SUNRISE

Strategies and Technologies for **Un**ited and **R**esilient Critical **I**nfrastructures and Vital **S**ervices in Pandemic-Stricken **E**urope

D3.4 Reference designs and best practices

Document Identification			
Status	Final	**Due Date**	31/07/2025
Version	1.0	**Submission Date**	31/07/2025

Related WP	WP3	**Document Reference**	D3.4
Related Deliverable(s)	D3.1, D3.2, D3.3, D4.6, D5.6, D6.6, D7.6, D10.1, D10.2	**Dissemination Level (*)**	PU
Lead Participant	INT	**Lead Author**	George Tsakirakis
Contributors	XLB, ATOS, SQD, IMA	**Reviewers**	Javier Andión and Jose Del Alamo (UPM)
			Tasos Gkamaris and Romeo Bratskas (SKYLD)

Keywords:
Reference designs, Best Practices, Key Features, Design Objectives, Critical Infrastructure Protection, AI-Enhanced Security, Cyber-Physical Security, Operational Resilience

This document is issued within the frame and for the purpose of the SUNRISE project. This project has received funding from the European Union's Horizon Europe Programme under Grant Agreement No.101073821. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the European Commission.

(*) Dissemination level: **(PU)** Public, fully open, e.g. web (Deliverables flagged as public will be automatically published in CORDIS project's page). **(SEN)** Sensitive, limited under the conditions of the Grant Agreement. **(Classified EU-R)** EU RESTRICTED under the Commission Decision No2015/444. **(Classified EU-C)** EU CONFIDENTIAL under the Commission Decision No2015/444. **(Classified EU-S)** EU SECRET under the Commission Decision No2015/444.

Document Information

List of Contributors	
Name	Partner
Nikos Avgeros, Stavros Theocharis	SQD
Jan Orlicky	IMA
Susana Gonzalez Zarzosa	ATOS
Daniel Vladušič	XLB

Document History			
Version	Date	Change editors	Changes
0.1	13/05/2025	George Tsakirakis (INT)	First ToC draft
0.2	18/06/2025	George Tsakirakis (INT)	XLAB contributions and merges
0.3	24/06/2025	George Tsakirakis (INT)	Consolidated SQD input
0.4	26/06/2025	George Tsakirakis (INT)	Consolidated ATOS input
0.5	02/07/2025	George Tsakirakis (INT)	Consolidated IMA input
0.6	06/07/2025	George Tsakirakis (INT)	INT – Updates, formatting fixes
0.7	10/07/2025	George Tsakirakis (INT)	INT revision and comments for partners
0.8	22/07/2025	George Tsakirakis (INT)	Consolidated 2nd round of inputs
0.9	25/07/2025	Jose Del Alamo & Javier Andion (UPM)	Internal review
0.9b	26/07/2025	Tasos Gkamaris & Romeo Bratskas (SKYLD)	Internal review
0.9c	30/07/2025	George Tsakirakis (INT)	Addressed review comments
0.9d	31/07/2025	Juan Alonso (ATS)	Quality Assessment
1.0	31/07/2025	Aljosa Pasic (ATS)	Final version

Quality Control		
Role	Who (Partner short name)	Approval Date
Deliverable leader	George Tsakirakis (INT)	30/07/2025
Quality manager	Juan Alonso (ATS)	31/07/2025
Project Coordinator	Aljosa Pasic (ATS)	31/07/2025

Table of Contents

SUNRISE

List of Tables

List of Figures

List of Acronyms

Abbreviation / acronym	Description
AC	Access Control
ACS	Access Control System
AI	Artificial Intelligence
AIRE	Atos Incident Report Engine
APT	Advanced Persistent Threats
AR	Application Requirements
BCMS	Business Continuity Management System
BCP	Business Continuity Plan
BR	Business Requirements
CER	Critical Entities Resilience Directive
CERCA	CybEr Risk assessment CAlculator
CI	Critical Infrastructure
CI/CD	Continuous Integration / Continuous Delivery
CIISI	Cyber Information and Intelligence Sharing Initiative
CIRAS	Cybersecurity Incident Report and Analysis System
CISCP	Cyber Information Sharing and Collaboration Program
CMS	Crisis Management System
CPR	Cyber-physical Resilience
CSRT	Computer Security Incident Response Teams
CSV	Comma-separated values
CTI	Cyber Threat Intelligence
D3.4	Deliverable number 4 belonging to WP 3
DoA	Description of Action
DPM	Demand prediction and management
DSA	Digital Services Act
DSM	Digital Surface Model
DSP	Digital Service Provider
DSS	Decision Support System
DTM	Digital Terrain Model
EC	European Commission
EMS	Energy Management System
EUDCC	European Digital Covid Certificate
GDPR	General Data Protection Regulation
GHG	Greenhouse Gases
GIS	Geographic Information System
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HER	Electronic Health Record
HPA	Horizontal Pod Autoscalers

ICT	Information and Communications Technology
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
LIDAR	Light Detection and Ranging
LLM	Large Language Model
MAAS	Mobility as a Service
MISP	Malware Information Sharing Platform
ML	Machine Learning
M-LLM	Multimodal Large Language Model
NAP	Neutral Access Point
NFC	Near Field Communication
NIR	Near-infrared
NIS2	Network and Information Security Directive
NISD	Network and Information Security Directive
ODISE	Open Data Infrastructure for Security and Emergency Management
OES	Operator of Essential Services
OODA	Observe, Orient, Decide, Act
PDCA	Plan-Do-Check-Act
PPE	Personal protective equipment
PT	Public Transportation
RAP	Regional Access Point
RiBAC	Risk-Based Access Control
SAR	Synthetic-aperture radar
SIEM	Security Information and Event Management
SOC	Security Operations Center
SSH	Social Sciences and Humanities
STLF	Short Term Load Forecasting
TIE	Threat Intelligence Engine
TINTED	Threat Intelligence Node with Data enrichment services
UAV	Unmanned Aerial Vehicle
UMCL	University Medical Center Ljubljana
VLM	Vision Language Model
VPN	Virtual Private Networks
VQA	Visual Question Answering
WP	Work Package

Executive Summary

The SUNRISE project was conceived to ensure that essential services remain uninterrupted when pandemics or other widespread crises challenge conventional operational models. Deliverable D3.4 translates the project's earlier specifications into a unified set of reference architectures, technical blueprints, and operational methodologies that can be directly applied across critical infrastructure domains—healthcare, energy, transportation, water management, and beyond. Drawing on pilot deployments and cross-sector collaboration, this document offers practical guidance for building systems that adapt dynamically to shifting conditions, maintain continuity under duress, and scale efficiently as demands evolve.

Central to this deliverable are four modular solution templates—addressing access control, demand forecasting, cyber-physical resilience, and remote inspection—that define clear component interfaces, containerized deployment strategies, and secure communication patterns. Each template is accompanied by detailed performance targets and validation workflows, ensuring implementers can replicate the high levels of prediction accuracy, throughput, and fault tolerance demonstrated in SUNRISE pilots. The reference designs integrate AI/ML pipelines for real-time forecasting and anomaly detection, together with orchestration scripts for automated scaling and failover. By standardizing these patterns, D3.4 enables rapid integration into existing networks, minimizing customization overhead while maximizing interoperability.

Beyond pure technical design, Deliverable D3.4 codifies a comprehensive operational playbook for infrastructure operators. It describes emergency provisioning procedures, role-based access workflows, continuous monitoring dashboards, and incident-response drills—each informed by lessons learned in early deployments. The emphasis on "privacy-by-design" and regulatory alignment ensures that data-sharing and AI components comply with GDPR, forthcoming AI regulations, and sector-specific standards. Governance recommendations outline agile certification processes, federated identity frameworks, and best practices for public-private partnership, thereby aligning technical resilience with organizational and legal requirements.

By weaving together strategic policy guidance, hands-on technical instructions, and real-world lessons, this deliverable equips policy-makers, technology providers, and operators with an end-to-end roadmap for resilient critical infrastructure. Whether facing pandemic lockdowns, cyber-attacks, or natural disasters, stakeholders who adopt the reference designs and best practices in D3.4 will be well-positioned to maintain vital services, accelerate recovery, and foster long-term societal resilience.

1 Introduction

1.1 Purpose of the document

Deliverable D3.4 aims to deliver the final reference designs and consolidated best practices for the SUNRISE Tools. While WP3 has, to date, bridged the gap between business and technology through D3.1 (initial requirements) [43], D3.2 (V2 designs) [44], and D3.3 (final specifications) [45], this document completes the picture by translating those specifications into production-ready architectures and operational guidelines for CI environments. It also integrates feedback gathered from the pilot activities, thereby closing the loop between design, implementation, and real-world validation.

1.2 Relation to other project work

Within the SUNRISE framework, WP3 serves as the pivotal bridge between the business-oriented work packages (WP1, WP2, WP8, and WP10) and the technically focused streams (WP4, WP5, WP6, and WP7). In its earlier deliverables, WP3 documented the evolving requirements and architectural designs through the initial, intermediate, and first piloting phases, thereby charting the continuous maturation of the SUNRISE tools. Deliverable D3.4 builds on this foundation by presenting the final, production-ready reference designs alongside a compendium of best practices derived from challenges encountered by tool owners during development and validation.

Contributions from the non-technical WPs have been instrumental in shaping these reference architectures. For example, D1.1 [41] provided the foundational business case analysis and stakeholder requirements that guided our prioritization of use-case scenarios and governance models. Similarly, outputs from D2.2 [42] and D10.2 [50] informed our policy alignment strategies and privacy-by-design principles embedded within each design template.

On the technical side, WP3's work has continuously leveraged successive deliverables from the tool-specific packages since month M6, using their evolving prototypes and test results as inputs to refine integration patterns and performance benchmarks. In particular, the M34 deliverables—D4.6 [46], D5.6 [47], D6.6 [48], and D7.6 [49]—have directly influenced the final content of this document, supplying the latest implementation insights, validation data, and deployment lessons that underpin the reference designs and operational guidelines set out in the following chapters.

1.3 Structure of the document

This document is divided into four main sections:

- **Chapter 1** defines the purpose of the document, relation to other project work and structure of the document.
- **Chapter 2** provides the High-level overviews of each tool's final solution (description, key features, system architecture).
- **Chapter 3** outlines the detailed reference designs (objectives, diagrams, performance benchmarks)
- **Chapter 4** presents the best practices (development, integration, testing, maintenance, lessons learned).

2 Overview of the final Solutions

This section provides a structured overview of the final SUNRISE tools. Each subsection introduces a specific solution, summarising its purpose, key features, and the system architecture behind it. The content outlines how each tool contributes to the overall support of Critical Infrastructure (CI) operators within the project.

2.1 Risk-based access control

2.1.1 High-Level Description

The Risk-based Access Control (RiBAC) tool (WP4) enhances physical access control in Critical Infrastructure (CI) facilities during pandemic or emergency situations while maintaining operational continuity. Unlike traditional systems focused solely on identity verification, RiBAC introduces dynamic risk assessment based on health and safety parameters.

RiBAC's unique value lies in seamlessly integrating pandemic-specific health screening with existing access control infrastructure. Rather than requiring complete system replacement, RiBAC serves as an intelligent extension operating in both normal and emergency modes. This dual-mode capability ensures CI operators maintain existing workflows while gaining enhanced protection when needed.

The tool addresses the critical business challenge of maintaining workforce safety without compromising operational efficiency. During COVID-19, many CI operators struggled with manual health screenings that created bottlenecks and increased infection risks. RiBAC automates these processes while providing comprehensive audit trails and compliance documentation required by health authorities.

2.1.2 Key Features

Multi-Modal Health Screening

Combines contactless temperature measurement, protective equipment detection (masks, PPE), and vaccination credential verification in a single streamlined process. Comprehensive screening completes in approximately 10 seconds per person while maintaining high accuracy rates (>98% for protective equipment detection, ±0.3°C temperature precision).

Adaptive Configuration

Supports flexible operational modes adjustable based on current risk levels. In normal operations, the system functions as standard access control. During pandemic conditions, additional health screening modules activate automatically based on predefined policies or manual triggers.

Privacy-Preserving Design

Implements privacy-by-design principles, processing health data locally without cloud transmission. Vaccination credentials are verified using privacy-preserving cryptographic protocols ensuring GDPR compliance.

Legacy System Integration

Designed to augment rather than replace existing access control systems. Supports standard interfaces (Wiegand, RS485) and integrates with major ACS platforms, minimizing deployment costs and training requirements.

Remote Management

Includes secure remote configuration and monitoring capabilities, allowing operators to adjust settings, update policies, and monitor system status across multiple locations from a central interface.

2.1.3 System Architecture

RiBAC employs a modular architecture centered around six core hardware/software modules designed for independent operation and easy replacement:

Processing Core

NVIDIA Jetson Nano platform [1] provides sufficient computational power for AI-based image processing with low power consumption. Includes 1GB RAM and expandable storage up to 128GB for offline operation.

Sensor Modules

Dual camera system comprising separate RGB camera for protective equipment detection and thermal camera for contactless temperature measurement (not stereo vision implementation). RGB camera includes AI acceleration via Intel Neural Compute Stick 2 [2] for real-time processing.

User Interface

7-inch capacitive touchscreen with intuitive graphical user interface (GUI) providing real-time feedback. Supports multiple languages and displays custom instructions based on operational policies.

Access Control Interface

Universal Radio Frequency Identification (RFID) reader supporting multiple standards (Mifare [3], DESFire, LEGIC, HID) with Bluetooth and Near Field Communication (NFC) [4] capabilities for mobile device integration.

Communication Module

Ethernet and WiFi connectivity with secure VPN support for remote management. Operates in offline mode with local data storage and synchronization capabilities.

Vaccine Credential Module

QR code scanner with integrated verification supporting national digital certificates and privacy-preserving credential systems.

A generalised depiction of RIBAC architecture can be seen in Figure 1.

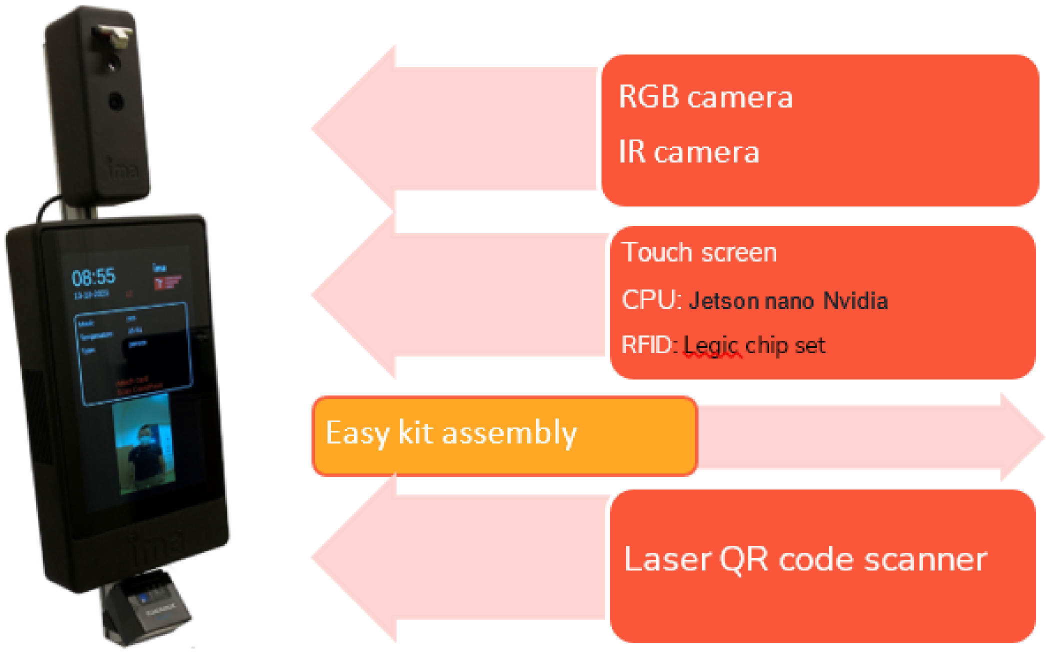

Figure 1: RiBAC Architecture

2.2 Demand prediction and management

2.2.1 High-Level Description

Demand prediction and management (DPM) Tool helps the operators of CI in handling changing demands of various sectors such as energy, transport, health and water sectors. With the help of the DPM AI-powered solution, short and mid-term forecasting is provided on the basis of historical data concerning the sector, weather, mobility trends, and popular events to be used in decision-making, clustering resources, and becoming more robust in crises situations like a pandemic. The final, third-year version of the DPM Tool that will be presented focuses on the predictive precision and the functionality versatility combining advanced machine learning algorithms with the user interface enhancements peculiar to CI. It also provides and adheres to multi-scenario forecasting (e.g., the level of pandemic impact), which is close to the requirements of the final users and thus it is ready to work both in a regular regime and crisis.

2.2.2 Key Features

Demand Prediction and Management Tool includes a set of comprehensive functionalities that assist operators of a critical infrastructure in various aspects, such as energy, water, health and transport. The tool leverages advanced AI-based, model-oriented forecasting—using methods such as N-BEATS [5], LSTM networks [6], and XGBOOST [7] which have been tailored to the specificities of each sector, including historical load behavior, mobility patterns, climatic anomalies, and crisis-related deviations.

To enhance the capabilities of the back-end system, a web-based interface that facilitates interaction of users with predictive services in an intuitive manner based on the UI/UX best practices is one of the distinguishing characteristics of the tool. The user can navigate through historical and estimated data and edit the parameters of predictions. One of the most important aspects of the web application is the intrinsic way of results' representation using a Grafana [8] based dashboard view. To support advanced interaction—particularly in health-related applications—the tool offers scenario-based forecasting. Users can choose from predefined pandemic impact levels (low, moderate, or high) to view corresponding demand forecasts. These scenarios are pre-calculated using different model configurations, allowing users to instantly compare how varying degrees of crisis affect demand. This helps decision-makers evaluate readiness, resource needs, and contingency plans under different operating conditions. The tool includes a robust primary data upload module that enables users to submit structured files (e.g., CSV) through an intuitive web interface, with immediate validation feedback for issues such as formatting errors, missing values, or schema mismatches. The same ingestion and validation pipeline is also accessible via a RESTful API, supporting integration into automated workflows. This dual-mode approach ensures consistent data handling while accommodating both manual uploads and system-to-system data transfers. It also supports the integration of data sources via API. The same functionality is provided through a special REST API endpoint to accommodate workflow automation and external system integration. This twofold solution means that both technical and non-technical users will be able to provide data efficiently through manual or automated pipelines. In addition, the use of a REST API provides the necessary means for future 3rd party extensibility of the system.

In addition to the forecasting features, the tool provides in-built performance indicators that show the accuracy of predictions over time and application; the tool has inherited indicators like MAE, RMSE and SMAPE. In the transport-oriented cases, the demand intensity visualized in the form of a dynamic heatmap provides useful information that can be used in the scope of logistics planning and responding to emergencies.

2.2.3 System Architecture

The Demand Prediction and Management Tool architecture is based on cloud-native and modular solutions that provide sources of flexibility, elasticity, and seamless integration with Critical Infrastructure operators. The proposed solution follows the microservices architecture leveraging technologies like Docker containers and Kubernetes orchestration capabilities on a managed cloud environment. This guarantees the ability to scale up resources, guarantee availability, and fault tolerance, while ongoing integration and delivery pipelines to update the components incrementally, eliminating the presence of downtime (Figure 2).

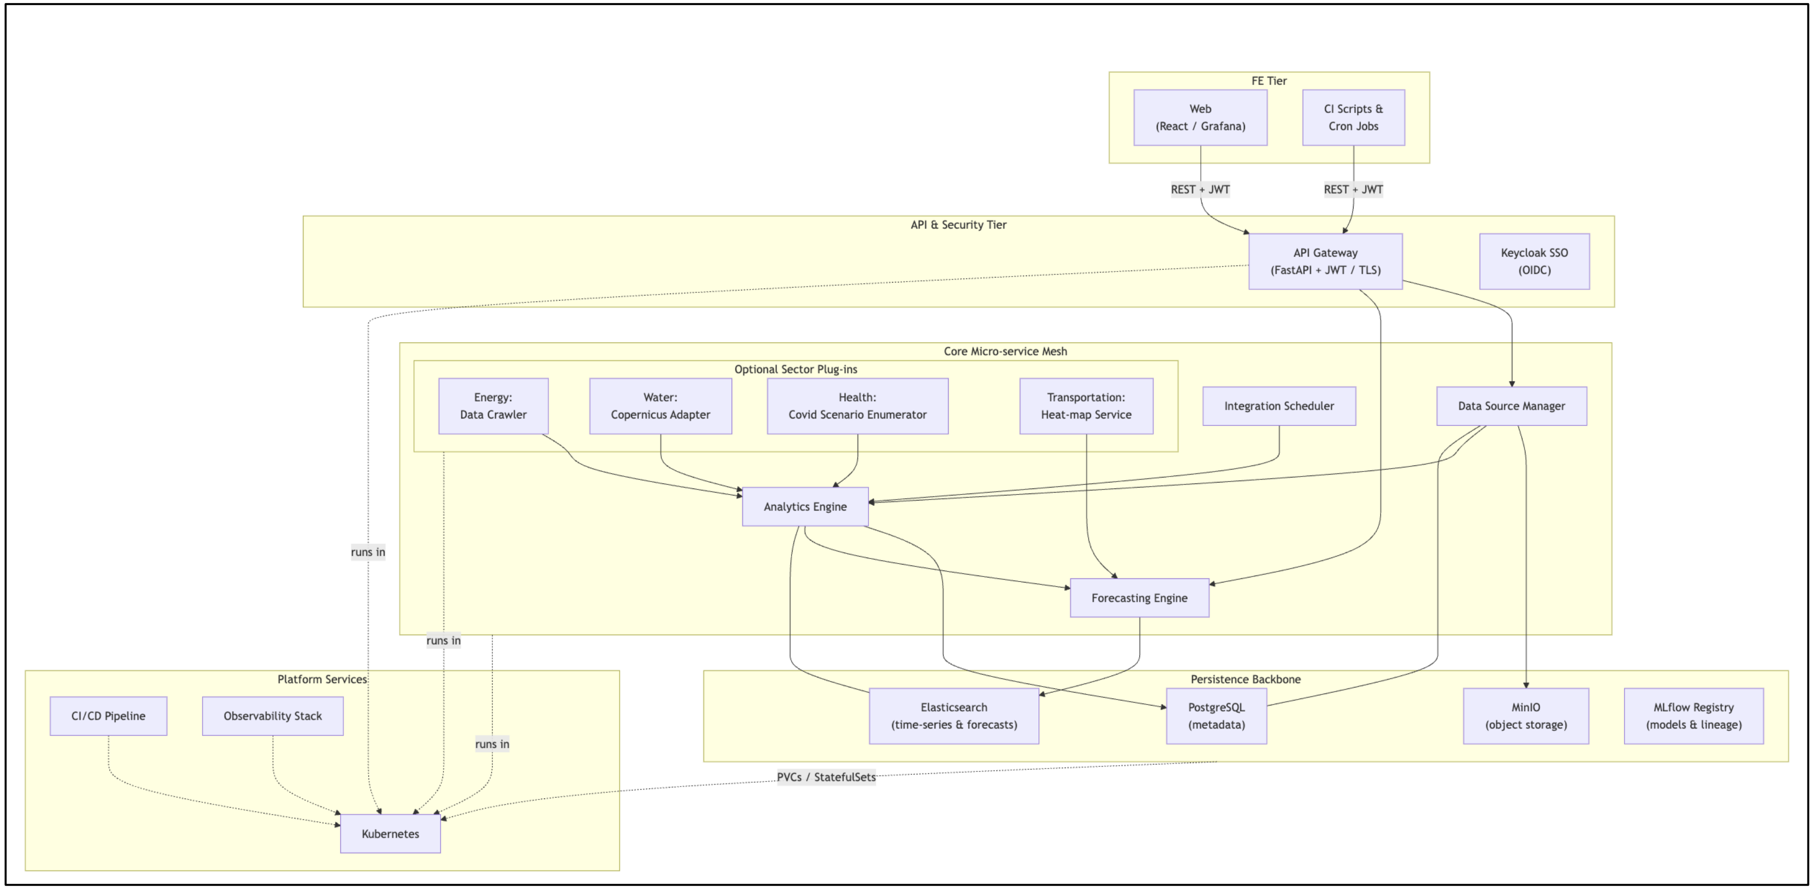

Figure 2: Demand Prediction and Management tool architecture

The DPM Tool can be divided in three distinctive entities Backend, Frontend and data pipelines. The Backend is actually heavily optimized backend engine including different services for each one of the sectors, one is user-friendly and smart frontend application, and another is a composite data-pipelines set to ensure secure, accurate, and timely processing of historical and real-time input streams. The tool provides the CI operators with a secure, web-based tool interface at the frontend, and as the primary point of interaction with the tool. It is user-friendly and secure in terms of the interface where the authentication and role-based access uses Keycloak [9]. When the user logs-in, he/she gets access to his or her organization-specific and operational-specific dashboard. These dashboards allow them to navigate through past trends, initiate new forecasts, review the metrics of predictions, and assess other possible scenarios in other operating circumstances.

Ingestion of data is done through two complementary mechanisms. First, there is an intuitive upload module implemented in UI, which enables the operators to upload structured files (e.g. CSV). The data are transferred to the system instantly and validated whether they have the correct format. Second, where organizations need automated data feeds, RESTful API offers the same capabilities, and other systems can push the data into the system itself. The two mechanisms share an object storage layer based on S3 MinIO [10], which is passed to the preprocessing and forecasting pipeline at the backend. It is built upon the FastAPI framework [11] and is deployed in a containerized environment.

The microservices, e.g. the forecasting engine, file validation logic, and the analytics layer have been encapsulated into Kubernetes [12] pods. Cross-services communication is established with the use of internal APIs channels and shared volumes of data, arranged with the use of Kubernetes native workflows. The operation of forecasting goes through a systematic pipeline where forecasting work is first performed, including data extraction and transformation and gets inferred into an AI model (orchestrated by MLflow [13]), and finally the forecasting result is obtained and stored using PostgreSQL [14] relational database and accommodated in Elasticsearch [15] to allow quick-search and visualization.

The tool is dynamic as it uses outside data-warehouses of information like Copernicus seasonal forecasts and COVID-19 scenarios and uses that to better its predictions. Since the analytics engine is decoupled to UI, forecasting tasks may be run asynchronously and continuous user interaction is supported even when the system is doing heavyweight work.

Hosted completely on Kubernetes, DPM Tool leverages the inherited capabilities of K8s such as resiliency, load balancing, and horizontally scaling. This infrastructure option makes the tool flexible on increased work demand. Moreover, this solution enables the integration of hybrid cloud native settings, and high service availability, hence a resilient and future-proof solution to critical infrastructure demand management.

2.3 Cyber-physical resilience

2.3.1 High-Level Description

Cyber-physical Resilience (CPR) tool arises from the need for Critical Infrastructures to have an integral and proactive security solution in their Security Operations Centres (SOCs) that can easily adapt to emergency and temporary situations and operational conditions, such as those faced during the COVID-19 pandemics. In circumstances that frequently change, cybersecurity operational teams in Critical Infrastructure operators realized they need to address issues related to human risks (e.g. lack of experienced staff for incident triage), evolution of threats (e.g. increase of COVID-19 related phishing messages) or change of priorities and security policies (e.g. related to remote work), derived from measures implemented at strategic level to contain pandemic spread (e.g. lockdown or restrictions in mobility). They had to address the increasing number of threats and cyber-attacks, while maintaining their cybersecurity efficient even with the reduced or psychologically affected staffing

Critical Infrastructures cybersecurity teams realized they had to enhance their decision-making systems by incorporating greater context awareness through aggregation of internal real time risk indicators, as well as external threat intelligence. The dynamic risk assessment provided by the CPR tool, with support for the application of temporary conditions and simulation of mitigations, enables them to adapt their procedures and mitigation measures to potential changes in the critical infrastructures' supply chains, to the introduction of new digital services, or changes in normal behaviour or connectivity patterns with an increased deployment of VPNs, employees working remotely, workforce absenteeism and a lack of security capacity and skills.

Cyber-physical Resilience tool also aims to address the need for critical infrastructures of tools that help them in the compliance with the increasing number of regulations at different levels (European, nationals and sectorial), such as GDPR, NIS2 or CER. Most of these regulatory frameworks require the timely generation and notification of incident reports to the Supervisory Authorities in different formats and deadlines, and with different levels of details about the security incidents. CPR tool integrates an Incident response and threat intelligence sharing module to support the CIs during the whole mandatory incident reporting process. Through secure data sharing with fine grained data sharing policies, where a piece of information can be encrypted or anonymized for a specific stakeholder or group, CPR tool also facilitates and fosters voluntary collaboration among critical infrastructures to combat the rise in the number of threats and cyberattacks they face. And thanks to its automatic processing of threat intelligence data and generation of source confidence scoring and threat scoring, it facilitates threat hunting tasks focusing the analyst efforts on trustworthy sources and information relevant for the monitoring infrastructure.

2.3.2 Key Features

CPR tool, through the use of state-of-the-art self-supervised machine learning models from the area of Natural Language Processing (NLP), it is able of an enhanced detection of anomalous behaviours based on logs collected from the monitored infrastructure that are usually not processed by Security Information and Event Management (SIEM) solutions. In this way, the tool complements the detection currently done by the monitoring solutions deployed in the Critical Infrastructures, such as Wazuh [16] or IBM QRadar [17], and provides a better anomaly detection for situations where existing detection rules and patterns are not suitable (e.g. zero-day attacks).

One of the key features of the CPR tool is the reduction of uncertainty during risk assessment improving decision-making process. This improvement has been achieved by the integration of temporary conditions (with a configurable percentage of impact in some risk indicators or changes in the asset value), physical risk indicators, such as availability of essential employees, health trend contextualization and indicators of compromise received from threat intelligence sources into the conventional risk assessment process (which usually relies on known vulnerabilities, a business profile and established risk models) to close the gap between strategy and operations. All these new inputs added in the CPR cyber-physical risk assessment module facilitates for example the management of absenteeism data, providing clearer insights into the availability of human resources during emergency situations and its impact on the risk assessment. The consideration of physical events can be also critical in specific situations, increasing in this way the preparedness and response capabilities during incidents. During the last period of the project, it has been also added a mitigation simulator that allows the risk assessment evaluation assuming/considering the application of specific countermeasures linked to the different indicators included in the risk models. All these dynamic risk assessment features implemented represent an important innovation in resilience management and have direct influence in the absorption of incidents and the overall resilience of the Critical Infrastructures. It is noteworthy that risk management is mostly pre-event (probabilities, scenarios), while crisis management focus is mostly during actions when risk materializes. Resilience, on the other hand, cuts across the timeline, focusing on system performance over time and this is why risk has to be dynamic and re-triggered also during the event evolution.

CPR tool provides some key features related to improving resilience of the Critical Infrastructures through collaboration and voluntary information sharing. CPR threat intelligence module provides a smooth integration of anonymization and encryption capabilities with the MISP [18] instance that can be already deployed in the Critical Infrastructures. Through the usage of tags, sensible information can be easily labelled for anonymization. Encryption is managed through data while a dashboard is provided to facilitate the sharing and visualization between specific users or groups Additionally, the CPR tool capacity to automatically process incoming data received from threat intelligence feeds and label these MISP events with granular threat scoring and source confidence scoring, and its automated mapping to MITRE ATT&CK attack patterns significantly enhance the efficiency of the Critical Infrastructures' response and security operation teams. It allows reducing their response time and improving the overall effectiveness of threat detection and response, ensuring they can proactively address potential threats concentrating their efforts on accuracy and relevant information.

The last core feature is connected to the support provided by the CPR tool to the Critical Infrastructures' legal and compliance teams for enhanced and more efficient incident response management, covering the existing gap in the harmonization of reporting procedures throughout the different applicable regulatory frameworks. CPR incident reporting module automates a flexible incident reporting workflow that can be easily adapted by the Critical Infrastructures' teams to different regulations in terms of deadlines for incident notifications, criteria used to consider an incident as significant, and templates that must be used for reporting (which can vary from one competent authority to other and at national level). The connection added with the CPR risk assessment module and the CPR threat intelligence module allows the integration of supply chain risk consideration into the incident reporting process, the consideration in the risk assessment of the application of countermeasures received as feedback from the CSIRTs after an early warning, and the consequently enrichment of the reports with all this additional information promoting at the same time the voluntary information sharing.

2.3.3 System Architecture

Cyber-Physical Resilience tool has a modular system architecture with a unified dashboard as shown in Figure 3. This allows to offer a wide range of complementary functionalities from a single access point, and it can be more easily integrated with the existing solutions in the Critical Infrastructures and scaled or deployed depending on the specific needs of the CIs.

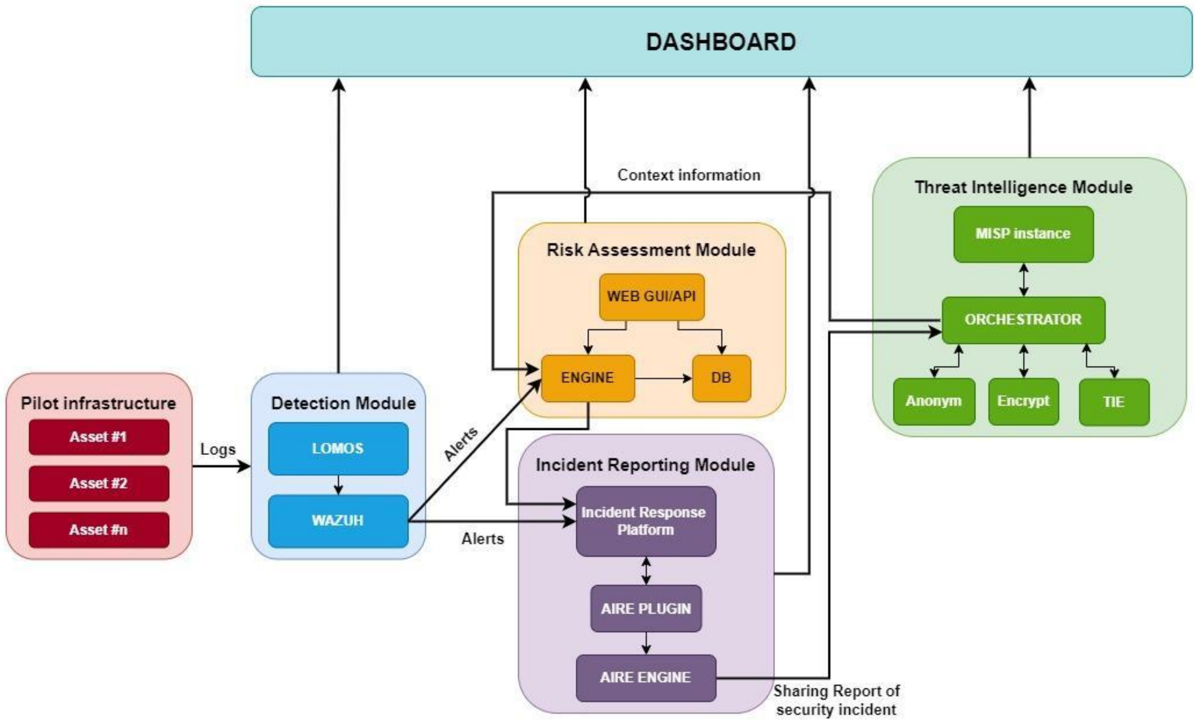

Figure 3: Cyber-Physical Resilience tool architecture extracted from D6.5[19]

As described in D6.5 [19] logs generated by the CIs' infrastructure are consumed by the CPR Anomaly Detection module, which is based on LOMOS (LOg MOnitorubg System) for log-based anomaly detection and Wazuh SIEM (selected for the pilots in the project although any other SIEM could be used) for generation of alerts. These alerts are received as real-time inputs for the CPR Risk Assessment module to trigger the generation of a new risk report. Context information used in the dynamic cyber-physical risk assessment performed by the CPR are provided by the CPR Threat Intelligence module. This information arrives filtered with a scoring generated by the TIE (Threat Intelligence Engine) so only relevant data will be considered.

The CPR Incident Reporting module is responsible for the incident management and generation of the incident reports that must be notified to the Supervisory Authorities. It integrates the open-source incident reporting platform TheHive [20] with AIRE (Atos Incident Reporting Engine) and consume alerts generated by the CPR Anomaly Detection module, risk reports generated by the CPR Risk Assessment module and Indicators of Compromise collected through the CPR Threat Intelligence module. The resultant reports generated are sent via email, but they can be also shared through MISP.

Secure data sharing through anonymization and encryption and with the use of sharing agreement is provided through the CPR Threat Intelligence module.

2.4 Remote infrastructure inspection

2.4.1 High-Level Description

The tool is intended for remote infrastructure inspection. Business sectors (energy transfer, energy production, transport, water management) in which the critical infrastructure operators, involved in the project, need such a tool to mitigate possible unforeseen circumstances and to ease their workload in normal circumstances.

The Remote Infrastructure Inspection (RII) tool eases the infrastructure inspection especially in cases where the infrastructure is hard to reach and vast. Therefore, we can expand the aforementioned business sectors to any sector that deals with such infrastructure.

The RII tool uses satellite and UAV technology, combined with the AI features. The addition of the AI features directly addresses the need to flag defects or possible defects only. This reduces the so-called information fatigue, experienced by operators of similar systems. However, as the tool is intended for use in the critical infrastructure, it must allow for manual inspection of the flagged defect. In the case of the RII tool, this means manual inspection of the footage taken with the UAV or the manual inspection of the satellite image. Thus, while the RII approach automates discovery of potential problems on the infrastructure, the failsafe is built in, enabling human decision on the state of the monitored infrastructure. The original footage—whether UAV-captured video or satellite imagery—is preserved to support human inspection and decision-making. This ensures situational awareness not only for repair operations but also for rapid emergency responses, such as fires or other urgent threats to critical infrastructure.

Having described the tool and the need for it, we turn to its innovative aspect, e.g., while the satellite imagery is already used in such applications (the satellite imagery providers have commercial offerings for monitoring linear infrastructure) and UAV footage is also used for monitoring, the innovation of the RII tool is in its use of AI modules and techniques, all whilst preserving the most important aspect – preserving the human ability to double-check the respective potential threats, defects or problems.

In the context of monitoring critical infrastructure through satellites and UAVs, the data processing detects potential threats and assesses the overall condition of the infrastructure. This processed information is then used to make informed decisions regarding maintenance, security measures and response strategies to mitigate risks and safeguard the critical assets. Importantly, it allows the critical infrastructure operator to prioritize repairs, based on the condition and availability of the repair crew and other, possibly unforeseen events.

The ability to prioritize is an important aspect of the outcomes provided by the tool. Namely, the critical infrastructure provides essential services to the public and in case of a defect, the priority to re-establish the service must be considered (e.g., in case of electricity supply defect, a hospital is always given higher priority than some residential area). It is also crucial to be able to give priority to the needed repairs in cases of unforeseen events – e.g., a weather event that causes mass degradation or loss of service by several critical infrastructures. In this case, the priority list may change, based on the needs and priorities of the public and other critical infrastructure operators.

The focus of the AI components, and thus our approach, is to reduce information/alert fatigue in operators and to reduce the need for human inspection, the latter being particularly important in cases of unforeseen temporary conditions where workforce availability may be restricted. This structured approach to data flow ensures that critical infrastructure operators have the necessary tools and insights to respond effectively to challenges and events within their operational environment.

High-resolution images and videos of critical infrastructure are captured by satellite and UAV systems. Subsequently, raw data undergoes processing using AI tools to enhance imagery quality and extract pertinent information.

The enhanced data is integrated into a user-friendly dashboard interface, empowering operators to visualize and analyze the information effectively. By offering operators a comprehensive view of the infrastructure, this process enables them to make well-informed decisions in response to various issues or events.

2.4.2 Key Features

One approach to remote inspection is the use of satellite imagery. The main advantage of this approach is the ability to monitor large areas in a continuous and non-invasive way. With a growing number of satellite providers, high resolution imagery can be collected every few days (the cost increases with the desired resolution, while low resolution images are free). Our satellite inspection tool can process optical and multispectral satellite imagery to detect overgrown vegetation and other changes in the area (i.e., change detection). This can support remote inspection of critical infrastructure by detecting events, which might pose danger to the CI. On the other hand, the main limitation in this approach is the resolution of the imagery, which restricts the level of detail and the scale of changes that can be detected. Therefore, this tool can also be used as a trigger for more localised and detailed UAV inspection, described in the next paragraph.

The UAV Inspection Tool is a key component of the SUNRISE remote inspection framework, designed to modernize the monitoring and maintenance of CIs through AI-driven automation and high-resolution aerial imaging. By integrating computer vision techniques, multimodal large language models (M-LLMs - computer vision and multimodal large language models; also referred to as vision-language models (VLMs)) and real-time data processing when needed, this tool enhances the ability of infrastructure operators to conduct non-invasive, frequent and risk-based inspections, minimizing reliance on traditional manual assessments.

The UAV inspection plays a crucial role in complementing satellite-based monitoring, offering detailed, localized inspections where satellite imagery lacks the necessary resolution. This combination of wide-area surveillance and precise UAV-assisted analysis provides a multi-layered inspection strategy, optimizing resource allocation to enable timely maintenance interventions. The tool's modular and adaptive architecture allows UAVs to be configured with different sensor types, adapting to the specific needs of energy, transportation and water infrastructure sectors.

We thus have an interplay of the methods, coarser satellite imagery-based and finer UAV-based, both assisted by the different AI modules. We also split them into different modules, as it will be shown below.

The satellite-based remote inspection tool consists of two components:

- Infrastructure change monitoring: Change detection in the vicinity of critical infrastructures is crucial – while some of the environmental events such as e.g. landslides, leaks can be discovered very efficiently, the human imposed events such as illegal build-ups are even more important. Namely, it is illegal to interfere with the environment near the critical infrastructure. Thus, it is crucial to detect such events in a timely manner in order to check for potential damages and to address potential safety concerns. Our component provides a general solution for detecting changes in the vicinity of the CI based on optical satellite imagery.

- Vegetation management: One of the potential threats to critical infrastructure is overgrown vegetation. This component supports vegetation monitoring by estimating the height of vegetation from freely available multispectral satellite imagery. Depending on the vicinity of the critical infrastructure, detected vegetation height can be used for threat estimation.

The UAV-based remote inspection tool addresses major challenges in infrastructure monitoring, such as aging assets, environmental risks and regulatory compliance requirements. Traditional inspection methods are often time-consuming, costly and constrained by personnel availability, limiting their frequency and effectiveness. By automating key inspection processes, the UAV tool enables the early detection of structural defects, vegetation encroachment and environmental hazards, significantly reducing operational risks and optimizing maintenance workflows.

UAV and AI-assisted assessment of critical infrastructure leverages major advances in computer vision and multimodal large language models (M-LLMs) (also referred to as vision-language models (VLMs)) to address real-world challenges in an innovative and effective way. This is reflected in its three main modules:

- object detection
- semantic segmentation
- Visual Question Answering (VQA)

which provide automated analysis of a broad range of diverse infrastructure conditions. Additionally, secondary modules such as 3D virtualization and anonymization enhance the tool's capabilities while ensuring compliance with privacy regulations and user data protection.

The described methodology (going from coarse to fine resolution, assisted by AI methods to reduce alert fatigue) is coupled with the intuitive user interface, which allows for integration of the legacy systems.

In the following paragraphs, we describe the use of the RII tools, end user benefits, performance enhancements and operational impact.

The coarse satellite imagery and the tools based on it, proved to be of limited use to CIs involved in the project. We invested considerable effort to extract the precise requirements and from the required tests to tailor the tools to the practical and financial needs of each CI. At HDE, the use of high-resolution satellite imagery to monitor issues like clogged grates or sediment accumulation proved ineffective - despite using advanced high-resolution satellite images, they proved to be still too coarse to show the potential threats (when checked manually), thus they were not usable even for the AI-based methods. For SZ, although a vegetation height model works and legally, the vegetation height is limited, they found it only marginal use for it, given their specific operational context. Additionally, they concluded that landslides could not be effectively detected via satellite due to resolution limitations and time constraints (i.e., much faster detection on the ground, using sensory grids).

ELES evaluated the vegetation monitoring tool, however given their legally protected vegetation free zones can be monitored in conjunction with UAVs regular checks of the infrastructure. Their primary interest shifted to identifying illegal structures such as hunting lodges within restricted areas - a use case where satellite imagery and change detection were considered valuable.

ACO found satellite imagery only useful for detecting very large water spills, which they already detect faster with their existing pressure-based systems. Consequently, they agreed with the project team to prioritize more practical and cost-effective alternatives, like UAV technology, for infrastructure monitoring in remote areas.

Given that vegetation height monitoring component in particular and the change detection component reached mature status (TRL7), the CIs reported high integration, development percentages (80<).

The business needs of the involved CI operators were better answered with the use of the UAV detection. All the involved CI operators reported high integration, development percentages of the tool (80<). They also provided estimates of the budget needed to adopt such a tool from at least the UAV perspective (50,000EUR for the equipment, and at least 2,000EUR for obtaining the piloting licences).

The CIs highlighted a significant increase in operational efficiency through early anomalies detection and efficient resource allocation.

The SUNRISE project's holistic approach ensures that this technology is not only a solution for crisis scenarios but also an essential component of everyday infrastructure management, improving safety, longevity and sustainability.

2.4.3 System Architecture

The overall system architecture of the RII Tool is modular (Figure 4 and Figure 17). Both, UAV and Satellite imagery are designed to be stand-alone components, but capable of secure interfacing and coordination with the back-end, which provides data to the graphical user interface (GUI).

Components Breakdown:

- Satellite Inspection Subsystem: Handles data from satellite imagery, performs vegetation height analysis and change detection (e.g., illegal buildings, landslides).
- UAV Inspection Tool: Equipped with cameras and sensors to capture detailed images. It includes modules for object detection, semantic segmentation, and 3D rendering. Anonymization and VQA (Visual Question Answering) are also integrated.
- Unified Backend: Aggregates and processes data from both subsystems and manages data flows.
- Dashboard/User Interface: Provides visual outputs, insights, and administrative tools for CI operators.

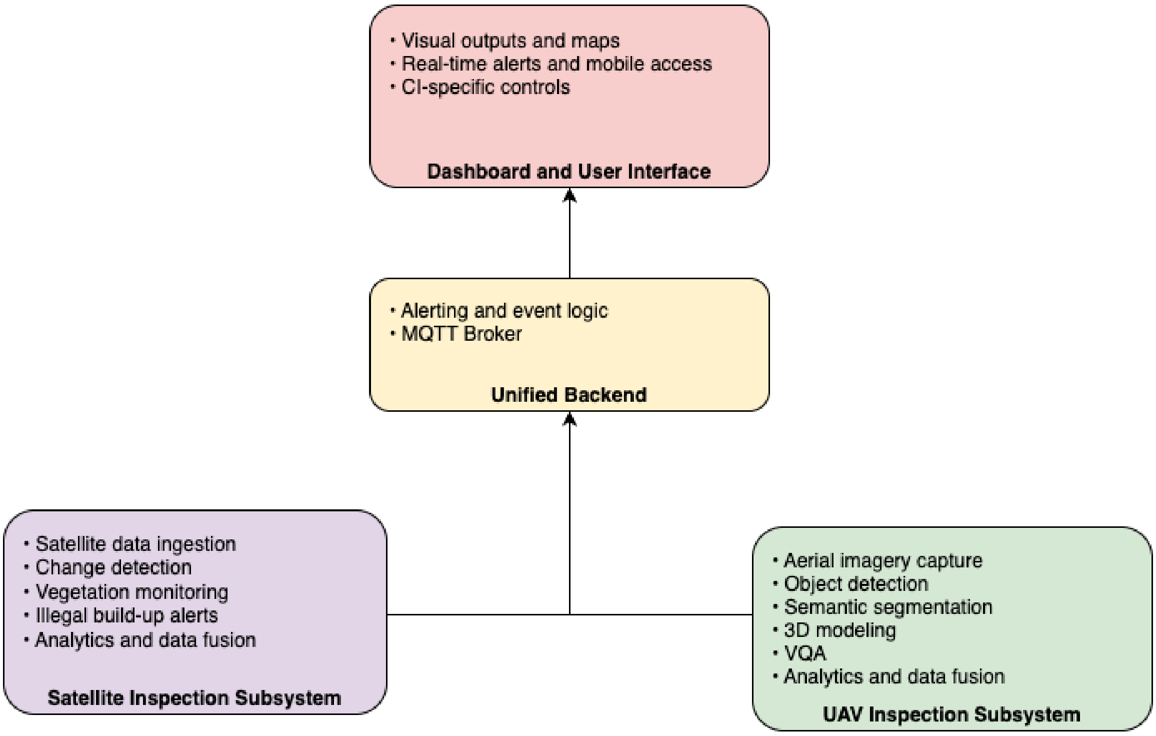

Figure 4: The overall high-level RII tool system architecture.

3 Reference Designs

3.1 Risk-based access control

3.1.1 Design Objectives

The primary design objective focused on ensuring health screening capabilities could be seamlessly integrated without disrupting normal CI operations. This required developing a system capable of switching between normal and enhanced screening modes with minimal user intervention while maintaining consistent access control functionality.

The system was designed to support deployment across diverse CI environments, from small facilities with single access points to large installations with dozens of terminals. The modular architecture allows selective deployment of components based on specific facility requirements and risk assessments, ensuring scalability and cost-effectiveness.

Regulatory compliance formed a core design requirement, with RiBAC meeting stringent data protection requirements including GDPR compliance and adopting privacy-by-design principles.

Technology readiness and reliability were prioritized through the use of proven technologies and standard interfaces to ensure high reliability and maintainability. Critical components utilize industrial-grade hardware with established supply chains and comprehensive documentation, reducing deployment risks and long-term maintenance costs.

3.1.2 System Diagrams

The RiBAC system architecture is illustrated in Figure 5, which shows the hierarchical organization of components and their interconnections. The terminal serves as the central processing unit, integrating multiple sensor inputs and communication interfaces to deliver comprehensive access control functionality

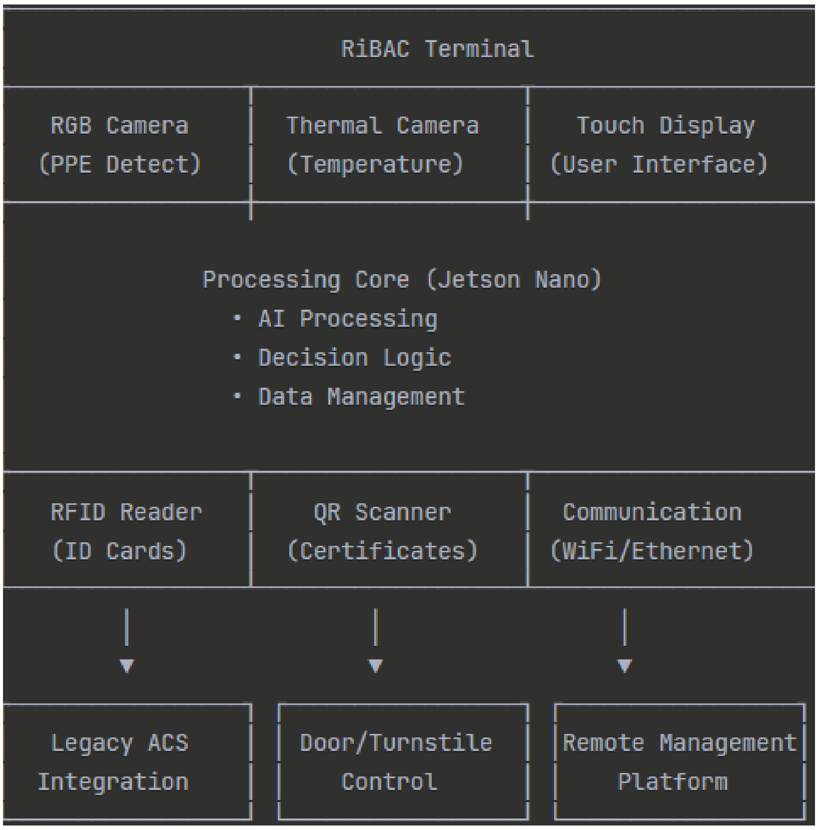

Figure 5: RiBAC Tool Block Diagram

As shown in Figure 6 , the system architecture emphasizes modularity and scalability. The upper tier contains sensor inputs and user interface components, the middle tier houses the central processing core, and the lower tier manages identification and communication functions. External integrations with legacy systems, physical access controls, and remote management platforms provide comprehensive operational capabilities.

Data flow follows a systematic sequence: user approach detection, parallel health screening processing, identity verification, risk-based decision processing, action execution, and secure reporting of anonymized access control data to management systems.

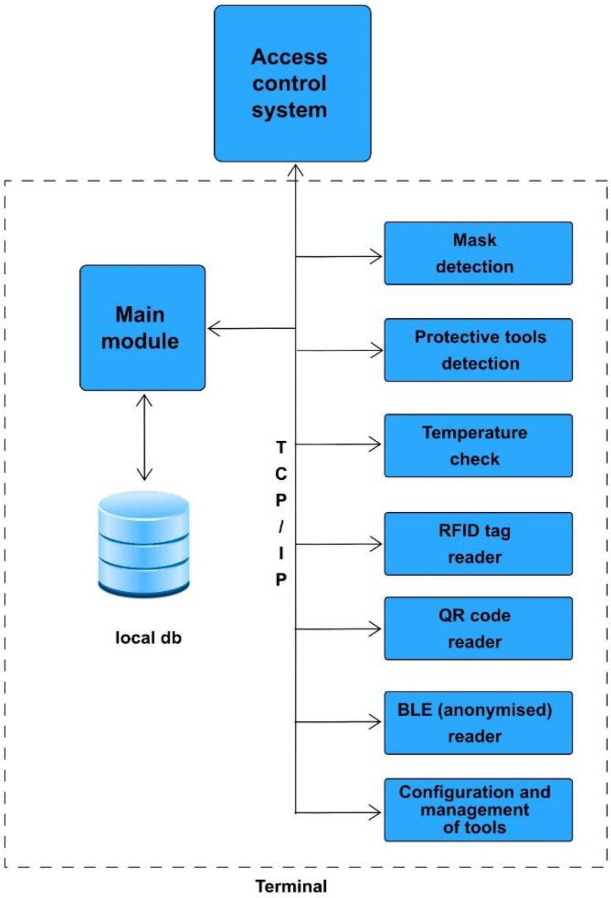

Figure 6: RiBAC Tool Modules

3.1.3 Performance Benchmarks

During comprehensive first round pilot testing, as documented in D4.4 [21], RiBAC achieved average screening times of 10 ± 4.14 seconds with 96% of users completing screening within 13 seconds. Temperature measurement accuracy reached ±0.3°C under controlled conditions, with compensation algorithms addressing environmental factors.

Protective equipment detection achieved 98.5% accuracy across diverse test conditions including different mask types, lighting conditions, and user demographics. The system maintained false positive rates below 2% while achieving zero false negatives for properly worn protective equipment.

System reliability during the piloting phase reached 99.7% uptime across 9 deployment sites with zero critical failures. Battery backup provides 4-8 hours of operation during power outages. Integration success was achieved with legacy access control systems in 7 of 9 pilot sites with minimal configuration changes, while remote management capabilities enabled 95% of configuration tasks to be completed without on-site visits.

3.2 Demand prediction and management

3.2.1 Design Objectives

The design of the DPM Tool follows a set of principles aimed at ensuring scalability, modularity, and ease of integration across varying CI operator environments. The primary objective was to deliver a sector-agnostic forecasting engine that can be tailored to specific infrastructure needs, such as hospitals, water utilities, energy operators, or transport. Flexibility was a core requirement, as each operator has different digital maturity levels, operational constraints, and data availability.

Key design goals included:

- Scalability, to handle increasing data volumes and prediction requests.
- Security and multi-tenancy, to allow isolated access for different CI operators.
- Resilience, through asynchronous and fault-tolerant forecasting pipelines.
- Interoperability, via standardized RESTful APIs and containerized deployment.
- User-centric operation, ensuring both technical and non-technical staff could use the system effectively.

These objectives translated into a modular design pattern where each core capability—data ingestion, preprocessing, forecasting, result visualization—is implemented as a loosely coupled, Kubernetes-managed microservice. This enables flexible deployments and continuous integration and delivery without interrupting service operations.

3.2.2 System Diagrams

The high-level reference design of the tool is composed of four main architectural layers (Figure 7): the data ingestion layer, the processing (Figure 8) and forecasting layer (Figure 9), the persistence and indexing layer, and the user interaction layer.

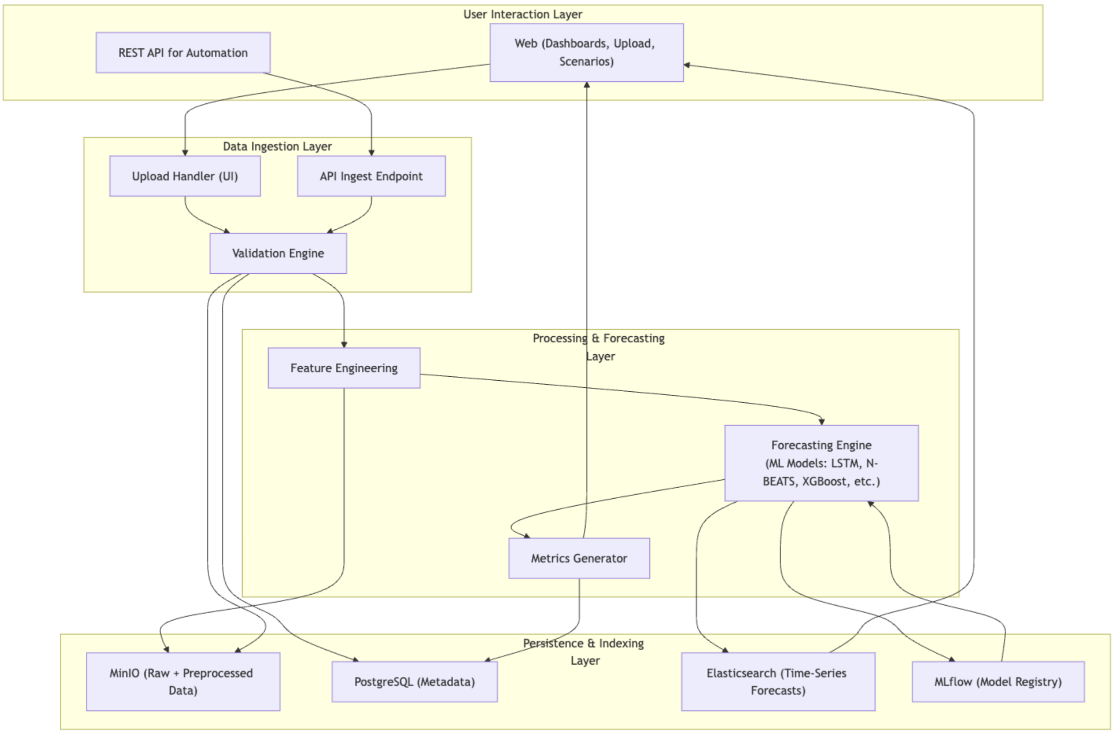

Figure 7: Interconnection of the main layers of DPM tool

The data ingestion layer supports both manual and automated input, enabling data to be uploaded directly by users or pushed into the system through external workflows and services. Data is uploaded via the web interface or pushed through the REST API, both of which store input files in the S3 MinIO object storage. A validation service verifies format, completeness, and structure before triggering the forecasting pipeline.

Sequence diagrams and data flow charts illustrate the transition from file upload to model execution and final result visualization. The architecture supports horizontal scaling through Kubernetes Horizontal Pod Autoscalers (HPA), enabling dynamic adjustment of resources under load. In addition, the final CI/CD Deployment Workflow for DPM tool is depicted in Figure 10.

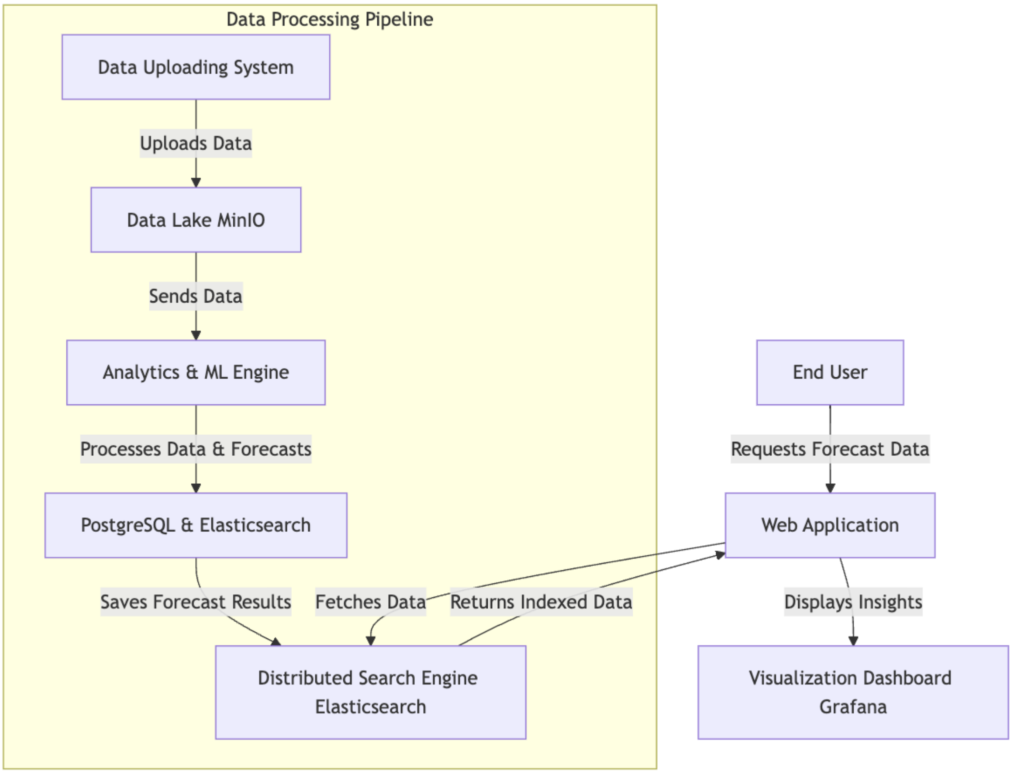

Figure 8: Final data processing pipeline architecture for DPM

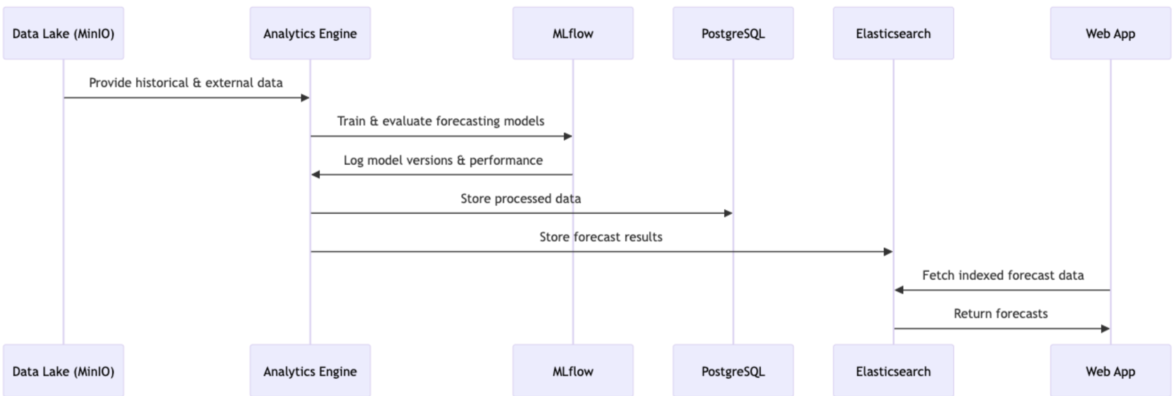

Figure 9: Forecasting Pipeline Execution Flow for DPM

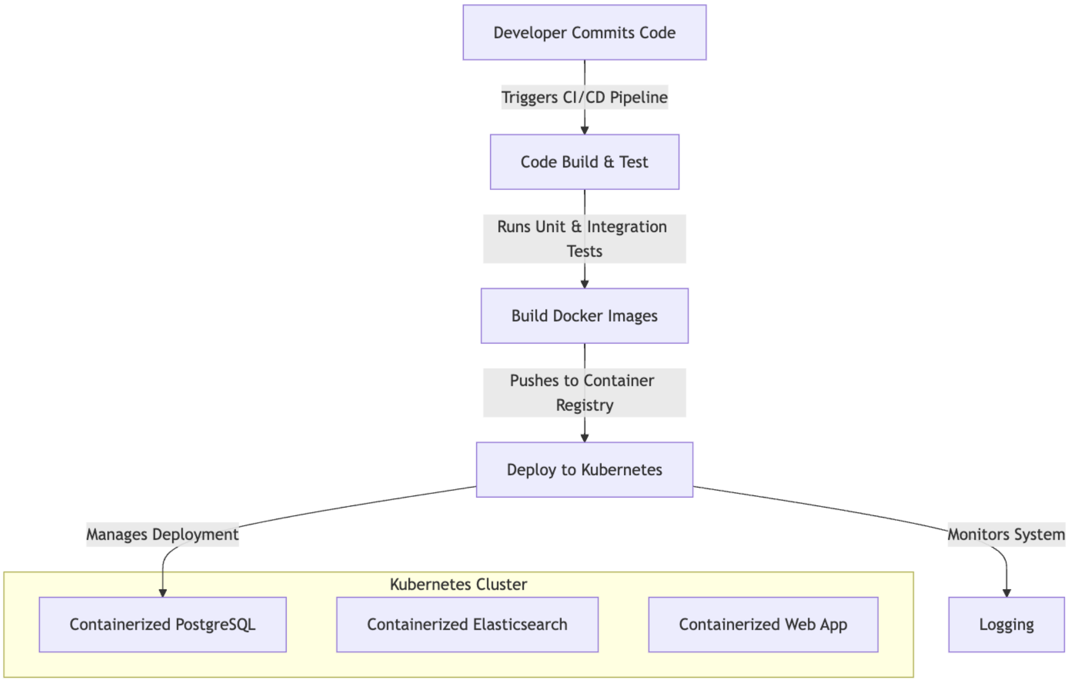

Figure 10: Final CI/CD Deployment Workflow for DPM

3.2.3 Performance Benchmarks

The performance of the DPM Tool has been validated through multiple pilot deployments across all target domains: energy, water, transport, and health. The evaluation focused on two main aspects: forecasting accuracy and system responsiveness.

In the energy domain, the Temporal Fusion Transformer [40] achieved a mean MAPE of 1.87% in national load forecasting for EKC, outperforming baseline models by up to 20%. Similarly, in the case of ELES, iterative modeling resulted in yearly accuracy improvements ranging from 7% to over 13%.

In healthcare forecasting, particularly for HQM during pandemic-related scenarios, the tool successfully delivered demand predictions under low, moderate, and high COVID-19 impact levels. These multi-scenario forecasts were validated by medical domain experts and formed the basis for strategic planning of medical resources and consumables.

In the transport sector, the tool's geospatial heatmap module rendered real-time demand projections with sub-second latency, leveraging Elasticsearch for fast indexing. This allowed operators to quickly identify high-pressure zones and respond dynamically to service needs. Backend response times for prediction requests remained consistently below 1.5 seconds for standard input sizes (~100,000 rows).

From a DevOps and deployment perspective, the tool is fully deployed in the cloud using Kubernetes. The stack includes PostgreSQL, Elasticsearch, FastAPI, MLflow, and a React-based frontend [22], all containerized for modular deployment. Full deployment can be completed in under 15 minutes in a clean environment, and updates are rolled out via Helm using zero-downtime rolling strategies.

Complementing pilot validations, the system is supported by an extensive suite of automated unit tests that cover critical services including file ingestion, validation logic, forecast execution, and API outputs. These tests run within the CI/CD pipeline on each commit, ensuring code reliability and preventing regressions during development iterations.

Overall, the DPM Tool has reached Technology Readiness Level 7 (TRL7), with over 80% functional integration achieved across participating pilot partners. It meets or exceeds the project's targets in forecasting accuracy, operational responsiveness, and end-user usability across all validated domains.

3.3 Cyber-physical resilience

3.3.1 Design Objectives

Considering the context of SUNRISE project with the impact of the COVID-19 pandemic on the cybersecurity solutions available in the market for the Critical Infrastructures as starting point and the business viewpoints presented by the different pilots covering various sectors, the first design objective was to provide a solution beyond the state of the art that could improve the overall resilience of the Critical Infrastructures and to have the ability to be adaptable and support the modification of operations based on changing conditions or emerging threats. By the combination of the Plan-Do-Check-Act (PDCA) [23] cycle (used as baseline of the SUNRISE Strategy process) with the OODA [24] (Observe, Orient, Decide, Act) loop decision-making which emphasizes the importance of contextualization of available information and an effective interpretation of new data in changing circumstances), the aim was to enhance our approach to risk management. The primary objective was to enable dynamic adjustment of risk baselines through observations of external environments and a cognitive orientation process. This approach should facilitate not only a best understanding of the current situation but also a better selection of the optimal decisions.

Provision of a modular and flexible solution for the enhancement of the Critical Infrastructure resilience from all these different perspectives was also a design goal. With this in mind, the Cyber-physical resilience tool was designed as a set of modules covering and improving a wide range of capabilities from real-time monitoring and detection of anomaly behaviours to automated responses to identified threats or security incidents, offering support to CI's decision-makers through a dynamic and complete risk assessment and not forgetting context awareness through the connection with threat intelligence data sources.

Considering that the ultimate goal is the deployment of an integral security solution into the Critical Infrastructures' premises, this disparity of functionalities provided by different modules should be capable of operating independently while also interacting with one another. Consequently, interoperability was also a design objective to ensure compatibility with existing security solutions available in the CIs and to facilitate integration and communication.

Cyber-physical resilience tool also has as design objective the usability and accessibility of the solution for the different type of Critical Infrastructure's end-users, with a user-centric design providing a unique dashboard that allows them to visualize at a glance the current status of the infrastructure but at the same time gives them access to more detailed information through specific dashboards for each of the modules.

3.3.2 System Diagrams

Figure 11 shows the information that is shared among the different CPR modules. Below the various steps are explained:

- Log data is received from the monitored infrastructure in the Critical Infrastructures as input for the Anomaly Detection module.
- Threat intelligence data received through MISP instances deployed in the Critical Infrastructures is the input for the Threat Intelligence module.
- Alerts and enriched indicators of compromises are inputs for the Risk Assessment module and the Incident Reporting module.
- The output generated by the Risk Assessment module is sent to the Incident Reporting module. Alerts generated during the application of countermeasures after early warning are also sent to the Risk Assessment module.
- All information generated by the different modules can be visualized by the different end-users' roles through the same CPR dashboard, which is connected with a Keycloak server for authentication and authorization management.

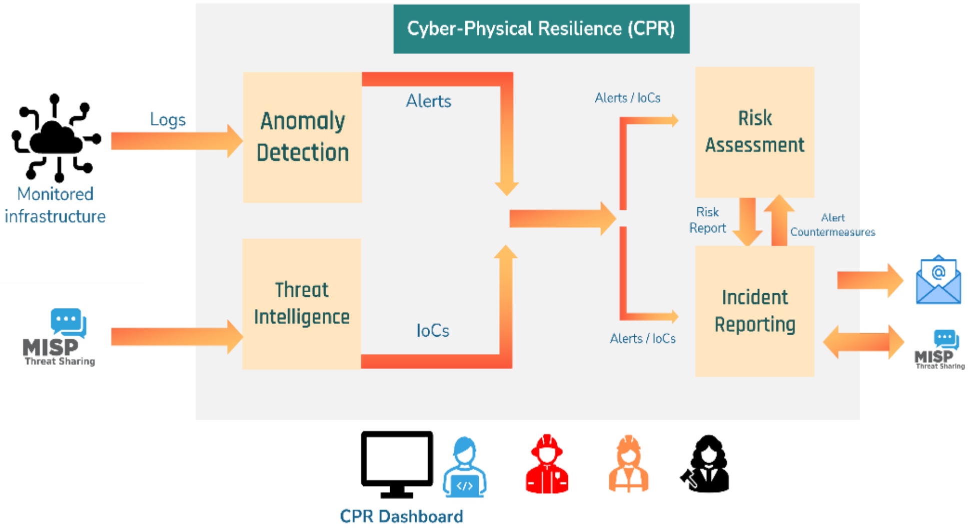

Figure 11: CPR System Diagram

Section 2.2.6 of D6.5 [19] describes with more details and an illustrative use case the workflow among the different components included in the CPR tool and the different users during a normal usage of the solution.

3.3.3 Performance Benchmarks

Since CPR tool is composed of different modules, validation of the performance of the CPR tool can has been done through the achievement of the Key Performance Indicators (KPIs) defined in the Grant Agreement [25] and reported in D6.4 [26] related to the improvement in more than 40% in the response time to a threat assessment and the improvement in more than 80% in the automation for threat assessment.

Through the execution of the second round of pilots and as reported in D6.6 [48], the CPR anomaly detection module was able to process datasets of almost 4 million log entries generated by a Slovenian Critical Infrastructure in a period of one year for training with the successful detection of anomaly behaviours. And it has been validated that the reception and triggering of the CPR risk assessment module with the generation of the risk report was done in near real-time. Although the processing time of the risk assessment depends on the number of risk models enabled, it could be demonstrated during the online demo session how the risk assessment (both, qualitative and quantitative) was increased when the alert was received. Automatic task generation and assignment throughout the full incident reporting workflow was provisioned by the CPR tool with minimal latency in the open-source incident response and management tool TheHive, and the response time is also on the order of seconds for the generation of incident report templates and their automatic notification via email to the service beneficiaries and competent-authorities channels configured. This was also validated during the demo online session and in the training videos generated about the CPR tool.

Deployment architecture was modified to reduce the deployment and docker build times. By reorganizing the dockerfile installation order, giving higher priority to the .R modules, which take a high percentage of the build time, when redeploying the CERCA tool after small modifications, configurations, re-deployment time is greatly reduced, thus contributing to the reduction of the response time to a threat assessment and automation for threat assessment.

3.4 Remote infrastructure inspection

3.4.1 Design Objectives

The primary goals and design principles of the RII Tool are clearly outlined below. With them, we describe what business objectives of the CIs they respond to. The latter were obtained with extensive research with CIs – additional meetings and thinking about their operation and needs in the completely new, unforeseen circumstances (like COVID-19) or otherwise.

The goal was to give the CIs a modular system and thus support easy technical integration into their respective environments. The modularity enables them to pick-and-choose the modules they find useful. Another goal was pursued: enabling the integration of legacy services, either into the SUNRISE system or integrating the modules into their own existing services. Given the CIs significant differences in terms of available computing and digital maturity, it was essential to pursue such an approach.

Modularity of the designed tool thus allows each module (satellite-based change detection/ vegetation monitoring, UAV-based object detection/3D reconstruction) to function independently or as part of an integrated pipeline. Tailored deployments depending on the digital maturity and needs of the CIs is enabled and ensures the tool remains practical and efficient in each specific context without burdening users with irrelevant features.

Given the CIs manage large, dispersed infrastructures it is crucial to enable wider and targeted views. The interplay, designed in the system, between satellite and UAV approaches, is thus essential and future-proofs the entire tool, despite the satellite methods not being useful at the time of this writing and limited to the resolution of the satellite images. The technology in the satellite imagery area improves and the RII tool already supports it, thus the CIs adopting RII tool will be already prepared. The same is true for the UAV-based methods – namely, with the advances in this technology, autonomous flights will become a matter of legislation and possibly, operational risk profile the CI will have to accept (i.e., what happens if the UAV accidentally fails during flight). This gives the CIs the necessary flexibility of the software and thus their operation, both, from integration and further development viewpoints.

The tool also emphasizes the user-centric design, with dashboards that provide infrastructure operators with intuitive controls, role-specific interfaces, and clear visual outputs, including alert overlays and timeline-based anomaly tracking, all while preserving the original inputs (image, video), for the operator to judge the possible defect and thus possible threat to business continuity, themselves. This directly answers also to the compliance of the tool, which requires for the tool, which uses the AI modules, human in the loop.

3.4.2 System Diagrams

The modular design of the RII Tool means that all of the three modules are connected with the industry standard MQTT bus. Their division is:

- The front-end GUI, operating via the MQTT bus to integrate satellite- and UAV-based modules
- The satellite imagery module
- The UAV module

The GUI login page is demonstrated in Figure 12.

Figure 12: The SUNRISE login page for the RII Tool.

The login supports two-factor authentication, needed for additional security, shown in Figure 13. We believe this is the most flexible way to add an additional layer of security, despite all modules and thus the whole RII Tool are to be installed at the CI premises.

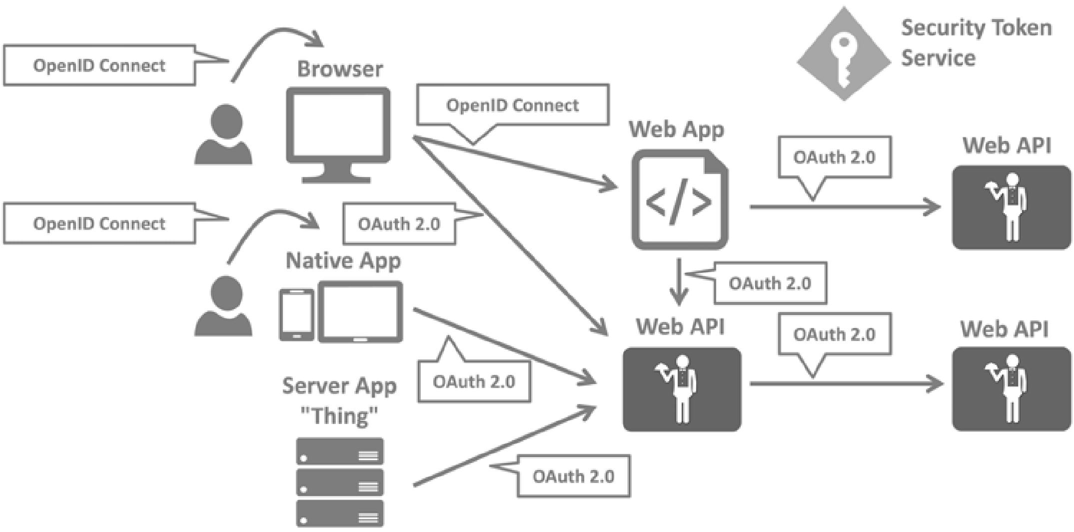

Figure 13: The two-factor authentication process.

The actual flow of messages between the web application (the GUI) and the other two modules is described (and numbered) below. The process is shown in Figure 14 and the description taken verbatim from D7.5 [27].

1. Incoming messages/events from UAV/Satellite systems are received. All this data is routed through an MQTT bus system. Within this system, the data is systematically queued, ensuring a sequential flow.
2. The Backend Coordinator processes all incoming messages/events. It retrieves the data at the front of the MQTT queue.
3. All incoming messages/events are internally stored in the Backend Inventory (MongoDB server).

4. The Backend Coordinator sends live or historical data to the Dashboard UI for visualization and responds to historical data requests from the Dashboard UI.
5. The Dashboard UI communicates with the Google Maps infrastructure to render maps, markers, points of interest, and heat maps, among other elements.
6. The Backend Coordinator sends requests to the Reporting Subsystem in order to compile the requested data and then receives the results.
7. The Reporting Subsystem and the Backend Inventory communicate with each other in order to process the requests and subsequently transmits the results to the Backend.
8. The Dashboard UI obtains an Access Token from the Identity Server to access backend APIs. Access to the UI is exclusively granted to authorized users, with authentication and authorization handled by a dedicated Authentication/Authorization unit, responsible for controlling user access and logging into the application.
9. Additional public services can offer crucial meteorological data, weather forecasts, maritime information, alerts, and more for visualization within the Dashboard UI.

It is important to note that the connection between the Backend Coordinator and the MQTT system is bidirectional. If any data needs to be transmitted from the application outward, the Backend Coordinator places it in MQTT, within the corresponding queue.

Figure 15 and Figure 16 show, how the two modules (presented in the high-level architecture diagram in Figure 14 connect to GUI in a completely distributed fashion.

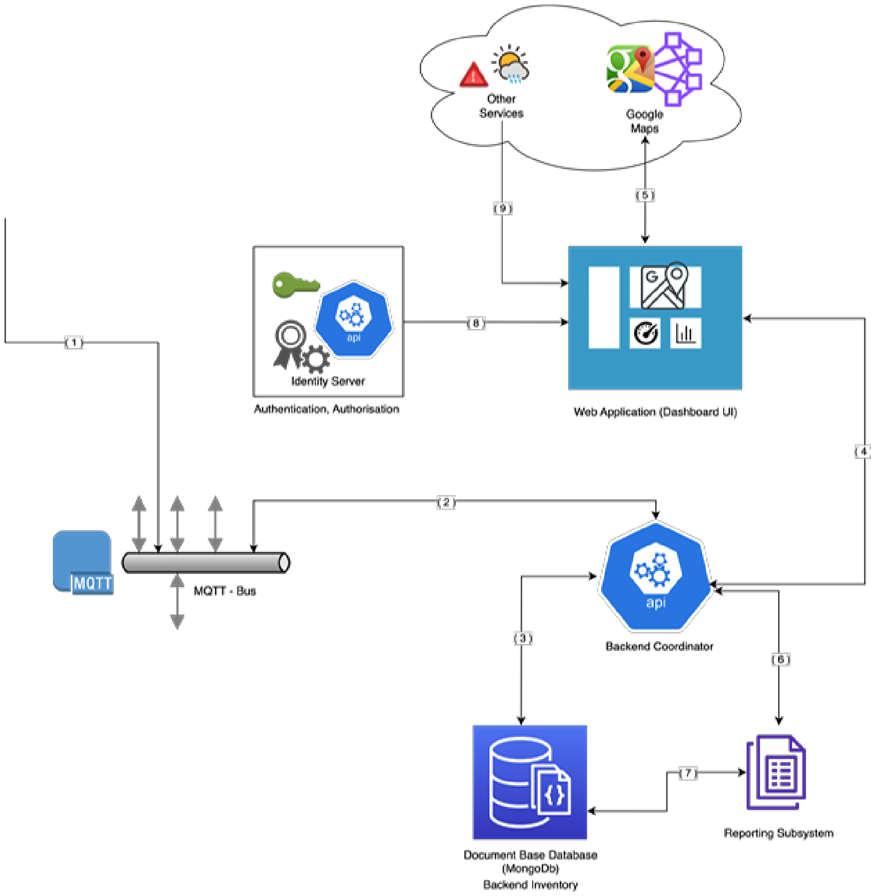

Figure 14: The GUI Architecture Diagram.

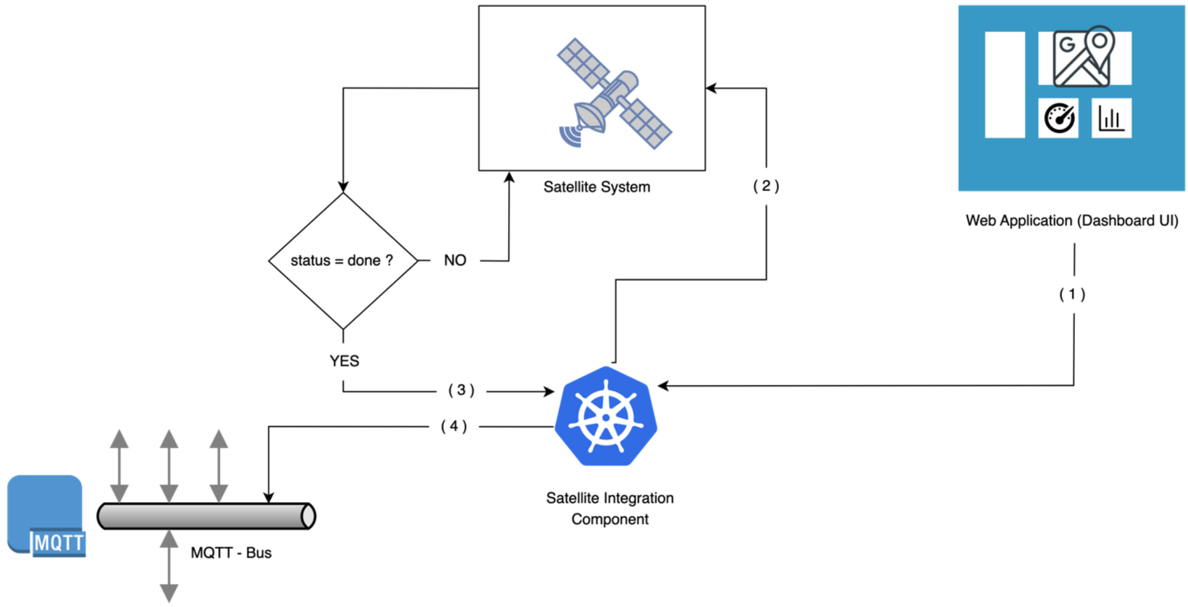

Figure 15: The connection, via MQTT, between GUI and the satellite module.

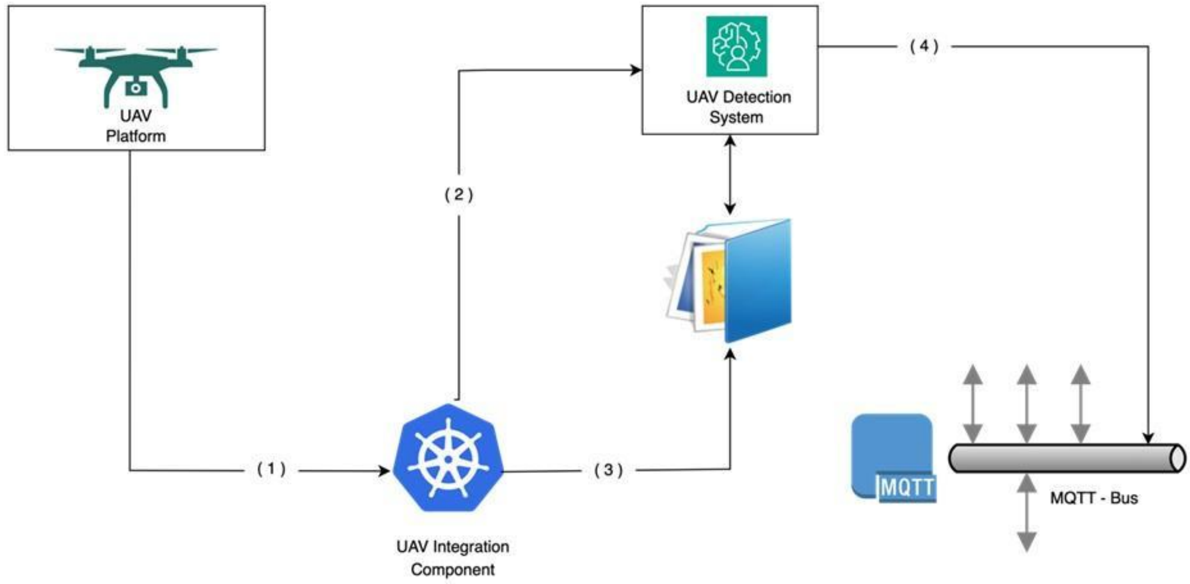

Figure 16: The connection, via MQTT, between GUI and the UAV module.

The internal diagrams of the two modules are presented in the following two paragraphs. They are still rather high-level, as their in-depth presentation is given in D7.5 [27].

Satellite module: The high-level presentation is given in Figure 17. More precise workflow, required for modelling the vegetation height and changes from the satellite images is shown in Figure 18.

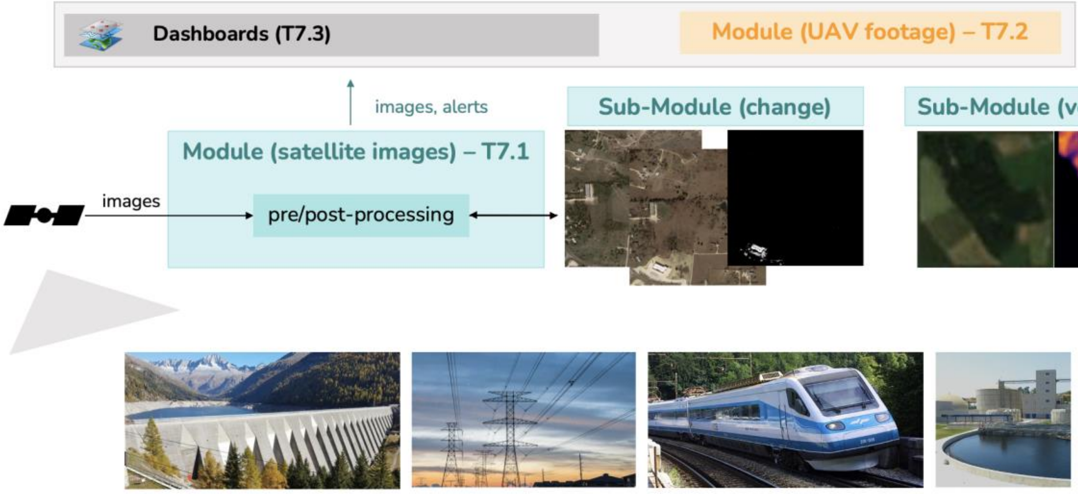

Figure 17: The high-level architecture of the satellite module.

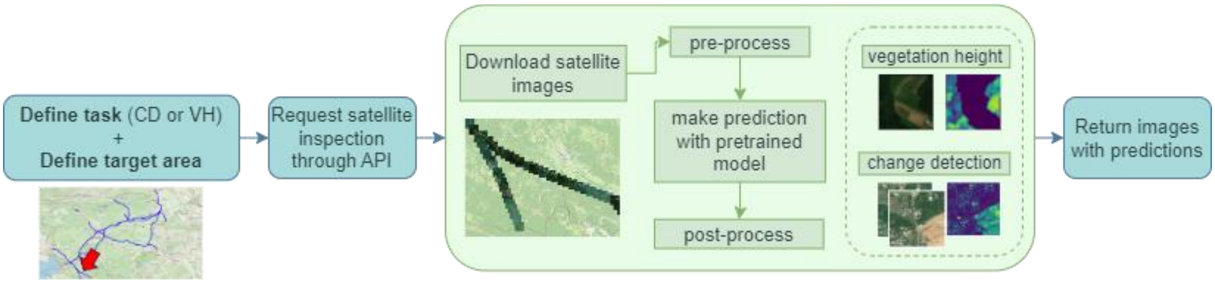

Figure 18: Detailed workflow for satellite images based AI submodules for detection of the vegetation height and change detection.

Figure 19 and Figure 20 demonstrate the UAV module based RII Tool which provides either a capability of real-time inspection (when there is a possibility to directly stream the video to the NVidia Jetson Orin device in the field) or to perform a batch check of the video, captured during several flights. More detail is found in the D7.5 [27].

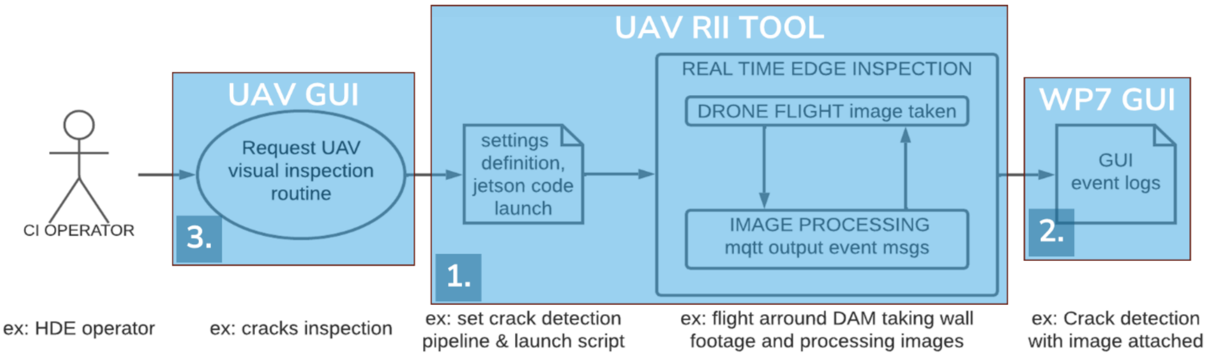

Figure 19: UAV based RII Tool, focusing on real-time inspection.

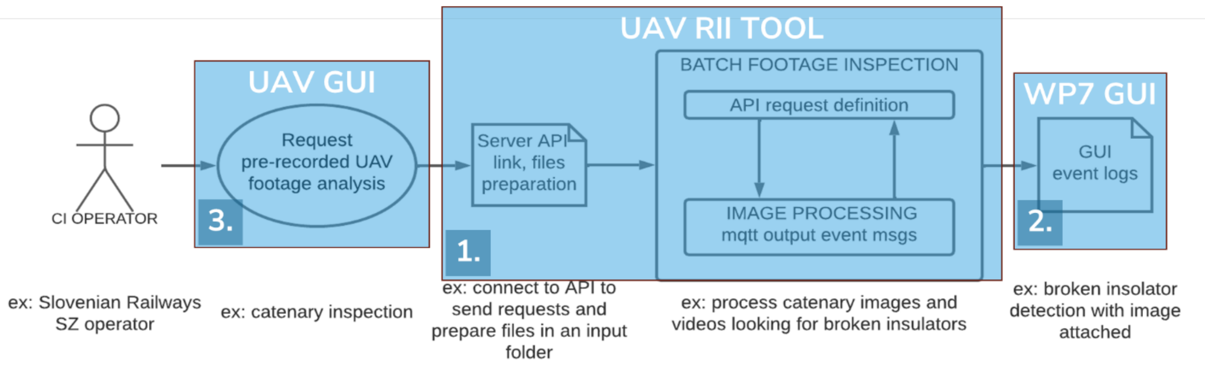

Figure 20: UAV based RII Tool, focusing on batch inspection.

3.4.3 Performance Benchmarks

In the development of the RII tool we had objectives that the tool must be:

- Flexible (allowing different adoption scenarios and integrations)
- Resource efficient,
- Trustable (we provide the results obtained on the known test data and also perform 2 pilots to validate the system and its modules)
- Mitigating the possibility to identify personnel on videos taken by UAV.

Here, we note that most of these results and tests were already provided as part of the D7.1 [29], D7.2 [28], D7.3 [30], D7.4 [31] and D7.5[27].

The identification of personnel is addressed with the implementation of a deep learning-based detection system TinaFace [32], an AI model recognized for its real-time state-of-the-art accuracy in face detection tasks. A comprehensive review of existing models confirmed that TinaFace offers the best performance across major benchmarks, making it the optimal choice for this solution. Specifically, TinaFace achieved an Average Precision (AP) of 0.97 on the WIDER Face [33] dataset (easy), 0.963 on WIDER Face (medium), and 0.934 on WIDER Face (hard). Tinaface ResNet-50-based architecture provides high efficiency for real-time applications, ensuring robust detection under various lighting conditions, angles and occlusions. This capability makes it highly reliable for UAV-based inspections, where environmental factors can vary significantly. As shown Figure 21, TinaFace ranks #1 in the WIDER Face benchmark, demonstrating its superior detection capabilities.

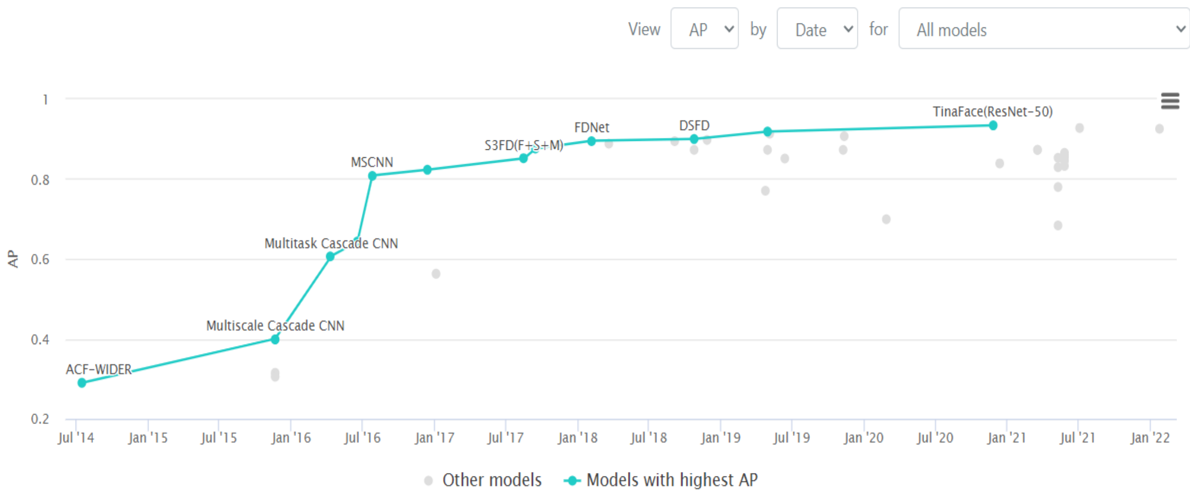

Figure 21: Face detection models leaderboard on WIDER Face (Hard)

The results of the tests on the Pilot 1 data are available below. For the results on the test datasets, please refer to the previous WP7 deliverables (D7.3 [30], D7.5 [27]).

Firstly, we address the applicability of the UAV based AI assistance method directly in the field. The use of the AI-based methods in the field requires stable internet connection to the operator's personnel in the field, which can be achieved either through standard 5G or using Star Link, when there is no mobile signal coverage. This has been achieved with the use of the modules on the UAV platform itself (see the D7.5 [27]). On the other end, we set out to achieve the resource-efficient edge deployment of the AI model, capable of segmentation and recognition of the possible infrastructure problems. This was achieved using the standard Jetson Orin AGX 32GB system. We thus achieved the practical ability to use AI-based reasoning about the state of the infrastructure, while executing the practical piloting in the field.

The performance of the VQA methods on the real data (real examples of pilot sequences) is shown in Figure 12 and Figure 13. All the data necessary to understand in detail each of the image analysis pipelines are developed in the previous deliverables, in particular in (D7.3 [30], D7.5 [27]). all solutions being based on the combination of the 3 AI submodules: detection, segmentation and VQA.

The tool has been tested by defining several pipelines dedicated to check the correct state of the power lines and catenary networks of trains, as well as specific elements of them, such as ceramic isolators. For this application, it is important to analyse both visible and infrared images, and therefore this is defined in the processing pipeline. This way, on the one hand, Figure 22 shows an example of a general state analysis of an electrical pylon using thermal and visible spectrum imaging as input to the VQA module.

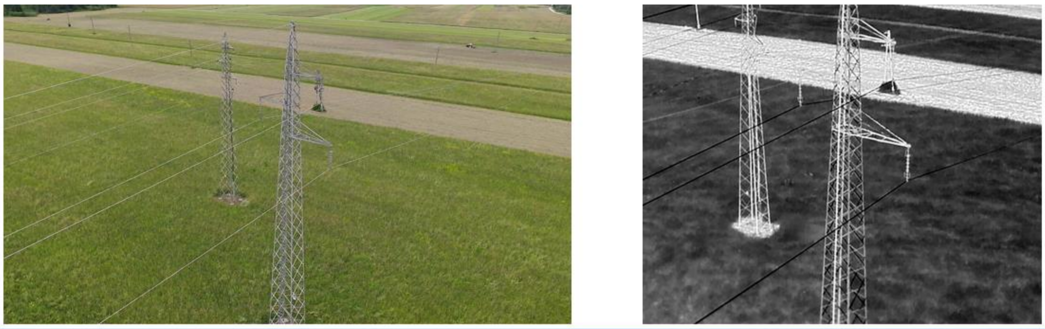

PROMPT: "Use RGB <|image_1|> and thermal <|image_2|> to describe the status of the electrical pylon. Please, justify your answer properly, step by step. Give a clear final conclusion of your analysis."
VQA: To assess the condition of the electrical pylon, I will carefully examine the image and look for any signs of wear, damage, or potential hazards. \n\n1. The pylon appears to be in good condition, with no visible signs of rust, corrosion, or structural damage.\n2. The wires are intact and properly connected to the pylon, indicating that the electrical system is functioning correctly.\n3. The pylon is standing upright and appears to be stable, suggesting that it has not been compromised or damaged.\n4. There are no visible signs of damage to the supporting structures or the ground beneath the pylon.\n5. The pylon is not leaning or tilted, which is a positive sign of structural integrity.\n\nBased on these observations, I can conclude that the electrical pylon is in good condition and does not require immediate attention. However, a thorough inspection should still be conducted to ensure its continued safety and reliability

Figure 22: UAV remote inspection tool example of analysis of an electrical pylon captured at the ELES facility.

As it can be seen from the response provided, the model is able to reasonably analyse the correct state of maintenance of the pylon, concluding that it is in good condition, with no rust, structural damage, damaged wires, or tilting.

On the other hand, Figure 23 shows two examples of the performance of the current solution applied to the inspection of specific elements of the electrical network. These images have been taken in the ELES pilot, but examples from SZ can also be seen in deliverable D7.3 [30].

PROMPT: "Is there any broken or damage insulator in <|image_1|>?"

VQA: "Upon close inspection, there is no visible damage to the insulators in the image. All insulators appear to be intact and undamaged."

PROMPT: "Is there any broken or damage insulator in <|image_1|>?"

VQA: "The image shows a close-up view of a metal structure, likely part of a larger infrastructure such as a power line or a communication tower. There is a bell-shaped insulator hanging from a vertical support, which is designed to support the tension of the cable or wire it is attached to. The insulator appears to be in good condition without any visible damage or signs of wear. The background is a grassy field, indicating that the structure is outdoors."

Figure 23: Two examples of UAV remote inspection tool applied to check the integrity of ceramic isolators in ELES infrastructure.

The previous examples, together with the rest of the analyses carried out and content included in previous deliverables, allow us to affirm that the tool developed fulfils the desired functionality in this first group of requirements.

For the satellite-based methods, we describe the tests executed in Pilot1 below, taken from D7.4 [31].

For the vegetation height, a meaningful evaluation requires comparison between the ground truth vegetation height and predictions, provided by the vegetation height model. During Pilot 1 two samples of LiDAR footage provided a point cloud model of the surface, from which we calculated true vegetation height. Using Whitebox Tools [34], we converted the point cloud into a Digital Terrain Model (DTM) representing the bare surface and a Digital Surface Model (DSM) representing canopy tops. We obtained a vegetation height model in meters by subtracting DTM from DSM. The ground truth was down-sampled to a spatial resolution of 10m to match the resolution of the Sentinel imagery.

To obtain predictions, we downloaded multiple Sentinel satellite images within 6 months of LiDAR capture date. Our pretrained model (described in D7.3 [30]) was used to make the prediction for each sample, and each prediction was then compared to the ground truth.

The first LiDAR sample was recorded along railway tracks at the Pilot 1 Trials of this project. The second was provided by ELES. See Table 3 for further details on the captured data.

Table 1 : Properties of the two LiDAR samples

	Sample 1: Railway Tracks (SŽ) – Pilot 1 Trial	Sample 2: Power Lines (ELES)
Date of LiDAR retrieval	June 21, 2024	2017
Size of area	700m long, 120m wide corridor	8km long, 170m wide corridor
Sentinel imagery available	4 images available in time frame April1, 2024 – July 1, 2024	5 images available in time frame April 1, 2017 – September 1, 2017

Figure 24 shows vegetation height predictions for Sample 1 imagery. Multiple satellite images were collected to evaluate the consistency of the model. Note that the standard deviation is, in general, low, with some exceptions. The average of all predictions was taken to compare to ground truth. Average absolute pixel-wise error is 4.21 and is presented in Figure 25. The error is extremely high (up to 20m at the edges), which can be evaluated as very likely, due to inaccuracies in the ground truth height further from the drone's flight line. This means that the drone's flight line must be aligned with the need of the LIDAR capturing, if it is performed.

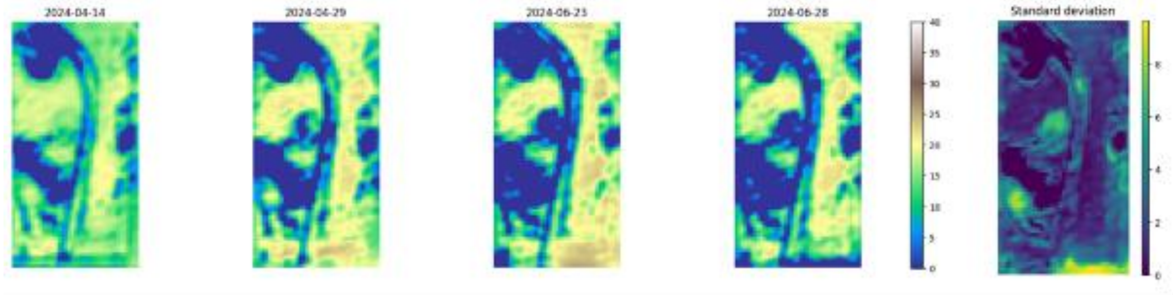

Figure 24: Vegetation height predictions of the 4 Sentinel images of Sample 1 area and standard deviation between them

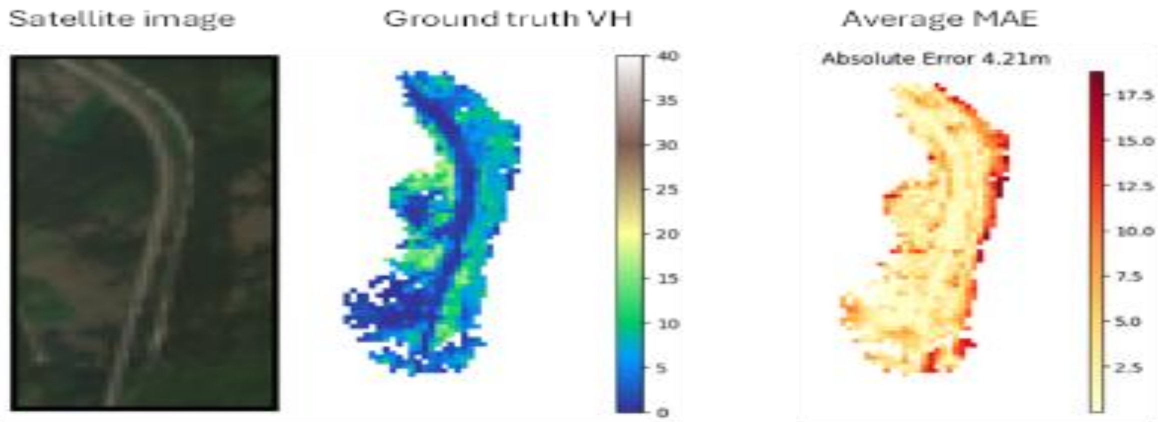

Figure 25: Sample 1 results. Left: Sentinel imagery, Middle: Ground truth vegetation height, Right: MAE of average prediction

Sample 1 covers a very small area with inaccuracies, while Sample 2 provides us with better insight into our model performance. In Sample 2, the evaluation was made for a narrow corridor (170m wide) with available LIDAR data, however, the prediction was done for a larger polygon.

In this case, 5 Sentinel images were available in the Summer of the year the sample was collected. Predictions for each available image are shown in Figure 26. The first image is a clear outlier, causing the standard deviation to be high. This might be due to the weather conditions at the time of image capture. This suggests that taking an average prediction over multiple collections provides a more accurate result. The average pixel absolute error is 2.8m.

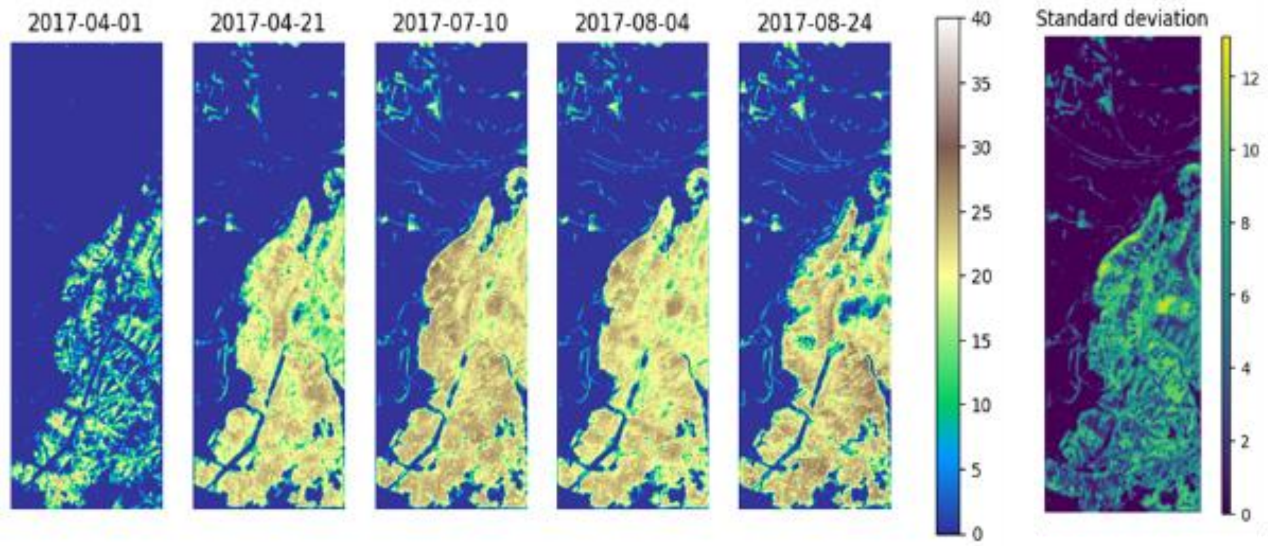

Figure 26: Predictions, using the vegetation height model, for all 5 available Sentinel images

Figure 27 gives further analysis into the error based on the tree height. The absolute error is very small (<2m) for low trees but increases for taller vegetation. Notably, the number of trees taller than 25m is very small, suggesting that the high error for these trees might not be representative of overall accuracy. The absolute error for trees up to 25m is below 6m. While mean absolute error is the better metric to evaluate the predictions, mean error gives a deeper understanding of the error. The model overestimates vegetation up to height 15m and underestimates higher trees.

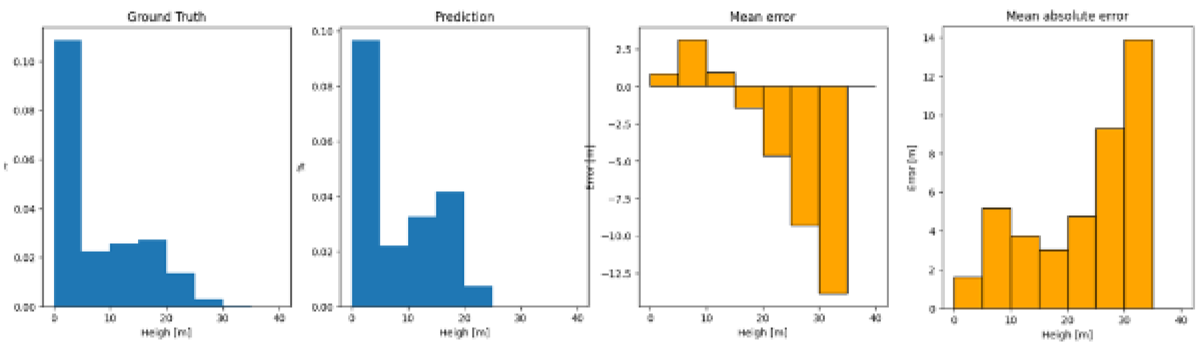

Figure 27: Comparison of ground truth, predictions and presentation of errors

For anomaly detection with the satellite imagery, we obtained a smaller dataset of detected changes by partner CI ELS. These were detected by their infrastructure inspection personnel and reported in a data log – there are around 100 entries. We removed cases, where precise location was not reported and then proceeded to obtain the before and after Sentinel satellite imagery. This imagery was used to evaluate the model.

In Figure 28, a pair of such images is showcased. The event date (event being the change detected by ELS personnel) was 14.02.2018. The before image, on the left, is one month before that date, and the after image, in the center, is one month after that date. The image on the right shows the prediction of Open Data Infrastructure for Security and Emergency Management (ODISE) model [35], given the before and after image, where the white pixels indicate locations with more than 20% predicted change.

Figure 28: Before and after Sentinel satellite images of location of detected change by ELS personnel

The model has predicted something, but it's difficult to determine if there's an actual change there—even by eye, it's difficult to see, mostly due to the low resolution of 10m. We came to a similar conclusion for other pairs of satellite images – namely, apart from having only a small set of usable (with usable meaning they don't have issues described in the previous section) satellite image pairs (around 25), the changes were difficult to notice, even by eye. Hence, the set of useful changes was small and is too small to declare the value of this KPI with sufficient statistical robustness. Thus, we must declare the value of this KPI as not yet completely assessed within Phase I piloting.

Similar to the case of vegetation monitoring, we explored the other aspects of assessing feasibility of **Change detection.** The satellite imagery providers are described in the previous section and hold also in case of change detection. Here, we evaluate suitability only for change detection. Only the SkySat [36] images were found to be suitable – smaller resolution isn't suitable for small-scale changes CIs are interested in detecting, as can be seen in the above discussion regarding changes detected by ELS personnel. However, as mentioned previously, SkySat imagery is not suitable for long, linear infrastructure. It is possible to use SkySat high-resolution imagery for regular monitoring, but it is very costly.

Thus, the recommended approach is to use the medium resolution PlanetScope [37] (3.7m) for wide area monitoring and then task a SkySat to investigate a suspect location in more detail. Furthermore, tasking long linear corridors takes weeks to months, not days. Another issue here is scarcity of data. The archives of SkySat are limited. Additionally, to train and evaluate the models, we need imagery of the same area over multiple collection times.

4 Best Practices

4.1 Risk-based access control

4.1.1 Development Practices

RiBAC development followed iterative cycles with regular stakeholder feedback integration through regular sprint reviews with CI operators. This approach ensured development priorities aligned with operational requirements and emerging health guidelines while maintaining technical excellence.

A comprehensive quality assurance framework included unit testing for individual components, integration testing for module interactions, and end-to-end validation in simulated CI environments.

Hardware-software co-design ensured optimal performance within power and processing constraints. AI model optimization specifically targeted the Jetson Nano platform capabilities while maintaining accuracy requirements, demonstrating the effectiveness of integrated development approaches.

4.1.2 Integration Guidelines

RiBAC integration follows a standardized approach beginning with legacy system assessment, interface mapping, and gradual feature activation. Standard protocols ensure broad compatibility while custom adapters address specific vendor requirements. Secure network integration uses VPN tunneling for remote management with local processing to minimize bandwidth requirements and ensure operation during network outages.

The system offers five distinct integration levels to accommodate varying requirements and existing infrastructure capabilities:

- Level 1 provides autonomous, offline operation compatible with company-native identification cards, requiring minimal infrastructure changes.
- Level 2 extends this with event generation capabilities including screen notifications, sound signals, and door lock actuator management for enhanced operational control.
- Level 3 introduces semi-integrated functionality with the ability to download access rights for existing identifiers and export captured transactions, enabling better coordination with legacy systems while maintaining operational independence.
- Level 4 represents full integration with legacy systems, providing seamless data exchange and unified management capabilities across the entire access control infrastructure.
- Level 5 offers specialized functionality demonstrated through implementations like the UKC Risk-based Exit Control (RiBEC) system, where RiBAC principles are adapted for specific operational requirements, such as controlling egress from sensitive areas like operating rooms.

This flexibility ensures that each pilot site can implement the most appropriate integration level based on their specific operational needs, existing infrastructure, and regulatory requirements.

4.1.3 Testing and Validation

Validation was conducted across 9 sites representing different CI sectors (healthcare, transportation, energy, telecommunications) ensuring broad applicability and identifying sector-specific requirements. Quantitative testing validated all functional requirements with specific KPIs including detection accuracy (>80%), processing speed (<15 seconds), and system reliability (>95% uptime) [21].

Structured user acceptance testing with 47 CI employees across different roles and technical backgrounds provided valuable feedback for usability improvements and training requirements. Comprehensive security testing including penetration testing, data protection validation, and third-party security audit confirmed adherence to cybersecurity frameworks and data protection regulations.

4.1.4 Maintenance Plan

The maintenance strategy centers on regular system health checks including camera calibration, temperature sensor verification, and software update deployment. Predictive maintenance algorithms monitor component performance and alert operators to potential issues before failures occur.

Remote support infrastructure enables 24/7 monitoring and proactive issue resolution. Secure remote access allows technical teams to provide support, configuration updates, and troubleshooting without on-site visits in most cases. The modular design enables selective component replacement and upgrade without system-wide impacts, while standardized interfaces ensure compatibility with evolving technologies.

Comprehensive training programs for CI operators include basic troubleshooting, routine maintenance procedures, and system administration. Online training modules and video guides supplement on-site training for ongoing staff development [21].

4.1.5 Lessons Learned

First round pilot testing revealed the importance of comprehensive site surveys before and after deployment.

CI operators particularly valued the ability to quickly adjust screening parameters based on changing health guidelines or risk levels, highlighting the need for operational flexibility in future versions. While technical performance met all requirements, pilot sites indicated that RiBAC brings clear value through reduced manual screening costs, improved compliance documentation, and enhanced security capabilities.

Large CI installations highlighted the importance of centralized management capabilities for configuration synchronization, reporting aggregation, and maintenance scheduling. These insights will inform future development priorities and business model optimization.

4.2 Demand prediction and management

4.2.1 Development Practices

The DPM Tool was developed using a modular, cloud-native architecture to ensure flexibility and maintainability. Development followed agile principles, with regular iteration cycles, pilot feedback sessions, and integration milestones. Each component—data ingestion, preprocessing, prediction engine, and visualization—was designed as a standalone microservice, containerized and deployed on Kubernetes in the cloud.

Development relied on modern, open-source technologies such as FastAPI, PostgreSQL, Elasticsearch, and MLflow. Code was managed in version-controlled repositories, with automated pipelines for testing and deployment. Continuous Integration (CI) ensured that each feature branch was tested against a comprehensive unit test suite before merging, while Continuous Deployment (CD) allowed fast iteration and updates in the live environment with minimal manual intervention.

The AI models, especially the Temporal Fusion Transformer and LSTM-based predictors, were developed and trained using historical CI-specific datasets. Each model was versioned using MLflow, allowing reproducibility, rollback, and side-by-side performance comparison during tuning phases.

4.2.2 Integration Guidelines

To simplify integration with CI operator environments, the tool was designed with a clear API-first philosophy. The RESTful backend exposes all key operations—data upload, prediction launch, forecast retrieval, and metrics querying—using standardized interfaces documented via OpenAPI.

CI operators can connect their own systems directly to the backend or interact via the intuitive web interface. Both manual and automated data ingestion paths are supported: CSV and JSON file uploads via the UI, and programmatic file transfer through a dedicated API endpoint. Uploaded files are stored in S3 MinIO object storage and validated before being processed.

Authentication and access control are handled via Keycloak, which enables integration with institutional identity providers and support for multi-tenant deployments. The cloud-hosted nature of the tool allows it to be accessed securely from CI operator environments without requiring extensive local infrastructure.

Deployment is fully containerized and delivered using Helm charts, making it easy to install or update on any managed Kubernetes cluster.

4.2.3 Testing and Validation

The solution is supported by a comprehensive testing framework. Unit tests were developed for every major component, including data ingestion services, preprocessing logic, model inference pipelines, and the REST API layer. These tests are run automatically in CI workflows and are designed to catch regressions, misconfigurations, and interface changes early in the development cycle. You may refer the reader to D5.4 [38] for further details, where all validation details are given.

In addition to automated unit tests, end-to-end integration tests simulate realistic user workflows—such as uploading a dataset, running a forecast, and evaluating results. The testing framework ensures consistency across versions and protects against unexpected behaviours during model updates or backend upgrades.

All forecasting models are validated against historical datasets provided by CI operators. Accuracy is measured using established metrics such as MAE, MAPE, RMSE, and SMAPE, and results are benchmarked against legacy systems or baseline statistical models. For each domain, at least two pilot deployments were used to validate system behaviour under real-world operational conditions.

4.2.4 Maintenance Plan

Maintenance of the DPM Tool is streamlined by its modular design and cloud-native deployment. Each service runs in its own container, which can be monitored, restarted, or updated independently. Kubernetes health checks and autoscaling policies ensure resilience and load balancing under varying demand.

Automated monitoring (via Grafana) and logging (via Elasticsearch) provide visibility into the system's performance and health. Alerts are configured for key failure points—such as data pipeline errors, prediction delays, or degraded service response times—allowing proactive incident management.

Model retraining is supported through a retrain pipeline and MLflow integration, allowing operators or data scientists to update predictive models periodically or as new data becomes available. All training runs are logged, versioned, and tested before deployment.

Comprehensive documentation is provided for deployment, operation, and extension. This includes user guides, configuration references, and training materials tailored for technical staff, analysts, and decision-makers across different CI domains.

4.2.5 Lessons Learned

Several lessons emerged from the development and deployment of the DPM Tool. First, ensuring high-quality historical data was more challenging than anticipated. Inconsistent formats, missing records, and varied data availability across CI operators required significant effort in data cleaning, standardization, and validation logic.

Second, while most CI operators initially preferred manual interaction via the dashboard, demand for automation and integration grew quickly. The dual design—combining a user-friendly interface with powerful backend APIs—proved essential for long-term adoption.

Third, cloud deployment brought substantial benefits in terms of scalability and maintainability but also introduced concerns around data privacy and access control. These were addressed by hosting the solution within secure, isolated namespaces, using tenant-specific authentication, and encrypting all data flows.

Finally, the inclusion of unit and integration testing from the early stages of development greatly improved confidence in the system's stability. Automated testing and cloud-native deployment allowed the team to release frequent updates without service interruptions—a major factor in the tool's successful pilot adoption.

4.3 Cyber-physical resilience

4.3.1 Development Practices

Cyber-Physical Resilience tool has been conceived as an integral solution with different modules developed independently to cover the different capabilities required and a unified CPR dashboard that provides access to all the summary information and the different dashboards offered by each module.

Communication among the modules is done using a Kafka message broker. In this way, standard MISP events JSON are sent by the CPR threat intelligence module to a specific Kafka topic and standard Wazuh alerts are sent from the CPR anomaly detection module. During the project the two partners involved in the development agree on the topics to be used and on these formats compatible with the different modules and exportable if required to integrate with other cyber security solutions present in the pilots.

A Continuous Integration environment based on the open-source GIT management software platform GitHub [39] has been used for the software development lifecycle, allowing improved code quality with automated testing and builds, collaboration and transparency among developers and better development integration of changes.

4.3.2 Integration Guidelines

As described in the previous section, the different modules included in the Cyber-Physical Resilience tool works as standalone and communication is done through a Kafka message broker. Alerts generated by the CPR anomaly detection module are sent to a specific topic and MISP events with threat intelligence data processed and enriched by the CPR threat intelligence module are sent to another topic. Both topics are consumed by the CPR risk assessment module and the CPR incident reporting module to trigger their respective processing. The resultant risk report generated by the CPR risk assessment module is also sent to a specific topic to be used by the CPR incident reporting module.

Communication among the different components included in each of the modules and between them and the open-source tools they interact with (e.g. the service aire-thehive-plugin included in the CPR incident reporting module with TheHive and with the services aire-reports-generator and aire-workflow-enforcement, or the service tinted-orchestrator included in the CPR threat intelligence module with the MISP instance and with tinted-privacy) is done through REST APIs.

All the components have been developed to be deployed as dockerized containers so they can be easily integrated with different environments in the Critical Infrastructures.

Considering software developed is proprietary, each partner involved in the development of the CPR tool has used its own GitHub repository for software configuration management.

4.3.3 Testing and Validation

The different modules in the Cyber-Physical Resilience tool have been tested and validated by the developers in their premises using logs generated and provided by the different CIs' pilots during piloting phases. Additionally, during the second piloting phase, CPR tool has been deployed in some of the pilots' premises to test and validate deployment and execution integrated with their infrastructures.

Validation of the KPIs defined for the CPR tool has been carried out through 4 different pilots across different CI sectors (including healthcare, water, transport, energy and telecommunications) to identify potential sector requirements while ensuring the applicability of the solution. The level of development and integration achieved during each piloting phases, along with the validation of the specific requirements associated to the CPR tool and end-user feedback (e.g. regarding tool usability) was gathered through surveys and reported in deliverable D6.5 [19].

4.3.4 Maintenance Plan

CPR tool has been designed as a flexible and modular tool for better future adaptation, extension and integration with different CIs' environments. The different CPR tool components are provided as docker images which facilitates its maintenance since any update required in some of the modules will only imply to update the corresponding image(s).

CPR tool integrates a set of risk patterns (with a specific set of indicators) and templates for incident reporting. They could be extended in the future to support new or more complex attack patterns as well as updates in the current regulatory framework or new directives. If they include the same indicators and information, they can be easily added through configuration, but if they require new ones, some consulting service could be required from the technical partners involved in the development of the tool.

Extensive documentation including tool deployment and operation has been provided together with online demo sessions and training videos explaining how to use the main features provided by the CPR tool. All this training material ensures that CI's stakeholders and end-users can effectively adopt, configure and use the tool within their specific operational contexts.

4.3.5 Lessons Learned

The usage of the CPR tool across the various pilot in different sectors has confirmed its effectiveness in enhancing cybersecurity situational awareness for the Critical Infrastructures analysts, as reported in D6.4 [26]. However, the complexity of the solution due to its own nature covering a wide range of functionalities and allowing adaptivity, makes that the usability needs to be improved with simplified configuration procedure. And additional training sessions can be required to increase the understanding on risk modelling ad assessment and achieve a more effective use of the tool and higher impact on an operational environment. Another important lesson learned from the deployment done in some of the operational environments with a huge number of devices in its infrastructure, is the need of assets aggregation. Unlike tests done in laboratory environments, more work needs to be done for improving the presentation of the results and for reducing the processing time during the evaluation of the risk assessment. In this sense, during the last pilot phase it has been already added a risk assessment simulator that provides a light execution of the R models.

Finally, it has been also learned that it is important to clarify from the beginning the operational requirements for the tools in order to obtain results. For example, since the CPR anomaly detection module (LOMOS) only focuses on human-readable logs [19], logs provided initially by the CIs were not compatible with tool and useful to identify anomalies.

4.4 Remote infrastructure inspection

4.4.1 Development Practices

All three partners developed the modules and submodules of the RII Tool independently, communicating only through APIs, thus creating a distributed application which satisfies the need for flexibility, meaning that the CIs can pick and choose what is the functionality that will be integrated into their processes.

The development was internal to all partners and included internal testing, mostly in the MLOps pipeline, as we were considering the accuracy of the AI modelling.

4.4.2 Integration Guidelines

Integration of all the modules is implemented through APIs, while the choreography for the front-facing GUI uses the MQTT.

To ensure secure, reliable, and low-latency communication between all system components—including the UAV control station, Jetson devices, centralized servers, and the GUI—a ZeroTier private VPN has been implemented. This VPN ensures encrypted and stable connectivity without additional manual authentication or configuration, contributing significantly to the ease of deployment and integration.

In all deployment scenarios, the process is standardized through Docker containerization, ensuring consistent and reliable operation across varied computing environments.

Software management and updates are centralized via a GitHub repository, enabling seamless synchronization, version control, and simplified deployment procedures without operational disruptions.

This results in a system that is easily transferrable among different infrastructures (either distributed or centralized on the CIs infrastructure).

4.4.3 Testing and Validation

The components were tested and validated internally. The integrated components were tested using sample videos, images and in the Pilot 1 and Pilot 2 iteration.

Testing and validation of the RII Tool was performed after each change that required changes of APIs between components, e.g. inclusion of different presentation for change detection or inclusion of different models, which required change. It is important to note that the field of AI is progressing at a tremendous rate, which caused requirements for some changes in APIs, however they were infrequent. Before each implementation of piloting, the internal tests were made to make sure that the piloting (in the field) was not hampered by connectivity or instability issues. This testing and validation entailed use of testing images and videos, which were used to test the correctness of connections between modules, the correct operation of MQTT bus and correctness of operation of GUI.

4.4.4 Maintenance Plan

The components, when adopted, require maintenance solely from the software versioning point of view. This means that the software components (not the AI-models) must be patched and upgraded when needed.

However, if there is a need to update the AI-models, more rigorous testing is needed, compliant with the then-available test sets and the version of AI-Act that is enforced.

4.4.5 Lessons Learned

The positive result in production of the RII Tool is that adoption the API-First Development results in a stable system that is flexible and robust. Importantly, adoption of containerization through Docker allows for the system to be easily transferrable among different infrastructures (either distributed or centralized on the CIs infrastructure).

5 Conclusions

Deliverable D3.4 completes WP3's work, turning SUNRISE's specifications into actionable blueprints and operational playbooks, transforming the SUNRISE consortium's conceptual designs and specification documents into concrete, field-tested blueprints and operational playbooks. By distilling four fully realized reference architectures—Risk Based Access Control, Demand Prediction & Management, Cyber-Physical Resilience, and Remote Infrastructure Inspection—this document offers a unified approach to deploying modular, containerized toolchains that can plug seamlessly into diverse CI environments. Each architecture has been refined through iterative validation, ensuring the patterns for secure messaging, data ingestion, AI-driven forecasting, and automated recovery are robust under real-world stress conditions.

Alongside these structural designs, the body of best practices encapsulates the consortium's accumulated expertise in managing complexity, ensuring regulatory alignment, and establishing robust public–private partnerships. It spans the discipline of API-first development and the automation of testing pipelines, as well as the nuanced orchestration of change management within multi-stakeholder environments. These operational guidelines guarantee that implementers not only deploy effective systems but also maintain them amid shifting threat landscapes and evolving policy frameworks. Privacy-by-design principles are woven throughout to ensure GDPR compliance and readiness for forthcoming AI regulations, while the prescribed governance framework promotes coordinated incident response across both public agencies and private operators.

Looking beyond immediate deployment, the methodologies and templates in D3.4 lay a lasting foundation for the SUNRISE project's exploitation and for future research into resilient infrastructure. CI operators can apply these blueprints to accelerate roll-out timelines, reduce integration overhead, and maintain service continuity when faced with pandemics, cyber-attacks, or natural disasters. Technology providers and policymakers, in turn, can leverage the documented lessons learned to refine certification processes, evolve regulatory frameworks, and invest in next-generation resilience technologies. As the SUNRISE consortium advances into its exploitation and dissemination phases, the insights gathered here will continue to guide practical implementations and cross-sector collaborations, ensuring critical services remain responsive, adaptive, and secure in the face of future disruptions.

References

[1] **NVIDIA**, " Jetson Nano Developer Kit," [Online]. Available: https://developer.nvidia.com/embedded/jetson-nano-developer-kit.

[2] **Intel**, "Intel Neural Compute Stick 2 – Specifications," [Online]. Available: https://www.intel.com/content/www/us/en/products/sku/182253/intel-neural-compute-stick-2/specifications.html.

[3] **NXP**, "MIFARE High Frequency (HF) Products," [Online]. Available: https://www.nxp.com/products/rfid-nfc/mifare-hf:MC_53422.

[4] **N. Forum,** "NFC protocols and standards," [Online]. Available: https://nfc-forum.org/what-is-nfc/.

[5] N. P. J. A. R. &. A. M. N. Singh, "**Short-term forecasting in smart electric grid using N-BEATS**," 2022.

[6] Y. W. Z. N. Y. Z. T. Z. S. L. Y. .. &. W. Q. Kong, "**Unlocking the power of LSTM for long term time series forecasting,**" 2025.

[7] S. N. H. C. G. L. X. &. L. J. Shan, "**A machine learning framework for enhancing short-term water demand forecasting using attention-biLSTM networks integrated with XGBoost residual correction**," 2023.

[8] G. Labs, "**Grafana Dashboards Library**," Grafana Labs, [Online]. Available: https://grafana.com/grafana/dashboards/.

[9] "Keycloak: **Open Source Identity and Access Management**," [Online]. Available: https://www.keycloak.org/.

[10] I. MinIO, "**High Performance, Kubernetes-Native Object Store**," [Online]. Available: https://min.io/.

[11] S. Ramírez, "**FastAPI – FastAPI framework, high performance, easy to learn, fast to code**," [Online]. Available: https://fastapi.tiangolo.com/.

[12] "**Kubernetes: Production-Grade Container Orchestration**," [Online]. Available: https://kubernetes.io/.

[13] "**MLflow: Open Source Platform for the Machine Learning Lifecycle**," [Online]. Available: https://mlflow.org/.

[14] "**PostgreSQL: The World's Most Advanced Open Source Relational Database,**" [Online]. Available: https://www.postgresql.org/.

[15] "**Elasticsearch: Distributed, RESTful Search and Analytics Engine**," [Online]. Available: https://www.elastic.co/elasticsearch.

[16] I. **Wazuh**, "Wazuh – Open source security platform," [Online]. Available: https://wazuh.com/.

[17] I. Corporation, "**IBM Security QRadar SIEM**," [Online]. Available: https://www.ibm.com/products/qradar-siem.

[18] M. Project, "**MISP – Open Source Threat Intelligence and Sharing Platform**," [Online]. Available: https://www.misp-project.org/.

[19] S. G. Z. e. al, SUNRISE **"D6.5 Cyber-physical resilience tool and training guide V3**", 2025.

[20] "**TheHive: Scalable, Open Source and Free Security Incident Response Platform**," [Online]. Available: https://strangebee.com/.

[21] T. T. e. al, SUNRISE**, "D4.4 Access control pilot report V1**", 2024.

[22] I. Meta Platforms, "React – A JavaScript library for building user interfaces," Meta, [Online]. Available: https://react.dev/.

[23] W. E. Deming, **Out of the Crisis**, Cambridge, MA: Massachusetts Institute of Technology, Center for Advanced Educational Services, 1986.

[24] T. Pearson, "**The ultimate guide to the OODA loop**," 2025. [Online].

[25] E. R. E. A. (REA), **SUNRISE** Project, **Grant Agreement** 101073821", 2021.

[26] G. Z. S. e. al, SUNRISE, "**D6.4 Cyber-physical resilience pilot report V1**", 2024.

[27] V. D. e. al, SUNRISE "**D7.5 Infrastructure inspection tool and training guide V3**," 2025.

[28] T. M. e. al, SUNRISE, "**D7.2 Infrastructure inspection tool and training guide V1**", 2023.

[29] Š. D. e. al, SUNRISE, "**D7.1 Infrastructure inspection conceptualization**", 2023.

[30] G. T. e. al, SUNRISE, "**D7.3 Infrastructure inspection tool and training guide V2**", 2024.

[31] T. M. e. al, SUNRISE, "**D7.4 Infrastructure inspection pilot report V1**", 2024.

[32] X. e. a. Zhu, "**TinaFace: Strong but Simple Baseline for Face Detection**," arXiv preprint, p. arXiv:2011.13183v3, 2020.

[33] P. w. Code, "**Face Detection Task Overview**," 2024. [Online]. Available: https://paperswithcode.com/task/face-detection.

[34] J. B. Lindsay, "**WhiteboxTools: An advanced geospatial data analysis platform**," [Online]. Available: https://www.whiteboxgeo.com/.

[35] J. R. C. European Commission, "**The ODISE Model: Open Data Infrastructure for Security and Emergency Management**," Publications Office of the European Union, Luxembourg, 2020.

[36] P. L. PBC, "**SkySat: High-resolution satellite imagery for Earth observation**," Planet Labs, [Online]. Available: https://www.planet.com/products/skysat/.

[37] P. L. PBC, "**PlanetScope: Daily, global imagery at 3–5 meter resolution**," Planet Labs, 2024. [Online]. Available: https://www.planet.com/products/planetscope/.

[38] G. D. Marco, SUNRISE, "**D5.4 Demand prediction and management pilot report V1**", 2024.

[39] **GitHub**, "GitHub: Where the world builds software," [Online]. Available: https://github.com/.

[40] A. S. A. K. S. A. S. &. K. A. Nazir, "**Forecasting energy consumption demand of customers in smart rid using Temporal Fusion Transformer (TFT)**," Results in Engineering, vol. 17, 2023.

[41] J. M. del Alamo, SUNRISE, "**D1.1 Local meetings with CI stakeholders**", 2022.

[42] S. S., SUNRISE, "**D2.2 Strategy for awareness and resilience of Cis V1**", 2024.

[43] G. T., SUNRISE, "**D3.1 Requirements and designs V1**", 2023.

[44] I. S. e. al, SUNRISE, "**D3.2 Requirements and designs V2**", 2023.

[45] G. T., SUNRISE, "**D3.3 Final technical specifications and designs**", 2024.

[46] G. D. Marco, SUNRISE, "**D4.6 Access control pilot report V2** ", 2025.

[47] J. A., SUNRISE, "**D5.6 Demand prediction and management pilot report V2**", 2025.

[48] G. D. Marco, SUNRISE, "**D6.6 Cyber-physical resilience pilot report V2**", 2025.

[49] D. V., SUNRISE, "**D7.6 Infrastructure inspection pilot report V2**", 2025.

[50] T. S., SUNRISE, "**D10.2 AI – Requirement No.2**", 2022.