NSF's International Research Network Connections Program

Kevin Thompson (CISE/ACI)

November 2015 for JETNet

Networking Programs in CISE/ACI

- Networking as a fundamental layer and underpinning of CI
- CC*DNI/CC-NIE/CC-IIE (Campus Cyberinfrastructure – Network Infrastructure and Engineering): joint with CNS
  - Campus networking upgrade (re-design to scienceDMZ at campus border and 10/100Gbps) and innovation program
- IRNC – International R&E Network Connections: joint with OIA/ISE
  - Scientific discovery as a global collaborative endeavor
  - Provide network connections linking U.S. research with peer networks in other parts of the world
  - Stimulate the deployment and operational understanding of emerging network technology and standards in an international context

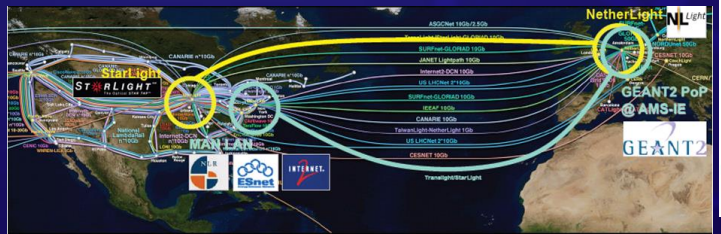

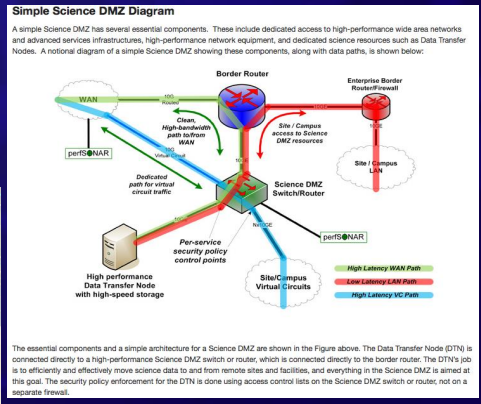

Its not all about bigger pipes

- re-architecting the campus border [longer term ~ re-engineering trust relationships]
- How your campus networking plans fit into the broader cyberinfrastructure strategy for your campus – campus CI Plan
- driving partnerships between scientists and campus IT – **enabling science and education**

* end-to-end perspective on networking and performance

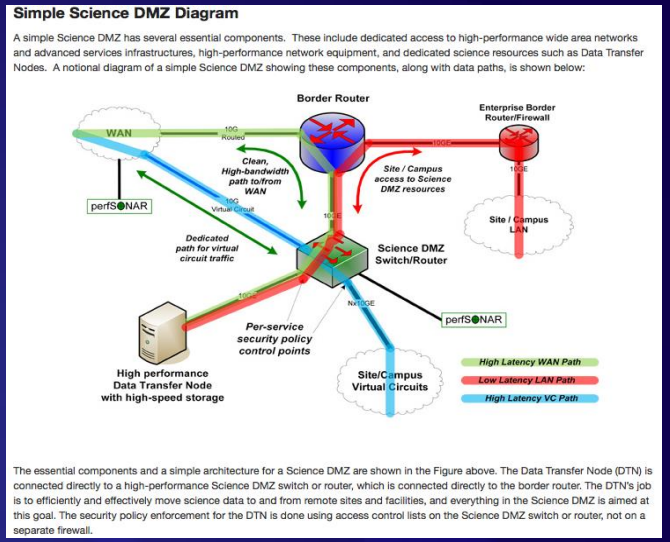

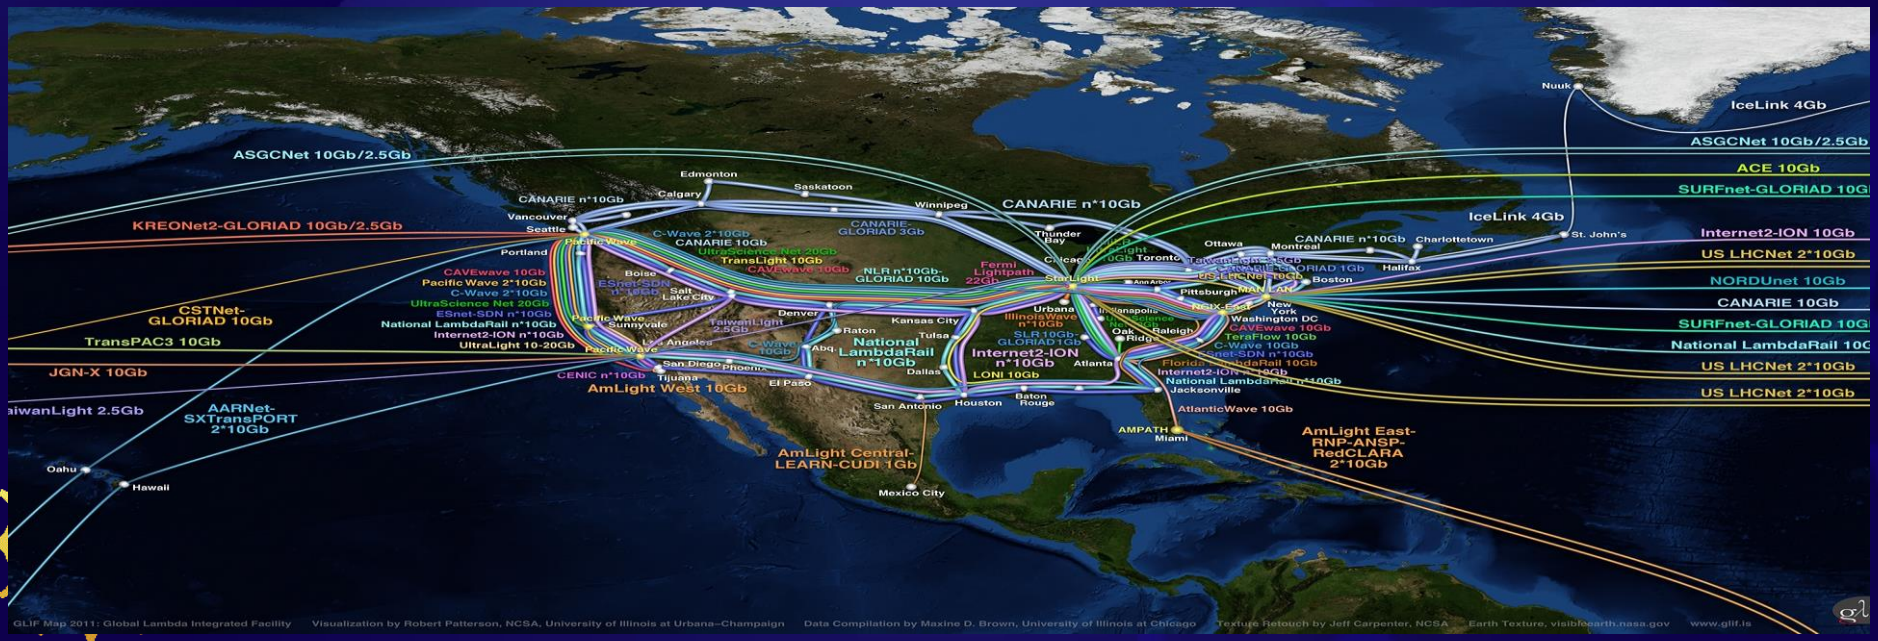

International R&E Networking Support at NSF

- NSF support dates back to 1990 and remains the only continuously funded network infrastructure program at NSF
- Connections funded in IRNC support all science and education network data flows (not just those involving NSF-funded projects)
- Current awards support multiple 10G connections and related activities
- New 5-year awards have been made in 2015 representing a move to 100Gbps+ connections
- Current program investment: ~$40M over 5 years

1st a quick reminder - Global topology r&e networking (NSF supports a part of this, topology shown is outdated)

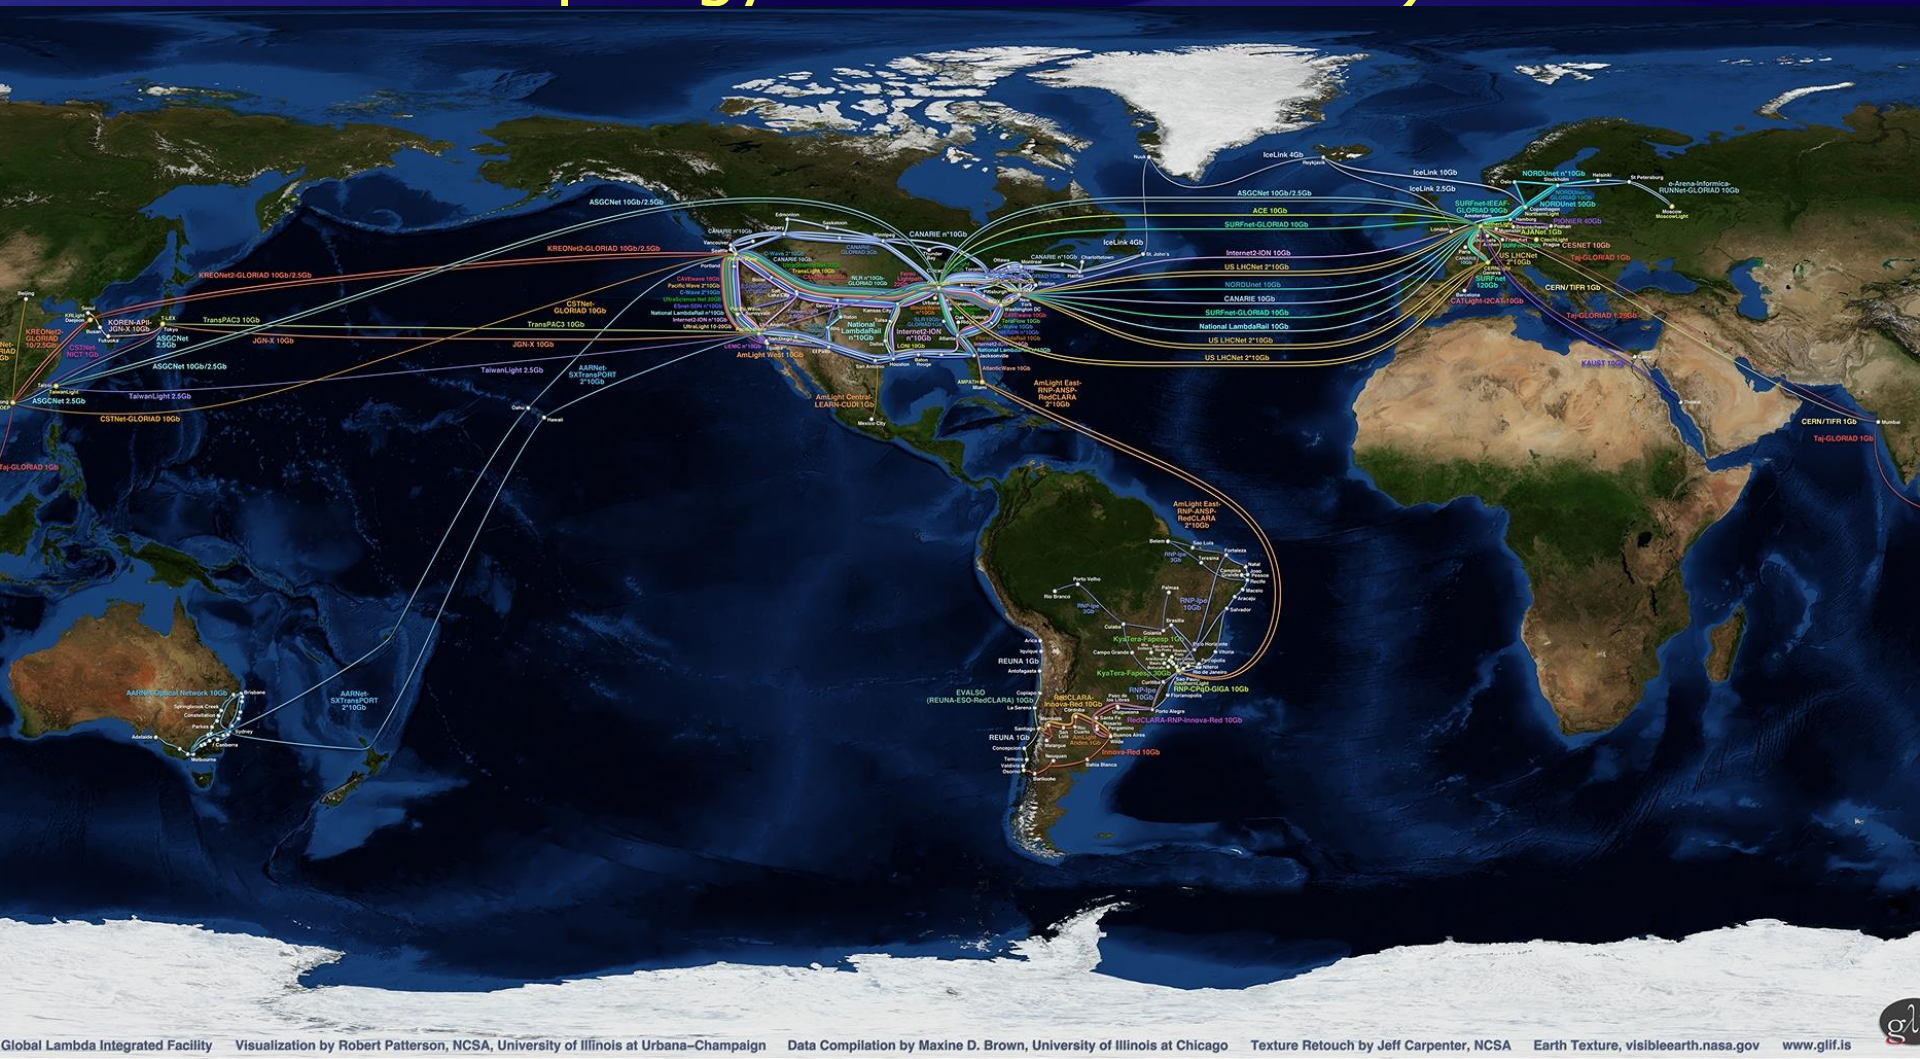

Update Oct 2015

French version now available. Cliquez ici pour la **version française**.

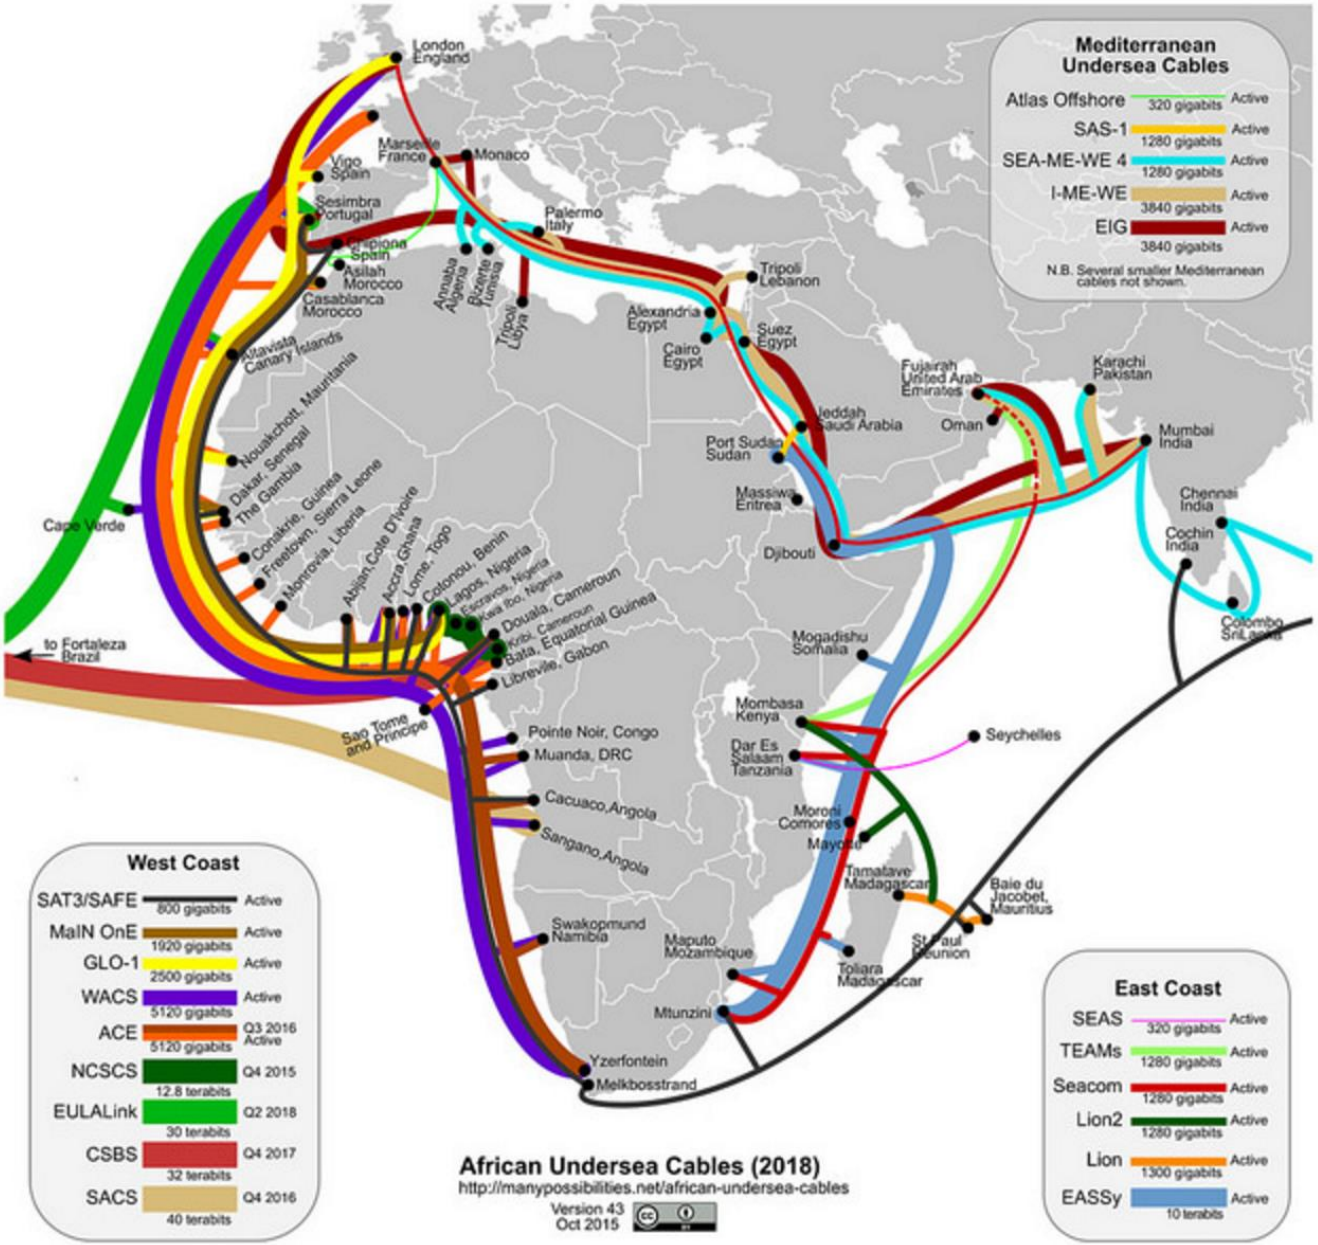

Cyberinfrastructure to Support LSST

Cerro Pachón

2 Day image storage

La Serena Base Facility

Alert Production, Chilean Data Access Center
Long-term storage (copy 1)
36 TFLOPS capacity
35 – 250 PB Storage building over survey

NCSA Archive Center

Nightly Reprocessing, Data Release Production Long-term storage (copy 2), Data Access Center, User Services
120 - 330 TFLOPS
35 – 250 PB Storage

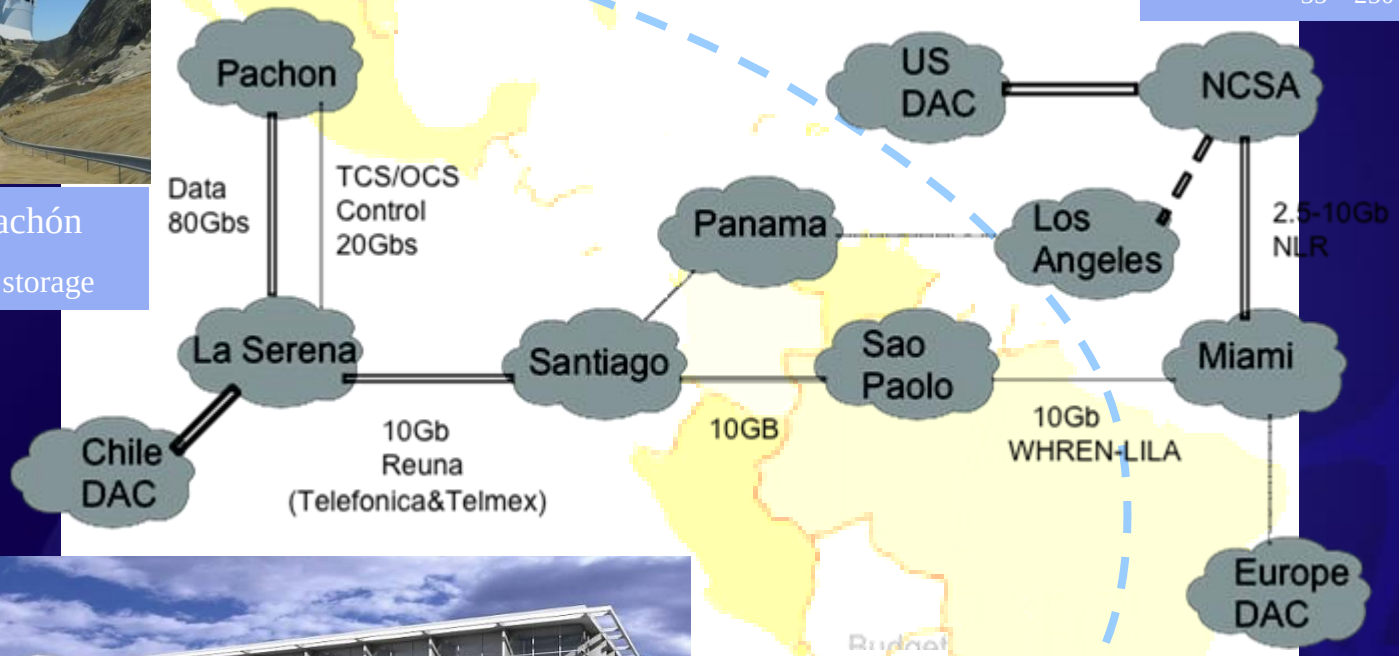

Data / Communications Network

New dedicated 10 x 10Gbs line from Summit to Base
Chilean, International, and US networks in place with LSST capacity
2.5 Gbps steady & 10 Gbps burst

IRNC 2015 Program Areas NSF 14-554

- IRNC: Backbone:
  - Must address high capacity R&E connectivity between the U.S. and either Asia or Americas (Europe and Africa to be addressed separately)
  - Anticipating the 100Gbps+ era
  - up to $1.2M per year for up to 5 years each
- IRNC: RXP:
  - up to $750k per year for 3-5 years
- IRNC: NOC:
  - up to $1M per year for up to 5 years
- IRNC: AMI:
  - up to $1M per year for up to 5 years
- IRNC: Engage:
  - Engagement for Training and Human and Network Capacity Building, Focus on local and regional network design and engineering engagements in developing areas of the world, up to 750k for up to 5 years

IRNC: RXP

- Direct support for international open exchange points in the U.S., recognizing their central role
- A platform for innovation
- SDN support encouraged, exploration allowed, e.g. experimental SDX
- Proposals should address concepts of an open exchange point, connectivity with I2 and Esnet
- Security plan required as supplemental doc
- See solicitation section C. for more details

IRNC: NOC

- A single centralized NOC function for IRNC network infrastructure
- Proposals should address a set of issues –
  - Passive and active measurement and perfSONAR
  - REN-ISAC communication
  - RouteViews
  - Innovative capabilities (e.g. viz, SDN monitoring)
  - Open source software development if applicable
- Interaction with AMI awardee expected but TBD
- See solicitation section D. for more details

IRNC: AMI

- Advanced passive measurement with flow granularity
- Proposals should address:
  - Line-rate measurement architecture/platform w/ 0-impact on user performance
  - Flow-level usage, analysis and reporting
  - Aggregate and summary flow reporting (e.g. AS matrix)
  - IPv4 required, IPv6 and SDN support encouraged
- Proposals should assume for costing the **initial** measurement of 2x40 and 6x10 Gbps links
- Privacy preservation strongly emphasized
- May optionally include other forms of active/passive measurement
- See solicitation section E. for more details

IRNC: Engage

- (F.1) Engagement for Training and Human and Network Capacity Building
  - Focus on local and regional network design and engineering engagements in developing areas of the world
- (F.2) Engagement for Global Coordination
  - Focused on efforts with the global r&e network engineering community
  - IREN principle adherence
  - Expected to represent the best interests of the NSF science community
  - Expected to coordinate fully with other U.S. leaders such as Esnet (i.e. not an exclusive role for U.S. r&e networking)
- All proposals should discuss plans/approaches to work with NSF science communities
- See solicitation section F for more details

IREN

- NSF effort to globally develop and coordinate future strategy for R&E networking
- Guiding principles created with European Commission over 2 years ago (http://fasterdata.es.net/nsf-iren)
  - Open Exchange points allowing bi-lateral peering all layers
  - Open shared transit
  - End-to-end interoperability
  - Close coordination/partnership with r&e networks
  - Resilient design in reducing/eliminating single points-of-failure
  - Emphasizing regional development and aggregation of demand
  - Allowance for different technology architectures, insertion, approaches
  - Open innovation

2015 IRNC Awards

- IRNC: Backbone
  - #1451018 PI: Julio Ibarra, FIU, "IRNC: Backbone: AmLight Express and Protect (ExP)"
  - #1451058 PI: David Lassner, U of Hawaii, "IRNC: Backbone: SXTransPORT Pacific Islands Research and Education Network"
  - #1450904 PI: Jennifer Schopf, Indiana U, "IRNC-BackBone- TransPAC4 - Pragmatic Application-driven International Networking

2015 IRNC Awards: RXP

- #1451024 PI: Julio Ibarra, FIU, "IRNC: RXP: AtlanticWave-SDX: a distributed intercontinental experimental SDX"
- #1450871 PI: Joe Mambretti, NWU, "IRNC: RXP: StarLight SDX A Software Defined Networking Exchange for Global Science Research and Education"
- #1451050 PI: Louis Fox, CENIC, "IRNC:RXP – Pacific Wave Expansion Supporting SDX & Experimentation"

2015 IRNC Awards

- IRNC: NOC
  - #1450934 PI: David Jent, Indiana U, "IRNC:NOC - Global Research Network Operations Center at Indiana University Enabling International Science and Innovation"
- IRNC: Engage
  - #1451045 PI: Steven Huter, U of Oregon, "IRNC: ENgage: Building Network Expertise and Capacity for International Science Collaboration"

2015 IRNC Awards: AMI

- #1540933, PI: Jennifer Schopf, Indiana U, "IRNC: AMI: NetSage - An Open, Privacy-Aware, Network Measurement, Analysis, and Visualization Service
  - Partners include U Hawaii and UC Davis
- #1450996, PI: Yan Luo, UMass Lowell, "IRNC: AMI: Collaborative Research: Software-Defined and Privacy-Preserving Network Measurement Instrument and Services for Understanding Data-Driven Science Discovery"
  - Non-lead collaborative awards:
  - #1450975, PI: Ghinita, Umass Boston
  - #1450997, PI: McGarry, UTEP
  - #1450937, PI: Bumgardner, U Kentucky
- #1450959: PI: Greg Cole, UT Knoxville, "IRNC: AMI: The InSight Advanced Performance Measurement System."

IRNC Links representation

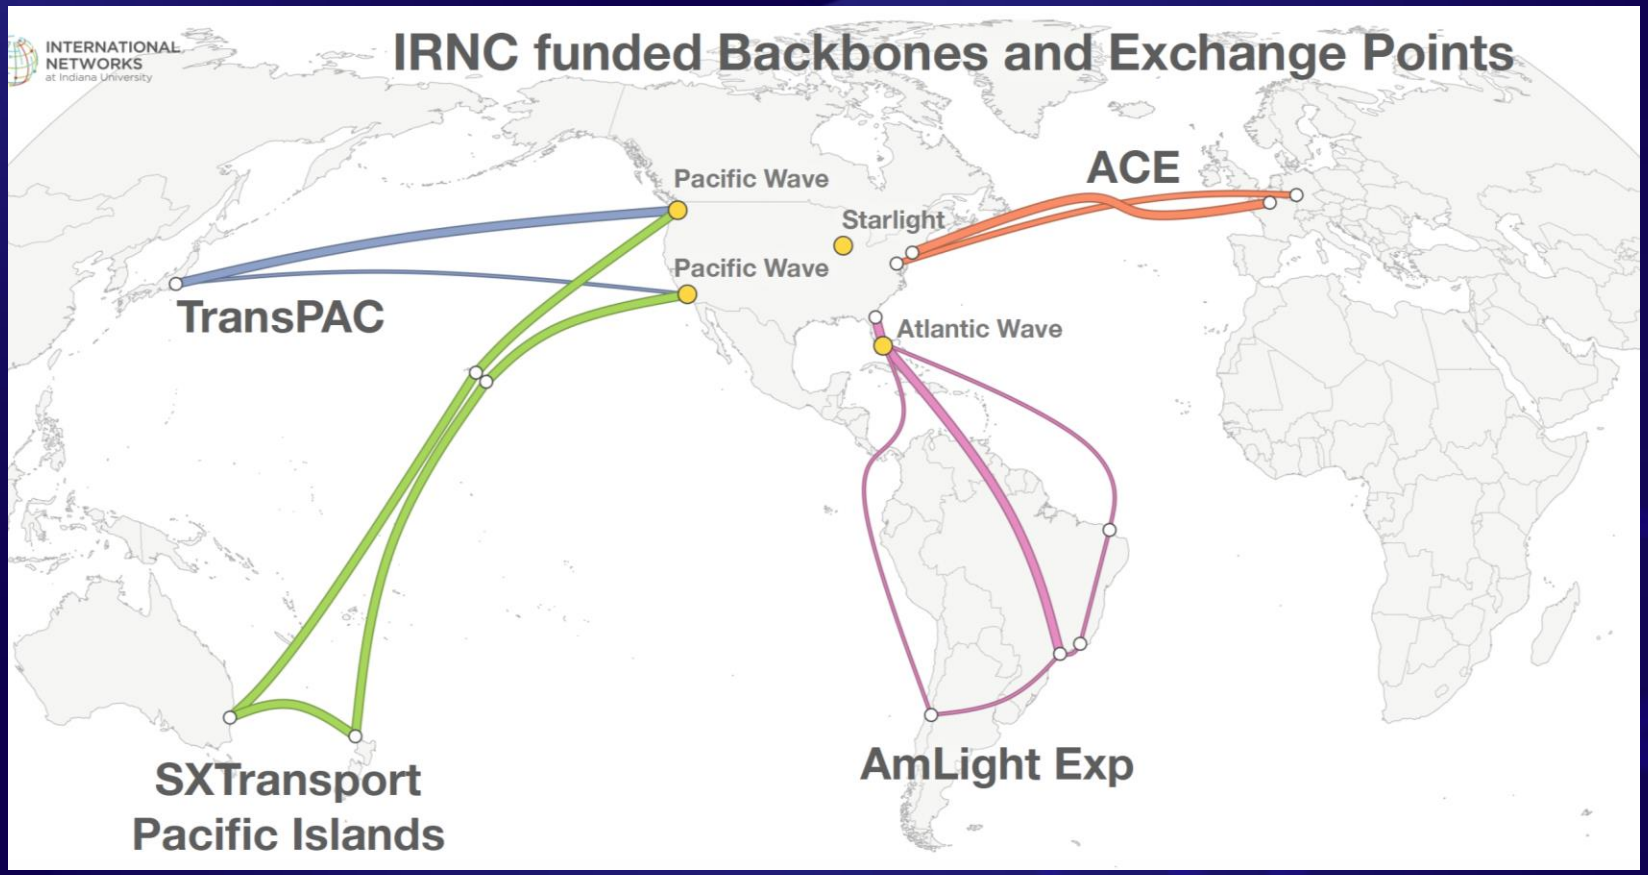

1st Trans-Pacific 100G R&E dedicated connection

- **"TransPac-PacificWave" Tokyo<->Seattle**
- **Partnership bet. TransPac4 and Pacific Northwest Gigapop (link provider)**
- IU Press Release: http://news.iu.edu/releases/iu/2015/10/transpac-pacific-wave.shtml
- PNWGP/CENIC Press Release: http://cenic.org/news/item/pacific-wave-announces-worlds-first-trans-pacific-100-gigabit-re-network
  - "Extension of the Pacific Wave open peering fabric to include a Tokyo point-of-presence, which directly interconnects with the WIDE/T-REX exchange in Tokyo, thus enabling direct R&E peering and exchange across the Pacific"

19

IRNC:Backbone - SX-TransPORT Support & Pacific Islands Research & Education Network (PIREN)

- Provide full domestic support for AARNET's current 2x40Gbps R&E circuits from Australia and New Zealand to the U.S., via Hawaii (including Mauna Kea) with upgrade to 2x100Gbps in 2016
- Continue to foster research and education (R&E) network capacity to interconnect Pacific Islands with each other and the global R&E network fabric by building on previous projects and relationships.
- Opportunistically connect Mauna Kea and Haleakala, sites of major international astronomy observatories
- Collaborate and cooperate with IRNC measurement, NOC, Engagement, and Open Exchange awardees

Key Partners

- AARNet (Australia's NREN)
- REANNZ (New Zealand's NREN)
- Pacific Wave (U.S. West Coast Distributed Open Exchange)
- NSRC (Engagement, Development, Training)

20

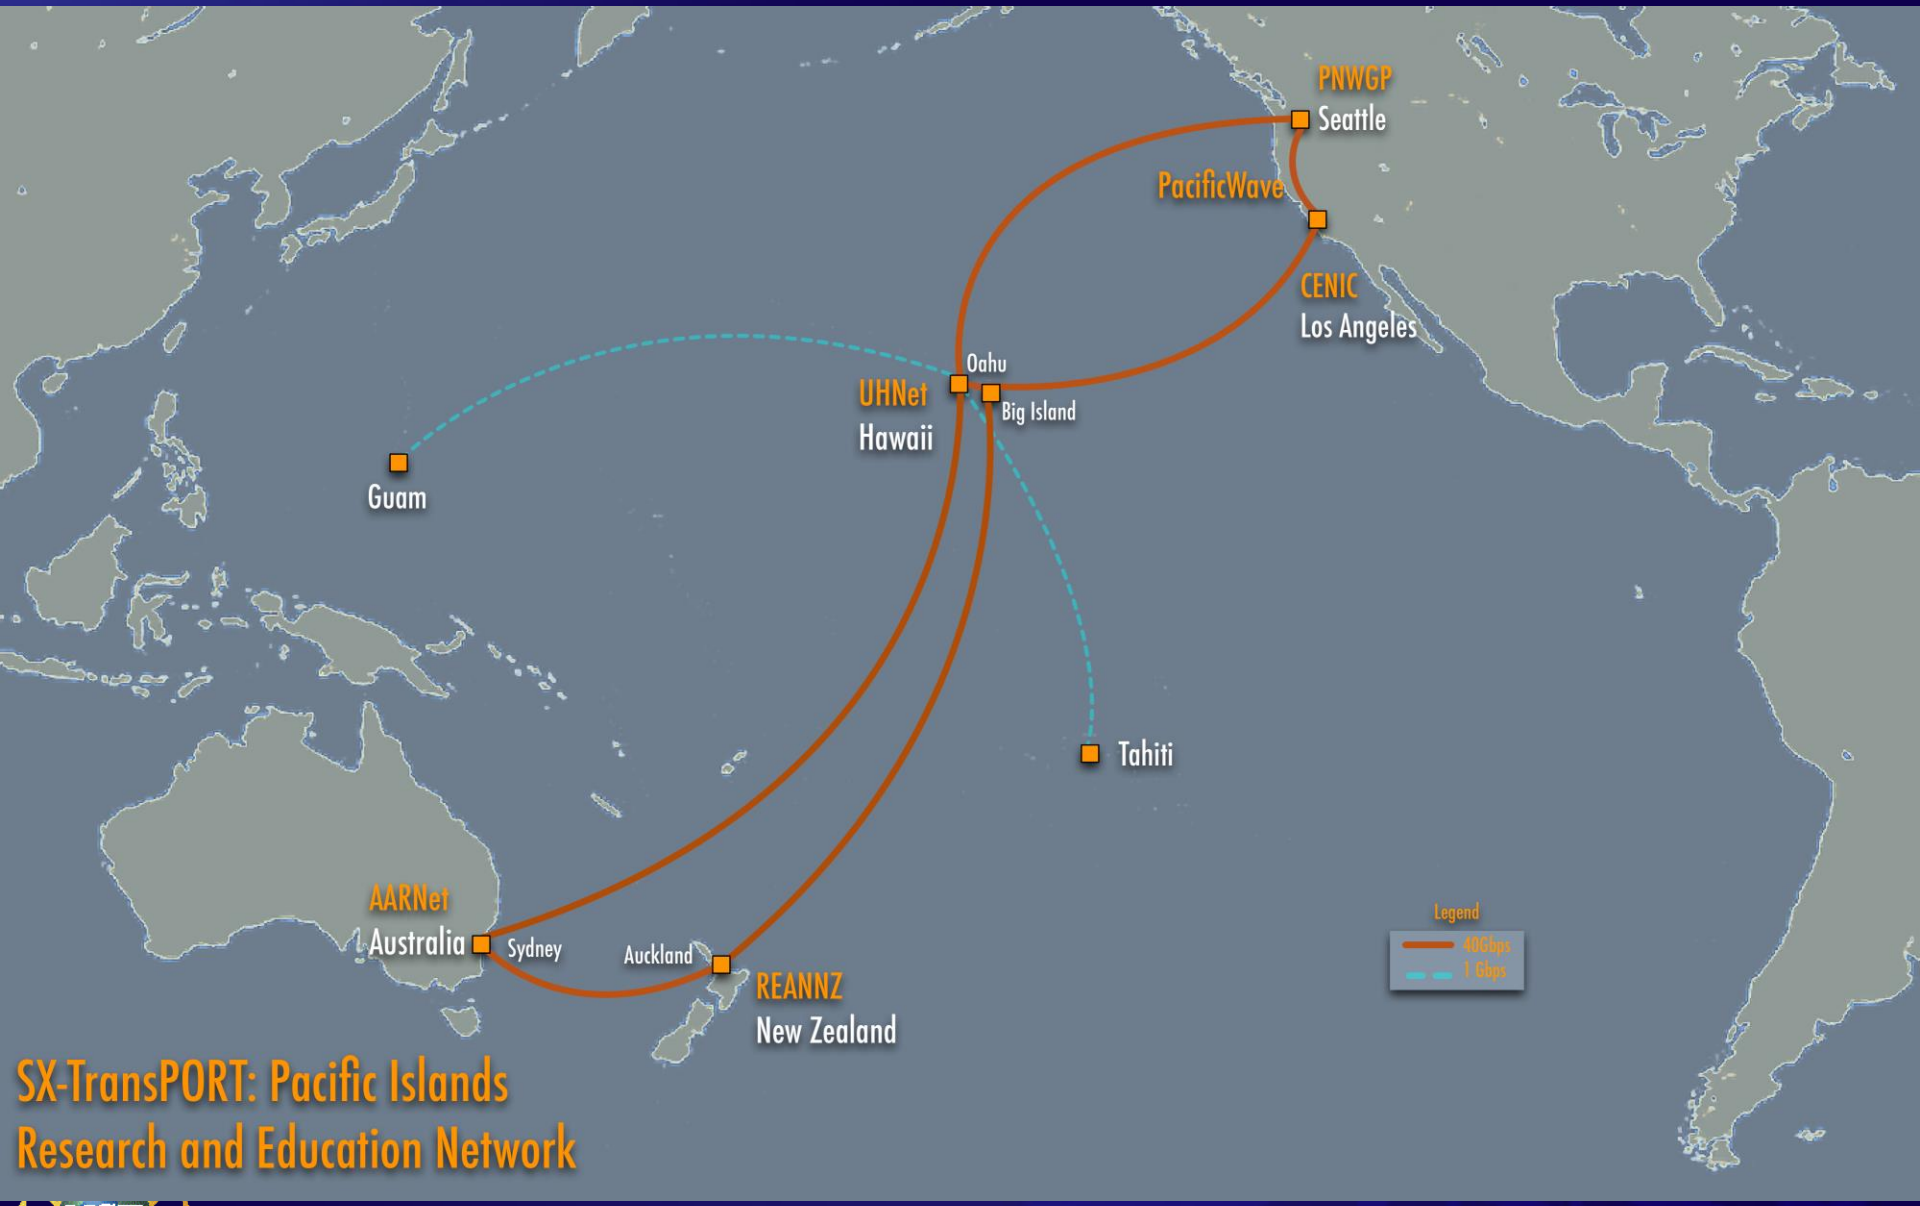

SX-TransPORT: Pacific Islands Research and Education Network

Notional Pacific Islands Research and Education Network (PIREN)

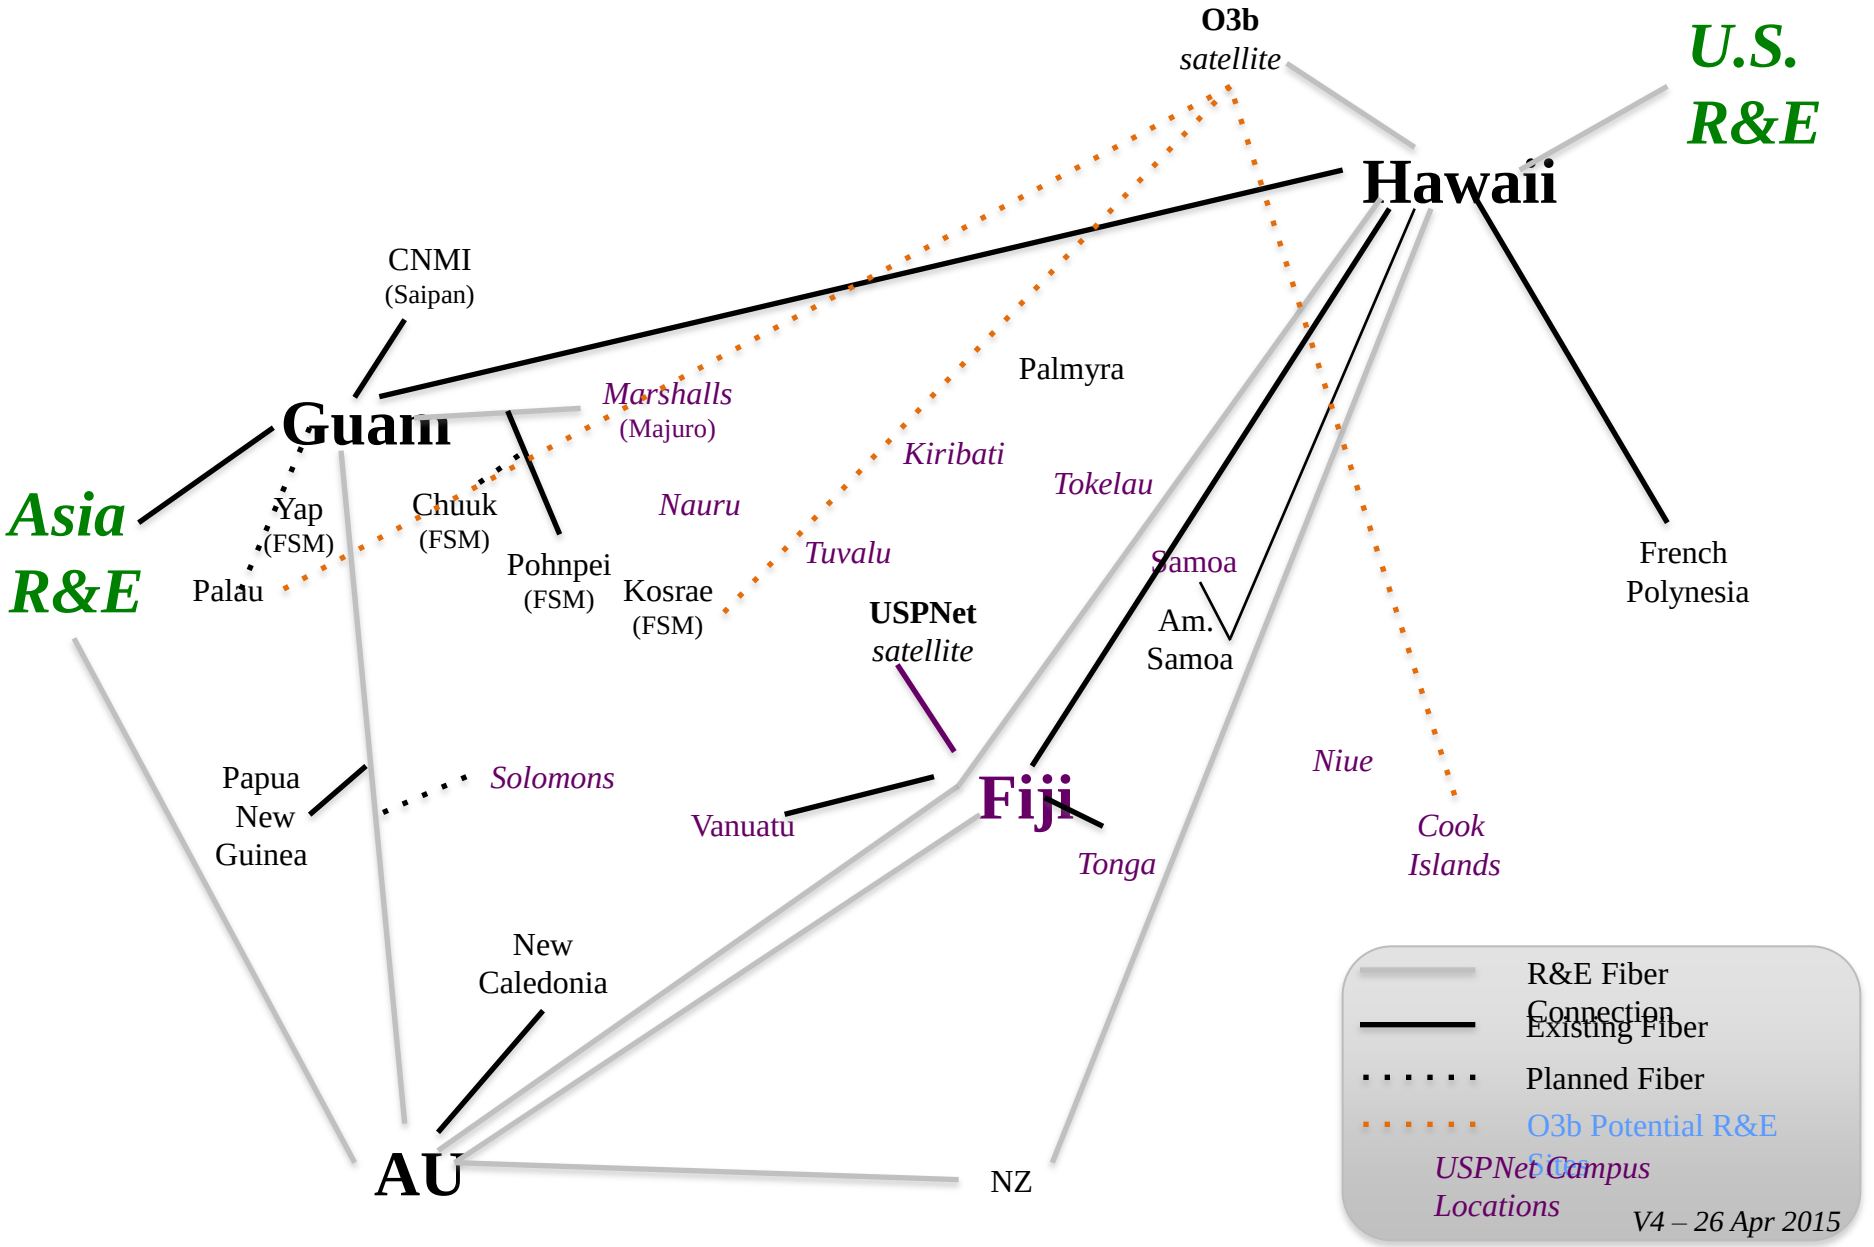

AmLight Today

- 4 x 10G links
  - Two topologies and
  - Two submarine cable systems to increase resilience and support for experimentation
- SDN Ring: Miami-São Paulo, São Paulo-Santiago, Santiago-Miami
  - 20G total capacity
  - Full Openflow 1.0 and network virtualization support
  - Uses Brocade devices
- MPLS Ring: Miami-Fortaleza, Fortaleza-Rio, Rio-São Paulo, São Paulo-Miami
  - 20G total capacity
  - Layer2 support via L2VPN
  - Uses Juniper devices
- Mutual redundancy between SDN and MPLS rings

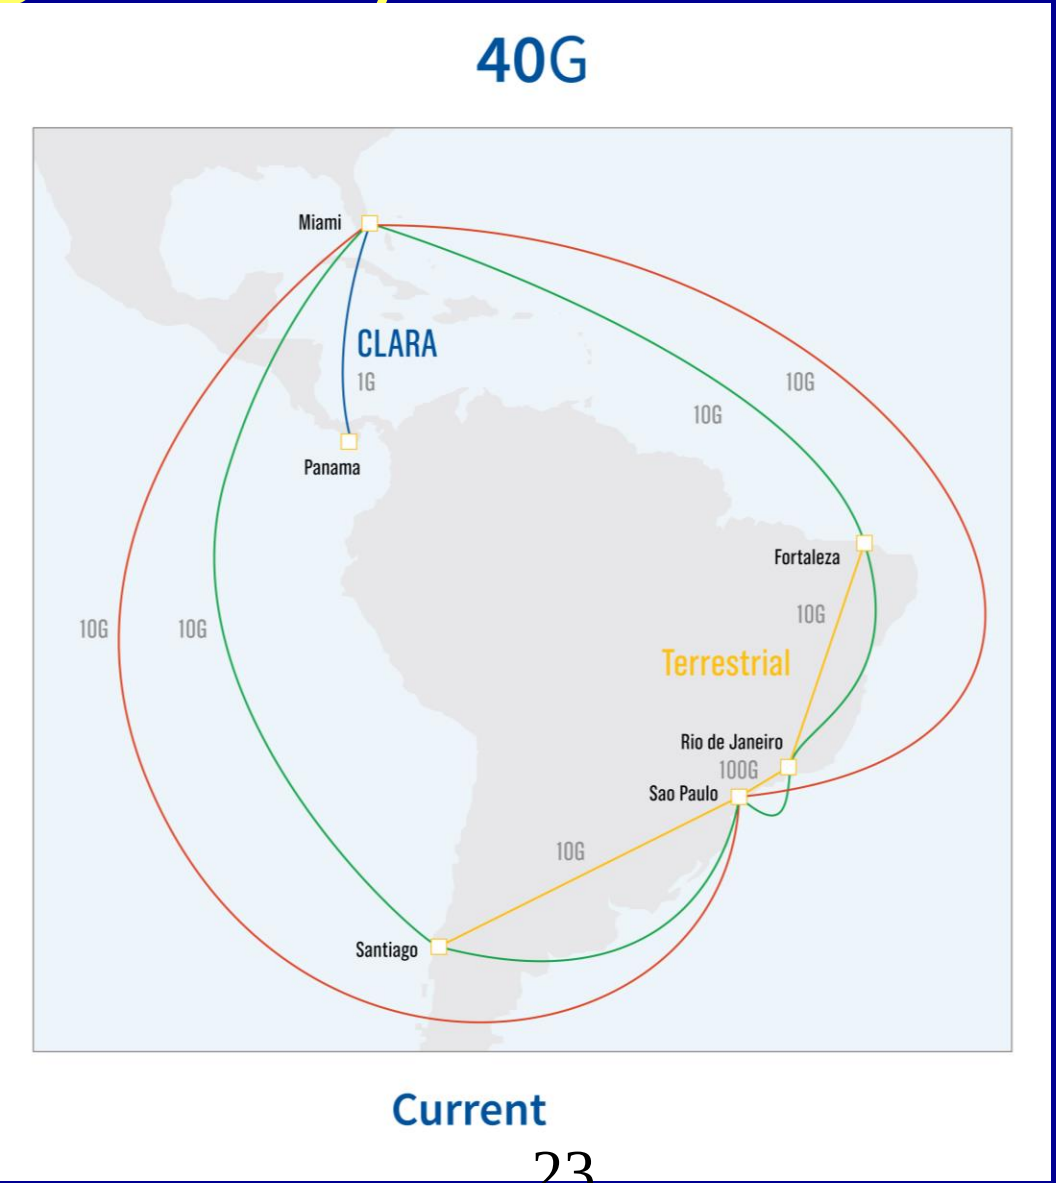

AmLight 2015-2017

- OpenWave 100G alien wave
  - U.S., Brazil, Latin America
  - Experimentation is initial focus
  - In the AmLight SDN domain
  - What we learn will enable our next 20 years
- 100G to AL2S, Miami-Jacksonville is operational
- 140G aggregate capacity using spectrum and leased circuits

AmLight Express and Protect (ExP) 2018-2031

- AmLight Express:
  - 300GHz of spectrum: Santiago-São Paulo, and São Paulo-Miami
  - Spectrum to be configurable by RENs to meet user/application requirements
- AmLight Protect:
  - 40G leased capacity ring
  - Miami, São Paulo, Santiago, Panama City, Miami
  - AMPATH, Southern Light, REUNA, and RedCLARA operated
- Potential for unprecedented regional resilience for U.S.-Latin America, and U.S.-Europe connectivity, supporting global science research

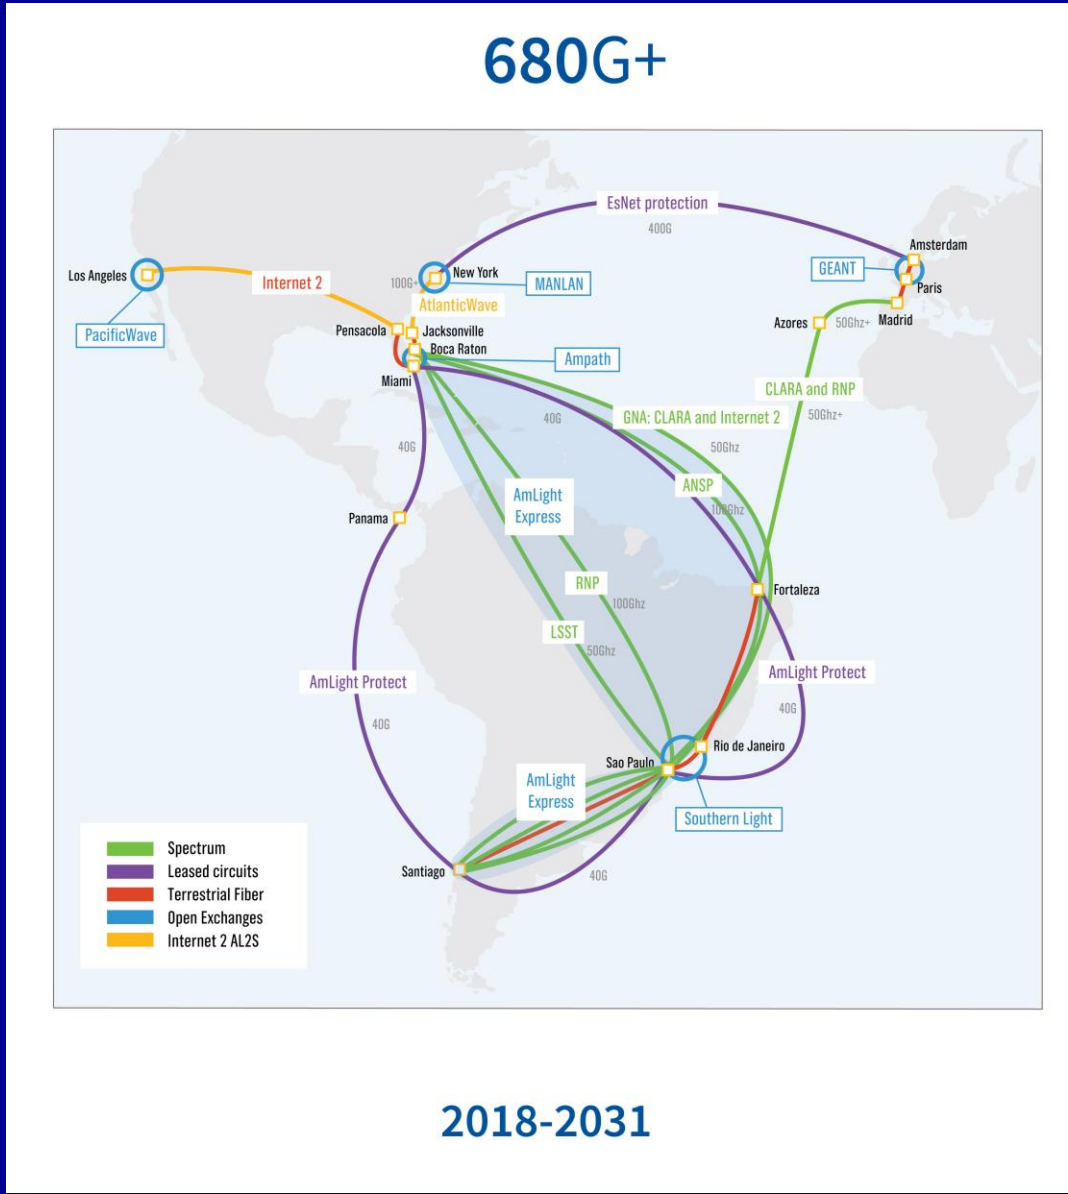

NetSage example

- Slide excerpt from TP4 talk at APAN 2015

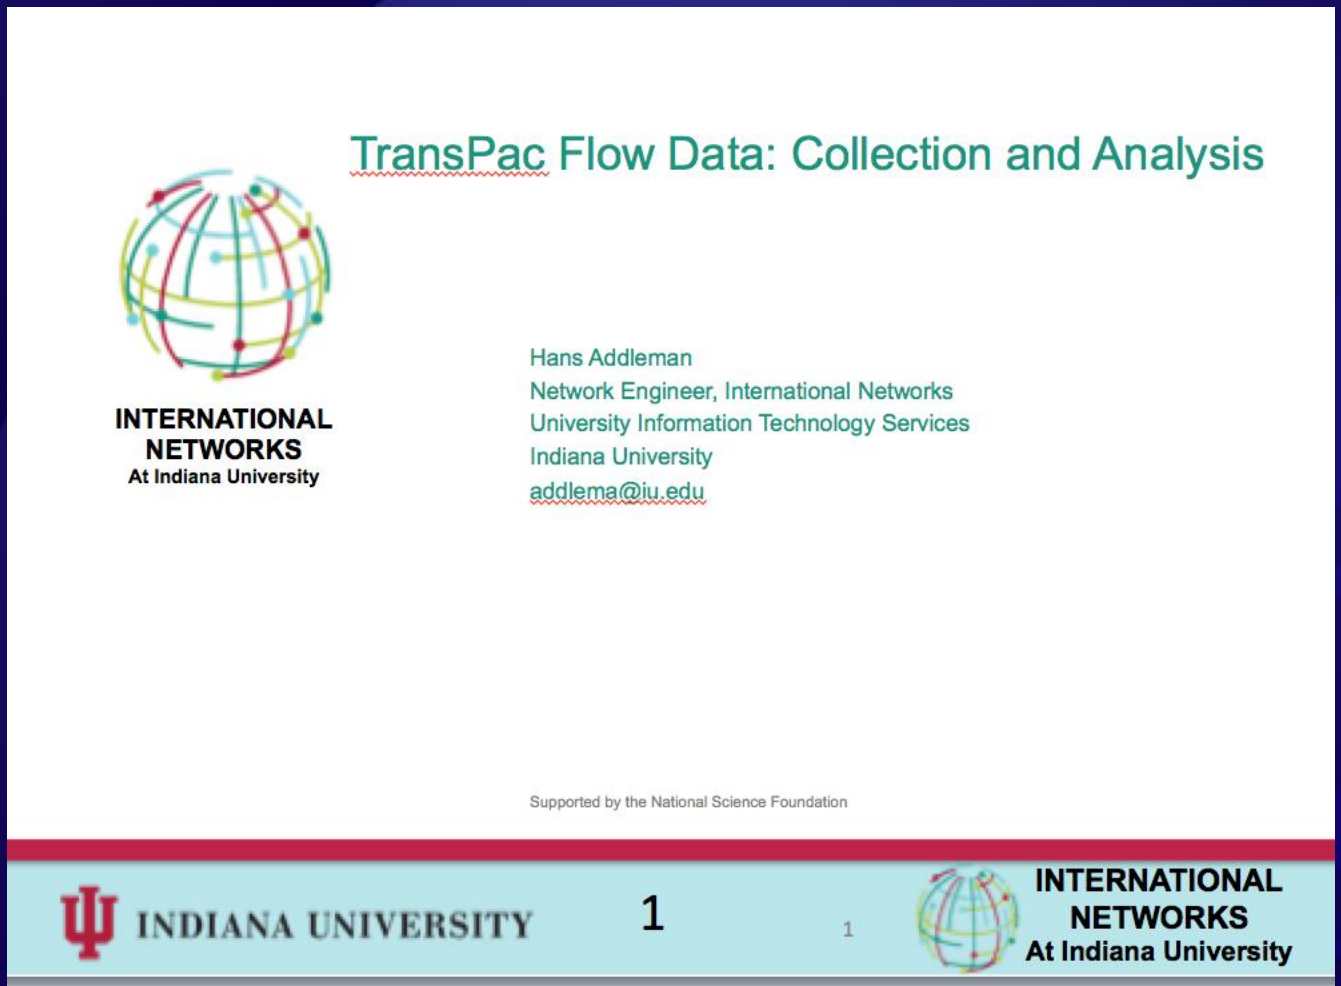

April 2015 Large Transfer

- 78% of the top 10 traffic in April was between 2 institutions.
  - The University of Tokyo
  - National Center for Atmospheric Research
  - Inbound towards NCAR

- Large file transfer between 2 institutions
- Can we help make this transfer better?
- Gives us a clue to what type of researchers are using our network.
- Reach out to other researchers in same field.
- Smaller transfers but still large noticed in March as well.

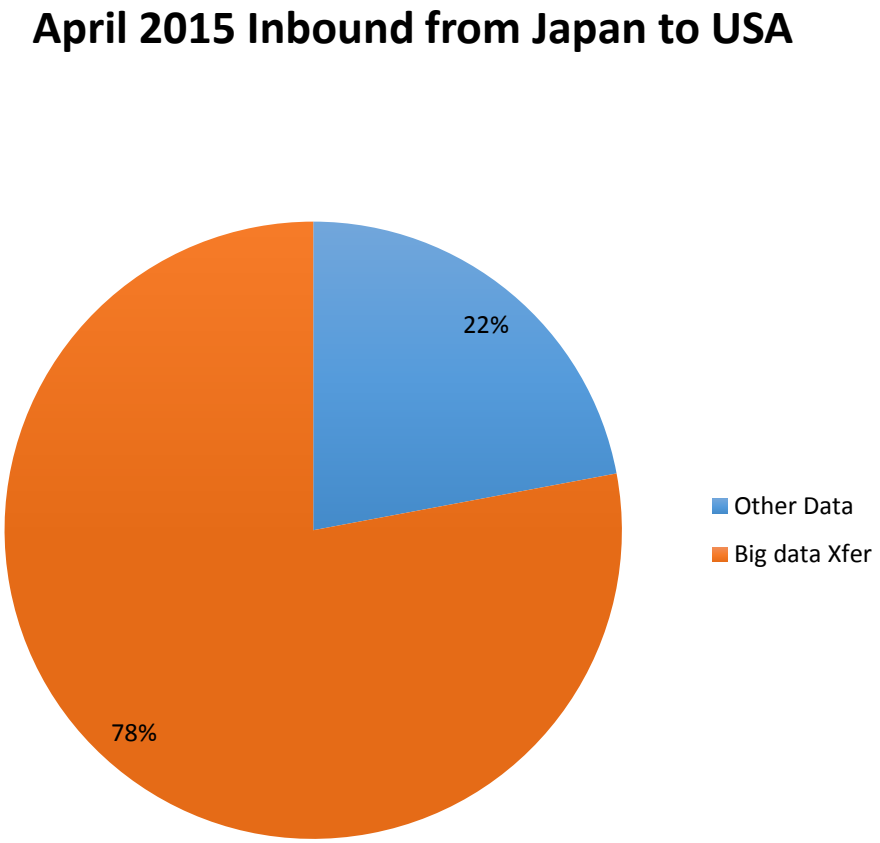

RXP Awards

- Starlight ->
- Next 2 slides represent CENIC and AtlanticWave-SDX

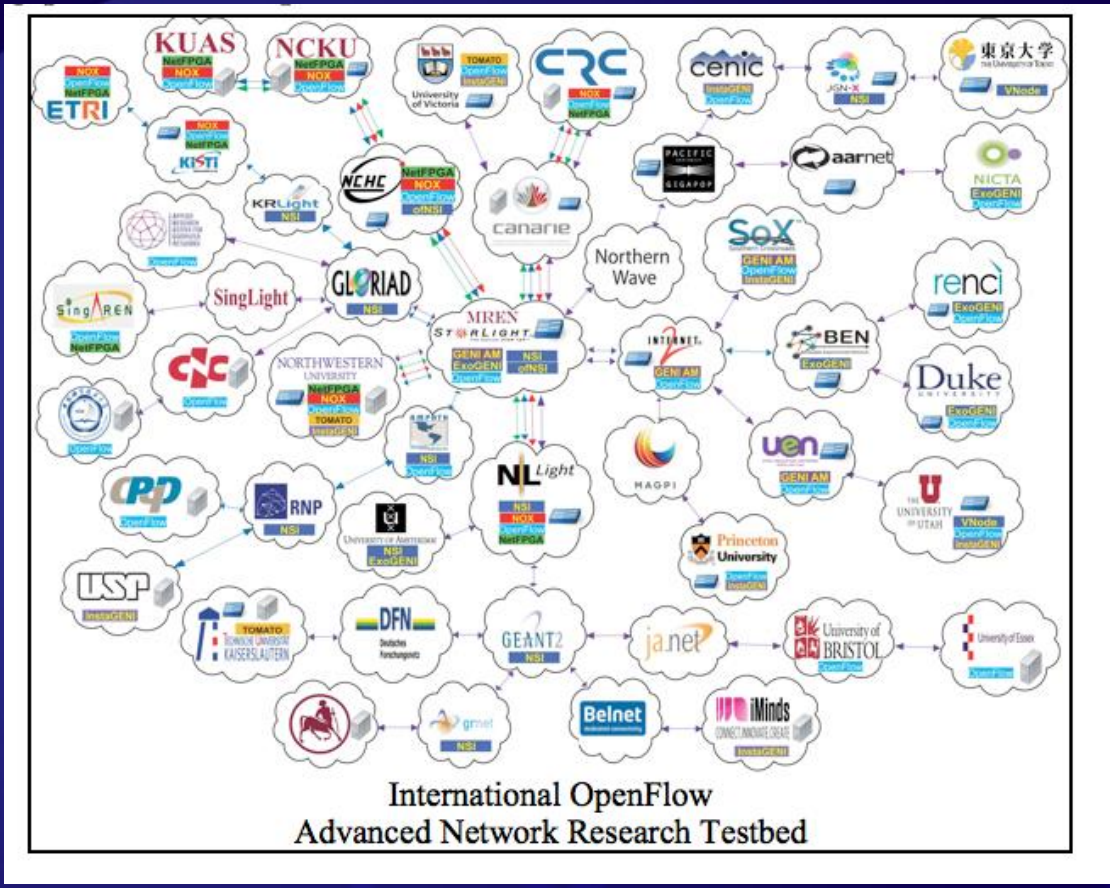

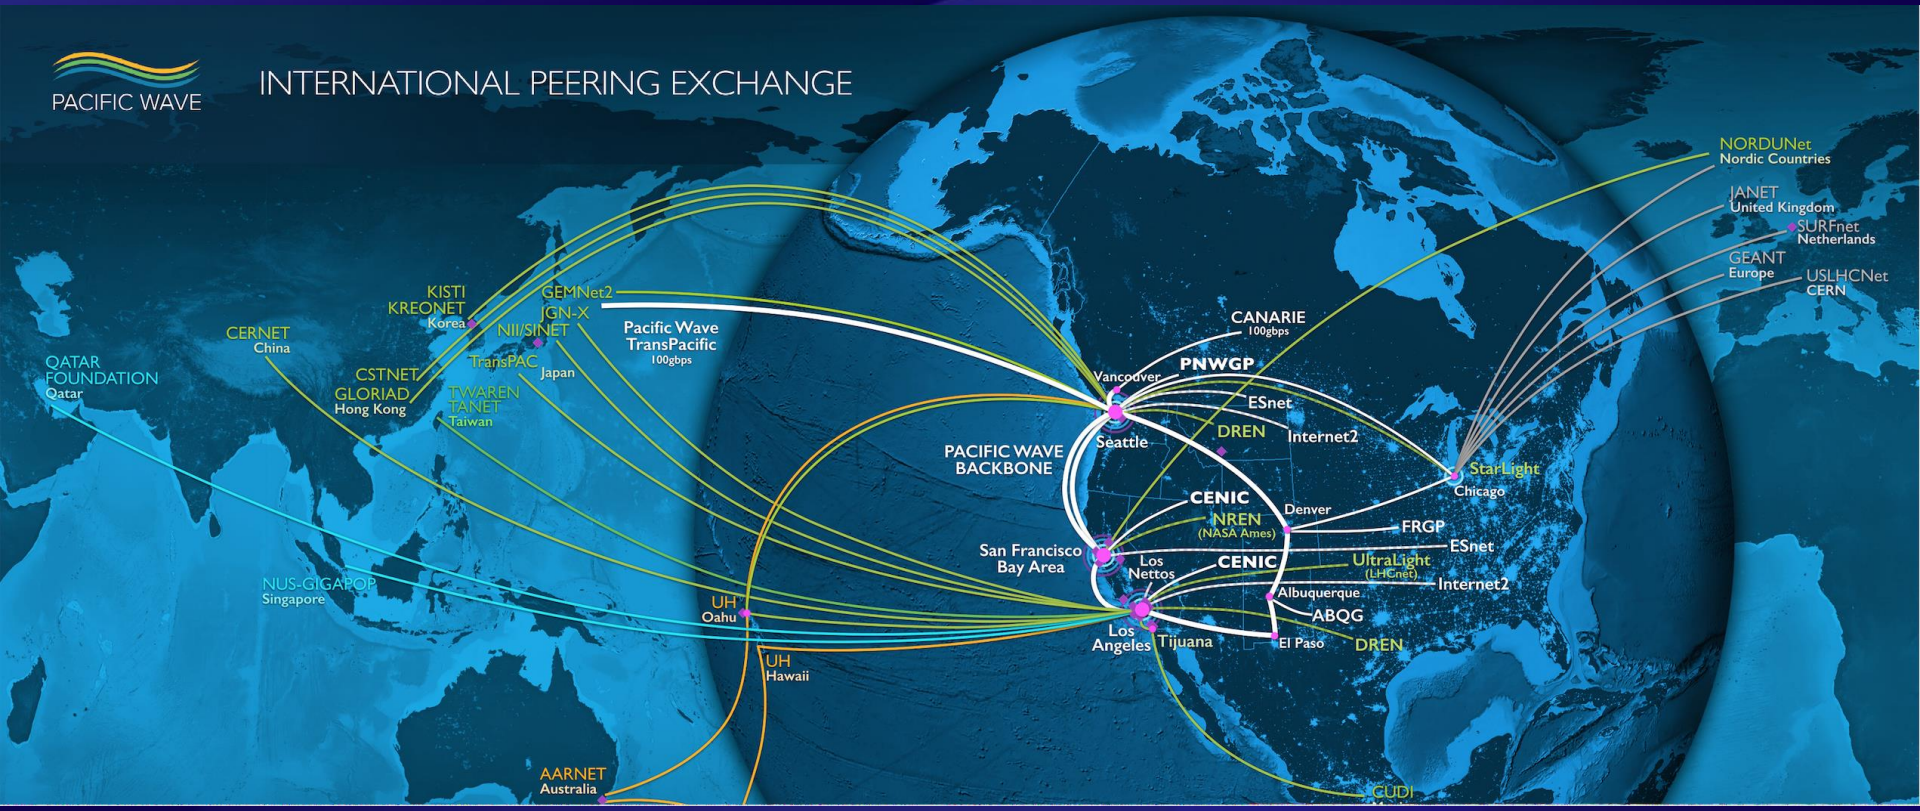
PACIFIC WAVE
INTERNATIONAL PEERING EXCHANGE
NORDUNet
Nordic Countries
JANET
United Kingdom
SURFnet
Netherlands
GEANT
Europe
USLHCNet
CERN
KISTI
KREONET
Korea
GEMNet2
JGN-X
NII/SINET
Japan
TransPAC
CERNET
China
QATAR
FOUNDATION
Qatar
CSTNET
GLORIAD
Hong Kong
TWAREN
TANET
Taiwan
Pacific Wave
TransPacific
100gbps
CANARIE
100gbps
PNWGP
Vancouver
ESnet
DREN
Internet2
Seattle
PACIFIC WAVE
BACKBONE
StarLight
Chicago
CENIC
Denver
NREN
(NASA Ames)
FRGP
ESnet
San Francisco
Bay Area
Los
Nettos
CENIC
UltraLight
(LHCnet)
Internet2
Albuquerque
ABQG
NUS-GIGAPOP
Singapore
UH
Oahu
Los
Angeles
Tijuana
El Paso
DREN
UH
Hawaii
AARNET
Australia
CUDI

Innovation Component

- Three options of deployment for SDX:
  - Option 1:
    - Single SDX controller managing entire IXP switch fabric
  - Option 2:
    - Intermediate slice manager
      - allows individual controllers to be handed a slice of network resources
      - While isolating resources from others
      - Most practical approach in near term
  - Option 3:
    - Creates a hierarchy of controllers with a local controller at each IXP managed by a separate higher-level controller

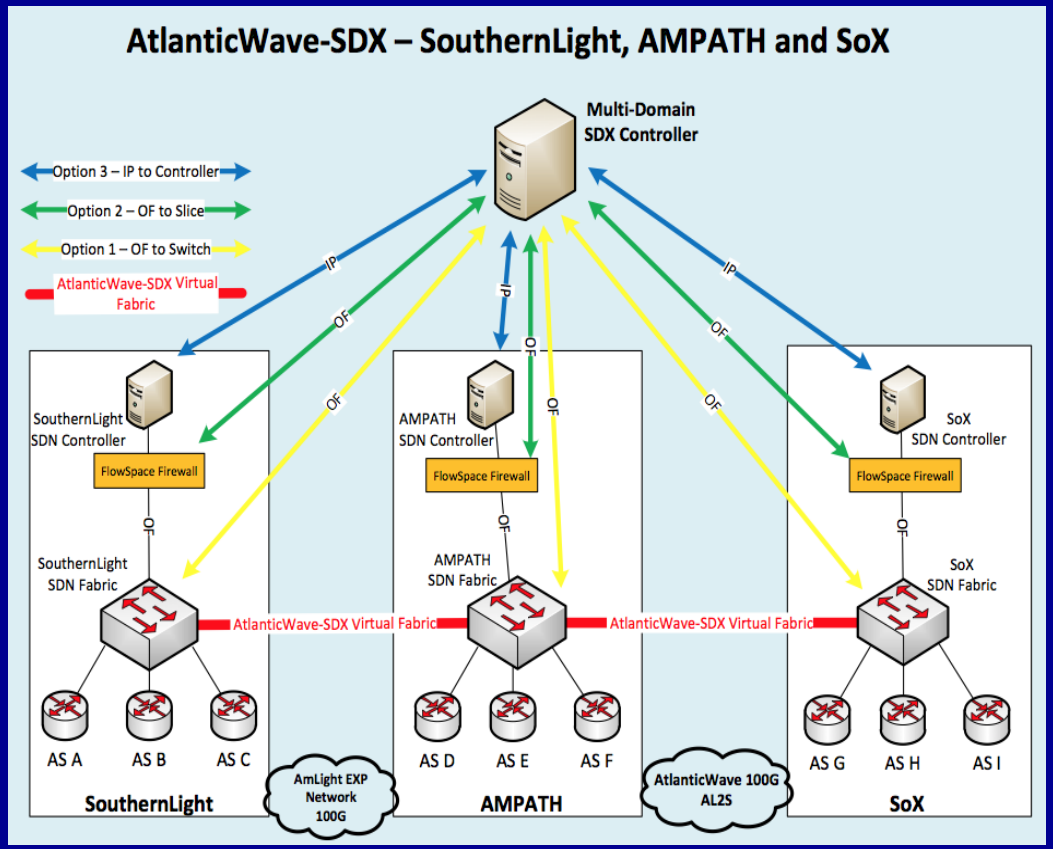

IRNC NOC (IU, PI: Dave Jent)

IRNC NOC: One Point Of Contact for IRNC Operations

The International Research Network Connections Network Operations Center (IRNC NOC) serves as a **cooperative point of contact and communications** for IRNC network management, providing **consolidated network monitoring, reporting, and operational visibility for the IRNC program**. The IRNC NOC facilitates a single set of operational expectations for all IRNC funded infrastructure programs; this enables greater availability of IRNC infrastructure and improves results in troubleshooting multi-domain network issues. A central data repository created by the IRNC NOC provides critical operational information; monitoring data and performance metrics in support of NSF funded science and research.

What we will do

Create the IRNC NOC presence:

- Dedicated instances of an IRNC ticket system, documentation system, monitoring system, and telephone lines.
- Provide IRNC infrastructure projects with a service desk that is available 24x7x365 to serve as a single point of contact to report and detect problems related to IRNC infrastructure projects, provide support for coordination and communications among the participants.
- Develop processes and documentation needed including notification, escalation and reporting processes and problems reporting.
- Create database of operational contact data and ops status including planned outages
- Turn-up of GlobalNOC systems for IRNC: NOC (NOC web presence, e-mail, ticketing, reporting, etc.)
- Develop and implement technical integration with IRNC: AMI projects for NOC monitoring/alerting
- Integrate RouteViews data into NOC monitoring systems to alert on metrics such as key route changes

What we will do

- **Provide:** End-to-end Performance Troubleshooting
  - Provide end-to-end performance engagement services for scientists working with international collaborators.
  - Direct assistance at solving performance problems.
  - Measure success of the effort via researcher satisfaction and effect technical change through recommendations.
  - Plan/implement perfSONAR enhancements for inter-domain performance troubleshooting

The NSRC Model

- Technical training and human resource development activities
- Direct engineering assistance
- Participatory development