

Gaia operations or 10 years of Big Data experience

Angélique Barbier^{a*}, Antoine Guerrier^b

^a *Centre National d'Etudes Spatiales (CNES), France, angelique.barbier@cnes.fr*

^b *Centre National d'Etudes Spatiales (CNES), France, antoine.guerrier@cnes.fr*

* Corresponding Author

Abstract

ESA's Gaia astronomy mission, launched on December 19th 2013, provides the most accurate 3D map of the Galaxy to date, observing 1.8 billion celestial objects in the Milky Way and beyond. The Gaia data processing is handled by a European Scientific Consortium, DPAC (Data Processing and Analysis Consortium); composed of around 430 scientists and engineers of 24 countries in Europe and beyond, and six DPC (Data Processing Centres) in Europe. CNES is a major actor of DPAC as it hosts one of the 6 processing centres: DPCC (DPC CNES) that is in charge of eight scientific chains dealing with spectroscopic processing (radial and rotational velocities), object characterization for solar system objects, binaries, galaxies and quasars and astrophysical parameter computations.

DPAC published data through regular releases. The last one was the Gaia Data Release 3 (DR3), which contains the result of the eight chains, processing 34 months of collected data. 80% of products published in Gaia Data Release 3 (DR3) were processed at DPCC. The Gaia catalogue is a real revolution and reference for all astronomers around the world. The Gaia contribution was awarded by the American Astronomical Society via the 2023 Berkeley Prize.

Since 10 years, the CNES team of 25 engineers work with around 80 scientists, spread over laboratories who provide the algorithms in Java code, that DPCC integrates in pipelines and validate in an operational centre, based on Big Data technologies.

DPCC deals with four operational challenges: data volume, complexity, performance and continuity. The input data volume is multiplied by 2 between 2 data releases. With 9 petabytes of dedicated storage and data exchanges ranging from 200 gigabytes to 3 terabytes per day with other centres, our data management concepts are the first key to the smooth running of the centre. Our success is also based on the adoption of a logic of continuous evolution, a kind of DevOps process, making the system robust to the increasing complexity. Then performance and continuity are both linked to the same root cause: timeliness constraints. A real expertise is therefore carried out by the engineers to optimize as finely as possible the use of all the resources of the computing and closely monitored cluster.

To conclude, DPCC maintains its effectiveness and achieves its objectives for each new data release only through a well-defined human organization and a pragmatic approach to operations and planning.

Keywords: data volume, complexity, performance, continuity, efficiency, pragmatism

Nomenclature

Not applicable.

Acronyms/Abbreviations

CNES: Centre National d'Etudes Spatiales
CCD: charge-coupled device
CU: Coordination Unit
DPAC: Data Processing and Analysis Consortium
DPACE: DPAC Executive committee
DPC: Data Processing Centre
DPCC: DPC CNES
DPCE: DPC ESAC
DR: Data Release
DRC: Data Reduction Cycle
ESA: European Space Agency
ESAC: European Space Astronomy Centre
ESOC: European Space Operation Centre
MDB: Main DataBase

MMI: Man-Machine Interface
NSS: Non Single-Star
RAM: Random Access Memory
SOC: Science Operation Centre
SAGA: System of Accommodation of Gaia Algorithms

1. Introduction

Gaia is a famous space mission of the European Space Agency (ESA), to provide a three-dimensional map of the Milky Way, for revealing the composition, formation and evolution of our Galaxy. The main objective is to map the sky in astrometry (stellar position, distances, and motion), photometry (brightness and colours) and spectroscopy (radial velocities and astrophysical parameters). It measures the characteristics of more than one billion objects in our galaxy –about one percent of its content- and beyond, with unprecedented precision (less than 0.00007 arcsecond and even better for bright stars). Gaia extends the previous catalogue from ESA’s Hipparcos mission by a factor 10 000, while measuring many additional parameters, thanks to the repetitive observations of each object during the last 10 years.

The Gaia satellite has been developed by Airbus Defence and Space (DS, formerly known as Astrium) for ESA and the control and mission centres are under ESOC responsibility.

The scientific data processing is delegated to the Data Processing and Analysis Consortium (DPAC), and relies on six Data Processing Centres (DPC) distributed all around Europe, which is in charge of elaborating the catalogue and publish it for free for the entire scientific community.

The paper aims at explaining how the Data Processing Centre CNES (DPCC) meets its four challenges - data volume, complexity, performance and continuity - thanks to a good balance between technical and organizational solutions, allowing it to meet the ESA and DPAC timing constraints while offering new possibilities for scientific calculations.

2. The Gaia project

2.1 Satellite operations

The Gaia satellite is a spacecraft (figure 1) built under an ESA contract by Airbus Defence and Space in Toulouse (France). It is equipped with a payload module, built around an optical bench that provides structural support for the two identical telescopes and the integrated single focal plane assembly, consisting of 106 CCDs with a total of 938 Million of pixels (figure 2).



Fig. 1. Gaia Deployable Sunshield Assembly fully deployed at Astrium Toulouse (Copyright Astrium)




Fig. 2. Gaia telescopes and integrated focal plan (Credit: ESA)

The Gaia satellite was orbiting around the Lagrange point L2, 1.5 million kilometers away from earth, in the opposite direction of the sun. It scans the sky with combined rotations, and records any object detected. A typical day for the satellite involved 70 million detected sources, 600 million measurements for astrometry, 155 million measurements by wide band spectrometry, 13 million measurements by narrow band spectrometry.

“Copyright 2025 by the Canadian Space Agency (CSA) on behalf of SpaceOps. All rights reserved. One or more authors of this work are employees of the government of France, which may preclude the work from being subject to copyright in France, in which event no copyright is asserted in that country.”

Launched on December 19, 2013 and initially planned for a 5-year mission, the system has finally doubled its exploration time, giving us 10 years of observations. Coming short of cold gas, observations stopped on January 15, 2025. The spacecraft retired on March 27, 2025, and finally passed to DPAC for the whole science processing. Some representative numbers of the overall mission are given in table 1.

Table 1. Number of data collected during the Gaia mission Operational data collected, ref. <https://www.cosmos.esa.int/web/gaia/mission-numbers>

Operational data collected	Value
Volume of science data collected (in GB) - compressed raw data	141,038
Number of object transits through the focal plane	267,336,261,480
Number of astrometric CCD measurements	2,635,171,720,297
Number of photometric CCD measurements	530,821,357,368
Number of spectroscopic CCD measurements	52,062,041,736
Total number of CCD measurements (astrometric + photometric + spectroscopic)	3,218,055,119,401

At the end, Gaia satellite has collected an average of 140 observations per star, which allows our scientific community to make amazing discoveries.

2.2 Data processing

During the course of the mission, Gaia scanned the sky producing a continuous stream of data that is unintelligible in its raw format to the scientists who want to use it. The task of converting the raw data into scientifically useful products is entrusted to the 430 scientists spread across 24 countries and software experts who form the Gaia Data Processing and Analysis Consortium (DPAC).

As shown in Figure 3, the Gaia DPAC has been split in 9 Coordination Units (CU) with an executive committee: the DPACE, in charge of the interface with ESA and to oversee and coordinate the DPAC activities, relying on a ‘Project Office’ team. Each CU is composed of members from different European laboratories, managed by a scientific leader and is attached to a DPC team which ensures the technical coordination of the software developments and execution. The scientific algorithms are developed in java language by all CU members, then delivered to the corresponding DPC for integration in so-called ‘chains’.

One unit (CU9) is responsible for the preparation, validation, documentation and distribution of the Gaia catalogue versions.

Located across Europe, the DPCs are in charge of processing one or more CU.

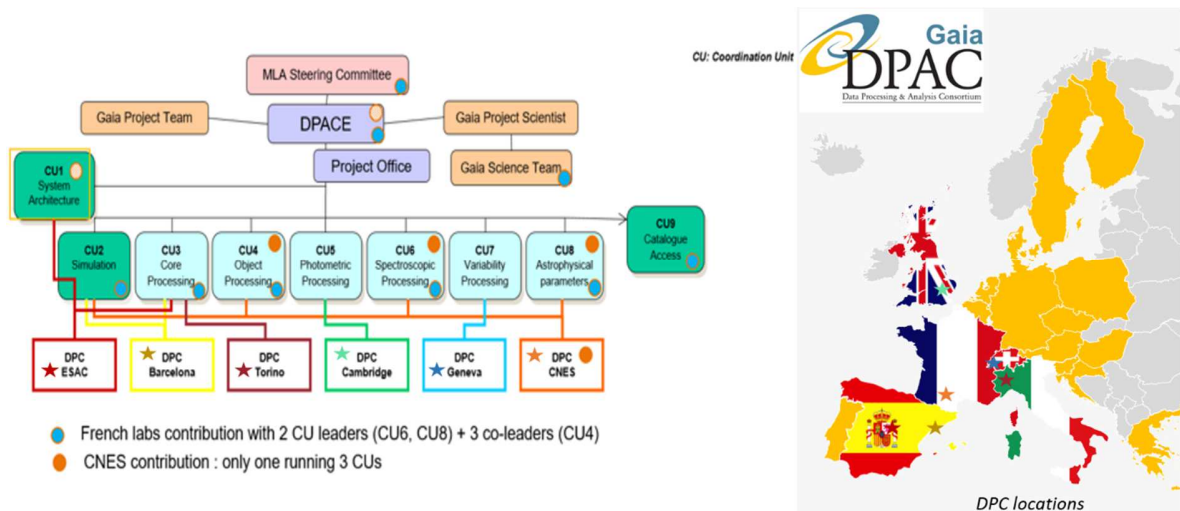


Fig. 3. DPAC organisation and DPC locations

The DPAC data processing requires to iterate several times over the data acquired during a predefined period - named ‘data segment’- in order to achieve high accuracy in both calibrations and results delivered with a Gaia data release. These are called Data Reduction Cycles (DRC) but often simply ‘cycles’. Segments and cycles are described in Fig. 4.




Fig. 4. Segments and Cycles in Gaia

These DRC always start from the beginning of the acquisitions that is the first measurement on July 25, 2014. Each DRC led to a new release at the end of each cycle. The last public Gaia data release, called DR3 covered 34 months of acquisition. The next one, called DR4 will cover 66 months of acquisition, the initial satellite mission duration.

After DR4 publication, DPAC will face off its last challenge: process 126 months of acquisition. This legacy is expected not before the end of 2030.

2.3 CNES involvement

The CNES's role in the Gaia project can be summarized by four key verbs: represent, support, build and produce.

CNES represents the French contribution as a member of the DPAC Executive Committee (DPACE), overseeing DPAC activities in consultation with the Coordination Unit and Data Processing Centre leaders, and as a member of the Gaia Steering Committee.

As a funding agency, CNES supports French laboratories and institutes. French laboratories and CNES together account for one-quarter of the consortium's participants.

As an expert in ground segment operations, CNES acts as deputy leader of the CU1 System Architecture and hosts a Data Processing Centre (DPC). The DPCC is one of the largest DPCs, as it hosts the processing of data for three major Coordination Units (CUs): Spectroscopy (CU6), Complex Objects (CU4: Solar System Objects, Multiple Stars, Extended Objects), and Astrophysical Parameters (CU8). The DPCC team develops processing pipelines based on hundreds of scientific algorithms and produces CU6/4/8 data for the next release thanks to a Big Data processing cluster. Forty percent of DPAC's scientific software is run at the DPCC.

The DPCC's involvement is also evident in the content of data releases. In the most recent public data release (DR3), 80% of the new products were generated at the DPCC.

3. The four challenges

3.1 Data volume

In order to refine our data reduction, each Data Release processing starts from the first Gaia measurement on July 25, 2014 to an end data, as shown in table 2. After 10 years of acquisitions (i.e. DR5), the data processing deals with several million and billion of values. In a strict computer science sense, all measurements are received in a 50GB/day telemetry via ESA Estrack network. This amounts around 140TB of compressed telemetry.

Table 2 : Number of raw data collected during each Data Release (G=giga)

Raw data	DR1	DR2	DR3	DR4	DR5
	until 16/09/2015	until 23/05/2016	until 28/05/2017	until 20/01/2020	until 15/01/2025

“Copyright 2025 by the Canadian Space Agency (CSA) on behalf of SpaceOps. All rights reserved. One or more authors of this work are employees of the government of France, which may preclude the work from being subject to copyright in France, in which event no copyright is asserted in that country.”

Number of Astrometric transits	30,3 G	52,4 G	77,9 G	142 G	265 G
Number of Astrometric CCD measurements	272,9 G	472,2 G	701,7 G	1278,5 G	2385,5 G
Number of Photometric transits	30,3 G	52,4 G	77,9 G	142 G	265 G
Number of Photometric CCD measurements	60,3 G	104,9 G	155,9 G	284 G	530 G
Number of Spectroscopic transits	1,9 G	3,1 G	4,8 G	9,2 G	17,3 G
Number of Spectroscopic CCD measurements	5,8 G	9,3 G	14,5 G	27,7 G	51,9 G
Compressed TM Volume to be processed	~16 TB	~28 TB	~41 TB	~ 75 TB	~ 140 TB

At each Data Release, DPCC has to manage more and more volume of input data.

But this input represents a small part of our processed data; throughout our processing, a lot of intermediate data is extracted and/or derived from this input raw data. For example, considering only the spectroscopic processing (CU6) performed at DPCC during DR4, DPCC has produced the following volumes for 50 TB of input data:

- 500TB of intermediate data, i.e. "less valuable" data, but which can be combined to produce major scientific results - e.g. extracted, cleaned and calibrated CCD spectra, internal information about transits and source, etc.;
- 50TB of output data, i.e. major scientific results passed to the other CU of our consortium - e.g. combined spectra, radial velocities, etc.

As volumetry is clearly dominated by the intermediate data, our main challenge is to manage and constrain this data. At each data release, we often have to find new solutions (new design, optimisations, etc.) in order to deal with the volume expansion.

In terms of output data, Gaia DR2 and DR3 are already available to all scientific communities from the Gaia Archive (<https://gea.esac.esa.int/archive/>). The total size of DR2 is 580GB and the total size of DR3 is about 10TB. These releases seem quite small. But it should be remembered that they only contain the golden nuggets of Gaia's data processing! Finally, the difference in volume between DR3 and DR2 is remarkable: DR3 is 17 times larger than DR2. The addition of a single measurement year cannot explain the difference in size between the two catalogue versions. In fact, this increase is due to the improvement of the scientific contributions. Between each release, the scientific algorithms become more complex and advance the understanding of Gaia's measurements. This allows for the production of new scientific results and therefore more volume to be managed by our DPCC.

At the end of the mission, DPAC will amass all the data produced by all DPCs -both the golden nuggets and some less valuable data- as well as all the telemetry or some intermediate data - into a legacy archive. The size of this final archive is currently estimated at 2 PB.

3.2 Complexity

Gaia DPAC operations plan is then very complex compared to many scientific space missions. The ground segment is centred at the Science Operation Centre (SOC) located in DPCE (Data Processing Centre ESAC), which acts as a 'hub' for the exchanges of data between all centres.

This centre gathered the raw and processed data in a Main Data Base (MDB) and sends it to each other DPC to start the processing of its CU, which can take from several days to several months according to the complexity of the software. At the end and after scientific validation of its own results, each DPC sends them back to ESAC. ESAC thus collects all the results, ingest them in the MDB-N and create the MDB-N+1 that will be the entry of the next cycle.

For a data release N, the work starts by grouping individual Gaia observations and matching them to sources on the sky, generating a working catalogue of sources, which is the basis for the downstream data processing chains and the heart of the MDB-N.

Each new release comes also with new or enhanced products, which means new chains to develop, integrate, validate and operate. Due to this iterative data reduction, the number of systems in operation in the different centres, the number of interfaces exchanged between them and the data volumes increase at each cycle.

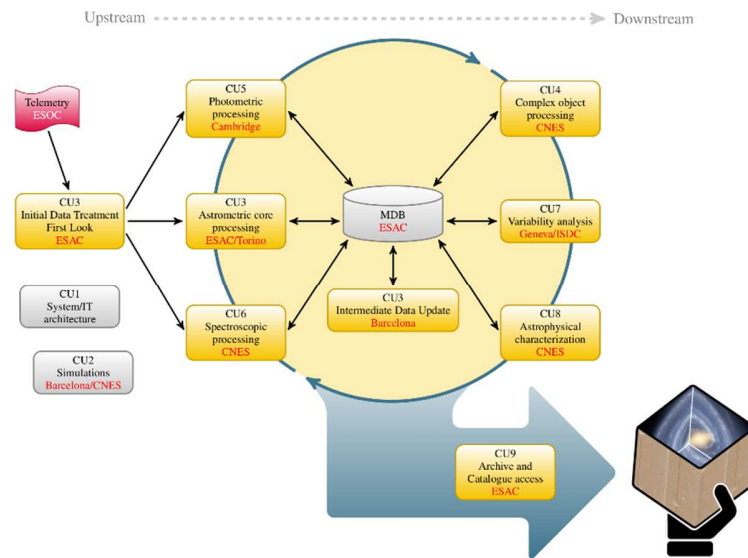


Fig. 5. The DPAC dataflow (DPACE courtesy)

Moreover, each time a new product enters the loop, its new interfaces (new tables in the MDB or parameters update in an existing table) are generally needed as input interface to the downstream chains: photometry (colours) computed by CU5 in Cambridge has been used by CU8 in Toulouse for computing temperatures. In fact, each chain improves its own results using results from the others, combining astrometry, photometry and spectroscopy parameters, so inter-dependencies between chains, CU and DPC are increasing at each cycle sometimes becoming circular:

- **Inter-cycle inter-dependency**
 - For example, CU6 data produced for MDB-N is used to produce CU8 data for MDB-N. Then CU8 data from MDB-N is used as an input data for CU6 chain during MDB-N+1 processing.
 - This strategy has no real impact on the operation plan on a processing cycle N.
- **Inner-cycle inter-dependency**
 - For example, during the cycle N, CU4 NSS chain has to run a first time to produced input data for CU6 chain. CU6 data are then an input data for a final run of CU4 NSS chain. All data are published in MDB-N.
 - This strategy is a constraint in the operation plan for the cycle-N.
 - This use case is managed only at DPCC.

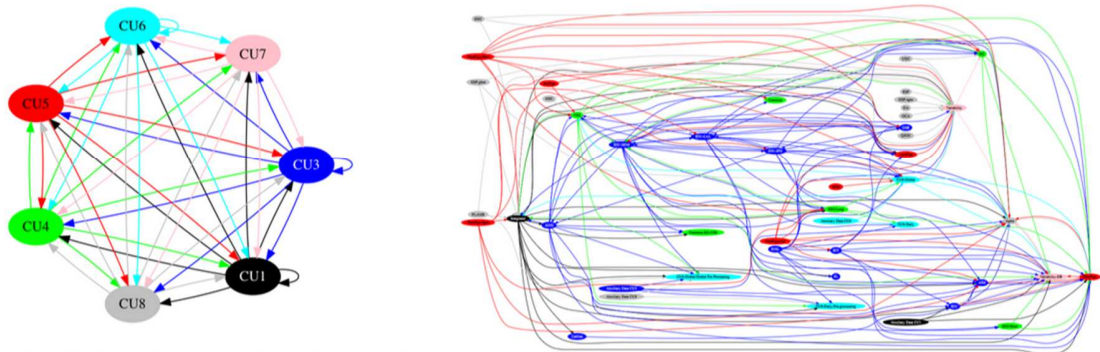


Fig. 6. Global and detailed top level CU inter-dependency for Cyclic Processing (Project Office courtesy)

Figure 6 gives an idea of the complexity at DPAC level, the Project Office team is in charge of the maintenance of global system inter-dependency and has to ensure the consistency of the system in a collaborative way between scientists of all the CU and DPC team. By hosting chains for three different CUs, DPCC has to handle a huge percentage of existing inter-dependencies.

3.3 Performance

Performance is at the centre of our concerns within the DPAC. The production of a data release for a DRC is clearly limited in time. Around 4 years have passed between the publication of DR2 and DR3. Around 4 years will pass between the publication of DR3 and DR4. Around 4 years are planned between DR4 and DR5. Yet the data acquisition period is almost doubled between each release: 34 months for DR3, 66 for DR4, 126 for DR5.

The planning should remain as much as possible on this window of about 4 years. It must not be affected by the increase in the number of records in entry or by new scientific use cases to be implemented.

The simple workaround would be to extend as much computing resources as necessary. However, like any project, Gaia is limited by cost.

The search for performance therefore revolves around:

- Schedule and operation plan at DPAC level and at DPCC level
 - at DPAC level, carried by the Project Office to manage interdependencies and the coordination of DPCs.
 - at DPCC level, to run in parallel all chains for CU4, CU6 and CU8.
- Technical improvements like more optimized data query or better allocations of resources for each processing pipeline.

The DPAC operation plan is carried by the Project Office. Its goal is to make sure that all the data are produced on time and that they comply with the needs of the downstream users. It is drawn up from the chronological sequence of chains (CU) for the current DRC, from the estimated duration of each of them, the integration in the MDB and data transfers durations, ending with catalogue extraction and validation (CU9). Microsoft Project software is used to gather all the information. The respect of delivery date is monitored by the Project office during a weekly web-conference with one representative from each DPC. Each DPC has its own provisional planning, used as an input to the DPAC schedule and vice-versa. Each DPC operational schedule is defined according to the maturity level of the three CU it manages and to the hardware and human resources it can rely on. Anyone shall inform if it is on-track or not as per the DPAC schedule. Because of inter-dependencies, any delay has a direct impact on the schedule of the activities in the other centres. Any major problem is discussed at DPAC level with all CU leaders and sometimes pragmatic choices are made and a science change can be withdrawn or postponed to the next processing cycle.

In a nutshell, the real performance of DPAC is to stick its operation plan in the DRC time frame as far as possible, despite the hazards. The real performance at DPCC is the ability to adapt its resources and teams to deliver data in time

3.4 Continuity

The ability to process twice the number of records in the same amount of time, without doubling the resources allocated, is only possible by maximising the use of all our resources. The challenge to be met, which goes hand in hand with the previous one "Performance", is continuity.

Our specification is therefore to have a data centre in operation 24h/24, 7d/7 with an availability rate approaching 95% over a year.

However, there are 2 temporal aspects to continuity:

- the day-to-day to achieve our short-term goals
- the long term to ensure that our cluster is still operational after 10 years.

On a day-to-day basis, the DPCC must be robust and resilient. These two characteristics are based on 3 fundamental pillars:

- The choice of technical architecture (described in section 4.1) and the associated redundancies were therefore crucial to achieving our availability target, despite our average of 1 incident per day.
- An organisation made up of an operational team to run the scientific pipelines and a maintenance team. There is constant communication between the two teams to share observations on the behaviour of the cluster. The ability to react quickly in the event of a major incident is optimal. Many activities and workarounds are now procedures. It should be noted that the teams are available on site only during the standard working hours of 5j/7 and off duty to reduce the costs.
- An adaptation of the operational plan according to the events observed and the presence of the teams. For example, during weekends, only validated and verified scientific pipelines are executed. Running only what is known in the absence of teams maximizes resource utilisation. Only a complete crash of the cluster could then have an impact on our scheduling optimisation goal. Given our availability goal and the stability of our resources, the risk is considered acceptable and does not require an on-call team. Adjustment of current runs is discussed in the Operation Control Group meeting every Thursday.

The DPCC has been operating for 10 years. This longevity would not have been possible if we had not kept our cluster up and running. Our system architect has therefore set up an annual roadmap to monitor hardware and software obsolescence.

This roadmap exercise has allowed us to remain aware of the essential evolutions to keep a system fully maintainable, possible evolutions to anticipate our future needs especially our performance needs and to take advantage of new functionalities or technologies. It has also allowed us to maintain the skills of our teams.

This roadmap covers both software, such as Java, and hardware. The DPCC can be seen as a technology hub that needs to be coherent and efficient.

“Copyright 2025 by the Canadian Space Agency (CSA) on behalf of SpaceOps. All rights reserved. One or more authors of this work are employees of the government of France, which may preclude the work from being subject to copyright in France, in which event no copyright is asserted in that country.”

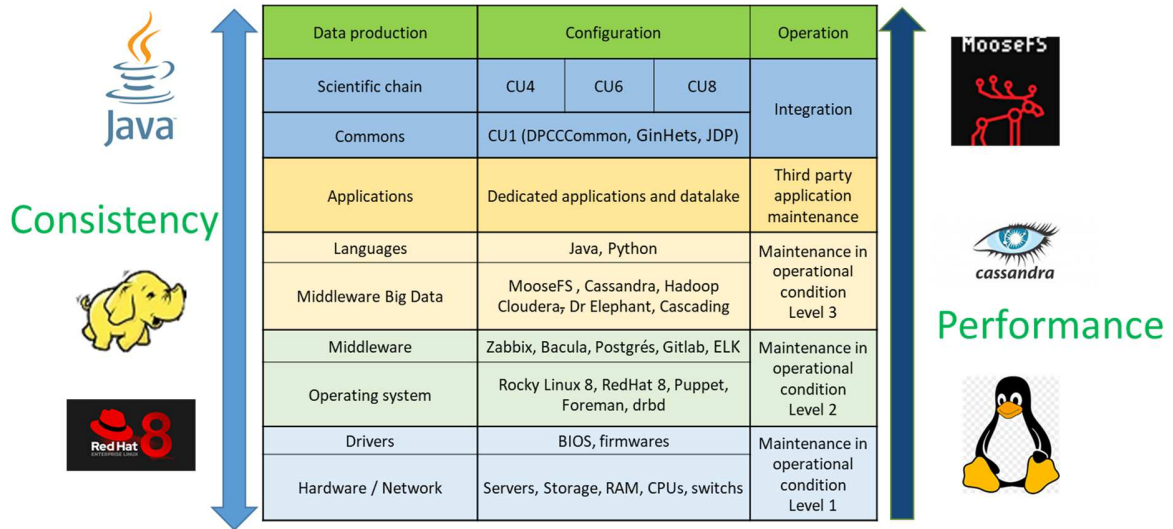


Fig. 7. DPCC technological « mille-feuille »

This follow-up has allowed us to limit the planned platform downtime, so that our availability target is not compromised, but is also taken into account in our operating plan.

The search for continuity is just as important as the continuity itself.

4. The keys of a success

4.1 Technological choices

4.1.1 A common programming language

The DPAC selected the Java programming language for processing Gaia satellite data in 2006. This decision was taken to ensure greater portability and efficiency in processing the complex data generated by the satellite. In addition to its popularity (c.f. figure 8), the use of an object-oriented language facilitated collaborative development on a software project as large as Gaia.




Fig. 8. Most popular programming language, ref. <https://www.tiobe.com/tiobe-index/>

Initially, DPAC people was reticent to adopt the new language, especially given the accumulated experience in other languages (FORTRAN or C) and the lack of libraries in Java. However, CU1 and DPCC people supported DPAC members through specific training and the implementation of simple environment (Ant, PMD, Checkstyle, etc.). Several years after, Java has been adopted by all DPAC members! Now, Java is still a viable programming language with various libraries and is mainly used in Big Data system.

As presented in section 2.3, the DPCC is in charge of the integration and operation of the following Java algorithms

- All algorithms used to characterize complex system (System Solar Objects, Extended Objects, Non-Single Stars, etc.) - CU4;
- All algorithms used to process wide band spectroscopic data - CU6;
- All algorithms used to characterize astrophysic objects - CU8.

These scientific algorithms are developed by several institutes spread across Europe (and even beyond). In order to be integrated at CNES, it was necessary to establish rules and development standards. In this way, the DPCC CUs were asked to develop their algorithms using the "facade pattern". The purpose of the facade was to hide internal complexity of the scientific module behind a common interface. In other hand, all data access layers are only managed by DPCC framework. Our philosophy has been to leave the Science to the scientists and to minimize the impact of technological changes on Science.

4.1.2 Our host framework

All scientific algorithms are integrated together to create jobs within a DPCC Host framework, namely SAGA (System of Accommodation of Gaia Algorithms). Its main functions are to provide the means to store and process a huge volume of astronomical objects, to orchestrate, chain and schedule complex scientific modules, and to provide a means for complex data queries. To deal with them, it is composed by

- **Phoebus product**, a generic CNES tool that gives the principal means for workflow orchestration and for process operation. SAGA has to handle and implement specific Phoebus interfaces and configuration. Phoebus is mainly composed of the supervision workflow engine coupled with a central MMI.
- **Gaiaproc component**, the main tooling developed to receive scientific modules. It defines the needed workflows and subsystems according to Phoebus language, provides the Job Framework library to execute scientific Facades and gives other means for controlling Hadoop jobs.
- **Platform level components**:
 - Softwares used by Gaiaproc in order to support storage and complex data querying needs;
 - A shared file system for scientific data and configuration storage;
 - Deployment, user management and monitoring tooling;
 - Scientific infrastructure softwares.
- **Gaiaweb**, a MMI developed for DPCC in order to gives a way for scientists to monitor their algorithm execution and check produced data.
- **CNES managed components**: Data Delivery Manager and Aspera client. These are tools used for file exchange between DPCs (presented in section 2.2), they rely on the platform components (Shared File system).

Since 2012, the SAGA architecture has been based on the Hadoop framework. The advantages of the Hadoop ecosystem are that it can easily handle petabytes of data and provides a framework for distributing processing across a cluster and managing cluster resources. The SAGA framework leads to the following implementation choices:

1. The use of distributed file system HDFS as data storage persistency,
2. Map/Reduce paradigm for job parallelisation and data manipulation. Its main drawback is the complexity of handling complex data structures and performing complex queries on the data.

3. The use of Cascading as Hadoop top layer for scientific job execution. It also allows simplification of data handling and processing, and provides tools to easily perform complex queries using the Hadoop framework.

Instead of writing directly process using mappers and reducers, Cascading provides a query API and query planner allowing the definition of data processing as flows. This technology has been selected to add an abstraction level on top of Hadoop, in order to simplify the writing of complicated queries similar to SQL joins. It has been preferred over other software like PIG and Hives, because it provides a Java API that allows all needed operations, and is easier to integrate into our SAGA software.

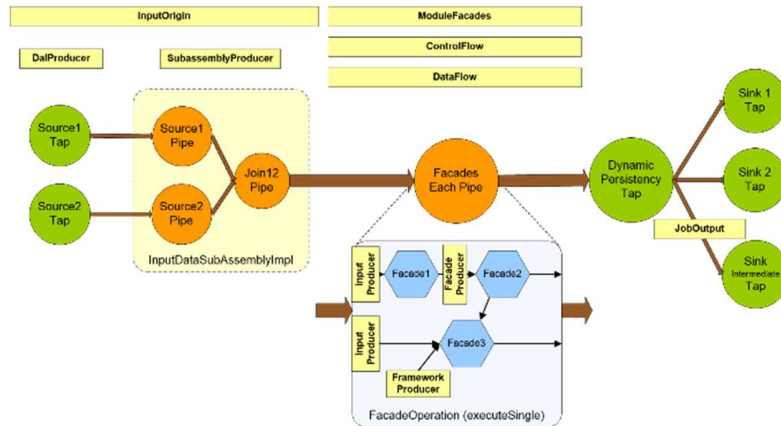


Fig. 9. Activity graph representing Cascading Flow at DPCC

An integrated scientific chain consists of several jobs. As illustrated in figure 9, each of these jobs is a dedicated cascading flow that can be divided into 3 elements:

1. **Data access** - it consists in Cascading flow input description (Source Tap), but also describes a part of the operation chain to build a coherent set of input data to be passed to façades.
2. **Façade chaining treatment** - this operation is generic and uses an automatically generated code (Java class) that chains scientific façade.
3. **Output data persistency** - it consists in the description of Cascading flow outputs (Sink Tap) thanks to an automatic persistency tool.

4.1.3 DPCC platforms

The DPCC hosts three computing clusters at Toulouse

- **DEV platform**, dedicated to integration activities (development and testing tasks);
- **VAL platform**, dedicated to acceptance and qualification activities;
- **OPS platform**, the biggest one, dedicated to qualification and operation activities.

As illustrated in figure 10, each of our Hadoop cluster is composed of Hadoop Master Nodes (NameNode and ResourceManager) and Hadoop Calculus Nodes. The network infrastructure is 10 Gbps.




Fig. 10. Schematic view of DPCC platform – picture of DPCC data centre

In 2024, the Operational platform consisted of:

- **Hadoop System** - 6 Hadoop Master Nodes and 225 Hadoop Calculus Nodes;
- **Distributed File System for data/application storages** - 4 MooseFS Manager servers and 4 MooseFS chunk servers;
- **No-SQL DB for Ephemeric storage** - 4 Cassandra Nodes dedicated to CU4 processing;

Totalling 6320 cores, 42 TB memory and 7625 TB disk capacity. After setting aside resources for the Operating system and for rotating maintenance of nodes, there remained roughly 4800 cores, 30 TB RAM. The Hadoop Job Tracker distributed the available resources to each chain in a fair way, respecting the maximum defined for each chain.

The total disk capacity is 7625TB, but 'only' 80% (6700TB) was to be used to avoid the local file system on the nodes becoming full during processing. In addition, the Hadoop file system security policy requires files to be tripled, which means we could store 2,200TB of usable data. Therefore, this disk space is allocated to the chain with the following quotas in 2024:

- 660 TB to store input data received from DPCE (daily and cyclic data reception);
- 160 TB for all CU4 runs;
- 650 TB for all CU6 runs;
- 500 TB for all CU8 runs;
- 500 TB for temporary Data (all chains).

Every year, we upgrade hardware to manage the obsolescence of the Calculus node and to meet the specific needs of each CU, especially when starting a new Data Release. In this way, we ensure that DPCC platforms are efficient and suited to future operations.

4.2 Organizational choices

Over the last 20 years, the DPCC's organisation has constantly adapted to the needs of the scientific teams, particularly in terms of development:

- At the start of the Gaia project at CNES, support was provided for the specification of processing chains and the development of scientific modules. We set up Java training meeting, recurrent quality control and a Hudson/Jenkins continuous integration platform. The aim was to deepen the scientific team's technical knowledge, while gaining an insight into their algorithms.
- Then, we set up pre-integration work by a team dedicated to each CU (design of the chains, more flexibility and back and forth with the scientists, first tests). These activities culminated in the integration of the resulting

chains into our DPCC framework (stricter, no evolution, etc.), carried out by an industrial company. During this period, we began to establish and apply an industrial process with the scientists to provide a better framework for our developments.

- Since 2021, the stabilisation/convergence of scientific chain designs has enabled us to merge development activities into a single integration phase.

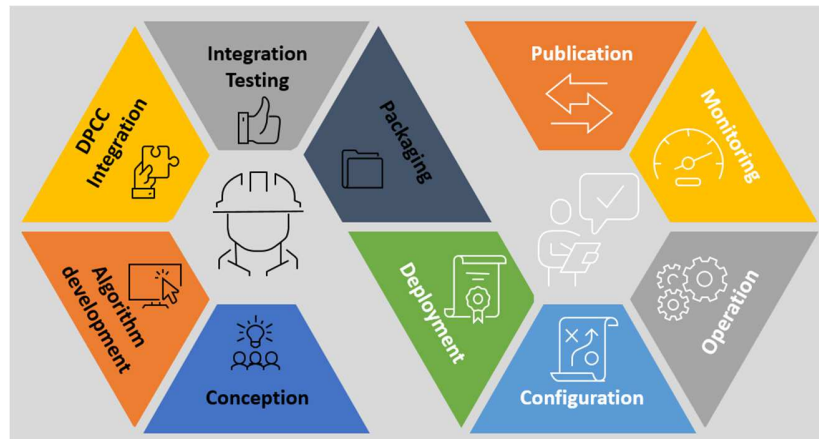


Fig. 11. DevOps at DPCC

As operations progress, our development cycle is shifting to a DevOps mode. As illustrated in figure 11, Development and operations teams work very closely together:

- **Development team**
 - Conception and design of the integrated chains at DPCC. Operation team experience gained from previous DRs is required to optimise/fix the design in line with the volume to be processed. This conception phase is only possible in collaboration with the scientific team.
 - Development of the technical and scientific algorithms. The DPCC team often implements some technical codes used to facilitate the integration of scientific algorithms. Each algorithm is delivered according to a integration schedule decided by CU leader and its CNES technical coordinator.
 - Integration of all algorithms in the DPCC Hadoop framework. It mainly consists in:
 - Deploy algorithm in the integration environment;
 - Implement Hadoop queries via Apache Cascading API;
 - Describe all algorithm interfaces in dedicated XML files.
 - Integration tests using JUnit (only on Hadoop queries) and using our DPCC workflow orchestrator (i.e. in real condition). These tests are completed with specific datasets (prepared by the development team) on a Hadoop cluster dedicated to the integration activities.
- **Operation team**
 - Deployment and configuration of the integrated chains on our Validation / Operation Hadoop clusters.
 - Acceptance / Qualification tests according to the qualification plan prepared by the CNES technical coordinator. During each test, operation team check performances and behaviours of the integrated chains. If there is any problem, the development team will be contacted to fix it (algorithm patch or integration patch required).
 - Operation and monitoring according to the operation plan prepared by the CNES technical coordinator. Here again, the development team will be contacted for any issue.

- Publish valuable results and execution reports from runs on our GaiaWeb platform in order to be downloaded by scientific team. These publications are crucial because they allow scientists to follow the progress of the runs at DPCC.

Our organisation allowed us to adapt to different interface changes and scientific algorithm configurations, and to be more flexible and responsive to any new ideas/needs. This flexibility means a large number of software patches to be managed at DPCC. Table 3 give an overview of these numbers ordered by CU and by DR. Our DPCC team manages approximately 70 patches (scientific or integration issues) per CU per year. This table also illustrates the increasing complexity of scientific chains (new functionality, optimization, etc.) between Data Release.

Table 3. Number of software patches managed at DPCC for DR3 and DR4¹

	CU4	CU6	CU8	Total
DR3	170	161	140	471
DR4	352	224	271	847

Today, a Data Release requires around 2-3 years of developments/qualifications before the operational phase at DPCC, which typically takes 1-2 years. As algorithms and designs stabilise and mature, the development phase can be accelerated to allow earlier qualification on new input data.

Over the years, DPCC team has built up a very strong relationship with the scientific team. This is mainly due to a two-headed CU organization shared between CNES and Scientific teams, illustrated in figure 12. For each DPCC CU, the scientific leader and his CNES technical coordinator are in constant contact to plan and coordinate all developments and tests.

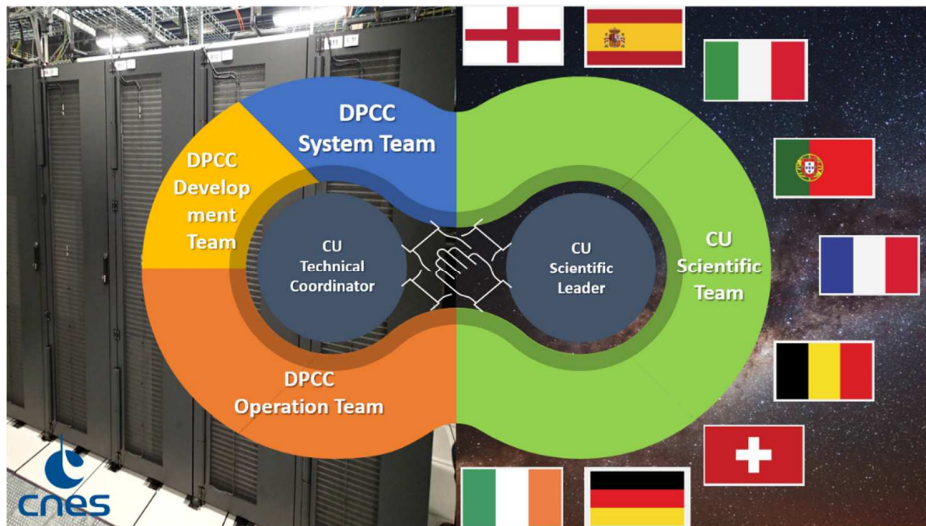


Fig. 12. Two-headed CU organization at DPCC

It's our recipe for success at DR2 and DR3. No doubt it will works for the next release!

4.3 The constant search of the right balance

Our work, in the end, is neither about technological choices nor organizational decisions that set a tempo or tone, but about the mash-up of five pieces that ultimately tell a story: the right balance.

¹ DR4 processing is still on-going at DPCC. So the number of patches will increase.

4.3.1 Technologic

For several years, we have been in the golden age of Big Data. This is especially true with the significant increase in daily data (due to human activities on networks). A study was carried out in 2011 to challenge the various data management technologies in relation to our Gaia requirements. This work evaluated the solutions according to 7 criteria: performance, ease of development, scalability, operability, reliability, durability and cost.

Technology	Performance	Development	Scalability	Operability	Reliability	Durability	Costs
Postgresql	★	★★★	☆	★★★	☆	★★★	★★★
Hadoop	★★★	★★	★★★	★★★	★★	★★★	★★
Caché	★★	★★★	★	★★	★★★	★★★	★
Cassandra	★★	★	★★★	★★	★★	★★★	★★
MongoDB	★	★★★	☆	NE	NE	NE	★★
PIProxy	☆	★★	☆	NE	NE	NE	★★
DB2	☆	★★★	☆	NE	NE	NE	★

Fig. 13. Global Synthesis

(NE = Not Evaluated because of no interest (destructive result on others criteria))

As shown in figure 13, the only technology that answered well all the criteria is Hadoop. It was both very scalable, and able to perform complex queries (even if it was not the simplest one). Hadoop had other advantages like the ability to control job distribution across the cluster. Its biggest drawback was the way data was stored in big files (no updates possible), but it was possible to adapt the current SAGA software in order to use Hadoop.

That’s why, CNES has decided to dive in Hadoop & Map/Reduce technology in 2012. Moreover, as shown in figure 14, Hadoop was really popular and consistent with our data volume challenge. DPCC was the first Big Data project at CNES!

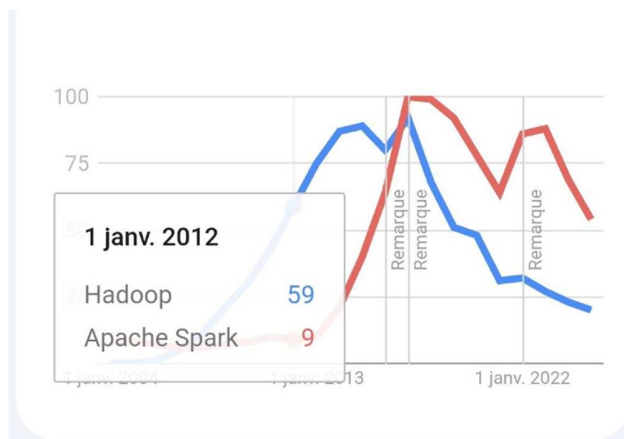


Fig. 14. - Hadoop .vs. Spark from Google trends

Since 2012, technologies have continued to improve in order to better manage these growing volumes. In parallel of the DR qualifications/operations, we carried out several proof of concepts (Apache Spark, Apache Flink and Apache Tez) to challenge our current system and to maintain the team's knowledge and mastery of Big Data. But the development costs were too high for the performance gains on our processing chain. What's more, we would have lost

our expertise in the Hadoop and Cascading technologies, particularly with regard to configuration tuning and the development of scientific queries. Changing our technology meant a major investment during Data Release production, which was too risky in term of DPAC schedule.

Today, DR4 is still on-going at DPCC and it is the second-last release. We are closer to the end of the DPAC processing. DPCC team is pragmatic. Our current Hadoop framework has proven to be efficient enough to produce Data Release data on time with the DPAC schedule.

4.3.2 Every resources it takes

To validate performance and control our schedule in line with the DPAC, we take advantage of the qualification phases to monitor our scientific chains. To do this, we use several Hadoop monitoring tools already described in SpaceOps-2021.5.x1298.

In addition, we have also developed 2 new tools at DPCC since 4 years:

- **Vigidoop** for having an overview of all processing chains launched on OPS platform. It gives us good indicators of the OPS platform use. An example of Core, RAM and Hadoop temporary monitoring is given in figure 15;
- **JobAnalysis** for analysing the computer performances of a given hadoop job. It builds several graphs very useful to find the optimized configuration needed by this job.

These DPCC tools are developed in Python and work from all hadoop statistics (number of reducers, physical memory, garbage collector, number of output data, etc.) directly on the OPS platform.

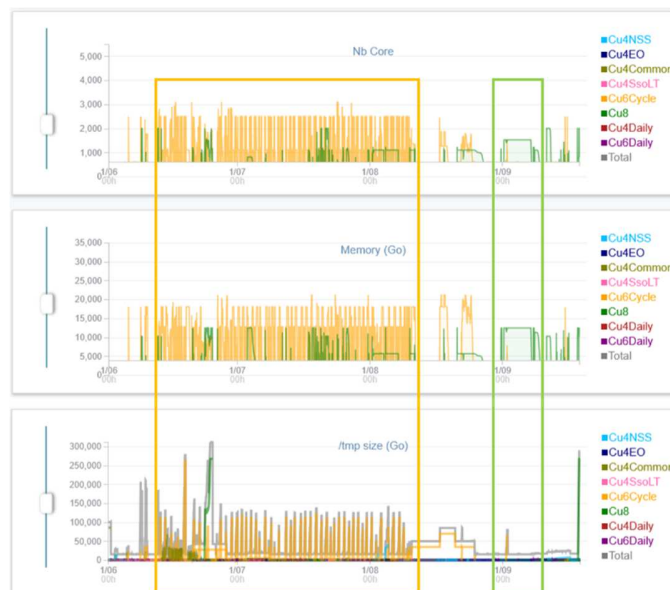


Fig. 15. Overview of CU6 and CU8 chains via Vigidoop
CU6 processing shown in orange box and CU8 in green box

As well as helping us on a day-to-day basis, all these tools allow us to assess the global performances of the processing chains and the use of the platform. As a result, we rely on them to refine our hardware/software needs for the coming years.

4.3.3 La Javanaise

As the French song "La Javanaise", the evolution of Java version is like a slow and languorous waltz torn between the desire of scientists to always have the latest version of Java and the nostalgia of the DPCC. But this evolution is mainly

driven by Java compatibility with our Hadoop distribution. Today, the Apache Hadoop distribution is only Java 11 compliant - evolution to Java 17 is not yet planned.

Java distribution is fixed at the beginning of each Data Release in order to stabilize the development scope of the scientific and technical teams. The higher usable version is fixed by each DPC. However, the version used on algorithms is particularly discussed on the scientific side. As the current level of coding doesn't take full advantage of Java's functionalities, and as it is still perfectible, this choice has not really impact on the performance of our chains. Anyway, with the data volume increase, our processing chains are more and more dominated by I/O (hard disk or network efficiency) and not by their real performances. Figure 14 perfectly illustrates it:

- **Orange box:** CU6 operations produces more than 100TB of temporary files created during Hadoop Map/Reduce queries (temporary means data cleaned at the end of the processing chain). Here, time is mainly spent in the read/write operation of these temporary files, instead of in scientific processing.
- **Green box:** CU8 qualifications mainly dominated by their scientific algorithms. Clipping effects seen in the Core and RAM graphs could indicate a performance issue in the scientific algorithms.

4.3.4 Hardware Better, Faster & Stronger

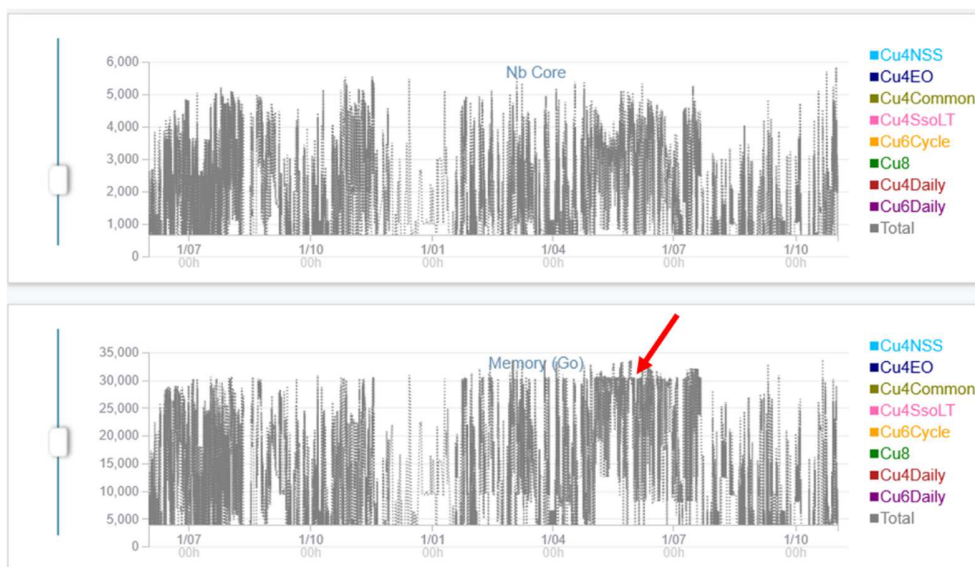


Fig. 16. Processing chains limited by RAM - red arrow illustrates clipping effect due to lack of RAM

To ensure a perfect efficiency of all DPCC platform, the DPCC System team works on the platform maintenance. An upgrade roadmap is yearly defined according to the monitoring results and the specific needs of each CU. After several years of qualification/operations, the goal of this roadmap is mainly to renew all obsolete nodes as follow:

- Reduce the number of nodes without major impact on the number of cores (new nodes have more cores than the oldest ones). At the end, it ensures less H/W maintenance;
- Increase disk space to deal with the increase of Gaia data volume presented in section 3.1;
- Increase RAM which is directly linked to the parallelization of scientific processing on our clusters (more RAM we need, less cores we use to parallelize our treatments). As shown in figure 16, scientific processing needs more and more RAM due to new functionality, new AI model used, etc.

4.3.5 Let it be

Since several years, we have proven that our technological choices make it possible to fulfil the DPCC mission:

- 7 Data processing chains have been successfully operated at CNES;
- All data have been produced on time with the DPAC schedule.
- Thanks to qualification plans, monitoring, etc., they have been operated with
 - the expected performances;
 - the expected numerical precision;
 - the allocated funds (the sinews of war).

As we are closer to the end of the mission than the beginning, each requirement that would entail a major change is studied in detail: benefit/risk analysis, cost, planning, possible workarounds, etc.

This guideline is best illustrated by the question of switching to Python for the next development cycle. As shown in figure 8, Python is currently the most popular language. However, a switch to Python would require a new interface definition for scientific algorithms, redevelopment of libraries, fine-tuning for operation and re-qualification of the cluster, and would imply a longer validation of results. All these steps are incompatible today with the objective of having the DR5 at the end of 2030. Moreover, the cost of such a modification would also be difficult to justify, given the long-term viability of the solution.

To sum up, our choices are now often geared towards the development of monitoring tools, back-up tools backed up by an operating procedure and therefore by a clear limitation on Hardware/Software evolutions at DPCC.

This pragmatic approach is summarized in the DPCC System Mantra « *No change, No risk, No delay!* » as long as our DPCC mission is still fulfilled.

6. Conclusions

The end seems near, but we're not yet facing the final curtain. The satellite has lived a life that's full: 10 years of scanning every source. And now, as tears subside after its passivation, the entire DPAC, and especially the DPCC, must rise the challenge of completing the mission with DR4 then DR5.

Throughout this paper, we have demonstrated our choices and our way of thinking—the keys to success rooted in efficiency and pragmatism as we tackled our four challenges: volume, complexity, performance, and continuity. Bad choices, we've had a few. But we planned each charted course, each careful step along the byway thanks to a pragmatic and efficient design-by-cost approach (days and euros)

By 2030, we hope to proudly declare, "We did it our way."

Acknowledgements



This article is the result of over 20 years of experience at the DPCC and collaboration with the DPAC. We would like to thank all the members of our team and the entire Gaia community.

References

- [1] Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., Brown, A. G. A., Vallenari, A., Babusiaux, C., Bailer-Jones, C. A. L., Bastian, U., Biermann, M., Evans, D. W., Eyer, L., Jansen, F., Jordi, C., Klioner, S. A., Lammers, U., Lindegren, L., Luri, X., Mignard, F., Milligan, D. J., ... Costigan, G. (2016). The Gaia mission. *\aap*, 595, eA1. <https://doi.org/10.1051/0004-6361/201629272>
- [2] Mercier, E., Els, S., Gracia, G., O’Mullane, W., Lock, T., & Comoretto, G. (2010). The project office of the Gaia Data Processing and Analysis Consortium. In G. Z. Angeli & P. Dierickx (Eds.), *Modeling, Systems Engineering, and Project Management for Astronomy IV* (Vol. 7738, p. e773810). <https://doi.org/10.1117/12.855764>
- [3] O’Mullane, W., Hernández, J., Hoar, J., & Lammers, U. (2009). Gaia Data Processing Architecture 2009. In D. A. Bohlender, D. Durand, & P. Dowler (Eds.), *Astronomical Data Analysis Software and Systems XVIII* (Vol. 411, p. 470).
- [4] O’Mullane, W., Lammers, U., Bailer-Jones, C., Bastian, U., Brown, A. G. A., Drimmel, R., Eyer, L., Huc, C., Katz, D., Lindegren, L., Pourbaix, D., Luri, X., Torra, J., Mignard, F., & van Leeuwen, F. (2007). Gaia Data Processing Architecture. In R. A. Shaw, F. Hill, & D. J. Bell (Eds.), *Astronomical Data Analysis Software and Systems XVI* (Vol. 376, p. 99). <https://doi.org/10.48550/arXiv.astro-ph/0611885>
- [5] Chaoul, L., & Valette, V. (2011). The CNES Gaia Data Processing Center: A Challenge and its Solutions. In L. Ouwehand (Ed.), *DASIA 2011 - Data Systems In Aerospace* (Vol. 694, p. e10).
- [6] Passot, X., & La Marle, O. (2010). CNES in GAIA data processing. In C. Turon, F. Meynadier, & F. Arenou (Eds.), *EAS Publications Series* (Vol. 45, pp. 89–94). <https://doi.org/10.1051/eas/1045014>
- [7] O’Mullane, W., Luri, X., Parsons, P., Lammers, U., Hoar, J., & Hernandez, J. (2011). Using Java for distributed computing in the Gaia satellite data processing. *CoRR*, abs/1108.0355. <http://arxiv.org/abs/1108.0355>
- [8] Guiraud, J., & Roux, W. (2022). New Questions Opened by the Big Data in the World of the Science Data Processing Centre for Gaia Mission in CNES. In C. Cruzen, M. Schmidhuber, & Y. H. Lee (Eds.), *Space Operations: Beyond Boundaries to Human Endeavours* (pp. 309–324). Springer International Publishing. https://doi.org/10.1007/978-3-030-94628-9_14