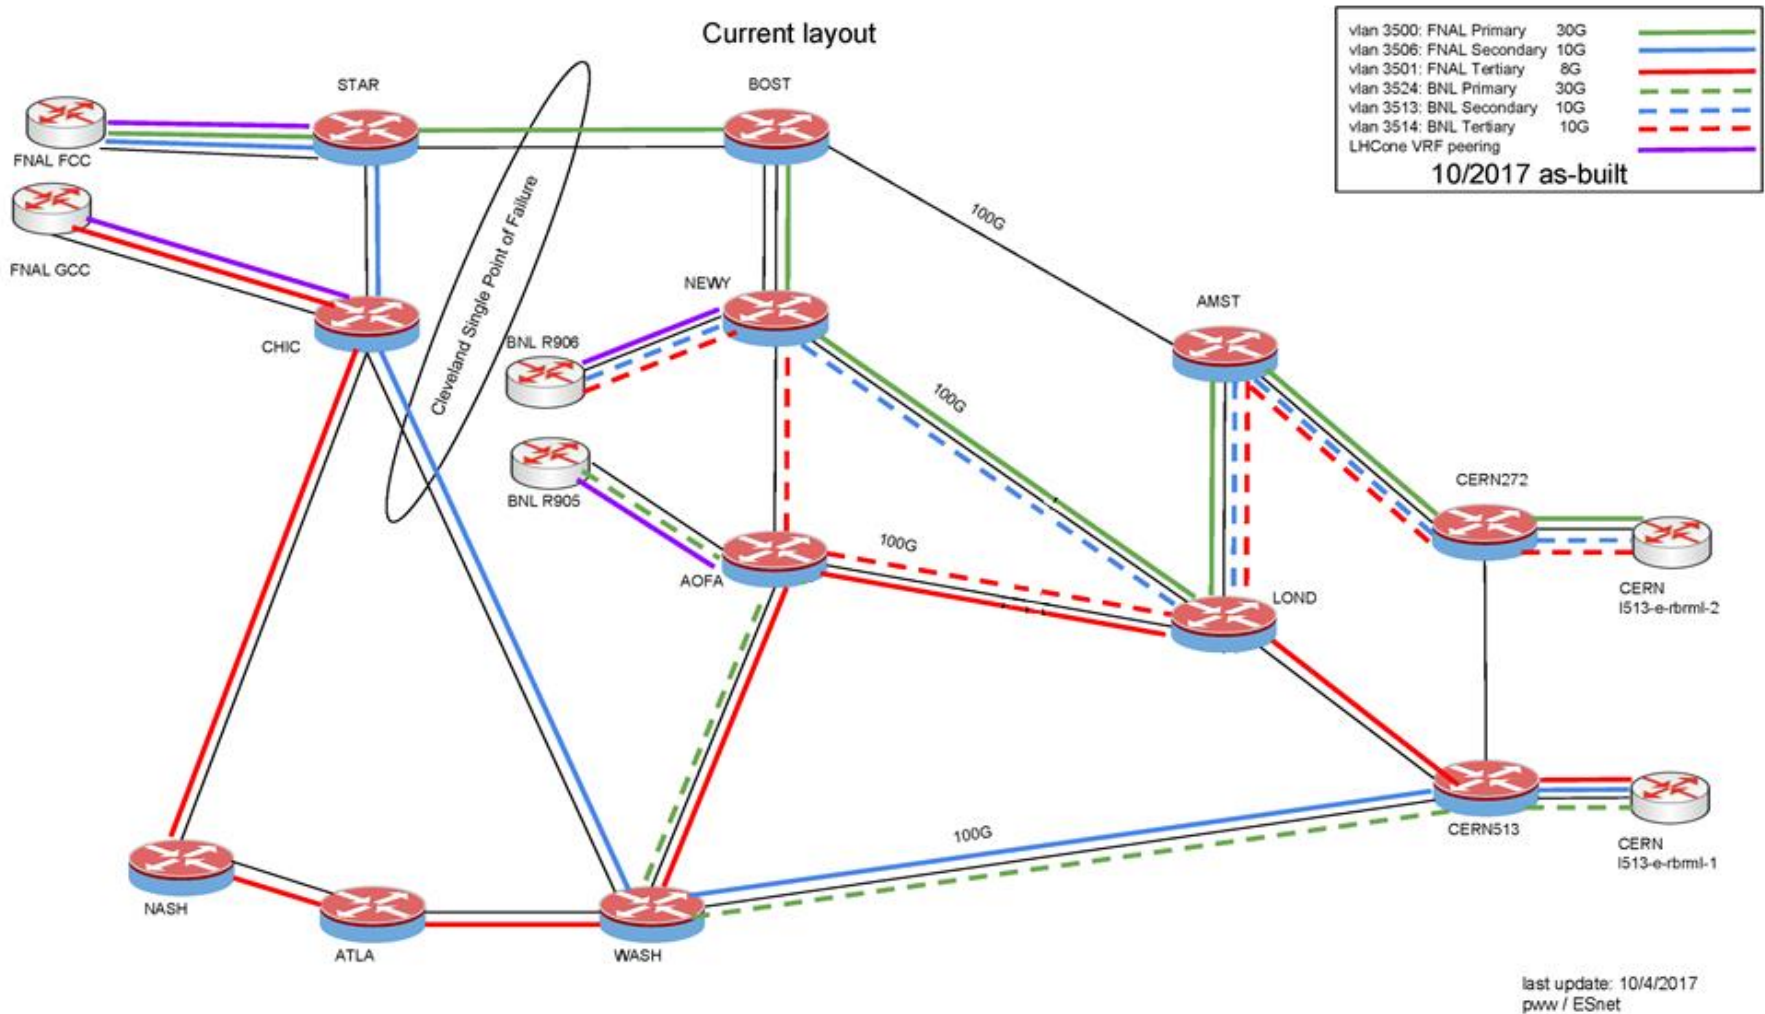




QoS for T0 \leftrightarrow T1 Traffic on the LHCOPN (w/ ProtoDUNE traffic request)

Phil DeMar (FNAL)
RAL LHCOPN/LHCONE Meeting
March 5, 2018

“Current” US-LHC T1 LHCOPN Config.




Current US-LHC T1 LHCOPN Config. (II)

- US-LHC T1 OPN Circuits:
 - Primary circuit: Default path; carries all T1's OPN traffic
 - 30GE b/w guarantee
 - Secondary circuit: Failover path when primary circuit is down
 - 10GE b/w guarantee
 - Tertiary circuit: (when primary & secondary are down...)
 - 8GE b/w guarantee
- FNAL T1 secondary OPN circuit carried on same TA link as BNL T1 primary OPN circuit
 - And visa versa...
 - Changes here expected to make more effective use of the 4th ESnet TA link

Discussions on US-LHC OPN Changes

- Two OPN design changes have been discussed:
 - Load-balancing model for OPN circuits:
 - Replacing existing primary/failover model
 - B/W guarantee reduction(s) on individual circuits
- OPN b/w guarantee cuts raise interest in QoS (for FNAL...)
 - T0 → T1 data is only OPN traffic with latency “concerns”
 - Complication for FNAL OPN services is potential for “new” traffic source (ProtoDUNE...)
- Current FNAL OPN traffic levels are ‘busy’, but OK



Previous QoS support on the OPN

- Priority for T0 → T1 traffic previously implemented by CERN on earlier generation of their LHCOPN routers (Force10)
 - Service not migrated to their current generation of LHCOPN routers (Brocade)
- Worth considering again?
 - If so, upon request or as a default configuration?

ProtoDUNE

ProtoDUNE

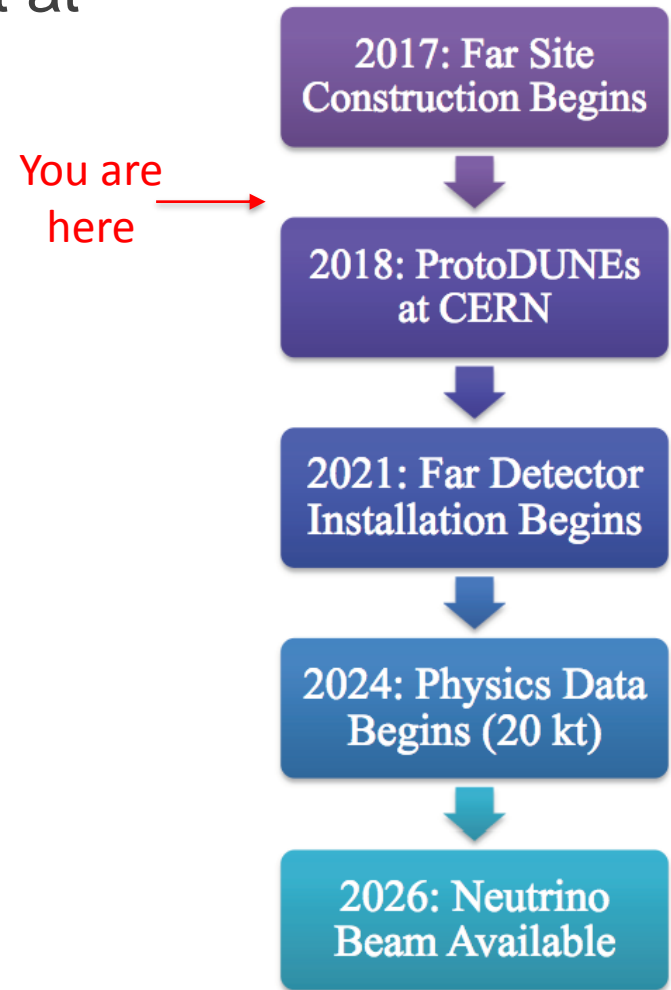
- Two prototype DUNE detectors built at the CERN Neutrino Platform:

The CERN Neutrino Platform



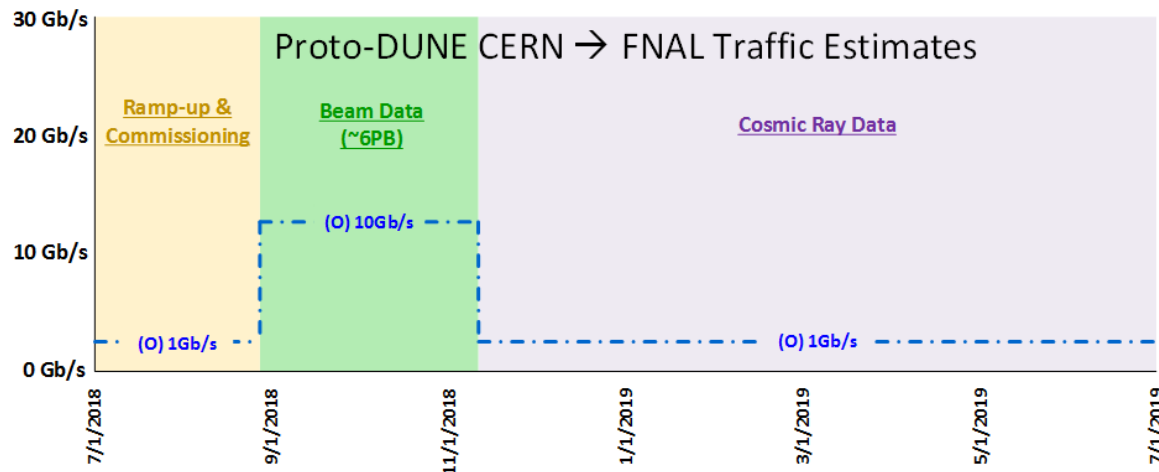
- Low-intensity beam provided by CERN SPS accelerator complex
 - Starting end of August (2018)
 - Run for ~10 weeks
- Cosmic Ray run will follow:
 - ~30-34 weeks
 - Much lower data rates

DUNE Timeline



Data Movement for ProtoDUNE

- Pre-processed data will be streamed from CERN EOS (T0) to FNAL dCache (not T1) for disk storage/archiving:
 - Disk/storage & archived copy retained at CERN as well
- Estimated data generated during beam run= ~6PB
 - Translates to (O)10Gb/s traffic level at peak intervals
 - (O) 1Gb/s during commissioning & cosmic ray runs
 - Latency not considered an issue



ProtoDUNE Traffic across the OPN?

- ProtoDUNE traffic could be routed to FNAL in many ways:
 - LHCOPN
 - LHCONE
 - P2P circuit
 - General routed path
- All “paths” use the same ESnet network infrastructure
- From a simplicity perspective; the OPN is preferred option
- Traffic would only impact CERN → FNAL segment of OPN:
 - FNAL would need to advertise it’s public dCache subnet (not part of the US-CMS T1)
 - No-export BGP Community would limit advertisement to CERN