This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101069732

aerOS
EUROPEAN IOT-EDGE-CLOUD

D5.5 – Technical evaluation, validation and assessment report (1)

Deliverable No.	D5.5	Due Date	31-AUG-2024
Type	Report	Dissemination Level	Public (PU)
Version	1.0	WP	WP5
Description	Document summarising the approach and methodology that the aerOS consortium has adopted for the software integration of the different components of the ecosystem, as well as the different KPIs that will be evaluated at the end of the project.		

Copyright

The aerOS consortium consists of the following 27 partners:

UNIVERSITAT POLITÈCNICA DE VALÈNCIA	ES
NATIONAL CENTER FOR SCIENTIFIC RESEARCH "DEMOKRITOS"	EL
ASOCIACION DE EMPRESAS TECNOLOGICAS INNOVALIA	ES
TTCONTROL GMBH	AT
TTTECH COMPUTERTECHNIK AG (*third linked party*)	AT
SIEMENS AKTIENGESELLSCHAFT	DE
FIWARE FOUNDATION EV	DE
TELEFONICA INVESTIGACION Y DESARROLLO SA	ES
ORGANISMOS TILEPIKOINONION TIS ELLADOS OTE AE - HELLENIC TELECOMMUNICATIONS ORGANIZATION SA	EL
EIGHT BELLS LTD	CY
INQBIT INNOVATIONS SRL	RO
FOGUS INNOVATIONS & SERVICES P.C.	EL
L.M. ERICSSON LIMITED	IE
SYSTEMS RESEARCH INSTITUTE OF THE POLISH ACADEMY OF SCIENCES IBS PAN	PL
ICTFICIAL OY	FI
INFOLYSIS P.C.	EL
PRODEVELOP SL	ES
EUROGATE CONTAINER TERMINAL LIMASSOL LIMITED	CY
TECHNOLOGIKO PANEPISTIMIO KYPROU	CY
DS TECH SRL	IT
GRUPO S 21SEC GESTION SA	ES
JOHN DEERE GMBH & CO. KG*JD	DE
CLOUDFERRO S.A.	PL
ELECTRUM SP ZOO	PL
POLITECNICO DI MILANO	IT
MADE SCARL	IT
NAVARRA DE SERVICIOS Y TECNOLOGIAS SA	ES
SWITZERLAND INNOVATION PARK BIEL/BIENNE AG	CH

Disclaimer

Authors

Name	Partner	e-mail
Ignacio Lacalle	P01 UPV	iglaub@upv.es
Rafael Vaño	P01 UPV	ravagar2@upv.es
Raúl San Julián	P01 UPV	rausanga@upv.es
Salvador Cuñat	P01 UPV	salcuane@upv.es
Raúl Reinosa	P01 UPV	rreisim@upv.es
Harilaos Koumaras	P02 NCSRD	koumaras@iit.demokritos.gr
Vasilis Pitsilis	P02 NCSRD	vpitsilis@iit.demokritos.gr
Andreas Sakellaropoulos	P02 NCSRD	asakellaropoulos@iit.demokritos.gr
Roger Briz	P03 INNO	rbriz@innovalia.org
Andreas Locatelli	P04 TTC	andreas.locatelli@ttcontrol.com
Francesca Flamigni	P04.1 TCAG	francesca.flamigni@tttech.com
Amparo Sancho-Arellano	P05 SIEMENS	amparo.sancho-arellano@siemens.com
Philippe Buschmann	P05 SIEMENS	philippe.buschmann@siemens.com
Renzo Bazan	P05 SIEMENS	renzo.bazan.ext@siemens.com
Ioannis Makropodis	P10 IQB	giannis.makropodis@inqbit.io
Ioannis Chouchoulis	P10 IQB	giannis.chouchoulis@inqbit.io
Vasiliki Sampazioti	P10 IQB	vasiliki.maria.sampazioti@inqbit.io
Katarzyna Wasielewska-Michniewska	P13 SRIPAS	Katarzyna.Wasielewska@ibspan.waw.pl
Przemysław Hołda	P13 SRIPAS	Przemyslaw.Holda@ibspan.waw.pl
Wiesław Pawłowski	P13 SRIPAS	Wieslaw.Pawlowski@ibspan.waw.pl
Tarik Taleb	P14 ICTFI	tarik.taleb@ictficial.com
Tarik Zakaria Benmerar	P14 ICTFI	tarik.benmerar@ictficial.com
Amine Taleb	P14 ICTFI	amine.taleb@ictficial.com
Vaios Koumaras	P15 INF	vkoumaras@infolysis.gr
Nikolaos Gkatzios	P15 INF	ngkatzios@infolysis.gr
Vasileios Mavrikakis	P15 INF	vmavrikakis@infolysis.gr
Eduardo Garro	P16 PRO	egarro@prodevelop.es
Álvaro Martínez Romero	P16 PRO	amromero@prodevelop.es
Kyriacos Orphanides	P17 ECTL	Kyriacos.orphanides@eurogate-limassol.com
Alessandro Cassera	P17 ECTL	Alessandro.Cassera@eurogate-limassol.com
Jon Egaña	P20 S21Sec	jon.egana@thalesgroup.com
Oscar Lopez	P20 S21Sec	oscar.lopez-perez@thalesgroup.com
Francesco Dellino	P25 MADE	francesco.dellino@made-cc.eu

Giulio Centi	P25 MADE	giulio.centi@made-cc.eu
Lucie Stutz	P27 SIPBB	Lucie.stutz@sipbb.ch

History

Date	Version	Change
26-APR-2024	0.1	ToC created and announced in plenary telco
31-MAY-2024	0.2	First content sent by partners and included in the document
28-JUN-2024	0.5	Second content sent by partners
19-JUL-2024	0.61	Merged content by editors and version submitted for Internal Review
31-JUL-2024	1.0	Final version submitted considering IR comments

Key Data

Keywords	KPI, evaluation
Lead Editor	P16 – PRO
Internal Reviewer(s)	UPV, TID

Table of contents

List of tables

List of figures

1. About this document

The aim of deliverable D5.5 is to provide an update of the evaluation process started in M9, and initially reported in Section 3 of D5.2. Hence, the ultimate goal of D5.5 is to iterate over the basis for a quantitative assessment of aerOS both from the standpoint of its developers through the technical KPIs, as well as its users and stakeholders, through the pilot KPIs and impact KPIs.

1.1. Deliverable context

Table 1 Deliverable context

Item	Description
Objectives	D5.5 contributes to **all objectives** set for aerOS by defining quantitative KPIs evaluation, as well as qualitative evaluation by means of the KeVIs envisioned in the action.
Work plan	D5.5 directly maps to T5.3 and T5.4 (mainly) as both are in charge of carrying out the evaluation results that are generated in the project. Consequently, D5.5 depends also in the work done by the aerOS development team along WP3 and WP4 tasks, the pilot integration works performed under T5.2, and the impact from communication, dissemination, standardisation, and exploitation team along WP6.
Milestones	There are not any specific milestones associated to the delivery of D5.5. However, it contributes to the achievement of: • MS 2 Use cases and requirements • MS 3 Components defined
Deliverables	The deliverables that are related with D5.5 are: • D5.2 - Integration, evaluation plan and KPIs definition (2) - (M18) • D5.6 - Technical evaluation, validation and assessment report (2) - (M36)
Risks	The risks that partially represent a background for D5.5 is: • Outcome does not meet expectations set by the stakeholders.

1.2. The rationale behind the structure

This deliverable is divided into seven sections. In the Introduction, objectives of the deliverable are presented. In the KPI Evaluation Methodology section, the notion of KPI as well as the evaluation methodology agreed, including the portioned into dimensions, is replicated (with some additions) from D5.2. Next, the technical evaluation of aerOS itself, and its surrounding components are provided. Section 4 constitutes the initial evaluation of the results from the pilots, and Section 5 recaps the impact evaluation conducted along WP6. A qualitative analysis from the KeVIs perspective is provided in Section 6, while the final section is reserved for conclusions

1.3. Deviation and corrective actions

This deliverable, initially to be submitted in M21, was moved to M24 in order to describe in detail the KPIs and their measurement metrics. This decision was based on the positive feedback received by the project reviewers during the mid-term review in April 2024 (M20).

2. KPI evaluation methodology

The evaluation methodology adopted has been inspired upon ideas derived from the [ASSIST-IoT H2020 Project.] This decision was made since that project analysed Next-Generation IoT platforms for different focus areas and application domains. It is thus very similar to the structure of aerOS, where five large scale pilots exist in addition to several Open Call projects, in various application domains of the Cloud-Edge-IoT continuum.

In particular, the aerOS KPI evaluation methodology follows a hierarchical approach, in which KPIs (either derived from project's Grant Agreement, requirements gathered in WP2 deliverables, or from insights obtained among different pilot group partners) are grouped by the definition of dimensions and fields of measurement, in accordance with project aspects. This hierarchical approach aims at providing a more tangible assessment of the project's success factors, comparing the result with respect to the Grant Agreement commitments (i.e., accomplished, not accomplished, partially accomplished), system gaps, as well as stakeholders feedback and satisfaction, as explained next.

2.1. KPI definition, dimensions, and fields

Technically, a KPI states for Key Performance Indicator. Thus, it is a type of performance measurement, which is done against a predefined set of values, called indicators. In aerOS it has been agreed that only quantitative indicators are used, meaning that each KPIs value shall be a number (or a percentage over a baseline). In general, data collection for KPI evaluation will be conducted through benchmark testing, desk research, on-line questionnaires or through one-to-one interviews. The latter will be conducted with coordinators of all pilots, as well as task leaders. In those cases, in which a KPI refers to users' satisfaction, a value from a set of numbers from 1 to 5, where 1 would represent complete dissatisfaction and 5 would represent complete satisfaction, will be assumed.

In particular, the KPI will be measured directly (for example messages/second, throughput in Mbps, or user satisfaction on the Likert scale) and it will be expressed as a fulfilment percentage of the prospected KPI target as follows:

- KPI value = no achievement → **KPI score = 0%**
- no achievement < KPI value < target → **0% < KPI score < 100%**
- KPI value >= target → **KPI score = 100%**

This approach will allow to interpret the value of an indicator against the original target set by aerOS partners.

To assess an overall aerOS evaluation, the KPI scores are grouped into **Field scores**, which are calculated as the average of all KPI scores for that field (i.e., aerOS service fabric KPIS, aerOS Data Fabric KPIs, Pilot 1 KPIs, Pilot 2 KPIs, Communication KPIs, Dissemination KPIs, etc.), and, in turn combined into a **Dimension score**, which is defined as the average of Field scores for that dimension, as presented in the two equations below.

$$FieldScore = \frac{\sum_{i=1,\dots,N} KPIscore_i}{N} \quad (1)$$

$$DimensionScore = \frac{\sum_{i=1,\dots,N} Fieldscore_i}{N} \quad (2)$$

Therefore, if all KPIs reach the target, which would manifest as all KPI scores being 100%, then Field score is also going to be 100%. Similarly, if all Field scores within the same dimension reach a 100% target, then Dimension score will also reach 100%. This approach would provide an indication of achievements for specific indicators, as well as overall results obtained by the aerOS project as a whole.

For aerOS, it has been decided to define three main dimensions (i) Technical aerOS platform KPIs; (ii) aerOS Pilot KPIs; and (iii) aerOS Impact KPIs. Several groups have been defined per dimension. The following table maps each dimension and group, with the different tasks of the project where actual assessment will take place.

Table 2 aerOS Dimensions, Fields and achieving tasks

Dimension	Filed of evaluation	Achieving tasks
1. Technical aerOS platform KPIs	1.1 aerOS network and compute fabric	T3.1 / T3.3
	1.2 aerOS Data Fabric	T4.1 / T4.2
	1.3 aerOS Service Fabric	T3.3
	1.4 aerOS aerOS cyber security and Trust components	T3.4 / T4.5
	1.5 aerOS self-* and monitoring	T3.5
	1.6 aerOS decentralised AI	T4.3
	1.7 aerOS common API	T3.2
	1.8 aerOS management framework	T4.6
	1.9 aerOS Embedded analytics	T4.4
	1.10 Stakeholder/user satisfaction/OpenCall	T1.1 / T1.4
2. aerOS Pilot KPIs	2.1 Pilot 1 Data-driven cognitive production lines	T5.2 / T5.4
	2.2 Pilot 2 Containerised edge computing near renewable energy sources	T5.2 / T5.4
	2.3 Pilot 3 High performance computing platform for connected and cooperative mobile machinery	T5.2 / T5.4
	2.4 Pilot 4 Smart edge services for the port continuum	T5.2 / T5.4
	2.5 Pilot 5 Energy Efficient, health safe and sustainable smart buildings	T5.2 / T5.4
	2.6 Overall pilots' engagement	T5.4
3. aerOS Impact KPIs	3.1 Communication	T6.1
	3.2 Dissemination	T6.2
	3.3 Stanardisation	T6.3
	3.3 Exploitation and business models	T6.4

2.2. KVI evaluation methodology

To supplement the existing KPI-based, quantitative strategy in aerOS, a qualitative evaluation for a societal value-driven approach to technology development will be used, built upon the idea of Key Value Indicators (KVIs). KVIs are markers of a pertinent social Key Value (KV) that is either affected or enabled by upcoming technology, particularly the new services offered by future networks [1]. These networks are expected to enable a wide range of innovative and advanced use cases and application domains, which will address significant societal demands and provide added value in several ways. Humans and machines will be able to connect in new ways that should benefit individuals, communities, and businesses. In addition to monitoring and reducing risks, these benefits must be maximized. In this deliverable, the methodology to be followed is described, as well as the initial analysis and results. Final results will be included in deliverable D5.6.

The methodology starts with the United Nations (UN) Sustainable Development Goals (SDGs) [2]. As seen in Figure 1, these goals offer a flexible framework for target priorities, as represented by the three related domains of environmental, social, and economic sustainability. But because these are state-specific, they must be interpreted for the ICT sector to assess their effects. Accordingly, the first step of the methodology is to define a set of societal KVs that are important to people and society and can be directly or indirectly addressed by future networks. Some examples of such KV are: Economical Sustainability and innovation, Digital Inclusion, Personal Freedom, Trust. However, the list is not restricted as may depend on specific cases.

Figure 1 Illustration of the UN SDGs ordered in three areas according to [2]

The second step is to attach KVIs to each KV. KVIs are primarily concerned with determining the societal value that results from technological advancements and the context in which they must be used. KVIs are based on a "societal readiness" approach that examines, rather than if society is ready for a new technology, if technologies are ready to be incorporated into society. A technology is considered important under the societal preparedness paradigm if it is useful, efficient, fosters practice innovation, and helps bringing systemic change for the good of society. Compared to KPIs, knowledge of human-related elements is enhanced by KVIs, which may necessitate discussions and the emergence of creative ideas. They are comparable to quality RRDs (Quality Requirements Documents) that support a "top-down" strategy, in which project goals must be defined before measurements are applied. KVIs need to be a set of quantifiable amounts or prerequisites that offer an approximation of an impacted KV in some way. As a result, targets should be able to be set using KVIs, such as the quantity of service users who satisfy a condition or the perceived satisfaction using a service.

The third step of the methodology involves the KV enablers for each KV. This step analyses the factors that determine a use case's usage, i.e., availability and popularity. This step defines the main elements that would either prevent the further construction of a use case or spread it and scale up the enabled value. These variables typically have to do with meeting the end user's expectations in terms of technical fulfilment, service coverage, ecosystem adaptation, and value proposition. While the precise exchange rate is not stated, better KV enablement should eventually produce value.

The final forth step aims at determining the KPIs that affect the related KVIs. This might not always be feasible or meaningful, depending on the KVI. When using KPIs, it should be evident how and to what degree a particular KVI is enhanced, without mentioning the exchange rate. To enable KVs, future networks can use the KPIs to set technical numerical targets. Ultimately, this stage ought to offer a foundation for a technological design that is motivated by values. All four steps are illustrated in Figure 2 [1].

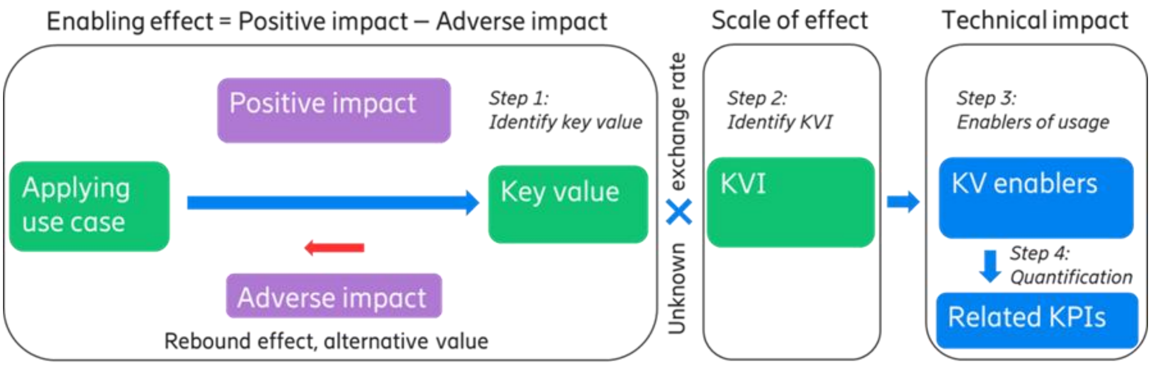

Figure 2 The four steps of the adopted KVIs methodology [1]

In the scope of aerOS, the aforementioned methodology is applied to study each of the identified use cases and conclude to a specific set of KVs, KVIs, KV enablers and KPIs per use case. The target is to conclude to a complete qualitative analysis of the performance gains achieved against use case performance indicators.

3. Technical KPIs for aerOS

The purpose of this dimension is to address the existence, functionality and availability of the technical components and features defined and implemented in the aerOS platform. To analyse the technical parameters, the assessment is builds upon the results and outcomes from integration, and testing activities. The main software modules under development to be evaluated, are the following:

1. aerOS network and compute fabric
2. aerOS data fabric
3. aerOS service fabric
4. aerOS cybersecurity and trust components
5. aerOS self-* and monitoring
6. aerOS decentralised AI
7. aerOS common API
8. aerOS management framework
9. aerOS embedded analytics
10. Stakeholder's satisfaction

3.1. aerOS network and compute fabric

3.1.1. KPI 1.1.1 Response time for the orchestration of IoT applications (KVI-1.1)

Table 3 KPI 1.1.1 Response time for the orchestration of IoT applications (KV-1.1)

KPI ID number .	KPI 1.1.1
KPI Name	Response time for the orchestration of IoT applications
Description	This KPI measures the time the orchestration system takes to achieve the target state of the blueprint of the IoT applications
Motivation	Whether achieving the initial state or transitioning states due to external conditions changes, the orchestration system should provide responsiveness for the IoT applications. A less responsive system would hinder the usefulness of such autonomous service for the end user and makes it less reactive to changing conditions.
Target value	<15% baseline
Prerequisites	aerOS installation ready in the concerned domains. aerOS installation implies here that, at least, self-awareness and self-orchestrator elements are functional in several IEs, that these (IEs) are organized in one (or more) domain(s) and that the HLO is capable of receiving implementation blueprints and allocating computing workloads.
aerOS components (task)	HLO (T3.3), LLO (T3.3), self-awareness (T3.5), self-orchestrator (T3.5)
Evaluation means	The evaluation process leverages the status of Data Fabric service components to monitor the deployment time effectively. This monitoring is crucial for understanding the time taken for various components to become operational. Additionally, the deployment time can be assessed more accurately from the Management Portal, providing a precise measure of response time from the user's perspective.

	At M24, given the current state of the demonstrator, the deployment time from the Management Portal that excludes the latency introduced by the HLO (High-Level Orchestration) AI service is being measured. This focused approach allows to gather baseline data on deployment efficiency without the additional complexity of AI processing delays. For D5.6, a more comprehensive measurement mechanism will be implemented. This advanced system will encompass all aspects of the deployment process, including HLO AI latency and other potential delays, ensuring a thorough and accurate evaluation of the deployment time across the entire service framework. This holistic approach will provide a detailed understanding of the deployment dynamics, enabling optimizations and enhancements to improve overall efficiency and user experience.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*10 s*[1]	***5 s (333%)***	*N/A*
Outcome elaboration (M24)	Performances based on the aerOS mid-review demonstration have been measured from the user perspective (https://youtu.be/UV4mnN4CrwI?si=gy1mRBYiEqeamdNS).		

3.1.2. KPI 1.1.2 Open-source components for aerOS to deploy and manage applications spanning the continuum (KVI-1.2)

Table 4 KPI 1.1.2 Open-source components for aerOS to deploy and manage applications spanning the continuum (KVI-1.2)

KPI ID number .	**KPI 1.1.2**
KPI Name	**Open-source components for aerOS to deploy and manage applications spanning the continuum**
Description	The KPI determines the number of components that aerOS generates and are able to deploy and manage applications in the continuum that have been shared with external communities through open-source contributions.
Motivation	This KPI is important to measure the impact aerOS has on the technological ecosystem, allowing the support of new technological business models and third-parties exploitation resulting from the project innovative work.
Target value	3
Prerequisites	The components are aerOS-created, or aerOS-enhanced, and are ***available for inspection, download, reuse and replication*** by the community outside of the project. A pre-requisite for final acceptation of this KPI is also the completion of the Open-Source Strategy that has been defined during the latest period of the project, and that will be put in place from M24 to M36 of aerOS.
aerOS components (task)	Any aerOS component subject of being open-source licensed (all WP3-WP4, and T6.4)

[1] The baseline is taken from the worst case in usual IoT applications in the literature [3]..

<table>
<tr><td>Evaluation means</td><td colspan="3">The number of public repositories related aerOS. With LICENSE and NOTICE files and with proper documentation for its usage.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>N/A</td><td>0 (0%)</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">aerOS has already contributed to diverse Open-Source repositories. More than 15 issues to public repos of open-source components to enhance the products with aerOS flavor.
• Orion-LD via various issues, and request for contributions to ETSI CIM.
• IOTA issues.
• Issue to KubeEdge repository.
• Contributions to Morph-kgc suite.
Therefore, although not aerOS-created components, the contribution from the project (that uses those technologies) is already evident.

About open source components of aerOS, the project has made huge efforts in the documentation and in creating an installation guide. The OSS components that are included in the Installation Guide, and that are custom provisions by aerOS are not yet public (therefore, not OSS yet), have been made available, notwithstanding, to Open Call #1 projects via the usage of specific GitLab tokens to access the public common deployments of aerOS.

Regarding the global public, repositories (out of the global list in the private GitLab) are being tagged as potential open-source components albeit no formal agreement has been reached yet.

This is the global list of components of aerOS. It is expected that most of them will be released as Open Source, with a permissive License for the Community. A formal, final list of the OSS components of aerOS will be submitted in D5.6, altogether with the proper Github and Zenodo links:
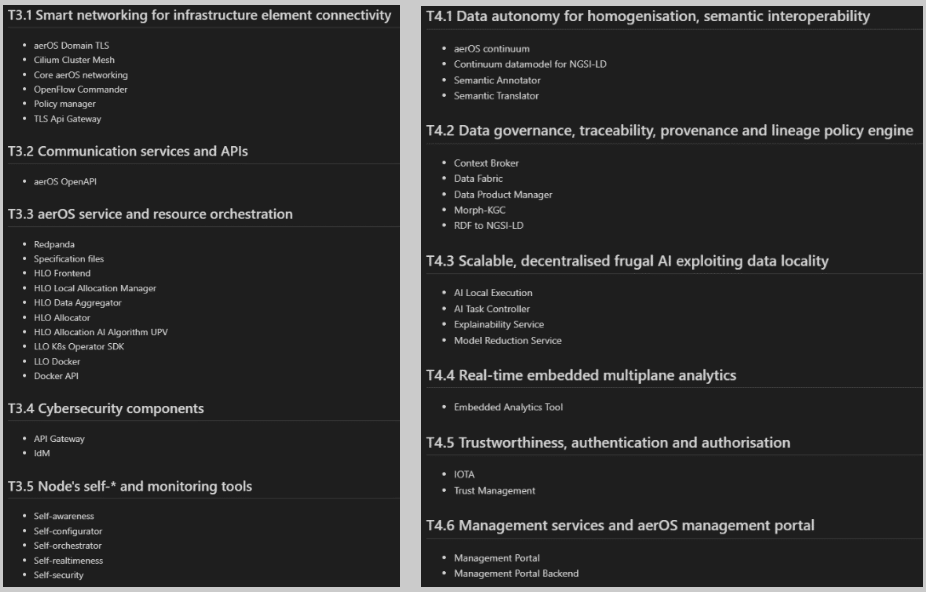
T3.1 Smart networking for infrastructure element connectivity
• aerOS Domain TLS
• Cilium Cluster Mesh
• Core aerOS networking
• OpenFlow Commander
• Policy manager
• TLS Api Gateway
T3.2 Communication services and APIs
• aerOS OpenAPI
T3.3 aerOS service and resource orchestration
• Redpanda
• Specification files
• HLO Frontend
• HLO Local Allocation Manager
• HLO Data Aggregator
• HLO Allocator
• HLO Allocation AI Algorithm UPV
• LLO K8s Operator SDK
• LLO Docker
• Docker API
T3.4 Cybersecurity components
• API Gateway
• IdM
T3.5 Node's self-* and monitoring tools
• Self-awareness
• Self-configurator
• Self-orchestrator
• Self-realtimeness
• Self-security
T4.1 Data autonomy for homogenisation, semantic interoperability
• aerOS continuum
• Continuum datamodel for NGSI-LD
• Semantic Annotator
• Semantic Translator
T4.2 Data governance, traceability, provenance and lineage policy engine
• Context Broker
• Data Fabric
• Data Product Manager
• Morph-KGC
• RDF to NGSI-LD
T4.3 Scalable, decentralised frugal AI exploiting data locality
• AI Local Execution
• AI Task Controller
• Explainability Service
• Model Reduction Service
T4.4 Real-time embedded multiplane analytics
• Embedded Analytics Tool
T4.5 Trustworthiness, authentication and authorisation
• IOTA
• Trust Management
T4.6 Management services and aerOS management portal
• Management Portal
• Management Portal Backend</td></tr>
</table>

3.1.3. KPI 1.1.3 Usage of 5G native APIs (3GPP NEF and CAPIF) (KVI-1.3)

Table 5 KPI 1.1.3 Usage of 5G native APIs (3GPP NEF and SEAL) (KVI-1.3)

KPI ID number .	**KPI 1.1.3**		
KPI Name	**Usage of 5G native APIs (3GPP NEF and CAPIF)**		
Description	5G Native APIs that have been specified by 3GPP allows the tight integration of services and applications in order to improve their performance and features, such as by retrieving information of the QoS level of the network and the location of the user.		
Motivation	The use of 5G native APIs can significantly enhance the performance of a service/application, providing additional context information and network awareness. Therefore, assessing the use of native 5G APIs as a KPI is important because it denotes the disruptive innovation of the developed services and applications.		
Target value	>50% aerOS scenarios using 5G network		
Prerequisites	Functional aerOS domain and aerOS APIs exposed		
aerOS components (task)	Ingress (T3.1), TLS (T3.1), OpenAPI (T3.2), HLO (T3.3), API Gateway (T3.4), Context Broker (T4.2)		
Evaluation means	Out of all the aerOS scenarios, the following two are going to use cellular network: • Pilot 4 – Predictive maintenance • Pilot 5 – Smart buildings As long as any one of those two scenarios show that the OpenCAPIF is integrated in, the KPI will be considered as fulfilled. To carry out the evaluation, there are three options: (i) reporting tools of OpenCAPIF will be used; (ii) exported report on discovered (aerOS) APIs; and (iii) POSTMAN endpoints with OpenCAPIF acting as consumer, getting all aerOS registered APIs.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***1/2 scenarios deployed, but not integrated (50%)***	*N/A*
Outcome elaboration (M24)	For M24, Pilot 5 has successfully deployed one OpenCAPIF instance in their development environment. The first two screenshots below (k9s services running on the domain, and capif service logs), as well as in the third figure show how the performance of the CAPIF is monitored in a Grafana instance.		

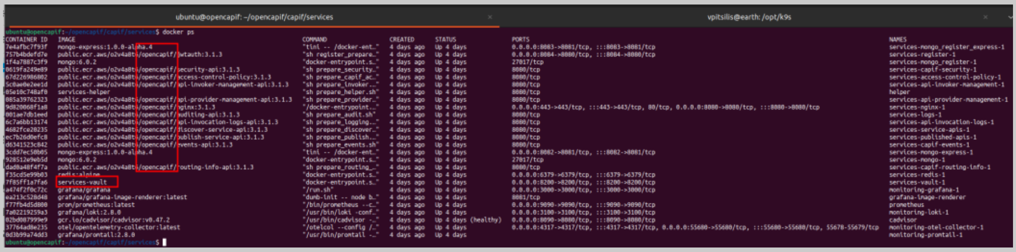

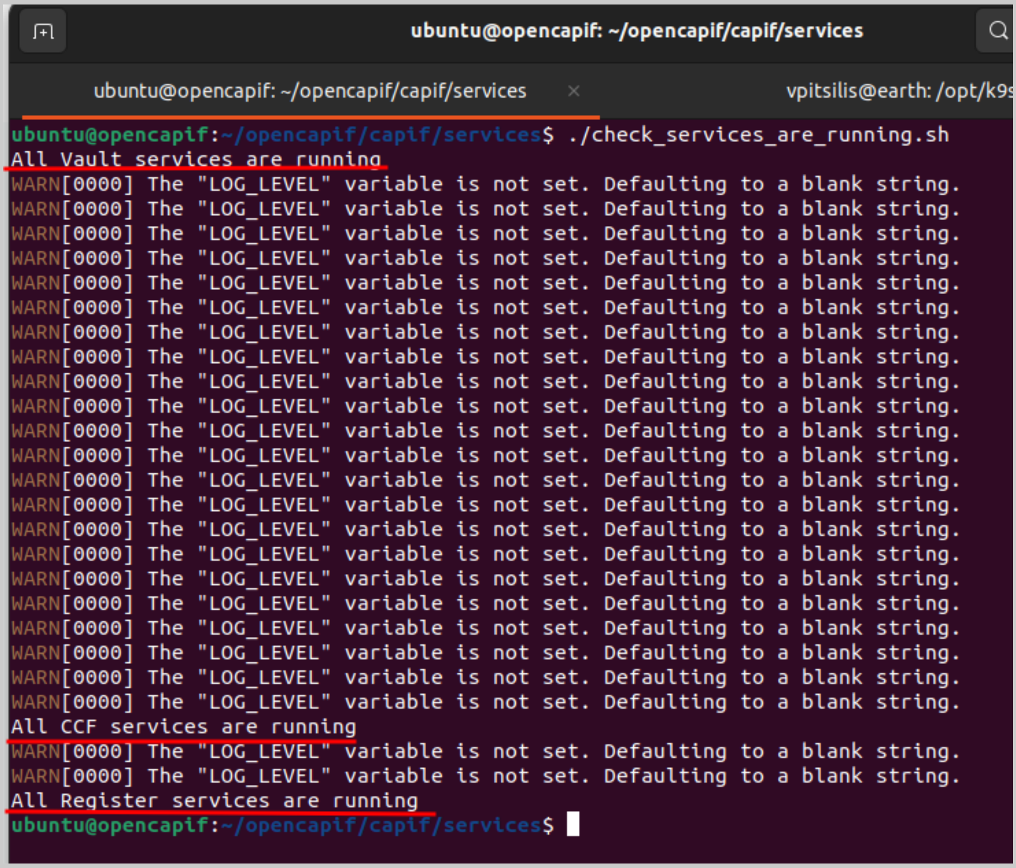

In M36, it will be connected and integrated with the aerOS APIs from the pilot pre-production environment.

3.1.4. KPI 1.1.4 Usage of TSN (KVI-1.4)

Table 6 KPI 1.1.4 Usage of TSN (KVI-1.4)

KPI ID number .	**KPI 1.1.4**		
KPI Name	**Usage of TSN**		
Description	This KPI measures the adoption rate of Time-Sensitive Networking (TSN) in more than 50% of the applicable scenarios within the project. TSN is a set of standards designed to improve the reliability, latency, and synchronization of standard Ethernet networks. The goal is to quantify the extent to which TSN is being utilized in scenarios where real-time, deterministic communication is critical.		
Motivation	The integration of TSN is crucial for scenarios that demand high levels of network determinism and reliability, such as in industrial automation, real-time control systems, and applications requiring precise synchronization. By targeting a greater than 50% adoption rate in relevant scenarios, this KPI encourages the advancement of network infrastructure towards more robust, latency-sensitive, and synchronized communication capabilities. This, in turn, supports the overall efficiency, safety, and performance of the systems relying on aerOS.		
Target value	>50% scenarios		
Prerequisites	To effectively evaluate the usage of TSN, it is essential to have a TSN-enabled infrastructure in place. This includes ensuring that the network infrastructure is equipped with TSN-compatible switches, routers, or other network devices. Additionally, it is crucial to have TSN-aware applications that can fully utilize TSN's capabilities.		
aerOS components (task)	Ingress (T3.1), TLS (T3.1)		
Evaluation means	The evaluation of this KPI aims to measure the number of scenarios where TSN is deployed. The process involves identifying and counting scenarios in which TSN-capable networking hardware such as switches or routers is utilized. Additionally, a pivot scenario that highlights a successful TSN integration will be examined. This scenario will provide valuable insights into challenges faced, benefits gained, and lessons learned.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***1 scenario (9%)***	*N/A*
Outcome elaboration (M24)	In Pilot 1 scenario "AGV swarm zero break-down logistics & zero ramp-up safe PLC reconfiguration for lot-size 1 production," a TSN-capable switch was successfully integrated to address latency and jitter issues in the communication between connected devices. The goal was to optimize the performance of the scenario. The TSN-capable switch used, was the Multiport Time Sensitive Networking (MTSN) Kit from SoCe. This switch offers scheduled, best-effort and reserved traffic, along with Time Aware Shaper and Credit Based Shaper functionalities. The setup involved connecting an		

	Industrial PC (SIMATIC Microbox PC) with an Intel Core i5-6442EQ processor and two Raspberry Pis 4 with a Quad core 64-bit ARM-Cortex A72 processor to the MTSN Kit switch. The switch, in turn, was connected to a router that facilitated internal traffic routing within the scenario. The figure below illustrates the overall setup, showcasing the connection between the devices and router. 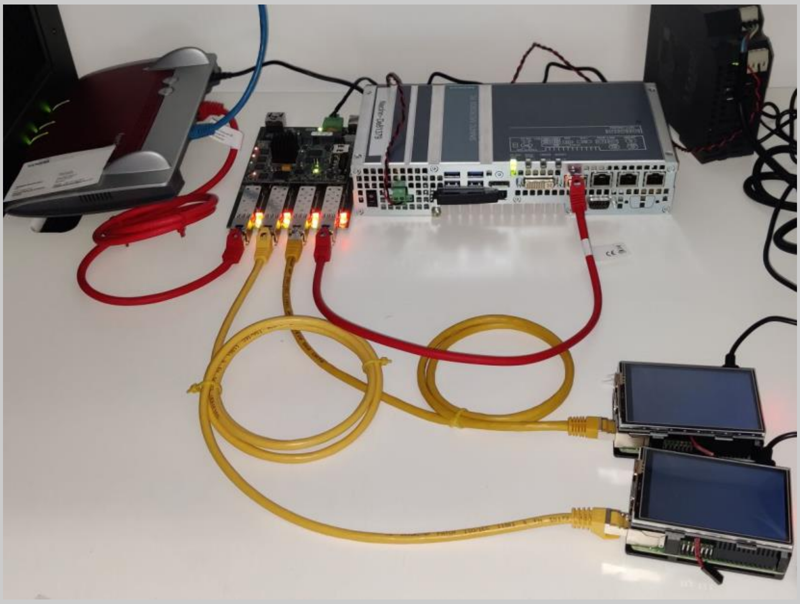 Additionally, a web-accessible administrator interface allowed for easy configuration and adjustment of the traffic between devices, as depicted in the following figure. 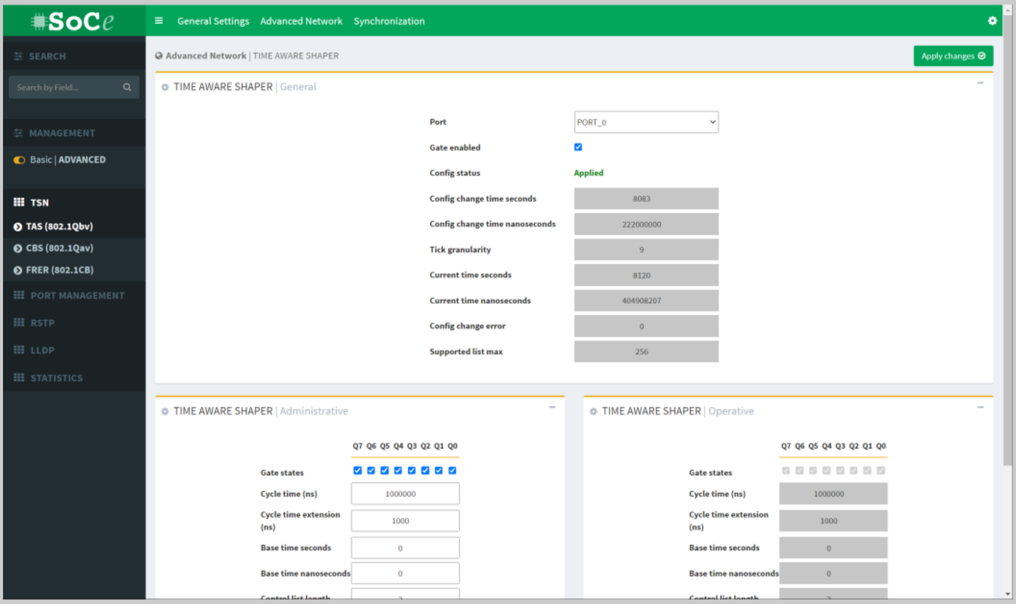

3.1.5. KPI 1.1.5 Number of old equipment units turned on actionable aerOS nodes

Table 7 KPI 1.1.5 Number of old equipment units turned on actionable aerOS nodes

KPI ID number .	**KPI 1.1.5**
KPI Name	**Number of old equipment units turned on actionable aerOS nodes**

Description	Devices that are incorporated into the aerOS continuum enabled by the Meta-OS		
Motivation	The operation of high-performance algorithms and highly efficient data transaction mechanisms depends on edge, IoT and cloud devices to orchestrate effectively the different services of the hyper distributed application workflows.		
Target value	20		
Prerequisites	aerOS self-components installed in an IE, and the IE integrated in a domain providing information and being able to accept workloads.		
aerOS components (task)	The minimum aerOS core services, namely, LLO (T3.3), Self-* (T3.5), Context Broker (T4.2)		
Evaluation means	KPI leader (INNO) will address the different aerOS pilots about the old equipment units that have become in new aerOS nodes. An internal description per pilot, including technical specifications, will be included within this table. The evidence of their use will be presented in D5.4.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***46 identified. The 6 IEs from Pilot5 up-and-running (30%)***	*N/A*
Outcome elaboration (M24)	The number of devices varies significantly between pilots: ***Pilot 1**: 19 devices* • *SIPBB (Scenario1)* – ***6 devices**:* Quality control, Manual workstation, Smart conveyor, AGV, PCB THT-Soldering, SMC AMS). This is the hardware needed to integrate into aerOS and whose data must be collected. • INNOVALIA (Scenario2) **- 0 devices**: There is no old equipment expected now to be turn on thanks to aerOS within the scope of the project • SIEMENS (Scenario3) **- 10 devices**: 4 existing AGVs and 1 industrial PC will be transformed into aerOS devices. Regarding the testbed, 4 existing Raspberry Pi's and 1 TSN switch will be also converted into aerOS devices. • MADE&POLIMI (Scenario4) – **3 devices**: MADE relies on their pre-existing on premise server. POLIMI relies on an existing Industrial PC and onboard PC of the AGV (the AGV was recently refurbished – addition of an SSD, new batteries, and upgrade from ubuntu 18.04). Considering AGV as old equipment there are 3. ***Pilot 2: 21 devices:*** Since the Cloud is public and out of context, is not considered. Within the Edge, decommissioned servers from another project are used. • 1x router node with following specification: 2 processors 12 core 2.3 GHz, 128 GB RAM, 2 x SSD Boot Disk 1TB • 1x control plane node with following specification: 2 processors 12 core 2.3 GHz, 256 GB RAM, 2 x SSD Boot Disk 1TB • 17x compute nodes with following specification: 2 processors 12 core 2.3 GHz, 256 GB RAM, 2 x SSD Boot Disk 120GB		

	• 2x storage nodes with following specification: 2 processors 12 core 2.3 GHz, 256 GB RAM, 2 x SSD Boot Disk 120GB ***Pilot 3:*** *Still to be determined* ***Pilot 4 - 0 devices*****:** All the hardware equipment for Pilot 4 has been procured as new, so there is no equipment turned on thanks to aerOS. ***Pilot 5 – 6 devices:*** 3 Up-boards used as IO GWs (7 years old based on the age of the *processor*. 2 NUCs (4 years old) and 1 HP-DL380 (5 years old)

3.1.6. KPI 1.1.6 Consistency of deployment compared to app blueprints

Table 8 KPI 1.1.6 Consistency of deployment compared to app blueprints

KPI ID number .	**KPI 1.1.6**		
KPI Name	**Consistency of deployment compared to app blueprints**		
Description	This KPI measures how consistent the consecutive deployments of the application compared its specified blueprint		
Motivation	The orchestration system is an autonomous system. It is important that this system keeps consistent its automatic deployments with respect to the blueprint and doesn't require manual oversight.		
Target value	>95%		
Prerequisites	aerOS installation ready in the concerned domains		
aerOS components (task)	HLO (T3.3), LLO (T3.3), Management Portal (T4.6).		
Evaluation means	Manual test of the applications by pilots and observability tools, such as K9s, for deployments verification. For now, only manual test has been done on the application mid-review demonstration. With continuous integration of the pilots, more data will be collected, and more observability tools will be integrated as part of the test.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***100% (105%)***	*N/A*
Outcome elaboration (M24)	With the limited tests conducted within the consortium (https://youtu.be/UV4mnN4CrwI?si=gy1mRBYiEqeamdNS) no consistency issues were noted. The TOSCA examples used for the demonstration are shown below:		

```yaml
description: This is a testing TOSCA composed of X nodes
node_templates:
  nginx:
    artifacts:
      application_image:
        file: nginx
        repository: registry.gitlab.aeros-project.eu
        type: tosca.artifacts.Deployment.Image.Container.Docker
    interfaces:
      Standard:
        create:
    implementation: application_image
          inputs:
            cliArgs:
            - -forwarding: ''
            - -brokerId: CFDomain
            envVars:
            - BROKER_ID: CFDomain
            - OTHER_ENV_VAR: another value
    requirements:
    - network:
        properties:
          ports:
            port1:
              protocol:
              - tcp
              source: 1026
            port2:
              protocol:
              - tcp
              source: 443
    - host:
        node_filter:
          capabilities:
          - host:
              properties:
                cpu_arch:
                  equal: x86_64
                realtime:
                  equal: true
                cpu_usage:
                  greater_or_equal: 0.3
                mem_size:
                  greater_or_equal: 500 MB
          properties: null
    type: tosca.nodes.Container.Application
  nginx2:
    artifacts:
      application_image:
        file: nginx
        repository: registry.gitlab.aeros-project.eu
        type: tosca.artifacts.Deployment.Image.Container.Docker
    interfaces:
      Standard:
        create:
          implementation: application_image
          inputs:
            cliArgs:
            - -forwarding: ''
            - -brokerId: CFDomain
            envVars:
            - BROKER_ID: CFDomain
            - OTHER_ENV_VAR: another value
    requirements:
    - network:
        properties:
          ports:
            port1:
              protocol:
              - tcp
              source: 1026
            port2:
              protocol:
              - tcp
              source: 443
    - host:
        node_filter:
          capabilities: null
          properties:
            id: urn:ngsi-ld:InfrastructureElement:ncsrd-w1
    type: tosca.nodes.Container.Application
tosca_definitions_version: tosca_simple_yaml_1_3
```

```yaml
tosca_definitions_version: tosca_simple_yaml_1_3

description:  TOSCA simple container application with with two service components ports and hosts requirements

node_templates:
  nginx_component:
    type: tosca.nodes.Container.Application
    requirements:
      - host:
          node_filter:
            properties:
              id: "urn:ngsi-ld:InfrastructureElement:ncsrd-w2"
      - network:
          properties:
            ports:
              nginx_port1:
                properties:
                  protocol: [tcp]
                  source: 80
              nginx_port2:
                properties:
                  protocol: [tcp]
                  source: 443
    artifacts:
      nginx_component_image:
        file: nginx:latest
        type: tosca.artifacts.Deployment.Image.Container.Docker
        repository: docker_hub

    interfaces:
      Standard:
        create:
          implementation: nginx_component_image
          inputs:
            envVars:
              - SOME_NGINX_ENV_VARIABLE: SOME_NGINX_ENV_VARIABLE_VALUE
              - SOME_MORE_NGINX_ENV_VARIABLE: SOME_MORE_NGINX_ENV_VARIABLE_VALUE

  nmosquitto_component:
    type: tosca.nodes.Container.Application
    requirements:
      - host:
          node_filter:
            capabilities:
              - host:
                  properties:
                    cpu_usage: { less_or_equal: 0.5 }
                    cpu_arch: { equal: x86_64 }  #data model supports : x86_64, arm64, arm32
                    mem_size: { greater_or_equal: 4096 MB } #MB used in the data model
                    realtime: { equal: false }
                    area:
                      coordinates:
                        - [-0.3411743020092746, 39.48103451588121]
                        - [-0.3411743020092746, 39.4103451588121]
                        - [-0.3411743020092746, 39.48103451588121]
                        - [-0.3411743020092746, 39.48103451588121]

      - network:
          properties:
            ports:
              mqtt_port1:
                properties:
                  protocol: [tcp]
                  source: 1883
              mqtt_port2:
                properties:
                  protocol: [tcp]
                  source: 9883
    artifacts:
      mqtt_component_image:
        file: eclipse-mosquitto:latest
        type: tosca.artifacts.Deployment.Image.Container.Docker
        repository: docker_hub

    interfaces:
      Standard:
        create:
          implementation: mqtt_component_image
          inputs:
            envVars:
              - SOME_MQTT_ENV_VARIABLE: SOME_MQTT_ENV_VARIABLE_VALUE
              - SOME_MORE_MQTT_ENV_VARIABLE: SOME_MORE_MQTT_ENV_VARIABLE_VALUE
```

3.2. aerOS Data Fabric

3.2.1. KPI 1.2.1 Full support for data pipelines in all use cases (incl. open calls), identified through requirements elicitation (KVI-5.1)

Table 9 KPI 1.2.1 Full support for data pipelines in all use cases (incl. open calls), identified through requirements elicitation (KVI-5.1)

KPI ID number .	**KPI 1.2.1**		
KPI Name	**Full support for data pipelines in all use cases (incl. open calls), identified through requirements elicitation**		
Description	aerOS Data Fabric exposes configurable tools that are used to build data transit and transformation workflows (data pipelines). This KPI will measure the coverage of pipelines, that are prepared within the project, and required by the use-cases		
Motivation	Verification of the Data Fabric tools, that support the creation of data pipelines in practice.		
Target value	>50% scenarios		
Prerequisites	Data Product Manager and Data Product Pipeline components have been implemented and deployed.		
aerOS components (task)	Semantic Annotator, Semantic Translator (T4.1), Data Fabric (T4.2)		
Evaluation means	One of the aerOS basic goals is to provide the flexibility and adequacy in handling the data pipelines that may be encountered in its future applications. To verify that this is achievable, data pipelines required by the use-cases (and open-callers) are being specified, configured and created using aerOS Data Fabric and associated tools. By taking advantage of its modular architecture and the flexibility of the data handling mechanisms it offers, as well as the comprehensive support already available for the most commonly used data formats, Data Fabric should meet these requirements. The evaluation process, while somewhat "binary" in nature, will be conducted at all stages of pilots (and open calls) development. It will apply to all scenarios where there is a need to create and handle data pipelines.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***0***	*N/A*
Outcome elaboration (M24)	Currently, aerOS Data Fabric related tools, developed within T4.1 and T4.2 allow for building semantic data pipelines by ingesting raw data in both batch and streaming fashion, annotating it (with Morph-KGC or Semantic Annotator), homogenizing semantically (with Semantic Translator), and exposing the result via Orion-LD Context broker. It has been tested for the MVP during the mid-term review. However, the use		

	case scenarios are still in the integration phase, so no evidence of its use can be provided yet.

3.2.2. KPI 1.2.2 Semantic and syntactic interoperability between all data producers and consumers in all use cases (KVI-5.2)

Table 10 KPI 1.2.2 Semantic and syntactic interoperability between all data producers and consumers in all use cases (KVI-5.2)

KPI ID number .	**KPI 1.2.2**		
KPI Name	**Semantic and syntactic interoperability between all data producers and consumers in all use cases**		
Description	To achieve interoperability when exchanging data all participants, data producers and data consumers must understand the data. This KPI will ensure, that the data is useful for producers and consumers, through either usage of a common syntax and semantics from the get-go, or by applying data transformations.		
Motivation	Semantic and syntactic interoperability of data for aerOS Data Fabric.		
Target value	>50% scenarios		
Prerequisites	Data Product Manager and Data Product Pipeline components have been implemented and deployed.		
aerOS components (task)	Semantic Annotator, Semantic Translator (T4.1), Data Fabric (T4.2)		
Evaluation means	As long as the Data Fabric and some of the semantic tools are deployed and there is a data producer and data consumer connected, it will be considered that the scenario/use case applies for the semantic and/or syntactic interoperability.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	In all use cases, whenever data interoperability is needed it is provided on both syntactic and semantic levels. The former is supported by the Semantic Annotator and/or Morph-KGC, the latter – thanks to the semantic homogenization mechanisms provided by Semantic Translator. All these solutions together ensure full interoperability in data production/consumption at the level of Data Fabric. Like previous KPI, since the use case scenarios are still in the integration phase, no evidence of semantic and syntactic interoperability use can be provided yet.		

3.2.3. KPI 1.2.3 Reference implementation for a data infrastructure supporting full user-control in the definition of data sources, consumers and flows (KVI-5.3)

Table 11 KPI 1.2.3 Reference implementation for a data infrastructure supporting full user-control in the definition of data sources, consumers and flows (KVI-5.3)

KPI ID number .	**KPI 1.2.3**		
KPI Name	**Reference implementation for a data infrastructure supporting full user-control in the definition of data sources, consumers and flows of deployment compared to app blueprints**		
Description	Number of aerOS use cases that have followed the standard Linked Open Terms (LOT) methodology for ontology development.		
Motivation	Ontologies enable integrating data in the knowledge graph that implements the Data Fabric (i.e., the data infrastructure).		
Target value	>=3		
Prerequisites	Ontology has been developed following the guidelines defined by the LOT methodology.		
aerOS components (task)	Semantic Annotator (T4.1), Data Fabric (T4.2)		
Evaluation means	The resulting ontology must be available in a GitLab repository, where the following ontology artifacts as recommended by the LOT methodology: • Ontology requirements (in CSV format) • Ontology diagram (based on Chowlk notation, created using draw.io tool) • Ontology code (OWL code programmed and validated with Protégé tool)		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***2 (66%)***	*N/A*
Outcome elaboration (M24)	Two use cases have been identified, namely, aerOS Continuum and aerOS Building. The respective ontologies for these use case have been development according to LOT methodology guidelines and recommended tools. The resulting ontologies are available on separate repositories in aerOS GitLab, including the ontology artifacts mentioned above. The following snapshot depicts the structure of the GitLab repository for the aerOS Building use case.		

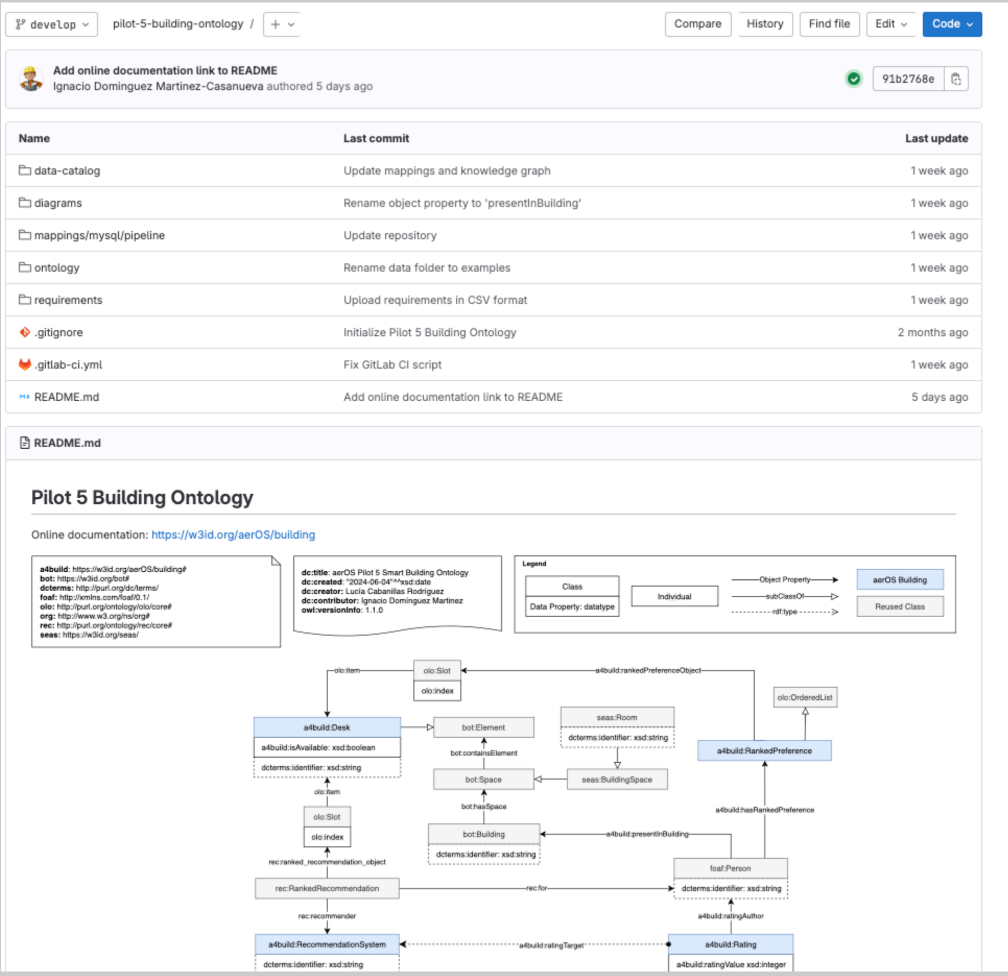

The GitLab repositories containing the ontology artifacts are listed below:

- **Continuum Ontology:** https://gitlab.aeros-project.eu/wp4/t4.1/pilot-5-building-ontology
- **Building Ontology:** https://gitlab.aeros-project.eu/wp4/t4.1/aeros-continuum

3.2.4. KPI 1.2.4 # of data sovereignty initiatives

Table 12 KPI 1.2.4 # of data sovereignty initiatives

KPI ID number .	**KPI 1.2.4**
KPI Name	**# of data sovereignty initiatives**
Description	Data sovereignty initiatives refer to efforts, policies, and frameworks with the main goal of ensuring data is subject to the laws and governance structure of where it is collected or processed. This KPI quantifies the number of data sovereignty initiatives that influence aerOS components.
Motivation	Data sovereignty is crucial to facilitate data sharing and trusted data transaction ensuring effective data usage control in distributed environments. The number of data sovereignty activities enhances the fidelity of AI models and effectiveness of autonomous control loops.

Target value	5		
Prerequisites	*N/A*		
aerOS components (task)	*N/A*		
Evaluation means	By tracking and reporting every data sovereignty in which aerOS partners have actively contributed. A valid initiative is considered each action that directly helps handling complies with regulations or legal issues, increasing control over the data and protecting it from unauthorized access		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	0	**9 (180%)**	N/A
Outcome elaboration (M24)	Software/components/applications that enable sovereignty: • Blueprint: 1 - DSSC v1.0: • Data Sovereignty Policy Management: 1 – ODLR • Data Space Communication Protocol: 1 - IDSCP 2.0 • Data Space Connector Supported: 2 – DSC, EDC • Data Catalogue: 1 – DCAT-AP • ID Management system: 1 – DAPS • Certificate Authority: 1 – X.509 • Digital Twin & Data Models: 1 – OPC-UA Work in progress • Data Space Interoperability: DSP (ISO/IEC JTC 1/SC 38) • Data Space Compatibility: 1 Eclipse TCK Further information on the enablers (software/components/applications) of sovereign data space related/due to aerOS can be found • https://gaia-x.eu/ • https://bdva.eu/ • https://www.fiware.org/ • https://dssc.eu/ • https://internationaldataspaces.org/		

3.2.5. KPI 1.2.5 aerOS data models in open markets

Table 13 KPI 1.2.5 aerOS data models in open markets

KPI ID number .	**KPI 1.2.5**
KPI Name	**aerOS data models in open markets**
Description	Number of data models used in aerOS publicly available to the open-source community.

Motivation	Promotion of open data models targeting the IoT-Edge-Cloud continuum for improving interoperability.		
Target value	5		
Prerequisites	Ontology artifacts, namely, requirements list, diagram, code, and documentation, available in the respective GitLab repository.		
aerOS components (task)	Semantic Annotator (T4.1), Data Fabric (T4.2)		
Evaluation means	The ontology must be publicly available, exposing an online documentation based on the WIDOCO tool[2]. Additionally, following LOT methodology best practices, the namespace URI of the ontology must be registered under the open w3id.org domain[3] and an entry of the ontology must be created in the Linked Open Vocabularies (LOV)[4] service for improved discoverability by the open-source community.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***2 (40%)***	*N/A*
Outcome elaboration (M24)	Currently, the aerOS Continuum Ontology and the aerOS Building Ontology have been published publicly online, using WIDOCO tool, and registered under the w3id.org domain, as shown in the following figure. 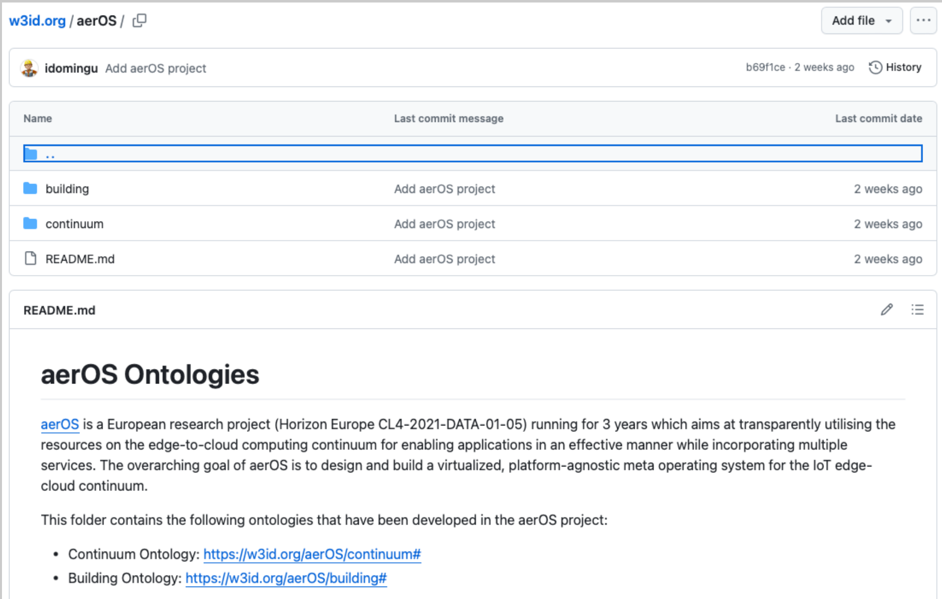 Online documentation for each ontology can be found in the following links: • **Continuum Ontology:** https://w3id.org/aerOS/continuum# • **Building Ontology:** https://w3id.org/aerOS/building#		

[2] https://github.com/dgarijo/Widoco
[3] https://w3id.org
[4] https://lov.linkeddata.es/dataset/lov

3.2.6. KPI 1.2.6 Semantic annotation support for commonly used data format

Table 14 KPI 1.2.6 Semantic annotation support for commonly used data format

KPI ID number .	**KPI 1.2.6**		
KPI Name	**Semantic annotation support for commonly used data formats**		
Description	Semantic annotation component transforms "raw data" into NGSI-LD based on specific annotation rules. This KPI will measure the number of data formats that aerOS can transform into NGSI-LD.		
Motivation	Data-level semantic interoperability and support for the unified data handling within aerOS Data Fabric.		
Target value	>=3		
Prerequisites	Data Product Manager and Data Product Pipeline components have been implemented and deployed.		
aerOS components (task)	Semantic Annotator (T4.1), Data Fabric (T4.2)		
Evaluation means	The evaluation of the KPI will be based on the number of "raw" data formats supported by the Semantic Annotator.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***3 (100%)***	*N/A*
Outcome elaboration (M24)	Current implementation of the Semantic Annotator supports 3 widely used "raw data" formats: XML, CSV, and JSON. The REST configuration interface of the Semantic Annotator, based on Swagger, is shown in the following figure. 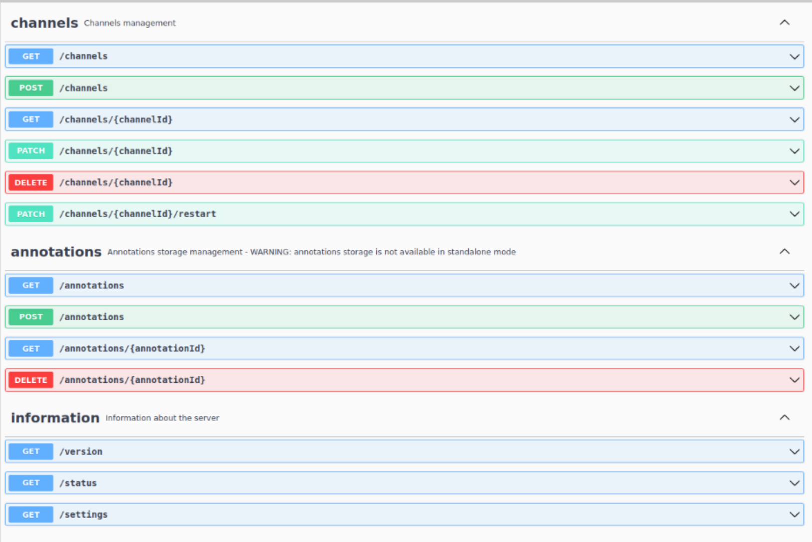		

The first example demonstrates annotation of JSON-based “personal data”. In the process, the Semantic Annotator utilizes annotation rules (depicted on the left and expressed in CARML format) telling how to transform the “raw data” into its semantic counterpart.

In the second example, the tool is used to semantically annotate a series of CSV-encoded “measurement data” coming from a temperature sensor.

The last example presents annotation of XML-based data coming from the same temperature sensor.

	

3.2.7. KPI 1.2.7 % data sources from aerOS scenarios to be semantically annotated and exposed via Data Fabric

Table 15 KPI 1.2.7 % data sources from aerOS scenarios to be semantically annotated and exposed via Data Fabric

KPI ID number .	**KPI 1.2.7**		
KPI Name	**% data sources from aerOS scenarios to be semantically annotated and exposed via Data Fabric**		
Description	Data ingested/manipulated by aerOS and expressed in NGSI-LD needs to use formal semantic models. This will enable semantic harmonization for heterogeneous data sources.		
Motivation	To achieve data shareability in aerOS.		
Target value	>50% scenarios		
Prerequisites	Data Product Manager and Data Product Pipeline components have been deployed.		
aerOS components (task)	Semantic Annotator, Semantic Translator (T4.1), Data Fabric (T4.2)		
Evaluation means	The evaluation takes into account the use cases which create/offer data and verifies if the data is properly served as NGSI-LD through aerOS Data Fabric.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*>20% scenarios*	***>45% scenarios (90%)***	*N/A*

Outcome elaboration (M24)	Data shareability in aerOS is obtained via semantically annotated data pipelines offered through Data Fabric infrastructure. All the essential elements for defining data pipelines are now available and their usage is taken into account by the aerOS use case scenarios. As long as the initial deployment tests are carried out, the percentage of aerOS scenarios will be quantified. • Pilot 1 will consider four kinds of data sources: ROS2 (Robot Operating System) DDS (Data Distribution Service), OPC UA (OPC Unified Architecture), PROFINET and MQTT broker(s). Both ROS2 DDS and MQTT will produce JSON to be semantically annotated. • Pilot 2 will provide data based on Kepler metrics, describing power consumption, resource utilization, and predicted availability of green energy (most, if not all will be semantically annotated). • Pilot 3 will handle four kinds of data sources: monitoring camera images (possibly semantically annotated through metadata), vehicle location data (including time stamp and geographic positioning; semantically annotated), field condition data (semantically annotated), and vehicle operation data (power level, engine status, vehicle configuration, etc.; semantically annotated). • Pilot 4 will consider port operational TOS data (batch-exported, possibly semantically annotated), information about the assets to be regularly maintained in the port (including equipment, materials, maintenance schedule and assigned employees; possibly semantically annotated), as well as live video feed data (most probably in "raw" format only). • Pilot 5 will involve two kinds of data sources, with data that originates either from MySQL database or MQTT broker. It will offer environmental parameters taken from sensors located in a smart building (such as temperature, humidity, pressure, or air quality). Additionally, based on the data coming from sensors, various monitoring, forecasting, recommendation/optimization data will be made available. All the data sources will be semantically annotated and exposed via Data Fabric.

3.2.8. KPI 1.2.8 Support for multiple types of data sources

Table 16 KPI 1.2.8 Support for multiple types of data sources

KPI ID number .	**KPI 1.2.8**
KPI Name	**Support for multiple types of data sources**
Description	The Data Fabric can support the ingestion of data from data sources based on different protocols and data formats such Files, RDBMS, Kafka or MQTT
Motivation	Demonstrates how the Data Fabric can cope with the heterogeneity of the continuum
Target value	>=3
Prerequisites	The Data Product Manager and Data Product Pipeline components have been implemented and deployed in a scenario with multiple heterogenous data sources.
aerOS components (task)	Semantic Annotator (T4.1), Data Fabric (T4.2)

<table>
<tr><td>Evaluation means</td><td colspan="3">Data Fabric’s support for a specific type of data source is validated when the data product owner can onboard a data product for such data source type and the Data Fabric builds a pipeline that retrieves raw data and eventually stores it in the NGSI-LD Context Broker.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>0</td><td>2 (66%)</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Currently, there is a total of two types of data sources supported in the aerOS Data Fabric: remote files and MySQL. The integration and extension of the Morph-KGC component has provided support for these two data sources. Additionally, the Data Product Manager REST API has been extended to enable onboarding data products based on these types of data sources.

The following snapshot depicts the documentation of the REST API to onboard data products from MySQL data sources:

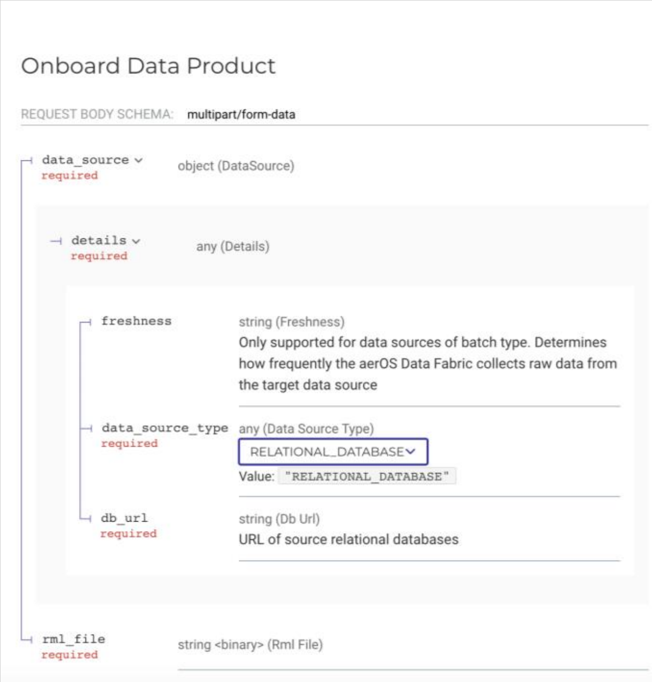

The following snapshot depicts the documentation of the REST API to onboard data products from remote file data sources:</td></tr>
</table>

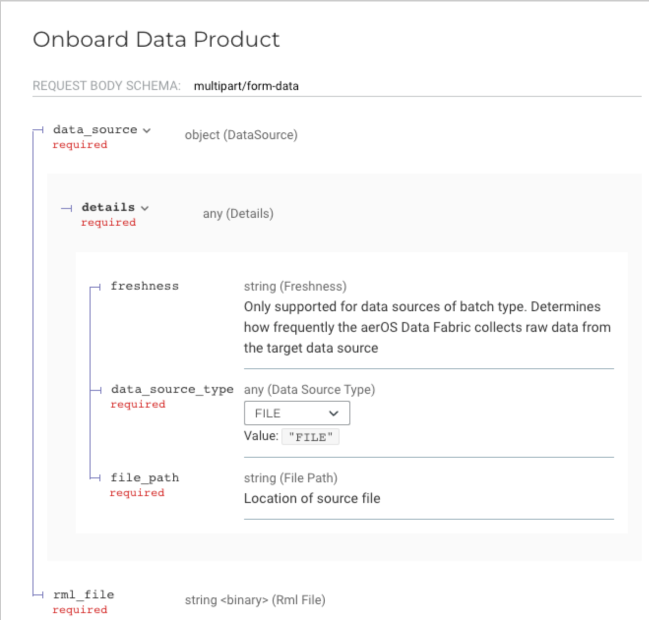

3.2.9. KPI 1.2.9 Data pipeline latency for data integration

Table 17 KPI 1.2.9 Data pipeline latency for data integration

KPI ID number .	**KPI 1.2.9**
KPI Name	**Data pipeline latency for data integration**
Description	The latency added by the Data Fabric when integrating from raw data into the knowledge graph
Motivation	High latency would limit the adoption in near real-time use cases
Target value	<1 s
Prerequisites	Data Product Pipeline components implemented and deployed in a scenario with a data source like MySQL.
aerOS components (task)	Semantic Annotator (T4.1), Data Fabric (T4.2)
Evaluation means	Using a custom developed tool, the end-to-end latency of a data product pipeline executed in the Data Fabric is measured. The total latency comprises the latencies introduced by the following steps: • Data mapping (*t1*) • RDF to NGSI-LD translation (*t2*) Therefore, the latency formula looks as follows: $t=t1+t2$ The latencies present during the ingestion of raw data in the mapping engine (e.g., Morph-KGC) as well as the materialization of the resulting NGSI-LD data in the NGSI-LD Context Broker.

Measurement period	Baseline	M24 (Deliverable D5.5)	M36 (Deliverable D5.6)
Measured value (% achieved)	N/A	**N/A**	N/A
Outcome elaboration (M24)	The tool responsible for measuring the end-to-end latency is still under development as of the writing of this deliverable. It is expected that the latencies will be measured by the next iteration of the document (D5.6).		

3.2.10. KPI 1.2.10 Simultaneous data pipeline execution

Table 18 KPI 1.2.10 Simultaneous data pipeline execution

KPI ID number .	KPI 1.2.10		
KPI Name	**Simultaneous data pipeline execution**		
Description	Maximum number of concurrent data pipelines running in the Data Fabric with guaranteed performance.		
Motivation	The Data Fabric is expected to simultaneously handle multiple data flows.		
Target value	5		
Prerequisites	Data Product Manager and Data Product Pipeline components implemented and deployed in a scenario with multiple batch data sources like MySQL or streaming data sources like Kafka.		
aerOS components (task)	Semantic Annotator (T4.1), Data Fabric (T4.2)		
Evaluation means	The Data Product Manager must handle the lifecycle of data product pipeline that ingest and integrate data from multiple data sources. To do so, the management of data products by data owners is enabled via a REST API implemented in the Data Product Manager. A preliminary version of this REST API is documented in the official aerOS documentation: https://docs.aeros-project.eu/en/latest/data/fabric/data_product_manager.html#interacting-with-the-data-product-manager		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	The data product lifecycle management feature in the Data Product Manager is still under development as of the writing of this deliverable.		

3.3. aerOS service fabric

3.3.1. KPI 1.3.1 Number of VNF/NetApps to improve performance and self-* network reconfiguration (KVI-2.3)

Table 19 KPI 1.3.1 Number of VNF/NetApps to improve performance and self- network reconfiguration (KVI-2.3)*

KPI ID number .	**KPI 1.3.1**		
KPI Name	**Number of VNF/NetApps to improve performance and self-* network reconfiguration**		
Description	The use of VNFs (or Network Apps as defined in 3GPP) aims at improving the performance of the self-configuration of the network. This is achieved by interfacing each Network App on the one hand with the native APIs of the 5G network, and, on the other hand with application itself through business APIs, enhancing either the performance of the service provision or the network configuration itself. This KPI follows the paradigm of the 3GPP SA6 standardisation activities, where specific Network Apps are realized as Vertical Application Enablers (VAEs), improving the performance for supporting services of vertical industries, or the network configuration.		
Motivation	The higher the number that a specific application or service is interfacing with NetApps, the more innovative the specific application/service is becoming, because it integrates features that are not currently available or possible with a simple OTT provision approach.		
Target value	> 6 Services/Applications that are interfacing with at least one NetApp		
Prerequisites	aerOS domain set up complete with at least one IE connected over 5G network.		
aerOS components (task)	Ingress (T3.1), TLS (T3.1), HLO (T3.3), LLO (T3.3), API Gateway (T3.4), Data Fabric (T4.2)		
Evaluation means	Logs showing 5G metrics (QoS and GPS location) exposed in aerOS Data Fabric will be monitored and presented.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***3 (50%)***	*N/A*
Outcome elaboration (M24)	In the MVP, three services are interfacing with one NetApp on top of the 5G-vehicle, which represents one IE in the NCSRD domain. In particular, as it can be seen below, the NetApp provides QoS and GPS location data to: • Command and Control (C2) application with the target to provide remote navigation		

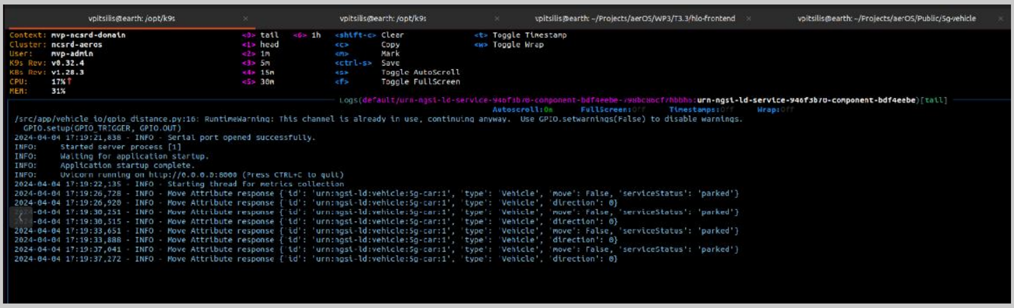

- Visualization application to provide state monitoring

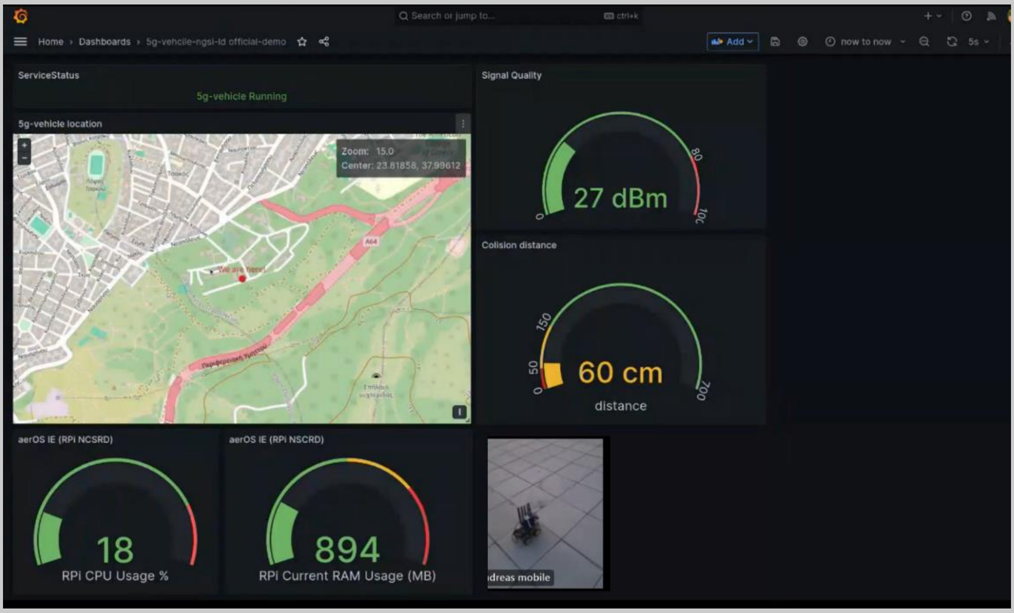

- Orion-LD to provide data to more consumers

```json
{
    "id": "urn:ngsi-ld:vehicle:5g-car:1",
    "type": "Vehicle",
    "category": {
        "type": "Property",
        "value": [
            "specialUsage",
            "Tracked"
        ]
    },
    "vehicleType": {
        "type": "Property",
        "value": "car"
    },
    "name": {
        "type": "Property",
        "value": "pi"
    },
    "vehiclePlateIdentifier": {
        "type": "Property",
        "value": "ncsrd"
    },
    "location": {
        "type": "GeoProperty",
        "value": {
            "type": "Point",
            "coordinates": [
                23.818836,
                37.996339
            ]
        }
    },
    "move": {
```


3.3.2. KPI 1.3.2 Total services delivered by aerOS

Table 20 KPI 1.3.2 Total services delivered by aerOS

KPI ID number .	**KPI 1.3.2**
KPI Name	**Total services delivered by aerOS**
Description	This KPI refers to the total number of basic and auxiliary services that aerOS consists of and delivers respectively to the users. This collection of services provides to the aerOS users the capability to interact with the Data Fabric plane of the Meta-OS, allowing them to discover IoT data, to deploy a distributed service across the continuum and to manage the aerOS nodes that can join the continuum.
Motivation	This KPIs denotes the complexity of the aerOS Meta-OS, but at the same time reflects also the complexity of the defined Meta-OS in terms of features and services offered to the aerOS users towards realizing the IoT-edge-cloud continuum in its full potential.
Target value	*> 50* aux and basic aerOS services deployed
Prerequisites	HW & Infrastructure integrated as aerOS IEs within aerOS domains
aerOS components (task)	All aerOS basic and auxiliary services from WP3 and WP4

<table>
<tr><td>Evaluation means</td><td colspan="3">Every aerOS continuum leader will provide a list of the aerOS basic and auxiliary services that are deployed along all their inside domains. The number will be listed in D5.5, and D5.6. The proof of this list will be shown as K9s screenshots in D5.4 per pilot.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>N/A</td><td>20 (40%)</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">As depicted in the following figures, for the MVP, all available aerOS services have been deployed on CF (first figure), NCSRD (second figure) and UPV (third figure) domains. These services are completely functional. In particular, the following aerOS services are already deployed:

LLO, HLO
Orion-LD CB, and Federator
Self-* (-awareness, orchestration, ...)
EAT
CI/CD (flux)
CNI, Certification Manager, Ingress, LoadBalancer
Krakend
LDAP, IdM
Management portal

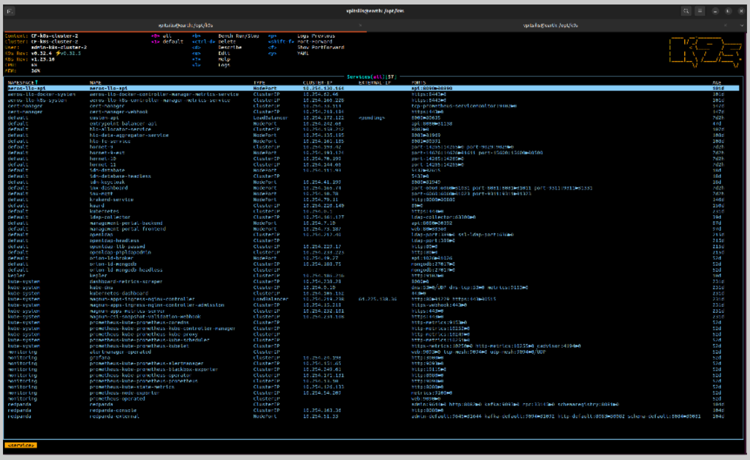
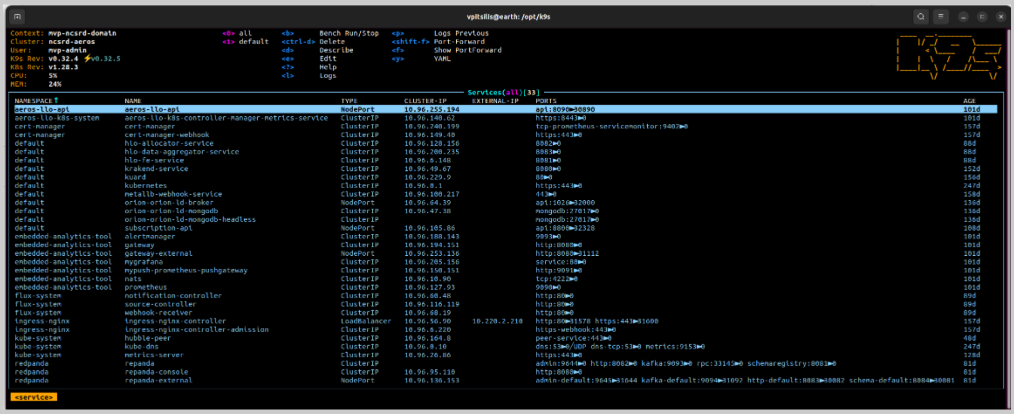</td></tr>
</table>

	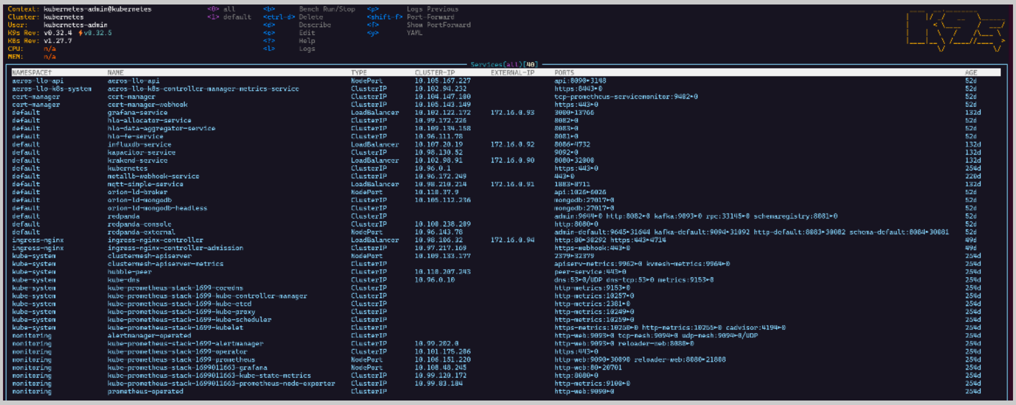

3.3.3. KPI 1.3.3 # of successful CI/CD pipelines implemented in the project

Table 21 KPI 1.3.3 # of successful CI/CD pipelines implemented in the project

KPI ID number .	**KPI 1.3.3**		
KPI Name	**# of successful CI/CD pipelines implemented in the project**		
Description	The number of repositories that have successfully completed the Continuous Integration (CI) / Continuous Deployment (CD) test designed for the project.		
Motivation	Successfully passing the CI/CD tests designed for the project leads to the conclusion that the developed code complies with the security and privacy requirements defined in the project and its correct functioning in the deployment environment.		
Target value	>4		
Prerequisites	Implementation of all the phases that were presented in the DevPrivSecOps methodology in D2.5.		
aerOS components (task)	All aerOS software components are invited to implement the DevPrivSecOps methodology presented in deliverable D2.4 with the tools and guidelines provided in D2.5.		
Evaluation means	The evaluation will assess whether all phases of the methodology have been successfully completed in the development of aerOS components. This can be seen in the GitLab pipeline of the repository, and all phases must be completed successfully (the tests of each phase must be completed successfully).		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***2 (50%)***	*N/A*
Outcome elaboration (M24)	The full pipeline methodology has now been successfully implemented in two aerOS components (two repositories):		

self-security (link to the pipeline):

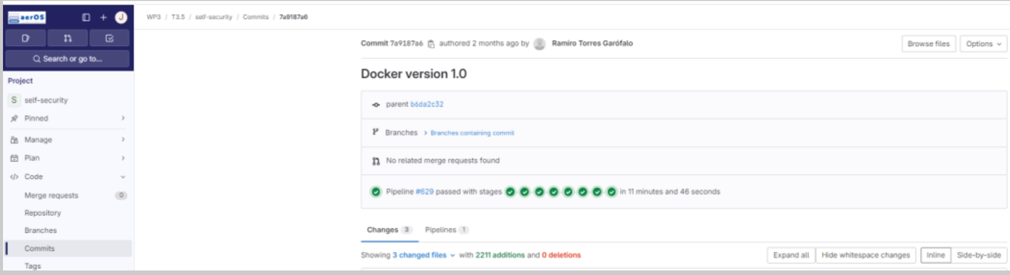

The CI/CD pipeline has been successfully completed, providing evidence that the code meets security and privacy needs.

self-orchestrator (link to the pipeline):

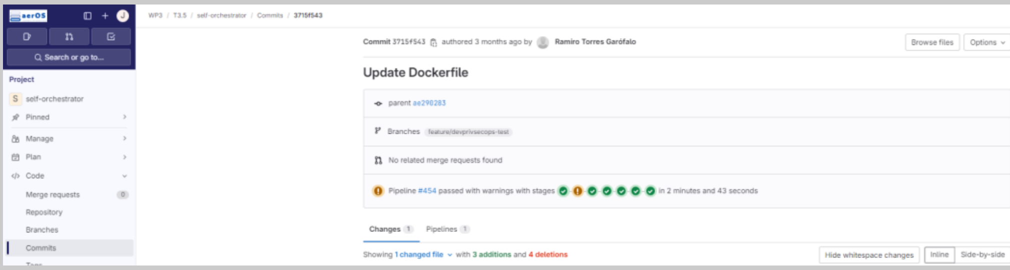

In this case it can be seen that the pipeline has a warning detected in the SAST state that needs to be analysed and remediated.

3.3.4. KPI 1.3.4 Number of different service components running in different domains that form functional services thanks to aerOS network components

Table 22 KPI 1.3.4 Number of different service components running in different domains that form functional services thanks to aerOS network components

KPI ID number .	KPI 1.3.4
KPI Name	Number of different service components running in different domains that form functional services thanks to aerOS network components
Description	This KPI refers to the number of independent software components, which if unified under the same continuum structure, then altogether form a distributed service.
Motivation	This KPI is important because it quantifies the number of components that a continuum realization connects in order to form a distributed service provision, which without the existence of the Meta-OS would not have been possible.
Target value	At least 4 components to be interfaced for the realization of a pilot.
Prerequisites	At least 2 aerOS domains setups complete.
aerOS components (task)	Ingress (T3.1), TLS (T3.1), HLO (T3.3), LLO (T3.3), API Gateway (T3.4), Self-configurator (T3.5), Data Fabric (T4.2)

Evaluation means	Screenshots, of management and reporting tools, which will explicitly show the deployment domains of service components. K9s will be used to provide evidence that service components are deployed beyond the borders of a single domain.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***2 (50%)***	*N/A*
Outcome elaboration (M24)	Although MVP has already provided evidence of independent software running as a unified service distributed over 2 domains as shown in the Demo Youtube video https://youtu.be/UV4mnN4CrwI?si=tOGQ6ATgbHWFN-Sn&t=324 (time 5:24, 6:17, 11:17), aerOS pilots are still integrating aerOS components and, thus have not yet deployed their use case services on their aerOS domains.		

3.3.5. KPI 1.3.5 Different types of networks managed by aerOS in pilot deployment

Table 23 KPI 1.3.5 Different types of networks managed by aerOS in pilot deployment

KPI ID number .	**KPI 1.3.5**		
KPI Name	**Different types of networks employed by aerOS in pilot deployment**		
Description	This KPI refers to the number of heterogeneous networks that aerOS is homogenizing within the IoT-Edge-Cloud continuum, offering to the users a unified and homogeneous experienced, independently of the underlying network technology.		
Motivation	This KPI quantifies that level of heterogeneity that that Meta-OS homogenizes and unifies under the same continuum.		
Target value	At least 2 network accesses (e.g. 5G, LAN, WiFi, ZigBee).		
Prerequisites	Computing resources integrated as IEs.		
aerOS components (task)	CNI (T3.1)		
Evaluation means	IE network layer and components reporting interface type and connectivity media used. Screenshots from within IEs explicitly stating their connectivity type will be provided.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	**2 (100%)**	*N/A*
Outcome elaboration (M24)	The IEs used for the MVP within NCSRD, CF and UPV aerOS domains were connected using either LAN, WiFi or 5G network media. The first screenshot shows the overall architecture regarding connectivity of the MVP. The following ones		

	evidence the use of different interfaces from the different IEs. 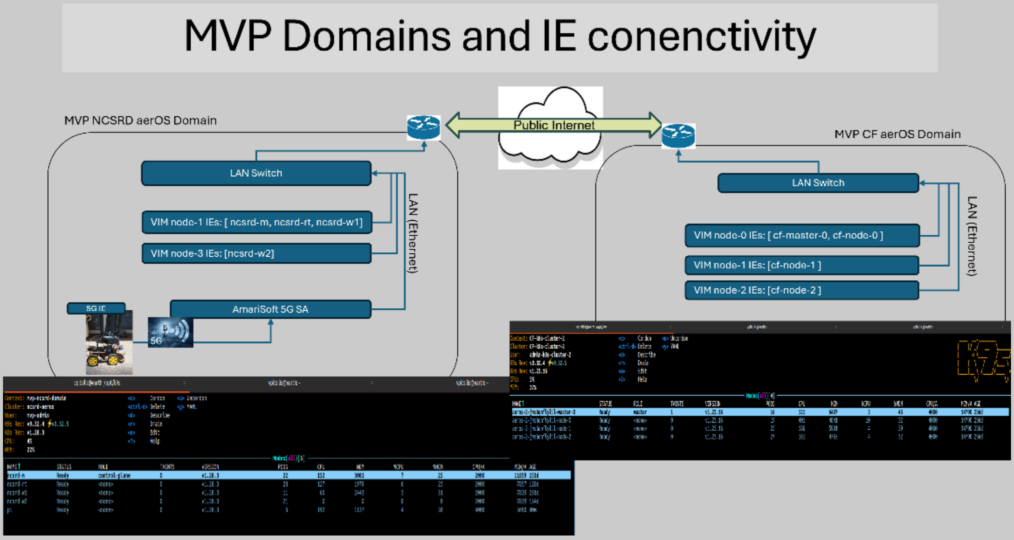  pi@pi:~$ lsusb Bus 003 Device 001: ID 1d6b:0002 Linux Foundation 2.0 root hub Bus 002 Device 001: ID 1d6b:0003 Linux Foundation 3.0 root hub Bus 001 Device 003: ID 1e0e:9001 Qualcomm / Option SDXPRAIRIE-MTP SN:0523BFAD Bus 001 Device 002: ID 2109:3431 VIA Labs, Inc. Hub Bus 001 Device 001: ID 1d6b:0002 Linux Foundation 2.0 root hub

3.4. aerOS cybersecurity and trust components

3.4.1. KPI 1.4.1 Delivery of dedicated aerOS components as Open-Source SW for cybersecurity, privacy and trust (KVI-3.1)

Table 24 KPI 1.4.1 Delivery of dedicated aerOS components as Open-Source SW for cybersecurity, privacy and trust (KVI-3.1)

KPI ID number .	KPI 1.4.1
KPI Name	Delivery of dedicated aerOS components as Open-Source SW for cybersecurity, privacy and trust

Description	The process of making components of aerOS regarding cybersecurity, privacy and trust available to the public as open-source software.		
Motivation	By delivering dedicated aerOS components as open-source software focused on cybersecurity, privacy, and trust, the initiative likely aims to contribute to the broader tech community by providing robust tools for building more secure and trustworthy digital environments		
Target value	*100% OSS services*		
Prerequisites	The integration of aerOS security, privacy and trust components in aerOS domain		
aerOS components (task)	aerOS IDM (T3.4), aerOS Secure API Gateway (T3.4), Self-security (T3.5), Trust monitoring component (T4.5)		
Evaluation means	All the cybersecurity services will make use of OSS licensing schemes (e.g., Eclipse, GPL, etc.). Moreover, to further boost the use of these tools by the community, the readme files and the aerOS official documentation of these services will include a brief guide about how to contribute to following releases.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***3 (75%)***	*N/A*

Outcome elaboration (M24)	• aerOS IDM (T3.4): https://gitlab.aeros-project.eu/wp3/t3.4/idm
	  ○ License: Apache license, version 2.0
	  ○ License/Readme file: License proof file pending
	• aerOS Secure API Gateway (T3.4): https://gitlab.aeros-project.eu/wp3/t3.4/api-gateway
	  ○ License: Unlicense license
	  ○ License/Readme file: https://gitlab.aeros-project.eu/wp3/t3.4/api-gateway/-/blob/main/README.md?ref_type=heads
	• aerOS Self-security (T3.5) https://gitlab.aeros-project.eu/wp3/t3.5/self-security/-/tree/main
	  ○ License: Unlicense license
	  ○ License/Readme file: https://gitlab.aeros-project.eu/wp3/t3.5/self-security/-/blob/main/LICENCE.TXT
	• aerOS Trust Monitoring component (T4.5) https://gitlab.aeros-project.eu/wp4/t4.5/trust-management
	  ○ License: Apache license, version 2.0
	  ○ License/Readme file: https://gitlab.aeros-project.eu/wp4/t4.5/trust-management/-/blob/main/LICENSE?ref_type=heads		

3.4.2. KPI 1.4.2 # scenarios with security, privacy and trust by design deployed (KVI-3.2)

Table 25 KPI 1.4.2 # scenarios with security, privacy and trust by design deployed (KVI-3.2)

<table>
<tr><th>KPI ID number .</th><th colspan="3">KPI 1.4.2</th></tr>
<tr><td>KPI Name</td><td colspan="3"># scenarios with security, privacy and trust by design deployed (KVI-3.2)</td></tr>
<tr><td>Description</td><td colspan="3">Quantifies the number of scenarios where principles of security, privacy, and trust have been integrated by design.</td></tr>
<tr><td>Motivation</td><td colspan="3">This KPI tracks the implementation of these foundational principles from the earliest stages of development, ensuring that each deployment is inherently secure, respects user privacy, and is trustworthy</td></tr>
<tr><td>Target value</td><td colspan="3">>50% scenarios</td></tr>
<tr><td>Prerequisites</td><td colspan="3">The integration of aerOS security, privacy and trust components in aerOS domain</td></tr>
<tr><td>aerOS components (task)</td><td colspan="3">aerOS IDM (T3.4), aerOS Secure API Gateway (T3.4), Self-security (T3.5), Trust monitoring component (T4.5)</td></tr>
<tr><td>Evaluation means</td><td colspan="3">A checklist or matrix has been developed to monitor the integration of each security, privacy, and trust component across various scenarios.
<table>
<tr><th>Pilot</th><th>Scenario Until M24</th><th>IDM Integration (Y/N</th><th>Secure API Gateway (Y</th><th>Trust Monitoring (Y/N</th><th>Self-security (Y/N)</th><th>Comments/Observations</th><th>Performance Details</th></tr>
<tr><td>Pilot1</td><td>Scenario 1: Green manufacturing (zero net-energy) and CO2 footprint monitoring (SIPBB)</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot1</td><td>Scenario 2: Automotive Smart Factory Zero Defect Manufacturing (INNOVALIA)</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot1</td><td>Scenario 3: AGV swarm zero break-down logistics & zero ramp-up safe PLC reconfiguration for lot-size-1 production</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot2</td><td>Scenario 4: Green Edge Processing</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot2</td><td>Scenario 5: Secure Federation of edge/cloud</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot3</td><td>Scenario 6: Cooperative large-scale producing</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot3</td><td>Scenario 7: Basis for CO2 neutral intelligent operation</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot4</td><td>Scenario 8:Predictive maintenance of Container Handling Equipment</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot4</td><td>Scenario 9: Risk prevention via Computer Vision in the edge</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot5</td><td>Scenario 10: Intelligent Occupational Safety and Health</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Pilot5</td><td>Scenario 11: Cybersecurity and data privacy in building automation</td><td></td><td></td><td></td><td></td><td></td><td></td></tr>
</table>
</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>N/A</td><td>0 scenarios (0%)</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">aerOS security, privacy and trust components have not been fully integrated by any pilot as of M24. Pilot 2 and Pilot 5 are two examples of pilots that have released components but have not yet integrated additional tools. Every detail from the pilots' presentations has been collected and included into the corresponding table. The function and ultimate outcome will be calculated using this data.</td></tr>
</table>

Pilot	Scenario Until M24	IDM Integration (Y/N	Secure API Gateway (Y	Trust Monitoring (Y/	Self-security (Y/N)	Comments/Observations	Performance Details
Pilot1	Scenario 1: Green manufacturing (zero net-energy) and CO2 footprint monitoring (SIPBB)	No	No	NO	NO		
Pilot1	Scenario 2: Automotive Smart Factory Zero Defect Manufacturing (INNOVALIA)	No	NO	No	No		
Pilot1	Scenario 3: AGV swarm zero break-down logistics & zero ramp-up safe PLC reconfiguration for lot-size-1 production	No	No	No	No		
Pilot2	Scenario 4: Green Edge Processing	Yes	Yes	No	No	KrakenD have been deployed in cloud(entrypoint)domain not in the container domain	
Pilot2	Scenario 5: Secure Federation of edge/cloud	Yes	Yes	No	No	KrakenD have been deployed in cloud(entrypoint)domain not in the container domain	
Pilot3	Scenario 6: Cooperative large-scale producing	No	No	No	No		
Pilot3	Scenario 7: Basis for CO2 neutral intelligent operation	No	No	No	No		
Pilot4	Scenario 8:Predictive maintenance of Container Handling Equipment	No	No	No	No		
Pilot4	Scenario 9: Risk prevention via Computer Vision in the edge	No	No	No	No		
Pilot5	Scenario 10: Intelligent Occupational Safety and Health	No	YES	No	No	IDM intergration is pending	N/A
Pilot5	Scenario 11: Cybersecurity and data privacy in building automation	NO	YES	No	No0	IDM intergration is pending	N/A

3.4.3. KPI 1.4.3 Delivery of a DevPrivSecOps cookbook and good practices manual (KVI-3.3)

Table 26 KPI 1.4.3 Delivery of a DevPrivSecOps cookbook and good practices manual (KVI-3.3)

KPI ID number .	**KPI 1.4.3**		
KPI Name	**Delivery of a DevPrivSecOps cookbook and good practices manual (KVI-3.3)**		
Description	Cookbook to guide aerOS developers to implement the DevPrivSecOps methodology		
Motivation	The DevPrivSecOps methodology designed in aerOS allows to guide the project developers (and the developer community) to develop secure and privacy aware code by design.		
Target value	The project expects to produce 3 cookbooks in different formats: a DevPrivSecOps methodology definition report, a methodology implementation manual and an interactive GitLab guide with an example implementation.		
Prerequisites	The only prerequisite is to have a code repository in the project's GitLab.		
aerOS components (task)	This cookbook is used to help aerOS tool owners implement the DevPrivSecOps methodology designed and presented in D2.4 and D2.5.		
Evaluation means	This KPI will be measured by the number of cookbooks that have been distributed to the consortium for this purpose.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***2 (66%)***	*N/A*
Outcome elaboration (M24)	With the completion of the task T2.4, as the end of this task, the deliverable D2.5 has been generated the implementation guides of the DevPrivSecOps methodology, and the configuration and use guides of the tools selected to implement this methodology have been added. A cookbook document has been generated (https://aeros-project.eu/wp-content/uploads/2024/07/aerOS_DevPrivSecOps_CB.pdf) as well as a		

	ReadTheDocs page has been generated with the implementation guidelines (https://docs.aeros-project.eu/en/latest/methodology/index.html) Finally, a video with an example of the implementation of the methodology will be generated and uploaded to the project's dissemination channels before M27

3.4.4. KPI 1.4.4 % of users/device/services properly authenticated

Table 27 KPI 1.4.4 % of users/device/services properly authenticated

KPI ID number .	**KPI 1.4.4**		
KPI Name	**% of users/device/services properly authenticated**		
Description	The percentage of users/device/services properly authenticated through the aerOS identity management (IdM) service.		
Motivation	Monitoring properly authenticated users' device/services (provided they have submitted the correct credentials) allows to verify the correct functioning of the aerOS IdM.		
Target value	> 95%		
Prerequisites	The prerequisite for this KPI is to have Keycloak and OpenLDAP installed, configured and federated to share data between them.		
aerOS components (task)	Keycloak and OpenLDAP (T3.4)		
Evaluation means	To evaluate this KPI, the Keycloak database will be accessed where the record of authentication attempts and which of these attempts have been successful is stored.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***30/31 –> 96,77% (100%)***	*N/A*
Outcome elaboration (M24)	At the moment the only Keycloak installed is the one in the CloudFerro cluster, which has been presented in the MVP1. The database of this component has been accessed to see the authentications that have been performed and it has been seen that out of 40 attempts only 1 has been unsuccessful, achieving a 97.5% of successful authentications. These authentication attempts have been made by the Management Portal and the KrakenD.		

	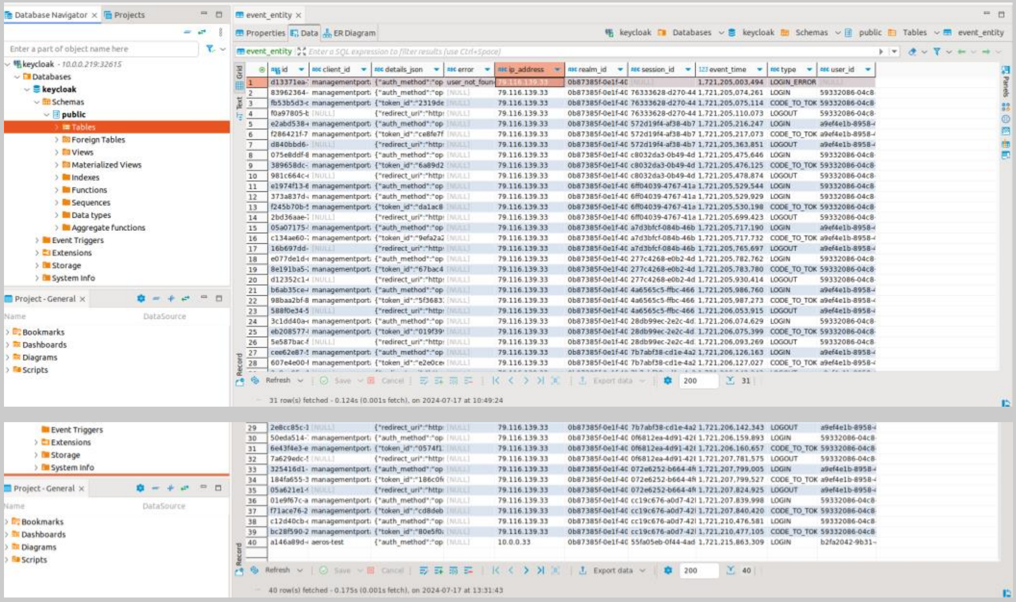 Once IdM has been installed in all project pilots, authentication attempts will be analysed to see if the KPI of >95% successful authentications is achieved in M36

3.4.5. KPI 1.4.5 # of parallel successfully authenticated user/devices/services

Table 28 KPI 1.4.5 # of parallel successfully authenticated user/devices/services

KPI ID number .	**KPI 1.4.5**
KPI Name	**# of parallel successfully authenticated user/devices/services**
Description	This KPI offers insights on the simultaneous load that the authentication system can proficiently manage
Motivation	By monitoring this KPI, potential bottlenecks are identified, and informed decisions about necessary upgrades or optimizations to accommodate growing demand are made. It also helps in stressing testing and capacity planning, ensuring that the aerOS remains responsive and secure even as the number of simultaneous authentication requests increases.
Target value	>150
Prerequisites	The prerequisite for this KPI is to have Keycloak and OpenLDAP installed, configured and federated to share data between them.
aerOS components (task)	Keycloak and OpenLDAP (T3.4)
Evaluation means	To evaluate this KPI, the Keycloak database where the record of authentication attempts and which of these attempts have been successful stored will be accessed and analysed.

Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***151/155 (94.2%)***	*N/A*
Outcome elaboration (M24)	Keycloak is currently installed only in the CloudFerro cluster. The Keycloak database was accessed through DBeaver using the appropriate credentials while conducting authentication tests to observe the results and derive an outcome. As it can be seen in the following screenshot, out of the 155 authentication attempts, 150 were successful. 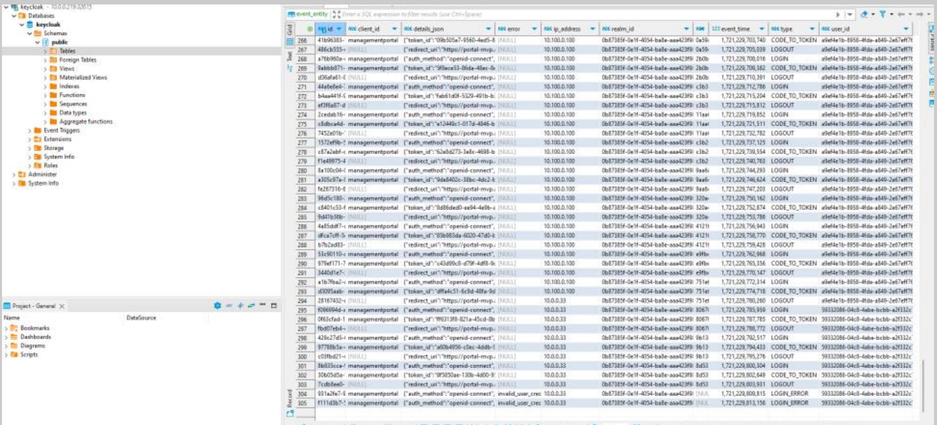 This is a very good initial sample regarding Keycloak's behaviour under a large number of authentication attempts. For M36, the results integrated into the pilots will be reviewed.		

3.4.6. KPI 1.4.6 % of users/device/services properly authorized

Table 29 KPI 1.4.6 % of users/device/services properly authorized

KPI ID number .	**KPI 1.4.6**
KPI Name	**% of users/device/services properly authorized**
Description	Percentage of users/devices/services successfully authorised through the aerOS Identity Management (IdM) service.
Motivation	Monitoring of properly authorised user devices/services (provided they have the permissions to access the target service/data) allows to verify the correct functioning of the authorisation component of the aerOS IdM.
Target value	*>95%*
Prerequisites	The prerequisite for this KPI is to have Keycloak and OpenLDAP installed, configured and federated to share data between them. Connection of Keycloak with the Management Portal for the users' authorization, and with KrakenD for the API access authorization.
aerOS components (task)	Keycloak, OpenLDAP and KrakenD (T3.4), and Management Portal (T3.6)

<table>
<tr><td>Evaluation means</td><td colspan="3">To evaluate this KPI, the Keycloak database will be accessed where the record of authentication attempts and which of these attempts have been successful is stored.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>N/A</td><td>100%</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">
At the moment, two different views have been enabled in the Management Portal, one for the domainadministrator and one for the iotservicedeployer. These views are authorised according to the role of the authenticated user.
Domainadministrator view:
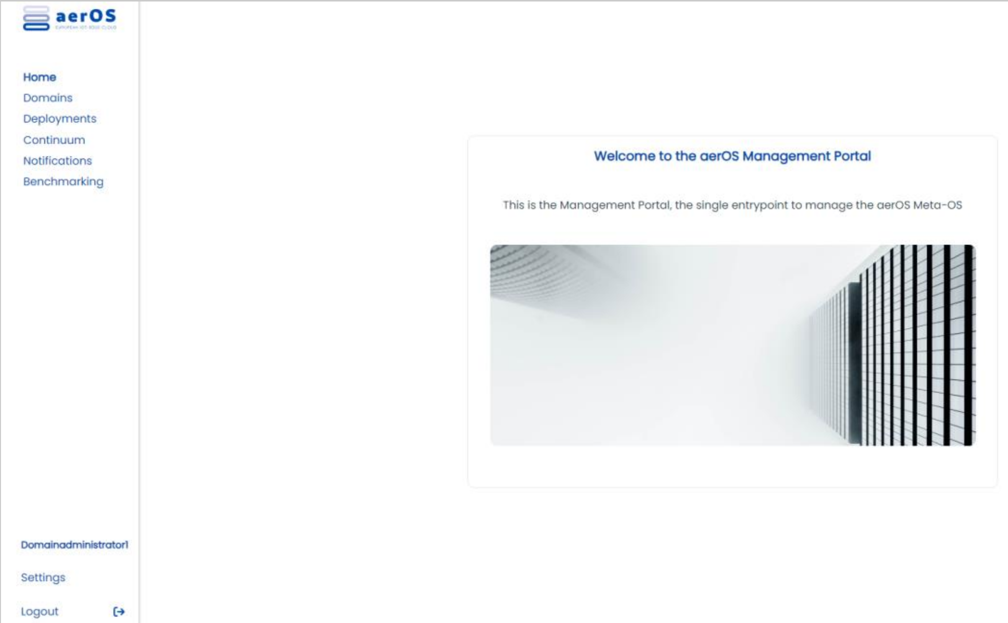

Iotservicedeployer view:
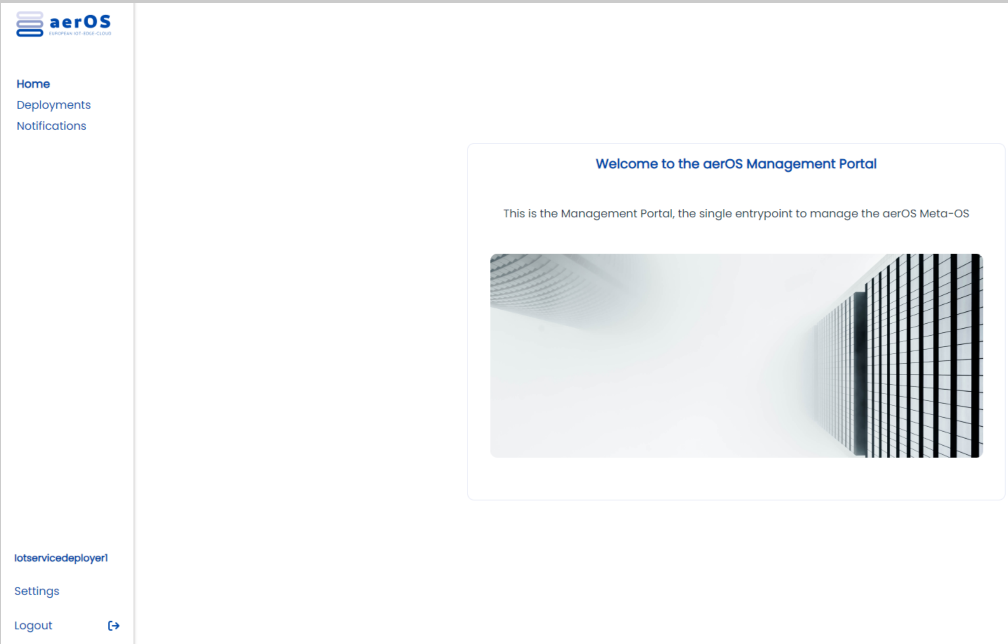

</td></tr>
</table>

	The correct functioning of the Management Portal with different roles has been validated in the MVP presented in M18. Real users have been generated for the pilots, now that the portal is being installed in the pilots, and test users are being used for the moment for each type of role.

3.4.7. KPI 1.4.7 # of petitions handled by the API Gateway per second

Table 30 KPI 1.4.7 # of petitions handled by the API Gateway per second

KPI ID number .	**KPI 1.4.7**		
KPI Name	**# of petitions handled by the API Gateway per second**		
Description	The total number of petitions that the API Gateway is capable of handling per second without getting overloaded.		
Motivation	It is important to set a minimum number of petitions that the API Gateway must be able to fulfil without overloading to handle the average operations of the aerOS platforms.		
Target value	15 petitions per second		
Prerequisites	*The aerOS* domain deployed and ready with KrakenD deployed as the sole entrypoint into the domain (Kubernetes Ingress, etc). The backends also must be deployed and ready to receive traffic from KrakenD.		
aerOS components (task)	*KrakenD API Gateway (T3.4), Keycloak IAM (T3.4), OpenLDAP (T3.4)*, Orion-LD (T4.2 and T4.6). All the necessary components to validate and authenticate a user as well as the backend to send the petitions to.		
Evaluation means	Using open-source benchmarking tools such as “autocannon,” an extremely high number of requests per second will be sent to the API Gateway to test its ability to withstand the load.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	10 petitions per second	***Average of 156Kb/s – see the different scenarios analysed)***	*N/A*
Outcome elaboration (M24)	KrakenD deployed in all aerOS domains alongside all other T3.4 tools, token caching done to reduce latency and the number of times a petition must call the backend. Improvements over the initial performance allowed for the latency to be greatly reduced, as the tests will show below. Two tests were performed on both the UPV testing domain and the CloudFerro testing environment, i) one where the user sends a petition with an invalid user token, ii) another where the token is valid. The first test showed these results on the UPV domain:		

```
Running 10s test @ http://10.108.77.245:8080/entities?type=IE
10 connections with 4 pipelining factor
```

Stat	2.5%	50%	97.5%	99%	Avg	Stdev	Max
Latency	1 ms	6 ms	21 ms	26 ms	7.29 ms	5.3 ms	76 ms

Stat	1%	2.5%	50%	97.5%	Avg	Stdev	Min
Req/Sec	4,651	4,651	5,139	5,583	5,133.2	263.15	4,651
Bytes/Sec	461 kB	461 kB	509 kB	553 kB	508 kB	26 kB	460 kB

```
Req/Bytes counts sampled once per second.
# of samples: 10

0 2xx responses, 51335 non 2xx responses
51k requests in 10.02s, 5.08 MB read
```

And these results on the CloudFerro domain:

```
Running 10s test @ https://cf-mvp-domain.aeros-project.eu/entities?type=IE
10 connections with 4 pipelining factor
```

Stat	2.5%	50%	97.5%	99%	Avg	Stdev	Max
Latency	59 ms	234 ms	482 ms	533 ms	237.6 ms	122.25 ms	688 ms

Stat	1%	2.5%	50%	97.5%	Avg	Stdev	Min
Req/Sec	400	400	427	512	463.6	46.53	400
Bytes/Sec	74.8 kB	74.8 kB	79.8 kB	95.8 kB	86.7 kB	8.73 kB	74.8 kB

```
Req/Bytes counts sampled once per second.
# of samples: 10

0 2xx responses, 4636 non 2xx responses
5k requests in 10.07s, 867 kB read
```

The second test showed these results on the UPV domain:

	Running 10s test @ http://10.108.77.245:8080/entities?type=IE 10 connections

Stat	2.5%	50%	97.5%	99%	Avg	Stdev	Max
Latency	8 ms	41 ms	113 ms	131 ms	45.77 ms	29.01 ms	162 ms

Stat	1%	2.5%	50%	97.5%	Avg	Stdev	Min
Req/Sec	203	203	216	222	215.1	5.45	203
Bytes/Sec	35.9 kB	35.9 kB	38.2 kB	39.3 kB	38.1 kB	965 B	35.9 kB

Req/Bytes counts sampled once per second.
# of samples: 10

2k requests in 10.02s, 381 kB read

And these results on the CloudFerro domain:

Running 10s test @ https://cf-mvp-domain.aeros-project.eu/entities?type=IE
10 connections

Stat	2.5%	50%	97.5%	99%	Avg	Stdev	Max
Latency	60 ms	63 ms	132 ms	246 ms	86.59 ms	37.4 ms	305 ms

Stat	1%	2.5%	50%	97.5%	Avg	Stdev	Min
Req/Sec	120	120	159	161	153.5	12.37	120
Bytes/Sec	31.8 kB	31.8 kB	42.1 kB	42.7 kB	40.7 kB	3.28 kB	31.8 kB

Req/Bytes counts sampled once per second.
# of samples: 10

2k requests in 10.08s, 407 kB read

As can be seen from both tests on both domains, KrakenD is very resilient to the load tests, being able of taking a load of over five thousand petitions per second if the token is invalid and thus KrakenD does not need to send any traffic to the backend.

The discrepancies between the number of petitions in the first test can be attributed to the distance between the UPV and CF domains. Since the tests were made from within the UPV it takes considerably less time for the petitions to reach the server.

As for the second test, the number of petitions is similar since most of the bottleneck here happens when KrakenD processes the valid petition into the backend and returns the valid response, not so much due to the distance, even though it is still relevant.

3.4.8. KPI 1.4.8 % trusted scenarios that make use of IOTA's DLT

Table 31 KPI 1.4.8 % trusted scenarios that make use of IOTA's DLT

KPI ID number .	KPI 1.4.8
KPI Name	% trusted scenarios that make use of IOTA's DLT

<table>
<tr><td>Description</td><td colspan="3">Being able to share continuum-wide relevant information using the IOTA DLT via Hornet node peer-to-peer message exchange.</td></tr>
<tr><td>Motivation</td><td colspan="3">One of the main tools that bring trust into the aerOS platform is IOTA, the number of messages shared by the different nodes is crucial for the platform to understand the global status of the continuum.</td></tr>
<tr><td>Target value</td><td colspan="3">5 data transactions per minute</td></tr>
<tr><td>Prerequisites</td><td colspan="3">The aerOS multiple interconnected domains</td></tr>
<tr><td>aerOS components (task)</td><td colspan="3">IOTA (T4.5), Trust score calculator (T4.5), Self* tools (T3.5), Use case tools (T5.2). Multiple elements of the aerOS continuum will generate events that need to be registered in the DLT and will share them in the IOTA Tangle to all the other IEs in the continuum.</td></tr>
<tr><td>Evaluation means</td><td colspan="3">The IOTA tools themselves will be used to monitor the traffic of any given deployment.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>1 data transaction per minute</td><td>3-4 transactions per minute in a single domain demonstrated (see below)</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Deployed 11 separate Hornet Nodes in a controlled testing environment among 4 different interconnected domains. All elements necessary deployed, testing integration with self* features and the trust score calculator.

For these tests the following architecture is used:

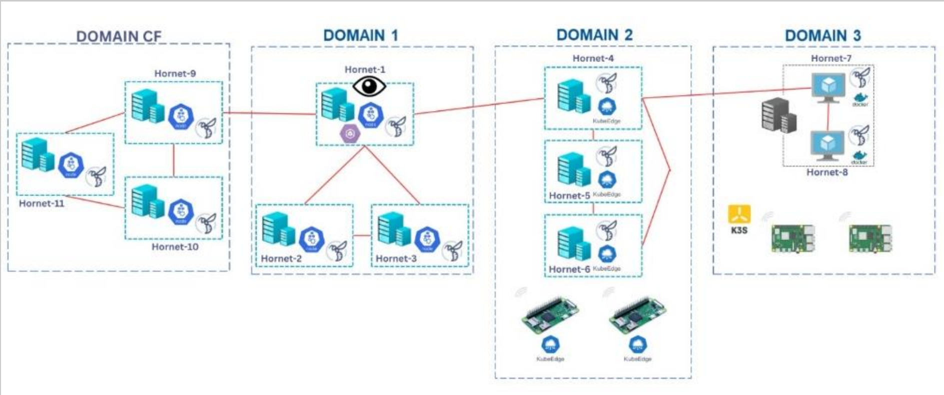

The events generated will be sent to the hornet nodes in the CloudFerro domain (hornet-9, hornet-10 and hornet-11). The inclusion of these data blocks will be verified by the Coordinator found in Domain 1 in the UPV infrastructure in the form of milestones, which are launched automatically by the coordinator every 20 seconds. The traffic received by node “hornet-9” in the span of a minute can be seen in the following images:</td></tr>
</table>

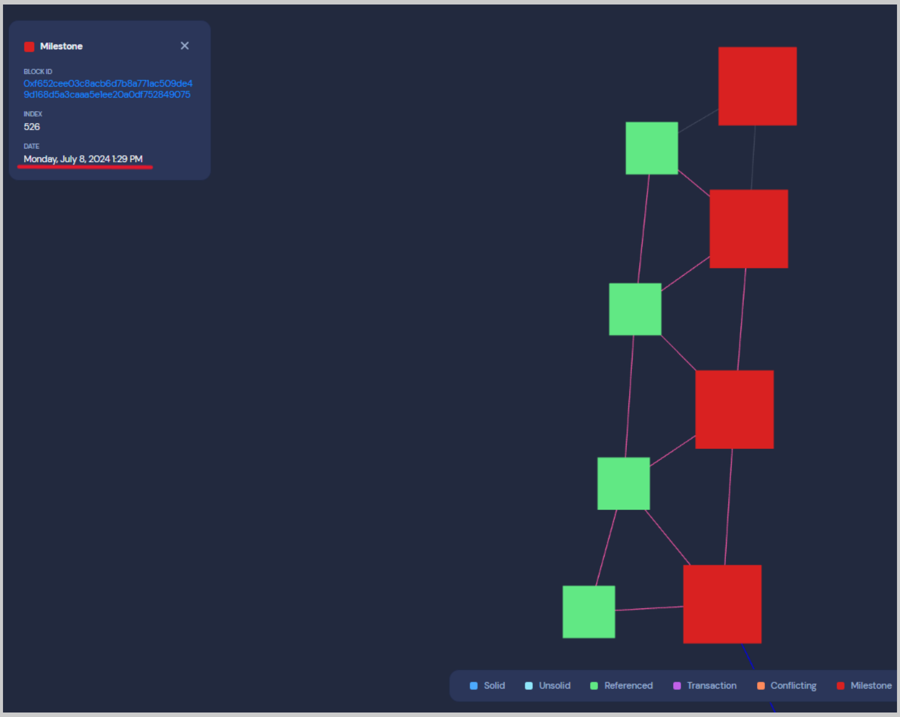

These show that in the span of a minute roughly 4 data transactions are made and that all of them are referenced by a milestone and thus included in the DLT (note the green color of the blocks, which showcases this reference).

An example of one such referenced block can be seen below:

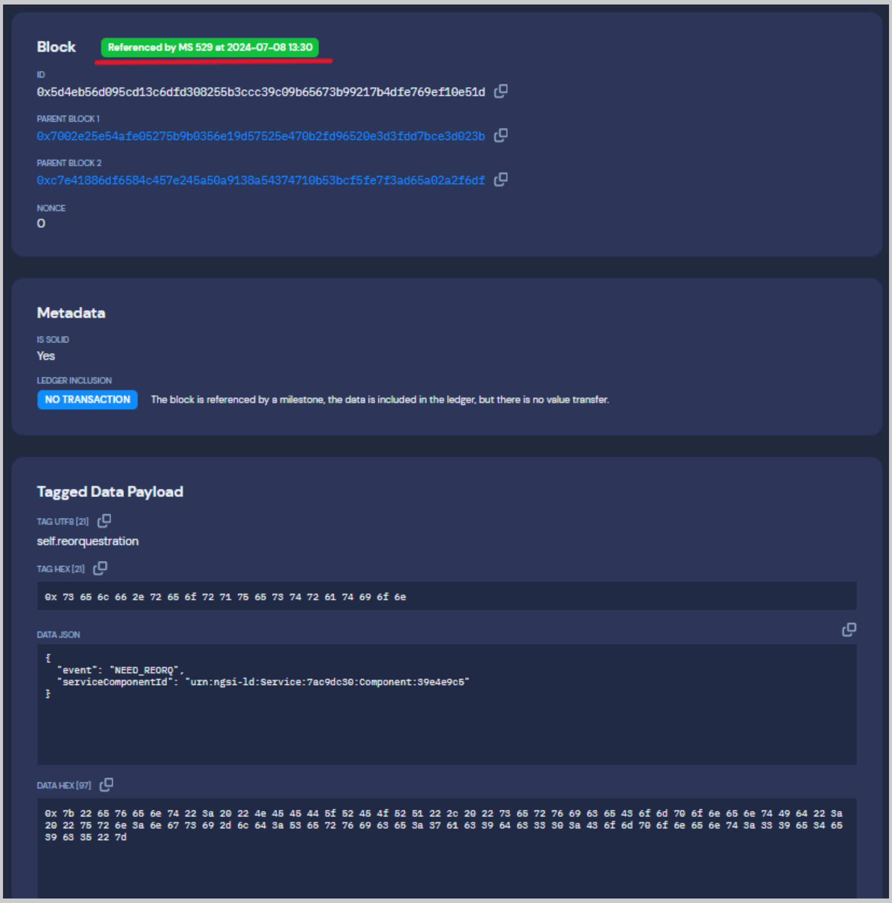

This block is referenced by milestone 529 seen in the previous image and is thus included in the DLT. Also, the contents of the block can be seen below, the tag "*self.reorquestation*" indicates that this is a reorquestation petition sent by the self* components, referencing the ID of such component.

3.4.9. KPI 1.4.9 Network overload limit due to the usage of IOTA and Tangle

Table 32 KPI 1.4.9 Network overload limit due to the usage of IOTA and Tangle

KPI ID number .	**KPI 1.4.9**
KPI Name	**Network overload limit due to the usage of IOTA and Tangle**
Description	Creating and implementing an IOTA Tangle network of nodes that share information between them without managing to overload the network.
Motivation	An IOTA Tangle network allows for peer to peer sharing of information between nodes, benefiting the entire continuum. However, it must be done in a way that does not completely overload the network.
Target value	aerOS private IOTA Tangle network deployed and running without increasing the network load by more than 30%
Prerequisites	The aerOS multiple interconnected domains

<table>
<tr><td>aerOS components (task)</td><td colspan="3">IOTA (T4.5), Trust score calculator (T4.5), Self-* components (T3.5), Use-cases tools (T5.2). Multiple elements of the aerOS continuum will generate events that need to be registered in the DLT and will share them in the IOTA Tangle to all the other IEs in the continuum.</td></tr>
<tr><td>Evaluation means</td><td colspan="3">The impact of IOTA on the network will be evaluated using the Kubernetes tools themselves, as well as the IOTA metrics plugins. A custom script will be used to generate an unusual load on the environment.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>aerOS services up and running as expected.</td><td>Minimal load increase with the expected network traffic.</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Deployed 11 separate Hornet Nodes in a controlled testing environment among 4 different interconnected domains. All elements necessary deployed, testing integration with self* features and the trust score calculator.

For these tests the following architecture is used:

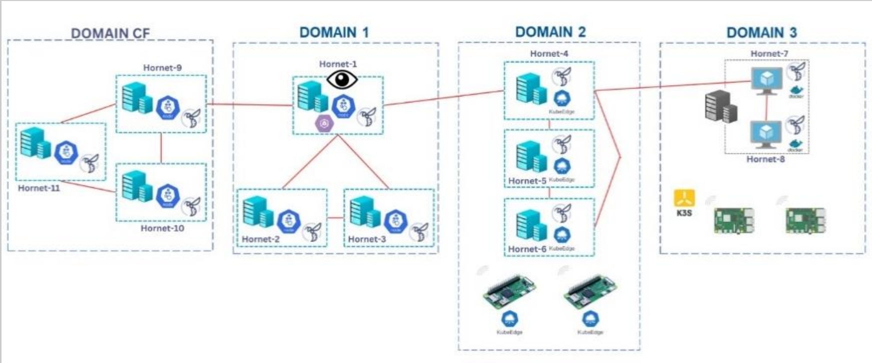

For these tests a custom script will be used to generate an unusually high amount of petitions (around 12 per minute, double of what is expected at the current time) and see if there’s a noticeable difference in the cluster metrics. The following images show 12 messages being received roughly in a minute:</td></tr>
</table>

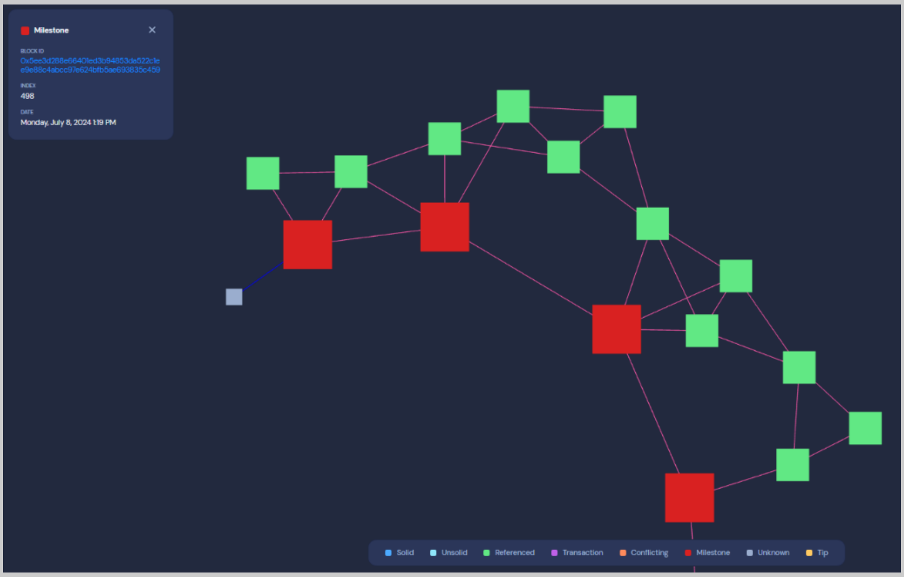

Note again that the red blocks represent milestones being launched every 20 seconds while the green ones represent the new blocks being added to the DLT.

As can be seen in the upper left corner of the image the timestamp for the test is 13:20, checking that timestamp in the cluster Grafana metrics component shows the following:

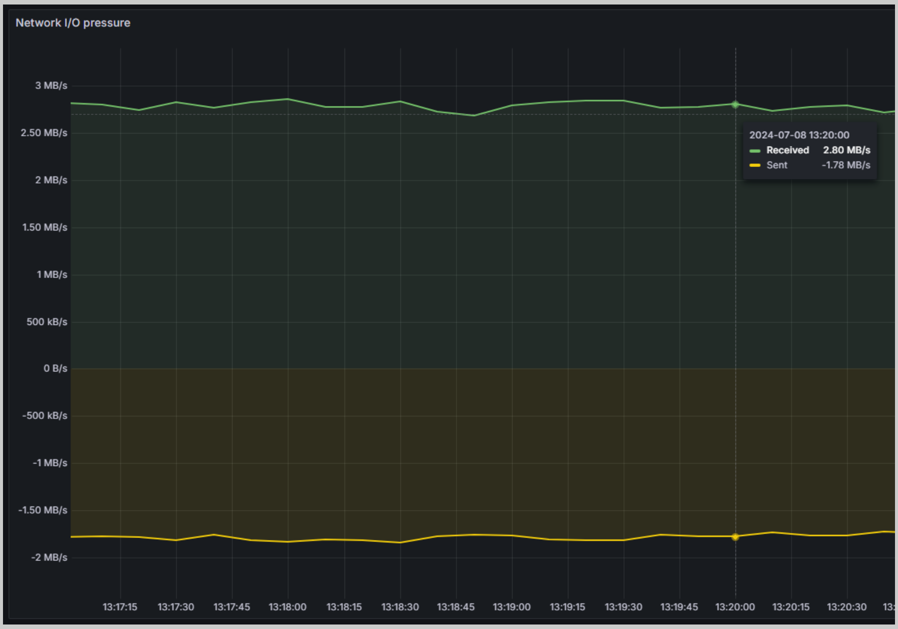

In the entire minute where the test was being made the difference in network load is practically non-existent, this can be further verified by checking the traffic in a wider time bracket:

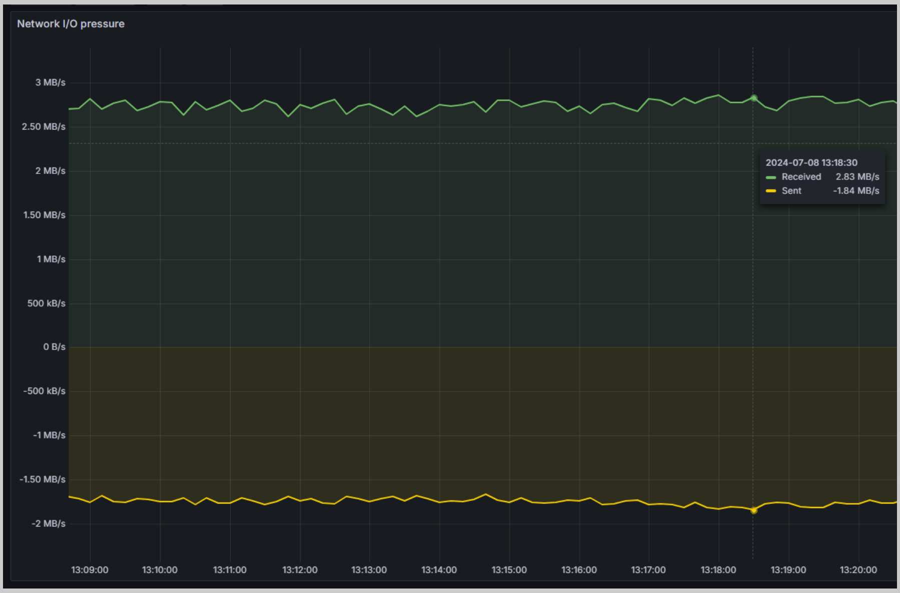

It can be safely said that there is no noticeable difference in the network load when the different IOTA Hornet nodes are sharing information.

3.4.10. KPI 1.4.10 Trust Score Recalculation and Resource Balance

Table 33 KPI 1.4.10 Trust Score Recalculation and Resource Balance

KPI ID number .	KPI 1.4.10
KPI Name	Trust Score Recalculation and Resource Balance

Description	This KPI evaluates the efficiency of the trust score recalculation process in relation to the consumption of aerOS resources, ensuring that the operational demands of maintaining updated trust scores do not lead to excessive use of resources.		
Motivation	This KPI is motivated by the imperative to harmonize the necessity for dynamic and robust trust management with the overarching need to preserve system performance and reliability		
Target value	Mean increase in resource usage due to trust score recalculation activities < 30% regular use		
Prerequisites	Trust manager and Orion-LD Context Broker running in a domain		
aerOS components (task)	Self-awareness (T3.5), Context Broker (T4.2), aerOS Trust Component (T4.5)		
Evaluation means	To evaluate this, the latency of queries to the Orion-LD Context Broker will be measured with and without the use of a trust manager in the domain. A custom script has been created to emulate multiple IEs and simulate the load generated by the self-awareness module. Additionally, another custom script has been used to calculate the latency of the queries to the Orion-LD Context Broker. The output of this script is a CSV file containing the latency and timestamps of the queries. This CSV file is later used as input to Grafana for visualizing the results.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	aerOS services up and running as expected	**<2% in 5 IEs scenario** **<1% in 20 IEs scenario**	*N/A*
Outcome elaboration (M24)	Two types of latency measurements were conducted (one without a trust manager and one with a trust manager) in two different scenarios: **1.** In the first scenario, the trust score was calculated in a cluster of 5 IEs **2.** In the second scenario, it was calculated in a cluster of 20 IEs. **Scenario 1 (5 IEs)**: In the first figure, the latency to the Orion-LD Context Broker when using the trust manager is observed, while in the second figure, the latency to the Orion-LD broker without the use of the trust manager is shown. 		

The statistics are as follows:

	With trust score	Without trust score
Mean Latency (ms)	3.93	3.86
Highest Latency (ms)	4.91	4.79
Lowest Latency (ms)	3.44	3.13

Comparing the results, the mean latency slightly increased from 3.86 ms without the trust manager to 3.93 ms with the trust manager, while the highest latency also decreased from 4.79 ms to 4.91 ms. Hence, there is no significant impact on the latency when querying the aerOS context broker with the trust score manager on a basic aerOS domain with 5 IEs (1.81%).

Scenario 2 (20 IEs): In the first figure, the latency to the Orion-LD Context Broker when using the trust manager is observed, while the second figure presents the latency to the Orion-LD broker without using the trust manager.

The statistics are as follows:

	With trust score	Without trust score

	Mean Latency (ms)	6.469	6.421
	Highest Latency (ms)	7.453	7.973
	Lowest Latency (ms)	5.071	5.095
	Upon comparing the results, it is observed that querying the aerOS Orion-LD Context Broker for calculating the trust score in a cluster of 20 IEs has no significant impact on latency (0.7%). For M36, the same tests will be conducted on the project's pilots.		

3.5. aerOS self-* and monitoring

3.5.1. KPI 1.5.1 Average overload time of IEs

Table 34 KPI 1.5.1 Average overload time of IEs

KPI ID number .	**KPI 1.5.1**		
KPI Name	**Average overload time of IEs**		
Description	The amount of time an IE has a system load above 80%, and therefore its performance may decrease considerably		
Motivation	Knowing how long an IE is overloaded allows actions to be taken to reduce its load, keeping it operational for longer		
Target value	Reduction of 20%		
Prerequisites	Hardware info submodule of the self-awareness module running on one IE of study.		
aerOS components (task)	Self-awareness (T3.5)		
Evaluation means	Using the hardware info sub-module of the self-awareness module of a node's self-* capabilities set, the node's performance will be obtained over time by measuring the percentage of use of both the CPU and the RAM memory. In this way, it will be possible to estimate the average time that the node remains in an overloaded state. For this purpose, those time slots will be selected in which the CPU or RAM memory usage exceeds 80 %. Several tests will be carried out over a certain period of time in which the aspects mentioned above will be analysed. Subsequently, the times obtained will be averaged in order to know with a certain degree of accuracy the average overload time of the node.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	This value shall be obtained by laboratory load tests to determine the actual average overload time of a node as a function of the actual workload. Through the self-awareness module, the % processor usage and	***37.5% of the total running time of a node***	*N/A*

	RAM load will be obtained. These values will indicate the total system load		

Outcome elaboration (M24)	Between the end of June and the beginning of July (27 June to 10 July) tests are being carried out to analyse how long the nodes remain in an overloaded state. For this purpose, a certain amount of workloads are being executed 3 times, during a period of 10 minutes (600 seconds) each time, on each of the nodes. For each period the loads will be different to simulate different scenarios that the aerOS IEs will have to face. Workloads will be combined with the execution of different services, such as server backend, databases, basic aerOS services, etc. During each period the self-awareness will measure the CPU and RAM usage, counting the amount of seconds that the node remains in an overloaded state through a test script in Python. Finally, the results obtained (the 3 periods of each node) will be averaged to obtain a realistic measurement of the average overload time of a node in the aerOS computing continuum.
	
	The tests were performed on two different sets of IEs, in order to cover the widest possible heterogeneity of architectures and technical specifications. The first set consisted of 3 nodes from CloudFerro cluster 2, which have a 4-core CPU and AMD64 architecture, 16 GB of RAM, 33 GB of storage and run Fedora 35. The second set consisted of 4 nodes from domain 3 of the UPV "continuum". This set consists of 2 virtual machines that have a 1-core CPU and AMD64 architecture, 8 GB of RAM, 60 GB of storage and run Ubuntu 22.04 LTS. In addition, it also consists of a Raspberry Pi 3 Model B+ and a Raspberry Pi 4 Model B. The RPi 3 runs Raspbian 12 and the RPi 4 runs Ubuntu 22.04 LTS.
	
	After running the tests on all the nodes described above, the following results (in seconds) were obtained:
	• CloudFerro cluster 2:
	  ○ Node 0: 106s, 252s and 176s. Mean = 178s (30 %).
	  ○ Node 1: 282s, 102s and 339s. Mean = 241s (40 %).
	  ○ Node 2: 150s, 160s and 195s. Mean = 169s (28 %).
	• Domain 3 of the UPV "continuum":
	  ○ VM 1: 217s, 180s and 44s. Mean = 147s (24 %).
	  ○ VM 2: 363s, 367s and 162s. Mean = 298s (49 %).
	  ○ RPi 3: 126s, 540s and 500s. Mean = 389s (65 %).
	  ○ RPi 4: 130s, 120s and 200s. Mean = 150s (25 %).
	The results indicate that the overload time of a node is between 25% and 65% of the execution time (on average). However, these values are highly dependent on both the power-related workloads assigned and the performance of the node itself. A lower-performing node (such as the RasPi 3) will overload more easily and for longer periods of time, but more powerful cloud nodes will have relatively short overload times. Considering the heterogeneity of aerOS compute nodes, the average overload time of a node is 225 seconds, equivalent to 37.5 % of the total running time of a node.		

3.5.2. KPI 1.5.2 Number of different topologies and hardware/software combinations of IEs supported

Table 35 KPI 1.5.2 Number of different topologies and hardware/software combinations of IEs supported

KPI ID number .	KPI 1.5.2
KPI Name	Number of different topologies and hardware/software combinations of IEs supported
Description	Indicates the number of different nodes on which self-* capabilities are capable of running depending on the type of IE, its operating system, hardware architecture or performance
Motivation	The aerOS computing continuum is composed by a multitude of domains, which belong to different organizations and companies. Each domain has different topologies and hosts very heterogeneous IEs. This variety in the nodes is due to differences in the hardware and software of each node. From variations in processor architectures to variations in specs, performance, operating system or task execution capabilities and more. Increasing the number of different types of supported nodes by aerOS will allow for greater heterogeneity.
Target value	10
Prerequisites	Basic set of self-* capabilities (self-awareness, self-orchestrator, self-diagnose and self-API) running on one IE of study.
aerOS components (task)	Self-awareness, self-orchestrator, self-diagnose and self-API (T3.5)
Evaluation means	The most representative combinations of future aerOS nodes forming the computing continuum will be selected by combining different architectures, operating systems, available resources and execution environments to test the flexibility of deployment of the set of self-* capabilities. At least the basic self-* capabilities will be installed in each selected combination and all functional combinations will be counted. To determine the heterogeneity of hardware-software combinations that will be able to support the set of self-* capabilities, the following combinations will be attempted:

Power	Node	Platform	Containerisation	OS
High-powered	Physical (SBC / laptop)	AMD64	Kubernetes (K3s…)	GNU/Linux (distro)
High-powered	Physical	AMD64	Docker	GNU/Linux
Low-powered	Physical	AMD64	Kubernetes	GNU/Linux
Low-powered	Physical	AMD64	Docker	GNU/Linux
Low-powered	Physical	ARM64	Kubernetes	GNU/Linux
Low-powered	Physical	ARM64	Docker	GNU/Linux

High-powered	Virtual	AMD64	Kubernetes	GNU/Linux
High-powered	Virtual	AMD64	Docker	GNU/Linux
Low-powered	Virtual	AMD64	Kubernetes	GNU/Linux
Low-powered	Virtual	AMD64	Docker	GNU/Linux
Low-powered	Virtual	ARM64	Kubernetes	GNU/Linux
Low-powered	Virtual	ARM64	Docker	GNU/Linux

Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*3*	***N/A***	*N/A*

Outcome elaboration (M24)

There are currently three domains within the aerOS computing continuum available:

1. The NCSRD domain consists of a Kubernetes cluster with nodes with both AMD64 and ARM64 architectures and different processing power and capabilities (CPU, RAM, storage, etc.).
2. The CloudFerro domain is formed by a Kubernetes cluster with 3 AMD64 machines with identical characteristics.
3. The UPV domain is in turn formed by three sub-domains, each one of them formed by machines of different architectures (AMD64 and ARM64), power, specifications, etc. Some of these machines are virtual machines belonging to Docker networks and other physical machines belonging to different Kubernetes clusters.

Until now, the following configurations have been tested:

Power	**Node**	**Platform**	**Containerisation**	**OS**
High-powered	Physical (SBC / laptop)	AMD64	Kubernetes (K3s…)	GNU/Linux (distro)
Low-powered	Physical	AMD64	Kubernetes	GNU/Linux
Low-powered	Physical	ARM64	Kubernetes	GNU/Linux
High-powered	Virtual	AMD64	Docker	GNU/Linux

To demonstrate the high installation capacity of the self-* components, up to three independent domains have been deployed within the so-called "continuum" of the UPV partner. Each of these domains is composed of nodes of different nature (physical, virtual, AMD64, ARM64, high-powered, low-powered, etc.) to test the installation and operation capacity of the different aerOS components. The following image shows in a simplified way the distribution of these nodes in the "continuum" of the UPV:

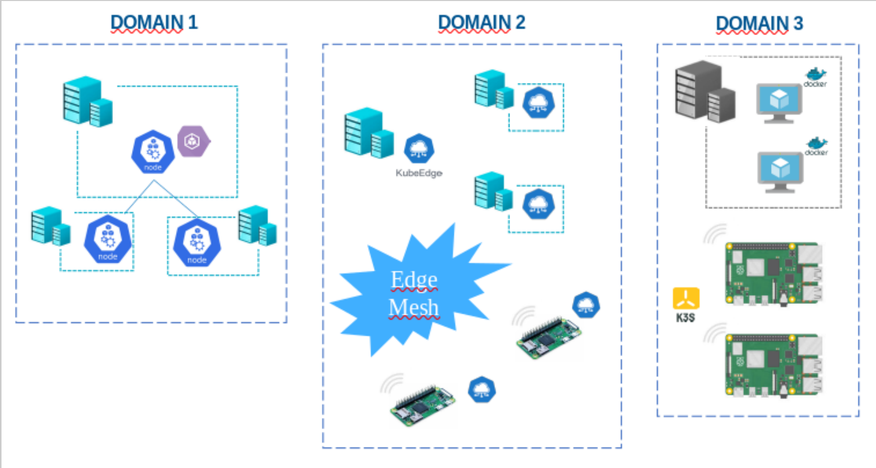

The following screenshots show examples of the execution of some self-* components on both UPV "continuous" machines:

```
satrd@continuum-docker-2:~$ uname -a
Linux continuum-docker-2 5.15.0-113-generic #123-Ubuntu SMP Mon Jun 10 08:16:17 UTC 2024 x86_64 x86_64 x86_64 GNU/Linux
satrd@continuum-docker-2:~$ docker ps
CONTAINER ID   IMAGE                                                                                                          COMMAND                  CREATED
 STATUS       PORTS                                      NAMES
b5af6aa6daf4   lbthomsen/docker-prime95:latest                                                                                "/bin/bash -c ./runp…"   4 days ago
 Up 4 days                                               service-component-dirpbgbv
56a35493f934   lbthomsen/docker-prime95:latest                                                                                "/bin/bash -c ./runp…"   4 days ago
 Up 4 days                                               service-component-lrltnexh
34d7464f36c3   806c5ce573b2                                                                                                   "docker-entrypoint.s…"   4 days ago
 Up 4 days    0.0.0.0:80->80/tcp, :::80->80/tcp          service-component-enysqzkd
ee81a470887c   registry.gitlab.aeros-project.eu/aeros-public/common-deployments/self-awareness/power_consumption_amd64:1.1.0-legacy-entity-id   "python3 script.py"   3 weeks ago
 Up 4 days                                               aeros_self-awareness_power-consumption_ctr
47f842971bd3   registry.gitlab.aeros-project.eu/aeros-public/common-deployments/self-awareness/hardware_info:1.1.0-legacy-entity-id   "python3 script.py"   3 weeks ago
 Up 4 days                                               aeros_self-awareness_hardware-info_ctr
49df1c4fee28   registry.gitlab.aeros-project.eu/aeros-public/common-deployments/self-orchestrator:1.0.0                       "node script.js"         3 weeks ago
 Up 4 days    0.0.0.0:8001->8001/tcp, :::8001->8001/tcp  aeros_self-orchestrator_ctr
4c50bdd2af7b   moby/buildkit:buildx-stable-1                                                                                  "buildkitd --allow-i…"   2 months ago
 Up 4 days                                               buildx_buildkit_mybuilder0
e15e5c2fe140   registry.gitlab.aeros-project.eu/aeros-public/common-deployments/llo/docker-api:1.0.0-mvp                      "/venv/bin/python3 -…"   3 months ago
 Up 4 days    0.0.0.0:8000->8000/tcp, :::8000->8000/tcp  aeros_llo_docker_api
satrd@continuum-docker-2:~$
```

And CloudFerro domain:

NAME	READY	STATUS	RESTARTS	AGE
entrypoint-balancer-api-5cd98b4d45-mgv92	1/1	Running	0	40d
orion-ld-broker-c5d946d64-krhtq	1/1	Running	5 (31d ago)	107d
orion-ld-mongodb-0	1/1	Running	0	107d
orion-ld-new-broker-7459bd464c-7lc2j	1/1	Running	0	2d18h
orion-ld-new-mongodb-0	1/1	Running	0	2d18h
self-awareness-hardwareinfo-4tk4p	1/1	Running	0	53d
self-awareness-hardwareinfo-mtnjd	1/1	Running	0	53d
self-awareness-powerconsumptionamd64-266ml	1/1	Running	0	53d
self-awareness-powerconsumptionamd64-lctbn	1/1	Running	0	53d
self-orchestrator-orchestrator-bl7qn	1/1	Running	0	53d
self-orchestrator-orchestrator-rmt49	1/1	Running	0	53d
self-security-gqnp8	2/2	Running	0	11d
self-security-kpf8z	2/2	Running	0	11d

In addition to virtual machines, components are also being tested on physical nodes, such as laptops and SBCs:

	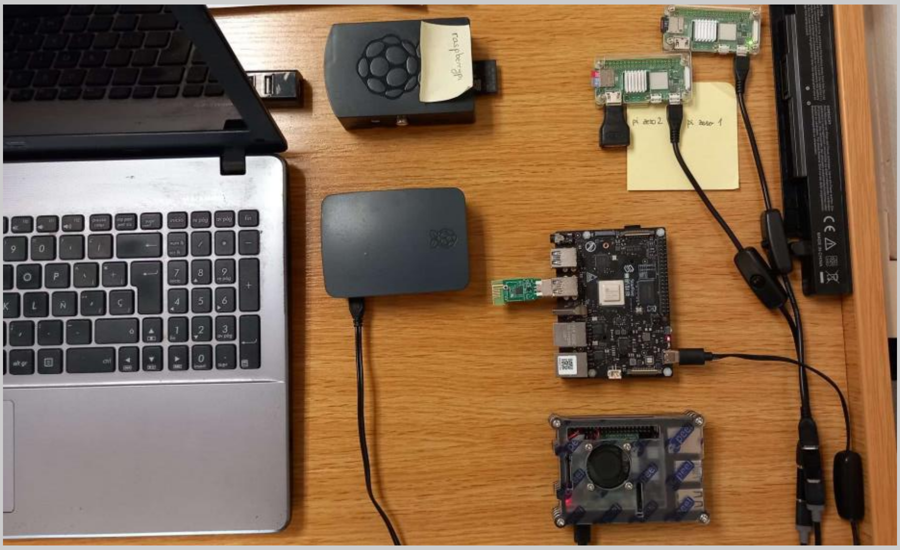

3.5.3. KPI 1.5.3 Number of metrics monitored from IEs

Table 36 KPI 1.5.3 Number of metrics monitored from IEs

KPI ID number .	**KPI 1.5.3**		
KPI Name	**Number of metrics monitored from IEs**		
Description	The amount of different information that the self-* modules are able to obtain on the characteristics, specifications, current performance and health status of the nodes where they run		
Motivation	The more metrics obtained, the more accurate will be the health indices of the nodes of aerOS		
Target value	15		
Prerequisites	Self-awareness module running on one IE of study		
aerOS components (task)	Self-awareness module (T3.5)		
Evaluation means	Each self-awareness submodule can measure a certain amount of data extracted from the node where it is running. By analysing the specifications of two modules (hardware metrics and energy consumption) it is possible to know how much information and metrics they can extract from each IE. The hardware info submodule is able to obtain information about the CPU (number of cores, architecture and current usage), the RAM memory (total amount, available and current usage), the storage (total amount, available and current usage), the operating system, the hostname, the internal IP address, the MAC address of the IE, the infrastructure element tier or if it is able to run real-time services. The energy consumption submodule is able to obtain information on the current and average energy consumption of the node.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**

Measured value (% achieved)	*N/A*	***20***	*N/A*
Outcome elaboration (M24)	Between the end of June and the beginning of July (27 June and 10 July), tests were carried out to determine the number of metrics and information that the self-awareness submodules are capable of obtaining. To do this, it was installed on a node (regardless of its characteristics) and all the metrics and information were obtained from the IE. Subsequently, it was compared with the specification of the module, and it was verified that it is indeed capable of obtaining all the requested information. Finally, the amount of information and metrics obtained have been counted. In order to obtain the values generated by the self-awareness, a query is realised by the name of the IE entity to the Orion-LD Context Broker of the corresponding domain, obtaining the following result: { "id": "urn:ngsi-ld:InfrastructureElement:1:0242ac1c0004", "type": "InfrastructureElement", "domain": "urn:ngsi-ld:Domain:1", "hostname": "aeros-self-test", "containerTechnology": "Kubernetes", "internalIpAddress": "172.28.0.4", "macAddress": "02:42:ac:1c:00:04", "lowLevelOrchestrator": "urn:ngsi-ld:LowLevelOrchestrator:1", "cpuCores": 4, "currentCpuUsage": 6, "ramCapacity": 7938, "availableRam": 6492, "currentRamUsage": 1446, "currentRamUsagePct": 19, "diskCapacity": 31066, "availableDisk": 3176, "currentDiskUsage": 26288, "currentDiskUsagePct": 90, "avgPowerConsumption": 6, "currentPowerConsumption": 7, "realTimeCapable": false, "trustScore": -1, "cpuArchitecture": "urn:ngsi-ld:CpuArchitecture:x64", "operatingSystem": "urn:ngsi-ld:OperatingSystem:Linux", "infrastructureElementTier": "urn:ngsi-ld:InfrastructureElementTier:Cloud", "infrastructureElementStatus": "urn:ngsi-ld:InfrastructureElementStatus:Ready",		

	"location": [0, 0] } Of all these values, the following are obtained by self-awareness: 1. *domain.* 2. *hostname.* 3. *internalIpAddress.* 4. *macAddress.* 5. *cpuCores.* 6. *currentCpuUsage.* 7. *ramCapacity.* 8. *availableRam.* 9. *currentRamUsage.* 10. *currentRamUsagePct.* 11. *diskCapacity.* 12. *availableDisk.* 13. *currentDiskUsage.* 14. *currentDiskUsagePct.* 15. *avgPowerConsumption.* 16. *currentPowerConsumption.* 17. *realTimeCapable.* 18. *cpuArchitecture.* 19. *operatingSystem.* 20. *infrastructureElementTier.* For a total of 20 attributes

3.5.4. KPI 1.5.4 Number of avoided service downgrade experience cases

Table 37 KPI 1.5.4 Number of avoided service downgrade experience cases

KPI ID number .	**KPI 1.5.4**
KPI Name	**Number of avoided service downgrade experience cases**
Description	All the different types of scenarios in which a continuum IE is prevented from not being able to respond to the requests made to it, either to obtain information or to request it to execute a certain workload. In other words, the trait of aerOS of reacting in advance (e.g., reorchestrating services that were running in the IE, or deactivating from being eligible for new services, or horizontally scaling replicas) so that the IE is still functional and operative in the mid-term.
Motivation	Reducing the number of situations in which a IE in the continuum stops responding to requests for information or becomes inoperative increases overall user satisfaction and provides a better image of a robust, reliable and fault-tolerant system
Target value	5 demonstrable scenarios

Prerequisites	Self-awareness and self-orchestration modules are functional. The KPI-1.5.5 has been demonstrated as VALID after measurements in M24, as the evaluation of KPI-1.5.4 depends to some extent of such KPI.		
aerOS components (task)	T3.3 HLO and LLOs, T3.5 self-awareness and self-orchestration.		
Evaluation means	The goal of the KPI is to demonstrate that aerOS can reduce the downgrade experience. In order to do that, the proposal is to make a comparison of a before/after (B/A) scenario: • First, the "Before" scenario is constructed. Here, the team aims at "profiling" which is the impact of several type of services in the continuum, enable to forecast at which point (and with which numbers, a IE would either collapse or downgrade the service experience). • Second, the "After" scenario will depart from the same type of services, and, after judiciously selecting specific thresholds, the self-orchestration mechanism will be put in place in the IEs of the continuum. There, by using the same evaluation means of KPI-1.5.5., the team will reflect whether these techniques avoiding the occurrence of the forecasted down situations. The methods used, and the assumptions taken, were: • Metrics: Trend forecasting of CPU usage metrics. • Assumption: Less current CPU usage per cores on a machine means it has more compute power available for the application. Therefore, this is a solid metric to consider when "degradation" is more likely to occur in a specific node (IE). • Methodology: To accurately determine if the task is going into a potential downgrade or is just experiencing a sudden spike in intensity, the task is monitored for a certain amount of time (2 min) before judging if there is a downgrade scenario or not. • After this time passes, a Neural Prophet regression model forecasts metrics for a certain amount of time in the future (1 min), and the predicted value is provided in the evaluation. This mechanism, in the runtime functioning of aerOS Meta-OS, would connect with the self-orchestration's reallocation trigger to offload to the HLO the picking of alternative best fit for the task.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***0 demonstrable scenarios (0%)***	*N/A*
Outcome elaboration (M24)	In this first period of evaluation, the team has generated the "Before" (in B/A) scenario, by delivering the following services: IE "A", hosting a task that performs some CPU-intensive computations. IEs "B" and "C" are idle in the system, running some low-demand background tasks. The IEs have the following system specifications, which illustrate two different types of nodes that often exist in aerOS continuums: • IE "A": Raspberry Pi 3b, 4-core @ 1.4GHz • IE "B": Virtual Host, 4-core @ 2.1GH		

The services that are used, which try to mimic 2 types of usual services in IoT deployments (thus, aerOS continuum) are:

1. “tarot-bench-linear“, a one-time, dummy service, hosting a mock task that performs some CPU-intensive computations, which would equal to sensor monitoring, actuation or updates of values across a distributed network.
2. “nginx”, a continuously running service that can receive requests and deliver a response. A web service has been used: nginx.

<u>Service 1 applied to IE “A”:</u>

The compute task is deployed in IE “A”. The degradation threshold is defined as CPU Usage > 30% after the 2 minutes run.

After the initial live time, time forecasting (using the mentioned FB Prophet model) is run for CPU Usage, and is predicted to surpass the threshold of 30% in the next minute.

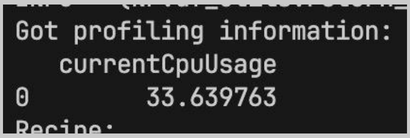

<u>Service 1 applied to IE “B”:</u>

The compute task is deployed in IE “B”. The degradation threshold is defined as CPU Usage > 35% after the 2 minutes run.

Then, for this second type of node (IE), forecasting is re-run. It can be seen that in the current IE, the task is not expected to have degraded performance, being able to profile this second type of node (B) with the following numbers.

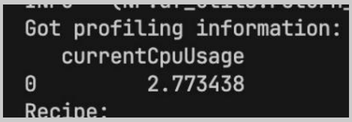

<u>Service 2 applied to IE “A”:</u>

On the next scenario, a NGINX service instance is deployed to a node type “A” that was already running several tasks.

The service is expected to be degraded, and after the monitoring time, a high, upward trend of the CPU consumption is reported:

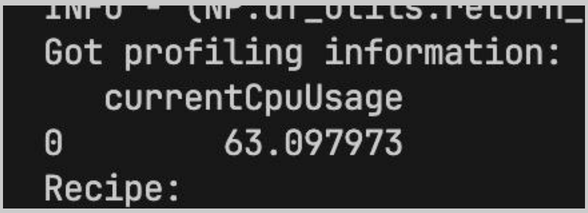

<u>Service 2 applied to IE “B”:</u>

The compute task (nginx service) is deployed in IE “B”. Also, the forecasting is re-run. It can be seen that in the current IE, the task is not expected to have degraded performance, being able to profile this second type of node (B) with the following numbers.

	INFO - (NP.df_utils.return_ Got profiling information: currentCpuUsage 0 5.549866 Recipe:

3.5.5. KPI 1.5.5 % of reorchestration requests issued by decentralized IEs

Table 38 KPI 1.5.5 % of reorchestration requests issued by decentralized IE

KPI ID number .	**KPI 1.5.5**
KPI Name	**% of reorchestration requests issued by decentralized IEs**
Description	Number of requests coming from decentralised IEs in the computing continuum to the main aerOS reorchestration systems based on their current or future workload to avoid failures in running services or system overloads that may generate unwanted situations
Motivation	A number of decentralised reorchestration requests provide insight into the actual performance and processing capacity of the IEs in the computing continuum
Target value	25%
Prerequisites	Self-awareness, self-diagnose, self-realtimeness, self-optimisation and adaptation, and self-orchestrator modules running on the continuum of study. First, on a continuum formed by 3 domains (MVP), and, later, in the continuum corresponding to the selected pilot.
aerOS components (task)	Self-awareness, self-diagnose, self-realtimeness, self-optimisation and adaptation, and self-orchestrator modules (T3.5).
Evaluation means	The structure of measuring this KPI will be tackled in two different stages. • First (to be covered in D5.5), a simulated scenario will be created departing from one running continuum. There, some services will be run, and a situation will be artificially generated to demonstrate that the reorchestration functionality is operative and that it indeed supports the KPI-1.5.4 in which the overload of an IE is reduced thanks to compensation in another part of the continuum. • Second (to be covered in D5.6), a running scenario will be observed during a certain timeframe (1 month, closer to the final date of aerOS). This running scenario will exist within one out of the 5 pilots of the project. The specific dates, pilot and timing will be described later in D5.6. The evaluation means here can be decomposed in two different methods: • Continuous observation of the services that are running in a continuum. Reporting if a service was originally allocated to a specific IE and then it ends up running in another IE (and providing evidence). • Checking the IOTA registries. As it has been designed, every time that the self-orchestration request is triggered an IOTA message will be immutably registered through the DLT. Therefore, checking the IOTA registries and the

<table>
<tr><td></td><td colspan="3">total amount of services deployed in a certain timeframe, this percentage will be extracted.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>This value will be obtained through laboratory tests. A domain with several IEs will be utilised and different workloads will be executed in each of them. The purpose will be to check of all the reorchestrations carried out, how many requests come directly from the IEs that compose the domain</td><td>1/3 service component is successfully re-orchestrated (33%).</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">
The tests carried out to analyse the correct functioning of two operations that are required for the final measurement of KPI-1.5.5:
• When a certain situation happens, the self-orchestrator module properly works, sending a triggering request to reallocate a specific service. For the measurement in D5.5, a threshold of 75% of CPU usage has been established.
For the D5.5 measurement, the continuum formed by CF (cloud domain), NCSRD (IoT domain) and UPV's mobile edge domain has been selected.


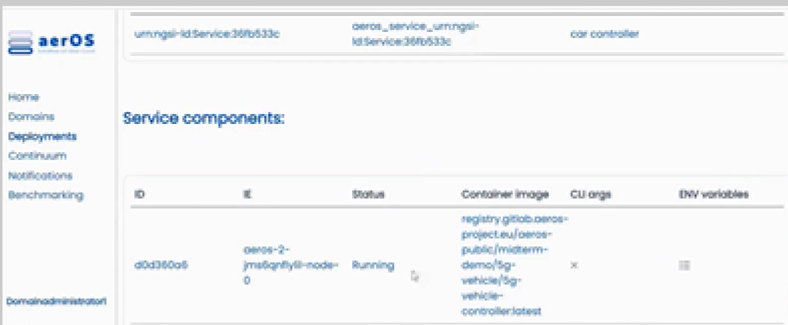

In the showcased scenario, one of the service components (gateway, controller, Grafana) had to be reorchestrated due to the overload in the Infrastructure Element where it was originally commissioned.
Thanks to the self-orchestration API, in the experimentation, a threshold of 75% of CPU usage was established. Therefore, automatically, when it reached such percentage (a K8s job was created to force the case), re-orchestration triggered.
</td></tr>
</table>

```
GET http://localhost:8001/rules:
{"statusCode":200,"statusName":"OK","rules":[{
    "name": "OverloadIERule",
    "conditions": {
        "any": [{
            "fact": "urn:ngsi-ld:InfrastructureElement:aeros-1-kvwgoc56nmap-node-0:currentRamUsagePct",
            "operator": "greaterThanInclusive",
            "value": 75
        }]
    },
    "event": {
        "type": "NEED_REORQ",
        "params": {
            "message": "urn:ngsi-ld:ServiceComponent:39e4e9c5"
        }
    }
}]}
```


As a result, the specific service that was running in the target IE, ends up being re-orchestrated, commissioned now to other IE (in another domain).

3.5.6. KPI 1.5.6 # of IoT healing scenarios covered

Table 39 KPI 1.5.6 # of IoT healing scenarios covered

KPI ID number .	**KPI 1.5.6**		
KPI Name	**# of IoT healing scenarios covered**		
Description	The KPI aims at measuring the potential situations that the self-healing procedures can take effect.		
Motivation	Self-healing crystalizes the capability of autonomously recovering affected parts of the system both at the hardware and software level caused by failures or abnormal states. Self-healing can also restart the system to pre-established routines scheduling, if necessary.		
Target value	5		
Prerequisites	IE's hardware setup is complete. Definition and implementation of scenarios are ready to be tested on IE hardware.		
aerOS components (task)	Self-healing (T3.5)		
Evaluation means	The importance of the self-healing functionality needs to be shown with specific scenarios. So far, the following "healing scenarios" have been identified: Sensor Failure, Device Power Alert, Network Protocol Violation, Link Quality Issues, Communication Failure Indication (no messages received by IE). In the first phase, tests ARE completed locally in FOGUS lab, running the defined scenarios on IE's hardware. In the second phase tests will be completed in the different Pilots of the project.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0 (no self-healing capabilities)*	***3/5 (60%)***	*N/A*
Outcome elaboration (M24)	Three of the defined scenarios have been tested locally under different conditions to analyse possible failures or abnormal states. Measurements are normal, any failures were correctly detected. The set up used for the measurements, featuring Raspberry Pi 4 Model B, DHT22 Digital Humidity & Temperature Sensor Module and Keyestudio Breakout kit for Raspberry Pi is shown below.		

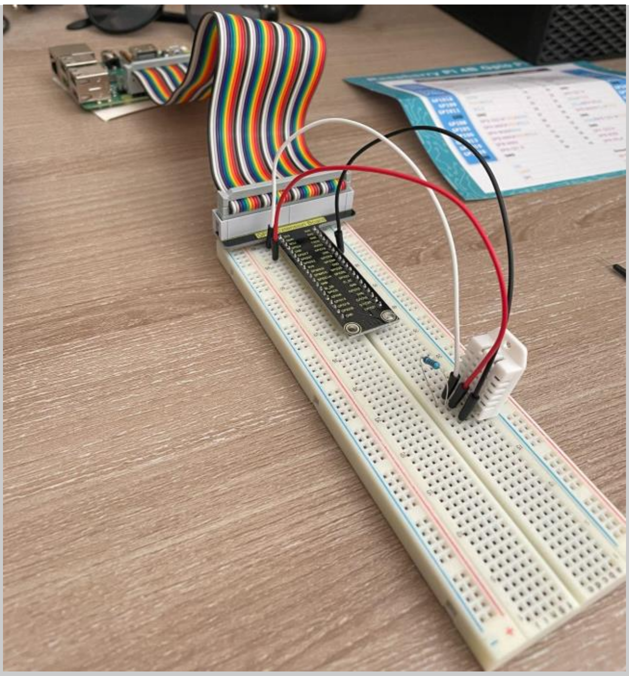

In addition, the scripts used for realizing the implemented scenarios (namely communication failure indication, device power alert, and sensor failure), as well all images of the log files that document the detection of failures or abnormal states in each scenario is shown in the following screenshots.

communication_failure_indication.py 1.46 KiB Blame Edit Replace Delete

```python
import os
import time
import logging
from datetime import datetime
from dotenv import load_dotenv

load_dotenv()

TARGET_IP = os.getenv('IE_IP', '192.168.1.15')
TIMEOUT = 5  # seconds

BASE_DIR = os.path.abspath(os.path.join(os.path.dirname(__file__), '../../'))

# Logging setup
env_log_dir = os.getenv('LOG_DIR', 'src/logs')
log_dir = os.path.join(BASE_DIR, env_log_dir)

if not os.path.exists(log_dir):
    os.makedirs(log_dir)

log_file = os.path.join(log_dir, os.getenv('COMMUNICATION_FAILURE_INDICATION_LOG_FILE', 'communication_failure_indication.log'))

logging.basicConfig(
    level=logging.INFO,
    format='%(asctime)s - %(message)s',
    handlers=[
        logging.FileHandler(log_file),
        logging.StreamHandler()
    ]
)

def is_reachable(ip, timeout=TIMEOUT):
    """Check if a given IP address is reachable."""
    response = os.system(f"ping -c 1 -W {timeout} {ip}")
    return response == 0

def monitor_communication():
    """Monitor the communication with the target IP."""
    while True:
        reachable = is_reachable(TARGET_IP)
        cur_time = datetime.now().strftime('%Y-%m-%d %H:%M:%S')
        if reachable:
            log_msg = f"[{cur_time}] Target IE is reachable."
            print(log_msg)
            logging.info(log_msg)
        else:
            log_msg = f"[{cur_time}] Target IE is NOT reachable."
            print(log_msg)
            logging.error(log_msg)
        time.sleep(10)

if __name__ == "__main__":
    monitor_communication()
```

cpu_power.py 1.90 KiB Blame Edit Replace Delete

```python
"""
Checks the CPU temperature periodically and detects any instances of
high temp. If high temp detected, appropriate action is taken.
"""

import os
import time
import logging
from datetime import datetime
from dotenv import load_dotenv

load_dotenv()

PERIOD = 10
HIGTHR_MSG = "\033[0;31mHIGH THRESHOLD\033[0m"
LOWTHR_MSG = "\033[0;33mLOW THRESHOLD\033[0m"
NORMAL_MSG = "\033[0;32mNORMAL OPERATION\033[0m"
HIG_THRESHOLD = 85.0
LOW_THRESHOLD = 80.0

BASE_DIR = os.path.abspath(os.path.join(os.path.dirname(__file__), '../../'))

# Logging setup
env_log_dir = os.getenv('LOG_DIR', 'src/logs')
log_dir = os.path.join(BASE_DIR, env_log_dir)

if not os.path.exists(log_dir):
    os.makedirs(log_dir)

log_file = os.path.join(log_dir, os.getenv('LOG_FILE', 'cpu_power.log'))

logging.basicConfig(
    level=logging.INFO,
    format='%(asctime)s - %(message)s',
    handlers=[
        logging.FileHandler(log_file),
        logging.StreamHandler()
    ]
)


def get_cpu_temp():
    """Returns CPU temperature in 'C."""
    temp = os.popen('vcgencmd measure_temp').readline()
    temp = temp.replace("temp=","").replace("'C\n","")
    return float(temp)


def status_report(cpu_temp):
    """Reports the status of CPU temp."""
    status = NORMAL_MSG
    if cpu_temp >= HIG_THRESHOLD:
        status = HIGTHR_MSG
    elif cpu_temp >= LOW_THRESHOLD:
        status = LOWTHR_MSG
    return status


def monitor_cpu_temp():
    """Monitors CPU periodically, to detect threshold violations."""
    while True:
        cpu_temp = get_cpu_temp()
        status = status_report(cpu_temp)
        cur_time = datetime.now().strftime('%Y-%m-%d %H:%M:%S')
        log_msg = f"[{cur_time}] {status}"
        print(log_msg)
        logging.info(log_msg)
        if status == HIGTHR_MSG:
            break
        elif status == LOWTHR_MSG:
            pass
        else:
            pass
        time.sleep(PERIOD)


if __name__ == "__main__":
    monitor_cpu_temp()
```

sensor_failure.py 1.39 KiB

Blame | Edit | Replace | Delete

```python
import time
import adafruit_dht
import board

from utils.outlier_detection import detect_outliers
from utils.self_healing import handle_failure
from config.settings import WINDOW_SIZE, Z_THRESHOLD

dht_device = adafruit_dht.DHT22(board.D4)
temperature_data = []
humidity_data = []

def read_sensor():
    try:
        temperature = dht_device.temperature
        humidity = dht_device.humidity
        return humidity, temperature
    except RuntimeError as error:
        return None, None


def monitor_sensor():
    while True:
        humidity, temperature = read_sensor()

        if humidity is None or temperature is None:
            handle_failure("No measurement at the device")
            continue

        temperature_data.append(temperature)
        humidity_data.append(humidity)

        if len(temperature_data) > WINDOW_SIZE:
            temperature_data.pop(0)
            humidity_data.pop(0)

        if len(temperature_data) == WINDOW_SIZE:
            temp_outliers = detect_outliers(temperature_data, Z_THRESHOLD)
            hum_outliers = detect_outliers(humidity_data, Z_THRESHOLD)

            if temp_outliers:
                handle_failure(f"Temperature outlier detected: {temp_outliers}")
            if hum_outliers:
                handle_failure(f"Humidity outlier detected: {hum_outliers}")


        time.sleep(10)

if __name__ == "__main__":
    monitor_sensor()
```

communication_failure_indication.log

src > logs > communication_failure_indication.log

```
2024-07-22 09:42:01,989 - [2024-07-22 09:42:01] Target IE is reachable.
2024-07-22 09:42:14,995 - [2024-07-22 09:42:14] Target IE is reachable.
2024-07-22 09:42:27,991 - [2024-07-22 09:42:27] Target IE is reachable.
2024-07-22 09:42:40,995 - [2024-07-22 09:42:40] Target IE is reachable.
2024-07-22 09:42:53,991 - [2024-07-22 09:42:53] Target IE is NOT reachable.
```

cpu_power.log

src > logs > cpu_power.log

```
2024-07-22 09:30:13,388 - [2024-07-22 09:30:13] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:30:23,399 - [2024-07-22 09:30:23] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:30:33,401 - [2024-07-22 09:30:33] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:30:43,413 - [2024-07-22 09:30:43] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:30:53,424 - [2024-07-22 09:30:53] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:31:03,435 - [2024-07-22 09:31:03] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:31:13,442 - [2024-07-22 09:31:13] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:31:23,453 - [2024-07-22 09:31:23] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:31:33,464 - [2024-07-22 09:31:33] ESC[0;32mNORMAL OPERATIONESC[0m
2024-07-22 09:31:43,476 - [2024-07-22 09:31:43] ESC[0;31mHIGH THRESHOLDESC[0m
```

	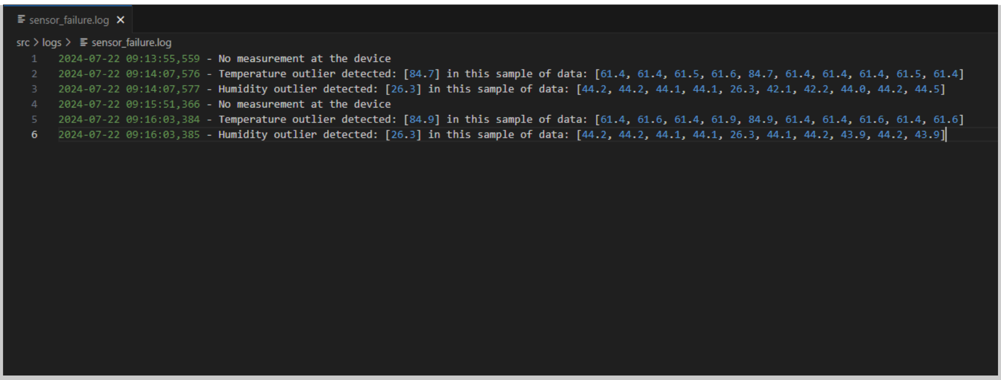

3.5.7. KPI 1.5.7 % of intrusion detected by the self-security

Table 40 KPI 1.5.7 % of intrusion detected by the self-security

KPI ID number .	**KPI 1.5.7**		
KPI Name	**% of intrusion detected by the self-security**		
Description	Indicates the ability to detect cybersecurity intrusions that have been made to the IE.		
Motivation	Measuring the percentage of intrusions that the self-security component has been able to detect allows the performance of the self-security component to be measured.		
Target value	>90% intrusions		
Prerequisites	Have the self-security component installed and running in IE		
aerOS components (task)	Self-security (T3.5)		
Evaluation means	In order to analyse this KPI, different attacks will be launched on the IE where self-security is installed and the ability of this component to detect will be tested. With this, the detection rate of the attacks will be calculated. Currently, the component is configured to detect "port scanning" and "denial of service (DoS)" attacks. In the remaining months, the portfolio of attacks that the tool is capable of detecting will be increased. At the same time, following the plan presented for the development of the component (where it is expected to have a version ready to operate in month M24), it will be installed in the use case environment so that its operation can be validated.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***0 (0%)***	*N/A*
Outcome elaboration (M24)	Self-security component has not been yet deployed. Consequently, it is not possible to obtain its performance in real scenarios of the project. Although not relevant, NMap attack tests have been additionally performed on a node of the CloudFerro domain,		

	where self-security has been installed and it can be seen in the following screenshot that it is able to detect nmap attacks as an example correctly: 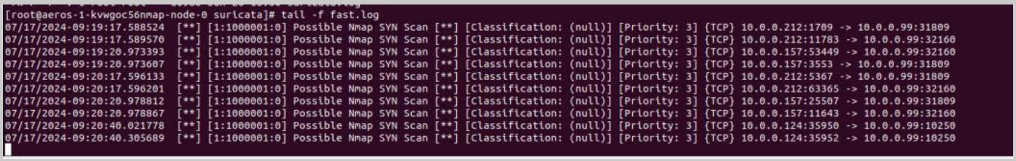

3.6. aerOS decentralised AI

3.6.1. KPI 1.6.1 Realising decentralized AI/ML with scalability comparable to centralized approach (KVI-4.1)

Table 41 KPI 1.6.1 Realising decentralized AI/ML with scalability comparable to centralized approach (KVI-4.1)

KPI ID number .	**KPI 1.6.1**
KPI Name	**Realising decentralized AI/ML with scalability comparable to centralized approach (KVI-4.1)**
Description	Scalability is the ability of AI algorithms, data, models, and infrastructure to operate at the size, speed, and complexity required. aerOS should operate with or be validated with at least three applications of decentralized AI.
Motivation	Decentralized AI/ML should not negatively influence operations of AI-based system compared with a centralized AI/ML.
Target value	>=3 applications
Prerequisites	aerOS deployment ready with final (or close to final) version of base components and aux AI components ready for evaluation
aerOS components (task)	AI Task Controller, AI Local Executor (T4.3)
Evaluation means	Three decentralized AI applications will be identified in and their scalability will be evaluated with proper tests or justification. One application will be based on experiments on decentralized vs centralized model training, the other two will be based on model inference in a decentralized and centralized approach. For model training evaluation metrics will cover time of training, trained model performance. For model inference metrics will cover: inference time and resource utilization (memory, CPU). Plan for a model training-based application: • Run model training at least 3 times in a centralized approach (can be outside aerOS) and measure evaluation metrics • Prepare a model to be trained with a decentralized approach on several aerOS IEs • Run at least 3 times the training process using IEs selected by aerOS and each time measuring evaluation metrics • Results from decentralized tests will be averaged to be compared with averaged results for a centralized approach

	Plan for a model inference-based application: • Run at least 3 times a set number of inferences on a model over a selected period of time in a centralized approach (can be outside aerOS deployment) and each time measure evaluation metrics • Run at least 3 times a set number of inferences on a model over a selected period of time using IEs selected by aerOS and each time measure evaluation metrics • Results from decentralized tests will be averaged to be compared with averaged results for a centralized approach		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	This KPI will be evaluated in M36 as final and stable versions of technical components are needed. So far scalability is considered during the design and technology choices. AI execution and algorithms should handle well scaling up.		

3.6.2. KPI 1.6.2 Energy consumption reduction due to moving AI from cloud to the edge (KVI-4.2)

Table 42 KPI 1.6.2 Energy consumption reduction due to moving AI from cloud to the edge (KVI-4.2)

KPI ID number .	**KPI 1.6.2**
KPI Name	**Energy consumption reduction due to moving AI from cloud to the edge (KVI-4.2)**
Description	Energy consumption should be decreased for AI being run closer to the edge, possibly on local data and with frugal adjustment.
Motivation	aerOS aims to establish what are the benefits and trade-off resulting from moving AI closed to the edge.
Target value	> 50% (on average on tested scenarios)
Prerequisites	aerOS deployment ready with final (or close to final) version of base components and aux AI components ready for evaluation in development/integration environment.
aerOS components (task)	AI Local Executor, AI Task Controller, frugal techniques (T4.3)
Evaluation means	Experiments will be conducted to measure energy consumption when running model inference on elements with different processing capabilities – cloud vs edge. For edge deployments frugal techniques will likely be applied. During analysis in M24-M36 a proper evaluation metrics will be selected. Evaluation plan:

	• Running at least 3 times a set of inferences over a period of time over a model deployed in the cloud where data needs to be sent from edge IEs to the cloud; each time measure evaluation metrics • Running at least 3 times a set of inferences over a period of time over a model deployed on edge IEs with local access to data; each time measure evaluation metrics • Running at least 3 times a set of inferences over a period of time over a frugal model (model from previous points after application of quantization/pruning) deployed on edge IEs with local access to data; each time measure evaluation metrics • Compared averaged results from three above options • The energy consumption of the running processes is foreseen to be established from information gathered using Kepler (a k8s monitoring tool).		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	From research literature: • BERT LM experiment used 8 V100 GPUs for 36 hours and used a total of 37 kWh. • Three sizes of DenseNets on MNIST lasted between 20 and 25 minutes and consumed between 20 and 38Wh • Energy reduction achieved with proposed methods for Microsoft Azure cloud compute platform was less than 27%. • Energy consumption on centralized vs distributed approach decreased, on average, less than 10%.	***N/A***	*N/A*
Outcome elaboration (M24)	This KPI will be evaluated in M36 as final and stable versions of technical components are needed to run a set of tests. The energy consumption of the running processes will be established based on information that will be gathered using Kepler (a k8s monitoring tool).		

3.6.3. KPI 1.6.3 Validation of comprehensive support, by aerOS, for distributed frugal AI components with explainability (KVI-4.3)

Table 43 KPI 1.6.3 Validation of comprehensive support, by aerOS, for distributed frugal AI components with explainability (KVI-4.3)

KPI ID number .	**KPI 1.6.3**
KPI Name	**Validation of comprehensive support, by aerOS, for distributed frugal AI components with explainability (KVI-4.3)**
Description	Identification of applications (within use cases/scenarios) in which distributed frugal AI components potentially supported by explainability should be applicable.

Motivation	aerOS will operate on heterogenous Infrastructure Elements with both internal decision-making and potentially AI-related services. This heterogeneity may require for application of frugal techniques to enable effective operations and interpretability/explainability to provide additional information about the decision making.		
Target value	>= 2 frugal applications, >= 2 XAI applications		
Prerequisites	Advanced or finalized design of aerOS internal operations and architectures for the pilots		
aerOS components (task)	AI Local Executor, AI Task Controller, frugality techniques, explainability techniques (T4.3)		
Evaluation means	An in-depth review of all the aerOS scenarios and continuum will be addressed. Every AI service that is trained with: (1) datasets sizes smaller by min 30% of the estimated full dataset, (2) using resources with limited capacities requiring application of frugal techniques, will be considered as a frugal application. A survey will be conducted with different end-users from those aerOS scenarios that claim that are making use of explainable AI. If the feedback obtained is higher than 50%, it will be considered as an acceptable XAI application.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	1 XAI application in progress, 1 identified to be done – 40% 0 frugal applications (but still under analysis) **Total = 20%**	*N/A*
Outcome elaboration (M24)	One XAI application comes from an internal use case (HLO allocations explanation, https://gitlab.aeros-project.eu/wp4/t4.3/explainability-service), the other is from external scenario and is being analysed with the pilot leader (pilot 5, in scenario Intelligent occupational safety and health). Frugal applications are being confirmed and analysed with pilot leaders.		

3.6.4. KPI 1.6.4 Delivery of cookbook/good practices manual for explainable frugal AI near the edge (KVI-4.4)

Table 44 KPI 1.6.4 Delivery of cookbook/good practices manual for explainable frugal AI near the edge (KVI-4.4)

KPI ID number .	**KPI 1.6.4**
KPI Name	**Delivery of cookbook/good practices manual for explainable frugal AI near the edge (KVI-4.4)**

Description	Description of good practices based on target value of application that would guide IoT developer to select the best approach.		
Motivation	Application of explainability and frugal techniques needs to be customized to a scenario being considered. Therefore, what is required is a set of good practices that would guide the IoT developer in selecting the best approach.		
Target value	>= 3 (pilot-specific)		
Prerequisites	Available results of evaluation of AI-based applications in the pilots and results of evaluation of aerOS deployment in the pilots. Results coming from all pilots in which AI was utilized, the most insightful will be selected for good practices preparation.		
aerOS components (task)	Explainability techniques, frugality techniques (T4.3)		
Evaluation means	Number of guidelines formulated, where a guideline is understood as a set of rules or remarks related to a specific feature or use case.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***N/A***	*N/A*
Outcome elaboration (M24)	KPI will be measured at M36. The implementation of the techniques used for explainability, and frugality are under development and need to be finished to formulate good practices/cookbook.		

3.6.5. KPI 1.6.5 Decentralized frugal AI techniques available

Table 45 KPI 1.6.5 Decentralized frugal AI techniques available

KPI ID number .	**KPI 1.6.5**
KPI Name	**Decentralized frugal AI techniques available**
Description	Techniques applied to provide frugality to AI in aerOS or aerOS-based deployments where AI operations in restricted conditions need to be supported.
Motivation	aerOS will operate on heterogenous Infrastructure Elements with both internal decision-making and potentially AI-related services. This heterogeneity may require for application of frugal techniques to enable effective operations.
Target value	>= 3 techniques
Prerequisites	N/A
aerOS components (task)	Frugality techniques (T4.3)

Evaluation means	Number of frugality techniques that have been evaluated for their effectiveness and for which applicability to aerOS scenarios was studied.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	**2 techniques (50%)**	N/A
Outcome elaboration (M24)	Two techniques have been implemented and experimented with: quantization and pruning. Quantization was used to reduce the precision of model weights from 32-bit floating-point down to 8-bit integer representation, to decrease model size and speed up inference (while maintain reasonable results). Pruning was used to remove less important weights from a model to improve computational efficiency, without significantly impacting model quality. So far, the techniques have been used in external scenarios and what is pending is to correlate them explicitly with aerOS use cases. Their current version is available in the GitLab repository https://gitlab.aeros-project.eu/wp4/t4.3/model-reduction-service.		

3.6.6. KPI 1.6.6 AI explainability techniques available

Table 46 KPI 1.6.6 AI explainability techniques available

KPI ID number .	**KPI 1.6.6**		
KPI Name	**AI explainability techniques available**		
Description	Techniques applied to provide interpretability/explainability in aerOS or aerOS-based deployments where AI interpretability/explainability can support the operations of the systems by enabling understanding of intelligent automatized decision-making.		
Motivation	System with automatic intelligent decision-making should provide means to monitor and verify its behaviour.		
Target value	>=2 techniques		
Prerequisites	N/A		
aerOS components (task)	Explainability techniques (T4.3)		
Evaluation means	Number of explainability techniques that have been evaluated for their effectiveness and for which applicability to aerOS scenarios was studied.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**

Measured value (% achieved)	*N/A*	***1 technique – Shapley values for RL (50%)***	*N/A*
Outcome elaboration (M24)	One technique based on Shapley values for service allocation was proposed. The method iteratively explains each decision made by the service allocator. The explainer is available in the GitLab repository: https://gitlab.aeros-project.eu/wp4/t4.3/explainability-service		

3.7. aerOS common API

3.7.1. KPI 1.7.1 % of aerOS core services exposed through OpenAPI

Table 47 KPI 1.7.1 % of aerOS core services exposed through OpenAPI

KPI ID number .	**KPI 1.7.1**		
KPI Name	**% of aerOS core services exposed through OpenAPI**		
Description	This KPI measures the proportion of aerOS's core services that are accessible through an OpenAPI. The goal is to ensure that a significant part of the system's functionality is available via well-defined and standardised interfaces.		
Motivation	Exposing services through OpenAPIs facilitates integration with other systems, encourages developer engagement, and supports a modular, scalable architecture. It enables third-party developers to easily connect with and build upon the aerOS platform, fostering innovation and expanding the system's capabilities.		
Target value	>50%		
Prerequisites	The exposed APIs of each aerOS component must be provided and can be reachable and interactive, providing the expected results.		
aerOS components (task)	HLO (T3.3), Context Broker (T4.2), and Data Fabric (T4.2)		
Evaluation means	The evaluation will involve identifying the total number of core aerOS services, involved in the project. At least 50% of these services must expose their components via Open API. The OpenAPI endpoints will be documented through screenshots and links to the public Swagger page for verification.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***25% of aerOS core services (50%)***	*N/A*
Outcome elaboration (M24)	The core services identified in aerOS are: HLO, Orion-LD Context Broker, IdM Keycloak, OpenLDAP, KrankedD API Gateway, Data Fabric, Self-* modules, Management Portal. To date, 25% of these core services have been successfully exposed through Open APIs. This includes key components such as the HLO		

Allocator, HLO Frontend, Context-broker, and Data Product Manager. The following snapshot depicts the structure of the GitLab repository for the aerOS OpenAPI:

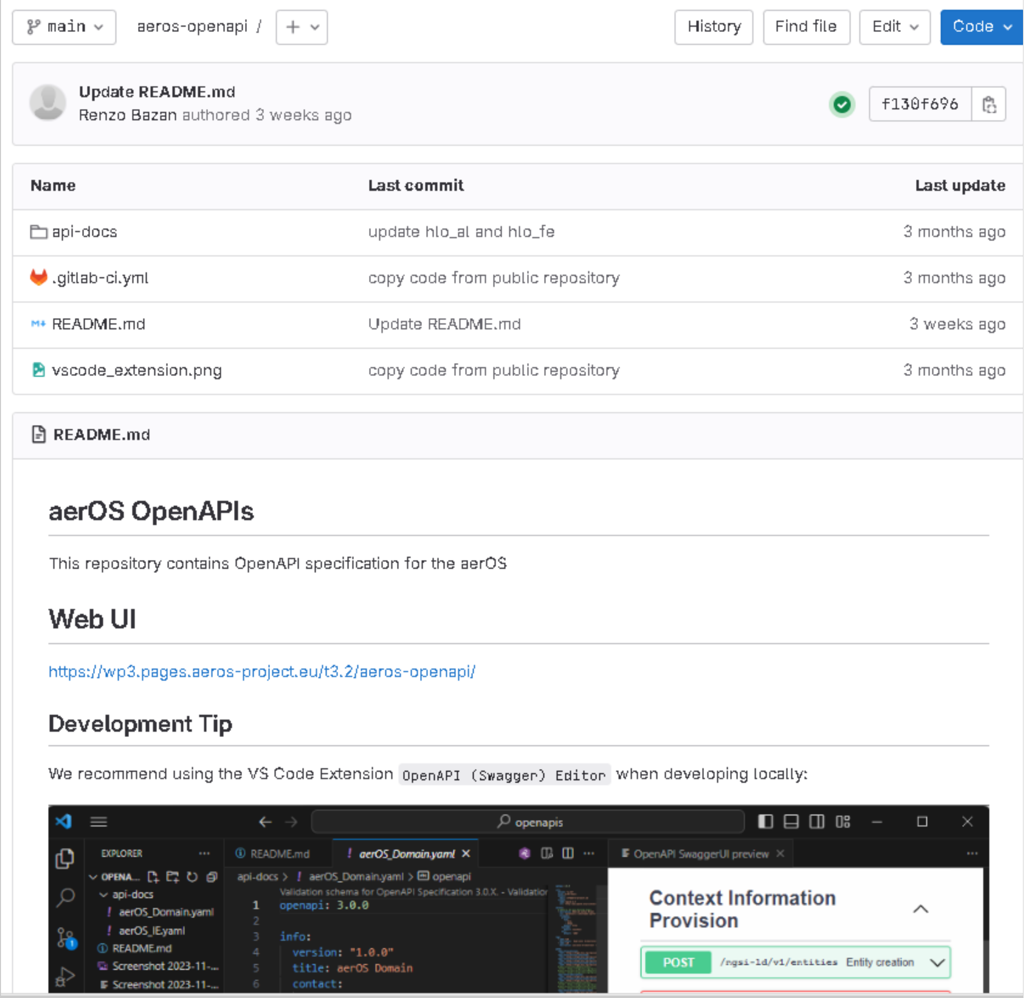

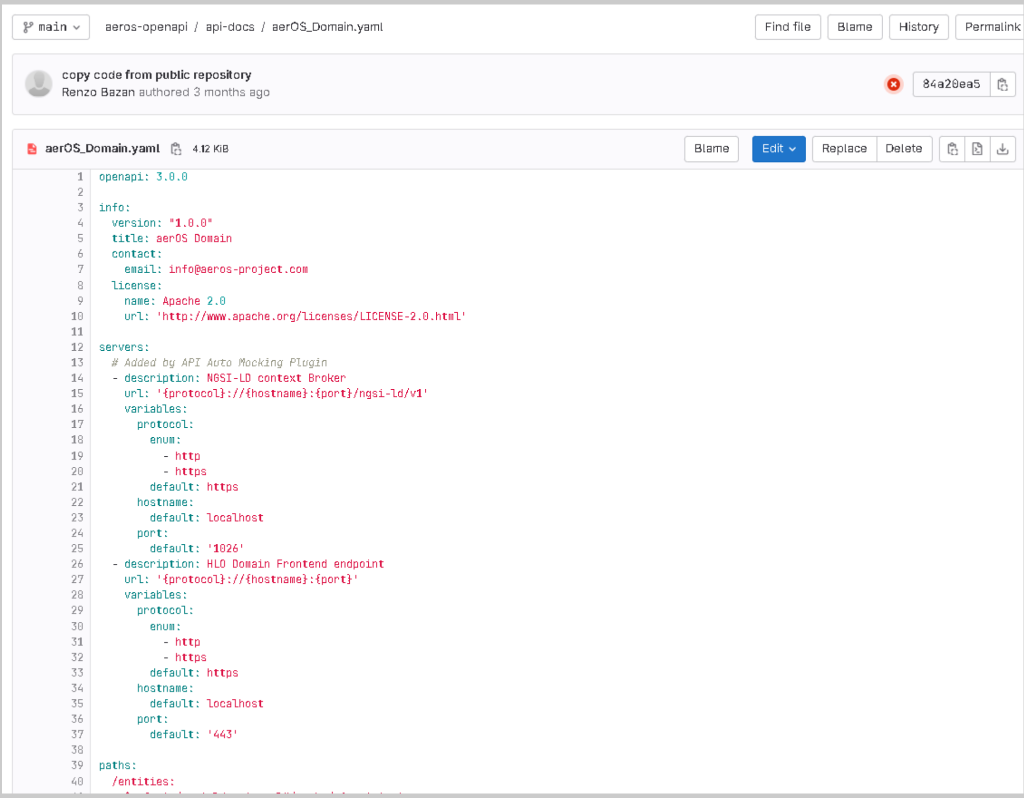

3.7.2. KPI 1.7.2 OpenAPI UIs for documenting APIs and generating code

Table 48 KPI 1.7.2 OpenAPI UIs for documenting APIs and generating code

KPI ID number .	**KPI 1.7.2**
KPI Name	**OpenAPI UIs for documenting APIs and generating code**
Description	These UIs will make it easier for developers to understand and work with the aerOS APIs. By providing clear documentation and tools for code generation, developers can more efficiently integrate their services with aerOS, reducing development time and potential errors
Motivation	The motivation behind this KPI is to enhance developer experience and productivity by providing comprehensive and accessible documentation of the aerOS APIs. With well-documented APIs and integrated code generation tools, developers can quickly grasp the functionality and implementation details of aerOS services. This leads to faster integration, fewer development errors, and a more streamlined development process. Ultimately, it supports the goal of creating a robust, developer-friendly ecosystem around aerOS.
Target value	2
Prerequisites	The aerOS domain OpenAPI must be provided
aerOS components (task)	HLO (T3.3), Context Broker (T4.2), and Data Fabric (T4.2).

Evaluation means	The evaluation of this KPI will involve several steps. First, it will require the identification of the main aerOS exposed APIs that will be documented using OpenAPI UIs. The success of these UIs will be measured by their completeness and usability. Evidence of successful implementation will be provided through access to the UIs, user guides, and examples of generated code.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***1 (50%)***	*N/A*
Outcome elaboration (M24)	By M24, significant progress has been made towards achieving the KPI for OpenAPI UIs for documenting APIs and generating code. The goal was to have at least one functional OpenAPI UI in place with complete documentation. To date, one OpenAPI UI has been successfully implemented, providing comprehensive documentation for the following aerOS components: HLO Allocator, HLO Frontend, Context Broker, Data Product Manager. While the current OpenAPI UI does not yet include code generation capabilities, the necessary tools have already been integrated to enable this feature in future iterations. Below is a screenshot of the Swagger page for the exposed services, showcasing the OpenAPI documentation: 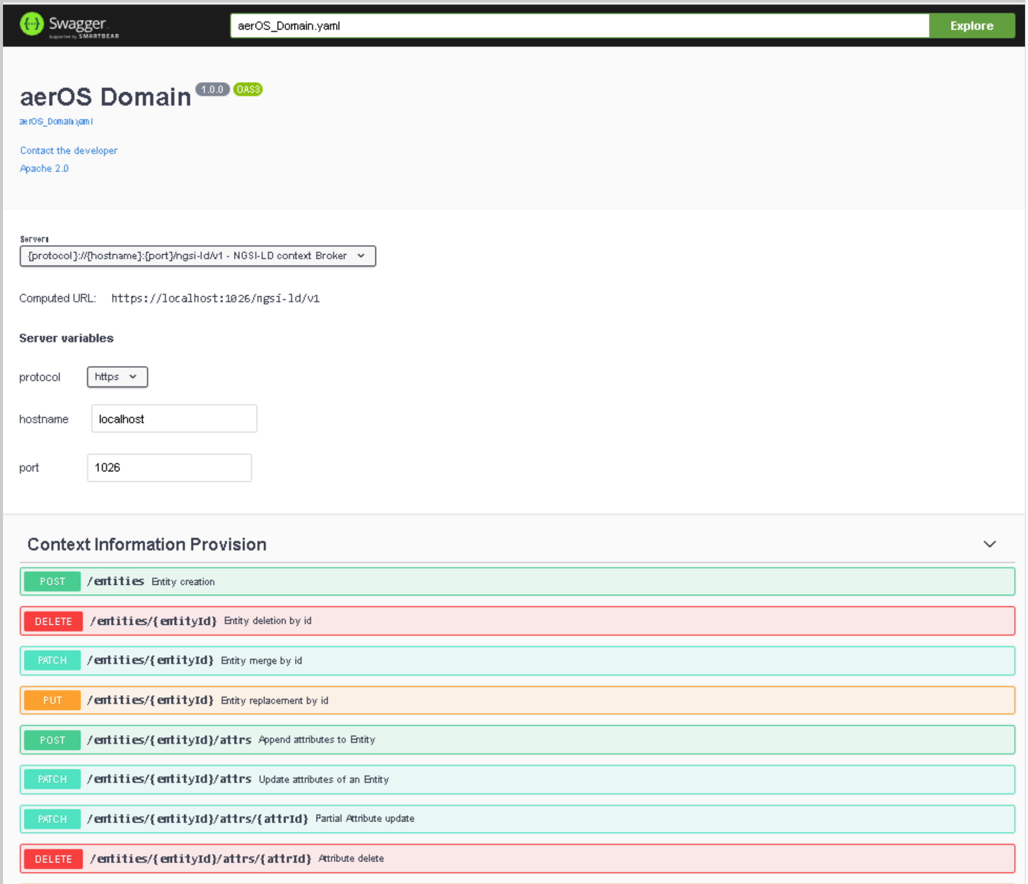		

3.7.3. KPI 1.7.3 Create Protocol Buffers definition for intra-orchestration module communication

Table 49 KPI 1.7.3 Create Protocol Buffers definition for intra-orchestration module communication

KPI ID number .	KPI 1.7.3		
KPI Name	Create Protocol Buffers definition for intra-orchestration module communication		
Description	This KPI aims to develop three Protocol Buffers definitions to facilitate communication between different modules within the aerOS orchestration layer. Protocol Buffers is a method of serializing structured data		
Motivation	Using Protocol Buffers enhances the efficiency and reliability of data interchange between modules. This approach ensures consistent, lightweight, and backward-compatible communication, crucial for maintaining the robustness and scalability of the orchestration layer.		
Target value	3		
Prerequisites	Identification and documentation of all modules within the aerOS orchestration layer that require Kakfa communication.		
aerOS components (task)	HLO (T3.3)		
Evaluation means	The evaluation will involve identifying and documenting the different modules within the aerOS orchestration layer that require Kafka communication. The next step is to develop Protocol Buffers definitions for each identified communication pathway. The definitions will be documented and reviewed to confirm they are consistent and lightweight. Evidence of successful implementation will be provided through code repositories and integration examples.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***4 (133%)***	*N/A*
Outcome elaboration (M24)	To date, four Protocol Buffers definitions have been successfully identified and documented, exceeding the original target. These definitions facilitate communication for the following components: • HLO Allocator • HLO Data Aggregator • HLO Deployment Engine • HLO Front End Four Protocol Buffers definitions have been developed, tested, and integrated, ensuring they are consistent, lightweight, and backward compatible. Below is a screenshot of the code repository showcasing the developed Protocol Buffers definitions:		

	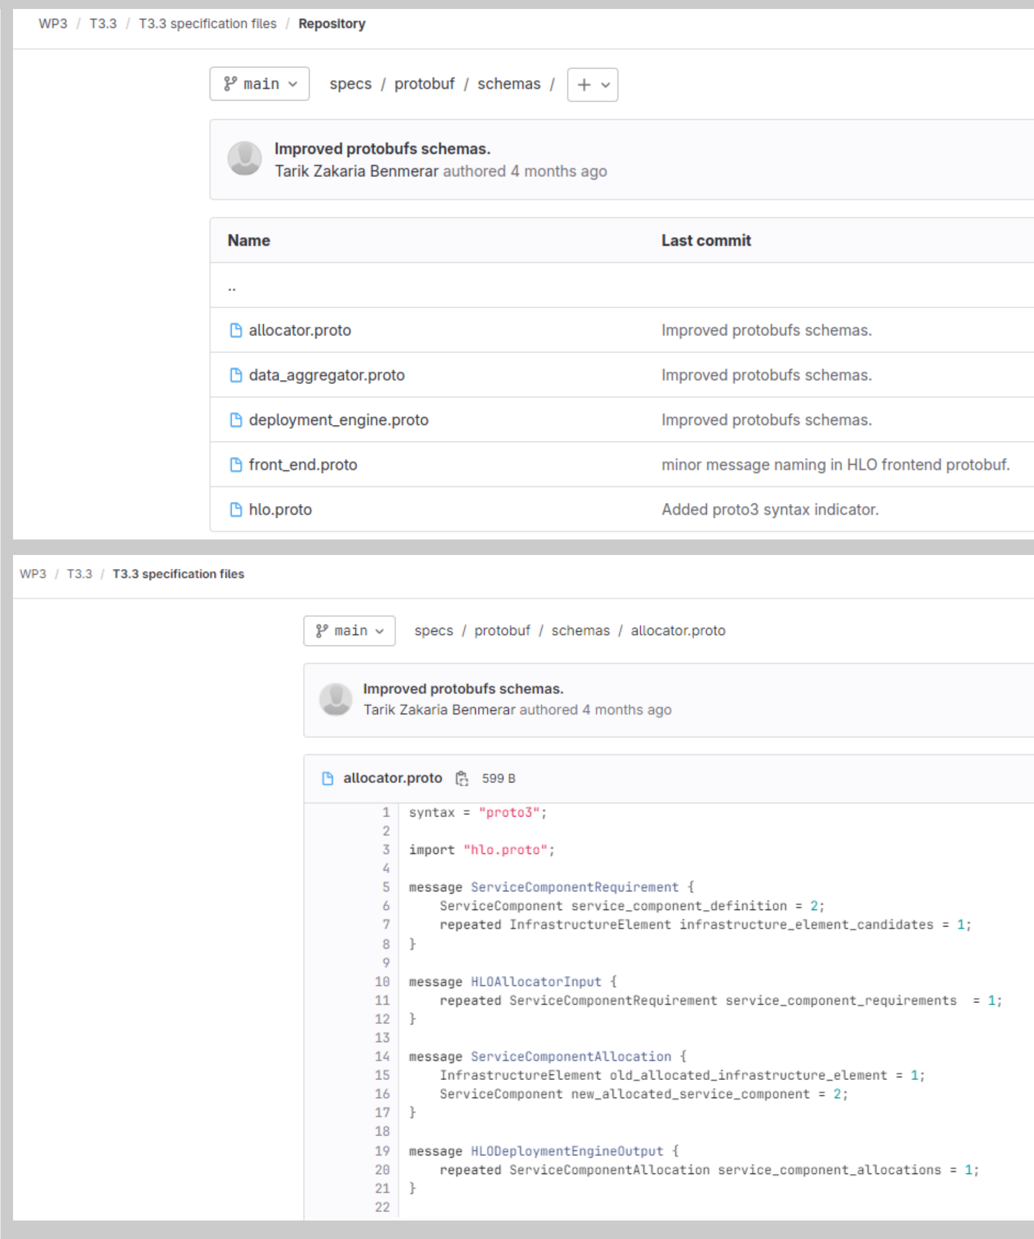

3.7.4. KPI 1.7.4 Reduce time to deploy service functions by non-technical team members using low code tool integrations

Table 50 KPI 1.7.4 Reduce time to deploy service functions by non-technical team members using low code tool integrations

KPI ID number .	**KPI 1.7.4**
KPI Name	**Reduce time to deploy service functions by non-technical team members using low code tool integrations**
Description	This KPI focuses on decreasing the time required for non-technical team members to deploy service functions within aerOS by over 40%, leveraging low-code tool integrations.

Motivation	Facilitating faster deployment of service functions by non-technical staff can significantly enhance operational efficiency. By integrating low-code tools, aerOS can democratize the deployment process, enabling a broader range of team members to contribute to service development and management, thus accelerating the project lifecycle and reducing dependency on technical specialists.		
Target value	Improvement of >40% over a baseline		
Prerequisites	Have a protocol compatible with the low-code tool to be used. ROS, MQTT and web-sockets are currently supported.		
aerOS components (task)	Low Code Tools and AsyncAPI (T3.2)		
Evaluation means	The evaluation of this KPI will involve several key steps. Initially, a baseline measurement of the time currently required for non-technical team members to deploy service functions without low-code tools will be established. Following this, appropriate low-code tools will be identified and integrated into the aerOS environment. The deployment time will then be measured and compared to the baseline to determine the percentage reduction achieved.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*40 seconds*	***20 seconds (120%)***	*N/A*
Outcome elaboration (M24)	To date, the integration and use of aerOS Low Code Tools and AsyncAPI have been tested in a pilot with Siemens, focusing on reprogramming the behaviour of multiple Automated Guided Vehicles (AGVs). This pilot involved two main AGV skills: navigation from point to point and lifting boxes. Traditionally, non-technical users needed 40 seconds to reconfigure an AGV using PLC methods, despite the HMI available for the PLC. With AsyncAPI and low-code tools, skills were defined in an AsyncAPI document, and plugin blocks for the low-code tool (behaviour trees) were generated. This allowed non-technical users to reprogram the AGV in runtime using an intuitive web GUI, reducing the time to 20 seconds and making training much shorter. This progress demonstrates our commitment to enhancing operational efficiency by empowering non-technical team members to deploy and manage service functions more quickly. A screenshot of the repository showcasing the integration of low-code tools and AsyncAPI, and the resulting behaviour tree configurations is shown below:		

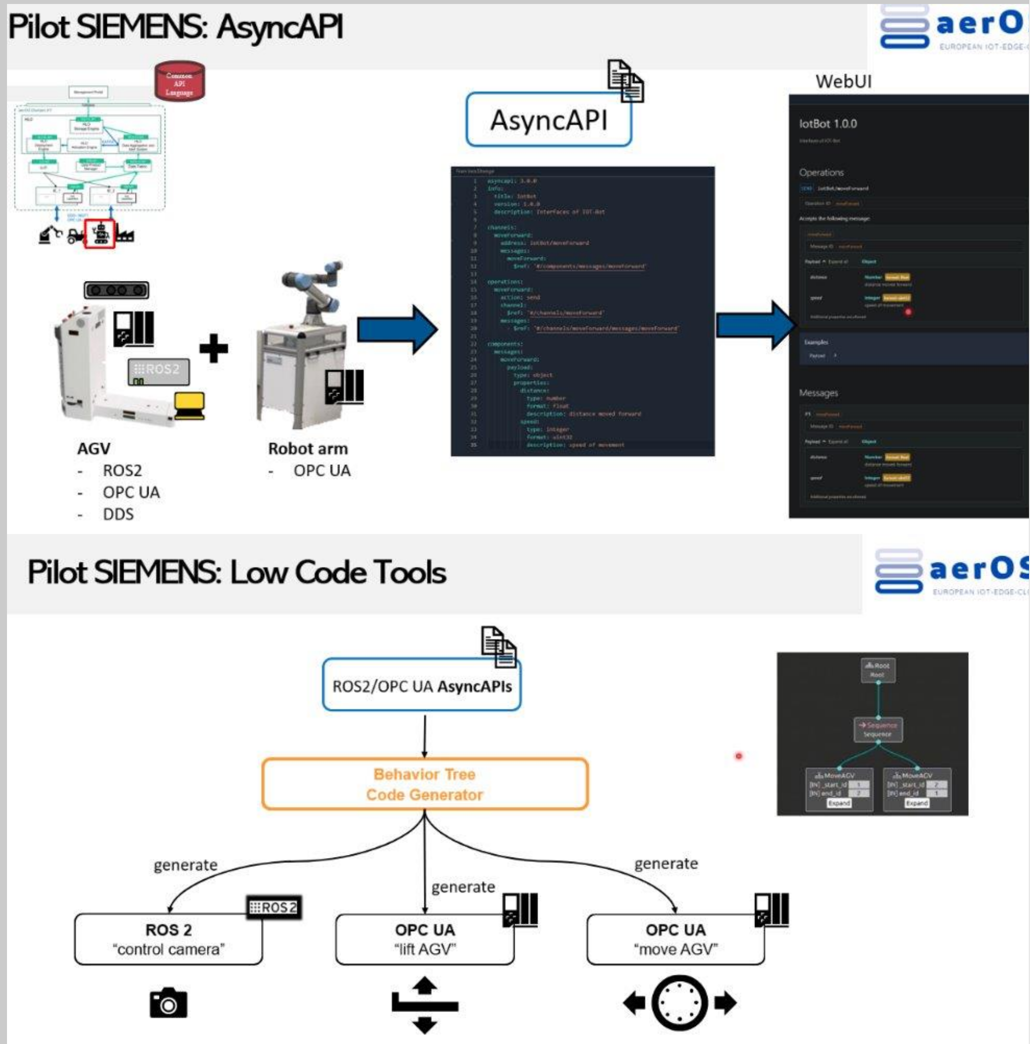

3.8. aerOS management framework

3.8.1. KPI 1.8.1 # of federated domains in all aerOS continuums

Table 51 KPI 1.8.1 # of federated domains in all aerOS continuums

KPI ID number .	KPI 1.8.1
KPI Name	**# of federated domains in all aerOS continuums**
Description	This KPI will quantify the total number of aerOS domains that have been created (and have been functionally deployed) in the project. It refers to all the domains that will have been federated. It gathers both the domains coming from specific pilots and those created for development or integration purposes, as long as they are federated with other domains.
Motivation	The motivation of this KPI is to represent the soundness of the design of "domain" concept. The goal is to be able to demonstrate that the theoretical design of IEs inside domains (i.e., designed and created by system administrators based on different

<table>
<tr><td></td><td colspan="3">criteria, such as topological sense, geographical proximity, container management framework technology, among others) are well translated into real deployments.</td></tr>
<tr><td>Target value</td><td colspan="3">15 (total domains) / (in) 8 (continuums)</td></tr>
<tr><td>Prerequisites</td><td colspan="3">A set of fully functional continuums (one for each pilot/testbench) composed by (at least) a functional domain with all the aerOS Basic Services running and all the IEs properly set up.</td></tr>
<tr><td>aerOS components (task)</td><td colspan="3">aerOS Management Portal (T4.6), aerOS Federator (T4.6), Context Broker (T4.2 and T4.6), aerOS AAA (T3.4), API Gateway (T4.2). Despite using the portal to check the number of federated domains, the action is performed in the Orion-LD instances of the entry point domains of each continuum</td></tr>
<tr><td>Evaluation means</td><td colspan="3">Total functional domains will be counted and indicated as the KPI. This will be endorsed by being able to connect to aerOS portal of the different continuums and checking the number of existing (and functional) domains in each.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>0 (nothing before aerOS)</td><td>6 domains (40%) in 2 continuums (25%).</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Currently, the installation of aerOS is being performed in the pilots following the first installation guide that is available in the aerOS Gitlab, so these continuums are not fully available yet. Therefore, the domains of two different testing environments have only been taken into consideration:
1. The internal testing environment of the UPV, with 3 domains.
Domains list
<table>
<tr><th>Id</th><th>Description</th><th>Public Url</th><th>Owner</th><th>Entrypoint</th><th>Status</th></tr>
<tr><td>1</td><td>UPV Domain 1</td><td>http://158.42.161.177</td><td>UPV</td><td>✓</td><td>Functional</td></tr>
<tr><td>2</td><td>UPV Domain 2</td><td>http://158.42.161.177</td><td>UPV</td><td>×</td><td>Functional</td></tr>
<tr><td>3</td><td>UPV Domain 3</td><td>http://158.42.33.247</td><td>UPV</td><td>×</td><td>Functional</td></tr>
</table>
2. The MVP continuum, which was used in the mid-term review, composed of 3 domains.
Domains list
<table>
<tr><th>Id</th><th>Description</th><th>Public Url</th><th>Owner</th><th>Entrypoint</th><th>Status</th></tr>
<tr><td>CloudFerro</td><td>CloudFerro Domain</td><td>https://cf-mvp-domain.aeros-project.eu</td><td>CloudFerro</td><td>✓</td><td>Functional</td></tr>
<tr><td>Mobile</td><td>Mobile domain for the demo</td><td>https://mobile-domain.aeros-project.eu</td><td>UPV</td><td>×</td><td>Functional</td></tr>
<tr><td>NCSRD</td><td>NCSRD aerOS MVP Domain</td><td>https://ncsrd-mvp-domain.aeros-project.eu</td><td>NCSRD</td><td>×</td><td>Functional</td></tr>
</table>
The followed procedure consists of accessing to the domains section of the Management Portal of each continuum to count the number of federated and available (with Functional status) domains</td></tr>
</table>

3.8.2. KPI 1.8.2 # of continuum functionalities available and operational through the Management Portal

Table 52 KPI 1.8.2 # of continuum functionalities available and operational through the Management Portal

KPI ID number .	**KPI 1.8.2**

<table>
<tr><th>KPI Name</th><th colspan="3"># of continuum functionalities available and operational through the Management Portal</th></tr>
<tr><td>Description</td><td colspan="3">This KPI will quantify the total number of available functionalities to operate the resources and services operating the continuum that can be managed via the aerOS Management Portal. The various functionalities will be associated to resources (e.g., creation of IEs, removal of IEs from the continuum, domains enabling…), users management (e.g., creation, roles assignation, …), services (e.g., deployment, monitoring…) and/or data (i.e., inspecting the Data Fabric…).</td></tr>
<tr><td>Motivation</td><td colspan="3">All continuum management functionalities will be available via APIs. This is managed by T3.2 and will serve as the basis for aerOS continuum establishment. However, some of them will also be available via an UI (in the aerOS portal). The motivation of this KPI is to represent how many of them will be usable via the UI, enlarging the human-oriented capacity of the Meta-OS (in this case, for system administrators' configurations, etc).</td></tr>
<tr><td>Target value</td><td colspan="3">10</td></tr>
<tr><td>Prerequisites</td><td colspan="3">The aerOS entry point domain with a Functional status. It means that the continuum is composed by (at least) a functional domain with all the aerOS Basic Services running and all the IEs properly set up.</td></tr>
<tr><td>aerOS components (task)</td><td colspan="3">The KPI targets the aerOS Management Portal (T4.6), but other components are needed: Orion-LD (T4.2 and T4.6), aerOS AAA (T3.4), HLO (T3.3), Self-* modules (T3.5).</td></tr>
<tr><td>Evaluation means</td><td colspan="3">At the end of the project, functionalities included in the portal will be endorsed via showcasing the possibilities of usage (screenshots, videos, demos).</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>0</td><td>6 (60%)</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Direct visual and functional testing by potential end users (aerOS technical developers and pilot users) after accessing to the portal. This includes the demo performed in the mid-term review, in which it was deployed a Management Portal that is still available for the aerOS partners.

Thus, a potential early accounting of this KPI would be:
Access with various user profiles validated by AAA.
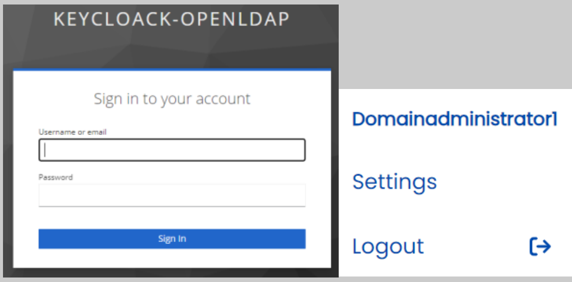

Observing all IEs and domains in the form of a table.</td></tr>
</table>

Domains list

Id	Description	Public Url	Owner	Entrypoint	Status
CloudFerro	CloudFerro Domain	https://cf-mvp-domain.aeros-project.eu	CloudFerro	✓	Functional
Mobile	Mobile domain for the demo	https://mobile-domain.aeros-project.eu	UPV	×	Functional
NCSRD	NCSRD aerOS MVP Domain	https://ncsrd-mvp-domain.aeros-project.eu	NCSRD	×	Functional

Domain detail

Id	Description	Public Url	Owner	Entrypoint	Status
NCSRD	NCSRD aerOS MVP Domain	https://ncsrd-mvp-domain.aeros-project.eu	NCSRD	×	Functional

Infrastructure elements:

ID	LLO	CPU arch	CPU cores	RAM capacity (MB)	Real time	Status	Metrics
ncsrd-m	NCSRD:01	x64	2	11960	×	Ready	
ncsrd-w1	NCSRD:01	x64	2	7940	×	Ready	
ncsrd-w2	NCSRD:01	x64	2	7940	×	Ready	
pi	NCSRD:01	arm64	4	3791	×	Ready	
ncsrd-rt	NCSRD:01	x64	2	7958	×	Ready	

- Visualizing the topology of the continuum in the form of a graph.

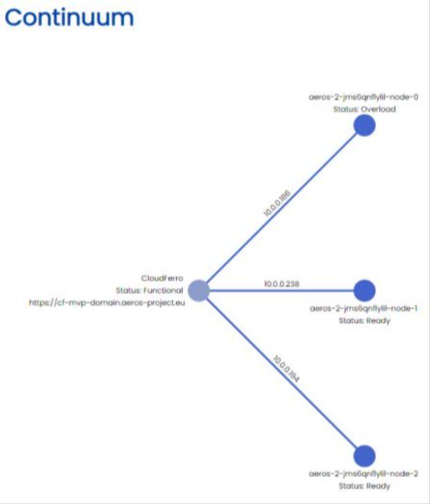

- Observing all deployed services and their underlying components in the form of a table

Deployed services New service deployment

ID	Name	Description	Actions
5bb4e56e	aeros_service_urn:ngsi-ld:Service:5bb4e56e	Gateway of the car	
7ac9dc30	aeros_service_urn:ngsi-ld:Service:7ac9dc30	car controller	

Service detail

ID	Name	Description
7ac9dc30	aeros_service_urn:ngsi-ld:Service:7ac9dc30	car controller

Service components:

ID	IE	Status	Container image	CLI args	ENV variables
39e4e9c5	None	Finished	registry.gitlab.aeros-project.eu/aeros-public/midterm-demo/5g-vehicle/5g-vehicle-controller:latest	×	

- Describing a service to be orchestrated using a guided form (first version with CLI and ENV_ARGs).

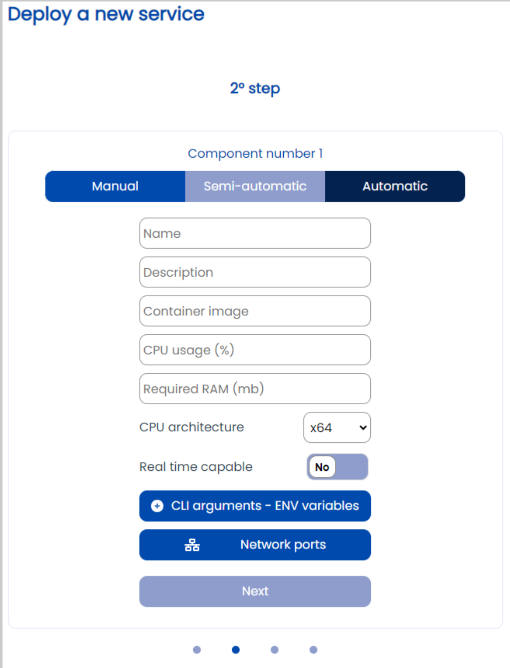

- Commissioning the orchestration of a service (connection with HLO).

Therefore, 6 relevant functionalities are already available through the platform

3.8.3. KPI 1.8.3 Performance of aerOS Federation Context Broker

Table 53 KPI 1.8.3 Performance of aerOS Federation Context Broker

KPI ID number .	**KPI 1.8.3**
KPI Name	**Performance of aerOS Federation Context Broker**
Description	This KPI represents the capacity of the aerOS Federator to withstand high querying and update loads. This will be a direct result of the capacity by the core element of such Federator, the ORION-LD context broker (release for aerOS). The metrics that will be assessed in this KPI will be: (a) number of simultaneous queries to get entities from the Broker–in the same second, (b) number of simultaneous updates of entities – in the same second.
Motivation	The performance of Orion-LD is crucial for earOS because it will allow to understand how many data can be processed, as well as the capacity of the federated network of brokers to support automated distribution of the state of the continuum. Orion-LD, by default, and for historical reasons, uses an old, deprecated driver for MongoDB (note that the associated database –MongoDB- needs more resources, but the restrictive element is the Context Broker). However, the entire database layer has been reimplemented, using the newest MongoDB driver (if requested via a CLI option). The old implementation is C++ while the new is pure C and thus, a boost in performance is expected. The differences in "queries per second" and "entity updates per second" between the default (old, deprecated MongoDB driver) and the new implementation (new MongoDB driver) are the basis for this KPI.

<table>
<tr><th>Target value</th><td colspan="3">5000 queries/s, 2500 updates/s</td></tr>
<tr><th>Prerequisites</th><td colspan="3">Orion-LD, as the Context Broker which is part of aerOS core services, must be functional in (at least) a cloud domain.</td></tr>
<tr><th>aerOS components (task)</th><td colspan="3">aerOS portal and the federation thanks to the Context Broker, allowing the DSNB (T4.6).</td></tr>
<tr><th>Evaluation means</th><td colspan="3">There is planned to realise a performance-measurement exercise in a development/integration environment This measurement exercise will be done at the end of the project (M36).

The procedure of measuring performance for this KPI is described here.</td></tr>
<tr><th>Measurement period</th><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><th>Measured value (% achieved)</th><td>N/A</td><td>3500 requests/s in the small scenario (70%)</td><td>N/A</td></tr>
<tr><th>Outcome elaboration (M24)</th><td colspan="3">For M24, the performance has been measured without the use of distributed operations, it is using purely local entities and no registrations involved. Evidences of the measurements that have been uses as basis can be found here.

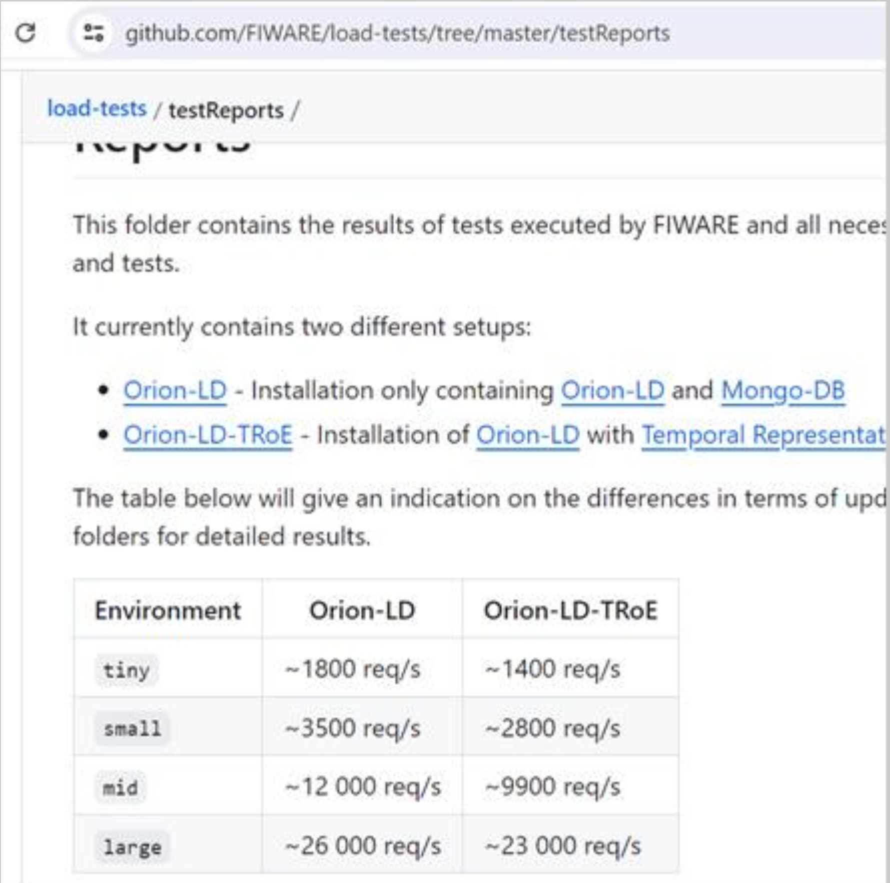
github.com/FIWARE/load-tests/tree/master/testReports
load-tests / testReports /
This folder contains the results of tests executed by FIWARE and all neces
and tests.
It currently contains two different setups:
• Orion-LD - Installation only containing Orion-LD and Mongo-DB
• Orion-LD-TRoE - Installation of Orion-LD with Temporal Representat
The table below will give an indication on the differences in terms of upd
folders for detailed results.
<table><tr><th>Environment</th><th>Orion-LD</th><th>Orion-LD-TRoE</th></tr><tr><td>tiny</td><td>~1800 req/s</td><td>~1400 req/s</td></tr><tr><td>small</td><td>~3500 req/s</td><td>~2800 req/s</td></tr><tr><td>mid</td><td>~12 000 req/s</td><td>~9900 req/s</td></tr><tr><td>large</td><td>~26 000 req/s</td><td>~23 000 req/s</td></tr></table>
As a footnote: aerOS is a highly distributed system and local-only performance is not too interesting. The performance of Orion-LD in aerOS is highly dependent on the network distance and quality between federated brokers.</td></tr>
</table>

3.8.4. KPI 1.8.4 Federation asymptote with minimum latency (domains)

Table 54 KPI 1.8.4 Federation asymptote with minimum latency (domains)

KPI ID number .	KPI 1.8.4
KPI Name	**Federation asymptote with minimum latency (domains)**
Description	The goal is to understand how many new domains can be added to a continuum (installation and addition procedure explained in deliverable D5.2) without trespassing a latency threshold. The latency will be associated to the retrieval of aerOS continuum monitoring information (e.g. the Infrastructure Elements or orchestrated Service Components) from all the available domains in the continuum. Here, the relevant aspect is to understand that aerOS federation is based on NGSI-LD's Context Source Registrations, that will automatically connect the ORION-LD of each domain with all the others (full connection network topology) for keeping updated the distributed state repository of the continuum. That connection is materialised in the exchange of distributed messages. The moment in which all the "update messages" will be completed since a new domain is added will be considered as the latency of the process.
Motivation	This KPI will represent the capacity of aerOS structure of federated domains in a single continuum to scale while maintaining a maximum decided latency.
Target value	≥ 4 domains
Prerequisites	The prerequisites will depend on the kind of tests performed to measure the value of the KPI as the evaluation will be composed of different scenarios. Thus, the prerequisites can be: (i) the installation in a local machine of the functional tests tool suite provided by the developers of Orion-LD; or (ii) a set of aerOS domains with, at least, Orion-LD and aerOS AAA components installed.
aerOS components (task)	Orion-LD (T4.2 and T4.6), aerOS AAA (T3.4) and Self-* modules (T3.5).
Evaluation means	Even though this measurement might be influenced by many factors not strictly related to aerOS traits (network delay, network load, processing time of messages in network nodes, etc.), data will be presented to gauge a relevant figure for this KPI. The measurement methodology planned in M24 (subject to future adjustments) has been: • A maximum "latency of the process" defined (when more federations are taking place in aerOS). • In a development/Integration continuum (not in pilots), tests are done. New domains will be iteratively added, till the "latency of the process" is surpassed. This will be the maximum capacity. The previous process will be repeated for a real case in, at least, one of the pilots of the project in M36. This will be the maximum achieved in TRL4. Three different scenarios have been envisaged to elaborate the outcome of this KPI because as explained before, the measurement of this KPI might be influenced by many external factors to aerOS (e.g. network latency due to the physical distance of domains and message exchange among heterogeneous networks).

	1. Creation of several Orion-LD instances in the same testing machine, leveraging the functional testing tool suite provided by Orion-LD: this scenario aims to provide the theoretical maximum value for the federation asymptote since the vast majority of external factors are avoided, but it presents a strong dependence with the testing machine hardware. 2. Laboratory or controlled environment use-case: the UPV internal continuum will be used because it contains IEs connected to the same UPV internal network, but through different technologies (Ethernet, Wi-Fi, virtualized networks, …) and distributed in diverse locations of the university. Furthermore, scenarios with and without the use of AAA modules can also be explored as this network presents a restricted exposure to the public internet. 3. Real distributed scenario: the aerOS MVP continuum is composed of 3 domains which are geographically distributed among Europe and also present a wide range of underlying IEs and network types.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	0	***Scenario 1: Unlimited, tested in 10 domains.*** ***Scenario 2: >3 domains*** ***Scenario 3: 3 domains***	*N/A*
Outcome elaboration (M24)	In all the scenarios the followed procedure is similar: the starting point is defined as a continuum composed of two domains, which actually means two Orion-LD instances with the proper CSR created. Then, some data that involves all the context brokers is gathered via a HTTP GET request and the response time is measured. After this, a new domain is added to the continuum (actually a new Orion-LD to the brokers registrations pool) and the procedure is repeated until the measurement value reaches or surpases the target maximum response time, which has been set to 5 seconds for M24. At this point, the number of Orion-LD instances is counted and it's actually the KPI value. For D5.5 (M24), only the first scenario has been extensively tested, providing the following results: in a medium-powerful VM (4 vCPU and 16GB RAM), the functional tests are run quickly with a response time less than 50 ms for the GET requests (the entire test is run in less than 1 second) in a scenario consisting of 10 Orion-LD instances. This means that (i) the results of this scenario are highly dependent on the hardware on which it is executed, (ii) it's an ideal scenario without external factors such as networking issues, and (iii) it clearly presents a strong relationship with KPI 1.8.3, which aims to evaluate the performance of the context broker. For the second scenario, the current response time for distributed requests is approximately less than 200ms for 3 domains in the UPV testbench. Therefore, this measured value can be potentially increased in further tests as it's far from the set target value. 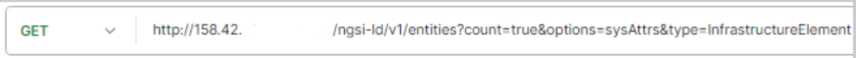  Finally, the measure value for the last scenario has been gathered from the results of the demo conducted in the aerOS mid-term review using the MVP continuum. The results of this demo confirmed the strong dependency of this KPI on the network conditions due to the fact that, in an initial scenario with 2 "static" domains (a pure		

cloud domain located in Poland and a domain with VMs and a mobile car, but with a strong 5G network connection, located in Greece), the response time for the requests was around 600-700ms.

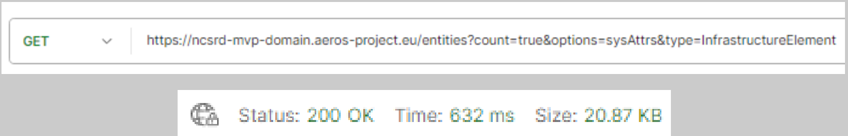

Nevertheless, when a new domain connected to a 4G router inside a building was added to the continuum, the response time quickly increased to roughly 2 seconds, with some peaks of around 3 seconds.

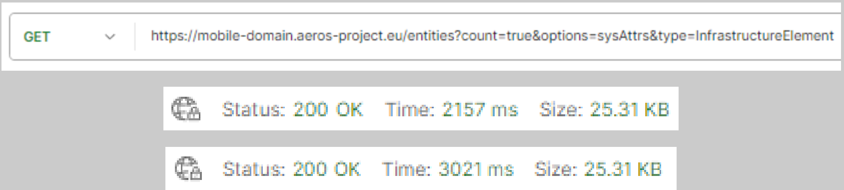

Finally, these data have been obtained so far, illustrating that with 3 domains the federation is completely functional even with less reliable network connections across different countries. Reflections on further number in other scenarios will come in D5.6 with the results obtained in M36.

3.8.5. KPI 1.8.5 Average offloading ratio of entrypoint balancing in aerOS scenarios

Table 55 KPI 1.8.5 Average offloading ratio of entrypoint balancing in aerOS scenarios

KPI ID number .	**KPI 1.8.5**
KPI Name	**Average offloading ratio of entrypoint balancing in aerOS scenarios**
Description	This KPI represents the effectiveness of the balancing algorithms selected and deployed in the Management Portal for achieving entry point diversity usage. The goal is to demonstrate that, on average, 30% of requests sent to EB are distributed for first processing to HLOs located in domains other than the one containing EB (entry point domain).
Motivation	Measure the effectiveness of EB in the distribution of requests between HLOs, which aims to minimize single HLO overutilization
Target value	30%
Prerequisites	Management Portal must be deployed on the aerOS and must correctly properly pass requests to EB; The information about HLOs and their domains must be present in Orion-LD Context Broker; The FE of HLOs should be accessible under addresses registered in Orion-LD Context Broker.
aerOS components (task)	API Gateway (T3.4), Orion-LD (T4.2), aerOS Management Portal (T4.6)
Evaluation means	The EB algorithm will be tested both in development/integration scenarios and in pilots. The tests will be performed separately on the continuum with 2 domains

	(preliminary tests) and 3 domains (advanced tests). For each test, 15 workloads will be used. In total, it is envisioned that 4 types of tests are going to be performed, aiming to encompass and represent various scenarios, in which the EB can be used: Test Case 1: There are no running services present in the continuum. The client requests are sent directly to the EB, which uses the weighting function based on CPU usage. The scenario aims to evaluate, whether the EB will be able to distribute the requests to different domains even when the score (i.e., number of running services divided by the weight) of each domain remains the same. Test Case 2: There are running services deployed on each of the domains: • For the continuum with 2 domains: 2 services running on *Domain1*, 1 service running on the *Domain2*. • For the continuum with 3 domains: 2 services running on *Domain1*, 1 service running for *Domain2*, 1 service running for *Domain3*. The client requests are directly sent to the EB, which uses the weighting function based on CPU usage. The scenario aims to evaluate, whether the EB will be able to distribute the requests when the continuum domains are in different states. Test Case 3: The state of the continuum is the same as in Test Case 2. However, this time, a different weighting function is going to be used. In particular, the RAM usage is to be considered along with CPU usage. The scenario aims to compare the effectiveness of different weighting functions, so that the one for which the better outcomes are achieved (i.e., offloading ratio) can be selected. Test Case 4: The state of the continuum is the same as in Test Case 2. However, in this scenario, part of the requests will have a target domain indicated by the client. For these requests, the execution of the load balancing algorithm should be omitted by EB. The scenario aims to evaluate the decision-making process of EB, as well as its ability of capturing dynamic changes of the continuum. For the states of the continuum described in Test Case 1 and Test Case 2, the tests are also going to be performed without the usage of EB, so that both of the obtained outcomes can be compared. To measure the offloading ratio, a strategy will be envisaged to quantify/catalogue whether a service deployment request originated in a certain domain and queried the HLO of a different domain as first processing option.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0% (when all clients select the target HLO explicitly)*	***Test Case 1 (preliminary): 33.3% offloading ratio (25%)***	*N/A*
Outcome elaboration (M24)	For the M24, the tests have been focused on the overall correctness of the EB algorithm, as well as EB integration with the Management Portal and HLO FEs. Therefore, before proceeding to the scenario tests of Test Cases 1-4, the automated unit and integration tests were conducted. For each of the integration tests, the offloading ratio was computed. These tests considered simplified versions of the scenarios described in Test Case 1 and 2, in particular:		

- load balancing performed for 4 upcoming client requests when 2 domains with no running services are present in the continuum (simplified Test Case 1) – **achieved offloading ratio: 50%**
- load balancing performed for 4 upcoming client requests when 2 domains with different number of running services are present in the continuum (simplified Test Case 2) – **achieved offloading ratio: 25%**

As presented on the following figure, 93% of EB functionalities have been covered by the performed tests.

Element	Class, %	Method, %	Line, %
org.aeros	100% (13/...	93% (46/49)	96% (120/124)
config	100% (1/1)	100% (2/2)	100% (8/8)
exception	100% (3/3)	100% (7/7)	100% (14/14)
feign	100% (0/0)	100% (0/0)	100% (0/0)
mapper	100% (1/1)	100% (1/1)	100% (6/6)
rest	100% (1/1)	100% (3/3)	100% (5/5)
service	100% (6/6)	94% (32/34)	96% (86/89)
EntryPointBalancer	100% (1/1)	50% (1/2)	50% (1/2)

In terms of scenario testing, for D5.5 (M24), only preliminary Test Case 1 was considered. The EB was deployed in the continuum with 2 domains. Then, the consecutive client requests were sent to the EB using its exposed REST API endpoint. Among 15 requests sent, 5 have been redirected to the domain other than the entrypoint domain. Hence **33.3% offloading ratio** was achieved. This is a promising result with respect to the established target and serves as a basis for further analysis in more complex Test Cases for M36.

3.8.6. KPI 1.8.6 QoE of Management Portal deployed on pilots

Table 56 KPI 1.8.6 QoE of Management Portal deployed on pilots

KPI ID number .	**KPI 1.8.6**
KPI Name	**QoE of Management Portal deployed on pilots**
Description	The idea of this KPI is to evaluate the QoE of the UI of the Management Portal. In particular, it is the intention fo quantify a metric for tracking both behavioural and attitudinal perception of the webapp.
Motivation	The rationale behind this KPI is to be able to assess the quality of experience of stakeholders using the aerOS Management Portal. Note that stakeholders (users) of the implemented UI will be both system administrators (for e.g., configuring domains) and 4owners (for e.g., monitoring data or KPIs).
Target value	>=68 SUS score
Prerequisites	aerOS runtime working, web service ready, HLO ready and Management Portal deployed in pilots

<table>
<tr><td>aerOS components (task)</td><td colspan="3">HLO (T3.3), Management Portal (T4.6)</td></tr>
<tr><td>Evaluation means</td><td colspan="3">For this first piloting phase, it was decided, as also anticipated in D5.2, to create a QoE survey based on the System Usability Scale (SUS), a widely recognized and standardised tool for assessing the usability of a system. The SUS provides a reliable measure of usability with a small number of questions, making it efficient for both respondents and analysts. It consists of 10 statements with alternating positive and negative phrasing to reduce response bias. Each item is scored on a 5-point Likert scale, ranging from Strongly Disagree (1) to Strongly Agree (5). This final score represents the usability of the system. By gathering feedback on key usability aspects, the survey aims to identify strengths and areas for improvement in the portal's design and functionality. In detail, At M29 it is planned the first QoE survey that effectively evaluates user experience in digital services and portals, particularly within the context of role-specific functionalities and interactions. This QoE survey will be designed based on several established QoE standards and frameworks to ensure comprehensive and relevant questions that cover key aspects of user experience:
ITU-T P.800 Series, methods for subjective determination of transmission quality, adaptable for interactive digital experiences.
ISO 9241-11:2018, defines usability based on effectiveness, efficiency, and satisfaction. Questions on usability, ease of navigation, and intuitiveness are based on this standard.
ISO/IEC 25010:2011, provides a quality model including characteristics such as functional suitability, performance efficiency, usability, reliability, security, compatibility, maintainability, and portability. The survey addresses aspects like performance, functionality, and overall satisfaction.
Nielsen's Usability Heuristics principles for user interface design, including visibility of system status, consistency, error prevention, and ease of use.
Finally, an updated and final QoE survey will be conducted at M35.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>N/A</td><td>N/A</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Although it has not been carried out yet, the SUS survey to be used for M29 is already defined as follows:

Your feedback is essential to help us improve the aerOS Management Portal. Please answer the following questions based on your experience with the portal. For each statement, indicate your level of agreement using the following scale (Strongly Disagree – 1, Disagree - 2, Neutral - 3, Agree - 4, Strongly Agree - 5):
I found it easy to navigate to the Home page and other sections like Domains, Deployments, and Continuum.
I found the layout and design of the portal confusing.
I felt confident using the portal to perform actions according to my assigned role (e.g., initiating orchestration requests, checking domain status).
I needed the support of a technical person to be able to use this portal.
The interactive continuum network graph effectively represented the computing resources in my domains.</td></tr>
</table>

	6. *I found inconsistencies in the way different sections of the portal operated.* 7. *I was able to quickly learn how to request a service orchestration in the Deployments section.* 8. *I found the portal cumbersome when trying to perform specific actions related to my role.* 9. *The portal provided all the necessary information and tools for managing my tasks effectively.* 10. *I had to spend a lot of time learning how to use different features of the portal.*

3.9. aerOS embedded analytics

3.9.1. KPI 1.9.1 # pre-packaged functions supported by Embedded Analytics Tool (EAT)

Table 57 KPI 1.9.1 # pre-packaged functions supported by Embedded Analytics Tool (EAT)

KPI ID number .	**KPI 1.9.1**		
KPI Name	**# pre-packaged functions supported by Embedded Analytics Tool (EAT)**		
Description	The Embedded Analytics Tool is a platform for the design, development and deployment of analytical functions. Several generalised functions are packaged with the Embedded Analytics Tool to address common operations to provide insights for management and AI/ML components.		
Motivation	The pre-packaged functions of the Embedded Analytics Tool provide basic operations for the aerOS system. These functions are leveraged by other components to provide insights such as data samples or highlight anomalies and data drifts. These functions are also generalised and can be customised through user parameters, allowing them to have “plug and play” characteristics in a range of different environments and scenarios.		
Target value	3		
Prerequisites	The Embedded Analytics Tool must be installed according to the instructions provided in the project repository. These instructions address security and privacy features through access tokens for downloading EAT components and credentials for dashboard login. EAT is considered successfully installed when all EAT components report “Running”.		
aerOS components (task)	Embedded Analytics Tool (T4.4)		
Evaluation means	The development of each function follows an incremental development approach with a unique predefined test data set for each function. Functions are evaluated based on expected versus actual results of function execution.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	0	***3 (100%)***	*N/A*

Outcome elaboration (M24)	3 pre-packaged functions have been delivered to the project repository and presented at the 5th Plenary Meeting Day 2 (Warsaw). They are available here: 1. Anomaly detection: https://gitlab.aeros-project.eu/wp4/t4.4/embedded-analytics-tool/-/tree/main/functions/anomaly-detection 2. Data drift: https://gitlab.aeros-project.eu/wp4/t4.4/embedded-analytics-tool/-/tree/main/functions/data-drift 3. Stratified sampling: https://gitlab.aeros-project.eu/wp4/t4.4/embedded-analytics-tool/-/tree/main/functions/stratified-sampling The testing of these functions on an internal testbed marks the first stage of validation. Next steps until M36 are to validate functions in an orchestrated aerOS environment.

3.9.2. KPI 1.9.2 # northbound wrappers designed for common operations with EAT

Table 58 KPI 1.9.2 # northbound wrappers designed for common operations with EAT

KPI ID number .	**KPI 19.2**		
KPI Name	**# northbound wrappers designed for common operations with EAT**		
Description	Wrappers will be designed and implemented for the creation of Embedded Analytics Tool functions. These wrappers will be available to all function authors who utilise the aerOS templates. These templates will be stored in the Embedded Analytics Tool repository and be available to all project partners.		
Motivation	Adopting a Function as a Service approach for the Embedded Analytics Tool enables flexibility for the function authors, but also introduces function design and communication challenges. Templating allows for function design and communication to be structured and guided. Function authors can focus on the core logic of their functions while using approved and tested implementations for common operations such as data retrieval or triggering external actions.		
Target value	3		
Prerequisites	Installation of the EAT functions repository is required. EAT specific applications such as faas-cli allow users to engage with EAT to create, deploy and remove functions. The aerOS template provides a structured model preconfigured with aerOS specific features such as visualization.		
aerOS components (task)	Embedded Analytics Tool (T4.4)		
Evaluation means	The aerOS template provides 3 defined operations within aerOS functions. These are Data Retrieval (e.g., requesting data from Data Fabric), aerOS component communication (e.g., forwarding data to HLO), and data visualization (e.g., exposing in-function metrics to EAT Grafana component). These operations are evaluated based on their successful execution as part of EAT pre-packaged functions, and the creation of use case specific functions.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**

Measured value (% achieved)	0	***3 (100%)***	*N/A*
Outcome elaboration (M24)	The aerOS template can be found on the project repository under aerOS-python3-flask template here: https://gitlab.aeros-project.eu/wp4/t4.4/embedded-analytics-tool/-/tree/main/functions/template The template provides the 3 operations listed above with flexibility to add additional operations in the event of future required changes. This template has provided the foundation of 3 currently deployed functions on the Embedded Analytics Tool.		

3.10. Stakeholder user satisfaction

3.10.1. KPI 1 Successful conduction of Open Calls (KVI-7.1)

Table 59 KPI 1.10.1 Successful conduction of Open Calls (KVI-7.1)

KPI ID number .	**KPI 1.10.1**		
KPI Name	**Successful conduction of Open Calls (KVI-7.1)**		
Description	This KPI will represent the combined number of applications received between round 1 and round 2 of Open Calls funding opportunity. Additionally, this KPI will refer to the successful implementation of Open Call projects (60k€, 9 months each).		
Motivation	To illustrate the capacity of the Consortium to engage dynamic, vibrant researchers in Universities, SMEs and RTOs to utilise aerOS technologies and provided added value to the pilots or the research strategy of the project.		
Target value	> 80 applicants		
Prerequisites	• Pre-requisites for M24: ○ The Process for the first Open Call award has been completed: publication, application window, evaluation, ranking, decision, announcement of winners and commencement of projects. ○ The Process for the 2nd Open Call is initiated: publication and opening of application window.		
aerOS components (task)	First Open Call projects are in execution, and so far, no specific components have been exploited yet.		
Evaluation means	On the one hand, the number of applications received will be published in aerOS website some days after the close of each of the two application windows. On the other hand, the final reports of the 15 OC projects to be funded will be summarised and included in deliverable D1.4. The KPI target will be then represented as follows: >80, and 15 out of 15 successfully completed projects that provide lessons learnt		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**

<table>
<tr><td>Measured value (% achieved)</td><td>0</td><td>38 applicants (47.5%)</td><td>N/A</td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">A total of 38 proposals were submitted for Open Call Application #1 of aerOS. The awarded proposals (now, projects) are:

Acronym	Submitter	Title	Pilot	Challenge	Country	Entity Type
HACER	Bytek	High Accuracy Cost Efficient Differential Positioning System using Real-Time Kinematics to optimise port logistics	P4	P4C1	Spain	SME
DAIMon	UPCT	Distributed AI-based Atmospheric Visibility Index Service for Agricultural Mobile Machinery within the aerOS framework	P3	P3C1	Spain	University
EcoQM	Nissatech	Framework for eco-quality monitoring and control supported by aerOS	P1	P1C1	Serbia	SME
ENERGETIC	Nextworks	Energy management and comfort living for green, healthy and productive offices	P5	P5C2	Italy	SME
ANEOSP	Secmotic Innovation SL	AI Nodes for Enhanced Occupational Safety in Ports	P4	P4C3	Spain	SME
IBRTEFC	The Data Cooks	IoT-Based Real-Time Environmental Footprint Calculator	P1	P1C1	Netherlands	SME
GreenAnalyzer	UCY	A framework for Geo-distRibuted Edge-cloud Energy consumption ANALYsis towards Zero Emission Rates	P2	P2C2	Cyprus	University
</td></tr>
</table>

3.10.2. KPI 1.10.2 # of stakeholders deploying aerOS

Table 60 KPI 1.10.2 # of stakeholders deploying aerOS

KPI ID number .	**KPI 1.10.2**
KPI Name	**# of stakeholders deploying aerOS**
Description	The number of entities public/private deploying aerOS Meta-OS components to support operation and implementation of advanced hyper distributed applications.
Motivation	The number of stakeholders deploying aerOS will generate the necessary evidence to support future adoption of Meta-OS assets
Target value	5
Prerequisites	aerOS ready to be deployed, with all the needs that implies.

aerOS components (task)	At least, the core aerOS services needed to set up an aerOS domain: HLO (T3.3), LLO (T3.3), Self-* (T3.5), Context Broker (T4.2), Data Fabric (T4.2)		
Evaluation means	Counting only the real stakeholders that have already deployed any aerOS domain, regardless of the number of IE involved or its purpose (testing or production environment).		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	0	***2 (40%)***	*N/A*
Outcome elaboration (M24)	**Pilot 1 – 0 Stakeholders** • NASERTIC (Infrastructure provider): 0 At this point aerOS components have not been deployed. • SSF (Scenario1) – 0 : aerOS has not been deployed yet in the SSF. However, APIs and Low-code tools are already operational on the production line, which allow to remotely access real-time information regarding the machine's status, environmental data (“temperature","humidity","vibration","atmospheric_pressure") and energy monitoring data ("power_current","power_cumulated","CO2_footprint_current","CO2_footprint_cumulated"). • INNOVALIA (Scenario2) – 0 aerOS has not been deployed yet any domain, but pre-production environment has been set up. • SIEMENS (Scenario3) – 0: aerOS has not been deployed yet any domain • POLIMI&MADE (Scenario4) 0: MADE is in the process of deploying several components, POLIMI has already deployed Orion-CB and created some of the entities, but still has not deploy the basic aerOS components. **Pilot 2 – 0 Stakeholders** All aerOS Basic Components have been installed in the test environment, but it fails to orchestrate and execute workloads. HLO is not ready for use because currently is not passing credentials (Not using Kraken). **Pilot 3 – 0 Stakeholder**. Depending on the complexity of the integration of aerOS. Ongoing work on the integration of several components, and currently going step by step with the instructions provided by UPV. **Pilot 4 – 1 Stakeholder**. The partners involved PRODEVELOP, EUROGATE, and CUT. PRODEVELOP has deployed an aerOS domain with a single IE. It is not deployed on EUROGATE for the port continuum pilot, but for the development and		

testing purposes. EUROGATE nor CUT have deployed any aerOS domain yet.

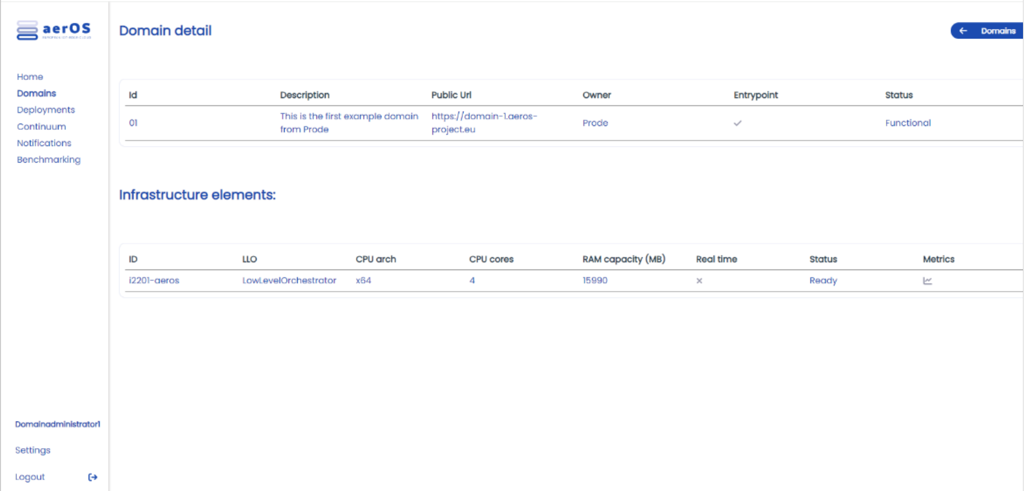

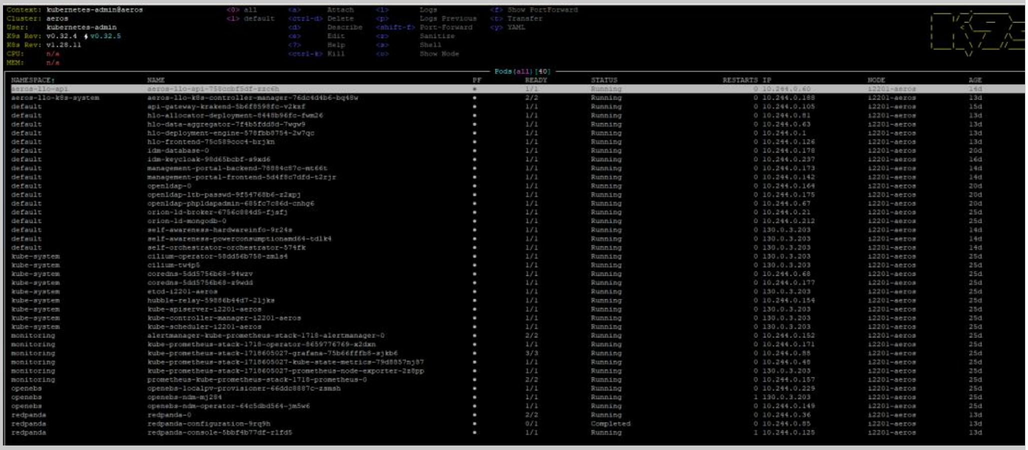

Pilot 5 – 1 Stakeholder: COSMOTE has two aerOS domains (one integrates KubeEdge also). The following aerOS components are already deployed in both domains: Smart Networking Components: CNI, MetalLB (loadBalancer), Ingress controller and Krackend (also security and APIs related); Data Fabric: Orion-LD Context Broker; IE Self-*: Self-orchestrator, Self-awareness (including power consumption); and Orchestration components: LLO and HLO

3.10.3. KPI 1.10.3 # Energy consumption & e-waste reduction in aerOS adopters

Table 61 KPI 1.10.3 # Energy consumption & e-waste reduction in aerOS adopters

KPI ID number .	KPI 1.10.3
KPI Name	# Energy consumption & e-waste reduction in aerOS adopters
Description	The adopters of the aerOS platform will be asked to provide information on energy consumption and e-waste compared to their baseline operation.
Motivation	aerOS will not only support productivity enhancement of European companies to increase economic growth across the EU, but will also support tackling relevant social challenges, including energy consumption and e-waste. As part of the ability to manage such urgent social challenges, it is important to show demonstrated figures towards this direction coming from the stakeholders adopting the platform.

Target value	2% to 10%		
Prerequisites	Integration is complete, and aerOS platform has been embraced by adopters.		
aerOS components (task)	All core aerOS components		
Evaluation means	Collecting data from adopters about energy consumption and e-waste before and after integrating the aerOS platform. To do so, a 4-step methodology per adopter is followed, namely: 1. Baseline Energy and waste consumption analysis based on literature (adopter specific) 2. Validation of the infrastructure that will be used for KPIs measurements 3. KPI measurement campaigns and 4. KPI collection and analysis (comparison with baseline and target values)		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*Use case specific value measured in the baseline scenario with no aerOS services*	***N/A***	*N/A*
Outcome elaboration (M24)	Values are not available yet from the adopters. They will be measured in the baseline scenario with and without aerOS services.		

4. aerOS pilot KPIs

The specific pilot developments throughout the project are evaluated and assessed over the pilot testbeds. This section describes how the evaluation is being managed and implemented during the pilot's trials execution. The results and analysis are part of the work being carried out in *Task T5.4 Continuous use cases analysis, evaluation, and assessment*. The objectives, status of the developments, and list of KPIs, etc. of each pilot are described next.

4.1. Pilot 1 Data-driven cognitive production lines

4.1.1. KPI 2.1.1 Production process accuracy

Table 62 KPI 2.1.1 Production process accuracy

KPI ID number .	KPI 2.1.1		
KPI Name	Production process accuracy		
Description	The accuracy of the process based on digital and virtual part analysis – online and offline.		
Motivation	The quality of the process is based on the adequate selection of the dimensional quality control instrumentation and the optimisation of the quality control strategy and configuration of the manufacturing equipment.		
Target value	10% increase		
Prerequisites	Historical data for each machine and the necessary aerOS components to be ready by M36.		
aerOS components (task)	N/A now. To define for M36.		
Evaluation means	From historical data and moment availability, several machines will be eligible to be selected based on each measurement specification. With aerOS and thanks to the automatization of the process, the optimal machine for the specific measurement is selected, increasing production process accuracy. The evaluation compares values with and without aerOS deployed. This is done by executing the measurement process at least 5 times for 2 different gages (Duplex and Spark), and then comparing these results with the ones obtained with aerOS selecting the optimal gage.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*Dependent on product GD&T complexity*	***N/A***	*N/A*
Outcome elaboration (M24)	Even though progress has been made, aerOS has not been integrated yet into the M3 workspace, and since progress tracking for this KPI follows a binary approach, there is still no progress to report.		

4.1.2. KPI 2.1.2 Digital service programming time

Table 63 KPI 2.1.2 Digital service programming time

KPI ID number .	**KPI 2.1.2**		
KPI Name	**Digital service programming time**		
Description	The time it takes to program quality control services and routines based on traditional monolithic (client/server) vs computing continuum platform.		
Motivation	Autonomous operation of zero-defect manufacturing services and dimensional quality control demand that automation services are highly available to ensure synchronised and safe.		
Target value	2 days		
Prerequisites	aerOS necessary components to be ready for M36		
aerOS components (task)	N/A now. To define for M36		
Evaluation means	Focused on the configuration speed of the testing environment. The metrology process consists of several sequential processes (Qualification, Registration, Alignment, Measurement and Reporting), and thanks to aerOS is possible to shorten the whole process by automating and/or making remote operation of some of the individual processes. To quantify the improvement, the total time of the process is compared with the baseline value.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*2 weeks*	***< 2weeks (10 business days)***	*N/A*
Outcome elaboration (M24)	On the way to achieve the goal of 2 days with aerOS, for M24 several changes have been made. The initial programming time of the routines has been slightly shortened (especially noticeable for short batches).		

4.1.3. KPI 2.1.3 Dimensional quality control productivity

Table 64 KPI 2.1.3 Dimensional quality control productivity

KPI ID number .	**KPI 2.1.3**
KPI Name	**Dimensional quality control productivity**
Description	The time to perform a quality control cycle (specification, programming and execution)
Motivation	Zero-defect manufacturing services and dimensional quality control are usually manually driven processes locked to users and machines being physically interacting.

	The productivity decoupling task programming, dispatching and execution of metrology routines can increase factory productivity.		
Target value	5 parts/hour		
Prerequisites	aerOS necessary components to be ready for M36 Special proprietary machine for parts faster pre-alignment.		
aerOS components (task)	N/A now. To define for M36.		
Evaluation means	Highly related with the previous KPI 2.1.2 and the goal to accelerate the overall process. In particular, with the loading and unloading of parts during the production metrology process (repetitive consequent measurements of the same piece). To quantify the improvement, the number of parts measured by hour and within the aerOS project, is compared against the baseline value outside aerOS.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*3 parts/hour (depending on GD&T complexity)*	***2 parts/hour (40%)***	*N/A*
Outcome elaboration (M24)	Within Production Metrology, a project containing a set of 64 parts has been considered. The part geometry is not simple, and for this reason the M24 value is under the predefined baseline. The total time of performing the quality control cycle for the first piece is 85 mins, followed by 30 mins per part for the next 63, plus another 5 mins for the final report. Total of 1980 mins (33hrs).		

4.1.4. KPI 2.1.4 Accuracy of the CO2-footprint prediction (%)

Table 65 KPI 2.1.4 Accuracy of the CO2-footprint prediction (%)

KPI ID number .	**KPI 2.1.4**
KPI Name	**Accuracy of CO2-footprint prediction (%)**
Description	This KPI represents the fidelity of the value obtained for the CO2 footprint
Motivation	Based on the methods used for CO2 footprint calculation and the data captured from the shopfloor the accuracy of the PCF value is higher and hence the impact and costs associated with product-related emissions lowered.
Target value	>80%
Prerequisites	To calculate the predicted and actual CO2 footprint for each product, the IEs setup must be completed, Node-Red configured for CO2 emissions calculation, the data collection tool set up and the network and APIs configured. To measure the overall accuracy of the CO2 footprint for the entire production, aerOS runtime must be running and web service must be ready.

aerOS components (task)	Networking (T3.1), API & Low Code tools (T3.2), Orchestration (T3.3), T3.5 (Self-*), T4.2 (Data Fabric), aerOS Portal (T4.6)		
Evaluation means	The actual CO2 footprint of each product is based on the production data sent by the various IEs involved in the scenario, which are then processed using the Node-Red tool to obtain the final PCF value. This value is collected for each new order, with at least one test carried out every two to three weeks (1 to 2 tests per month). It is then compared with the predicted CO2 emissions calculated before production of each product, to measure the evolution of the PCF accuracy in grams. The overall accuracy of the prediction of the CO2 footprint for the entire production is estimated after aggregating the accuracy for each product and displayed in the SSF using visualization tools.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	• ***CO2-footprint prediction for F330 Model and High Edition: 450 grams.*** • ***Actual CO2 Emissions for F330 Model and High Edition: 459 grams.*** ***Accuracy of the predicted value is within 50% tolerance.***	*N/A*
Outcome elaboration (M24)	The predicted CO2 value is calculated before each drone is produced. For the value measured in M24, predicted and actual data were collected for a specific drone, the F330 model (the smallest drone produced), and in High Edition version (the largest producer of CO2 emissions). Comparison of the two values gives an accuracy within the 50% tolerance.		

4.1.5. KPI 2.1.5 CO2-footprint measurement (% products)

Table 66 KPI 2.1.5 CO2-footprint measurement (% products)

KPI ID number .	**KPI 2.1.5**
KPI Name	**CO2-footprint measurement (% products)**
Description	This KPI assess the number of products that can be assessed for CO2 footprint calculation (unit, batch, family level)
Motivation	The PCF calculation and the associated DPP information for Scope 3 environments demand different level of granularity in terms of data collection and product-level calculation to meet with EU regulation.
Target value	10% - 100%

Prerequisites	To calculate the number of products that can be assessed for the CO2 footprint (at unit, batch or family level), the IEs configuration must be complete, Node-Red configured, the data collection tool set up and the network and APIs configured. To measure the global percentage of products that can be assessed, the aerOS runtime must be running and the web service must be ready.		
aerOS components (task)	Networking (T3.1), API & Low Code tools (T3.2), Orchestration (T3.3), T3.5 (Self-*), T4.2 (Data Fabric), aerOS Portal (T4.6)		
Evaluation means	To measure the number of products assessed for CO2 footprint calculation the count of products at unit, batch, and family levels are tracked. Data is collected monthly, noting how many products are assessed in each category. ERP systems and data visualization tools are used for tracking and reporting. The performance is measured by the total counts and percentage coverage of assessments.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***1 out of 20 drones assessed (5% at unit level)***	*N/A*
Outcome elaboration (M24)	At M24, a test was carried out to establish the number of products assessed for CO2 footprint calculation, at unit level. Out of a total of 20 drones tested, 1 was assessed for CO2 footprint calculation, i.e. 5%.		

4.1.6. KPI 2.1.6 CO2 emissions reduction (kg)

Table 67 KPI 2.1.6 CO2 emissions reduction (kg)

KPI ID number .	**KPI 2.1.6**
KPI Name	**CO2 emissions reduction (kg)**
Description	aerOS system should permit with AI/ML component contribution to optimize travels of AGV and infer CO2 emission reduction
Motivation	The PCF should contribute to industrial competitiveness
Target value	<20%
Prerequisites	Fully integrated system
aerOS components (task)	KrakenD (T3.1), Ingress (T3.1), Ingress&KrakenD conf (T3.1), CertManager&LetsEncrypt (TLS) (T3.1), FDQN (T3.1), NAT Capable (T3.1), OpenAPI (T3.2), AsyncAPI (T3.2), Low-Code (T3.2), HLO (T.3.3), LLO (k8s, Docker) (T3.3), KeyCloack (T3.4), KrakenD (T3.4), OpenLDAP (T3.4), Self-awareness (T3.5), Self-orchestration (T3.5), Self-diagnose (T3.5), Self-security (T3.5), Self-healing (T3.5), Self Configuration (T3.5), Self-API (T3.5), Data-Interoperability (T4.1), LDAP (T4.2), EAT (T4.4), Trust SCO. (T4.5), IOTA (T4.5), Portal (T4.6).

Evaluation means	The methodology used to evaluate this KPI will be as follow: 1. Measure the energy consumption of the AGV per travel in kWh. 2. Find correlation with CO2 emission based on the electricity grid emission factor. This factor represents the average amount of CO2 emitted per unit of electricity produced in the region where the AGV operates, in our case Lombardy, Italy. CO2 emissions (kg) = Electricity consumption (kWh) * Grid emission factor Lombardy(IT) (kg CO2/kWh).		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0%*	***N/A***	*N/A*
Outcome elaboration (M24)	CO2 emission reduction is not yet measurable given that the integration will be finalized after M24. KPI results, based on the plan of activities, will be available at M32.		

4.1.7. KPI 2.1.7 AGV usage

Table 68 KPI 2.1.7 AGV usage

KPI ID number .	**KPI 2.1.7**
KPI Name	**AGV usage**
Description	AGV use above 80%
Motivation	The AGV usage should be optimized to exploit as much as possible its work in manufacturing areas. The AGV usage optimization will impact also AGV availability
Target value	>80%
Prerequisites	The AGVs must be fully operational and integrated into the manufacturing workflow. An efficient orchestrator and load balancing system must be established to distribute tasks evenly among all AGVs. Additionally, staff should be trained to manage AGV operations and use the monitoring systems effectively to ensure balanced AGV usage.
aerOS components (task)	HLO (T3.3), Data Fabric (T4.2), Low Code Tools (T3.2)
Evaluation means	The evaluation of this KPI will involve continuous monitoring of AGV usage data to ensure balanced workload distribution. This includes collecting data on the operational hours and tasks completed by each AGV within a given time frame. The data will be analysed to calculate the percentage of time each AGV is in use compared to its total available time. The goal is to ensure that all AGVs are utilized evenly, preventing any single AGV from being overused. Comparisons will be made against the target value of 80% to ensure optimal usage. Any imbalances in AGV usage will be investigated, and adjustments to the orchestrator system will be made to address them.

Measurement period	Baseline	M24 (Deliverable D5.5)	M36 (Deliverable D5.6)
Measured value (% achieved)	*30%*	***50%***	*N/A*
Outcome elaboration (M24)	Data collected over this period will be analysed to determine the current usage levels of each AGV and ensure balanced workload distribution. The report will highlight the percentage of time each AGV was in use and compare it to the target value of 80%. Any challenges faced in optimizing AGV usage and the strategies implemented to address them will be described.		

4.1.8. KPI 2.1.8 AGV availability

Table 69 KPI 2.1.8 AGV availability

KPI ID number .	KPI 2.1.8		
KPI Name	AGV availability		
Description	AGV availability above 95%		
Motivation	The AGV availability should be increased to make the manufacturing process leaner and more responsive.		
Target value	>95%		
Prerequisites	There should be a enough number of AGVs in the fleet that have no issues and connected to the aerOS continuum. Fully integrated system		
aerOS components (task)	KrakenD (T3.1), Ingress (T3.1), Ingress&KrakenD conf (T3.1), CertManager&LetsEncrypt (TLS) (T3.1), FDQN (T3.1), NAT Capable (T3.1), OpenAPI (T3.2), AsyncAPI (T3.2), Low-Code (T3.2), HLO (T.3.3), LLO (k8s, Docker) (T3.3), KeyCloack (T3.4), KrakenD (T3.4), OpenLDAP (T3.4), Self-awareness (T3.5), Self-orchestration (T3.5), Self-diagnose (T3.5), Self-security (T3.5), Self-healing (T3.5), Self-Configuration (T3.5), Self-API (T3.5), Data-Interoperability (T4.1), LDAP (T4.2), EAT (T4.4), Trust SCO. (T4.5), IOTA (T4.5), Portal (T4.6).		
Evaluation means	The evaluation of this KPI will involve continuous monitoring of AGV operational status. This includes collecting data on the total available time and the actual operational time of each AGV. The data will be analysed to calculate the availability percentage, ensuring it meets or exceeds the target value of 95%.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0%*	***N/A***	*N/A*
Outcome elaboration (M24)	AGV availability results is not yet measurable given that the integration will be finalized after M24. KPI results, based on the plan of activities, will be available at M32.		

4.1.9. KPI 2.1.9 AGV travel saved/valve

Table 70 KPI 2.1.9 AGV travel saved/valve

KPI ID number .	KPI 2.1.9		
KPI Name	AGV travel saved/valve		
Description	aerOS system should permit with AI/ML component contribution to optimize travels of AGV and improve the ratio travel/valve.		
Motivation	The AGV travel/valve improved makes the manufacturing processes more lean, responsive and energy demanding.		
Target value	<20%		
Prerequisites	Fully integrated system		
aerOS components (task)	KrakenD (T3.1), Ingress (T3.1), Ingress&KrakenD conf (T3.1), CertManager&LetsEncrypt (TLS) (T3.1), FDQN (T3.1), NAT Capable (T3.1), OpenAPI (T3.2), AsyncAPI (T3.2), Low-Code (T3.2), HLO (T.3.3), LLO (k8s, Docker) (T3.3), KeyCloack (T3.4), KrakenD (T3.4), OpenLDAP (T3.4), Self-awareness (T3.5), Self-orchestration (T3.5), Self-diagnose (T3.5), Self-security (T3.5), Self-healing (T3.5), Self-Configuration (T3.5), Self-API (T3.5), Data-Interoperability (T4.1), LDAP (T4.2), EAT (T4.4), Trust SCO. (T4.5), IOTA (T4.5), Portal (T4.6).		
Evaluation means	The methodology for assessing the KPI will be based on measuring the following parameter on a fixed number of cycles: (Travel Factor) TF = Number of AGV travels / Number of valves carried		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0%* *(1 Travel per Valve)*	***N/A***	*N/A*
Outcome elaboration (M24)	AGV travel save is not yet measurable given that the integration will be finalized after M24. KPI results, based on the plan of activities, will be available at M32.		

4.1.10. KPI 2.1.10 Definition of the calculation model

Table 71 KPI 2.1.10 Definition of the calculation model

KPI ID number .	KPI 2.1.10
KPI Name	Definition of the calculation model
Description	Time required to define the calculation model for a specific product
Motivation	Faster definition of the calculation model for each specific product improves the overall efficiency of the production and demonstrates the effectiveness of aerOS regarding real-time data processing.

Target value	> 30%-time reduction		
Prerequisites	To calculate the time required to define the calculation model for a specific product, the IEs configuration must be complete, Node-Red configured, the data collection tool set up and the network and APIs configured.		
aerOS components (task)	Networking (T3.1), API & Low Code tools (T3.2), Orchestration (T3.3), T3.5 (Self-*), T4.2 (Data Fabric)		
Evaluation means	To measure the time required to define the CO_2 footprint calculation model for a specific product, start and end times of the model development process are recorded. This time is continuously tracked for each product using tracking tools. The total and average time taken are then calculated. Dashboards and reports are used to visualize and monitor these times, identifying areas for process improvement.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*120 minutes*	***90 minutes (25% reduction)***	*N/A*
Outcome elaboration (M24)	To define the model for calculating the CO_2 footprint of a specific product, the process takes 90 minutes at M24, a time reduction of 25% compared with the reference value of 120 minutes.		

4.1.11. KPI 2.1.11 Transparency of CO2/PCF data (minutes)

Table 72 KPI 2.1.11 Transparency of CO2/PCF data (minutes)

KPI ID number .	**KPI 2.1.11 (SIPBB)**
KPI Name	**Transparency of CO2/PCF data (minutes)**
Description	Time required to access CO2/PCF data for a specific product.
Motivation	Faster access to CO2/PCF data for each specific product allows greater transparency for customers and real-time control of the factory.
Target value	< 2 minutes
Prerequisites	To calculate the time required to access CO2/PCF data for a specific product, the IEs configuration must be complete, Node-Red configured, the data collection tool set up and the network and APIs configured, as well as the aerOS runtime and the web service to access the data.
aerOS components (task)	Networking (T3.1), API & Low Code tools (T3.2), Orchestration (T3.3), Cybersecurity tools (T3.4), T3.5 (Self-*), T4.2 (Data Fabric), aerOS Portal (T4.6)
Evaluation means	To measure the time required to access CO2/PCF data for a specific product, request and access times for each data retrieval are recorded. These times are continuously tracked using data management systems and time tracking tools. The total and average time taken to access the data are then calculated. Dashboards and reports are used to

	visualize and monitor these times, identifying areas for process improvement and reduce access times.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	Prerequisites have not been met: aerOS runtime and web service have not yet been deployed.		

4.2. Pilot 2 Containerised edge computing near renewable energy sources

4.2.1. KPI 2.2.1 Consumed renewable energy based on decision making process of aerOS

Table 73 KPI 2.2.1 Consumed renewable energy based on decision making process of aerOS

KPI ID number .	**KPI 2.2.1**		
KPI Name	**Consumed renewable energy based on decision making process of aerOS**		
Description	The total amount of renewable energy consumed on monthly basis.		
Motivation	KPI shows that the absolute energy usage is big enough to consider the pilot as representable.		
Target value	20 MWh/month		
Prerequisites	Containers need to be connected to power source.		
aerOS components (task)	All		
Evaluation means	Monitoring of power consumption energy meter.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0 MWh/month*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met, as the containers have not yet been mounted in their final location. This is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.2.2. KPI 2.2.2 Effectiveness of task distribution through aerOS to nodes

Table 74 KPI 2.2.2 Effectiveness of task distribution through aerOS to nodes

KPI ID number .	**KPI 2.2.2**		
KPI Name	**Effectiveness of task distribution through aerOS to nodes**		
Description	KPI shows the share of scheduled task completed on time.		
Motivation	The proper on-schedule job handling is crucial for the overall trust in the compute solution.		
Target value	99.5% of tasks executed on schedule		
Prerequisites	Container connected to power, connected to the network. aerOS continuum installed and ready to use.		
aerOS components (task)	HLO (T3.3), LLO.		
Evaluation means	Due date of each workload will be compared with actual end date of processing.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met, as the containers have not been mounted in their final location. It is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.2.3. KPI 2.2.3 Scalability of task distribution and management through aerOS

Table 75 KPI 2.2.3 Scalability of task distribution and management through aerOS

KPI ID number .	**KPI 2.2.3**
KPI Name	**Scalability of task distribution and management through aerOS**
Description	The amount of task scheduled by aerOS in Pilot 2 compute edges.
Motivation	KPI shows the flexibility and scalability of aerOS. The task might contain more than one job. Tasks might be batch or interactive type.
Target value	10k tasks executed/month

Prerequisites	Container connected to power, connected to the network. aerOS continuum installed and ready to use.		
aerOS components (task)	HLO (T3.3), LLO (T3.3), IdM (T3.4)		
Evaluation means	Task done will be counted based on logs and visualized with Grafana.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met, as the containers have not been mounted in their final location. This is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.2.4. KPI 2.2.4 CPU utilization efficiency

Table 76 KPI 2.2.4 4 CPU utilization efficiency

KPI ID number .	**KPI 2.2.4**		
KPI Name	**CPU utilization efficiency**		
Description	KPI shows the average CPU consumption by worker nodes (excluding master nodes control and network devices).		
Motivation	For energy saving it's very important to have a proper autoscaling solution. KPI shows the capability of disabling unused nodes to save energy.		
Target value	80%		
Prerequisites	Container connected to power, connected to the network. aerOS continuum installed and ready to use.		
aerOS components (task)	HLO (T3.3), LLO (T3.3)		
Evaluation means	Metrics from nodes will be reported in Prometheus and displayed on Grafana.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0%*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met, as the containers have not been mounted in their final location. This is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.2.5. KPI 2.2.5 Carbon awareness share of green energy

Table 77 KPI 2.2.5 Carbon awareness share of green energy

KPI ID number .	**KPI 2.2.5**		
KPI Name	**Carbon awareness share of green energy**		
Description	KPI shows the green energy share for jobs with green energy preference label.		
Motivation	System shall support the choice of the green energy when scheduling job. Some urgent jobs might be launched regardless of energy source, and some can strongly prefer/require green energy.		
Target value	60%		
Prerequisites	Containers are connected to green energy. Energy meters and Data Logger connected to appropriate places.		
aerOS components (task)	HLO (T3.3), LLO (T3.3)		
Evaluation means	Monitoring consumption of green energy at energy meter.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0%*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met, as the containers have not been mounted in their final location. It is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.2.6. KPI 2.2.6 Number of edge nodes connected in the aerOS continuum

Table 78 KPI 2.2.6 Number of edge nodes connected in the aerOS continuum

KPI ID number .	**KPI 2.2.6**
KPI Name	**Number of edge nodes connected in the aerOS continuum**
Description	The total count of pilot's edge nodes (physical locations).
Motivation	KPI shows more than one edge node with the different energy supply, it gives the opportunity to show advantages of the aerOS job distribution subsystem.
Target value	2
Prerequisites	Container is ready to Host RACK and HW. HW is available. HW is installed in container and properly configured.

aerOS components (task)	All		
Evaluation means	Count Number of containers serving as Edge Node.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met. The containers have not been mounted in their final location, which is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.2.7. KPI 2.2.7 Number of batch processing jobs successfully distributed and executed by the system

Table 79 KPI 2.2.7 Number of batch processing jobs successfully distributed and executed by the system

KPI ID number .	**KPI 2.2.7**		
KPI Name	**Number of batch processing jobs successfully distributed and executed by the system**		
Description	The number of batch jobs scheduled, orchestrated and executed by aerOS continuum.		
Motivation	Significant amount of batch jobs shows the ability of handling complex parallel computing tasks.		
Target value	300k		
Prerequisites	Container connected to power, connected to the network. aerOS continuum installed and ready to use.		
aerOS components (task)	HLO (T3.3), LLO.		
Evaluation means	Task done will be counted based on logs and displayed on Grafana.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met. The containers have not been mounted in their final location, which is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.2.8. KPI 2.2.8 Precision of the Future Price prediction algorithm

Table 80 KPI 2.2.8 Precision of the Future Price prediction algorithm

KPI ID number .	**KPI 2.2.8**		
KPI Name	**Precision of the Future Price prediction algorithm**		
Description	Precision of the price value predicted by running MLOps microservice compared to the actual price value published by the energy exchange the next day.		
Motivation	Significant deviation in predicted and actual value eliminates the usability of the microservice.		
Target value	85%		
Prerequisites	Containers infra operational, Electrum microservice for price estimation, aerOS runtime working, access to TGE and PSE work platform.		
aerOS components (task)	HLO and LLO (T3.3)		
Evaluation means	Price from TGE will be compared with estimated price of microservice. Grafana will be utilized for visualizing the estimations.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0%*	***N/A***	*N/A*
Outcome elaboration (M24)	The prerequisites have not been met. The containers have not been mounted in their final location, which is scheduled for September 2024. The KPI will be measured after the physical mounting and connection to the green energy power on site.		

4.3. Pilot 3 High performance computing platform for connected and cooperative mobile machinery

4.3.1. KPI 2.3.1.a (and KPI 2.3.2.b) Performance and connectivity capabilities improvement (single vehicle)

Table 81 KPI 2.3.1 Performance and connectivity capabilities improvement (single vehicle)

KPI ID number .	**KPI 2.3.1.a and KPI 2.3.2.b**
KPI Name	**Performance and connectivity capabilities improvement (single vehicle)**
Description	This complex KPI may include 2 measures of capabilities of a single vehicle: 1. KPI 2.3.1.a Performance without using AI-supported application(s), where the improvement should be more than 20%

	2. KPI 2.3.1.b Performance of the connectivity with temporary network infrastructure, meaning that high bandwidth connectivity e.g., 5G, should be available in rural environment to achieve the connectivity in so called dead areas with GPS (i.e., no connectivity at the moment)		
Motivation	For suggested KPI 2.3.1.a: As applications become more complex, they require more computing capabilities on the edge device. Mobile machinery for agriculture and construction applications poses hard challenges to developers of computers, because of the rugged environment and conditions in which they must operate. Measuring computing capabilities gives an indication on the innovation and engineering efforts expended in making the computer suitable for the aforementioned use-case. For suggested KPI 2.3.1.b: Connectivity is needed in order to realize the edge to cloud continuum which is the research topic in aerOS. Measuring the availability and sustained speed of network connectivity gives an indication of efforts spent in realizing the needed infrastructure both at the edge and at the cloud.		
Target value	For performance: GPU: 12.6 FP16 TFLOPS; CPU: SPEC int 2k6: 22, SPEC int rate: 140 Gflops. For connectivity: 4G/5G network available.		
Prerequisites	• Assembly and test of the prototype HW platform to be used in the pilot. • Integrating and testing of the required OS and libraries. • Availability of required interfaces between HW platform and other components with the target prototype vehicle		
aerOS components (task)			
Evaluation means	For performance: The performance is evaluated by integrating the TTControl platform, HPCP prototype extended with the NVIDIA-based packages, running aerOS software on prototype John Deere machines and executing the lab and field tests for John Deere's scenarios, while at the same time monitoring the computing resources utilization such as CPU, GPU and memory. Only with the execution of John Deere's applications, e.g. the sustainability impact can be measured. For connectivity: The connectivity using the temporary network will be tested by the aerOS SW and John Deere's applications (e.g. running operational instructions) on John Deere's prototype machines.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	For performance: GPU: 2x128 GFLOPS FP 16 CPU: 26000 DMIPS. For connectivity: No network available.	For performance: GPU: 12.6 FP16 TFLOPS; CPU: SPEC int 2k6: 22, SPEC int rate: 140 Gflops. For connectivity: 4G / 5G network available	N/A because was delivered M18

Outcome elaboration (M24)	The HW prototype are available for the laboratory use and testing and are currently used by both pilot partners TTControl and John Deere. The expected performance is achieved when executing the target SW in a laboratory setup.		

4.3.2. KPI 2.3.2 Swarm of vehicle performance improvement

Table 82 KPI 2.3.2 Swarm of vehicle performance improvement

KPI ID number .	**KPI 2.3.2**
KPI Name	**Swarm of vehicle performance improvement**
Description	Performance using AI-supported application(s) to monitor and optimize the integration of AI-based solutions to enhance vehicle efficiency, and overall performance. This KPI helps to identify areas for improvement, to fine-tune the AI algorithms, and ensure a seamless operation experience for end users.
Motivation	Applications based on Artificial Intelligence methodologies (Deep neural networks) will be developed to process images coming from sensors on the mobile machinery (cameras). The metric of frames per second indicates how efficient the AI algorithms are as well as how powerful the hardware is that has been developed for the use on the mobile machine (which is subject to the same constraints mentioned in KPI 2.3.1.1 with respect to accommodating powerful processing in the challenging environment of agriculture and construction machines). By leveraging aerOS, the goal is to improve this frame rate by at least 20%, enabling faster processing and subsequently increasing the tractor's operation speed. This improvement is possible due to more frequent updates on the field status, resulting from the higher frame rate.
Target value	Target value: 6 FPS pro Camera and 18 km/h
Prerequisites	Finalization of the preparation and setup of two electric prototype tractors and the associated implements. • Integration of High-Performance ECUs in tractors to convert them to IEs. • Establishment of a 4G/5G private network on the test field. • Setup of on-premises and cloud IEs such as computing nodes and VMs. • Ensuring the aerOS runtime is integrated and fully operational.
aerOS components (task)	aerOS Basic Services (Service Fabric (T3.3), Data Fabric (T4.2), Federated Orchestration (HLO/LLO) (T3.3), Management Portal (T4.5), Services and configuration API (T3.2), aerOS Auxiliary Services (e.g. Self-Capabilities (T3.5), Embedded Analytics (T4.3)), Virtualization
Evaluation means	Lab and field tests will be conducted to assess the AI's performance in ensuring accurate and efficient field work operations. Multiple AI models, such as a model designed to optimize tillage, will be employed to enhance different aspects of agricultural practices. The performance of the IEs, the prototype machines, the AI models, and the network infrastructure, such as achieved machine speed, compute

	resource utilization and required computation time, will be tracked by the embedded control software of the prototype machines and applications across the different IEs. The system's ability to adapt to real-time environmental and operational changes will be tested, along with scalability assessments to determine the capacity of aerOS to handle varying farm sizes and complexities. Impact evaluations will measure improvements in resource utilization efficiency.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*The baseline frame rate of 4 frames per second (FPS) per camera represents the current processing capacity for the exemplary task in the use case.*	***N/A***	*N/A*
Outcome elaboration (M24)	N/A, because this KPI was defined for M34.		

4.3.3. KPI 2.3.3 CO2 emissions reduction thanks to platooning

Table 83 KPI 2.3.3 CO2 emissions reduction thanks to platooning

KPI ID number .	**KPI 2.3.3**
KPI Name	**CO2 emissions reduction thanks to platooning**
Description	CO2 indicators to measure and track the CO2 emissions and subsequent reduction due to the utilization of electric tractors and the aerOS services. Here in particular for the cultivating/grubbing activity during stubble cultivation.
Motivation	The motivation behind the CO2 emissions reduction KPI in the aerOS project is to quantify and assess the environmental impact of deploying the aerOS solution and transitioning from diesel-powered tractors to electric tractors. Climate change and environmental conservation are increasingly important global concerns, and the reduction of CO2 emissions is a crucial step towards addressing these challenges.
Target value	A reduction of 80% - 17,9 kg CO2/ha
Prerequisites	• Finalization of the preparation and setup of two electric prototype tractors and the associated implements. • Integration of High-Performance ECUs in tractors to convert them to IEs. • Establishment of a 4G/5G private network on the test field. • Setup of on-premise and cloud IEs such as computing nodes and VMs. • Ensuring the aerOS runtime is integrated and fully operational.

aerOS components (task)	aerOS Basic Services (Service Fabric, Data Fabric, Federated Orchestration (HLO/LLO), Management Portal, Services and configuration API, aerOS Auxiliary Services (e.g. Self-Capabilities, Embedded Analytics), Virtualization		
Evaluation means	Lab and field tests will be conducted to evaluate capabilities of aerOS in optimizing resource utilization and real-time adaptation of machine operations. These evaluations focus on optimizing resource use and dynamically adjusting operations to changing conditions. The impact on reducing CO2 emissions through sustainable practices is measured by monitoring the power consumption and operational performance of the prototype machines during field and overall operations. The data is tracked by the embedded control software of the machines, ensuring operational quality remains high.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	89,31 kg CO2/ha (33,7 l Diesel/ha)	N/A, because this KPI was defined for M34.	
Outcome elaboration (M24)	N/A, because this KPI was defined for M34.		

4.4. Pilot 4 Smart edge services for the port continuum

4.4.1. KPI 2.4.1 Reduction of CHE idle time due to failures

Table 84 KPI 2.4.1 Reduction of CHE idle time due to failures

KPI ID number .	**KPI 2.4.1**
KPI Name	**Reduction of CHE idle time due to failures**
Description	The preventive maintenance tool and approach used in EUROGATE Container Limassol plans the maintenance task according to some number of working hours. This sub-optimal approach is frequently not enough for removing any unexpected failure of Container Handling Equipment (CHE) components, and undesired idle times at operational hours occur. aerOS predictive maintenance models are expected to reduce these IDLE times.
Motivation	It will show how the predictive maintenance on the edge service to be deployed in the project provides a relevant benefit to EUROGATE operational efficiency.
Target value	20-30%
Prerequisites	All sensors are deployed in the CHEs under tests, and their maintenance associated data is acquired and collected for AI modelling. Access to the CMMS system of EUROGATE is also required.
aerOS components (task)	OpenAPI (T3.2), AAA (T3.4), Context Broker (T4.2), Data Fabric (T4.2)

<table>
<tr><td>Evaluation means</td><td colspan="3">A report of the idle time of CHEs under study in the project due to maintenance tasks will be extracted for the first two years of the project from the Computerised Management System (CMMS) of EUROGATE. Once the predictive maintenance service from aerOS is deployed, an analysis between the original idle times with the preventive maintenance and the new ones with the predictive maintenance will be carried out. If the idle time of the use cases of the pilot is reduced at least 20%, the KPI will be considered as fulfilled.</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>Total 2023 downtime of 4 straddle Carriers: 900h
Total 2023 downtime of 2 STS: 297.70h</td><td>Q1-Q2 2024 downtime of 4 straddle Carriers: 514h
Q1-Q2 2024 downtime of 2 STS: 69.1h</td><td></td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Models are not ready, but monitoring system is on place. Thanks to this, abnormal situations were spotted, and EUROGATE maintenance procedures have efficiently changed to a more predictive approach, reducing time to identify failures and idle times (i.e., 30mins repairing vs original 2h).
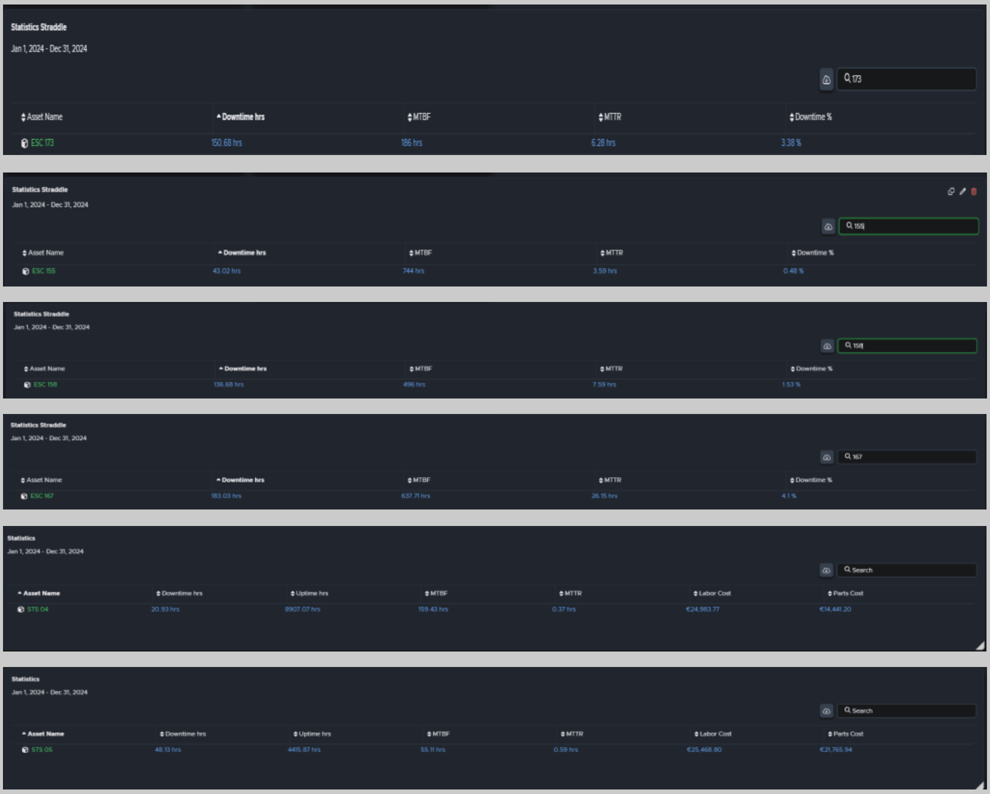
</td></tr>
</table>

4.4.2. KPI 2.4.2 Increase on detection of equipment malfunctions (from manual to automatic)

Table 85 KPI 2.4.2 Increase on detection of equipment malfunctions (from manual to automatic)

<table>
<tr><th>KPI ID number .</th><th colspan="3">KPI 2.4.2</th></tr>
<tr><td>KPI Name</td><td colspan="3">Increase on detection of equipment malfunctions (from manual to automatic)</td></tr>
<tr><td>Description</td><td colspan="3">The maintenance of EUROGATE CHEs is based on preventive inspection. The predictive maintenance service will be able to detect equipment malfunctions more precisely.</td></tr>
<tr><td>Motivation</td><td colspan="3">It will show how the predictive maintenance on the edge service to be deployed in the project provides a relevant benefit to EUROGATE operational efficiency.</td></tr>
<tr><td>Target value</td><td colspan="3">30-40% with respect to 2023</td></tr>
<tr><td>Prerequisites</td><td colspan="3">All sensors are deployed in the CHEs under tests, and their maintenance associated data is acquired and collected for AI modelling. Access to the CMMS system of EUROGATE is also required.</td></tr>
<tr><td>aerOS components (task)</td><td colspan="3">OpenAPI (T3.2), AAA (T3.4), Context Broker (T4.2), Data Fabric (T4.2), Explainability service (T4.3)</td></tr>
<tr><td>Evaluation means</td><td colspan="3">A comparative analysis between the manual equipment malfunctions reported for the four CHEs under study in the project versus the automatic ones provided by the predictive maintenance service will be conducted. If the proper identification of malfunctions is increased at least 30%, the KPI will be considered as fulfilled</td></tr>
<tr><td>Measurement period</td><td>Baseline</td><td>M24 (Deliverable D5.5)</td><td>M36 (Deliverable D5.6)</td></tr>
<tr><td>Measured value (% achieved)</td><td>In 2023: 30 unplanned failures detected (manual), 0 predictive (automatic)</td><td>Q1-Q2 2024: 15 unplanned failures detected (manual), 0 predictive (automatic)</td><td></td></tr>
<tr><td>Outcome elaboration (M24)</td><td colspan="3">Like previous KPI, models are not ready, but monitoring system is on place. Thanks to this, abnormal hydraulic situations were spotted, and EGCTL maintenance scheduled maintenance tasks for the associated CHEs.

Before and After Time required for 4000h Maintenance Service

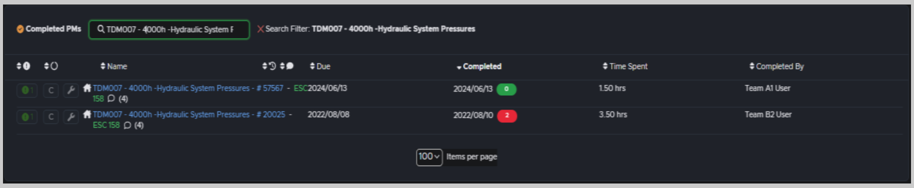

Hydraulic Issues detected by observation

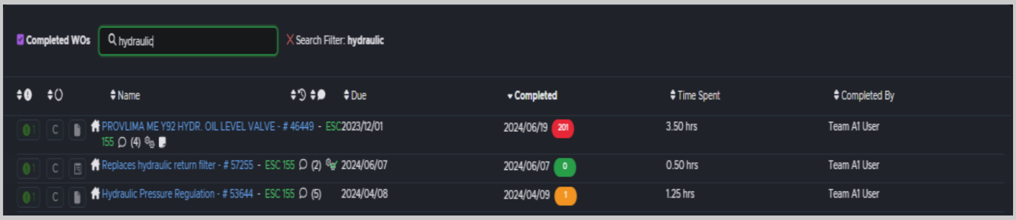

</td></tr>
</table>

4.4.3. KPI 2.4.3 Increase of number of actual damaged containers (manually reported by staff vs automatic system-reports)

Table 86 KPI 2.4.3 Increase of number of actual damaged containers (manually reported by staff vs automatic system-reports)

KPI ID number .	**KPI 2.4.3**		
KPI Name	**Increase of number of actual damaged containers (manually reported by staff vs automatic system-reports)**		
Description	When loading/discharging containers to/from vessels by Ship-to-Shore (STS) cranes, a manual inspection by the port stevedores is carried out in order to confirm there are no wrong seals and damages generated during the manoeuvre. These reports will be more accurate if an automatic system which makes use of cameras and computer vision functionalities is deployed by aerOS.		
Motivation	It will show how the Computer Vision (CV) on the edge service to be deployed in the project provides a relevant benefit to EUROGATE business.		
Target value	30-40%		
Prerequisites	All cameras are deployed under STS crane operations, their video streams are received in an AV server, and the accurate enough CV models are available and in execution. In addition, access to the ERP system data of EUROGATE is also need for data comparison.		
aerOS components (task)	Context Broker (T4.1), Data Fabric (T4.2), HLO and LLO (T3.3), Model Reduction service (T4.3), Management portal (T4.6)		
Evaluation means	A historical analysis of EUROGATE customers complaints, which demand penalties for not complying with the SLAs will be collected for the first two years of the project from the ERP system of EUROGATE. A comparison between the number of penalties received before and after the CV service from aerOS is deployed will be carried out. If the proper detection of actual damaged containers leads to a reduction of 30% of complaints procedures, the KPI will be considered as fulfilled.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*350 damaged containers reported by terminal staff + 30 damaged containers not reported and claimed*	***Q1-Q2 2024: 248 damaged containers reported by terminal staff + 15 damaged containers not reported and claimed (not using CV models)***	*N/A*
Outcome elaboration (M24)	During the period from January to June 2024, a total of 248 containers were discharged with damage from ships and were identified by the terminal staff. A review of the damages that were identified related to side dents or holes, which stemmed from violent or negligent handling. The monthly figures of containers identified is the following:		

January: 43

February: 41

March: 48

April: 34

May: 41

June: 41

Furthermore, another 15 containers cases that were discharged with damage and not identified by terminal staff were recorded during the period. A relevant complaint to investigate the matter was by the customer that owns the container. Given that the CV model has not yet been deployed as its still in development all checks for damages were performed by terminal staff and this inadvertently led to mistakes

Reviewing the cases that have been missed, the oversight of reporting damages by terminal staff does not seem to pertain to a specific type of damage but rather to fluctuations in personnel attention, which tends to decrease during specific times of the day.


08-06-2024 Sat 18:54:31
stinis
MEDITERRANEAN
SHIPPING CO
QC4 Gantry Cam1


APZU 337144 7
22G1
MAX. GROSS 30,480 KGS
67,200 LBS
TARE 2,220 KGS
4,890 LBS
NET 28,260 KGS
62,310 LBS
CU.CAP. 33.2 CU.M.
1,170 CU.FT.
CIMC

4.4.4. KPI 2.4.4 Performance evaluation metrics of regression AI models (R2)

Table 87 KPI 2.4.4 Performance evaluation metrics of regression AI models (R2)

KPI ID number .	**KPI 2.4.4**
KPI Name	**Performance evaluation metrics of regression AI models (R2)**
Description	Different regression AI models will be developed and deployed at the edge nodes of the Port Continuum domain. These models shall be accurate enough to predict equipment malfunctions before any failure occur. In regression models, R-square (R2) corresponds to the squared correlation between the observed outcome values and the predicted values by the model. The Higher the R-squared, the better the model.
Motivation	An accurate regression model should be provided in order to replace the current preventive maintenance for the new one developed in the project.
Target value	0.8

Prerequisites	Large amount of time-series data from the CHEs shall be collected for ML model training. A Python script in charge of accuracy validation shall also be available.		
aerOS components (task)	Data Fabric (T4.2), Explainability service (T4.3)		
Evaluation means	From the different CHEs' telemetry dataset collected, a portion of them will be used for validation purposes. Python-based Juptyer notebooks will be used for evaluating the R2 metric of the developed model against these validation datasets. As long as the model surpassed R2>=0.8, the KPI will be fulfilled.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	Despite having already a large amount of data, the data scientist team is analysing the different parameters that must be used for correlation and as ML features. Hence, the models are still under development, without having an initial R2 value. The plan is to speed up during the last year of the project with results associated to R2 of the models every quarter.		

4.4.5. KPI 2.4.5 Performance evaluation metrics of regression AI models (MAE/RMSE) for predictive maintenance of CHEs

Table 88 KPI 2.4.5 Performance evaluation metrics of regression AI models (MAE/RMSE) for predictive maintenance of CHEs

KPI ID number .	**KPI 2.4.5**
KPI Name	**Performance evaluation metrics of regression AI models (MAE/RMSE) for predictive maintenance of CHEs**
Description	Different regression AI models will be developed and deployed at the edge nodes of the Port Continuum domain. These models shall be accurate enough to predict equipment malfunctions before any failure occur. Similarly to R2, other commonly used evaluation metric for regression models is Mean Average Error or Root Mean Squared Error (MAE/RMSE). Both metrics refer to refers to the mean of the absolute values of each prediction error on all instances of the test dataset. The lower the MAE/RMSE, the better the model.
Motivation	An accurate regression model should be provided in order to replace the current preventive maintenance for the new one developed in the project.
Target value	20%
Prerequisites	Large amount of time-series data from the CHEs shall be collected for ML model training. A Python script in charge of accuracy validation shall also be available.

aerOS components (task)	Data Fabric (T4.2), Explainability service (T4.3)		
Evaluation means	From the different CHEs' telemetry dataset collected, a portion of them will be used for validation purposes. Python-based Jupyter notebooks will be used for evaluating the MAE/RMSE metric of the developed model against these validation datasets. As long as the error of the model <20%, the KPI will be considered as fulfilled.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	Despite having already a large amount of data, the data scientist team is analysing the different parameters that must be used for correlation and as ML features. Hence, the models are still under development, without having an initial MAE/RMSE value. The plan is to speed up during the last year of the project with results associated to MAE/RMSE of the models every quarter.		

4.4.6. KPI 2.4.6 Performance evaluation metrics of classification AI models (accuracy) for damaged containers

Table 89 KPI 2.4.6 Performance evaluation metrics of classification AI models (accuracy) for damaged containers

KPI ID number .	**KPI 2.4.6**
KPI Name	**Performance evaluation metrics of classification AI models (accuracy) for damaged containers**
Description	Different classification AI models will be developed and deployed at the edge nodes of the Port Continuum domain. These models shall be accurate enough to detect and classify damages identified at containers' surfaces. Accuracy is one of the most common metrics used for the evaluation of these classification models. AI accuracy is the degree to which an AI system produces correct outputs or predictions based on the given inputs or data. Therefore, if the AI system classifies damages, its accuracy is the percentage of images that the model correctly labels as dents, etc.
Motivation	An accurate classification model should be provided in order to replace guarantee the proper detection and classification of surfaces across the loaded/unloaded containers.
Target value	60%
Prerequisites	Large amount of video streams recorded needed for CV models ML training shall be available, especially with damages visible on containers' surfaces as part of the data set.
aerOS components (task)	N/A
Evaluation means	From the different videos collected from the cameras deployed in the dock area, a portion of them will be used for validation purposes. TensorFlow evaluation tool that

	offers various libraries for model validation, testing, and evaluation will be used. As long as the accuracy of the model > 60%, the KPI will be considered as fulfilled.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*N/A*	***mAP50: 53%***	*N/A*
Outcome elaboration (M24)	Two models were developed, one for container detection and one for damage detection. Both models utilize a pre-trained YOLOv8 model. The container detection model was trained on 533 images and validated on 100 images, while the damage detection was trained on 470 images and validated on 70 images. The container detection model achieved very strong validation results with a mAP50 score of 96%. Note that mAP50 measures the mean average precision at an intersection over union (IoU) threshold of 0.5. On the contrary, the damage detection model achieved a mAP50 score of 53%, which is lower than the desired 60% score. Active efforts are focused on improving this score by exploring more advanced image pre-processing techniques, collecting more training data, and performing more extensive hyper-parameter tuning. Validation results for damage detection: 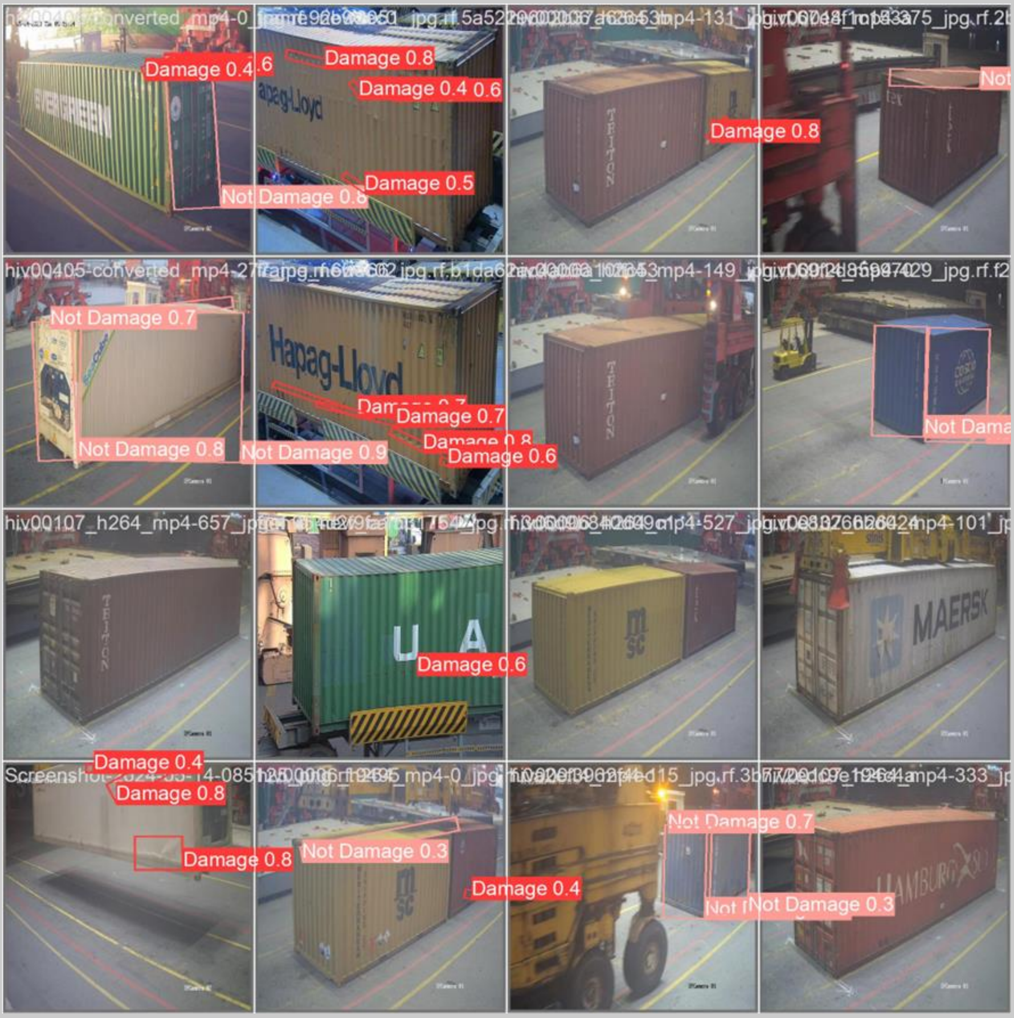 mAP50 and mAP50-95 for damage detection:		

	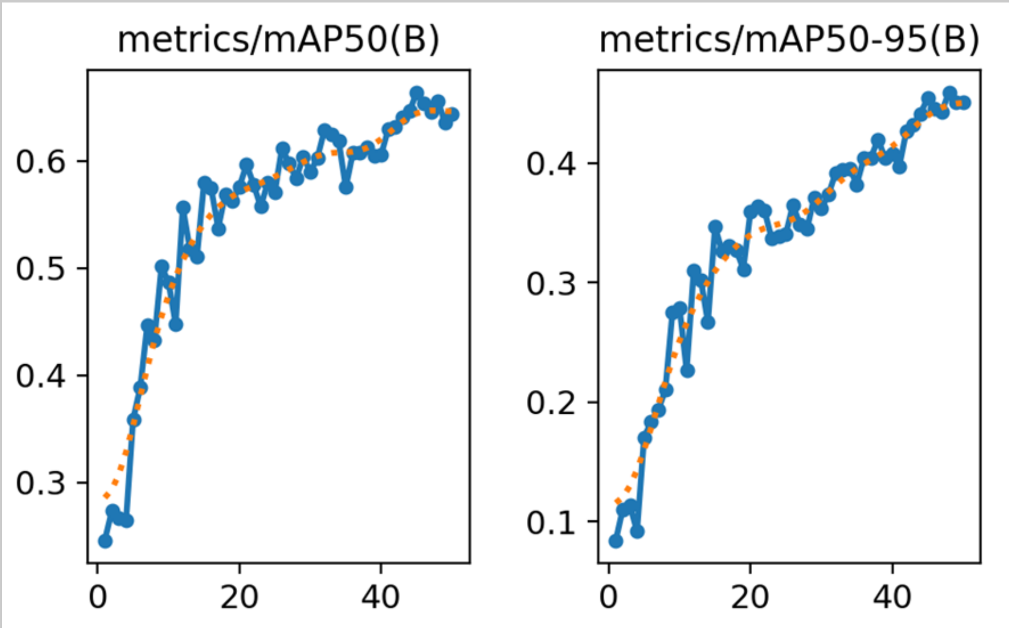

4.4.7. KPI 2.4.7 Performance evaluation metrics of classification AI models (F1) for damaged seals

Table 90 KPI 2.4.7 Performance evaluation metrics of classification AI models (F1) for damaged seals

KPI ID number .	KPI 2.4.7
KPI Name	**Performance evaluation metrics of classification AI models (F1) for damaged seals**
Description	While accuracy is often used as a primary indicator of the quality and effectiveness of an AI system, there are other metrics like precision and recalls that help to evaluate the quality of a model. This KPI will evaluate the F1 score. F1 balances the trade-off between precision and recall, which can vary depending on the model and the data.
Motivation	An accurate classification model should be provided in order to replace guarantee the proper detection and classification of surfaces across the loaded/unloaded containers.
Target value	60%
Prerequisites	Large amount of video streams recorded needed for CV models ML training shall be available, especially with wrong or damaged seals included in the data set.
aerOS components (task)	N/A
Evaluation means	From the different videos collected from the cameras deployed in the dock area, a portion of them will be used for validation purposes. TensorFlow evaluation tool that offers various libraries for model validation, testing, and evaluation will be used. As long as the F1 of any of the model developed in the project is > 60%, the KPI will be considered as fulfilled.

Measurement period	Baseline	M24 (Deliverable D5.5)	M36 (Deliverable D5.6)
Measured value (% achieved)	*N/A*	***N/A***	*N/A*
Outcome elaboration (M24)	Despite having already a large number of videos, it has been found extremely challenging to identify wrong seals spots. Due to this lack of ground truth, the CV team has prioritized the final training of the ML models associated to the detection of damage on containers' surfaces. The plan is to continue collecting data with wrong seals until end of 2024 and focus on the training during 2025.		

4.4.8. KPI 2.4.8 Number of models executed on edge nodes

Table 91 KPI 2.4.8 Number of models executed on edge nodes

KPI ID number .	KPI 2.4.8		
KPI Name	Number of models executed on edge nodes		
Description	This KPI will evaluate the scalability capabilities of the models that are going to be developed in the port continuum pilot of the project. Since the goal is to have as lightweight as possible AI models, the way to confirm that approach is by confirming that these developed models can perform their inference process properly at the edge, without requiring high computational resources.		
Motivation	The IEs / nodes that are being used in Port Continuum pilot do not provide high processing capabilities. Frugal and lightweight AI models shall be developed in order to guarantee that they are run under these low-processing conditions.		
Target value	5		
Prerequisites	The Infrastructure Elements of Pilot 4 are commissioned and available for ML models deployment.		
aerOS components (task)	Self-* (T3.5), Model reduction service (T4.3), Embedded Analytics Tool (T4.4)		
Evaluation means	The Pilot 4 models will be deployed in the edge IEs of the port continuum (either the IEs of the predictive maintenance use case, or the IEs attached to the cameras of the damaged detection through CV use case). Logs from the OpenCV instance running on these IEs will be collected, proving if the new models are deployed and under successful execution. As long as 5 models in total are running, the KPI will be fulfilled.		
Measurement period	Baseline	M24 (Deliverable D5.5)	M36 (Deliverable D5.6)
Measured value (% achieved)	N/A	N/A	

Outcome elaboration (M24)	Since the models are still under training phase, they are just deployed in the development environment, and not on the IEs of the port continuum pilot. As long as they are accurate enough, they will be deployed.

4.5. Pilot 5 Energy efficient, health safe and sustainable smart buildings

4.5.1. KPI 2.5.1 Energy use reduction

Table 92 KPI 2.5.1 Energy use reduction

KPI ID number .	**KPI 2.5.1**		
KPI Name	**Energy use reduction**		
Description	20% Energy use reduction, using frugal AI and real-time processing in aerOS rather than in the cloud.		
Motivation	Energy consumption is a significant operational cost factor that all enterprises seek to reduce. Furthermore, energy efficiency is a strategic sustainability target for most enterprises, and especially for MNOs that maintain many sites.		
Target value	20% reduction of the daily baseline consumption.		
Prerequisites	IoT Domain ready, AI deployment for inference complete.		
aerOS components (task)	T3.1 (Networking), T3.3 (Orchestration), T3.4 (Self-*), T4.2 (Data fabric), T4.1 (Semantic Translation & Annotation)		
Evaluation means	Energy utilization is selectively measured by the pilot using smart metering devices and related data are collected and stored for post-processing. Furthermore, energy-related AI forecasting is expected through open calls to produce accurate approximation of future values for comparison.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*Baseline measurements vary per room, but some indicative consumptions to be reported without the aerOS optimisation range from 40Kwh - 150Kwh.*	***N/A***	*N/A*
Outcome elaboration (M24)	aerOS intelligence gains can be measured upon the completion of the first end-to-end pilot integration which is targeted for M30. Thus, the results can only be provided by M36.		

4.5.2. KPI 2.5.2 Edge processing performance gains

Table 93 KPI 2.5.2 Edge processing performance gains

KPI ID number .	**KPI 2.5.2**		
KPI Name	**Edge processing performance gains**		
Description	Edge processing and IoT performance gains, by evaluating the performance characteristics of the solution.		
Motivation	An extensive number of IoT sensors are deployed in the Smart Buildings ecosystem, generating/processing huge amount of data that are only valid for the location they originate from, yielding their transmission and collection for central processing meaningless and wasteful. The distinctive infrastructure characteristics of each building rationalize the autonomous and decentralized decision-making at the edge with the use of the aerOS nodes intelligence, and the effects are instantaneous and tactile.		
Target value	The measurement of the Edge processing performance gains is a composite KPI that can be approximated by collecting the following sub-KPIs 1. Exhibit average E2E Communication Latency < 100 ms for the aerOS nodes deployed locally (in the edge), measured through ping tools. 2. Demonstrate the gains of KubeEdge vs. K8 deployments utilising light devices at the far edge gaining 20 % less memory resources consumption comparing the cluster reported average measurement values. 3. Demonstrate the gains of Kube Edge for service resilience, measuring the service recovery time under various disruptive conditions showcasing 90% increase in recovery time (Kube-edge vs. K8)		
Prerequisites	IoT Domain Ready, 2 aerOS IE running and aerOS runtime working		
aerOS components (task)	T3.1 (Networking), T3.3 (Orchestration), T3.5 (Self-*)		
Evaluation means	Use of aerOS self-* capabilities for nodes monitoring and measurement tools through network protocols (e.g., ping)		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*The pilot is implemented on premises and dedicated networks already and typical values monitored include:* *Latency: 2-3 ms* *Memory: 1.5 Gbps.*	***100%*** • ***Latency of communication between the pilot5 aerOS nodes (ms): Average: 0.919 ms*** • ***Memory utilization when deploying IoT Application in a KubeEdge node: 730 Mbytes***	*N/A*

		• ***Time to recover IoT application when master node is down***	

Outcome elaboration (M24)

In M24 the basic aerOS Infrastructure Element capabilities are in place and the relevant transformation of the IoT Pilot components is completed, and relevant measurements are collected.

Baseline Latency Average (2.4 ms)

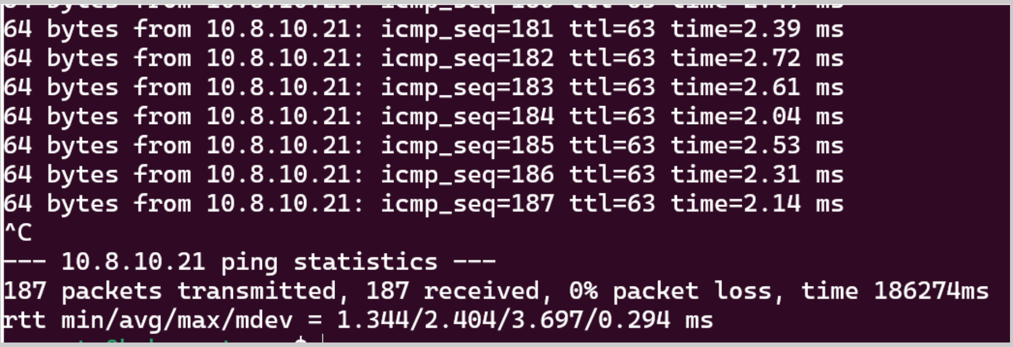

M24 Latency Average (0.919 msec)

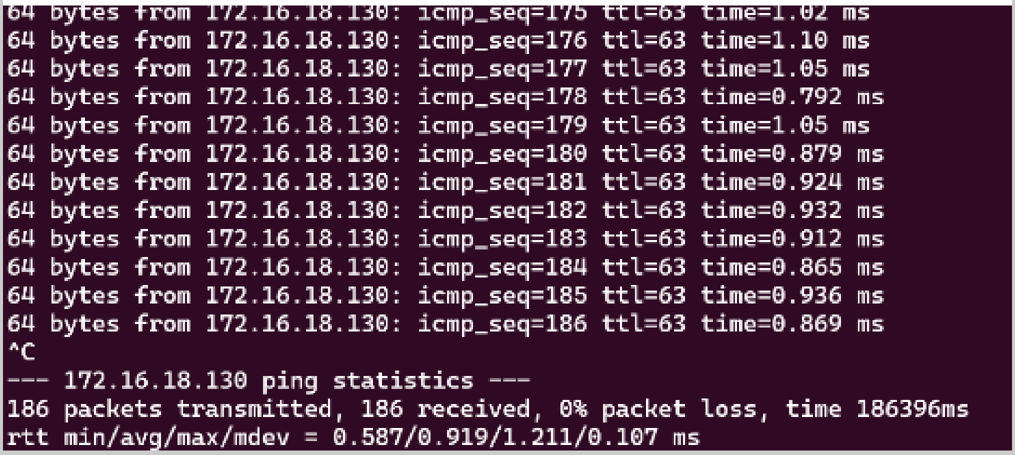

Baseline Memory of IoT K8 (1.5 Gbps)

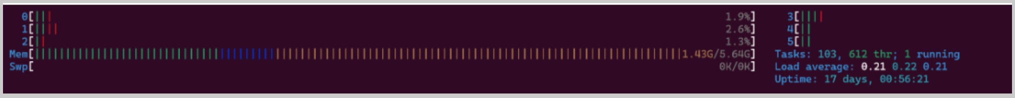

M24 Memory of IoT KubeEdge Node (730 Mbytes)

Baseline Service recovery (k8) – IoT Application is NOT running (right window) when k8s master node or network is down (left window).

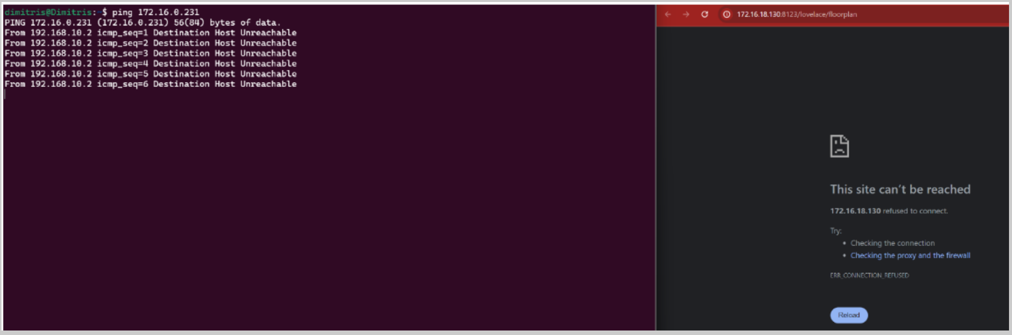

M24 Service Recovery (KubeEdge) – IoT Application (right window) is running when K8 master node or network is down (left window).

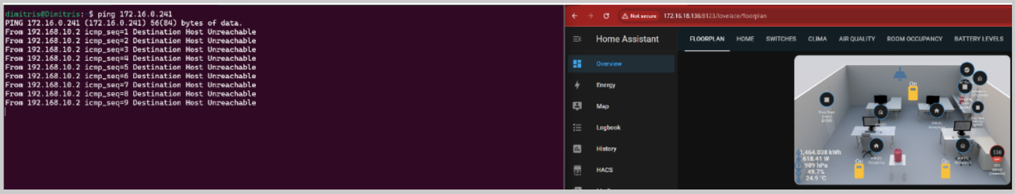

4.5.3. KPI 2.5.3 5G capabilities to execute security and privacy functions

Table 94 KPI 2.5.3 5G capabilities to execute security and privacy functions

KPI ID number .	**KPI 2.5.3**		
KPI Name	**5G capabilities to execute security and privacy functions**		
Description	Development of VNFs/CNFs in the 5G network to be integrated in aerOS to execute certain security and privacy functions will be evaluated		
Motivation	Leveraging niche network technologies and the 5G capabilities is an important tool to enhance the secure and reliable communication of the IOT system as well as to enhance the end-users' interactions.		
Target value	5		
Prerequisites	5 IEs setup complete, aerOS runtime working5G connectivity		
aerOS components (task)	T3.1 (Networking), T3.3 (Orchestration, HLO), T3.4 (Cybersecurity components)		
Evaluation means	List of CNFs deployed within aerOS domains will be provided. Monitoring capabilities of K9s tools will be used to export screenshots demonstrating CNFs deployment.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***N/A***	*N/A*

Outcome elaboration (M24)	The focus up to M24 is to identify the subgroup of aerOS API needed to facilitate the integration of security and privacy functions for 5G network. The development depends on a complete aerOS API specification which is currently under finalization and is also based on advanced HLO capabilities which are now under development. Combining the above with the identification of the appropriated APIs which is now an ongoing task will allow the development of security and privacy functions for 5G network. This is not expected to be finalised before M32 as development of these CNFs will follow requirements fulfillment. Thus, report on the results will be available in M36.

4.5.4. KPI 2.5.4 Service availability

Table 95 KPI 2.5.4 Service availability

KPI ID number .	**KPI 2.5.4**		
KPI Name	**Service availability**		
Description	The aerOS automation responds to failures by instantly re-deploying failed nodes with minimum interruption time.		
Motivation	Due to the distributed characteristics of the smart buildings IoT deployment, with vast number of sensors managed by nodes locally deployed per room and building it is important that automation systems ensure that all nodes are running with minimum interruption time.		
Target value	99.99% in the service window of operations		
Prerequisites	At least one IEs setup complete, aerOS runtime working		
aerOS components (task)	T3.1 (Networking), T3.3 (Orchestration), T3.5 (Self-*)		
Evaluation means	The availability of the pilot's IE and service application measured through the node's uptime in the service window period of 1 month for at least 3 consecutive months following the final installation of all the aerOS meta-OS intelligence. The service window is defined to be the actual expected window of operation, that exclude known maintenance periods.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*Manual operation*	***100%*** ***99,9999% in the service window a period of one month for at least one pilot node.*** ***Uptime: 25 days in the service window of 1 month***	*N/A*

<table>
<tr><td>Outcome elaboration (M24)</td><td>Demonstrate the uptime of at least one pilot node, for a period of one month using raw data from the host. A known maintenance activity is included in this period explaining the uptime to be 25 days.
<pre>
cosmote@node01:~$ uptime
 11:38:43 up 25 days, 3 min, 1 user, load average: 0.92, 0.89, 1.16
</pre>
<pre>
cosmote@aerosmaster:~$ k get pods
NAME   READY   STATUS   RESTARTS   AGE
grafana-7c96dc94cd-5hn42   1/1   Running   0   24d
hlo-allocator-deployment-5f4dbcf5-rbjhp   1/1   Running   3 (24d ago)   46d
hlo-data-aggregator-7454bb88c4-d2qr7   1/1   Running   2 (24d ago)   28d
hlo-deployment-engine-ddff75f96-8sh8p   1/1   Running   2 (24d ago)   28d
hlo-frontend-76fc48bb64-zvm2n   1/1   Running   3 (24d ago)   46d
homeassistant-hacs-fsd7z   1/1   Running   4 (24d ago)   126d
influxdb-66bf49cf84-z2c2h   1/1   Running   0   24d
kapacitor-844f5c9fd4-qtlkv   1/1   Running   2575 (65s ago) 24d
krakend-cb9b9ffbf-vwg5r   1/1   Running   2 (24d ago)   28d
metallb-controller-6cb58c6c9b-rt84b   1/1   Running   17 (24d ago)   179d
metallb-speaker-97xgk   4/4   Running   8 (131d ago)   165d
metallb-speaker-dxzds   4/4   Running   75 (24d ago)   213d
metallb-speaker-rcx9c   4/4   Running   40 (24d ago)   213d
mqtt-simple-6c6896dbb4-g7799   1/1   Running   0   24d
orion-ld-broker-7c4d9f7f65-w4xsq   1/1   Running   0   21d
orion-ld-mongodb-0   1/1   Running   0   21d
redpanda-0   2/2   Running   6 (24d ago)   46d
redpanda-console-6c8ddcff97-ksk74   1/1   Running   8 (24d ago)   46d
self-awareness-hardwareinfo-tcpkv   1/1   Running   3 (21d ago)   21d
self-awareness-hardwareinfo-zpdq7   1/1   Running   3 (21d ago)   21d
self-awareness-powerconsumptionamd64-krxtj   1/1   Running   3 (21d ago)   21d
self-awareness-powerconsumptionamd64-m47zj   1/1   Running   3 (21d ago)   21d
self-orchestrator-orchestrator-g95mj   1/1   Running   2 (24d ago)   46d
self-orchestrator-orchestrator-n9g7w   1/1   Running   3 (24d ago)   46d
telegraf-mqtt-54555fc8cb-fl9qm   1/1   Running   1 (21d ago)   24d
</pre></td></tr>
</table>

4.5.5. KPI 2.5.5 Service creation / scalability

Table 96 KPI 2.5.5 Service creation / Scalability

KPI ID number .	KPI 2.5.5
KPI Name	**Service creation / scalability**
Description	Demonstrate the capability of dynamic provisioning of the service as well as scaling in and out of buildings
Motivation	As new rooms, floors, buildings, sites are added in the Smart Buildings ecosystem per enterprise, it is important that the process to incorporate these is dynamic, transparent, and easy.
Target value	< 10 min end-to-end
Prerequisites	aerOS runtime working
aerOS components (task)	T3.1 (Networking), T3.3 (Orchestration), T3.5 (Self-*)
Evaluation means	Measure the time-to-deploy one IoT GW (a core pilot-5 service) leveraging the aerOS orchestration capabilities using the aerOS self-* capabilities as well as the OS system commands (e.g. time) to retrieve the clock time of start and end deployment.

Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*Manual*	*100%* *Time-to-deploy: 34 secs*	*N/A*
Outcome elaboration (M24)	Measure the time to deploy using basic system tools 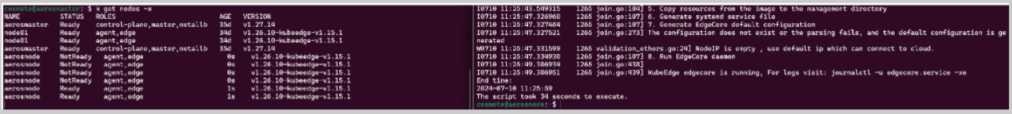		

4.5.6. KPI 2.5.6 Services directly managed by the aerOS orchestrator

Table 97 KPI 2.5.6 Services directly managed by the aerOS orchestrator

KPI ID number .	**KPI 2.5.6**		
KPI Name	**Services directly managed by the aerOS orchestrator**		
Description	Number of services/workloads directly managed by the FOM and deployed along the IEs		
Motivation	Efficient use of available computational resources and dynamic migration of workloads to maximise performance is enabled through the operations of federation automation as developed by aerOS. All application components and services of the smart buildings pilot must be managed by the federation orchestration (HLO/LLO) capabilities so that to always operate on the most appropriate Infrastructure Element at a given time.		
Target value	3		
Prerequisites	3 IEs hosting distinct pilot services (IoT GWs) complete, aerOS runtime working		
aerOS components (task)	T3.1 (Networking), T3.3 (Orchestration)		
Evaluation means	Exhibit the management of 3 pilot services through the aerOS monitoring (self-*) dashboards.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***Exhibit the operation of 3 pilot5 services in the aerOS-capable infrastructure (k8s/KubeEdge) (100%)***	*N/A*
Outcome elaboration (M24)	Install three (3) pilot 5 application components (IoT GWs) on the aerOS-capable infrastructure.		

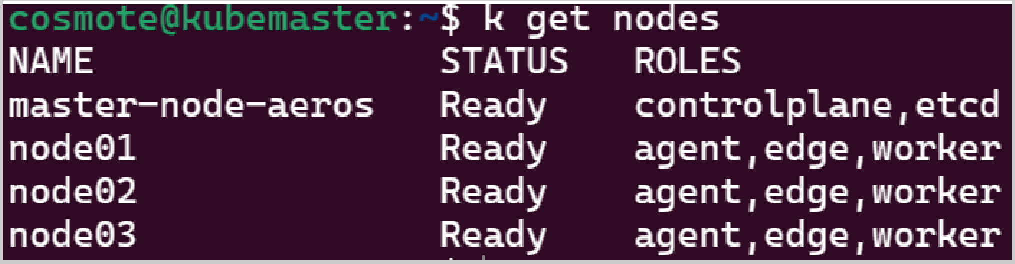

```
cosmote@kubemaster:~$ k get nodes
NAME                 STATUS   ROLES
master-node-aeros    Ready    controlplane,etcd
node01               Ready    agent,edge,worker
node02               Ready    agent,edge,worker
node03               Ready    agent,edge,worker
```

4.5.7. KPI 2.5.7 Improvement of air quality

Table 98 KPI 2.5.7 Improvement of air quality

KPI ID number .	**KPI 2.5.7**		
KPI Name	**Improvement of air quality**		
Description	Reduction of CO_2 levels (or other gasses) because of using frugal AI and real-time processing in aerOS to achieve an efficient distribution of workers in the office.		
Motivation	Health safety at office buildings is a societal requirement following the pandemic.		
Target value	A typical acceptable target is set to be 400-600 ppm per room for the demo, average > 20% improvement. Especially for the rooms of the pilot, and the specific demo situation, the target is set to me to reduce the max CO_2 lower than 1000 ppm in all cases.		
Prerequisites	IoT Domain ready, AI deployment for inference complete, IoT Actuation finalized		
aerOS components (task)	T3.1 (Networking), T3.3 (Orchestration), T3.4 (Self-*), T4.2 (Data Fabric), T4.1 (Semantic Translation & Annotation)		
Evaluation means	The evaluation can be achieved by measuring the ppm values from the sensors of a room with a certain number of employees for the first half of the day. For the second half of the day, activate the aerOS system and observe the improvements in the ppm values. Provisionally this can be extended to measuring the ppm values in a specific room over the course of one week, and assuming that the exact conditions can be recreated, measure with the aerOS intelligence activated to compare the results.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*Relative value per room. Spike values in the range of 1200-1500 ppm are measured.*	***N/A***	*N/A*
Outcome elaboration (M24)	aerOS intelligence and gains can be measured upon the completion of the first end-to-end pilot5 integration which is targeted for M30. Thus, the results can only be provided by M36.		

4.5.8. KPI 2.5.8 Number of AI models used/adapted for the pilot

Table 99 KPI 2.5.8 Number of AI models used/adapted for the pilot

KPI ID number .	**KPI 2.5.8**		
KPI Name	**Number of AI models used/adapted for the pilot**		
Description	Number of AI models which has been used in the pilot or specifically adapted to its requirements.		
Motivation	The pilot is addressing a wide range of parameters that need to be optimised, from health-related indicators to energy consumption metrics. Due to this diversity, many AI models need to be evaluated, and through the appropriate configuration and calibration the most suitable models to be identified and used.		
Target value	6 models in total for the AI part of the components Forecasting and Health-Energy		
Prerequisites	aerOS runtime working, IoT Sensors deployed, collect and persistently store data		
aerOS components (task)	T4.2 (Data Fabric)		
Evaluation means	Can be deduced by the number of trained AI models saved in the Pilot 5 database		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*No AI is used*	***4 AI models adapted for forecasting (70%)***	*N/A*
Outcome elaboration (M24)	Four AI models have been employed for the pilot's needs. Specifically, four machine learning-based models are used for the forecasting component. These models are built on the XGBoost Regressor. They take as input information like the time, room of interest, and occupancy, and they predict the future Temperature, Humidity, CO, and CO2 levels for each scenario. Each one of the 4 models, works separately from the others. Specifically: 1) XGBoost Regressor for Temperature forecasting. 2) XGBoost Regressor for Humidity forecasting. 3) XGBoost Regressor for CO2 forecasting. 4) XGBoost Regressor for CO forecasting.		

4.6. Overall pilots engagement

4.6.1. KPI 2.6.1 Validation of aerOS in different use cases (KVI-6.1)

Table 100 KPI 2.6.1 Validation of aerOS in different use cases (KVI-6.1)

KPI ID number .	**KPI 2.6.1**
KPI Name	**Validation of aerOS in different use cases (KVI-6.1)**

Description	Confirm that the aerOS platform has been validated with the committed number of use cases		
Motivation	The consortium has specific commitments as described in the DoA to validate the aerOS platform with the predefined use cases. More may be derived during the project execution and discussions.		
Target value	*>5*		
Prerequisites	Pilot needed per use case must be implemented and running.		
aerOS components (task)	Depending on the use cases, pilots 1-5 (T5.2).		
Evaluation means	Coordination with the pilots to confirm use case examination and validation.		
Measurement period	***Baseline***	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***N/A***	*N/A*
Outcome elaboration (M24)	Pilots are not operational yet. This KPI needs pilots in operational mode exploiting the advantages of the aerOS platform. This is expected to be done closer to the project end.		

4.6.2. KPI 2.6.2 Enable fast-track development of new use cases through external partners (e.g., open call third parties) based on aerOS' Open-Source Software components and tools from O1 (KVI-6.2)

Table 101 KPI 2.6.2 Enable fast-track development of new use cases through external partners (e.g., open call third parties) based on aerOS' Open Source Software components and tools from O1 (KVI-6.2)

KPI ID number .	**KPI 2.6.2**
KPI Name	**Enable fast-track development of new use cases through external partners (e.g., open call third parties) based on aerOS' Open Source Software components and tools from O1 (KVI-6.2)**
Description	The project has commitments for two open calls allowing new partners to join the Consortium and develop applications and/or aerOS components
Motivation	It should be measured that the open calls were successful, and the anticipated number of new use cases has been reached.
Target value	14
Prerequisites	Open calls announced and new use cases selected.

aerOS components (task)	Required components per pilot depending on the new use cases		
Evaluation means	Open calls organized and executed successfully, and evaluation of the new use cases completed and validated.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	*0*	***7***	*N/A*
Outcome elaboration (M24)	First round of Open Calls aimed at funding innovative proposals that will enhance aerOS's objectives framed (mandatorily) within one (out of its five) pilot(s). The project run a detailed first open call from which seven new third party contributions have been selected. A detailed guide for applicants was available, with eligibility and evaluation criteria, as well as template for proposals. The call received 38 proposals that have been sent to external reviewers. 15 surpassed the threshold and proceeded with further consideration by the project committee, to conclude to the 7 selected proposals: HACER (Pilot 4), DAIMON (Pilot 3), EcoQM (Pilot 1), ENERGETIC (Pilot 5), ANEOSP (Pilot 4), IBRTEFC (Pilot 1), GreenAnalyzer (Pilot 2).		

4.6.3. KPI 2.6.3 Identification of new application domains to deploy aerOS architecture (KVI-6.3)

Table 102 KPI 2.6.3 Identification of new application domains to deploy aerOS architecture (KVI-6.3)

KPI ID number .	**KPI 2.6.3**
KPI Name	**Identification of new application domains to deploy aerOS architecture**
Description	Analysis of potential new application domains (out of the ones already tackled by aerOS pilots) where aerOS benefits would be clear.
Motivation	Scalability and uptake potential of aerOS.
Target value	3
Prerequisites	Some pre-requisites identified are: • Architecture is complete (D2.7) • Open Calls (round #1 and #2) are selected and have started validating aerOS' components • T6.4 has performed several Business Analysis tools, revealing new application domains where such a Meta-OS would be of interest.
aerOS components (task)	The intention is to detail here the most relevant aerOS components that would be transferrable to further domains other than the currently covered in aerOS pilots. Up to now M24, the most relevant components that are being required by Open Call projects, and by other initiatives such as EUCEI and the project SAFE-6G are:

	• Self-* tool suite (T3.5) • Orchestration (HLO and LLOs) (T3.3) • Data Fabric and Federation (T4.2) • AAA and cybersecurity around all those (T3.4)		
Evaluation means	A report will be done (included in D5.6) referring to domains coming from Open Calls, other identified, and the potentialities of aerOS adoption in those sectors.		
Measurement period	**Baseline**	**M24 (Deliverable D5.5)**	**M36 (Deliverable D5.6)**
Measured value (% achieved)	N/A	1 (33%)	N/A
Outcome elaboration (M24)	Pre-requisites have not been met yet. However, the project SAFE-6G expressed interest in applying aerOS technologies into different uses cases that rely on 6G field' research. As a matter of fact, aerOS has already performed a technical workshop to SAFE-6G partners, showcasing the capacities of aerOS to support deployments in such scenarios. Therefore, the team is considering a different "application domain" can already be counted into this KPI's evaluation. It is worth mentioning that all active Open Calls (7 projects) cannot be included as new domains, due to their involvement and enclosure within aerOS' already existing pilots.		

5. aerOS impact KPIs

The aerOS impact KPI dimension evaluation is a bit different from the technical and pilot ones. While the previous two ones aim at (i) evaluating tangible software and hardware results and their performance; and (ii) evaluating the outcomes of the application of such technologies into the selected pilots, the impact evaluation aims at gathering those KPIs that are relevant to understand the impact of the project: i.e., towards the community, and towards future exploitation. However, while the two previous dimension KPIs have been thoroughly described and analysed in this report, impact KPIs are carefully monitored and reported in WP6 and its impact deliverables. A list of all the KPIs that belong to this dimension is presented in the next table in this deliverable. For more details, the reader may refer to D6.2 and subsequent WP6 deliverables.

Table 103 List of aerOS Impact KPIs

Field	KPI id	Name	Target	M24(23)
3.1 Communication	KPI.3.1.1	# of Website unique visitors / page views	4000/10000	5,115/20,505
	KPI.3.1.2	# of aerOS posts in social networks/ #of newsletters issued	1000/12	736/7
	KPI.3.1.3	# of aerOS social-media community members across all-sites	1000	1,018
	KPI.3.1.4	# of videos delivered about aerOS technical and global advances / webinars-workshops organised	20 / 6	11/15
	KPI.3.1.5	# of interviews/articles/press releases with external relevant dissemination targets	30	19
	KPI.3.1.6	# of liaison with other projects of the cluster including CSA events	35 actions	40
3.2 Dissemination	KPI.3.2.1	# of scientific papers published in conferences / Q1-Q2 journals	20 / 8	5/11
	KPI.3.2.2	# of activities towards Education institutions (courses, lectures, PhDs)	15	4
	KPI.3.2.3	# of presentations and other activities in events/conferences/fairs by aerOS partners	35	39
	KPI.3.2.4	# of workshops organised / average participants in each workshop	3 / 60	10/20
	KPI.3.2.5	# of PhD and MSc theses started about aerOS	6	10
3.3 Standardisation	KPI.3.3.1	Contributions to standardisation bodies	12	18
	KPI.3.3.2	Exploitation to entry-points into standardisation bodies	25	15
	KPI.3.3.3	aerOS contributions to European pre-normatives	3	2
	KPI.3.3.4	aerOS contributions to data-related clusters and initiatives	10	1

3.4 Exploitation and business models	KPI.3.3.5	# of contributions to relevant data spaces (GAIA-X, IDSA)	10	2
	KPI 3.4.1	Contribution to OSS projects	12	11
	KPI.3.4.2	Business plans for exploitable assets, stakeholders and key alliances identified and contacted	100,00%	0
	KPI 3.4.3	New business lines on aerOS by partners	2	0
	KPI 3.4.4	# of startups adopting aerOS results as technological baseline for business	1	0
	KP 3.4.5	# of tech-transfer contracts signed based on aerOS (from Universities/RTOs)	1	0
	KPI 3.4.6	Private investments in aerOS and related open technologies	10 M€	0
	KPI 3.4.7	Market share in edge-cloud-computing of Europe vs world	32,00%	0

The Exploitation KPIs (KPI 3.4.1 to KPI 3.4.7), which significantly depend on business analysis and the successful outcomes of aerOS technologies, cannot be evaluated yet. Nevertheless, the business analysis is undergoing, and final reports will be submitted at the end of the project.

6. Initial KeVI analysis

Following the methodology described in Section 2.2, a detailed qualitative analysis is provided for each aerOS pilot/use-case. The analysis is drawn around a table including the following:

- **Affected KeVs**, i.e., values important to people and society, directly addressed or indirectly impacted by the specific pilot's use case),
- **Key Value Indicators (KeVIs)**, measurable quantities or requirements that provide estimates of the KeVs),
- **KeV Enablers**, technological advancements needed for fulfilling the KeVs,
- **Key Performance Indicators (KPIs)**, used for measuring the impact of the pilot towards addressing the KeVs).

A discussion follows every table for explanatory purposes.

6.1. Pilot 1 Data driven cognitive production lines

Use case: Data-Driven Cognitive Production Lines

KeV	KeVIs	KeV Enablers	Pilot KPIs
Environmental sustainability (Addressing **SDG#13**: Climate Action)	• Reduce the CO2-emission through the calculation of an optimized production	• Edge intelligent services • Flexible analytics	• CO2 emissions reduction (**KPI#2.1.6**)
Economical sustainability and innovation (Addressing **SDG#9**: Industry, innovation and infrastructure, **SDG#12**: Responsible consumption and production)	• Energy efficiency in manufacturing workflows • Process stability	• Smooth interaction between the quality control intelligence engine and various dimensional instrumentation equipment • Edge intelligent services	• Production process accuracy (**KPI#2.1.1**) • Dimensional quality control productivity (**KPI#2.1.3**) • Reduction of the average production time of the final product

Pilot 1 focuses on the evolution of modular manufacturing systems that requires strategies for mass customization to handle a diverse range of products. This entails creating highly flexible, sustainable (green) digital production lines for low-volume, highly customized manufacturing. Implementing smart rapid response features for self-optimization, reconfiguration, and adaptation of production lines is crucial. Fast data processing is essential for making intelligent, automated, human-centered decisions. aerOS will introduce new autonomous, flexible edge layer services with intelligent orchestration to enhance production autonomy to Level 4, as outlined in EFFRA CF2 pathways. This distributed, edge-powered modular approach will support IoT edge-cloud integration, modular green production, and digital AI/ML workflow orchestration, enabling real-time closed-loop processes, active energy monitoring, and self-adaptive scheduling.

This pilot can benefit both the environment and the economy by enhancing sustainability. Environmentally, it aims to reduce CO2 emissions by optimizing production, with flexible analytics and edge intelligent services limiting emissions. Economically, it promotes sustainability and innovation through process stability and energy efficiency in manufacturing workflows. The smooth interaction between the quality control intelligence engine and dimensional instrumentation equipment, supported by edge intelligent services, helps achieve economic goals. Specifically, production process accuracy can improve, dimensional quality control productivity too, and reduce average production time, enhancing overall workflow efficiency and reducing energy consumption.

6.2. Pilot 2 Containerised edge computing near renewable energy sources

Use case: Containerized Edge Computing near Renewable Energy Sources

KeV	KeVIs	KeV Enablers	Pilot KPIs
Societal sustainability (Addressing **SDG#3**: Good Health and Well-being, **SDG#11**: Sustainable cities & Communities)	• Better public health • Better air quality in neighbouring areas	• Federated edge nodes and a private cloud located directly at renewable energy premises • Connection different smart devices and data sources from wind and PhV farms operated by ELECT	• Carbon awareness share of green energy **(KPI#2.2.5)**
Environmental sustainability (Addressing **SDG#13**: Climate Action)	• Reduced overall energy usage • Enhancement of energy and resource optimization efforts • Decarbonation	• Deployment of edge intelligence services • Analytics for efficiency and environmental footprint	• Carbon awareness share of green energy **(KPI#2.2.5)** • Consumed renewable energy based on decision making process of aerOS **(KPI#2.2.1)**
Economical sustainability and innovation (Addressing **SDG#8**: Decent Work and Economic Growth)	• Reduced system's capital intensity	• Operational cost efficiency • Low-cost scalability and expandability • Dynamic coordination of available resources	• Effectiveness of task distribution through aerOS to nodes (**KPI#2.2.2**) • Scalability of task distribution and management through aerOS (**KPI#2.2.3**)

Pilot 2 addresses the growing need for sustainable solutions in the edge-cloud industry. The energy-intensive nature of traditional data centres makes expanding operations by purchasing or building more space increasingly impractical. Typically, resource-demanding services, such as AI, have been deployed in the cloud. However, decentralizing these services to edge devices, like those supported by aerOS, can enhance performance. By leveraging the capabilities of the aerOS platform, the European cloud industry can benefit from numerous edge nodes, both near and far, located at renewable energy production sites. This approach allows for rapid scaling of operations while maintaining computing and storage capacity through advanced orchestration capabilities.

This pilot benefits society, the environment, and the economy. From a societal sustainability perspective, it promotes good health and well-being, along with the development of smart cities and communities. By utilizing federated edge nodes and a private cloud located at renewable energy sites and connecting various smart devices and data sources from wind and PV farms operated by ELECT, it improves health and air quality in nearby areas. This can increase the carbon awareness share of green energy.

Environmentally, the pilot supports climate action by implementing containerized edge computing near renewable energy sources. This reduces overall energy usage, enhances energy and resource optimization, and contributes to the decarbonization movement. Deployment of edge intelligence services and efficiency analytics help in achieving these goals based on aerOS decision-making processes.

Economically, the pilot provides decent work and economic growth by reducing system capital intensity. It achieves operational cost efficiency, dynamic coordination of available resources, and low-cost scalability and expandability. Task distribution effectiveness through aerOS ensures that most tasks are executed on schedule, with the system capable of handling many tasks per month, demonstrating the scalability of task distribution and management.

6.3. Pilot 3 High performance computing platform for connected and cooperative mobile machinery

Use case: High Performance Computing Platform for Connected and Cooperative Mobile Machinery to improve CO2 footprint

KeV	KeVIs	KeV Enablers	Pilot KPIs
Environmental sustainability (Addressing **SDG#13**: Climate Action)	• Reduce energy consumption • Reduce CO2 and resource wastage	• Tasks implemented at the edge reducing the latency and reaction time • AI and real-time embedded analytics • Increasing operation efficiency of mobile machinery in the field	• CO2 emissions reduction thanks to platooning (**KPI#2.3.3**)
Economical sustainability and innovation (Addressing **SDG#8**: Decent Work and Economic Growth)	• Efficiency of a large-scale production system	• Data autonomy with semantics for a fleet of vehicles • IoT edge and cloud technologies to orchestrate AI/ML-based services	• Swarm of vehicle performance improvement (**KPI#2.3.2**)

Societal sustainability & Digital inclusion (Addressing **SDG#11**: Sustainable cities & Communities)	• Improvement on connectivity capabilities in rural areas	• Use of low-latency networks like 4G or 5G • Increasing communication efficiency due to data preprocessing	• Performance and connectivity capabilities improvement (single vehicle) (**KPI#2.3.1**)
Simplified life (Addressing **SDG#11**: Sustainable cities & Communities)	• Automated safe and secure execution of tasks at the edge node of the vehicles swarm	• Vehicle will be equipped with its own far edge node running aerOS • Vehicles will be interconnected with smart devices and sensors onboard	• Swarm of vehicle performance improvement (**KPI#2.3.2**)

Pilot 3 addresses the rapid digital transformation in agriculture, with a focus on precision farming as a means to reduce inputs and maximize yields and product quality. Digitalization enables integrated control of production machinery while also necessitating interaction with other production systems and information services within the food production and value chains. Edge computing, combined with limited or temporary networks (due to connectivity constraints in rural areas), allows for the deployment of intelligence without requiring permanent cloud connectivity, promoting self-reliance. This approach is crucial for synchronizing and optimizing tractor operations, paving the way for more productive and sustainable farming practices. Current systems, such as connected and cooperative agricultural machinery, face resource limitations in tasks like data access and processing, ensuring data privacy and security, and maintaining cloud continuity. In-vehicle edge nodes (e.g., JD edge), interacting with smart devices, networking components, and the compute continuum, will benefit from the support provided by an IoT-Edge-Cloud continuum solution.

This pilot aims to enhance society, the economy, the environment, and simplify human life. Environmentally, it focuses on reducing energy consumption, CO2 emissions, and resource wastage by implementing tasks at the edge, reducing latency and reaction time, utilizing AI and real-time embedded analytics, and increasing the operational efficiency of mobile machinery.

Economically, the pilot promotes sustainability and innovation by enhancing large-scale production system efficiency. Data autonomy with semantics for vehicle fleets, along with IoT edge and cloud technologies orchestrating AI/ML-based services, can improve vehicle swarm performance.

From a societal perspective, the pilot supports digital inclusion and the creation of sustainable cities and communities. It improves connectivity in rural areas using low-latency networks like 4G or 5G, increasing communication efficiency through data pre-processing. This translates to improved performance and connectivity for individual vehicles.

Lastly, the pilot aims to simplify life in sustainable cities and communities. Automating task execution at the edge node of vehicle swarms enhances safety and security. Vehicles equipped with their own far edge nodes running aerOS, interconnected with smart devices and sensors, can achieve improvement in frame rate performance.

6.4. Pilot 4 Smart edge services for the port continuum

Use case: Smart edge services for the Port Continuum

KeV	KeVIs	KeV Enablers	Pilot KPIs
Personal health and protection from harm (Addressing **SDG#3:** Good Health and Well-being)	• Maintenance team takes faster and better decisions lowering human mistakes and safety risks • Machine comes for maintenance only when its required reducing occurrences and incident risk	• Analytics and AI tools	• Increase on detection of equipment malfunctions (from manual to automatic) **KPI#2.4.2**) • Performance evaluation metrics of regression AI models (MAE/RMSE) for predictive maintenance of CHEs (**KPI#2.4.5**) • Reduction of CHE idle time due to failures **(KPI#2.4.1)**
Economical sustainability and innovation (Addressing **SDG#9**: Industry, Innovation and Infrastructure)	• Reduction of cost and time in decision making • Early detection of issues • Extending the lifespan of industrial assets • Cost efficiency • Higher asset utilization in operations during peak days • Reduction in resource consumption (lubricants, oils etc) & wastes as maintenance occurs only when its necessary • Reduction in unexpected break down should result in reduction in clean ups	• Secure and trusted environment for TOS and CMMS • Analytics and AI tools	• Performance evaluation metrics of classification AI models (accuracy) for damaged seals (**KPI#2.4.6**) • Performance evaluation metrics of classification AI models (F1) for damaged seals (**KPI#2.4.7**)
Simplified life (Addressing **SDG#3**: Good Health and Well-being)	• Container seals autonomously, without requiring human intervention • Reduction in safety hazards as for checking seals personnel needs to be in close proximity to heavy machinery	• Intelligent orchestration of distributed applications • Analysing video streams at the edge • Self-orchestrated IoT edge-cloud continuum	• Increase of number of actual damaged containers (manually reported by staff vs automatic system-reports) (**KPI#2.4.3**)

Pilot 4 aims to enhance cargo operations at EGCTL through smart edge services for the Port Continuum. Currently, Quay and Yard cranes at EGCTL rely on multiple PLC controllers, which are the most accurate source of data on crane status. However, Big Data, AI/ML, and IoT technologies are primarily based on remote servers or cloud platforms, creating a gap between the precise data from PLCs and the KPIs used for analysis and predictions. This gap results in real-time observability challenges and latency issues, hindering terminal efficiency and potentially causing operational disruptions. As physical expansion of terminals is not feasible, improving operational performance necessitates adopting the Industry 4.0 (I4.0) digitalization paradigm. This approach enhances decision-making by improving information availability and presentation. While first-generation IoT architectures cannot support advanced computer vision and predictive maintenance services directly at the edge, aerOS enables the orchestration of smart edge services. This allows maritime companies to react more quickly without relying on high-performance cloud processing.

The integration of smart edge services in port operations not only enhances personal safety and operational efficiency but also contributes to achieving sustainable development goals related to health, innovation, and simplified living conditions within industrial environments. This holistic approach leverages advanced technologies to create a safer, more efficient, and sustainable port environment.

More specifically, personal health and protection from harm can be addressed by the implementation of analytics and AI tools that enables the maintenance team to make faster and more accurate decisions, thereby reducing human errors and safety risks. As a result, a general improvement can be noted in automatic detection of equipment malfunctions, performance metrics of AI models for predictive maintenance, and reduction of idle time due to failures.

From an economical sustainability and innovation perspective, the focus is on reducing costs and time in decision-making processes, extending the lifespan of industrial assets, and optimizing resource consumption. This is supported by a secure environment for terminal operating systems (TOS) and computerized maintenance management systems (CMMS), along with the deployment of analytics and AI tools. These can be achieved by the accuracy and performance in AI models for detecting damaged seals, thereby reducing maintenance costs and resource consumption significantly.

For simplifying life, the emphasis is on autonomous container seal inspections to minimize safety hazards associated with human proximity to heavy machinery. The use of intelligent orchestration of distributed applications and edge computing allows for real-time analysis of video streams, contributing to operational efficiency and safety. The reduction complaints through automated container damage detection highlights the effectiveness of this approach in improving overall safety and operational efficiency.

6.5. Pilot 5 Energy Efficient, health safe and sustainable smart buildings

Use case: Energy Efficient, Health Safe & Sustainable Smart Buildings

KeV	KeVIs	KeV Enablers	Pilot KPIs
Environmental sustainability (Addressing **SDG#13**: Climate Action)	• Reduce energy consumption, the CO2 and resource wastage • Accurate sensor-based environmental perception	• Energy-efficient monitoring sensors • Flexible analytics services and network automation • Smart actuator systems to implement AI-empowered ad-hoc decision-making	• Energy use reduction (**KPI#2.5.1**) • Edge processing performance gains (**KPI#2.5.2**)

Societal sustainability (Addressing **SDG#3**: Good Health and Well-being, **SDG#11**: Sustainable cities & Communities)	• Reduced health incidents in workplaces. • Increased operational efficiency of interventions in workplaces.	• Calculate and monitor the Health Index of a working environment • Flexible service fabric with dynamic network and service orchestration and automation	• Improvement of air quality (**KPI#2.5.7**) • Services directly managed by the FOM (**KPI#2.5.6**)
Economical sustainability and innovation (Addressing **SDG8**: Decent Work and Economic Growth)	• Cost Efficiency of working environments	• Operational cost efficiency. • Low-cost scalability and expandability • Dynamic coordination of available resources	• Services directly managed by the FOM (**KPI#2.5.6**) • Service Creation/Scalability (**KPI#2.5.5**)
Personal health and protection from harm (Addressing **SDG#3**:: Good Health and Well-being, **SDG#13**: Climate Action)	• Reduce energy consumption, the CO2 and resource wastage • Employ health monitoring measures • Ensure employees' preferences & behavioural data protection	• Network and service automation for sensor-data analytics • Secure & trustworthy AI • Resilient & reliable services. • Calculate and monitor the Health Index of a working environment	• Service Availability (**KPI#2.5.4**) • Improvement of air quality (**KPI#2.5.7**)
Privacy and confidentiality (Addressing **SDG #16**: Peace, Justice & Strong Institutions)	• Employ robust authentication and authorization mechanisms to safeguard the entire IoT edge-cloud ecosystem	• Utilization of 5G capabilities, Virtual Network Functions (VNFs for executing specific security and privacy functions) • System E2E privacy and security • Decentralized processing / offloading to devices, edge, etc	• 5G capabilities to execute security & privacy functions (**KPI#2.5.3**) • Edge processing performance gains (**KPI#2.5.2**)

Pilot 5, "Energy Efficient, Health Safe & Sustainable Smart Buildings," addresses the complexities introduced by the COVID-19 pandemic, which has rendered the maximization of workspace occupancy without regard for safety unacceptable. Ensuring proper employee placement, social distancing, and energy efficiency has become increasingly complex. Real-time data processing and decision-making via aerOS can provide autonomous solutions for safe and sustainable workplaces. This use case demonstrates aerOS by optimizing efficiency and safety in Smart Buildings, improving energy efficiency through real-time processing and AI, and utilizing 5G and smart network components to enhance aerOS capabilities. The expected benefits include intelligent edge decision-making and adaptable solutions that integrate diverse data and platforms within the IoT edge-cloud continuum. Based on the above analysis, this pilot demonstrates potential benefits for environmental, economic, and societal sustainability, as well as specific advantages in personal health, and protection from harm, and privacy and confidentiality.

Environmental sustainability is a key aspect of this pilot's contribution. In the context of climate action, reducing CO2 emissions and minimizing energy and resource waste is crucial. Pilot 5 achieves this by using energy-efficient monitoring sensors and flexible analytics and network automations. These are part of a smart, AI-powered system that optimizes office temperature, determines appropriate occupancy levels for employee health and safety, and reduces power consumption for air conditioning and lighting.

In promoting a safe and healthy working environment, the pilot also advances societal sustainability by reducing workplace health incidents through improved air quality and enhanced operational efficiency of interventions. Technically, this is achieved by calculating and monitoring a Health Index of the work environment and employing a flexible service fabric with dynamic network and service orchestration and automation features. Additionally, personal health and protection from harm, a subcategory of societal sustainability, are addressed. This involves reduced CO2 emissions and resource waste, employee health monitoring, and enabling employees to set personal preferences for their workspace while respecting personal data protection. Good health is supported by the automated monitoring of the Health Index and sensor analytics.

Economic sustainability and innovation are logical outcomes of the pilot's approach. A cost-efficient working environment leads to lower operational costs and easier management and monitoring. Pilot 5 proposes a dynamic coordination system to monitor and map available resources, using them in an environmentally friendly, health-conscious, and cost-effective manner. The system offers low-cost scalability and expandability, capable of managing multiple offices and buildings.

Ensuring reliability and security is vital, and the pilot emphasizes privacy and confidentiality. Robust authentication and authorization mechanisms safeguard the entire IoT edge-cloud ecosystem. The utilization of 5G capabilities and Virtual Network Functions (VNFs) for specific security and privacy functions further enhances security. The pilot also employs decentralized processing, edge computing, and an end-to-end (E2E) privacy and security system.

The aerOS project aims to create a sustainable and innovative platform, tailored to various pilots with different use cases that prioritize environmental, economic, and societal respect. This is evident from the detailed analysis provided. Each pilot emphasizes sustainability and enhancing human life, underscoring the importance of a holistic approach in technological innovations. These initiatives are designed with a single goal: to help and improve human lives. Ultimately, the aerOS project exemplifies how technology can be harnessed to address the pressing challenges of our time, balancing the needs of the environment, economy, and society. Its holistic approach ensures that every innovation contributes to a more sustainable and equitable future for all.

7. Conclusions

This document has presented a thorough update of the three-dimensions structured KPIs of the project.

With respect to D5.2, technical KPIs of aerOS have been enhanced by adding prerequisites, evaluation procedures and early results. Although some measurements have been tackled, given the work in WP3 and WP4 is still on-going, only partial results have been obtained. In general, the results are looking good so far, and now that integration among pilots has begun, evaluation of some KPIs that depend on the number of percentage scenarios are closer to be conducted.

With regards to pilot evaluation, the KPIs have been also improved in terms of description, measurement means and outcome elaboration. WP5 in general, and T5.2 in particular, is now at an intense stage of devotion (trials being executed right now), thus most of the KPI results cannot be measured yet. This deliverable paves the way for a final evaluation during the last months of the project, to be documented in D5.6 Impact KPIs are also in place. Some of those (such as communication, dissemination, and standardisation target numbers) have been updated. Others (e.g., exploitation KPIs) are still being tackled under the works of WP6. Therefore, many are still at zero, but business analysis is ongoing, and final updates will be provided in D5.6.Finally, based on a methodology around the UN Sustainable Development Goals (SDGs), the deliverable includes a detailed qualitative analysis per pilot/use-case, identifying related Affected Key Values (KeVs), i.e., values important to people and society, directly addressed or indirectly impacted by the specific pilot's use case), Key Value Indicators (KeVIs), i.e. measurable quantities or requirements that provide estimates of the KeVs), KeV Enablers, i.e., technological advancements needed for fulfilling the KeVs, and Key Performance Indicators (KPIs) used for measuring the impact of the pilots towards addressing the KeVs.

In summary, the document meets the goal of establishing an intermediate control point of the advance on KPIs evaluation. The final part will be assessed and reported during the last year of project.

References

[1] “What societal values will 6G address”, 6G IA White Paper, 2022, url: https://5g-ppp.eu/wp-content/uploads/2022/05/What-societal-values-will-6G-address-White-Paper-v1.0-final.pdf
[2] “The 17 Goals”, United Nations, 2020, url: https://sdgs.un.org/goals

[3] H. Lee, H. Mun and Y. Lee, "Comparing Response Time of Home IoT Devices with or without Cloud," 2020 IEEE International Conference on Consumer Electronics (ICCE), Las Vegas, NV, USA, 2020, pp. 1-6.