Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

OPS4 Gaia Legacy project: Hybrid data management, dynamic data modeling and dedicated data reduction sw for Gaia-like big data exploitation

Deborah Busonero, Raffaella Buzzi, Sara Gelsumini, Mario G. Lattanzi, Enrico Licata, Roberto Morbidelli, Alberto Vecchiato

Spoke 3 Technical Workshop, Trieste October 9 / 11, 2023

Scientific Rationale – The goal beyond the ICSC

- Generation of a deep and complete sky, on 4π sterad, as a reference tool and therefore interoperable for the integration of multiband data (from radio to high energies) and multimessenger data (e.g. sources of gravitational waves, neutrinos,...) for efficient data mining aimed at fast multidimensional scientific data exploitation;
- Capacity for ad hoc recalibrations of astrometric and photometric data for the reclassification and redetermination of the fundamental properties (motions and magnitudes) of classes of objects of particular astrophysical interest;
- Operations of telescopes from Earth and space and support for the study of new missions/projects;
- Interoperability and integration of metadata from non-astronomical databases, i.e. engineering and orbital data, data from service modules or payloads, or data coming, e.g., from Space Weather and/or surveillance of space debris (space debris surveillance).

“From Research to Industry

From Gaia DPCT to INAF OPS4 Legacy project – A Big Data/ HTC Facility for future exploitation of the mission data

DPCT Overall HW infrastructure

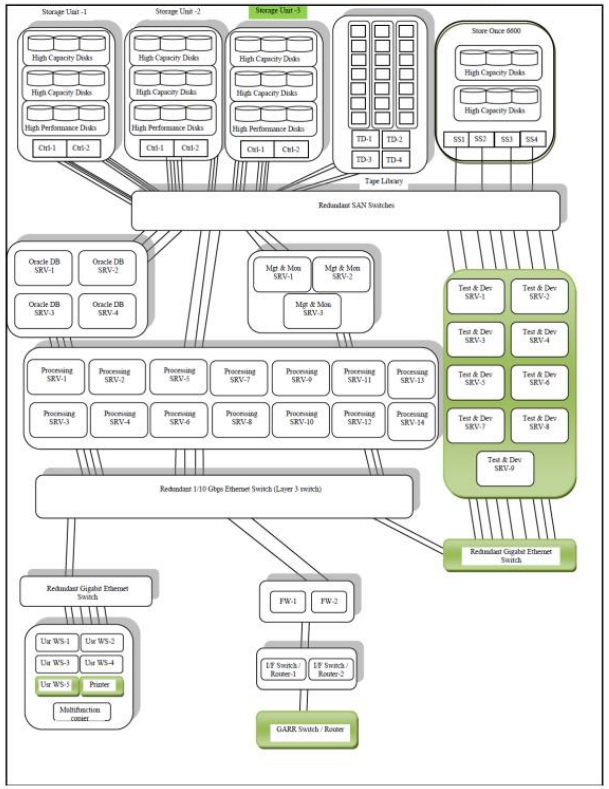

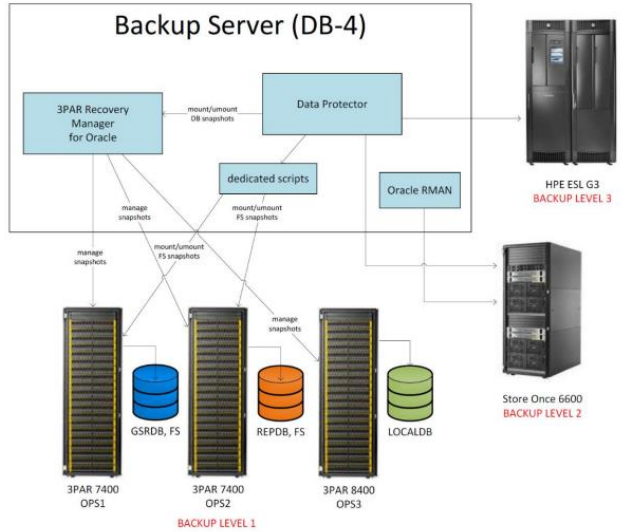

Credits: ALTEC

14 servers with 600 CPU cores and 5 TB of RAM
Several PB of storage capacity
Oracle RAC technology with DBMS Oracle
3 levels of backup one on primary storage, one on disks, one on tape library

Direct connection with CINECA HPC systems

- Accomplished Work, Results

NOW: OPS4

HW & Logical Design

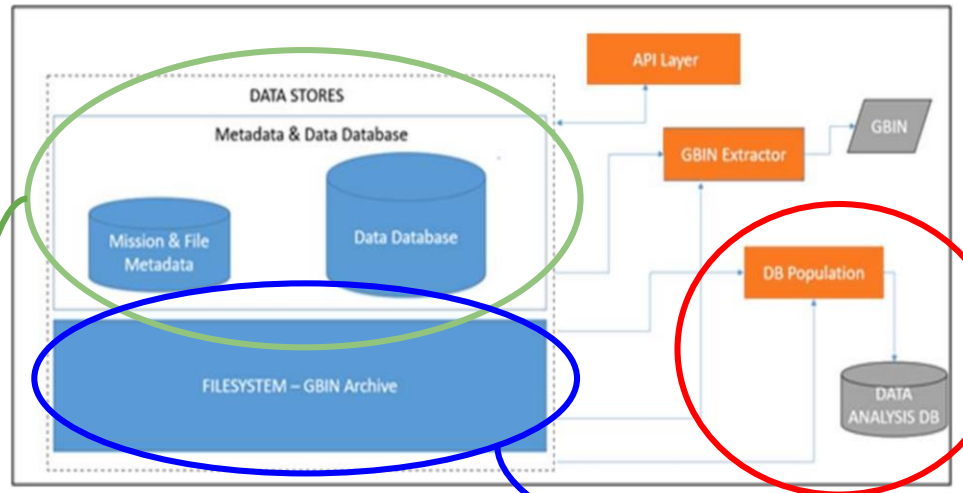

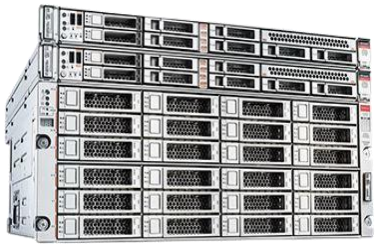

Oracle Database Appliance X8-2-HA, TAA

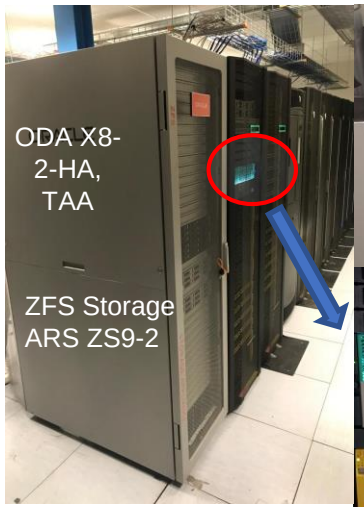

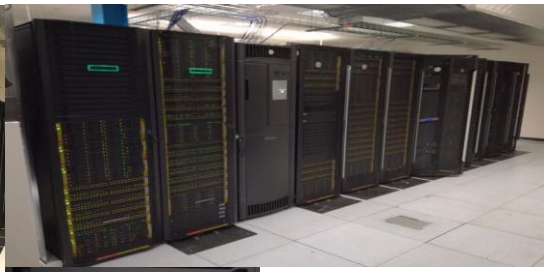

HPE Nimble HF40

Oracle ZFS Storage Appliance Racked System ZS9-2

THE FUTURE: OPS4 legacy

- Technical Objectives, Methodologies and Solutions

Creation of a HW and SW infrastructure for the management, manipulation and analysis of large amounts of data relating to the whole sky starting from the Gaia mission data, satisfying the requirements of interoperability and data preservation

Moving from an approach where data is stored and processed in real-time to an online analytical processing which uses complex queries to analyze aggregated historical data from the OLTP system.
We need fast queries and analysis of data from different perspectives:

- Run queries at billions of rows (sources) per second for each CPU core
- Switching between a Source oriented search by row (space), to a columnar search, by Transit (time) leveraging both indexing methods without the need to duplicate the DB volume;

❑ We need also to pre-aggregating and pre-calculating the information in the database before delivering to the users

Solutions

❑ Hybrid data management

A paradigm shift is required to move from operations to verification and then exploitation (legacy)

- ***Energy efficiency:*** cold data / cold place → Long term data storage on file system
  - not tape, to fulfill scientific data requests, responsive access is still needed
- ***Fast response times:*** hot data / hot place → 2 DBs
  - 1 Data Database: with a subset of data (vertical cut) to respond quickly to most frequent requests
  - 1 Metadata Database: to store mission metadata and to relate requested data to their respective location on the filesystem: larger and/or less frequent requests
- **Scalability:**
  - easier to expand ZFS storage, keeping high level of performances leveraging Metadata DB
  - DataAnalysis DBs are hosted on engineered machine, pluggable, lift and shift…

- Actual Solutions

❑ Oracle technologies for Gaia-like explotation

In-Memory

- Oracle Database In-Memory implements state-of-the-art algorithms for in-memory scans, joins, and aggregation. These optimizations, enable to run queries at billions of rows per second for each CPU core so becomes the perfect application for switching between a Source oriented search by row (space), to a columnar search, by Transit (time) leveraging both indexing methods without the need to duplicate the DB volume

Spatial

- leverage GIS to create a spatial index (the full GAIA sky): defined using orthodromic distance
- "Target" object is quickly placed in relation to the "candidate" objects present in the catalog, allowing only those satisfying the requirement to be identified among all distances.

Relational Duality

Using Duality Views, data is still stored in relational tables in a highly efficient normalized format but is accessed by apps in the form of JSON documents. Developers can thus think in terms of JSON documents for data access while using the highly efficient relational model for data storage, without having to compromise simplicity or efficiency.

- Possible solutions

❑ Dynamic Data Modeling

The GAIA operations/pipeline DM is not suitable for technical/scientific exploitation. The implementation of TLS Integrator Prototype was effective to complete the PoCs but reusability and maintainability can be an issue.

This experience highlights the importance of a DM as dynamic as possible.

Investigation on the JSON fields and collections inside a Oracle DBMS for populating the DataAnalysis DBs

- JSON fields are inherently schema-less but can be queried, partitioned and indexed as primitive fields (also supporting Oracle Spatial)
- “document like” objects can be created from the relational database, and using SQL syntax , used to create Tables on the DataAnalysis DBs tailored for a specific use case
- This operations could be scripted to automate the construction of the DataAnalysis DBs

Starting from Oracle 23c ***Relational Duality*** will allow the creation of JSON views based on relational tables, allowing the process layer to manage “Object like” data structures, which will be automatically mirrored by relational tables on the DB side

- Timescale, Milestones and KPIs

April 2024 | November 2024 | April 2025 | September 2025

- GW use case on different platform (investigation within the WP4 efforts)

DM convertion from Gaia Data Model to the suitable one

OPS4 system online to support the Gaia operations

SPIE 2024 on June

OPS4 system online to support data exploitation for the Verification

ADASS on November 2024

OPS4 system available to users, with its fundamental set of features

- Support on-site data exploitation and analysis
- Depending on the new HW acquisition

Accomplished Work, Results

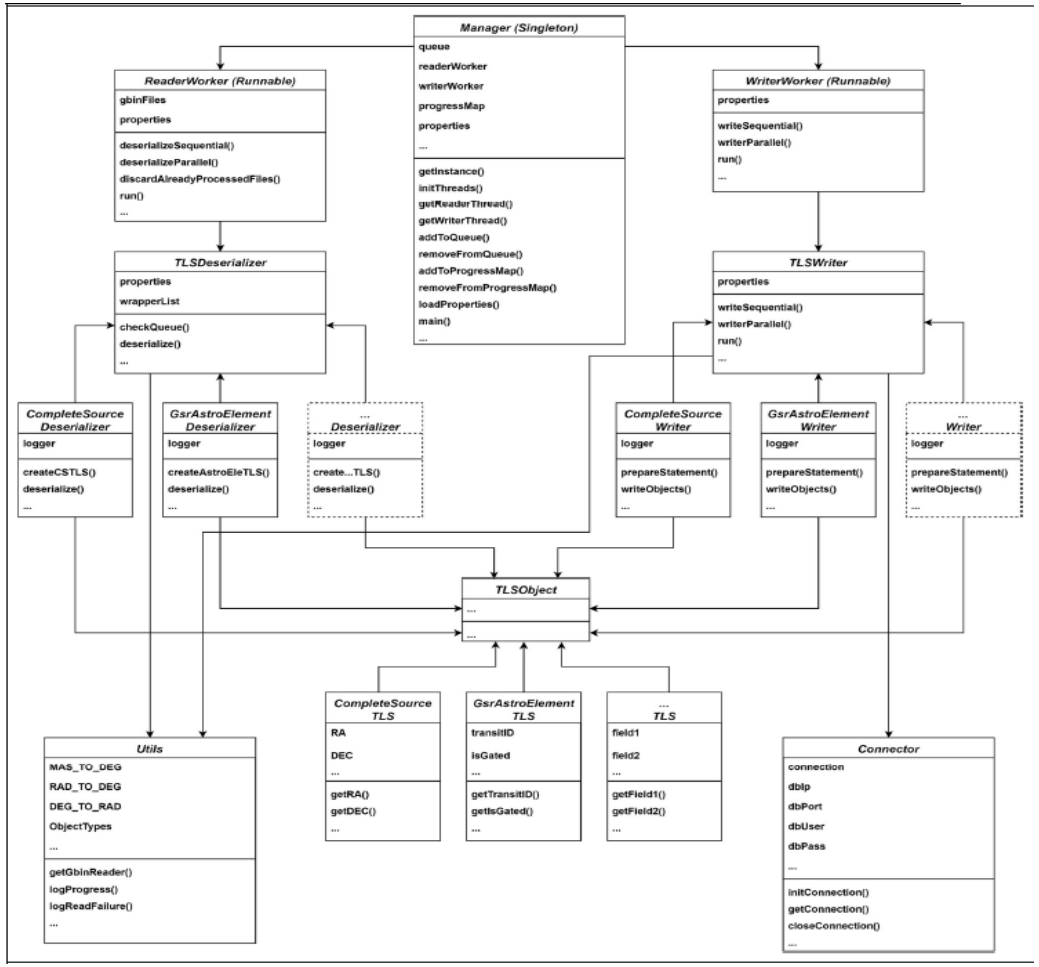

High level Software design of the TLS Prototype Integrator used to populate the DB

Experience on TLS Prototype SW Infrastructure

Manager: entry point , coordinator, ***Singleton*** design

Read/WriteWorker: ***Parallel implementation,*** queue

TLSDeserializer: Abstract, I/O optimized

TLSWriter: Abstract, **Java Reflection**

TLSObject: Supertype, **ObjectModel**

Connector: ***ojdbc7,*** manage CPU by limiting connections/processes

- Accomplished Work, Results

Prototype based on an Oracle DBMS able to provide an impact assessment, in terms of hardware resources and software tools, of specific scientific test cases, in particular, to explore the possibility of carrying out operations currently not foreseen or not feasible on the "Mission Database" of the Gaia mission, or more generally, with the tools available in the astronomical community, starting from (but not limited to) the VO

REQUIREMENTS

1. **PoC0:** indexing and cone search on the entire GAIA sky (DR2)
   - Cross Match the Complete Sources
   - DM for scientific exploitation (legacy)

DONE

2. **Poc1:** GSR Residual computation
   - provide a R&D environment, to perform scientific analysis of results of GSR pipeline
3. **PoC2:** Gravitational Waves: analyze portions of the sky (tenth of square degrees) to identify significant variations of some sources, to verify the effective presence of gravitational waves.

- Accomplished Work, Results

HIGHLIGHTS

1. Indexing via Oracle Spatial & graph based on the orthodromic formula

$$\gamma = arccos(cos(ra_2 - ra_1) * cos(dec_2) * cos(dec_1) + sin(dec_2) * sin(dec_1)$$

with ra_1, ra_2, dec_1, dec_2 the right ascension and declination of bodies 1 and 2 and γ their angular distance

2. TLS Integrator to transform the DM from the Hybrid relational (operations) to a purely relational to allow direct access to scientific data

3. Develop a data Visualization layer to query the data and analyze it by Source and by Time

N.B.: The amount of data required to complete PoC2 was beyond the capabilities of the prototype HW

Poc0 complete reliability was verified with respect to the corresponding VO methods demonstrating quantitatively that it can operate without some "quantitative" limits (**VO crash for CS with over 2 million objects extracted**) and without time-loss with respect to the corresponding method of the VO (the execution of the Cone Search was often even faster).

- Accomplished Work, Results

3 main use case categories

by position, by time and by a combination of the two.

Caveat: accessing data by position or by time requires the creation of substantially different indexing strategies and DM design.

By Position

- Indexing on healpixes proved ineffective
- Leveraging the capabilities of Oracle Spatial proved to be extremely more effective
- Partitioning strategies, based on healpixes instead could prove to be a useful strategy

By Time

- Usually full table scans indexing on transitIDs
- Leverage a time based partitioning strategy
- Suitable for short time frames (days / months) or full mission limited by XM
- More suited t o perform calibrations and analysis about the instrument's response

The use case of the detection of gravitational waves using GAIA Astrometry is the most demanding since fall in the third category, requiring to access data both by time and position. Highly correlated on huge angular and temporal scales.

Next Steps and Expected Results (by next checkpoint: April 2024)

- GW use case on different platform (investigation within the WP4 efforts)

DM convertion from Gaia Data Model to the suitable one

But

The Data are covered by an NDA so we need to think about it

- In parallel we continue the implementation of the hybrid data menagement described above with the goal to put the OPS4 system online, to support the Gaia operations and data exploitation for the Verification.