

SpaceOps-2025, ID # 378

## **Bridging ESA Innovations with Commercial Space Operations: Integration of an Ops Companion for Satellite Decision Support**

**Satish Madhogaria<sup>a\*</sup>, Mahesh Nidugala<sup>a</sup>, Edoardo Cocci<sup>a</sup>, Bodo Werner<sup>a</sup>, Mahsa Taheran<sup>b</sup>,  
Evridiki V. Ntagiou<sup>c</sup>**

<sup>a</sup> *Telespazio Germany GmbH, Darmstadt, Germany, \*Satish.Madhogaria@telespazio.de*

<sup>b</sup> *constellr GmbH, Freiburg, Germany*

<sup>c</sup> *European Space Operations Centre, European Space Agency, Darmstadt, Germany*

\* Corresponding Author

### **Abstract**

As the space industry continues to expand with an influx of satellite missions, the efficient management and retrieval of information in critical operational scenarios becomes increasingly challenging. With each mission generating vast volumes of data and documentation stored in disparate formats, efficient information access and navigation tools are paramount. This paper presents the envisaged integration of OCAI, the Operations CompAnIon AI-based tool developed by the European Space Agency, into the commercial EASE-Rise mission control system developed by Telespazio Germany.

**Keywords:** virtual assistant, decision support, machine learning, mission operations

### **1. Introduction**

The fundamental problem addressed by this work lies in the amount, complexity, and heterogeneity of information associated with satellite missions. Typically, information regarding the satellite platform's design, including AIT test campaign reports, for example, and the operational aspects, such as user manuals and flight operation procedures, is encoded in diverse formats and stored across various platforms. Accessing pertinent information promptly is crucial for effective decision-making and problem resolution. However, as technology evolves rapidly and the size and diversity of mission elements increases, the sheer volume and heterogeneity of information available can overwhelm operators, potentially leading to delays in response and severe mission impacts. With the growing complexity of systems and data fragmentation, virtual assistants are increasingly valuable. The European Space Operations Centre (ESOC) has addressed this need by developing the Operations CompAnIon (OCAI) [1, 2, 3], which streamlines data retrieval, correlation, and analysis across multiple heterogeneous systems. Its core capabilities include advanced search and data correlation across sources of diverse data types (TM/TC, procedures, events, logs, space safety data, etc.), supporting queries based on human language, as well as specific space operations terminology. The tool is already operationally validated for 14 missions and expanded to diverse mission operations domains. In order to make use of the benefits and further demonstrate the generality of the tool, OCAI has been integrated with the EASE-Rise platform, Telespazio's native End-to-End solution for the smart operation of large satellite constellations. Thus, enhancing the decision-making processes for operations teams that utilize EASE-Rise for commercial satellite missions.

AI-based tools tailored to the needs of satellite mission operators represent a significant step towards addressing the challenges posed by the exponential growth of information in the space industry. Expanding the usage of tools like OCAI beyond institutional missions will drive further enhancements and development, supported by additional use cases and valuable feedback. This paper aims to provide valuable insights into the collaborative potential between ESA innovations and commercial space operations use cases, underscoring the enhanced decision-making capabilities and operational efficiencies achieved through this integration.

One of the missions currently being operated with the EASE-Rise mission control system is constellr's Copernicus Contributing Mission "HiVE". We envisage to use this mission as pilot mission for validation.

### **2. Overview of Each Component**

The advantages of integrating OCAI into EASE-Rise can be better understood by exploring the architecture of these components. Additionally, to better understand EASE-Rise as a mission management platform, an overview of one of its targeted missions - the pilot mission constellr HiVE - is provided.

## 2.1 EASE-Rise

EASE-Rise is Telespazio’s innovative software solution to disrupt ground segment and operations processes that have been the traditional method of mission control for decades. Using the latest and most efficient IT technologies, it provides an end-to-end solution that can exchange data between all parts of the platform to improve the efficiency of the system. An architectural overview of EASE-Rise is shown in Fig. 1. EASE-Rise provides three primary components that work together to enable efficient mission operations:

- Configuration Application – Used to define mission parameters, its assets and set up various operational elements for different services.
- Mission Control Application – Provides real-time monitoring, control and status tracking of mission and its assets.
- Mission Planning Application – Focuses on automating tasks, optimizing asset utilization, and improving mission execution efficiency.

Each of these components can potentially generate and consume large volumes structured and unstructured data (depending on the missions) making data retrieval and analysis increasingly complex as mission operations scale.

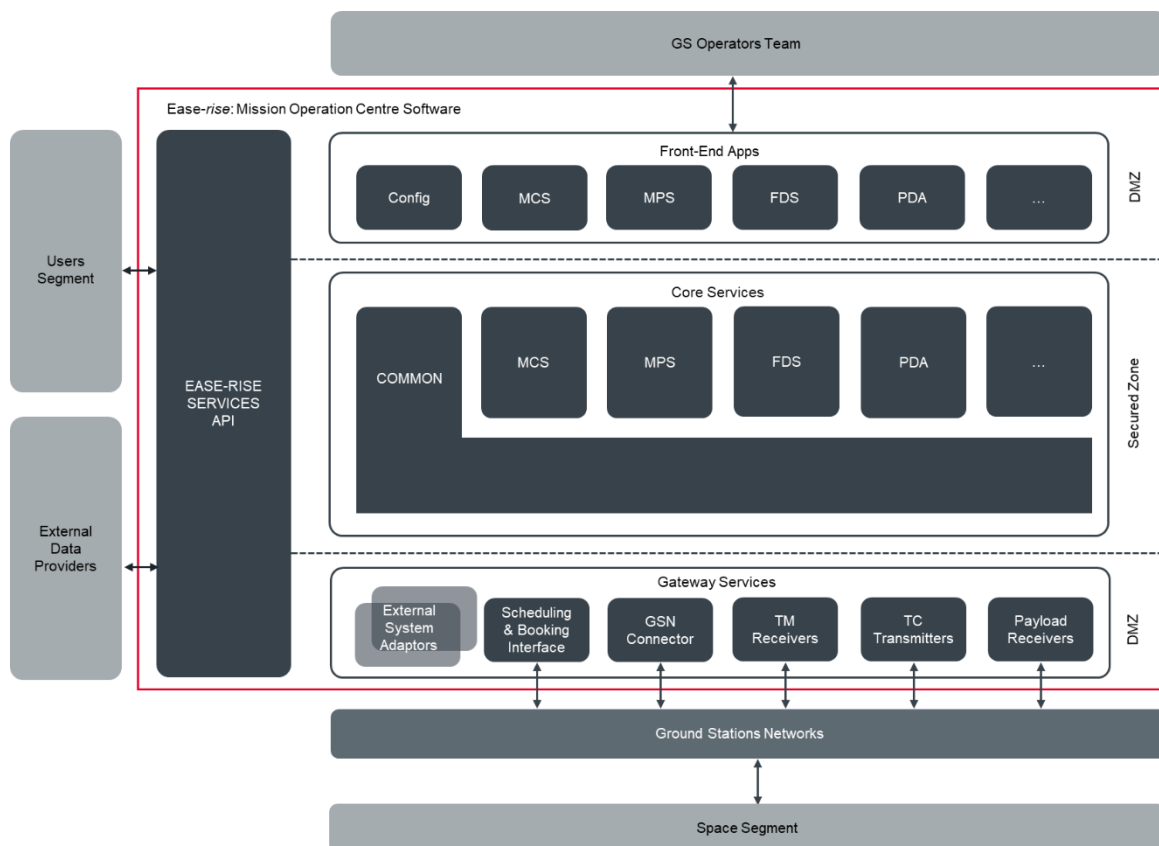


Fig. 1. EASE-Rise - Architectural Overview

## 2.2 Envisaged Demo Mission: Constellr's Copernicus Contributing Mission 'HiVE'

HiVE (High-Precision Versatile Ecosphere Monitoring Mission) is constellr’s microsatellite constellation, delivering high-resolution thermal intelligence for commercial and institutional applications. With its first satellite launched in January and the second one planned for a launch in June 2025, HiVE is gradually becoming operational, providing detailed thermal infrared (TIR) data for agriculture, infrastructure monitoring, urban planning, and environmental applications. At the core of HiVE’s capabilities is its cryocooled long-wave infrared (LWIR) imaging system, which enables sub-30m spatial resolution and high radiometric accuracy — previously achievable only with larger, more expensive satellite platforms.

Constellr operates its own satellites, with an advanced chain of tasking architecture spanning across different teams (see Fig. 2). The composite mission operations segment integrates various in-house developed and procured tools. EASE-Rise is used as the mission control system, with a number of custom services developed for HiVE TM/TC pipelines.

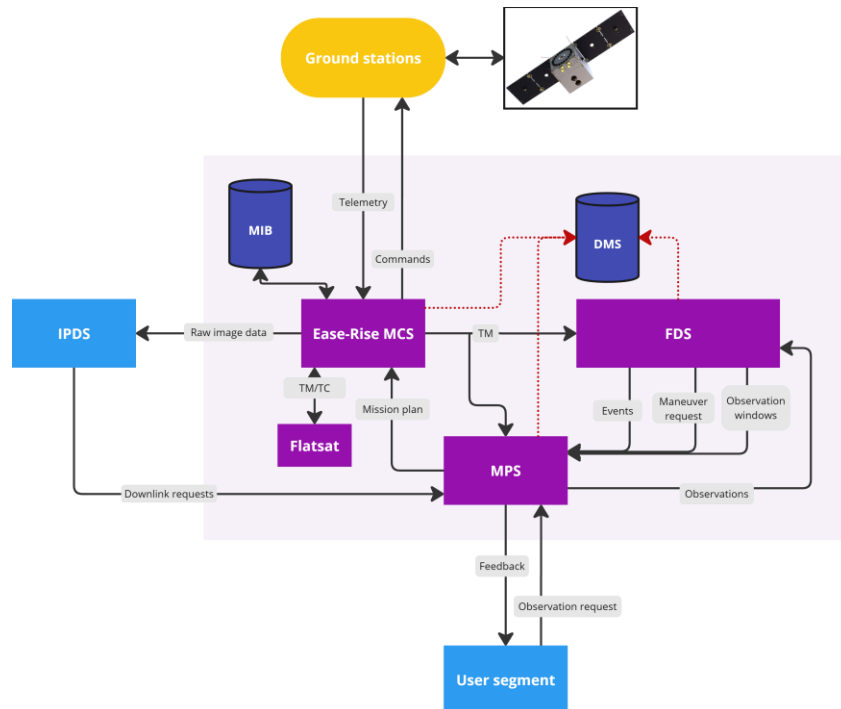


Fig. 2. Constellr HiVE Ground Segment architecture

### 2.3 OCAI

The European Space Operations Centre (ESOC) has developed the Operations CompAnIon (OCAI), an AI-driven tool designed to enhance decision-making for flight control teams by simplifying the complexity and fragmentation of data systems in operational control centres. Created to address the challenges of handling multiple heterogeneous systems, OCAI applies Natural Language Processing and Knowledge Engineering techniques to serve as a centralized access point, making Mission Operations information readily available to operators and engineers.

Tailored to support both operations and training, the tool consolidates all necessary insights within a single, intuitive application. It enables seamless connections between different sources of information and provides advanced correlation and search functionalities, as well as automated metadata extraction—including related and inferred terms. OCAI processes a wide variety of data types, from telemetry and telecommand (TM/TC) archives and operational procedures to anomaly reports, mission manuals, and both system-generated and human-generated logs.

Currently in use across 14 ESOC missions to facilitate rapid information retrieval, knowledge discovery, and new Flight Control Teams’ training, OCAI is continuously evolving with new AI capabilities. These include the integration of Large Language Models to generate precise responses and references throughout the entire knowledge base, automate reports, and streamline specific tasks, all developed in close collaboration with its users. Thanks to its modular structure, OCAI is also being tested in various fields, such as Earth Observation, Scientific Research, and Human-Robotic Exploration. Deployed in an operational setting, the system also incorporates advanced platform functionalities that enhance its reliability, including monitoring dashboards, data segregation mechanisms, and authorization controls.

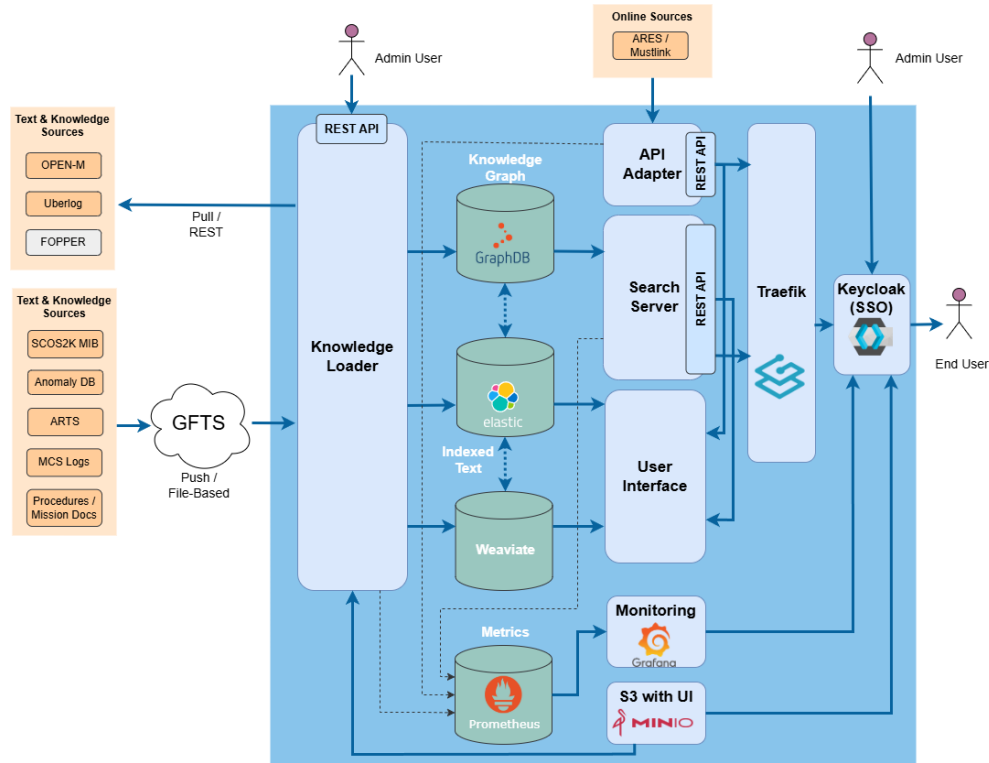


Fig. 3. OCAI - High-level architecture

### 3. Integration Benefits

When fully operational and handling multiple missions, EASE-Rise can generate massive amounts of data, including

- **Configuration Data** – Mission setups, asset configurations, operational parameters, service-specific details etc.
- **Telemetry and Logs** – Sensor readings, system health reports, control logs, mission status updates and operational anomalies.
- **Historical Data & Analytics** – Past missions’ data, performance metrics, resource utilization trends, failure analysis reports.
- **Procedures and Manuals** – Standard operating procedures (SOPs), mission guidelines, technical manuals, emergency protocols etc.

Traditional search mechanisms often fail to provide quick, contextual and accurate retrieval of this data leading to slower decision-making where Users struggle to find relevant information in time-sensitive situations. Additionally, information is spread across different applications making it difficult to correlate data from multiple sources.

Integrating OCAI, an AI-powered search engine into EASE-Rise can significantly enhance data retrieval and decision-making by enabling AI-powered context-aware search across multiple data sources.

### 4. Technical Integration Approach

Both OCAI (Fig. 3) and EASE-Rise (Fig. 1) follows micro-service architecture and therefore, integration of OCAI into EASE-Rise can be realised through RESTful APIs with only a few adaptations. Below, we describe few scenarios on how EASE-Rise can interface with OCAI’s backend modules:

#### 4.1 EASE-Rise interfacing with OCAI's Knowledge Loader

OCAI's Knowledge Loader (see Fig. 3) is responsible for creating knowledge graph that captures domain information and represented as graph. In semantic search applications such as OCAI, ontologies [4] play a crucial role in providing a structured representation of knowledge, helping search engines deliver more relevant and context-aware results. Currently, OCAI provides a specific set of ontologies predominantly used in the mission operations domain. However, following the OCAI Ontology approach, these ontologies are designed in a modular way, allowing arbitrary extensions with entities from multiple domains. There are no constraints on extending them to other domains, missions, or products.

Since the ontologies within OCAI were built for mission operations, EASE-Rise can leverage the existing ontologies or expand them according to the naming convention entities in EASE-Rise. Fig. 4 depicts the ontology representation for different entities in EASE-Rise. For example

- **Configuration Data** may contain entities such as *Mission*, *Metrics*, *Parameters*, *Commands* and *Components*.
- **Telemetry and Logs** can include *Activity*, *Command*, *Telemetry Packet*, *Telemetry Parameter*, and *Event*.
- **Historical Data** may contain *Missions*, *Telemetry Parameters*, *Telemetry Packets*, and *Procedures or Command Stacks*
- **Manuals** and documentation reference *Procedures*.

Concept expansion in OCAI establishes in-depth relationships between these entities deriving a knowledge graph using lexical analysis. Any additional ontology mappings can be extended alongside the predefined ones to enhance adaptability and interoperability.

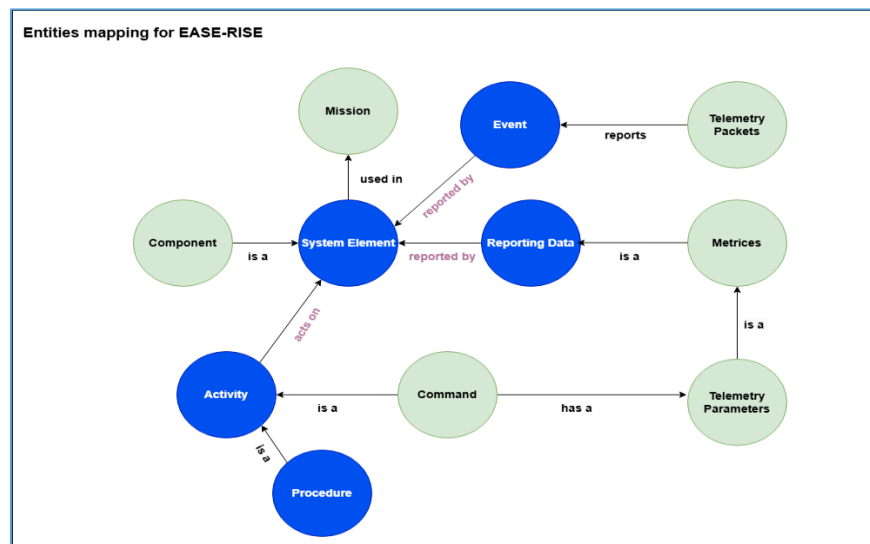


Fig. 4. Entities mapping for EASE-Rise

The integration of OCAI into EASE-Rise requires proper data handling from the EASE-Rise data source, which involves data fetching, data extraction, and mapping each related data point using concept expansion. The processed data is then uploaded into a graph database by deriving a knowledge graph and indexing it in ElasticSearch followed by NLP related tasks such as concept expansion [4, 5] and entity extraction for efficient retrieval.

OCAI supports multiple fetchers and extractors tailored to ESOC's mission operations which retrieves text files and other form of data via API calls. It ensures seamless handling of different data formats—structured (CSV, JSON), unstructured (PDF, text), and streamed (logs, telemetry)—while enabling effective transformation,

storage, and semantic search capabilities. EASE-Rise can leverage these existing fetchers and extractors while also implementing new ones to accommodate its diverse data sources. The final step in document processing involves reading data from previous streams and performing various NLP-related operations before indexing the processed data into *ElasticSearch* and *Weaviate*. Fig. 5 shows how the data generated by EASE-Rise are indexed and stored for efficient information retrieval via OCAI's knowledge loader pipeline.

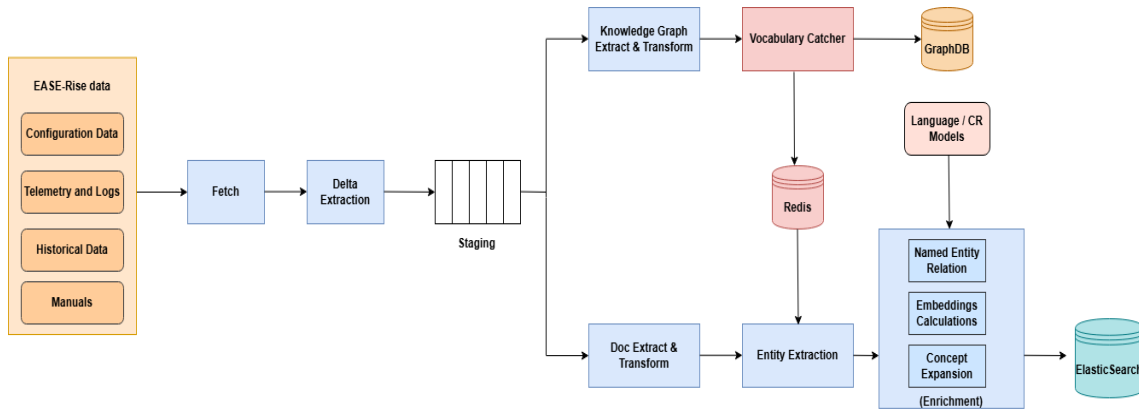


Fig. 5. Data transformation using OCAI's existing functionalities for efficient information retrieval.

#### 4.2 EASE-Rise Interfacing with OCAI's Search Server and API Adapter

The search mechanism plays a crucial role in semantic applications. To support efficient information retrieval, OCAI provides two Elasticsearch indexing methods. The first is keyword indexing, which maintains references to documents containing specific keywords. This enables the efficient application of the BM25 algorithm—a variant of the Term Frequency-Inverse Document Frequency (TF-IDF) weighting scheme—for relevance scoring. The second is K-nearest neighbour (KNN) indexing, which stores document embeddings and facilitates approximate nearest neighbour searches using cosine similarity. This approach allows the system to identify documents that are semantically similar to a given query vector.

OCAI supports several search mechanisms built on top of these indexing strategies. The simple keyword search relies directly on Elasticsearch's BM25-based scoring function and is particularly effective for structured queries, such as those involving document identifiers, due to its speed and relevance. The keyword-enhanced search begins with a traditional keyword-based retrieval and then re-ranks the results by calculating the cosine similarity between the query embeddings and the stored document embeddings. This method is especially useful when keyword search returns a large number of results, helping to bring the most semantically relevant documents to the top. Lastly, semantic search operates purely on embedding similarity, using KNN to retrieve documents that are close to the query vector in the embedding space. This approach is most effective when keyword-based methods produce sparse results, or when the query represents abstract or complex concepts not directly expressed through keywords.

The search logic combines results from multiple engines, ensuring consistent scoring to enable easy integration of new techniques and support post-processing or enrichment. OCAI uses three engines: the Vocabulary Engine enriches search terms using the Knowledge Graph and document properties from Elasticsearch; the Document Engine retrieves documents based on keyword or hybrid techniques, with indexing-time features like concept expansion handled by the Knowledge Loader; and the API Calls Engine, a pluggable component, enriches results by mapping Knowledge Graph entities to predefined API calls via the API Adapter. Each API call inherits the score of its corresponding Knowledge Graph entity.

The repositories in OCAI provide higher-level APIs to Elasticsearch and the Knowledge Graph, abstracting the complexity of their respective clients for the search engines. There are two main repositories: *DocStore* Repository, which bridges the Elasticsearch client with the OCAI search API. It manages document formatting, filter conversion, result highlighting, and avoids returning excessively long documents to preserve performance. It supports keyword and embedding-based search, recalculates score using embedding distances, handles regex

queries, cleans concept expansion artefacts, supports document lookup by ID, and retrieves stored embeddings. The GraphDB Repository maps operations to SPARQL queries and supports node lookup by ID or vocabulary term, subclass retrieval, and extraction of node metadata for UI presentation. The API adapters are mission operation-specific modules. They are direct REST endpoints to different domains, retrieving information directly integrated with the search server.

## 5. Status and Future Enhancements

Currently, we are working on the integration of the tools as described in chapter 4. To the submission date of this paper, EASE-Rise is being optimized in an agile approach to satisfy evolving operational demands of its operated missions, constellr HiVe being one of them. This envisaged demo mission has been launched in Jan 2025 leading to a growing data set in terms of both telemetry and documentation which is planned to be used for validation of the integrated system.

Also, OCAI itself is undergoing a continuous enhancement - driven by operator's experience and evolving demands of institutional and commercial operation environments. Focus of the near future will lie on establishing a multi-agent architecture. Multi-agent architectures allow AI-driven systems to be more modular, scalable, and adaptable. By integrating insights from diverse sources of a ground segment system, the system can form a more accurate and comprehensive understanding of the mission's status for the operator. This upcoming multi-agent architecture optimally aligns with a micro-serviced based missions control system like EASE-Rise.

Further ongoing efforts prioritize supporting advanced approaches such as Retrieval-Augmented Generation (RAG) [6], which enhances information retrieval by providing generative summaries of document results. RAG is a widely used approach and in the light of expanding OCAI's impact to complex and multi-agent like architectures, it can help to integrate and contextualize new information efficiently.

Not only for institutional but also for commercial space operations, explainability and qualification are essential to ensure user trust and adoption of LLM-based systems. Explainability provides transparency into how decisions are made, aiding in the identification of biases or errors, while qualification ensures that data and models meet established standards for accuracy and reliability. Industrial contracts are currently focused on ensuring OCAI's processes comply with these explainability and qualification standards.

## 6. Conclusion

As the space industry evolves toward increasingly complex and data-intensive missions, integrating AI-powered decision support systems is no longer a luxury but a necessity. The integration of OCAI with the EASE-Rise mission control platform marks a pivotal step in bridging ESA's innovative operations solutions with the dynamic demands of commercial space operations. By streamlining information retrieval, enhancing interoperability, and improving situational awareness, this collaboration demonstrates the tangible benefits of AI-driven tools in mission-critical environments.

Looking ahead, the evolution of OCAI and its integration into commercial operations environments will continue to push the boundaries of intelligent automation in space operations. Advancements in multi-agent architectures, Retrieval-Augmented Generation, and explainability frameworks will ensure that AI-based solutions remain both powerful and trustworthy. As these technologies mature, they will pave the way for more autonomous, resilient, and efficient satellite operations, reinforcing AI's role as a cornerstone of the next generation of space exploration and commercial mission management.

By embracing these innovations and fostering collaboration between institutional and commercial players, we are not just optimizing operations—we are shaping the future of spaceflight itself.

## References

- [1] E. Ntagiou, Intelligent Operations Assistant – ESA Statement of Work, 2021.
- [2] G. De Canio, J. Eggleston, A. M. Palowski, M. Schmidt, P. Collins, B. Sousa, D. Mesples, R. Santos, E. Ntagiou, OCAI: the Operations CompAnIon to support decision making of flight control teams, 17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023.
- [3] E. Ntagiou, T. Nogueira, M. Baab, L. Palladino, P. Leitão, L. Medina, B. Werner, C. Texeira, J. Eggleston, Smart Space Operations: OCAI's Contribution to Operational Excellence, IAF Space Operations Symposium, Milan, Italy, 14-18 October 2024.
- [4] G. Konstantinidis, G. Flouris, G. Antoniou, V. Christophides, A Formal Approach for RDF/S Ontology Evolution, 18th European Conference on Artificial Intelligence, Patras, Greece, July 21-25, 2008, DOI: 10.3233/978-1-58603-891-5-70
- [5] Nils Reimers, Iryna Gurevych, Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks, 2019, DOI: <http://arxiv.org/abs/1908.10084>
- [6] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Kuttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-augmented generation for knowledge-intensive NLP tasks, Advances in Neural Information Processing Systems (NeurIPS), Online, 6 - 12 December 2020, DOI: <https://arxiv.org/abs/2005.11401>