

## Unlocking the Future of Data Access at EUMETSAT Meeting the Needs of a Growing and Diverse User Community

Daniel Lee<sup>a\*</sup>, Michael Schick<sup>b</sup>

<sup>a</sup> EUMETSAT, Germany, [Daniel.lee@eumetsat.int](mailto:Daniel.lee@eumetsat.int)

<sup>b</sup> EUMETSAT, Germany, [Michael.Schick@eumetsat.int](mailto:Michael.Schick@eumetsat.int)

\* Corresponding Author

### Abstract

EUMETSAT, Europe's operational satellite agency for weather, oceanography, and climate, has experienced significant growth in its satellite fleet and user base. This expansion has led to a diversification of user needs, ranging from real-time access for weather warnings to long-term data series for climate research and specialized applications. Historically, data access was limited to real-time dissemination via satellite (EUMETCast) or offline retrieval from tape archives primarily designed for long-term preservation. This model served the initial user community of weather agencies focused on real-time warnings, but the expanding capabilities and diverse needs of new users, along with the availability of new technologies, demanded a more flexible and accessible approach.

To address these evolving demands, EUMETSAT has transformed its data access services from traditional satellite broadcasts and offline tape archives to a comprehensive portfolio leveraging cloud technologies and APIs. The new portfolio delivers the equivalent of half a petabyte of data per day to users in real-time via EUMETCast. The Data Store service based on data available online delivers over 100 terabytes across millions of downloads per month to hundreds of active users, and its imagery service EUMETView contains more than 100 configured and prepared dataset visualizations supplying over 2 million images per day to users around the world.

The new data access service portfolio provides a variety of access options, including satellite and terrestrial networks for real-time data, cloud-based services to access & process any EUMETSAT data, including tailoring and visualization capabilities for data exploitation. This diversified approach ensures a level playing field for users with different setups and needs, fostering wider utilization of EUMETSAT's data products. The use of standardized APIs internally and externally enhances the portfolio's extensibility and interoperability, positioning it well to support emerging initiatives like the WMO Information System 2.

The portfolio is unique in many aspects that may be worth imitating for other agencies. The portfolio's services can be used from any cloud and are agnostic of the downstream infrastructure leveraged by users. Users on any continent who require that the data is available on their premises (e.g., agencies that run their own numerical weather prediction models in near-real-time) can self-subscribe to massive amounts of time-critical data with high reliability and low latency via EUMETCast, while others who lack the infrastructure have the option to perform their computation using EUMETSAT services near the data holdings. For users focused on human interpretation of products rather than data assimilation, EUMETSAT's view service greatly reduces the upfront costs of adopting our data.

This paper details the rollout approach behind EUMETSAT's evolving data access services, highlighting the lessons learned during this transformative journey. It explores the current state of the portfolio and outlines future directions in the rapidly changing landscape of data access and utilization. By sharing these insights, EUMETSAT aims to provide valuable guidance to other agencies facing similar challenges in meeting the needs of their growing and diverse user communities.

**Keywords:** cloud technologies, data access, data management

## Acronyms/Abbreviations

EUMETSAT - European Organisation for the Exploitation of Meteorological Satellites  
API - Application Programming Interface  
WMO - World Meteorological Organization  
WIS2 - WMO Information System 2  
EUMETCast - EUMETSAT's real-time satellite-based multicast data dissemination system, available both for Satellite and via Terrestrial networks  
ICSI - Information-Centric Services Infrastructure  
WMS - Web Map Service  
WFS - Web Feature Service  
WCS - Web Coverage Service  
OGC - Open Geospatial Consortium  
GUI - Graphical User Interface  
CLI - Command Line Interface  
FAIR - Findable, Accessible, Interoperable, Reusable  
AI - Artificial Intelligence  
ML - Machine Learning  
NWP - Numerical Weather Prediction  
ORR - Operations Readiness Review

## 1. Introduction

EUMETSAT is Europe's operational satellite agency dedicated to weather, oceanography, and climate monitoring, operating continuously to deliver critical satellite data services not only to European member states but to a global user community. Historically, data provision relied solely on real-time satellite dissemination via EUMETCast for numerical weather prediction and offline retrieval from the Data Centre's tape archives for long-term preservation and climate applications. However, driven by evolving user needs spanning immediate weather warnings to extensive climate research, and empowered by technological advancements, EUMETSAT has revolutionized its data access model. This paper outlines the approach taken to transition to a diversified, cloud-based portfolio, enabling broad, flexible, and equitable data access tailored to the diverse requirements of today's rapidly expanding user community.

## 2. Expanding User Needs and Technological Opportunities

Historically, EUMETSAT operated a relatively limited number of satellite missions, primarily catering to numerical weather prediction where data was provided to users via real-time data dissemination via satellite (EUMETCast) and long-term data storage was on tape archives within our Data Centre. However, recent years have marked significant growth in our mission portfolio, most notably through our involvement in Europe's Copernicus programme. The addition of Sentinel-3, and later Sentinel-6, introduced an array of novel capabilities and instruments, greatly broadening our monitoring scope across weather, oceanography, and climate disciplines. Consequently, our user community diversified, expanding into new application domains and increasing demand for more varied data products.

As our mission responsibilities and data production scaled, the existing dissemination technologies risked falling short of evolving user expectations. Satellite-based real-time distribution, while reliable, incurs substantial costs that escalate proportionally with data volume. Furthermore, the traditional tape-based archives were not well-suited to meeting the expanding user need for comprehensive and accessible time series data.

Moreover, societal recognition of the importance and utility of our data spurred interest from new communities outside our historical user base. These emerging users often required streamlined, flexible access to data tailored to specific application needs, presenting further challenges to our historical dissemination model. This growing demand emphasized the necessity for innovative capabilities and more versatile access methods.

Advances in cloud-based technologies presented a timely opportunity to bridge this gap. Cloud infrastructure offered the agility to rapidly deploy, evolve, and interconnect services through standardised APIs, facilitating loosely coupled, highly adaptable solutions. Recognising this potential, EUMETSAT embarked on a pathfinding initiative to find the future of data access, employing a phased approach. Early initiatives explored cloud infrastructure, data transformation

software, the creation of a data lake, development of visualization services, and modernization of our real-time data distribution systems. Successful pathfinders evolved into operational pilots, setting the stage for today's comprehensive, user-focused operational portfolio.

### **3. Transformation to Cloud and API-based Portfolio**

When EUMETSAT embarked on its modernization journey, it was evident that cloud-based solutions represented the future of data access. We decided to use an on-premises cloud environment, the newly established Information-Centric Services Infrastructure (ICSI), because it was distinct from our operational infrastructure, safeguarding our essential satellite operations. It also gave us complete control over our technology stack and eliminated the technical difficulties that would have been inherent in transferring large data volumes into a commercial cloud. Venturing into this domain required the organization to adopt new approaches and skill sets.

Our primary objective was to explore how to transition from the traditional model of satellite-based real-time data dissemination and offline tape archives toward a more flexible, accessible portfolio driven by cloud technologies and standardised APIs. To achieve this, we developed a series of systems in the ICSI. For clarity, these systems are referred to by the operational names they eventually adopted:

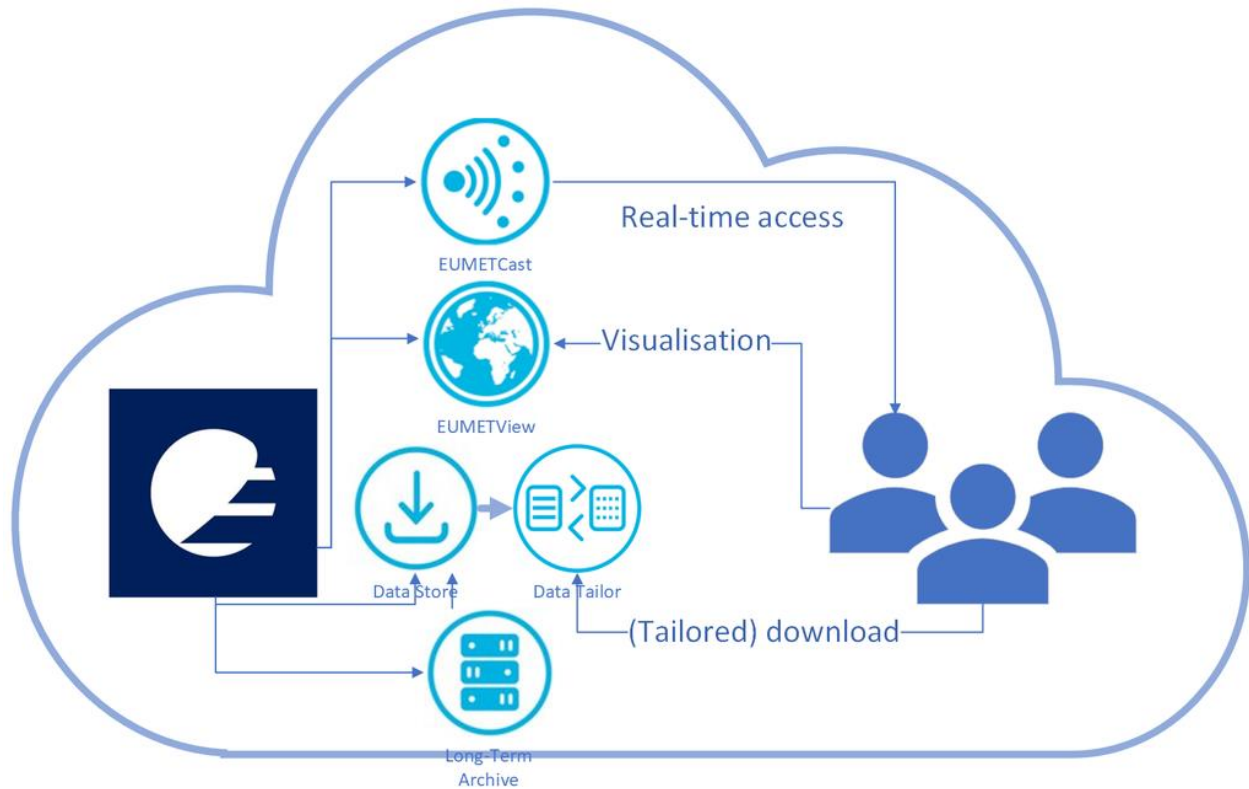
1. **Data Store:** Provides rapid, online access to extensive historical time series data and exposes the full data catalogue via standardised APIs.
2. **EUMETView:** An interactive visualization platform serving a wide array of satellite data products via OGC interfaces, facilitating qualitative data analysis and sharing.
3. **Data Tailor:** A toolbox to convert bespoke EUMETSAT data formats into user-specific formats, significantly facilitating integration into broader workflows, notably within the geospatial community.
4. **EUMETCast Evolution:** Continuing the satellite-based multicast service for real-time, high-reliability data dissemination but with modernised capabilities, ensuring continuity and improved integration into user workflows.
5. **European Weather Cloud:** A joint initiative providing a scalable computing environment co-located with our extensive data holdings, enabling users from member states to perform data-intensive analyses efficiently without extensive local infrastructure investments.

Adopting cloud technology allowed us to rapidly prototype, experiment, and deploy, greatly shortening our learning cycle. Standardised APIs formed the backbone of the portfolio. This let the services interact with each other whilst remaining loosely coupled, and also facilitating the integration of the new data services into diverse user workflows and software ecosystems.

While our implementation was based on an internal platform, the underlying architectural principles are portable. Our choice to build on-premises was influenced by our need to retain control over data location and system design and ensure proximity of processing to our data holdings. This incurred significant up-front investment and high running costs in operations, but aligned with our mandate and long-term strategic goals. Other organisations may find public cloud environments more suitable, depending on regulatory obligations, funding models, or integration needs. Ultimately, the decision between public or private cloud should be guided by a careful assessment of data governance, user requirements, and total cost of ownership.

This transformation represented a significant technological leap for EUMETSAT, dramatically broadening the utility of our data for existing users while significantly lowering barriers for new user communities to access and utilize our satellite data products effectively.

#### 4. Implementation and Operational Rollout



EUMETSAT's longstanding expertise resides in developing satellite systems where failures can have high impacts. However, the services we introduced via the ICSI were decoupled from these operational satellite systems, allowing us to accept greater risk and uncertainty in the early stages of development. This flexibility enabled us to rapidly prototype, solicit user feedback, and iteratively refine solutions, even though initial requirements were not fully defined. To accommodate this exploratory mode of working, we actively collaborated with Quality Assurance and security teams to streamline evaluation processes and reduce delivery overhead.

Automation proved instrumental in managing this process. With automated deployment and testing mechanisms, we sustained rigorous quality control without compromising our ability to rapidly iterate, refine, and enhance the services. Indeed, automation allowed EUMETSAT to simultaneously increase agility and maintain high quality, eventually reaching a deployment cadence of multiple releases per day for some services during intensive development phases. This accelerated feedback cycle was crucial to the timely evolution of services.

Ultimately, as services matured and approached full operational readiness, our rapid development pace shifted towards stabilization. This stabilization phase culminated in an Operations Readiness Review (ORR)—a comprehensive assessment to confirm service reliability, robustness, documentation quality, monitoring readiness, and personnel preparedness. After the ORR, services transitioned into EUMETSAT's operational configuration management system, introducing additional overhead. Deployments also needed to be scheduled across operational systems, which slowed the pace of change. Despite this increased formality, the maturity and robustness gained through continuous, iterative development enabled the new data services portfolio to sustain a regular monthly release cycle post-transition, still let us rapidly incorporate new learnings and needs into our operations.

This implementation journey also let the organisation gain early experience with new engineering practices such as DevOps, significantly enriching our internal capabilities and setting a precedent for future improvements. (A detailed discussion of EUMETSAT's DevOps adoption is available in the separate paper, "Don't Fear the DevOps – We Already do it".)

Following the conclusion of the development phase, the services in the data access portfolio underwent an Operations Readiness Review (ORR) to mark the transition into sustained operations. In this context, "operational" does not imply uniform 24/7 availability—it signifies that the services are considered stable, well-documented, and suitable for handover to dedicated operations teams rather than remaining solely under the stewardship of their development teams. For real-time dissemination services such as EUMETCast Terrestrial, we provide 24/7 on-call support and have implemented automated monitoring aligned with their criticality. In contrast, the other services in the portfolio do not require round-the-clock coverage. This makes it essential to define and communicate differentiated availability expectations and support models to both internal stakeholders and external users.

Monitoring of the data access services is supported by a mixture of enterprise-wide mission performance tools and service-specific telemetry. While our traditional tooling is designed around satellite mission metrics—such as product generation rates from geostationary platforms—these are not always sufficient for assessing the behaviour of the cloud-native services in the portfolio. Traditional enterprise monitoring at EUMETSAT typically focuses on counting inputs and outputs—how many products are ingested and how many are disseminated—which aligns well with the characteristics of satellite mission operations. However, this approach falls short when applied to cloud-native data access services. These systems require more differentiated monitoring, including performance metrics such as query latency, failure rates, and container-level telemetry (e.g., the number of running pods). To address this, we have supplemented existing tools with ELK and other custom-built solutions that provide the additional observability needed to operate these more dynamic, service-oriented platforms effectively. Automated alerting is currently limited to EUMETCast, which is our only real-time service. This targeted use of alerting enables our staffing model to remain sustainable by limiting 24/7 support to systems with strict timeliness requirements, as continuous coverage is resource-intensive and only justified where operational necessity demands it.

Service management for the portfolio is a collaborative effort. During office hours, analysts monitor system performance and escalate issues to a mix of operations and development engineers, depending on the nature of the issue. Our escalation procedures are evolving in tandem with our operational experience. Across the data access services, we target approximately one release per month, which translates roughly into one release per system per quarter, following a planned and coordinated process. While some deployment steps are automated, integration with existing manual configuration management introduces constraints, and improving release efficiency and traceability remains an area of focus.

Operational documentation—including standard operating procedures and training material—is a prerequisite for transitioning any service into operations. These resources form the foundation for the handover. Looking ahead, we see potential value in more gradual transitions, where development teams retain partial ownership post-ORR to support early operational phases. This would enable better feedback loops between developers and operators. Although current responsibilities are distributed across organisational units, we maintain close collaboration to ensure continuity and resilience in service delivery.

The evolution of our data access services is also helping us to consider established operational norms at EUMETSAT in new light. While our organisation has traditionally prioritised stability in long-lived satellite systems, modern data services often benefit from iterative change and rapid responsiveness to user needs. In this context, frequent updates are not a sign of immaturity, but of an architecture that can safely accommodate evolution. Adopting controlled, continuous change strategies helps reduce overall operational risk—particularly for user-facing services—as compared to large and infrequent releases. This highlights the importance of tailoring operational models to match the lifecycle and pace of each system.

## **5. Outcomes and Lessons Learned**

The new data access services have successfully transitioned into operations and have been widely adopted across the user community. The following outcomes reflect both technical performance and key lessons learned during this transition—insights that may inform similar efforts at other organisations.

EUMETCast Terrestrial, in particular, has seen significant uptake, with 77 institutional users currently relying on this real-time terrestrial multicast distribution system, delivering more than a petabyte of data per day to users located on almost every continent in real-time.

Migration trends clearly indicate a shift away from legacy tape-based archives toward the Data Store, affirming users' strong preference for the enhanced accessibility and responsiveness of the new portfolio. Internally, the infrastructure supporting the Data Store provides access to over one thousand distinct data collections—more than one hundred of which are publicly available—and handles millions of newly added data products daily. Presently, the Data Store consistently delivers over 100 terabytes of data each month to approximately one thousand active users. The strategic decision to prioritize standardized API-based access has clearly resonated with users, with the majority integrating our data directly into custom software solutions via API call. Access via our command-line interface (CLI) EUMDAC is the second most utilized option, while graphical user interface (GUI)-based access remains comparatively low, probably because it is used mostly for explorative workflows that are later automated.

EUMETView, our visualization service, now offers nearly 100 data layers accessible via standardized web service interfaces (WMS, WFS, and WCS). The service supports visualizations going back to 2020, with continuous expansion of both historical and newly collected data. The platform serves around two million requests per day and also powers EUMETSAT's popular Earthview livestream on YouTube.

A crucial lesson was the benefit realized through the implementation of a loosely coupled, scalable cloud-based architecture—namely, the ability to quickly and cost-effectively respond to evolving requirements. For example, the recent activation of the World Meteorological Organization's Information System 2 (WIS2) required only straightforward extensions of existing capabilities, avoiding costly redesigns or substantial new investments. This architectural agility continues to allow EUMETSAT to efficiently engage new user communities without jeopardizing the integrity or manageability of our infrastructure.

This positive user feedback reaffirms several of our strategic and technological choices—particularly the emphasis on modular services, standardised APIs, and automation as foundational principles. Importantly, the experiences and insights gained from this initiative are valuable beyond the scope of data services, offering beneficial lessons for improvements across the wider landscape at EUMETSAT.

## **6. The way forward**

EUMETSAT's vision for Data Access Services is to continuously evolve, providing unified, user-friendly, and cost-effective access to all EUMETSAT products with appropriate timeliness. Our goal is to empower users to exploit data at scale, turning information into actionable knowledge to safeguard lives, protect property, and enable informed decision-making in the context of climate change.

In the near term, our existing data access services portfolio will integrate more closely with our evolving cloud computing offerings. Users of these cloud-based services will experience enhanced edge computing capabilities, ensuring data processing is highly efficient and occurs proximate to the data holdings. Additionally, we are pursuing substantial enhancements to our data services, specifically focusing on making our datasets more readily analyzable and suitable for Artificial Intelligence (AI) and Machine Learning (ML) workflows. A significant shift is anticipated from a traditional load-and-transform approach towards implementing a comprehensive data fabric paradigm, facilitating easier, more integrated, and scalable data exploitation. Concurrently, we are preparing to migrate our extensive long-term archive into the cloud environment, further integrating historical datasets seamlessly with the broader portfolio of data services.

Looking to the mid-term, we anticipate AI and ML to become increasingly essential components in satellite data utilization. Our approach is centered on adhering closely to the FAIR principles (Findable, Accessible, Interoperable, Reusable), ensuring that our data is structured to facilitate easy discoverability and efficient use by automated analytical systems. Efforts will focus particularly on improving how our data can be streamed using community tools and standards, and improving our data taxonomy and metadata. This should enable AI-driven identification and retrieval of relevant datasets, thereby reducing reliance on traditional human-centric approaches to data discovery - the degree of human involvement should be driven by business needs and intent, not by technology.

Ultimately, by proactively addressing both near-term challenges and future opportunities, EUMETSAT will ensure its data access services portfolio remains adaptive, responsive, and optimally positioned to support the expanding needs of our diverse user community for years to come.

## 7. Lessons for Other Organizations

EUMETSAT's experience in evolving its data access portfolio offers a number of concrete takeaways that may benefit other organizations. While we do not expect other agencies to replicate our approach exactly—each operates under different constraints and priorities—the underlying principles and practices we followed can inform efforts to modernise data access in ways that best fit their specific needs.

### 7.1 Transferable design principles

- **Loose coupling through standardised APIs** has proven essential in building a resilient and flexible ecosystem. It enables independent evolution of services while maintaining interoperability.
- **Microservice-oriented architecture** should be favoured over large, tightly integrated monoliths. This approach facilitates agile development, simplifies maintenance, and allows for targeted scaling.
- **Clear separation between experimental and operational systems** was a key enabler for innovation. By decoupling development from mission-critical infrastructure, we were able to prototype, iterate, and adapt rapidly without risking core operations.
- **Automation throughout the development lifecycle**, including continuous integration, deployment, and quality assurance, significantly reduced time-to-deployment while preserving service reliability.
- **Data proximity**, whether achieved through on-premises or commercial cloud deployments, is critical. Placing processing close to data minimises latency and bandwidth costs, enabling efficient access and use, particularly at scale.

### 7.2 Constraining costs

While EUMETSAT's approach was enabled by a dedicated, internal cloud platform (ICSI), replication does not require equivalent infrastructure. The platform itself is complex and resource-intensive to establish, but many of the benefits we realised stem from architectural and procedural choices rather than platform specifics.

For organizations looking to make cost-effective improvements, initial steps might include:

- Exposing data catalogues or subsets via API using community standards.
- Offering downloadable products via basic, hosted endpoints.
- Adopting cloud-native principles (statelessness, containerisation, modular deployments) incrementally within existing systems.

These efforts can deliver tangible improvements in accessibility and flexibility, even without large-scale infrastructure investment.

### 7.3 Practical implementation

- **Start small with user-driven pathfinders.** Early engagement with users helps focus development efforts on high-impact needs and provides critical feedback loops.
- **Invest in DevOps skills and automation early.** Embedding these practices supports continuous delivery and scalability from the outset.
- **Integrate QA and security into agile workflows.** These disciplines should be embedded throughout development, enabling quality and compliance without becoming delivery bottlenecks.

By sharing these lessons, we aim to support other agencies in making practical, informed choices as they modernise their own data services.

## 8. Conclusions

The successful transformation of EUMETSAT's data access services portfolio represents a significant achievement in adapting to the evolving landscape of user needs and technological opportunities. By strategically

embracing new engineering approaches, cloud technology, standardized APIs, and a flexible, scalable architecture, EUMETSAT has significantly expanded accessibility and usability of its satellite data, enabling a rapidly growing and diverse global user community to leverage our products more effectively.

The outcomes underscore the importance of close collaboration with users and across siloes, agile and iterative development processes, and the crucial role of automation and standardised interfaces in enabling swift, reliable, and efficient service evolution. It also highlights the need for swift deployment to expose our ideas to reality as quickly as possible and calibrate our implementation against that. Looking ahead, EUMETSAT remains committed to ongoing improvement, particularly in enhancing data readiness for advanced analytical techniques including AI and Machine Learning, and fully realizing the potential of modern cloud-based technologies.

By openly sharing the experiences, challenges, and insights gained throughout this transformative journey, EUMETSAT hopes to provide a valuable reference point for other satellite and data-focused organizations embarking on similar paths, fostering wider advancements in operational satellite data utilization worldwide.