SRO RTDP

Streaming Readout Real-Time Development and Testing Platform

D. Lawrence (PI), V. Gyurjyan, C, May, A. Roy, J. Tsai

Why SRO?

- Optimize experimental equipment to minimize the need for delay cables.
- Enable full integration of all detector components for comprehensive event identification.
- Implement real-time data processing to minimize costly IO latencies.
- Facilitate future event identification processes to support verifying and validating new physics theories.

Why Now?

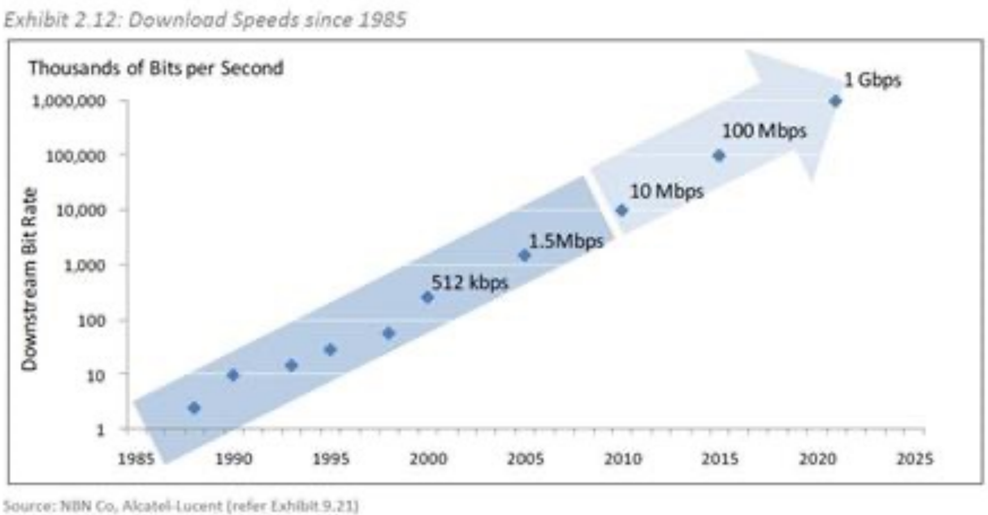

Network speeds have grown 10X every five years since 1990

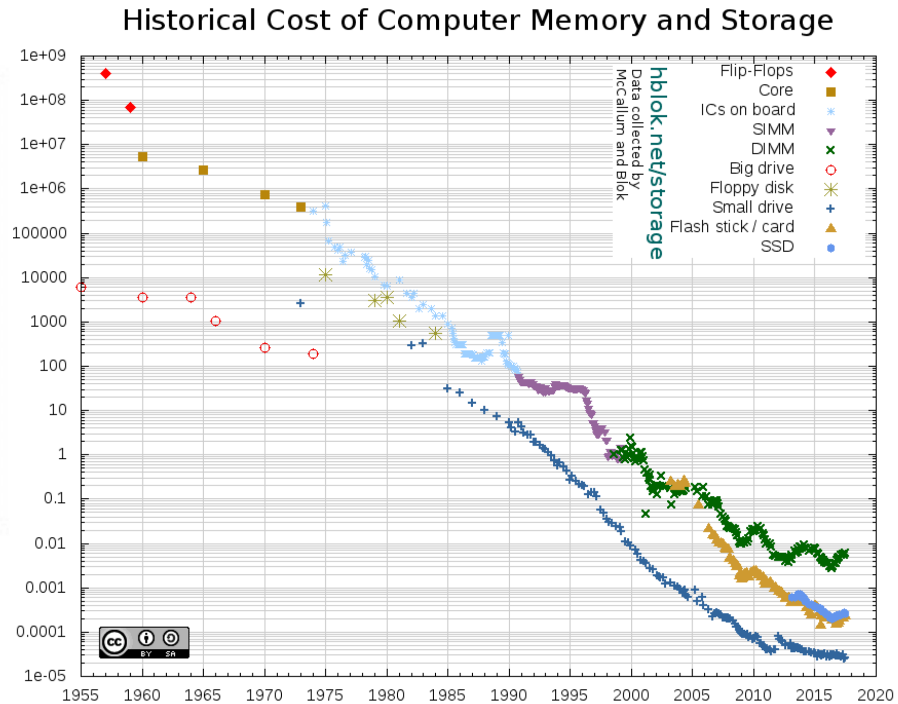

Why RTDP?

Grand Challenge in Readout and Analysis for Femtoscale Science

Grand Challenge in Readout and Analysis for Femtoscale Science

Amber Boehnlein, Rolf Ent, Rik Yoshida November, 2018

Introduction

Micro-electronics and computing technologies have made order-of-magnitude advances in the last decades. Combined with modern statistical methods, it is now possible to analyze scientific data to rapidly expose correlations of data patterns and compare with advanced theoretical models. While many existing nuclear physics and high-energy physics experiments are taking advantage of these developments by upgrading their existing triggered data acquisition to a streaming readout model, these experiments do not have the luxury of an integrated systems from DAQ through analysis. Hence, we aim to remove the separation of data readout and analysis altogether, taking advantage of modern electronics, computing and analysis techniques in order to build the next generation computing model that will be essential for probing femto scale science.

Integrated Whole-Experiment Model

An integrated whole-experiment approach to detector readout and analysis towards scientific output is summarized in the following figure.

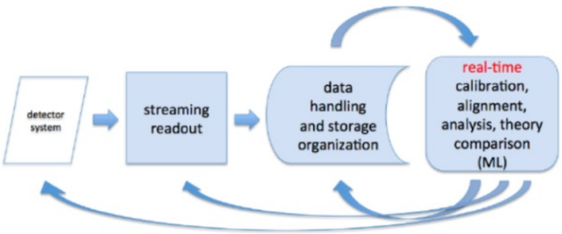

Key Elements

An integrated whole-experiment approach to detector readout and analysis towards scientific output will take advantage of multiple existing and emerging technologies. Amongst these are:

- "Streaming readout" where detectors are read out continuously.
- Continuous data quality control and calibration via integration of machine learning technologies.
- Task based high performance local computing.
- Distributed bulk data processing at supercomputer centers.
- Modern statistical methods that can detect differences among groups of data or associations among variables even under very small departures from normality.

Existing and Proposed Efforts

Several of the current LDRD proposals as well as separate on-going efforts naturally fit into the framework of the integrated whole-experiment model of data handling and analysis. They are

- Jefferson Lab EIC science related activities
  - Web-based Pion PDF server
- Jefferson Lab and related part of the Streaming Consortium proposal to the EIC Detector R&D committee including
  - Crate-less streaming prototype
  - TDIS streaming readout prototype
  - EM Calorimeter readout prototype
  - Computing workflow - distributed heterogeneous computing
- LDRD proposals
  - JANA development 2019-LDRD-8
  - Machine Learning MC 2019-LDRD-13
  - Streaming Readout 2019-LDRD-10

Grand Challenge

Develop a proof of concept of quasi-instantaneous high-level nuclear physics analysis based on modern statistics from a self-calibrated matrix of detector raw data synchronized to a reference time, without intermediate data storage requirements, with production systems developed for late stage 12 GeV analysis and the Electron Ion Collider. We propose organizing some of the LDRD proposals and other exploratory work around these themes to achieve proof of concept.

GC Related Projects

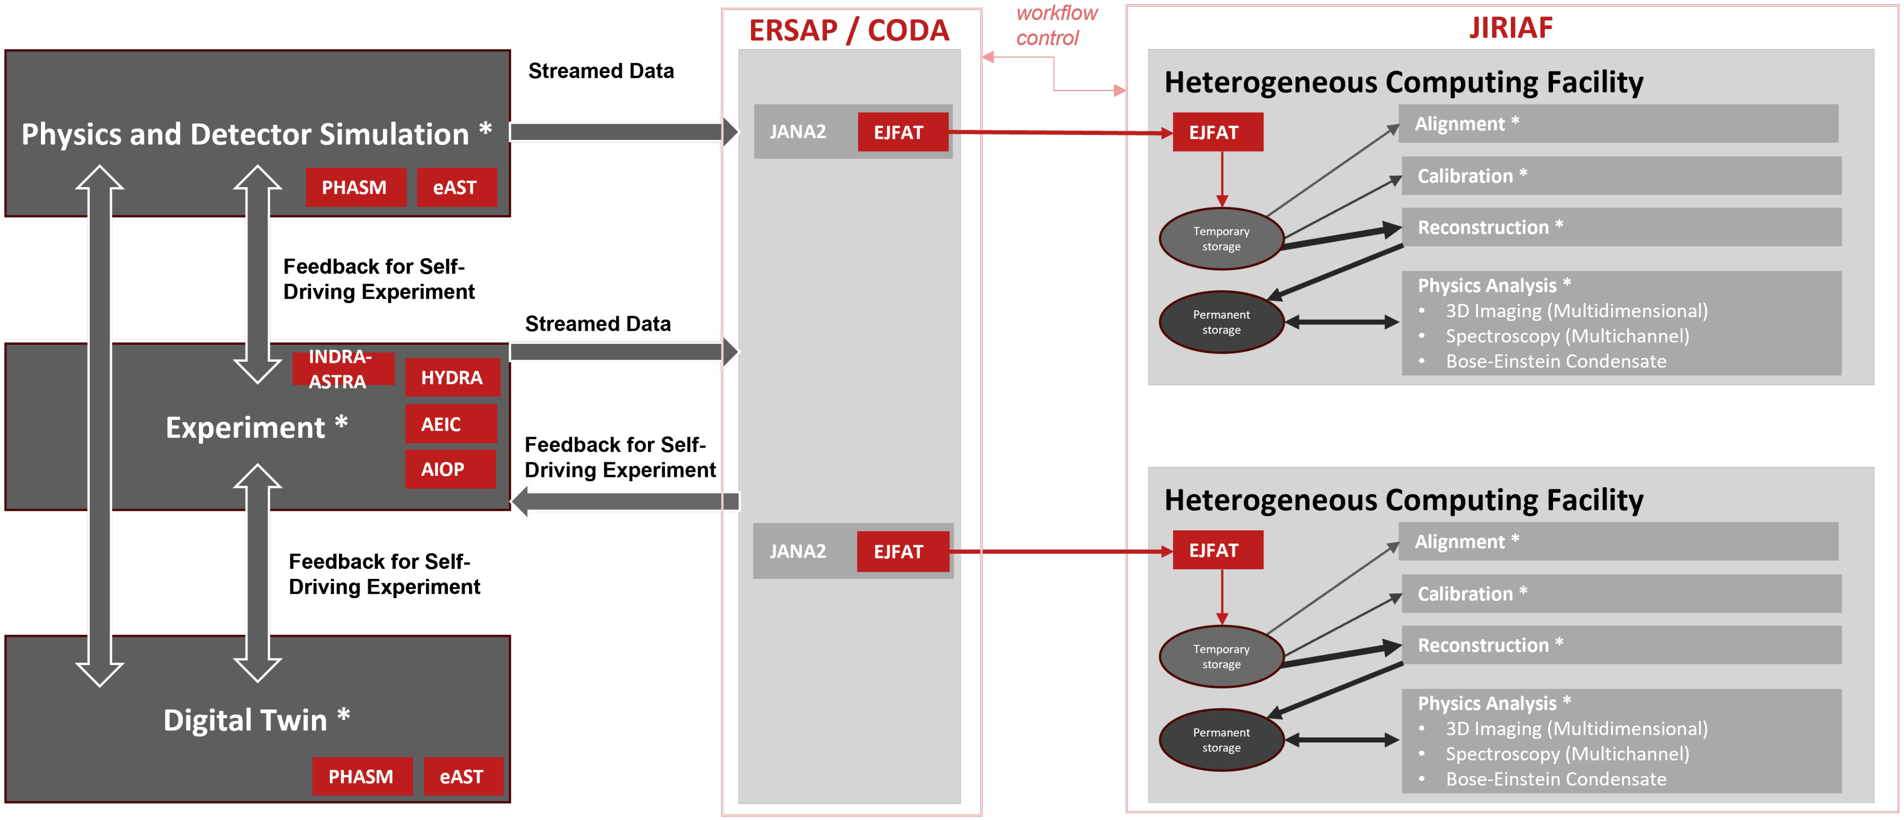

* AI/ML

GC Timetable

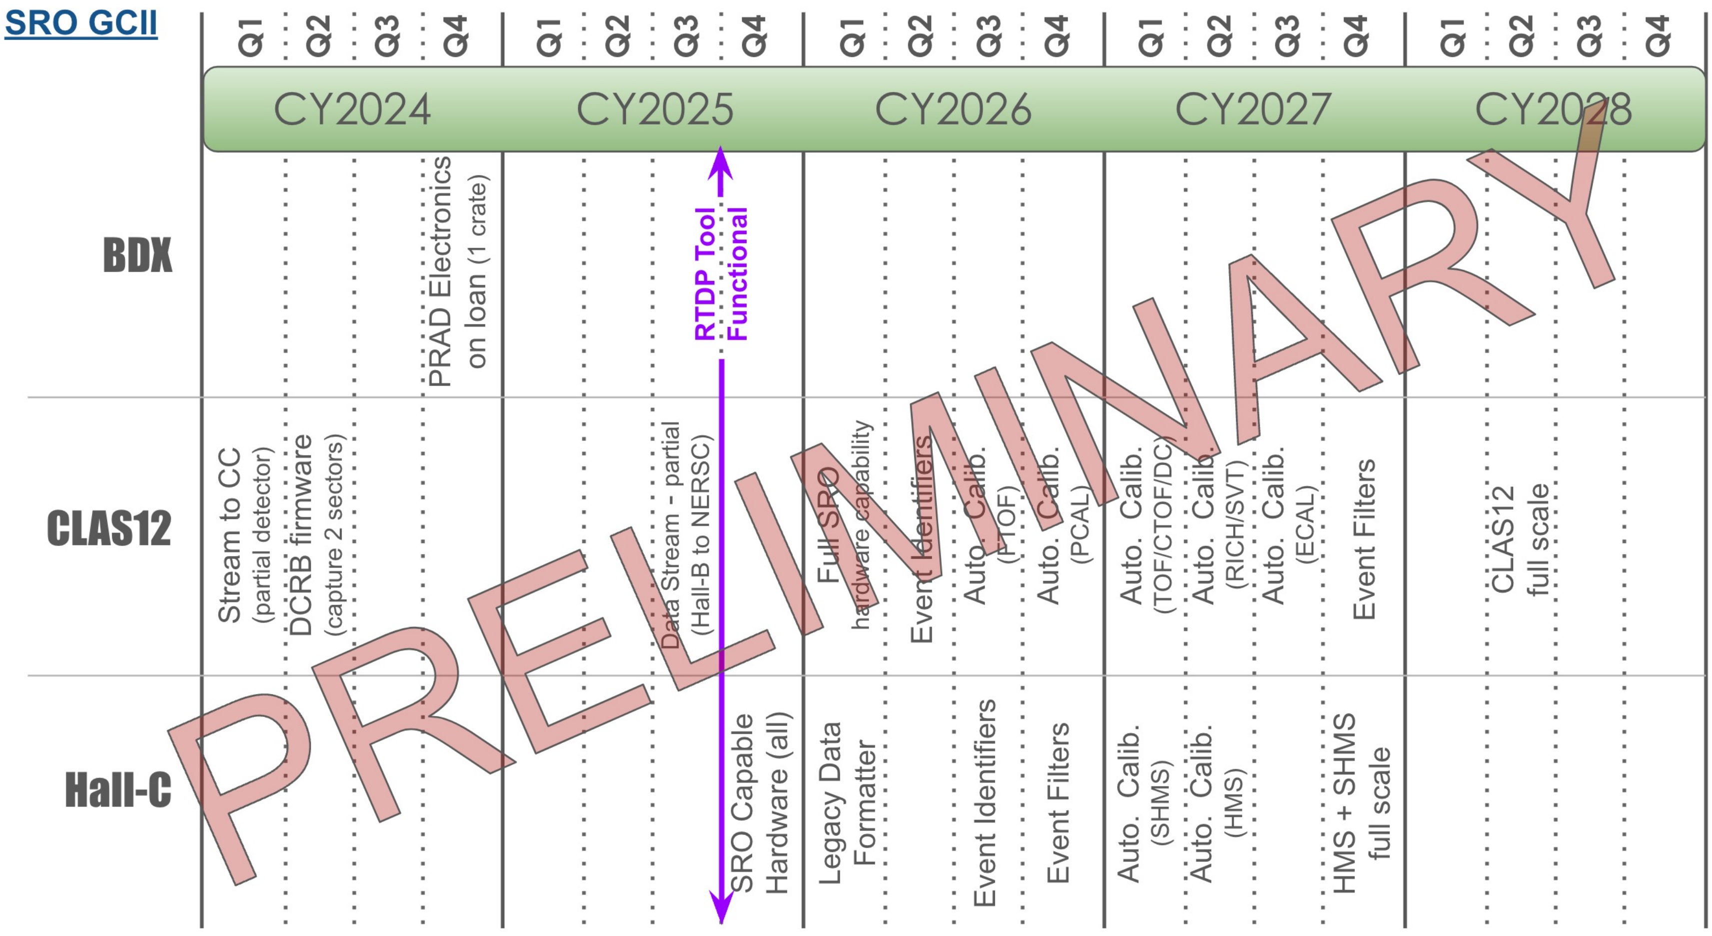

Milestones Achieved

ID	Task	status	Comments
M01	Create prototype ERSAP configurations for INDRA and CLAS12 test systems	■	A CLAS12 example and "Hello World" example have been placed in Github. INDRA has not been done yet.
M02	Identify or capture SRO formatted data from CLAS12 and INDRA test systems with data tag/filtering capability (output data ready for further offline processing)	✓	Data was captured at various beam currents from CLAS12 on Dec. 17. INDRA data capture done using pulser inputs to SAMPA setup.
M03	Evaluate existing solutions for configuring and launching remote distributed processes	✓	see evaluations in document on EPSCI wiki.
M04	Establish code repository(s), project site, and method of documentation	✓	This has been done here: https://github.com/JeffersonLab/SRO-RTDP
M05	Create stream splitter program for EVIO or HIPO data formatted files	✓	Created for GlueX. (See text for details on HIPO)
M06	Create stream splitter program for simulated data in PODIO for ePIC	✓	Prototype tested using FABRIC testbed. Simulated ePIC data sent from CERN to 8 different US sites.
M07	Create VTP emulator using files produced by stream splitter	■	Mostly done for raw data. Not started for simulated data.
M08	Create controller program to synchronize multiple VTP emulators	✓	Satisfied through alternate design using synchronized system clocks.
M09	Determine appropriate schema for all aspects of monitoring	■	Monitoring info. extracted as JSON records from both docker and /proc sources on Linux. Display in Grafana prototyped, but not yet complete.
M10	Establish databases for monitoring system using existing JLab servers.	×	This work has not begun
M11	Integrate Hydra as monitoring component.	■	Work done to containerize GlueX online monitoring in order to allow full test with Hydra. The Hydra is nearly complete containerizing Hydra (for off project purposes) which we will use.
M12	Integrate off-line data analysis framework into platform for CLAS12 data	-	planned for FY24Q4
M13	Integrate off-line data analysis framework into platform for ePIC or GlueX simulated data	-	planned for FY24Q4
M14	Integrate example JANA2 analysis into platform	-	planned for FY24Q4

Major Highlights:

- Containerize GlueX online
- Second CLAS12 Data Capture exercise
- DPPUI Development
- Dynamic I/O monitoring of processes from user space
- ePIC PODIO data sent via FABRIC Testbed from CERN to 8 different sites across US

Computing Infrastructure

EXPERIMENTAL NUCLEAR PHYSICS

TEST DEMONSTRATES FULL DATA-STREAMING CAPABILITY OF CLAS12 FORWARD DETECTOR

On Thursday, May 16, a successful beam-test of a streaming readout system comprised of 24 data streams covering one-third of the CLAS12 Forward Detector was conducted. The data were streamed from a real physics experiment that was running at various beam currents on the dual LD2-Pb target in Hall B as part of Run Group E. The data streamed directly from the front-end crates to the Data Center in CEBAF Center. It was captured to preserve the detailed time structure of the network traffic, so that it can be replayed in the same way during accelerator-down times, when the beam is unavailable. This is a significant advance beyond earlier tests that were limited in scope and not performed under realistic conditions. The test demonstrated that full streaming is possible for the CLAS12 Forward Detector. Streaming readout systems are advanced systems that are able to collect and process all data generated by the detector in real time, while traditional systems rely on highly specialized hardware to select events of interest and discard the rest of the data.

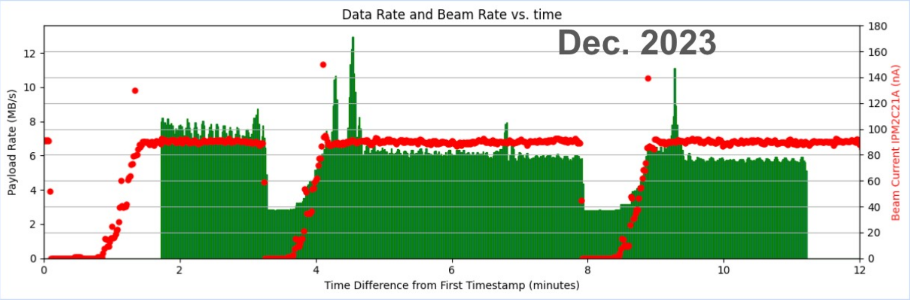

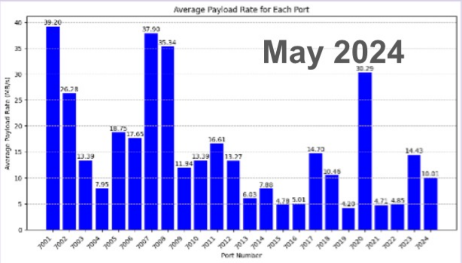

Current State

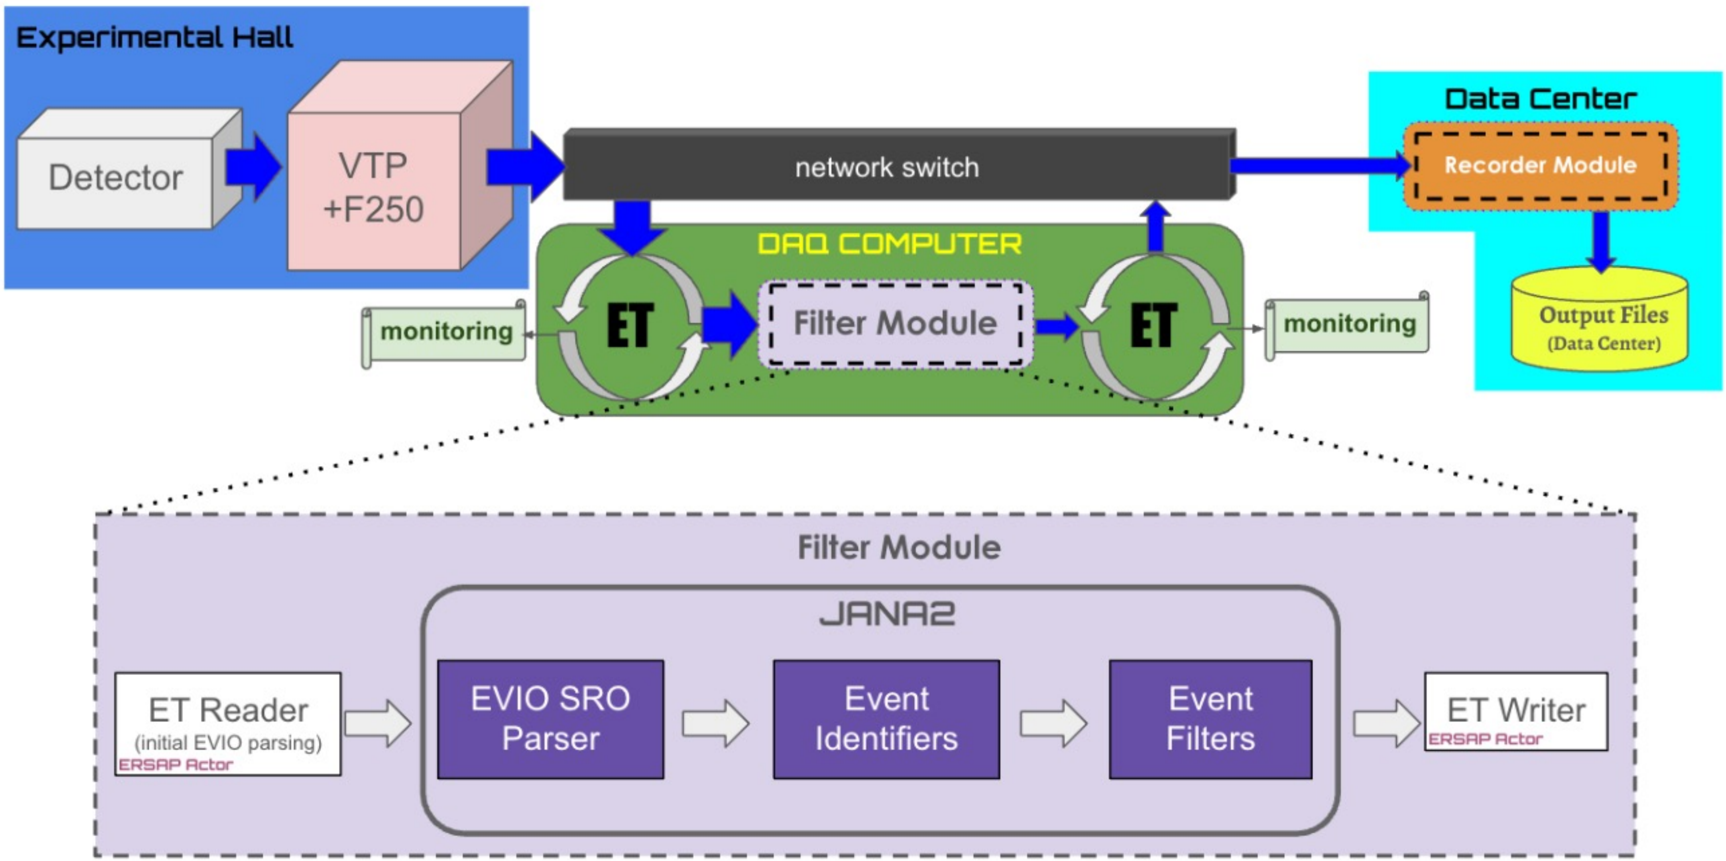

Final Deliverable

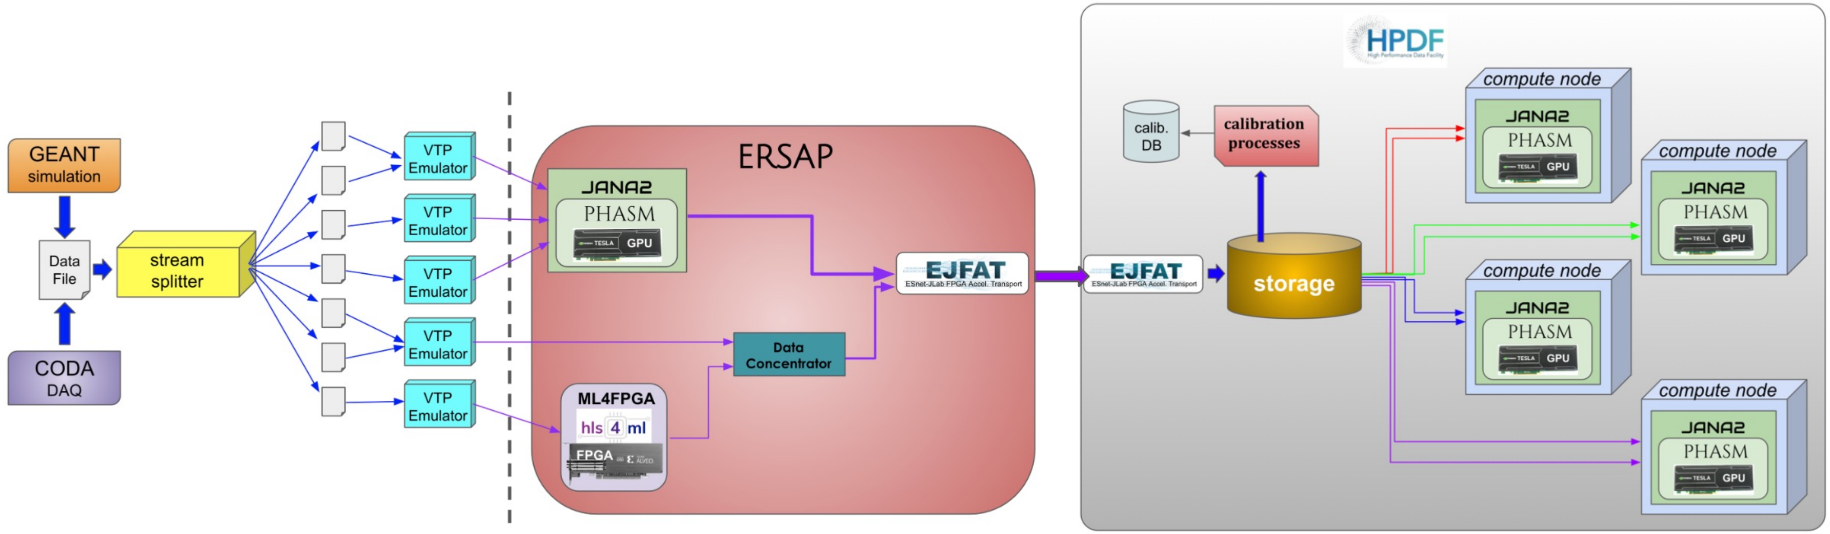

Highly configurable multi-stream source allows realistic streaming simulations

Onsite components will implement first stages of data filtering/reduction

Offsite processing must incorporate built-in calibration latencies and storage. This will also help inform HPDF design

RTDP FE

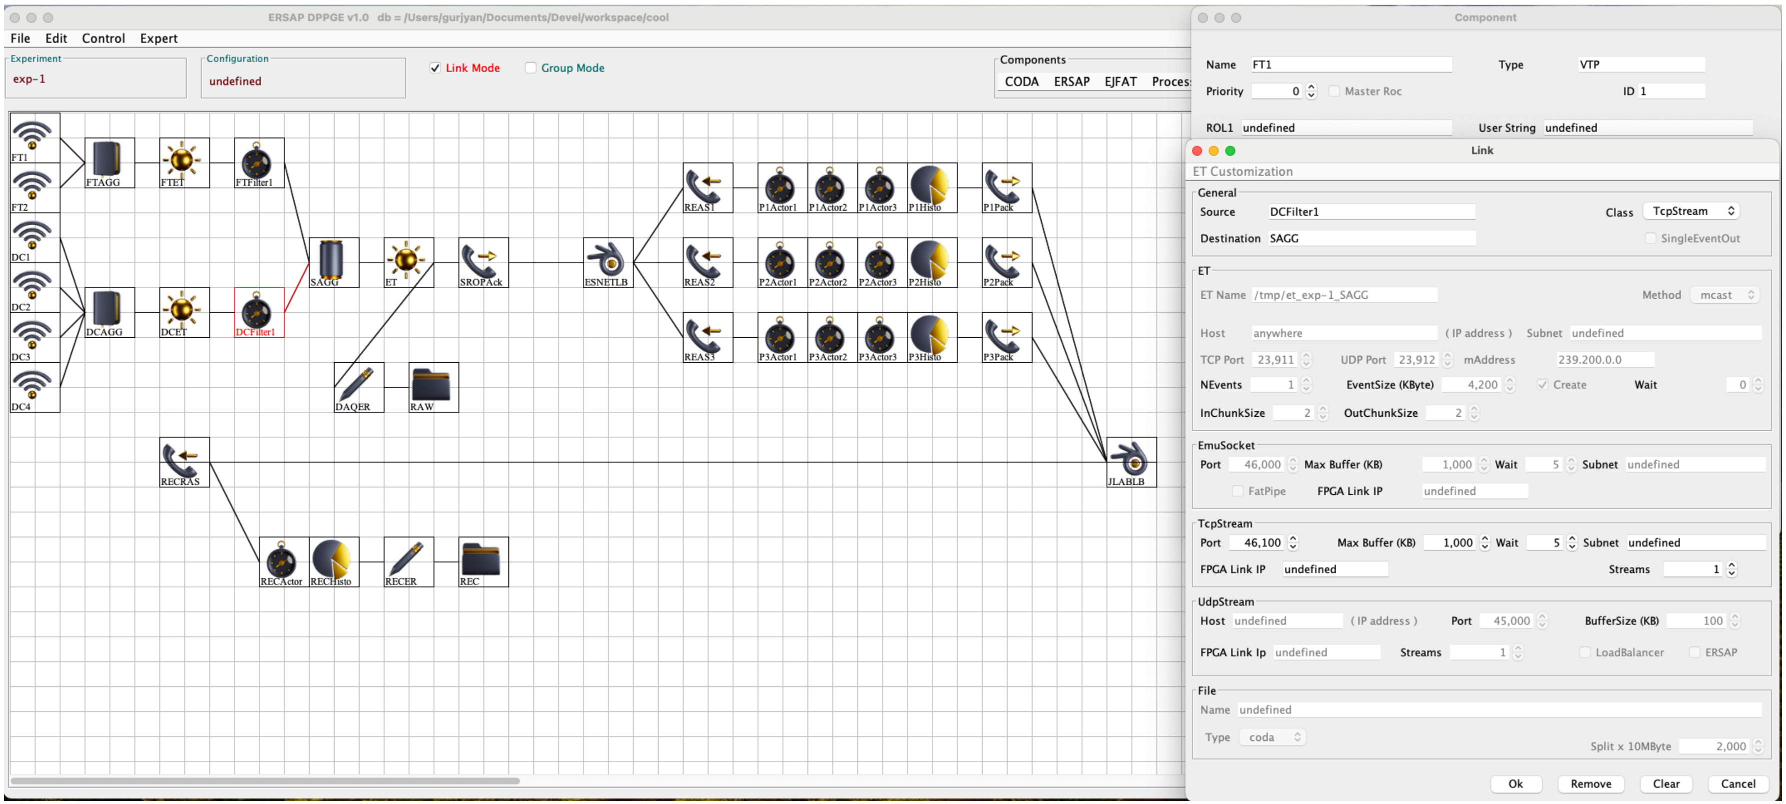

ERSAP DPPGE v1.0 db = /Users/gurjyan/Documents/Devel/workspace/cool
File Edit Control Expert
Experiment
exp-1
Configuration
undefined
Link Mode
Group Mode
Components
CODA ERSAP EJFAT Proces
FT1
FT2
DC1
DC2
DC3
DC4
FTAGG
FTET
FTFilter1
DCAGG
DCET
DCFilter1
SAGG
ET
SROPAck
ESNETLB
DAQER
RAW
REAS1
REAS2
REAS3
P1Actor1
P1Actor2
P1Actor3
P1Histo
P1Pack
P2Actor1
P2Actor2
P2Actor3
P2Histo
P2Pack
P3Actor1
P3Actor2
P3Actor3
P3Histo
P3Pack
JLABLB
RECRAS
RECActor
RECHisto
RECER
REC
Component
Name FT1
Type VTP
Priority 0
Master Roc
ID 1
ROL1 undefined
User String undefined
Link
ET Customization
General
Source DCFilter1
Class TcpStream
Destination SAGG
SingleEventOut
ET
ET Name /tmp/et_exp-1_SAGG
Method mcast
Host anywhere
(IP address)
Subnet undefined
TCP Port 23,911
UDP Port 23,912
mAddress 239.200.0.0
NEvents 1
EventSize (KByte) 4,200
Create
Wait 0
InChunkSize 2
OutChunkSize 2
EmuSocket
Port 46,000
Max Buffer (KB) 1,000
Wait 5
Subnet undefined
FatPipe
FPGA Link IP undefined
TcpStream
Port 46,100
Max Buffer (KB) 1,000
Wait 5
Subnet undefined
FPGA Link IP undefined
Streams 1
UdpStream
Host undefined
(IP address)
Port 45,000
BufferSize (KB) 100
FPGA Link Ip undefined
Streams 1
LoadBalancer
ERSAP
File
Name undefined
Type coda
Split x 10MByte 2,000
Ok
Remove
Clear
Cancel

Budget and Summary

- SRO experiments will have a **diversity** and **complexity** of components **beyond** that of **traditional** DAQ systems.

Budget	FY24	**FY25**	Total
($K)	$271K	**$285k**	$556k*

**increase of $11k from original*

- A Real-Time Development and Testing Platform for Streaming Readout DAQ is a **tool** that will help **develop** the **technical aspects** of **future** JLAB **experiments**

David Lawrence (PI)	SSCS	25%
Vardan Gyurjyan	SSIII	15%
Cissie Mei	SCSII	20%
Jeng-Yuan Tsai	Postdoc	50%
Michael Goodrich	SSIII	20%

- This platform will **benefit** the **HPDF** in being flexible enough to provide streaming profiles for even non-NP experiments

Thank You

Reviewer Questions

1. Are additional in-beam SRO tests with CLAS12 or GlueX anticipated during FY25 given the issues with the Hall B test carried out to date?

   There are no specific plans at this time. Much of the work for RTDP itself can be done using the existing captured data. The late schedule start for FY25 beam (end of Jan.) will put additional pressure on the schedule so we will want to show show a significant benefit will be gained from the exercise being done now as opposed to the post-LDRD era (e.g. Fall 2025) when RTDP as a tool can be used.

2. What are your plans for reconstructing the collected SRO data to demonstrate integrity and offline SRO data “assembly”?

   The most viable path is likely to work with Hall-B to couple this into the AI/ML based fast recon L3 trigger system they are already developing. This will require less time from the RTDP developers who can focus more on stream integrity monitoring.

3. How will SRO tests be integrated with level-3 testing to develop a complete SRO system?

   From the RTDP standpoint, L3 is treated as a filter for L1 data so can use existing L1 data sets. We will focus first on GlueX data and then CLAS12 to ensure generality. The plan is to make use of Hall compute resources during beam down periods for some testing of longer data flows.