AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

AtlanticWave-SDX

NSF #OAC-2029278

Americas Research Platform (AmRP) Working Group Meeting September 24, 2021

Julio Ibarra, FIU, PI
Jeronimo Bezerra, FIU, Co-PI
Heidi Morgan, USC-ISI, Co-PI
Yufeng Xin, RENCI, Co-PI
Lisandro Granville, UFRGS, Co-PI

Outline

- Background and examples of OXPs
- AtlanticWave-SDX 2.0 project
- Use Case

What is an Open Exchange Point (OXP)?

- An *eXchange Point* refers to a facility where networks meet to peer and exchange their traffic
- In the IP world, eXchange Points (IXPs) are Ethernet domains where networks exchange traffic based on destination MAC adresses
- An **Open eXchange Point (OXP)** usually refers to academic IXPs where networks exchange traffic without complex user policies or traffic shapping
- AMPATH, SouthernLight, SAX, ZAOXI, and SOX are examples of OXPs
- An OXP typically offers services: Access and Aggregation, Compute and Storage, Monitoring and Measurement, Peering, etc.

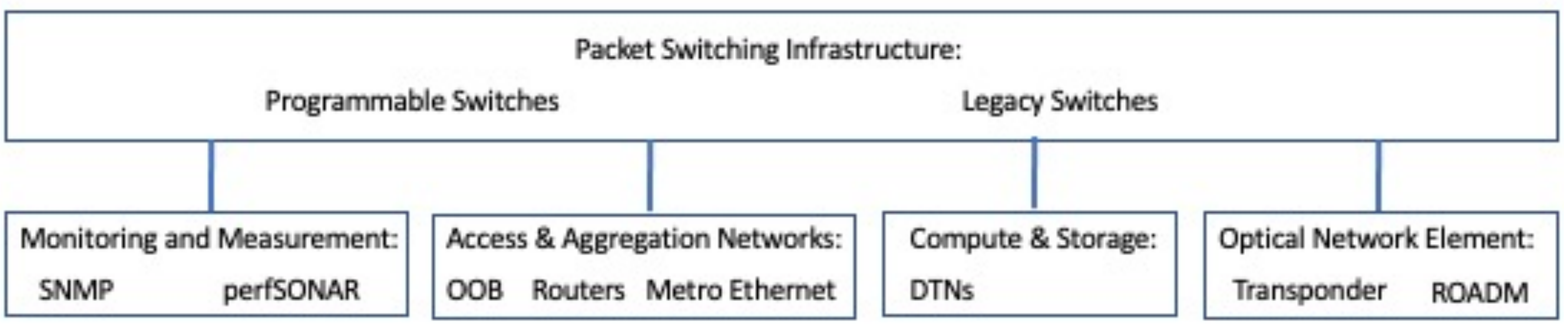

Characteristics of a Software-Defined Exchange

- A Software-Defined Exchange (SDX) builds upon an OXP by adding Software-Defined Networking (SDN) functions on compute, storage and networking resources
- SDX enhances the OXP forwarding capabilities by
  - enabling OXP users to create complex traffic engineering policies for both incoming and outgoing traffic
- SDX can be particularly useful for R&E Networks where multi-domain services are very popular, especially L2VPNs
- The AtlanticWave-SDX project was created to enhance the provisioning functions of the AtlanticWave distributed OXPS

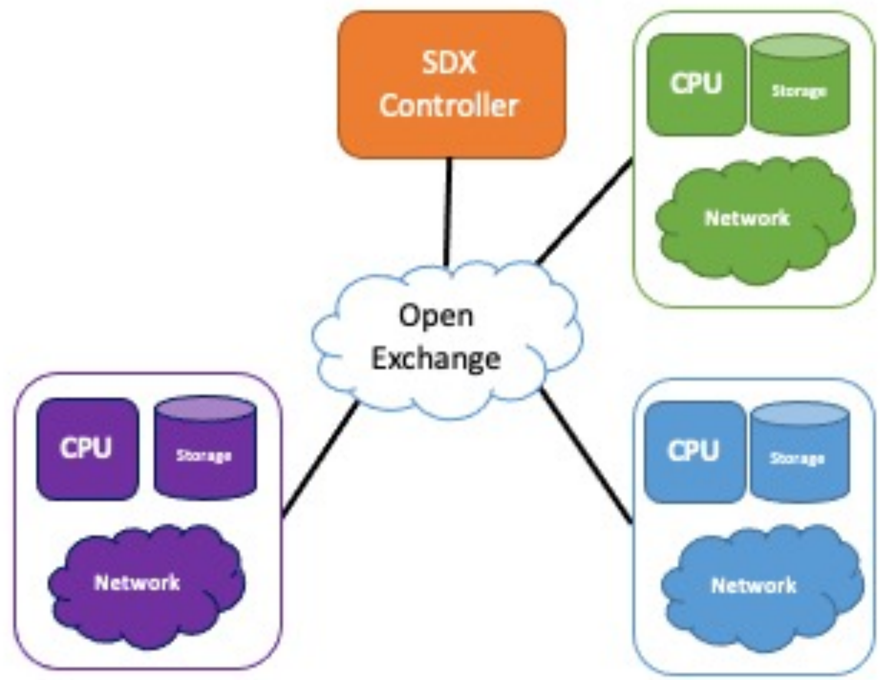

OXP: AMPATH

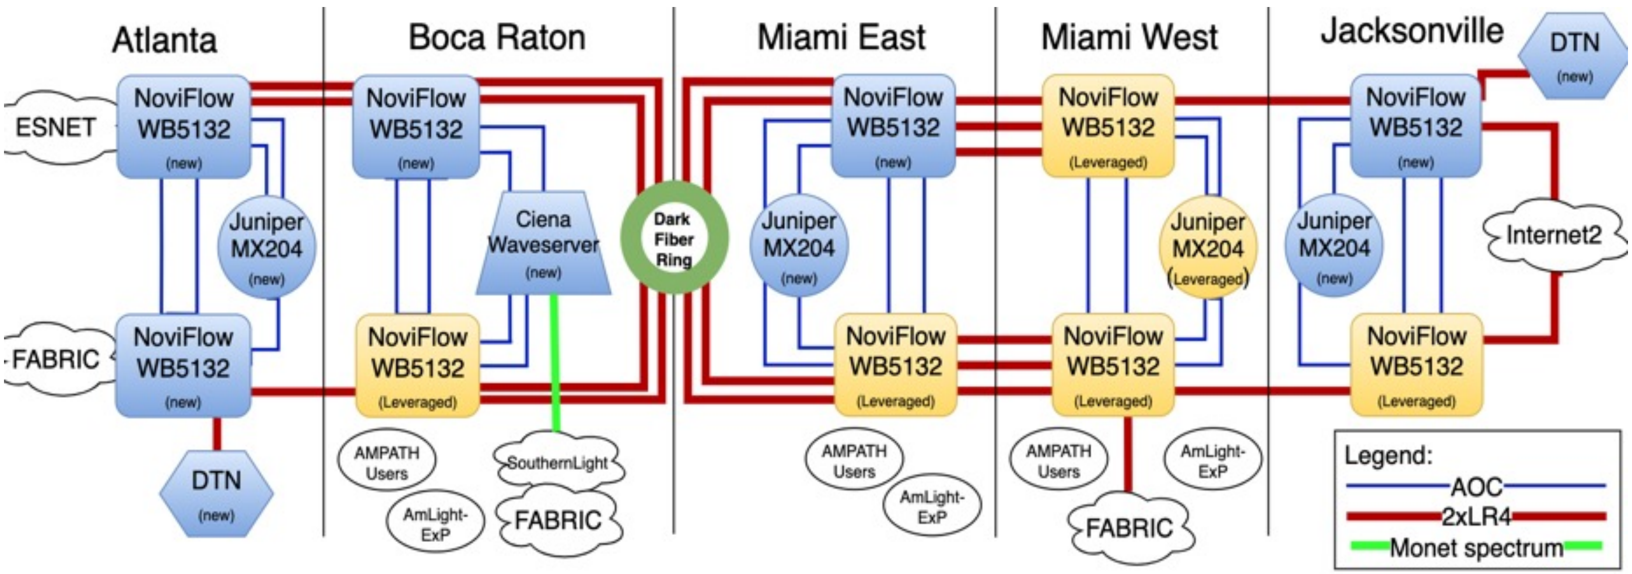

- Distributed OXP with colocation in Miami, Boca Raton, Jacksonville and Atlanta (soon)
- Access ports from 10G and 100G
- **Switching Fabric**: L2VPNs using SDN controllers with user interfaces that enable users to provision their own circuits
- **Routing Services**: IPv4, IPv6, Multicast IPv4, and VRF services. Networks can peer with AMPATH using MP-BGP
  - Mutually Agreed Norms for Routing Security (MANRS) best practices, Resource Public Key Infrastructure (RPKI)
- **Compute and storage**: One 100G DTN with NVMe cards available to the community
- **Monitoring and Measurement services**: Dedicated 10G perfSonar nodes at every site. In-band Network Telemetry (INT) enabled at every colocation site

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

OXP: AmLight

- Distributed OXP connecting the U.S. to South America and Africa
- OXPs: (3) Florida: Miami, Boca Raton, Jacksonville; (2) Brazil: São Paulo, Fortaleza; (1) Chile: Santiago, (1) Panama: Panama City, (1) Puerto Rico: San Juan, and (1) South Africa: Cape Town
- Carries academic and commercial traffic
- Operates as a production SDN network
- Capacity between the U.S. and S. America: **600G**
- Capacity between Africa, U.S. and Brazil: **100G**

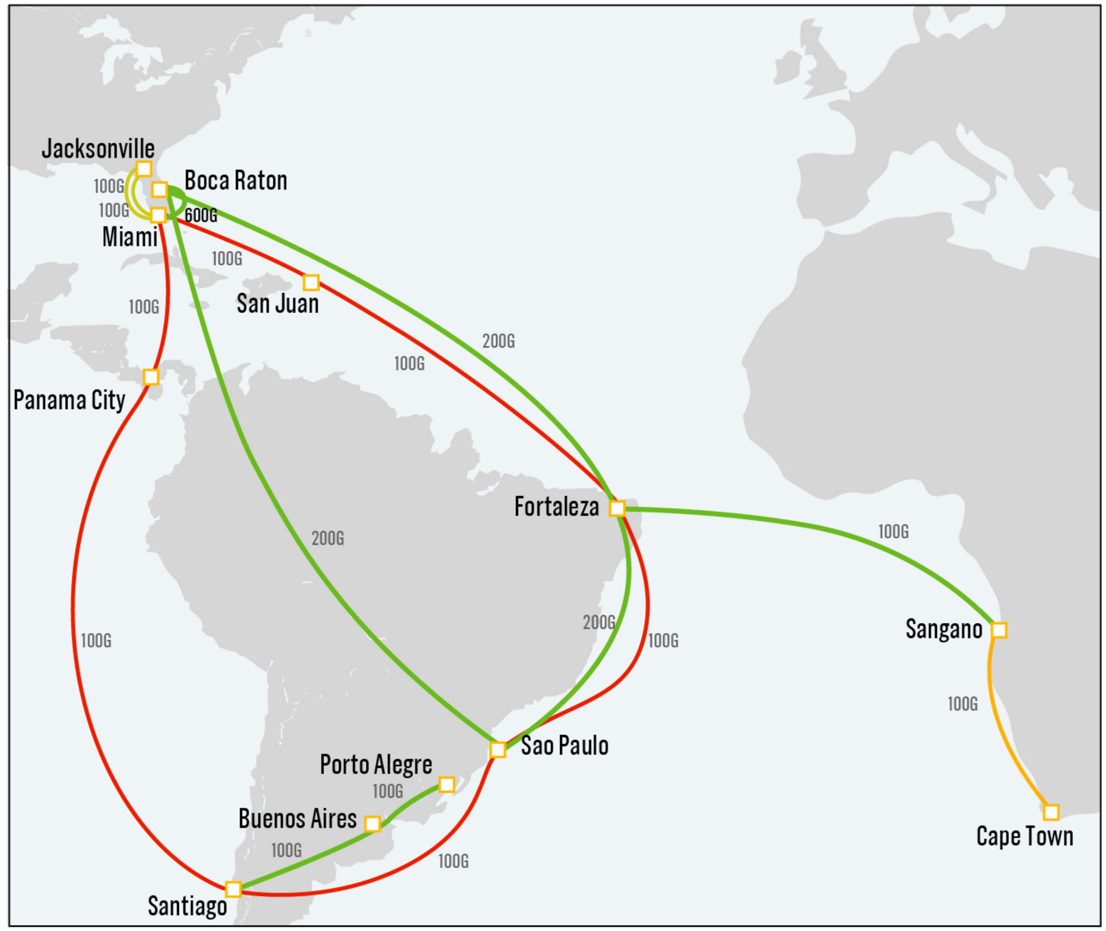

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

OXP: South America Exchange (SAX)

- South America eXchange (SAX)
- Based in Fortaleza, Brazil
- Operated by RNP
- Connects two cable landing stations and RNP’s points of presence
- SAX connects to the following submarine cable systems:
  - SAC (South Atlantic Crossing): Latin America, 100Gbps
  - SACS (South Atlantic Cable System): Africa, 100Gbps
  - Monet: USA, 200Gbps
  - EllaLink: Europe, 2x100Gbps
- Carries academic and comercial traffic
- perfSONAR node with 100G capability

AtlanticWave-SDX 2.0

AtlanticWave-SDX: A Distributed Production SDX, supporting research, enhancing operations, and interoperability testing at national and international scales. NSF Award# OAC-2029278

Management & Development team: FIU, RENCI, USC-ISI, RNP, UFRGS

Goals of the AtlanticWave-SDX:

- **Core goal**: Enhance the AtlanticWave-SDX with *Autonomic Network Architecture* concepts and designs
  - Self-management, resilient, scalable, and secure
- **Network-driven goals**:
  - Autonomic operation by leveraging network telemetry
  - New network services (L3VPN and Cloud)
- **User-driven goals**:
  - Enhance user experience via CILogon
  - Enable integration with scientific workflows, including Pegasus and OSG
  - Integrate with interdomain orchestrators: SENSE/AutoGOLE and FABRIC

Architecture and Stakeholders

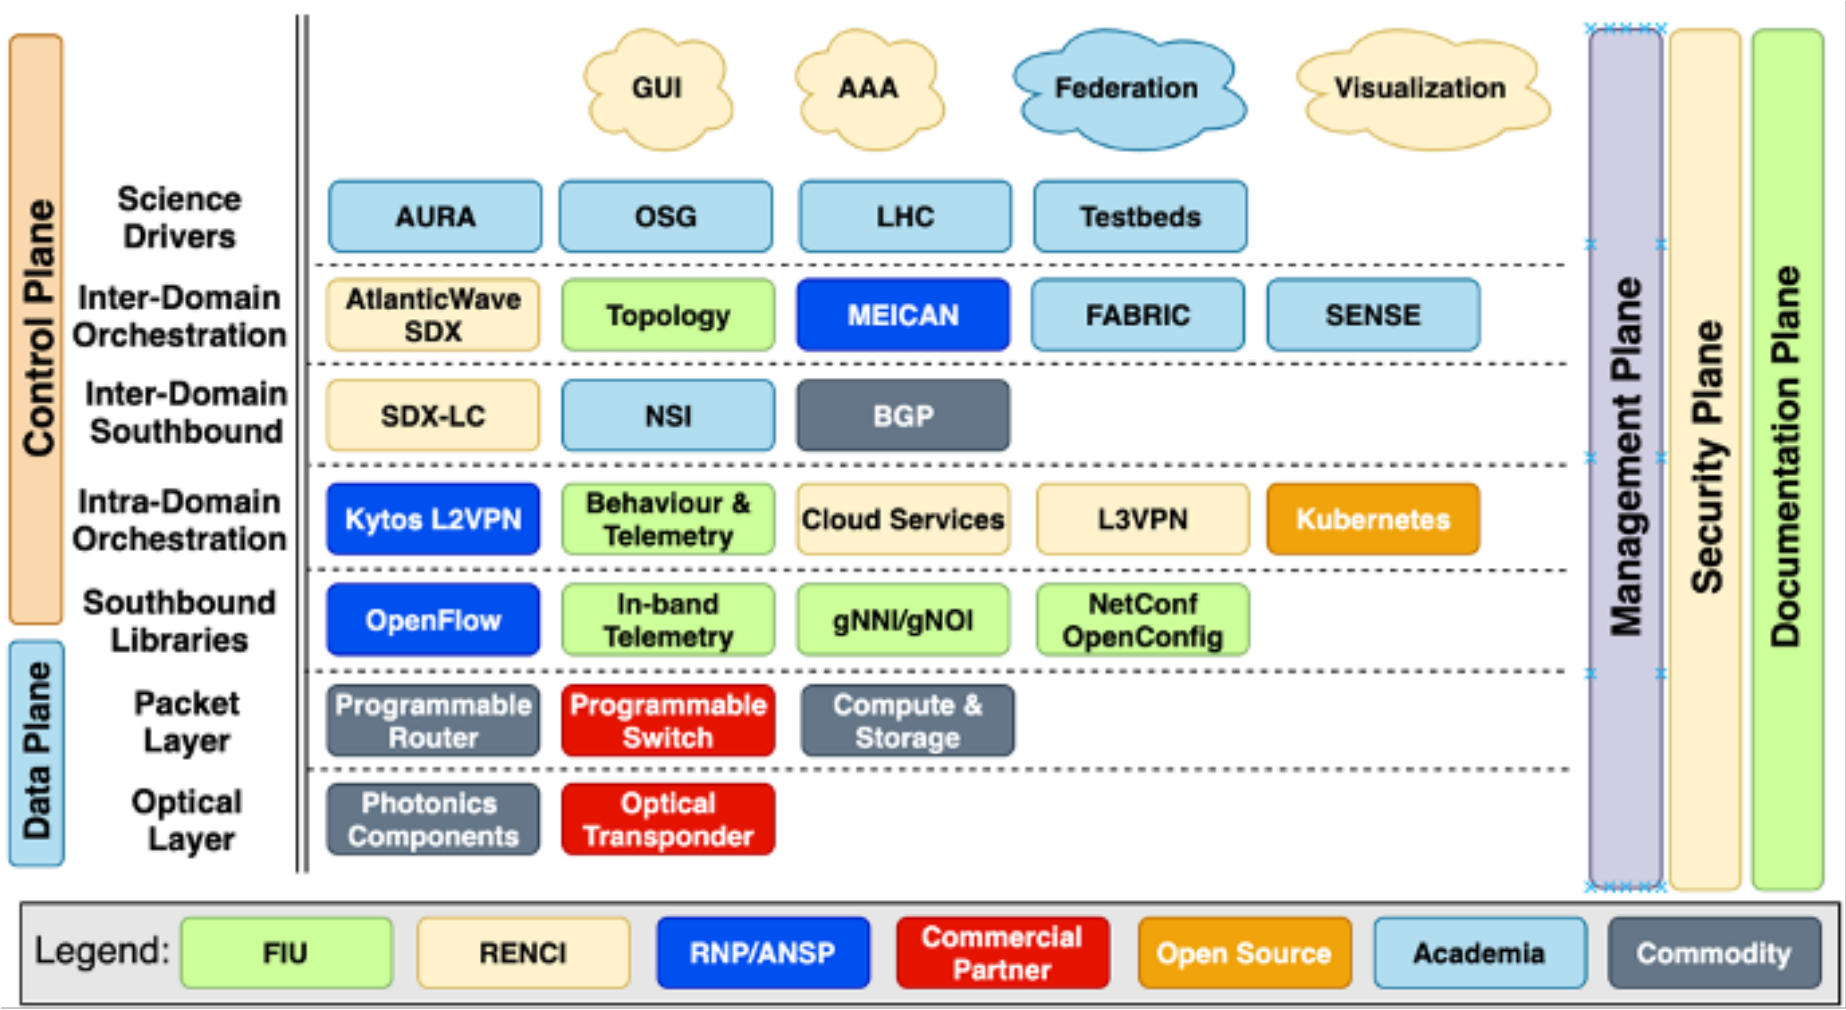

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Leveraging In-band Network Telemetry (INT)

- INT is a P4 application
  - It records network telemetry information in the packet
  - while the packet traverses a path between two points in the network
- As telemetry is exported directly from the Data Plane, Control Plane is not affected:
  - *Translating: you can track/monitor/evaluate EVERY single packet at line rate and in real time*
- Examples of telemetry information added:
  - *Time stamp, ingress port, egress port, queue buffer utilization, sequence#, and many others*
- SDX interprets telemetry data to identify and respond to network anomalies at the packet and optical layers

Out Time: 123144143 ns	
In Time: 123132143 ns	
Queue: 2	Occ: 15MB
Hop Delay: 12 us	
In: Port 1	Out: Port 2
Switch: 1	
Out Time: 124145243 ns	
In Time: 124144143 ns	
Queue: 0	Occ: 10KB
Hop Delay: 1.1 us	
In: Port 1	Out: Port 4
Switch: 2	
Out Time: 125146343 ns	
In Time: 125145243 ns	
Queue: 0	Occ: 10KB
Hop Delay: 1.1 us	
In: Port 31	Out: Port 28
Switch: 3	
Out Time: 126147443 ns	
In Time: 126146343 ns	
Queue: 0	Occ: 10KB
Hop Delay: 1.1 us	
In: Port 12	Out: Port 13
Switch: 4	
Out Time: 127187443 ns	
In Time: 127147443 ns	
Queue: 0	Occ: 21MB
Hop Delay: 40 us	
In: Port 1	Out: Port 7
Switch: 5	

INT Metadata

Adding Self-management functionality

- **Self-configuration**: Functions do not require configuration, by either an administrator or a management system. They configure themselves, based on self-knowledge.
- **Self-healing**: Autonomic functions adapt on their own to changes in the environment and heal problems automatically.
- **Self-optimizing**: Autonomic functions automatically determine ways to optimize their behavior against a set of well-defined goals.
- **Self-protection**: Autonomic functions automatically secure themselves against potential attacks.

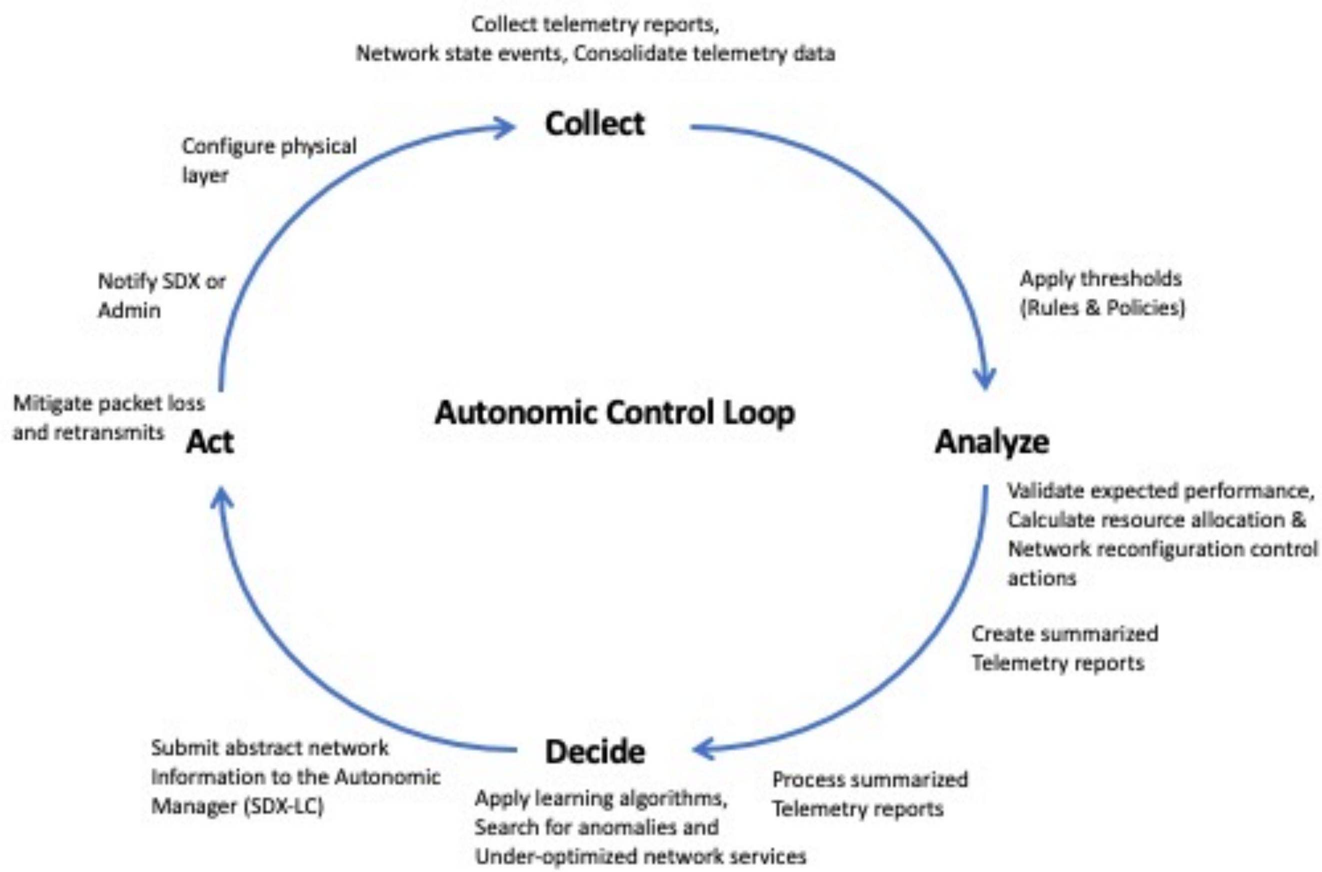

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

SDX Sub-second Closed-loop Network Orchestration

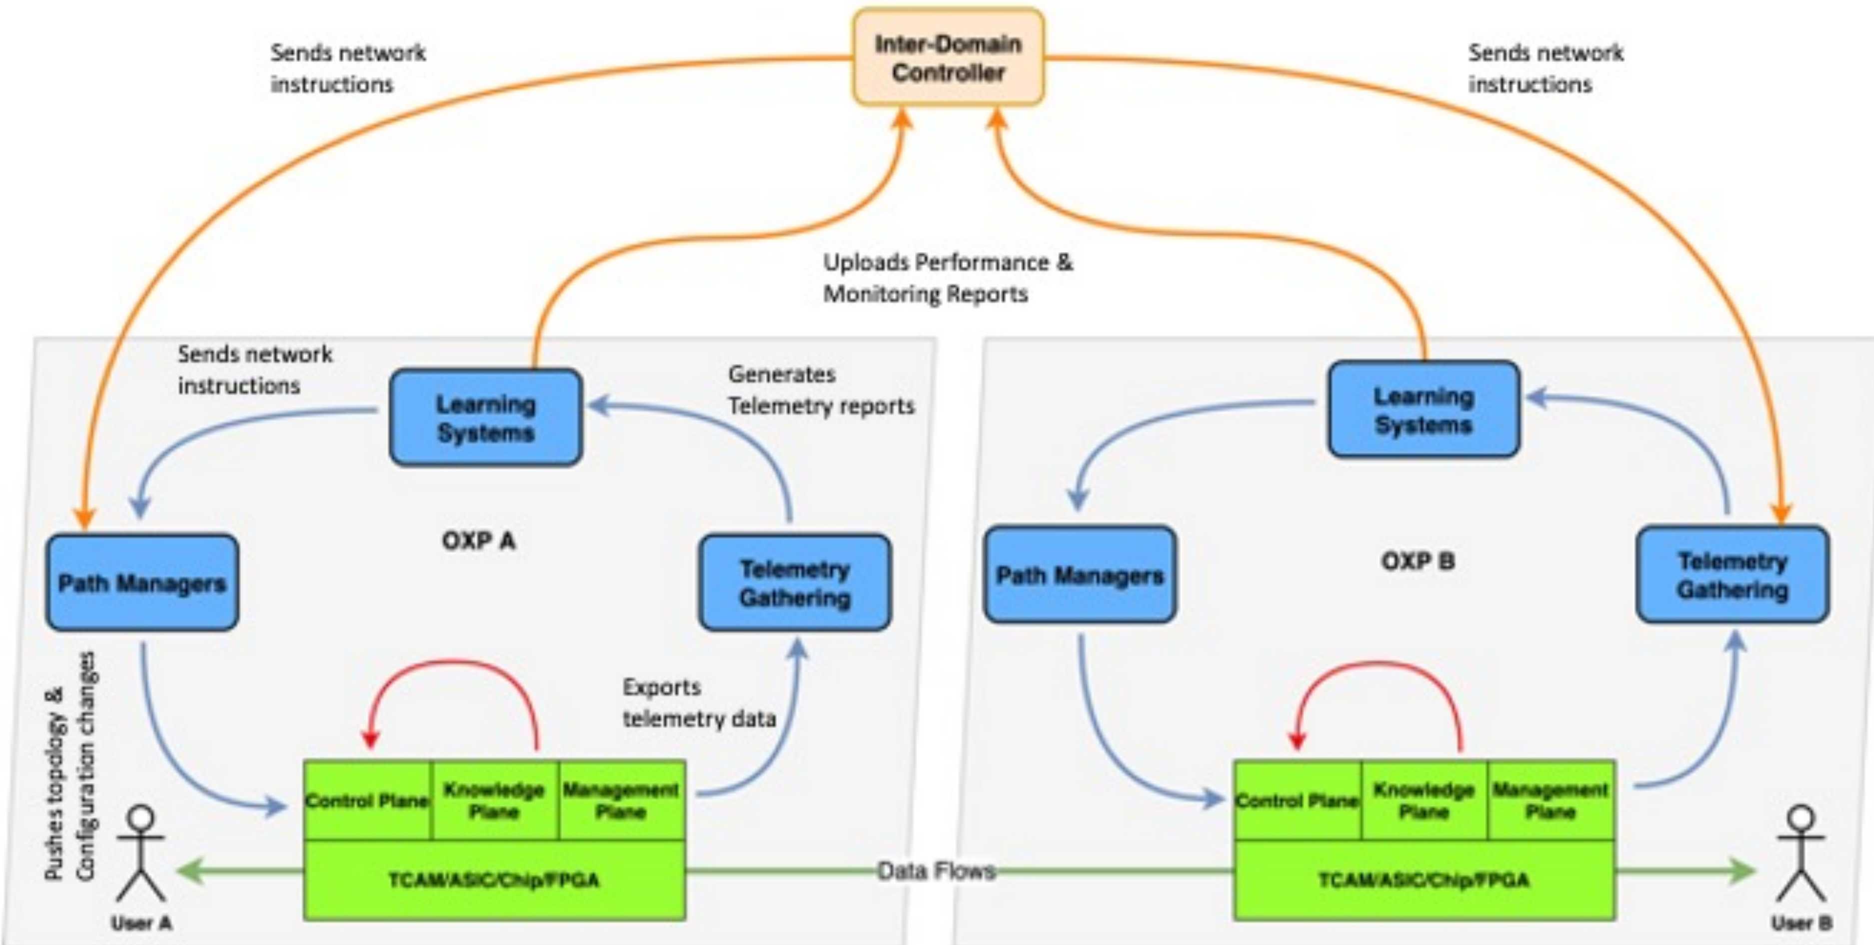

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Use Case: Vera Rubin Observatory

- Vera Rubin 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, and produce a 13 Gigabyte image
- Each image must be transferred to the archive site at SLAC, Menlo Park, CA, within 5 seconds, inside the 27 seconds window
- Multi-traffic types with different priorities (db sync, control, general Internet traffic) must also be supported
- Full network visibility is required to mitigate issues in real time

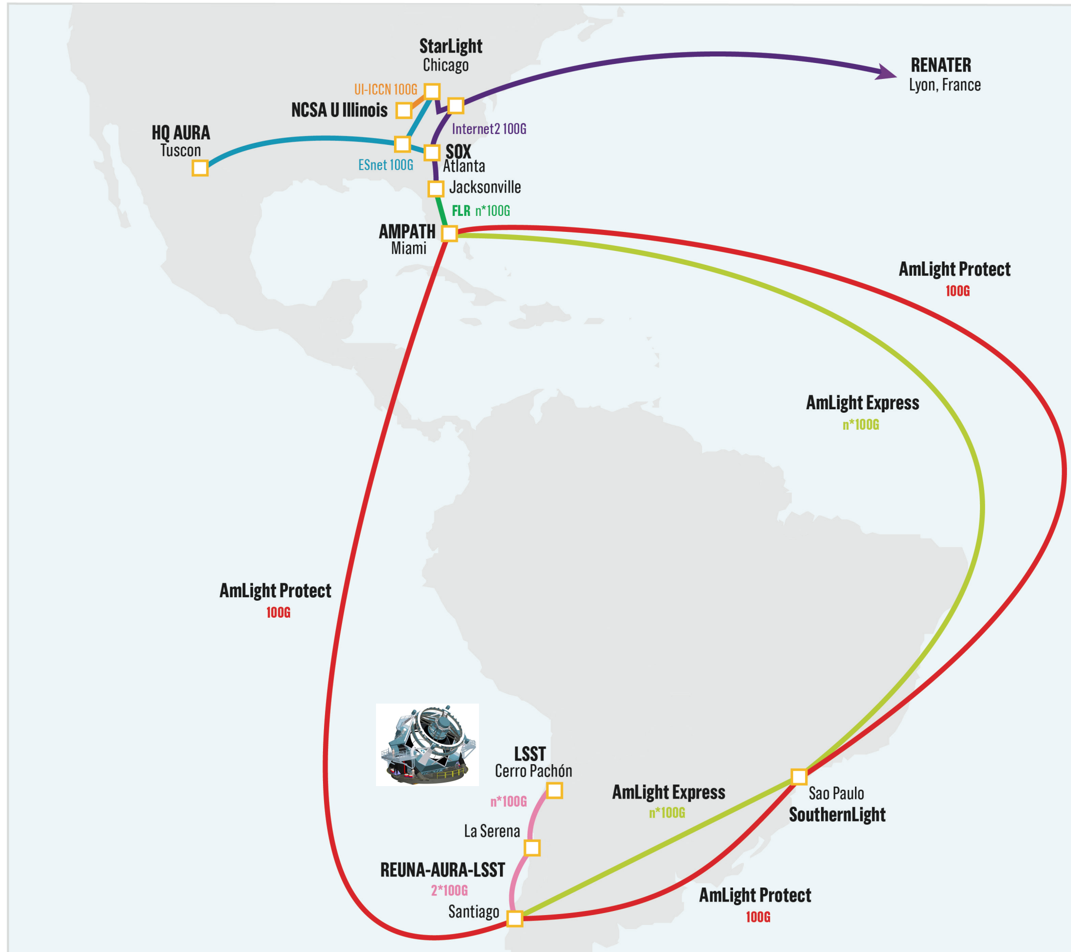

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Use Case: Vera Rubin Observatory [2]

- The AW-SDX Controller will be programmed with Vera Rubin application requirements
- Sub-second closed-loop orchestration focuses on traffic engineering and network tuning in real-time
  - Network policies steer high-priority network flows; e.g.,
  - By setting thresholds for outgoing interface queue occupancy over 80% for more than 500 milliseconds, or
  - Links with bandwidth utilization above 75Gbps for more than 1 second
- Sub-second closed-loop network orchestration aims
  - to lower the risk of having packet drops that could lead to poor data transfer performance

Thank You

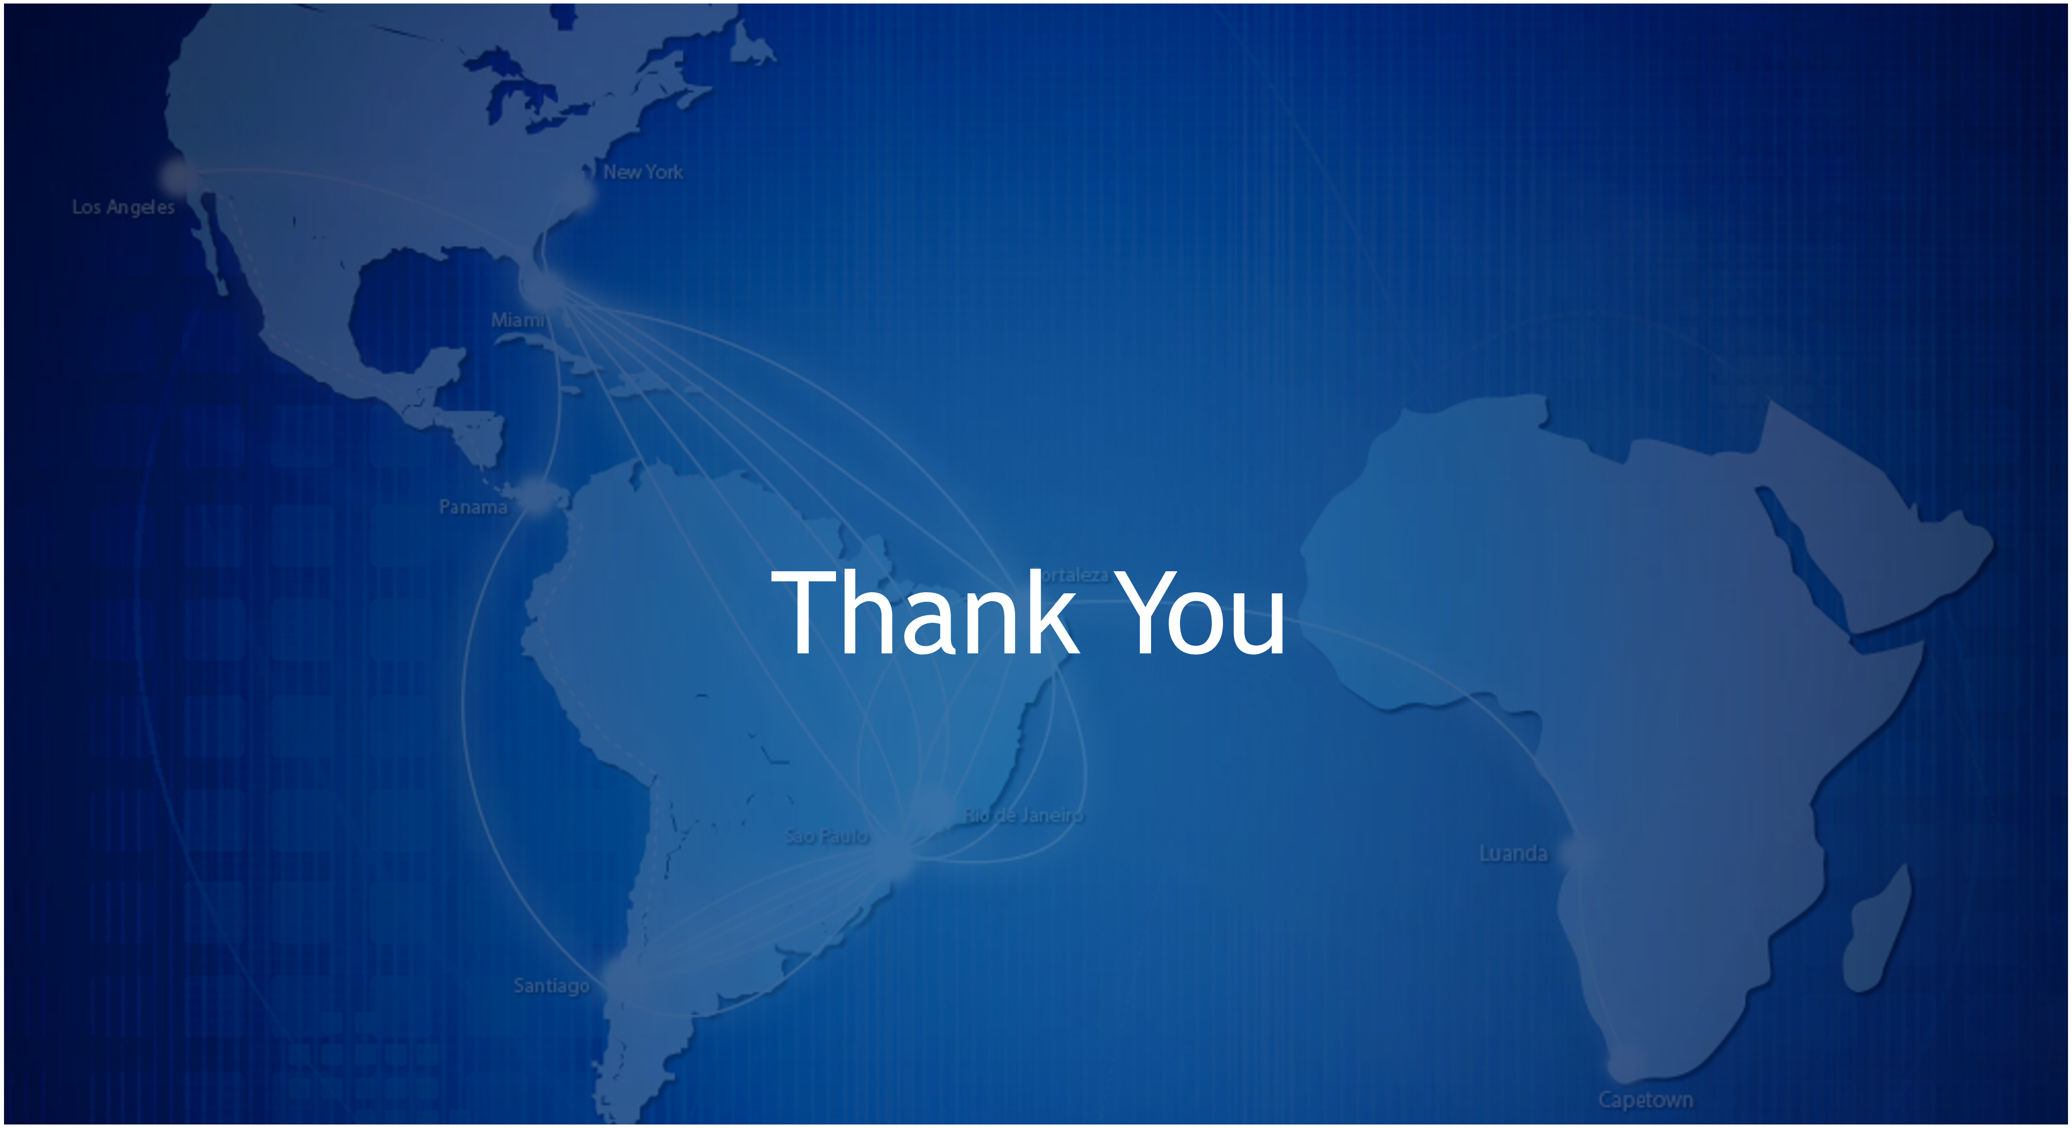

NSF

FIU
FLORIDA INTERNATIONAL UNIVERSITY

USC University of Southern California

renci

UFRGS
UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL

RNP

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE