SC26 GNA-G Network Research Exhibition and Partner NREs at the Caltech/ESnet Booth
A Next Generation Multi-Terabit/sec Campus and Global Network System for Data Intensive Sciences

The Global Network Advancement Group
Data Intensive Science and AutoGOLE/SENSE Working Groups

Abstract Submitted on behalf of the teams by:
Harvey Newman, Caltech, newman@hep.caltech.edu



Executive Summary

The Global Network Advancement Group (GNA-G), through its Data-Intensive Sciences (DIS) and SENSE/AutoGOLE Working Groups, is developing a next-generation, multi-terabit/sec global system to support the rapidly escalating demands of data-intensive science. Scientific programs such as the LHC, Rubin Observatory LSST, SKA, Cryo-EM, genomics, and distributed AI/ML workflows now generate unprecedented data volumes, requiring coordinated global infrastructure capable of predictable, high-throughput, multi-domain operation.

Building on the progress of previous Supercomputing conferences, especially SC25 and further developments since then, this NRE will deploy a state of the art set of end-systems capable of multi-Tbps throughput supporting AI-integrated workflows through many partner NREs in a wide range of disciplines.

This system integrates programmable networks (P4, SRv6, PolKA, BGP-CT), multi-domain orchestration (SENSE/AutoGOLE), real-time telemetry (INT, perfSONAR, ALTO, GradientGraph), ML-driven intelligence (NetPredict, Hecate, GradientGraph), high-throughput end systems (400G–800G DTNs, GPU nodes), and regional caches/data lakes into a unified, composable architecture. It is validated through global testbeds such as GP4L, FABRIC, NRP, and GRP, and deployed across multi-Tbps optical infrastructures including APOnet, AmLight, Rednesp/RNP, ESnet, Internet2, GEANT, and StarLight.

At SC26, this NRE and more than a dozen partners NREs will demonstrate this architecture at scale, showcasing multi-domain orchestration, programmable data planes, AI-driven optimization, Digital Twin simulation, and multi-Tbps global data movement. The system supports a progressive deployment model spanning R&D → advanced production → full production, ensuring coexistence with general R&E traffic while enabling the next decade of scientific discovery.

New and Other Key elements this year include:

- Integration with the Genesis Mission and the American Science Cloud, through our partnerships with the ESnet SENSE, National Research Platform, and Oak Ridge National Lab and teams. These teams have demonstrated and continue to support and develop multiple AI-driven applications with a uniform interface based on the DOE Integrated Research Infrastructure (IRI) API, spanning multiple sites, various topologies and working environments.

  By SC26, this should include the "Minimum Viable Product" (MVP) which is a major near term milestone of the Genesis Mission.

- Partnership with AmLight in supporting data transmission among the Vera Rubin Observatory, the main data processing center at SLAC, and end sites in the VRO scientific collaboration. This will be complemented by the latest high throughput demonstrations between the SC26 and the Rednesp links to and among the Sao Paulo universities, including the use of RDMA across intercontinental distances. The groups in South Africa and at UC Santa Cruz will also present the latest results in distributed cryo-EM workflows, supported by AmLight's Americas-Africa Lightpaths Express and Protect (AmLight-ExP) project.

- 800G native networking, including a pair of Celestica DS5000 64 port 800G switches (one at the SC26 venue and one at the Caltech campus), an Arista 7280R4 switch at the CENIC PoP in LA and one at the Caltech campus. This will be complemented by multiple 400G and other switches from Dell, Arista and Juniper at the venue, in SCinet and at multiple remote sites in the US, Europe, Latin America, Asia, the Middle East and South Africa.

- A new scale of local and wide area connectivity, including 12 400G links between the Caltech booth and SCinet. This will include 3-4 400G links to the SCinet cluster. 800G links to SCinet are also planned, as well as several wide area 800G links between SC26 in Chicago and Los Angeles. The associated target thoughput during the data tsunami will correspondingly be in the range of 10 Tbps to and from the Caltech booth alone.

- State of the art servers and network equipment supporting the multi-Tbps demonstrations described here, including several Gen 5 (PCIe 5.0) servers with dual 400G (NIVIDIA ConnectX-7) interfaces, and Bluefield 3 and Gen 6 servers with ConnectX-9 800G native interfaces as they become available.

- This has been enabled by strong industry support and collaboration, most notably from Supermicro, AMD, Ciena, Arista, Celestica, Dell and Juniper.

- The scope of this SC26 NRE can be appreciated by considering the wide range of results shown at SC25, summarized in Figure 1 below, along with the major new elements summarized above.

SC25 Network Research Exhibition NRE-125 and Partners
17 NREs Hosted at or Partnering with Caltech Booth 3224

NRE-101	Eli Dart (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	Justas Balcas, Xi Yang (SENSE) et al.	Multi-Resource Cyberinfrastructure Services for Science Domain Workflows via SENSE
NRE-117	Tom Lehman (SENSE, FABRIC) et al.	FABRIC: Adaptive ProgrammaBle Research Infrastructure for Computer Science and Science Applications
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-124	Jessa Westclark and Rafael Coelho Lopes de Sa (MGHPCC/UMass) et al.	Advanced capabilities in high-bandwidth research network 800G MGHPCC to ESnet
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, Jeronimo Bezerra (FIU) et al.	AmLight 2.0: Flexible control, deep visibility, and programmability @ Tbps!

Figure 1 NRE-125 and 16 partner Network Research Exhbitions at SC25

Abstract: Impacts

The Global Network Advancement Group (GNA-G) and its Data Intensive Sciences (DIS) and SENSE/AutoGOLE Working Groups, a worldwide collaboration bringing together major science programs, research and education networks, and advanced network R&D projects spanning the U.S, Europe, Asia, Latin America and Oceania, are developing a next generation network-integrated system that will meet the challenges of data intensive sciences, showing the way towards the next generation of intelligent operations of R&E networks, as well as new methods and modes of network operation and capacity management that will benefit both network research and production teams.

Building on the wide range of developments and results shown at SC25, this SC26 network research exhibition (NRE) will demonstrate, in partnership with more than a dozen 16 NREs covering specific areas, many aspects of the next generation system which the GNA-G is developing both for data intensive sciences and to help set the course for future R&E network operations.

As shown, several of the teams' demonstrations include state of the art programmable open source network operating systems, layer 2 multi-domain virtual circuit overlays coupled to extensive structured use of IPv6 address spaces, novel source routing approaches to efficient flow distribution, steering and load balancing, and stateful traffic engineering and workflow acceleration taking network segment and site state, policy and priority and real-time SLAs into account.

These advances, being developed within the GNA-G framework and its SENSE/AutoGOLE and Data Intensive Sciences (DIS) working groups, the Global Research Platform, and its many science, computer science and R&E network teams aim to address the key challenges including:

(1) global data distribution, processing, access and analysis, (2) the coordinated use of massive but stilllimited computing, storage and network resources, (3) coordinated operation and collaboration within global scientific enterprises each encompassing hundreds to thousands of scientists, and (4) enabling the science programs to make efficient use of the available network and site infrastructures, while simultaneously accommodating and in concert with the network operations required to supporting the worldwide academic and research community across national, regional and transoceanic boundaries.

In order to meet these high-level goals, the teams are working towards deploying dynamic and adaptive programmable software-driven system which coordinates worldwide networks as a first class resource along with computing and storage, across multiple domains. We are following a systems design approach to create a global dynamic fabric that flexibly allocates, balances and conserves the available network resources, while negotiating and working with the site-resident systems that aim to accelerate workflow.

Reinforcement and other model-based machine learning techniques are also planned to be used to optimize system operations, according to objective functions that take priority, policy, network and site state-changes, workflow objectives and other constraints into account.

An overarching concept is "Consistent Network Operations," where stable load balanced high throughput workflows crossing optimally chosen network paths, up to preset *high water marks* to accommodate other traffic, are provided by autonomous site- resident services dynamically interacting with network-resident services, in response to demands from the science programs' principal data distribution and management systems.

The demonstrations in this and partner NREs are based on a wide range of ongoing R&D projects: from regional caches/data lakes to intelligent control and data planes to machine-learning based based optimization, including: SENSE/AutoGOLE and its integration with the Rucio, FTS and XRootD data access and management systems, AmLight, FABRIC; NOTED, GEANT/RARE freeRtr and the GP4L Lab, PolKA, ESNet High Touch, NetPredict and Hecate, Qualcomm GradientGraph, among others.

A recently developed capability in SENSE is the use of ECMP, allows flexible sharing and load balancing of ensembles of flows along multiple paths. This facilitates the sharing of the network capacity on major R&E production network, among the flows generated by data intensive science programs and the traffic that supports the at-large academic and research communities.

Many of the aspects composing this NRE and its partners benefit from the Open Science Grid, National Research Platform (NRP) led by UCSD, as well as the Global Research Platform (GRP) led by iCAIR, and the Americas Research Platform (AmRP) led by FIU.

We are also leveraging the worldwide move towards a fully programmable ecosystem of networks and end- systems, bringing together the use of P4, NPL and other programmable network services; segment routing including SRv6 and SR MPLS along with PolKA; and other emerging core concepts and components such as BGP Classful Transport (BGP- CT) that enables multidomain overlays with defined service levels and other attributes.

Continued themes this year include: (1) the GP4L Lab (GP4L), its architecture and capabilities, (2) a broader view of the programmable networking ecosystem including open-source operating systems (P4, SONIC, NPL, BGP-CT, gNMI and other mainstream programmable approaches, and (3) the synergy between GEANT/RARE freeRtr and the major vendors' OSes and tools. A key recent development associated with the GP4L effort is the development of a Digital Twin that incorporates high fidelity instances of

containerized as well as VM-based instances of network OSes, and which thereby enables a smooth transition from a complex simulated network to an actual production network deployment.

System Elements and Concepts

The cornerstone system concepts, components and developments demonstrated include:

- Integrated operations and orchestrated management of resources: interworking with and advancing the site (Site-RM) and network resource managers (Network-RM) and the Orchestrator developed in the SENSE program.
- Fine-grained end-to-end monitoring and data collection, with a focus on the edges and end sites, enabling data analytics-assisted intelligent and automatic decisions driven by applications supported by optimized path selection and load balancing mechanisms driven by machine learning.
- 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.
- An overlay network with PolKA tunnels forming virtual circuits. The overlay will be built by integrating persistent resources from GNA-G AutoGOLE/SENSE and GEANT RARE testbeds. Underlay congestion can be detected by tunnel monitoring and signalized to the overlay so that the traffic is steered from congested tunnels to other paths. Comparisons between segment routing and PolKA regarding controllability and performance metrics are also Planned. This provides offering a new range of network functionality including: multipath routing, in-network telemetry and proof-of-transit with path attributes to support higher level stateful traffic engineering decisions.
- A paragon network at several our partners' sites Composed of P4 programmable devices, Smart NICs and Xilinx FPGA-based network interfaces providing packet-by- packet inspection, agile state tracking, real-time decisions and rapid reaction as needed.
- High throughput platform demonstrations in support of workflows for the science programs mentioned. This will include reference designs of NVMe server systems to match a 400G – 800G network core, as well as servers with multi-GPUs and programmable smart NICs with FPGAs.
- Integration of edge-focused extreme telemetry data (from P4 switches and end hosts) and end facility/application caching stats and other metrics data to facilitate automated decision- making process.
- Development of dynamic regional caches or "data lakes" that treat nearby site facilities as a unified data resource, building on the successful petabyte cache currently in operation between Caltech and UCSD and in ESNet based on the XRootD federated access protocol
- The integration of Qualcomm Technology's GradientGraph (G2) with 5G / OpenALTO / SRv6 to demonstrate how applications can benefit from the intelligent routing, rate limiting, and service placement decisions.

Development Trajectory

The teams' several branches of advanced development, driven by a diverse set of challenging use cases within the GNA-G framework and associated R&D projects, and enabled by the National Research Platform and the Global Research Platform, as well as the GP4L, are embedded and inter- operate within an emerging "composable architecture" of subsystems, components and interfaces organized into several areas:

- **Visibility:** monitoring and information tracking and management including IETF ALTO/OpenALTO, BGP-LS, sFlow/NetFlow, Perfsonar, Qualcomm GradientGraph congestion information, Kubernetes statistics, Prometheus, P4/Inband telemetry
- **Intelligence:** NetPredict, Hecate, RL-Gradient Graph, stateful decisions using composable metrics (policy, priority, network- and site- state, SLA constraints, responses to 'events' at sites and in the networks, ...), Elastiflow/Elastic Stack
- **Controllability:** SENSE/AutoGOLE, SUPA, P4/PINS, segment routing (SRv6, PolKA), BGP/CT, DOCA 3.0
- **Network OSes and Tools:** SONIC, GEANT RARE/freeRtr, Kytos-ng, Calico VPP, Bstruct- Mininet environment, etc.
- **Orchestration**: SENSE, Kubernetes, dedicated code and APIs for interoperation and progressive integration relationship between R&D and Production Operations, with the joint mission to consistently serve and meet the needs of both the major science programs and the at-large academic and research communities, and a sustained commitment to deploy and leverage the latest developments to meet the evolving goals of these communities, while advancing the missions of the network science and engineering communities alike.

A key concept is a new framework of traditional production, advanced production and R&D facilities coexisting side-by-side, to provide an efficient on-ramp for new services at increasing scales moving continually into production use. At SC26 (and beyond) this is exemplified by the Caltech Tier2 facility which includes 8 400G DTNs and 4 X 100G (advanced production) and 4 X 400G (R&D) network connections between the campus and the CENIC PoP in LA, along with the 2 X 100G shared links provided by the Caltech IMSS campus network team.

\

Foundations

The Global Network Advancement Group (GNA-G) and Its Working Groups

These groundbreaking efforts build on the common global infrastructure, aimed at both production network operations and the progressive deployment of new advanced services, that the GNA-G is developing through its unique partnership of national, regional, continental and transoceanic networks, global exchange points and other R&E network providers, as illustrated in Figure 2.



Figure 2 The GNA-G mission and its transoceanic partner consortia spanning the Atlantic and Pacific regions.

The GNA-G organization and architecture shown in Figure 3 illustrates both the synergetic relationship between R&D and Production Operations, the joint mission to consistently serve and meet the needs of both the major science programs and the at-large academic and research communities, and a sustained commitment to deploy and leverage the latest developments to meet the evolving goals of these communities, while advancing the missions of the network science and engineering communities alike.

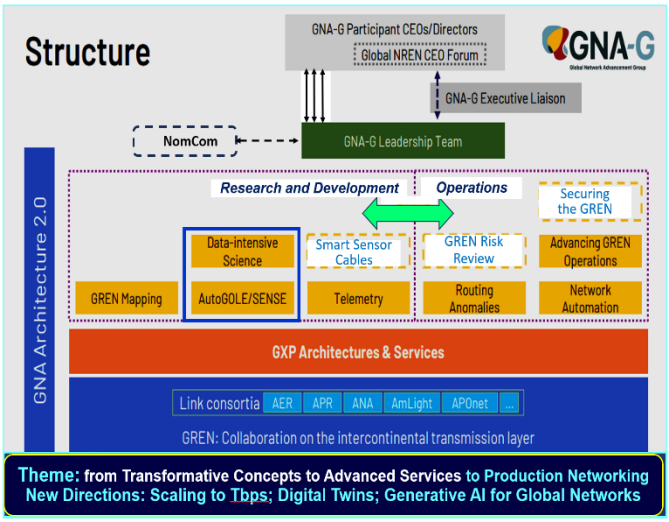


Figure 3 The GNA-G organizational architecture and working groups, emphasizing the Data Intensive Sciences and SENSE/AutoGOLE groups.

Since 2020, the GNA-G and its partners have deployed the Global SENSE/AutoGOLE, FABRIC and GP4L Lab programmable testbeds, and have established a persistent development trajectory harnessing the teams' own advancements in software defined and Terabit/sec network technologies and methods, intelligent global operations.

One important focal area this year, in addition to the science programs mentioned above, is a demonstration showcasing a distributed Cryo-EM workflow connecting the UC Santa Cruz and the University of Cape Town (UCT), enabled by high-performance research and education networks such as AmLight, SANReN, and TENET. By leveraging the advanced FABRIC infrastructure, the National Research Platform's Nautilus, and the Open Science Grid, the project enables remote Cryo-EM data acquisition, preprocessing, and analysis, helping researchers overcome the barriers related to location and computational resources. By embedding computation into the network and enabling transcontinental data transfers of up to 10 TB, this initiative demonstrates how collaborative infrastructure can democratize access to cutting-edge structural biology research.

Resources

This demonstration set, and its many partners, will use:

- 800G switches from Celestica and Arista and 400G switches from Dell in the Caltech booth and in SCinet. This will include multiple 800G, 800G 2 X FR4 and 2 X DR4 transceivers.
- Approximately 12 400G links at SC26 in Chicago between SCinet and the Caltech Booth, including 3-4 400G links to the SCinet cluster. In addition, approximately 2-3 800G links if and when available. Additional links as needed to accommodate the partners' demos will include approximately 6 X 400G links between the SC26 venue and Los Angeles, 2-3 X 400G links from Miami, as well as connections to the Vera Rubin Observatory, Rednesp and the Sao Paulo universities in Latin America as well as sites in South Africa via AmLight Express and Protect, FABRIC, NA-REX 2.0, and the AutoGOLE/SENSE and GP4L global testbeds.
- Gen5 (PCIe 5.0) servers with multiple 400G ConnectX-7 network interfaces in the Caltech booth and the SCinet cluster. This will also include Gen 6 servers, ConnectX-9 800G network interfaces, Bluefield3 and possibly Bluefield4 DPUs, if and when available.
- A key continuing concept this year is a framework and environment encompassing traditional production, advanced production and R&D facilities coexisting side-by-side, to provide an efficient on-ramp for new services at increasing scales moving continually into production use. At SC26 (and beyond) this will be exemplified by the integration of the new SCinet/NRP Fujitsu Monaka Cisco cluster in SCinet representing an advanced site in a global network ensemble. This is also exemplified by the Caltech Tier2 facility which includes 8 400G DTNs and 4 X 100G (advanced production) and 4 X 400G (R&D) network connections between the campus and the CENIC PoP in LA, along with the 2 X 100G shared links provided by the Caltech IMSS campus network team.

Acknowledgements

This work has been partially supported by the US National Science Foundation (NSF) Grants OAC-2030508, OAC-2019012, OAC-1841530, OAC-1836650, OAC-1659403, 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, DE-SC0015528, DE-SC0016585, and FP-00002494. Finally, this work would not be possible without the outstanding 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 and the Global Platform for Labs led by the RNP Brazilian National Network, RENATER and GEANT, and many industry partners, notably Ciena, Supermicro, AMD, Arista, Celestica and Dell.

Partners: Group Leads and Participants, by Team

- **Caltech HEP:** Harvey Newman (newman@hep.caltech.edu), Justas Balcas (jbalcas@caltech.edu), Preeti Bhat, Andres Moya, Sravya Uppalapati, Raimondas Sivinskas
- **Caltech IMSS:** Jin Chang (jin.chang@caltech.edu), Azher Mughal (azher@caltech.edu), Michael Sparacio (sparacio@caltech.edu) , Joel Monaly
- **Ciena:** Scott Kohlert (skohlert@ciena.com), Marc Lyonnais, Rod Wilson
- **UCSD/SDSC/NRP:** Frank Wuerthwein (fkw888@gmail.com), Tom deFanti (tdefanti@eng.ucsd.edu), Larry Smarr, John Graham, Tom Hutton (hutton@ucsd.edu), Diego Davila
- **Fermilab:** Oliver Gutsche (gutsche@fnal.gov), Phil Demar (demar@fnal.gov), Syed Asif Shah
- **ESnet/SENSE:** Inder Monga (imonga@es.net), Chin Guok (chin@es.net), Tom Lehman (tlehman@es.net), John MacAuley, Xi Yang, Justas Balcas (jbalcas@es.net), Dale Carder (dwcarder@es.net)
- **LBNL/NERSC:** Alex Sim (asim@lbl.gov)
- **Oak Ridge National Lab:** Mariam Kiran (kiranm@ornl.gov , Anees Al Najjar
- **KAUST:** Alex Moura (alex.moura@kaust.edu.sa), Kevin Sale (kevin.sale@kaust.edu.sa)
- **Nebraska/UNL:** Ken Bloom (kbloom@unl.edu), Garhan Attenbury (garhan.attebury@unl.edu)
- **Vanderbilt:** Andrew Melo, andrew.m.melo@accre.vanderbilt.edu
- **CERN:** Edoardo Martelli (edoardo.martelli@cern.ch), Carmen Misa (carmen.misa@cern.ch)
- **Qualcomm Gradient Graph:** Jordi Ros-Giralt (jros@qti.qualcomm.com)
- **UFES:** Magnos Martinello (magnos.martinello@ufes.br), Moises R.N. Ribeiro (moises@ele.ufes.br), Cristina K. Dominicini (cristina.dominicini@ifes.edu.br), Everson Borges (everson@ifes.edu.br), Rafael Guimarães
- **RNP:** Marcos Schwarz (marcos.schwarz@rnp.br), Renata Frez, Leandro Ciuffo (leandro.ciuffo@rnp.br)
- **Rednesp:** Joao Eduardo Ferreira **(USP)**, jef@ime.usp.br, Ney Lemke **(UNESP)** ney.lemke@unesp.br, Carlos Antonio Ruggiero **(USP)** toto@ifsc.usp.br, Rogerio Motitsuki, Jorge Marcos de Almeida **(USP)** (jorge@usp.br), Bruno Baldim (bruno@rednesp.br), Paulo Simionato simionato@rednesp.br, Daniel M Batista **(USP)** batista@ime.usp.br, Ricardo Dahab **(UNICAMP)** rdahab@unicamp.br, Geraldo Teodoro (geraldo@rednesp.br, Eduardo A Trettel **(UNICAMP)** trettel@unicamp.br, Sandro Rigo **(UNICAMP)** srigo@unicamp.br
- **Federal University at Sao Carlos (UFSCAR):** Erick L Melo erick@ufscar.br, Fabio R da Silva fabiors@ufscar.br Jose Teixeira S. Junior jtsjunior@ufscar.br, Paulo Matias matias@ufscar.br
- **Universidade Mackenzie:** Calebe P Bianchini Univ Mackenzie calebe@mackenzie.br, Pedro Calindo Junior pedro@mackenzie.br,
- **UNESP (SPRACE NCC UNESP):** Sergio Novaes (Sergio.Novaes@cern.ch), Rogerio Iope (rogerio.iope@unesp.br)
- **UERJ:** Alberto Santoro (Alberto.Santoro@cern.ch), Eduardo Revoredo
- **CENIC:** Louis Fox (lfox@cenic.org), Josh Dickman (jdickman@cenic.org), Gary Kazaryan (gkazaryan@cenic.org), Sana Bellamine, Tony Nguyen
- **RENATER/GEANT/RARE:** Frédéric LOUI (frederic.loui@renater.fr)
- **Pacific Wave/USC:** Celeste Anderson (celestea@usc.edu)
- **Starlight/MREN/iCAIR:** Joe Mambretti (j-mambretti@northwestern.edu), Jim Chen, Fei Yeh
- **Internet2:** Chris Wilkinson (cwilkinson@internet2.edu), Matt Zekauskas (matt@internet2.edu), James Deaton
- **AmLight/FIU:** Julio Ibarra (Julio@fiu.edu), Jerónimo Bezerra, Vasilka Chergarova, Anthony Ulloa, Italo Brito, Vinicius Arcanjo
- **Amlight/ISI:** Heidi Morgan (hlmorgan@isi.edu)
- **TENET:** Paul Sullivan (paul@tenet.ac.za)
- **SANReN, CSIR:** Kasandra Pillay (kpillay@csir.co.za), Johann Hugo (johann@sanren.ac.za), Ajay Makan (ajay@sanren.ac.za)
- **Xiamen:** Qiao Xiang (xiangq27@gmail.com), Chenyang Huang, Ridi Wen, Yuxin Wang, Jiwu Shu
- **KISTI/KREONET:** Buseung Cho bscho@kisti.re.kr, Mazahir Hussain, Changjin Park
- **CANARIE:** Thomas Tam (Thomas.Tam@canarie.ca)
- **DE-KIT:** Bruno Hoeft (bruno.hoeft@kit.edu)
- **Oligo:** Lee, Carlyn-Ann (carlyn@oligo.space)
- **NIST:** Davide Pasavento (davide.pasavento@nist.gov)
- **SURFNet:** Hans Trompert (hans.trompert@surfnet.nl)

- **CESNET:** Michal Hažlinský, (hazlinsky@cesnet.cz)
- **Clemson:** Cole McKnight (cbmckni@g.clemson.edu)
- **NCHC/TAWREN:** Li-Chi Ku, lku@narlabs.org.tw
- **ETH Zürich:** Adrian Perrig (adrian.perrig@inf.ethz.ch), Jonghoon Kwon (jong.kwon@inf.ethz.ch), François Wirz (wirzf@inf.ethz.ch)
- **Magdeburg:** David Hausheer (hausheer@ovgu.de), Marten Gartner (marten.gartner@ovgu.de)
- **Northeastern University:** Edmund Yeh (eyeh@ece.neu.edu), Yuanhao Wu, Volkan Mutlu, Yuezhou Liu
- **Tennessee Tech:** Susmit Shannigrahi (sshannigrahi@tntech.edu), Sankalpa Timilsina
- **UCLA:** Lixia Zhang (lixia@cs.ucla.edu), Jason Cong (cong@cs.ucla.edu), Michael Lo, Sichen Song
- **GNA-G AutoGOLE / SENSE WG Members**: https://www.gna-g.net/join-working-group/autogole-sense
  GNA-G Data Intense Science WG Members: https://www.gna-g.net/join-working-group/data- intensive-science/