Are the Clouds Clearing on the Ground? Redesigning Ground Segments for Secure Constellations

Edoardo Cocci[a], Stewart Hall [a].

[a] *Telespazio Germany GmbH, Europaplatz 5, 64293 Darmstadt, Germany,*
edoardo.cocci@telespazio.de, stewart.hall@telespazio.de

Abstract

Cloud-native ground segments have played a central role in enabling scalability, flexibility, and automation for modern satellite constellations. By leveraging elastic infrastructure and cloud-based services, operators have been able to support increasingly complex missions with smaller teams and shorter development cycles. However, cloud-native design is often implicitly equated with cloud-dependent deployment, a coupling that is increasingly being questioned as space systems assume defence and security-critical roles.

As sovereignty, security, and regulatory requirements intensify, many constellation operators are reconsidering where and how their ground infrastructure runs. In several cases, this has led to renewed interest in on-premises or sovereign deployments, even for ground segments originally designed to operate in public cloud environments. This shift introduces a key architectural challenge: how can operators preserve scalability, automation, and operational agility while operating outside traditional cloud platforms?

This paper argues that these capabilities are not inherent properties of the cloud itself, but rather outcomes of specific architectural design choices. Drawing on lessons learned from a cloud-to-on-premises migration of a constellation ground segment, the paper examines which design principles enable portability across deployment environments and which assumptions create hidden dependencies. It discusses the role of modular architectures, clear system boundaries, and API-first design in decoupling operational logic from infrastructure, allowing ground segments to adapt to security-driven constraints without sacrificing automation or scalability.

By framing cloud-native capabilities as architectural outcomes rather than deployment characteristics, the paper provides guidance for designing ground segments that remain flexible across cloud, on-premises, and hybrid environments. This perspective is particularly relevant for future constellations operating at the intersection of commercial scalability and security-critical requirements.

Keywords: Cloud-native architecture, ground segment, satellite constellations, deployment portability, microservices, sovereignty, DevSecOps, infrastructure-as-code, containerization, Kubernetes

Acronyms:

API	Application Programming Interface
AWS	Amazon Web Services
DevSecOps	Development, Security, and Operations
HA	High Availability
IaC	Infrastructure-as-Code
TRL	Technology Readiness Level

1. Introduction

Cloud computing is a profound change to IT solutions forcing abandonment of legacy software architectures. Cloud computing has reached near-universal adoption across industries with some forecasts of $947 Bn in 2026 being spent on cloud services [REF1]. The volumes of data moving through cloud computing are staggering, estimated at 200 zettabytes. Cloud computing has formed part of the strategies of business modernisation with 96% using it and 94% reporting improved security as a result. By 2032, the market could hit USD $2 Tn.

Adoption across various technical and critical domains is widespread: The technology industry is a ubiquitous user: financial services are a major use at 84%; manufacturing usage is 60%; healthcare is more cautions at 41%; government down at 16.8% (US). The common use cases are software-as-a-service, storage and backup, and DevOps.

In contrast, the adoption of the cloud computing has been slower in the historically conservative domain of satellite manufacturing and operations. Major US cloud providers, such as AWS, Azure and Google Cloud have partnerships with satellite operators to provide on-demand virtualised ground stations. With the development of more advanced rad-tolerant processors there is a trend toward on-board edge-computing needed for software defined satellite in domains such as satcom and earth observation. Space-based data centres are being proposed. In some respects, cloud computing adoption has been slow in the satellite control segment.

In recent years the evolution of cloud-native ground segments has demonstrated real benefits achieved with cloud architectures: elasticity, DevOps agility, cost scaling. However, their close association with public cloud deployments is now being challenged by security-driven requirements, particularly by European governments, especially use of non-EU services.

There is a tension between cloud-native architectures and deployment sovereignty, driven largely by an assumption often implicating that "cloud-native" means "public cloud". Together with rising security and sovereignty pressures, this is now putting the need of re-framing the architectural question.

Note that the effects described in this paper arise from the experiences of Telespazio Germany in developing its commercial ground segment product known as EASE-Rise. Prototype versions of were code-named ENABLE and based on AWS components. The commercialised version, which achieved Technology Readiness level 9 in January 2025, continued to rely on AWS.

Throughout 2025, there were many conversations with governments who expressed their concerns about "cloud" and sovereignty including reliance on AWS. Since then, EASE-Rise has been migrated to cloud agnostic technology (Kubernetes cluster). New enquires may be satisfied with EASE-Rise running on any sovereign cloud or on-prem.

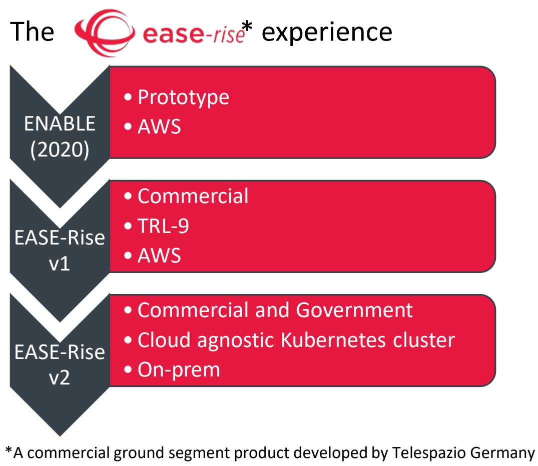

2. Cloud Assumption Is Being Revisited - The Architect's Dilemma

The decades of relevant experience of Telespazio Germany enables some reflection on the dilemma of the ground segment architect. Classically, the architect has focussed on the functional requirements of a ground system, decomposing it into functional sub elements that can be procured and developed. But enterprise architecture asks more profound questions of a non-technical nature. For example: how should a system architecture address the short-term needs and risks of a short duration or experimental satellite mission? Contrastingly, how should it address the long-term needs of the operators should the experimental mission lead-on to sustained operational missions? In the former case, there is questionable value buying on-premises hardware when public cloud computing can provide the compute. In the latter case, the ongoing operational costs of public cloud must be calculated carefully against the periodic costs of on premises hardware re-procurement and software re-integration.

Recent years have seen significantly more proposals for new government defence and dual use satellite missions, especially constellation missions. As mentioned in the introduction, governments are among the slowest adopters of cloud computing. In part this may be due to previous experience and ideology, but when asked, they cite legitimate concerns about data sovereignty, data location, assurance and ultimate control of the supply chain i.e. industrial dependency. Government missions, and government influence of commercial missions, are more likely to be the reason why cloud computing is not and may never be ubiquitous in satellite systems.

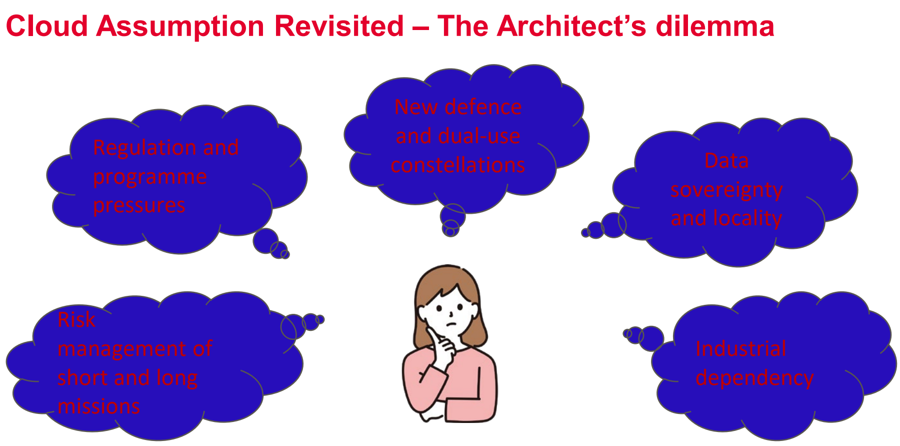

3. Architectural Challenges and benefits of Cloud

Does cloud computing automatically solve all the problems of lower costs, deploy-ability, scalability and automation [REF3]? In our experience, no it doesn't. Cloud computing and public clouds are an enabler of these properties, but these are properties of software architecture. In our examination of legacy monolithic systems, we observed that "hidden coupling" was the real obstacle to using them in the cloud. In effect, it was impossible to containerise services since these were internalised within the architecture; individual services were deeply coupled to all the others resulting in technical rigidity. Consequently, we have observed that cloud deployment of legacy monolithic system aka lift-and-shift, are typically entire virtual machines.

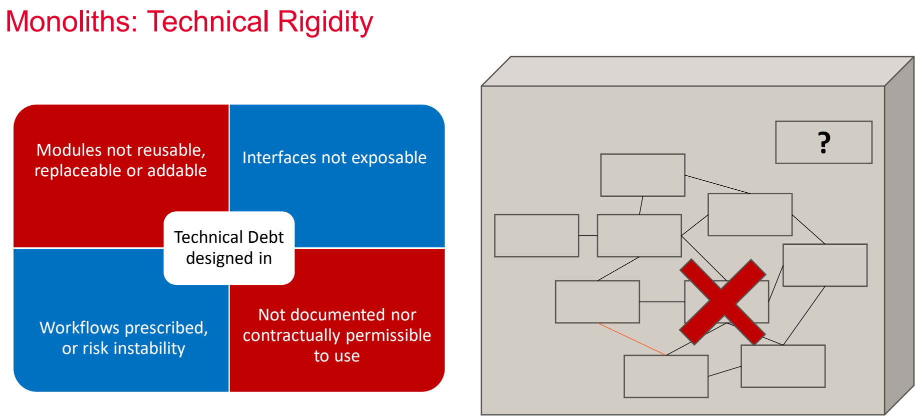

The lift-and-shift approach that relies on virtual machines fails to deliver the service orientation or deployment flexibility that user wants, it results in consistently high resource allocation (and cost), and embeds mission data in inaccessible ways. In this section, we outline the architectural principles from infrastructure choices.

Stateless micro-services and data lake architectures unlock all the potential of cloud computing. But these need to be orchestrated to deliver ground segment functions. Orchestration requires a deep understanding of all the potential user workflows. The data-lake needs to support replication and seamless load balancing and synchronisation across replicas, potentially in different geographic locations, e.g. Kafka.

Resilience is almost for-free: By instantiating stateless replica services everywhere**,** together with health-checks, retries and circuit breaker policies (if services fail). Replication and data synchronisation works over large geographies, being responsive across multiple location within a European country. It does not break over even larger geographies e.g. Germany to Australia has a typical latency of 200-300ms.

Scalability: On public cloud, the bundles of micro-services that make up the operational domain of a new satellites can be simply scaled with a single click. However, on premises requires physical architecture and more complex network design and deployment.

Availability: With on-premises systems, availability is calculated based on hardware and software mean-time-to-fail and mean-time-to-recover**.** In cloud computing, compliant physical hardware availability is never an issue, and software exists in multiple replicas with fast restarts. High availability is therefore achieved by specific software architecture that protects and dedicates cloud resources to the mission. We always recommend regular availability cloud services for satellite testing and integration, and high availability for satellite operations.

Obsolescence: With on-premises systems, hardware obsolescence is a perennial problem, albeit partially solved with machine virtualisation which has simplified the insertion of new servers. With cloud computing, the migration of micro-services from old server to new server is even more seamless.

Security: Physical security provided automatically by the public cloud, but how is it provided on-premises? Software and data security, however, can rely on standard authentication and authorisation solutions such as Cognito and Keycloak and either cloud provided or mission specific crypto services.

4. A Case of Cloud-to-On-Prem Migration

Essentially, this is the lessons learned from the EASE-Rise experience. The EASE-Rise migration to cloud agnostic technology demonstrates that when architecture is modular and API-driven, migration becomes adaptation, not a complete redesign.

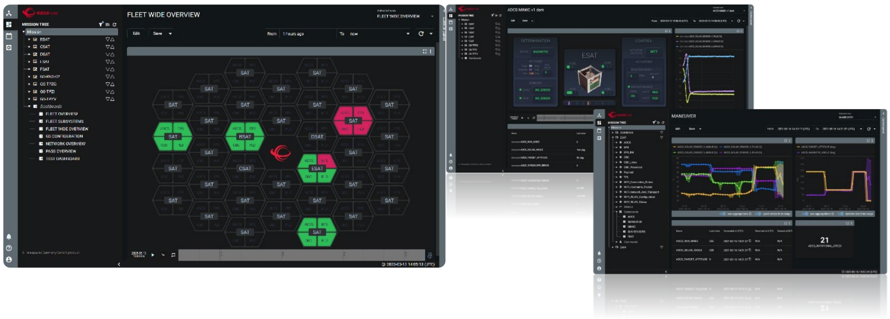

Technological and industrial revolutions as profound as cloud computing can't be ignored by the space industry. Cloud-native infrastructure is ideal for the elastic, scalable ground segments that are needed for constellations and distributed operations. Leveraging domain knowledge and starting from scratch can be faster than re-architecting legacy monolithic software.

With the benefit of hindsight, it was a good decision to take the quickest technology path to market. In our case this was starting with AWS since it offered a more comprehensive and coherent set of cloud technologies at that time. Using AWS it established EASE-Rise as TRL-8/9 as early as possible which has created considerable presence in the market. It is never possible to predict for geo-political changes five years in advance. It is geo-politics that has created demand for sovereign solutions; migration from AWS to Kubernetes cloud-agnostic infrastructure was a necessary response to the market changes**.**

We worked with a German specialist MaibornWolff [REF4] to co-engineer the migration. They are infrastructure specialists who work alongside you product experts. This teaming significantly eases the task of migration and ensures right first time, every time.

5. Architectural Principles for Secure and Portable Ground Segments

Based on the EASE-Rise migration experience, a simple set of practical design guidelines emerges which can be used to guide future space projects of this type. The fundamentals of secure and portable ground segments rest on four principles: **Micro-Services Architecture**; **Open Standards**; **Containers**; and **DevSecOps, Infrastructure-as-Code**. These principles demonstrate that cloud-level operational capabilities can be retained in on-prem or hybrid environments.

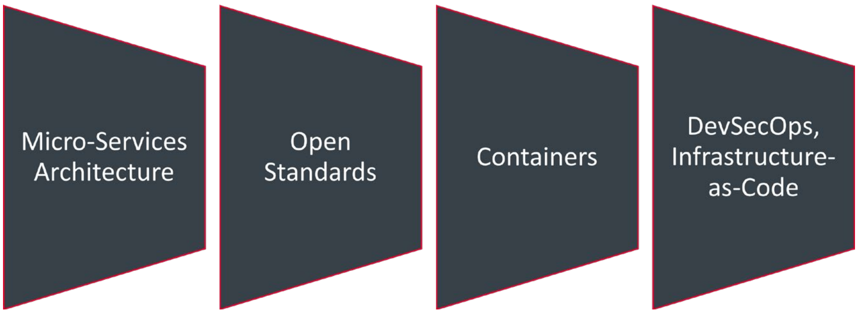

6. Conclusion

Ground segment maturity lies in decoupling operational capability from infrastructure choice. Cloud native architecture solved the problems of hardware, cost of ownership, software deployment, obsolescence, and inflexibility. However, sovereignty and security needs from expansion of government programmes drove the migration to cloud-agnostic architecture, revealing an important distinction: cloud native does not mean public cloud dependent. Many of the benefits of cloud also help with on premises deployment. It is service-oriented architecture ultimately determines cloud and on-premises portability.

7. References

REF1	55 Cloud Computing Statistics for 2024 – https://spacelift.io/blog/cloud-computing-statistics
REF2	**Cloud Computing Market Size, Share & Growth Report, 2034**
REF3	**ESAW 2026 Orchestrating the Ground Segment E. Cocci**
REF4	**https://www.maibornwolff.de/**