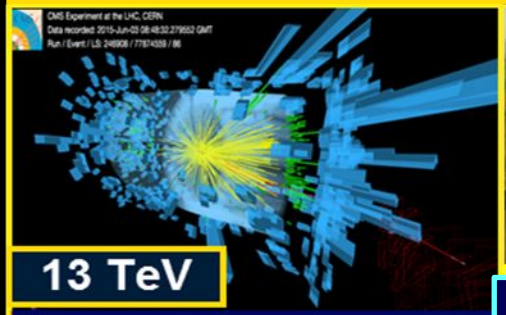
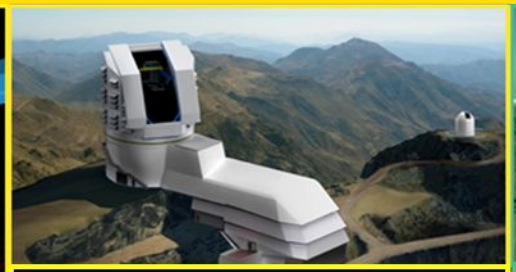
Global Network Advancement Group

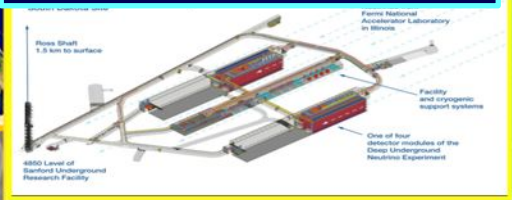
A Next Generation Multi-Tbps Campus and Global Network System for Data Intensive Sciences

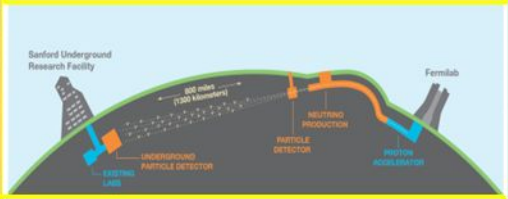


LHC

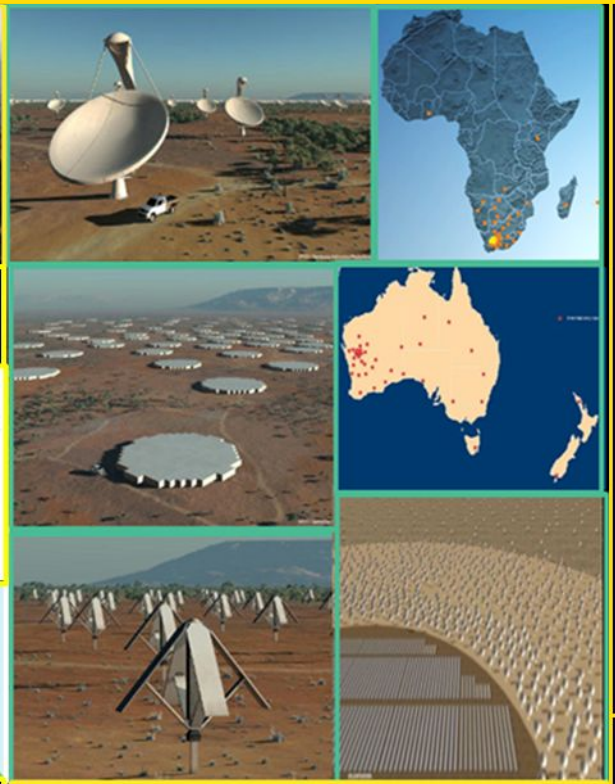


Rubin Observatory LSST





LBNF/DUNE



SKA

LHC Run 3 and HL-LHC

Rubin Observatory

SKA

BioInformatics

Earth Observation

Gateways to a New Era

SC26: NRE-118 and 15 Partner NREs at Booth 5438

INDIS 13th Workshop November 15, 2026. On behalf of the GNA-G and its Data Intensive Sciences and AutoGOLE/SENSE Working Groups

Full Slide Set Here

LHC: Discovery of the Higgs Boson and Beyond; 75 Years of Exploration !



Physicists Find Elusive Particle Seen as Key to Universe



2013 Nobel Prize

Englert	Higgs



	Energy Frontier	Intensity Frontier	Cosmic Frontier
Higgs Boson	●		
Neutrino Mass		●	●
Dark Matter	●	●	●
Cosmic Acceleration			●
✶ Explore the Unknown	●	●	●

48 Year Search; 78 Year Exploration

Theory (1964): 1950s – 1970s;

LHC + Experiments Concept: 1984

Construction: 2001; Operation: 2009

Run1: Higgs Boson Discovery 2012

Runs2,3 and Going Forward: Precision Measurements and BSM Exploration: 2013 - 2041



Advanced Networks Were Essential to Higgs Discovery and Every Ph.D Thesis; They will be Essential to All Future Discoveries

- **NOTE: ~85% of Data Still to be Taken**
- **Greater Intensity: Upgraded detectors for more complex events**
- ***To 10X Data Taking Rate at HL LHC***

Global Data Flow: LHC Grid Hierarchy
A Worldwide System Invented by Caltech (1999)



Experiment

~PBytes/sec

Online System

~1000 MBytes/sec

13 Tier1, *170 Tier2*
+ 300 Tier3 Centers:
1k to N X 10k Cores

Tier 0

CERN Center
100 PBs of Disk;
Tape Robot

Tier 1

N X100 Gbps

London

Paris

Taipei

Chicago

10 to 100 Gbps

Tier 2

CACR

Center

Center

Center

Tier 3

Institute

Physics data cache

Workstations

10 to 100 Gbps

Tier 4

Tier0: Real Time Data Processing
Tier1 and Tier2: Data and Simulation Production Processing
Tier2 and Tier3: Data Analysis

A Global Dynamic System
Increased "Elastic" Use of Additional HPC and Cloud Resources
Fertile Ground for Control with AI/ML

WLCG Transfers Dashboard: Throughput Sept. 2015 – August 2025

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

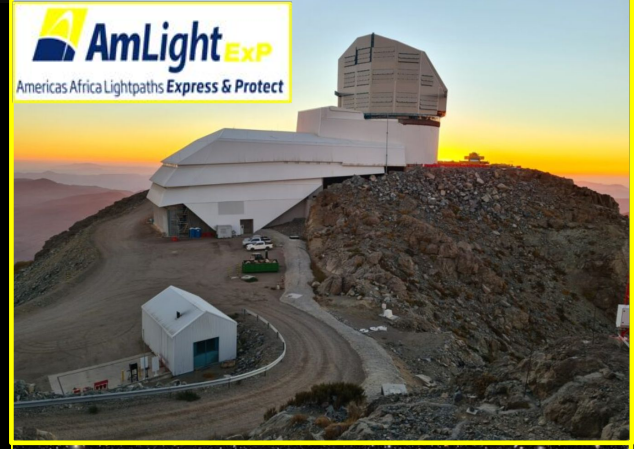
- ***Another Sea Change in 2024-5***
- **1.8X in Transfer Rates in 12 months**
- **DC24 (25% HL LHC): to 300 GB/s ⇨ ~10 Tbps by the HL LHC era**

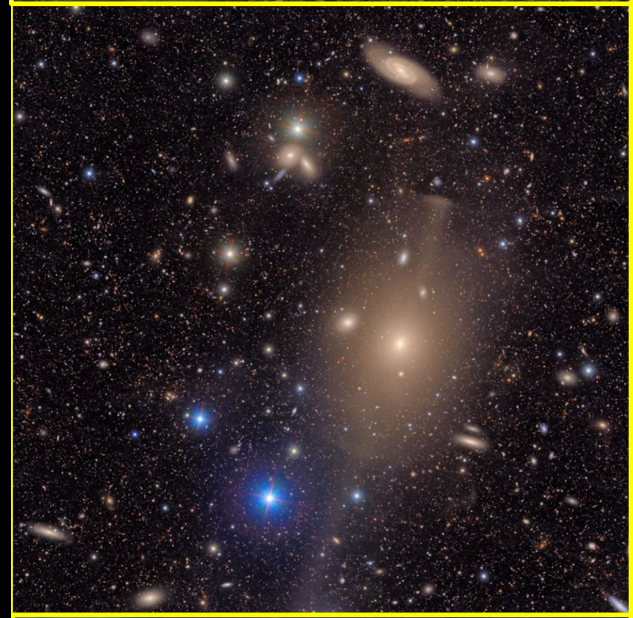
~17X Growth in Throughput 2015-2025: +33%/Yr; + Much Faster Growth Bursts

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

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

- **Rubin Observatory will take hundreds of 3.2 Gigapixel images of the Southern Hemisphere sky, every night for ten years:** 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 survey: 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 !***
- **Dataset by the end of the survey will be enormous —about 60 petabytes!** Most of the astronomers who make discoveries will access the data using an online portal: **the Rubin Science Platform**
- **No need for 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 !**

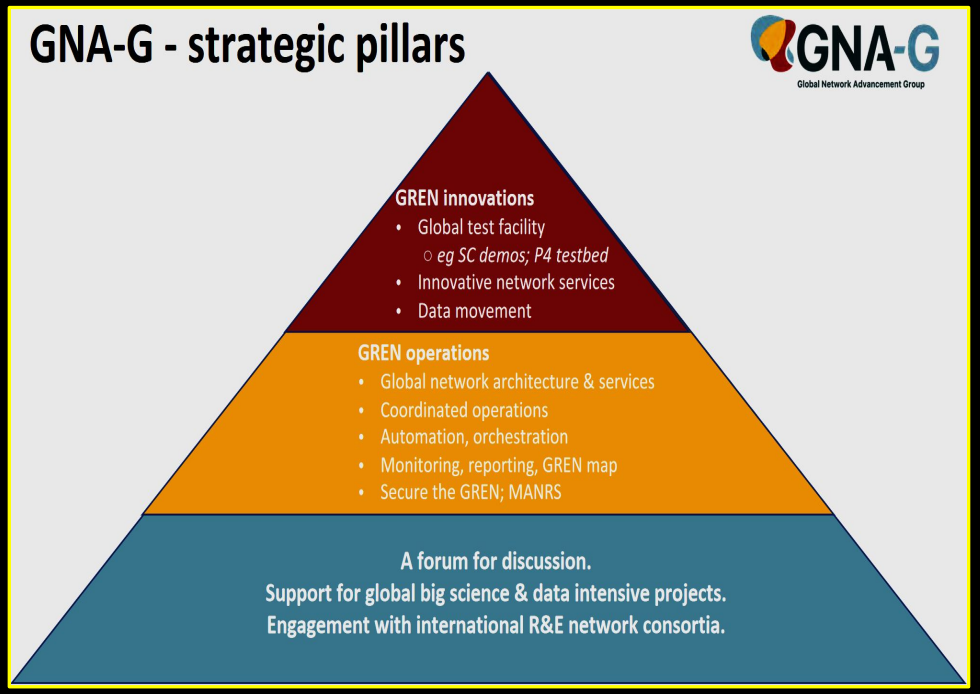






Global Network Advancement Group (GNA-G)
Since September 2019: In Our 7th Year

- **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**

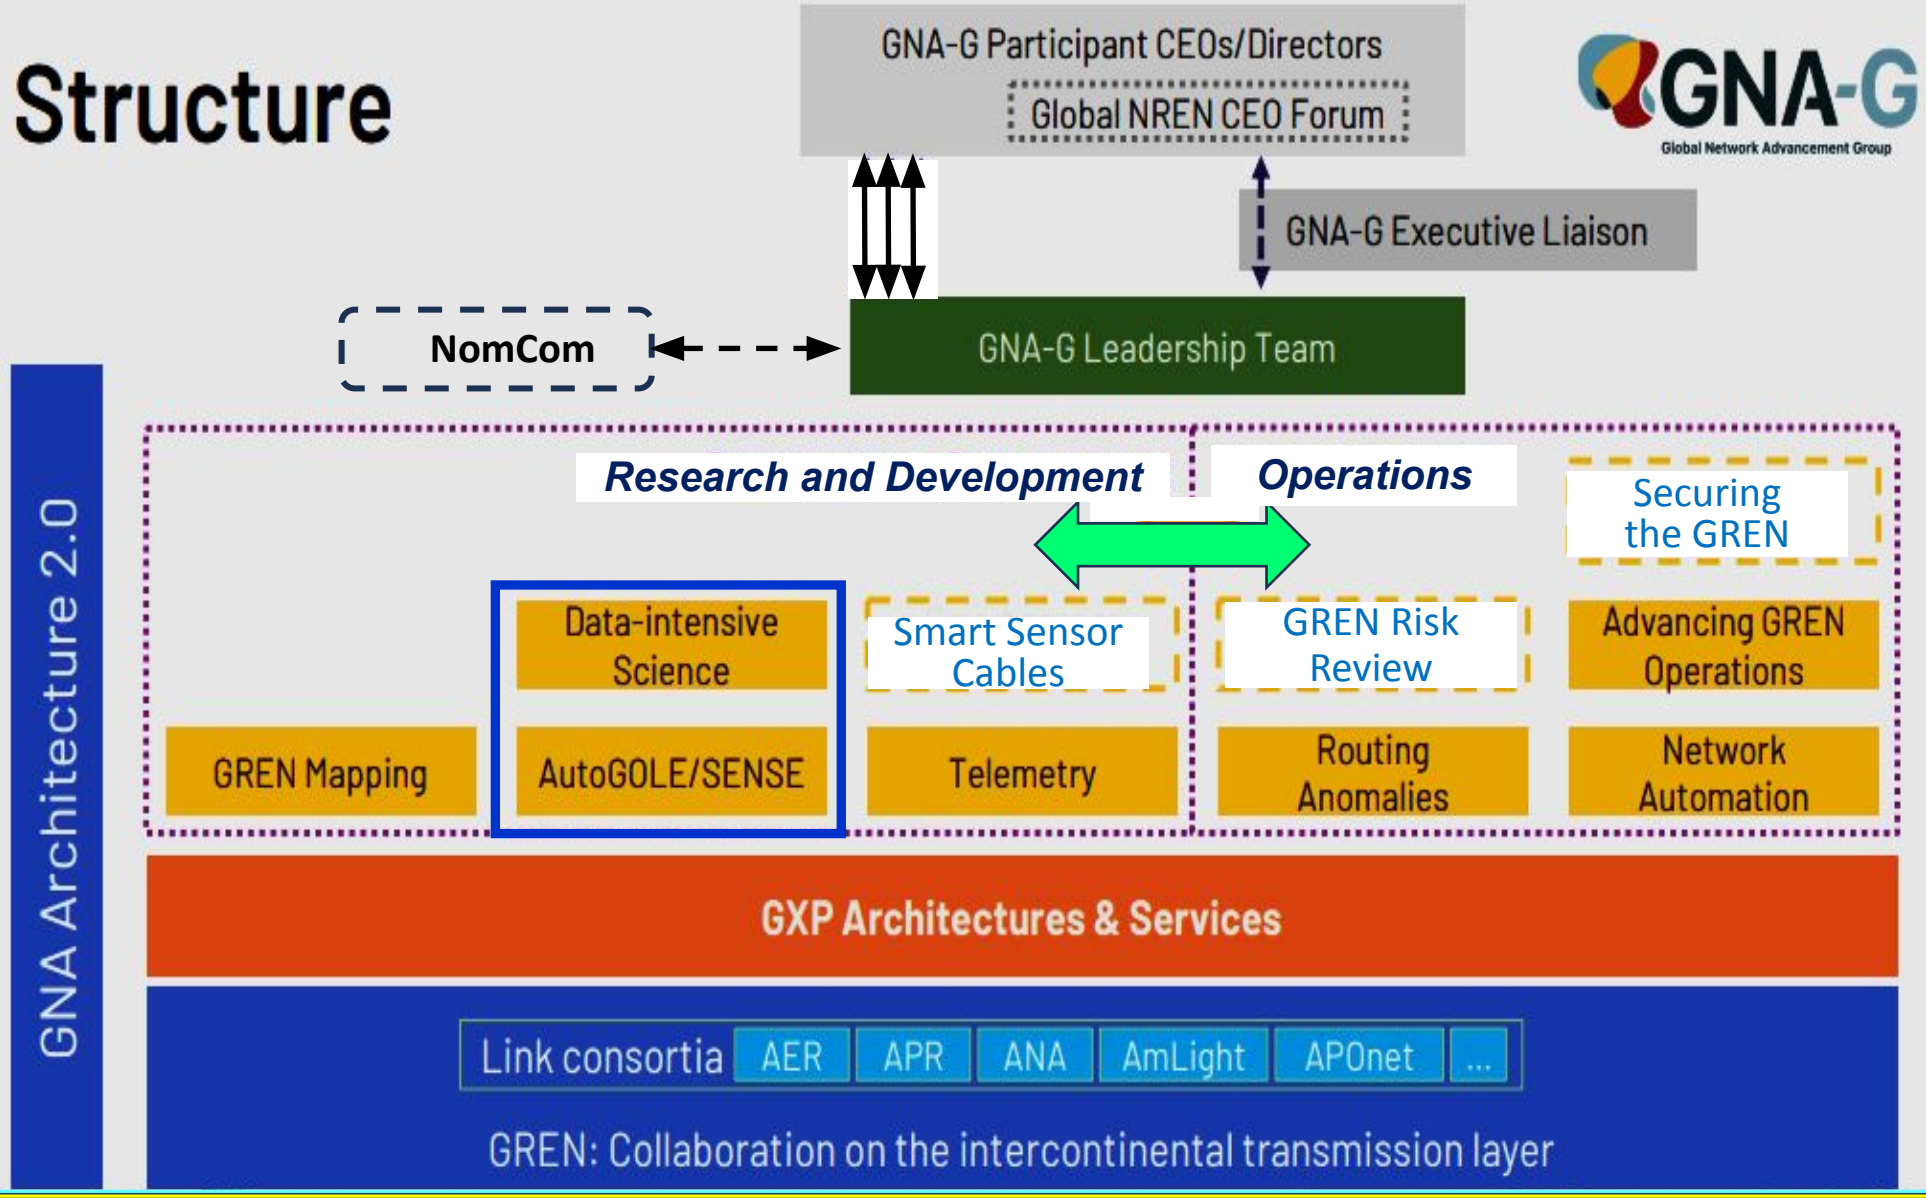
See https://www.dropbox.com/s/qsh2vn00f6n247a/GNA-G%20Meeting%20slides%20-%20TechEX19%20v0.8.pptx?dl=0

The Global Network Advancement Group

Progress and Ongoing Impact 2019-2026

- **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 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***

Structure



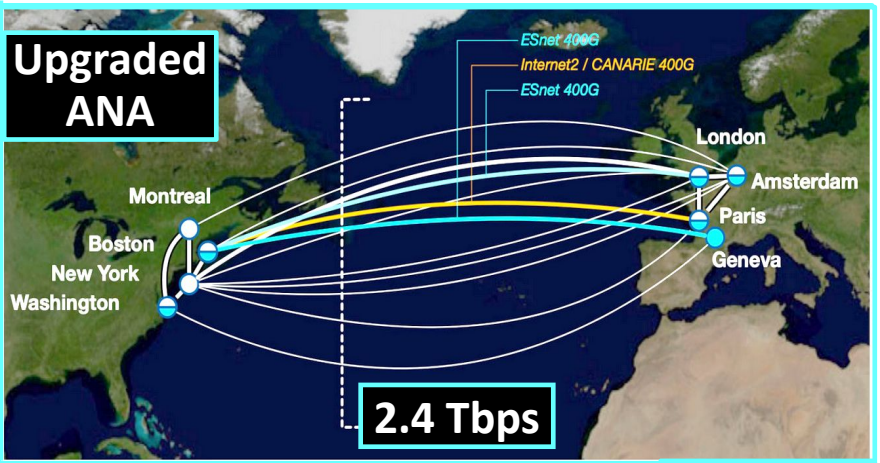
Theme: from Transformative Concepts to Advanced Services to Production Networking
New Directions: Scaling to Tbps; Digital Twins; Generative AI for Global Networks

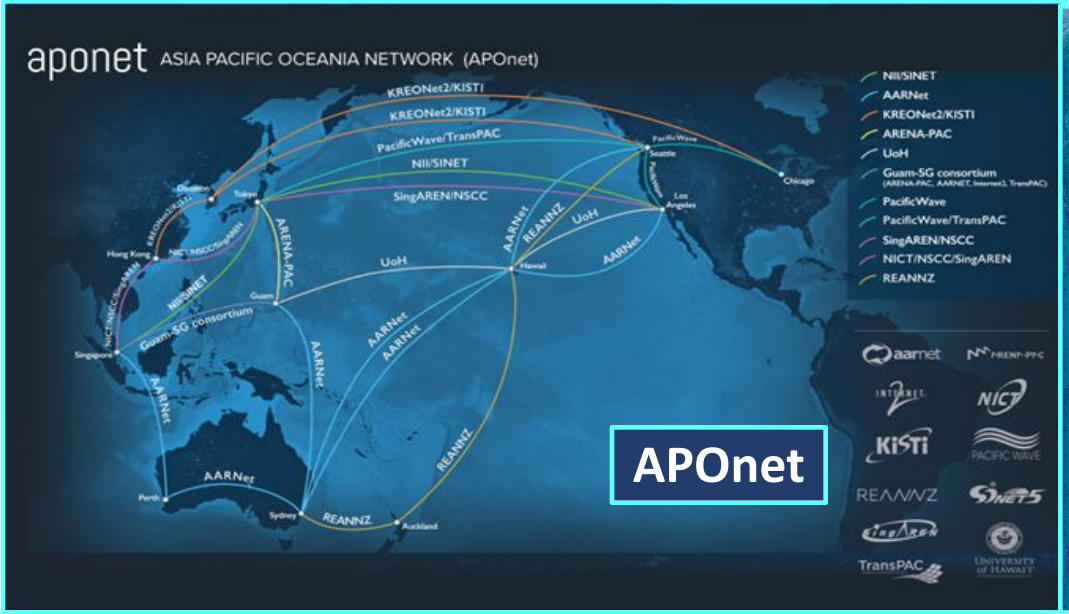
GNA-G
Global Network Advancement Group

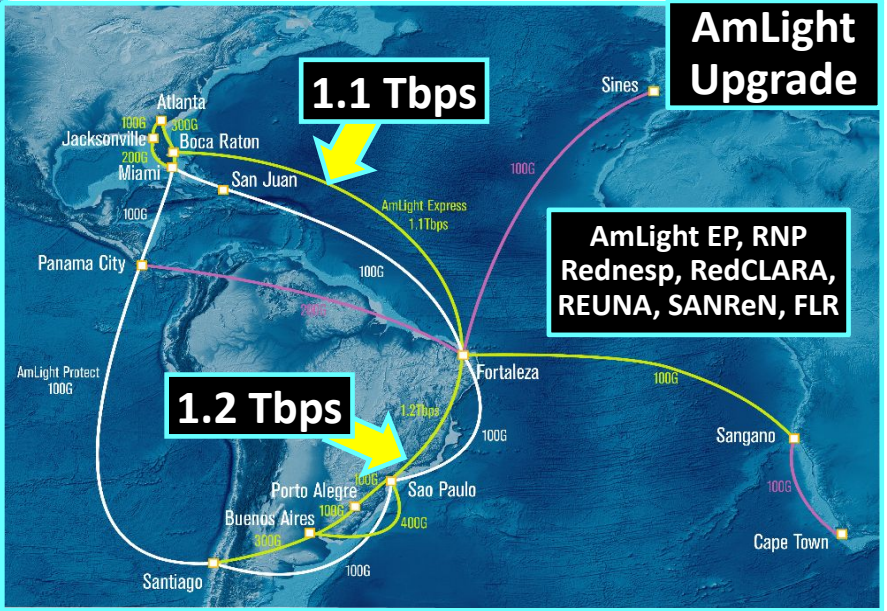
https://www.gna-g.net/

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



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 Meet the challenges faced by leading-edge data intensive programs in HEP, astrophysics, genomics, 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
- Developing a wide range of the 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 concepts and technologies in programmable networks and global-scale distributed systems
- Hallmarks: Progressive multidomain integration; compatibility internal + external; *A comprehensive systems-level approach*



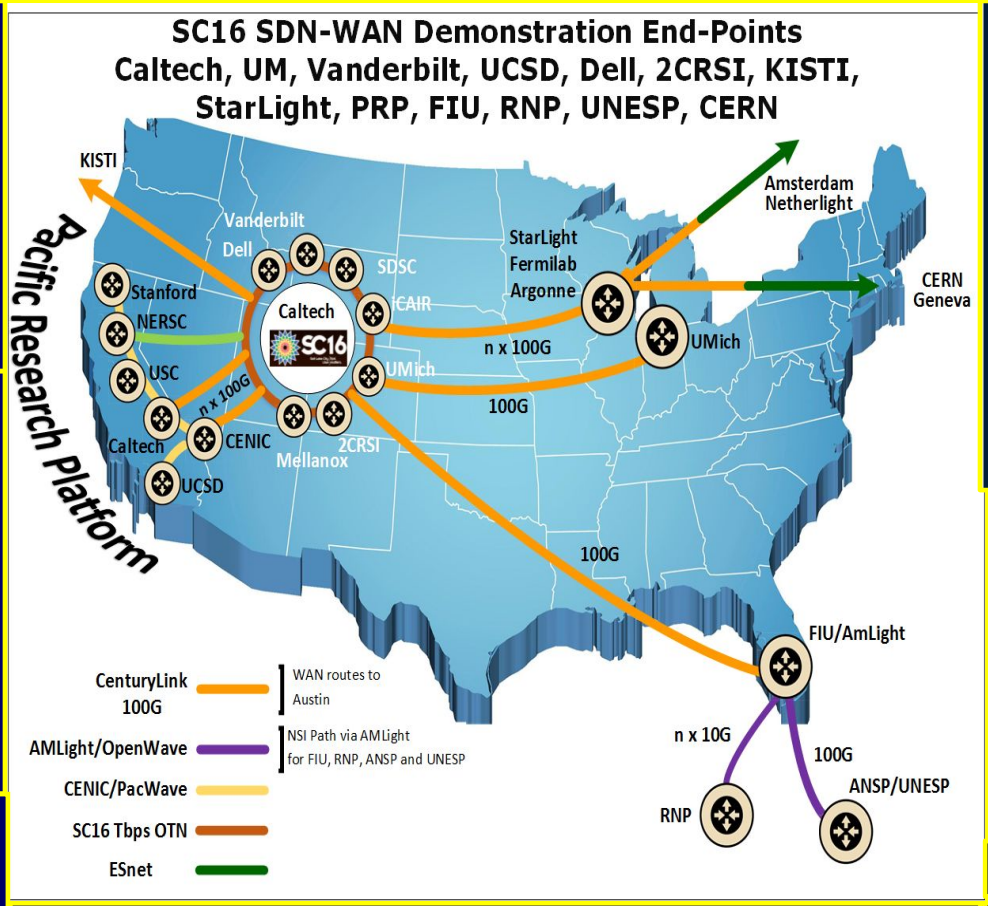
SC15-26: SDN Next Generation Terabit/sec Ecosystem for Exascale Science

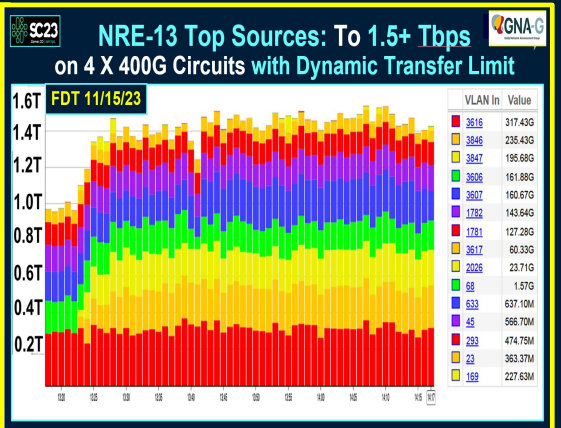
supercomputing.caltech.edu

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

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

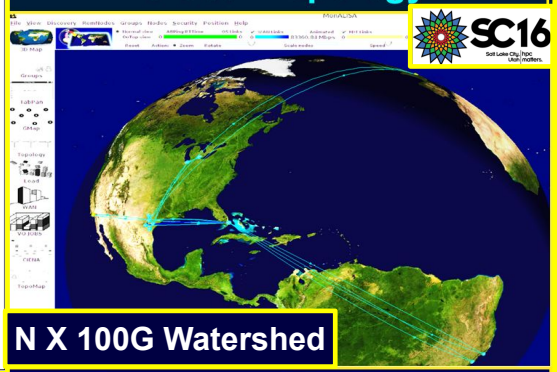
Preview PetaByte Transfers to/ from Sites With 100G - 1000G DTNs





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

Global Topology



N X 100G Watershed

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

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

Tbps Rings for SC18-26: Caltech, Ciena, SCinet, StarLight + Many HEP, Network, Vendor Partners

NRE-118: 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 flexible *high water marks* to accommodate other traffic
- Negotiating with site demands from 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

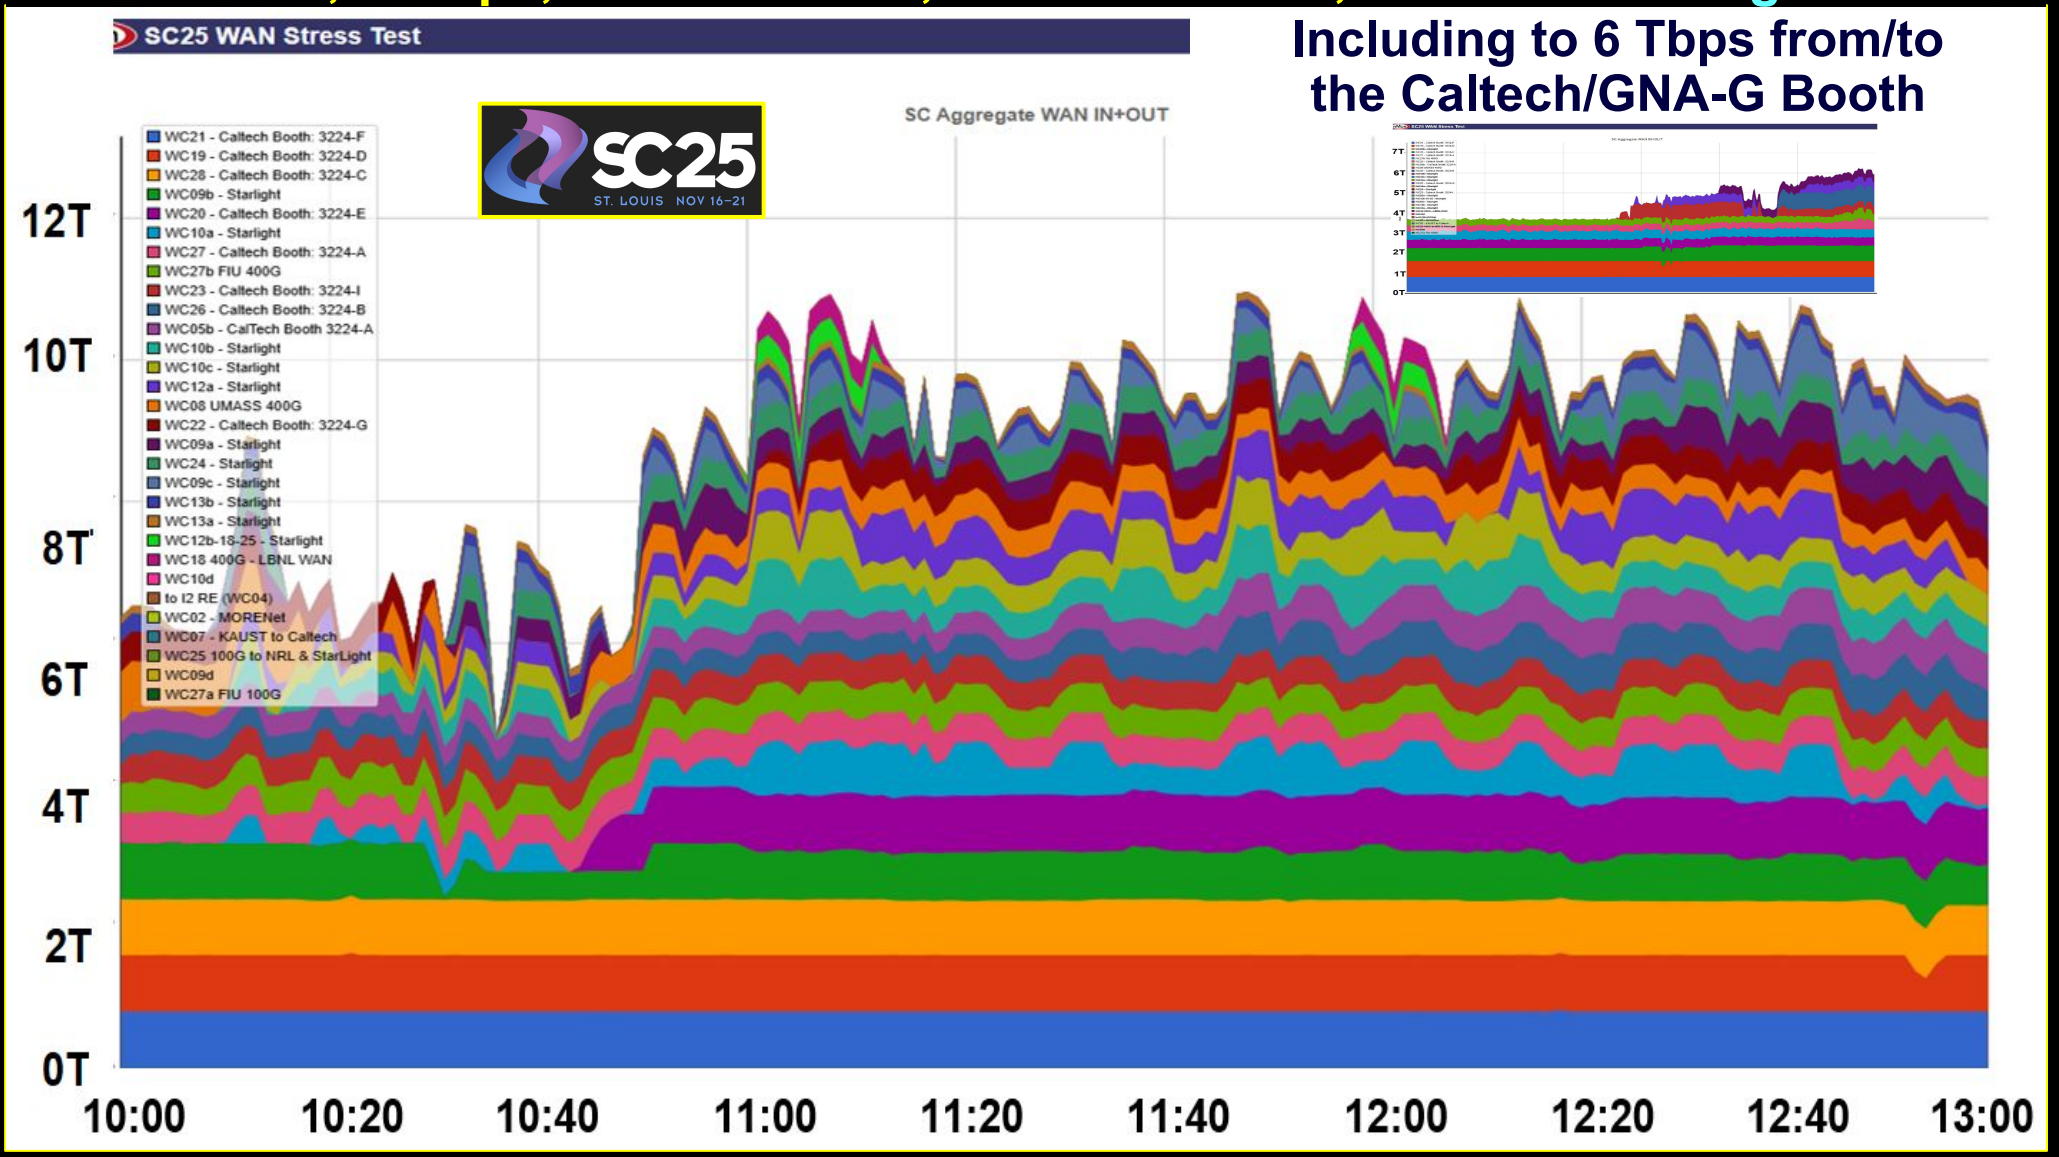
SC26 Network Research Exhibition NRE-118 and Partners
15 NREs Hosted at or Partnering with Caltech Booth 5834

NRE-118	Harvey Newman et al	The Global Network Advancement Group: A Next Generation System for Data Intensive Sciences
NRE-106	Mohammad Sada et al	Agentic WAN-Native LLM Orchestration: From User- Driven Inference to Geographically Distributed Model Instances
NRE-107	Mohammad Sada et al	Bringing Together a Kubernetes-Native Network Operator and P4 Programmable Testbed
NRE-108	Mohammad Sada et al	The National Research Platform & SCinet: Live, MultiDomain AI/HPC Transnational Workflows
NRE-112	Rafael Coehlo Lopez de Sa et al	Reimagining advanced capabilities in high-bandwidth research networks
NRE-115	Carlyn Lee et al	Federated Edge AI for Distributed Cosmic Ray Sensor Networks
NRE-119	Vasilka Chergarova et al	Network-Enabled Transcontinental CryoEM Over AmLight and FABRIC
NRE-122	Tom Lehman et al	Customized End-to-End Services for Science Workflows via SENSE
NRE-123	Tom Lehman et al	FABRIC - Programmable Cyberinfrastructure with External Resource Integration
NRE-128	Marcos Schwarz et al	GP4L: Global Platform for Labs
NRE-129	Carlos Ruggiero et al	High Performance Networking with the São Paulo Backbone SP Linking 10 Academic Institutions
NRE-130	Carlos Ruggiero et al	High Performance Networking with the Backbone SP High-Throughput RDMA across Intercontinental Circuits
NRE-139	Magnos Martinello et al	High-Performance Scientific Data Paths with PolKA and Programmable DTNs
NRE-142	Mariam Kiran et al	SciNet: Science-Aware Network Operations for Research Networks
NRE-143	Alex Szalay et al	High Performance Networking with IDIES: Data Intensive Engineering and Science Enabling Discoveries

Data Tsunami: Unprecedented Throughput Levels

Sustained Flows of 10 Tbps with 11 Tbps Peaks

Including flows across the Atlantic and Pacific, between Saint Louis & sites in the US, Europe, Latin America, Asia and Africa, monitored using Inmon

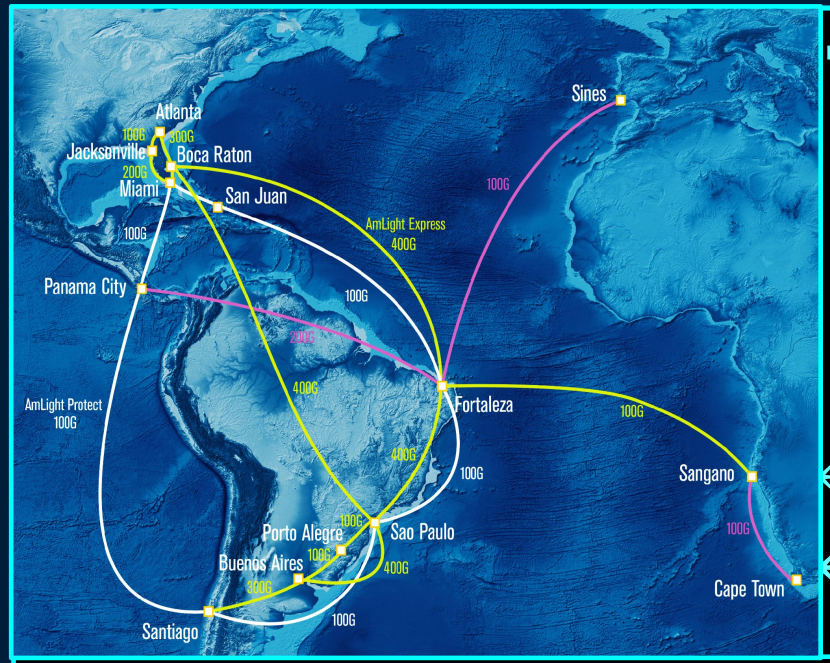


Rednesp, RNP and AmLight: Expanding Capacity Among Latin America, US, Europe and Africa

Total 1.7 Tbps capacity between Brazil and the USA:

- **Rednesp 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; then to San Juan & Miami
- ✳ **AmLight Link (with Rednesp and RNP):** São Paulo to Florida through Fortaleza, on the Monet cable: **Upgraded from 400G to 1.1 Tbps by October 2025**

100G Links to Europe: Bella link between São Paulo and Sines in Portugal, **and to South Africa:** São Paulo to Sangano to Capetown on Angola Cable



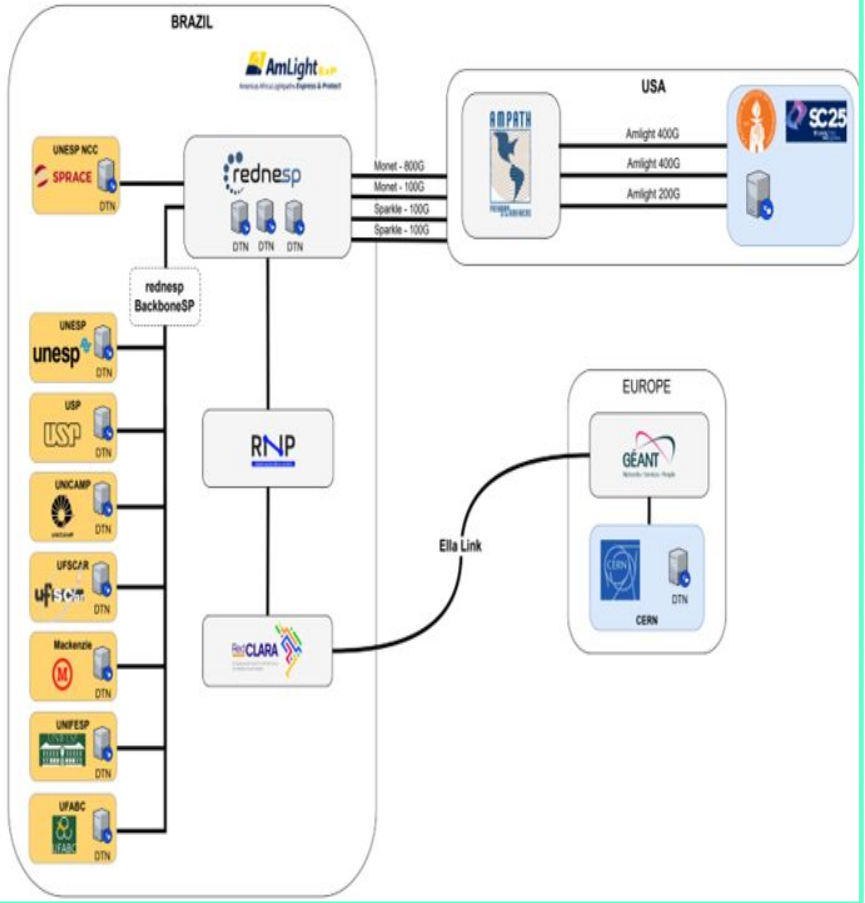
- "Backbone SP" connecting most universities in São Paulo State with 100G links was completed in May 2023
  - UNICAMP (University of Campinas),
  - UNIFESP (Federal University at São Paulo)
  - UNESP (State University of São Paulo)
  - USP (University of São Paulo)
  - UFSCAR (Federal University of São 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 being added by Rednesp** from its main Equinix PoP to USP, UNESP and UNICAMP

✳ **USP dark fiber infrastructure completed in 2025**

✳ **Four main USP campuses connected at 400G-800G**

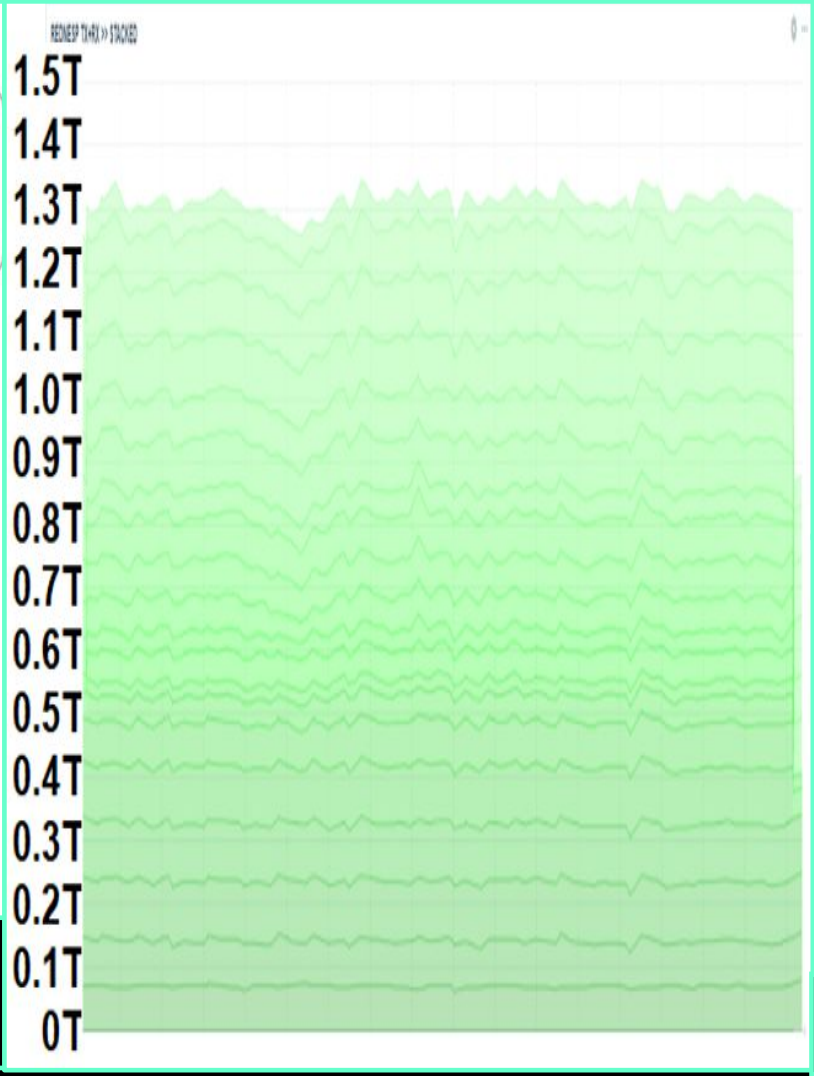
Data Tsunami: Wide Area 1+ Tbps Flows over SCinet, AmLight, Rednesp between the Caltech booth and 8 Sao Paulo Universities

Simplified Setup for Rednesp at SC25

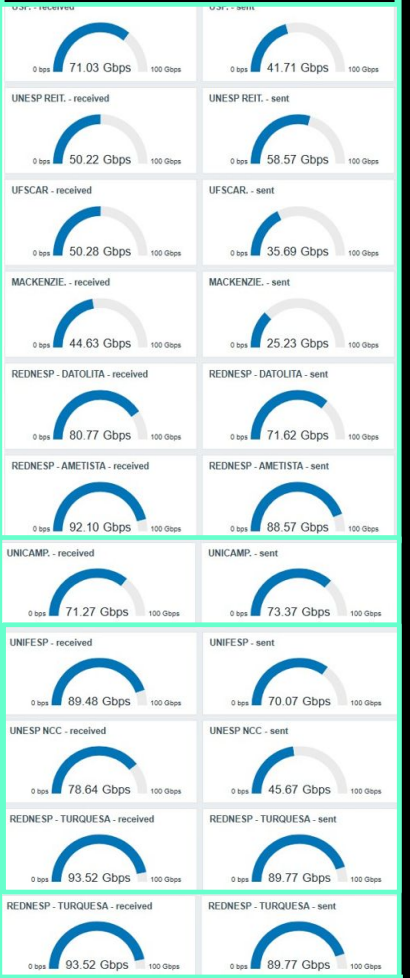


Twice the SC25 Target Achieved; RDMA Over 146 msec RTT

Max Throughput >1.3 Tbps Using RDMA (ROCEv2)



To 92 Gbps Individually



80G Cap at Production Sites

NRE-142 SciNet: Science-Aware Network Operations for Research Networks

ORNL/Caltech/HPE/UCR: Science Intent Becomes Verified Network Action

Issue

While AI-driven autonomous laboratories are emerging, existing research networks remain passive communication infrastructures that work to optimize packet delivery without understanding the scientific workflows they support. This disconnect creates bottlenecks in distributed experimentation, multi-facility workflows, and scientific decision making.

New Paradigm

Science-Aware Autonomous Networks (SciNet) is a new networking paradigm where artificial intelligence (AI) is integrated into the network control plane to continuously observe scientific workflows, predict future network demands, reason about workflow dependencies, and autonomously optimize network behavior.

Solution Path

SciNet combines multimodal data fusion, knowledge graphs, graph neural networks, causal inference, foundation models, digital twins, and multiagent reinforcement learning into a unified agentic architecture. This transforms research networks into active participants in scientific discovery, optimizing scientific productivity through workflow-aware routing, predictive resource allocation, and autonomous mitigation of failures.

Goals and Impact

Rather than optimizing for packets or flows, SciNet optimizes for scientific experiment outcomes: short experimental latency, fewer workflow failures, faster discovery cycles across multiple science domains. This is a foundational capability for AI-driven autonomous laboratories and the agentic workflows that drive them: A science-aware network measured by the SLAs it maintains and the experiments it accelerates rather than by the bits it moves.

SciNet: Science Intent Becomes Verified Network Action

SciNet-Extensions

JSON/YAML Infrastructure/ Hardware APIS

SCiNet Monitoring and Evaluation Layer

Science Workflow Characterization

Verify Intent, Network Policies, Service Level Agreements for Science

Benchmarked Science Workflows

SciNet Flow Engine

Task and Distributed Job Queues, SciNet Databases

Dynamic Branching of Tasks

Agents: Network Resource Manager, Route Optimization, Security Enforcement

Configuration Management, Conflict Resolution

SciNet Agentic Workflow

Make, Plan, Execute Tools, Learn, Optimize

Network Digital Twin

Testing, Configuration Checks

Agents: P4 runtime, SmartNIC Control, OLCF Job Submitter, MPQUIC Agent

Programmable Network Infrastructure Layer

Distributing workflows across multiple sites with databases, HPC, streaming and Cloud

PHASE-1 TECHNICAL CORE

1 · SEMANTIC INTENT

Telemetry + workflow + instrument metadata → tagged graphs

2 · PREDICTIVE

Scheduling: stage / policy / priority/ congestion / deadline

3 · VERIFIED CONTROL

Digital twin → MPQUIC / P4 SmartNIC / DPU actions

HP Enterprises PRIOR RESULT

10× lower AI-training latency

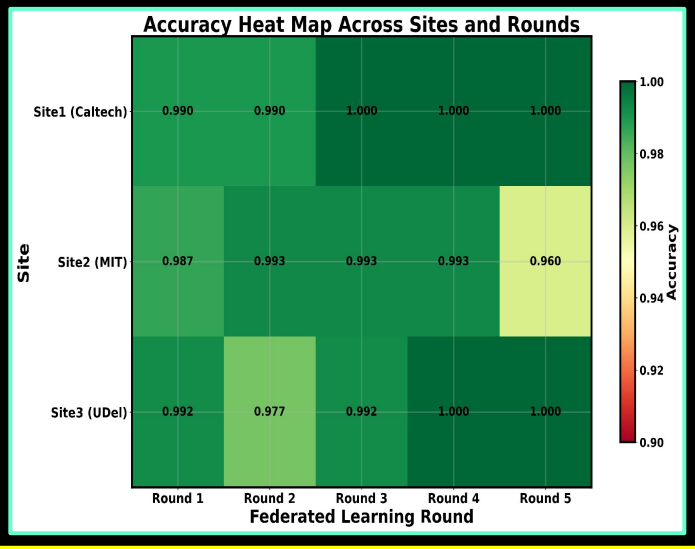
PHASE 1: HEP/Rubin + OLCF → **PHASE 2 Planned:** ESnet + NERSC production

NRE-115: Federated AI for Distributed

- **Scalability:** Architecture designed to support distributed sites with varying data volumes and processing capabilities
- **Protocols:** IPv6, Layer 2/3 switching optimized for data transmission and federated learning
- **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

Validation: classification efficiency of >98% for ten image types

Caltech, MIT, U. of Delaware

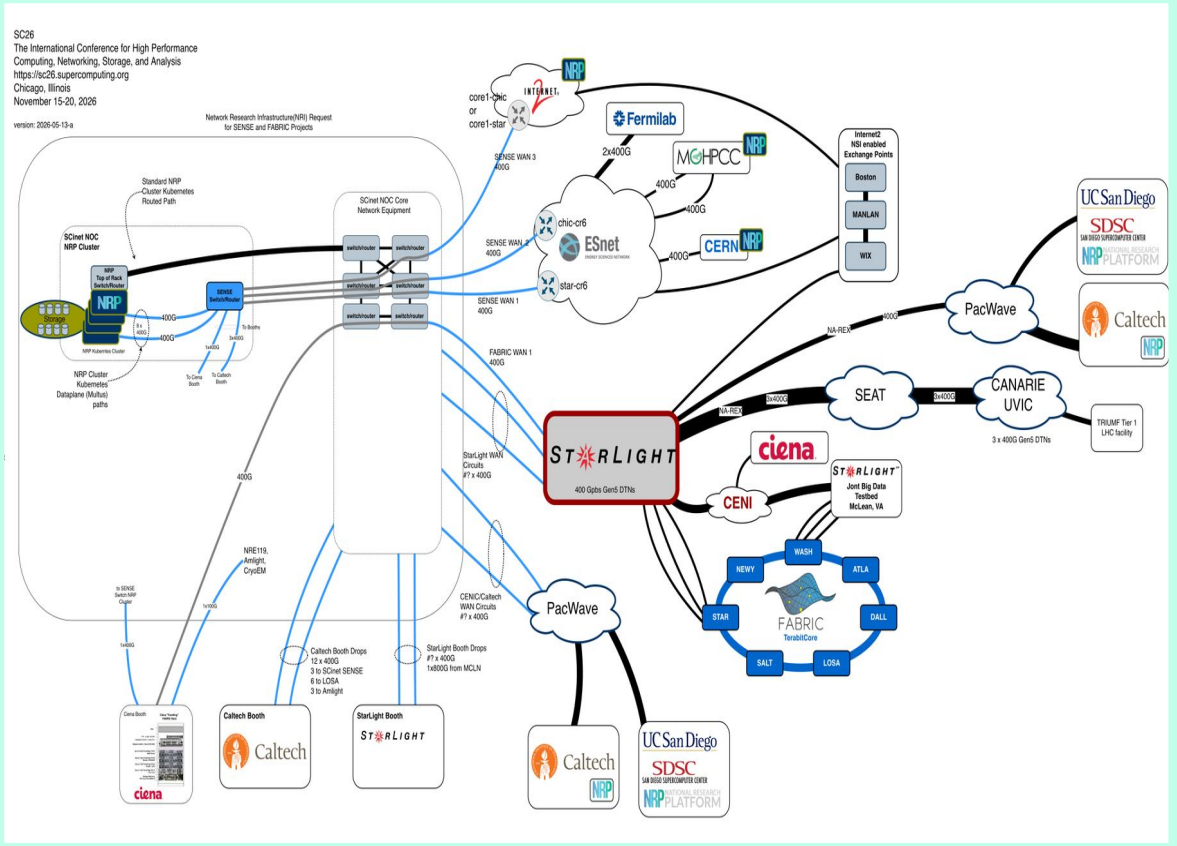


IMPACT: Results represent initial validation of federated learning concepts for distributed cosmic ray science, with ongoing efforts to characterize bandwidth requirements for multimodal data (images and sensor sequences) from edge devices including smartphone applications and low-power detectors

Using federated learning strategies including parameter exchange and selective data transmission based on model confidence, at SC25 accurate classification was achieved with data transmission of ~100 kbytes per round, out of Gigabytes of raw data

NRE-122: Customized end-to-end Multi-Resource Services for Science Domain Workflows via SENSE

- **IMPACT:** This demonstration will show SENSE in operation across a variety of infrastructures including: (1) Networks and Exchange Points (ESnet, Internet2, GNA-G AutoGOLE, CENIC, PacificWave, AMPATH/AMLight, StarLight, NetherLight Exchange, CERN); (2) Edge Systems (NRP Kubernetes Cluster, DTNs, LHC XRootD/dCached systems, Amazon Web Services); and (3) the FABRIC infrastructure which includes network embedded compute/storage resources



- **SENSE services and features which benefit multiple science workflows that will be shown include:**
  - Rucio/FTS/XRootD/dCache LHC CMS/ATLAS
  - DOE Integrated Research Infrastructure (IRI)
  - GNA-G AutoGOLE Applications
  - National Research Platform (NRP) Integration
  - Realtime workflow specific monitoring and workflow topology visualization

NRE-106: National Research Platform and SENSE: WAN Native LLM Orchestration Across Internet2, NRP and SCinet

GOALS:

- Demonstrate WAN-orchestrated LLM serving by routing live inference requests from SCinet-attached Nautilus nodes to LLM backends on PNRP (SDSC, MGHPCC) and Internet2-connected caches, with real-time path control via SENSE
- Showcase deterministic QoS for AI workloads using ESnet SENSE, establishing dedicated 100Gbps Layer 2 VLANs with guaranteed bandwidth among distributed inference endpoints, overcoming best-effort internet variability
- Highlight the NRP's Envoy AI Gateway and vLLM/SGLang backends operating across multiple administrative domains, with per-model concurrency policies and scavenging replicas that yield zero-reservation, low-preemption WAN access for AI experiments
- Run live, attendee-driven LLM experiments including cross-site model inference, fine-tuning, and RAG (retrieval-augmented generation) workloads that transparently span SCinet, PNRP, and Internet2 resources
- Validate the SENSE Kubernetes Operator as a production-grade orchestrator for AI workflows, building on prior cross-testbed integrations with FABRIC and AutoGOLE to enable programmable, intent-based WAN provisioning for LLM services

NRE-139: High-Performance Scientific Data Paths with PolKA and Programmable DTNs

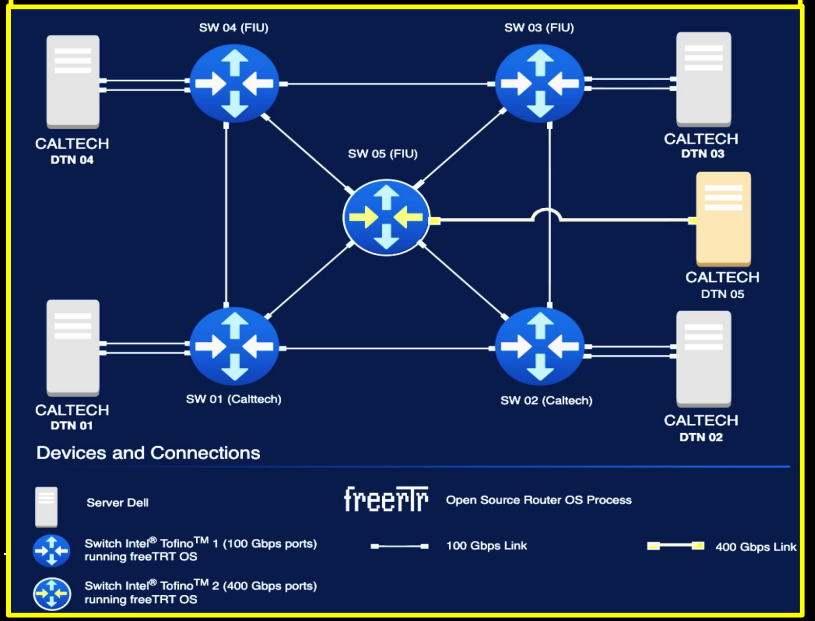
OBJECTIVES:

- **End-to-End 400 Gbps Data Transfers:** Sustained high-throughput transfers by combining PolKA tunnels with optimized DTN processing pipelines
- **Joint Network and Endpoint Optimization:** Demonstrating how path-aware routing and eBPF-based congestion control and telemetry feedback in a closed loop can dynamically adapt to traffic behavior
- **Programmable Congestion Control at the DTN:** Implementing adaptive transport mechanisms using eBPF/XDP, overcoming the rigidity of kernel-based congestion control algorithms (e.g. CUBIC, BBRv3)
- **Global Testbed Validation:** Executing long-distance experiments across AmLight-ExP andGP4L, including in “data tsunami” scenarios between Brazil & the US
- **Enhanced Observability and Resilience:** Integrating fine-grained telemetry in both network and host layers to enable rapid fault detection, mitigation, and traffic steering with quick path

IMPACT

By tightly integrating network programmability with endpoint intelligence, this approach advances the state of the art in high-performance data movement, providing a practical, deployable solution for next-generation scientific infrastructures

Experimental Topology at the Caltech Booth

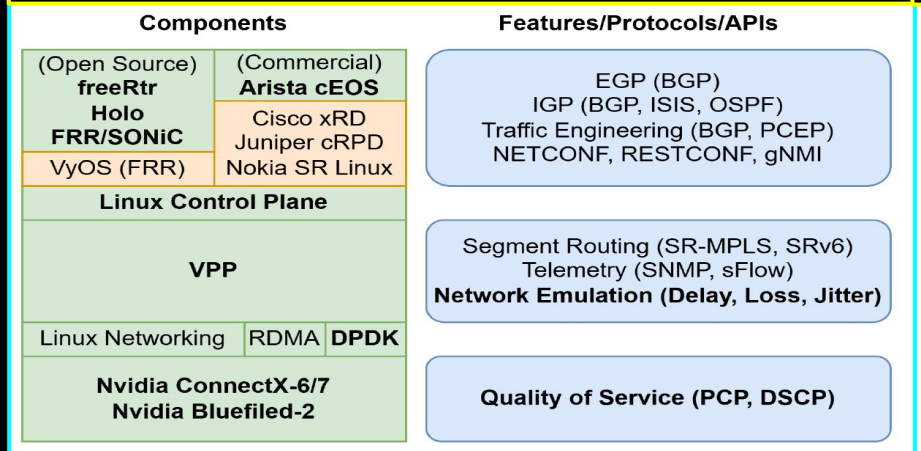


NRE-128 Global Platform for Labs (GP4L)

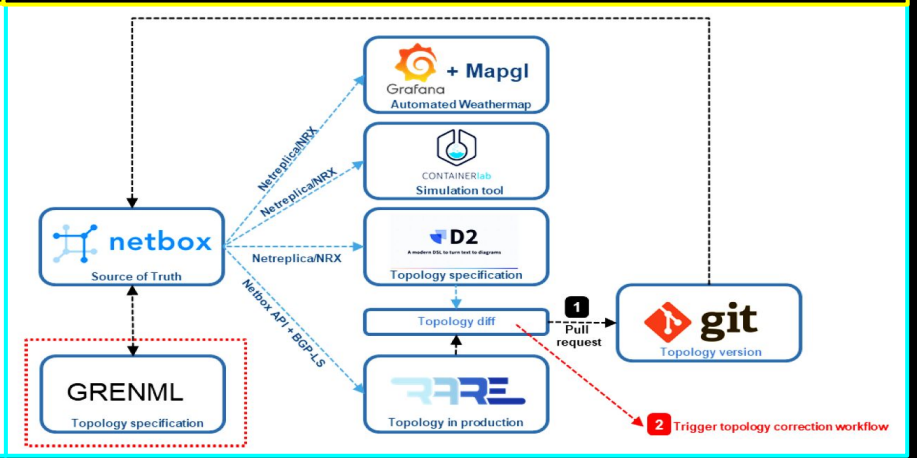
Marcos Schwarz RNP, Frederic Loui, RENATER

- GP4L (Global Platform For Labs) is a global overlay network composed of 100G/400G programmable switches and SmartNICs, spanning the GNA-G AutoGOLE/SENSE Persistent Multi-Resource Testbed, NRP, FABRIC & the GEANT RARE initiative
- It provides routing services using multiple open network operating systems which implement traditional networking protocols and can be extended to integrate next generation features as an opt-In strategy
- This showcase will demonstrate (1) Advancements in a high performance multi-OS software dataplane based on SmartNICs, VPP and multi control plane options (Left) (2) New integration capabilities, generating a Digital Twin from GNA-G GREN Map topology files (Right) and (3) Support for other NREs experimenting with at-scale advanced/novel networking protocols and use cases:
- Such as PolKA (Polynomial Key-based Architecture) a stateless source routing protocol based on arithmetic operations over a polynomial encoded route label

Software Dataplane Architecture



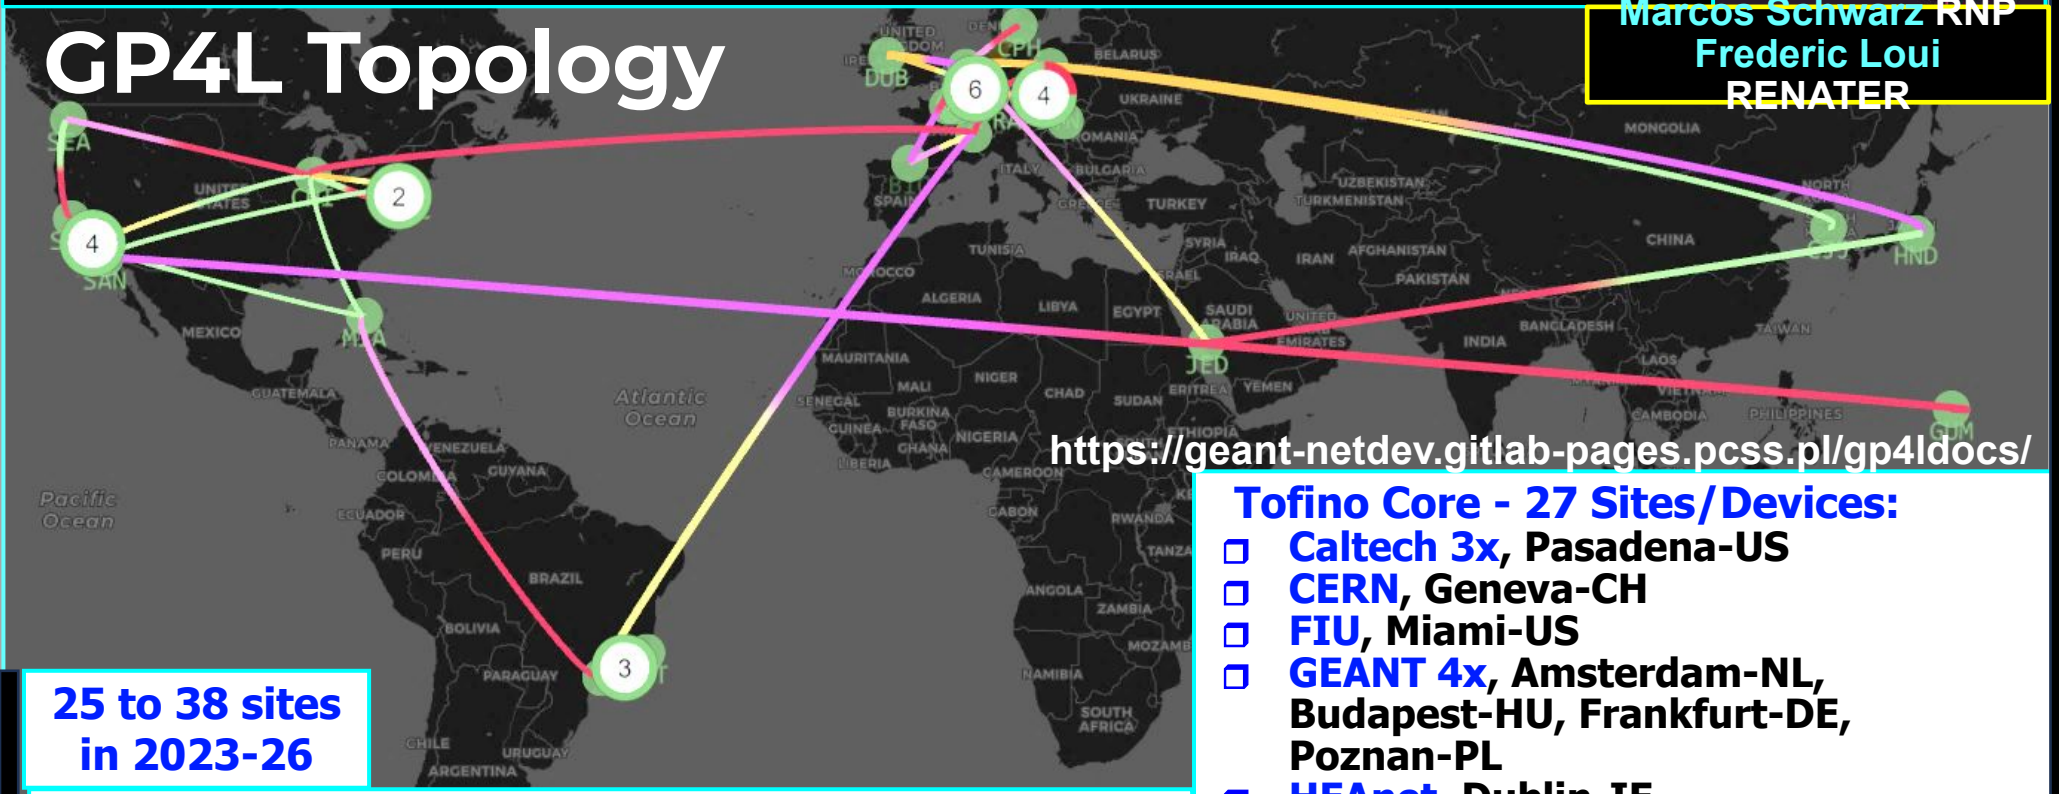
GP4L Automation Use Case



Global Platform for Labs (GP4L)

Marcos Schwarz RNP
Frederic Loui
RENATER

GP4L Topology



25 to 38 sites in 2023-26

BlueField-2/DPDK Islands – 7 Sites/Devices:

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

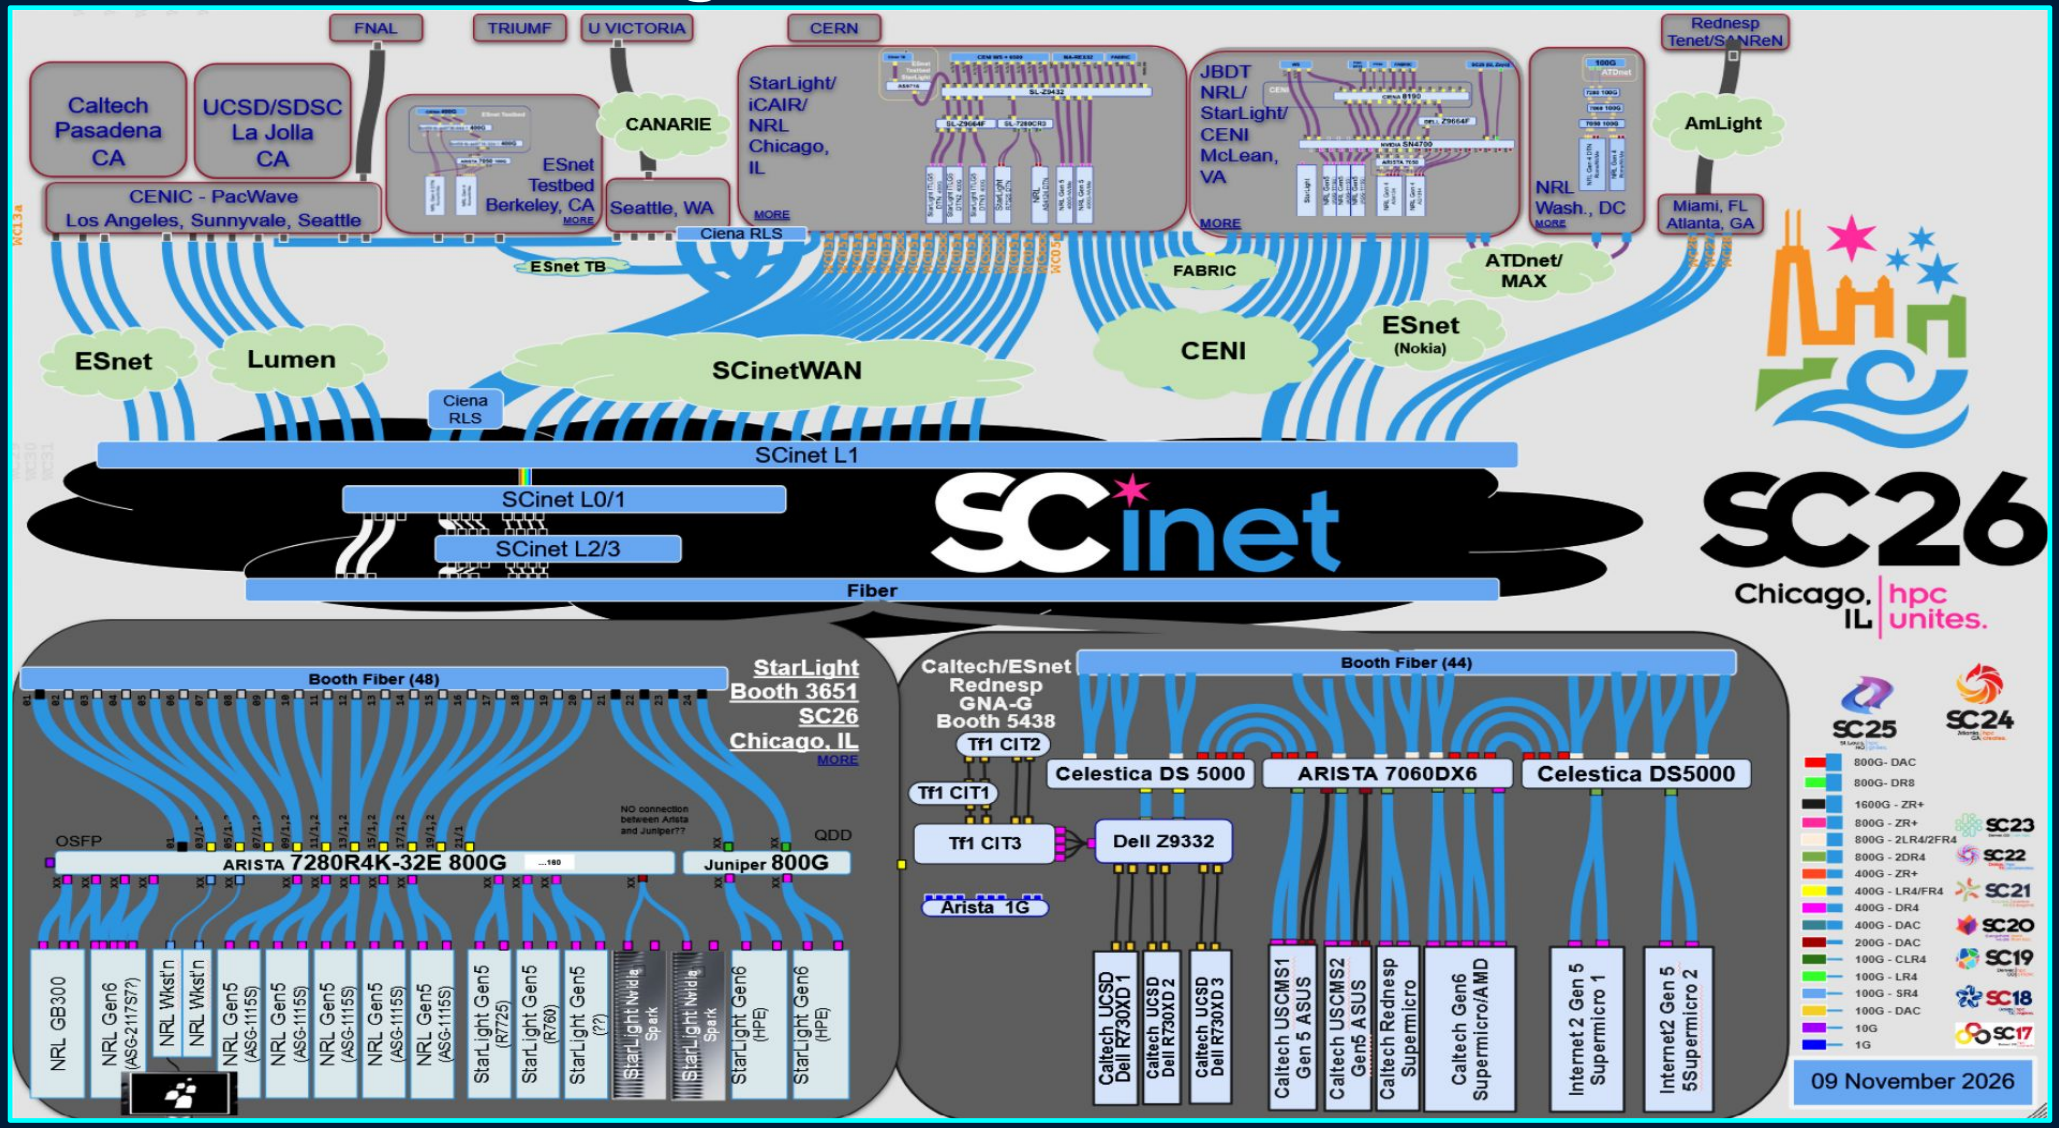
x86/DPDK Islands – 4 Sites/Devices:

- FABRIC, Miami-US
- 2x GEANT, Paris-FR, Prague-CZ
- KAUST, 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
- SC26, Chicago-US
- SouthernLight, São Paulo-BR
- StarLight, Chicago-US
- SWITCH 6x, Geneva-CH
- Tennessee Tech, Cookeville-US
- UFES, Vitória-BR
- UFMG, Belo Horizonte-BR

Caltech and StarLight/NRL Booths and SCinet at SC26



SC26: Global footprint. Multi-Terabit/sec Triangle Starlight – McLean – Chicago
22 X 400G to the Caltech Booth: 8 X 400G LA-Chicago;4 X 400G + 6 X 100G to Caltech Campus
3 X 400G to Latin America and South Africa; 2 X 400G to Fermilab; 4 X 300G to CERN
with CENIC, Ciena, Internet2, ESnet, StarLight, AmLight, Rednesp, RNP, US CMS & Partners

CENIC, ESnet and Internet2 at the LA PoP

2 Tbps to Caltech HEP
4 X 400G + 4 X 100G
CENIC-CIT via WS Ais

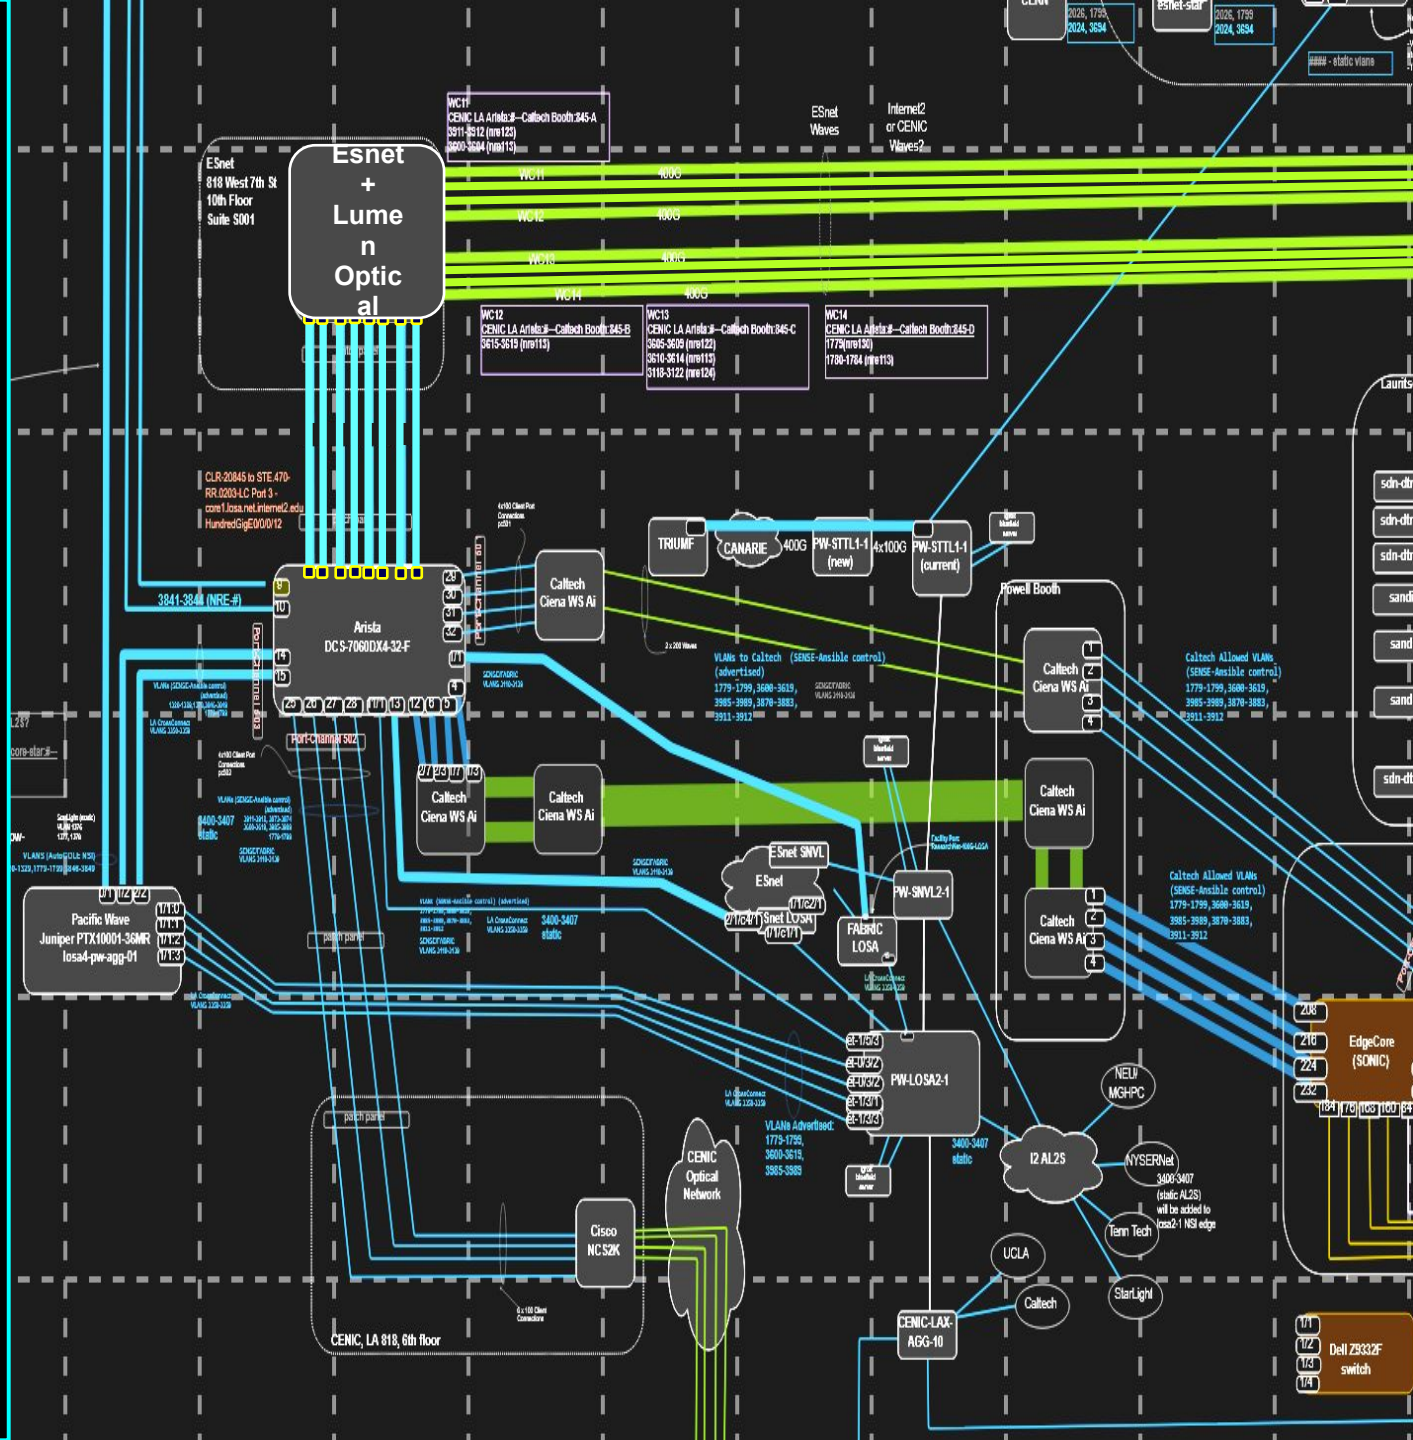
8 X 400G LA-Chicago
via Esnet, Lumen

2 X 100G to Caltech
Campus IMSS

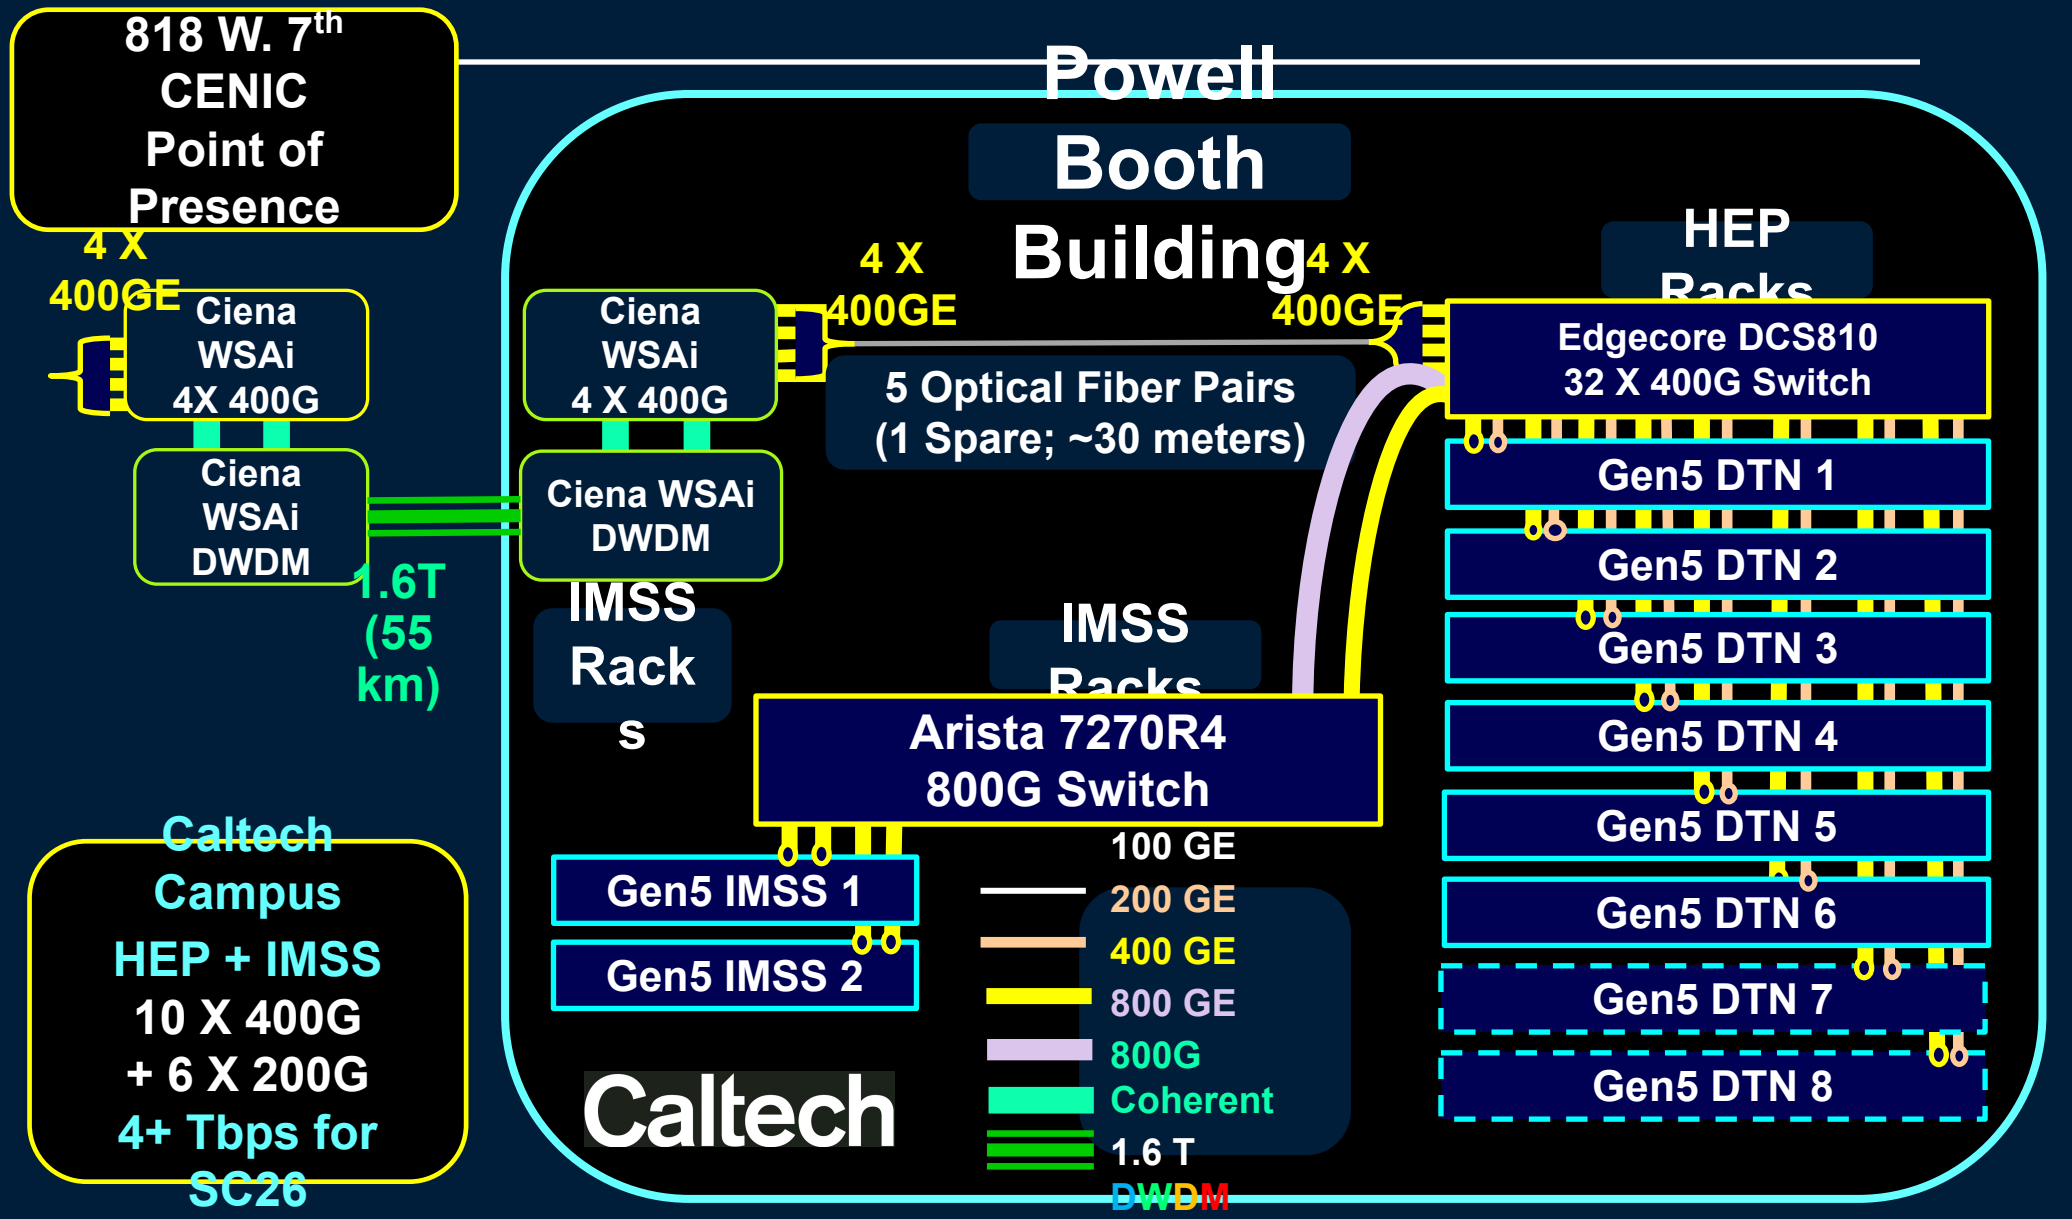
4 x 100G
to UCSD/SDSC

400G NA-REX

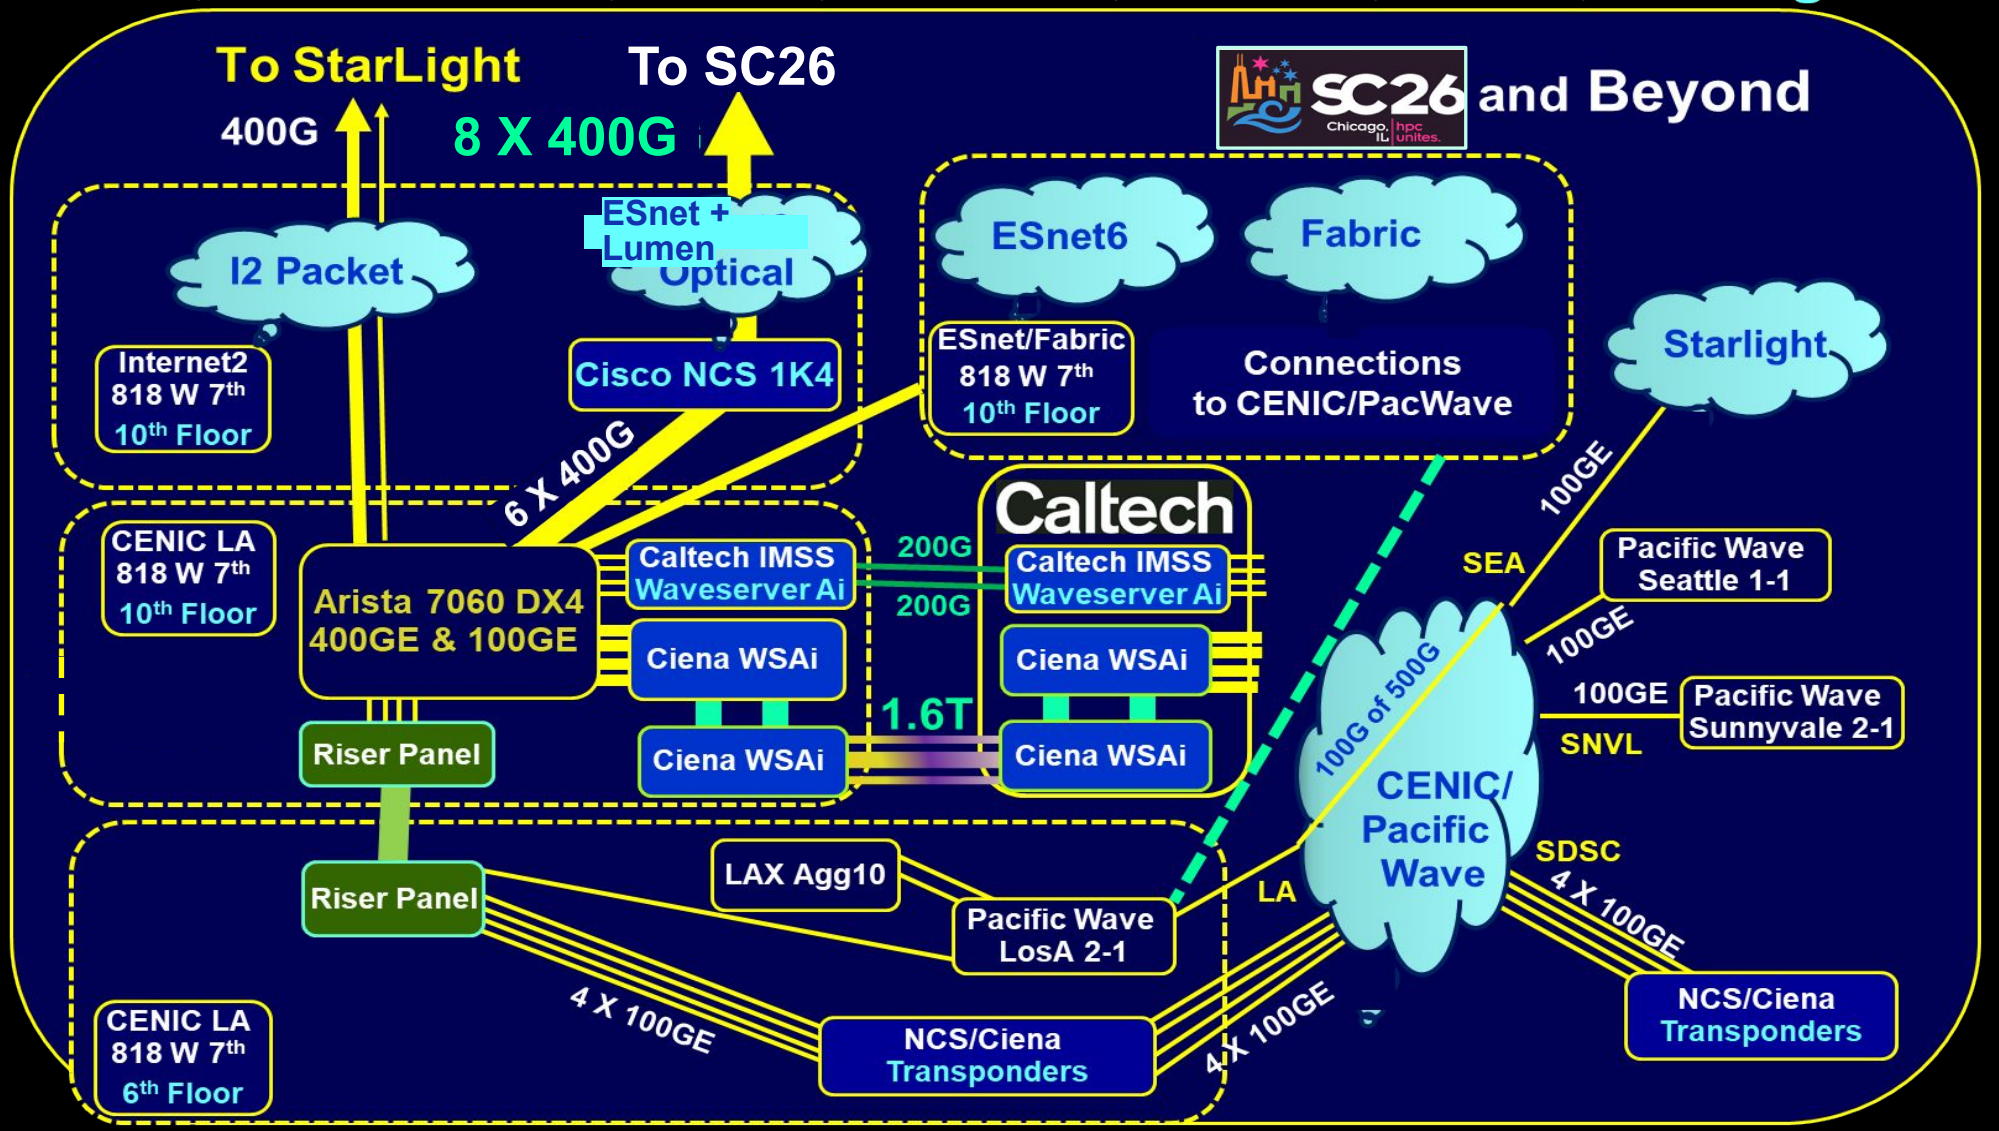
400G to ESnet
Production



Simplified Caltech – LA 4 X 400GE Layout



A New Generation Persistent 2Tbps Super-DMZ: Ciena, Arista, CENIC, Pacific Wave, ESnet, Internet2, Caltech, UCSD, StarLight++



SC26: 8 X 400G on ESnet, Lumen LA-CHI: Ciena, Caltech and CENIC Using WS Ais and a dark fiber pair. Bringing 4 X 400GE via 2 800G Waves direct to the campus

Gen 5 & 6 Supermicro + ASUS DTNs

Servers: Supermicro + AMD Venice, EPYC CPUs to 5 GHz GHz; Total ~1800 Threads

Supermicro AS-2127H7-N

2 Venice SP7 9996 CPUs

2 ConnectX-9s Bluefield 3

8 Solidigm P1010 NVMes



5 Gen5 + 1 Gen6 Server at SC26 with

~15 400GE ConnectX-7s++: 6 Tbps

+ 3 Gen3 Servers with 2 X 100G: ~0.6 Tbps

Tbps+ Storage: 14 Gen5 NVMe SSDs

Tofino1 CIT3
Tofino1 CIT2
Tofino1 CIT1
Celestica DS5000 64 X 800G Switch

Supermicro/AMD Venice Gen6 5 X 400G
Supermicro(REDNESP) Gen5 2 X 400G

Arista 7060X6 64 X 800G Switch

Supermicro (Internet2) 2 X 400G Server 1
Supermicro (Internet2) 2 X 400G Server 2

Celestica DS5000 64 X 800G Switch
Console
PFSense SOPHOS XG 201
Dell Z9332F 32 X 400G Switch
Dell Z9100 32 X 100G Switch
Arista 7010T 1G Management Switch
Celestica DS5000 64 X 800G Switch

ASUS520A Gen5 2 X 400G Caltech/USCMS 1
ASUS520A Gen5 2 X 400G Caltech/USCMS 2

Dell 730XD DTN 2 X 100G UCSD 1 (2U)
Dell 730XD DTN 2 X 100G UCSD 2 (2U)
Dell 730XD DTN 2 X 100G UCSD 3 (2U)

To ~7 Tbps in 1.3 Racks

The Road Ahead: 2026 and Beyond

- **SC26 and Beyond: Gen5 to Gen6 Servers: 400G to 800G per port; 1.6T/Server**
  NVMe Storage: 10 → 20 GBytes/sec/SSD (e.g. Micron 9550 → 9650)
- **N x 400G → N X 800G Switches: Examples Celestica DS5000, Arista 7060 X6**

Arista 7060 X6-64PE	Celestica DS5000
	

64 X 800G OSFP; Runs SONIC;
Support for 800G ZR/ZR+;
800G → 2 X 400G Transceivers
for ConnectX-7 or 8; Bluefield 3;
800G → ConnectX-9, Bluefield 4
Caltech Campus + CENIC PoP; SC26

Note: 64 X 1.6T at SC27 ?

- **2026 in Production: 200G per lane (224 Gbps signaling):**
  e.g. 800G DR4. Or DR8 (100G per lane)
- **OFC26: The Race to 400G per lane (448 Gbps signaling):**
  e.g. 1.6T DR4 (400G per lane)
- **OIF Interoperability Demos: 224 Gbps and 448 Gbps**
  https://www.oiforum.com/wp-content/uploads/CEI-Rolling-Deck.pdf
- **New modulation schemes: PAM-4 → PAM-6 → PAM-8**
  More bits per symbol, but needs higher S/N for the same Symbol Error Rate (SER)
- **Technical Paper: Keysight Race to 448 Gbps:**
  https://www.keysight.com/us/en/assets/7125-1060/white-papers/Race-to-448-Gbps.pdf

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**
- ✳ **Leveraging the worldwide move towards a fully programmable ecosystem of networks and end-systems** plus **operations platforms** (NRP, OSG; global SENSE Testbed, GP4L, FABRIC)
- ✳ **Building on ongoing R&D projects: from regional caches/data lakes to intelligent control and data planes to AI/ML-based optimization**
  [SENSE/AutoGOLE, SciNet, ESNet HT, GEANT/RARE, AmLight, FABRIC, GP4L; PolKA, Kytos, NetPredict, Hecate, DeepRoute, Qualcomm GradientGraph, ...]
- ✳ **Key milestones:** integration of SENSE + network services **with: Rucio/FTS/ XRootD/dCache for LHC; Globus; EJFAT;** 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 worldwide communities**

Building an AI-native Research Ecosystem for Experimental Particle Physics: A Community Vision

A whitepaper presenting a vision for how Artificial Intelligence (AI) can accelerate science discovery: 500 Authors, ~120 FTEs



arXiv https://arxiv.org/abs/2602.17582

High Energy Physics - Experiment

[Submitted on 19 Feb 2026]

Building an AI-native Research Ecosystem for Experimental Particle Physics: A Community Vision

Thea Klaeboe Aarrestad, Alaa Abdelhamid, Haider Abidi, Jahred Adelman, Jennifer Adelman-McCarthy, Shuchin Aeron, Garvita Agarwal, Usman Ali, Cristiano Alpigiani, Omar Alterkait, Mohamed Aly, Oz Amram, Saeed Ansari Fard, Aram Apyan, John Arrington, Marvin Ascencio-Sosa, Mohammad Atif, Aneesha Avasthi, Muhammad Bilal Azam, Bhim Bam, Joshua Barrow, Rainer Bartoldus, Amit Bashyal, Aashwin Basnet, Ayse Bat, Lothar A. T. Bauerdick, John Beacom, Chris Bee, Michael Begel, Matthew Bellis, Rene Bellwied, Rakitha Beminiwattha, Gabriele Benelli, Douglas Benjamin, Catrin Bernius, Binod Bhandari, Avinay Bhat, Meghna Bhattacharya, Saptaparna Bhattacharya, Prajita Bhattarai, Sudip Bhattarai, Wahid Bhimji, Jianming Bian, Burak Bilki, Mary Bishai, Kevin Black, Kenneth Bloom, Brian Bockelman, Johan Sebastian Bonilla Castro, Tulika Bose, Nilay Bostan, Othmane Bouhali, Dimitri Bourilkov, Dominic Brailsford, Gustaaf Brooijmans, Elizabeth Brost, Maria Brigida Brunetti, Quentin Buat, Brendon Bullard, Jackson Burzynski, Paolo Calafiura, Rodolfo Capdevilla, Fabian Andres Castaño Usuga, Raquel Castillo Fernandez, Fabio Catalano, Viviana Cavaliere, Flavio Cavanna, Giuseppe Cerati, Aidan Chambers, Maria Chamizo-Llatas, Philip Chang, Andrew Chappell, Arghya Chattopadhyay, Sergei Chekanov, Jian-ping Chen, Yi Chen, Zhengyang Chen, J. Taylor Childers, Hector Chinchay, Yuan-Tang Chou, Tasnuva Chowdhury, Neil Christensen, Wonyong Chung, Rafael Coelho Lopes de Sa, Simon Corrodi, Kyle Cranmer, Matteo Cremonesi, Roy Cruz, Mate Csanad, Mariarosaria D'Alfonso, Carlo Dallapiccola, Daine Danielson, Sridhara Dasu, Gavin Davies, Kaushik De, Patrick de Perio, Klaus Dehmelt, Marco Del Tutto, Carlos Ruben Dell'Aquila, Sarah Demers et al. (359 additional authors not shown)

"We outline grand challenges that must be addressed to enable transformative break-throughs and describe how current and planned experimental facilities can implement this vision to advance our understanding of the universe

We show how facilities currently under construction, such as the HL-LHC, DUNE and soon EIC, can benefit from and serve as proving grounds for this vision, while also enabling a longer-term goal for how future experiments -- like FCC-ee at CERN, IceCube-Gen2, a Muon Collider in the US, and smaller-scale projects

– Can Be Fully AI-Native

We describe how a truly national-scale and int'l collaboration, jointly managed across large funding partners, and involving both laboratories and universities, can make this happen"

An invitation to join this community effort

Tulika Bose, Peter Elmer, Gordon Watts, et al.

AI Native Experimental Lifecycle

Four Grand Challenges (GC):

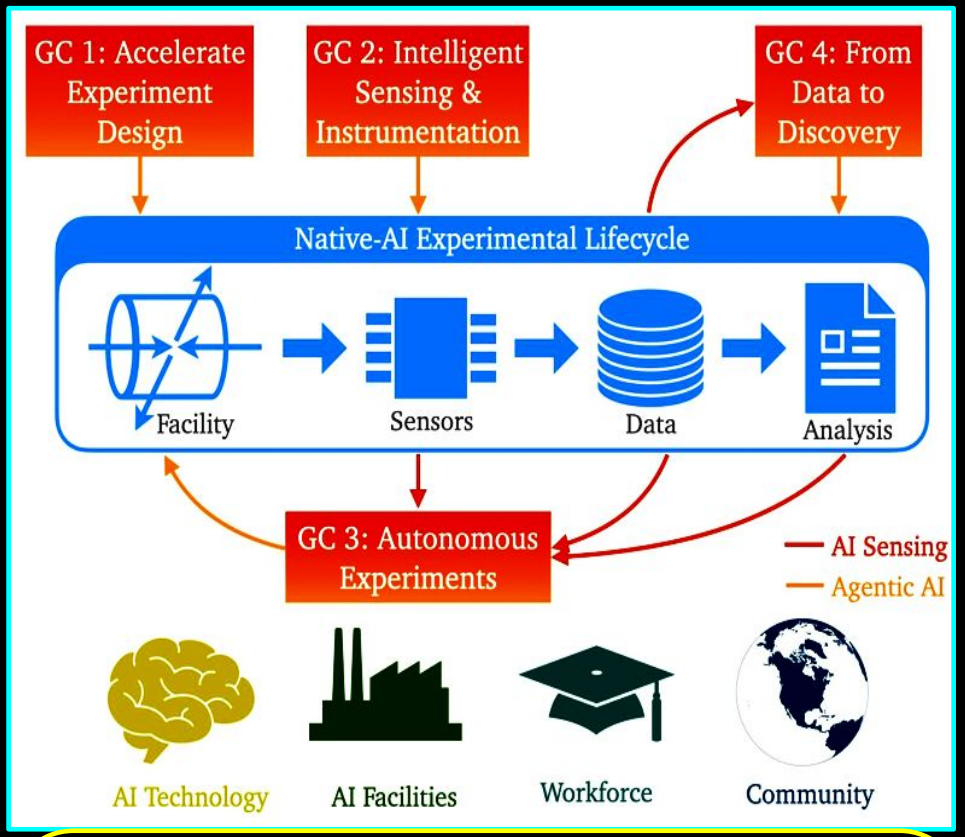
1. Accelerated Experimental Design
2. Intelligent Sensing & Instrumentation
3. Autonomous Experiments
4. From Data to Discovery

Are inter-connected and create a self-reinforcing pipeline:

A natural organizing principle for

A national-scale, multi-institutional collaboration that unites universities, laboratories, industry

Shared technologies and infrastructure



Facilities under construction (e.g. HL-LHC) can be testing grounds for this vision while future facilities (such as a muon collider) can be fully AI-native

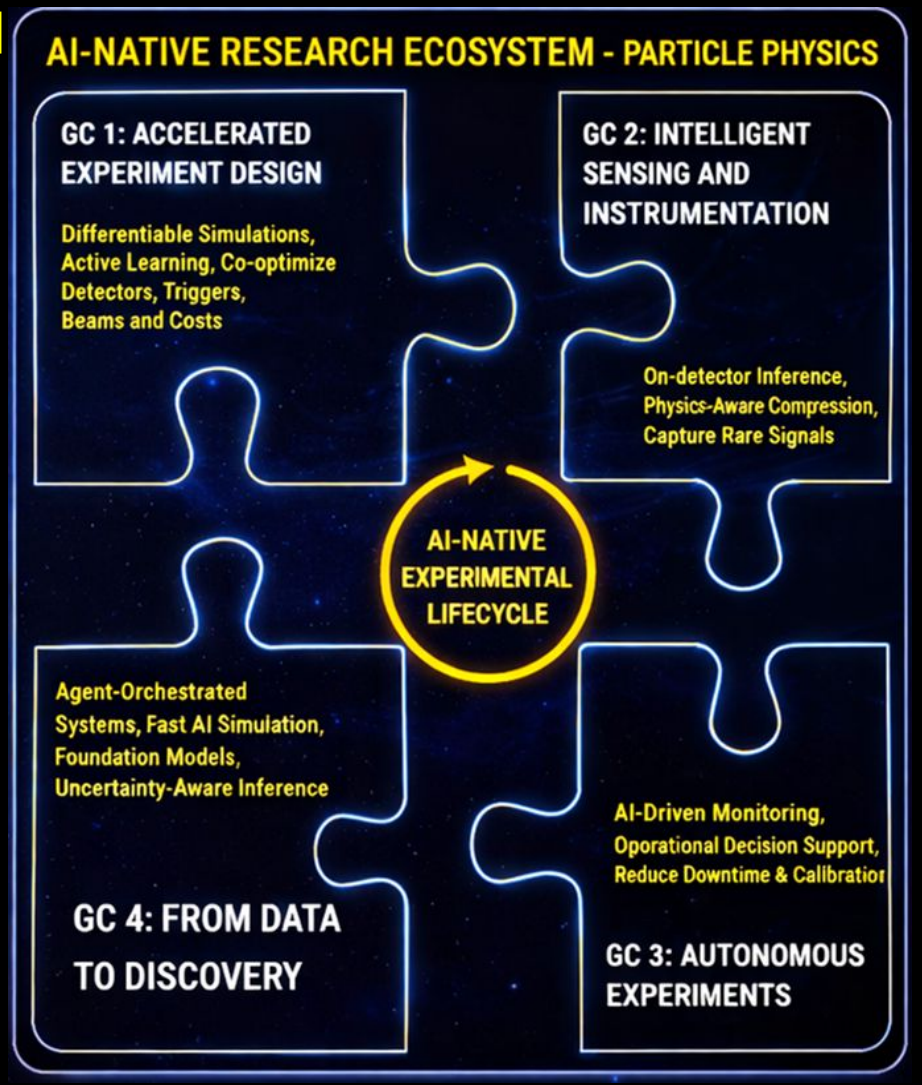
Tulika Bose

Closing the Loop: from Design to Discovery

The Grand Challenges are Connected

- ✳ **Accelerated design** informs intelligent instrumentation
- ✳ **Intelligent sensing** enables autonomy
- ✳ **Autonomous operation** feeds higher-quality data into inference
- ✳ **From data to discovery** closes the loop, **back to design**

Together these challenges define an end-to-end AI-Native experimental paradigm



Note: there are some cross-cutting themes and emerging opportunities that do not quite fit in these GC

GC3: Autonomous Experiments

Objective

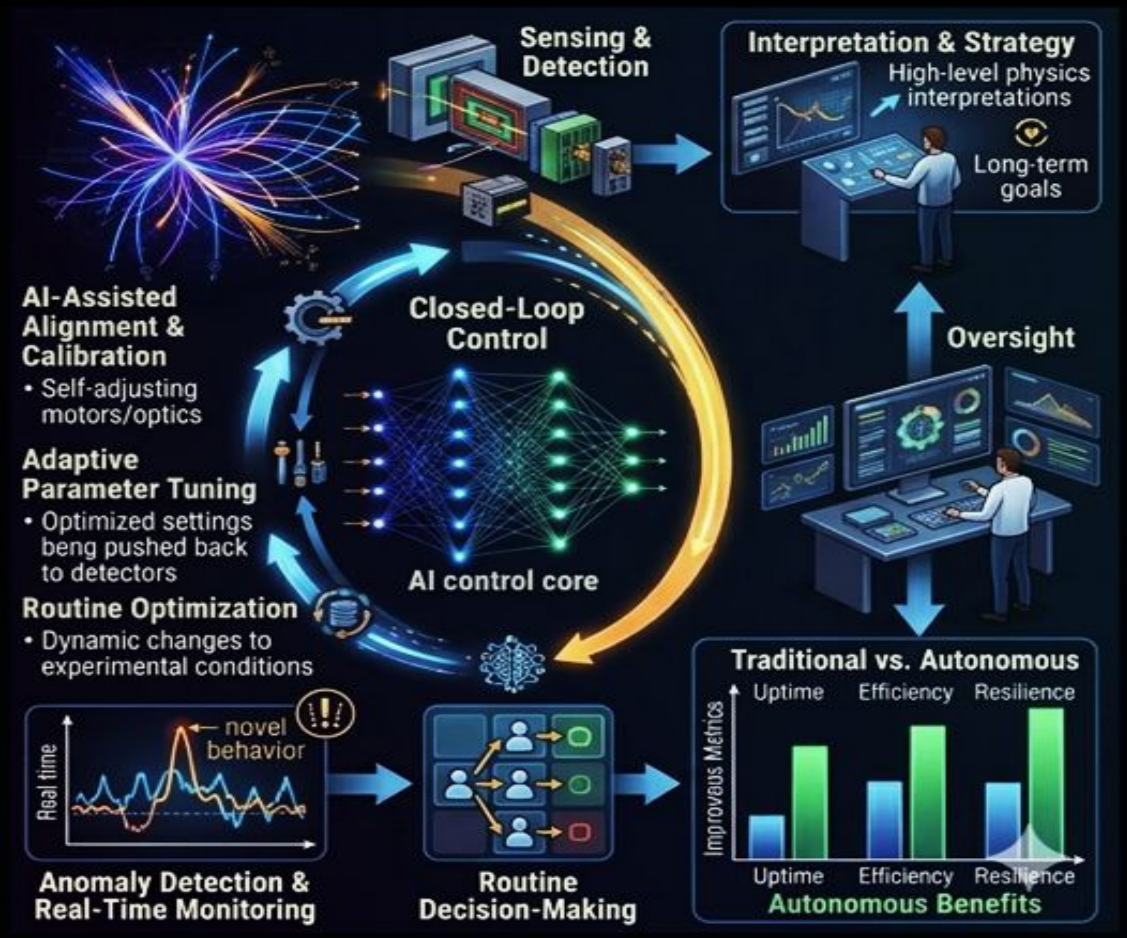
- Enable Experiments to monitor, adapt and optimize themselves with appropriate human oversight

Key Elements

- AI-assisted calibration, alignment and monitoring
- Closed-loop control of experimental conditions
- Real-time detection of anomalies and novel behavior
- Routine decisions handled by AI; strategy trained by people [++]

Why It Matters

- Autonomy increases resilience, scientific uptime, and consistent data quality



Proactive operations • less downtime • preserved expertise

Transform labor-intensive, reactive operations into proactive, resilient, and continuously calibrated systems, capturing more high-quality data with less downtime and preserving institutional knowledge over decades-long experiment lifetimes

Eric Moreno,
Phil Harris et al

"Just Furnish Context": AI Agents Can Already Autonomously Perform Experimental HEP Analysis



Eric A. Moreno[1,2,*] — RESEARCHER WITH LHC PHYSICS PHD @ MIT

Samuel Bright-Thonney[1,2,*] — IAIFI FELLOW @ MIT

Andrzej Novak[1,2,*] — POSTDOC @ MIT HEP-EX & ML

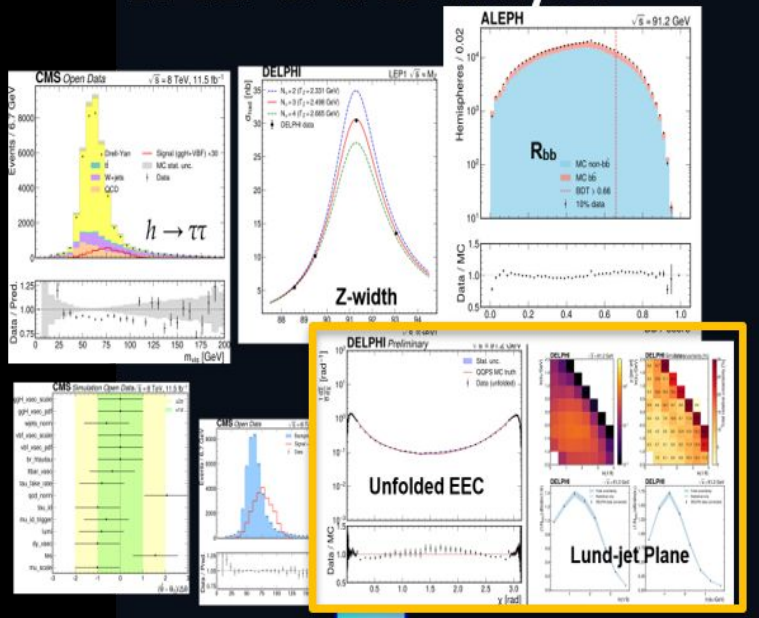
Dolores Garcia[3] — FELLOW @ CERN

Philip Harris[1,2] — ASSOCIATE PROFESSOR OF PHYSICS @ MIT

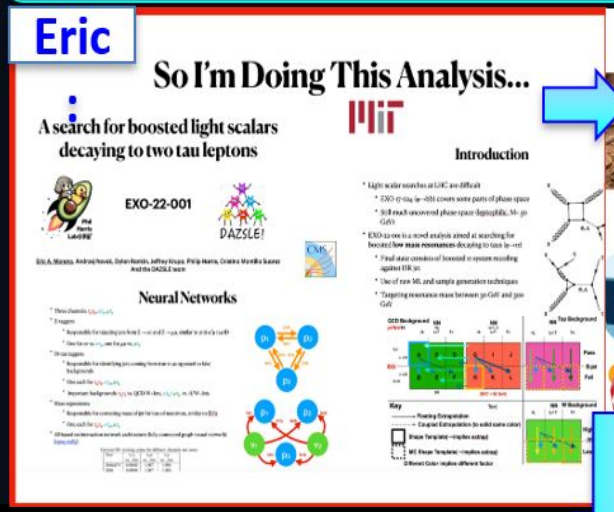
arXiv:2603.20179v3 [hep-ex] 23 Jun 2026

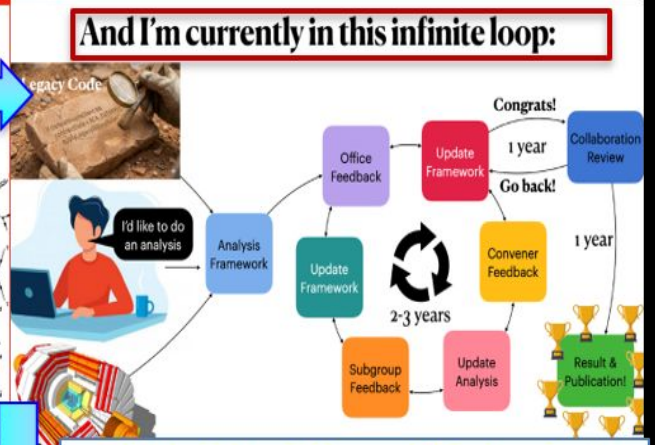
https://jfc-mit.github.io/

2 Days, then 3 weeks of Running;
10 LEP or LHC Analyses



Eric





Phil : Like Many Students !

How to Get Out of the Infinite Loop ?
Eric and Phil: *Let's Go For It !*

Phil's Message: *This is Not Hard*

✸ What and How ?

Review Is a System — Not a Single Critic

Six Reviewers + One Arbiter; Humans Control Unblinding

Moreno, Harris

MULTI-AGENT REVIEW STACK

EXECUTION

EXECUTOR ×1
Generalist agent plans + runs all phases

SPECIALIZED REVIEW PANEL

PHYSICS
Independent ARC Physics completeness Checklist-blind

CRITICAL
8-step review 14-point figure audit Regression + upstream

CONSTRUCTIVE
Alternatives + cross-checks Clarity + Notation Effort estimate

DECISION + AUTHORITY

ARBITER ×1
A / B / C synthesis PASS · ITERATE · ESCALATE

SPECIALISTS ×4
Fixer
Investigator
Note-writer
Typesetter

PLOT VALIDATOR
8 figure checks
11 physics + 6 consistency
10 red flags ⇒ A

RENDERING
Phase 5
8 PDF/layout checks
No physics judgment

BIBTEX
Phases 4a / 4b / 5
Format + real-record checks; INSPIRE cross-check

HUMAN GATE
After fixed-seed 10% data
APPROVE · CHANGE · HALT
Owns full unblinding

Automation catches large classes of errors/omissions early
Humans retain authority over unblinding, interpretation, and publication

The Coming Challenges: Reliable New Analyses for Discovery and/or Precision

Beyond CMOS

The Next Computing Paradigm Beyond Conventional Transistor Scaling

IEEE International Roadmap for Devices and Systems™
2024 Edition – Beyond CMOS

WHY BEYOND CMOS?

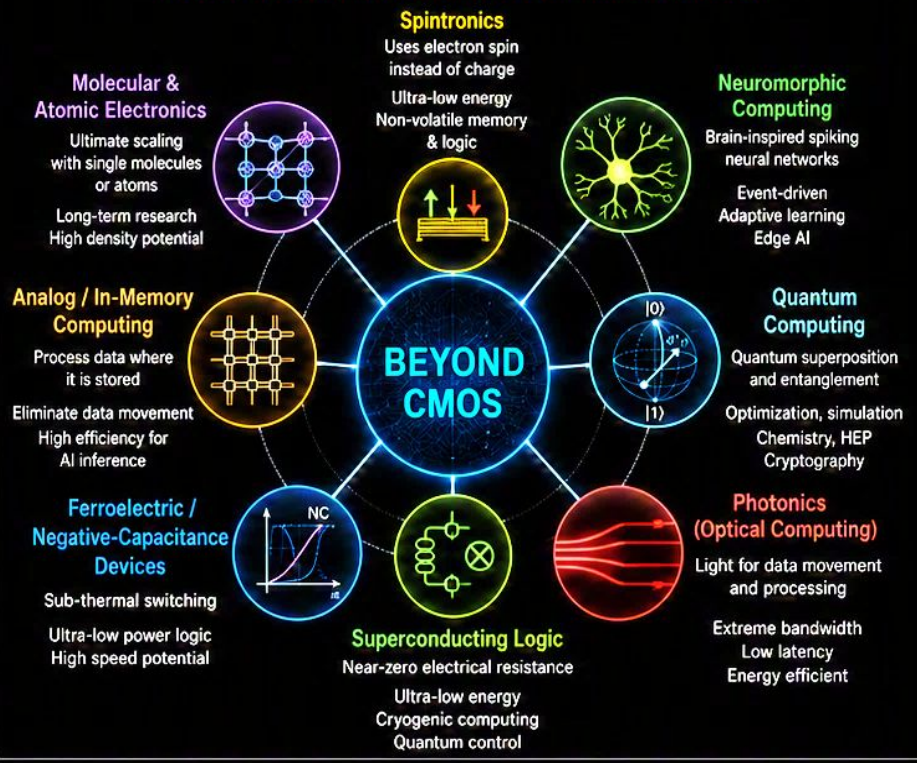
As conventional CMOS approaches atomic dimensions, future gains increasingly come from new physical state variables, new materials, new architectures, and new computing paradigms rather than continued transistor scaling alone.

The 2024 IEEE IRDS emphasizes co-optimization of devices, architectures, and applications—not a single replacement, but a diverse portfolio of complementary technologies.

GUIDING PRINCIPLES (IRDS 2024)

- Information state variables beyond electron charge (spin, phase, polarization, flux, molecular states)
- Non-equilibrium computation for lower switching energy
- Novel energy-transfer mechanisms for communication
- Nanoscale and quantum enabled thermal management
- Sub-lithographic, self-assembly and novel manufacturing
- Alternative architectures co-designed with emerging devices and systems

A PORTFOLIO OF COMPLEMENTARY TECHNOLOGIES

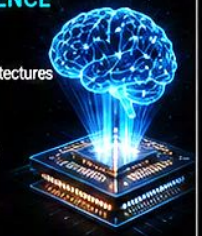


MAJOR TECHNOLOGY DIRECTIONS

Technology	Primary Advantage	Best Applications
Spintronics	Ultra-low energy, non-volatility	Memory, AI accelerators
Neuromorphic	Event-driven, adaptive learning	Edge AI, robotics, adaptive systems
Quantum	Massive parallelism for certain problems	Optimization, chem., HEP, simulation
Photonics	Extremely high bandwidth	AI clusters, HPC interconnects
Superconducting	Near-zero dissipation	Cryogenic computing
Ferroelectric / NC	Sub-thermal switching	Ultra-low-power logic
Analog / In-Memory	Eliminate data movement	AI inference, signal processing
Molecular / Atomic	Atomic-scale devices	Long-term research, ultra-high density

WHY IT MATTERS FOR AI-NATIVE SCIENCE

- Orders-of-magnitude improvement in energy efficiency
- New computational capabilities beyond von Neumann architectures
- Massive parallelism and new algorithms
- Specialized accelerators for AI, simulation, optimization, and scientific discovery
- Co-design of devices → architectures → algorithms → applications is the dominant paradigm





Beyond CMOS is not a single successor to silicon—
it is a portfolio of complementary technologies
that together define the future of
AI-native computing, high-performance computing,
and scientific discovery.

IEEE IRDS™

CMS

Source: IEEE International Roadmap for Devices and Systems™ (IRDS™) 2024 Edition – Beyond CMOS

Co-optimized Devices • Architectures • Applications

Takeaway Messages for the 21st Century

- We are ready to address the challenges of the HL LHC, DUNE, Future Colliders…
  But the capabilities, opportunities, solutions and challenges, will be very different
- Research in the 2030s-40s Will Be Transformed by Beyond CMOS Technologies and Physics Paradigms: from 2D and Metamaterials to Neuromorphic to Thermodynamic Computing
- The AI Revolution has barely begun; Accelerating and transforming society through:
  - Autonomous, Persistent, Agile Systems
  - Generative AI: The ability to synthesize, imagine (and Hallucinate)
  - Agentic AI Architectures Matched to Complex Systems
    with Multi-objective Optima & Constraints
  - Able to Plan, Co-Design, Structure, Execute Complex Tasks; Learn and Self-Improve
  - Remembering Information and Context: Personal, Enterprise, Community-Shared
  - Evolving from AI Assistant to Associate; to Collaborative Partner … to Mentor ?
- In order to plan, meet the challenges, and maximize our discovery reach,
  we must engage with, learn and take advantage of:
  - State of the art and emerging technologies and new paradigms;
    Including the AI revolution: in both its present and *future* trajectories
  - Learn and adapt, as future paradigms, opportunities and challenges appear
- ***Be ready to operate safely and maintain the principles of science and education in whatever new environments emerge, while fulfilling our missions of discovery and benefiting society***

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 has been 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, Celestica, Supermicro and Dell.

