BNL HDRMX
Network & Computing

Robert Petkus, May 13, 2016

Overview

- Computational cluster, storage, and networking
- Beamline architecture
- Current work and near-term plans
- Lab-wide computing initiative and our role

Cluster, Storage, and Networking

- NSLS-II deployed a small cluster and storage system with the following objectives:
  - Create an environment to develop tools for:
    - Data transfer and acquisition
    - Job scheduling and workflow management
    - Computational analysis
  - Assess actual beamline data rates and expectations
  - Provide resources for:
    - GPU computing
    - Parallelization
    - Mid-term storage

Cluster, Storage, and Networking

Prototype environment (10 nodes):

- Computation (10G) (HP DL360p Gen 9)
  - 4x high-clock CPU (2x 6C 3.4GHz)
  - 5x high-core CPU (2x 18C 2.3GHz)
  - 1x graphics system (2x Telsa K80)
- Local node resources:
  - SSD
  - SanDisk Fusion ioMemory (SS PCI-e)
  - RAMdisk
- Job scheduling
  - HTCondor + Univa
- 40G + IB

Cluster, Storage, and Networking

- Storage (General purpose configuration)
  - ~800TB storage (Netapp E5500)
    - 18x RAID 6 (8+2) 6TB SAS
  - 3x storage cluster nodes (NSDs)
    - 40G network → clients
    - 56G IB → storage
    - GPFS filesystem

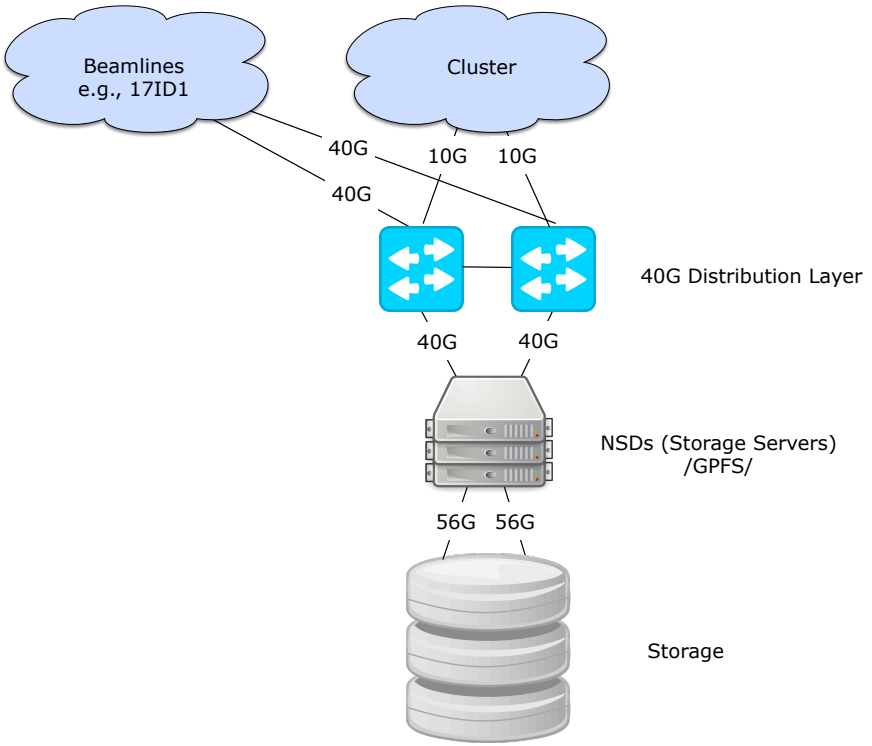

Beamline Architecture

Typical beamline today:

- 2x links to network core (2x 1/10/40)
- 2x analysis/storage servers + 100-300TB local storage (GPFS)
- Data buffered 1 day → 3 months then streamed to central storage

FMX:

- 2x 40G links to core
- Limited local analysis / storage
  - Data direct from detector → central storage

Current work, near-term plans

- While the prototype cluster was configured as a baseline to support general needs, FMX has shaped it's development
  - GPFS and kernel optimizations
  - Job scheduling with Univa
- Near term plans include:
  - 40G/IB compute nodes (10G is saturated)
  - Solid state write-read buffering
  - Storage expansion and reconfiguration
  - Storage specialization / tiered storage
  - Inter-institutional data transfer
  - Use of the institutional cluster

BNL Computing Initiative

- Plans are underway to build a class-leading data center in the early 2020s for the whole lab community
- NSLS-II has space reserved to meet expected workloads
- The Center for Data-Driven Discovery is staging a general purpose cluster
  - NSLS-II will participate in data transfer and computation studies to ensure success
  - Beamlines will be able to utilize these resources