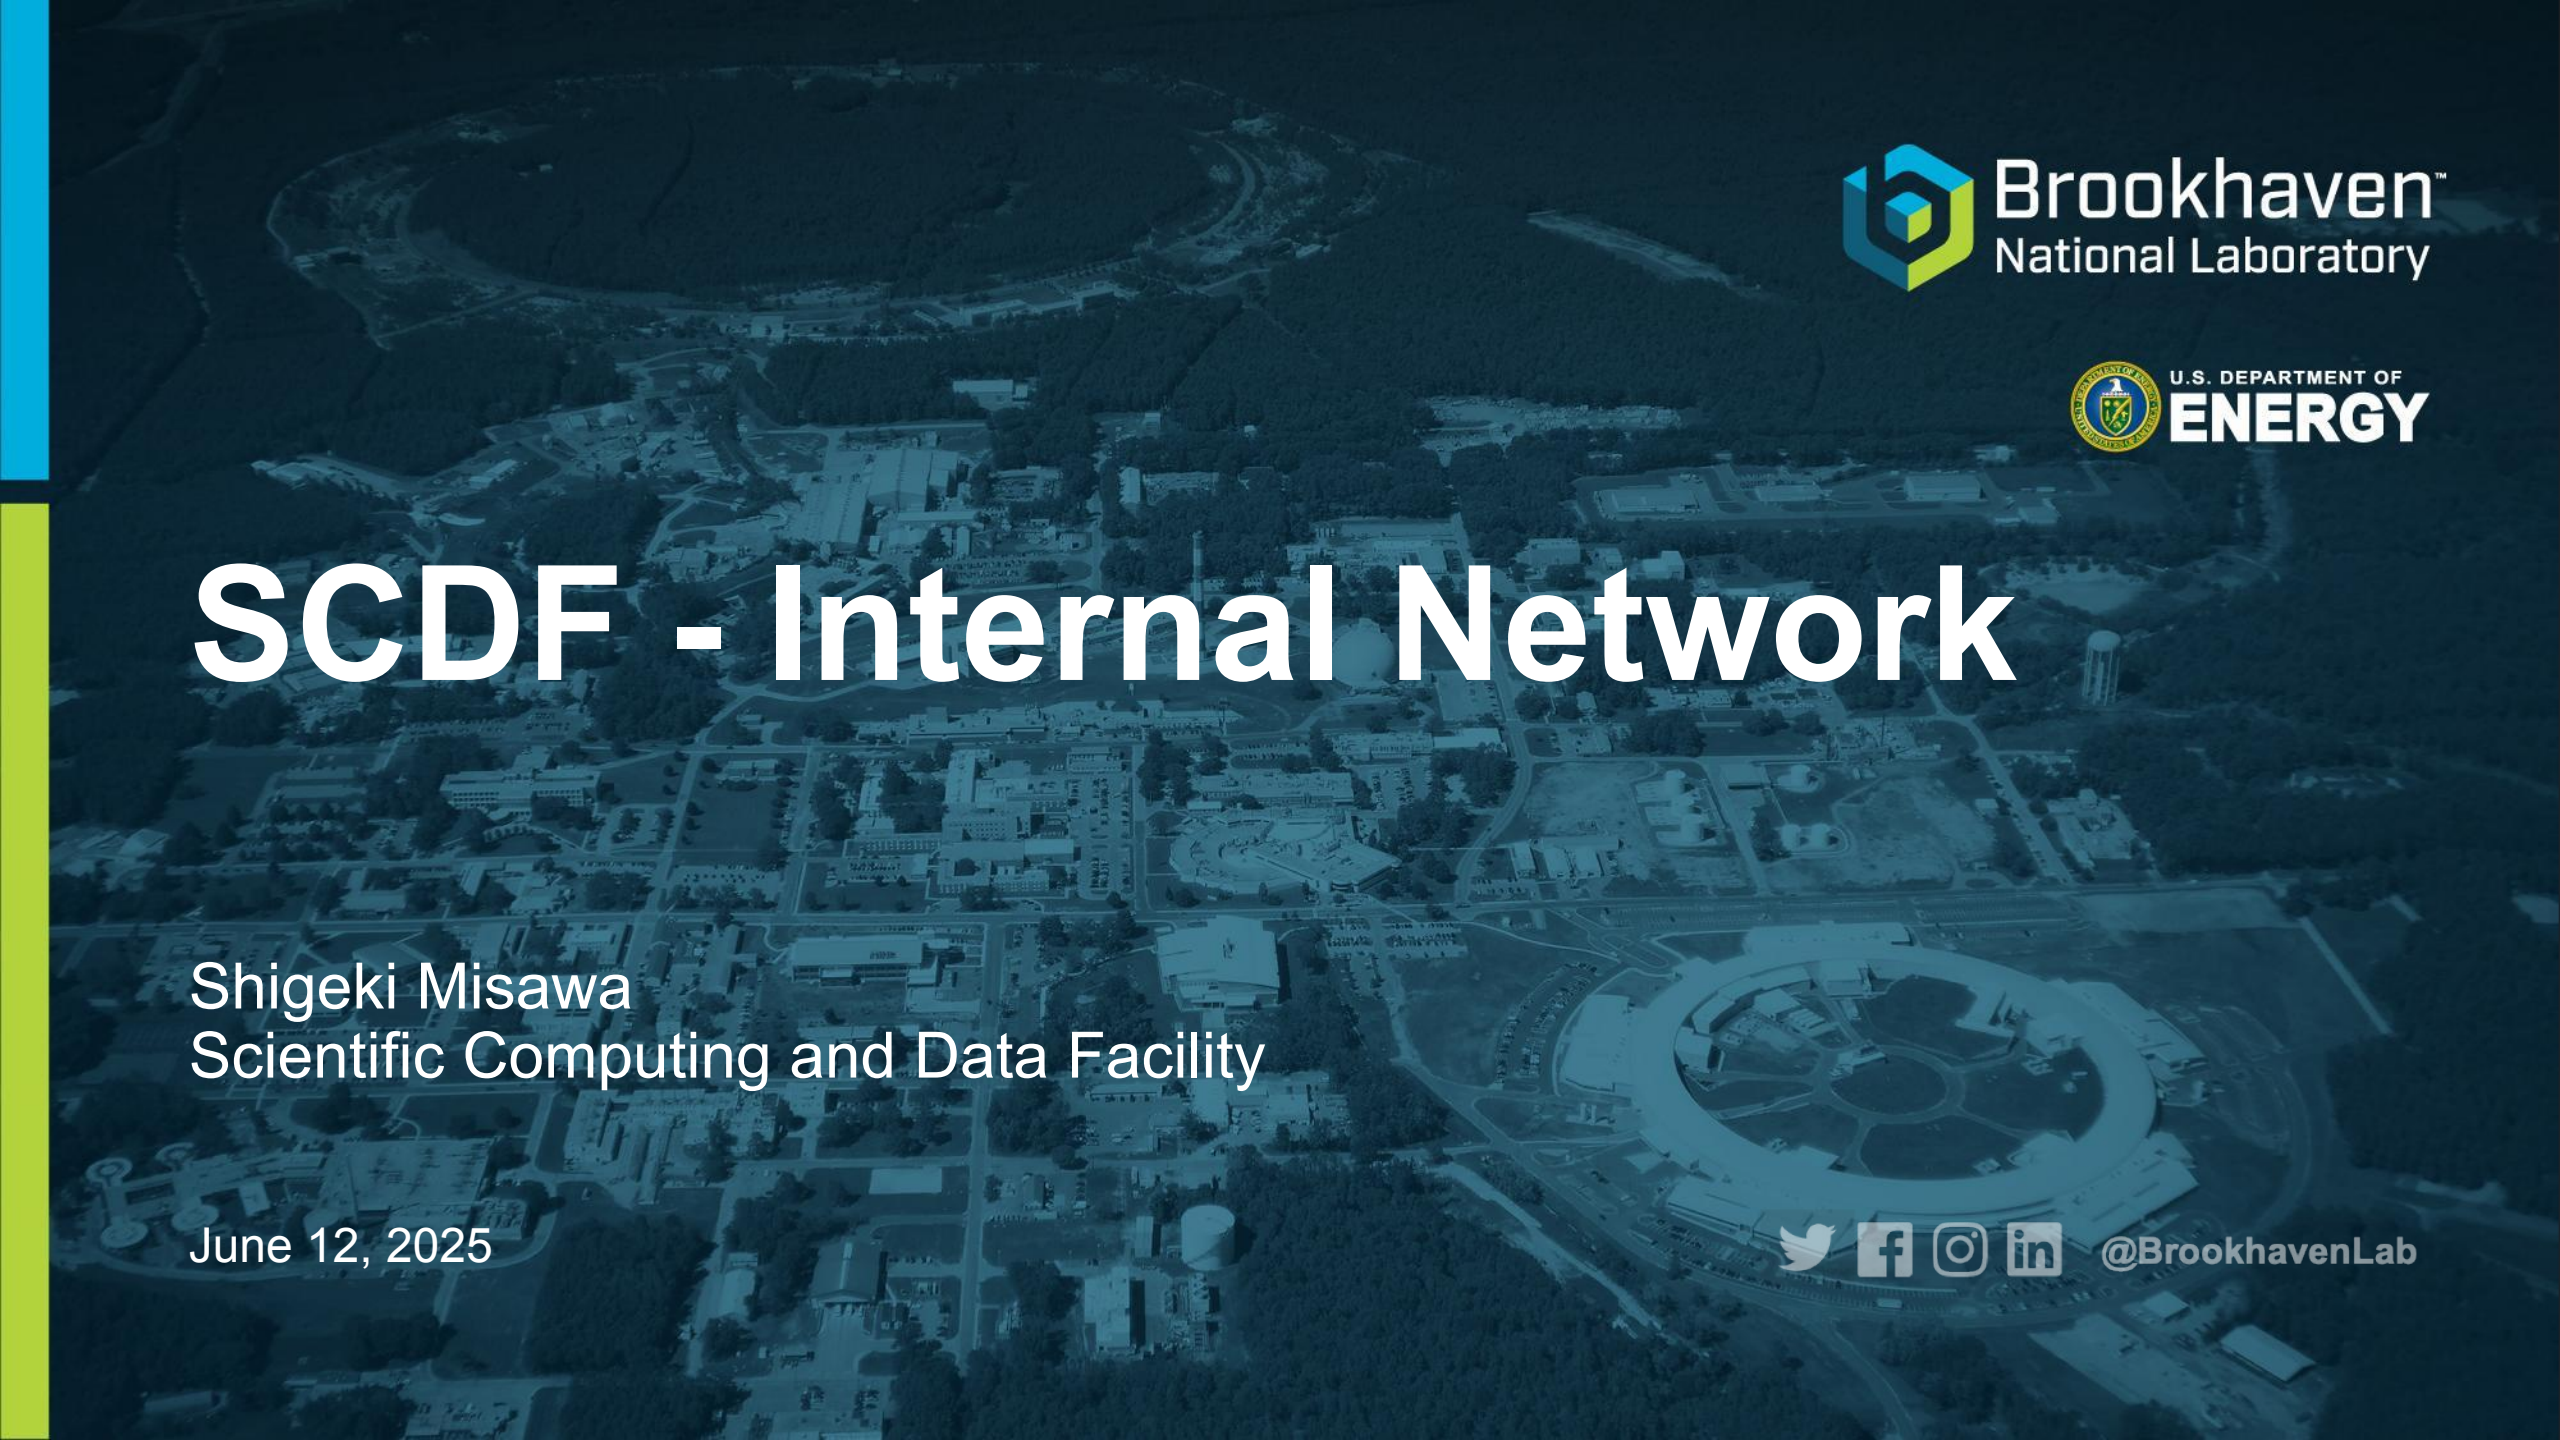

Brookhaven National Laboratory

U.S. DEPARTMENT OF ENERGY

SCDF - Internal Network

Shigeki Misawa
Scientific Computing and Data Facility

June 12, 2025

@BrookhavenLab

Next Gen Data Intensive Science

- 10x increase in data volume/rate from next generation NP (sPHENIX) and HEP High Luminosity-LHC experiments (HL-ATLAS)
  - 100s of petabytes generated per year
- Proliferation of data intensive experiments from new groups (e.g. NSLS-II, CFN)
  - Smaller and more numerous than NP/HEP collaborations
  - Geographically dispersed in the BNL campus
  - Limited infrastructure to support in situ compute/storage resources

Network fo Data Intensive Research

- High Throughput Science Network (HTSN) architecture, conceived in 2013 to support data intensive research.
  - Science DMZ
    - Termination point for high bandwidth WAN connectivity
  - Science Core
    - High bandwidth, “frictionless” network to carry scientific data within the BNL campus
    - Connects scientific instruments directly to data center compute/storage resources
    - Interconnects compute/storage resources within the data center

Science DMZ WAN Connectivity

- Science DMZ and Core are completely decoupled from the BNL campus network
  - Science and campus connect only at the BNL perimeter
- Science DMZ supports IPv4 and IPv6
  - Critical for international groups
- Dedicated 100 Gbps firewall protects Science Core network from WAN
- DTN's on the DMZ enable high bandwidth (100s of Gbps) data transfers to/from the WAN

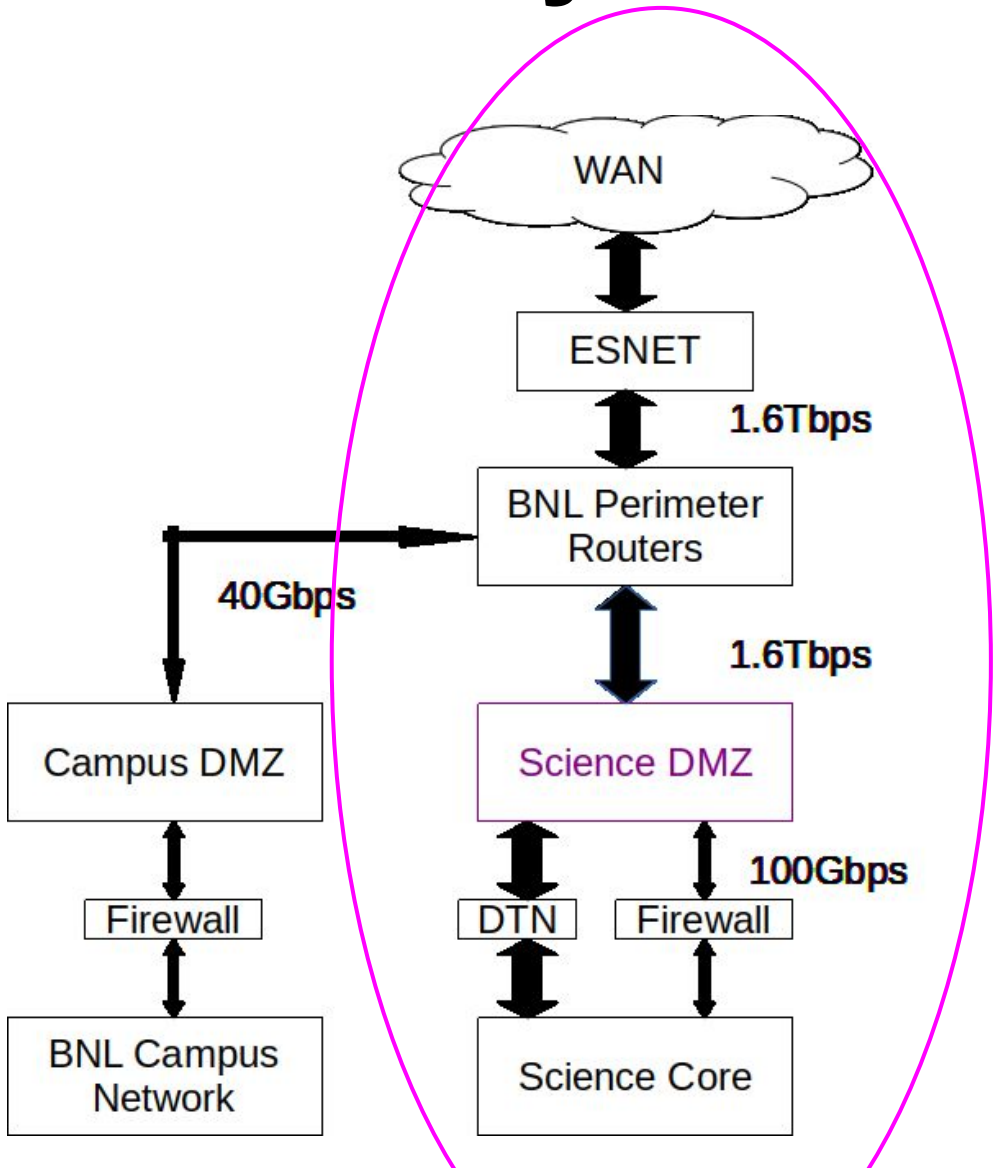

Science DMZ Rationale

- High bandwidth WAN connectivity is needed to access resources outside of BNL
  - e.g. DOE/NSF HPC facilities, public cloud, other institutes
- Computing trends are making access to external resources more important
  - Proliferation of new services
  - Increasingly targeted hardware resources, e.g. GPU systems
- Science DMZ is heavily used by ATLAS and other HEP programs. Use by EIC experiments is expected in the future

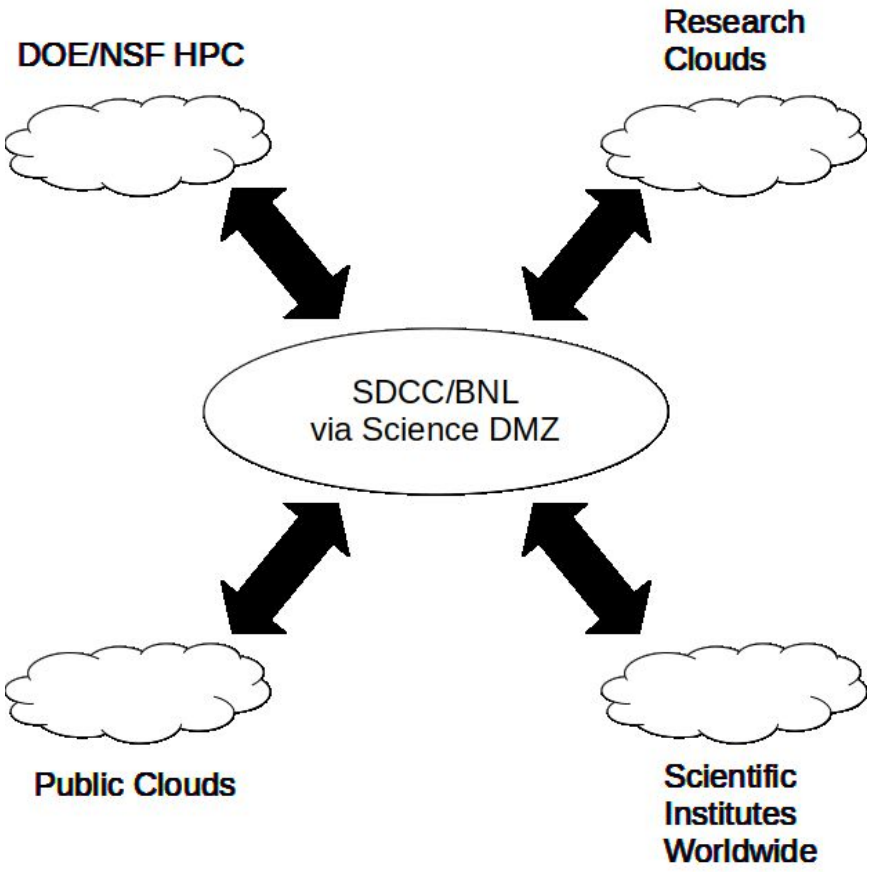

Science Core Rationale

- Connect “remote” sites to the data center
- Eliminate NxN inter-site visibility to enhance security
  - Assume remote sites are “untrustworthy”
- Keep network “frictionless”
  - No firewalls
  - Avoid use of DTNs
  - Network throughput limited by link speed

[1] Upgradable to 400 Gbps per fiber pair

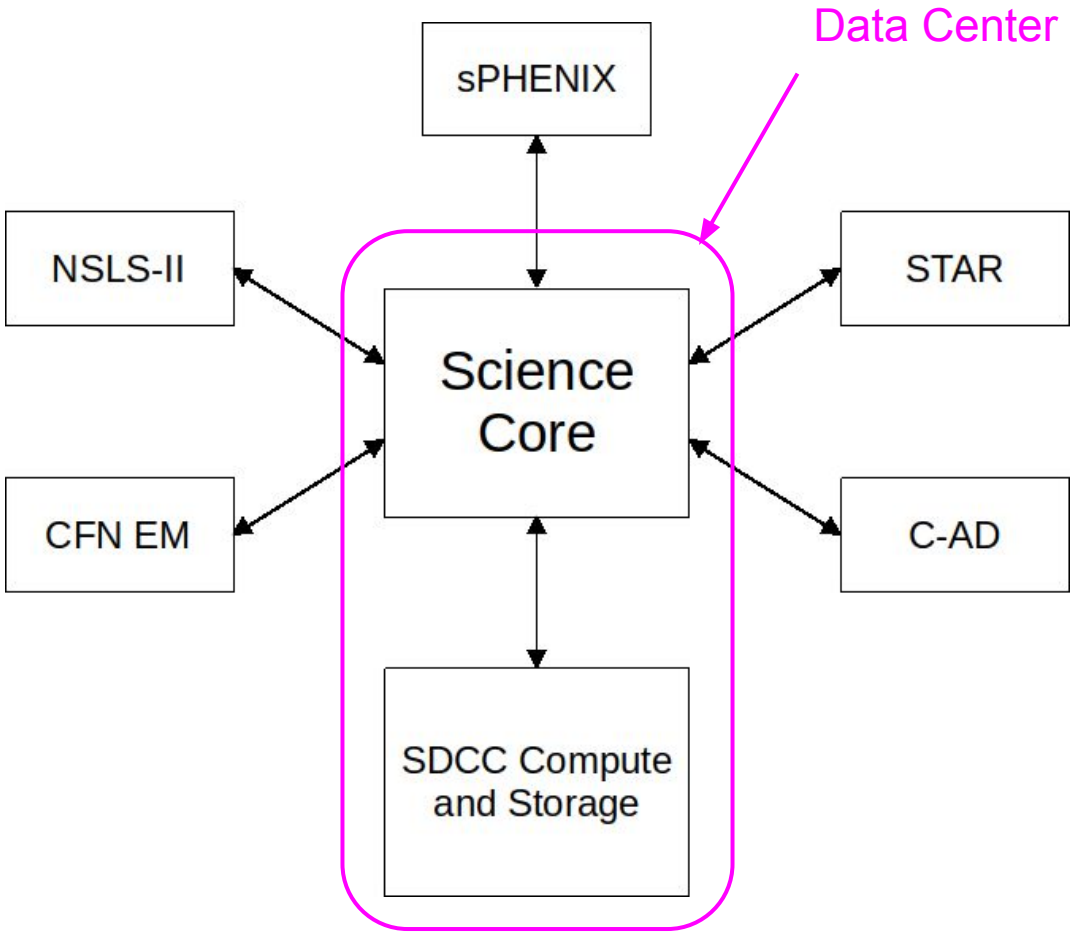

Remote Site Connectivity

Science Core Connectivity

- Each remote site is an AS (autonomous system)
- Use BGP route advertisements to limit reachability between different sites and to resources at the SCDF
  - Eliminates NxN visibility
  - Allows control of bi-directional communication between remote site and data center resources

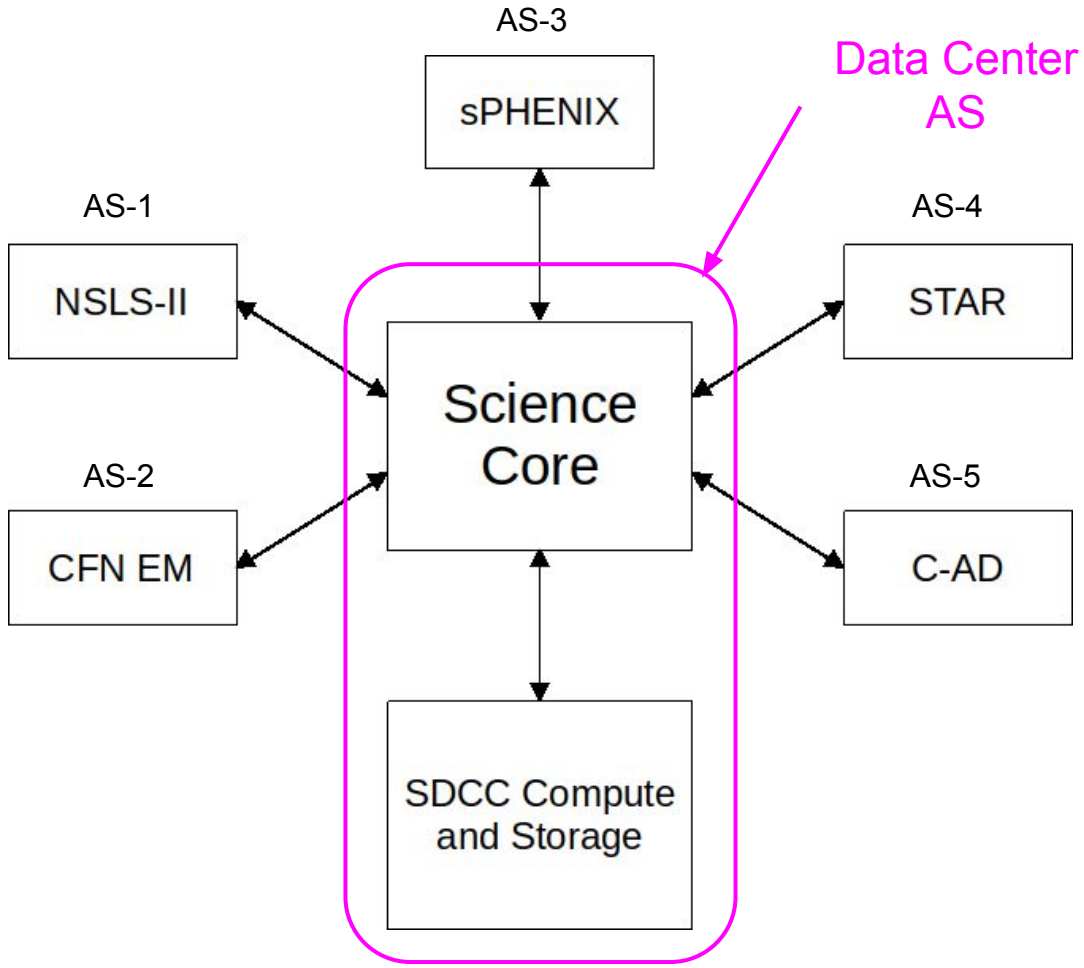

External Connectivity

Science Core (Data Center Connectivity)

- Provides full connectivity between resources inside the data center.
  - Data flows between storage and compute are isolated within the SCDF spine and leaf network
  - Data flows to other, mostly lower bandwidth SCDF services routed through Science Core
- SCDF internal network is IPv4 with gradual introduction of dual stack IPv4/IPv6 just starting (for ATLAS)

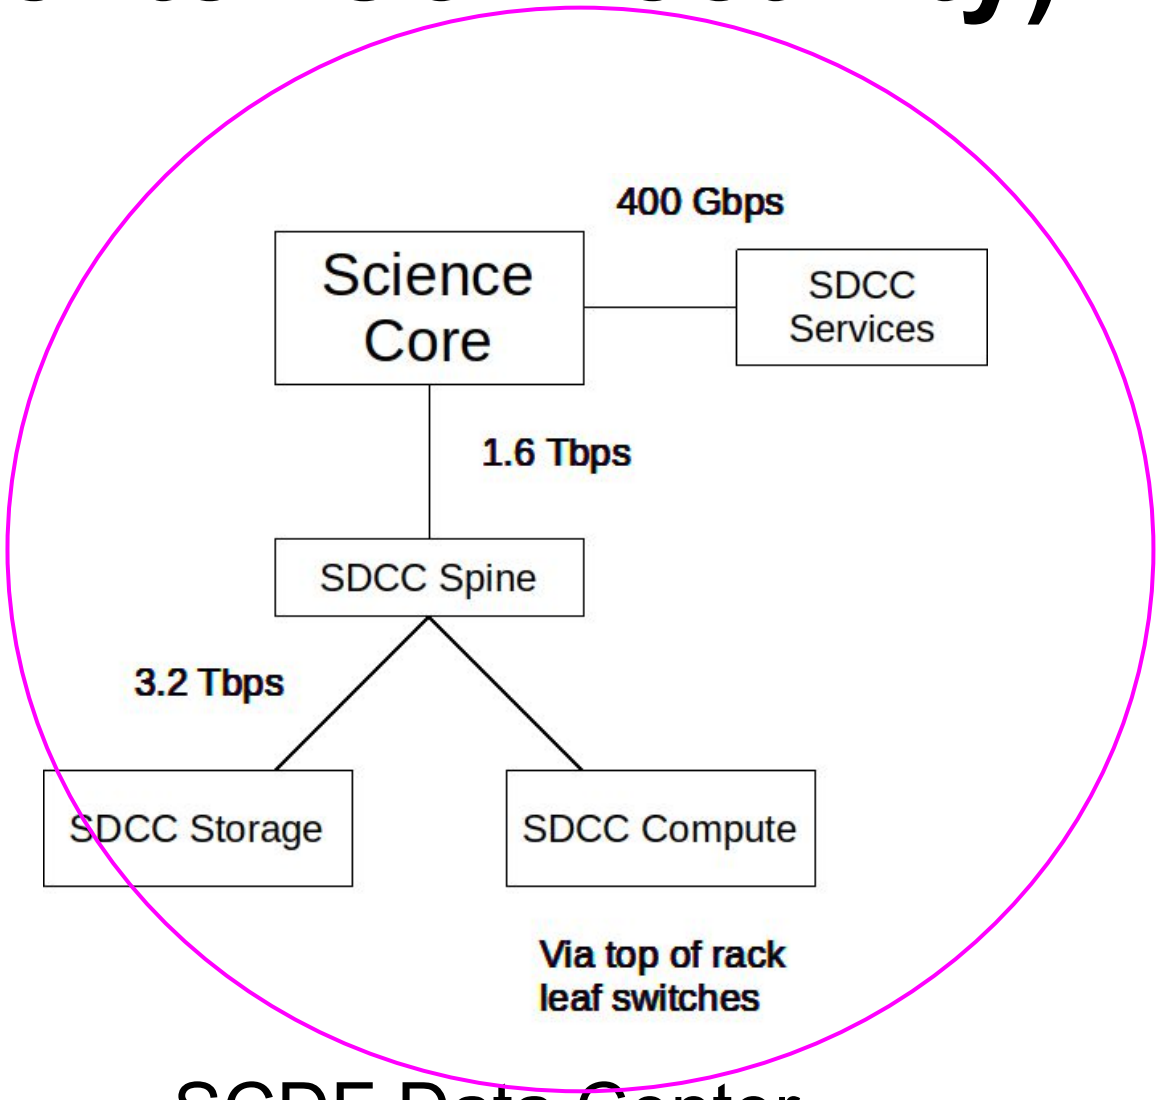

SCDF Data Center

Science Core Proof of Concept (~ FY16)

- CFN E-TEM
  - Direct writes to GPFS at SDCC from CFN systems
  - Analysis of GPFS resident data on BNL Institutional Cluster
  - WAN transfers via sftp/Globus

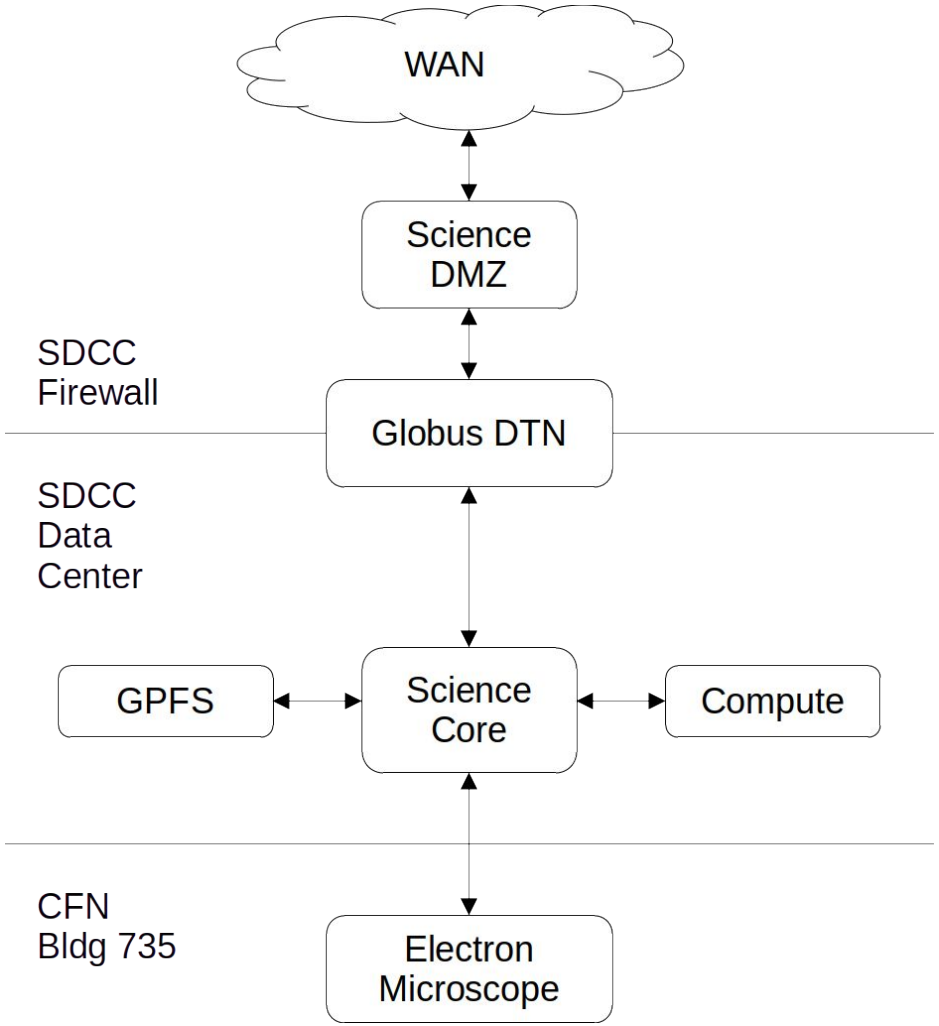

Remote Access to Storage Now in Production

- Direct access to data center storage through Science Core from remote sites now widespread
- Storage provided varies depending on requirements
  - Industry standard file sharing (NFS/SMB)
    - NSLS-II, CFN EM
  - Scale out parallel file systems (disk)
    - sPHENIX, CFN EM, NSLS-II
  - Scale out bulk data storage (disk)
    - ATLAS (LAN/WAN)
  - Nearline and archival storage (tape)
    - sPHENIX, STAR, ATLAS

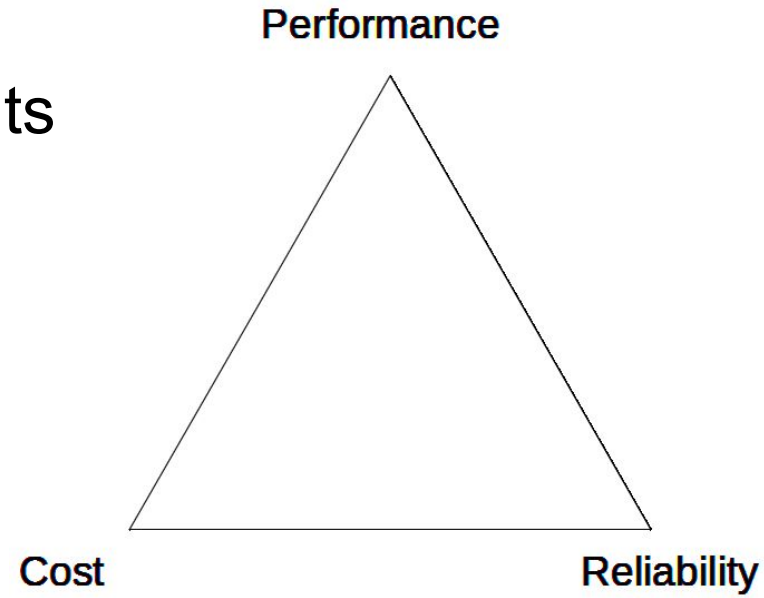

Requirements

Concluding Remark

- Science Core and DMZ creates new opportunities
  - Allows computing resource to be brought closer to the experimental apparatus
  - May make resources available to research group that might not otherwise be able to afford if they were worked independently
  - Makes resources outside of BNL accessible to experiments
  - Potentially opens up avenues of research that would otherwise not be possible
- Mission of the SDCC is to partner with research group to enable data intensive research.