

COR-Virtu: A path towards virtualization on operational Ground Center

François-Xavier Mouy^{a*}, Hélène Ruiz^b

^a *COR-VIRTU Technical Manager, CNES DTN/TSA/LMC, francois-xavier.mouy@cnes.fr*

^b *Head of "Shared Software and Facilities" Department, CNES DTN/TSA/LMC, helene.ruiz@cnes.fr
Centre National d'Etudes Spatiales, Toulouse, France*

* Corresponding Author

Abstract

The French Space Agency, CNES, has been operating an S+X-band Ground Stations Network for more than 30 years. This network is a multi-mission asset that provides Tracking, Telemetry and Command (TT&C) services to French and Multinational governmental satellites (currently around 10 in flight) in Launch and Early Orbit Phase (LEOP), in mission phase and in the End Of Life operations (EOL). On occasions, the spare capacity of the network provides equivalent services to other satellite or network operators in the frame of cross supports or commercial agreements.

The COR (the Operational Network Center) Facility, set in Toulouse France, implements all the services between Satellites Control Center or Mission Centers and WorldWide Antennas. The center is composed of several systems, affected to several goals. The center is 24/7 manned.

Since several years CNES NOC start, mainly for manage hardware obsolescence, to virtualize (mostly in p2v) server's systems. Benefits were very good and improved the reliability and scalability of systems.

However, the provisioning was not really managed and the hardware used to be sized up by a sum of old physicals systems characteristics (CPU, RAM, Disk). That method led to very expensive massive over provisioned architecture (hardware and licensing).

CNES decided to improve the sizing of the hardware but without taking the risk of loss of performance issue. Indeed, we operate several software's and platforms from multiples providers (COTS, internal handmade tools, or specific ordered products). For some software, the requirements were defined a long time ago for platforms that don't exist anymore.

After a brief recall of the COR architecture, the paper will present the steps that lead to a better hardware provisioning management.

The paper will explain how we start with the improvement of the monitoring of systems in operation to build a testing method on an offline representative platform of several hardware architectures to meet real software needs with the best performances.

Finally, the paper will present the choices done and benefits founds. The path to our Ground center virtualization with our own timeline out of the virtualization fashion way. We had some surprises comparing some software's prerequisites editors versus real needs in our context.

Keywords: (Virtualization, Provisioning, Sustainability, Operational availability)

Acronyms/Abbreviations

Physical to Virtual method (P2V)

Network Operation Center /Centre d'Opérations Réseau (NOC / COR)

1. Introduction

1.1 The Ground Station Network

Since 1984 the CNES Ground Stations Network has been providing ground stations supports to satellites from the Low Earth Orbit (LEO) to Geostationary Earth Orbit (GEO) altitudes range and in all mission phases from Launch and Early Orbit Phase to End of Life. This Network is currently composed of:

- 5 dual-band (S and X) and 2 S-band multimission antennas compatible with LEO, MEO and GEO satellites requirements (link budgets, antenna speeds and accelerations...)
- A Network Operations Center (NOC) managing the operations of the ground stations and the interfaces with the user's control and mission centers
- A Scheduling Office in charge of the elaboration of the ground stations schedule and of statistics on the performance of the services
- An Orbit Computation Center (OCC) in charge of the processing of the localization measurements performed by the ground stations, the restitution of the orbits of the satellites and the provision of predicted pointing data to the ground stations.

The perimeter of the CNES Ground Stations Network and its interfaces with internal and external customers as well as with external suppliers is shown on the following picture.

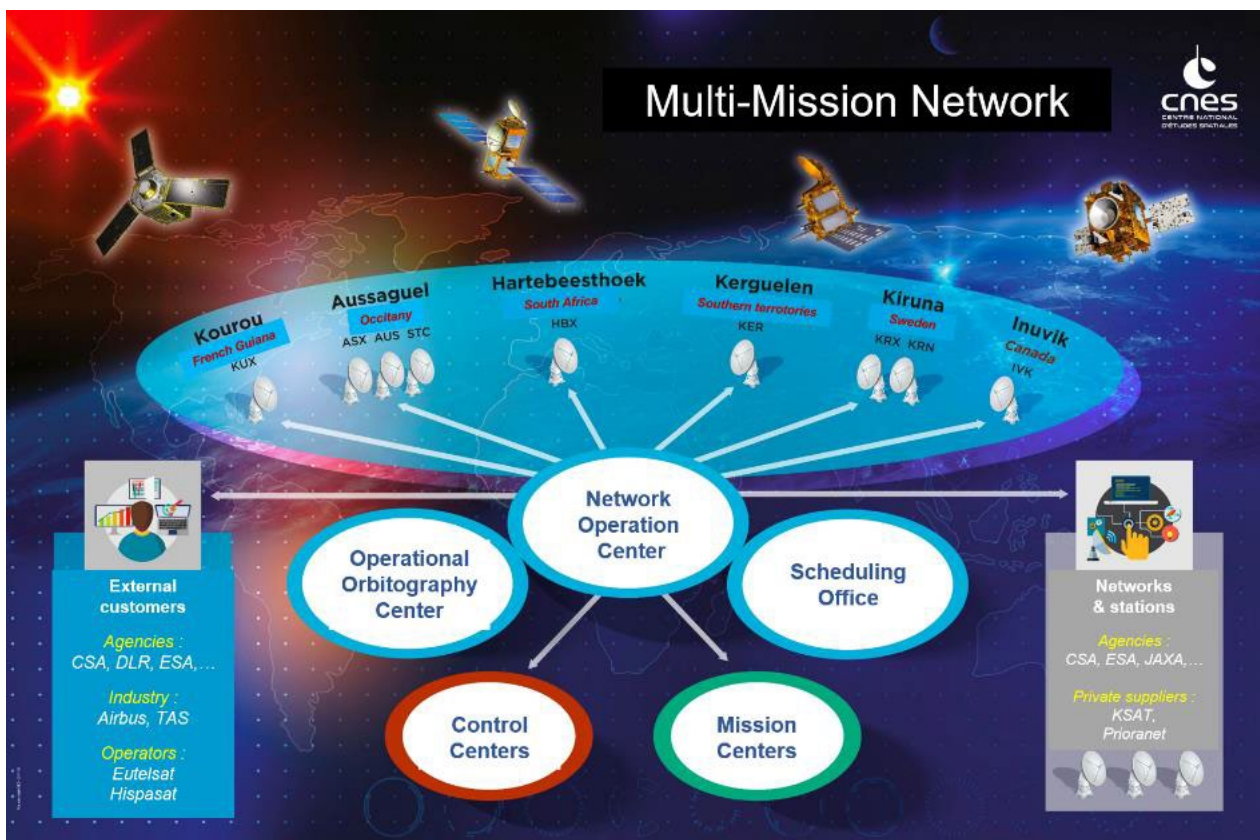


Fig1. Perimeter of the CNES Ground Stations Network

The Network Operations Center is the central point of the global Network in terms of architecture but also for the provision of the services required by the control and mission centers of the supported missions. On a 24/7 basis, the main role of the NOC operator is to:

- Monitor the status of the ground stations and operate them.
- Monitor and operates all systems and interfaces involved in the data exchanges between the stations and the customer's centers for the provision of the services:
- Platform and/or payload telemetry provision in real time and/or post-pass
- Real-time sending of commands
- Collection and provision post-pass of localization data (ranging, 1-way or 2-way Doppler and antenna angles measurements)

All those systems have their own software lifecycle and for the long term stability needed, maintaining, integrating new versions, just keeping everything running is a real challenge

1.2 Virtualization beginning

In recent years, CNES (Centre National d'Études Spatiales) Network Operations ground Center (NOC) first adopted virtualization as a strategic solution to manage the growing challenge of hardware obsolescence. The first major step in this direction was the introduction of virtualization across our servers. The initial goal was to address the need for efficient **remote management** of operating system updates, which needed to be applied simultaneously across our infrastructure spread around the world. This became especially critical given the diverse geographic locations and mission-critical nature of the systems in operation.

Virtualization allowed CNES to roll out updates to **remote operating systems** in a controlled manner, ensuring consistency and minimizing the risk of inconsistencies between different environments. Moreover, one of the key advantages was the ability to implement **rollback mechanisms**. If an update caused unforeseen issues or impacted system performance, the rollback option provided an effective safety net, enabling quick recovery to a previous, stable configuration without disruption to ongoing operations.

This approach marked the beginning of a broader effort to virtualize our infrastructure. As the benefits became clear in terms of reliability and flexibility, we began transitioning more systems into virtualized environments, initially focusing on **Physical-to-Virtual (P2V)** migration to extend the lifespan of aging hardware. However, over time, we faced new challenges in hardware provisioning and system performance that required further refinement to optimize resource allocation.

2. The sizing challenge

While virtualization delivered substantial improvements, the provisioning strategy used during its initial implementation was far from optimal. The process involved sizing up hardware resources by simply aggregating the characteristics of various older physical systems, which were often inconsistent in terms of CPU, RAM, and disk configurations. This method resulted in a heavily over-provisioned architecture—leading to costly investments in both hardware and licensing. With resources allocated generously across the board, many systems operated underused capacity, thus inflating operational costs unnecessarily.

To comply with the need for improved hardware sizing without sacrificing performances, CNES set out to refine the hardware provisioning approach to better match actual software requirements. The challenge was twofold: we needed to improve the accuracy of hardware sizing while avoiding the risk of performance degradation. At CNES, our systems support a wide range of software and platforms, including commercial off-the-shelf (COTS) solutions, internally developed tools, and customized products. The requirements for many of these software systems were defined years ago, often based on legacy platforms that no longer exist.

To address this gap, CNES sought to establish a more precise and efficient way of determining hardware specifications, accounting for the real operational needs of each system.

2.1 Enhancing System Monitoring

Our journey towards better provisioning began with a focused effort to enhance the monitoring of systems in operation. We recognized that accurate, real-time insights into resource utilization (such as CPU usage, memory consumption, and disk I/O) were crucial in making informed decisions about hardware requirements.

Lucky with a dedicated offline platform built to simulate a variety of hardware architectures and configurations, we improved monitoring data, turned to a testing methodology that would allow us to validate our assumptions and refine our provisioning approach. By deploying advanced monitoring tools, we could track how each virtual machine and application performed under typical loads, providing valuable data that would drive the refinement of our provisioning process. By running detailed performance tests on this platform, we could compare different hardware architectures against the real-world demands of our software, ensuring that each configuration met performance expectations without unnecessary over-provisioning.

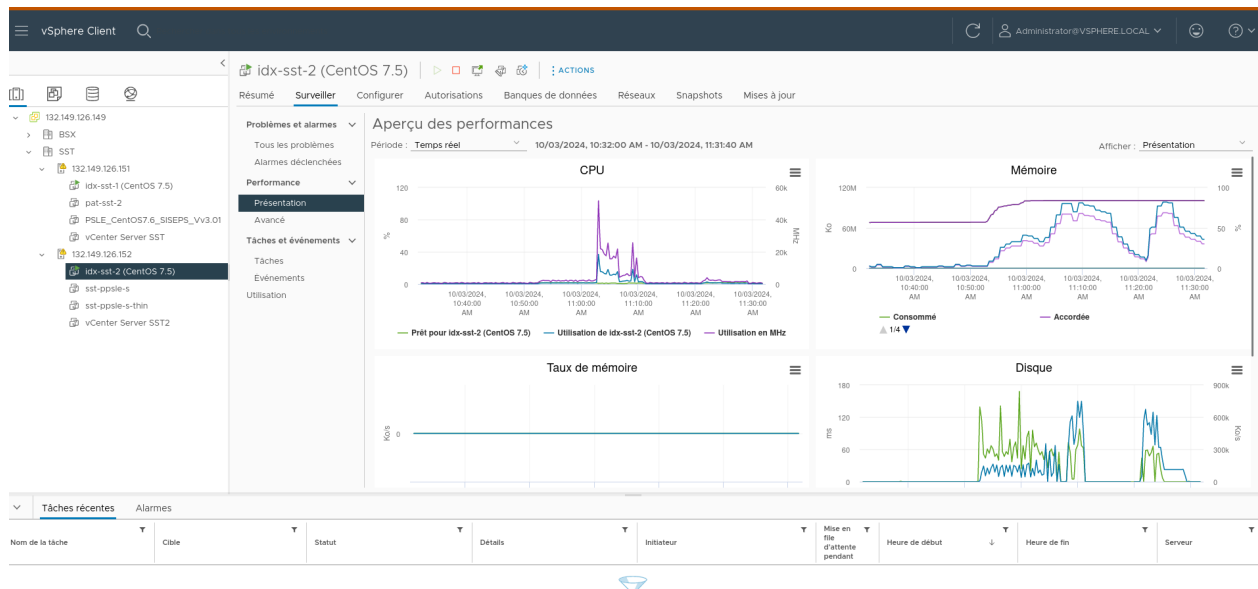


Fig. 2. Vcenter Monitoring during satellite passage processing session

2.1 Improving offline platform representativeness

The offline pre-production platform can't be an integral duplication of the whole system but even if the operational network partitioning, for security aspect, cannot be virtualized, this constrains is not applicable for offline platform. Virtualization allowed us to simulate not just physical hardware but entire systems, even those that had not yet been virtualized in the production environment. This is particularly beneficial for systems where hardware resources (such as network interfaces, specific drivers, or specialized I/O requirements) are unique or complex. Virtualized systems can emulate many of the configurations of non-virtualized environments, ensuring that software is tested under conditions that closely resemble the production environment, even if those conditions are not yet fully virtualized.

We have the ability to do it for improve representativeness and test the whole processing chain. This platform played a critical role in ensuring that our systems were tested under realistic conditions. Virtualization, in particular, provided significant flexibility by allowing us to simulate environments that were not yet virtualized. Many legacy systems or certain specialized software platforms had not yet been moved to virtualized environments, and traditional testing methods would have required physical hardware. However, virtualization made it possible to create accurate simulations of these systems, thereby improving the overall representativeness of our tests.

Simulating network conditions within a virtualized platform also provided us with several key benefits:

Real-World Network Replication: Virtualization allowed us to replicate complex network topologies without needing to set up multiple physical machines. We could simulate the impact of network delays, packet loss, and varying bandwidth constraints on system performance, even before moving the actual systems into the virtualized environment.

Testing Network-Dependent Software: For software applications with strict network requirements, such as distributed systems or real-time communications software, virtualization helped us evaluate their performance

without the need for dedicated physical resources. We could simulate different network configurations, such as high-bandwidth, low-latency, or geographically distributed networks, to observe how these configurations impacted the software’s operation.

Virtualized Network Interfaces: The use of virtual network interfaces allowed for detailed testing of network-bound performance in a controlled, reproducible manner. We could fine-tune network parameters and test how software responded to different configurations (such as VLANs, DNS settings, or IP configurations), ensuring that any network-related issues would be identified before full deployment.

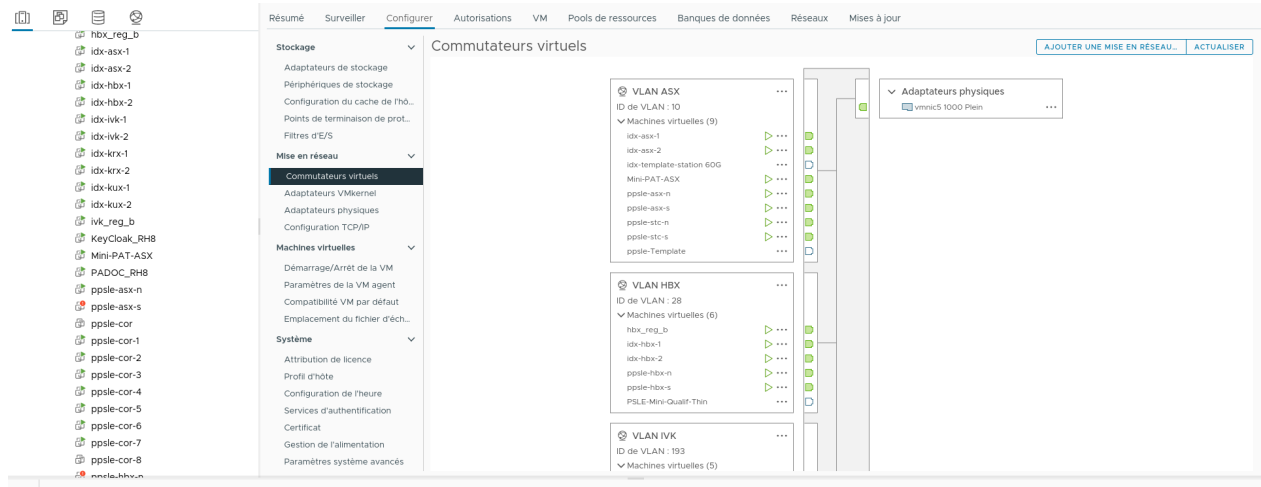


Fig. 3. Virtual Lan platform configuration example for antenna station simulation

3. Analysis of systems behaviour

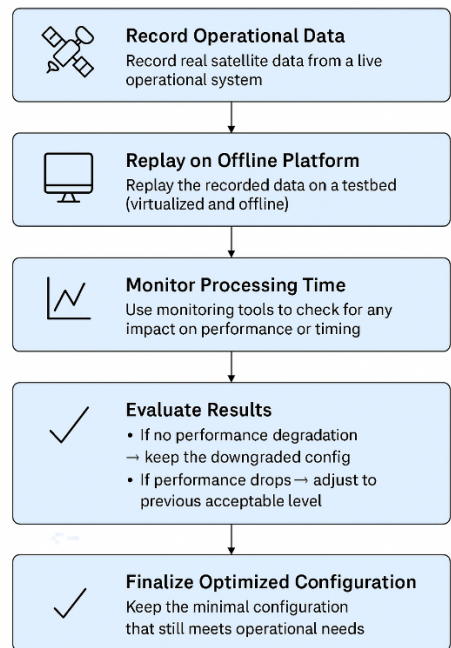
To evaluate the actual resource needs of each system, we adopted a **pragmatic and reproducible testing methodology** using real operational data. Rather than relying solely on theoretical sizing or outdated specifications, we recorded satellite data processed by our operational systems during routine activity. This data set served as a realistic workload, capturing the true computational profile of our applications.

We then **replayed this data** on our offline virtualized test platform, gradually **reducing allocated resources**—such as CPU or memory—in a controlled environment. At each step, we monitored the time required for processing and assessed system behavior using enhanced performance monitoring tools. As long as there was no measurable degradation in processing time or functionality, we continued to reduce the resources incrementally.

This approach allowed us to define the **real lower bound of required resources** without compromising performance. It also helped us identify over-provisioned systems and confidently adjust their sizing, based on measurable outcomes rather than assumptions. The method is simple in principle, but highly effective when supported by **good observability and monitoring practices**.

Ultimately, this replay-based approach ensured that our provisioning strategy was both **data-driven and directly connected to operational reality**, aligning our infrastructure with actual needs rather than conservative estimations.

Replay-Based Resource Testing Methodology



4. Key Decisions and Benefits

The final step involved translating the results of our testing into actionable decisions that improved hardware provisioning at CNES NOC. Based on our refined testing methodology, we found several key adjustments to our virtualized infrastructure, leading to more efficient and effective resource allocation.

Optimized Resource Allocation:

Based on the data collected during our testing phase, we recalibrated the resource requirements of virtualized system to more closely align with actual needs. By fine-tuning CPU, RAM, and disk configurations, we ensured that hardware resources were fully utilized, without over-provisioning. This optimization ensured that we were not wasting valuable resources—whether it was excess power, memory, or storage capacity.

Cost Reduction:

Through the more accurate sizing of resources, we significantly reduced both hardware and licensing costs. With fewer physical machines needed, and with virtual resources allocated precisely, we avoided the unnecessary expenses associated with massive over-provisioning. This change helped us optimize our IT budget while still maintaining high system performance.

Performance Confidence:

Having validated the optimal hardware configurations through rigorous testing, we were able to implement these changes with confidence. Each virtualized system is now properly supported by its underlying infrastructure, providing reliable and predictable performance for both critical operations and everyday tasks.

Flexibility:

Our approach to hardware provisioning has increased our ability to scale and adapt more easily as new software requirements emerge. Rather than rigidly adhering to the old model of over-provisioning resources, the new system allows us to fine-tune configurations on the fly. This means that as new tools and software are introduced, we can adjust our hardware resources accordingly, ensuring that performance needs are met without any excess.

4.1 Ecological Benefits of Tailored Hardware Sizing

In addition to the tangible technical and financial benefits, the improvements in hardware provisioning at CNES have also led to significant ecological advantages:

Energy Efficiency: Optimized hardware configurations typically consume less power, as systems are not running under unnecessary loads. Reducing the number of over-provisioned systems meant that fewer servers were required, leading to lower overall energy consumption, which directly contributes to a smaller carbon footprint for our data centers.

Lower Cooling Demands: Optimized systems not only consume less power but also produce less heat. As a result, our cooling requirements were reduced, leading to a decrease in the energy needed to maintain proper operating temperatures. This translates into lower energy consumption not only for the servers themselves but also for the cooling systems that maintain ideal operational conditions.

Extended Lifecycle of Physical Hardware: Tailoring our provisioning strategy has also helped extend the life of physical hardware. Systems that are properly sized are less likely to experience overheating or excessive wear and tear. This extends the operational lifespan of the physical infrastructure, delaying the need for costly hardware replacements and reducing electronic waste. The careful management of hardware resources aligns with sustainability goals by reducing the frequency of hardware upgrades and promoting the reuse of existing assets.

Reduced Manufacturing and Disposal Impact: By optimizing hardware provisioning, we've minimized the need to acquire additional physical servers, which in turn reduces the demand for raw materials and energy involved in manufacturing new hardware. This decreases the environmental impact associated with hardware production and contributes to a lower rate of electronic waste generation. Furthermore, by making use of existing equipment for a longer period, we decrease the frequency at which older servers must be decommissioned and disposed of, helping to reduce the environmental impact associated with the disposal of obsolete technology.

6. **Conclusion:** Why We Keep Control Over Our Infrastructure

Virtualization has brought CNES significant operational, financial, and ecological benefits, it has never been our intention to follow industry trends blindly, particularly when it comes to outsourcing infrastructure to the cloud.

One of the main reasons we choose to keep our own hardware is that **operational availability is non-negotiable** in our context. Our systems support critical operations, often in coordination with international partners and space assets. These operations require guaranteed uptime, deterministic behaviors, and full control over recovery scenarios conditions that are not always compatible with cloud services, which may involve shared resources, uncertain latencies, scheduled maintenance, or provider-side outages beyond our control.

Additionally, we prefer to rely on well-known, proven, and stable platforms systems we have operated, tested, and hardened over time. In our field, predictability and long-term support are more valuable than fast innovation cycles. Our platforms must remain stable for years, sometimes even decades, and many newer cloud-native tools or platforms do not offer the lifecycle, transparency, or determinism we require.

By maintaining our own infrastructure and carefully selecting technologies we fully understand and control, we ensure that our operations are resilient, secure, and optimized for our very specific constraints. This strategy has proven effective, especially as we continue to evolve our virtualized environment at our own pace, based on operational need, not market trends.

In addition, while modern trends often push toward container orchestration platforms like **Kubernetes**, we have not identified a concrete operational benefit in adopting this technology for our current needs. Kubernetes is a powerful solution for managing large-scale, dynamic, cloud-native applications with frequent deployments and horizontal scalability needs. However, our systems are designed for **stability, predictability, and long-term operation**, often with low change frequency and strong lifecycle constraints. Introducing Kubernetes would add significant complexity, both in terms of deployment and maintenance, without providing clear functional gains for our operational environment. For us, **robustness and simplicity** remain higher priorities than aligning with industry fashions.