



The SVOM French Science Center Infrastructure

H. Louvin¹ , D. Corre¹ , A. Formica¹ , L. Jouvin¹, K. Tazhenova¹, T. Sadibekova² , J. Palmerio¹ , N. Dagoneau¹ , A. Sauvageon¹, A. Claret¹ , F. Agneray³, C. Moreau³, T. Fenouillet³, Y. Roehly³, J-C. Thome³, L. Michel⁴ , P. Maggi⁴ , L. Kleiver⁴, P. Maeght⁵, U. Jacob⁵, F. Piron⁵ , P. Bacon⁶ , N. Bellemont⁶, C. Cavet⁶ , F. Daigne⁷ , F. Lacreu⁷, G. Tcherniatskiy⁷ , J. Wang⁷, Y. Canton⁸ , I. Jegouzo⁹, N. Leroy¹⁰ , S. Du¹⁰, S. Lion¹⁰ , R. Le Montagner¹⁰ , C. Van Hove¹⁰, J. Alaux¹¹, M. Boiziot¹¹ , and T. Auphan¹²

¹ CEA Paris-Saclay, Institut de Recherche sur les lois Fondamentales de l'Univers, 91191 Gif sur Yvette, France; henri.louvin@cea.fr

² Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191 Gif-sur-Yvette, France

³ Aix Marseille Univ, CNRS, CNES, LAM, Marseille, France

⁴ Université de Strasbourg, CNRS, Observatoire Astronomique de Strasbourg, UMR 7550, 67000 Strasbourg, France

⁵ Laboratoire Univers et Particules de Montpellier, Université Montpellier, CNRS/IN2P3, 34095 Montpellier, France

⁶ Université Paris Cité, CNRS, CEA, Astroparticule et Cosmologie, 75013 Paris, France

⁷ Sorbonne Université, CNRS, UMR 7095, Institut d'Astrophysique de Paris, 75014, Paris, France

⁸ LUX, Observatoire de Paris, Université PSL, Sorbonne Université, CNRS, 92190 Meudon, France

⁹ UNIDIA, Observatoire de Paris-PSL, CNRS, 92190 Meudon, France

¹⁰ Université Paris-Saclay, CNRS, IJCLab, 91405, Orsay, France

¹¹ IRAP, Université de Toulouse, CNRS, CNES, Toulouse, France

¹² Aix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France

Received 2025 December 17; revised 2026 March 4; accepted 2026 March 17; published 2026 October 6

Abstract

At the heart of the Space-based multi-band Variable Objects Monitor (SVOM) French ground segment, the French Science Center (FSC) is a cloud-based platform which provides services and tools for the management, storage, scientific processing and visualization of SVOM data for the French community. This digital center is a critical node of the SVOM system since it is the single point of access to SVOM scientific data for the French community and the only component of the ground system connected to the SVOM VHF network allowing near-real time communication from the satellite. Scientific processing pipelines are fully integrated into its infrastructure, allowing the automated production of high-level scientific data and the generation and broadcast of alerts to the scientific community at large. The software components of the FSC are running 24/7 and thus require a high level of automation, which led to the development of dedicated software relying on modern technical solutions such as micro-services, application containerization, infrastructure-as-code and continuous integration and deployment. In this paper, we describe the FSC infrastructure design, technological choices, and the process of SVOM data ingestion, archiving, and automated processing by the FSC scientific pipelines. We present analysis on the FSC performances in terms of availability, amount of data processed as well as processing speed.

Key words: astronomical databases: miscellaneous – methods: data analysis – (stars:) gamma-ray burst: general

1. Introduction

Space-based multi-band Variable Objects Monitor (SVOM) is a French-Chinese space-based mission dedicated to the detection, localization and multi-wavelength study of high energy transient astronomical events, with a focus on the most distant stellar explosions: gamma-ray bursts (GRBs). The mission is composed of two sets of instruments: the first is integrated onboard a satellite, while the second is located on Earth (Cordier et al. 2026b).

Figure 1 shows the global organization of the SVOM mission ground segment elements. The French Mission Ground Segment, composed of the blocks shown in the left side of the figure, is divided in six main components:

1. the Alert network, the global network of VHF-band radio receivers allowing real-time reception of satellite data, specifically but not limited to GRB detection alerts (Cordier et al. 2026a);

2. the French Payload Operation Center (FPOC), in charge of the management of the French payload;
3. the three Instrument Centers (ICs), each dedicated to a French instrument, in charge of its monitoring, control, configuration and calibration:
 - a. the ECLAIRs Instrument Center (EIC) for the ECLAIRs onboard telescope (Godet et al. 2026);
 - b. the MXT Instrument Center (MIC) for the MXT onboard telescope (Götz 2026);
 - c. the French-Mexican Ground Follow-up Telescope (FM-GFT) Instrument Center (GIC) for the COLIBRÍ follow-up telescope (Basa et al. 2026).
4. The French Science Center (FSC), which is the subject of this paper.

As the central node of the French Mission Ground Segment, the FSC is expected to carry out the following tasks:




Figure 1. French-Chinese scientific ground segment.

1. Handle the reception and storage of the real-time monitoring and alert data coming from the VHF alert network in near-real time for all instruments (Cordier et al. 2026a);
2. Handle the reception and storage of the complete observation and housekeeping telemetry data from the French instruments;
3. Automatically produce high-level SVOM scientific data for all SVOM observation programs;
4. Generate and broadcast alerts to the scientific community following the onboard detections of GRBs or other transient sources of interest;
5. Provide services, interfaces and tools that allow the other SVOM centers, the members of the SVOM collaboration and the scientific community at large to visualize and access the SVOM data that they are privy to.

According to the French-side SVOM mission requirements, any software developed to perform these tasks must comply with technical requirements with regards to development workflow, versioning, continuous integration, code quality and security. The software infrastructure itself must provide high availability and storage replication. Performance-wise, the FSC is expected to distribute SVOM alert coordinates to the SVOM ground-based telescopes less than 30 s after detection for 65% of the alerts and less than 20 minutes for 95% of the alerts.

The definition of the FSC technical perimeter started 9 yr prior to SVOM launch. Since then, the development and maintenance of the FSC infrastructure and software components is under the responsibility of a team of approximately 40 software engineers and scientists distributed in eleven French national laboratories:

APC, CEA/Irfu/DAP, CEA/Irfu/DEDIP, CPPM, IAP, IJCLab, IRAP, LAM, LUPM, LUX and ObAS.

The software solution that was chosen consists of a cloud-based platform centralizing all the software modules required to perform the FSC activities, including but not limited to: a database server, services handling raw data ingestion, the scientific processing pipelines producing high-added value scientific products, orchestration services handling the automation of data processing chains, and web-based visualization interfaces. This solution allows the system to reach the high level of automation required to meet the requirements and allows the SVOM collaborators to interact with ease with all the scientific data and tools available at FSC. The software infrastructure, the FSC components, their interactions and their performances are described in the following sections of this paper.

2. FSC Infrastructure

2.1. Architecture Principles

The FSC architecture is service-oriented: FSC applications are dedicated to single tasks, which might be of high complexity but are always as limited as possible in size and scope. All software components are fully containerized using Docker¹³ for easy deployment in the cloud, where they are orchestrated in clusters of Docker engines using Docker Swarm mode. There are three FSC cloud instances:

1. the *integration* instance, where all code modifications are deployed continuously for testing;
2. the *pre-production* instance, where the new stable versions of software are tested in a controlled environment before being authorized to be put in production;
3. the *production* instance, where the latest FSC version is running and used by the scientific SVOM community.

2.1.1. Infrastructure as Code

The FSC instances are deployed on OpenStack¹⁴ clouds provided by national research centers. The integration and pre-production instances are hosted by IJCLab (Irene Joliot Curie Laboratory in Orsay France) on their Virtual Data¹⁵ and the production instance on the cloud services of the French national research computing center CC-IN2P3.¹⁶

The software infrastructures for all the FSC instances are provisioned and managed through scripts (rather than physical hardware configuration or interactive configuration tools). This process is known as “infrastructure as code” (IaC) and allows to fully automatize the creation of the software infrastructure for each of the FSC instances, including the creation,

¹³ <https://docs.docker.com/build/>

¹⁴ <https://www.openstack.org/>

¹⁵ <https://openstack.ijclab.in2p3.fr/activites/cloudvd/>

¹⁶ <https://doc.cc.in2p3.fr/en/Hosting/cloud-iaas.html>

instantiation and configuration of the virtual machines (VMs), networks, security, storage, databases, etc.

This IaC approach is implemented at FSC using a combination of Terraform¹⁷ and Ansible¹⁸ scripts, the former for provisioning the immutable OpenStack cloud infrastructure (network, VMs, volumes, security groups) and the latter for configuring the VMs.

2.1.2. Software Development

The software modules composing the SVOM FSC were created and are maintained by development teams distributed in several French national laboratories: APC, CEA/Irfu/DAP, CEA/Irfu/DEDIP, CPPM, IAP, IJCLab, IRAP, LAM, LUPM, LUX and ObAS. In addition, the FSC integrates a data processing module that is provided by a team from NAO, China. In order to ensure sharing, control and homogeneity among the productions of the different teams, the source code of all FSC software is saved and versioned on a GitLab¹⁹ instance.

The number of programming languages used is limited: Python3²⁰ for processing pipelines and most backend services, Java²¹ for raw data management modules, and JavaScript (React²² and Angular²³) for web-based applications. In addition, C++ is also used as a built-in module in one of the Python libraries developed by the ECLAIRs trigger team, which allows to use the methods extracted from the ECLAIRs onboard software.

We make an intensive use of the GitLab CI/CD²⁴ continuous development tools, and follow a common git workflow as well as common continuous integration procedures to automatize the build, test, package delivery and deployment of FSC applications. All FSC software development teams are therefore using the same tools with the same configurations for:

1. code linting, i.e., proper formatting (pylint²⁵, ruff²⁶, black²⁷);
2. unit testing (pytest²⁸, jest²⁹);
3. code quality and vulnerability scans (SonarQube³⁰, Trivy³¹);

¹⁷ <https://terraform.io>
¹⁸ <https://www.ansible.com/>
¹⁹ <https://about.gitlab.com/solutions/source-code-management/>
²⁰ <https://www.python.org/>
²¹ <https://www.java.com/>
²² <https://react.dev/>
²³ <https://angular.dev/>
²⁴ <https://about.gitlab.com/solutions/continuous-integration/>
²⁵ <https://www.pylint.org/>
²⁶ <https://docs.astral.sh/ruff/>
²⁷ <https://pypi.org/project/black/>
²⁸ <https://docs.pytest.org/>
²⁹ <https://jestjs.io/>
³⁰ <https://www.sonarsource.com/products/sonarqube/>
³¹ <https://trivy.dev/latest/>
³² <https://www.docker.com/>

4. container building and pushing (Docker³², Kaniko³³, Harbor³⁴);
5. documentation generation (Sphinx³⁵);
6. changelogs generation (towncrier³⁶).

2.1.3. Data Storage

Each FSC cloud instance has a dedicated PostgreSQL³⁷ database server, where all data received at FSC (i.e., raw telemetry and calibration data) and all data generated by FSC processings are stored. In addition to SVOM data, SQL databases are also used to store FSC software images, authentication data and software code quality profiles. In addition to the SQL databases, a persistent storage disk space is used to save file-like data that are supposed to be kept during the whole lifespan of the mission and beyond. This space is stored on a Ceph file system.³⁸

The PostgreSQL and the persistent storage are backed up every 24 hr, and a replication of the PostgreSQL database is in place in another facility than the production site. The volume of data saved on FSC databases and persistent storage at the time of this paper's redaction on 2025 September 16 is presented in Table 1.

2.2. Centralized Messaging

The FSC is composed of a multitude of distinct services. Those services have two ways of communicating with each other, depending on the number of services involved in the communication:

1. one-to-one communications are performed through REST API interfaces using standard HTTP protocol;
2. one-to-many communications are made using a centralized messaging system.

For one-to-many communications, the FSC centralized messaging system is a publish/subscribe message broker hosted at FSC and relying on the message-oriented middleware NATS.³⁹

NATS is an open-source publish-subscribe messages broker:

1. publishing services send messages to the server on pre-defined subjects identified by a string;
2. subscribing services are connected as listeners to one or several subjects;
3. the server broadcasts all messages received on a subject to all subscribers of that subject.

³³ <https://github.com/GoogleContainerTools/kaniko>
³⁴ <https://goharbor.io/>
³⁵ <https://www.sphinx-doc.org/en/master/index.html>
³⁶ <https://towncrier.readthedocs.io/>
³⁷ <https://www.postgresql.org/>
³⁸ <https://ceph.io/en/>
³⁹ <https://nats.io/>

Table 1

FSC Data Usage (estimated 2025 September 16). File-like Objects are Stored on Disk, while All Other Data is Stored in SQL Databases

Data Type	DB Size	Disk Usage
Science data (SDB)	905 MB	6.1 TB
X-Band raw data	609 GB	1.2 TB
VHF raw data	102 GB	31 GB
Instruments configurations and housekeeping data	94 GB	83 MB
Instruments calibrations	...	99 MB
SVOM detection alert notices	30 MB	...
System monitoring data	87 MB	190 GB
Scientific pipelines orchestrator	4.6 GB	...
Burst Advocate web interface	2.6 GB	1.1 GB
Applications Docker images	926 MB	130 GB
FSC Authentication	6.5 GB	...
Code quality analysis data	327 MB	...

In order to ensure the distribution of messages to subscribers even if they are not connected at the time the message was sent, the FSC NATS server makes use of the built-in distributed persistence layer of NATS called JetStream.⁴⁰ This configuration allows the FSC services subscribed to the NATS messaging system to be safely restarted, for example when deploying a software update. All notifications missed during the service downtime will be recovered upon restart.

The Chinese Science Center uses the same architecture, but with a different message-oriented middleware called RabbitMQ⁴¹ that relies on the MQTT protocol. The messaging systems are key components of the interface between FSC and the Chinese centers. A couple of services running at FSC are connected to the CSC MQTT server in order to be notified about the availability of new data (more on that in Section 3.2).

3. Raw Data Management

3.1. VHF Raw Data Management

The FSC is the only SVOM component connected to the VHF alert network, which allows it to receive and analyze limited-size communications from the satellite in near-real time (Cordier et al. 2026a). The other SVOM centers such as the ICs or Chinese centers do not have a direct connection to the VHF network and therefore retrieve VHF data from the FSC.

The VHF data stream consist of 94-byte encoded messages or “packets”, broadcasted every 1.83 s from the onboard VHF radio transmitter. The messages are received on the ground by a network of VHF radio antennas which demodulate the signal and send the contents of the digitized binary messages to the FSC through a REST API. The packets are received by the so-

called “vhf-manager” service, which saves them in a dedicated database. The VHF raw data ingestion entails several steps:

1. filtering of duplicated messages, which are due either to the simultaneous reception by multiple antennas of a single packet or to the satellite repeating the broadcast of the same packet several times;
2. validation of the packet content integrity via checksum verifications;
3. unpacking of the binary VHF packet content using a dedicated FSC service that is responsible for reading the binary VHF packets and converting them in JSON format;
4. computing of on-ground unified identifiers depending on the packet content (GRB-related data only);
5. saving of the interpreted VHF packet content in the VHF database;
6. broadcast of notifications regarding the availability of new VHF data to other FSC components using the centralized messaging system (see Section 2.2).

In order to improve the reliability of SVOM communications for alert-related data, some of the messages sent through the VHF network are also downloaded using the Chinese navigation satellite system Beidou. This communication subsystem has a slower rate than VHF and a much smaller bandwidth but it allows to speed up the on-ground reception of alert data in some cases, and could even be vitally important in case of VHF network temporary unavailability.

The FSC provides an authenticated REST API which allows French and Chinese SVOM centers as well as authorized SVOM collaborators to access the VHF data stored in the VHF database.

3.2. X-Band Raw Data Management

SVOM uses X-Band communications to download the complete observation and housekeeping telemetry data from the satellite. Unlike SVOM VHF telemetry, the X-Band telemetry reception and distribution is under the responsibility of the Chinese National Space Science Data Center (SSDC). The SSDC receives all telemetry data that the satellite dumps through X-Band antennas several times per day—every 2–6 hr depending on the satellite orbit. The telemetry concerning French instruments is then uploaded to a dedicated FTP server.

The availability of French payload X-Band data is notified to FSC through the MQTT messaging server (see Section 2.2). Upon reception of the notification by the dedicated service at FSC, the X-Band telemetry data is downloaded from the SSDC server and sent to the “xband-manager” service for storage and processing. The data files are stored on the persistent data storage file system, and their metadata are saved in a dedicated database.

⁴⁰ <https://docs.nats.io/nats-concepts/jetstream>

⁴¹ <https://www.rabbitmq.com/>

The X-Band files that are received at FSC contain chains of 1022 byte binary messages, which are extracted from each file by the xband-manager. Due to the sheer volume of data that is received through the X-Band channel, it is inconceivable for the xband-manager to interpret the data on the fly and store JSON values in the database in the same way that it is done for VHF data. Only the first few bytes of each X-Band packet are decoded to extract metadata (such as the type of data, the origin instrument, the time of packet production and the current observation number) after which both the metadata and the binary packet are saved in the database, where all the other FSC components can safely access them through an authenticated REST API. Notifications of availability are sent through the centralized messaging system (see Section 2.2) so that the automated X-Band data processing can start.

3.3. Ancillary Data Management

In addition to the raw SVOM instrument data, the SVOM ground centers also exchange mission ancillary data files (mostly in JSON, CSV and TXT formats). These files include the mission observation work plan, statistics on executed observations, predicted ephemeris for the spacecraft as well as derived information such as planned X-Band passes and South Atlantic Anomaly flybys.

Such files are regularly sent to FSC by the FPOC using CREST, an FSC-hosted service in charge of their reception, archiving and provision. CREST is a file-indexing database allowing for an accurate tracking of time-dependent data, with an associated REST API for uploading and downloading files. The codebase of this service originates from a development in the High Energy Physics community for the management of Conditions Data (Laycock et al. 2018).

The FSC CREST service is also used as the main file exchange interface between the French ICs and FPOC for instrument configuration files, health and alarm reports, onboard software updates, and more.

4. Scientific Data Processing

Scientific data processing at FSC is performed by dedicated services referred to as “pipelines.” There are currently 13 data processing pipelines running continuously at FSC, totaling 73 execution modes and about 180 different output types. In this section we describe how these services are interfaced with the other parts of the system and how they are automatically triggered to process data in real-time, with a focus on the alert processing chain and its performances. The inner workflow and data analysis algorithms of the highest-level scientific pipelines operated at FSC are described in dedicated papers (Goldwurm et al. 2026; Maggi & Michel 2026; Piron et al. 2026; Wu & Palmerio 2026).

4.1. Data Models

The FSC workflow comprises many modules that have been developed separately by different contributors. These modules must be able to communicate with each other and process the scientific products that circulate within the infrastructure. To ensure consistency in such a heterogeneous architecture, a high level of interoperability is achieved between components at three levels, all of which are based on JSON resources (Michel et al. 2020).

1. *vocabulary*. All properties contained in the messages exchanged by different modules are named according to a common vocabulary defined in a JSON schema.
2. *message format*. Messages exchanged via NATS/JetStream or REST endpoints are validated against specific JSON schemas.
3. *product descriptors*. All FITS files that may be stored in the Science Database (SDB) are described in JSON files. These descriptors provide accurate descriptions of quantities that can be used to retrieve specific products as well as various metadata. Each individual FITS file is validated against its descriptor before being ingested by the SDB and any discrepancy interrupts the ingestion process.

In addition, most services expose a common REST API, as detailed in Section 4.4.1.

4.2. Science Products Database

The SDB is the database storing all scientific products—in FITS format (Pence et al. 2010)—and related ancillary files produced by the FSC scientific pipelines for all SVOM observation programs. The SDB provides REST APIs allowing users to import and export scientific products.

New data is sent to the SDB by the FSC pipelines using the SDB import API, which includes REST access points for importing new products and for updating existing ones. All submitted products are analyzed before being recorded, in order to validate their format and version according to the corresponding product data model (see Section 4.1). Ancillary files can only be imported if they are properly associated to data products already present in the database. Whenever a scientific product is successfully imported into the SDB, its availability is notified to the other FSC services using the NATS/JetStream centralized messaging system.

The data provision of the SDB relies heavily on the Astronomical Information System (ANIS; Agneray & Moreau 2019), a generic tool aimed at facilitating and homogenizing the implementation of astronomical data of various kinds and catalogs in dedicated Information Systems. ANIS includes a REST API allowing the FSC scientific pipelines to extract data from the SDB according to various search criteria, and a dedicated web interface through which SVOM collaborators can search, visualize and download data from the SDB.

4.3. Calibration Database

The FSC Calibration Database (CalDB) is built upon a regular HEASARC CalDB instance for FITS files (George & Corcoran 2005) extended with a specific capability for supporting YAML files in certain conditions. It provides a REST API to import and export SVOM calibration files, and a dedicated web interface allowing SVOM collaborators to query it.

The HEASARC's CalDB system stores and indexes datasets associated with the calibration of high energy astronomical instrumentation. The system can be accessed by users and software alike to determine which calibration datasets are available, and which should be used for data reduction and analysis. The FSC CalDB instance is restricted to a single mission—SVOM—and four instruments—ECLAIRS, MXT, GRM and VT.

In addition to the regular FITS files storage, the SVOM instance of the CalDB can also manage YAML format. YAML are all stored in a specific node contained onboard (one per instrument). The use case for the YAML storage is to trace the configuration files that have been uploaded onboard and used by the onboard software. It is not possible for ICs to predict the moment at which a given configuration is actually applied onboard. The ICs have to rely on the telemetry to know when the onboard configuration changes. When this is the case, the IC operators upload the corresponding YAML to the CalDB with a time stamp corresponding to the onboard application time.

4.4. Data Processing Automation

Scientific data processing at FSC is for the most part automated. The sequential execution of the various processing pipelines composing the chains is under the supervision of the so-called orchestrator service. All FSC services communicate with each other either via direct REST API calls or through notifications sent using the centralized messaging server described in Section 2.2.

The complete sequence of services involved in the automated processing of VHF telemetry data is schematically shown in Figure 2. Figure 3 displays the same diagram for the automated processing of X-Band telemetry data.

4.4.1. Data Processing Pipelines

The FSC data processing pipelines are the FSC services responsible for processing SVOM scientific data and generating output products. These output products are most of the time FITS files that are sent to the SDB for storage as soon as they are produced. All the FSC pipelines are continuously running services in the FSC cloud. They are mostly idle and will start processing data upon reception of a processing request, either from the automated FSC orchestrator or by an FSC user that can request a processing manually through the FSC website.

In order to ensure that all pipelines and their REST API are consistent with the common data model (see Section 4.1), a

generic task scheduler has been developed at FSC under the name “pipeline-bricks.” All FSC pipelines use this scheduler and the associated REST API as a wrapper around the actual pipeline software. The pipeline-bricks principles are as follows:

1. One specific bricks instance always runs the same task sequence (namely the workflow).
2. This sequence can be marginally modified in reprocessing mode.
3. Each task must be made available as an individual executable command. The underlying programming language has no impact.
4. Bricks instances can only run one sequence at a time. Any processing request received while a sequence is running will return an error.
5. Tasks are executed one by one without parallelization.
6. A processing sequence executes all tasks of the workflow. There is no conditional execution.
7. Bricks is stateless in the sense that it has no memory of previous processing sequences or knowledge of the global context.

Bricks has no control on the task activities or the files that are imported or exported. It can however get some information from a runtime file that is shared between all tasks, provided that the core software writes information into it. This mechanism can allow bricks for example to provide through its REST API the list of generated science products for a given processing. This also allows bricks to automatically add VO data in some science products.

4.4.2. Orchestrator Service

In order for data to be dynamically processed when available, automated processing requests must be sent to FSC pipelines through the pipeline-bricks REST API. The choice has been made to centralize these executions in a dedicated service developed at FSC, dubbed the FSC orchestrator. This central service is in charge of:

1. reacting when new data is available to be processed;
2. sending processing request(s) to the pipeline(s) in charge of processing the data;
3. handling processing queues if a pipeline is currently busy processing other data;
4. keeping track of which pipeline has been triggered, when and with which data;
5. monitoring ongoing processings, retrieving their logs and storing them in a database;
6. updating the logs and the final status of processings when they are over.

Whenever new data are available in any of the various databases at FSC, a data notification is broadcasted on a




Figure 2. Complete automated FSC processing chain for SVOM VHF alert telemetry, from the VHF telemetry reception to the final scientific products and alert broadcast.




Figure 3. Complete automated FSC processing chain for SVOM X-band telemetry, from the X-band telemetry reception to the final scientific products.

specific NATS/JetStream subject with all the information required for the orchestrator to identify the type of data and the scientific observation to which it relates. Each database in the FSC ecosystem has its own publication subject on the NATS/JetStream server. A dedicated orchestrator service is connected to each of those message streams. Upon reception of a data notification, the corresponding orchestrator:

1. determines if one (or several) processing operation(s) should be launched to treat the new data;
2. sends the corresponding processing request(s) to the relevant pipeline(s) or service(s);
3. monitors the evolution of the launched processing(s);
4. takes actions upon the processing(s) termination: return status and log retrieval, retry scheduling, etc.;
5. stores relevant data concerning the processing(s) in a database.

Most orchestrated processings terminate by sending processed data to the SDB or another FSC database. If and when this happens, the corresponding database sends a data

notification on the NATS/JetStream server which might in turn trigger a new orchestrated processing. This allows sequentially execution of processings with a fine-grain monitoring on their inputs and outputs, producing higher level scientific products each step of the way.

The pipeline execution process by the orchestrator services is illustrated in Figure 4.

4.5. Alert Data Processing

The so-called “core” observation program (CP) of SVOM is dedicated to GRB detection and study. GRBs are astronomical transients that can be very short lived, and the global scientific community must be warned as soon as possible of the detection of such phenomena in order for their study to be as extensive as possible. The onboard detection of GRBs by either one of the SVOM instruments having triggering abilities is received on the ground through so-called “alert” VHF messages providing location and other information on the detection. After that, “alert sequence” VHF messages are sent

Table 2
FSC Data Processing Statistics from 2024 June 22 (SVOM Launch) to 2025 September 30

Data Type	Count Since Launch	2025 Monthly Average
VHF Packets received and processed	30,645,621	1,862,060
X-Band raw data files received	31,850	2,242
X-Band raw packets processed	883,039,648	64,927,121
Instruments housekeeping telemetry data	158,666	10,541
Orchestrated processings launched	405,761	31,115
Scientific products generated	938,250	74,019
VOEvent notices generated (since 2025-01)	2420	242
VOEvent notices public broadcasts (since 2025-01)	447	45




Figure 4. Orchestrated pipeline processing flowchart.

by one or more of the SVOM instruments containing more detailed observation data of the event.

Since the FSC is the only SVOM center directly connected to the VHF network, the processing pipelines that are dedicated to real-time analysis of SVOM VHF alerts are hosted at FSC. One of the responsibilities of FSC is to automatically generate “notices” in VOEvent format regarding the SVOM GRB detections and broadcast them to the worldwide scientific community in near-real time. The alert data processing at FSC consists of the following steps:

1. the VHF manager receives the alert message through the VHF network;
2. the packet decoder service decodes the alert message content;
3. the VHF manager sends a notification on NATS/JetStream to signal availability of the alert data;

4. the orchestrator receives the notification and triggers a first set of pre-processing pipelines;
5. the pre-processing pipelines generate the first scientific products and import them in SDB;
6. the SDB saves the products and sends a notification on NATS/JetStream to signal their availability;
7. the orchestrator receives the notification and triggers the notices generation pipeline;
8. the notices pipeline generates the alert notice and sends them to the alert manager service;
9. the alert manager service saves the notice in a dedicated database;
10. the alert manager broadcasts the notice to the follow-up telescopes and the worldwide community using the NASA GCN network.

Overall, a dozen distinct FSC services are involved in the processing of alert data and the generation and broadcast of the VHF alert notices. The public SVOM VOEvent notices are broadcasted using the GCN network. Details can be found on the GCN website⁴² and all public SVOM VOEvents are available on the FSC public alerts page.⁴³

4.6. Statistics and Performances

4.6.1. Data Processing Statistics

A few statistics on the quantity of data received and produced at FSC are gathered in Table 2. The numbers presented in this table cover the period from the launch on 2024 June 22–2025 September 30, with the exception of VOEvent notices production which only started in 2025 January. The corresponding volume of data in bytes can be found in Table 1, discussed in Section 2.1.3.

The second column of Table 2 shows the number of items received or generated between 2024 June 22 and 2025 September 30. The rightmost column lists the monthly average

⁴² <https://gcn.nasa.gov/missions/svom>

⁴³ <https://fsc.svom.org/alerts/>




Figure 5. SVOM notices broadcast delays for the alerts from 2025 September colored by instrument.

since the end of the FSC commissioning period on 2025 January 1st. Since then, the system has been running continuously and processing data at a high rate, with on average 31 thousand orchestrated processings generating 74 thousand scientific products every month. The efficiency of the data processing performance is studied in the following section.

4.6.2. Data Processing Performance

As presented in Section 4.5, the alert processing at FSC involves a lot of services, among which the centralized messaging system, the orchestrator services and several pipelines and databases. Therefore, the study of the FSC alert processing performance allows us to evaluate the overall FSC architecture efficiency.

Figure 5 shows the notices broadcast delays for the SVOM alerts from 2025 September. The delays are computed as the time difference between the alert message generation onboard and the broadcast of the first automated SVOM notice to the SVOM follow-up telescopes, which covers all the processing steps presented in the previous section. We can observe that for more than 50% of the alerts, the generated notice is broadcasted less than 30 s after the alert is detected onboard, and all the alerts were broadcasted in less than 15 minutes. The secondary peak centered around 2 minutes is due to alerts that were received in a delayed manner at FSC because they were not transmitted through the VHF network but through the secondary Beidou communication subsystem which is a bit slower. This is visible in Figure 6.

In order to evaluate the actual FSC processing time and disregard the impact of board-to-ground communication delays, we show in Figure 7 the time difference between the FSC reception time of the alert message and the notice broadcast for the alerts displayed in Figure 5. The processing time at FSC is very stable and never exceeded 45 s, with a median value around 14 s. We observe a significant difference between GRM alerts and ECLAIRs alerts which are due to




Figure 6. SVOM notices broadcast delays for the alerts from 2025 September colored by reception channel.




Figure 7. FSC notices processing times for the alerts from 2025 September.

differences in the data processing. The ECLAIRs alerts undergo a few more processing steps before the notice is created and broadcasted, which allow for extra on-ground filtering and automated identification of false alerts.

5. Data Visualization Tools

The current state of FSC and its performances are a reflection of the work of the software development teams participating in the FSC development. Aside from the services already described in the previous sections, a significant amount of software developed for FSC are web-based user interfaces that allow FSC developers, FSC administrators and SVOM scientists to visualize SVOM data, FSC processings and monitor the FSC infrastructure. A total of 23 dedicated user interfaces have been




Figure 8. The FSC website home page [internal].




Figure 9. The FSC orchestrator web interface [internal].




Figure 10. The FSC CalDB web interface [internal].

developed and are maintained at FSC. A few selected examples of such interfaces are described in the following sections.

5.1. FSC Website

All the web interfaces developed for FSC operation are gathered under a single proxy with a centralized authentication system provided by a Keycloak⁴⁴ instance. All web interfaces are deployed like any other FSC services on the production

cloud and are accessible to all SVOM/FSC members through a common portal shown in Figure 8.

Some of those web interfaces are dedicated to visualizing data related to FSC operations and the behavior of common services such as the centralized messaging system or the orchestrator services (see Figure 9). Each FSC database has an associated web interface service that allows visualizing the database content, downloading data and in some cases and with the proper access rights even uploading data (see Figure 10). Some specialized web pages are focused on SVOM instrument health monitoring by the ICs.

5.2. Infrastructure Monitoring Interfaces

Most of the FSC user interfaces that are meant to be used by FSC operators and SVOM scientists were specifically developed from scratch to fit the SVOM mission needs. However, the user interfaces used by the FSC administration team to monitor the infrastructure mostly rely on existing tools, such as Grafana⁴⁵ which is used for graphical displays of data coming from application logs aggregated using Grafana Loki⁴⁶ or time-series system metrics collected using Prometheus.⁴⁷ Figure 11 shows an example of an FSC Grafana dashboard allowing administrators to monitor the health of the FSC instance and services.

The FSC Grafana instance is also used to display data stored in the various FSC databases, which allows to create real-time monitoring user interfaces for the mission such as illustrated by Figure 12.

One of the upsides of the use of Grafana as a monitoring tool is the use of Grafana “Alert rules” which are used to regularly check specific metrics related to FSC dashboards and fire notifications when thresholds are reached. These are used as the FSC main alarm system in the event of disk saturation, memory or CPU overload, missing data, services interconnection loss, etc.

5.3. Science Data Visualization

The main objective of FSC is to generate scientific products and provide tools and services to access and visualize them, and a great deal of web user interfaces developed and maintained at FSC are related to the display of scientific data from the SVOM mission. Those web interfaces and associated tools are described extensively in a dedicated paper (Claret et al. 2026). An example of one the central web interfaces for those tasks is shown in Figure 13.

Even if most of the web interfaces developed and maintained at FSC are internal and restricted to the SVOM collaborators, a few of them are also publicly accessible to the scientific community at large. Links to those public pages are available on the FSC homepage at <https://fsc.svom.org>.

⁴⁵ <https://grafana.com/>

⁴⁶ <https://grafana.com/docs/loki/latest/>

⁴⁷ <https://prometheus.io/>

⁴⁴ <https://www.keycloak.org/>




Figure 11. FSC web interface displaying application logs, powered by Grafana [internal].




Figure 12. FSC web interface displaying real-time mission monitoring data powered by Grafana [internal].

6. Conclusion

The SVOM FSC cloud-based platform is the result of the collaborative work of 10 national French research laboratories, and the central node of the SVOM French Mission Ground Segment. It allows, since the launch of SVOM in 2024 June, to perform the archiving, monitoring and scientific analysis of SVOM data for the French community. All the output products and associated visualization tools that were developed for FSC are accessible to all co-investigators of the SVOM collaboration.

The innovative technological choices that were made in developing the FSC proved most efficient in dealing with the

flow of SVOM data. With a routine average of 31 thousand automated processings generating 74 thousand scientific products every month, and web interfaces accessible to all SVOM collaborators to access the processed data in real-time, the FSC fulfils its role and allows the scientific needs of the SVOM mission to be met. Thanks to the dedication of the developer teams involved in the FSC development, maintenance and evolution, the FSC already showcases a level of automation that is unprecedented at the beginning of an astrophysical space observatory lifespan.




Figure 13. The FSC real-time GRB data visualization web interface, dubbed “iFSC-Tools” [internal].

Contributors

CEA/Irfu¹: A. Claret, M. Bocquier, B. Cordier, D. Corre, N. Dagoneau, A. Formica, D. Götz, L. Jouvin, J-P. LeFèvre, H. Louvin, J. Palmerio, T. Roland, A. Sauvageon, S. Schanne, K. Tazhenova, D. Turpin AIM²: T. Sadibekova, J. Rodriguez LAM³: F. Agneray, S. Basa, C. Moreau, T. Fenouillet, Y. Roehlly, J-C. Thome ObAS⁴: L. Kleiver, A. Lorang, P. Maggi, L. Michel LUPM⁵: U. Jacob, P. Maeght, F. Piron, T. Maiolino APC⁶: P. Bacon, N. Bellemont, F. Cangemi, C. Cavet, A. Coleiro, A. Goldwurm, C. Lachaud, S. Le Stum IAP⁷: F. Daigne, L. Domisse, F. Lacreu, G. Tcherniatsky, J. Wang LUX⁸: Y. Canton, S. Verdani UNIDIA⁹: I. Jegouzo IJCLab¹⁰: S. Du, N. Leroy, S. Lion, R. Le Montagner, C. Van Hove IRAP¹¹: J. Alaux, M. Boiziot, L. Bouchet, O. Godet CPPM¹²: M. Ageron, T. Auphan, H. Benamar, D. Dornic.

Acknowledgments

The Space-based multi-band Variable Objects Monitor (SVOM) is a joint Chinese-French mission led by the Chinese National Space Administration (CNSA), the French Space Agency (CNES), and the Chinese Academy of Sciences (CAS).

The software development activities in all FSC contributing laboratories are supported by CNES.

ORCID iDs

H. Louvin <https://orcid.org/0009-0006-1908-4592>
D. Corre <https://orcid.org/0000-0001-6354-8514>
A. Formica <https://orcid.org/0000-0001-8308-2643>
T. Sadibekova <https://orcid.org/0000-0002-5162-4222>
J. Palmerio <https://orcid.org/0000-0002-9408-1563>

N. Dagoneau <https://orcid.org/0000-0002-1361-2562>
A. Claret <https://orcid.org/0000-0003-2901-8223>
L. Michel <https://orcid.org/0000-0001-5702-0019>
P. Maggi <https://orcid.org/0000-0001-5612-5185>
F. Piron <https://orcid.org/0000-0001-6885-7156>
P. Bacon <https://orcid.org/0000-0003-1350-2037>
C. Cavet <https://orcid.org/0000-0001-8988-9800>
F. Daigne <https://orcid.org/0009-0005-7119-4716>
G. Tcherniatsky <https://orcid.org/0009-0002-9399-7606>
Y. Canton <https://orcid.org/0009-0002-7205-3788>
N. Leroy <https://orcid.org/0000-0002-2321-1017>
S. Lion <https://orcid.org/0000-0003-0164-8354>
R. Le Montagner <https://orcid.org/0000-0002-6099-8939>
M. Boiziot <https://orcid.org/0000-0003-0298-9820>
T. Auphan <https://orcid.org/0009-0008-9988-3646>

References

- Agneray, F., & Moreau, C. 2019, *ASPC*, **521**, 723
Basa, S., Lee, W. H., Watson, A. M., et al. 2026, *RAA*, **26**, 104031
Claret, A., Turpin, D., Moreau, C., et al. 2026, *RAA*, **26**, 104033
Cordier, B., Jeannin, L., Lafabrie, P., et al. 2026a, *RAA*, **26**, 104006
Cordier, B., Wei, J.-Y., Zhang, S., et al. 2026b, *RAA*, **26**, 104001
George, I. M., & Corcoran, M. F. 2005, The HEASARC Calibration Database, Tech. rep., OGIP Calibration Memo CAL/GEN/91-001
Godet, O., Atteia, J.-L., Schanne, S., et al. 2026, *RAA*, **26**, 104007
Goldwurm, A., Bacon, P., Bellemont, N. G. G., et al. 2026, *RAA*, **26**, 104011
Götz, D., Crepaldi, S., Doumayrou, E., et al. 2026, *RAA*, **26**, 104013
Laycock, P. J., Dykstra, D., Formica, A., et al. 2018, *JPhCS*, **1085**, 032040
Maggi, P., Michel, L., Gotz, D., et al. 2026, *RAA*, **26**, 104017
Michel, L., Formica, A., Le Fèvre, J.-P., Meesen, C., & Zurbach, C. 2020, *ASPC*, **522**, 599
Pence, W. D., Chiappetti, L., Page, C. G., Shaw, R. A., & Stobie, E. 2010, *A&A*, **524**, A42
Piron, F., Daigne, F., Maiolino, T., et al. 2026, *RAA*, **26**, 104012
Wu, C., Palmerio, J., Sadibekova, T., et al. 2026, *RAA*, **26**, 104027