ESnet
ENERGY SCIENCES NETWORK

ESnet Status Update

Paul Wefel
Network engineer
Lawrence Berkeley National Laboratory

SAACC - Virtual meeting

April 2021

U.S. DEPARTMENT OF ENERGY
Office of Science

BERKELEY LAB

Agenda

- What is ESnet
- ESnet6 Project Implementation Update
  - Optical Core
  - Packet Core & Low Touch Edge Details
- Automation
- High Touch services
- Questions

ESnet

ESnet: DOE's High Performance Network (HPN) Scientific User Facility and derives its mission from Office of Science

Mission of DOE Office of Science:
...delivery of scientific discoveries and major scientific tools to transform our understanding of nature...

Mission of Energy Sciences Network:
Science network user facility designed to **accelerate** scientific research and discovery.

ESnet

Connects all of the DOE national labs, many DOE sites, and hundreds of research and commercial networks internationally.

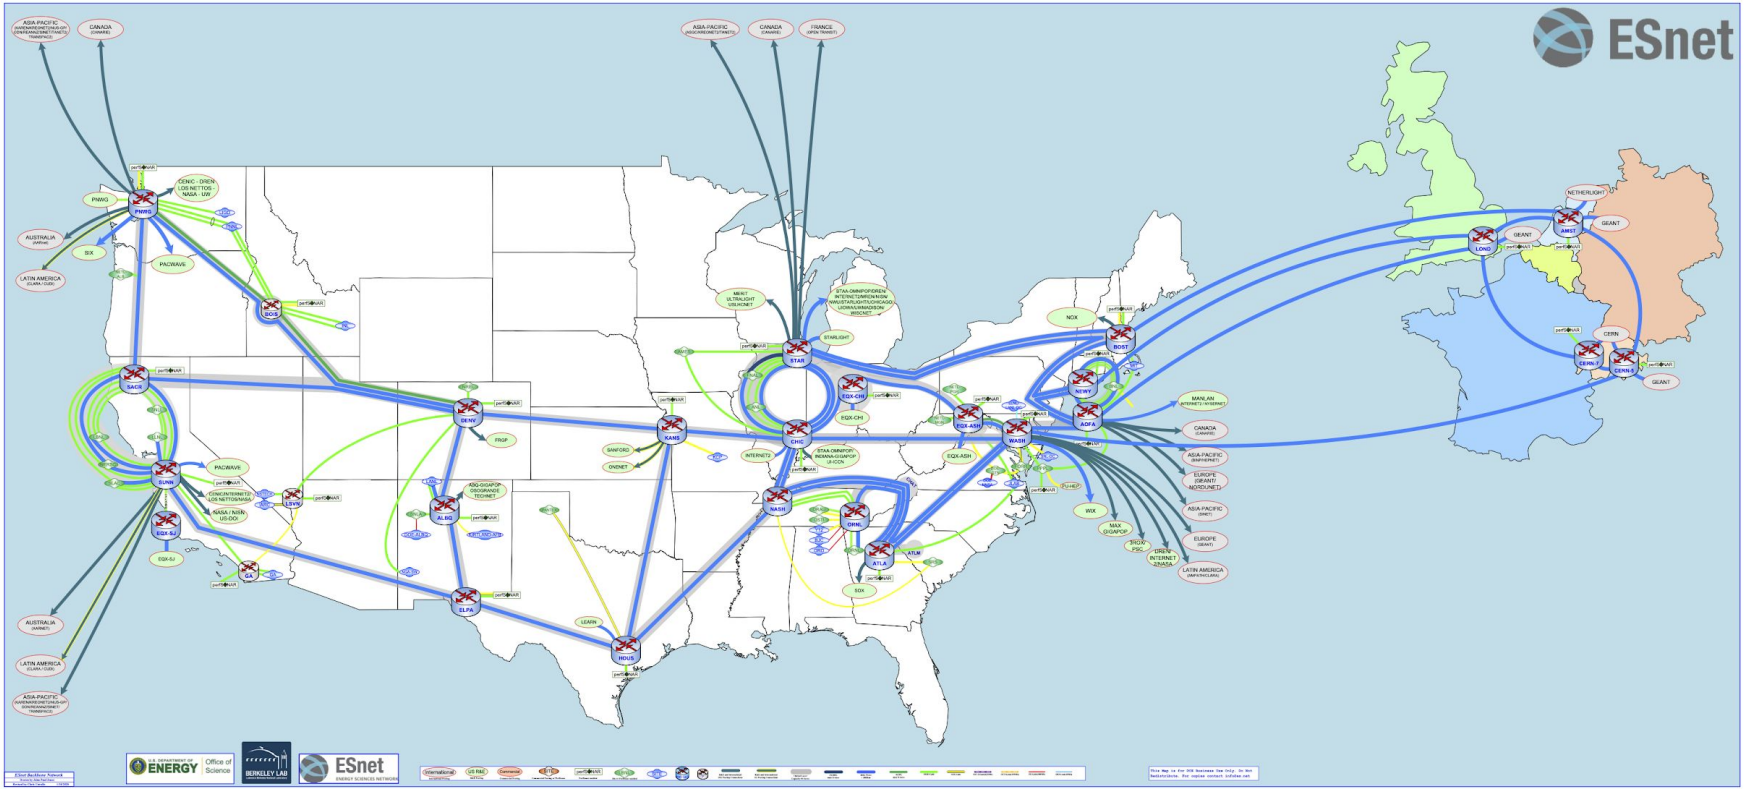

An Exabyte Network Today

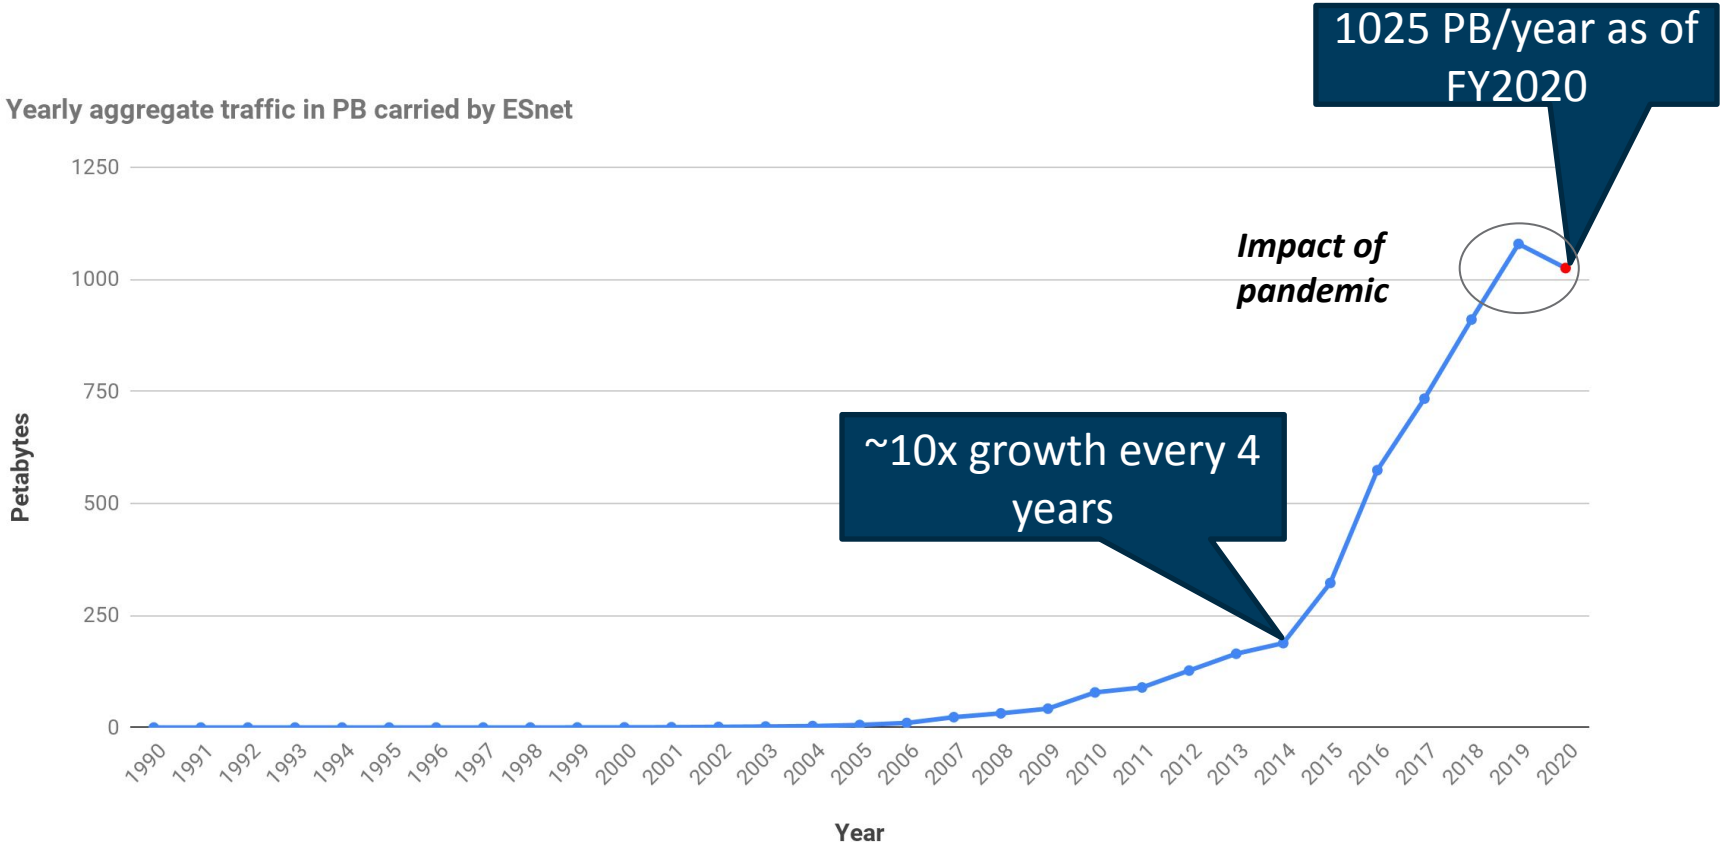

All Sites Supported in Their Transition to Work-from-home

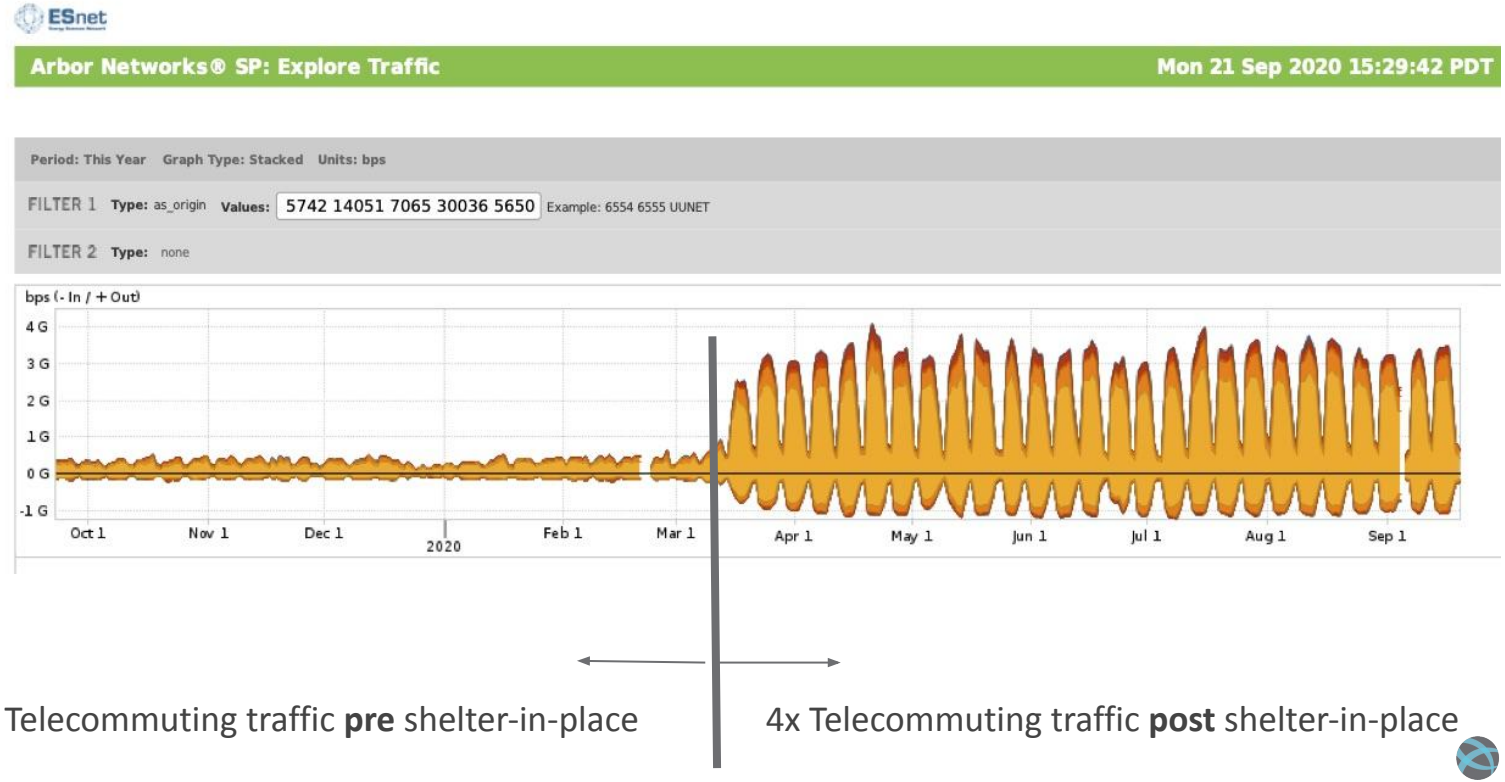

Telecommuting traffic **pre** shelter-in-place

4x Telecommuting traffic **post** shelter-in-place

The evolution of the ESnet network

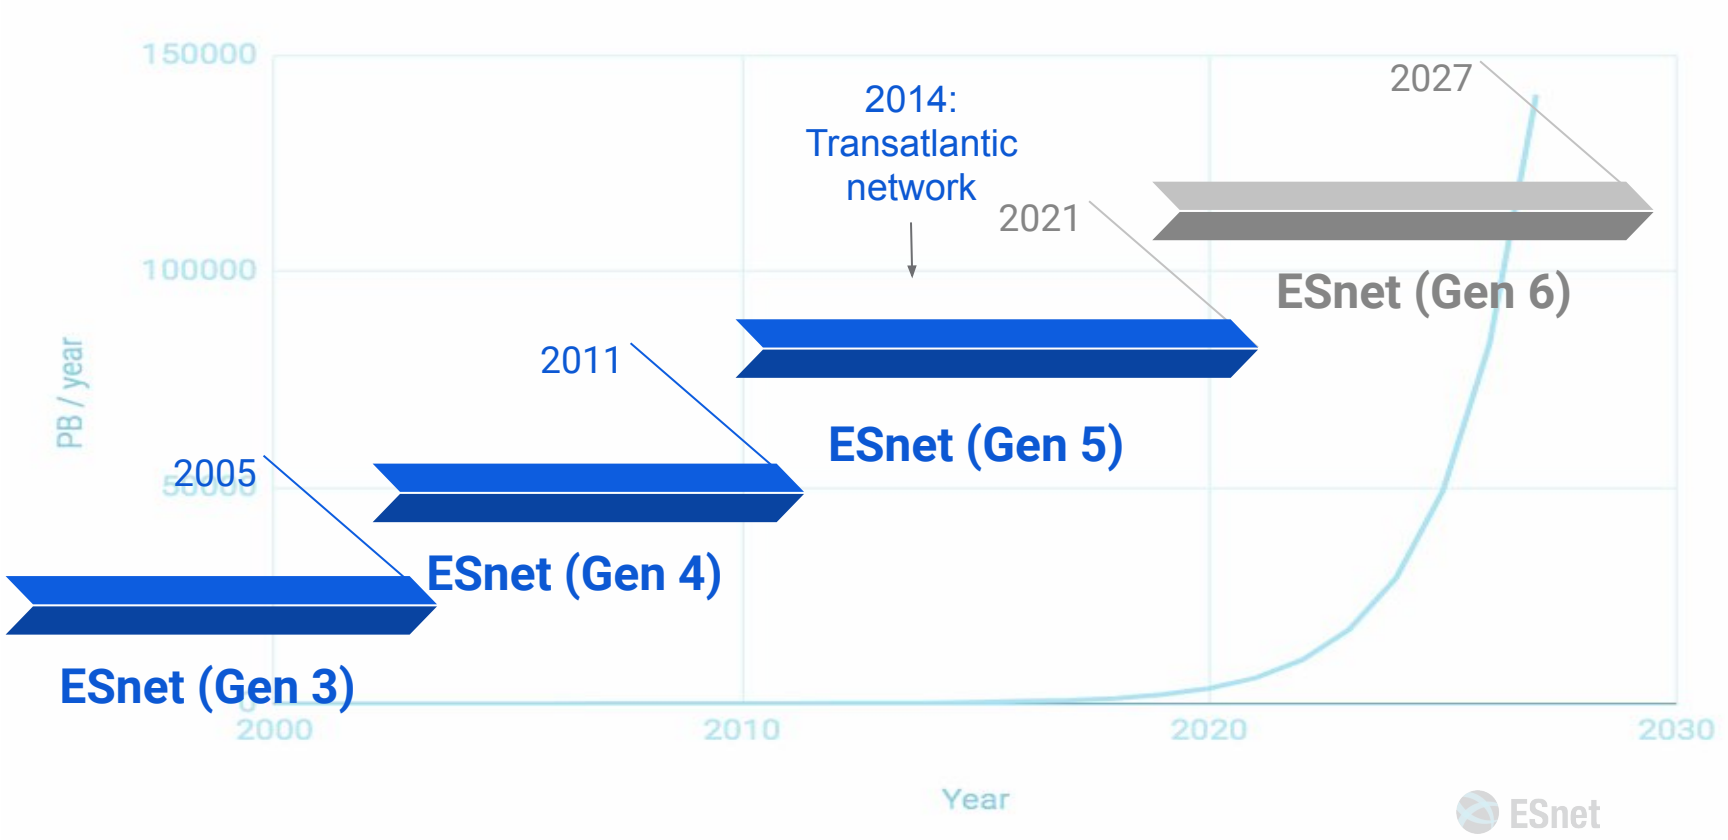

Implementation Context

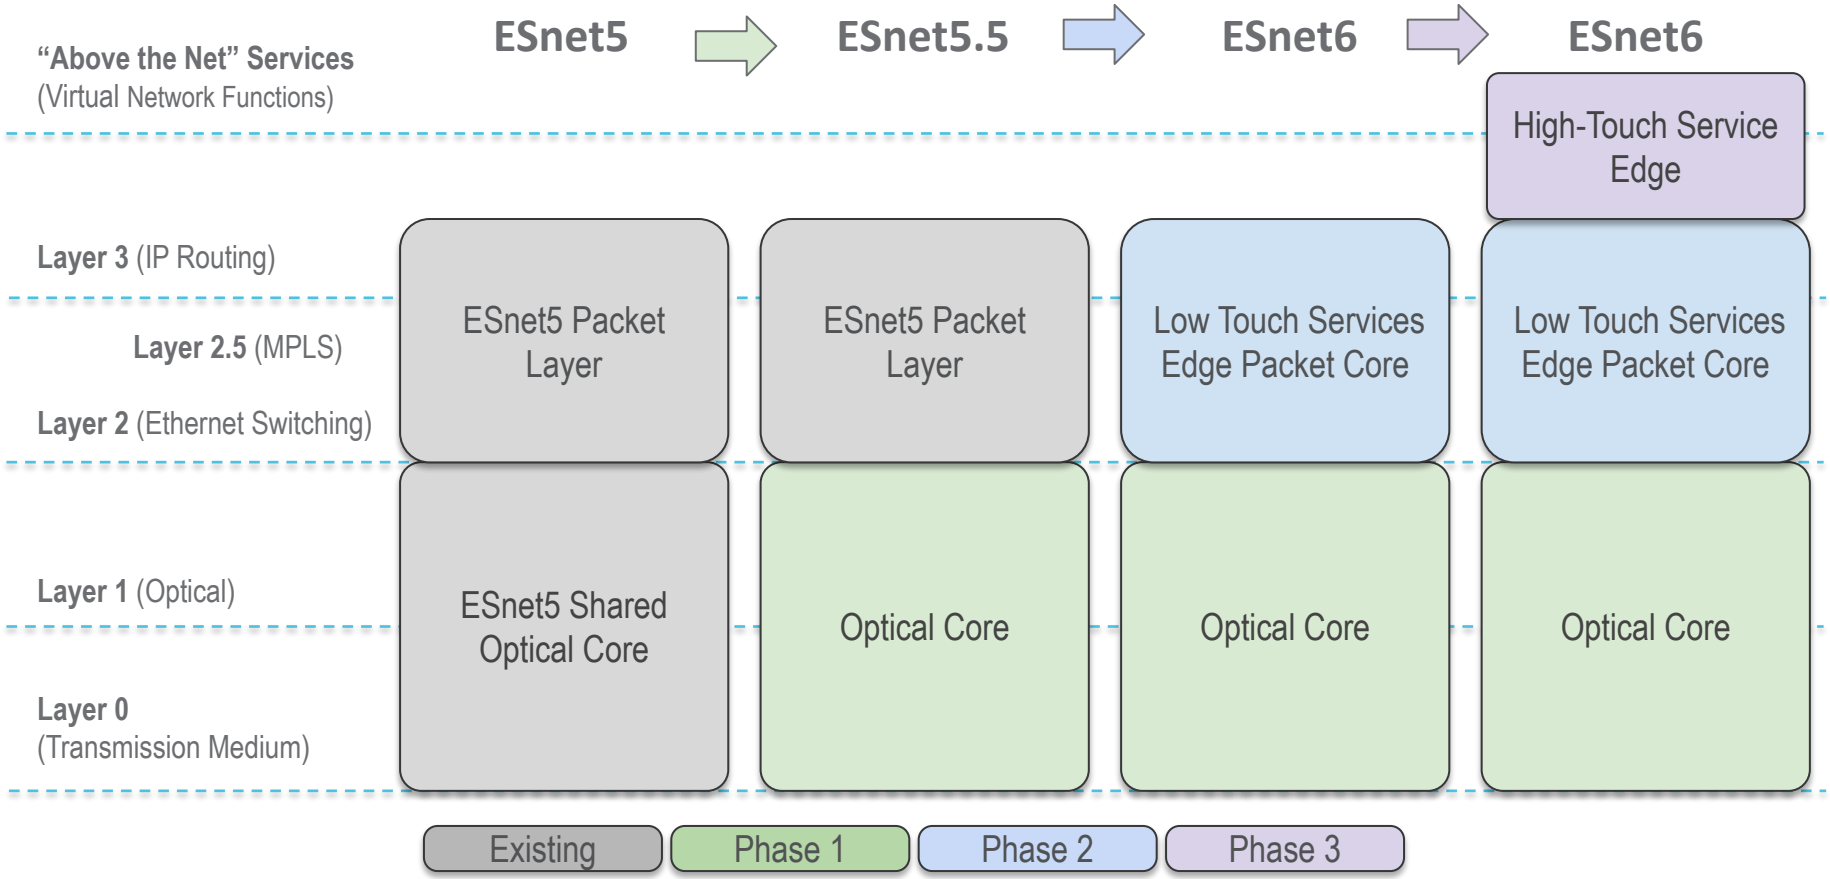

“High-Touch” vs “Low-Touch” Hardware

“High-Touch”
Programmable data-plane

Pros:
- Flexible to customize for specialized use cases

Cons:
- Complexity of designing and implementing specialized use cases
- Higher cost

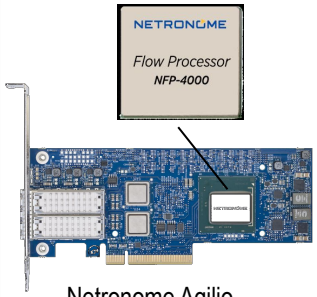

Netronome Agilio CX SmartNIC (NPU)

Barefoot Torfino2 Switch (P4)

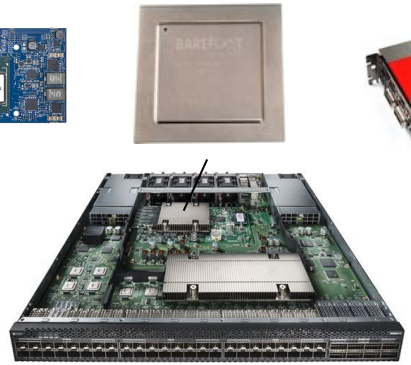

Xilinx Alveo U280 (FPGA)

“Low-Touch” (and “No-Touch”)
Application-Specific Integrated Circuits (ASIC) based data-plane

Pros:
- Optimized for specific tasks
- Lower cost

Cons:
- Inflexible

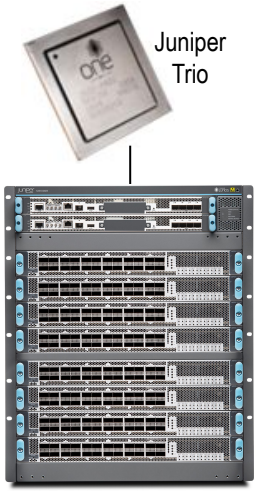

Juniper Trio

Juniper MX10008

Nokia FP4

Nokia 7750 SR-7s

Broadcom Jericho

Arista 7504R

ESnet

Phase 1 - Done!

ESnet 5 → ESnet 5.5 (Current)

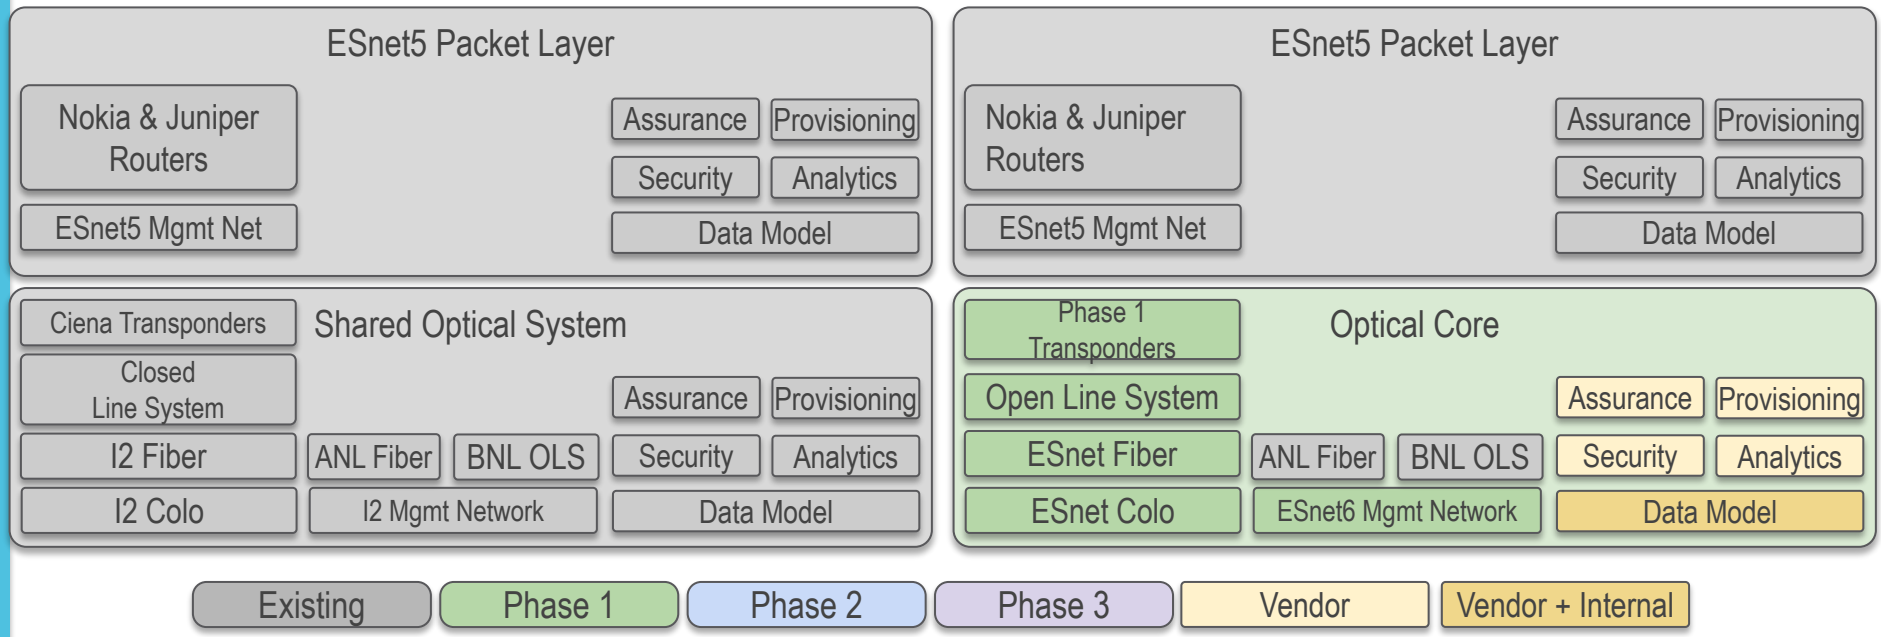

OLS install snapshot 4/13/2020

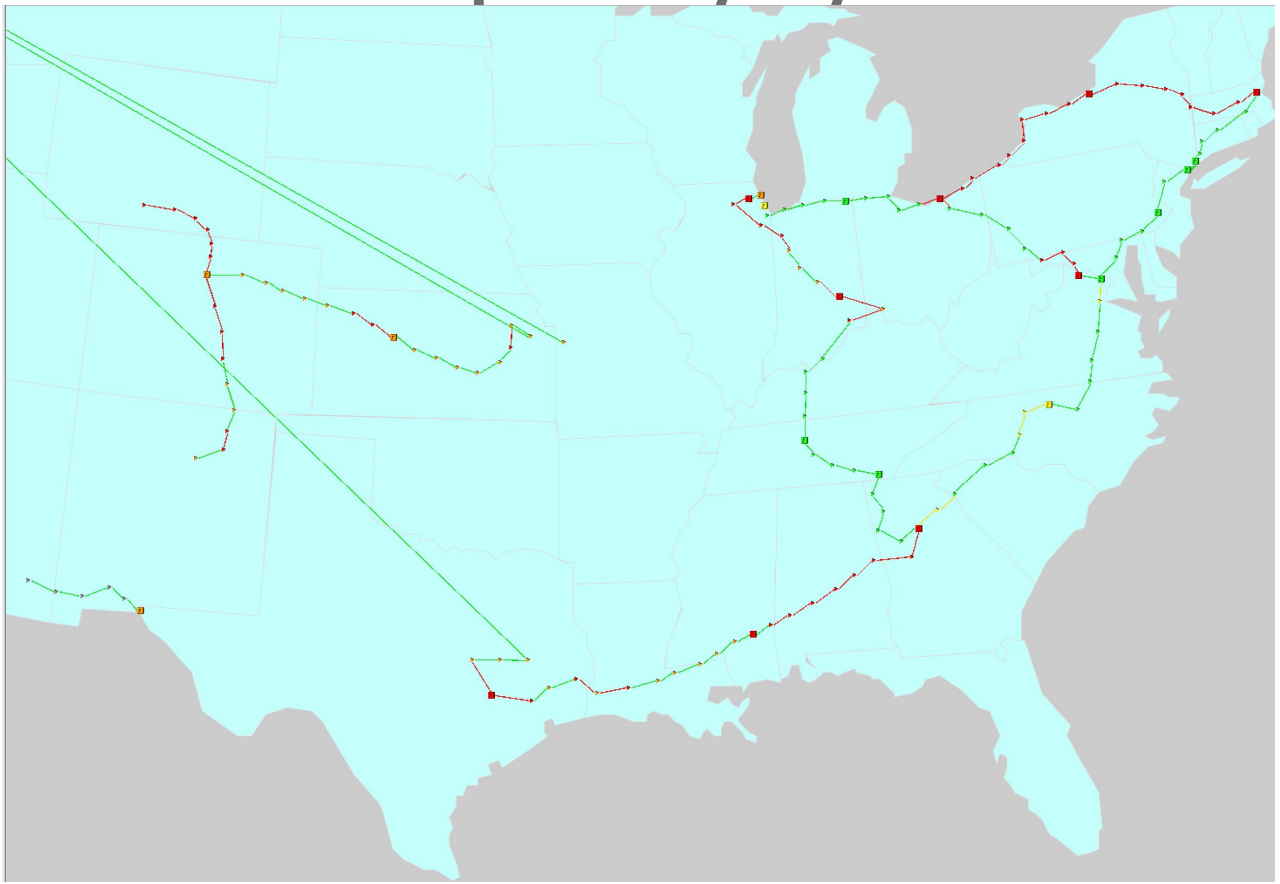

ESnet

Now a fully deployed optical network

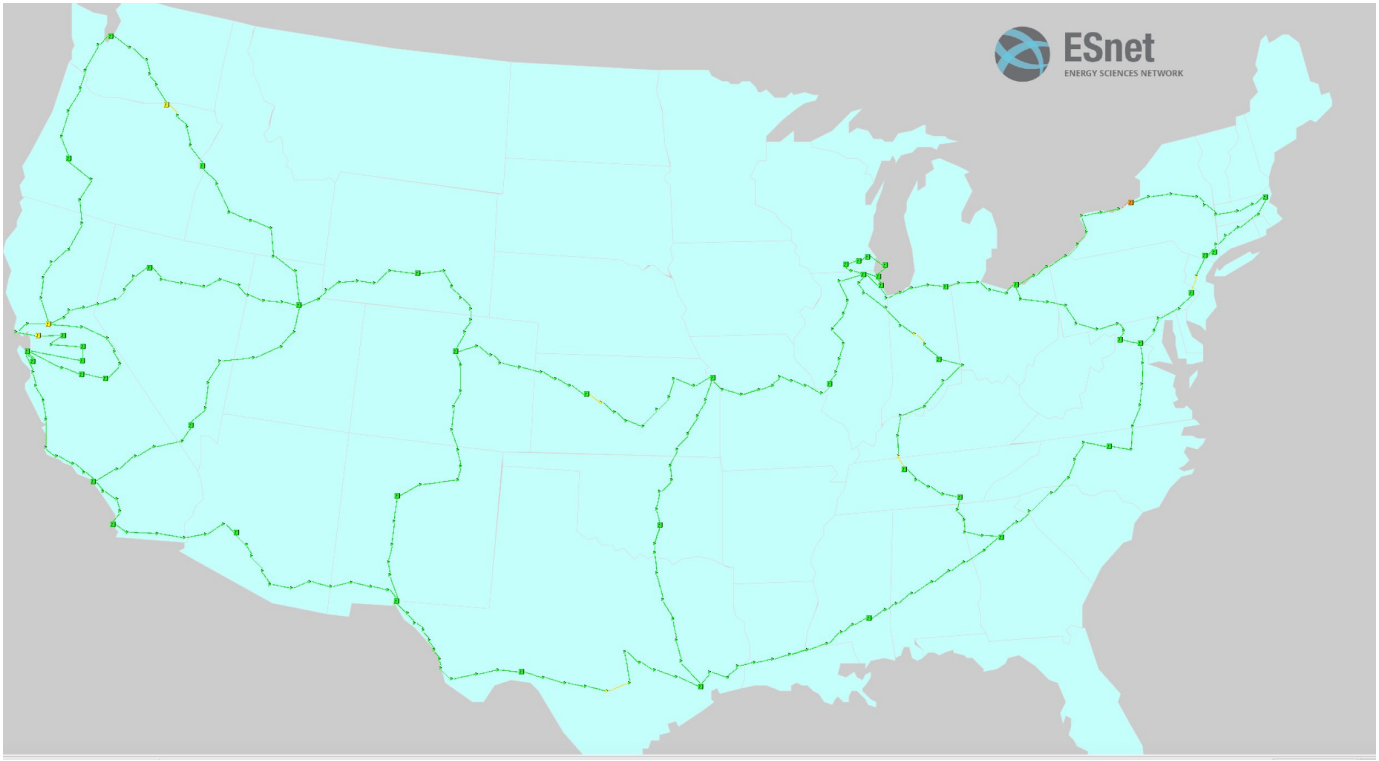

ESnet

Phase 1 - Done!

15,000+ Miles of Fiber

Infinera Flex-ILS Optical System

272 Amplifiers

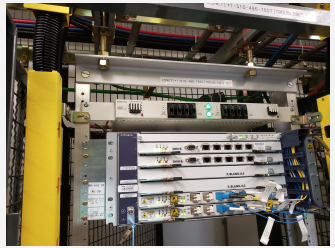

37 ROADMS

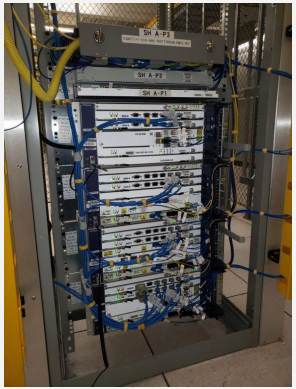

300+ Colo Facilities

72 Groove Transponders

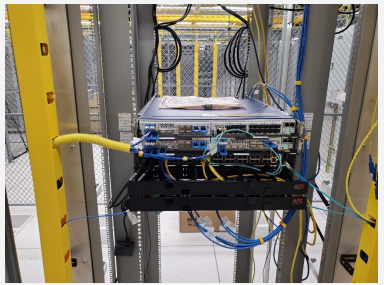

Layer

Layer

Assurance

Provisioning

Security

Analytics

Data Model

Data Model

Snet5 Mgmt Net

Phase 1 Transponders

Open Line System

ESnet Fiber

ANL Fiber

ESnet Colo

ESnet6 Mgr

Phase 2

Phase 3

Vendor

Phase 2 - In Progress

ESnet 5.5 (Current)

ESnet 6 - Low Touch

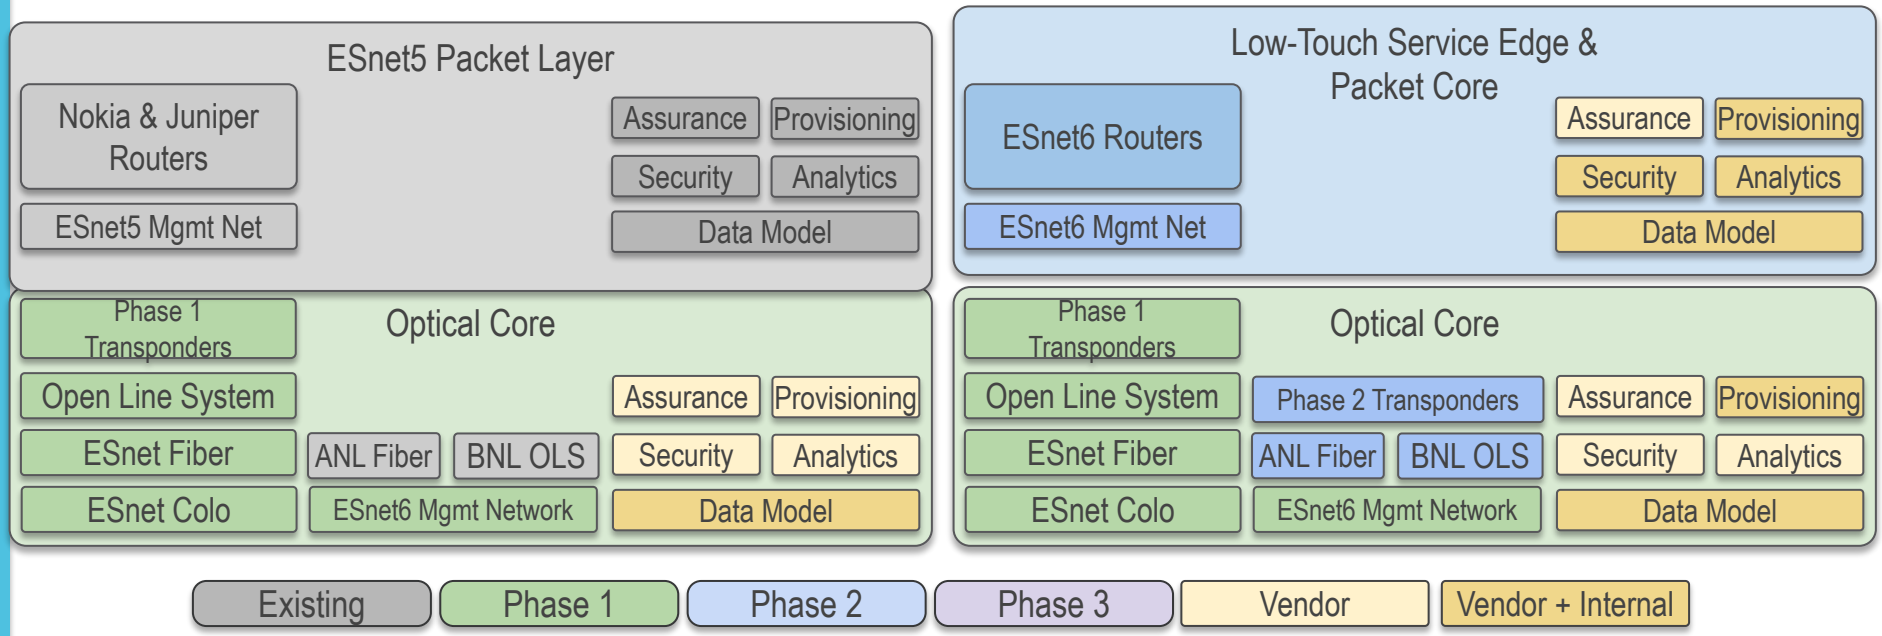

Phase 2 - In Progress

67 Routers
Nokia 7750-SR2

ESnet 6 - Low Touch

41 Transponder Chassis
Ciena Waveserver 5

27 Management Network

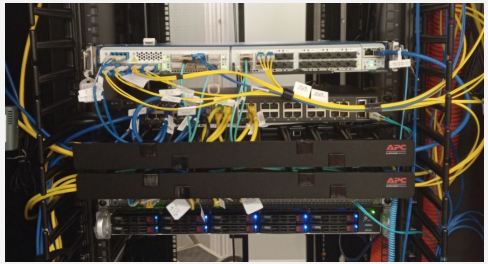

OLS to Brookhaven

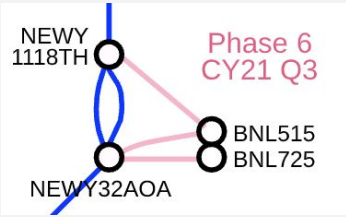

ANL Fiber

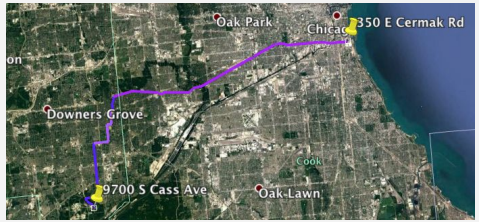

Layer
Assurance
Provisioning
Analytics
ESnet6 Routers
ESnet6 Mgmt Net
Packet
Optical
Phase 2 Transponde
ANL Fiber
BNL OLS
Assurance
Provisioning
Security
Analytics
ESnet Colo
ESnet6 Mgmt Network
ESnet6 Mgmt Network
Data Model

Existing
Phase 1
Vendor
Vendor + Internal

Routing Technology

7750 SR-s series

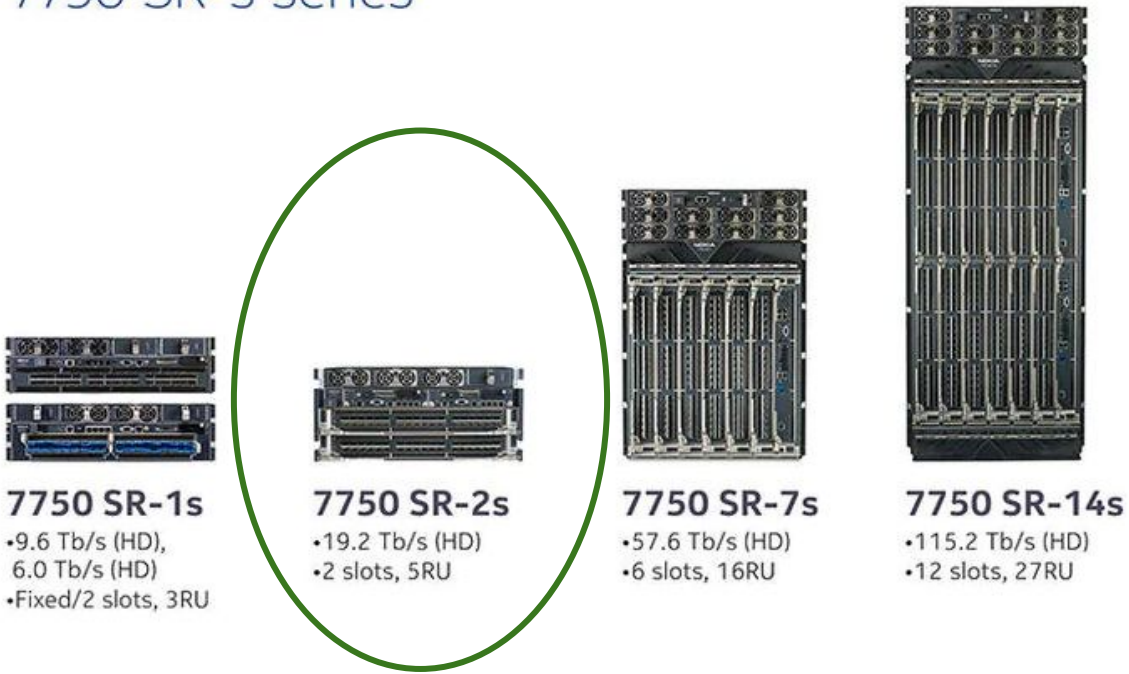

7750 SR-1s
- 9.6 Tb/s (HD), 6.0 Tb/s (HD)
- Fixed/2 slots, 3RU

7750 SR-2s
- 19.2 Tb/s (HD)
- 2 slots, 5RU

7750 SR-7s
- 57.6 Tb/s (HD)
- 6 slots, 16RU

7750 SR-14s
- 115.2 Tb/s (HD)
- 12 slots, 27RU

ESnet

Nokia SR-2s Large Router Configuration

- Two 36-connector XMA-s cards licensed for the full 4.8 tbps
- Max supported per Router:
  - 24 x 400GE
  - 96 x 100GE
  - Other combinations of 10GE, 100GE & 400GE totaling 9.6 tbps

- QSFP-DD connectors support:
  - QSFP+: 4x10GE or 1x 40GE
  - QSFP28: 10x10GE or 1x100GE
  - QSFP28-DD: 2x100GE
  - QSFP56-DD: 1x400GE

ESnet

17

Nokia SR-2s Small Router Configuration

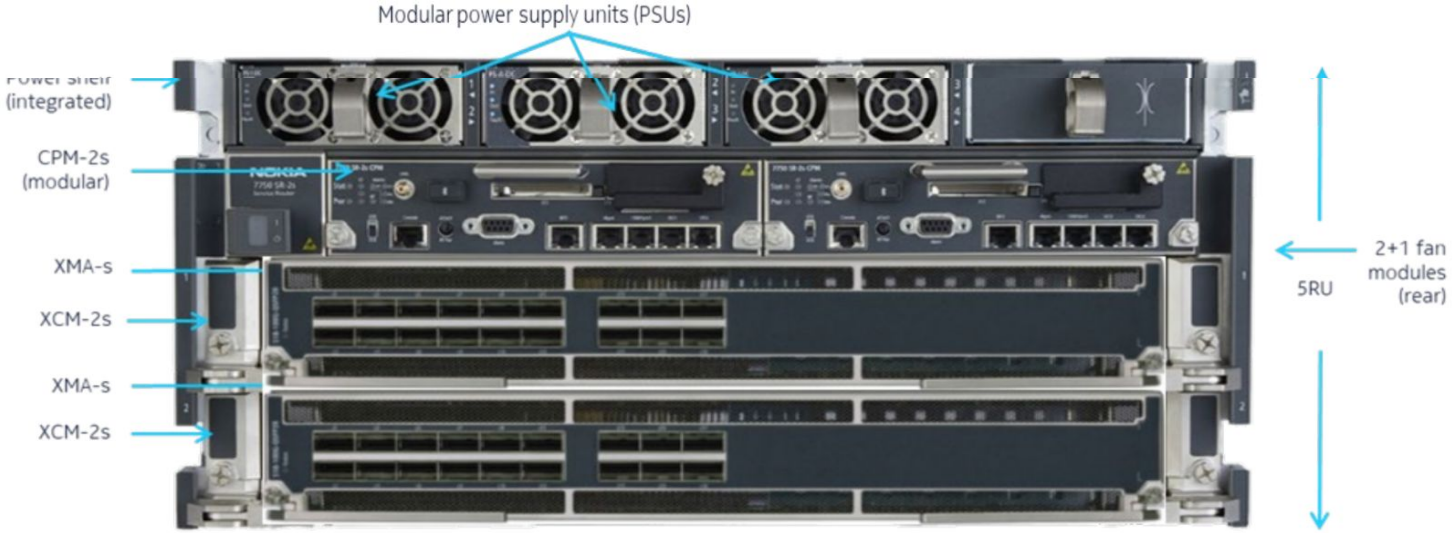

- Two 18-connector XMA-s cards licensed for 6 connectors & 600 gbps
- Max supported per Router:
  - 12 x 100GE
  - 120 x 10GE
  - Other combinations of 10GE and 100GE totaling 1.2 tbps

- QSFP28 connectors support:
  - QSFP+: 4x10GE or 1x 40GE
  - QSFP28: 10x10GE or 1x100GE
- License upgrade-able to 2.4 tbps per card

ESnet

Phase 3

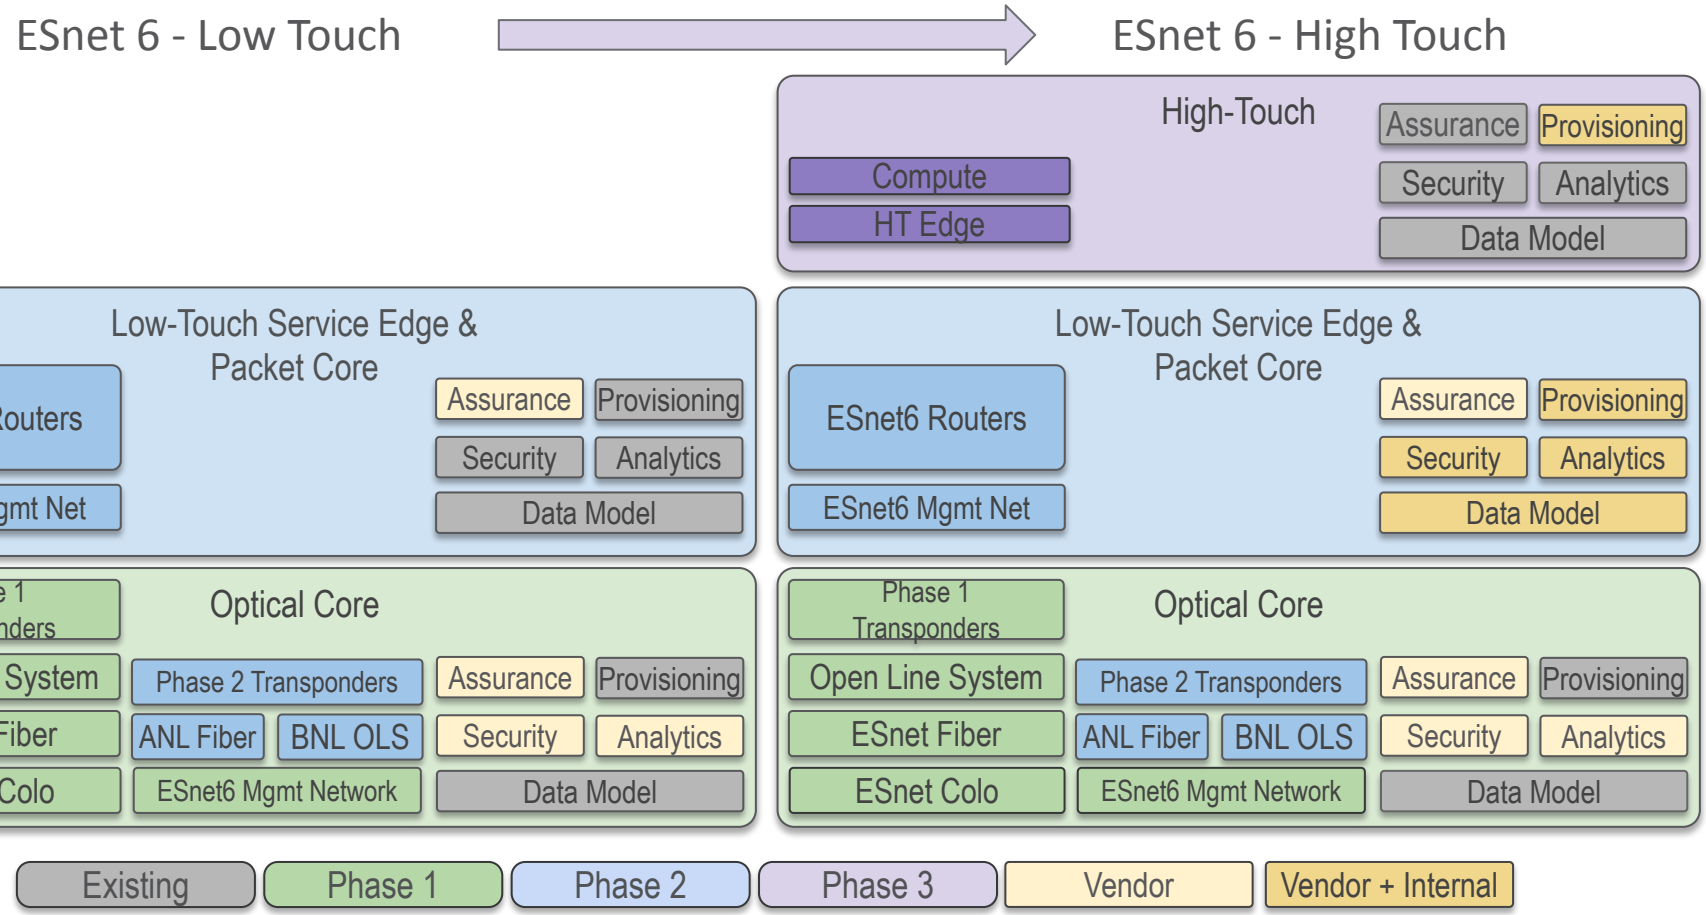
ESnet 6 - Low Touch
ESnet 6 - High Touch
High-Touch
Compute
HT Edge
Assurance
Provisioning
Security
Analytics
Data Model
Low-Touch Service Edge & Packet Core
ESnet6 Routers
ESnet6 Mgmt Net
Assurance
Provisioning
Security
Analytics
Data Model
Low-Touch Service Edge & Packet Core
ESnet6 Routers
ESnet6 Mgmt Net
Assurance
Provisioning
Security
Analytics
Data Model
Phase 1 Transponders
Optical Core
Open Line System
ESnet Fiber
ESnet Colo
Phase 2 Transponders
ANL Fiber
BNL OLS
ESnet6 Mgmt Network
Assurance
Provisioning
Security
Analytics
Data Model
Phase 1 Transponders
Optical Core
Open Line System
ESnet Fiber
ESnet Colo
Phase 2 Transponders
ANL Fiber
BNL OLS
ESnet6 Mgmt Network
Assurance
Provisioning
Security
Analytics
Data Model
Existing
Phase 1
Phase 2
Phase 3
Vendor
Vendor + Internal

Phase 3

ESnet 6 - Low Touch → ESnet 6 - High Touch

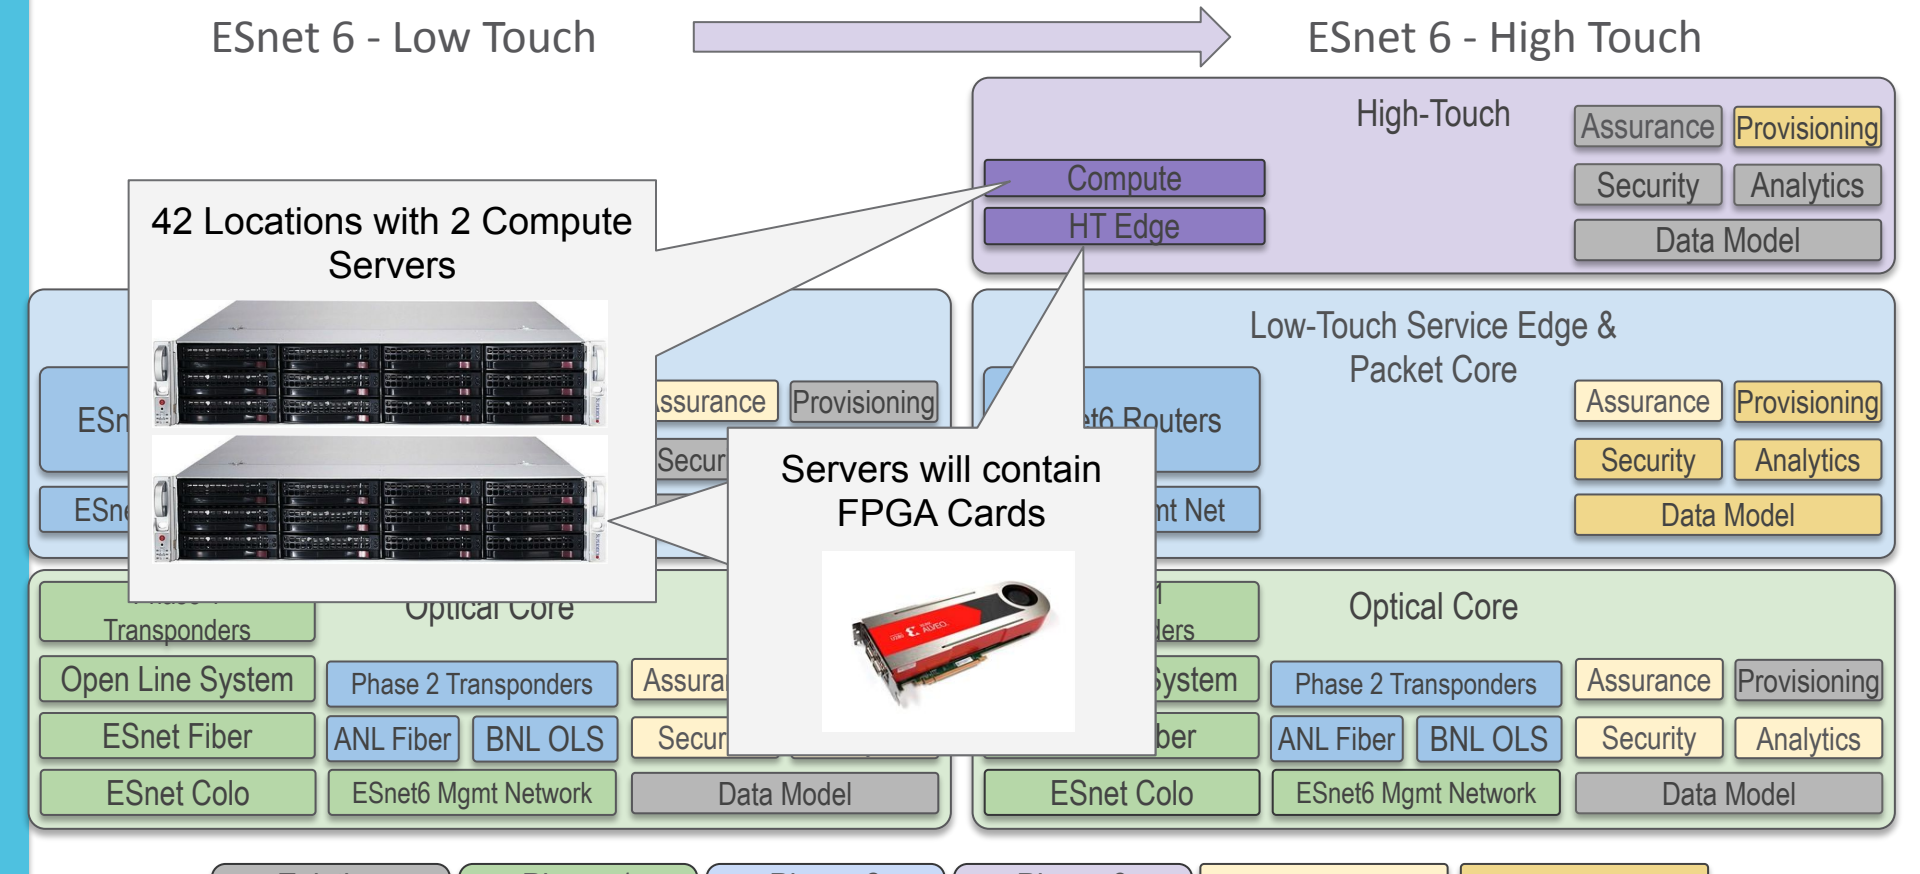

Existing | Phase 1 | Phase 2 | Phase 3 | Vendor | Vendor + Internal

What is High Touch?

- A programmable data plan for development and deployment of innovative science data services.
  - Provides unprecedented flexibility to customize packet processing at scale.
  - Developing new services to support emerging science application workflows not currently supported by mainstream packet edge routers without the huge expense of contracting router vendors to develop (and support) the capability.
  - Platform for research into features and functions that will educate the design and architecture of ESnet7.
- First High Touch Service: **Precision Network Telemetry**
  - We can get detailed insights into how the network is behaving.
  - We can profile how flows are performing in our network and take proactive action.
  - We can use the detailed flow information for traffic engineering, capacity planning, or anomaly detection (e.g., AI/ML applications)

ESnet

Precision Network Telemetry Services

Flow (Feature) Distillation - Provides flow summaries of packets entering the network, thereby enabling full visibility of network traffic without large storage requirements.

Network Microscope - Provides the capability to dynamically select network flow(s) for replication (of only the packet header and not user data contents), augmented with timestamps, to be redirected to compute resources for further processing, e.g., security analysis, feature extraction, etc.

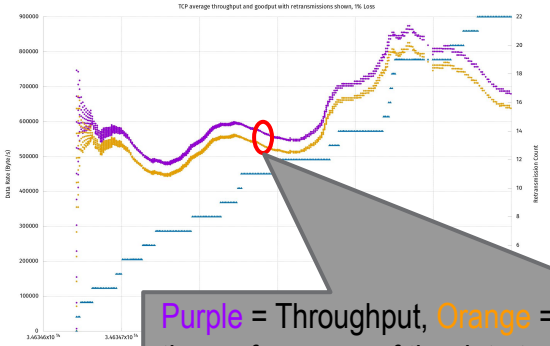

Purple = Throughput, Orange = Goodput. The difference between the lines represents the performance of the data transfer, convergence is good, divergence is bad.

Benefits

- We can get detailed insights into how the network is behaving.
- We can profile how flows are performing in our network and take proactive action.
- We can use the detailed flow information for traffic engineering, capacity planning, or anomaly detection (e.g., AI/ML applications)

ESnet

Delivering ESnet6 Automation

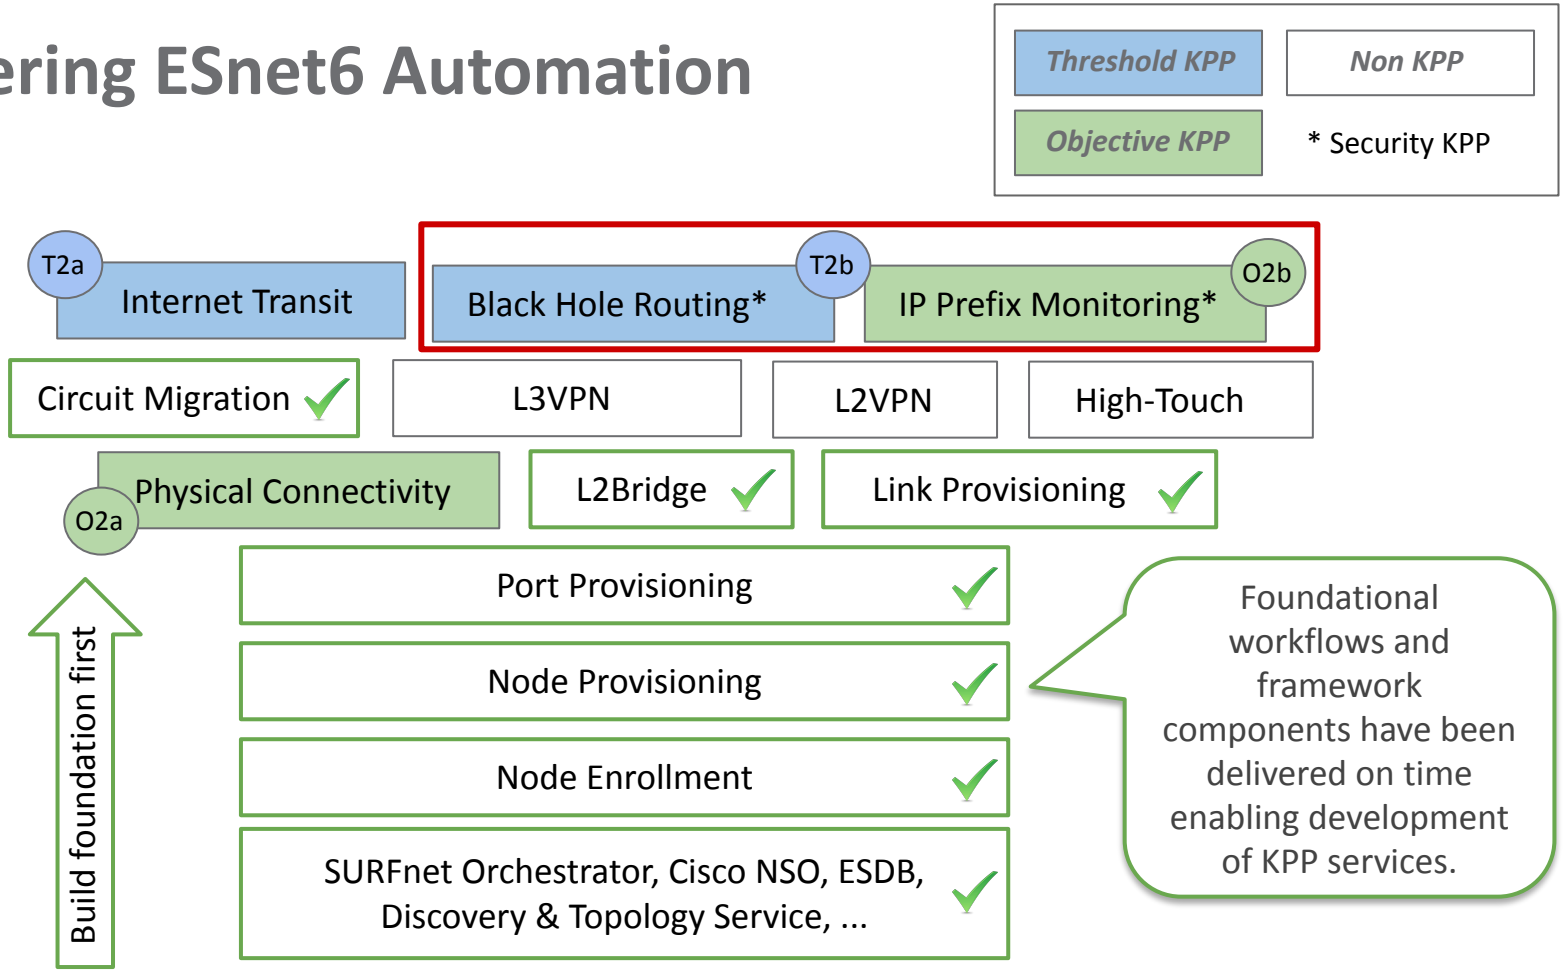

Orchestration Implementation

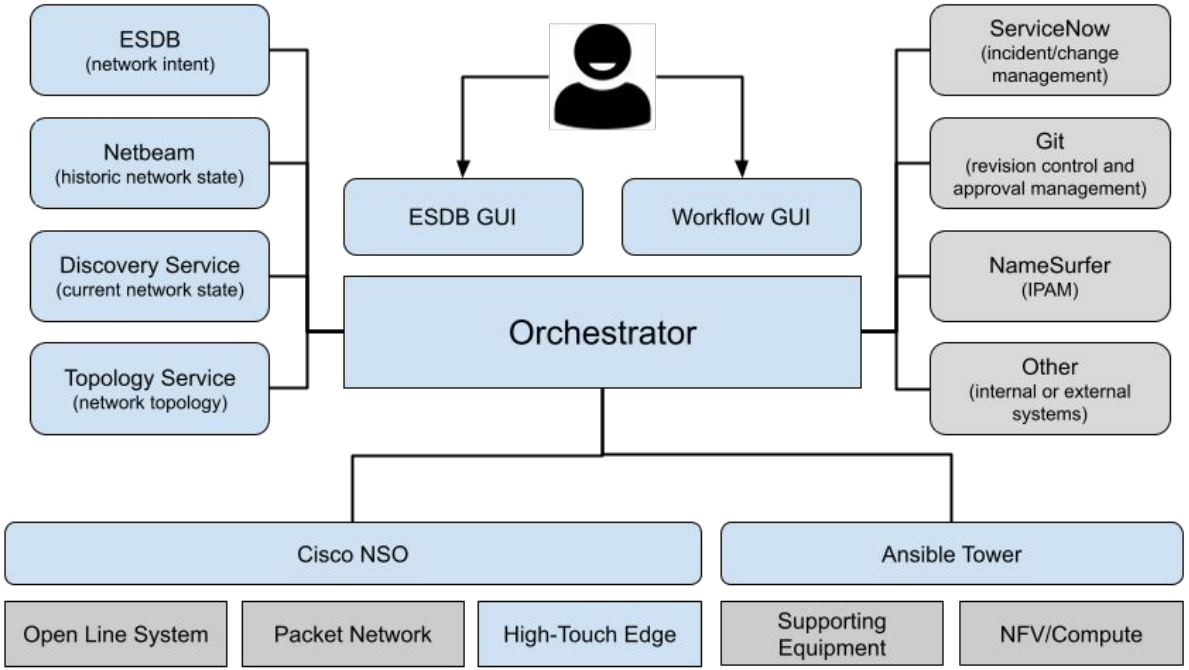

The five key components to a cohesive orchestration solution:

- Workflow management;
- Automated provisioning;
- Network intent;
- Network discovery;
- Network topology.

ESnet

System Management

- **INOC**
  - Monitors the ESnet optical layer
  - Monitors the fiber layer by virtue of monitoring the optical layer
    - Providing the ENOC ticketed alerts. Not currently taking action
  - Manages problems including simple HW failures & fiber cuts
- **ESnet NOC**
  - Monitors the ESnet packet layer
  - Takes certain mitigating actions in case of routing instability (costing flapping links, etc.)
  - Interacts with ESnet sites for coordinating site access for problem resolution
  - Monitors the OLS by virtue of monitoring services it provides to the packet layer
    - Receives ticketed alerts from INOC through ServiceNow integration (alerts fed to INOC by TNMS)
  - Manages some more complex problems, including span loss
  - Coordinates planned maintenance events, handles trouble tickets, and escalates to Neteng
- **ESnet Neteng**
  - Protocol design and operationalization (e.g. Segment Routing configuration and label space definition, etc.)
  - System specification with Planning and Architecture Group
  - Automation planning and specification (with Security and Software Engineering Groups)
  - Support day to day operations
  - Provide on-call support
  - Receive escalations from ENOC, diagnose and triage problems
  - Service-owner support for PCE
  - Escalate to optical team if needed (subgroup of engineers who are optical experts)

ESnet

ESnet6 Network System

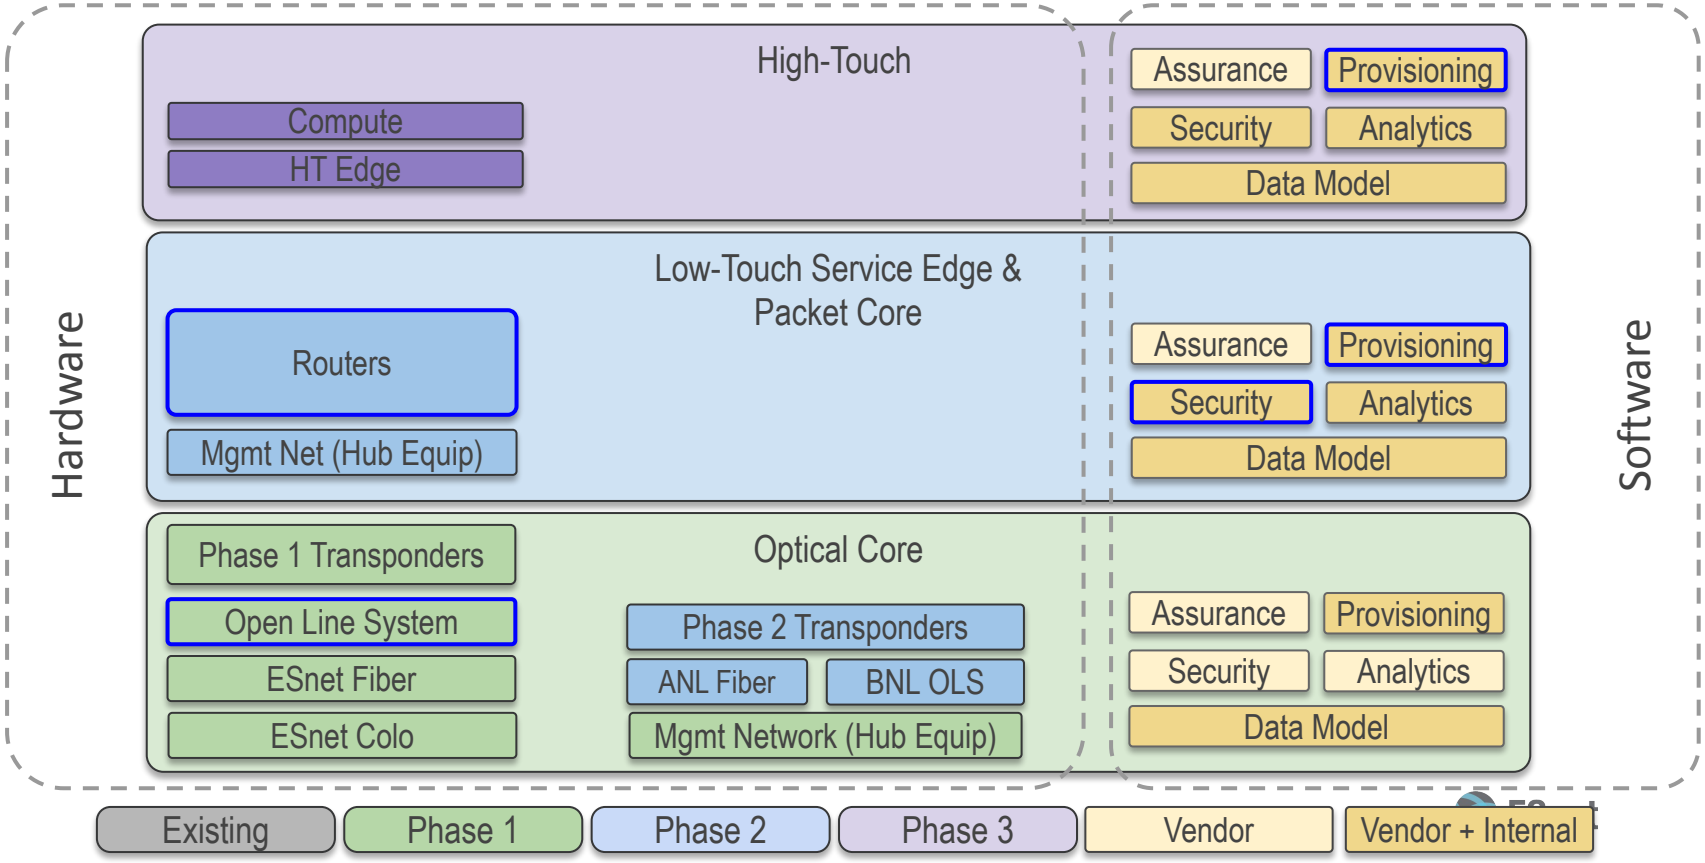

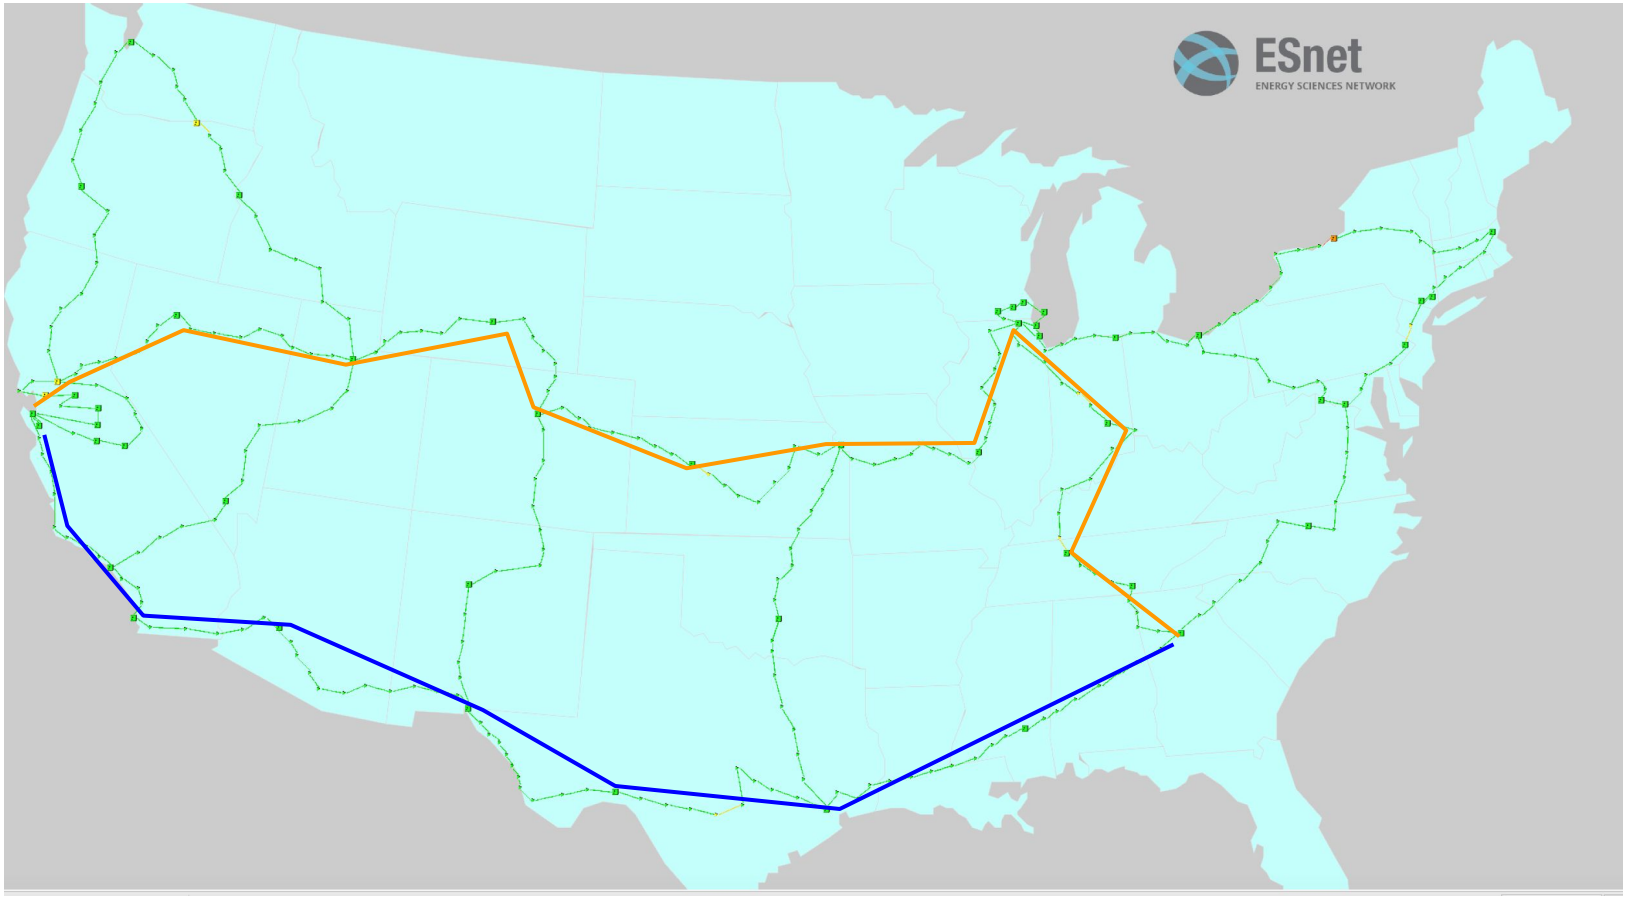

ESnet6 optical footprint with planned 100G paths between ATLA and SLAC in support of Rubin Observatory Long Haul Network

ESnet

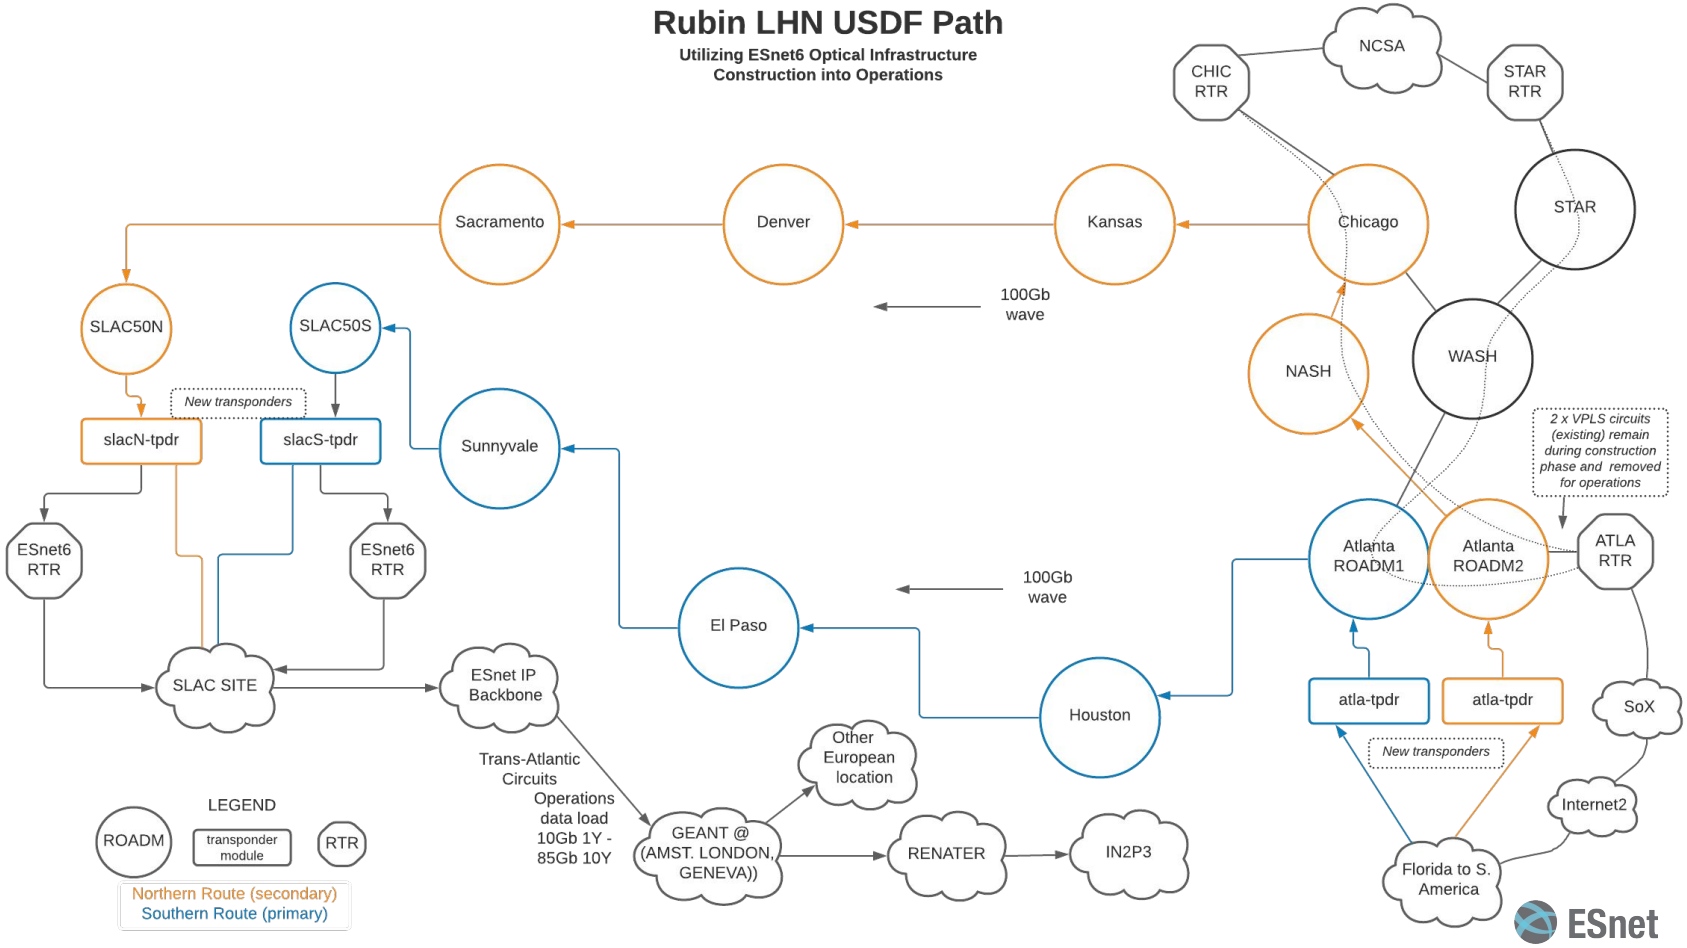
Rubin LHN USDF Path
Utilizing ESnet6 Optical Infrastructure
Construction into Operations
NCSA
CHIC RTR
STAR RTR
STAR
Sacramento
Denver
Kansas
Chicago
100Gb wave
SLAC50N
SLAC50S
New transponders
slacN-tpdr
slacS-tpdr
Sunnyvale
NASH
WASH
2 x VPLS circuits (existing) remain during construction phase and removed for operations
ESnet6 RTR
ESnet6 RTR
Atlanta ROADM1
Atlanta ROADM2
ATLA RTR
100Gb wave
El Paso
SLAC SITE
ESnet IP Backbone
Houston
atla-tpdr
atla-tpdr
SoX
New transponders
Trans-Atlantic Circuits
Operations data load 10Gb 1Y - 85Gb 10Y
Other European location
Internet2
GEANT @ (AMST. LONDON, GENEVA))
RENATER
IN2P3
Florida to S. America
LEGEND
ROADM
transponder module
RTR
Northern Route (secondary)
Southern Route (primary)
ESnet

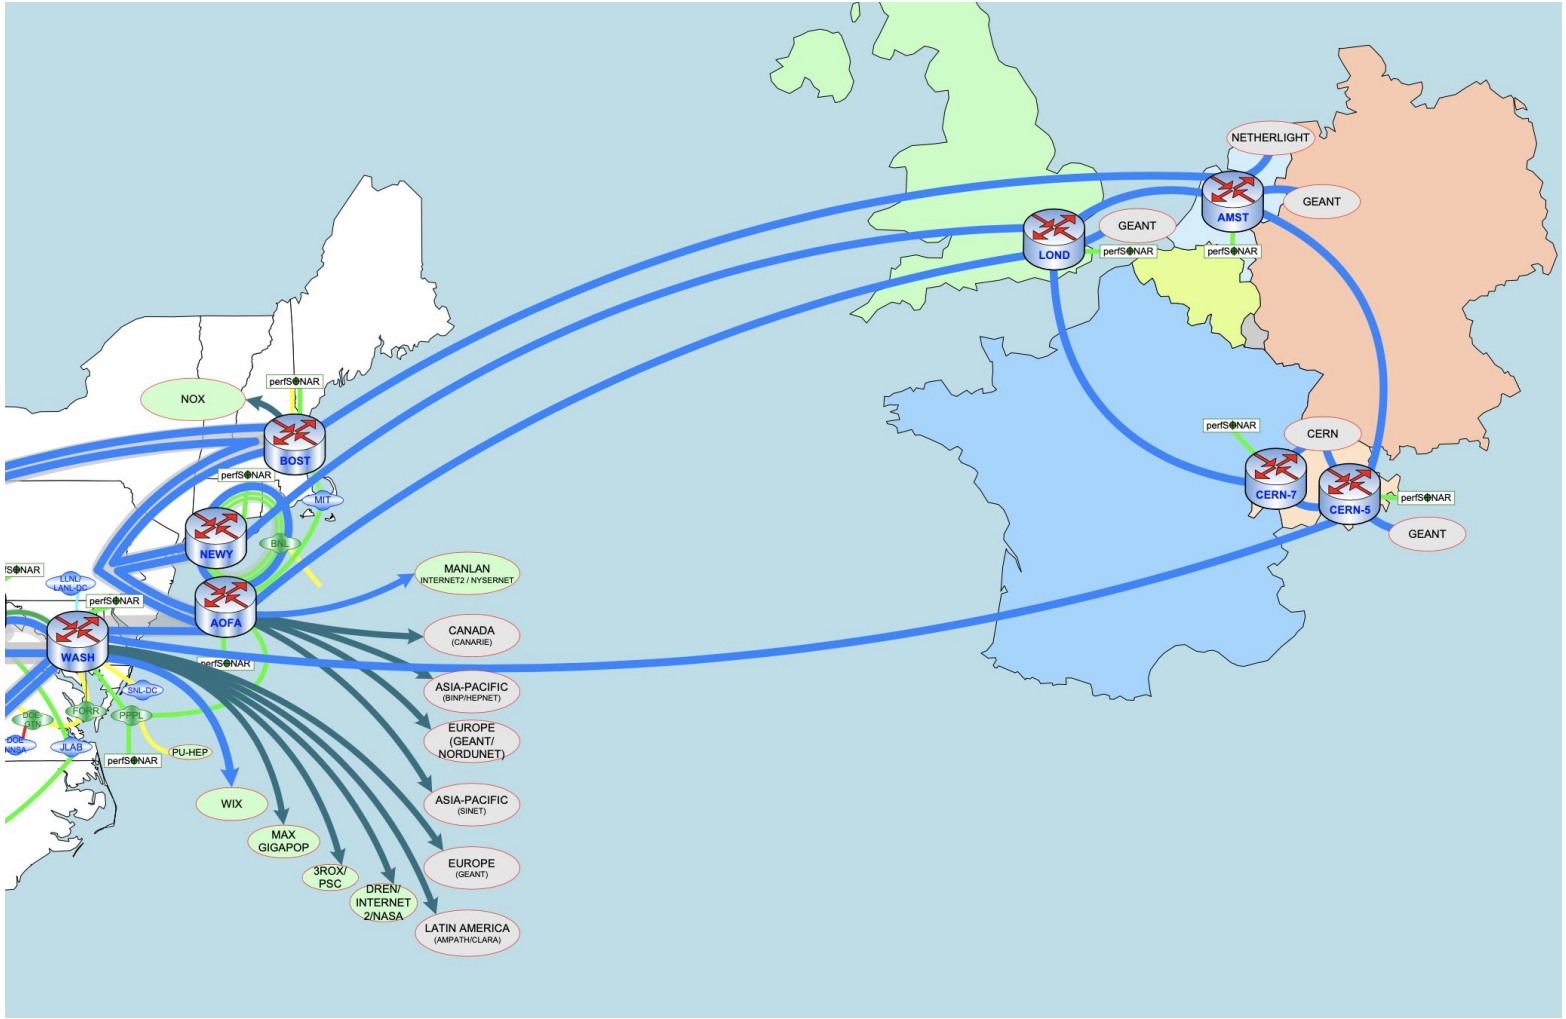
NETHERLIGHT
GEANT
GEANT
AMST
LOND
GEANT
perfSONAR
perfSONAR
perfSONAR
CERN
CERN-7
CERN-5
perfSONAR
GEANT
NOX
perfSONAR
BOST
perfSONAR
MIT
NEWY
BNL
AOFA
WASH
perfSONAR
perfSONAR
LNL
LANL-DC
perfSONAR
SNL-DC
FCHR
PPPL
DOE GEN
JLAB
PU-HEP
perfSONAR
MANLAN
INTERNET2 / NYSERNET
CANADA
(CANARIE)
ASIA-PACIFIC
(BINP/HEPNET)
EUROPE
(GEANT/
NORDUNET)
ASIA-PACIFIC
(SINET)
EUROPE
(GEANT)
LATIN AMERICA
(AMPATH/CLARA)
WIX
MAX
GIGAPOP
3ROX/
PSC
DREN/
INTERNET
2/NASA

Questions...

ESnet

Transatlantic Forecasting

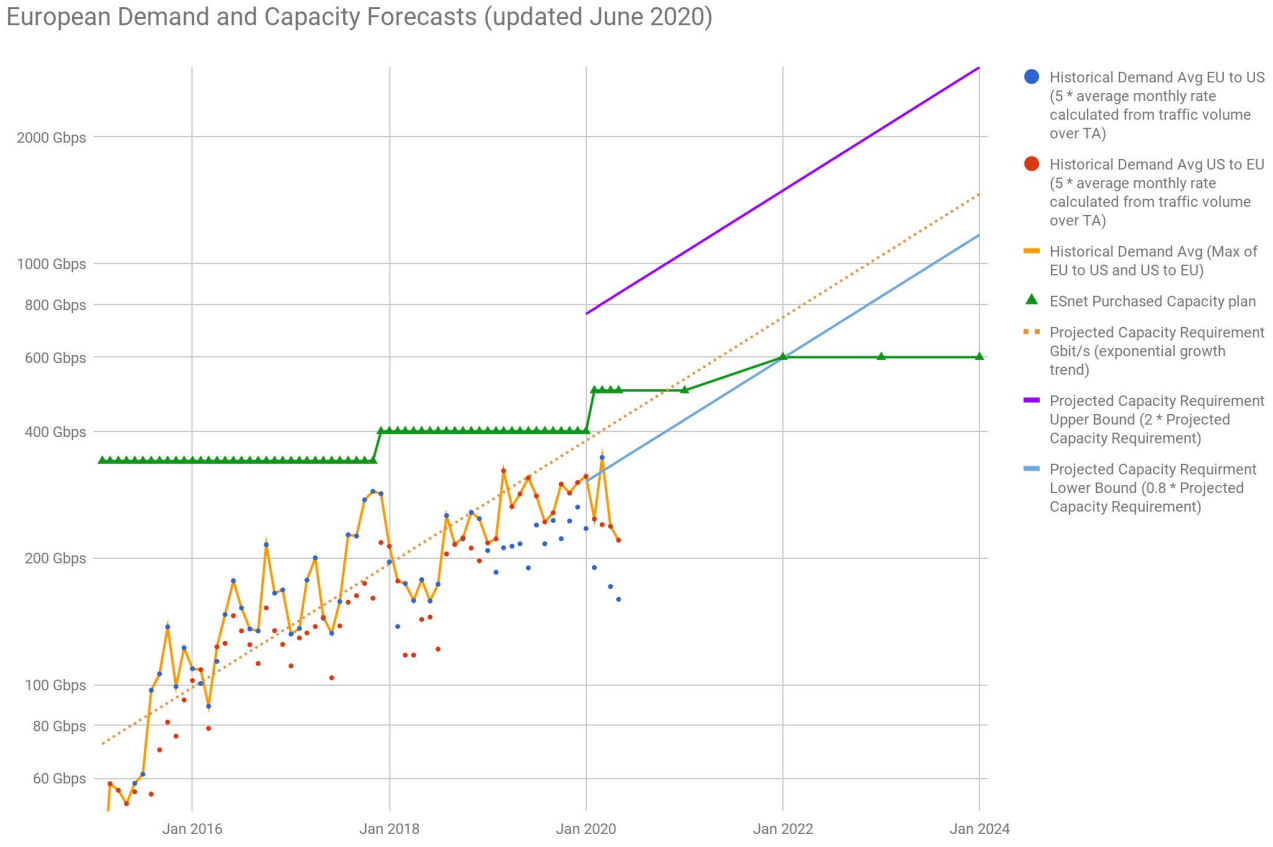

Note: Projected Capacity Requirement is calculated as 5 times the average utilization (notes presented on a separate slide).

ESnet

Transatlantic Forecasting

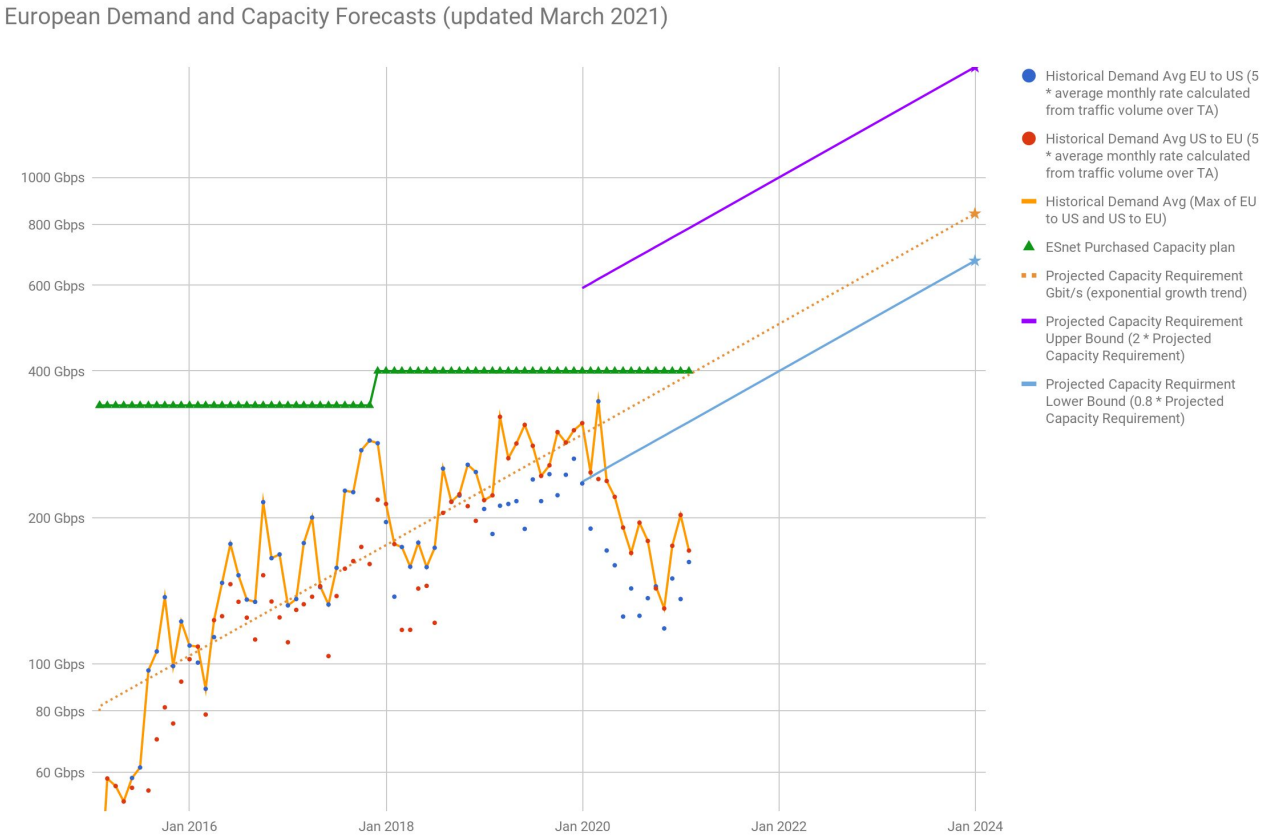

Note: Projected Capacity Requirement is calculated as 5 times the average utilization (notes presented on a separate slide).

ESnet