Ref. Ares(2026)7460529 - 29/07/2026

Artificial Intelligence in Manufacturing leading to Sustainability and Industry5.0

Project Acronym:

AIMS5.0

Grant agreement no.: 101112089

Deliverable no. and title	**D6.6 – Summarizing public report about Tier 2 Use Cases, demonstrations and validation towards sustainability, productivity, resilience, human centricity and potentials for the EU Industry**	
Work package	WP6	Artificial Intelligence for Operational Excellence in the ECS (Tier 2 Use Cases)
Tasks	T6.1-6.3 (all Tasks)	
Subtasks involved	ST6.1.1-6.3.4 (all Subtasks)	
Lead contractor	Infineon Technologies AG Thomas Gutt, thomas.gutt@infineon.com	
Deliverable responsible	Infineon Technologies Dresden AG & Co. KG Germar Schneider, germar.schneider@infineon.com	
Version number	V1.0	
Due Date	29/07/2026	
Status	Final	
Dissemination level	PU	

Copyright: AIMS5.0 Project Consortium, 2026

Authors

Participant no.	Part. short name	Author name	Chapter(s)
1	IFAG	Nicola Pieri	2.4
5	FMX	Karli Hantzschmann, Karl-Benedikt Reith	2.3
6	FUH	Lars Mönch	2.4
9	HSZG	Fabian Lindner, Pascal Herfort	all
10	IFD	Germar Schneider	all
11	SYSTEMA	Gerhard Luhn	2.1
12	TUD	Sebastian Rank	2.3
14	AIT	Abdelkader Shaaban, Oliver Zendel	2.2
16	IFAT	Silke Braun	2.4
18	ZelosPlant	Ralph Baldrian	2.2
19	RSA	Belal Abu Naim	2.2
20	TIAG	Francesca Flamigni	2.2
46	GUT	Mateusz Groth	2.2

Executive Summary

Deliverable 6.6 presents the status of the Use Cases (UCs) in Work Package 6 “Artificial Intelligence for Operational Excellence in the ECS (Tier 2 Use Cases)” by the end of the third and last project year of AIMS5.0. Deliverable 6.6 focuses on the achievements of the UCs whose content shall be made publicly available to the European public and industry. Confidential Use Cases are evaluated in Deliverable 6.4.

The UCs describe their progress during Year 3 of the project and their contribution to the defined goals concerning sustainability, productivity, resilience, and human centricity. Therefore, each UC describes the baseline from the first project year, the target KPI and its approaches and results in the last project year. Overall, the UCs achieved their goals to a large extent, producing a multitude of exploitable results and contributions to the different goal dimensions.

Table of Contents

1 Introduction 6
1.1 Objective and scope of the document 6
1.2 Structure of the Deliverable report 6
2 Use Case Status 7
2.1 Use Case 10 (Task 6.1.2): MES concept based on AI Improving Production Cycle (AI-Prod-MES) 7
2.1.1 Use Case Characteristics 7
2.1.2 Results 8
2.1.3 Learnings and Outlooks towards Project End and Beyond 10
2.2 Use Case 11 (Task 6.1.3): AI-supported Industrial IoT for Indoor Food Production 11
2.2.1 Baseline and validation 11
2.3 Use Case 14 (Task 6.2.1/6.2.2): Speeding-up the Benefit of Automation by Virtual Commissioning through Digital Twins 27
2.3.1 Use Case Characteristics 27
2.3.2 Approach/Method 28
2.3.3 Results 28
2.3.4 Learnings and Outlooks towards Project End and Beyond 30
2.4 Use Case 18 (Task 6.1.5): Using semantic Data Model as base for AI optimized economic & ecologic/societal supply chain improvements 30
2.4.1 Use Case Characteristics 30
2.4.2 Approach/Method 31
2.4.3 Results 32
2.4.4 Learnings and Outlooks towards Project End and Beyond 33
3 Conclusion 34
Appendices 35
Appendix A - Abbreviations 35

List of Figures

Figure 1: detailed challenges of the existing MES baseline, which in summa leads to the conclusion, that the current MES baseline is not ready for a required AI-supporting production automation infrastructure 8
Figure 2: Architecture and methodology for exploring previously invisible production potentials; basis: one year of production data from the entire factory of project partner IFD. 9
Figure 3: Improvement in Overall Equipment Effectiveness (OEE) and cycle time based on production data from project partner IFD. 9
Figure 4: Distribution of measured end-to-end latencies (standard OS) 13
Figure 5: Distribution of measured end-to-end latencies (RT-OS) 13
Figure 6: Received signal strength achieved for the Simulated Annealing algorithm 14
Figure 7: Received signal strength achieved for the genetic algorithm 14
Figure 8: Autonomic Management of a Hydroponics System for Industrial Food Production 16
Figure 9: LLM in Action 17
Figure 10: The hydroponics demonstrator 18
Figure 11: Humidity recovery and exhaust-fan response after a mist disturbance. Both relative humidity and fan speed are shown on the same percentage scale. Solid lines denote relative humidity, dashed lines denote fan-speed commands, and the dotted horizontal line indicates the relative-humidity threshold. Passive recovery is shown with the fan held at 0 %. 19
Figure 12: AI-based vs. non-AI-based Power consumption of the demonstrator 20
Figure 13: Per device energy consumption 21
Figure 14: ThreatGet high-level architectural model of UC11 23
Figure 15 Asset-to-damage-scenario mapping for UC11 in ThreatGet 23
Figure 16 ThreatGet methodology applied in UC11 24
Figure 17: Attack Tree, Environmental Sensor compromise: Bottom-up kill chain from five Shared-Medium-Interface entry points to the Environmental Sensor via Sensor Data Processing 25
Figure 18 ThreatGet Analysis View 26
Figure 19: Overall ThreatGet coverage of the CRA essential cybersecurity requirements 26
Figure 20: Task relation 28
Figure 21: Illustration of why semantically connected data can support ATP and OC modeling. 31
Figure 22: UC18 overview, workflow and achievements. 32

List of Tables

Table 1: Abbreviations 35

1 Introduction

1.1 Objective and scope of the document

The goal of Deliverable 6.6 is to provide a detailed status report about Tier 2 Use Cases (UC), demonstrations and validation towards sustainability, productivity, resilience and human centricity for year 3 (Y3) and potentials for the European Industry including decarbonisation. Deliverable 6.6 focuses on the achievements of all UCs in Work Package 6 (WP6) whose content shall be made publicly available to the European public and industry. Deliverable 6.6 represents the public version of Deliverable 6.4.

By the end of the project, the final iterations of improvements in the UCs are finished, the results in terms of level of achievement of the envisioned KPIs are reported and more general implications for the European Industry are derived. The overall implementation journeys of AI in the UCs of WP6 are summarized and highlighted.

The Deliverable also refers to the beginning of the project. The first Deliverable 6.1 included the lists of requirements and specifications for each UC in WP6. Additionally, Deliverable 6.1 linked these requirements and specifications to the overall goals of AIMS5.0: Sustainability, Productivity, Resilience, and Human Centricity.

Now at the end of the project, each UC evaluates its specific contribution to the defined goals, referencing the problem in Y1 and the assessment of target achievement.

1.2 Structure of the Deliverable report

This deliverable report is structured along the selected public WP6 UCs. The report consists of three sections for each UC in Section 2: (1) the characteristics of the UC including the baseline from Y1 and the target KPI from Y1, (2) the approaches and methods to achieve the goals in Y3, (3) the concrete results in Y3 with regard to the goal dimensions of Deliverable 6.1 and (4) the learnings and outlooks of this UC towards the project end and beyond.

Finally, the overall status of the WP6 UCs is summarized in Section 3. Also, the contribution of this document to the progression and success of the overall Project and adjunct WPs is highlighted.

2 Use Case Status

2.1 Use Case 10 (Task 6.1.2): MES concept based on AI Improving Production Cycle (AI-Prod-MES)

Lead: SYSTEMA; **Partners:** IFD, IFAT, IFAG

2.1.1 Use Case Characteristics

Problem/Baseline

The aim of this activity is to research and demonstrator a new, AI-enabling version of the Manufacturing Execution System (MES), in order to achieve an overall, balanced and effective information flow orchestration. This orchestration is to enable further integration of innovative, transparent AI components in the area of semiconductor manufacturing automation.

In view of the challenging "Myriad Product Mix", the core task is to integrate a new level of adaptability and "smartness" of the production system by embedding AI methods in an explorable, redesigned and real-time capable MES and high automation framework, so that the requirements of the end customers are met.

Goal/Target KPIs

UC10 aims to improve the manufacturing process and validate the results based on detailed analysis of productivity characteristics. Those characteristics are based on Factory Physics and cover main productivity indicators, such as production cycle time, OEE, flow factor and others.

Those improvements reflect customers' requirements and have been mapped to research, develop and implement an AI-enabling version of the Manufacturing Execution System (MES). Accordingly, following 6 Subtasks have been worked out:

- Subtask 1: Research into new requirements for databases to harmonize information in the existing data landscape; a first core component is the ***database layer***; different available solutions have been analyzed and first evaluation tasks have been done.
- Subtask 2: Research distributed SW Eng. Proc. ("MES Airbus"): Based on the requirements analysis, the ***Real-time Adaptive Manufacturing Information Infrastructure (RAMII)*** architecture has been chosen to research, develop and implement representative components within an "MES Airbus" framework
- Subtask 3: Phasing out old/harmonization new core components;
- Subtask 4: Design of migration scenarios, risk analysis: From within the MES architecture, functional patterns form the components ***WIP***, ***Equipment*** and ***Factory Communication Services*** have been created and researched
- Subtask 5: MES Interface to the holistic model HIM;
- Subtask 6: Improved data visualization from MES.

In more detail, following challenges have been identified and analyzed (Figure 1):

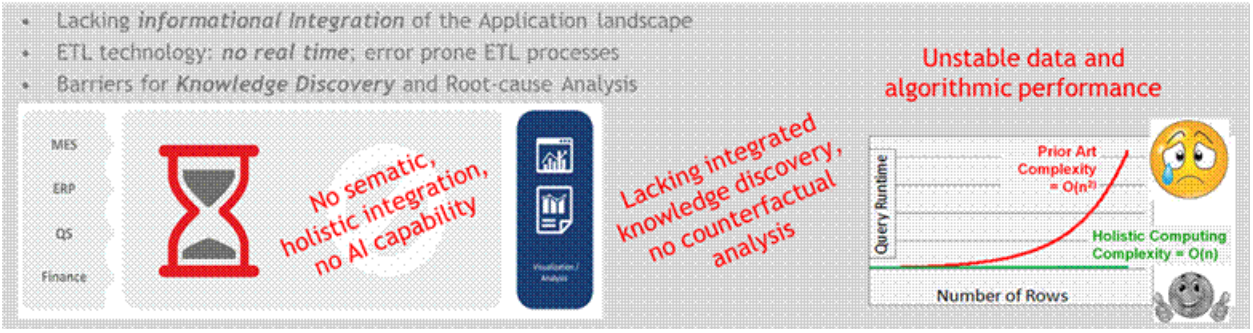

Figure 1: detailed challenges of the existing MES baseline, which in summa leads to the conclusion, that the current MES baseline is not ready for a required AI-supporting production automation infrastructure

Approach/Method

SYSTEMA and partners researched the six subtasks, designed an appropriate target architecture, developed functional pattern and integrated an overall demonstrator in the context of UC10. In more detail:

- Subtask 1: Evaluation tasks regarding the selection and assessment of the database layer were successfully completed. A database component was selected and incorporated into a functional demonstrator.
- Subtask 2: Further integration of the database layer was carried out based on the Real-time Adaptive Manufacturing Information Infrastructure (RAMII) architecture. On this basis, an integrated functional demonstrator of the "MES Airbus" framework was implemented and investigated in accordance with the project plan. Initial test results confirm that this component meets the requirements.
- Subtasks 3, 4: Within this MES architecture, functional requirements and patterns for the WIP, Equipment, and Factory Communication components were incorporated into a functional demonstrator, and the integrated functionality of the services was successfully investigated.
- Subtasks 5, 6: Building on the semantic mapping and the import of one year of IFD production data into the holistic information model (HIM), previously unidentified potential for production optimization was discovered and analyzed.

2.1.2 Results

Fulfillment of Requirements and Specifications

These subtasks are the basis for further detailed analyses and production optimizations within the scope of specific use cases. Development and investigation of initial functional demonstrator for subtasks 1–6: Analyses were conducted and an initial integrated functional demonstrator was created (focusing on the database component and MES components: WIP, equipment, and factory communication services; see Figure 2). Using these components, the transition of the software engineering process—specifically the MES—from a local setup to a distributed or distributable structure ("MES-Airbus") was concretely validated and successfully tested. The goal is to characterize lead time, production throughput, and equipment utilization in greater detail in order to optimize production, which is still characterized by high and—in part—uncontrolled fluctuations. Close cooperation with partners took place. Based on this, specific requirements were analyzed and functional demonstrator were created:

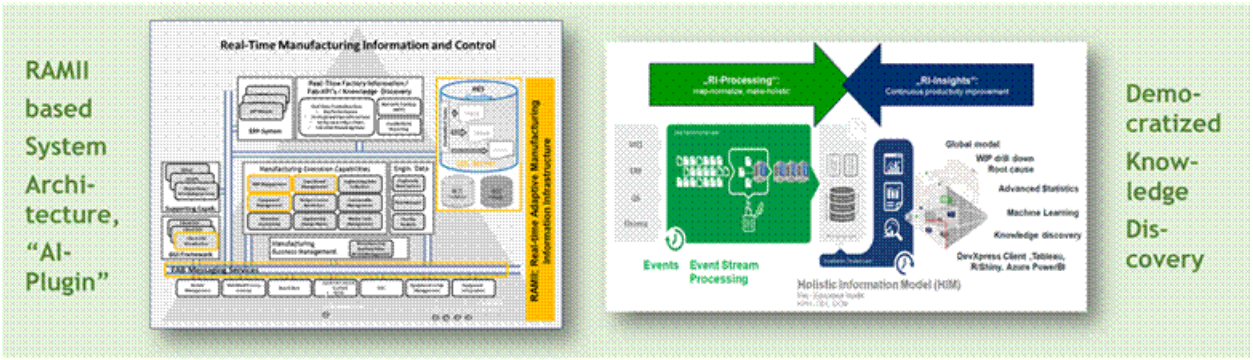

Figure 2: Architecture and methodology for exploring previously invisible production potentials; basis: one year of production data from the entire factory of project partner IFD.

The new MES architecture for analyzing and optimizing productivity potential in critical work areas was investigated in greater detail, implemented in a demonstrator, and validated. New planning approaches were employed to achieve a significant increase in Overall Equipment Effectiveness (OEE). This result was made possible, in particular, by the fine granularity of the factory representation within the holistic model. To this end, the multi-level visualization concept was further refined and expanded.

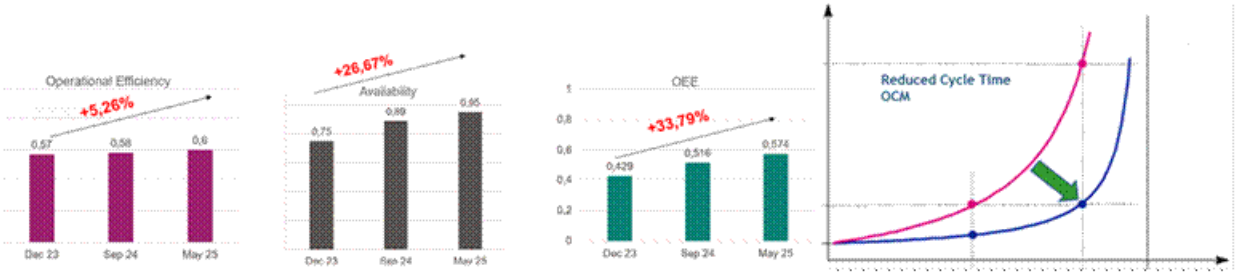

Figure 3: Improvement in Overall Equipment Effectiveness (OEE) and cycle time based on production data from project partner IFD.

In more detail, following highlights were achieved:

- 1 Year of production data has been successfully transferred into the holistic information model HIM
- Anomaly detection, functional pattern available Y2, capability ~ 5 % shown
- RAMII architecture specification done
- First AI-based trainings for adapted/enhanced production scenarios performed (UC17)
- Improved throughput characterization based on new anomaly detection capability
- First results on development of new methods and processes as enabler to enhance operational efficiency shown
- WipFlowShare-degree of utilization of equipment is characterized from the perspective of the dynamics of the entire production: new functionality to detect possible bottleneck usage of one-of-a-kind equipment
- Building efficient structures in MES as basis for an improved material & production flow, digital processes for production supply production lines à This makes the highly material and cost-intensive production of semiconductors more environmentally friendly

Achievement of KPIs

Summary: The defined KPIs (AI based MES capabilities) and MES acts as a main enabler for 10% productivity increasement.

In more detail:

- IT-Perspective:
    - AI based MES capability > 10 % (more than 10 % of architectural core components of the MES have been researched, designed and integrated in an autonomous demonstrator)
- Production optimization perspective:
    - Proven Self-adjusted/-adapted manufacturing Processes > 15 % (functionalities which enable of self-adjusted manufacturing flows: such as intelligent coupling between production processes which have to fulfill time constraints; example: time coupling between wet clean process and furnace process: when cleaning takes place, a time window opens and the furnace process has to take place within this window)
    - CT reduction of critical workcenter and new product ramps > 10% (for example to identify and remove production bottlenecks)

Impact on Sustainability and Decarbonization:

The Reference Architectural Model Industrie 4.0 (RAMI 4.0) is currently being extended to manage environmental footprints. An example is the concept of energy efficiency, which is an important topic in AIMS50. Following RAMI 4.0 subconcepts deal with this aspect and are part of different AIMS50 activities:

- **Life Cycle & Value Stream:** Tracks the environmental footprint of physical products and assets (e.g., carbon emissions and material recycling) from raw materials to final disposal.
- **Hierarchy Levels:** Integrates sustainability metrics across all organizational levels, from individual factory field devices up to connected, global supply chains.
- **IT Layers:** Uses the upper **Business** and **Functional** layers to orchestrate green Key Performance Indicators (KPIs) and drive energy-aware production processes.

In AIMS50, those aspects are part of the digital twin/holistic information model.

2.1.3 Learnings and Outlooks towards Project End and Beyond

The MES serves as the production backbone and has been proven to be a synergistic partner for the UC9, UC12, UC15, and UC17.

In summary, the MES solution under investigation demonstrates the expected technological progress in line with project objectives by integrating AI-based data analysis, real-time capabilities, and semantic interoperability within a unified architectural framework (RAMI). Achieving these goals enables the transition from static, transaction-driven production control to dynamic, self-optimizing manufacturing processes. The resulting system supports continuous performance monitoring and adaptive process adjustments based on real-time data streams, thereby

2.2 Use Case 11 (Task 6.1.3): AI-supported Industrial IoT for Indoor Food Production

Lead: TIAG; **Partners:** RSA, AIT, PYT, GUT

Problem/Baseline
This use case focuses on improving automation across the plant lifecycle through a flexible, decentralized indoor food production framework based on interconnected IoT devices, AI-driven automation which enables security, safety, and scalability. This introduces cybersecurity challenges, including compliance with latest regulations such as Cyber Resilience Act (CRA).

The approach for optimizing plant growth and production efficiency while ensuring secure and scalable operations, includes the following sub-scenarios:

1. Visual Plant Mass Estimation
2. AI-driven antenna beam switching for wireless communication optimization
3. System-of-systems architectures and microservices for AI-enabled production
4. CRA support through threat modeling and risk assessment

2.2.1 Baseline and validation

Sub-Scenario 1 – Visual Plant Mass Estimation

The current baseline for plant inspection is a manual process where workers have to periodically visually inspect the full plantation. Due to the time and cost constraints this is done between 3-7 times a week. The current incident response can thus be up to 2 days.

The validation for the plant instance module uses annotated instance segmentation data from the data campaign which was not used during training. This validation set contains similar plant scenes and associated plant instance masks. The evaluation is done by comparing the learned model‘s predictions (instance mask per plant) vs. the provided ground truth masks. The standard metrics accuracy and recall are used. The validation of the plant stress detection incident response time uses dedicated plant stress sequences. Here plant stresses are introduced deliberately (e.g. pump failure, wrong pH settings). The response time is calculated as the time until the visual plant mass estimation flags the subsequent growth under stress as abnormal.

Sub-Scenario 2 – Wireless communication optimization by using AI-driven antenna beam switching

At present, device-to-gateway communication in Indoor Food Production facilities is largely dependent on wired networks. The proposed approach introduces a scalable ESPAR antenna system with adaptive beam-switching, eliminating the need for physical cabling. By dynamically adjusting its radiation pattern, the system delivers reliable, flexible, and efficient wireless connectivity across the entire facility improving the reliability of the connection over the typical monopole antenna that radiates uniformly around it.

Baseline/problem

Initially the control application was executed on a dedicated MFN100 edge device running a standard Linux operating system (Ubuntu 24.04) without real-time operating system support. The evaluated workload comprised AI-based vector estimation, signal processing, and UART monitoring/configuration applications representative of an industrial edge-computing environment. End-to-end latency was measured as the elapsed time between the reception of a valid UART packet and the transmission of the corresponding steering vector generated by the

application. The measured latency under this configuration served as the reference baseline for following performance assessments.

This baseline represents a conventional non-real-time deployment scenario for industrial edge applications.

The wireless connectivity problem was the limited reliability of standard omnidirectional wireless links between distributed plant/soil sensors and the edge processing unit in an indoor farming environment. The baseline solution was therefore defined as communication using a regular omnidirectional antenna without spatial adaptation or intelligent control of the radiation pattern.

Goal/Target KPIs

The KPI defined for UC11 targeted **"Improve end-to-end latency by at least 10%"** compared to the baseline implementation.

The validation was planned through experimental measurements performed on a real control workload running on the MFN100 platform. The objective was to evaluate the extent to which the adoption of real-time operating system capabilities could improve system responsiveness, reduce latency variability, and enhance timing determinism under realistic operating conditions.

The goal was to improve wireless communication quality by using a switched-beam ESPAR antenna and AI-based optimisation methods. The main target KPI was the improvement of signal quality compared with the omnidirectional antenna baseline, expressed through SNIR and supported by PER-based link quality assessment. In addition, the solution was expected to support fast and repeatable steering-vector selection, allowing the gateway to react to changing propagation conditions.

Approach/Method

During Year 3, TTTech implemented and validated a real-time execution environment based on Nerve RT-OS and compared it against the baseline Ubuntu configuration.

The Experimental Setup consisted of of the same hardware platform and execution of the same control workload. TTTech compared between Ubuntu 24.04 (baseline is presented in Figure 4) and Nerve RT-OS in stress-test conditions generating at least 80 % CPU utilisation to emulate realistic industrial workloads. The evaluation was based on 35,000 statistically relevant measurements.

In Year 3, the work focused on validating the wireless connectivity improvement approach in a dedicated, remotely controlled testbed developed within the project. The setup uses 36 BLE-based transmitting agents arranged around the ESPAR antenna, enabling repeatable evaluation of different signal directions and communication conditions. Two optimisation methods were used for steering-vector selection: Simulated Annealing and Genetic Algorithm.

Methodology and results

The end-to-end execution time was measured from reception of a valid UART package to generation and transmission of a new steering vector. Latency distributions were collected for both operating systems and analysed with respect to average latency, maximum latency, standard deviation and number of outliers. The outliers were defined as values exceeding ±2 standard deviations from the average latency.

The validation task demonstrated that the improved RT-OS configuration significantly outperformed the baseline implementation. Based on 35,000 measurements collected under high-load conditions, the following improvements compared to the baseline were obtained:

Average latency:	35 % reduction
Maximum latency:	56 % reduction
Standard deviation:	75 % reduction
Outlier count:	48 % reduction

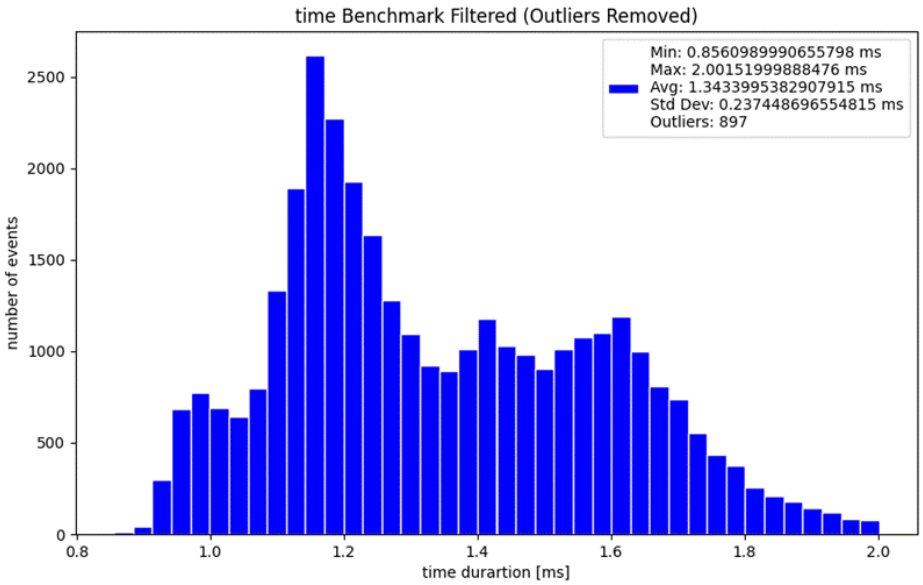

Figure 4: Distribution of measured end-to-end latencies (standard OS)

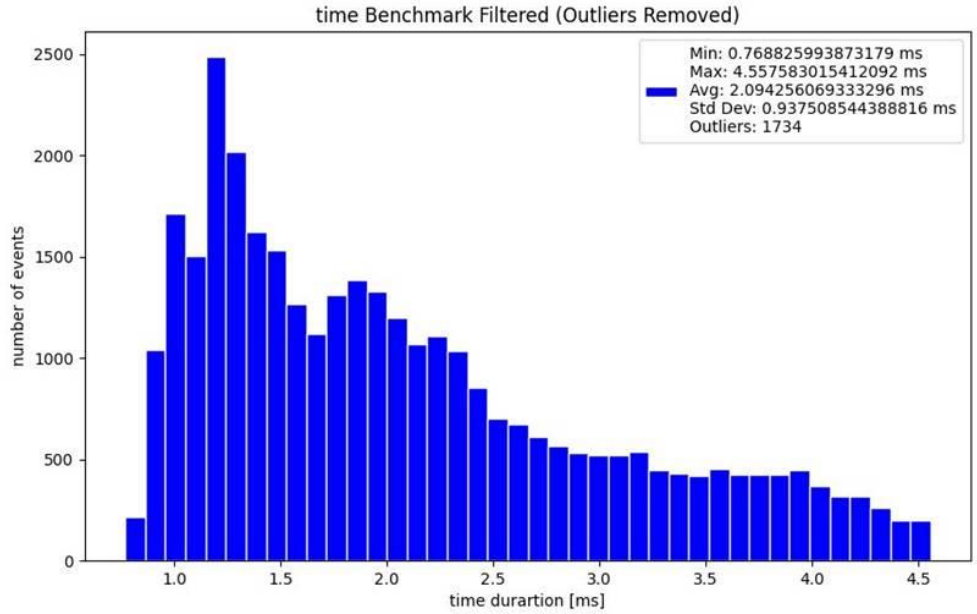

Figure 5: Distribution of measured end-to-end latencies (RT-OS)

The achieved 35 % reduction in average end-to-end latency exceeded the targeted KPI of 10 %. For the wireless connectivity improvement, the methodology consisted of comparing the performance of the AI-optimised switched-beam antenna against the omnidirectional baseline in repeated test runs. Each optimisation method was tested in 1000 runs, using RSSI/SNIR as the main quality indicator. Both Simulated Annealing and Genetic Algorithm achieved an improvement of approximately 8 dB in received signal quality compared with the baseline, with typical convergence after around 135 iterations for Simulated Annealing and around 120 iterations for the Genetic Algorithm leading to 400 % improvement over the omnidirectional-antenna-based baseline. The observed standard deviation was approximately 1.2 dB, indicating stable and repeatable behaviour.

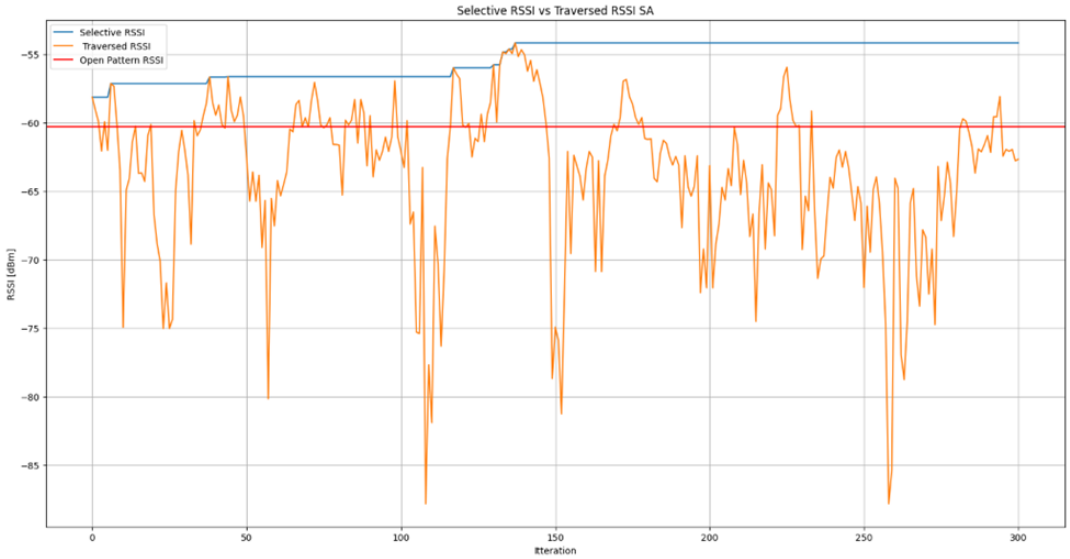

Figure 6: Received signal strength achieved for the Simulated Annealing algorithm

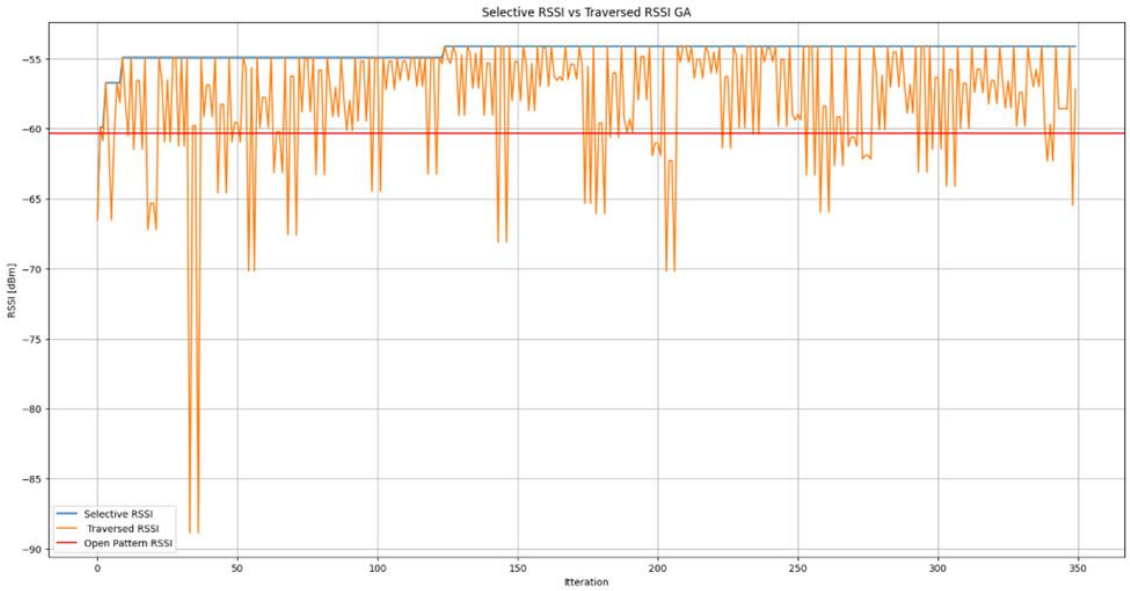

Figure 7: Received signal strength achieved for the genetic algorithm

Learnings and Outlooks towards Project End and Beyond

The obtained results demonstrated that introducing RT-OS on the MFN100 edge platform significantly improves latency compared to a standard Linux deployment.

The results support the adoption of real-time edge computing technologies for industrial applications requiring fast and predictable responses. The improvements, which are illustrated in Figure 5, shows that edge computing could effectively be used for future AI-enabled control systems for enhanced reliability, responsiveness, and operational efficiency.

Sub-Scenario 3 – SoS-based architectures & microservices for AI-enabled sustainable production

This sub-scenario demonstrates how repetitive manual activities—such as continuously observing environmental conditions, interpreting trends, selecting appropriate climate control actions, and adjusting equipment settings—can be replaced by an autonomic management system operating directly at the network edge. In the current operation, several critical processes remain manual, including pH and electrical conductivity (EC) analysis, which is performed only once per day for each closed system without automated or real-time monitoring, and power consumption management, where energy usage is neither optimized nor supported by AI-based actuation or load balancing mechanisms. The proposed approach reduces operator workload, enables continuous real-time optimisation of growing conditions, and establishes an extensible platform for future autonomous management of additional indoor farming processes, including irrigation, nutrient dosing, pH and EC regulation, lighting optimisation, and energy management. By integrating Industrial IoT devices, edge computing, and explainable AI planning into a unified management architecture, the demonstrator provides a practical realization of automated planning and control for next-generation indoor food production systems.

Baseline and Problem

Baseline 1, this scenario relies on daily manual supervision and conventional rule-based control. Environmental parameters are monitored by operators, who interpret sensor data and manually configure ventilation, lighting, and other actuators to maintain suitable growing conditions. While effective, this approach requires continuous human involvement, offers limited adaptability to changing conditions, and does not exploit AI-assisted decision-making or autonomous edge-based control.

Baseline 2, this scenario represents the conventional operation of the indoor farming demonstrator without AI-assisted control. Environmental control devices, including lighting, ventilation, and auxiliary equipment, operate using predefined control logic without global optimization of their energy consumption. Although the system continuously monitors power usage through distributed power sensors, no autonomous coordination is performed to optimize the overall facility energy consumption. Consequently, devices operate independently, leading to higher overall energy usage while maintaining the required environmental conditions.

Goal/Target KPIs

The goal of sub-scenario 3 is to demonstrate how an AI-driven autonomic management system can improve the operational efficiency and sustainability of industrial indoor farming by replacing repetitive manual tasks with intelligent edge-based automation and by optimizing resource consumption through continuous autonomous planning.

The sub-scenario targets the following Key Performance Indicators (KPIs):

Improved efficiency ratio by 15 %: This improvement is achieved by reducing manual operational tasks associated with monitoring the hydroponic system. In particular, the autonomic management framework continuously collects and analyses measurements from the water quality monitoring system (e.g., pH and electrical conductivity (EC)), eliminating the need for operators to perform routine measurements and manually interpret the results.

Automated data collection, analysis, and decision support significantly reduce operator workload while enabling continuous supervision of the production process.

Improved resource usage and output by 10 %: This KPI is achieved through AI-based optimization of the facility's energy consumption. The MAPE-AI framework continuously monitors environmental conditions and the operational state of climate control equipment, including ventilation, lighting, and other controllable devices. Based on real-time sensor data and historical observations, the AI planner determines energy-efficient operating strategies that maintain optimal growing conditions while minimizing unnecessary power consumption. By coordinating multiple actuators through intelligent planning rather than independent rule-based control, the system reduces the overall energy usage of the indoor farming facility without compromising crop production performance.

Together, these improvements demonstrate how edge computing and AI-enabled autonomic management can increase operational efficiency, reduce labour requirements, optimize energy consumption, and contribute to more sustainable and scalable industrial indoor food production.

Approach/Method

During Year 3, RSA developed and validated an edge-native MAPE-AI (Monitor–Analyze–Plan-Execute with Artificial Intelligence) framework for autonomous environmental management of an indoor hydroponic farming system, depicted in Figure 8:. The approach combines Industrial IoT devices, edge computing, and Large Language Models (LLMs) to enable intelligent planning and control without relying on cloud infrastructure. Sensor data from climate and water quality monitoring devices are continuously collected, analysed against predefined target ranges, and used to construct structured system states for AI-based decision making.

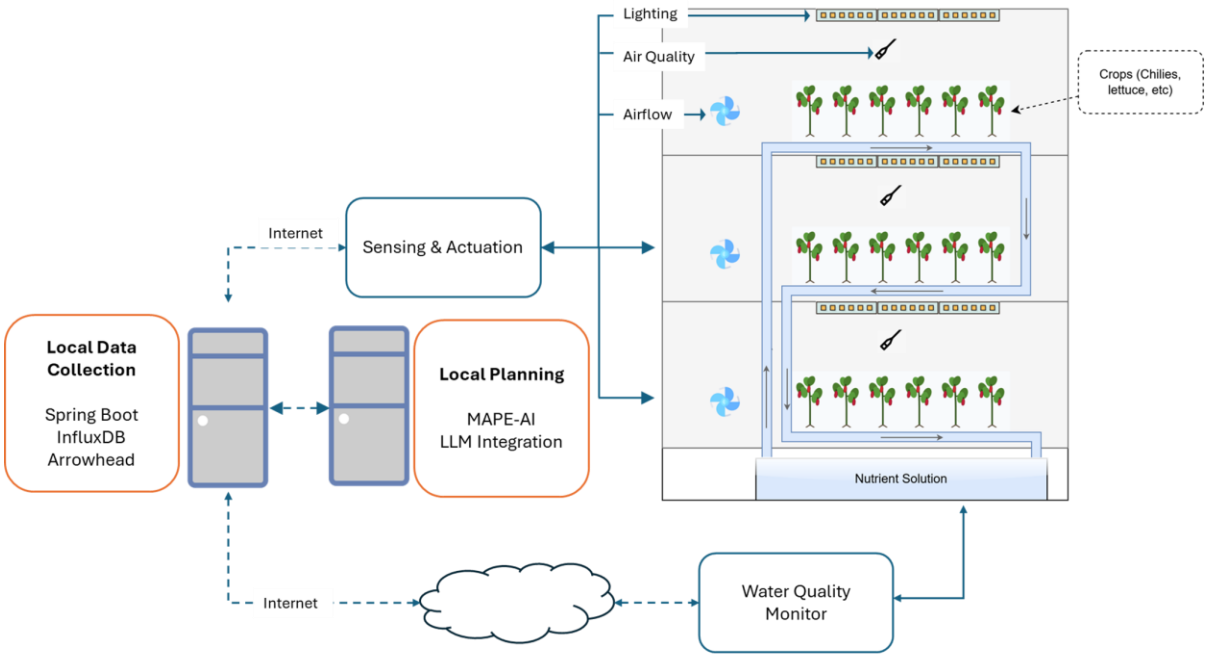

Figure 8: Autonomic Management of a Hydroponics System for Industrial Food Production

A locally deployed, INT4-quantized Qwen3-4B model, optimized using OpenVINO and executed on an industrial edge PC, serves as the planning component of the MAPE-AI loop, see Figure 9. The model receives structured prompts containing current sensor measurements, historical observations, actuator states, control rules, and target environmental conditions, and

generates explainable JSON-formatted control actions for ventilation and other environmental controls. The architecture was designed to be modular and extensible, allowing additional sensors, actuators, or AI models to be integrated with minimal changes to the management framework.

- The LLM states that there is a gradual increase in humidity and temperature
- The values are within ideal ranges but showing a trend toward less optimal conditions
- The LLM decides to take action by increasing the fan speed to 60%
- The LLM realizes that its action was successful in reducing humidity

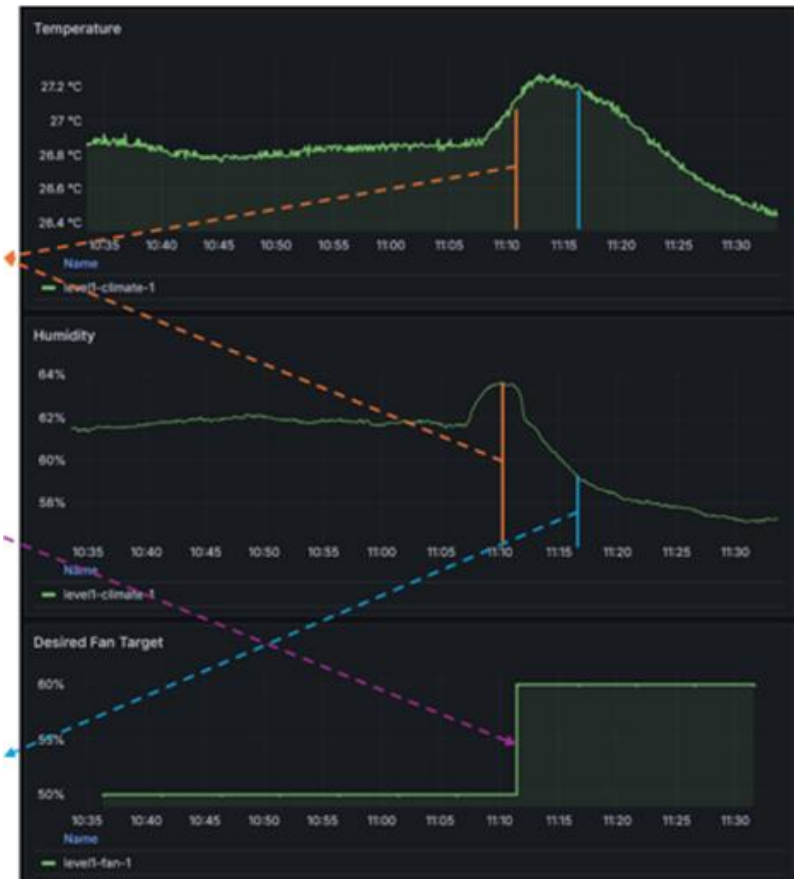

Figure 9: LLM in Action

The proposed solution was implemented on a physical hydroponic lettuce demonstrator (Figure 10) and evaluated through repeatable disturbance-response experiments. Multiple LLM prompt configurations were compared with conventional PI controllers to assess recovery time, inference latency, action consistency, and overall control performance. The evaluation demonstrated the feasibility of edge-based AI planning for soft real-time environmental control while providing a foundation for autonomous optimisation of resource usage and reduced operator intervention in industrial indoor farming systems.

Methodology and results

The Year 3 activities followed an experimental methodology to design, implement, and validate an AI-based autonomic management framework for indoor hydroponic farming. A distributed cyber-physical testbed was developed, integrating Industrial IoT devices for climate sensing, water quality monitoring, lighting control, and ventilation with an industrial edge computing platform. The system implements a MAPE-AI (Monitor–Analyze–Plan–Execute with Artificial Intelligence) control loop, where a locally deployed and quantized Qwen3-4B Large Language Model (LLM) performs the planning function based on structured sensor measurements, historical observations, actuator states, and predefined environmental targets.



Figure 10: The hydroponics demonstrator

The proposed methodology was based on the design and experimental validation of an edge-deployed MAPE-AI (Monitor–Analyze–Plan–Execute with Artificial Intelligence) framework for autonomous environmental control of a hydroponic lettuce cyber-physical system. The experimental platform consisted of an ESP8266-based sensing and actuation layer, an industrial Vecow SPC-5600 Edge PC, and an InfluxDB time-series database. The ESP8266 periodically acquired temperature and relative humidity measurements from an SHT31 sensor and controlled an exhaust ventilation fan through an HTTP interface. The edge PC executed the MAPE-AI control loop, stored historical measurements, and communicated actuator commands back to the microcontroller. The AI planning component used a locally deployed INT4-quantized Qwen3-4B Large Language Model, optimized with OpenVINO to enable inference on resource-constrained hardware.

The MAPE-AI control loop comprised four phases. During the Monitor phase, sensor measurements and actuator states were collected and stored as structured system snapshots. In the Analyze phase, the current environmental conditions were evaluated against predefined target ranges (22–25°C temperature and 50–70 % relative humidity) and recent trends were derived from historical data. During the Plan phase, the LLM received a structured prompt containing target ranges, sensor history, actuator states, control rules, and a required JSON output format, and generated explainable ventilation commands. Finally, the Execute phase translated the generated JSON output into actuator commands that were sent to the hydroponic ventilation system.

The evaluation employed a repeatable disturbance-response protocol. Each experiment started from a quasi-steady operating condition, where the exhaust fan remained unchanged for at least 2 minutes, temperature varied by less than 0.2°C, relative humidity by less than 2 percentage points, and both variables remained within the defined target ranges. A humidity disturbance was then introduced by spraying approximately 60 mL of room-temperature water into the growing chamber for 30 seconds, increasing the relative humidity to approximately 70–80 %. The controller was subsequently allowed to autonomously restore the environment to the predefined steady-state conditions.

Seven experimental scenarios were evaluated: two classical PI controller configurations (PI-M and PI-A), four LLM-based planning strategies (LLM-B, LLM-R, LLM-RH, and LLM-RHF), and a passive recovery scenario without active control. Each scenario was executed three times using

the same disturbance protocol. Performance was assessed using quantitative metrics including recovery time and LLM inference latency, while qualitative analysis evaluated the correctness of environmental state classification, consistency of generated control actions, adherence to the physical effect model, and the influence of prompt engineering strategies such as target-range repetition, extended sensor history, and fan-minimization instructions.

To evaluate the impact of AI-based energy optimization, power consumption data were collected from multiple power sensors installed throughout the indoor farming demonstrator. The dataset consisted of timestamped power measurements for each monitored device and was divided into two operational phases: baseline operation (without AI control) and AI-controlled operation, where the MAPE-AI framework autonomously optimized the operation of climate control devices.

For each timestamp, the power consumption of all monitored devices was aggregated to obtain the total system power. Energy consumption was then computed by integrating power over time and converting the results to watt-hours (Wh). Since the two experimental phases had different durations, energy consumption was normalized by operating time, resulting in normalized energy consumption (Wh/h), which served as the primary performance indicator. Additional performance metrics included mean, median, maximum, and 95th-percentile power consumption. A complementary per-device analysis was also conducted to evaluate how AI control affected the energy consumption of individual components while maintaining equivalent observation periods for both scenarios.

Results

The experimental evaluation compared two classical proportional-integral (PI) controllers, four LLM-based planning configurations, and a passive (no-control) baseline using a repeatable humidity disturbance protocol. Performance was assessed in terms of recovery time (time required to return the system to a quasi-steady state) and, for the LLM-based approaches, average inference time, see Figure 11:.

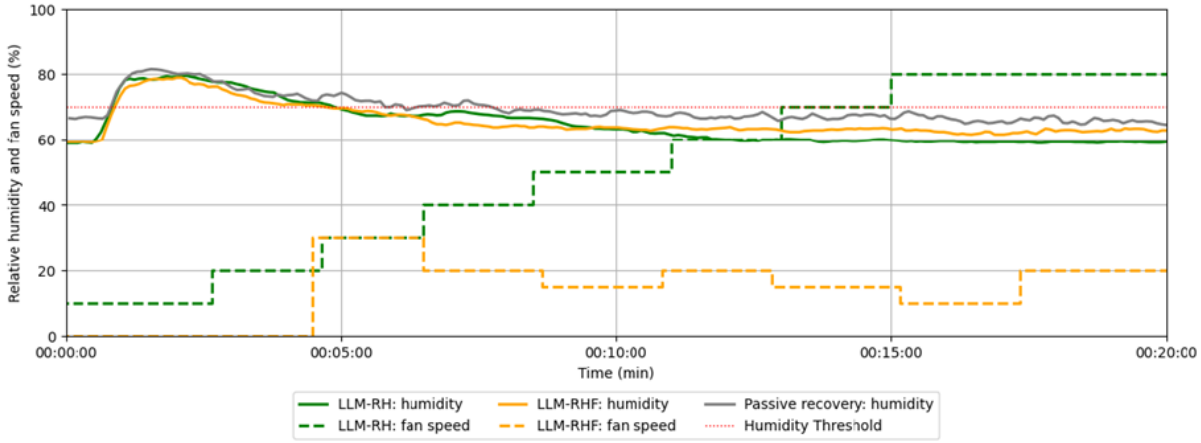

Figure 11: Humidity recovery and exhaust-fan response after a mist disturbance. Both relative humidity and fan speed are shown on the same percentage scale. Solid lines denote relative humidity, dashed lines denote fan-speed commands, and the dotted horizontal line indicates the relative-humidity threshold. Passive recovery is shown with the fan held at 0 %.

The baseline LLM configuration (LLM-B) delivered the most consistent AI-based performance, achieving recovery times between 5 min 46 s and 7 min 22 s, with an average inference time of 1 min 13 s. These recovery times were comparable to the classical PI controllers, where the moderate PI controller (PI-M) recovered in 6 min 05 s to 10 min 50 s, and the aggressive PI controller (PI-A) in 6 min 33 s to 9 min 52 s. In contrast, the passive recovery scenario, where

no active ventilation control was applied, required an average of 25 min 03 s, demonstrating the significant benefit of autonomous control.

Three additional LLM prompt variants were evaluated to investigate the impact of prompt engineering. The LLM-R configuration, which required explicit repetition of target climate ranges, reduced state-classification errors but increased average inference time to 1 min 38 s and produced more variable recovery times ranging from 5 min 11 s to 14 min 00 s. The LLM-RH configuration further extended the historical sensor context from two to eight snapshots, increasing inference time to 1 min 59 s while recovery times ranged from 10 min 52 s to 15 min 14 s. Although the longer history improved the planner's interpretation of environmental trends, it did not improve control performance. Finally, the LLM-RHF configuration introduced a fan-minimization strategy to reduce unnecessary ventilation. This resulted in the highest average inference time (2 min 05 s) and the largest variation in recovery time (9 min 31 s to 28 min 21 s). While the strategy successfully reduced actuator usage after recovery, one trial showed delayed corrective action because the planner interpreted a slowly decreasing humidity trend as recoverable without increasing fan speed.

Overall, the results demonstrate that a locally deployed LLM can achieve control performance comparable to conventional PI controllers while providing greater flexibility and explainability. They also highlight the trade-off between prompt complexity and operational performance: richer prompts improved reasoning quality and trend awareness but increased inference latency and, in some cases, reduced recovery consistency. The study confirms the feasibility of edge-based LLM planning for autonomous environmental management and establishes a scalable foundation for extending the system to additional sensors, actuators, and indoor farming control functions.

The evaluation demonstrated a clear improvement in energy efficiency when the MAPE-AI framework was enabled. The normalized system energy consumption decreased from 3.97 Wh/h in the baseline configuration to 3.04 Wh/h under AI control, corresponding to an absolute reduction of 0.93 Wh/h and a relative energy reduction of 23.4 %. This improvement significantly exceeds the project KPI target of 10 % improved resource usage and output, see Figure 12.

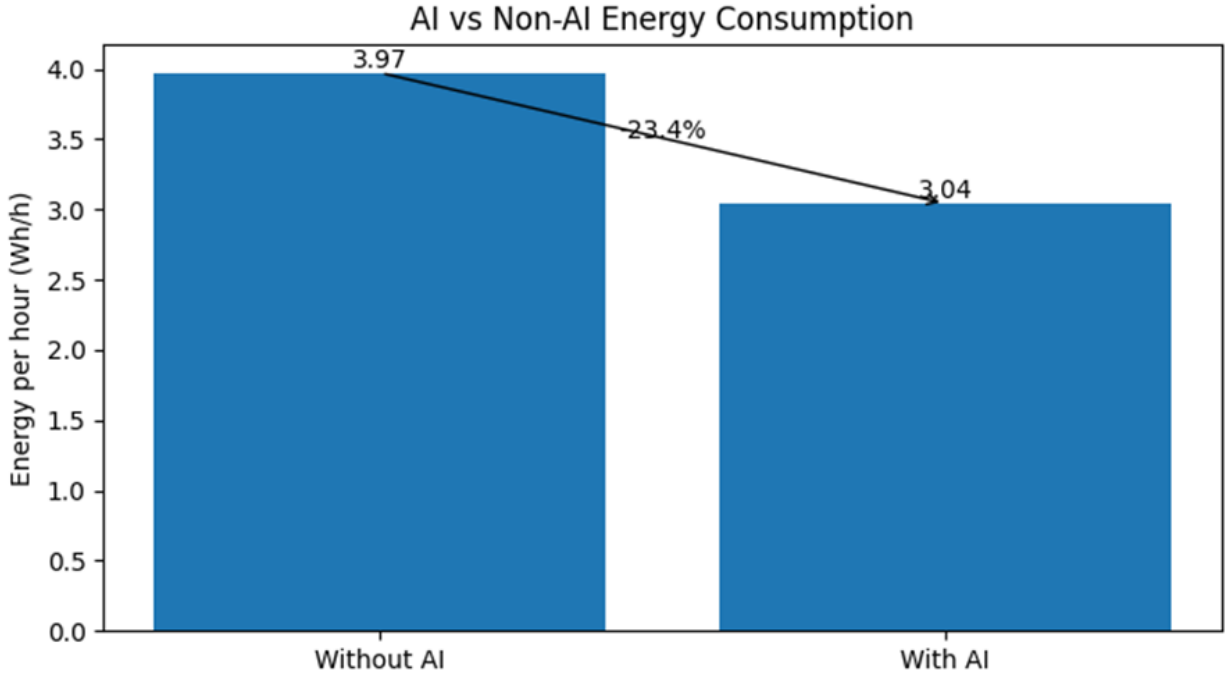

Figure 12: AI-based vs. non-AI-based Power consumption of the demonstrator

In addition to reducing total energy consumption, the statistical analysis showed lower average power consumption and more efficient system operation during AI-controlled execution. The per-device analysis further revealed that the energy savings resulted from coordinated optimization across multiple components rather than improvements in a single device. While some devices consumed less energy due to optimized control strategies, others exhibited slightly higher consumption to support targeted environmental regulation, such as increased ventilation when required. Overall, the AI planner redistributed energy consumption across the system to achieve more efficient operation while maintaining the desired environmental conditions, see Figure 13:.

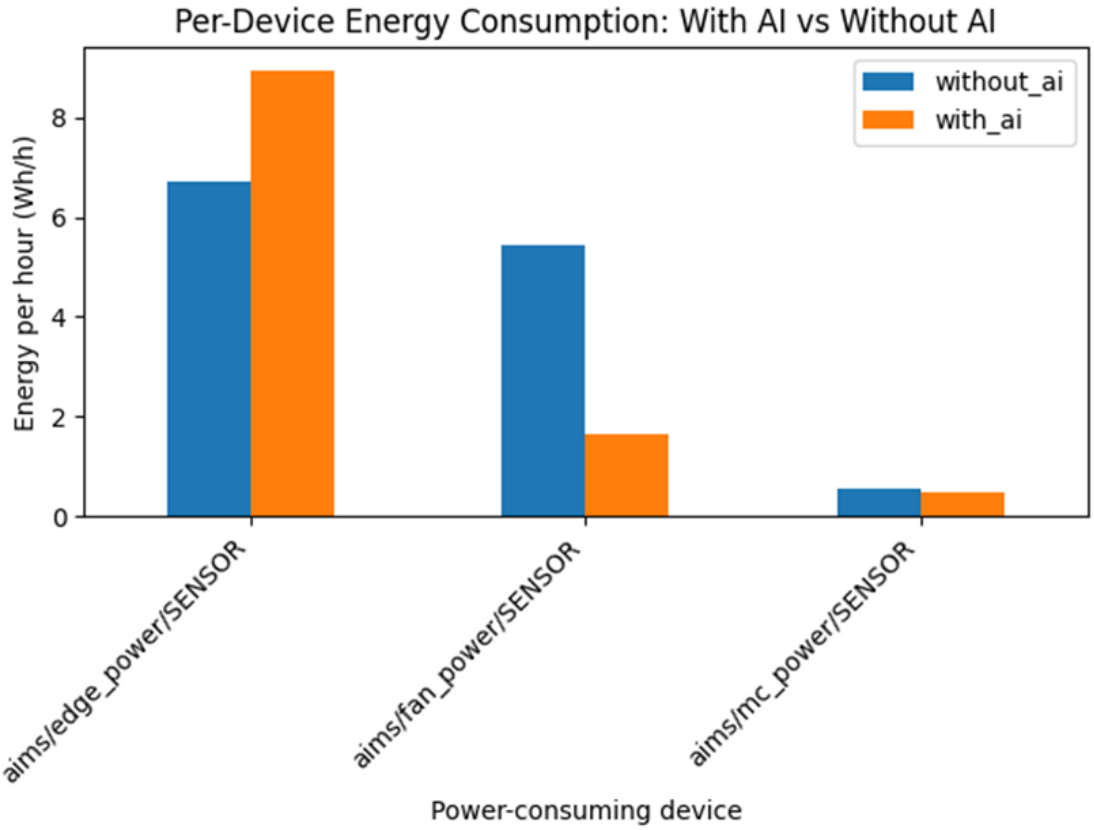

Figure 13: Per device energy consumption

These results demonstrate that the MAPE-AI autonomic management framework can optimize the operation of distributed indoor farming equipment, significantly reducing the facility's overall power consumption while maintaining autonomous environmental control, thereby contributing directly to the project's resource efficiency objectives.

The surveys conducted across WP7 provide a baseline for regulatory awareness in the consortium. The use-case-level survey in D7.1 (Table 1) showed that only one of ten use cases could identify specific formal AI standards applicable to its operations, while three reported no applicable regulations, standards, or guidelines at all. The subsequent persona-based questionnaire in D7.2 (N=59) confirmed this pattern at the individual level: 49% of respondents were unsure whether their organizations follow formal AI guidelines, and only 24% reported applying any standardised guidelines or checklists — of whom just 36% referenced official standards rather than internally developed ones. Compliance-related risks were also under-recognised, cited by only 37% of respondents compared with 71% for data-related and 68% for people-related challenges. Together, these findings indicated significant fragmentation in both awareness and practical adoption of the AI compliance landscape, reinforcing the need for the self-assessment tool developed in T7.4.

Sub-Scenario 4 – Support CRA requirements through systematic threat modelling and risk assessment

UC11 applies the ThreatGet methodology to an Industrial IoT stack for indoor food production and evaluates the results against the essential cybersecurity requirements of the EU Cyber Resilience Act (CRA). The analysis produced different threat scenarios and formulised multiple related attack steps across the full architecture.

Baseline and Problem

Indoor food production is a highly connected cyber-physical environment. Sensors, controllers, and remote servers jointly manage the critical agronomic variables, climate, irrigation, and nutrients, on which the crop depends. In this setting, a single compromised signal can cascade into concrete real-world damage: crop loss, food-safety incidents, or plant downtime.
At the start of the project, UC11 lacked a systematic, model-based view of the cybersecurity exposure of this Industrial IoT (IIoT) stack. Threat information was fragmented across components, and no explicit link existed between component-level threats and emerging European regulatory expectations, most notably the EU Cyber Resilience Act (CRA). This baseline gap, technically credible cybersecurity analysis on one side, regulatory language on the other, is the problem UC11 set out to close.

Goal and Target KPIs

The Year-1 objective was to establish a structured threat-analysis model for UC11 that captures the main architectural elements, assets, interfaces, and damage scenarios relevant to the real Indoor Food Production System. The intention was not only to perform an initial cybersecurity analysis, but also to build a model structure that reflects the actual UC11 architecture and can serve as a basis for systematic and reusable threat assessment. In this way, Year 1 focused on creating the analytical foundation needed to connect the technical system view of UC11 with compliance-relevant cybersecurity considerations under the CRA.

The concrete Year-1 KPI targets were:

- Complete a threat model of the UC11 architecture, including its critical components, assets, and interfaces.
- Cover the main threat categories relevant to Industrial IoT environments.
- Establish a traceable link between identified threats and the CRA essential cybersecurity requirements, in particular Article 13 and Annex I.
- Develop a reusable methodology that can be applied not only to UC11 but also to other AIMS5.0 use cases and, more broadly, to Industrial IoT deployments.

The success criterion was therefore not limited to technical completeness. It also included regulatory traceability, that is, the ability to show which CRA requirements are already supported by the analysis and where the remaining gaps still exist. Based on the current baseline assessment, ThreatGet covers approximately 20 % of the essential CRA requirements for products with digital elements. This indicates that ThreatGet already provides meaningful support for CRA-oriented threat and risk analysis, while also highlighting the need for complementary methods or extensions to address the remaining requirements that fall outside its present scope

Figure 14 presents the initial high-level ThreatGet model of UC11. It represents the main structural elements of the Indoor Farming Production System, divided into the Operational Environment and the Remote Monitoring and Controlling environment. The model includes the

Sensors Hub, Edge Device, Sensor ECU, Actuators Hub, Network Router, Data Server, and Data Store, all connected through the relevant communication paths. The purpose of this model in Year 1 was to establish a realistic architectural representation that contains the elements necessary for threat analysis and reflects the actual structure of UC11 as a distributed Industrial IoT system.

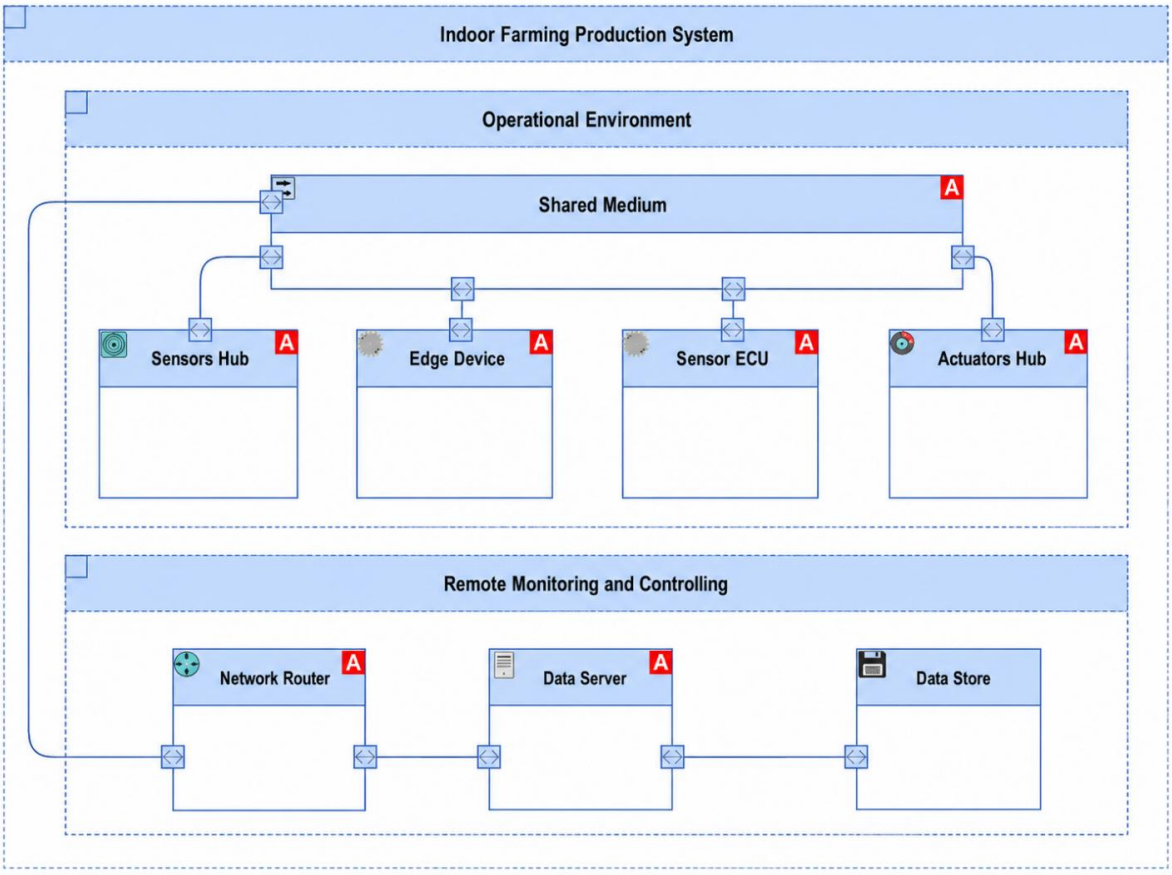

Figure 14: ThreatGet high-level architectural model of UC11

In ThreatGet, each component is associated with one or more assets. The "A" label shown on a component indicates that the element contains at least one defined asset. These assets are then analysed against relevant damage scenarios.

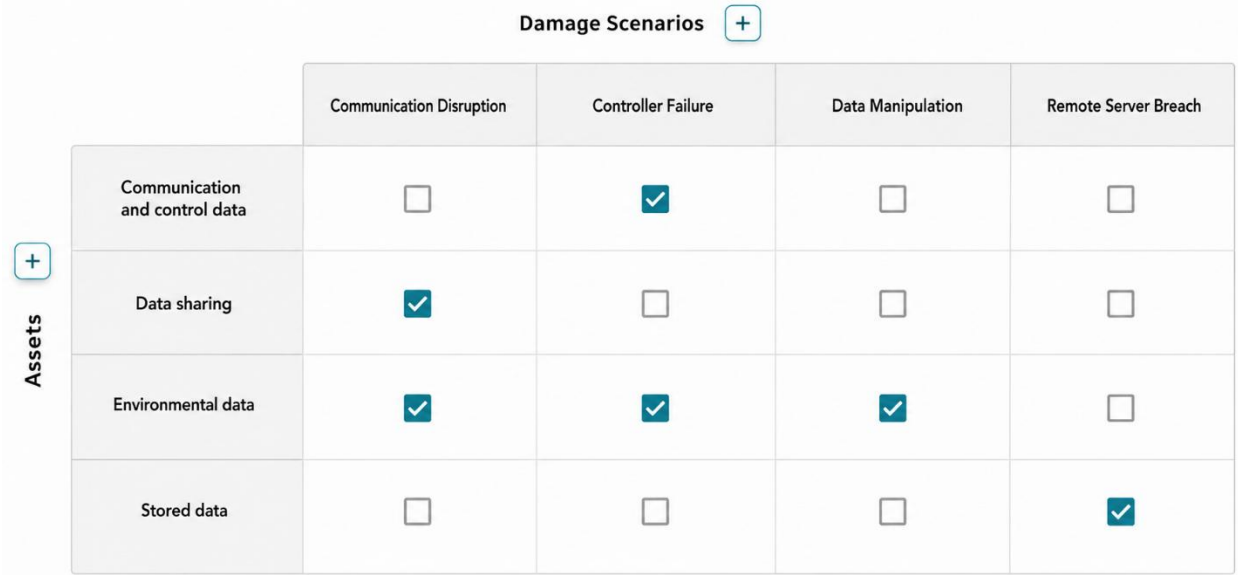

Figure 15 Asset-to-damage-scenario mapping for UC11 in ThreatGet

Figure 15 illustrates this principle. It shows how assets such as communication and control data, data sharing, environmental data, and stored data are mapped against damage scenarios including communication disruption, controller failure, data manipulation, and remote server breach. This asset-to-damage-scenario mapping is essential because it provides the analytical basis for identifying which parts of the system are exposed to which types of harm, and it supports the subsequent derivation of threats, risks, and mitigation needs.

Taken together, these Year-1 results establish the initial ThreatGet-based cybersecurity model for UC11 as a realistic and compliance-relevant starting point. They show how the actual system architecture can be translated into a structured threat-analysis model, how assets and damage scenarios can be linked systematically, and how this can support an evidence-based interpretation of selected CRA cybersecurity requirements in practice.

Approach and Method

By Year 3, the analysis approach was consolidated around the ThreatGet methodology developed at AIT and applied through a structured four-step workflow: Model, Analyse, Attack Trees, and Assess. This workflow supports a traceable progression from architectural representation to threat generation, attack-path reasoning, and risk evaluation. Each step produces outputs that feed directly into the next, creating a compact but systematic method for cybersecurity analysis in UC11.

As illustrated in Figure 16, the methodology begins with modelling the system architecture and its trust boundaries. This is followed by rule-based threat generation, which identifies relevant threats based on the architectural structure and defined assets.

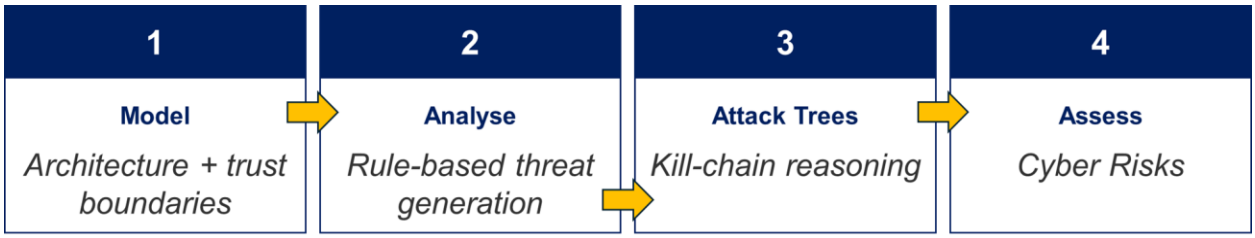

Figure 16 ThreatGet methodology applied in UC11

In the third step, these threats are analysed through attack trees and kill-chain reasoning in order to understand how an attacker could progress through the system. In the final step, the resulting threat scenarios are assessed in terms of risk, expressed as the combination of likelihood and impact. The outputs of this workflow are then interpreted against the CRA essential cybersecurity requirements.

Results

The ThreatGet-based analysis of UC11 follows a structured and traceable workflow that links architectural modelling, threat generation, attack-path reasoning, and risk assessment to CRA-oriented cybersecurity analysis.

A concrete example of this process is shown in Figure 17, which illustrates the attack tree for Environmental Sensor compromise. In this scenario, multiple entry points located at the Shared Medium Interface converge on the Sensor Data Processing Interface, enabling an Initial Access step. This is followed by Privilege Escalation and Defense Evasion against Sensor Data Processing, ultimately leading to compromise of the Environmental Sensor. The figure therefore demonstrates how ThreatGet captures a bottom-up kill chain across several architectural interfaces and components.

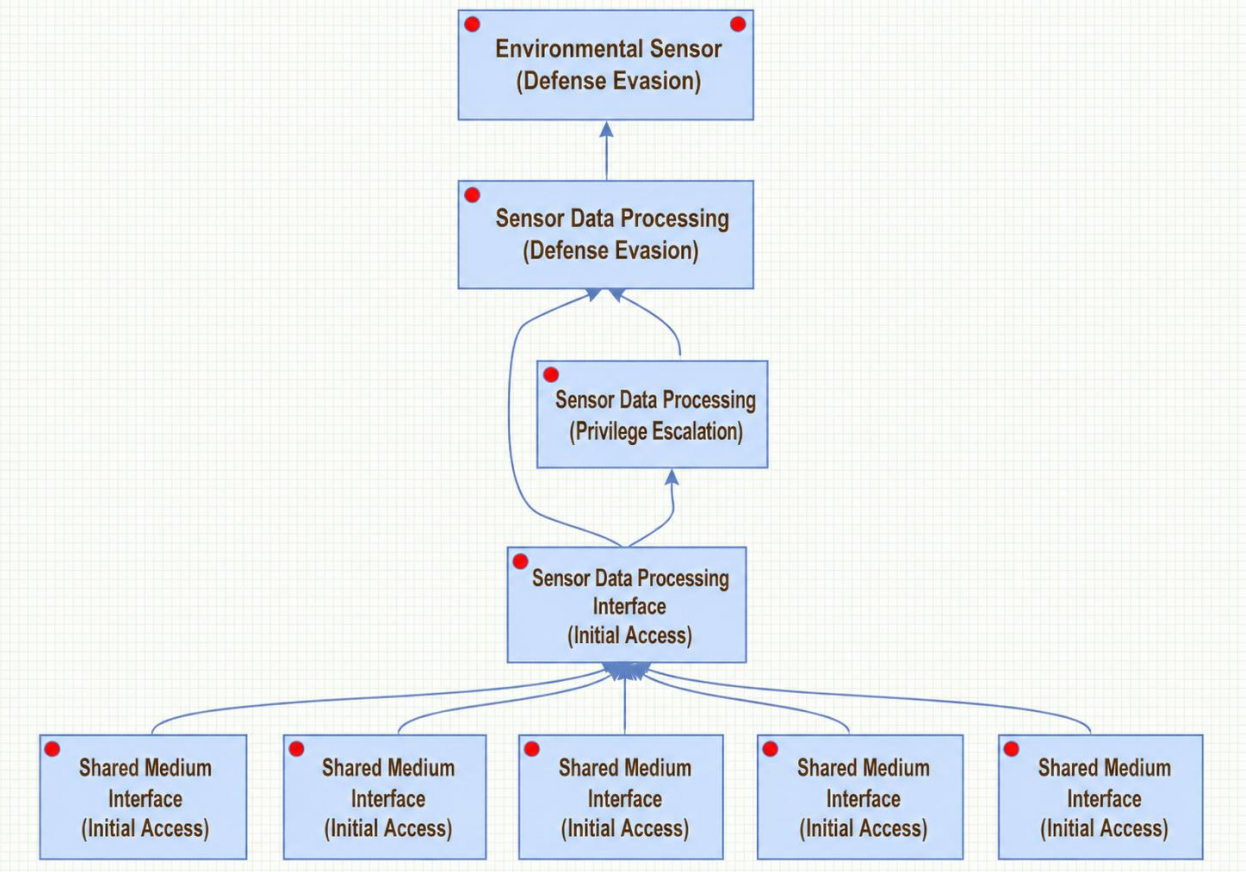

Figure 17: Attack Tree, Environmental Sensor compromise: Bottom-up kill chain from five Shared-Medium-Interface entry points to the Environmental Sensor via Sensor Data Processing

The resulting evidence generated by the tool is shown in Figure 18, where the ThreatGet analysis view provides a traceable overview of the workflow stages and the generated threat scenarios. In the current example, the interface displays the progression from Specifications and Definitions to Architecture, Analysis, and Summary, together with the list of generated threat scenarios. This consolidated analysis view supports systematic review of the identified threats and provides the practical basis for linking the technical analysis to CRA-oriented compliance considerations.

Application of the methodology to UC11 produced a substantial evidence base. In total, 109 threat scenarios and 115 attack steps were modelled end-to-end, covering all eight UC11 asset types. Across the generated scenarios, Tampering emerged as the dominant STRIDE category, followed by Information Disclosure.

On the regulatory side, ThreatGet was assessed against the 15 essential cybersecurity requirements derived from CRA Article 13 and Annex I. The overall result was 58 % coverage, distributed across three bands, as shown in Figure 19.

Status	ID	Asset	Rule Title	Target	Threat Types	Cybersecurity Properties	Risk ↓	Risk Treatment	
	30	Environmental data	Change Operating Mode	Sensors Hub	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	31	Environmental data	Change Operating Mode	Sensors Hub	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	71	Environmental data	Remote System Discovery	Sensors Hub	INFORMATION DISCLOSURE	Confidentiality, Integrity, Availability		NONE	
	79	Environmental data	Remote System Discovery	Sensors Hub	INFORMATION DISCLOSURE	Confidentiality, Integrity, Availability		NONE	
	80	Communication and control data	Remote System Discovery	Edge Device	INFORMATION DISCLOSURE	Confidentiality, Integrity, Availability		NONE	
	53	Environmental data	Command-Line Interface	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	54	Communication and control data	Command-Line Interface	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	56	Environmental data	Command-Line Interface	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	57	Communication and control data	Command-Line Interface	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	59	Environmental data	Modify Controller Tasking	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	60	Communication and control data	Modify Controller Tasking	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	62	Environmental data	Modify Controller Tasking	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	63	Communication and control data	Modify Controller Tasking	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	34	Environmental data	Change Operating Mode	Sensors Hub	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	35	Environmental data	Change Operating Mode	Sensors Hub	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	36	Environmental data	Change Operating Mode	Sensors Hub	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	37	Environmental data	Change Operating Mode	Sensors Hub	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	38	Communication and control data	Change Operating Mode	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	
	39	Communication and control data	Change Operating Mode	Edge Device	TAMPERING	Confidentiality, Integrity, Availability		NONE	

Figure 18 ThreatGet Analysis View

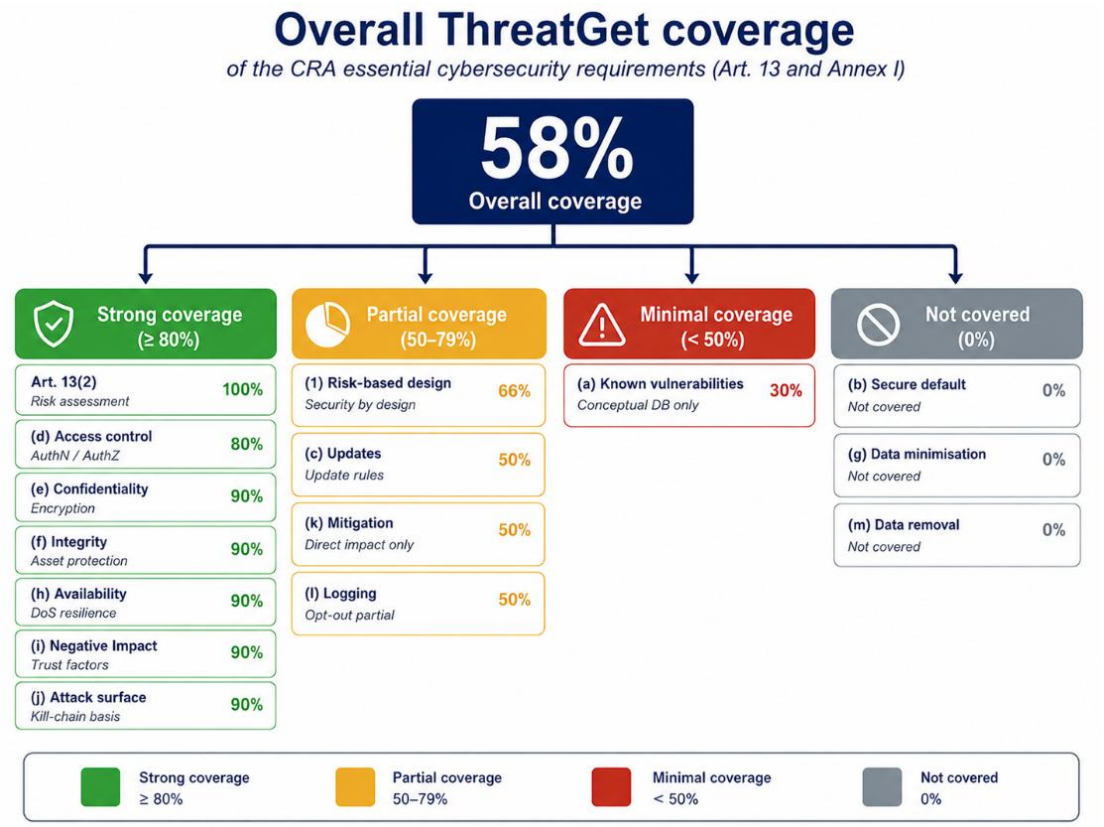

Figure 19: Overall ThreatGet coverage of the CRA essential cybersecurity requirements

Strong coverage was achieved for seven requirements: Risk Assessment, Access Control (Authentication and Authorisation), Confidentiality (Encryption), Integrity (Asset protection), Availability (DoS resilience), Negative Impact (Trust factors), and Attack Surface (Kill chain basis). Partial coverage was found for four requirements: Risk-Based Design (Security by design), Updates (Update rules), Mitigation (Direct impact only), and Logging (Opt-out partial).

Minimal or no coverage was observed for the remaining four requirements: Known Vulnerabilities (Conceptual DB only), Secure Default Configuration (Not covered), Data Minimisation (Not covered), and Data Removal (Not covered).

A key strength of the analysis is its traceability from technical findings to specific CRA requirements. Attack-Surface Reduction informs requirement: Attack Surface with 90 % coverage; Access and Interface Security with 80 %; Confidentiality, Integrity, Availability, Negative Impact, and Attack surface each achieve 90 % against their corresponding CRA requirements; and the Risk-Driven Analysis Workflow fully supports Article 13(2) with 100 % coverage. At the same time, Updates, Mitigation, and Logging remain only partially covered at 50 %, Known Vulnerabilities is covered conceptually at 30 %, and Secure Default Configuration, Data Minimisation, and Data Removal are not addressed in the current ThreatGet scope.

Overall, the UC11 results demonstrate that the ThreatGet methodology generates an evidence base that is both technically meaningful and regulatorily interpretable, thereby supporting CRA-aligned threat and risk analysis in Industrial IoT systems.

2.3 Use Case 14 (Task 6.2.1/6.2.2): Speeding-up the Benefit of Automation by Virtual Commissioning through Digital Twins

Lead: FMX; **Partners:** IFD, HOST, TUD, IFAT

2.3.1 Use Case Characteristics

Problem/Baseline

FMX offers a range of material handling automation components including corresponding local control software and higher-level control systems which are functional and proven but lack integrated AI capabilities. Commissioning still follows a traditional, linear workflow: after product development and manufacturing, commissioning happens first in-house, and only afterwards on the customer site. That linear workflow is long and often causes delays.

Goal/Target KPIs

UC14 (including related tasks 6.2.1 and 6.2.2, see Figure 20) aimed to solve the two mentioned shortcomings: missing AI functionality in automation components and control systems and the absence of virtual commissioning.

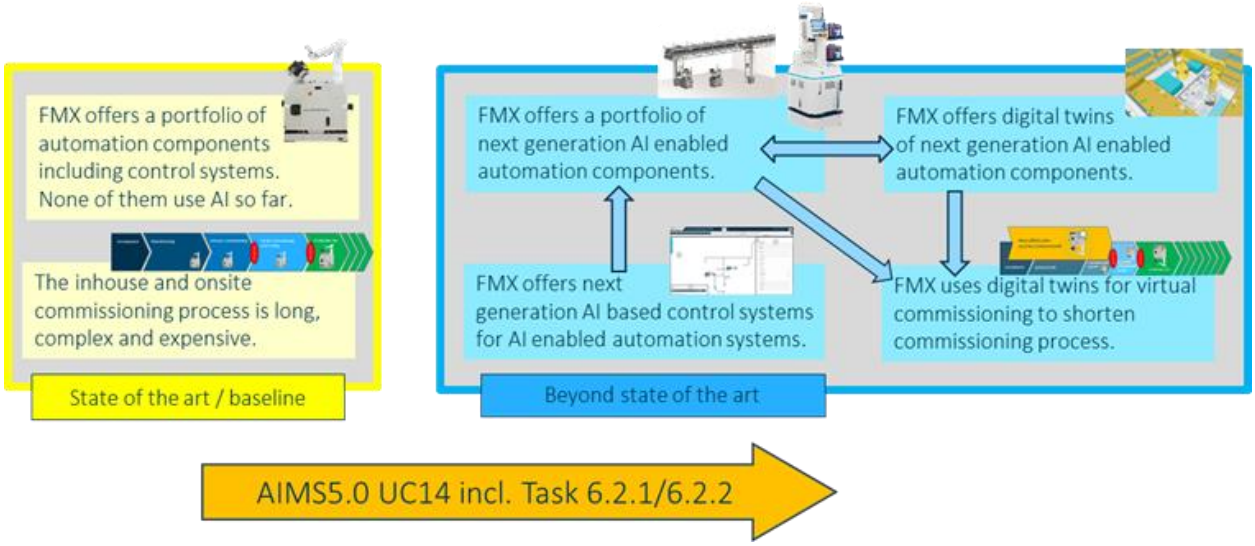

Figure 20: Task relation

For target KPIs see table in Achievement of KPIs section below.

2.3.2 Approach/Method

FMX and partners designed and developed a completely new generation of AI-enabled intralogistics automation components for material storage, transport and handling based on detailed requirement specifications jointly developed between FMX and microelectronics manufacturers. In addition to the AI enablement, FMX also tested the application of an AI-based controller for AMRs using AI-based travel time prediction. For this purpose, FMX successfully applied a supervised learning approach based on historical log files combined with neural network models to improve prediction accuracy and operational efficiency.

For the Digital Twin activities, FMX tested the demonstrator, which was developed in Y2 using representative simulation and validation scenarios. Furthermore, FMX investigated how Virtual Commissioning could be integrated into the existing software development and testing procedures, including Continuous Integration and Continuous Deployment (CI/CD) cycles. Applying Digital Twins to support and optimize commissioning activities led to the shifting of significant parts of the testing and validation processes from on-site commissioning to virtual commissioning.

2.3.3 Results

Fulfillment of Requirements and Specifications

All material handling automation systems (e.g. HERO-SCOUT mobile robot, FlatStocker3.0, Purge Monitoring FOUP, etc.) fulfill the requirement specifications jointly developed with customers and have been successfully tested either at FMX in Dresden or at customer sites. Also, the AI-based controller was successfully tested against representative example scenarios and demonstrated improved accuracy in travel time prediction compared to conventional approaches.

The Digital Twin was used for initial software tests to validate the general workflow and applicability of virtual commissioning methods. In particular, the demonstrator was successfully evaluated against requirements defined in Y1 such as:

- Applicability and accuracy, ensuring sufficient correspondence between simulation behavior and real-world application behavior;
- Adjustability, enabling integration with evolving project requirements and system modifications;
- Integration capability, ensuring compatibility with existing software development and commissioning processes.

Also, FMX with partners successfully applied the first digital twins for commissioning. The results confirmed the applicability of virtual commissioning approaches for semiconductor intralogistics environments and demonstrated clear benefits regarding development efficiency and project execution (see KPIs in table below).

Achievement of KPIs

These are the KPIs of UC14 with their original targets and the current status:

- **The time required for in-house setup, commissioning, and testing** could be significantly reduced, measured as the time from final assembly of the automation equipment in-house until it is ready for shipment. For the TWC digital twin, a reduction of approximately 25 % was achieved despite its demonstrator status. For the LUNA robot cell, a reduction of approximately 33 % compared to the original project plan is expected.
- **The time needed for on-site setup, commissioning, and testing** could also be substantially reduced, measured as the time from final installation of the automation equipment on-site until it is ready for production. Initial customer deployments of the TWC demonstrated a reduction of approximately 20 %, despite both the automation system and the digital twin still being at demonstrator stage. For the LUNA robot cell, a reduction of approximately 33 % is expected.
- **The scheduled downtime of customer process equipment** required for setup and testing activities is expected to be reduced by approximately 50 %, measured as the total time during which the process tool is unavailable for production due to testing of the automation equipment. This benefit will be demonstrated during the ongoing on-site commissioning of the LUNA robot cell using its digital twin.
- As a consequence of the reduced in-house and on-site commissioning efforts, the **handover of the automation system for productive operation** can be significantly accelerated, measured as the time from project start (purchase order) until readiness for production. Current results indicate a reduction of approximately six weeks compared to conventional project execution, while the target of handing over the system eight weeks earlier is expected to be achieved by the end of the AIMS5.0 project.
- In addition, the project successfully delivered several **next-generation AI-enabled solutions**, all of which were successfully validated against their respective requirement specifications. The FlatStocker 3.0 AI-enabled storage system was successfully tested, the HERO-SCOUT mobile robot was successfully evaluated at Infineon, and a demonstrator implementation of an AI-based fleet management function for travel-time estimation using machine learning was successfully developed and validated.

Impact on Electronics Components Supply Chain

Strengthening the productivity and competitiveness of semiconductor electronics manufacturers such as Infineon leads directly to a strengthening of supply chains for electronic components.

Impact on Sustainability and Decarbonization

Intralogistics automation has recognized positive impacts on sustainability and decarbonization, as it significantly increases the efficiency of energy-intensive semiconductor fabs.

2.3.4 Learnings and Outlooks towards Project End and Beyond

The planned goals of the project are already reached or will be achieved by end of the project.

2.4 Use Case 18 (Task 6.1.5): Using semantic Data Model as base for AI optimized economic & ecologic/societal supply chain improvements

Lead: IFAG; **Partners:** FUH, NXP, IFD

2.4.1 Use Case Characteristics

Use Case 18 (UC18) addresses one of the key challenges facing the semiconductor industry: improving **transparency** regarding the economic and environmental performance of manufacturing and supply chain operations. The increasing demand for sustainable semiconductor production requires organizations to understand the impact of production decisions on both operational efficiency and carbon emissions. UC18 therefore focuses on creating a semantic framework that enables data-driven optimization of sustainability and business performance across the **semiconductor value chain**.

The use case leverages Semantic Web technologies, industry standards, and the **Digital Reference (DR) ecosystem**, as well as the **WP4's Open Access Platform (OAP)**. The core objective is to model semiconductor manufacturing, supply chain, operating curve (OC), and ecological operating curve (eOC) concepts in a machine-readable manner to support intelligent decision-making processes. The semantic framework enables the representation and analysis of relationships among production efficiency, resource utilization, customer demand, and environmental impact.

At the business level, the use case aims to deliver improvements in three strategic areas:

- **Economic and ecological optimization** through the integration of operating curve and ecological operating curve concepts.
- **Improved Available-to-Promise (ATP)** performance by incorporating production-related concepts such as Statistical Process Control (SPC), Yield, and Overall Equipment Effectiveness (OEE).
- **Enhanced transparency of CO_2 emissions** and sustainability impacts across manufacturing and supply chain operations.

Beyond its sustainability objectives, UC18 also contributes to knowledge retention, workforce upskilling, and the broader adoption of semantic technologies within the semiconductor domain. The use case therefore combines **environmental**, **operational**, and **human-capital** perspectives into a unified semantic modeling initiative.

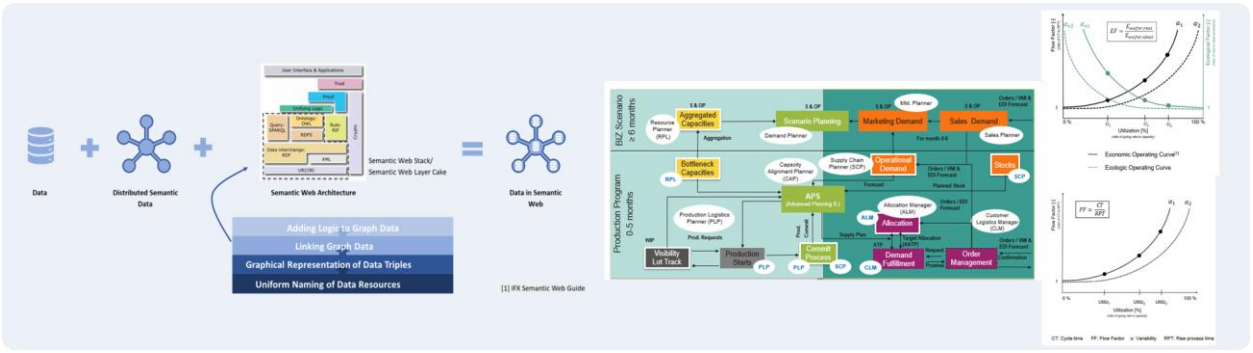

Figure 21: Illustration of why semantically connected data can support ATP and OC modeling.

2.4.2 Approach/Method

The UC18 methodology combines domain expertise, semantic modeling, and knowledge graph technologies to create a reusable and extensible representation of **semiconductor manufacturing and supply chain knowledge**.

The starting point was the development of a DrOWLing model derived from the use case requirements and IFAG's simulation work. The semantic model is centered around three key entities — **Customer**, **Supplier**, and **Available to Promise (ATP)** — and includes additional concepts representing the inputs, outputs, and decision variables relevant to the operating curve and ecological operating curve models. Existing Digital Reference concepts were reused whenever applicable to maximize interoperability and alignment with the broader DR ecosystem.

The semantic engineering process followed a structured workflow consisting of eight phases:

1. **Ground Truth Definition** – Domain modelers created a DrOWLing representation of the concepts to be semantically formalized.
2. **Expert Review** – Industry experts and stakeholders validated the conceptual representations and relationships.
3. **Ontology Generation** – The DrOWLing model was transformed into a formal ontology and validated using semantic reasoning tools.
4. **Data Validation** – Example datasets were mapped against the ontology to verify correct interpretation and linkage of information.
5. **Information Inference** – Reasoning mechanisms were applied to validate the ontology's capability to derive meaningful conclusions.
6. **Formal Standardization** – Entities and concepts were aligned with relevant industry standards.
7. **User Feedback and Refinement** – Stakeholder feedback was incorporated to improve usability and accuracy.
8. **Continuous Improvement** – An iterative update cycle ensured long-term relevance and quality of the semantic model.

The resulting ontology was implemented using established **Semantic Web standards**, including OWL 2 DL for formal representation, Turtle (.ttl) for serialization, SKOS for annotation, and SPARQL for querying. The final semantic model consists of 34 classes, 30 object properties, and 29 data properties, while maintaining strong integration with several **Digital Reference domains**, including Demand Fulfilment, Supply Chain Planning, Order Management, Planning, Generic Data Model, and Semiconductor Operations.

A significant innovation of UC18 is the integration of semantic technologies with sustainability modeling. Unlike traditional operating curve approaches, the semantic representation enables

improved interoperability, knowledge reuse, explainability, and future AI-supported optimization capabilities. Furthermore, knowledge graphs and semantic standards provide a shared understanding among experts, reducing knowledge fragmentation and facilitating collaboration across organizational boundaries.

2.4.3 Results

UC18 successfully demonstrated the value of semantic technologies for representing and analyzing sustainability and performance factors within semiconductor manufacturing and supply chains.

One major outcome was the development of expert-validated semantic models for **Economic and Ecological Alphas**, enabling improved visibility of CO_2 emissions relative to production performance. The implementation contributed seven additional classes to the Digital Reference ecosystem and established a reusable foundation for sustainability-oriented decision support. The use case also extended ATP-related modeling capabilities by introducing concepts such as SPC, Yield, and OEE into the semantic framework. These additions are expected to contribute to **ATP improvements** in the low single-digit percentage range while simultaneously improving customer service performance. The ATP semantic extension contributed eight new classes to the Digital Reference.

Another important achievement was the development of a dedicated Carbon Handprint semantic model. This model enables organizations to evaluate and communicate the positive environmental contribution of semiconductor technologies and associated applications. The **Carbon Handprint ontology** introduced 34 additional semantic classes and was successfully implemented as an OAP use case. The model supports the broader objective of reducing CO_2 emissions and advancing carbon-neutral manufacturing strategies.

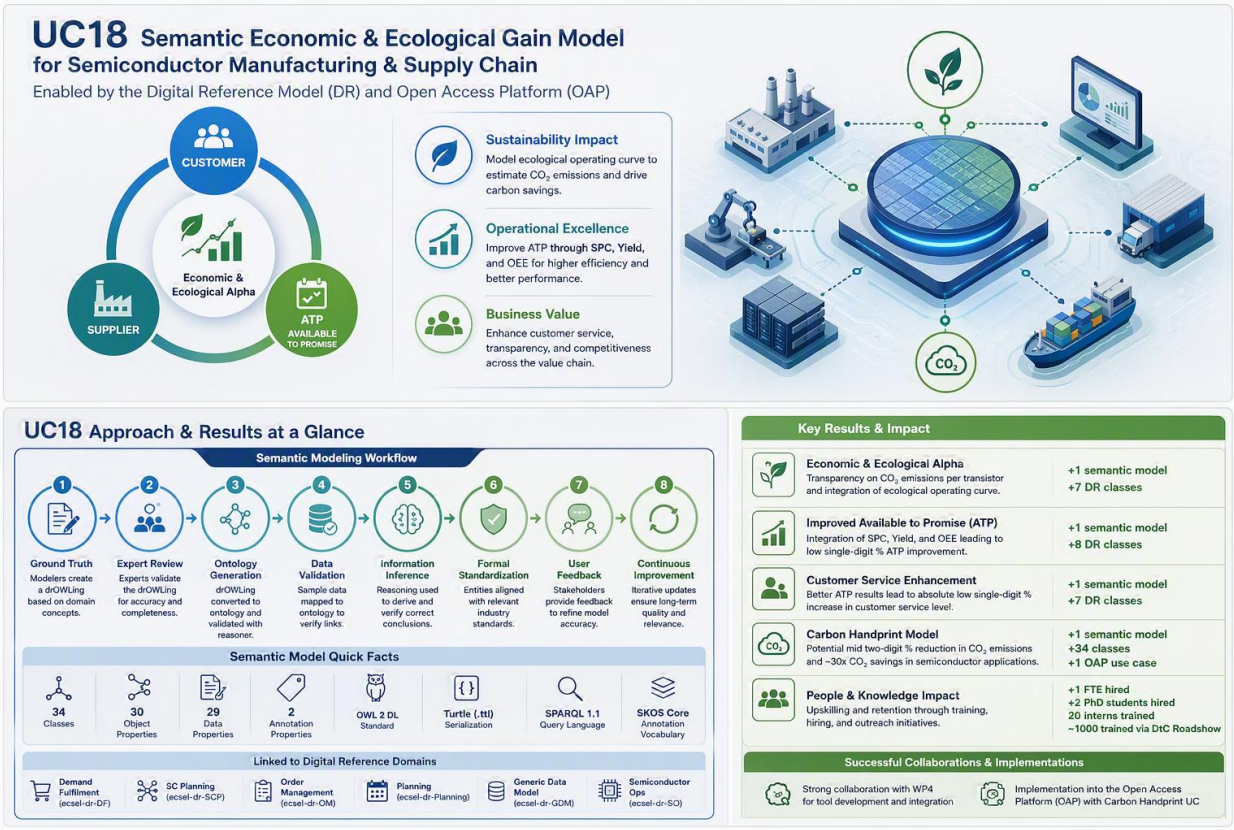

Figure 22: UC18 overview, workflow and achievements.

From a Digital Reference perspective, UC18 delivered multiple **expert-validated semantic models**, extended existing reference domains, and strengthened integration between sustainability, production, and supply chain concepts (Figure 22). The collaborative work with WP4 further enabled practical implementation within the OAP environment, transforming conceptual semantic models into operational assets accessible to the project ecosystem.

The use case also generated significant impact in terms of knowledge transfer and **workforce development**. One full-time employee and two PhD students were trained and hired, approximately twenty intern modelers received training, and around one thousand individuals were reached through Digital Twin-related dissemination and roadshow activities.

2.4.4 Learnings and Outlooks towards Project End and Beyond

The semantic foundation established by UC18 provides a scalable basis for future **sustainability-oriented digitalization initiatives** within the semiconductor industry. The ontology-based representation of operating curve and ecological operating curve concepts can be further enriched with additional manufacturing, logistics, and environmental datasets, enabling more accurate sustainability assessments and optimization scenarios.

A particularly promising direction is the application of **semantic-powered Artificial Intelligence** for optimizing ecological operating curves and supporting transparent sustainability decisions. By combining knowledge graphs, reasoning techniques, and AI methods, future solutions may deliver explainable recommendations for balancing productivity, resource utilization, customer service, and environmental impact.

The use case also creates opportunities for stronger cross-project integration. The developed DrOWLing and ontology can be expanded to incorporate concepts from related AIMS5.0 use cases, particularly **UC1** (Green AI Lifecycle Management Approaches), **UC10** (AI-driven Manufacturing Execution Systems), and **UC12** (Human-Centered AI for Semiconductor Manufacturing Networks). Such extensions would increase interoperability across use cases and further strengthen the Digital Reference as a common semantic backbone for the semiconductor domain.

In the longer term, UC18 demonstrates how semantic technologies can transform fragmented domain knowledge into **reusable digital assets** that support sustainability, operational excellence, and collaborative innovation. Through continued integration into the Digital Reference ecosystem and Open Access Platform, the use case establishes a pathway toward more transparent, data-driven, and environmentally responsible semiconductor manufacturing and supply chain management.

3 Conclusion

Deliverable D6.6 presents a consolidated summary of the Use Cases developed within Work Package 6 (WP6) that can be made publicly available, as they do not contain confidential information requiring special protection. This distinguishes it from Deliverable D6.4, where the nature of the work results precluded public disclosure. Out of the 12 Use Cases addressed in WP6, three were identified whose content and outcomes could be documented and shared with the public, thereby making the results of the AIMS5.0 project accessible to the wider industry. The overarching goal in doing so is to contribute to improvements in productivity, sustainability, and the decarbonization of global industry.

The structure of this Deliverable follows the same framework as D6.4, providing a comprehensive summary of the results achieved across the three selected Use Cases over the course of three years. Each Use Case is introduced with a description of its objectives and goals, followed by an outline of its specific characteristics. References are made to the targets defined in the original project proposal, and the degree to which these targets were achieved is discussed in detail. The requirements and specifications developed during the first project year are also addressed. Furthermore, a description of the current baseline is provided for each Use Case, alongside a thorough account of the methods and technologies employed in their development and implementation. The technologies utilized span the domains of factory automation, digitalization, machine learning, and artificial intelligence — all of which represent core focus areas of the AIMS5.0 project. All obtained results are presented and mapped against the originally defined objectives.

Reviewing the key outcomes, it is noteworthy that in several cases the results achieved not only met but exceeded the targets defined in the original project proposal. The Deliverable opens with UC10 (MES Airbus), which forms an important foundation for Use Cases UC12, UC13, UC14, UC15, UC17, and UC20. With regard to AI-based MES capability, an improvement of more than 10 % was achieved, while new functionalities such as self-adjusting manufacturing processes demonstrated improvements of more than 15 %. Within UC11 (Indoor Food Production), significant savings in energy consumption and operational effort were realized in plant production through the application of machine learning and AI. In the context of UC14, virtual commissioning — enabled by the use of digital twins — led to substantial reductions in time and effort during the commissioning phase. Finally, the results from UC18 and UC19 demonstrated that the application of Semantic Web ontologies enabled meaningful reductions in CO_2 footprint, contributing to the project's sustainability and decarbonization objectives.

Appendices

Appendix A - Abbreviations

Table 1: Abbreviations

Abbreviation	Meaning
AI	Artificial Intelligence
AIMS5.0	Artificial Intelligence in Manufacturing leading to Sustainability and Industry5.0
ANKO	Equipment Control (German: Anlagenkontrolle)
AMC	Airborne Molecular Contamination
AMHS	Automated Material Handling Systems
ATP	Available To Promise
AuDi-EC	Automated Digitized Equipment Control
CeRULEo	Comprehensive utilities for Remaining Useful Life Estimation methods
CI/CD	Continuous Integration and Continuous Deployment
CMOS	Complementary Metal-Oxide-Semiconductor
CO_2	Carbon Dioxide
CRA	Cyber Resilience Act
D	Deliverable
DB	Database
DES	Discrete Event Simulation
DR	Digital Reference
DRO	Digital Reference Ontology
DRL	Deep Reinforcement Learning
EC	electrical conductivity
ECS	Electronic Components and Systems
eOC	ecological Operating Curve
FOUP	Front Opening Unit Pod
FTIR	Fourier-transform infrared spectroscopy
GUI	Graphic User Interface
HCl	Hydrogen Chloride
HF	Hydrogen Fluoride
HIM	holistic information model
I5.0	Industry 5.0
IC	Integrated Circuit
IEPE	Integrated Electronics Piezo-Electric
IFAG WAR	Infineon AG Warstein
IIoT	Industrial Internet of Things
KPI	Key Performance Indicator
LLM	Large Language Model
MAPE-K	Measure, Analyze, Plan, Execute, Knowledge
MES	Manufacturing Execution System
MEMS	Micro-Electro-Mechanical System
ML	Machine Learning
MLOps	Machine Learning Operations
MPC	Multi Party Computation
M	Month
NH_3	Ammonia
O	Objective
OC	Operating Curve
OCM	Operating Curve Management
OAP	Open Access Platform
OEE	Overall Equipment Effectiveness
OEM	Original Equipment Manufacturer
OHT	Over-Head Transportation

PI	proportional-integral
R	Requirement
RF	Radio Frequency
RAMII	Real-time Adoptive Manufacturing Information Infrastructure
S	Specification
SiC	Silicium-Carbid
SPC	Statistical Process Control
SPDT	Single Pole Double Throw
SHAP	SHapley Additive exPlanations
UC	Use Case
WBG	Wide Band Gap
WIP	Work In Progress
WP	Work Package
XAI	eXplainable Artificial Intelligence
Y	Year