



SC19

Denver, CO | hpc
is now.

scinet

SC19 SCinet INDIS Workshop

Architecture, Technical Challenge, XNET and
NRE

November 17th, Denver, CO

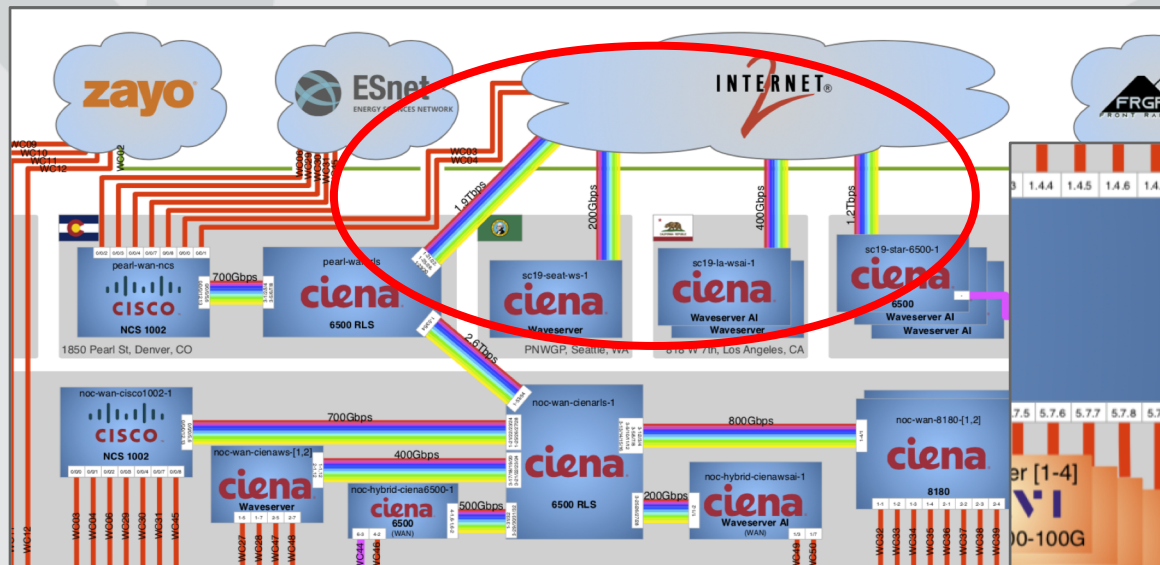
SCinet Agenda

- Architecture
- Technical Challenge
- Experimental Networking
- Network Research Exhibition

SCinet Architecture

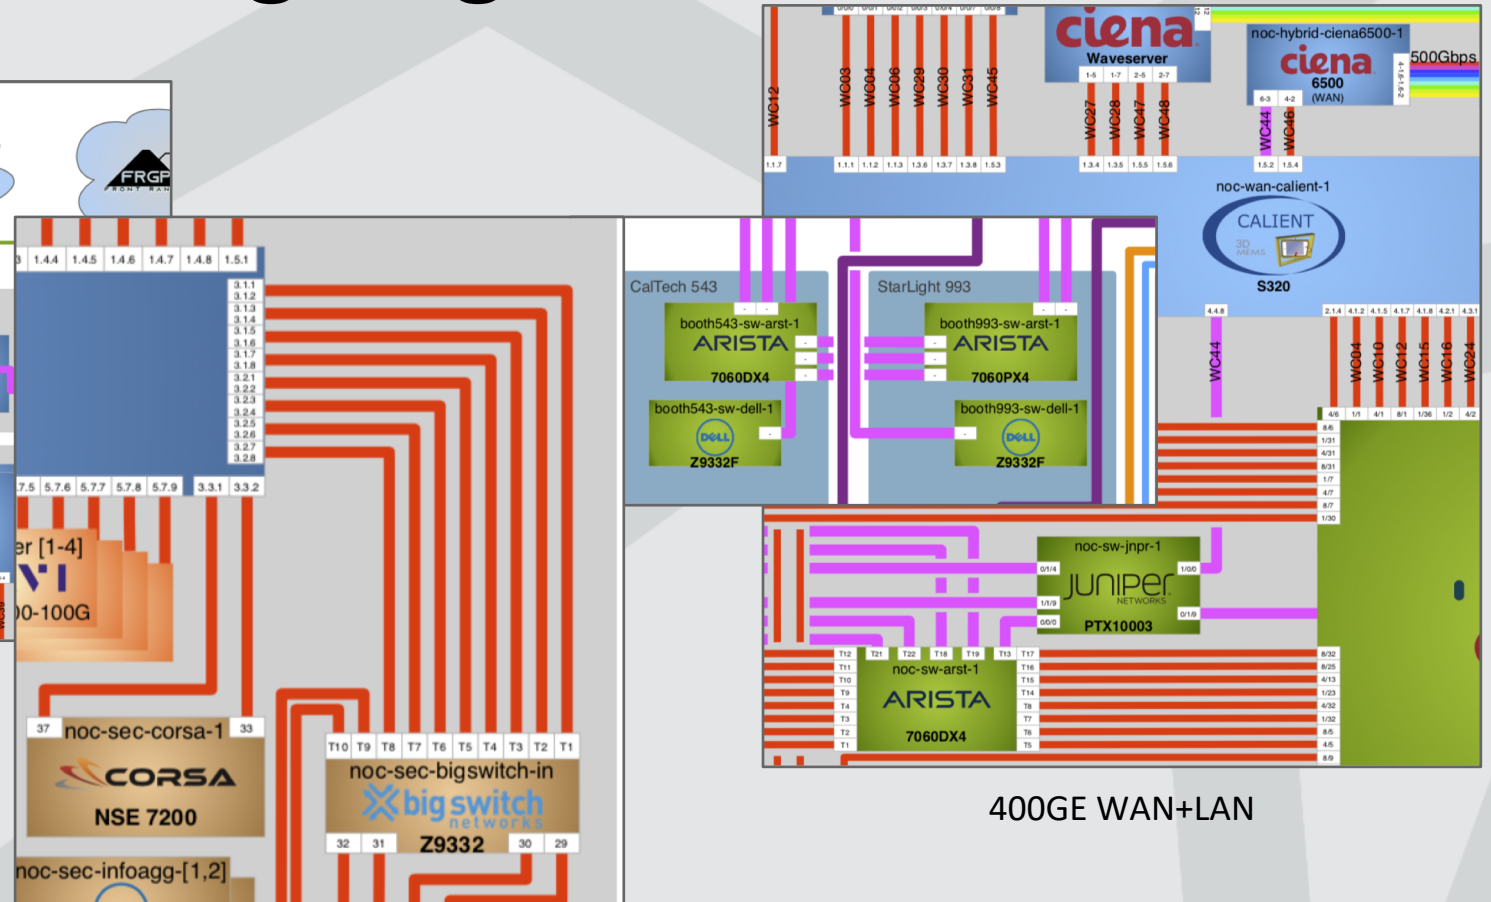
SC19/SCinet Architecture Team:
Lance Hutchinson, Sandia National Laboratories
Scott Richmond, ESnet
JP Velders, Universiteit van Amsterdam

Packet Architecture high-lights



Open Line System:

- across I2 Spectrum
- Chicago (StarLight):
 - 8x100GbE
 - 1x400GbE
- Seattle (PNWGP):
 - 2x100GbE
- Los Angeles:
 - 4x100GbE



400GE WAN+LAN

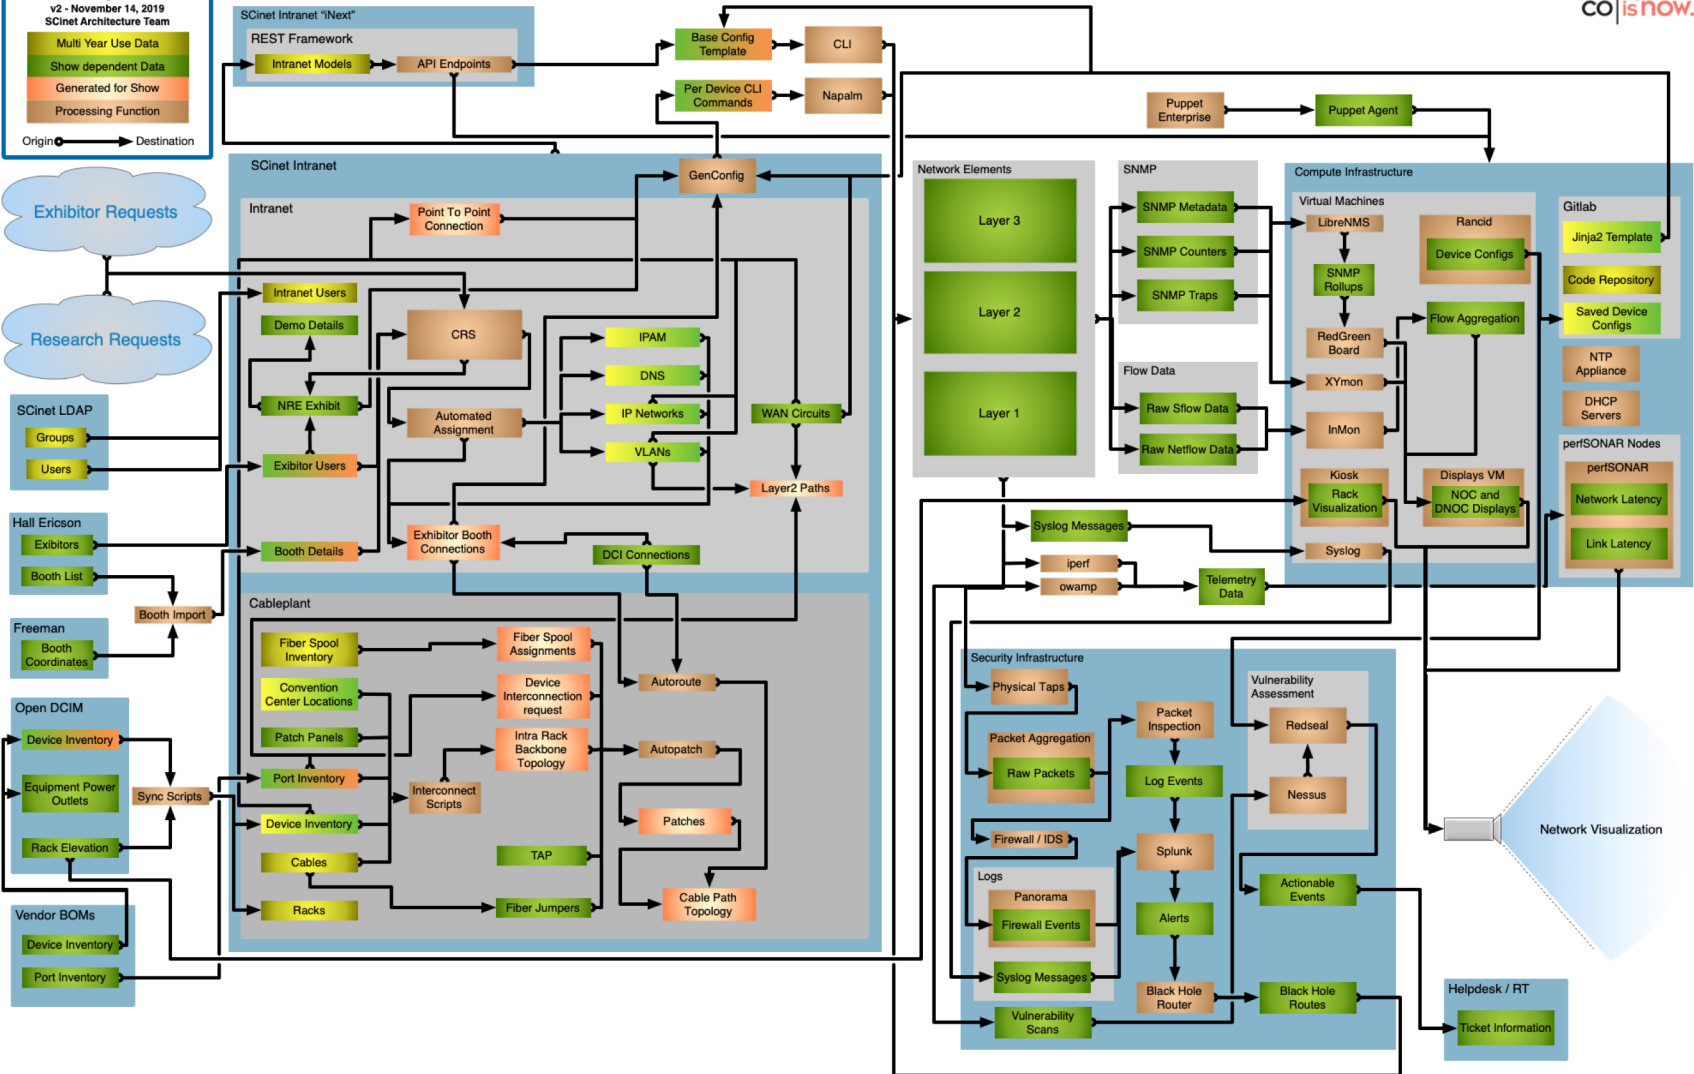
Optical CrossConnect inserts taps on circuits and feeds Security Infrastructure.

V2:

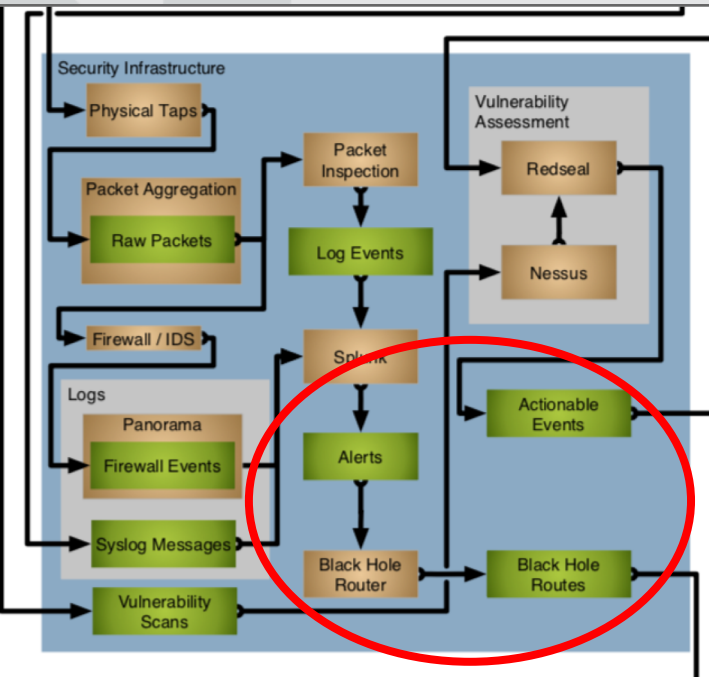
SCinet
Information Ecosystem
v2 - November 14, 2019
SCinet Architecture Team

- Multi Year Use Data
- Show dependent Data
- Generated for Show
- Processing Function

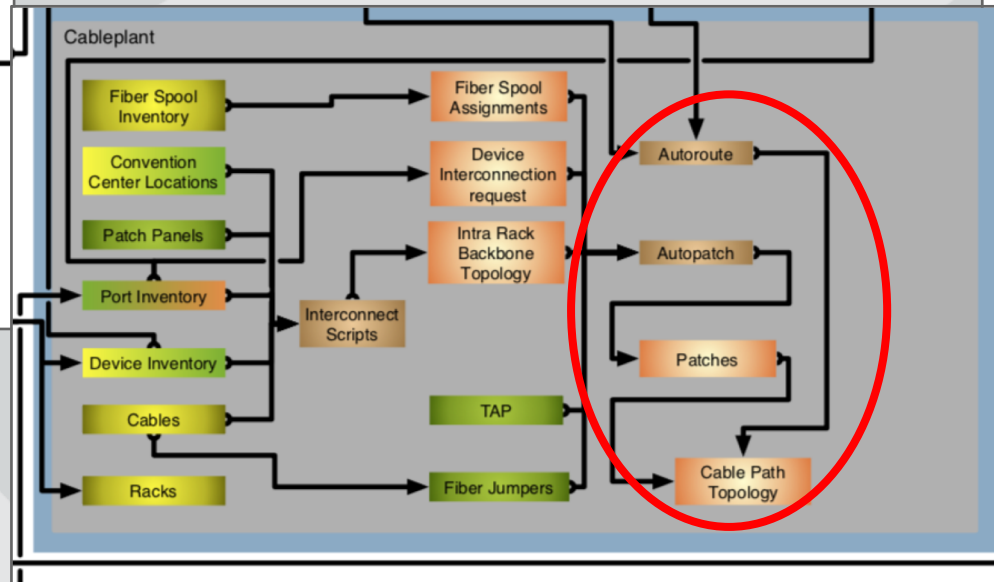
Origin → Destination



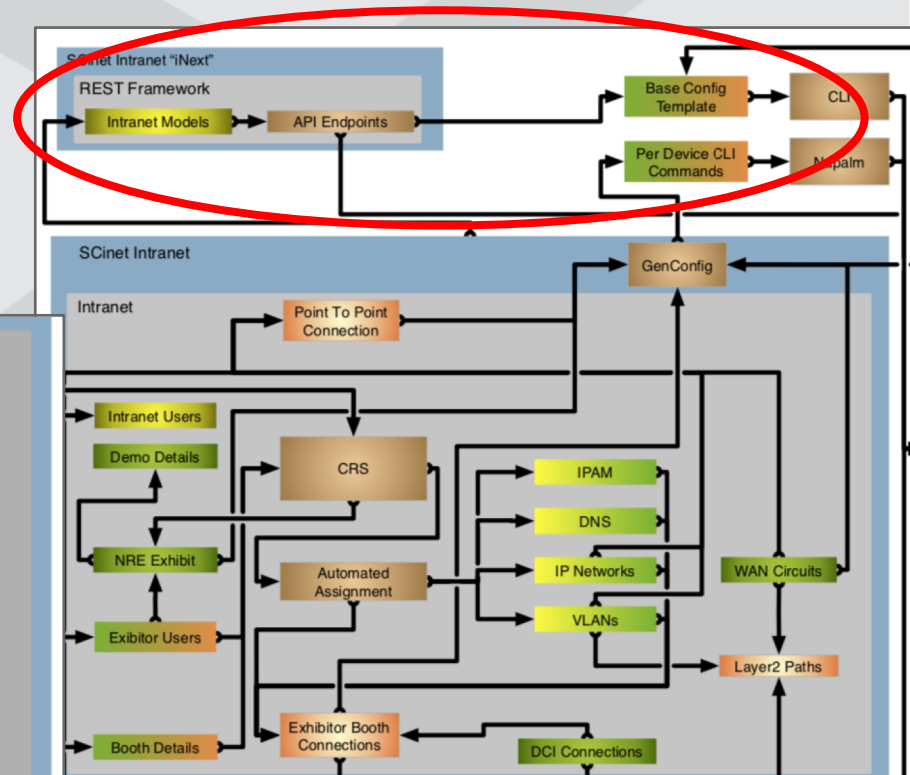
Information Ecosystem high-lights



From data via actionable information to actual improved (Network) Security



Automated generation of Layer 0 / Physical (fiber) topology



Automated configuration towards Orchestration

SCinet Technical Challenge TC


Presented by Cees de Laat for

Organizers:

Ilya, Marc, Paola, Sarah, JP, Michelle, JP, Christer, Troy, NRE team,
others

Jury:


Jim Rogers, Kathy Yellick, Dan Stanzione,
Inder Monga, Rodney Wilson, Cees de Laat



AIM

Bring Supercomputing and Networking together!

- **Networks** become **commodity**
- **SuperComputers dwarf** in the Clouds
- **Quantum around the corner** (very big corner ;-))
- **The art of doing Science is changing!**
- **Machine Learning** and **Artificial Intelligence** coming up
- Data poor to Data rich and **streaming/near real time** & **decision support**



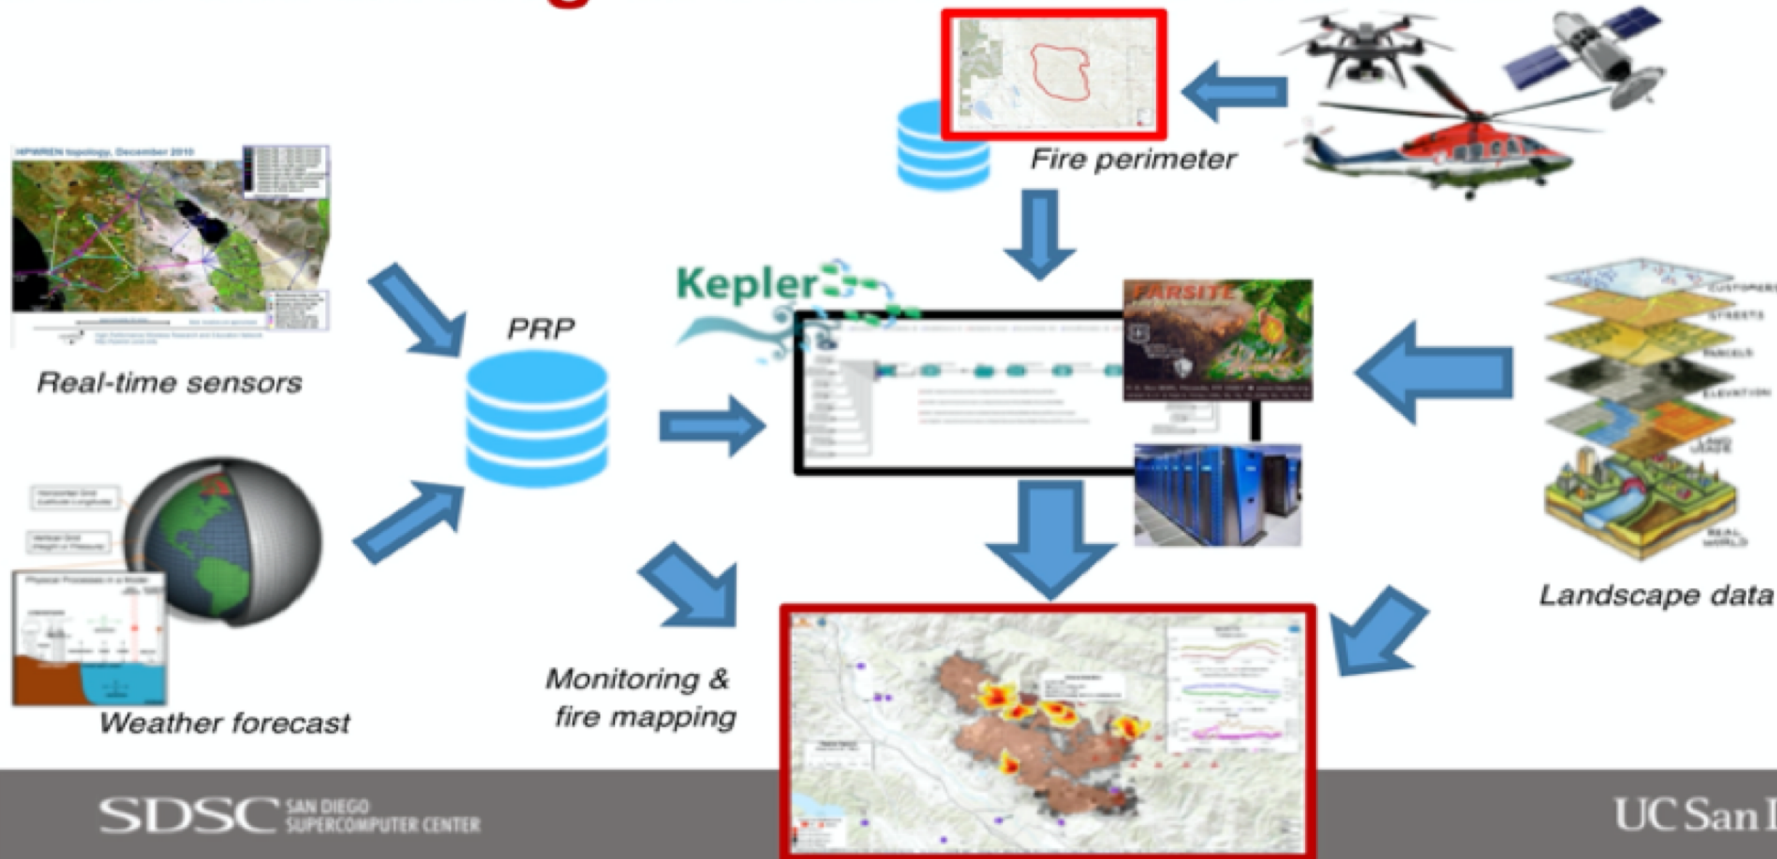
COMPETENCES:


The conversation ICT - Scientist is profoundly changing

- Discussion is on **data & method level**
- **Competences and methods** needed in **ICT & Data Science support** evolve.
- And that will change again if we get our own **SIRI 4 Science**
 - For example **The DoE AI townhalls** <https://www.sciencemag.org/news/2019/10/department-energy-plans-major-ai-push-speed-scientific-discoveries>
- Application **requirements** are **changing**
- **PRP-NRP-GRP** pave the way to **make data flow** to the AI absorbers
- We are well **beyond Bandwidth Challenges!**

EXAMPLE:

Fire Modeling Workflows in WIFIRE





So:

The first annual SCinet Technology Challenge (TC) will feature scientific demonstrations and experiments that highlight **interdependent operations** of **sophisticated networking, computing** and **storage** infrastructures.

The goal of the challenge is to demonstrate that **both networking and high-performance computing resources are essential** elements of the cyber infrastructure required to advance modern data-driven scientific applications.



What:

Teams of researchers will demonstrate their advanced **scientific applications**, while taking advantage of **distributed storage** and **compute** resources located on SC floor and elsewhere such as **university data centers**, **national labs**, **public clouds**, all linked by **the ultra-high-performance SCinet network!**



WHEN:

Tuesday, November 19, 10:30-12:00 and 13:15 - 14:00, SC Theatre next to NOC

- Technology Challenge Overview and Jury Introduction
- **5G Citizens Broadband Radio Service (CBRS) Proof-of-Concept for Scientific Applications**
 - University of Utah Center for High Performance Computing, Murray School District and the Utah Education and Telehealth Network
- **Dynamic Network-Centric Multi-Cloud Platform (DyNamo) for Real-Time Weather Forecasting Workflows**
 - RENCI/UNC Chapel Hill, USC/ISI, UMass Amherst and Rutgers University
- **Real-Time Analysis of Streaming Synchrotron Data**
 - Argonne National Laboratory, Northwestern University, Starlight, Northern Illinois University, University of Chicago



AND THEN:

Thursday, November 21, 10:30 am The SC Theatre in front of the NOC

- **Technology Challenge Recognition Ceremony**
 - **2019 SCinet Technical Challenge Most diverse resource set**
 - For highest diversity of resource types and geographic distribution, their utilization and degree of orchestration/automation
 - **2019 SCinet Technical Challenge Best use of data**
 - For greatest data velocity/variety/volume
 - **2019 SCinet Technical Challenge Most original technical approach**
 - For innovation and the originality of the technical approach
 - **2019 SCinet Technical Challenge Best presentation and visualization**
 - For best human interactivity and effectiveness of the presentation and best quality/originality of visualization
 - **2019 SCinet Technical Challenge Winner**
 - Brings networking, computing and storage together, top award

Experimental Networks

XNET

TeamLeads: Kevin Quire (UEN)
Nick Buraglio (Esnet)
Jerry Sobiexski (NORDUnet)


What is “*XNET*”?

Experimental Networks - That time between the research Proof of Concept and a full Production capability ... often 10 years or more (!)



SC *XNET* tries to deploy promising experimental technologies that require at-scale deployment and real-world experience to drive their evolution.

XNET is [part of] a pipeline for introducing innovative technologies into SCinet.



XNET Projects for SC19

Data Transfer Nodes as a Service – Jim Hao Chin (iCAIR), & Anna Giannakou(LBNL), { Fei Yeh, Se-Young Yu, Xiao Wang } (iCAIR), Eric Pouyoul(Esnet)

Streaming Telemetry – Kevin Quire (UEN) & Matt Mayes (UEN)

NetBox -GeorgeRobb(ESnet)

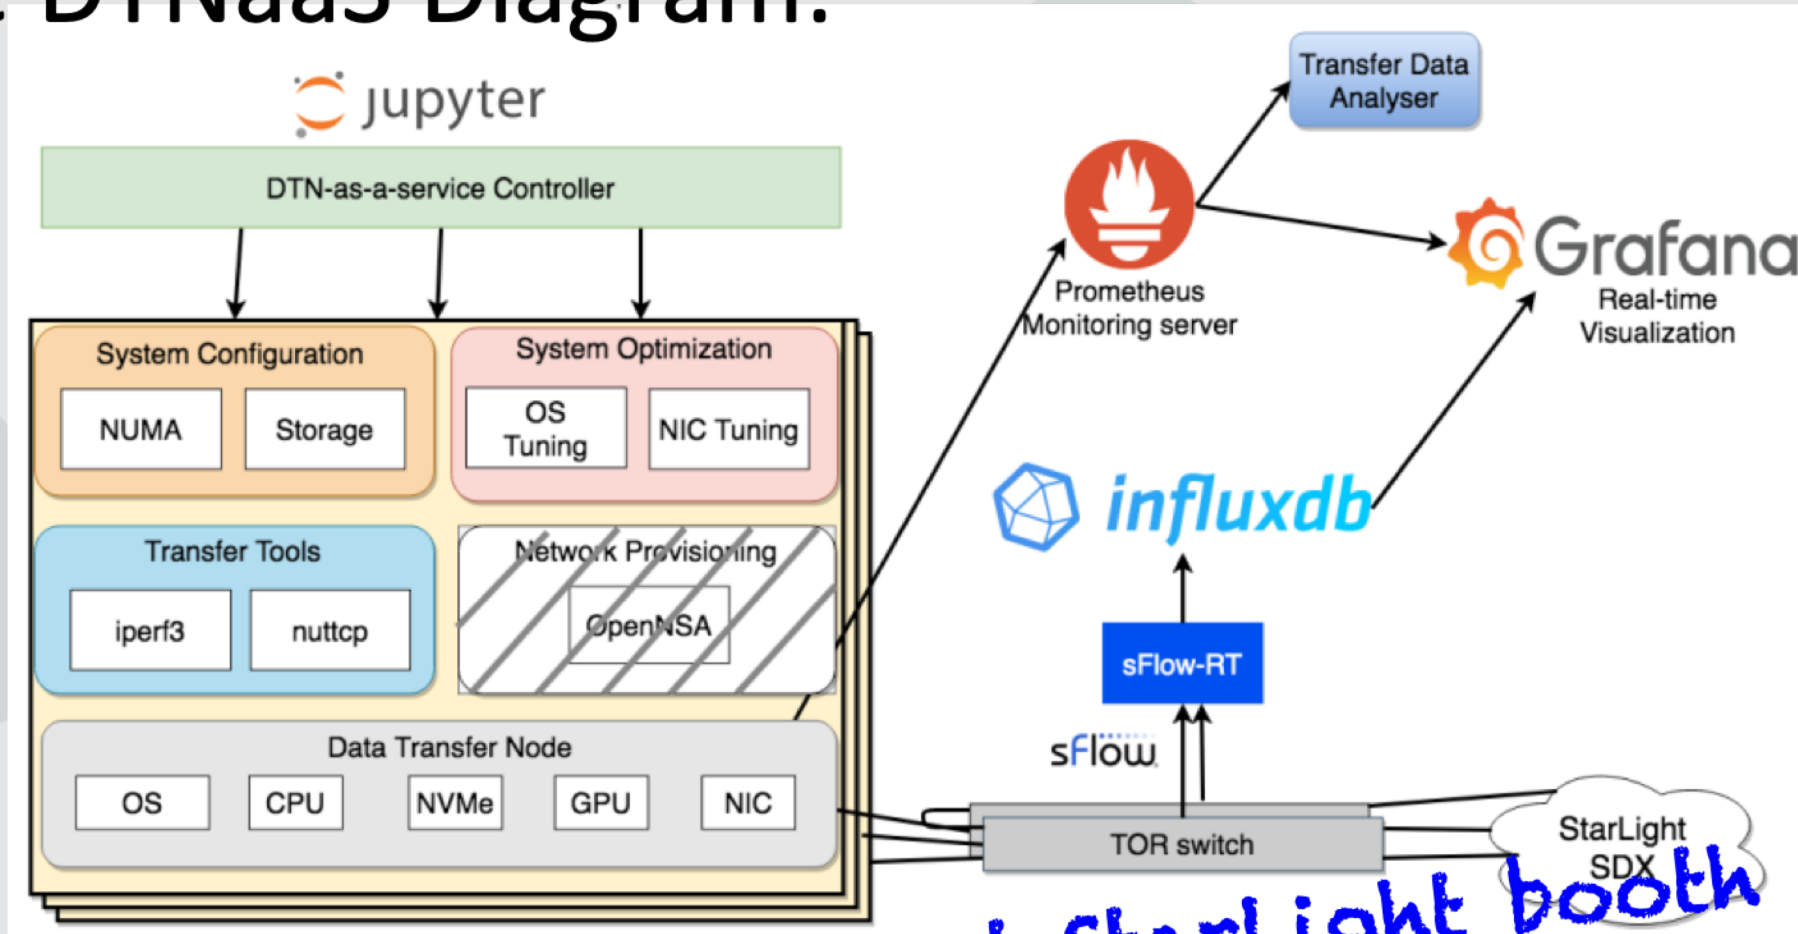
Global Virtualization Services – Jerry Sobieski(NORDUnet),Susanne Naegele-Jackson(FAU/DFN),Michal Hazlinsky(CESNET), Andy Forrester, Hussein Al-Azzawi(UNM), GEANT, Internet2




Data Transfer Nodes

- SC19 SCinetDTN-as-a-Service is a 3rd year X-NET/NRE project.
- The project provides Data Transfer Node software and hardware platform as prototype services to support SC SCinetcommunity
- For SC19, new prototype services include: Kubernetes, NVMeoFand 400G LAN/WAN experiments.
- For SC20, SCinetDTN plans to establish as part of DevOps Services
- Please see SC19 INDIS Workshop paper: “SCinetDTN-as-a-Service Framework” for detail.

The DTNaaS Diagram:



FFI: Jim Chin at StarLight booth 993



What is the “netbox” project?

- NetBox is an open source web application designed to help manage and document computer networks.
- NetBox was developed specifically to address the needs of network and infrastructure engineers.
 - **IP address management (IPAM)** - IP networks and addresses, VRFs, and VLANs
 - **Equipment racks** - Organized by group and site
 - **Devices** - Types of devices and where they are installed
 - **Connections** - Network, console, and power connections among devices
 - **Virtualization** - Virtual machines and clusters
 - **Data circuits** - Long-haul communications circuits and providers
 - **Secrets** - Encrypted storage of sensitive credentials




Netbox progress at SC19

- Elegant Integration in to Authentication and Authorization.
- Clean and usable API interfaces that are compatible with existing infrastructure and business process.
- Expands upon features of present DCIM (including cable-plant and house models with pictures).
- Expands upon ability to rapidly model multiple sites and preserve previous sites.
- Trials are still in progress. [SC20's scheduled code sprint could bring in additional automation and expansion.]

FFI:

George Robb (ESnet)

see him at SCinet



Streaming Telemetry

- What is it?

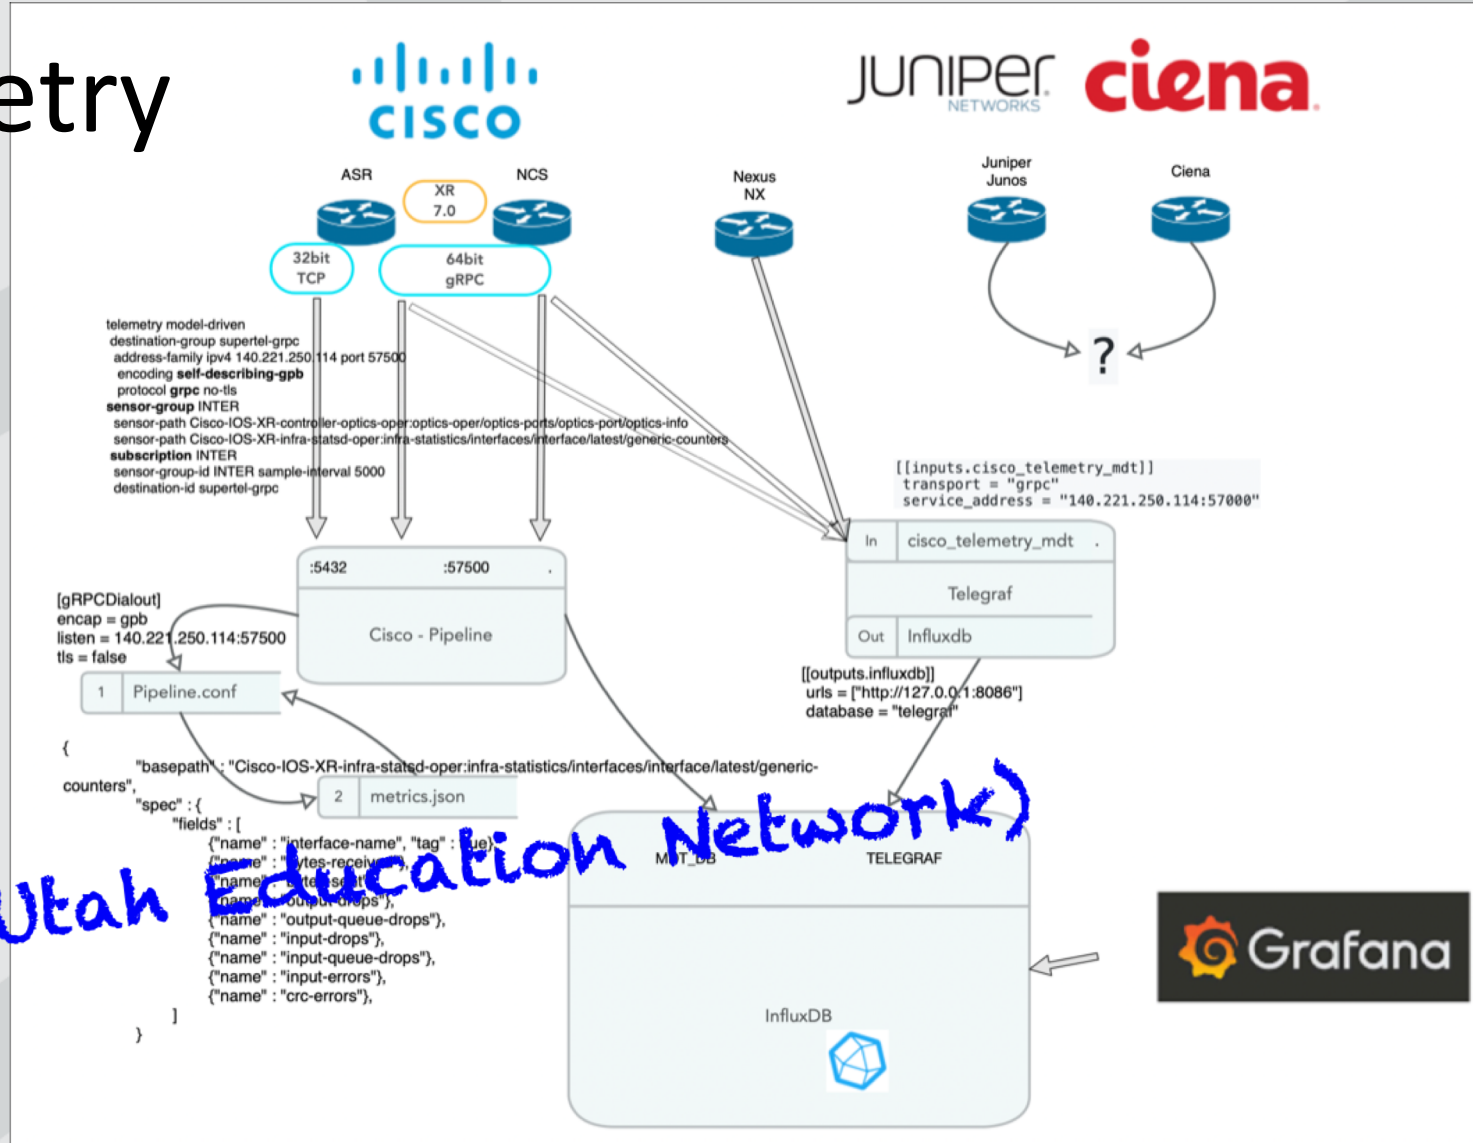
 - Bandwidth Interface statistics

 - SNMP pollers sample on (typically) 5 minute interval

A lot can happen in 5 minutes!

- Changing from 5 min to 1 min intervals increases telemetry by 30% (!)
- Going to 1 minute...or less
- Lots going on to grab data, store it, analyze it, and display in real time

Streaming Telemetry



FFI: Kevin Quire (Utah Education Network)
see him at SCinet



Global Virtualization Service (GVS)

- Virtualization of cyber-infrastructure (compute, transport, switching, storage, functions, mobile spectrum, instruments, sensors...)

 - Supports a common object model and common lifecycle API

 - Rapid global provisioning of virtual resources

 - Orchestrated, Automated reservations and provisioning

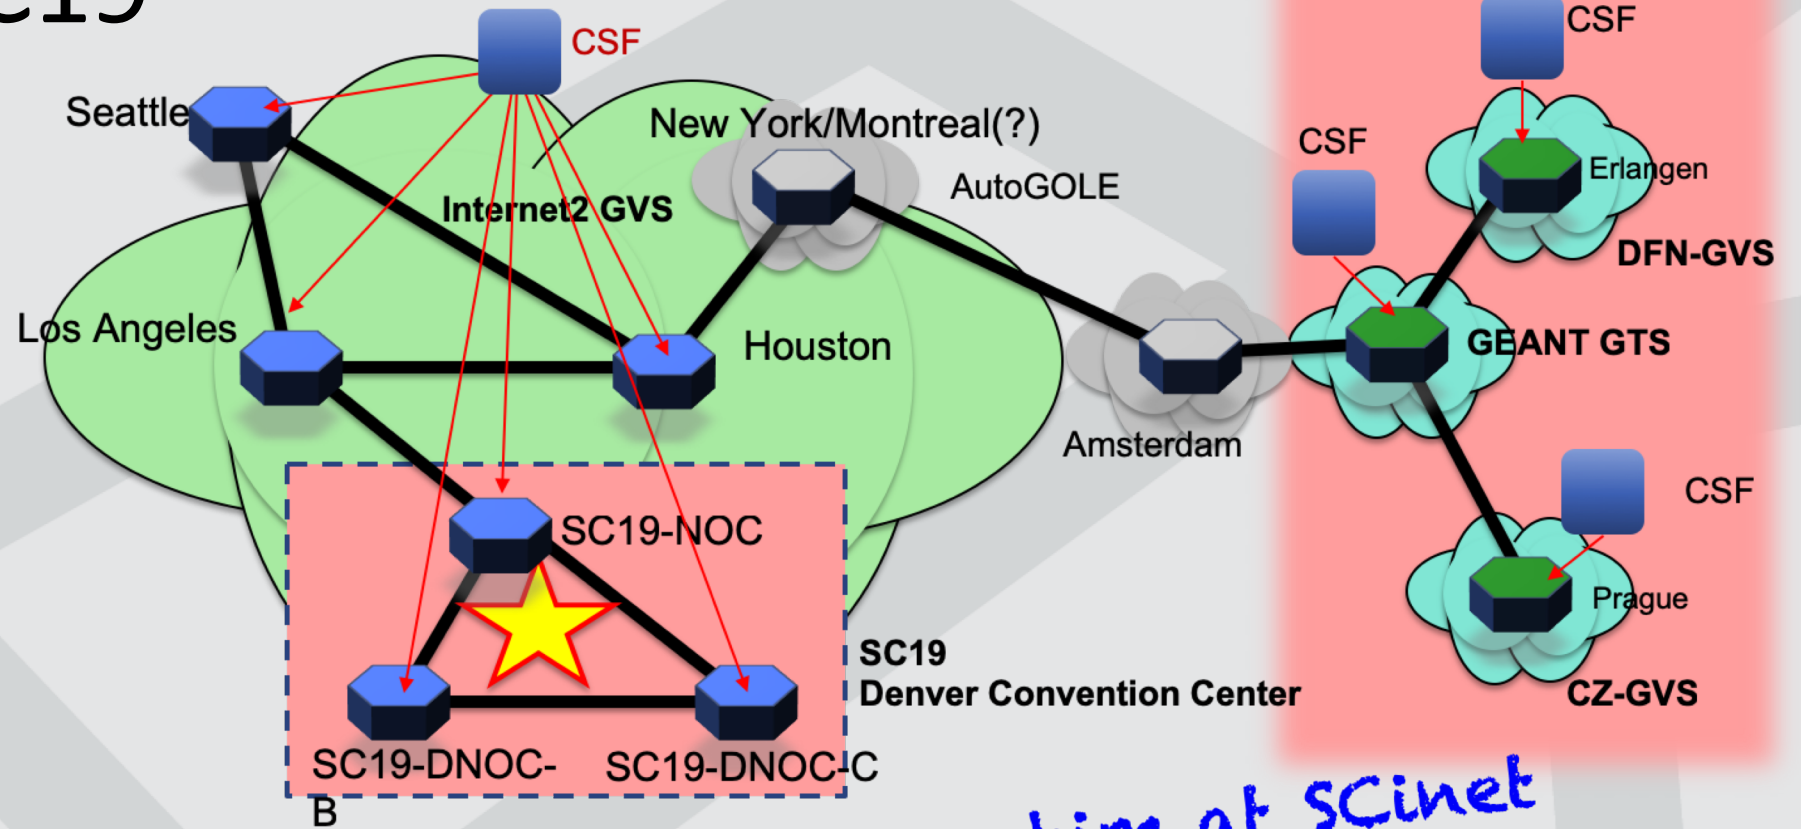
 - Line rate (100G) performance

 - Deterministic performance

 - OpenSource (developed by GEANT network community)

- Great potential for rapidly and easily establishing SC class global demonstration environments

GVS at SC19




FFI: Jerry Sobieski (NORDUnet) see him at SCinet

Network Research Exhibition NRE

David Wheeler (NCSA)
Corby Schmitz (ANL)

<https://sc19.supercomputing.org/scinet/network-research-exhibition/>



NRE Process

The Network Research Exhibition (NRE) team provides outreach and ongoing engagement with exhibitors undertaking experiments at SC

- Preliminary Abstract
 - provides an idea of where the research is headed
- Network Requirements
 - NRE team works with the experimenter to determine the technical needs of the experiment
 - Bundled and provided to the SCinet team for provisioning
- Final Publishable Abstract
 - NRE team collects and prepares the final abstracts for publication
- Publishable Results
 - Encourage the experimenters to share their results with the wider community
 - Showcase at the following year's INDIS workshop



NRE Outcomes

30 accepted demonstrations

91 hours of meetings gathering data

241 VLANs involved

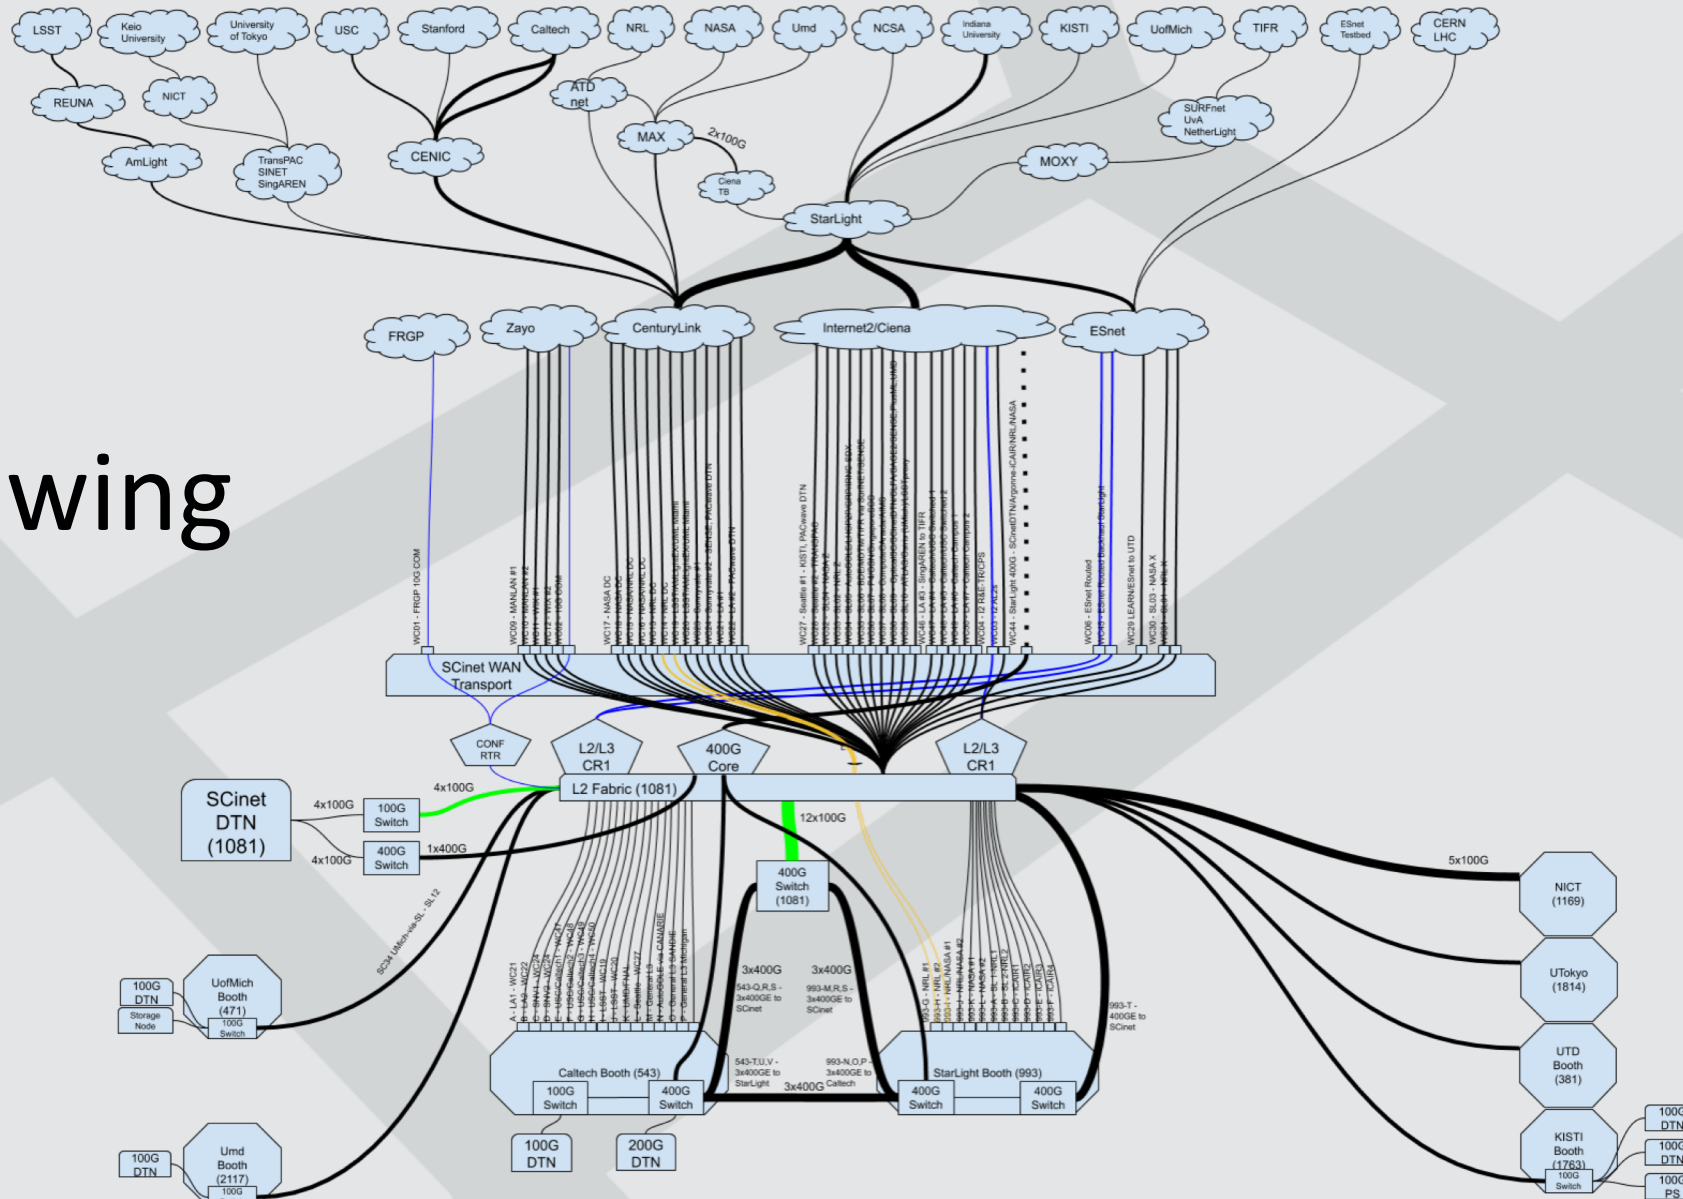

34 WAN connections involved (Including 400GE from Chicago to Denver)

52 Booth drops involved

8 SCinet DTN connections involved

Weekly communication with researchers and involved parties

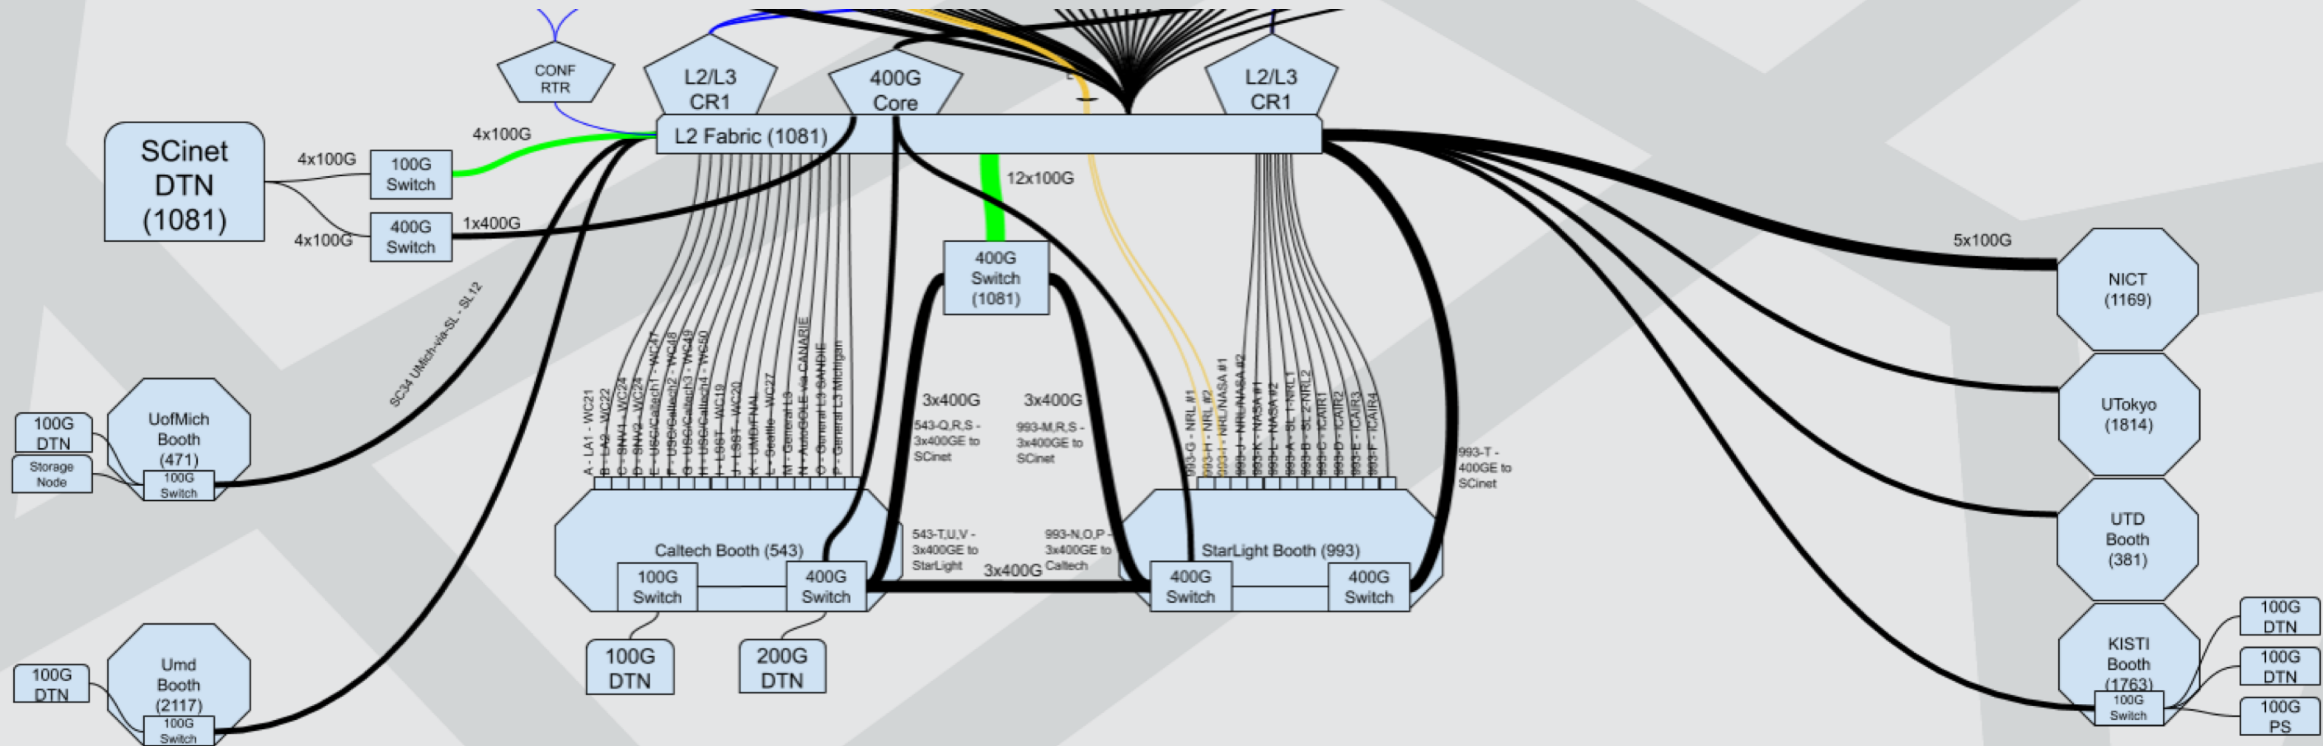





NRE Drawing



NRE Drawing - Local Focus





NRE Highlights

NRE-004 - 400GE WAN Service - Chicago to Denver

Multiple vendors and research participants came together to deliver

NRE-008 - Tracking Network Events with Data Structures

Near real-time tracking of networking events for security or orchestration

Two years of NRE experiments created basis for larger LDRD support

NRE-014 MMCFTP

5x 100GE into Denver (trans-Atlantic and trans-Pacific) for traffic flows

NRE-030 and NRE-031 - Ceph and OSiRIS and Traffic Management

Ceph Storage platform enhancements with intelligent

Dynamic adaptation to bottlenecks and async paths