Global Network Advancement Group
Next Generation Network-Integrated System for Data Intensive Sciences

LHC | LBNF/DUNE | SKA

LHC Run 3 and HL-LHC

Rubin Observatory

SKA

BioInformatics

Earth Observation

Gateways to a New Era

Harvey Newman (Caltech)

6th Global Research Platform Workshop, Chicago

on behalf of the GNA-G and its DIS and AutoGOLE Working Groups

Global Network Advancement Group (GNA-G)
Since September 2019

- An open volunteer group devoted to developing the blueprint to make using Global R&E networks both simpler and more effective
- Its primary mission is to support global research and education using the technology, infrastructures and investments of its participants.

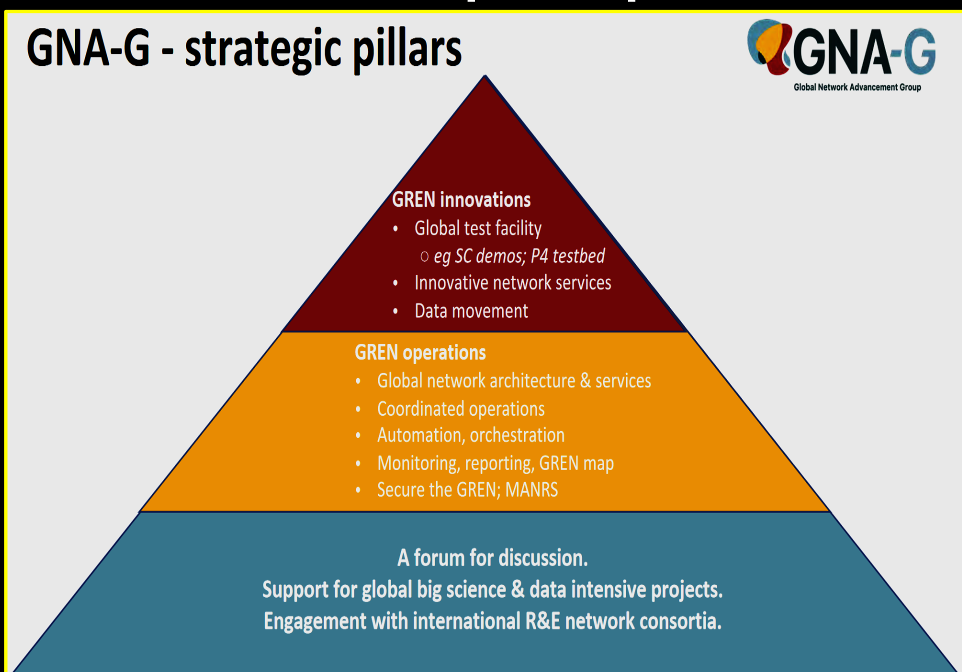

- The GNA-G is a data intensive research & science engager that facilitates and accelerates global-scale projects by

  (1) enabling high-performance data transfer, and

  (2) acting as a partner in the development of next generation intelligent network systems that support the workflow of data intensive programs

✸ In a manner compatible with the traffic and requirements of the worldwide Academic and Research communities

See https://www.dropbox.com/s/qsh2vn00f6n247a/GNA-G%20Meeting%20slides%20-%20TechEX19%20v0.8.pptx?dl=0

GNA-G
Global Network Advancement Group

https://www.gna-g.net/

Mission: Support global research and education using the technology, infrastructures and investments of its participants

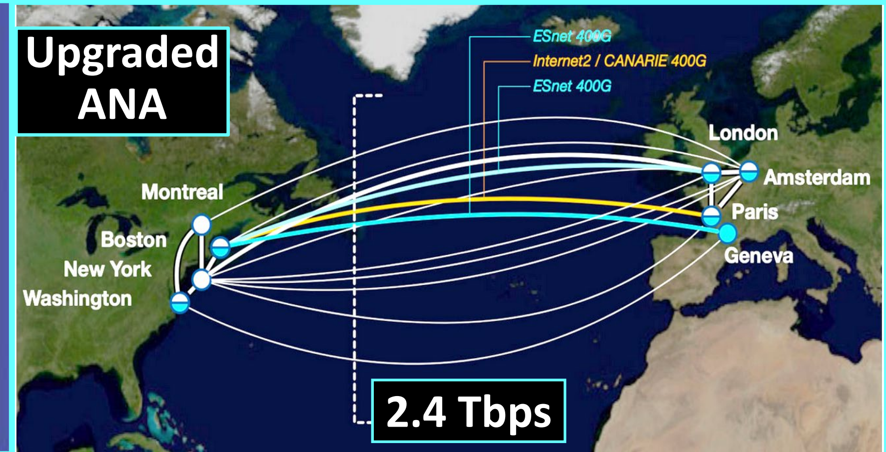

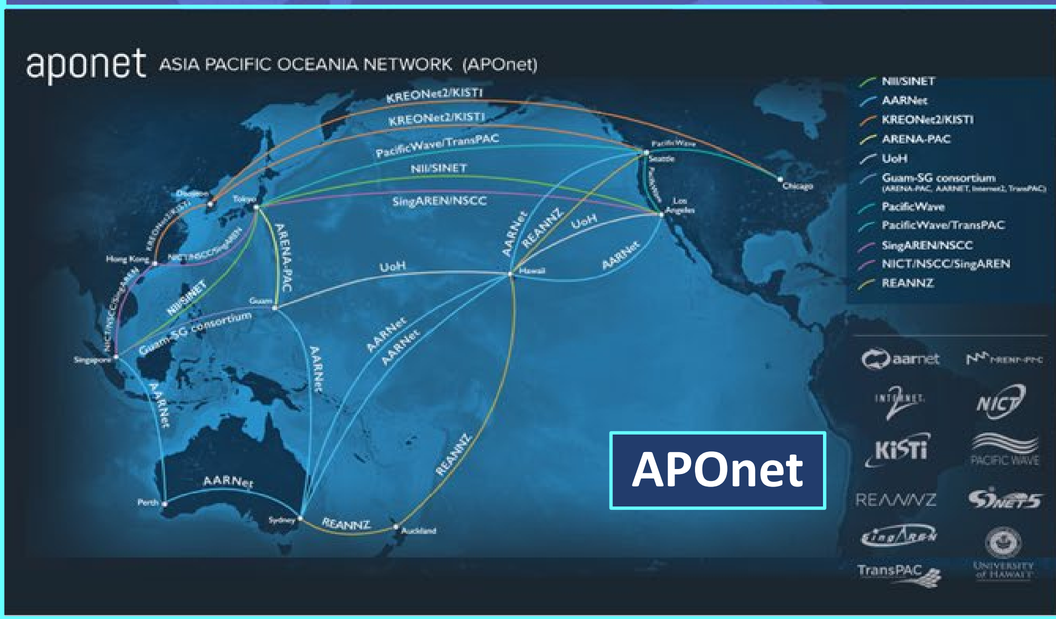

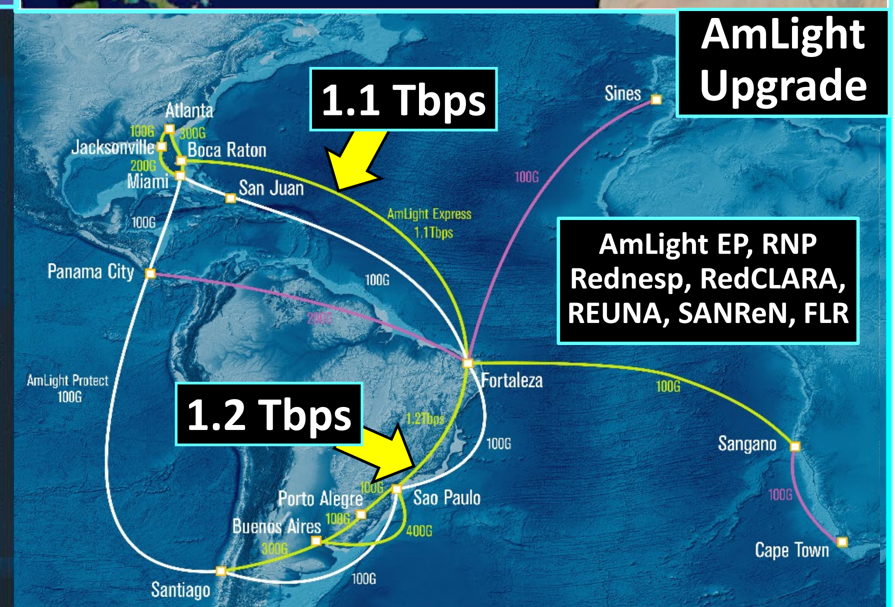

The GNA-G exists to bring together researchers, National Research and Education Networks (NRENs), Global eXchange Point (GXP) operators, regionals and other R&E providers, in developing a common global infrastructure to support the needs

Rednesp and RNP: Expanding Capacity Among Latin America, US, Europe and Africa

Total 1 Tbps capacity between Brazil and the USA:

- Rednesp (formerly ANSP) has 4 links to the USA, connecting Sao Paulo to Florida.
- **Atlantic links:** São Paulo- Miami: 400G via Boca Raton and 100G via San Juan
- **Pacific 100G link:** São Paulo to Santiago (Chile) and Panama; from there to San Juan & Miami

✸ **AmLight Link (with Rednesp and RNP):** São Paulo to Florida through Fortaleza, on the Monet cable (Angola cables). 400G Now; Upgrade to 1.1 Tbps by this fall

Transatlantic Links: Bella link between São Paulo and Sines in Portugal, and a link connecting São Paulo, Angola and South Africa

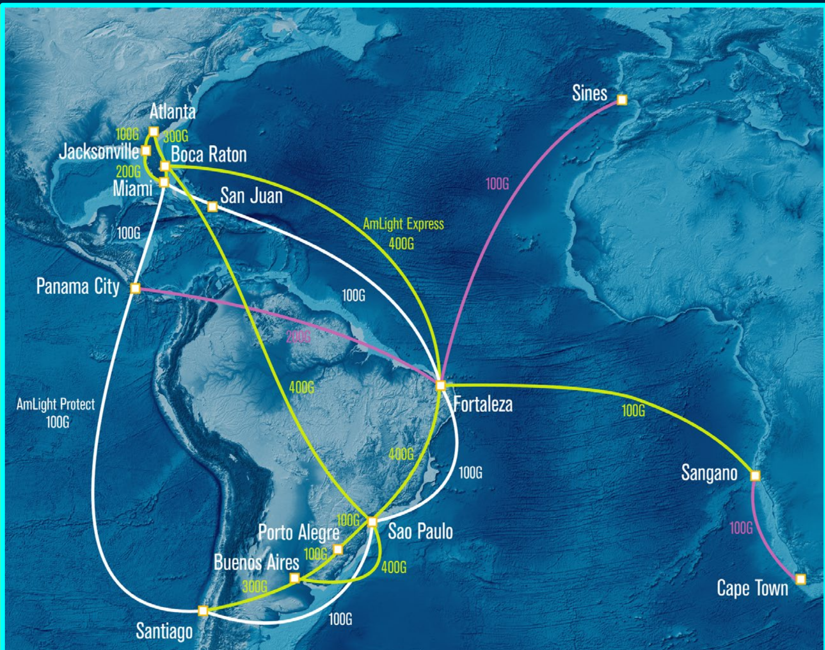

- "Backbone SP" connecting most universities in São Paulo State with 100G links was completed in May 2023
  - ❒ UNICAMP (University of Campinas),
  - ❒ UNIFESP (Federal University at São Paulo)
  - ❒ UNESP (State University of São Paulo)
  - ❒ USP (University of São Paulo)
  - ❒ UFSCAR (Federal University of São Carlos)
  - ❒ Mackenzie University
  - ❒ ITA (Aeronautics Institute of Technology)
  - ❒ UFABC (Federal University of ABC) to SP4

✸ The SIRIUS synchrotron light source has been connected to the Backbone via UNICAMP since 2024

✸ 400G dark fiber links are being added by Rednesp from its main Equinix PoP to USP, UNESP and UNICAMP

✸ USP dark fiber infrastructure nearing completion: 3 campuses already connected at 200G

✸ Four main USP campuses are expected to be connected at 400G-800G by the end of 2025

9/2025 Snapshot: ESnet Traffic Growth +114% in the Last Year with LHCONE and OSCARs ~60% of the Total

	Bytes/Mo	Pct of Total	One Month Change	One Year Change
OSCARS	45.6 PB	19.5%	+1.9%	+130%
LHCONE	93.8 PB	40.1%	+4.9%	+116%
Normal Traffic	94.3 PB	40.3%	-15.7%	+105%
Total	233.8 PB		- 3.5%	+114%

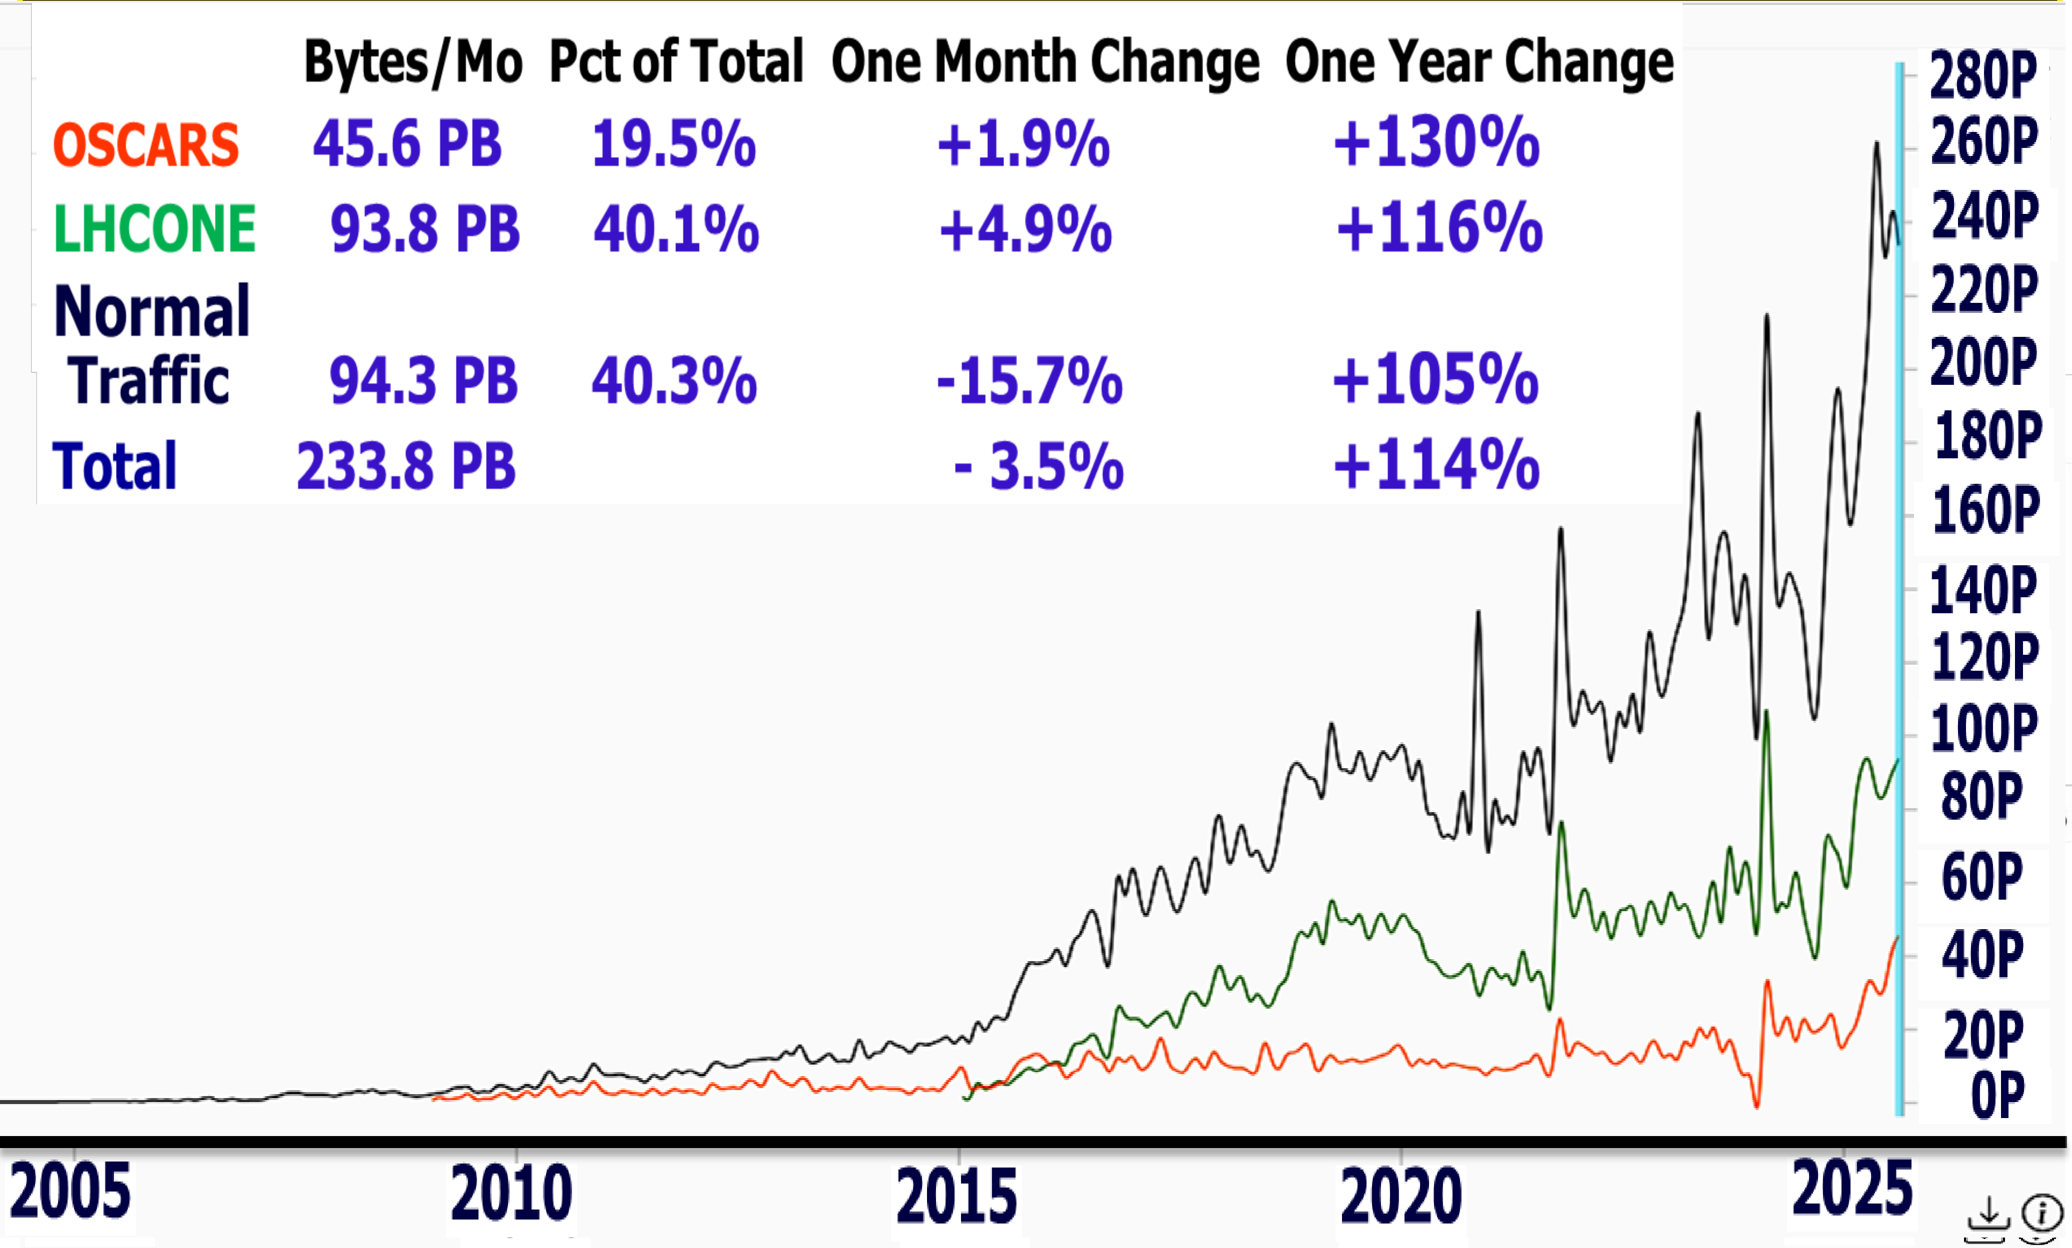

WLCG Transfers Dashboard: Throughput Sept. 2015 – August 2025

GNA-G
Global Network Advancement Group

LHC Data Flows Have *Increased* in Scale and Complexity since the start of LHC Run2 in 2015

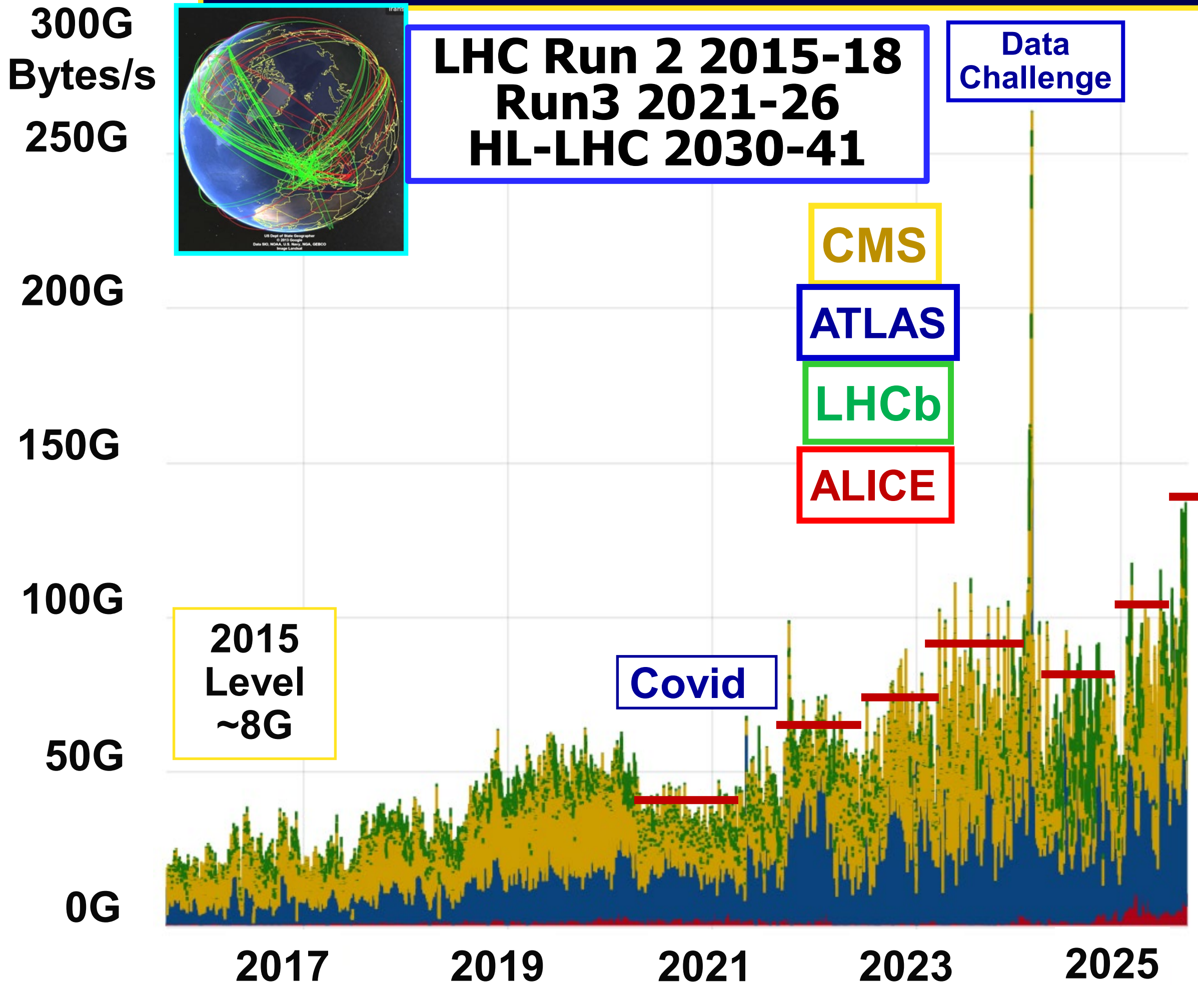

75-130 GBytes/s Weekly Avg
Larger Hourly/Daily Bursts

Complex Workflow

- To ~2 M jobs (threads) simultaneously
- Multi-TByte to Petabyte Transfers
- To ~30 M File Transfers/Day
- Millions of remote connections

- *Another Sea Change in 2024-5*
- 1.8X in Transfer Rates in 12 months
- DC24 (25% HL LHC): to 300 GB/s

~17X Growth in Throughput 2015-2025: +33%/Yr; + Much Faster Growth Bursts

https://monit-grafana.cern.ch/d/AfdonlvGk/wlcg-transfers?orgId=20&from=now-10y&to=now

Structure

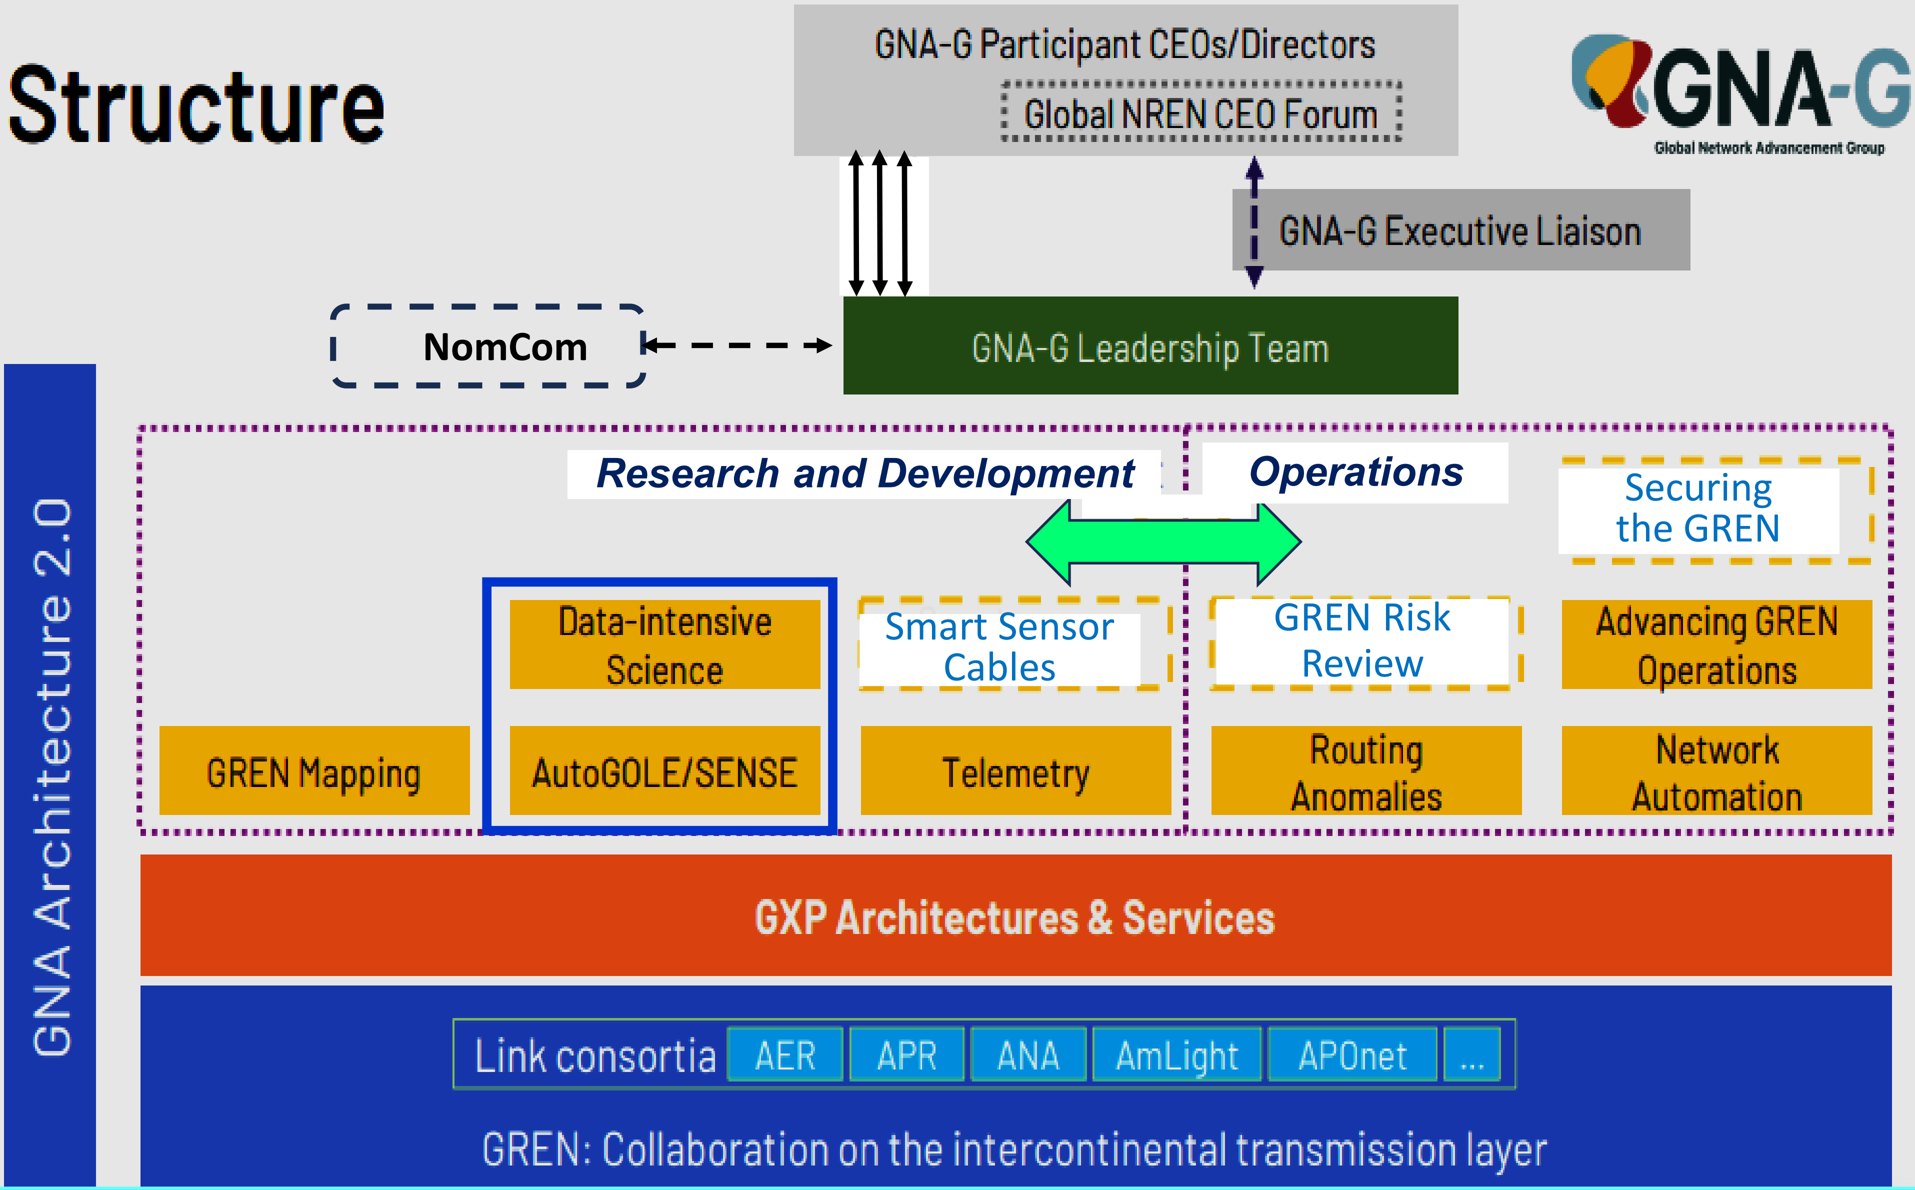

Theme: from Transformative Concepts to Advanced Services to Production Networking
New Directions: Scaling to Tbps; Digital Twins; Generative AI for Global Networks

The Global Network Advancement Group DIS WG: Next Generation System for Data Intensive Sciences

- A Vast *Partnership* of Science and Computer Science Teams, R&E Networks and R&D Projects; Convened by the GNA-G DIS WG; with GRP, NRP, AmRP
- Mission: Demonstrate the road ahead; Develop and deploy next generation systems
  - Meeting the challenges faced by leading-edge data intensive programs in HEP, astrophysics, genomics and other fields of data intensive science;
    - ✸ *Compatible with other use by the at-large A&R communities*
  - **Accelerating the path to the next round of discoveries**
- Demonstrating a wide range of latest advances in:
  - State of the art long distance data transfer methods and tools, local and metro optical networks and server designs
  - Software defined and Terabit/sec networks
  - Intelligent global operations and monitoring systems
  - Workflow optimization methodologies with real time analytics
  - Emerging technologies and concepts in programmable networks and global-scale distributed systems
- **Hallmarks: Progressive multidomain integration; compatibility internal + external;** ***A comprehensive systems-level approach***

GNA-G
Global Network Advancement Group
SC25
ST. LOUIS NOV 16-21

The Global Network Advancement Group

Progress and Impact Since 2019-2025

- Creating a forum to share and coordinate knowledge and resources across multiple world regions in support of both leading-edge science programs and the broader academic and research community that rely on networks for their daily work
- Developing key services, methods and system-level concepts to meet the challenges faced by distributed science collaborations as they share, process and collaboratively analyze exponentially growing datasets across the globe
- Creating a persistent R&D network infrastructure spanning world regions, and a development trajectory providing an effective pathway to bring advanced network and associated data services progressively into production use
- Deploying transformative approaches to intelligent network operations and the acceleration of science workflows: from multidomain virtual circuits to new forms of source based routing, to machine-learning based methods to optimize path selection, flow balancing, and the levels of service provided to multiple traffic classes with varying requirements and constraints
- Bringing the above themes together to develop a next generation paradigm of network operations and intelligent agent-based services, to accelerate our progress towards the next round of science discoveries, and advance our knowledge of the universe
- *While ensuring the broad A&R community's continued excellent service*

Global Petascale to Exascale Workflows for Data Intensive Sciences: Development Trajectory

- Advances Embedded and Interoperate within a 'composable' architecture of subsystems, components and interfaces, organized into several areas; coupled to rising Automation
  - Visibility: Monitoring and information tracking and management including IETF ALTO/OpenALTO, BGP-LS, sFlow/NetFlow, InMon, Perfsonar, Traceroute, Qualcomm Gradient Graph congestion information, Kubernetes statistics, Prometheus, Inband telemetry (P4, AmLight Kytos)
  - Controllability: SENSE/AutoGOLE/SUPA, P4, segment routing with SRv6 and/or PolKA, BGP/PCEP
  - Network OSes and Tools: SONIC, GEANT RARE/freeRtr; Calico VPP, Bstruct-Mininet environment, ...
  - Orchestration+Automation: SENSE, Kubernetes (+k8s namespace), dedicated code and APIs for interoperation and progressive integration
  - Intelligence: Stateful decisions using composable metrics (policy, priority, network- and site-state, SLA constraints, responses to 'events' at sites and in the networks, ...), using NetPredict, Hecate, GradientGraph, Elastiflow/Elastic Stack, Yale Bilevel optimization, Coral

SC15-25: SDN Next Generation Terabit/sec Ecosystem for Exascale Science

supercomputing.caltech.edu

SDN-driven flow steering, load balancing, site orchestration Over Terabit/sec Global Networks

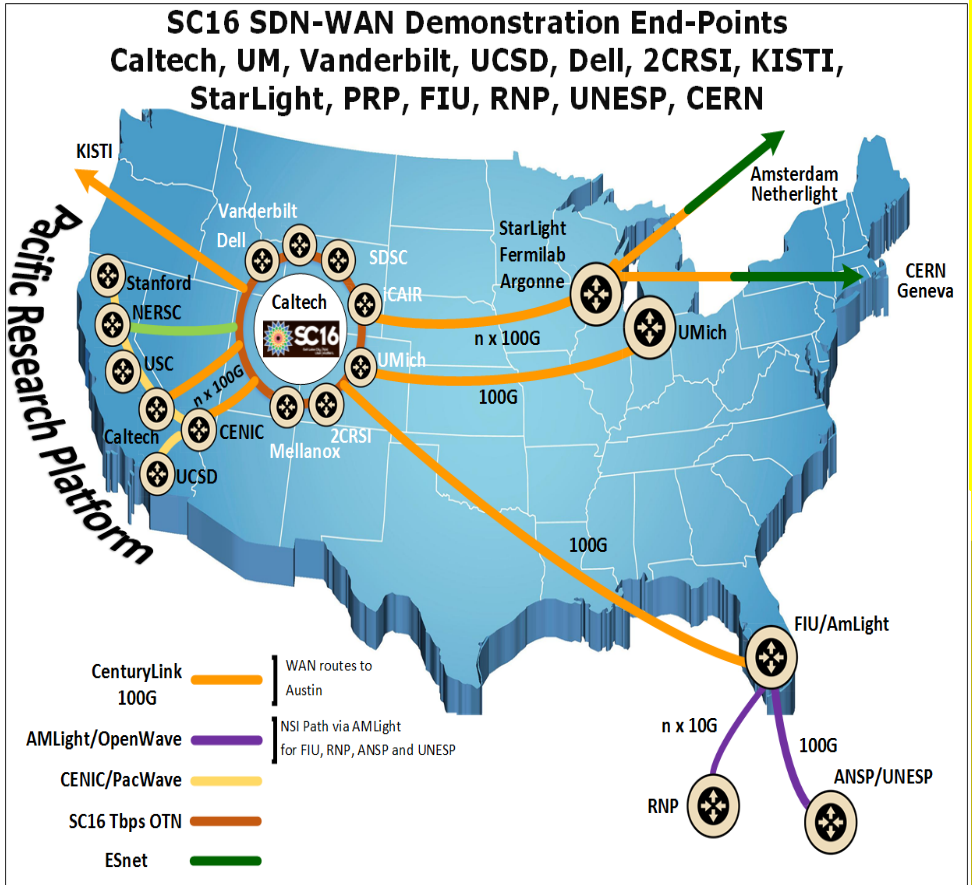

NRE-13 Top Sources: To 1.5+ Tbps on 4 X 400G Circuits with Dynamic Transfer Limit

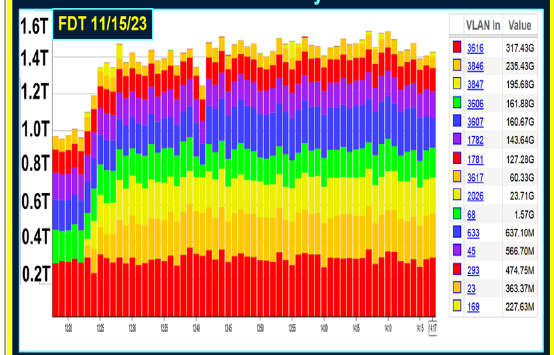

With Just 2 Gen5 + 2 (of 6) Gen3 Servers at SC23 and 3 Gen5 Servers at Caltech

SC16+: Consistent Operations with Agile Feedback Major Science Flow Classes Up to High Water Marks

Global Topology

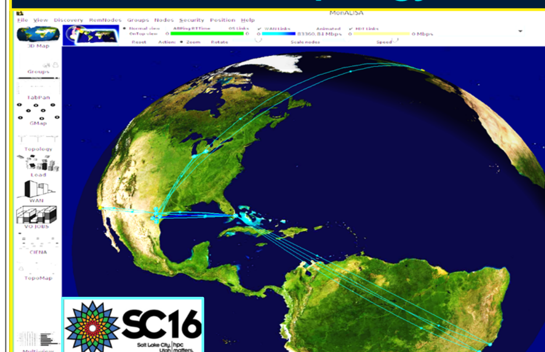

29 100G NICs; Two 4 X 100G and Two 3 X 100G DTNs; 1.5 Tbps Capability in one Rack; 9 32 X100G Switches

Preview PetaByte Transfers to/ from Sites With 100G - 1000G DTNs

LHC at SC15: Asynchronous Stageout (ASO) with Caltech's SDN Controller

Tbps Rings for SC18-24: Caltech, Ciena, SCInet, StarLight + Many HEP, Network, Vendor Partners

SC25 Network Research Exhibition NRE-125 and Partners
NREs Hosted at or Partnering with Caltech Booth 2824

NRE-101	Tom Lehman (ESnet) et al.	DOE Integrated Research Infrastructure (IRI) Compute Job Workflow Portability
NRE-102	Alex Moura (KAUST) et al.	Middle-East 800Gbps Transcontinental eScience Bridge
NRE-106	Tom Lehman (ESnet) et al.	Multi-Resource Cyberinfrastructure Services for Science Domain Workflows via SENSE
NRE-117	Tom Lehman (FABRIC) et al.	FABRIC
NRE-118	Vasilka Chergarova (AmLight) et al.	Cryo-EM Without Borders: Empowering Distributed Research via AmLight and Collaborative (Human) Networks
NRE-119	Mohammad Sada (UCSD/SDSC) et al.	National Research Platform and SCinet: Enabling Live, Multi-Institutional Scientific AI/ML and HPC Workflows
NRE-120	Mohammad Sada (UCSD/SDSC) et al.	Live High-Precision Per-Packet Kubernetes Data Telemetry from NRP to StarLight via ESnet SENSE
NRE-121	Mohammad Sada (UCSD/SDSC) et al.	Real-Time In-Network Machine Learning and P4 Testbed Deployment on FPGA SmartNICs, DPUs, and Switches
NRE-122	Mohammad Sada (UCSD/SDSC) et al.	Agentic AI with Qualcomm Cloud AI 100 Ultra Cards for HPC Cluster Management and Resource Provisioning
NRE-125	Harvey Newman (Caltech) et al.	The Global Network Advancement Group: A Next Generation System for Data Intensive Sciences
NRE-126	Marcos Schwarz (RNP) et al.	GP4L: Global Platform for Labs
NRE-127	Carlyn Ann Lee (Oligo/Caltech) et al.	Onboard AI and Federated Learning for Cosmic Ray Event Detection
NRE-128	Marcos Schwarz (RNP), Carlos Ruggiero (USP/Rednesp) et al.	High Performance Networking with the Sao Paulo Backbone SP Linking 8 Universities and the Bella Link
NRE-129	Carlos Ruggiero (USP/Rednesp) et al.	University of Sao Paulo (USP) High Performance Dark Fiber Infrastructure
NRE-130	Everson Borges (IFES), Magnos Martinello (UFES) et al.	PolKA Routing Approach to Support Traffic Engineering for Data-intensive Sciences: VRO Use Case
NRE-132	Julio Ibarra (FIU) et al.	AmLight 2.0: Flexible control, deep visibility, and programmability @ Tbps!

SC25 Network Research Exhibition
STARLIGHT NREs Partnering with Caltech Booth 2824

NRE-107	Joe Mambretti (STARLIGHT) et al.	**Software Defined Exchange (SDX)** Multi-Services for Petascale Science
NRE-108	Joe Mambretti (STARLIGHT) et al.	**Prototype 1.2 Tbps WAN Infrastructure** Architecture, Technology and Control Systems
NRE-109	Joe Mambretti (STARLIGHT) et al.	**Prototype 1.2 Tbps WAN Services** Architecture, Technology and Control Systems
NRE-110	Joe Mambretti (STARLIGHT) et al.	**400G E2E WAN Services** Architecture, Technology and Control Systems
NRE-111	Joe Mambretti (STARLIGHT) et al.	**Global Research Platform**

NRE-128: High Performance Networking with the Sao Paulo Backbone SP Linking 8 Universities and the Bella Link

- The total bandwith from Brazil to the US is expected to rise to 1.2 Tbps by the fall of 2025
- There is also a connection to Chile, to the Vera Rubin Observatory's Large Synoptic Survey Telescope (LSST)), and a connection to Buenos Aires in Argentina (RNP)

Demonstration GOALS

- Improve on last year's bandwidth between Sao Paulo and the USA when 550 Gbps (total bandwidth) was reached at SC24
- At SC25, the target will be 700 Gbps: a new record throughput connecting Brazil and the USA using academic networks
  - Simultaneous with ~80 Gbps to CERN
- 400 Gbps equipment also is being purchased and deployed, and it should be possible to demonstrate 400G connections in the state of Sao Paulo

Rednesp Intercontinental Connections

NRE-102: King Abdullah University of Science and Technology (KAUST) : Middle-East 800Gbps Transcontinental eScience Bridge

- The KAUST Global Exchange is an important transcontinental eScience bridge that helps address critical infrastructure resilience challenges, and supports the Global Research and Education Networks' (GREN) academic collaborations
- KAUST is developing a new Data Transfer Service (DTS) platform, aimed at accelerating data-intensive science between the Middle East region and global collaborators
  - Preliminary lab tests indicate the tuned DTS stack sustains ≥95 % line-rate over 400 Gbps, foreshadowing similar efficiency up to ~800 Gbps, while revealing TCP/QUIC parameter sets that minimize buffer-bloat over three-continent spans
- The demonstration objective is to validate 400 Gbps disk-to-disk transfers and possibly above, in cooperation with partner Ciena, provisioning equipment to enable a temporary additional international optical circuit between the KAUST main campus and its PoPs in NetherLight, Amsterdam and the SingaREN Open Exchange (SOX) in Singapore
- KAUST is seeking Saudi providers of a new 400G or 800G circuit that will be evaluated using the DTS, as a proof of concept for improving KAUST's network resilience and increase its capacity

NRE-132 AmLight 2.0: Flexible Control, Deep Visibility, and Programmability @ Tbps!

AmLight-ExP Team, FIU/Rednesp/VeraRubin/RNP <sdn@amlight.net>

- ❐ AmLight 2.0 uses optical spectrum and leased capacity to build a reliable, leading-edge R&E network infrastructure to enable research and collaboration among the US, Latin America, the Caribbean, and South Africa
- ✸ In 2021-3 AmLight deployed its custom-made SDN programmable fabric to enable network experimentation and flexible forwarding pipelines, fully controlled by AmLight's SDN controller named Kytos-ng [https://kytos-ng.io]
- ❐ With 20 programmable switches in production, AmLight generates millions of telemetry reports per second, which are used for daily troubleshooting, monitoring, and Machine Learning research
- ✸ The P4 In-band Network Telemetry (INT) solution built via Kytos-ng supports the team's work by identifying under-provisioned buffers, microbursts, evaluating perfSONAR test results, and packet tracing to evaluate QoS policies

✸ In 2024/2025, AmLight 2.0 expanded its footprint: to Argentina with a 400G optical channel, and to Brazil with an amazing 1.1Tbps capacity over its optical spectrum on the Monet submarine cable

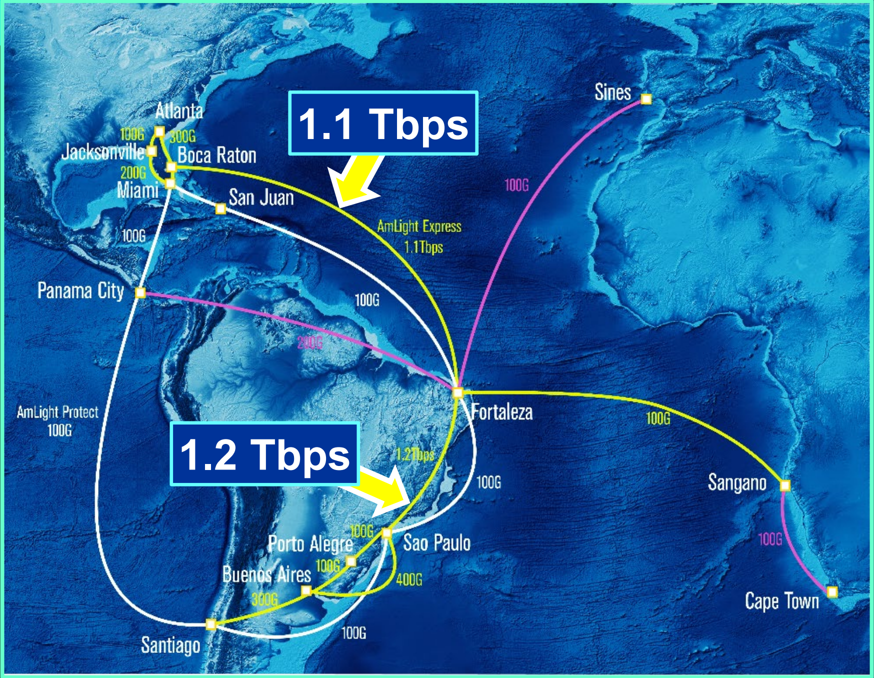

NRE-125: A Next Generation Multi-Terabit/sec Campus and Global Network System for Data Intensive Sciences

GNA-G Data Intensive Sciences and AutoGOLE/SENSE Working Groups

- **Development Trajectory:** Parallel developments + mission-driven progressive interfacing and system-level integration
- **Overarching Concept: Consistent Network Operations:**
  - Stable load balanced high throughput workflows cross optimally chosen network paths
  - Provided by site-resident services dynamically interacting with network-resident services
  - Up to preset or fliexible *high water marks* to accommodate other traffic
- Responding to (or negotiating with) site demands from the science programs' principal data distribution and management systems
- **Architecture: Data Center Analogue**
  - Classes of "Work" (work = transfers, or overall workflow), defined by task parameters and/or priority and policy
  - Adjust rate of progress in each class to respond to network or site state changes, and "events"
  - Moderate/balance the rates among the classes to optimize a multivariate objective function with constraints

NRE-125: Global Network System for Data Intensive Sciences with Multi-Tbps Campuses and Regionals (II)

- Integrated operations and orchestrated management of resources: interworking with and advancing the site (Site-RM) and network resource managers (Network-RM) developed in the SENSE program
- An ontological model-driven framework with integration of an analytics engine, API and workflow orchestrator extending work in the SENSE project, enhanced by efficient multi-domain resource state abstractions and discovery mechanisms
- Fine-grained end-to-end monitoring, data collection and analytics, supporting applications with path selection + load balancing mechanisms optimized by machine learning.
- The integration of Qualcomm Technology's GradientGraph (G2) with Source-based segment routing (PolKA, SRv6). These will be used to demonstrate how applications can benefit from the intelligent routing, rate limiting, and service placement decisions
- AI/ML-based optimization via Hecate (ORNL); Integrated with PolKA
  Multi-Agent LLMs, Generative AI for Network Management
- Demonstrating how the integration with science networks such as integration of the LHC Rucio / FTS data management system, and AutoGOLE / SENSE virtual circuit and orchestration services / and comprehensive monitoring (e.g. Inmon) may be used towards optimizing large-scale data transfers, and sharing data among scientists at sites across the globe

NRE-132: AmLight & the Vera Rubin Observatory
The Legacy Survey of Space and Time (LSST)

- **Rubin Observatory will take hundreds of images of the Southern Hemisphere sky, every night for ten years,** for a survey called the Legacy Survey of Space and Time (LSST)
- **Data from these images will be used by astronomers around the world to make countless discoveries,** but VRO was specifically designed to advance four science areas:
  - **Understanding the nature of dark matter and dark energy**
  - Creating an inventory of the Solar System
  - **Mapping the Milky Way**
  - Exploring the transient optical sky, i.e. **studying objects that move or change in brightness**
- **Rubin Observatory will produce about 20 terabytes of data every night** during the ten-year survey. **A new high resolution (3.2 Gigapixel) image read out every 4 seconds.** ***Alerts !***
- **By the end of the survey, the data set will be enormous — about 60 petabytes!** Most of the astronomers who make discoveries will access the data using an online portal called **the Rubin Science Platform**
- **They won't need expensive equipment or computing power, just an internet browser.** This is the first time this much astronomical data will be available to so many people: **there's no telling what discoveries scientists will make using VRO !**

NRE-132 AmLight 2.0: Flexible Control, Deep Visibility, and Programmability @ Tbps!

- Amlight 2.0 now offers opportunities to the academic community, including:
  - Real-time telemetry to detect packet drops and microbursts that could limit science data flows
  - Dynamic provisioning, complex pathfinding, per-packet network programmability and telemetry
  - Integration: with academic orchestrators such as AutoGOLE/SENSE, and with testbeds such as FABRIC and RARE

GOALS:

- Showcase the newest open SDN architecture, programmable fabric, and the capabilities of Kytos-ng and Amlight's Inline real-time telemetry
- Demonstrate APIs to integrate with science workflows
- Present the challenges, implementation, lessons learned, and Opportunities for collaboration
- In collaboration with AmLight members Vera Rubin Observatory, Rednesp, RNP, FLR, and SPRACE, demonstrate Amlight's network capabilities by generating up to 1 Tbps of traffic, while leveraging our DTNs in Chile and Sao Paulo to the SC25 showfloor, landing at the Caltech booth

PolKA: An Efficient Source Routing Approach to Meet the Requirements of Data Intensive Sciences

GNA-G
Global Network Advancement Group

No tables in the core	Open source/ Interoperable	Support in prog. switches	Fixed length header	Topology agnostic multipath routing

PolKA: Polynomial Key-based Architecture for Source Routing Implementation

- A single user-defined encoded/decoded label defines the path: identifying each switch and port along the way
- Polynomial Residue Number System (RNS)
- Chinese Remainder Theorem (CRT)
- Packet forwarding based on mod operation: remainder of division, using switch CRC hardware for speed
- Transparently traverses fixed function switches in the path as needed

- Easy Setup of paths/tunnels using a standard CLI
- Open Source Implementation in RARE/freeRtr
- Many powerful network applications: Proof of transit, PBR, multipath, multicast, failure protection, telemetry, ...

How does PolKA work?

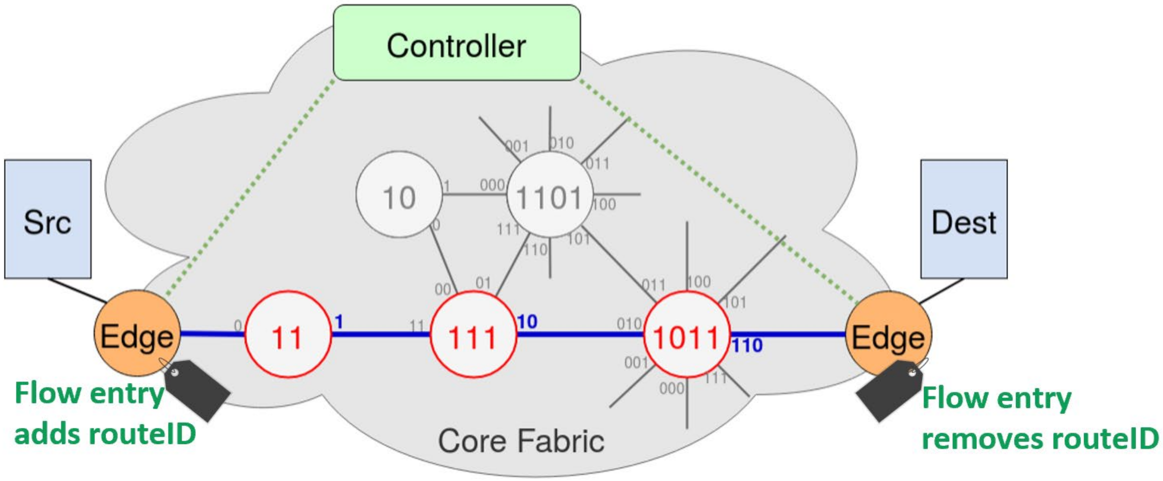

- The **Controller** installs **flow entries** at the edges to add/remove ***routeIDs***.

- The **Controller** calculates the ***routeID*** using CRT:
  - Complexity: $\mathcal{O}(len(M)^2)$, where $M(t) = \prod_{i=1}^{N} s_i(t)$

R = 10000

routeID

- Forwarding:

portID = < routeID >$_{\text{nodeID}}$

1	=	$<10000>_{0011}$
10	=	$<10000>_{0111}$
110	=	$<10000>_{1011}$

nodeID polynomials

$$s_1(t) = t + 1 = 11$$
$$s_2(t) = t^2 + t + 1 = 111$$
$$s_3(t) = t^3 + t + 1 = 1011$$

portID polynomials

$$o_1(t) = 1$$
$$o_2(t) = t = 10$$
$$o_3(t) = t^2 + t = 110$$

Calculate routeID with CRT

$$t^4 \equiv 1 \mod (t+1)$$
$$t^4 \equiv t \mod (t^2+t+1)$$
$$t^4 \equiv (t^2+t) \mod (t^3+t+1)$$

$$t^4 = 10000$$

NRE-130: A Case Study: Data-Intensive Traffic from the Vera Rubin Observatory (Chile) to SLAC Supported by the Path Aware PolKA Network, AmLight and GP4L

UFES/IFES/RNP/FIU/Caltech Team

- Presents a collaborative case study integrating PolKA (Polynomial Key-based Architecture) path-aware networking to meet the stringent performance requirements of the Vera Rubin Observatory Long-Haul Network (LHN)
- The LHN delivers 13 GB astronomical images every 27 seconds from Chile to the US Data Facility at SLAC via a purpose-built, SLA-driven 400 Gbps infrastructure
  - Traffic steering and fault resilience are essential to ensure continuous sub-7-second image delivery under the 180ms+ round-trip delay
- Path-aware routing represents a paradigm shift, from traditional routing which primarily relies on "shortest path" metrics;
  - Instead it empowers network devices to select routes based on a richer set of criteria including: real-time performance, available bandwidth, latency, jitter, and even state and policy-based constraints
- For the Vera Rubin LHN, PolKA's source routing with path encoding could enable fine-grained path selection, real-time telemetry, and rapid rerouting capabilities that align with the LHN's operational objectives and guarantees

NRE-130: A Case Study: Data-Intensive Traffic
from the Vera Rubin Observatory (Chile) to SLAC
Supported by the Path Aware PolKA Network and AmLight

- **GOAL:** Demonstrate that a path-aware PolKA protocol, deployed on a global programmable network testbed (GP4L), can meet the stringent performance requirements of the Vera C. Rubin Observatory use case by:
  1. Showcasing mock data transfers and dynamic steering of VRO workloads using the PolKA Path-Aware Networking architecture
  2. Enabling agile and resilient path reconfiguration for data-intensive flows across multiple intercontinental 100G network domains, ensuring sustained performance even under link failure; and
  3. Stress-testing PolKA's ability to handle the VRO dynamic workflow, from UFES (Brazil) to the Caltech booth via an intercontinental network under the background traffic generated by the data tsunami challenge

- **IMPACT**
  - ❐ By leveraging the AmLight-ExP programmable infrastructure — including telemetry collectors, perfSONAR, and SDN-controlled 100G links spanning Latin America and the U.S, this NRE shows how PolKA and AmLight together deliver higher throughput with delivery guarantees for data-intensive science workflows
  - ❐ The demonstration highlights packet steering over multiple paths, flow isolation for the Rubin transfers, and visualizations of path utilization and failover recovery
  - ❐ This work advances the efforts to build programmable, path-aware networks that meet the needs of emerging Tbps-scale scientific workflows

NRE-118 Cryo-EM Without Borders:

Empowering Distributed Research via AmLight and Collaborative (Human) Networks

- ❒ Cryo-Electron Microscopy (Cryo-EM) is revolutionizing structural biology, enabling near-atomic resolution of macromolecules
- ❒ Access remains limited in Sub-Saharan Africa and South America due to cost, infrastructure, and expertise gaps
- ❒ Centralized imaging core facilities exist in Europe/US, but distance creates barriers for Global South researchers
- ✸ AmLight's advanced network infrastructure now enables remote Cryo-EM workflows, supporting real-time data transfer, analysis, and collaboration across continents

Use Case: South African Researchers prepare an anti-HIV antibody which can become a vaccine candidate

1. Sample is purified at Univ. of Cape Town (UCT)
2. Data is acquired/preprocessed at Univ. of California Santa Cruz
3. Datasets are transferred back to UCT for analysis

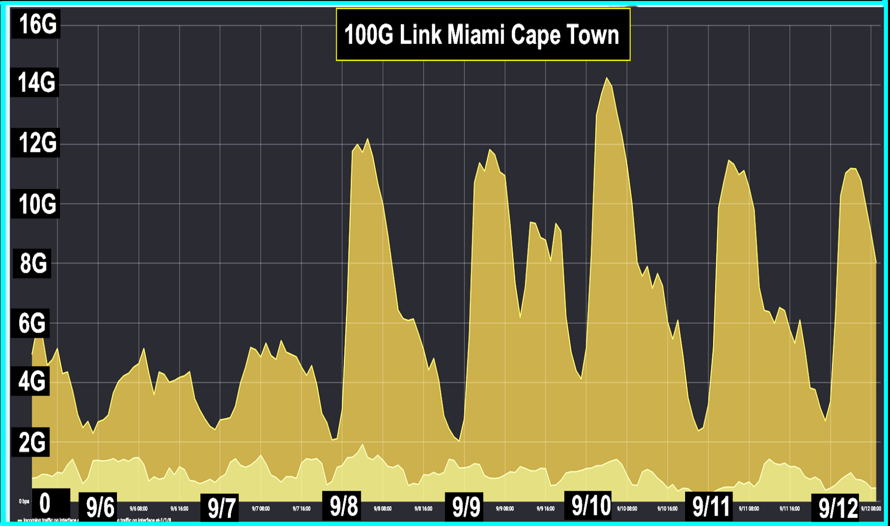

NRE-118 Demonstration: AmLight in Action

- ❒ **Goal:** Demonstrate high-throughput trans-continental Cryo-EM data processing "on-the-fly" with AmLight's Terabit Network
- ✸ Removing the barriers to groundbreaking research by intercontinental collaborations
- ✸ With diverse areas of expertise, while leveraging high end instumentation
- ❒ **Setup:** Disk-to-disk transfer between the USA ↔ Cape Town, South Africa
- ❒ **Tools:** Data Transfer Nodes (DTNs), Globus for end-user workflow emulation
- ✸ **Targets:** Exceed prior 100Gbps benchmark by the SANReN team; showcasing real-world performance tuning
- ✸ **Network Path:** UCT (Cape Town) ↔ UC Santa Cruz (California) via the AmLight/Miami 400G link over the Angola ACS submarine cable

✸ **IMPACT:** Demonstrating how AmLight and global NRENs empower distributed, collaborative science, reducing inequities in access and enabling cutting-edge biomedical research

Cryo-EM Workflow

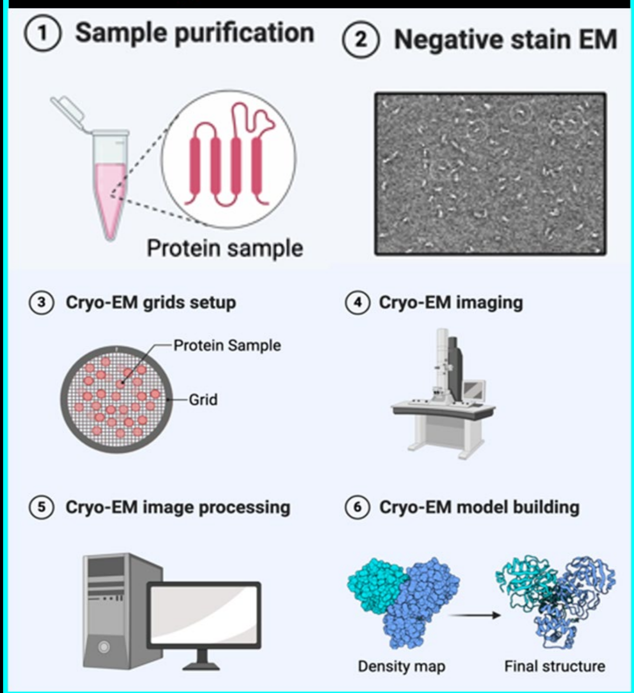

FIU Florida International University · USC University of Southern California · AmLight ExP Americas Africa Lightpaths Express & Protect · University of California Santa Cruz · RNP · nicis SANReN · TENET Tertiary Education and Research Network of South Africa · CSIR Touching lives through innovation · science & innovation Department: Science and Innovation REPUBLIC OF SOUTH AFRICA

NRE-127: Onboard AI and Federated Learning
for Cosmic Ray Event Detection
Caltech, U. Delaware, MIT, Oligo Collaboration

Project Background and Scope

- ❐ Institutions worldwide collect vast amounts of cosmic ray data, but geographical distribution and data management challenges limit effective collaboration
  - ❐ Traditional centralized approaches limit scalability & real-time processing capabilities
- ✴ We are developing a prototype Multi-site federated learning framework for Cosmic Ray particle classification across distributed institutions, with real-time processing capability and space deployment potential
- ❐ Institutions maintain local data while collaboratively building comprehensive AI models through parameter sharing

IMPACT

- ❐ Cosmic rays provide crucial insights into high-energy physics, space weather, and astrophysical phenomena
- ✴ Collaborative analysis with federated learning across institutions could accelerate discoveries in astroparticle physics and space science, and help track geomagnetic storms in real-time
- ❐ As a bonus, this approach provides data privacy benefits for applications where data sensitivity is a concern

(1) CosmicWatch: Desktop Muon Detector (v3X)
(2) CREDO: Cosmic Ray Extremely Distributed Observatory

NRE-127: Onboard AI and Federated Learning for Cosmic Ray Event Detection

Federated Learning Architecture: Three-Site Demo to follow simulations at three sites: Caltech, MIT, U. of Delaware

- ✴ Collaborative Learning Emulation Study
  - ❐ Each site trains on their specialized data
  - ❐ Local data processing is maintained – no raw data needs to be transmitted
  - ✴ Model parameters are shared via federated averaging: All sites learn to recognize all |particle types through collaborative knowledge sharing
- ✴ Results: Simulation shows each site achieves >98% accuracy classifying all particle types
- ✴ Demonstrating potential for effective colla-borative learning through distributed processing

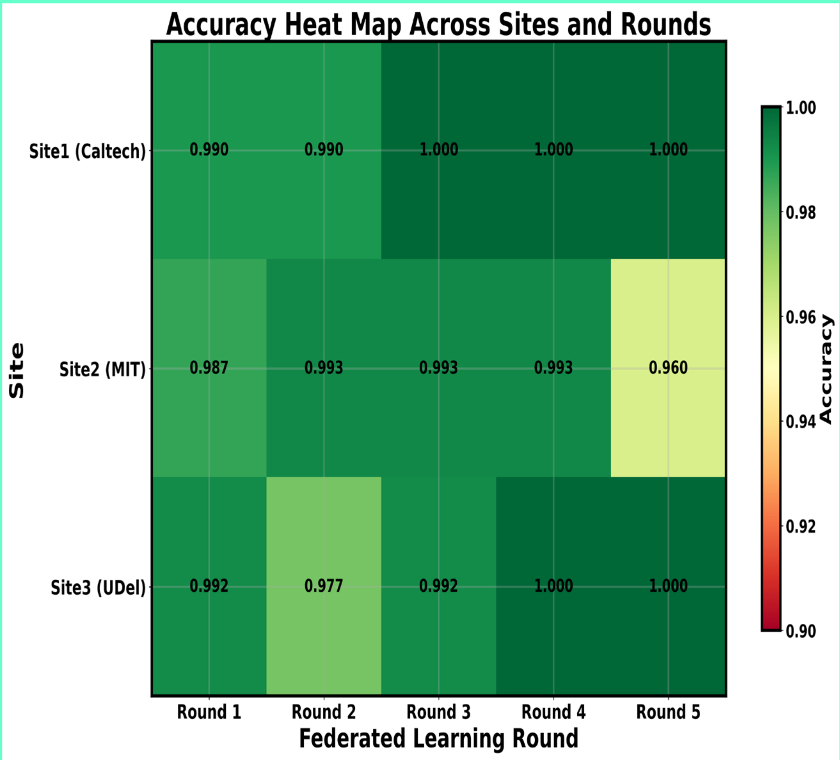

✸ SCinet Network Integration

- ❐ Primary Circuit: 100 Gbps connection for real-time data transmission and model coordination
- ❐ Scalability: Architecture designed to support 3+ distributed sites with varying data volumes and processing capabilities
- ❐ Protocols: IPv6, Layer 2/3 switching optimized for data transmission and federated learning

NRE-127: Onboard AI and Federated Learning
for Cosmic Ray Event Detection

Tasks Completed to Date

- ❒ **Framework Development:** of the federated learning architecture
- ❒ **Network Integration:** 100 G network requirements specified for SCinet integration
- ❒ **Scalability Simulation: Federated learning simulation across 3 sites using 2,354 cosmic ray images** on NRP Nautilus QAIC, validating the distributed processing framework
- ❒ **Space Deployment Concept:** Framework designed for autonomous operation in space environments

✸ SC25 Live Demonstration Goals:

Real-time federated learning demonstration with actual cosmic ray detectors at 3 sites, leveraging data from the CREDO smartphone application's citizen science network

Components:

- ✸ **Live Data Processing:** Real-time cosmic ray detection and image classification **using the CREDO smartphone app**
- ❒ **Federated Learning:** Multi-site model training through **parameter sharing across institutions**
- ❒ **Network Performance:** High-bandwidth data transmission & model parameter exchange via SCinet
- ❒ **Interactive Visualization:** Real-time display of collaborative learning progress across sites
- ✸ **Audience Experience:** See live federated learning in action, with real cosmic ray detectors collaborating across three institutions

NRE-106: Multi-Resource Cyberinfrastructure Services for Science Domain Workflows via SENSE

- The SENSE system enables the automated provision of Cyberinfrastructure services with integrated network, compute, and storage resources.
  - The SENSE API enables science workflow agents to initiate, monitor, and optimize the lifecycle of their workflow
- SENSE includes a model-based orchestration system which operates between the SDN layer controlling the individual networks/end-sites, and science workflow agents/middleware.
- The key SENSE components are the "Orchestrator", "Network Resource Manager" and "Site Resource Manager" which together enable advanced features including multi-resource services, automated control, real time optimizations, and full life-cycle monitoring.

- This demonstration will show SENSE in operation across a variety of infrastructures including:
  - Networks and Exchange Points: ESnet, Internet2, GNA-G AutoGOLE, CENIC, PacificWave, AMPATH/AMLight, StarLight, NetherLight Exchange
  - Edge Systems: National Research Platform Kubernetes Cluster, Data Transfer Nodes, LHC XRootD/dCached systems, Amazon Web Services, and
  - the FABRIC infrastructure which includes network embedded compute/storage resources.

NRE-106: Multi-Resource Cyberinfrastructure Services for Science Domain Workflows via SENSE

SC25 NREs, Scinet, SENSE, AutoGOLE, NRP Multi-Services Automation Map

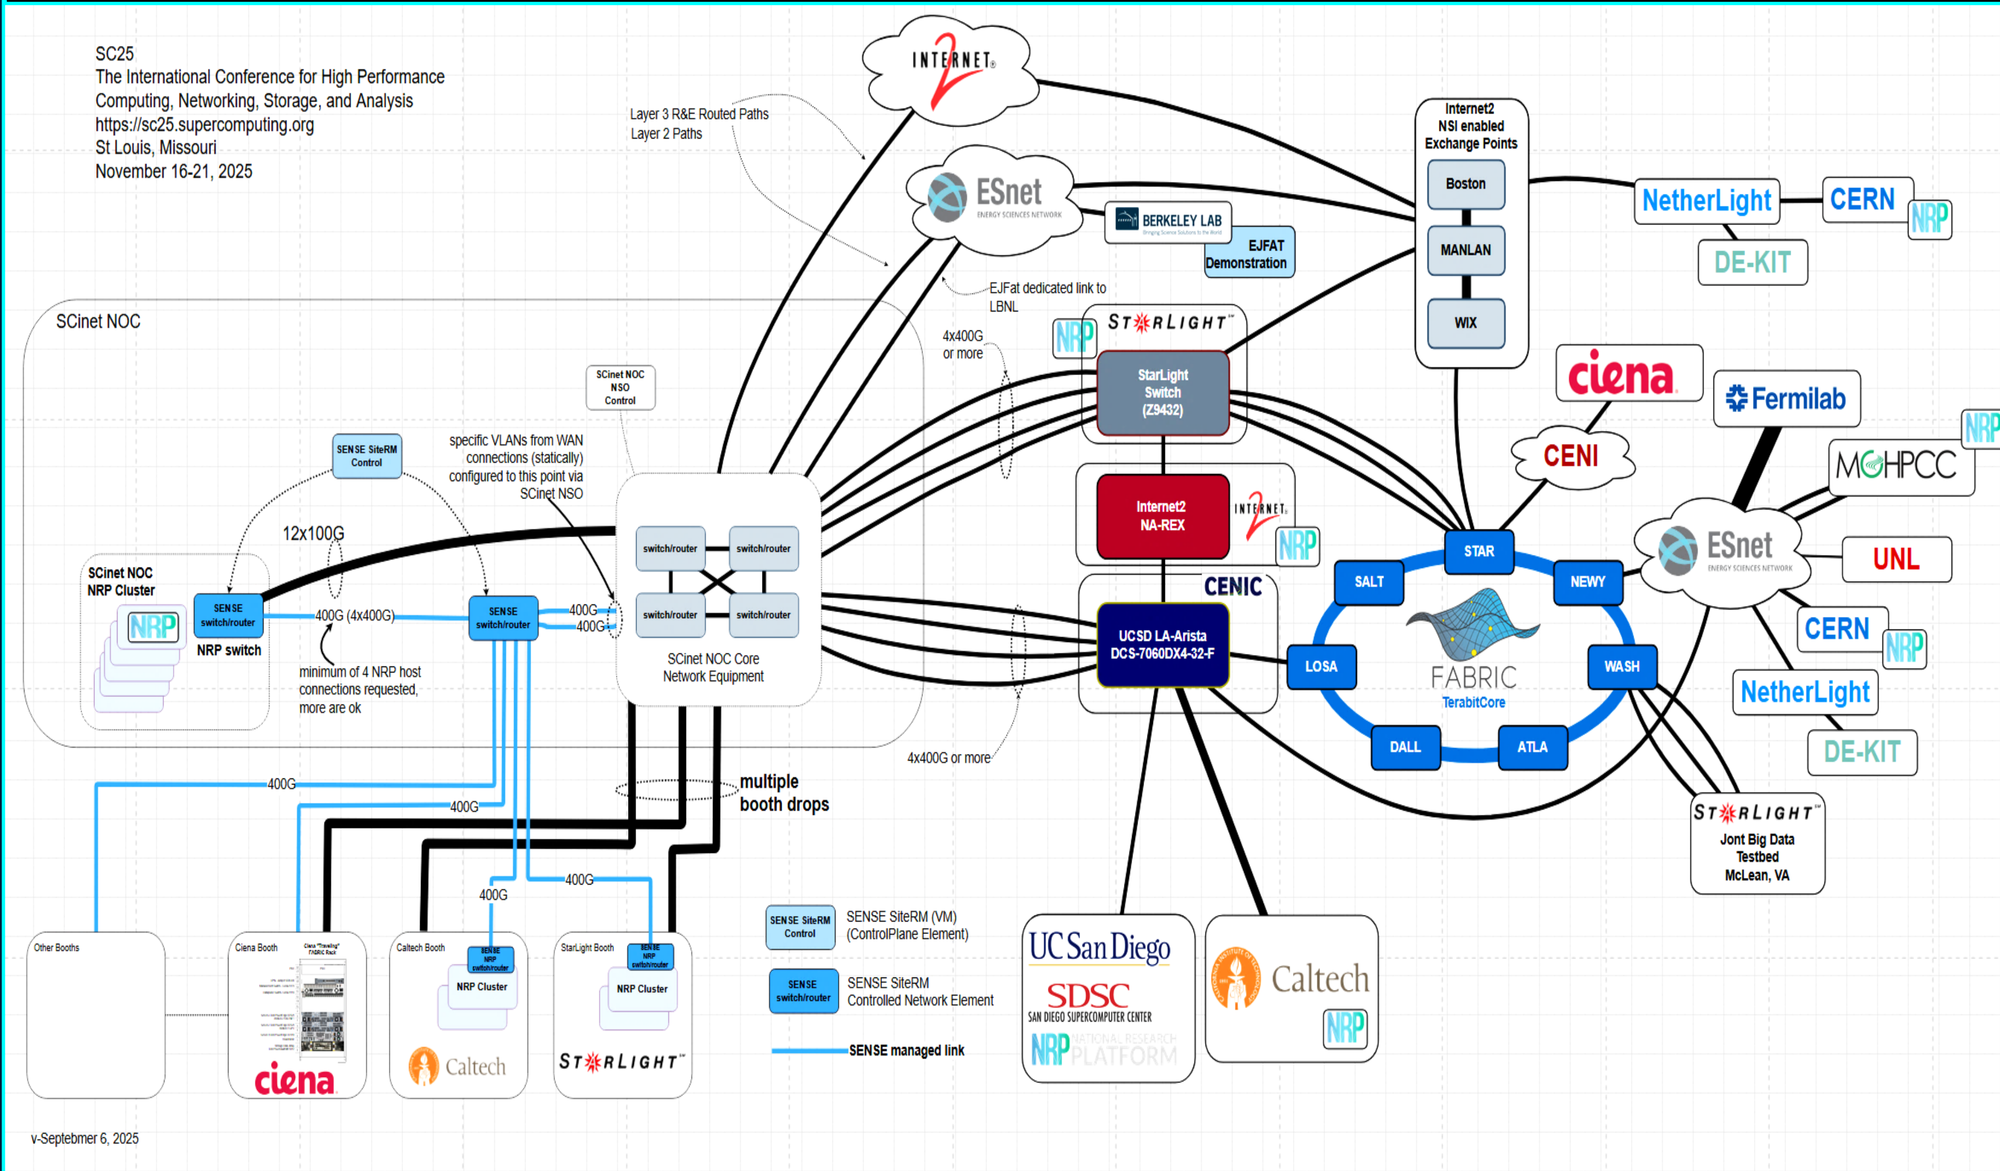

NRE-117: FABRIC: An Adaptive programmaBle Research Infrastructure for Computer Science and Science Applications

https://portal.fabric-testbed.net

- ❒ FABRIC enables cutting-edge experimentation and research at-scale in the areas of networking, cybersecurity, distributed computing, storage, virtual reality, 5G, machine learning, and science applications
- ❒ FABRIC has created a rich environment to support this wide variety of activities:
  - ✸ FABRIC includes a distributed set of equipment at commercial collocation spaces, national labs and campuses, spanning 33 sites across the US, CERN, and the Universities of Amsterdam, Bristol, Tokyo
  - ❒ Each site has large amounts of compute and storage, interconnected by high speed, dedicated optical links; It also connects to specialized testbeds: 5G/IoT PAWR, NSF Clouds; the Internet and several HPC facilities
  - ✸ FABRIC programmable networking includes automated network provisioning, P4 controlled elements, and realtime monitoring
  - ✸ The Ciena Traveling FABRIC Rack will be connected to the SCinet network infrastructure, local exhibit booth, and the rest of the FABRIC via SCinet [*] WAN circuits, and will be part of the demonstrations and research activities

[*] SCinet, an advanced cutting-edge large scale infrastructure, is designed, custom-built, operated and administered every year for the SC conferences by expert volunteers from industry, academia, and governments.

NRE-101: DOE Integrated Research Infrastructure (IRI)
IRI Compute Workflow Portability

- The DOE Integrated Research Infrastructure (IRI) initiative, aims to establish a seamless, secure, programmable scientific ecosystem connecting experimental instruments, HPC centers, storage platforms, and network infrastructures
- Overarching goal is to enable integrated, cross-facility scientific workflows that reduce the time-to-insight and enhance collaborative discovery
- The IRI architecture plans to support multi-site experimentation and data-intensive workflows by standardizing APIs and protocols for compute, storage, data movement, orchestration, and identity management
- Core efforts include the Pathfinder projects, which demonstrate practical integrations across DOE light sources, fusion experiments, climate simulations, and the HPC and data centers at ALCF, OLCF, and NERSC
- The IRI design and testing strategy foresees integration with advanced academic infrastructures such as the National Research Platform (NRP) and FABRIC testbed. These platforms offer flexible, distributed, and high-performance compute and networking resources that mirror DOE facility use cases, enabling early prototyping and experimentation at scale
- This NRE showcases cross-domain resource federation, real-time network provisioning, and data mobility — all orchestrated through a unified framework

NRE-119: National Research Platform and SCinet: Enabling Live, Multi-Institutional Scientific AI/ML and HPC Workflows

GOALS

- ❒ Demonstrate interactive AI workflows and scientific visualization live on the SCinet show floor via Nautilus nodes
- ❒ Highlight reproducible, multi-user JupyterHub, Kubernetes, and SLURM-based access, leveraging NRP's user-facing infrastructure for real-time research
- ❒ Showcase heterogeneous hardware orchestration with Kubernetes, including GPUs, DPUs, AMD/Xilinx Alveo and Intel Agilex and Stratix FPGAs
- ❒ Establish full production network connectivity with deterministic QoS across SCinet, StarLight, FABRIC, Caltech, and Chameleon Cloud using ESnet SENSE
- ❒ Run scientific workloads on the SCinet cluster at SC25 with guaranteed low preemption, zero resource reservation, and by using underutilized resources

Mohammad Sada et al
NREs 119-222

NRE-119: National Research Platform and SCinet: Enabling Live, Multi-Institutional Scientific AI/ML and HPC Workflows

Mohammad Sada et al
NREs 119-222

✸ Interactive Experience at SC25

- ✸ Attendees will interact with live research workflows provisioned on the Nautilus nodes at the SCinet NOC; SCinet-connected compute nodes on the exhibit floor will be directly incorporated to enable seamless execution of high-performance interactive workloads
- ❒ These will be visually demonstrated via Grafana, Prometheus, Observable, Jupyter and inMon dashboards, and immersive visualizations, and public web interfaces with full access to underlying code, AI models, and real-time outputs
- ❒ The Nautilus cluster [*] guarantees low-preemption on the SCinet nodes, ensuring these workloads maintain dedicated compute and network resources to other NREs throughout the conference

[*] With > 1400 GPUs, 50+ FPGAs (35 P4 Programmable and 7 DPUs across ~425 Nodes

✸ Nautilus supports a wide range of scientific workflows:

- ✸ 3D Cell Reconstruction with AI
- ✸ Digital Twins for Manufacturing and Space Exploration
- ✸ Wildfire Monitoring and Immersive Visualization
- ✸ LLM as a Service and GenAI Novel Hardware Architectures for Scientific Research
- ✸ Open Science Grid (OSG) Experiment Hosting: high-throughput computing workloads

NRE-120: Live High-Precision Per-Packet Kubernetes Data Telemetry from NRP to StarLight via ESnet SENSE

- This NRE links the NRP Kubernetes cluster resources at SDSC to the StarLight facility in Chicago via the ESnet SENSE Orchestrator
  - Using dedicated 100Gbps direct links integrated into the AutoGOLE global topology, NRP transmits live Kubernetes network data packets to P4-programmable Tofino2 SmartNIC endpoints at StarLight for advanced packet inspection, telemetry, and analysis.
  - Each packet originating from SDSC, where Kubernetes workloads operate, using umerous NRP resources and is routed over SENSE using orchestrated L2 paths with a dedicated QoS VLAN tag
  - Upon arrival at StarLight, programmable P4 logic running on Tofino2 returns a header-only data stream to a Xilinx Alveo FPGA running P4 and DPDK acceleration for high-throughput packet processing and offload, ingesting it into a Kafka streaming database in real time
- This enables live monitoring and analysis of the Kubernetes cluster network behavior and workload communication patterns

NRE-120: Real-Time In-Network Machine Learning and P4 Testbed Deployment on FPGA SmartNICs, DPUs, and Switches

GOALS: To Demonstrate and Showcase:

- Scientific workflows deployed on the NRP Kubernetes cluster generate significant network traffic among containers and microservices. Real-time telemetry of packet headers collected at StarLight enables dynamic detection of network congestion, packet loss, or anomalous traffic patterns
- Seamless, high-throughput (100Gbps) connectivity between the NRP Kubernetes cluster at SDSC and StarLight in Chicago, orchestrated by ESnet SENSE within the AutoGOLE topology that provision QoS-enabled VLAN Layer 2 paths across multi-domain research infrastructures
- Advanced real-time packet telemetry extraction and analysis by leveraging P4-programmable Tofino2 SmartNICs at StarLight, which trim payloads and extract header metadata from live Kubernetes network traffic. This telemetry is accelerated via FPGA and DPDK and streamed into Kafka for scalable, low-latency monitoring and analysis.
- Highlight practical scientific and AI/ML workflow use cases enabled by live telemetry, such as dynamic network performance monitoring, and anomaly detection, and Kubernetes workload optimization across geographically distributed sites
- Integration of programmable data planes, FPGA acceleration, and intelligent network orchestration, validating the capabilities of ESnet SENSE and FABRIC to provision Layer 2 paths both within datacenters and over wide-area R&E networks

NRE-121: Real-Time In-Network Machine Learning and P4 Testbed Deployment on FPGA SmartNICs, DPUs, and Switches

- ❐ SDSC and NRP are deploying a large-scale programmable testbed integrating P4-programmable SmartNICs, DPUs, and switches for real-time in-network computing R&D and systems development

- ❐ Featuring 7 IGROK Nodes each with a Bluefield2 2X100G DPU, integrated in the SENSE topology to enable distributed monitoring & control
  - ✸ This highlights two closely related efforts:

  (1) Real-time machine learning (ML) inference embedded directly in the network data plane using fixed-point arithmetic and Taylor series approximations on AMD/Xilinx Alveo U55C FPGA SmartNICs, and

  (2) Deployment of a flexible, scalable P4 experimentation environment spanning geographically distributed sites

GOALS: To Demonstrate

- ✸ An ML framework enabling model execution by combining
- ❐ Control-plane integration via Kubernetes with efficient, low-latency packet processing in the P4 data plane
- ❐ Programmable telemetry and sFlow-based observability

NRE-121: Real-Time In-Network Machine Learning and P4 Testbed Deployment on FPGA SmartNICs, DPUs, and Switches

Environment

- These inference workloads run live on the Nautilus cluster, which connects to global research infrastructure including SCinet, FABRIC, Chameleon Cloud, AMPATH, and StarLight through the AutoGOLE topology
- The broader testbed features 32 Alveo U55C FPGAs, Bluefield2 DPUs located at PacWave peering points, U280 FPGAs at FABRIC sites, a 400G Tofino2 switch running SONiC PINS, and virtual P4 switches built on BMv2.

Interactive Experience at SC25

- Live ML inference on FPGA SmartNICs processing packet streams over SCinet , the Scinet cluster and the NRP testbed
- Visualizations in Grafana, JupyterHub, and Prometheus displaying inference latency, model adaptation, and network throughput
- Dynamic topology control, with L2 paths established and torn down in Real time between SCinet and remote FABRIC, AMPATH, and PacWave sites
- Hands-on use of the OpenNIC shell and ESnet SmartNIC framework for programming/monitoring

Agile Traffic Migration to Better Paths

Al-Najjar et al., Framework for Integrating Machine Learning Methods for Path-Aware Source Routing, INDIS, SC24

Anees Al-Najarr et al

- Multiple AI tools that can predict network performance
- Integrate with network tools for implementing action on networks
- Integrate Path-Aware Source Routing (PolKA) for Selection Decisions
- Collaborations across network providers for building data sets
  - For training, intelligence; Foundation model construction

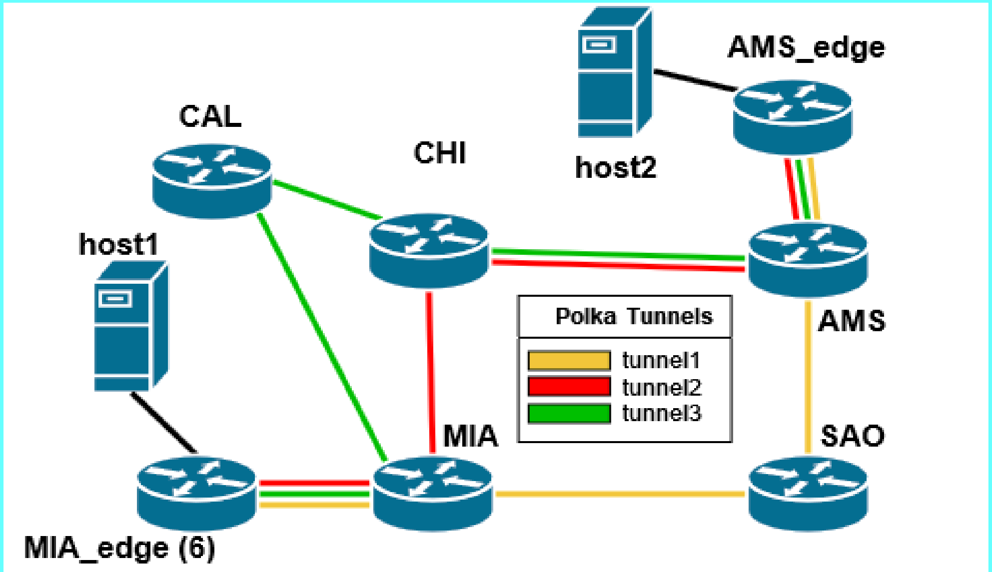

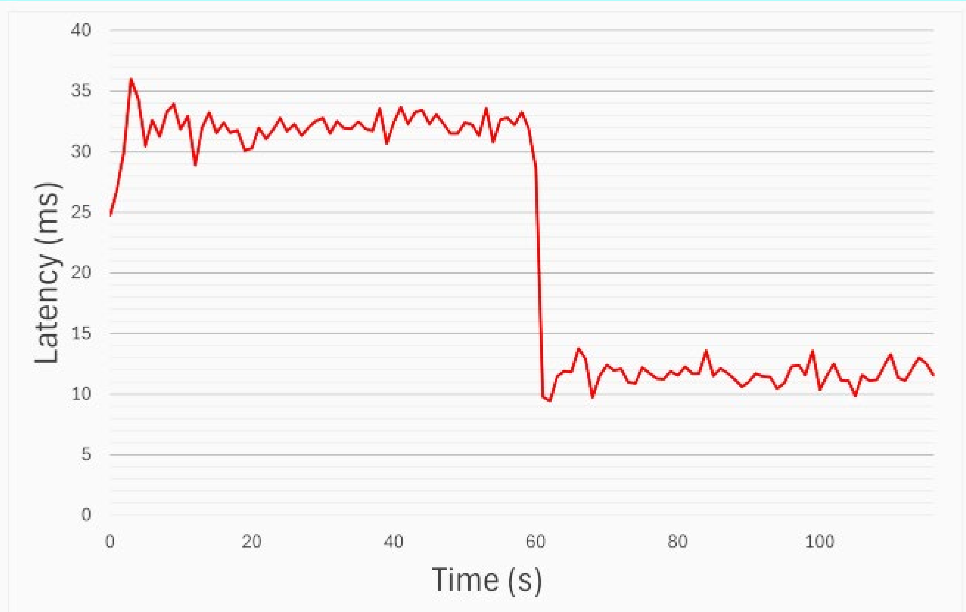

Multi-Agent LLMs for Network Management

Mariam Kiran et al

OFCnetLLM: Large Language Model for Network Monitoring and Alterness

Yoon et al., https://www.arxiv.org/abs/2507.22711

Generative Agentic AI

- ✸ Overall Situational Awareness, Analysis and Problem Solving via:
- ✸ Multi-agent reasoning capability
- Pattern recognition across multiple network databases
- Chain-of-thought problem diagnosis
- Agents operate on local nodes
- Data remains secure and isolated
- Improved database query efficiency

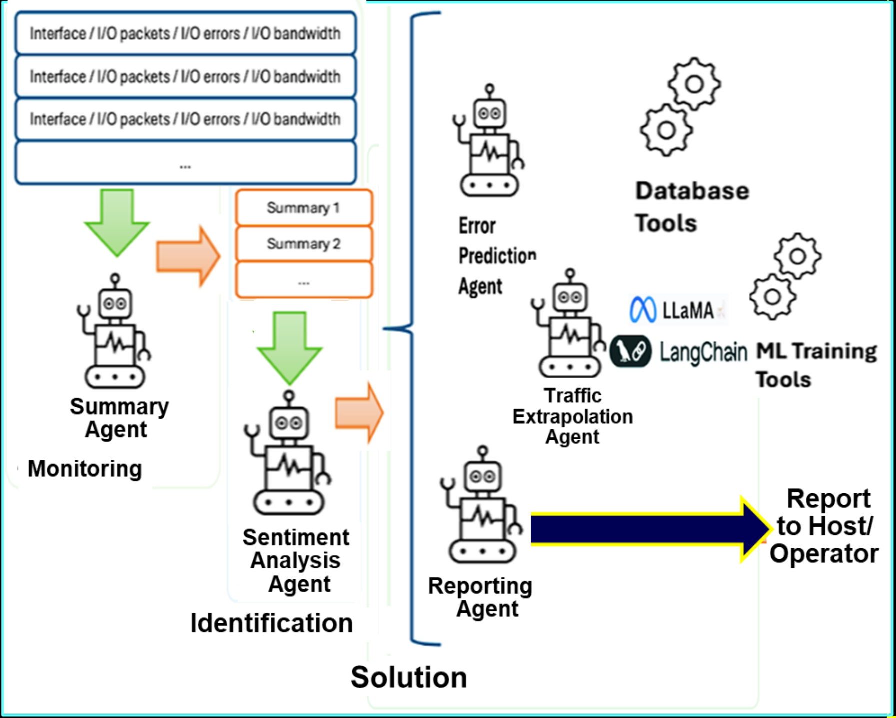

ChatGPT for Networks

USER: No error ?

OFCnetLLM: That's great news ! I am glad that all systems are functioning normally

However I'd like to provide some additional context and insights: Currently our network is experiencing zero input and output errors, which indicates a high level of reliability and stability.
This is excellent news for our users and applications

❒ To further validate this information, I can run a quick diagnostic to ensure everything is running smoothly

Would you like me to perform a:

1. System Health Check: A comprehensive review of all components to identify any issues
2. Network Traffic Analysis: A detailed examination of traffic patterns to detect any anomalies or bottlenecks
3. Performance Monitoring: Real-time monitoring of key performance Indicators (KPIs) such as latency throughput and packet loss

➡ Tell me which of the above to run. All ?

NRE-122: Agentic AI with QAIC Ultra Cards for HPC Cluster Management and Resource Provisioning

GOALS

- Enable autonomous resource provisioning in HPC clusters by leveraging agentic AI running on Qualcomm AI 100 Ultra cards
- Showcase retrieval-augmented generation for dynamic documentation and operational support, eliminating manual documentation burdens
- Integrate AI agents tightly with Nautilus Kubernetes orchestration to optimize workload scheduling and cluster efficiency
- Demonstrate network debugging capabilities using LLMs with agentic tools from measuring and monitoring data flows

Interactive Experience at SC25

- Attendees will experience how agentic AI and LLM tools using Qualcomm Cloud AI 100 Ultra cards, can be launched on NRP in seconds
- Through a browser-based interface, visitors can interact with AI agents to provision resources, troubleshoot networks, and retrieve documentation using natural language
- The session offers hands-on access to real-time demos, allowing participants to explore how LLMs can enhance HPC cluster operations and intelligent network management

Global Platform for Labs (GP4L)

Motivation ⇨ Mission

Marcos Schwarz RNP

- **How Can We Increase the rate of progress of R&E Networks without interfering with production ?**
- **Develop and operate end-to-end/multi-domain orchestration services ?**
  - **Resource reservation (guaranteed bandwidth)**
  - **Resource provisioning (Circuits, VRFs)**
  - **Underlay observability**
  - **Dynamic traffic steering/flow control/engineering**
  - **Dynamic creation of L3 VPNs**
  - **Closed loop multi-domain visibility/intelligence/controllability**
- **How can we create/sustain an integration initiative/platform to propose and validate next generation protocols and services?**

Proposition: **Build programmable platforms to experiment/develop**
On persistent pre-production networks
While leveraging industry/R&E open ecosystems

Global Platform for Labs (GP4L)

SC25 ST. LOUIS NOV 16–21

Marcos Schwarz RNP

GP4L Topology

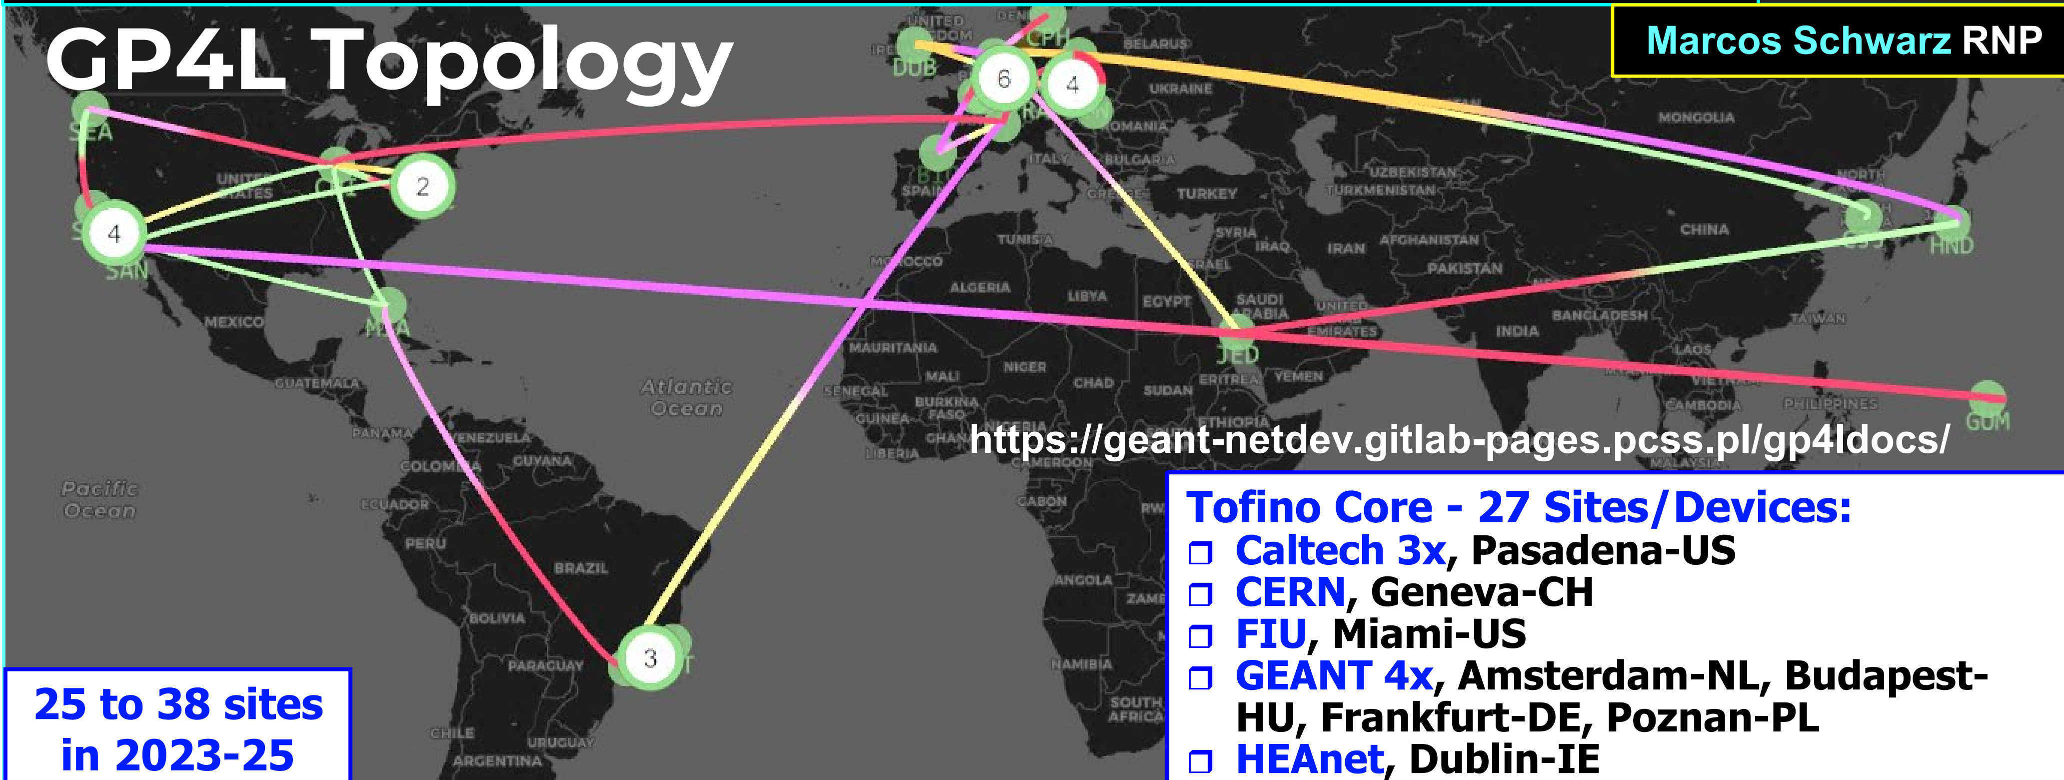

https://geant-netdev.gitlab-pages.pcss.pl/gp4ldocs/

25 to 38 sites in 2023-25

BlueField-2/DPDK Islands – 7 Sites/Devices [New]:

- **Pacific Wave/UCSD**, Chicago-US, GUAM-GU, Los Angeles-US, New York-US, San Diego-US, Seattle-US, Sunnyvale-US

x86/DPDK Islands – 4 Sites/Devices:

- **FABRIC**, Miami-US
- **2x GEANT**, Paris-FR, Prague-CZ
- **KAUST [New]**, Saudi Arabia-SA

Tofino Core - 27 Sites/Devices:

- **Caltech 3x**, Pasadena-US
- **CERN**, Geneva-CH
- **FIU**, Miami-US
- **GEANT 4x**, Amsterdam-NL, Budapest-HU, Frankfurt-DE, Poznan-PL
- **HEAnet**, Dublin-IE
- **KDDI**, Tokyo-JP
- **KISTI**, Daejeon-KR
- **RENATER**, Paris-FR
- **RNP**, Rio de Janeiro-BR
- **SC25 [New]**, Saint Louis-US
- **SouthernLight**, São Paulo-BR
- **StarLight**, Chicago-US
- **SWITCH 6x**, Geneva-CH
- **Tennessee Tech**, Cookeville-US
- **UFES**, Vitória-BR
- **UFMG [New]**, Belo Horizonte-BR
- **UMd/MAX**, College Park-US

Global Platform for Labs (GP4L)
SC25 NRE-126 Goals

Marcos Schwarz RNP

Demonstrate a global multi-vendor, multi-OS, multi-controller, multi-service programmable network testbed and its capabilities:

To accelerate the development, deployment and adoption of new protocols, operations and management paradigms

- **Support intercontinental high capacity transfers (100G to 400G) of next generation protocols, such as PolKA**
- **Support for 100G+ software dataplane [*] based on multiple control plane options options including the addition of Holo Routing**
- **Showcase the integration of existing GP4L automation and orchestration use cases with the GNA-G GREN Map initiative**

[*] 100G+ Software Dataplane

- Validate VPP integration with multiple open/commercial control planes
- Develop a reference blueprint for exploring traffic engineering & telemetry
- Allow integration of existing [1] or new IBN solutions leveraging well established protocols/APIs

Components			Features/Protocols/APIs
(Open Source) freeRtr Holo SONiC (FRR) VyOS (FRR)	(Commercial) Arista cEOS Cisco xRD Juniper cRPD Nokia SR Linux		EGP (BGP) IGP (BGP, ISIS, OSPF) Traffic Engineering (BGP, PCEP) NETCONF, RESTCONF, gNMI
Linux Control Plane			
VPP			**Segment Routing (SR-MPLS, SRv6)** Telemetry (SNMP, sFlow)
Linux Networking	RDMA	**DPDK**	
Nvidia ConnectX-6 Nvidia Bluefiled-2			Quality of Service (HTB)

[1] Dellsperger, Severin, and Daniel Politze. "Leveraging Intent-Based Networking and SRv6 for Dynamic End-to-End Traffic Steering." NOMS 2025-2025 IEEE Network Operations and Management Symposium. IEEE, 2025.

Caltech and StarLight/NRL Booths at SC24

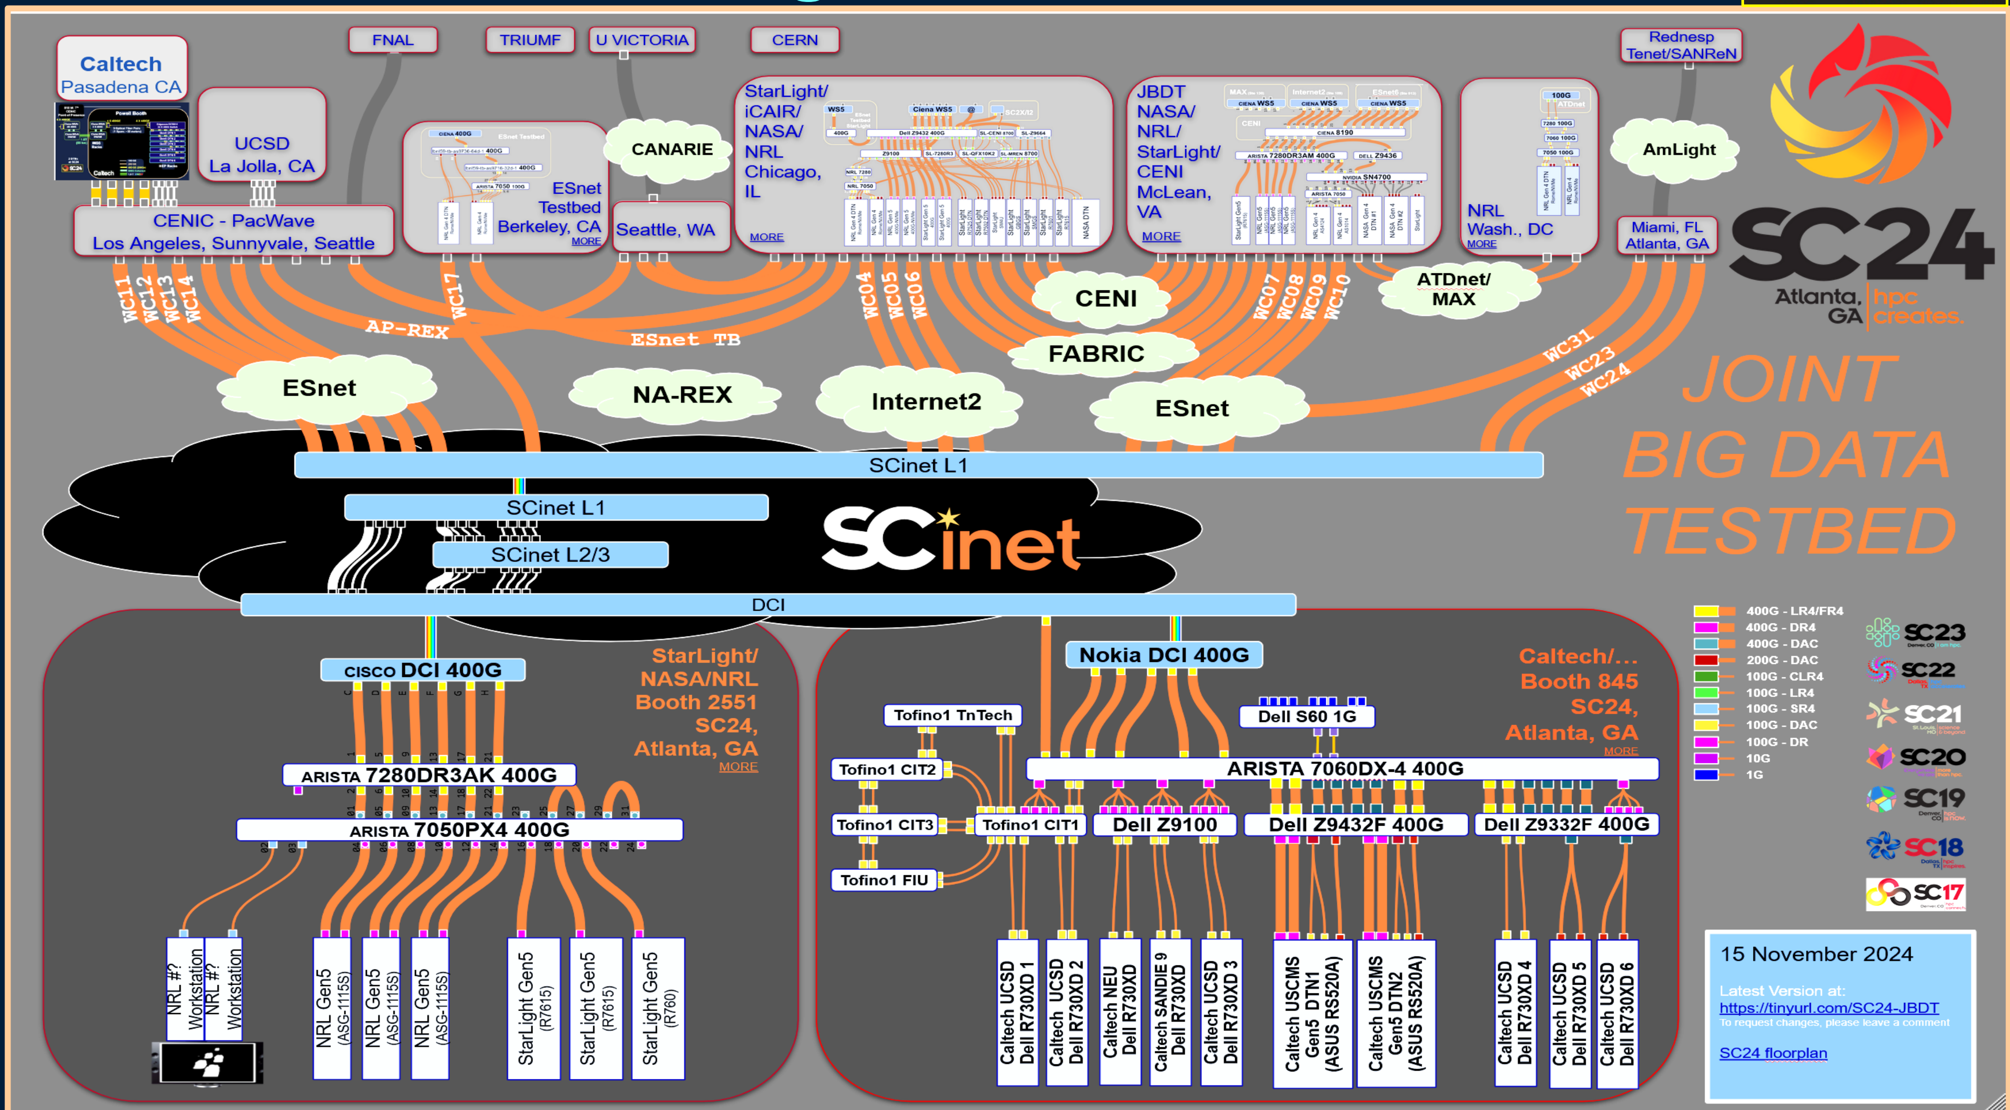

SC24: Global footprint. Terabit/sec Triangle Starlight – McLean – Atlanta;
6 X 400G to the Caltech Booth: 2X 400G LA-ATL; 4 X 400G to the Caltech Campus,
2 X 400G + 100G to Latin America and South Africa; 400G to Fermilab; 400G to CERN
with CENIC, Ciena, Internet2, ESnet, StarLight, US CMS and Network Partners

SC25

DTNs: Supermicro AS21x5-HNTR

PCIe 5.0 Ports: Each Two x16, 1 OCP 3.0 x16

CPUs: EPYC 9375F, 9474F
3.6-3.85 GHz, to 4.3 GHz 224 Cores

Gen5 Servers at SC25: 3 X 1 Tbps

Two ConnectX-7: 2 X 400GE
One Broadcom OCP3.0: 200G
+ 8 Gen3 Servers ~1.6 Tbps

Tofino1 CIT1	42
Tofino1 CIT2	41
Tofino1 FIU	40
Tofino1 TnTech	39
Dell Z9432F 32 X 400G Switch	38
Arista 7050 DX4 32 X 400G Switch	37
	36
Supermicro Gen5 DTN1 2 X 400G + 200G	35, 34
Supermicro Gen5 DTN1 2 X 400G + 200G	33, 32
	31
Supermicro Gen5 DTN1 2 X 400G + 200G	30, 29
	28
Dell 730XD DTN 2 X 100G UCSD 1 (2U)	27, 26
Dell 730XD DTN 2 X 100G UCSD 2 (2U)	25, 24
	23
Dell S60 1G Management Switch	22
Console	21
Dell Z9100 32 X 100G Switch	20
	19
Dell 730XD DTN 2 X 100G UCSD 3 (2U)	18, 17
Dell 730XD DTN 2 X 100G UCSD 4 (2U)	16, 15
	14
Dell 730XD DTN 2 X 100G UCSD5 (2U)	13, 12
Dell 730XD DTN 2 X 100G UCSD6 (2U)	11, 10
	9
Dell 730XD DTN 2 X 100G NEU 1 (2U)	8, 7
Dell 730XD DTN 2 X 100G SANDIE 9 (2U)	6, 5
PDU, cables etc.	4, 3, 2, 1

To ~4.6 Tbps in a single rack

On 11/21/2024 12:47 PM, Harvey Newman wrote:

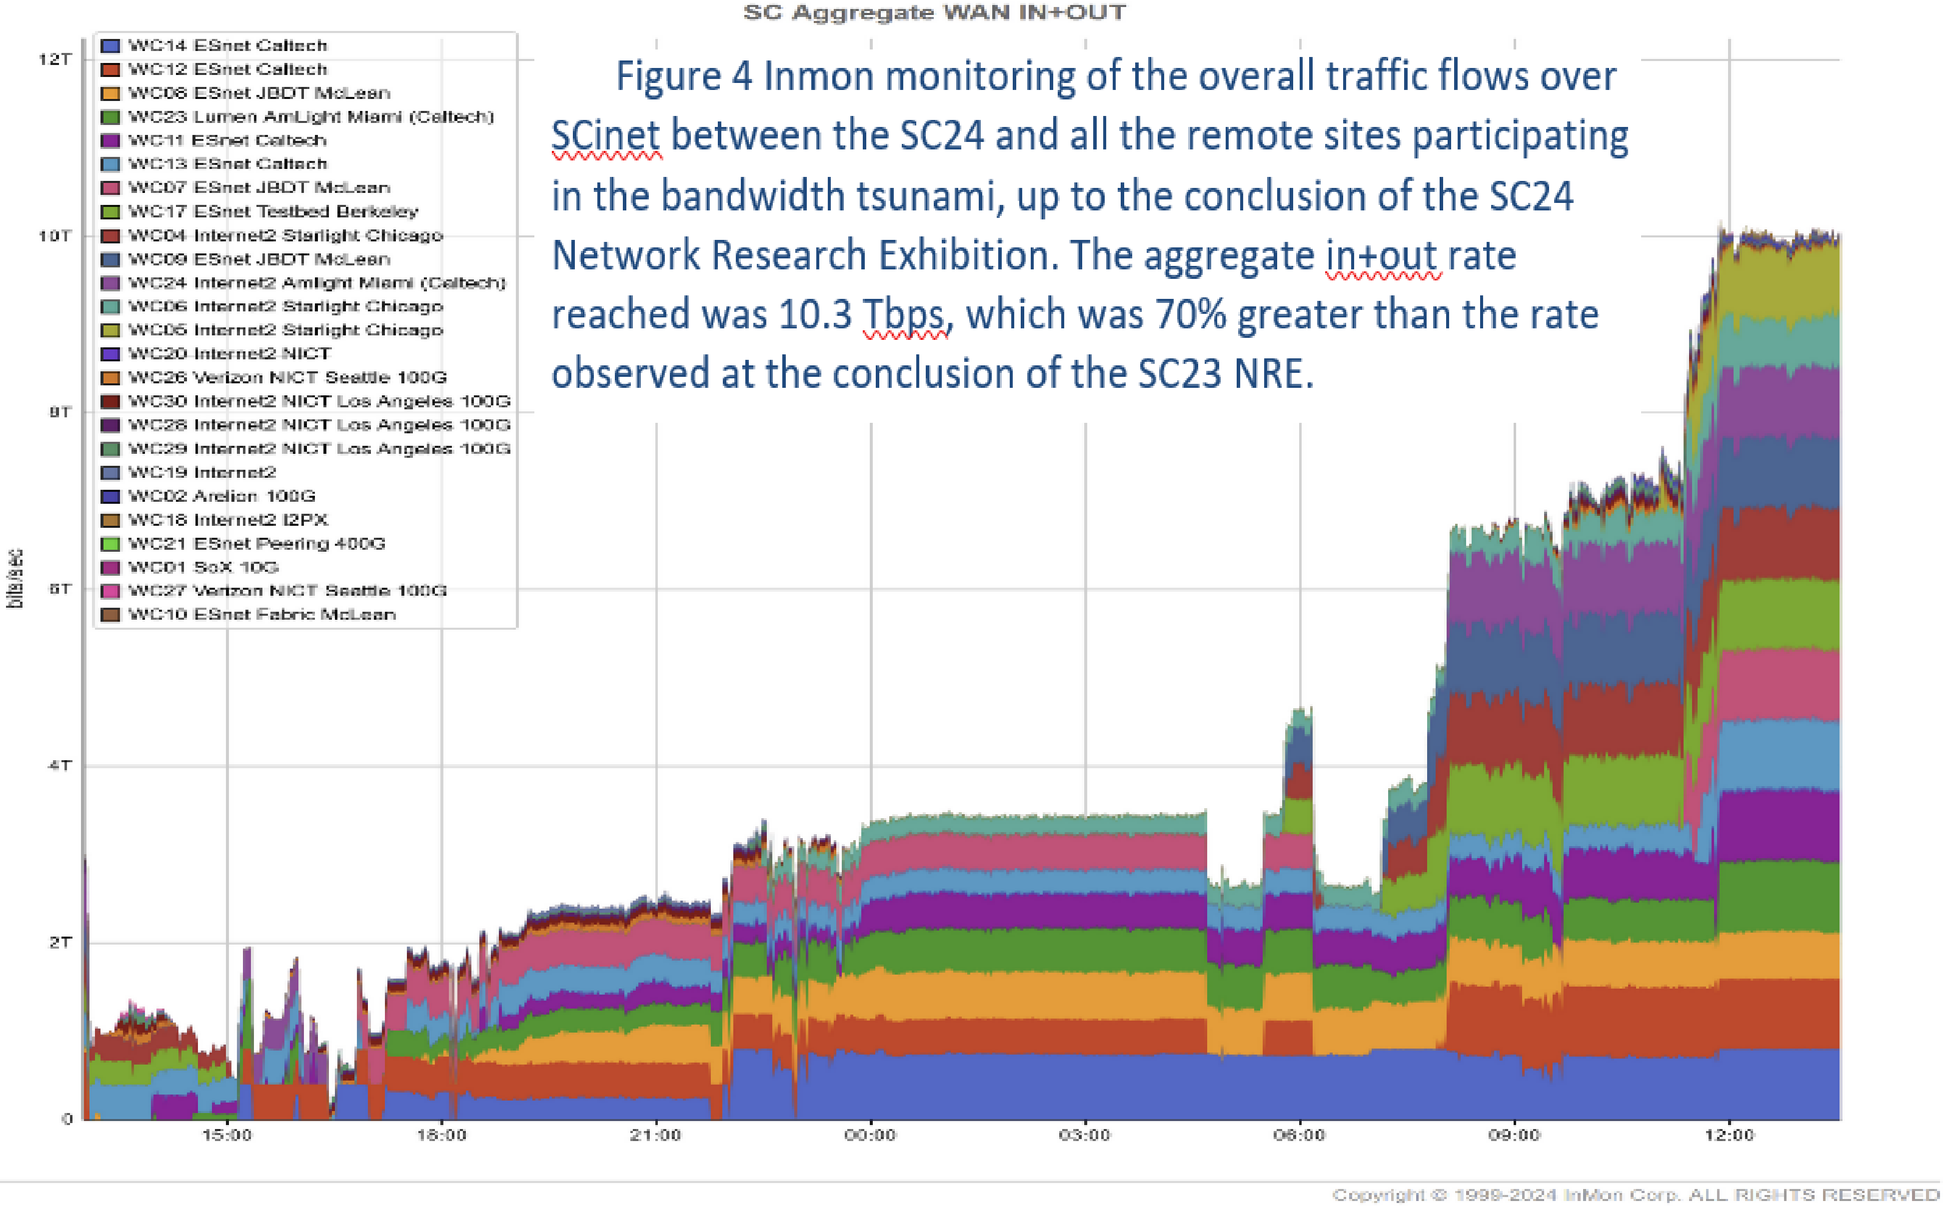

Figure 4 Inmon monitoring of the overall traffic flows over SCinet between the SC24 and all the remote sites participating in the bandwidth tsunami, up to the conclusion of the SC24 Network Research Exhibition. The aggregate in+out rate reached was 10.3 Tbps, which was 70% greater than the rate observed at the conclusion of the SC23 NRE.

A New Generation Persistent 400G Super-DMZ: Ciena, Arista, CENIC, Pacific Wave, ESnet, Internet2, Caltech, UCSD, StarLight++

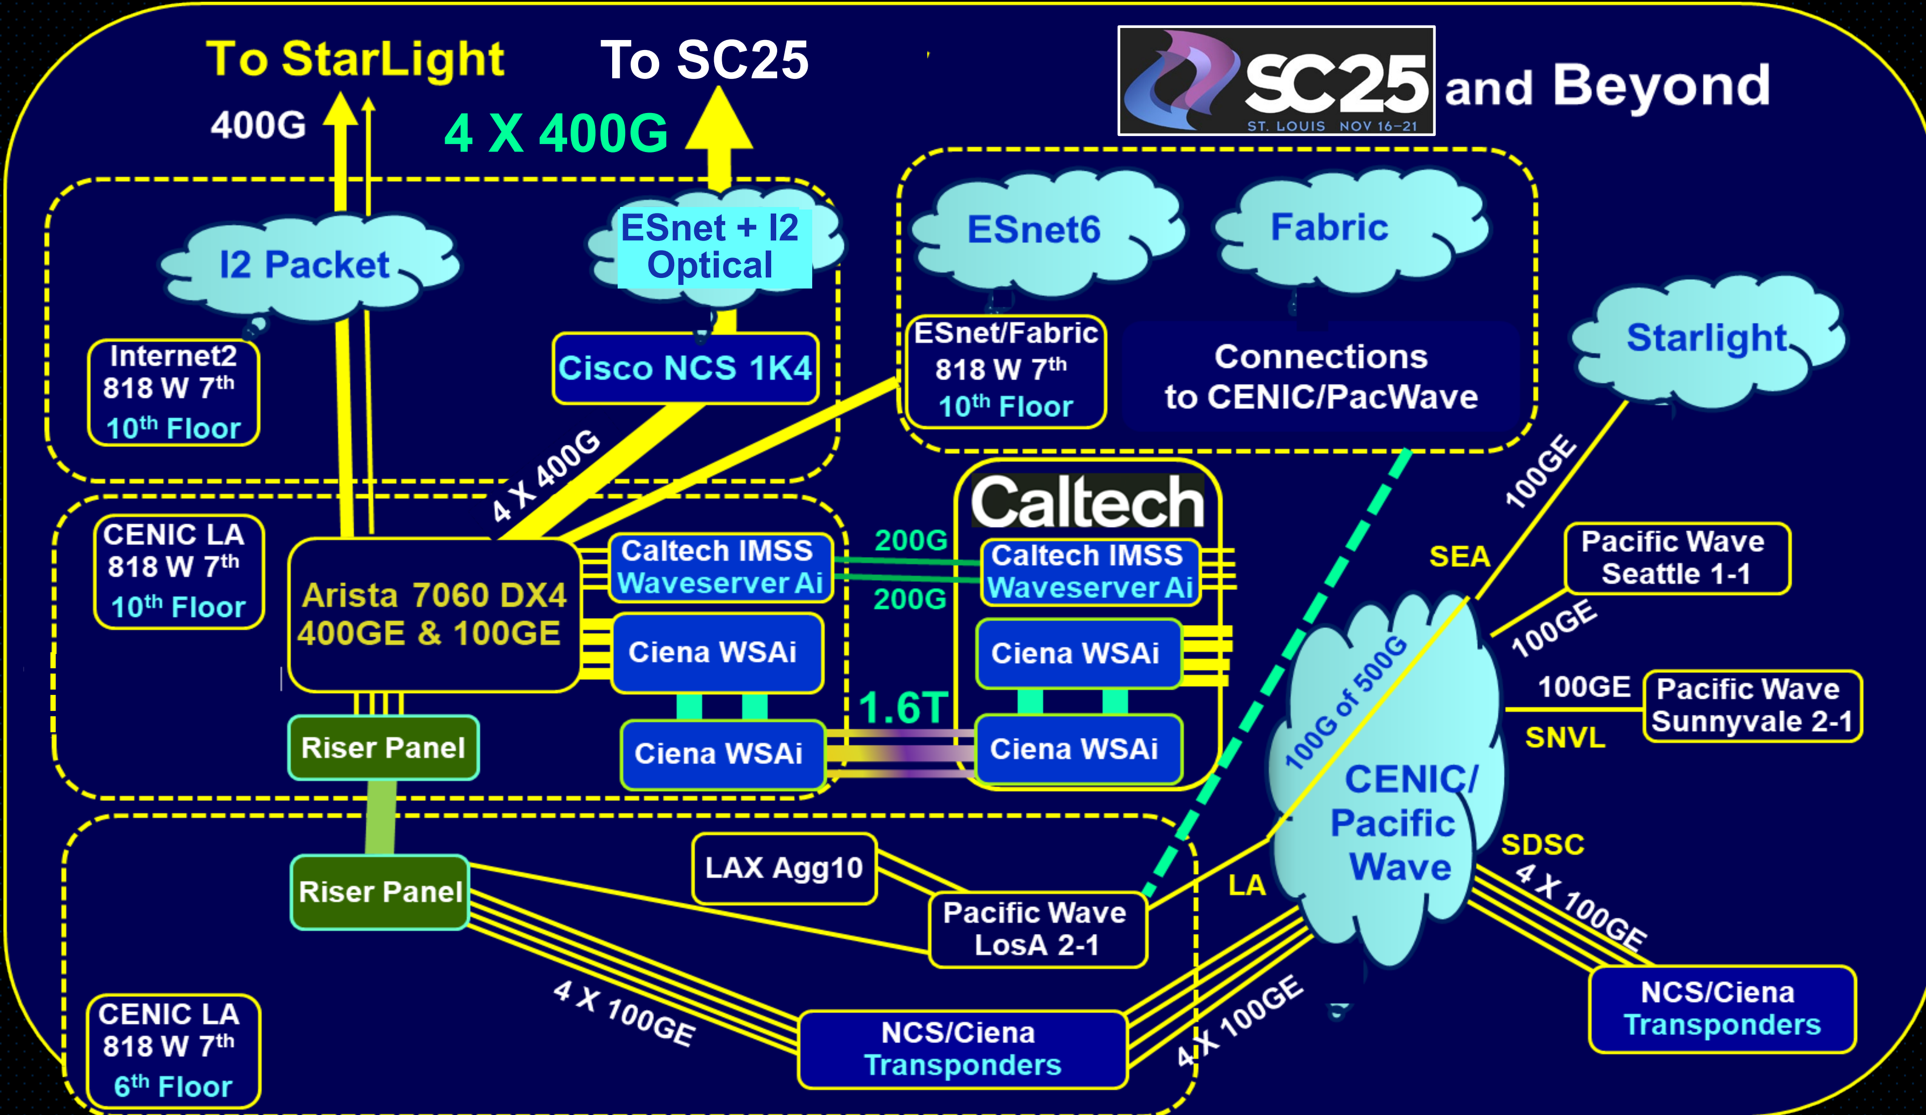

SC24: 4 X 400G on ESnet, I2 Atlanta-LA: Ciena, Caltech and CENIC Using WS Ais and a dark fiber pair. Bringing 4 X 400GE via 2 800G Waves direct to the campus

CENIC, ESnet and Internet2 at the LA PoP

400G + 4 X 100G to Caltech via WS Ais

4 X 400G LA-Atlanta via ESnet, Internet2

4 x 100G to UCSD/SDSC

2 X 400G to Pacific Wave via CENIC

Permanent:
400G NA-REX Prototype
400G to ESnet Production

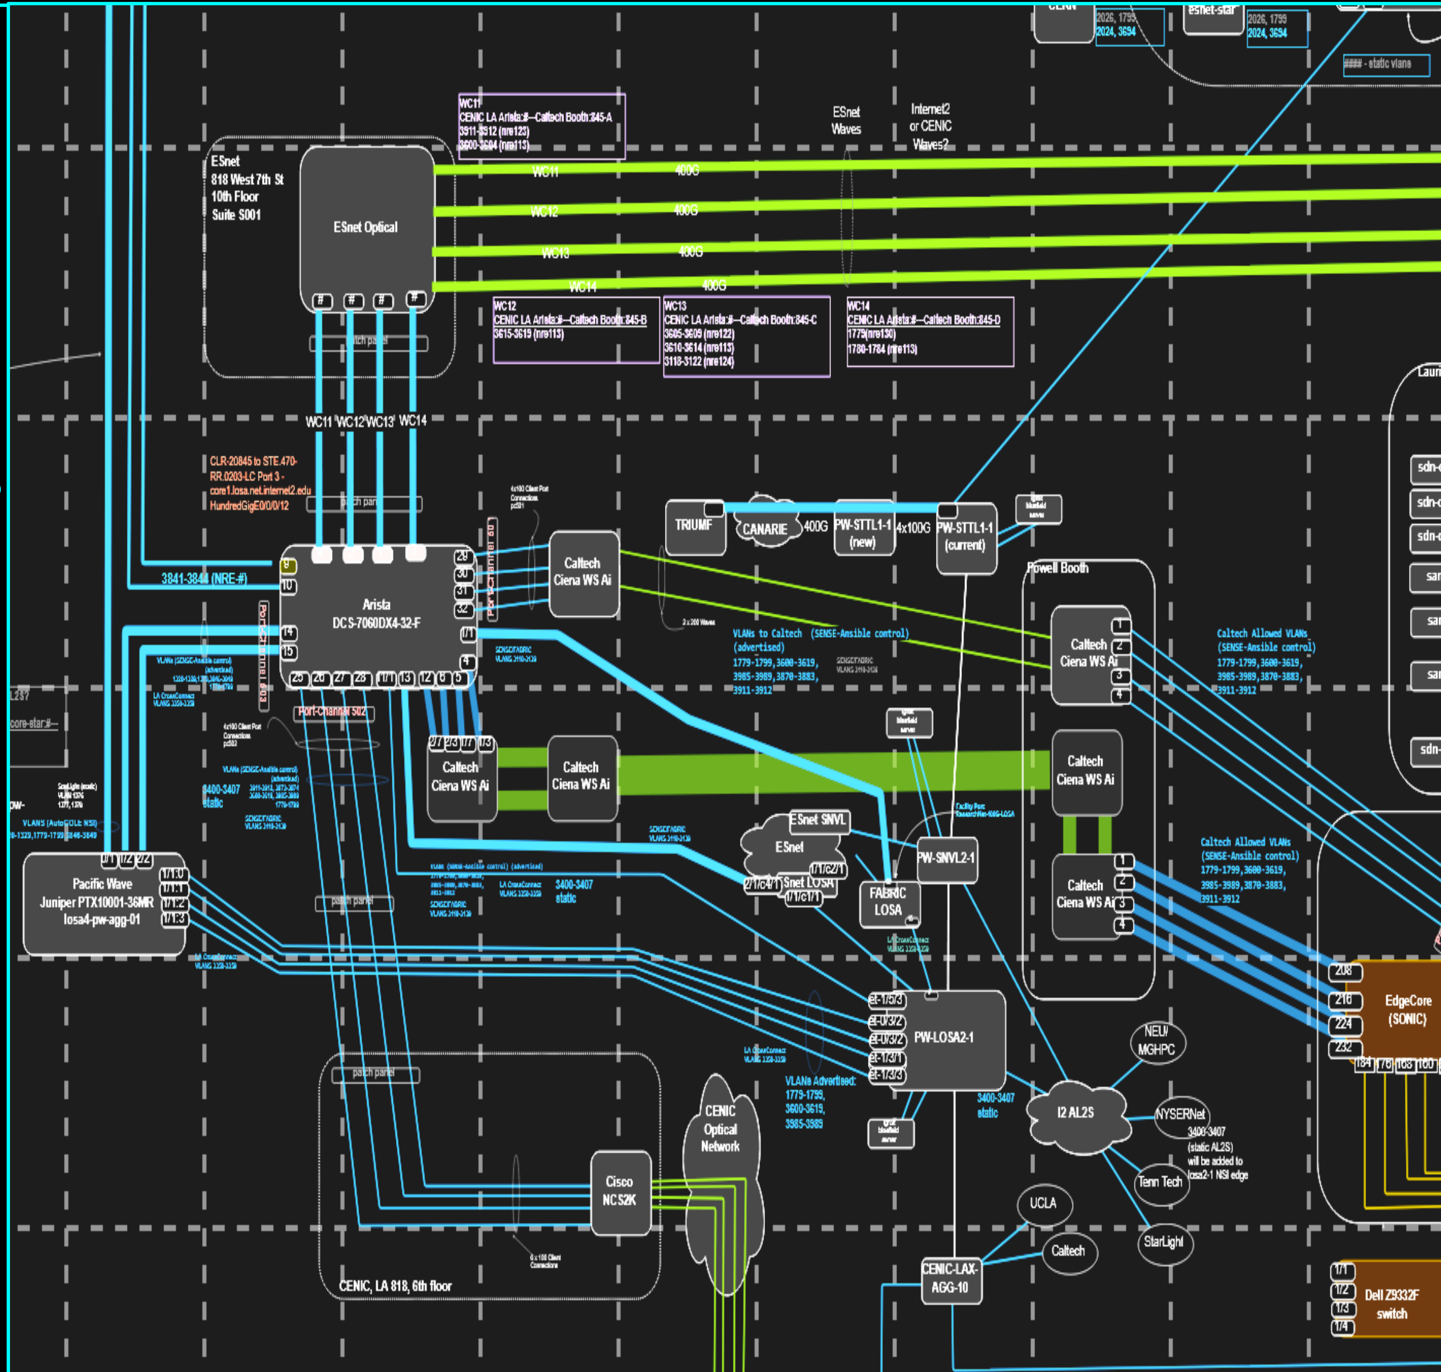

Next Generation Network-Integrated System

- **Top Line Message:** To meet the challenges & realize the discovery potential of data intensive science programs, we are developing a system which:
  - Coordinates worldwide networks as a first class resource along with computing and storage, across multiple world regions
  - Simultaneously supports major Data Intensive Science programs and the worldwide academic and research community
- Follows a Systems Approach: A global dynamic fabric that flexibly allocates, balances and conserves the available network resources
  - Negotiating with site systems that aim to accelerate workflow
- We are leveraging the worldwide move towards a fully programmable ecosystem of networks and end-systems (P4, PINS; SRv6; PolKA), plus operations platforms (NRP, OSG; global SENSE Testbed, GP4L, FABRIC)
- We are building on ongoing R&D projects: from regional caches/data lakes to intelligent control and data planes to AI/ML-based optimization [E.g. SENSE/AutoGOLE, NOTED, ESNet HT, GEANT/RARE, AmLight, FABRIC, GP4L; Qualcomm GradientGraph, Kytos, NetPredict, Hecate, DeepRoute, PolKA, ALTO ...]
- Key milestones: integration of SENSE + network services with: Rucio/FTS/XRootD/dCache for LHC; Globus Data Movement and Management; EJFAT (ESnet/JLaB FPGA Accelerated Transport); Across agencies with the DOE IRI
- The GNA-G and its Working Groups, the GRP, AmRP, KRP, and R&E network teams are key partners: developing the vision, empowering the worldwide communities

Acknowledgements

This ongoing work has been partially supported by the US National Science Foundation (NSF) Grants OAC-2030508, OAC1841530, OAC-1836650, MPS-1148698, and PHY-1624356, along with research grants from many international funding agencies and direct support from the many regional, national, and continental network and industry partners mentioned. The development of SENSE is supported by the US Department of Energy (DOE) Grants DE-SC0015527, DESC0015528, DE-SC0016585, and FP-00002494.

Finally, this work would not be possible without the significant contributions and the collaboration of the many HEP, network and computer and research teams partnering in the Global Network Advancement Group, in particular the GNA-G Data Intensive Sciences and AutoGOLE/SENSE Working Groups, CENIC, Internet2, ESnet, STARLIGHT, NRP, Rednesp, USP, the GP4L and the RNP Brazilian National Network, together with the LHC experiments and many industry partners, most notably Ciena, Arista and Dell.