



EUROPEAN
SPALLATION
SOURCE

Data Management Group

13.4.6

Tobias Richter

2016-07-19

Data Management Group



Core Team:

Tobias Richter - Group Leader
in post, staff cost committed
DMSC core person #1
frozen, uncommitted
DMSC core person #2
frozen, uncommitted



In Kind Agreements:

ISIS – Running
PSI – Agreed & signed & on hold

In Kind ratio
signed 57%
potential 73%

Delivery strategy:

- Small local core team
 - Coordinate development, integrate and test
 - Retain knowledge
 - Smooth transition to operations
- In Kind for main development
- Evolve from existing projects where possible
- BrightnESS to cover orphaned scope

What do we do? -- Overview



① Neutron Event Formation

Bridge the gap between detectors and storage disk or Mantid.
Neutron data is important.

② Data Streaming, Aggregation and Recording

Record entire state of the instrument and sample environment so the neutron data makes sense.

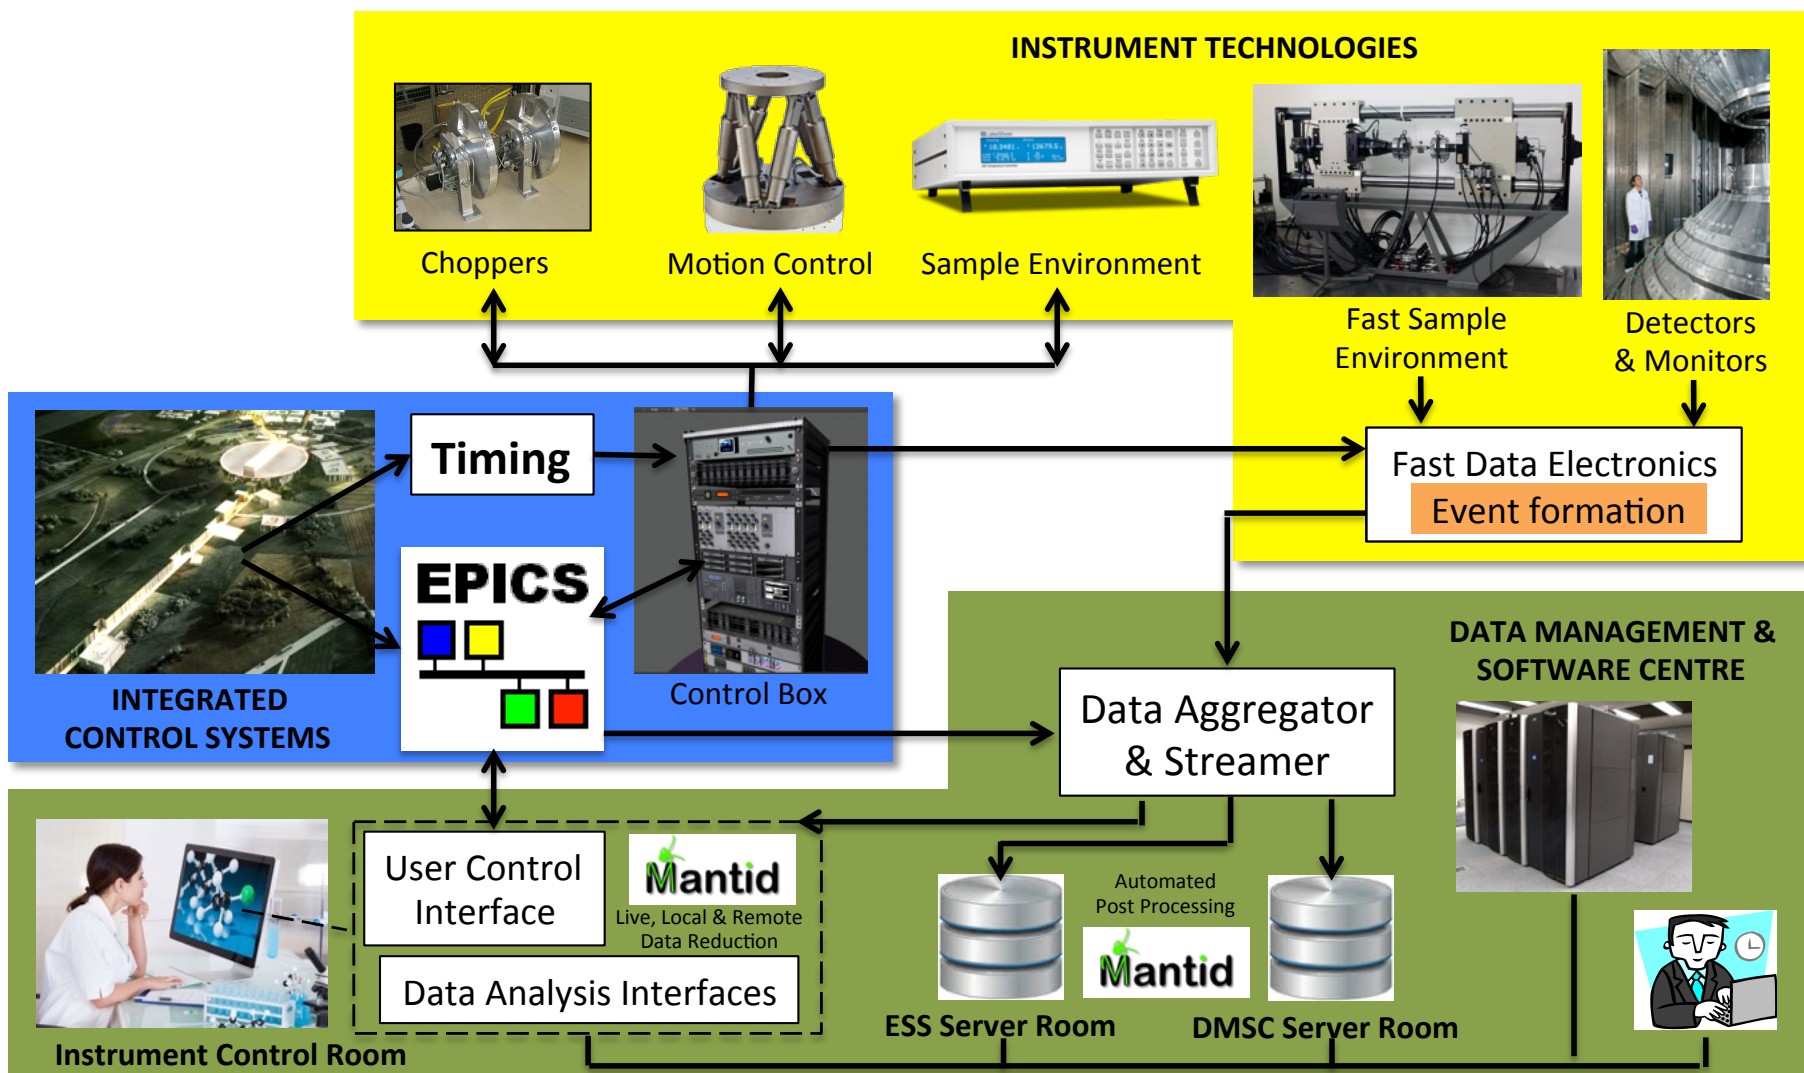
③ Data Curation and Processing

Keep a record of what sample was used, under what condition, by whom, when, so you can find the raw, reduced and potentially analysed data for future use in publications or to re-evaluate and re-analyse.

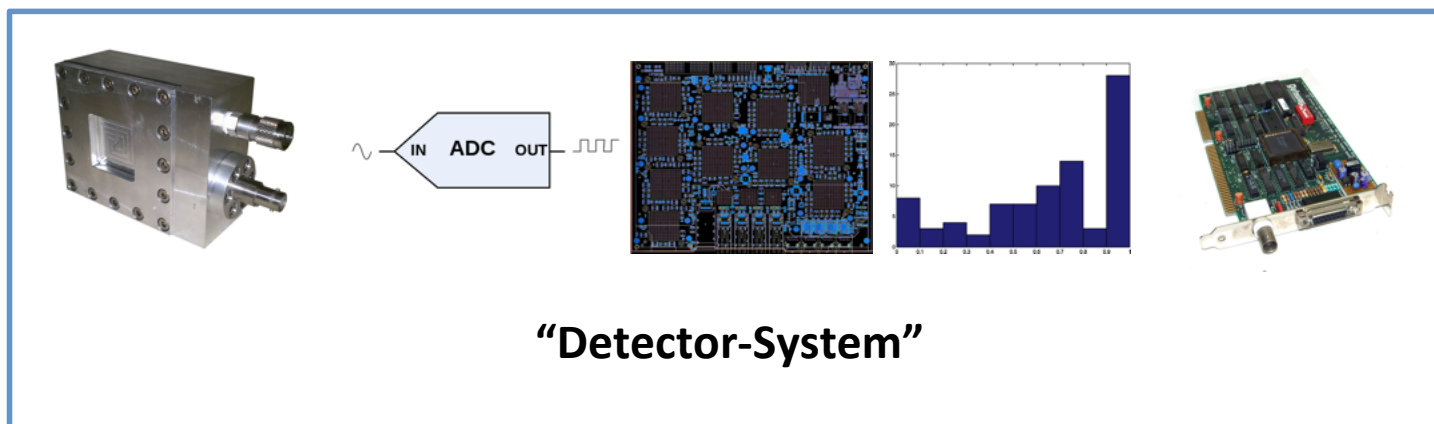

④ Testing and Integration

Integrate and test the whole lot within itself and other systems for problems with performance, integration of different devices and setups.

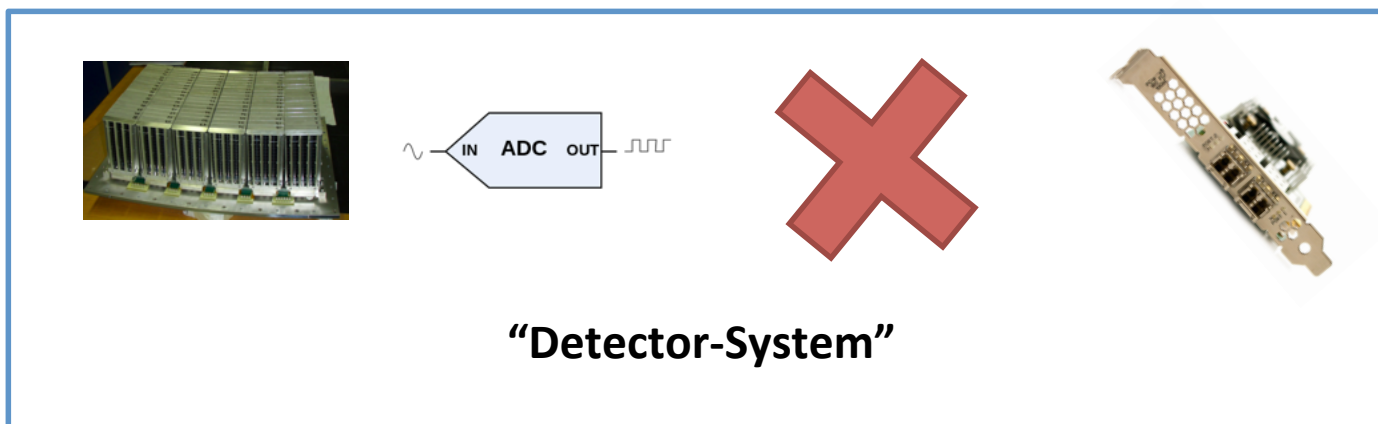
DMSC Domain and Interfaces



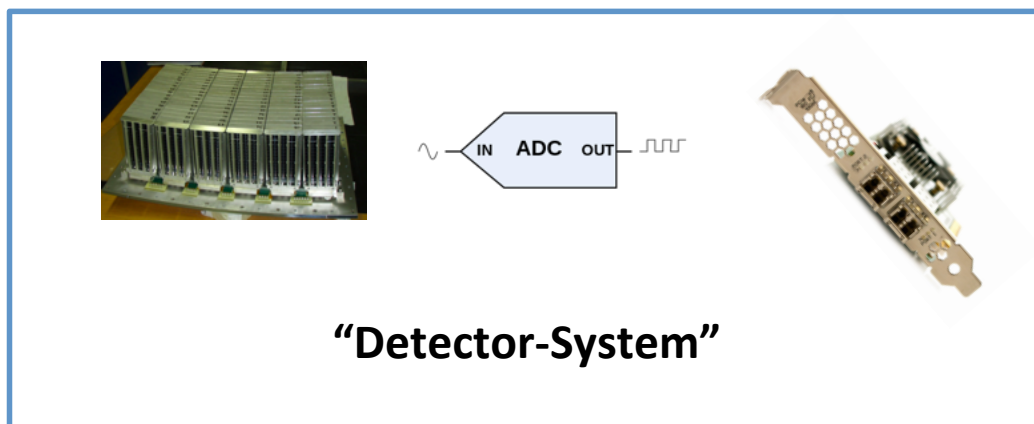

Traditional Counting Chain



TDR Counting Chain



BrightnESS Counting Chain



“Detector-System”

The image shows a stack of detector modules, a schematic of an ADC (IN, ADC, OUT) with a sine wave input and a square wave output, and a printed circuit board with various connectors.



BrightnESS WP5 – Live Data Management



Funded by EU Horizon 2020 INFRADEV-3

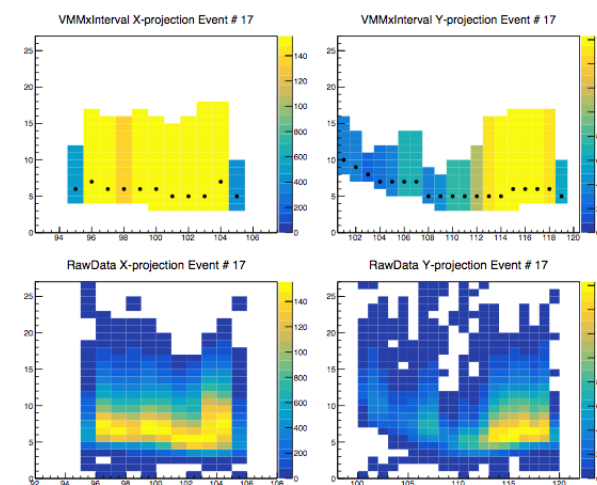
- Projects to help ESFRI's with start-up
- Projects are ~50/50 business & technical

BrightnESS Work Package 5:

- Captures previously orphaned scope of event formation in detectors & capture of fast sample environment data
- Partners are ESS/DMSC, PSI, U. Copenhagen, Elettra
- 5.1: Software event formation for new B and Gd detectors
- 5.2: Fast sample environment capture
- 5.3: Integration with data streaming in 13.4.6

BrightnESS funding ends August 2018!

Staff (and knowledge) leaves, so scope will be again orphaned unless this is captured in the core budget.



Tasks		
Task 5.1	Creating a standard neutron event data stream for different detector types	ESS, KU
Task 5.2	Creating a standard method for streaming meta-data for fast applied fields	ESS, PSI
Task 5.3	Software to aggregate and make available the neutron event data and sample meta-data	ESS, PSI, Elettra ⁸

Event Formation

Scope

- **Processing algorithms are different for each detector type!**

Not all of them can be covered in the BrightnESS time frame.

- Data rates are huge and need to be processed at that rate on the fly.
- Software processing is more flexible, more easily scalable and can be better monitored for correctness than electronics.




Per hit/strip (raw data):

- VMM3 data: up to **38 bits** (depending on exact mode)
- Channel-ID
- Amplitude (10 bit)
- Time stamp (12 bit)

Per neutron (intermediate data):

- Average: 10 strips in x, 10 strips in y (**3.8 Kbit**)
- VMM3 data decoded into:
 - Strip ID (16 bit):
Detector number (0-2),
Module number (0-7),
chip number (0-24),
strip number (0-63)
 - Amplitude (10 bit),
Time stamp (12 bit)

Additional task: Clustering (also over several chips)



Neutron Event:
timestamp, pixelID

Resources

- BrightnESS (KU & DMSC)
- At least Core person (#1) required after end of BrightnESS to capture knowledge
- No neutron facility benchmark comparison because this is a unique new development at ESS, CERN operates a similar scheme with immense effort

Streaming and Aggregation Work Package



ESS runs EPICS as the distributed network control system.



The data aggregator collects the time stamped data from all sources and makes a unified data stream available for online processing and visualization (within the Instrument Data Group scope).

Dynamic changes to the configuration of instrument and sample environment need to be supported.

Outside of our scope:

- Control of the experiment (Instrument Data Group)
- Hardware device support (within ICS and technology groups).

BrightnESS has some funding for fast sample environment data capture.



Data Streaming, EPICS readout, File Writing



Resources

- STFC/ISIS in kind 12 PY, comparable to SNS or DLS staff levels
- one core team member #1 (3 PY) for integration and testing
- Some BrightnESS assistance until August 2018


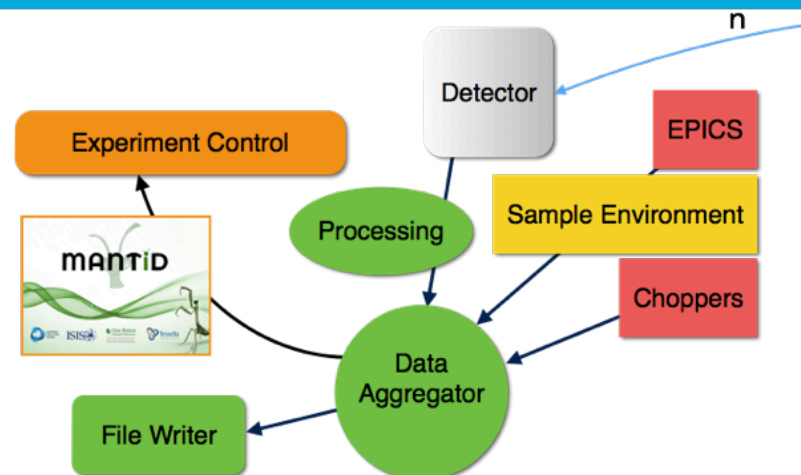
Scope

- EPICS readout non-event detectors and metadata
- Data Aggregator - unique data sink and source
- NeXus File Writing

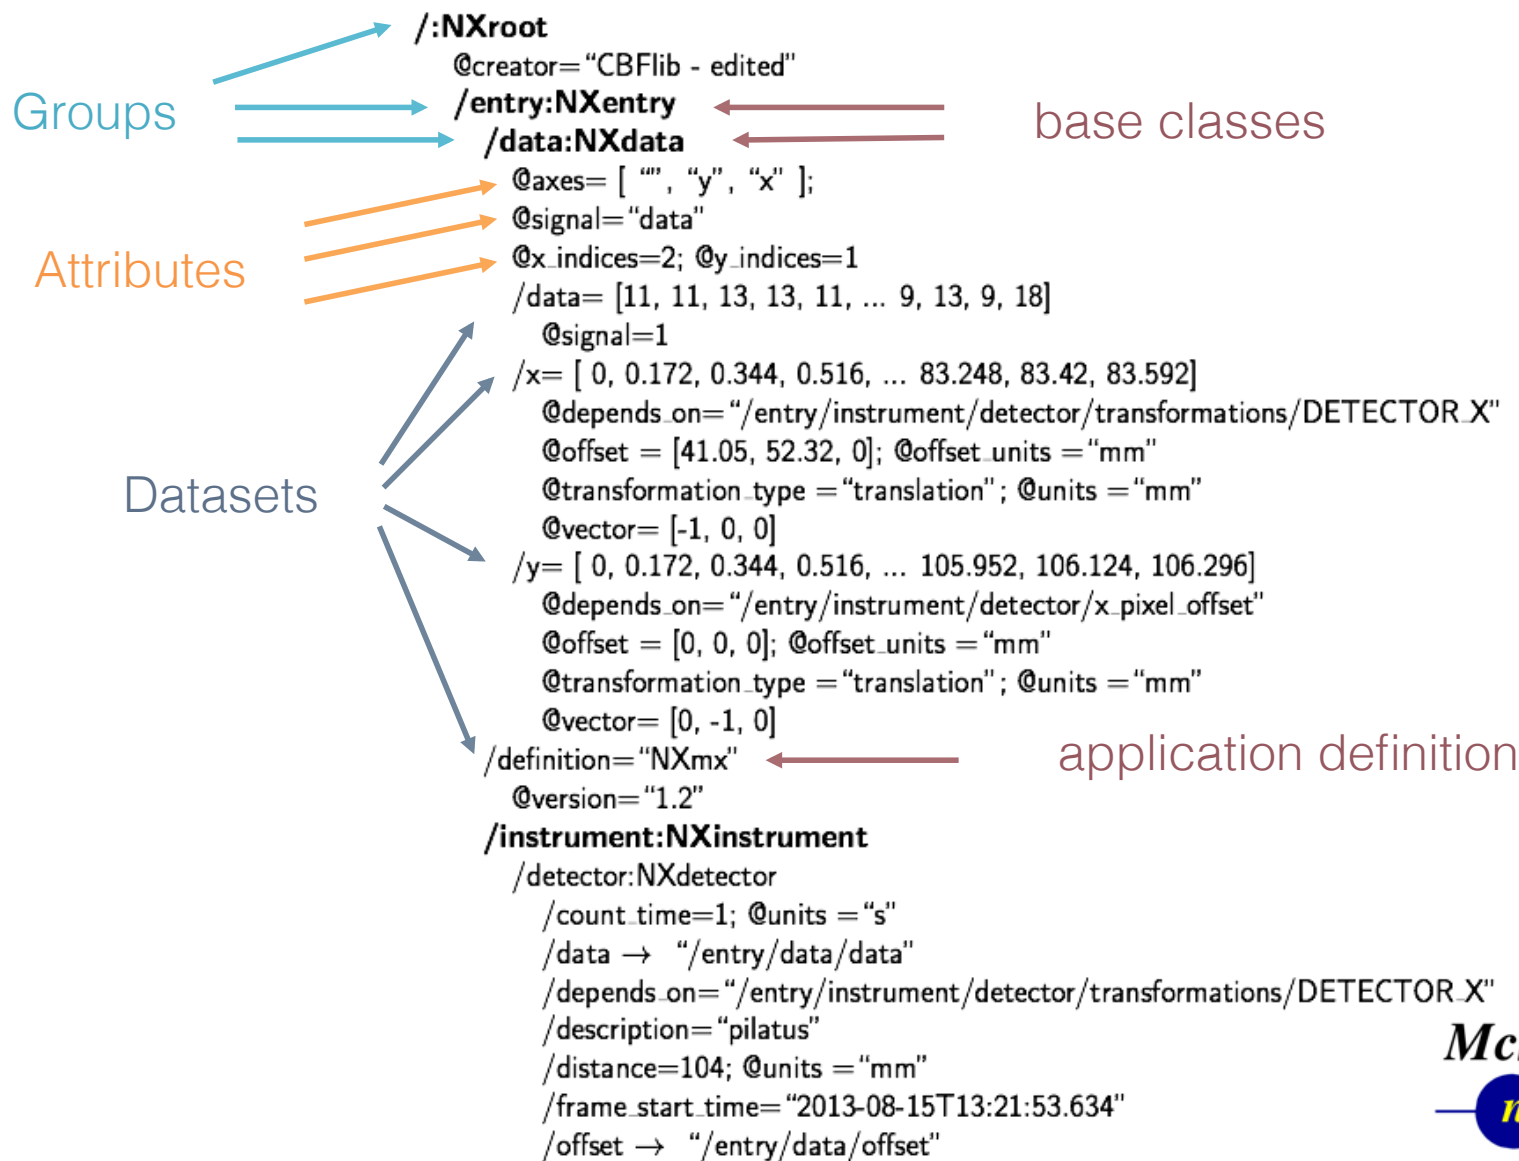


Requirements

- reliable operation
- high availability
- resilience against hardware and network problems
- long term maintainability



NeXus Data files



Files are written in a community agreed, high performance data format to a central location.



The HDF Group



Data Curation and Processing



This WP includes capturing the information required for automatic data reduction and analysis as well as the software infrastructure for running it.

Providing users with a convenient overview of their data:

- Where is the raw data for sample X?
- Have I measured each of my samples?
- Give me all reduced data for samples measured at 5K!

Makes sensible use of the data more likely and opens the door for more automation and higher sample throughput.

This effort also ties together many systems. That means it holds a large proportion of integration.

Data Curation and Processing In Kind

Resources

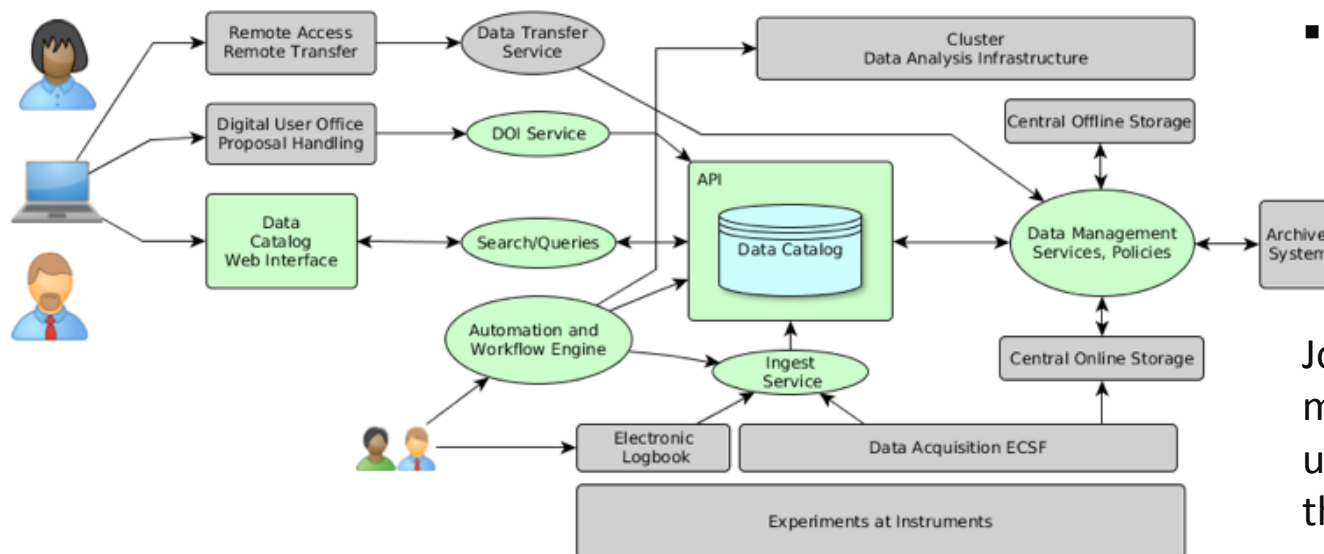
- PSI in kind 6 PY, signed by PSI and agreed by Swiss in-kind committee
- DMSC Core Person #2 to drive requirements and lead integration
- Levels in line with effort at steady state facilities like ISIS or SNS



Objectives

- keep track of the location of data files at any point in time
- Capture associated information about calibration and scientific intent
- automation and workflow for data processing and handling
- recording of experimental parameters for data reduction
- making data findable and accessible for users on site, remotely and long term

Scope



A Swiss-funded curation project just started.

Joining now with in kind and a matching core staff member puts us in the best position to drive the requirements for our science needs.

With many stakeholders integration is a major task. Instruments and critical equipment are being delivered by different partners, but need to work together seamlessly in commissioning.

At the same time software can not be delivered bug free.

With a sufficient amount of testing, the integration will benefit and the number of bugs found in commissioning and operations can be reduced.

Simple Testing Can Prevent Most Critical Failures: An Analysis of Production Failures in Distributed Data-Intensive Systems

CONNECT WITH US



Authors:

Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Renna Rodrigues, Xu Zhao, Yongle Zhang, Pranay U. Jain, and Michael Stumm, *University of Toronto*

Open Access Content

USENIX is committed to Open Access to the research presented at our events. Papers and proceedings are freely available to everyone once the event begins. Any video, audio, and/or slides that are posted after the event are also free and open to everyone. [Support USENIX](#) and our commitment to Open Access.

 [Yuan PDF](#)


 [View the slides](#)

 [BibTeX](#)

Abstract:

Large, production quality distributed systems still fail periodically, and do so sometimes catastrophically, where most or all users experience an outage or data loss. We present the result of a comprehensive study investigating 198 randomly selected, user-reported failures that occurred on Cassandra, HBase, Hadoop Distributed File System (HDFS), Hadoop

Integration and Testing Facilities



ESSIIP Lab Lund

Main integration point with NSS technology groups, ICS, in kind and BrightnESS partners, and within DMSC.
Procurement of three 10GB/s servers initiated for realistic performance evaluation.

Other Testing Facilities

DMSC

virtual machines for automated builds and deployment, static verification and simple functional tests

PSI

Independent 10 GB/s performance test setup for their needs

STFC/ISIS

mini verification cluster

HZB test beamline

Potential for real functional demonstrator

Software Status



A number of code repositories have been created, some for short term evaluation projects. They are kept in a common location on Bitbucket. The Atlassian tool suite (Jira and Confluence) is used to integrate workflows of the disperse teams.

















Continuous test and integration is being setup for the longer living projects.

Test deployments planned in 4 countries some involving real hardware and some with simulated devices.



European Spallation Source / Data Management

Repositories

Repository	Last updated
 damacluster	2016-05-26
 data-aggregator	2016-06-15
 data-aggregator-ioc-sim	2016-06-15
 data-management-dev-env	2016-06-15
 data-management-essip-deployment	17 minutes ago
 data-management-essip-test	3 days ago
 eventdata_via_EPICSv4_prototype	2016-04-14
 kafka-example	2016-04-28
 kafkaGenerator	2 days ago
 kafka-minimal	2016-05-24
 mcstasGenerator	2016-06-02
 NMX	2016-05-17
 sinq-amorsim	2016-06-17
 stub-config-service	an hour ago
 uTPC	2016-05-23
 zmq_protocol_buffer_example	2016-05-16

Progress in BrightnESS and In-Kind

Event Formation

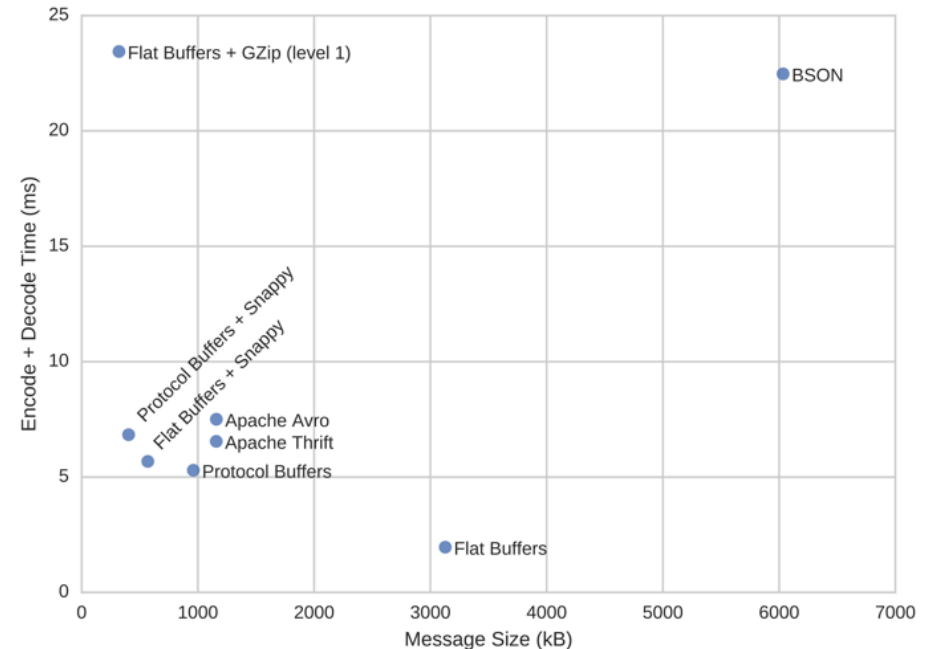
- Data from detector trials at IFE/R2D2 are being classified and analysed
- Baseline algorithms for NMX investigated
- Capabilities of ESS readout chain being explored

Data Steaming

- Completed project to evaluate Apache Kafka, zeroMQ, and EPICS v4 as transport technology
- Decision made on wire format serialization (Google Protocol Buffers, BSON/JSON, Flat Buffers)
- Reviewed first results from event generator tests
- Initiated procurement of 10GB network test equipment for ESSiIP integration lab project
- Steaming part of BrightnESS with PSI/Elettra initiated

Data Curation

- Refined Work Package Structure and in kind delivery with PSI
- Agreement signed by Swiss Partners and ready for submission to IKRC



Group Strategy and Work Units

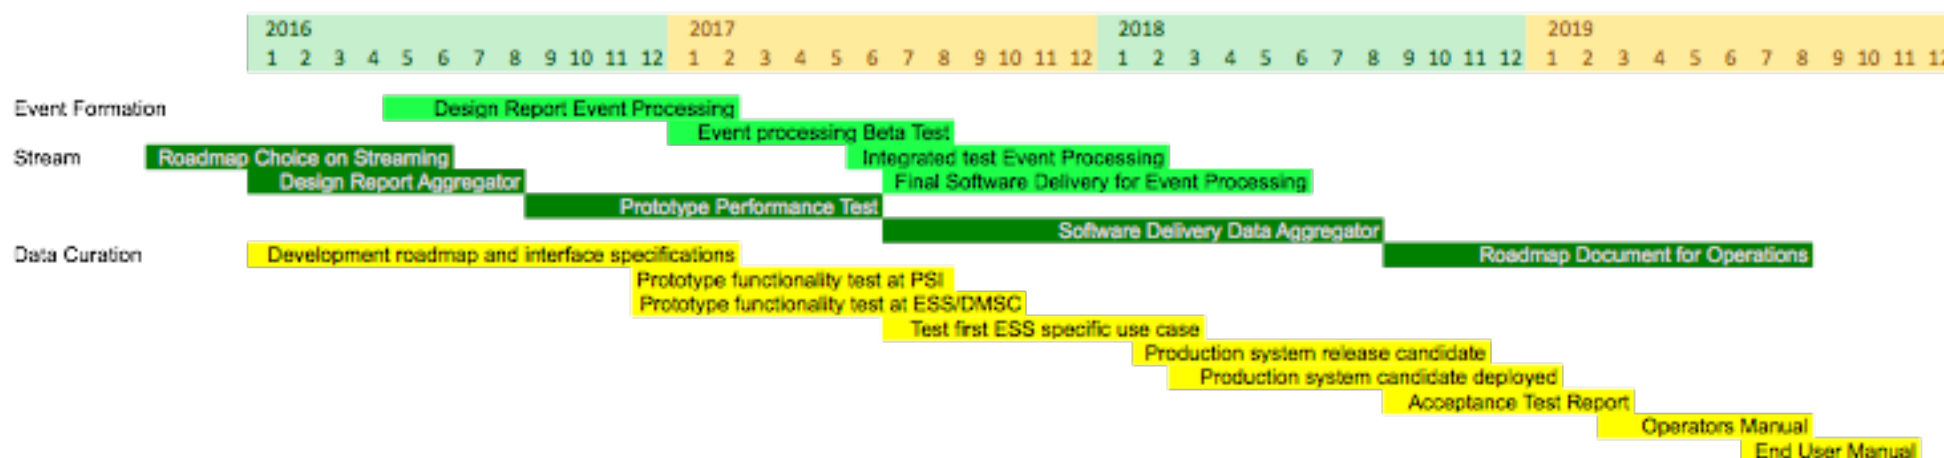


Strategy

- Core people in main areas to coordinate work and to preserve knowledge for operations (mitigating a primary risk)
- Forge strategic partnerships for in kind and collaborative projects
- Cover orphaned scope with third party money (BrightnESS)

Scope	Core Person	In Kind Partner	BrightnESS Partners
Event Formation	#1 (after end of BrightnESS)	-	ESS, KU
Fast Sample Environment	-	-	ESS, PSI
Data Streaming	#1 frozen	ISIS (started)	ESS, PSI, Elettra
EPICS readout	#1 frozen	ISIS (started)	-
File Writing	#1 frozen	ISIS (started)	-
Data Curation and Processing	#2 frozen	PSI (signed, on hold)	-

High-level Schedule



Notes:

- Schedule constructed for first instruments in 2019
- Still on track for first instruments in 2019
- Frozen recruitments are eating up the float in the schedule
- BrightnESS ends in August 2018

BrightnESS scope is essential to the delivery of science at ESS.

Most of the Data Management work is “horizontal” – that means it caters equally for all instrument classes. With two exceptions:

- Event Processing algorithms are dependent on detector type
- Imaging uses non-event recording detectors, that require special treatment

Core team frozen, so not committed staff cost at the moment, except for group leader. This needs to change, so there is staff to lead the local integration and have a route into operations.

Not being able to transfer into operations as planned in 2019 leads to added costs and problems in retaining knowledge. These can only be partially mitigated by adapting plans.

Creative choice of an achievable milestone for 2019 to move into operations might be a solution.