BNL and JLAB Connectivity via ESnet

- Best path between BNL and JLAB is:
  - NYC <-> Washington DC <-> Ashburn (Green Path).
  - ~11ms of latency between ESnet routers
- The next diverse path between BNL and JLAB is:
  - Boston <-> Chicago <-> Nashville <-> Atlanta (Orange and Red)
  - ~60ms of latency between ESnet routers
- ESnet's backbone is Nx400G links on their optical system with the caveat/bottleneck being the last mile to JLAB is limited to two 100G procured Segra circuits. ESnet is pursuing the potential for dark fiber into JLAB but its still in the investigatory phase.

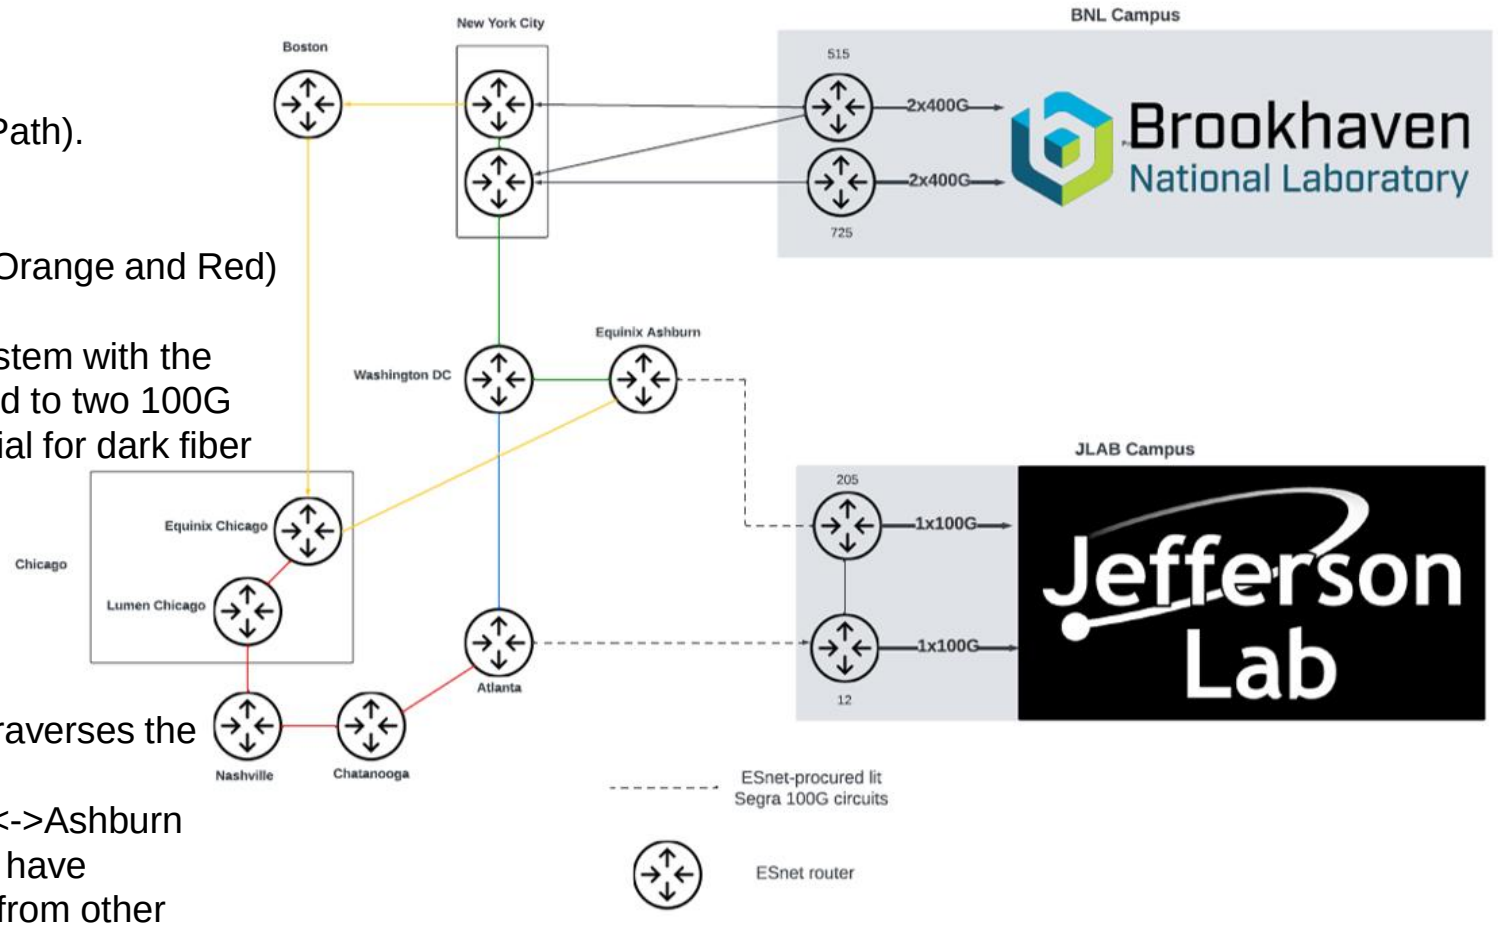

- As of today, traffic between BNL and JLAB already traverses the green path due to best path selection.
- Since the primary path to JLAB is via NYC<->Wash<->Ashburn if a virtual circuit is utilized, it must be prioritized and have bandwidth guarantees even with other priority flows from other collaborations such as LHC (e.g., LHCOPN).

Brookhaven National Laboratory

BNL High Throughput Science Network (HTSN)

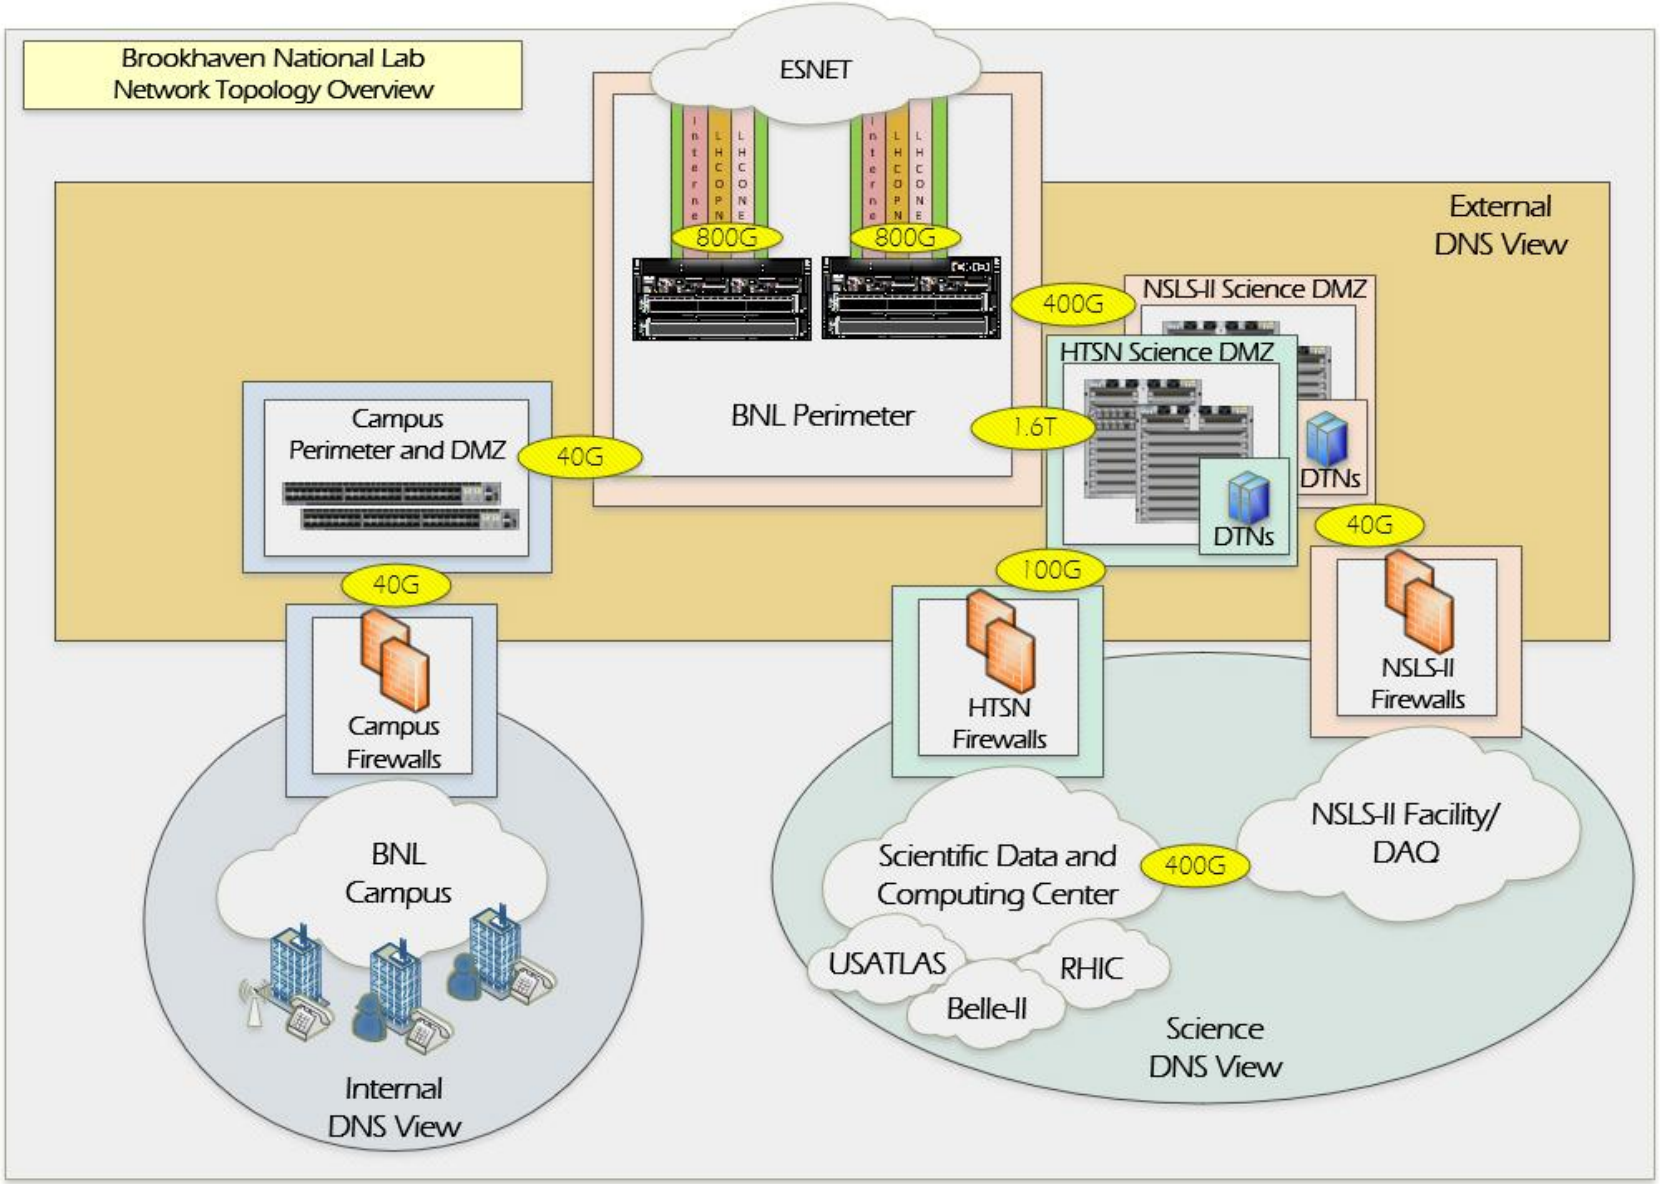

Brookhaven National Laboratory

HTSN/SDCC Network Infrastructure

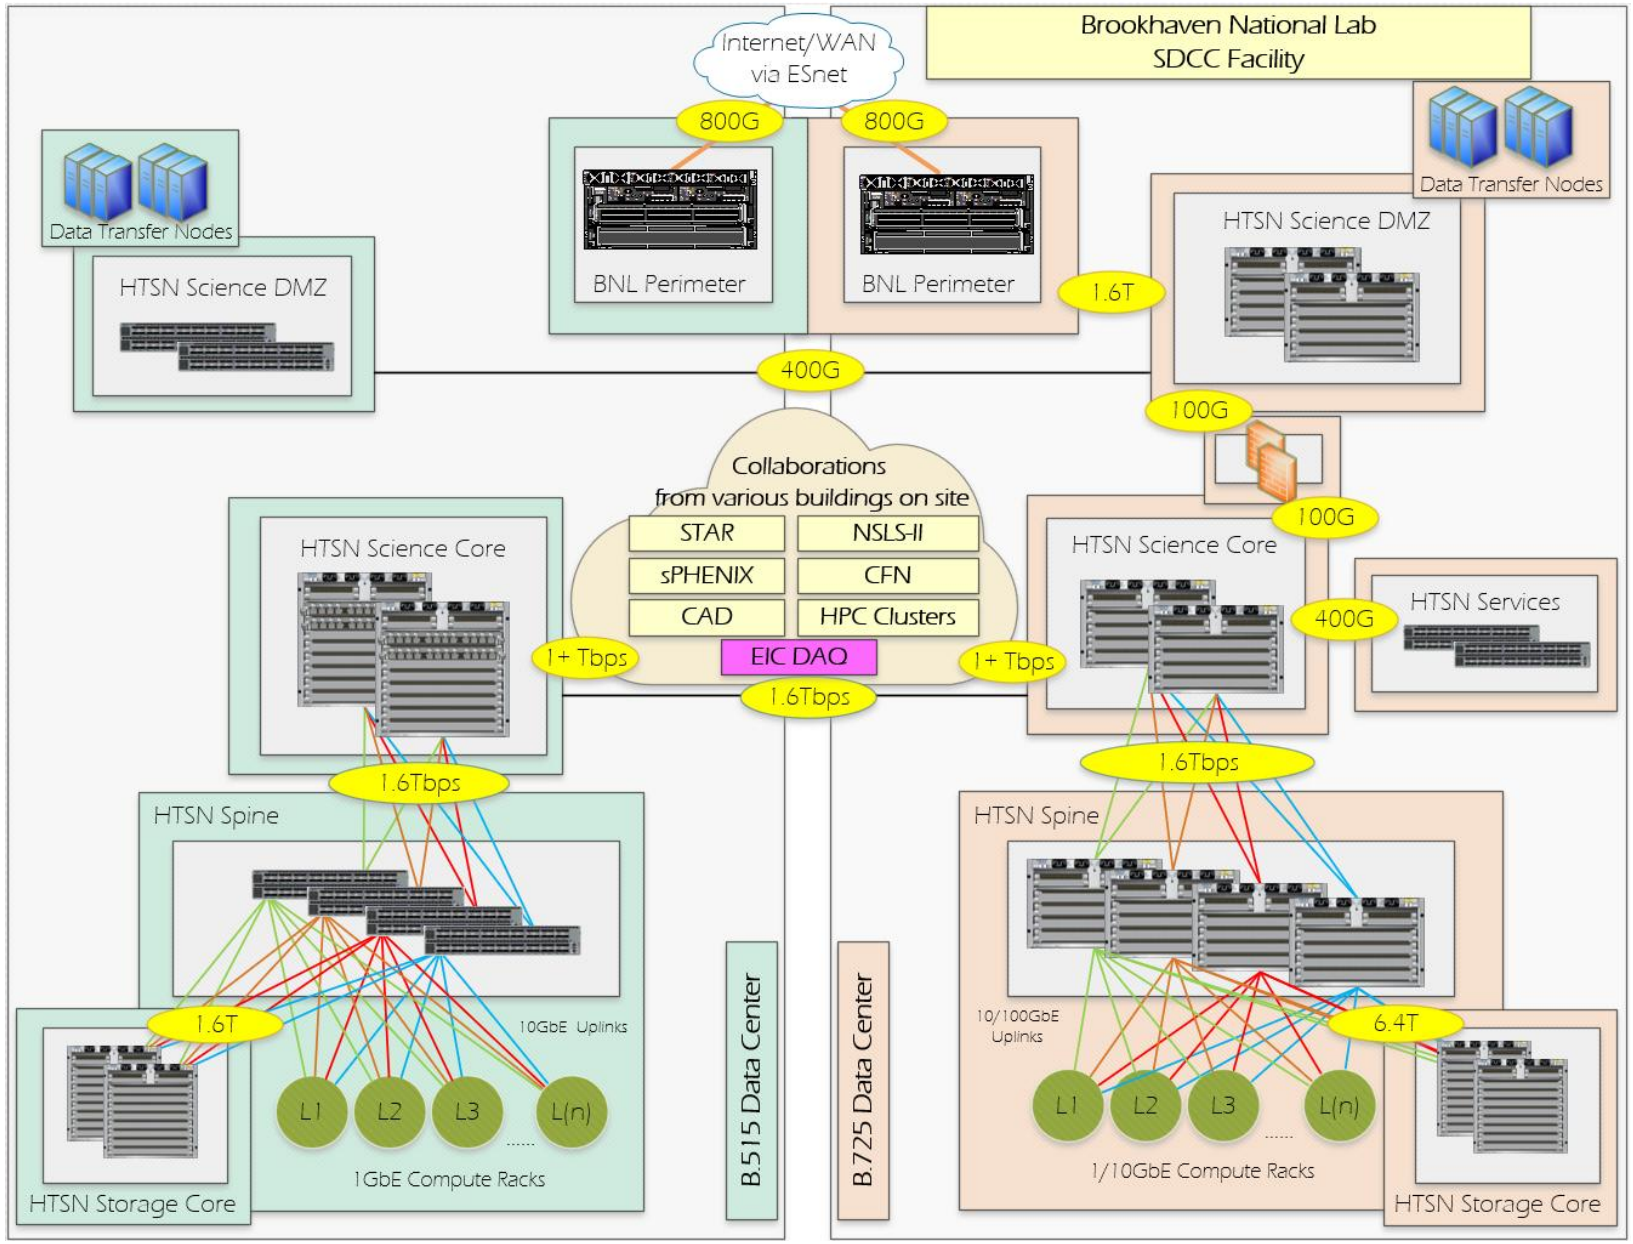

Brookhaven National Laboratory

ePIC Internet/WAN Network Topology Option 1 - (Most Preferred)

- Network topology is very similar to the LHC where VPLS/OSCARS circuits are terminated on BNL's Network Perimeter which extend point to point connectivity with JLAB via EBGP.
- No new physical circuits would be required from ESnet.
- This architecture continues to provide a single point of entry into BNL from ESnet which allows Cyber Security and ITD Network Engineering to utilize already deployed monitoring and mitigation infrastructures.
- Allows both the ePIC DAQ and BNL E1 site/doors to transfer data to JLAB via the VPLS/OSCARS circuits. This is similar to the LHCOPN.
- Places a high throughput firewall between the ePIC DAQ and the Internet and JLAB. Which would allow the ePIC DAQ to utilize ESnet general backbone (Internet) to reach JLAB or any other destinations if the VPLS circuits were to fail.
- Connects the ePIC DAQ to the SDCC via the HTSN Science Core with route filtering/ACLs (like sPHENIX).

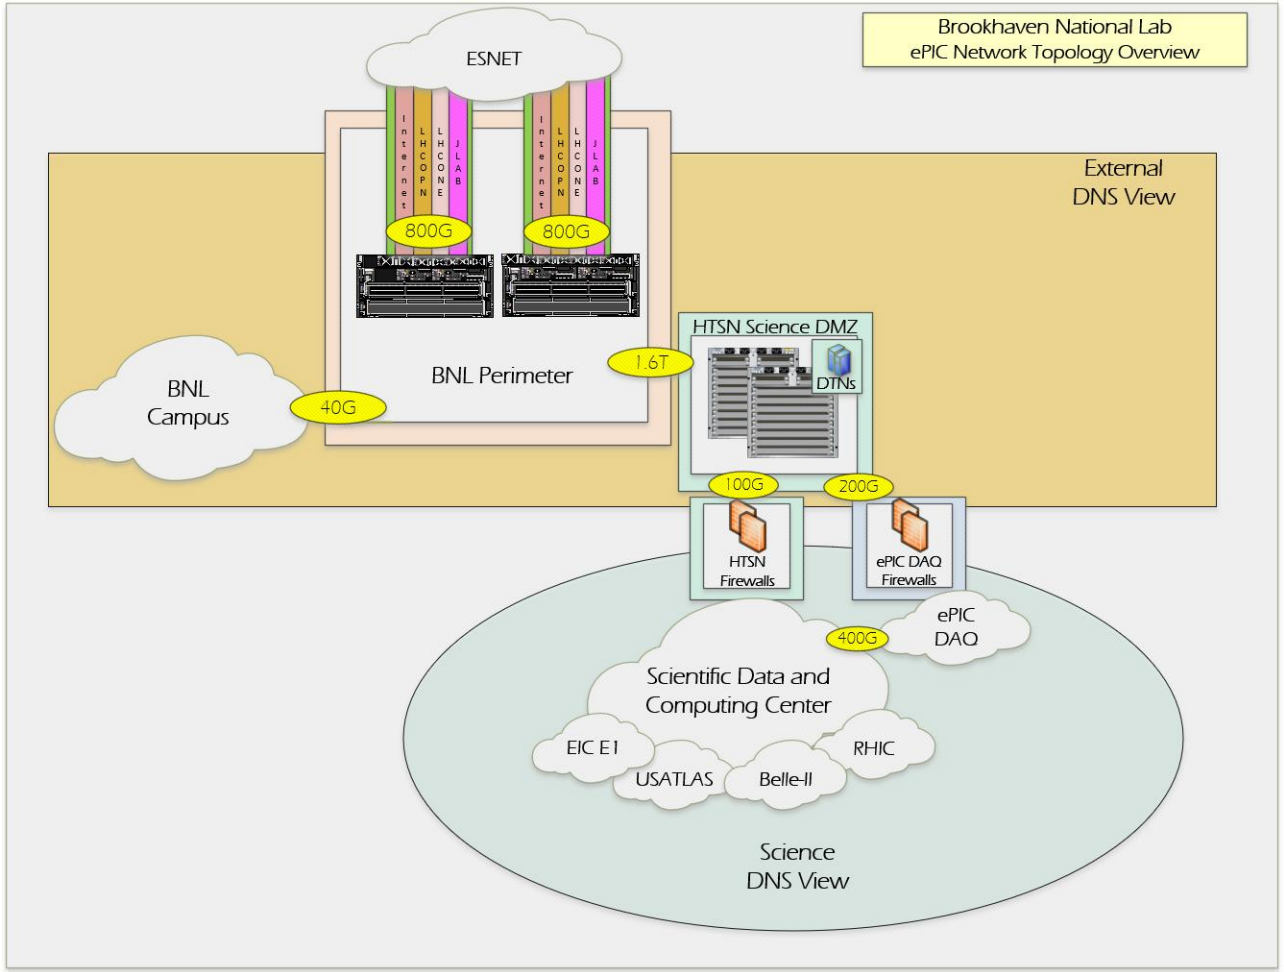

Brookhaven National Laboratory

ePIC Internet/WAN Network Topology Option 2

- Network topology is similar to Option 1 but would require multi-homing the buffer boxes to the BNL Science DMZ.
- Buffer boxes could reside on their own subnet and router ACLs could be utilized along with Iptables on the buffer boxes.
- Since the main data path between e0 and JLAB isn't through the firewall a lower performance firewall could be utilized.

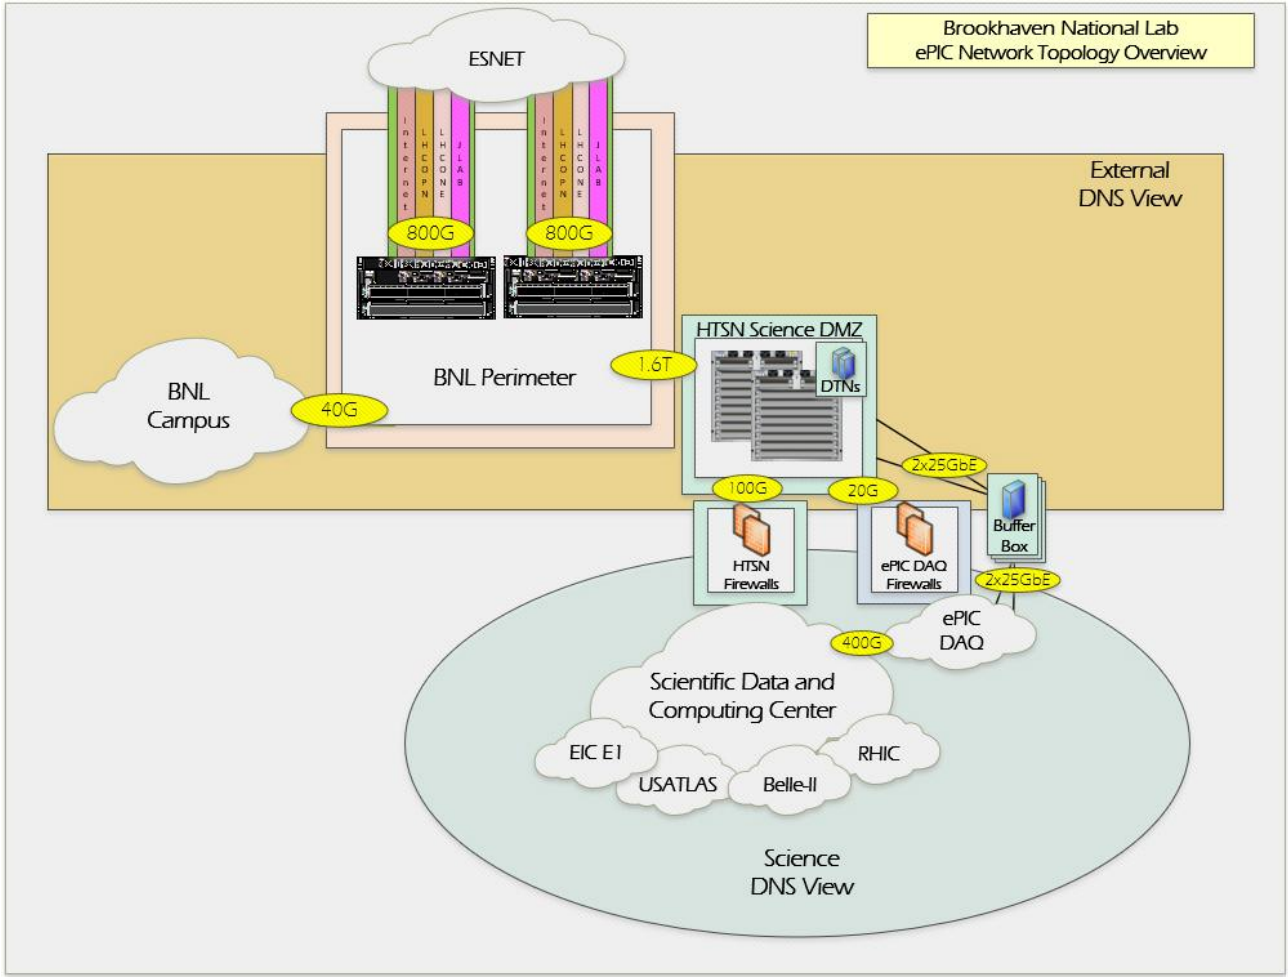

Brookhaven National Laboratory

ePIC Internet/WAN Network Topology Option 3 - (Least Preferred)

- Least preferrable option due to inherit security concerns.
- This topology would require two additional 400GbE circuits from ESnet to connect directly to the ePIC DAQ switches circumventing BNL's Network Perimeter and Science DMZ.
- Would require additional monitoring and mitigation infrastructures from Cyber Security and ITD Network Engineering.
- Allows the ePIC DAQ to transfer data to JLAB via a non-firewalled path but would still require high throughput firewalls in the event the VPLS circuits to JLAB fail.
- Would not allow the BNL E1 site/doors to utilize the virtual circuit to JLAB.
- Connects the ePIC DAQ to the SDCC via the HTSN Science Core with route filtering/ACLs (like sPHENIX)… Same as option 1.

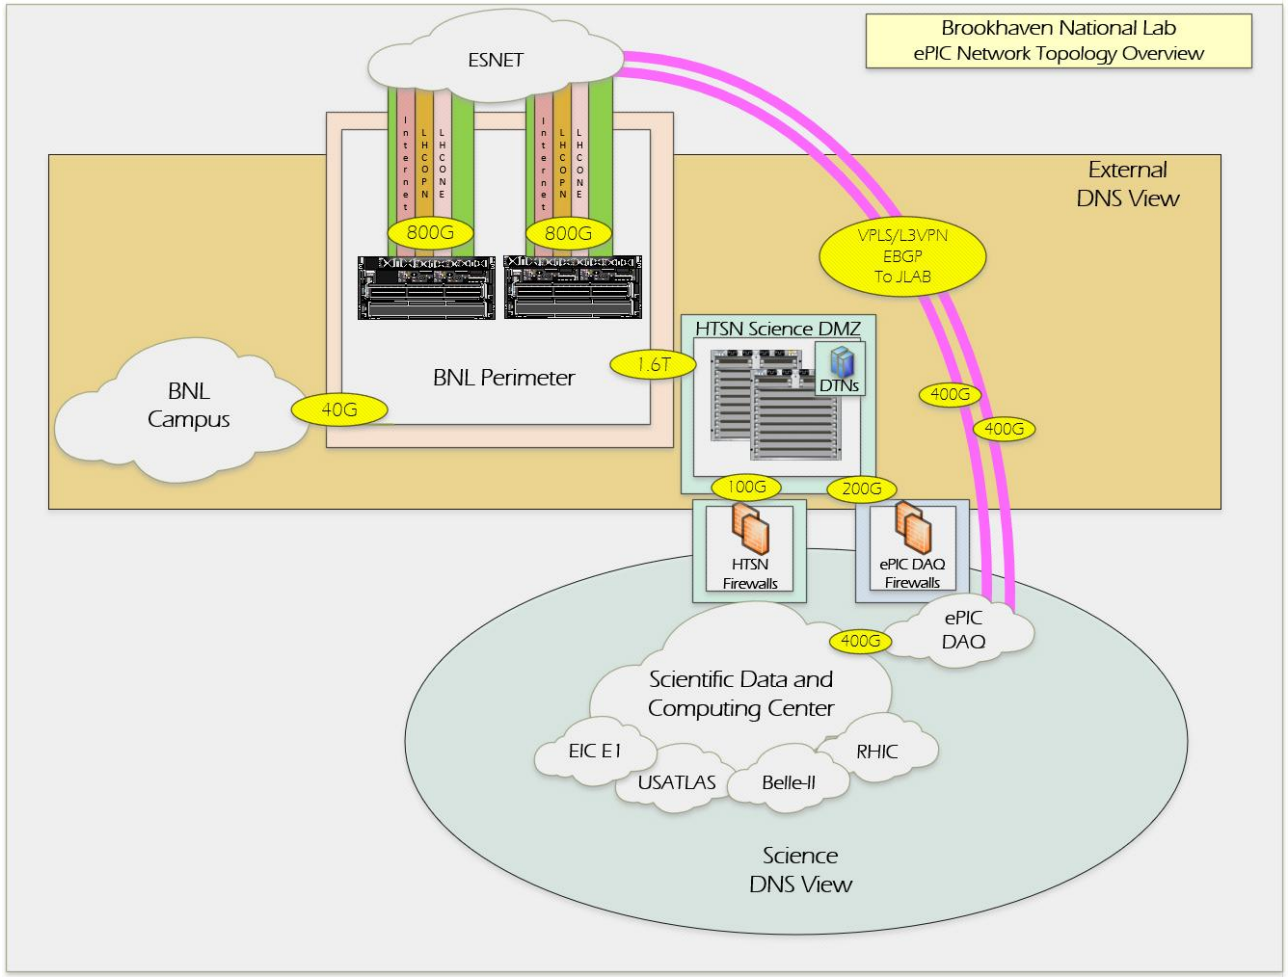

Brookhaven National Laboratory

Internal ePIC DAQ Infrastructure Placement Options

(As it relates to network infrastructure)

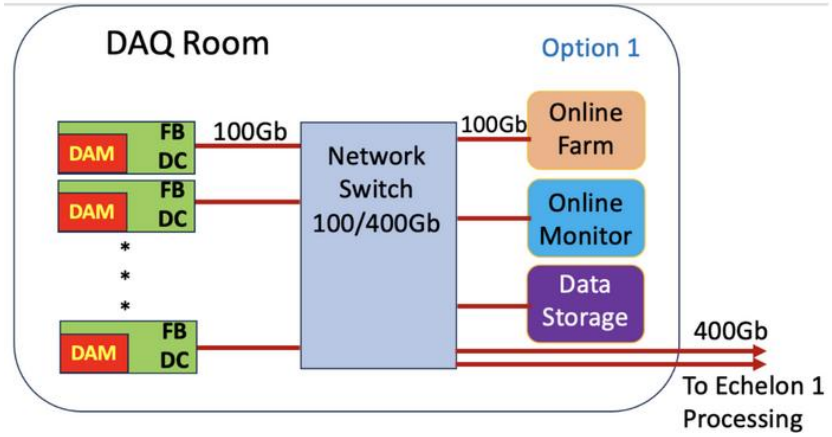

Option 1: Fully in DAQ room in IP6

- From a network infrastructure perspective, we have no major concerns aside from ensuring the building can adequately support the additional infrastructure.

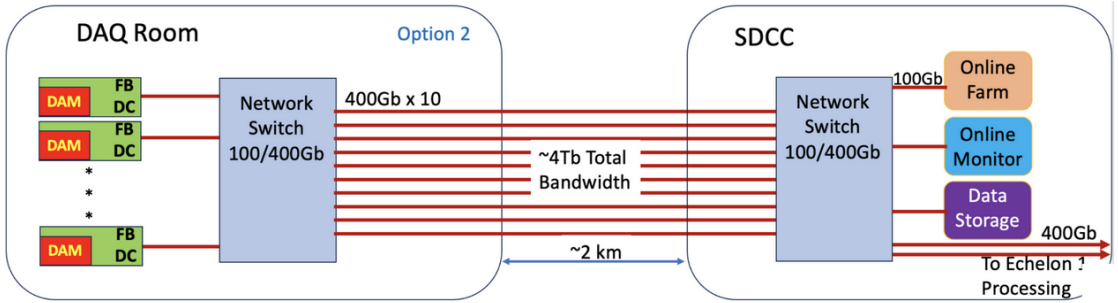

Option 2: Split DAQ Enclave

- From a network infrastructure perspective, this option is feasible; however, currently there are no diverse conduit paths between buildings 1006A and 725 (SDCC).
- If buffering or storage capabilities are unavailable at IP6, we recommend establishing a secondary, diverse path to Building 725.

Brookhaven National Laboratory

Fiber Path Between IP6 and SDCC

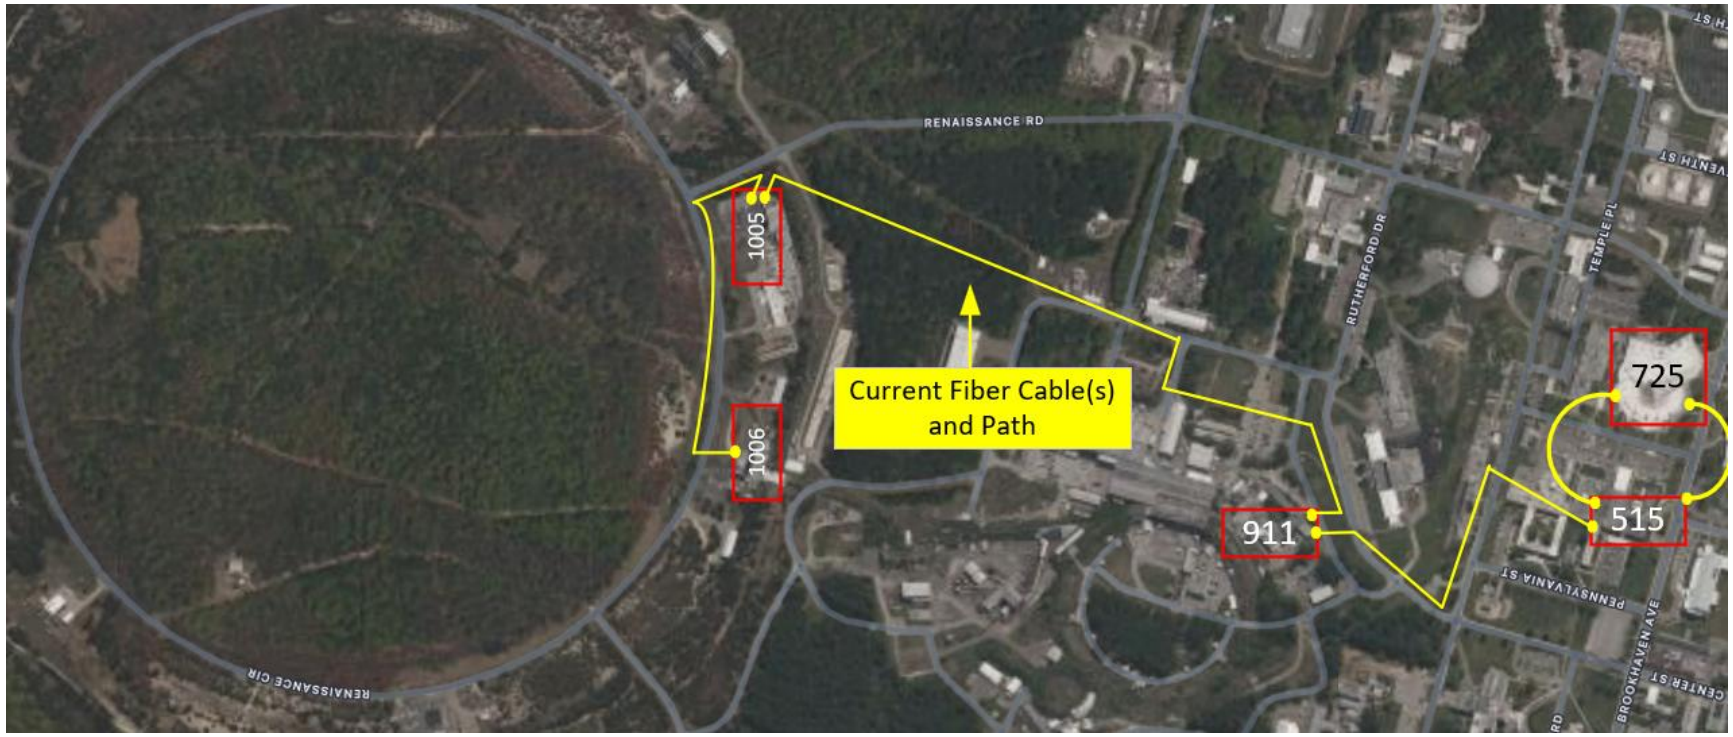

- Consists of a series of single multi-fiber cable runs, shown in yellow, connecting buildings 515, 911, 1005, 1006.
- The cables are terminated in patch panels in each of the buildings, and patch cables are used to connect the fibers between the cables.
- A break in any one of these cables would result in loss of connectivity from RHIC to the data center.
- Currently, approximately two-thirds of the fiber capacity between buildings 515 and 1005 has been utilized in support of RHIC

Internal ePIC DAQ Network Infrastructure – Option 1

- **IP-6 DAQ Network Infrastructure**
  - Could consist of two 4 slot Arista chassis' with:
    - 2x 24 port 400GbE line cards.
    - 2-4x 36 port 100GbE line cards.
    - ~3km from building 1006 to 515 on current fiber path. Will require the use of LR transceivers.

- **Building 725 DAQ Network Infrastructure**
  - Spine could consist of two 4 slot Arista 7804 chassis' with:
    - 2x 36 port 400GbE line cards.
  - Compute racks should utilize Top of Rack (ToR) Arista 7280CR3-32D4 or similar switches:
    - Switches provide 32x100GbE and 4x400GbE
    - This would allow us to connect the compute ToRs in any location within the 725 MDH or even external locations.
  - Storage head nodes will be redundantly connected to a centralized MLAG pair of Arista 7280CR3-32D4 or similar switches.
    - If the storage head nodes are 25GbE we may want to use a different type of fixed chassis switch.
    - Head nodes will be connected via fiber cabling back to the centralized Arista switches. This will allow the storage racks to be located anywhere within the 725 MDH.

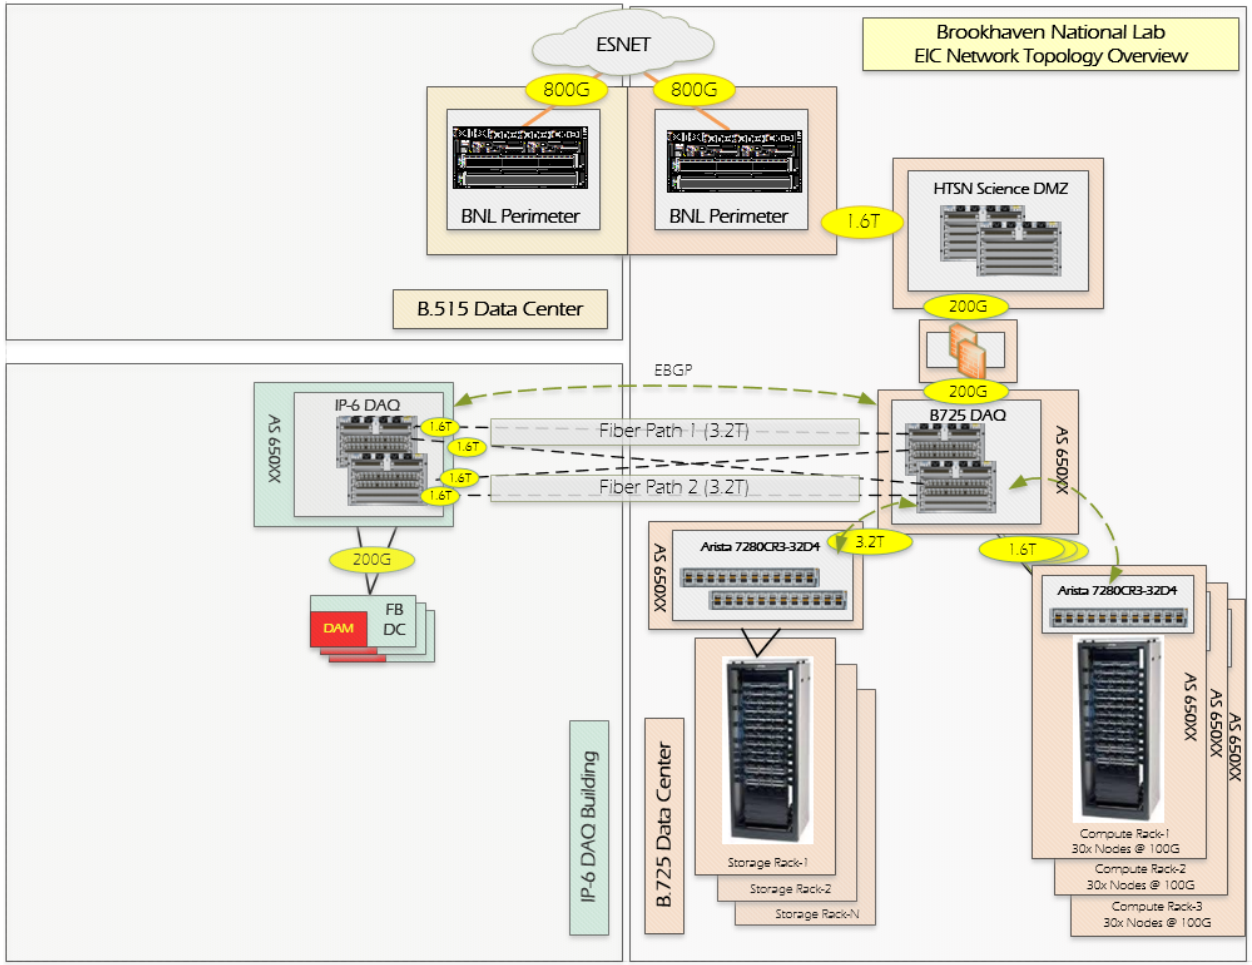

Brookhaven National Laboratory

Internal ePIC DAQ Network Infrastructure – Option 2

- **IP-6 DAQ Network Infrastructure**
  - Two Arista fixed chassis switches each supporting 24 port 400GbE.
  - Each rack containing DAM servers will receive a ToR Arista 7280CR3-32D4 (32x100GbE and 4x400GbE)

- **Building 725 DAQ Network Infrastructure (same as option 1)**
  - Spine could consist of two 4 slot Arista 7804 chassis' with:
    - 2x 36 port 400GbE line cards.
  - Compute racks should utilize Top of Rack (ToR) Arista 7280CR3-32D4 or similar switches:
    - Switches provide 32x100GbE and 4x400GbE
    - This would allow us to connect the compute ToRs in any location within the 725 MDH or even external locations.
  - Storage head nodes will be redundantly connected to a centralized MLAG pair of Arista 7280CR3-32D4 or similar switches.
    - If the storage head nodes are 25GbE we may want to use a different type of fixed chassis switch.
    - Head nodes will be connected via fiber cabling back to the centralized Arista switches. This will allow the storage racks to be located anywhere within the 725 MDH.

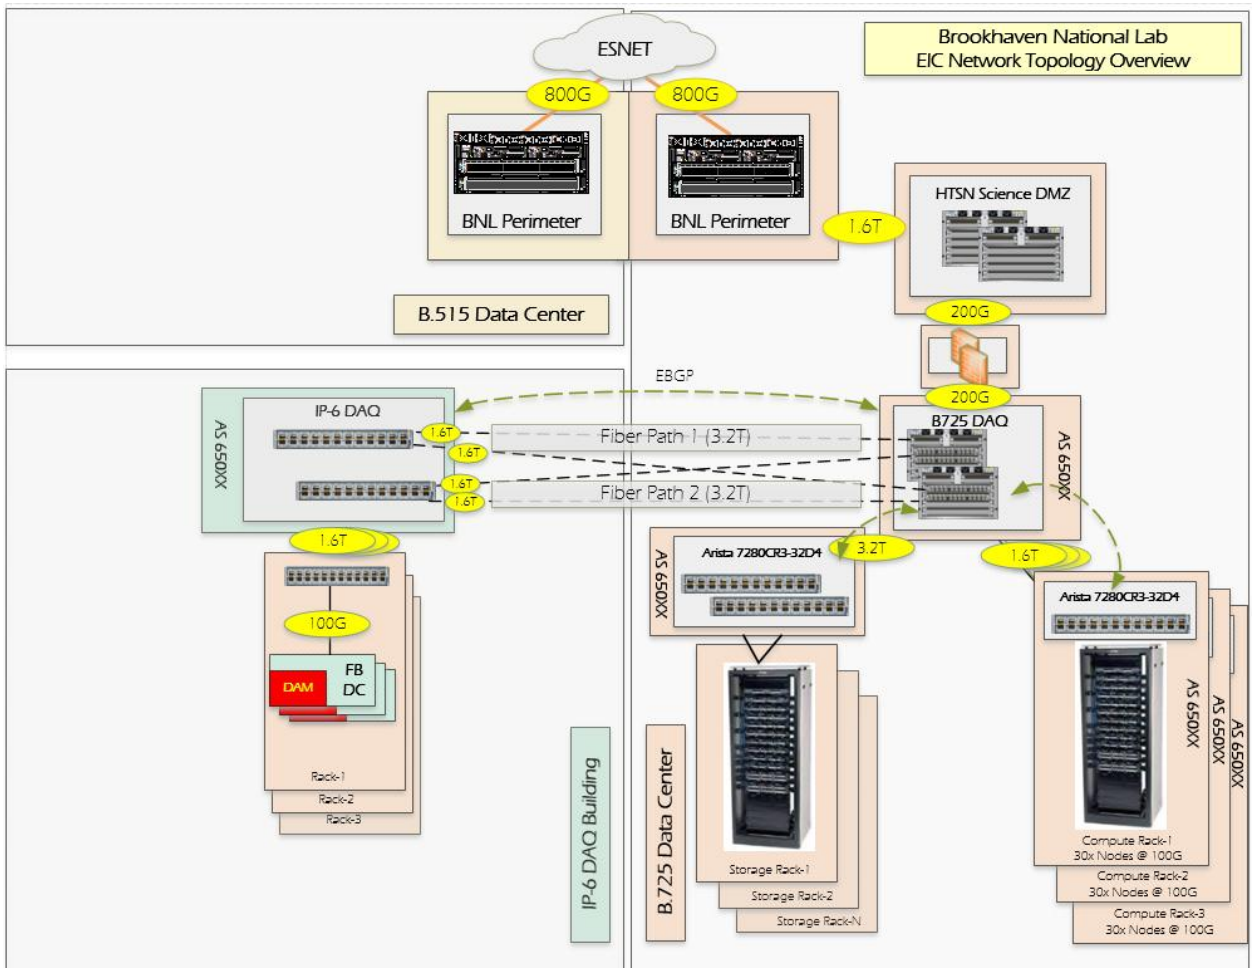

Brookhaven National Laboratory