AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Innovating the Network for
Data Intensive Science (INDIS)


# Relevance of Software Defined Exchange Points in the Interdomain Path of Scientific Workflows

*INDIS, November 17, 2019*


**Julio Ibarra, PhD**
Florida International University

# Network connections & capacity in the Atlantic is increasing

**AmLight Express:**
- 6x100G links
- Boca-Raton-Miami, Fortaleza, Sao Paulo
- Operational

**AmLight-SACS:**
- 1x100G link
- Miami, Fortaleza, Cape Town
- RFS Dec. 2019

**Bella-S:**
- 1500GHz acquisition
- RFS 2020
- 2x100G waves initially

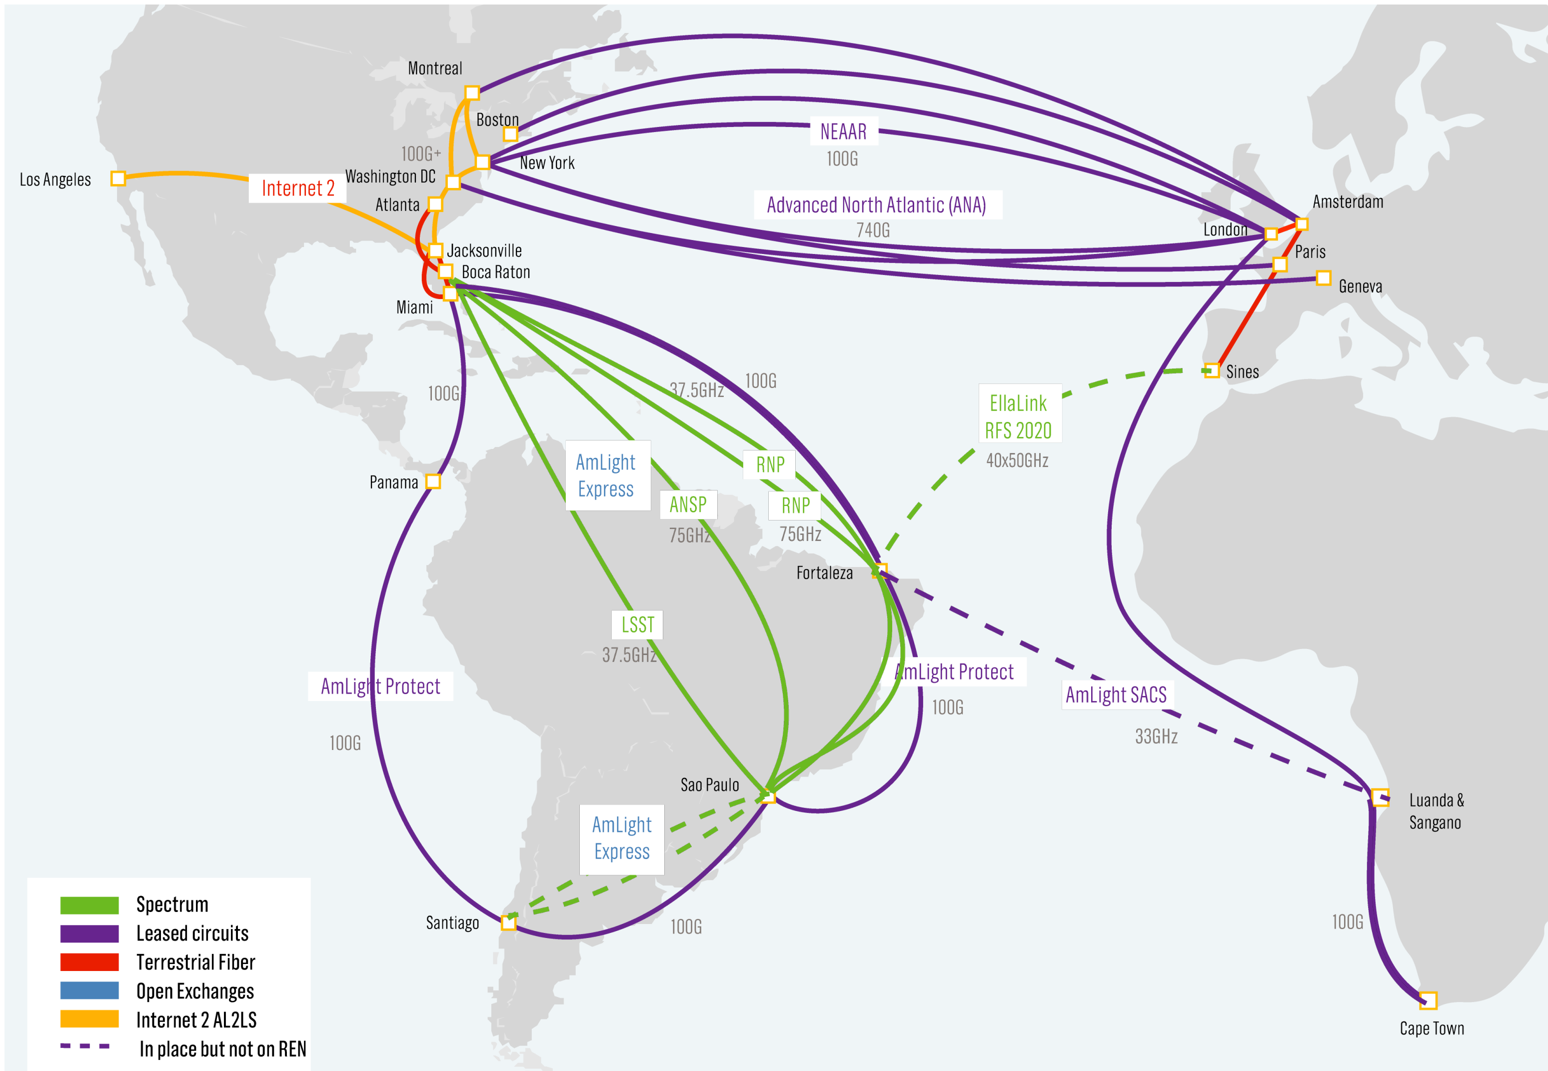




AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Open Exchange Points (OXPs) are a core component of the international connections

There are 11 OXPs connecting circuits across the Atlantic

ANA:

- MOXY
- MANLAN
- WIX
- NetherLight
- Paris OXP
- London OXP

AmLight:

- AMPATH
- SouthernLight
- SAX
- AndesLight

AmLight-SACS:

- AMPATH
- SAX
- ZAOXI

Bella-S:

- SAX
- TBD

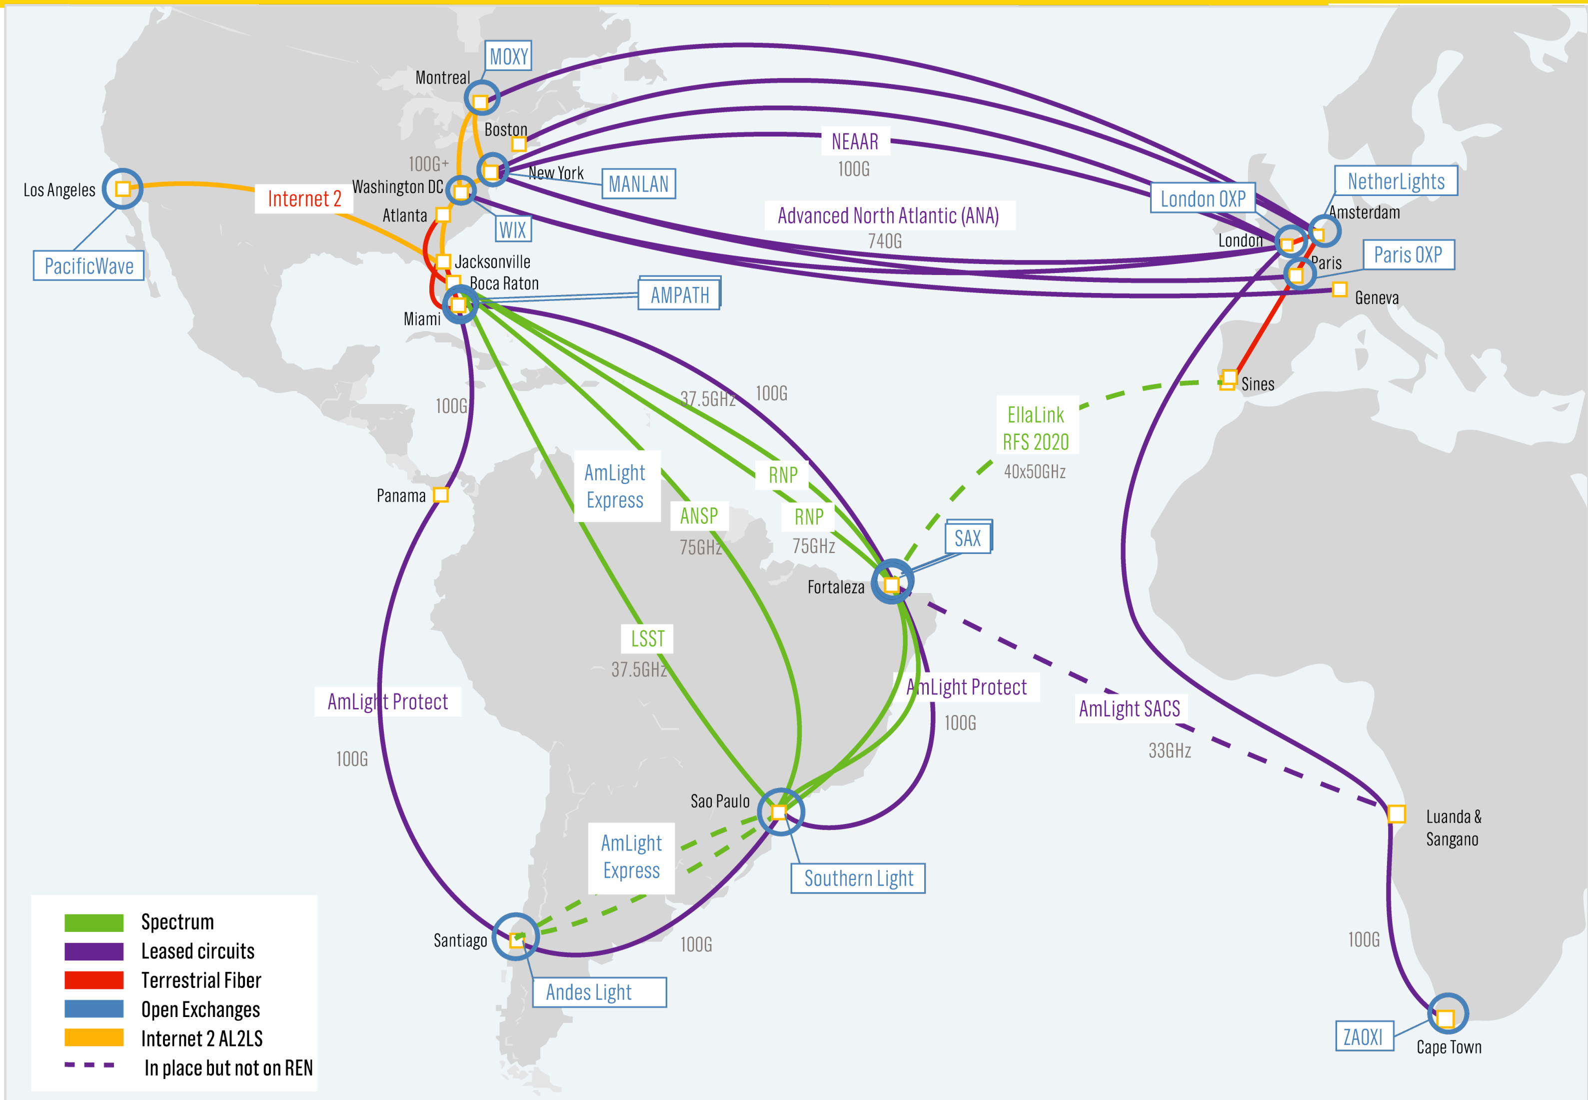






AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Science applications are dependent on international research network services

- Distributed High Throughput Computing applications
  - Latency sensitive
  - Open Science Grid

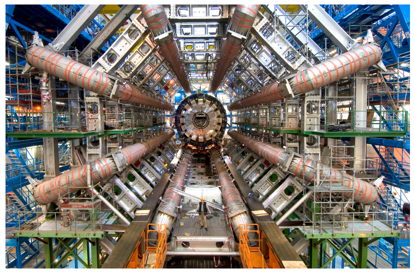

- Real-Time, high-throughput, high-resilience applications
  - Strict SLA
  - LSST

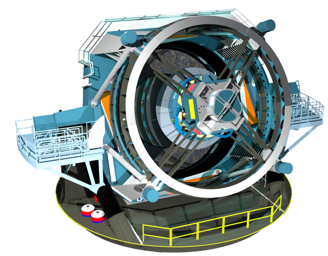

- International Research Testbeds
  - At-scale experimentation
  - FABRIC
  - AmLight-ExP
  - Others


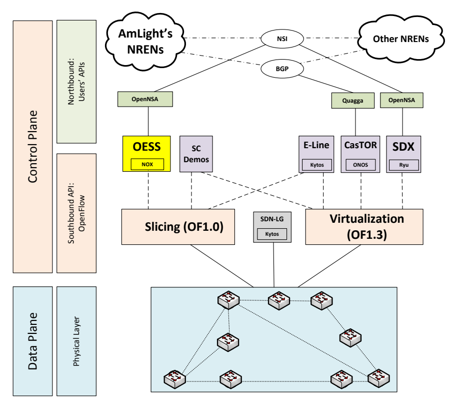




AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Impact to network and OXP operators?

- **Interdomain Provisioning** effort increases
  - # Circuits, DTNs, clusters, and types of devices are increasing
  - High degree of **coordination** and **communication**
    - Engaging NOCs, opening tickets, many conference calls and emails [1]
    - In the order of days or sometimes weeks [1, 3]

- **Building Resiliency** is essential, but it has a price [2, 4]
  - Pre-defined static paths
  - Consumes resources
  - Adds complexity

- **Troubleshooting more interdomain paths is messy**
  - Lack of visibility across administrative domains makes it difficult to know where to start
  - Tools for troubleshooting the data plane are well known, but their application can be ad hoc

# Multiple NRENs, OXP, and WAN end-to-end example

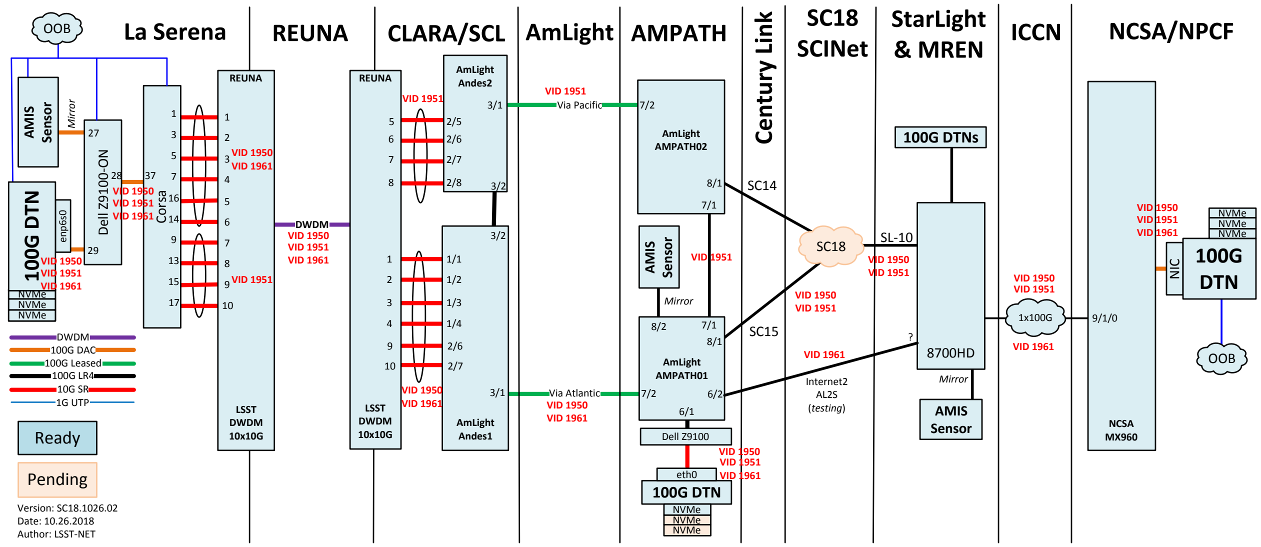


- Provisioning, resiliency and troubleshooting activities increase dramatically as more international end-to-end paths are added
- 9 network operators and 4 OXPs built the network for the LSST NRE at SC18
- Approximately 6 months of communication and coordination



AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Relevance of Software-Defined Exchanges (SDX) in the end-to-end path: An example

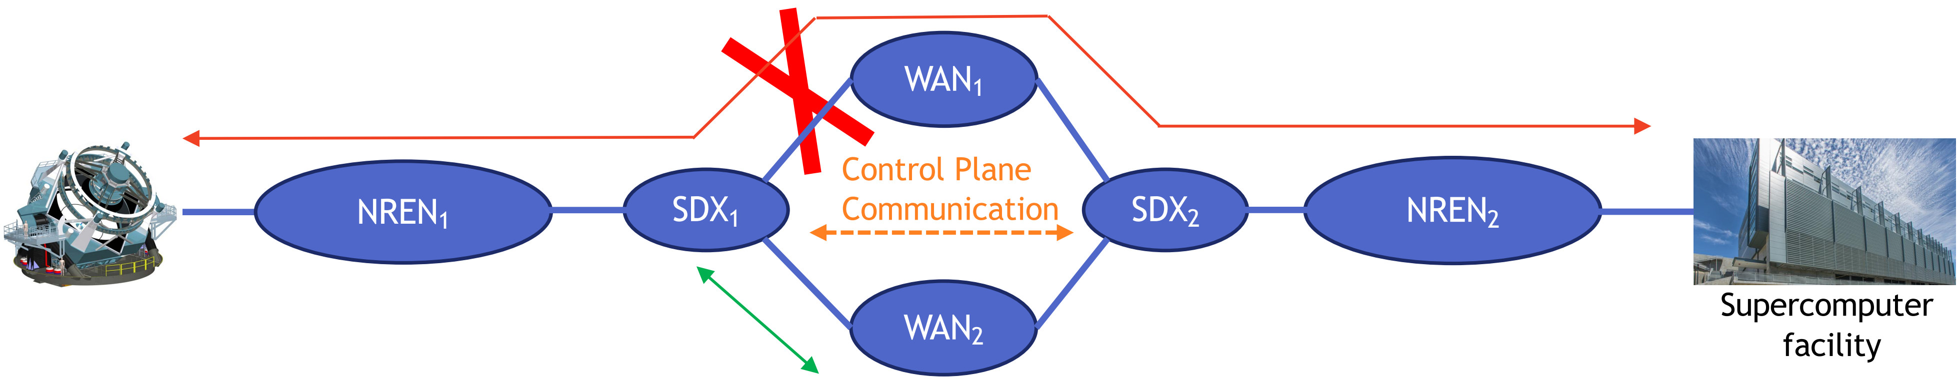


- Replace OXPs with SDXs and a Control Plane communication channel in the interdomain path
- Upon failure with connection to WAN$_1$, SDX$_1$ notifies SDX$_2$ via Control Plane Communication channel
- SDX$_1$ requests use of secondary path
  - CP layer computes a new path, then propagates rules to SDX$_1$ and SDX$_2$
  - Traffic is then dynamically rerouted across WAN$_2$
- End points can save time and money
  - No longer involved in alternate path selection
  - No longer require expensive routers with protocols for link continuity
  - Can save network resources (e.g. VLANs)

# What is a Software-Defined Exchange (SDX)?

- A SDX refers to a meet-me point where independent-neutral administrative domains can exchange computing, storage and networking resources [4, 11]
- A SDX can support network aware applications to achieve end-to-end programmability and control [10]

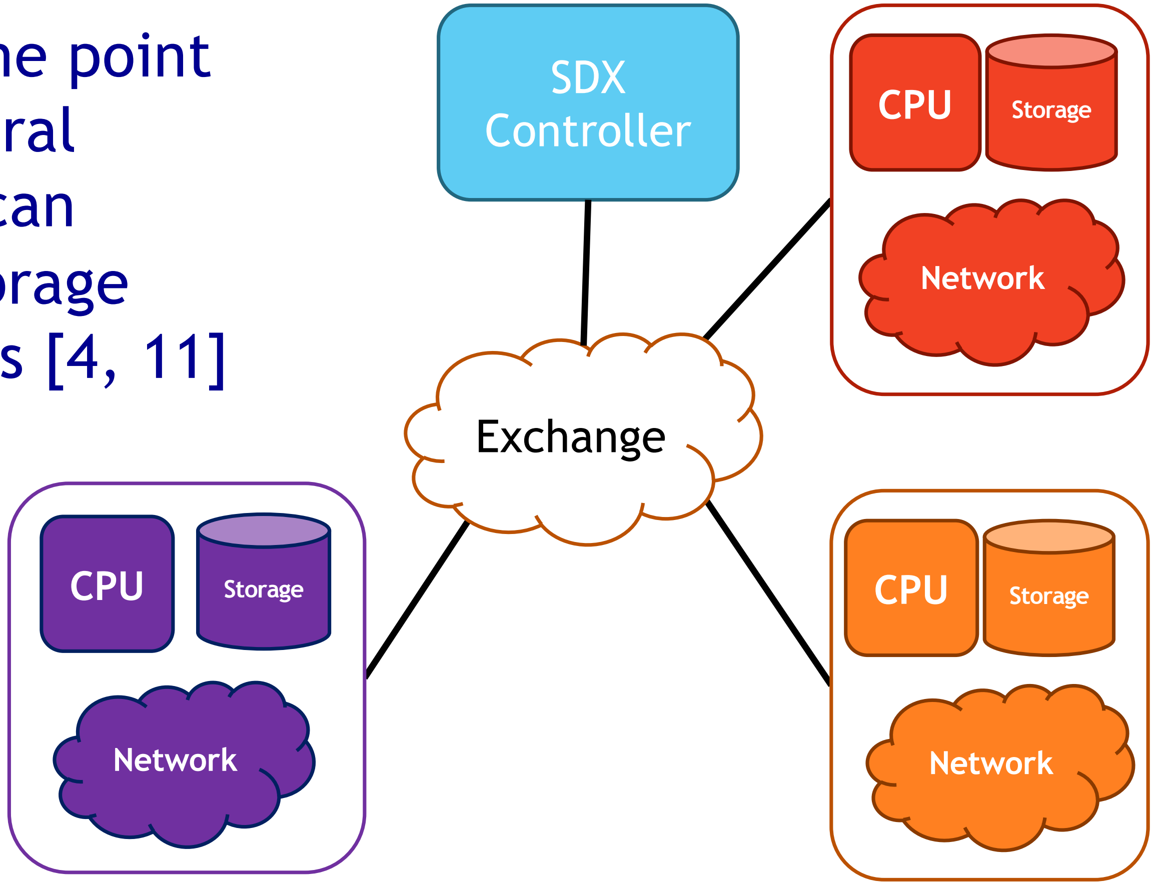

# Active SDX Projects

- NSF’s International Research Network Connections (IRNC) program funded three SDX projects to
  - Conduct research, development, and experimentation of multi-domain SDN networks
  - Interconnect SDN peers both nationally and internationally
  - Participate and play a leadership role in the planning, instantiation, coordination, and prototyping support of global-scale SDN exchanges
- StarLight SDX: A Software-Defined Exchange for Global Science Research and Experimentation, Award# OAC-1450871. PI: Joe Mambretti.
- PacificWave Expansion Supporting SDX and Experimentation: Award# OAC-1451050. PI: Louis Fox.
- AtlanticWave-SDX: A Distributed Intercontinental Experimental SDX, Award# OAC-1451024. PI: Julio Ibarra.

# AtlanticWave-SDX project goals

- To enable domain scientists to reserve network resources through a multi-domain SDX by
  - Simplifying the interface for domain scientists to request network resources
  - Providing interfaces to program the forwarding plane to respond to application requirements

- To build a distributed SDX between the U.S. and S. America
  - To support a dramatic increase in south-north science flows
  - To integrate the SDN infrastructures at AMPATH, SoX, SouthernLight, and AndesLight open exchange points

# AtlanticWave-SDX leverages the AmLight Network

- Express Ring: Boca Raton-Miami, Fortaleza, Sao Paulo
  - 6 (green lines) x 100G links
- 100G Protect Ring: Miami-Fortaleza, Fortaleza-Sao Paulo, Sao Paulo-Santiago, Santiago-Panama, and Panama-Miami (solid orange)
- 10G ring from Miami-Sao Paulo-Miami for protection (red dashed)
- 10G Miami-Santiago for protection (orange dashed)
- Open Exchange Points in Miami, Fortaleza, Sao Paulo, and Santiago

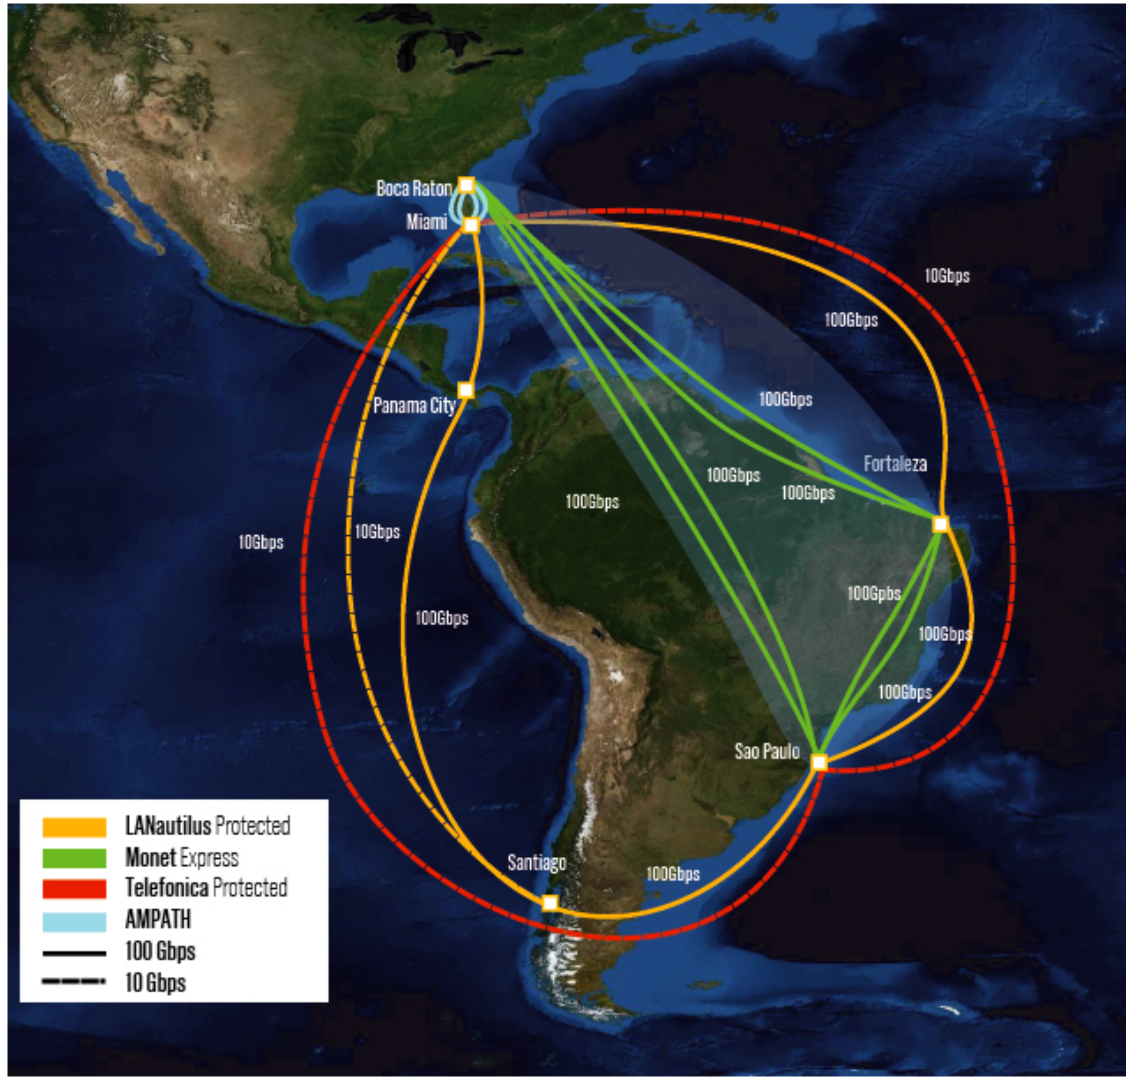




AtlanticWave SDX INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Partners and Collaborators

- Florida International University (FIU)
- University of Southern California Information Sciences Institute (USC-ISI)
- Georgia Institute of Technology (GT)
- Renaissance Computing Institute at UNC (RENCI)
- Academic Network of Sao Paulo (ANSP)
- Association of Universities for Research in Astronomy (AURA)
- Rede Nacional de Ensino e Pesquisa (RNP, Brazil)
- Red Universitaria Nacional (REUNA, Chile)
- Florida LambdaRail
- Internet2

# AW-SDX Implementation

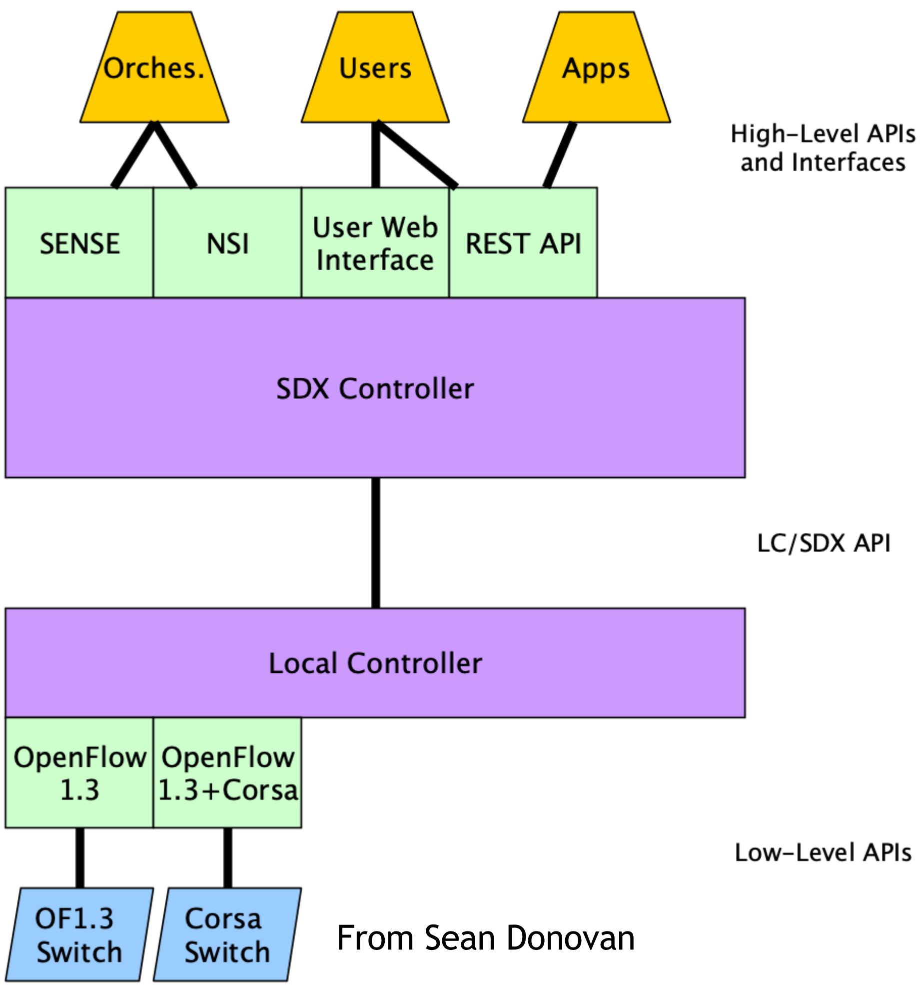


From Sean Donovan

## SDX Controller Responsibilities

- Global view
  - Implements specific high- to mid-level rules
- Provides Northbound Interfaces
- Breaks down User-level policies to Local Controller (LC) rules

## Local Controller Responsibilities

- Converts LC rules to resource-specific protocols
- Relays status from resource to SDX Controller



AtlanticWave SDX INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Example: Real-Time high-throughput, high-resilience application

- 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 image
- Each image must be transferred to the archive site at NCSA in Champaign, Illinois, within 5 seconds, inside the 27 seconds window
- Multi-traffic types with different priorities (db sync, control, general Internet traffic) must also be supported

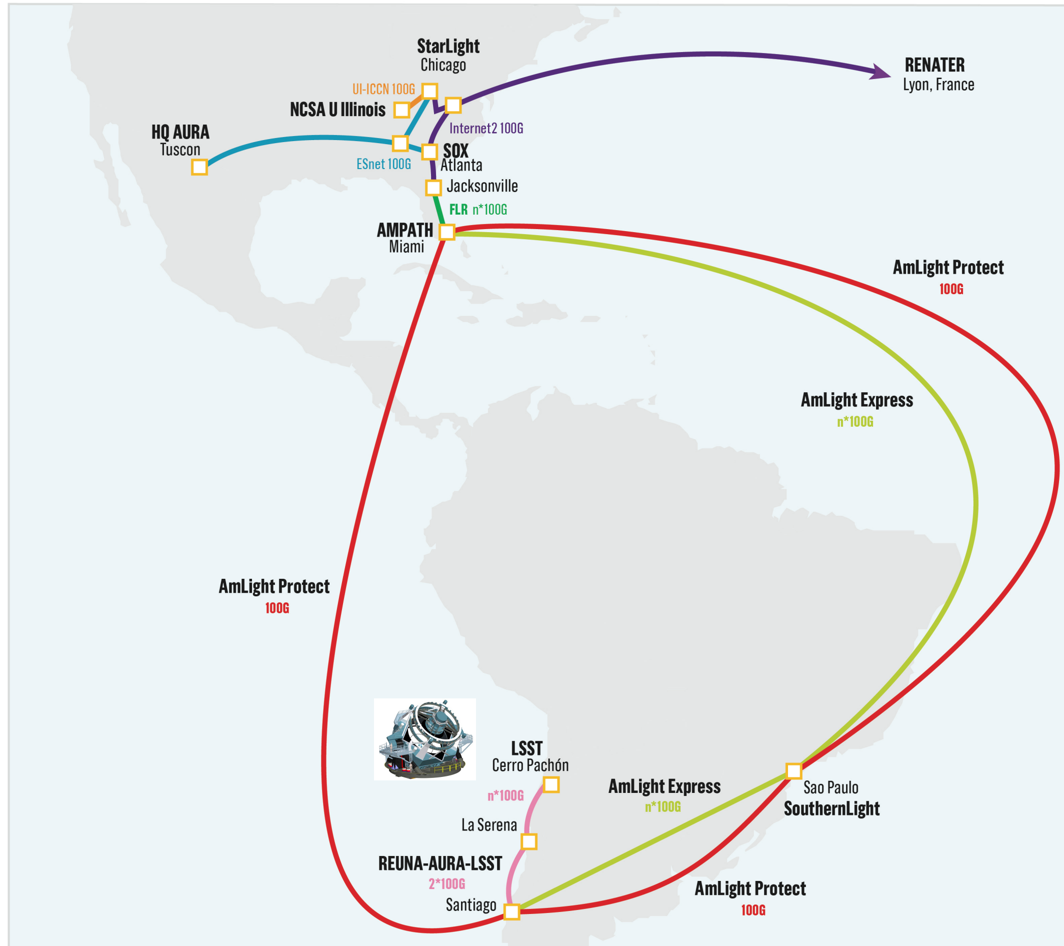




AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Example: SDX for LSST on AmLight-ExP

- AW-SDX LC running at 4 exchange points on the AmLight-ExP network, SDX Controller managing the CP channels to each SDX LC
- Request arrives to transfer an image from DTN1 and DTN2
- Programmed with LSST application requirements, the SDX Controller finds a path on the Express (green) network
  - SDX Controller deploys mid-level LC-Rules to the Local Controllers (LCx)
  - LCs translate LC-Rules to low-level (OpenFlow 1.3) rules
  - OF1.3 rules installed into each switch in the forwarding plane

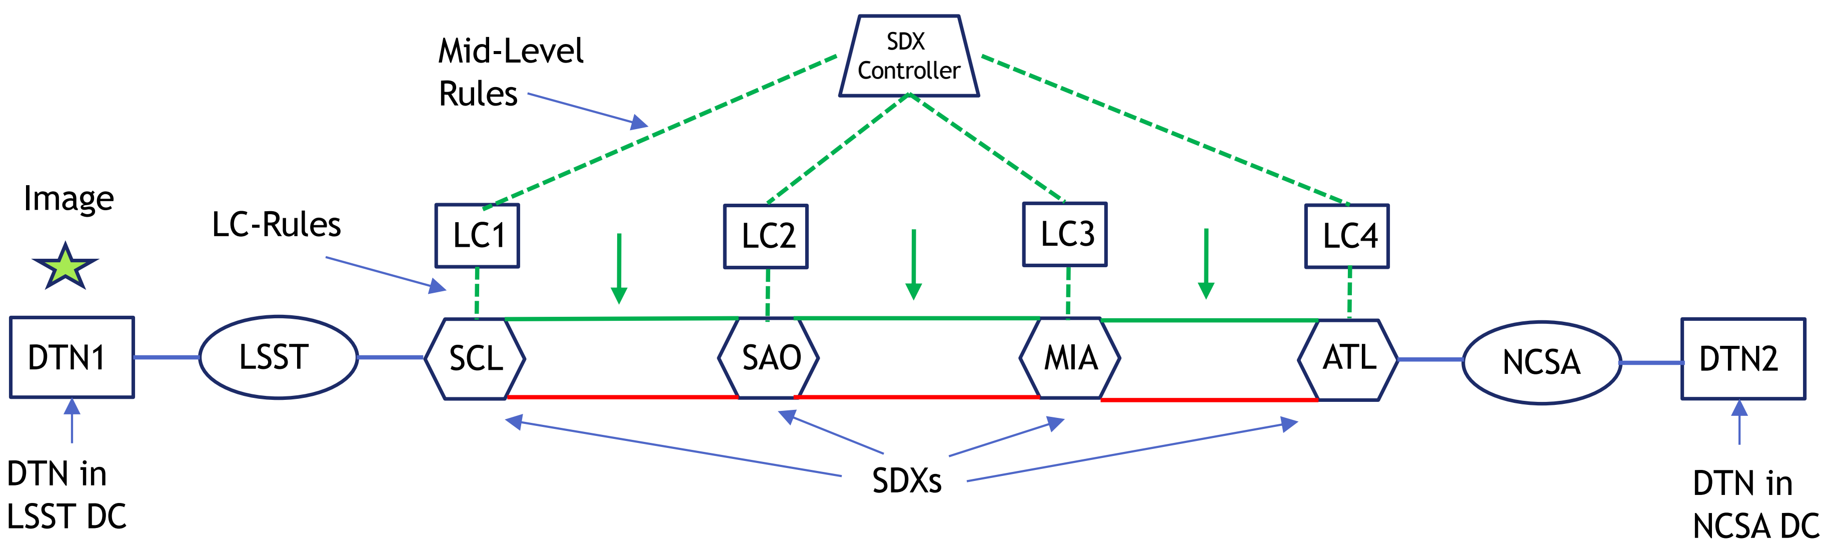

# Example: SDX for LSST on AmLight-ExP

- Local Controllers (LCx) are programmed with LC-Rules and Mid-level Rules for the Express network
- SDX Controller has a full topology from SCL to ATL
- Image is in transit from DTN1 to DTN2

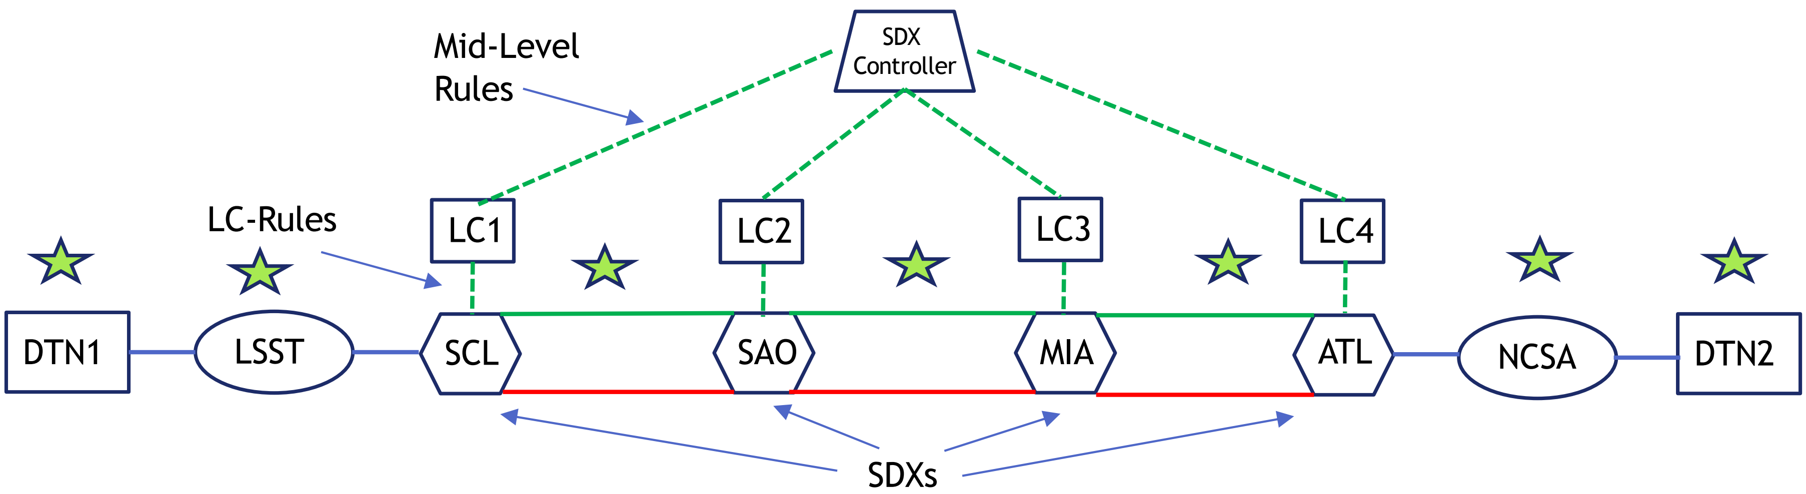

# Example: SDX for LSST on AmLight-ExP

- Fiber cut on the Express path between Santiago and Sao Paulo
  - Forwarding plane notifies LC1 and LC2 of port status change
  - LC1 and LC2 notify the SDX Controller of the event

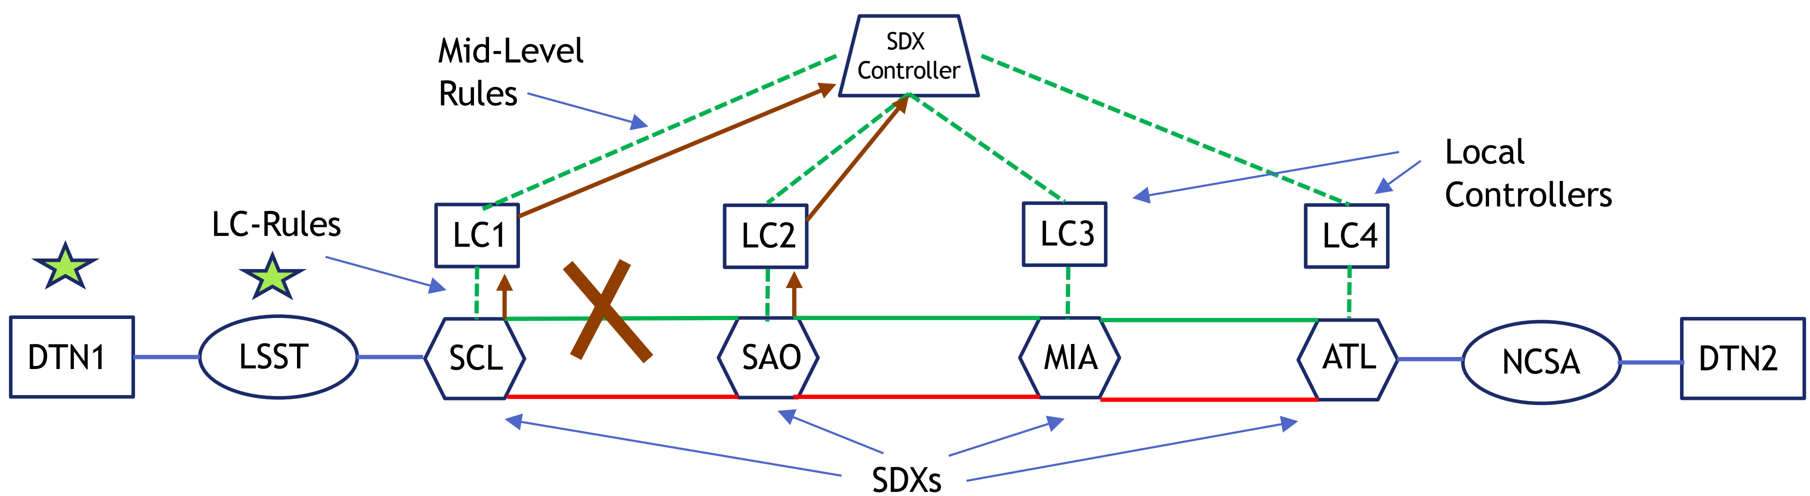

# Example: SDX for LSST on AmLight-ExP

- SDX Controller evaluates application requirements for LSST
- SDX Controller computes a new path in response to requirements
  - Deploys Mid-level, LC-rules, and OF 1.3 rules to use Protect path (red) from SCL to SAO
  - Remaining path from SAO to ATL remains unchanged
- End points were not involved in the alternate path selection

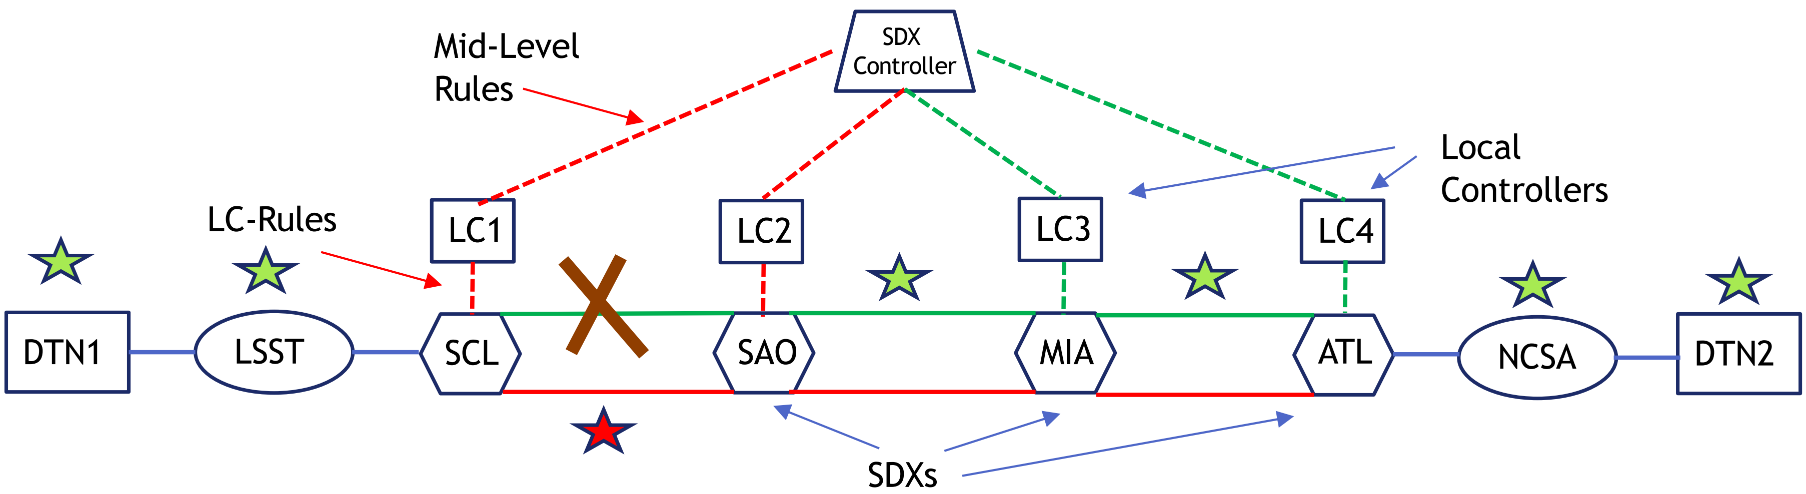

# Next Steps and Future Work

- Deploy the AtlanticWave-SDX on the AmLight-ExP network
- Continue developing SDX to better support requirements of dHTC, HA science applications, and international testbeds
  - Programming application requirements into the SDX
- Improving resiliency through traffic engineering at the SDX using programmability to steer traffic via alternate paths
- Facilitating data plane troubleshooting effort by providing counters and measurement data across domains
  - Exploring In-band Network Telemetry (INT) to provide more visibility of science flows in the data plane

# References


1. N. Charbonneau, V.M. Vokkarane, C. Guok, I. Monga, Advance reservation frameworks in hybrid IP-WDM networks, IEEE Commun. Mag. 49 (5) (2011) 132–139.
2. I. Monga, C. Guok, W.E. Johnston, B. Tierney, Hybrid networks: Lessons learned and future challenges based on ESnet4 experience, IEEE Commun. Mag. 49 (5) (2011) 114–121, http://dx.doi.org/10.1109/MCOM.2011.5762807.
3. Tepsuporn, Scott, et al. "A multi-domain SDN for dynamic layer-2 path service." Proceedings of the Fifth International Workshop on Network-Aware Data Management. ACM, 2015.
4. Chung, Joaquin, et al. "Orchestrating intercontinental advance reservations with software-defined exchanges." Future Generation Computer Systems 95 (2019): 534-547.
5. Gerola, Matteo, et al. "Icona: Inter cluster onos network application." Proceedings of the 2015 1st IEEE Conference on Network Softwarization (NetSoft). IEEE, 2015.
6. S. H. Yeganeh, A. Tootoonchian, and Y. Ganjali, “On scalability of software-defined networking,” in IEEE Communications Magazine 51 (2), pp. 136-141, 2013.
7. Agarwal, Kanak, et al. "SDN traceroute: Tracing SDN forwarding without changing network behavior." Proceedings of the third workshop on Hot topics in software defined networking. ACM, 2014.
8. Bezerra, J., Galiza, H., Ibarra, J. and Schwarz, M., “AmLight’s OpenFlow Sniffer dissected: Troubleshooting production networks”, Brazilian Symposium on Computer Networks and Distributed Systems (SBRC) Workshop on Future Internet Research and Experimentation (WPEIF),pages 33-36, Salvador Bahia Brazil, June 2016.
9. Bezerra, J., Ibarra, J., Schwarz, M., Freitas, H., Morgan, H., “Mitigating the Risks of Supporting Multiple Control Plans in a Production SDN Network: A Use Case”, In Anais do SBRC/WPEIF 2017, Belem, PA, Brazil, (p.34-37).
10. Chung, Joaquín, et al. "Atlanticwave-sdx: An international sdx to support science data applications." Software Defined Networking (SDN) for Scientific Networking Workshop, SC'15. 2015.
11. Report of the NSF workshop on software defined infrastructures and software defined exchanges, Washington, DC, 2016.

# Acknowledgment

- AtlanticWave-SDX Team Members
  - Jeronimo Bezerra, Vasilka Chergarova, Sean Donovan, Joaquin Chung, others
- Come see the AtlanticWave-SDX demo
  - Georgia Tech booth #1809, Tuesday, at 3PM
  - Caltech booth #543, Wednesday, at 10AM

- Harvey Newman and Team at the Caltech booth 543
  - Many years of collaboration in advanced network experiments to South America using the AmLight network
  - Global Petascale to Exascale Workflows for Data-Intensive Science, Tuesday, at 2PM



AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Thank You!

- NSF AmLight-ExP, AtlanticWave-SDX, AmLight-SACS,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



AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

# Questions?


NSF

FIU FLORIDA INTERNATIONAL UNIVERSITY

Georgia Tech

USC University of Southern California

renci

AtlanticWave SDX INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE