CyberSecDome

An innovative Virtual Reality based intrusion detection, incident investigation and response approach for enhancing the resilience, security, privacy and accountability of complex and heterogeneous digital systems and infrastructures.

Best Practice Guidelines for AI-powered Cybersecurity Solutions: Lessons Learned from the CASPER Pilot and Future Directions for CyberSecDome

CyberSecDome has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101120779.

Abstract

CASPER validated selected CyberSecDome tools within the INTEL ANN smart parking environment, a representative smart-city ecosystem consisting of interconnected services, communication infrastructures, operational assets, and data-driven processes. Through the deployment of CyberSecDome components, the execution of predefined attack scenarios, the collection of evidence, and the validation of cybersecurity capabilities, the pilot generated practical insights regarding the adoption of AI-powered cybersecurity solutions in operational environments. The lessons learned provide actionable recommendations for organisations seeking to deploy, validate, and operationalise AI-driven cybersecurity technologies within complex and interconnected environments.

This document presents the key best practices derived from the CASPER pilot expertise and proposes future directions for the evolution of the CyberSecDome ecosystem. The recommendations focus on deployment preparation, validation methodologies, evidence management, tool interoperability, and operational cybersecurity processes that can support future adopters of AI-assisted cybersecurity technologies.

1. Introduction

The increasing digitalisation of smart city services creates new cybersecurity challenges. Smart parking systems, intelligent transportation services, IoT infrastructures, and connected urban platforms rely on continuous data exchange and operational interoperability between multiple components. As a result, organisations require cybersecurity solutions capable of monitoring complex environments, identifying emerging threats, supporting risk assessment activities, and assisting security operators in decision-making processes.

Within this context, CASPER deployed and validated CyberSecDome tools in the INTEL ANN smart parking environment. The pilot followed an evidence-oriented methodology in which assets, attack scenarios, evidence artifacts, validation activities, and performance indicators were systematically linked throughout the project lifecycle.

The document is intended for cybersecurity practitioners, infrastructure operators, researchers, and organisations considering the adoption of AI-assisted security solutions in smart-city and critical-infrastructure settings. Its purpose is not to repeat the technical results of the pilot, but to capture the practical lessons that emerged from the deployment and validation process and translate them into recommendations for future deployments of AI-powered cybersecurity solutions.

2. Best Practice Guidelines for AI-Powered Cybersecurity Solutions

BP1. Establish Asset to Attack Scenario Traceability Before Technology Deployment

One of the most important observations from the CASPER pilot was that effective cybersecurity validation requires a clear understanding of the assets that participate in the operational environment and how those assets are affected by potential cyber threats.

Rather than treating cybersecurity monitoring as an isolated technical activity, the CASPER methodology established traceability between operational assets, attack scenarios, validation activities, collected evidence, and project KPIs. This approach allowed the project to understand not only whether a cybersecurity tool generated alerts, but also how those alerts related to specific operational assets and cybersecurity objectives.

Recommendation

Before deploying AI-powered cybersecurity solutions, organisations should establish and maintain a traceable mapping between operational assets, relevant threats, attack scenarios for evaluation, generated evidence, and validation objectives.

BP2. Use Evidence-Oriented Validation Instead of Activity-Oriented Validation

A recurring lesson from the CASPER pilot was that cybersecurity validation should not focus solely on performing activities such as penetration tests or attack simulations. Instead, validation should focus on the evidence generated by those activities.

Throughout the pilot, scenario execution was designed to produce measurable artifacts, including logs, alerts, PCAP files, vulnerability findings, incident records, and visualisation outputs. This allowed validation conclusions to be supported by traceable evidence rather than subjective observations.

Recommendation

Validation processes should be designed around evidence generation and evidence traceability rather than around the completion of predefined testing activities.

BP3. Link Cybersecurity Scenarios to Measurable Validation Objectives

The CASPER pilot demonstrated that cybersecurity scenarios become significantly more valuable when they are explicitly linked to measurable validation objectives. Rather than treating attack scenarios as standalone testing exercises, the project associated each scenario with expected artifacts, validation activities, and KPI assessment processes. Measurable outcomes may include detection rates, response times, false positive reduction, incident resolution metrics, or evidence completeness indicators. This approach enabled objective evaluation of cybersecurity capabilities and improved the reproducibility of validation results.

Recommendation

Organisations should ensure that cybersecurity scenarios are designed to support specific validation objectives and measurable outcomes. Scenarios should generate evidence that can be directly used to assess cybersecurity performance and validate the effectiveness of deployed solutions.

BP4. Integrate Risk Assessment, Threat Detection, and Security Testing

The pilot highlighted the importance of combining multiple cybersecurity perspectives instead of relying on individual tools operating independently.

The CyberSecDome ecosystem allowed the project to combine asset information, dynamic risk assessment outputs, penetration testing results, and threat detection evidence. This integrated approach provided a more complete view of the cybersecurity posture than any individual component could provide independently.

Recommendation

Risk assessment, threat detection, and security testing capabilities should be treated as interdependent capabilities within a unified cybersecurity workflow. Organisations should ensure that outputs generated by these functions can be correlated and jointly analysed in order to support more informed cybersecurity decisions and a comprehensive understanding of the security posture.

BP5. Treat Interoperability as a Cybersecurity Requirement

The CASPER experience showed that the value of cybersecurity platforms depends not only on individual tool performance but also on the ability of different components to exchange information and contribute to a common operational picture.

Throughout the pilot, the interaction between risk assessment, detection, incident management, investigation, visualisation, and collaborative intelligence functions played a significant role in supporting the overall validation process.

Recommendation

Interoperability should be considered a core cybersecurity requirement from the earliest stages of solution deployment and integration. Cybersecurity components should be designed to exchange information efficiently and contribute to a shared operational picture that supports coordinated security monitoring, investigation, and response activities. Open standards, common data models, and interoperable interfaces should be prioritised whenever possible.

BP6. Establish Traceable KPI-to-Evidence Validation Chains

The CASPER pilot demonstrated that cybersecurity validation becomes significantly more reliable when quantifiable success metrics are defined from the beginning of the validation process. The project methodology linked validation objectives, cybersecurity scenarios, generated artefacts, tool outputs, and KPIs through a traceable validation chain. In this context, KPIs were not treated as generic indicators, but as measurable criteria connected to expected evidence, acceptance criteria, and performance thresholds. These thresholds may include detection accuracy, response time, false

positive reduction, incident resolution performance, or evidence completeness. By defining these metrics early, organisations can assess whether AI-powered cybersecurity capabilities meet their intended objectives in a transparent, comparable, and reproducible way.

The project methodology linked assets, scenarios, evidence packages, validation activities, and KPIs throughout the validation lifecycle. This approach enabled structured assessment of CyberSecDome capabilities using measurable and reproducible evidence rather than subjective observations.

Recommendation

Organisations deploying AI-powered cybersecurity solutions should establish traceable links between:

- validation objectives,
- Cybersecurity attack scenarios,
- generated evidence,
- tool outputs and
- performance indicators.

A structured KPI-to-evidence validation chain improves transparency, supports reproducibility and strengthens confidence in cybersecurity evaluation results.

BP7. Design Validation Activities with Future Evaluation and Exploitation in Mind

A key strength of the CASPER methodology was the early definition of KPIs, evidence requirements, validation criteria, and reporting structures.

This ensured that the outputs generated during deployment and validation could later support evaluation activities, dissemination actions, exploitation planning, and future adoption recommendations.

Recommendation

Validation activities should be planned from the beginning to support not only technical assessment, but also future evaluation, reporting, dissemination, exploitation, and knowledge-transfer objectives. Knowledge-transfer activities, training materials, and stakeholder engagement plans should also be considered during validation planning.

3. Future Directions for CyberSecDome

The CASPER pilot identified several opportunities for the future evolution of the CyberSecDome ecosystem. These directions build upon the deployment, validation, and operational experience gained during the pilot and aim to further enhance the effectiveness, adaptability, and applicability of AI-powered cybersecurity solutions.

FD1. Towards Autonomous Evidence Correlation and Security Intelligence

Future versions of CyberSecDome could further strengthen the correlation and consolidation of cybersecurity evidence generated across different components. Improved evidence management and correlation capabilities would enable security operators to establish stronger links between risks, vulnerabilities, alerts, incidents, and validation outcomes, supporting more informed and evidence-based cybersecurity decision-making.

FD2. Towards Adaptive Risk Assessment

Future risk assessment capabilities could become more dynamic and adaptive by continuously incorporating operational changes, newly identified vulnerabilities, emerging threats, and contextual information in an automatic manner. AI-assisted risk assessment mechanisms could support a more accurate representation of the evolving cybersecurity posture of complex smart city environments. Future implementations could incorporate automatic asset identification, predictive analytics and threat intelligence feeds to proactively adapt risk profiles.

FD3. Towards Integrated Detection and Response Workflows

The integration between threat detection, incident investigation, and response activities could be further strengthened in future CyberSecDome deployments. Closer alignment between these functions would facilitate faster incident handling, improved situational awareness, and more efficient cybersecurity operations.

FD4. Towards Broader Smart City Adoption

The concepts and capabilities validated within the INTEL ANN smart parking environment could be extended to additional smart city domains, including intelligent transportation systems, municipal digital services, IoT infrastructures, and other critical urban environments. This expansion would increase the applicability and impact of CyberSecDome across diverse operational contexts.

FD5. Towards Federated and Collaborative Cybersecurity Intelligence

Future CyberSecDome deployments could further enhance collaborative cybersecurity operations through improved cross domain knowledge sharing, collaborative intelligence mechanisms, and coordinated decision-making processes. Strengthening collaboration among stakeholders would contribute to a more resilient and proactive cybersecurity ecosystem.

4. Conclusions

The CASPER pilot demonstrated that the successful adoption of AI-powered cybersecurity solutions depends not only on the technical capabilities of individual technologies, but also on the strength of the methodologies used to deploy, validate, evaluate, and operationalise them.

The experience gained during the pilot highlighted the importance of establishing clear traceability between assets, threats, scenarios, evidence, and performance indicators. It also demonstrated the strategic value of evidence-oriented validation, measurable validation objectives, integrated cybersecurity workflows, interoperability between cybersecurity components, and structured KPI-to-evidence validation processes. Together, these elements strengthen the credibility, transparency, and reproducibility of cybersecurity validation activities.

The best practices presented in this guideline provide practical and actionable recommendations for organisations seeking to deploy and evaluate AI-powered cybersecurity solutions in smart-city and critical-infrastructure environments. At the same time, the proposed future directions provide a strong foundation for the continued evolution of CyberSecDome towards more adaptive, evidence-driven, integrated, and collaborative cybersecurity capabilities.

By building upon these principles, future deployments of CyberSecDome can further strengthen cybersecurity resilience, support faster and more informed decision-making, and contribute to the secure operation of increasingly complex digital ecosystems.