

Italian Gaia Data management, processing infrastructure and Legacy

A joint collaboration between INAF and ALTEC



Enrico Licata
on behalf of INAF & ALTEC teams



INAF USC-C - General Assembly
March 09 – 13, 2026
Savoia Excelsior Palace - Riva del Mandracchio, 4, 34124 Trieste TS

The Team



Lorenzo Bramante
Ruben De March
Daniele Gontero
Rosario Messineo
Luigi Squillante
Leonardo Tolomei

Former Members

Fabio Filippi
Andrea Fonti
Enrico Pigozzi
Marco Vaschetto



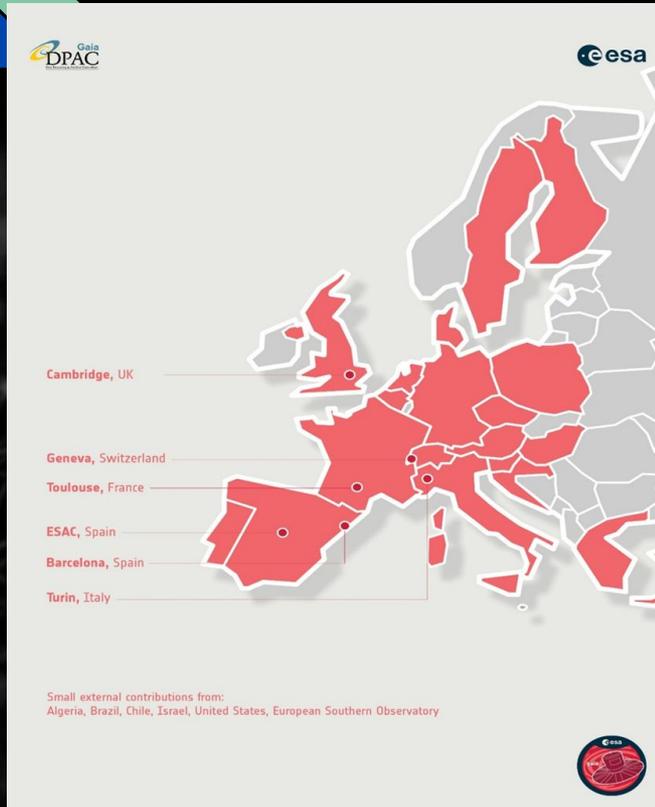
INAF
ISTITUTO NAZIONALE
DI ASTROFISICA

Ummi Abbas
Beatrice Bucciarelli
Deborah Busonero
Raffaella Buzzi
Mario G. Lattanzi
Enrico Licata
Roberto Morbidelli
Alberto Vecchiato

Former Members

Luca Bianchi
Marco Pecoraro

DPCT & The DPCs



Each DPC has its own infrastructure and it's responsible for the execution of specific software components, covering all or part of the need of a single or multiple Coordination Unit (CU). These are:

- DPCB (Barcelona)
- DPCC (CNES)
- DPCE (ESAC)
- DPCG (Obs. Geneva / ISDC)
- DPCI (IoA, Cambridge)
- DPCT (Turin)

DPCT: one of the 6 data processing centers (DPCs) of the Gaia SGS, hosted in ALTEC in Turin. Under dedicated ASI industrial and scientific contracts, its construction and operation is the result of the work of an integrated INAF-OATo / ALTEC team.

The DPCT provides the infrastructure and operations support to the Astrometric Verification Unit (AVU) activities for CU3

CU3 and the CUs

DPCT provides the infrastructure support (in terms of HW, DB and software framework) to run the CU3 sw systems that are part of the AVU (Astrometric Verification Unit): AIM, BAM and GSR

AVU is the unit responsible for the development and maintenance of the following three CU3 software systems:

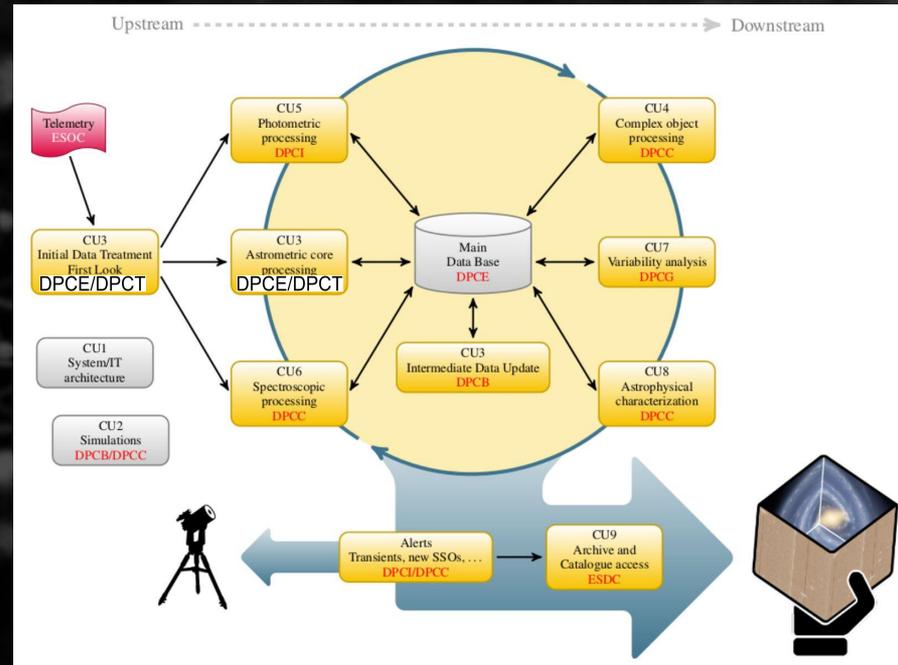
AIM - Astrometric Instrument Model Data Analysis Software System, in charge of processing the Astro data telemetry in order to monitor and analyse the Astro instrument response over the mission lifetime.

BAM/AVU - Basic Angle Monitoring software system, in charge of processing the BAM telemetry in order to monitor and analyse the BAM behaviour over time.

GSR - Global Sphere Reconstruction, the mathematical and numerical framework that shall be used to verify the astrometric mission products produced by AGIS

CU1: System Architecture
CU3: Core Processing
CU5: Photometric Processing
CU7: Variability Processing
CU9: Catalogue Access

CU2: Data Simulations
CU4: Object Processing
CU6: Spectroscopic Processing
CU8: Astrophysical Parameters



Language, Versioning, Repository



To create a **uniform environment** and **improve the software integration between partners** JAVA was chosen as the standard language for the project across all DPCs and CUs to leverage its portability due to the JVM



The chosen versioning system is SVN, **using the standard structure of Trunk, Branch and Tags** to manage versioning and releases of sow across all CUs



All libraries required by the software developed in the scope of the project is maintained on a Nexus repository hosted by ESA. **Only the libraries available on the repository can be used in the released code**

Data Model & Dictionary Tool

MDB DM which is used by all DPCs different DMs for each of the DPCs, giving flexibility to each operational center.

The DM itself is a versioned SW, with trunk, branches and tags

The dictionary Tool:

- allows the creation and editing of DM objects
- is shared by all DPCs and allows consultation of all the DMs of all the DPCs
- Each entry/table contains: fields, descriptions, types, multiplicity, units
- Each table has a table level description and a table level history
- All these informations can be edited and committed to a remote server

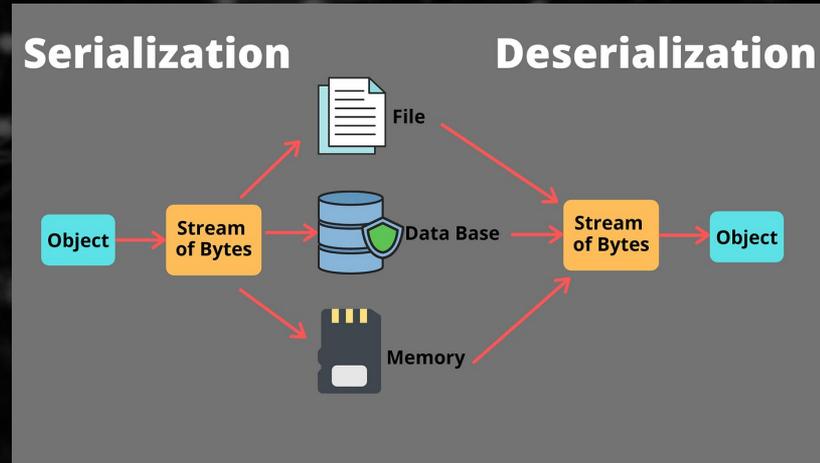
The screenshot shows the 'DPC/UC3/DPCT Trunk Data Model' application. The left pane displays a tree view of the data model structure, including 'MDB DM', 'DPC/UC3/DPCT DM', and 'AIM' elements. The main pane shows a table description for 'DPC/UC3/DPCT/AVU/AIM/AimWrapper'. The table lists various metadata fields with their descriptions, data types, and units.

#	Name	Description	Det. Desc.	Type	Multiplicity	Units
1	solutionId	Solution Identifier	View	long		
2	creationDate	Creation Date of the Data	View	double		Time[Julian Date (day)]
3	gClass	Window class used in SM, AF an...	View	byte		
4	priority	Priority of the observation for its tr...	View	short		
5	objectType	Object type	View	byte		
6	aocsUpdate	AOCS update occurred in AF2-9	View	byte		
7	gMag	G magnitude estimated by the VPU	View	short		Magnitude[mag]
8	ccdRow	The ccd Row in the processed As...	Add	int		
9	isFov1	Identify the Fov (0/false or 1/true)	Add	boolean		
10	smFlux	SM Flux in G band	View	double		Flux[e-s]
11	smFluxError	Error in SM Flux	View	double		Flux[e-s]
12	smAlCentroid	Along scan coordinate in SM	View	double		Time[TDI Period]
13	smAlCentroidError	Error in along scan coordinate in ...	View	double		Time[TDI Period]
14	afAcWinCoord	Across scan coordinate of the AF...	View	short	[CCD_NUMBER_AL[0]]	Length & Distance[pixel]
15	afBackground	Background estimated in AF Cent...	View	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
16	afBackgroundError	Background error estimated by A...	Add	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
17	afBckgMetadata	Background estimated by AF Cen...	Add	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
18	afBckgMetadataError	Background error estimated by A...	Add	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
19	afFlux	AF Flux in G band	View	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
20	afFluxError	Error in AF Flux	View	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
21	afFluxMetadata	AF Flux in G band metadata	Add	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
22	afFluxMetadataError	Error in AF Flux metadata	Add	double	[CCD_NUMBER_AL[0]]	Flux[e-s]
23	afCentroid4d	Along scan coordinate in AF	View	double	[CCD_NUMBER_AL[0]]	Time[TDI Period]
24	afCentroid4dError	Error in along scan coordinate in ...	View	double	[CCD_NUMBER_AL[0]]	Time[TDI Period]
25	chi2	Chi2 of the fitter	Add	double	[CCD_NUMBER_AL[0]]	
26	chi2Metadata	Chi2 metadata of the fitter	Add	double	[CCD_NUMBER_AL[0]]	
27	doF	Degree of freedom of the fitter	Add	int	[CCD_NUMBER_AL[0]]	
28	doFMetadata	Degree of freedom metadata of L...	Add	int	[CCD_NUMBER_AL[0]]	
29	goF	Goodness of Fit	Add	double	[CCD_NUMBER_AL[0]]	
30	goFMetadata	Goodness of Fit Metadata	Add	double	[CCD_NUMBER_AL[0]]	
31	afAcCoord	across-scan coordinate in AF (for...	View	double	[CCD_NUMBER_AL[0]]	Length & Distance[pixel]
32	afAcCoordError	standard error in across scan co...	View	double	[CCD_NUMBER_AL[0]]	Length & Distance[pixel]
33	waveNumber	Effective wavenumber	View	float		Frequency[nm ⁻¹]
34	waveNumberError	Error in Effective wavenumber	View	float		Frequency[nm ⁻¹]
35	ssc	Spectrum Shape Coefficients	View	float	[8]	

Data Format - GBIN

GBIN is a specialized, Java-serialized file format used by the Gaia Data Processing and Analysis Consortium (DPAC) to store astrometry data. It requires specific Gaia data model classes on the classpath at runtime for decoding. It is not a standard, general-purpose file format, but rather a niche, data-driven structure.

- Developed for handling data from the Gaia astrometry satellite during operations
- **Compressed serialized java objects**
- Requires a specific java sw (Gaia Tools) with the correct Data Model classes to be able to access the data.
- **Can also be accessed via reflection in tools like TOPCAT**
- Lightweight and suitable for operations
- **Not suitable for scientific exploitation and dissemination**



ECSS standard

European Cooperation for Space Standardization: is a collaboration between the European Space Agency (ESA), the European space industry represented by Eurospace, and several space agencies, to develop and maintain a coherent, single set of user-friendly standards for use in all European space activities

It is a comprehensive system of standardization that includes a tailoring process to help in elaborating the ECSS Applicability Requirements Matrix

For example with respect to the software production, GAIA project adopted the following documents:

Defined at the planning stage:

- Software Design Document (SDD)
- Software Requirements Specifications (SRS)
- Software Test Specifications (STS)

Updated with each sw release:

- Software Release Notes (SRN)
- Software Test Report (STR)
- Software User Manual (SUM)



AIM 22.0 Software Release Note

22.0.6 patch release

prepared by: E. Licata, D. Busonero
approved by: DPCT CCB
authorized by: M. G. Lattanzi
reference: GAIA-C3-SP-INAF-ELC-005-03
issue: 03
revision: 5
date: 2020-10-20
status: Issue

More than 10 years of continuous science Operations !!



SKY-SCANNING COMPLETE FOR ESA'S MILKY WAY MAPPER GAIA

From 24 July 2014 to 15 January 2025, Gaia made more than three trillion observations of two billion stars and other objects, which revolutionised the view of our home galaxy and cosmic neighbourhood.

580 MILLION

Accesses of Gaia catalogue so far



13 000

Refereed scientific publications so far



2.8 MILLION

Commands sent to spacecraft



142 TB

Downlinked data (compressed)



500 TB

Volume of data release 4
(5.5 years of observations)



3 TRILLION

Observations



2 BILLION

Stars & other objects observed



938 MILLION

Camera pixels on board



15 300

Spacecraft 'pirouettes'



55 KG

Cold nitrogen gas consumed



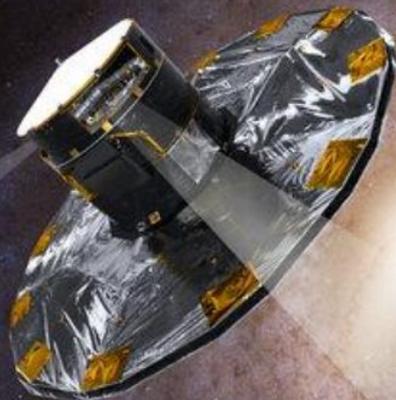
3827

Days in science operations



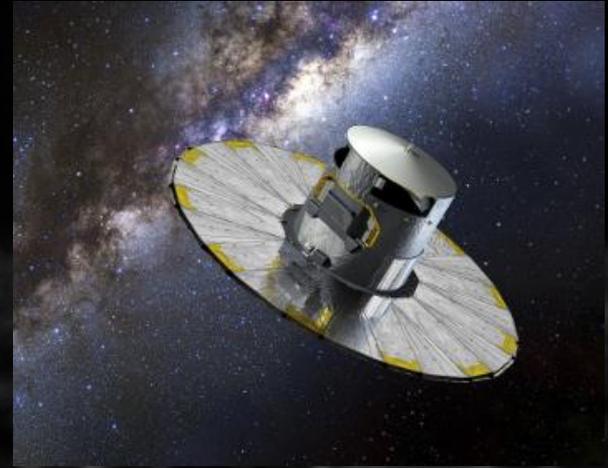
50 000 HOURS

Ground station time used



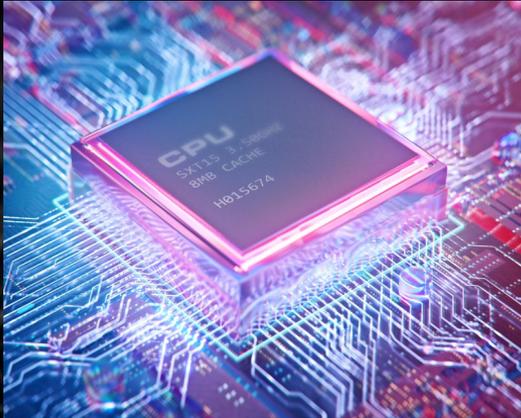
A few info on Gaia and its Data

- Launched on 19/12/2013
- Expected duration ~5 years
- Actual duration 10+ years
- End of Gaia scientific data collection: March 2025
- Observed and measured 2+ billion sources
- Each source has been observed on average ~200 times
- Each source transit is comprised by
 - telemetry data
 - actual raw data (6.5k ~ 35k windows, 18x12 px, every 4.41s for every one of the 62 AF CCDs)
 - data and metadata generated by multiple processing pipelines across europe (9 CUs)
 - and much more...
- Estimated size of the full dataset @DPCT (including outputs) ~ 3PB
- Estimated size of the full dataset at the end of data reduction in 2030 ~ tenth of PB



*Raw data are
NOT images*

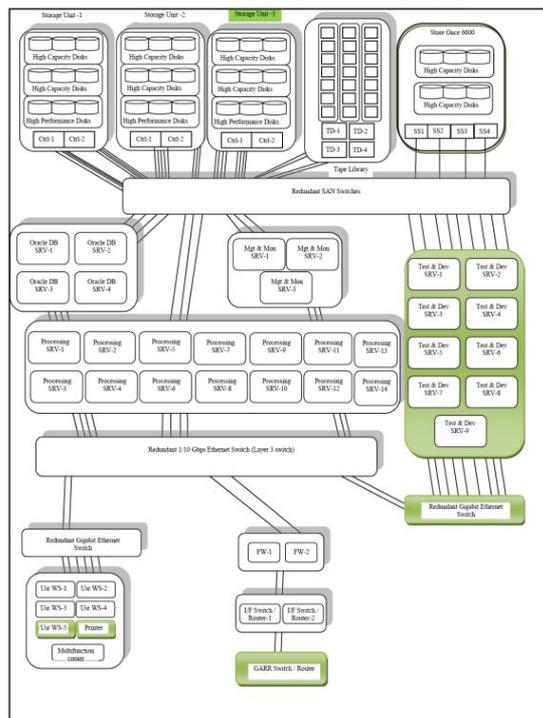
Processing Requirements @ DPCT



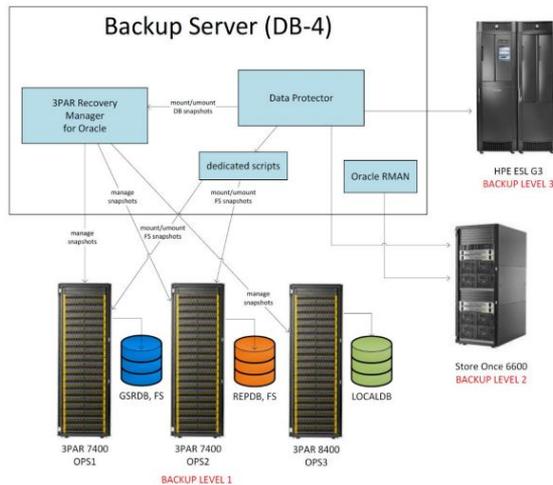
1. Handle a massive dataset: $10^6 \sim 10^7$ transits / day
2. Pipeline integrated DBMS capable of supporting OLTP (online transactional processing) and OLAP (online analytical processing) operations
3. Max processing time for Daily pipeline < 24h
4. Heterogeneous workloads
 - Computation
 - Monitoring
 - Visualization
5. Store all mission data and all produced outputs of both daily and cyclic processing

HW infrastructure 1/2

DPCT Overall HW infrastructure



Backup Server (DB-4)



Credits: ALTEC

INTERNET LINK: 10Gbps (300 Mbps guaranteed) via GARR

STORAGE CAPACITY: about 2.6 PB overall rawdisk space distributed among 2 HP P7400 and 1 HPE P8400, HPE Nimble HF40, Oracle ZFS and Oracle ODA

COMPUTING: 14 servers HP DL580 G7/G9 with a total of about 600 CPU cores and 4.5TB RAM.

DEV & TEST: 7 servers HP

DB SERVERS: 3 HP DL580 G10 dedicated to the database cluster based on Oracle RAC technology

NETWORK CONNECTION: LAN network up to 10 Gbps. SAN network redundant at 8/16 Gbps.

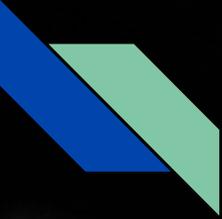
SECURITY SERVICE: redundant firewall on 2 HPE DL380 G10 servers for replacement. based on pfSense, enabling secure remote access via VPN.

INFRA MONITORING AND MANAGEMENT: services based on VMWare virtual environment configured with 2 HP DL580 G7 servers clustered and managed by vCenter Server.

BACKUP SERVERS: HP DL580 G7 dedicated to database and file system backups from data volume snapshots and HPE DL580 G10 servers for replacement.

3 LEVELS BACKUP: L1 on primary storage array, L2 on disks (StoreOnce 6600) and L3 on tape libraries (HP ESL G3).

HPC INTERCONNECTION: access to HPC supercomputer at CINECA for dedicated processing



HW infrastructure 2/2

The DPCT hardware infrastructure is based on a **distributed environment**, including a **computational grid**, a **database grid** and a **storage area network**.

The **DBMS choice is Oracle** that provides advanced availability and scalability features. Oracle allows multiple computers to run the Oracle DBMS software simultaneously while accessing a single database, thus providing a **clustered database**.

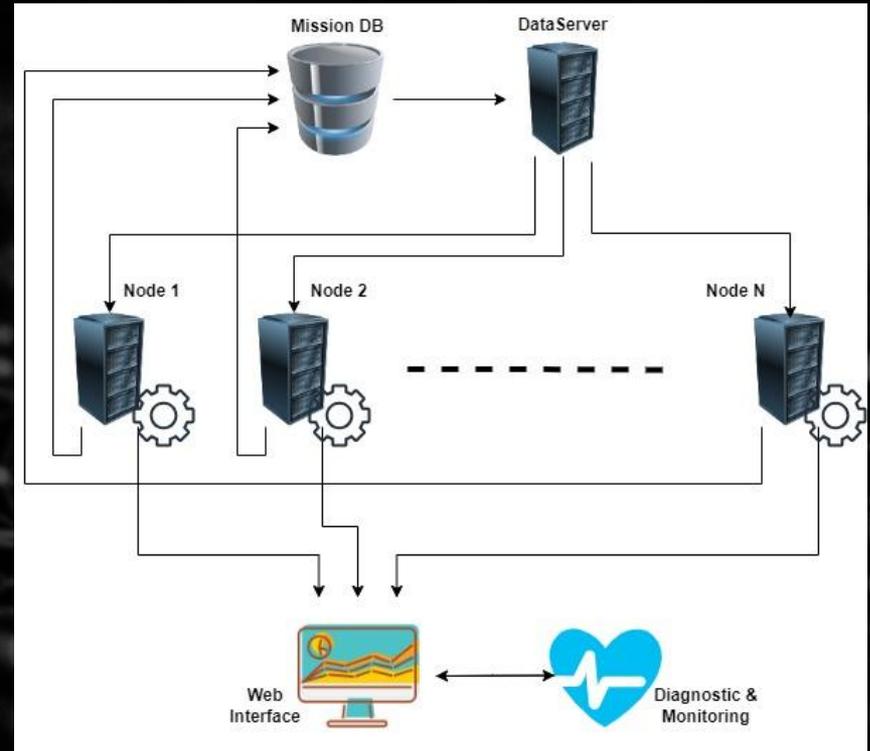
The database grid will use the following Oracle products:

- Oracle Server
- Oracle RAC
- Oracle Partitioning
- Oracle ASM to manage storage used by database
- Oracle Spatial

Simplified Functional Schema

For the execution of each production run:

- Data is retrieved from **Mission DB** and sent to the **DataServer (DS)**
- The **DS** performs checks, selects data, creates data packages to be sent to the processing nodes
- **Each node runs the scientific SW.**
 - Statistical data collection is performed at this level
- Outputs produced by the nodes is used to populate the **Web Interface**, and to perform **Monitoring and Diagnostic Operations**, while being **stored back to the Mission DB**



Production & Test Environments

To be able to develop your project you also need an appropriate dev and testing environment, together with a solid SW life-cycle



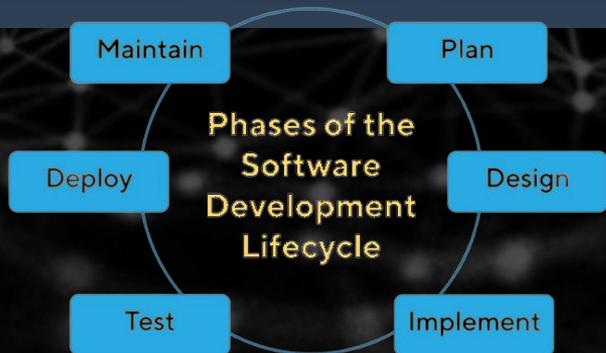
PROD ENV: 14 servers with a total of about 600 CPU cores and 4.5TB RAM.

DEV & TEST: half as much (is it enough?)

- Ensure consistency with sw baseline and configuration
- Unit Tests → Use Cases → System Test → Integration Test → Prod
- SRN → STR → Release
- ALL documentation adhere to ECSS standards (European cooperation for space standardization)
- Test automation via Jenkins + Ant Tasks
- 3 SW systems competing for HW resources

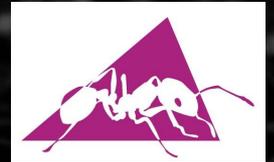
Considerations

- Is it reasonable / necessary to have a full scale test env?
- Plan for a procedure to keep Prod and Test envs configuration aligned



Workflow - CD/CI - Issue Tracking

- Centralized management utility VMware vCenter: deployment of VMs across the servers
- Ingestion, Archive, and daily Workflow, are automated using custom GAIA SW system developed in Java
- A monitoring web interface updated in real time (in house sw)
- All main project's branches on SVN (both for infrastructure and scientific SW) are monitored by a Jenkins instance that manages
 - svn checkouts
 - compilation
 - tests with apache ant tasks
- Daily pipelines programmatically generate a report each day
- Issue tracking is managed with JIRA



GAIA
TOOLS

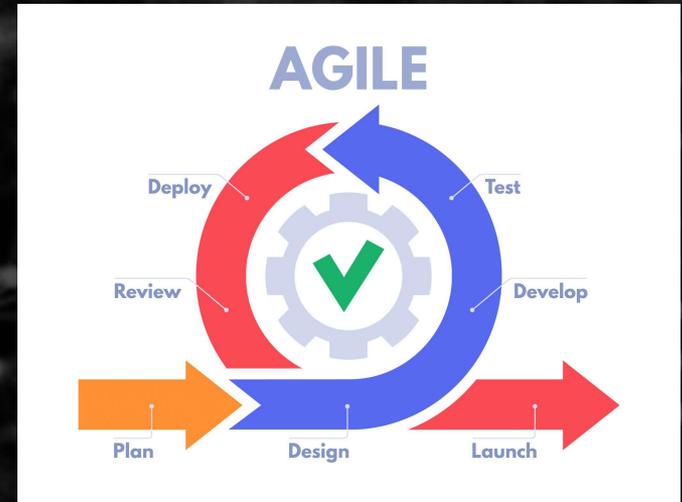


Agile Project Management

A project management approach that involves breaking the project into phases and emphasizes continuous collaboration and improvement.

- Periodic “STAFF Meetings” with whole team: 1/month to keep everyone engaged and discuss procurement and contractualization between all parties involved
- **Operations briefing:** checking the status of operations, bi-weekly
- **Scrum / Dev meetings:** among INAF and ALTEC developers to tackle specific technical issues

The frequency of each of these meetings has been adjusted during the project’s life-cycle (Planning /Commissioning / Early Operations / Operations) with the intent of maintaining minimal overhead



Infrastructure Software

- DPCT_SoftwareSystem
 - manages workflow
 - node deployment
 - db interaction
- DPCT_SoftwareConfiguration
 - manages the conf of daily and DRC workflows
 - separate conf for all project modules
 - plot descriptors
- DPCT_ConnectorLibrary
 - Integration between the Infra SWs and the scientific SWs

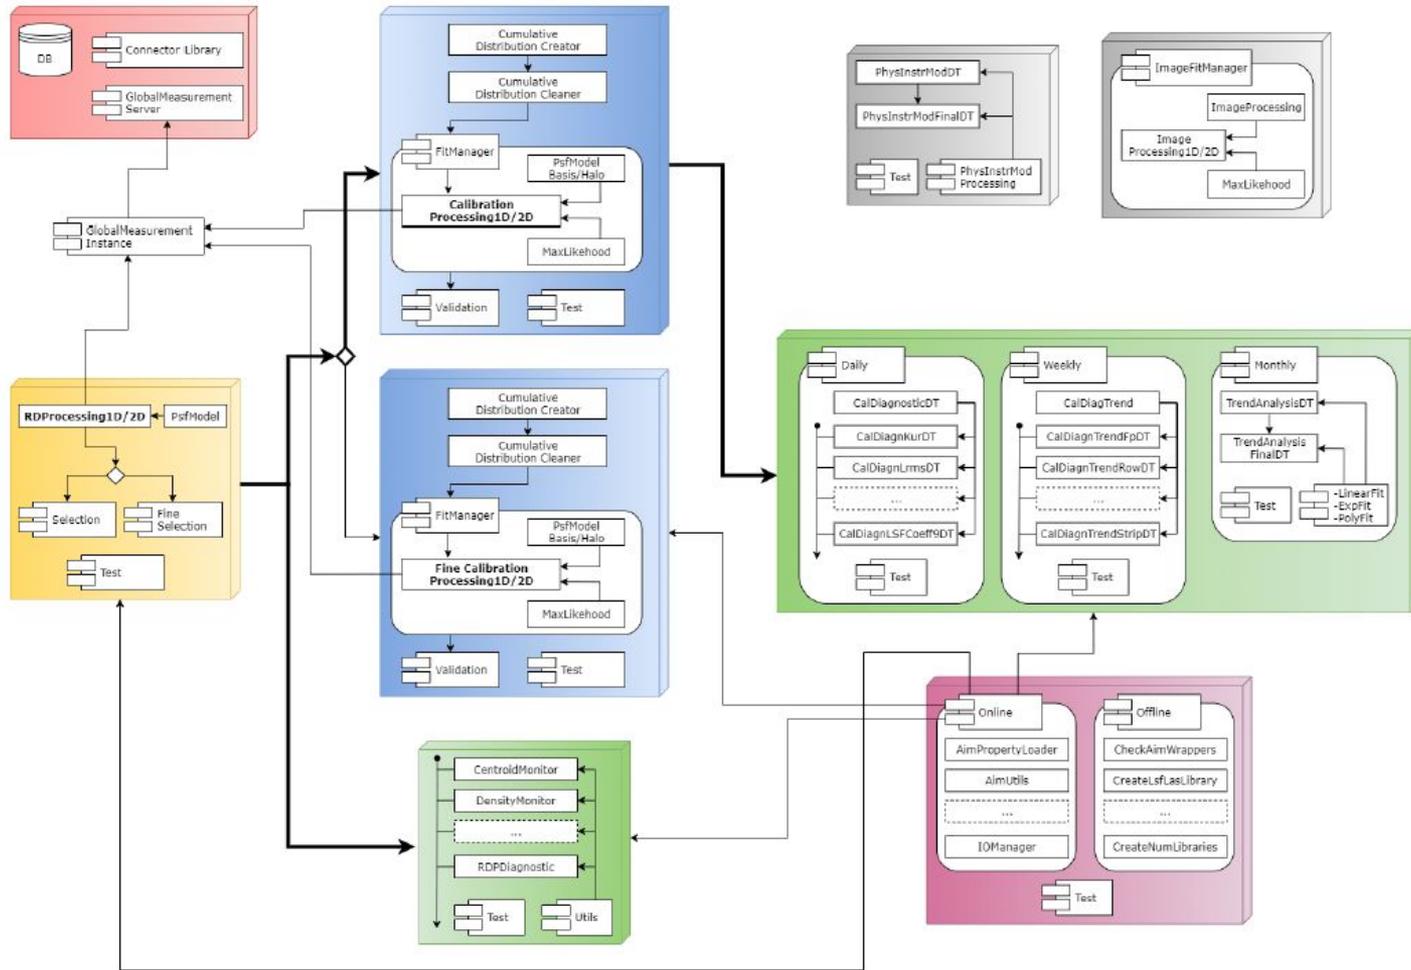
The integration between Infra and Scientific code managed by the ConnectorLibrary required:

- A Shared effort between INAF and ALTEC on specific parts of the source code
- Required tight cooperation and coordination between research institution and industry for the development and maintenance

Scientific Code - Overview

- 5 pipelines:
 - **AIM** - Astrometric Instrument Modeling - daily/cyclic
 - **BAM** - Basic Angle Monitoring - daily/cyclic
 - **GSR** - Global Sphere Reconstruction - cyclic
- Parallelization is managed by the infrastructure
 - but many parts have been parallelized to allow testing of large batches of data
- Source code is under **version control using SVN** (chosen by ESA and imposed to the Consortium)
- Implements established design patterns such as: Chain of Responsibility / Singletons / Abstract Factory
- **The data packages send to the processing nodes are not stored "as is"**: they are runtime items, built from an underlying db and **following a specific DM**
- the whole AIM and BAM pipelines are performed locally on the DPCT infrastructure
- Computational requirements for GSR exceeds the capability of the DPCT
 - the SOLVER is executed on LEONARDO

HIGH LEVEL ACTIVITY DIAGRAM of AIM PIPELINE:

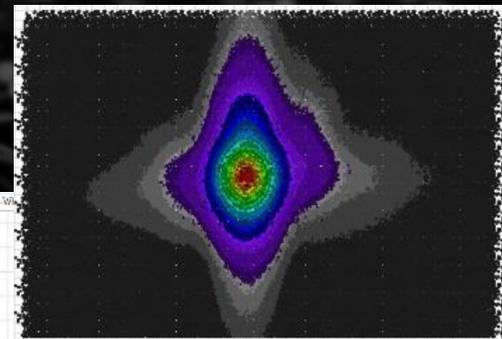
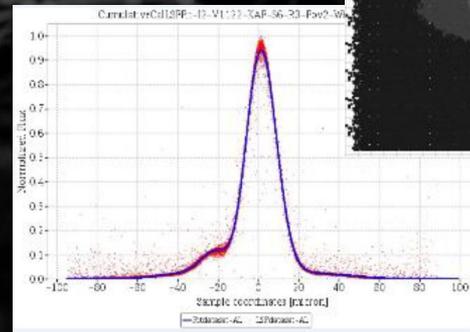
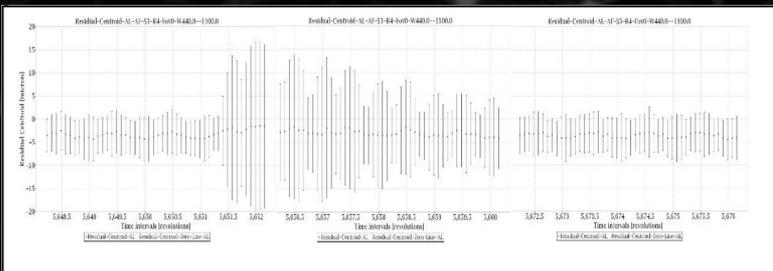


AIM - Astrometric Instrument Modeling

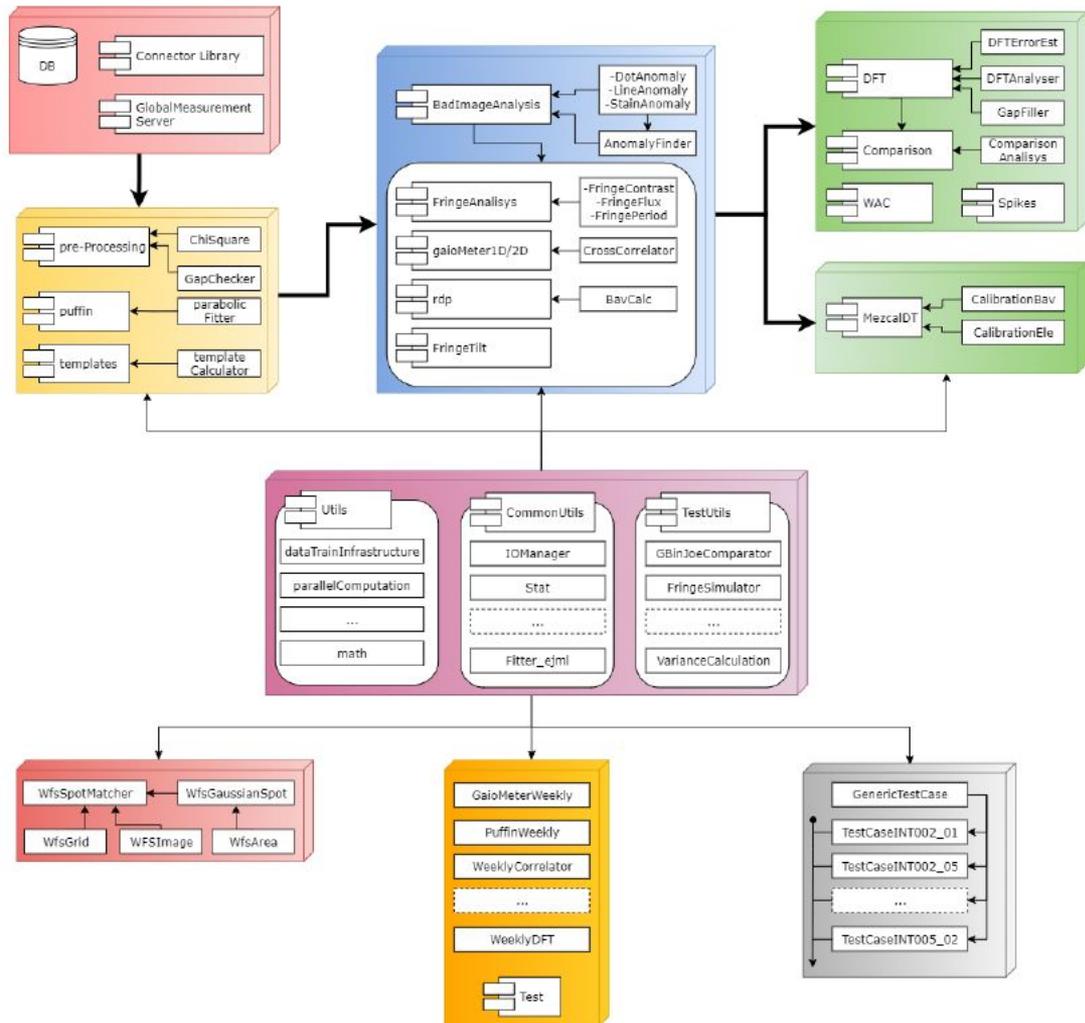
AIM daily pipeline: Raw data processing, image parameter determination, LSF/PSF modelization and calibration, astrometric instrument monitoring and diagnostics throughout the mission lifetime

- Daily sw version 100.000+ lines of code
- Cyclic sw version about 100.000+ lines of code
- 24 hours of raw data each run : from 2×10^6 to 15×10^6 raw images transits
- Complex structure of the pipeline: 10 sw modules managed by an unsupervised coordinator, the output of one run become the input of the next one
- ~6 hours of time execution on the DPCT Operation platform for each run

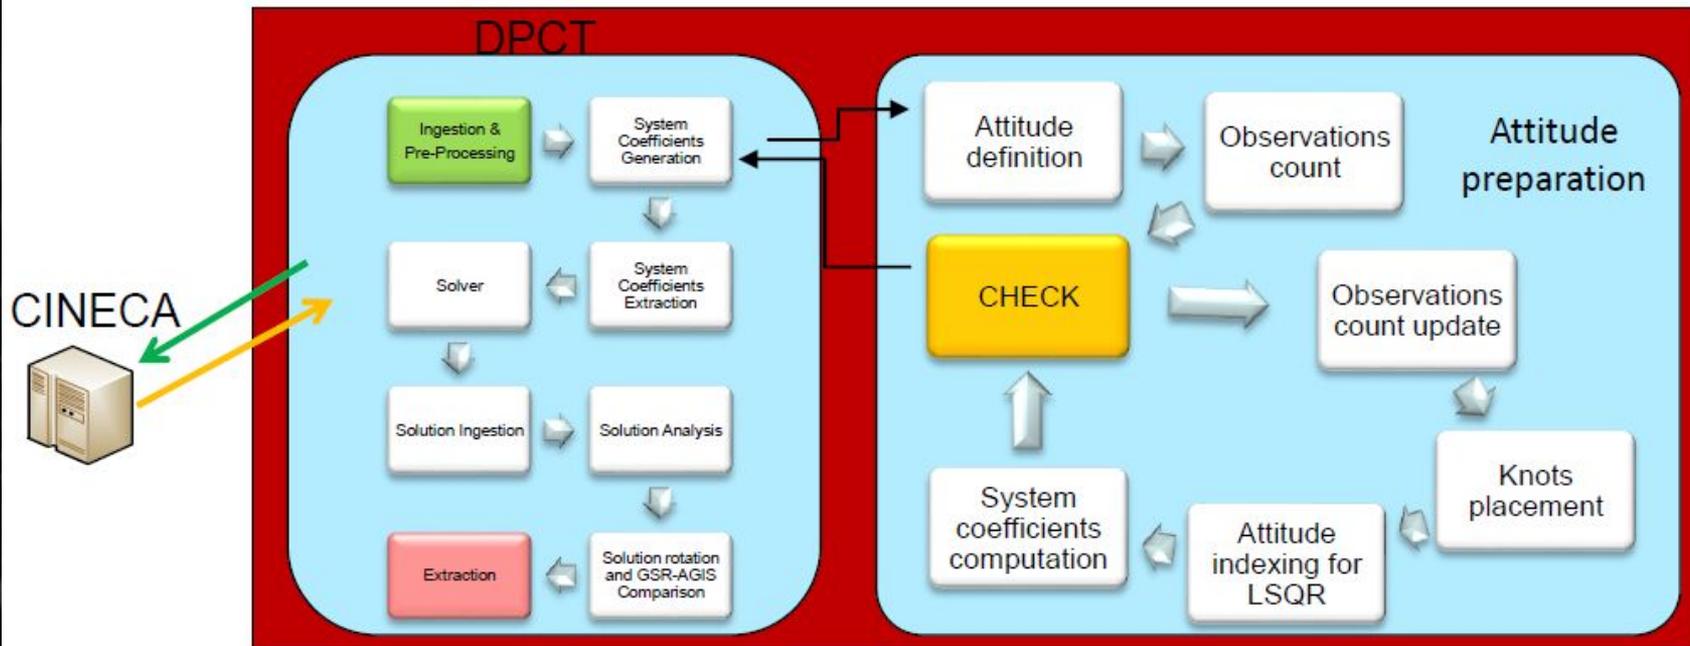
Cyclic version of AIM the software aims to reprocess images for calibration improvement and refined image parameters calculation



HIGH LEVEL ACTIVITY DIAGRAM of BAM PIPELINE:



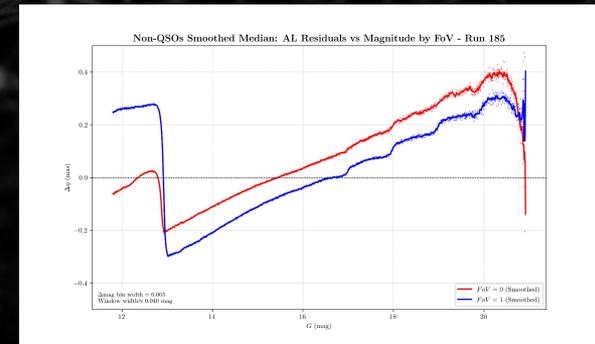
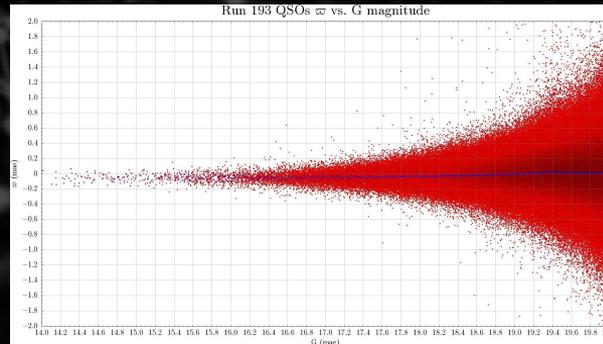
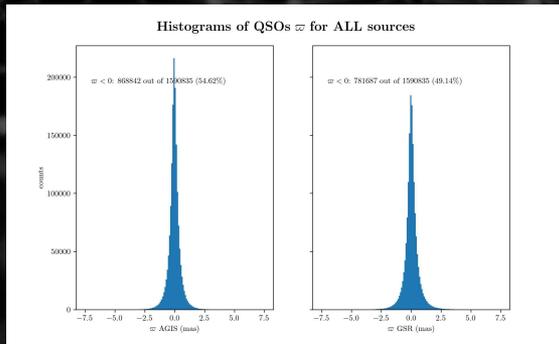
HIGH LEVEL DIAGRAM OF GSR PIPELINE



GSR - Global Sphere Reconstruction

The Global Sphere Reconstruction (GSR) solves a linearized system of equations whose result gives the global astrometric reference system (position, parallax, proper motions).

- Starting from 10^7 to 10^8 objects for each run
- Very complex pipeline structure
- 130.000+ lines of code in Java for data preparation, selection and analysis, and...
- **The Solver running at CINECA**
 - See any of *Valentina Cesare's* presentations on the GPU implementation of the solver
- Final GSR output sent to DPCE in the MDB
- The Solver module run at CINECA which is managed as one processing node of the DPCT
- The whole process could be iterated for Non-Linearity



Transitioning towards a new DMS 1/3

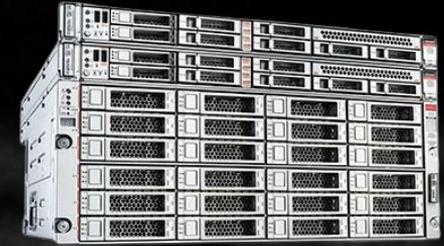
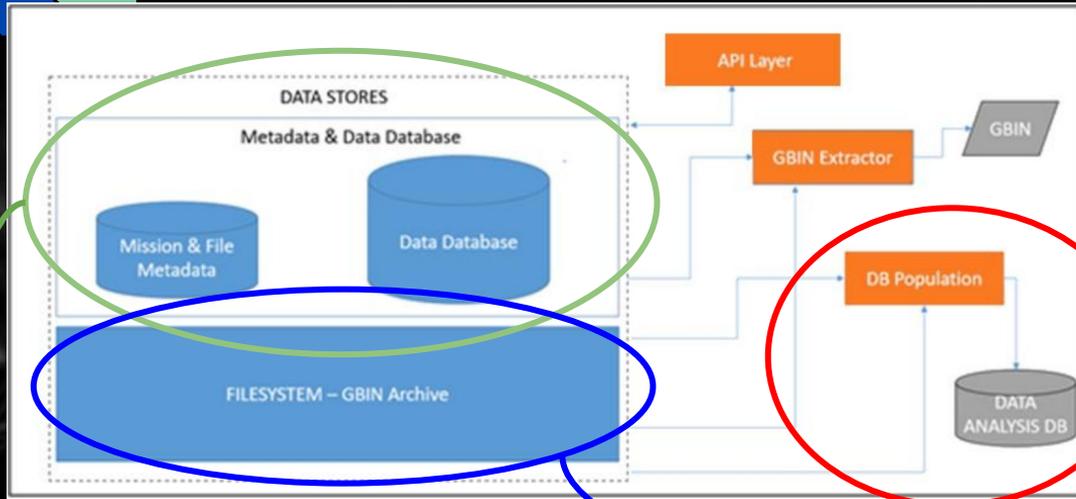
Why

1. Mission extended far beyond the expected timeframe
2. Massive databases/tablespaces (~TB) with detrimental effect on DMBS performance and maintenance
3. Daily / Cyclic operations require different DBMS paradigms
4. Heterogeneous requirements for the data model and its indexing/partitioning:
 - Hybrid DM: Relational with serialized objects (BLOB)

How

1. New Data Management System : Hybrid FS (ZFS) and DB solution
 - data lake with all mission data as GBINS
 - 2 different databases for metadata and mission metadata
 - ODA dedicated to the scientific exploitation of the data
2. New data model:
 - Vertical cut of the original tables
 - removed / exploded every BLOB

Transitioning towards a new DMS 2/3 HW & Logical Design



Oracle Database Appliance
X8-2-HA, TAA



Oracle ZFS Storage
Appliance Racked
System ZS9-2



HPE Nimble HF40

Transitioning towards a new DMS 3/3

Original GAIA DM and Data format

GAIA DM it's an hybrid between relational and document based:

- Every object is represented as a table but some fields are objects as well
- Impossible to query on some of its data/metadata
- Access to the data requires specific sw tools

GAIA Data Format

- Gbin data format is a zipped serialized java object
- very limited interoperability
- very lightweight (structure not included)

New DM and Data Format

OPS4 DM is fully relational

- all fields are query-able with standard SQL
- data can be accessed directly on the DB (vertical cut) or retrieved fully from the archive

Data Format:

- Experimenting with HDF5
- greatly improved interoperability
- increased storage requirements (by a factor of 2)

Spoke 3 - WP4 IDL and IGUC

Interoperability Data Lake for the Gaia Use Case

The purpose of Spoke 3 is the development of innovative applications and software capable of fully **exploiting cutting-edge HPC technologies** and big data storage solutions, to achieve excellence in the areas of astronomy, high-energy astrophysics, astroparticle physics, and cosmology.

WP4

WP4 builds upon **best practices** and already implemented frameworks for managing data and software with FAIR and Open Science principles, to develop **innovative frameworks** capable of addressing the Big Data Challenge

IGUC

The IGUC **innovation grant** originates within the National Center for HPC, Big Data and Quantum Computing and is developed through **a joint collaboration between INAF and Leonardo S.p.A.** to study various technological solutions, such as DMS and the **use of alternative DBMS**, to manage, store, and access big data for the GAIA use case.

The Team in IGUC



Deborah Busonero
Giacomo Coran
Massimo Costantini
Sara Gelsumini
Cristina Knapic
Mario G. Lattanzi
Enrico Licata



Filippo Balla
Gennaro Chiorazzo
Fabrizio Lupi
Daniel Procopio
Sonia Regis



Chiara Francalanci
Paolo Giacomazzi



Carolina Berucci

IGUC Objectives

Technological testing of various DBMS and Data Management Systems, starting from the GAIA use case, with the objective of **estimating and comparing the performance of the different systems**, in a way that is as hardware and scale invariant as possible

OPS4@DPCT

- Oracle DBMS
- ZFS Filesystem
- HDF5 File format
- Oracle ODAx8

Spoke 3 - WP4 IDL

- Postgres DBMS,
- Rucio DMS,
- HDF5 File format,
- INFN Data Lake machines

IGUC - Leonardo/Cherrydata

- AyraDB,
- ext4 Filesystem,
- HDF5 File format,
- INAF infrastructure (bare metal and virtualized)



GAIA Test Case: Dataset 1 - Cone search + Meridian

Given a direction defined by (α, δ) and a radius ϵ , we have that a generic source of coordinates (α', δ') is inside the cone search if:

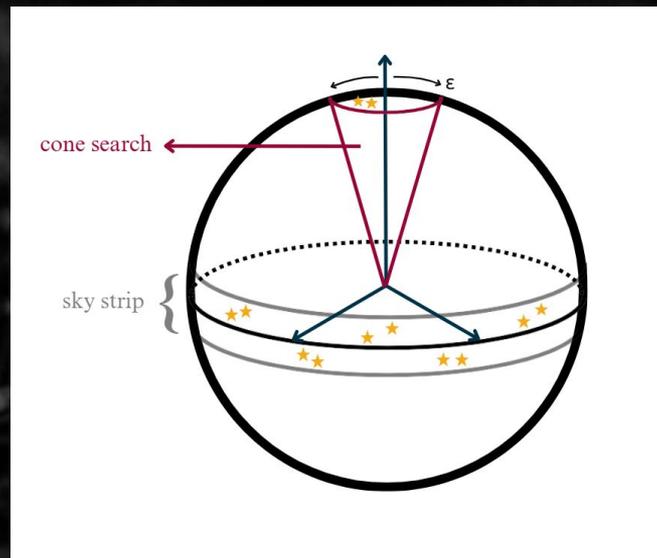
$$\cos(\theta) \geq \cos(\epsilon)$$

where

$$\cos(\theta) = [\cos(\delta) \cos(\delta') \cos(\alpha' - \alpha) + \sin(\delta) \sin(\delta')]$$

While it falls within the plane of semi-width ϵ and perpendicular to the direction (α, δ) if:

$$- \sin(\epsilon) \leq \cos(\theta) \leq \sin(\epsilon)$$



GAIA Test Case: Dataset 1 - Cone search + Meridian

Select all sources and related transits for the specified regions of space and the specified timeframe

Search details:

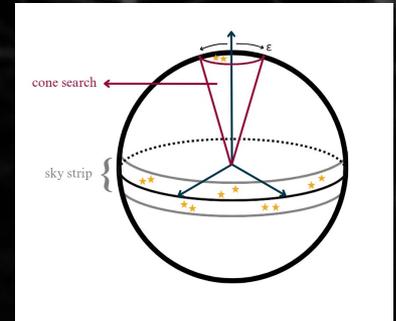
- Cone search direction:
 - $\alpha = 0$ [rad],
 - $\delta = \text{PI}/4$ [rad] = 45 (deg)
- Cone radius & semi-width of meridian band:
 - $\varepsilon = 0,002182$ [rad] = 1/8 (deg)
- Time range: 1 month of data
- Identified $\sim 4.5 \cdot 10^6$ Sources

Cardinality of the dataset:

- CompleteSources (CS) : 148.958.258
- AstroElementaries (AE) : 2.391.111.378
- CrossMatch (XM): 5.467.623.512

Data Volume:

- Gbin Format (serialized compressed java objects):
 - CS ~ 91 GB, AE ~ 1.1 TB, XM ~ 76 GB
- HDF5 Format
 - CS ~ 9.6 TB, AE ~ 34.7 TB, XM ~ 22.6 TB
- HDF5 compressed (7z, LZMA2):
 - CS ~ 425.6 GB, AE ~ 2.3 TB, XM ~ 721.9 GB



GAIA Test Case: Dataset 2 - 20k Cone searches

Select all sources and related transits for the specified regions of space over the entire mission

Search details:

- **Cone search direction:**
 - $\alpha_k, \delta_k = 1$ of 20.000 directions equally spaced along an homogeneous spiral from celestial north pole
- **Cone radius**
 - $\varepsilon = 4,71 \cdot 10^{-4}$ [rad] $\sim 96,7$ arcsec
- **TransitID range:**
 - all available mission data

Expected search results:

- Identified $\sim 2.0 \cdot 10^6$ Sources

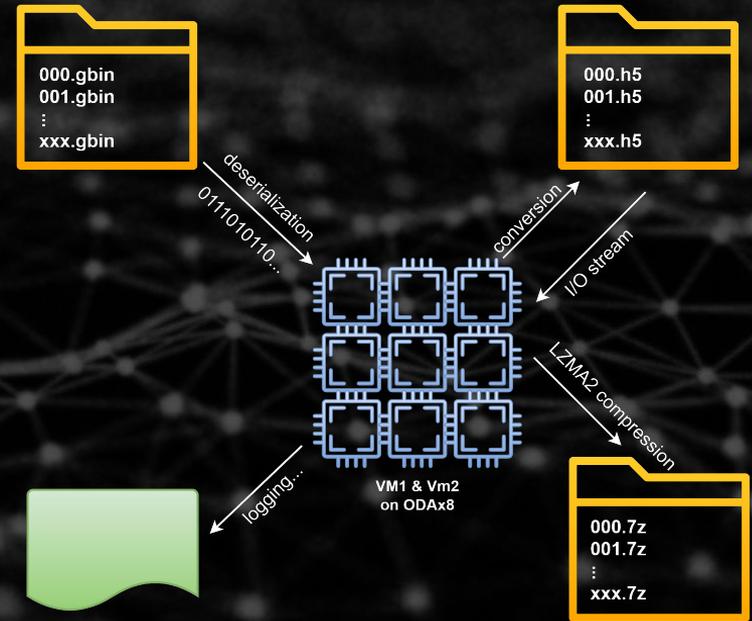
The extraction of this dataset required the development of a specific sw solution based on Apache NiFi.

Without it the number of gbins selected by this dataset and their total volume, **covered almost the entire GAIA datalake at DPCT**

HDF5 Converter

The software heavily leverages reflection and recursion to explore each input objects, regardless of its type, structure or complexity, and creates a corresponding HDF5 file mapping: Java objects, arrays and primitives are mapped respectively to HDF5 Groups, Datasets, and Attributes.

- Gbin files are read from a folder and all its subfolders
- Using **MDBDM (Gaia Datamodel)** and **HDFQL** (library) an .h5 version of each gbin is created and stored in a temp folder
- Each h5 is then read through a bufferedStream, compressed using **LZMA2 compression algorithm** and stored into an output folder
- The hierarchical structure of the original gbin is conserved into the final HDF5

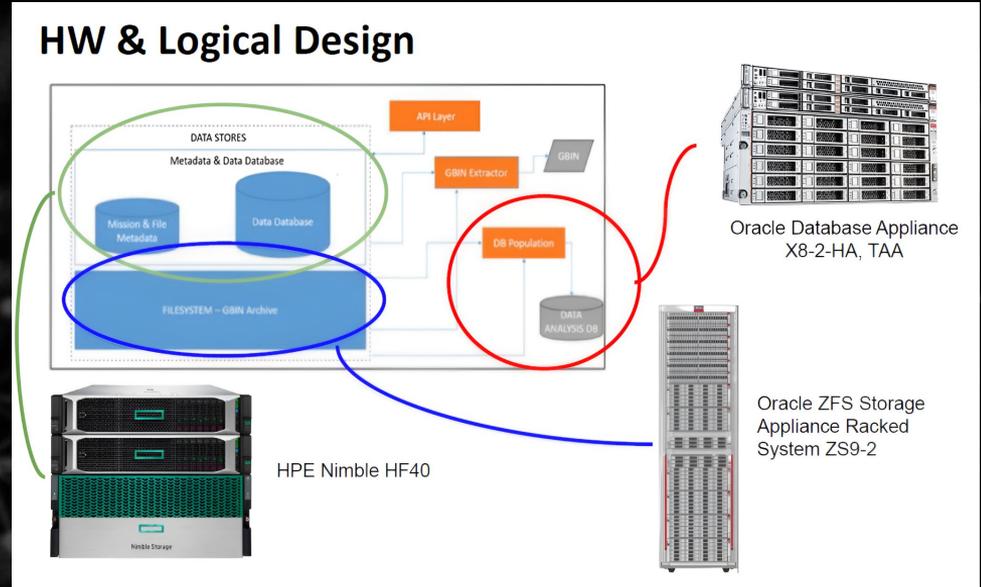


OPS4 @ DPCT - Oracle + ZFS

The legacy project transitions from operational real-time data processing (OLTP) to a hybrid data management approach (OLAP + FS)

New DMS based on the interaction between Data Lake (GBIN format) managed by ZFS and two different DBs hosting metadata and mission metadata implementing a FAIR paradigm.

The submission of query/Data request is managed via an API layer, and outputs are delivered as files (GBINS / FITS / HDF5) or are used to populate a specific Data Analysis DB hosted on the ODA



Spoke 3 WP4 IDL

- DBMS: **postgreSQL**
 - Datamodel an **indexing based on the metadata** (ATTRIBUTES) of the provided HDF5 Data-Lake
- DMS: Rucio instance hosting the Data-Lake
- **Cut and Merge custom SW:**
 - Able to retrieve specific parts of the original HDF5 file and provide only the required GROUPS / DATASETS / ATTRIBUTES to the user
- **Web interface** to allow the submission of queries from the users

Jportal Retrieval System

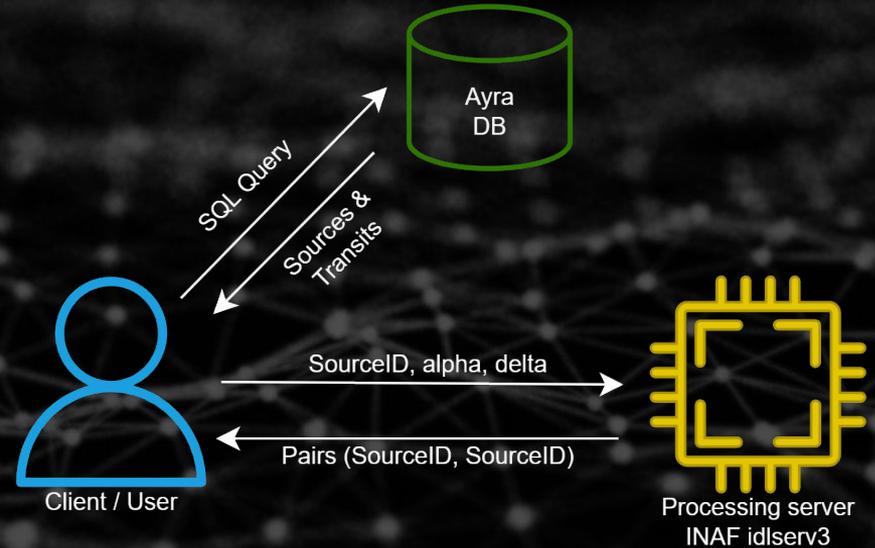
- Query on objectName, sourceId, or other metadata
- **Results pair metadata** retrieved from PostgreSQL **to the corresponding files** from Rucio Data Lake
- Data downloadable thanks to **MinIO**

Tap Endpoints following IVOA standards

- enables the execution of queries from other clients such as TOPCAT

IGUC - Leonardo & Cherrydata

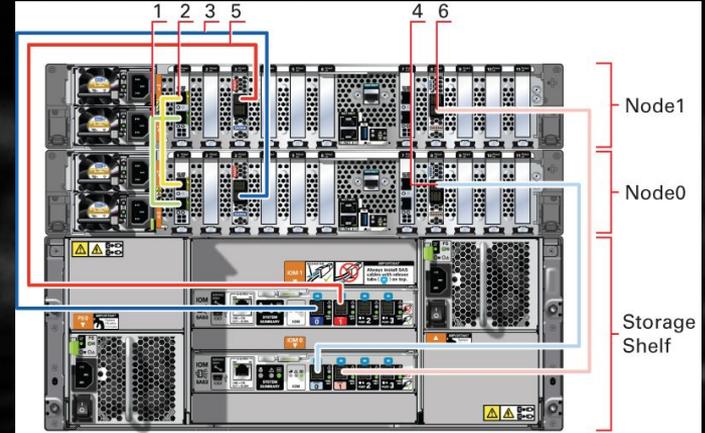
- DBMS: **AyraDB**
 - database developed by Cherrydata,
 - Key-Value core
 - SQL operations
- **Custom SW for data extraction** from the HDF5 files
- Data-Lake: **HDF5 files**
- **Deployment** of a test AyraDB cluster **on INAF** infrastructure@OATs



Results on Oracle DBMS + Oracle Spatial + ZFS

ODAx8 integrated system:

- 2 processing servers with a total of 64 Intel® Xeon® Gold 5218, and 384 GB of memory per server.
- Storage shelf with 6 x 7.68TB SSDs and 18 x 14TB HDDs
- connected via dual 12Gb/s SAS interface
- DBMS Oracle 19c + Oracle Spatial and Oracle BI
- Storage ZFS



Data are ingested directly into the DB tables using TLS Integrator: a parallel software developed inside the project, that reads from gbin and populates the database. Sample ingestion times (subset):

1 Reader Thread
10 Writer Threads
RAM used: ~12GB

CS entries: 3.105.011, time: 101s, ~32 μ s per entry, 30 krecord/s
AE entries: 128.407.793, time: 1696s, ~13 μ s per entry, 75 krecord/s
XM entries: 129.062.767, time: 206s, ~1.6 μ s per entry, 626 krecord/s

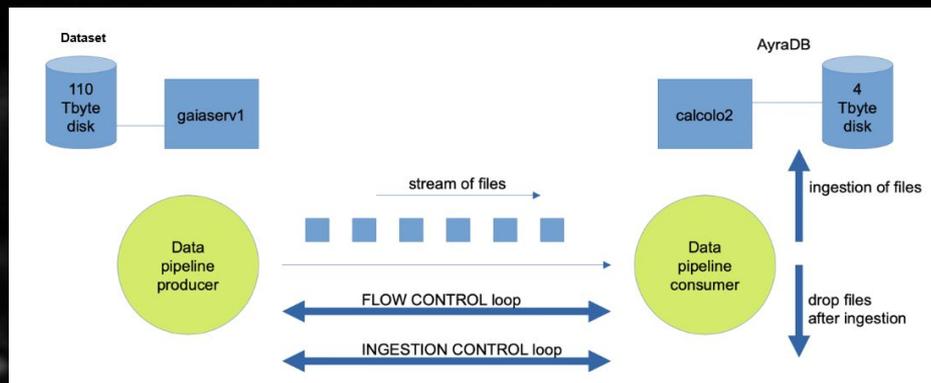
Results on AyraDB + KNN + ext4 fs

gaiaserv1:

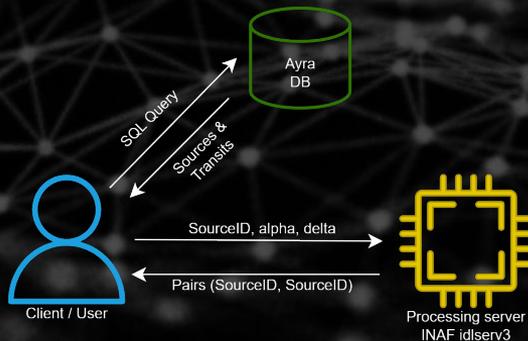
CPU: 16
RAM: 64 GB
Disk: 110 TB, virtual disk

calcolo2:

CPU: 128 Intel Xeon gold 6548Y+
RAM: 206 GB
Disk: 4 TB, local SSD



The dataset is on **ingested on calcolo2**, streaming from **gaiaserv1**.
Calcolo2 then manages the DBMS and client queries.
KNN computation if offloaded to **INAF idlserv3**



Dataset	N. files	Gross size	N. records	Ingestion time	Throughput
AstroElementary	199648	5.4 Tbyte	2,434,219,917	16 hours	42 krecord/s
CompleteSource	15949	432 Gbyte	159,440,000	2 hours	22 krecord/s
CrossMatch	60710	1.1Tbyte	6,068,000,000	6 hours	280 krecord/s

Results on PostgreSQL + RUCIO

- DBMS: **postgreSQL**
 - Datamodel an **indexing based on the metadata** (ATTRIBUTES) of the provided HDF5 Data-Lake
 - “standard” environment and “sharding” environment based on HEALPIX lvl6 (64 shards) to speed up source identification and healpix lvl12 to build spatial index in each shard
- DMS: Rucio instance hosting the Data-Lake
- **Cut and Merge custom SW:**
 - Able to retrieve specific parts of the original HDF5 file and provide only the required GROUPS / DATASETS / ATTRIBUTES to the user

Jportal Retrieval System

- **Web interface** to allow the submission of queries from the users
- Query on objectName, sourceId, or other metadata
- **Results pair metadata** retrieved from PostgreSQL **to the corresponding files** from Rucio Data Lake
- Data downloadable thanks to **MinIO**

Tap Endpoints following IVOA standards

- enables the execution of queries from other clients such as TOPCAT

IGUC results Summary

- Testing on the **Oracle solution** showed **great results (constant time) when leveraging the Spatial index**, but we found a clear degradation of the performance when the change in geometry lead to the full table scans
- Testing of the **AyraDB solution**, showed **good results on cone searches** and **surprising results on meridian searches**. However meridian searches showed a clear dependency on the dataset size, and which suggest **possible scalability issues**
- For the GAIA test case, both solutions proved to be appropriate
- Having completed these assessments in the scope of **IGUC will surely allow us to make informed choices** about the technologies we will use in future projects
- **The goals of the IGUC project can be considered achieved**

Lesson Learned & Friendly advices

- **Plan for the future:**
 - space missions can outlive the planned lifetime and you and your infrastructure need to be ready for it
- **Do not skimp on the test environments!**
 - Make sure to scale them properly and plan procedures (automated if possible) to keep the envs consistent
- **A broad amount of different expertise is required for the success of such projects:**
 - leverage the expertise and tools of your colleagues (MIRTA :D)
- **Share information and foster communication:**
 - this keeps people involved and motivated, since these projects are a **loooong** commitment
- **Be ready for the turnover:**
 - Having good documentation and even better procedures to bring the new members up to speed is paramount!

The 20 years GAIA SGS experience, similarly to other mission like Euclid, Planck etc formed a community of experts, that is now an asset for INAF in perspective of future large missions and projects

Lesson Learned... continued

- **HDF5: great for accessibility**, maybe not optimal for the Gaia dataset
 - Rich hierarchical format
 - great potential for chunking and lookup
 - a lot more volume overhead than expected (probably due to the data model)
- **Do not underestimate the importance of the correct data format for your project**
 - Starting from gbins and the development and optimization of the HDF5 Converter resulted a task way more complex and time consuming than expected
- **Always have contingency plans for unexpected setbacks**
 - The intermittent availability of the DPCT, especially during the first part of the project, due to GAIA Mission constraints, and infrastructure maintenance, posed a few challenges
- We would like to **continue experimenting** on this ideas, testing other technologies and data formats like **Parquet** and and DBMS like **DuckDB**
- Use the **knowledge acquired** through IGUC, to improve our infrastructure **to support the italian scientific community**

Bibliography

partial... there are too many contributions to cite them all

1. **Italian Gaia data management and processings infrastructure for the AVU systems at DPCT**, Enrico Licata, Deborah Busonero, Sara Gelsumini - USCVIII General Assembly 14-18 Ott 24
2. **Gaia use case in Spoke 3**, Sara Gelsumini, Enrico Licata, Deborah Busonero - Spoke 3 Final Workshop, Sexten, 14-19 Dec 2025
3. **Spoke 3 Archive Infrastructure** - Giacomo Coran, Massimo Costantini - INAF OATs - Spoke 3 Final Workshop, Sexten, 14-19 Dec 2025
4. **INAF - Relazione Attività delivery 02** - Filippo Balla - 29 Luglio 2025
5. **INAF - Relazione Attività delivery 03** - Filippo Balla, Valentina Nalesio - 7 Novembre 2025
6. **CherryData - Relazione finale Attività Gaia** - Chiara Francalanci, Paolo Giacomazzi - Settembre 2025
7. **TN-IGUC-INAF-EL-001, TN-IGUC-INAF-EL-002, TN-IGUC-INAF-EL-003 TN-IGUC-INAF-EL-004**
Enrico Licata, Deborah Busonero, Mario Lattanzi - Settembre 2024 - Settembre 2025