RTDP: Streaming Readout Real-Time Development and Testing Platform

Authors: Ayan Roy, **David Lawrence**, Jeng-Yuan Tsai, Marco Battaglieri, Markus Diefenthaler, Vardan Gyurjyan, Xinxin (Cissie) Mei, Nathan Brei, Sergey Boiarinov

MOTIVATION

Streaming Readout (SRO) Data Acquisition is a new paradigm in which data read from a detector is streamed directly into a large, High Throughput Computing (HTC) farm. This enables things like triggerless DAQ, sophisticated event filtering, and fast calibrations which lead to shorter delays in producing reconstructed data that is ready for physics analysis.

The RTDP project aims to build a software platform that can be used to simulate the many components of complex SRO systems to aid in the prototyping and design of such systems. The ability to swap real components such as GPUs and FPGAs in for their simulated counterparts will add a level of hardware testing that improves the realism of such simulations.

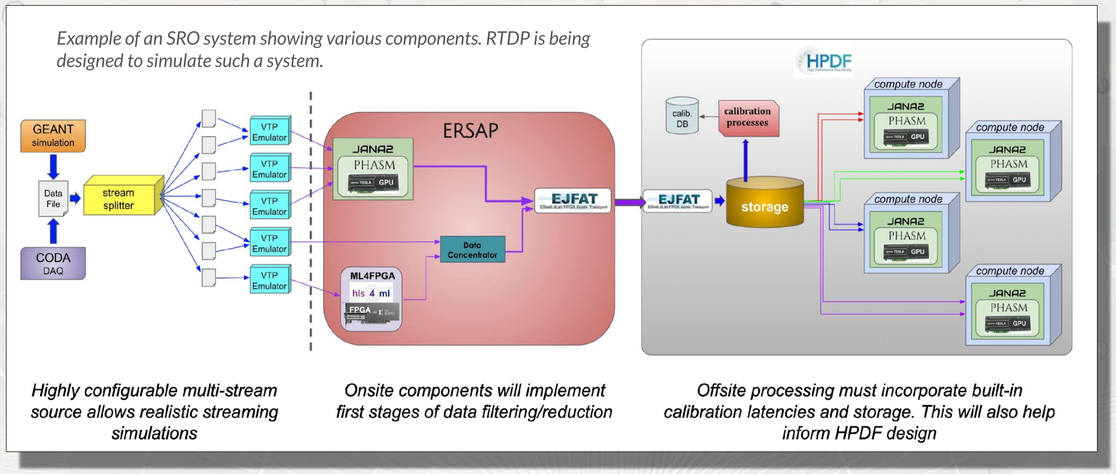

Highly configurable multi-stream source allows realistic streaming simulations

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 OBJECTIVES

- **Process Launcher:** The software that will be used to configure and launch each component
- **Monitoring System:** Advanced monitoring for continuous data validation (this will include Hydra, among other tools)
- **Proxy Components:** Allow effects-based simulations to be developed using a performant language such as C++
- **Multi-stream Event Source:** Read data from file(s) and stream into system as though coming from a DAQ system.
- **High Bandwidth Test:** Stream data from an experimental hall that saturates the available bandwidth to the Data Center in CEBAF Center and process on heterogeneous hardware using the RTDP system.
- **Insights for HPDF:** The platform will be a tool for developing, testing, and validating SRO systems that utilize remote compute facilities similar to what HPDF will do..

STREAMING ePIC DATA FROM CERN TO US

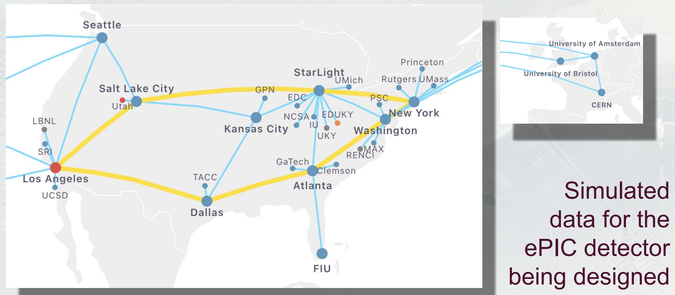

Name	ElCreconTCPmulti							
Lease Expiration (UTC)	2024-04-15 11:54:13 +0000							
Lease Start (UTC)	2024-04-14 11:54:14 +0000							
Project ID	a7818636-1fa1-4e77-bb03-d171598b0862							
State	StableOK							
ID	**Name**	**Cores**	**RAM**	**Disk**	**Image**	**Image Type**	**Host**	**Site**
Nodes								
19bdb4d0-619a-4abe-8d98-6ad35a986b43	server_CERN	4	16	100	docker_rocky_8	qcow2	cern-w1.fabric-testbed.net	CERN
abc1026b-fc73-4ed7-9759-e37c29ded40b	worker_ATLA	4	16	100	docker_rocky_8	qcow2	atla-w1.fabric-testbed.net	ATLA
eeb8ae21-e7c7-4cc1-8940-6cbe0fe9a3ba	worker_DALL	4	16	100	docker_rocky_8	qcow2	dall-w1.fabric-testbed.net	DALL
ac0936ca-b44e-4388-aa58-0a4c775234b9	worker_KANS	4	16	100	docker_rocky_8	qcow2	kans-w1.fabric-testbed.net	KANS
710cdd4d-1c2a-4c6d-9126-3aadf1ecc076	worker_LOSA	4	16	100	docker_rocky_8	qcow2	losa-w2.fabric-testbed.net	LOSA
1b3e6ea2-1f39-467e-9bf9-cd756356a0f7	worker_NEWY	4	16	100	docker_rocky_8	qcow2	newy-w1.fabric-testbed.net	NEWY
709bf2f5-e684-4671-816e-042b50c94751	worker_SALT	4	16	100	docker_rocky_8	qcow2	salt-w2.fabric-testbed.net	SALT
ff915fef-a6ee-466e-9917-8e53a158b6e8	worker_SEAT	4	16	100	docker_rocky_8	qcow2	seat-w2.fabric-testbed.net	SEAT
99cf0a22-0da4-4c9b-a6b3-dc0ede261950	worker_WASH	4	16	100	docker_rocky_8	qcow2	wash-w1.fabric-testbed.net	WASH

Simulated data for the ePIC detector being designed for the EIC was modified into a stream form that could be distributed to multiple sites. The graphics on the left show an exercise where the FABRIC testbed was used to send events from CERN to 8 different sites across the US. RTDP is a general tool that will support designs using the ePIC/SoLID software stack as well as ones using software stacks from CLAS12 or GlueX.

STREAMING DATA FROM CLAS12 TO DATA CENTER

In Dec. 2023 and then again in May 2024 beam data was taken with the CLAS12 detector and streamed directly to CEBAF Center. The TCP data packets were captured into industry standard .pcap files using a high speed NIC with hardware timestamps. This allowed the capture of data at various beam intensities that could be played back with exactly the same network time structure and TCP packets as it had when the beam was on. In the top plot on the left, the payload rate vs. time in green plotted against the left hand axis. The red indicates the beam current for this same time period plotted against the right hand axis.

The bottom plot on the left (blue) shows the payload rates for individual streams. The CLAS12 drift chambers are the bins on the left half of the plot while the right half corresponds to F250 ADC modules used for calorimeters.

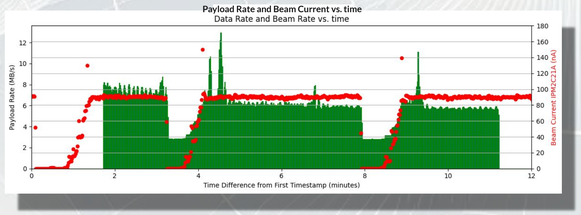

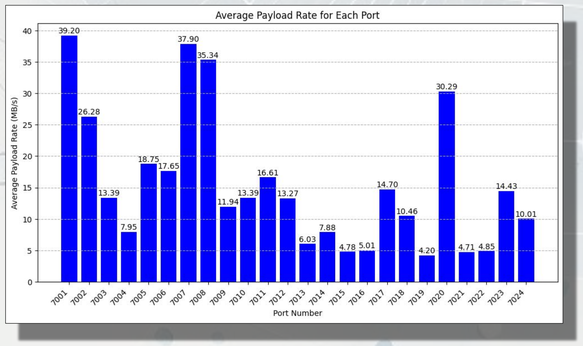

GRAPHICAL CONFIGURATION + ONTOLOGY

Configuring a RTDP workflow can become quite complex due to the number of components. Extending the CODA graphical interface to support RTDP will allow users to more easily build and visualize very complex configurations in a format already familiar to them.

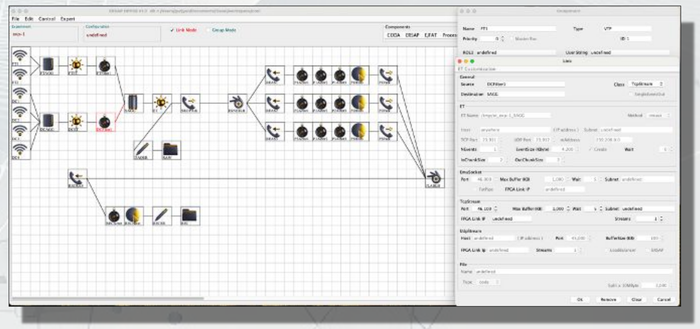

ACKNOWLEDGEMENT

This project is funded through the Thomas Jefferson National Accelerator Facitlity LDRD program. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under contract DE-AC05-06OR23177.