



Contribution ID: 108

Type: **Presentation**

From instrument to publication: A First Attempt at an Integrated Cloud for X-ray Facilities

Tuesday 28 January 2020 10:20 (20 minutes)

Large scientific X-ray instruments such as MAX IV [1] or XFEL [2] are massive producers of annual data collections from experiments such as imaging sample materials. MAX IV for instance has 16 fully funded beamlines, where 6 of which can produce up to 40 Gbps of experimental data during a typical 5 to 8 hour time-slot, resulting in up to 90 to 144 TBs for a particular beamline experiment.

Scenarios like this calls for solutions that can manage petabytes of datasets in an efficient manner, while enabling scientists with a path of least resistance to define on the fly and subsequent batch processing that often seeks to find needle answers in the data haystack. General outlier detection, pattern recognition and basic statistics just as bin counting are some of the typically tasks conducted during the post analysis phase. To enable scientist with such capabilities, the current challenges calls for an integrated solution that is both able to scale horizontal in terms of available storage, but also be able to make on the fly informed decisions that could potentially either reduce the experimental data stream before it is persistently stored, or enable feedback mechanisms to the instrument itself about which data is of interest to the scientists and that which has no or little value.

The continuous collaboration between the eScience group [3] at the Niels Bohr Institute and the MAX IV facility through their Data Storage and Management Project (DataSTaMP) [4] and European Open Science Cloud (EOSC) [5] participation aims to provide just such an integrated cloud solution to elevate the combined data services available to researchers in general.

The architecture design to enable this is made of two distinct services. The High Throughput Storage System (HISS) and the Electronic Research Data Archive (ERDA) [6]. HISS is a developing distributed system that is designed as a high speed I/O gateway of storage nodes for stream oriented data collections. The system does this through temporary buffering during storage and retrieval of high bandwidth streams, acting in a sense as a front proxy to a subsequent persistent storage location such as a PFS or tape archive system. In addition to being a mere set of buffer nodes that allows for temporal storage reservations, the system is also being designed to allow for an on the fly scheduling of operations to be conducted during the I/O of datasets by scheduling preprocessing tasks on an FPGA accelerator. This enables for both in situ decisions about particular data points mid stream or general data reduction/prefiltering as specified by a user defined kernel, that may also introduce feedback streams to the data provider itself. A provider in this instance could be a beamline instrument at MAX IV.

The system enables access through a REST API that is inspired by and aims to be compatible with the AWS S3 [7] service and commandline tools. To define the computational kernels that are targeted for accelerated computation the proposal is to transpile python kernels into VHDL through an eScience developed toolchain that consists of Bohrium [8], SMEIL [9] and SME [10].

The target of the HISS offloading in our instance, is the ERDA system that is subsequently responsible for retaining the incoming collections, that is either stored GPFS or tape archives. The archive then on top of this provides a rich set of features for both managing and post-processing of data upon being stored. This includes Dropbox like sharing and synchronization in addition to efficient data access to home and collaborative datasets through standard secure protocols like WebDAV over SSL/TLS (WebDAVS), FTPS and SFTP. For processing the service enables processing of existing datasets through a JupyterHub [11] environment with container based JupyterLab [12] sessions for interactive executions of personal or collaborative resources.

It is the aim of this integrated cloud solution to enable both the receipt of instrumentation data streams directly from the source, while allowing user defined decision making to take place before the data is persistently stored. For instance, the user could specify a reduction or statistics kernel that would alleviate the need to schedule such processing upon finishing the experimental phase. Enabling them to immediately interpret the results generated from the computed metadata.

[1] <https://www.maxiv.lu.se>

[2] <https://www.xfel.eu>

[3] <https://www.nbi.ku.dk/Forskning/escience/>

[4] <https://www.maxiv.lu.se/accelerators-beamlines/technology/kits-projects/datastamp/>

[5] <https://ec.europa.eu/research/openscience/index.cfm?pg=open-science-cloud>

[6] <http://www.erda.dk>

[7] <https://docs.aws.amazon.com/cli/latest/userguide/cli-services-s3.html>

[8] <https://bohrium.readthedocs.io>

[9] <https://github.com/truls/libsme>

[10] <https://github.com/kenkendk/sme>

[11] <https://jupyterhub.readthedocs.io/en/stable/>

[12] <https://jupyterlab.readthedocs.io/en/stable/>

Primary authors: Mr MUNK, Rasmus (Niels Bohr Institute); Prof. VINTER, Brian (Niels Bohr Institute); Mr MATEJ, Zdenek (MAX IV Laboratory); Mr BARCZYK, Artur (MAX IV Laboratory)

Presenters: Mr MUNK, Rasmus (Niels Bohr Institute); Prof. VINTER, Brian (Niels Bohr Institute)

Session Classification: Fabric and platforms for Global Science

Track Classification: User Voice: Novel Applications, Data Science Environments & Open Data