From Data Silos to Intelligent Operations

An AI-Based Approach to Ground Station Incident Investigation

Dimitri Accad
Nieves Salor Moral
Starion Deutschland GmbH, Darmstadt, Germany

Andrea Di Luca
Fondazione Bruno Kessler, Trento, Italy

Presenter: Dimitri Accad
d.accad@stariongroup.eu

STARION

About the Presenter

Dimitri Accad

Aeronautic and Aerospace engineer and AI Engineer | Starion Deutschland GmbH

I am an Aeronautic and Aerospace Engineer and an Artificial Intelligence Engineer specialising in Generative AI. After studying across France, Scotland, and the United States, I built my career across the space industry, from large organisations like ArianeGroup and Arianespace to innovative start-ups such as ThrustMe. I now work at Starion as the principal Generative AI specialist, supporting AI projects for the European Space Agency.

d.accad@stariongroup.eu | Starion Deutschland GmbH, Darmstadt, Germany

The AISHGO Project

AI for Automation of Satellite Health Monitoring and Ground Operations

AISHGO combines two work packages: **SMAI** and **AIIA**, the latter led by Starion.

Objective

Develop AI tools to support daily satellite operations at **ESOC** and beyond. User-friendly, versatile solutions for anomaly detection, long-term analysis, and intelligent investigation.

Approach

Open-source, on-premise deployment with small LLMs (8B parameters) using agentic architectures, RAG, and robust quality controls to meet accuracy requirements.

Scope

Ground station incident investigation, anomaly detection, predictive maintenance, and intelligent assistants. Paving the way for autonomous space operations.

The Challenge

Ground station operators face overwhelming daily data from isolated systems, making incident investigation slow, error-prone, and dependent on individual expertise.

Data Silos

Information scattered across isolated environments (STCs, ARTs, OSPMS, emails, chat logs, and technical documents)

Manual Investigation

Operators rely on subjective judgment and manual log searches across fragmented sources

Knowledge Loss

Staff turnover causes institutional memory loss; inconsistent formats lead to data duplication

Security Constraints

On-premise deployment required; external proprietary models precluded by confidentiality

System Overview

Our solution provides a suite of AI tools empowering operators throughout the incident investigation lifecycle.

Data Pipelines

Continuously extract, harmonise, and ingest data from heterogeneous sources

Knowledge Visualization

Vector + relational databases with Similarity & knowledge graph navigation

Multi-Expert AI Agents

Specialised AI agents for incident, facility, visual, and general queries with quality controls

Conclusion Generation

Autonomous agent for root cause analysis and mitigation actions

Tech Stack: Qwen3:8B | LangGraph | Weaviate | PostgreSQL | ColPali | Gemma3 | dspy | React

Data Pipeline Architecture

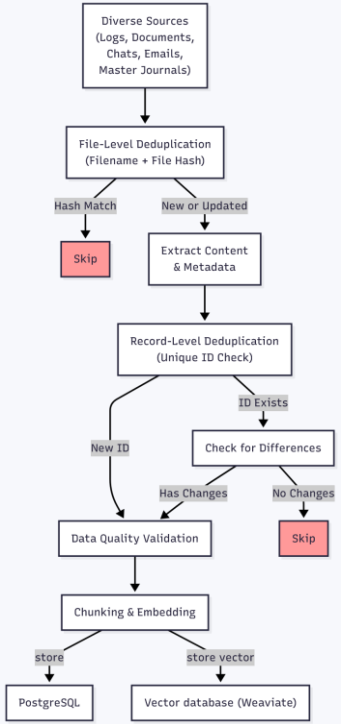

Data Pipeline Architecture

1. **Source Extraction**

   Scheduled cron jobs retrieve data daily from STCs, ARTs, documents, and communications

2. **File-Level Deduplication**

   Hash-based algorithms detect duplicates; only new or updated files proceed

3. **Record-Level Deduplication**

   CSV rows checked individually; changes trigger updates, unchanged records skipped

4. **Data Quality Validation**

   Rule-based approach: discard entries missing critical fields, retain those with non-critical gaps

5. **Chunking & Embedding**

   Overlapping chunks with metadata (timestamps, source IDs, position) for temporal coherence

6. **Dual Storage**

   Vectors in Weaviate, structured data in PostgreSQL, linked by unique identifiers

Advanced Data Processing

Embedding Strategy

- Overlapping chunks preserve temporal & context coherence
- Metadata enrichment: timestamps, source IDs, position in text
- Separate vector store collections per document type
- Enables transparent sourcing for each query
- Handles inconsistencies in notation and schema changes across years of human-created data

Multimodal Processing

- User-uploaded PDFs, images, and illustrations require special handling
- ColPali framework avoids lossy text-only segmentation
- Images, diagrams, and tables retrieved alongside text passages
- Gemma3 vision LLM creates embeddings from visual content
- Enables retrieval of schematics and charts during investigation

ColPali and dspy

ColPali

Efficient Visual Document Retrieval with Vision Language Models

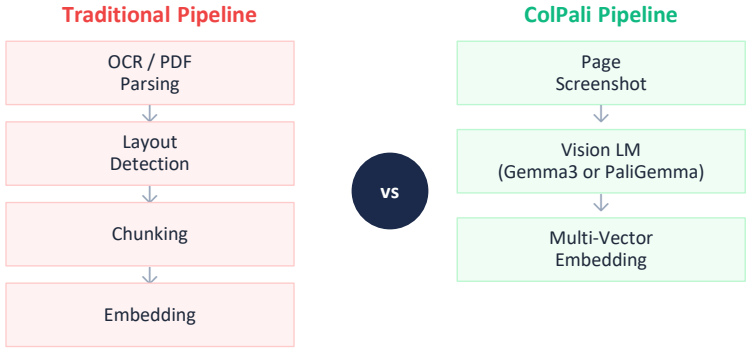

Why ColPali?

- Captures text, figures, tables, and layout in one pass
- No OCR errors, no lost diagrams, no brittle extraction
- We use it with Gemma3 vLLM for user-uploaded docs

dspy

Declarative Self-improving Python (Stanford NLP)

Programming LMs instead of prompting them

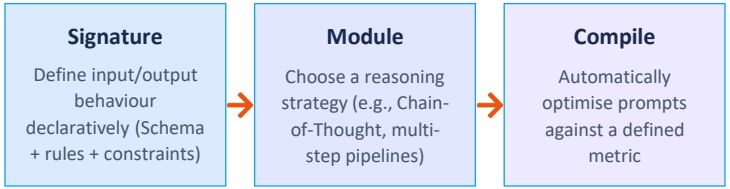

Our Usage in this System

Automatic Filter Generation:
Agents use dspy signatures to produce consistently formatted metadata filters (dates, stations, missions) for vector search

Structured JSON Output:
Ensures reliable schema compliance for search parameters, eliminating format errors in the agentic workflow

Data Lifecycle

From diverse sources to AI-powered insights

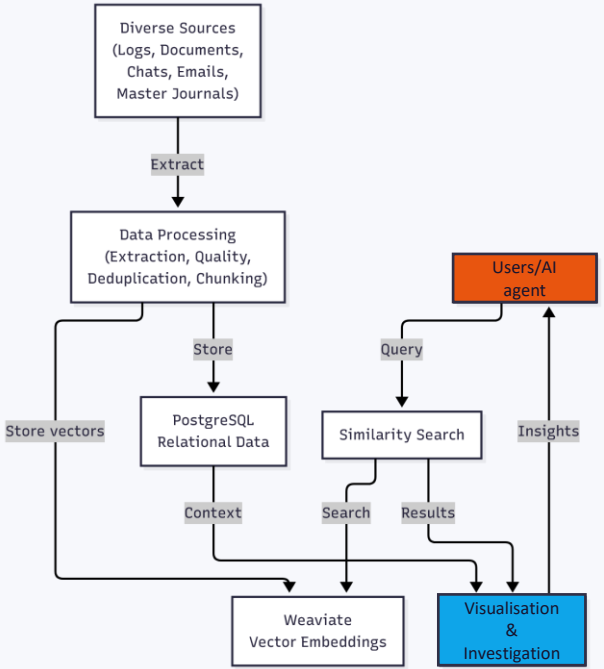

End-to-end data lifecycle

Common Process

- Search for Similarities in the vector Database (may use filters)
- Enrich them with the relational database data

Agent

Return the aggregated results to the agent which will further process the results (Re-Ranking)

Visualisation

Process the result for visualisation (t-SNE or graph construction)

Similarity Chart Visualization

Similarity Chart

Similarity Graph (t-SNE)

Dimensionality reduction projects high-dimensional vectors onto 2D maps. Clusters related ideas together, boosting discovery of similar incidents.

Knowledge Graph

Built using metadata on related incidents as links. Fully browsable and interactive: clicking on nodes reveals detailed incident information. Supports rapid, evidence-driven investigations across missions.

Key Capabilities

- Navigate relationships between incidents and anomaly reports
- Filter by dates, ground station origin, mission
- Combines AI-generated similarity with manual operator correlations

Similarity Chart Visualization

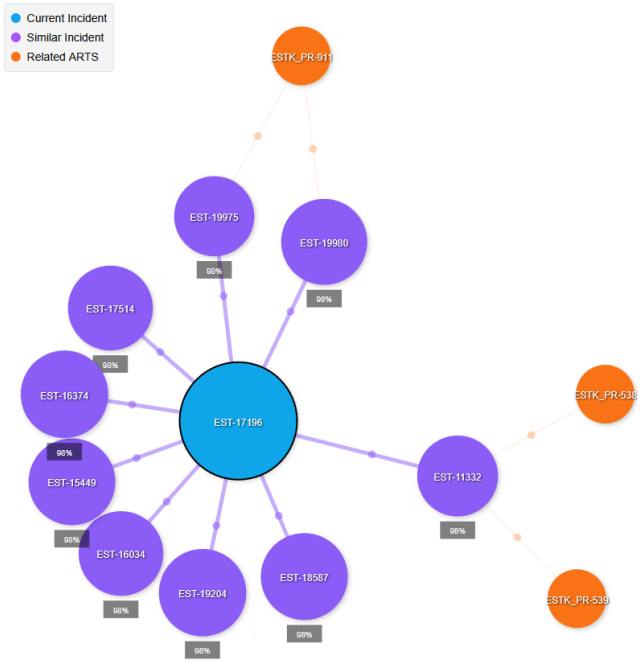

Knowledge Graph

Similarity Graph (t-SNE)

Dimensionality reduction projects high-dimensional vectors onto 2D maps. Clusters related ideas together, boosting discovery of similar incidents.

Knowledge Graph

Built using metadata on related incidents as links. Fully browsable and interactive: clicking on nodes reveals detailed incident information. Supports rapid, evidence-driven investigations across missions.

Key Capabilities

- Navigate relationships between incidents and anomaly reports
- Filter by dates, ground station origin, mission
- Combines AI-generated similarity with manual operator correlations

Agentic AI Architecture

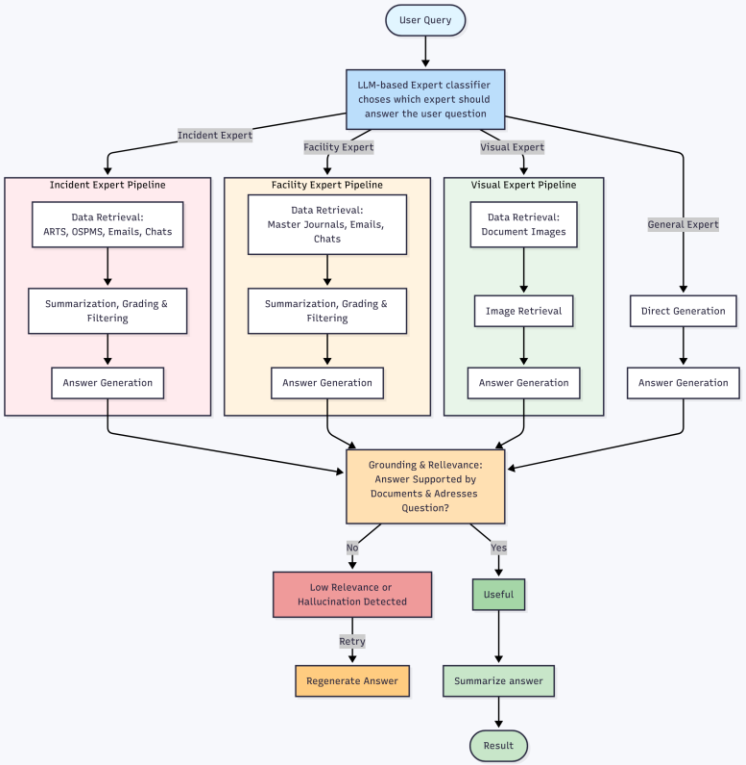

AI workflow Architecture

1 **Acronym Expansion**

Acronym expansion step using a predefined dictionary

2 **Specialist Routing**

LLM-based classifier that selects which specialist experts should handle the question

3 **Advanced Retrieval Step**

Retrieve relevant documents from different data sources

4 **Answer Quality Gate**

Grounding and relevance check; is the answer supported by the documents? Does it address the user's question ?

5 **Summarisation**

Summarise the answer to the user

Multi-Expert Investigation System

Workflow: Acronym expansion > LLM-based routing > Retrieval > Execution > Quality check > Summary

Incident Expert

ARTs, OSPMs, emails, chats

Ground Station Expert

Master journals, emails, chats

Visual Document Expert

PDFs, schematics, diagrams (ColPali)

General Expert

Fallback for general questions

Retrieval Mechanisms

Hybrid Search

Combines semantic similarity (vector) with BM25F keyword scoring for technical identifiers

Automatic Filtering

dspy-generated filters on metadata (dates, stations) for focused and consistent searches

Re-ranking

LLM grades retrieved chunks against user query; retry loop rewrites query if results insufficient

Quality Control Mechanisms

Space operations demand accuracy. Multiple validation layers ensure integrity and transparency.

1

Retrieval Quality Control

- Summarisation of retrieved chunks
- Re-ranking against user query
- Retry with rewritten query if insufficient documents

2

Answer Quality Control

- Hallucination check: answer grounded in documents?
- Relevance check: does it address the question?
- LLM-as-a-judge validation methods

3

Output Quality Control

- Source links provided with every answer
- Confidence scoring for generated conclusions
- Structured JSON validation for UI integration

Agent UI & Source Control

The UI helps the user assess the result quality by providing interactive data sources and confidence scores

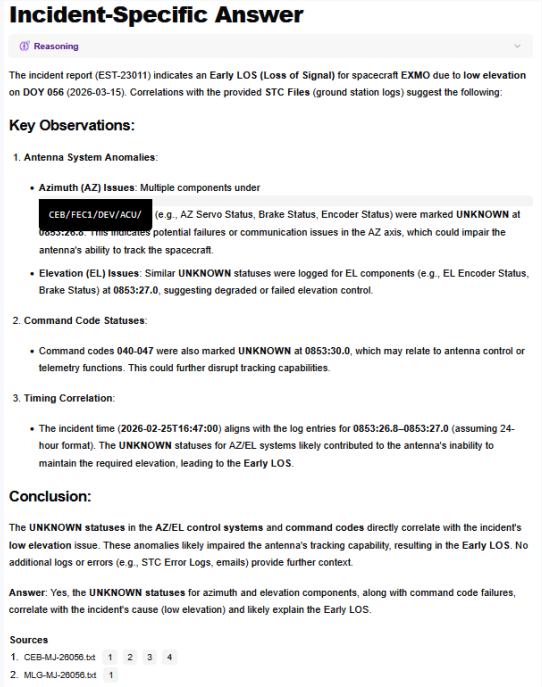

Incident-Specific Answer

Reasoning

The incident report (EST-23011) indicates an **Early LOS (Loss of Signal)** for spacecraft **EXMO** due to **low elevation** on **DOY 056** (2026-03-15). Correlations with the provided **STC Files** (ground station logs) suggest the following:

Key Observations:

1. **Antenna System Anomalies**:
   - **Azimuth (AZ) Issues**: Multiple components under `CEB/FEC1/DEV/ACU/` (e.g., AZ Servo Status, Brake Status, Encoder Status) were marked **UNKNOWN** at **0853:26.8**. This indicates potential failures or communication issues in the AZ axis, which could impair the antenna's ability to track the spacecraft.
   - **Elevation (EL) Issues**: Similar **UNKNOWN** statuses were logged for EL components (e.g., EL Encoder Status, Brake Status) at **0853:27.0**, suggesting degraded or failed elevation control.
2. **Command Code Statuses**:
   - Command codes **040-047** were also marked **UNKNOWN** at **0853:30.0**, which may relate to antenna control or telemetry functions. This could further disrupt tracking capabilities.
3. **Timing Correlation**:
   - The incident time (**2026-02-25T16:47:00**) aligns with the log entries for **0853:26.8–0853:27.0** (assuming 24-hour format). The **UNKNOWN** statuses for AZ/EL systems likely contributed to the antenna's inability to maintain the required elevation, leading to the **Early LOS**.

Conclusion:

The **UNKNOWN statuses** in the **AZ/EL control systems** and **command codes** directly correlate with the incident's **low elevation** issue. These anomalies likely impaired the antenna's tracking capability, resulting in the **Early LOS**. No additional logs or errors (e.g., STC Error Logs, emails) provide further context.

Answer: Yes, the **UNKNOWN statuses** for azimuth and elevation components, along with command code failures, correlate with the incident's cause (low elevation) and likely explain the Early LOS.

Sources

1. CEB-MJ-26056.txt 1 2 3 4
2. MLG-MJ-26056.txt 1

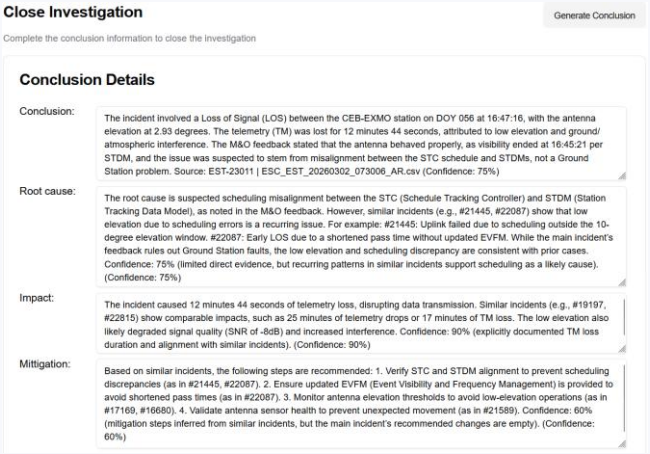

Close Investigation

Generate Conclusion

Complete the conclusion information to close the investigation

Conclusion Details

Conclusion:

The incident involved a Loss of Signal (LOS) between the CEB-EXMO station on DOY 056 at 16:47:16, with the antenna elevation at 2.93 degrees. The telemetry (TM) was lost for 12 minutes 44 seconds, attributed to low elevation and ground/atmospheric interference. The M&O feedback stated that the antenna behaved properly, as visibility ended at 16:45:21 per STDM, and the issue was suspected to stem from misalignment between the STC schedule and STDMs, not a Ground Station problem. Source: EST-23011 | ESC_EST_20260302_073006_AR.csv (Confidence: 75%)

Root cause:

The root cause is suspected scheduling misalignment between the STC (Schedule Tracking Controller) and STDM (Station Tracking Data Model), as noted in the M&O feedback. However, similar incidents (e.g., #21445, #22087) show that low elevation due to scheduling errors is a recurring issue. For example: #21445: Uplink failed due to scheduling outside the 10-degree elevation window. #22087: Early LOS due to a shortened pass time without updated EVFM. While the main incident's feedback rules out Ground Station faults, the low elevation and scheduling discrepancy are consistent with prior cases. Confidence: 75% (limited direct evidence, but recurring patterns in similar incidents support scheduling as a likely cause). (Confidence: 75%)

Impact:

The incident caused 12 minutes 44 seconds of telemetry loss, disrupting data transmission. Similar incidents (e.g., #19197, #22815) show comparable impacts, such as 25 minutes of telemetry drops or 17 minutes of TM loss. The low elevation also likely degraded signal quality (SNR of -8dB) and increased interference. Confidence: 90% (explicitly documented TM loss duration and alignment with similar incidents). (Confidence: 90%)

Mitigation:

Based on similar incidents, the following steps are recommended: 1. Verify STC and STDM alignment to prevent scheduling discrepancies (as in #21445, #22087). 2. Ensure updated EVFM (Event Visibility and Frequency Management) is provided to avoid shortened pass times (as in #22087). 3. Monitor antenna elevation thresholds to avoid low-elevation operations (as in #17169, #16680). 4. Validate antenna sensor health to prevent unexpected movement (as in #21589). Confidence: 60% (mitigation steps inferred from similar incidents, but the main incident's recommended changes are empty). (Confidence: 60%)

Agent UI & Source Control

The UI helps the user assess the result quality by providing interactive data sources and confidence scores

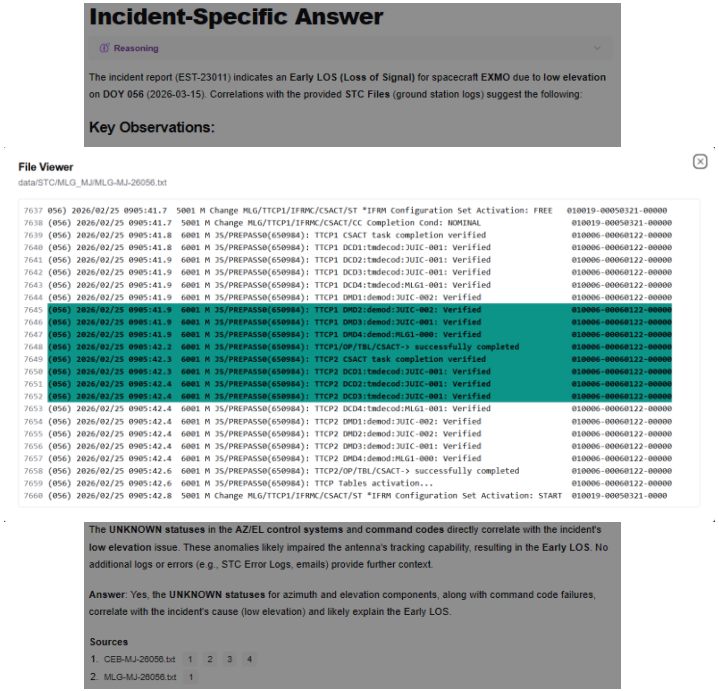

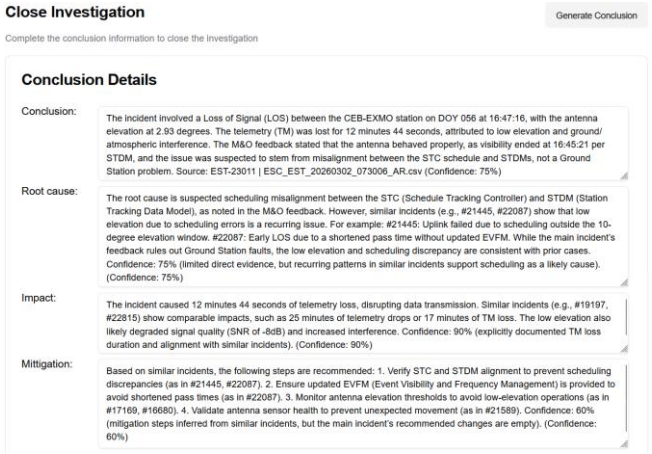

Automated Conclusion Generation

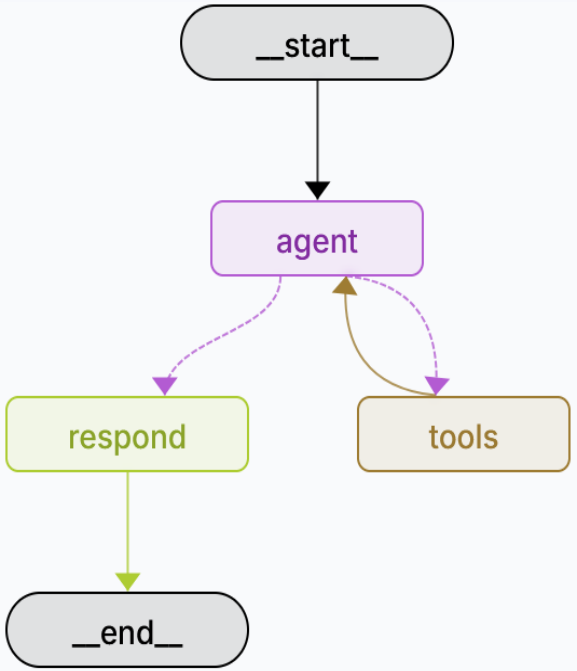

Conclusion Agent Architecture

Fully Autonomous Agentic Loop

Architecture: Tool-calling agentic loop with Qwen3:8B in thinking mode. Model reflects on actions and steers investigation direction.

Search Autonomy: Single search tool used autonomously with vector, keyword, or hybrid modes based on investigation needs.

Context Engineering: Agent state ingested at each loop to prevent context expansion. Duplicate filtering across subsequent queries.

Outputs: Proposes conclusions, root cause, system impact, and mitigation actions. Self-evaluates confidence scores.

Response Agent: Separate agent reformats findings into structured JSON for UI integration, ensuring flawless schema accuracy.

Human-in-the-loop: Generated conclusions are reviewed by operators before entering the incident database.

API-in-the-Loop

Enable autonomous agents in network-restricted secure environments

The Problem

Agent deployed on a secure server with restricted outgoing requests. Cannot reach PostgreSQL on another server to enrich vector DB results with relational metadata.

The Solution

Repurposed the Human-In-The-Loop (HITL) protocol: the API call arguments are embedded in the response body. The calling service (with DB access) processes it and resumes the agent conversation.

Why it works?

The HITL protocol does not rely on emitting a request from the server. Instead, the request is embedded in the response body of the initial call, enabling the external service to fulfil it and resume the conversation seamlessly.

Key Contributions

Unified Data Pipeline

Continuously harmonizes heterogeneous operational data sources into queryable knowledge bases

Multi-Expert Agentic System

RAG with quality controls (re-ranking, hallucination detection) tailored for mission-critical environments

API-in-the-Loop Protocol

Novel approach enabling autonomous agents within network-restricted secure environments

Deployed & Validated

Operational system validated through real-world use at ESOC ground stations with positive user feedback

Conclusion & Future Work

What We Achieved

An operational AI-driven incident investigation framework addressing data fragmentation and institutional knowledge loss in ground station environments. Integrates automated pipelines, interactive visualization, multi-expert AI agents, and autonomous reasoning, all within hardware and security constraints.

What Comes Next

- Unified Agentic Architecture: Migrate multi-expert system to fully autonomous tool-calling
- Real-Time Collaborative Investigation: Surface agent reasoning in real-time for operator guidance
- Co-Pilot Experience: Transform from answer-delivery tool to collaborative investigation partner

Thank you!

Dimitri Accad | d.accad@stariongroup.eu

STARION