

# Project GAINS

(Generative AI for Network Sustainability)

Deliverable D3.3

Technical Report on Design and Implementation of Knowledge Manager

Giordano, A. – Basile, L. – De Giacomo A. – Haider, B. – Fruncillo, D.  
2026-01

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

Status: Final | Version: 1.0

Classification

## Abstract

This work describes the design and prototypical implementation of the Knowledge Manager, a knowledge-plane component oriented toward operations for telco and cloud environments, with the objective of transforming documentation, policies, and technical artifacts into a system that is queryable, applicable, and verifiable. The Knowledge Manager is conceived as a containerized service, reproducible in Docker, capable of operating with an external OpenAI-compatible endpoint for embeddings and generation, maintaining decoupling from the provider and control over the information perimeter effectively transmitted externally. The design is centered on three non-negotiable requirements: contextual applicability, grounding with verifiable citations, and security against contamination and exfiltration. The solution adopts a layered architecture: deterministic ingest pipeline, knowledge store with hybrid search and structured filters, and agentic runtime for retrieval and synthesis. The ingest is implemented in Langflow through versioned flows, exportable and importable, which perform parsing, normalization, deterministic chunking, canonicalized hashing, mandatory redaction for snapshots, and anti-leakage controls with fail-closed rejection. Indexing uses Weaviate with BYOV configuration, in which vectors are computed by the application and associated with persisted objects, in order to preserve portability and governance of the embeddings lifecycle. Retrieval combines keyword search and vector search through hybrid search, with selectable fusion algorithms and inspectable scoring metadata, and applies binding filters based on flattened applicability properties such as vendor, device role, service, and version constraints, reducing the probability of semantic drift in operational contexts. Agentic behavior is orchestrated through Elysia, adopting a decision tree that makes explicit the transitions between request classification, gating, multi-source retrieval, deterministic re-ranking, construction of the evidence pack, grounded synthesis, and final validation. The operational mode imposes strict constraints: no prescriptive claim is emitted without at least one valid reference, and conflicts between sources are resolved with a policy-first approach, degrading the response to non-operational when evidence is insufficient. To support auditability, the system produces an Explainability Pack that can be persisted and exported, which includes claim-refs mapping, retrieval parameters, artifact hashes, and quality indicators such as citation coverage and unsupported claim rate, enabling repeatable verification and diagnosis of regressions.

Consistent API contracts are defined for ingest, cited or structured query, service blueprint, snippet discovery, and explain pack, with a common envelope, error model, RBAC, and export controls. The prototypical implementation includes a reference repository, full docker-compose, automatic bootstrap of schema and flows, secret management via Docker secrets, and reproducible smoke tests. Evaluation is structured on retrieval and grounding metrics, including precision and stability of top-k, applicability violation rate, citation coverage, and absence of unsupported claims in operational mode, with a regression pipeline on a locked dataset. Security is treated as an end-to-end requirement: minimization of payloads toward the LLM, defenses against direct and indirect prompt injection through evidence-pack isolation, export control policies, rate limiting, and fail-closed behavior, in alignment with the main risk taxonomies for LLM-integrated applications. Overall, the Knowledge Manager provides a solid implementation foundation for integrating technical knowledge and operational governance in an agentic context, with industrial reproducibility, traceability, and verifiable quality controls.

## **Table of Contents**

### 1. Purpose, Perimeter, and Operating Context

#### 1.1 Objectives of the Knowledge Manager

#### 1.2 Boundaries of responsibility and interactions with other agents

#### 1.3 Assumptions, constraints, and operational definitions

#### 1.4 Modes of Operation: PoC & Extended Profile

#### 1.5 Industrial replicability requirements and NRRP criteria

### 2. Project requirements and criteria

#### 2.1 Functional requirements (ingest, retrieval applicable, blueprint, snippet, citations)

#### 2.2 Non-functional requirements (security, auditability, portability, observability)

#### 2.3 Acceptance criteria and quality gates

#### 2.4 Architectural choices and trade-offs

### 3. Knowledge Manager Logical Architecture

#### 3.1 Overview and components (Elysia, Weaviate, Langflow, OpenAI-compatible LLM)

#### 3.2 Main flows: ingest, cited query, blueprint, explain pack

#### 3.3 Security and governance boundaries (RBAC, audit, segregation of roles)

#### 3.4 Error model and controlled degradation

3.5 Observability: trace\_ids, metrics, structured logs

#### 4. Canonical Data Model and Schemas

4.1 Principles of the data model (applicability, traceability, deduplication)

4.2 Cross-domain entities (Knowledge Ref, Taxonomy Applicability, Response Meta)

4.3 Knowledge Objects (Doc Source, Doc Chunk, Policy, Service Definition, Config Snippet, Blueprint, Config Snapshot Sanitized)

4.4 Hashing rules, idempotency, and versioning

4.5 Full JSON examples and required fields

#### 5. Knowledge Store with Weaviate

5.1 Persistence and indexing strategy

5.2 Weaviate Class Schema and Indexed Properties

5.3 Filters structured by applicability (vendor, device\_role, os\_version, service\_id)

5.4 Hybrid search and ranking parameters (keyword + semantics)

5.5 Index deduplication, updates, and rebuilds

5.6 Conservative indexing policies (sanitized snapshots)

#### 6. Ingest pipeline with Langflow

6.1 Ingest profiles and treatment policies

6.2 Parsing, normalization, and deterministic chunking

6.3 Redaction and anti-leakage controls

6.4 Metadata enrichment and validation: applicability

6.5 Calculation of embeddings via OpenAI-compatible endpoint

6.6 Upsert in Weaviate and idempotency management

6.7 Receipt, reason codes, and audit events

6.8 Examples of Langflow Flow (Export) and runtime configuration

#### 7. Applicable retrieval and agentic RAG with Elysia

7.1 Gating policy and query classification (operational vs informational)

- 7.2 Candidate set construction and binding filters
- 7.3 Deterministic ranking and boosting (policy-first)
- 7.4 Grounded synthesis with citations and claim-refs validation
- 7.5 Calculation of confidence, warnings, and management of conflicts between sources
- 7.6 Structured output mode for integration with other agents
- 8. Knowledge Manager API Contracts
  - 8.1 Conventions (versioning, envelope, trace\_id, idempotency)
  - 8.2 Endpoints: /query, /blueprint, /snippets, /ingest, /explain, /health
  - 8.3 Request schemas and canonical error codes
  - 8.4 RBAC for endpoints and export policies
  - 8.5 OpenAPI (excerpt) and compatibility rules
- 9. Explainability pack and document assembly
  - 9.1 ExplainPack canonical structure
  - 9.2 Claim-refs mapping and granularity criteria
  - 9.3 Export to JSON and markdown\_bundle
  - 9.4 Persistence, bundle\_uri, and correlation between
  - 9.5 Assembly security checks
- 10. Prototype implementation (code and components)
  - 10.1 Repository and module structure
  - 10.2 OpenAI-compatible clients (embeddings and chat)
  - 10.3 Weaviate integration (init schema, queries, filters)
  - 10.4 Langflow (flows, secrets, variables, environments)
  - 10.5 Elysia (decision tree, tool bindings, output contracts)
  - 10.6 Logging, tracing, and metrics
  - 10.7 Testing strategies (unit, contract, e2e)
- 11. Deployment in Docker

- 11.1 Docker-compose: Weaviate, Langflow, Elysia, Knowledge API, optional dependencies
- 11.2 Dockerfile and build strategy
- 11.3 Secrets management and configurations (env + docker secrets)
- 11.4 Bootstrap schema Weaviate and Langflow import flow
- 11.5 Repeatable end-to-end smoke testing
- 11.6 Minimum hardening and production profile
- 12. Evaluation, Quality, and Regression
  - 12.1 Test and coverage datasets (documents, policies, snippets)
  - 12.2 Retrieval metrics (precision@k, nDCG, hit-rate for applicability)
  - 12.3 Grounding metrics (citation coverage, unsupported claim rate)
  - 12.4 Regression and knowledge drift tests
  - 12.5 PoC acceptance criteria and thresholds
- 13. Security, compliance, and privacy
  - 13.1 Sanitized data minimization and snapshot processing
  - 13.2 Threat model (retrieval abuse, leakage, prompt injection)
  - 13.3 Countermeasures (redaction, allowlist, policy-first, export control)
  - 13.4 Auditability and traceability for review
- 14. Limitations and evolutionary roadmap
  - 14.1 Limitations of the PoC and residual risks
  - 14.2 Evolutions: approval workflow, "approved" status for snippet/policy, advanced governance
  - 14.3 Multi-vendor extension and advanced OS version management
  - 14.4 Performance and cost optimizations
- 15. Conclusion

Bibliography

Appendices