Brookhaven National Lab

LHCOPN/ONE Fall 2017

KEK

Mark Lukasczyk

BROOKHAVEN NATIONAL LABORATORY

a passion for discovery

U.S. DEPARTMENT OF ENERGY | Office of Science

BNL Multipurpose Scientific Lab

- BNL is a multifaceted laboratory that hosts numerous scientific programs including:
  - **Basic Energy Sciences (BES)**
    - National Synchrotron Light Source II (NSLS-II)
    - Center for Functional Nanomaterials (CFN)
  - **Biological and Environmental Research (BER)**
    - Cryo-EM (FY18-19)
  - **High-Energy Physics (HEP)**
    - ATLAS Tier-1 Computing Facility
    - Belle-II Tier-1 Computing Facility
  - **Nuclear Physics**
    - Relativistic Heavy Ion Collider (RHIC)
      - STAR
      - PHENIX
      - RHIC Computing Facility

BNL's Unified Network Architecture

- Unified Network Perimeter
  - Two Diverse 100G circuits that peer with ESnet in New York City.
  - These 100G circuits are utilized by all scientific and administrative communities at BNL.
  - All traffic to and from BNL flows through either of these two circuits.
- Unified Science DMZ
  - Procured and architected by BNL to support open, high-speed WAN access for **all** scientific collaborations throughout the BNL campus.
  - All collaborations connecting to the BNL Science DMZ must have approval from BNL's Cyber Security group and have a valid cyber security plan which follows The Federal Information Security Management Act (FISMA).

Brookhaven National Lab's High Throughput Science Network

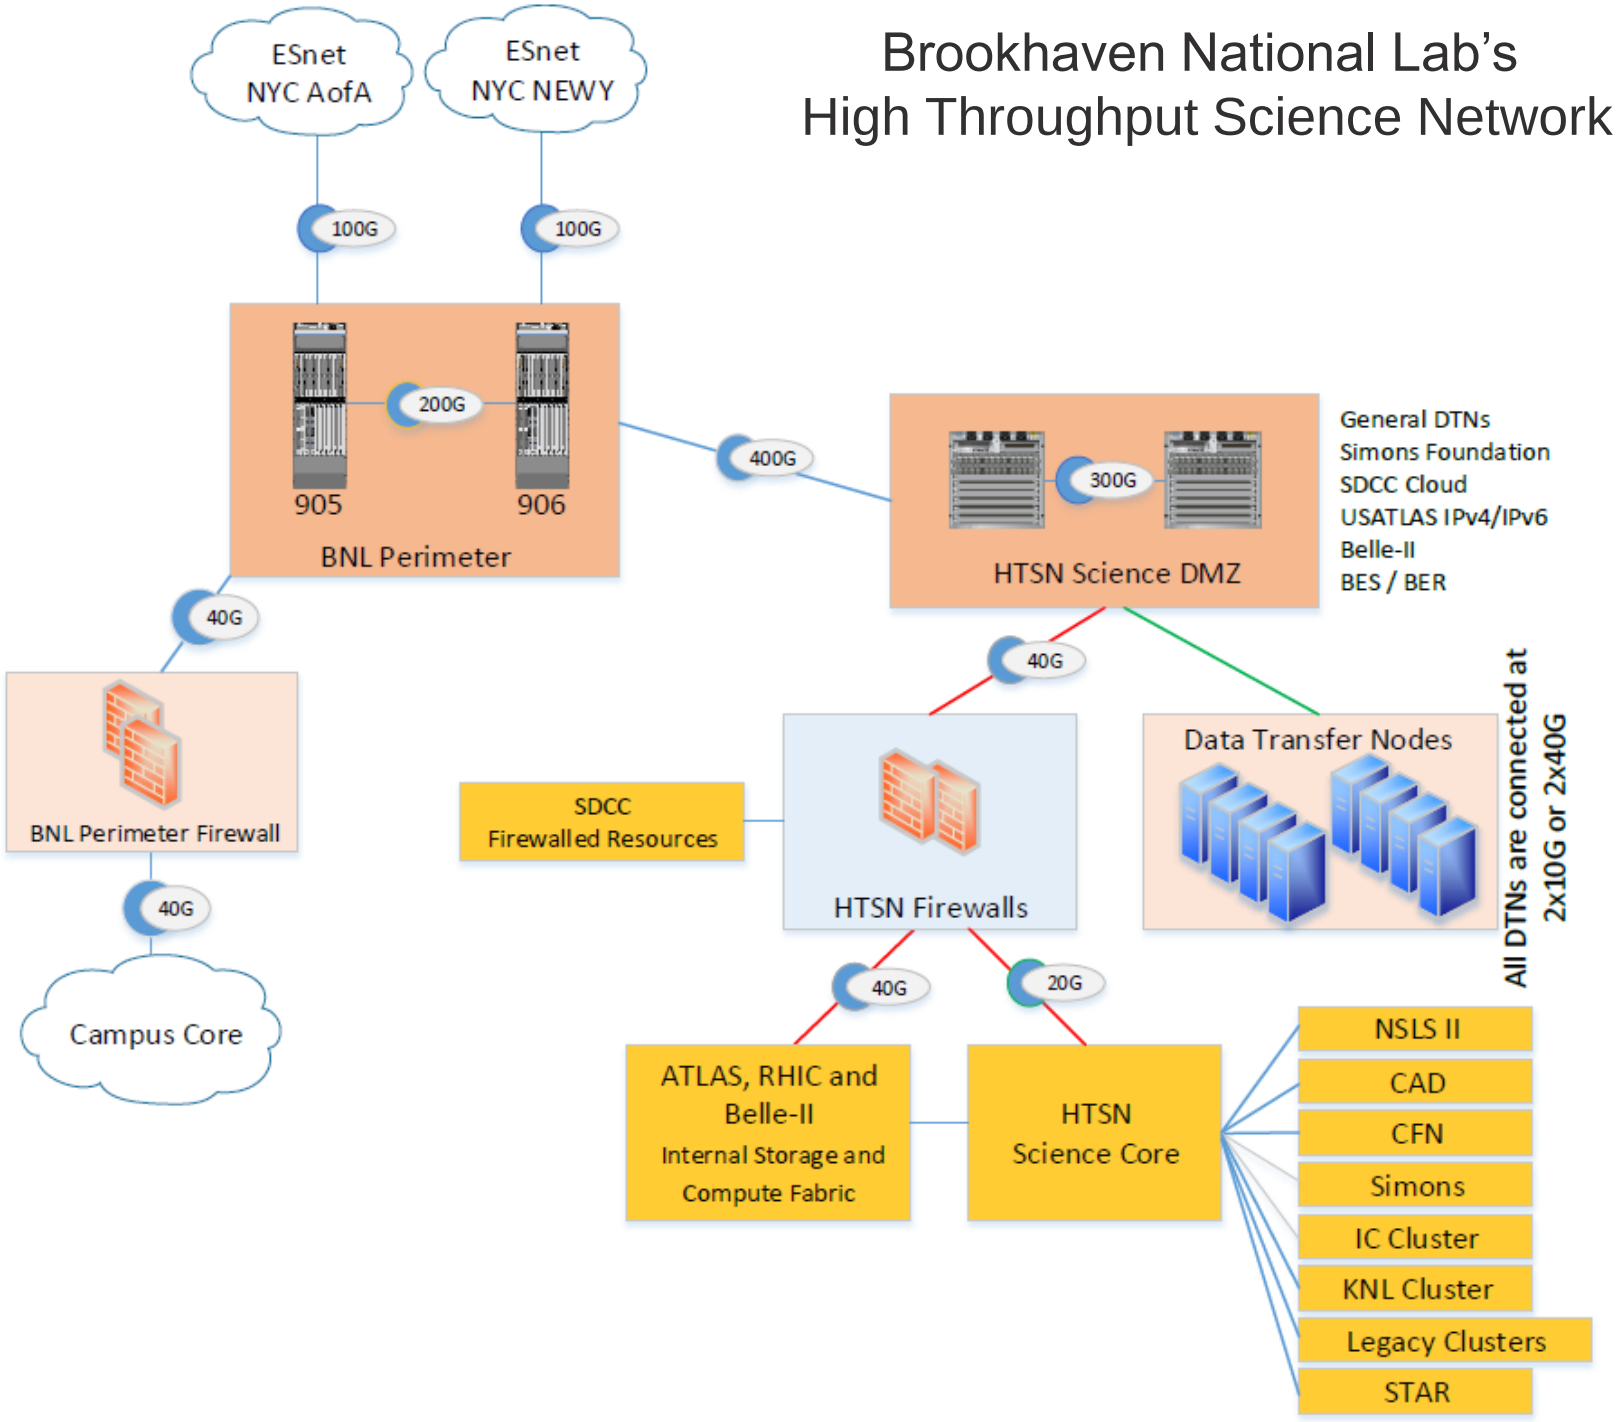

Complexities Operating Unified Network Perimeter and Science DMZ

- Adherence to the LHCOPN/ONE AUPs in a multipurpose lab utilizing a unified network perimeter becomes exponentially complex as scientific programs want exclusivity over a Virtual Private LAN Service (VPLS) or L3VPN circuits while utilizing BGP (e.g. LHCONE and LHCOPN).
- Complexities include:
  - Implementation of separate routing tables
  - Implementation of policy based routing
  - Increased cost in border routers
  - Increased operational complexities
  - Limits the number of vendors to choose from
- The addition of Belle-II to BNL demonstrates some of these complexities.

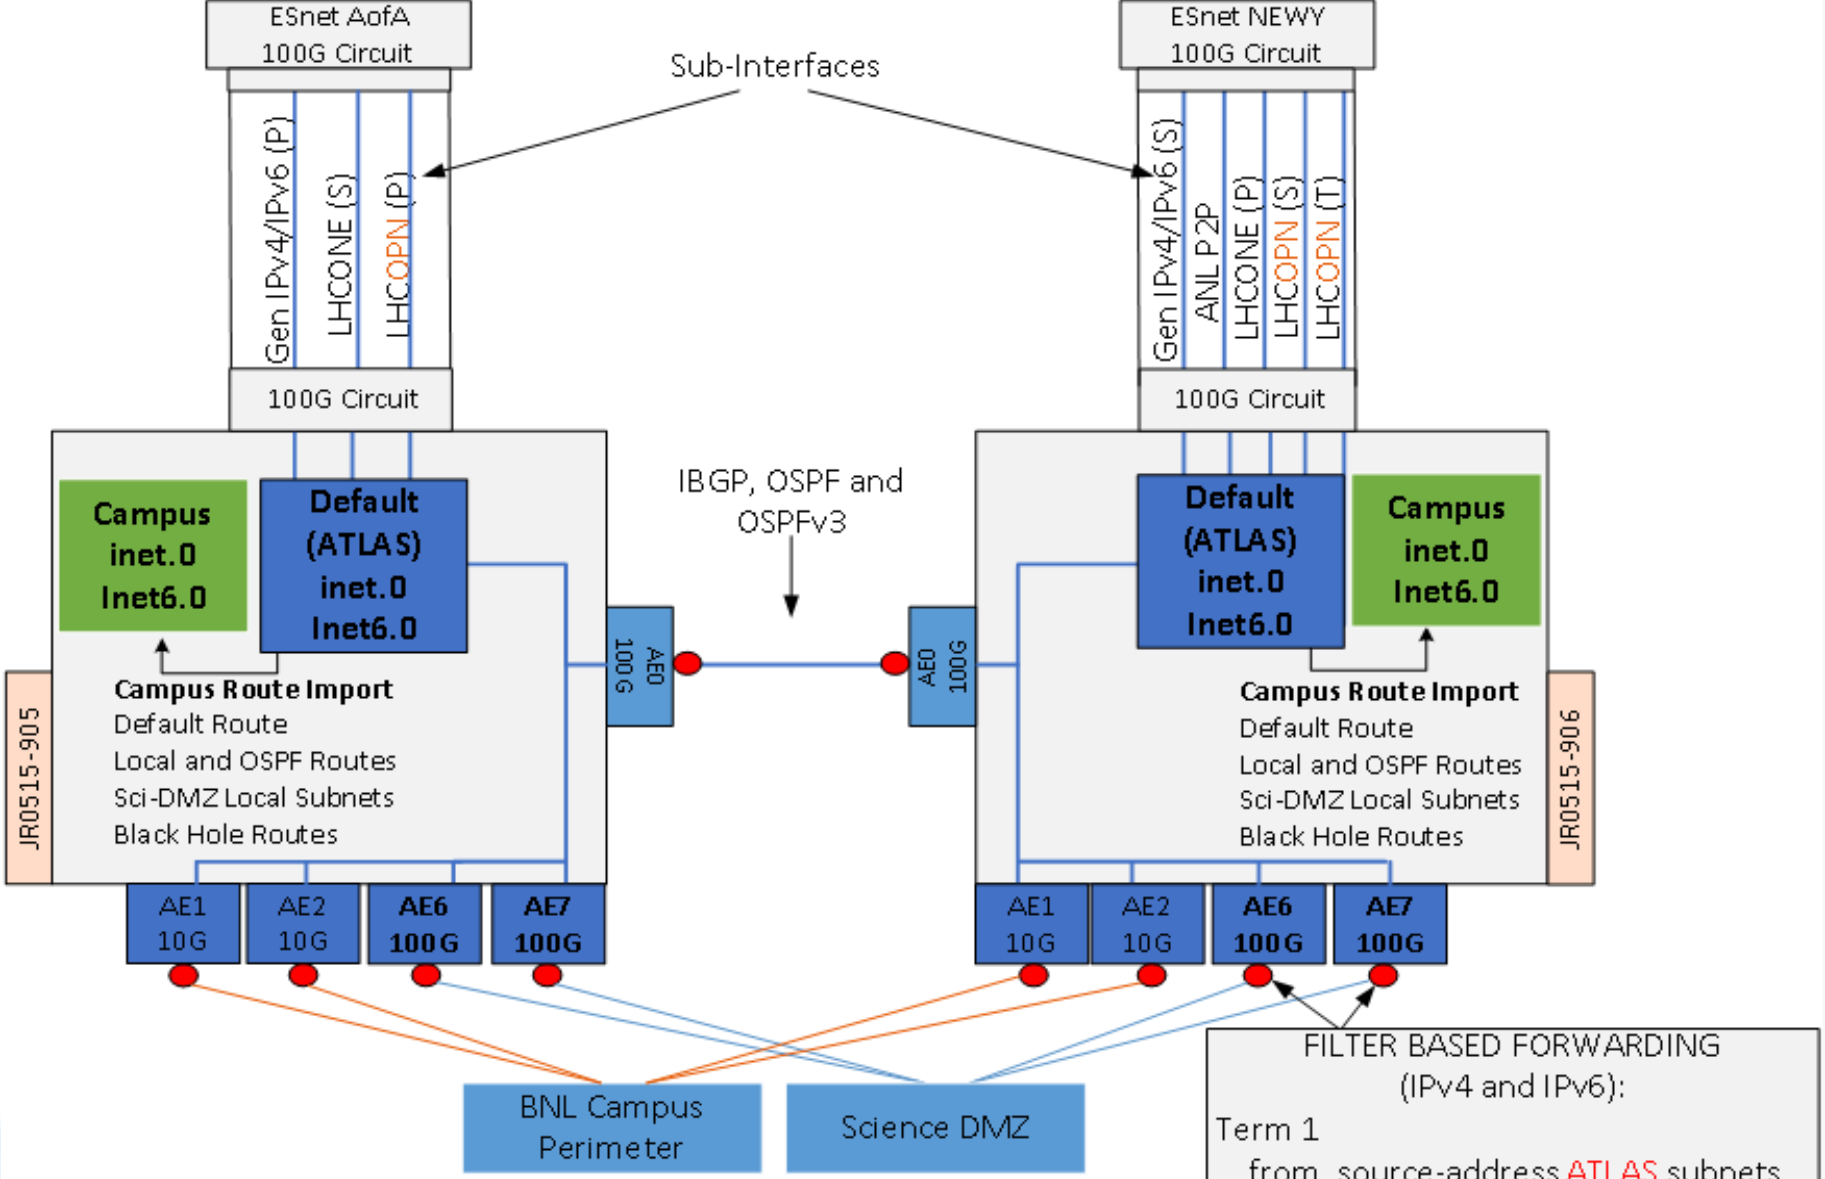

Current BNL Network Perimeter VRF Architecture

Belle-II Computing Center At BNL

- BNL is currently transitioning to a Tier-1 computing center for Belle-II.
- Belle-II requires access to general Internet and LHCONE
- Adhering to the LHCOPN/ONE AUP's while implementing Belle-II resources creates a multitude of complexities.

Belle-II Complexities Due to AUPs

- IP Addressing Complexities
  - Because current general purpose DTN subnets are not HEP they are not permitted on LHCONE or LHCOPN.
  - Can’t utilize USATLAS subnets since Belle-II isn’t permitted on LHCOPN.
  - BNL had to allocate dedicated direct assignment IPv4 space.
  - This WAN complexity also leads to system administrators needing to follow a complex matrix to determine what subnet their DTN’s should be assigned too.

Belle-II Complexities Due to AUPs

- Belle-II Routing Complexities at BNL
  - Since Belle-II isn't permitted on LHCOPN it can't utilize the current "Default" routing table as USATLAS does.
  - Since Belle-II requires LHCONE access it can't utilize the current "Campus" routing table.

Adhering to the current AUP is Complicated & Difficult

- Continue adding Perimeter VRFs to split traffic
  - Increases complexity. Requires additional EBGP (LHCONE) and IBGP/IGP peering's.
  - Still requires separate routing tables and Policy Based Routing.
  - Longer time to deploy
  - Locks us into very expensive hardware
- Use point solutions per Scientific program
  - Provide dedicated equipment, circuits and peering's
  - Adds significant costs for procurement, maintenance and management.
  - Longer time to deploy
  - Moves in the opposite direction of a unified network
  - Who is supposed to pay for this?

Belle-II's Current State at BNL

- At what point does the cost and complexity of these solutions no longer justify strict compliance to the AUPs?
- These AUPs are forcing BNL to migrate away from an "end site" architecture and begin to mimic a service provider environment.
- BNL has made an internal decision not incur the additional complexities of creating another Perimeter VRF or procuring a point solution.
- Given the short time frame for implementation, advertising Belle-II's address space through the LHCOPN was the only viable cost effective and manageable option.

BNL Scientific FY17 Traffic

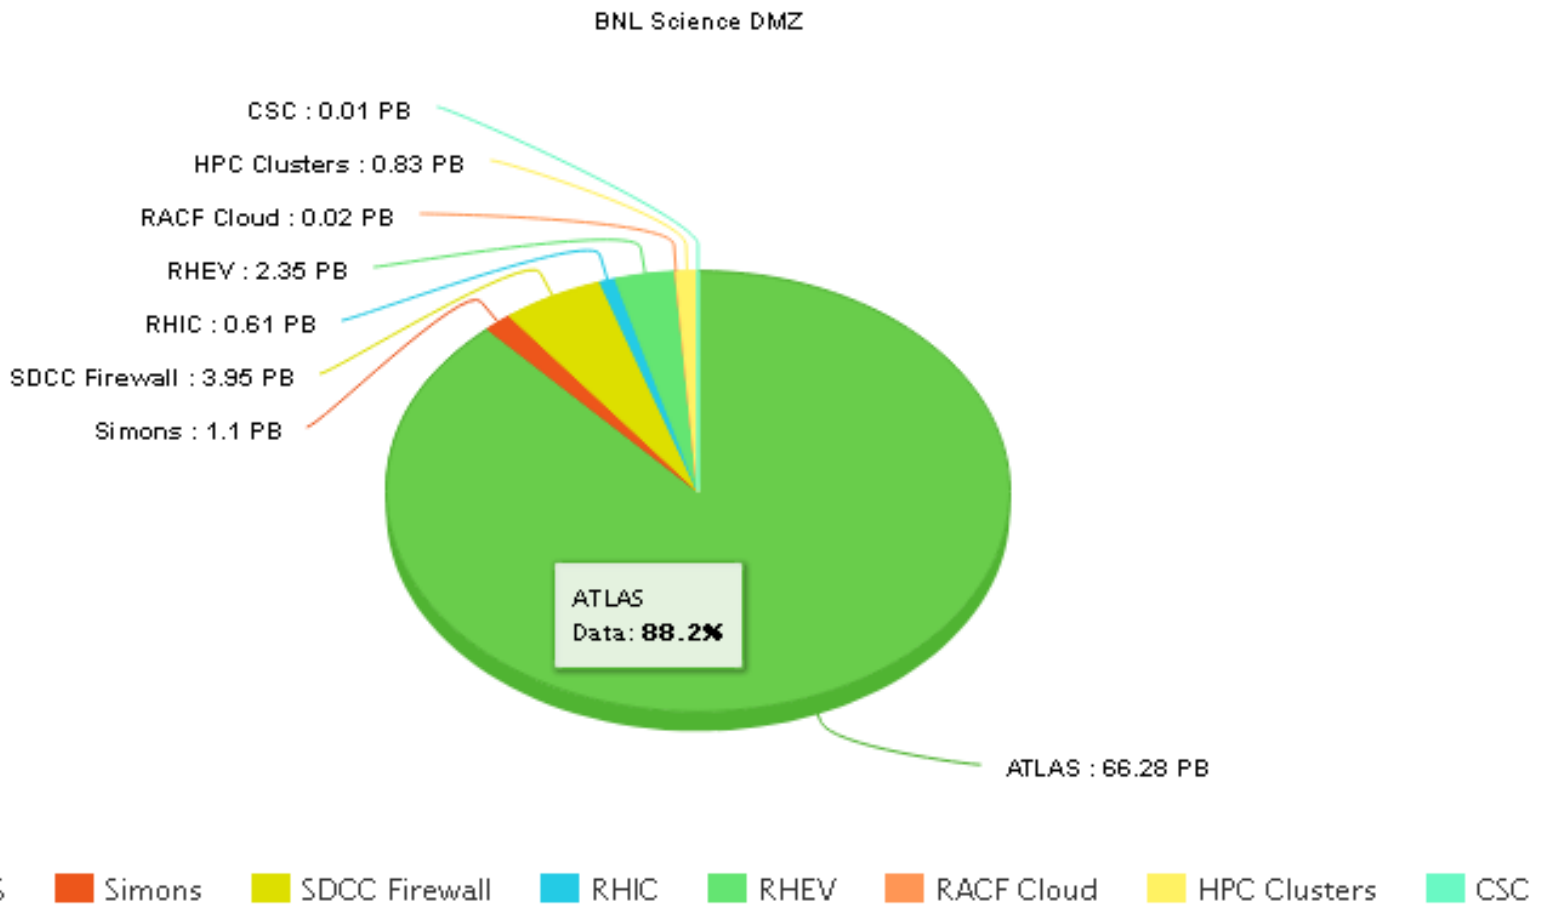

Costs Incurred to Support AUPs

- Given that commodity ASICs can now support full Internet routing tables (e.g. Arista), BNL is strongly investigating the possibility to reduce costs by utilizing these appliances on its Network Perimeter.
- Given BNL's FY17 scientific traffic, it isn't practical to keep procuring expensive network hardware to detour BNL's 12% of non HEP scientific data away from possibly utilizing LHCOPN/ONE.
- Given the fiscal climate and the growth of additional scientific programs at BNL, these AUPs will only continue to increase operational complexities and drive hardware costs higher.

BNL's Aspiration

- Remove the complexities of source based routing and revert to destination based routing which will eliminate operational complexity and reduce costs.