BEAMLINE DATA MANAGEMENT AT THE SYNCHROTRON ANKA

A. Vondrous#, T. Jejkal, D. Ressmann, W. Mexner, R. Stotzka
KIT, 76344 Eggenstein-Leopoldshafen, Germany

Abstract

We present a data management architecture consisting of beamline data management (BLDM) and data repository to enable data management at the synchrotron facility ANKA. Nearly each measurement device writes data with a different format, size and speed on storage devices that are distributed over the synchrotron facility. The operators perform some data management tasks manually and individually for each measurement method. In order to support the operators, users and data analysts to manage the datasets, it is necessary to collect the data, aggregate metadata and to perform ingests into the data repository. The data management layer between the measurement devices and the data repository is referred to beamline data management, which performs data collection, metadata aggregation and data ingest. Shared libraries contain functionalities like migration, ingest or metadata aggregation and form the basis of the BLDM. The workflows and the current state of execution are persisted to enable monitoring and error handling. After data ingest into the data repository, archiving, content preservation or bit preservation services are provided for the ingested data. The data repository is implemented with the KIT Data Manager. In summary, BLDM can connect the existing infrastructure with the data repository without major changes of routine processes to build a data repository for a synchrotron.

INTRODUCTION

Data management is essential for science in the information age. Preserving primary data produced in scientific investigations is not only important to establish scientific integrity [1, 2, 3]. The commodity for gaining knowledge and competitive advantages is data. Extracted relations from data stocks provide insights and help to understand phenomena. Beside good scientific practice and efficient data analysis methods, science also benefits from sharing data.

Astronomy is the textbook example for increasing the scientific outcome by sharing data. As an example, according to the Hubble space telescope bibliography [4] about 50% of the published papers related to measurements of the Hubble space telescope are based on publically available data.

To create the basement for responsible data preservation, the European commission [5] introduces data management requirements for funding scientific investigations. Large facilities like synchrotrons produce a not negligible amount of valuable data, such that we try to find a convenient way to establish data management capabilities for synchrotrons.

The following two examples of biology and materials science illustrate the reasons for an elaborated data management at a synchrotron facility.

The value of measurement data produced in a synchrotron is not only determined by hardware, experience and operational methods, it is also determined by the application. As an example, measurements of biological specimens reveal insights into the biomechanical processes of insects as shown e.g. by van de Kamp et al. in [6]. The value to the biology community is higher than the actual costs because the insights answer questions, confirm argumentation chains and create new perspectives. Storing the raw and derived data is of interest to preserve the findings for further analysis.

Sometimes, the scientific application does not exist yet and the value of the measurement data cannot be properly estimated. Measurements in the field of materials science have the potential to be of use for science and industry in the future. The crystallography open database [7] is one example for preserving crystallographic data and providing access to structured datasets.

Storing data in a structured manner for long-term usage is a challenge in the synchrotron context. The broad spectral range of the produced light and the number of measurement huts (beamlines) enable the usage of diverse measurement methods simultaneously.

Manual data management by beamline operators with logbooks or with spreadsheets is common practice. Tasks as search, retrieval, analysis or conversion are becoming demanding considering large files, many files or distributed storage locations.

To preserve valuable data, to automate data management tasks and to support operators and users at the synchrotron ANKA, a data repository is going to be implemented. The data repository keeps track of all datasets and provides data services like search, preservation, analysis, publication, curation, processing or migration.

Those benefits of a data repository are not for free. It is necessary to agree on data structures, to define metadata schemes and to aggregate the metadata. In addition, the data management has to be aware of the synchrotron specific infrastructure properties. Finally, the data management should interfere as less as possible to enable smooth beamline operation.

In this work we present data management from the source to the data repository considering synchrotron specific requirements.

METHOD

We divide the data flow from the data source (camera or detector) to the archive into two management areas. The first area is beamline data management (BLDM),

Corresponding author: alexander.vondrous@kit.edu

which is responsible to perform onsite data management services. The data repository manages the second area and is responsible for data services and user interaction. The separate management areas are shown schematically in the overview Figure 1 from a data flow perspective. Some components like the proposal system for metadata acquisition, the web fronted for controlling the data flow or visualization and computing resources are neglected.

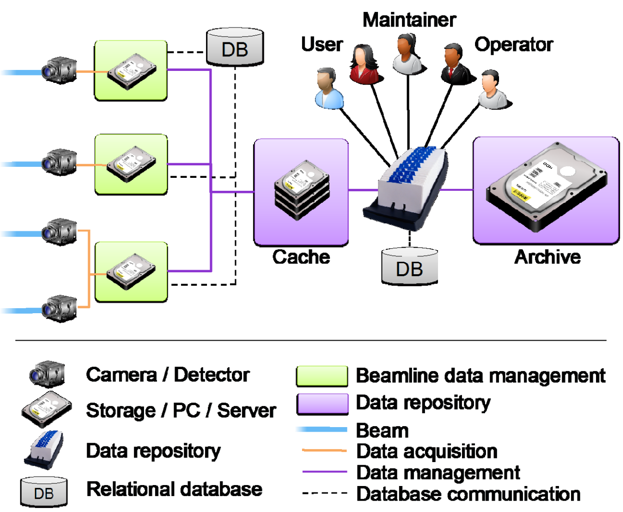

Figure 1: Data management components and management areas are separated into beamline data management and data repository.

Beamline Data Management

In order to manage data, we group measurement results (files) into datasets. Each dataset should contain the results of one measurement and their corresponding meta data. The structure of a dataset has to be known to enable data services like automated analysis, metadata extraction or curation.

Because the datasets are valuable, we protect them against accidentally manipulation by users through changing the ownership after data acquisition to a service account referred to data repository user and provide read only access to the investigating user group.

Before the datasets are ingested, metadata is aggregated and checksums of the files are computed to detect data transfer errors. The final steps are to register the dataset in the data repository and to copy the dataset content to the cache location. The cache is managed by the data repository, which checks file integrity and copies the files to the archive storage. The metadata of the dataset are used for efficient search and data organisation. The following list represents an expected workflow at the beamline.

1. Create empty dataset
2. Measure data
3. Store dataset
4. Change owner
5. Change mode
6. Compute checksum
7. Aggregate metadata
8. Register dataset
9. Copy files to cache
10. Release for archiving workflow
11. Remove dataset from BLDM storage

All BLDM tasks related to the measurement data like ownership change, checksum computation or file transfer are implemented in a shared library. The tasks are managed by, and communicate through a relational database. With this database, the tasks are combined to a workflow, which is executed stepwise to ingest a dataset in the data repository. The idea of using tasks is inspired by the UNIX philosophy to “Make each program do one thing well.” [8].

To connect the data sources with the data repository for convenient data management, we provide an interface of data management tasks to the control system. The control system used at ANKA for user operation is TANGO [9]. With this interface, data management processes can be integrated into the daily measurement tasks of beamline operators. Figure 2 illustrates the structure of one task handler instance with some tasks.

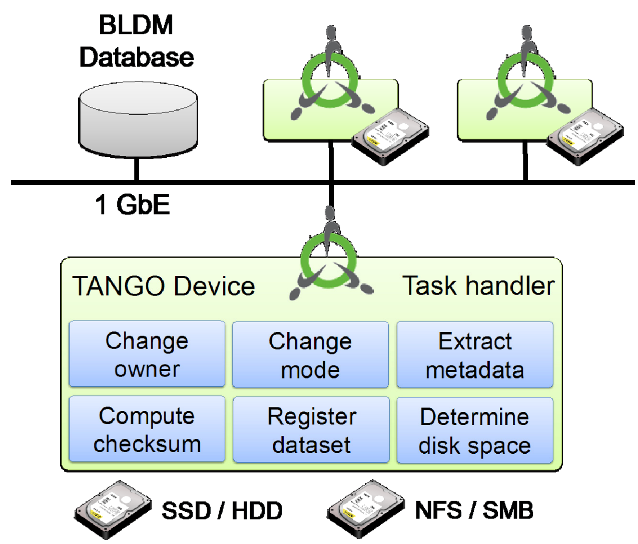

Figure 2: One task handler instance (green) using shared library tasks (blue) with a TANGO interface to integrate data management functionality into the synchrotron control system. Each task has access to the BLDM database.

The tasks and the TANGO device are implemented in C++ with the POCO [10] and Boost [11] libraries for logging and shared library access. It is necessary to provide operations on files after the data acquisition on the PC.

Data Repository

The data repository is implemented with the KIT Data Manager [12]. After the task handler registered and copied the dataset to the cache location, the data repository archiving workflow is triggered. The archiving process consists mainly of file integrity checks, file copy operations and metadata indexing. A representational state transfer (REST) interface can be used to build a user

ISBN 978-3-95450-146-5

interface or to access data with other programs or computers for further investigation.

Ingest Setup

We evaluated the transfer rate during the ingest process from a workstation to the data repository cache. The ingest process consists of 21 REST calls to register the dataset (8th workflow step) and the dataset upload via WebDAV (9th workflow step) over a one gigabit Ethernet connection between two Linux machines. Each dataset consists of one metadata file (382 byte) and one measurement data file with random content. The file size of the measurement data file is varied from one byte to 100 GB. Three ingests for each dataset are performed. The average value is used to estimate the total transfer rate.

RESULTS

The data repository and the BLDM are under construction but the key components are implemented:

Create empty dataset; Store dataset; Change owner; Change mode; Compute checksum; Register dataset; Copy files to cache

The evaluation results of the ingest processes are depicted in Table 1. The transfer rate is increasing with the file size. A file transfer of a 1 byte and a 1 KB measurement data file takes 7.28 and 7.26 seconds, such that we define the time to register the data to be 7.26 seconds. This registration time is used to estimate the transfer rate of the data upload. The estimated transfer rates for 1 MB and 10 MB without registration time do not reach a high transfer rate. Above 100 MB of file size the estimated transfer rates are greater 100 MB/s. The registration time of about 7 seconds is negligible for the large 100 GB dataset.

Table 1: Mean Ingest Runtime for Different File Sizes

File size	Mean time	Total transfer rate	Estimated data transfer rate*
1 byte	7.28 s	0.00 MB/s	-
1 KB	7.26 s	0.00 MB/s	-
1 MB	7.32 s	0.14 MB/s	16.67 MB/s
10 MB	7.58 s	1.32 MB/s	31.25 MB/s
100 MB	8.13 s	12.30 MB/s	114.94 MB/s
1 GB	16.75 s	61.13 MB/s	107.90 MB/s
10 GB	103.84 s	98.61 MB/s	106.03 MB/s
100 GB	955.85 s	107.13 MB/s	107.95 MB/s

The measurement results are performed with a simple test case consisting of one file. A more complex evaluation under production conditions with high performance hardware should be conducted.

* Establishing the connection and data registration are considered as registration time, such that 7.26 s are removed from the Mean time to estimate the data transfer time.

CONCLUSION

The separation into two data management areas enables provides a convenient implementation of the interface with the synchrotron specific control system for seamless integration into beamline operation.

The presented architecture to execute data related tasks is one component for data management to hide complexity and serve modern data management requirements.

ACKNOWLEDGMENT

The authors thank the German Helmholtz Association for funding the project, the Large Scale Data Management and Analysis project, and the Large Scale Data Facility for providing storage resources. We also thank Halil Pasic for his contributions to the project.

REFERENCES

[1] Deutsche Forschungsgemeinschaft, "Safeguarding Good Scientific Practice", WILEY-VCH 2013

[2] Max-Planck-Gesellschaft, "Regeln zur Sicherung guter wissenschaftlicher Praxis", MPG 2009, http://www.mpikg.mpg.de/23731 (visited Oct. 2014)

[3] European Science Foundation, "Good scientific practice in research and scholarship", ESF 2000, http://www.esf.org/fileadmin/Public_documents/Publications/ESPB10.pdf (visited Oct. 2014)

[4] J. Lagerstrom, "Measuring the Impact of the Hubble Space Telescope: open data as a catalyst for science", World Library and Information Congress: 76th IFLA 2010 in Gothenburg, Sweden

[5] European Parliament and Council, "establishing Horizon 2020 – the Framework Programme for Research and Innovation", Official Journal of the European Union 2013

[6] T. van de Kamp et al., "A Biological Screw in a Beetle's Leg" Science 1 July 2011: 333 (6038), 52. [DOI:10.1126/science.1204245]

[7] S. Gražulis et al. "Crystallography Open Database (COD): an open-access collection of crystal structures and platform for world-wide collaboration". Nucleic Acids Research 40, p. 420-427

[8] E.S. Raymond, *The Art of UNIX Programming*, (Addison-Wesley Professional 2003)

[9] TANGO Control System, http://www.tango-controls.org/ (visited Oct. 2014)

[10] POCO, http://pocoproject.org/ (visited Oct. 2014)

[11] BOOST, http://www.boost.org/ (visited Oct. 2014)

[12] C. Jung, A. Streit, Large-Scale Data Management and Analysis (LSDMA) - Big Data in Science, (Karlsruhe 2014) page 9

ISBN 978-3-95450-146-5