IN2P3 Scientific Council - Computing and Data Processing

Report on Rubin/LSST, Euclid and CTA Real Time Analysis

Dominique BOUTIGNY (LAPP)

with many contributions acknowledged in the text

June 2022

In this report we will address the main characteristics of the Rubin/LSST and Euclid computing models and their implementations at IN2P3. We will also describe the “real time analysis” component of the CTA / LST project.

1. Rubin/LSST[1]

The Rubin Observatory will be a central instrument for cosmology in the coming decade. The Rubin Observatory hosts a 8-m class telescope being built in Chile, which will perform a photometric survey (called LSST, for Large Survey of Space and Time) of the visible sky in six optical bands (u, g, r, i, z,y) over the course of ten years. With its 3.2 billion pixel camera and its large field of view (~10 deg²) it will cover the whole southern hemisphere sky in ~4 days by taking 1 picture every 40 seconds. By doing so every night for 10 years, it will produce an unprecedented amount of data in the astronomy field.
IN2P3 is the only French institute involved in Rubin construction and operation. There are a few INSU members associated to the Science Collaborations.
The following table summarizes some key numbers for 10 years of operation:

Number of images collected	5.5 millions
Number of data collected per 24h period	~20 TB
Raw data + calibration after 10 year	60 PB
Total amount of data including intermediary products	2 eB (2000 PB)
Number of objects in the final catalog	37 billions (20 B galaxies, 17 B stars)
Final database size	15 PB

The images acquired by Rubin will follow 2 different paths:

1. Each image will be transferred and promptly processed at the US Data Facility (USDF) in such a way that any difference with respect to a set of reference images will trigger an alert in less than 60 seconds after acquisition. It is expected that ~10

[1] Credit to: F. Hernandez (CC-IN2P3), E. Ishida (LPC), J. Peloton (IJCLAB)

million alerts will be generated every night. This flow will be distributed toward a set of officially approved alert brokers in charge to classify, enrich and redistribute the alerts to the scientific community. In France, the Fink collaboration is building such a broker that will be deployed and maintained at CC-IN2P3 (see Fink section below).

2. On a regular basis (2 times during the first year and then once per year), all the data collected since the beginning of the project will go through the Data Release Processing (DRP) pipeline in order to generate catalogs of objects (star, galaxies, asteroids, light curves of transient detections,...) and science-ready images. The set of consistent catalogs and images will constitute a Data Release that will be made available for data rights holders. The DRP will be handled by the US (25%), UK (25%) and France (50%) (see DRP section below).

It is important to note that the Rubin project is responsible to produce the catalogs and to provide an analysis platform to access the data but it is not in charge of their scientific exploitation. The science is under the responsibility of 8 science collaborations. IN2P3 is mainly in the Dark Energy Science Collaboration (DESC) which has its own demands in terms of computing and data processing.

1.1. The Data Release Processing (DRP)

The annual DRP requires a complex workflow with multiple steps and heavy data access. It will be split among 3 sites: the US Data Facility (operated by SLAC) for 25%, the UK Data Facility for 25% and the FRench Data Facility (operated by CC-IN2P3) for 50%. The whole workflow will be centrally orchestrated by the USDF. The role of IN2P3 was formalized in 2015 with a memorandum of agreement (MOA) but as a result of a Rubin re-organization a new formal agreement is still to be written and signed. The agreement also foresees that CC-IN2P3 will maintain an integral copy of the published data. This will guarantee that the French community keep privileged access to the Rubin data, thus maximizing the scientific return.

The DRP will run on heterogeneous systems. While not yet fully decided, the USDF may rely on the Google cloud and the Kubernetes container orchestration system. CC-IN2P3 will continue to use its “on-premise” system and UKDF will probably run on an in-house OpenStack cloud. The job orchestration among the processing facilities will probably use the PanDA system originally developed for ATLAS and the data management will rely on Rucio which is another tool from the HEP world.

At the moment the image processing system is being exercised on a simulated dataset independently at an Interim Data Facility (IDF) in the US (fully implemented on Google cloud) and at CC-IN2P3. This large scale test is critical to debug the whole workflow and to validate the processing tools as well as the underlying hardware and software infrastructure.

From the international network point of view, CC-IN2P3 is working with RENATER and GEANT in order to have a 100 Gb/s link available between the USDF and CC-IN2P3 before the start of the survey (the current bandwidth is 2x20 Gb/s).

The following picture shows the foreseen Rubin computing architecture which is being deployed at CC-IN2P3. It also indicates the expected growth between 2022 and 2032. *These numbers are still under discussion and are evolving as the Rubin computing model is being refined.*

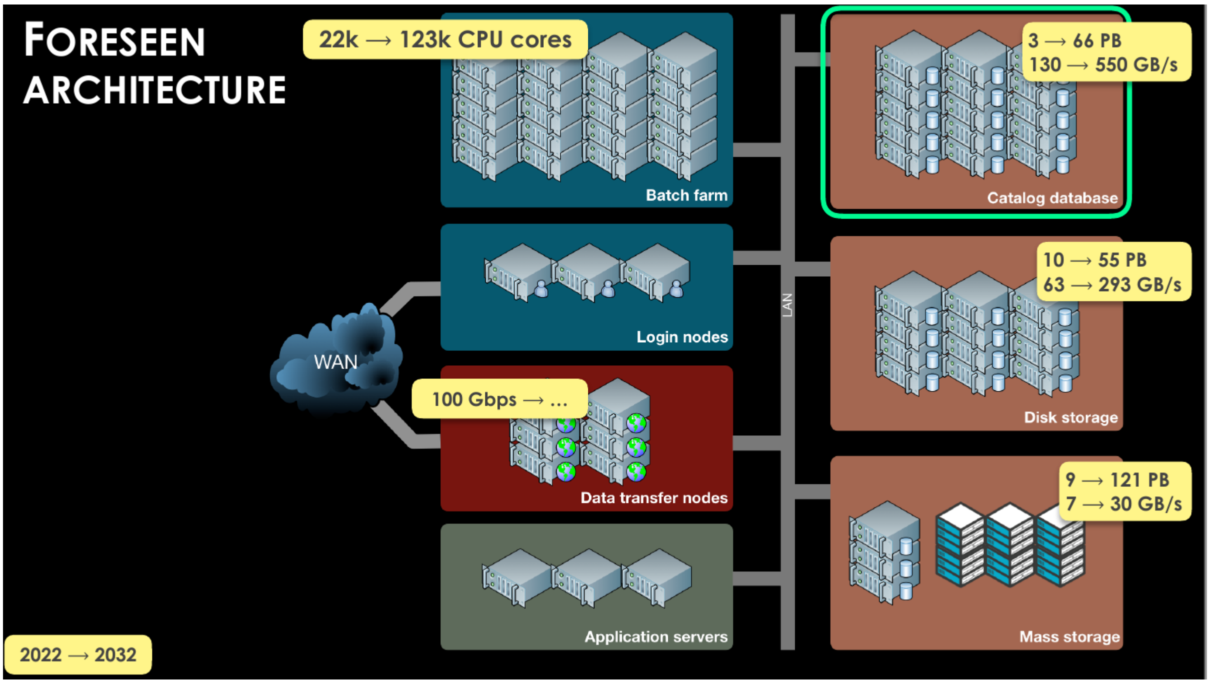

1.2. The Qserv database

In order to store and to give access to the catalog of objects, Rubin has developed ***Qserv***, a relational SQL database system distributed over a large number of nodes (~1000). Each node in the cluster runs a local MariaDB database and can handle subsets of complex queries. The global result is then aggregated from all the individual nodes and returned to the user. Qserv is optimized for astronomical queries and is natively supporting a 3D spherical geometry. Qserv is also able to cache and reuse the results of partial queries in order to speed-up the processing time. LPC has provided one developper to the Qserv team allowing to build considerable expertise in France. A Qserv instance has been deployed at CC-IN2P3 on a 20 node Kubernetes cluster which is already used for DESC science analyses.

1.3. The Rubin Science Platform (RSP)

In order to give access to up to 6000 users, Rubin is developing the RSP, an interactive science analysis platform offering programmatic access to survey data using the Python programming language and based on Jupyter notebooks, querying tools compatible with Virtual Observatory (IVOA) standards, the Firefly image visualization and running on a scalable Kubernetes platform. IN2P3 is deploying an instance of such a platform. LAPP, through the ESCAPE European project, has contributed to the deployment and assessment of the RSP and is currently checking whether it can be reused by other collaborations.

1.4. The Fink broker

The Fink collaboration has developed a prototype broker deployed on the VirtualData Cloud (UPSaclay), which uses OpenStack. It is based on Apache Kafka for the communication layer, Apache Spark for the processing and Apache HBase for the science database. It is being tested on the ZTF alert stream and has been proved to scale up to the expected

Rubin/LSST rate (10k alerts / 30s). It is expected that 10 years of raw alerts will require 3 PB of storage. Fink is implementing sophisticated machine learning algorithms based among others, on adaptive learning and Bayesian neural networks. The enriched and filtered ZTF alerts are already distributed through a science portal and corresponding API.
The Fink collaboration is very active and new functionalities and algorithms are regularly enriching the broker.
Fink is currently being deployed on the CC-IN2P3 cloud. This will ensure to meet the requirements for processing the stream from the 10 years of LSST, and maximize the availability and the stability of this important service.

1.5. DESC Computing

The Dark Energy Science Collaboration (DESC) will of course rely on Rubin processed data but will also need some specific data processing in order to match the extreme precision on galaxy measurements or supernovae light curves required for cosmology. DESC will also need to run complex pipelines to derive the cosmology parameters from the various cosmology probes.
The bulk of the computing resources needed by DESC will be provided by NERSC (Berkeley) supercomputers. IN2P3 will also maintain the DESC analysis tools at CC-IN2P3 to maximize the scientific return in France and also to contribute its fair share to the DESC collaboration.
Significant resources have been provided to DESC in the past years by IN2P3, especially for the 2nd DESC Data Challenge for which all the image processing has been completely performed at CC-IN2P3 by IN2P3 personnel.

1.6. Computing activities related to Rubin Commissioning

The LSST survey is scheduled to start in spring 2024 but some commissioning data are already available and are being transferred at CC-IN2P3. This commissioning activity is going to ramp up quickly and its success will be crucial for the French contribution to Rubin. Today we already have full focal plane images from the camera test bed at SLAC and on-sky data from the auxiliary telescope at Cerro Pachón (summit).

1.7. Core human-power

CC-IN2P3	2 dedicated FTE (3 persons) + ~ 5 FTE shared between several engineers for general operation and infrastructure support	General Infrastructure - DRP - Qserv DB - RSP
LAPP	1.5 FTE	DRP - RSP - Qserv DB - DESC computing
LPC	1.2 FTE	Qserv - Fink
IJCLAB	1.6 FTE	Fink
LPSC	0.1 FTE	DRP

1.8. Conclusions on Rubin computing

1. As can be seen from the previous sections, IN2P3 is deeply involved in the Rubin computing and will bring a lot of resources for the annual Data Release Processing over the operations phase. This effort will guarantee a strong scientific return, first by the expertise acquired on the Rubin software and second by the availability of the data and analysis tools in France.
2. We notice that the Rubin computing is relying on tools that are developed and validated outside the HEP / Astroparticle community. The most striking example is the containers' technology including Kubernetes which is directly derived from the GAFAMs. We are no longer driving these developments but we need more and more sharp experts to implement and maintain them.
3. It should be emphasized that the expertise and flexibility provided by CC-IN2P3 is invaluable in the Rubin and DESC context. The fact that CC-IN2P3 engineers are now fully involved in Rubin and are part of the decision processes is essential for the success of the project.

2. Euclid[2]

Euclid can be seen as the in-orbit Rubin counterpart for cosmology. The telescope's smaller diameter (1.2 m) and field of view (0.5 deg^2) is compensated for by its location outside the atmosphere at the L2 Lagrange point, allowing it to provide much sharper images and to get a sensitivity extended in the infrared region of the spectrum. Euclid is equipped with two instruments: a high quality, panoramic, visible imager (VIS) for gravitational lensing measurements, and a near infrared, 3-filter (Y, J and H) photometer and a slitless spectrograph (NISP) dedicated to galaxy clustering measurements.
Euclid was scheduled for launch in early 2023 on a Soyuz space launcher. The international situation is forcing the collaboration to consider the Ariane 6 backup launcher as the new baseline which will result in a yet to be estimated delay. After launch, Euclid will survey ⅓ of the total sky (15 000 deg^2) during 6 years.
In order to fully reach the mission's core science goal on dark energy, Euclid should be complemented with external ground data including Rubin-LSST in order to get accurate photometric redshift measurements. The processing of these external data is fully part of the Euclid computing model.
The following table summarizes some key numbers on Euclid data:

Single exposure time	90 s in photo mode (30 s for Rubin) and 550 s in spectro mode
Total amount of data including intermediary and external products	150 PB
Number of measured galaxies	2 10^9
Number of galaxies with spectroscopic measurements	30 - 50 10^6

[2] Credit to Q. Le Boulch (CC-IN2P3), Y. Copin (IP2I), S. Escoffier (CPPM), K. Ganga (APC), S. Kermiche (CPPM)

Raw data / 24 hours	140 GB

Compared to Rubin, the dataset size is smaller but still very significant.
A community of more than 1500 scientists and engineers is currently active in the Euclid Consortium. France contributes to 30% of Euclid construction with contributions from CNES, IN2P3, INSU and IRFU. CC-IN2P3 will provide 30% of the storage and processing capacity.

2.1. Organization of the Euclid Ground Segment

The main elements of the Euclid Ground Segment are the Ground Station, the Mission Operation Center (MOC), the Science Operation Center (SOC) and the Euclid Consortium Science Ground Segment (ECSGS). The Ground Station and the MOC compose the Operations Ground Segment (OGS) and are provided by ESA; the SOC and the ECSGS compose the Science Ground Segment (SGS).

The following figure shows the various components of the Euclid data processing infrastructure and their connections

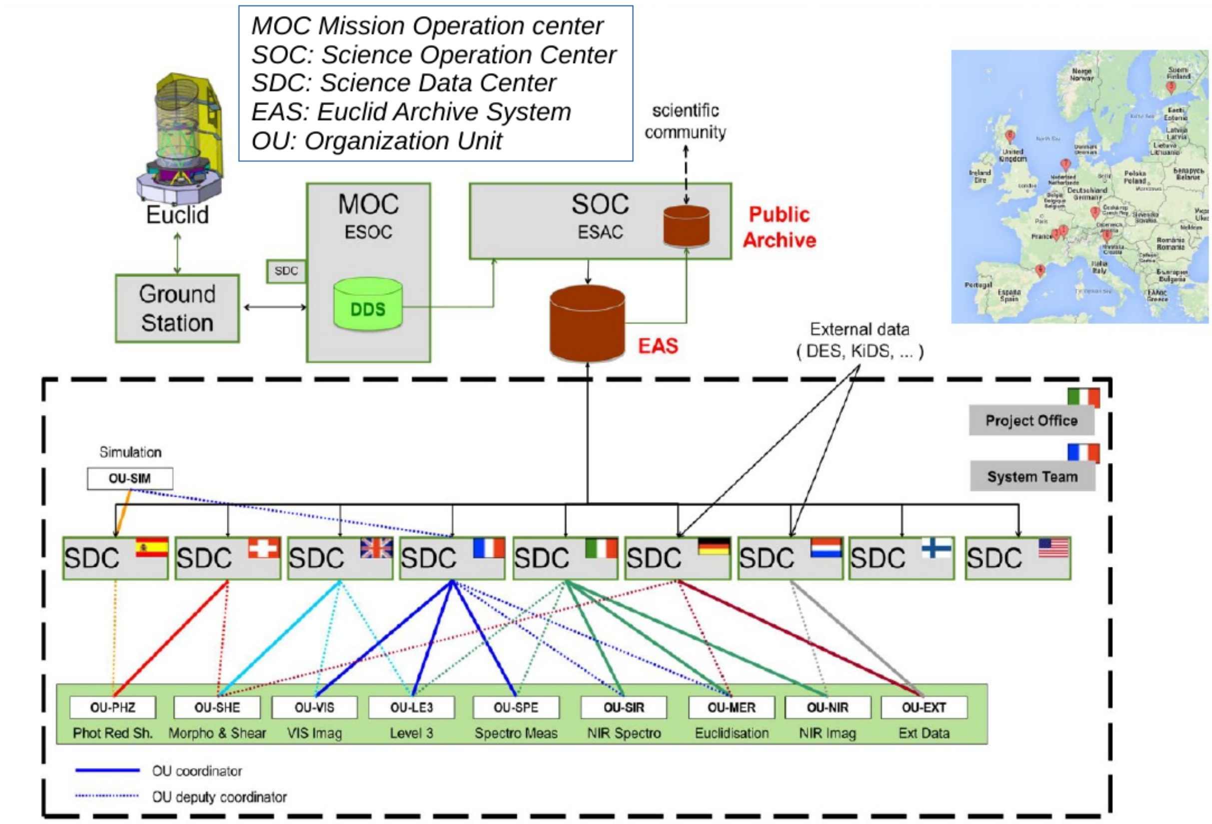

Figure: Organization of the Euclid Ground Segment

The major objective of the SGS is the delivery of high-quality science products for validation and release suitable to meet the science objectives of the mission. These deliverables resulting from the mission, which are made available to the general scientific community and which form the basis for scientific research and publications, include not only processed data but also the associated quality control information, which ensures traceability of input data sets as well as the processing steps applied. Managing the amount of data that the mission will generate and the heavy processing that is needed to go from the raw data to the science products pose challenges for the SGS. These characteristics enforce the concept of a data-centric approach: all the operations of the SGS will revolve around a distributed Euclid Archive System (EAS) that plays a central role in the data products storage and their metadata inventory.

The SGS is in charge of designing, developing, integrating and operating the scientific data processing. Organisation Units (OUs) are organizational groups that are mapped upon the organization of processing functions. They are in charge of the definition of data processing algorithms and of the development of code prototypes.

The backbone of the ground segment is formed by a series of Science Data Centers (SDC), one in each participating country of the Euclid Consortium Consortium, which will provide large and reliable computing resources to the SGS. These SDCs are in charge of running the Euclid SGS pipelines on designated sky patches. The SDCs are also in charge of instrument-related processing and will host the Instrument Operation Teams (IOTs), one for each instrument.

2.2. Organization Units (OU)

Euclid is a complex mission and its major task is to derive important statistical results by minimizing and controlling systematic uncertainties in the experiment. This calls for a well-controlled flow of sophisticated simulations which are needed to evaluate and monitor performances, to validate the data reduction procedures and, in the end, to demonstrate the full control of the aforementioned systematic uncertainties.

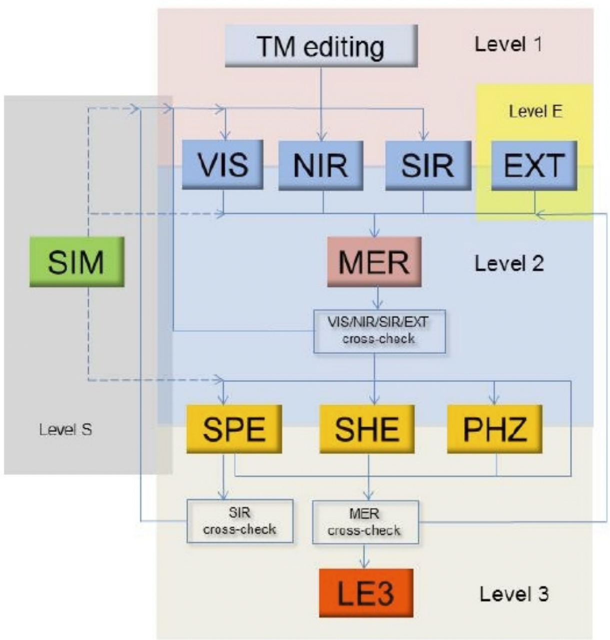

Data processing levels, functions and flow in the Euclid SGS

The SGS data products, and consequently the data processing, are categorized in six levels (levels S, 1, E, 2, 3 and Q). The logical data processing flow is represented in the figure above, where the different data processing levels are connected with logical data processing functions. These logical functions, or Organisation Units (OU), have been defined by considering that they represent self-contained processing units. These OUs are listed hereafter.

2.2.1. Level S (pre-mission simulated data)

- **OU-SIM: Simulations:** Simulations play a key role in Euclid science, in order to discriminate between an actual signal of interest and instrumental or data reduction artefacts. The OU-SIM integrates or interfaces with the simulation activities that are already taking place in the Instrument Development Teams, so as to provide the VIS, NIR and SIR OUs with input data.

2.2.2. Transition between level 1 (raw VIS and NISP images) to level 2 (calibrated data)

- **OU-VIS: Visible imaging data:** Processing of the Visible imaging data from edited telemetry to produce fully calibrated images, as well as source lists (for quality check purposes only).
- **OU-NIR: Near-Infrared imaging data:** Processing of the Near-Infrared imaging data from edited telemetry to produce fully calibrated images as well as source lists (for quality check purposes and to allow spectra extraction),
- **OU-SIR: Extraction of spectra in NISP:** Processing of the Near-Infrared imaging data from edited telemetry to produce spectrograms and extract fully calibrated spectra from the slitless spectroscopic frames taken by the NISP.
- **OU-EXT: External data (level E):** External data are data needed by the SGS which are obtained with instruments other than those aboard the Euclid satellite. The data is typically obtained by external consortia, either publicly available or following agreements with EC. This is essentially multi-wavelength data for photo-z estimation, but also spectroscopic data to validate the spectrometric redshift measurement tools.

2.2.3. Level 2 (calibrated data)

- **OU-MER: Merging:** Merging of all the Level 2 information. It is in charge of providing stacked images, source catalogs and calibrated photo-z's where all the multi-wavelength data (photometric and spectroscopic) are aggregated.

2.2.4. Transition between level 2 (calibrated data) to level 3 (catalogs)

- **OU-SPE: Spectroscopic redshift:** Extracts spectroscopic redshifts
- **OU-PHZ: Photometric redshifts:** Computes photometric redshifts from the multi-wavelength imaging data.
- **OU-SHE: Shear measurements:** Computes shape measurements on the visible imaging data.

2.2.5. Level 3 (catalogs)

- **OU-LE3: high-level science:** Computing all the high-level science data products, from the fully processed shape and redshift measurements (and any other possibly needed Euclid data).

2.2.6. Level Q (Quick-Release data)

IN2P3 is involved in OU-SIM activities, particularly in SIM-NIS (CPPM) and SIM-EXT (APC), OU-EXT (APC) and OU-LE3 (LPSC), and is co-lead of OU-SIR (IP2I).

The principal input providers for the data processing tasks are the Science Working Groups, who specify the scientific tasks that the SGS must perform to generate the necessary data products for the scientific exploitation of the mission. The Mission Operations Center is responsible for downloading the science data generated by the instruments and providing that data to the SGS with those other data required for the interpretation of the science data. The external data providers supplement the Euclid science data with data from other surveys, effectively forming a third instrument alongside the VIS and NISP instruments.

2.3. The French Science Data Center (SDC-Fr)

The SDCs are in charge of the science data processing software (pipelines) and simulation software development, the science data processing and the generation of the data products. There are 8 SDCs over Europe and 1 in the US (IPAC), sharing the processing and storage on a grid-like infrastructure. Each SDC will run the complete processing pipeline on its allocated dataset. Once processed, the output data products are referenced in a centralized metadata database located at ESAC (Madrid). The critical data are replicated on several SDC.

The pipeline developments which are also a responsibility of the SDCs are handled by IN2P3, INSU and CEA. Contrary to Rubin-LSST, the development of the science pipelines is fully integrated within the global computing model (in Rubin, science pipelines are handled by separate scientific collaborations).

The French SDC (SDC-Fr) production infrastructure that runs the pipelines is the CC-IN2P3, and is referred to as SDC-PROD. The computing infrastructure at CC-IN2P3 is funded by IN2P3, and CNES is funding 3 fixed term engineer contracts.

The SDC-Fr also provides a software development infrastructure as a contribution to the SGS System Team Common Tools. It is under the responsibility of APC and is referred to as SDC-DEV.

The development of the Common Tools is managed and performed as a main contribution of the French SDC in collaboration and with contributions from other SGS System Team members and SDCs. Among the collaborative tools, the software development team uses:

- Git to share code and track code changes;
- CODEEN (COmmon DEvelopment ENvironment) facilities for testing the build and test on the target platforms;
- Redmine wiki to report about bugs and features.

CODEEN is managed by APC and is hosted at CC-IN2P3, a mirror will be set up at SDC-Uk (ROE) as a failover.

2.4. Scientific Challenges (SC)

Extensive scientific end-to-end (E2E) tests (so-called Scientific Challenges, SC) will demonstrate the ability of the Euclid mission to reach the scientific goals of the mission; specifically for the SGS they will demonstrate the ability of the SGS to process the data to required level of accuracy such that the Science Working Groups (SWGs) can extract the scientific results from the data.
Simulations are a core element of the SCs, providing the simulated data based through which the ability of the SGS to recover the input parameters is demonstrated. The inputs to OU-SIM will be provided by the SWGs for the catalogs which form the basis of the simulations and the instrument parameters through the Mission Database. At IN2P3 we are major contributors to all Euclid scientific challenges, from SC1 in 2015 with 0.5 deg^2 to SC8 in 2021 with 150 deg^2, and we coordinate production of Euclid simulations at CC-IN2P3. Below is the ratio of different processing functions for CPU and storage in OU-SIM for SC8.

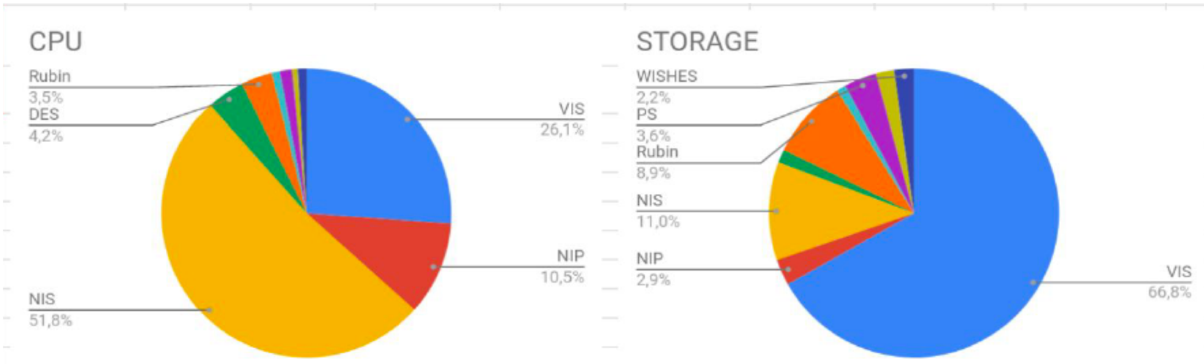

CC-IN2P3 is also a major contributor to the processing and storage resources used during these Scientific Challenges, for both simulations and processing. SC8 has required about 250 TB of storage and a million CPU hours at CC-IN2P3.

2.5. Instrument Teams (IDT and IOT)

Instrument Operations Teams (IOTs) are in charge of the operations of the Euclid instrument modes and are an integral part of the SGS. IOT members will be in charge of the instruments maintenance and operations and will be composed of instrument scientists and engineers, expected to come from the teams developing and testing on-ground the instruments (IDTs for Instrument Development Teams), and by data processing experts coming from the relevant OUs. The IDTs will maintain responsibility on their instruments until the end of the in-flight Performance Verification phase after launch. The IOTs will be in charge during routine operations.
During the PV phase, The IDTs/IOTs will deliver to the SOC a routine calibration plan for each instrument (VIS and NISP) the operational catalogs containing calibration targets and fields, avoidance targets and regions, priority fields, and so on.
During routine operations, the IOTs will prepare instrument-commanding instructions, the instruction database and procedure updates for their respective instruments and provide them to the SOC via the EAS. The IOTs expect to be able to submit instrument commanding requests and pointed calibration requests on a regular basis to the SOC.

The SOC and IOTs will jointly monitor the health of the payload (including verification of the instrument status, monitoring instrument calibration and performing trend analysis), providing instrument health reports on a regular basis.
IN2P3 is fully involved in IDT and future IOT teams, with the responsibility of the NISP Instrument Scientist, the NISP spectroscopic calibration Scientist and the (IR) detector scientist. Staff from the instrument-oriented SDCs (the IOTs) shall be involved in the on-ground calibration activities from the earliest phase possible.

2.6. Science Working Groups

Science Working Groups (SWGs) are gathered along the science objectives of the mission, both core and legacy. The computing and storage needs for the scientific exploitation developed in SWGs are not included in the SGS, the identification of the needs in each SWG is underway within EC.

2.7. Euclid resources at CC-IN2P3

The following figures show the expected evolution of CPU, disk and tapes resources at CC-IN2P3 as a function of time. We see that Euclid is considerably less demanding than Rubin in terms of resources. As usual, this kind of resource extrapolation over a decade has some limitations and we should expect the model to evolve, especially the ratio between disk and tapes.

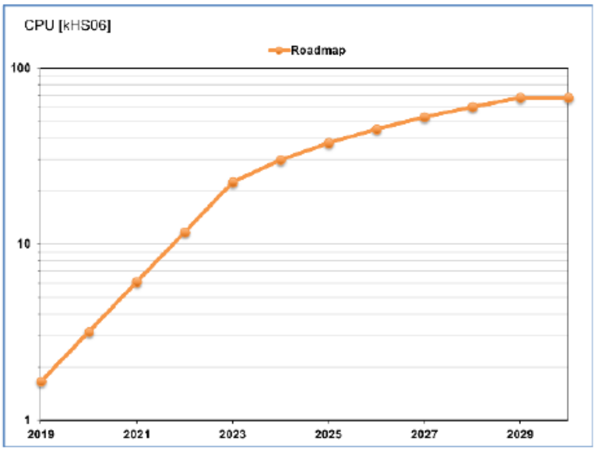

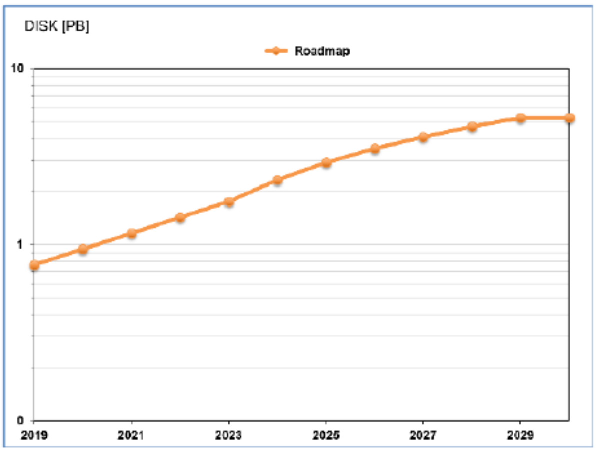

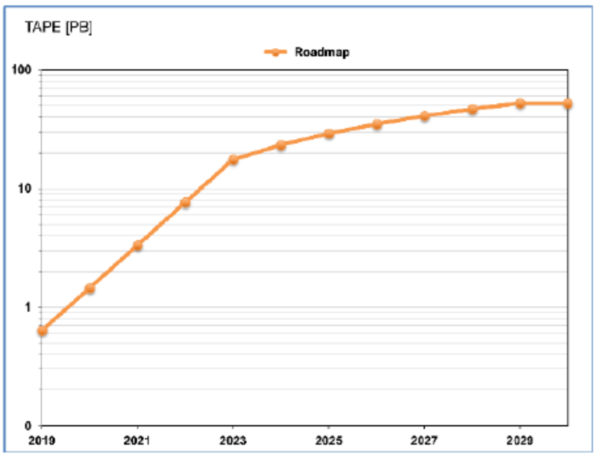

These resource needs are funded through a CNES / IN2P3 convention until the first Data Release (DR1). These figures are currently being re-estimated and a new convention is under discussion to cover the whole mission.

2.8. Human-power

CC-IN2P3	3 FTE (CDD) + 0.3 FTE	SDC-FR Production Infrastructure
CPPM	2 FTE (CDD) + 2.5 FTE	OU-SIM / IOT NISP
IP2I	1 FTE (CDD) + 1.6 FTE	OU-SIR
APC	2.4 FTE (CDD) + 0.75 FTE	OU-EXT / OU-SIM-EXT
LPSC	0.4 FTE	

*In this table, CDD are all funded by the Centre National des Etudes Spatiales (CNES).

3. CTA Real Time Analysis [3]

This section is largely borrowed from "The Science Alert Generation system of the Cherenkov Telescope Array Observatory" - A. Bulgarelli et al.

The Cherenkov Telescope Array (CTA) Observatory, with dozens of telescopes located in both the Northern and Southern Hemispheres, will be the largest ground-based gamma-ray observatory and will provide broad energy coverage from 20 GeV to 300 TeV. The large effective area and field-of-view, coupled with the fast slewing capability and unprecedented sensitivity, make CTA a crucial instrument for the future of ground-based gamma-ray astronomy. To maximize the scientific return, the array will send alerts on transients and variable phenomena (e.g. gamma-ray burst, active galactic nuclei, gamma-ray binaries, serendipitous sources). Rapid and effective communication to the community requires a reliable and automated system to detect and issue candidate science alerts. This automation will be accomplished by the Science Alert Generation (SAG) pipeline, a key system of the CTA Observatory. SAG is part of the Array Control and Data Acquisition (ACADA) working group.

Candidate science alerts will be sent to the Transients Handler that will evaluate the results of the SAG to generate the final science alert to the community within 5 s of receiving it. If we consider also 5 s required to acquire data, CTA is capable of issuing science alerts with a total latency of less than 30 s.

[3] Credit to: S. Caroff (LAPP)

3.1. General Architecture

The following picture describes the general architecture of the system and emphasizes the central role of the SAG pipeline

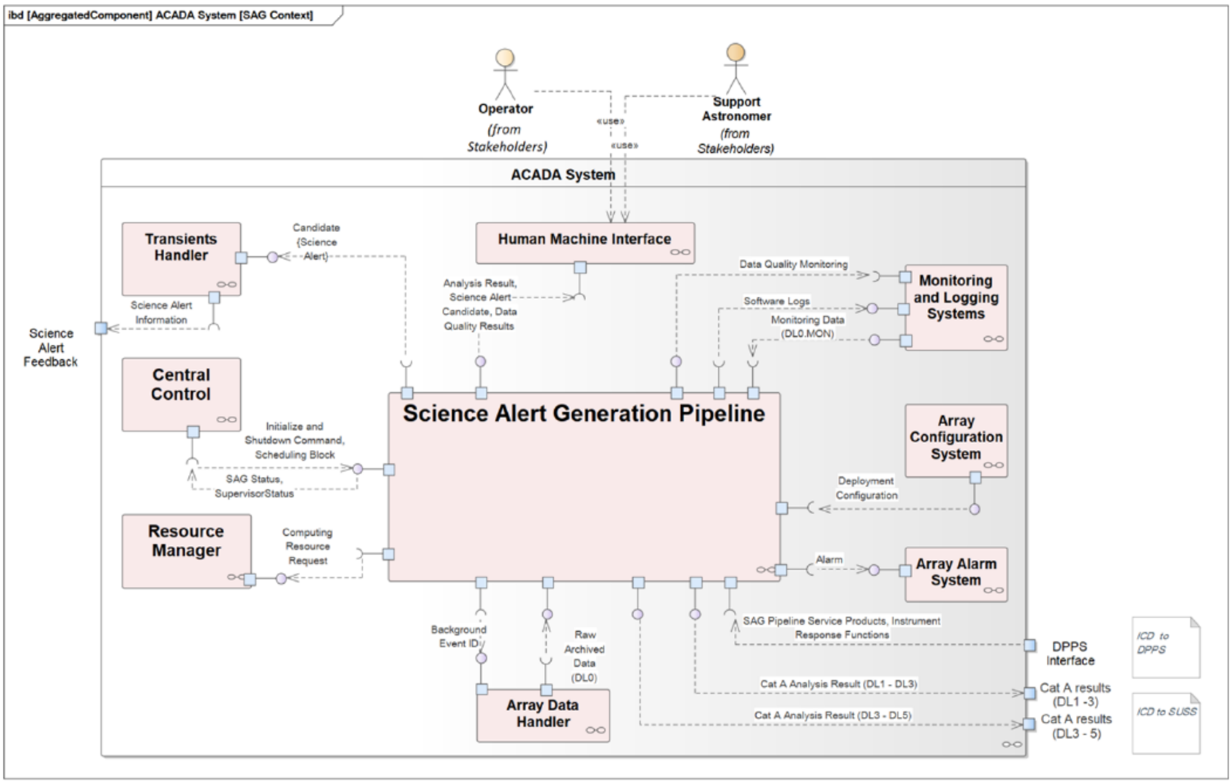

The SAG pipeline is composed of the following components:

1. **sag-reco:** the Low-Level Reconstruction pipeline for fast reconstruction to process individual DL0 events and produce the corresponding DL3 (the gamma-ray like event list)) data with a latency of less than 15 s.
2. **sag-dq:** The SAG system also performs an online data quality analysis to assess the instruments' status during the data acquisition with a latency of less than 5 s: this analysis is crucial to confirm good gamma-ray sources and transient detections.
3. **sag-sci:** The High-Level analysis scientific pipelines is able to execute scientific analysis and generate candidate science alerts with a latency of less than 5 s.

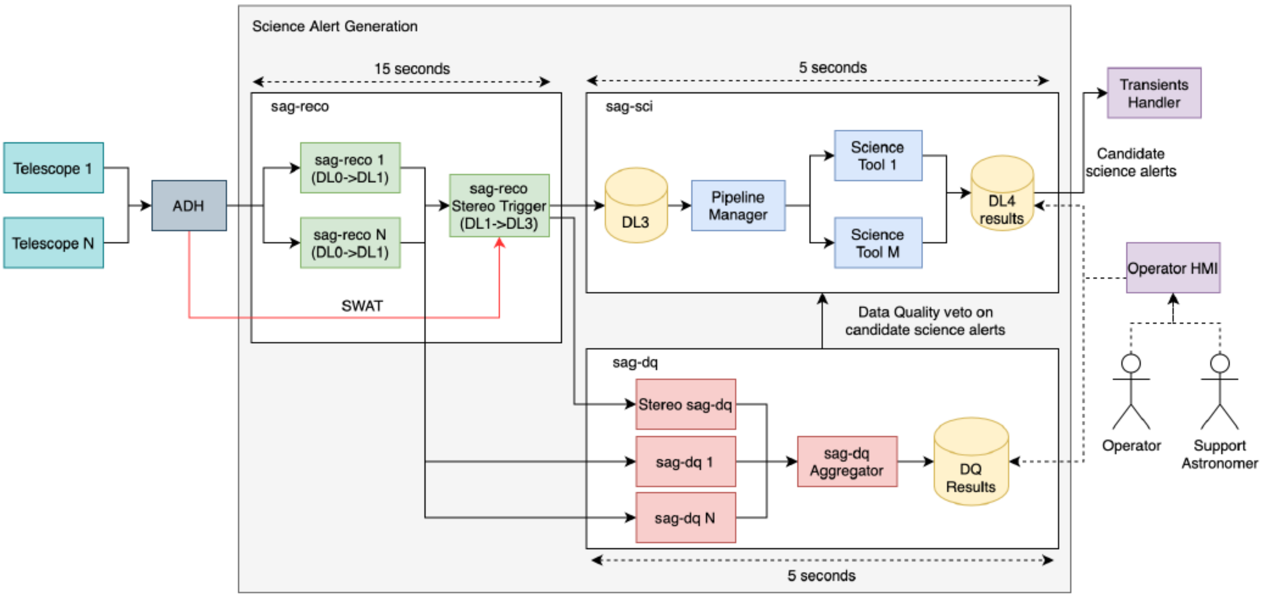

The SAG pipeline design and implementation is performed within the LST collaboration and will be delivered to the CTA Observatory (CTAO) upon completion. The expected delivery date is for 2026, with several releases every 6 months supposed to handle more and more functionalities and telescopes. The current release is the release 1 with the objective to run a simple mono hillas analysis with the LST-1 telescope, the delivery date for this release is the end of the year 2022.

3.2. Implementation

The sag-reco pipeline is under the LAPP responsibility. In order to fulfill the very strict timing constraints for alert generation while minimizing the size and the cost of the cluster to execute it, the SAG pipeline is designed using a highly optimized High Performance Computing (HPC) approach exploiting the most advanced features of modern CPUs (vectorization, caching, etc.)
The current sag-reco algorithm is based on the Hillas parameterization of the shower images to identify and reconstruct gamma rays. Other algorithms based on templates or on the Machine Learning suite GammaLearn are also being considered.

3.3. Human-power

LAPP	0.3 FTE (CDD) + 1.15 FTE	SAG-RECO
To be hired at LAPP	1 FTE (CDD)	SAG-RECO