

# Project GAINS

(Generative AI for Network Sustainability)

Deliverable D3.1

Technical Report on Network Monitor Design and Implementation

Giordano, A. – Basile, L. – De Giacomo, A. – Haider, B – Paesani, M.  
2026-01

Work Package: WP3 Design, Prototype Implementation and Evaluation of Agents

Status: Final | Version: 1.0

Classification: internal project use; controlled sharing with partners as needed.

## Abstract

This deliverable presents the design and implementation blueprint of the Network Monitor of the GAINS project, a software agent aimed at transforming network telemetry into verifiable, replicable, and operational knowledge ready for industrial adoption. In telco and datacenter infrastructures, the availability of data does not coincide with the ability to make reliable decisions: the sources are heterogeneous, the signals are often not comparable over time, and the summaries produced in emergencies lack traceability. The Network Monitor was created to bridge this gap, imposing a canonical data model, an explicit transformation chain, and an evidence governance that allows audit, reproducibility, and continuous improvement. Overall, the Network Monitor is proposed as an engineering foundation for advanced and sustainable observability: an agent capable of making the network not only measured, but also interpretable, verifiable, and governable.

The architecture is "read-only by design": the agent does not change configurations or network state and operates exclusively through connectors in read-only mode to monitoring systems, databases, and historical time-series. Each analytical result includes a `trace_id` and evidence referable to the sources and time windows used, enabling transparent diagnostics and measurable technical accountability. The analytical pipeline is intentionally deterministic: normalization of metrics and units, quality gates, calculation of robust baselines (p50/p95 quantiles), health scoring, and detection of anomalies with versioned detectors. This choice prioritizes reliability and controllability, reducing false positives and making results repeatable with the same input and configuration.

The system strictly separates high-cardinality raw telemetry from long-term memory: the history remains in the native systems, while the knowledge store retains only low-cardinality derivatives (cardinality, summary, anomaly events, operator notes, registry of evidence, and inputs for sustainability analysis).

In particular, the Network Monitor introduces the canonical object Footprint Inputs, which produces load proxies and efficiency signals that can be attributed by device and aggregated by PoP, accompanied by method, confidence, declared assumptions, and references to evidence.

This enables a progressive enhancement approach: even in the absence of physical device-level power measurements, the agent provides comparable and auditable outputs, which can be improved when direct signals become available.

Integration with an external OpenAI-compatible LLM endpoint is expected but not authoritative: it is used in an optional way for embedding and grounded explanations, keeping deterministic mechanisms separate that guarantee health, severity, and detection.

The deliverable also includes stable API contracts with standard envelopes and consistent error models, as well as a reproducible deployment blueprint in Docker, with segregation of secrets, mount read-only for time-series sources, conservative caching and retention policies, and a testing strategy based on contract testing and golden tests to ensure determinism and controlled regression.

For the next PoC phase, LibreNMS will be adopted as the reference NMS, already operational in the target network with real measurements. This choice accelerates integration and validation on authentic data, mainly impacting connectors and mapping to the canonical model, without altering architectural invariants and contracts. [1][2][3][4][5][19]

## Table of Contents

1. Introduction
  - 1.1 Purpose of the deliverable and scope
  - 1.2 Operating context and safety principles
  - 1.3 Recruitment and external dependencies
2. Network Monitor Requirements
  - 2.1 Traceability requirements and criteria
  - 2.2 Functional requirements
  - 2.3 Non-functional requirements
  - 2.4 Climate requirements and measurability constraints
3. Network Monitor Agent Architecture
  - 3.1 Architectural principles
  - 3.2 Logical view of components
  - 3.3 Component model and responsibilities
  - 3.4 Boundaries of responsibilities and interfaces
  - 3.5 Main Workflows
4. Data model and agent memory
  - 4.1 Canonical data model adopted by the Monitor
  - 4.2 Data domains and memory levels
  - 4.3 Persisted objects, indexing and retention policies

- 4.4 Evidence Reference, trace\_id and audit trail
- 5. Acquisition, normalization, and ETL pipelines
  - 5.1 Read-only sources and methods
  - 5.2 Internal envelope and unit normalization/naming
  - 5.3 Mapping device identity and enrichment metadata
  - 5.4 Version able pipelines and data quality criteria
- 6. Deterministic analysis: baseline, scoring and anomaly detection
  - 6.1 Time windows and steps
  - 6.2 Baseline (p50/p95) and drift
  - 6.3 Quality gates and missing data management
  - 6.4 Detectors: signals, rules and severity
  - 6.5 Anomaly Event Output and Correlation
- 7. Climate support: production of Footprint Inputs
  - 7.1 Physical signal chain and logical proxies
  - 7.2 Confidence, assumptions and comparability
  - 7.3 Evidence and limitations of estimation
- 8. API contracts and response models
  - 8.1 Endpoint families
  - 8.2 Response envelopes and invariants
  - 8.3 Standardized error model
  - 8.4 Open API and contract testing
- 9. Docker deployment blueprints
  - 9.1 Objective and scope of the blueprint
  - 9.2 Containerized components and external dependencies
  - 9.3 Repository structure and artifact management
  - 9.4 Docker compose reference
  - 9.5 Bootstrapping and smoke testing
- 10. Implementation reference: construction details

10.1 Implementation choices and internal modules

10.2 Connectors (API/DB/RRD) and read-only policy

10.3 Versioned configurations (detectors.yml, mappings.yml)

10.4 Init schema for knowledge store and "bring your own vectors"

10.5 Structured logging, metrics and tracing

10.6 Deterministic testing and golden tests

10.7 Security: secrets, redaction, minimization

11. Conclusions and lessons learned

Bibliography

Appendices