



OFCnet: High Performance Research Networks

Moderator: Scott Kohlert, Ciena

Joe Mambretti
ICAIR - Northwestern
University

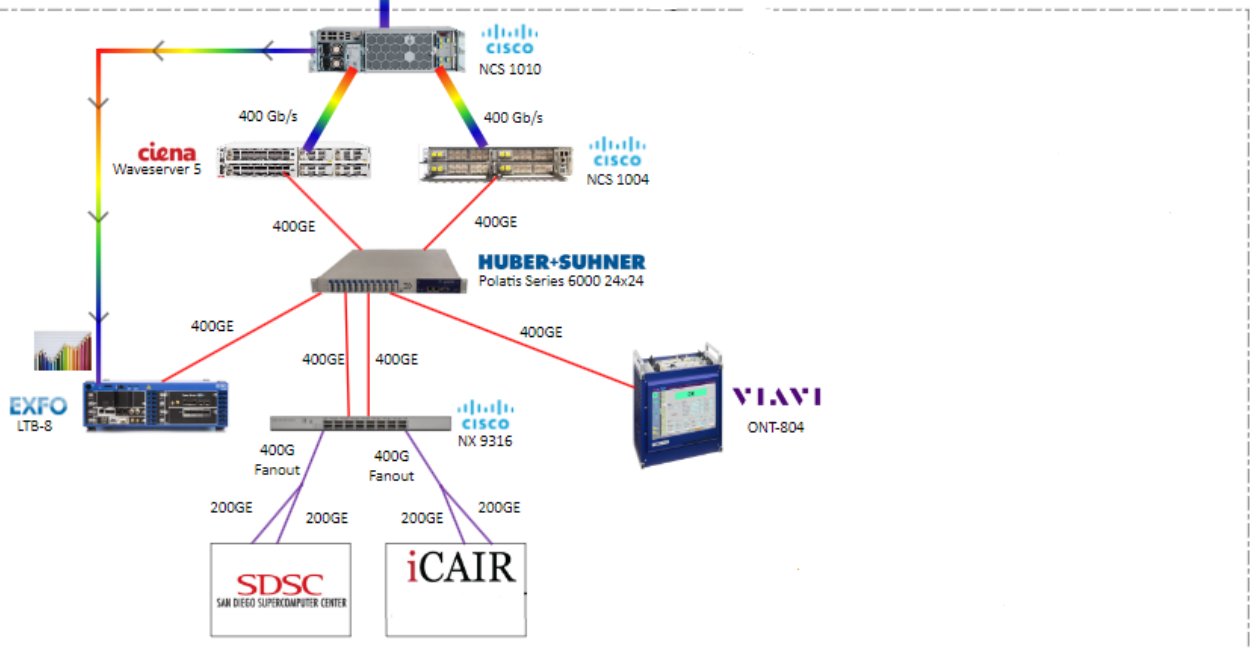

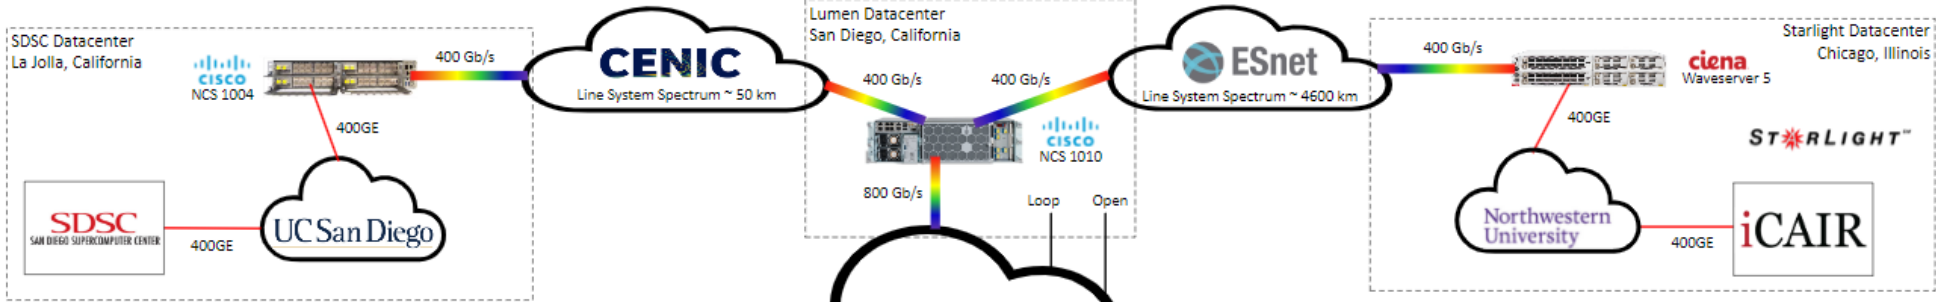
Chris Tracy
Energy Sciences
Network

Thomas Hutton
San Diego
Supercomputer Center

Cees de Laat
University of
Amsterdam

OFCnet Research Network Connectivity

- OFCnet provides 2x 400GE circuits between the show floor and remote research network locations
 - San Diego Supercomputer Center – La Jolla UCSD Campus
 - StarLight – Chicago Northwestern Campus
- A significant undertaking, but we have some experience in doing this.
- Generous loans of dark fiber, spectrum, and equipment from Lumen, ESnet, CENIC, Cisco, Ciena, Huber-Suhner, EXFO, and Viavi helped to make this a reality.



Panel Discussion

- Four members of organizations within the R&E community that are running demonstrations on the OFCnet network.
- First three will discuss aspects of their demonstration.
- Final presenter will discuss the future of R&E involvement in OFCnet.
- Questions at the end of the panel presentations.



Joe Mambretti
ICAIR/Northwestern University

iCAIR/Northwestern/StarLight OFCnet 400 Gbps WAN Demonstrations


Joe Mambretti, Director, (j-mambretti@northwestern.edu)

**International Center for Advanced Internet Research (www.icaair.org)
Northwestern University**

Director, Metropolitan Research and Education Network (www.mren.org)

**Co-Director, StarLight (www.startap.net/starlight), Director, StarLight
International/National Communications Exchange Facility
(www.startap.net/starlight),**

PI IRNC: RXP: StarLight SDX, Co-PI Chameleon, PI-iGENI, PI-OMNINet



**OFC Conference
San Diego, California
March 7-9, 2023**



Selected Applications



GENI
www.geni.net



GLEON
www.gleon.org



USGS EROS
www.usgs.gov/centers/eros



NEON
www.neonscience.org



Open Storage Network
www.openstorage.network.org



OSIRIS
www.osris.org



XSEDE
www.xsede.org



Blue Waters
bluewaters.ncsa.illinois.edu



PRAGMA
www.pragma-grid.net



CENTRA
www.globalcentra.org



OSG
www.opensciencegrid.org



GRP
theglobalresearchplatform.net/



PRP
pacificresearchplatform.org



CHASE-CI
www.calit2.net/newsroom/article.php?id=2910



SAGE2
sage2.sagecommons.org



Polar Geospatial Center
www.pgc.umn.edu




IceCube
icecube.wisc.edu



Chameleon
www.chameleoncloud.org




Jetstream
www.jetstream-cloud.org



Genomic Science Program
genomicscience.energy.gov



LSST
www.lsst.org



Pierre Auger Observatory
www.auger.org



Belle II
www.belle2.org



LBNF/DUNE/ProtoDUNE
lbnf.fnal.gov



ISS
www.nasa.gov/station



SKA
www.skatelescope.org



XENON
xenon.astro.columbia.edu



NOVA
novaexperiment.fnal.gov



Virgo
www.virgo-gw.eu



LIGO
www.ligo.caltech.edu



SDSS
www.sdss.org



ALMA
www.almaobservatory.org



LHC
home.cern/science/accelerators/large-hadron-collider



LHCONE
twiki.cern.ch/twiki/bin/view/LHCONE/WebHome



LHCOPN
twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome



IVOA
www.ivoa.net


Instruments: Exebytes Of Data



High Luminosity LHC



SKA Australia Telescope Facility




Vera Rubin Observatory



KSTAR Korea Superconducting Tokamak



Next Gen Advanced Photon Source



Bioinformatics/Genomics



High Performance E2E WAN Optical Transport of High-capacity Science Research Data at 400Gbps

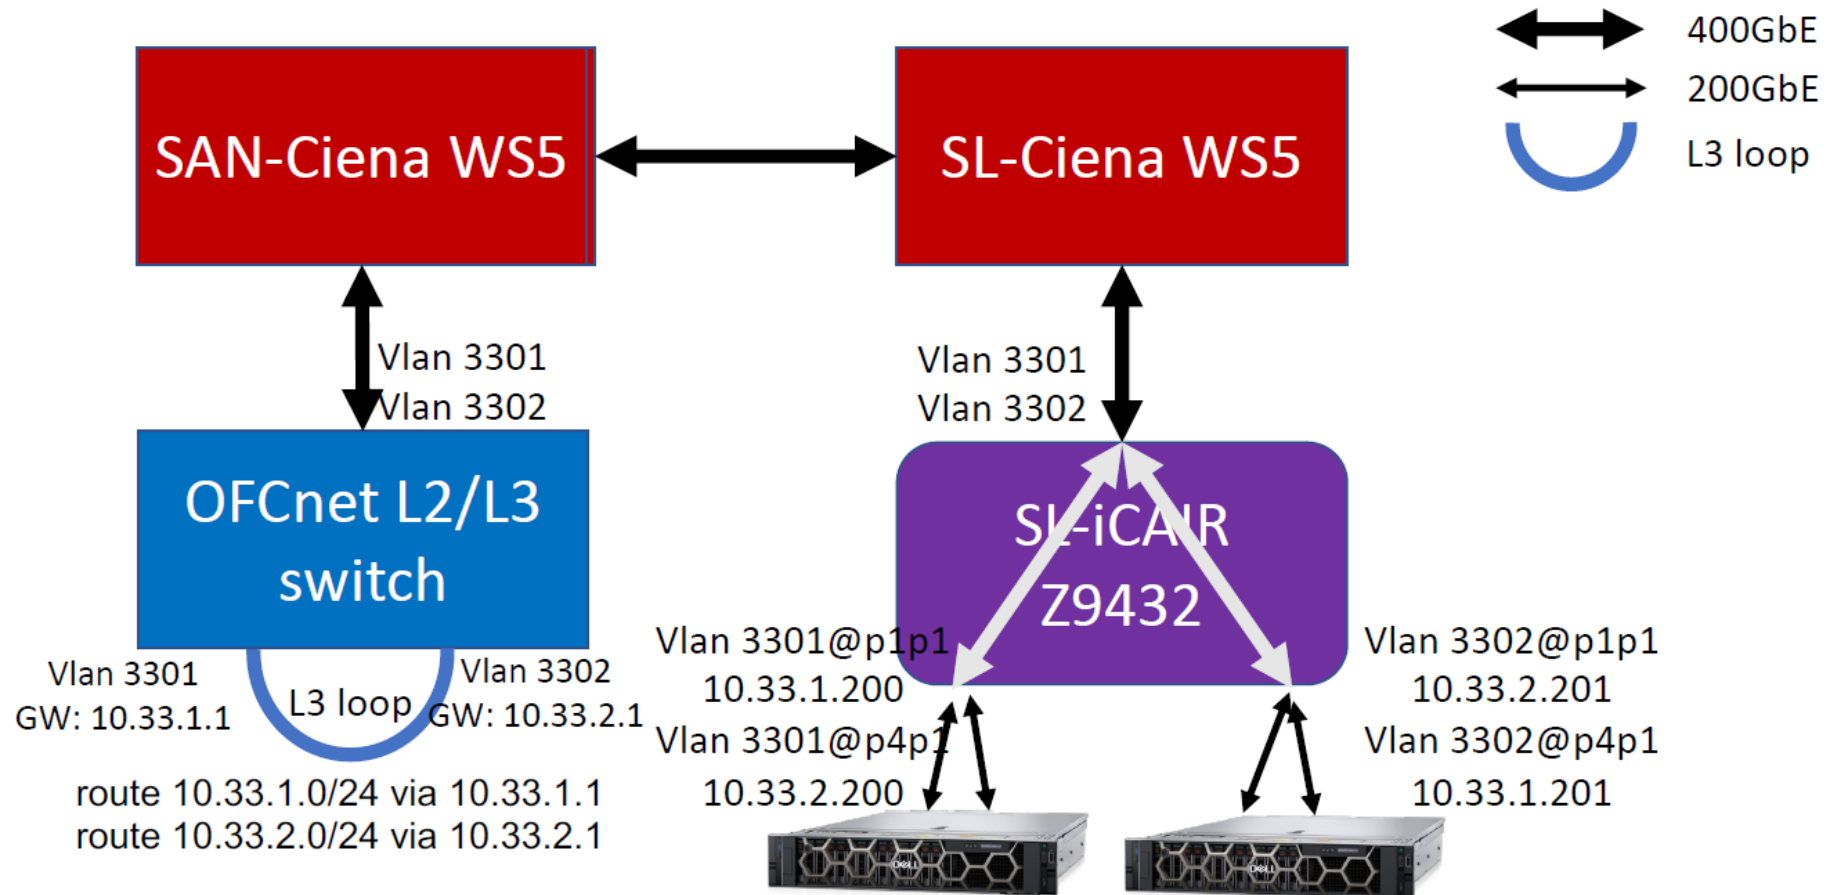
Joe Mambretti
ICAIR/Northwestern
University

Se-Young Yu
ICAIR/Northwestern
University

Fei Yeh
ICAIR/Northwestern
University

Jim Chen
ICAIR/Northwestern
University

StarLight Configuration Test plan A

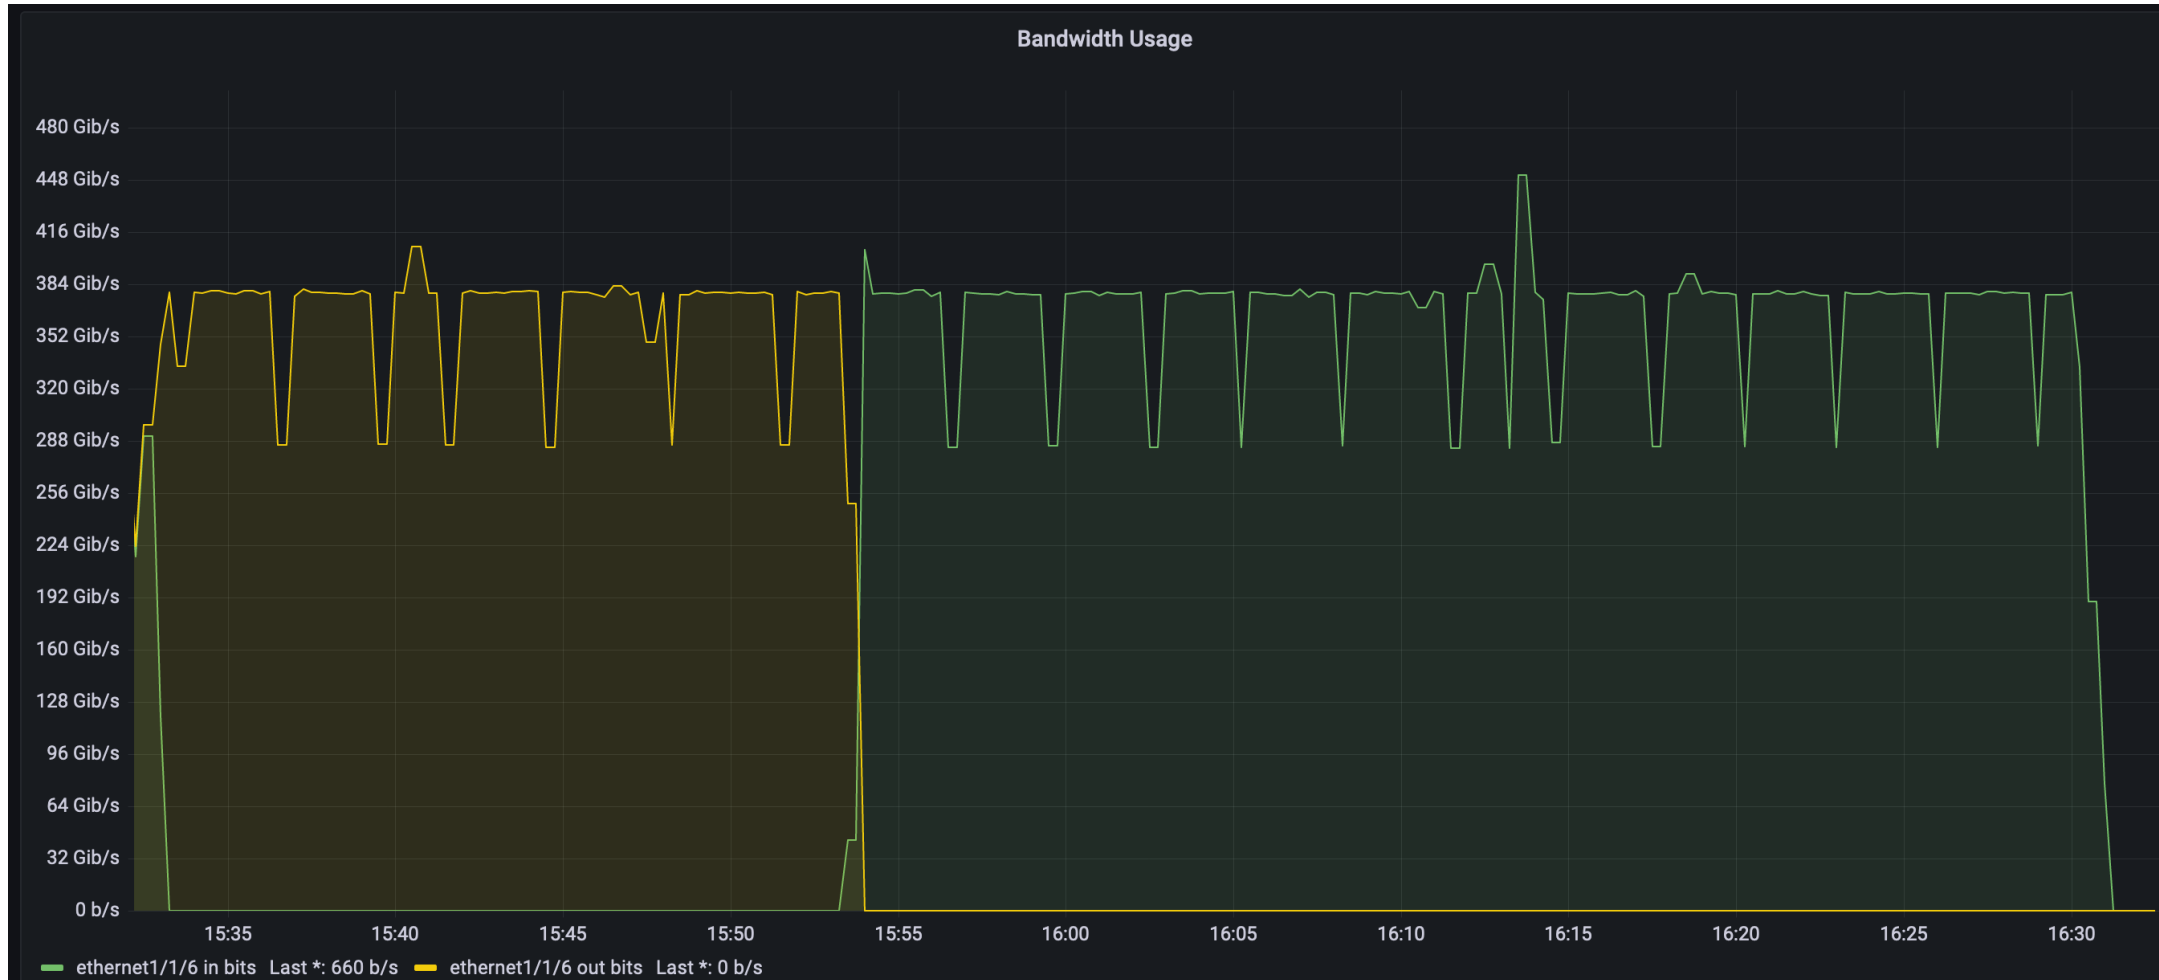


9300 km loop

~ 46 ms RTT

2x Unidirectional
200G Flows

~ 380G+



9300 km loop
~ 46 ms RTT
4x Unidirectional
200G Flows
~ 300G+ Each
Direction





Chris Tracy
ESnet



Flow Monitoring in Research & Education (R&E)

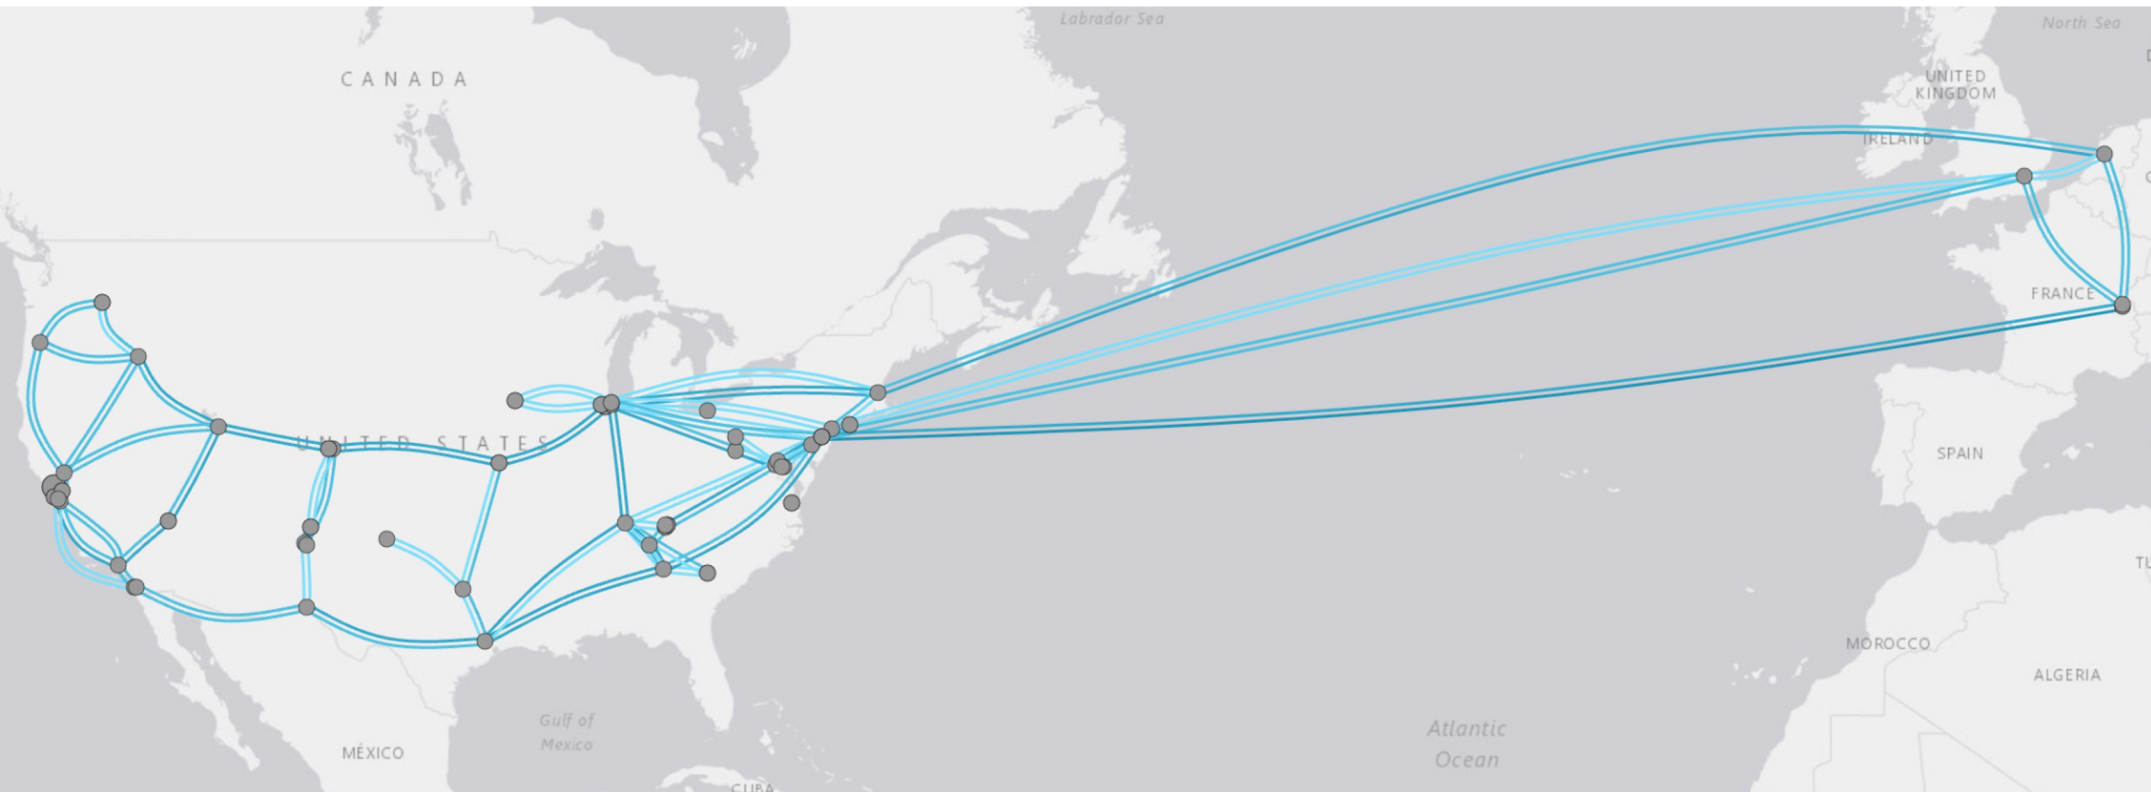
*Chris Tracy
Energy Sciences Network (ESnet)
Planning and Architecture Group
Lawrence Berkeley National Laboratory*

*Special Acknowledgements:
Andy Lake, Eli Dart, and the
development teams for both
Stardust and High Touch*

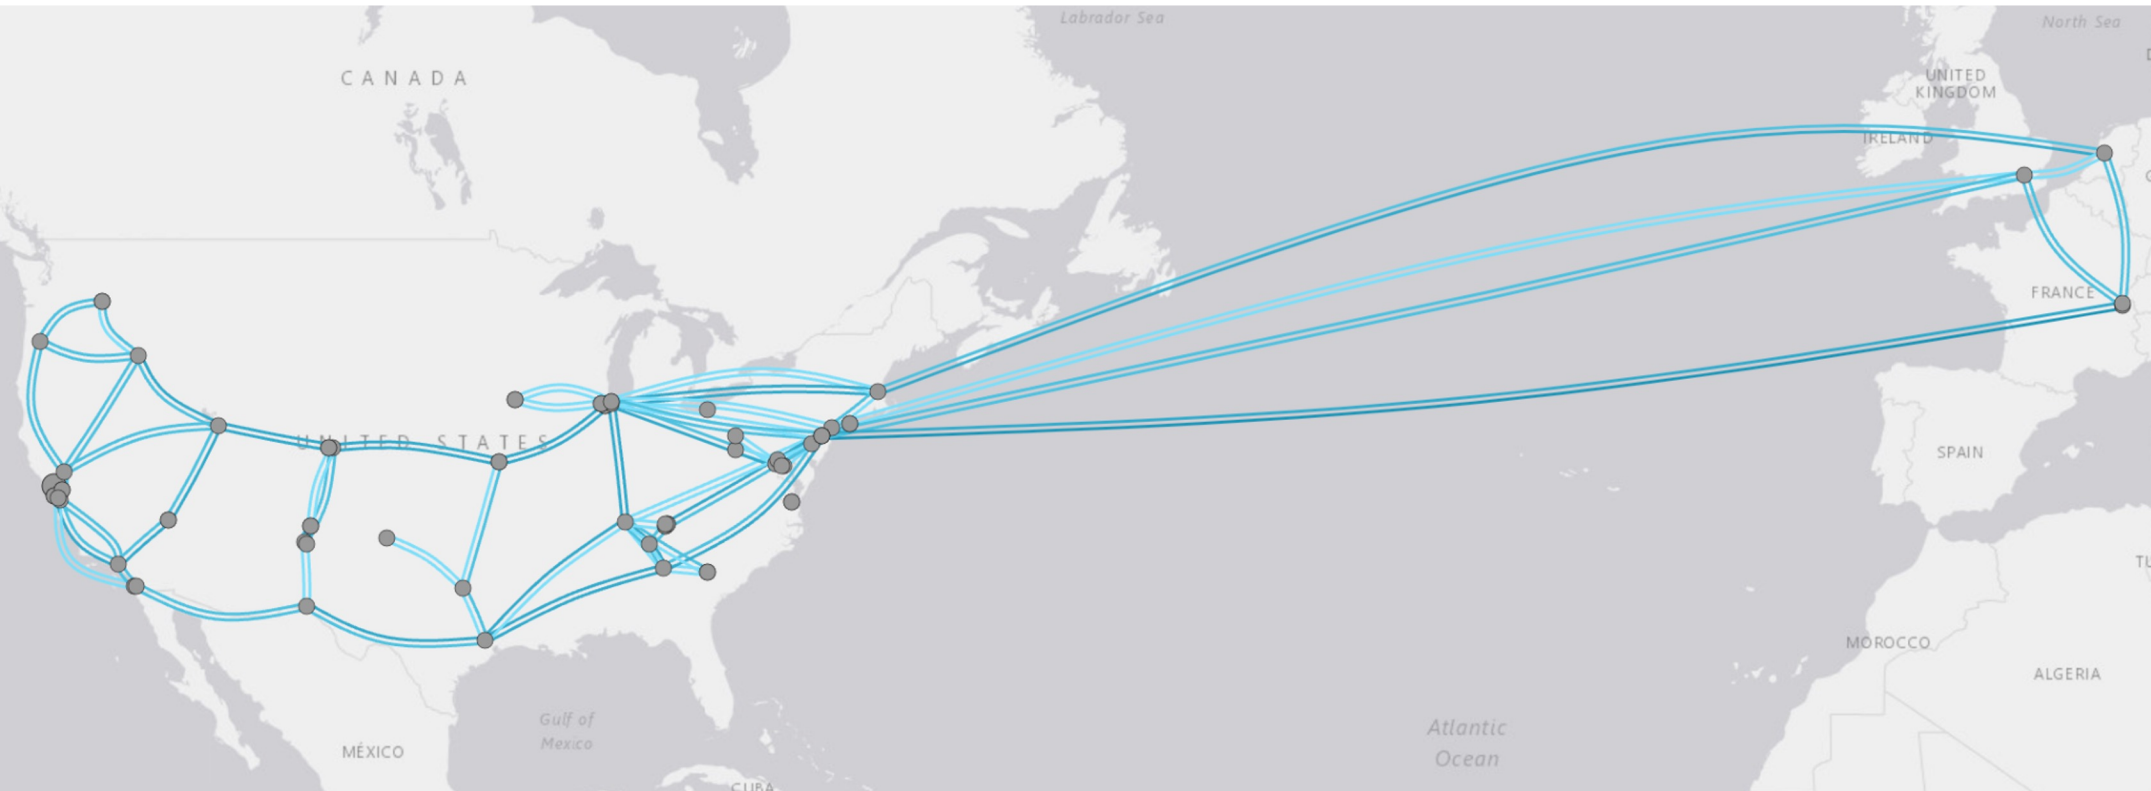
ESnet Background

- ESnet is a special-purpose mission network, funded by the US Congress to support scientific goals of the Department of Energy.
- We see networking as a means to an end: scientific productivity.
- We aim to create a world in which discovery is unconstrained by geography.

This is not an ISP.



It's a DOE user facility designed to overcome the constraints of geography.



We do this by offering unique capabilities, and optimizing the facility for data acquisition, data placement, data sharing, data mobility.

Picture worth 1,000 words

Goal is to present data in ways that allows us to visualize and contextualize related measurements together

Not This



This



```
> Oct 12, 2021 @ 15:09:12.794
@timestamp: Oct 12, 2021 @ 15:09:12.794
@version: 1
@version.keyword: 1
end: Oct 12, 2021 @ 15:08:35.453
meta.app_name: https
meta.app_name_detail: https:443/udp
meta.app_name_detail.keyword: https:443/udp
meta.app_name.keyword: https
meta.app_port: 443
meta.as_names: 16:LBL, 7922:CONCAST-7922
meta.as_names_str: 16:LBL, 7922:CONCAST-7922
meta.as_names.keyword: 16:LBL, 7922:CONCAST-7922
meta.bgp.as_hop#: 6, 939
meta.bgp.as_hop_name: 6939:HURRICANE
meta.bgp.as_hop_name.keyword: 6939:HURRICANE
meta.bgp.as_hop_padding: 1
meta.bgp.as_hop#: 7, 922

> Oct 12, 2021 @ 15:09:12.794
@timestamp: Oct 12, 2021 @ 15:09:12.794
@version: 1
@version.keyword: 1
end: Oct 12, 2021 @ 15:08:35.463
meta.app_name: https
meta.app_name_detail: https:443/tcp
meta.app_name_detail.keyword: https:443/tcp
meta.app_name.keyword: https
meta.app_port: 443
meta.as_names: 45:LLL-TIS-AS, 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.as_names_str: 45:LLL-TIS-AS, 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.as_names.keyword: 45:LLL-TIS-AS, 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.bgp.as_hop#: 8, 675
meta.bgp.as_hop_name: 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.bgp.as_hop_name.keyword: 8675:MICROSOFT-CORP-MSN-AS-BLOCK


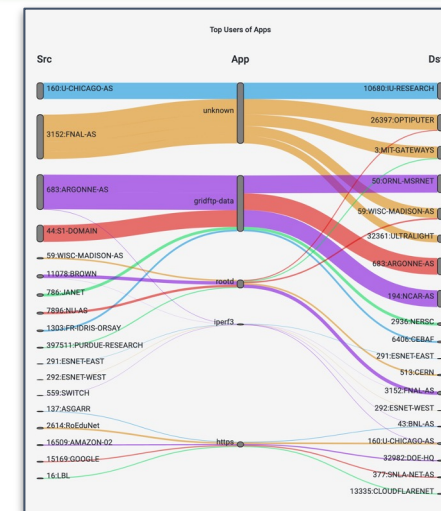
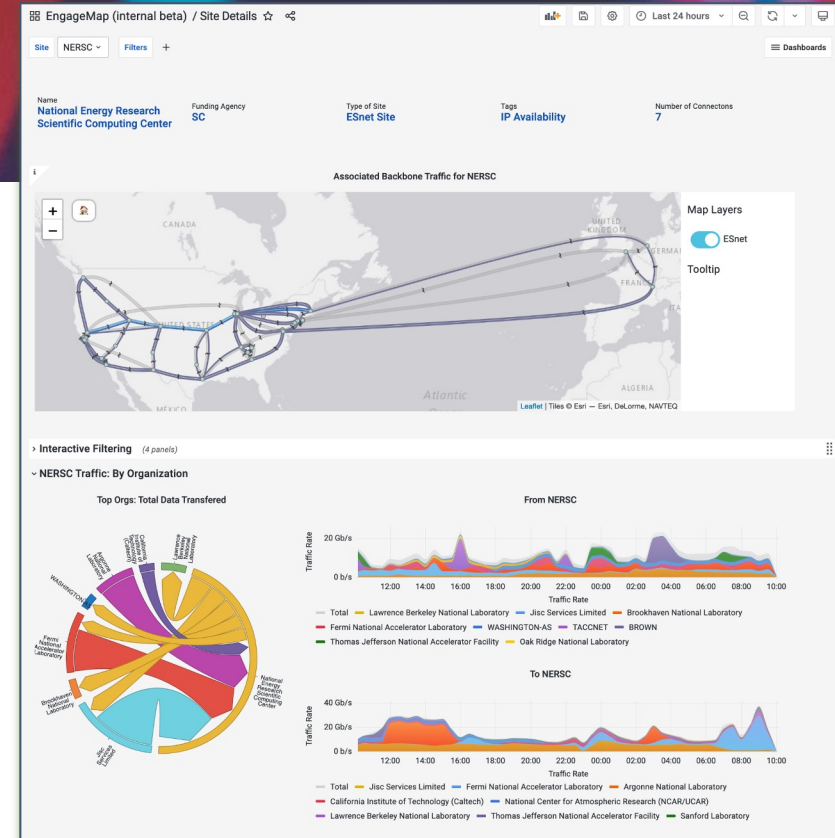
> Oct 12, 2021 @ 15:09:12.793
@timestamp: Oct 12, 2021 @ 15:09:12.793
@version: 1
@version.keyword: 1
end: Oct 12, 2021 @ 15:08:35.373
meta.app_name: https
meta.app_name_detail: https:443/tcp
meta.app_name_detail.keyword: https:443/tcp
meta.app_name.keyword: https
meta.app_port: 443
meta.as_names: 43:BNL-AS, 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.as_names_str: 43:BNL-AS, 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.as_names.keyword: 43:BNL-AS, 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.bgp.as_hop#: 8, 675
meta.bgp.as_hop_name: 8675:MICROSOFT-CORP-MSN-AS-BLOCK
meta.bgp.as_hop_name.keyword: 8675:MICROSOFT-CORP-MSN-AS-BLOCK

> Oct 12, 2021 @ 15:09:12.792
@timestamp: Oct 12, 2021 @ 15:09:12.792
@version: 1
@version.keyword: 1
end: Oct 12, 2021 @ 15:08:35.343
meta.app_name: domain
meta.app_name_detail: domain:53/tcp
meta.app_name_detail.keyword: domain:53/tcp
meta.app_name.keyword: domain
meta.app_port: 53
meta.as_names: 19115:CHARTER-19115-DC, 3443:ESNET-AS
meta.as_names_str: 19115:CHARTER-19115-DC, 3443:ESNET-AS
meta.as_names.keyword: 19115:CHARTER-19115-DC, 3443:ESNET-AS
meta.bgp.as_hop#: 65, 292
meta.bgp.as_hop_name: 65292:AS65292
meta.bgp.as_hop_name.keyword: 65292:AS65292
meta.bgp.as_hop_padding: 1
meta.bgp.as_hop#: 3, 443

> Oct 12, 2021 @ 15:09:12.789
@timestamp: Oct 12, 2021 @ 15:09:12.789
@version: 1
@version.keyword: 1
end: Oct 12, 2021 @ 15:08:35.879
meta.app_name: gnutel
meta.app_name_detail: gnutel:2886/tcp
meta.app_name_detail.keyword: gnutel:2886/tcp
meta.app_name.keyword: gnutel
meta.app_port: 2, 886
meta.as_names: 598722:AS598722, 291:ESNET-EAST
meta.as_names_str: 598722:AS598722, 291:ESNET-EAST
meta.as_names.keyword: 598722:AS598722, 291:ESNET-EAST
meta.bgp.as_hop#: 291
meta.bgp.as_hop_name: 291:ESNET-EAST
meta.bgp.as_hop_name.keyword: 291:ESNET-EAST
meta.bgp.as_hop_padding: 1
meta.bgp.as_hop#: 291

> Oct 12, 2021 @ 15:09:12.788
@timestamp: Oct 12, 2021 @ 15:09:12.788
@version: 1
@version.keyword: 1
end: Oct 12, 2021 @ 15:08:35.131
meta.app_name: psync-https
meta.app_name_detail: psync-https:8443/tcp
meta.app_name_detail.keyword: psync-https:8443/tcp
meta.app_name.keyword: psync-https
meta.app_port: 8, 443
meta.as_names: 45:LLL-TIS-AS, 7922:CONCAST-7922
meta.as_names_str: 45:LLL-TIS-AS, 7922:CONCAST-7922
meta.as_names.keyword: 45:LLL-TIS-AS, 7922:CONCAST-7922
meta.bgp.as_hop#: 6, 939
meta.bgp.as_hop_name: 6939:HURRICANE
meta.bgp.as_hop_name.keyword: 6939:HURRICANE
meta.bgp.as_hop_padding: 1
meta.bgp.as_hop#: 7, 922

> Oct 12, 2021 @ 15:09:12.787
@timestamp: Oct 12, 2021 @ 15:09:12.787
@version: 1
@version.keyword: 1
end: Oct 12, 2021 @ 15:08:35.861
meta.app_name: https
meta.app_name_detail: https:443/udp
meta.app_name_detail.keyword: https:443/udp
meta.app_name.keyword: https
meta.app_port: 443
meta.as_names: 4181:TDS-AS, 293:ESNET
meta.as_names_str: 4181:TDS-AS, 293:ESNET
meta.as_names.keyword: 4181:TDS-AS, 293:ESNET
meta.bgp.as_hop#: 45, 848
meta.bgp.as_hop_name: 45848:AS6948
meta.bgp.as_hop_name.keyword: 45848:AS6948
meta.bgp.as_hop_padding: 1
meta.bgp.as_hop#: 45, 848
meta.bgp.as_hop_padding_len: 1
```



Powering external websites

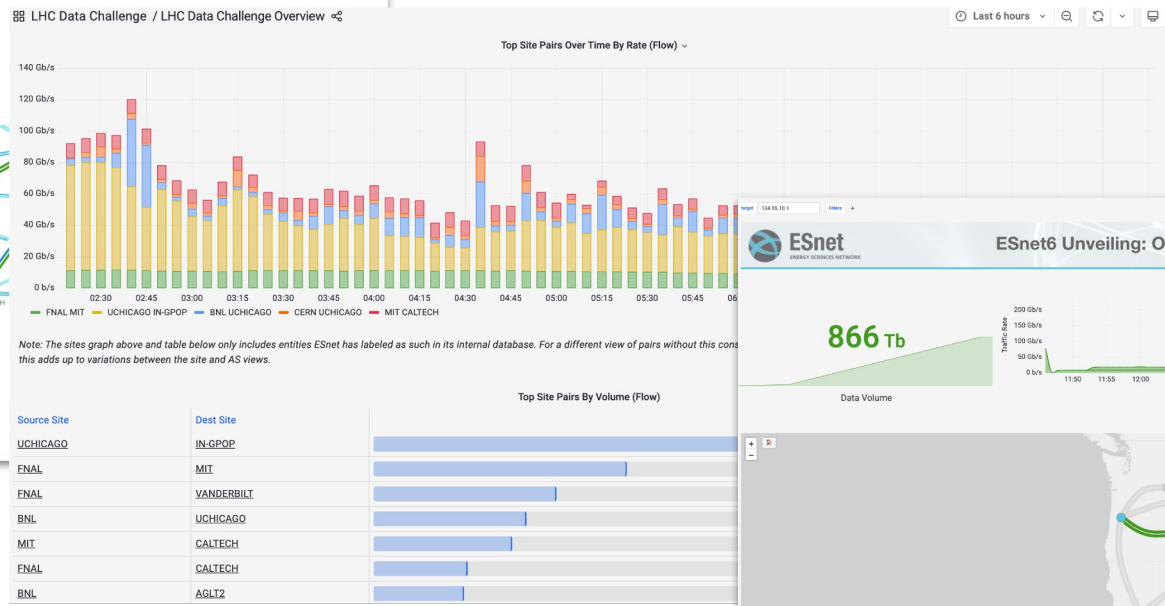
ESnet
ENERGY SCIENCES NETWORK

Customer Service Portal Login


About ESnet

ESnet is a high-performance, unclassified network built to support scientific research. Funded by the U.S. Department of Energy's Office of Science (SC) and managed by Lawrence Berkeley National Laboratory, ESnet provides services to more than 50 DOE research sites, including the entire National Laboratory system, its supercomputing facilities, and its major scientific instruments.

my.es.net Portal



Data Challenges



Live ESnet6 Unveiling Demo



Stardust

Network Measurement and Analysis for ESnet

Extensible / Open Architecture

NSF NetSAGE project derived Approach

- Integrate where we can, innovate where it makes a difference.
 - ◆ metadata and viz
- Loose coupling to avoid lock in

Authenticated access methods for many user groups.

Dashboards, Indexed APIs and “Raw”

- Grafana user editable dashboards
- Elasticsearch Query API access
- Kafka feeds
- RBAC with 2 Factor

Multi Datasource

Low and High Cardinality

- Network Traffic Flows
- Interface Usage
- Optical Line System Performance
- perfSONAR
- High Touch measurements
- LHC Firefly measurements

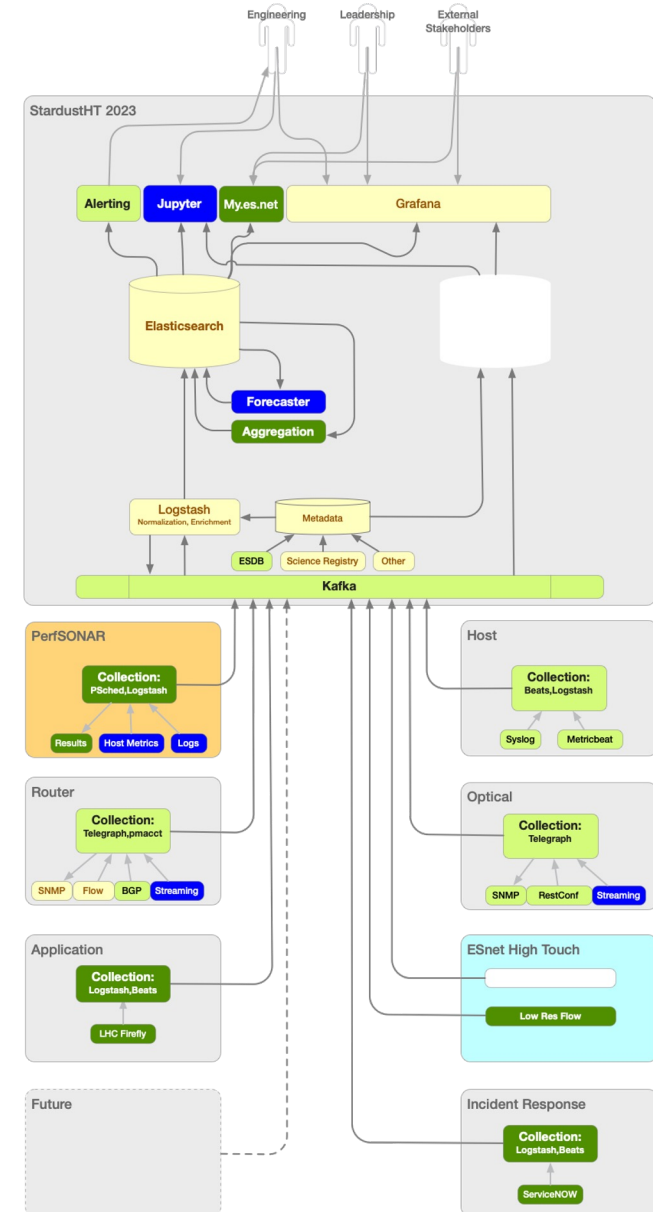
Flexible aggregation

Variable time buckets and dimension reduction

- Summarize in time
 - ◆ Hourly summary
- Summarize by dimension
 - ◆ All NERSC traffic
- Variable retention
 - ◆ hourly data for a year
 - ◆ 30 sec data for 90 days

Stardust Design Evolution

NetSAGE derived Component	2016
Stardust Version 1 Component	2021
Stardust Version 2 Component	2022
StardustHT Component	2023



Introducing High Touch - Line rate packet telemetry with FPGAs

FPGA (AMD Xilinx
Alveo U280)

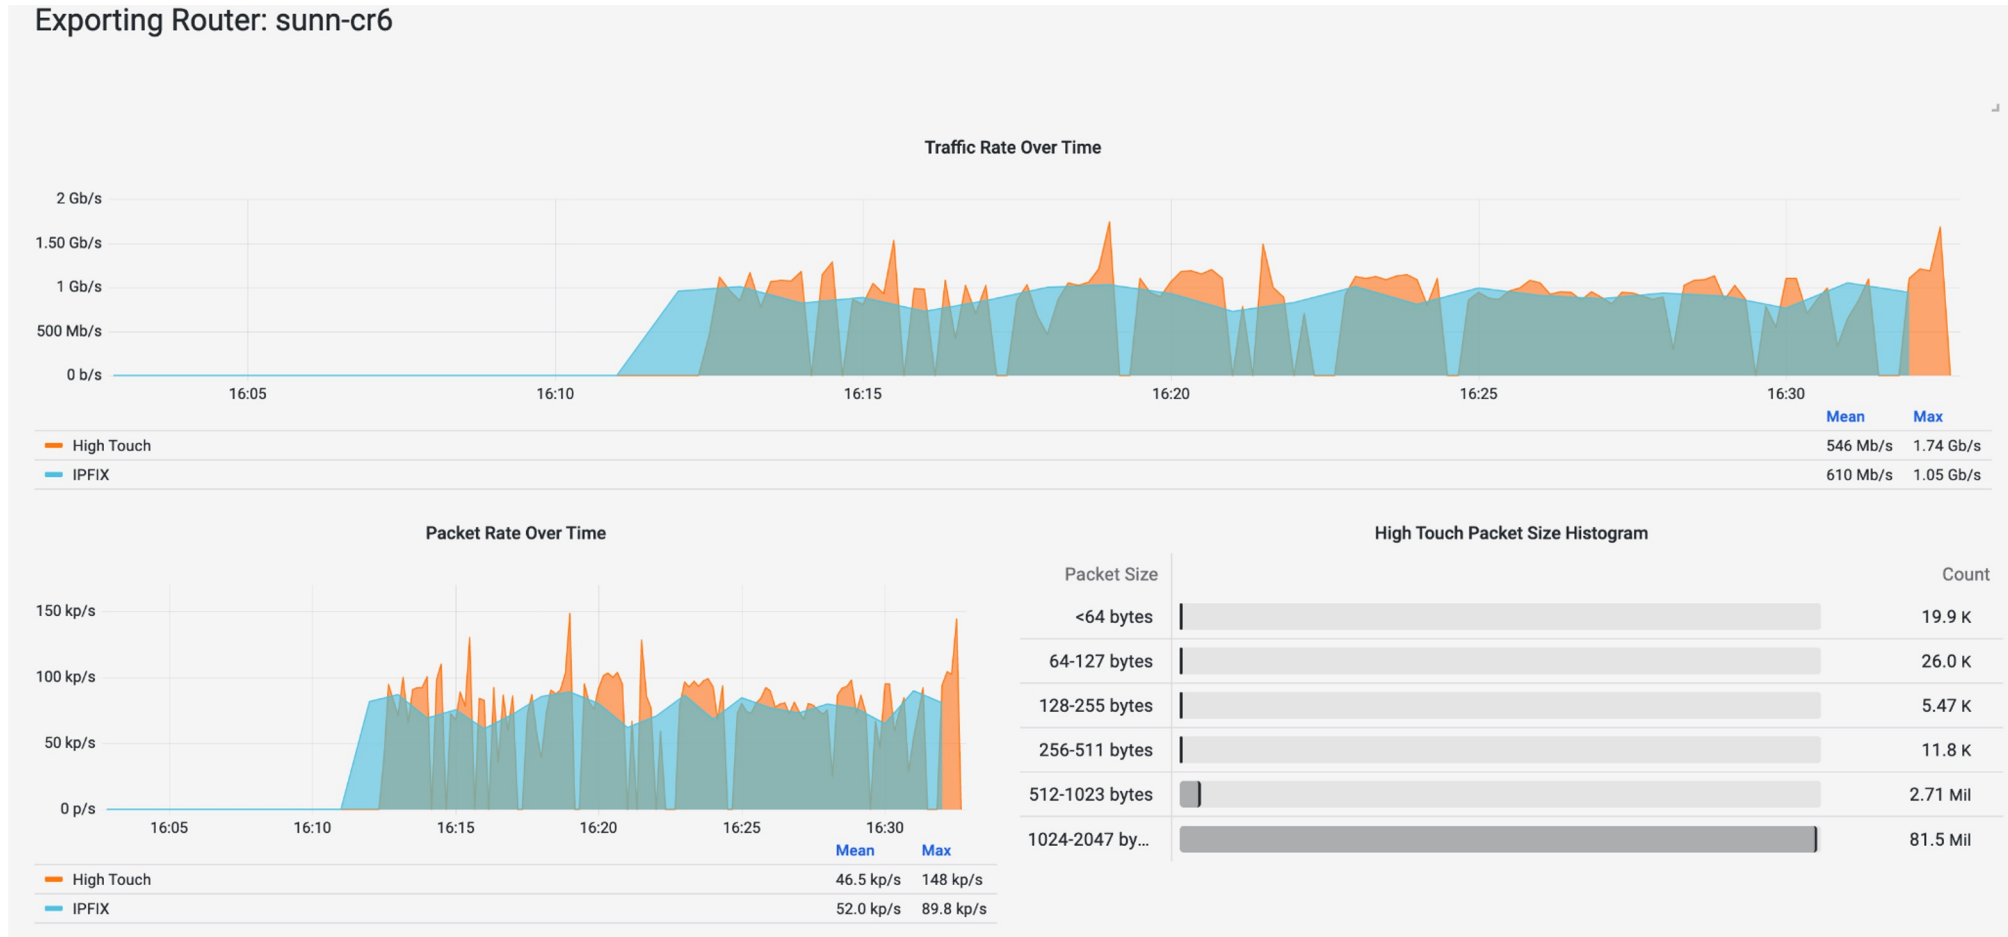
High Touch Servers

Router



See our live OFCnet demo in booth #6109 to learn more or ask questions!

High Touch provides very fine-grained visibility into individual flows



See our live OFCnet demo in booth #6109 to learn more or ask questions!



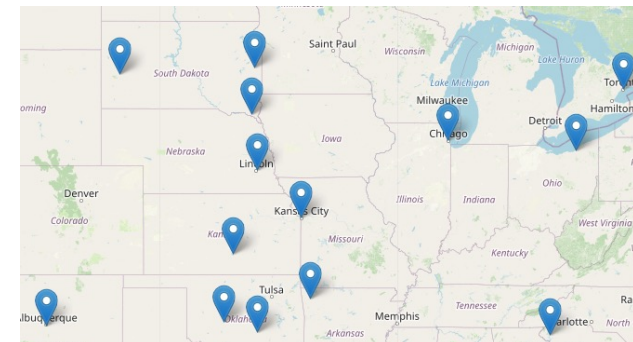
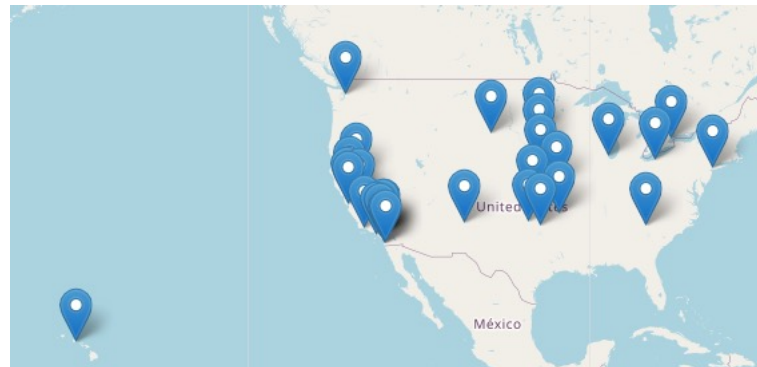
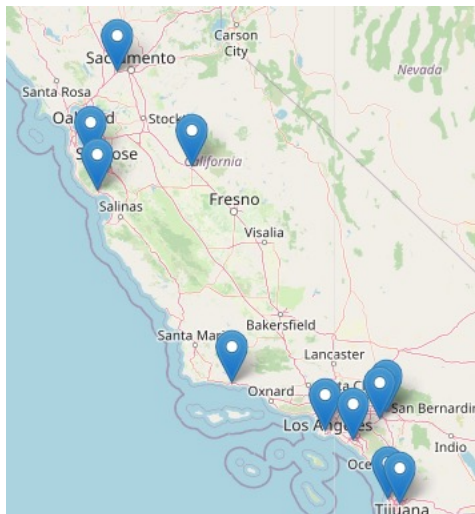
Tom Hutton
SDSC



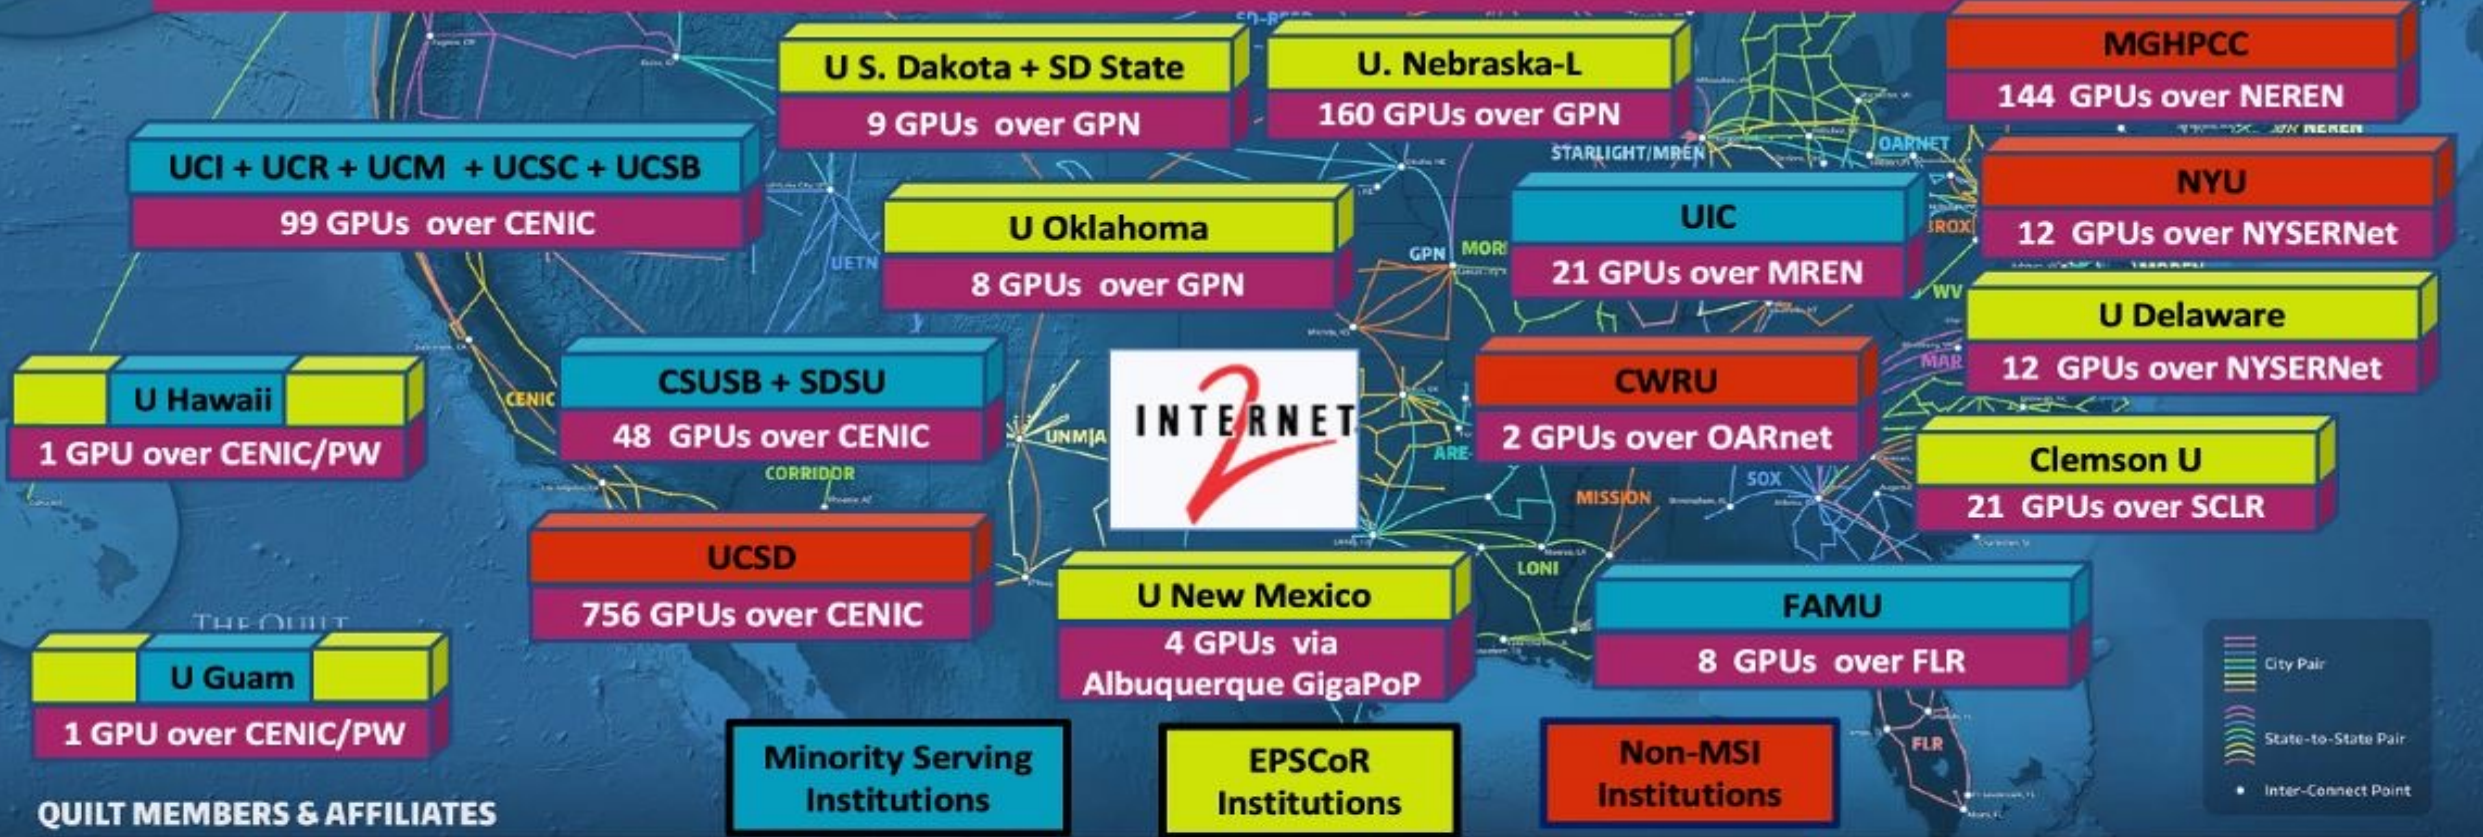
National Research Platform
an Open National Cyberinfrastructure

*Thomas Hutton
San Diego Supercomputer Center
University of California, San Diego
hutton@sdsc.edu*

Hardware on NRP is Global



Nautilus' >19,000 CPU Cores and ~1,300 GPUs
PRP/BYOR Partnerships Provided GPUs and CPUs 2017-22
PNRP added 288 A10s at UN-L and MGHPCC and 72 A100s at UCSD in 2022/23



QUILT MEMBERS & AFFILIATES




NRP Cyberinfrastructure Stack

NRP operates at all layers of the stack, from IPMI up

- IPMI reduces TCO and lower threshold to entry
- Kubernetes allows service deployments
 - Also the natural layer for application container deployment
- Admiralty allows K8S federation with folks who want control
 - Including cloud integration to access TPUs & other cloud only architectures
- HTCondor allows NRP to show up as a “site” in OSG

The layer you integrate at depends on

- Control you want
- Effort you can afford



NRP is both Vertically and Horizontal Scalable



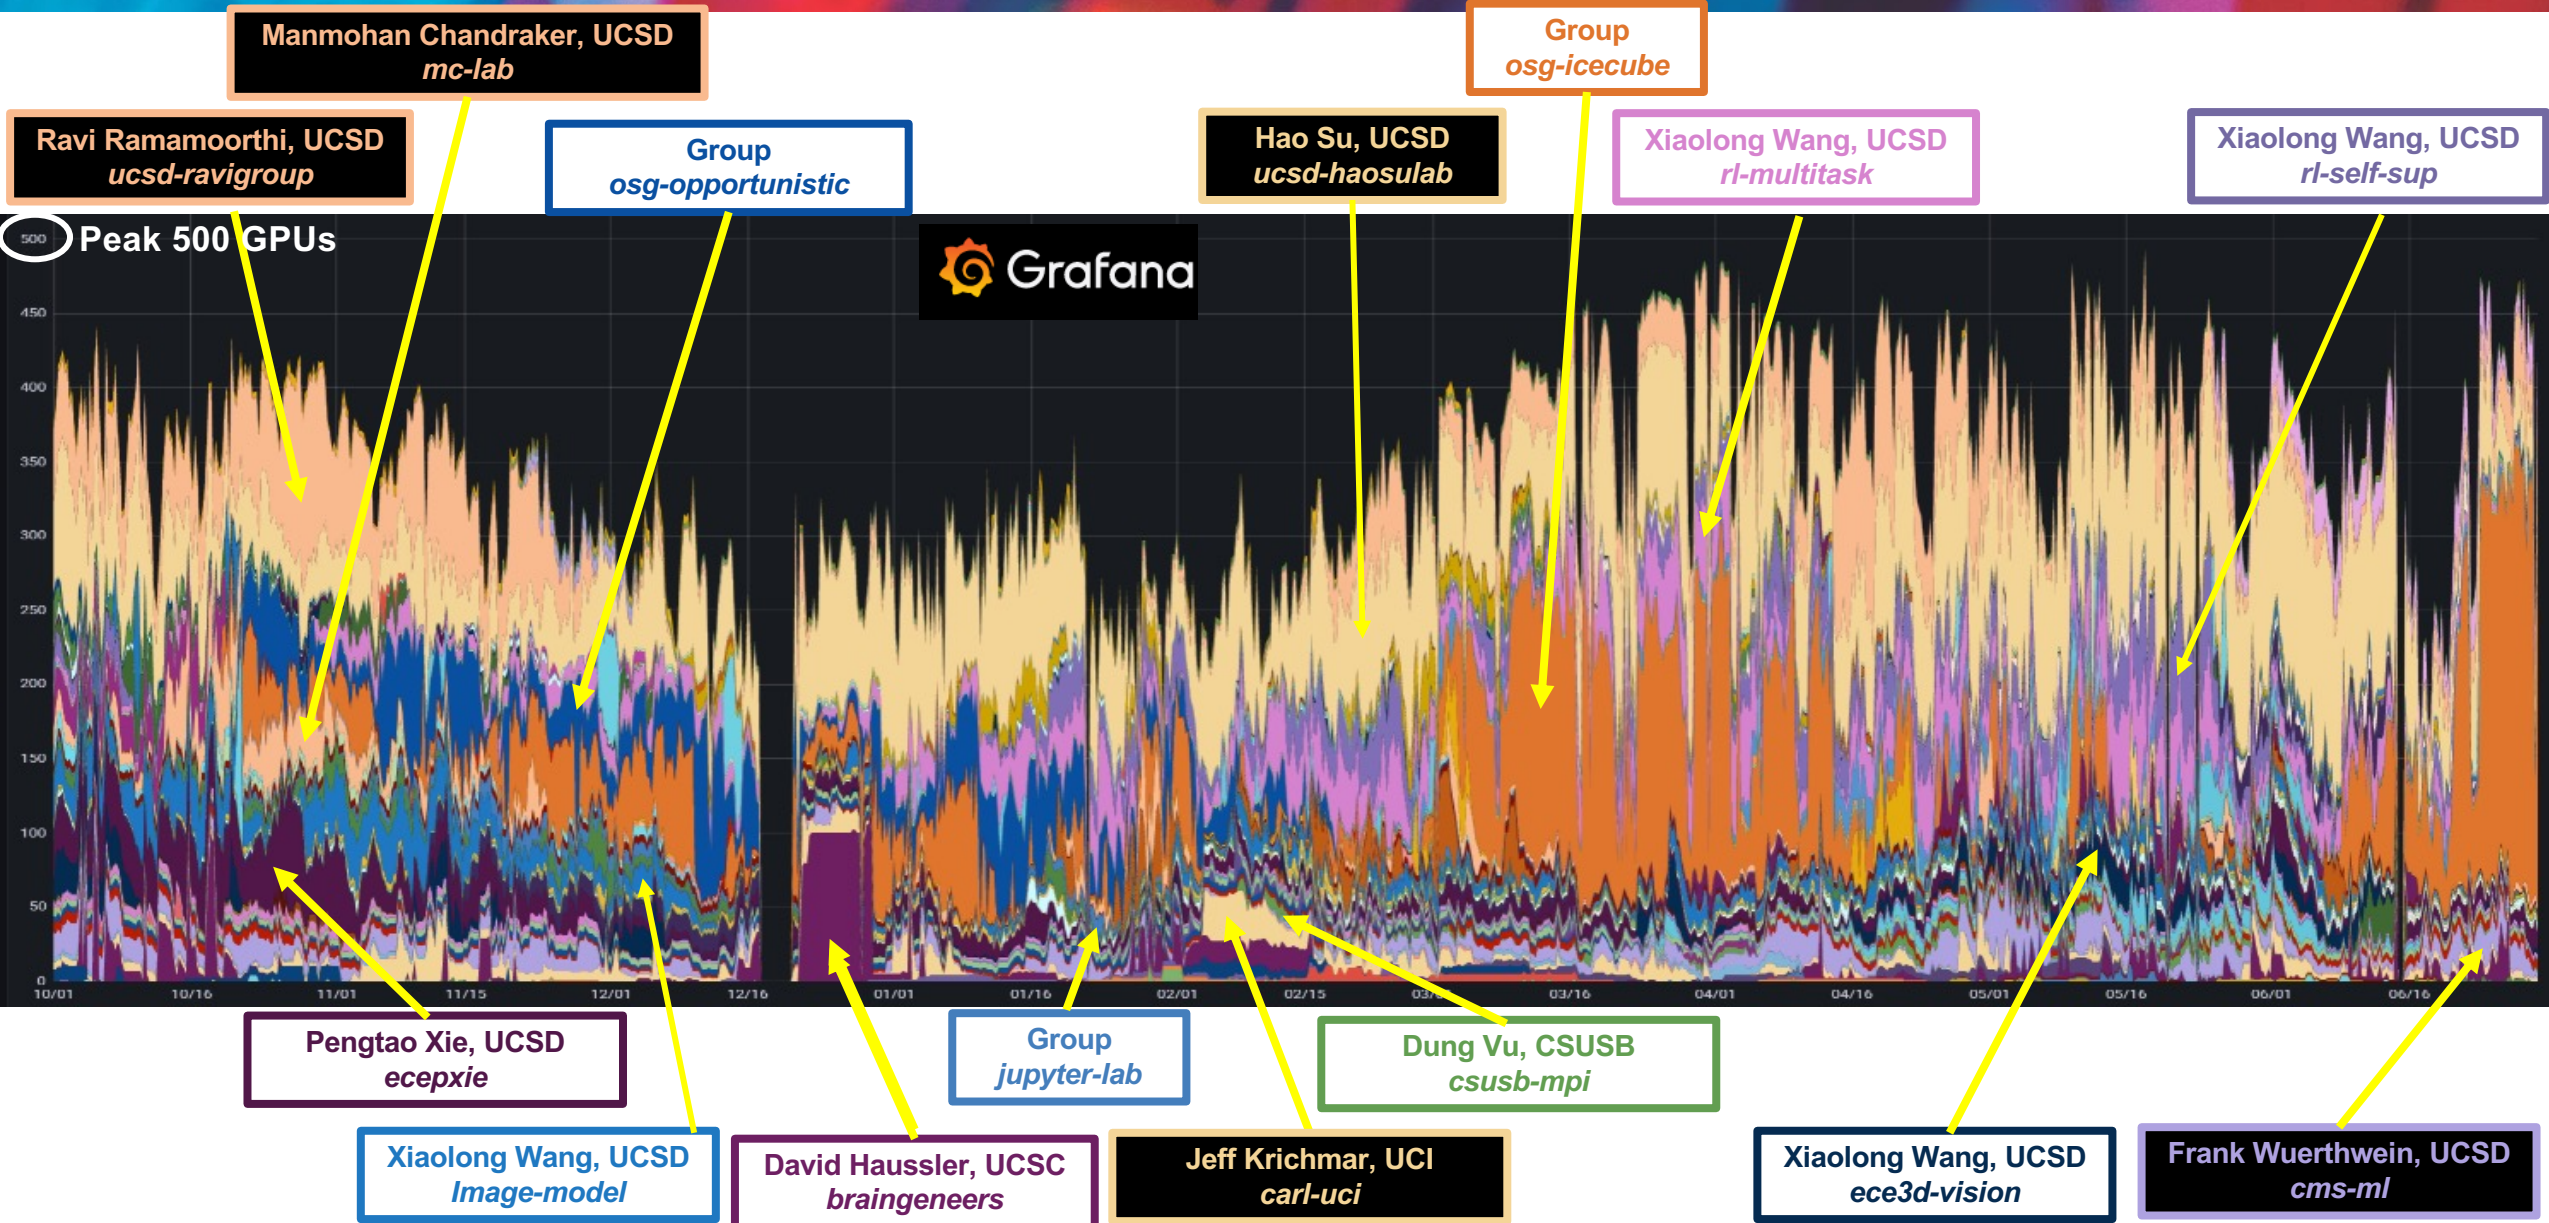
Supporting Nautilus for the next decade

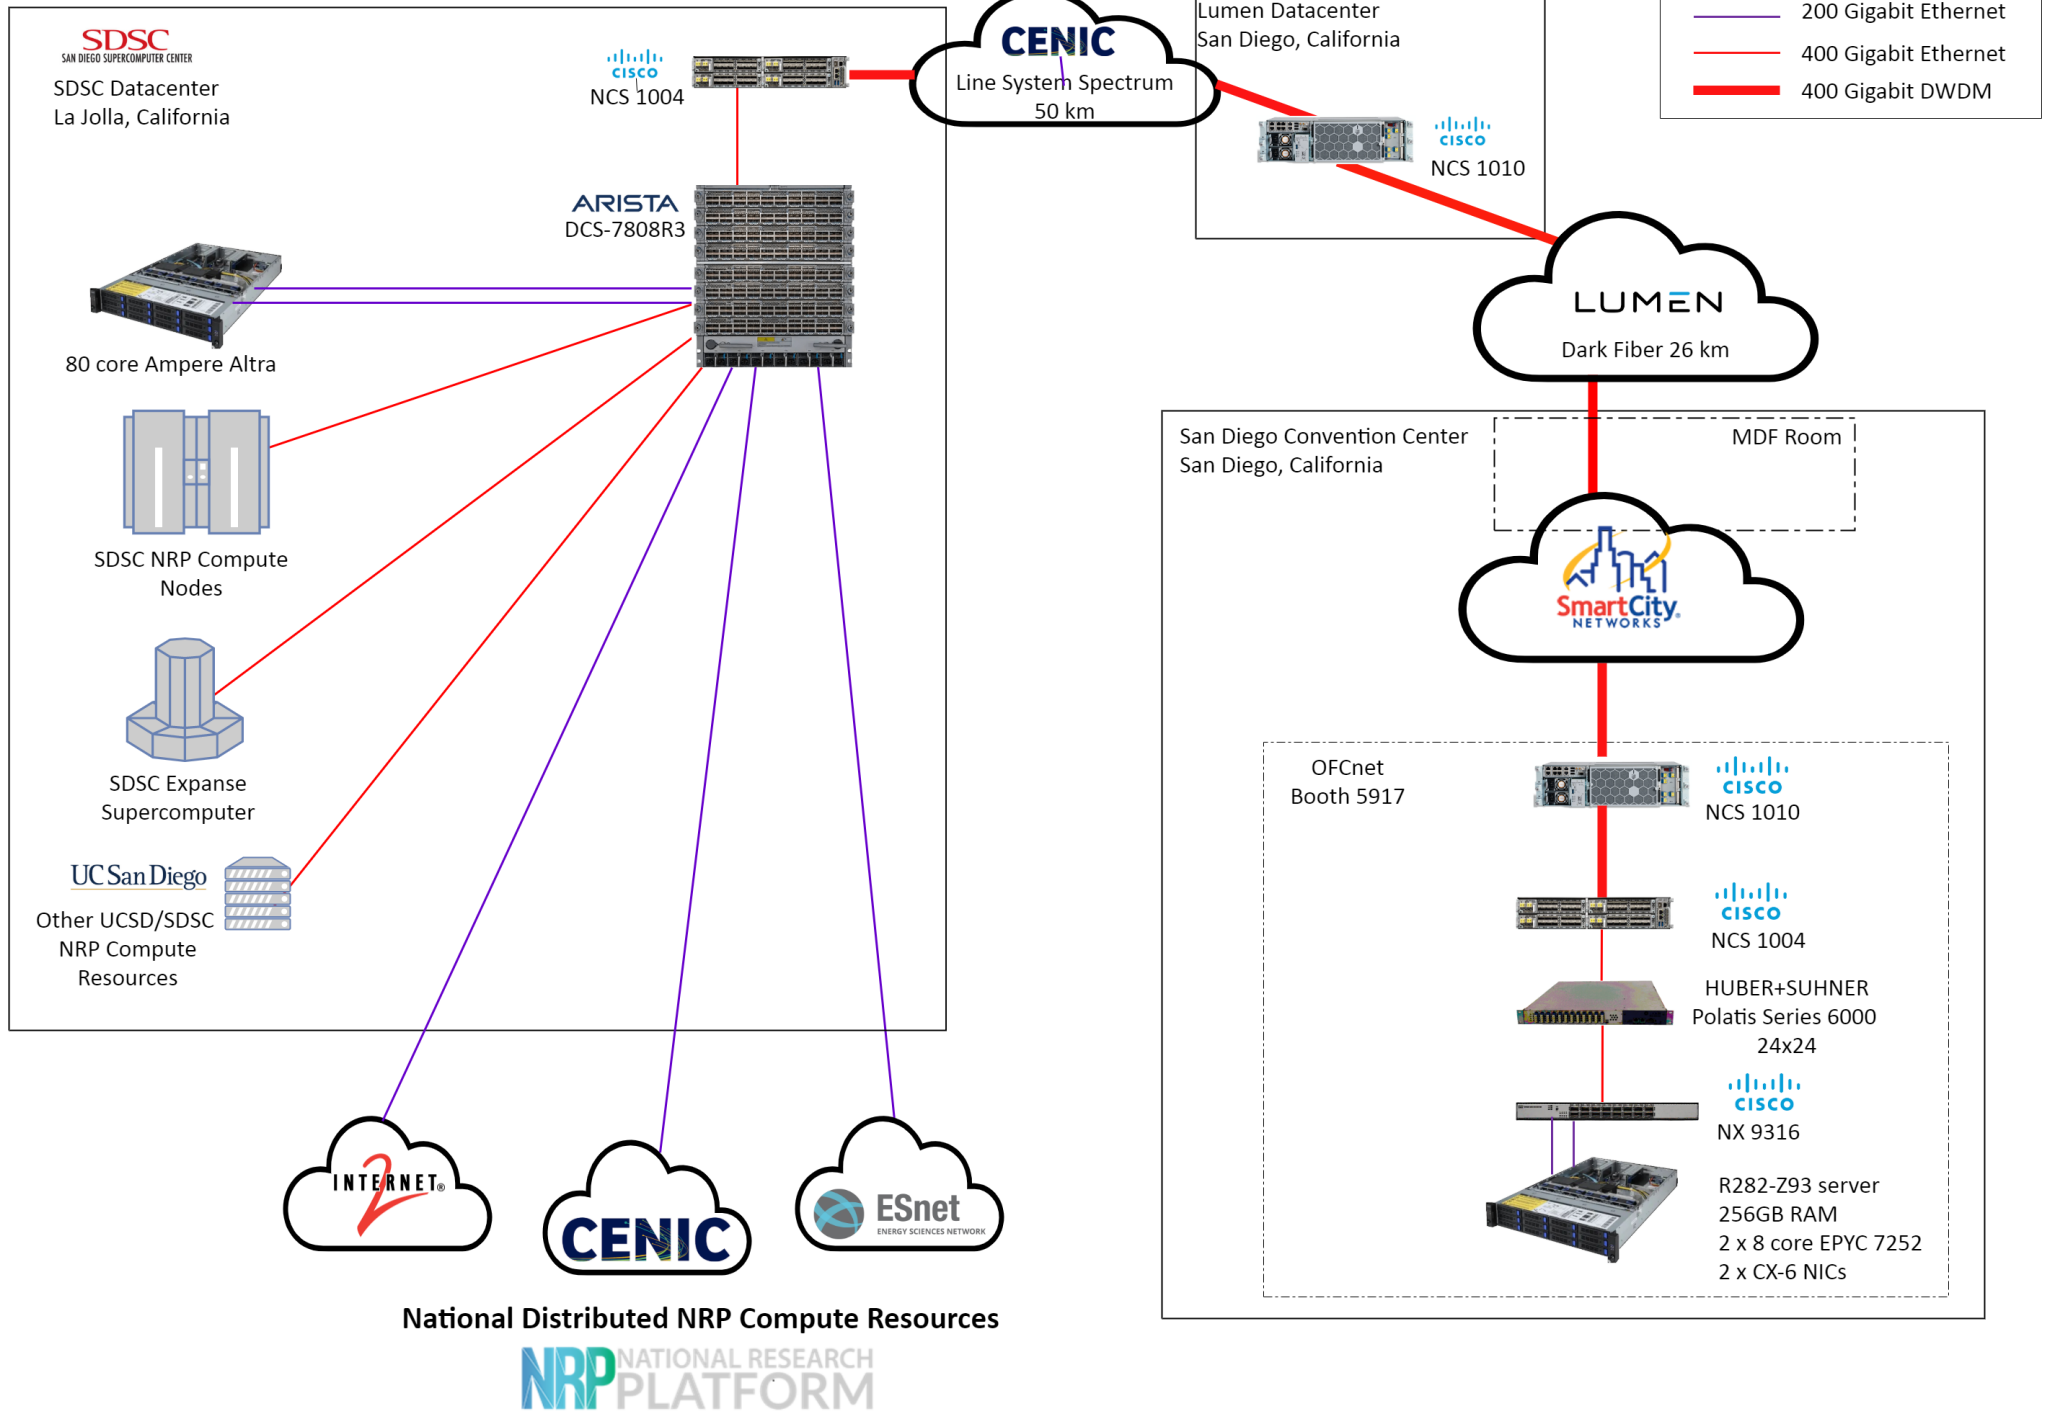
Nautilus = K8S infrastructure of PRP for the last 5+ years

Nautilus = K8S of NRP for the next 10 years

K8S is the defacto standard these days for devOps service deployments.

By being K8S native, people can join hardware anywhere in the world into Nautilus with extremely low TCO (decentralized bring your own hardware) and then use that hardware for global service deployments





Acknowledgements

This work was partially supported by the NSF grants OAC-1541349, OAC-1826967, OAC- 2030508, OAC-1841530, OAC-2005369, OAC-21121167, CISE-1713149, CISE- 2100237, CISE-2120019, OAC-2112167



Cees de Laat
University of Amsterdam



OFCnet Birds of a feather: Designing and Operating the Next Generation Optical Photonic Networks.

Moderators: Cees de Laat, University of Amsterdam
Reza Nejabati, University of Bristol

Program

Time	Title	Presenter
0:00	Welcome, introduction	Cees de Laat
0:10	Introduction to OFCnet	Marc Lyonnais, chair OFCnet
0:20	OFC Demozone	Marco Ruffini, Ben Puttnam, DemoZone
0:30	Panel introduction	Reza Nejabati
0:35	Panel	
	Andrew Lord	BT
	Hübel Hannes	AIT
	Richard Murray	Orcacomputing
	Daniel Kilper	TCD
	Inder Monga	ESnet
	Jörg-Peter Elbers	ADVA
	Each gets 5 minutes to present, 5 minutes to discuss	
1:35	General discussion with audience on outcomes & next steps	
2:00	close	

Goal for this BOF

In this BOF we propose a workshop series that solicits papers and demonstrators reports on all aspects of building networks out of components and using those networks for the whole range of commodity to extreme applications. The aim of the workshop series is to bridge and expand between the Technical Programs Demo Zone and the newly created OFCnet. Furthermore, we will solicit input on possible challenges and awards for demonstrating novel new architectures, technologies and implementations.

- a clear direction, scope and format for a workshop series to start in 2024
- identify co-chairs for such workshop
- publication venue
- potential challenges and awards to be formalised in a call for participation for 2024

Outcome of the BOF

- lots of positive remarks and suggestions were voiced regarding the OFCnet exhibit
- several attendees mentioned Hackatons & attention to software
- challenges & awards
- dedicated EXPO sessions
- possibly posters, call for abstracts/short papers
- work on the complementary positioning regarding Demo-Zone
 - should we equip networking to the demo zone?
- no major request to start a workshop series at this time

Actions

- create a program committee to:
 - define, oversee and judge challenges
 - call for short papers
 - for expo
 - for posters (with demo-zone)
- software oriented team for organizing hackathons
- Any other suggestions?



Questions/Comments