Next Generation Cyberinfrastructures
for LHC, HL LHC and Data Intensive Sciences

+ INDUSTRY

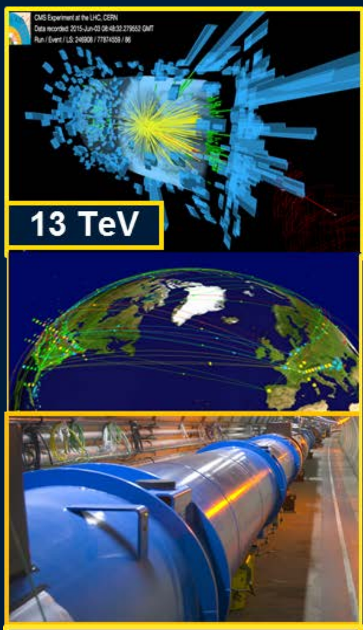

LHC

LSST

SKA

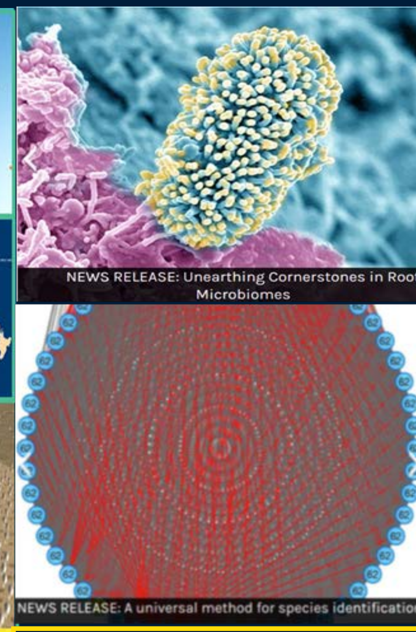

Joint Genome Institute

LHC Beyond the Higgs Boson

LSST SKA

BioInformatics

Earth Observation

Gateways to a New Era

Harvey Newman, Caltech

Americas Research Platform Workshop

UCSD CalIT2 Qualcomm Institute

September 17, 2019

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; 75 Year Exploration

Theory (1964): 1950s – 1970s;

LHC + Experiments Concept: 1984

Construction: 2001; Operation: 2009
Run1: Higgs Boson Discovery 2012

Run2 and Going Forward:
Precision Measurements and BSM
Exploration: 2013 - 2039

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

- NOTE: ~98% of Data Still to be Taken
- Greater Intensity: Upgraded detectors for more complex events
- *To 10X Data Taking Rate in 2027-39*

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

WLCG Transfers Dashboard: Throughput Aug. 2018 – Aug. 2019

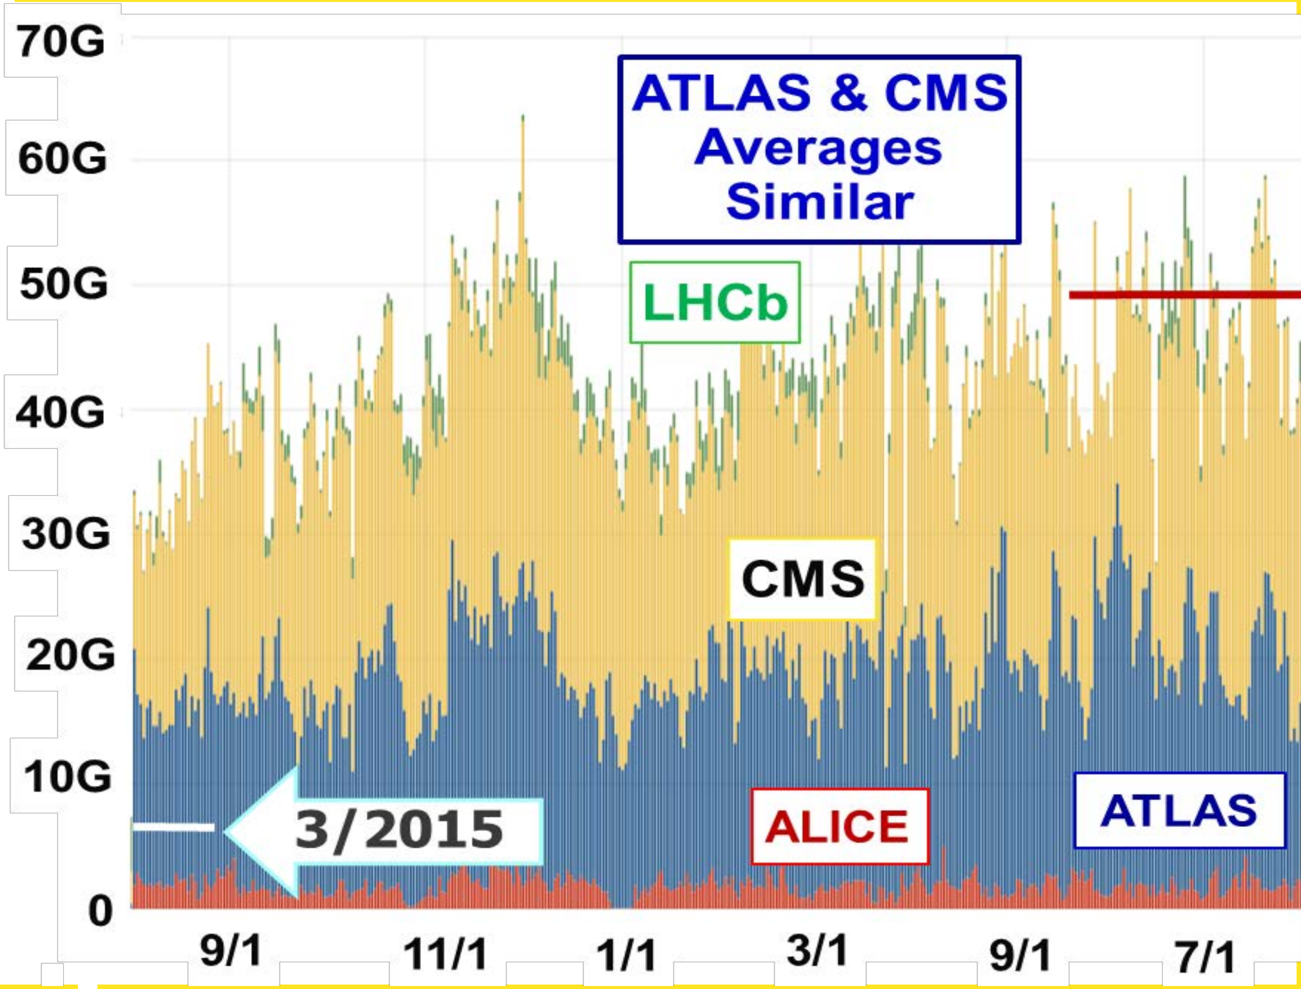

49 GBytes/s Sustained

60+ GBytes/s Peaks

Complex Workflow

- **Multi-TByte Dataset Transfers;**
- **6-17 M Transfers/Day**
- **>100k of remote connections (e.g. AAA) simultaneously**

7X Growth in Sustained Throughput in 4.3 Years: +57%/Yr; 90X per Decade

Core of LHC Networking LHCOPN, LHCONE, GEANT, ESnet, Internet2, CENIC…

LHCOPN: Simple & Reliable Tier0+1 Ops

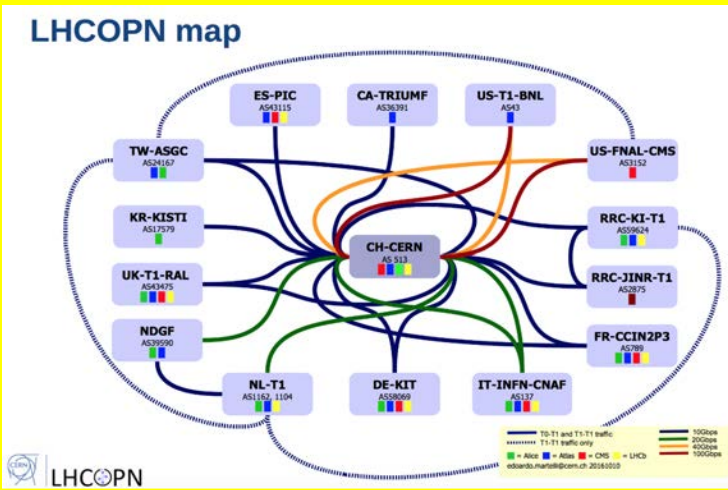

Internet2

GEANT

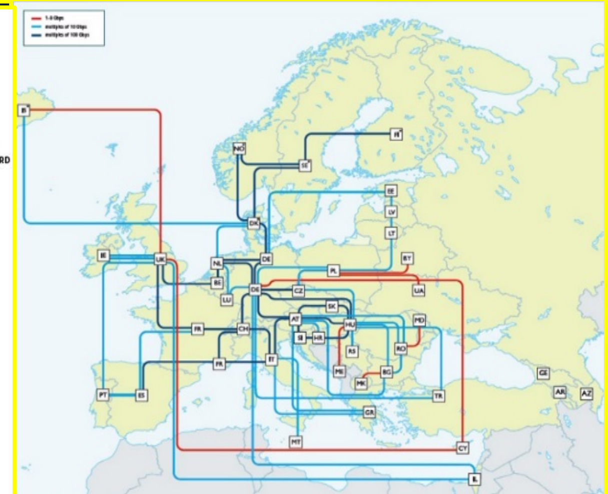

LHCONE VRF: 170 Tier2s

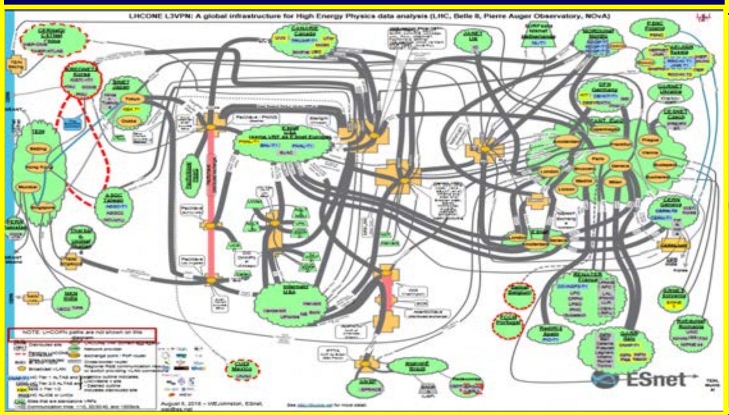

ESnet (with EEX)

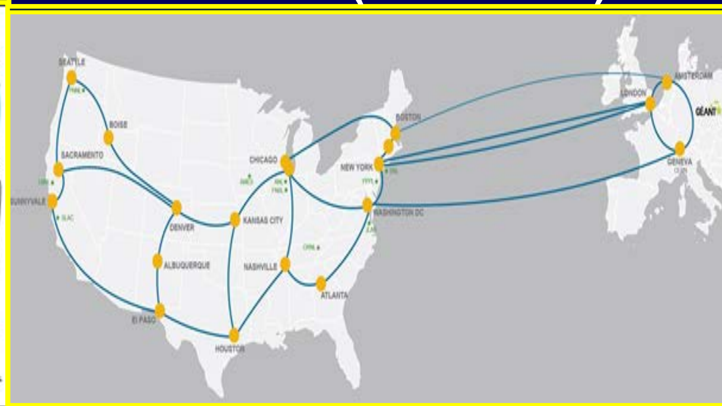

CENIC and PRP

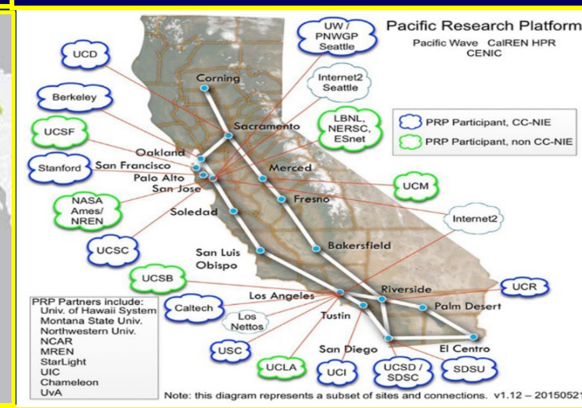

+ NRENs in Europe, Asia, Latin America, Au/NZ; *US State Networks*

Facts of Life: Towards the Next Generation Computing Model (CMS Example)

"Naïve" Extrapolations: Daunting!

- HL-LHC scales for CMS computing
  - Exa-byte scale disk and tape storage (x50 w/r to now)
  - CPU needs 5M cores (x20 w/r to now)
  - transfer of exa-byte-sized data samples across the Atlantic at 250-500 Gbps (ESnet now has allocated 40Gbps transatlantic for the LHC)
- These estimates got DOE's attention…

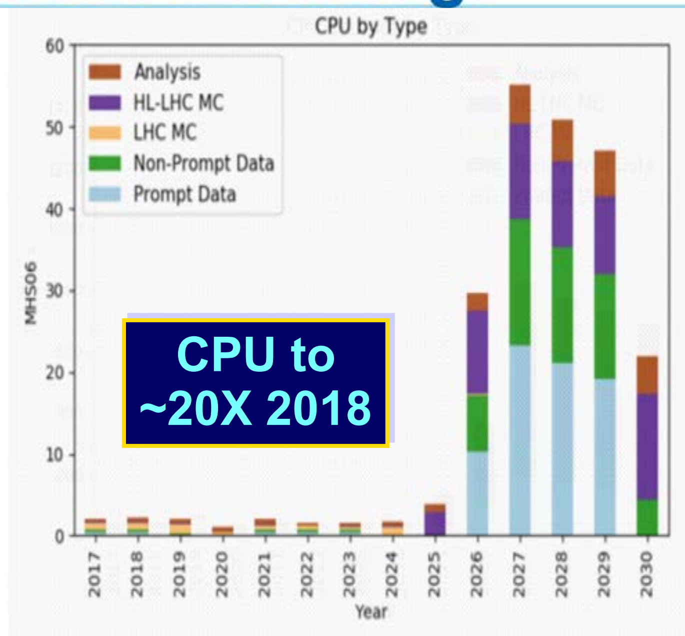

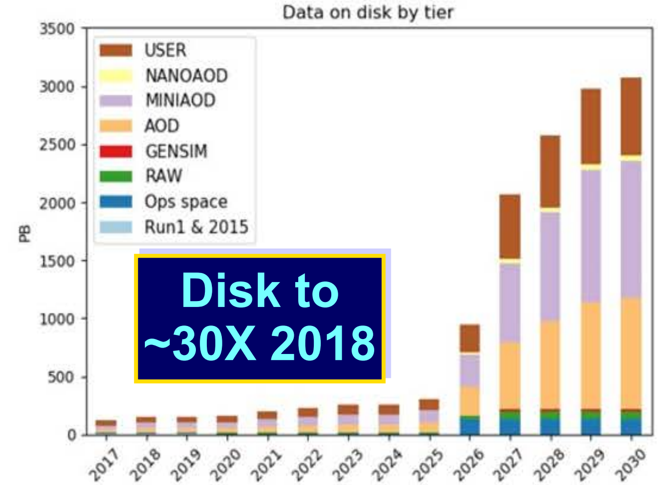

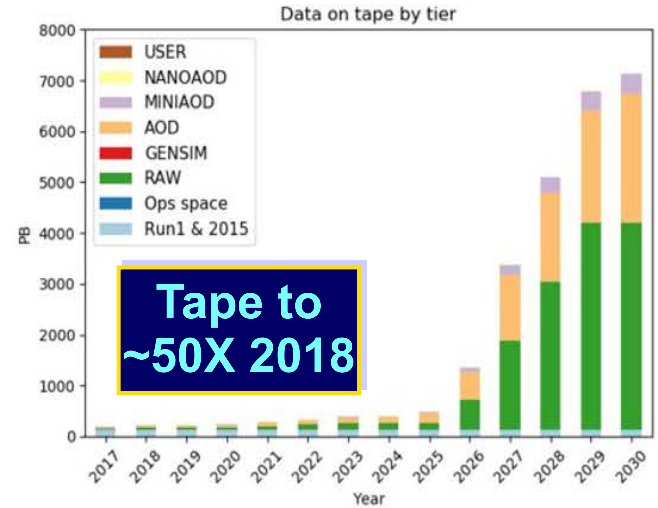

A New Era of Challenges: Global Exabyte Data Distribution, Processing, Access and Analysis

- **Exascale Data for the LHC Experiments**
  - **13 Tier1s and 170 Tier2s: ~1 Exabyte Now; ~5 to 50 EB during HL LHC Era**
- **Network Total Flow of >1 EB this Year**
  - **>1 Exabyte flowed over WLCG in 2018**
- **Emergence Now of 400G in Hyper-Data Centers, 200G in Wide Area**
  - **800G; 400G in Wide Area by 2021-22**
- **Network Dilemma: Per technology generation (~10 years)**
  - **Capacity at same unit cost: 4X**
  - **Bandwidth growth: 35-70X in Internet2, GEANT, ESnet**
- ***During LHC Run3***
  We will likely reach a network limit
- **Unlike the past: Optical and switch need to move beyond evolutionary; or**
  Physics Limits by ~HL LHC Start

New Levels of Challenge

- **Global data distribution, processing, access and analysis**
- **Coordinated use of massive but still limited *diverse* compute, storage and network resources**
- **Coordinated operation and collaboration *within and among* scientific enterprises**


- **HEP will experience increasing Competition from other data intensive programs**
  - **Sky Surveys: LSST, SKA**
  - **Next Gen Light Sources**
  - **Earth Observation**
  - **Genomics**

Responding to the Challenges

New Overarching "Consistent Operations" Paradigm

- **VO Workflow Orchestration systems that are**
  Deeply network aware, reactive, adaptive and pro-active
- ***Network Orchestrators with similar, real-time character***
- **Together responding in real-time to:**
  - **State changes in the networks and end systems; anomalies**
  - **Actual versus estimated transfer progress, access IO speed**
- **Prerequisites:**
  - **End systems, data transfer applications and access methods capable of high of throughput [➡ e.g. FDT]**
  - **Realtime end-to-end monitoring systems [End sites + networks]**
- **Elements for efficient operations within the limits:**
  SDN-driven bandwidth allocation, load balancing, flow control at the network edges and in the core
- **SDN + Ai-Driven Workflow Optimization: Success Metrics that balance**
  throughput, resource use, policy/priority per VO, fair sharing …
- **Beyond Deep Learning: Classical Ai +Game theory for Stable Solutions**

SENSE Solution Approach

End-to-End model-based distributed resource reasoning and intelligent service orchestration

- Hierarchical service resource architecture
- Unified network and end-site resource modeling and computation
- Model based realtime control
- Application driven orchestration workflow
- End-to-end network data collection and analytics integration

Ai for Science Berkeley Town Hall Breakout: I. Monga et al.

AI is essential for facilities, and facilities are essential for AI

Without the integration of facilities in the AI workflow, AI for science is impossible.

Steps in the AI workflow

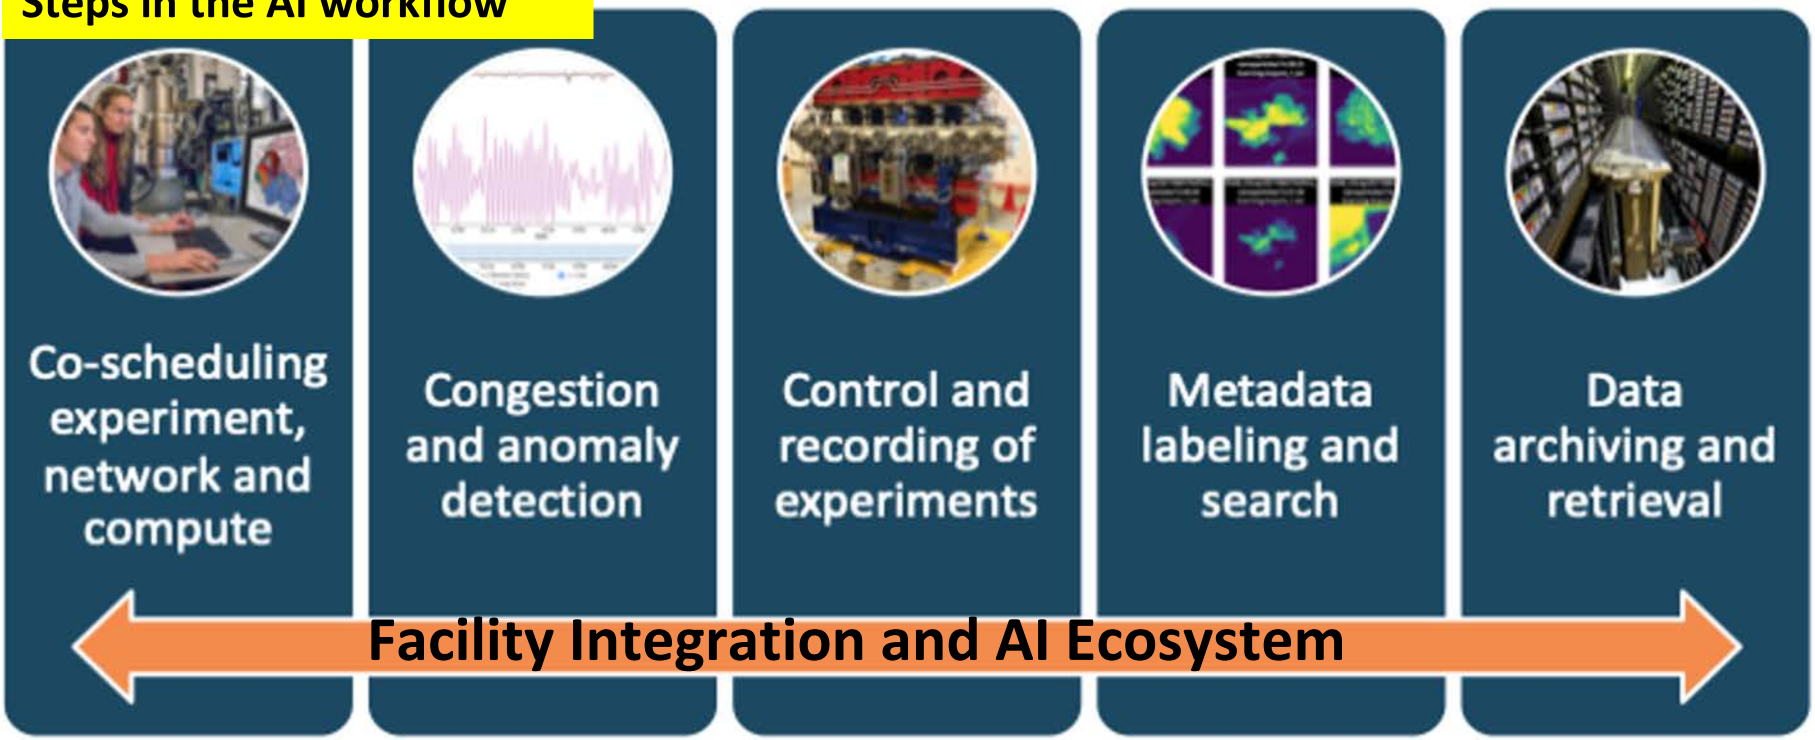

Automation is Key in Facility Integration: AI for AI

SC15-19: 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 Site Edges of Exascale Facilities With 100G -1000G DTNs

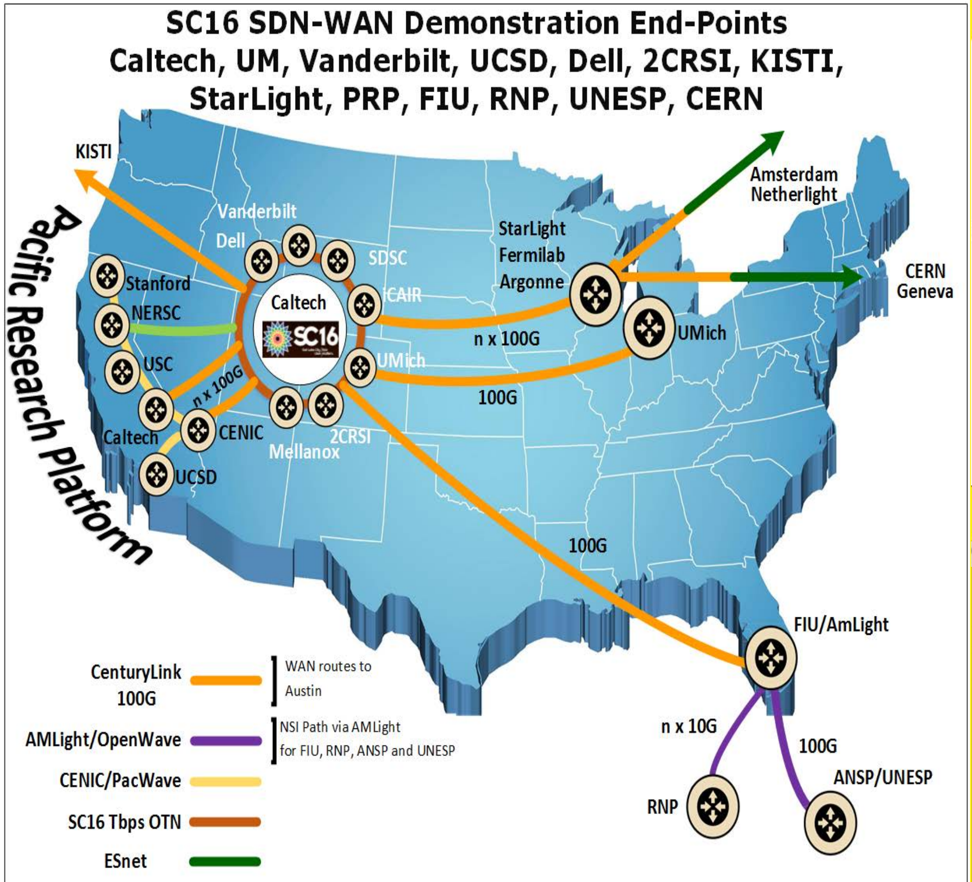

900 Gbps Total
Peak of 360 Gbps in the WAN

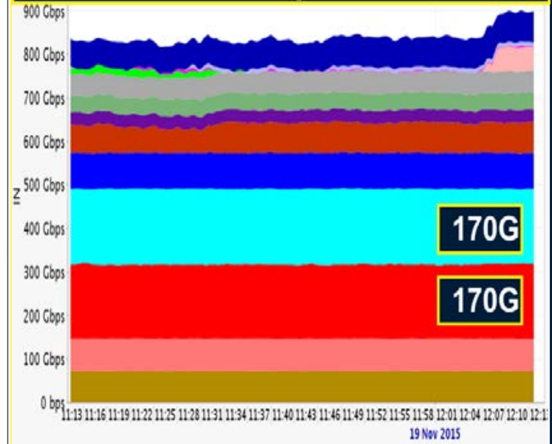

Global Topology

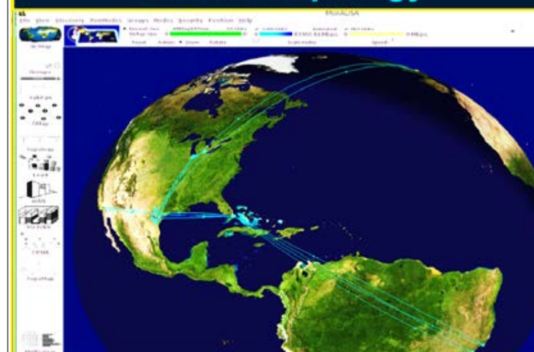

LHC at SC15: Asynchronous Stageout (ASO) with Caltech's SDN Controller

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 Ring for SC17: Caltech, Ciena, Scinet, OCC/ StarLight + Many HEP, Network, Vendor Partners at SC16

SC18

Global Petascale to Exascale Science Workflows

Accelerated by Next Gen SDN Architectures and Applications

+ INDUSTRY

Caltech and Partners at SC19

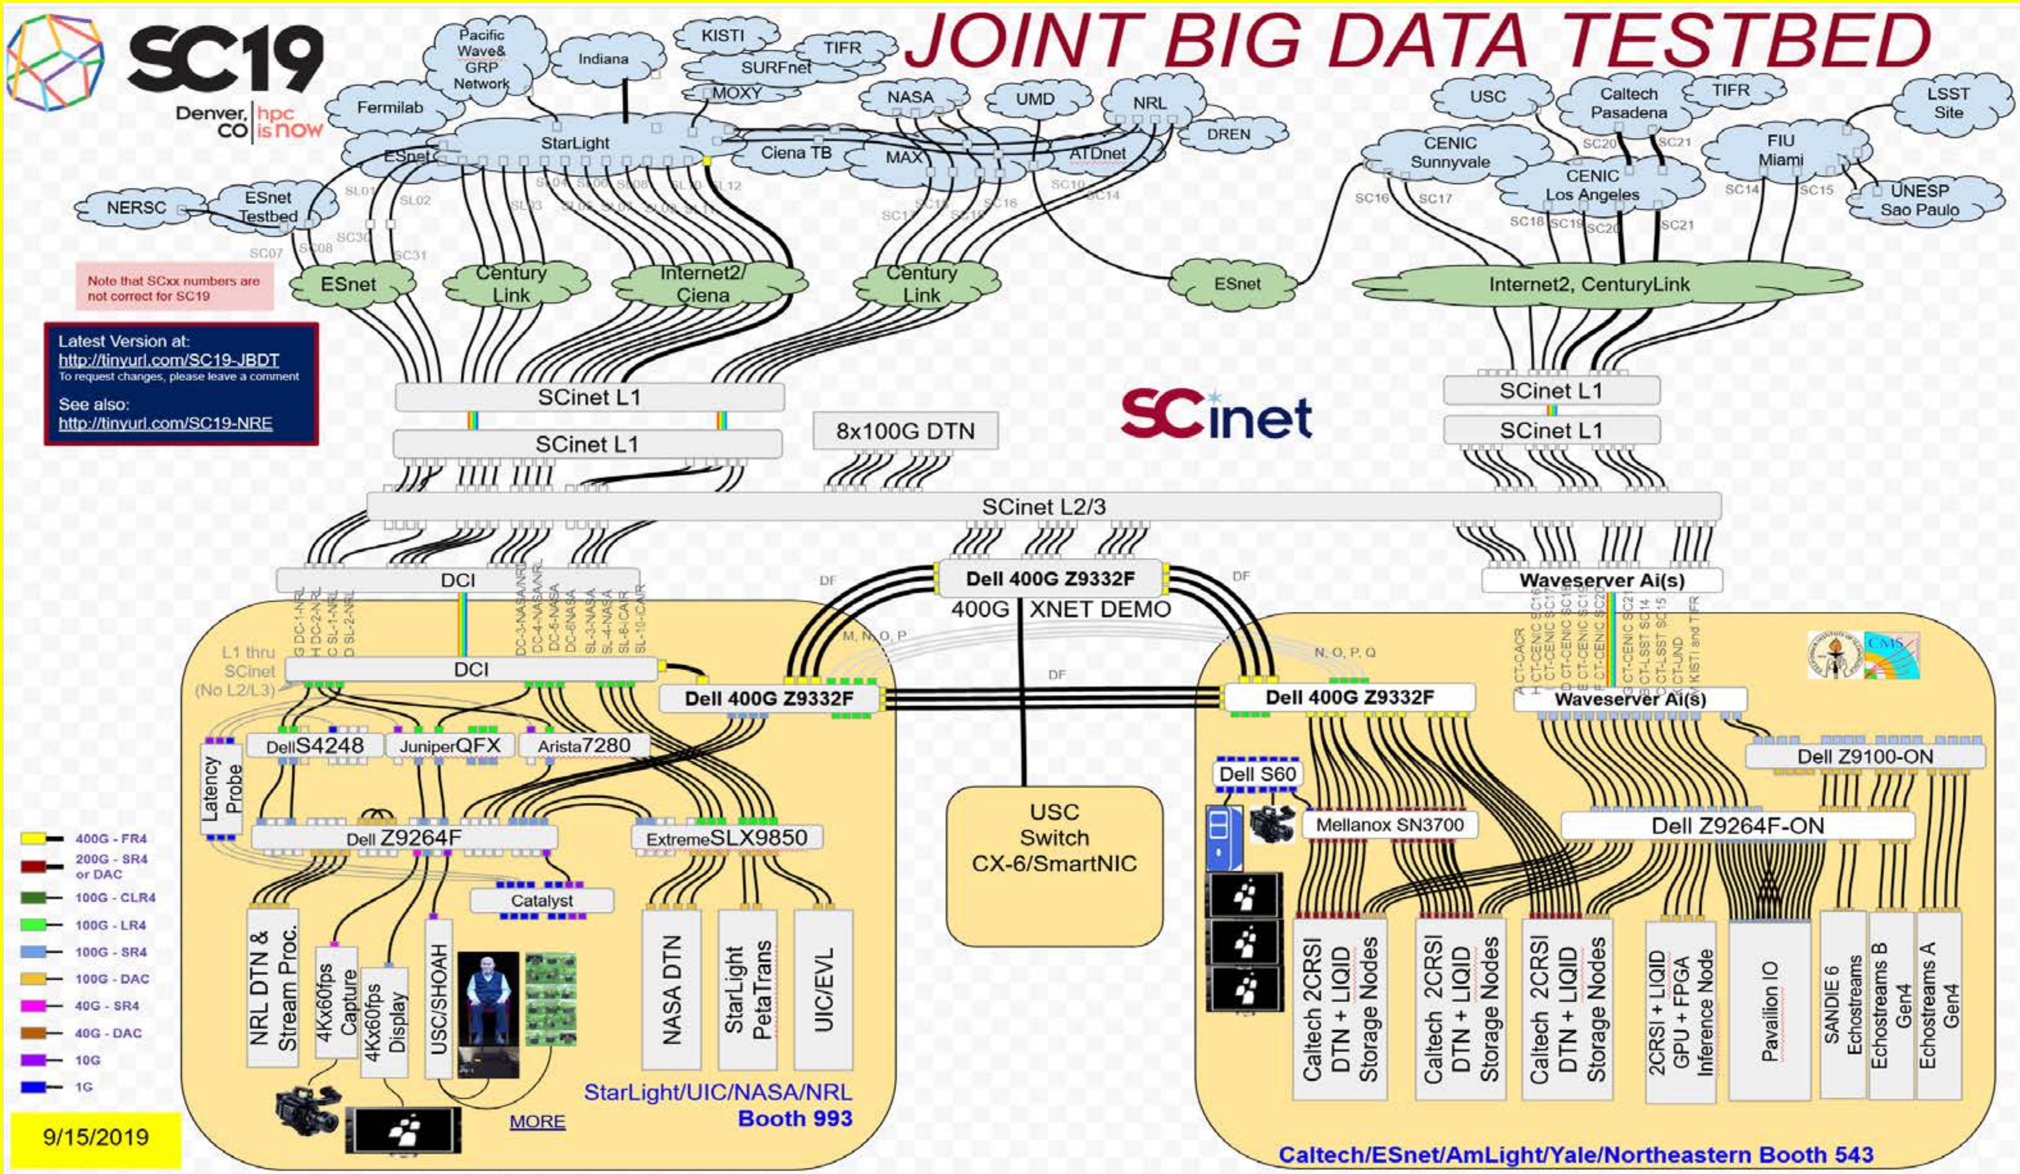

Microcosm: Creating the Future of SCinet and of Networks for Science