Pacific Northwest Gigapop Meeting
October 23, 2018

Large Synoptic Survey Telescope (LSST) Scaling Issues and Network Needs

Heidi Morgan, Ph.D.
University of Southern California

Outline

- What is LSST?
- Data Management Facilities
- Network Requirements
- Deployment Calendar
- First Optic Light Demonstration 2017
- Demonstration SC18

Telescope Construction August 2018

AmLight ExP Americas Lightpaths Express & Protect

Large Synoptic Survey Telescope

- This telescope will produce the largest evolving image of the Universe:
  - Expecting 10 million alerts, 15-30 Terabytes of data... every night!
  - First motion picture of our universe
  - This 10 year survey of the sky will catalog 37 billion objects (20B galaxies, 17B stars)
  - Each image is the size of 40 full moons
  - 3200 megapixels camera with 2-second readout
  - 8.4-m mirror, the width of a single tennis court

Webcam: Cerro Pachón Summit 2015-2017 (https://gallery.lsst.org/bp/#/)

Large Synoptic Survey Telescope(2)

WIDE:

- World’s Widest Digital Camera

FAST:

- Rapidly scan the sky
- Charting - exploding supernovae to potentially hazardous near-Earth asteroids

DEEP:

- Trace billions of remote galaxies
- Probes of the mysterious dark matter and dark energy

Site (http://www.lsst.org/lsst/gallery/site)

Mirror Casting (http://www.lsst.org/lsst/gallery/mirror-casting)

Telescope (http://www.lsst.org/lsst/gallery/telescope)

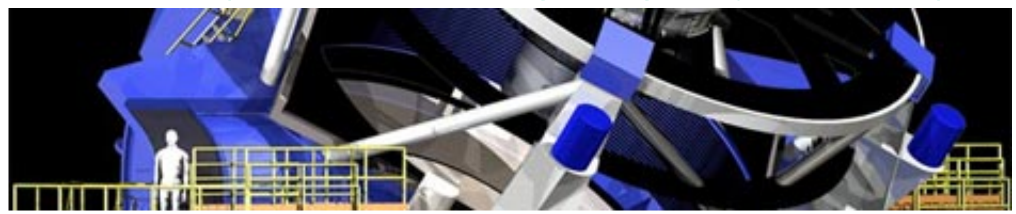

Data (http://www.lsst.org/lsst/gallery/data)

LSST AmLight Trailer (https://youtu.be/T6eRshTuU88)
LSST & AmLight Video (https://youtu.be/c16h1yyS-78)
Seeing the Beginning of Time (https://tinyurl.com/ybk27qm7)

LSST Data Management Facilities

LSST Network Requirements

The LSST operation will consist of two main channels:

- Control Channel
  - Remotely operated from Tucson, AZ.
  - Requires low bandwidth, low latency, and high priority,
  - Bandwidth needs are about 2 Mbps
- Data Channels
  - Multiple data streams summing up 90 Gbps:
    - Each 12.7GB data set (6.4GB picture + 6.3GB metadata) must be transmitted to the U.S. in 5 seconds
    - Database synchronization, etc.
- The end-to-end path must provide high resilience, low delay, multiple paths, and an efficient control plane to act in all network status changes

LSST End-to-End Path (1)

- Some R&E networks can accommodate some of the LSST requirements:
  - Multiple paths with multiple 100G links (including AmLight-ExP)
  - Dynamic provisioning, bandwidth reservation
- But R&E networks are interconnected through Academic Exchange Points that have challenges:
  - Heterogeneous configurations and services
  - Almost no support for network programmability (SDN)

LSST End-to-End Path (2)

- High demand end-to-end applications like LSST require that all networks in the path support QoS and Programmability
  - Including the Academic Exchange Points
  - Software Defined Exchanges (SDX) within the exchange point are a possible solution
  - AtlanticWave SDX is aiming to overcome those limitations that exchange points have
  - LSST is a primary use case for Awave-SDX, a collaboration between FIU, Gatech, and SouthernLight in Sao Paulo to be extended to AndesLight in Santiago, Chile this October

LSST End-to-End Path (3)

AtlanticWave-SDX

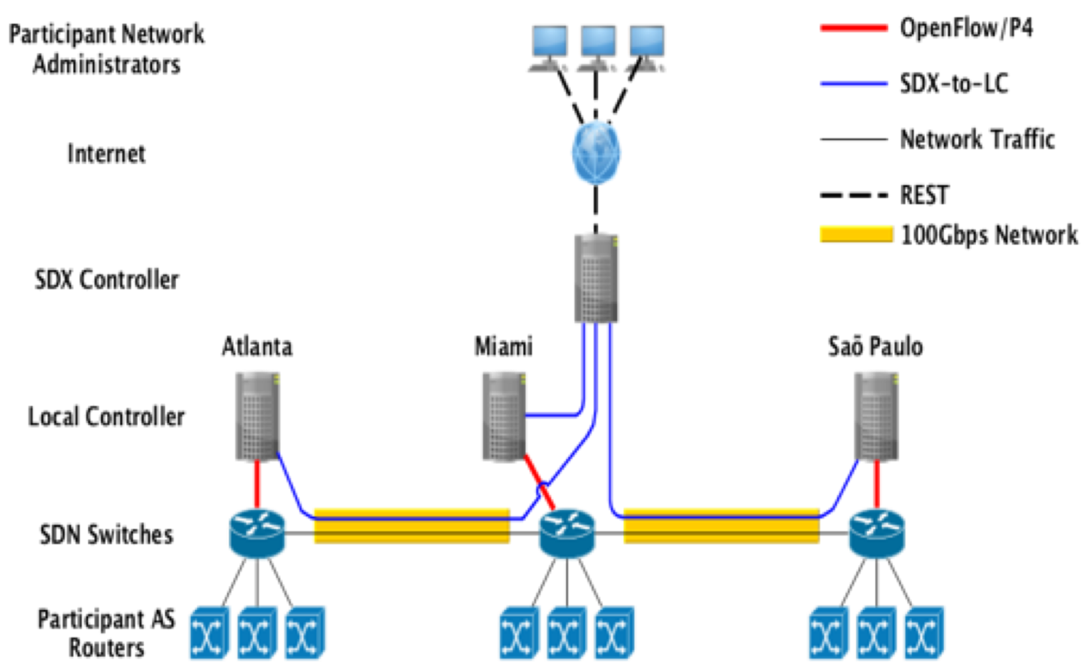

- Updated: now with four locations to cover: *La Serena*, Sao Paulo, Miami, and Atlanta.
- Thousands of KM of fiber between each location
- Split controller design
  - Central controller for interacting with users
  - Local controllers at each location

AmLight ExP
Americas Lightpaths Express & Protect

LSST Deployment Calendar

- From 2017 to 2020, the LSST network will be installed and fine-tuned (mainly for evaluations and simulations)
- Engineering ”first light” anticipated in 2019
- Science ”first light” in 2021
- Full operations for a ten-year survey commences in Q4 2022
- The LSST Network Engineering Team must guarantee all network requirements by 2021
  - Bandwidth guaranteed and resilience
  - Effective and efficient operation

<table>
<tr><td colspan="3">Fiscal 2017</td><td colspan="4">Fiscal 2018</td><td colspan="4">Fiscal 2019</td><td colspan="4">Fiscal 2020</td><td colspan="4">Fiscal 2021</td><td colspan="4">Fiscal 2022</td><td colspan="4">Fiscal 2023</td><td colspan="5">Fiscal 2024</td></tr>
<tr><td colspan="4">Calendar 2017</td><td colspan="4">Calendar 2018</td><td colspan="4">Calendar 2019</td><td colspan="4">Calendar 2020</td><td colspan="4">Calendar 2021</td><td colspan="4">Calendar 2022</td><td colspan="4">Calendar 2023</td><td colspan="4">Calendar 2024</td></tr>
<tr><td>Q1</td><td>Q2</td><td>Q3</td><td>Q4</td><td>Q1</td><td>Q2</td><td>Q3</td><td>Q4</td><td>Q1</td><td>Q2</td><td>Q3</td><td>Q4</td><td>Q1</td><td>Q2</td><td>Q3</td><td>Q4</td><td>Q1</td><td>Q2</td><td>Q3</td><td>Q4</td><td>Q1</td><td>Q2</td><td>Q3</td><td>Q4</td><td>Q1</td><td>Q2</td><td>Q3</td><td>Q4</td><td>Q1</td><td>Q2</td><td>Q3</td><td>Q3</td></tr>
<tr><td colspan="23">LSST Construction & Commissioning</td><td colspan="9" rowspan="3">LSST Full Operations</td></tr>
<tr><td colspan="7">Net Deployment - Phase 1</td><td colspan="8">Net Deployment - Phase 2</td><td colspan="8"></td></tr>
<tr><td colspan="8"></td><td colspan="15">LSST Early Operations</td></tr>
</table>

AmLight ExP
Americas Lightpaths Express & Protect

LSST Network Topology

- AmLight-Express (spectrum) is the primary
  - MONET cable system is Ready For Service (RFS) Fortaleza, Brazil to Boca Raton, Florida
  - South Atlantic Cable System (SACS) is lunched and will be RFS by the end of 2018
- AmLight-Protect 100G ring is the backup

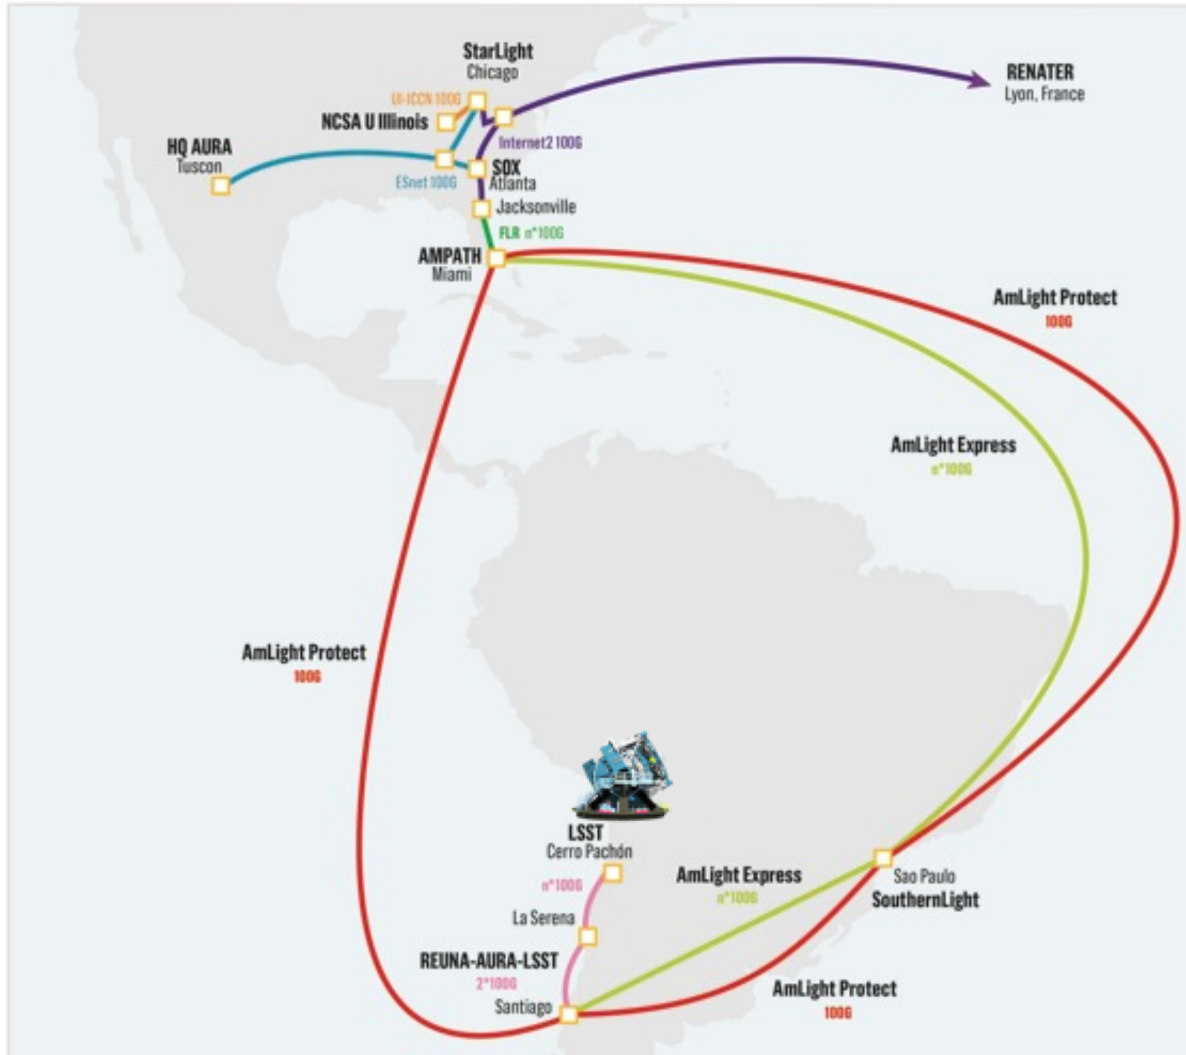

AmLight ExP
Americas Lightpaths Express & Protect

SLA requirements for networking

- General SLA:
  - Mean Time Between Failures (MTBF) - **180 days**
  - Mean time To Repair (MTTR) - **48 Hours**
  - Different QoS profiles for different network conditions (chart below)
- Bandwidth Required:
  - Summit (Cerro Pachon) to Base(La Serena, Chile): 200Gbps
  - Base (La Serena, Chile) to Archive (NSCA Champagne, Illinois, USA): 200Gbps
  - Archive to Data Archive Centers - DACs (ex. DAC Lyon, France): 10Gbps

	Chile to Miami		Miami to Illinois	
	Full Connectivity	Minimal Connectivity	Full Connectivity	Minimal Connectivity
Traffic Types	Bandwidth Available 300G	Bandwidth Available 100G	Bandwidth Available 200G	Bandwidth Available 100G
Science & Data Backbone Transport	90	70	90	70
Other Behavior Aggregates	45	10	45	10
Chilean Academic Usage	Best Effort	Best Effort	Best Effort	Best Effort
Bandwidth Reservation for LSST	45%	80%	68%	80%

First Optic Light Demonstration Dec 2017

- The LSST Network Engineering Team (NET) completed the first successful transfer of digital data over LSST/AURA fiber optic networks from the Summit Site on Cerro Pachon, Chile to the Base Site in La Serena, Chile and on to the Archive Site at NCSA in Champaign.
  - A set of 6 x 10 Gbps Data Transfer Nodes (DTN)
  - Sustained data rate of approximately 44 gigabits per second over a period of 24 hours

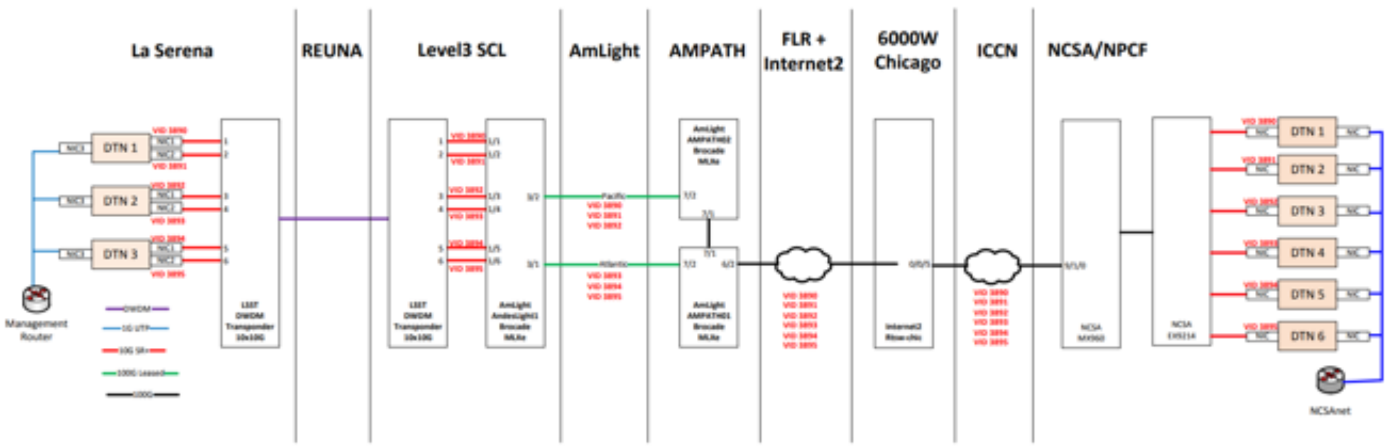

AmLight ExP
Americas Lightpaths Express & Protect

Demonstration Supercomputing Conference SC18

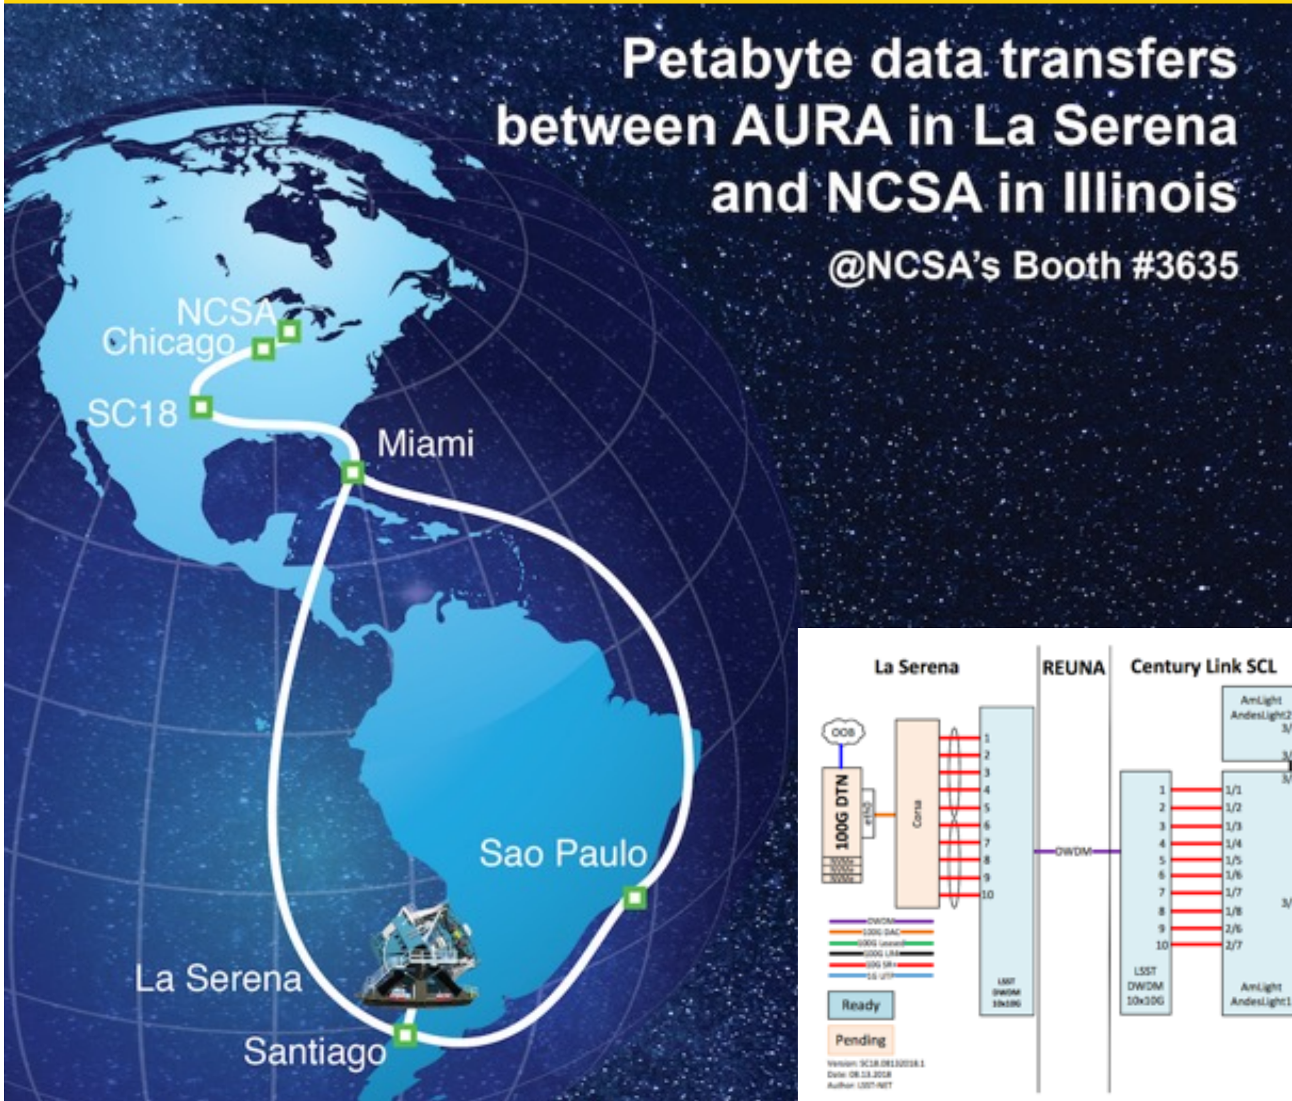

- *Turbo* version of the First Optic Light demo:
  - Instead of memory-to-memory tests, we will run 100G disk-to-disk tests
  - 100G FIONA DTNs installed in La Serena and NCSA
    - DTNs will be added to PRP's Maddash after SC18
  - Real files from (GEMINI or DECam) project will be used

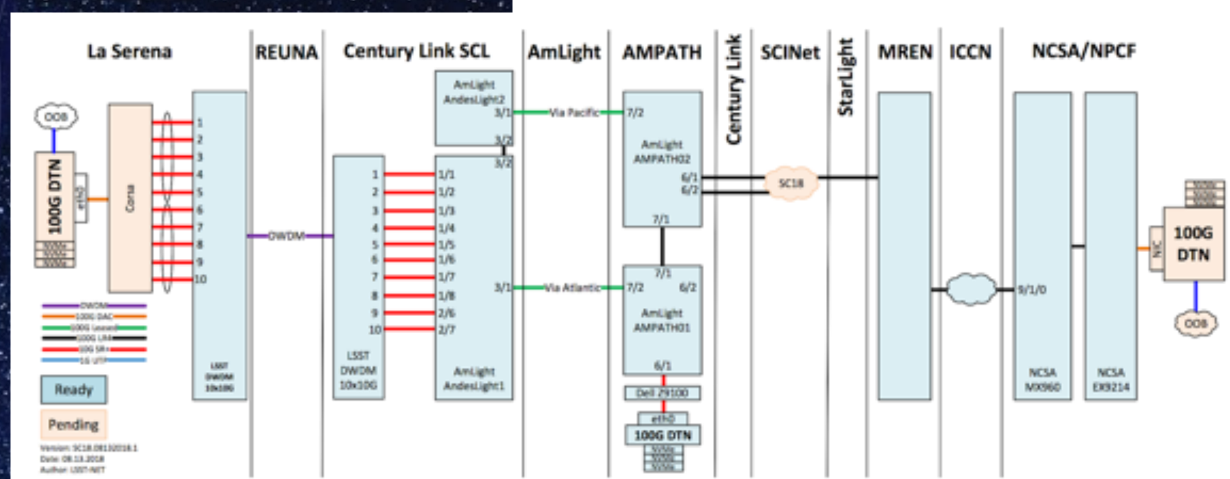

AmLight ExP
Americas Lightpaths Express & Protect

Network Monitoring/Performance

- The *LSST Management Plan* defined a centralized Network Management System (NMS) for easy monitoring
  - To be used by LSST NOC and LSST Engineering teams
  - A single NMS will see all links, devices, CPUs, power supplies, etc. of all networks in the path
- perfSonar instrumenting the path: currently each site has two 10G ports dedicated for monitoring and performance evaluation; 100G nodes are planned for the future.
- End-to-end Performance Metrics: Latency, Packet Loss, Jitter, Maximum Payload Throughput
- Other performance and monitoring instruments being evaluated, such as the IRNC AMIS nodes (UMass)

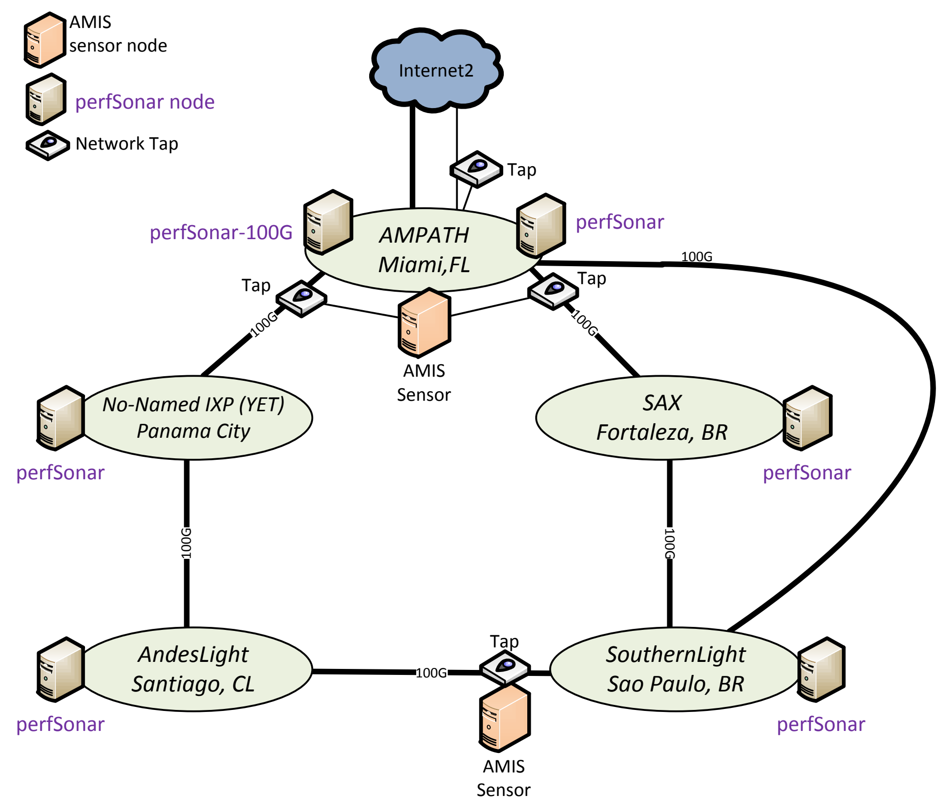

Performance Instruments at AmLight

AmLight ExP
Americas Lightpaths Express & Protect

In-band Network Telemetry

- INT is an implementation to record operational information in the packet while the packet traverses a path between two points in the network:
  - Complements current out-of-band OAM mechanisms based on ICMP or other types of probe packets.
  - 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, WiFi link power, CPU utilization, Battery Utilization, Sequence #, and many others

- As metadata is exported directly from the Data Plane, Control Plane is not affected:
  - Translating: *you can track/monitor/evaluate EVERY single packet at line rate.*

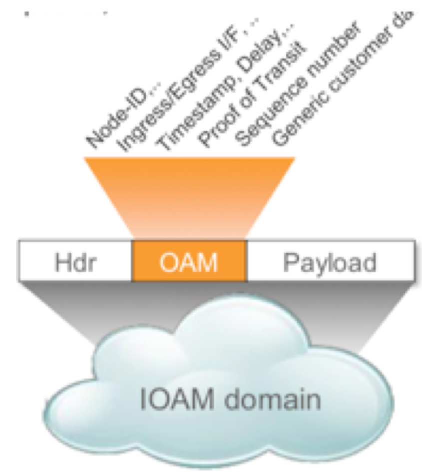

AmLight ExP
Americas Lightpaths Express & Protect

Telemetry at AmLight: LSST

- What if the LSST doesn’t manage to send its data in its 5-seconds transfer window?
  - For instance, because of packet loss, lack of capacity, lack of buffers, microburst, DoS attacks?

- If the data transfer window is missed, will AmLight engineering team be able to fix whatever it is happening before the next window (in 22 seconds)?

- How many windows are we going to miss if we have to troubleshoot it manually?
  - Don’t forget our engineers will be sleeping!
    - ***AmLight-INT Project might be the solution!***

Telemetry at AmLight: A Use Case

- NSF IRNC: Backbone: AmLight In-band Network Telemetry (AmLight-INT), Award# OAC-1848746
- Focuses on LSST and Big Data needs
- AmLight-INT Project Plan:
  - Deploy P4/INT-capable switches
  - Deploy INT sinks (100G hosts) to collect metadata
  - Develop a new methodology to collect and export INT data in real time to feed SDN controllers and users with monitoring information
  - Create a Network Telemetry Design Pattern to be used by other R&E networks

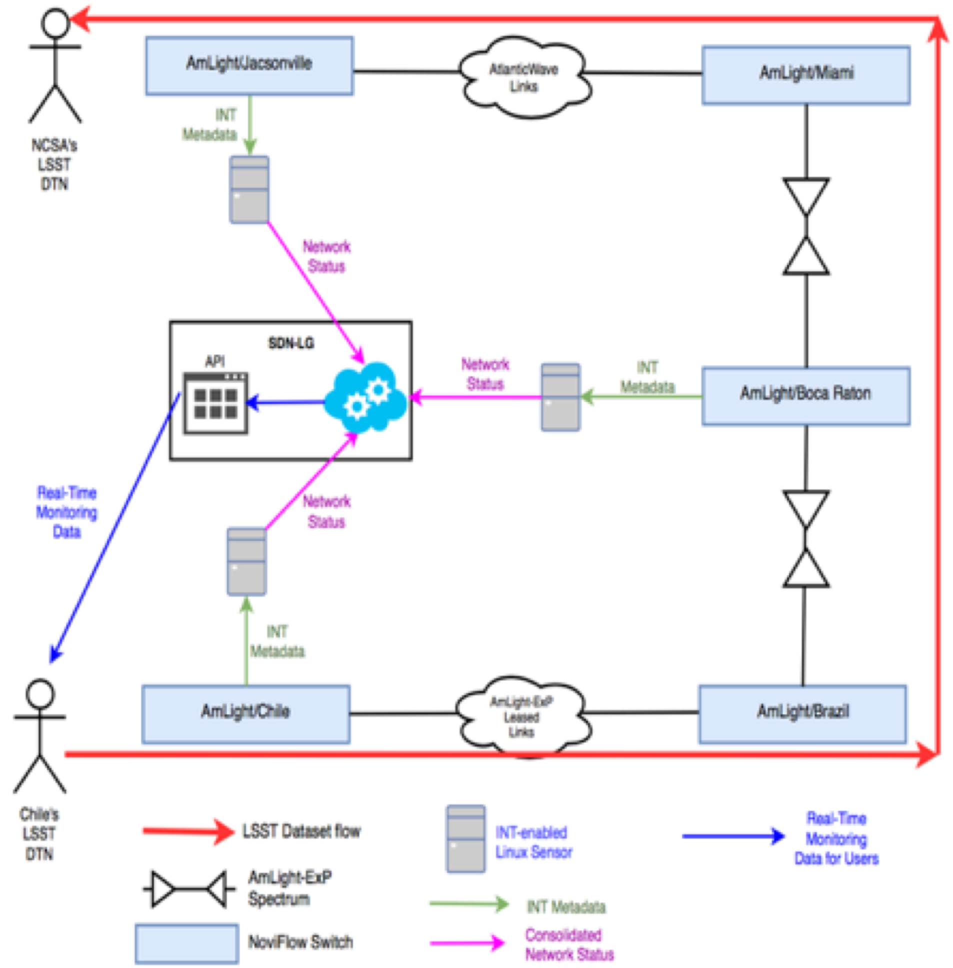

AmLight ExP
Americas Lightpaths Express & Protect

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
- Fortaleza is a landing point for all cables, except for SABR

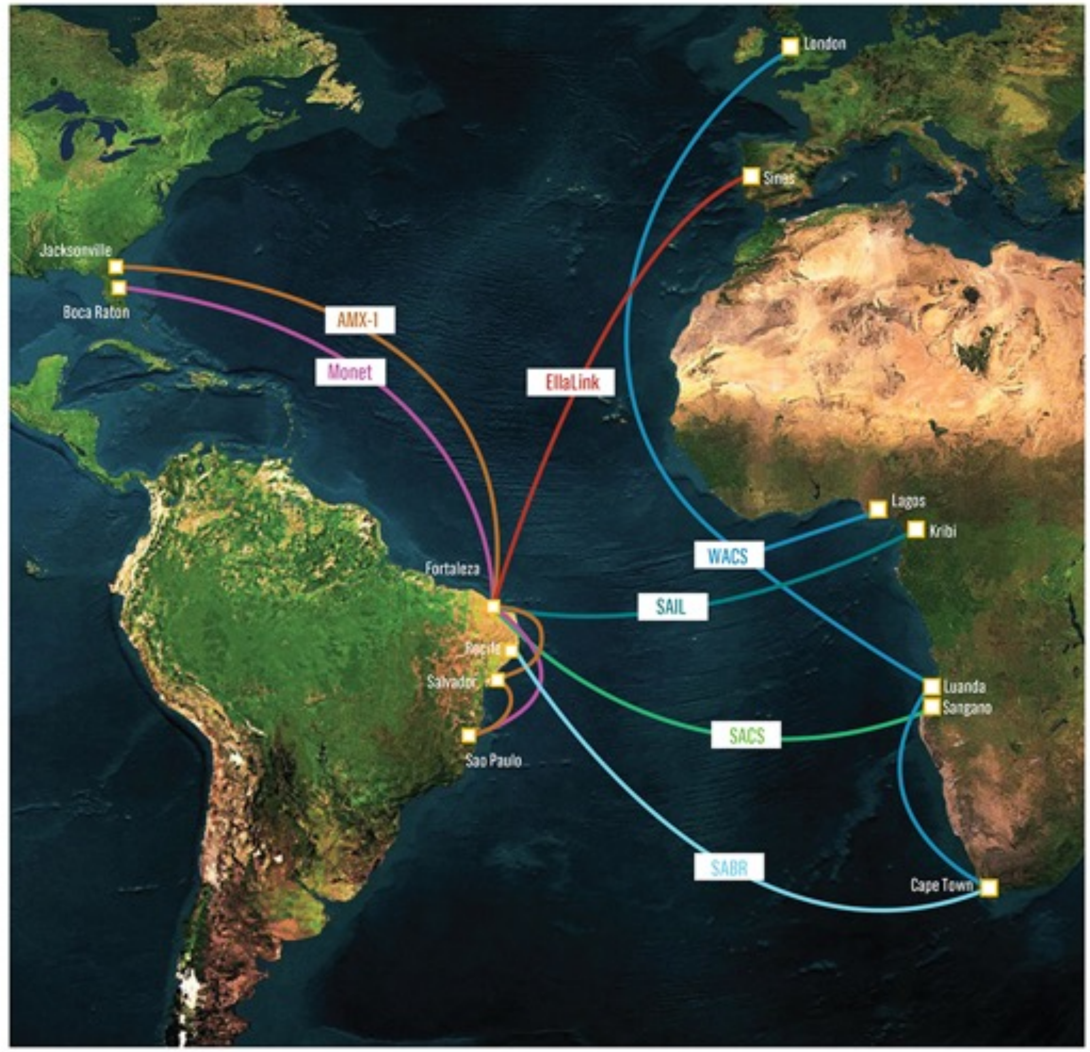

AmLight ExP
Americas Lightpaths Express & Protect

AARCLight: Americas Africa Research and eduCation Lightpaths, NSF Award #OAC-1638990

- Planning activity that aims to
  - Define a strategy for research and education network connectivity between the US and West Africa along the South Atlantic
  - Coordinate planning efforts among stakeholders in the U.S., Africa, and Brazil
  - Create economies of scale
  - Understand the potential impact
    - From the use of the offered spectrum
    - Towards serving the broadest communities of interest in research and education

AmLight ExP
Americas Lightpaths Express & Protect

MeerKAT Radio Telescope

- MeerKAT is a 64-dish system
- MeerKAT will be integrated into the mid-frequency component of SKA
- Generates data at a rate of 4.7Gbps
- Data is transported either by tapes, or scientists travel to South Africa
- Data is collected and digitized at each antenna and then streamed via optical link to the Karoo Array Processor Building (KAPB) for science processing
- Clock signal is distributed from a central point to remote dishes

- Western Cape has many international collaborations with LSST, Very Large Array (VLA), and Arecibo observatory. South Africa LSST Committee members are Markus Bottcher (NWU), Fernando Camilo (SKA SA; Chair), Roger Deane (RU), Russ Taylor (UCT, UWC), Patricia Whitelock (SAAO, UCT), and Ted Williams (SAAO).

AmLight ExP
Americas Lightpaths Express & Protect

SKA Data Transport Requirements

- SKA will be the world's largest radio telescope
- Australia and South Africa are the host countries
- SKA Regional Science and Engineering Centers (SRC) will provide access to
  - SKA data products
  - Computational resources for processing and analysis
  - A long-term archive for SKA science data products
- NRENs will be critical to support data movement to SRCs and users

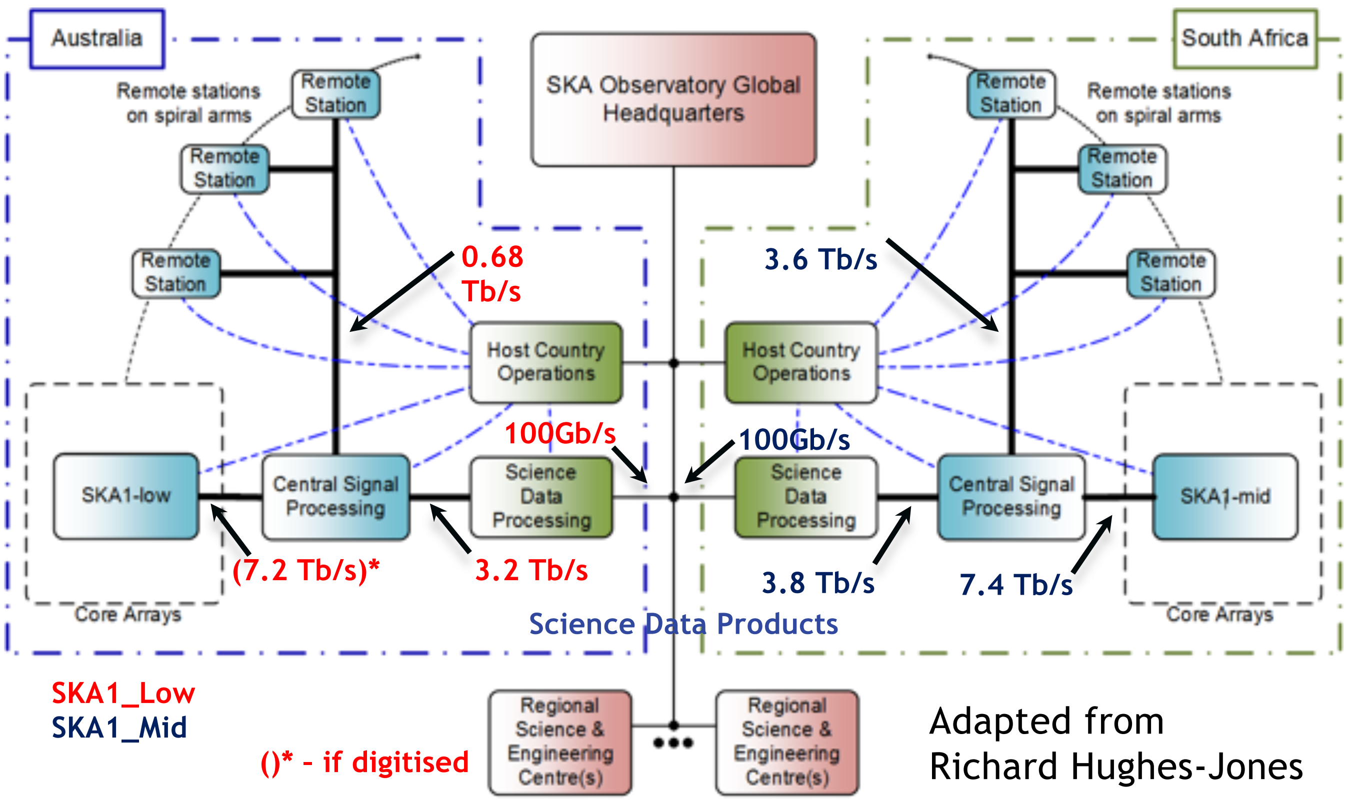

Adapted from Richard Hughes-Jones

AmLight ExP
Americas Lightpaths Express & Protect

Challenges & Ongoing Efforts

- “LSST NOC” is a virtual NOC composed of multiple NOC teams
- Bi-annual AmLight-ExP South American Astronomy Coordination Committee (SAACC) meetings focus on networking needs for science & technology
- Monthly LSST Network Engineering Team (NET) calls and NET meetings at the SAACC
- Bandwidth Calendaring: Different network utilization profiles for night and day are being discussed
- Security and Applications are being handled by separate Association of Universities for Research in Astronomy (AURA) Data Management Teams

THANK YOU!

Heidi Morgan, Ph.D.
hlmorgan@isi.edu

Infrastructure, science application support, education, outreach and community building efforts are made possible by funding and or support from:

- National Science Foundation (NSF) awards: LSST #ACI-1202910, AmLight ExP #ACI-1451018, AtlanticWave SDX #ACI-1451024
- FAPESP, ANSP – grant no. 2008/52885-8
- Rede Nacional de Ensino e Pesquisa (RNP)
- REUNA
- CLARA
- Association of Universities for Research in Astronomy (AURA)
- Florida International University
- Internet2 and GRNOC
- ESnet
- Florida LambdaRail (FLR)
- National Center for Supercomputing Applications (NCSA)
- Latin American Research and Education community
- The many national and international collaborators who support our efforts