



EuroHPC-01-2019



IO-SEA

**IO - Software for Exascale Architectures
Grant Agreement Number: 955811**

D6.2

**Dissemination, Exploitation and Training Report and Future Plans
Year 2**

Final

Version: 1.0
Author(s): Sai Narasimhamurthy (Seagate)
Contributor(s): Jean-Robert Bacou/Philippe Couvee (ATOS), Katerina
Slaninova(IT4I)
Date: 31.03.2023

Project and Deliverable Information Sheet

IO-SEA Project	Project Ref. №: 955811	
	Project Title: IO - Software for Exascale Architectures	
	Project Web Site: http://www.iosea-project.eu	
	Deliverable ID: D6.2	
	Deliverable Nature: Report	
	Deliverable Level: PU	Contractual Date of Delivery: 31 / 03 / 2023
		Actual Date of Delivery: 31 / 03 / 2023
EC Project Officer: Daniel Opalka		

* - The dissemination levels are indicated as follows: PU = Public, fully open, e.g. web; CO = Confidential, restricted under conditions set out in Model Grant Agreement; CI = Classified, information as referred to in Commission Decision 2001/844/EC.

Document Control Sheet

Document	Title: Dissemination, Exploitation and Training Report and Future Plans Year 1	
	ID: D6.1	
	Version: 0.4	Status: Draft
	Available at: http://www.iosea-project.eu	
	Software Tool: Microsoft Word	
	File(s): IO-SEA_Deliverable_WP6_D6.2v0.2-COPY.docx	
Authorship	Written by:	Sai Narasimhamurthy (Seagate)
	Contributors:	Jean-Robert Bacou/Philippe Couvee (ATOS), Katerina Slaninova(IT4I)
	Reviewed by:	Eric Gregory (Juelich Supercomputing Center), Jenny Wong (ECMWF)
	Approved by:	Project Co-ordinator

Document Status Sheet

Version	Date	Status	Comments
0.1	02/02/2023	Draft	Table of Contents
0.2	02/10/2023	Draft	Initial draft
0.3	03/03/2023	Draft	Draft for Internal Review
1.0	31/03/2023	Draft	Incorporate internal reviewer comments

Document Keywords

Keywords:	IO-SEA, HPC, Exascale, Software
------------------	---------------------------------

Copyright notice:

© 2021-2024 IO-SEA Consortium Partners. All rights reserved. This document is a project document of the IO-SEA project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the IO-SEA partners, except as mandated by the European Commission contract 955811 for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

Table of Contents

Executive Summary	7
1 Introduction.....	8
2 Overall Strategy for Generating Impact.....	10
2.1 Expected impacts with the IO-SEA project.....	10
2.2 Strategy and Plans.....	11
2.3 Definitions of KPIs for successful dissemination.....	17
2.4 Alignment with Evolving European Objectives.....	18
3 Dissemination, Exploitation and Collaboration Activities	19
3.1 Dissemination, Exploitation and Collaboration Activities	19
3.2 Collaborations.....	34
4 Exploitation	36
5 Training.....	39
6 Conclusion.....	41
7 References	42
8 List of Acronyms and Abbreviations	43

List of Figures

Figure 1 Progression towards European Exascale with IO-SEA (Past/Present/Future)	11
Figure 2 Website Statistics (Year 1).....	28
Figure 3 Website Statistics (Year 2).....	28
Figure 5 IO-SEA Twitter Page.....	29
Figure 6 Twitter impressions showing a spike of activity around the time of the SC 22 conference. (Quarter 7 of the project.).....	30
Figure 7 IO-SEA LinkedIn page	31
Figure 8 IO-SEA Dissemination Snapshots.....	32
Figure 9 Example giveaway - Portable Storage Device at the SEA projects booth at ISC'22.....	33
Figure 10 European Exascale projects website.....	33
Figure 10 IO-SEA Poster with a common SEA theme	35
Figure 12 IO-SEA github page	37

List of Tables

Table 1 Dissemination Strategy	13
Table 2 Collaboration Strategy.....	15
Table 3 Exploitation Strategy.....	16
Table 4 Training Strategy	17
Table 5 Key Performance Indicators	18
Table 6 Dissemination Activities	24
Table 7 Publications/Newsletters.....	26
Table 8 IO-SEA Press Releases.....	27
Table 9 Twitter impressions by quarter.....	30
Table 10 Organised training activities (M24).....	40
Table 11 Organised training activities – attendance (M24)	40
Table 12 Training activities – updated plan (M24)	41

Executive Summary

The goal of this deliverable is to provide an update for Year 2 on the Dissemination, Exploitation, Collaboration and Training activities for IO-SEA, updating the work done in Year 1, and to discuss the future plans for the course of the project. Please note that we have used the acronym D&E throughout the document to include all these activities.

1 Introduction

At the end of Year 2 of the IO-SEA project, we reiterate the main objectives of Dissemination, Exploitation and Training Tasks (6.1, 6.2 and 6.3 as per the IO-SEA “Description of Action”) and provide the current Status Summary of the work.

Dissemination: *The primary goal of dissemination is to reach wider audiences within HPC/AI communities in specific, and the computing community generally, to impart the outcomes and to provide the community an understanding of the methods, tools and techniques developed within the project. We also aim to cover communities specific to the different use cases used in the project.*

This work starts from the very beginning of the project and ramps up as the project progresses.

Status Summary: We have covered a broad range of audiences in the areas of HPC/AI and related communities based on all the activities done until the end of Year 2. The audiences included HPC administrators, application experts, technology experts, HPC policy makers and last but not the least, students who are interested to pursue their careers in the upcoming areas of HPC in Europe. Details of the dissemination are covered in Section 3.1. The dissemination has specifically focused on showing the connections between the SEA projects and how they are working together very closely to achieve Europe’s ambitions in the area of European HPC executed through EuroHPC. The goal has been to show how the SEA projects have developed the necessary basic building blocks towards Modular Supercomputing Architectures (MSAs), with IO-SEA focusing on storage, DEEP-SEA focusing on compute, and RED SEA focusing on the network.

We are also disseminating into complementary areas, such as Big Data/AI and cloud-based computing on very early feedback we got from the experts.

Collaboration: *The primary goal of this is to identify synergies with other projects and initiatives and work with them to achieve European HPC/AI objectives.*

This work starts from the beginning of the project.

Status Summary: We have very closely collaborated with the SEA projects throughout the second half of Year 1 and the whole of Year 2, project in putting forth a unified vision of the SEA projects. We believe we have been quite successful at achieving that aim, thanks to the early feedback that we got from the experts during Year 1 of the project.

We have also established collaborations outside of the SEA projects and the individual organisations have also used their own collaboration channels. We will discuss that in section 3.2.

Exploitation: The primary goal of this is to get market adoption of technologies developed in the project and gain usage for them within the HPC/AI and other related communities.

This activity is more back loaded within the project as exploitation objectives are targeted once concrete outcomes are available.

Status Summary: This work is now going in earnest. We have strong exploitation pathways towards European Exascale ambitions through SEA projects linkages that have been

strengthened with the EUPEX¹ EU project. We discuss it more in Section 4. The individual partners are also working on their respective exploitation pathways.

Training: *The primary goal of training is to educate the community on the usage of the technologies developed within the project.*

We expect some of these tasks to also contribute towards de-facto standards in some areas of technology that IO-SEA is relevant to.

Status Summary: We have had a clear training plan from the beginning of the project and have been executing it very satisfactorily, as we will describe in Section 5, evidenced by audience participation. The recordings of the training, and the Youtube channel associated with the SEA projects that we have created have been fully exploited. We have used a mix of face-to-face, hybrid, and online training events. We have focused on technologies that are starting to be mature for use by the community.

The rest of this document is organized as follows: Section 2 details the main impacts desired from the project, strategic objectives in achieving those impacts, and how our progress toward them measures up at the end of Year 2. We also define some Key Performance Indicators (KPIs) for success based on the feedback we have received from the experts. We also highlight evolving European objectives towards the end of Year 2. Section 3 details all the D&E activities. Section 4 concludes the report.

The intended audience for this deliverable is the public/community who is interested in knowing more about efforts in the areas embarked upon by IO-SEA.

¹ <https://eupex.eu/>

2 Overall Strategy for Generating Impact

We discuss in this chapter the overall strategy for generating impact through D&E activities and how we measure up to it towards the end of Year 2.

2.1 Expected impacts with the IO-SEA project

We reiterate below the primary expected impacts of IO-SEA planned at the beginning of the project:

- Contribution to the realisation of the specific and general objectives of EuroHPC²
- Strengthening scientific leadership as well as the competitiveness and innovation potential of European industry
- Contributing to a sustainable Exascale HPC supply ecosystem in Europe and ensuring European technological autonomy in this field
- Leveraging the efforts on the European low-power processing technologies (in particular the European Processor Initiative) and contributing to the realisation of future exascale system architectures based on such technologies
- Creation and promotion of European intellectual property, maturity of solutions and potential for commercial exploitation in future European exascale HPC systems
- Creating impact for many scientific communities (weather forecasting, bio-informatics, etc.) by addressing these highly demanding, data-intensive applications through data management tools and methods
- Delivering valuable data management technologies to not just HPC, but also the AI, cloud ecosystems and other related areas
- Creating market-exploitation pathways for commercial partners within the project

Progress toward expected impacts at the end of Year 2:

We continue to progress toward the mission of EuroHPC. Objectives. The tools that are now implemented in the different work packages (Ephemeral Services implementations in WP2, I/O instrumentation framework in WP3, the HSM & Phobos developments in WP4 and the Data Access and Storage Application Interface, or, “DASI” API in WP5) are ideally positioned towards **development, deployment, extension and maintenance** of world-leading, connected supercomputing and data infrastructure. For example, the Rados Gateway (RGW) S3³ API, used on top of the object store, helps in the federation of data infrastructures and hence connected supercomputing. The tools developed help to manage the immense deluge of data that occurs in large connected supercomputing infrastructures. The IO-SEA work directly flows into the EUPEX Pilot, which contributes to EuroHPC’s EPI (European Processor Initiative) strategy. EUPEX is one of Europe’s key programs targeting future European exascale systems. EUPEX is also ideally targeted towards Europe’s second exascale system, the first one being JUPITER⁴, which will be deployed at Juelich Supercomputing Centre. The second exascale system will exploit European HPC technology, to which IO-SEA and EUPEX will contribute. The tools and methods developed, especially the first version of the DASI API, are well-positioned to impact data-intensive HPC applications. The first versions of these tools implementations are now available. The APIs such as RGW S3 indeed help to make similar object storage infrastructure useful for both HPC and the cloud. Seagate, ATOS and ParTec,

² <https://eurohpc-ju.europa.eu/discover-eurohpc-ju>

³ <https://docs.ceph.com/en/latest/cephadm/services/rgw/>

⁴ https://eurohpc-ju.europa.eu/new-call-tender-procurement-jupiter-exascale-supercomputer-2023-01-17_en

the industrial partners, are already working on their market exploitation plans, as will be highlighted in Section 4.

With Dissemination, the story we wish to highlight is how initiatives starting all the way back to 2015 are starting to converge to European Exascale through IO-SEA and its sister projects, DEEP-SEA and RED-SEA.




Figure 1 Progression towards European Exascale with IO-SEA (Past/Present/Future)

The D&E work aims to address these expected impacts through the following focus areas:

- **Community Adoption:** Make sure that the technologies developed in the IO-SEA project will be adopted by the community.
- **HPC Strategy Alignment:** Place the project technologies at the heart of the European HPC strategy.
- **Industry and Research Support:** Support the European industry and research in its goal of developing an autonomous technological HPC environment
- **Applications' expertise:** Improve European competitiveness in application areas through a highly relevant data-centric approach.
- **Commercial Solutions:** Introduce and promote commercial solutions built on project technologies.
- **Exascale platforms:** Increase the return on investments made in PRACE⁵ supercomputers and in EC supported efforts towards Exascale platforms.
- **Information sharing:** Spread information on the project advantages (reduction of power consumption, efficiency, ability to build a data management system for an exascale platform) to the European public

Next, we compare our progress against our goals in these focus areas.

2.2 Strategy and Plans

2.2.1 Dissemination

Focus Areas	Original Plan	Status as of Y2
-------------	---------------	-----------------

⁵ <https://prace-ri.eu/>

Community Adoption	The plan is to increase participation in trade shows, conferences and major venues for presenting IO-SEA Technologies. This will be a forerunner for new collaborations, training events and, eventually, community adoption and participation in growing the community around IO-SEA technologies. The project will provide open source for almost all of the work and this will be a critical central piece. We will also provide metrics to track, on a periodic basis, the community adoption of foundational platform components such as CORTX Motr ⁶ . The website and social media efforts will also be geared towards this effort.	We have increased participation in these events, as will be discussed in Section 3, leveraging also the joint efforts of the other SEA projects. We now have an IO-SEA github ⁷ portal that provides community access to all the tools and methods developed in the project. We will also participate in events such as Big Data Value Forum ⁸ targeting the DAIRO ⁹ community.
HPC Strategy Alignment	This will be achieved by hosting joint workshops with other EuroHPC projects and of course the other SEA projects. This type of dissemination in international venues such as SC (Super Computing) will also increase outreach for European HPC technologies as a whole.	We are extremely well-aligned with the other SEA projects in hosting joint events and disseminations, as will be discussed in the next section. In putting together a unified “SEA Projects” front to the HPC community we provide the necessary building blocks for MSAs.
Industry and Research Support	Industry and Research organisations will share the IO-SEA developments with their respective communities, bolstering support for IO-SEA solutions.	The industrial partners have disseminated the work into their communities using their open source channels (for example, Seagate using CORTX Motr). The Research organisations have also promoted the work with their respective communities (for example, ECMWF has shared DASI with the weather and climate communities)
Applications ⁷	There will be heavily focus on	There have been multiple

⁶ <https://github.com/Seagate/cortx-motr>

⁷ <https://github.com/io-sea>

⁸ <https://european-big-data-value-forum.eu/>

⁹ <https://www.bdva.eu/DAIRO>

Expertise	applications discussing their experiences as part of the dissemination – and joint workshops between EuroHPC and SEA projects will focus on cross co-design.	meetings focusing on cross co-design between the SEA projects. Additionally, there have been a workshops between the EuroHPC projects (for example, the First EuroHPC19 Projects Collaboration Worrkshop in September 2022 in Madrid) as well and individual workshops between IO-SEA and ADMIRE.
Commercial Solutions	Please see Section 4, Exploitation	Please see individual exploitation pathways and plans of the commercial partners Seagate, ATOS and ParTec.
Exascale Platforms	<p>The project will continue to investigate the possibility of including IO-SEA components as part of the Modular Supercomputing Architecture (MSA)¹⁰ developed at Juelich, and also as part of the EUPEX pilot.</p> <p>The previously used SAGE prototype coming out of the SAGE project¹¹ will be repurposed to work with MSA further bolstering the collaboration there. This part is being finalised.</p>	The SEA projects are actively working towards the realisation of MSAs. The SAGE prototype has now been readied to connect to the DEEP-SEA system at Juelich. The outcomes of IO-SEA will feed into EUPEX as discussed with the EUPEX project.
Information Sharing	<p>To summarize, this will be achieved through a combination of:</p> <ol style="list-style-type: none"> 1. Project Website 2. Social Media (Twitter & LinkedIn) 3. Dissemination materials 4. Participation in Exhibitions 5. Publication of Papers 6. Press Releases 7. Social Media 	As will be discussed in Section 3.

Table 1 Dissemination Strategy

¹⁰ <https://juser.fz-juelich.de/record/862856>

¹¹ <https://sagestorage.eu/>

2.2.2 Collaborations

Focus Areas	Original Plan	Status as of Year 2
Community Adoption	Collaboration with SEA Projects to start with, and then further collaboration with other EuroHPC projects will aim towards spearheading community-building and adoption of the IO-SEA solution.	Very active collaboration with SEA projects on dissemination, and joint exploitation and usage of each others assets (e.g.: SAGE prototype usage in DEEP-SEA). Collaboration with EuroHPC projects such as ADMIRE are also ongoing.
HPC Strategy Alignment	Joint workshops, Birds of a Feather (BoF) events, etc with other EuroHPC projects & SEA projects will be a collaborative endeavour. These will hopefully also catalyse deeper technical engagements.	To be discussed in Section 3
Industry and Research Support	We will seek collaborations with the European Processor Initiative (EPI) and associated projects such as EUPEX.	Outcomes of IO-SEA will feed into EUPEX and that is already discussed, planned and co-ordinated with EUPEX. A partnership agreement that is now ready between the 10 EuroHPC projects provides a strong framework for building an ecosystem of indigenous European technologies supported by European industry and research.
Applications' Expertise	Collaborations between the EuroHPC and SEA projects will focus on cross co-design between the use cases and the other (non-IO-SEA) platforms further bolstering the application portfolios and exposing them to new challenges.	Applications are being re-used across projects , for example, the NEK5000 use cases used in DEEP-SEA will be assessed in IO-SEA. More re-use of applications across projects will be sought going forward.
Commercial Solutions	Please see Section 4, Exploitation	Please see Section 4, Exploitation
Exascale Platforms	The project dissemination will target organisations hosting pre-Exascale and upcoming Exascale machines. The project will also actively work and	This has been part of our dissemination strategy. The venues that has been picked up (ISC, SC,

	disseminate into EuroHPC projects working on Exascale Pilots and the European Processor (eg: the EUPEX project)	Teratec Forum, etc) targets stakeholders in organisations hosting these machines. Many discussions with these entities have taken place at these venues. We expect this active dissemination to continue to continue. Strong linkage with EUPEX already exists. Also, Juelich Supercomputing Centre (JSC) will host Europe's first exascale supercomputer, which bolsters the IO-SEA, DEEP-SEA and RED-SEA outcomes as JSC is a partner in all these projects.
Information Sharing	Please see Section 3, Dissemination and Collaboration	Please see Section 3, Dissemination and Collaboration

Table 2 Collaboration Strategy

2.2.3 Exploitation

Focus Areas	Original Plan	Status as of Year 2
Community Adoption	Community adoption through dissemination and collaboration will eventually lead to new exploitation pathways for IO-SEA technologies by helping to build and design entirely new solutions in the marketplace, supported by the community.	Open source outcomes from IO-SEA continues to be published on github. For further information see Section 4 for individual exploitation
HPC Strategy Alignment	Please see Section 3, Dissemination and Collaboration	Please see Section 3, Dissemination and Collaboration
Industry and Research Support	IO-SEA will help European Industry and Research participants (supported by SEA and EuroHPC project collaborations) in these projects to build new liaisons and explore new exploitation pathways	There are already strong exploitation pathways from the individual partners at the end of Year 2. We will discuss that in the exploitation section (Section 4).
Applications'	Please see Section 3, Dissemination and	Please see Section 3,

Expertise	Collaboration	Dissemination and Collaboration
Commercial Solutions	IO-SEA commercial partners (ATOS, ParTec and SEAGATE) have solid exploitation plans to build commercial solutions based on their contributions. This will be detailed in the Exploitation plan discussed later in the project.	Seagate will continue to work on the object store and S3 strategies (including their usage in real cloud deployments). ATOS continues to use the outcomes of IO-SEA in the Smart Data Management Suite. An SME in the IO-SEA project. ParTec continues to embellish the ParaStation Healthchecker and the ParaStation Management suite for use in Exascale HPC based on IO-SEA outcomes. For more details please see the exploitation section (Section 4).
Exascale Platforms	Usage in working exascale platforms is also a longer term exploitation pathway for IO-SEA technologies. This is supported by collaboration activities that are initiated.	The EUPEX project is the natural pathway and next stage for the outcomes of IO-SEA to move towards Exascale. EUPEX is a route for upcoming exascale procurements in Europe.
Information Sharing	Please see Section 3, Dissemination	Please see Section 3, Dissemination

Table 3 Exploitation Strategy

2.2.4 Training

Focus Areas	Original Plan	Status as of Year 2
Community Adoption	A solid training plan is now available that will propel community adoption. The training plan will be “backloaded” time-wise so that training activities start in earnest once early implementations start to become available. The training plan will be discussed later in the report in Section 5.	The training activities are proceeding as planned. We have had very good audiences for each of the training initiatives due to the publicizing through Twitter and LinkedIn channels and through the individual project partners

		More details in Training section, Section 5.
HPC Strategy Alignment	Please see Section 3, Dissemination and Collaboration	Please see Section 3, Dissemination and Collaboration
Industry and Research Support	IO-SEA will help European Industry and Research participants work towards training their respective communities – helping future exploitation and usage.	Individual partners such as Seagate have had their own tutorials/hackathon events to train the community as discussed in Section 4.
Applications' Expertise	Please see Dissemination and Collaboration in Section 3.	Please see Dissemination and Collaboration in Section 3
Commercial Solutions	Training (especially of all open source outcomes) is a pre-cursor to community building and later exploitation.	See "Industry and Research Support" above
Exascale Platforms	Training activities will focus on trying to encourage the use of IO-SEA technologies in very new environments and deployments targeting exascale computing in particular – as investigated by other projects such as EUPEX.	The training events to date have targeted the usage of tools (HSM, Phobos, Ephemeral Services, CORTX, etc.) on exascale platforms, and focused on giving the community a true hands-on experience.
Information Sharing	Please see Section 3, Dissemination and Collaboration	Please see Section 3, Dissemination and Collaboration

Table 4 Training Strategy

2.3 Definitions of KPIs for successful dissemination

Based on feedback we have received from the experts we have the following Key Performance Indicators or KPIs for successful D&E by the end of the project, which we verbally discussed with the experts during the M18 review. This is based on our past experiences of successful projects and also our goal of positioning ourselves with the EUPEX program for European Exascale.

Key Performance Indicator, or KPI	Goal
Exascale	Use of software components in EUPEX
Booth participation	10 events
Workshops and BoFs	15 events
Publications	8 publications

Hackathons	3 hackathons
Promo Videos	2 videos
Social Media/ Twitter	50K impressions
Website	3000+ unique visitors
Training	~10 events with 200+ total participants
Collaborations with other orgs	5 one-to-one collaborations
Press Releases	4 in total (including in co-ordination with other SEA projects)

Table 5 Key Performance Indicators

2.4 Alignment with Evolving European Objectives

The latest annual report from the EuroHPC Joint Undertaking was released in June 2022:

<https://eurohpc-ju.europa.eu/system/files/2022-06/Annual%20Activity%20Report%202021.pdf>

We have reviewed the missions and objectives of the JU and concurred that all the work that we continue to do is very much in alignment in the areas of promoting scientific excellence (that data-intensive applications in IO-SEA targeting Exascale), development of innovative supercomputer and data ecosystem (through the tools, methods, and techniques that we continue to build in IO-SEA), and also in the area of Federation of Supercomputing and Data Infrastructures, enabled through S3 gateways and object stores developed in IO-SEA.

Also the hosting agreement for JUPITER, Europe's first exascale supercomputer has now been signed between EuroHPC and Juelich Supercomputing Center, one of our IO-SEA partners, and the primary location of our IO-SEA prototype. We also continue to work very closely with DEEP-SEA that is co-ordinated by Juelich Supercomputing Center and look at how some of the technologies that were developed through previous collaborations by SEA project partners have come together in the JUPITER procurement.

As we work towards exploiting the results of IO-SEA into EUPEX, we will monitor the requirements of Europe's second Exascale deployment sometime towards the end of the SEA projects' timeframe.

It is also worth mentioning that major cloud initiatives in Europe such as GAIA-X¹² -- which is a Franco-German initiative to build a federated cloud infrastructure in Europe with data that will reside in European servers, with data-sharing abiding by the European data-sharing rules – is actively continuing. HPC remains a major piece in the GAIA-X architecture. The IO-SEA project tools and methods, e.g., RGW, Phobos,

¹² www.gaia-x.eu

etc., are also applicable for HPC in the cloud. Some IO-SEA partners participate in this initiative and assess the development of this ecosystem.

Further, IO-SEA partners ATOS, CEA, Juelich and Seagate continue to be on the Steering Board of the ETP4HPC¹³ organisation that is part of the RIAG, or, Research and Innovation Advisory Group of EuroHPC. The above partners and the University of Mainz have been active contributors of the ETP4HPC Strategic Research Agenda (SRA) which is one of the key documents in the European HPC ecosystem to help describe and develop the roadmap for HPC. Seagate and University of Mainz led the latest work done by the storage and I/O work group for the ETP4HPC SRA (SRA5)¹⁴. Knowledge and experience gained in the IO-SEA project (Cloud APIs, HSM, tape libraries, telemetry, etc) is brought to these efforts.

3 Dissemination, Exploitation and Collaboration Activities

We discuss the on-going Dissemination, Collaboration and Exploitation activities that have been pursued until the end of Year 2.

3.1 Dissemination, Exploitation and Collaboration Activities

We continue to heavily stress dissemination and training in this period of the project. The exploitation activities have started and is back-loaded with a full program for the third year of the project. For all dissemination activities, the target audience continues to be:

- Existing and future customers for our partners, products, and services
- Key influencers of the HPC and European exascale strategy (EuroHPC) who can provide a multiplier effect for our dissemination
- The HPC community in Europe and throughout the world.
- AI/Deeo Learning and Big-data communities in general, with increasing HPC demands
- Scientific research communities of the Centres of Excellence and those of our applications areas
- Early career scientists who will form the future workforce in HPC.
- EuroHPC projects as part of “Transition to Exascale” calls (10 projects including the SEA projects) in “Advanced Exascale Pilots” calls in 2020
- The European Processor Initiative, The European PILOT and the European Pilot for Exascale (EUPEX) interlinked projects helping to develop sovereign European Exascale pilot systems
- Wider European public
- Weather and Climatology communities (through ECMWF and ESIWACE)
- Cloud Computing community

¹³ www.etp4hpc.eu

¹⁴ <https://www.etp4hpc.eu/sra.html>

The audiences reached through the end of Year 2 can be categorized as follows:

- Web audience (1000+ unique visitors, as recorded in server statistics) accessing our website: <https://iosea-project.eu/> as will be discussed later in this section.
- Social Media:
 - Twitter audience (~227 followers, including large organisations and other research programs in the US and the EU) through our twitter handle @io-sea.
 - LinkedIn with 100+ followers
 - YouTube channel which is recently created¹⁵
- Audience from workshops, conference talks, etc. and we aim to obtain rough record estimates of the audience for each event.
- Targeted audience captured from publications, blogs etc is much harder to estimate, but is directly linked to the readership.
- Audience from booth activity – primarily from Supercomputing and International Supercomputing Conference.

We next list all the dissemination activities undertaken by the IO-SEA project.

3.1.1 Dissemination Activities

The following is a list of dissemination activities held during the first and second years. New events in the second year of the project are marked in RED below. The focus has been on introducing the technologies and the providing an overall picture of the IO-SEA project.

Date	Event Name	Title or Content	Venue	Lead	Audience type (& Number if Known)
8/6/2021	<i>Teratec Forum 2021 [1]</i>	<i>Extreme Data Challenges in Numerical Weather Prediction</i>	<i>virtual</i>	<i>ECMWF</i>	<i>HPC/HPDA</i>
22-24/6/2021	<i>Teratec Forum 2021 [1]</i>	<i>Presentation of the IO-SEA project (logo, website) at the booth of the LEXIS project</i>	<i>virtual</i>	<i>IT4I</i>	<i>N/A</i>
24/6/2021	<i>Teratec Forum 2021 (storage workshop)[1]</i>	<i>Stockage HPC : de nouveaux paradigmes pour s'attaquer aux défis des architectures Exascale</i>	<i>virtual</i>	<i>CEA</i>	<i>around 50 participants (online)</i>
25/6/2021	<i>The 30th International Symposium on High-Performance Parallel and Distributed Computing: "Emerging Open Storage Systems and</i>	<i>CORTX: An Object Store for the Data Intensive Era</i>	<i>virtual</i>	<i>Seagate</i>	<i>Remote, about 20 (Storage experts, research institutions, Industry)</i>

¹⁵ https://www.youtube.com/channel/UCRIN8VDHwMNC_uCKVGFI-MA

D6.2. Dissemination, Exploitation and Training Report and Future Plans Year 2

	<i>Solutions for Data Intensive Computing”, EMOSS 2021[2]</i>				
25/6/2021	<i>The 30th International Symposium on High-Performance Parallel and Distributed Computing: “Emerging Open Storage Systems and Solutions for Data Intensive Computing”, EMOSS 2021[2]</i>	<i>Global Memory Abstraction Solutions for Emerging Storage Systems</i>	<i>virtual</i>	<i>ATOS/Seagate</i>	<i>Remote, about 20 (Storage experts, research institutions)</i>
24/6 - 2/7/2021	<i>ISC’21 [3]</i>	<i>Presentation of the IO-SEA project (logo, website) at the booths of the LEXIS project and IT4I</i>	<i>virtual</i>	<i>IT4I</i>	<i>N/A</i>
2/7/2021	<i>The International Supercomputing Conference 2021: Workshop on the In Situ Co-Execution of High-Performance Computing & Data Analysis [4]</i>	<i>Advancing Object storage architectures for Extreme scale</i>	<i>virtual</i>	<i>Seagate</i>	<i>Remote, about 20 (Storage experts, research institutions, Industry)</i>
24/8/2021	<i>Summer school on Effective HPC for Weather & Climate 2021 [5]</i>	<i>Modern Storage</i>	<i>virtual</i>	<i>Seagate</i>	<i>Remote, around 50 (Students in HPC storage and I/O, Industry and Academia)</i>
24/6 - 2/7/2021	<i>ISC High Performance 2021 [3]</i>	<i>Presentation of the IO-SEA project (logo, website) at the booths of the LEXIS project and IT4I</i>	<i>virtual</i>	<i>IT4I</i>	<i>N/A</i>
5/11/2021	<i>Motr Interfaces Workshop</i>	<i>Motr Interfaces</i>	<i>virtual</i>	<i>Seagate</i>	<i>Remote, around 50 - HPC EU Projects</i>
15-18/11/2021	<i>Supercomputing Conference 2021 [7]</i>	<i>Presentation of the IO-SEA project (logo, printed leaflet, presentation) at the booth of IT4I</i>	<i>America’s Center Convention Complex, St. Louis</i>	<i>IT4I</i>	<i>> 200 attendees of the booth</i>

D6.2. Dissemination, Exploitation and Training Report and Future Plans Year 2

16/11/2021	SC21 BoF [6]	<i>Object Stores for HPC: a Devonian Explosion or an Extinction Event?</i>	virtual	CEA/Seagate/ECMWF	around 60 participants (online)
3/3/2022	CORTX Meet An Architect community event	IO-SEA Ephemeral Services	Remote	Seagate	Remote, about 20 (Storage experts, research institutions, Industry)
5/2022	Joint SEA-BoF at ISC22 [8]	Smart resource management beyond compute nodes	ISC 2022, Hamburg	CEA/Seagate	50+ attendees, HPC data centre and application community
29/5 - 2/6/2022	High-Performance Computing in Science and Engineering (HPCSE) conference 2022 [9]	Presentation of the poster "Cryo-Electron microscopy image processing on remote HPC clusters" by Martin Golasowski	Hotel Solan, Roznov pod Radhostem, Czech Republic	IT4I	80 participants of the conference
1/6/2022	The 31st International Symposium on High-Performance Parallel and Distributed Computing: Emerging Open Storage Systems and Solutions for Data Intensive Computing, 2022[10]	IO-SEA Talks - TBD	HPDC'22 Stockholm	Seagate	50 attendees
29/5 - 2/6/2022	ISC High Performance 2022 [3]	Presentation of the IO-SEA project (logo, printed leaflets, presentation) at the booth of IT4Innovations	Congress Center Hamburg, Germany + online	IT4I	> 135 registered attendees at the booth
29/5- 2/6/2022	ISC High Performance 2022 [3]	IO-SEA at the SEA projects booth	ISC, Hamburg '22	CEA and partners	Attendance by 100+ participants across the HPC ecosystem
20 - 22/6/2022	HiPEAC 2022 [11]	Presentation of the IO-SEA project within the talk regarding the LEXIS Platform by	Novotel Budapest City, Hungary	IT4I	20 attendees

D6.2. Dissemination, Exploitation and Training Report and Future Plans Year 2

		<i>Martin Golasowski</i>			
<i>13-18/11/2022</i>	<i>Supercomputing Conference 2022 [12]</i>	<i>IO-SEA project at ACROSS project booth</i>	<i>SC'22 in Dallas, TX</i>	<i>CEA and partners</i>	<i>Attendance by 100+ participants across the HPC ecosystem</i>
<i>13 - 18/11/2022</i>	<i>Supercomputing Conference 2022 [12]</i>	<i>Presentation of the IO-SEA project (logo, printed leaflets, presentation, stickers, beer coasters) at the booth of IT4I</i>	<i>Dallas, USA</i>	<i>IT4I</i>	<i>300 registered attendees at the booth</i>
<i>13/11/2022</i>	<i>BoF 1 at SC'22 [13]</i>	<i>The Storage Tower of Babel? . . . Not! Actually, maybe?</i>	<i>SC'22</i>	<i>CEA</i>	<i>Around 60 (Students in HPC storage and I/O, Industry and Academia)</i>
<i>13/11/2022</i>	<i>BoF 2 at SC'22 [13]</i>	<i>SC22 BoF131: Disaggregated Heterogeneous Architectures</i>	<i>SC'22</i>	<i>CEA w/ Sea projects</i>	<i>Around 50 (Students in HPC storage and I/O, Industry and Academia)</i>
<i>21 - 23/11/2022</i>	<i>European Big Data Value Forum 2022 [14]</i>	<i>IO-SEA booth at EBDVF</i>	<i>EBDVF'22, Prague, Czech Republic</i>	<i>Seagate</i>	<i>400 participants of the conference</i>
<i>2 - 3/12 2022</i>	<i>Hackathon at U of Edinburgh</i>	<i>CORTX hackathon at U of Edinburgh</i>	<i>Edinburgh, Scotland</i>	<i>Seagate</i>	<i>Attendance of 20 students</i>
<i>01- 2023</i>	<i>IO-SEA at SEA projects booth [15]</i>	<i>HiPEAC'23</i>	<i>Toulouse, France</i>	<i>CEA and partners</i>	<i>200+ attendees from HPC & Embedded Computing domains</i>
<i>01- 2023</i>	<i>IO-SEA presentation at HiPEAC'23 [15]</i>	<i>HiPEAC'23</i>	<i>Toulouse, France</i>	<i>ATOS</i>	<i>40+ attendees from HPC and embedded computing domains</i>

23-25 March, 2023	XIX Discussions in Structural Molecular Biology, CIISB ¹⁶	Poster	Nove Hrady, Czech Republic	IO-SEA Application Partners	150, researchers
March 2023	EUPEX session	Presentations about Phobos & Robinhood	online	CEA	Remote, around 30 EUPEX members
08/03/202 3	ETP4HPC General Assembly [16]	Presentation on Destination Earth and References to IO-SEA	Sassenhe im- Leiden, Netherla nds	ECMWF	100 Attendees from HPC technology community
20 – 23/03 2023	EuroHPC Summit Week'23 [17]	IO-SEA poster presentation	Gothenb urg, Scotland	Juelich	TBD
06/23	PASC Conference 2023 [18]	IO-SEA talk at "Data Management across the continuum" session	Basel, Switzerla nd	ATOS	TBD

Table 6 Dissemination Activities

3.1.2 Publications/Newsletters

The following is the status of the publications towards the end of the Year 2. Heeding feedback from the Month 18 Review in Luxembourg, we are focused on increasing the number of publications and have an action plan to generate them, with publication expectations set for each partner. New publications in Year 2 are marked in RED.

Date Submitted	Title of the article	Lead	IO-SEA Authors	Name of the publication	Status/Link
9/4/2021	Record number of H2020 projects in IT4I	IT4I	Markéta Dobiašová/IT4I	Internal newsletter IT4Innovations 14/2021	http://mailchi.mp/205401830f78/intern-zpravodaj-8113642
23/4/2021	IT4I is collaborating on a new H2020 project	IT4I	Markéta Dobiašová/IT4I	Internal newsletter IT4Innovations 16/2021	http://mailchi.mp/b55d31ea16b2/intern-zpravodaj-8113654

¹⁶ <https://www.ciisb.org/about-ciisb/about-ciisb>

4/5/2021	Storage facilities for the most powerful supercomputers in Europe will be developed in Ostrava	IT4I	Markéta Dobiašová/IT4I	Lupa.cz	https://www.lupa.cz/aktuality/v-ostrave-a-brne-budou-vyvijet-uloziste-pro-nejvykonnejsi-superpocitace-v-evrope/?utm_source=FB&utm_medium=post&utm_campaign=Lupa&fbclid=IwAR0-OpLhu2ohn4anwcNxNWzhmj3t3_Cl4ICj9-aPtF19pYINJ-KD00-yhm0
7/5/2021	Lupa.cz informs about the IO-SEA project	IT4I	Markéta Dobiašová/IT4I	Internal newsletter IT4Innovations 18/2021	http://mailchi.mp/a05490b6dfe6/intern-zpravodaj-8113666
4/6/2021	IT4Innovations participates in implementing the ACROSS and IO-SEA projects (CZ version)	IT4I	Markéta Dobiašová/IT4I	Newsletter IT4Innovations 1/2021	https://mailchi.mp/e269a1ac7050/it4innovations-newsletter-12021
4/6/2021	IT4Innovations participates in implementing the ACROSS and IO-SEA projects (EN version)	IT4I	Markéta Dobiašová/IT4I	Newsletter IT4Innovations 1/2021	https://mailchi.mp/955ea0aaa113/it4innovations-newsletter-8120109
15/7/2021	Storage facilities for the most powerful supercomputers in Europe will be developed in Ostrava (CZ version)	IT4I	Markéta Dobiašová/IT4I	Newsletter IT4Innovations 2/2021	https://mailchi.mp/36035502eb36/it4innovations-newsletter-8122029?e=af22d43958
16/7/2021	Storage facilities for the most powerful	IT4I	Markéta Dobiašová/IT4I	Newsletter IT4Innovations 2/2021	https://mailchi.mp/83846adaa318/it4innovations-newsletter-

	<i>supercomputers in Europe will be developed in Ostrava (EN version)</i>				8124029?e=af22d43958
14/09/2022	<i>Protected Functions: User Space Privileged Function Calls</i>	<i>U of Mainz</i>	<i>Moti, N., Salkhord eh, R., Brinkmann, A</i>	<i>International Conference on the Architecture of Computing Systems, 13-15/08/2023</i>	https://link.springer.com/chapter/10.1007/978-3-031-21867-5_8
15/03/2022	<i>An Enterprise-Grade Open-Source Data Reduction Architecture for All-Flash Storage Systems.</i>	<i>U of Mainz</i>	<i>Mohammadamin Ajdari, Patrick Raaf, Mostafa Kishani, Reza Salkhord eh, Hossein Asadi, André Brinkmann (2022).</i>	<i>Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS).</i>	https://dl.acm.org/doi/abs/10.1145/3489048.3530963

Table 7 Publications/Newsletters

3.1.3 Press Releases

We have had 3 press releases covering IO-SEA.

Date Submitted	Press Release	Lead	IO-SEA Authors	Status/Link
----------------	---------------	------	----------------	-------------

8/4/2021	<i>IT4Innovations excels in the Horizon 2020 projects (CZ version)</i>	IT4I	Markéta Dobiašová	https://www.it4i.cz/o-it4i/infoservis/tiskove-zpravy/it4innovations-soucasti-16-projektu-horizont-2020-a-jednim-z-klicovych-inovatoru
8/4/2021	<i>IT4Innovations excels in the Horizon 2020 projects (EN version)</i>	IT4I	Markéta Dobiašová	https://www.it4i.cz/en/about/infoservice/press-releases/it4innovations-part-of-16-horizon-2020-projects-and-one-of-the-key-innovators
23/8/2021	<i>Pan-European project IO-SEA for Exascale Data and Storage Technologies</i>	ICHEC	Venkatesh Kannan	https://www.ichec.ie/news/pan-european-project-io-sea-exascale-data-and-storage-technologies

Table 8 IO-SEA Press Releases

IO-SEA press releases have been picked up by very important venues such as HPCWire.¹⁷ We plan to make a couple of additional press releases towards the end of the project.

3.1.4 Website Update

The IO-SEA Website (<https://iosea-project.eu/>) is a valuable resource introducing the concepts developed within the project and is a platform for sharing project materials, publications, public deliverables, blogs, and event coverage. We continue to update it regularly.

The website highlights our links to the joint SEA projects. We have gathered some website usage statistics during the first year of the project.

Year 2 saw a 3.8-fold increase in new users compared to Year 1.

The new versus returning visitors is still relatively low, so building up returning customers is an area to look at going forward.

¹⁷ <https://www.hpcwire.com/off-the-wire/pan-european-project-io-sea-for-exascale-data-and-storage-technologies-kicks-off/>

D6.2. Dissemination, Exploitation and Training Report and Future Plans Year 2

The Y1 & Y2 results are shown below. Please note that the numbers for March 2023 are seen as low since the snapshot was collected only in early March:

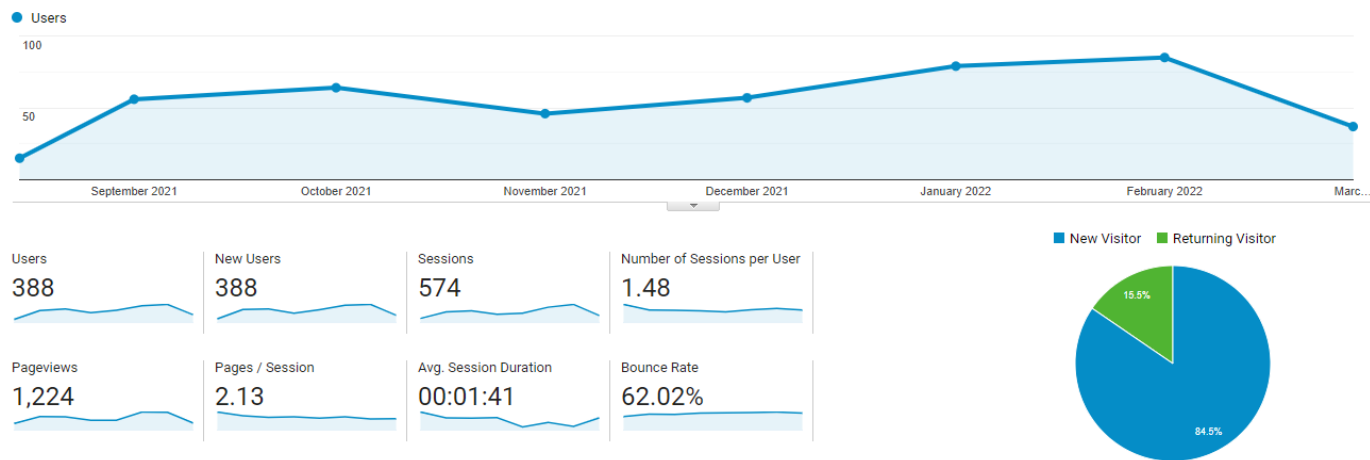


Figure 2 Website Statistics (Year 1)

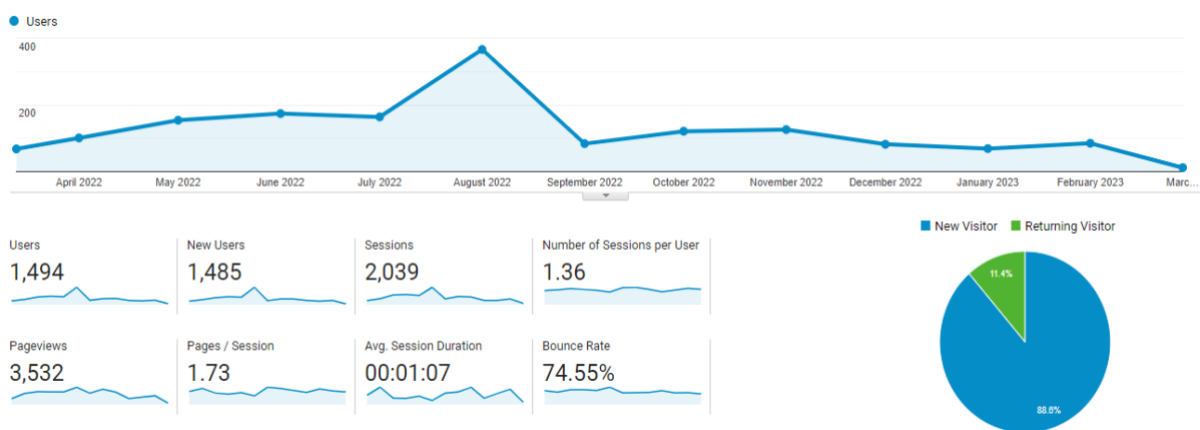


Figure 3 Website Statistics (Year 2)

3.1.5 Social Media Update

IO-SEA uses three channels for its Social Media strategy: Twitter, LinkedIn and YouTube. We use Twitter to engage audiences on important project highlights, and the focus is on content. We use LinkedIn as a mechanism to link with communities, and the focus is on people that can connect to the project.

We have also recently created the YouTube¹⁸ channel on behalf of all the SEA projects helping to disseminate video content, presentations, and Tutorials. We continue to provide videos on

¹⁸ <https://www.youtube.com/@seaprojects6122/videos>

a regular basis and are in a phase where we plan to include video materials on a more regular basis during Year 3.

Twitter

We have used Twitter since the beginning of the project. On an average, we have had about one tweet a week. The 226 followers include major EU projects such as EUPEX, EuroCC¹⁹, EsiWACE²⁰, etc and key influencers within the HPC community in Europe. This is increased from 151 followers we had at the end of Year 1.




Figure 4 IO-SEA Twitter Page

We have had a total of approximately 30 thousand Twitter impressions in the first two years of the project. Impressions are the number of times the tweet has been viewed. We summarize below the impressions we have had in one of the quarters, where we can see more activity during the times of major events such as Supercomputing (SC) in November.

¹⁹ <https://www.eurocc-access.eu/>

²⁰ <https://www.esiwace.eu/>

Your Tweets earned **4.6K impressions** over this **91 day** period



**More activity around events
(SC'22)**

Figure 5 Twitter impressions showing a spike of activity around the time of the SC 22 conference. (Quarter 7 of the project.)

Project Quarter	Impressions
Q1 (Apr 2021 – Jun 2021)	12.2K
Q2 (Jul 2021 – Sep 2021)	5.7K
Q3 (Oct 2021 – Dec 2021)	3.2K
Q4 (Jan 2022 – Mar 2022)	1.1K
Q5 (Apr 2022 – Jun 2022)	2.6K
Q6 (Jul 2022 – Sep 2022)	1K
Q7 (Oct 2022 – Dec 2022)	4.6K
Q8 (Jan 2023 – Mar 2023)	1.9K (Until end of Feb 2023)
Total	32.3K

Table 9 Twitter impressions by quarter.

LinkedIn

We set up the IO-SEA LinkedIn profile in Q3 of the project, with the aim of connecting the project with relevant people in the HPC community. This complements our Twitter presence and is used for dissemination and also attracting appropriate audiences linked to project partners (for example, members of past projects).

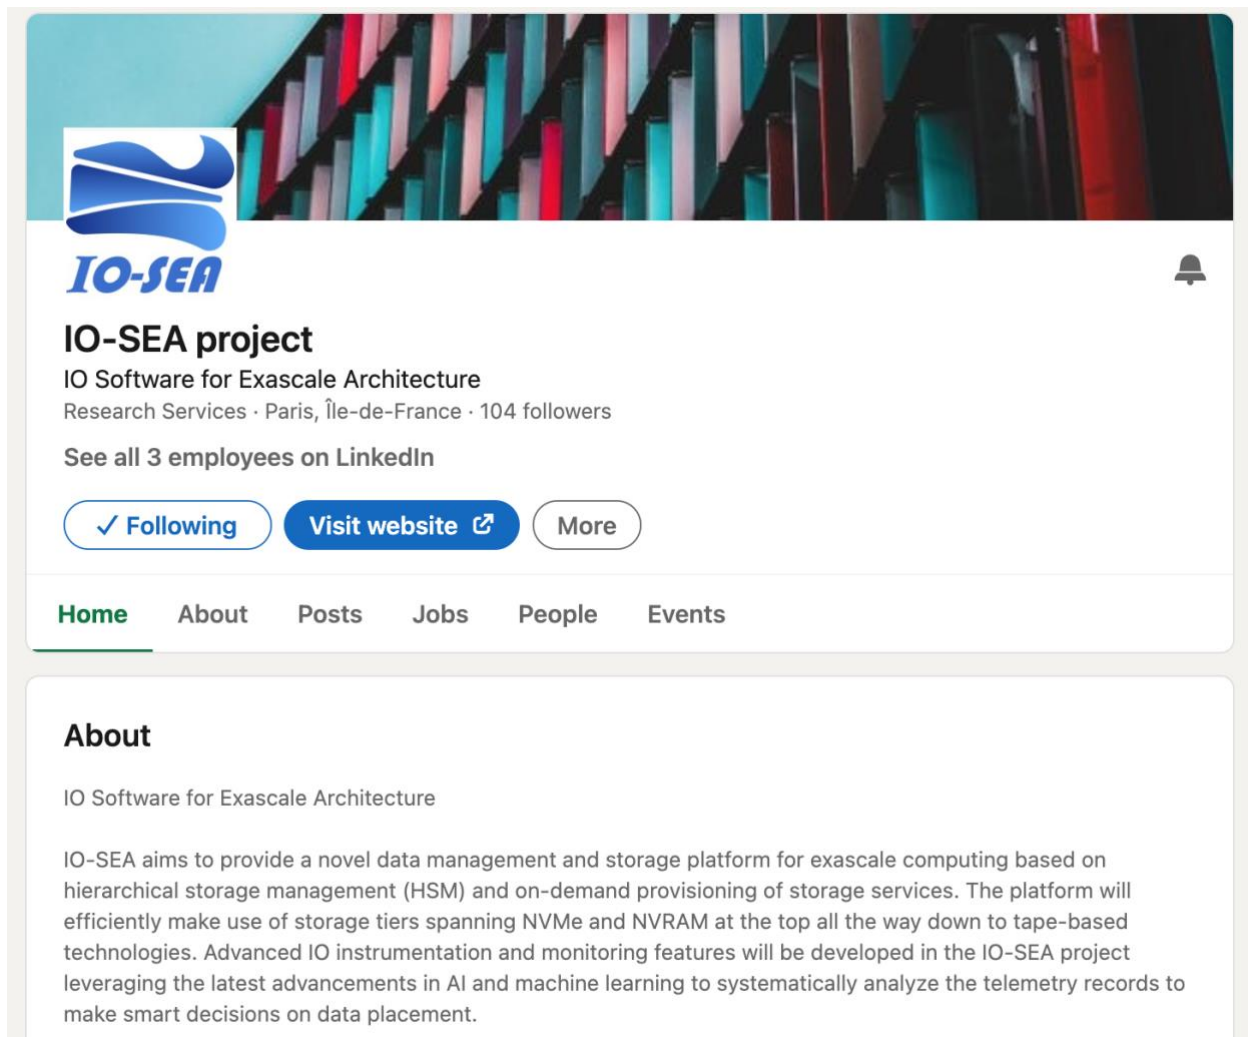


Figure 6 IO-SEA LinkedIn page

3.1.6 Snapshots

As seen in the dissemination log, IO-SEA has participated in major events such as the Supercomputing Conference (SC) in the United States, International Supercomputing Conference (ISC) and the Teratec Forum in Europe, and also major conferences such as HPDC²¹, SIAM²², etc. We provide some recent snapshots below from Year 2 of the project. We provide here a brief collage of pictures from the various events.

²¹ <https://www.hpdc.org/2022/>

²² <https://www.siam.org/conferences/cm/conference/cse23>



Figure 7 IO-SEA Dissemination Snapshots

3.1.7 Dissemination Materials

The dissemination materials included flyers and the introductory presentation deck (introducing the IO-SEA project) which have been continuously updated during the second year. We have also had prize giveaways for raffle drawings from IO-SEA partners. A sample of one of them is shown below in Figure 8.



Figure 8 Example giveaway - Portable Storage Device at the SEA projects booth at ISC'22

3.1.8 Contribution to joint websites

IO-SEA contributes actively to two other websites, the joint SEA projects website²³, which IO-SEA was managing on behalf of the other SEA projects, and the EuroHPC projects website. The joint SEA Projects website is now being revamped and will be managed by DEEP-SEA. The SEA Projects website also an excellent example of co-ordination between the SEA Projects.

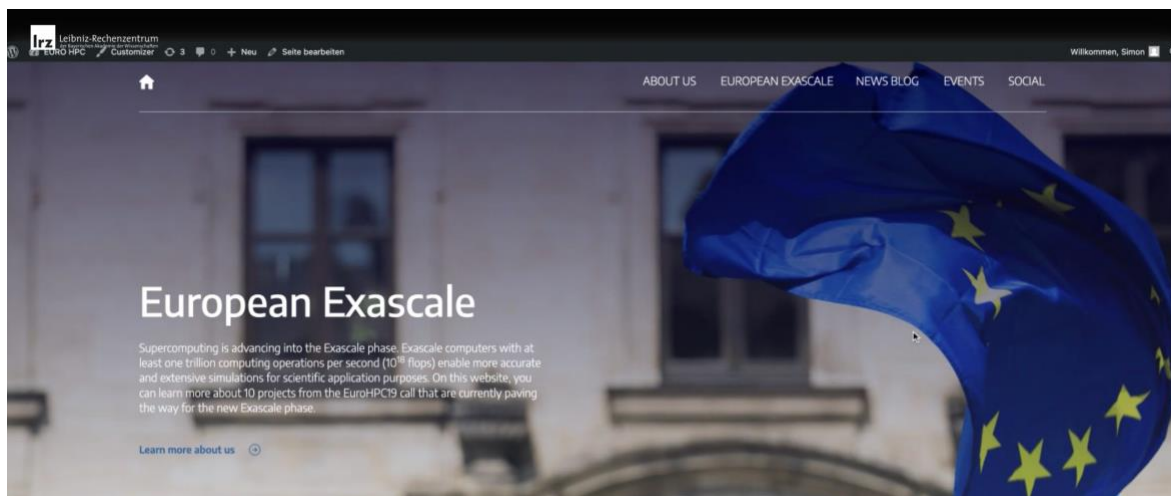


Figure 9 European Exascale projects website

²³ <https://www.sea-projects.eu/>

3.2 Collaborations

Our strategic objective for inter-project collaboration has been to identify complimentary value propositions for IO-SEA and jointly develop pieces of the Exascale HPC for the wider community going forward – primarily focusing on collaborations with the ten EuroHPC projects, including the three SEA projects.

We have collaborated with closely with RED-SEA and DEEP-SEA throughout the project and we have only increased these in the second half of the project. We have regular monthly co-ordinator sync-ups and a separate monthly dissemination sync up between SEA projects. This has enabled us to have joint activities (Booths participation, BOFs and events) in many major forums such as ISC, SC and the Teratec forum. Booth participation is always co-ordinated between the SEA projects optimizing on overall costs, demonstrating joint synergies to the community, and leveraging the personnel from the three SEA projects. We have also developed flyers and posters with a common “SEA theme”. An example below in Figure 10 is an IO-SEA poster presented at the recently concluded HiPEAC²⁴ event in Toulouse, France:

²⁴ <https://www.hipeac.net/2023/toulouse/#/>

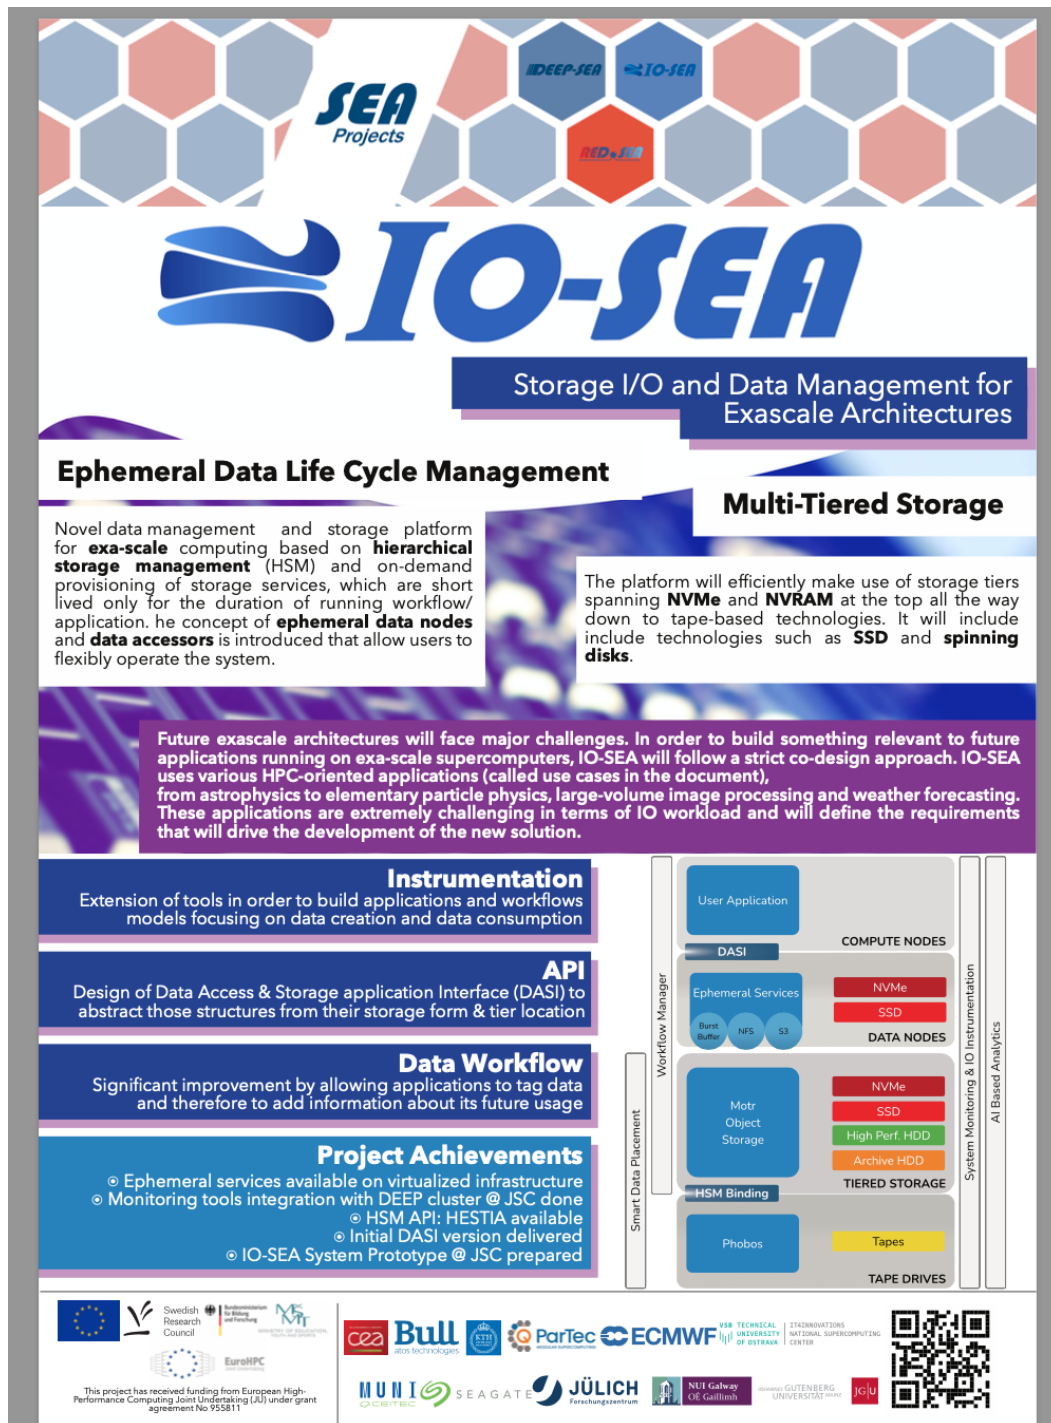


Figure 10 IO-SEA Poster with a common SEA theme

The SEA Projects YouTube channel²⁵ is also an example of joint collaboration between the SEA projects.

The SAGE prototype at Juelich forms the connecting link between the IO-SEA and the SEA projects, especially DEEP-SEA. The SAGE (also called the IO-SEA prototype) is now up and running and is being connected to the DEEP-SEA compute cluster hence putting the two big

²⁵ https://www.youtube.com/channel/UCRIN8VDHwMNC_uCKVGF1-MA

modules of the MSA together! We are also involved in using each others applications hence enabling “cross co-design”.

We have an active collaboration with ADMIRE project on IO traces, this is lead by JGU. This collaboration is tracked by EuroHPC as a technical workgroup (“I/O traces and performance analysis”). IO-SEA has a representative in the workgroup. Details of the members of the working group and early documentation is available here²⁶. The workgroup involves IO-SEA, ADMIRE, MAELSTROM, RED-SEA and DEEP-SEA projects and is led by ADMIRE. The IO-SEA project is actively contributing traces. We also have a collaboration with RED-SEA focused on making the network and storage mutually aware of each other.

DEEP-SEA and IO-SEA also have a common benchmarking team developing a common benchmarking framework.

The three SEA projects have spearheaded the development of a collaboration agreement with the ten EuroHPC projects. The ambitious collaboration agreement between 60 different partner organisations was finalised early in Year 2 of the projects.

IO-SEA partner Seagate collaborated with the EsiWACE2 project for installation of CORTX at DKRZ²⁷ weather and climate centre as a test system. Seagate also continues to collaborate with Intel on developing a solution with CORTX Motr and DAOS²⁸ with the RGW API.

4 Exploitation

The consortium continues to build a novel I/O environment for exascale-class supercomputers, based upon extending many existing products and components, as well as several components developed in the context of the project. There are strong exploitation pathways, particularly for the industrial partners ATOS, Seagate and Par-tec, and also for the open source technology components used within the project. We will focus on exploitation more strongly in Year 3 as the project developments reach maturity.

The EUPEX project is of particular interest, paving the way for a self-reliant European HPC industry, capable of delivering exascale-class supercomputers designed in Europe.

One of the objectives of EUPEX is to provide a software ecosystem for the pilot based on European technologies. Its design will not only take into consideration the needs of the key applications identified in EUPEX, but also those of the system operators for the management of large-scale Modular Supercomputing Architecture (MSA) systems.

The EUPEX software stack will address 4 objectives:

- **Management:** Define a management software stack to support the administration of modular systems while being versatile enough to meet the requirements of upcoming architectures.
- **Execution environment:** Integrate different components forming the execution environment that will enable the efficient utilisation of all available resources on the modular architecture of the EUPEX platform.

²⁶<https://fz-juelich.sciebo.de/s/5S9Jzub4DTRQNML>

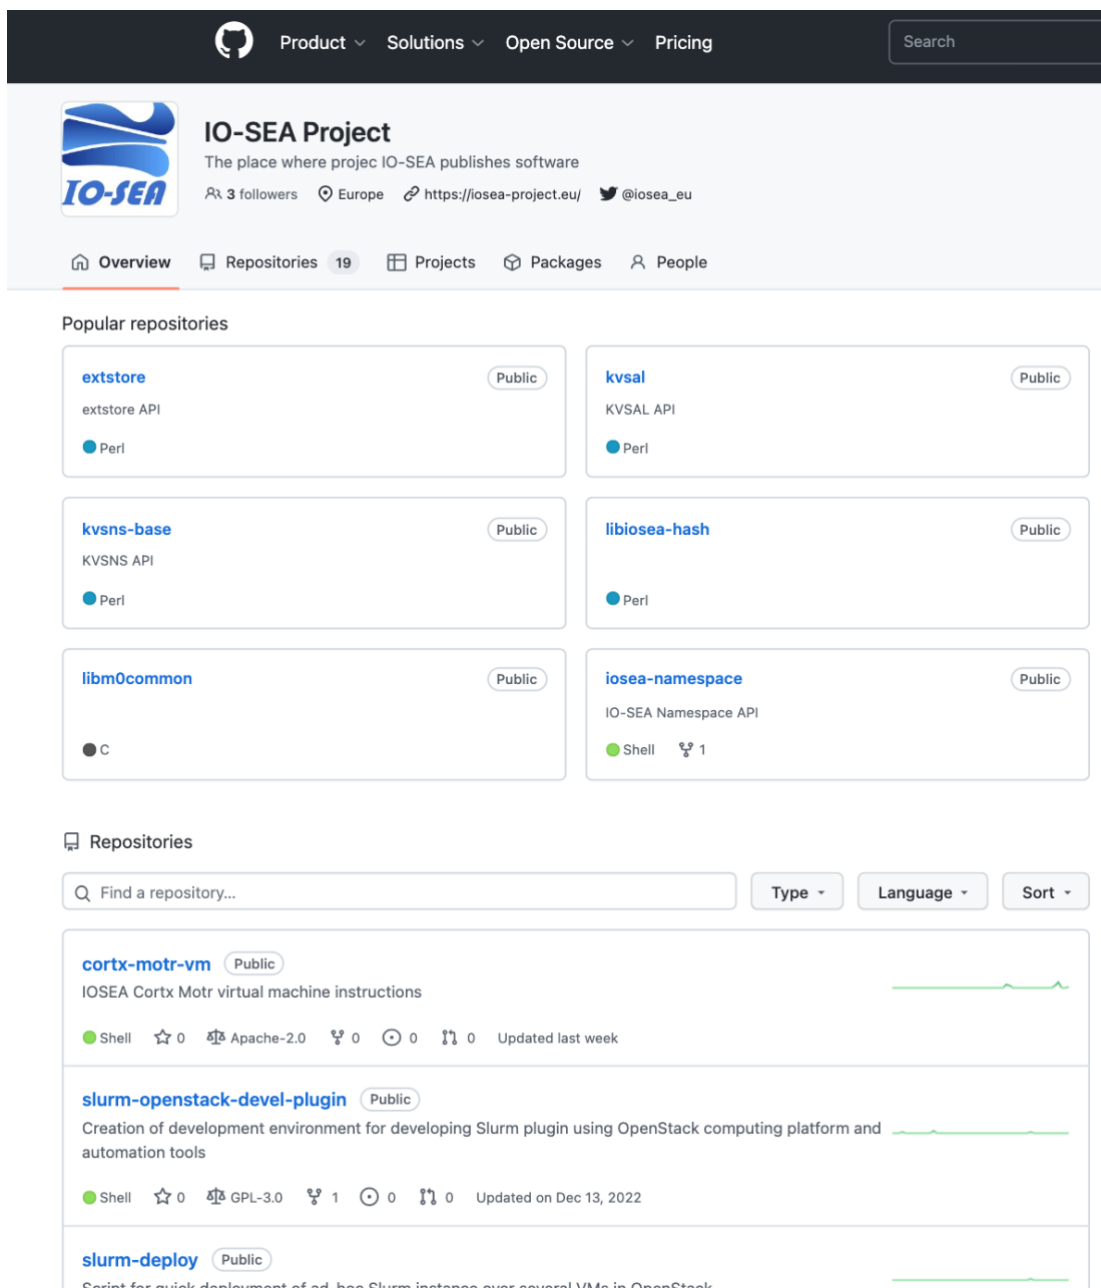
²⁷ https://www.dkrz.de/en/about-en/aufgaben/dkrz-and-climate-research/DKRZ_Rolle

²⁸ <https://docs.daos.io/v2.2/>

- **Tools:** Provide a set of tools to aid application developers and system operators in optimising the efficiency with respect to performance and energy, i.e., to maximise system utilisation.
- **Storage architecture:** Define a multi-tier storage architecture to meet the I/O demands of large-scale MSA systems, to transparently integrate fast storage technologies, and to minimize data movement.

The last objective, pertaining to storage architecture, encompasses the integration of the IO-SEA multi-tier storage platform.

In the previous Year 1 Report [19], we discussed plans for leveraging open source assets from the partners to develop the different tools and methods in the project. We are now fully exploiting those assets and it is now available at a single point through our Github page, which is public:



The screenshot displays the GitHub profile for the IO-SEA Project. At the top, there are navigation links for Product, Solutions, Open Source, and Pricing, along with a search bar. The profile header includes the IO-SEA logo, the project name, a description, and social media links. Below this, there are tabs for Overview, Repositories (19), Projects, Packages, and People. The 'Popular repositories' section features six cards for: extstore (Perl), kvsal (Perl), kvsns-base (Perl), libiosea-hash (Perl), libm0common (C), and iosea-namespace (Shell). The 'Repositories' section has a search bar and filters for Type, Language, and Sort. It lists three repositories: cortx-motr-vm (Shell, Apache-2.0), slurm-openstack-devel-plugin (Shell, GPL-3.0), and slurm-deploy (Shell).

Figure 11 IO-SEA github page

We expect these to be available to the other new EuroHPC Exascale pilot projects especially the EUPEX project. In Year 3 we will continue to collaborate closely with EUPEX to exploit the outcomes of IO-SEA towards European Exascale platforms.

Additionally, each of the industrial partners is working to exploit the outcomes of the work within their markets.

For the industrial partner ATOS, the features developed within the IO-SEA project have already been released in the Smart Data Management Suite and the effort will be pursued in order to make it ready for deployment on the EUPEX platform:

- Version 4, released in July 2022, includes the workflow support
- Version 5, planned for June 2023 release, will include extensions related to instrumenting ephemeral services and the support of the Ganesha ephemeral service
- Version 6, planned for early 2024, will include the workflow manager and a new user interface
- Version 7, planned for late 2024 (post end of the project), will include most of the IO-SEA features developed in WP2 and WP3

Seagate is continuing to develop new products and services based on innovations within the project. The innovations will be extended to provide a broad ecosystem for customers including those from the Cloud. Seagate's Lyve Cloud (which is now deployed) is an example of a new product offering that has had its roots in CORTX Motr²⁹. CORTX Motr continues to be positioned for use in future exascale prototypes and installations, and is the European object store of choice, especially for the research community.

IO-SEA continues to provide a strong exploitation pathway for European SMEs, exemplified by ParTec. Further developments of ParaStation HealthChecker and ParaStation Management³⁰ within the IO-SEA project will be integrated directly in ParTec's product portfolio. These advanced software packages should be published in production maturity and ready for use in the European exascale computers that are coming up. The new functionality regarding data management and storage will also enable improved I/O support on heterogeneous modular supercomputer architectures, and simplify their use for application developers.

The IO-SEA project enables ParTec to continue the development of ParaStation. For this purpose, ParaStation Management is expanded to include the ability to dynamically assign data nodes in heterogeneous modular supercomputers. Additionally, ParaStation HealthChecker will be further developed to take this into account in node monitoring.

The new functionalities implemented in IO-SEA will advance ParaStation and strengthen ParTec's position as an HPC service and software provider.

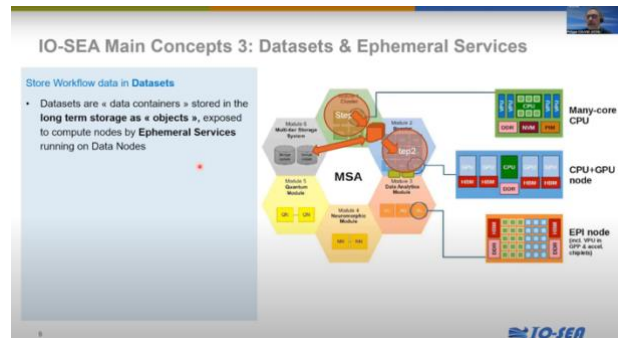
²⁹ <https://www.seagate.com/gb/en/services/cloud/storage/>

³⁰ <https://par-tec.com/software/>

5 Training

Training and education are a important functions of the project. These focus on the use of the IO-SEA on-demand runtime environment by data intensive applications outside the project, and the use of the individual tools and techniques developed within the project for by applications outside the project. During the first six months of the project, potential topics for training workshops, webinars and hands-on events were identified and preliminary training plan was set up. This preliminary plan was published in M6 in D6.1 [19]. The training activities themselves started in M19 and were held online or in hybrid mode, due to on-going effects of the COVID-19 pandemic, and the geographical dispersal of the trainees and tutors.

The first training event was organised as online webinar “Software for exascale architectures driven by a hierarchical storage management approach” on 13 October 2022. The webinar was divided into two parts; in the first part, Philippe Couvée (ATOS) presented the novel IO-SEA approach to data management, including the storage platform for exascale computing based on hierarchical storage management and ephemeral data life-cycle management.




The second part focused on how these will be used by our five data-intensive use cases, which were introduced by James Hawkes (ECMWF), Damien Chapon (CEA), Jiří Nováček (CEITEC), Eric Gregory (FZJ), and Ghazal Tashakor (FZJ).



The second training event, “CORTX Motr Object Storage” was held on 6 December 2022 as a hybrid event at IT4I in Ostrava, Czech Republic. The hands-on training was led by Ganesan Umanesan and David Vasiliauskas with the introduction to the IO-SEA project and the context of the training topic within the project done by Sai Narasimhamurthy (all Seagate). This hands-on training introduced the CORTX Motr object storage suitable for extreme-scale computing with a demo sample running on virtual machine.

Attendees also practiced a deployment and usage of CORTX and Sample apps/S3 with prepared examples.

The two-day online hands-on training event “Phobos – Object Store on Tapes” was held on 20-21 February 2023 by Sebastien Gougeaud and Guillaume Courier (CEA). The tutors described storage architectures in HPC and motivated the need for developing Phobos. The attendees had the opportunity to learn more about tape storage technologies, about I/O libraries in general, and how to manage them in a virtual environment. The second part of the training event included the use of Phobos, and its advanced features, such as resource tags, versioning, and object deletion.



The list of the organised training events including the links to the IT4I Indico system is provided by Table 10.

Topic no.	Event name	Responsible partner	Venue	Date
1	Software for exascale architectures driven by hierarchical storage management approach https://events.it4i.cz/event/148/	ATOS, ECMWF, CEA, CEITEC, FZJ	Online	13 October 2022
2	CORTX Motr object storage https://events.it4i.cz/event/156/overview	Seagate	Hybrid (IT4I)	6 December 2022
3	Phobos - object store on tapes https://events.it4i.cz/event/161/	CEA	Online	20-21 February 2023

Table 10 Organised training activities (M24)

Training materials were prepared for each training event organised as hands-on training workshops or webinars. The materials are available at IT4I website <https://events.it4i.cz> and are always available also after the end of the training events. The recordings from the training events can be found at SEA projects' YouTube channel as described earlier in the document.

The events were promoted to the wider community through social media, the networks of EuroCC, BDVA/DAIRO, HiPEAC, SEA projects community, other projects from EuroHPC-01 call, and through the partners' networks. Table 11 shows the audience of the organised training events:

Topic no.	Event name	Audience			Total
		Academic	Commercial	Government, non-profit	
1	Software for exascale architectures driven by hierarchical storage management approach	24	10	3	37
2	CORTX Motr object storage	22	4	2	28 (onsite 12, online 16)
3	Phobos - object store on tapes	13	14	2	29

Table 11 Organised training activities – attendance (M24)

IT4I continues to organise the training activities. The next training event is planned to be on 10 May 2023 and will focus on infrastructure monitoring tools. The training will be led by ParTec as an online event. An updated list of topics for the development of next training events including responsible partners for each topic and expected month of delivery for the rest of the project is shown in Table 12.

Topic	Description	Responsible	Preliminary
IO-SEA - 955811			

no.		partner	Deadline
4	Infrastructure monitoring tools	ParTec	M18 (to be organized in 10 May 2023)
5	Long term data storage in IO-SEA and Object storage in HPC applications	CEA, ICHEC, Atos	M24
6	JUBE Benchmarking Use Cases + Synthetics	FZJ	M29
7	Scientific data curation and DASI	ECMWF	M31
8	General technical session - IO-SEA interface, usage, features from all WPs	CEA	M30
9	EM - success story	CEITEC, IT4I	M33
10	RAMSES - success story	CEA	M33
11	ECMWF - success story	ECMWF	M33
12	TSMP - success story	FZJ	M33
13	LQCD - success story	FZJ	M33

Table 12 Training activities – updated plan (M24)

6 Conclusion

This deliverable summarises the overall strategy for dissemination, exploitation, collaboration and training and activities accomplished by the end of the second year of the IO-SEA project.

In Year 3, we will:

- active work on publications as per feedback we have received from the EuroHPC reviewers.
- continue to focus on collaboration events and workshops
- continue to work on joint-participation events with the other SEA and EuroHPC projects.
- focus on exploitation and links of outcomes developed in IO-SEA with the European exascale roadmap.
- continue to expand the social media footprint as we head into Year 3 of the project.
- monitor the European ecosystem around HPC, AI, and cloud computing and take steps to engage with other communities and examine how we can collaborate to meet the new challenges.

7 References

- [1] Teratec Forum, https://teratec.eu/gb/forum_2021/index.html, Accessed March 2023
- [2] EMOSS'21, <https://www.hpdc.org/2021/workshops/>, Accessed March 2023
- [3] International Supercomputing Conference (ISC), <https://www.isc-hpc.com/>, Accessed March 2023
- [4] *The International Supercomputing Conference 2021: Workshop on the In Situ Co-Execution of High-Performance Computing & Data Analysis*, <https://hpcda.github.io/>, Accessed March 2023
- [5] EsiWACE Summer School, <https://hps.vi4io.org/events/2021/esiwace-school>, Accessed March 2023
- [6] Supercomputing 2021 BoFs, <https://sc21.supercomputing.org/program/birds-of-a-feather/>, Accessed March 2023
- [7] Supercomputing 2021, <https://sc21.supercomputing.org/>, Accessed March 2023
- [8] BoF sessions SC 22, <https://www.isc-hpc.com/bof-sessions-2022.html>, Accessed March 2023
- [9] HPCSE'22 Conference, <https://www.it4i.cz/en/events/event-archive/high-performance-computing-in-science-and-engineering>, Accessed March 2023
- [10] HPDC EMOSS'22, <https://www.hpdc.org/2022/workshops/>, Accessed March 2023
- [11] HiPEAC 2023, <https://www.hipeac.net/2022/budapest/#/>, Accessed March 2023
- [12] SC 2022 Exhibits, <https://sc22.supercomputing.org/exhibits/>, Accessed March 2023
- [13] SC 2022 BoFs, <https://sc22.supercomputing.org/program/birds-of-a-feather/>, Accessed March 2023
- [14] EBDVF 2022, <https://www.bdva.eu/european-big-data-value-forum-2022>, Accessed March 2023
- [15] HiPEAC 2023, <https://www.hipeac.net/2023/toulouse/#/>, Accessed March 2023
- [16] ETP4HPC 2023 GA, <https://www.etp4hpc.eu/events/etp4hpc-conference-sassenheim-near-amsterdam-630.html>, Accessed March 2023
- [17] EuroHPC Summit Week 2023, https://eurohpc-ju.europa.eu/events/eurohpc-summit-2023-2023-03-20_en, Accessed March 2023
- [18] PASC Conference 2023, <https://pasc23.pasc-conference.org/>, Accessed March 2023
- [19] Dissemination, Exploitation and Training Report and Future Plans Year 1, IO-SEA Submitted Deliverable, March 2022

8 List of Acronyms and Abbreviations

AI – Artificial Intelligence

BDVA – Big Data Value Association

BSD – Berkley Software Distribution

D&E – Dissemination and Exploitation

DAIRO - Data AI and Robotics

DAOS - Distributed Asynchronous Object Storage

DASI – Data and Storage Application Access Interface

DKRZ - Deutsches Klimarechenzentrum

DL – Deep Learning

EsiWACE - Centre of Excellence in Simulation of Weather and Climate in Europe

ETP4HPC – European Technology Platform for High Performance Computing

EUPEX – European Pilot for Exascale

EUROCC - European National Competence Centers

GPL – General Public License

HiPEAC - High Performance and Embedded Architecture and Compilation

HPDC - High Performance Parallel and Distributed Computing

KPIs – Key Performance Indicators

MSA - Modular Supercomputing Architecture

PRACE – Partnership for Advanced Computing in Europe

RGW – Rados Gateway

SIAM - Society for Industrial and Applied Mathematics

SME – Small and Medium Enterprise