AtlanticWave-SDX Presentation and Demonstration

FIU: *Julio Ibarra, Jeronimo Bezerra, Vasilka Chergarova*
USC-ISI: *Heidi Morgan, George Papadimitriou*
UNC RENCI: *Yufeng Xin, Mert Cevik, Cong Wang, Michael Stealey*
Georgia Tech: *Russ Clark*

Introduction

- Open Exchange Points (OXPs) are core components of international research network connections
  - There are at least 11 OXPs connecting long-haul circuits across the Atlantic
- Like data centers, OXPs host computation and storage resources
- OXPs provide proximity to upstream transit and peering providers, and to content and cloud providers
- OXPs are uniquely positioned to steer network traffic depending on network conditions

Motivation for a Software-Defined Exchange Point

- Provisioning multi-domain end-to-end network paths is challenging
  - Coordinating multiple network and OXP operators in the path could take days, weeks, or months
- Failures in the network path can be costly to users and operators
  - Time to repair could be days (terrestrial) to weeks (submarine)
  - High coordination cost, involving multiple engineers when
    - Troubleshooting across multiple network domains
    - Time reconfiguring the network to use an alternate path

Software-Defined Exchanges (SDX) use case: Failure Handling

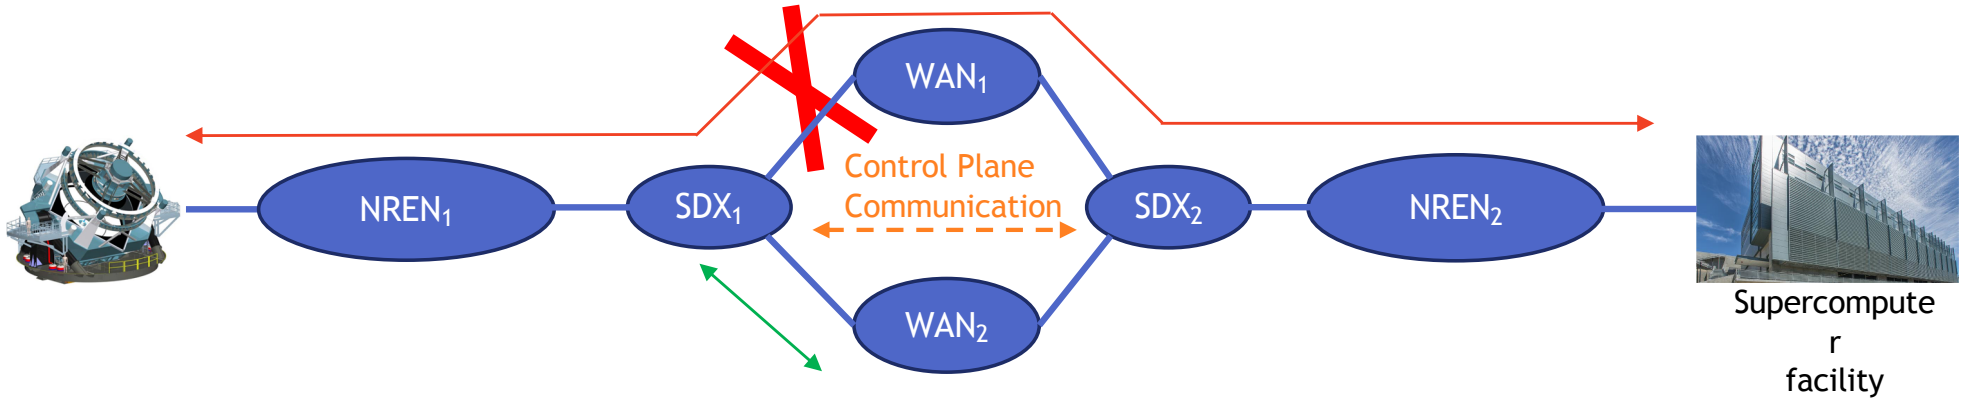

- Evolve OXPs with SDXs by adding a Control Plane communication channel in the interdomain path
- Upon failure with connection to WAN_1, SDX_1 notifies SDX_2 via Control Plane Communication channel
- SDX_1 requests use of the secondary path
  - CPCC computes a new path, then propagates rules to SDX_1 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)

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

AtlanticWave-SDX Project and its Goals

- AtlanticWave-SDX: A Distributed Intercontinental Experimental SDX, NSF Award# OAC-1451024. PI: Julio Ibarra. Goals are:
- To build a distributed SDX between the U.S. and S. America
  - To support planned increases in south-north science flows
  - To integrate the SDN infrastructures at AMPATH, SoX, SouthernLight, and AndesLight open exchange points
- 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

AtlanticWave-SDX leverages the AmLight network

- AmLight provides the underlay network for AtlanticWave-SDX
- AmLight Express network (green), 600Gbps in service:
  - 200G from Boca Raton to Sao Paulo
  - 200G from Boca Raton to Fortaleza
  - 200G from Sao Paulo to Fortaleza
- 100G AmLight Protect ring (solid orange): Miami-Fortaleza, Fortaleza-Sao Paulo, Sao Paulo-Santiago, Santiago-Panama, Panama-San Juan, and San Juan-Miami
- SDXs are shown in circles

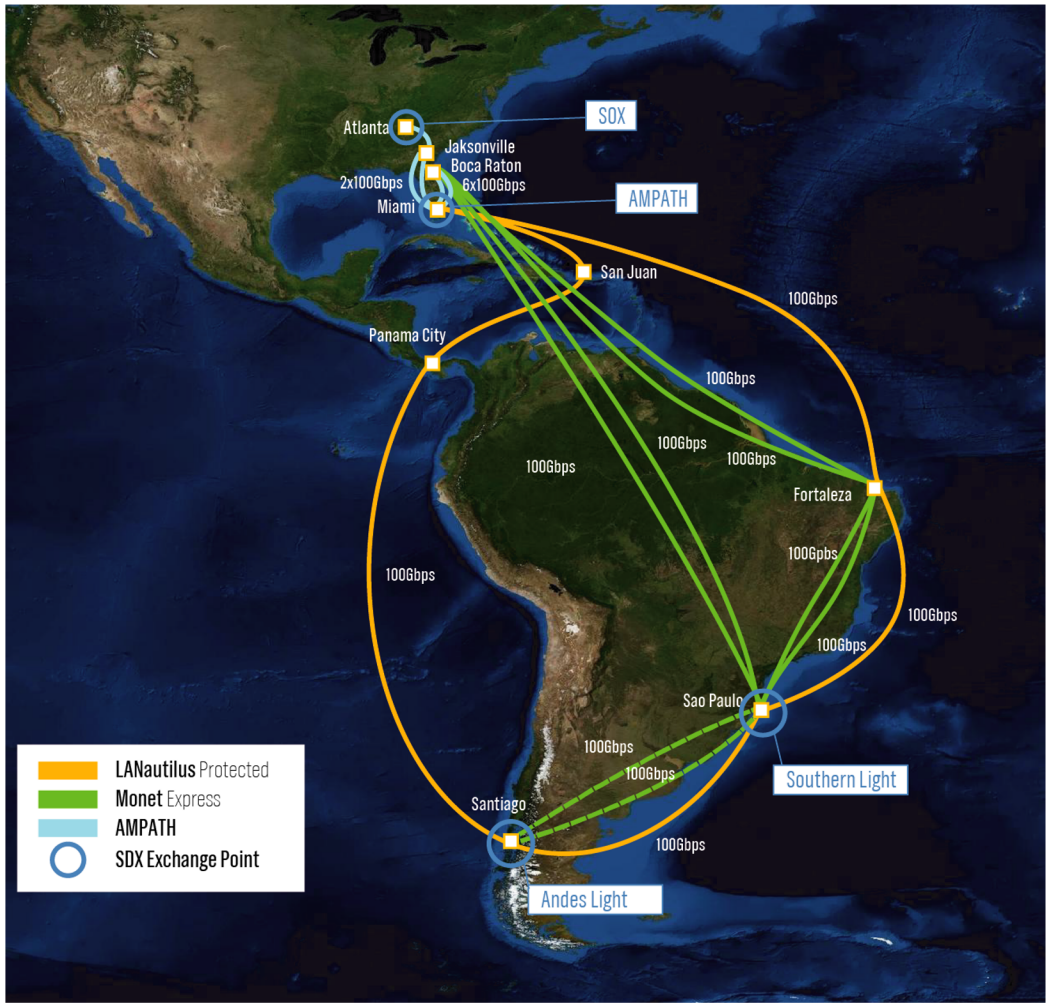

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Project Partners and Collaborators

- Florida International University (FIU)
- University of Southern California Information Sciences Institute (USC-ISI)
- Renaissance Computing Institute at UNC (RENCI)
- Georgia Institute of Technology (GT)
- 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

AtlanticWave-SDX Demonstration

- SDX Features
  - In-band Management Plane
  - Authentication via CILogon
  - Point-to-Point Connections
  - Point-to-Multipoint Connections
  - Bandwidth Reservation
  - Failure Handling
  - REST API
- Flexibility, scalability and reliability
- Demonstrate SDX Features using Pegasus workflow application

AtlanticWave-SDX Network Topology

- AmLight network is the underlay optical network with four OXPs
- An overlay ring topology is configured with Corsa switches deployed at the OXPs.
- A SDX Local Controller and a DTN are deployed at each OXP.
- Duplicated the overlay network at BEN testbed.
- CI/CD environment

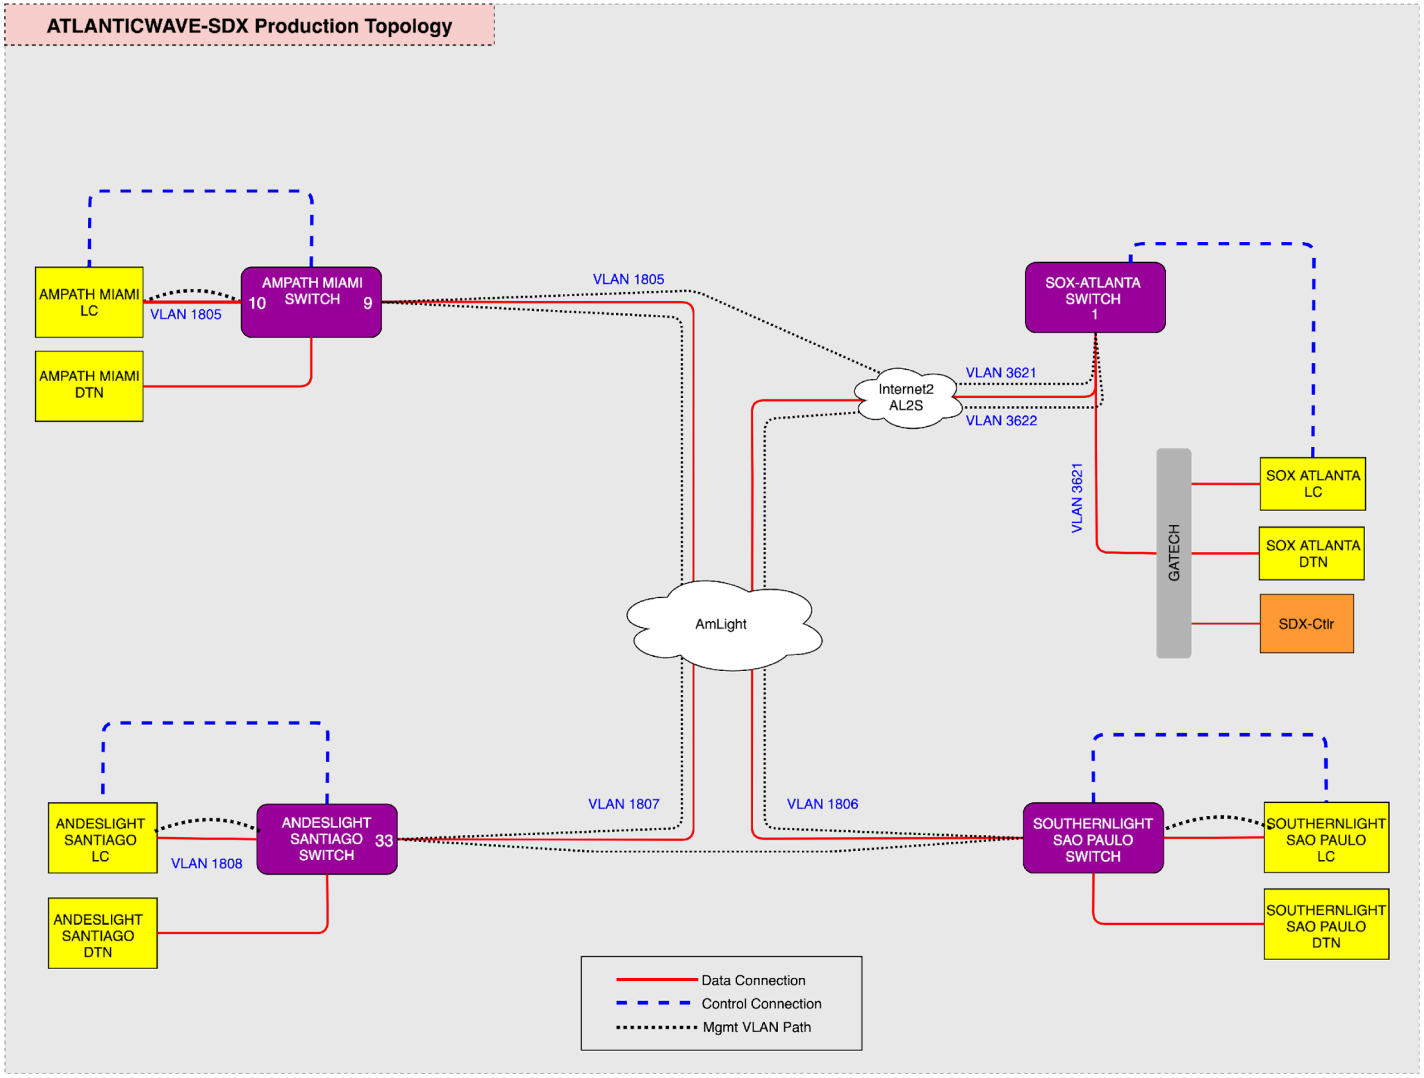

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

AtlanticWave-SDX Portal and CILogon

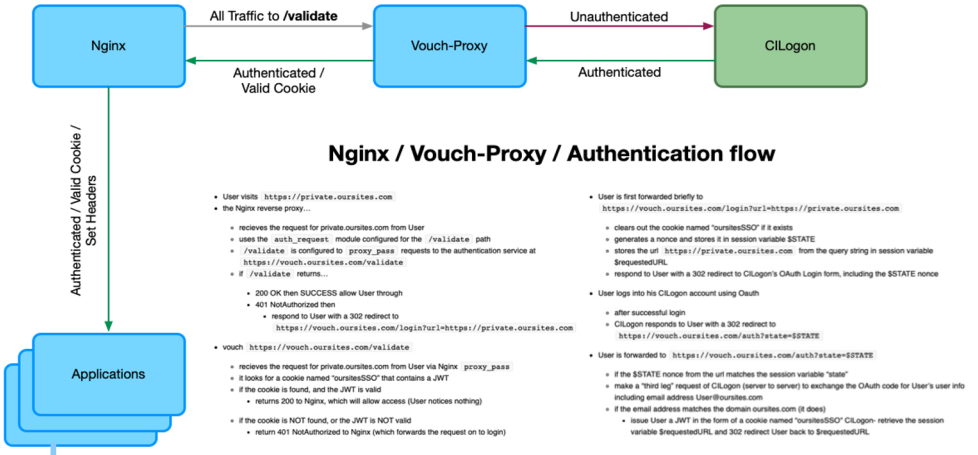

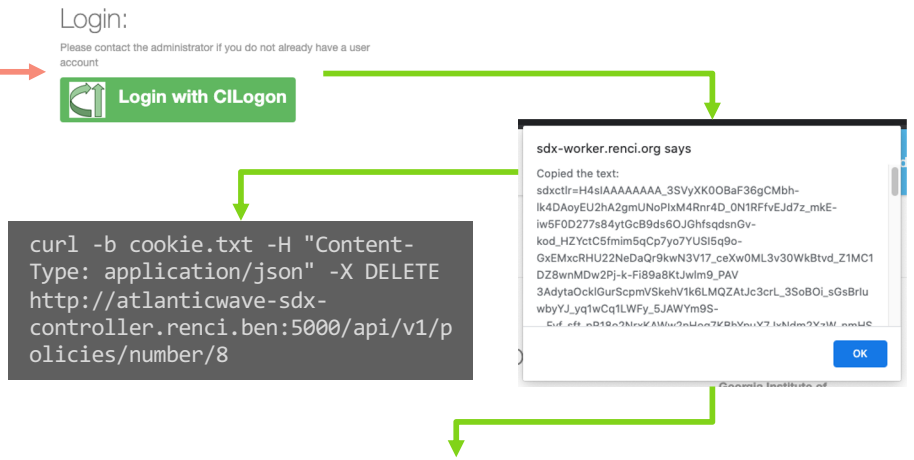

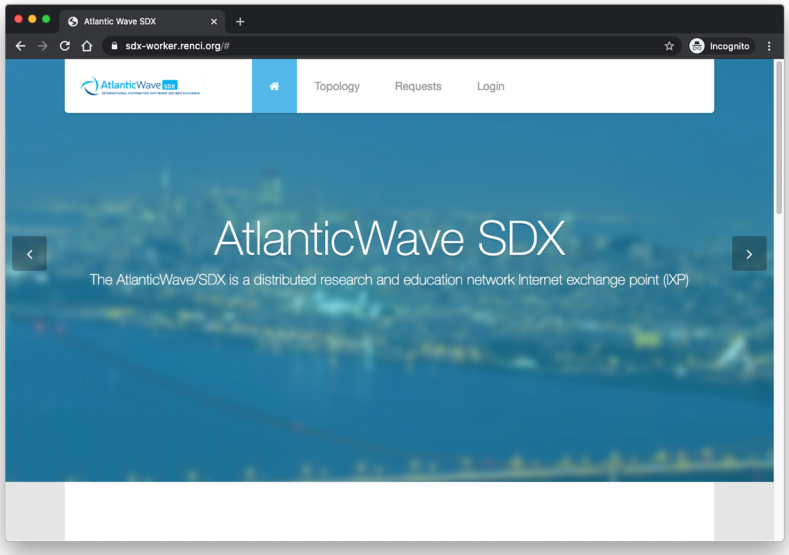

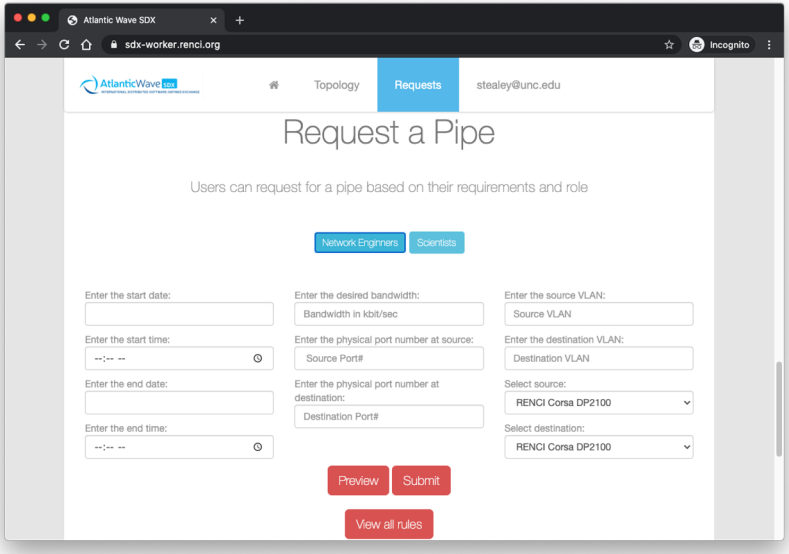

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Failure Handling

- Two phase failure handling
  - In-band management plane
  - Data plane
- Single link failure
  - Atlanta --- Miami
- Auto recovery
  - Management link recovery
  - Data link recovery
- Robust stateful recovery of the controllers after crash

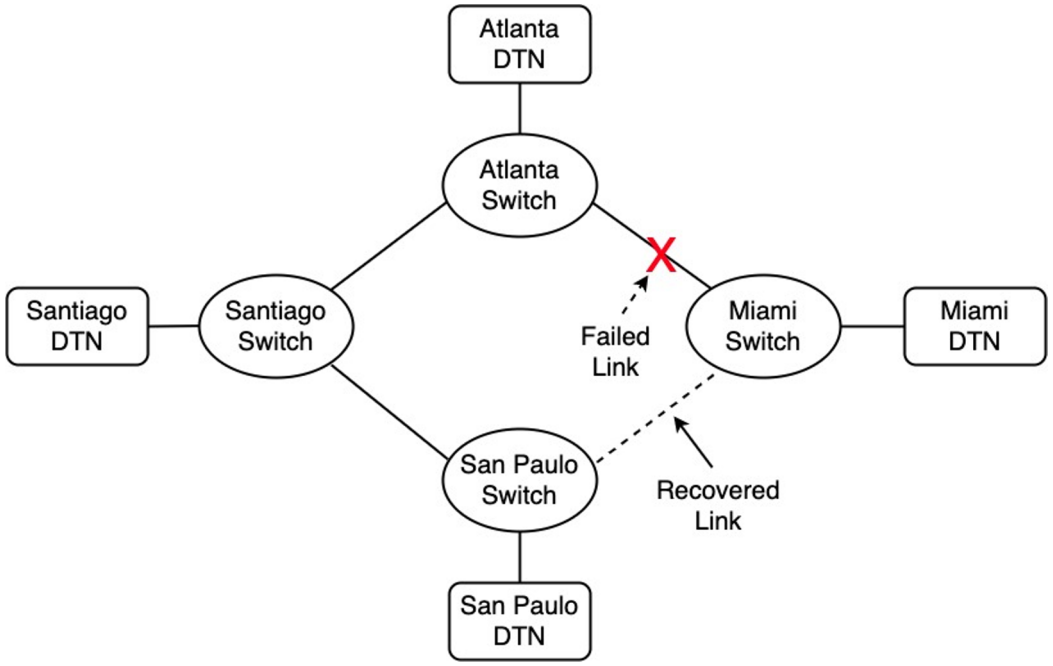

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

AW-SDX Demo

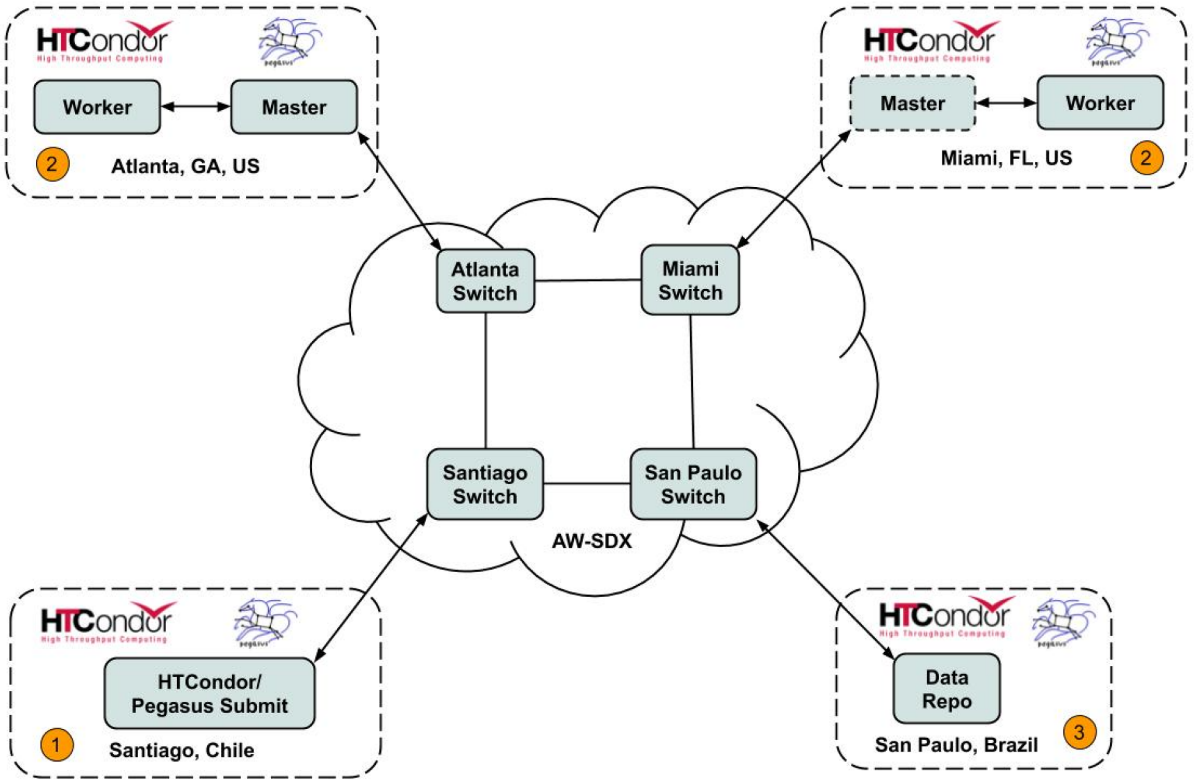

- Demonstrate SDX features
- Pegasus/HTCondor
  - Santiago, Chile (Job submit)
  - Atlanta, GA, US (worker 1)
  - Miami, FL, US (worker 2)
  - San Paulo, Brazil (Data repo)

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Experiment Workflows

- 1000Genome
  - Provides a reference for human variation
  - Reconstructed the genomes across 26 different populations
- Collaborative Adaptive Sensing of the Atmosphere (CASA)
  - NowCasting: short-term advection forecasts (30 mins)
  - 31 grids of predicted reflectivity per minute

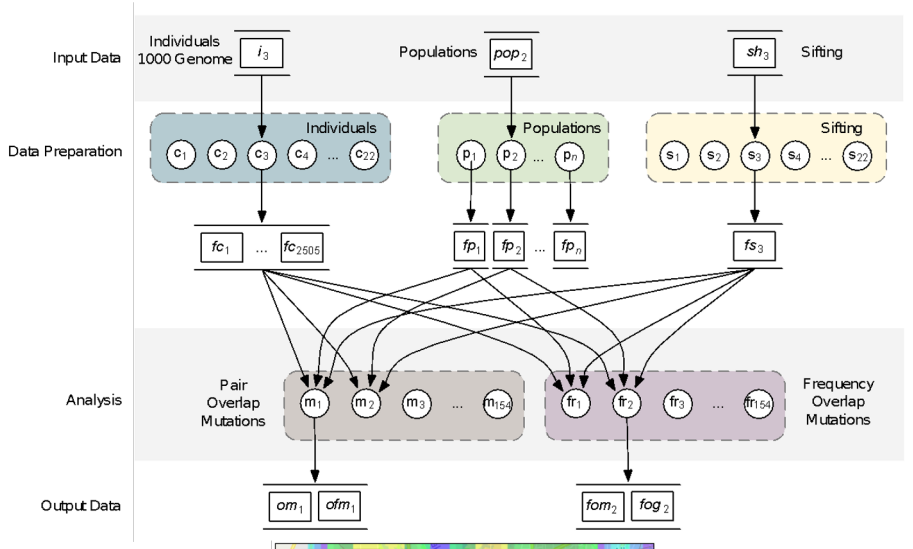

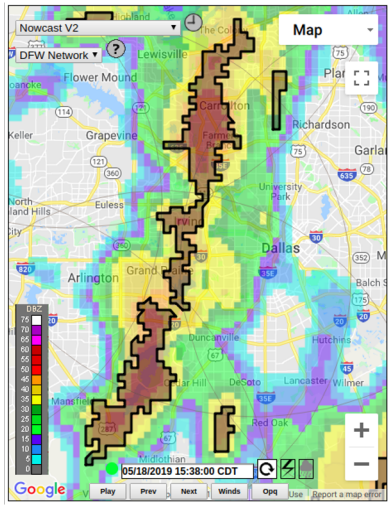

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Data Visualization

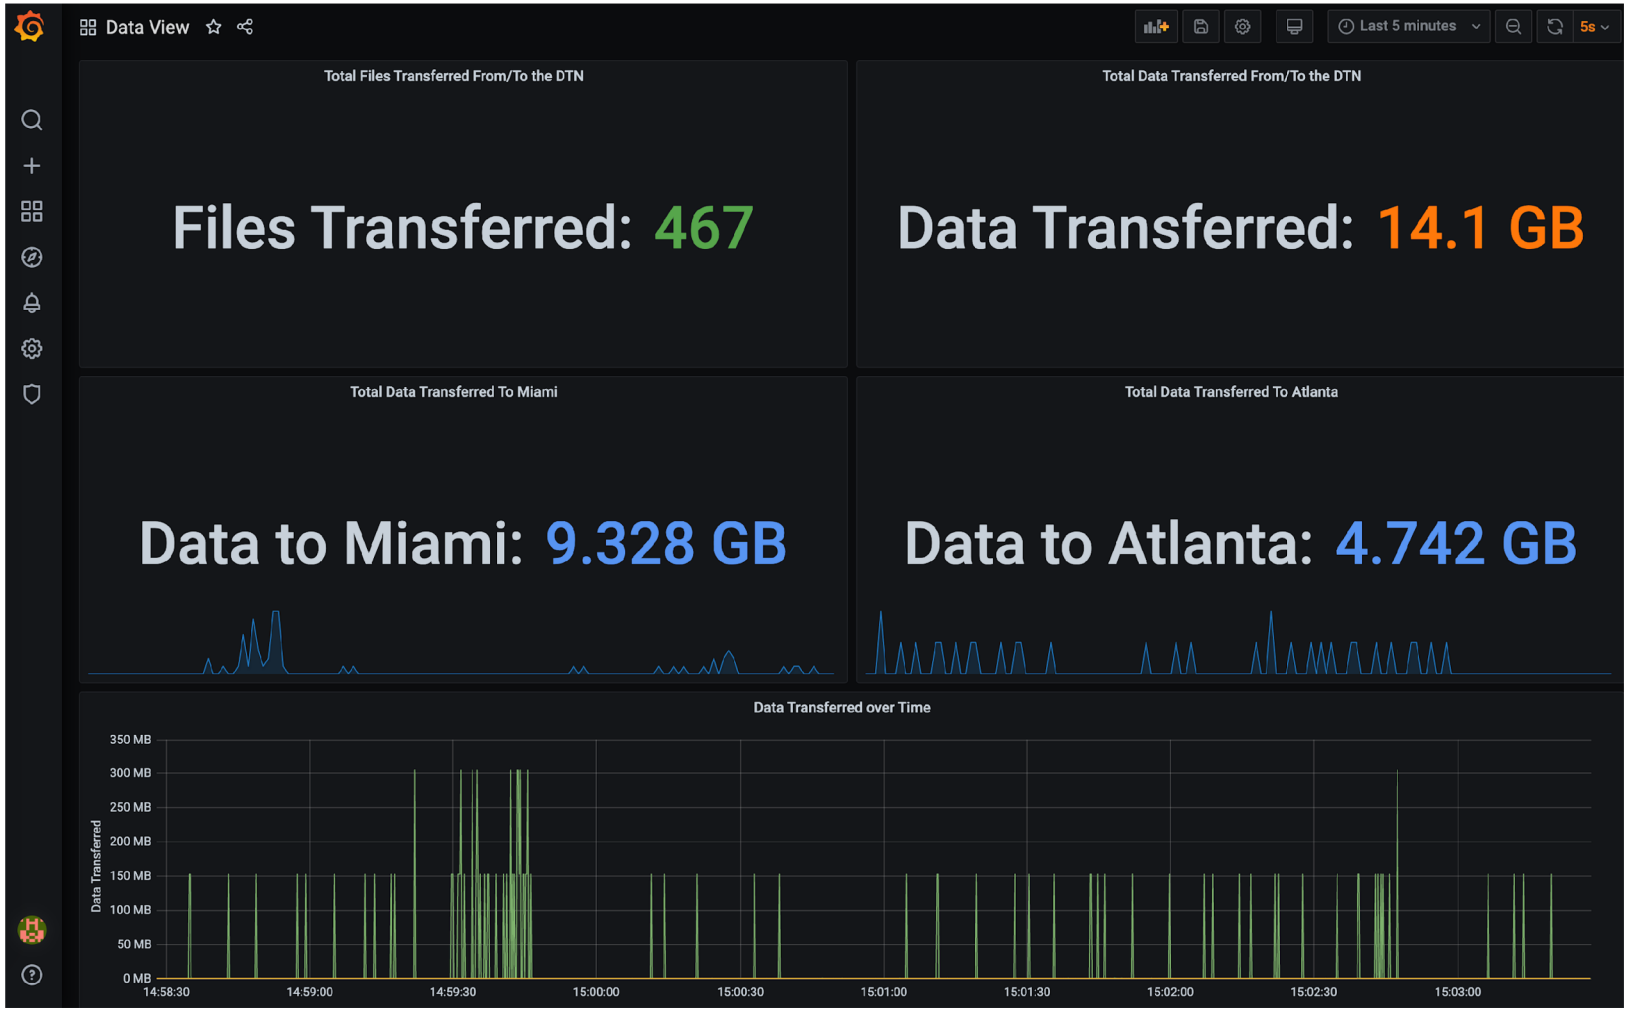

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE

Thank You

NSF

FIU FLORIDA INTERNATIONAL UNIVERSITY

Georgia Tech

USC University of Southern California

renci

AtlanticWave SDX
INTERNATIONAL DISTRIBUTED SOFTWARE-DEFINED EXCHANGE