INTERNET2 TECH ex 24 BOSTON DEC 9-13

Dynamic Research Connections Across Borders

North American Research and Education Exchange Collaboration (NA-REX)

December 10, 2024

Panelists

- **Jonah Keough**
  Managing Director
  Pacific Wave
- **Chris Wilkinson**
  Senior Director, Network Infrastructure and Operations
  Internet2
- **Hans Addleman**
  Technical Director
  IN@IU
- **Kate Robinson**
  Network Engineer
  ESnet
- **Julio Ibarra**
  Research Professor, Technology Augmented Research
  FIU (AMPATH-AmLight)
- **Thomas Tam**
  Chief of Network Engineering
  CANARIE (MOXY)

Agenda

- NA-REX Overview
- Internet2 and MANLAN/WIX
- Pacific Wave
- ESnet
- International Networks @ Indiana
- AmLight
- MOXY
- Recent Collaborations Supporting Science at SC24
- What's Next

NA-REX Program Elements

Collaboration, Coordination

Improve coordination between all North American international exchange point and capacity operators

Regular meetings for project coordination, system planning, grant opportunities, communications

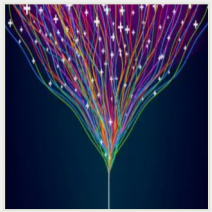

> 400 Gbps GXP Interconnection

Provide dedicated bandwidth between all exchange points

A minimum of 400 Gbps today, 2x or 3x 400 Gbps coming online (part of footprint) early in 2025

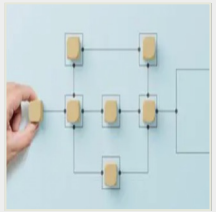

Uniform Provisioning Standards

Create & maintain common operating principles

Currently NSI based to support AutoGOLE / SENSE

Measurement, Analysis, Visualization

Create & maintain centralized repository of utilization data which is persistent and easy to consume

Currently based on NetSage

NA-REX Use Cases

Network Research

Leveraging the shared platform to support research into network technologies

SDN, DTNs, Interdomain SR, Hecate, P4 - what's next?

Testbed Development

Supporting testbeds for science with dedicated networks

Mission Specific
Overlay Networks

Distributed Hybrid Quantum Computing with PQC/QKD secured links

Transit for Specific Applications

Enabling advanced science researchers through advanced connectivity

Data Intensive Science
SKA
NREN Backup Paths
& more

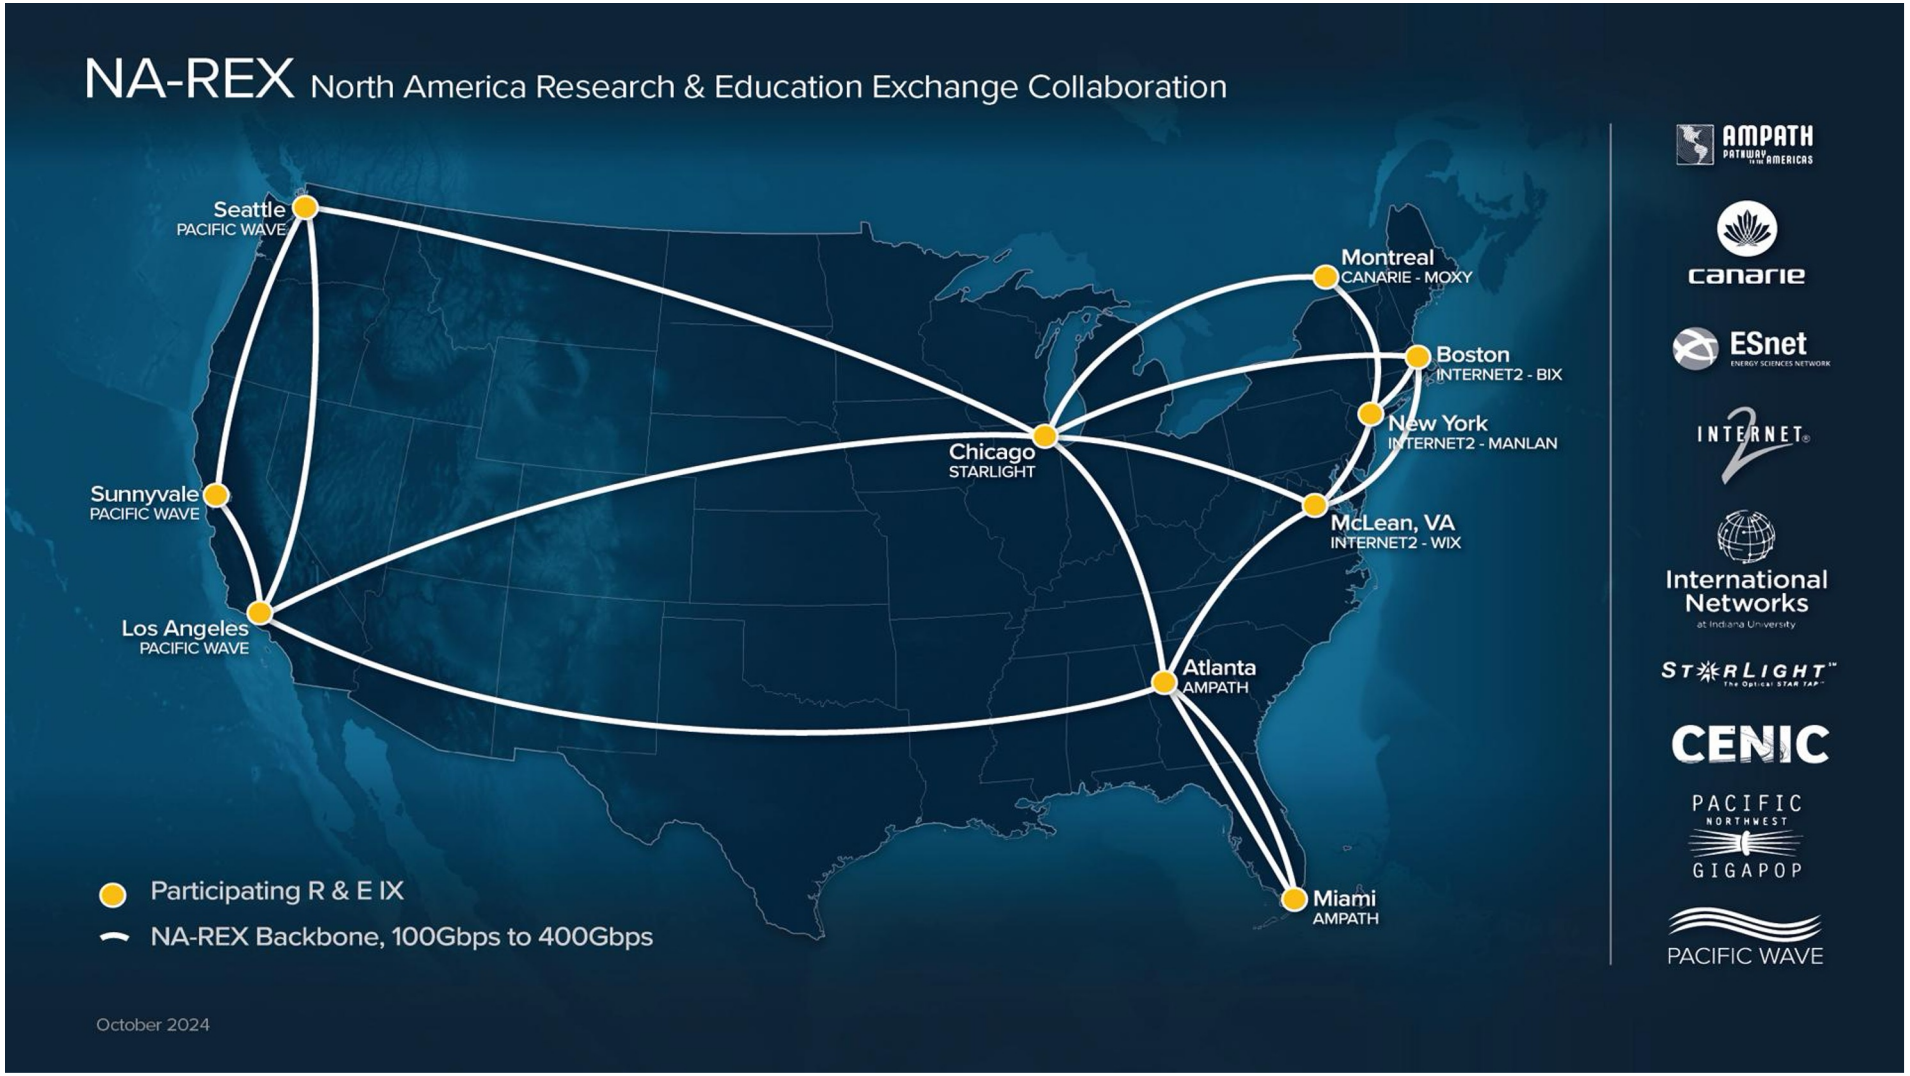
NA-REX North America Research & Education Exchange Collaboration
Seattle
PACIFIC WAVE
Sunnyvale
PACIFIC WAVE
Los Angeles
PACIFIC WAVE
Chicago
STARLIGHT
Montreal
CANARIE - MOXY
Boston
INTERNET2 - BIX
New York
INTERNET2 - MANLAN
McLean, VA
INTERNET2 - WIX
Atlanta
AMPATH
Miami
AMPATH
Participating R & E IX
NA-REX Backbone, 100Gbps to 400Gbps
October 2024
AMPATH
canarie
ESnet
INTERNET2
International Networks
STARLIGHT
CENIC
PACIFIC NORTHWEST GIGAPOP
PACIFIC WAVE

Internet2 Exchange Point Fabric

- R&E Open Exchange on East Coast of US

  Boston (BIX), New York (MAN LAN), and Washington (WIX)

- Operate independently from Internet2 network
- Currently Layer 2 only
- Currently working to enable AutoGOLE/SENSE via NSI / SuPA (SURF ultimate Provider Agent)
- Based on Arista 7280PR3K-24 switches (OSFP)

  moving to Arista 7280DR3K-24 switches (QSFP28-DD)

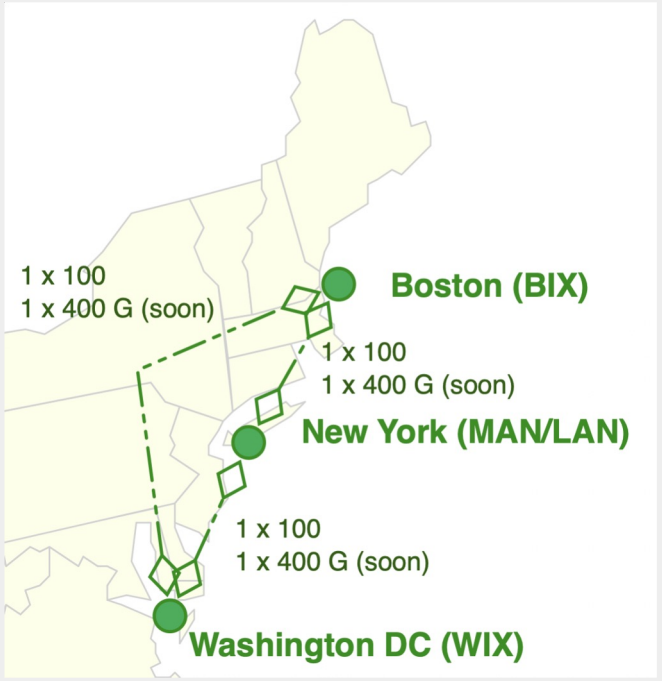

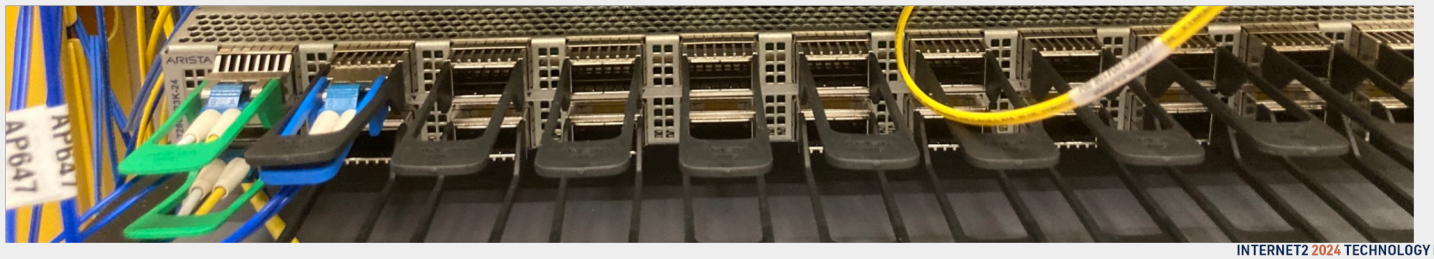

Internet2 Exchange Point Fabric

Primarily Support Advanced North Atlantic (ANA)

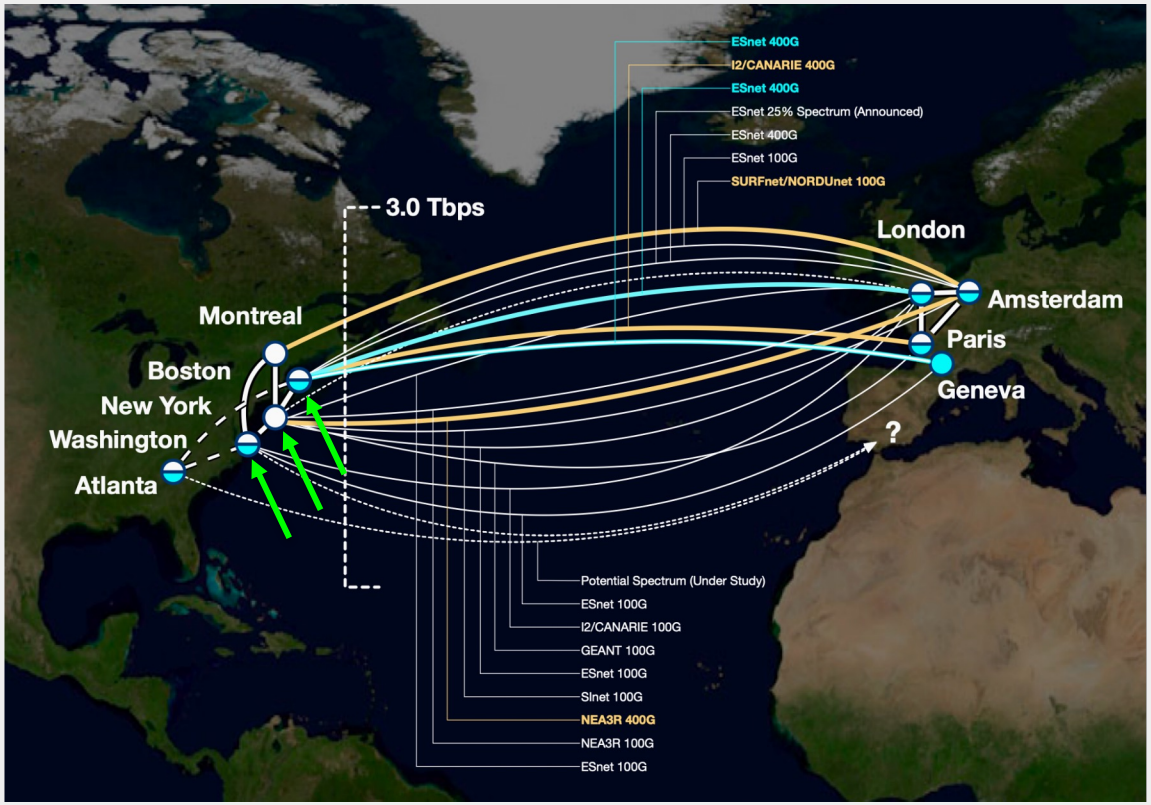

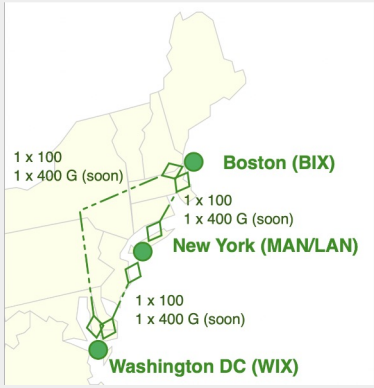

NA-REX Transport provided by Internet2 Line System

- 15,000 km of terrestrial links dedicated to NA-REX
- Based primarily on Ciena WL5e Modules; transitioning to coherent pluggable optics for shorter spans

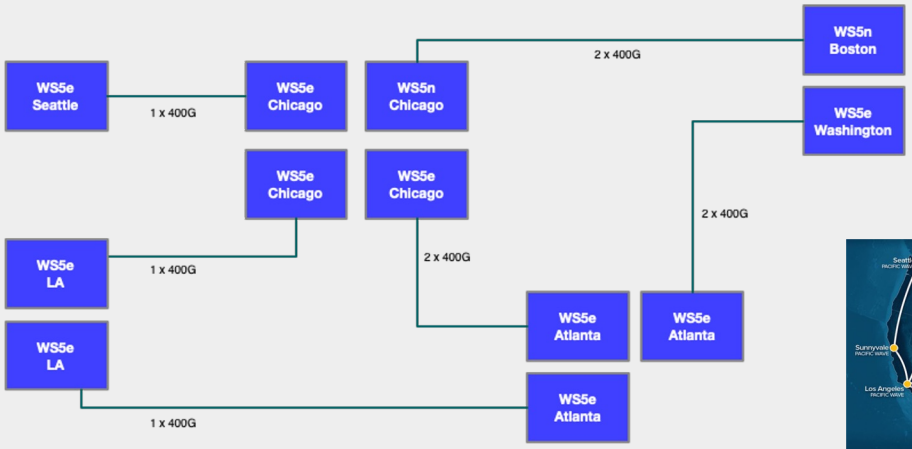

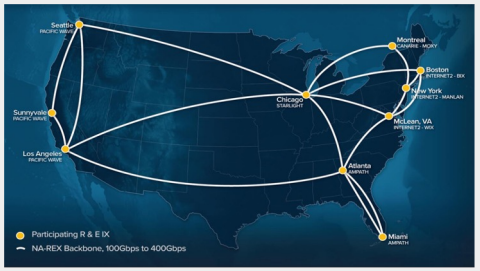

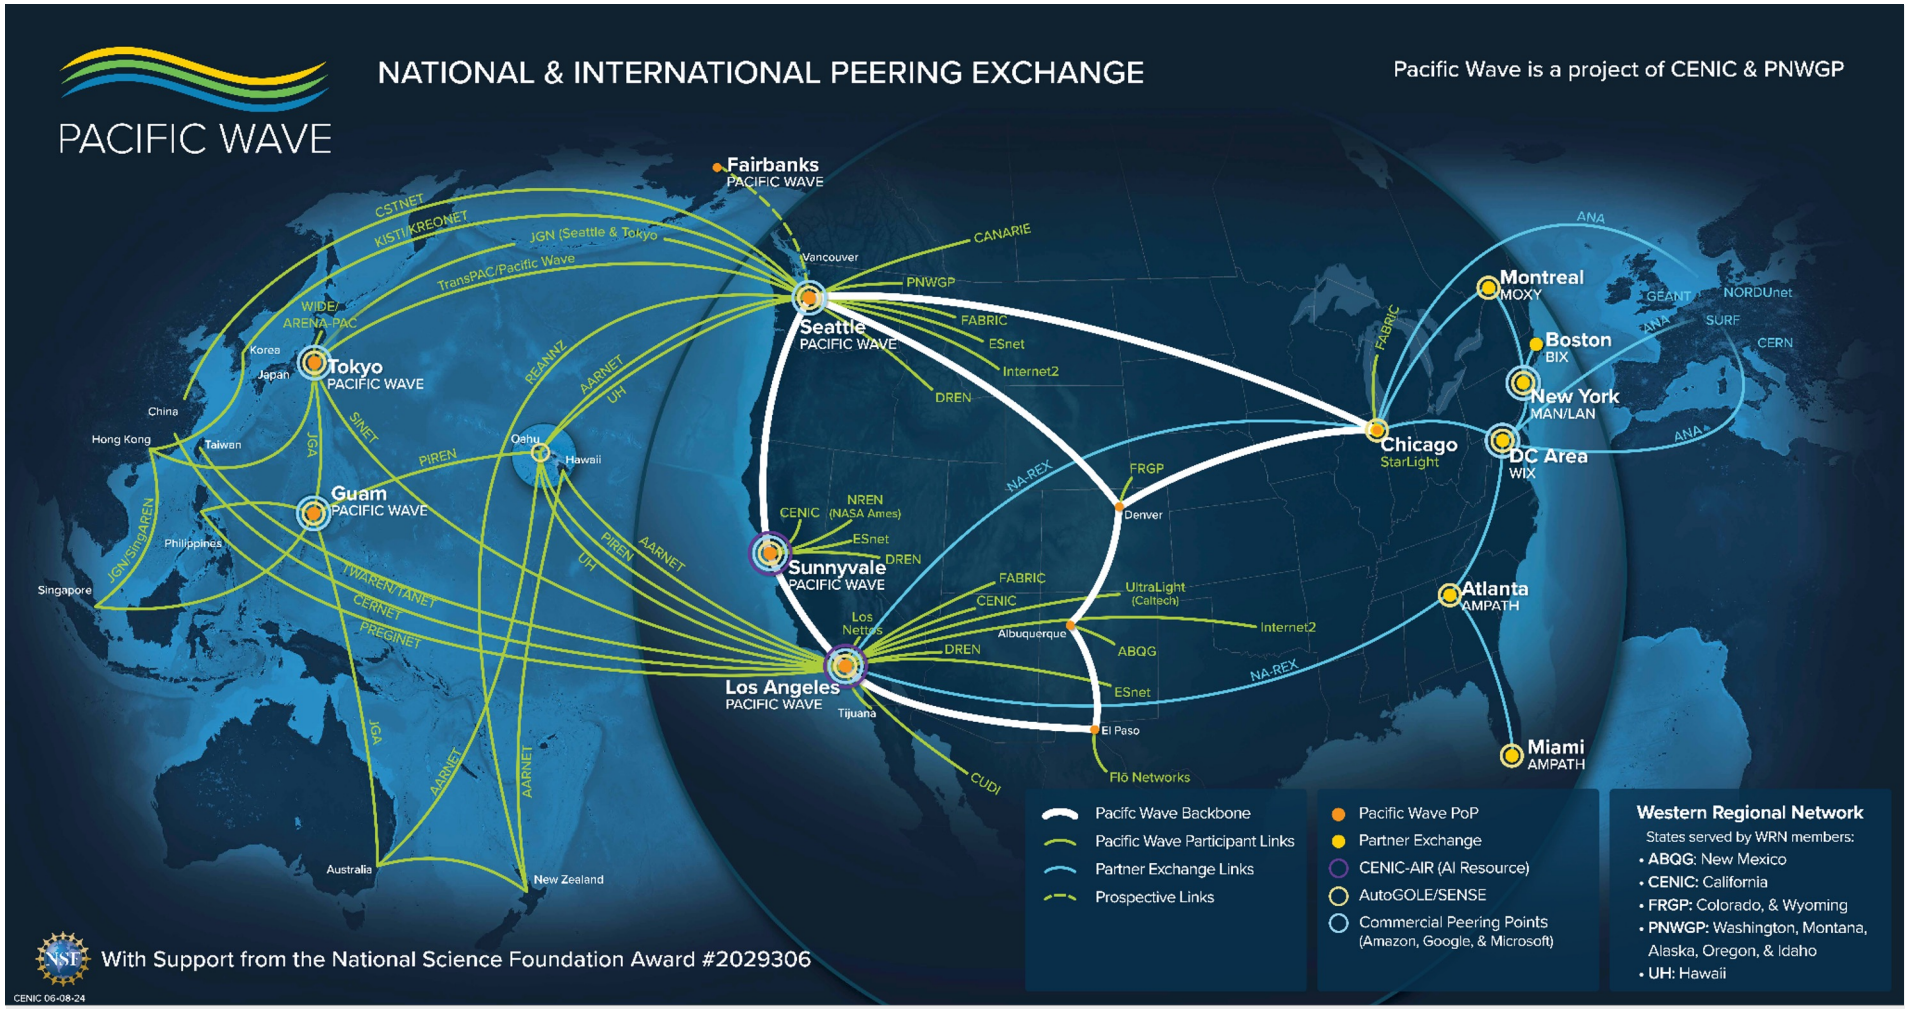
PACIFIC WAVE
NATIONAL & INTERNATIONAL PEERING EXCHANGE
Pacific Wave is a project of CENIC & PNWGP
Fairbanks
PACIFIC WAVE
CSTNET
KISTI/KREONET
JGN (Seattle & Tokyo
TransPAC/Pacific Wave
WIDE/
ARENA-PAC
Korea
Japan
Tokyo
PACIFIC WAVE
China
Hong Kong
Taiwan
Philippines
Singapore
Guam
PACIFIC WAVE
JGA
SINET
PIREN
JGN/SingAREN
TWAREN/TANET
CERNET
PREGINET
Oahu
Hawaii
REANNZ
AARNET
UH
Australia
New Zealand
Vancouver
Seattle
PACIFIC WAVE
CANARIE
PNWGP
FABRIC
ESnet
Internet2
DREN
NREN
CENIC
(NASA Ames)
ESnet
Sunnyvale
PACIFIC WAVE
DREN
FABRIC
CENIC
Los
Nettos
Albuquerque
DREN
Los Angeles
PACIFIC WAVE
Tijuana
CUDI
Flō Networks
El Paso
ESnet
ABQG
UltraLight
(Caltech)
Internet2
Denver
FRGP
NA-REX
Chicago
StarLight
FABRIC
Montreal
MOXY
ANA
GÉANT
NORDUnet
SURF
CERN
Boston
BIX
New York
MAN/LAN
DC Area
WIX
Atlanta
AMPATH
Miami
AMPATH
Pacifc Wave Backbone
Pacific Wave Participant Links
Partner Exchange Links
Prospective Links
Pacific Wave PoP
Partner Exchange
CENIC-AIR (AI Resource)
AutoGOLE/SENSE
Commercial Peering Points
(Amazon, Google, & Microsoft)
Western Regional Network
States served by WRN members:
• ABQG: New Mexico
• CENIC: California
• FRGP: Colorado, & Wyoming
• PNWGP: Washington, Montana, Alaska, Oregon, & Idaho
• UH: Hawaii
With Support from the National Science Foundation Award #2029306
CENIC 06-08-24

Pacific Wave 2025 Roadmap

- Maintaining Current Juniper MX10008 Platform
- AutoGOLE for Circuit Provisioning Automation
- Complete final 400 Gbps Backbone Upgrades (Seattle - Los Angeles)
- Deploy New Node in Alaska at Fairbanks
- Continue coordination with GOREX (Guam Open R&E Exchange)
- iGROK Monitoring Nodes & Dashboards
- Enabling FABRIC Connectivity
- Access to CENIC AI Resource (CENIC-AIR)

ESnet is the DOE'S data circulatory system…

- ESnet supports the DOE scientific research ecosystem.
- Interconnects all national labs and user facilities
- Provides reliable, high-performance connectivity to global research collaborations, the Cloud, and the larger Internet.

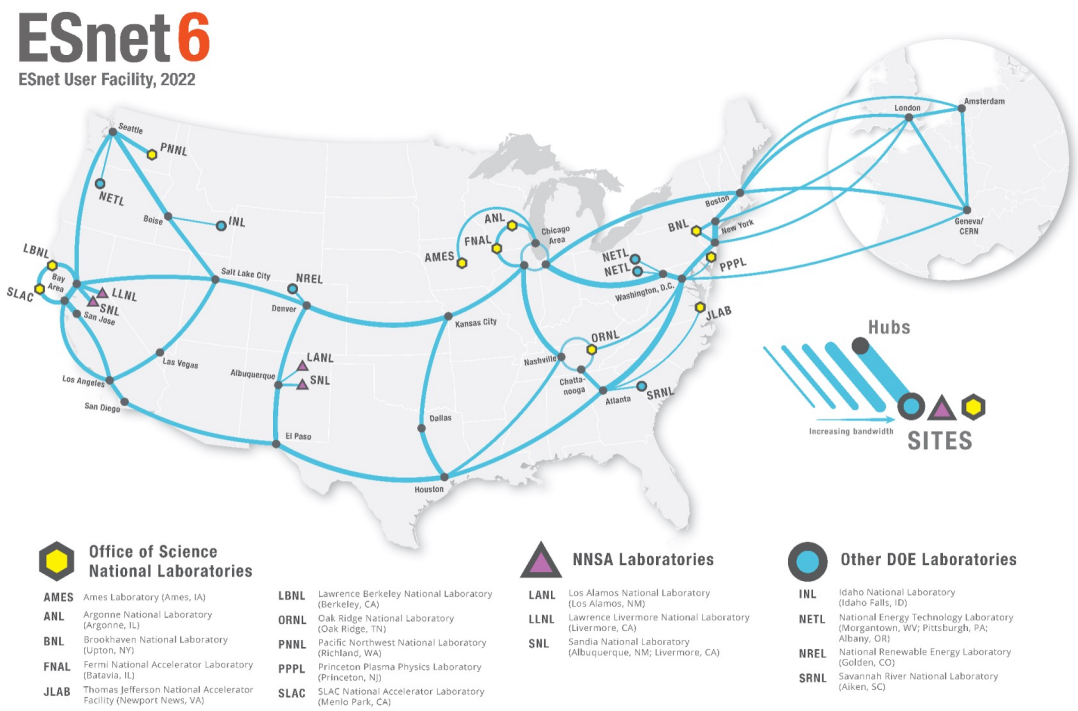

ESnet

…and the stage for a global science laboratory.

ESnet's Vision

Scientific progress will be completely unconstrained by the physical location of instruments, people, computational resources, or data.

ESnet's Mission

Networking that accelerates science.

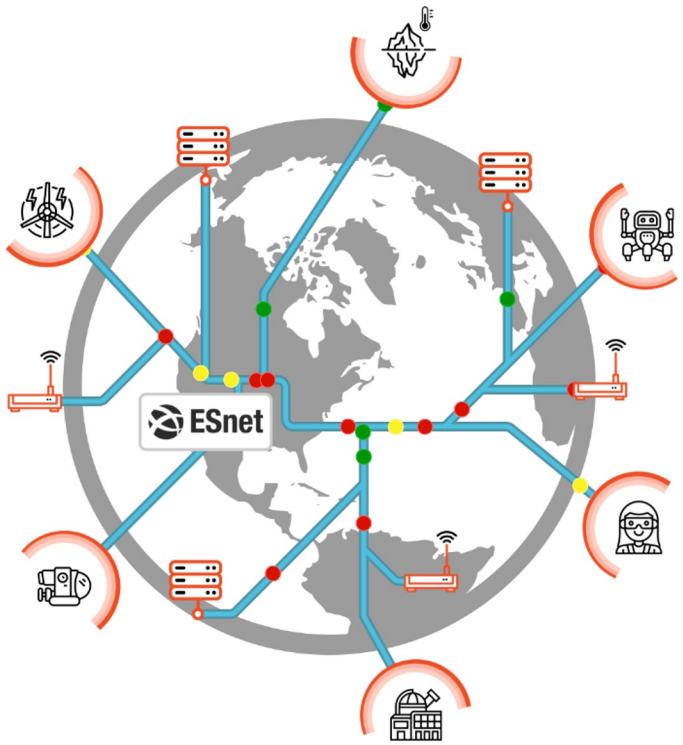

ESnet

ESnet - Peering Exchange Points

Public Peering Exchange Points

Filter

Exchange AŻ IPv4	ASN IPv6	Speed Port Location	RS Peer	BFD Support
AMS-IX 80.249.213.7	293 2001:7f8:1::a500:293:1	10G	✓	○
Equinix Ashburn 206.126.236.137	293 2001:504:0:2::293:1	100G	✓	○
Equinix Chicago 208.115.136.58	293 2001:504:0:4::293:1	100G	✓	○
Equinix Dallas 206.223.119.28	293 2001:504:0:5::293:1	100G	✓	○
Equinix San Jose 206.223.116.137	293 2001:504:0:1::293:1	100G	✓	○
MASS-IX 206.53.143.5	293 2001:504:47::125:0:1	100G	✓	○
Pacific Wave 207.231.240.164	293 2001:504:b:20::164	100G	○	○
Pacific Wave 207.231.244.2	293 2001:504:b:30::2	100G	○	○
Pacific Wave 207.231.242.13	293 2001:504:b:11::13	100G	○	○
Pacific Wave 207.231.240.13	293 2001:504:b:10::13	100G	○	○
SIX Seattle 206.81.81.102	293 2001:504:16::125	100G	✓	○
SIX Seattle (Jumbo)	293	100G	✓	○

https://www.peeringdb.com/net/940

FABRIC + FAB Testbed

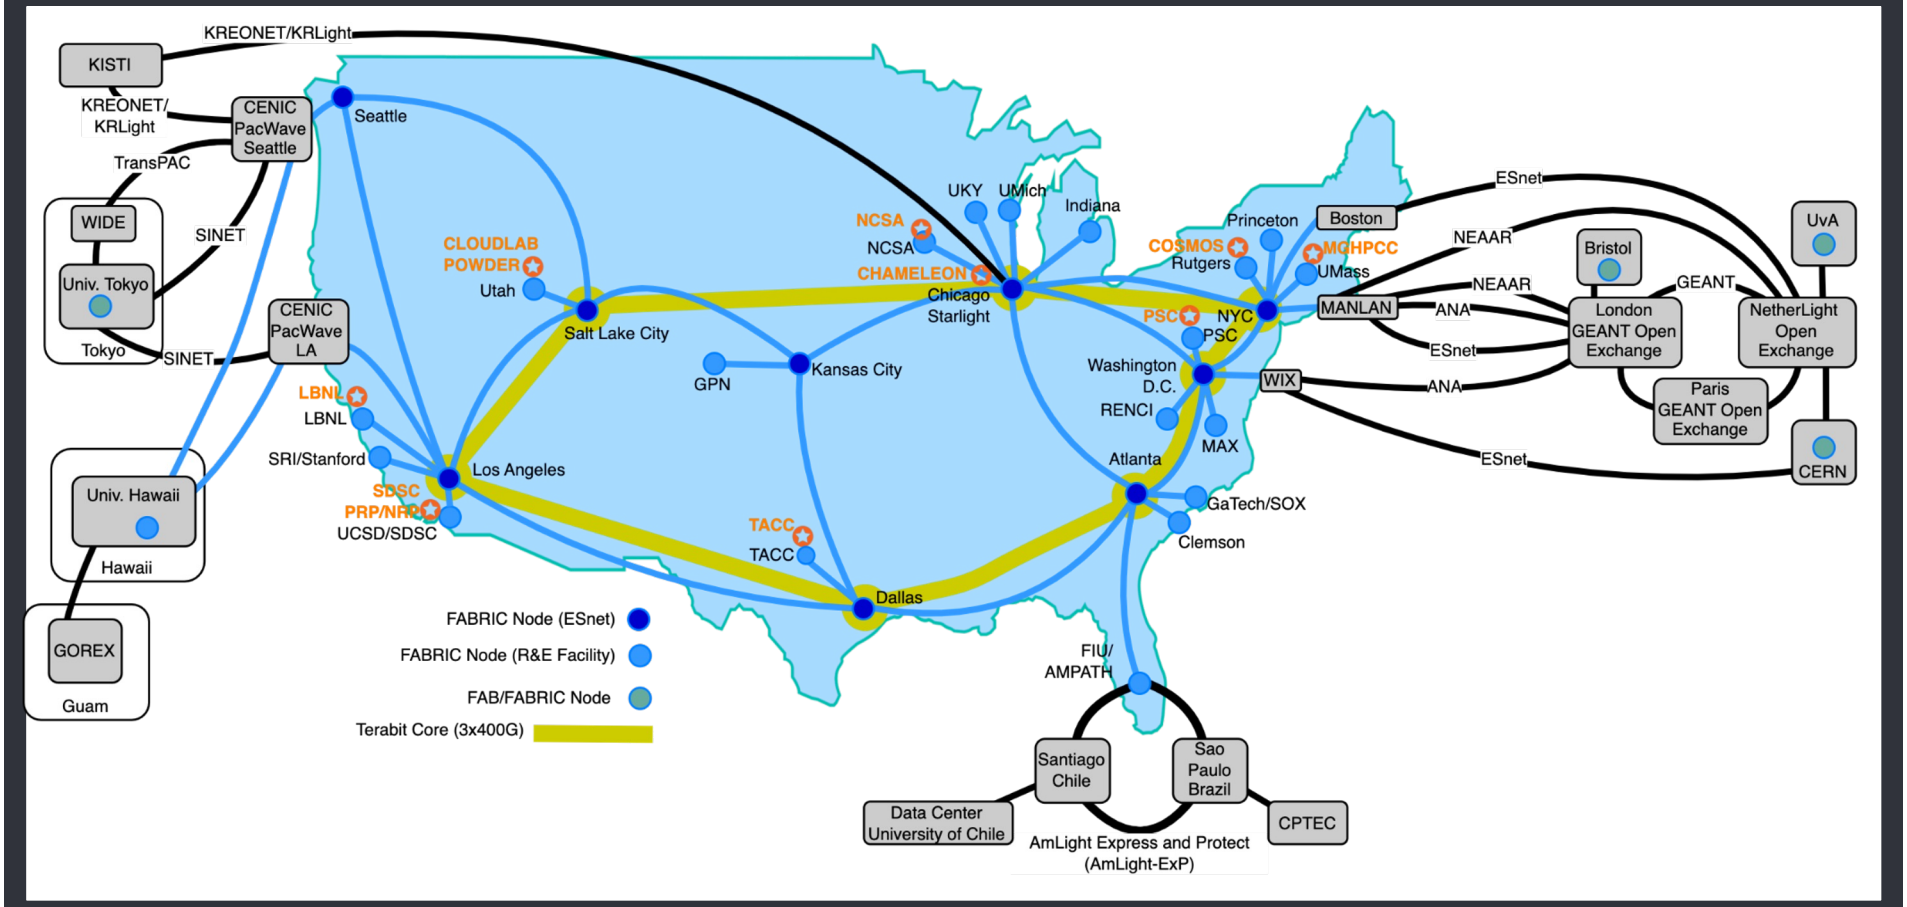

Americas Lightpaths (AmLight), NSF Award #OAC-2029283

- 600G of upstream capacity between the U.S., Latin America, Caribbean and 100G to Africa
- OXPs: Florida(3), Georgia, Brazil(2), Chile, Argentina, Puerto Rico, Panama, and South Africa
- Production SDN Infrastructure since 2014
- Deeply programmable across the network stack
  - Programmable P4 Data Plane
  - Open Source SDN Controller (kytos-ng)
  - Fine-grained telemetry
- Highly instrumented
  - PerfSonar, sFlow, Juniper Telemetry Interface (JTI), In-band Network Telemetry (INT)

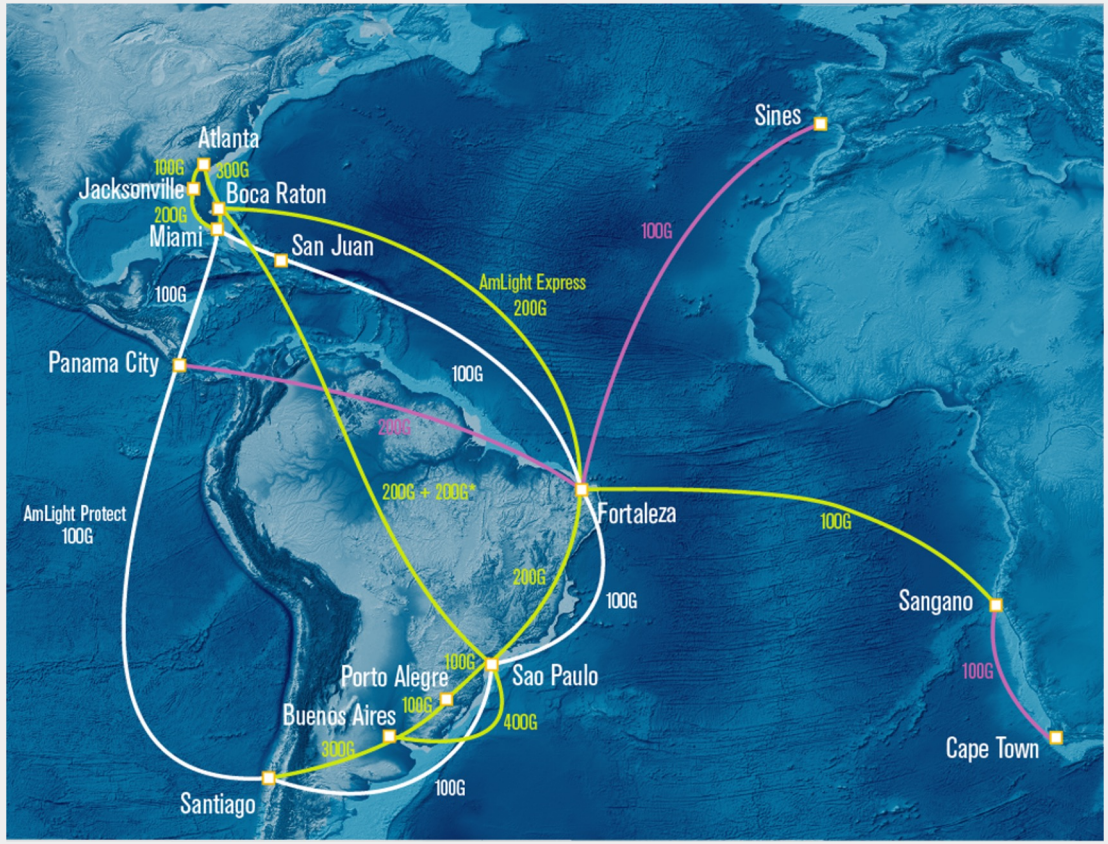

AmLight Roadmap for 2025

- Deploy a Ciena Waveserver 6E at Sao Paulo, Fortaleza, and Boca Raton to activate a total of 1.1Tbps (currently we have 400G using Waveserver Ai)
- Upgrade 1x100G from Miami to Jacksonville to 1x400G
- Upgrade 1x100G from Jacksonville to Atlanta to 1x400G
- Activate the NA-REX connectivity: 1x400G to StarLight and 1x400G to PacificWave

MOXY

- R&E Open Exchange in Canada
- Established in mid 2017
- Located in the heart of Montreal
- Supporting ANA system in improving physical route diversity and research activities across North Atlantic
- Operate independently from CANARIE network

MOXY Platform and Service

Current:

- Juniper QFX10002
- Support 10G and 100G Connectivity
- ANA 100G link to Amsterdam
- 100G links to Starlight and MANLAN
- L2 only services
- NSI - SuPA(SURF ultimate Provider Agent) deployed

Looking forward - Summer 2025

- Deploy a new 400G/800G switching platform
- Upgrade ANA from 100G to 400G to Amsterdam
- Establish 400G NA-REX links to Starlight and MANLAN
- Deploy L2 only services, evolve services if needed
- Maintain SuPA - NSI support

International Networks at Indiana University (IN@IU) TransPAC5 and NEA3R

- Currently support 1 400G and 3 100G links
  - 100G NY-London
  - **400G NY-Amsterdam**
  - 100G TransPAC-Pacific Wave Seattle - Tokyo
  - 100G GSCC Guam - Singapore
- Coordinating globally to ensure US researchers have worldwide high speed access to collaborators and projects.
- NSF Awards: NEA3R #2028495 and TransPAC5 #2028501

IN@IU links carry data from researchers in 84% of the world!

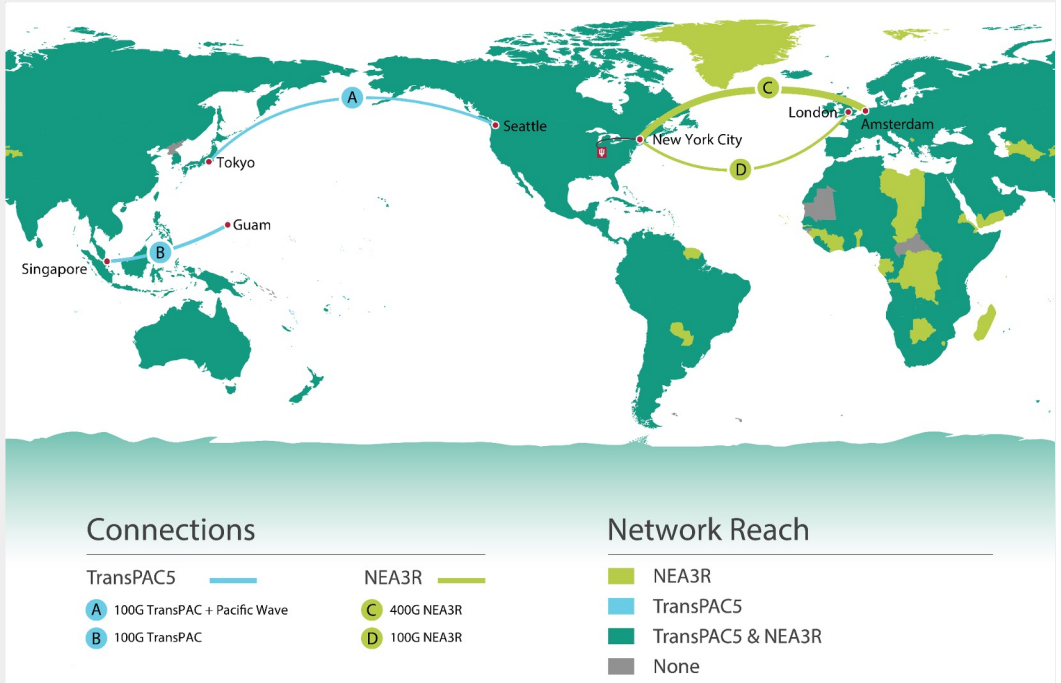

NEA3R, TransPAC5, and FABRIC Across Borders (FAB)!

- FABRIC Across Borders (FAB) is an extension of the FABRIC testbed connecting the core North America infrastructure to four nodes in Asia, Europe, and South America. By creating the networks needed to move vast amounts of data across oceans and time zones seamlessly and securely, the project enables international collaboration to speed scientific discovery
- **NEA3R** and **TransPAC** provide the international connectivity for FABRIC as part of FAB
- FAB is built around science uses cases in Smart Cities, Weather, Physics, Space, and Computer Science

NA-REX at SC24 - Networked Research Experiments

- SC is a test bed and showcase for cutting-edge developments in high-performance networking, computing, storage, and analysis.
- Networked Research Exhibition (NRE) demonstrations leverage the advanced capabilities of SCinet, the conference's dedicated high-capacity network.
- Network researchers from government, education, research, and industry are invited to submit proposals for demonstrations and experiments at the SC Conference that display innovation in emerging network hardware, protocols, advanced network-intensive scientific applications, and more!

- SCinet
  - 200 Volunteers - 9 countries and 34 states
  - $47m of contributed equipment
  - 1 year to design, 1 month to build, 1 week to operate, 1 day to teardown

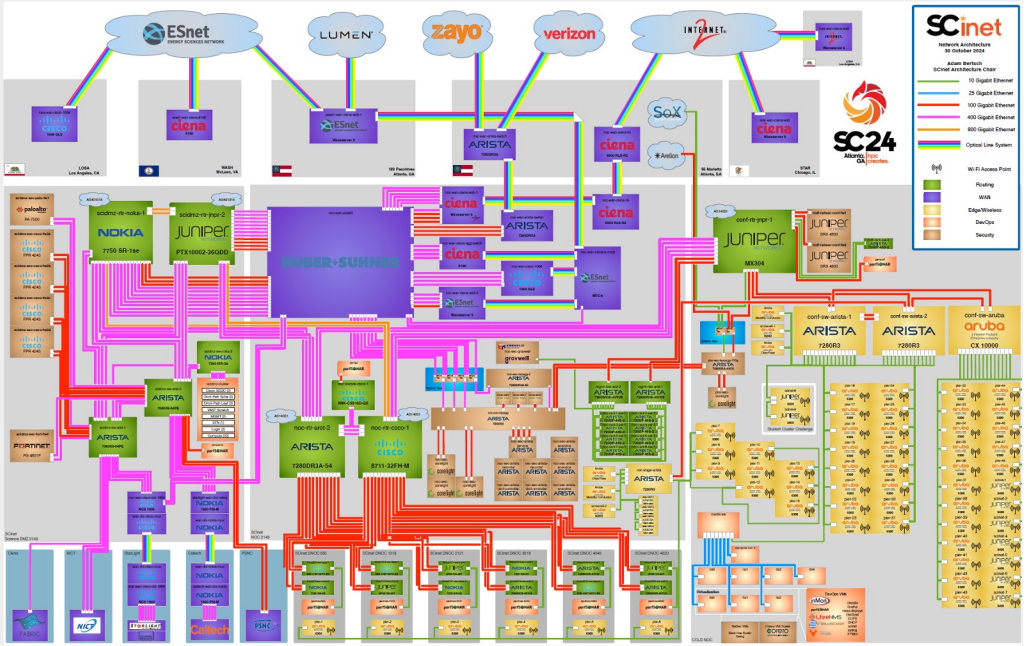

NA-REX at SC24 - Networked Research Experiments (2)

- 36 Networked Research Experiments
- https://sc24.supercomputing.org/scinet/network-research-exhibition/accepted-nre-demos/
- Over 10Tbps of traffic in and out of Atlanta.

- NRE participants use cases
  - High Speed File Transfers (MMCTFP)
  - Connecting and using international testbeds.
    - **FABRIC** Node/Rack on the show floor!
  - Science flow tagging
  - Sensor networks
  - Programmable packet processing
  - Novel routing protocols
  - Testbed integration demonstrations
  - Network resiliency research
  - Data Intensive Science applications
  - Automated network service orchestration

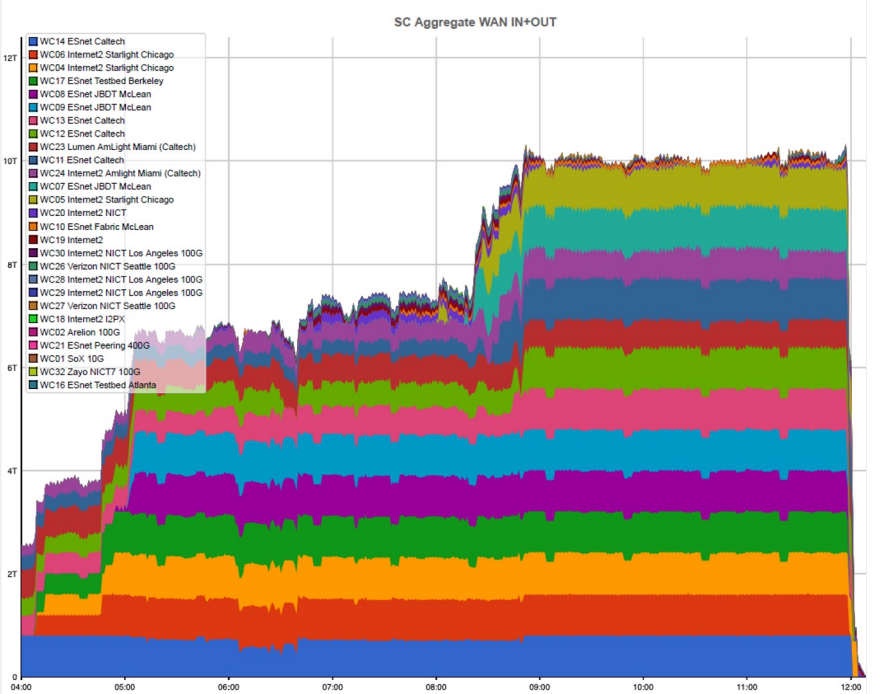

NA-REX at SC24 - International NRE's using NA-REX

- **At least 12 of the 36 NRE's used NA-REX for some or all bandwidth!**
- NRE007 - 400 Gbps E2E WAN Services: Architecture, Technology and Control Systems
- NRE012 - FABRIC And Data Intensive Science Prototype Services
- NRE014 - MMCFTP's Data Transfer Experiment Using Ten 100 Gbps Lines Between JAPAN and USA
- NRE017 - Cross-Site Network Telemetry Based on Programmable Network Technology
- NRE020 - Toward Terabit-Scale Anonymous Communication Leveraging Programmable Switches
- NRE021 - Floating-Cyber Physical System for Local Production and Consumption of Data
- NRE022 - AutoGOLE/SENSE: Edge Site Resource Integration with Network Services
- NRE024 - FABRIC
- NRE029 - Multi Domain Experiments Using ESnet SENSE on the National Research Platform / PacWave / FABRIC
- NRE035 - Hybrid CPU, GPU, QPU Infrastructure for Hybrid Quantum – Classical Computing Use Cases Development with Secure QKD/PQC
- NRE036 - Streaming Event Horizon Telescope Data to MIT-Haystack
- NRE037 - Disk-to-Disk Data Transfer Performance Investigation on Long-haul Networks Towards Practical Application of Research-EnhanceD ONION

Quantum Communications – SC24 NRE35

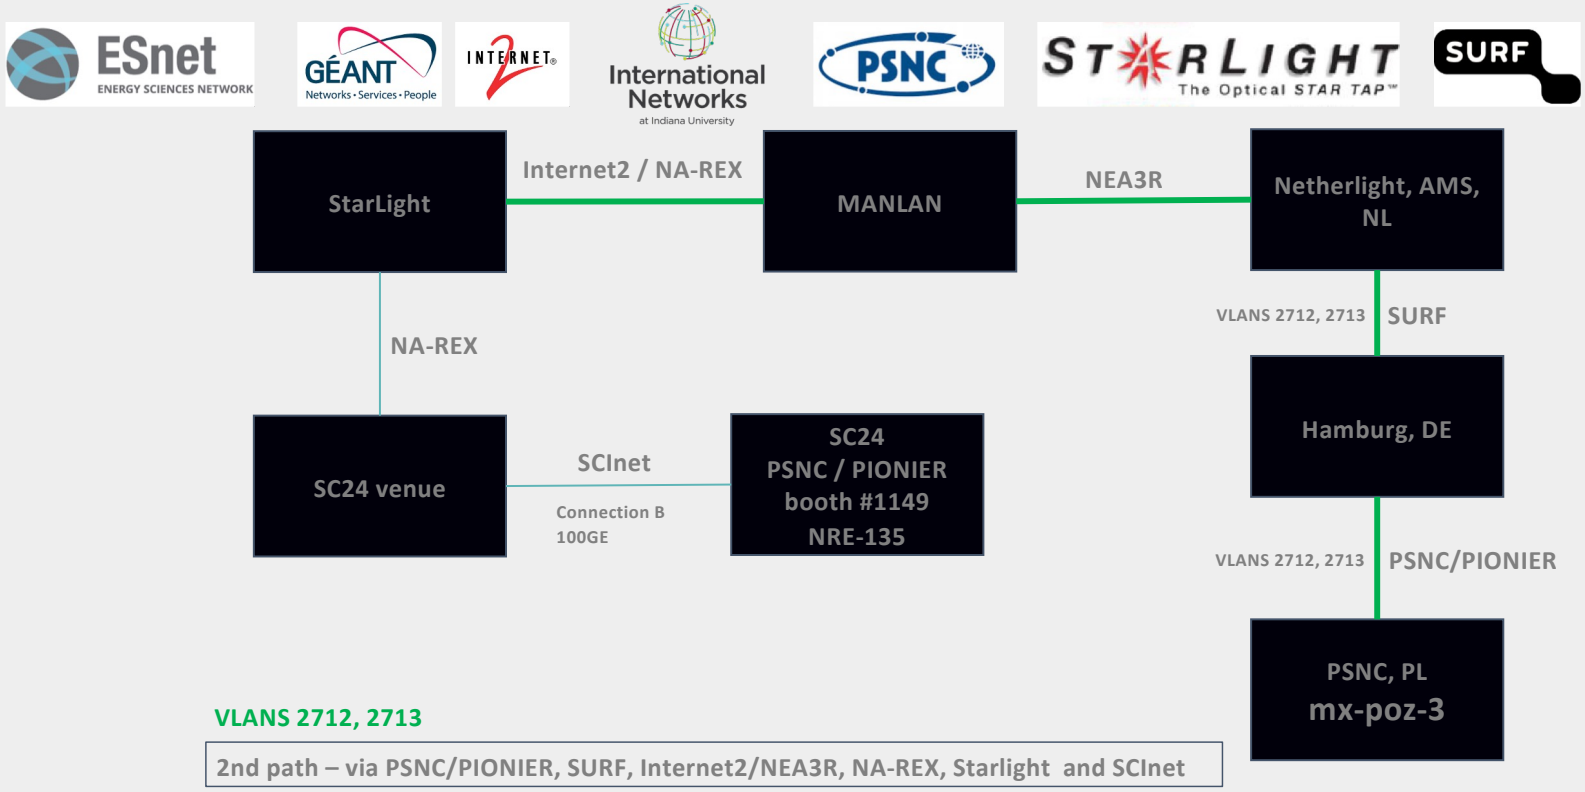

https://www.psnc.pl/distributed-hybrid-quantum-classical-computing-in-a-post-quantum-cryptography-world/

What's next?

- Augment NA-REX Shared Exchange Link Capacity - 400Gbps and beyond
  - Prepare for new subsea capacity increases (2 x 400Gbps+)
- Continuous Improvement for Exchange Operational Coordination
- Increase Programmatic Support for Experimental Use Cases
  - e.g., rapid link topology changes for SCinet, OFCnet, network experimentation
- Review Next Generation Network Orchestration & Automation

INTERNET2

TECH ex 24

BOSTON DEC 9-13