

Project GAINS

(Generative AI for Network Sustainability)

Deliverable D3.2

Technical Report on Network Orchestrator 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

The operational management of networks and digital infrastructures of the new generation requires increasingly fast, controllable, and auditable change processes, in the face of growing complexity due to technological heterogeneity, variable traffic dynamics, security constraints, and sustainability requirements. In this context, traditional automation based on playbooks and static rules shows structural limitations: it struggles to generalize across domains and vendors, does not robustly make decision assumptions explicit, and rarely natively integrates an accountability model that links decisions, evidence, configurations, and impact. This deliverable presents the design and reference implementation of the Network Orchestrator, a software agent oriented toward intent-based networking and change governance, designed to transform operational intents into candidate configurations and execution plans, integrating knowledge grounding and estimation of energy impact in a traceable and reproducible manner. The proposed approach combines deterministic control components with assistive components based on language models, constrained by schemas and validation rules. The Network Orchestrator operates as a stateless API service with external persistence of state and artifacts, introducing an explicit and idempotent job lifecycle, capable of managing end-to-end the submit–context–generation–validation–impact–documentation–approval chain. The architectural core is an agentic pipeline that: normalizes the intent into a canonical representation; performs applicable retrieval from a knowledge store using hard filters on vendor, role, and version; produces structured artifacts that include explicit references to the sources used; enforces deterministic security and compliance policies; performs syntactic and semantic validations on configuration and change plan; estimates energy impact as-is/to-be/delta via telemetry and models, making explicit confidence, method, and assumptions; and finally generates an operational explain pack oriented toward review, audit, and human decision-making. The use of an external OpenAI-compatible LLM endpoint is confined to non-authoritative tasks such as assisted parsing, variable binding, and controlled enrichment of the plan, with minimization of context, preventive drafting, and schema-based post-validation, mitigating risks of prompt injection and data leakage. The main contribution is a replicable implementation model that makes change generation “decidable”: each output is packaged into versioned and hashed artifacts, linked to knowledge references and evidence references, and accompanied by validation reports and verifiable rollback conditions. This approach enables change

governance consistent with industrial requirements of traceability, operational risk reduction, and compliance control, and introduces a first level of sustainability by design thanks to the systematic production of impact estimates and the explicit management of uncertainty. Evaluation is defined through a multi-level suite of unit tests, schema contract tests, integration and regression tests, including quantitative metrics on correctness, security, auditability, and performance, as well as adversarial tests for robustness. The Network Orchestrator constitutes a key building block for the industrialization of agent-based architectures applied to networks and infrastructures, in line with digital and sustainable transition goals: it reduces time-to-change, increases the quality and verifiability of decisions, and prepares controlled integration toward execution systems and closed-loop monitoring, while keeping reproducibility, accountability, and data protection at the center. [15][22][23][27]

Table of Contents

1. Purpose and Scope of the Network Orchestrator

1.1 Functional objectives

1.2 Recruitment and operational constraints

1.3 Interfaces with other platform components

2. Project requirements and criteria

2.1 Functional requirements

2.2 Non-functional requirements: safety, audit, resilience, quality

2.3 Criteria of industrial replicability and portability

2.4 Sustainability principles and attributable footprint

3. Network Orchestrator Logical Architecture

3.1 Internal Components: Intent Normalizer, Planner, Validator, Artifact Manager

3.2 Main flows: intent, plan, approval, artifacting, degrade

3.3 Knowledge and telemetry integration

3.4 Boundaries of responsibility and threat model

4. Data model and canonical schemas

4.1 Core Objects: Orchestration Job, Candidate Configuration Artifact, Change Plan Artifact, Footprint Estimate Artifact

4.2 Versioning, hashing, immutability, and lineage

4.3 References: knowledge refs and evidence refs

- 4.4 Full JSON payload examples
- 5. Network Orchestrator API Contracts
 - 5.1 Public endpoints and authentication
 - 5.2 Error model, idempotency, and correlation required
 - 5.3 Job lifecycle states, transitions, degradation policies
 - 5.4 Call and Answer Examples
- 6. Agent pipeline and build logics
 - 6.1 Intent parsing and normalization in the canonical model
 - 6.2 Applicable retrieval and grounding on blueprints, snippets, and policies
 - 6.3 Candidate generation: configuration and change plan
 - 6.4 Static validations and consistency checks
 - 6.5 Estimation of impact and sustainability: as is, to be, delta, confidence
 - 6.6 Human in the loop: approve, reject, and audit trail
 - 6.7 Final outputs, explain packs, and rollback conditions
- 7. Software implementation
 - 7.1 Repository and module structure
 - 7.2 Runtime configuration: env, secrets, profiles
 - 7.3 Controlled prompting strategy and templates
 - 7.4 Validation and render drivers for vendor, device_role, OS
 - 7.5 Logging, tracing, and audit trail
- 8. Deployment in Docker
 - 8.1 Prerequisites and execution profile
 - 8.2 Docker compose: complete services, volumes, networks
 - 8.3 Bootstrapping, initialization, schemas, and smoke tests
 - 8.4 Minimal hardening for enterprise environments
- 9. Testing and evaluation

9.1 Unit test and contract test

9.2 Integration testing with knowledge and telemetry

9.3 Robustness testing, regression, safety

9.4 Metrics of quality: correctness, applicability, reproducibility, timing

10. Limitations, risks, and evolutionary roadmap

10.1 Failure modes and mitigations

10.2 Expected developments and industrial generalization

Bibliography

Appendices

A. Examples: end-to-end intent, plan, approval output

B. Dictionary of fields and conventions

C. Deployment and verification worklist

D. Essential acronyms