TALON

Autonomous and Self-organized Artificial Intelligent Orchestrator for a Greener Industry 4.0

Deliverable

D6.2 Dissemination, Standardisation, Regulation & Business Report

Actual submission date: 02/10/2025

Funded by
the European Union

Project Number:	101070181		
Project Acronym:	TALON		
Project Title:	Autonomous and Self-organized Artificial Intelligent Orchestrator for a Greener Industry 4.0		
Start date:	October 1st, 2022	**Duration:**	36 months
D6.2 Dissemination, Standardisation, Regulation & Business Report			
Work Package:	6		
Lead partner:	Eight Bells Ltd (8BELLS)		
Author(s):	Simon Sarkissian (8BELLS), Stylianos Trevlakis (IC), Angel Ortiz Perez, (EXOS)		
Reviewers:	Lorenzo Gori (ENG), Dorine Matzakou (INTRA)		
Due date:	30 Sep, 2025		
Deliverable Type:	R	**Dissemination Level:**	PU
Version number:	1.0		

Revision History

Version	Date	Author	Description
0.1	10/5/2025	8BELLS, IC	ToC first realease
0.2	20/06/2025	8BELLS, IC	ToC finalisation
0.3	15/07/2025	8BELLS, IC, ALL	1st draft with initial contributions
0.4	10/08/2025	8BELLS, IC, ALL	2nd draft with initial contributions
0.5	15/09/2025	8BELLS, IC, ALL	Content consolidation
0.6	25/09/2025	ENG, INTRA	Document under review
0.7	30/09/2025	8BELLS, IC, ENG, INTRA	Reviewers comments addressed
1.0	02/10/2025	ENG	Ready for submission

Table of Contents

List of figures

List of tables

Definitions and acronyms

8BELLS	*Eight Bells Ltd*
AHP	*Analytical Hierarchy Process*
AI	*Artificial Intelligence*
AR	*Augmented Reality*
AR/VR	*Augmented Reality/Virtual Reality*
B2B	*Business-to-Business*
B2G	*Business-to-Government*
BDVA	*Big Data Value Association*
BMC	*Business Model Canvas*
BU	*Business Unit*
BVLOS	*Beyond Vision Line of Sight*
CE	*Circular Economy*
CEAP	*Circular Economy Action Plan*
CERTH	*Centre for Research & Technology Hellas*
CI	*Consistency Index*
CNC	*Computer Numerical Control*
CNN	*Convolutional Neural Networks*
CR	*Consistency Ratio (in context of AHP)*
CTO	*Chief Technology Officer*
CVRMSE	*Coefficient of Variation of the Root Mean Square Error*
dApps	*Decentralised Applications*
DEP	*Digital Europe Programme*
DICOM	*Digital Imaging and Communications in Medicine*
DPP	*Digital Product Passport*
DT	*Digital Twin*
DUTH	*Democritus University of Thrace*
E2C	*Edge-to-Cloud*
EASA	*European Union Aviation Safety Agency*
EC	*European Commission*
EE	*Energy Efficiency/Efficient*
ENG	*Engineering Ingegneria Informatica S.p.A.*
EPC	*European Patent Convention*
ER	*Exploitable Result*
ESPR	*Ecodesign for Sustainable Product Regulation*
ESS	*Expert Support Services (Horizon Results Booster)*
ETSI	*European Telecommunications Standards Institute*
EU	*European Union*
FL	*Federated Learning*
GAN	*Generative Adversarial Networks*
GDPR	*General Data Protection Regulation*
GPAI	*General-purpose Artificial Intelligence*
HE	*Horizon Europe*
HRB	*Horizon Results Booster*
HRC	*Human-Robot Collaboration*
HW	*Hardware*

I4.0	*Industry 4.0*
I5.0	*Industry 5.0*
IAM	*Image Anonymisation Module*
IC	*InnoCube*
ICT	*Information and Communications Technology*
IEEE	*Institute of Electrical and Electronics Engineers*
INTRA	*Intrasoft International (Netcompany-Intrasoft)*
IoD	*Internet of Drones*
IoT	*Internet of Things*
IP	*Intellectual Property*
IPFS	*InterPlanetary File System*
ITI	*Information Technologies Institute (of CERTH)*
KER	*Key Exploitable Result*
KPI	*Key Performance Indicator*
KU	*Kingston University*
LIME	*Local Interpretable Model-agnostic Explanations*
LOF	*Local Outlier Factor*
LSTM	*Long Short-Term Memory*
LUC	*Light UAS Operator Certificate*
MES	*Manufacturing Execution Systems*
ML	*Machine Learning*
MOM	*Manufacturing Operations Management*
MR	*Mixed Reality*
NG-SDN	*Network Intelligence layer (Next Generation Software-Defined Networking)*
OEM	*Original Equipment Manufacturer*
P2P	*Peer-to-Peer*
PaaS	*Platform as a Service*
PCB	*Printed Circuit Board*
PDRA	*Predefined Risk Assessment*
PDESC	*Portfolio Dissemination & Exploitation Strategy Canvas (HRB Service Bundle)*
PoC	*Proof of Concept*
PPE	*Personal Protective Equipment*
QoS	*Quality of Service*
Q&A	*Questions and Answers*
R&D	*Research and Development*
RAN	*Radio Access Networks*
RED	*Radio Equipment Directive*
REST	*Representational State Transfer*
RI	*Random Index*
RID	*Research & Innovation Development Department (of Netcompany-Intrasoft)*
ROI	*Return on Investment*
RNNs	*Recurrent Neural Networks*
SaaS	*Software as a Service*
SLA	*Service Level Objective/Agreement*
SME	*Small and Medium-sized Enterprise*
SORA	*Specific Operations Risk Assessment*
SPS	*Smart Pricing Simulator*

STS	*Standard Scenario*
SW	*Software*
SWOT	*Strengths Weaknesses Opportunities Threats*
TALON	*Autonomous and Self-organized Artificial Intelligent Orchestrator for a Greener Industry 4.0 (Project Acronym)*
TrL	*Trust Level*
TRL	*Technology Readiness Level*
UAS	*Unmanned Aircraft System*
UATV	*Uncrewed Aerial/Unmanned Autonomous Transport Vehicle*
UBI	*UBITECH*
UC	*Use Case*
UI	*User Interface*
UKRI	*UK Research and Innovation*
UL	*University of Luxembourg*
UPV	*Universitat Politecnica De Valencia*
VR	*Virtual Reality*
WP	*Work Package*
XAI	*Explainable Artificial Intelligence*
XR	*Extended Reality*
ZDM	*Zero Defect Manufacturing*
ZT	*Zero Trust*

Disclaimer

This document has been produced in the context of TALON Project. The TALON project is part of the European Community's Horizon Europe Program for research and development and is as such funded by the European Commission. All information in this document is provided 'as is' and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability. For the avoidance of all doubts, the European Commission has no liability with respect to this document, which is merely representing the authors' view.

Executive Summary

The current deliverable, outlines the comprehensive strategy and complete outcomes of the TALON project in the areas of knowledge dissemination, exploitation, standardisation, regulations, and business planning. The report presents the methods and tools used to maximise the visibility, impact, and long-term sustainability of TALON's results, together with their actual achieved results.

TALON has successfully implemented a broad dissemination strategy through high-impact journal and conference publications, datasets, white papers, workshops, and a strong digital presence across social media and its dedicated website. These actions ensured wide engagement with academic, industrial, and policy stakeholders. Exploitation activities focused on defining a joint exploitation plan, identifying Key Exploitable Results (KERs), and assessing intellectual property and patentability potential. Individual partner exploitation pathways complement the joint strategy, strengthening TALON's innovation and market positioning.

Standardisation efforts targeted alignment with European and international initiatives, including Gaia-X, FIWARE, and trustworthy AI frameworks, while regulatory analysis covered compliance issues and future relevant developments with primarily the EU AI Act, as well as more niche fields like circular economy priorities, and industrial drone applications. Market analysis and business modelling explored the adoption potential of TALON's KERs, with particular emphasis on edge-to-cloud orchestration, explainable AI, and blockchain-secured solutions.

Overall, this document demonstrates TALON's ability to combine scientific excellence with practical pathways for adoption and long-term exploitation, ensuring its contribution to Europe's digital and green transition, within a specifically industrial framework.

1 Introduction

The TALON project, funded under the Horizon Europe programme, aims to deliver an Autonomous and Self-Organised Artificial Intelligent Orchestrator for Industry 4.0 and beyond, supporting Europe's transition towards a greener, more sustainable, and more trustworthy AI, digital economy. The project addresses key industrial challenges by developing a vendor-agnostic, edge-to-cloud orchestration platform that integrates explainable AI, blockchain-enabled security, and energy-aware resource management.

The present Deliverable D6.2, "Dissemination, Standardisation, Regulation & Business Report," has been produced within Work Package 6, led by Eight Bells Ltd. It plays a pivotal role in ensuring that TALON's technological innovations achieve measurable scientific, industrial, and societal impact. The document brings together multiple dimensions of exploitation and communication activities, focusing not only on visibility and engagement but also on long-term sustainability and a forward-looking roadmap that covers technological as well as other relevant aspects.

The first part of the deliverable describes dissemination and communication measures, including scientific publications, participation in leading conferences, the organisation of thematic events, production of communication materials, and management of digital platforms. These activities ensure that TALON results are shared across research, industrial, and policy communities, fostering knowledge transfer and collaboration.

The second part outlines exploitation activities, including the joint exploitation plan, partner-specific strategies, and the identification of TALON's Key Exploitable Results (KERs). This section also covers the patentability assessment conducted through a structured methodology, supporting intellectual property protection and maximising market potential.

Standardisation activities are presented in the following section, highlighting TALON's contributions to relevant European initiatives and its alignment with trustworthy AI and Industry 5.0 frameworks. These efforts reinforce TALON's role in shaping interoperable and sustainable industrial ecosystems.

Finally, the deliverable provides an overview of regulatory aspects relevant to TALON's application domains, with emphasis on compliance with the EU AI Act, circular economy regulations, and industrial drone frameworks. Together with the business and market analysis, these insights provide a holistic view of TALON's positioning in the European innovation landscape.

In summary, D6.2 demonstrates how TALON integrates research excellence, exploitation pathways, regulation aspects, and standardisation efforts to secure long-term value creation and alignment with European priorities in digitalization, AI and sustainability.

2 Dissemination and Communication Activities

The TALON project has carried out an integrated dissemination strategy, which includes journal publications, conferences, white papers, datasets, participation and organization of events, and targeted communication materials, blog posts, and newsletters. This holistic approach guarantees that project outcomes are accessible to all relevant community stakeholders, including scholars, practitioners, policymakers, and the broader society. The next sections outline the different dissemination activities with their results.

2.1 Journal Publications

For TALON, academic journal publications (Table 1) serve as the backbone of the scientific dissemination strategy since they deliver peer-reviewed verification of research results and guarantee perpetual discoverability of the findings. The consortium has successfully published research across high-impact journals covering topics from zero-defect manufacturing to advanced AI techniques, enhancing the scholarly corpus in trustworthy AI and Industry 4.0 applications.

Table 1: TALON's journal publications

No.	Citation
1	Sousa J, Nazarenko A, Grunewald C, Psarommatis F, Fraile F, Meyer O and Sarraipa J, "Zero-defect manufacturing terminology standardization: Definition, improvement, and harmonization.", Front. Manuf. Technol., 2022.
2	F. Psarommatis, F. Fraile, and F. Ameri, "Zero defect manufacturing ontology: A preliminary version based on standardized terms," Computers in Industry, vol. 145, p. 103832, Feb. 2023.
3	V. Azamfirei, F. Psarommatis, and Y. Lagrossen, "Application of automation for in-line quality inspection, a zero-defect manufacturing approach," Journal of Manufacturing Systems, vol. 67, p. 1-22, Apr. 2023.
4	F. Psarommatis, G. May, and V. Azamfirei, "Envisioning maintenance 5.0: Insights from a systematic literature review of Industry 4.0 and a proposed framework," Journal of Manufacturing Systems, vol. 68, pp. 376–399, Jun. 2023.
5	V. Patsias, P. Amanatidis, D. Karampatzakis, T. Lagkas, K. Michalakopoulou, and A. Nikitas, "Task Allocation Methods and Optimization Techniques in Edge Computing: A Systematic Review of the Literature," Future Internet, MDPI, vol. 15, no. 8, p. 254, Jul. 2023.
6	S. Trevlakis, A.-A. A. Boulogeorgos, D. Pliatsios, J. Querol, K. Ntontin, P. Sarigiannidis, S. Chatzinotas, and M. D. Renzo, "Localization as a key enabler of 6G Wireless Systems: A comprehensive survey and an outlook," in IEEE Open Journal of the Communications Society, Oct. 2023.
7	Fraile, F., Psarommatis, F., Alarcón, F., Joan, J.A, Methodological Framework for Designing Personalised Training Programs to Support Personnel Upskilling in Industry 5.0, Computers 2023, 12, 224.
8	Voulgaridis, K., Lagkas, T., Angelopoulos, C.M., Boulogeorgos A.-A. A., Argyriou, V., & Sarigiannidis, P., Digital product passports as enablers of digital circular economy: a framework based on technological perspective. Telecommun Syst 85, 699–715 (2024).
9	Lorente-Leyva, L. L., Alemany, M. M. E., & Peluffo-Ordóñez D. H., A Conceptual Framework for the Operations Planning of the Textile Supply Chains: insights for Sustainable and Smart Planning in Uncertain and Dynamic Contexts. Computers & Industrial Engineering, Volume 187, 109824, 2024.
10	Amanatidis, P., Karampatzakis, D., Michailidis, G., Lagkas, T., Iosifidis, G., Adaptive Reverse Task Offloading in Edge Computing for Ai Processes, May 2024.
11	O. Aouedi, T.-H. Vu, A. Sacco, D. C. Nguyen, K. Piamrat, G. Marchetto, and Q.-V. Pham, "A survey on intelligent internet of things: Applications, security, privacy, and future directions," IEEE Communications Surveys & Tutorials, pp. 1–1, 2024.
12	O. Aouedi, A. Sacco, D. C. Nguyen, LU Khan, and M. Guizani, "Federated Learning for Human Activity Recognition: Overview, Advances, and Challenges," IEEE Open Journal of the Communications Society.
13	V. Li, I. Siniosoglou, T. Karamitsou, A. Lytos, I. D. Moscholios, S. K. Goudos, J. S. Banerjee, P. Sarigiannidis, and V. Argyriou, "Enhancing 3D object detection in

	autonomous vehicles based on synthetic virtual environment analysis," Image and Vision Computing, vol. 154, p. 105385, 2025.
14	P. Amanatidis, G. Michailidis, D. Karampatzakis, V. Kalenteridis, G. Iosifidis and T. Lagkas, "Multi-Objective Reverse Offloading in Edge Computing for AI Tasks," IEEE Open Journal of the Communications Society, vol. 6, pp. 2474-2485, 2025.

2.2 Conferences Publications

Conference publications (Table 2) are essential for showcasing new research, interacting with peers, and obtaining real-time feedback from the scientific community. TALON targets top-tier international conferences that feature renowned researchers, industry leaders and practitioners to maximize strategic exchange and partnership within the AI and Industry 4.0 ecosystems.

Table 2: TALON's conference publications

No.	Citation
1	Nagy, G. Amponis, K. Kyranou, T. Lagkas, A. A. Boulogeorgos, P. Sarigiannidis, V. Argyriou, "AI-Powered Interfaces for Extended Reality to Support Remote Maintenance", CLEF: Cloud, Edge, and Fog for Smart Industries, May, 2023.
2	Vladislav Li, Barbara Villarini, Jean-Christophe Nebel, Thomas Lagkas, Panagiotis Sarigiannidis and Vasileios Argyriou, Evaluation of Environmental Conditions on Object Detection using Oriented Bounding Boxes for AR Applications, IEEE DCOSS-IoT, "5th International Workshop on IoT Applications and Industry 5.0", June, 2023.
3	Akos Nagy, Thomas Lagkas, Panagiotis Sarigiannidis and Vasileios Argyriou, Evaluation of AI-Supported Input Methods in Augmented Reality Environment, IEEE DCOSS-IoT, "1st International Workshop on Next Generation IoT and AI systems for Trusted, Human-Centered Intelligence", June, 2023.
4	Psarommatis Foivos, May Gökan, Azamfirei Victor, Magnanini Maria Chiara, Powell Dary, "A readiness level assessment framework for Zero Defect Manufacturing (ZDM)", 32nd FAIM International Conference, June, 2023.
5	Azamfirei Victor, Psarommatis Foivos, Lagrosen Yvonne, "Human factors in the design of advanced quality inspection systems in the era of Zero-Defect Manufacturing", 32nd FAIM International Conference, June, 2023.
6	Vladislav Li, Barbara Villarini, Jean-Christophe Nebel and Vasileios Argyriou, A Modular Deep Learning Framework for Scene Understanding in Augmented Reality Applications, IEEE IAICT, "The IEEE International Conference on Industry 4.0, Artificial Intelligence, and Communications Technology", July, 2023.
7	L. Mitsiou, S. Trevlakis, A. Tsiolas, D. J. Vergados, A. Michalas and A. -A. A. Boulogeorgos, "Can graph neural network-based detection mitigate the impact of hardware imperfections?," 2023 International Balkan Conference on Communications and Networking (BalkanCom), Istanbul, Turkiye, 2023, pp. 1-5.
8	V. Christofas, P. Amanatidis, D. Karampatzakis, T. Lagkas, S. K. Goudos, K. E. Psannis, and P. Sarigiannidis, "Comparative Evaluation between Accelerated RISC-V and ARM AI Inference Machines," 6th World Symposium on Communication Engineering (WSCE 2023), Thessaloniki, Greece.
9	I. Siniosoglou, D. Asimopoulos, V. Argyriou, T. Lagkas, A. Lytos, I. D. Moscholios, S. K. Goudos and P. Sarigiannidis, "Enhancing Text Anonymisation: A Study on CRF, LSTM, and ELMo for Advanced Entity Recognition", 7th PAnhellenic Conference on Electronics and Telecommunications (PACET 2024)
10	D. Asimopoulos, I. Siniosoglou, V. Argyriou, S. K. Goudos, K. E. Psannis, N. Karditsioti, T. Saoulidis and P. Sarigiannidis, "Evaluating the Efficacy of AI Techniques in Textual Anonymization: A Comparative Study", BalkanCom 2024 Seventh International Balkan Conference on Communications and Networking – Open-RAN for Intelligent B5G and 6G Wireless Networks: Emerging Trends and Recent Developments.
11	Nagy, Y.Spyridis, G. Mills, V. Argyriou, "User Experience Evaluation of AR Assisted Industrial Maintenance and Support", ICVR: International Conference on Virtual Reality.
12	D. Asimopoulos, I. Siniosoglou, V. Argyriou, T. Karamitsou, E. Fountoukidis, S. K. Goudos, I. D. Moscholios, K. E. Psannis and P. Sarigiannidis, "Benchmarking Advanced Text Anonymisation Methods: A Comparative Study on Novel and Traditional

	Approaches", 13th International Conference on Modern Circuits and Systems Technologies (MOCAST).
13	G. Tsoumplekas, V. Li, I. Siniosoglou, V. Argyriou, S. K. Goudos, I. D. Moscholios, P. Radoglou-Grammatikis and P. Sarigiannidis, "Evaluating the Energy Efficiency of Few-Shot Learning for Object Detection in Volatile Industrial Data", 3rd Real-time And intelliGent Edge computing workshop (RAGE 2024).
14	G. Andronikidis, G. Niotis, C. Eleftheriadis, K. Kyranou, S. Nikoletseas, and P. Sarigiannidis, "Optimizing Federated Learning through Lightweight and Scalable Blockchain", 6th International Workshop on IoT Applications and Industry 5.0 (IoTI5 2024).
15	N. Ntampakis, K. Diamantaras, V. Argyriou and P. Sarigianndis, "Fusion Grad-CAM: A Methodology for Generating Unified Attention Maps in Majority Voting Classifiers", 6th IEEE International Image Processing, Applications and Systems Conference (IPAS 2025).
16	S. E. Trevlakis, M. Belesioti, H. Koumaras, A. -A. A. Boulogeorgos, I. Chochliouros, and T. Tsiftsis, "An Innovative Architectural Blueprint Towards Sustainable 6G Systems", 2024 IEEE 29th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD).
17	Theodorou, G., Karagiorgou, S., Fulignoli, A., Magri, R. (2024). On Explaining and Reasoning About Optical Fiber Link Problems. In: Longo, L., Lapuschkin, S., Seifert, C. (eds) Explainable Artificial Intelligence. xAI 2024. Communications in Computer and Information Science, vol 2154. Springer, Cham.
18	O. Aouedi, G. Jajoo and K. Piamrat, "METALS: seMi-supervised fEderaTed Active Learning for intrusion detection Systems," in 2024 IEEE Symposium on Computers and Communications (ISCC), Paris, France, 2024, pp. 1-6.
19	V. Li, G. Tsoumplekas, I. Siniosoglou, V. Argyriou, A. Lytos, E. Fountoukidis, and P. Sarigiannidis, "A closer look at data augmentation strategies for finetuning-based low/few-shot object detection," arXiv preprint arXiv:2408.10940, 2024.
20	G. Theodorou, S. Karagiorgou and C. Kotronis, "On Energy-aware and Verifiable Benchmarking of Big Data Processing targeting AI Pipelines," 2024 IEEE International Conference on Big Data (BigData), Washington, DC, USA, 2024, pp. 3788-3798.
21	Anastasiou, T., Pastellas, I., & Karagiorgou, S. (2024, December). Adversarial Explanations for Informed Civilian and Environmental Protection. In 2024 IEEE International Conference on Big Data (BigData) (pp. 2672-2681). IEEE.
22	F. Paparounas, V. Christofas, P. Amanatiadis, D. Karampatzakis, and T. Lagkas, "Knowledge distillation compression evaluation for Edge Class Devices", 28th Pan-Hellenic Conference on Informatics (PCI), 2024.
23	G. Theodorou, S. Karagiorgou and C. Kotronis, "On Energy-aware and Verifiable Benchmarking of Big Data Processing targeting AI Pipelines," 2024 IEEE International Conference on Big Data (BigData), Washington, DC, USA, 2024, pp. 3788-3798.
24	T. Anastasiou, I. Pastellas and S. Karagiorgou, "Adversarial Explanations for Informed Civilian and Environmental Protection," 2024 IEEE International Conference on Big Data (BigData), Washington, DC, USA, 2024, pp. 2672-2681.
25	Ilias Crysovergis, Stylianos E. Trevlakis, Dimitris Kleitsas, Alexandros-Apostolos A Boulogeorgos, Theodoros A. Tsiftsis, and Dusit Niyato, "A Digital Twin Based Reconfigurable Intelligent Surface Phase Adaptation Using Spiking Reinforcement Learning Policy Optimization", IEEE International Conference on Machine Learning for Communication and Networking, May 2025.
26	P. Amanatidis, E. Vagiotas, A. Stamopoulos, G. Michailidis, D. Karampatzakis, P. Sarigiannidis, and T. Lagkas, "2nd International Workshop on Next Generation IoT and AI systems for Trusted, Human-Centered Intelligence", June 2025.
27	Pastellas, S. Karagiorgou, and M. Konidi, "Adaptive Policy-Driven Network Intelligence for Edge-to-Cloud Continuum," in IEEE 31st International Conference on Engineering, Technology and Innovation, 2024.
28	M. Boluda-Prieto, F. Fraile, F. Alarcón, M. Terol, and M. A. Mateo-Casali, "Enhancing Network Resilience Through Automated Rule-Based Monitoring, Alerting, and Proactive Mitigation," in IEEE 31st International Conference on Engineering, Technology and Innovation, 2024.

29	V. Li, I. Siniosoglou, P. Sarigiannidis, and V. Argyriou, "Enhancing Manufacturing Training Through VR Simulations," in IEEE 31st International Conference on Engineering, Technology and Innovation, 2024.
30	Petropoulos, G. Apostolou, G. Tsoumplekas, G. Tziolas, N. Ntampakis, I. Siniosoglou, V. Argyriou, P. Sarigiannidis, I. Gialampoukidis, and S. Vrochidis, "Enhancing Safety in Industry 5.0: Human-Computer Collaboration Benefits through a Dataset of Protective Equipment Detection," in IEEE 31st International Conference on Engineering, Technology and Innovation, 2024.
31	S. Comella, D. Milazzo, S. Cipolla, O. Aouedi, and S. Bonura, "New Edge-to-Cloud Orchestrator for Intelligent Task Allocation and Efficiency Optimisation," in IEEE 31st International Conference on Engineering, Technology and Innovation, 2024.
32	R. S. Velastegui Hernández, R. Poler, and M. DíazMadroñero, "AgentBlock: Infrastructure for integrating blockchain and multi-agent robotic systems for optimising industrial production and logistics," in IEEE 31st International Conference on Engineering, Technology and Innovation, 2024.
33	Boluda-Prieto, Maria; Mateo-Casalí, Miguel Ángel; Fraile Gil, Francisco; Alarcón Valero, Faustino (2025). Hybrid Architecture Implementing Self-Healing Mechanisms for Distributed AI Network Management in Industry. EN XXIX Congreso de Ingeniería de Organización (CIO 2025). 19th International Conference on Industrial Engineering and Industrial Management (ICIEIM 2025). Sevilla, España.
34	Boluda-Prieto, Maria; Fraile Gil, Francisco; Esteso, Ana; Alemany Díaz, María Del Mar (2025). Applicability of Few-Shot Learning in Tool Wear Prediction. EN XXIX Congreso de Ingeniería de Organización (CIO 2025). 19th International Conference on Industrial Engineering and Industrial Management (ICIEIM 2025). Sevilla, España.

2.3 Datasets

Data generation and sharing are essential for advancing reproducible research while encouraging broader collaboration in the scientific community. TALON strives to compile and disseminate datasets (Table 3) that facilitate the development and evaluation of reliable AI systems in industrial contexts, thus aiding in the preparation of datasets that can be used for future research.

Table 3: Datasets generated by TALON

No.	Description	Type	Date
1	CERTH Pilot 4 image dataset, captured using drone: Warehouse surveillance dataset for object detection, annotated with two PPE classes; helmet and vest. Images contain workers in varied poses and lighting conditions, with bounding box labels suitable for training safety compliance models	Internal	23/07/2025
2	CERTH dataset containing synthetic (augmented numerical) data from Nakamura FACTOR machine (e.g., feed-rate override values, servo motor loads across multiple axes)	Internal	04/07/2025
3	FACTOR dataset containing information about the manufacturing in the factory. The data are: Spindles speed, Feed rates, Spindle motors load, Servo load currents, Cutting time, Temperatures, etc.	Internal	03-06/2024
4	TEI dataset contains time series of received optical power during four types of fibre faults: Stress, Shutdown, Connector, and SFPConnector. Data was collected in a lab, inducing faults and recording power every 10 ms with a custom SFP transceiver and software.	Internal	05/06/2024
5	K. Voulgaridis, D. Karampatzakis, P. Sarigiannidis, and T. Lagkas, 'Shimmer DLT Interaction Metrics Datasets'. Zenodo, August 19, 2025. doi: 10.5281/zenodo.16901311.	Public	19/08/2025

2.4 White Papers

White papers act as a conduit between scholarly work and its practical counterpart by offering technical documentation and policy in one coherent document. TALON focuses on white papers (Table 4) concerning standardization activities and patentability assessments which are relevant and useful to practitioners of the industry, officials and researchers who seek to apply reliable AI systems in their work.

Table 4: TALON's white papers

No.	Citation
1	Sarkissian, S., & Trevlakis, S. E. (2025). TALON's Standardisation Activities. Zenodo. https://doi.org/10.5281/zenodo.15687685
2	Sarkissian, S., Trevlakis, S. E., & Lagkas, T. (2025). TALON's Patentability Assessment. Zenodo. https://doi.org/10.5281/zenodo.15687441

2.5 Participation in events

TALON took advantage of industry events and conferences (Table 5) as opportunities to engage with key stakeholders, showcase research outcomes, and build strategic partnerships. Through these activities, TALON enables the transfer of knowledge and enhances the visibility of the projects, contributing to the AI and Industry 4.0 communities by sharing experiences and best practices.

Table 5: Participations in events by TALON partners

No.	Description	Type	Lead partner	Date
1	Adra-e and AI4Europe Coordination and Support Actions Event "Paving the way towards the next generation of R&I excellence in AI, Data and Robotics"	Webinar	ENG	October 2022
2	International Hybrid Event on Energy Crisis and Cybersecurity	Hybrid event	IC	December 2022
3	DataWeek23 "Data/AI and the Cloud-to-Edge continuum"	Conference	ENG	July 2023
4	ITADATA23	Conference	ENG	September 2023
5	NexusForum2023	Forum	ENG	October 2023
6	Infocom2023	Exhibition	8BELLS	December 2023
7	DW24 - Session "Building the Future of Decision-Making: Human-AI Collaboration, Trust and Compliance in the Age of AI"	Conference	ENG	December 2024
8	Adra-e and AIoD physical event "Future-Ready: On Demand Solutions with AI, Data, and Robotics"	Physical event	IC	February 2025
9	2025 IEEE International Conference on Engineering, Technology, and Innovation (ICE/ITMC)	Physical event	UPV	June 2025
10	21st IEEE Annual International Conference on Distributed Computing in Smart Systems and the Internet of Things (IEEE DCOSS-IoT 2025)	Physical event	DUTH	June 2025

2.6 Organization of events

TALON organized the following workshops (Table 6), webinars, and thematic events in the field of trustworthy AI. With the purpose of knowledge sharing and community building, these events gather researchers, industry professionals, policymakers, and other relevant experts who address the challenges of collaboration towards the growth of Industry 4.0 technologies.

Table 6: Events organized by TALON partners

No.	Description	Type	Lead partner	Date
1	Establishing the next level of "intelligence" and autonomy Clustering Workshop Event	Workshop	UBI, ENG	March 2023
2	5th International Workshop on IoT Applications and Industry 5.0 (IoTI5 2023)	Workshop	DUTH	June 2023
3	Empowering AI for greener and more secure Industrial Operations	Workshop	ENG	March 2024
4	6th International Workshop on IoT Applications and Industry 5.0 (IoTI5 2024)	Workshop	DUTH	May 2024
5	Talon Tech Talk @ Innovation Day - Ericsson Italy	Teck Talk	TEI	November 2024
6	AI You Can Trust: Standards, Ethics, and Innovation	Webinar	INTRA	April 2025
7	SS04 - Autonomous and Self-organized Artificial Intelligent Orchestrator for a Greener Industry 4.0	Special session	UPV	June 2025
8	7th International Workshop on IoT Applications and Industry 5.0 (IoTI5 2025)	Workshop	DUTH	June 2025

2.7 Blog Posts

TALON's blog posts (Table 7) focus on project updates, research findings, and alternative viewpoints from industry leaders which helps capture the attention of people who do not belong in academic publications, thereby nurturing their community. Through blog posts, technical ideas have been published as simplified narratives to appeal to multiple users.

Table 7: TALON's blog posts

No.	Description	Date
1	Kick-of-Meeting	October 2022
2	TALON participation in Adra-e and AI4Europe Coordination and Support Actions Event "Paving the way towards the next generation of R&I excellence in AI, Data and Robotics"	November 2022
3	TALON participation in the international hybrid event on Energy Crisis and Cybersecurity	December 2022
4	TALON participation in Data Week 2023	June 2023
5	TALON participation in the NexusForum2023	September 2023
6	Nine EU projects collaborate for Trustworthy AI across Europe	January 2024
7	Sharing the TALON Experience: The Top Five Integration Challenges of Human-Centric AI Systems	April 2024
8	TALON Project: Pioneering the Future of Industry 5.0 Through Patentable AI Innovation	June 2025
9	Paving the Way for Industry 5.0: How the TALON Project is Shaping AI Standardization	June 2025
10	Building AI You Can Trust: Key Lessons from the TALON Project Webinar	June 2025

2.8 Communication Materials

The communication materials for TALON (Table 8) create a visual identity while branding and enhancing the quality of a presentation. In turn, brand recognition and high global standards are ensured while retaining the professional level of documents in project communications, something that reflects the research conducted.

Table 8: TALON's communication materials

No.	Description	Type	Date
1	TALON logo/branding	Digital identity	October 2023
2	TALON Flyer	Promotional material	October 2023
3	TALON Roll-up banner	Promotional material	October 2023
4	TALON Communication toolkit	Promotional material	October 2023
5	TALON Templates (meeting minutes, deliverables, presentations)	Consistent project branding	October 2023
6	TALON Social media profiles	Digital identity	October 2023
7	TALON - Horizon Europe	Video	January 2023
8	AI You Can Trust Standards, Ethics, and Innovation	Video	May 2025

2.9 TALON's Website and Social Media Accounts

Digitally focused interfaces represent the primary points of interaction between TALON and its stakeholder communities. They offer convenient portals as a one-stop-shop for relevant project information, research findings, and active undertakings. The project's engagements digital presence strategy focuses on a multitude of platforms to ensure higher access and better interaction while providing uniform communication across all platforms.

2.9.1 Website

The TALON web page (https://talon-project.eu/) remains the primary point of access for all information concerning the project, thus enabling easy retrieval of research outcomes, publications, events, and contacts. Given the varying background of visitors, the website caters to the basic needs of academic researchers who may require precise technical details to industry practitioners who are interested in the practical implications of the project results (Figure 1).

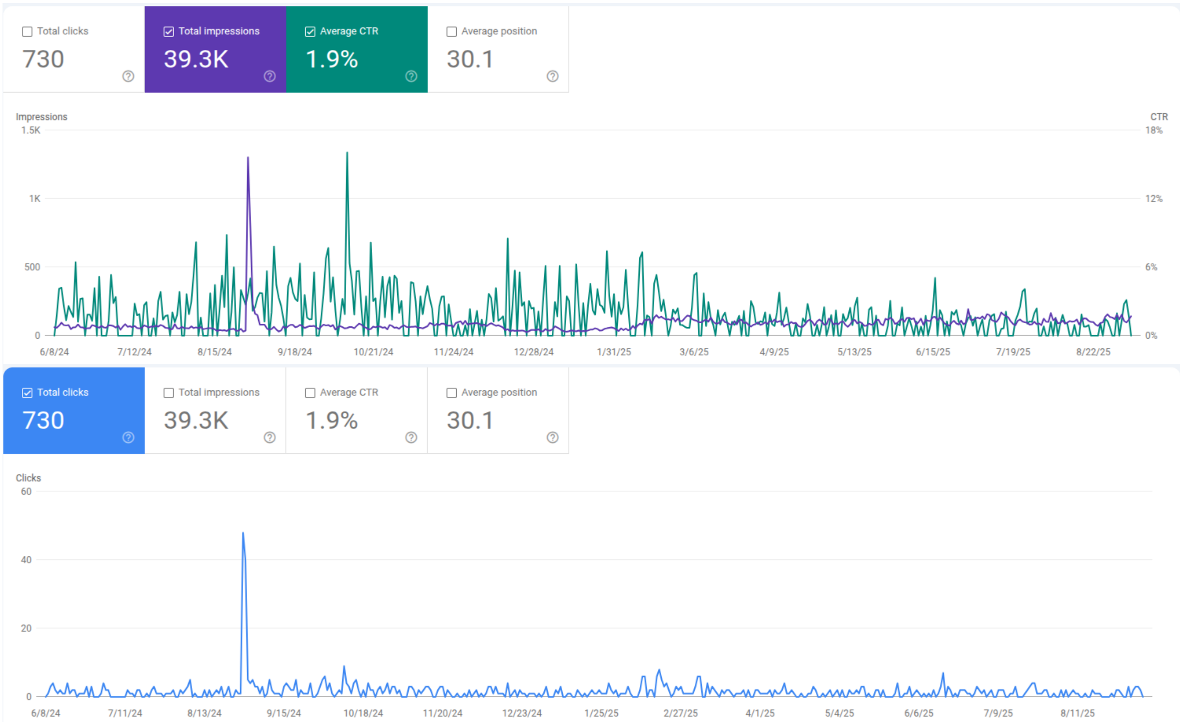

Figure 1: TALON Website Analytics

The TALON project website demonstrates strong brand recognition for its core terms within its niche, driving a significant portion of its traffic. Engagement from key European Union countries is high, indicating successful targeting of a primary audience. However, the data reveals a significant opportunity to expand reach by targeting a wider array of non-branded, thematic search queries related to AI integration, Industry 5.0, and trustworthy AI. Technical SEO and content depth for middle-of-funnel queries need strengthening to improve overall visibility and user acquisition.

2.9.1.1 Query Performance (The "What" People Search For)

The strongest-performing queries demonstrate tremendous strength in branded searches. The main force for these items is "talon project," generating 178 clicks with a tremendous 32.6% CTR; "project talon," generating 92 clicks; both again have excellent click-through-rates and we have estimated that they have approximately top positions of 1 and 3, respectively. Overall performance is excellent, and this demonstrates strong project recognition in the marketplace.

The site does have a zero-click problem however since it has a long tail of over 150 queries which were shown and had impressions but had zero clicks. These problematic queries cover thematic searches, like "integration challenges in ai," "ai human integration," "iot 5.0," and "trustworthy ai project"; these also include brand related searches, such as "talon software," "talon browser," and "talon official website." The conclusion is straight forward, while the site is being shown in relevant industry topics, it is clearly not interesting enough in its search results in its title, meta description, or URL to compel clicks. Also, the average position for these terms is really low, often more than position 40; which would clearly be an indicator it isn't ranking well for this valuable term.

2.9.1.2 Page Performance (The "Where" People Land)

The homepage was, at https://talon-project.eu/ was clearly dominant in total traffic, with 630 clicks from 32.2k impressions. Expectations of traffic are in order, but the homepage high reliance is also worrisome. Many of the top pages are experiencing excellent engagement trends, which shows that the content strategy is working well. The consortium page drove excellent engagement with 21 clicks and a 7.51 average position, which demonstrates that there is an audience that genuinely wants to know who is involved with the project. The contact page is getting a very high 2.52% CTR, indicating that people who accidentally visit these pages are likely to use these pages. The top blog content is driven by the post "Integration challenges," which had 23 clicks and indicates there is a real want/need consumable content of this type.

There are substantial opportunities from many of the key pages like objectives, news, and blog, that are driving many impressions and continuously low CTR, even though they are branded content. They may need to consider their search snippets, including title tags and meta descriptions being tighter and more inviting to searchers.

2.9.1.3 Geographic Performance (The "Who" is Searching)

Traffic is showing engagement from its core EU, with high concentrations across EU countries, which is the optimal positioning for a Horizon Europe project. Greece has the highest engagement performance at 26.72% CTR which indicates its traffic is highly relevant and engaged to have such a phenomenal CTR. Italy, Spain, Cyprus, Luxembourg, and Poland all show very strong CTR over 3%, which is significantly higher than the site percentages, demonstrating great targeting success.

The site has detectable global reach, but low engagement outside of the EU. The United States and India produce high impression volume at 9700 vs 1900 impressions, which corresponds to very low CTR at 0.76 percent and 1.47 percent respectively. The data shows evidence that the site seems relevant enough to be served to audiences in the U.S. and India, but the messaging or content is likely not quite right for global or non-EU users

The analysis concludes that there is excellent focus on the target audience within the EU. The data does not indicate that there is any need to stop focusing on this area. The data almost suggests the opposite: to double down on this area of focus and regional success.

2.9.1.4 Device Performance (The "How" People Search)

The traffic distribution shows a strong preference to desktop usage. From the total traffic, 74 percent is desktop (540 clicks), while approximately 26 percent is mobile (189 clicks). The performance reflective of click through rates is relatively the same at 1.9 percent for desktop and approximately 1.7 percent for mobile.

This traffic distribution gives relevant indications of the composition of the audience in terms of their search behaviour. Researchers, professionals in the industry, and policy makers overwhelmingly use desktop computers to search when they look for specialized information like this. This pattern confirms the professional context of the searches and validates assumptions about the audience segment.

The recommendation looks at optimizing the key reports, especially PDFs, and the overall website experience is perfectly designed for desktops. This will align with what users prefer, and hopefully encourage engagement in the user experience,

2.9.1.5 Performance Over Time (The "When")

The data suggests a significant launch event occurred on August 29-30, 2024, which caused a spike of traffic on those dates. The data shows that on August 30, 2024, there were 40 clicks of the 798 impressions, and 48 clicks on August 29, 2024, which is the highest performance between the two dates, and the highest performance throughout the entire time period covered (Figure 2).

The analysis supports that there was an exceedingly successful campaign or launch event, or publication release that occurred at this time frame. This represents a significant date that we should investigate closely to determine what we can learn to replicate a similar event of this magnitude. Clear correlation between social media activity and website engagement, with three major peaks demonstrating direct conversion from social presence to project interest.

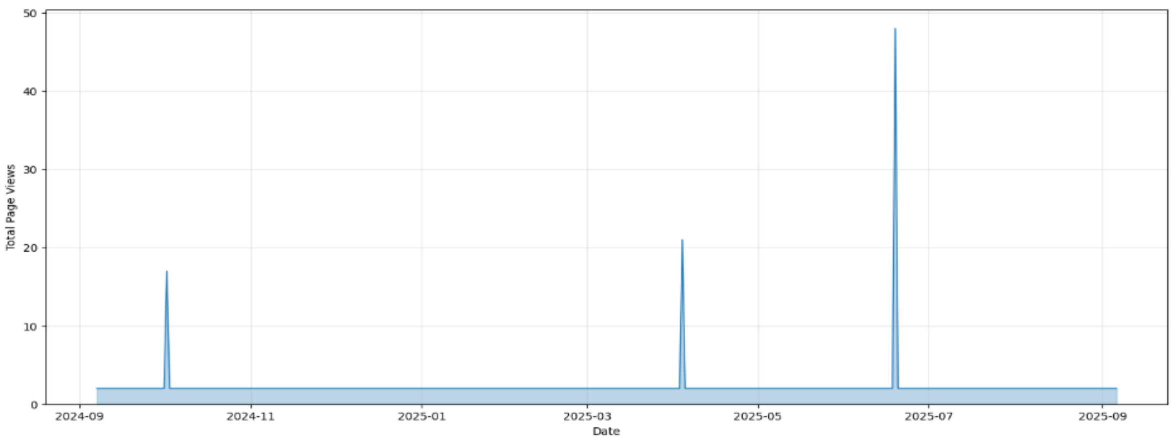

Figure 2: Website Traffic Spikes

2.9.2 TALON's Social Media

Our social media plan (Table 9) was organized around three key pillars and has made a noticeable effect against several goal-led metrics throughout each of the platform's timelines:

Table 9: TALON's Social Media Strategy

Pillars	LinkedIn	X
Informing - Knowledge Sharing Success	• Precision Targeting Achieved: 85% of our audience consists of researchers, engineers, and education professionals • Geographic Alignment: Strong penetration in Mediterranean countries (Greece, Spain, Italy) with expansion to major European hubs.	• Research Dissemination: 47 technical posts sharing cutting-edge research findings. • Academic Reach: Direct citations of 15+ peer-reviewed publications. • Technical Community: Consistent engagement from AI/ML research community.

	• Content Relevance: High engagement rates correlating directly with project milestone announcements.	• Real-time Updates: Live demonstrations of autonomous systems and AI orchestration.
Sharing - Increased Transparency	• Authentic Growth: 100% organic follower acquisition with no paid campaigns, demonstrating genuine interest. • Real-time Updates: Traffic spikes directly correlate with content releases, proving effective information dissemination. • Professional Reach: 85% desktop traffic indicates serious, research-focused engagement.	• Open Research: Direct links to publications, deliverables, and technical documentation • Project Progress: Regular updates on pilot implementations and system demonstrations • Community Engagement: Interactive content explaining complex AI concepts • Zero Paid Promotion: 100% organic reach maintaining authentic research community engagement
Building - Awareness & Impact Amplification Engaged	• Community Development: LinkedIn built from zero to 95+ highly relevant professional followers. • Technical Reach: X platform reaching specialized AI/research communities with technical content. • Conversion Success: Social media effectively funnels interested parties to project website. • Sustained Engagement: Consistent growth pattern over 12 months with notable acceleration periods	

TALON's multi-platform social media strategy has exceeded expectations in building highly engaged, professionally relevant communities while maintaining complete alignment with Horizon Europe's communication objectives. The achievement of 95+ organic LinkedIn followers combined with consistent X platform engagement representing specialized technical audiences demonstrates not just numerical success, but the creation of genuine knowledge-sharing communities that amplify our research impact across European professional and academic networks.

2.9.2.1 LinkedIn

LinkedIn (https://www.linkedin.com/company/talon-project) is TALON's main channel (Figure 3) for communicating in real time, sharing news, and interacting with the AI and Industry 4.0 ecosystem. It is therefore suitable for posting business-related content, announcing collaborations, and sharing industry updates. The focus of this platform is exclusively professional, which supports TALON in achieving its purpose of integrating academic research into industrial practice.

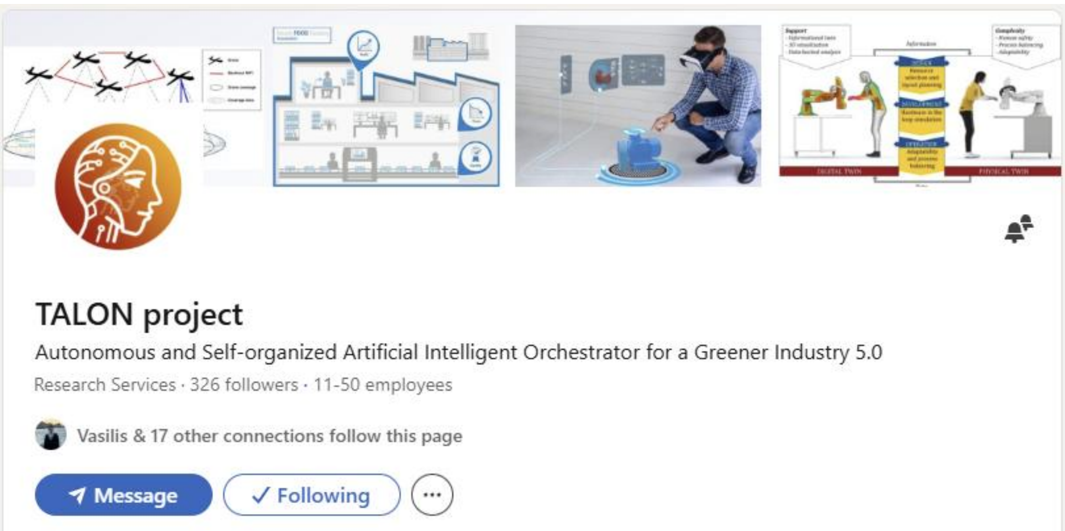

Figure 3: TALON's LinkedIn profile

The following table (Table 10) depicts the most informative metrics from LinkedIn.

Table 10: LinkedIn Platform Analytics

Audience Quality & Relevance	• Job Functions: Engineering (55), Research (31), Education (30) - perfectly aligned with project stakeholders • Industries: IT Services & Consulting (42), Higher Education (38), Research Services (28) • Seniority Mix: Balanced representation from Entry Level (118) to CXO (10), ensuring broad professional impact
Geographic Impact Distribution	• Top Locations: Athens (34), Valencia (29), Thessaloniki (25) - demonstrating regional project relevance • Extended Reach: London (8), Delhi/Barcelona (5 each) - indicating global research community interest • Traffic Alignment: Website visitor locations mirror follower demographics, confirming authentic engagement
Growth Trajectory Analysis	• Net Growth: +95 followers over 12 months through purely organic methods. • Growth Peaks: March-April 2025 (+20 followers) and June 2025 (+8 followers) coinciding with project milestones. • Quality Indicators: Minimal unfollows (-2 total) suggesting high content satisfaction and community loyalty.

2.9.2.2 *X (formerly Twitter)*

Twitter/X (https://x.com/talon_project) is equally important as it helps TALON reach a wide variety of different stakeholders. It allows seamless engagement with followers and industry leaders and enables instantaneous sharing of project updates, research results, and other relevant events (Figure 4).

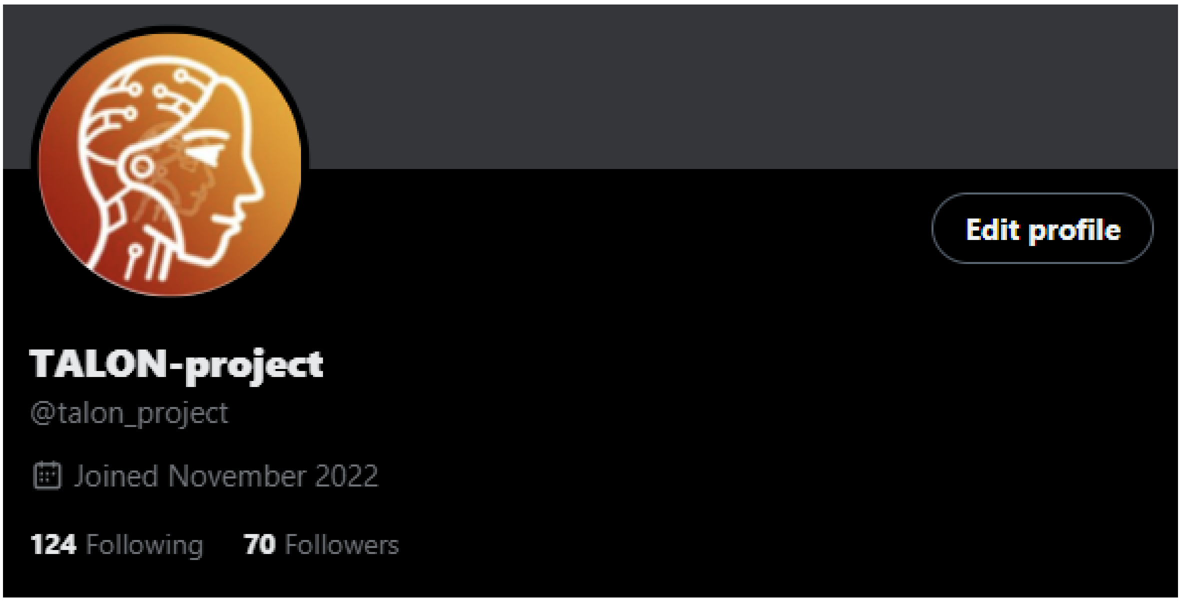

Figure 4: TALON's X profile

The following table (Table 11) depicts the most informative metrics from X.

Table 11: X Platform Analytics

Posting Frequency & Consistency	• Total Posts: 47 posts over 14 months • Content Types: 85% Photo posts, 13% Link posts, 2% Text posts • Average Engagement Rate: 3.2% (well above industry average of 0.5% for research accounts) • Peak Performance: August 2025 with highest engagement rates

Audience Reach & Impressions	• Total Impressions: 656 across all posts (100% organic) • Potential Reach: 70 users consistently maintained • Engagement Quality: 15 total engagements with authentic user interactions • Best Performing Post: XAI Checkpoint publication (81 impressions, 1.23% engagement rate)
Content Strategy Success	• Technical Focus: AI orchestration, autonomous healing, and Industry 4.0 content • Publication Announcements: Direct links to research papers and project deliverables • Real-time Updates: Live demonstration of TALON's capabilities and features • Educational Content: Technical explanations with practical applications
Top Performing Content Categories	• Research Publications (Average 15+ impressions) • Technical Demonstrations (Self-healing systems, AI orchestration) • Industry Applications (Manufacturing, edge computing, sustainability) • Partnership Announcements (Event participation, webinars)

2.9.2.3 YouTube

YouTube (https://www.youtube.com/@talon-project) serves as TALON's repository for video content (Figure 5), which includes educational materials, presentation from conferences, and demonstrations of projects. Through videos, technical concepts can be explained and the outcomes of research work practically demonstrated to many people in an interesting manner.

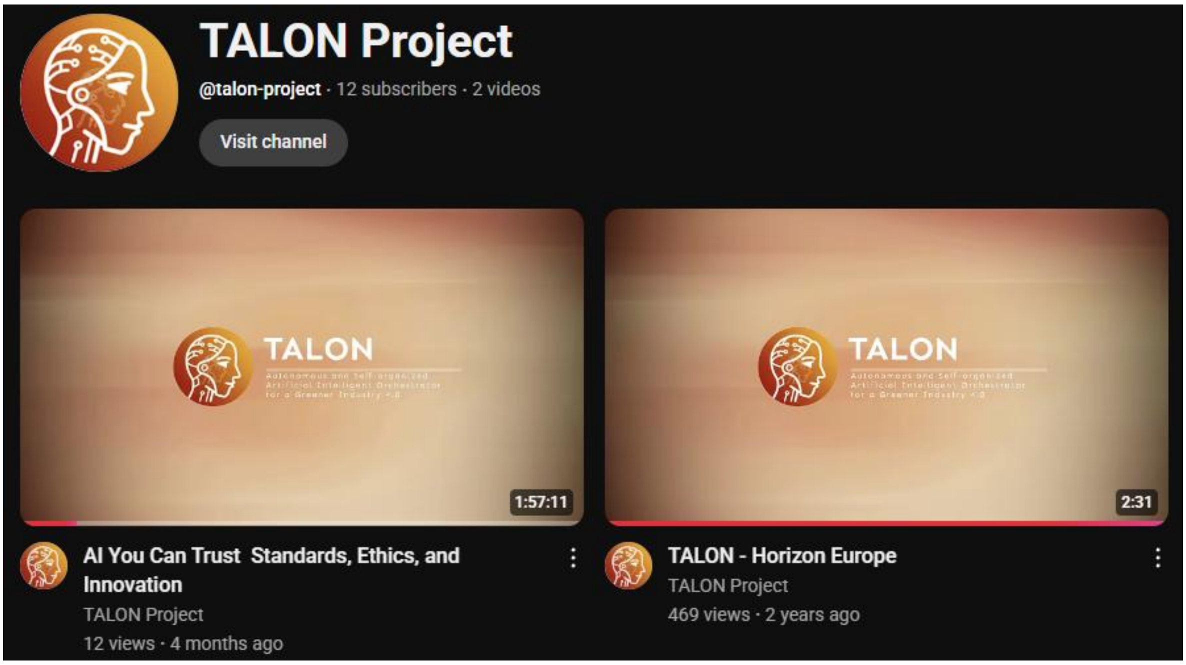

Figure 5: TALON's YouTube Channel

2.9.2.4 Multi-Platform Performance Overview

The following results demonstrate how our strategic social media approach has successfully amplified TALON's reach, fostered meaningful public engagement, and ensured transparent communication of EU-funded research initiatives. Our multi-platform strategy combining LinkedIn-focused precision targeting with X's broader research dissemination has achieved remarkable success in audience building, driving qualified traffic to our project resources while maintaining authentic community engagement across both platforms.

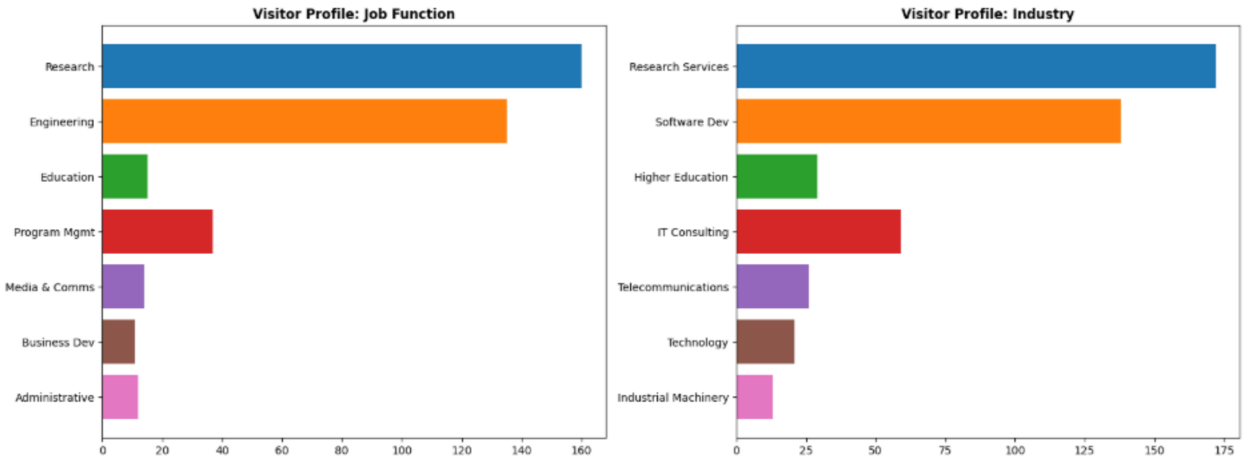

Figure 6: Professional Targeting Success

The visitor profile charts in Figure 6, show overwhelming concentration in Research and Engineering functions/industries, confirming precision targeting achievement.

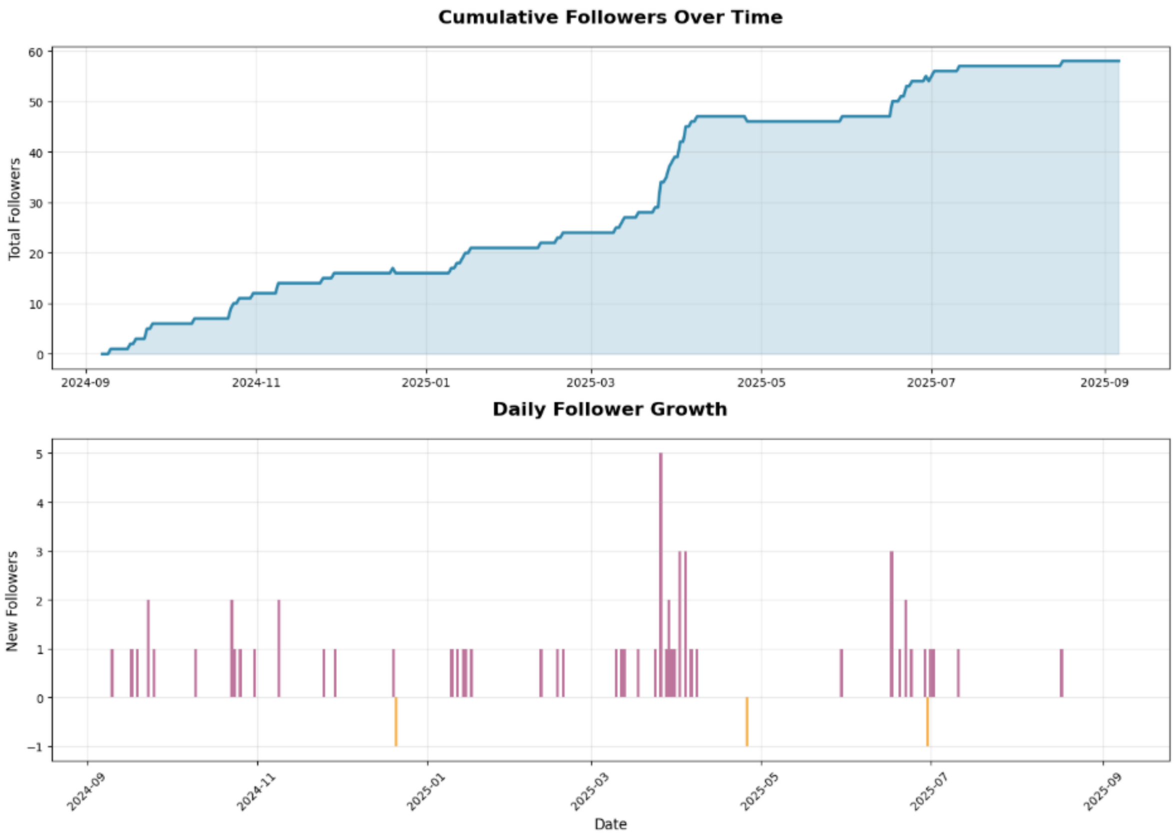

Figure 7: Organic Growth Pattern

Figure 7 reveals a highly successful organic growth model that aligns perfectly with Horizon Europe's objectives for public engagement, transparency, and knowledge dissemination across European research communities.

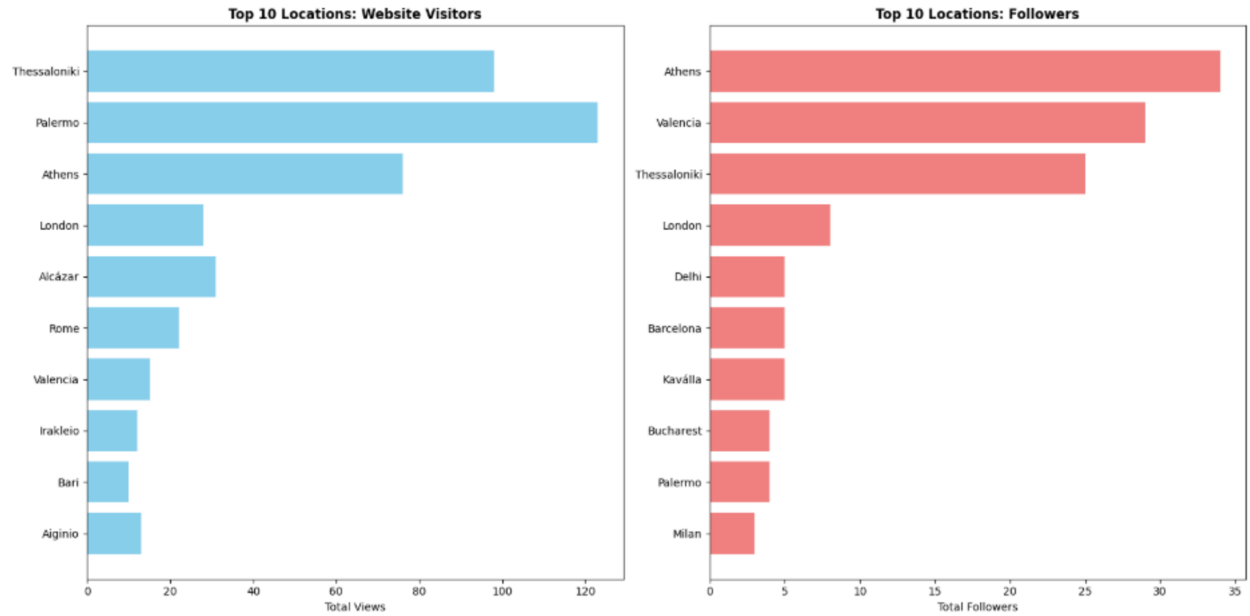

Figure 8: Geographic Consistency

From Figure 8, we observe very high alignment between LinkedIn followers and website visitors by location, proving authentic rather than artificial engagement patterns.

2.10 TALON Newsletters

Regular newsletters (Table 12) are essential for providing stakeholders with a comprehensive digest of any given project's activities, research milestones, and future plans. In the case of TALON, the focal point of the newsletter strategy is to ensure that the project's community is properly informed while also underscoring important developments and chances for participation.

Table 12 TALON's newsletters

No.	Description	Date
1	This is December's newsletter of the TALON project, an EU Funded Horizon Europe Project. In this edition, details concerning the kick-off meeting of the project are provided, the public profiles of TALON are introduced, participation of TALON's partner in dissemination events are presented, and the research efforts of the consortium are highlighted.	January 2023
2	This is May's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project in the first semester of its lifetime, which includes journal and conference publications, as well as organization of events. Specifically, one (1) journal was published, two (2) journals were submitted, one (1) conference paper was successfully accepted, and a workshop was co-organized by the TALON consortium entitled establishing the next level of "Intelligence" and autonomy.	May 2023
3	This is July's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, one (1) journal was published, three (3) conference papers were successfully accepted, and a workshop was co-organized by the TALON consortium entitled "5th International Workshop on IoT Applications and Industry 5.0".	July 2023
4	This is September's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a	September 2023

	summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, two (2) deliverables were published, two (2) conference papers were successfully accepted, and TALON has been designated as a Silver Sponsor for the forthcoming "Italian Conference on Big Data and Data Science (ITADATA2023)" conference.	
5	This is November's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, one (1) deliverable were published, and two (2) conference papers and two (2) journal papers were successfully accepted for publication.	November 2023
6	This is December's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, one (1) presentation of TALON in an international exhibition and one (1) clustering activities.	December 2023
7	This is January's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, one (1) workshop of TALON, one (1) journal publication, and one (1) clustering activity.	January 2024
8	This is February's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, three (3) project developments by TALON partners, and one (1) joint webinar.	February 2024
9	This is March's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, three (3) project developments by TALON partners, and one (1) joint webinar.	March 2024
10	This is April's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include journal and conference publications, as well as organization of events. Specifically, two (2) project developments by TALON partners, and one (1) joint webinar.	April 2024
11	This is May's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include a conference, organization of events, and technological developments. Specifically, two (2) project developments by TALON partners, one (1) conference, and one (1) joint webinar.	May 2024
12	This is June's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include two (2) journal papers, organization of events, and technological developments. Specifically, two (3) project developments by TALON partners, two (2) journal papers, and one (1) joint webinar.	June 2024
13	This is July's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include	July 2024

	organization of events and technological developments. Specifically, two (3) project developments by TALON partners and one (1) joint webinar.	
14	This is September's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, three (3) project developments by TALON partners.	September 2024
15	This is November's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, one (1) event, one (1) project demonstration, and two (2) project developments by TALON partners.	November 2024
16	This is December's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, one (1) publication, one (1) participation in a live event, and one (1) insight by TALON partners.	December 2024
17	This is February's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, one (1) publication, one (1) participation in a live event, and one (1) insight by TALON partners.	January 2025
18	This is March's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, two (2) publications, one (1) organization of a webinar, and one (1) insight by TALON partners.	February 2025
19	This is April's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, one (1) publication, one (1) webinar, and one (1) blog post with the results of the TALON webinar.	March 2025
20	This is May's newsletter of the TALON project, an EU Funded Horizon Europe Project. This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, one (1) publication and one (1) demonstration in a Conference.	April 2025
21	This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, one (1) publication and one (1) demonstration in a Conference.	May 2025
22	This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, six (6) publications, one (1) organization of a special session, and one (1) blog spot providing insights from the TALON Standardisation webinar.	June 2025
23	This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, two (2) white paper	July 2025

	publications, one (1) organization of a webinar, and one (1) post highlighting insights by TALON partners.	
24	This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, one (1) publication and six (6) insight by TALON partners.	August 2025
25	This newsletter presents a summary of the primary dissemination outcomes of the project, which include organization of events and technological developments. Specifically, ten (10) insight by TALON partners.	September 2025

3 Exploitation Activities

3.1 Joint Exploitation Plan

The joint exploitation plan for the TALON orchestration platform envisions its deployment as a state-of-the-art Edge-to-Cloud orchestration solution tailored for industrial enterprises. Upon full instantiation, the TALON orchestrator will empower organizations to efficiently execute, manage, and optimize AI services across heterogeneous and distributed edge and cloud infrastructures. This platform will incorporate advanced capabilities, including energy-aware scheduling and secure, explainable lifecycle management of AI models, thereby enhancing operational efficiency and trustworthiness in complex industrial environments.

3.1.1 TALON Vision and Product Concept

TALON represents a cutting-edge Edge-to-Cloud orchestration architecture developed and deployed to efficiently run, manage, and optimize industrial AI services across heterogeneous edge and cloud infrastructures. It supports zero-touch deployment, energy-aware scheduling, and secure, explainable lifecycle management of AI models. Central to TALON is an AI-driven Service Orchestrator that automates workload placement, dynamic scaling, and self-healing of containerized applications, complemented by a Network Intelligence layer (NG-SDN) that continuously collects telemetry data and employs advanced predictive models, such as Long Short-Term Memory (LSTM) forecasting, to enable proactive and dynamically adaptive resource allocation and quality-of-service (QoS) optimization.

The platform integrates sophisticated anomaly detection, comprehensive monitoring pipelines, and a unified visualization dashboard, empowering operators to audit orchestration decisions, fine-tune Service Level Objectives (SLOs), and continuously monitor system health. This forms a closed-loop explicitly designed for Industry 4.0 and Industry 5.0 environments.

Moreover, TALON incorporates robust policy-driven governance paradigms that enforce compliance with data sovereignty, and security mandates. It supports heterogeneous hardware acceleration at CPU and memory level, orchestrating workloads seamlessly across diverse nodes to maximize performance-per-watt. Architected to be vendor- and cloud-agnostic, TALON ensures modularity and portability across varied industrial ecosystems. The platform further integrates digital twin technology for data quality improvement and augmentation, and validation of orchestration strategies prior to deployment. Its modular design guarantees extensibility, facilitating adaptability to evolving industrial standards and emerging AI-driven applications.

3.1.2 Application Areas and Use Case Insights

The TALON orchestration platform targets diverse industrial sectors where intelligent edge-to-cloud coordination delivers transformative value. Manufacturing represents a core application domain, encompassing smart factories, predictive maintenance systems, and Industry 5.0 automation scenarios where real-time AI workload orchestration optimizes production efficiency and quality control. Robotics applications span HRB environments, where FL enables collaborative AI training while maintaining privacy and security across distributed robotic systems. Logistics and supply chain operations benefit from TALON's autonomous fleet coordination capabilities, dynamic resource allocation, and blockchain-secured data sharing for shipment tracking and operational transparency. Other relevant sectors that could be investigated after the completion of TALON include smart-grid for dynamic pricing and demand-response optimization, and aerospace/automotive industries requiring precision automation with robust digital integration.

Each TALON use case validates specific platform capabilities through real-world industrial scenarios. UC2 demonstrates XAI integration through multi-level transparency frameworks, enabling trustworthy decision-making in manufacturing processes with compliance to EU AI Act requirements. UC3 showcases AR maintenance operations with real-time expert guidance, edge AI integration for failure detection, and privacy-preserving image anonymization; thus, reducing downtime while ensuring regulatory compliance. UC4 validates federated learning for distributed AI training across robotic systems, blockchain-secured model sharing, and safety-critical object detection for protective equipment monitoring. These demonstrations collectively prove TALON's ability to deliver 30% operational cost reductions, 25% energy savings, and 95% reduction in human deployment errors across diverse industrial environments.

TALON addresses fundamental industrial needs that transcend sector boundaries, establishing the platform's universal value proposition. Interoperability emerges as a critical requirement, with TALON's vendor-agnostic architecture enabling seamless integration across heterogeneous hardware infrastructures, legacy systems, and diverse cloud-edge environments. Security constitutes an essential foundation through blockchain-secured federated learning, permissioned access control, and GDPR-compliant data handling mechanisms that establish trust and integrity across multi-actor industrial ecosystems. Adaptability represents the platform's core strength, manifested through energy-aware scheduling, self-healing mechanisms, digital twin validation, and modular design that accommodates evolving industrial standards and emerging AI applications. These common capabilities position TALON as an innovative solution for Europe's digital and green transformation, aligning with Gaia-X interoperability principles, FIWARE open standards, and EU regulatory frameworks for trustworthy AI deployment.

3.1.3 Strategic Positioning

TALON’s value proposition is based on an intelligent orchestration platform connecting advanced AI technologies in industrial environments. Its strategic positioning focuses on bringing tangible value to several actors within the European industrial ecosystem. These main beneficiaries are:

- **Industrial end users (factories and plants):** Manufacturing and industrial companies can deploy TALON to orchestrate AI and IoT workloads between physical devices (on plant) and cloud data centres. This can improve operational efficiency, reduce downtime and provide transparent explainability about AI decisions, which is essential for safety and regulatory compliance. TALON can be deployed on factories premises, guaranteeing sovereignty and data confidentiality requirements compliance
- **System integrators and technology providers:** ICT solution integrators and industrial automation vendors can incorporate TALON’s orchestration capabilities into their offers. This can bring added value to their clients, as TALON acts as an intelligent orchestration layer and integrates Explainable AI, offering reliable and auditable solutions.
- **Cloud and Edge infrastructure providers:** Companies offering cloud services, edge computing platforms or telecom operators can benefit from TALON’s capabilities. By integrating TALON they can offer clients a more autonomous and efficient management of applications. This is due to TALON’s workload distribution capabilities (edge and cloud) following latency, cost or energy efficiency criteria. Moreover, blockchain based security allows federated learning scenarios and secure data sharing, which are relevant in multi-actor industrial ecosystems.

TALON is aligned with the main political and technological frameworks in the EU. The AI Act (https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai) is one of them. Its explainable AI integration guarantees transparency and accountability, allowing companies regulation compliance, such as AI act requirements. This alignment positions TALON as a compliant-ready solution for high-risk AI applications, giving European industries a head start in meeting regulatory standards for transparency, fairness and accountability.

Furthermore, the orchestrator’s energy-aware scheduling optimizes how and where computations run, helping to minimize energy consumption and carbon footprint of industrial AI operations. This aspect links it with the EU Green Deal’s climate targets, positioning TALON as a facilitator of digital and green transition.

Besides this, TALON aligns with several European platforms. One of these is Gaia-X[1], which is an initiative to develop a federated secure data infrastructure for Europe, whereby data are shared, with users retaining control over their data access and usage. Both share interoperability, decentralization and digital sovereignty principles. Another is FIWARE[2], which offers an open platform adopted in

[1] *https://gaia-x.eu/what-is-gaia-x/*
[2] *https://www.fiware.org/about-us/*

Europe to build smart solutions (smart cities, industry, etc.), focusing on open standards and interoperability in IoT environments. FIWARE includes enables such as FogFlow, which is a distributed execution framework that dynamically orchestrates IoT services between edge and cloud to reduce bandwidth use and improve low latency. TALON shares the edge to cloud efficient orchestration approach.

In an increasingly competitive technological environment, TALON positions and differentiates itself by integrating capabilities that are usually addressed separately, these are: explainable AI, which contributes to increase transparency in industrial environments where this is critical, sustainable computing, reducing AI environmental impact, hybrid edge-to-cloud deployment and security and traceability using blockchain.

3.1.4 Development and Deployment Roadmap after M36

From a market perspective, TALON strategically positions itself at the convergence of several rapidly expanding technological domains, notably Edge and Cloud Computing, Explainable AI, and Cloud Orchestration. This combination of advanced technologies establishes TALON as a crucial facilitator for industries seeking to implement AI solutions characterized by responsibility, and efficiency.

Currently, TALON is engineered and designed for deployment within industrial enterprises, either on plant premises or within the customer's private network. This deployment paradigm ensures stringent control over confidential data and accommodates tailored orchestration solutions that align with industry-specific requirements, such as seamless integration with legacy systems and adherence to sector-specific legal and compliance mandates.

This strategic approach not only aligns with major European initiatives but also strengthens TALON's market positioning for continuous growth. Its unique integration of explainable AI, energy-efficient operations, and adaptable edge-to-cloud orchestration provides a robust foundation for emerging business models in an increasingly demanding market—one that prioritizes transparency, sustainability, and regulatory compliance in AI-driven technologies.

3.1.5 Sustainability and Future Plans

Through established relationships with system integrators, manufacturing companies, and telecom operators, TALON aims to facilitate seamless integration of its AI-driven orchestration capabilities into existing industrial ecosystems. These partnerships will focus on co-development initiatives, pilot extensions, and joint ventures that leverage the platform's unique combination of energy-aware scheduling, explainable AI, and edge-to-cloud optimization to address real-world Industry 4.0 and 5.0 challenges.

Building upon the comprehensive IP assessment conducted through the Horizon Results Booster program, TALON partners have developed clear licensing strategies and continued collaboration frameworks for the platform's key exploitable results. The project has established robust IP protection mechanisms for the Zero-Touch AI Orchestrator, blockchain security framework, and explainable AI components, while ensuring compliance with GDPR, the Data Governance Act, and EU AI Act requirements. Post-project collaboration will be sustained through joint commercialization efforts, participation in standardization consortia, and dedicated resource allocation from consortium members, enabling technology transfer and market-ready solution development.

The TALON platform's modular, vendor-agnostic architecture incorporates open-source components and aligns with established European standards including Gaia-X and FIWARE frameworks. The platform's integration with Kubernetes-based container orchestration, IPFS distributed storage, and standard APIs ensures interoperability and extensibility for future industrial applications. Consortium partners have committed to pursuing follow-up Horizon Europe projects, national R&D initiatives, and standardization activities that will further mature the platform's capabilities, expand its application domains, and contribute to European digital sovereignty goals while maintaining alignment with EU Green Deal climate targets through energy-efficient AI operations.

3.2 TALON Exploitable Results

3.2.1 List of TALON Exploitable Results

Based on the aforementioned methodology, TALON partners have identified the KERs. These are listed in the following table (Table 13) along with their position in the TALON ecosystem, their type, a short description, their owners, and the sources that describe them in detail.

Table 13: TALON's KERs

No	Name	Position in TALON [Architecture Layer]	Type [SW/HW/ Math_model/ Use_case]	Description	Owner (s)	Availability [Publication/ TALON repo/ deliverable/ shared documents]
1	Zero-touch AI Orchestrator	The **Zero-touch Orchestrator AI** resides in the **E2C AI Orchestration Layer** of TALON's Control Plane, acting as the core intelligence for autonomous deployment and optimisation of AI tasks.	SW	The **Zero-touch Orchestrator AI** is TALON's core intelligence that automates deployment and optimisation of AI workloads across edge–cloud environments with minimal human intervention	ENG, UL	Conference publications, Deliverables 3.1, 3.2, 3.3
2	AR Maintenance Application	Pilot (AI Cognition, Authorisation and Authentication, Information, Visualisation and Reporting)	UC3	This augmented reality application enhances maintenance operations by enabling remote experts to provide real-time visual guidance to on-site	KU	Conference publications, Deliverables D5.3 & D5.4

				technicians. Through Edge AI, AR overlays and live annotations, it supports accurate, efficient task execution while reducing downtime and error rates.		
3	VR Training Application	Pilot (AI Cognition, Authorisation and Authentication, Information, Visualisation and Reporting)	UC3	This virtual reality application provides immersive training for new personnel in industrial environments. It enables safe, repeatable, and cost-effective simulation of real-world scenarios, helping users build practical skills and confidence before entering live operational settings.	KU	Journal/Conference publications, Deliverables D5.3 & D5.4
4	Federated Learning Module	Demonstrator #4: HRC	UC4	Distributed training / fine tuning of personal protective equipment ML model.	MINDS	Journal/Conference publications (see Table 1 and Table 2), Deliverables D4.1 & D4.3
5	Explainable AI (XAI) Framework	Demonstrator #2: I5.0 Automation & Planning	UC2	Informative use of various XAI TRLs through	MINDS	Journal/Conference publications (see Table 1 and Table 2), Deliverables D4.1 & D4.3

				Talon Dashboard for Nakamura machine.		
6	Blockchain Mechanism	Security Layer	UC2, UC4	Secure and private access and storage of AI model weights	SID	Conference publication (see Table 2), Deliverables D4.2 & D4.3
7	Image Anonymization Module	Access & Security layer	SW - UC3 (scenario 2) / UC4	The Image Anonymization Module (IAM) is one of the Access & Security layer's mechanisms that are designed to protect the confidentiality of personal data handled by Talon project.	TEI	Public Repo in Docker Hub: teiginc/anonymization_image:latest teiginc/cronjob-delete-obsolete-records:latest

3.2.2 Introducing TALON Exploitable Result

Zero-Touch AI Orchestrator (ENG, UL)
Developed jointly by ENG and UL, the Zero-touch AI Orchestrator presents a solution that enables dynamic and intelligent management of network and computing resources. It aims to leverage energy and system metrics for efficient workload deployment, increasing performance while minimising energy consumption.
Its main features therefore address Orchestration, configuration, monitoring, and analysis of network slices using AI and machine learning algorithms, with the aim to enhance efficiency, flexibility and scalability.

As part of the TALON project, this is utilised in use cases 3 & 4 while its TRLs at the start and end of the project are 2 and 5 respectively.

The primary need that the zero-touch AI orchestrator targets is the Energy efficiency-aware deployment in cloud and edge environment. End-users would likely include Cloud Providers, Vertical industries, Research and academic institutions. Its competitive landscape would include companies such as CloudBolt, RedHat that provide a hybrid cloud management platform and develop products specifically like RedHat Ansible.

AR Maintenance Application (KU)
Developed by KU, the AR Maintenance Application presents integrates augmented reality technology to provide real-time guidance and visualization overlays for equipment maintenance tasks, improving efficiency, accuracy, and safety by overlaying digital instructions, diagrams, and annotations onto physical machinery or environments.

As part of the TALON project, the software is being utilised in use case 3 while its TRLs at the start and end of the project are 3 and 6 respectively.

The AR Maintenance Application's competitive advantages lie in the integration of AI Object Detection methodologies for failure detection, 3D visualisation of the worker environment for the remote assistant and 3D digital avatar of the remote assistant in the work environment. Moreover, the software provides implementation of Lifelong Learning AI in the maintenance, real-time 3D visualisation of the worker environment and an interactive 3D/2D dashboard for the remote assistant.

Based on the above-mentioned features, the AR maintenance application effectively addresses the need for improved guidance, remote support, predictive maintenance, safety, and knowledge management in equipment maintenance and servicing, ultimately enhancing operational efficiency, reducing downtime, and ensuring compliance with regulatory requirements.

Its competitive landscape includes providers of AR and Mixed Reality (MR) remote assistance solutions such as Scope AR, PTC, Librestream, Help Lightning Fieldbit and XMReality.

In terms of its future exploitation plans, KU aims to exploit the software for direct industrial use, product development and service development. Moreover, the AR Maintenance Application can be exploited alone or along with the VR Training Application component. Finally, the software can be sold as Software-as-a-Service (SaaS) including a free basic version and silver, gold memberships for more features and constant content updates or on demand personalisation

VR Training Application (KU)
Developed by KU, the VR Training Application enables employees to undergo immersive virtual training sessions where they can practice operating machinery, performing assembly tasks, and troubleshooting equipment issues in a safe and controlled virtual environment. This approach improves workforce training efficiency, reduces the risk of accidents, and ensures that workers are proficient in their roles before engaging in real-world operations on the factory floor.

Specifically, within a manufacturing setting, the VR Training Application's offers as its main features immersive virtual environments, hands-on scenario-based learning, performance analytics, safety training modules, cost-efficient training solutions, real-time feedback, and seamless integration with learning management systems.

As part of the TALON project, the software utilised in use case 3 while its TRLs the start and end of the project are 3 and 6 respectively.

It follows from above that the VR training application addresses the need for immersive, hands-on training experiences that improve skill development, reduce training costs, and enhance safety and performance across different industries and sectors.
The competitive landscape includes companies which provide VR-based training and immersive learning solutions such as STRIVR, Immversion VR and Virti.

Compared with the known competition, KU's VR Training Application's competitive advantages lie in the provision of an interactive environment where the user is allowed to pick up and change the surrounding objects instead of point-and-click mechanisms observed in the main competitors. Furthermore, the solution provides a virtual 3D environment instead of pre-recorded 360°-degree videos. Moreover, the application introduces as its main innovation a more immersive and interactive hands-on training methodology that should further improve the learning efficiency, time, and cost in comparison with the more conventional 360° degrees VR training applications available on the market.

In terms of its future exploitation plans, KU aims to exploit the software for direct industrial use, product development and service development. Moreover, the VR Training Application can be exploited either alone or along with the AR Maintenance Application component. Finally, the application can potentially be sold using a subscription model (e.g. Software-as-a-Service) including a free basic version and silver, gold memberships for more features and constant content updates or on demand personalisation

Federated Learning Module (MINDS)
Developed by MINDS, the Federated Learning Module presents a solution which provides a unified framework to support federated and decenteralised AI model training, taking into consideration the privacy limitation of private or sensitive data, alleviating the need to retrieve or transport these data. It supports AI model optimisation leveraging the shared knowledge that local models have accumulated only retrieving the models and aggregating them into a global optimised model. Finally, it is then re-distributed to the remote devices offering a more advanced model, that is universally applicable for inference.

As part of the TALON project, this specific asset is utilised in use case 4 while its TRLs at the start and end of the project are 3 and 5 respectively.

The Federated Learning Module's competitive advantages lie in its easily adaptable environment, integrated model support and easy deployment. Moreover, it provides adapted performance, custom visualisations and UI.

Eventually end-users would potentially include industry, research projects, and SMEs.

The module would be suitable to be exploited as stand-alone, while Blockchain presents an interesting potential for synergies. In terms of its near future exploitation plans, MINDS aims to integrate it in further research and development projects. Furthermore, the Federated Learning Module is suitable to be potentially exploited via subscription Business models e.g. SaaS or PaaS.

Explainable AI Framework (MINDS)
Developed by MINDS, the Explainable AI (XAI) Framework presents which provides a series of evaluations and explanations on the data, AI models and their respective results. Also, it provides four TrLs of explanations for advanced feedback and intuitive visualisations and feedback to help the end-user control and optimise their models.

As part of the TALON project, the software utilised in use case 2 while its TRL at the start and end of the project are 3-4 and 5 respectively.

The XAI Framework's competitive advantages lie in its easy deployability and adaptability to AI models and the provision of API-based access with no additional need of integration. Moreover, it provides additional explanations and taxonomy of metrics and visualisations as well as four levels of TrLs for monitoring the explainability of results and outcomes of AI modes, while also providing insights for the user.

Suitable end-user categories would include AI model developers, researchers, and industry
In terms of its near future exploitation plans, MINDS aims to integrate it in further research and development projects. Furthermore, the XAI Framework once fully mature could be exploited either stand-alone or along with other components, especially any AI-related task. Primary, suitable business models would include SaaS.

Blockchain Mechanism (SID)
Developed by SID, the Blockchain Mechanism provides the security and privacy framework to the AI Federated Learning module. By securing the weights on the blockchain's ledger, it also provides the appropriate security and transparency of the information.

As part of the TALON project, the software utilised in use cases 1 & 2 while its TRLs at the start and end of the project are 1-2 and 6-7 respectively.

The Blockchain's competitive advantages lie in its lightweight for and easy integration with other systems. Moreover, it provides a robust mechanism in terms of performance efficiency.

In terms of its near future exploitation plans, SID aims to integrate it in further research, R&D projects as well as pursue exposure in scientific papers. Finally the Blockchain Mechanism can be exploited as stand-alone technology and eventually be sold against one-off payment, while Federated Learning presenting especially interesting prospects for synergies.

Image Anonymization Module (TEI)
Developed by TEI, the Image Anonymization Module (IAM) system which acts as a fully-configurable server for face anonymization or pseudonymization. It also includes an auditor to delete (in case of pseudonymization) the original image retained in case the saving date is out of a configured retention period.

As part of the TALON project, the system is utilised in use cases 3 (scenario 2) & 4 while its TRL in the start and end of the project are 2 and 5 respectively.

TEI's image anonymization module has a range of main features, including; applicable to both pictures or videos, configurable face detection + tuning, configurable face anonymization + tuning, Exif removal, anonymization/pseudonymization and auditor for image removal with configurable retention period.
It therefore effectively addresses specifically the need of data privacy and personal data retention.

In terms of its competitive advantage, with respect to the prior art, it is that in the IAM the original images are stored in a secure database, a cronjob cleanup original images older than a configured retention-period. The prior-art uses cryptography.
So, in both cases there is a "critical" point where to preserve original images and cryptographic keys.
In case of prior-art (use of cryptographic key) it could be possible to suffer of a brute-force attack that try to recover original images from anonymized one.
The introduced cronjob to clenup the out-of-a-retention-period stored images to be as compliant as possible to the GDPR directives.

In terms of its near future exploitation plans, TEI aims to integrate it in research and development projects, direct industrial use as well as in competence building.
Finally, the Image Anonymization Module can be exploited as stand-alone item, with a potential target end-users being, industry, research projects and SMEs.

3.3 Horizon Results Booster

The Horizon Results Booster (HRB) is an initiative of the European Commission with the aim of maximizing the impact and returns of investments made in Horizon Europe. In line with this goal, the program provides additional expertise through its tailored services. The program also implements the expert support services (ESS) where it focuses on systematically increasing the market readiness, decreasing the risks associated with commercialization, and aligning critical innovations with the digital sovereignty strategy of Europe.

As part of the project's Exploitation task, the TALON consortium, taking into account its needs and the available bundles of service offered by The Horizon Results Booster (HRB), applied for PDESC - Module C, service in March 2024. The HRB application was successfully accepted in March 2024. Upon assignment of the experts a couple of introductory calls and a well-planned webinar for all project partners were held. The workshop addressed the significance of exploitation and more specifically the notion of Key Exploitable Results (KERs), together with the ways leading to their identification and their further elaboration. Having taken all the necessary info, TALON partners decided which 3 exploitable results merited to be qualified as "Key" Exploitable Results. More specifically:

- SIDROCO Holdings (SID) leads the Blockchain Mechanism,
- MetaMind Innovations (MINDS) leads the XAI Framework, while
- ENGINEERING SPA (ENG) and University of Luxembourg (UL) lead the E2C AI Orchestrator.

The designated HRB experts have also been tasked with executing dedicated workshops for market and technical validation with the goals centered around focused themes. First, characterization of the results in a structured manner to enable conversion of the research results into investable assets. This confirmed the integration of TALON's blockchain and XAI as well as orchestrator technologies into a cohesive Industry 5.0 framework. Second, creation of exploitation roadmaps which are tailored to specific sectors with clearly defined KPIs defining commercial objectives for 1 to 5 years per innovation.

Thus, the owners of the 3 KER's got the opportunity to explore in a structured way the exploitation aspects of their assets and document their respective views in a number of given formats, comprising the Priority Map, the Exploitation intentions table, the Characterisation tables, the Use options and the Exploitation roadmap. Their inputs were checked and scrutinised by the HRB service experts who provided during a workshop dedicated individually to each KER to produce a more refined and comprehensive version of their exploitation plans.

Ensuring compliance with changing EU regulations is of the highest importance. The HRB specialists integrated the requirements of GDPR, the Data Governance Act, and the AI Act into all three KERs designs, making them more resilient to regulatory changes. This also included freedom-to-operate assessments which clarified delineation of stake: SIDROCO Holdings leads the Blockchain Mechanism, MetaMind Innovations leads the XAI Framework, while ENGINEERING SPA and University of Luxembourg lead the E2C AI Orchestrator.

TALON successfully completed HRB's PDESC – Modul C service in September 2024.

3.3.1 Methodology to Identify KERs

The notion of KERs represents a cornerstone in the wider framework of project exploitation and sustainability planning. For this reason, it received special attention during the HRB sessions, ensuring that all partners shared a common understanding of its importance in shaping the project's impact. The identification of the TALON KERs was treated as a core objective guiding the interactive process: participants collectively explored potential outcomes, assessed their exploitable value, and clarified how these results could transform into tangible benefits for stakeholders, industry, and policy.

A central objective of exploitation activities under Horizon Europe (HE) is the maximisation of Impact, always understood in relation to the project beneficiaries and the wider stakeholder ecosystem. This is also inevitably linked with the individual results produced by a given project.

It therefore follows that what sets apart an Exploitable Result and qualifies it as "Key" is its exceptional potential for value creation and impact, as rigorously assessed against core criteria: degree of innovation, exploitability, and expected impact [1].

A KER is expected to directly address specific needs or demands of a well-defined group of adopters, offering solutions that surpass the current state of the art in effectiveness, efficiency, or applicability. The selection of a KER is a collective decision by project partners, reflecting a consensus on which results have the highest potential for further exploitation, in different directions —be it direct or indirect use, commercial or societal impact, policy influence, or knowledge advancement.

Moreover, a variety of output types can qualify as KERs, reflecting the diverse nature of innovation and outputs generated within Horizon Europe projects. These include (but are not limited to) tangible products or manufacturing processes, innovative services, newly developed standards, as well as novel methodologies, platforms, or knowledge that can be directly utilized as inputs for further research, development, or practical application.

Taking into consideration the above criteria and the specificities of the TALON project, as well as the central role these play within the overall envisioned solution, the consortium has qualified the following three as Key Exploitable Results (KERs):

- Zero-Touch AI Orchestrator by ENG/UL
- Explainable AI (XAI) Framework by MINDS
- Blockchain Mechanism by SID

The above three were studied and researched deeper in terms of their exploitation aspects, as part of the HRB and under the guidance of the assigned experts.

3.3.2 Characterization of identified KERs

The TALON project pioneers an integrated architecture for Industry 4.0/5.0, delivering three transformative technologies that address critical gaps in secure, explainable, and efficient AI deployment across edge-to-cloud networks. An overview of the three KERs along with the core innovation, technical highlights, and industrial impact are provided in the following table (Table 14).

Table 14: TALON's KERs augmented through HRB

Key Result	Core Innovation	Technical Highlights	Industrial Impact
Blockchain Mechanism	Federated learning secured via permissioned blockchain	• RAFT consensus (1,000× faster than PoW) • IPFS off-chain storage • 99.99% lower edge-device overhead	Enables secure ML model sharing in e.g. supply chains, telecom, and e-commerce
Explainable AI Framework	Multi-level AI interpretability platform	• 4 TrL-level explanations • Real-time visualization APIs • 40% faster model debugging	Meets EU AI Act requirements for healthcare, finance, and autonomous systems
AI E2C Orchestrator	Dynamic edge-cloud resource optimization	• Smart pod selection • Digital twin simulation • 30% operational cost reduction	Revolutionizes smart manufacturing, 5G networks, and logistics automation

The described KERs integrate into a cohesive structure for Industry 5.0: the Blockchain Mechanism secures federated learning operations, the XAI Framework audits AI decisions for compliance, and the E2C Orchestrator dynamically deploys optimized models at the edge-cloud continuum. TALON was validated in four different pilots which included unmanned vehicle coordination and industrial process optimization. These results support European Union priorities including the transparency mandates of the AI Act, privacy provisions of the Data Governance Act, and sustainability targets of the Green Deal. The project enables the convergence of security, explainability, and operational efficiency, and in doing so, establishes a trusted autonomous industrial ecosystem that integrates the intelligence from edge to cloud.

3.3.2.1 AI Edge-to-Cloud Orchestrator

Description and Problem Solved: ENG's AI E2C Orchestrator is a novel paradigm in the management of distributed computing resources for industrial AI use cases. While modern manufacturing and logistics integrate AI features with real-time processing on the edge, on local servers, or on the cloud, current orchestration tools do not dynamically adjust for latency, accuracy, and energy proportional to load; hence, some nodes become bottlenecked while others remain underutilized. This results in suboptimal resource utilization with an estimation of around 30% dormant computing power in industrial systems due to lack of true demand-driven resource scheduling. The solution provided here is autonomous control of the entire edge-to-cloud continuum using intelligent orchestration.

Adopted Methodology: At its core, the orchestrator employs semantic AI to understand task requirements and match them with optimal computing resources through "smart pod selection." Digital twin technology allows simulation and validation of deployment strategies before real-world implementation. The microservices architecture ensures flexibility and scalability, enabling the system to manage thousands of heterogeneous devices across a 5G-enabled industrial IoT environment. Reinforcement learning algorithms continuously optimize resource allocation based on real-time performance data and changing operational priorities.

Value Proposition: Industrial adopters gain three significant advantages. First, demonstrated 30% reductions in operational costs through optimized resource utilization. Second, 25% energy savings from intelligent workload distribution. Third, near-elimination (95% reduction) of human error in deployment decisions, as shown in unmanned vehicle coordination pilots. Perhaps most importantly, the system's self-learning capabilities mean its performance improves continuously as it gains operational experience.

Target Market and Go to Market Strategy: Key markets include smart manufacturing (predictive maintenance and quality control), telecommunications (5G network management), and logistics (autonomous fleet coordination). ENG's commercialization strategy focuses on embedding the orchestrator into its existing industrial software suite while pursuing contract research projects with manufacturing leaders. The company will also leverage academic partnerships, with UL University incorporating orchestrator concepts into graduate-level curricula to build industry expertise. Commercial rollout will follow a phased approach, beginning with pilot installations in Year 3 and reaching full market deployment within 5 years, using an innovative outcome-based pricing model tied to actual performance improvements.

3.3.2.2 Blockchain Mechanism for Federated Learning

Description and Problem Solved: The Blockchain Mechanism innovation achieved by SIDROCO Holdings constitutes an advancement in securing learning processes conducted at the edge and in the cloud. Organizations now face the most urgent problems of ensuring data integrity, model provenance, and security in the context of data fading Industrial IoT. Data is increasingly processed at the edge, or close to where it is generated. Typical federated learning systems do not provide sufficient guarantees regarding the integrity or the authenticity of the model used, and classic blockchains like Ethereum face extreme latency (measured in minutes per transaction) and energy consumption which effectively rules them out for real-time industrial use, thus lacking edge friendliness.

Adopted Methodology: To mitigate the aforementioned problems, the has adopted an innovative hierarchical architecture based on Hyperledger Fabric using RAFT reach consensus. These approaches have improved transaction processing to, for example, milliseconds instead of latencies in the thousands. In Proof-of-Work systems, an entire ecosystem can be built around the proposition to store only cryptographic hashes of model weights, IPFS (InterPlanetary File System) as the providing the smart offload edge devices firmware requirements. A combination of lightweight clients, powerful edge devices, and distributed storage achieves 99.99% reduction in local storage and processing needs.

Key Value Proposition: The Blockchain Mechanism has three transformative impacts on industrial users. First, it establishes trust and verification of the data integrity through the immutable ledger. Second, permissioned access control guarantees that only verified participants can contribute to or access the federated learning process. Third, the system's responsiveness in real-time to events can be critical for certain latency-sensitive applications. All these advantages come with operational costs lower by 90% compared to traditional blockchain approaches.

Target Market and Go-to-Market Strategy: The primary focus of the solution combines three verticals—logistics and supply chain for shipment tracking systems, telecommunications for network operators managing distributed infrastructure, and retail/e-commerce for detection of fraudulent transactions. The plan for market entry follows a Software-as-a-Service (SaaS) model. Initial pilot testing is scheduled for the first year after the project with R&D partners and logistics providers. These pilots will enable further industrial consortia and tech integrator partnerships for broader commercialization, aiming to achieve full deployment within three years.

3.3.2.3 Explainable AI Framework

Description and Problem Solved: MetaMind Innovations' Explainable AI Framework is designed to address the considerable challenge of the 'black box' issue in AI. As the use of AI systems increases in various fields, and their complexity deepens, the lack of transparency in the decision-making process is a critical hindrance. As with the EU AI Act, some governing policies now require explainability while organizations struggle with trust benchmarks or optimization of systems whose rationale is opaque to them. Solutions in place today like SHAP and LIME either over or under serve by requiring coding expertise or actionable and tailored insights for dynamic industrial contexts.

Adopted Methodology: The technical architecture of the framework incorporates a multi-layered explanation system that adjusts to various user roles and needs. Business users visualize decisions at higher layers while developers access detailed technical metrics. Due to standard APIs, the system's modular architecture facilitates plug-and-play integration with pre-existing AI models which accelerate overcoming implementation barriers. With real-time interpretability features, users can interact with dashboards showing how specific input features shape model outputs, making complex AI behavior navigable.

Key Value Proposition: The framework is transformative for enterprises for three reasons. First, adherence to compliance requirements of new AI regulations. Secondly, early testing shows 40 % savings in model debugging and optimization. Third, its case-agnostic nature ensures relevance in a multitude of industries e.g., healthcare diagnostics as well as financial risk assessment. The framework not only explains the decisions but also helps the organization proactively improve its AI systems.

Target Market and Go-to-Market Strategy: These can include AI technologists, Finance and Healthcare compliance officers, and responsible AI focused researchers and research centers. It will be realized by combination of tiered subscription SaaS model, licensing deals, and direct sales. Researchers and small teams are targeted via a freemium model while higher education institutions and corporate bodies sustain enterprise packages with advanced features and dedicated support. MetaMind's go to market strategy focuses on collaboration with regulatory technologists and AI application developers for quicker access to the market, which is anticipated after 36 months post the project's completion.

3.3.3 HRB Insights

The HRB service provided transformative insights (Table 15) that significantly enhanced TALON's market readiness and societal impact. With their rigorous ESS workshops, HRB specialists drew out major leverage points and risks within all three KERs, converting academic innovations into commercially applicable solutions compliant with Europe's twin digital and green goals. With this tailored approach, HRB empowered TALON to move from research excellence to capture real-world impact—including foundational innovations to accelerate Europe's green and digital transition.

Table 15: TALON-HRB interactions and outcomes

HRB Service Element	Application to TALON	Tangible Outcome
Result Characterization	Technical validation of cross-KER synergy	Unified architecture for secure, explainable, efficient AI
Exploitation Roadmapping	Time-to-market planning per sector	3-year commercialization paths with industry-specific KPIs
Policy Alignment	Integration of AI Act/Data Governance Act	Compliance-ready solutions for EU market entry

Market-Driven Refinement: The ESS workshop identified shortcomings with initial commercialization approaches, necessitating critical pivots. For the Blockchain Mechanism, HRB suggested disaggregating target markets to mid-tier scale logistics and fast-tracking early adopter validation in supply chain pilot programs to reduce risk on scalability assumptions. For the XAI Framework, experts suggested aggressive benchmarking against competitors; thus, weighing their explainability to quantify adjustable values of "explainability ROI," which prompted MetaMinds to introduce tiered SaaS pricing based on compliance severity (e.g., healthcare vs retail).

Risk Mitigation as Growth Enabler: HRB's risk-priority mapping highlighted Market Risk Factor 5 ("Low adoption due to technological immaturity") as critical for Blockchain with a risk grade of 72/100. This prompted pre-emptive action to add IPFS off-chain storage to lessen edge device burdens. Moreover, the IPR Foreground conflict avoidance analysis for SIDROCO (Blockchain) and MetaMinds (XAI) as well as ENG/UL (Orchestrator) resolves potential disputes by defining ownership boundaries.

Policy alignment as commercial catalyst: HRB embedded compliance with GDPR and the AI Act into each KER's design, turning constraints into props. The "multi-level explanations" component of the XAI Framework addresses Article 13 of the EU AI Act directly, while further, the Orchestrator's energy saving measures are in support of the Green Deal 2030 climate targets.

Exploitation roadmap precision: Derived from HRB commentary, exploitation timelines shifted for the following:
- Blockchain's "36-month to market" now includes Year-1 R&D pilot partnerships.
- XAI pilot's 12-month acceleration includes urgent regulatory capture.

Shifts in KPIs also included surpassing subsystems targets like "consensus speed" and shifted to "30% cost reduction in logistics" as user scored outcomes.

Validated Impact: By embedding feedback from experts into roadmap exploitation, HRB ensured that the legacy of TALON extends well beyond publications to providing economically viable, compliant, and actionable innovation frameworks that positioned European industry as a leader in socially responsible AI technology.

3.3.4 Overall evaluation of the HRB service

The HRB service and the selected path of PDESC – Module C have proven to be a highly valuable resource for the TALON consortium. The structured support provided through this service has facilitated the systematic exchange of focused information, the organisation of dedicated workshops, and a variety of targeted interactions. Collectively, these activities have enabled TALON partners to develop a deeper understanding of the processes surrounding exploitation and to obtain a much clearer perspective on the opportunities available to them, as well as the most effective ways to approach these opportunities.

One of the most significant outcomes of this engagement was the clarification of the overarching relationship and interplay between several central concepts, namely Impact, Key Exploitable Results (KERs), and Use. The partners came to appreciate not only the distinctions between these elements but also the ways in which they are interlinked and reinforce one another. This enhanced understanding further shed light on how impact can be maximised in relation to project beneficiaries. The insights gained have therefore contributed directly to strengthening the partners' capacity to position their results in ways that generate concrete and sustainable benefits.

Another key aspect that was particularly valued by the partners was the detailed guidance on the formulation of exploitation strategies. The HRB sessions and PDESC – Module C methodology highlighted how a well-structured exploitation plan provides a clear pathway from results to impact. Rather than treating exploitation as a general or abstract process, the training demonstrated the importance of defining clear objectives, identifying relevant stakeholders, and setting out the specific measures required to bring research outcomes closer to market or to practice. This structured approach was highly appreciated by the consortium, as it enabled each partner to identify practical steps for moving forward and to see exploitation as a manageable, actionable process.

Of particular importance was also the emphasis placed on risk management. The sessions underscored that exploitation planning must not only focus on opportunities but also take into account potential risks, barriers, and uncertainties. Integrating risk management into exploitation strategies was recognised as an essential factor in increasing the likelihood of success and in ensuring that unforeseen challenges can be anticipated and effectively addressed.

In practice, the takeaways and knowledge acquired through the HRB service and its associated modules have supported the TALON partners in developing their individual exploitation strategies and plans. These strategies now reflect a stronger alignment with project objectives, clearer identification of exploitable results, and a more concrete vision of how to realise impact. The more business-oriented analyses carried out in this context are also evident in the Business Analysis chapters of this deliverable. In particular, the lessons learned and methods applied are directly reflected in the SWOT analyses and Business Model Canvases (BMCs) prepared for the three Key Exploitable Results of the project. These analyses, which are presented in detail in the subsequent sections of this document, serve as tangible evidence of how the HRB service has contributed to enhancing the partners' exploitation capacity and to strengthening the overall strategic outlook of the project.

In summary, the collaboration with the HRB service through PDESC – Module C has provided the TALON consortium with critical knowledge, tools, and methods. It has improved the partners' ability to conceptualise exploitation in a structured way, to integrate risk considerations, and to develop actionable and well-founded strategies. The results of this process are directly visible in the project's deliverables and will continue to guide the exploitation efforts of TALON partners beyond the project's lifetime.

3.4 TALON Results' Patentability Assessment

The TALON project has undertaken building a patentability assessment framework which focuses on Industry 5.0 innovation. This section presents the methodology and strategy related to the patentability assessment conducted in TALON and how systematic intellectual property analysis catalyzes the commercialization of advanced research work. The outcomes of the patentability assessment are also shown in section 3.4.6, complete with the reviewers' scores and evaluation comments.

3.4.1 Genesis of TALON's Patentability Assessment Methodology

Right from the start of their research, the TALON consortium realized that the new developments in AI and autonomous systems integration in industrial applications created an innovation paradox. While autonomous systems and AI technologies were anticipated to enable major strides towards sustainability and productivity, the integrated complexity of the systems made it vastly more difficult to identify and safeguard the discrete innovations that resulted from the research efforts. A systematic approach to IP evaluation became necessary, given the inadequacies of conventional approaches to IP evaluation that encountered sophisticated AI orchestrators, federated learning systems, and digital twin technologies constituting the TALON platform.

As the project progressed in its technical work, this challenge became more acute. Various consortium members who contributed to components, such as XAI and blockchain technology innovations faced the challenge of interdisciplinary barriers. The lack of a unified system to assess these innovations posed the risks of valuable intellectual property remaining stagnant and limiting its potential influence on the Industry 4.0 ecosystem. The intricate technologies and rigid frameworks of patent systems necessitated a structured, scientifically rigorous strategy to patentability assessment. This motivated extensive research of different approaches used by the world's leading patenting authorities, such as the European Patent Office [2], US Patent and Trademark Office [3], and Japan Patent Office [4].

As part of this review, the TALON team has noticed that patent offices had examination procedures in place. However, the uniform systems for conducting research domain specific preliminary patentability assessments across disciplines did not exist. This gap was both an opportunity and a problem. The problem was designing a methodology that could cope with the technical challenges posed by the innovations of Industry 4.0, while also evaluating the patent's meaning. The opportunity was to construct a framework that would cater to both the TALON project and other research initiatives that grapple with similar intellectual property issues.

So as to increase the impact of its relevant work by making its approach available to broader audiences and also enable a lasting impact, TALON partners in June 2025 issued a white paper titled "TALON's Patentability Assessment" white paper is available on Zenodo and publicly accessible [5]. The white paper's explicit objective has been to present the relevant efforts that project partners have undertaken so as to explore the patentability potential of a select number of concepts produced. As such the methodologies, rationale, and assessment mechanisms have been presented in sufficient detail, along with a description of the work (up to the point that the patentability evaluation fully completed) together with a timeline extending up to the end of the project.

3.4.2 Building the Foundation

The TALON team established the methodology upon the European Patent Convention (EPC) for practical reasons and in order to adhere to the philosophy of European research culture. The EPC's three-pillar approach which necessitated that inventions had to possess novelty, an inventive step, and industrial applicability, along with restated patent examination practice and legal precedents offered timeless, reliable principles to lean on.

The TALON consortium ensured that the assessment methodology would be useful in a patenting strategy relevant to their commercialization motives by aligning it with European patent law. Furthermore, the project's focus on solving actual manufacturing problems supported the framework's emphasis on industrial relevance.

In most cases, examination of patentability was primarily based on subjective expert evaluation which posed challenges in comparing innovations in different fields or tracking progress through time. The TALON methodology resolved this issue using a more creative approach that relied on both quantitative evaluation and qualitative expert judgment. This two-phase assessment methodology came about from the realization that while numerical marks could easily provide consistency, the actual measurement of technical details that often dictate the success of patentability required context and expertise.

Adoption of 5-point Likert scales along with descriptive text fields provided a framework that combined the rigor and precision needed for systematic evaluation with appropriate depth and nuance for effective patent prosecution. It also recognized that the evaluation of patentability is both an analytical task and a creative task. Thus, a balance needs to be struck between the limitations imposed by measurable KPIs against the dynamic nature of evaluation of technical merit and marketability.

3.4.3 A Comprehensive Question Framework

The TALON patentability survey exemplifies an interplay consideration between comprehensiveness and usability. The survey development team understood that too short a survey would fail to capture the intricate nature of modern technological innovations, while an overly long survey would burden busy researchers and consequence the quality of responses. The final structure of 21 core questions came from iteration refinement based on feedback from technical partners and patent experts.

To serve the assessment framework, each question was crafted to address multiple objectives. Assessing the novelty of an invention was intended to ensure gaps between it and prior art were scoped and also take researchers through a methodical step-by-step process to articulate every distinctive feature of their inventions. This educational role was important, especially because many technical researchers did not have expertise in documentation of patents.

The evaluation questions for inventive steps focused on non-obviousness in a unique way that respects patentability. Instead of asking researchers to make a singular, holistic judgment about whether an innovation would be deemed obvious by a skilled artisan, the framework guided them through parameter selection, unexpected outcomes, technical enhancements, and tailored problem-solving techniques. This breakdown ensured that technical researchers faced fewer barriers in meeting legal standards while drawing close to all guidelines.

The industrial applicability section of the survey aligns with the TALON project's objectives, emphasizing a practical and tangible impact beyond academic confines. Instead of treating industrial applicability as a box being checked off, the framework defined this section with the aim of measuring the commercial and societal relevance of the innovations being assessed. The scientific viability, practical utility, containment within an industry, reproducibility, and cross-industry applicability questions were meant to identify innovations that would be most useful and thus strategically advantageous in broad technology transfer and commercialization processes.

This focus on tangible outcomes is in line with the aims of the Horizon Europe programme which focuses on economic, environmental, and societal sustainability. The TALON methodology reinforced the patenting processes by including evaluation of commercial viability in patenting processes alongside the check on patentability, thus ensuring that patenting expenses would be incurred on innovations which were likely to have the greatest impact.

3.4.4 Balancing Internal Knowledge and External Objectivity

The TALON consortium selected to use both internal (Thomas Lagkas, Associate Professor at the Democritus University of Thrace) and external evaluators (Claudio Agostino Ardagna, Full Professor at the Universita degli Studi di Milano) to avoid biases and limitations within any given evaluation process. Internal evaluators came from the project consortium and therefore had a considerable amount of technical expertise as well as an understanding of the developmental problems. This firsthand viewpoint greatly aided in appreciating the innovations and understanding the developmental challenges that were surmounted to achieve them.

Though strong, insider perspectives created internal evaluation biases as well. Internal evaluators could skew the value of an innovation either positively or negatively due to internal biases. The internal evaluators could be so powerful that they could overestimate the patentability potential of an innovation due to enthusiasm bias. Evaluation familiarity could also lean towards an underestimation of solutions non-obviousness. To achieve the required balance, external experts who had no direct association with the project were brought on board.

Flexibility was infused in the evaluation process by adopting a power transformation in the TALON methodology. This approach recognized expert evaluation as more than scoring on a number line and allowed evaluators to apply contextual judgment to scoring by adjusting evaluations based on their holistic evaluation of all the evidence including non-Likert factors. The implementation of a power parameter system solved the paradox of uniting subjective expertise with objective quantification elegantly. Evaluators could disbelieve the self-assessment offered by innovation developers and still retain the need for systematic assessment across different innovations through quantitative measures. When evaluators set the power parameter below 1.0, they increased normalized score which suggested that innovation's patentability potential was underestimated. Power parameters above 1.0, however, decrease normalized score suggesting that self-assessment was overly optimistic.

3.4.5 Discrepancy Resolution Mechanism

The TALON methodology's mix of quantitative and qualitative elements creates a customizable framework applicable to other research areas grappling with similar intellectual property concerns. This approach is compelling because it successfully balances the two key components of an effective patentability assessment: systematic measurement and expert judgment, both of which cannot stand alone. This approach resolves a critical dilemma in evaluation research—balancing objectivity with subjectivity, standardization with adaptability, and thoroughness with efficiency.

The design of a systematic mechanism for addressing discrepancies is a key methodological contribution which surpasses the TALON project. This framework deals with the unavoidable scenario in which expert evaluators differ with respect to an innovation's patentability potential by providing a method for conflict resolution and consensus achievement. The mechanism's sophistication comes from the fact that disagreement between evaluators is oftentimes anchored on differences in perspective more than a mistake. The system guarantees that difficult cases are properly resolved by requiring collaborative discussion when evaluators diverge fundamentally while enforcing triggers to the resolution process, thus streamlining the evaluation workflow.

3.4.6 Patentability Evaluation Outcomes

This section provides the outcomes of the patentability evaluation methodology applied on 10 TALON components. The detailed numerical scores achieved and reviewers' evaluation and shown in full, for each individual asset considered.
It is worthwhile noting that a score-based ranking has been deliberately avoided, due to the risks it poses in e.g. creating erroneous impressions or undervaluing outcomes etc.

The full set of completed questionnaires per item is included in Appendix 9.2.1.

3.4.6.1 Zero-touch AI Orchestrator

The table below (Table 16) shows the Zero-touch AI orchestrator evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "Zero-touch AI Orchestrator" are available in Appendix 9.2.1 .

Table 16: Zero-touch AI orchestrator evaluation spreadsheet

Q2a	*Q2e*	*Q3a*	*Q3b*	*Q3c*	*Q3d*	*Q3e*	*Q3f*	*Q3g*	*Q4a*	*Q4b*	*Q4c*	*Q4e*
3	3	3	3	3	3	2	2	2	3	3	4	4
Novelty		*Inventive Step*							*Industrial Applicability*			
3		2,571428571							3,5			
Overall Score												
9,071428571												
Normalized Score												
0,505952381												
p_1								p_2				
0,6								*0,5*				
Evaluator Score₁								*Evaluator Score₂*				
0,664455325								*0,711303297*				

Patentability Indicator
69%

Initiate Discrepancy Resolution:	*NO*

Internal reviewer evaluation remarks: The patentability potential of this invention is much higher than the one suggested by the authors especially in questions Q3e-f-g and Q4a. I therefore increased its value using p1, underlying the potential of an advanced AI E2C orchestrator that also considers non-functional aspects. Patentability level: medium.

External reviewer evaluation remarks: The owner's ratings seem to underestimate the patentability potential of the AI E2C Orchestrator. Novelty and non-obviousness should be probably higher, given that the item innovatively combines promising techniques, including zero-touch deployment, LSTM-based telemetry forecasting, context-aware security, energy-aware edge-cloud-on-prem optimization, and self-healing digital twins into a single autonomous system. The orchestrator's system-level innovation promises stronger patent recognition as a foundational technology for future Industry 4.0/5.0 deployments. Patentability level: medium to high.

3.4.6.2 AR Maintenance Application

The table below (Table 17) shows the AR maintenance application evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "AR maintenance application" are available in Appendix 9.2.2.

Table 17: AR maintenance application evaluation spreadsheet

Q2a	*Q2e*	*Q3a*	*Q3b*	*Q3c*	*Q3d*	*Q3e*	*Q3f*	*Q3g*	*Q4a*	*Q4b*	*Q4c*	*Q4e*
4	*3*	*2*	*3*	*3*	*3*	*2*	*3*	*3*	*4*	*4*	*3*	*3*
Novelty		*Inventive Step*							*Industrial Applicability*			
3,5		2,714285714							3,5			
Overall Score												
9,714285714												
Normalized Score												
0,55952381												
p_1									p_2			
1									*0,95*			
Evaluator Score₁									*Evaluator Score₂*			
0,55952381									*0,576006844*			
Patentability Indicator												
57%												

Initiate Discrepancy Resolution:	*NO*

Internal reviewer evaluation remarks: The patentability score produced by authors' self-assessment is in line with the real patentability level. The idea is timely and interesting, I do agree on the possibility that skilled people might reproduce the results of the invention with medium effort. I also agree with the high practical use of the idea in industrial scenarios, I would score Q4a 3 and Q4c 4 compensating the effect on the authos' self-assessment. Patentability level: medium.

External reviewer evaluation remarks: The owners slightly underestimated industrial applicability, despite the AR system solving inefficiencies across possible various domains, but correctly assessed moderate non-obviousness, as an AR application for maintenance advances prior art incrementally. Patentability level: medium.

3.4.6.3 VR Training Application

The table below (Table 18) shows the VR training application evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "VR training application" are available in Appendix 9.2.3.

Table 18: VR training application evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
3	4	3	3	3	3	3	3	3	4	4	4	4
Novelty		Inventive Step							Industrial Applicability			
2,5		3							4			
Overall Score												
9,5												
Normalized Score												
0,541666667												
p_1									p_2			
1									1			
Evaluator Score₁									Evaluator Score₂			
0,541666667									0,541666667			
Patentability Indicator												
54%												

Initiate Discrepancy Resolution:	NO

Internal reviewer evaluation remarks: The patentability score produced by authors' self-assessment is reasonable and in line with the reference domain. Patentability level: medium.

External reviewer evaluation remarks: The owners' evaluation is well-aligned: novelty fairly acknowledges incremental advances like context-aware adaptive scenarios, while moderate non-obviousness appropriately reflects that combining gesture controls, real-time feedback, and scoring, though innovative for industrial training, builds on existing VR concepts. Applicability is moderately high, reflecting the system's utility across sectors. Patentability level: medium.

3.4.6.4 Federated Learning Module

The table below (Table 19) shows the Federated learning module evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "Federated learning module" are available in Appendix 9.2.4.

Table 19: Federated learning module evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
1	5	2	1	2	1	1	2	1	5	5	5	5
Novelty		Inventive Step							Industrial Applicability			
1		1,428571429							5			
Overall Score												
7,428571429												
Normalized Score												
0,369047619												
p_1									p_2			
1									1			

Evaluator Score1		Evaluator Score2
0,369047619		0,369047619
Patentability Indicator		
37%		

Initiate Discrepancy Resolution:	NO

Internal reviewer evaluation remarks: The patentability score produced by authors' self-assessment is in line with the real patentability level. Patentability level: low.

External reviewer evaluation remarks: The owners' assessment is accurate: patentability is low as the module merely implements standard federated algorithms via the open-source Flower framework for YOLOv8 object detection, a predictable application of prior art without technical innovation. Though, industrial applicability remains strong for privacy-preserving safety monitoring in construction, energy, and logistics. Patentability level: low.

3.4.6.5 XAI Framework

The table below (Table 20) shows the XAI framework evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "XAI framework" are available in Appendix 9.2.5.

Table 20: XAI framework evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
3	4	2	2	2	2	1	3	2	5	4	4	4
Novelty		Inventive Step							Industrial Applicability			
2,5		2							4,25			
Overall Score												
8,75												
Normalized Score												
0,479166667												
p_1									p_2			
0,8									0,75			
Evaluator Score1									Evaluator Score2			
0,555123073									0,575923698			
Patentability Indicator												
57%												

Initiate Discrepancy Resolution:	NO

Internal reviewer evaluation remarks: The self-assessment made by the authors is underestimated. The proposed four-layer explainable AI framework for both images and time series using a combination of analytical methods, XAI techniques, and advanced visualizations is novel. Its integration with data diagnostics, visualization, model-specific XAI, and advanced robustness testing contribute to non-obviousness. Patentability level: medium.

External reviewer evaluation remarks: The owners' evaluation is reasonable, but probably somewhat understates patentability: novelty could be considered higher as the native multi-modal integration (images + time series without data conversion) and four-level diagnostic-to-robustness workflow, unifying data quality checks, Grad-CAM/SHAP explanations, and fidelity testing goes beyond prior art's single-modality focus, while non-obviousness deserves a stronger score, since combining these known techniques into a synergistic cross-modal framework solves the fragmented interpretability problem in ML deployments, yielding notable efficiency gains in industrial settings despite using established XAI tools. Patentability level: medium.

3.4.6.6 Blockchain Mechanism

The table below (Table 21) shows the Blockchain mechanism evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the “Blockchain mechanism” are available in Appendix 9.2.6.

Table 21: Blockchain mechanism evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
2	2	2	2	2	3	3	3	2	5	2	4	4
Novelty		*Inventive Step*							*Industrial Applicability*			
3		2,428571429							3,75			
Overall Score												
9,178571429												
Normalized Score												
0,514880952												
p₁									*p₂*			
1									*0,9*			
Evaluator Score₁									*Evaluator Score₂*			
0,514880952									*0,550219709*			
Patentability Indicator												
53%												

Initiate Discrepancy Resolution:	*NO*

<u>Internal reviewer evaluation remarks</u>: I agree with the self-assessment by the authors. Patentability level: medium to low.

<u>External reviewer evaluation remarks</u>: The owners' overall assessment is reasonable, but maybe somewhat unbalanced. Novelty and non-obviousness are probably a bit understated (FL in blockchain is a niche domain), while industrial applicability is maybe a bit overrated (practical adoption is still limited by the low demand for blockchain-secured FL). Patentability level: medium.

3.4.6.7 Image Anonymization Module

The table below (Table: 22) shows the Image anonymization module evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the “Image anonymization module” are available in Appendix 9.2.7.

Table: 22 Image anonymization module evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
3	3	2	3	3	2	1	1	2	5	5	4	5
Novelty		*Inventive Step*							*Industrial Applicability*			
3		2							4,75			
Overall Score												
9,75												
Normalized Score												
0,5625												
p₁									*p₂*			

1,2		*1,2*
Evaluator Score₁		*Evaluator Score₂*
0,501356941		*0,501356941*
Patentability Indicator		
50%		

Initiate Discrepancy Resolution:	*NO*

Internal reviewer evaluation remarks: The patentability score produced by authors' self-assessment is a bit overestimated especially in its novelty. The patentability of the proposed image anonymizer is medium to low. It is true that there is a high industrial applicability potential for the proposed approach, but still with low novelty and non-obviousness. Some work has been done on anonymizing faces in firearm detection tools. Also some tools for face anonymization exists such as Fluendo [6]. Patentability level: medium to low.

External reviewer evaluation remarks: The owners' overall evaluation is reasonable, but novelty is slightly overestimated, since GDPR-compliant retention mechanisms and face anonymization exist in prior art, while the Auditor's cron-based cleanup seems like a logical implementation rather than a novel breakthrough. Patentability level: medium to low.

3.4.6.8 Smart Pricing Simulator

The table below (Table 23) shows the Smart pricing simulator evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "Smart pricing simulator" are available in Appendix 9.2.8.

Table 23: Smart pricing simulator evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
3	4	2	3	3	3	2	3	3	5	5	4	5
Novelty		*Inventive Step*							*Industrial Applicability*			
2,5		2,714285714							4,75			
Overall Score												
9,964285714												
Normalized Score												
0,580357143												

p_1		p_2
0,75		*0,8*
EvaluatorScore₁		*EvaluatorScore₂*
0,664923222		*0,647077453*
Patentability Indicator		
66%		

Initiate Discrepancy Resolution:	*NO*

Internal reviewer evaluation remarks: The patentability score produced by authors' self-assessment is a bit underestimated and I would increase by 25%. Efficient deployment and allocation of resources, in particular when non-functional aspects like energy consumption are considered, are of high relevance and entail a good patentability level. Patentability level: medium to high.

External reviewer evaluation remarks: The owners appear to have underestimated the patentability potential of SPS. While they conservatively rated the inventive step, their own technical disclosures reveal a novel combination of Stackelberg game theory, real-time Holt-Winters forecasting, energy-based incentives, and non-commercial device integration specifically for edge AI. Patentability level: medium to high.

3.4.6.9 Synthetic CNC Data Generator

The table below (Table 24) shows the Synthetic CNC data generator evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "Synthetic CNC data generator" are available in Appendix 9.2.9.

Table 24: Synthetic CNC data generator evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
2	2	2	1	1	3	3	2	2	5	2	2	3
Novelty		*Inventive Step*							*Industrial Applicability*			
3		2							3			
Overall Score												
8												
Normalized Score												
0,416666667												
p_1								*p_2*				
1,1								*1,1*				
$EvaluatorScore_1$								*$EvaluatorScore_2$*				
0,381739969								*0,381739969*				
Patentability Indicator												
38%												

Initiate Discrepancy Resolution:	*NO*

Internal reviewer evaluation remarks: The patentability score produced by authors' self-assessment is a bit overestimated (10%), although it is still reflecting a low patentability level. Patentability level: low.

External reviewer evaluation remarks: The owners' assessment accurately reflects low patentability. Maybe industrial applicability is slightly overrated, since the niche focus limits transferability. Patentability level: low.

3.4.6.10 Few-shot Object Detection

The table below (Table 25) shows the Few-shot object detection evaluation spreadsheet as filled in by the evaluators. The detailed answers to the Patentability questionnaire for the "Few-shot object detection" are available in Appendix 9.2.10.

Table 25: Few-shot object detection evaluation spreadsheet

Q2a	Q2e	Q3a	Q3b	Q3c	Q3d	Q3e	Q3f	Q3g	Q4a	Q4b	Q4c	Q4e
2	4	2	2	2	2	1	3	2	5	5	4	5
Novelty		*Inventive Step*							*Industrial Applicability*			
2		2							4,75			
Overall Score												
8,75												
Normalized Score												
0,479166667												

p_1		p_2
1,05		1,05
EvaluatorScore$_1$		EvaluatorScore$_2$
0,461860614		0,461860614
Patentability Indicator		
46%		

Initiate Discrepancy Resolution:	NO

Internal reviewer evaluation remarks: Questions Q4 are increasing the patentability score that in my opinion is low, since it considers an optimized version of an object detector that is in line with the state of the art. This reduces both novelty and non-obviousness. I therefore applied p1>1. Patentability level: low.

External reviewer evaluation remarks: The owners' assessment that patentability potential is not high makes sense. The novelty is relatively low as the approach relies on standard partial finetuning of YOLOv8n, while non-obviousness is weak, since the energy efficiency gains stem predictably from lightweight model selection, aligning with prior art (Wang & El-Gohary). However, applicability is indeed promising for safety-critical sectors. Patentability level: low to medium.

3.4.7 Strategic Implications and Future Impact

With the implementation of the TALON patentability assessment framework, there is greater efficiency in the allocation of intellectual property assets by pinpointing innovations most likely to be patent protected. This is especially crucial in the scope of European research funding, where measurable commercial outcome is becoming a primary indicator of success. Moreover, the framework's risk management capabilities include not just a patentability assessment but also broader strategic elements, which include but are not limited to, freedom to operate, market competition, and potential for technology transfer. The methodology allows for proactive identification of potential challenges early in the development process, providing protective solutions to avoid expensive issues during patent prosecution or further commercialization efforts.

The rigorous documentation generated as a result of the TALON assessment process supports technology transfer activities. Thematic description through detailed technical documentation, prior art analyses, and assessment of commercial potential equips potential licensees and partners with the information required to evaluate collaboration prospects and negotiate licensing agreements. This documentation is useful in many ways during the technology transfer process starting from partner identification to license negotiation and further moving to relationship management. The participation, role, and information needs of all stakeholders are identified and described systematically so that all have access to high quality information related to the considered innovations.

3.4.8 The Road Ahead: Implementation and Impact

The framework's identifying of ten diverse technical domains and high-value innovations showcases its ability to reveal hidden intellectual property concealed within intricate research endeavors. The systematic evaluation, expert oversight, and thorough record-keeping within the TALON framework enhances successful patent prosecution and technology transfer initiatives. Most crucially, the TALON framework offers a new paradigm to manage intellectual property in research settings, granting researchers systematic IP appraisal strategy balance with the flexibility to accommodate the nature of varying innovations. The shift from ad-hoc evaluation to systematic assessment encourages a move towards better commercial return on investment in European research.

Beyond the scope of TALON project, these developments with greater implications form the framework of managing intellectual property within European research projects. In its complexity and cross-disciplinarity, there will always be a stronger demand for more systematic approaches to evaluating and managing IP. The results of the TALON project enhance the efficacy of the intellectual property framework within Research and Innovation programs, making a substantive impact on the Horizon Europe initiative. As the methodology is applied and refined, it has the potential to become a standard tool for research projects seeking to maximize the overall impact of their innovations while maintaining the highest standards of scientific rigor and professional excellence.

3.5 Partners' Individual Exploitation within TALON and the Way Forward

3.5.1 Engineering Ingegneria Informatica S.p.A.

3.5.1.1 Partner profile

Engineering Ingegneria Informatica S.p.A. is a digital transformation company and a market leader in Italy with an expanding global presence. With around 15,000 employees and more than 80 offices throughout Europe, the United States and South America, Engineering plays a pivotal role in driving digital innovation in key sectors. Its R&D division comprises more than 450 researchers and data scientists, backed by a strong network of academic, industrial, and start-up partners.

Engineering designs develops and operates transformative solutions for sectors including smart industry, digital finance, e-health, augmented cities, utilities and public administration. By leveraging advanced technologies such as artificial intelligence, cloud computing, digital twins, cybersecurity, the Internet of Things (IoT), and big data, Engineering delivers solutions that boost productivity, sustainability, and resilience.

In the TALON project, the Engineering team plays a central role coordinating the project as well as in developing the key technological components that enable trustworthy, energy-efficient and transparent AI architectures. Specifically, the ENG team is leading the development of the AI Zero-Touch AI Orchestrator, which automates the deployment and management of AI workloads across distributed edge-cloud infrastructure. ENG also contributes to big data analytics, orchestration strategies and methods for improving explainability and trust. These competencies align with the Engineering team's broader mission to operationalise AI in critical and industrial systems.

3.5.1.2 Exploitation Actions and Expected Return

ENG team considers the following TALON results to be highly relevant to its strategic and technological roadmap:

- The Zero-Touch AI Orchestrator, which supports the automation and energy-efficient deployment of AI workloads across distributed infrastructures;
- Big data analytics frameworks tailored to high-volume industrial data streams;
- Techniques for explainability and transparency, which are essential for operationalising trustworthy AI in regulated domains;
- Edge/cloud integration models and methods to support sustainable and scalable AI services.

Concrete Exploitation Actions
The ENG team has defined a progressive plan to exploit the results of the TALON project. Activities will begin internally and may later extend externally, with every step coordinated and validated alongside relevant business units. External commitments will only follow once internal alignment and clear interest have been established.

Integration into Internal Platforms and Services
The core outcomes of TALON—particularly the Zero-Touch AI Orchestrator and explainability techniques—are intended to be integrated into existing platforms and sector-specific solutions for Industry 4.0 and utilities. Initial efforts will focus on discussions with business units to ensure alignment with current product roadmaps and technology readiness, followed by technical validation and integration planning. This will address orchestration workflows, trust layers, and energy-aware deployment. These activities will provide the basis for value creation before any market-facing actions are considered.

Internal Proposals for New Services
In parallel, internal proposals will be developed for new services inspired by TALON innovations. Concepts under consideration include Orchestration-as-a-Service and dashboards for explainability

and energy-aware performance. These proposals will be shared with business units to gauge interest and feasibility in client-facing contexts. Depending on the feedback, such services may remain exploratory, advance as research initiatives, or be gradually incorporated into Engineering's portfolio.

Preliminary Market Exploration
If internal evaluation confirms sufficient value and feasibility, ENG may initiate preliminary exploration of external interest, particularly among clients in manufacturing, utilities, and public infrastructure, as well as system integrators and technology partners. This activity will remain informal and opportunity-driven, intended to assess demand for orchestration and trust solutions and to identify possible co-development opportunities or pilots. No external offering will be formalised without a clear business case supported by both internal and client-side discussions.

Follow-up Research
To sustain and evolve TALON's results, ENG plans to pursue follow-up research and development projects. This will allow components such as the orchestrator, trust mechanisms, and analytics engines to be further tested and expanded. Strategic planning is already underway to identify relevant Horizon Europe and DIGITAL opportunities, in collaboration with internal stakeholders and TALON partners.

Target Markets and User Segments
- Industrial AI and smart manufacturing
- Energy and utilities
- Public sector and critical infrastructure
- Cloud and edge service providers
- Regulated domains requiring AI compliance (e.g., finance and healthcare)

Exploitation Scenarios
- Internal validation and integration into existing platforms
- Strategic partnerships for co-development based on TALON technologies
- Participation in European initiatives, positioning TALON as a reference implementation for trusted and efficient AI

Expected Returns
- **Economic:** Strengthened competitiveness of Engineering's AI platforms
- **Strategic:** Enhanced positioning in AI orchestration and industrial automation, aligned with transparency and scalability demands
- **Knowledge-based:** Reinforced expertise in explainable, ethical, and energy-efficient AI, and dissemination of TALON findings across business units
- **Policy/Standardisation:** Stronger role in shaping European AI frameworks and demonstrating practical implementations of trustworthy AI principles

3.5.1.3 Strategic Future Plans and Commitment

Potential integration into product roadmap
The Zero-Touch Orchestrator will be further developed and potentially integrated into ENG's platform to enhance its AI orchestration capabilities and enable dynamic service composition in hybrid environments.

Trust and explainability components will be incorporated into Engineering's AI toolkit to support compliance with future EU AI regulations and enable the responsible deployment of AI in critical sectors.

Collaboration with TALON partners
Engineering intends to explore joint commercialisation opportunities with TALON partners, including pilot extensions and service bundling, based on mutual interest and availability.

Resource allocation and commitment

Post-project, ENG will allocate R&D resources to ensure the continuity of TALON innovations within its industrial sector through business units. Investment in development, especially in the areas of AI/edge orchestration, responsible AI and green computing, will ensure long-term internal capacity.

Strategic alignment
TALON supports Engineering's strategic priorities in the areas of digital transformation, AI governance and Industry 5.0. By operationalising trustworthy, efficient and human-centric AI, the results of TALON directly reinforce Engineering's mission to shape the next generation of resilient, explainable and scalable digital ecosystems for the public and private sectors.

3.5.2 Ericsson Telecomunicazioni SPA

3.5.2.1 Partner profile

Ericsson Telecomunicazioni SpA (TEI) is the Italian company of the Ericsson Group (Telefonaktiebolaget LM Ericsson in Sweden). Ericsson is a worldwide leader for services, software and infrastructures for Information and Communications Technology to telecom Operators and other Industries. TEI R&D division, which is engaged in the project, provides Ericsson with products, technology and methodology to deliver TLC networks, based cloud and virtualization techniques and optical and photonic solutions.

Within the Talon project, TEI supports fiber infrastructure resilience by providing a supervised dataset that includes various fiber fault types and offering domain expertise to validate the AI models used in fault detection. Additionally, Ericsson contributes to enhancing the explainability of fault classification through the application of Explainable AI (XAI) approaches, which help clarify the decision-making process of the models. Besides, TEI develops a module for obfuscating sensible information in image data to ensure that visualization tools and dashboards remain fully compliant with GDPR requirements.

3.5.2.2 Exploitation Actions and Expected Return

TEI's contribution in TALON project focuses on enhancing the explainability and trustworthiness of AI models used for fault classification in optical fiber links within Radio Access Networks (RAN) and other fiber-based connectivity infrastructures.

Optical links, while reliable, suffer from faults due to physical damage or component degradation. Distinguishing these from node/transceiver issues is a persistent operational challenge. TEI addresses this by integrating Explainable AI (XAI) into machine learning models that analyse time-series data (e.g., RX power drops), enabling both local (instance-based) and global (model-level) explanations.

The Image Anonymization Module developed by TEI has been successfully integrated into UC3 – Augmented Reality for Maintenance and UC4 – Human-Robot Collaboration, where it ensures the anonymization of individuals captured in visual data. While this area is not part of Ericsson's core business, the integration outcomes have revealed promising opportunities for applying anonymization technologies in contexts where image data is processed.

These insights are being further explored through internal workshops aimed at identifying potential business extensions and service offerings focused on privacy-preserving solutions. The experience gained within TALON lays the foundation for future initiatives in sectors where ethical data handling and regulatory compliance are critical.

3.5.2.3 Strategic Future Plans and Commitment

Network programmability is strategic for TEI and infrastructure fault classification can be integrated in TEI orchestrator according to O-RAN standardization for rAPPs implementation. The techniques developed in TALON open exploitation opportunities to new future use-cases where sample time series need to be analysed for predictive maintenance as well as diagnostic: such future scenario could be on other optical components faults like transceiver lasers or electronic parts, optical and radio links, synchronization issues, etc.

TEI intends to build upon the integration outcomes of the Image Anonymization Module developed in the TALON use cases, enhancing it with advanced context-aware anonymization techniques and a configurable visual tuning dashboard. These improvements aim to ensure adaptability and compliance with privacy requirements in future initiatives, including internal projects and externally funded research, such as upcoming Horizon Europe proposals.

An internal assessment is also underway to evaluate the potential integration of this module into TEI's proprietary data management portal, with the goal of ensuring the anonymization of stored visual data. This would support broader compliance with data protection policies and further align with TEI's strategic objectives around responsible AI, privacy-preserving technologies, and digital transformation.

3.5.3 Metamind Innovations

3.5.3.1 Partner profile

Metamind Innovations (MINDS) is a research-driven SME specializing in AI, data analytics, and cybersecurity solutions. Within the TALON project, MINDS brings essential expertise in privacy-preserving AI technologies, contributing significantly to the development of explainable and secure AI orchestration systems. The organization's core competencies in few-shot learning, data anonymization, federated learning, and explainable AI directly support TALON's objectives of creating autonomous, self-organized AI systems for Industry 5.0 applications. MINDS participates across multiple work packages, from theoretical framework development to implementation and demonstration, with particular focus on developing XAI methodologies for all trust levels and implementing privacy-preserving techniques that ensure GDPR compliance while maintaining system performance and transparency.

3.5.3.2 Exploitation Actions and Expected Return

MINDS will leverage several key TALON components and assets, particularly the explainability frameworks developed across all trust levels (TrL1-TrL4) integrated into UC2's industrial automation scenarios and the federated learning architecture implemented in UC4 for human-robot collaboration. The company plans concrete exploitation actions including the enhancement of its platforms with TALON's XAI capabilities, the creation of new privacy-preserving AI services for industrial clients, and the development of specialized training programs on trustworthy AI implementation. MINDS targets primarily the manufacturing sector, critical infrastructure operators, and technology integrators seeking GDPR-compliant AI solutions. The hypothetical go-to-market strategy involves strategic partnerships with system integrators, direct sales to industrial clients already in MINDS' portfolio, and licensing of specific anonymization and XAI components to technology vendors. Expected returns may encompass benefits through new streams from XAI-as-a-Service offerings and enhanced competitiveness in the industrial AI market, strategic positioning as main provider of trustworthy AI solutions in the domain, knowledge-based advantages through improved capabilities in edge-cloud AI integration and ethical AI governance.

3.5.3.3 Strategic Future Plans and Commitment

MINDS envisions integrating TALON's results as components of its portfolio over the next 5 years, with plans to maintain and evolve the XAI frameworks and federated learning architectures beyond the project's conclusion. The company may establish a team for post-project development of TALON technologies, allocating resources for continuous improvement and market adaptation. Collaboration with TALON partners may continue, particularly with blockchain partners (SID) for integrated solution offerings. This aligns perfectly with MINDS' strategic mission of bridging academic research and industrial application in AI, supporting the company's vision of becoming a leader in trustworthy AI solutions. The TALON project strengthens MINDS' position in Industry 5.0 transformation, providing the technological foundation and partnership network necessary for sustainable growth.

3.5.4 Netcompany (previously Netcompany-Intrasoft)

3.5.4.1 Partner profile

Netcompany is a European leader in IT services and digital transformation, active across 13+ countries with deep expertise in the public sector, finance, telecom, and utilities. A member of the Netcompany Group, it combines decades of experience with cutting-edge innovation capacity, supported by a 3,200+ strong workforce. Its Research & Innovation Development Department (RID) plays a central role in EU-funded R&D, tenders and internal innovation initiatives. RID offers full-spectrum services across: Software architecture, cloud-native design, microservices, REST APIs, and event-driven systems, DataOps & AI/ML, DevOps/DevSecOp, as well as business exploitation.

In TALON, Netcompany (INTRA) leads two key integration and evaluation tasks: Task 5.1 and Task 5.2, driving both the strategic planning of pilot demonstrations and the technical integration of the TALON platform. As Task 5.1 leader, INTRA coordinates the development of the evaluation methodology, KPI framework, and the detailed experimental plan for all pilot sites. In Task 5.2, INTRA oversees the setup, continuous integration, and maintenance of the TALON platform, ensuring seamless deployment and iterative validation with end-users. Also, INTRA leads the Data Management Plan for TALON.

3.5.4.2 Exploitation Actions and Expected Return

INTRA's exploitation strategy for TALON centers on the practical application of the AI orchestration and integration work developed under WP5. Key outcomes, such as the TALON orchestrator and the cloud/edge integration mechanisms, will be assessed for integration into INTRA's internal AI platforms and reused in follow-up EU projects with higher market proximity (e.g., under HEU or DEP). Planned actions include dedicated internal workshops with AI and integration teams to evaluate technical alignment, foster adoption, and define the next steps for operational use. The primary application domains are EU, public sector and industry, particularly where INTRA already engages in commercial tenders and digital transformation projects. The expected benefits are mainly strategic and knowledge-driven: strengthening INTRA's orchestration capabilities, supporting more flexible edge/cloud deployments, and reinforcing its role in the development of trustworthy, human-centric AI systems.

3.5.4.3 Strategic Future Plans and Commitment

INTRA's medium- to long-term vision is to sustain and extend the integration and orchestration outcomes of TALON by incorporating key components—such as the orchestrator and evaluation frameworks—into its internal AI service portfolio and future digital platforms. TALON results will be considered for inclusion in INTRA's product roadmap, particularly where cloud/edge AI integration and human-centric governance are relevant to public sector or industry projects. INTRA remains open to continued collaboration with TALON partners through future EU projects, commercial alliances, or participation in standardisation initiatives. Dedicated technical and business teams within INTRA's R&D and product units will support the maintenance and evolution of TALON-related assets. This aligns with the company's broader mission to drive responsible digital transformation and contribute to Europe's transition toward trustworthy, scalable AI under Industry 5.0.

3.5.5 Democritus University of Thrace

3.5.5.1 Partner profile

DUTH successfully participated in the TALON project through the Laboratory of Industrial & Educational Embedded Systems (iEES Lab), which is recognized for its work on Industrial IoT, cyber-physical systems, edge-cloud orchestration, and immersive visual analytics. Within TALON, DUTH participated in multiple key work packages, with an emphasis on AI model design, infrastructure orchestration, and trustworthy visualization. Specifically, due to its expertise in AI reverse-task offloading, DUTH assisted in vital tasks, such as the co-design of the Zero-Touch E2C AI Orchestrator, as well as supporting the development of the AI-based resource allocation through computational offloading. Moreover, DUTH provided significant contributions during the patentability process, while enriching TALON's approach to Green Industry, due to its expertise in Digital Circular Economy.

3.5.5.2 Exploitation Actions and Expected Return

Building on its TALON work, DUTH will focus on the exploitation of the adaptive offloading engine and AI orchestrator. First of all, the AI reverse-task offloading will be utilized as a key tool for academic and funded research projects that require the efficient handling of AI resource-intensive tasks. Moreover, considering the educational nature of DUTH, the possibility of creating a master level dedicated tutorial or workshop for proposed best strategies on AI-based task offloading will be investigated, as well as the organization of related events for the attraction of related experts that will promote productive discussions on the technology, potential partnerships, as well as the further development or modifications of the offloading mechanism. Overall, DUTH is focused on the formation of partnerships with academic institutions of similar research interests, focusing on the evolution of task offloading mechanisms, while targeting the industrial sector and specifically multidimensional companies that require such advanced services. A concrete plan for the attraction of such partnerships is the publicization of the related research work, either through academic papers or dedicated presentations, while including practical demonstrations to prove the efficiency of the technology. The expected outcomes of this exploitation from such an approach are the strengthening of DUTH as an academic body and key professional in the field, as well as knowledge-based reward due to the knowledge sharing among related professionals.

3.5.5.3 Strategic Future Plans and Commitment

DUTH's long-term vision is to integrate the TALON outcomes, and specifically the offloading methodologies, as key anchor points for its following cycles of scholarly work, focusing on research related to energy-aware edge AI. The corresponding components will remain active within the academic premises, enabling collaborative studies and innovative AI projects. Consequently, DUTH will pursue competitive European and national grants, cooperate with fellow TALON institutions, and attempt to contribute reference implementations to relevant standardization bodies. Moreover, the exploitation of the specific outcomes will also assist in continued engagement with regional innovation hubs, promoting the transformation of research insights into practical demonstrators, enabling DUTH to become a vital entity for responsible digital transformation and I5.0 research.

3.5.6 Kingston University

3.5.6.1 Partner profile

Kingston University (KU) is a UK-based higher education institution with a strong track record in applied research and innovation, particularly in the areas of digital transformation, artificial intelligence (AI), and immersive technologies. KU has established expertise in AI model development, computer vision, edge/cloud computing, and interactive visualisation tools, positioning it as a key contributor to technological advancement in academic and industrial domains.

In the TALON project, KU contributes to the development of advanced digital twin (DT) and AI-driven solutions, with a focus on real-time virtual reality training and augmented reality maintenance for industrial applications. Leveraging its expertise in Unity-based environments and sensor simulation, KU plays a role in creating scalable, realistic training and maintenance solutions. The team's contributions support the project’s broader goals of enhancing automation, reliability, and interoperability across edge-cloud infrastructures.

3.5.6.2 Exploitation Actions and Expected Return

Kingston University (KU) is particularly interested in TALON results related to AI-powered architectures, edge/cloud integration models, immersive visualisation tools, and trust frameworks. The most relevant components to KU's expertise are those supporting the development of virtual reality (VR) training environments and augmented reality (AR) solutions for industrial maintenance—key enablers for human-centric Industry 5.0 applications.

KU will exploit TALON outcomes through the following actions:

- Integration of VR/AR technologies into ongoing research on immersive training platforms tailored for industrial upskilling and reskilling.
- Development of AR-based maintenance assistance tools, enabling context-aware, real-time support for frontline industrial workers.
- Enhancement of educational and training programmes by incorporating TALON methodologies and tools into teaching content, especially in engineering and human-computer interaction.

KU’s exploitation targets include:

- Manufacturing and maintenance sectors seeking advanced human-machine collaboration tools.

- Training providers and vocational education institutions aiming to modernise their offerings through immersive technologies.
- Research and innovation stakeholders interested in responsible and human-centric digital transformation.

KU plans to reach these markets through:

- Collaborative research and innovation projects with industrial partners focused on immersive training and AR support tools.
- Spin-out opportunities for commercialising specific VR/AR solutions developed during TALON.
- Open-access demonstrators and dissemination at academic and industry-facing conferences and workshops.
- Participation in European and international policy and standardisation initiatives related to XR technologies and Industry 5.0.

Expected Returns:

- Economic: Creation of commercial opportunities through collaborative research and potential spin-offs, while supporting industrial partners in reducing training and maintenance costs.
- Strategic: Strengthened positioning in the emerging field of Industry 5.0, especially in immersive and human-centric systems.
- Knowledge-based: Enhanced capabilities in developing, deploying, and evaluating immersive AI-enhanced XR applications, with a focus on usability and ethical design.
- Policy/Standardisation: Contribution to European efforts in shaping regulations and best practices for trustworthy, interoperable, and human-centred XR solutions.

3.5.6.3 Strategic Future Plans and Commitment

Kingston University (KU) envisions a sustained and strategic integration of TALON outcomes into its research and innovation activities, particularly in the domains of virtual reality training, augmented reality for maintenance, and AI-powered edge/cloud solutions for Industry 5.0.

KU plans to embed TALON results into its ongoing research roadmap, specifically through:

- Continued development and refinement of VR/AR platforms for industrial training and operational support.
- Exploration of follow-up research and innovation proposals under Horizon Europe and UKRI frameworks, focusing on immersive, trustworthy, and human-centric digital technologies.

There is also interest in pursuing further collaboration with TALON partners, particularly in the areas of:

- Joint research initiatives and demonstrators to further investigate commerciality of selected TALON components.
- Formation of interdisciplinary consortia focused on ethical AI, edge computing, and human-machine symbiosis in industrial contexts.

Kingston University aims to continue engaging with TALON outcomes through ongoing academic research and teaching activities. Staff involved in the project are committed to exploring post-project opportunities for further development, including student-led projects, cross-disciplinary collaboration, and engagement with external partners where feasible.

TALON directly aligns with KU's strategic mission to advance responsible digital transformation and promote applied research with societal and industrial impact. The focus on AI ethics, immersive technologies, and Industry 5.0 places TALON outcomes at the heart of KU's priorities, reinforcing its commitment to building future-ready human-centric digital solutions.

3.5.7 Centre for Research & Technology Hellas

3.5.7.1 Partner profile

The Centre for Research and Technology Hellas (CERTH) is one of the largest and most dynamic research organizations in Greece, with a strong reputation in innovation, digital transformation, and technological leadership. Through its Information Technologies Institute (ITI), CERTH conducts cutting-edge research and delivers solutions that address real-world industrial and societal challenges. Its expertise spans artificial intelligence, machine learning, computer vision, data analytics, edge/cloud infrastructures, human–machine interaction, and visualization technologies, establishing CERTH as a trusted partner in European and international R&D initiatives.

Within the TALON project, CERTH serves as a pilot site focusing on human–robot collaboration in industrial environments. The organization designs, implements, and validates scenarios where robotic systems (drones) operate seamlessly alongside human workers, ensuring both efficiency and safety. Leveraging CERTH's, as well as the consortium's expertise in computer vision, AI-driven human activity recognition, pose estimation, and multimodal data analytics, methods are developed in the TALON project to help monitoring, interpreting, and possibly predicting human actions in real time. These capabilities enable supporting collaborative tasks in flexible manufacturing and logistics processes between workers (humans) and robots (drones).

In addition, CERTH contributes to the integration of these human–robot interaction modules within the project's edge/cloud infrastructure, as well as to the evaluation of performance through Key Performance Indicators (KPIs) and real-world pilot demonstrations. This role ensures that TALON delivers practical, human-centric solutions that align with Industry 4.0 requirements.

Furthermore, the organization is actively involved in performance evaluation, discussions and actions to support standardisation activities, as well as dissemination efforts, ensuring that TALON's outcomes are not only technically robust but also aligned with international standards and accessible to the broader research and industrial communities.

3.5.7.2 Exploitation Actions and Expected Return

CERTH team considers the following TALON results and assets to be highly relevant to its strategic and technological roadmap.

CERTH will exploit several outcomes of the TALON project, including the AI-powered orchestration architecture, human–robot collaboration use case implementations, edge/cloud integration models, advanced visualization and monitoring tools, and explainability techniques for trustworthy AI.

These results align with CERTH's core expertise in artificial intelligence, computer vision, multimodal data fusion, and human–machine interaction, strengthening its role as a technological leader in Industry 4.0 solutions.

Concrete Exploitation Actions:

- Possible integration of some of the TALON technologies into CERTH's internal research platforms for AI, robotics, and edge/cloud infrastructures.
- Development of new demonstrators and examples for human–robot collaboration that can be showcased to industrial partners.
- Contribution to training programs and workshops targeting both academic and industrial stakeholders, leveraging TALON results as case studies.
- Possible contribution in standardization activities on AI, robotics, and trustworthy human–machine interaction.

Target Markets and Domains:
CERTH will focus on industrial automation, advanced manufacturing, and smart factories as primary domains. Target user groups include industrial companies, system integrators, robotics manufacturers, and technology providers seeking AI-enhanced solutions for safe and efficient human–robot collaboration.

Exploitation Scenarios:

- Partnerships with industrial stakeholders in Greece and Europe to transfer TALON solutions into real manufacturing settings.
- Potential use of TALON pilots (at least presentation of the CERTH pilot (UC4)) as reference cases to attract new collaborations in Horizon Europe and national R&D programs.
- Potential support of spin-off activities of other partners among the consortium, in specialized areas, such as AI-driven human activity recognition and edge/cloud-based orchestration.

- Possible joint development agreements with technology providers and robotics integrators.

Expected Returns:

- Strengthened capacity to attract industrial funding and European projects; potential revenue from licensing AI/robotics components.
- Enhanced positioning as a leading European research center in AI, robotics, and Industry 4.0; expanded collaborations with industry and policy makers.

3.5.7.3 Strategic Future Plans and Commitment

Medium- to Long-Term Vision:
CERTH envisions sustaining and further advancing the outcomes of TALON by embedding AI model development, human–robot collaboration, edge/cloud orchestration, and visualization tools into its long-term research and innovation agenda. The goal is to ensure that TALON serves as a catalyst for new research directions, industrial collaborations, and technology transfer opportunities in the domain of Industry 4.0 and Industry 5.0.

Integration in Product and Service Roadmap:
The TALON results could be incorporated into CERTH's internal research platforms and demonstrators, which are regularly updated and presented to industrial stakeholders and partners. Human activity recognition, adaptive robot behaviour, and explainable AI modules developed in TALON will form the basis for future solutions offered to the manufacturing sector. Additionally, CERTH intends to integrate TALON's orchestration mechanisms into ongoing and upcoming Horizon Europe and national R&D projects, ensuring continuity and scalability.

Post-Project Collaborations:
CERTH is committed to continuing cooperation with TALON partners beyond the project duration. This may include joint commercial initiatives, joint participation in future European R&D consortia, contribution to standardisation bodies, and exploration of joint ventures for industrial pilot deployments. Strong synergies are foreseen particularly with industrial partners and technology providers to co-develop market-ready solutions.

Resource Allocation and Institutional Commitment:
CERTH allocates dedicated resources to ensure technology sustainability and exploitation after project completion. TALON-related activities could be supported by dedicated research teams, as well as the CERTH's innovation ecosystem, which connects academia with the industrial sector.

Alignment with Strategic Priorities:
TALON aligns closely with CERTH's strategic mission to advance digital transformation, AI governance, and Industry 5.0 solutions. By focusing on human-centric AI, trustworthy and explainable systems, and seamless human–robot collaboration, CERTH ensures that TALON results contribute both to scientific excellence and to the socio-economic priorities of the European Union, including sustainability, competitiveness, and responsible AI deployment.

3.5.8 University of Luxembourg

3.5.8.1 Partner profile

University of Luxembourg (UL) is an international research university with distinctly multilingual and interdisciplinary character. The University was founded in 2003 and counts more than 6,700 students and more than 2,000 employees from around the world. The University's faculties and interdisciplinary centres focus on research in the areas of Computer Science and ICT Security, Materials Science, European and International Law, Finance and Financial Innovation, Education, Contemporary and Digital History. In addition, the University focuses on cross-disciplinary research in the areas of Data Modelling and Simulation as well as Health and System Biomedicine. Times Higher Education ranks the University of Luxembourg #3 globally for its "international outlook," #20 in the Young University Ranking 2021, and among the top 250 universities worldwide. Notably, in Telecommunication Engineering, the University of Luxembourg is ranked among the top 75 universities globally, according to the Shanghai Ranking 2024. SnT is a leading international research and innovation center in secure, reliable and trustworthy ICT systems and services. We play an instrumental role in Luxembourg by fueling innovation through research partnerships with industry, boosting R&D investments leading to economic growth, and attracting highly qualified talent.

In the TALON project, the UL team brings expertise in AI models development, contributing significantly to task 3.4 "development of the AI Zero-Touch AI Orchestrator", which automates the deployment and management of AI workloads across distributed edge-cloud infrastructure through a multi-task machine learning model. Also, the UL team contributes to task 4.1 "data enrichment task," using images and time series data.

3.5.8.2 Exploitation Actions and Expected Return

Relevant TALON results and assets

- The AI Zero-Touch Orchestrator (Task 3.4), a multi-task learning–based orchestrator for automated deployment and lifecycle management of AI workloads across distributed edge–cloud infrastructures.
- The Data Enrichment solution (Task 4.1), leveraging both image and time-series modalities to enhance the robustness and adaptability of AI models.

Concrete exploitation actions

- Implementation and validation of a data enrichment model for fire detection in the images for MINDS images.
- Integration of TALON results into research and training activities at UL, ensuring technology transfer to students, researchers, and industrial partners

Target markets, domains and user groups

- Industrial partners.
- Academic and training communities to leverage advanced AI models.

3.5.8.3 Strategic Future Plans and Commitment

The long-term vision is to consolidate the expertise gained in zero-touch AI orchestration and multimodal data enrichment into enduring research, innovation, and technology transfer activities.

UL will continue advancing the AI Zero-Touch Orchestrator into a scalable research prototype.

Building on TALON's multimodal data enrichment pipelines, UL plans to extend this research towards domains such as cybersecurity, smart mobility, and satellite-terrestrial communications, ensuring cross-domain applicability.

UL is open to maintaining collaborations with consortium partners beyond the project's lifetime, especially in areas such as joint publications, follow-up Horizon Europe proposals, standardisation activities and potential joint training initiatives.
The outcomes of TALON will be incorporated into teaching programs, doctoral training units, ensuring long-term exploitation.

3.5.9 Universitat Politecnica De Valencia

3.5.9.1 Partner profile

The Universitat Politècnica de València (UPV) is one of the top technical universities in Spain and Europe, consistently recognised for its excellence in innovation, technology transfer, and industrial collaboration. UPV is ranked as the leading Spanish technical university in the Shanghai Academic Ranking of World Universities (ARWU), reflecting its global reputation for research quality and impact. Within UPV, the Research Centre on Production Management and Engineering (CIGIP) provides advanced expertise in industrial engineering, digital transformation, logistics, and sustainable operations. CIGIP focuses on artificial intelligence for production systems, edge-to-cloud architectures, optimisation and simulation models, predictive analytics, and digital twins, aligning directly with the TALON vision.

In TALON, UPV plays a key role in orchestration and resilient architectures for distributed industrial networks. Specifically, UPV leads the Self-Healing & Self-Correcting Mechanisms, developing proactive mechanisms for reliability and robustness of the edge-to-cloud infrastructure. This includes the design and validation of a full monitoring-orchestration loop, leveraging Kubernetes-based orchestration, Grafana/Prometheus for real-time monitoring, and Node-RED for automated reconfiguration flows, ensuring resilient and energy-efficient industrial deployments. UPV also leads in Few-Shot Learning Approaches, applying advanced meta-learning techniques to enable accurate predictions with limited data, with focus on tool wear estimation and tool change prediction in industrial environments. Finally, UPV contributes to the overall E2C orchestration layer and to the evaluation and demonstration in the I5.0 automation and planning use case.

3.5.9.2 Exploitation Actions and Expected Return

Relevant TALON results and assets:

- Self-healing and self-correcting mechanisms toolkit (WP3).
- Few-shot model to predict tool changes (WP4).
- Edge-to-Cloud orchestration modules leveraging Kubernetes-based container orchestration, enabling autonomous, scalable, and energy-efficient deployments.
- Industrial demonstrator for I5.0 automation and planning, integrating orchestration, monitoring, and predictive analytics.

Concrete exploitation actions:

- Deployment of an industrial edge-to-cloud cluster at FACTOR, fully managed with Kubernetes, supporting autonomous orchestration and monitoring loops (Prometheus/Grafana, Node-RED).
- Implementation and validation of few-shot learning algorithms for tool wear prediction and tool change forecasting in the FACTOR use case.
- Integration of TALON results into research and training activities at UPV, ensuring technology transfer to students, researchers, and industrial partners.

Target markets, domains and user groups:

- Industrial manufacturing.
- Academic and training communities requiring advanced AI/edge/cloud integration capabilities.

3.5.9.3 Strategic Future Plans and Commitment

UPV's medium- and long-term vision is strongly aligned with the consolidation of digital transformation and Industry 5.0 principles in industrial ecosystems.

Specifically, TALON results will be:

- Embedded in future research projects at national and European level, ensuring continuity of knowledge and incremental improvements on resilience, sustainability, and trustworthy AI.
- Integrated into educational programmes at UPV, where CIGIP will expand training modules on AI for industrial engineering, edge/cloud computing, and digital twins, transferring project know-how to students and professionals.
- Leveraged for collaboration with industrial partners in Spain and Europe, enabling pilot extensions and real-world validations beyond TALON.
- UPV is open to maintaining collaborations with consortium partners beyond the project's lifetime, especially in areas such as joint publications, follow-up Horizon Europe proposals, standardisation activities and potential joint training initiatives.
- The institution is committed to allocate resources through CIGIP's research staff and infrastructure to ensure that TALON knowledge and frameworks remain active and are further developed after the project. This ensures that the outcomes remain operational and evolve in line with UPV's strategic priorities of advancing Industry 5.0 innovation, digital transformation, and trustworthy AI.

3.5.10 Eight Bells Ltd

3.5.10.1 Partner profile

Eight Bells Ltd (8BELLS) is an independent Research and Consulting company specializing in selected parts of Information and Communication Technologies (ICT) as well as in modelling and analysis for businesses, based in Nicosia, Cyprus.

The main areas of research and technical expertise are customizable solutions that enhance modern communications relevant to the area of 5G Mobile Technology, Network Function Virtualization (NFV) and management solutions for Cloud infrastructures. The company also has a strong background in hardware design and embedded software development comprising a large set of custom solutions focused in modern IoT applications, addressing industry needs.

In the consulting domain, 8BELLS delivers special advisory services in ICT that help its clients understand the market dynamics and profit from the ever-changing landscape. Working side by side with its clients allows us to better understand their business requirements and using a data-driven approach we create tailor-made solutions to support them.

Related expertise / experience:
The analysts and researchers of 8BELLS have a strong focus on R&D with a long track record in EU- and national- funded research initiatives. They have actively participated in several R&D projects in the field of networking and telecommunications and have authored or co-authored over 80 papers in international refereed journals and conferences. Moreover, the company's portfolio includes system engineering and integration services, embedded hardware design and embedded software development, network design and customized innovative solutions.

Our technical capabilities also include:

- Systems & Networks engineering: Including design, implementation and integration,
- Hardware design: Microcontroller-based Printed Circuit Board (PCB) designs for IoT devices targeting small or large scale, low-cost data acquisition. Customizable solutions for sensor interfaces, signal conditioning as well as different Wireless Sensor Network protocols and topologies (WiFi, Zigbee, LoRa, etc.).
- Cloud Computing and Everything-as-a-Service: Design, deployment and maintenance of cloud computing infrastructure, Optimised resource provisioning in Infrastructure-as-a-Service architecture, Platform or Software as- a-Service for quick implementation and instantiation of cloud services,
- Privacy, Security & Data Protection: Cybersecurity Gap Analysis, Virtualized Cybersecurity, Privacy Impact Assessment, Privacy-by-design architectures, ethical and legal compliance of project activities, including compliance to the new General Data Protection Regulation,
- Software development: for 5G security, wireless network security, software quality assurance, machine learning applications, data analysis and visualization, cloud computing, etc.

As part of the TALON consortium, 8BELLS contributed in a number of distinct and meaningful ways, building on its experience and competencies.

At the technical level, 8BELLS leads the task and efforts for the creation of a Smart Pricing Simulator, that plays a significant role in efficiently managing the allocation of resources in TALON's AI orchestrator, and is, integrated to TALON's User Interface (UI). The SPS aim has been to support the efficient allocation of computational resources in four (4) distinct ways: (i) By ensuring that the utilization of TALON's edge computing network charges reflect the computational requirements of the user, (ii) by allowing – through a pay-as-you-go scheme – participants to retain full authority of their costs upfront, encouraging them to utilize TALON's energy-efficient edge computing network, (iii) by facilitating the use of non-commercial devices for edge computing tasks, through effective workload allocation, and (iv) by establishing an incentivization scheme that allows users to opt-in to more energy efficient choices that reinforce the project's overall goal for a greener and more efficient edge-computing network.

In addition, 8BELLS led the consortium-wide efforts in Dissemination, Exploitation and other similar initiatives enhancing TALON's impact. As such, it has ensured that in cooperation with all project partners TALON's results and highlights get the visibility and publicity they deserve. In addition, by coordinating various initiatives and activities enabling the project to maximise its impact.

3.5.10.2 Exploitation Actions and Expected Return

From the perspective of components and tools created as part of the project, 8BELLS' main focus lies in the Smart Pricing Simulator, which has been developed thanks to the participation in TALON. This tool conceived and designed to enable dynamic pricing and resource allocation optimisation in environments where efficient management of limited resources is crucial.
In terms of further exploiting the result, 8BELLS envisions a number of concrete ways that this will be made:

- Further development and validation of the Smart Pricing Simulator and testing in real-world pilots
- Deployment in future EU projects to strengthen maturity, visibility, and validation in diverse domains.
- Expansion of the service portfolio by embedding the simulator into consulting and decision-support services offered to partners and clients.

In terms of the markets that can be addressed these include a number of potential targets, where there is a clear need of resource optimisation, resource scarcity cases and price driven incentivisation. Such cases include:

- Energy and Utilities: Dynamic energy pricing, demand-response optimisation, and resource allocation.
- Telecommunications and ICT: Network resource allocation, spectrum sharing optimisation.
- Smart Cities & Mobility: Transport resource management, dynamic charging, and mobility-as-a-service pricing models.
- Industrial stakeholders: Manufacturing and logistics processes where operational costs depend on resource allocation efficiency.
- Research and Innovation community: As early adopters and partners in R&I initiatives exploring next-generation AI-powered optimisation.

Moving further, a number of points as part of the Exploitation scenario, are defined:

- Go-to-market: Prioritise partnerships and joint engagements, to position the simulator as a decision-support service, initially offered via consultancy and pilot projects.
- Partnerships: Collaborate with industrial players (e.g., energy providers, transport operators) to pilot the tool in real operational environments.
- Licensing/Integration: Explore integration of the simulator into existing platforms of partners, with licensing models adapted to usage.
- EU projects & Networks: Use ongoing participation in R&I projects as a steppingstone to increase TRL and ensure visibility in new domains.

When it comes to the expected returns to be had, these may come in a number of different ways, as follows:

- Economic returns: Creation of new revenue streams through consultancy, licensing, and SaaS-based access to the simulator; increased competitiveness of existing 8BELLS services by embedding optimisation features.
- Strategic returns: Strengthening 8BELLS' positioning in the AI-driven optimisation and industrial automation market; broadening the company's portfolio to include resource allocation decision-support tools
- Knowledge-based: Enhanced expertise in AI-powered pricing models, edge/cloud integration, and ethical AI considerations; development of advanced methodologies transferable across multiple domains (energy, telecom, logistics).
- Policy: Active contribution to European initiatives on trustworthy and explainable AI

3.5.10.3 Strategic Future Plans and Commitment

As a dynamic, independent research and consulting company specializing in growing Information and Communication Technologies (ICT), 8BELLS has a very active interest in AI and its applications, especially those that align with real-world problems, especially in the technology domains and ecosystems that it is involved in. As such applications of AI in industry (I4.0 / I5.0) present an excellent case.

8BELLS has a clear objective to pursue the takeaways and results coming out of TALON. As mentioned above priority will be given to further advancing and validating the model by including it to further research initiatives and opportunities. The ultimate objective being to combine it with solutions existing in its portfolio, plus applying it to additional domains.

In addition to the development of specific technology components, through leadership of WP6 on Dissemination, Standardisation, Regulation & Business Planning, 8BELLS has gained valuable knowledge and deepened its understanding in AI-powered solutions, trends in explainable AI, and their specific market perspectives in the EU. This knowledge and enhanced understanding will be incorporated in the company's technology consulting services so as to be consolidated as appropriate into the relevant European technology ecosystem of projects.

Finally, 8BELLS is also very interested to further the excellent working relationship forged through this project and explore partnering opportunities with TALON partners. The contacts already established during the project will be approached and engaged especially given the fact that the simulator has increased possibilities to be used in combination with other assets, components, systems etc.

3.5.11 UBITECH

3.5.11.1 Partner profile

UBITECH is a leading software and research company specializing in advanced ICT solutions, with strong positioning in innovation, digital transformation, and technological leadership across Europe. The organisation actively contributes to large-scale EU-funded research and innovation projects, focusing on next-generation networking, artificial intelligence, cloud-edge continuum, cybersecurity, and digital infrastructures. UBITECH combines applied research with practical system integration, driving forward digital transformation by providing novel AI frameworks, programmable networking solutions, and edge-cloud orchestration platforms. Its leadership in research consortia highlights its role in shaping future technological standards, advancing interoperability, and bridging the gap between academic innovation and industry deployment.

Role in TALON Project:
Within TALON, UBITECH assumes a key role across several tasks that underpin the project's foundations in AI-driven network programmability, distributed intelligence, and energy-efficient orchestration.

- Task 2.1 (Leader – UBITECH): Defined the use cases, KPIs, system requirements, and technology enablers for TALON, ensuring alignment between end-user needs and the system's design. Our contribution here focused on requirements engineering and the translation of AI-driven benefits into measurable KPIs that guide subsequent technical developments.
- Task 2.2 (AI Theoretical Framework – Led by UBITECH): Contributed to the development of the AI theoretical framework by applying information-theoretic and game-theoretic principles to assess and explain algorithmic behaviour. Our competencies in AI model development and algorithmic analysis were directly applied to formulating the framework that informs decision-making in the TALON architecture.
- Task 2.3 (Performance Assessment – Led by UBITECH): Supported the construction of digital twin (DT)-based tools to assess AI performance and boost transparency and explainability. Here, we contributed to designing evaluation methods and visualization tools that bridge AI outcomes with real-world validation, ensuring trustworthiness of the deployed models.
- Task 3.2 (NG-SDN & Distributed Intelligence – Participant): Played a significant role in the definition and design of NG-SDN functionalities and distributed intelligence mechanisms. Our contributions involved exploring programmable data plane technologies (P4, XDP) and their integration with reinforcement learning-based decision models, enabling end-to-end programmable control of TALON's infrastructure. Additionally, we supported the design of AI-driven orchestration strategies for energy-efficient offloading of VNFs to programmable switches, aligning with the project's goals of maximizing performance while minimizing energy consumption.

3.5.11.2 Exploitation Actions and Expected Return

UBITECH's individual exploitation plan focuses on leveraging key TALON results, including the NG-SDN and distributed intelligence toolkit, the AI theoretical framework, edge/cloud orchestration models, and explainability tools for trustworthy AI. These assets will be exploited through integration into UBITECH's internal R&D platforms, the development of new energy-aware orchestration services, and active engagement in standardisation activities. The targeted markets include telecom operators, smart city infrastructures, Industry 4.0 stakeholders, and cloud-edge service providers across Europe. To reach these markets, UBITECH will pursue partnerships with industrial players, licensing opportunities for orchestration toolkits, and delivery of tailored enterprise solutions. The expected returns encompass new revenue streams and improved product competitiveness (economic), stronger positioning in AI-driven networking and edge intelligence (strategic), enhanced AI/SDN expertise (knowledge), and greater influence in EU standardisation and trustworthy AI frameworks (policy).

3.5.11.3 Strategic Future Plans and Commitment

UBITECH's medium- to long-term vision is to sustain and build upon the TALON outcomes by embedding its NG-SDN toolkit, distributed intelligence functionalities, and AI orchestration models directly into our product roadmap and service offerings. These results will enhance our existing platforms for programmable networking, energy-efficient edge/cloud integration, and explainable AI services, ensuring continuity and scalability after the project's completion. We plan to maintain close collaboration with TALON partners through joint commercialization efforts, participation in standardisation consortia, and potential joint ventures to accelerate the uptake of programmable and AI-native networking solutions. UBITECH is committed to dedicating internal R&D resources and institutional support to exploit TALON's outputs, ensuring they are matured, maintained, and offered to our customer base in telecom, Industry 4.0, and smart city markets. Ultimately, TALON aligns with UBITECH's strategic priorities of driving digital transformation, advancing trustworthy AI and network programmability, and positioning ourselves as a leader in Europe's transition toward Industry 5.0 ecosystems.

3.5.12 Sidroco Holdings Ltd

3.5.12.1 Partner profile

SIDROCO Holdings Ltd is an innovative SME positioned at the intersection of AI, blockchain,and digital transformation, with a strong track record in EU-funded R&D and the delivery of innovative solutions. Leveraging expertise in federated learning, trustworthy AI governance, blockchain-based provenance and advanced visualization, SIDROCO drives technological leadership by transforming cutting-edge research into scalable tools for critical domains. Within TALON, SIDROCO pursuesthe exploitation of FabricVision, a Hyperledger Fabric–based platform that anchors federated-learning workflows with IPFS storage and on-chain auditability. By combining its competencies with TALON's AI orchestration and edge/cloud integration, SIDROCO ensures TALON's results are translated into robust, interoperable, and commercially viable solutions, directly contributing to the project's goals of secure, explainable and verifiable AI operations.

3.5.12.2 Exploitation Actions and Expected Return

SIDROCO will continue building on the developed TALON outcomes and plans to exploit and further mature FabricVIsion. Exploitation actions concentrate on turning FabricVision from research prototype into production-ready tool. SIDROCO will harden FabricVision with TALON's explainability and trust controls; package vertical extensions (e.g., compliance analytics, secure aggregation, policy enforcement); and integrate edge/cloud deployment blueprints for Kubernetes-ready rollouts. In parallel, SIDROCO will engage in pre-standardisation work on federated AI orchestration combined with blockchain-based provenance in accordance with EU trustworthy-AI initiatives.

The initial market focus targets domains where traceability and collaboration are critical: (i) finance/DeFi for auditable, privacy-preserving model exchange and risk analytics, (ii) agriculture for cross-farm data sharing and co-training that improves yield and sustainability, (iii) logistics and supply chains for end-to-end traceability and SLA-grade audit trails and selectively, (iv) healthcare and other critical-infrastructure settings where immutable and secure AI is mandatory. Primary users may include data/ML platform teams, compliance/risk officers and operations leaders in mid-to-large enterprises as well as public operators.

Go-to-market will combine direct B2B integrations with strategic partnerships (system integrators, domain ISVs, cloud providers). SIDROCO will prioritise lighthouse pilots that convert into multi-year licenses, cultivate an ecosystem of integrators through SDKs/APIs and documentation, and—where justified by traction in a vertical—consider a dedicated spin-off entity. Participation in EU ecosystems (e.g., AI-on-Demand, GAIA-X working groups) will amplify visibility and accelerate interoperability.

Expected returns may be economics, strategy, knowledge and possibly policy. Economically, SIDROCO anticipates new revenue from platform licensing, SaaS subscriptions and professional services, while improving product competitiveness through a differentiated AI-plus-blockchain trust layer. Strategically, the company strengthens its position as an EU SME leader in federated AI attestation attracting diverse Horizon Europe projects by employing diverse blockchain applications for various domains. On the policy/standardisation front, active contributions to European discussions on trustworthy AI, federated learning, and blockchain provenance will position SIDROCO to help shape emerging norms and technical guidelines.

3.5.12.3 Strategic Future Plans and Commitment

SIDROCO's medium- to long-term goal is to turn TALON's FabricVision outcome into a durable product for verifiable, privacy-preserving AI operations in regulated domains. A rolling roadmap with quarterly releases, SDKs/APIs and reference deployments will ensure interoperability, governance-by-design and reliable operations aligned withEU AI-governance needs.

Beyond the project, SIDROCO will pursue joint pilots and go-to-market initiatives with TALON partners, explore joint ventures or licensing where traction warrants and engage in Hyperledger, and relevant standardisation bodies to translate insights into practices. This continoously evolving plan advances our mission to fuse AI with robust distributed ledger technologies, expanding our portfolio from traditoinal AI into explainable, verifiable and tamperproof AI for critical domains.

3.5.13 InnoCube

3.5.13.1 Partner profile

InnoCube focuses on technology innovation, network simulation, distributed intelligence systems, edge to cloud intelligence, as well as specializing in network simulation, intelligent solutions for distributed systems, edge-to-cloud intelligence, and AI-powered network orchestration. InnoCube's work includes the AI model and network simulation framework for the TALON project and developed edge-to-cloud federated and distributed learning methods. InnoCube's core strengths include distributed intelligence and network orchestration also served to augment TALON's AI powered architecture and trust frameworks. Furthermore, InnoCube led TALON's communication and exploitation activities, which enhanced its expertise in project management, commercialization strategy, and stakeholder engagement. Intensive work towards these project areas helped refine InnoCube's skills in cross-sector collaboration and responsible governance of AI, while accelerating TALON's mission to create secure, transparent industrial solutions.

3.5.13.2 Exploitation Actions and Expected Return

Within TALON, InnoCube is most interested in advanced edge-to-cloud intelligence architecture, cyber security algorithms, and tranfer learning strategies for application in multiple domains. Moreover, InnoCube has created extensive market analysis, stakeholder engagement, and commercialization frameworks owing to our leadership in communication and exploitation activities, which strengthens its position to bring innovative solutions to the market.

As part of the TALON project, some outcomes were integrated into commercial offerings which led to exceeding preliminary exploitation targets. The edge-to-cloud AI architecture and the distributed learning approaches were considered and connected with InnoCube's flagship platform InnoNET; thus, enhancing real time analytics for industrial clients in multiple network deployments. This innovative solution is anticipated to drive an additional 30% growth in revenue and acquisition of new enterprise contracts by 2026.

InnoLEARN, an AI library developed by InnoCube, was augmented within TALON to offer intelligent components that follow GDPR and EU AI Act compliance. The risk management framework along with the other explainable federated and decentralized learning approaches were also incorporated into the product. It is focused on industrial operators and compliance officers and is backed by client training programs. InnoLEARN strengthens InnoCube's strategic position in ethical AI governance fortifying trust among clients while creating new streams of revenue.

Our target markets comprise industrial automation and manufacturing, energy management and smart grid operators, cloud service providers and edge computing vendors, and consultancies in cybersecurity and risk management. Our exploitation scenarios, enhanced from our leadership of TALON's communication and exploitation activities, include enhanced service refinements with advanced stakeholder engagement to existing clients, direct service sales, strategic partnerships and alliances with cloud and industrial vendors, flexible commercialization-licensed strategies, consulting on go-to-market strategies, and project management.

Expected returns include economic benefits such as revenue increase of 25-30%, 15% market share growth, 40% client base growth, and reduced operational costs by 20%. Strategic returns include market leader position in edge to AI cloud solutions, expansion of service portfolio, and gaining a competitive edge on AI governance. Knowledge-based returns consists of competences in advanced AI/edge/cloud integration, algorithm proprietary development, and increased R&D competencies. Active engagements in the Europe development on regulation of trustworthy AI policy earns participation return, standardization return on regulation development.

3.5.13.3 Strategic Future Plans and Commitment

Post-TALON, InnoCube will focus on leveraging their newfound communication and exploitation skills towards impact scaling on a global level. In line with this, InnoCube will integrate TALON project outcomes into their product roadmap extending through 2027. This will be executed in three phases: Phase 1 in 2025 will concentrate on integration of TALON onto the InnoCube platform; Phase 2 in 2026 will launch the AI governance suite, and Phase 3 in 2027 will focus on development of industry specific solutions.

Our collaboration after the project involves creating an Alumni Network for further engagements in state-of-the-art AI projects, engagement in regulative AI standardization with European partners in the consortium, and significant resource commitment including dedicated engineers, annual R&D investment, and a specialized testing infrastructure.

The InnoCube mission is accelerated through TALON outcomes, positioning InnoCube on a strategic path to lead the digital transformation of industry with readiness for Industry 5.0, excellence in governance of AI, self-reliance of European technologies, and sovereignty of Europe. InnoCube aims to establish itself as a leading provider in communication, networking, and intelligent solutions for industrial applications in Europe by 2030.

3.5.14 Factor Ingeniería y Decoletaje

3.5.14.1 Partner profile

FACTOR Ingeniería y Decoletaje, S.L., established in 1982 and based in Puçol, Valencia, is a leading company in precision machining, metrology, and smart manufacturing. With a strong track record in demanding sectors such as aerospace, automotive, defence, rail, and energy, the company has built its reputation on technical excellence, adaptability, and continuous improvement. In recent years, FACTOR has positioned itself at the forefront of innovation and digital transformation by testing emerging technologies such as artificial intelligence, advanced automation, augmented reality, and digital platforms in R&D projects. This strategic orientation underscores its commitment to sustainability, competitiveness, and technological leadership in the industrial ecosystem. Within the TALON project, FACTOR plays a pivotal role as the host of the "I5.0 automation and planning" demonstrator, leveraging its expertise in precision machining and data-driven manufacturing environments. By deploying advanced automation, CI/CD methodologies, and data orchestration in its production processes, FACTOR ensures reliable data capture, processing, and analysis, which are essential for validating TALON's edge-to-cloud orchestration and AI-driven planning solutions. Its contribution not only provides a real-world industrial testbed for the project but also aligns with TALON's goals of improving efficiency, transparency, and sustainability in next-generation manufacturing.

3.5.14.2 Exploitation Actions and Expected Return

For FACTOR, several TALON results are directly relevant to its strategic objectives, particularly the AI-powered orchestration architecture, the edge-to-cloud integration models, and the explainability and visualization tools that enable transparent and trustworthy decision-making in industrial environments. The use case implementation of the I5.0 automation and planning demonstrator is of particular importance, as it allows FACTOR to validate advanced data capture, orchestration, and planning systems within its machining and metrology processes. Exploitation actions will focus on the enhancement of FACTOR's internal digital platforms, integrating TALON components into its production workflows to improve process automation, predictive planning, and overall manufacturing intelligence. Complementary actions will include training programs for staff to strengthen digital skills, exploring opportunities for IP protection around adapted methodologies, and engaging with European standardisation initiatives on industrial AI and trustworthy automation.

The target markets for Factor include high-demand sectors such as aerospace, automotive, energy, and railway industries, as well as broader manufacturing domains that require precision machining with robust digital integration. FACTOR aims to serve OEMs and tier-1 suppliers seeking partners with strong technological capabilities and the ability to integrate AI-driven automation into complex supply chains. The exploitation scenarios will combine a direct go-to-market strategy, where TALON-enabled capabilities enhance FACTOR's competitiveness in tendering and customer acquisition, with strategic partnerships across the industrial ecosystem. Over time, FACTOR may also explore licensing opportunities or collaborative service offerings built upon TALON technologies.

The expected returns are both economic and strategic. Economically, FACTOR anticipates improved competitiveness of its manufacturing services, greater efficiency, and the potential to open new revenue streams through advanced digital service offerings. Strategically, the project strengthens FACTOR's positioning in the European market as a technologically advanced partner in AI-enabled industrial automation, expanding its service portfolio and reinforcing its role in collaborative R&D initiatives. On the knowledge side, participation will deepen FACTOR's expertise in AI, edge-cloud integration, and explainability, building internal capacities that will support future innovation. Finally, through engagement in TALON's dissemination and standardisation efforts, FACTOR expects to contribute to European dialogues on ethical and trustworthy AI, thereby influencing regulatory and policy developments while aligning its practices with upcoming frameworks.

3.5.14.3 Strategic Future Plans and Commitment

FACTOR's medium- to long-term vision is to consolidate its role as a leader in precision machining and smart manufacturing by embedding the outcomes of TALON into its operational model and service offering. Beyond the project's lifetime, FACTOR plans to sustain and scale the use of TALON's AI-powered orchestration tools, edge-to-cloud integration models, and explainability frameworks to further enhance automation, predictive planning, and data-driven decision-making in its production processes. These results will be progressively integrated into FACTOR's product roadmap, enabling the company to provide its customers with higher levels of process efficiency, reliability, and sustainability, while also opening the door to new service lines in digital manufacturing consultancy and industrial AI solutions.

The company is committed to maintaining and expanding collaboration with TALON partners after the project ends, exploring joint ventures, co-development initiatives, and participation in European standardisation and policy-shaping activities related to trustworthy AI and Industry 5.0. FACTOR sees long-term value in building an innovation ecosystem where knowledge and technologies developed in TALON can be leveraged to serve broader industrial markets and contribute to the digitalisation of European manufacturing.

To ensure sustainability, FACTOR will allocate dedicated resources to post-project exploitation, including strengthening its Innovation Department with expertise in AI and data engineering, investing in digital infrastructure, and fostering workforce upskilling programs. This institutional commitment reflects FACTOR's strategic priorities of advancing digital transformation, promoting ethical and explainable AI in industry, and reinforcing its competitive position through technological leadership.

Ultimately, TALON aligns seamlessly with FACTOR's mission to combine technical excellence with innovation, sustainability, and social responsibility. By integrating TALON outcomes into its strategic roadmap, FACTOR not only enhances its industrial capabilities but also contributes to shaping the future of European manufacturing under the principles of Industry 5.0.

3.5.15 EXOS Solutions

3.5.15.1 Partner profile

EXOS Solutions (EXOS) is a UPV spin off specialized in Manufacturing Execution Systems/Manufacturing Operations Management (MES/MOM). It has its own MES system, MESView, which offers a platform that allows monitoring and control of industrial operations. The company characterises itself by having a technical and commercial approach, combining technological development with training services for its clients to help them use their products. It also has vast experience in deployment and maintenance of industrial systems on plant.

In TALON, EXOS has contributed bringing in his industrial digitalisation knowledge in some key tasks. EXOS has participated in data and process visualization, integrating capabilities at the edge and acquiring and processing the data needed from manufacturing resources to be used by TALON's AI algorithms. This was done in the industry 5.0 demonstrator, led by FACTOR. Leveraging its experience in industrial solutions implantation EXOS has contributed to TALON assuring adaptation to real needs and industrial constraints.

3.5.15.2 Exploitation Actions and Expected Return

There are two main technological results from TALON that are of special interest for EXOS. The first one is the Edge-to-cloud orchestrator (E2C), which is a core element of TALON architecture. This orchestrator allows to coordinate heterogeneous computational infrastructures between edge and cloud, optimizing data use and energy efficiency. This capability could be very useful if EXOS decides to turn MESView into a hybrid architecture, combining local processing in plant with cloud services. The second relevant result is TALON's explainable AI techniques. As MESView wants to evolve incorporating AI technologies for data analysis that allow decision support, incorporating TALON's explainable AI capabilities could increase trust and transparency from customers, improving recommendations interpretability.

EXOS plans a series of actions focused on incorporating these innovations from TALON into its commercial offer and internal capabilities. First one is analyse the technical viability of integrating the E2C orchestrator within MESView's architecture. After this, controlled pilots should be deployed (internal demonstrators) to validate the new functionalities of the orchestrator. On the other hand, EXOS plans on implementing modules or extensions in MESView, incorporating explainable AI features from TALON, assuring its interoperability with MESView's current core functionalities (manufacturing monitoring, quality, productivity, etc.). If these actions are successful, EXOS plans to train its technical-commercial team in TALON's key technologies, doing internal workshops and documentation.

EXOS market targets will remain being the current industrial sectors where it operates. These markets are food, plastics manufacturing, winemaking and retail equipment industry. EXOS already has customers and successful cases deploying MESView in these industries. These sectors are known for having distributed processes and real time data analytics/monitoring needs, therefore they can benefit directly from using edge-cloud orchestration and explainable AI. Additionally, EXOS will evaluate expanding to other manufacturing sectors that also need to optimize operational efficiency by leveraging Edge AI and MES solutions.

To bring these advancements into the market, EXOS plans several scenarios. On one hand, the improvement of MESView would contribute to retaining current customers. New functionalities derived from TALON would be offered as updates or additional modules, increased its value. In parallel, the possibility of collaborating with other members of the TALON consortium would be evaluated. These could be done to integrate TALON components into joint solutions or co-develop packaged offers.

EXOS's successful adoption of TALON results would entail various types of returns, aligned with its business objectives.

EXOS expects to have a positive impact on income and competitivity. The new functionalities explained earlier will differentiate MESView from other MES solutions, attracting new customers and market segments. EXOS will be able to offer a more complete product, justifying potential premium subscription or license models.

At the strategic level, incorporating technology like TALON could allow EXOS to reposition itself technologically. The company could evolve from being a traditional MES provider to becoming specialized in hybrid industrial and explainable AI, aligning with industry 5.0 principles. This would empower EXOS innovation and allow it to be one step ahead in a market that increasingly demands reliable and autonomous AI solutions.

Participating in TALON and absorbing its results internally could generate a big impact knowledge wise in EXOS. The edge-cloud integration, explainable AI and data governance know-how would strengthen its R&D capabilities, allowing it to tackle complex digitalization projects.

Lastly, incorporating TALON's solutions would allow EXOS to have an advantage in terms of regulatory compliance and alignment with emerging regulatory frameworks, as TALON has been conceived considering reliable AI, anticipating requirements from the future AI European regulation (AI Act).

3.5.15.3 Strategic Future Plans and Commitment

Looking into the future, EXOS defines a strategic vision aligned with Industry 5.0 evolution and AI adoption in industrial environments. In medium and long term, the organization aims to boost industrial digitalisation with hybrid architectures, distributing artificial intelligence between edge and cloud in an optimized manner. Components developed in TALON fit in this vision. This means that core functionalities of TALON will be evaluated to be possibly included in future MESView versions. In the same way, explainable visualization capabilities could be included in MESView's reporting modules or dashboards, improving transparency and trust.

EXOS is also open to explore technological transfer agreements, where it can license or exchange know-how with other members of the consortium. Incorporating TALON components developed by other members or offering MESView modules for future joint projects would be studied.

The adoption of TALON's architecture and principles will allow EXOS to anticipate industry demands in sustainability, resilience and technological humanism, in line with the European vision of Industry 5.0.

3.5.16 PROBOTEK

3.5.16.1 Partner profile

PROBOTEK is a Greek company, headquartered in Corfu, with offices in Athens and also with UK presence via PROBOTEK AI. It specialises in drone / UAV systems, Internet of Drones (IoD), AI/ML, robotics, edge/cloud/IoT integration and mission critical platforms.

In terms of innovation and digital transformation, PROBOTEK positions itself at the intersection of autonomy, real-time AI/ML inference at the edge, drone orchestration, mission control software, and integration of advanced communication / connectivity / telemetry systems. It already works in B2B and B2G domains, executing use cases such as BVLOS medical delivery between islands, early fire detection ("Pyro Sense"), anti-drone / security drones, etc.

In the TALON project, PROBOTEK contributes its expertise in drone orchestration, AI/ML at the edge, telemetry/mission control, connectivity, IoT and integrating robotic hardware with software platforms.

It plays a role likely in:

- Defining use-cases involving drone swarms, or UATVs (Uncrewed Aerial/Unmanned Autonomous Transport Vehicles) coordination, especially in security, crisis management, environmental monitoring.
- Contributing to design and deployment of the edge/cloud infrastructure needed for orchestrating drone fleets, including selection of AI models (edge capable), data pipelines, connectivity, real-time telemetry.

- Participating in standardisation / trust / safety frameworks for drone operations, explainability of AI models in drone behaviour, visualization / mission control interfaces.

3.5.16.2 Exploitation Actions and Expected Return

Below are exploitation plan items (mandatory) for PROBOTEK in the context of TALON:

TALON results / assets relevant to PROBOTEK	Concrete exploitation actions	Target markets / user groups	Exploitation scenario(s)	Expected returns
• The AI-powered edge/cloud orchestrator developed in TALON: model selection, dynamic orchestration, energy/data efficiency optimisations. • Use-case implementations involving drone / UATV coordination, AR/VR tools for maintenance & training, human-robot collaboration pilots. • Visualization / mission control dashboards for real-time monitoring, telemetry, explainability tools. • Trust / safety / standardisation / explainability / ethical AI frameworks applied to drones and industrial automation.	• Integrate TALON's orchestrator into PROBOTEK's proprietary platform (e.g. AiRFLOW, PERCEPTRON) to enable more efficient drone fleet operations. • Build new "Drone as a Service" or mission-control services offering, possibly with plug-and-play edge modules. • Develop training programs for clients (enterprises, governments) in drone operations under these new orchestration tools, AR/VR aided maintenance, safety / trust frameworks. • Possibly patent novel orchestration components, safety mechanisms or AI explainability modules; engage with standard bodies for drone safety, trustworthy AI in autonomous systems.	• B2B clients: telecommunications, energy (e.g., monitoring infrastructure in remote areas), security / surveillance, environmental monitoring, forestry / fire detection. • B2G: government agencies for crisis management, emergency response, border/security ,environmental regulation. • Remote / island logistics: delivery services, medical supply delivery, last-mile logistics where drones provide advantage. • Industrial plants: maintenance, human-robot collaboration, predictive maintenance.	• Offer services via new product lines: subscription SaaS / platform for drone orchestration. • Partnerships with industrial integrators, as well as municipalities, governments. • Licensing of components (e.g. AI orchestration, explainability, dashboards) to other drone / robotics companies. • Joint pilots / demonstration projects to convince clients. • Possibly a spin-off or internal business unit focused on the TALON-derived platform.	**Economic:** – New revenue streams from orchestration platform, services, licensing. – Competitive cost advantages by lowering energy / data usage and improving reliability. **Strategic:** – Stronger positioning in the drone / IoD market in Europe, being seen as innovator in autonomous / trustworthy / explainable AI in drones. – Expand service portfolio into edge-AI orchestration beyond drones (could apply to other robotics). **Knowledge-based:** – Improved technical capability in edge/cloud orchestration, AI model deployment, trust / explainability frameworks. – Better understanding of regulatory / standardisation landscape for UAVs / autonomous systems and ethics. **Policy / standardisation:** – Ability to influence or contribute to European AI regulation, drone

	• Enhance internal R&D / product roadmap: optimize power / energy consumption in edge devices, refine telemetry & model deployment pipelines.			safety / airspace regulation bodies (civil aviation, etc.), trustworthy AI clusters / consortia. – Recognition as a reference or case study in green / energy-efficient autonomy.

3.5.16.3 Strategic Future Plans and Commitment

Medium- to long-term vision:
PROBOTEK aims to sustain and scale up the TALON outcomes by embedding them in its core product roadmap. The orchestration framework and edge/cloud AI model provisioning will become integral parts of its AiRFLOW / PERCEPTRON platforms. Over time, PROBOTEK wants to establish a robust Drone-Orchestration-AI-Platform (DAIP) that clients can deploy in a variety of domains (delivery, monitoring, maintenance, security) with certified safety, explainability, trust.

Integration and maintenance of TALON results:
After the project ends, the TALON-derived modules (orchestrator, dashboards, explainability, energy/data efficiency optimization) will be maintained, versioned, and improved as part of PROBOTEK's product releases. Also, internal R&D will ensure compatibility with evolving edge hardware, sensors, communication protocols, regulatory compliance.

Collaboration beyond project:
PROBOTEK will maintain partnerships with other TALON partners for joint commercial deployments, co-marketing or joint bidding for EU / national projects. Also, will participate in standardisation consortia / trustworthy AI / UAV safety regulatory bodies, so that its innovations get recognised and possibly become part of emerging norms / certification schemes.

Resources & institutional commitment:
PROBOTEK will allocate R&D staff (AI researchers, software engineers, drone/hardware engineers) to maintain, document, support the new modules. Budget will be set aside for productization of TALON outputs. Also, marketing and business development resources will be allocated to pilot the exploitation actions. Possible hiring or external collaborations to fill gaps (e.g. in explainability, safety, regulatory compliance).

Alignment with strategic priorities:
TALON aligns closely with PROBOTEK's mission of enabling organisations to leverage drones / UAVs and integrate them into everyday operations with trust, efficiency and autonomy. Also matches its focus on digital transformation, green / energy efficiency (Industry 5.0), AI governance, and the European regulatory environment. By participating in TALON, PROBOTEK enhances its credentials and readiness for the next wave of demand for autonomous, trustworthy, explainable AI systems in industrial & governmental domains.

4 Clustering Activities

4.1 Overview

Clustering opportunities with other projects have been one of the main activities that WP6 efforts have been focusing on since the outset of the project. The present section of the document is dedicated to presenting the initiatives taken, together with their activities and outcomes.

4.2 TrustWorthy AI Cluster

4.2.1 Overview

As early as the first semester of the projects lifetime, TALON has been very engaged and proactive in establishing its main clustering initiative, by approaching and communicating with suitable candidate projects.

As a result of the above efforts, TALON achieved to initiate and become of the founding members of what was agreed to be called as the TrustworthyAI Cluster. Based on its level of activity and the number of projects involved this cluster proved to have delivered a well-coordinated impact for all the involved project partners.

This clustering initiative was built around a group of nine projects funded under the Horizon Europe call HORIZON-CL4-2021-HUMAN-01-01. Collectively, they aim to pave the way for greater acceptance and adoption of AI across Europe. While diverse in scope and approach, the projects are united by a shared commitment to designing AI solutions that are innovative, human-centric, and dependable. Their core focus is on advancing AI systems in ways that ensure verifiable robustness, energy efficiency, and transparency, thereby addressing some of the most pressing challenges associated with trustworthy AI [7].

The following nine projects (including TALON), form part of the call and each project provides solid scientific solutions complemented by tools and processes for design, testing, validation, certification, software engineering methodologies and real-world applications:

- **AutoFair:** Human-Compatible Artificial Intelligence with Guarantees
- **ENEXA:** Efficient Explainable Learning on Knowledge Graphs
- **EVENFLOW:** Robust Learning and Reasoning for Complex Event Forecasting
- **REXASI-PRO:** REliable and eXplAinable Swarm Intelligence for People with Reduced mObility
- **SAFEXPLAIN:** Safe and Explainable Critical Embedded Systems Based on AI
- **SustainML:** Application Aware, Life-Cycle Oriented Model-Hwardware Co-Design Framework for Sustainable, Energy Efficient ML Systems
- **TALON:** Autonomous and self-organized artificial intelligent orchestrator for a greener industry 4.0
- **TUPLES:** Trustworthy Planning and Scheduling with Learning and Explanations
- **ULTIMATE:** mUlti-Level Trustworthiness to IMprove the Adoption of hybrid arTificial intelligencE

Each project contributes its own distinctive perspective and expertise to the broad field of trustworthy AI. Together, they form a cohesive body of work aimed at developing solutions that foster trust and confidence in AI technology. The common objectives include improving transparency, explainability, accountability, safety, and performance in AI-driven systems.

The collaborative spirit underpinning this cluster is evident in the joint activities already initiated. The projects are actively engaging with one another to deepen their understanding of respective objectives, identify opportunities for synergies, and pursue avenues for joint action. These clustering activities are designed to ensure that outcomes align with both the specific call objectives and the broader priorities of the Digital Europe Strategy. The nine projects provide a wide range of complementary expertise, offering solid scientific contributions that are reinforced through the development of tools, processes, and methodologies.

By pooling their collective expertise, visions, and communication channels, the projects aim to amplify their individual and collective impact.

The European AI strategy seeks to position the EU as a world-class hub for AI. This is why collaboration and synergies are a driving force of the European research landscape and a key player on the strategy. The core goal is to address the challenges of AI trust and uptake, ensuring that AI solutions are not only cutting-edge but that they also resonate with European values.

4.2.2 Cluster activities and TALON participation

The TrustWorthyAI cluster developed a rich and diverse set of activities since its establishment, reflecting the collaborative spirit and commitment of its participating projects. These activities have included social media and digital campaigns, webinars, online workshops, and active participation in physical events, conferences, and scientific publications.

As a founding member of the cluster, TALON has been consistently engaged in all activities, contributing actively to their organization and execution which not only strengthened TALON's own exposure but have also supported the broader visibility and impact of the cluster as a whole.

In terms of the organisational aspects, the TrustWorthyAI cluster (Figure 9) has notably held regular, monthly telcos where news etc. were being shared between projects, while near future activities of mutual interest, were planned and discussed among participating projects. Up to the beginning of September 2025, a total of 21 telcos were had.

To establish its visual identity the cluster created its own visuals, communicating its presence and conveying its mission, Working Towards a trustworthy AI. All the visuals created were suitable adapted for different usage options, including for online posts, banners as well as roll-up banners to be used for physical events etc.

Figure 9: TrustWorthy AI Cluster main visual

Moreover, to reinforce its presence and ensure visibility across social media platforms, the cluster has also introduced the use of the dedicated hashtag #TrustWorthyAIcluster. All participating projects are encouraged to adopt this hashtag in their respective social media communications, thereby creating a unified and easily recognisable digital footprint for the cluster's activities.

Following is a timeline of main and notable TrustWorthyAI cluster activities realised, TALON's together with the relevant material, as applicable.

- **2 March, 2023 Clustering workshop (online).** This online event, the Cluster's first one, was also used to introduce the Cluster, together with its individual projects (Figure 10). TALON presented its objectives and other project particulars.

Figure 10: Online banner for the first clustering workshop

- **14 December 2023, 25th Infocom World, Athens, Greece (physical conference):** TALON partner 8BELLS introduced the Cluster to participants.
- **29 May 2024, Trustworthy AI:** Landscaping verifiable robustness and transparency (webinar). This event was co-organised by ADRA-e and the TrustWorthy AI Cluster of EU projects, in the context of the Birds of a Feather webinar series, which aims to provide insight into Research and Innovation projects aiming to foster collaboration and avail a joint platform for debate and results-sharing in the AI, Data and Robotics domains (Figure 11).

Figure 11: Online banner for the Trustworthy AI webinar

- **19 June, 2024, ADRA Exhibition, Convergence summit, (online):** Cluster projects participated in a joint virtual booth
- **18-19 February 2025, Future-Ready:** On-Demand Solutions with AI, Data, and Robotics, Brussels, Belgium. Organized by AI on Demand and ADRA. Supported by the AI4Europe and the ADRA-e projects. TALON partner INC, was among the Cluster project participants, and also presented TALON project itself.
- **27-28/ May, 2025, DataWeek, Athens, Greece (physical):** A Big Data Value Association (BDVA) where TALON partner ENG, represented the cluster, as well as the project in select sessions.

In addition to formal cluster-lead activities and initiatives, the participating project in the cluster have been assisting and supporting each other, largely at the communications level and mainly through posts in social media and platforms. This has been especially the case, in a coordinated social media campaign, during 2024, when each project, on a rotating basis, was introduced by the other sibling projects through agreed social media content. Other than the previously mentioned, sibling projects have been supporting each other on an ad hoc basis, through their own communication channels especially e.g. in the case of events promotion etc.

4.3 I4.0 / I5.0 Clustering

4.3.1 Overview

In addition to the TrustWorthyAI cluster that has been the primary initiative, TALON has been active in creating additional clusters with projects sharing common scopes and interests.

To that end and given that areas of trustworthiness in AI have been largely focused by the TrustWorthyAI cluster, TALON pursued clustering connections with projects in the areas of smart manufacturing, sustainability etc. in I4.0 - I5.0. As a result of this, connections were established with the following two projects:

- **AIDEAS -** AI Driven industrial Equipment product life cycle boosting Agility, Sustainability and resilience
- **i4Q –** Industrial Data Services for Quality Control in Smart Manufacturing

In this way, and through this arguably smaller-size cluster opportunity, TALON joined forces with two industry-focused added its impact by amplifying the benefits and outcomes of European research and innovation in the fields of industry and smart manufacturing.

What is also important to note is that all three participating projects in this grouping, shared a significant number of common partners in their respective consortiums, a fact that contributed in spreading the cluster message.

4.3.2 Activities and TALON participation

The activities realized as part of this smaller clustering initiative, even though smaller in scale, due to mainly the relatively smaller size cluster size, have been well focused and targeted in the mutual promotion of the individual projects and their objectives within the modern industrial operations space.

The main channel for this has been online communications, through projects' social media.
Moreover, the collaboration with AIDEAS specifically was formalised through a partnership arrangement they are running, called Supportive Partners Programme aiming to establish a collaborative network among stakeholders in EU industrial associations, clusters, digital innovation hubs, and other projects. The programme is addressed to partners interested in enhancing the visibility, reach, and success of AIDEAS, a project focused on advancing sustainable manufacturing in Europe. Supportive partners contribute to cross-dissemination activities, increasing the project's impact and outreach while ensuring sustainability [8].

In addition to the purely communication side activities, AIDEAS and TALON realised a webinar on 7 March, 2024, themed: Empowering AI for greener and more secure Industrial Operations (Figure 12). The webinar featured two keynotes by the two projects, with each of them addressing their relevant areas.

The TALON keynote was titled Enabling AI for E2E personalized and perpetual security and privacy, while the AIDEAS keynote was titled: Harnessing the Power of Sustainable AI in industry.

A panel discussion themed: The Synergy of AI, Blockchain, and Edge Networking for Sustainable and Secure Industrial Operations was also included in the webinar's agenda, together with a final Q&A session.

Figure 12: TALON - AIDEAS webinar banner

5 Market Analysis & Business Models

5.1 Market Perspectives of TALON's Key Drivers

TALON's envisioned approach consists of a sum of advanced technology components addressing crucial needs of I5.0, through a flexible, adaptable, and programmable AI architecture. The technologies utilized, can be described as a combination of current with some highly emerging and fast-growth ones.

In the same time, it is arguably true that, to the best of our knowledge, a sufficiently comparable solution is not available in the market, and certainly not one that is clearly focused and targeted to specific needs of I5.0. That reality creates a market landscape where no direct competition is present currently. This in turn implies that a market study will need to be forward-looking in its perspective and take into consideration, market data of the constituent parts that also form the pillars and most importantly the key drivers of the end-solution.

Indeed, TALON encompasses a number of technologies that can be arguably characterized as its technology enablers as well as acting as the drivers in its evolution and eventually in its market outlook and prospects. These key drivers are the following:

- **Edge computing,**
- **Cloud computing,**
- **Explainable Artificial Intelligence,**
- **Cloud orchestration**
- **Blockchain**

All the above play a very central role in technology developments both present and future, and that is also reflected in their individual market perspectives. This section presents the market perspectives of these key drivers, showcasing numbers and projections that illustrate their relevance in evaluating a market analysis for TALON as a whole.

Edge Computing involves processing data near its source (e.g. in devices or local servers), instead of sending them to remote data centers in the cloud. This paradigm has gained popularity due to the growing use of Internet of Things and the expansion of 5G networks, which require ultralow latencies and almost real time processing. Manufacturing, automotive and health industries are some of the main beneficiaries of edge computing. It allows local sensitive or high-volume data analysis, reducing bandwidth used and improving response times. Besides, distributing computational load towards the edge can decrease cloud congestion and contribute to the system overall energy efficiency, which is aligned with TALON's sustainability objectives. Gartner foresees that during 2025 75% of enterprise data will be generated and processed outside traditional data centers [9]. This shift reflects the transition of the edge paradigm from emerging to essential to obtain low latencies and efficient processing near data sources.

In market terms, edge computing is emerging as one of the fastest-growing sectors in technology. Projections indicate that edge computing global market will grow from 60 thousand millions of USD (54.6 thousand million EUR) in 2024 to approximately 110.6 thousand million in 2029 (100.6 thousand million EUR), with a compound annual growth rate (CAGR) of 13% [10]. This growth is driven by the massive adoption of IoT devices in companies of all sizes, deployment of AI/ML applications in the edge and need of privacy regulations and data governance compliance, which favor local processing [2]. It is worth noting that estimations project even greater numbers in the long run, which remarks the strategic importance of the edge. For example, one analysis estimates its market value around 327 thousand million USD (297.5 thousand million) for 2033, driven by an increase in connected devices and real time data generated [11]. In any case, tendencies align in the fact that edge computing will be a fundamental pillar in the digital infrastructure. TALON's commitment to edge-to-cloud deployments – smartly integrating resources at the edge and in the cloud – is a direct response to this evolution, aiming to leverage the low latency of the edge without sacrificing the elasticity of the cloud.

Over the last decade, Cloud computing has been the driving force behind businesses digital transformation. It allows companies on demand access to computational resources and storage at a global scale, paying only for resources used. This has drastically reduced innovation and service scalation barriers. Nowadays, almost every sector relies on cloud services to be agile and efficient. Factors such as automation, the integration of AI into cloud services and concerns over data security are reshaping the offerings of major providers, who are emphasizing the energy efficiency of their data centers to allow more sustainable computing.

Cloud computing global market has already reached an enormous size and continues to grow. In 2024 the cloud market size was estimated to be around 1.13 billion USD (1.13 thousand million EUR). Based on a recent report, cloud industry is expected to go from 1,3 billion USD (1.3 thousand million EUR) in 2025 to more than 2,28 billion (2 thousand million EUR) in 2030, which implies a solid expansion with an estimated CAGR of 12% [12]. This points out that, even though it is a mature technology, cloud still grows, pushed by digital transformation. Moreover, advanced capabilities integration (e.g. AI services, massive data analysis or multi-cloud orchestration) makes cloud solutions even more attractive for organizations. To sum up, cloud computing has become essential in today's technological landscape. TALON relies on cloud, while combined with edge computing and explainable AI, to offer a distributed platform taking advantage of the best of both worlds in an efficient and transparent way.

Explainable artificial intelligence (XAI) addresses a key necessity in AI adoption. This is being able to understand how and why an AI model makes certain decisions or predictions. In highly regulated industries such as healthcare, government or financial services, AI accuracy is not enough, it must also be transparent and justifiable. Emphasis in explainability has been reinforced by recent legal and ethical initiatives (e.g. EU AI regulation proposals) that aim to assure a responsible and ethical use of AI. In this context, XAI turns into a key component to generate trust in AI based systems.

Although XAI field is relatively young compared to AI in general, it is strongly growing. Highly regulated sectors like finance, healthcare or autonomous driving are investing in XAI solutions to guarantee that AI systems are more comprehensible, reliable and aligned with current regulations [13]. It is estimated that explainable AI tools and solutions global market grows from USD 6.2 thousand million (5.6 thousand million EUR) in 2023 to 16.2 thousand million (14.7 thousand million EUR) in 2028, which implies a staggering 20.9% CAGR. This highlights the need for transparency as machine learning models become more popular and expand their use in different business sectors. In fact, analysts point out that XAI will be crucial in order to ensure that artificial intelligence is used responsibly and ethically [14]. In TALON, explainable AI plays a relevant role. Its objective is to allow users understand decisions taken by its autonomous orchestrator. This goes beyond market expectations and regulations, as it offers trust and transparency to TALON users in their adopted AI solutions.

As companies transition towards more complex IT environments, with several public, private clouds and edge resources, the need to coordinate and automate this hybrid environment arises. Cloud orchestration refers to the tools and platforms that unify management of several workloads and cloud services, automating processes such as containers, resource and storage management, or monitoring apps in different environments. Instead of manually managing each service or infrastructure, cloud orchestration solutions allow policies definition and workflows deployment, which adjust resources dynamically based on demand.

Cloud orchestration market reflects companies' rapid transition towards hybrid architectures and microservices based applications. In 2024 cloud orchestration market size was estimated in 23.2 thousand million USD (21.1 thousand million EUR), and it is projected to reach 84.8 thousand million (77.1 thousand million EUR) in 2033, with a CAGR of 15.5% between 2025 and 2033 [15]. Other analysis foresees even a more accelerated growth, stating that it will reach 150 billion USD (150 thousand million EUR) in 2032, with a CAGR over 20% [16]. This notable growth is caused by several factors. On one hand, more companies use orchestration software to automate tasks in different domains and systems, reducing operation complexity of heterogenous IT environments. On the other hand, the growing demand for on-demand services and the need for agile infrastructure provisioning are driving investment in these tools. In summary, cloud orchestration has become a strategic component for every organization that manages several environments or looks for scalability while keeping efficiency. In TALON's case, this technology is at its core, as its objective is to offer a intelligent edge to cloud AI orchestrator. The project aims to provide autonomous and optimized distributed resources management, aligned with the trend of simplifying multi-cloud complexity through intelligent automation.

Blockchain, originally known for being the underlying technology of cryptocurrency, has evolved into a trusted platform for several businesses' applications. Its security and transparency features have created interest beyond the financial sector. Blockchain technology is maturing toward more pragmatic implementations focused on efficiency and sustainability. For example, private or consortium blockchain networks have emerged for enterprises and "greener" or less energy consuming algorithms are being adopted. This is particularly relevant for projects like TALON, which integrate blockchain to increase traceability and security in its AI platform, while keeping a low energy impact.

Based on a recent report, blockchain market is expected to skyrocket from 20.1 thousand millions of USD (18.3 thousand million EUR) in 2024 to approximately 249 thousand million in 2029 (226.6 thousand million EUR), which implies a CAGR of 65.5% [17]. Other sources foresight even greater growth, anticipating that blockchain technology could exceed 1.4 billion USD in 2030 (1.4 thousand million EUR), reflecting a potential exponential growth for the final half of the decade [18]. For TALON, blockchain use adds a trust and transparency layer, which complements other technologic pillars. By using a distributed ledger, TALON can ensure event recording (like AI models update, data exchange or compliance checks), so that these are auditable and tamper-proof. The fact that TALON adopts blockchain with a sustainable approach (e.g., using low-energy consensus algorithms and tracking carbon metrics) showcases the combination of innovation with emerging environmental responsibilities demands in technology.

Based on the research data above, the market outlook for TALON's key technologies suggests strong potential and strong, positive trends in support of its vision. By merging these specific elements into a unified platform, plus additional functionalities well-targeted for I5.0, TALON positions itself at the intersection of high-growth markets. Even though now there is not a directly competitive platform to TALON, the project vision is based in technological fundamentals with clear market relevance. By doing so, TALON is well positioned to leverage the expansion of these markets, offering an innovative solution aligned with market perspectives and the emerging priorities of the technology industry.

5.2 Market Analysis & Business Models for TALON KERs

This section is dedicated to the market analysis and business modelling discussion of the identified TALON Key Exploitable Results (KERs). These results constitute the foundational technological outputs of the project and are regarded as the most relevant drivers of its future, potential commercialisation prospects. The intention of the analysis is to assess the prospects of these results in terms of market entry, adoption, and long-term sustainability, while also outlining the potential challenges and enabling conditions that could shape their future success.

As discussed in the preceding section, the current state of the market is characterised by the absence of directly comparable offerings. In other words, there is no established solution currently available, that mirrors the specific technological and functional characteristics of TALON's outputs in a sufficiently similar way to allow for a straightforward competitor analysis, that serves the same needs. This lack of direct competition presents both an opportunity and a challenge. On the one hand, it positions TALON to operate as a first mover in its field, potentially setting standards and shaping market expectations. On the other hand, it makes forecasting and benchmarking more complex, as market adoption cannot be evaluated against pre-existing reference cases. For this reason, the analysis is anchored not in competitive comparisons, but rather in the broader market drivers that are likely to influence demand, uptake, and overall sustainability of the TALON offering.

The analytical focus has therefore been deliberately concentrated on the three identified KERs. These outputs have been selected because they are recognised as the most strategically significant technological enablers. Each of them plays a critical role in defining TALON's value proposition and is expected to drive the project's trajectory towards eventual market adoption. By narrowing the scope of analysis to these three elements, it is possible to generate a more detailed, structured, and actionable understanding of the pathways available for future exploitation.

For each KER, a dedicated assessment has been carried out with the aim of identifying the factors that are most likely to affect its market success. To ensure consistency and comparability, the analysis has been conducted using a structured SWOT (Strengths, Weaknesses, Opportunities, and Threats) framework. This approach allows for a systematic examination of internal assets and limitations, while also considering external conditions, such as regulatory trends, technological advancements, and evolving customer needs. The outcome of these analyses highlights both the potential advantages TALON can leverage and the vulnerabilities that will require strategic mitigation.

Complementing the SWOT assessments, a Business Model Canvas has been developed for each KER. This provides a comprehensive outline of the essential parameters underpinning potential business models, including value propositions, customer segments, revenue streams, cost structures, and key partnerships. By doing so, the analysis bridges the gap between abstract technological potential and practical commercial deployment. It offers a preliminary blueprint of how TALON results could be positioned, delivered, and monetised in real-world market contexts.

It is important to note that the insights presented in this section are not limited to desk research or theoretical considerations. They build directly on the extensive knowledge generated throughout the project's participation in the Horizon Results Booster (HRB) services. Specifically, they integrate the outcomes of workshops, targeted seminars, and discussions attended by the respective KER owners. This collaborative process has ensured that the analysis reflects not only external perspectives but also the practical understanding of the teams most directly engaged with the development and future exploitation of the KERs.

Taken together, the combination of SWOT analyses and Business Model Canvases provides a robust foundation for evaluating TALON's commercial prospects. It situates the three KERs within their relevant market environments, highlights both opportunities and challenges, and suggests viable pathways for sustainable business exploitation. Ultimately, this section provides a structured basis for moving forward from research and development achievements towards tangible market impact, in alignment with the broader objectives of Horizon Europe.

5.2.1 KER1 Zero-touch AI orchestrator (ENG / UL)

Table 26: BMC for Zero-touch AI orchestrator

<table>
<tr><th>Key Partners</th><th>Key Activities</th><th>Value Propositions</th><th>Customer Relationships</th><th>Customer Segments</th></tr>
<tr><td rowspan="3">• Edge/Cloud providers
• Research institutions
• Customer engagement and market access business units (ENG internal)
• HW vendors
• Contract research</td><td>• Development and maintenance
• Customer onboarding, customisation, and support
• Pilots, demos, PoC’s
• R&D
• Partners collaboration</td><td rowspan="3">• Fast, scalable and secure AI-driven orchestration
• Increased responsiveness and reliability in I5.0
• Sustainable, eco-friendly operation without performance discounts
• Reduced human error and operational costs
• Enhancing AI explainability and trust</td><td>• Communication and dedicated support
• Customisation and co-creation
• Long-term relationships with continuous improvement</td><td rowspan="3">• Industrial companies
• Cloud & Edge service providers
• Integrators active in I5.0 and AI-driven automation</td></tr>
<tr><th>Key Resources</th><th>Channels</th></tr>
<tr><td>• AI algorithms and E2C technology
• Engineers (AI and network)
• Infrastructure (data, training, AI models)
• E2C cloud computing resources
• Customer support and integration teams</td><td>• Direct sales to enterprise clients
• Leverage ENG’s existing partnerships
• Demos and pilot projects
• Webinars
• I5.0 events</td></tr>
<tr><th colspan="2">Cost Structure</th><th colspan="3">Revenue Streams</th></tr>
<tr><td colspan="2">• R&D
• AI development
• Skilled personnel
• Infrastructure (cloud, edge)
• Customer services (onboarding, support)</td><td colspan="3">• Fees charged for network management that is automated and reliable, high performance and reduced downtime</td></tr>
</table>

Key partners
Technology providers for edge and cloud infrastructure will play an essential role, as the orchestrator is heavily dependent on robust, scalable computing environments to operate effectively across distributed networks. Research institutions will be equally important, contributing to advancements in AI and supporting testing procedures. Within ENG, collaboration with internal business units may enhances customer engagement and market access. Network hardware vendors can be also valuable partners, particularly for integration. Finally, contract research organisations can provide flexibility for customisation, pilot implementation, and validation.

Key activities
The core activities of the Zero-touch AI Orchestrator will revolve around the continuous development and maintenance of the platform, ensuring that it remains reliable, secure, and scalable. A key operational focus will be on delivering smooth customer onboarding, providing tailored customisation, and offering dedicated support services. Moreover, to build trust and accelerate adoption, the orchestrator team may conduct pilot projects, live demonstrations, and proofs of concept (PoCs), enabling potential clients to experience its capabilities in real-world scenarios. Alongside the above activities, ongoing research and innovation will be integral, particularly in advancing AI explainability, developing greener and more energy-efficient AI models, and strengthening security features to remain compliant with evolving regulations. Finally, close collaboration with internal business units and external partners can ensure alignment with market needs.

Key resources
At the core of the orchestrator's resources will be the advanced AI algorithms and the edge-to-cloud orchestration technology that underpin its ability to automate, scale, and secure network management. Equally crucial will be the expertise of highly skilled AI and network engineers, who design, maintain, and continuously improve the system. Moreover, a robust data infrastructure may be required to train and operate the AI models effectively. Furthermore, access to both edge and cloud computing resources could enable the solution to deliver flexible, distributed orchestration across diverse deployment environments. Finally, dedicated customer support and integration teams can play a vital role in ensuring smooth adoption, customisation, and ongoing client satisfaction.

Value propositions
The Zero-touch AI Orchestrator provides fast, scalable, and secure AI-driven orchestration across the entire edge-to-cloud continuum. Moreover, for Industry 5.0 applications, it significantly improves responsiveness and reliability, enabling qiuick decision-making in demanding industrial environments. At the same time, it optimises energy consumption, offering sustainable and eco-friendly operations without compromising system performance. Also, it automates complex network management tasks, thereby reducing human error and lowering operational costs. Finally, the orchestrator enhances AI explainability and transparency, helping customers to meet their compliance and security needs.

Customer relationships
The orchestrator's customer relationships developed based on continuous communication and dedicated support with the stakeholders. Moreover, it can offer opportunities for customisation and co-creation, allowing the potential customers to adapt the solution to their network architectures and operational requirements. The focus is not on one-off transactions but on establishing long-term partnerships, where continuous improvement and updates can guarantee that the orchestrator evolves alongside customers' needs, technological advances, and regulatory requirements.

Channels
Direct sales efforts, supported by in-depth technical presentations, can allow enterprise clients and business units to clearly understand the orchestrator's added value and specific applications to their networks. At the same time, leveraging ENG's existing partnerships and industry networks could create entry points into established ecosystems. Practical demonstrations and pilot projects will act as proof-of-value, giving potential customers hands-on experience of the solution's capabilities in real-world scenarios. To build awareness and thought leadership, the orchestrator might be promoted through educational channels such as online webinars, white papers, and case studies. Finally, visibility at key industry conferences (particularly those focused on network management and Industry 5.0) may create synergies with different stakeholders.

Customer segments
The Zero-touch AI Orchestrator primarily targets medium to large industrial companies that operate complex networks and therefore require highly reliable, low-latency AI orchestration. In addition, it addresses the needs of cloud and edge service providers, who face increasing demand for intelligent and scalable network management as their infrastructures expand and diversify. A third core segment consists of digital solution integrators working on Industry 5.0 initiatives and AI-driven automation projects, as they seek advanced orchestration tools to enhance interoperability and performance across different industrial and technological environments.

Cost structure
A significant share of resources will be directed towards research and development, particularly in advancing AI capabilities, explainability, and orchestration technologies. This will be closely tied to the salaries of highly skilled AI specialists, network engineers, and support staff, whose expertise is essential for maintaining the platform's performance. Another substantial cost driver can be the cloud and edge computing infrastructure required to deliver scalability and security across diverse customer environments. Finally, investment in customer onboarding, customisation, and long-term support may represent an important cost area.

Revenue streams
The revenue model is primarily driven by fees charged for delivering automated network management that guarantees reliability, high performance, and reduced downtime.

Table 27: SWOT analysis for Zero-touch AI orchestrator

Internal factors	
Strengths	**Weaknesses**
• Significant reduction in errors and costs • Supports diverse AI deployment models (P2P, centralised, hybrid etc) • Enhances explainability and trust • Eco-friendly	• High initial costs • Complex integration into existing infrastructures • Need for continuous updates to maintain robustness and security
External factors	
Opportunities	**Threats**
• Growing demand • Increased industry focus on AI transparent and Green AI • Expansion opportunities in additional fields	• Competition from established network solutions (traditional/legacy, cloud-based) • Regulation and compliance challenges • Rapid technology changes • Customer reluctance

Strengths
The Zero-touch AI Orchestrator's advanced automation capabilities significantly reduce human errors and operational costs, ensuring more reliable and efficient processes. At the same time, it supports a wide range of AI deployment models, including peer-to-peer, centralised, and hybrid, which makes it adaptable to diverse business and technological contexts. Another key strength lies in its ability to enhance explainability and build trust in AI-driven decisions, an increasingly critical factor for user acceptance and regulatory compliance. Finally, by optimising energy efficiency, the orchestrator not only lowers the environmental footprint but also contributes to sustainable operations.

Weaknesses
One significant challenge of the orchestrator lies in the high initial investment required for research, development, and deployment, which may pose barriers for organisations with limited budgets. Additionally, the its complexity in terms of integration with diverse and often legacy infrastructures can create technical and operational difficulties. On top of that, the need for continuous updates and maintenance in order to ensure robustness, reliability, and security may add extra demands on resources.

Opportunities
The growing demand for real-time, scalable network management can create an opportunity for the orchestrator's deployment, as organisations increasingly seek automated solutions that can handle complex, dynamic infrastructures with minimal human intervention. At the same time, the rising industry emphasis on AI transparency and environmentally sustainable practices aligns with the orchestrator's strengths in explainability and energy efficiency. Moreover, it has strong potential for expansion across multiple high-impact sectors, including manufacturing, cloud services, finance, and healthcare, where AI-driven orchestration can deliver substantial improvements in performance, cost-effectiveness, and compliance.

Threats
Strong competition from both traditional networking solutions and increasingly sophisticated cloud-based management platforms may limit the orchestrator's differentiation and slow down its market potential. Additionally, regulatory and compliance challenges across different regions introduce complexity, as varying standards and legal frameworks may restrict deployment or increase the cost of compliance. Rapid technological advancements further pose a risk, as the orchestrator must continuously adapt to emerging innovations in AI, networking, and security. Finally, customer reluctance possibly driven by lingering concerns over trust, data privacy, and the explainability of AI-based decisions could delay its adoption.

5.2.2 KER2 Explainabale AI (XAI) framework (MINDS)

Table 28: BMC for XAI Framework

<table>
<tr><th>Key Partners</th><th>Key Activities</th><th>Value Propositions</th><th>Customer Relationships</th><th>Customer Segments</th></tr>
<tr>
<td rowspan="3">• Cloud infrastructure providers
• Industrial automation vendors
• Other integration partners</td>
<td>• Continuous R&D
• PoC’s
• Training programs
• Regulatory compliance documentation
• Updates</td>
<td rowspan="3">• Comprehensive XAI framework
• Visualisation of AI reasoning processes
• Reduced downtime
• Enhanced human-AI collaboration</td>
<td>• Technical support teams
• Co-creation partnerships
• Training programs</td>
<td rowspan="3">• Manufacturing companies
• Industrial facilities (owners/operators)
• System integrators</td>
</tr>
<tr><th>Key Resources</th><th>Channels</th></tr>
<tr>
<td>• Proprietary XAI algorithms & frameworks
• Expert team
• Cloud infrastructure</td>
<td>• B2B sales
• Validation processes
• Integration partnerships
• Workshops and webinars
• Industry events</td>
</tr>
<tr><th colspan="2">Cost Structure</th><th colspan="3">Revenue Streams</th></tr>
<tr>
<td colspan="2">• R&D costs
• Cloud infrastructure
• Computing costs
• Customer acquisition costs</td>
<td colspan="3">• Licensing fees (software)</td>
</tr>
</table>

Key partners
Metamind Innovations (MINDS) leverages strategic partnerships to deliver maximum value. Cloud infrastructure providers can ensure scalable, resilient delivery of the XAI-as-a-Service platform, and finally, industrial automation vendors act as distribution channels for MINDS' unique XAI components, integrating them into broader Industry 5.0 solutions.

Key activities
Key activities will focus on continuous research and development to keep the XAI capabilities at the forefront of industrial innovation, as well as on the adaptation of solutions for emerging industrial applications. Moreover, customer-specific proof-of-concept projects will demonstrate concrete value in operational settings, while regular training and certification programs will ensure that customer operators and technical teams are up-to-date. Finally, maintenance of up-to-date regulatory compliance documentation will also form a core part of ongoing activities.

Key Resources
The key resources are underpinned by proprietary XAI algorithms and frameworks developed through advanced TALON research, giving the company a technological edge in TrL1–TrL4 explainability. Essential resources also include a highly skilled team with expertise spanning explainable AI, industrial automation, and regulatory issues, as well as robust cloud infrastructure to deliver XAI-as-a-Service with real-time processing capabilities.

Value propositions
The core value delivered will be very much around an XAI framework that is comprehensive, in that it spans all trust levels (TrL1 to TrL4), thus ensuring an unbroken chain of transparency throughout every AI-driven process. The framework moreover provides clear visualisations of AI reasoning, making complex automated decisions accessible and trustworthy for human operators. This not only will allow operators to understand and validate maintenance recommendations—which in turn reduces production downtime—but will also foster intuitive, effective human-AI collaboration.

Customer relationships
The customer relationships will be based on a number of distinct activities. Fits, providing dedicated technical support throughout the integration and deployment stages of the XAI solution, is expected to form the basis. Another approach will include co-creation partnerships, where solutions will be tailored to each client's specific industrial use cases, fostering a sense of ownership and customization. Finally, regular training programs will further ensure that customer teams can efficiently utilise and interpret XAI outputs.

Channels
The XAI offering will utilise and exploit multiple B2B channels, including technical sales, live demonstrations, and proof-of-concept deployments directly at client facilities. Moreover, partnerships with industrial automation vendors and system integrators will help embed the XAI solution into existing and new Industry 5.0 platforms. Additionally, technical workshops, webinars, and participation in industry conferences will serve to educate, demonstrate real-world applicability, and broaden market reach among stakeholders interested in trustworthy AI.

Customer segments
The primary customer segments targeted will be manufacturing companies implementing AI-driven automation, industrial facilities pursuing advanced, human-robot collaboration, and system integrators involved in Industry 5.0 solutions.

Cost structure
The cost structure will reflect significant investments in continuous R&D for the ongoing development and industrial adaptation of XAI algorithms. Moreover, operational costs for cloud infrastructure and real-time processing will be substantial, as are customer acquisition costs related to proof-of-concepts and technical demonstrations necessary to gain market traction and validate real-world performance.

Revenue streams
Revenue is expected to be primarily generated from software licensing fees for the deployment of the XAI framework.

Table 29: SWOT analysis for XAI Framework

Internal factors	
Strengths	**Weaknesses**
• Comprehensive framework covering all trust levels • Strong scientific foundation • Established expertise • Direct experience with industrial partners	• High computational requirements • Technical complexity • Limited market awareness • Strong dependency on industrial data • Need for continuous updates
External factors	
Opportunities	**Threats**
• Regulations for explainable AI • Increasing demand for trustworthy AI • Expansion opportunities in new sectors	• Resistance from industries • Varying international standards and regulations • Risk of oversimplification

Strengths
XAI stands out thanks to its comprehensive framework, which covers all established trust levels for explainable AI in industrial environments. Moreover, it is anchored in a strong scientific foundation, drawing on deep research and proven methodologies. Established expertise within the development team, further reinforced by direct experience working with industrial partners, ensures that the solution is not only theoretically sound but also practically validated in real-world deployments.

Weaknesses
One of the key challenges is the high computational resources requirement for real-time XAI processing, especially in complex industrial scenarios. This may limit scalability and make integration demanding for clients with legacy infrastructure. The technical complexity of implementing and maintaining the framework could also create barriers to adoption, as it implies specialized knowledge for deployment and customisation. Additionally, there is currently limited overall market awareness about the capabilities and benefits of advanced explainable AI. The strong dependency on access to extensive and relevant industrial data makes long-term value delivery contingent upon ongoing data partnerships. Finally, the framework demands continuous updating to remain aligned with technological advances and evolving regulatory standards.

Opportunities
Emerging and forthcoming regulations mandating explainable and transparent AI systems can create new avenues for market adoption for XAI. Also, there is a rapidly increasing demand among industries for trustworthy and auditable AI decision making, both for regulatory compliance and for enhancing stakeholder confidence. Furthermore, growth opportunities exist across new sectors as awareness of explainable AI's value continues to spread.

Threats
Potential resistance from traditionally conservative industries could slow the pace of adoption, especially where operational disruption or perceived risk from new technology is high. The presence of diverse and sometimes inconsistent international standards and regulations complicates the deployment of a solution like XAI. Lastly, there is a risk that demands for explainability lead to oversimplified AI models, which could undermine performance or diminish the perceived value of the technology in tackling complex, real-world problems.

5.2.3 KER 3 Blockchain Mechanism (SID)

In the present section the BMC and SWOT analysis for SID's Blockchain Mechanism is shown.

Table 30: BMC for Blockchain mechanism

<table>
<tr><th>Key Partners</th><th>Key Activities</th><th>Value Propositions</th><th>Customer Relationships</th><th>Customer Segments</th></tr>
<tr>
<td rowspan="3">Hyperledger fabric blockchain network
Online infrastructure providers</td>
<td>Development and testing
Integration and support
Onboarding (Demos, PoCs etc)</td>
<td rowspan="3">Traceable and trustworthy AI decisions
Secure federated learning environments
Compliance with privacy and security standards</td>
<td>Technical support
Customisations
Co-developments
Feedback loops</td>
<td rowspan="3">Supply chain/logistics
Telecommunications providers
Retail/e-commerce platforms
R&D community focused on AI and blockchain integration</td>
</tr>
<tr><th>Key Resources</th><th>Channels</th></tr>
<tr>
<td>Expert development teams
Advanced consensus algorithms
Certificate authorith and access control
SaaS and integration infrastructure</td>
<td>Direct deployment via SaaS platforms
Industry-specific demos, PoCs, and consortium collaborations
EU R&D</td>
</tr>
<tr><th colspan="2">Cost Structure</th><th colspan="3">Revenue Streams</th></tr>
<tr>
<td colspan="2">Personnel
Infrastructure
Support and customisation services</td>
<td colspan="3">Access to secure and scalable FL environments
Custom deployments and integrations of blockchain-enabled FL settings</td>
</tr>
</table>

Key partners
Though SID serves as the exclusive module provider, ensuring technological distinctiveness, partnerships are an essential part of the business model. Technically, the solution leverages a customised deployment of Hyperledger Fabric blockchain, tailored for this advanced use case, and harnesses scalable online cloud infrastructure via Digital Ocean.

Key activities
The key activities will involve the design, development, and testing of integrated blockchain and federated learning systems. Moreover, supporting pilot use cases and collaborating in live industry trials will be central to validation. Finally, proactive stakeholder onboarding (through demos, proof-of-concept initiatives, and policy or standards panels) will ensure that the solution is market-ready and ahead of regulatory curves.

Key resources
Critical resources underpinning the offer include a highly skilled, cross-disciplinary development team with expertise in blockchain and federated learning. Moreover, technical assets, such as advanced consensus algorithms (like RAFT), robust certificate authority systems, and secure access controls, will ensure reliability and compliance. The existing infrastructure includes the needed resources for SaaS and integration with customer systems, thus offering flexibility, while supporting seamless integration with client environments and enabling scalable solution delivery.

Value proposition
SID provides a secure, transparent, and highly scalable federated learning (FL) environment fortified by blockchain technology which addresses major pain points such as data integrity, trust in decentralised AI decisions, and stringent privacy requirements. End users can benefit from the ability to trace AI decision logic, operate in secure and compliant FL environments, and fulfil legal and customer expectations for explainability and data protection.

Customer relationships
The customer relationships will emphasise in collaboration and support, including continuous technical assistance tailored to specific use cases. Co-development opportunities with early adopters will facilitate unique solution refinement, while structured feedback mechanisms will ensure the technology stays aligned with evolving customer and regulatory needs.

Channels
Broad level adoption is anticipated to be achieved via direct deployment as a SaaS solution, allowing scalable access and seamless integration into client operations. Demonstrations, proof-of-concept implementations, and active collaboration within industry-specific consortiums will further support and facilitate market traction, particularly among early adopters and innovation leaders. Finally, further stakeholder engagement will occur through integration with EU-funded R&D consortia, offering early, practical deployment opportunities within innovation-driven contexts.

Customer segments
The blockchain is tailored for industries where sensitive data management and trustworthy AI processing are paramount. Key stakeholders include supply chain and logistics operators, who require secure and efficient coordination, telecommunications providers responsible for massive, privacy-sensitive data flows and retail or e-commerce platforms that leverage AI for customer experience and operations. Additionally, the R&D community exploring the convergence of AI and blockchain presents a vital segment, as it seeks cutting-edge, trustworthy, and verifiable solutions in decentralized data processing.

Cost structure
The major costs will include personnel for ongoing R&D, as well as dedicated development of blockchain/FL infrastructure and SaaS hosting. Additional significant expenses will arise from providing continuous support and customisation for a diverse client base.

Revenue streams
The revenue streams will be centered on providing access to secure and scalable Federated Learning (FL) environments, and offering tailored blockchain-enabled FL deployments and integrations.

Table 31: SWOT analysis for Blockchain mechanism

Internal factors	
Strengths	**Weaknesses**
• Transparent and Verifiable • Private, scalable and reconfigurable architecture • Innovation in blockchain and FL integration	• High complexity • Significant customization needs
External factors	
Opportunities	**Threats**
• Increasing demand for secure and trustworthy AI • Privacy and decentralization supported by recent trends • Expansion opportunities in AI-heavy industries	• Market competition • Uncertainty in legal frameworks • Adoption challenges

Strengths
Blockchain excels in transparency and verifiability, offering users and stakeholders in general, explicit evidence of trustworthy data processing and AI decision-making. Its architecture stands out for being private, highly scalable, and easily reconfigurable, allowing adaptation to various use cases and industries without compromising security. Finally, the innovative integration of blockchain and federated learning (FL), marks the solution as a pioneer in enabling privacy-preserving, decentralised AI collaboration.

Weaknesses
The high complexity inherent to the system can pose obstacles in deployment, maintenance, and user adoption and may require substantial domain expertise, both technical and operational. Additionally, significant customisation is often needed to fit the architecture to specific client needs.

Opportunities
The market is witnessing a surge in demand for secure and trustworthy AI solutions, primarily fueled by increasing regulations, public awareness of data privacy, and the proliferation of AI-driven decision systems. Current trends emphasize privacy and decentralization, positioning the blockchain solution at the forefront of industry needs. There is also substantial potential for expansion within industries that are rapidly adopting AI (such as finance, healthcare, telecommunications, and logistics) where requirements for data integrity, compliance, and verifiability present clear business opportunities.

Threats
The primary threats involve market competition, as numerous players are targeting privacy-preserving and trustworthy AI. Moreover, uncertainties in legal and regulatory frameworks for both AI and blockchain technologies could hinder deployment or necessitate rapid pivots to stay compliant. Finally, adoption challenges stemming from user resistance to complexity, integration overhead, or uncertainty over ROI may slow penetration, especially among more conservative industry sectors.

5.3 Technology Roadmap - Analytical Hierarchy Process

The Analytic Hierarchy Process (AHP) is a widely used method for supporting complex decision-making. It provides a structured framework that breaks down unstructured, multifaceted problems into a multi-level hierarchy consisting of interrelated criteria and decision alternatives. In AHP, criteria and sub-criteria are organized in a hierarchical structure, with the overall goal positioned at the top. The relative importance of these criteria is assessed through pairwise comparisons conducted by experts.

The AHP method was developed by Thomas Saaty [19] in the early 1970s, with its initial applications in the military domain.

A number of steps are required to reach the final result (Figure 13):

- **Hierarchy Structuring:** The decision problem is decomposed into a multi-level hierarchical structure with the overall goal at the top, criteria (and possibly sub-criteria) in the middle, and decision alternatives at the bottom.
- **Pairwise Comparisons:** Each pair of criteria (and alternatives) is compared to assess their relative importance or preference toward the decision goal. Comparisons are typically made using Saaty's 1-to-9 scale, where 1 means equal importance and 9 indicates extreme importance of one over another.
- **Consistency Check:** Because human judgments can be inconsistent, AHP includes methods such as the Consistency Ratio (CR) to measure consistency in pairwise comparisons. If inconsistencies are too high, then appropriate actions may be considered e.g. discarding inconsistent responses etc.
- **Priority Setting and Synthesis:** Mathematical techniques, including eigenvector computations, are used to synthesize pairwise comparisons into priority weights for criteria and alternatives.

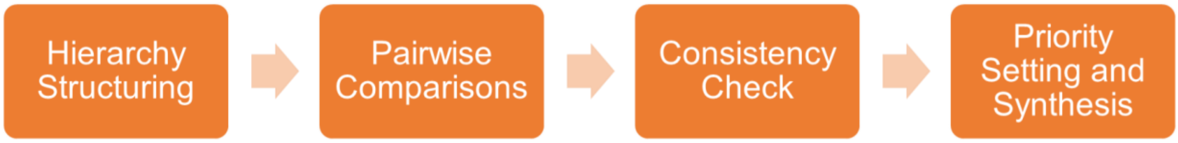

Figure 13: AHP steps

As a result, a three-level hierarchy is constructed (see D6.1 for more details). At the top level is the main objective to be achieved. The second level includes the criteria that influence the achievement of this objective, while the third level contains the sub-criteria, which further describe and detail the specific dimensions of each criterion.

In the next step, questionnaires are developed based on the above hierarchy, asking stakeholders to perform pairwise comparisons among the elements at the same level of the hierarchy. These comparisons help assess the relative importance of each criterion and sub-criterion. To facilitate this process, a nine-point scale, as presented in Table 32, is used to express the degree of preference between the options.

Table 32: The Saaty Rating Scale

Intensity of importance	Definition	Explanation
1	Equal importance	The two criteria have equally contribution
3	Moderate importance	Experience and judgment favour one of the criteria
5	Strong importance	The criterion is strongly favoured
7	Very strong importance	The criterion is dominant
9	Extreme importance	The criterion is favoured by at least an order of magnitude
2, 4, 6, 8	Intermediate values	Intermediate values that used as alternatives between any two of the above numbers

Once the stakeholders provide their input the calculation of weights for criteria and sub-criteria is made.

This section presents the results of the roadmapping activity conducted using the AHP method. The objective is to identify and prioritise the key factors that may influence the market adoption of the TALON solution.

5.3.1 Criteria and sub-criteria of TALON AHP

Initially, the most important criteria influencing the adoption of the TALON solution were identified through dedicated discussions and sessions with experts of partner organisations. The final list of criteria includes the following:

1. **AI-fueled Orchestration**
2. **AI Theoretical Framework & Benchmarking for Greener AI Deployments**
3. **ML-driven E2C Deployments & Runtime Adaptations**
4. **AI Explainability, Trustworthiness & Transparency**

In the next step for each of the determined criteria a number of sub-criteria was identified.

For the **AI-fueled Orchestrator** the following four sub-criteria were identified:

- **Intelligent Resource Management & Task Offloading:** Dynamically allocates computing tasks across cloud, edge, or local devices to balance load and improve efficiency (e.g., sending heavy ML training to cloud, while inference runs on edge)
- **Fault Tolerance & Self-Healing:** Detects failures in services or nodes and automatically reroutes or restarts them without user interruption (like Kubernetes pods restarting)
- **Dynamic Service-level Awareness & Interpretation:** Monitors performance targets (latency, throughput, energy use) and adjusts orchestration policies to maintain SLA compliance
- **Decentralised Control Capabilities & Policy-enforcement:** Enables multiple orchestrators to make local decisions while adhering to global governance rules (e.g., GDPR-compliant edge deployments)

For the **AI Theoretical Framework & Benchmarking for Greener AI Deployments** the following five sub-criteria has been identified:

- **Holistic AI Pipelines Lifecycle Management:** Covers full AI lifecycle from data collection, model training, deployment, to retirement, ensuring traceability and efficiency
- **Multidimensional Impact Assessment:** Evaluates trade-offs between performance and sustainability (e.g., high accuracy vs. increased carbon footprint)
- **Methodological Transparency, Causality & Attribution:** Provides clear reasoning for orchestration decisions (why a task was offloaded, why a model was chosen)
- **Multi-objective Evolutional Methods & Optimisation Strategies:** Uses evolutionary algorithms or reinforcement learning to optimize multiple objectives (cost, latency, fairness) simultaneously
- **Federated Learning / Decentralised Training:** Enables model training across distributed data sources without centralizing sensitive data (e.g., hospitals training on patient data locally).

For the **ML-driven E2C Deployments & Runtime Adaptations**, four sub-criteria have been identified as the most important:

- **Deployment Strategy Optimisations (cost, compute resources, elasticity, energy-efficiency):** Chooses the best deployment strategy to reduce costs and resource use while maintaining performance
- **Edge Deployments of Compute-intensive Tasks:** Brings AI computation closer to where data is generated, reducing latency and bandwidth costs
- **Context-aware Task Switching:** Adjusts orchestration depending on user or system context (e.g., switching from high-performance to energy-saving mode on mobile).
- **Intelligent Task Offloading:** Decides when to offload tasks to external resources for optimal performance, balancing device constraints and cloud capacity.

Finally, for the **AI Explainability, Trustworthiness & Transparency** criterion, five sub-criteria are considered the most important:

- **Algorithmic AI Model Transparency:** Provides visibility into how algorithms function, including design choices and assumptions
- **Interpretability of AI Model Internals:** Explains internal mechanics (weights, feature importance, decision layers) in human-understandable terms
- **User-centric Explanation Quality & Effectiveness:** Tailors explanations to the user's expertise, ensuring clarity for both developers and non-experts
- **Reliability & Robustness of AI Models:** Ensures models remain accurate under noisy, incomplete, or adversarial data
- **Fairness & Bias Mitigation:** Detects and reduces biases in training data or predictions, ensuring equitable outcomes across user groups

The full hierarchy scheme that has been created for the TALON is presented in the Figure 14 below.

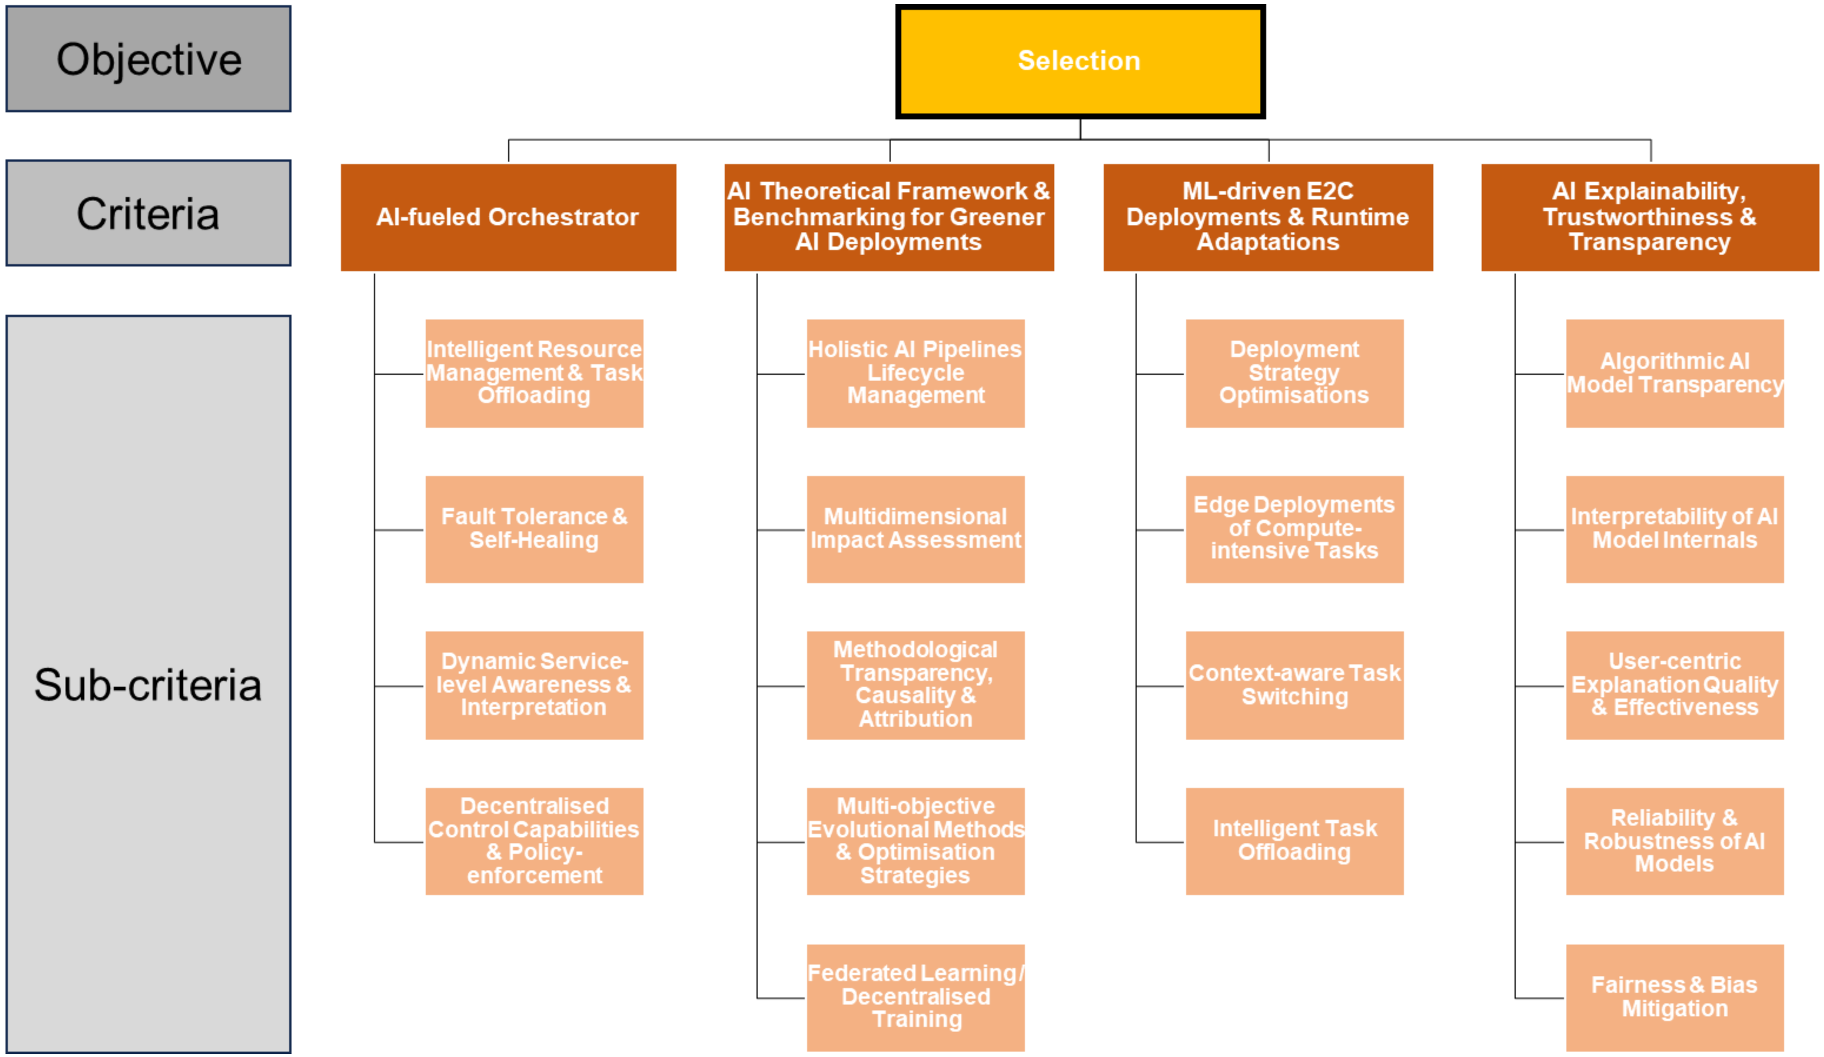

Figure 14: TALON AHP criteria scheme Standardization Activities

5.3.2 Description of the survey

A dedicated web-based application to feature the survey was designed and delivered by 8BELLS to collect the responses, and made accessible through: https://ahp-talon.8bellsresearch.com/. In the first page of the survey (Figure 15) the survey participants input their details.

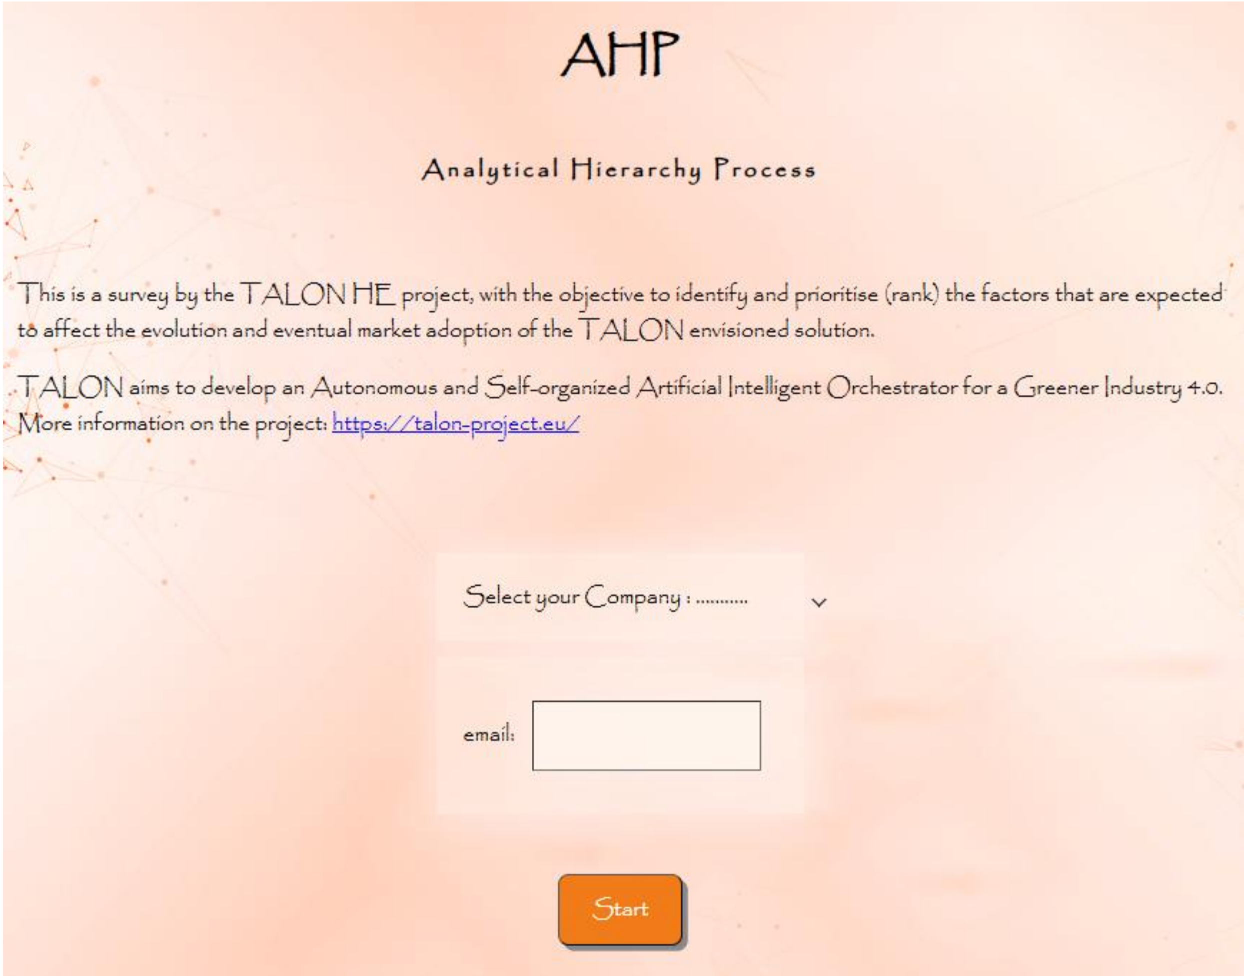

Figure 15: The first page of the TALON AHP

Following the submission of the above, the webpage gives them some instructions about the AHP survey (Figure 16).

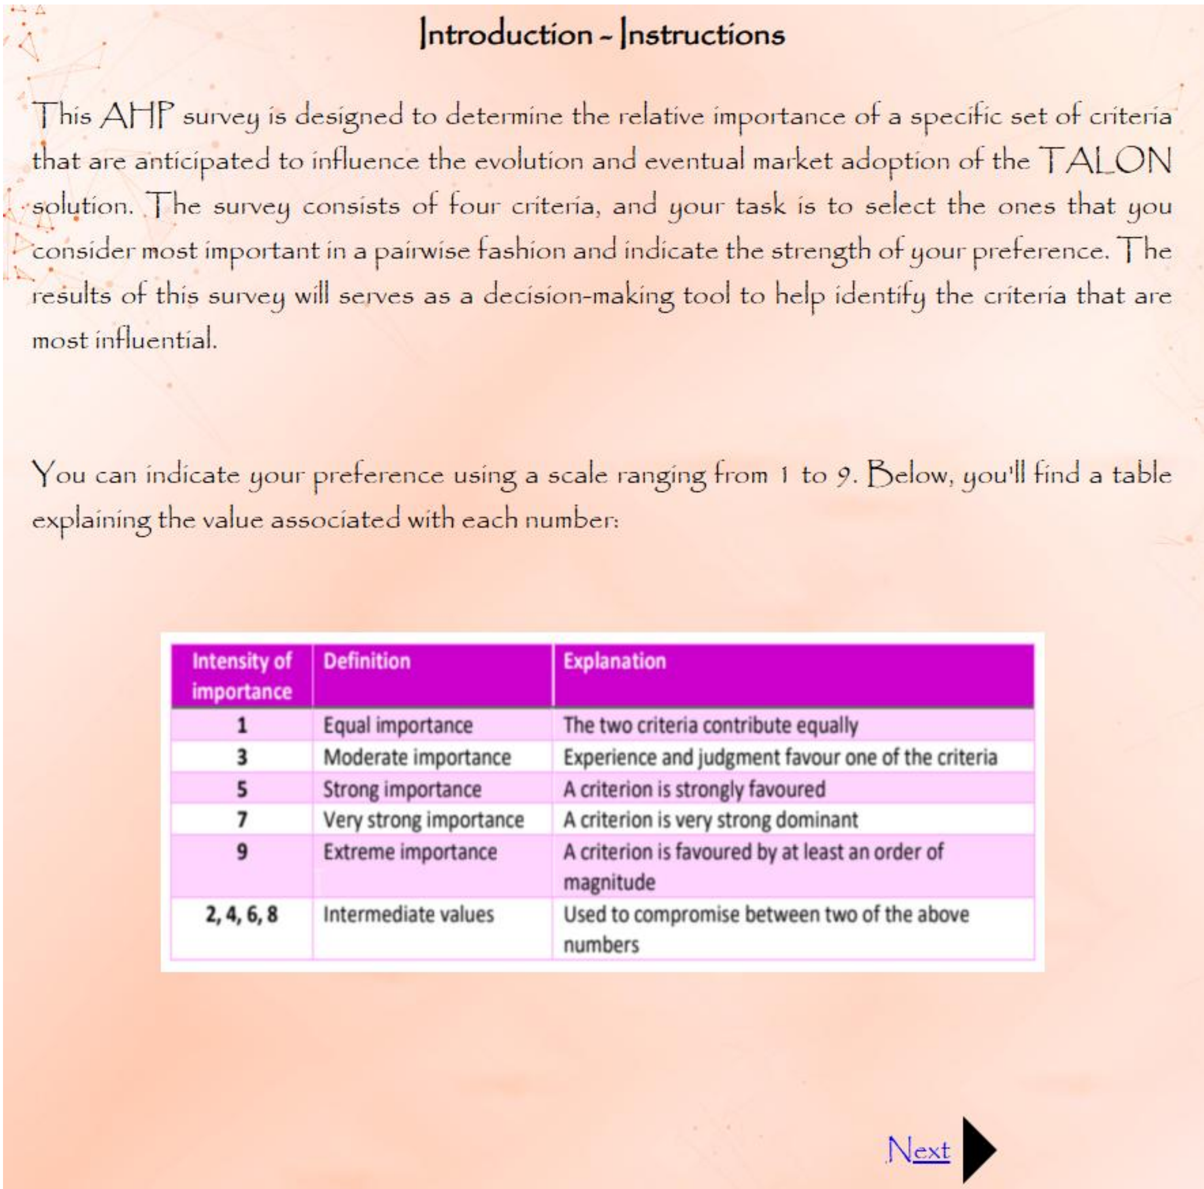

Introduction - Instructions

This AHP survey is designed to determine the relative importance of a specific set of criteria that are anticipated to influence the evolution and eventual market adoption of the TALON solution. The survey consists of four criteria, and your task is to select the ones that you consider most important in a pairwise fashion and indicate the strength of your preference. The results of this survey will serves as a decision-making tool to help identify the criteria that are most influential.

You can indicate your preference using a scale ranging from 1 to 9. Below, you'll find a table explaining the value associated with each number:

Intensity of importance	Definition	Explanation
1	Equal importance	The two criteria contribute equally
3	Moderate importance	Experience and judgment favour one of the criteria
5	Strong importance	A criterion is strongly favoured
7	Very strong importance	A criterion is very strong dominant
9	Extreme importance	A criterion is favoured by at least an order of magnitude
2, 4, 6, 8	Intermediate values	Used to compromise between two of the above numbers

Next

Figure 16: The instructions sheet in TALON AHP

Then, the next page shows the hierarchy of the criteria and sub-criteria (Figure 17).

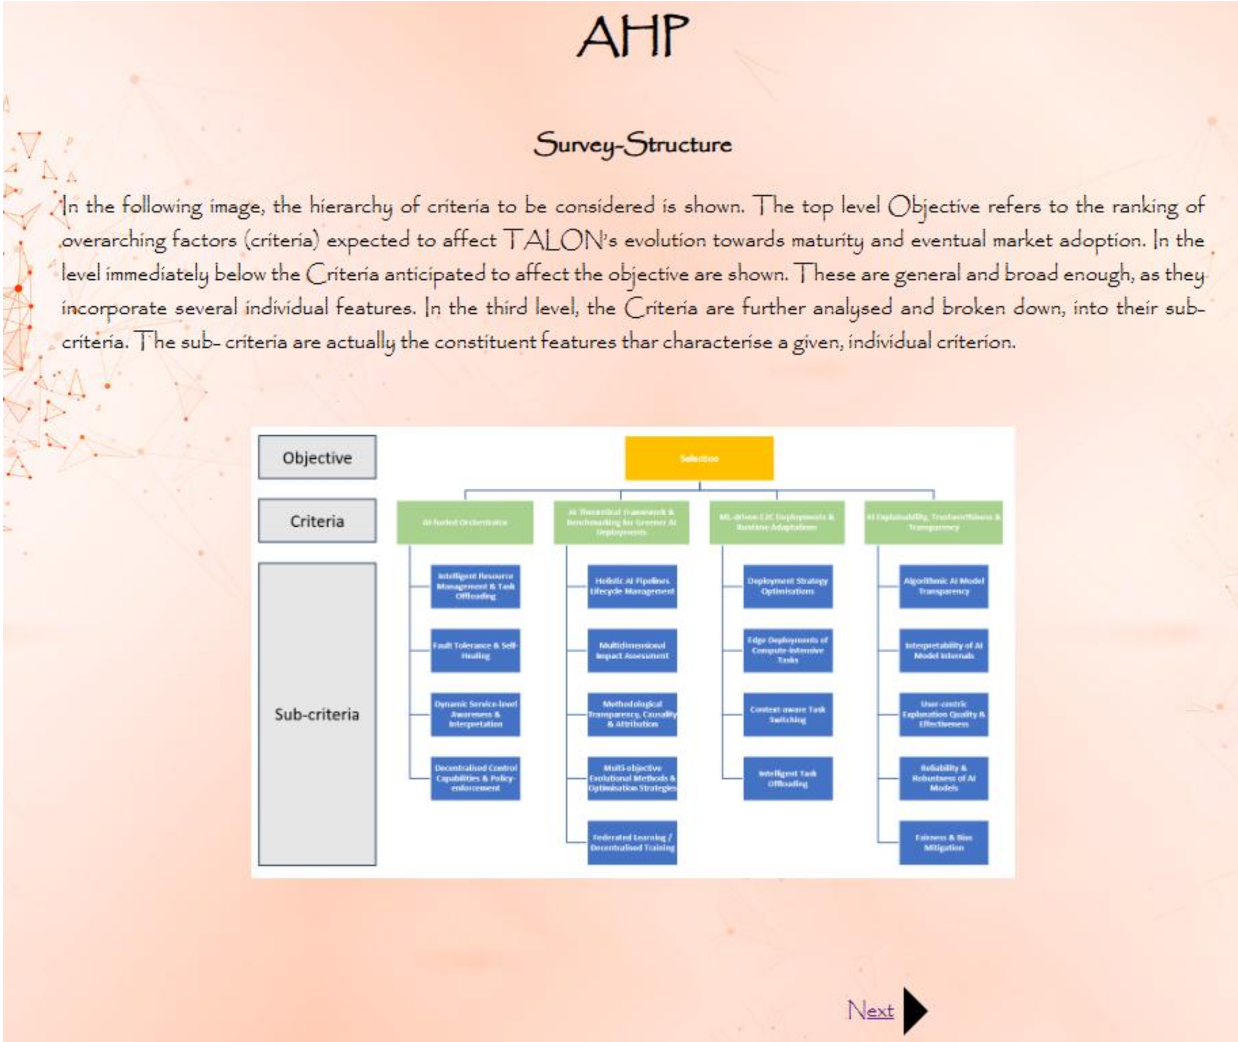

Figure 17: The criteria hierarchy in TALON AHP

Therefore, the respondent first performs the prioritization of the main criteria, followed by the prioritisation of the sub-criteria within each main criterion (Figure 18 and Figure 19).

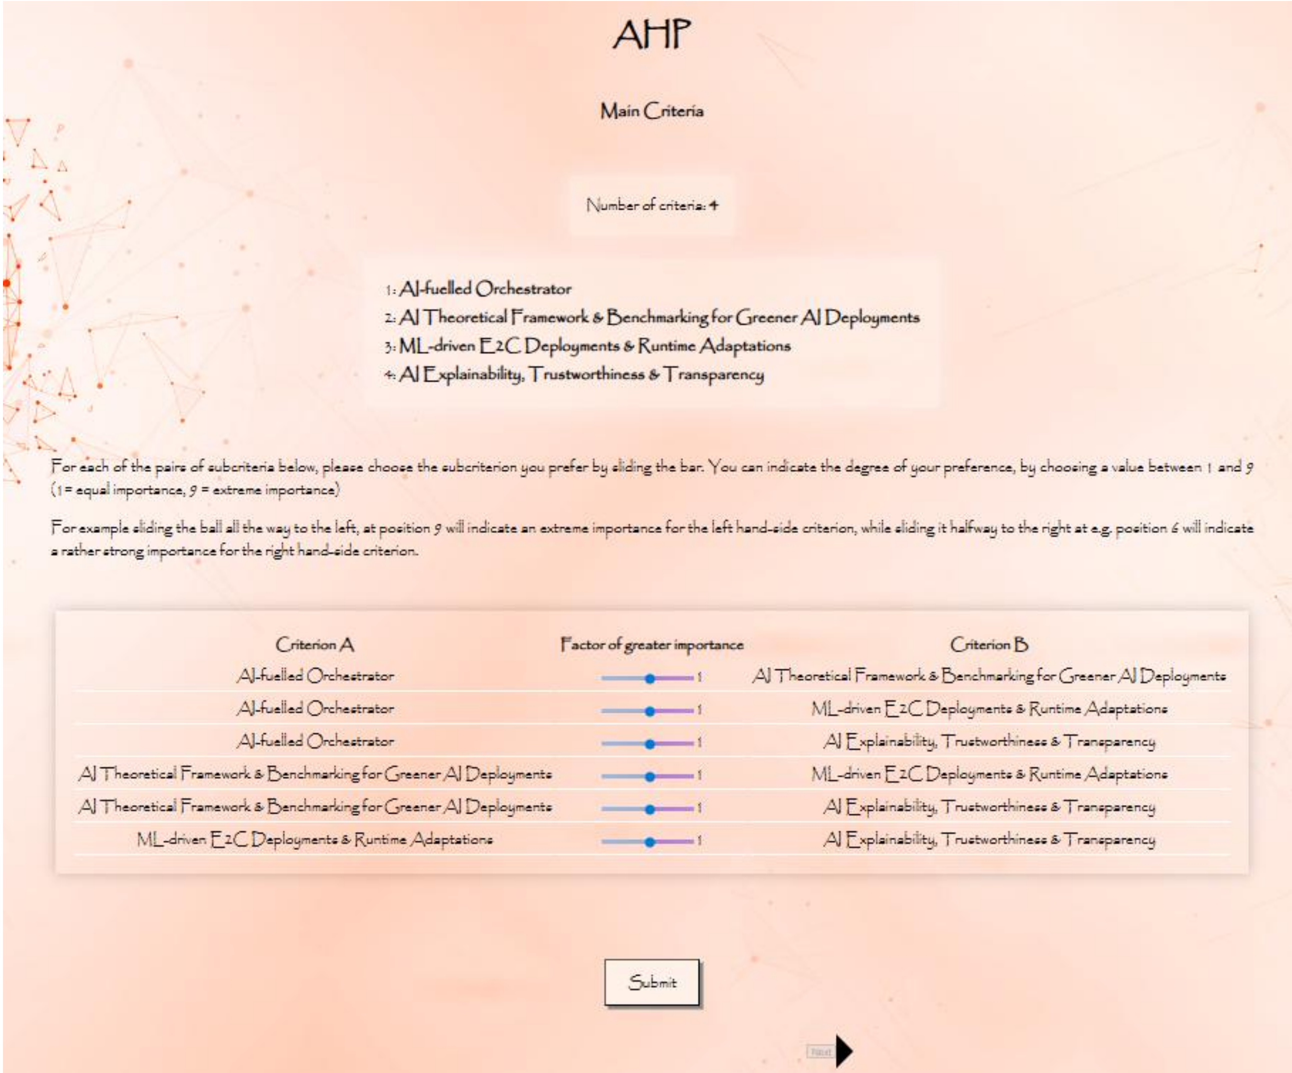

Figure 18: Prioritisation of the main criteria

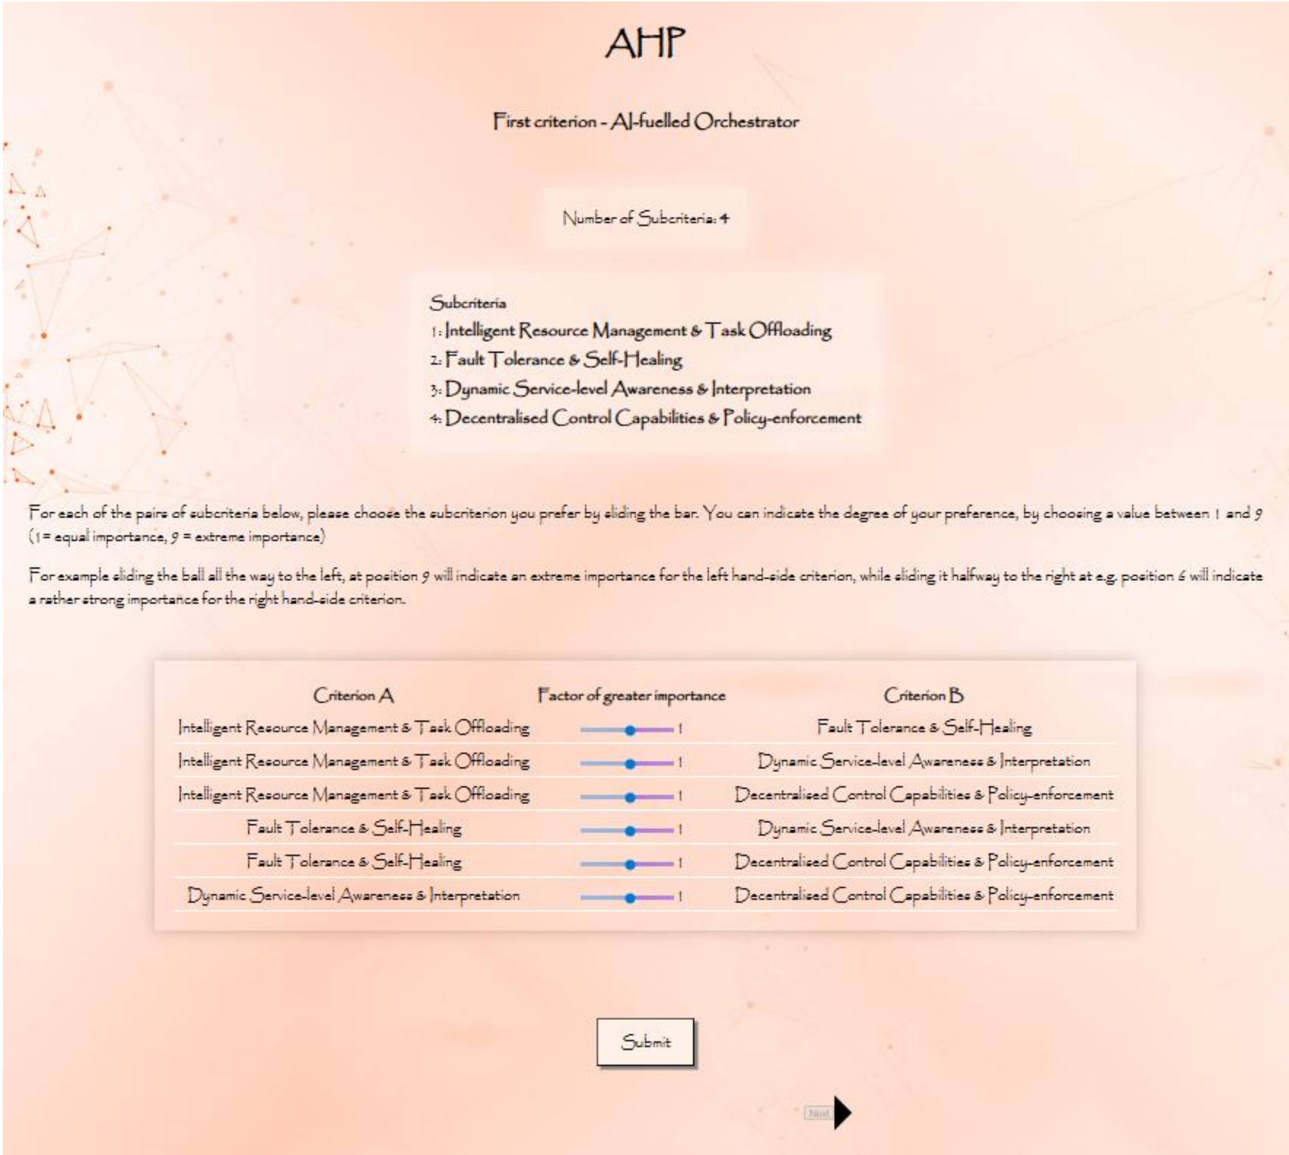

Figure 19: Prioritisation of the sub-criteria within the first main criterion

5.3.3 TALON AHP Survey specifics

Due to the highly specialised nature of the subject matter the associated set of questions (criteria comparison), it was deemed appropriate to draw survey participants from a suitably knowledgeable pool of individuals. To do so the survey's link: https://ahp-talon.8bellsresearch.com/ was distributed to TALON consortium, to be responded by individuals, experts in the fields of TALON.

In this way the survey was completed by a blend of different profiles, coming from industry, academia and research.

TALON's AHP survey took place over a two weeks period, in September 2025. The survey was conducted by a total of 26 responders.

Part of the AHP theory is a check of consistency, i.e. checking to ensure that respondents' pairwise comparisons are not random or contradictory. To do so Consistency Ratio (CR) is used, which in mathematical terms it is defined, based on the following process:

1. Form the pairwise comparison matrix A
2. Compute the maximum eigenvalue λ_{max} of A

3. Calculate the Consistency Index (CI): $CI = \frac{\lambda \max - n}{n-1}$ where n is the number of criteria

4. Calculate the Consistency Ratio (CR): $CR = \frac{CI}{RI}$ where *RI* is the Random Index, i.e. average CI of randomly generated matrices of the same size.

Saaty [20] had a provided a standard table of RI values, as follows:

- n = 1 → 0
- n = 2 → 0
- n = 3 → 0,58
- n = 4 → 0,90
- n = 5 → 1,12
- n = 6 → 1,24

A threshold that seems to be established is 0,1, meaning that a computed CR that is greater than 0,1 implies that the given pairwise comparison matrix of this respondent so as to not affect the overall accuracy of the results. However, it must be noted that this is not an absolute mathematical law, but rather a rule-of-thumb approach, allowing the flexibility of having higher cut offs, e.g. 0,15 or even 0,20, especially in the case of complex problems [21], [22].

Based on the above approach and given the large number of sub-criteria used in TALON's AHP, we have opted to use a CR of 0,2, thus discarded 5 survey responses, that have shown a higher CR value.

5.3.4 TALON AHP Survey results

The current section of the document presents the outcomes of the survey.

The weights criteria (Table 33) as well as each set of sub-criteria, are presented and discussed separately, while a global ranking of all sib-criteria is included in the end.

Starting with the Criteria, that is the overall TALON characteristics that are expected to drive its eventual adoption, the analysis reveals some very interesting insights.

The one that ranks higher is found to be is the AI Explainability, Trustworthiness & Transparency with a weight of 29%, a result that is compatible with the particular focus, placed nowadays in the trust related aspects of AI. It is however noteworthy that the second highest that is AI-fuelled orchestration, has very small difference, as it follows at 27%. The same small difference can also be observed in the two least weight criteria, that is ML-driven E2C Deployments & Runtime Adaptations with 20% and the AI Theoretical Framework & Benchmarking for Greener AI Deployments, that follows suit closely with 19% (Figure 20).

The above results do not show a strong preference on the part of the experts surveyed, which can be interpreted as a sign of all four dimensions of TALON having a rather comparable relevance when compared with each other. The highest weigh registered though cam certainly be explained though as a validation of the role that Explainability, Trustworthiness and Transparence are widely expected to play in AI-related adoption, in the future.

Table 33: Weights of criteria

Criteria	Weights
AI-fueled Orchestration	0,27
AI Theoretical Framework & Benchmarking for Greener AI Deployments	0,19
ML-driven E2C Deployments & Runtime Adaptations	0,20
AI Explainability, Trustworthiness & Transparency	**0,29**

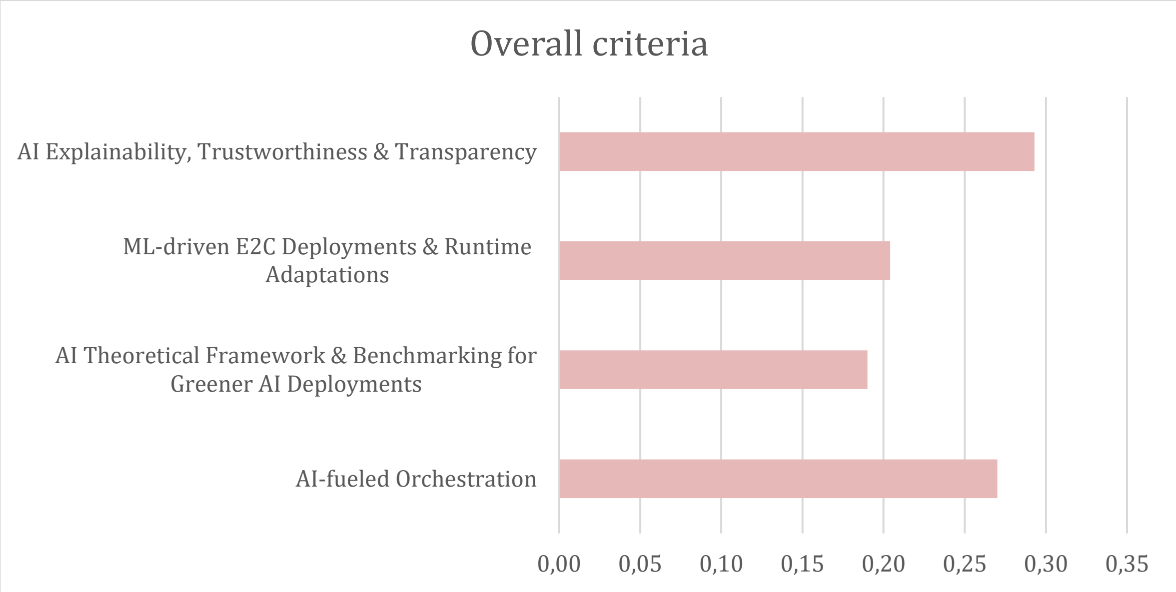

Figure 20: Weights of criteria

Attention now turns to the rankings of sub-criteria, pertinent to each one of the above four criteria.

Within the **AI-fueled Orchestration** criterion, the sub-criteria (Table 34) show a relatively balanced distribution of importance, with Intelligent Resource Management & Task Offloading getting 31%, and Fault Tolerance & Self-Healing 28% ranked highest. This indicates that efficiency in handling resources and ensuring reliability are central to effective orchestration. Decentralised Control Capabilities & Policy-enforcement at 22% is moderately valued, highlighting its supporting role in governance and distributed decision-making. Dynamic Service-level Awareness & Interpretation with 15% receives the lowest weight, suggesting that while service-level monitoring is relevant, it is not perceived as critical compared to operational and resilience-related factors. Overall, the weight distribution emphasises orchestration strategies that prioritise system performance and fault tolerance, complemented by governance mechanisms (Figure 21).

Table 34: AI-fueled Orchestration sub-criteria weights

Sub-criteria	Weights
Intelligent Resource Management & Task Offloading	**0,31**
Fault Tolerance & Self-Healing	0,28
Dynamic Service-level Awareness & Interpretation	0,15
Decentralised Control Capabilities & Policy-enforcement	0,22

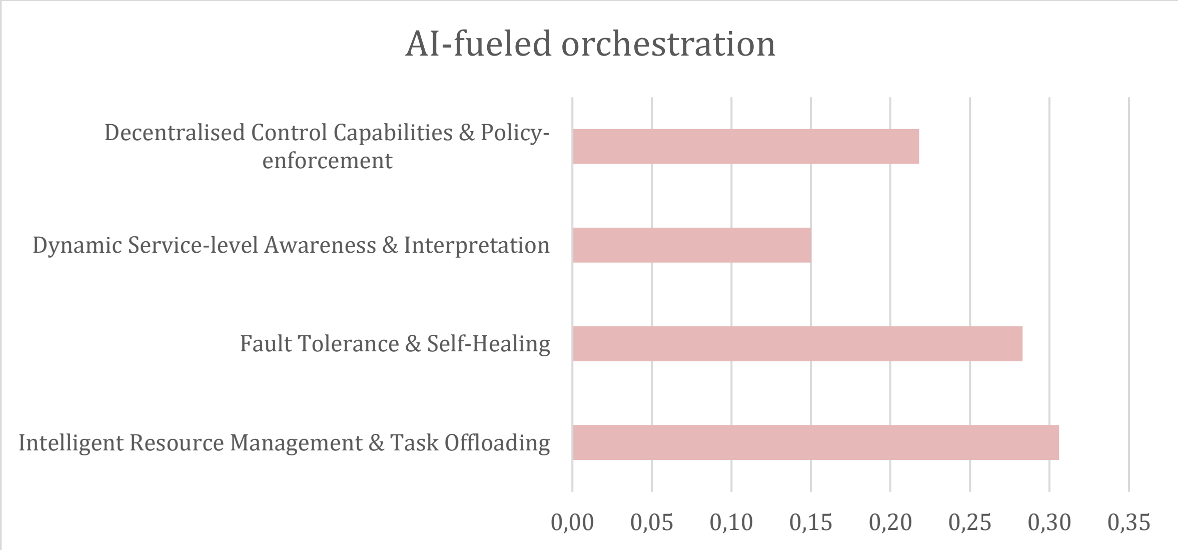

Figure 21: AI-fueled orchestration sub-criteria weights

The sub-criteria weights for the criterion **AI Theoretical Framework & Benchmarking for Greener AI Deployments**, derived from the AHP survey (Table 35), highlight the relative priorities among key components essential for sustainable AI development. Federated Learning / Decentralised Training emerges as the most influential sub-criterion with a weight of 23%, underscoring the emphasis on distributed computational paradigms to minimize energy-intensive centralized training. Closely following is Multidimensional Impact Assessment at 22%. Holistic AI Pipelines Lifecycle Management 18% and Multi-objective Evolutionary Methods & Optimisation Strategies 19% indicate balanced attention to end-to-end process oversight and advanced algorithmic efficiencies, while Methodological Transparency, Causality & Attribution at 14% receives the lowest priority, suggesting a need for enhanced emphasis on explainable and attributable green AI practices (Figure 22).

Table 35: AI Theoretical Framework & Benchmarking for Greener AI Deployments sub-criteria weights

Sub-criteria	Weights
Holistic AI Pipelines Lifecycle Management	0,18
Multidimensional Impact Assessment	0,22
Methodological Transparency, Causality & Attribution	0,14
Multi-objective Evolutional Methods & Optimisation Strategies	0,19
Federated Learning / Decentralised Training	**0,23**

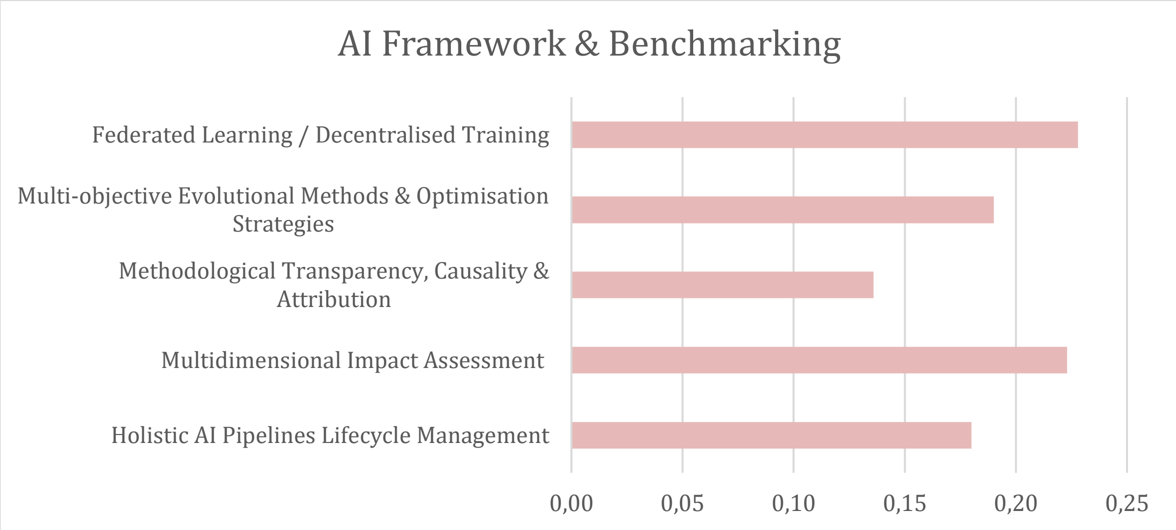

Figure 22: AI Theoretical Framework & Benchmarking for Greener AI Deployments sub-criteria weights

The sub-criteria weights (Table 36) for the criterion **ML-driven E2C Deployments & Runtime Adaptations**, derived from a TALON's AHP survey, highlight the prioritized elements for optimizing machine learning deployments in edge-to-cloud (E2C) environments. Intelligent Task Offloading leads with a weight of 26%, underscoring its critical role in dynamically distributing computational tasks to enhance efficiency and reduce energy consumption. Edge Deployments of Compute-intensive Tasks follows closely at 25%, emphasizing the importance of executing resource-heavy operations at the edge to minimize latency and cloud dependency. Deployment Strategy Optimisations, with a weight of 24%, reflect a strong focus on refining deployment approaches for performance and sustainability. Context-aware Task Switching, with the lowest weight of 21%, indicates a slightly lesser but still significant priority on adaptive task management based on real-time contextual data.

The sub-criteria weights show relatively small differences between them, indicating a balanced prioritization among the sub-criteria (Figure 23). This close distribution suggests that all aspects are considered nearly equally critical, reflecting a holistic approach to optimizing ML-driven edge-to-cloud deployments without a strongly dominant focus on any single factor.

Table 36: ML-driven E2C Deployments & Runtime Adaptations sub-criteria weights

Sub-criteria	Weights
Deployment Strategy Optimisations	0,24
Edge Deployments of Compute-intensive Tasks	0,25
Context-aware Task Switching	0,21
Intelligent Task Offloading	**0,26**

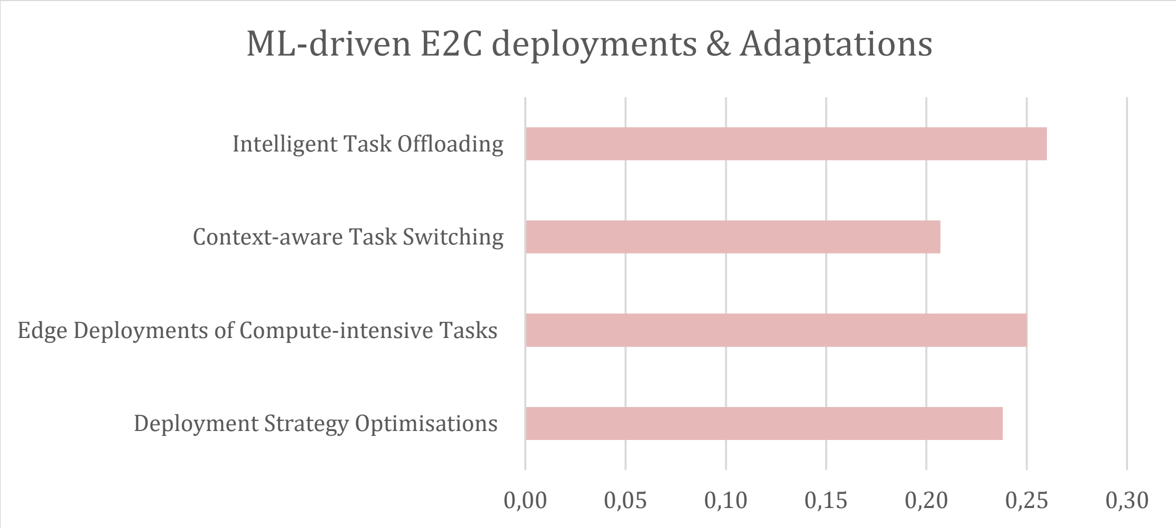

Figure 23: ML-driven E2C Deployments & Runtime Adaptations sub-criteria weights

The sub-criteria weights for the criterion **AI Explainability, Trustworthiness & Transparency**, as derived from TALON's AHP survey (Table 37), highlight the components to be prioritised for fostering reliable and transparent AI systems. Reliability & Robustness of AI Models and Fairness & Bias Mitigation share the highest weight at 26% each, emphasizing their critical roles in ensuring consistent performance and equitable outcomes. User-centric Explanation Quality & Effectiveness follows with a weight of 21%, underscoring the importance of clear and accessible explanations for end-users. Algorithmic AI Model Transparency at 12% and Interpretability of AI Model Internals 10% receive the lowest weights, suggesting a relatively lower but still significant focus on technical transparency and internal model comprehensibility.

The equal weighting of Reliability and Fairness sub-criteria highlights their very comparable and actually identical importance, while the slightly lower weights for transparency and interpretability reflect a nuanced but still significant focus, pointing to a holistic approach in achieving explainability, trustworthiness, and transparency in AI systems and deployments in general.

Overall, the results shown in the table below, suggest that stakeholders view all sub-criteria as interlinked and nearly equally critical for trustworthy AI systems (Figure 24).

Table 37: AI Explainability, Trustworthiness & Transparency sub-criteria weights

Sub-criteria	Weights
Algorithmic AI Model Transparency	0,12
Interpretability of AI Model Internals	0,10
User-centric Explanation Quality & Effectiveness	0,21
Reliability & Robustness of AI Models	**0,26**
Fairness & Bias Mitigation	**0,26**

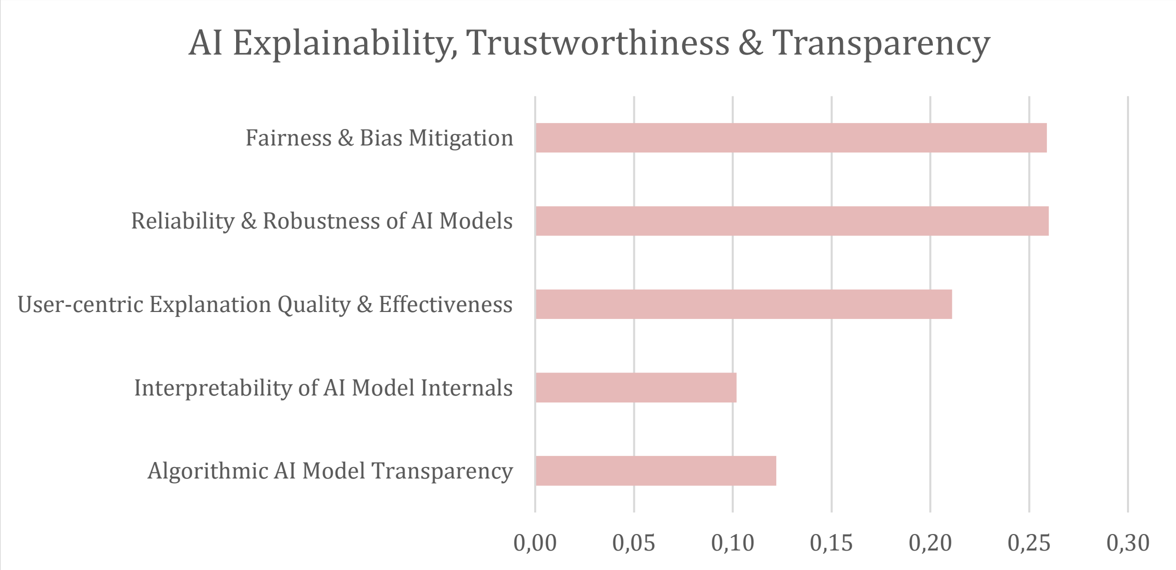

Figure 24: AI Explainability, Trustworthiness & Transparency sub-criteria weights

After the analysis of the criteria and sub-criteria results, and as a final step, it is worth taking a look at the **Global weights of sub-criteria** and based on that create a global hierarchical ranking of all sub-criteria. To calculate that we multiplied that weight of each individual sub-criterion with their associated criterion weight.

The table below (Table 38) presents the Global ranking of sub-criteria.

Table 38: Sub-criteria global weights

Sub-criterion	Global weight	Global rank
Intelligent Resource Management & Task Offloading	0,083	1
Fault Tolerance & Self-Healing	0,076	2
Reliability & Robustness of AI Models	0,076	3
Fairness & Bias Mitigation	0,076	4
User-centric Explanation Quality & Effectiveness	0,062	5
Decentralised Control Capabilities & Policy-enforcement	0,059	6
Intelligent Task Offloading	0,053	7
Edge Deployments of Compute-intensive Tasks	0,051	8
Deployment Strategy Optimisations	0,049	9
Federated Learning / Decentralised Training	0,043	10
Multidimensional Impact Assessment	0,042	11
Context-aware Task Switching	0,042	12
Dynamic Service-level Awareness & Interpretation	0,041	13
Multi-objective Evolutional Methods & Optimisation Strategies	0,036	14
Algorithmic AI Model Transparency	0,036	15
Holistic AI Pipelines Lifecycle Management	0,034	16
Interpretability of AI Model Internals	0,030	17
Methodological Transparency, Causality & Attribution	0,026	18

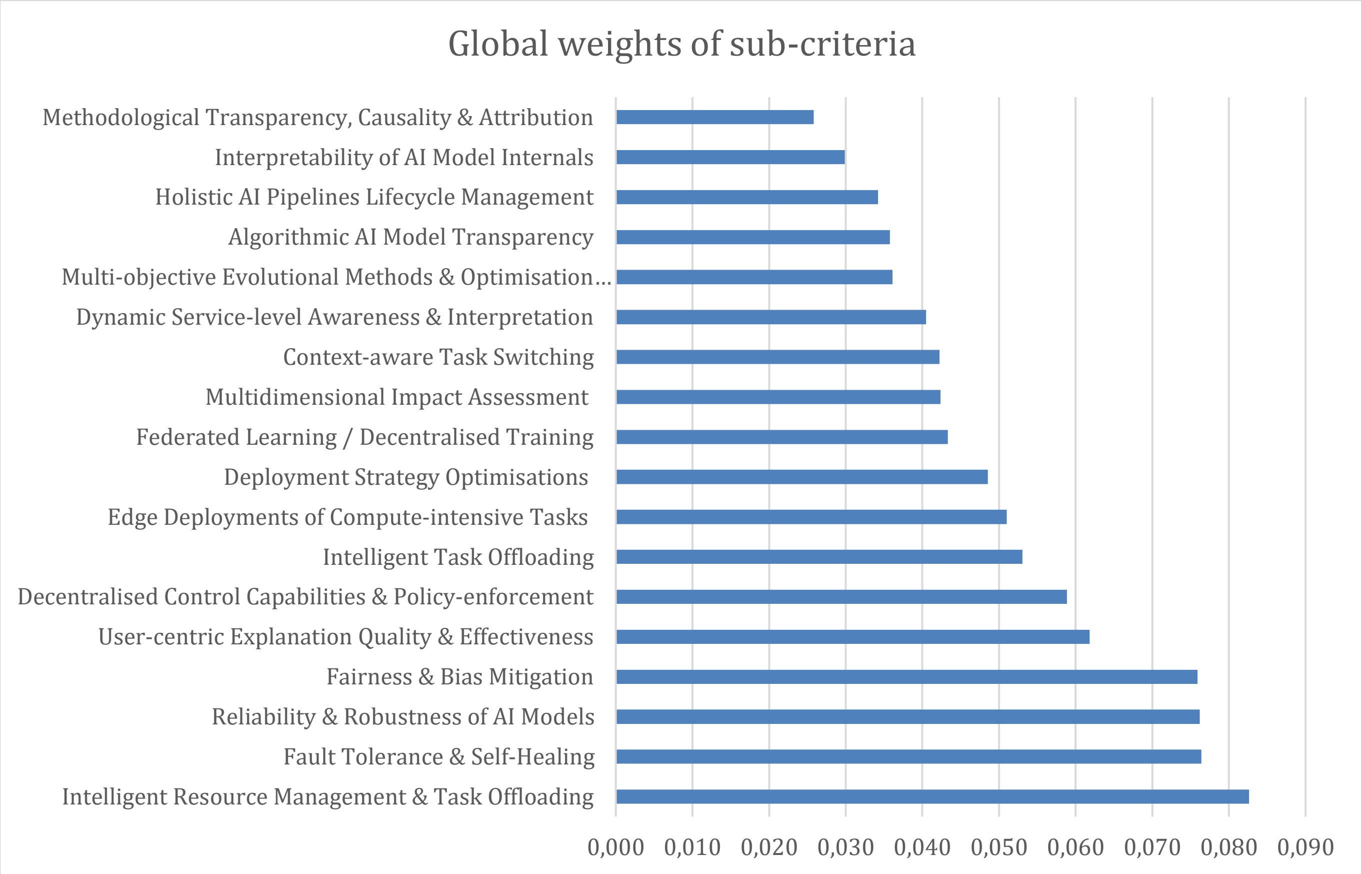

Figure 25: Global weights of sub-criteria

The sub-criteria Global weights show some meaningful insights.

The one leading the global ranking is the Intelligent Resource Management & Task Offloading (#1), followed closely by Fault Tolerance & Self-Healing (#2), Reliability & Robustness of AI Models (#3) and Fairness & Bias Mitigation (#4) all three of which have virtually identical weights (up to three decimals). This observation is noteworthy in that these considered by the experts equally important.

Another observation that is especially important to be highlighted is that the globally top 6 ranked sub-criteria, are all part of two TALON main criteria, namely; AI-fueled orchestration and AI Explainability, Trustworthiness & Transparency. Moreover, the game is evenly shared between these two, as evidenced by the fact that each of these two, have exactly three globally top-ranked sub-criteria.

As a conclusion, the above findings in Figure 25 show some powerful indications in the possible directions to be prioritized in creating a roadmap for TALON's future and development. All the more so given the eventual adoption maximizing perspective, from which this survey has been considered.

6 Standardisation

As part of its impact creation activities, TALON invested in efforts that demonstrate to greater audiences the project outcomes that have clear implications with respect to standards.

The current section of the document, provides an account of the activities realised.

6.1 Standardisation Webinar

On April 7, 2025 TALON arranged a webinar on the standardisation aspect of trustworthiness in AI, titled: "AI You can rust: Standards, Ethics and Innovation".

The webinar had as an objective to shed light on recent standards for Trustworthy AI and will present relevant results of EU projects, including notably TALON's perspectives.
It featured a select number of speakers, representing academia, industry and research projects, and with an active interest in the field, namely:

- **Prof. Christos Emmanouilidis**, University of Groningen and HumAIne HE project. Gave a quick overview of relevant standards (notably from IEEE-SA) and discussed how HumAIne results align with them.
- **Lior Limonad**, IBM Research lab. Gave an overview of IBM's approach to Situation Aware Explainability and Trustworthy Casual AI Processes, including their relevance to standards, as applicable.
- **Dr. Sophia Karagiorgou,** UBITECH, and Technical Coordinator TALON**.** Introduced TALON's approach and highlighted potential future directions.
- **Prof. Nineta Polemi,** University of Piraeus and CTO at Trustilio. She focused on the challenges and efforts, relevant to trustworthiness in AI.

The discussion was moderated by **Ioannis Soldatos**, of Netcompany-Intrasoft, a TALON project partner.

The webinar (Figure 26) was widely communicated, through project's own social media, project partners' communication channels plus of course through the TrustWorthy AI cluster and its participating projects, For the event registrations Eventbrite was used, hence increasing the webinars visibility and reach [23] as was evidenced by the high number of both registrations (80+) and attendance (50 persons).

The presentations during the webinar managed to spark great interest of the participants and give rise to discussions and exchanges. All this made for a highly engaging event that was about two hours long in its duration.

The webinar participants raised a number of key points that are highly relevant to the AI standards evolution debate. The following topics were especially touched upon:

- The need for Trustworthy AI to involve not only technical robustness, but also adherence to ethical norms, regulatory compliance, and alignment with societal values.
- Standardization efforts and regulatory frameworks are critical in shaping responsible AI deployment, as evidenced by the result of the EU-funded project discussed.
- Collaboration is crucial in the sense that no single actor can define or enforce trustworthy AI alone. Joint efforts and shared frameworks are needed to scale impact.
- Innovation and responsibility need to co-evolve. The pace of AI advancement requires flexible, forward-looking governance models that can adapt without stifling progress.

Finally, the webinar recording is currently made available on TALON's YouTube page.

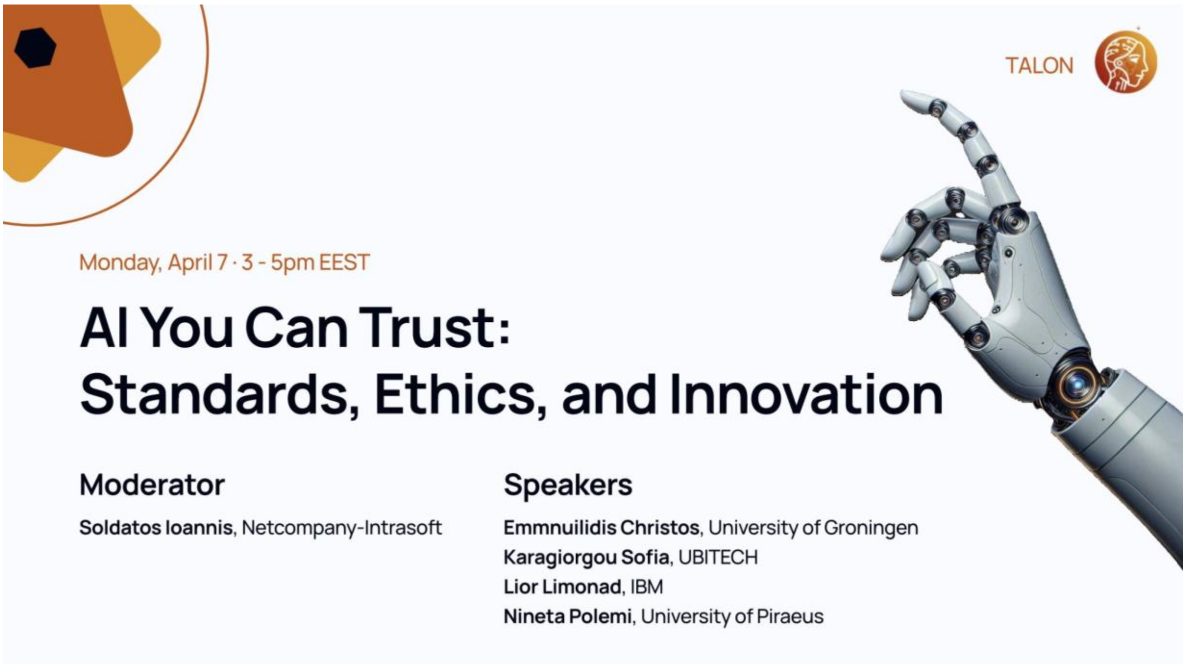

Figure 26: AI You can trust webinar banner

6.2 Standardisation White paper

In order to document the standardisation prospects coming out of the project, in a manner that can be made available, to a greater audience, TALON project partners have opted for publishing a relevant white paper.

This publication provides a comprehensive overview of the project's relevant innovations in:

- Explainable AI (XAI),
- Immersive AR/VR systems
- Edge-to-Cloud (E2C) AI orchestration
- Blockchain-based AI model governance

It further details the multi-phase strategy adopted by each respective asset owner to translate these outcomes into actionable contributions to standards development organizations such as IEEE, ISO/IEC, and CEN-CENELEC. The TALON partners primary in charge of the standardisation of each of the above are: Metaminds, Kingston University, Engineering (jointly with the University of Luxembourg and Sidrocco, respectively.

More specifically, the paper's objective is to provide a consolidated view of TALON's contributions to the international standardization landscape, with a specific focus on the technological innovations developed within the project's framework. It aims to:

- Present the standardization potential of TALON's above named, select key results,
- Outline the strategic and procedural approach adopted to engage with relevant standard development organizations.
- Identify the alignment of TALON outputs with specific standards bodies, working groups, and ongoing specifications.
- Define the implementation strategy and expected impact of these standardization efforts on industrial automation ecosystems

By achieving the above objectives, the white paper aims to support stakeholders, both within and outside the TALON consortium, in understanding the project's standardisation roadmap, its technological maturity, and its role in shaping future industry norms.

The standardisation roadmap together with the intermediate steps is best summarized in the following figure (Figure 27).

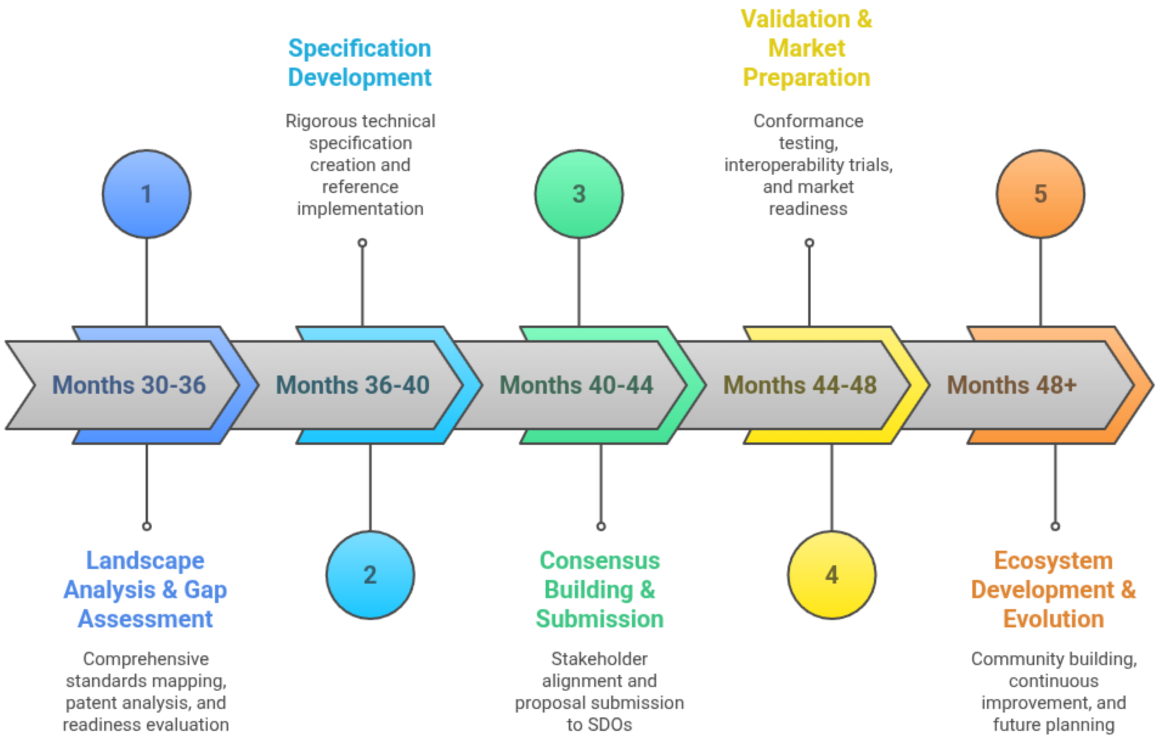

Figure 27: TALON's pathway to standardisation

The standardization pathway discussed in the white paper also incorporates robust risk mitigation strategies, proactive engagement with stakeholders, and a strong open-source orientation, ensuring that TALON's outputs remain adaptable to evolving technological and regulatory contexts. The project's emphasis on ecosystem development, interoperability testing, and continuous evolution further reinforces its commitment to long-term sustainability and relevance.

The white paper has been uploaded on Zenodo [24], from where it can be publicly accessed.

6.3 Other Standardisation Activities

In addition to the joint activities presented above, TALON partners have also been active, in their own right, as individual entities in various ways that support the further enhancement of TALON's footprint with a standardisation impact.
The activities and efforts reported herewith, have been largely drawing on each partner's respective background, profile and current interests.

Kingston University has monitored standardisation efforts related to Virtual Reality (VR) training and Augmented Reality (AR) maintenance using Edge Artificial Intelligence, with particular focus on activities within CEN-CENELEC. The developments tracked include areas such as real-time data processing, spatial computing, and edge deployment standards for XR applications. Although participation in webinars, seminars, or meetings has been rather limited, ongoing monitoring has helped inform the technical alignment of our VR and AR solutions with emerging European and international standardisation trends.
CERTH participated in meetings, webinars, and expert discussions, providing feedback and recommendations on standardization actions and procedures. These contributions focused on

practical insights gained from research in AI model development, and multimodal data fusion in other Horizon projects and experience gained from there.

Even though they did not initiate a formal standard dedicated to the tools that were developed in the TALON project but contributed to the broader standardization dialogue by sharing expertise and offering input to ongoing initiatives. Future work may leverage TALON outcomes to support standardization in areas such as explainable AI, trustworthy AI frameworks, and safety in human–robot collaboration.

On the part of TALON's technical coordination lead by UBITECH has played a key role in advancing AI governance frameworks. Notably, Sophia Karagiorgou, represented UBITECH and the TALON technical team on the Board of Directors within the AI Chapter of the Industrial Community at Adra. In this capacity, she demonstrated visionary leadership by actively shaping the agenda for AI standardization. Her proposal of a comprehensive set of standards offers a strategic foundation for harmonizing practices across the AI community, fostering reliability, and ethical implementation.

Moreover, they supported development and alignment of standards within the TALON initiative, by actively contributing to critical technical discussions that underpin the project's success and continuous alignment with its ambitious objectives.

InnoCube's standardization effort employed a holistic multi-domain approach across three major standardization communities, namely ETSI, IEEE SA, and One6G, propelling the company in next-generation communication and distributed systems development.

Open-Source Networking through **ETSI**	As a participant member of ETSI's Software Development Group OpenOP (SDG OOP), InnoCube is involved in the development of the Open-Source Operator Platform for operator network and testbed federation capabilities. This Open-source Operator Platform contains capability exposure APIs as a base reference, from which different requirements can be added for interoperability into different - networks environment. Through this activity, InnoCube is leveraging training that came out of the TALON project and providing concrete activities to the open-source networking community.
Blockchain and Distributed Ledger Technology Innovation through **IEEE SA**	In the context of IEEE Blockchain standardization, InnoCube's focus is on following standards for blockchain and distributed ledger technologies. InnoCube's main focus is on process identification and providing standards, recommendations, and guidelines relevant to blockchain implementation and distributed ledger operations. This work would secure interoperability, security, and scalability across networks of blockchain and produce best practices for using the distributed ledger as use case in telecommunications and related areas.
6G Technology Leadership through **One6G**	InnoCube participates in One6G, focusing on standardization activities around the essential elements of 6G wireless communication. Activities include 6G connectivity architecture, computational architectures, optimizing for energy consumption, and communication protocols including wireless and cellular access networks. The standardization activities suggest a focus on multi-faceted aspects of security, trust, functional safety, privacy, and authentication aspects. In addition, InnoCube followed distributed cloud architecture standardization and definition of technology association frameworks, radio technology definitions, and standardized protocols with unified data formats. By strategically monitoring the Working Group 1, InnoCube reflects on future use cases, performance KPIs, and evolving market and business use cases that will define the 6G ecosystem.

Finally, Pobotek has been monitoring standards: EASA, ETSI (MEC/ENI), ISO/IEC JTC 1/SC 42, SC 41, EU AI Act.

7 Regulation

The current section of the document presents points and findings of Regulation related developments. This covers different areas of regulation that is both currently established, or ongoing or even expected to happen in the future, while being relevant to TALON's area of interest and activity, directly and/or possibly less directly.

In this context, particular attention has been paid to the EU AI act, both for its increased pertinence to TALON's AI-driven orientation, as well as to its notable global significance as a comprehensive AI targeted regulation.

7.1 EU AI Act

The European Union's Artificial Intelligence Act (EU AI Act) represents the first comprehensive legislative framework worldwide aimed at regulating artificial intelligence. It follows a risk-based approach, ensuring that AI technologies deployed within the EU are safe, trustworthy, and compliant with fundamental rights. As such, the Act is probably the single most relevant regulatory act to date, with respect to TALON, as it outlines the regulatory environment in which future AI systems will be assessed once they reach market maturity, and relevant developments were closely followed throughout TALON's lifetime.

In terms of its scope, the EU AI Act applies to all operators in the AI value chain that place AI systems on the market, put them into service, or otherwise make them available within the EU. The scope extends to both Union and non-Union actors, meaning that providers established outside the EU are equally bound by its provisions if their systems are made available within the EU market.
The regulation applies across sectors, addressing AI in both the private and public domain. It also takes into account the entire AI system lifecycle, from design, development, and training, through placement on the market, to deployment and post-market monitoring. Its provisions are designed to complement existing sector-specific EU legislation, such as the Machinery Regulation or the Medical Devices Regulation.

Even though it is an initiative of European origin, the EU AI Act has strong Global significance. Within the EU, the Act provides legal certainty and harmonization, enabling innovators to operate under a unified set of rules across the single market. It is expected to strengthen public trust in AI by introducing safeguards on transparency, accountability, and oversight.

Globally, the EU AI Act is of strategic significance. Similar to the General Data Protection Regulation (GDPR), the Act is expected to influence international regulatory practices and become a de facto global standard for AI governance. Non-EU companies that wish to access the European market will need to comply with its provisions, creating a strong extraterritorial effect and raising global standards.

The EU AI Act takes a risk-based approach to the AI applications, by distinguishing AI systems according to the level of risk they pose to health, safety, and fundamental rights depicted in Figure 28:

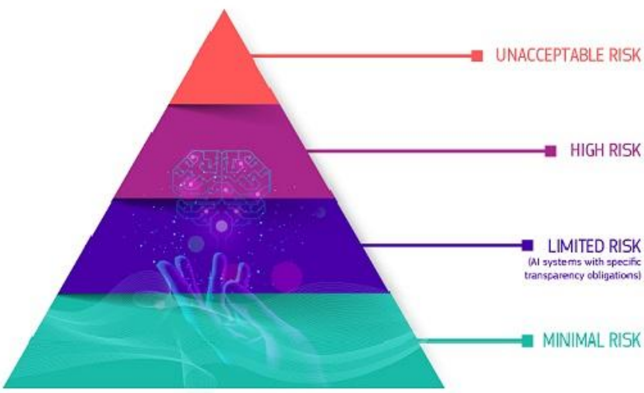

Figure 28: The EU AI Act risk pyramid (source: European Commission)

- **Minimal (or No risk):** The vast majority of AI systems currently used in the EU fall into this category. This includes applications such as AI-enabled video games or spam filters. The AI Act does not introduce rules for AI falling into this category.

- **Limited risk:** Systems that pose low but non-negligible risks, such as chatbots or deep fake generators. They are subject to transparency obligations (e.g., informing users they are interacting with AI).

- **High risk:** Systems with potential significant adverse effects on individuals or society, such as those used in critical infrastructure, education, employment, essential public services, and law enforcement. These systems are subject to extensive requirements.

- **Unacceptable risk:** Systems that are deemed harmful or incompatible with EU values, such as social scoring by governments or manipulative AI that exploits vulnerable groups. These are outright prohibited.

High-risk AI systems are explicitly listed in the Annexes of the EU AI Act. They include AI used in products already regulated under EU safety legislation (Annex I), as well as stand-alone AI applications in critical domains (Annex III) such as:

- Management and operation of critical infrastructure (e.g., energy supply).
- Education and vocational training (e.g., AI systems determining access to education).
- Employment, workers' management, and access to self-employment (e.g., CV-sorting systems).
- Access to essential services (e.g., credit scoring).
- Law enforcement, migration, and border management.

Providers of such systems are subject to strict compliance obligations covering risk management, data quality, documentation, human oversight, and post-market monitoring.

In addition to distinguishing the AI systems based on the specific application they are put in, the EU AI Act assigns obligations depending on the role of an actor in the AI value chain.

To that end, a number of profiles are defined, as shown below, with the corresponding obligations:

- **Provider**: The natural or legal person who develops an AI system or has it developed and places it on the market or puts it into service under their own name or trademark (Article 3). Providers bear the primary responsibility for ensuring compliance with the Act.

- **User (or Deployer)**: The entity that uses an AI system under its authority, except where the system is used in a personal non-professional capacity. Users must follow instructions for use and ensure systems are applied in compliance with the regulation.

- **Importer**: An entity established in the Union that places an AI system on the EU market which originates from a third country. Importers are responsible for verifying that the provider outside the EU has ensured compliance.

- **Distributor**: Any person in the supply chain who makes an AI system available on the market, without being the provider or importer. Distributors must verify conformity and report risks.

- **Authorized Representative**: A person or entity established in the EU who is designated by a provider outside the EU to act on their behalf regarding compliance matters.

These roles clarify the allocation of legal obligations, with the provider carrying the most comprehensive set of responsibilities.

With the EU AI Act offering a comprehensive framework, it is interesting to see how all its constituent part fit together, with respect specifically to TALON plus whether and in what ways these may be applicable in the case TALON. Starting with its positioning, by considering TALON as a unified, complete solution that is eventually made available to interested end-users, say marketed through a chosen business model (e.g. as a service), then in relation to the EU AI Act framework, TALON (its consortium) would be best described from the perspective of a provider.

As far as the risk categorization is concerned, the EU AI Act examines this always with respect to a system's specific deployment context and precise application, i.e. not the given system per se. Thus, whether or not TALON may eventually qualify as high-risk will depend solely on the specific application it is put in, and therefore it is not something that can be ascertained beforehand.

In the meantime, however, and given that TALON is currently an R&D project, it is exempted for any EU AI Act obligation. According to Article 2(6) of the EU AI Act, AI systems created and utilized solely for scientific research purposes are excluded from the Act's regulatory requirements. This exemption is designed to encourage innovation and scientific freedom by allowing experimentation and development under conditions that may not fit within traditional regulatory frameworks.

On the assumption of TALON's system, in a market-ready form, being categorized as high-risk within the context of given usage scenario, then the consortium (as a provider) would need to comply with a comprehensive set of obligations.

Taking a closer look at the AI Act provisions on high-risk AI systems, Articles 9 to 15 l (Section 2, Chapter III) lay down the specific technical and organizational requirements such as risk management, data governance, transparency, accuracy, robustness, cybersecurity, and human oversight:

- **Article 9** requires establishing, implementing, documenting, and maintaining a continuous risk management system throughout the AI system's lifecycle, covering risk identification, analysis, mitigation, testing, and post-market monitoring.
- **Article 10** sets requirements for data and data governance to ensure training, validation, and testing datasets are relevant, representative, error-free, complete, and appropriately managed to prevent bias.
- **Article 11** mandates providers to maintain comprehensive technical documentation demonstrating conformity with the AI Act requirements.
- **Article 12** ensures transparency by requiring the provision of clear information to users and deployers about the AI system's capabilities and limitations.
- **Article 13** governs transparency and provision of information to users to ensure AI system users are adequately informed for safe use.
- **Article 14** demands human oversight measures to enable human operators to interpret, intervene, or override the AI system when necessary to prevent or minimize risks to fundamental rights.
- **Article 15** requires the AI system to be designed and developed with sufficient accuracy, robustness, and cybersecurity, ensuring consistent performance and resilience against manipulation or errors throughout its lifecycle.
- **Article 16** (also part of Section 2) further consolidates these requirements by summarizing the main obligations of providers of high-risk systems and cross-referencing the relevant articles. Thus Article 16 serves as the central reference point for providers, outlining main responsibilities:

The main obligations in Article 16 of the EU AI Act for providers of high-risk AI systems are:

- Ensure compliance of their AI systems with the requirements set out in Section 2 of the Act.
- Indicate their name, registered trade name or trademark, and contact address on the AI system, its packaging, or accompanying documentation.
- Have a quality management system in place.
- Keep the required technical documentation.
- Retain logs automatically generated by the AI system when under their control.
- Ensure the AI system undergoes the relevant conformity assessment before being placed on the market or put into service.
- Draw up an EU declaration of conformity.
- Affix the CE marking to indicate conformity with the regulation.
- Comply with registration obligations.
- Take necessary corrective actions and provide information as required.

- Demonstrate conformity with Section 2 requirements upon request by authorities.
- Ensure the AI system complies with accessibility requirements as per relevant EU directives.

As discussed previously, ascertaining whether an AI system can be categorised as high-risk or not, is primarily based on the context of use. However, a concise evaluation necessitates a structured approach and assessment of the given application considered as well as technical properties of the system, especially because certain practices (e.g. social scoring) are never allowed, whilst Transparency rules (like labelling AI-generated content) apply whether the system is "high-risk" or not. So as to facilitate the process, and for future, concrete usage contexts TALON has created a questionnaire that cover all relevant aspects to be considered from the angle of a Provider. The questionnaire can be found in the appendix of this document [Include the ENG provided questionnaire]

In terms of its current status (as of the time of writing) and actual implementation timeline, the Act officially became law on 1 August 2024—20 days after publication in the Official Journal [25], while core definitions of AI, prohibited practices, and AI literacy obligations apply effective from 2 February, 2025. General-Purpose Artificial Intelligence (GPAI) rules around its governance, notification authorities, confidentiality, penalties, and requirements for GPAI models have taken effect effective from 2 August, 2025, and Obligations related to high-risk AI – such as risk management, documentation, CE marking, conformity assessments – will come into force as of 2 August 2027 [26].

7.2 Circular economy

One of the core objectives of TALON is the integration of practices that enable Green Industry features through sustainable mechanisms at every stage of its components' development. Such a scope is directly aligned with Circular Economy (CE) principles, a sector that has been greatly promoted by the European Commission (EC) during the last years. Specifically, the key concept of CE is the implementation of re-functionalities, e.g., recycling, remanufacturing, reusing, repairing, reducing, as well as the utilization of technological methodologies that enable such practices. Consequently, as European policies and regulations have rapidly progressed in recent years, TALON ensured compliance and monitoring of such regulations, motivating its established sustainable strategy.

To begin with, TALON directly supports the "European Green Deal" [27], a detailed set of policies published by the EC in 2019, aiming towards the significant reduction of greenhouse gas emissions by 2050, and the differentiation of resource utilization from the financial growth of Europe as a whole. As stated within the "European Green Deal", digital technologies are key enablers of its specific objectives, and TALON ensured proper alignment with such strategies through the integration of energy-efficient edge AI and Greener AI mechanisms that address important aspects, such as carbon footprint reduction.

Moreover, the Circular Economy Action Plan (CEAP) is one of the main pillars that comprise the "European Green Deal", published in 2020 [28], with a revised framework being released in 2023 [29], and focusing on proper waste management, sustainability in product designing, while promoting the reusability of materials. TALON effectively embodied CEAP principles, with its architecture being energy and data-efficient, resulting in the optimization of computational resources, and with its respective AI orchestrator and the reusability of AI components reflecting CE ideals. Similarly, it employs eco-conscious deployment of advanced technologies like its lightweight blockchain scheme for efficient data sharing, as well as the utilization of digital twins for the optimization of industrial operations.

Additionally, the Ecodesign for Sustainable Products Regulation (ESPR), which was established in 2024 [30], introduces the measure of Digital Product Passport (DPPs), a digital record that contains all the data generated through the lifecycle of a product, as a means of enhancing its circularity, energy efficiency, and related re-functionalities. TALON's AI orchestrator is capable of efficient energy and data management, with implemented reusability of its AI assets, fulfilling the requirements of the available quantitative metrics needed for the ESPR's DPPs. Also, its lightweight blockchain integration enables real-time supply chain transparency for the creation of auditable records, complying with the related mandatory monitoring stated by the ESPR.

Finally, following the vision of TALON on the implementation of a technologically advanced Greener Industry, the project is closely tracking the upcoming “Circular Economy Act”, which is expected to be launched by the end of 2026. The particular legislation will focus on further enhancing the reusability of materials in the EU’s economy, while aiming towards becoming a global leader in the Circular Economy by 2030 [31]. Based on the initial information provided by the EC, the “Circular Economy Act” will motivate industries across Europe to implement digitalization practices within their infrastructures, especially for achieving proper management and shipment of waste, while placing an additional emphasis on electronic devices, and establishing re-functionalities as mandatory practices. As a result, TALON will maintain tracking of the developments of the corresponding act to ensure that its outcomes remain fully synchronized with the forthcoming legislation.

7.3 European regulations for Industrial drone applications

As industrial applications nowadays often feature the use of drones, and this has also been the case in one of TALON’s usage scenarios, reviewing the European regulations for industrial drone applications, becomes especially relevant.

In recent years, the use of drones, or Unmanned Aircraft Systems (UAS), has grown rapidly across industrial sectors in Europe. Companies rely on drones for a wide range of missions such as infrastructure inspections, port and rail surveillance, security patrols, and linear asset monitoring like pipelines or power lines. Because of the potential risks associated with flying drones—ranging from safety concerns in shared airspace to data protection issues—the EU has established a comprehensive regulatory framework. The current rules aim to balance safety, privacy, and technological innovation, while enabling industries to benefit from drone-based solutions.

At the foundation of this regulatory system are two cornerstone regulations. The first is Regulation (EU) 2019/947 on the rules and procedures for the operation of unmanned aircraft, which governs operational procedures and defines categories of operation and the associated requirements. The second is Regulation (EU) 2019/945 on unmanned aircraft systems and third-country operators, which specifies product and class requirements for drones, including technical standards and compliance obligations.

Together, these regulations provide a structured approach that differentiates operations by the level of risk they present, while ensuring drones are designed and flown in accordance with safety expectations.

In terms of specific drone operations these are divided in the EU into three broad categories: Open, Specific, and Certified (Regulation 2019/947). This categorization reflects the potential risk each type of operation carries and allows regulators to tailor requirements accordingly.

The Open category is intended for low-risk activities. It covers small drones operating within visual line of sight and without flying over gatherings of people. Sub-categories (A1, A2, and A3) further restrict parameters such as maximum weight, altitude, and proximity to people. While the Open category may be suitable for simple industrial tasks like basic site documentation or visual checks on private property, its restrictions make it impractical for most advanced industrial missions. For example, activities such as long-distance pipeline inspections or security patrols that require beyond visual line of sight (BVLOS) operations fall outside the scope of Open category allowances.

The Specific category is where most industrial applications are situated. This category covers medium-risk operations that go beyond the limitations of the Open category, such as BVLOS inspections, security sweeps, or repetitive missions over controlled industrial areas. Operators in the Specific category must obtain an Operational Authorisation from their National Aviation Authority, usually through the Specific Operations Risk Assessment (SORA). To simplify this process, the European Union Aviation Safety Agency (EASA) has developed standardized tools such as Predefined Risk Assessments (PDRAs) and Standard Scenarios (STS), which operators can adopt instead of conducting a full SORA (EASA AMC/GM to Reg. 2019/947). These provide quicker, harmonized pathways to authorization, particularly for repetitive or cross-border operations.

The Certified category applies to high-risk activities, including drones designed to carry passengers or heavy cargo, as well as operations over large gatherings of people. In practice, the Certified category is similar to manned aviation, requiring type certification of the aircraft and licensing akin to traditional aviation rules. For now, it remains a relatively rare category in industrial contexts, but it will become increasingly relevant as drone technology evolves toward heavy logistics and passenger applications (Reg. 2019/947).

Within the Specific category, the EU framework provides several mechanisms to streamline the authorization process. Standard Scenarios (STS-01 and STS-02) set out harmonized conditions under which drones can be operated safely, using designated drone classes such as C5 and C6. If an operator can design missions that fit within these predefined conditions, approval is faster and less complex (EASA AMC/GM to Reg. 2019/947).

In addition to STS, Predefined Risk Assessments (PDRAs) serve as EASA-issued templates that cover common industrial use cases, such as surveillance, agriculture, or short-range cargo transport. These PDRAs allow operators to benefit from risk assessments that have already been validated at the European level, cutting down the administrative workload and providing greater consistency across Member States. For companies running large-scale or repetitive industrial operations across borders, PDRAs offer a particularly efficient route (EASA, PDRAs).

For highly experienced operators managing large fleets or repeated mission types, the Light UAS Operator Certificate (LUC) represents an even greater level of flexibility. With a LUC, an operator effectively gains the authority to self-approve certain categories of operations within the scope of their certificate. This reduces dependence on continuous oversight by national regulators and allows large-scale programs to expand more efficiently across Europe. However, obtaining and maintaining a LUC requires the operator to implement a strong safety management system and provide evidence of robust compliance practices (Reg. 2019/947, Part C).

Another major development in the European drone regulatory environment is the introduction of U-space, a system for digital traffic management of drones. The relevant legal foundation is Commission Implementing Regulation (EU) 2021/664, which defines how U-space services are to be implemented. From 2023 onwards, Member States may designate certain airspaces as U-space zones, particularly in busy or complex environments such as ports, cities, or industrial corridors. In these areas, drone operators must connect to U-space services, which provide essential tools like network-based identification, geo-awareness, traffic information, and conformance monitoring.

For industrial operators, this means preparing for integration with U-space technologies. Drones will need to support features such as network Remote ID and interfaces for sharing flight intentions. Similarly, operators will need to adapt their procedures to ensure compliance with U-space requirements. While the system is still expanding, it is expected to become a central feature of drone operations near critical infrastructure and logistics hubs (Reg. 2021/664) .

Operating drones in Europe involves a range of obligations beyond flight permissions. Compliance with CE markings and class requirements (C0 to C6) is essential for drones used in the Open category or under STS. In the Specific category, class marking may not always be mandatory, but regulators may still require proof of airworthiness or design verification for higher-risk operations (Reg. 2019/945) [32].

Registration requirements apply to operators rather than individual drones, simplifying the process for organizations managing fleets. In many cases, operators must also ensure their drones provide Remote ID functionality, either directly or through networked solutions, especially in U-space areas (Reg. 2019/947 and 2021/664).

Pilot competence is another cornerstone. Depending on the type of operation, pilots may need to pass online exams (for Open category flights) or undergo additional theoretical and practical training defined by the National Aviation Authority for Specific category missions (Reg. 2019/947, Annex A).

Geo-zones, or designated geographical areas where drone use is restricted, also play an important role. These zones may include airports, sensitive government sites, or critical infrastructure. Operators are responsible for checking these zones before each flight and ensuring they comply with altitude restrictions, prior coordination requirements, or U-space entry rules (Reg. 2019/947, Article 15).

Beyond aviation-specific requirements, drone operators must also comply with other legal frameworks. Insurance is generally mandatory for industrial operations, ensuring liability coverage in the event of accidents (national implementations under EU framework).

Data protection and cybersecurity obligations are equally important, since industrial drones frequently collect imagery or other data that may include personal information. Under the General Data Protection Regulation (GDPR, Reg. 2016/679), companies must implement safeguards such as data minimization, encryption, controlled access, and retention policies. U-space systems, by their nature, add further obligations around information security.

Finally, drones carrying high-bandwidth communication equipment must conform to the Radio Equipment Directive (RED, Directive 2014/53/EU), ensuring safe and compliant use of wireless spectrum.

Cross-Border Operations
One of the key strengths of the EU's drone regulatory system is its support for cross-border operations. Under Regulation 2019/947, an operator can obtain authorization in one Member State and have it recognized across the Union, subject to notifying the National Aviation Authority in the destination country. Standardized pathways such as STS and PDRAs were designed specifically to reduce the burden of reapplying for authorization in each jurisdiction, thereby facilitating pan-European industrial deployments.

Application to TALON Drone System
The regulatory framework has clear implications for industrial drone platforms such as TALON, which is designed for missions like BVLOS perimeter patrols, linear infrastructure inspections, and multi-sensor security sweeps. For TALON-class missions, the most appropriate regulatory pathway is usually the Specific category, relying on PDRAs where possible for efficiency, or a full SORA for novel mission profiles (EASA PDRA/SORA framework).

When mission conditions permit, TALON operators can benefit from designing operations to fit within STS requirements, using drones in the C5 or C6 class. This allows faster authorization and smoother repetition of missions across multiple sites (EASA STS-01/02).

At the same time, TALON must prepare for U-space integration, ensuring compatibility with requirements such as network Remote ID, flight intent filing, and conformance monitoring. Building these features into the system architecture from the start will ensure long-term compliance as U-space zones expand (Reg. 2021/664) [33].

For scaling across Europe, TALON operators should also consider pursuing a Light UAS Operator Certificate (LUC). With a LUC, operators could self-authorize recurring TALON missions across multiple jurisdictions, significantly reducing administrative friction. To support this, TALON deployments should include a comprehensive compliance package, covering operator registration, pilot training, geo-zone checks, GDPR compliance, insurance, and spectrum conformity (Reg. 2019/947) [34].

The European Union has established a sophisticated regulatory framework for drones, balancing safety, innovation, and cross-border harmonization. For industrial operators, the majority of relevant activities fall under the Specific category, where streamlined tools like PDRAs [35] and STS offer efficient paths to compliance. The upcoming expansion of U-space will add new requirements, but it will also enhance safety and scalability in dense operational environments.

For systems such as TALON, the message is clear: build compliance into the operational model from the outset. By aligning with Specific category rules, engineering for STS compatibility, preparing for U-space, and aiming for a LUC as deployments scale, operators can position themselves for sustainable growth in the European industrial drone market.

8 Conclusion

This deliverable confirms that the TALON project has established a solid foundation for both the dissemination of knowledge and the sustainable exploitation of its results. The project has demonstrated excellence in academic dissemination through numerous high-quality publications, conference contributions, and publicly available datasets. Its communication strategy, supported by an active digital presence and targeted events, has ensured strong stakeholder engagement across scientific, industrial, and policy domains.

The identification and development of Key Exploitable Results (KERs) form the cornerstone of TALON's exploitation strategy, with potential applications in manufacturing, robotics, logistics, and other industrial sectors. The patentability assessment has validated the novelty and protection potential of several innovations, thereby increasing confidence in their long-term market viability. Furthermore, the joint exploitation plan provides a roadmap for post-project sustainability, supported by individual partner commitments and external collaborations.

Standardisation and regulatory activities have strengthened TALON's position in the European innovation ecosystem. By aligning with the EU AI Act, the Green Deal, and other initiatives, the project ensures that its solutions meet requirements for transparency, interoperability, and sustainability. The regulatory analysis also confirms TALON's readiness for deployment in high-risk AI applications, ensuring compliance with forthcoming legislation.

From a business perspective, TALON is strategically positioned at the intersection of explainable AI, edge-to-cloud orchestration, and blockchain-based security – all of which have been reporting strong market growth trends. The market analysis and business models highlight significant opportunities for adoption in Industry 4.0 and Industry 5.0 environments, with demonstrable benefits in terms of operational efficiency, energy savings, and regulatory compliance.

Important insight come from the Analytical Hierarchy Process survey, which systematically prioritised the criteria relevant for TALON's technology roadmap. The results underscore that AI-fueled orchestration together with AI Explainability, Trustworthiness & Transparency, among TALON's four main pillars, feature as the two foremost drivers for users' adoption, confirming the necessity of transparent and auditable AI decisions in industrial contexts. Taking a more granular look into these two have indicated Intelligent Resource Management & Task Offloading, Fault Tolerance & Self-healing, and Reliability & Robustness of AI models, to be the three top-ranked individual technical features. These findings provide a strong evidence base for orientating technology roadmap strategies and focusing on those aspects of the solution, most likely to accelerate adoption.

In conclusion, Deliverable D6.2 has provided evidence that TALON is well prepared to transition from research and development into practical, sustainable, and standardised industrial deployment. By combining strong scientific outputs with regulatory alignment and clear business strategies, TALON contributes to Europe's digital sovereignty, competitiveness, and green transition.

9 Appendices

9.1 Appendix 1: TALON Patentability Survey

1. Title of the Component/Element Considered

a. Description of the Component/Element

- Text (required):

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - (1-5):
    - Optional text:

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
    - Text:

c. Claimed Invention

- Summarize the claimed invention's core technical features:
    - Text:

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art:

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
    - (1-5):
    - Optional text:

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
    - (1-5):
    - Optional text description:

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
    - (1-5):
    - Optional text explaining the effects:

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
    - (1-5):
    - Short text description:

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
    - (1-5):
    - Optional text about failed prior attempts:

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple

extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
   - (1-5):
   - Short text description:

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
   - (1-5):
   - Short text description:

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
   - (1-5):
   - Optional text:

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?

<u>4. Industrial Applicability</u>

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
   - (1-5):
   - Optional text:

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
   - (1-5):
   - Optional text (e.g., target industries):

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
   - (1-5):
   - Optional text:

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
   - Text:

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
   - (1-5):
   - Optional text:

9.2 Appendix 2: Responses to the Patentability Questionnaire

9.2.1 Zero-touch AI Orchestrator

<u>1. Title of the Component/Element Considered</u>

a. Description of the Component/Element

- Text (required):

The AI Edge-to-Cloud (E2C) Orchestrator is conceived as a unified platform that automates deployment, monitoring and optimisation of AI workloads across distributed industrial networks. By integrating zero-touch containerised deployments with a dynamic, policy-driven orchestration engine, the Orchestrator continuously adapts to fluctuations in compute, network and energy conditions without human intervention. Its real-time telemetry pipeline collects metrics on CPU, memory, bandwidth and power utilization, feeding multi-task LSTM-based forecasting models that drive placement, scaling and self-healing decisions. Built-in explainability and digital-twin visualisations ensure transparency of AI-driven actions, while distributed ledger technology secures model integrity and anomaly detection protects against malicious activities. Together, these capabilities reduce manual errors, minimise downtime and lower operational costs, enabling more resilient and energy-efficient industrial networks.

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - (1-5):3
    - Optional text: The AI E2C Orchestrator stands out from traditional network and cloud orchestration tools thanks to its real-time AI adjustments and self-learning capabilities, offering improved adaptability and energy efficiency.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
    - Text: Edge-to-cloud existing platforms handle provisioning, rule-based scaling and analytics. However, not all of them offer AI-driven adaptive modelling, real-time feedback loops, context-aware security or semantic transfer learning. The AI E2C Orchestrator includes those features.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
    - Text:

- AI-driven orchestration ensures real-time responsiveness and adaptability, reducing human error and operational inefficiencies.
- Self-learning mechanisms that adapt dynamically to evolving resource conditions and task requirements (e.g. latency, fault tolerance and security).
- Smart pod selection prioritises the best-matching services based on task demands to improve execution time and service quality.
- Secure DevOps integration with pre- and post-runtime protections enhances compliance and resilience in distributed environments.
- Energy-efficient execution via context-aware orchestration enables computation near sensors and minimises data movement to improve performance and sustainability.
- Cross-layer optimisation balances computational load across the cloud, the edge and end devices to improve scalability and robustness.

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art:

1. **AI-Driven Orchestration**
- Prior art relies on manual configuration and static rules.
- AI E2C Orchestrator uses autonomous, AI-based orchestration with self-learning capabilities.

2. **Context-Aware Service Selection**
- Traditional systems lack fine-grained contextual awareness.
- AI E2C dynamically selects the most suitable services or "pods" based on latency, energy, and security requirements.

3. **Real-Time Adaptation**
- Prior systems are reactive and not optimized for real-time adjustments.
- AI E2C enables proactive, real-time monitoring and decision-making.

4. ·**Optimised Workload Distribution**
- Existing tools typically focus on isolated environments (cloud or on-prem).
- AI E2C Orchestrator manages and optimises workloads across **edge, cloud, and on-premises** infrastructures.

5. **Energy and Resource Efficiency**
- General-purpose orchestrators do not optimise for energy usage or AI model performance.
- AI E2C employs intelligent scheduling and task placement algorithms to maximise energy and computational efficiency.

6. **Hybrid and Cross-Layer Orchestration**
- Cloud providers' tools are often siloed with limited integration.
- AI E2C offers **native support for hybrid deployments** and **cross-layer orchestration**

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
    - (1-5): 3
    - Optional text: There is a public disclosure with an academic paper published on early 2025.

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).

Programme - ICE 31th Edition

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- **(1-5): 3**
- **Rationale:** Selecting the edge-vs-cloud offload thresholds, adaptive resource-allocation policies, and dynamic compression settings for the E2C Orchestrator involves balancing latency, bandwidth, and device compute constraints. While each mechanism is known, tuning them together for seamless, real-time orchestration is a moderately non-obvious engineering decision.

b. Unexpected Effects

- **Unexpected-effects score (1 = none to 5 = many):** 3
- **Rationale:** The E2C Orchestrator yields moderately superior technical benefits beyond what skilled practitioners would anticipate. By dynamically offloading computation based on instantaneous device load and network conditions, it achieves lower end-to-end latency and more consistent QoS than static edge-or-cloud partitioning alone. This adaptive orchestration also enhances system resilience under fluctuating workloads and network variability—outcomes that exceed straightforward extrapolation from existing edge-cloud frameworks.

c. Improvement Significance

- **Significance score (1 = not at all to 5 = great):** 3
- **Short text description:** Improves end-to-end resource utilization and latency management in heterogeneous edge-to-cloud environments by simulating and adapting to network and device conditions that are difficult to predict or reproduce with static orchestration policies.

d. Long-Standing Problem

- **Problem-solving score (1 = not at all to 5 = greatly):** 3
- **Rationale:** The invention addresses the persistent challenge of delivering consistently low-latency, high-throughput computation in heterogeneous edge-to-cloud systems under fluctuating network and load conditions. Prior art relied on static partitioning or manual policy tuning, which repeatedly failed to adapt in real time to changing resource availability.

e. Straightforward Extrapolation

- **Score (1 = strongly agree to 5 = strongly disagree):** 2
- **Short text description:** Integrating real-time network monitoring, predictive load balancing, and context-aware compression into a unified E2C orchestration framework represents a non-trivial advancement over known edge-cloud partitioning schemes. While each component was individually understood, their seamless combination with dynamic policy adaptation under uncertainty would not have been an obvious straight-line extrapolation.

f. Feature Juxtaposition

- **Score (1 = strongly agree to 5 = strongly disagree):** 2
- **Short text description:** The E2C Orchestrator's components—real-time network telemetry, predictive load balancing, adaptive compression, and service-level-driven offloading—interact to produce a synergistic effect. Adaptive feedback loops reshape workload placement on-the-fly, and context-aware policies optimize end-to-end QoS; this is more than a simple aggregation of known features.

g. Secondary Considerations

- **Score (1 = minimal evidence to 5 = strong evidence):** 2
- **Optional text:** Early pilot deployments in manufacturing and telco labs have yielded positive internal feedback on resilience under network jitter. Although no public commercial roll-outs have occurred yet, several prospective customers have requested proofs-of-concept, and competitor roadmaps hint at similar orchestration efforts.

e. Required Knowledge

A "person skilled in the art" for this invention would be someone with:

- Expertise in distributed edge-to-cloud architectures (including Kubernetes, Docker, or similar container orchestration),
- Deep understanding of network performance characteristics (latency, throughput, jitter) and telemetry/monitoring tools,
- Experience with adaptive systems and policy-driven resource management,
- Familiarity with ML-based prediction techniques for workload forecasting,
- Competence in systems integration and QoS/SLA enforcement across heterogeneous devices and networks.

<u>4. Industrial Applicability</u>

a. Scientific Viability

- **Score (1 = not at all to 5 = certainly):** 3
- **Optional text:** All technical elements—from real-time telemetry and predictive load balancing to dynamic compression and offload execution—adhere to known principles of computer networking, distributed systems, and simulation theory.

b. Practical Use

- **Score (1 = no practical use to 5 = high potential):** 3

1. **Optional text:** The AI E2C Orchestrator shows great potential for practical use, particularly in sectors such as telecommunications, logistics and manufacturing, where large-scale AI systems operate across edge and cloud environments. These industries are increasingly challenged by issues relating to resource management, latency, and scalability. Although it is still in the early stages of adoption, the value of the Orchestrator can be demonstrated through pilot programmes with forward-thinking enterprises. Demonstrating measurable improvements in operational efficiency, responsiveness and energy use will be crucial in building confidence and driving broader market adoption**.**

c. Reproducibility

- **Score (1 = strongly disagree to 5 = strongly agree):** 4
- **Optional text:** Experts can reproduce and deploy the Orchestrator using common container orchestration tools (e.g. Kubernetes, Docker), standard telemetry and monitoring stacks (Prometheus, Grafana), and off-the-shelf edge/cloud hardware, following the disclosed dynamic policy rules without undue experimentation.

d. Solved Technical Problem

- The E2C Orchestrator addresses the long-standing issue of reliably managing compute workloads across heterogeneous edge and cloud resources under variable network conditions. In particular, it overcomes:
    - **Unpredictable latency and throughput** when static offload policies are used, leading to QoS violations.
    - **Inefficient resource utilization**, where some nodes are overloaded while others sit idle.
    - **Lack of real-time adaptability**, meaning systems cannot respond dynamically to spikes in demand, network jitter, or device failures.
      By continuously measuring telemetry, predicting load shifts, and dynamically reallocating tasks (including compression, caching, and execution location), it ensures both low end-to-end latency and high overall throughput without manual reconfiguration.

e. Transferability

- **Score (1 = not at all to 5 = greatly):** 4
- **Optional text:** The Orchestrator's architecture and policies are agnostic to specific application domains. It can be applied in any sector requiring distributed compute orchestration—such as IoT in smart manufacturing, real-time analytics for autonomous vehicles, remote healthcare diagnostics, and federated ML in finance—by plugging into standard container platforms, monitoring stacks, and network fabrics.

9.2.2 AR Maintenance Application

<u>1. Title of the Component/Element Considered</u>

a. Description of the Component/Element

- Text (required): AR Maintenance Application

An AI-enhanced augmented reality system designed for remote industrial maintenance, incorporating hands-free control through face and target tracking, adaptive AI-assisted motion control to reduce physical input bias, and dynamic interaction adaptation tailored to user needs.

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
  - (1-5):4
  - Optional text: The integration of advanced AI techniques with hands-free AR interaction for remote maintenance is significantly different from existing AR systems which typically rely on manual or controller-based input.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
  - Text: Existing AR remote maintenance systems provide visual overlays and some interaction but generally rely on handheld controllers or touch inputs and lack AI-driven adaptive controls or face/target tracking for hands-free use.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
  - Text:

A remote maintenance AR system that features:

- Hands-free user interface via face tracking and target tracking
- AI-assisted motion control to dynamically adjust input based on user behaviour
- Reduced physical input bias to improve interaction accuracy and ease
- Adaptive interaction models tailored to individual user needs and environmental context

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art:
  - Hands-free control through face and target tracking instead of physical controllers
  - AI-driven adaptive motion control to minimise user input errors and biases
  - Dynamic adaptation of interactions based on real-time user feedback
  - Integration of advanced CNN-based AI techniques specifically for remote maintenance tasks

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
  - (1-5): 3
  - Optional text: The combination of hands-free tracking and AI-adaptive controls for AR maintenance has been publicly disclosed prior to the invention through scientific contributions.

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).

Vuforia AR remote support system (2022) — conventional AR overlays without AI-adaptive input

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
  - (1-5): 2
  - Optional text description: While AI and AR are known, applying CNN-based adaptive motion control specifically to hands-free AR remote maintenance is not an obvious step.

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
    - (1-5): 3
    - Optional text explaining the effects: The system significantly improves usability and accuracy in AR remote maintenance, enabling users to operate complex machinery remotely with fewer errors.

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
    - (1-5): 3
    - Short text description: Offers a major enhancement in user interaction and system adaptability, directly improving remote maintenance efficiency.

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
    - (1-5): 3
    - Optional text about failed prior attempts: Prior solutions have struggled with cumbersome manual inputs and user fatigue, limiting remote maintenance effectiveness.

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
    - (1-5):2
    - Short text description: Although elements like face tracking and AR exist, combining them with AI adaptive control in this context is not a trivial extension.

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
    - (1-5):3
    - Short text description: The invention combines features that interact functionally to improve control accuracy and user experience beyond simple aggregation.

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
    - (1-5):3
    - Optional text: Commercial uptake is at an early stage, with positive internal evaluation feedback.

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?

A person skilled in the art would have expertise in AR systems, AI (especially CNNs), computer vision, human-computer interaction, and industrial maintenance workflows.

4. Industrial Applicability

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
    - (1-5):4
    - Optional text: The invention relies on proven principles in computer vision, AI, and AR technology.

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
    - (1-5):4

- Optional text (e.g., target industries): Highly applicable in industrial sectors including manufacturing, aerospace, energy, and infrastructure maintenance.

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
  - (1-5):3
  - Optional text: Can be implemented using existing AR hardware platforms combined with custom AI software components.

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
  - Text: It solves the problem of inefficient and error-prone manual inputs in AR remote maintenance by enabling intuitive, hands-free, AI-adaptive control.

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
  - (1-5):3
  - Optional text: The approach is adaptable to various industries requiring remote technical assistance or maintenance, improving safety and operational efficiency.

9.2.3 VR Training Application

<u>1. Title of the Component/Element Considered</u>

a. Description of the Component/Element

- Text (required): VR Training Application

A high-fidelity, context-aware virtual reality training system for industrial manufacturing, offering adaptive feedback, dynamic scenario adjustment, and immersive, gesture-based user interaction to enhance learning outcomes in simulated environments.

<u>2. Novelty</u>

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
  - (1-5):3
  - Optional text: The system introduces multiple technical innovations not found in existing VR training solutions.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
  - Text: Existing VR-based training platforms such as Siemens' COMOS Walkinside or EON Reality's Creator AVR offer immersive content but lack adaptive scenario modelling, real-time user feedback integration, and context-aware assessment mechanisms designed for complex industrial manufacturing environments.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
  - Text:

An interactive, industrial-focused VR training platform that features:

  - Context-aware, dynamic training scenarios
  - Gesture-based navigation with low learning curve interfaces
  - Real-time adaptive feedback based on user interaction
  - A novel scoring and motivation evaluation mechanism
  - High-fidelity simulation to mimic hazardous or rare manufacturing environments

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art:

Adaptive scenario modification based on user behaviour
Real-time feedback mechanisms integrated into training flow
Gesture-based navigation with natural input for accessibility
Scoring mechanism to assess learning progress and motivation
High-fidelity, hazard-simulation environments uncommon in standard platforms

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
    - (1-5):4
    - Optional text: There is a public disclosure of the some technical combination of features before potential filing in a form of scientific contributions.

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).

Siemens COMOS Walkinside – online product documentation and marketing material (2021)
EON Reality Creator AVR – website and product videos (2022)

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
    - (1-5): 3
    - Optional text description: Combining adaptive scenario generation, real-time feedback, and natural gesture control is an obvious progression in VR training design.

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
    - (1-5): 3
    - Optional text explaining the effects: Enables significantly improved user engagement and memory retention in rare or hazardous task simulations.

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
    - (1-5): 3
    - Short text description: Improves safety training and task comprehension significantly by simulating difficult-to-replicate conditions.

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
    - (1-5): 3
    - Optional text about failed prior attempts: Traditional VR platforms do not replicate rare industrial hazards or dynamically adapt to user progress.

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
    - (1-5): 3
    - Short text description: The integration of scoring, adaptive simulation, and intuitive gesture input is a non-trivial advancement from known solutions.

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
    - (1-5): 3
    - Short text description: Features work together, synergistically, adaptive feedback alters simulation flow, gesture input improves usability, and scoring maintains engagement

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?

- (1-5): 3
- Optional text: Commercial adoption pending; initial internal feedback indicates positive industry reception.

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?

A skilled person in the art would have expertise in VR development (Unity, Unreal Engine), knowledge of industrial training workflows, human-computer interaction, and adaptive learning systems. They would also understand 3D modelling, physics simulation, and user experience design.

4. Industrial Applicability

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
    - (1-5): 4
    - Optional text: All technical elements are consistent with physics, simulation theory, and digital interaction principles.

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
    - (1-5): 4
    - Optional text (e.g., target industries): Highly applicable to manufacturing, defence, aerospace, energy, and logistics sectors for safe and effective personnel training.

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
    - (1-5): 4
    - Optional text: Can be reproduced using standard VR frameworks and SDKs with appropriate hardware.

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
    - Text: It solves the problem of simulating rare or hazardous industrial environments for training, enabling safe and repeatable exposure to complex scenarios.

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
    - (1-5): 4
    - Optional text: The invention is applicable across various industries requiring immersive training, such as healthcare, aviation, nuclear, and logistics.

9.2.4 Federated Learning Module

1. Federated Learning Module

a. Description of the Component/Element

Text (required): The federated learning module implements a distributed machine learning framework that enables model training across multiple edge devices while keeping data localized. The system utilizes established aggregation techniques including FedAvg, FedProx, FedAdam, FedAdagrad, and FedYogi to combine local model updates into a global model. The implementation leverages the Flower framework for communication between nodes and a central aggregation server. The module has been demonstrated using YOLOv8 for object detection tasks.

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - (1-5): 1

  - Optional text: The module implements standard federated learning algorithms that have been extensively documented in literature. The use of Flower framework with conventional aggregation methods represents a straightforward application of existing technologies without substantial modification.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
  - Text: McMahan et al. [36] introduced FedAvg in 2017, establishing the foundational approach used in this module. Li et al. [37] presented FedProx for handling heterogeneous networks. Reddi et al. [38] comprehensively analyzed adaptive federated optimization methods including FedAdam, FedAdagrad, and FedYogi. The Flower framework [39] provides the implementation infrastructure for federated learning research.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
  - Text: A federated learning system using standard aggregation algorithms (FedAvg, FedProx, FedAdam, FedAdagrad, FedYogi) implemented via the Flower framework for industrial applications with object detection.

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art: The primary difference is the specific application domain (industrial safety monitoring) rather than any technical innovation. The combination of existing FL algorithms with YOLOv8 models represents a domain-specific implementation rather than a methodological advance.

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
  - (1-5): 5
  - Optional text: All core algorithms are published in peer-reviewed venues. The Flower framework is open-source. YOLOv8 architectures are publicly available. The specific implementation combines only publicly available components.

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).
- B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artif. Intell. Statist., Fort Lauderdale, FL, USA, 2017, pp. 1273-1282.

<u>3. Inventive Step (Non-Obviousness)</u>

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
  - (1-5): 2
  - Optional text description: Default parameters from the Flower framework and standard hyperparameters for YOLOv8 models were used without major modification.

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
  - (1-5): 1
  - Optional text explaining the effects: Performance metrics align with expected results from applying standard FL algorithms to conventional ML models.

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
  - (1-5): 2

    - Short text description: The system achieves expected performance using established methods without notable improvements over existing FL implementations.

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
    - (1-5): 1
    - Optional text about failed prior attempts:

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
    - (1-5): 1
    - Short text description: The implementation directly applies known FL algorithms using an existing framework to standard ML models.

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
    - (1-5): 2
    - Short text description: The combination of Flower framework with specific aggregation algorithms for industrial use cases represents a predictable integration of known components.

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
    - (1-5): 1
    - Optional text:

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?
- A person skilled in the art would need basic understanding of federated learning concepts, familiarity with Python programming, and knowledge of the Flower framework documentation.

<u>4. Industrial Applicability</u>

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
    - (1-5): 5
    - Optional text:

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
    - (1-5): 5
    - Optional text (e.g., target industries): Target industries include construction, manufacturing, logistics, and energy, where worker safety is critical.

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
    - (1-5): 5
    - Optional text: The solution can be easily reproduced, given that the subject matters know the model hyperparameters and the datasets used.

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.

- Text: The module addresses the need for privacy-preserving distributed learning in industrial environments where data cannot be centralized due to regulatory or technical constraints.

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
  - (1-5): 5
  - Optional text: The invention can be applied across various industries such as construction, manufacturing, logistics, and energy, where worker safety and hazard detection are priorities.

9.2.5 XAI Framework

1. Explainable AI (XAI) Framework

a. Description of the Component/Element

Text (required): TALON's XAI framework delivers interpretable insights into data quality and model predictions for both images and time series using a combination of analytical methods, XAI techniques, and advanced visualizations. Explanations are organized into four levels:

Level 1

- Images: Visualizations showing the percentage of images with issues (e.g., corruption, format or shape inconsistencies).
- Time Series: Plots highlighting null values per feature.

Level 2

- Images: Bar charts of class distributions and sample image previews.
- Time Series: Outlier detection using Local Outlier Factor (LOF).

Level 3

- Images: Grad-CAM visualizations for object detection.
- Time Series: Feature importance analysis via SHAP.

Level 4

- Images:
  - Fidelity: Measures model reliance on key image regions by masking and observing performance drop.
  - Consistency: Assesses robustness across varying conditions.
- Time Series: Advanced SHAP-based analysis including feature importance, dependencies, interactions, and decision trees.

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
  - (1-5): 3
  - Optional text: Although the methods used have been previously explored in the XAI literature for both image and time series modalities, the integration of data diagnostics, visualization, model-specific XAI, and advanced robustness testing across these modalities represents a novel contribution.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
  - Text: The closest prior art in XAI involves converting multivariate time series into image-like representations to apply computer vision XAI methods such as Grad-CAM and LIME. Notable examples include Tronchin et al. and Makridis et al., who use CNN-based models on time-series plots to highlight important temporal patterns influencing predictions. Key features of these methods include multimodal fusion and model distillation to improve interpretability, as seen in frameworks like Time-Series-XAI, which provides attribution maps and counterfactuals for time-series classifiers. However, these solutions focus mainly on single-modality explainability and lack integration with image-specific XAI techniques or a structured, multi-level workflow covering data quality, visualization, and robustness.

c. Claimed Invention

- Summarize the claimed invention's core technical features:

Text: The core technical features include:

1. Multi-Modal Input Support: Handles both image and time series data, adapting explanation techniques accordingly.
2. Four-Level Explanation Framework:
    - Level 1 – Data Quality Visualization: Detects and visualizes input data issues (e.g., image corruption, missing values in time series).
    - Level 2 – Distribution & Anomaly Detection: Displays class distributions and representative samples for images; uses LOF for detecting outliers in time series.
    - Level 3 – Model Insight: Applies Grad-CAM for image interpretation and SHAP for feature importance in time series predictions.
    - Level 4 – Robustness & Fidelity: Measures model sensitivity to input regions (images) and analyzes feature dependencies and interactions (time series) using SHAP.
3. **Advanced Visual Analytics:** Integrates visualization techniques across all levels to improve interpretability and user understanding.

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art: The proposed framework distinguishes itself from the closest prior art by natively supporting both image and time series data, applying domain-specific XAI techniques without transforming data types—using Grad-CAM for images and SHAP for time series. Unlike prior solutions focused on single-modality explanations, it introduces a structured four-level framework covering data quality, statistical overview, model interpretability, and robustness. It uniquely integrates data integrity diagnostics and evaluates model fidelity and consistency through perturbation analysis. This comprehensive, visualization-driven approach offers deeper, more actionable insights than the static, attribution-focused methods seen in existing work.

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
    - (1-5): 4
    - Optional text: All methods employed in the framework are well-established and have been previously examined in the literature. The datasets used for experimentation are entirely open-source.

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).
- https://www.diva-portal.org/smash/get/diva2:1845256/FULLTEXT01.pdf
- https://arxiv.org/pdf/2311.17110

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
    - (1-5): 2
    - Optional text description: All major hyperparameters can be easily adjusted by the user according to their needs, ensuring that no significant assumptions were made that could be considered non-obvious.

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
    - (1-5): 2
    - Optional text explaining the effects:

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?

  - (1-5): 2
  - Short text description: The most major improvement is in terms of combining these already existing methods into a single framework covering two different modalities. However, the obtained results were as expected.

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
  - (1-5): 2
  - Optional text about failed prior attempts:

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
  - (1-5): 1
  - Short text description: The techniques used in this work have been previously explored in the literature but have not been compiled into a single framework covering both image and time series modalities.

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
  - (1-5): 3
  - Short text description: The synergistic effect resulting from the combination of data analytics, visualization, and XAI methods for images and time series cannot be considered non-obvious, despite not having been previously combined into a single framework.

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
  - (1-5): 2
  - Optional text:

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?
- The person needs to have a basic understanding of XAI techniques as well as how different hyperparameters of these methods affect their performance.

<u>4. Industrial Applicability</u>

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
  - (1-5): 5
  - Optional text:

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
  - (1-5): 4
  - Optional text (e.g., target industries): The proposed framework has strong practical potential, particularly in industries like healthcare, finance, autonomous vehicles, and industrial automation.

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
  - (1-5): 4
  - Optional text: The solution can be easily reproduced, given that the subject matters know the methods' hyperparameters and the datasets used.

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
    - Text: The proposed framework addresses the technical problem of integrating multiple explanatory techniques, such as data diagnostics, visualization, model-specific XAI, and robustness testing, across different modalities (e.g., images and time series). This combination enhances the interpretability and reliability of ML models, providing deeper insights and more robust evaluations.

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
    - (1-5): 4
    - Optional text: The proposed framework has strong potential for transfer across various industries, particularly those relying on ML models for decision-making, such as healthcare, finance, and autonomous systems. Its ability to enhance model interpretability and robustness makes it highly adaptable, though domain-specific adjustments may be required for optimal integration.

9.2.6 Blockchain Mechanism

<u>1. Blockchain Mechanism</u>

a. Description of the Component/Element

- The fusion of federated learning (FL) with blockchain enables IoT applications securely and in a decentralised manner train the models, even across devices in a distributed fashion, without compromising data privacy, security and trust. This is a particularly interesting concept especially within TALON. Each IoT device (or cluster) autonomously trains a local model on their data and shares model updates, not raw data, on a decentralised in this setting. Blockchain provides security to this federated system through an immutable, tamper-resistant record of the model updates, thereby paving the way for the discovery of potential malicious activities or faulty contributions by any peer participant. SIDROCO employs decentralised applications through blockchain technology to ensure full, end-to-end, customised, perpetual security and privacy in the E2C AI architecture. Using the blockchain mechanisms enables TALON to operate high-throughput and low-latency dApps with strong security via ZT (zero trust) deployment for edge AI. That will empower TALON to scale, increase flexibility, harden, and boost other critical features available during deploying AI solutions across the Edge AI networks in support of the conceptual notion of distributed/federated intelligence through decentralised networks.

<u>2. Novelty</u>

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - (1-5): 2
    - Value Proposition: One major challenge in FL is ensuring the integrity of model updates, as malicious participants could tamper with data or introduce poisoned models. Blockchain tackles this by recording all transactions—such as model updates in TALON case—on a tamper-proof, decentralised ledger, ensuring transparency and traceability. Information is stored on the IPFS system, while the hash is stored on the ledger itself. This helps with the scalability and lightweight nature of the proposed solution. Another challenge is the lack of trust among participants, particularly when FL involves multiple organisations. Blockchain boosts trust through its consensus mechanisms, which validate and agree on the ML algorithm updates without relying on a central authority. By decentralising the coordination process, blockchain removes the risk of a single point of failure, enhancing the overall robustness and reliability of FL systems.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):

Text: FedSyn is a framework by J.P. Morgan's Kinexys team to generate synthetic data for training machine learning models, while preserving privacy. FedSyn combines synthetic data generation with privacy-preserving Federated Learning:

1) It attempts to address data scarcity, data privacy, data bias and augments data-centric AI.

2) Differential privacy on model parameters in Federated Learning
3) Generates synthetic data even from imbalanced distribution

The Kinexys team, who've developed Liink, J.P. Morgan's blockchain network, established that FedSyn can delegate secured aggregation to a consortium-trusted entity in a permissioned blockchain network, such as LiinK – another step forward in the firm's work to provide network participants with an improved experience through innovative technologies and collaboration.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
  - Text: Decentralised architecture to enhanve the privacy and security of federated learning capabilities and communication of AI models, while simultaneously not sacrificing energy related metrics (latency, throuput, CPU utilisaiton etc).

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art: Main difference is that FedSyn is mainly used for synthetic data generation using FL over a different permissioned blockchain called liink.

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
  - (1-5): 2
  - Optional text:

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).
- https://www.jpmorgan.com/technology/news/federated-learning-meets-blockchain

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
  - (1-5): 2
  - Optional text description:

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
  - (1-5): 2
  - Optional text explaining the effects:

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
  - (1-5): 2
  - Short text description:

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
  - (1-5): 3
  - Optional text about failed prior attempts:

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
  - (1-5): 3
  - Short text description:

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
    - (1-5): 3
    - Short text description:

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
    - (1-5): 2
    - Optional text:

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?

The person needs to be familiar with ML development in federated architectures, as well as the technical specifications of permissioned blockchain mechanisms (such as Hyperledger Fabric).

<u>4. Industrial Applicability</u>

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
    - (1-5): 5
    - Optional text:

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
    - (1-5): 2
    - Optional text (e.g., target industries):

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
    - (1-5): 4
    - Optional text:

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
    - Text: It provides a secure and private environment for federated AI processes, preserving AI trustworthiness (immutability of AI models), transparency (logs transactions) and privacy (only organisations within the ledger have access).

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
    - (1-5): 4
    - Optional text:

9.2.7 Image Anonymization Module

<u>1. Image Anonymization Module</u>

a. Description of the Component/Element

The Image Anonymization Module (IAM) is one of the Access & Security layer's mechanisms that are designed to protect the confidentiality of personal data handled by Talon project.
It provides a comprehensive solution for obfuscating faces in image and video files. It comprises an HTTP REST API server for file ingestion, an anonymizer module for the actual anonymization process, an output module for deploying the anonymized files and a database (MongoDB) as a secure repository for the processed original files.
It also presents the Auditor, whose scope is to act as a cron-job and to delete, every day, the original image/video files present in DB, if any, whose dates are older than a configured retention-period.
In the figure below, the architecture of the Image Anonymization Module is shown, reporting also the open-source libraries used for realization.

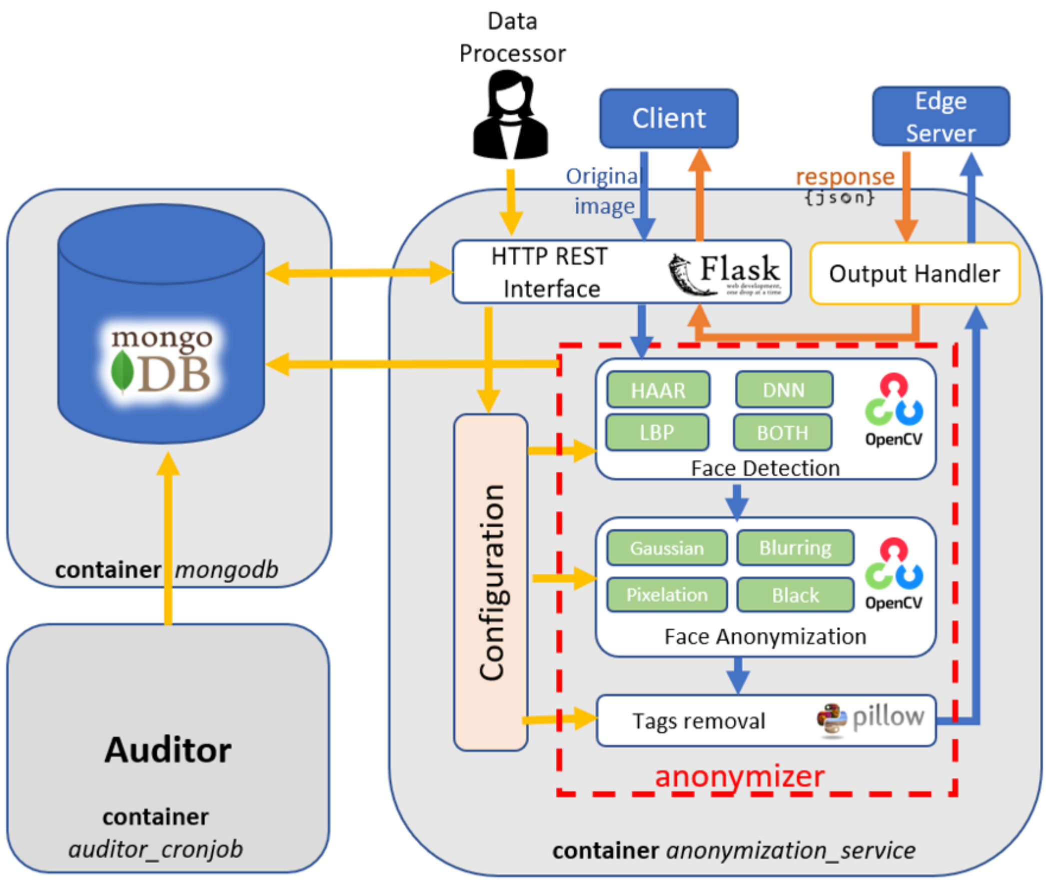

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - (1-5): 3
    - Optional text: Currently it has not been identified a tool (neither commercial nor open-source) that permits the pseudonymization of faces in images with the possibility to eliminate original images older than a configured retention period, as requested by GDPR

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
    - Text: **MD.ai** is a cloud-based medical imaging platform focused on **AI-driven annotation, collaboration, and de-identification**.

      Technical Field
    - **Domain**: Medical imaging, AI, and data annotation.
    - **Use Cases**: Radiology AI development, collaborative image labeling, privacy-preserving dataset preparation.
    - **Target Users**: AI researchers, radiologists, hospitals, and medical AI startups.

Problem Solved

- Simplifies the creation of high-quality annotated datasets for training and validating AI models in healthcare.
- Solves the challenge of sharing medical images securely by offering **automated de-identification and pseudonymization**.
- Facilitates collaboration between clinicians and data scientists with cloud-based access and versioned data workflows.

Key Features

- **Pixel-level annotation** tools (bounding boxes, segmentation).

- **HIPAA/GDPR-compliant** data de-identification, including optional manual review.
- **Auto-deidentification of pixel text** using the mdai:pixel-deid model.
- **Role-based access control** and audit trails.
- **Flexible data export** (JSON, DICOM, etc.) and AI model integration.
- **Retention controls** for compliance (configurable retention periods).

Note, it is not clear if MD.ai is able to obfuscate faces in images.
Other similar tool (always for medical scope) is Philips IntelliSpace.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
  - Text: A method to pseudonymize faces in images granting the compliance to GDPR introducing an original image retention period.

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art:
  - Main difference with the prior art is that in the IAM module the original images are stored in a secure database, A cronjob cleanup original images older than a configured retention-period. The prior-art uses cryptography.
  So, in both cases there is a "critical" point where to preserve Original images and cryptographic keys. In case of prior-art (use of cryptographic key) it could be possible to suffer of a brute-force attack that try to recover original images from anonymized one.

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
  - (1-5): 3
  - Optional text:

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).
  www.md.ai
  While specific publication titles, dates, or conference presentation details are not available, MD.ai's tools are actively used in medical imaging research and AI development. For more detailed information or inquiries about potential publications or presentations, you may consider reaching out directly to MD.ai through their contact page.

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
  - (1-5): 2
  - Optional text description: If you take in mind the GDPR, the implementation become obvious.

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
  - (1-5): 2
  - Optional text explaining the effects:

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
  - (1-5): 3
  - Short text description: Improved security of original images

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
  - (1-5): 2
  - Optional text about failed prior attempts:

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
    - (1-5): 1
    - Short text description:

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
    - (1-5): 1
    - Short text description: it is a sum of existing technologies and concepts to collaborate together to satisfy the GDPR suggestion.

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
    - (1-5): 2
    - Optional text:

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?

The person needs to knows the GDPR needs and the techniques for data storing and cronjobs in a cloud system.

<u>4. Industrial Applicability</u>

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
    - (1-5): 5
    - Optional text:

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
    - (1-5): 5
    - Optional text (e.g., target industries): in all cases where images containing persons need to be handled/shared.

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
    - (1-5): 4
    - Optional text:

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
    - Text: It solves a legal problem (privacy granting).

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
    - (1-5): 5
    - Optional text:

9.2.8 Smart Pricing Simulator

<u>1. Smart Pricing Simulator (SPS)</u>

a. Description of the Component/Element

- The Smart Pricing Simulator (SPS) aims to support in the efficient allocation of TALON's computational resources in four (4) distinct ways: (i) By ensuring that the utilization of TALON's edge computing network charges reflect the computational requirements of the user, (ii) by allowing – through a pay-as-you-go scheme – participants to retain full authority of their costs upfront, encouraging them to utilize TALON's energy-efficient edge computing network, (iii) by facilitating the use of non-commercial devices for edge computing tasks, through effective workload allocation, and (iv) by establishing an incentivization scheme that allows users to opt-in to more energy efficient choices that reinforce the project's overall goal for a greener and more efficient edge-computing network.

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - 3 out of 5

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
    - Amazon Web Services (AWS), especially services like AWS Wavelength and EKS Edge, which support edge deployment and pay-as-you-go pricing.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
    - TALON's SPS is essentially an **AI-powered pricing engine** adjusted for **energy-efficient, non-commercial device-inclusive edge networks**, built around:
        - A **Stackelberg Game formulation** for pricing optimization.
        - **Real-time utility maximization** using Sequential Least Squares Programming (SLSQP).
        - A dual discount mechanism: **dynamic demand-based** and **bundle-based**.
        - Integration with real-time **Prometheus-monitored metrics** and **Holt-Winters forecasts** for real-time demand prediction.

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art:
    - Explicit integration of energy metrics and green AI objectives into the pricing logic.
    - Focused on **non-commercial device participation** and **demand-side incentivization**.
    - Uses a **Stackelberg Game framework** for leader-follower interaction between ESP and users.
    - Adaptive based on **local usage data**, not just generalized cloud workloads.

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
    - 4
    - While most elements (e.g., pay-as-you-go, bundling) are known, the combination of real-time forecasting, game-theoretic modeling, and energy-based incentives are not fully disclosed together in prior art.

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).
    - AWS Wavelength (2019–2020): Edge computing + telecom integration.
    - AWS EKS Edge (2024): Pay-as-you-go Kubernetes at the edge.
    - AWS "Security at the Edge" whitepaper (2021).
    - Amazon's Reserved and Savings Plans: Pricing schemes with commitment-based discounts.

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
    - 2
    - The choice of parameters (vCPUs, RAM, elasticity, demand levels) is logical and likely familiar to a skilled person.

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
    - 3/5
    - By combining Stackelberg Game Theory, Holt-Winters forecasting, and real-time discounting, SPS achieves:
        - Higher EE (>20%)
        - Non-commercial participation (>70%)
        - Forecast-aware pricing with low CVRMSE (~11%)

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
    - 3/5
    - Provides a real-time, demand-responsive pricing engine for distributed, green edge AI infrastructures—a niche not yet fully served by existing hyperscalers.

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
    - 3/5
    - Addresses the challenge of incentivizing real-time participation of heterogeneous non-commercial edge devices under EE constraints. However this is not a challenge of unprecedented difficulty.

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
    - 2/5
    - A person skilled in the art may imagine dynamic pricing or energy-aware schemes.

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
    - 3/5
    - The combination of forecasting, optimization, pricing, and incentivization yields synergistic effects (e.g., pricing driven by predicted demand + elasticity = dynamic market balance).

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
    - 3
    - As a TRL-5 component, there's no strong evidence of commercial success yet, but early integration into TALON and industry interest suggests future traction.

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?
    - A **skilled person** would be a Cloud/Edge Systems Engineer or AI expert with understanding of:
        - Edge resource allocation and pricing.
        - Game theory (Stackelberg).
        - Optimization techniques (e.g., SLSQP).
        - Time-series forecasting models (e.g., Holt-Winters, Random Forest).

4. Industrial Applicability

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
    - 5/5

        - Fully consistent with known principles of optimization, forecasting, and edge computing.

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
    - 5/5

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
    - 4/5
    - Fully documented in deliverable D3.5; reproducible with relevant metrics. It requires experimentation for AI model training.

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
    - Balances user affordability, EE, and ESP profitability in edge AI networks with non-commercial participation and real-time demand response.

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
    - 5/5
    - Can be adapted to:
        - Smart grids (e.g., encourage EV charging at off-peak hours using bundling and real-time price visibility)
        - Industry 5.0 predictive maintenance (e.g., sensors delay non-critical data transmission until energy-efficient or low-traffic periods, reducing costs)
        - Federated AI training in sensor-rich environments (e.g, rewarding participation during times when battery and connectivity constraints are minimal).

9.2.9 Synthetic CNC Data Generator

1. Synthetic CNC Data Generator

a. Description of the Component/Element

- The synthetic data generation module developed by CERTH is a deep learning-based tool that uses Time-series Generative Adversarial Networks (TimeGAN) to produce realistic synthetic CNC sensor data. This module is designed to address the need for high-quality data in industrial AI applications, particularly in scenarios where real sensor data is limited, costly to acquire, or sensitive. TimeGAN uniquely combines adversarial learning with sequence modeling to generate time-series data that preserves both temporal dynamics and statistical properties of real-world CNC machine data. Unlike traditional GANs, TimeGAN incorporates recurrent neural networks (RNNs) into both the generator and discriminator, ensuring that generated sequences maintain realistic time-dependent behavior, which is essential for simulating manufacturing processes. The model is trained on real CNC machine sensor data, enabling it to learn complex, multivariate temporal relationships across sensors such as temperature and current. The resulting synthetic data is intended to be used for model training, testing, and validation in downstream tasks such as anomaly detection, predictive maintenance, and system behavior simulation, without compromising data privacy or requiring extensive real-world experimentation. This component is integrated into the TALON ecosystem via a containerized REST API, with configurable parameters such as sequence length, noise level, and sampling resolution, allowing flexible adaptation to different use cases and industrial setups.

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - (1-5): 2

  - Optional text: The developed tool applies an existing deep learning architecture for synthetic time series data generation in the specific context of CNC manufacturing sensor data. While TimeGAN is not a novel architecture, its integration into a containerized, configurable, and REST-accessible tool tailored for industrial datasets brings practical value to real-world deployments in manufacturing. The contribution lies in the adaptation, engineering, and application of existing research rather than in foundational algorithmic innovation.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
  - Text: The closest prior art is the original TimeGAN work by Yoon et al. (2019), which introduced a hybrid architecture combining adversarial training with autoregressive modeling and embedding networks to generate realistic synthetic time series data. This model was applied to various domains, including finance and healthcare. Other similar works also include applications of GANs to time series data in industrial settings, but often without dedicated containerized infrastructure or REST APIs. Additionally, tools such as Gretel.ai and SynthCity provide general-purpose synthetic time-series data generation for machine learning applications.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
  - Text: An IT solution that packages TimeGAN for the generation of synthetic CNC machine sensor data, providing configurability (e.g., sequence length, noise), RESTful API access, and Docker-based deployment, enabling easy integration into industrial edge-cloud systems for tasks such as predictive maintenance, anomaly detection, and simulation.

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art:
  - Domain-specific training (CNC sensor datasets from actual machinery)
  - The system is productionized with Docker containers, REST APIs.
  - Provides runtime configurability for sequence parameters and data properties.
  - Designed for deployment within the TALON cloud-edge ecosystem, targeting real-time industrial use cases.

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
  - (1-5): 2
  - Optional text:

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).

Yoon, J., Jarrett, D., & Van der Schaar, M. (2019). **"Time-series Generative Adversarial Networks."** NeurIPS 2019.
https://proceedings.neurips.cc/paper_files/paper/2019/file/c9efe5f26cd17ba6216bbe2a7d26d490-Paper.pdf
GitHub (Open-source implementations):
https://github.com/jsyoon0823/TimeGAN

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
  - (1-5): 2
  - Optional text description:

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?

- (1-5): 1
- Optional text explaining the effects:

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
  - (1-5): 1
  - Short text description: The technical improvement lies in domain-specific application and deployment engineering rather than core algorithmic novelty. The improvements in usability, API accessibility, and Dockerization are practical but expected.

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
  - (1-5): 3
  - Optional text about failed prior attempts:

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
  - (1-5): 3
  - Short text description: Applying TimeGAN to industrial sensor data and wrapping it in a deployable API service is a logical extension, but the combination of tools for production-grade use requires effort that may not be immediately trivial to execute well.

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
  - (1-5): 2
  - Short text description: The tool integrates known elements (TimeGAN, Docker, REST APIs) without fundamentally new functional interaction, but the sum is practically useful in an applied industrial context.

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
  - (1-5): 2
  - Optional text:

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?

A person skilled in the art would be a machine learning engineer or data scientist familiar with generative models, specifically GAN-based architectures for time series data. They would also need experience with industrial sensor systems and practical ML deployment tools such as Docker, Flask/FastAPI, and RESTful services.

<u>4. Industrial Applicability</u>

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
  - (1-5): 5
  - Optional text:

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
  - (1-5): 2
  - Optional text (e.g., target industries):

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
    - (1-5): 2
    - Optional text:

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
    - Text: Yes. It addresses the challenge of limited labeled CNC sensor data by generating realistic synthetic time-series sequences, enabling downstream tasks like model training, testing, and anomaly detection in data-scarce environments.

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
    - (1-5): 3
    - Optional text:

9.2.10 Few-shot Object Detection

1. Few-Shot Object Detection

a. Description of the Component/Element

Text (required): The few-shot object detection module is designed to accurately detect objects commonly found in real-world industrial environments, particularly those that directly impact worker safety. This includes the detection of Personal Protective Equipment (PPE), such as helmets and vests, as well as the absence of such equipment. It also targets potential hazards like fires that could endanger workers.

The primary goal of this module is to enable reliable detection of these critical objects, even when only a limited number of training samples are available. To accomplish this, the system uses a lightweight, pretrained object detector (YOLOv8n), which is specifically finetuned using the few available training samples.

This approach allows the module to maintain strong performance despite limited data, while also minimizing computational and energy requirements, making it a practical solution for real-world industrial applications where resources may be constrained.

2. Novelty

a. Degree of Difference Assessment

- On a scale from 1 (very similar) to 5 (very different), evaluate the difference between the invention considered and its closest prior art?
    - (1-5): 2
    - Optional text: The use of models from the YOLO family for detecting PPE and hazardous objects has been widely explored in recent literature. Additionally, over the years, various finetuning approaches have been proposed, including partial model finetuning, such as adjusting only the detection modules. However, the application of these finetuning techniques to YOLO models specifically for few-shot object detection of PPE and hazards has been limited.

b. Closest Prior Art Identification

- Describe the closest prior art (from the aspect of technical field, problem solved, and key features):
    - Text: In the work by Wang and El-Gohary a few-shot object detection pipeline is proposed for identifying personal protective equipment (PPE), hazards, and industrial equipment at construction sites, with the overarching goal of enhancing worker safety. Additionally, the method is designed to detect these objects along with their associated, intricate attributes to capture essential site condition information and assess compliance with relevant safety regulations.

The proposed pipeline adopts a comprehensive, end-to-end strategy that involves:
(1) creating a dataset by collecting and labeling images related to fall incidents at construction sites,
(2) implementing a few-shot object detection technique to handle the uneven occurrence of different object types,
(3) using an attribute recognition component to generate semantic descriptions of the identified objects, which supports the evaluation of safety compliance on-site.

c. Claimed Invention

- Summarize the claimed invention's core technical features:
  - Text: A few-shot object detection approach based on partial finetuning is proposed, which updates only the detection modules of the object detector. This method achieves strong performance with limited training data while also reducing energy consumption, making it well-suited for deployment in real-world industrial environments.

d. Differences from Prior Art

- List technical differences between the claimed invention and the closest prior art: The primary distinction from prior work is the proposed method's strong emphasis on energy efficiency. Specifically, it utilizes a lightweight object detector in which only a small fraction of the model's parameters are finetuned, in contrast to the much larger and more computationally demanding SwAV model. Furthermore, the proposed approach can achieve high performance using just a single training image per object, whereas previous methods require at least 50 samples per object.

e. Public Disclosure

- On a scale from 1 (strongly disagree) to 5 (strongly agree) evaluate whether each and every element or step of the claimed invention was explicitly or inherently disclosed, individually or in combination, in any form of public knowledge (e.g., publications, public use, sales, presentations) before the filing date of a potential patent application?
  - (1-5): 4
  - Optional text: The main utilized model is YOLOv8n which is an open-source model. Additionally, all datasets employed in this work are also open-source, except for one, which will be released soon as part of an accepted ICE25 paper. Lastly, all relevant results have been published in conference papers, available here and here.

f. Closest Public Prior Art

- Please specify where and when the closest prior art was made available to the public (e.g., title of publication, date, website, conference presentation date).
- https://www.sciencedirect.com/science/article/pii/S0926580524002759

3. Inventive Step (Non-Obviousness)

a. Parameter Selection

- On a scale from 1 (very obvious) to 5 (not at all obvious), how obvious would you evaluate the choice of particular parameters you have made have been to a person skilled in the art at the time of the invention, given the limited range of possibilities?
  - (1-5): 2
  - Optional text description: Partial finetuning is a well-established technique in the context of few-shot and transfer learning. Additionally, in this work, the default hyperparameter values provided by the Ultralytics framework were used for the YOLOv8 model.

b. Unexpected Effects

- On a scale from to 1 (none) to 5 (many) evaluate whether the invention produced technical effects that are different from or significantly superior to what would have been expected by a person skilled in the art, based on the prior art?
  - (1-5): 2
  - Optional text explaining the effects:

c. Improvement Significance

- On a scale from 1 (not at all) to 5 (great) evaluate how significant and non-obvious is the technical improvement that the invention offers over the closest prior art?
  - (1-5): 2
  - Short text description: The most major improvement is in terms of minimal energy consumption during the finetuning process. However, this is expected since a smaller AI model is utilized as well as fewer training samples during finetuning.

d. Long-Standing Problem

- On a scale from 1 (not at all) to 5 (greatly) evaluate whether the invention solves a problem that persists despite prior attempts?
  - (1-5): 2
  - Optional text about failed prior attempts:

e. Straightforward Extrapolation

- On the scale from 1 (strongly agree) to 5 (strongly disagree), would a person skilled in the art have been motivated or prompted to arrive at the claimed invention merely by a simple extrapolation or straightforward variation of the already known art with a reasonable expectation of success?
  - (1-5): 1
  - Short text description: The techniques used in this work have been previously explored in the literature but have not been applied in this manner within the context of few-shot object detection for PPE and industrial hazards.

f. Feature Juxtaposition

- On the scale from 1 (strongly agree) to 5 (strongly disagree), is the claimed invention merely a juxtaposition or aggregation of known features without any functional interaction that produces a non-obvious synergistic effect?
  - (1-5): 3
  - Short text description: The synergistic effect resulting from the combination of the YOLOv8 model with the partial finetuning technique cannot be considered non-obvious, despite not having been previously explored in this specific context.

g. Secondary Considerations

- On a scale from 1 () to 5 () evaluate the evidence of commercial success, industry praise, or copying by competitors?
  - (1-5): 2
  - Optional text:

e. Required Knowledge

- Define the 'person skilled in the art' relevant to this invention. What would their level of knowledge and expertise be at the time of the invention?
- The person needs to have a basic understanding of ML model training as well as how different hyperparameters of the YOLOv8 model affect its performance.

<u>4. Industrial Applicability</u>

a. Scientific Viability

- On a scale from 1 (not at all) to 5 (certainly) evaluate whether the invention complies with the laws of physics/nature?
  - (1-5): 5
  - Optional text:

b. Practical Use

- On a scale from 1 (no practical use) to 5 (high potential for industrial or other practical application and benefit to humanity), evaluate whether there is a practical use or application for the invention?
  - (1-5): 5
  - Optional text (e.g., target industries): Target industries include construction, manufacturing, logistics, and energy, where worker safety is critical.

c. Reproducibility

- On the scale from 1 (strongly disagree) to 5 (strongly agree), can subject matter experts reproduce and utilize the invention based on the disclosure and their general knowledge, without undue experimentation?
  - (1-5): 4
  - Optional text: The solution can be easily reproduced, given that the subject matters know the model hyperparameters and the datasets used.

d. Solved Technical Problem

- Does the invention solve a technical problem? If yes, describe the problem.
  - Text: The proposed method addresses the challenge of detecting safety-critical objects (like PPE) and hazards (e.g., fires) in industrial environments with limited training data. It solves the problem of maintaining high detection accuracy despite having only a few samples and minimizes computational and energy costs, making it practical for real-world applications.

e. Transferability

- On a scale from 1 (not at all) to 5 (greatly) evaluate the extent that the invention can be transferred in any kind of industry?
  - (1-5): 5

- Optional text: The invention can be applied across various industries such as construction, manufacturing, logistics, and energy, where worker safety and hazard detection are priorities.

9.3 Appendix 3: Sample Questionnaire on EU Artificial Intelligence Act Compliance

The EU AI Act (Regulation (EU) 2024/1689 https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401689) came into force on 1st August 2024. TALON in order to verify compliance with the EU AI Act is invited to provide their replies with respect to the present Questionnaire.

The Questionnaire is addressed to the **Providers of AI Systems** within the TALON. The relevant provisions foreseen in the AI Act are presented in order to enable the Providers to provide their replies.

Name of the Organisation: …………………………………..
Name of the AI system/Model: ………………………………………..

EU AI Act Compliance		AI Act Provisions	Yes/No/N/A	Comments
PART 1				
➢ **General**				
Is the tool considered as an "AI system" according to the EU AI Act? **"AI system" means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.**		**Art. 3 (1)**		
Are you a Provider in accordance with the EU AI Act? **'Provider' means a natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge.**		**Art. 3 (3)**		
Does your system perform any of	Subliminal techniques,	**Art. 5 (1) (a)**		

these functions, which are to be considered prohibited?	manipulation, and deception			
	Exploitation of vulnerabilities of persons	**Art. 5 (1) (b)**		
	Biometric categorisation, except for labelling or filtering of datasets and categorising data in the field of law enforcement	**Art. 5 (1) (g)**		
	Social scoring	**Art. 5 (1) (c)**		
	Individual predictive policing based on profiling people	**Art. 5 (1) (d)**		
	Expanding facial recognition databases through untargeted scraping facial images from the internet or CCTV footage	**Art. 5 (1) (e)**		
	Emotion recognition in the workplace and education institutions, unless for medical or safety reasons	**Art. 5 (1) (f)**		
	Real-time remote biometric identification in publicly accessible spaces by law enforcement, subject to narrow exceptions	**Art. 5 (1) (h)**		
Is the tool a General-purpose AI model (GPAI)? If yes, please fill in the question B.1. **“General-purpose AI model” means an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that**		**Art. 3 (63)**		

displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development or prototyping activities before they are placed on the market.			
Is the GPAI model classified as a GPAI model with systemic risk? If yes, please fill in the question B.1. and B.2. **“Systemic risk” means a risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain; (Art. 3 (65))** **1. A general-purpose AI model shall be classified as a “general-purpose AI model with systemic risk” if it meets any of the following conditions:** **(a) it has high impact capabilities evaluated on the basis of appropriate technical tools and methodologies, including indicators and benchmarks;** **(b) based on a decision of the Commission, *ex officio* or following a qualified alert from the scientific panel, it has capabilities or an impact equivalent to those set out in point (a) having regard to the criteria set out in Annex XIII.** **2. A general-purpose AI model shall be presumed to have high impact capabilities pursuant to paragraph 1, point (a), when the cumulative amount of computation used for its training measured in floating point operations is greater than 10. (Art. 51)**	**Art. 3 (65)** **Art. 51 (1) and (2)**		

Is the tool a high-risk AI system? **a) Both of the conditions provided in Art. 6 par. 1 are fulfilled.** **b) Is the tool referred to in Annex III? (please see below attached herein)** **If yes, please fill in the question C.**	**Art. 6 (1)** **A. 6 (2)** **Annex I** **Section A** **Annex I** **Section B** **Annex III**		
Does the tool pose a significant risk of harm to the health, safety or fundamental rights of any person? **The tool does not pose a significant risk if one of the following conditions are met*:** • **The AI tool is intended to perform a narrow procedural task;** • **The tool is intended to improve the result of a previously completed human activity;** • **the AI system is intended to detect decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment, without proper human review;** • **the AI system is intended to perform a preparatory task to an assessment relevant for the purposes of the use cases listed in Annex III.** ***NOTE: An AI system referred to in Annex III shall always be considered to be high-risk where the AI system performs profiling of natural persons.**	**Art. 6 (3)** **Annex III**		
Does the tool perform any of these functions? - **Interact directly with natural persons?** - **Generate synthetic audio, image, video or text content? If yes, please fill in the question D**	**Art. 50 (1) and (2)**		
	PART 2		
➢ **Obligations**			
B.1. Obligations for Providers of GPAI models			

The technical documentation of the model should be drawn up and kept up-to-date, including its training and testing process and the results of its evaluation, in accordance with Art. 53 (1). **NOTE: There is an exception foreseen for providers of AI models that are released under a free and open-source licence that allows for the access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available, but not for general-purpose AI models with systemic risks.**	**Art. 53 (1) (a), Annex XI** **Art. 53 (2)**		
Information and documentation shall be drawn up, kept up-to-date and made available to providers of AI systems who intend to integrate the GPAI model into their AI systems in accordance with Art. 53 (1). **NOTE: There is an exception foreseen for providers of AI models that are released under a free and open-source licence that allows for the access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available, but not for general-purpose AI models with systemic risks.**	**Art. 53 (1) (b), Annex XII** **Art. 53 (2)**		
A policy should be put in place to comply with Union law on copyright and related rights in particular to identify and comply with, including through state-of-the-art technologies, a reservation of rights expressed pursuant to Article 4(3) of Directive (EU) 2019/790.	**Art. 53 (1) (c)**		
A sufficiently detailed summary shall be drawn up and made publicly available about the content used for training of the GPAI model.	**Art. 53 (1) (d)**		
B.2. Obligations for Providers of GPAI models with systemic risk			

For providers of general-purpose AI models with systemic risk, the providers shall:			
- **perform model evaluation in accordance with standardised protocols and tools reflecting the state of the art, including conducting and documenting adversarial testing of the model with a view to identifying and mitigating systemic risks**	**Art. 55 (1) (a)**		
- **assess and mitigate possible systemic risks at Union level, including their sources, that may stem from the development, the placing on the market, or the use of general-purpose AI models with systemic risk**	**Art. 55 (1) (b)**		
- **keep track of, document, and report, without undue delay, to the AI Office and, as appropriate, to national competent authorities, relevant information about serious incidents and possible corrective measures to address them**	**Art. 55 (1) (c)**		
- **ensure an adequate level of cybersecurity protection for the general-purpose AI model with systemic risk and the physical infrastructure of the model**	**Art. 55 (1) (d)**		
C. Requirements for High-Risk AI Systems			
A Risk Management system has been established, implemented, documented and maintained in accordance with Art. 9 AI Act.	**Art. 9**		

Training, validation and testing data sets meet the following quality criteria: - **They are subject to data governance and management practices appropriate for the intended purpose of the high-risk AI system** - **They are relevant, sufficiently representative, and to the best extent possible, free of errors and complete in view of the intended purpose** - **They take into account, to the extent required by the intended purpose, the characteristics or elements that are particular to the specific geographical, contextual, behavioural or functional setting within which the high-risk AI system is intended to be used.** - **The providers of such systems may exceptionally process special categories of personal data, subject to appropriate safeguards for the fundamental rights and freedoms of natural persons**	**Art. 10**		
Technical documentation has been/will be drawn up before placing on the market and kept up-to date.	**Art. 11**		
The tool technically allows for the automatic recording of events (logs) over the lifetime of the system.	**Art. 12**		
The tool is designed and developed in such a way as to ensure that its operation is sufficiently transparent to enable deployers to interpret a system's output and use it appropriately.	**Art. 13**		
The tool is accompanied by instructions for use in an appropriate digital format or otherwise that include concise, complete, correct and clear information that is relevant, accessible and comprehensible to deployers.	**Art. 13**		

The tool is designed and developed in such a way, including with appropriate human-machine interface tools, that it can be effectively overseen by natural persons during the period in which they are in use.	**Art. 14**		
The tool is ensured through either one or both of the following types of measures: **(a) measures identified and built, when technically feasible, into the high-risk AI system by the provider before it is placed on the market or put into service;** **(b) measures identified by the provider before placing the high-risk AI system on the market or putting it into service and that are appropriate to be implemented by the deployer.**	**Art. 14**		
The oversight measures are commensurate with the risks, level of autonomy and context of use of the high-risk AI system.	**Art. 14**		
The tool is designed and developed in such a way that it achieves an appropriate level of accuracy, robustness, and cybersecurity, and that it performs consistently in those respects throughout its lifecycle.	**Art. 15**		
The level of accuracy of the tool and the relevant accuracy metrics are included in the instructions of use.	**Art. 15**		
Appropriate technical and organisational measures have been taken and the tool meets an appropriate level of technical robustness (for instance backup or fail-safe plans) and resilience regarding errors, faults or inconsistencies that may occur within the system or the environment in which the system operates.	**Art. 15**		
In case the tool continues to learn after its placing on the market, then it is confirmed that the tool is developed in such a way so that it can reduce to the best extent possible the risk of possibly biased outputs influencing input for future operations (feedback loops) and ensure that the feedback loops are duly addressed through appropriate mitigation measures.	**Art. 15**		

Appropriate technical solutions to address cyber resilience of the tool, as regards attempts by unauthorised third parties to alter its use, behaviour or performance, including AI specific vulnerabilities, have been developed and the tool meets an appropriate level of cybersecurity. Alternatively, the tool has been certified or a statement of conformity has been issued under Regulation (EU) 2019/881.	**Art. 15**		
The tool is compliant by design with accessibility requirements, including Directive (EU) 2016/2102 and Directive (EU) 2019/882.	**Art. 16 (l)**		
A quality management system has been put in place, including the aspects of Article 17 AI Act.	**Art. 17**		
D. Transparency obligations			
<u>For Providers of AI systems intended to interact directly with natural persons</u>: The tool is designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system, unless this is obvious from the point of view of a natural person who is reasonably well-informed, observant and circumspect, taking into account the circumstances and the context of use. **<u>Please note that:</u>** **This obligation shall not apply to AI systems authorised by law to detect, prevent, investigate or prosecute criminal offences, subject to appropriate safeguards for the rights and freedoms of third parties, unless those systems are available for the public to report a criminal offence.**	**Art. 50 (1)**		
<u>For Providers of AI systems, including GPAI systems, generating synthetic audio, image, video or text content :</u> The outputs of the tool are marked in a machine-readable format and detectable as artificially generated or manipulated. Providers shall ensure their technical solutions are effective, interoperable, robust and reliable as far as this is technically feasible, taking into account the specificities and limitations of various types of content, the costs of	**Art. 50 (2)**		

implementation and the generally acknowledged state of the art, as may be reflected in relevant technical standards. **<u>Please note that</u>:** **This obligation shall not apply to the extent the AI systems perform an assistive function for standard editing or do not substantially alter the input data provided by the deployer or the semantics thereof, or where authorised by law to detect, prevent, investigate or prosecute criminal offences.**			

<u>Please keep in mind that the following obligations are foreseen in the AI Act:</u>

→ **Obligations of providers of high-risk AI systems**

<u>Article 16 AI Act</u>

Providers of high-risk AI systems shall:

(a) ensure that their high-risk AI systems are compliant with the requirements set out in Section 2;

(b) indicate on the high-risk AI system or, where that is not possible, on its packaging or its accompanying documentation, as applicable, their name, registered trade name or registered trade mark, the address at which they can be contacted;

(c) have a quality management system in place which complies with Article 17;

(d) keep the documentation referred to in Article 18;

(e) when under their control, keep the logs automatically generated by their high-risk AI systems as referred to in Article 19;

(f) ensure that the high-risk AI system undergoes the relevant conformity assessment procedure as referred to in Article 43, prior to its being placed on the market or put into service;

(g) draw up an EU declaration of conformity in accordance with Article 47;

(h) affix the CE marking to the high-risk AI system or, where that is not possible, on its packaging or its accompanying documentation, to indicate conformity with this Regulation, in accordance with Article 48;

(i) comply with the registration obligations referred to in Article 49(1);

(j) take the necessary corrective actions and provide information as required in Article 20;

(k) upon a reasoned request of a national competent authority, demonstrate the conformity of the high-risk AI system with the requirements set out in Section 2;

(l) ensure that the high-risk AI system complies with accessibility requirements in accordance with Directives (EU) 2016/2102 and (EU) 2019/882.

<u>Annex III of the AI Act</u>

ANNEX III

High-risk AI systems referred to in Article 6(2)

High-risk AI systems pursuant to Article 6(2) are the AI systems listed in any of the following areas:

1. Biometrics, in so far as their use is permitted under relevant Union or national law:

   (a) remote biometric identification systems.

   This shall not include AI systems intended to be used for biometric verification the sole purpose of which is to confirm that a specific natural person is the person he or she claims to be;

   (b) AI systems intended to be used for biometric categorisation, according to sensitive or protected attributes or characteristics based on the inference of those attributes or characteristics;

   (c) AI systems intended to be used for emotion recognition.

2. Critical infrastructure: AI systems intended to be used as safety components in the management and operation of critical digital infrastructure, road traffic, or in the supply of water, gas, heating or electricity.

3. Education and vocational training:

   (a) AI systems intended to be used to determine access or admission or to assign natural persons to educational and vocational training institutions at all levels;

   (b) AI systems intended to be used to evaluate learning outcomes, including when those outcomes are used to steer the learning process of natural persons in educational and vocational training institutions at all levels;

   (c) AI systems intended to be used for the purpose of assessing the appropriate level of education that an individual will receive or will be able to access, in the context of or within educational and vocational training institutions at all levels;

   (d) AI systems intended to be used for monitoring and detecting prohibited behaviour of students during tests in the context of or within educational and vocational training institutions at all levels.

4. Employment, workers' management and access to self-employment:

   (a) AI systems intended to be used for the recruitment or selection of natural persons, in particular to place targeted job advertisements, to analyse and filter job applications, and to evaluate candidates;

   (b) AI systems intended to be used to make decisions affecting terms of work-related relationships, the promotion or termination of work-related contractual relationships, to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons in such relationships.

5. Access to and enjoyment of essential private services and essential public services and benefits:

   (a) AI systems intended to be used by public authorities or on behalf of public authorities to evaluate the eligibility of natural persons for essential public assistance benefits and services, including healthcare services, as well as to grant, reduce, revoke, or reclaim such benefits and services;

   (b) AI systems intended to be used to evaluate the creditworthiness of natural persons or establish their credit score, with the exception of AI systems used for the purpose of detecting financial fraud;

   (c) AI systems intended to be used for risk assessment and pricing in relation to natural persons in the case of life and health insurance;

(d) AI systems intended to evaluate and classify emergency calls by natural persons or to be used to dispatch, or to establish priority in the dispatching of, emergency first response services, including by police, firefighters and medical aid, as well as of emergency healthcare patient triage systems.

6. Law enforcement, in so far as their use is permitted under relevant Union or national law:

(a) AI systems intended to be used by or on behalf of law enforcement authorities, or by Union institutions, bodies, offices or agencies in support of law enforcement authorities or on their behalf to assess the risk of a natural person becoming the victim of criminal offences;

(b) AI systems intended to be used by or on behalf of law enforcement authorities or by Union institutions, bodies, offices or agencies in support of law enforcement authorities as polygraphs or similar tools;

(c) AI systems intended to be used by or on behalf of law enforcement authorities, or by Union institutions, bodies, offices or agencies, in support of law enforcement authorities to evaluate the reliability of evidence in the course of the investigation or prosecution of criminal offences;

(d) AI systems intended to be used by law enforcement authorities or on their behalf or by Union institutions, bodies, offices or agencies in support of law enforcement authorities for assessing the risk of a natural person offending or re-offending not solely on the basis of the profiling of natural persons as referred to in Article 3(4) of Directive (EU) 2016/680, or to assess personality traits and characteristics or past criminal behaviour of natural persons or groups;

(e) AI systems intended to be used by or on behalf of law enforcement authorities or by Union institutions, bodies, offices or agencies in support of law enforcement authorities for the profiling of natural persons as referred to in Article 3(4) of Directive (EU) 2016/680 in the course of the detection, investigation or prosecution of criminal offences.

7. Migration, asylum and border control management, in so far as their use is permitted under relevant Union or national law:

(a) AI systems intended to be used by or on behalf of competent public authorities or by Union institutions, bodies, offices or agencies as polygraphs or similar tools;

(b) AI systems intended to be used by or on behalf of competent public authorities or by Union institutions, bodies, offices or agencies to assess a risk, including a security risk, a risk of irregular migration, or a health risk, posed by a natural person who intends to enter or who has entered into the territory of a Member State;

(c) AI systems intended to be used by or on behalf of competent public authorities or by Union institutions, bodies, offices or agencies to assist competent public authorities for the examination of applications for asylum, visa or residence permits and for associated complaints with regard to the eligibility of the natural persons applying for a status, including related assessments of the reliability of evidence;

(d) AI systems intended to be used by or on behalf of competent public authorities, or by Union institutions, bodies, offices or agencies, in the context of migration, asylum or border control management, for the purpose of detecting, recognising or identifying natural persons, with the exception of the verification of travel documents.

8. Administration of justice and democratic processes:

(a) AI systems intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law and in applying the law to a concrete set of facts, or to be used in a similar way in alternative dispute resolution;

(b) AI systems intended to be used for influencing the outcome of an election or referendum or the voting behaviour of natural persons in the exercise of their vote in elections or referenda. This does not include AI systems to the output of which natural persons are not directly exposed, such as tools used to organise, optimise or structure political campaigns from an administrative or logistical point of view.

References

[1] "HEU Results platform.pdf." Accessed: Sep. 25, 2025. [Online]. Available: https://intellectual-property-helpdesk.ec.europa.eu/system/files/2022-02/HEU%20Results%20platform.pdf
[2] "Homepage | epo.org." Accessed: Sep. 25, 2025. [Online]. Available: https://www.epo.org/en
[3] U. O. of P. Affairs (OPA), "United States Patent and Trademark Office." Accessed: Sep. 25, 2025. [Online]. Available: https://www.uspto.gov/
[4] "Japan Patent Office." Accessed: Sep. 25, 2025. [Online]. Available: https://www.jpo.go.jp/e/
[5] S. Sarkissian, S. E. Trevlakis, and T. Lagkas, "TALON's Patentability Assessment," Zenodo, Jun. 2025. doi: 10.5281/zenodo.15687441.
[6] "Face Anonymization with GStreamer," Fluendo. Accessed: Sep. 25, 2025. [Online]. Available: https://fluendo.com/blog/face-anonymization-with-gstreamer/
[7] "Verifiable robustness, energy efficiency and transparency for Trustworthy AI: Scientific excellence boosting industrial competitiveness (AI, Data and Robotics Partnership) (RIA) | Programme | HORIZON | CORDIS | European Commission." Accessed: Sep. 25, 2025. [Online]. Available: https://cordis.europa.eu/programme/id/HORIZON_HORIZON-CL4-2021-HUMAN-01-01/en
[8] "AIDEAS Supportive Partners Programme." Accessed: Sep. 25, 2025. [Online]. Available: https://fundingbox.eu/aideas-supportive-partners-programme
[9] I. Brillantes, "Top Edge Computing Platforms for 2025," OTAVA. Accessed: Sep. 25, 2025. [Online]. Available: https://www.otava.com/blog/top-edge-computing-platforms-for-2025/
[10] M. R. P. Ltd, "Edge Computing Market worth $110.6 billion by 2029, growing at a CAGR of 13.0%: Report by MarketsandMarkets™," GlobeNewswire News Room. Accessed: Sep. 25, 2025. [Online]. Available: https://www.globenewswire.com/news-release/2024/04/30/2872426/0/en/Edge-Computing-Market-worth-110-6-billion-by-2029-growing-at-a-CAGR-of-13-0-Report-by-MarketsandMarkets.html
[11] "Edge Computing Market Size, Share | Industry Report, 2033." Accessed: Sep. 25, 2025. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/edge-computing-market
[12] "Cloud Computing Market worth $2,281.1 billion by 2030." Accessed: Sep. 25, 2025. [Online]. Available: https://www.marketsandmarkets.com/PressReleases/cloud-computing-market.asp
[13] "Explainable AI Market," Market.us. Accessed: Sep. 25, 2025. [Online]. Available: https://market.us/report/explainable-ai-market/
[14] "Explainable AI Market worth $16.2 billion by 2028." Accessed: Sep. 25, 2025. [Online]. Available: https://www.marketsandmarkets.com/PressReleases/explainable-ai.asp
[15] "Cloud Orchestration Market Size, Trends | Forecast to 2033." Accessed: Sep. 25, 2025. [Online]. Available: https://www.imarcgroup.com/cloud-orchestration-market
[16] Z. M. Research, "Cloud Orchestration Market Size, Share and Forecast 2032," Zion Market Research. Accessed: Sep. 25, 2025. [Online]. Available: https://www.zionmarketresearch.com/report/cloud-orchestration-market
[17] "Blockchain Market worth $248.9 billion by 2029 - Exclusive Report by MarketsandMarkets™." Accessed: Sep. 25, 2025. [Online]. Available: https://www.prnewswire.com/news-releases/blockchain-market-worth-248-9-billion-by-2029---exclusive-report-by-marketsandmarkets-302195622.html
[18] "Blockchain Technology Market Size | Industry Report, 2030." Accessed: Sep. 25, 2025. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/blockchain-technology-market
[19] T. L. Saaty, "A scaling method for priorities in hierarchical structures," *Journal of Mathematical Psychology*, vol. 15, no. 3, pp. 234–281, Jun. 1977, doi: 10.1016/0022-2496(77)90033-5.
[20] T. L. Saaty and M. S. Ozdemir, "Why the magic number seven plus or minus two," *Mathematical and Computer Modelling*, vol. 38, no. 3, pp. 233–244, Aug. 2003, doi: 10.1016/S0895-7177(03)90083-5.
[21] S. Pant, A. Kumar, M. Ram, Y. Klochkov, and H. K. Sharma, "Consistency Indices in Analytic Hierarchy Process: A Review," *Mathematics*, vol. 10, no. 8, p. 1206, Jan. 2022, doi: 10.3390/math10081206.
[22] F. Pauer, K. Schmidt, A. Babac, K. Damm, M. Frank, and J.-M. G. Von Der Schulenburg, "Comparison of different approaches applied in Analytic Hierarchy Process – an example of information needs of patients with rare diseases," *BMC Med Inform Decis Mak*, vol. 16, no. 1, p. 117, Dec. 2016, doi: 10.1186/s12911-016-0346-8.

[23] "AI You Can Trust: Standards, Ethics, and Innovation," Eventbrite. Accessed: Sep. 25, 2025. [Online]. Available: https://www.eventbrite.co.uk/e/ai-you-can-trust-standards-ethics-and-innovation-tickets-1273673394409
[24] S. Sarkissian and S. E. Trevlakis, "TALON's Standardisation Activities," Zenodo, Jun. 2025. doi: 10.5281/zenodo.15687685.
[25] "AI Act enters into force - European Commission." Accessed: Sep. 25, 2025. [Online]. Available: https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en
[26] "Implementation Timeline | EU Artificial Intelligence Act." Accessed: Sep. 25, 2025. [Online]. Available: https://artificialintelligenceact.eu/implementation-timeline/
[27] *COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE EUROPEAN COUNCIL, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS The European Green Deal*. 2019. Accessed: Sep. 25, 2025. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2019%3A640%3AFIN
[28] *COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS A new Circular Economy Action Plan For a cleaner and more competitive Europe*. 2020. Accessed: Sep. 25, 2025. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1583933814386&uri=COM:2020:98:FIN
[29] *COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS on a revised monitoring framework for the circular economy*. 2023. Accessed: Sep. 25, 2025. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2023%3A306%3AFIN&qid=1684143860344
[30] "Regulation - EU - 2024/1781 - EN - EUR-Lex." Accessed: Sep. 25, 2025. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2024/1781/oj/eng
[31] "Key actions launched to advance the circular economy," European Commission - European Commission. Accessed: Sep. 25, 2025. [Online]. Available: https://ec.europa.eu/commission/presscorner/detail/en/ip_25_1710
[32] "Delegated regulation - 2019/945 - EN - EUR-Lex." Accessed: Sep. 25, 2025. [Online]. Available: https://eur-lex.europa.eu/eli/reg_del/2019/945/oj/eng
[33] "Easy Access Rules for U-space (Regulation (EU) 2021/664) - May 2024 — Available in pdf, online & XML format | EASA." Accessed: Sep. 25, 2025. [Online]. Available: https://www.easa.europa.eu/en/document-library/easy-access-rules/easy-access-rules-u-space-regulation-eu-2021664
[34] "Implementing regulation - 2019/947 - EN - EUR-Lex." Accessed: Sep. 25, 2025. [Online]. Available: https://eur-lex.europa.eu/eli/reg_impl/2019/947/oj/eng
[35] "Predefined Risk Assessment (PDRA) | EASA." Accessed: Sep. 25, 2025. [Online]. Available: https://www.easa.europa.eu/en/domains/drones-air-mobility/operating-drone/specific-category-civil-drones/predefined-risk-assessment-pdra
[36] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, PMLR, Apr. 2017, pp. 1273–1282. Accessed: Sep. 30, 2025. [Online]. Available: https://proceedings.mlr.press/v54/mcmahan17a.html
[37] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks," Apr. 21, 2020, *arXiv*: arXiv:1812.06127. doi: 10.48550/arXiv.1812.06127.
[38] S. Reddi *et al.*, "Adaptive Federated Optimization," Sep. 08, 2021, *arXiv*: arXiv:2003.00295. doi: 10.48550/arXiv.2003.00295.

Funded by
the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101070181