Fermilab WAN Capabilities

Phil DeMar, Ray Pasetes, Andrey Bobyshev

ICAC Review

10/15/2019

Two Perimeter Security zones

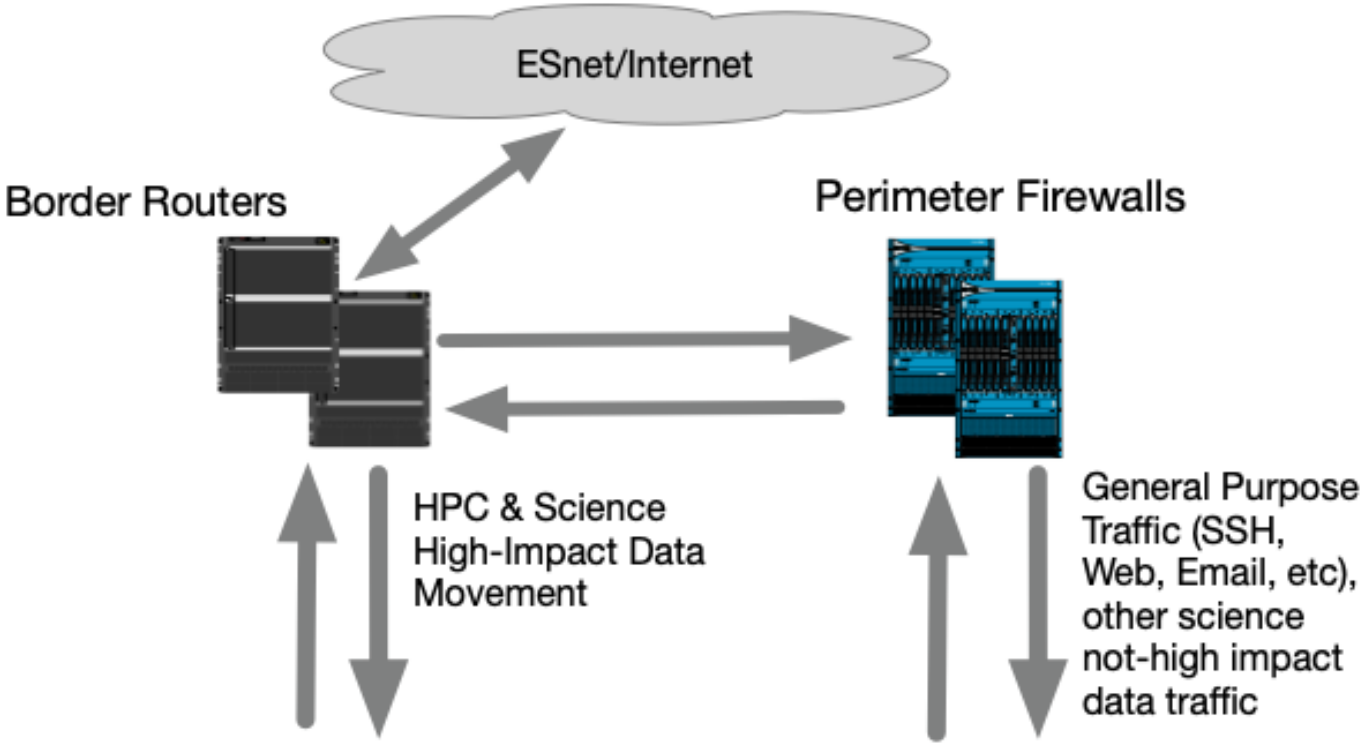

FY19: 34% completed (131 subnets)

FNAL Campus LAN & Off-Site Network Paths

External

Network Perimeter

Internal Network

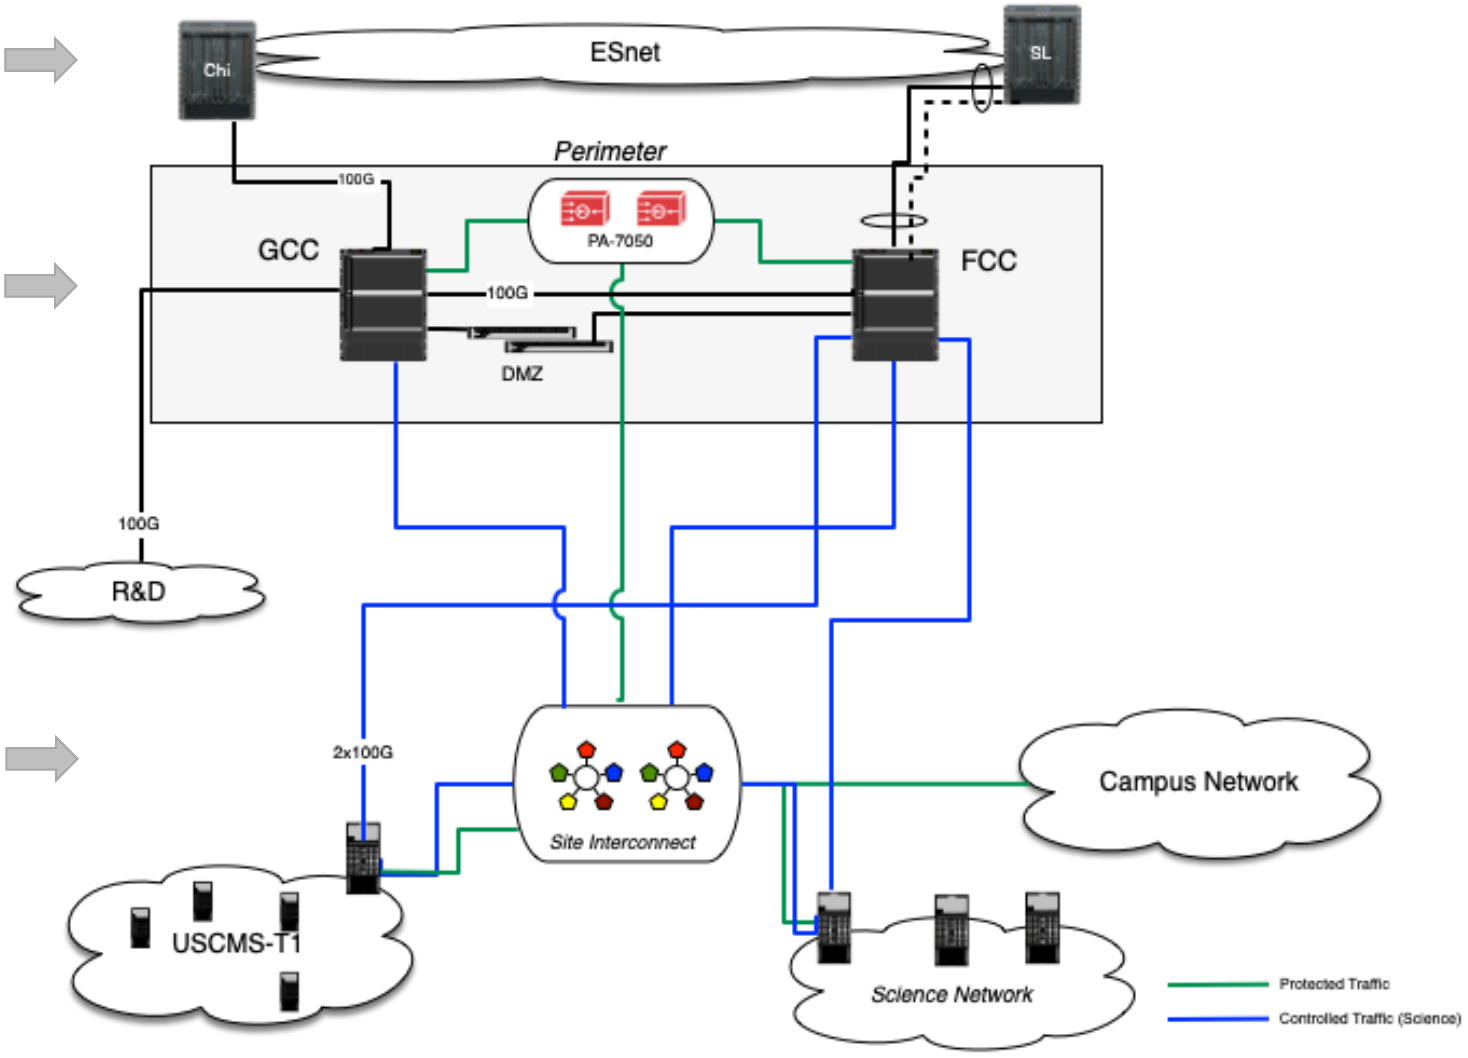

FNAL Network Perimeter

Perimeter has geographic diversity for both devices & fiber

Aggregate off-site bandwidth is currently 300Gb/s (3 x 100GE)

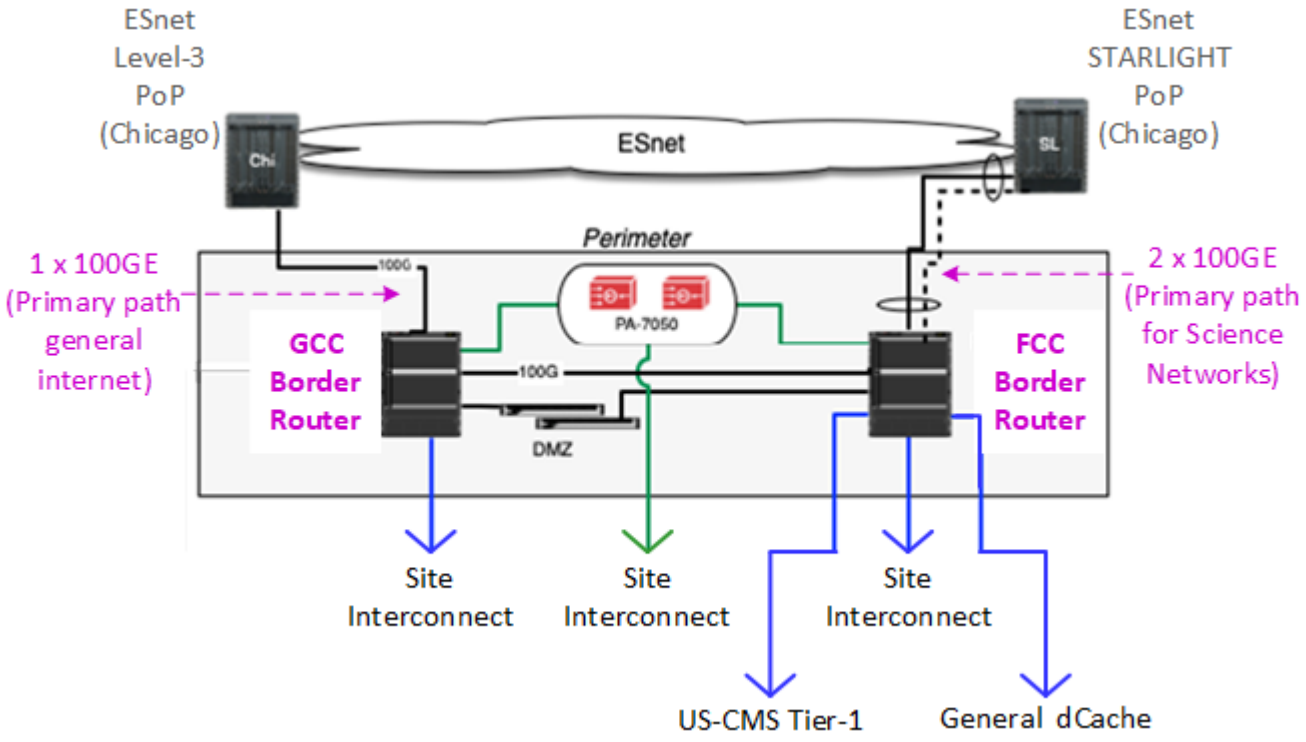

Virtual Routing & Forwarding (VRF) technology on border routers

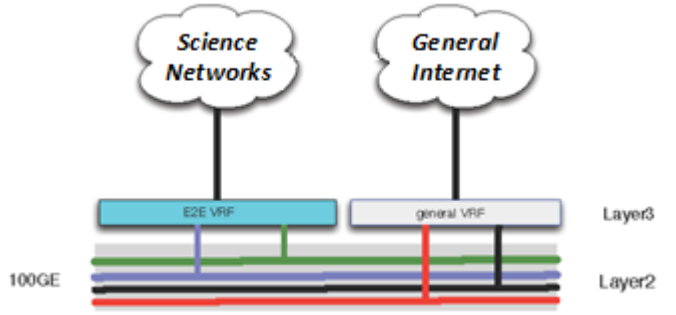

ESnet – FNAL's ISP

- ESnet is FNAL's only offsite internet service provider
- ESnet backbone is extremely capacious & highly redundant:
  - Including trans-Atlantic services to CERN & GEANT
- Metropolitan area network (MAN) services for last mile(s)

ChiExpress – FNAL's Connectivity to ESnet's PoPs

- A local optical network operated by ESnet
- Ethernet-based, not SONET-based:
  - Wave reconfiguration within (O) 10 minutes
- Site capacity upgrades involve only hardware

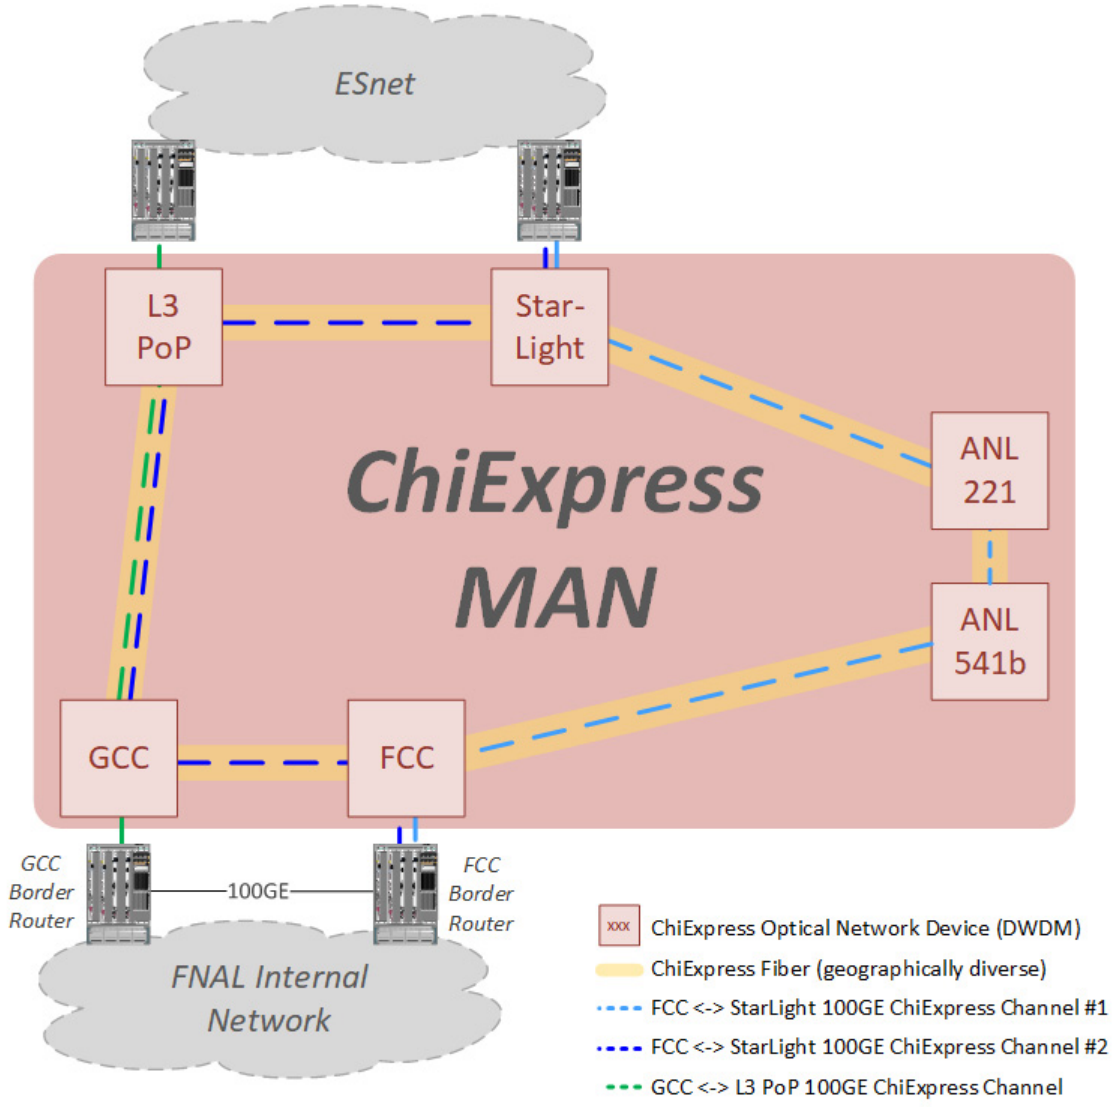

Special-Purpose Science Data Networks

- General strategy – keep science data traffic logically separate, if feasible:
  - Optimize performance over WAN science data network paths
  - Internal LAN upgrades for high-impact science requirements easier to target
  - More flexible perimeter security models
  - Limit interaction with interactive traffic:
    - Don't want this to be your Zoom meeting data stream…

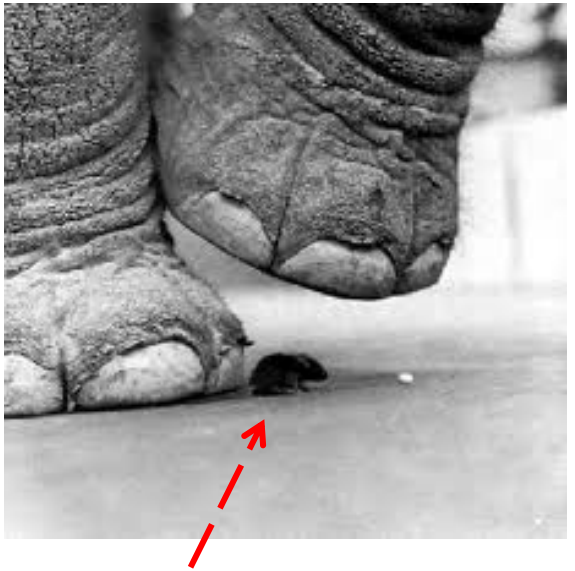

- Examples – LHCOPN, LHCONE, NOvA Far Detector

LHC Optical Private Network (LHCOPN)

Star-shaped “private” network:

- CERN (LHC Tier-0) at the center
- “Links” to the LHC T1s
- T0 <-> T1 and (some) T1 <-> T1 traffic

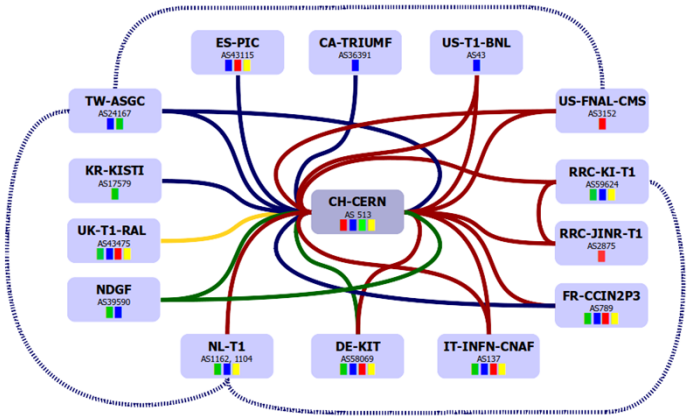

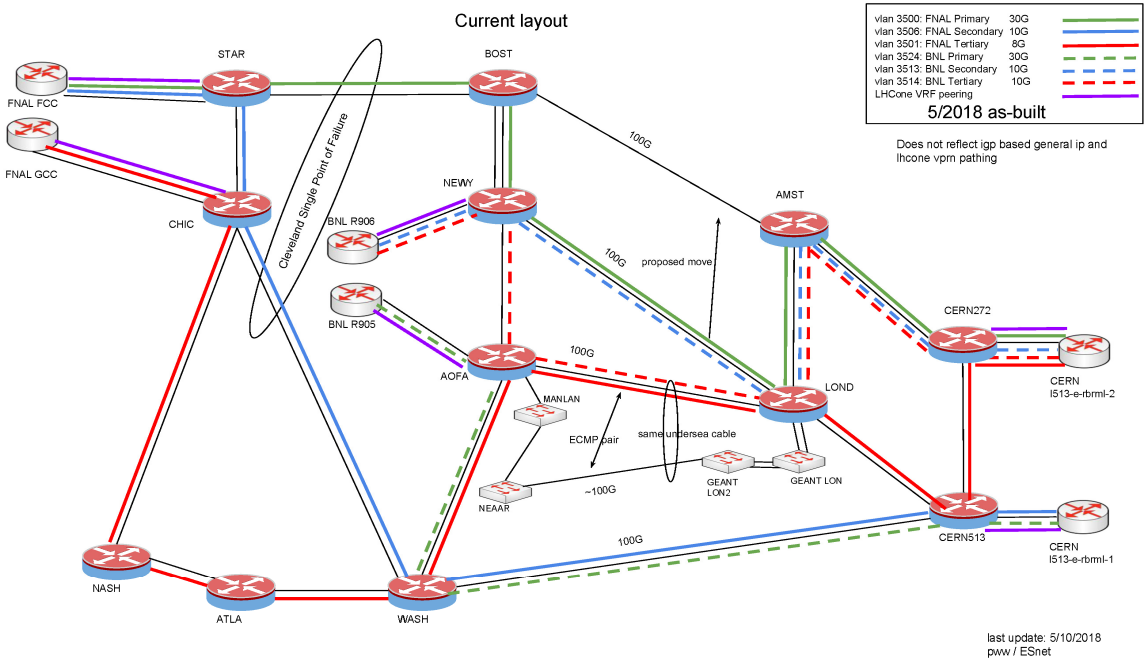

FNAL LHCOPN links are virtual circuits over ESnet:

- B/W guarantees
- Primary, secondary, & tertiary circuits

Fermilab

LHC Open Network Environment (LHCONE)

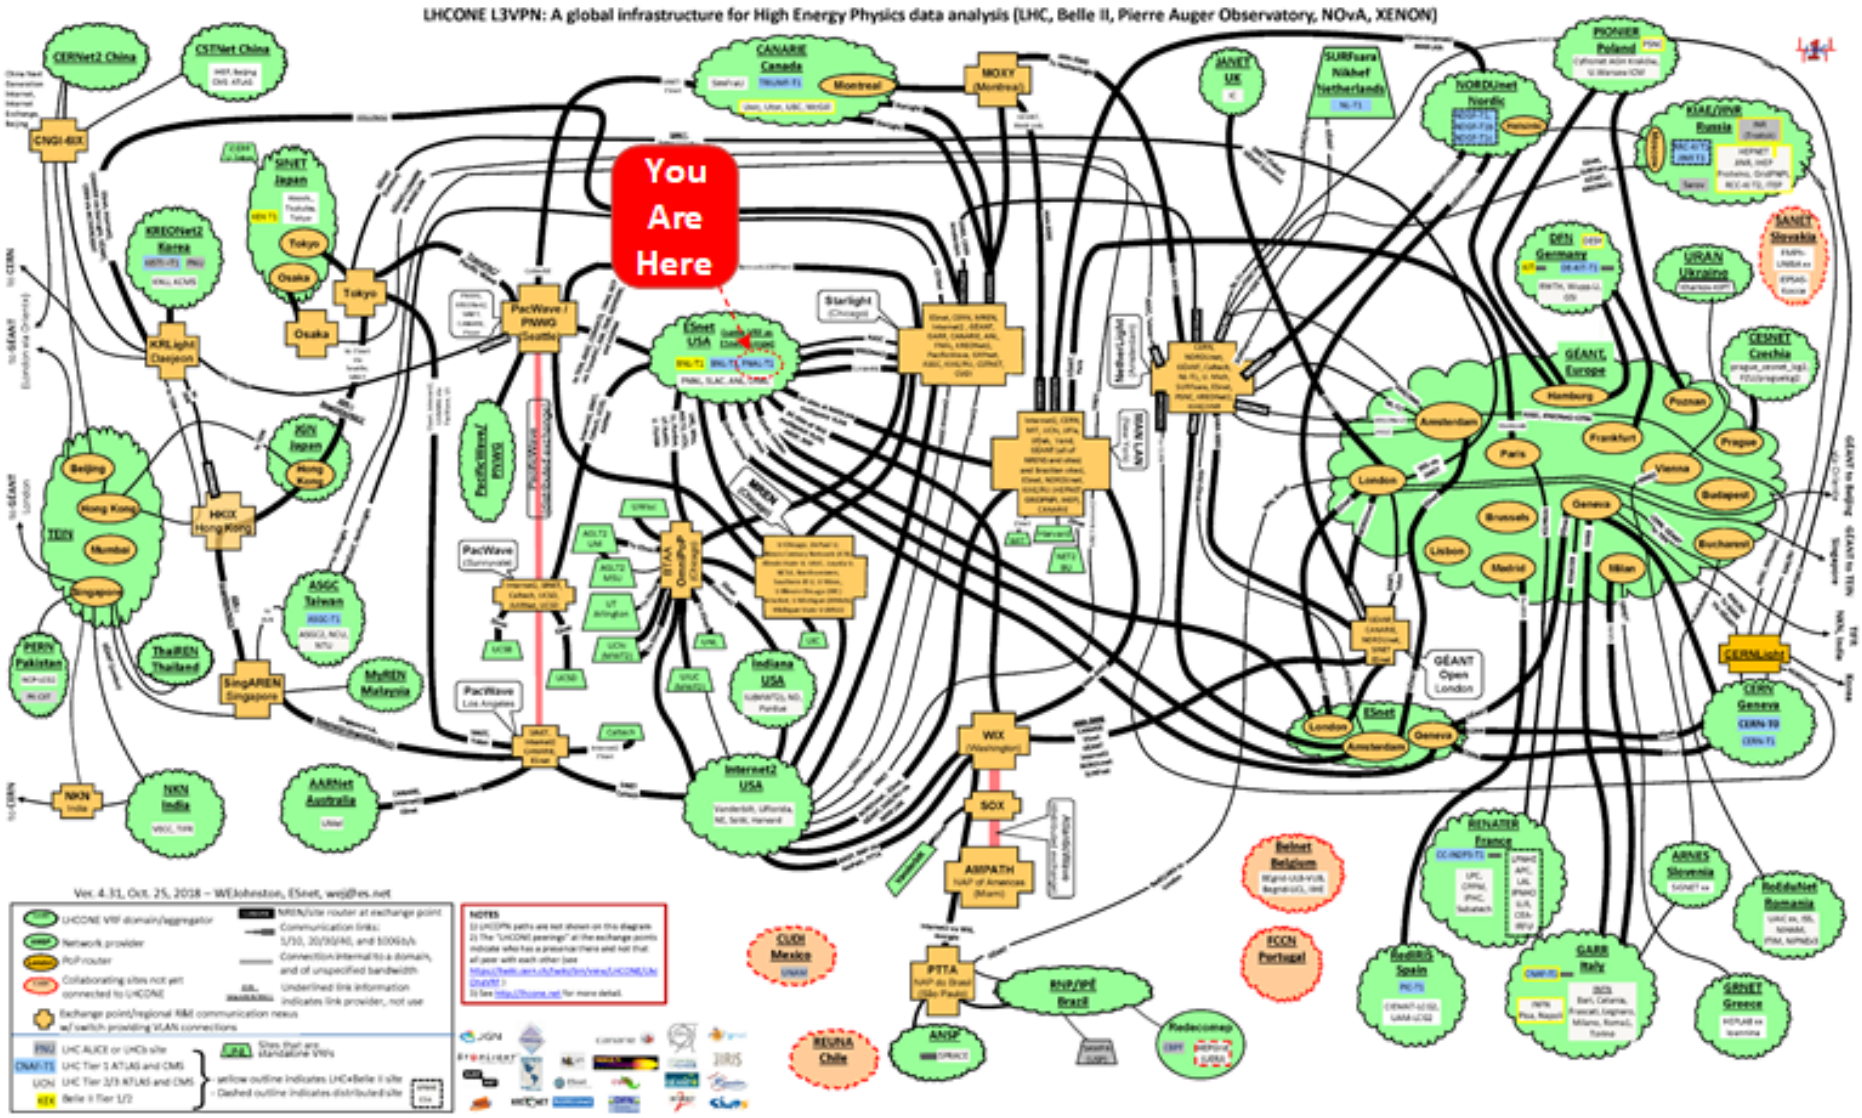

Fermilab

LHCONE Translated into Understandable Terms

- LHCONE is an overlay network (in most places…):
  - VRF technology typically used to separate routing
- Logically separates LHC traffic from general R&E network traffic for traffic management purposes:
  - Also can lower security risk profile at sites for LHC traffic
- End sites are responsible for separating traffic:
  - At FNAL, we use Policy Routing techniques to accomplish that
- There (currently…) are no QoS capabilities within LHCONE
- Several smaller HEP/Astro collaborations have been granted authorization to utilize LHCONE for their data movement:
  - Belle-II, Pierre Auger, NOvA, XEON

Point-to-Point Network Services

- NOvA Far End Detector (FD):
  - Redundant low-bandwidth pt-to-pt connection(s)
    - vLANs with B/W guarantees plumbed through multiple service providers
  - Logically, NOvA FD lies inside FNAL's security perimeter

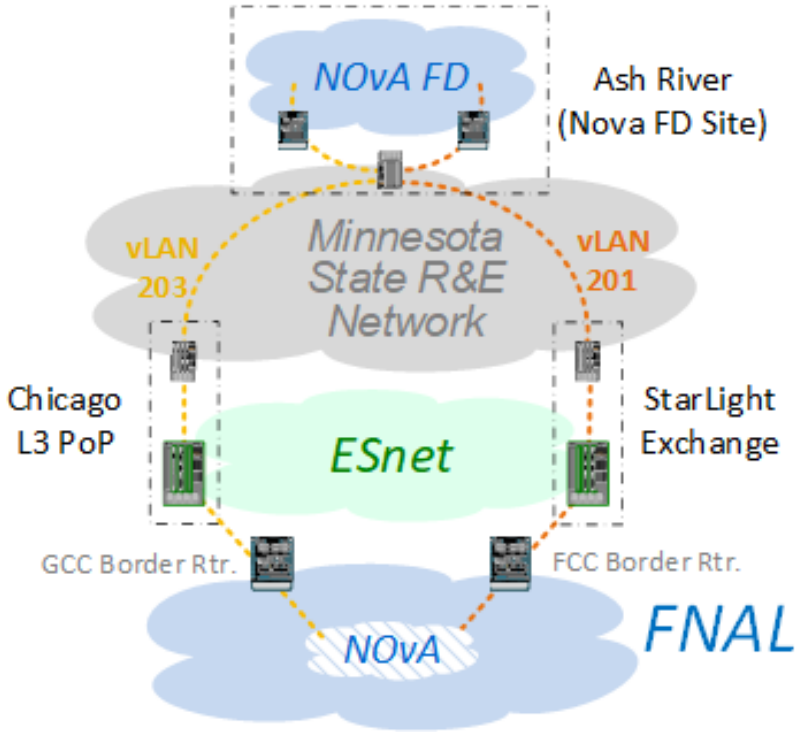

- Pt-to-pt circuits to US-CMS Tier-2 facilities:
  - UFL, MIT, Nebraska, Purdue, UCSD, Wisconsin
  - Being phased out,; largely replaced by LHCONE

DUNE Far Detector (in South Dakota)

- Far Detector (FD) connectivity to FNAL:
  - FD local network will lie within FNAL security perimeter
  - WAN support model currently under discussion w/ ESnet
  - B/W under study; targeting 100GE path
    - Redundant path expected to be much, much less

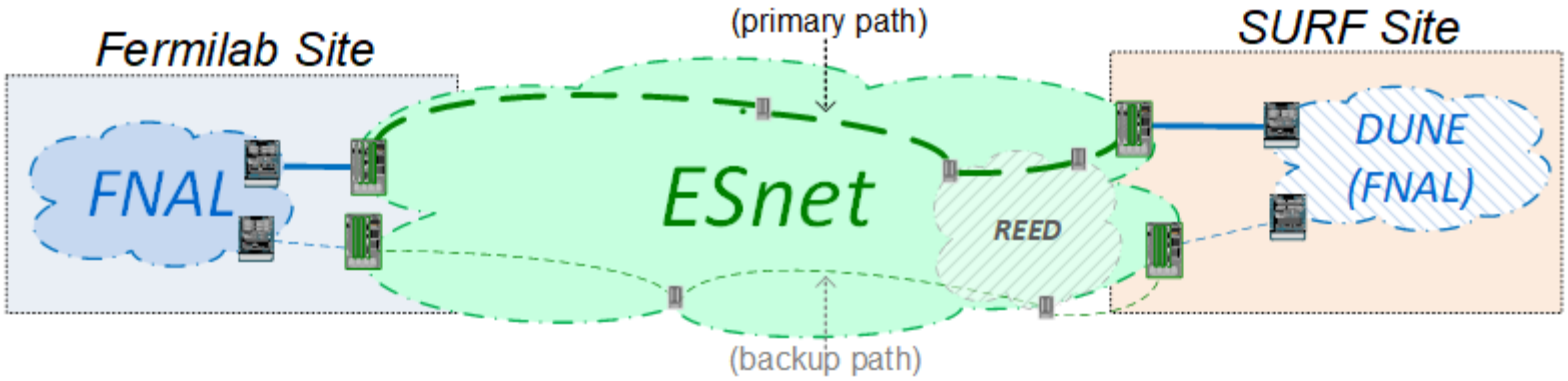

- Off-line DUNE network expected to look similar to LHC:
  - LHCONE-like service (but probably not an LHCOPN-like service)

Large Synoptic Survey Telescope (LSST)

- Dedicated network infrastructure from telescope (Chile) to archive site (NCSA)
- FNAL working with NCSA on operational support for the long-haul network:
  - Experiences from LHC networking believed to be helpful

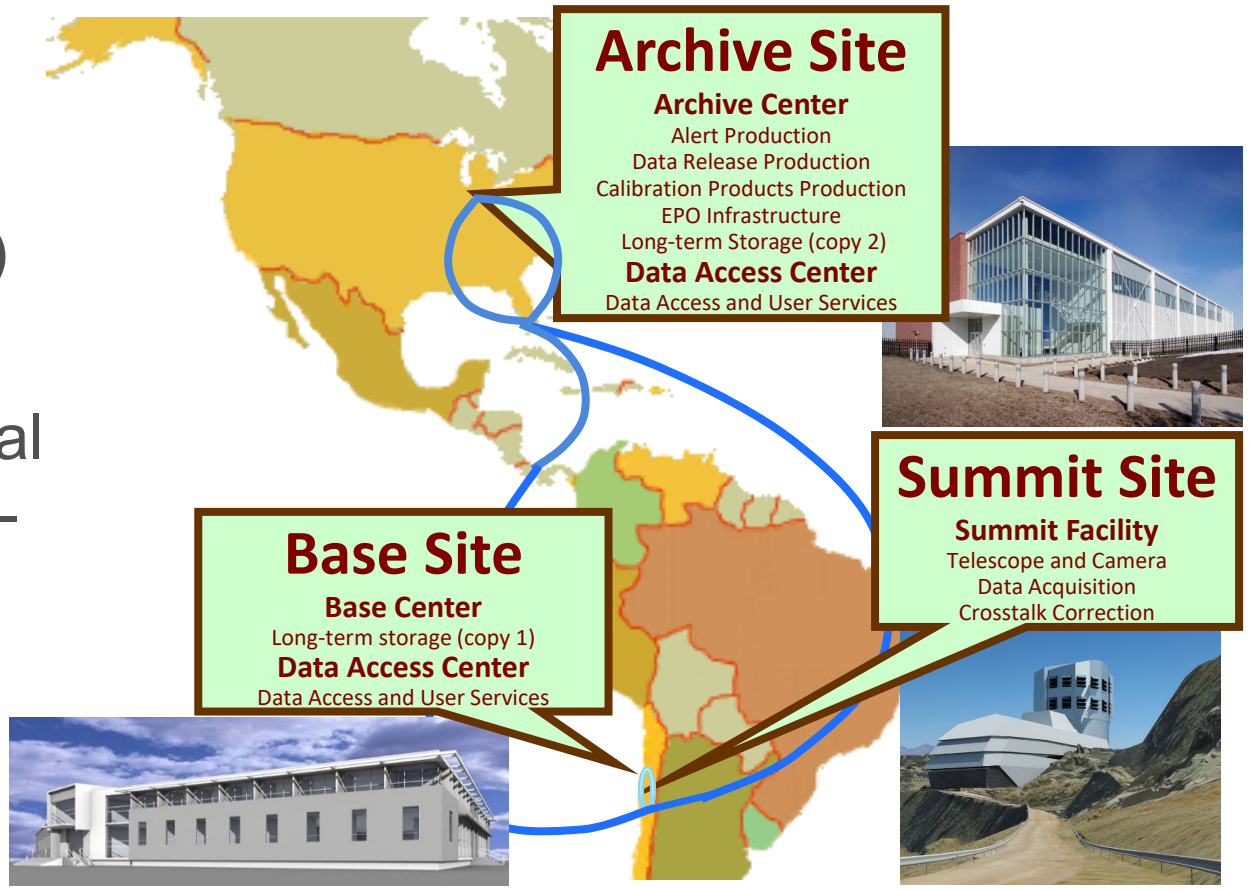

Fermilab

Metropolitan Research & Education Network (MREN)

- Consortium of Chicago area research universities & National Labs (FNAL & ANL):
  - Like a mini-regional network
  - Universities directly connect
  - The Labs indirectly via ESnet

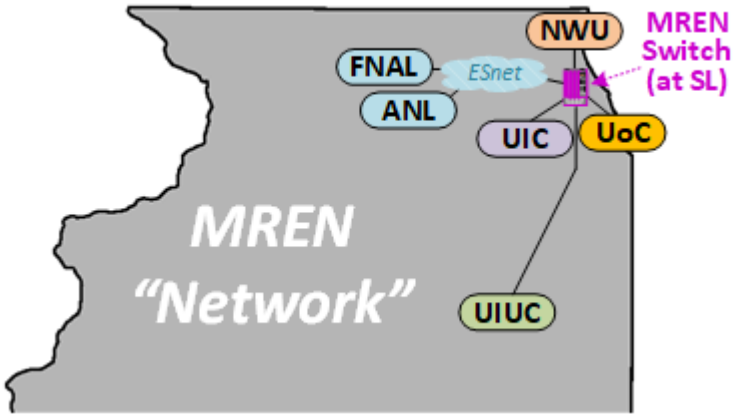

- Maintains network switch(es) at the StarLight International Exchange:
  - Potential connectivity path to StarLight-connected networks

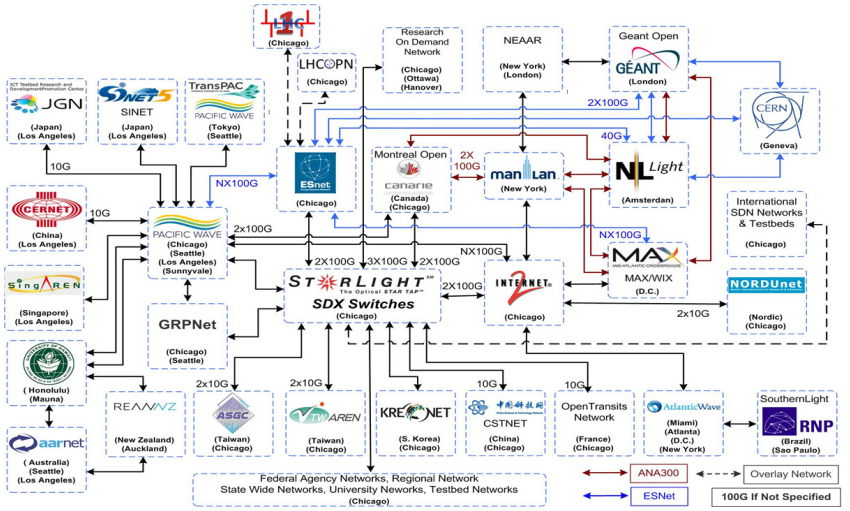

Fermilab

ESnet6

- The next generation of ESnet:
  - 2021-ish
- ESnet will “own” fiber, optical infra., and routers
- Topology similar to ESnet5
- Trans-Atlantic & intra-European links still leased services

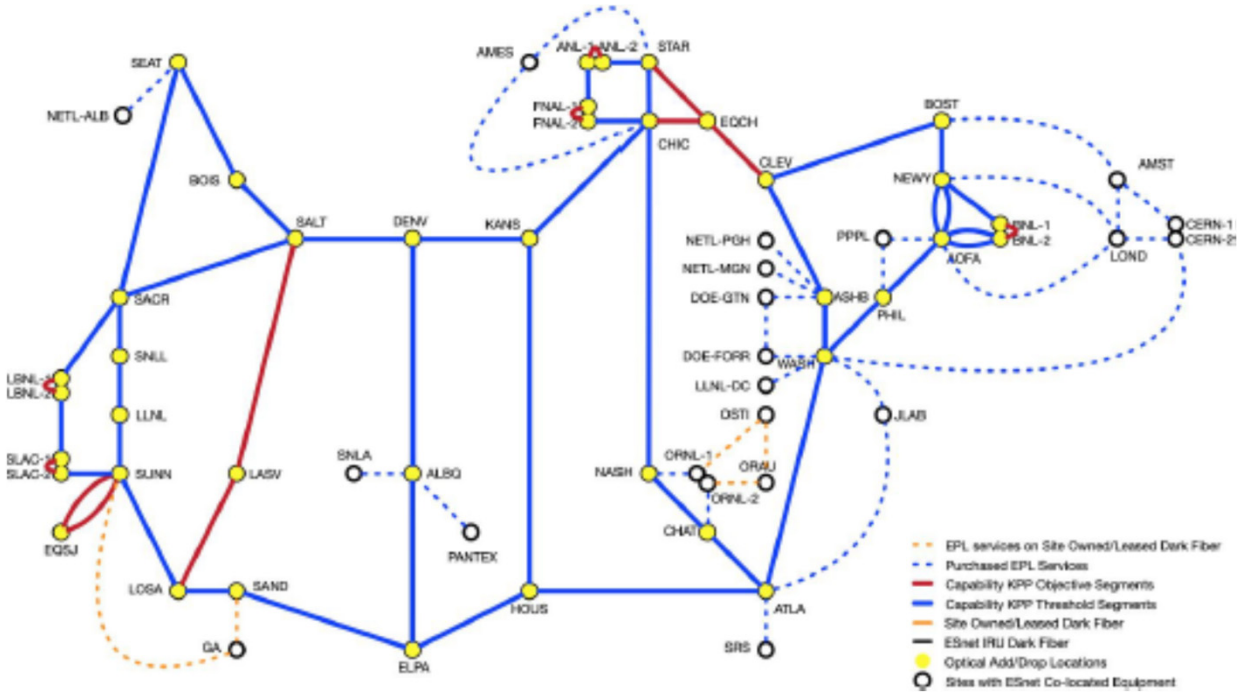

Fermilab

ESnet6 – New Networking Model

- "Smart" edge and "hollow" core
- Network services pushed to sites
  - Service & path decisions made at the edge
- Core's function is packet forwarding
  - One routing hop across ESnet

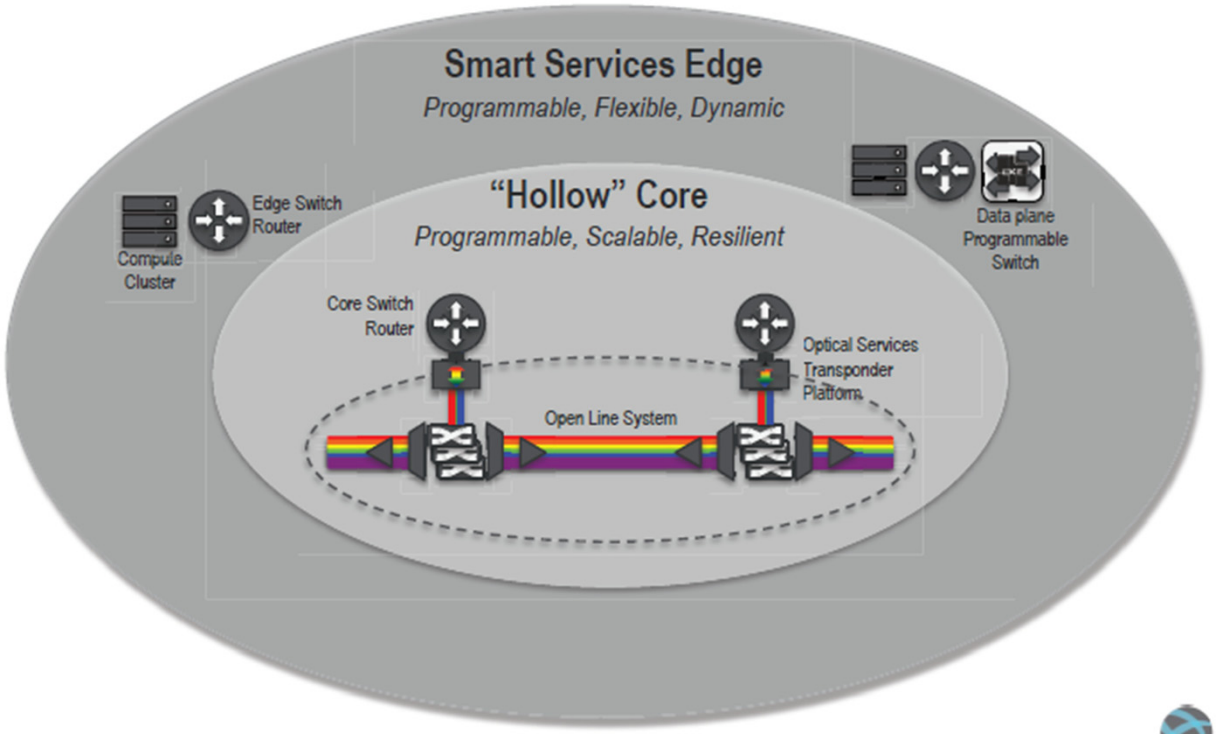

ESnet6 – Multiple Classes of Services

- "No touch" services are essentially default IP routing
- "Low touch" services are today's LHCOPN & LHCONE-like services:
  - More flexible, dynamic, & more automated
- "High touch" services allow per packet tagging & manipulation
- Service classes isolated from each other

No Touch Services

Low Touch Services

High Touch Services

Summary

- FNAL WAN services have capacious bandwidth & high levels of resiliency
  - Needs for more bandwidth can be satisfied relatively easily
  - But trans-Atlantic bandwidth demands present challenges
- "Special" handling for science data traffic is provided:

Network Service	Service Characteristics	Example(s)
Routed IP	General internet access	ESnet base service
"Private" network	Small # of tightly-restricted sites Bandwidth guarantees	LHCOPN
Science discipline network	Large # of loosely-connected sites No bandwidth guarantees	LHCONE
Pt-to-Pt circuits	Deterministic, secure path(s) Bandwidth guarantees	NOvA FD; DUNE FD StarLight (ie., Intl. networks…)

Questions