Infrastructure for Integrating Experimental and HPC Facilities

Mario Cromaz, LBNL

NSAC LRP Town Hall - Nuclear Structure, Reactions and Astrophysics

Distributed Experimental Readout and Analysis

- Changes in landscape:
  - Advances in optical networks -100Gb local switched networks common-place, 400Gb wide area networks emerging
  - HPC, commercial clouds provide unprecedented compute capability

- The Idea - Real time network-based detector readouts/analysis pipelines:
  - Push high-performance IP networks into detector electronics - take advantage of rapid industry development, reduce costs and development time
  - Streaming readout - send events (not files) with network-level latencies (10's - 100's ms)
  - Opportunity to do higher-level analysis of experiments in real time - digital twins where experiment / simulation-theory calculations run concurrently

Network-Based Data Pipelines

Example: GRETA data pipeline (480k evt / s, 4 GB /s)

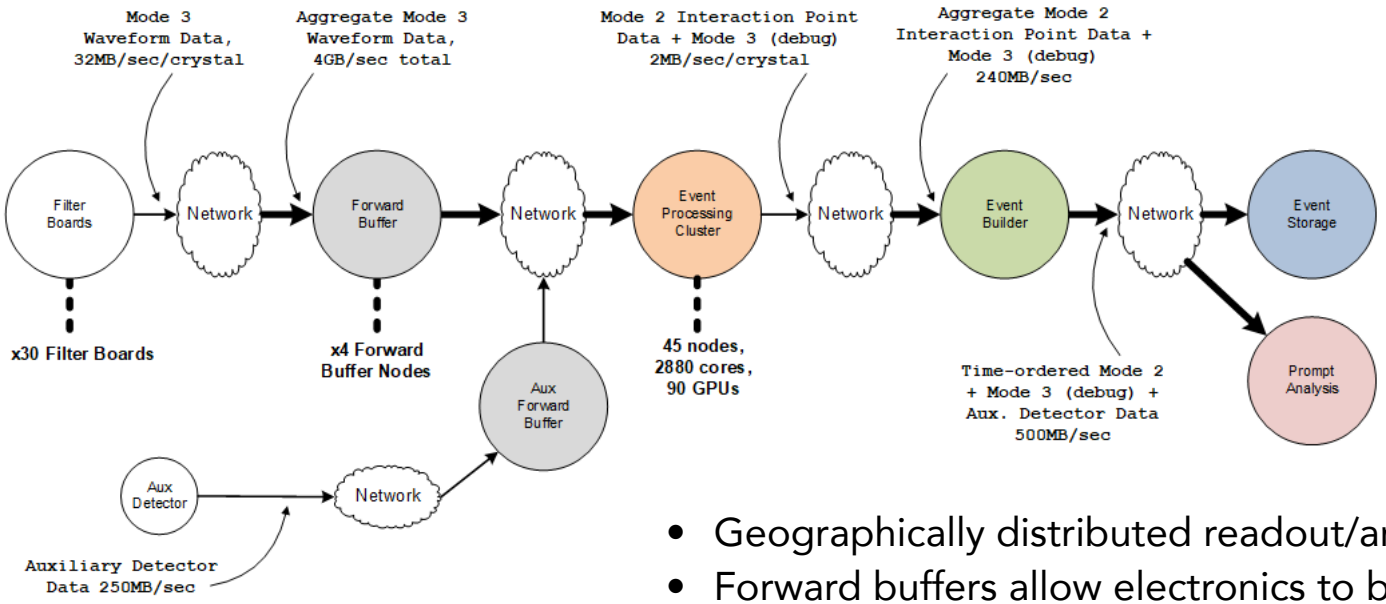

Message size	Maximum message rate (Single stream)	Maximum streams at 20k messages / sec
100 bytes	≥ 400k*	50
500 bytes	≥ 400k*	50
1 kB	270k	50
5 kB	210k	50
8 kB	172k	50
8.8 kB	170k	50
9.2 kB	110k	20
20 kB	80k	18
40 kB	50k	16
64 kB	28k	*

- Geographically distributed readout/analysis
- Forward buffers allow electronics to be free running and downstream computing resources to self-schedule
- Discrete, container-based pipeline components deployable locally, on HPC facilities or in the commercial cloud
- Goal - real time, large scale processing with 100's of ms of latency

Opportunities / Status

- Provides a mechanism for experiments - even small ones - to easily take advantage of distributed HPC and cloud resources
- Synergistic with ASCR Integrated Research Initiative (IRI)
- Architecture in keeping with the 'cloud-based' orientation of next-gen HPC facilities (deployable container sets, API-based)

- What it takes:
  - API standardization
  - Software development across multiple disciplines
  - Policy framework between institutions

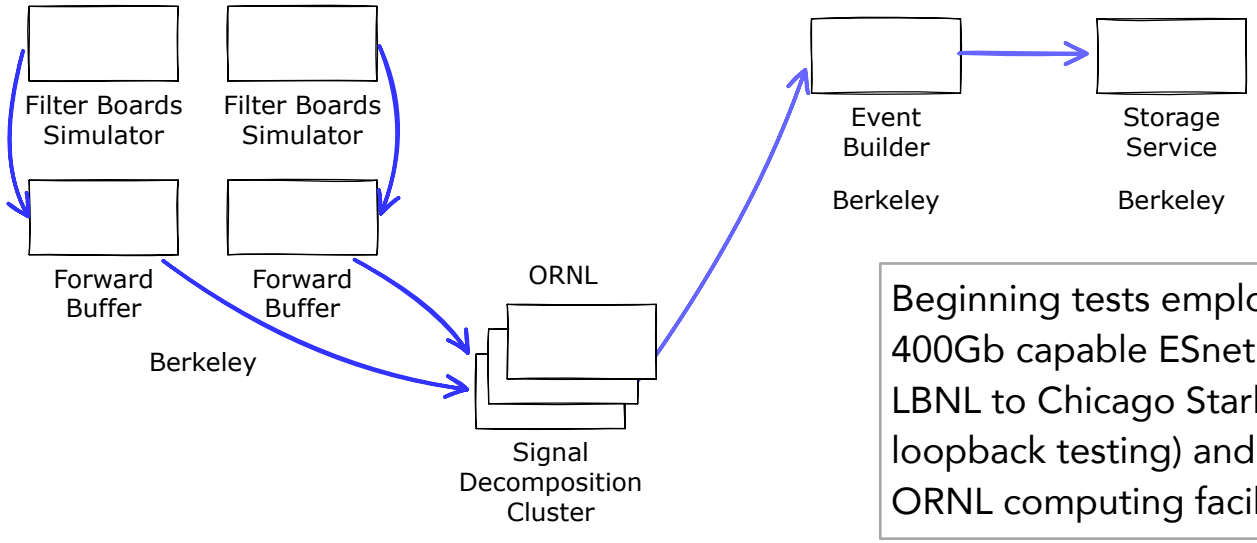

Beginning tests employing 400Gb capable ESnet testbed at LBNL to Chicago Starlight (initial loopback testing) and OLCF/ ORNL computing facilities.