



IMPACT ENHANCEMENT, COMMERCIAL ROADMAP, EXPLOITATION AND STANDARDIZATION REPORT V1.0

DELIVERABLE NUMBER: D7.4

DUE DATE: 30.06.2024

DATE OF SUBMISSION: 19.07.2024

NATURE: R

DISSEMINATION LEVEL: PU

WORK PACKAGE: WP7

LEAD BENEFICIARY: CPLAY



DOCUMENT CONTROL SHEET

DELIVERABLE TITLE:	IMPACT ENHANCEMENT, COMMERCIAL ROADMAP, EXPLOITATION AND STANDARDIZATION REPORT V1.0
AUTHORS:	MAX BRUNTON (CPLAY)
CONTRIBUTORS:	KONSTANTINOS NIKAS (ICCS), GIANNIS LEDAKIS, CHRISTOS-ALEXANDROS SARROS (UBI), HERVE PONCE (SIPEARL), MICHAEL WURZER (KTM), FOIVOS ZAKKAK (RHAT-IE), POLYVIOS PRATIKAKIS (FORTH), MICHELE PAOLINO (VOSYS), SERGIO SAPONARA (UNIPi), CLAUDE CHAUDET (UNIGE), DANIEL KREFL (SED), CHRISTOS KOTSELIDIS, JUAN FUMERO, ATHANASIOS STRATIKOPOULOS (UNIMAN)
REVIEWERS:	MICHAEL WURZER (KTM), KRZYSTOF NIENARTOWICZ (SED)
APPROVED BY:	CHRISTOS KOTSELIDIS (UNIMAN), DIONISIOS PNEVMATIKATOS (ICCS)

DOCUMENT HISTORY

Version	Date	Status	Description/Comments
0.1	20.05.2024	Draft	ToC
0.2	21.06.2024	Draft	Initial content added
0.3	30.06.2024	Draft	Added partners' exploitation plans
0.4	08.07.2024	Draft	Released for internal review
0.5	15.07.2024	Draft	Updated draft based on the reviewers' comments
1.0	19.07.2024	Final	Final version submitted to EC



DISCLAIMER

AERO has received funding from European Union's Horizon Europe research and innovation programme under Grant Agreement No 101092850. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the granting authority. Neither the European Union nor the granting authority can be held responsible for them.

This document contains material and information that is proprietary and confidential to the AERO consortium and may not be copied, reproduced or modified in whole or in part for any purpose without the prior written consent of the AERO consortium.

Although the material and information contained in this document is considered to be precise and accurate, neither the Project Coordinator, nor any partner of the AERO Consortium nor any individual acting on behalf of any of the partners of the AERO Consortium make any warranty or representation whatsoever, express or implied, with respect to the use of the material, information, method or process disclosed in this document, including merchantability and fitness for a particular purpose or that such use does not infringe or interfere with privately owned rights.

In addition, neither the Project Coordinator, nor any partner of the AERO Consortium nor any individual acting on behalf of any of the partners of the AERO Consortium shall be liable for any direct, indirect or consequential loss, damage, claim or expense arising out of or in connection with any information, material, advice, inaccuracy or omission contained in this document.



TABLE OF CONTENTS

1	Introduction	7
2	Exploitation and Commercial Roadmap.....	8
2.1	The open-source AERO exploitation strategy	8
2.2	The commercial AERO exploitation strategy	9
2.3	SWOT analysis.....	9
2.4	Expected project outcomes	10
2.5	Preliminary market analysis.....	10
3	Exploitation efforts	13
3.1	Academic partners	13
3.2	Industrial partners	15
4	Standardization Activities.....	18
4.1	RHAT	18
4.2	UNIMAN	19
4.3	CPLAY	20
5	Summary	21



Executive Summary

This deliverable will outline the commercial strategy for the AERO Project. This strategy is split into two parts. Firstly, the open-source strategy which will be applied to the open-source project outcomes. These account for the majority of the AERO outcomes, as this is a fundamentally open-source endeavour. However, there are also some proprietary solutions that complement these open efforts, led primarily by the industrial AERO partners (CPLAY, RHAT, SIPEARL, KTM, VOSYS, SED and UBI). These outcomes will have their own strategy that assists the open-source strategy of AERO.

An updated SWOT (Strengths, Weaknesses, Opportunities, Threats) analysis has been carried out to reflect the current strengths and weaknesses of the project's exploitation efforts - namely that the very technical nature of the project will require some work to simplify, in order to then create convincing value propositions and benefits for both the open-source and commercial strategy of AERO. Additionally, the project outcomes in Section 2.4 have been updated to account for changes over the first 18 months of the project; specifically, the addition of Maestro and VOSySmonitor as two proprietary, commercial outcomes.

Finally, Section 4 of the deliverable reports on the efforts of CPLAY, RHAT and UNIMAN towards contributions to standards.



List of Abbreviations & Acronyms

Abbreviation/Acronym	Meaning
CTS	Conformance Test Suite
DoA	Description of Action
HPC	High Performance Computing
MPP	Massively Parallel Processing
OCK	oneAPI Construction Kit
PEST	Political, Economic, Social and Technological
RVV	RISC-V Vector
SIG	Special Interest Group
SWOT	Strengths, Weaknesses, Opportunities and Threats
TG	Target Group
UDF	User Defined Functions



1 Introduction

This deliverable presents the ongoing progress of WP7 as regards impact enhancement, our planned commercial roadmap, exploitation activities and standardization efforts. This M18 report outlines how AERO members will maximize the opportunity to successfully exploit the project's outcomes. This includes commercial success, but also the desired impacts across the European Processor Initiative (EPI) developer ecosystem.

The report makes close reference to the EPI, and specifically the EU Processor. For the purpose of the AERO Project, the 'EU Processor' refers to the ARM-based Rhea-1 Processor being developed by SIPEARL, the first-generation of EPI's General Purpose Processor.

The remainder of this deliverable is organized as follows:

- **Section 2** describes the AERO exploitation plan and commercial roadmap by examining each of the expected project outcomes. It details the general strategy of the project as outlined in the DoA. This section also includes an updated SWOT analysis for the project, and some preliminary market analysis – though this work will feature much more prominently in the M36 report, per the AERO Exploitation and Business Plan.
- **Section 3** presents the exploitation plans of the academic partners and focuses on the commercial interests and plans for each industry partner. This is necessary because, although the project has overall expected outcomes, certain areas hold more investment and interest for industry partners. Industry partners therefore seek to maximise the opportunities for their efforts related to the AERO project with their own plans for exploitation.
- **Section 4** reports the targeted standardization activities of the project until M18.
- **Section 5** summarises the deliverable.

2 Exploitation and Commercial Roadmap

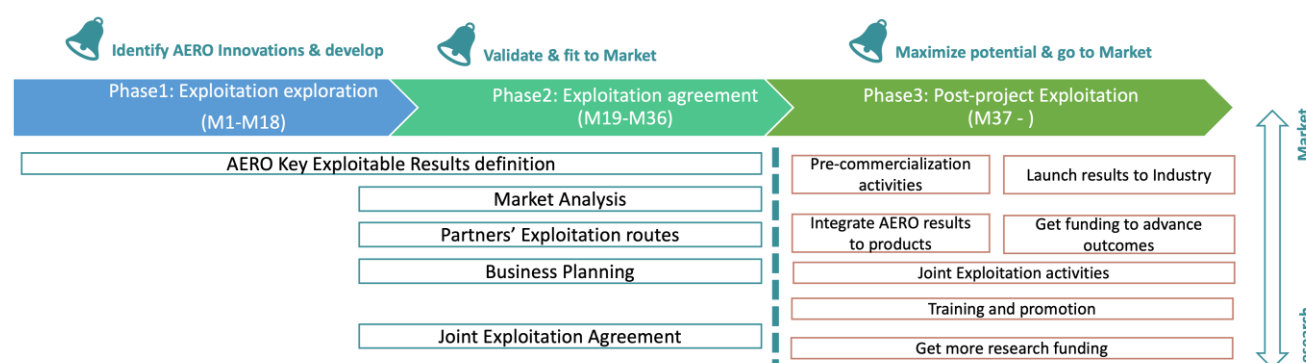


Figure 1 The AERO exploitation strategy

According to the AERO exploitation strategy depicted in Figure 1, this M18 report marks the conclusion of **Phase I** of the AERO exploitation plan. This phase has involved agreeing on the expected project outcomes and outlining the commercial plans for the rest of the project. This will lay the groundwork for the Go-to-Market strategy and market analysis that will be featured in the M36 report. Hence, this deliverable presents the identified key exploitable outcomes of AERO and presents a preliminary market analysis along with exploitation interests of all partners.

2.1 The open-source AERO exploitation strategy

As highlighted in the DoA, the AERO consortium has identified three methods for the exploitation of project results. As a fundamentally open-source project, the **open-source shared benefits exploitation model** is the primary way by which project outcomes will be exploited. In a 2021 report on the impact of open-source software, it was estimated that this was worth between €65-€95 billion to the European economy¹. Therefore, by contributing to this open-source ecosystem, AERO can directly benefit the EU economy, especially as the report also concluded that a 10% increase in open-source contributions could generate an additional 0.4-0.6% per year in EU GDP. In this way, AERO aligns strongly with the open-source software strategy of the European Commission².

A few initial exploitation activities have already been identified, tied to this open-source exploitation model. For the academic partners of AERO these include:

- Undergraduate, postgraduate and technical training. It can be measured by the number of students trained per year.
- MSc and PhD dissertations/theses tackling AERO-related problems and challenges. It can be measured by the number of research students involved.
- New research grants related to Cloud and High Performance Computing (HPC) or other areas of key European interests. It can be measured by the number of applications.

¹ <https://ec.europa.eu/newsroom/dae/redirection/document/79021>, p.14

² https://commission.europa.eu/about-european-commission/departments-and-executive-agencies/digital-services/open-source-software-strategy_en



2.2 The commercial AERO exploitation strategy

While predominantly open source, there exist also some proprietary AERO outcomes. These outcomes will still benefit the open-source ecosystem and encourage development with open software, but will have proprietary and/or protected features. Some examples of these are VOSySmonitor and Maestro developed by VOSYS and UBI respectively.

In general, AERO's commercial aims are linked closely to the open-source shared benefits model, as all commercial/proprietary options stem from the open-source activities. By reaching our target audiences and building the number of users engaging with the open-source AERO outcomes, partners can then upsell where necessary/appropriate to proprietary solutions, such as the oneAPI Construction Kit (OCK) or VOSySmonitor.

All the proprietary AERO outcomes come from the industrial partners, and, therefore, the relevant partners will be responsible for monitoring the potential value of their exploitable outcomes and establishing a suitable pricing model tailored to their clients. For instance, VOSYS will offer paid services for custom porting and for developing custom solutions on top of VOSySmonitor.

Besides exploiting each component separately, the AERO partners have been discussing the commercial exploitation of AERO as a whole, based on its value proposition that can be summarised as follows:

“AERO offers a complete software stack for enabling heterogenous cloud infrastructure leveraging the first EU processor”.

As the preliminary market analysis presented in Section 2.5 reveals, there exists a significantly large customer base that could be attracted to the complete AERO offering, such as cloud service and infrastructure providers and operators, research organisations and academic institutes, individual researchers, etc. Potential exploitation routes could include subscriptions to the AERO service over a SaaS model, consulting services, and customisation fees or support contracts.

The commercial exploitation plan for AERO, either as a complete suite or separate components, will be finalised and presented in Deliverable D7.5 in M36. It will include a market analysis together with a complete Business Plan towards the commercialisation of AERO, identifying the potential customer segments, income streams, and the appropriate channels to cultivate business relationships.

2.3 SWOT analysis

To complement the initial PEST analysis included in the AERO DoA, a SWOT analysis has been carried out to reflect the exploitation progress until M18. It is presented in Table 1.

Table 1. AERO SWOT analysis

Strengths	Weaknesses
<ul style="list-style-type: none"> ➤ Essential infrastructure for digital supply chains ➤ Open-source project, integrating with EU's open-source strategy ➤ Breadth of expertise across partners, key players in the EPI ecosystem ➤ Partners involved in standardization bodies with good contacts throughout open-source ecosystems ➤ Some components already in use (e.g., TornadoVM being integrated in Gaia) 	<ul style="list-style-type: none"> ➤ Very technical, hard to explain benefits to broad audience ➤ Cost of integration vs benefits not always simple to assess
Opportunities	Threats
<ul style="list-style-type: none"> ➤ Synergies with other projects in the EPI ecosystem (RISER, Vitamin-V) ➤ Sponsorship and exhibition opportunities to connect with target audiences at key events 	<ul style="list-style-type: none"> ➤ Delays in the delivery of the project's hardware ➤ Software bugs/critical errors

2.4 Expected project outcomes

Table 2 presents the identified project outcomes, along with their licensing and exploitation routes. The outcomes can be grouped in 6 general categories, all of them targeting the SIPEARL Rhea platform.

Table 2. AERO outcomes

Category	Outcomes	Licensing	Exploitation Routes
Managed programming languages & application ecosystem	OpenJDK, Quarkus, GraalVM, TornadoVM, TeraHeap	Open-source	Support, Training, Consulting, Extension
Native programming languages and application ecosystem	SYCL, LLVM	Open-source	Support, Training, Consulting, Extension
Cloud software stack components	Knot, FunctionBench, FaasRail, Firecracker extensions	Open-source	Support, Training, Consulting, Extension
	Maestro	Proprietary	Paid support, custom porting and development services
Drivers	OpenCL driver for the oneAPI Construction Kit	Open-source	Support, Training, Consulting, Extension
Security	Optimised drivers for the security extensions of Rhea, Homomorphic Encryption	Open-source	Support, Training, Consulting, Extension
Virtualization	VOSySmonitor	Proprietary	Paid support, custom porting and development services

2.5 Preliminary market analysis

While this endeavour has only just begun and will feature more prominently in the M36 report, some preliminary research of the open-source software ecosystem within Europe already shows good opportunities for the AERO project.



Firstly, it is worth reaffirming the AERO Target Groups (TGs) in Table 3. Considering the HPC Stakeholders, research in 2023 has shown that 60.4% of developers across Europe, the Middle East and Africa (EMEA) plan to be involved in projects that utilize HPC resources within the next twelve months or are already involved in such³. While this study does cover areas beyond Europe, in this preliminary stage it is nonetheless useful to acquire a sense of scale for the upcoming emphasis of software development with supercomputing projects.

Similarly, across the EMEA 45.6% of developers create heterogeneous applications – with 28.8% of these developers then using OpenCL, and 14.7% using SYCL⁴. This indicates that there is a significant audience of HPC Stakeholders (system software companies, HPC users, application vendors, etc.) that AERO can target with its project outcomes, and especially the first two categories presented in Table 2.

Table 3. AERO Target Groups

AERO Target Groups	
Industrial Cloud Stakeholders	Cloud Service Providers, Industrial cloud users
HPC Stakeholders	Industrial HPC users, HPC manufacturers, HPC application vendors, System software companies, Storage facilities
Industry Association & Technology clusters	European initiatives and clusters (EPI, NESSI, SEA4SA), associations (PRACE, ETP4HPC), federations (ACM, IEEE, NEM)
Academia & Research Community	Research organisations and academic institutes, individuals engaged in research initiatives and/or working in research/academic organisations/institutes
H2020/HEU Green cloud computing projects and international projects	Participants of EPI, EuroCC and EuroHPC projects, EuroCloud, Green Flash of the Berkeley's Laboratory, etc.
Policy Makers, Standardisation Organizations	Policy makers (EC Directorates and Units, EuroHPC JU, Ministries and Governments, Regulatory Agencies), Standardization Organisations (ETSI, Khronos, OpenJDK, CNF)
Wider Public	Civil society representatives, general public and anyone interested in the project

In addition, looking at the Industrial Cloud Stakeholders, it is evident from research that the software ecosystem around Cloud computing is continuously growing. Error! Not a valid bookmark self-reference. and Table 5 presenting some of the research performed by Evans Data Corp in 2023 are particularly relevant for AERO, which has the mission of building and supporting the cloud infrastructure for Rhea, the first European processor. It is evident that there exists a large audience of developers that AERO can target during the implementation of its exploitation strategy, highlighting especially its security, cloud and virtualization outcomes (presented in the third, fifth and sixth categories in Table 2).

³ Evans Data Corporation (2023) Global Development Survey Report 23.1, Santa Cruz, p. 176

⁴ Evans Data Corporation (2023) Global Development Survey Report 23.1, Santa Cruz, p. 174



Table 4. The percentage of organisations already using Cloud and how soon the rest of them are considering using Cloud⁵

(% within row)		What are your plans for deploying apps to a Cloud environment?					
		Done / In progress	1-6 months	7-12 months	13-24 months	> 24 months	No plans
Region	APAC	42.7	28.2	16.8	19.9	4.7	7.3
	EMEA	35.9	29.5	20.3	13.9	5.0	15.6
	North America	25.9	43.5	28.9	18.9	9.0	8.7
	Latin America	42.7	31.4	20.0	11.9	2.7	6.5

Table 5. Distribution of organisations based on their usage of virtualised instances⁶

(% within row)		Does your organisation deploy virtualised instances of hardware, software or infrastructure?		
		Yes	No	Not aware
Region	APAC	67.9	27.9	4.2
	EMEA	58.2	32.9	8.7
	North America	71.3	20.4	8.2
	Latin America	68.6	22.7	8.6

⁵ Evans Data Corporation (2023) Global Development Survey Report 23.1, Santa Cruz, p.234

⁶ Evans Data Corporation (2023) Global Development Survey Report 23.1, Santa Cruz, p. 237



3 Exploitation efforts

3.1 Academic partners

AERO comprises several academic partners (research institutes and universities) who have their own exploitation efforts underway. These efforts relate closely to the open-source strategy shared benefits exploitation model of AERO, whereas the industry partner interests are more commercial in nature. The following subsections elaborate on the efforts and plans of each academic partner.

3.1.1 ICCS

As ICCS is a non-profit research organisation, its main target is to enhance its technological know-how through its participation in AERO in order to utilise it in future research activities.

In the context of AERO, ICCS focuses in particular on the upbringing and optimisation of various components of the serverless software stack. Serverless computing is currently one of the top trends in Cloud computing. The serverless market attained in 2023 a value of approximately 8.79 billion USD and is projected to grow at a CAGR of 30% between 2024 and 2032⁷. Hence the expected publications (ICCS has already published 2 relevant papers) and the release of the relevant open-source code will improve ICCS' scientific position at the international level, reinforce already established research collaborations, and create new synergies that will lead to new research grants.

Further, in the first 18 months of the project, the ICCS team includes two PhD students, who are actively pursuing their PhD in the fields of system software and cloud and serverless computing.

3.1.2 FORTH

FORTH has released several versions of Knot and TeraHeap, as both are ongoing projects. The TeraHeap ARM port has been merged in the main TeraHeap repository, and has been presented to industrial and research partners. Further research within AERO has produced TeraHeap extensions and functionality that is expected to be accepted for publication in the near future. Moreover, FORTH has done work on LLVM and RISC-V within AERO, which has resulted in two accepted talks in EuroLLVM, accepted and merged pull requests in LLVM, as well as connections with the LLVM and MLIR communities. Finally, work on Knot, TeraHeap and LLVM within AERO has involved several MSc theses and one PhD candidate's work.

3.1.3 UNIFI

UNIFI plans to exploit the results of the activities in the AERO project in several ways.

First, by making available to partners the numerous HPC platforms available at its Green Data Center (GraceHopper, Ampere Altra, etc.), UNIFI will reinforce its network of academic and industrial collaborations to ease participation in new industry-funded or institutional funded projects. Additionally, the results of the research on cloud-based user data privacy thanks to efficient implementation of homomorphic encryption will also be a key asset to achieve new industry-funded

⁷ Source: <https://www.expertmarketresearch.com/reports/serverless-architecture-market>



or institutional funded projects. UNIMI will also evaluate patenting some solutions and/or implementing knowledge-transfer and technology-transfer strategies via one of the 14 spin-offs already activated by the Department of information Engineering. In case a decision is made to not cover the idea with a patent, or after the patent, the results will be also disseminated to the technical community via presentation at conferences and/or publication in journals.,

Furthermore, the results will be exploited also in terms of high-level education by involving in AERO activities master thesis students and PhD candidates. Finally, UNIMI will facilitate synergies among the activities in AERO and the activities in the Italian Center for Supercomputing (ICSC), where UNIMI leads Spoke 6 on Engineering Applications.

3.1.4 UNIGE

UNIGE is preparing the 4th data release of the Gaia stellar catalogue. The preparation phase uses several outcomes of the AERO project to improve computation efficiency, notably UNIMAN's TornadoVM and SED's MPP Database. The radical speedup possible thanks to the porting of algorithms to TornadoVM opened a new dimension of the scientific output related to classification of variable stars. This has already affected the planning and projected results of the project.

Once the development, evaluation and production phases (producing the variable stars catalogue for Gaia Data Release 4) are over, various publications are expected in order to disseminate the adaptations that were done to the computation methods, guidelines and results within the community and to broader non-astronomical audiences.

3.1.5 UNIMAN

In the context of AERO, significant technological know-how has been acquired related to the new optimizations that are being developed in TornadoVM and are tailored to the AERO hardware platform and the AERO use cases.

Since the beginning of the project, TornadoVM has released the following versions: v0.15.0, v0.15.1, v0.15.2, v1.0.0, v1.0.1, v1.0.2, v1.0.3, v1.0.4 and v1.0.5. The continuous development of TornadoVM has tracked the attention from industry, resulting in being invited in numerous industrial-based events, such as DevovxBE'24, JVMLS'23, DevovxUK'24, Gaia Data Processing Consortium '24, to present the latest updates on the technology. Notably, UNIMAN has been one of the few universities invited to present TornadoVM in industry-leading conferences. Furthermore, the traction of the new optimizations of the technology has resulted in reaching more than 1k stargazers in GitHub.

Additionally, UNIMAN has used TornadoVM as the core technology for the successful fulfilment of numerous master thesis and bachelor projects.

Finally, research staff from UNIMAN are working to initiate the commercialization of TornadoVM. Their objective is to incorporate a start-up company that will productise the TornadoVM for the AI market.



3.2 Industrial partners

3.2.1 UBI

In the scope of AERO, UBI is improving its MAESTRO framework for usage in the envisioned EU cloud. MAESTRO helps with the orchestration of cloud and edge applications and their whole lifecycle, through a Kubernetes compatible and proprietary orchestration logic that can also integrate with 5G telco and cloud providers. The improvements in MAESTRO, along with the adoption of an open-source approach will allow MAESTRO to increase its customer base and offer UBI increased opportunities to address the cloud (and cloud-to-edge) orchestration market within Europe. Specifically, through the developments in AERO, MAESTRO will provide ready to use application templates that can be deployed seamlessly to the envisioned AERO platform (SIPEARL's Rhea with ARM Neoverse V1 cores, together with RISC-V and GPU accelerators). In addition, it will offer support to efficient serverless applications (through integration with Knative).

3.2.2 SIPEARL

In the dynamic exascale supercomputing landscape, SIPEARL is commercializing a family of high-performance, low-power European microprocessors dedicated to supercomputing and AI inference, starting with the RHEA microprocessor.

Development is initially focused on HPC performance, as part of EuroHPC projects, to address major challenges in medical research, security, energy management and climate with a reduced environmental footprint. In addition, the aim is also to demonstrate performance across Cloud configurations that will impact the next generation cloud and HPC centres. As part of the AERO project, a number of new functionalities will be implemented and verified, such as multitenancy, virtualization/containerization, confidential compute requirements, and environment stacks enabling us to expand the Rhea ecosystem with open-source contributions developed by our partners. This can impact traditional cloud application but also addresses cloud functionality moving to HPC centres to support future hybrid HPC-cloud and federated HPC solutions.

SIPEARL works closely with its partners, as well as leading names from the scientific community who are its stakeholders, future customers and end-users.

3.2.3 KTM

KTM and the Pierer Mobility Group operate digital twins in the area of Powered-Two-Wheelers at large scales (both motorcycles and e-bikes). At the core of the digital twins is the interoperability of the on-vehicle compute elements with the back-end cloud services. In the context of AERO, KTM is exploring the compatibility and viability of EU-derived architectures for transitioning its back-end compute cloud stack. The exploitation and future commercial transition of all KTM's backend services on the EU cloud will reduce the operational costs of the whole company while paving the way for further adoption of other EU-developed technologies focusing on on-board automotive hardware.

3.2.4 RHAT

In the context of AERO, RHAT is able to bring up its existing "Red Hat build of Quarkus" commercial product to officially support AArch64 architectures for the first time both in JVM-mode and native-



mode. The upstream Quarkus project, which is mainly driven by RHAT, has been constantly gaining popularity in the software industry since its first release. Quarkus' success is attributed to its high quality and great performance, as well as the great developer experience it offers. As cloud service providers adopt ARM based servers in their offerings, it's in RHAT's best interest to ensure its clients can use Quarkus on these more cost-efficient platforms as well.

3.2.5 VOSYS

The development activities carried out by VOSYS in the AERO project provide an opportunity to the company to penetrate the Cloud computing and High Performance Computing markets with the VOSySmonitor product, and hardware accelerators virtualization technologies and security/safety solutions.

The VOSySmonitor product extensions built during the project are about new features (secure virtualization, hardware acceleration, etc.) specifically targeting the above markets. Such features will then be exposed also to the markets that today VOSySmonitor is targeting (i.e., automotive and energy), thus providing additional exploitation possibilities.

For what concerns the hardware accelerators and security/safety virtualization, VOSYS has identified an increasing interest by the Satellite Communication (SATCOM) market, who are moving from specific purpose hardware to a cloud-oriented architecture where hardware accelerators are key. As a matter of fact, the nascent IEEE Waveform Architecture for Virtualized Ecosystem (WAVE) consortium initiative⁸ proposes to run satellite hardware/software on hardware accelerated servers like the ones that AERO envisions. As a result, VOSYS considers SATCOM as a very important exploitation direction, to be addressed with accelerator virtualization solutions for FPGAs, as well as with ARM-based server solutions.

In conclusion, all VOSYS research and development activities will contribute to increasing the company know-how and dissemination activities, leading to new exploitation opportunities for design and development services.

3.2.6 SED

In the context of AERO, SED is exploring the feasibility and utility of GPU acceleration of User Defined Functions (UDF) and internal database processes of a Massive Parallel Processing (MPP) Postgres database cluster, and the potential of moving such a distributed MPP database to ARM-based architectures. As SED is a provider of business to business and business to research organisation commercial services, with main expertise in development and maintenance of open-source petabyte scale data processing solutions, the insights and developments made during the AERO project are of direct relevance for our commercial offerings. SED envisages that the developments made will allow offering to customers more energy efficient data processing solutions via switching to ARM-based server clusters, and more cost-effective scaling via offering GPU acceleration of functions directly on the database level. Last but not least, SED has recognised the possibility that some of the software developments made, in particular its new Postgres to Java interface offering unique capabilities not

⁸ <https://www.waveconsortium.org>



present in the current industry standard pJava, could be extended to an independent offering with focus on customisation and support.

Finally, SED already supports the ESA Gaia mission variability studies via UNIGE liaison and have capability to test AERO developments on real-world astronomical big-data.

3.2.7 CPLAY

As part of AERO, the oneAPI Construction Kit (OCK) is being developed and positioned as the leading open-source toolchain that will bring SYCL and OpenCL to ARM and RISC-V hardware. This work has significant opportunity for commercialisation, with CPLAY positioned to provide different levels of commercial support - onboarding, training, implementation support and bespoke development work for larger companies and organisations requiring this. CPLAY is well-positioned to provide commercial, custom optimizations for businesses. This work will in turn support the longevity of AERO and support the open-source exploitation model by growing the size of open-source, standards based heterogeneous development across RISC-V and ARM hardware.



4 Standardization Activities

This section reports the standardization activities that took place until M18. As we are only half-way through the project, this report entails also upstream work as part of these standardization activities.

4.1 RHAT

In the context of AERO, RHAT contributes to three upstream open-source projects, namely OpenJDK, Quarkus and GraalVM.

4.1.1 OpenJDK

OpenJDK is the most mature among the three projects RHAT contributes to, with its first release being in 2007. As a result, OpenJDK has already gone through a number of refinements in terms of standardization and is considered pretty stable. OpenJDK defines a clear set of rules⁹ regarding the governance, the development, and maintenance of the project. RHAT has been actively involved in the project for many years in various roles. During the AERO project, RHAT employees hold the following, related to AERO, roles in the OpenJDK project:

- Andrew Haley is appointed as a member of the governing board.
- Severin Gehwold is appointed as the lead maintainer for the OpenJDK 21 updates project¹⁰, i.e., the project ensures long term support for the JDK21 version, which is the one used for the evaluation of the project.
- Andrew Haley and Andrew John Hughes are appointed as maintainers of the OpenJDK 21 updates project.
- Martin Balao Alonso and Andrew Dinn are members of the OpenJDK Vulnerability Group¹¹.

These appointments demonstrate both the standardized processes of the OpenJDK project as well as the strong involvement of RHAT in these processes.

4.1.2 GraalVM

GraalVM (the upstream project RHAT bases its Mandrel builds on) is a newer project than OpenJDK, with its first release being in 2019. Note, however, that GraalVM is essentially the successor of the open-source research virtual machine Maxine, which dates back to 2005. That said, the code base is quite mature and stable, but the way the project is governed and maintained is still being refined. During the AERO project, the GraalVM project has made several steps in standardizing its processes as well as offering more stable releases. The most significant ones are the following:

- The GraalVM team kicked off the OpenJDK project Galahad¹², with the primary goal to contribute Java-related GraalVM technologies to the OpenJDK Community and prepare them for possible

⁹ <https://openjdk.org/census>

¹⁰ <https://openjdk.org/projects/jdk-updates/>

¹¹ <https://openjdk.org/groups/vulnerability/>

¹² <https://openjdk.org/projects/galahad/>



incubation in a JDK main-line release. Foivos Zakkak and Severin Gehwolf, working for RHAT in the AERO project, are appointed as committers to the project.

- The creation of upstream maintenance repositories^{13,14} for long term support versions that will be maintained by the community. Foivos Zakkak, working for RHAT in the AERO project, following the upstream “Call for Lead Maintainers”¹⁵ nominated himself for and was appointed the role of the Lead Maintainer for these backport repositories.
- The move to the OpenJDK release schedule, which is both better suited for long-term support and familiar to Java developers.

4.1.3 Quarkus

Quarkus is the youngest of all three projects RHAT contributes to, with its first release being in 2019. Despite being fairly young, Quarkus has done a great job in creating a community and attracting external contributors. Furthermore, Quarkus builds on top of existing standards¹⁶ making its adoption easier, which helped significantly in making it quite popular in a pretty short time. However, on the governance, development, and maintenance side there are still things that can be improved and thus Quarkus is planning to move to a foundation¹⁷. This is expected to allow Quarkus to scale further as well as strengthen its longevity by detaching it from RHAT and making it vendor-neutral. We anticipate that this move will:

- Increase transparency on how decisions are made and where the project is heading to.
- Enhance participation by doing even more things in the open.
- Improve communication.

4.2 UNIMAN

UNIMAN has been involved in two standardization activities. More specifically:

- UNIMAN participates in the oneAPI Language Special Interest Group (SIG) to contribute to the compatibility of the oneAPI programming model for the Java ecosystem. Juan Fumero (Research Fellow at UNIMAN), an Intel Innovator, represents UNIMAN and has been involved in those activities.
- UNIMAN is in close collaboration with two OpenJDK projects, namely, Babylon¹⁸ and HAT, which are initiated by Oracle to design an easy and programmer friendly way for the smooth integration of GPU compute workflows into Java. UNIMAN contributes all the lessons learnt by implementing TornadoVM and aims to ensure that TornadoVM will be compliant with the outcome of those projects.

¹³ <https://github.com/graalvm/graalvm-for-jdk21-community-backports>

¹⁴ <https://github.com/graalvm/graalvm-for-jdk17-community-backports>

¹⁵ <https://github.com/oracle/graal/issues/8935>

¹⁶ <https://quarkus.io/standards/>

¹⁷ <https://quarkus.io/blog/quarkus-in-a-foundation/>

¹⁸ <https://github.com/openjdk/babylon>



4.3 CPLAY

CPLAY is committed to open standards and is a member of the Khronos Group and the SYCL Working Group.

Through AERO, CPLAY aims for DPC++/OCK to become the go-to open-source toolchain for SYCL and OpenCL acceleration on AArch64 and RISC-V. Towards this end, CPLAY is pursuing multiple open-source activities:

- Generic OpenCL support (via OCK) for AArch64, RISC-V CPUs and RISC-V Vector (RVV) accelerator.
- SYCL support on AArch64 and RISC-V via 2 paths; through the aforementioned OpenCL and through the SYCL Native CPU target.
- Improve targeting AArch64 and RISC-V through adding and/or improving support for cross-compilation.
- Maintain RISC-V compatibility with different RVV versions for portability with previous implementations/releases.
- Ensure OpenCL/SYCL conformance through frequent/nightly CTS testing.



5 Summary

This deliverable has outlined the exploitation strategy and commercial roadmap of AERO. By taking advantage of an open-source, shared-benefits exploitation model, AERO can contribute effectively to the economy of Europe via open-source contributions. More than this, industry partners can upsell where appropriate for specific optimizations and implementations.

A SWOT analysis has been carried out to examine the current strengths and weaknesses of the project, before conducting a preliminary market analysis that will also feature in the M36 report. It is apparent that AERO operates with an effective understanding of the market, with a large existing audience of Industrial Cloud Stakeholders and more that can be targeted.

Last but not least, AERO partners have been active in the areas of standardization and work upstream. These include CPLAY's membership of the Khronos Group and open-source work with OCK, RHAT's work with OpenJDK, GraalVM, and Quarkus, and UNIMAN's participation in the oneAPI Language SIG.