EuroHPC-01-2019

IO-SEA

IO – Software for Exascale Architectures

Grant Agreement Number: 955811

D7.3

Periodic Report - Year 3

Final

Version:	1.0
Author(s):	M. Gilliot (CEA)
Contributor(s):	
Date:	March 29, 2024

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:	D7.3
	Deliverable Nature:	Report
	Deliverable Level: PU *	**Contractual Date of Delivery:** 31 / March / 2024
		Actual Date of Delivery: 31/ March / 2024
	EC Project Officer:	René Chatwell

*– 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:** Periodic Report - Year 3	
	ID: D7.3	
	Version: 1.0	**Status:** Final
	Available at: https://www.iosea-project.eu/	
	Software Tool: LaTeX	
	File(s): IO_SEA_D1.2.pdf	
Authorship	**Written by:**	M. Gilliot (CEA)
	Contributors:	
	Reviewed by:	A. Ph. Deniel (CEA) B. I. Zhukov (JSC)
	Approved by:	Exec Board/WP7 Core Group

Document Status Sheet

Version	Date	Status	Comments
0.1	25/01/2024	Outline and first elements sketched	for internal discussion
0.2	26/03/2024	Draft ready for internal review	
0.3	28/03/2024	Version including comments from the internal review	
1.0	29/03/2024	Last changes and corrections for submission	approved by IO-SEA

Document Keywords

Keywords:	IO-SEA, HPC, Exascale, Software

Contents

List of Figures

List of Tables

Executive Summary

This is the final deliverable of WP7. WP7 is in charge of the scientific coordination and management of the IO-SEA project. This report presents the work performed by WP7 from April 2023 until March 2024 (M25-M36), building on the elements provided in D7.1 [1] (issued in M12), the first Interim Report (issued in M18), as well as D7.2 [2].

The report shows how WP7 provided organisational support to the other WPs for conducting their technical and scientific work in a structured and well-organised manner, and how WP7 ensured that the work conducted is in line with the DoA and the Consortium Agreement. It also shows how WP7 ensured the overall coherence of IO-SEA's effort and liaised efficiently with our funding authority, as well as our SEA-friends.

1 Introduction

WP7 covers all tasks related to the Management and the Scientific Coordination of the IO-SEA project. This deliverable reports on the activities from M25 to M36 performed in WP7 within the IO-SEA project. This report is the follow-up document of deliverable D7.1 and D7.2, which were issued in M12 and M24 respectively. In addition, the activities covering the first 18 months have already been described in the First Periodic Report, issued after M18. This report will therefore focus on the specific activities performed in the last year of the project, referencing D7.1, D7.2 and the Periodic Interim Report where relevant.

The activities of WP7 are divided into six tasks, for each of which the work performed will be presented in the next chapters.

- The management and communication activities are covered by the Task 7.1 and 7.5, respectively. These activities (as presented in Chapter 2) cover all aspects of running the project on a day-to-day basis, ensuring the progress of the work within the project with respect to DoA. In addition, these tasks also cover communication with our Project Officer.
- Task 7.2 focuses on controlling the quality of the project's output - in particular its deliverables. As a reminder, we provide details on the overall IO-SEA quality control process in [1]. In this deliverable, we focus in Chapter 3 on the deliverables and milestones of the period M25-M36.
- Task 7.3 (presented in Chapter 4) is coping with Risk Management. We describe the risks having materialised and the actions taken to mitigate these risks. The overall risk assessment process set up in IO-SEA is described in [1]. A full and detailed list of all risks identified in the DoA is also part of the Final Project Report, which is currently being prepared.
- Related to Task 7.4, Chapter 5 presents details on the financial status of the project. This data will cover the period M1-M36, but will be based on *preliminary* data only. A more detailed analysis, along with explanations and justifications on deviations from the planned effort, are also part of the Final Project Report, as not all financial data is available at the time of writing.
- The collaborations with RED-SEA, DEEP-SEA and other R&I projects funded under the same EuroHPC call (such as the ADMIRE project) is covered - from an organisational point of view - by Task 7.6 ("Task for the Complementary Grants"), detailed in Chapter 6. This task has been completed in M12, and the technical and scientific elements of these collaborations are presented as part of the WP6 reports.

2 Management and Communication

The management of IO-SEA is based on the “management by exception” approach with delegated responsibility [3]. In the context of IO-SEA, this means that the decision-taking body — the Project Board — is only solicited when the project deviates significantly from the planned activities and tasks, or encounters difficulties in fulfilling the project’s objectives. Changes to the Description of Action are also approved by the Project Board. The day-to-day management is done by the Executive Board, with the support of the Project Management Office (also referred to as WP7 Core Group). For reference, Figure 1 gives an overview of the different management bodies in IO-SEA[1].

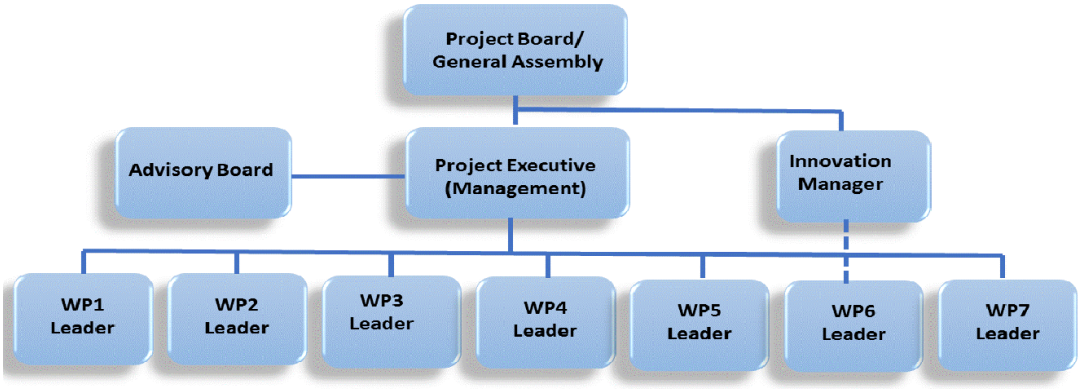

Figure 1: IO-SEA Management Structure

2.1 Project Board

The Project Board is the final decision taking body for the IO-SEA project (cf. [1] and [2] for more details). It comprises one representative from each partner and meets approximately every six months or upon request of the Project Coordinator or one of the partners[2] (cf. Table 2).

Since the writing of D7.2 in M24, three other Project Board meetings have taken place (cf. Table 1). The meetings in April 2023 and in October 2023 focused mainly on the amendments to the Description of Action, which were needed to adapt during the project lifetime (cf. Section 4.2). The main topic of the meeting in March 2024 was an internal assessment of our project and its results, as well as the preparation of the final project review.

2.2 Project Executive (or "Executive Board")

The Executive Board guides on an operational level the IO-SEA project. Its role is to oversee all quality assurance procedures and guidelines, to ensure the progress towards the technical objectives of the

[1] cf. WP6 reports for details on the Innovation Manager and its activities

[2] ParTec is a Linked Third Party to FZJ, but is regarded as "partner" to the IO-SEA project and also takes part in the Project Board.

Date	Location	Main Topics and Decisions
17/05/2021	Remote: Kick-off	Approval of the members of the Executive Board; Approval of the partner's contribution to joint communication costs
06/10/2021	Remote	Update on the project's progress
05/04/2022	Remote	Feedback from the M9 informal Check-up, some changes at Executive Board level to be approved
10/10/2022	Face-to-face in Paris	Update on the project's progress and M18 review preparation
26/04/2023	Face-to-face in Grenoble	Approval of the 2nd amendment to the Grant Agreement)
11/10/2023	Face-to-face in Ostrava	Preparation of the 3rd amendment (formal approval by e-mail after the meeting itself)
22/03/2024	Remote	Status on Review preparation, internal IO-SEA assessment

Table 1: Meetings of the IO-SEA Project Board

Partner	Project Board Representative	Project Board Proxy
CEA	Jacques-Charles Lafoucrière	Jean-Philippe Nominé
Atos-Bull	Cornel Crisan	Jean-Robert Bacou
FZJ	Wolfgang Frings	Eric Gregory
ECMWF	James Hawkes	Tiago Quinto
Seagate	John Forgan	Mark Wiggins
ICHEC	Buket Benek Gursoy	Venkatesh Kannan
IT4I	Jan Martinovic	Martin Golasowski
KTH	Stefano Markidis	Artur Podobas
CEITEC	Jirka Novacec	
JGU	André Brinkmann	Reza Salkhordehhaghighi
ParTec	Hugo Falter	Ina Schmitz

Table 2: Members of the Project Board (as of March 2023)

project, and to an effective cooperation amongst the project partners. On average, the Executive Board met once a month (cf. Table 4), with some flexibility and adjustments when the Cross-WP Sessions are getting more frequent or when All-Hands meetings are occurring. For example, upon the decision of the Executive Board, no Executive Board meetings took place in April and Mai of 2023, but in total nine Cross-WP Sessions took place between March and June 2023. During this time in the project, we were focusing on the Pilot and its deployment and used the Cross-WP Sessions for synchronising amongst the consortium partners.

Since M20, we have operated with an "Extended Executive Board", bringing together the Scientific Coordinator, the Project Management Office, the WP leaders, as well as one representative per partner (cf. Table 3). This setup has proven efficient, especially when coping with technical and with organisational adaptations. The scope of the "Extended Executive Board" is unchanged compared to the initial mission of the Executive Board as described in the DoA. Its topics are mainly:

- The activities within the different WPs,
- Cross-WP issues such as the IO-SEA Pilot, or collaboration related topics, and
- Topics related to the project management, such as the planning of All-Hands meetings, or the review of the risk and their mitigation actions.

Name	**Role**	**Proxy**
Philippe Deniel (CEA)	Scientific Coordinator	Maike Gilliot
Eric Gregory (JSC)	WP1 Leader	
Celine Lemarinier (Atos)	WP2 Leader	Philippe Couvée
Philippe Couvée (Atos)	WP3 Leader	Celine Lemarinier
Sebastien Gougeaud	WP4 Leader	
James Hawkes (ECMWF)	WP5 Leader	
Sai (Seagate)	WP6 Leader	
Maike Gilliot (CEA)	WP7 Leader	Philippe Deniel
Stefano Markidis (KTH)	KTH representative	
Stefan Krempel (ParTec)	ParTec representative	
Buket Benek Gursoy	ICHEC representative	Katie O'Connor
Martin Golasowski	IT4I representative	
André Brinkmann	JGU representative	Reza Salkhordehhaghighi
Jirka Novacec	CEITEC representative	

Table 3: Members of the Executive Board (as of March 2024)

2.3 Project Management Office

The Project Management Office (also referred to as "WP7 Core Group") runs the day-to-day activities of the project, under the control of the Executive Board. The WP7 Core Group consists of the

Meetings in 2021	Meetings in 2022	Meetings in 2023	Meetings in 2024
16/04	21/01	20/01	12/01
21/05	18/02	20/02	22/03
18/06	18/03	17/03	
28/07	15/04	29/06	
17/09	20/05	25/08	
15/10	17/06	10/11	
19/11	22/07	08/12	
17/12	16/09		
	16/12		

Table 4: Meetings of the IO-SEA Executive Board

Scientific Coordinator (Philippe Deniel, CEA), the WP7 Leader (Maike Gilliot, CEA), Jean-Robert Bacou (contributing to WP7, Eviden) and Cornel Crisan (contributing to WP7, Eviden).

The WP7 Core Group coordinates its activities in all tasks of WP7. Most effort in WP7 is provided by the WP7 Core Team. Its main activities are:

- It prepares the meetings and the minutes of the different bodies (Project Board, Executive Board, Advisory Board), including our All-Hands meetings. It plans, jointly with the Executive Board our Cross-WP Sessions (no preparation of minutes for our Cross-WP Sessions, but we collect the presentations and the recordings).
- It coordinates the preparation of the different elements for the reviews (presentations, demos, and reports).
- It makes sure that the project's internal quality assurance procedures are respected, such as controlling that the different deadlines of the review process are met.
- It acts as a liaising element towards the Project Officer and our partner projects and represents the IO-SEA project in CBCP meetings and the joint All-SEA-coordinator meetings.
- It supports WP6 in all communication matters which require coordination with our SEA-friends.

2.4 Advisory Board

IO-SEA's Advisory Board comprises people from other EuroHPC projects and qualified persons in the HPC/Data management domain. Its role is to give advice on the orientations and on the implementation of the IO-SEA project. It is composed of members from industry and academia, from Europe and from the US (cf. [4], [5] for more details on our Advisory Board). In 2022, Etienne Walter joined the Advisory Board, now its members are:

Figure 2: The members of the IO-SEA Advisory Board

- **Estela Suarez**, a senior scientist and leader of the DEEP series of EU-funded projects, she has driven the development of the Cluster-Booster and the Modular Supercomputing Architectures.
- **Johann Lombardi**, now one of the leading scientific advisors for the DAOS Foundation and also involved in Enakta Labs.
- **Jalil Boukhobza**, Professor at the ENSTA-Bretagne, with a focus on storage system design, performance evaluation and energy optimisation.
- **Robert Ross**, a Senior Computer Scientist at Argonne National Laboratory and the Director of the DOE SciDAC RAPIDS Institute for Computer Science, Data, and Artificial Intelligence.
- **Soraya Zertal** is an Associate Professor in Computing at the University of Paris Saclay-UVSQ and a member of Li-PaRAD lab, with research interests in data storage systems and performance evaluation.
- **Etienne Walter** is senior expert and project director in the Atos BDS (Big Data & Security) R&D Division and is currently coordinating, as General Manager, the phase 2 of the European Processor Initiative (EPI), as well as the RED-SEA project.

The Advisory Board met in May 2022, and again in October 2022, also involving all technical WP Leaders to allow for sound, technical discussions with our Advisory Board. The first meeting was mainly about presenting our approach for our I/O software stack. The discussions of the October meeting (21/10/2022) focused on the DASI, on what is to be done to ensure its uptake by the user communities, and on how DASI aligns with existing "interfaces," such as MPI-IO or HDFS. The Board also suggested looking into the ADIOS framework of Oak Ridge, which might be comparable to DASI

in some respects [6]. Regarding the Ephemeral Services, the members raised implementation-related questions, trying to see how the ephemeral services would be allocated and scheduled. This meeting helped to refine the architectural choices for the IO-SEA software stack and to understand which features to focus on, which influenced also the work on the IO-SEA Pilot system.

The collaboration with the advisors went beyond these formal meetings. For example, at SC21 and SC22, IO-SEA held Birds-of-a-feather sessions, with some of our advisors being part of our speakers and panellists, and joint publications were submitted. On top of this, the link with the projects DEEP-SEA (coordinated by E. Suarez), RED-SEA and EUPEX (both coordinated by E. Walter) was followed up upon via the SEA Coordinator meeting and our joint SEA events and activities.

2.5 Cross-WP Sessions

These regular cross-work package sessions were not foreseen in the initial proposal, but were set up in an ad-hoc fashion as the need emerged. Probably amplified by the lack of face-to-face meetings in the first year of the project, the need for regular exchanges between the work packages became clear during the architecture-definition phase.

This effort was kicked off by a full-day workshop and has now become a regular meeting (approximately every other Wednesday from 16:00–17:00, depending on the needs of the partners). Based on the need within the project, the members of the Executive Board decide on the topic for each session in advance and appoint for each session a leader (in most cases one of the WP leaders or the Scientific Coordinator takes this role). The full list of sessions and their topics is given in Table 5.

Today, these Cross-WP Sessions fulfil two purposes: First, they allow discussion and exchange on topics of common interest between the WPs. Second, they serve as "mini-all-hands meetings" for sharing widely information with all project participants. This simplifies communication and has proven very useful, for example for the review preparation and the work on our joint Pilot. They thus serve as a complement to the All-Hands meeting that takes place regularly every six months.

Overall, 20 Cross-WP Sessions took place in the third year of the IO-SEA project, with the main focus on our IO-SEA Pilot system.

Date	Title	Leading WPs
11/01/2023	How to access the IO-SEA Pilot System	WP1, WP2, WP3
25/01/2023	Io-SEA software packaging (as .rpm)	Scientific Coordinator
08/02/2023	Accessing the IO-SEA Pilot	Seagate
08/03/2023	Annexe Storage components in WP4: Phobos	WP4
22/03/2023	Collaboration with ADMIRE: The IO-Traces Initiative	JGU
03/05/2023	Pilot: how to integrate NFS-Ganesha	All WPs
10/05/2023	Pilot: update on V1.0	All WPs
24/05/2023	Pilot update: issues on DEEP	All WPs
31/05/2023	Update on Pilot and first use case runs	All WPs
07/06/2023	Update on use case runs on IO-SEA Pilot/DEEP system	All WPs
14/06/2023	Pilot update: WFM 1.1 deployed	All WPs
21/06/2023	Pilot and use case updates	All WPs
05/07/2023	Digging into use case issues on the Pilot system	All WPs
12/07/2023	Connecting Hestia with WFM V1.6: Implementation details	All WPs
26/07/2023	.rpm discussion	All WPs
09/08/2023	WFM V1.4 as reference version for D1.4	All WPs
27/09/2023	Pilot status and strategy beyond V1.6	All WPs
22/10/2023	As part of the All-hands meeting in Ostrava: Discussion: update of the deployment strategy for for the remaining months in IO-SEA	All WPs
08/11/2023	Upgrading to V1.4	All WPs
06/12/2023	Pilot: preparing upgrade to V1.6	All WPs
03/01/2024	Pilot updated and joint SEA demo	All WPs
24/01/2024	Joint demo including BXI partition (for RED-SEA)	All WPs
31/01/2024	Remote All-hands meeting: Discussion V2.0 and beyond	All WPs
21/02/2023	Upcoming deliverables and milestones: status	All WPs
13/03/2024	Upcoming review elements (including recommendation the section of the Final Project Report)	All WPs

Table 5: Overview of Cross-WP Sessions (2nd Reporting Periode)

3 Quality Control

Within IO-SEA, quality management focuses on the quality assessment of the work and the deliverables produced within the project. The process is overseen by the WP7 Core Group, but relies on all partners, all work package leaders, and all project participants.

The review process for deliverables foresees a timeline with intermediate steps to ensure that the final deliverables are submitted on time. Its main steps are:

1. The WP7 Core Team suggests internal reviewers for the deliverables to come, taking into account possible conflicts of interest, the technical background, and seeking for an overall balanced effort between all partners in the review effort. This selection is then amended and approved by the Executive Board.
2. The outline (table of contents) of the document is to be provided 2 months before the due date for submission.
3. 15 working days before the due date: a first complete version has to be submitted into IO-SEA's shared storage space and made available to the project-internal reviewers.
4. 10 working days before the due date: the project-internal reviewers have completed their reviews and submitted their comments to the author, who prepares a new version of the document, taking into account the reviewer's comments and suggestions. The document annotated by the reviewers has been uploaded to the shared storage space.
5. Five working days before the due date: the Scientific Coordinator and the WP7 Core Group review the document. The main author takes the comments and suggestions into account and prepares the final version of the document.
6. One to two working days before the due date: the WP7 Core Group submits the document to EuroHPC via the EC portal.

This process had proven useful and efficient in the first 2 years of the project, and has been applied in the same way in the third year of the project.

3.1 Deliverables M25-M36

The Table 3 provides an overview of the deliverables submitted or currently under preparation (with a deadline in M36). It shows the date we initially targeted when writing the proposal, as well as the reviewed submission dates.

Overall, for the period M25- M36, nine deliverables were planned for in the initial DoA. The deliverable, D1.4, was shifted to M30, thus with a total of 10 deliverables in the third year of the project. Six out of these ten deliverables have been submitted in time, the remaining four are - at the time of writing - under preparation (and due by the end of March 2024). The following deviation from the initial planning occurred:

- Deliverable D1.4 was scheduled for M24. It required the use cases to run on the IO-SEA Pilot. However, as the Pilot has only been ready with some delay (cf.Chapter 4), we requested to shift this deliverable to M30. This shift has been granted by the PO, and formalised via an amendment to the DoA.
- Due to the extended joint work and effort on the IO-SEA Pilot, all the final deliverables of the technical WPs have been shifted from M31 and M32 to M35, allowing for some more time to finish the work on the Pilot and progress on the developments. Again, this shift has been discussed with the PO upfront before being formalised via an amendment to the Grant Agreement.

Del. No.	WP	Title	Lead Ben.	Nature	Est. Del. Date	Rev. Due Date	Receipt Date	Status
D1.4	WP1	First application port and evaluation	FZJ	Other	24	30	30	Submitted
D2.2	WP2	Ephemeral Data Access Environment: First version of the solution	BULL	Report	22	24	24	Submitted
D3.2	WP3	Instrumentation and Monitoring: First version of the solution	BULL	Report	24	24	24	Submitted
D4.2	WP4	Hierarchical Storage Management Feature: First version of the validations	CEA	Report	24	24	24	Submitted
D5.3	WP5	Adaptations of Application workflows to the DASI	KTH	Report	24	24	24	Submitted
D6.2	WP6	Dissemination, Exploitation and Training Report and Future Plans Year 2	SEAGATE	Report	24	24	24	Submitted
D7.2	WP7	Periodic Report Year 2	CEA	ORDP	24	24	24	Submitted
D2.3	WP2	Ephemeral Data Access Environment: Final version of the solution	BULL	Report	31	35	35	Submitted
D3.3	WP3	Instrumentation and Monitoring: Final version of the validations	BULL	Report	31	35	35	Submitted
D4.3	WP4	Hierarchical Storage Management: : Final version of the validations	CEA	Other	31	35	35	Submitted
D5.4	WP5	Demonstrations of data-centric workflows using DASI for in-situ processing and visualisation	JGU MAINZ	Report	32	35	35	Submitted
D5.5	WP5	Public release of Data Access and Storage Interface	ECMWF	Other	32	35	35	Submitted
D1.5	WP1	Final report on applications experience	FZJ	Report	36	36		Under preparation
D5.6	WP5	Final Application workflow adaptations to the DASI	KTH	Report	36	36		Under preparation
D6.3	WP6	Dissemination, Exploitation and Training Report and Future Plans Year 3	SEAGATE	Report	36	36		Under preparation
D7.3	WP7	Periodic Report Year 3	CEA	ORDP	36	36		Under preparation

Figure 3: Overview of IO-SEA deliverables for the period M25 - M36

3.2 Milestones M25 - M36

As in Year 1 and Year 2 of the project, we have opted for a reduced number of milestones, since most of the software components developed in the technical work packages have strong dependencies. It thus makes sense to regroup most of the development effort under a single milestone. Table 4 gives an overview of the milestones for the periode M24- M36, indicating the intial due date (according to the intial DoA), as well as the updated due date (after the second amendment).

Milestone number	Milestone name	Related WPs	Due date	Updated due date	Means of verification
MS 4	IO-SEA Infrastructure is instrumented	WP2	M24	M24	D3.2
MS 7	IO-SEA: 1st version of the solution	WP1, WP2, WP3, WP4, WP5	M24	M24	D1.4, D2.1, D3.2, D4.2, D5.2
MS 10	Year 2 activities and reporting complete (Incl (a) Project Management activities & (b) Dissemination and exploitation and training activities)	WP6, WP7	M24	M24	D6.2, D7.2
MS 5	AI Based result available	WP2	M31	M35	D3.3
MS 8	IO-SEA: final version of the solution	WP1, WP2, WP3, WP4, WP5	M36	M36	D2.3, D3.3, D4.3, D5.3
MS 11	Year 3 activities and reporting complete (Incl (a) Project Management activities & (b) Dissemination and exploitation and training activities)	WP6, WP7	M36	M36	D6.3, D7.3

Figure 4: Overview of IO-SEA milestones for the second reporting period

As detailed in [1], we have opted for a light-weight milestone assessment process, focusing on three main elements for assessing whether a milestone is achieved:

- Have the milestone's objectives been achieved?
- Have the documents and deliverables related to the milestone been reviewed according to the IO-SEA internal rules and submitted to the EC in time?
- Have the milestone results been presented to the IO-SEA consortium?

The possible assessments are "YES", "NO" or "PARTIALLY". The achievement of the milestones is assessed by the Executive Board. The milestones for the period M25 - M36 are the following:

- **MS 4, MS 7**, and **MS 10** have been presented in D7.2 [2].

- **MS 5** describing the AI-based solution and its results and outcomes for smart data placement. This milestone is specific to WP3. It was initially planned for M31 but shifted to M35 - again to allow for some more time and to compensate for the effort needed to bring up the Pilot system. This milestone was met.
- **MS 11**, which closes the Year 3 activities of WP6 and WP7 was also due in M36. The Executive Board approved this milestone in its last Executive Board meeting on 22/03/2024.
- **MS 8** is related to the technical WPs (WP1-WP5) and marks the final version of the IO-SEA software stack. This milestone is due in M36. The Executive Board approved this milestone in its last Executive Board meeting on 22/03/2024.

4 Risk Management

Within the IO-SEA project, the Executive Board oversees the management of possible risks. The starting point of the risk assessment is the list of risks with their associated risk-mitigation measures, as presented in the DoA. Regularly, the Executive Board re-assesses the risks identified so far and adds newly identified risks to the list. The full list of risks is part of the Final Project Report. This section will focus on the main risks we identified and mitigated in Year 3.

4.1 Pending Risk: Covid Pandemic

The first risk detailed in M12 and M24 is linked to the **Covid pandemic** and thus reduced opportunities for face-to-face events with possible impact on (1) the project management for team-building and for making the partners engage in the project, and (2) for our communication and dissemination activities.

This risk did not materialise. From M24 onwards, face-to-face events became again more frequent. As of today, it seems that overall some reluctance to travel remains due to its overhead with respect to time and cost. In addition, also environmental concerns are raised more often. We have thus developed a mixture of remote and in-person events in the past years, that worked out well for our project.

- We held face-to-face All-Hands meeting in October 2022 (Paris), in April 2023 (Grenoble, hosted by Eviden), and in October 2023 (Ostrava, hosted by IT4I). All other meetings were held remotely to allow a wide audience to attend.
- The final joint SEA-workshop took place in January 2024 at LRZ as an in-person event too. By co-location of this event with the HiPEAC conference, we minimised the overhead related to travel activities.
- For communication, outreach and dissemination, we opted for a mix of online/offline events: we were present at many conferences and events, but offered many of our training sessions online (cf. [7] for details).

4.2 Pending risk: Deviation from the planned Use of Resources

The analysis in M12, M18 and M24 indicated a considerable underspending with respect to the planned effort for some of the partners. The main issues for the reduced effort in the project are mainly the following:

- The "deviation" is assessed by assuming a linear use of resources and a linear workload during the project lifetime. This is a simplified assumption, that does not hold in all cases. In particular, in WP5 some of the tasks (and thus the related effort) are planned for the second part of the project.

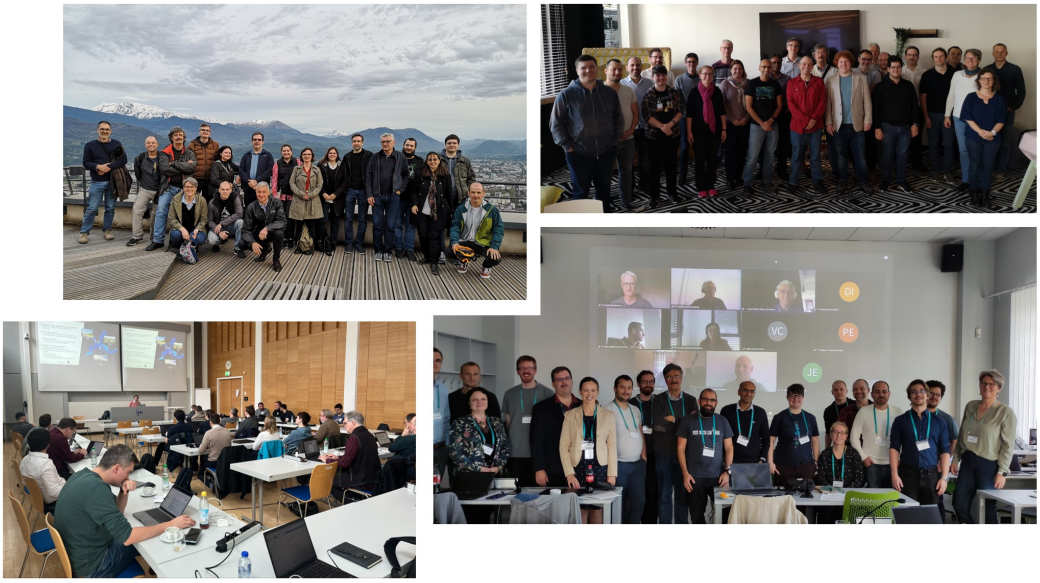

Figure 5: Face-to-face meetings for IO-SEA

- Some partners faced some issues in hiring at the beginning of the project. The generally difficult situation in hiring has been accentuated by the Covid pandemic. As a consequence, some effort was made later than planned.

Despite the lower effort reported in the first months of the project, none of the WP Leaders reported on missing engagement from the partners. On top of this, all partner sites reported on actions to speed up the hiring of new employees to ramp up the expected effort by the end of the project.

Now, by the end of the project, we see that indeed the planned effort (in PM) has been provided by the partners. We observe some underspending with respect to the planned cost for the action, but this is not due to less engagement or effort from the partners. Details on this matter are given in Section 5. It must be stressed that all figures and the analysis given in this deliverable are based on *preliminary* figures. At the time of writing of this deliverable, the final numbers are not available. These numbers will be updated and refined for the Final Project Report, amended with additional explanations and justifications where needed.

4.3 Pending risk: Discontinuation of the NVRAM technology

The NVRAM technology, which is at the heart of T2.2 is discontinued. It would thus not make much sense to conduct the task as initially planned. We suggest replacing it with an equivalent task, as we see strong interest in the community and industry for very fast non-volatile storage technologies, as testifies the CXL (Compute Express Link) initiative [8].

During the M18 review, we suggested focusing this task on an extensive exploratory study of “Disaggregated Memory and I/O Architecture for Data Nodes”, looking into the usage of disaggregated

memory and I/O architectures for data nodes on which ephemeral I/O and storage services can be deployed, and evaluating the CXL protocol.

This has given rise to an amendment which was approved by EuroHPC in August 2023. The suggested changes have been implemented and the work has been conducted as planned by the amendment.

4.4 Pending risk: Lack of a common IO-SEA Pilot System

The DoA did not foresee any physical HPC infrastructure for integrating and testing the IO-SEA software stack. However, in the first months of the project the lack of such a common Pilot System has been identified as a risk, (1) possibly hindering seamless integration into a solid and coherent IO-SEA storage solution, and (2) lacking vehicle for the use cases to test and to assess the IO-SEA software stack in a reproducible and comparable way. In [5] different strategies have been described to mitigate the risk:

1. Single VMs: As image for deployment in a Virtual Machine (VM) has been configured, so that all project participants can execute their developments in a shared and common environment.
2. A Cluster of VMs on IT4I's OpenStack cluster, with dedicated data nodes where the image of the IO-SEA-VM can be launched.
3. DEEP System: Access to the DEEP-SEA system in Jülich, mainly for instrumenting the use cases and for benchmarking
4. IO-SEA Pilot: The hardware of the SAGE2 prototype is being re-purposed to serve as IO-SEA data nodes and to connect the IO-SEA data nodes to the DEEP system [9]. Thus, the applications can then use the compute nodes of DEEP for their computation, whereas the IO-SEA pilots host all IO and data-related services and nodes.

The first three options were operational rather soon in 2022. Their limitations, however, became quickly obvious. These were mainly limitations related to the software dependencies between the WPs/teams, and in particular the problem of keeping the interfaces coherent and the dependencies consistent.

The project thus engaged in refurbishing the Pilot of the SAGE project and using it for IO-SEA's purposes [9]. Overall, this required considerable unplanned effort. The first version of our Pilot was available in April 2023, and in January 2024 the final version V1.6 has been reached. The work on the Pilot has led to the delay of D1.4 (as the use cases needed the Pilot system for running their first runs), as well as of the deliverables D2.3, D3.3, D4.3, D5.3, and D5.4, as well as of the corresponding milestone MS 8 (shifted from M31 - M36). We are convinced that the work on the Pilot has contributed significantly to the quality of the software delivered. We have shown its integration and thanks to the Pilot system the use cases were able to run comparable and reproducible runs using the IO-SEA storage approach. Thus, this common Pilot contributes considerably to our assessment of the IO-SEA software stack by the use cases.

4.5 Pending risk: Changed Business Strategy at Seagate

We were informed in the summer of 2023 that the entire EU R&D team of Seagate is affected by their organisational cost-cutting and headcount reduction measures due to severe economic pressures in the last couple of financial quarters. This resulted in Seagate's IO-SEA contributors being let go from Seagate.

WP7 has immediately initiated a dialogue with Seagate as to understand their position and the role they wish to play in the project for the remaining months. Seagate's main goal was to avoid a critical impact on the project. Nevertheless, Seagate has clearly indicated that they will not be able to perform all pending technical tasks. Mainly affected by the changes on the Seagate side are WP6 (Dissemination), T2.2 and T5.3. This gave rise to a set of discussions with all project partners to come up with solutions to mitigate the "loss" of Seagate's contribution. This gave rise to an amendment of the Grant Agreement, which is currently being finalised. The main changes are the following:

- **Task 2.2:** Seagate is the main contributor to T2.2 (study on "Disaggregated Memory and I/O Architecture for Data Nodes"), with some support from Atos/Eviden for experiments on real hardware. Seagate started (with Eviden's support) the work on CXL evaluation and emulation. The work done until Q1 2023 is described in D2.2. Eviden cannot take over the work still to be performed by Seagate, but it will perform its part of the study and the work as planned. The outcome will be described as part of D2.3.
- **Task 5.3:** Seagate's Cortx-Motr technology was the main storage technology to be exploited in T5.3. Cortx-Motr will continue to exist in the public domain, as open-source software, but many of the technological goals are unachievable without Seagate's support in this project. ECMWF has succeeded in building a Cortx-Motr backend for DASI, but recognised that this is less desirable than planned. Therefore, ECMWF will additionally develop DASI with the same backend interface as the ECMWF FDB software, which will allow using multiple backends with DASI, including POSIX, Ceph, DAOS and NVRAM. The POSIX backend will additionally be integrated with the workflow manager and data nodes to provide a valuable storage backend on which to demonstrate the two in-situ processing usecases of T5.3. JGU will also provide a backend to GekkoFS, linking with the ADMIRE project, as an alternative to Cortx-Motr. A new task text is provided to reflect these change.
- **WP6 Leadership and Effort:** Seagate is leading the work package 6 in collaboration with KTH, and leading Task 6.2 (Exploitation and Innovation) jointly with Eviden. Seagate also is the lead beneficiary for D6.3 - Dissemination, Exploitation and Training Report and Future Plans Year 3, due at M36. KTH will take over the leadership of the WP from Seagate and become the lead beneficiary for D6.3. Also, ParTec hired Sai Narasimhamurthy, who was leading WP6 from the Seagate side from M1 - M24. With Sai Narasimhamurthy now being part of ParTec, ParTec intends to spend more effort in WP6 than originally planned and will be assisting KTH in WP6 activities.

These changes have been agreed on by the IO-SEA Project Board in November 2023 and have been implemented as such. From a formal point of view, these changes gave rise to an amendment, which is currently being finalised.

5 Financial Management

The WP7 Core Group also traces the actual effort declared on the project, thus monitoring the deviations and the evolution of the deviations over time, as part of T7.4. At the beginning of the project, we observed - as detailed during the M18 review - some severe deviations. Whereas on a project-wide level, the use of effort per WP and per partner was in line with the expected numbers, there were at M18 severe discrepancies for some WPs and some of the partners. This section shows some *preliminary figures* covering the period M1 - M36. Detailed M36 figures with explanations and justifications in case of stronger deviations will be provided as part of the Final Project Report. Moreover, for some partners, the expected effort (in PM) is about to change with the current amendment. This is not reflected in the current analysis either (as the amendment is formally speaking not yet fully approved). These elements will be detailed in Section 5 of the Final Project Report.

5.1 Efforts in PM per partner: Preliminary Analysis

As indicated earlier, we observed in M18 and M24 deviations for some partners in some work packages. Overall, the reasons can be roughly summarised as follows:

- **Issues in hiring:** Some partners reported delays in hiring. The generally difficult situation for hiring, especially in the public sector, has been even more difficult with the Covid pandemic. The partners have taken measures (such as hiring head-hunters) to overcome the situation and reported that from 01/2023 onwards all positions will be filled.
- **Non-linear effort within the project:** As pointed out, we assume a linear use of the resources over the project lifetime. The input from the partners shows that this hypothesis is not always correct, for example, some training activities which are requiring more effort in the second half of the project.
- **Deviations due to unplanned activities:** As all R&I projects, IO-SEA also faced its share of unexpected situations, whose resolutions required effort, which was not planned for, for example for Eviden in WP2 and KTH for some of its WP5 related work. Detailed justifications will be provided as part of the Final Project Report.

Despite the lower effort reported, none of the WP leaders reported on lacking engagement or contributions from partners. Moreover, all partners reconfirmed that the final use of resources on their side will be in line with the overall planned effort. Now, in M36, we observe that the actual use of resources has indeed balanced out. The figures given in this report are not final, and in particular, the last changes planned by the ongoing amendment are not yet integrated into the expected consumption. All of this will be addressed in the Final Project Report.

5.2 M36: Preliminary Analysis of the Effort by Partner

As shown in Figure 6, a total of 791 PM have been provided for the IO-SEA Project, which is in line with the expected effort (826 PM planned initially by the DoA). For most partners, the deviations are around 2%, with some exceptions (cf. also Figure 7:

Partner	Expected consumption (PM)	Consumed (PM)	Deviation (%)
CEA	130,00	114,88	-11,63
Atos-Bull	165,00	182,29	10,48
FZJ	112,00	115,11	2,78
ECMWF	65,00	66,92	2,95
Seagate	96,00	71,65	-25,37
ICHEC	36,00	30,00	-16,67
IT4I	82,00	90,60	10,49
KTH	36,00	36,33	0,92
CEITEC	20,00	20,42	2,08
JGU	48,00	43,25	-9,90
ParTec	36,00	20,23	-43,79
SUM	**826,00**	**791,67**	**-4,16**

Figure 6: Effort - PM per partner (as of M36 - preliminary figures)

- A considerable "underspending" by Seagate: As detailed in Chapter 4, the current amendment revises the effort and the contribution of Seagate to the IO-SEA project. This reduced implication is not yet taken into account here.
- Eviden and IT4I have both declared some more effort than planned, which will be explained and justified in the Final Project Report.
- An underspending around 10 - 25 % by CEA, JCG, Partec and ICHEC related to the fact that the figures reported so far are still missing January - Mars 2024. Some more contributions are expected to be reported.

5.3 M36: Preliminary Analysis of the Effort by WP

As shown in Figure 8, overall the effort per WP has been respected. The major deviation we observe at this point is related to WP5, with 20% less effort than initially planned. This is caused by the latest changes to WP5 and its scope (with less effort for technical work from Seagate on these tasks). The graphical representation of the effort per WP is given in Figure 9.

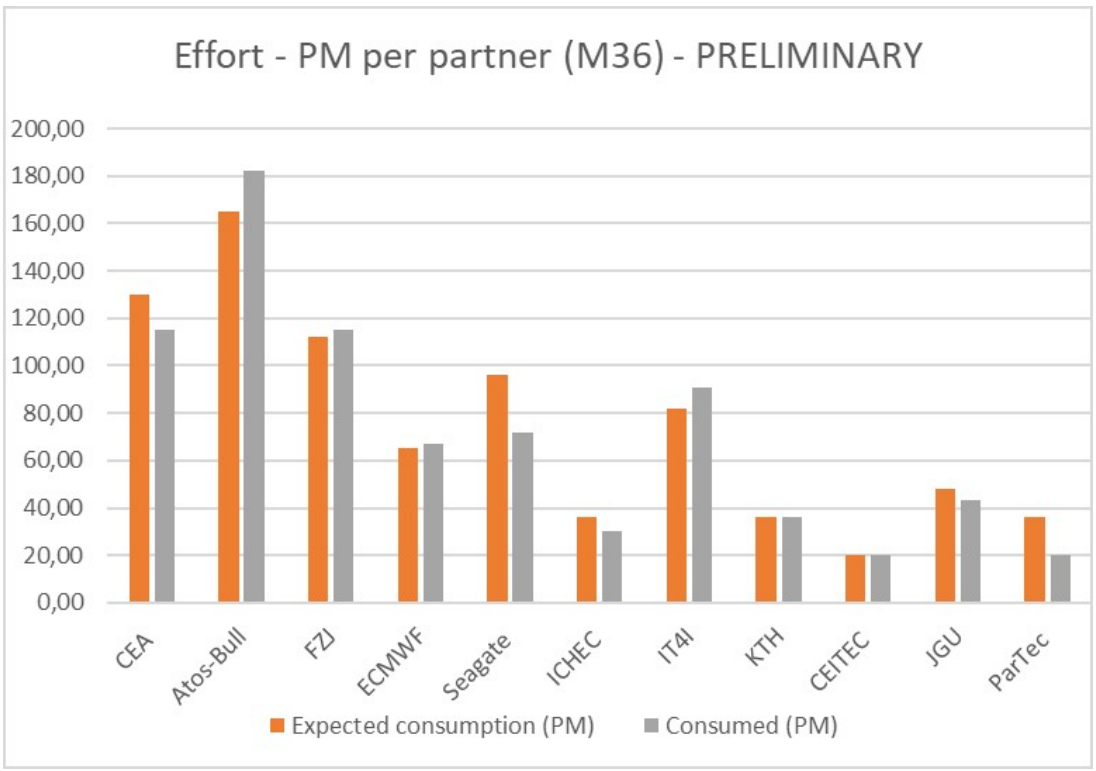

Figure 7: Effort - PM per partner (as of M36 - preliminary figures)

The first line of Figure 8 shows that the total effort of 791 PM sums up to a total of 5,1 MioEuro of Direct Personnel Cost compared to the planned 7 MioEuro of Direct Personnel Cost. It must be noted that the project will globally end with some underspending. The setup based on co-financing of European and national funds makes it even more difficult and complex to adapt the planned and allocated resources to specific activities or to shift the budget between partners and tasks. For the Final Project Report, we will dig into more detail on where this underspending occurred.

WP	Expected consumption PM)	Consumed (PM)	Deviation (%)
TOTAL COST	**7 062 252**	**5 193 185**	**-26,47**
WP1	165,00	169,26	2,58
WP2	139,00	139,43	0,31
WP3	136,00	150,09	10,36
WP4	157,00	137,94	-12,14
WP5	158,00	129,00	-18,35
WP6	35,00	34,33	-1,93
WP7	36,00	31,63	-12,14
TOTAL	**826,00**	**791,67**	**-4,16**

Figure 8: Effort - PM per WP (as of M36 - preliminary figures)

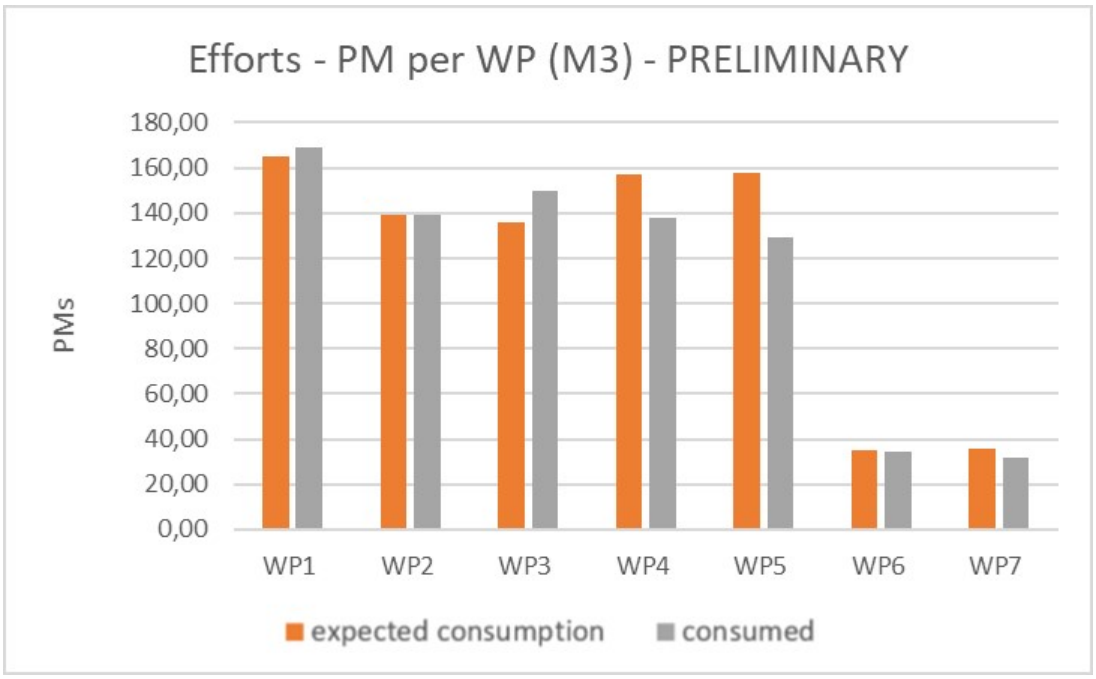

Figure 9: Effort - PM per WP (as of M36 - preliminary figures)

6 Common task for complementary grants

The role of this task was to formally install, kick-off and coordinate the collaboration with the other SEA-projects and the other projects funded under the EuroHPC-19 call. A milestone of this task was the signature of the **Complementary Grant Agreement**, which was completed in M12.

The signature of the Complementary Grant Agreement closed this task officially. In practice, the collaborations continue:

- The joint SEA-dissemination meetings, as well as all joint dissemination and outreach activities are followed up by WP6. The activities with respect to this are described in D6.3 [7].
- The Scientific Coordinator and the leader of WP7 take part in the monthly SEA-coordinator calls (cf. Section 2).
- The Scientific Coordinator and the leader of WP7 also take part in the regular calls of the SEA-coordinator with the PO. They act as liaison elements towards the funding agency to keep us mutually informed, to identify synergies and to assess the current status of our different collaborations (cf. Section 6).
- The Scientific Coordinator and the leader of WP7 also act as entry point to the IO-SEA project for our EuroHPC-19 partners: taking part in the Cross-Project Collaboration Board (CPCB) meetings and contributing to the joint effort in the different work streams and our common outreach and dissemination activities (such as the joint session during the EuroHPC Summit in March 2024 in Antwerp).
- Our technical collaboration with DEEP-SEA, RED-SEA and ADMIRE is described in the deliverables of the technical WPs (WP1-WP5). In addition, D6.3 [7] provides in annexe elements documenting the synergies and the collaborations with our partners.

7 Summary

WP7 of IO-SEA is in charge of the scientific coordination and management of the IO-SEA project. This report presents the work performed by WP7 from April 2023 until March 2024 (M25-M36), building on the elements provided in D7.1 [1] (issued in M12), the first Interim Report (issued in M18), as well as D7.2 [2].

For this third year of the IO-SEA project, WP7 managed and operated the IO-SEA project with the tools, methods and internal procedures that had been set up at the beginning of the project and that have proven useful. The management bodies and in particular the Extended Executive Board and the Cross-WP sessions worked well and contributed to the overall success of the project. The technical work performed is in line with the scientific expectations as described in the DoA. Some adaptations to the DoA have been necessary. These amendments have been conducted jointly with the PO and the IO-SEA Project Board. Overall, the project's effort spending is in line with the foreseen effort. Risks - such as the lack of a common Pilot system - have been addressed and mitigation actions have been set up. A special attention in the final phase of the project was drawn to our dissemination, training and exploitation activities. While this is the focus of WP6, WP7 provided support and acted as a liaison element with our SEA-friends.

Overall, WP7 provided organisational support to the other WPs for conducting their technical and scientific work in a structured and well-organised manner, in line with the DoA and the Consortium Agreement. WP7 ensured the overall coherence of IO-SEA's effort and liaised efficiently with our funding authority, as well as our SEA-friends.

List of Acronyms and Abbreviations

H

HDFS Hierarchical Data Format (HDF) is a set of file formats (HDF4, HDF5) designed to store and organize large amounts of data .

I

IT4I IT4Innovations National Supercomputing Centre at VSB Technical University of Ostrava, Czech Republic.

M

MPI-IO Message Passing I/O.

S

SAGE2 R-and-I project funded under H2020, whose outcome is used and developed further in IO-SEA.

V

VM Virtual Machine.

Bibliography

[1] M. Gilliot. IO-SEA D7.1: Periodic report for wp7 - year 1. Technical report, IO-Software for Exascale Architectures, march 2022.

[2] M. Gilliot. IO-SEA D7.2: Periodic report for wp7 - year 2. Technical report, IO-Software for Exascale Architectures, march 2023.

[3] Management by Exception. https://en.wikipedia.org/wiki/Management_by_exception.

[4] IO-SEA Advisory Board. https://iosea-project.eu/advisory-board/.

[5] E. B. Gregory, P. Couvée, and M. Golasowski. IO-SEA D1.1 Application and co-design input. Technical report, IO-Software for Exascale Architectures, jul 2021.

[6] Oak Ridge The Adaptable IO System (ADIOS). https://www.olcf.ornl.gov/center-projects/adios/.

[7] S Medeiros, D; Markidis. IO-SEA D6.3: Periodic report for wp6 - year 3. Technical report, IO-Software for Exascale Architectures, march 2023.

[8] Compute Express Link. www.computeexpresslink.org.

[9] SAGE2. https://cordis.europa.eu/project/id/800999/fr.