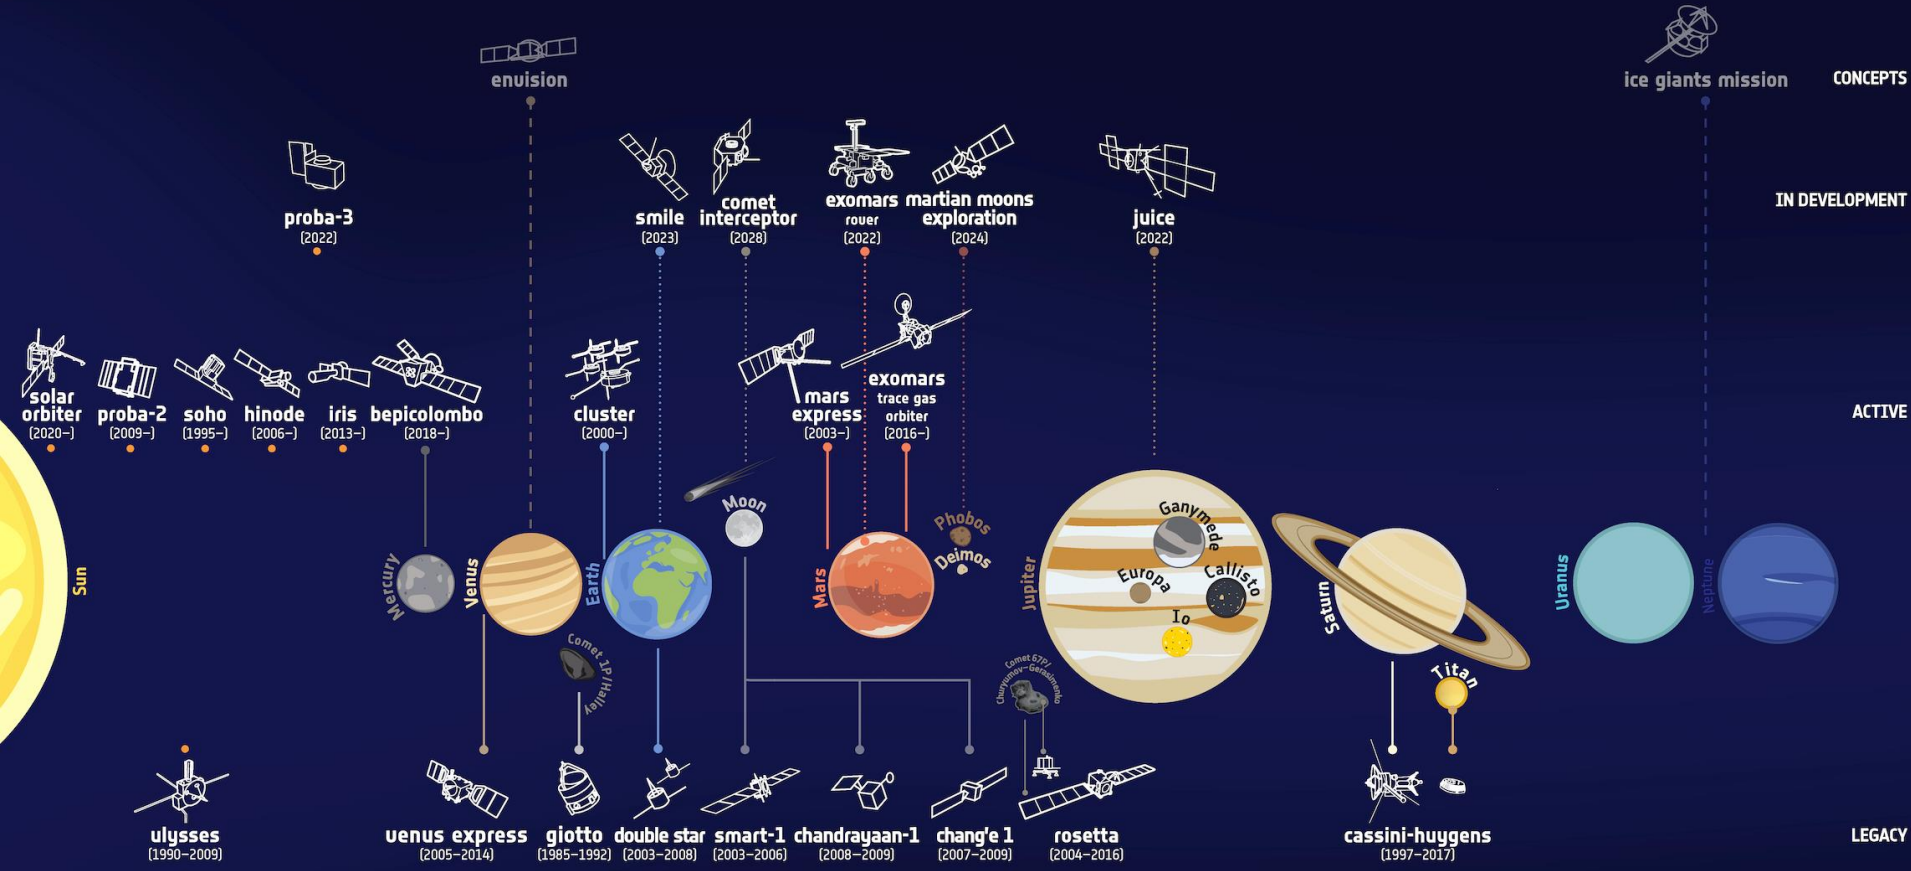


# IT storage solutions for big data space missions


Rubén Alvarez


28/10/2020

# → SOLAR SYSTEM EXPLORERS



# → COSMIC OBSERVERS





50 Gbytes of data per  
day (Gaia)

# DATA GENERATED PER DAY

100 PB 100000

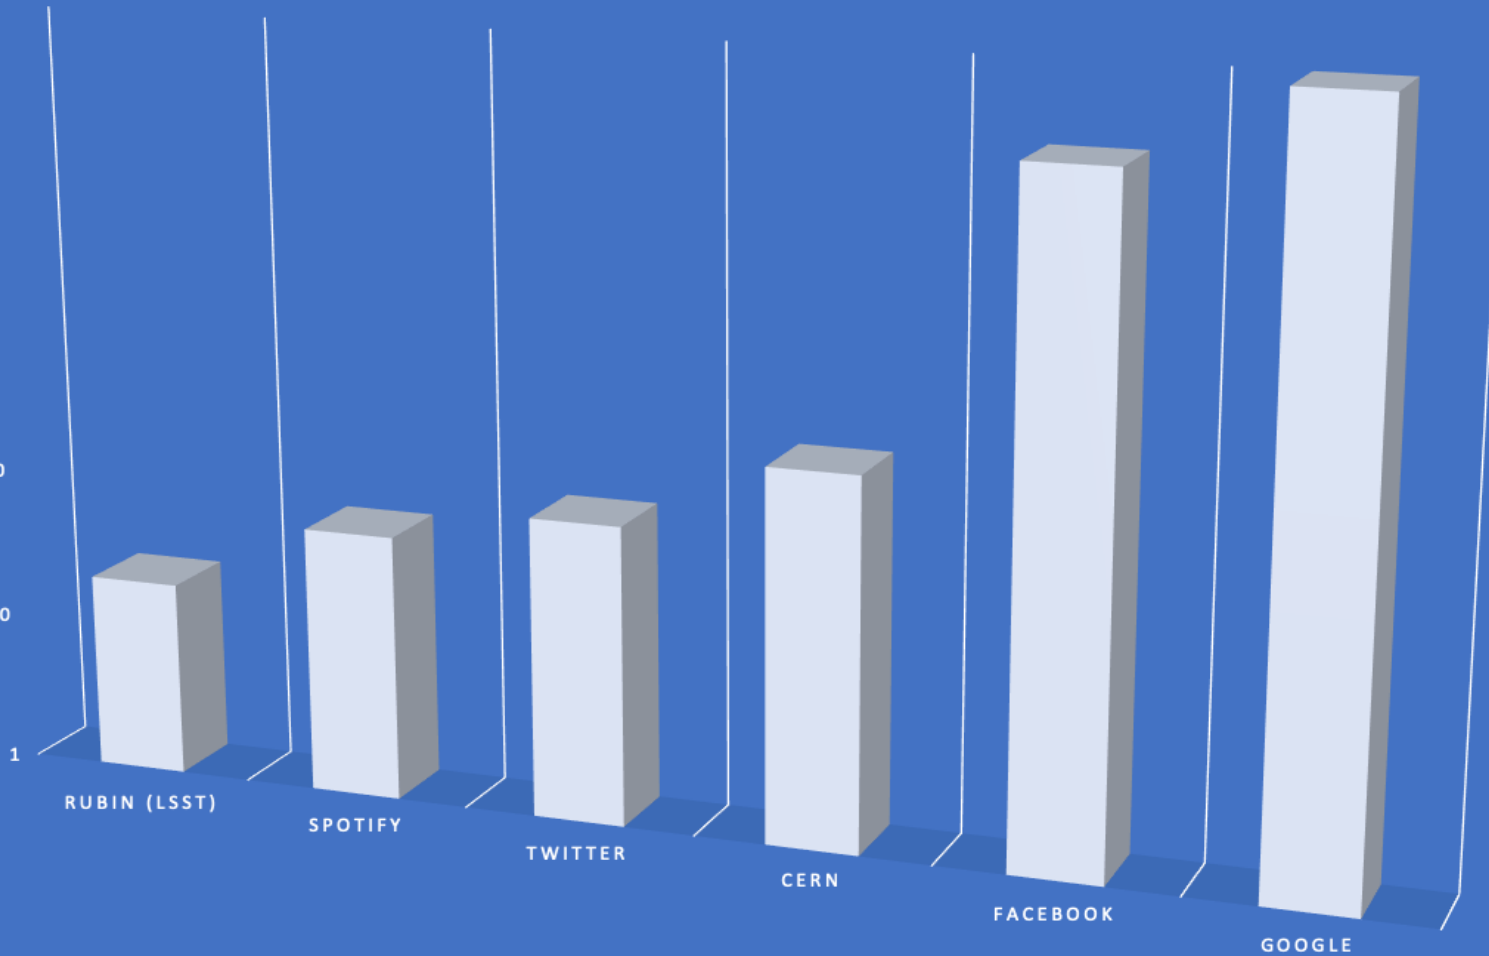
10000

1 PB 1000

100

10

1 TB 1



RUBIN (LSST)

SPOTIFY


TWITTER

CERN

FACEBOOK

GOOGLE

# ESA Space Science Archives - Volume (PB)






Moore's law: Computer processing power would double every 18 months

Kryder's law: Disk capacity would double every 13 months

Nielsen's law: Network bandwidth would double every 21 months

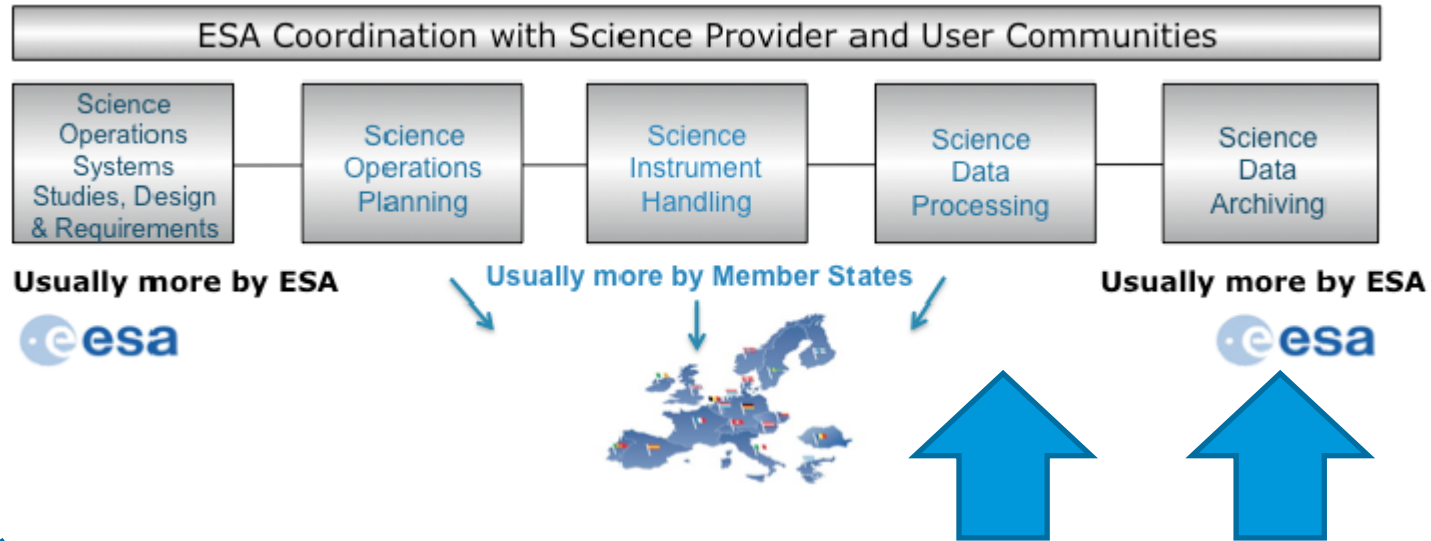
# Space Segment




# OGS



# Science Ground Segment (SGS)



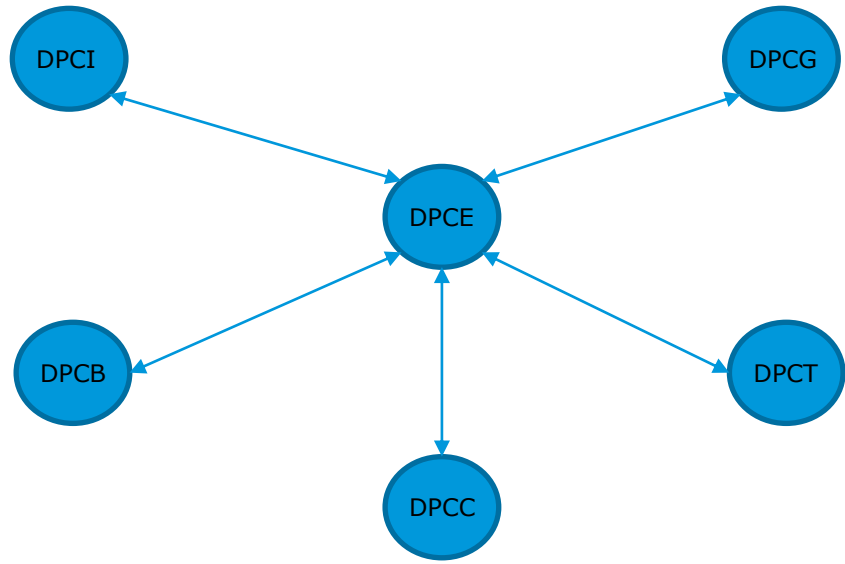




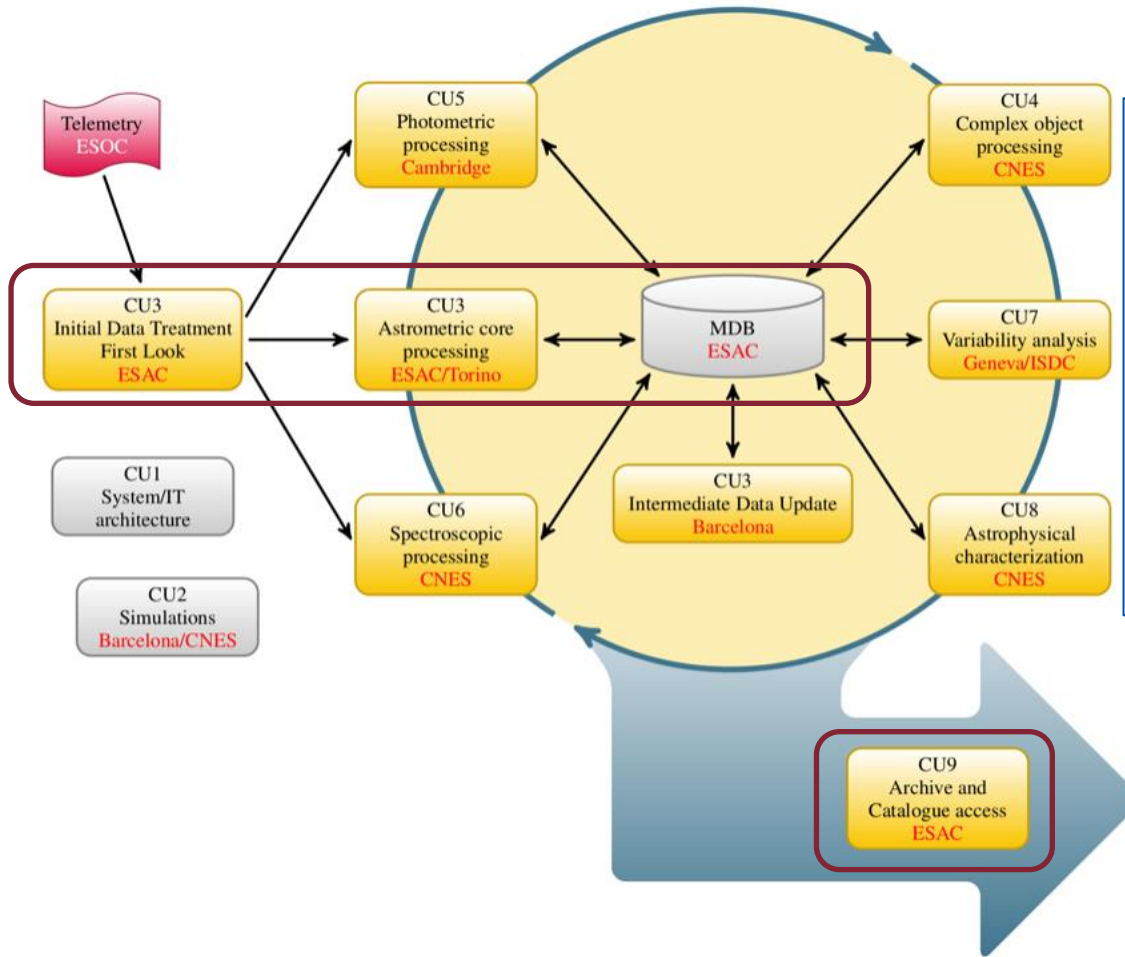
Launched on 19.12.2013

Objective: to make the largest, most precise 3D map of our galaxy by surveying more than 1000 million stars

160 TB raw TM -> 1.5 PB science products

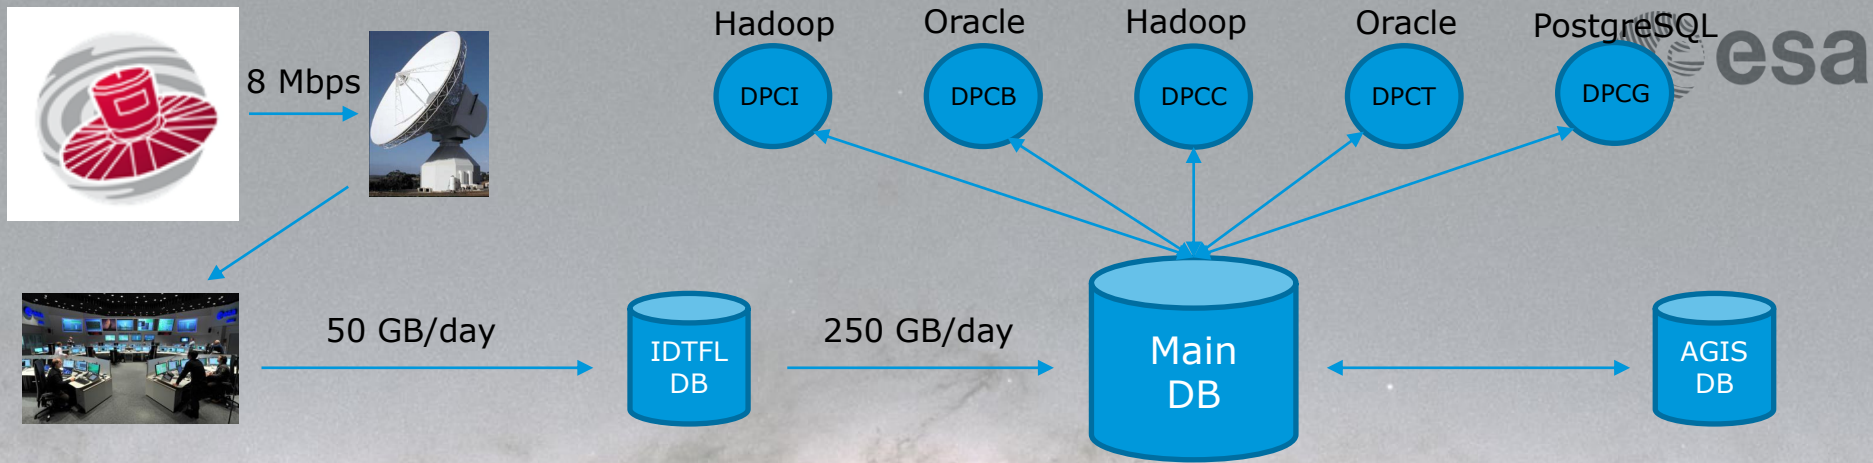


- Between 250 GB and 450 GB transferred every day
- Aspera (IBM) is the software used for data transfers
- Very efficient: UDP with fault tolerance, near-zero latency, can use all available bandwidth if needed



- Data transfer software common for all DPCs
- Each DPC can use their own storage systems, DBs and compute clusters. Data produced has to be consistent with MDB data model
- Daily processing + cyclic processing





## ESAC

- DBMS - Intersystems Cache
- Main driver – Performance (store tens of thousands of Java objects per second)
- IT storage – DB servers SSD local disks + network storage based on spinning disks with SSD cache (NetApp)
- Capacity – MainDB currently 775 TB. It stores metadata. Data is stored in NAS

# Welcome to the Gaia Archive at ESA

Gaia is a European space mission providing astrometry, photometry, and spectroscopy of more than 1000 million stars in the Milky Way. Also data for significant samples of extragalactic and Solar system objects is made available. The Gaia Archive contains deduced positions, parallaxes, radial velocities, and brightnesses. Complementary information on multiplicity, photometric variability, and astrophysical parameters is provided for a large fraction of sources.



- DBMS – PostgreSQL
- Main driver for the DB – performance
- PostgreSQL has specific functions for Astronomy
- Current size of the science archive – tens of TBs

## Top Features



How to cite and acknowledge Gaia.



Query for Gaia sources using an ADQL (Astronomical Data Query Language) interface in





Direct download of Gaia data files.



For questions, suggestions or problem reports, contact the Helpdesk.

## gaia archive visualisation





Launch date:  
July – December 2022

Objective: to study the  
geometry and the nature of  
the dark universe

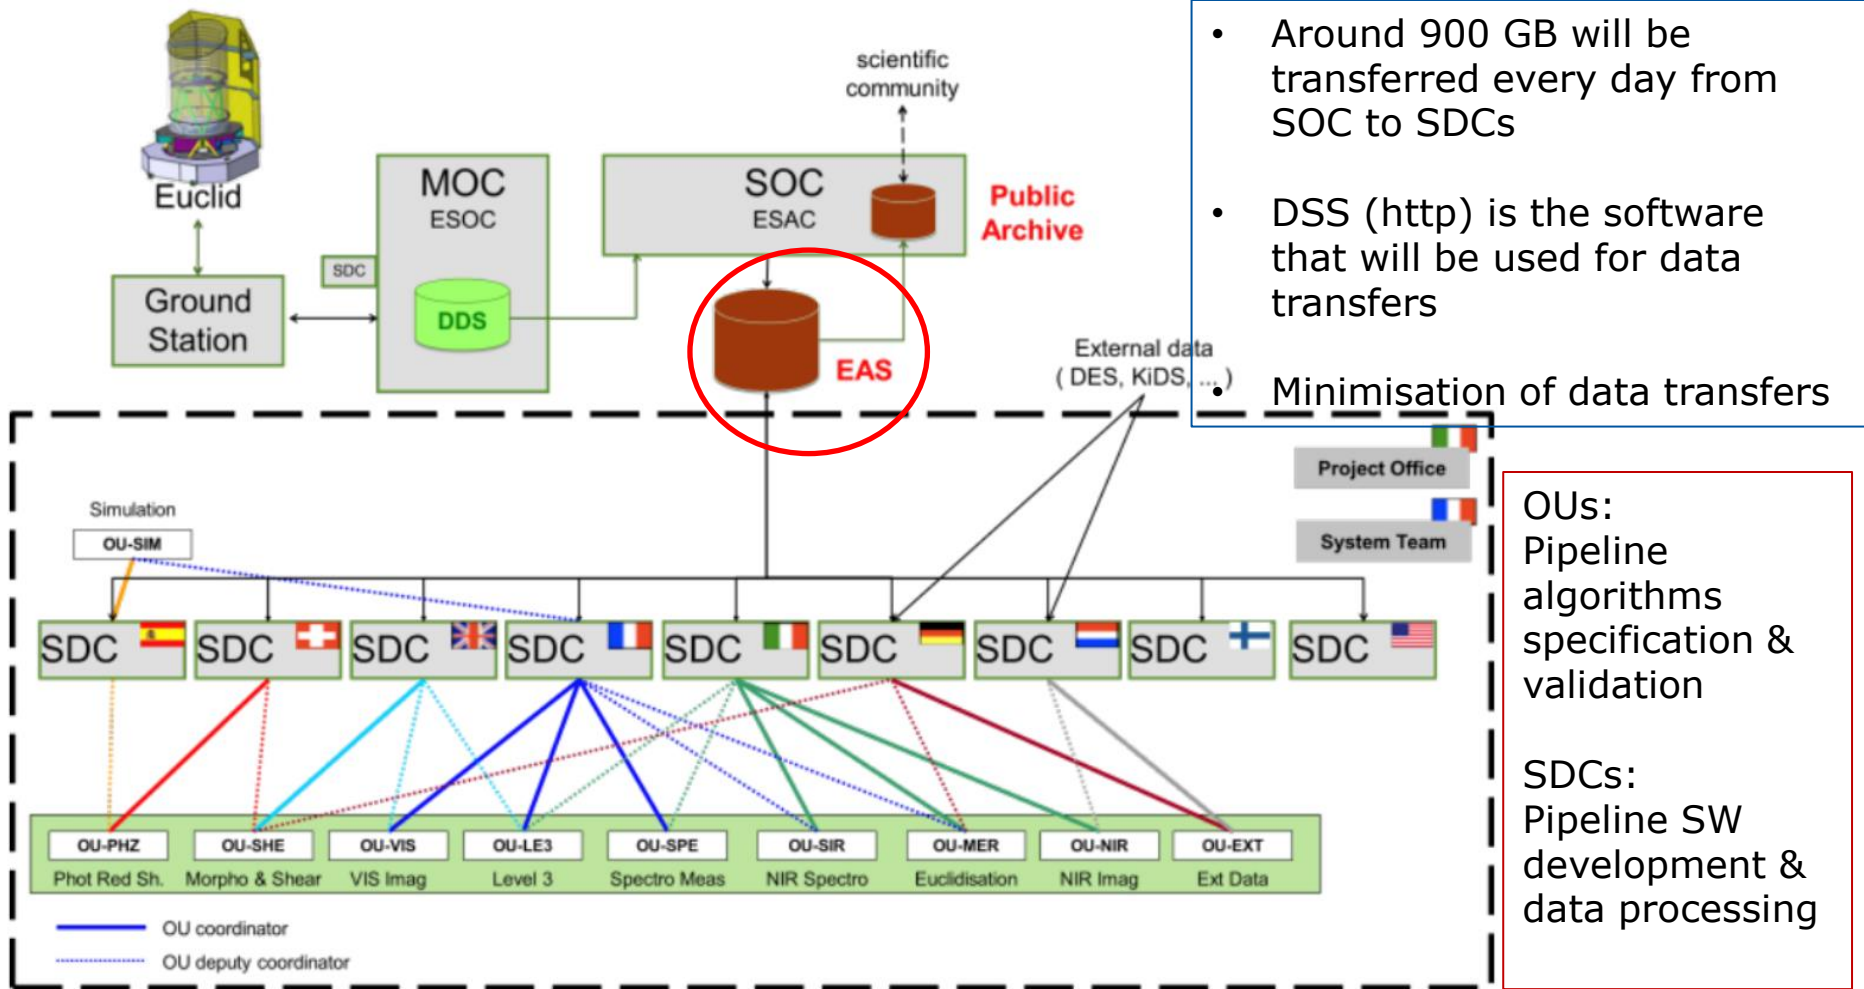
$10^{10}$  galaxies will be observed

200 TB raw TM

10 PB Science Archive

10 – 100 PB images of ground based  
observatories

The mission could generate 26 PB per  
year



Organisation of the Science Data Centers (SDC) under responsibility of the Euclid Consortium. The SDC are computing centers in charge of implementing and running the data processing pipelines – Courtesy Euclid Consortium/ESA/SGS Team.



## Euclid Archive System – EAS

Distributed data sets across SDCs and centralized metadata (ESAC)

Registers any operation with data

Stores dependencies of data products

Keeps information on each stage of the processing

### DPS

Supports the data processing activities

### SAS

Supports data releases and long-term data preservation

### DSS

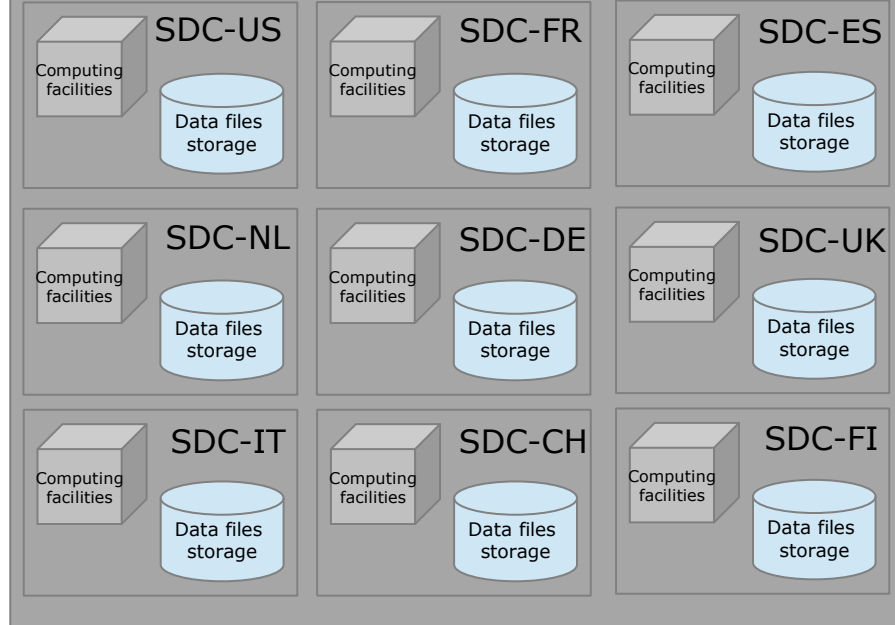
Used to store data files

File location in metadata provides a “global filesystem”

Data Processing System  
(EAS-DPS)

Science Archive System  
(EAS-SAS)

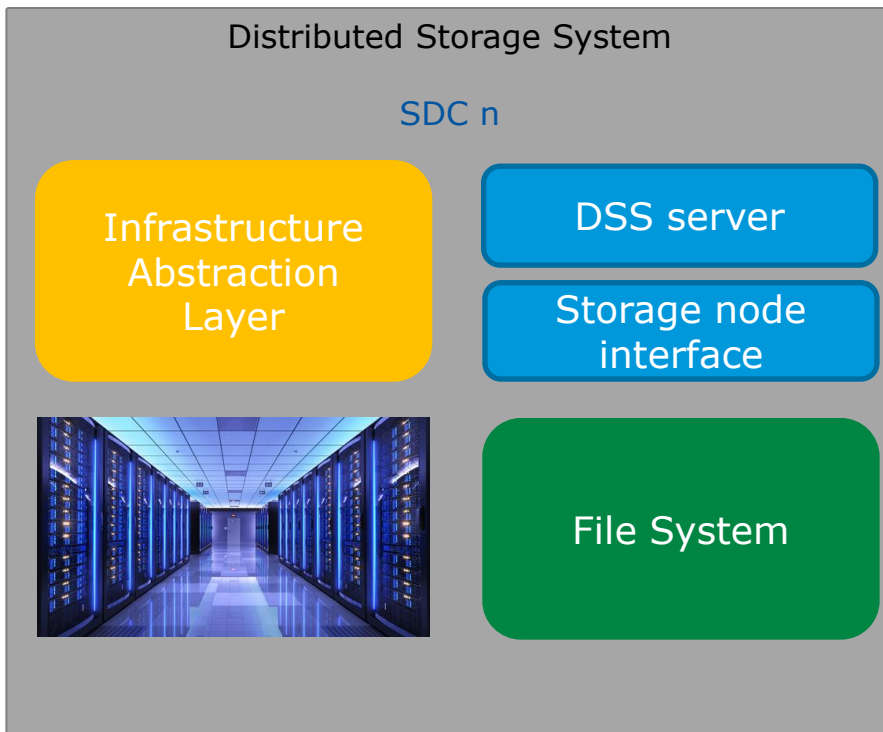

## Distributed Storage System (EAS-DSS)





Data Processing System  
(EAS-DPS)

Science Archive System  
(EAS-SAS)



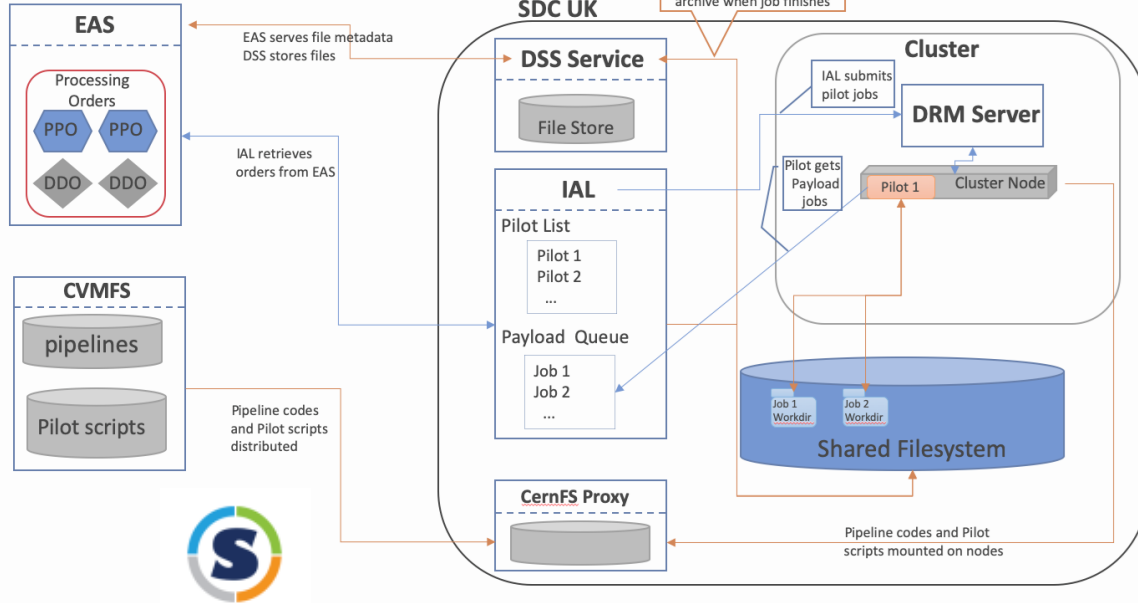
Allows the data processing software to run on any SDC independently of the underlying IT infrastructure

Filesystems can be specific per SDC, e.g. GPFS, Ceph, BeeGFS, etc

Large number of small files and small IO patterns

### Euclid Architecture

#### SGS Consortium Services



Parallel filesystems for the pipelines in order to minimize data moves

Spinning disks + SSDs for cache

Ethernet, in some cases infiniband

Slurm as job scheduler for most of the clusters

Singularity as the container technology

Mark S. Holliman



## WELCOME TO THE EUCLID SCIENCE ARCHIVE SYSTEM

Euclid is an ESA mission to map the geometry of the dark Universe:

- Discover the origin of the Universe's accelerating expansion.
- Discover the nature of 95% of the Universe: dark energy and dark matter.
- Measure shapes of galaxies distorted by gravitational deflection due to dark matter.
- Measure non-random distribution of galaxies resulting from the action of gravity.



## DBMS – PostgreSQL

Main driver for the DB – several reasons: good performance, astronomical functions, big community behind ...

Moving towards Greenplum for the Euclid catalogue in order to achieve the needed performance and scalability

Estimated final size of the science archive – 10 PB (by 2030)



«YOU CAN EITHER MOVE YOUR QUESTIONS OR THE DATA. [...] OFTEN IT TURNS OUT TO BE MORE EFFICIENT TO MOVE THE QUESTIONS THAN TO MOVE THE DATA.»

Jim Gray, eScience: A Transformed Scientific Method

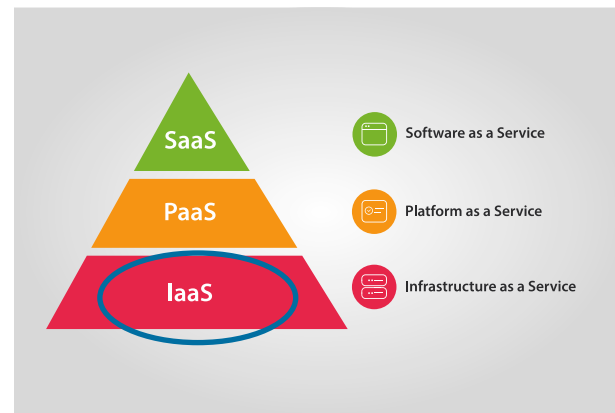

Vicente Navarro

## BRING YOUR QUESTIONS TO THE DATA

There is a new paradigm, opening completely new opportunities for discovery – a data-intensive approach to science. In many domains, we have entered what could be called the golden age of surveys, with several large-scale projects, spanning decades, between finished, ongoing, and planned activities. ESA is responsible, or is a major partner, in several of these initiatives.

There is, however, a new profound change: data has become a major technological challenge. Increases by multiple orders of magnitude in dataset size means that transferring data to a scientist is often unfeasible.

ESA datalabs gives you a privileged position; bring your code directly to ESA's infrastructure – there is a great set of tools and programming languages are flexible – and execute it with direct access to ESA's archives.



ESDC JupyterLab  
tbd

Euclid DPS JupyterLab  
tbd

GSSC JupyterLab  
tbd

Herschel JupyterLab  
tbd

- Persistent Area
- Persistent Shared Area
- Volatile Area

esa | datalabs

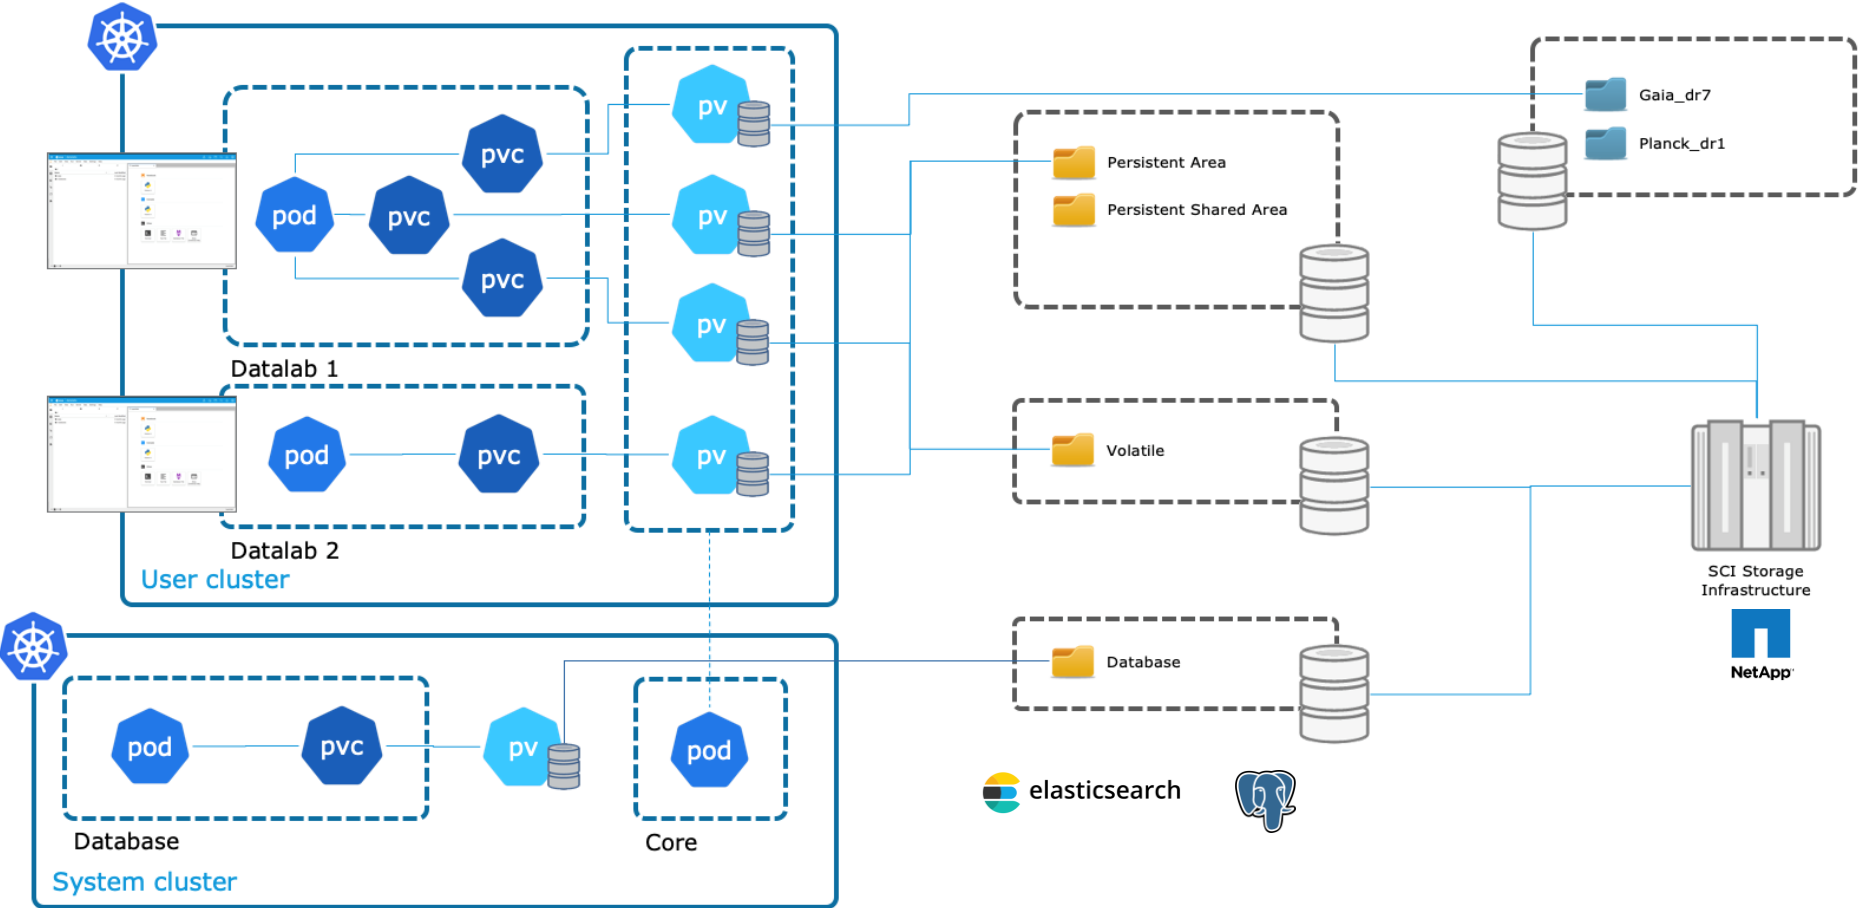
## Data Volumes

+ add from catalog   + add custom

- My data selection from Planck** (ftp icon)  
/media/data/planck1  
ftp://guest:\*\*\*\*@planck.esa.int/user\_sel/3...
- My Dropbox account** (dropbox icon)  
/media/data/dropbox  
dropbox://myusr:\*\*\*\*@dropbox
- Gaia data release v7** (nfs icon)  
/media/data/gaia\_data  
nfs://nfs-server2.esa.int/gaia/v7
- Some random webdav data** (webdav icon)  
/media/data/webdav\_test  
webdav://user:\*\*\*\*@xyz.esa.int/some/path
- My Google Drive account** (gdrive icon)  
/media/data/gdrive  
gdrive://user:\*\*\*\*@gdrive
- My home server** (sshfs icon)  
/media/data/myhome  
sshfs://user:\*\*\*\*@home.dyndns.example/home/...

- Data Collections
- + Gaia data release

- Read-write access
- Read-only access



ESA UNCLASSIFIED - For Official Use



## Thanks to:

Bruno Merin, European Space Agency

Christophe Arviset, European Space Agency

Guillermo Buenadicha, European Space Agency

Mark. S. Holliman, Royal Observatory, Edinburgh

Miguel Angel Diego, European Space Agency

Rocio Guerra, European Space Agency

Sara Nieto, European Space Agency

Vicente Navarro, European Space Agency