

Multi-site data streaming orchestration with SciStream

Joaquin Chung, Jamil Hasibul, Rohan Vardekar, Flavio Castro, Se-young Yu

Abstract

Modern big data science requires coordinating multi-site scientific instruments to generate and process data at high rates. Such a study involves the orchestration of online processing for data reduction, feature detection, and experiment steering over memory-to-memory streaming from multiple source instruments to remote high-performance computers (HPC). Yet, a systematic approach to orchestrating online processing over multiple data sources and processors is lacking. SciStream, our proposed solution, is an embedded middlebox-based architecture with control protocols to enable efficient and secure memory-to-memory data streaming between producers and consumers without direct network connectivity. SciStream operates at the transport layer to be application agnostic, supporting well-known protocols such as TCP, UDP, and QUIC. This demonstration will emulate multi-site online data processing using ESnet testbed, StarLight, FABRIC and SC show floor. The mem-to-mem data streaming over the Wide-Area Network (WAN) between ESnet Testbed, StarLight, FABRIC and SC showfloor will be enabled by SciStream.

Goals

1. Demonstrate on-demand memory-to-memory streaming of scientific data utilizing SCinet facilities, ESnet Testbed, FABRIC, Argonne, and StarLight Testbed.
2. Show SciStream's various transport protocol support, including TCP, UDP, and QUIC.
3. Evaluate SciStream on a real WAN with an emulator of the scientific data streaming and multiple WAN connections.

Resources

We will use compute resources hosted at the StarLight booth on the SC show floor for our demonstration. First, we require a 400 Gbps circuit from the StarLight booth in Denver to the StarLight IXP in Chicago. Then, we will extend another circuit from StarLight to ESnet Testbed, FABRIC, and Argonne. In addition, at least 4 compute nodes, with at least 32 CPU cores, memory of 128 GB, and storage of 400GB, are needed.

No need for further physical network devices. ESnet, StarLight, Argonne, and FABRIC can provide all equipment. No software tools are needed. Need of a monitor.

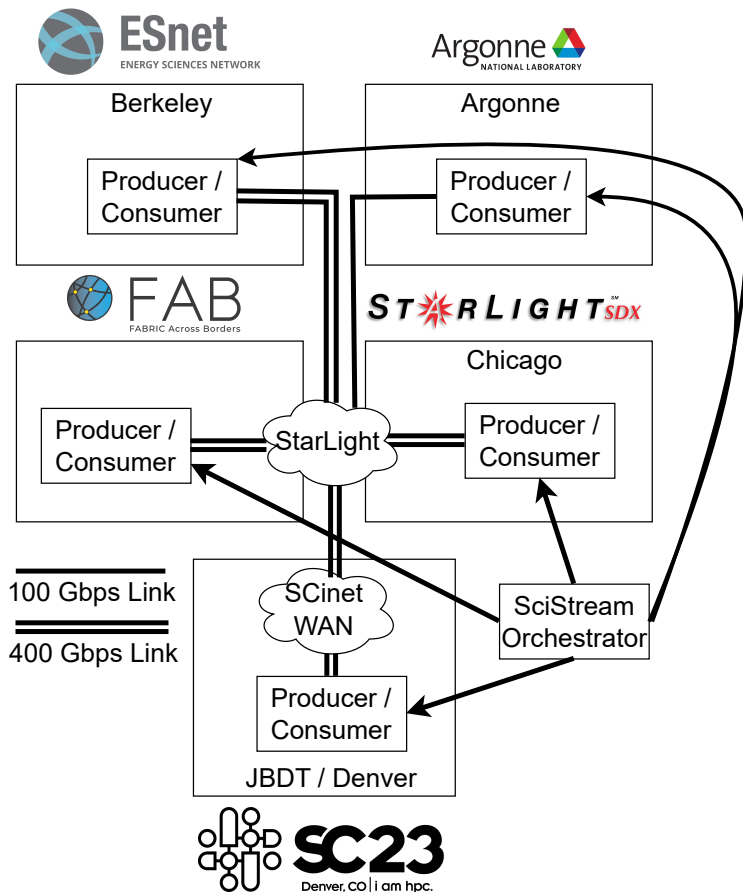


Figure 1. Testbed design

Involved Parties

[List of other institutions, researchers and entities involved in the planning and execution of this demonstration. This should include names and contact information]

- Joaquin Chung, Argonne National Laboratory, chungmiranda@anl.gov
- Jamil Hasibul, Argonne National Laboratory, mjamil@anl.gov
- Rohan Vardekar, Argonne National Laboratory, rvardekar@anl.gov
- Flavio Castro, Argonne National Laboratory, fcastro@anl.gov
- Se-young Yu, ESnet, youf3@es.net
- Jim Chen, iCAIR, jim-chen@northwestern.edu
- Joe Mambretti, iCAIR, j-mambretti@northwestern.edu
- Paul Ruth, RENCi, pruth@renci.org