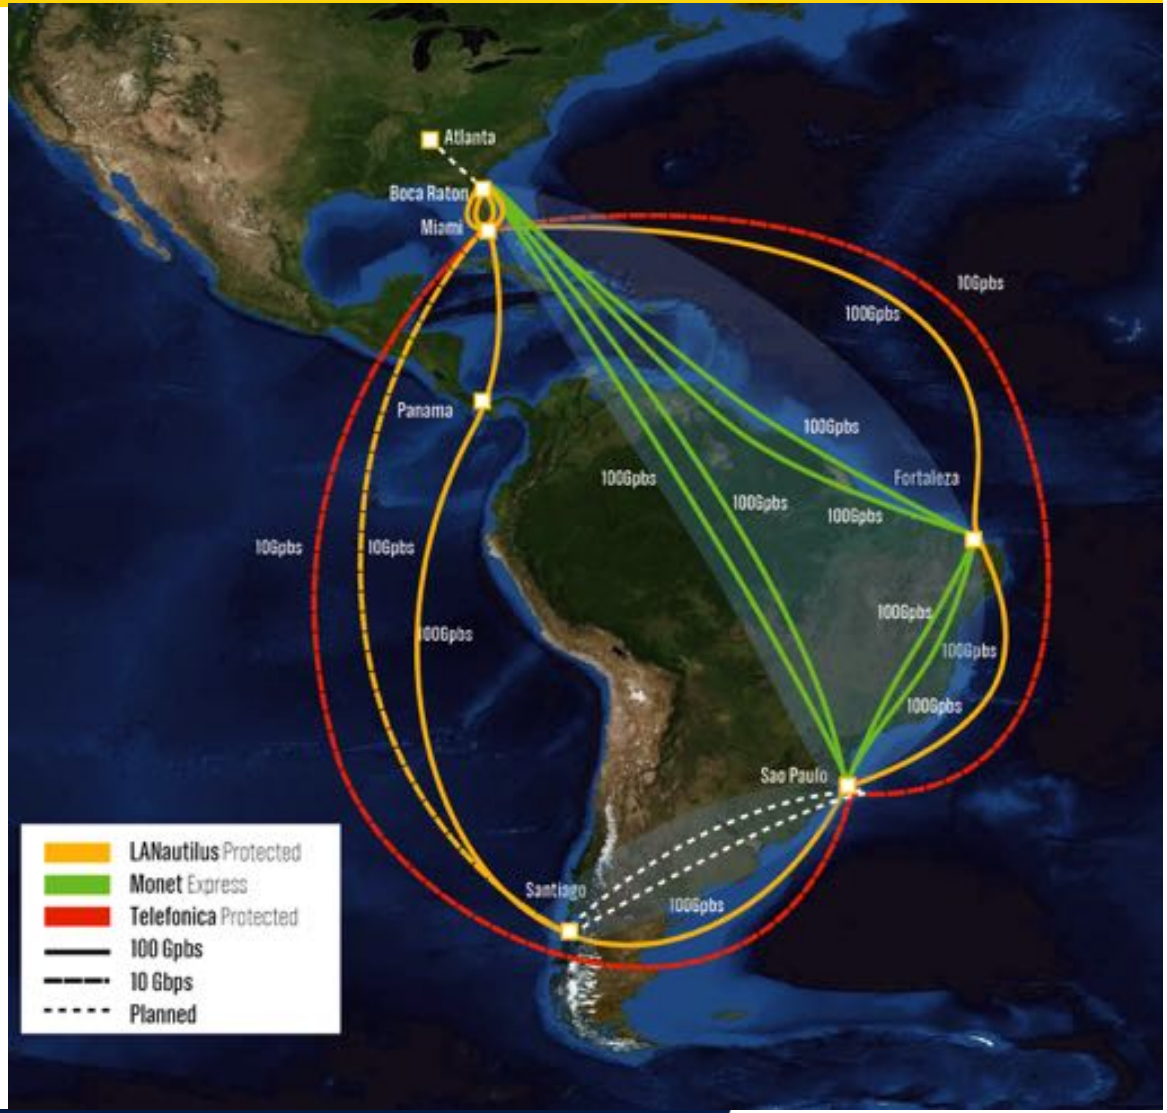

AmLight ExP
Americas Lightpaths Express & Protect

Other activities in 2019 for AmLight-ExP

- Network Programmability:
  - New SDN controller (Kytos E-Line) to manage AmLight-ExP links
  - Replacement of the current southbound APIs (OF 1.0 and 1.3) for a OpenFlow 1.4 version
  - Support for QoS, per Bidirectional Forwarding Detection (BFD) link, among other benefits
- Network Connectivity:
  - 600Gbps between Boca Raton and Miami
  - 200Gbps between Miami and Jacksonville to connect to Internet2

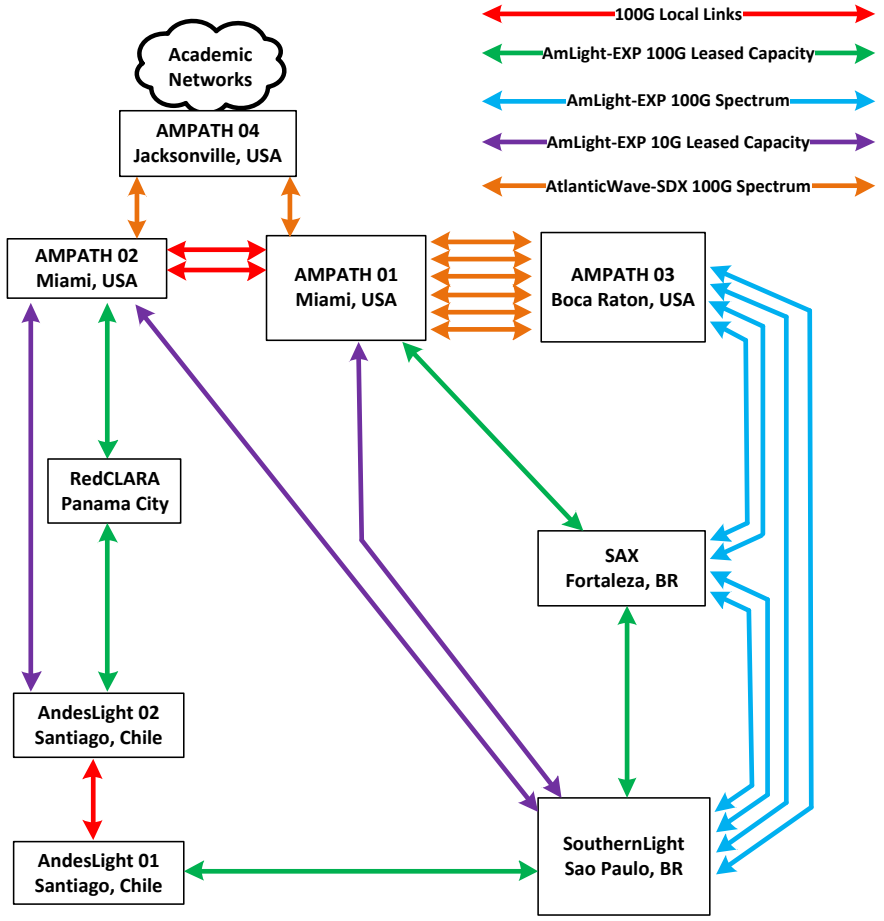

AmLight ExP
Americas Lightpaths Express & Protect

SAACC and LSST NET Spring meetings

- AmLight South American Astronomy Coordination Committee (SAACC) Meeting Spring 2019, scheduled for April 17, 2019 in Santiago, Chile
- LSST NET meeting, scheduled for April 18, in Santiago
- Participants can register at https://saacc-2019-chile.eventbrite.com
- Logistics and video conference details will be provided later
- More details can be found on AmLight website and SAACC Facebook page

Outline

- Overview
- Current Status
- Planned activities for 2019
- Science Drivers
- AmLight-INT
- AARCLight
- Conclusion

Large Synoptic Survey Telescope (LSST)

- LSST is a large-aperture, wide-field, ground-based optical telescope under construction in northern Chile
- The 8.4 meter telescope will take a picture of the southern sky every 27 seconds
  - producing a 13 Gigabyte data set
- Each image must be transferred to the archive site at NCSA in Champaign, Illinois, within 5 seconds
  - Reason: for processing to be completed in time to generate "transient alert" notifications to the worldwide astronomical community within 60 seconds
- In response to LSST's transient alert requirements, the network must be scalable, highly available, reliable, and provide high-throughput and guaranteed bandwidth

AmLight ExP
Americas Lightpaths Express & Protect

Challenge: LSST Use Case

- LSST will be remotely operated from Tucson, Arizona
- Every 27 seconds throughout the night, LSST will take a 6.4GB picture of the sky, process it, generate transient alerts (6.3GB) from this picture, and then send a 13GB data set to NCSA, at Urbana-Champaign, Illinois
  - From the 27-seconds window, only 5 seconds are available for data transmission
  - Multi-traffic types with different priorities (db sync, control, general Internet traffic) must be supported

- A 0.001% of packet loss will compromise the LSST application
- Packet loss isolation will have to be handled in real-time
- Strict SLA:
  - MTBF (180 days in a year)
  - MTTR (48 hours)

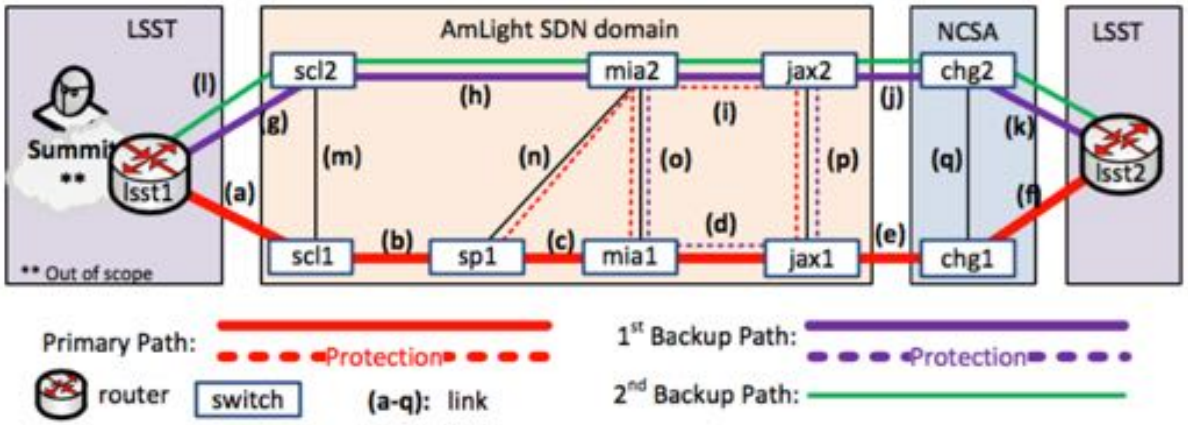

AmLight ExP
Americas Lightpaths Express & Protect

Super Computing 2018

- AmLight participated in 4 Network Research Exhibitions (NRE) at SC18
- Data transfers of LSST during Operations was a NRE
  - Figure shows the testbed built from La Serena to NCSA
- Networks supporting NRE were: AURA, REUNA, CLARA, FIU-AmLight, SCInet/CenturyLink, StarLight/MREN, NCSA
- Peaks of 100G were achieved
- Sustained rates of 80G were achieved
- Press release: https://www.amlight.net/?p=3345

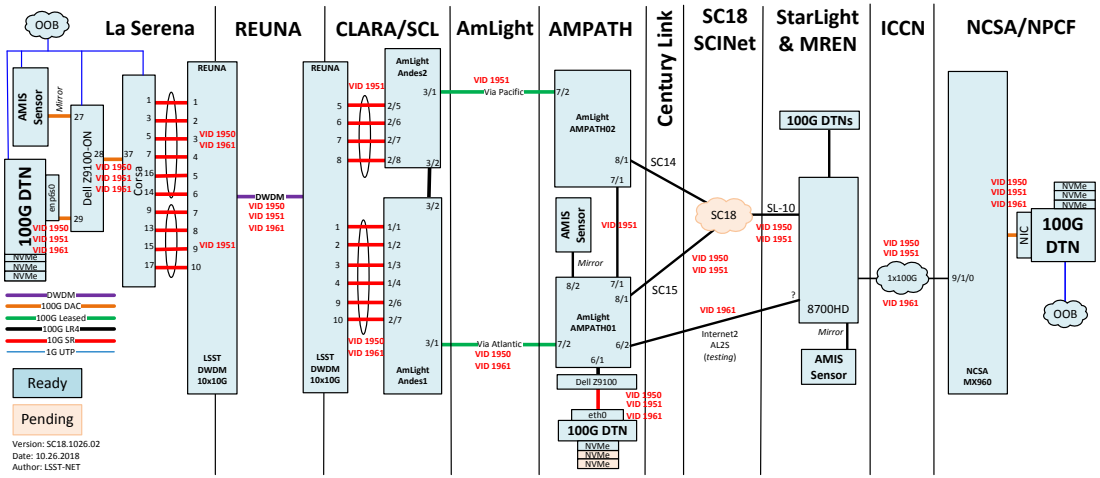

AmLight ExP
Americas Lightpaths Express & Protect

UNESP Center for Scientific Computing

- LHC Data Analysis & processing
- Processing for Open Science Grid (OSG), and other international instruments
- Received 2,614 TB in flows from U.S. sources in 2017
- Transmits 792 TB in flows to U.S. destinations

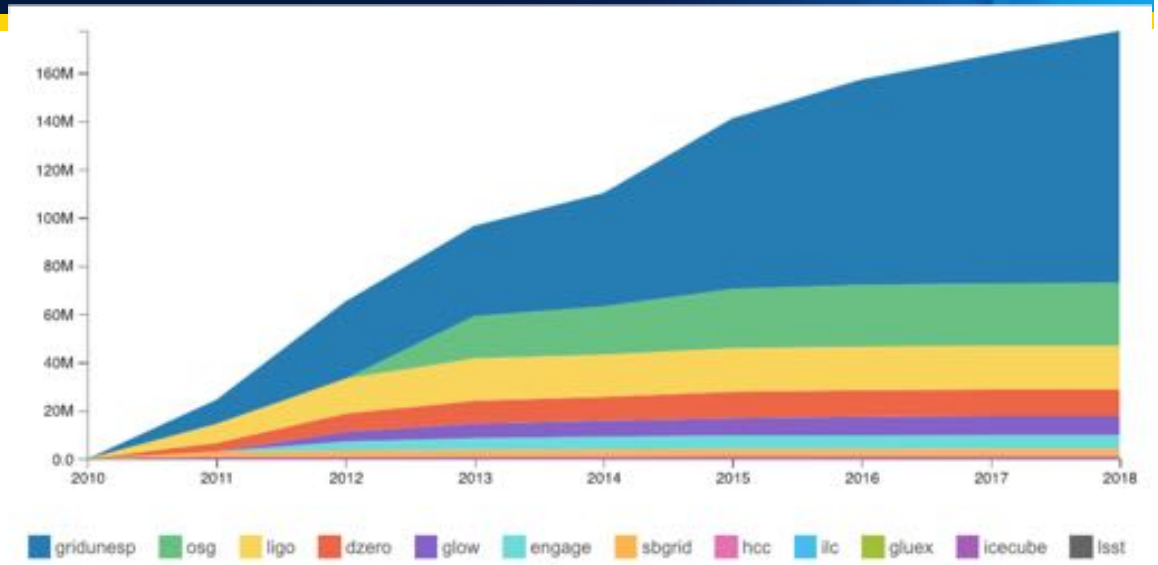

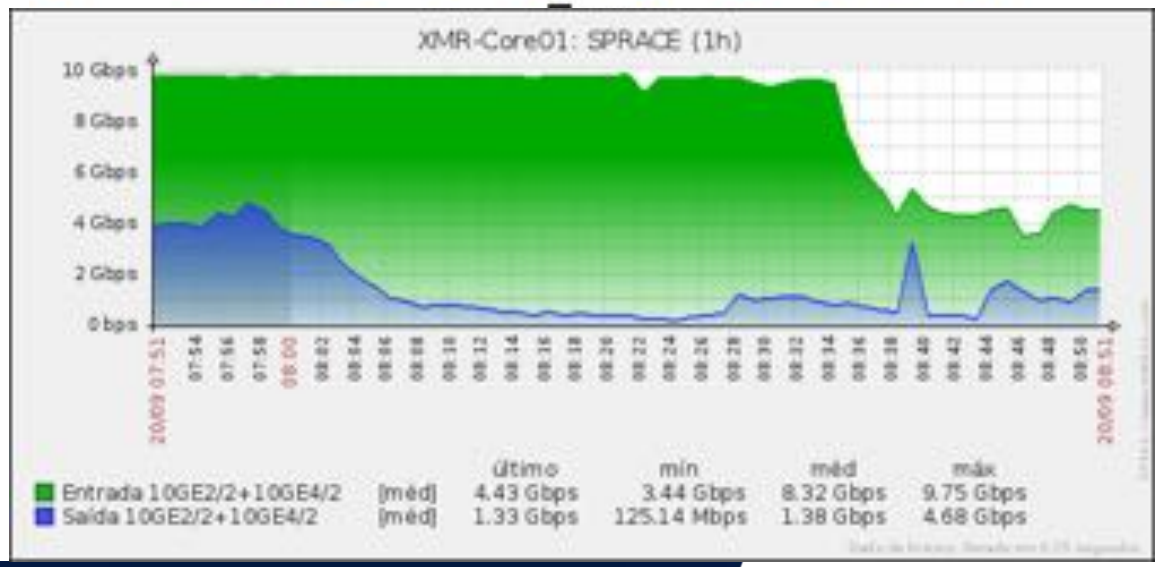

AmLight ExP
Americas Lightpaths Express & Protect

Outline

- Overview
- Current Status
- Planned activities for 2019
- Science Drivers
- AmLight-INT
- AARCLight
- Conclusion

What is In-band Network Telemetry (INT)?

- INT records network-state information in the packet while the packet traverses a path between two points in the network:
  - In-band refers to mechanisms that add tracing or other types of telemetry information to data traffic
  - Complements current out-of-band mechanisms based on ICMP or other types of probe packets, such as Ping, Traceroute and perfSonar.
  - Basically, INT adds metadata to each packet with information that could be used later for troubleshooting activities.
- Example of information added:
  - Timestamp, ingress port, egress port, pipeline used, queue buffer utilization, CPU utilization, Battery Utilization, Sequence #, and many others
- As metadata is exported directly from the Data Plane, Control Plane is not affected:
  - *Can track/monitor/evaluate EVERY single packet at line rate!*

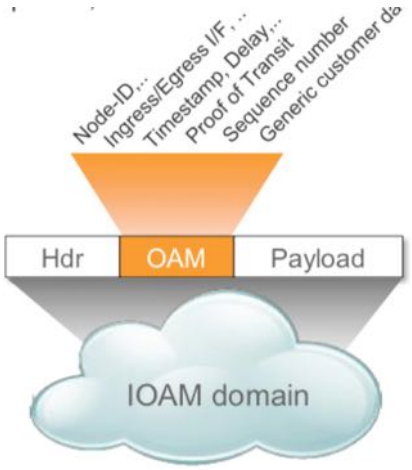

AmLight ExP
Americas Lightpaths Express & Protect

AmLight-INT Workflow

- INT-capable SDN switches will export INT metadata of pre-selected SDN circuits defined by the AmLight-ExP SDN controller
- INT-enabled Linux Sensors, installed at multiple sites, will collect the INT metadata directly from the INT-capable switches
- Once INT metadata is collected, it will be processed, analyzed, consolidated, and exported to the AmLight SDN-LG
- The AmLight SDN-LG will receive feeds from all INT-enabled Linux Sensors then consolidate, store the network status data
- Users interested in the network status data will be able to import it using the SDN-LG API

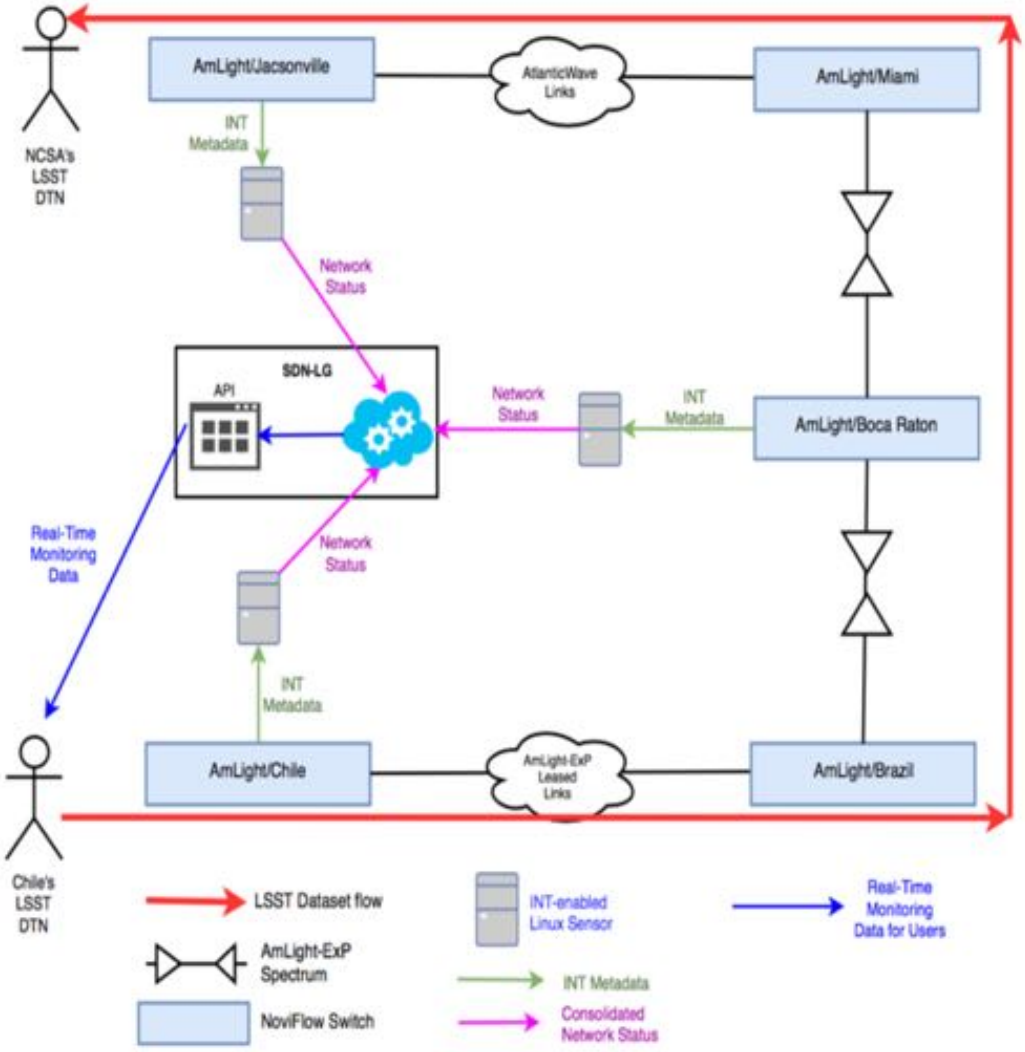

AmLight ExP
Americas Lightpaths Express & Protect

Outline

- Overview
- Current Status
- Planned activities for 2019
- Science Drivers
- AmLight-INT
- AARCLight
- Conclusion

New submarine cables in the South Atlantic

- **Monet:** Boca Raton, FL-Fortaleza, BR. Operational
- South Atlantic Cable System (SACS): Fortaleza, BR- Sangano, Angola. Q3 2018
- South Atlantic Inter Link (SAIL): Fortaleza, BR – Kribi, Cameroon. Q3 2018 (TBD)
- EllaLink: Fortaleza, BR – Sines, Portugal. RFS 2020
- **America Movil (AMX-1):** Fortaleza, BR – Jacksonville and Hollywood, FL. Operational
- SABR: Cape Town, SA – Recife, BR. RFS 2019
- Seabras-1: Praia Grande, BR – Wall Township, NJ.
- Fortaleza is a landing point for all cables, except for SABR & Seabras

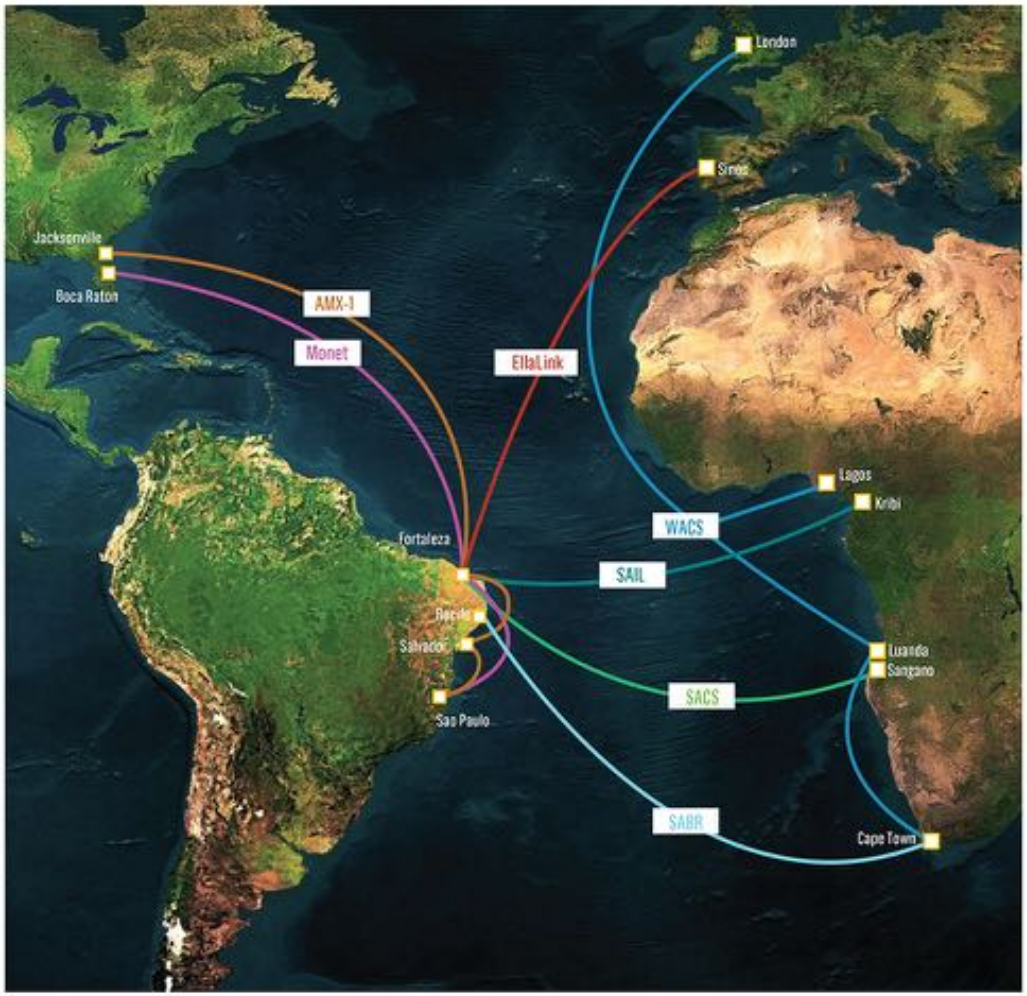

AmLight ExP
Americas Lightpaths Express & Protect

AARCLight: Network infrastructure resources in the Southern Hemisphere (potential leveraging)

- 225GHz of spectrum on Monet committed in AmLight-ExP project
- 40GHz of spectrum on SACS is available to the R&E community
- TENET operates 220G of capacity on WACS
- South Atlantic eXchange point (SAX) is under development in Fortaleza, led by RNP
- R&E exchange point in Cape Town operated by SANREN and TENET
- R&E exchange point in Lagos, operated by WACREN

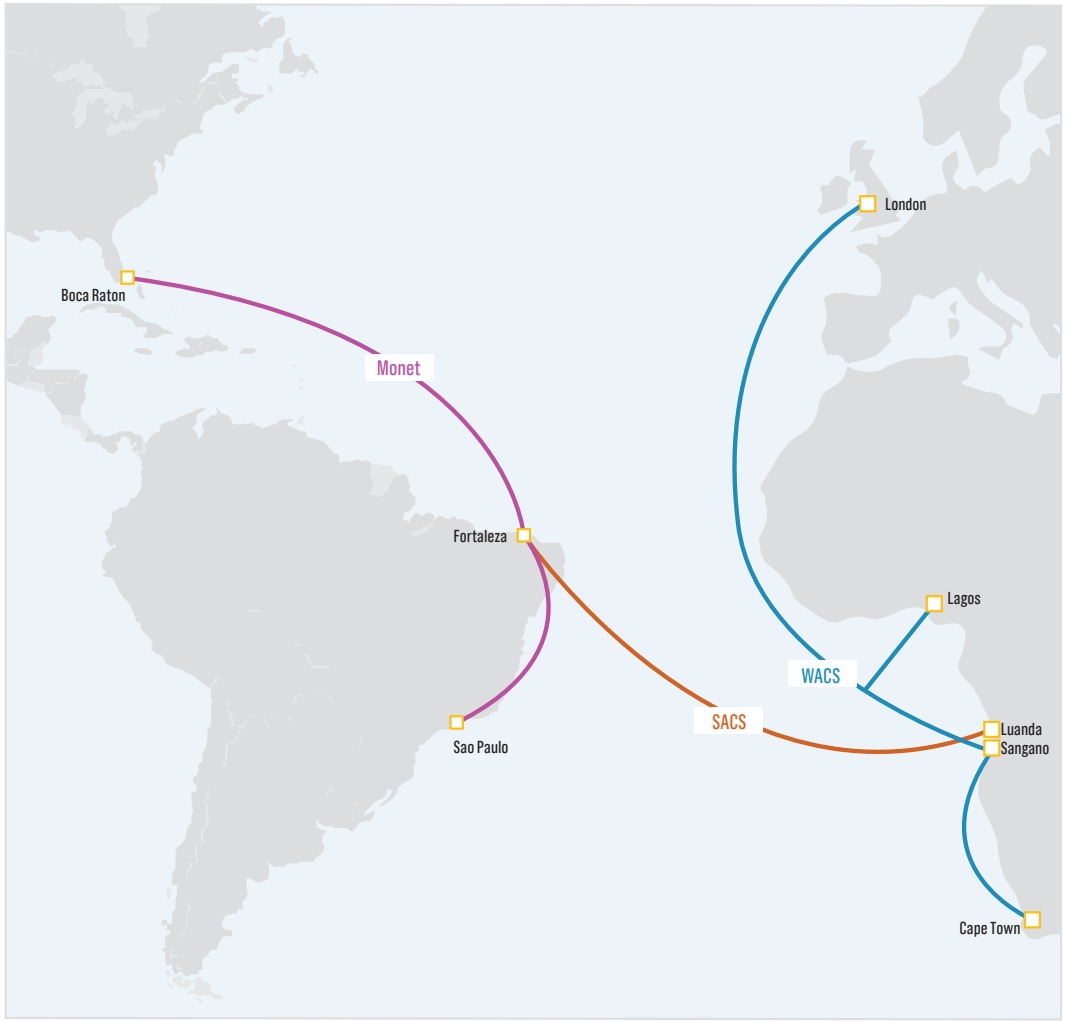

AmLight ExP
Americas Lightpaths Express & Protect

About AARCLight

- Project of the U.S. National Science Foundation, Award #OAC-1638990, at FIU
- Aims to define a strategy for research and education network connectivity between the US and West Africa
- Aims to coordinate planning efforts among stakeholders in the U.S., Africa, and Brazil
- Aims to create economies of scale
  - Making use of the offered spectrum
  - Towards serving the broadest communities of interest in research and education

Planned Activities:

1. Establish collaborative partners
2. Establish a contract with Angola Cable for Spectrum
3. Develop a network design in coordination with AARCLight partners
4. Develop a plan to integrate the science disciplines with the network strategy
5. Develop a community engagement plan
6. Preparing a case study to present to the communities of interest and National Science Foundation (NSF)

Opportunities in Angola

- SACS and WACS are connected at Sangano Cable Station in Angola, with possibility of being extended to Luanda via AC backhaul
- SACS creates an opportunity to connect West Africa to Brazil, and the USA

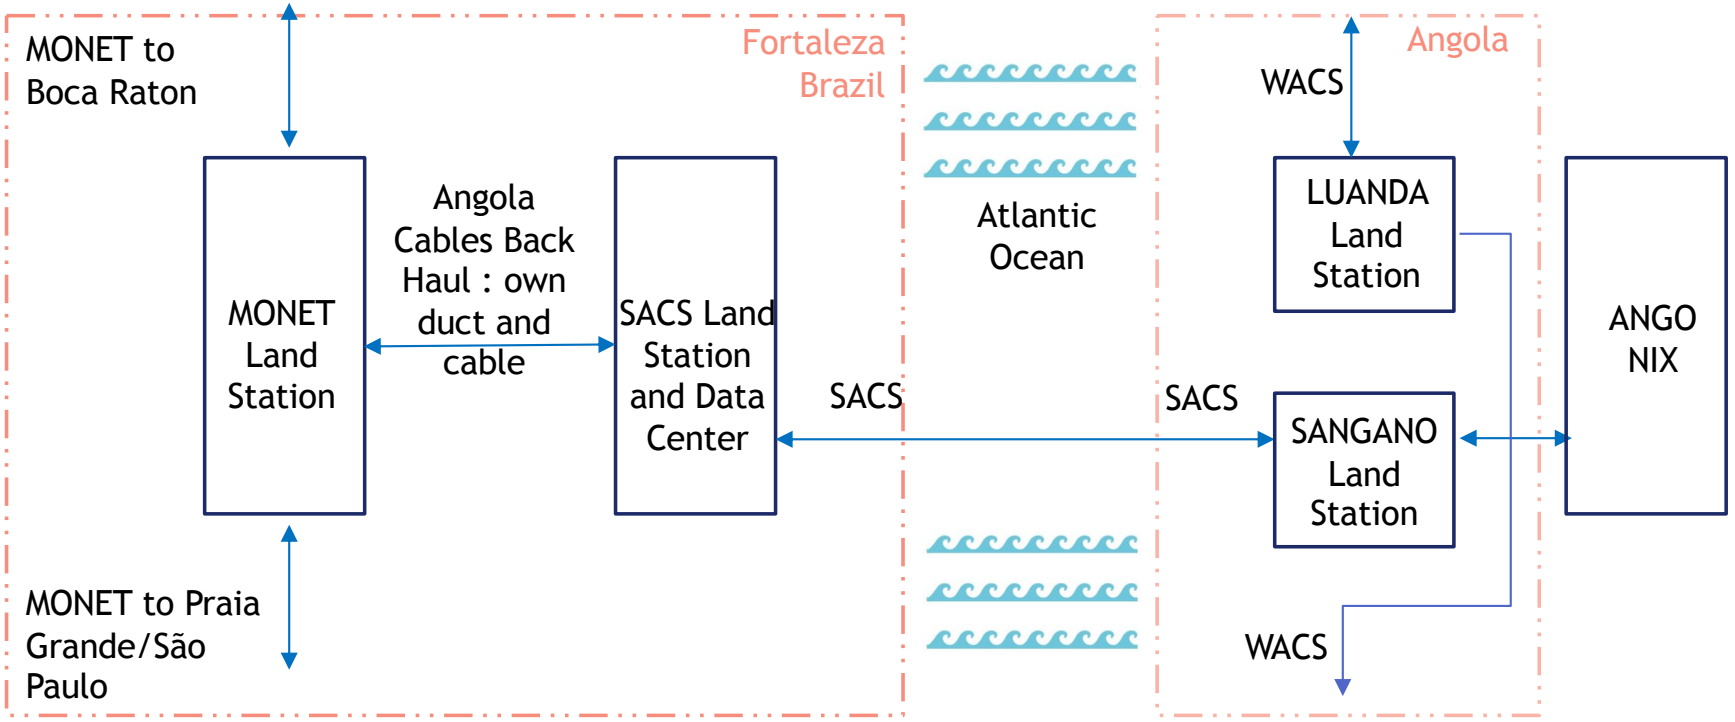

AmLight ExP
Americas Lightpaths Express & Protect

New GOLE/GXP in Fortaleza: South Atlantic Crossroads (SAX)

- Fortaleza is a key aggregation point in the South Atlantic
  - 2018: Monet (to US)
  - 2019: SACS (to Angola)
  - 2020: Ellalink (to EU)
- RNP's Global Exchange Point, SAX, to be deployed in Fortaleza, CE, Northeast Brazil as a hub for multiple international connections in the region
- SAX is designed to be compliant with both GNA and GLIF initiatives
- International partners, such as CLARA, AmLight and GEANT, are crucial to providing the new circuits reaching SAX.

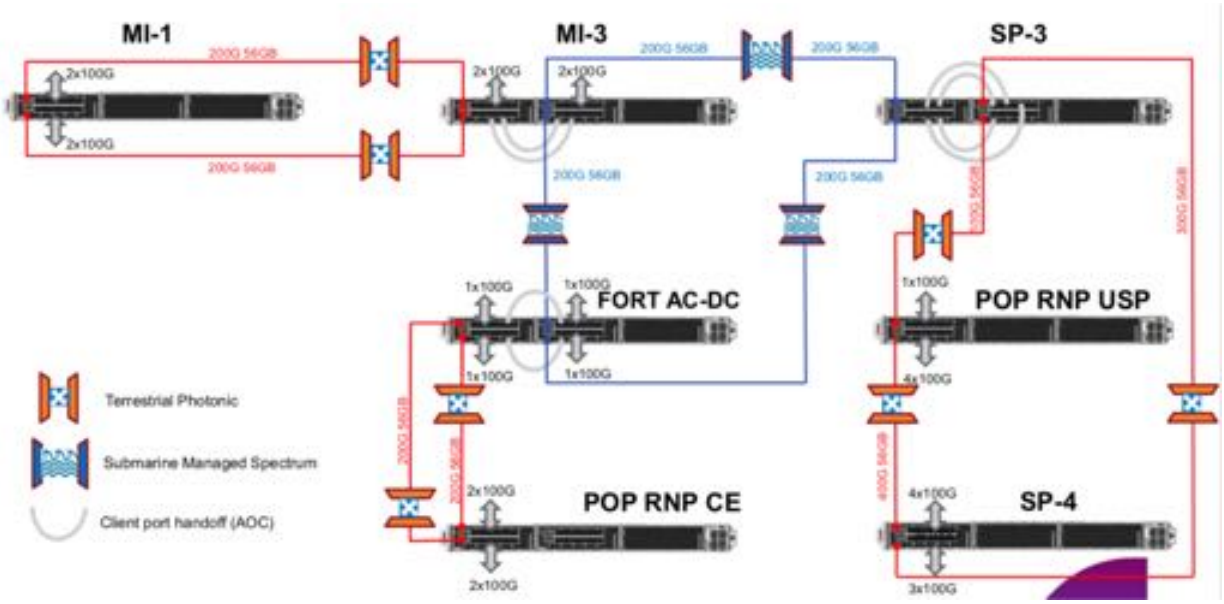

AmLight ExP
Americas Lightpaths Express & Protect

Outline

- Overview
- Current Status
- Planned activities for 2019
- Science Drivers
- AmLight-INT
- AARCLight
- Conclusion

Conclusion

- AmLight continues to
  - Develop a network infrastructure that interconnects North America with key aggregation points in South and Central America (Brazil, Chile, Panama)
  - Evolve as a rational R&E network infrastructure that is scalable, effective and efficient
  - Facilitate effective peering among academic networks and communities of interest
  - Respond to the network requirements from research communities
  - Continua su exito a traves de cooperacion y colaboracion
  - Achieve success in its goals through cooperation and collaboration

Thank You!

- NSF AmLight-ExP, AtlanticWave-SDX, CC*IIE, AARCLight,AmLight-INT, AMPATH infrastructure, science application support, education, outreach and community building efforts are made possible by funding and support from:
  - National Science Foundation (NSF) awards OAC-1451018, OAC-1451024, OAC-1541402, OAC-1638990, OAC-1848746
  - FAPESP, ANSP – grant no. 2008/52885-8
  - Rede Nacional de Ensino e Pesquisa (RNP)
  - Cooperación Latino Americana de Redes Avanzadas (CLARA)
  - Association of Universities for Research in Astronomy (AURA)
  - Florida International University
  - Latin American Research and Education community
  - The many national and international collaborators who support our efforts

THANK YOU!


NSF


FIU
FLORIDA
INTERNATIONAL
UNIVERSITY


AMLIGHT


AMPATH
PATHWAY OF THE AMERICAS


AmLight ExP
International Distributed Exchange Point


AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE


AARC Light
Americas Africa Research & eduCation
Lightpaths

Opportunity: Hybrid Network Operation

- Intra-domain coordination is changing
  - Integrating spectrum and leased capacity into the SDN operation
  - More complexity in the SDN Controller configuration

- Inter-domain coordination across multiple network operators is becoming more complex
  - Social engineering challenge
    - Impact to application will involve ALL network operators in the path
  - Multiple vendor technology interoperability challenges
  - Inter-domain troubleshooting challenges

Gracias!

- NSF OpenWave, AmLight, OSDC-PIRE, CC-NIE, CC*IIE, AMPATH, AtlanticWave infrastructure, science application support, education, outreach and community building efforts are made possible by funding and support from:
  - National Science Foundation (NSF) awards ACI-1451018, ACI-1451024, ACI-1440728, ACI-0963053, ACI-1140833, ACI-1246185, ACI-1341895
  - FAPESP, ANSP – grant no. 2008/52885-8
  - Rede Nacional de Ensino e Pesquisa (RNP)
  - Cooperación Latino Americana de Redes Avanzadas (CLARA)
  - Association of Universities for Research in Astronomy (AURA)
  - Florida International University
  - Latin American Research and Education community
  - The many national and international collaborators who support our efforts

Current Status: Network Virtualization and Programmability

- AmLight became an SDN network in 2014
- Slicing for network testbeds using OpenFlow 1.0
- Virtualization using Corsa switches for network testbeds using OpenFlow 1.3
- Researchers use slicing/virtualization to prototype network-aware applications
  - Can implement testbeds with real network devices
  - Can validate their research in a production environment, and at scale
- Current testbeds:
  - ONOS/CasTOR, FIBRE, Kytos E-Line, Awave-SDX
- SDN Looking Glass tool for troubleshooting the data plane

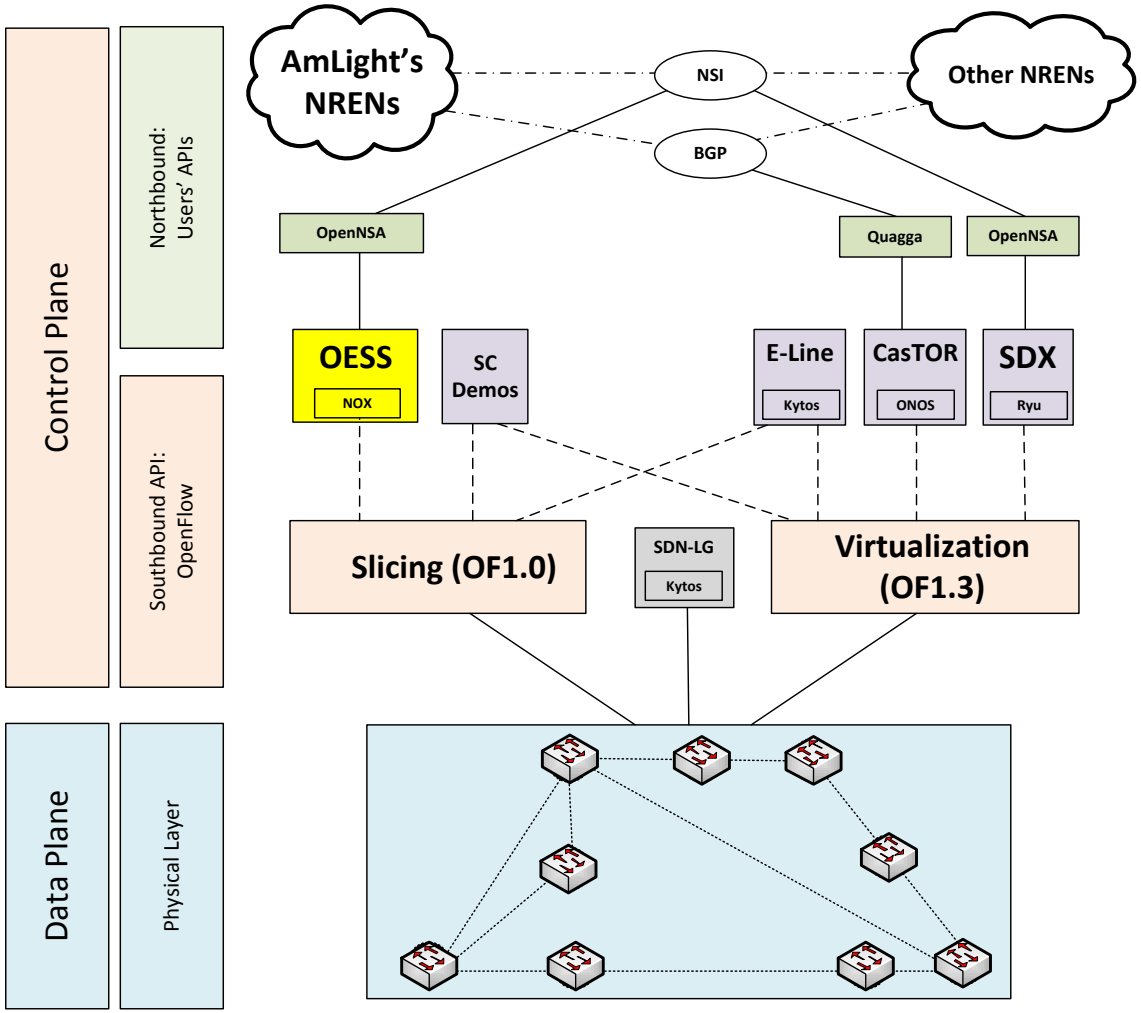

AmLight ExP
Americas Lightpaths Express & Protect

Opportunities: Fortaleza and Southern Hemisphere

- New data-intensive science instruments in the southern hemisphere:
  - SKA in South Africa
  - Large Synoptic Survey Telescope (LSST) in Chile
- New submarine cables in South Atlantic
  - SACS submarine cable to Angola (RFS Q3 2018)
  - South Atlantic Inter-Link (SAIL) to Cameroon (2018)
  - EulaLink submarine cable from Fortaleza to Portugal (2020)
- Network aggregation in the southern hemisphere
  - Fortaleza as a south Atlantic hub
- Leveraging network infrastructure linking the U.S., South America, and Africa

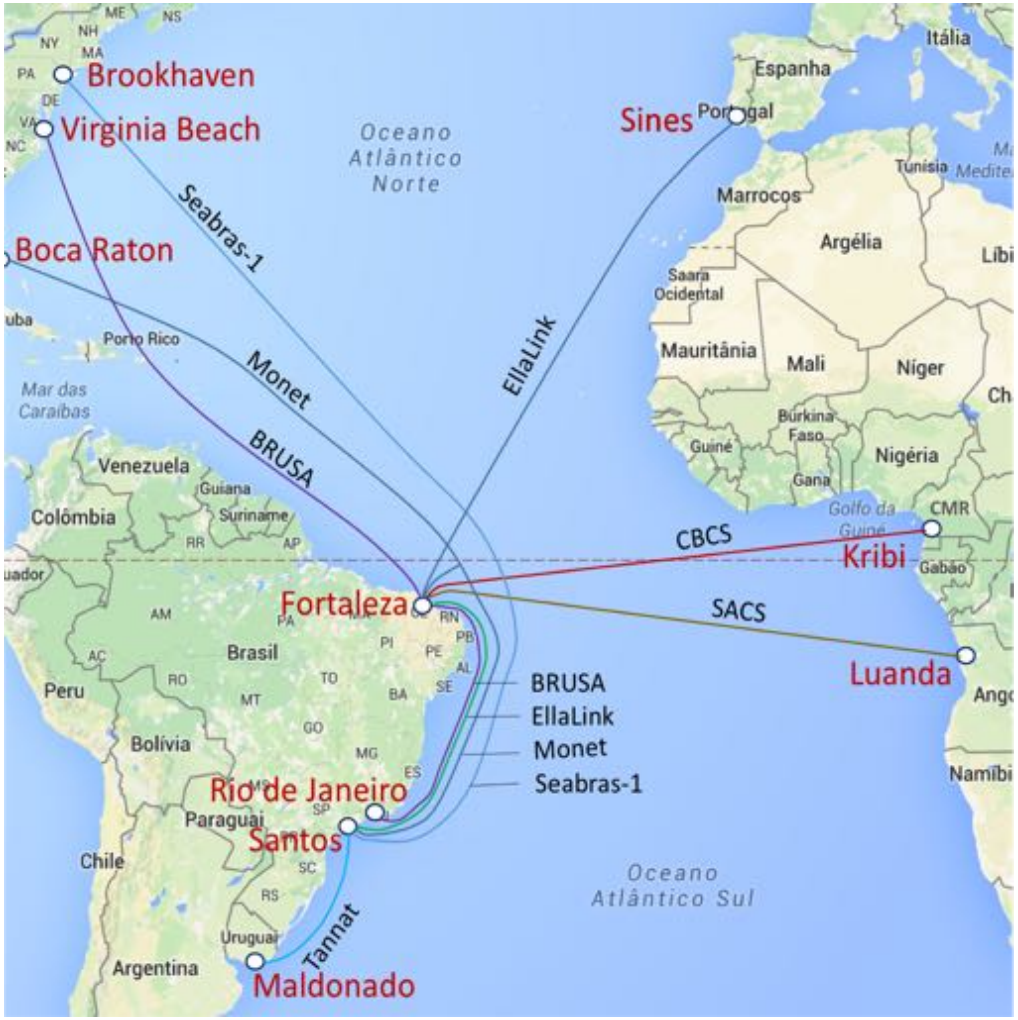

Americas Summit of Academic Networks 2019

AmLight ExP
Americas Lightpaths Express & Protect

Opportunities: Bandwidth abundance

- AmLight-ExP will have access to 225GHz of linear spectrum on Monet
  - Goal is to channelize this spectrum in 6x37.5GHz channels
  - Resulting in 6x100G ethernet circuits, initially
  - 200G of Protect (leased) capacity
- Improving international transit in Florida to Internet2 by adding protection to 100G transport

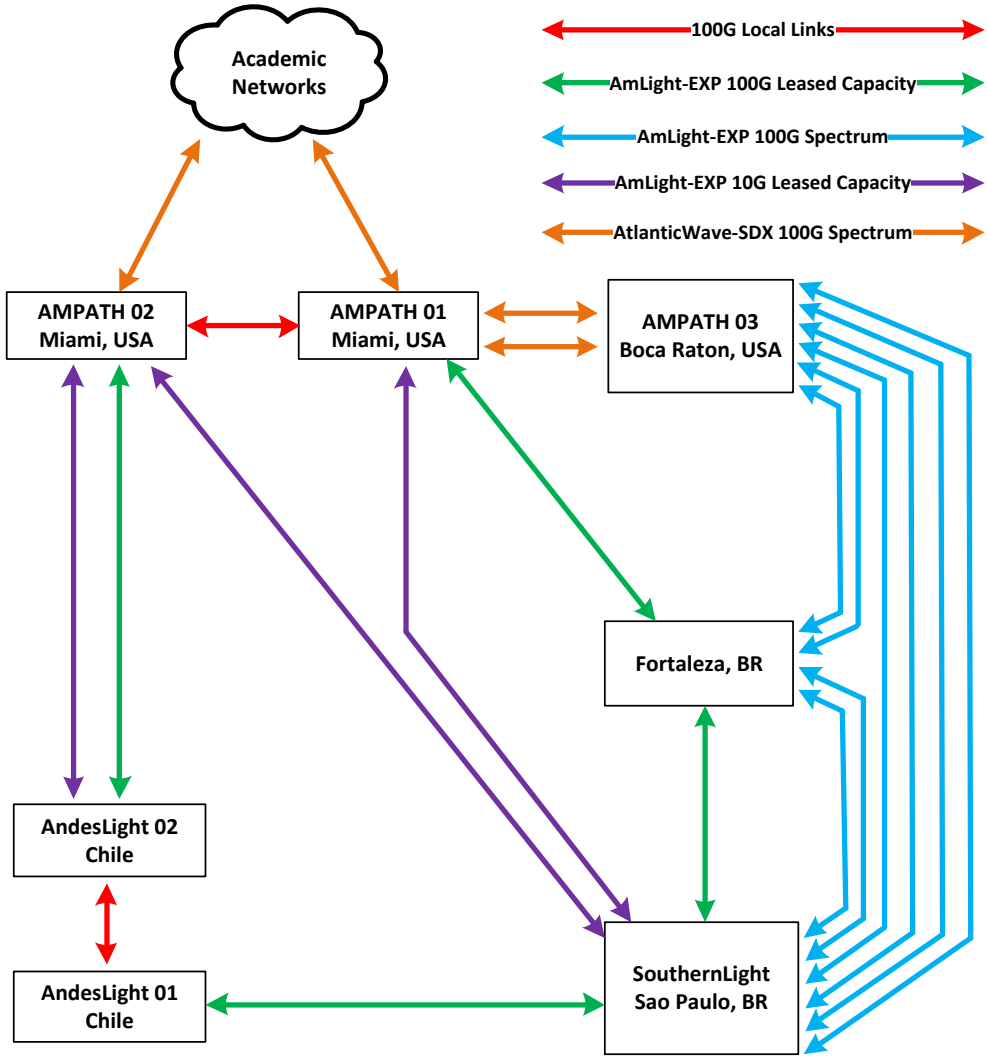

AmLight ExP
Americas Lightpaths Express & Protect

Opportunity: Improving Network Troubleshooting

- Troubleshooting network transient events is complex and time consuming using legacy technologies (ICMP, SNMP, NetFlow, etc.)
  - Transient events refer to network anomalies that are difficult to trace, but cause disruption, such as micro bursts, intermittent congestion, etc.
- Tracking transient events in real time is complex and costly, because
  - Monitoring flows in real time overloads network resources
  - Polling SNMP or OpenFlow counters is costly in sub 30 second intervals, because of CPU consumption
- LSST promises to change the paradigm of how network troubleshooting will be conducted
  - LSST data transfer must complete in less than 5 seconds
  - A 0.001% of packet loss will compromise the LSST application
  - Packet loss isolation will have to be handled in real-time
- Network operators need new tools to monitor flows in real time without impacting network performance

Opportunity: In-band Network Telemetry (INT)

- AmLight-INT is a new project, funded by NSF, that will include telemetry functions in the AmLight network links
- In-band Network Telemetry (INT) 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
- INT enables new network troubleshooting functionality:
  - (1) Network troubleshooting Micro-burst detection, trajectory sampling;
  - (2) Advanced congestion control;
  - (3) Advanced routing, for instance, utilization-aware routing; and
  - (4) Network dataplane verification
- AmLight-INT will respond to the LSST SLA requirements by enabling real-time network troubleshooting functions

Challenge: LSST Use Case

- Inter-domain coordination challenge example
- Hybrid networks, multiple network paths, multiple vendors equipment, multiple network operators
- Strict SLA:
  - MTBF (180 days in a year)
  - MTTR (48 hours)
  - 5 sec 13GB image transfer time, every 27 seconds

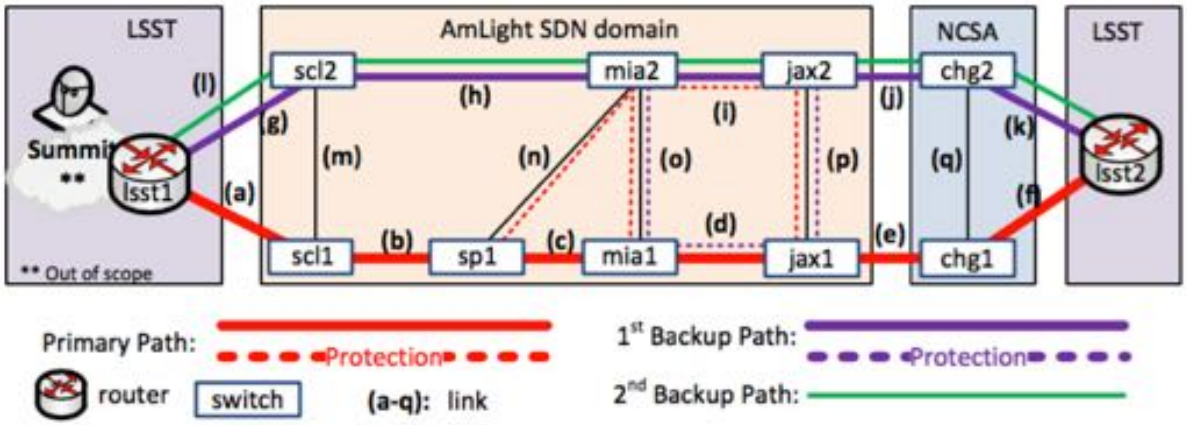

AmLight ExP
Americas Lightpaths Express & Protect

~~Monitoring and Measurement~~

- Each AmLight PoP has a 10G perfSonar node with two NICs (BWCTL and OWAMP)
- Two 100G network taps installed in Miami to support the IRNC AMIS project
- One 100G network tap being installed between Sao Paulo and Santiago
- Maddash portal available

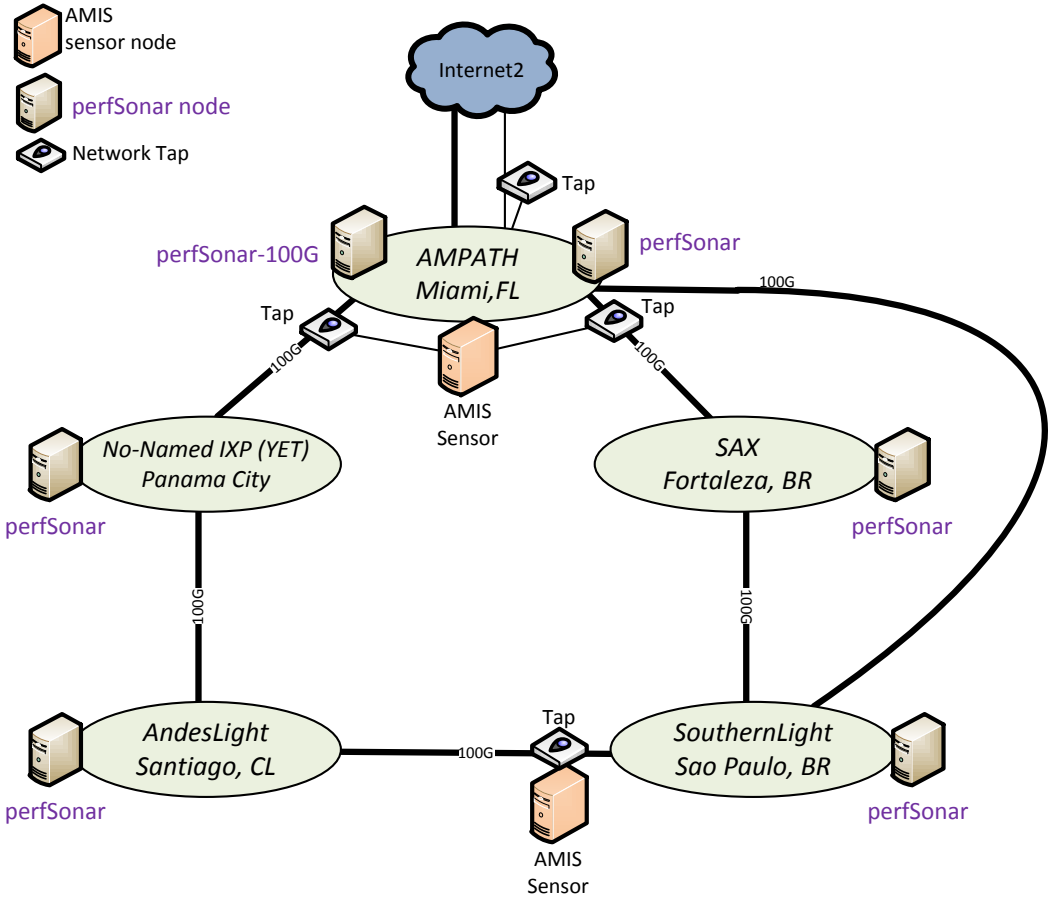

AmLight ExP
Americas Lightpaths Express & Protect

What is In-band Network Telemetry (INT)?

- INT 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
    - Buffers, counters, What else?