



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani

PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC,  
Big Data and Quantum Computing

# *Data Lake Spoke3 and IDL*

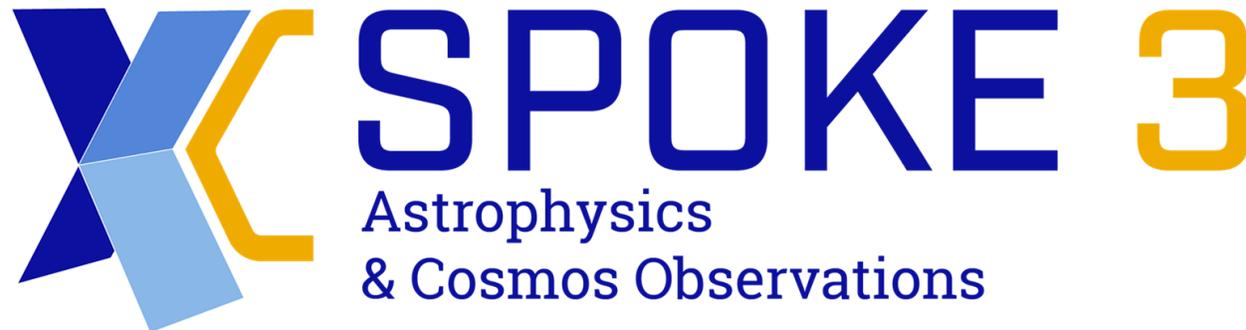
*G. Coran, M. Costantini*

*INAF - OATs*

Archives and Data Management Systems, Bologna Feb 26-27, 2025

## Spoke 3 project

- Spoke 3 is a project supported by European funds for the Fondazione ICSC for astrophysics and cosmological observations;
- The main objectives of the Spoke 3 are the exploitation of cutting-edge solutions in HPC and Big Data processing and analysis for problems of interest in various astrophysical areas;
- It also undertakes a specific outreach and training program to boost the use of high performance and high throughput, cloud solutions in research centres, academia and private companies;
- Within this project we participate in the Work Package 4, which starts from already implemented best practices and frameworks to manage data and software with FAIR and Open Science principles to develop innovative frameworks able to satisfy the Big Data Challenge.



## IDL project

- **The IDL project is one of the projects funded by the ICSC Foundation with industrial leadership, co-participated by Spoke 2 and Spoke 3;**
- **The project involves the construction of an Interoperable Data Lake, hence the name, that is, an archiving system capable of satisfying both industrial needs and the interoperability paradigms supported by the scientific field;**
- **At the moment the project is divided into two sections, an industrial-led section that deals with archiving data on space debris provided by TASI and our scientific section led by INAF for which we are testing the archiving also in this case of both observational and simulation data.**

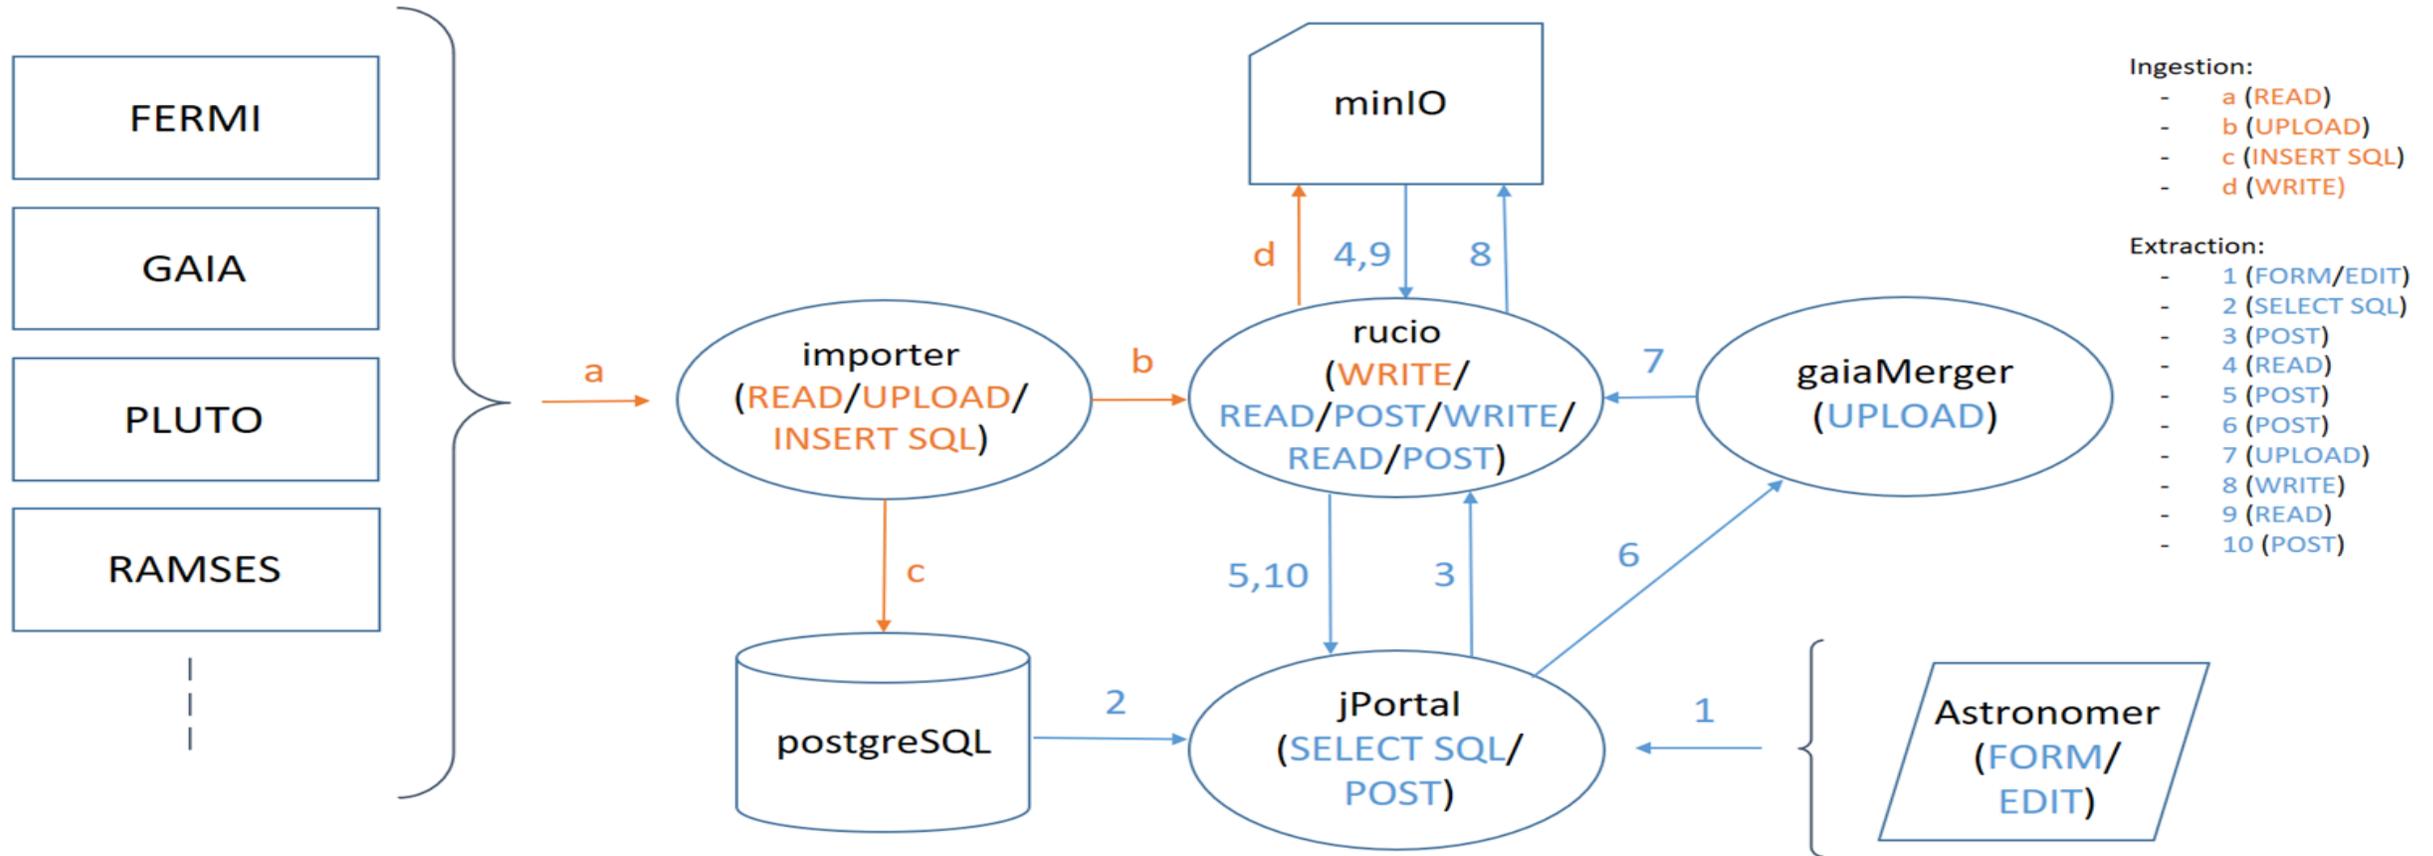
## Spoke 3 Data Lake

- **For the Spoke 3 project we are focusing our work on the development of a shared archive able to support the products of the other WPs' work and to provide an effective and rapid archiving of the data collected in relational databases according to the international standards dictated by IVOA and recovery services for such data both according to common channels such as IVOA clients, with the association of a clear TAP service to the archive, and with a customized portal specifically for this project;**
- **For the duration of the project, we have focused on a couple of observational use cases, which, we will show you shortly, are already well underway, and we are trying to repeat ourselves for another couple of simulation cases;**
- **Spoke 3 WP4 works integrated with a dedicated Rucio instance within an infrastructure provided by INFN.**

## IDL Data Lake

- **For the IDL project we are expanding the ingestion solutions to be able to handle numerous different data files formats and metadata structures and to be integrated with a custom Rucio implementation;**
- **As scientific data we are using a dataset of Birales data provided by colleagues Germano Bianchi and Pierluigi Di Lizia, data from the Fermi satellite and a dataset of public pulsar data present in the IRA web archive;**
- **The data lake is provided as infrastructure by machines managed by INFN, supports a decentralization based on the blockchain system, also managed by INFN, and relies on a private DBMS, AyraDB, developed by CherryData, a start-up from Milan.**

# Ingestion and Retrieval in the Spoke 3 Archive Infrastructure



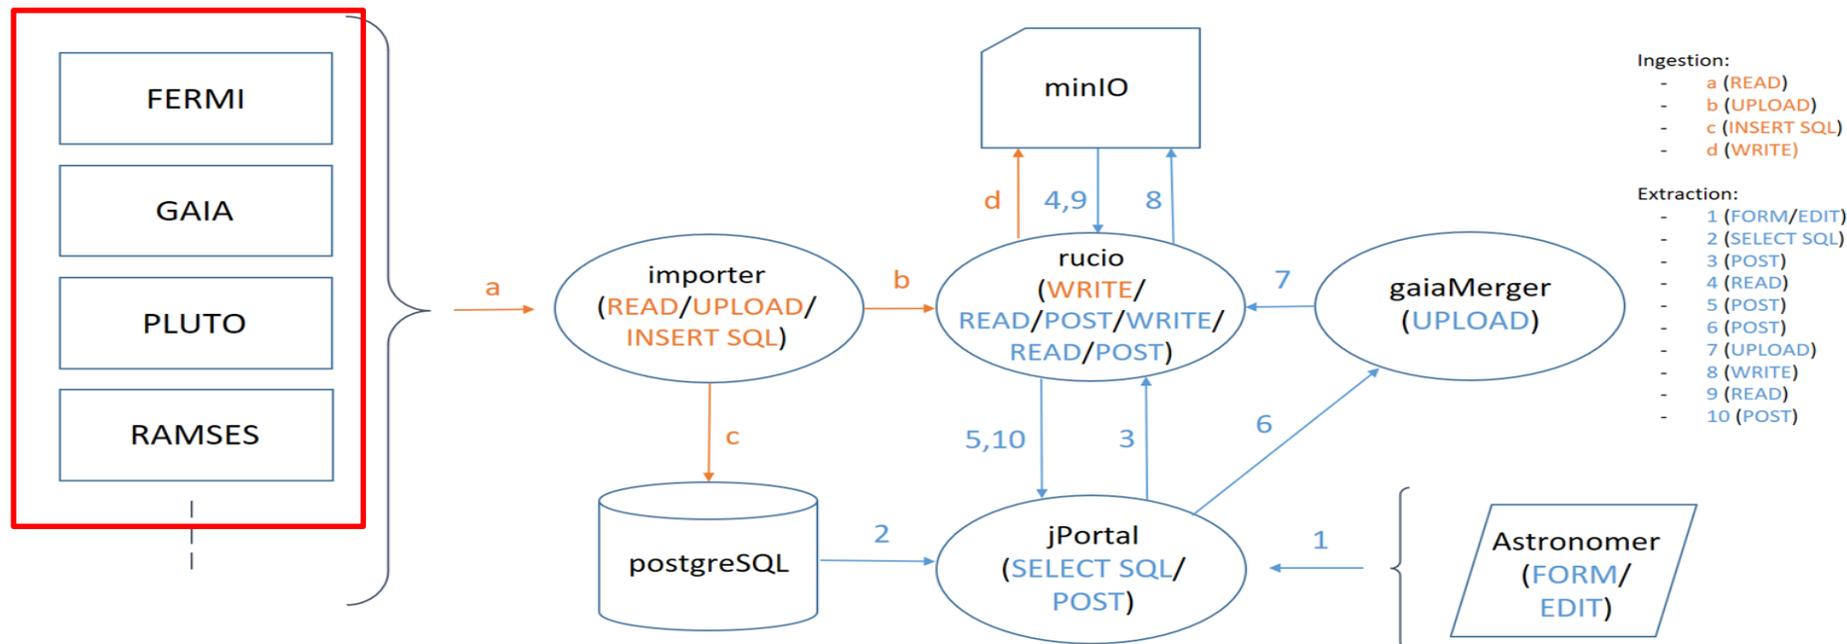
# Ingestion Basis

- The ingestion system is based on NADIR, the solid IA2 infrastructure that allows the ingestion and transfer of data to the archive in Trieste in a fast and efficient way; the system is used for various Italian data catalogues and is the basis of the archiving system of structures such as TNG, LBT and the Asiago Observatory;
- An updated version of the NADIR software ported in Python has been developed in the last year and a half and is already active in Asiago and is being initialized for LBT.



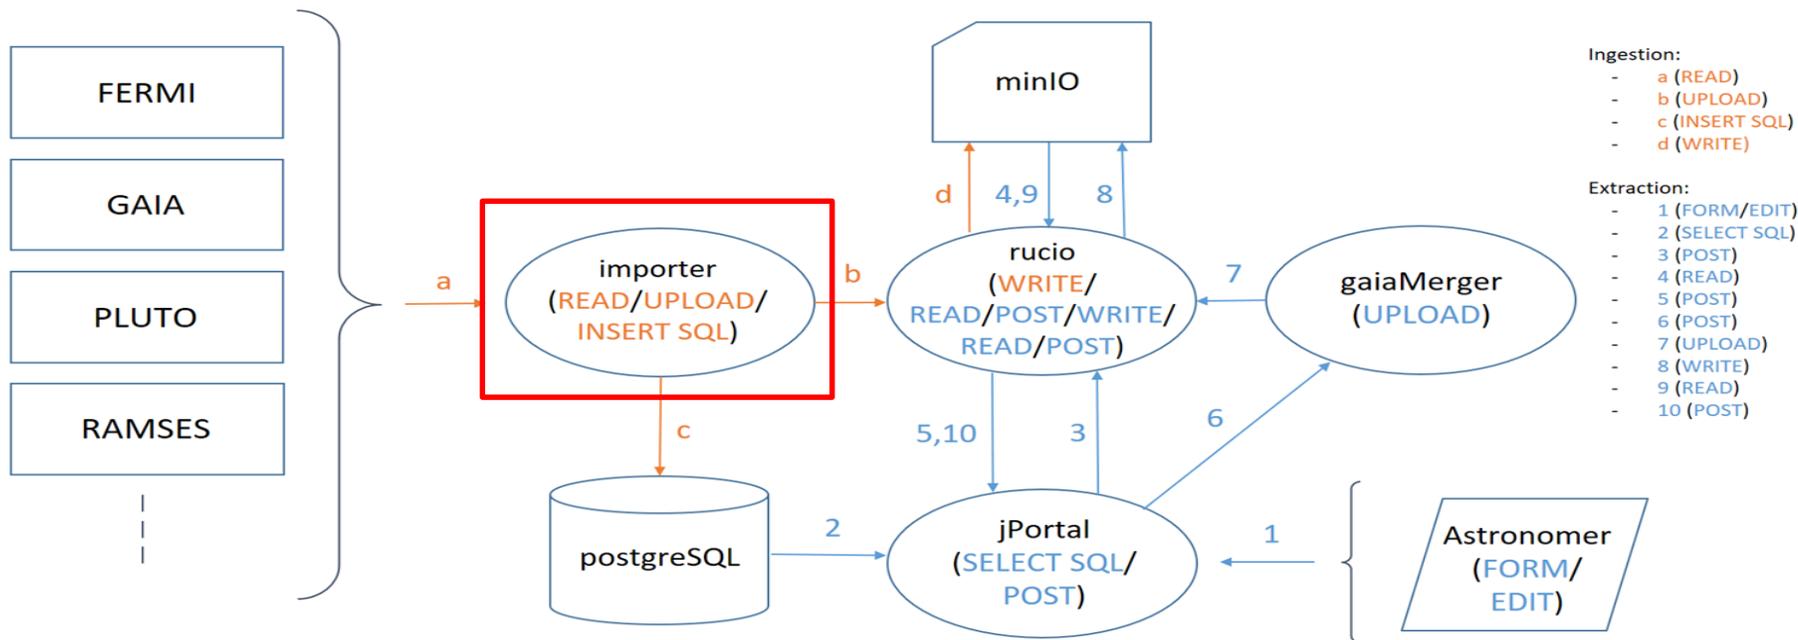
# Data and Metadata Ingestion

- For the Spoke 3 project we are testing the infrastructure with data from the Fermi and Gaia space telescopes as observational use cases, for which the tests carried out are already meeting the targets initially defined in WP4, while we are managing data from Pluto and Ramses as simulation use cases.



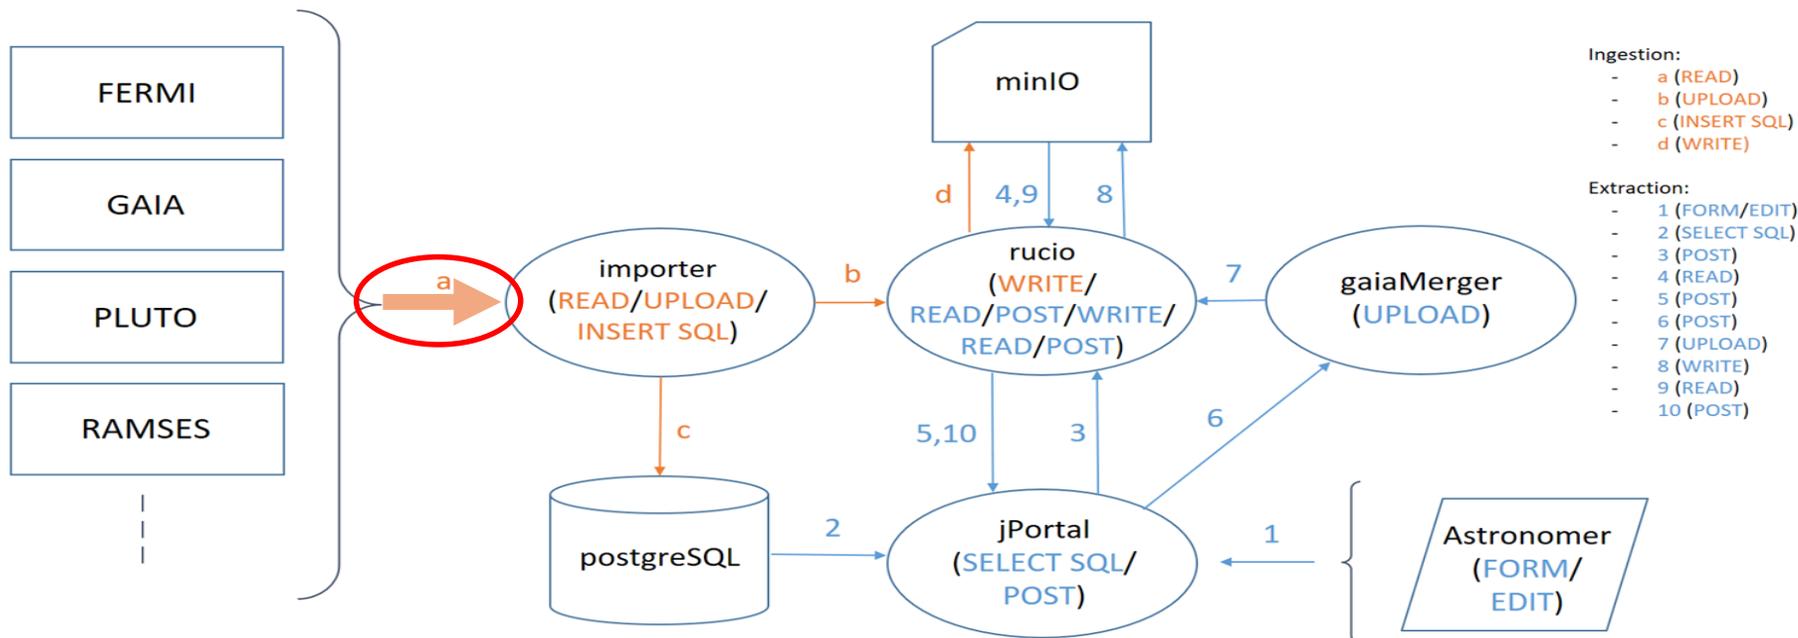
# Data and Metadata Ingestion

- The ingestion process begins with the extraction of metadata by the "Importer", a software implemented as a device in the TANGO Controls distributed control system, with which it is also possible to monitor it.



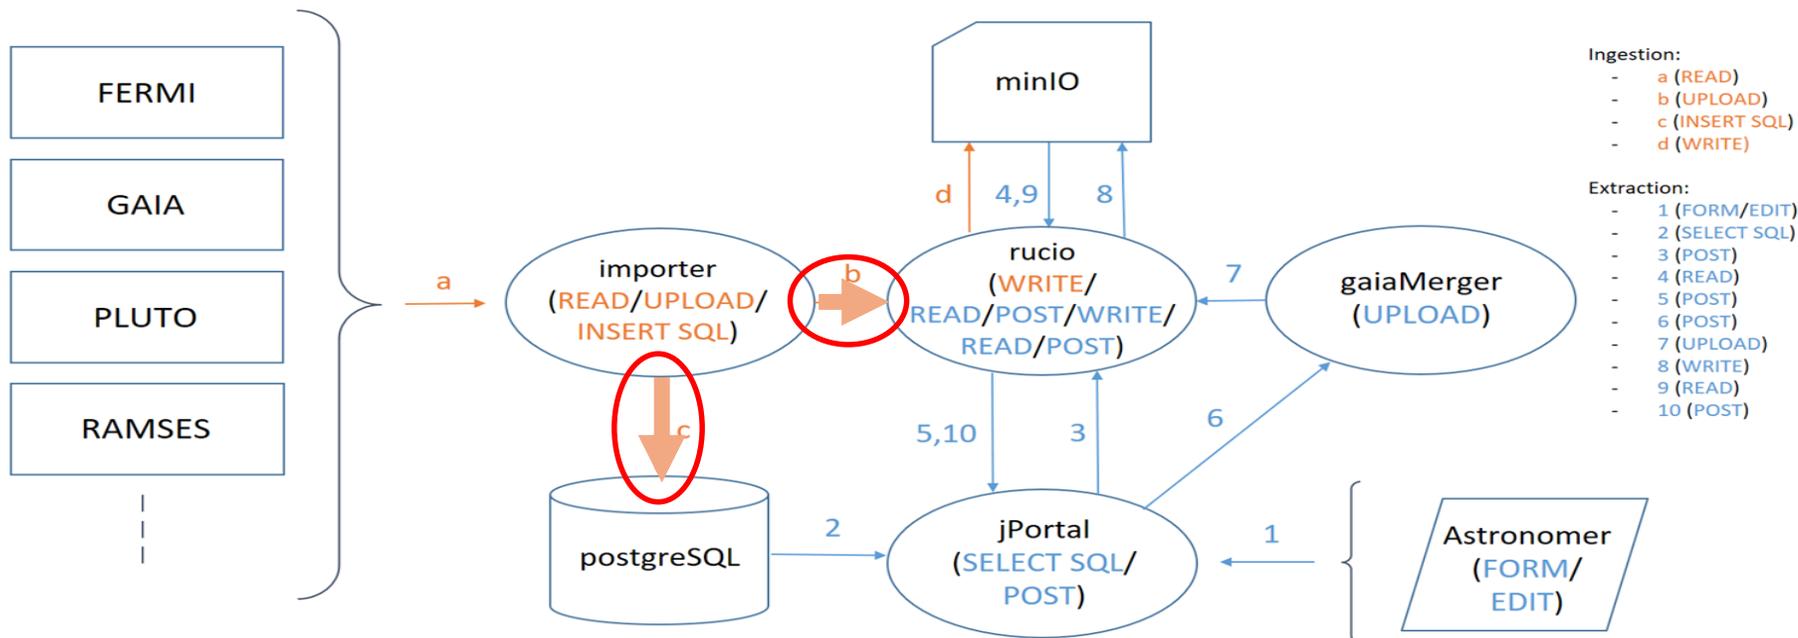
# Data and Metadata Ingestion

- The initialized "Importer" instance retrieves the data descriptive model from a local database, in our case MariaDB, and through the use of mapping tables retrieves the key-value pairs with extraction functions developed specifically for the various types of formats used by our use cases.



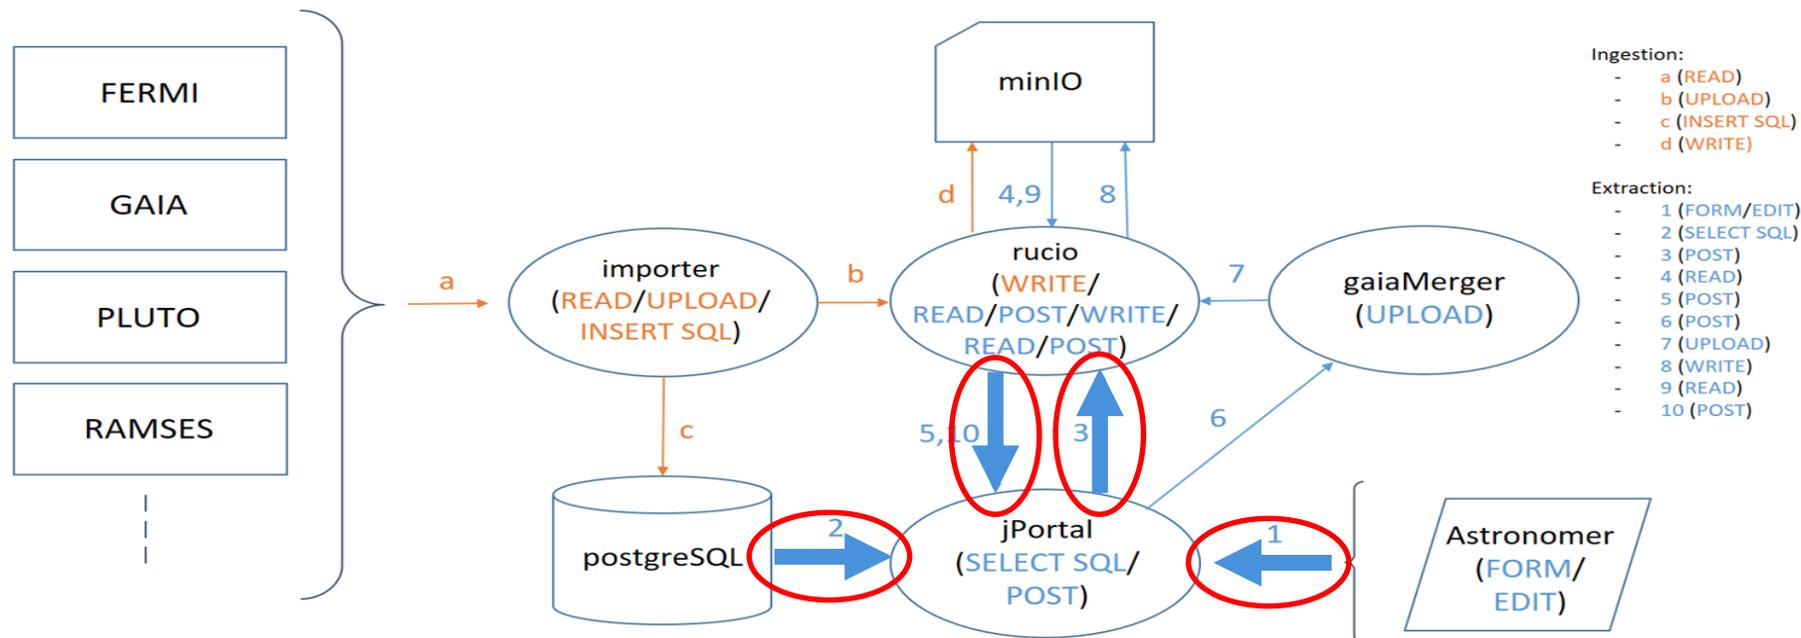
# Data and Metadata Ingestion

- The metadata thus extracted is written to a PostgreSQL database on a dedicated machine while the data is passed to a specially prepared Rucio infrastructure that manages data storage and replication in an S3 object storage managed via MinIO.



# Data and Metadata Retrieval

- As for data recovery, we have created a dedicated portal that can interface with the underlying database and communicate with Rucio for file recovery;
- Communication between the portal and the database therefore passes through a service that has been specifically developed, but which also acts as a TAP service, responding in a standard manner to calls received from commonly used clients that comply with the IVOA directives.



# JPortal Retrieval

≡ jPortal

Help

Settings

Observation

Fermi

Gaia

Simulation

Pluto

Ramses

Resolve Edit Search Cut & Merge Reset

completesource ▾

Form

objectName

sourceId

alpha

delta

epsilon

alpha

Min: 1.61364

Max: 1.63109

alphaStarError

Min: 0.0119842

Max: 971.737

Console

- The astronomer can create a query by selecting an objectName, a sourceId, or any other available metadata;



# JPortal Retrieval

≡ jPortal

Help Settings

Observation  
Fermi Gaia

Simulation  
Pluto Ramses

Resolve Edit Search Cut & Merge Reset

Sources

<input type="checkbox"/>	file_name	source_id	alpha	alpha_star_error	delta
<input type="checkbox"/>	<a href="#">CompleteSource_130097_0003</a>	3425114617586286848	1.618320372472243	0.0252838789312494	0.40193937806911
<input type="checkbox"/>	<a href="#">CompleteSource_130097_0003</a>	3425114617586288128	1.6182774194347807	0.08808795660406121	0.4018892840683867
<input type="checkbox"/>	<a href="#">CompleteSource_130097_0003</a>	3425114617587515008	1.6182973429850582	2.864232984901267	0.4019187701892775
<input type="checkbox"/>	<a href="#">CompleteSource_130097_0003</a>	3425114823742974720	1.6182802384704114	20.517780727768972	0.4020303469978408
<input type="checkbox"/>	<a href="#">CompleteSource_130097_0003</a>	3425114823745120512	1.6182794166233065	0.534499046224391	0.4020499999301908

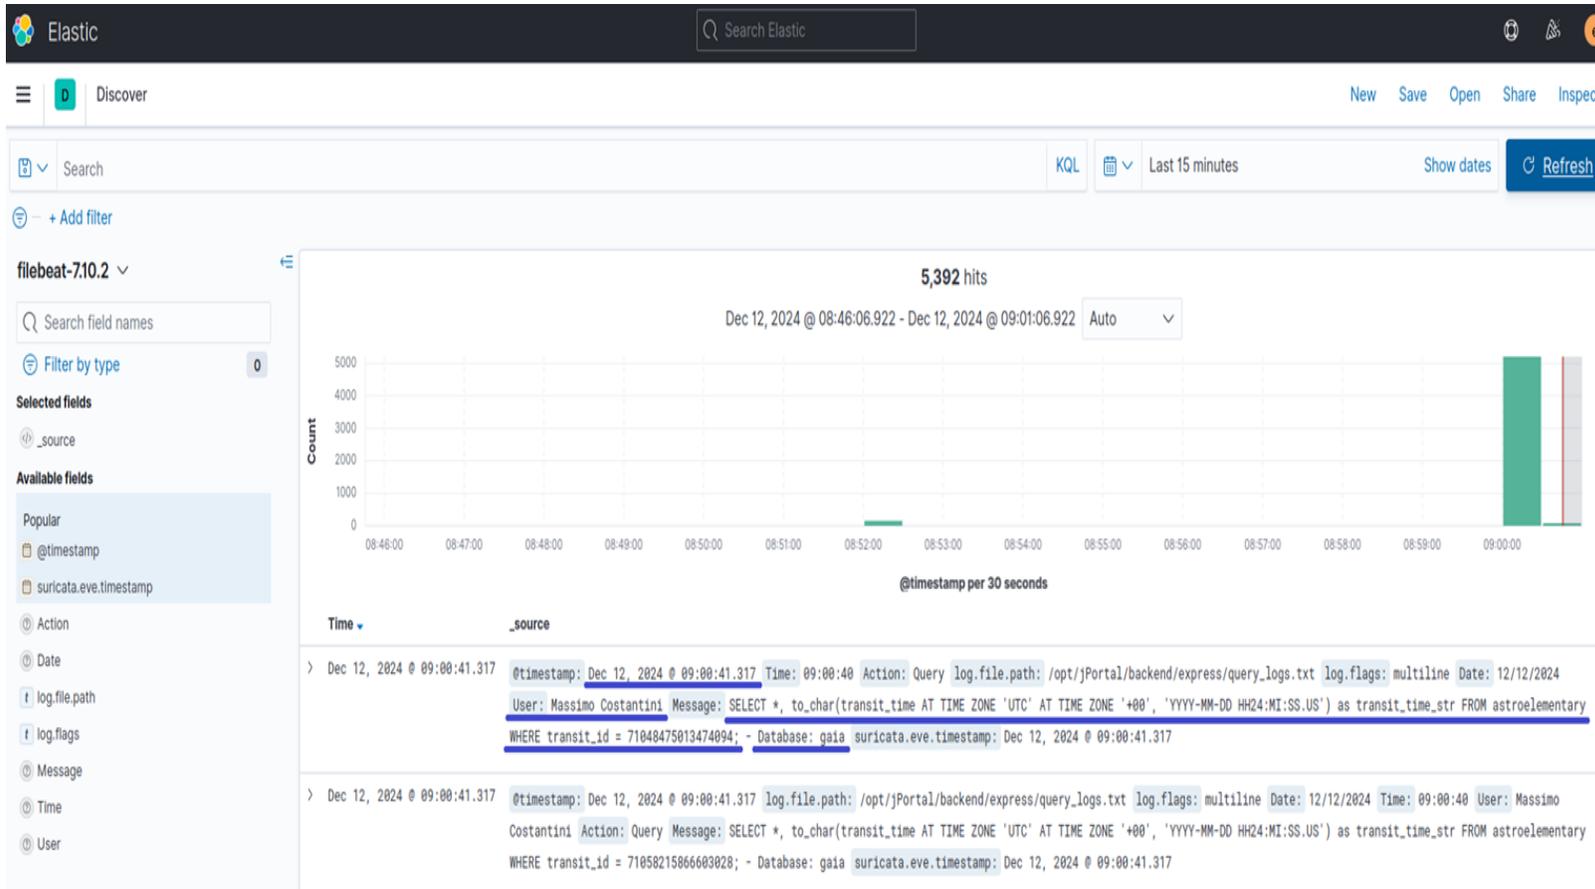
Transits

<input type="checkbox"/>	file_name	source_id	transit_id	ac_win_coord	transit_time
<input type="checkbox"/>	<a href="#">AstroElementary_134755_0006</a>	3425114617586286848	71044382629193047	[ 6008, 10765, 7849, 23375, 5877, 2960, 16268, 27836, 28160, 4277 ]	2017-05-05 15:27:27.522425
<input type="checkbox"/>	<a href="#">AstroElementary_134756_0006</a>	3425114617586286848	71048474888300196	[ 9139, 14738, 21180, 27859, 3839, 32384, 20020, 18213, 13261, 19935 ]	2017-05-05 17:14:01.677219
<input type="checkbox"/>	<a href="#">AstroElementary_134757_0006</a>	3425114617586286848	71058215738283403	[ 32345, 14328, 19099, 11576, 21558, 31161, 18624, 26578, 20740, 5842 ]	2017-05-05 21:27:41.755278

Console

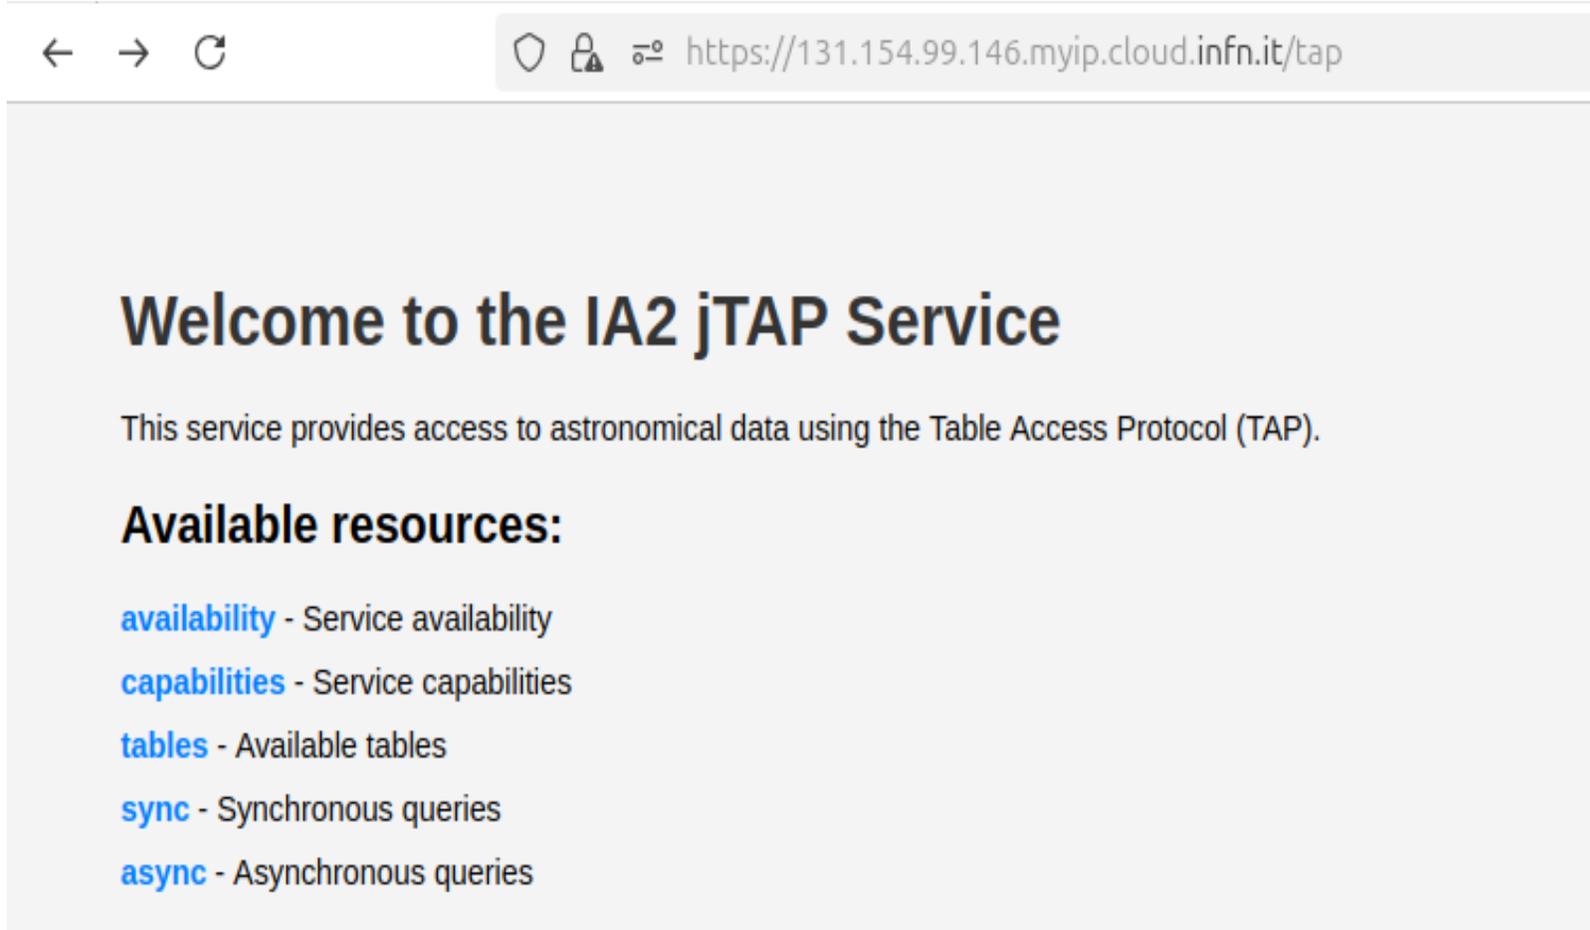
- The astronomer can create a query by selecting an **objectName**, a **sourceId**, or any other available metadata;
- An editor also allows for creating or simply copying/pasting more specific queries;
- The result is a pair of tables that represent not only the metadata retrieved from PostgreSQL but also the corresponding files obtained from Rucio and downloadable from MinIO.

# Elasticsearch / Kibana



- All queries and downloads are logged using Elasticsearch and visualized via Kibana, including details such as timestamps, users and executed queries.

# TAP Service Retrieval



The screenshot shows a web browser window with the address bar containing the URL `https://131.154.99.146.myip.cloud.infn.it/tap`. The main content of the page is titled "Welcome to the IA2 jTAP Service" and includes a description: "This service provides access to astronomical data using the Table Access Protocol (TAP)." Below this, there is a section titled "Available resources:" with a list of links: [availability](#) - Service availability, [capabilities](#) - Service capabilities, [tables](#) - Available tables, [sync](#) - Synchronous queries, and [async](#) - Asynchronous queries.

- **Adding a /tap endpoint to the URL, you access a 'jTAP,' the Table Access Protocol service that, as specified by the IVOA standard, provides a series of additional endpoints;**

# TAP Service Retrieval

```
-<vos:tableset xsi:schemaLocation="http://www.ivoa.net/xml/VODDataService/v1.1 http://www.ivoa.net/xml/VODDataService/v1.1 http://www.ivoa.net/xml/VOSITables/v1.0 http://vo.ari.uni-heidelberg.de/docs/schemata/VOSITables-v1.0.xsd">
-<schema>
  <name>TAP_schema</name>
-<table>
  <name>TAP_schema.columns</name>
  <description/>
-<column>
  <name>utype</name>
  <description/>
  <dataType xsi:type="vod:TAPType">varchar</dataType>
</column>
-<column>
  <name>description</name>
  <description/>
  <dataType xsi:type="vod:TAPType">varchar</dataType>
</column>
-<column>
  <name>column_size</name>
  <description/>
  <dataType xsi:type="vod:TAPType">integer</dataType>
</column>
-<column>
  <name>table_name</name>
  <description/>
  <dataType xsi:type="vod:TAPType">varchar</dataType>
</column>
-<column>
  <name>column_name</name>
  <description/>
  <dataType xsi:type="vod:TAPType">varchar</dataType>
</column>
-<column>
  <name>datatype</name>
  <description/>
  <dataType xsi:type="vod:TAPType">varchar</dataType>
</column>
```

- Adding a /tap endpoint to the URL, you access a 'jTAP,' the Table Access Protocol service that, as specified by the IVOA standard, provides a series of additional endpoints;
- For example, clicking on 'tables', you get a VOTable containing all the metadata in the standard XML format;

# TAP Service Retrieval

The screenshot shows the 'Table Access Protocol (TAP) Query' application window. The interface includes a menu bar (Window, TAP, Authentication, Registry, Edit, Interop, Help) and a toolbar with icons for search, refresh, help, and close. Below the toolbar are tabs for 'Select Service', 'Use Service', 'Resume Job', and 'Running Jobs'. The main area is divided into two panes. The left pane, titled 'Metadata', contains a search field and a tree view showing a hierarchy of services: 'TAP Service (10)', 'TAP\_schema (5)', 'fermi (2)', and 'gaia (3)'. The right pane, titled 'ADQL', shows a table with columns: Name, Type, Unit, Indexed, Descrip..., UCD, Utype, Xtype, and Flags. Below the panes is a 'Service Capabilities' section with 'Query Language: ADQL-2.0', 'Max Rows: [empty]', and 'Uploads: unavailable'. The 'ADQL Text' section has a 'Query Mode: Synchronous' dropdown and a text area containing the query: `SELECT * FROM gaia.completesource WHERE alpha < 1.61835 AND alpha > 1.61815 AND delta < 0.40205 AND delta > 0.40185;`. At the bottom, there is an 'Examples' section and a 'Run Query' button.

- Adding a /tap endpoint to the URL, you access a 'jTAP,' the Table Access Protocol service that, as specified by the IVOA standard, provides a series of additional endpoints;
- For example, clicking on 'tables', you get a VOTable containing all the metadata in the standard XML format;
- This also means that it is now possible to execute the same queries from a client other than jPortal, such as TOPCAT...

# TAP Service Retrieval

TOPCAT(1): Table Browser

Window Rows Help

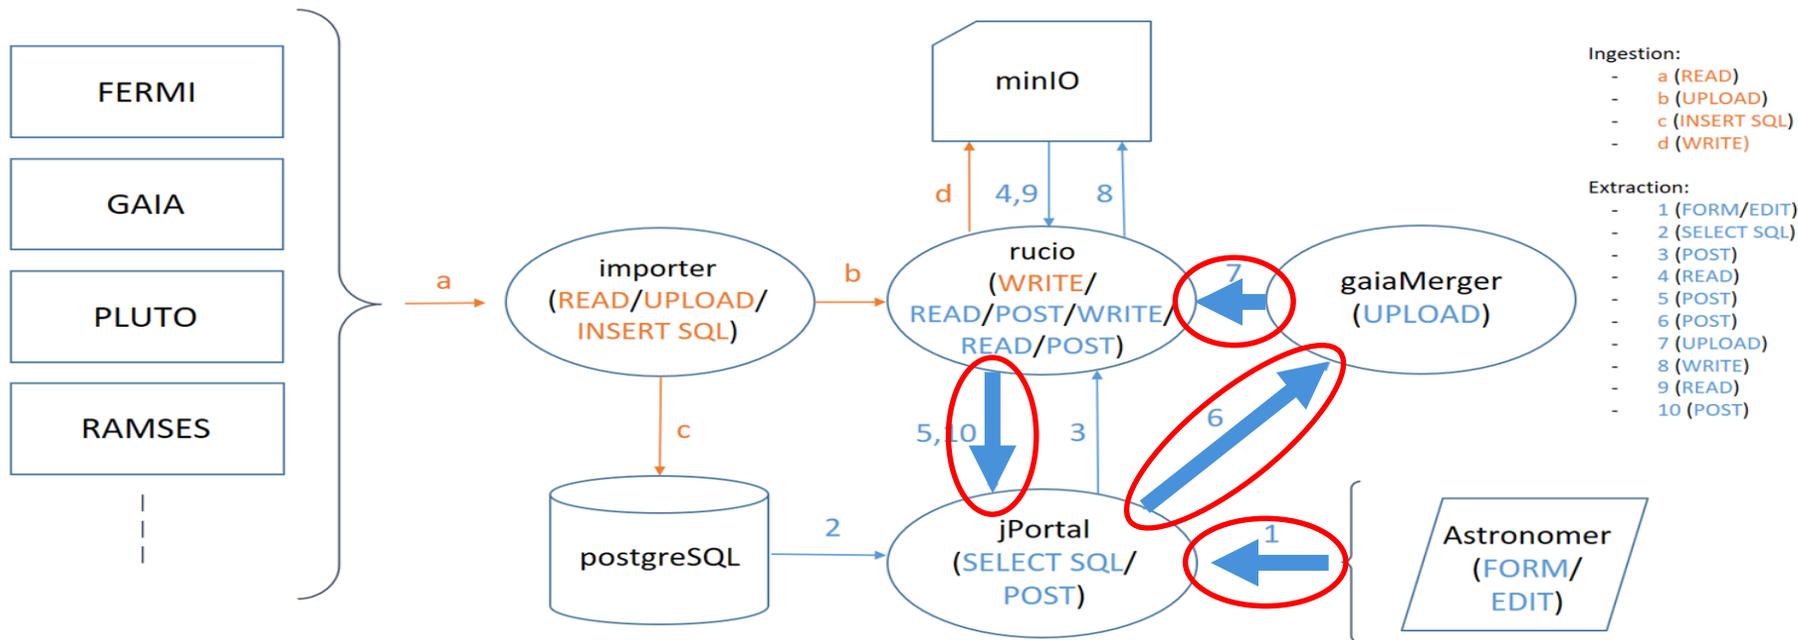
Table Browser for 1: TAP\_1\_gaia.completesource

	id	did_rse	did_scope	did_name	checksum	file_version	file_name
1	39432	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003
2	39438	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003
3	39441	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003
4	39635	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003
5	39640	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003
6	39642	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003
7	40112	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003
8	40118	TEST_USERDISK	test	CompleteSource_test_2025-02-16_0_CompleteSourc...	df04ee8881a1be3d3762ed41fe3551ca	0	CompleteSource_130097_0003

- Adding a /tap endpoint to the URL, you access a 'jTAP,' the Table Access Protocol service that, as specified by the IVOA standard, provides a series of additional endpoints;
- For example, clicking on 'tables', you get a VOTable containing all the metadata in the standard XML format;
- This also means that it is now possible to execute the same queries from a client other than jPortal, such as TOPCAT...
- ...obtaining the same metadata.

# Additional Services

- We are also working to integrate additional services to the archive, among these the GaiaMerger service is already available, which allows you to perform a Cut&Merge of Gaia files to generate a much lighter and easier to manipulate output than those present in the archive.



# Additional Services

≡ jPortal

- Help
- Settings
- Observation
  - Fermi
  - Gaia
- Simulation
  - Pluto
  - Ramses

Resolve Edit Search Cut & Merge Reset

Cut & Merged

file\_name

[gaia\\_merger\\_23629d157c73b4a8d31f87722fb5d2b2\\_1739889698.73819](#)

Console

- The GaiaMerger service is currently available only from jPortal;
- The file is immediately downloadable and is preserved only for a limited time;
- We are also working to integrate services such as DataLink that will allow standardized calls to this type of services as well.

# Conclusions and Next Steps

## Conclusions:

- **As you can see, it is possible to build a data lake that meets scientific standards;**
- **NADIR's modularity has been improved, allowing code variations to be minimized when the format and structure of the input files vary;**
- **The main challenges encountered are the obsolescence of software and libraries, the high demand for storage resources and most importantly the complexity in communication among stakeholders;**

## Next Steps:

- **We are working to structure simulated data in a coherent and, most importantly, standardized manner to ensure correct integration into the Spoke 3 database;**
- **As mentioned, we are developing the necessary standard services for proper integration with IVOA clients, as for the DataLink;**
- **We aim to integrate analysis services such as FermiTools for Fermi data, making them accessible both from JPortal and standard clients.**



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani

PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC,  
Big Data and Quantum Computing

*Thank you for your attention!*  
*Questions are welcome.*

Archives and Data Management Systems, Bologna Feb 26-27, 2025