EuroHPC-01-2019

IO-SEA

IO – Software for Exascale Architectures

Grant Agreement Number: 955811

D7.5

Data Management Plan

Final

Version: 1.0
Author(s): Maike Gilliot (CEA)
Contributor(s): Philippe Deniel (CEA)
Date: 30.09.2021

Project and Deliverable Information Sheet

IO-SEA Project	**Project ref. No.:**	955811
	Project Title:	IO – Software for Exascale Architectures
	Project Web Site:	https://www.iosea-project.eu/
	Deliverable ID:	D7.5
	Deliverable Nature:	Report
	Deliverable Level: PU *	**Contractual Date of Delivery:** 30/09/2021
		Actual Date of Delivery: 30/09/2021
	EC Project Officer:	Daniel Opalka

*– The dissemination levels are indicated as follows: **PU** - Public, **PP** - Restricted to other participants (including the Commissions Services), **RE** - Restricted to a group specified by the consortium (including the Commission Services), **CO** - Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

Document	**Title:** Data Management Plan	
	ID: D7.5	
	Version: 1.0	**Status:** Final
	Available at: https://www.iosea-project.eu/	
	Software Tool: LaTeX	
	File(s): IO-SEA_D7.5_DMP.pdf	
Authorship	**Written by:**	Maike Gilliot (CEA)
	Contributors:	Philippe Deniel (CEA)
	Reviewed by:	André Brinkmann (JGU) Artur Podobas (KTH)
	Approved by:	Exec Board/WP7 Core Group

Document Status Sheet

Version	Date	Status	Comments
0.1	15.08.2021	Document setup and first outline	Shared internally for first comments and feedback
0.2	15.09.2021	Integration of first comments	
0.3	28.09.2021	Shared for internal review and the WP7 Core Group	Comments of the internal review have been integrated, plus format correction and language issues
1.00	30.09.2021	Final version for submission	

Document Keywords

Keywords:	IO-SEA, HPC, Exascale, Software, Data Management Plan, Intellectual Property Rights (IPR), IP Management, Innovation Management

Contents

Executive Summary

This deliverable presents IO-SEA's Data Management Plan (DMP). It describes the strategy for coping with data generated, used, modified and stored as part of the project activities. In order to make this data findable, accessible, interoperable and reusable according to the FAIR principles [1], this DMP details different aspects of data management strategy of IO-SEA:

- **The overall principles and the overall approach for handling data during and after the end of the project** are based on the principles of our consortium agreement. In addition, a project-internal approach for managing data has been set up, in particular for controlling the release into a central, public "data hub" for all IO-SEA related data that can be shared publicly.
- **The description of the different data sets that will be generated, processed and stored**, including information on standards, whether data will be shared, and how data will be curated and preserved during the project life time and beyond. The key data sets identified so far are (1) applications and other executables, (2) data related to the monitoring of the IO-performance of our use case applications, (3) data related to the benchmarking for assessing the performance of the storage architecture.

This DMP is not a final document. In particular, the data sets will evolve during the project life time. Updates will be provided via the yearly project reports.

1 Introduction

This deliverable presents IO-SEA's Data Management Plan (DMP). It describes how data generated, used, and modified within IO-SEA will be managed. Moreover, it describes how this data will be made findable, accessible, interoperable and re-usable according to the FAIR principles [1] in order to be of benefit for the wider research community. Section 2 presents the project itself and its objectives as to gain an overview of the scientific data sets arising from the project's activities.

Section 3 then describes the general principles for using and sharing this data based on concepts laid out in the consortium agreement. It also describes the process that has been set up in order to control the release of data into a central, public "data hub" for all IO-SEA related data that can be shared publicly.

The different data sets as planned as of today are detailed in Section 4. It must be noted that this DMP is not a final document. In particular, the data sets and their usage and storage policies will evolve, and other data sets may arise during the project life time. This will be integrated step by step. We close this deliverable with our conclusions in Section 5.

2 Project Description

IO-SEA aims to provide a novel data management and storage platform for exascale computing based on object stores, hierarchical storage management (HSM) and on-demand provisioning of storage services. The platform will efficiently make use of storage tiers spanning NVMe and NVRAM at the top all the way down to tape-based technologies. System requirements are driven by data intensive use-cases, in a very strict co-design approach. The concept of data nodes and data accessors is introduced that allow users to flexibly operate the system, using various well-known data access paradigms, such as POSIX namespaces, S3-Swift interfaces, MPI-IO and other data formats and protocols. These resources eliminate the problem of treating storage resources as static and unchanging system components – which is not a tenable proposition for data intensive exascale environments. The methods and techniques are applicable to exascale class data intensive applications and workflows that need to be deployed in highly heterogeneous computing environments.

Critical aspects of intelligent data placement are considered for extreme volumes of data. This ensures that the right resources among the storage tiers are used and accessed by data nodes as close as possible to compute nodes, optimising performance, cost, and energy at extreme scale. Advanced IO instrumentation and monitoring features will be developed to that effect leveraging the latest advancements in AI and machine learning to systematically analyse the telemetry records to make smart decisions on data placement. These ideas coupled with in-storage-computation remove unnecessary data movements within the system.

It leverages technologies developed by the SAGE and the NextGen-IO projects and strengthens the TLR[1] of the developed products and technologies ([2], [3]).

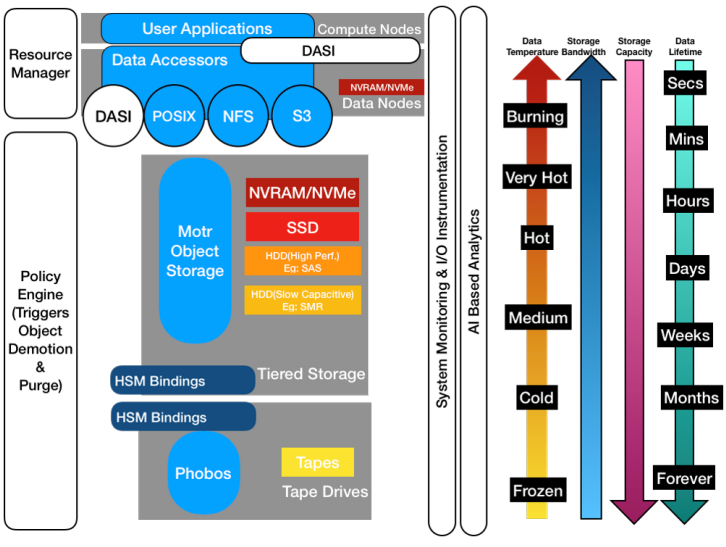

Figure 1: Overview IO-SEA architecture

[1] Technology Readiness Level

Figure 1 shows the different elements of the IO-SEA storage architecture, with at its heart the tiered storage module (based on application Motr, developed within SAGE2) and the tape storage module Phobos. The Data Access and Storage Interface (DASI) will provide an application and user-centric layer for accessing the data via the different I/O services (POSIX, NFS, S3, ..). As transverals elements, we have first all monitoring relatated elements: They alllow to observe the data movements and to assess of the data placement strategies ex-post. The second transversal element is the Policy Engine deciding on data placement and on data movements, and third all benchmarking related components for assessing the overall impact of the data storage policy on the runtime of the use case applications.

3 General principles regarding data management in IO-SEA

The overall data management and intellectual property (IP) strategy is derived from some underlying elements. First of all, of course the grant agreement and the consortium agreement, which has been signed by all project partners. The project partners have designed the project proposal with the intentions to (1) leverage technologies developed in previous storage-related EU-funded projects, and (2) to release - where possible - its results as open source. These principles are detailed in the remainder of this section.

3.1 General principles regarding data ownership and conditions for sharing data

A consortium agreement was negotiated and signed by all project partners. It specifies *inter alia* the terms and the conditions of ownership, access rights to background and results, and exploitation rights, based on the DESCA model [4]. Regarding the ownership of results generated within the project, the consortium agreement foresees:

- Results are owned by the party or parties that generated them. In the case of joint ownership, the respective ownership shares will be decided on the basis of the resources utilised by each party for the generation of the result, the background and knowledge contributed by each party, and the respective inventive contribution of each party.
- The owner(s) decide(s) of the licence terms and the conditions for use and for sharing the result, in accordance with the principles of the grant agreement and the consortium agreement.
- The owner(s) also decide(s) on possible measure for IP protection.
- Each party remains the owner of its background.

All decisions with respect to the conditions for using and for sharing data will follow these stipulations. As part of IO-SEA's IP management, a process will be set up to formally trace the decisions taken by the partners regarding the licence terms of the results. On a day-to-day basis, this process will be followed-up by the Executive Board, which brings together all work package leaders. The Project Board, consisting of one representative of each project partner, shall be informed as well.

3.2 Responsibility for data

Every partner is responsible for the data they are collecting. The data will be collected, combined, stored and transmitted according to the relevant national, European and institutional regulations. At project level, the responsibility lies with the scientific coordinator. Regarding IPR issues of applications and other pieces of code used or developed within IO-SEA, he is supported by the innovation manager.

Personal data as defined by the General Data Protection Regulation (GDPR) [5] is not needed for the execution of the research activities as described in the Description of Action. Only for the purpose of coordinating and managing the research activities some personal data is required (for project-internal mailing lists, for example).

3.3 Overall IO-SEA IP strategy

The IP approach of IO-SEA's storage architecture is based on two pillars. First, it will leverage existing technologies. For example, the storage engine Cortx-Motr [6] has been developed within SAGE2 [2], and the code Phobos has been developed at CEA for many years now [7]. By doing so, this project federates the competences and the communities of the existing technologies and can achieve higher TRLs.

Second, IO-SEA relies strongly on open source codes. For example, Motr-Cortx is licensed under Apache2.0, Phobos under LGPL2.1. In addition, proprietary software will also be used, as for example the IO-Pattern Analyzer of Atos [8].

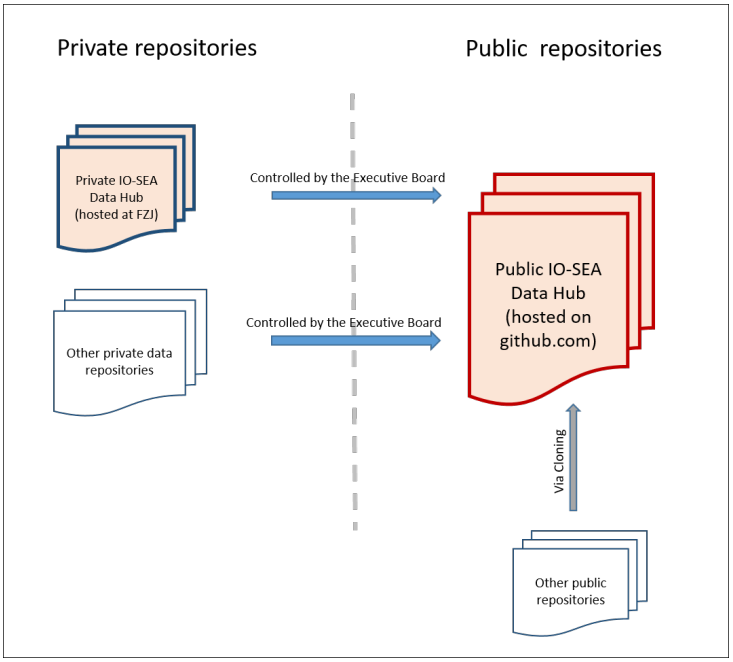

Figure 2: Overview on the different data repositories

Thus, the goal is not to achieve a monolithic entity under one single licensing scheme, but rather a composition of different bricks, which are interoperable on a technical level, and also from an IP point of view. The innovation manager keeps track of the different components and their IPR and ensures a coherent exploitation of the project results. To this end, all applications and other executables are tracked via the IO-SEA software licensing and source code repository. This repository is under the control of the innovation manager, who ensures that it is updated regularly by all project partners.

3.4 Central IO-SEA data hubs

As mentioned before, IO-SEA will build upon existing applications, stored on different sites, some of which are public, others are private.In order to make the project outcome easily accessible for dissemination and for exploitation, we will collect and reference all data elements of IO-SEA in two central data hubs. As shown in Figure 2, we opted for two different gitlab instance: a private one (hosted at the Jülich Computing Centre), and a public instance (hosted on github.com). Whereas the first will remain private and only accessible by IO-SEA partners, the second gitlab project will be made public as soon as first data will be collected. The transfer of data from the private gitlab instance to the pubic one is controlled by the Executive Board. This only applies to data sets for which the data owner(s) agree to share it. Other data sets - such as for example some proprietary use case application - will remain on the private repositories.

4 Data Set Description

The main data sets can be derived from IO-SEA's overall architecture. First, we have the **use case applications and other executable**, including all pieces of code modified or developed as part of the IO-SEA project activity. A second set relates to the **monitoring activity**, including the logs of the observed system activities. The logs and traces generated during the **benchmarking of the use case applications** form another data set, along with the **scientific data need for benchmarking**.

Some other data sets focus less on the data generated while performing the research activities, but more on annex data sets related to the communication and the dissemination of the outcome of the research activities, as well as the management and the coordination of this research activities. These data sets gather all kind of **documents and reports**, **graphics, pictures and videos**, and on **training material**.

Data set #1	**Applications and executables**
Type of data	Source code and executables of IO-SEA's use case applications, as well as any other piece of code developed or modified within IO-SEA.
Format	C, Python, makefiles
Data source	Many of these applications have been developed in previous projects or collaboration and will be improved and enhanced as part of the IO-SEA project. Nevertheless, some will be developed from the scratch within IO-SEA.
Expected size of data	1-2 GB
Data property	The stipulations of the consortium agreement apply, in particular the sections on "Background" and "(Joint) Results".
Licence terms and conditions for use and for sharing	This will depend on the IPR scheme of each specific application. The innovation manager assures that for all applications and other executables which are a result of the research activity a licence is agreed on by the owner(s). Note that many of the applciations used and further developed in IO-SEA are open source.
Storage medium and storage location	Today, the applications are stored in different repositories, some of which are public, and some of which are private. As described in Figure 2, the data will be made accessible via the IO-SEA data hub.

Data set #2	**System Monitoring**
Type of data	All types of logs and traces of (1) the observed system itself, and (2) of the data movements between the data nodes (for a given use case application). In second step, once the AI-based recommender system is in place, the recommendations themselves will be part of this data set too.
Format	Plain text files(.txt, .rtf, ...)
Data source	This data will be generated within the IO-SEA project.
Expected size of data	Several hundreds of GB
Data property	The consortium agreement and its stipulations for results apply.
Licence terms and conditions for use and for sharing	Terms for sharing and for using this data (all of it or partially) are not settled yet, and will be decided by the owners in accordance to the general rules in the consortium agreement.
Storage medium and storage location	In a first step, this data will be stored on the plate form where the applications are executed and monitored. To what extend a longer term storage strategy is needed will be defined in the coming months.

Data set #3	**Benchmarking: Scientific input/output data of the use case applications**
Type of data	To assess the IO-SEA storage framework, the use cases will regularly be executed on a default platform to assess their IO performance via the JUBE benchmarking tool. For each use case application, two or three "standard" input data sets will be used repeatedly in order to get comparable performances. In the context of IO-SEA, the scientific input and output data to the use case applications is not of particular interest.
Format	Plain text (.txt, .json, .xml, ...)
Data source	The input data will be provided by the use case owners and remains under its original licence scheme. The scientific output data will be generated by the use case application.
Expected size of data	These data sets can become very large (up to 50-80 GB for a single run).
Data property	The data sets are background, and are provided by the use case owners, and remain their property. Also the scientific output data is their property.

Licence terms and conditions for use and for sharing	Some of these data sets have been made available in the past by the use case owners as part of their open source strategy for their application. The project partners will investigate how this data sets can be shared in order to allow the community at large to check and validate the terms for sharing and for using this data (all of it or partially) are not settled yet, and will be decided by the owners in accordance to the general rules in the consortium agreement.
Storage medium and storage location	In a first step, this data will be stored on the platform where the applications are executed. To what extend a longer term storage strategy is needed will be defined in the coming months.

Data set #4	**Benchmarking configuration files, as well as raw and aggregated benchmarking data**
Type of data	This data set includes the configuration files, the raw and the aggregated benchmarking data, as well as all other auxiliary output files.
Format	Plain text (.rtf, .txt, ...)
Data source	This data will be generated within the IO-SEA project
Expected size of data	Very small, with at most 1 MB per file.
Data property	To be decided based on the principles laid out in the consortium agreement.
Licence terms and conditions for use and for sharing	The aggregated benchmarking data could be of interest to the scientific community. The owners will jointly agree on an appropriate licence, with the aim to be as open as possible.
Storage medium and storage location	In a first step, this data will be stored within the JUBE benchmarking system, only the relevant data will be kept and moved tonto the IO-SEA data hubs.

Data set #5	**Documents and Reports**
Type of data	This data set includes text-based documents for internal use (such as meeting minutes and internal presentations), as well as confidential documents (grant agreement, the IO-SEA consortium agreement, exchanges with the funding agency as well as reports and presentations related to the project reviews). It also covers all documents for external use, such as deliverables, reports, white papers, public presentations, or the content of the IO-SEA web page [9].
Format	.pdf, .docx, .tex, .pptx, .rtf, .odt, .xlsx, .txt
Data source	This data will be generated within the IO-SEA project

Expected size of data	A total size of 80-150 MB is expected.
Data property	This data is the property of the IO-SEA consortium, unless explicitly stipulated differently.
Licence terms and conditions for use and for sharing	All internal documents (such as meeting minutes, meeting presentations, the grant agreement or the consortium agreement) will not be shared, and remain for internal use only. However, most deliverables will be made publicly available on our web page, once they deliverable have been approved by EuroHPC.
Storage medium and storage location	In a first step, this data will be stored a shared work space server hosted at FZJ, where the data is accessible to all IO-SEA participants. Elements approved for publication, will then either be directly accessible via the IO-SEA web page or referenced on the web page (for example for scientific publications).

Data set #6	**Graphics and pictures**
Type of data	This data set includes graphics and images representing graphically technical aspects of our project and its related use case applications, which might be either for internal use only or approved for publication as part of a deliverable, a report, or a scientific publication. This data set also includes pictures taken during IO-SEA meetings, or on other events.
Format	.pdf, .jpeg, .png, .bmp, .tiff, .gif, .jif
Data source	Apart from some graphics related to existing applications (such as the use case applications), most of this data will be generated within the IO-SEA project.
Expected size of data	Up to 100 GB
Data property	Graphics are under the same regime as the document using these graphics, except if otherwise stated.
Licence terms and conditions for use and for sharing	In particular for pictures, where GPDR may apply, the use will be determined on a case-by-case basis. Graphics representing technical aspects are under the same regime as the document using these graphics, except if otherwise stated.
Storage medium and storage location	Pictures will be stored on the shared work space server hosted at FZJ, where the data is accessible to all IO-SEA participants.

Data set #7	Training material
Type of data	As part of the IO-SEA project, material for training will be developed. This may include slides, other presentations and written teaching material, videos, and documentation (including exercises).
Format	.pdf, .docx, .tex, .pptx, ;mp4, .avi, .mpg, .m4v, flv, .txt, .odt, .xlsx
Data source	The teaching material associated to IO-SEA will not be completely from the scratch, it might be partially based on existing material for training.
Expected size of data	Not expected to exceed 100 MB
Data property	To be defined on a case-to-case basis in accordance with the general principles described in the Grant Agreement.
Licence terms and conditions for use and for sharing	Generally, IO-SEA plans to make its training material widely available to all users and to ensure the possibility of "reuse" of the training material in the future. However, the precise terms of licence will be decided by the (joint) the owner(s), based on the principles laid out in the Grant Agreement.
Storage medium and storage location	During the project life time, this data will be stored a a shared work space server hosted at FZJ, where the data is accessible to all IO-SEA participants.

5 Conclusion

This deliverable describes the data management strategy of IO-SEA and its different components: The overall principles and the overall approach for handling data during and after the end of the project and also a description of the different data sets that will be generated, processed and stored within IO-SEA.

One key element are the two IO-SEA data hubs. The first hub is hosted by one of the project partners and will store data reserved to the project partners. A second hub (hosted at github.com) has been set up for sharing data with a wider community.

The elements provided within this report are only a first element towards a FAIR data usage. On a day-to-day basis, the innovation manager and the Executive Board of IO-SEA will follow up on these IP related issues, in strong collaboration with the work package on dissemination and exploitation. We will therefore provide regular updates via the yearly project reports.

Bibliography

[1] Directorate-General for Research & Innovation European Commission. H2020 guidelines on fair data management in horizon 2020 – version 3.0. https://ec.europa.eu/research/participants/data/ref/h2020/grants_manual/hi/oa_pilot/h2020-hi-oa-data-mgt_en.pdf, Online; accessed 29/09/2021.

[2] Seagate. Sage-2 project web site. https://sagestorage.eu, Online; accessed 29/09/2021.

[3] Nextgenio project web page. http://www.nextgenio.eu, Online; accessed 29/09/2021.

[4] Desca 2020 model consortium agreement. https://www.desca-agreement.eu/desca-model-consortium-agreement, Online; accessed 29/09/2021.

[5] European general data protection regulation. https://gdpr-info.eu, Online; accessed 29/09/2021.

[6] Seagate. Cortx-motr github. https://github.com/Seagate/cortx-motr, Online; accessed 29/09/2021.

[7] CEA Website. Phobos - short description. http://www-hpc.cea.fr/en/red/opensource.htm, Online; accessed 29/09/2021.

[8] Atos. Io pattern analyzer brochure. https://atos.net/wp-content/uploads/2018/07/CT_J1103_180616_RY_F_TOOLSTOIMPR_WEB.pdf, Online; accessed 29/09/2021.

[9] Iosea project web site. http://aiweb.techfak.uni-bielefeld.de/content/bworld-robot-control-software/, Online; accessed 29/09/2021.

[10] The Message Passing Interface Forum. Mpi: A message-passing interface standard – version 3.1, June 2015.

[11] Springer, editor. *SLURM: Simple Linux Utility for Resource Management*, volume 2862, 1962.

[12] FZJ Web Site. Jube benchmaking suite. https://www.fz-juelich.de/ias/jsc/EN/Expertise/Support/Software/JUBE/_node.html, Online; accessed 29/09/2021.