NGenIA: Next Generation SDN Integrated Architecture for LHC and Exascale Science

SC17 Denver, CO hpc connects.

+ INDUSTRY

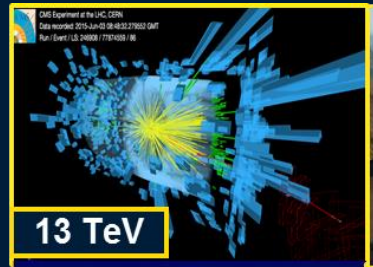

LHC

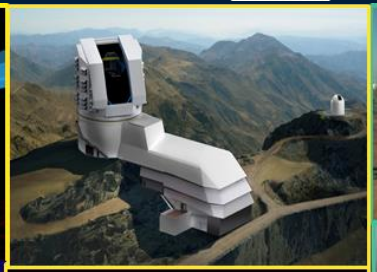

LSST

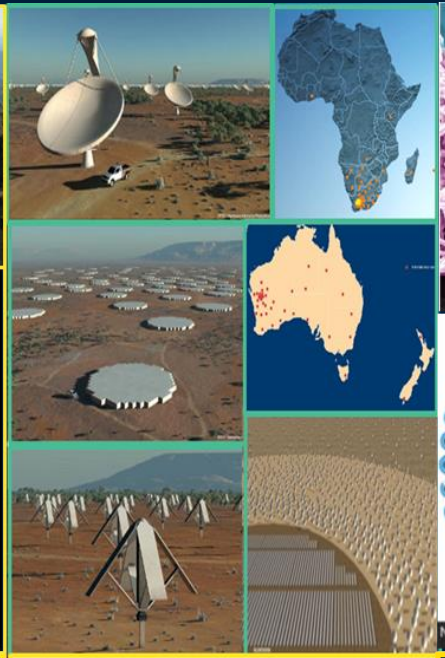

SKA

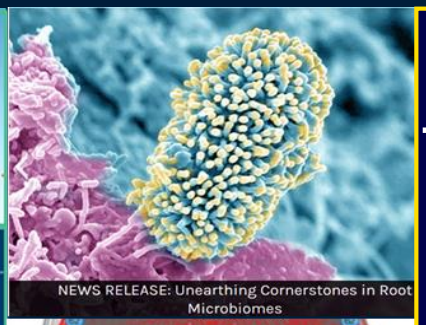

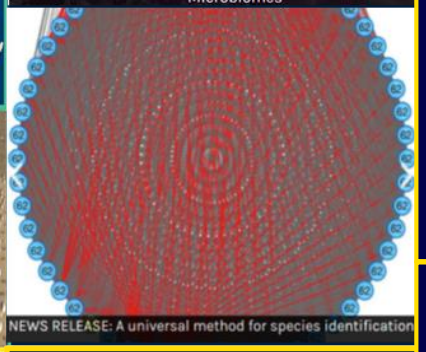

Joint Genome Institute

LHC Beyond the Higgs Boson

LSST SKA

BioInformatics

Earth Observation

Gateways to a New Era

Harvey Newman, Caltech

INDIS Panel on Data Intensive Sciences

Denver, November 12 2017

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

- Exascale Data
  - 0.5 EB now, 1 EB by end of LHC Run2
  - To ~100 EB during HL LHC Era
- Network Total Flow of >1 EB this Year
  - 850 Pbytes flowed over WLCG in 2016
- Projected Shortfalls by HL LHC
  - CPU ~4-12X, Storage ~3X, *Networks*
- Network Dilemma:
  Per technology generation (~8 years)
  - Capacity at same unit cost: 4X
  - Bandwidth growth: 30X (Internet2); 50X (GEANT), 70X (ESnet)
- *During LHC Run3 (~2022)*
  we will likely reach a network limit
- This is unlike the past
  - Optical, switch advances are evolutionary; *physics barriers ahead*

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

SC15-17: SDN Next Generation Terabit/sec Ecosystem for Exascale Science

supercomputing.caltech.edu

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

SC17 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

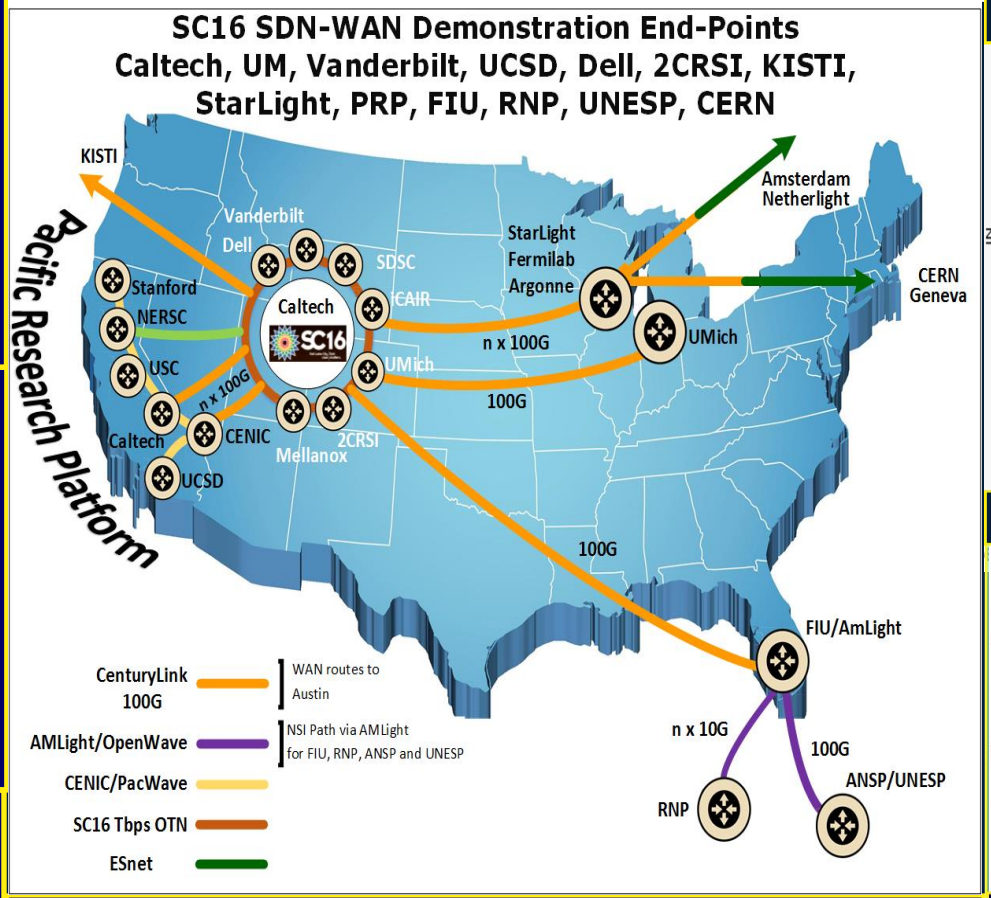

LHC at SC15: Asynchronous Stageout (ASO) with Caltech's SDN Controller

900 Gbps Total
Peak of 360 Gbps in the WAN

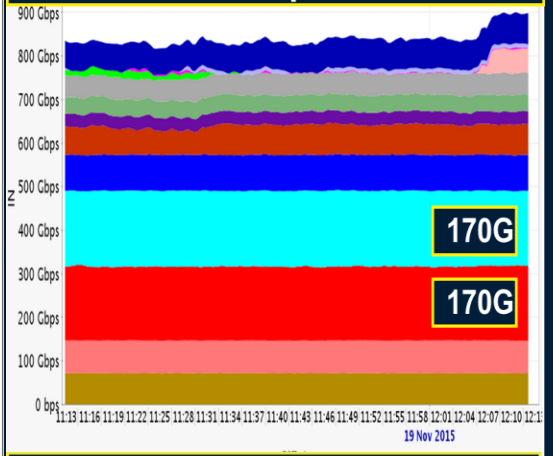

Global Topology

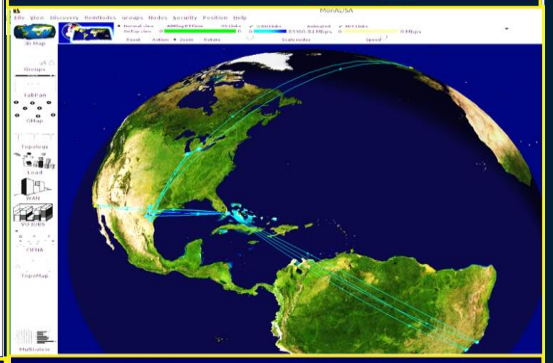

29 100G NICs; Two 4 X 100G and Two 3 X 100G DTNs; 1.5 Tbps Capability in one Rack; 9 32 X100G Switches

Tbps Ring at SC17: Caltech, Ciena, Scinet, OCC/ StarLight + Many HEP, Network, Industry Partners

Caltech and Partners at SC17

A Global Collaborative Scientific Enterprise

- ~3 Tbps each at the Caltech and Starlight/OCC booths
- 1.8 Tbps between the Booths and ~3Tbps out to the WAN
- *Connections to the 2CRSI, SDSC, UNESP booths*
- Caltech Booth: 200G dedicated to the Caltech campus; 300G to PRP (UCSD, Stanford, UCSC, et al); 200G to Brazil + Chile via Miami; 200G to ESnet
- Waveserver Ai + Coriant DCIs in the booths: *N X 100GE to 400G, 200G, 100G waves*
- *Intel State of the Art SSDs (P4608, P4500, P4800X) at the heart of it*

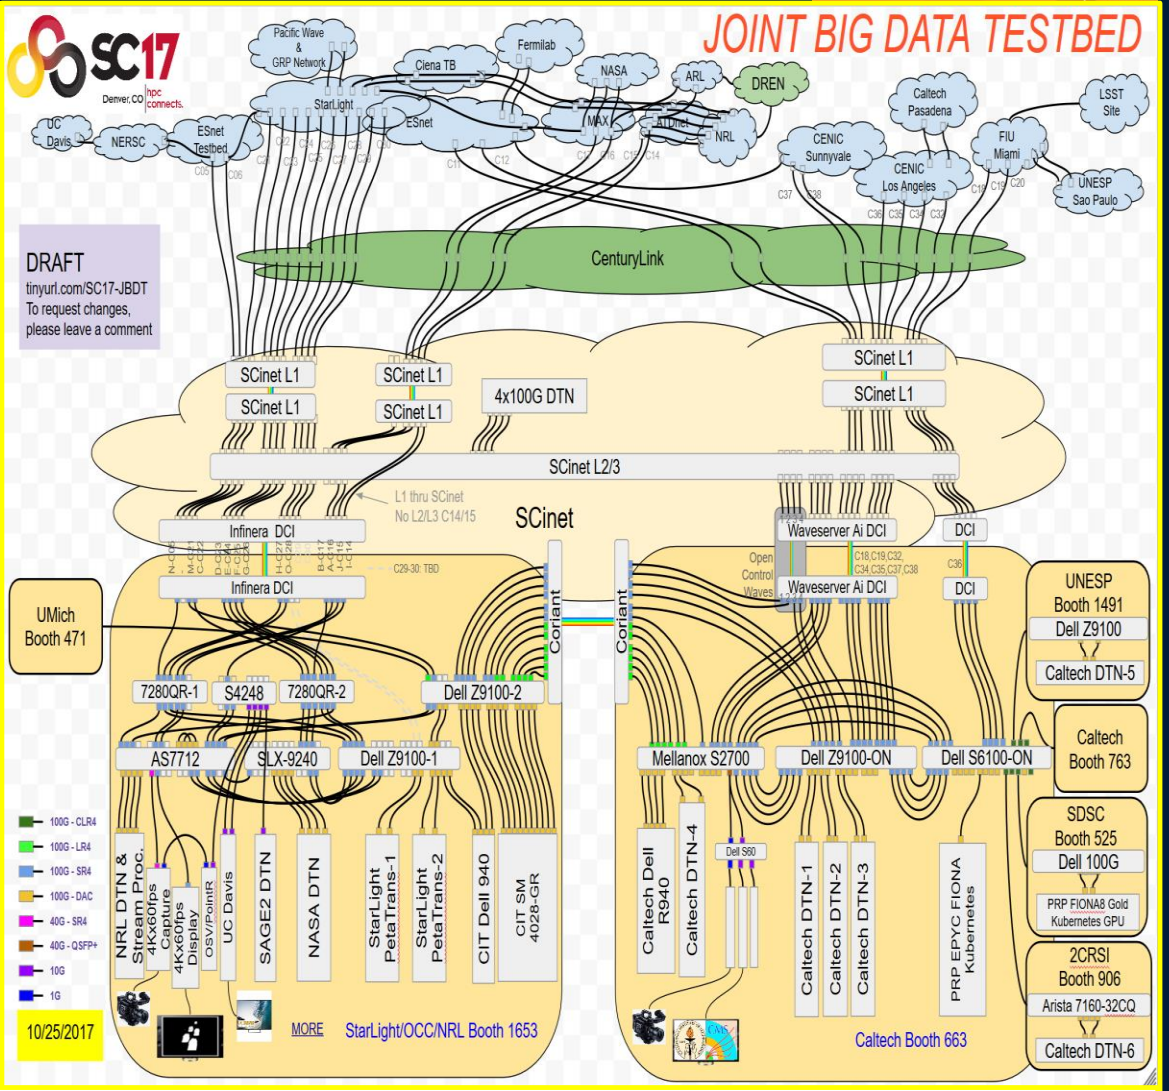

Microcosm: Creating the Future SCinet and the Future of Networks for Science

New Paradigms for Exascale Science

- Tbps SDN Intercontinental Steered Flows for LHC and LSST
  - Kytos Controller: LSST
  - Multicontroller Multi-Domain SDN
- SDN NGenIA and SENSE
  - LHC And Network Orchestrators
- HEPCloud with Caching; Dynamic Paths
- Steering with Real-time Analytics + Machine Learning
- Immersive VR with Ai Guide for HEP

Partners: Fermilab, PRP, USCD, Argonne, StarLight, NRL, FIU, UNESP, NERSC/LBL, CERN, ATT; SCinet, Ciena, ESnet, Internet2, CENIC, AmLight, RNP, ANSP; Dell, 2CRSI, Arista, Supermicro

Visit Caltech Booths 663, 763

SC17 Denver, CO hpc connects