Ref. Ares(2026)6517945 - 29/06/2026



Enhancing Copernicus Security Services – EU governmental crisis management hub for forced population displacement

EO Data Processing Pathway Selection Guide, D5.4

WP5 – Gap analysis towards maximising the practical utility of the Copernicus Security Service

Funded by the European Union

This project has received funding from the European Union's Horizon Europe research and innovation programme under GA 101190051. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

Lead Contributor	Evangelos Michos (C3I)
Contributors	Natalia Czortek (CREO)
Reviewers	Charalampos Zafeiropoulos (ICCS), Ioannis Kavouras (ICCS)

Due Date	30/09/2025
Delivery Date	29/09/2025
Type	R - Document, report
Dissemination Level	PU – Public
Keywords	On-board Processing, Ground-based Processing, Hybrid Processing, Data Fusion, Communication Requirements, Latency, Bandwidth, Cost-Benefit Analysis, Mission Design, Decision Framework, Risk Assessment

Document History

Version	Date	Description	Author	Description/ Action	Validation
0.1	02/06/2025	Internal draft with ToC	E. Michos (C3I)	Request for partners' input	C3I
0.2	10/06/2025	Request input in Literature Review (Chapter 3)	E. Michos (C3I)	Request assistance in Chapter 3 from partners	Partners
0.3	22/07/2025	Circulate Decision-Making framework (Chapter 4)	E. Michos (C3I)	Send the framework to the whole consortium for review	C3I
0.4	25/07/2025	CREO input	E. Michos (C3I)	CREO adds their input in Chapter 5 and cost-economic analysis	CREO
0.5	05/09/2025	First internal version	E. Michos (C3I)	Peer Review from ICCS. Review from GSH (WP Leader)	ICCS, GSH
0.6	11/09/2025	Peer Review	E. Michos (C3I)	Receive peer review	C3I
0.7	12/09/2025	Revised version	E. Michos (C3I)	Inclusion of peer review comments	C3I
0.8	17/09/2025	SAB Review	SAB	SAB security check.	SAB
0.9	25/09/2025	Final Draft	E. Michos (C3I)	Integration of SAB comments	C3I
1.0	29/09/2025	Final version	L. Panagiotopoulou (GSH)	Submission to EC	
1.1	16/06/2026	Revised version after Review	P. Pierettis	Added use case mapping and improvement (section 4.4.4)	C3I
1.1	26/06/2026	Updated Final version	L. Panagiotopoulou (GSH)	Submission to EC	

Legal Disclaimer

This document reflects only the views of the author(s). Neither the European Commission nor the Granting Authority (European Health and Digital Executive Agency) is in any way responsible for any use that may be made of the information it contains.

The information in this document is provided “as is” and no guarantee or warranty is given that the information is fit for any particular purpose. The above referenced consortium members shall have no liability for damages of any kind including without limitation direct, special, indirect, or consequential damages that may result from the use of these materials subject to any liability which is mandatory due to applicable law.

This document and the information contained within may not be copied, used, or disclosed, entirely or partially, outside of the THEIA consortium without prior permission of the project partners in written form.

© 2026 by THEIA Consortium.

Contents

Executive Summary 6
List of Tables 7
List of Figures 7
1. Introduction 10
1.1 Purpose and scope of the deliverable 10
1.2 Structure of the deliverable 10
1.3 References 11
2. EO Data Processing Context 12
2.1 What is EO Data Processing 12
2.2 EO Mission Workflows: From Data Acquisition to Insight 13
2.3 EO Data Characteristics and Processing Demands 16
2.3.1 EO Data Types 16
2.3.2 Latency, bandwidth, accuracy and volume considerations. 16
2.3.3 EO Data Implications 17
3. On-board vs Ground-based processing: Comparative Evaluation 19
3.1 Definition and Scope 19
3.1.1 What is On-board processing? 19
3.1.2 What is Ground-based processing? 20
3.2 Review of Existing Processing Architectures and Approaches 21
3.2.1 Review of EO Data Processing Literature 21
3.2.2 Summary of Key Missions using On-board/Edge Processing 24
3.2.3 Industry Trends and Technology Roadmaps 25
3.2.4 Research Gaps and Opportunities 28
3.3 Comparative Evaluation: Technical, Operational and Economic Analysis 30
3.3.1 Technical Comparison 30
3.3.2 Operational Considerations 31
3.3.3 Economic Evaluations (Cost-Benefit Analysis) 33
3.4 Cost Benefit Summary 36
3.5 Summary of findings 37
4. Criteria and Guidelines for Processing Pathway Selection 39
4.1 Mission-Specific Factors 41

4.1.1 Type of mission 41
4.1.2 Urgency and latency needs 42
4.1.3 Geographic coverage and revisit rates 44
4.1.4 Decision Matrices and Architecture Trade-offs 46
4.2 Decision-Making Framework 49
4.2.1 Structured flowchart / decision tree 50
4.2.1.1 Decision tree 51
4.2.1.2 Reasoning for the decision tree 52
4.2.1.3 Flowchart 55
4.2.2 Thresholds and scoring mechanisms 56
4.3 Tactical/Operational Risk Assessment 59
4.4 Example Decision Scenarios 61
4.4.1 Real-time coastal monitoring scenario 61
4.4.2 Post-event environmental damage assessment scenario 63
4.4.3 Hybrid mission profile scenario 65
4.4.2 Mapping of Processing Pathways to THEIA Use Cases 68
5. Lessons from EO Mission Case Studies 69
5.1 OPS-SAT (ESA) – On-Board Processing as a Technology Testbed 69
5.2 PRISMA (ASI) – Hyperspectral Imaging and Adaptive Processing 69
5.3 Sentinel-1 (ESA) – Ground-Based Processing at Scale 70
5.4 ECOSTRESS (NASA) – Thermal Data and Downlink Bottlenecks 71
5.5 HYPSO-1 (Norway) – Maritime Observation 72
5.6 Summary 72
6. Strategic Recommendations and Considerations 74
6.1 Recommendations for EO Mission Architects 74
6.2 Investment Priorities for On-Board Processing 76
6.3 Roadmap for Hybrid Processing Models 79
6.4 Interoperability and Risk Considerations 82
7. Conclusions 84
References 86

Executive Summary

D5.4 presents a comprehensive guide for selecting optimal data processing pathways (on-board versus ground-based) in Earth Observation (EO) missions. In response to the increasing complexity of EO technologies, the surge in data volumes and the critical need for timely information delivery, this document explores the technical, operational and economic trade-offs associated with each processing approach. Through detailed analysis, it provides mission architects with a structured decision-making framework tailored to the specific requirements of diverse EO scenarios, including factors such as latency sensitivity, mission type, bandwidth limitations and system scalability.

Drawing from recent advancements in space-based processing capabilities and grounded in a review of existing EO missions, the guide consolidates key insights into a set of selection criteria and strategic recommendations. The document also includes a series of literature-based case studies that demonstrate real-world implications of processing choices on data integrity and dissemination speed. Ultimately, D5.4 aims to support evidence-based decision-making and inform future investments in on-board processing technologies, aligning with broader goals of enhancing the performance, responsiveness and autonomy of European EO systems.

List of Tables

Table 1. EO Data Types 16
Table 2. EO Data Types (extended) 17
Table 3. Cost-benefit analysis summary 36
Table 4. On-board vs Ground-based EO Data Processing Comparative Matrix 37
Table 5. Alignment of EO Processing Architectures with Mission Types 46
Table 6. Factor-Based Comparison Table 47
Table 7. SWOT Analysis for on-board processing (extended) 48
Table 8. SWOT Analysis for ground-based processing (extended) 48
Table 9. Decision-framework structured pathway 51
Table 10. Reasoning behind the Decision-framework pathway 52
Table 11. Classifying the Weighted Score Average 57
Table 12. Scoring Criteria and Rationale 57
Table 13. Risk categories and mitigation techniques 60
Table 14. Evaluation of Processing Approaches for Real-Time Coastal Monitoring 62
Table 15. Scoring-Based Evaluation of Processing Suitability for an Early Warning Satellite Mission 62
Table 16. Evaluation of Processing Approaches for Post-Event Environmental Damage Assessment 64
Table 17. Scoring-Based Evaluation of Processing Suitability for a Post-Event Damage Assessment 65
Table 18. Evaluation of Processing Approaches for Border Surveillance using Hybrid Processing 66
Table 19. Scoring-Based Evaluation of Processing Suitability for a Border Surveillance Using Hybrid Processing 67
Table 20. Satellite systems specification (see [61]- [72]) 72
Table 21. Gap Analysis of Current On-Board Processing Capabilities 77
Table 22. Key interoperability elements supporting THEIA's integration with Copernicus services and EU EO infrastructure. 83

List of Figures

Figure 1. EO preservation workflow [4] 15
Figure 2. EO Data Generation Rates forecast [7] 18
Figure 3. On-board processing in Earth Satellites [11] 20
Figure 4. Comparison between on-board processing flow and the traditional ground processing flow [57] 40
Figure 5. High-level EO decision framework 50
Figure 6. Flowchart for the decision-framework 55
Figure 7. Roadmap for Hybrid Processing Models 81

List of Acronyms / Abbreviations

Acronym / Abbreviation	Explanation
AI	Artificial Intelligence
AIS	Automatic Identification System
API	Application Programming Interface
ASIC	Application-Specific Integrated Circuit
AUV	Autonomous Underwater Vehicles

CAPEX	Capital Expenditures
CNN	Convolutional Neural Networks
COTS	Commercial Off-The-Shelf
CSS	Copernicus Security Services
DBI	Downlink Bottleneck Index
DTN	Delay-Tolerant Networking
E2E	End-to-end
ECOSTRESS	Ecosystem Spaceborne Thermal Radiometer Experiment on Space Station
EMS	Emergency Management Service
EMSA	European Maritime Safety Agency
EO	Earth Observation
EOSDIS	Earth Observing System Data and Information System
ESA	European Space Agency
FPGA	Field-Programmable Gate Array
GA	Grant Agreement
GDPR	General Data Protection Regulation
GRACE	Gravity Recovery and Climate Experiment
HPC	High-Performance Computing
HySIS	Hyperspectral Imaging Satellite
ISO	International Organization for Standardization
KPI	Key Performance Indicator
LANCE	Land, Atmosphere Near real-time Capability for EOS
LEO	Low Earth Orbit
LT	Latency Tolerance
ML	Machine Learning
NASA	National Aeronautics and Space Administration
NRT	Near Real Time
NTNU	Norwegian University of Science and Technology
OGC	Open Geospatial Consortium
OPEX	Operating Expenditures
OTA	Over-the-Air
R&D	Research and Development
RF	Radio Frequency
RNN	Recurrent Neural Networks
SAR	Synthetic Aperture Radar
SATCEN	European Union Satellite Centre
SoC	System on Chip
STAC	SpatioTemporal Asset Catalog
SWaP	Size, Weight and Power
UAS	Unmanned Aircraft System
VPU	Vision Processing Unit
WCS	Web Coverage Service
WFS	Web Feature Service

WMS	Web Map Service
WP	Work Package

1. Introduction

This deliverable, **D5.4 – EO Data Processing Pathway Selection Guide**, corresponds to **T5.4** under **Work Package 5 (WP5)**, which focuses on system architecture, technical specifications and application scenarios. The aim of this report is to provide a structured and comprehensive decision-making guide to help EO mission architects choose between on-board and ground-based data processing approaches. The guide is grounded in a comparative evaluation of processing technologies, informed by technical, operational and economic considerations and supported by case-based insights drawn from existing EO missions to **improve the operational value chain of the Copernicus Security Service (CSS)**, strengthening the EU's autonomy in space-based intelligence systems.

1.1 Purpose and scope of the deliverable

The purpose of this deliverable is to **analyse and evaluate the trade-offs between on-board and ground-based data processing approaches in EO missions**, with a specific focus on operational relevance, mission efficiency, cost, latency, data integrity and future technological scalability. This guide:

1. **Defines and contrasts** the two data processing paradigms.
2. Explores the **implications of each choice** across a wide set of mission-specific factors, including data volume, bandwidth constraints, real-time needs and system complexity.
3. Introduces a **structured decision-support framework**, including scoring mechanisms and selection criteria, to assist EO architects in choosing the most appropriate processing pathway for their operational scenario.
4. Presents **lessons learned from literature-based case studies** and mission profiles that highlight the real-world consequences of different processing strategies.

The output of T5.4 is expected to feed into the technical and architectural decisions of WP5, supporting the integration of adaptable processing schemes in the evolving THEIA platform architecture.

1.2 Structure of the deliverable

The remainder of this deliverable is structured as follows:

Section 2 sets the scene by providing an overview of Earth Observation (EO) data processing, including mission workflows and key characteristics such as data types, latency, volume and accuracy. It also introduces the rationale for considering processing location-on-board or ground-

based-as a critical decision in EO mission design. **Section 3** presents a conceptual overview of on-board and ground-based data processing, offering clear definitions, historical context and an in-depth literature review. It outlines key technological trends, highlights relevant missions and identifies gaps and opportunities emerging from current research and development efforts. It also delivers a comparative analysis of on-board versus ground-based processing. It examines the technical, operational, economic and security aspects of each approach, providing a structured assessment of their respective advantages and limitations. The section concludes with a strengths-and-weaknesses matrix to support quick reference by mission architects. **Section 4** introduces a criteria-based selection framework tailored for EO mission planning. It discusses mission-specific decision factors, presents a decision tree and scoring mechanisms and provides example scenarios that demonstrate how to apply the framework in practice. **Section 5** showcases lessons learned from real-world EO missions. Five case studies illustrate diverse processing strategies and offer valuable insights into latency management, bandwidth constraints, data quality preservation and processing trade-offs. **Section 6** offers strategic recommendations to mission designers and EU stakeholders. It highlights areas for future investment, suggests directions for R&D in on-board technologies and outlines a roadmap for adopting hybrid processing models. The section also aligns THEIA outcomes with Copernicus Security Service priorities and EU technological sovereignty goals. Ultimately, **Section 7** concludes the deliverable by summarizing key findings and decision-making guidelines, reflecting on the broader implications for EO system design and future mission readiness.

1.3 References

▪ Project GA with No101190051.

▪ THEIA Partners CA

2. EO Data Processing Context

2.1 What is EO Data Processing

EO data processing refers to the systematic series of operations applied to raw data collected from various remote sensing platforms - including satellites, aircraft, drones and ground-based sensors - to transform it into meaningful information about Earth's physical, chemical and biological systems [1]. To manage and interpret the vast amount of data collected by EO systems, space agencies like NASA and ESA have established standardized processing levels [2] [3]. These levels define how much transformation the raw data undergoes before it becomes usable by scientists, engineers, or decision-makers. The levels are as follows:

- **Level 0 (raw data):** This is the unprocessed data directly received from the satellite's sensors. It includes everything the sensor "saw" but still contains instrumental noise, distortions and other artifacts. It's mainly used by system engineers and is not ready for scientific or operational use.
- **Level 1 (Calibrated and Georeferenced Data):** Time-referenced data annotated with ancillary information, including radiometric and geometric calibration coefficients and georeferencing parameters. At this stage, the raw data is corrected for sensor and platform effects - this includes radiometric calibration (making sure the brightness levels are accurate) and geometric correction (aligning the image correctly with Earth's surface coordinates).
- **Level 2 (Geophysical Measurements):** Derived geophysical variables at the same resolution and location as Level 1 source data. Here, the data is transformed into meaningful physical quantities - for instance, converting pixel brightness into measurements of vegetation cover, chlorophyll concentration, or atmospheric temperature. These values are aligned with the satellite's original resolution but now carry direct scientific meaning.
- **Level 3 (Spatially Organized Data):** Variables mapped on uniform space-time grids, with missing points interpolated and complete regions mosaicked together from multiple orbits. This data is aggregated, gridded and harmonized over space and time. It often combines observations from multiple passes or orbits to create consistent global or regional maps.
- **Level 4 (Model-Derived or Interpolated Products):** Model outputs or results from analyses of lower-level data, representing variables that were not directly measured but derived from these measurements. The final level goes beyond observation-it includes data products generated by models, simulations, or data fusion.

For THEIA, these levels are relevant because the framework must decide where in the chain (*on-board or on the ground*) different levels of processing should occur. For example, cloud masking (*Level 1–2*) can be done on-board, while full climate analysis (*Level 3–4*) remains ground-based.

2.2 EO Mission Workflows: From Data Acquisition to Insight

EO mission workflows represent the end-to-end (E2E) chain of activities that transform raw data collected by EO sensors into actionable insights. These workflows are designed to support a wide range of applications-from rapid disaster response to long-term climate monitoring-by ensuring the accuracy, accessibility and usefulness of the data. Each EO mission follows a structured set of phases, from sensor deployment to data processing, dissemination and ultimately, long-term preservation [4].

1. Mission Planning and Sensor Selection

EO workflows begin during the mission planning phase, where mission objectives are defined, areas of interest are selected and decisions are made about the types of sensors and platforms needed. This planning also involves orbit design, revisit rates, spatial resolution and data volume forecasts. These elements fundamentally shape the type and complexity of the data processing chain that follows [4].

2. Data Acquisition and Downlink

Once operational, the satellite collects data as it orbits the Earth. Depending on the platform's design, it may acquire data in bursts (e.g. over targeted areas) or continuously (e.g. for global monitoring). The raw sensor data (Level 0) is stored temporarily on-board until a communication window opens with a ground station, at which point the data is transmitted to Earth. Real-time transmission is possible for missions with high bandwidth or urgent use cases, but in many cases, on-board storage and queued downlink are necessary due to bandwidth constraints [4], [5].

3. Ground Reception and Pre-processing

Once received by a ground station, the data enters the preprocessing phase, where it is transformed into usable formats. This includes:

- ✓ Radiometric correction, to normalize sensor brightness values;
- ✓ Geometric correction, to align images with geographic coordinates;
- ✓ Atmospheric correction, to remove distortions caused by clouds, aerosols and gases;
- ✓ Data formatting and metadata tagging.

These procedures result in Level 1 and Level 2 data, depending on the amount of correction and geophysical interpretation applied. For EO missions with latency-sensitive applications (e.g.,

emergency management), this stage may occur partially or entirely on-board, enabling quicker response times [4].

4. Advanced Processing, Analytics and Dissemination

After preprocessing, the data is ingested into EO data centers or cloud infrastructures for higher-level processing. This includes data fusion, time-series analysis and application-specific processing (e.g., flood detection, land-use classification). Increasingly, these steps involve machine learning algorithms and AI pipelines, which automate complex pattern recognition tasks. The processed data is then made available to users via web portals, APIs, or direct distribution mechanisms, depending on user access levels and licensing [4], [6].

5. Archiving and Long-Term Preservation

A crucial, often underemphasized phase in the EO mission workflow is data. According to [4], EO data must be curated across the entire mission lifecycle to ensure that it remains accessible, interpretable and usable decades into the future. This requires:

- ✓ Data appraisal and consolidation (removing corrupted files, ensuring completeness),
- ✓ Documentation of processing pipelines and metadata,
- ✓ Integration of associated tools (e.g. visualization, calibration routines) and
- ✓ Assignment of persistent identifiers to support discoverability and citation.

Each phase involves activities such as defining the designated user community, estimating preservation costs and risks and preparing a comprehensive archive of both data and related documentation.

6. Reprocessing and System Feedback

The EO workflow is not strictly linear. As new algorithms, processing tools, or scientific requirements emerge, archived data can be reprocessed to improve quality or extract new insights. This is particularly relevant for long-term climate datasets, where consistent time-series and inter-sensor harmonization are critical. Reprocessing may also be triggered by mission updates, hardware degradation, or new calibration standards [4].

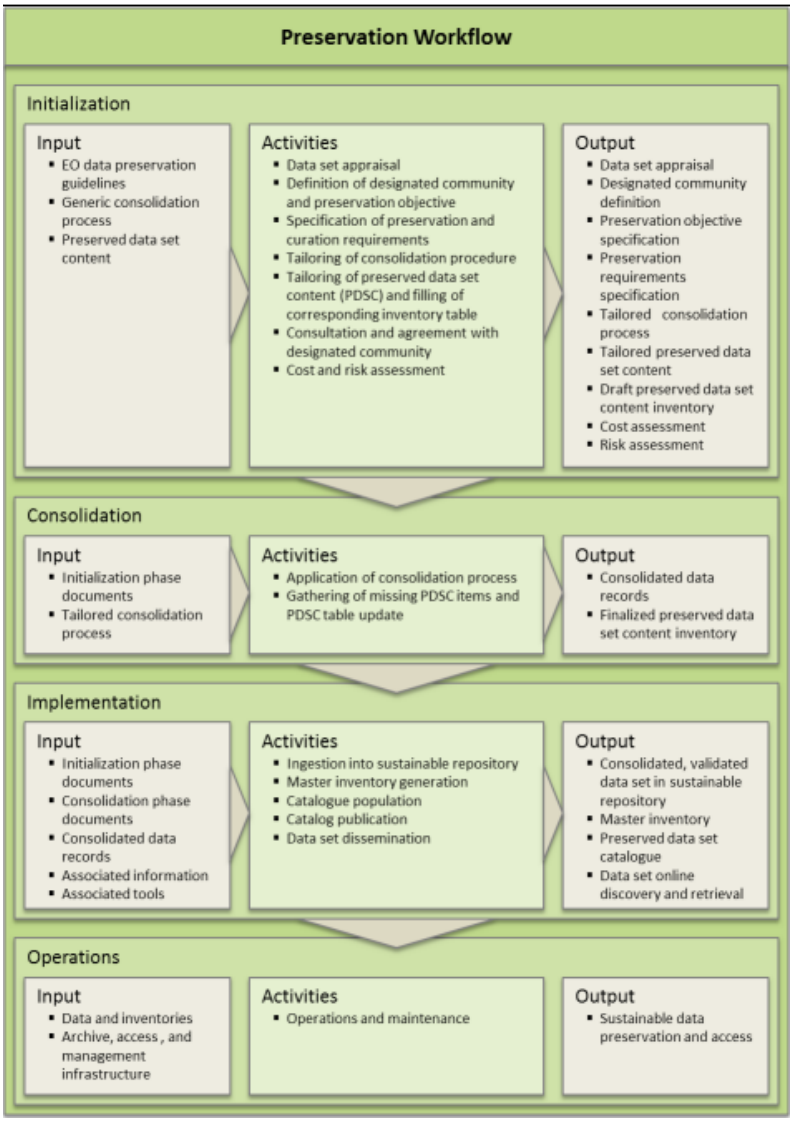


Figure 1. EO preservation workflow [4]

In more detail, as seen above in ***Figure 1***, the preservation workflow for EO space data is a structured four-phase process designed to ensure that EO datasets remain accessible, usable and trustworthy over the long term. It begins with the **Initialization Phase**, where data managers evaluate the value of preserving a dataset, define its intended users (the designated community), set preservation goals, identify requirements and estimate associated costs and risks. Next is the **Consolidation Phase**, during which data records are cleaned, completed and organized into a standardized format, producing a finalized inventory of what will be preserved. The **Implementation Phase** follows, involving the ingestion of these consolidated records into sustainable repositories, cataloging them for discovery and enabling online access and retrieval. Finally, the **Operations Phase** ensures the continued maintenance, monitoring and updating of

both the data and the systems that store them, guaranteeing long-term preservation. This workflow not only secures historic and current EO data but also lays the foundation for efficient preservation in future missions.

2.3 EO Data Characteristics and Processing Demands

2.3.1 EO Data Types

With information from [2], EO data encompasses a variety of sensor types, each tailored to capture specific aspects of Earth's environment.

Table 1. EO Data Types

Data Type	Information
Optical Data	Optical sensors capture reflected sunlight in the visible and near-infrared spectra, producing images akin to conventional photography. These sensors are invaluable for applications such as land-use mapping, vegetation monitoring and urban planning. However, their reliance on sunlight means they cannot operate effectively during night time or in cloudy conditions.
Synthetic Aperture Radar	SAR is an active sensing technology that emits microwave pulses toward the Earth's surface and records the reflected signals. This capability allows SAR to penetrate cloud cover and collect data regardless of weather conditions or time of day. SAR data is particularly useful for monitoring surface deformation, detecting changes in land cover and assessing disaster impacts [8], [9].
Thermal Data	Thermal sensors detect emitted infrared radiation, measuring surface temperatures. This data type is essential for applications such as monitoring volcanic activity, assessing soil moisture levels and detecting thermal pollution in water bodies. Thermal imaging is effective both day and night, providing continuous monitoring capabilities.
Radio Frequency Data	RF sensors monitor radio emissions from Earth's surface and atmosphere. They are employed in applications like tracking maritime and aerial traffic, detecting illegal transmissions and studying ionospheric phenomena. RF data processing requires specialized techniques to filter and interpret the complex signals received [10].

2.3.2 Latency, bandwidth, accuracy and volume considerations.

The performance and usability of EO data are not defined solely by sensor quality or mission scope, but also by key operational parameters such as latency, bandwidth, accuracy and data volume. These factors critically influence data transmission strategies, processing infrastructure requirements and ultimately, the timeliness and reliability of insights delivered to end-users.

Table 2. EO Data Types (extended)

Data Type	Information
Latency	Latency refers to the time delay between data acquisition by the satellite and its availability to end-users. Low-latency data is critical for applications requiring near real-time information, such as disaster response and weather forecasting. Achieving low latency necessitates efficient data processing pipelines and rapid data transmission methods.
Bandwidth	Bandwidth pertains to the data transmission capacity between the satellite and ground stations. High-resolution sensors generate large volumes of data, demanding substantial bandwidth for timely downlink. Limitations in bandwidth can lead to data bottlenecks, necessitating the development of data compression techniques and selective data transmission strategies to optimize the use of available bandwidth.
Accuracy	Accuracy in EO data encompasses both positional accuracy and measurement precision. High accuracy is vital for reliable analysis and decision-making. Ensuring accuracy involves rigorous calibration of sensors, application of correction algorithms and validation against ground truth data. For instance, geometric corrections are applied to align images with geographic coordinates, while radiometric corrections adjust for sensor and atmospheric influences.
Volume	The volume of EO data is influenced by factors such as spatial, spectral and temporal resolutions. High-resolution, multi-spectral and frequent observations result in substantial data volumes. Managing this data requires robust storage solutions, efficient data processing capabilities and scalable infrastructures. Cloud computing platforms and distributed processing systems are increasingly utilized to handle the vast amounts of data generated by modern EO missions.

2.3.3 EO Data Implications

The exponential growth of EO data is reshaping both the opportunities and challenges in the space sector. As highlighted by [7], EO data is expected to reach a staggering 2 exabytes by 2032, driven by a compound annual growth rate of 11%. This increase is largely fuelled by advances in sensor technologies, including hyperspectral imaging and improvements in sensor optics, which now enable satellites to capture data with unprecedented spatial, spectral and temporal resolution. These technological developments have allowed satellites to generate larger and more complex datasets than ever before. At the same time, the proliferation of satellite constellations with frequent revisit capabilities ensures continuous and dense streams of information. While this brings enormous potential for real-time monitoring and data-driven applications, it also places heavy demands on processing infrastructures, data transmission systems and analytics capabilities.

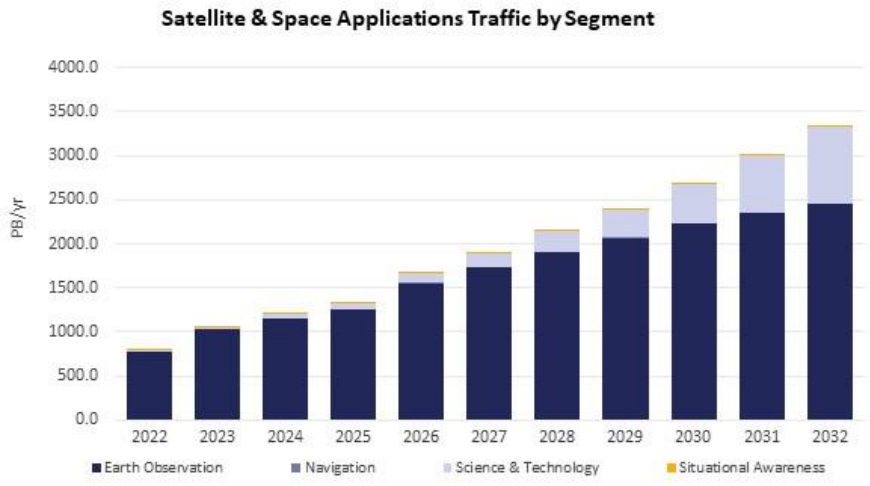


Figure 2. EO Data Generation Rates forecast [7]

Managing this surge in data requires significant investment in ground-based infrastructure, including next-generation ground stations, high-bandwidth communication systems and cloud-based processing platforms. Commercial satellite operators and government agencies alike must address mounting challenges related to data storage, algorithmic analysis and personnel expertise. It is emphasized that as EO data becomes more integral to sectors such as emergency response, climate monitoring and national security, stakeholders need to consider robust data management frameworks and spectrum management strategies to ensure reliable, high-speed transmission. These developments reinforce the importance of scalable, interoperable and real-time EO data processing pipelines, whether implemented on-board satellites or through advanced ground-based infrastructure.

3. On-board vs Ground-based processing: Comparative Evaluation

This chapter provides a comprehensive analysis of on-board and ground-based data processing for EO missions. It begins by defining the two approaches and their operational scope, followed by an in-depth review of existing literature, mission experiences and industry trends. The second half of the chapter delivers a comparative analysis based on technical, operational and economic criteria, guiding the selection of the most suitable processing architecture for future EO missions.

To ensure that the literature review captured the full breadth of relevant work across technical, operational and economic dimensions, C3I actively sought contributions from all project partners. Requests for assistance were circulated during the preparation of this chapter, inviting partners to share domain-specific references, recent studies and institutional knowledge. This collaborative approach was intended to extend the knowledge base beyond C3I's own survey, ensuring that the review reflects a comprehensive and balanced perspective. The inclusion of multiple viewpoints strengthens the credibility of the comparative analysis that follows and demonstrates the collective effort of the consortium in building a shared evidence base for subsequent methodological and strategic recommendations.

3.1 Definition and Scope

3.1.1 What is On-board processing?

On-board processing refers to the execution of data processing tasks directly on a satellite, utilizing its integrated computing systems. This approach includes the acquisition, selection, compression and storage of sensor data, as well as its transmission to ground stations. By processing data in situ, satellites can reduce the volume of information that needs to be downlinked, thereby optimizing bandwidth usage and enhancing the efficiency of EO missions. The significance of on-board processing has grown in tandem with the increasing complexity of satellite payloads and the vast amounts of data they generate. Traditional EO satellites often produce more raw data than can be transmitted to ground stations in real-time. To address this challenge, on-board processing employs various signal processing and compression techniques to reduce data volume before transmission. This capability is particularly crucial for missions requiring rapid data analysis and response, such as disaster monitoring and environmental assessments [11].

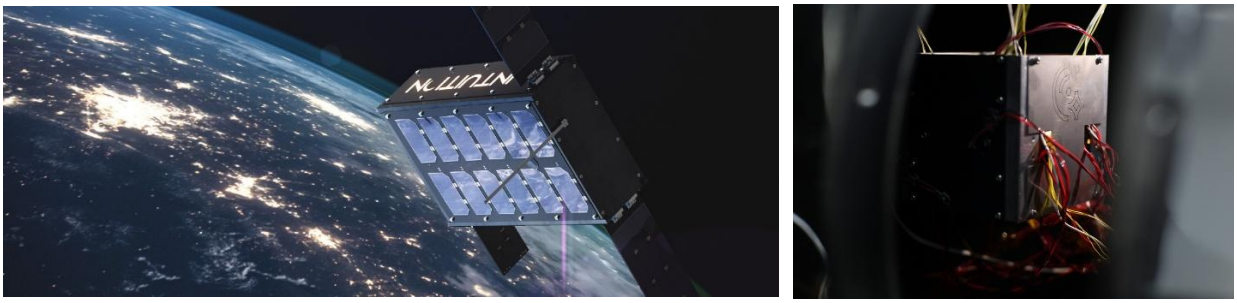


Figure 3. On-board processing in Earth Satellites [11]

Advancements in AI and ML have further expanded the potential of on-board processing. For instance, the Phi-Sat-1 mission demonstrated the use of AI algorithms to filter out images obscured by clouds, ensuring that only high-quality data is transmitted to Earth [12]. Similarly, the Intuition-1 satellite employed on-board AI to process hyperspectral images, significantly reducing the data volume sent to ground stations and facilitating timely analysis [13].

3.1.2 What is Ground-based processing?

On the other hand, **ground-based processing** refers to the handling and analysis of data received from EO satellites at terrestrial facilities. Once raw data are transmitted from satellites to ground stations, they undergo a series of processing stages to transform them into usable information products. These stages are categorized by NASA's EOSDIS into levels ranging from Level 0 (*raw data*) to Level 4 (*model output or analysis results*) [3]. Moving on, the infrastructure supporting ground-based processing includes satellite tracking ground stations equipped with antennas and receivers to collect data from orbiting satellites. For instance, the Alaska Satellite Facility operates multiple antennas to downlink data, which are then processed and distributed for various scientific and operational applications [14]. The ground system is responsible for collecting and distributing the mission's most valuable asset: the data. Key considerations for ground systems include data volume to satisfy mission requirements, location of ground assets relative to mission orbit parameters, budget limitations and latency requirements [15]. Ground-based processing offers several advantages, including access to substantial computational resources, facilitating complex data analyses and long-term storage solutions. However, it also presents challenges such as latency in data availability due to transmission times and potential bottlenecks in processing large data volumes.

3.2 Review of Existing Processing Architectures and Approaches

3.2.1 Review of EO Data Processing Literature

The growing complexity of EO missions has intensified the debate between the use of on-board and ground-based processing approaches. In their foundational roadmap, the authors of [16] provide a systems-level perspective on the evolution of data handling technologies aboard spacecraft. Their study, originating from ESA's Data Systems Division, systematically highlights how the exponential increase in sensor resolution and data rates has outpaced the downlink capacity of traditional communication systems. They argue that space-qualified computing components, while improving slowly, remain years behind terrestrial equivalents, creating a performance gap. To address this, the authors forecast a shift toward heterogeneous system-on-chip (SoC) architectures, including RISC-V cores and hardware accelerators, to enable more powerful on-board data processing in the medium term. This reference serves as a conceptual anchor for understanding the hardware-driven limits that influence where processing should occur.

The work conducted in [17] takes a more algorithmic approach, proposing a hybrid architecture where hyperspectral image compression is executed on-board the satellite, while a deep CNN is used on the ground to reconstruct the full-resolution image. Their work demonstrates that by applying learned priors, it is possible to compress data aggressively in orbit - achieving bandwidth efficiency - without significantly degrading image quality. The paper evaluates their system on real hyperspectral scenes and compares the reconstruction performance with traditional compression schemes. The key contribution lies in optimizing the processing distribution across the data pipeline, showing that a computationally light encoder can run on-board while relying on ground-based AI for sophisticated inference. This dual-architecture approach reinforces the idea that processing architecture should not be binary but adaptive to mission needs.

A detailed survey by [18] expands on ESA's strategic shift toward edge intelligence in space. The authors identify the latency and bandwidth bottlenecks associated with traditional EO systems and argue that even preliminary data screening and prioritization conducted in orbit can dramatically reduce processing and delivery times. They highlight that while data rates from modern sensors are growing by factors of 100× per generation, downlink capabilities only increase by 3–5×. Furano et al. also introduce the concept of "usefulness-aware" data filtering-where ML models perform basic object detection or change detection on-board to determine which data is worth transmitting. The study lays out technical prerequisites for this transition, including low-power hardware, fault-tolerant software and AI algorithm miniaturization. It is an important piece that frames the operational incentives for shifting parts of the EO workflow into space.

The authors of [19] focused on real-world processing trade-offs in high-resolution optical missions such as the Pléiades series. They analyze an on-board/ground split processing chain where satellites transmit compressed panchromatic and multispectral imagery and the high-quality color images are reconstructed on the ground through pan-sharpening techniques. The study compares the signal degradation introduced by in-orbit compression against the gains made through complex restoration and fusion techniques on the ground. The conclusion is that processing design must be co-optimized: on-board tasks like compression must be tuned to ensure compatibility with ground-based algorithms. This is a strong empirical example showing that EO processing decisions are not isolated but interdependent across the system architecture.

The authors of [20] present an end-to-end architecture designed for time-sensitive applications such as rapid disaster monitoring. Developed under the EU EO-ALERT project [51]), their architecture utilizes high-performance COTS hardware to perform on-board compression, encryption and prioritization, achieving sub-second latency for urgent EO products. The authors provide benchmarks comparing their latency-oriented pipeline with conventional ground-based systems, which can take minutes to process and deliver actionable intelligence. Their results show how architectural decisions (e.g. as in-memory pipelining and task scheduling) can be tuned to minimize time-to-information, an increasingly critical metric in emergency response scenarios. This reference demonstrates how mission type (e.g., real-time vs. scientific observation) dictates the processing architecture.

Continuing the deep review, the Φ-Sat-1 mission, detailed in the work of [21], which is one of the first operational examples of on-board artificial intelligence applied in orbit. The authors describe how a convolutional CNN was deployed on the Myriad 2 Vision Processing Unit (VPU) aboard the satellite to detect and filter cloud-covered images in real-time. This AI-based filter prevented the transmission of roughly 30% of images deemed unusable, freeing up bandwidth and reducing the downstream burden on ground-based systems. Importantly, the paper evaluates the reliability of the model's inference and confirms no significant degradation in the quality of retained data. This mission proves that on-board AI can autonomously increase data value while minimizing transmission and post-processing time, which is critical in bandwidth-constrained EO architectures.

The authors of [22] present a broader technical review of AI-powered on-board image recognition systems tailored for small EO satellites. Their study identifies key challenges related to deploying machine learning models in low-power, radiation-hardened environments. The authors evaluate different neural network architectures and edge computing chips, arguing that the limited energy and computational capabilities of smallsats necessitate highly optimized model deployment. They also outline latency trade-offs between processing data in orbit and sending it for full analysis to the ground. The authors conclude that although on-board inference offers significant advantages for real-time applications, its benefits are best realized in

constellation-based systems where multiple nodes can coordinate data acquisition and prioritize downlink content collectively.

In [23], the authors turn their attention to Synthetic Aperture Radar (SAR) systems and the specific challenges of SAR data processing. Historically, SAR data have been downlinked as raw signal bursts and processed entirely on Earth, requiring high bandwidth and introducing significant delays before usable imagery is available. The authors explore emerging techniques to shift pre-processing stages-such as data compression, initial focusing and quality checks-into orbit. They show how doing so can reduce data volume by up to 70% without compromising data quality, enabling faster delivery of analysis-ready SAR products. The paper also highlights that SAR image formation is computationally intensive and requires optimized on-board hardware (*e.g., Field-Programmable Gate Arrays (FPGAs) or neuromorphic processors*) to be viable. Their work reinforces the domain-specific nature of EO processing decisions: in high-volume sensing like SAR, even partial on-board workflows can yield major efficiency gains.

The research performed in [24] offers a state-of-the-art survey on the integration of AI in EO satellites, covering both academic advances and commercial implementations. The authors categorize dozens of missions based on their processing approach-ranging from legacy satellites that downlink all raw data, to modern AI-driven platforms capable of event detection, image enhancement and prioritization. They also assess the landscape of space-qualified hardware accelerators and software frameworks optimized for in-orbit inference. One of the key conclusions is that the traditional architecture of centralized ground-based processing is increasingly unsustainable given projected data growth. Instead, they advocate for distributed EO processing architectures, where task delegation is dynamic based on satellite capabilities, mission priority and bandwidth availability. Their survey is an essential resource that demonstrates how architecture and AI strategy are becoming coupled in EO system design.

To conclude this section and move on to the next sub-sections of Chapter 3, from our deep analysis and research, we clearly see that the literature reflects a clear evolution in EO system design - away from the traditional model of centralized ground-based processing toward more dynamic and distributed processing architectures. The decision to process data on-board or on the ground is increasingly context-specific, shaped by factors such as mission latency requirements, sensor data volume, platform constraints and available ground infrastructure. **What emerges from recent studies is a growing consensus that hybrid models, which strategically split the EO workflow between satellite and ground systems, offer the best performance-to-resource trade-off**. These architectures capitalize on the immediacy of in-orbit filtering and compression, while still taking advantage of the flexibility and computational power of terrestrial processing chains. The next section will illustrate how these principles have been implemented in real EO missions, providing concrete examples of architecture decisions in practice.

3.2.2 Summary of Key Missions using On-board/Edge Processing

Real-world EO missions provide valuable insights into how on-board and ground-based processing strategies are applied to meet operational needs. One of the most prominent examples is the Φ-Sat-1 mission, launched by the European Space Agency (ESA) in 2020. As the first satellite to include an on-board AI processor (Intel's Myriad 2 Vision Processing Unit), Φ-Sat-1 employed a CNN to filter cloud-covered optical images in real time. This strategy reduced the volume of non-useful imagery transmitted to ground stations by up to 30%, conserving bandwidth and minimizing post-processing workload. According to [21], the on-board AI performed reliably throughout the mission, validating the feasibility of deploying deep learning models on radiation-tolerant hardware in orbit. **Φ-Sat-1 is widely regarded as a milestone in space-based edge computing, demonstrating both the technical viability and operational benefit of AI-enabled on-board processing**.

The OPS-SAT mission, launched by ESA in 2019, was designed as a testbed for advanced space software, including on-board data analytics and AI. Although not a science mission per se, OPS-SAT's 1U CubeSat platform included an experimental high-performance computer that allowed in-orbit execution of algorithms for image processing, compression and autonomy [18]. ESA reports that partners used the system to validate real-time image enhancement and autonomous fault recovery, marking it as a milestone in software-defined space systems. OPS-SAT illustrates the potential of adaptive, reprogrammable satellite platforms where processing strategies can evolve during the mission lifespan.

The HYPSO-1 satellite, launched in 2022 by the Norwegian University of Science and Technology (NTNU), focuses on hyperspectral imaging for marine and coastal applications. The mission implements on-board cloud screening and data prioritization using an FPGA-based architecture, reducing the volume of data before transmission. As discussed by the authors of [22], such designs are particularly effective for small satellites with limited bandwidth and contact windows. The selection of only high-value hyperspectral frames for downlink allowed the HYPSO-1 to improve the effective use of on-board resources while enabling faster delivery of useful data, a trade-off that is especially critical in latency-sensitive ocean monitoring missions.

Germany's EnMAP mission offers a hybrid processing chain in which both on-board and ground-based segments play defined roles. The work conducted in [23] explains that EnMAP conducts initial radiometric correction and formatting on-board, enabling rapid ingestion of Level-1 data by the ground segment. More computationally intensive operations, such as atmospheric correction and biophysical modeling, are then executed on Earth using standardized Copernicus-compliant chains. This architecture balances speed and flexibility, demonstrating how modular processing splits can reduce delivery latency without compromising scientific integrity.

In contrast, the Sentinel-2 mission (*part of the EU Copernicus programme*) represents a traditional ground-centric processing model. The authors of [24] explain that Sentinel-2 collects high-resolution multispectral imagery and transmits it without on-board filtering or compression, prioritizing data completeness over bandwidth efficiency. While this supports long-term archiving and high-fidelity reprocessing, it introduces significant delays, sometimes spanning hours or days, before data products reach users. The Sentinel architecture is ideal for scientific consistency and retrospective analysis, but less suited to missions requiring immediate situational awareness.

3.2.3 Industry Trends and Technology Roadmaps

The work in [18] in an influential ESA-backed publication, argue that AI deployed "at the edge"-meaning on-board spacecraft-is crucial for managing the growing data bottlenecks in EO. They note that while sensor data rates are increasing by a factor of 100 with each generation, downlink bandwidth is only improving at 3–5×, leading to an unsustainable disparity. Their work outlines the technical and operational challenges of running AI models in orbit-such as power constraints, thermal limits and radiation hardening-but emphasizes that real-time in-space inference (like cloud filtering) could significantly reduce transmission costs and improve responsiveness. The authors of [24] survey recent advancements in satellite-based AI, focusing on technical feasibility and deployment architectures. They argue that sending raw data to Earth is no longer viable and that inference workloads-such as object detection or anomaly classification-must shift on-board. Their literature review synthesizes best practices in model compression, use of VPUs and FPGAs and radiation fault tolerance. They conclude that on-board AI, while still maturing, is essential for managing latency and resource constraints in real-time EO missions. The authors of [25] provide a sweeping overview of how AI is transforming EO workflows across Europe. The authors advocate for the integration of AI across the entire EO chain-from image acquisition to semantic understanding. They highlight key trends such as self-supervised learning, explainable AI and federated models that could better suit edge environments. Acknowledging ESA's leadership, the paper identifies edge-AI as a core enabler for shifting EO systems from image providers to insight services.

The European Commission's Joint Research Centre [26] has formally recognized on-board processing as a strategic priority in the latest Copernicus Strategic Research & Innovation Agenda. Annex III identifies AI, edge computing and latency-sensitive processing as high-priority topics for future investment. It calls for innovations that will allow satellites to filter, analyze, or fuse data on orbit-enhancing responsiveness in areas like disaster response and environmental monitoring. This institutional framing aligns with Horizon Europe goals for strategic autonomy and digital sovereignty. ESA's Φ-lab has become a central hub for nurturing AI-in-space innovations. Through its InCubed programme [27], the Φ-lab funds mission concepts and technologies that embed AI closer to the sensor. This report highlights its EO MakerSpace initiative, which supports the

development of reprogrammable, AI-enabled payloads. **Φ-lab sees edge intelligence not just as a performance boost, but as a transformative force enabling autonomous, taskable satellites**. Their support for software-defined payloads, in particular, aims to foster a new generation of "thinking sensors." ESA's Φ-sat-2 mission, launched in 2024, builds directly on Φ-sat-1 and pushes the boundaries of on-board AI [28]. According to this source, Φ-sat-2 is equipped with a multi-application AI computer capable of real-time fire detection, vessel monitoring, cloud masking and image compression-all running on-board a CubeSat. It also supports post-launch reprogramming, demonstrating the shift toward software-defined constellations. This mission exemplifies the Φ-lab vision of agile, smart EO infrastructure.

In a 2023 trend analysis, ESA highlighted edge computing as one of the "six trends to watch" in commercial Earth observation [29]. The report outlines how commercial operators are beginning to invest in on-board intelligence to reduce costs, lower latency and differentiate their offerings. ESA notes the rise of data-as-a-service models enabled by smarter satellites, especially in sectors like maritime surveillance and agriculture, where time-to-insight is critical. The World Economic Forum [30], in collaboration with Deloitte, emphasizes the global explosion in EO data - hundreds of terabytes per day - and points to AI as a critical enabler of scalability. Their white paper argues that Europe must adopt AI-infused data workflows to manage this information tsunami. It identifies AI on the edge, supported by secure on-board computing, as a linchpin for building more resilient and responsive EO systems - especially for sustainability and climate missions.

Unibap's SpaceCloud® platform [31] offers a commercial implementation of ESA's edge computing vision. Described in ESA's InCubed overview, SpaceCloud transforms satellites into mini data centers that can host containerized apps, run AI models and adapt missions in orbit. Unibap's iX5 and iX10 hardware units have flown on test missions, demonstrating that Earth imaging satellites can analyze scenes in real time, apply AI filters and transmit only relevant results. This marks a shift from data collection to insight delivery, all from orbit. KP Labs, a Polish space tech company, has achieved Europe's first successful demonstration of AI-processed hyperspectral imagery on-board a CubeSat [32]. Intuition-1, launched in late 2023, uses a custom AI accelerator called Leopard to run deep learning models on hyperspectral scenes in orbit. The result: real-time feature classification and data compression on-board, reducing transmission costs and latency while maintaining scientific utility. This aligns perfectly with Horizon Europe's goals for smart, compact and autonomous EO systems. LuxSpace and Klepsydra presented an early but robust approach to modular on-board processing [33]. In their 2019 proof-of-concept, they showed that by networking specialized computing nodes on-board a satellite, one could achieve scalable data handling without overloading a single processor. Their CANopen-based architecture supports real-time data compression and distribution, making it a practical design for smallsat constellations requiring reconfigurability and resilience. Planetek Italia's AI^{ix} initiative envisions satellites as real-time service providers [34]. Their upcoming AIX-1p demonstrator

supports uplinked AI tasks and user-defined analytics, enabling “cognitive cloud” processing directly in orbit. Planetek’s architecture allows for customized workflows per customer, dynamically changing as mission needs evolve. If successful, this could lead to a flexible EO infrastructure where satellites can process and deliver custom insights on demand.

The TerraWatch Space’s EO outlook forecasts that 2025 will be a turning point for operational adoption of edge-AI in satellites [35]. They cite the success of test missions like Φ-sat-1 and Intuition-1 as signs that the technology has matured past pilot phase. They predict that major constellations will begin integrating on-board analytics as standard, shifting the paradigm of EO from data acquisition to insight generation. Finnish satellite maker reOrbit emphasizes that speed is now a differentiator in EO [36]. Their 2024 technical brief outlines how inter-satellite links, AI-based filtering and on-board analytics can shorten the gap between data collection and user action. They note that national and commercial constellations alike are adopting on-board processing to meet customer demands for immediacy. This shift is not only technical but business-driven-faster insights equal competitive advantage. Finally, Unibap’s 2025 white paper argues that edge computing enables “value creation in orbit” [37]. Instead of sending data home for interpretation, a smart satellite can run multiple service apps in parallel-delivering wildfire alerts, vessel tracking and agriculture maps in real time. This redefinition of EO missions-turning satellites into multi-service platforms-is what Unibap sees as the path forward for European satellite operators seeking efficiency, responsiveness and scalability.

To conclude, the current landscape of EO data processing in Europe reflects a decisive shift toward intelligent, decentralized and adaptive architectures. Across institutional roadmaps, industry platforms and mission demonstrators, the emphasis is increasingly placed on integrating AI and edge computing into the EO pipeline-not as experimental add-ons, but as foundational elements for next-generation satellite operations. ESA’s Φ-lab, Copernicus strategic agendas and commercial actors like Unibap, KP Labs and Planetek Italia collectively point to a future where satellites are not just data collectors but autonomous agents capable of filtering, analyzing and responding to environmental phenomena in real time. These developments represent a critical rethinking of how EO missions are designed-from the software-defined nature of payloads to the dynamic reconfiguration of analytic tasks in orbit. Moreover, this transition supports broader European objectives for digital sovereignty, sustainability and strategic autonomy in space. By investing in smart processing technologies, Europe is not only reducing dependence on ground infrastructure and external data pipelines, but also positioning itself at the forefront of responsive, high-value EO service delivery. The convergence of modular hardware platforms, AI-driven analytics and flexible software ecosystems is paving the way for EO constellations that are mission-adaptive, resource-efficient and deeply integrated into decision-making workflows across sectors like agriculture, maritime security, climate monitoring and emergency response. As these technologies mature and the ecosystem continues to align around common standards

and architectures, the role of satellites will evolve from passive sensors into proactive, real-time decision-support assets. In this context, Europe is uniquely placed to lead the way in shaping the future of Earth observation as a seamless, edge-enabled digital infrastructure serving both public interest and commercial innovation.

3.2.4 Research Gaps and Opportunities

Despite rapid advancements in on-board data processing for EO, several technical and strategic gaps continue to challenge the operational integration of such technologies. A critical area where research remains limited is the systematic benchmarking of latency across EO mission architectures. While studies like those by [20] and [23] demonstrate that on-board processing can drastically reduce time-to-information, there is no widely adopted methodology for measuring latency gains under standardized mission profiles. Moreover, latency is rarely isolated from other performance indicators such as energy consumption, data quality preservation and prioritization accuracy. As [24] points out, these variables are often studied independently, preventing a comprehensive understanding of the true value proposition of low-latency, in-orbit analytics.

In parallel, the deployment of real-time analytics on-board satellites remains in its infancy, especially in missions beyond demonstrators like Φ-Sat-1 and Intuition-1. While ESA's Φ-sat-2 mission aims to expand this frontier with multiple concurrent AI apps in orbit (ESA, 2024b), existing models still rely heavily on pre-launch training and static decision thresholds. The lack of adaptive or self-updating AI architectures in space represents a major bottleneck. Current research does not yet offer robust pathways for continual learning or federated learning schemes that could allow satellites to evolve and refine their analytic capabilities based on changing operational conditions or user needs [25]. This creates an opportunity for further research into trustworthy, updatable AI frameworks that comply with the computational and security constraints of in-orbit systems.

Another research gap lies in data prioritization methodologies. While missions like HYPSO-1 and Intuition-1 have successfully implemented basic filtering and cloud screening techniques as explained in [22] and [32], the criteria used to define "high-value" data are often heuristic or application-specific. [18] and [21] mention the promise of "usefulness-aware" filtering but stop short of offering a unified metric for data relevance in operational terms. This opens a significant opportunity to develop standardized information value functions or multi-objective ranking schemes that can dynamically adjust prioritization strategies based on mission context, user profiles, or downstream processing capabilities.

On the policy and regulatory front, standardization and interoperability remain underexplored, especially for hybrid processing models. While platforms explained in the works of [34] and [38] propose flexible, software-defined payloads, there is no clear framework for certifying,

validating, or interconnecting these systems across agencies and commercial operators. This lack of standardization not only raises concerns about data trustworthiness and reproducibility but also complicates integration with EU-wide services like Copernicus. Additionally, questions around data sovereignty, especially for on-board pre-processed or AI-augmented imagery, require further clarification in the context of the EU's digital strategy and the General Data Protection Regulation (GDPR). As EO missions evolve toward more autonomous, in-orbit processing pipelines, new regulatory and ethical challenges are emerging. Three key concerns stand out:

- Data Sovereignty: As AI-enabled satellites perform processing and filtering in orbit, questions arise regarding ownership and jurisdiction of the processed data. For instance, if imagery is analyzed and compressed before reaching the ground, who determines its classification, priority, or deletion. This has implications for national security, international data exchange and public accountability.
- Transparency of On-board Decisions: On-board AI models (e.g., for cloud masking or anomaly detection) are often black-box in nature. Without robust explainability protocols, users may not be able to trace how or why certain observations were prioritized or discarded. This undermines trust, especially in missions involving defense, emergency response, or environmental monitoring.
- GDPR and Data Privacy Compliance: Although EO data is typically considered non-personal, location-based analytics or imagery of private assets (e.g., vehicles, homes, critical infrastructure) can intersect with personal data regulation. If a satellite autonomously flags, filters, or transmits such data without human validation, this could trigger GDPR compliance requirements-especially under Articles 22 (*automated decision-making*) and Article 5 (*data minimization and accountability*) [45].

Meanwhile, as hybrid architectures gain traction, these concerns underline the need for transparent processing chains, verifiable on-board logic and harmonized regulatory standards. Future missions should consider embedding compliance-by-design principles in both on-board firmware and ground-level data governance protocols. Finally, a persistent and underdeveloped domain is the economic modeling of on-board versus ground-based architectures. While qualitative benefits of on-board analytics (*e.g., faster insights, reduced downlink costs*) are well-acknowledged, there is a lack of quantitative frameworks for assessing cost-effectiveness across the mission lifecycle. [33] and [37] touch on modular and upgradable architectures that could improve long-term return on investment, but they do not model total cost of ownership or opportunity costs in multi-mission scenarios. This leaves a substantial gap for studies that combine technical performance with economic trade-off analyses, particularly under constraints such as launch mass, energy budgets and software development cycles.

To conclude, while the EO sector is undoubtedly undergoing a transformative shift fueled by advances in on-board processing and intelligent edge architectures, this evolution remains constrained by several critical research and implementation gaps. **The absence of standardized performance evaluation frameworks-particularly in measuring latency, energy efficiency and data value - makes it difficult to compare different mission architectures and to quantify the real-world advantages of in-orbit analytics**. Without such benchmarks, both mission planners and funding bodies lack the evidence needed to make data-driven decisions on whether, when and how to adopt on-board processing solutions. In parallel, scalable and adaptive AI systems in space remain largely experimental. Most current approaches rely on static, pre-trained models that cannot adapt to unforeseen environmental conditions or evolving mission priorities. The lack of on-board learning capabilities, versioning protocols and trust-assurance mechanisms limits the autonomy and resilience of future EO missions. Furthermore, interoperability and standardization - especially in relation to Copernicus services and the broader European EO ecosystem - are still at a nascent stage. As more satellites become software-defined and AI-enabled, the absence of shared data formats, validation schemes and integration protocols threatens to fragment rather than unify EO capabilities. Equally pressing is the lack of holistic economic modeling to assess total mission cost-benefit when shifting from ground-centric to hybrid or fully on-board architectures. Although qualitative gains like reduced latency and bandwidth savings are clear, decision-makers still lack quantitative tools to model return on investment, lifecycle costs and opportunity trade-offs-particularly for small satellite operators or new entrants to the EO market.

3.3 Comparative Evaluation: Technical, Operational and Economic Analysis

3.3.1 Technical Comparison

The technical dimension is perhaps the most fundamental when assessing the viability of on-board versus ground-based data processing in EO missions. Differences in processing power, energy consumption, data storage capacity and real-time computational capabilities determine what kinds of analytics can be performed where and how reliably. On-board processing is inherently constrained by the physical and environmental limitations of spacecraft systems-such as size, mass, power availability and radiation exposure-while ground-based systems benefit from virtually unlimited computational resources and infrastructure. Nevertheless, advancements in low-power AI chips, space-hardened FPGAs and software-defined payloads are closing this gap, enabling satellites to perform increasingly complex tasks in orbit. This section draws on concrete mission data and technology evaluations to explore how these technical differences manifest in practice and what trade-offs they imply for mission architects.

The work in [18] highlights the fundamental technical disparities between on-board and ground-based processing for EO. They note that modern EO satellites generate massive data volumes that far outpace available downlink bandwidth, since "the growth in the volume of data being produced by sensors is not being matched by a corresponding increase in data download capacity". On-board computing must contend with strict power and processor limitations (*often using older radiation-hardened hardware*) and limited storage in contrast to virtually unlimited resources on the ground. These constraints make it challenging to perform heavy analyses in orbit. However, [18] points out that shifting certain processing on-board can yield significant benefits: it reduces the data volume that must be transmitted (*alleviating bandwidth pressure*), enables faster responsiveness (*since less data needs to be downlinked before acting*) and provides flexibility to filter or select mission data in real time. In essence, while satellites cannot rival ground stations in raw computing power or energy supply, even limited on-board processing can pre-screen and compress data, mitigating the downlink bottleneck and speeding up the delivery of critical information [18].

In another work, [39] illustrates these technical trade-offs through the example of ESA's Φ-Sat-1 experiment. Φ-Sat-1 carried an Intel Movidius Myriad 2 vision processing unit – a low-power AI accelerator – to perform cloud detection on-board. The motivation was that each generation of EO sensors can produce 100× more data, while downlink capacity grows only 3–5×, creating an unsustainable gap. By processing imagery in space, Φ-Sat-1's on-board AI could discard useless images (*e.g. fully cloudy scenes*) before transmission. This resulted in about a 30% bandwidth savings, as those cloudy images never had to be sent to ground. In other words, the satellite effectively "freed" 30% of its downlink capacity by doing simple filtering on-board, a task that would normally be done on the ground after consuming precious downlink resources. Technically, this demonstrated how even limited AI inference on-board (*constrained by a <2W low-power chip*) can dramatically cut downlink requirements and optimize data flow. The trade-off is that the satellite's computer must be robust enough to handle the processing task within its power/energy budget, but the Φ-Sat-1 success showed that modern COTS AI chips can meet this challenge. This example underscores the technical advantage of on-board processing: by intelligently reducing data at the source, satellites mitigate their own downlink and storage limitations, something ground-only processing cannot achieve due to its need to receive all raw data first.

3.3.2 Operational Considerations

Beyond technical specifications, the choice between on-board and ground-based processing carries significant operational implications. These include mission responsiveness, latency in data delivery, control autonomy and the complexity of ground segment interactions. While ground-based processing centralizes control and simplifies satellite design, it introduces delays in mission workflows due to dependency on downlink availability and ground station schedules. Conversely,

on-board processing can support faster decision-making, mission autonomy and dynamic tasking-but at the cost of increased system complexity and the need for highly reliable in-orbit computing. This section examines how operational models evolve when processing capabilities are moved closer to the sensor, highlighting examples from recent European missions and experiments that underscore the real-world benefits and challenges of shifting intelligence into orbit.

The authors of [40] analyze different mission operation approaches and their impact on data latency – the delay from data acquisition on orbit to delivery of information to the end user. They compare traditional ground-downlink-centric operations with newer concepts like inter-satellite relay networks and edge (on-board) computing. Their findings indicate that on-board processing alone does not necessarily improve latency, because even if a satellite processes data immediately, it still must wait for a downlink opportunity to deliver results. In other words, a processed image cannot reach users until a ground station pass occurs (unless other communication changes are made). By contrast, adding real-time relay links (*e.g. optical inter-satellite links to a geo-stationary relay or a network of ground stations*) can cut latency significantly – their study showed a ~2× latency reduction with relay networks in a Low Earth Orbit (LEO) constellation, versus using only ground stations. Interestingly, they note that achieving equivalent low latency with just ground stations would require many more ground sites and around 5× higher operational costs per year. Thus, from an operational standpoint, on-board processing is valuable primarily for reducing data volume, but not sufficient on its own to solve latency issues. **Timely data delivery still depends on communication architecture**. The implication is that mission designers must combine on-board data processing with improved downlink strategies (more ground stations or inter-satellite links) to meet real-time or low-latency requirements. On-board "edge" computing helps by prepping and compressing data so that when a link is available the transfer is faster, but it does not eliminate the fundamental latency of orbital communications [40].

Unibap AB's in-orbit demonstration of a "SpaceCloud" computer in 2021 showcased how on-board processing can transform operations by enabling a form of cloud computing in space [41]. In this experiment (*hosted on D-Orbit's ION platform*), a high-performance computer node was installed on the satellite to run AI algorithms and even allow uploading new software after launch. The project's goal was to achieve "extremely low-latency decision support" by processing imagery on-board and quickly delivering insights to users. By connecting the satellite's computer securely to a cloud environment on the ground, the mission could downlink processed information (*alerts, analytics*) immediately when available, rather than raw data later. This approach reduces the need for complex real-time decision-making on the ground because the satellite can autonomously analyze events (*e.g. detect a wildfire or ship*) and send an alert as soon as it has a communications link. Operationally, this points to greater mission autonomy and

responsiveness – satellites could potentially task themselves or prioritize certain data without waiting for ground commands. Unibap's demonstration also emphasized flexibility: the SpaceCloud platform was capable of receiving "additional applications and experiments" uploaded during the mission to enhance or change its functionality. Such in-orbit reconfigurability is a major operational shift; traditionally, processing algorithms are fixed at launch and any change requires ground reprocessing, but here the satellite's "brain" can be updated like a remote server [41]. This flexibility improves control – operators can tweak on-board apps – but also demands robust design to ensure new software doesn't jeopardize the mission. Overall, the Unibap case suggests that on-board processing can simplify operations by handling data on the satellite (*reducing reliance on downlink scheduling and ground centers*) and add operational agility through software updates, albeit with added complexity in satellite software management.

Another operational advantage of on-board processing is the faster access to actionable information, as demonstrated by the recent Intuition-1 satellite as cited in [42]. Intuition-1, a 6U CubeSat launched in 2023 with a hyperspectral sensor, includes a powerful Leopard data processing unit running AI algorithms. KP Labs reported that this on-board AI system can detect, classify and analyze surface phenomena in orbit, which "not only optimizes data relevance and precision but also minimizes the volume of data transmitted to Earth, speeding up the availability of vital information". In practice, this means Intuition-1 can, for example, spot a specific event or feature (*such as an algae bloom or deforestation indicator*) and send down just the relevant result or a small subset of processed images, rather than waiting to downlink heaps of raw hyperspectral data for ground processing. The latency from observation to insight is reduced because initial analysis happens immediately during the satellite pass. KP Labs describes this as a "new benchmark for efficiency and functionality in satellite operations," indicating that mission timelines can be compressed – analysts get results faster and may even receive alerts during the same orbit pass when data was collected (Zachara, as cited in [42]). Importantly, this operational model also eases the burden on ground control: instead of continuous manual intervention to download and process everything, the ground segment can focus on handling high-priority insights streamed from the satellite. However, it requires high confidence in the on-board algorithms' accuracy and the ability to verify and trust those autonomous decisions. The Intuition-1 case nonetheless exemplifies how on-board processing enables a more efficient operational workflow, with satellites becoming smarter actors that deliver targeted information in near real-time.

3.3.3 Economic Evaluations (Cost-Benefit Analysis)

This cost-benefit analysis explained below addresses T5.4' objective of comparing the economic implications of on-board versus ground-based processing, demonstrating how capital

investments, operational expenditures, and mission scalability trade-offs should inform architecture selection.

Economic factors are central to any EO system design decision, particularly when evaluating cost trade-offs between on-board and ground-based processing architectures. These trade-offs extend beyond initial development expenses to encompass long-term operational costs, infrastructure demands and mission scalability. **On-board processing can reduce recurring costs by minimizing data volume, thus saving bandwidth and decreasing dependency on extensive ground infrastructure**. However, this must be weighed against the capital costs of advanced processors, integration complexity and the potential need for redundancy and reprogrammability in space. Recent economic models and case studies (explained in upcoming paragraphs) demonstrate that while on-board processing can yield substantial savings over the mission lifecycle, the optimal balance depends on mission goals, data types and urgency requirements. This section reviews such findings to provide a nuanced understanding of the economic implications shaping EO mission architecture decisions.

On-board vs. ground processing involves significant cost trade-offs in both data handling and infrastructure. One clear economic benefit of on-board data reduction is the savings in communication costs. For example, Ubotica's Φ-sat-1 AI experiment (see [43]) showed that filtering data in space can "significantly reduce aggregated downlink costs" by transmitting only useful images. In traditional missions, downlinking large volumes of raw data requires extensive ground station contact time (or costly relay satellite usage) and big data pipelines on Earth – all of which contribute to operational expenses. In fact, performing cloud detection on-board, Φ-sat-1 avoided sending roughly 30% of images that were cloud-obstructed, directly translating to 30% less downlink bandwidth used (*see [39] for more info on this*). Over the mission life, this cuts the amount of ground station time needed (which can be rented or allocated at high cost) and reduces the required size of ground data centers. Thus, even though adding an AI chip on the satellite incurred additional upfront cost, it yielded recurring cost savings by lowering daily communication needs. This is especially relevant for constellations or high-data missions, where communication bills (*or service fees for commercial ground station networks*) accumulate rapidly. Ubotica and ESA essentially demonstrated a cost-efficiency case: investing in smarter satellites can reduce the data pipeline infrastructure and bandwidth expenditures on the ground. Conversely, a purely ground-based approach might save on satellite manufacturing costs (*simpler payload electronics*) but would likely spend more on downlink bandwidth and ground processing infrastructure throughout the mission.

The work in [44] provides a quantitative perspective on these economic trade-offs by modeling optimal resource allocation in a satellite network for EO data processing. They consider a scenario with a LEO constellation and formulate how to best allocate computing tasks (on-board vs on-ground) and downlink bandwidth. A key insight from their study is that allowing satellites to

preprocess data (*an "orbital edge computing" paradigm*) can avoid wasting "precious bandwidth" on downlinking irrelevant data. In their optimized solution, each satellite processes as much as it efficiently can in orbit – for instance, compressing imagery or extracting features – so that only high-value results consume downlink capacity. This approach minimizes the total bandwidth required and can reduce the need for large ground station networks. In economic terms, it improves throughput per cost: more useful information is delivered for the same communication expense. Their work in [44] implicitly shows scalability benefits as well – as satellite fleets grow, a ground-centric approach would require proportionally more ground stations and backhaul capacity (*significant capital and operating costs*), whereas an on-board processing approach scales more gracefully by handling a portion of the load on the satellites themselves. Furthermore, on-board processing contributes to lower OPEX by decreasing routine data transmission volumes, which is in line with other analyses that found ground networks without on-board intelligence must spend much more to achieve similar performance. However, Valente et al. also acknowledge the upfront investment in more capable on-board computers and advanced software. This CAPEX must be weighed against the operational savings. The economic optimum will differ by mission: for some, the cost of flying high-performance processors (*and potential shorter satellite lifespans or higher risk*) is justified by big savings in downlink costs; for others, cheap satellites with robust ground analytics might still be more economical if latency is not critical. **Recent studies converge on the idea that on-board processing, when used judiciously, can reduce total mission costs by trimming the expensive ground segment and bandwidth needs, making the entire EO data chain more cost-effective.** Ultimately, from: a) technical, b) operation and c) economical points of view, we can say the following:

- **From a technical standpoint, on-board processing is constrained by limited processing power, energy availability and radiation-hardened hardware**. While current space-grade AI chips and Field-Programmable Gate Arrays (FPGAs) have enabled basic inference tasks-such as cloud masking or anomaly detection-they still lag behind terrestrial systems in terms of versatility and raw computational throughput. Ground-based processing, on the other hand, benefits from high-performance servers, cloud infrastructure and upgradable software environments that can support complex and evolving algorithms. However, the gap is narrowing with the success of missions like Φ-Sat-1, HYPSO-1 and Intuition-1, which demonstrate that even compact, low-power satellites can perform meaningful processing in orbit.
- **Operationally, the most significant benefit of on-board processing lies in its potential to reduce data latency and enhance mission autonomy**. Satellites equipped with edge computing capabilities can detect and prioritize events in real time, reducing the volume of data that needs to be transmitted and accelerating the delivery of critical insights. This is particularly valuable in time-sensitive applications such as disaster response, maritime

surveillance, or tactical defense scenarios. Nevertheless, real-time responsiveness still depends heavily on downlink opportunities and communications architecture - on-board processing is an enabler, but not a complete solution to latency without complementary improvements in inter-satellite or ground-based communication networks.

- **Economically, on-board processing offers recurring operational savings through reduced bandwidth consumption and more efficient ground infrastructure usage**. However, it requires higher upfront investments in computing hardware, algorithm optimization and system validation. The break-even point for these investments varies by mission profile. For data-intensive missions with frequent revisits or large constellations, the savings in downlink costs and ground station time can justify the added complexity and hardware costs. Conversely, missions with lower data rates or long-term scientific objectives may still find ground-centric processing more cost-effective.

3.4 Cost Benefit Summary

As thoroughly explored in **Section 3.3.3**, the economic analysis reveals that **the most significant cost differentiator between on-board and ground-based processing lies in the trade-off between up-front capital investment and long-term operational expenditure**. On-board processing demands higher initial costs due to the inclusion of advanced space-qualified processors, algorithm tuning and mission validation. However, these investments are often offset over time by lower recurring costs - most notably through bandwidth savings, reduced reliance on costly ground station networks and minimized data infrastructure needs.

The comparative evaluation is presented in the cost-benefit summary table below. **On-board processing is most economically advantageous in responsive missions where communication opportunities are limited or ground infrastructure costs are high. Ground-based processing, while offering flexibility and lower up-front hardware costs, entails ongoing operational expenses due to full data transmission and infrastructure maintenance. Hybrid models strike a balance - enabling data prioritization and compression on-board while retaining analytical depth and storage on the ground**. Overall, the findings from Section 3.3.3 support a shift toward adaptive architectures where cost is optimized not in absolute terms, but relative to mission urgency, data volume and infrastructure availability.

Table 3. Cost-benefit analysis summary

Architecture	Upfront Cost	Recurring Cost	Scalability	Bandwidth Savings	Mission Type Suitability
On-board	High	Low	Moderate	High	Tactical, Time-Critical

Ground-based	Moderate	High	High	Low	Long-term, Scientific
Hybrid	Moderate	Moderate	High	Moderate / High	Mixed-use, Responsive

3.5 Summary of findings

The comparative analysis conducted in **Chapter 3** sheds light on the evolving landscape of EO mission architecture. **Table 4** below merges these insights into a comparative matrix across key mission architecture dimensions:

Table 4. On-board vs Ground-based EO Data Processing Comparative Matrix

Dimension	On-board processing	Ground-based processing
Processing Power	Limited (radiation-hardened, low-power chips)	High-performance, scalable and adaptable computing infrastructure
Energy Efficiency	Power-constrained	Virtually unlimited
Storage	Limited capacity; temporary buffer for prioritized data	High-capacity storage for complete raw datasets
Latency	Low if combined with proper communication relay architecture	Dependent on satellite pass schedule and ground station availability
Responsiveness	Real-time analytics possible (e.g. anomaly alerts, cloud masking)	Delayed insights; requires post-pass analysis
Mission Autonomy	Supports in-orbit decision-making and tasking	Requires ground-based intervention for all mission logic
Communication Load	Significantly reduced by pre-filtering and data compression	High; full datasets must be downlinked regardless of relevance
Infrastructure Cost	Reduces recurring comms costs and data center load	Higher OPEX due to bandwidth and processing pipeline
Scalability	Software-defined payloads enable flexibility; hardware is constrained	Fully scalable with cloud and ground systems
Upgradability	Increasingly possible via reprogrammable payloads	Easy to update, maintain and optimize

This comparison clearly indicates that there is no one-size-fits-all model. The optimal processing pathway must align with mission objectives, technical readiness and economic constraints. For instance, short-lived tactical missions or rapid response tasks (*e.g. fire detection, border surveillance*) benefit from lightweight on-board processing and faster decision loops, whereas high-fidelity environmental monitoring or climate change datasets may still prioritize post-mission ground-based analytics for consistency and reusability. **To analyze further, the comparative matrix underscores a fundamental trade-off between processing proximity and processing depth**. On-board systems excel when timeliness, bandwidth conservation and autonomy are critical - particularly in missions where decisions must be made in almost real-time or where communication opportunities are limited. However, these benefits are inherently constrained by power, thermal and computational budgets, making them ideal for pre-processing, filtering, or basic analytics rather than complex interpretation tasks. In contrast, ground-based systems offer greater analytical flexibility and scalability, supporting the full spectrum of data products, historical comparisons and multi-mission integration. Yet, their dependency on downlink schedules and infrastructure introduces latency and higher operating costs. **Ultimately, hybrid architectures - where satellites perform lightweight, event-driven analytics on-board while offloading heavy processing and archiving to the ground - emerge as the most adaptive, cost-efficient and mission-responsive approach, particularly for multi-satellite constellations in the European context**. The most important key takeaways from **Chapter 3** are summarized as follows:

- ✓ **On-board processing is not a replacement for ground processing but a complementary capability** that should be strategically deployed where it adds mission value (e.g. low-latency alerts, data preselection, tasking autonomy).
- ✓ **Ground processing remains essential for complex and iterative analytics**, long-term storage and archival-quality products.
- ✓ **Hybrid models offer the best of both worlds**, combining on-board filtering and classification with high-throughput ground analysis and flexible reprocessing.
- ✓ **Mission type, urgency and data volume should drive processing location decisions**. For example, a wildfire monitoring mission may benefit from real-time on-board inference, while a climate change observation mission may prioritize full ground-based processing for archival accuracy.
- ✓ **Economic feasibility is mission-specific**. Small CubeSat constellations benefit disproportionately from on-board processing due to bandwidth savings, while larger institutional missions may afford heavier ground infrastructure investments.
- ✓ **Investments in software-defined and upgradable payloads** (*e.g. reprogrammable AI processors*) **are key enablers** of scalable and adaptable on-board processing in the European EO ecosystem.

4. Criteria and Guidelines for Processing Pathway Selection

Building upon the detailed comparative analysis conducted in **Chapter 3**, this chapter aims to translate those findings into practical guidelines for mission designers, system architects and decision-makers involved in planning EO missions. Rather than presenting a rigid framework, the goal here is to offer a flexible decision-support toolset that accounts for the varying objectives, constraints and operational environments of different EO scenarios. With EO missions becoming increasingly diverse - from rapid-response CubeSat constellations to high-fidelity, climate-focused observatories - the choice between on-board, ground-based, or hybrid processing architectures must be made with mission specificity and contextual nuance in mind.

The following sections outline the key factors that should inform processing architecture decisions. These include not only the technical and operational considerations covered earlier (*e.g., latency tolerance, data volume*), but also mission purpose, geographic coverage needs and temporal revisit requirements. Each of these factors introduces different processing demands and trade-offs.

*During the development of the decision-making framework in **Chapter 4**, the framework was disseminated across the consortium for feedback and validation. Partners were invited to review both the qualitative decision tree and the quantitative scoring model, with the goal of testing their clarity, applicability and actual feasibility against different mission contexts. The feedback received was used to refine the structure of the framework, clarify the parameter definitions and strengthen the consistency of the scoring thresholds. This partner-driven review process ensured that the framework is not only theoretically sound but also practical and aligned with the operational realities and priorities identified across the project.*

Traditional EO data processing architectures typically follow a sequential ground-centric flow: raw data is captured by the satellite and transmitted in bulk to Earth, where it undergoes multiple stages of processing-radiometric correction, georeferencing, analysis and dissemination. While this model offers centralized control and access to powerful computational resources, it also introduces several bottlenecks. These include delays in product delivery due to limited downlink windows, bandwidth saturation and increased operational overhead for data sorting and prioritization. As explained in [57], the conventional flow relies heavily on ground infrastructure to transform raw observations into actionable insights, which may not be ideal for time-sensitive applications (*see **Figure 4***). In the following sections, we explore how different mission characteristics challenge this model-and how emerging processing paradigms such as on-board intelligence or distributed hybrid architectures offer viable alternatives tailored to specific operational contexts.

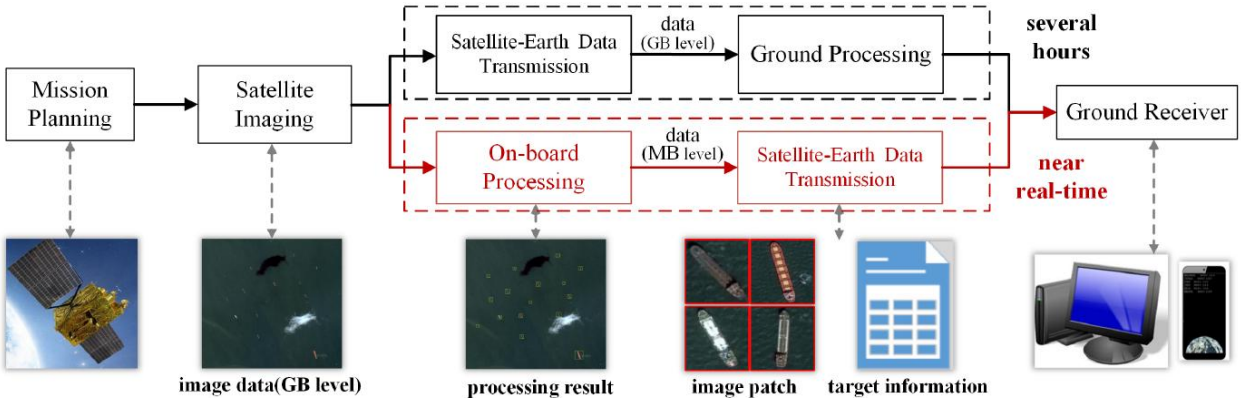


Figure 4. Comparison between on-board processing flow and the traditional ground processing flow [57]

To keep the framework consistent and auditable, we define the following terms once:

- Latency Tolerance (LT): operational time sensitivity of the product.
  - $\begin{cases} LT = 1 & \text{real} - \text{time} \ (\approx \leq 5 \ \text{min}) \\ LT = 2 & \text{rapid} \ (\leq 15 \ \text{min}) \\ LT = 3 & \text{delayed} \ (\geq 24 \ \text{h}) \end{cases}$
- Downlink Bottleneck Index (DBI): DBI = $\frac{Expected\ daily\ data\ volume\ (Gb/day)}{Expected\ daily\ downlink\ capacity\ (Gb/day)}$
  - DBI < 1: bandwidth comfortable • 1–3: tight • ≥2: structurally bandwidth-limited.
- Downlink capacity = $\sum_{i=1}^{N}(t_{contact,i} \times R_{net,i} \times \eta_{link})$, where:
  - $t_{contact,i}$ is the duration of ground pass i (in sec)
  - $R_{net,i}$ is the net throughput during pass i (in Gb/sec)
  - η_{link} is the efficiency factor (protocol + coding + availability)
- Interpolation bands:
  - $\begin{cases} DBI < 1 & comfortable\ (no\ bandwith\ stress) \\ 1 \leq DBI \leq 2 & tight\ (bandwidth\ constrained) \\ DBI \geq 2 & bandwidth\ limited\ (backlog\ expected) \end{cases}$
- On-board Compute TRL: flight maturity of the target on-board HW/SW stack.
  - If TRL < 6, downgrade on-board options to hybrid or ground.
- Update Path Readiness: existence of a secure OTA mechanism (*signed bundles, verification and rollback*) for models/pipelines.
- Traceability/Compliance Rigor: required strength of provenance, reproducibility, auditability (*e.g., CSS workflows*).
- Scoring scale mapping: Criteria are scored 0..10 and combined via a weighted average, then mapped to 0..100 by multiplying by 10 (*see **Section 4.2.2***).
  - Bands: 0–33 Ground, 34–66 Hybrid, 67–100 On-board.

4.1 Mission-Specific Factors

Selecting an optimal EO data processing pathway depends strongly on the specific mission context. Different mission types impose distinct requirements on data processing in terms of timeliness, data quality and volume. In particular, whether a mission is designed for rapid response or for long-term monitoring will shape the processing strategy. Similarly, the urgency of information needs (*latency requirements*) and the geographic scope of the mission (*coverage area and revisit frequency*) are critical factors. This section examines these mission-specific factors in detail, drawing on institutional and academic sources to highlight how each consideration influences EO data processing choices.

4.1.1 Type of mission

EO missions vary widely in their objectives, which can range from acute disaster response to continuous environmental monitoring. Missions oriented toward disaster response or crisis management prioritize rapid situational awareness. In such cases, satellites are tasked to capture and deliver imagery as quickly as possible to support emergency decision-making. For example, the Copernicus Emergency Management Service (EMS) is explicitly designed to provide "timely and accurate geospatial information" in the immediate aftermath of disasters. The Sentinel satellite missions operated by the European Union illustrate this focus: they support emergency management by providing data on a near-real-time basis, supplying responders with up-to-date imagery during floods, earthquakes, or other crises [46]. In practice, this means that during an event, the availability of current EO information is often valued more than the highest possible resolution or exhaustive accuracy. A recent study on flood response found that getting current satellite data to responders within hours was deemed more crucial than awaiting higher-resolution imagery that might come days later [47]. This reflects a common trade-off in disaster-oriented missions: speed and coverage take precedence over slight gains in precision. Notably, operational programs use all-weather, day-night sensing capabilities (*like SAR*) to ensure data can be collected regardless of conditions and they may employ pre- and post-event image comparisons to quickly map impacts. In summary, for event-driven missions such as emergency response, the processing pathway is optimized for fast turnaround and robust, "good-enough" information that can save lives and guide relief efforts [48].

In contrast, long-term monitoring missions focus on tracking environmental changes, climate variables, or other phenomena over extended periods. These missions support science and policy by building consistent data records rather than immediate actionable alerts. As a result, they emphasize data quality, calibration and continuity. For example, climate monitoring programs require highly accurate and stable measurements to detect subtle trends over years or decades. Satellites like NASA's GRACE (Gravity Recovery and Climate Experiment) provide valuable long-term datasets (*e.g. of water reserves or ice mass*) but with coarse resolution and higher latency,

making them unsuitable for real-time crisis response (*their monthly data cycles are too slow for disaster management*). Instead, GRACE and similar missions excel at long-term environmental assessment, where delays of days or weeks have little impact on utility [47]. Likewise, the European Copernicus program includes the Climate Change Service, which ingests satellite observations into climate data records; such efforts tolerate longer processing times to ensure accuracy and homogeneity across years of observations. Long-term EO missions often reprocess data archives when algorithms improve, underscoring the priority on scientific rigor over immediate use. As NOAA notes, polar-orbiting environmental satellites contribute to "long-term climate and environmental data records" even as others support short-term forecasting. In practical terms, a long-term monitoring mission's data processing pathway can afford to be ground-centered and batch-oriented, producing "standard" science-grade products with thorough validation [49]. The type of mission thus dictates whether the EO data processing should be geared toward rapid delivery (*as in disaster response*) or toward creating high-quality, consistent datasets (*as in climate or resource monitoring*). Many EO initiatives balance both aims, but clarifying a mission's primary type is a first step in tailoring the processing approach.

4.1.2 Urgency and latency needs

The urgency with which information is needed - often expressed as allowable data latency - is a pivotal factor in EO data processing design. Latency is defined as the time lag between an observation being made (*satellite image capture*) and the data or derived product being delivered to the end user [50]. Missions that require low latency generally demand an EO processing pipeline that can deliver useful information in near real-time. For instance, in wildfire or flood emergencies, receiving satellite analysis even a few hours sooner can significantly improve response outcomes [48]. The authors demonstrated that routing emergency vehicles with the aid of flood extent maps available within hours of a satellite overpass can reduce delays and economic losses in a disaster scenario. In essence, when human lives or critical infrastructure are at stake, the speed of data delivery becomes as important as the data itself. Recognizing this, agencies have developed specialized systems for low-latency data. NASA's Land, Atmosphere Near real-time Capability for EOS (LANCE) provides certain satellite products within ~3 hours of observation to support tactical applications like wildfire management, volcanic eruption alerts and extreme weather tracking [52]. User communities have indicated that a few hours' latency is the threshold for these applications to be effective - for example, wildfire responders can work with a 3-hour old satellite fire map, but a 24-hour delay would render it far less useful [52] [49]. In the European context, the emphasis on low latency is evident in initiatives like EO-ALERT, an Horizon-2020 project that prototyped an "intelligent satellite" system for real-time civil alerts. Performing on-board processing and filtering of imagery in space enabled EO-ALERT to demonstrate end-to-end delivery of warnings in near-real-time, dramatically reducing latency compared to traditional approaches [51]. Such edge-processing architectures show that when

urgency is paramount, shifting some processing from ground to on-board can enable users to "see what they want, when they want, close to real time" (*as the EO-ALERT team described, highlighting the importance of speed for emergency response*) [51].

Not all missions demand the lowest possible latency and there is often a trade-off between speed and data refinement. Missions with less urgent needs may opt for more extensive ground processing, generating higher-level, calibrated products on a longer timeline. In EO data management, a distinction is made between Near Real-Time (NRT) data streams and standard or delayed data products. NRT data are typically available within 1–3 hours after acquisition, serving applications that require immediate awareness [49]. These rapid products may be generated with automated algorithms and preliminary calibrations. In contrast, science-grade or "standard" products might only be released 1–2 days (*or even weeks*) after acquisition, after undergoing rigorous processing, quality control and integration of additional data for accuracy [49]. For example, Sentinel satellite data centers produce near real-time imagery for quick viewing and analysis, but final analytic products (with geometric precision corrections or atmospheric calibrations) are delivered later, once all ancillary information is applied. The appropriate latency level is dictated by mission needs: an agricultural monitoring service tracking crop health over a season can accept a delay of a day or two to receive higher-quality imagery, whereas a flood mapping service for first responders cannot. Copernicus EMS addresses this by offering different service levels - Rapid Mapping for fast (*hours-to-day*) products in emergencies, versus risk recovery mapping for less time-critical analysis. In designing the processing pathway, mission planners must therefore weigh how critical timeliness is relative to other factors. When low latency is essential, strategies include deploying accelerated processing pipelines, using high-bandwidth downlinks and possibly simplifying algorithms to run quickly. An extreme approach is on-board data processing, wherein satellites process raw data in situ and transmit only key results or alerts to the ground. This approach, increasingly feasible with modern on-board computers, can bypass delays associated with downlinking large raw datasets and queuing for ground processing.

Commercial and institutional R&D in Europe have highlighted on-board processing as a way to meet latency requirements for security and disaster missions (KP Labs, 2025). For instance, experimental small satellites like ESA's OPS-SAT have tested real-time image processing in orbit to immediately detect events [53]. The clear lesson across these efforts is that urgency drives the need for low-latency data workflows. Missions must define acceptable latency thresholds (*e.g. minutes, hours, days*) and the EO data pipeline – whether on ground, in the cloud, or on the satellite – is then tailored to meet those timelines. Crucially, this must be achieved without overwhelming users with false or unverified information, so even low-latency systems incorporate checks to maintain reliability. Overall, aligning the processing pathway with the

mission's urgency ensures that end-users receive information in time to act upon it, which is a cornerstone of EO utility in operational settings [53].

4.1.3 Geographic coverage and revisit rates

The geographic coverage required by a mission and the revisit rate (*frequency of observations over a given area*) are additional mission-specific factors that fundamentally influence EO data processing. Different missions demand coverage at different scales – some focus on local or regional areas, while others require global surveillance. Likewise, some applications need observations of the same location multiple times per day, whereas others can work with weekly or monthly revisit. These requirements will determine the choice of satellite or constellation and the resulting data volume and processing load. Missions with broad geographic coverage (*continental or global*) typically rely on satellites in low Earth polar orbits that can systematically scan the Earth over time, or on constellations of many satellites working in concert. The Copernicus programme, for example, aims for global coverage to support its services, deploying multiple Sentinel satellites to ensure no area of the world is left unmonitored for long. A wide coverage mission, such as global land cover mapping or climate observation, generates large volumes of data that must be ingested and processed, often in near-continuous streams. Data processing architectures for such missions need to be highly scalable and automated to handle the sheer amount of imagery. For instance, NASA's MODIS instruments capture the entire Earth's surface every 1–2 days, producing terabytes of data that are processed into global products like vegetation indices and cloud masks.

Ensuring timely processing of global datasets may involve distributed ground segment networks and cloud processing to keep up with the data flow. One consequence is that data downlink and storage become critical considerations: global missions often push the limits of bandwidth and processing pipelines must include data thinning, compression, or on-board preprocessing to manage the load [54]. As noted in this work, the explosion in small satellite constellations and the increasing revisit of commercial EO systems have led to an exponential growth in data volumes. While this offers unprecedented analytical opportunities (*e.g., daily monitoring of every point on Earth*), it also requires more efficient processing techniques *(like real-time filtering or AI-based analysis*) to extract useful information without human analysts being overwhelmed [54]. Thus, missions with expansive coverage must architect their data processing with big data in mind – using high-performance computing and intelligent algorithms to turn raw imagery into insight across wide areas. On the other hand, missions with limited or targeted geographic coverage (such as monitoring a specific city, critical infrastructure site, or a narrow border region) can tailor their data processing more narrowly. Because the area of interest is confined, these missions might utilize higher-resolution sensors or more frequent collects over that area without needing to handle global-scale data. The processing system in this case can be optimized for detail and accuracy over volume - for example, an urban air quality monitoring mission might

integrate satellite data with local sensors in a focused way. However, limited coverage missions face the challenge of ensuring the satellite is overhead when needed.

This brings the concept of revisit rate into play. Revisit rate is the interval at which a given location is observed by the satellite(s). A high revisit (*say multiple times per day*) is essential for missions like weather nowcasting or certain security surveillance tasks, where phenomena can change hourly. Geostationary satellites illustrate this: positioned over a fixed region, they provide continuous or frequent (*e.g., 5-15 minute*) imagery of that same area. For example, according to [55], EUMETSAT's Meteosat geostationary satellites deliver rapid-repeat images for Europe and Africa, enabling real-time weather monitoring, but they only cover their fixed hemisphere. Conversely, polar-orbiting satellites like Sentinel-2 have a global reach but revisit a specific location on the order of days. Sentinel-2's twin satellites achieve roughly a 5-day revisit for any given spot at the equator, which is sufficient for tracking gradual land changes but not for immediate event detection. To increase revisit frequency, operators can deploy constellations of satellites. Commercial providers (*e.g., Planet with its PlanetScope Doves*) have exemplified this by fielding dozens of mini-satellites to image the entire Earth daily (or even multiple times per day in mid-latitudes) [56]. Such high-revisit constellations drastically improve the temporal resolution of observations, ensuring that transient events (*like sudden flooding, illegal mining activities, or rapid crop growth stages*) are caught in at least one of the passes. From a data processing perspective, higher revisit means more frequent data ingest cycles and often more automated change-detection or alerting algorithms – since comparing images from consecutive passes becomes a primary analysis mode. The system must be capable of updating outputs (maps, statistics, alerts) at the same pace as the data arrives. In addition, frequent revisits can enable near-continuous monitoring, but only if the processing can keep up and deliver insights continuously as well. In summary, the spatial and temporal coverage demands of an EO mission set the stage for its processing strategy.

Ultimately, we would say that missions requiring global coverage and high revisit tend to produce "continuous" streams of data, benefiting from robust, scalable processing pipelines (*often in the cloud or through distributed ground stations*) and sometimes employing selective data handling (*to focus on areas of change*). Missions with more focused coverage might concentrate on higher resolution data and can be more targeted in processing, but they must consider how to achieve sufficient revisit - possibly via multiple satellites or complementary platforms (*like UAVs or high-altitude platforms for gap-filling*). Ultimately, aligning the processing pathway with geographic coverage and revisit requirements ensures that the mission delivers data at the right place and right time. For example, a mission tasked with daily tropical deforestation monitoring will combine a wide coverage sensor with daily processing cycles and automated alerts, whereas a mission monitoring polar ice caps may integrate multiple satellite sensors to achieve both broad coverage and the necessary revisit interval for seasonal changes. All of these result in the

following: accounting for these factors in the design phase, mission planners can select a data processing approach (*on-board vs on-ground processing, real-time vs batch processing, etc.*) that best matches the mission's spatial scope and temporal observation rhythm.

4.1.4 Decision Matrices and Architecture Trade-offs

To consolidate the diverse mission-specific drivers discussed in the previous sections, this part introduces structured comparison tools to aid EO system designers in selecting the most suitable data processing architecture. Rather than offering prescriptive solutions, the following matrices reveal insights from real-world missions and institutional research into practical, comparative frameworks. By mapping processing models to mission types, system constraints and operational contexts, these decision-support tools provide a bridge between theoretical considerations and implementable strategies. The SWOT analysis further deepens this perspective by identifying the strengths, weaknesses, opportunities and risks associated with on-board processing-now emerging as a key enabler of responsive and autonomous EO systems.

Table 5. Alignment of EO Processing Architectures with Mission Types

Dimension	On-board processing	Ground-based processing
Disaster Response / Emergency	On-board or Hybrid	Prioritizes low latency; enables rapid alerts and bandwidth-efficient downlink (*e.g., Φ-Sat-1, EO-ALERT*).
Climate Change Monitoring (Long-Term)	Ground-based	Requires high accuracy, consistency, calibration; tolerates high latency (*e.g., GRACE, Copernicus C3S*).
Maritime Surveillance (Continuous Watch)	Hybrid	Frequent observations + some real-time filtering (*e.g., Unibap's vessel tracking*).
Agricultural Monitoring (Seasonal Trends)	Ground-based / Hybrid	Weekly revisit is sufficient; focus is on multispectral calibration and trend analysis.
Urban Air Quality or Infrastructure	Targeted On-board / Hybrid	Area-specific missions can use on-board filtering + ground correlation with sensor networks.
Illegal Activity Detection (e.g. mining)	On-board or Hybrid	High revisit + rapid detection needed; constellations benefit from edge analytics.

Table 5 offers a mission-type centric overview of how processing models align with real-world operational needs. It highlights the critical insight that architecture decisions should always follow mission purpose. For example, on-board processing clearly benefits time-critical missions such as disaster response and illegal activity monitoring, where near-real-time analytics can

trigger alerts or actions before the next ground pass. Conversely, climate change monitoring and environmental trend tracking favor ground-centric models due to their demand for long-term data consistency, recalibration and archival accuracy. The table also shows that many mission categories fall into a hybrid processing space, combining lightweight on-board analytics with detailed ground-based post-processing. **This reinforces the idea that a "*one-size-fits-all*" approach is suboptimal and that tailoring the architecture to mission urgency, resolution needs and revisit constraints is essential**.

Table 6. Factor-Based Comparison Table

Dimension	On-board processing	Ground-based processing
Latency	Enables immediate detection and alerts from orbit	Dependent on downlink windows and processing pipeline
Data Volume	Reduces burden via compression, pre-filtering (e.g. discard cloudy scenes)	Handles full datasets; best for post-hoc analysis
Geographic Coverage	Less scalable; limited storage and compute on-board	Supports global processing via distributed systems and HPC
Revisit Frequency	Helps prioritize high-frequency images (e.g., flag changes), enabling bandwidth efficiency	Handles all data regardless of revisit; best for retrospective analysis
Processing Flexibility	Limited to pre-uploaded models or software-defined logic	Highly flexible, supports algorithm refinement and reprocessing
Regulatory Control	Less transparent; challenges with on-board AI explainability and GDPR	Easier to audit, log and manage user access under EU data protection frameworks

Table 6 reframes the discussion from a mission-type view to a design factor perspective, comparing how different processing models handle key system-level considerations. **One of the most revealing dimensions here is latency: on-board systems clearly excel in minimizing delay between data acquisition and insight generation, while ground-based models dominate in post-hoc flexibility and data richness**. The table also emphasizes the trade-offs around data volume: while ground-based processing handles complete raw datasets for retrospective analysis, on-board filtering minimizes unnecessary downlink, enhancing efficiency. Furthermore, regulatory control emerges as a growing differentiator, with ground systems offering clearer audit trails and GDPR compliance workflows. **This comparison reinforces that selecting a processing strategy is not just about technology - it also reflects risk tolerance, compliance needs and the degree of mission flexibility required**.

Table 7. SWOT Analysis for on-board processing (extended)

Strengths	Weaknesses
➢ *Low latency / real-time capability* ➢ *Data pre-filtering reduces bandwidth* ➢ *Enables autonomy* ➢ *Efficient for tactical missions*	➢ *Limited compute and power on-board* ➢ Harder to update or re-train models ➢ Explainability and trustworthiness issues ➢ Regulatory gaps (e.g. GDPR compliance)
Opportunities	**Threats**
➢ *Software-defined payloads (e.g. SpaceCloud)* ➢ *AI advances enable more on-board analytics* ➢ *Federated learning (future potential)*	➢ *Misclassification or data loss on-board* ➢ *Regulatory scrutiny around decision-making* ➢ *Incompatibility across satellite constellations*

Table 8. SWOT Analysis for ground-based processing (extended)

Strengths	Weaknesses
➢ *Flexible, scalable and upgradeable* ➢ *Easier compliance and audit* ➢ *Suitable for data fusion and archival*	➢ *High bandwidth and infrastructure dependency* ➢ *Higher latency; slow in time-sensitive use* ➢ *Bottlenecks in crowded downlink environments*
Opportunities	**Threats**
➢ *Cloud scalability and AI integration* ➢ *Centralized control and versioning*	➢ *Operational delays during emergencies* ➢ *Overwhelmed analysts with too much data*

The SWOT analysis in **Table 7** and **Table 8** provide a balanced strategic lens for evaluating the architectural trade-offs between on-board and ground-based processing in EO missions. **On-board processing demonstrates clear strengths in reducing data volume, enabling mission autonomy and delivering rapid or real-time insights** - qualities that are particularly valuable in tactical or time-sensitive contexts such as disaster response, illegal activity detection, or maritime surveillance. Its strengths align closely with trends in software-defined payloads, AI acceleration and federated learning potential, which are reinforced by European innovation programs like Φ-lab and Horizon Europe. However, **the weaknesses and threats are non-trivial: limited computational power, difficulty updating algorithms post-launch, regulatory** ambiguity (e.g. GDPR) **and the challenge of explainability in AI models on-board pose significant risks**. Conversely, the SWOT for ground-based processing highlights its stability, flexibility and maturity. **Strengths include scalability, ease of algorithm updates, regulatory transparency and suitability for high-fidelity archival processing and long-term environmental monitoring**. Ground-based workflows excel in scientific missions where latency is less critical but analytic depth, version control and traceability are essential. Nonetheless, **weaknesses emerge in high-throughput or urgent missions, where latency and bandwidth dependencies introduce significant operational**

delays. Opportunities such as cloud-native architectures and centralized AI integration offer promise, but the threats remain - especially during emergencies where bottlenecks or infrastructure downtime can severely hamper performance.

4.2 Decision-Making Framework

So far in D5.4, we have analyzed the technical, operational and economic characteristics of EO data processing approaches - and having mapped these against mission-specific drivers - this section will be introducing a structured decision-making framework designed to support system architects, mission designers and institutional stakeholders in selecting the most appropriate data processing pathway. **This framework does not impose a one-size-fits-all solution but offers a modular, flexible structure adaptable to different mission types, operational tempos and user requirements.** At its core, the framework comprises two interrelated components:

1. **A structured decision tree or flowchart**, which provides step-by-step guidance by asking key binary or scaled questions that lead to suggested processing architectures (on-board, ground-based, or hybrid);
2. **A set of scoring thresholds and evaluation criteria**, which allow users to quantify key mission parameters and derive a weighted recommendation.

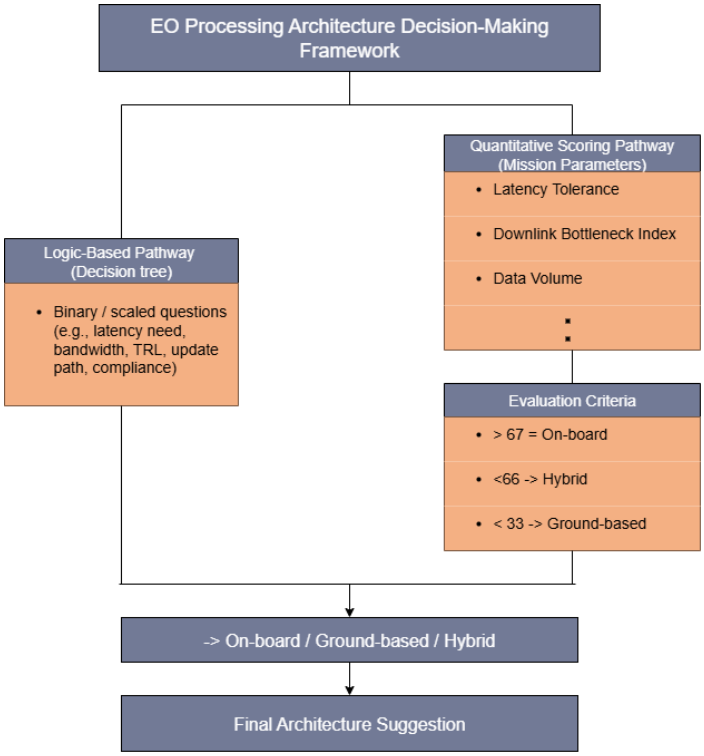


Figure 5. High-level EO decision framework

In the next two subsections, we define how these tools operate. **Section 4.2.1** presents the decision flow logic, capturing high-level architecture choices based on mission purpose, latency tolerance and autonomy needs. **Section 4.2.2** introduces a threshold-based scoring methodology, allowing mission planners to assign weights to critical factors and receive a numerically supported output to guide internal deliberations and stakeholder consultations.

4.2.1 Structured flowchart / decision tree

The structured flowchart presented in this section constitutes the **first core pillar of our EO Processing Decision-Making Framework**. Its objective is to guide mission planners, system architects and decision-makers through a logic-based evaluation pathway, allowing them to determine whether a mission is best suited to on-board processing, ground-based processing, or a hybrid model. **Unlike generic checklists, this flowchart applies a context-sensitive and branching logic, making it ideal for early-stage architecture decisions**. It starts with high-level questions related to mission urgency and operational complexity and then narrows down through binary and scaled queries that progressively define the most suitable processing mode. The structure ensures that no critical dimension is overlooked - including latency tolerance, on-board compute capacity, regulatory exposure and bandwidth constraints.

4.2.1.1 Decision tree

The decision tree consists of a step-by-step sequence of decision points, typically phrased as yes/no (binary) or scaled questions. Each response progressively filters the possible processing approaches and brings the user closer to a justified architectural suggestion.

Table 9. Decision-framework structured pathway

Step	Decision Question	Path Response
Q1	Is Latency Tolerance ≤1 (≤5 minutes)?	Yes → Go to Q2
		No → Proceed to Q3
Q2	Is there a need for real-time alerts or tactical decision support?	Yes → Go to Q3
		No → Proceed to Q4
Q3	Can the satellite support a low-power AI chip for in-orbit analytics (e.g., Myriad VPU, FPGA)?	Yes → If LT ≤1 and DBI ≥2 → On-board. Otherwise → Hybrid.
		No → Hybrid
Q4	Is the mission focused on long-term environmental monitoring or climate record building? (*i.e., Latency Tolerance ≥3; ≥24h acceptable*).	Yes → Likely Ground-based
		No → Proceed to Q5
Q5	Will the sensors generate high data volumes (e.g., hyperspectral or SAR payloads)?	Yes → Go to Q6
		No → Proceed to Q7
Q6	Is DBI ≥2 (structurally bandwidth-limited)?	Yes → If LT ≤1 and DBI ≥2 → On-board. Else → Hybrid.
		No → If DBI <1 after contact scheduling → Favor Ground.
Q7	Do we need to frequently revisit the same area (e.g., multiple times per day)?	Yes → Consider Hybrid with constellations or on-board prioritization
		No → Proceed to Q8
Q8		Yes → Go to Q9

	Is the compute power available on-board sufficient to run basic inference or preprocessing (e.g., filtering, compression)? (*On-board compute TRL ≥6 required*).	No → Ground (compute TRL <6).
Q9	Is the satellite intended to be reprogrammable or updatable post-launch?	Yes → On-board if LT ≤1 and DBI ≥2. Hybrid otherwise
		No → Proceed to Q10
Q10	Is Traceability/Compliance Rigor required? (*e.g., GDPR, CSS-class workflows*)	Yes → Ground, unless on-board can provide provenance logs and explainability
		No → Proceed to Q11
Q11	Will the mission run for more than 5 years, requiring high reliability and supportability?	Yes → Ground (long duration & high reliability).
		No → Proceed to Q12
Q12	Is the mission operating under tight budget constraints (e.g., academic or small commercial)?	Yes → If latency is critical, opt for On-board-lite (e.g., compression only)
		No → All options open

4.2.1.2 Reasoning for the decision tree

This interaction between the two logic streams is clearly depicted in ***Figure 5***, which visualizes the two evaluation branches feeding into a common conclusion - the final processing architecture recommendation. The figure emphasizes how both the **logic-based pathway** (*left side of the diagram*) and the **quantitative scoring pathway** (*right side*) ultimately converge toward a unified outcome, potentially enriched by confidence levels or probability weights. The reasoning behind selecting the appropriate questions is explained below:

Table 10. Reasoning behind the Decision-framework pathway

Question	Reasoning
Q1	This initial sets the latency threshold that separates time-sensitive missions (such as disaster response, security, or tactical intelligence) from those that can afford slower processing timelines. A Latency Tolerance ≤1 (≤5 minutes) indicates missions that require near-real-time outputs, while higher tolerance values move the evaluation toward less time-critical architectures.

Q2	This question tests whether the mission requires real-time alerts or tactical decision support. If yes, the architecture must support low-latency outputs, favoring on-board or hybrid models. If no, latency is less critical and ground or hybrid approaches remain suitable.
Q3	This question checks whether the satellite can host a low-power AI accelerator (e.g., VPU, FPGA). If available and combined with high latency sensitivity and bandwidth limitations, on-board processing becomes viable. If no such hardware is available, hybrid models provide a more practical balance by combining limited pre-processing in orbit with deeper analysis on the ground.
Q4	Missions focused on long-term environmental monitoring or climate record building typically tolerate higher latency (≥24 hours). These missions emphasize consistent calibration, archival quality and multi-year trend analysis, which are best supported by ground-based processing.
Q5	High-volume payloads such as hyperspectral or SAR instruments generate large data flows. This step identifies whether data throughput is a critical factor that could strain ground infrastructure and therefore requires consideration of bandwidth constraints in the next step.
Q6	DBI quantifies the relationship between daily data volume and available downlink capacity. A DBI ≥2 indicates structural bandwidth limitations, favoring on-board or hybrid models. A DBI in the range of 1–3 supports hybrid solutions, while DBI <1 suggests that ground processing is sufficient.
Q7	Missions that require frequent revisits of the same area benefit from hybrid or on-board prioritization. This avoids overwhelming the downlink by enabling pre-filtering or prioritization of data at the source.
Q8	This step evaluates whether sufficient compute power is available on-board to run basic inference or preprocessing tasks. A Technology Readiness Level (TRL) of at least 6 is required to ensure operational feasibility. If this threshold is not met, ground-based processing is recommended.
Q9	The availability of reprogrammable or updatable payloads significantly increases the feasibility of on-board or hybrid processing, as algorithms and models can be updated post-launch. If this capability is absent, hybrid or ground-based architectures are safer to ensure mission adaptability over time.
Q10	Missions subject to traceability or compliance requirements, such as GDPR or CSS-class workflows, require auditable and reproducible outputs. In such cases, ground-based processing is generally favored unless on-board systems can provide verifiable provenance logs and signed outputs.
Q11	Missions with durations longer than five years place emphasis on long-term reliability and supportability, effectively met by ground-based architectures.
Q12	Missions operating under tight budget constraints may opt for simplified “on-board-lite” approaches if latency is critical. In such cases, lightweight on-board functions (e.g., compression, cloud masking) can provide responsiveness at minimal cost, while other missions may retain all options depending on priorities.

One of the most significant advantages of the logic-based evaluation method lies in its alignment with mission context awareness. Unlike pure numerical models, **the flowchart structure enables the incorporation of qualitative judgments that are difficult to translate into hard scores** - such as "*how much does it matter if delivery is delayed by six hours?*" or "*how critical is local autonomy to the mission?*" While the decision tree helps generate qualitative architectural suggestions, the scoring pathway applies a structured numerical evaluation that quantifies key mission characteristics - such as latency sensitivity, bandwidth availability and on-board compute resources. **If the two pathways converge on the same outcome, it provides high confidence. If they diverge, it can flag areas where trade-offs or further discussion may be necessary - such as technical feasibility versus regulatory risk, or low latency needs versus compute limitations**.

4.2.1.3 Flowchart

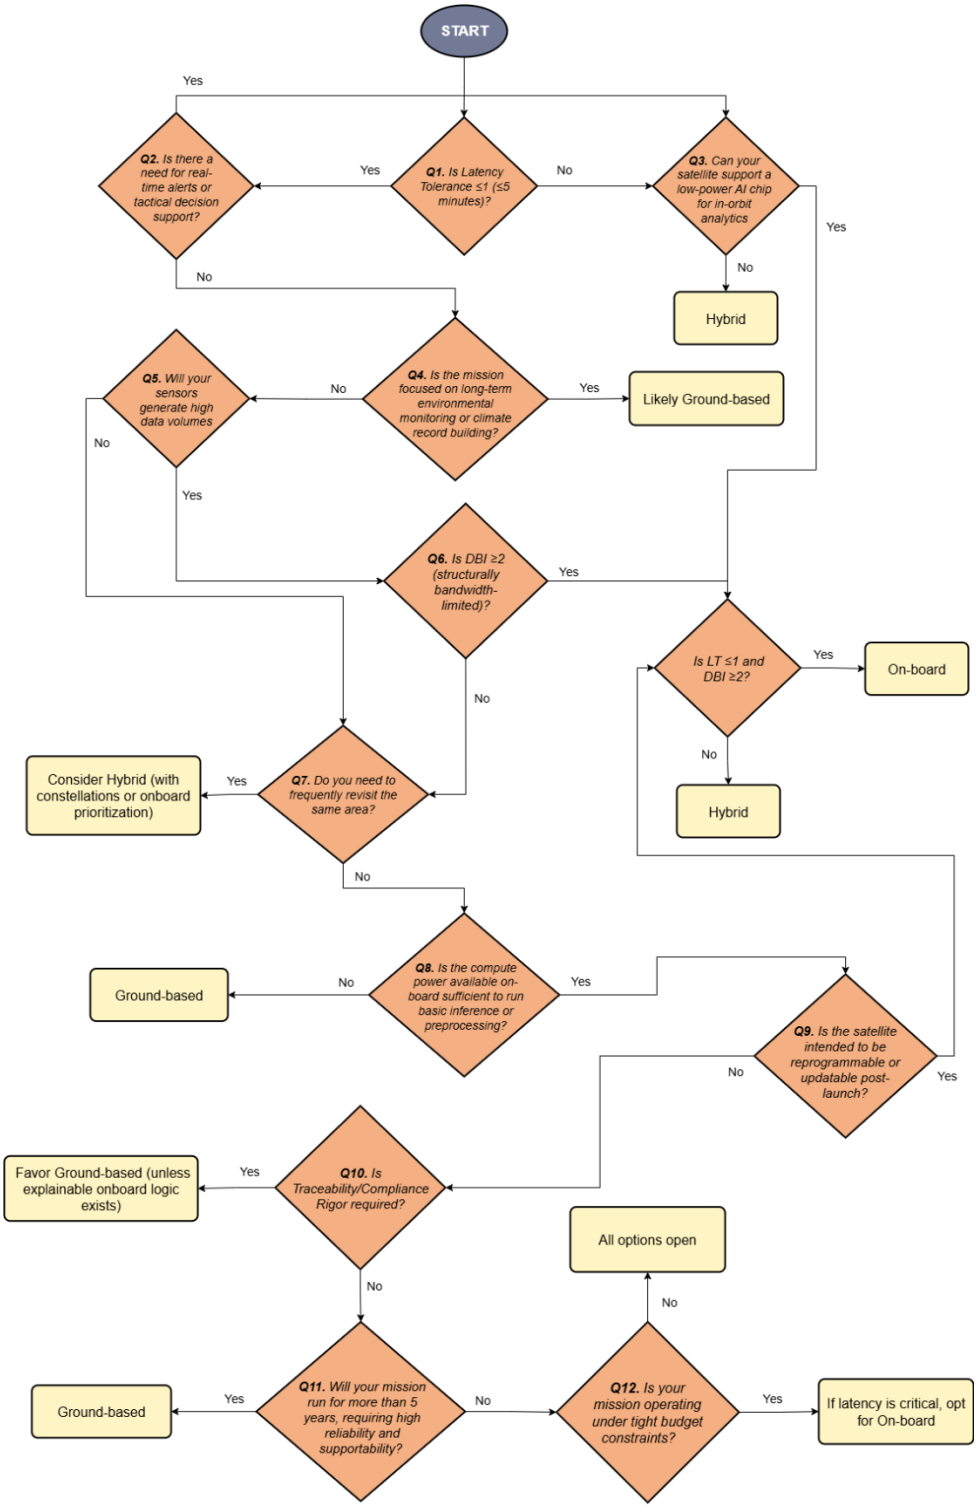


Figure 6. Flowchart for the decision-framework

The flowchart presented in ***Figure 6*** serves as a practical and intuitive way to navigate the complexity of the architecture decisions. The flowchart depicts and allows a vast space of possible configurations into a focused set of architectural options by guiding users through a series of well-sequenced questions (*e.g. ranging from latency needs and bandwidth constraints to on-board compute capability and data governance concerns*). When used alongside with the quantitative scoring model in the next section (**Section 4.2.2**), it becomes a powerful element of a comprehensive architecture selection framework.

4.2.2 Thresholds and scoring mechanisms

Following the qualitative logic-based framework introduced in **Section 4.2.1**, this section introduces a complementary quantitative scoring methodology that enables EO mission designers and system planners to numerically assess mission requirements and derive a confidence-weighted recommendation for selecting the optimal processing architecture, a) on-board, b) ground-based, c) or hybrid. Whereas the flowchart provides rapid decision-pathing through binary and context-specific questions, the scoring system operates as a weighted checklist that allows to assign scores to key operational factors based on mission specifications. The cumulative score can help validate the output of the decision tree - or, in cases where the decision is borderline or complex, it can offer granular insight into which factors are driving the recommendation and where trade-offs exist. The framework evaluates **eight (8) core parameters**, each representing a key driver of processing architecture choice. These are:

1. **Latency tolerance**
2. **Bandwidth constraints**
3. **Data volume (daily)**
4. **On-board compute availability**
5. **Mission urgency**
6. **Revisit frequency**
7. **Regulatory/sensitivity concerns**
8. **Need for real-time analytics**

Each parameter is scored on a scale of 0 to 10, where higher values favor on-board or hybrid processing and lower values indicate suitability for ground-based processing. Each factor may also be assigned a weight (W) between 1 and 3 depending on its criticality to the mission, allowing the score to reflect mission-specific priorities (*e.g., a flood response mission may assign W=3 to latency and urgency*). The final architecture suggestion is based on the weighted average of all parameters, which yields a total score between 0 and 100. To avoid ambiguity, the framework always uses a 0–10 scale for individual parameters. The weighted average is multiplied by 10 to normalize to a 0-100 band, ensuring consistency across all case studies. The system then interprets this result based on pre-defined scoring thresholds:

Table 11. Classifying the Weighted Score Average

Weighted Score Average	Recommended Architecture
0-33	Ground-based processing
34-66	Hybrid (consider on-board augmentation)
67-100	On-board or edge-oriented processing

Each of the eight parameters is briefly explained below, including the logic for score interpretation and examples.

Table 12. Scoring Criteria and Rationale

Parameter	Explanation	Rationale	Scoring
Latency Tolerance (LT)	How quickly must the data be delivered to end users?	Missions requiring fast turnaround naturally benefit from on-board or hybrid filtering, which reduces dependency on downlink schedules	0 = Acceptable latency >24 hours
			5 = Latency under 6–12 hours is ideal
			10 = Mission requires near-real-time insights *(≤5 minutes*)
Bandwidth Constraints	Is downlink bandwidth limited by orbital position, contact time, or relay infrastructure?	Limited bandwidth pushes toward on-board filtering or compression to reduce transmitted data volumes	0 = Mission has access to frequent, high-speed downlinks
			10 = Mission operates with limited downlink windows
			DBI mapping for this criterion: when available, score Bandwidth constraints by DBI: DBI < 1 → score 0–2, DBI 1–<2 → score 5–7, DBI ≥ 2 → score 8–10
Data Volume			0 = <1 GB/day

	What is the estimated daily data output of the payload(s)?	High-volume missions benefit from on-board reduction to avoid data congestion and delays	5 = 10–20 GB/day
			10 = >100 GB/day
On-board Compute Availability	Can the satellite host capable edge processors (e.g., Myriad VPU, FPGA, GPU)?	On-board processing requires a minimum threshold of computing resources and software support to be feasible	0 = No capacity on-board
			5 = Limited processing (*e.g., basic compression or anomaly flags*)
			10 = Equipped with programmable, AI-enabled processors (TRL ≥6).
Mission Urgency	Does the mission involve time-sensitive applications?	Urgency pushes toward low-latency architectures and favors edge processing	0 = Low-urgency mission
			10 = Tactical or emergency applications requiring immediate awareness
Revisit Frequency	How frequently must the same location be observed?	High revisit rate implies more data collection and often shorter windows to extract insight-favoring on-board filtering and tasking	0 = Monthly or seasonal coverage
			5 = Weekly revisit desired
			10 = Multiple revisits per day required
Regulatory and Sensitivity Concerns	Are there GDPR or national data sovereignty implications?	Sensitive missions may favor ground processing (score low) unless explainability and traceability are embedded into on-board logic	0 = Open-access, non-sensitive data
			10 = Data with location privacy, identifiable assets or defense implications
Need for Real-Time Analytics	Does the mission involve on-board inference (e.g., object detection)?	Missions needing autonomous analytics must lean toward edge-enabled payloads	0 = All analysis happens on Earth
			10 = Requires real-time on-board classification or alert generation

Note #1: "Mission urgency" captures the contextual importance of delivering products quickly (e.g., disaster response versus seasonal monitoring), while "Need for real-time analytics" captures the technical requirement for inference to run on-orbit (e.g., on-board object detection). It should be treated them as complementary but non-identical: urgency is about mission value, analytics is about system function.

Note #2: While the eight parameters above remain the core of the scoring model, architects should also record two "gating conditions" outside the weighted score: (i) On-board Compute TRL - architectures relying on components below TRL 6 should be downgraded unless risk-mitigated; and (ii) Update Path Readiness - missions without a secure, verifiable over-the-air update and rollback capability should bias toward hybrid or ground-based approaches. These gates do not alter the numeric score but provide essential context for interpreting borderline results.

Note #2: Interpretation gates (non-scored):

Gate 1 - On-board Compute TRL: if < 6, downgrade pure on-board recommendations unless a documented risk-mitigation plan exists.

Gate 2 - Update Path Readiness: if no secure OTA with rollback, favor Hybrid or Ground even when the numeric score leans on-board.

One of the primary strengths of the quantitative scoring system lies in its transparency. The approach is also highly reusable: because the framework is modular, it can be adapted across different mission profiles simply by adjusting the input scores and weights. This makes it applicable in early concept phases, trade studies, or architecture reviews across a wide range of EO scenarios. Another major advantage is its **ability to foster team alignment**. Because multiple mission stakeholders can independently contribute scores and weights, the scoring process becomes a collaborative tool for surfacing priorities and constraints.

4.3 Tactical/Operational Risk Assessment

The decision of whether to implement on-board or ground-based data processing must not rely solely on performance and cost evaluations. Tactical and operational risk-defined here as the likelihood and consequence of a system failure, performance degradation, or data loss during the mission lifecycle-represents a third critical pillar in the selection framework**. Assessing the risk profile of each processing approach helps mission designers better anticipate failure modes, system resilience challenges and context-specific vulnerabilities that could compromise mission success**.

On-board processing introduces new types of operational dependencies. Processing hardware on the satellite is exposed to harsh space environments, where radiation-induced errors, thermal stress and power constraints may reduce system robustness. Moreover, limitations in update flexibility *(e.g., difficulty in patching firmware post-launch)* present longer-term risks, particularly in missions expected to evolve over time or rely on adaptive AI-based decision-making. Failures in on-board processing often lead to irreversible data loss, since raw data is typically not transmitted back to Earth. By contrast, **ground-based processing centralizes computational resources in controllable environments**, which significantly reduces hardware exposure and allows for easier upgrades, maintenance and recovery actions. However, this model introduces its own vulnerabilities, particularly around communication latency, bandwidth availability and potential transmission failures. In tactical or time-critical scenarios-such as near-real-time emergency monitoring-these dependencies may present unacceptable risks if communication infrastructure is degraded or contested. Missions operating in denied or disrupted environments *(e.g., disaster zones or areas with contested satellite links)* may face heightened communication risks that favor on-board autonomy. Conversely, missions with guaranteed downlink infrastructure but requiring complex or iterative data analysis may benefit from the reduced system complexity and improved control of ground-based systems.

Table 13. Risk categories and mitigation techniques

Risk Category	On-Board Processing	Ground-Based Processing	Typical Mitigations
Hardware Reliability	Radiation effects, thermal fatigue, embedded HW failure	Ground infrastructure aging, server hardware failure	Redundancy, error-correcting memory, COTS qualification
Software Vulnerability	Limited debugging, AI model drift, firmware corruption	Centralized fault propagation, update-induced regressions	Simulation testing, OTA updates, containerized deployments
Communication Disruption	N/A (self-contained)	Downlink outage, bandwidth contention, latency	Multi-band redundancy, prioritization protocols
Data Integrity	Loss due to on-board filtering or buffer overflow	Loss due to corrupted transmission or incomplete downlink	Checksum verification, partial recovery protocols
Maintainability	Very limited post-launch maintainability	Full lifecycle flexibility for updates and scaling	Modular SW design, predictive maintenance tools

Ultimately, the decision to favor one path over the other should be weighted by mission class, timeline sensitivity, environmental volatility and tolerance for data loss or delay. For instance, a LEO disaster-response satellite might favor on-board processing to ensure real-time delivery under unstable links, whereas a climate-monitoring mission with persistent coverage and no urgent latency requirements may benefit from the flexibility and scalability of ground-based

processing. As this risk assessment shows, no single pathway is universally superior. Risk-aware architecture selection must align system resilience with operational priorities. The insights offered here complement the cost-performance evaluation and scoring methodologies in **Chapter 4**, forming a more holistic basis for mission-specific decision-making.

4.4 Example Decision Scenarios

This section provides a series of contextualized example scenarios that illustrate how the decision-making framework can be applied in practice. Building upon these scenarios, it further provides an explicit linkage between EO data processing pathways and the operational needs of the THEIA pilot use cases. Each scenario is designed to reflect real-world mission requirements, priorities and constraints, offering EO system designers and mission architects a reference point for evaluating trade-offs between on-board and ground-based data processing options.

4.4.1 Real-time coastal monitoring scenario

Real-time coastal monitoring missions often require rapid detection and reporting of vessel activity, marine anomalies, or environmental violations. In such time-sensitive scenarios, the choice of processing pathway *(on-board, ground-based, or hybrid)* significantly impacts response time and data efficiency. This use case is especially relevant for regions facing high volumes of maritime activity or irregular migration, where rapid response is essential to mitigate risks. From a processing architecture standpoint, real-time coastal monitoring strongly favors on-board processing capabilities. When EO satellites or UAVs are equipped with edge computing platforms, initial data filtering, object detection *(e.g., vessel recognition)* and threat prioritization can occur directly in-flight. This reduces the time-to-alert, enabling near-instantaneous transmission of only the most relevant results – such as geolocated detections or confidence-scored anomalies-back to ground control or command centers. The bandwidth savings are also substantial, particularly when monitoring large coastal areas where most of the captured imagery may be non-actionable.

However, relying exclusively on on-board processing introduces certain trade-offs. These include limited model complexity due to computational constraints, challenges in retraining or updating AI algorithms on orbit and a higher risk of missing edge cases or contextual cues that could be identified more robustly with high-performance ground infrastructure. In contrast, ground-based processing allows for the use of more sophisticated models, ensemble methods and data fusion with external intelligence sources *(e.g., AIS, radar, crowdsourced alerts)*, enhancing overall analytical depth. The optimal approach for real-time coastal monitoring may therefore involve a hybrid solution: performing time-critical inference *(such as vessel detection or thresholding)* on-board, followed by more advanced post-analysis on the ground for strategic evaluation or historical pattern recognition. This scenario highlights the importance of aligning processing strategies with mission objectives-balancing immediacy against completeness and computational

richness. The following table summarizes the relative suitability of each processing pathway *(on-board and ground-based)* based on the key evaluation criteria for this scenario.

Table 14. Evaluation of Processing Approaches for Real-Time Coastal Monitoring

Criterion	On-Board Processing	Ground-Based Processing
Latency	Very low (real-time alerts possible)	Higher (pass/ground-contact dependent)
Processing depth	Limited (edge inference models)	High (complex models, fusion with external data)
Bandwidth usage	Low (transmit only filtered results)	High (requires full raw data transmission)
Adaptability	Lower (harder to update on orbit)	High (easily adaptable models and methods)
Use case fit	Immediate detection & alerting	Deep analysis, forensics and mission planning

In this scenario, a CubeSat constellation is designed for near-real-time wildfire detection in remote areas, where latency, revisit frequency and mission urgency are paramount. This worked example illustrates how the scoring and weighting methodology introduced in Section 4.1 can be applied to a concrete mission concept. The table below scores and weights key evaluation parameters to determine the suitability of on-board vs ground-based processing. The resulting weighted average can help mission architects justify architectural choices in a traceable and criteria-based manner.

Table 15. Scoring-Based Evaluation of Processing Suitability for an Early Warning Satellite Mission

Parameter	Score	Weight	Weight score
Latency Tolerance	10	3	30
Bandwidth Constraints	8	2	16
Data Volume	6	2	12
On-Board Compute Availability	7	2	14
Mission Urgency	10	3	30
Revisit Frequency	9	2	18
Regulatory Sensitivity	4	1	4

Real-Time Analytics Need	9	3	27
Total			151
Weighted Average (151/18)			8.39 (84/100)

Note: Scores shown here use the 0–10 anchors from Section 4.2.2, normalized to a 0–100 weighted average as described in Section 4.2.2.1.

Based on the high urgency of the mission, limited latency tolerance and strong real-time analytics needs, the scoring model suggests a clear preference for on-board processing. Although some trade-offs exist, such as in model complexity and data volume handling, the weighted result supports an edge-processing architecture. This example showcases how the framework enables mission-specific architectural decisions using transparent, score-driven analysis. The coastal monitoring scenario highlights the critical role of low-latency processing in time-sensitive maritime applications. On-board processing offers clear advantages in responsiveness and bandwidth efficiency, enabling rapid situational awareness and early warning capabilities. However, these benefits must be balanced against platform limitations and mission constraints. A hybrid approach-where initial detection and prioritization occur on-board, while more intensive analytics and data archiving are handled on the ground-emerges as a practical and scalable model. This ensures operational robustness while maximizing the value of EO assets in complex coastal environments.

4.4.2 Post-event environmental damage assessment scenario

Post-disaster environmental damage assessment is a recurring use case in EO missions, particularly within the framework of Copernicus EMS and civil protection agencies. These assessments occur in the aftermath of events such as floods, wildfires, earthquakes, or oil spills, with the goal of producing comprehensive spatial analyses to inform recovery operations, environmental rehabilitation and damage attribution reporting. Unlike scenarios that require near-instantaneous detection, post-event analysis emphasizes completeness, accuracy and consistency over immediacy.

In this context, ground-based processing is typically favored. High-resolution, multi-spectral or SAR data is collected from one or more satellite passes and is subjected to extensive calibration, co-registration, atmospheric correction and integration with ancillary data sources *(e.g., cadastral datasets, baseline imagery)*. These tasks are computationally intensive and benefit from cloud or HPC environments, which provide scalability and flexibility to support comparative analytics over time and space. Moreover, the ability to refine processing algorithms iteratively *(e.g., through damage classification thresholds or supervised segmentation techniques)* enhances

analytical depth-a feature that is difficult to achieve in constrained on-board environments. That said, certain pre-processing steps *(such as cloud masking or basic radiometric corrections)* can still be performed on-board to reduce the volume of downlinked data. For example, imagery flagged as obstructed or irrelevant can be discarded in orbit, ensuring that only actionable post-event scenes are prioritized for transmission. This hybrid approach conserves bandwidth and accelerates the time to analysis, especially when communication windows are constrained or multiple satellites are collecting data concurrently.

In this scenario, a constellation of EO satellites is activated in the aftermath of a severe flooding event affecting multiple European countries. The mission's objective is to provide detailed post-disaster damage maps to civil protection agencies within 72 hours, supporting recovery operations, insurance assessments and environmental impact reporting. The emphasis is on completeness, accuracy and reproducibility of geospatial products, including land cover change, infrastructure damage and flood extent analysis. Data from optical and SAR sensors must be processed and fused with cadastral and baseline imagery to produce validated map outputs suitable for legal and planning use. The objective is to deliver accurate, multi-layer environmental impact maps to civil protection authorities within 48 hours. The scenario also demonstrates how mission urgency shifts in the post-event phase. While the initial hours of an emergency may require rapid mapping *(as covered in **Section 4.4.1**)*, damage assessments typically unfold over days or even weeks, allowing time for full dataset retrieval, human-in-the-loop validation and detailed report generation. As such, latency is a secondary factor and decision-makers prioritize data quality, reproducibility and documentation. To support decision-makers, the following table summarizes how on-board and ground-based processing approaches score against key evaluation criteria for this specific scenario.

Table 16. Evaluation of Processing Approaches for Post-Event Environmental Damage Assessment

Evaluation Criteria	On-Board Processing	Ground-Based Processing
Latency	Lower latency possible for quick prioritization	Less time-critical; analysis occurs post-event
Processing Depth	Limited, suitable for pre-processing and filtering	High-depth processing with full-scene reconstruction and calibration
Bandwidth Usage	Efficient – only relevant data can be downlinked	Requires full dataset downlink, potentially bandwidth-intensive
Adaptability	Limited flexibility to update models in orbit	High – easy to adjust workflows and reprocess as needed
Use Case Fit	Useful for initial data triage	Optimal for accurate, high-quality damage mapping and reporting

Table 16 highlights the comparative strengths of ground-based processing in this context. While on-board processing offers advantages in reducing latency and initial bandwidth usage, it is

ground-based processing that excels in resolution, reproducibility and integration with broader datasets-factors that are essential for post-disaster analysis.

Table 17. Scoring-Based Evaluation of Processing Suitability for a Post-Event Damage Assessment

Parameter	Score	Weight	Weight score
Latency Tolerance	6	2	12
Bandwidth Constraints	8	2	16
Data Volume	10	3	30
On-Board Compute Availability	6	2	12
Mission Urgency	7	2	14
Revisit Frequency	5	2	10
Regulatory Sensitivity	7	2	14
Real-Time Analytics Need	5	3	15
Total			132
Weighted Average (132/18)			7.33 (73/100)

While the weighted score falls in the On-board band (73/100), the qualitative context (high data volume, need for reproducibility, moderate latency tolerance) points to ground processing as the safer and more appropriate pathway. This demonstrates that the scoring output should always be interpreted alongside the implementation gates of Chapter 6.

4.4.3 Hybrid mission profile scenario

Some EO missions require both low-latency data products and high-fidelity analysis, which cannot be effectively delivered by either on-board or ground-based processing alone. In such cases, a hybrid processing architecture-where part of the data is processed directly on-board the satellite or UAV, while the rest is transmitted for deeper analysis on the ground-offers the most effective solution. A hybrid profile is particularly useful when timely alerts are needed, but also when detailed and accurate data products must be generated later. For instance, an EO satellite might perform on-board detection of thermal anomalies *(e.g., fire hotspots)* and transmit alert signals immediately, while storing the full-resolution imagery for later download and in-depth

processing on the ground. This approach ensures both immediate situational awareness and long-term analytical value. The hybrid model is increasingly adopted in missions where bandwidth constraints, processing complexity and urgency vary across the data lifecycle. It allows system designers to tailor workflows: prioritizing edge computing for rapid decisions and delegating heavy computation *(such as data fusion or deep learning segmentation)* to cloud-based platforms. It also supports redundancy and resilience, since on-board pre-processing can serve as a fallback when communication links are degraded.

In this scenario, a hybrid EO system composed of small satellites and long-endurance UAVs monitors migration flows and cross-border movement across a remote EU border zone. The mission must detect thermal anomalies and moving objects in near-real time to trigger alerts for border patrol units, while also retaining full-resolution scene data for strategic analysis, legal documentation and cross-border coordination. On-board AI modules flag activity hotspots during flight and transmit metadata alerts, while ground-based systems perform object recognition, context fusion and historical comparison to support long-term monitoring and threat assessment. This hybrid approach balances rapid operational awareness with forensic depth.

Table 18. Evaluation of Processing Approaches for Border Surveillance using Hybrid Processing

Evaluation Criteria	On-Board Processing	Ground-Based Processing	Hybrid Processing
Latency	Real-time alerting possible from the edge	Delayed insights due to downlink and batch processing	Enables immediate alerts with follow-up analysis from ground segment
Processing Depth	Limited by hardware; good for simple object detection	High – allows for fusion, classification and contextual interpretation	Balanced – quick detection on-board, detailed analysis on the ground
Bandwidth Efficiency	High – only key metadata or detections are transmitted	Low – full data downlinked regardless of relevance	Efficient – filters data on-board, saves bandwidth while retaining full data when needed
Adaptability	Harder to update AI models or logic in orbit	Highly flexible and easy to maintain	Moderate – model updates can be pushed periodically to edge units
Use Case Fit	Suitable for rapid detection of activity	Best for strategic post-hoc intelligence analysis	Ideal – supports both rapid alerts and deeper situational awareness

As shown in **Table 18**, hybrid processing offers balanced advantages for border surveillance missions. It combines the low-latency alerting and bandwidth savings of on-board filtering with the adaptability and processing power of ground-based systems. While it introduces additional complexity, the operational flexibility and mission alignment justify this approach in use cases requiring both real-time responsiveness and analytical rigor.

Table 19. Scoring-Based Evaluation of Processing Suitability for a Border Surveillance Using Hybrid Processing

Parameter	Score	Weight	Weight score
Latency Tolerance	9	3	27
Bandwidth Constraints	9	3	27
Data Volume	6	2	12
On-Board Compute Availability	7	2	14
Mission Urgency	9	3	27
Revisit Frequency	8	2	16
Regulatory Sensitivity	5	1	5
Real-Time Analytics Need	8	3	24
Total			152
Weighted Average (152/19)			8.0 (80/100)

The weighted score (80/100) leans toward On-board. However, because the mission requires both immediate tactical alerts and archival-quality forensic products, a hybrid configuration is the more appropriate interpretation. This illustrates how scoring provides a baseline, while mission context refines the final architecture choice.

This scenario demonstrates how hybrid processing architectures can effectively address the competing demands of speed, precision and resource optimization in complex EO missions. By intelligently distributing tasks between on-board and ground segments, hybrid systems offer mission designers the ability to respond quickly in dynamic environments while still benefiting from in-depth analysis and long-term data retention. As EO technologies evolve, hybrid architectures are likely to become increasingly standard in operations that demand both agility and analytical robustness.

.

4.4.4 Mapping of Processing Pathways to THEIA Use Cases

This section provides the explicit linkage between EO data processing pathways and the operational needs of the THEIA pilot use cases. Within THEIA, three main Pilot Use Cases (PUCs) have been defined, addressing different aspects of forced population displacement. PUC1 focuses on enhanced border and maritime surveillance capabilities in response to migration flows, PUC2 addresses terrestrial surveillance of population displacement and flows, while PUC3 explores displacement driven by climate change and extreme events.

In the context of PUC1, where near real-time vessel detection and tracking are essential, on-board or hybrid processing approaches are preferred in order to reduce latency and optimize bandwidth usage. In PUC2, which involves monitoring refugee camps, border regions and inland displacement dynamics, hybrid or ground-based processing is more suitable, as these applications require multi-temporal analysis, data fusion and more complex analytical processing.

Finally, in PUC3, which focuses on crisis and environmental events such as floods or wildfires, a hybrid approach is recommended, combining rapid on-board detection with more detailed ground-based analysis. Overall, this mapping shows that adaptive processing architectures are necessary to effectively support the diverse operational requirements of THEIA.

5. Lessons from EO Mission Case Studies

Chapter 5 presents selected EO missions and satellite programs that exemplify the use of either on-board or ground-based data processing strategies. Each case study highlights key decisions, capabilities and observed outcomes related to data latency, integrity, bandwidth usage and mission responsiveness.

5.1 OPS-SAT (ESA) – On-Board Processing as a Technology Testbed

This technology demonstrator nanosatellite, led by European Space Agency (ESA), was designed to enable in-orbit constant experimentation for researchers in an accessible way. This was a response to the problem of stagnation in the development of certain space technologies caused by space companies mainly using outdated, albeit proven and reliable, technologies. From the launch of the mission in 2019 until its re-entry in 2024, 284 experiments were conducted by 134 teams [59]. This small (0.3 m x 0.1 m x 0.1 m) satellite was equipped with a processing platform that was: *"10 times more powerful than any other preceding ESA satellite"*, special communication devices, GNSS receiver, an optical camera with a resolution of 80 meters and special space-to-ground interface [60]. That has made the OPS-SAT mission a good place to test software for space technologies and has enabled further testing of various functions from different distant fields, including artificial intelligence, finance and data compression [60].

The significant limitation of OPS-SAT was its miniature design. CubeSats have several drawbacks associated with their small size, including limited power delivery and payload mass and volume. However, this reduced launch costs and did not have a significant negative impact on the performance of the on-board computer. During the mission, experiments related to real-time data processing on board were conducted and the satellite demonstrated its ability to process data on board in near real time. An example is the embedded cloud segmentation mechanism described in the article by [61]. The mission revealed the huge potential of on-board processing and low latency operations, paving the way for future nanosatellite earth observation missions. The significance of utilized solutions is clear given the evaluation criteria described in **Chapter 4**; they also show their importance in the decision pathway. This demonstrates on board processing possibilities and that nanosatellites can be used in low-latency missions under certain assumptions and conditions.

5.2 PRISMA (ASI) – Hyperspectral Imaging and Adaptive Processing

Another satellite mission is PRISMA (it. *PRecursore IperSpettrale della Missione Applicativa*, eng. *Hyperspectral Precursor of the Application Mission*). This is an Italian Earth-observation mission

operated by ASI (it. *Agenzia Spaziale Italiana*, eng. *Italian Space Agency*) launched on 22 March 2019, which focuses on acquiring hyperspectral imagery. The hyperspectral images have ~250 spectral bands (*spectral range: 400 - 2500 nm*) at 30 m spatial resolution with a 30 km swath. Additionally, PRISMA collects panchromatic imagery at 5 m spatial resolution [62]. The satellite has a PDHT (*Payload Data Handling and Transmission*) component on board, which is responsible for data management, including temporary image storage and data transmission coordination [62], [63].Thus, the processing takes place mostly in the ground segment. In Matera, Italy there is located the IDHS (*Instrument Data Handling System*) which is responsible for payload data reception and main data processing [62].

The PRISMA system has found applications in some fields, including environment monitoring. This satellite has a re-look time of ~29 days, with the possibility of reducing to less than 7 days by manoeuvring with camera angles [62]. What is more, the response time is less than 14 days, due to a data acquisition latency of less than 9.5 days and data processing latency for more processed, final products of less than 4.5 days (*from acquisition*), as reported in [64] and [65]. This allows PRISMA to be used for near-real-time solutions, albeit with some reservations and limitations. From the decision pathway point of view, PRISMA demonstrates that the hybrid approach to data processing can be a reasonable choice in specific cases. However, all aspects of the system and its projected applications must be considered during mission preparation process.

5.3 Sentinel-1 (ESA) – Ground-Based Processing at Scale

One of the most popular satellite missions is the Sentinel missions. The Sentinel-1 mission is an important part of the Copernicus Earth Observation program and focuses on delivering SAR (*Synthetic Aperture Radar*) satellite imagery for numerous purposes. This technology enables observation even on cloudy days and at night. Since 2014, three satellites (*Sentinel-1 A, B and C*) were launched, of which two (A and C) remain active. The characteristics of their polar, sun-synchronous orbit allow them to cover almost entire globe with satellite imagery. What is more, acquired data have a relatively short revisit period of 2-6 days (for constellation of two satellites) and high ground resolution of 5 meters in *Strip Map Mode* [66].

The *Copernicus Ground Segment* plays a crucial role in Sentinel-1 data processing, as the data collected by the space segment is sent to ground stations (the X-band downlink provides a data rate of 520 Mbps [67] and then properly processed. It should be mentioned that derived raw SAR data undergo processing to generate various products, including SLC (*Single Look Complex*) data (e.g. for interferometric applications), GRD (*Ground Range Detected*), special geophysical products etc. [66]. Finally, once these processes are completed, the data are publicly released free of charge via the Internet. This presented approach has several benefits. By moving complex algorithms and powerful computing resources to ground segment, analysts can apply advanced

techniques that would be impractical on board the satellite. Importantly, Sentinel-1 products are offered in three timeline levels: NRT (*Near Real-Time*), STC (*Short Time-Critical*), NTC (*Non Time-Critical*) and at various processing levels. Although Sentinel-1 is not on-board based processing system, the products are delivered within 10 minutes (regional - within the station coverage, NRT, processing level 0) to ca. 24 hours [67]

Summarizing, the temporal resolution sufficient for many applications and the long mission duration of over 10 years make Sentinel-1 mission data useful for many purposes and play a vital role in various EO projects. Thus, while Sentinel-1 delivers frequent, high-resolution data worldwide, the quick-access mode is limited to the coverage of the ground stations and provides only level-0 products [67]. Compared to other satellites, the Sentinel-1 mission is a very good source of data which can be used especially in long-term and complex EO projects and can be applicated for numerous fields, including geophysics, agriculture, oil pollution monitoring, emergency management etc. All in all, although Sentinel-1 bases on on-ground processing, it allows comprehensive analyses in near-real time. Processed and raw data products are offered. However, an important limitation is resolution of imagery which sometimes can be insufficient. Furthermore, the mass of the satellite makes it more expensive to launch into orbit than nanosatellites. In the end, the downlink data rate, short latency and short revisit time are the advantages of this system.

5.4 ECOSTRESS (NASA) – Thermal Data and Downlink Bottlenecks

The ECOSTRESS (*Ecosystem Spaceborne Thermal Radiometer Experiment on Space Station*) mission, led by NASA, aims to monitor evapotranspiration. This mission differs from the other missions described in this chapter, as it does not use own satellite platform and its aperture is set on the International Space Station. Placing a radiometer on the International Space Station was a reasonable decision because the existing ISS infrastructure could be used. The ECOSTRESS radiometer has been on the ISS since 2018 and from this time it has collected data which is applicable in drought and wildfire monitoring [68]. It is worth noting that the ISS follows an irregular orbit and the revisit time varies from less than 1 day to 5 days [68]. ECOSTRESS captures the multispectral TIR imagery at medium resolution, with pixel dimensions at nadir 38 x 69 meters; however, the images are resampled to 70-meter resolution [68].Space component allows to transmit the data to stations on ground. The data is distributed free of charge and allow analyses in near-real time. That makes ECOSTRESS imagery very useful in numerous fields, particularly when compared to other satellite data, as MODIS has lower spatial resolution and Landsat has lower temporal resolution [68]. What is more, ECOSTRESS has important advantage, as it allows to monitor diurnal changes, for example ground temperature in cities [69].

5.5 HYPSO-1 (Norway) – Maritime Observation

A relevant example of satellite EO mission, which significantly utilizes on-board processing is HYPSO-1. This Norwegian 6U CubeSat acquires hyperspectral imagery since 2022. The mission is developed by the Norwegian University of Science and Technology (NTNU) and its role is to deliver images useful in ocean monitoring. Images are delivered in near real time, which was an important objective of this system – it was shown that raw data latency can be lower than 4 hours [70]. It should be noted that HYPSO-1 uses hybrid processing, as special algorithms are utilized on board to reduce the volume of data and cope with data transmission limitations [72]. The CCSDS-123 compression algorithm was implemented into FPGA chips in HYSPO-1, what decreased compression "*from 5-6 minutes to less than a second"* [71]. *"The FPGA-based image processing algorithms on the OPU are key to enable faster download and distribution of data while at the same time relieving HYPSO-1's power budget"*. The satellite uses S-band transceiver for data downlink with a data rate of 1 Mbps, which is relatively low therefore, on-board processing can be crucial for system's proper operation.

It is worth noting that this satellite system uses multiple innovative solutions. One of them is that surveys conducted with HYPSO-1 can be supported by apertures on the ground, such as UAVs and autonomous underwater vehicles (AUVs) that collect data in situ. As [70] notes: "*This assimilation of different sensor systems should give a sounder insight into the oceanographic area of interest than the sum of their contributions".* What is more, as previously mentioned, system utilizes on-board processing – this mission demonstrates the importance of it in the decision pathway. Processing on board helps the HYPSO-1 system quickly deliver the data, what can be very useful in missions which require quick image access.

5.6 Summary

Table 20. Satellite systems specification (see [61]- [72])

Parameter	OPS-SAT	PRISMA	Sentinel-1	ECOSTRESS	HYPSO-1
Launch	2019	2019	2014 (-1A) 2016 (-1B) 2024 (-1C) 2025 *plan* (-1D)	2018	2022
Status	Finished (2024)	Active	Active (-1A, -1C)	Active	Active
Earth Observation Instrument	Optical	Optical	SAR	Radiometer (vis/IR)	Optical
Spatial Resolution	~80 m	5m (PAN) 30 m (HYC)	Up to 5 m (*Stripmap*)	69 m x 38 m	~60 - 70 m (close to nadir)

					~400 – 500 m (off-nadir)
Temporal Resolution	N/A	≤ 7 days	≤ 6 days (constellation)	≤ 5 days	≤ 3 days (achievable in selected scenarios)
Latency	N/A	≤ ~4.5 days (depends on location and configuration)	≥ 10 min	~4 hours (radiance), ~1 day (surface temp.), ~3-5 days (evaporation)	4 hours best (raw data for special scenarios)
Bandwidth Efficiency	50 Mbps	155 Mbps	520 Mbps	4.5 Mbps	1 Mbps
Suggested Architecture	On-board demonstrator (edge experiments, CubeSat 3U)	Ground-based (traditional chain, PRISMA HYC/PAN)	Ground-based (bulk SAR processing, HPC centers)	Ground-based (science-driven, limited pre-processing)	Hybrid (on-board reduction + ground post-processing)

This indicates that some of these satellites can potentially be used both for long-term analyses of archived data and for quick decision-making based on near-real time data. However, all of them demonstrate that modern space technologies enable researchers to monitor various phenomena worldwide, offering new possibilities that are competitive to traditional earth observation methods. EO mission case studies discussed in this chapter point out different approaches to data processing. Systems like Landsat-1 and PRISMA (ASI) base on ground-based processing. The strong advantage of this approach is that it allows to utilize the computational platforms that are more powerful than computers which could be mounted on-board. However, on-board processing offers new perspectives for the use of satellite systems. This was demonstrated by the OPS-SAT mission, during which numerous experiments were conducted on-board and the HYPSO-1 mission, where the specification of a system required a reduction in the amount of data due to the low bandwidth of the downlink.

All in all, the satellite-system technology is evolving and offers new possibilities, such as on-board processing. The trend toward miniaturization of satellite platforms is also evident, as demonstrated by the growing use of CubeSats - good examples include OPS-SAT (3U) and HYPSO-1 (6U). This technology enables cost reduction for a satellite mission, including launch costs, but it also comes with trade-offs such as power-consumption limits etc. However, advances in earth observation systems allow the use of satellite imagery from multiple sources, enabling accurate analyses.

6. Strategic Recommendations and Considerations

This chapter turns the criteria and trade-offs developed in **Chapter 4** and the scenarios outlined in **Chapter 5** into clear, testable rules. The recommendations aim to support both current system deployments and the design of future hybrid architectures that balance real-time responsiveness with analytical depth. Moreover, this section reflects on broader technological and operational considerations, including the implications of on-board AI, evolving ground infrastructure capabilities and the role of mission-specific customization in maximizing system value.

6.1 Recommendations for EO Mission Architects

EO mission architects today operate in a landscape marked by growing complexity: diverse sensor payloads, tighter latency requirements, evolving user needs and increasingly hybrid mission objectives. Based on the analysis presented in this guide, it is clear that processing pathway decisions must be considered as a core design dimension-integrated from the earliest stages of mission planning rather than treated as secondary to platform design or data product definition.

A key recommendation is to **ensure that the processing strategy is aligned with the mission's urgency profile**. For time-critical applications such as disaster response, illegal activity detection, or real-time environmental monitoring, low-latency architectures are essential. On-board data processing-particularly when augmented with AI/ML capabilities-can support near-instantaneous alerting and triaging of high-priority events. This is especially important when ground access is delayed or limited by bandwidth. At the same time, the limitations of both purely on-board and purely ground-based models suggest **that hybrid architectures should be treated not as a special case, but as the default consideration**. By distributing tasks intelligently between on-board filtering, compression, or initial detection and ground-based analysis for high-fidelity processing and long-term storage, mission flexibility is increased without compromising on data quality. This dual-pathway strategy allows responsiveness and analytical depth to coexist. Architects must also consider **long-term adaptability**. On-board systems are inherently less flexible once deployed, so the capacity to update AI models or processing pipelines-whether through reprogrammable logic, modular firmware, or containerized deployments-can greatly extend mission relevance over time. Similarly, decisions around data downlinking should account for the relative value of the data versus the bandwidth and energy costs involved. Intelligent prioritization schemes, such as relevance-based filtering or confidence-thresholded alerts, can reduce unnecessary transmission and help preserve limited resources. Importantly, **risk tolerance** should inform the final architectural choice. As explored in **Chapter 4**, each processing strategy comes with its own risk profile-ranging from missed detections to hardware failure or degraded model accuracy in changing conditions. Processing pathways must be selected with these operational risks in mind, particularly in missions where reliability and accuracy are critical

to success. Finally, architects are encouraged to **look beyond data delivery and focus** on insight delivery. EO missions are ultimately decision-support systems. Prioritizing architectures that enable faster, clearer and more targeted outputs-especially those that integrate with existing alerting, visualization, or decision-making tools-will ensure that the system delivers operational value, not just raw imagery.

To make these recommendations actionable, we translate the scoring model of **Chapter 4** into explicit criteria and thresholds that can guide architectural decisions. The weighted average of the eight parameters (*Latency Tolerance, Bandwidth, Data Volume, On-board Compute Availability, Mission Urgency, Revisit Frequency, Regulatory Sensitivity, Real-time Analytics*) remains the formal decision output, producing a 0–100 score that classifies missions as Ground (0–33), Hybrid (34–66), or On-board (67–100). In addition to this overall score, we recommend two diagnostic indicators be monitored explicitly throughout the design cycle:

- **Downlink Bottleneck Index (DBI)**: ratio of expected daily data volume to expected daily downlink capacity. A DBI ≥ 2 indicates a structurally bandwidth-limited mission.
- **Latency Tolerance**: the maximum acceptable delay for delivering mission outputs. A Latency Tolerance ≤ 1 corresponds to near-real-time needs (*e.g., < 5 minutes*).

These indicators do not replace the full scoring algorithm but provide "early warning" conditions. In practice, if DBI is consistently high and LT is consistently low, then even before a full scoring exercise, the mission will tend to classify toward On-board or Hybrid solutions. Conversely, a low DBI combined with high LT usually biases the weighted score toward Ground solutions.

On-board processing is strongly favored when the combined scoring framework (Section 4.2.2) yields values ≥ 67 and the diagnostic checks confirm DBI ≥ 2 with LT ≤ 1. In this situation, pre-filtering, compression, or inference on the platform prevents downlink congestion, cuts mean time-to-alert and allows the ground segment to focus on the most relevant data. OPS-SAT and Φ-Sat-1 demonstrated the value of this approach, with on-board AI reducing data volume by more than 30% and delivering tactical alerts significantly faster than ground-only chains.

Ground-based processing remains the better option when the weighted score falls ≤ 33, or when the diagnostic checks confirm LT ≥ 3 (≥ 24 h acceptable delay) or on-board compute maturity is below TRL 6. Here, the advantages of flexible algorithms, archive correlation and reproducibility dominate, as illustrated by Sentinel-1's SAR interferometry products.

Hybrid processing corresponds to intermediate cases (scores between 34–66), especially when LT falls in the rapid (≤ 15 min) range and DBI is between 1–3. ECOSTRESS, for example, generated rapid fire-detection alerts on-board while still downlinking complete spectral datasets for later climate analysis. Hybrid architectures amortize risk: if on-board models underperform, the ground chain still guarantees end-product quality.

Before finalizing any decision, architects should apply three implementation gates. These do not replace the weighted scoring framework of **Chapter 4**, but act as safeguards to ensure that the chosen architecture remains robust and compliant:

- ✓ For CSS and similar workflows, products must be auditable and reproducible. If on-board cannot provide provenance logs and signed outputs, bias toward ground or hybrid.
- ✓ Require a secure, documented over-the-air (OTA) update and rollback mechanism for models and pipelines. Without it, purely on-board is unsafe and should be downgraded.
- ✓ Recompute DBI using realistic contact schedules and link rates. If actual DBI < 1 after scheduling, bandwidth is not a true constraint, reducing the case for heavy on-board processing.

On a final note, mission urgency profiles (*e.g., immediate alert, T+24 h reporting, T+7 d analysis*) should be explicitly linked to pathway selection rather than assumed late in the design. On-board heavy designs trade flexibility for responsiveness; ground heavy designs do the opposite; hybrids deliberately combine both. Whichever pathway is chosen, model governance, provenance and graceful degradation strategies must be integrated from the outset. Only then can EO missions evolve from data delivery pipelines to decision-support systems that maximize operational value.

6.2 Investment Priorities for On-Board Processing

As sensor resolution increases and data volumes multiply, transmitting all raw data to the ground becomes increasingly impractical - especially in missions constrained by limited downlink bandwidth, intermittent connectivity, or latency-sensitive use cases. To address these constraints, targeted investments in on-board data processing offer a transformative opportunity to pre-filter, analyze and prioritize EO data before it ever leaves the satellite or aerial platform.

One of the primary areas of investment should focus on edge computing hardware that balances high performance with strict limitations in size, weight and power *(SWaP)*. Devices such as radiation-hardened FPGAs, low-power GPUs *(e.g., Jetson-class modules)* and dedicated AI accelerators are increasingly capable of running complex inference models in orbit. These platforms allow for real-time object detection, anomaly identification and adaptive data prioritization, directly supporting mission-critical operations. **Equally important is the development of modular software frameworks that can be updated post-deployment**. This includes on-board AI pipelines that can be retrained or fine-tuned via uplinks, enabling satellites to evolve in response to changing mission needs. Such adaptability is essential for long-duration missions or those operating in dynamic environments, where static models may quickly become obsolete. **Another key investment area lies in telemetry-aware optimization**, where processing routines are tuned to data transmission schedules, bandwidth availability and satellite orbital dynamics. Systems capable of adjusting their processing intensity or data selection logic based

on current transmission opportunities will ensure better utilization of ground station windows and minimize unnecessary data retention. Finally, **investment should also target on-board diagnostics and decision-support features**, which can flag anomalies in sensor performance, identify corrupted data segments, or trigger autonomous mission reconfiguration. These capabilities are especially valuable in missions where ground-based operators cannot intervene in real time, such as deep-space EO or fast-moving constellations with low revisit times.

To guide future development efforts and investment strategies, it is essential to identify the most pressing capability gaps that currently hinder the broader deployment of intelligent on-board processing solutions. The following table summarizes a targeted gap analysis across key technical areas. Each row identifies a current limitation, outlines the direction for improvement and assigns a qualitative priority level based on its expected mission impact.

Table 21. Gap Analysis of Current On-Board Processing Capabilities

Capability Area	Current Limitation	Needed Improvement / Future Focus	Priority
AI model updateability	Static models, limited reprogramming	Modular firmware or in-orbit re-trainable AI pipelines	High
Compute performance	SWaP-constrained AI inference only	Support for multitask, multi-sensor parallel processing	High
Storage & buffering	Limited solid-state buffer, no long-term storage	Dynamic caching and smart data discard policies	Medium
Power efficiency	High power draw for on-board GPUs	Ultra-low power AI accelerators with adaptive logic	High
Software stack	Vendor-locked or non-interoperable tools	Open, reconfigurable, containerized pipelines	Medium
Security & integrity	Basic checks, vulnerable firmware layers	Secure boot, anomaly detection, encrypted model updates	High
Provenance & auditability	Minimal on-board lineage	On-board product lineage (W3C PROV-like) and signed artifacts	High

These findings highlight the areas where on-board processing currently faces structural or functional bottlenecks. The highest-priority gaps are those that directly impact real-time operations, mission flexibility, or security. Notably, the lack of model updateability constrains long-duration missions in changing environments, while compute limitations restrict the use of richer inference algorithms. Power efficiency is also a persistent barrier, particularly in small platforms such as CubeSats or UAS, where battery life is critical. Storage strategies and software modularity are recognized as medium-priority gaps, as they tend to impact usability and

performance scalability rather than fundamental mission success. Security, however, is an essential baseline for all EO missions. Any future on-board processing capability must embed cryptographic safeguards, secure communication protocols and autonomous fault-detection mechanisms to ensure trustworthiness and resilience. Addressing these hardware and software gaps is not merely a matter of optimization-it is a prerequisite for the next generation of autonomous, adaptive EO systems. Investments that close these gaps will directly improve mission responsiveness, reduce data overload and enable EO platforms to deliver value closer to the point of data collection. Future on-board architectures must therefore be designed with flexibility, security and upgradability at their core.

A central R&D focus should be the **development of re-trainable on-board AI systems**. While current models are typically static and trained pre-launch, future missions would benefit from AI architectures that can be updated in-orbit, either autonomously or via secure uplink. This requires not only modular machine learning pipelines, but also new methods for on-board validation, fallback handling and fail-safe switching - all of which are under-explored areas in space systems engineering. In parallel, **energy-efficient AI computing remains a critical need**. Research should prioritize the co-design of hardware and software to reduce power consumption while preserving inference accuracy. This includes lightweight neural network architectures (e.g., spiking neural networks or quantized models) and hardware-aware optimization techniques that match algorithm demands to the capabilities of radiation-hardened FPGAs or low-power GPUs.

Another promising direction is the **integration of adaptive task orchestration engines**, enabling satellites to dynamically allocate computing tasks based on mission priorities, sensor input, or bandwidth availability. Such orchestration systems could allow the same hardware to serve multiple processing roles, adjusting in real time to operational needs - for example, shifting from image compression to event detection during a crisis scenario. Furthermore, R&D should explore in-orbit collaboration between assets, such as federated learning or distributed AI processing across satellite constellations. These techniques would allow missions to benefit from pooled data and compute resources without relying solely on ground-based coordination. This vision aligns with the trend toward smarter constellations and real-time coordination across platforms. Finally, research must **address robustness and explainability in AI models** deployed on-board. As on-orbit AI becomes more autonomous, ensuring that outputs are verifiable, traceable and trustworthy becomes not just a technical concern, but a regulatory one. Techniques such as model introspection, anomaly detection and interpretable machine learning will be key enablers for compliance with EU digital trust and data governance frameworks.

6.3 Roadmap for Hybrid Processing Models

Hybrid data processing architectures, which combine on-board and ground-based capabilities, have emerged as a flexible and scalable solution to address the diverse needs of modern EO missions. As demonstrated in **Chapter 4**, hybrid models offer a balanced response to the trade-offs between latency, processing depth and bandwidth limitations. However, fully deploying hybrid systems at scale requires a structured, phased approach - one that aligns technological maturity, system integration and operational workflows over time. This roadmap outlines a strategic progression for adopting hybrid processing, highlighting actions and priorities across short-, medium- and long-term horizons.

Thresholds used in this roadmap: "near-real-time" = ≤ 5 min, "rapid" = ≤ 15 min, "delayed" = ≥ 24 h; DBI bands: low < 1, medium 1–3, high ≥ 3.

1. Short-Term Phase *(0-3 Years)*: Foundations and Pilot Integration

In the initial phase, the priority is to enable basic hybrid workflows using currently available technologies. Missions can begin by embedding lightweight AI or rule-based filtering on-board, allowing only the most relevant data segments to be downlinked for deeper ground-based analysis. This stage does not require full autonomy, but rather focused task delegation - such as cloud masking, thermal anomaly detection, or confidence-scored object flagging in orbit. Namely, the key short-term actions include:

- Integrating existing edge-AI accelerators *(e.g., NVIDIA Jetson, Intel Movidius)* into EO satellite or UAV platforms
- Deploying containerized processing modules that allow separation of mission-critical tasks from experiment-specific models.
- Developing mission-specific processing policies *(e.g., when to switch from full image downlink to cropped segment transmission)*.
- Validating hybrid data pipelines in small-scale demonstrations, such as pilot missions or tech demonstrators like Φ-Sat-1 or Intuition-1.

This phase also requires refining the ground segment to handle variable input volumes, metadata-enriched alerts and mixed-resolution data streams. While full coordination between on-board and ground systems may be limited at this stage, the goal is to prove value and identify bottlenecks in operational conditions.

2. Medium-Term Phase *(3-6 Years)*: Expansion, Interoperability and Standardization

As initial hybrid workflows mature and pilot integrations demonstrate value**, the medium-term focus should shift toward scaling hybrid processing across mission classes, sensor types and operating environments**. This phase emphasizes multi-mission adaptability, interoperability and the development of shared frameworks that support more autonomous task coordination between on-board and ground segments.

One of the main priorities is the implementation of adaptive processing pipelines, where the satellite or aerial platform can modify its behavior based on mission context, environmental triggers, or real-time input from the ground. For example, a wildfire monitoring constellation might autonomously shift from normal scanning mode to high-frequency regional tracking when thermal anomalies are detected. At the same time, model and algorithm update pipelines must be formalized. This includes secure mechanisms for uplinking new AI models or inference parameters, as well as verification protocols to ensure reliability post-deployment. The ability to refresh on-board processing logic dynamically-while maintaining data integrity and mission safety-will be a defining capability for hybrid platforms in this timeframe. Ground segment systems should also evolve to support multi-resolution fusion, enabling the combination of quick-look on-board summaries with full-resolution ground-based reanalysis. This will enhance analytics pipelines and provide mission operators with both rapid alerts and deep situational insight in a unified environment. Another key element of this phase is the standardization of metadata formats, tasking interfaces and hybrid architecture APIs. Without a common language between edge and ground systems, hybrid workflows risk fragmentation and inefficiency. EU-led coordination in these areas can accelerate the maturity of hybrid processing and ensure compatibility across Copernicus and commercial platforms.

3. Long-Term Phase (6+ Years): Autonomous Orchestration and Scalable Intelligence

In the long term, hybrid EO processing is expected to evolve beyond static division of tasks into a fully orchestrated, adaptive architecture, where satellites, aerial assets and ground systems act as coordinated nodes in a real-time sensing and analytics network. The goal of this phase is not only to optimize data flows and processing efficiency, but to enable dynamic decision-making at the system-of-systems level, supporting autonomous mission reconfiguration based on evolving conditions. A key capability in this phase will be the integration of on-board learning or federated learning mechanisms, where AI models can improve over time based on in-situ experience or shared insights from peer platforms. This would allow distributed intelligence across constellations or fleets, reducing reliance on central ground control and accelerating mission responsiveness. Hybrid EO systems will also be expected to support context-aware orchestration-deciding in real time whether a task should be executed on-board, relayed to the

ground, or distributed across platforms based on priority, resource availability, or user-defined thresholds.

From an infrastructure perspective, hybrid systems will likely become natively interoperable with digital twin environments, enabling continuous simulation, monitoring and predictive analysis. Real-world data collected and processed at the edge will be looped into digital replicas of observed environments, feeding not only analytics but also policy and scenario planning engines. Ultimately, long-term success depends on sustained investment in modular, standards-based design, cybersecurity by design and tight integration between EO platforms and the broader EU Space and Digital Strategies. By this stage, hybrid processing will not be an add-on - it will be the default architecture for scalable, autonomous and user-centric EO missions.

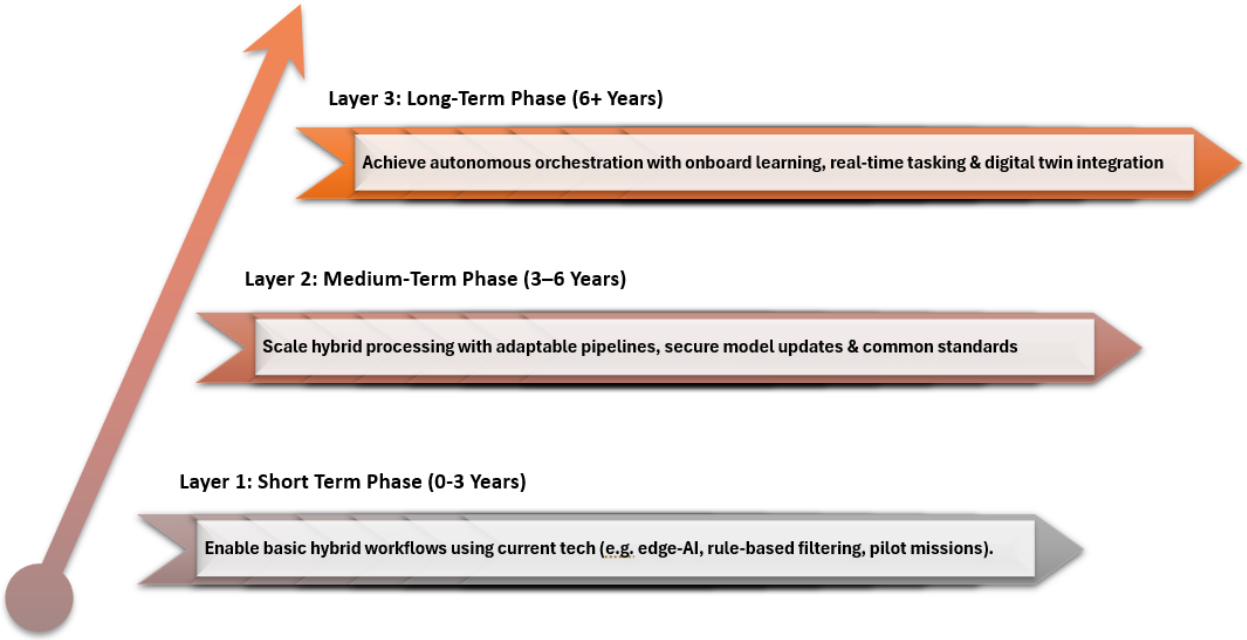


Figure 7. Roadmap for Hybrid Processing Models

As shown in the table below, **hybrid processing is not a static configuration but an evolving framework that matures over time**. The short-term phase focuses on integrating lightweight on-board capabilities to address urgent latency and bandwidth challenges. The medium-term expands this foundation by enabling model updates, multi-mission adaptability and standardized interfaces between edge and ground components. Finally, the long-term phase envisions a fully autonomous, orchestrated environment where EO assets operate collaboratively, learn from operational feedback and integrate seamlessly with digital twin ecosystems.

6.4 Interoperability and Risk Considerations

While THEIA's dual-pathway approach enables flexibility in EO data processing design, the long-term strategic risks associated with system architecture choices must be explicitly considered-particularly in the context of European sovereignty, technological independence and policy alignment. EO mission planners, infrastructure investors and institutional decision-makers must evaluate not only short-term technical performance, but also systemic exposure to external dependencies and geopolitical uncertainties.

A key area of concern lies in the reliance on proprietary hardware platforms and non-EU cloud infrastructures for critical processing workloads. Many current EO applications depend heavily on U.S.-based GPU and FPGA vendors and on hyperscale cloud providers operating under jurisdictions with differing data protection regimes. This introduces risks related to supply chain fragility, vendor lock-in and extraterritorial data access (*e.g., under legislation like the U.S. CLOUD Act*). In addition, long-term sustainability risks arise from deploying rigid or closed systems that may not support future upgrades, reusability, or migration to open standards. Processing chains that are not modular or easily portable may create maintenance burdens and limit future innovation. This is particularly relevant when EO platforms are deployed for 5–10 year mission lifespans, during which both hardware and algorithmic landscapes evolve rapidly.

Table 22. Strategic Risk Considerations for EO Processing Architectures

Risk Category	Description	Potential Impact	Suggested Mitigation
Dependence on Non-EU Cloud Infrastructure	Use of hyperscale cloud providers outside EU jurisdiction may expose mission data to extraterritorial control.	Loss of data sovereignty; reduced control over access and compliance.	Prioritize EU-based or sovereign cloud alternatives; use encryption and delay-tolerant networking (DTN).
Reliance on Proprietary Chips (Non-EU vendors)	Dependence on U.S. GPUs/ASICs increases exposure to export restrictions and limited long-term availability.	Mission delays, cost overruns, or upgrade bottlenecks.	Invest in EU chip R&D; use hardware-agnostic pipelines and open architectures.
Lack of Modularity / Upgrade Path	Closed or rigid systems hinder future reusability, scalability, or algorithm portability.	Obsolescence; limited adaptability for evolving mission needs.	Adopt modular, open standards and containerized workflows.
Regulatory and Legal Misalignment	Processing platforms governed by non-EU laws may conflict with EU data protection (e.g., GDPR).	Full autonomy, system-of-systems coordination, predictive analytics	Adopt modular, open standards and containerized workflows.
Supply Chain Fragility	Global crises or trade restrictions may disrupt access to critical EO hardware (e.g., semiconductors, sensors).	Legal exposure; compromised compliance with EU sovereignty and privacy regulations.	Use EU-compliant platforms; establish clear contractual terms for data governance.

To ensure long-term integration with Copernicus components and other EU-led infrastructures, all architectural and processing pathway decisions must be rooted in openness, compatibility and standardization. A key aspect of this i lies in the use of standardized data formats. THEIA's processing architecture supports widely adopted geospatial formats such as GeoTIFF and NetCDF, enabling seamless ingestion and dissemination across EO platforms. Moreover, the integration of emerging metadata standards like STAC (SpatioTemporal Asset Catalog) enhances discoverability and cataloging across distributed archives, increasing the reusability of data products. Interoperability is also ensured through the use of standardized APIs and access protocols, particularly those defined by the Open Geospatial Consortium (OGC), such as WMS (Web Map Service), WFS (Web Feature Service) and WCS (Web Coverage Service).

Table 22. Key interoperability elements supporting THEIA's integration with Copernicus services and EU EO infrastructure.

Interoperability Aspect	Standard / Method	Purpose / Benefit
Data Formats	GeoTIFF, NetCDF, STAC	Standardized geospatial data handling and metadata compatibility
APIs & Service Access	OGC WMS, WFS, WCS	Uniform access to EO layers and products; integration with Copernicus tools
Metadata & Cataloging Standards	STAC, ISO 19115	Discoverability and archiving across distributed systems
Regulatory and Legal Misalignment	Format and protocol fallback support	Enables operation with existing Copernicus ground segment infrastructure

The table above reveals the core interoperability considerations relevant to THEIA's system architecture and processing choices. Each row highlights a critical domain-such as data standardization, API-level integration, or legacy system compatibility-where alignment with established norms ensures technical coherence with Copernicus and other EU infrastructure.

7. Conclusions

This guide tackled a central architectural choice in modern EO: where to process- on-board, on the ground, or in a hybrid split. With data volumes rising and response windows shrinking, the choice directly affects latency, bandwidth, cost, auditability and EU sovereignty. We translated the problem into operational criteria so mission teams can make reproducible, defensible decisions. Through a structured, comparative analysis of on-board and ground-based processing pathways, D5.4 has explored the trade-offs, performance parameters and contextual variables that shape this decision. The proposed evaluation framework integrates technical, operational and strategic dimensions, enabling mission planners to make informed, tailored choices. The inclusion of practical use cases and risk-based scoring models further strengthens the applicability of the guide in real-world design scenarios. Ultimately, D5.4 serves as both a practical reference and a strategic tool-supporting system architects, funding agencies and EU policy stakeholders as they navigate an increasingly hybrid EO ecosystem. With the emergence of AI-enabled platforms, sovereign hardware requirements and Copernicus integration goals, **selecting the appropriate processing pathway is not just a technical choice-it is a strategic imperative**.

The findings of this guide can be condensed into a set of operational rules that complement the weighted scoring framework presented in **Chapter 4**. **On-board processing demonstrates the greatest value when missions are simultaneously bandwidth-limited and time-critical**. This situation is typically characterized by a DBI of two or higher - meaning that the expected daily data volume exceeds the available downlink capacity by at least a factor of two - and a Latency Tolerance of one or lower, indicating the need for near-real-time alerts (≈ five minutes). Under such conditions, pre-filtering or inference in orbit, with buffering of the full-resolution data for later transfer, ensures that urgent insights are delivered without overwhelming ground infrastructure. **Ground-based processing continues to be essential when latency requirements are relaxed or when the maturity of on-board technology is insufficient**. Missions that can tolerate delivery times of 24 hours or more (LT ≥ 3) with on-board hardware or software below TRL 6 (*not yet flight-proven*) should downgrade on-board options and favor hybrid or ground-only solutions. This approach maximizes flexibility, reproducibility and scalability without burdening the spacecraft with functions better handled on Earth. **Hybrid architectures emerge as the default solution when mission needs are mixed**. These designs are particularly relevant when tactical cues must be generated in under fifteen minutes while full-resolution products are retained for science or forensic analysis, or when DBI values fall between one and three. Hybrid processing allows lightweight detection and prioritization to occur on-board while deep analysis and archival are conducted on the ground, thereby balancing responsiveness with depth.

The framework was illustrated through multiple example scenarios-including real-time coastal monitoring, post-event environmental damage assessment and hybrid mission profiles-demonstrating its adaptability across diverse use cases. In each case, the scoring and prioritization of architectural features enabled mission planners to assess the suitability of on-board, ground-based, or hybrid configurations based on operational constraints and user needs. **Importantly, the framework is designed to be flexible and extendable**. It can accommodate evolving technologies, emerging mission types, or policy changes, making it a living tool for EO architecture planning. Embedding this decision-making process early in the design cycle allows missions to reduce risk, optimize performance and ensure long-term compatibility with EU infrastructure and strategic objectives.

Looking forward, the evolution of EO missions will continue to demand architectures that balance responsiveness, analytical depth and sovereignty considerations. **The analysis in this guide confirms that hybrid configurations will often be the most effective, combining on-board responsiveness with the depth and auditability of ground-based infrastructures. However, the path chosen for any mission must remain rooted in quantitative evaluation, risk assessment and strategic alignment.** As a practical call to action, pathway evaluation should be formalized at multiple points in the design lifecycle. At Mission Concept Review, teams should compute the DBI, set explicit LT targets and assess the maturity of on-board compute platforms. At Preliminary Design Review, these metrics should be updated with refined link budgets and weighted criteria and the scoring framework should be executed to document the rationale for the chosen pathway. At Critical Design Review, the analysis should be repeated with validated flight hardware and software evidence. The results of all three runs should be archived in the design history file, providing a clear record of how architectural choices were made and ensuring traceability for both technical audits and policy oversight.

References

[1] European Space Agency. (n.d.). What is Earth Observation? Retrieved from https://www.euspa.europa.eu/eu-space-programme/copernicus/what-earth-observation

[2] https://en.wikipedia.org/wiki/Remote_sensing

[3] NASA Earthdata. (n.d.). Data Processing Levels. Retrieved from https://www.earthdata.nasa.gov/learn/earth-observation-data-basics/data-processing-levels

[4] CEOS-WGISS Data Stewardship Interest Group. (2023). Preservation workflow for Earth Observation space data (Version 1.1). Committee on Earth Observation Satellites. Retrieved from http://www.ceos.org/ourwork/workinggroups/wgiss/interest-groups/data-stewardship/

[5] https://en.wikipedia.org/wiki/Copernicus_Programme

[6] Garg, A., Patil, A., Sarkar, M., Moorthi, S. M., & Dhar, D. (2024). Advancements in data processing and calibration for the Hyperspectral Imaging Satellite (HySIS). arXiv. https://doi.org/10.48550/arXiv.2411.08917

[7] Kawade, P. (2023, October 31). The implications of increasing Earth observation data. NSR - Northern Sky Research. https://www.nsr.com/the-implications-of-increasing-earth-observation-data/

[8] https://www.earthdata.nasa.gov/learn/earth-observation-data-basics/sar

[9] Wang, Y., & Zhu, X. X. (2017). Earth observation using SAR and social media images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (pp. 94–103). https://openaccess.thecvf.com/content_cvpr_2017_workshops/w18/papers/Wang_Earth_Observation_Using_CVPR_2017_paper.pdf

[10] https://www.cailabs.com/aerospace-defense/laser-communications/earth-observation

[11] KP Labs. (2025, January 28). What is on-board data processing? https://www.kplabs.space/news/what-is-on-board-data-processing

[12] https://en.wikipedia.org/wiki/Phi-Sat-1

[13] https://en.wikipedia.org/wiki/Intuition-1

[14] https://en.wikipedia.org/wiki/Alaska_Satellite_Facility

[15] NASA. (n.d.). 11.0 Ground Data Systems and Mission Operations. Retrieved from https://www.nasa.gov/smallsat-institute/sst-soa/ground-data-systems-and-mission-operations/

[16] Furano, G., & Menicucci, A. (2018). Roadmap for on-board processing and data handling systems in space. In S. Pontarelli (Ed.), Dependable Multicore Architectures at Nanoscale (pp. 253–281). Springer.

[17] Valsesia, D., & Magli, E. (2019). High-throughput on-board hyperspectral image compression with ground-based CNN reconstruction. IEEE Transactions on Geoscience and Remote Sensing, 57(12), 9544–9553.

[18] Furano, G., Menicucci, A., & Mottola, L. (2020). Towards the use of artificial intelligence on the edge in space systems: Challenges and opportunities. IEEE Aerospace and Electronic Systems Magazine, 35(12), 44–56.

[19] Cucchetti, T., Fressin, F., Mathé, M., & Mailhes, C. (2021). On-board/on-ground image processing chain for high-resolution Earth observation satellites. International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLIII-B3, 755–762.

[20] Caon, G., Pessina, A., Caltagirone, F., & Dell'Erba, D. (2021). Very low latency architecture for Earth observation satellite on-board data handling, compression and encryption. In Proceedings of IGARSS 2021 - IEEE International Geoscience and Remote Sensing Symposium.

[21] Vercruyssen, V., Innocenti, L., Vanneschi, M., & Krutz, D. (2021). The Φ-Sat-1 mission: The first on-board deep neural network demonstrator for satellite Earth observation. IEEE Transactions on Geoscience and Remote Sensing, 60, 1–14.

[22] Chintalapati, S., Dutt, S., & Ramamoorthy, A. (2024). Opportunities and challenges of on-board AI-based image recognition for small satellite Earth observation missions. Advances in Space Research, 73(4), 1752–1768.

[23] Parra, L., González, R., Mateos, J., & Masini, A. (2024). Advancements in on-board processing of synthetic aperture radar (SAR) data: Enhancing efficiency and real-time capabilities. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 17, 1313–1327.

[24] Duggan, T., Liu, J., Singh, R., & Zhao, X. (2025). Advancing Earth observation: A survey on AI-powered image processing in satellites. arXiv preprint arXiv:2501.12030. https://arxiv.org/abs/2501.12030

[25] Tuia, D., Schindler, K., Demir, B., Camps-Valls, G., Zhu, X. X., Kochupillai, M., … Schneider, R. (2024). Artificial intelligence to advance Earth observation: a perspective. IEEE Geoscience and Remote Sensing Magazine, 12(3), 5063–5078.* (Advance online publication May 2023 on arXiv:2305.08413).

[26] European Commission, Joint Research Centre. (2024). Earth Observation Strategic Research and Innovation Agenda (Copernicus Sentinel R&I Roadmap). Luxembourg: Publications Office of the European Union.

[27] European Space Agency (ESA). (2024a, September 27). EO MakerSpace: crafting the future of smarter Earth observation systems. Φ-lab / InCubed Program News.

[28] European Space Agency (ESA). (2024b, July 2). New satellite to show how AI advances Earth observation. ESA Observing the Earth – Φ-sat-2 Mission News.

[29] European Space Agency (ESA). (2023, October 26). Six trends to watch in commercial Earth observation. ESA Observing the Earth – News.

[30] World Economic Forum. (2024, January 11). The catalytic potential of artificial intelligence for Earth observation. Cologny/Geneva: WEF. (White Paper in collaboration with Deloitte).

[31] Unibap AB. (2025, January 31). Why add an edge computer to your satellite constellation?. [White paper / Blog post]. Uppsala, Sweden: Unibap News.

[32] KP Labs. (2024, April 12). KP Labs releases first hyperspectral images processed by AI on-board Intuition-1 satellite. [Press release].

[33] Parache, Y. I., Ghiglino, P., & Perzo, N. (2019). High performance on-board image processing using CANopen for Earth observation satellites. In Proceedings of the ESA On-Board Payload Data Processing Workshop (OBDP 2019). ESA/ESTEC, Netherlands.

[34] Planetek Italia. (2025). AI^{ix} – Smart In-Orbit Data Processing (Cognitive Cloud Computing in Space). [Project overview]. Retrieved 2025, from Planetek.it.

[35] Ravichandran, A. (2024, Jan 1). Earth Observation in 2024 and outlook for 2025. TerraWatch Space Newsletter. (Issue: "Ten Trends for EO").

[36] reOrbit. (2024). Earth observation today: Fast data from space unlocking new business opportunities and supercharging environmental science. [Company blog post].

[37] Unibap AB. (2025, January 31). Why add an edge computer to your satellite constellation?. [White paper / Blog post]. Uppsala, Sweden: Unibap News.

[38] European Space Agency (ESA). (n.d.). Unibap SpaceCloud® Framework – In-Orbit Payload Data Processing. [Product portfolio page]. Retrieved 2025, from ESA InCubed website.

[39] Halfacree, G. (2020, October). Intel's Movidius Myriad 2 VPU takes artificial intelligence into space aboard the PhiSat-1. Hackster.io News. Retrieved from: https://www.hackster.io/news/intel-s-movidius-myriad-2-vpu-takes-artificial-intelligence-into-space-aboard-the-phisat-1-af8b6e0b5c5b#:~:text=,raw%20data%20back%20to%20Earth

[40] Kling, U., Moser, S., Hiebl, D., & Ziegler, S. (2023). Reducing data latency with inter-satellite links in LEO constellations: Trade-off analysis and impact on concept of operations. In Proceedings of the 17th International Conference on Space Operations (SpaceOps 2023).

[41] Unibap AB. (2021, March 4). SpaceCloud makes its maiden spaceflight with D-Orbit in Q2 2021 [Press release]. Retrieved from: https://news.bequoted.com/newsroom/unibap/pressreleases/spacecloud-makes-its-maiden-spaceflight-with-d-orbit-in-q2-2-80828

[42] Schreiber, R. (2024, April 12). AI-driven hyperspectral imaging breakthrough by Intuition-1 satellite. SpaceDaily. Retrieved from https://www.spacedaily.com/reports/AI_driven_hyperspectral_imaging_breakthrough_by_intuition_1_satellite_999.html

[43] Ubotica Technologies. (2020, September 28). Satellite successfully applies AI to process Earth observation imagery in-flight in historic first for space [Press release]. Retrieved from https://ubotica.com/satellite-successfully-applies-ai-to-process-earth-observation-imagery-in-flight-in-historic-first-for-space

[44] Valente, F., Eramo, V., & Lavacca, F. G. (2023). Optimal bandwidth and computing resource allocation in low Earth orbit satellite constellations for Earth observation applications. Computer Networks, 232, 109849

[45] European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Official Journal of the European Union, L119, 1–88. https://eur-lex.europa.eu/eli/reg/2016/679/oj

[46] European Space Agency (ESA). (n.d.). Copernicus Sentinels supporting emergency management. Retrieved April 2, 2025, from https://sentinels.copernicus.eu/copernicus-data/services/emergency

[47] Kantamaneni, K., Ravan, S., & Campos, L. C. (2023). A systematic review investigating the use of Earth observation for the assistance of water, sanitation and hygiene in disaster response and recovery. Sustainability, 15(4), 3290. https://doi.org/10.3390/su15043290

[48] Oddo, P. C., & Bolten, J. D. (2019). The value of near real-time Earth observations for improved flood disaster response. Frontiers in Environmental Science, 7, 127. https://doi.org/10.3389/fenvs.2019.00127

[49] NASA Earthdata. (n.d.). Near real-time versus standard data products. NASA Earth Observing System Data and Information System. Retrieved March 28, 2025, from https://www.earthdata.nasa.gov/learn/earth-observation-data-basics/near-real-time-versus-standard-products

[50] Kerr, M., & Koudelka, O. (2022, June 6). Real-Time Earth Observation for Responsive Disaster Management. Technical Session of the 65th United Nations Committee on the Peaceful Uses of Outer Space (COPUOS). Retrieved from https://www.unoosa.org/documents/pdf/copuos/2022/3._Austria_-_EOALERT_DEIMOS_UN_Kerr_Koudelka_2022_final.pdf

[51] European Commission. (2022). Real-time Earth observation products to save lives (EO-ALERT). CORDIS Results in Brief. Retrieved April 2, 2025, from https://cordis.europa.eu/article/id/436188-real-time-earth-observation-products-to-save-lives

[52] Tian, B., Lin, J., Wolfe, R., Murphy, K., Mehta, A., & Ju, J. (2022). NASA LANCE: Near Real-Time Data for Decision Making and Societal Benefits. NASA Goddard Space Flight Center.

Retrieved from https://lance.modaps.eosdis.nasa.gov/docs/LANCE_NRT_Decision_Making_Tian_2022.pdf

[53] KP Labs. (2025, January 28). What is on-board data processing? Retrieved from https://kplabs.space/blog/what-is-on-board-data-processing/

[54] Kawade, P. (2023, October 31). The implications of increasing Earth observation data. Analysys Mason. Retrieved from https://www.analysysmason.com/research/content/articles/the-implications-of-increasing-earth-observation-data/

[55] EUMETSAT. (2023). Meteosat Second Generation (MSG). European Organisation for the Exploitation of Meteorological Satellites. https://www.eumetsat.int/meteosat-second-generation

[56] Planet Labs. (2020). Planet imagery and archive: Monitoring the Earth daily. https://www.planet.com/products/planet-imagery/

[57] Yao, Y., Jiang, Z., Zhang, H., & Zhou, Y. (2019). On-Board Ship Detection in Micro-Nano Satellite Based on Deep Learning and COTS Component. Remote Sensing, 11(7), 762. https://doi.org/10.3390/rs11070762

[58] SpatioTemporal Asset Catalog (STAC) specification: https://stacspec.org

[59] Oerther, T. (2024, 10). OPS-SAT End-of-Life. ESA. Retrieved August 20, 2025, from https://opssat.esa.int/docs/opssat1/2024_ESA_Cleanspacedays_OPS-SAT__EOL.pdf

[60] OPS-SAT Operations nanoSatellite. (2025, February 11). Retrieved August 22, 2025, from eoPortal: https://www.eoportal.org/satellite-missions/ops-sat

[61] Kervennic, E., Louis, T., Benguigui, M., Bobichon, Y., Avaro, N., Grenet, I., . . . Girard, A. (2023). Embedded cloud segmentation using AI : Back on years of experiments in orbit on OPS-SAT. 2023 European Data Handling & Data Processing Conference (EDHPC), 1-8. doi:10.23919/EDHPC59100.2023.10396625

[62] PRISMA Hyperspectral. (2023, July 31). Retrieved August 25, 2025, from eoPortal: https://directory.eoportal.org/satellite-missions/prisma-hyperspectral

[63] Vangi, E., D'Amico, G., Francini, S., Giannetti, F., Lasserre, B., Marchetti, M., & G., C. (2021). The New Hyperspectral Satellite PRISMA: Imagery for Forest Types Discrimination. Sensors (Basel), 21(4):1182. doi:10.3390/s21041182

[64] Loizzo, R. (2018, March 6). EO activities at the Italian Space Agency and PRISMA mission. Retrieved August 26, 2025, from IOCCG Website: https://ioccg.org/wp-content/uploads/2018/03/1.5-prisma_ioccg_6marzo2018-s.pdf

[65] Galeazzi, C., Carpentiero, R., De Cosmo, V., Garramone, L., Longo, F., Lopinto, E., & Varacalli, G. (2009). The PRISMA system and PAN/HYP instrument. Proceedings of the EARSeL Workshop 2009, Tel Aviv. Retrieved August 26, 2025, from https://old.earsel.org/workshops/IS_Tel-Aviv_2009/PDF/earsel-PROCEEDINGS/3020%20Galeazzi.pdf

[66] Cisło-Lesicka, U., & ProGea4D. (2020). 3.1. Program Copernicus. In B. Hejmanowska, & P. Wężyk, Dane satelitarne dla administracji publicznej (1st ed., pp. 131-149). Gdańsk: POLSA (pol. Polska Agencja Kosmiczna. eng. Polish Space Agency). Retrieved August 22, 2025, from https://polsa.gov.pl/wp-content/themes/polsa/files/Podrecznik.pdf

[67] Sentinel-1 Team. (2013). Sentinel-1 User Handbook. European Space Agency. Retrieved August 22, 2025, from https://ftp.itc.nl/pub/Dragon4_Lecturer_2018/ESA%20EEs%20and%20Sentinels%20Brochures%20pdf/Copernicus%20Sentinels%201,2,3/Sentinel-1%20User%20Handbook.pdf

[68] Fisher, Lee, Purdy, Halverson, Dohlen, Cawse-Nicholson, . . . Hook. (2020). ECOSTRESS: NASA's Next Generation Mission to Measure Evapotranspiration From the International Space Station. Water Resources Research, 56. doi:10.1029/2019WR026058

[69] ISS: ECOSTRESS. (2015). Retrieved August 26, 2025, from eoPortal: https://www.eoportal.org/satellite-missions/iss-ecostress

[70] Bakken, Henriksen, Birkeland, Langer, Oudijk, Berg, . . . Johansen. (2023). HYPSO-1 CubeSat: First Images and In-Orbit Characterization. Remote Sensing, 15(3), 755. doi:10.3390/rs15030755

[71] Eine, S. L., Langer, D. D., Birkeland, R., & Orlandic, M. (2024, October 14-18). Reconfiguration of FPGA during operation of small satellite for flexible hyperspectral data compression. 75th International Astronautical Congress (IAC). Retrieved September 1, 2025, from https://rogerbi.folk.ntnu.no/publ/2024_Reconfig_IAC_paper.pdf

[72] NTNU. (2025). HYPSO-1 is a tiny satellite trying to make big waves in ocean research. Retrieved August 25, 2025, from NTNU Website: https://www.ntnu.edu/smallsat/hypso-1

END OF DOCUMENT