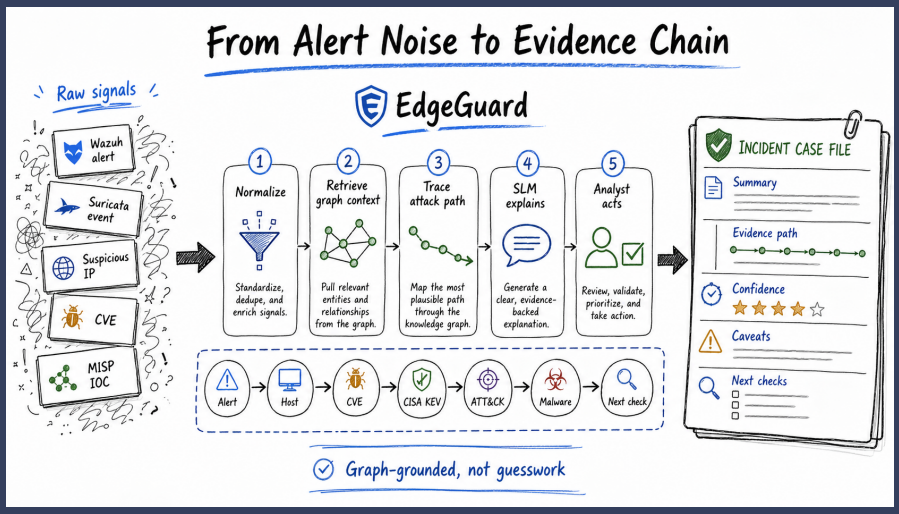
Resilmesh
securing cyber infrastructures

Project Spotlight: EdgeGuard Finalizes Sprint 1 for Explainable Edge Cybersecurity

In modern Security Operations Centers (SOCs), analysts rarely get a complete story during an incident - they get fragments. A lateral-movement alert lands first, followed by a vulnerability record, a MITRE ATT&CK technique, and an isolated threat intelligence feed. Piecing these messy fragments together under intense time pressure is one of the industry's steepest challenges.



EdgeGuard is stepping in to solve this exact problem. Accepted under the EU-funded **Resilmesh project** (Open Call 2, Challenge C2), EdgeGuard combines Graph-Augmented Retrieval (GraphRAG) and specialized Small Language Models (SLMs) to automatically turn fragmented alerts into traceable, evidence-backed case files right at the edge.

The project recently hit its first major milestone with the official finalization of Sprint 1, establishing a powerful graph substrate ready for operational testing. A complete, in-depth look at the architecture, plumbing, and timeline has been detailed by the team in Ratio1's technical blog post, EdgeGuard: SLM/GraphRAG Edge Intelligence.

A Powerful Research & Industry Collaboration

As outlined in the project's foundational framework, EdgeGuard's spine is built on a balanced collaboration between academic research and specialized edge infrastructure:

- IICT-BAS leads the research-heavy foundation, owning the graph schema design, Cyber Threat Intelligence (CTI) normalization, reasoning-path extraction, and the strict evaluation of SLM explanation faithfulness.

- Ratio1 delivers the edge reality, managing containerized deployment on edge nodes, node cryptographic identity, NATS-style messaging routing, and the decentralized architecture required to move the project from the lab to a viable commercial service.

Inside the Sprint 1 Milestone

While the machine-learning layer is slated for Sprint 2, the team spent Sprint 1 building the unglamorous but vital data plumbing required to power the AI. Key achievements from the deliverable include:

- Data Ingestion at Scale: The architecture successfully integrated 11 threat collector modules. An internal data run across 7 core sources (including NVD, CISA KEV, MITRE ATT&CK, and AlienVault OTX) populated the graph with over 251,000 zone-tagged nodes and 252,000 verified relationships.

- Rapid Graph Traversals: Early smoke tests demonstrated that complex, four-hop investigative paths (tracing an Indicator $\rightarrow$ Malware $\rightarrow$ Threat Actor $\rightarrow$ Technique) execute in under 40 milliseconds on internal test instances.

- Resilmesh Integration: EdgeGuard implemented a shared-database integration pattern within Neo4j and established real-time NATS messaging lines for seamless alert enrichment.

Local, Traceable, and Bounded AI

Rather than relying on massive, cloud-dependent LLMs that are prone to data leaks and hallucinations, EdgeGuard targets edge-optimized SLMs (1B to 8B parameters). In this design, the graph carries the strict truth, and the model's narrow job is to translate analyst queries and summarize the evidence path.

By keeping sensitive inference local to the edge, the system preserves strict data sovereignty. Furthermore, the architecture leverages blockchain principles to maintain a tamper-evident record of the exact model version and graph snapshot used during an investigation, ensuring every AI-generated case note is legally and operationally defensible.

The Road to Sprint 2

With the foundational graph architecture securely in place, the EdgeGuard team is moving directly into its next phase. Immediate milestones include launching a 730-day historical production-test baseline on Ratio1 infrastructure, curating the specialized Q&A training dataset, fine-tuning the SLM, and kicking off the first round of formal latency, quality, and explainability benchmarks.

For a comprehensive technical analysis of these milestones and the underlying data pipelines, read the full article on the Ratio1 Blog.

Funding Acknowledgement: EdgeGuard is developed in the context of the Resilmesh project, funded by the European Union. This post reflects the authors' interpretation of the project work and does not represent the official position of the European Union or the Resilmesh consortium.

Follow Resilmesh on Social Media:

Funded by the European Union
The content of this website represents the views of the authors only and is their sole responsibility. The European Commission does not accept any responsibility for use that may be made of the information it contains.