Computing in Euclid-IT Science (Cosmology)

Carmelita Carbone (on behalf of many people)

ISTITUTO NAZIONALE DI ASTROFISICA
NATIONAL INSTITUTE FOR ASTROPHYSICS
INAF
Istituto di Astrofisica Spaziale e Fisica cosmica di Milano

Meeting "INAF USCVIII – General Assembly", Galzignano, 14-18 Oct 2024

Motivations for HPC computing in Euclid

- Computing resources for SWGs (Melita's talk) and SGS (Daniele's talk): different needs for different data analyses of upcoming big data

- Different cosmological probes (GC, WL, Galaxy Clusters, CMBX, …): computational resources for cosmological N-body simulations/mock catalogues, parameter inference, modelling, and single/joint-probe analyses

Computing for Euclid Galaxy Clustering

(P. Monaco, E. Sefusatti, M. Lepinzan, T. Castro, L. Tornatore, E. Sarpa, G. Parimbelli ++)

1. Computing resources to run **PINOCCHIO simulations** for the GC covariance
   (3500 smaller simulations + 1000 big simulations that cover half of the sky):
   INFN computing time (euclid project), Galileo100@CINECA and
   ISCRA-B grant, Galileo100@CINECA and
   UniTS-CINECA agreement
   -> total of ~5M core-hr.
   + Pleiadi system for smaller runs (700,000 core-h).

2. Development: ICSC/Spoke3, 700,000 gpu-hr on Leonardo@CINECA.
3. Postprocessing (creation of galaxy catalogs, application of systematics, mixing matrix, power spectrum measurement): -> Pleiadi system.

The 1000 EuclidLargeMocks are the largest collection of DM halos on the past light cone ever produced! (paper + press release in preparation).

PINOCCHIO Mocks: adding systematics to ideal catalogs

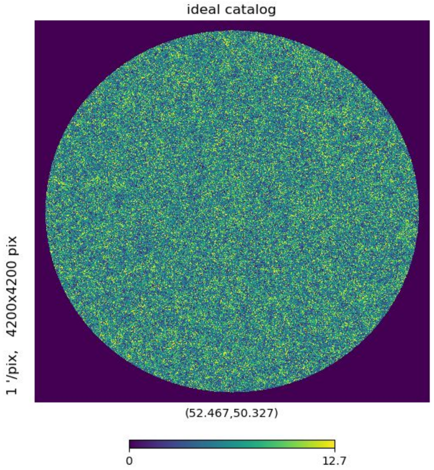

30deg circle is ~ the same area of DR1 (pre-launch)

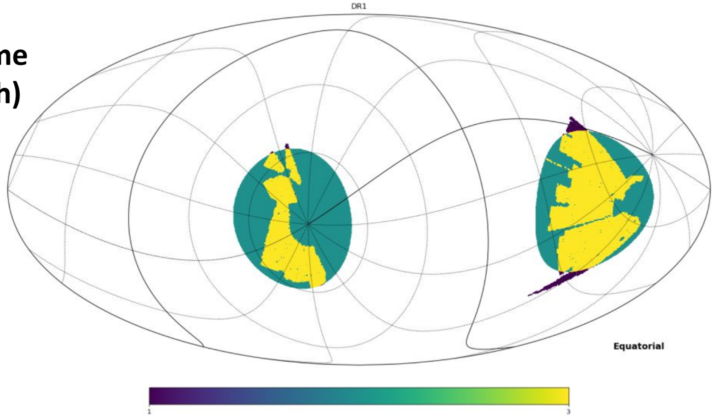

ideal catalogs, average P(k) over 1000 mocks

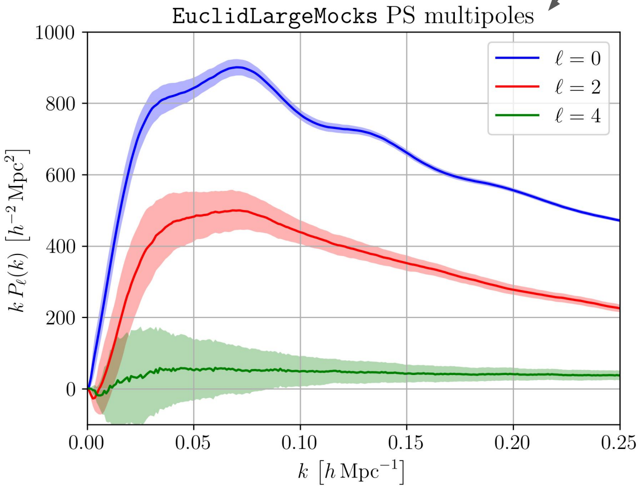

angular systematics: random catalog with varying exposure time due to survey strategy

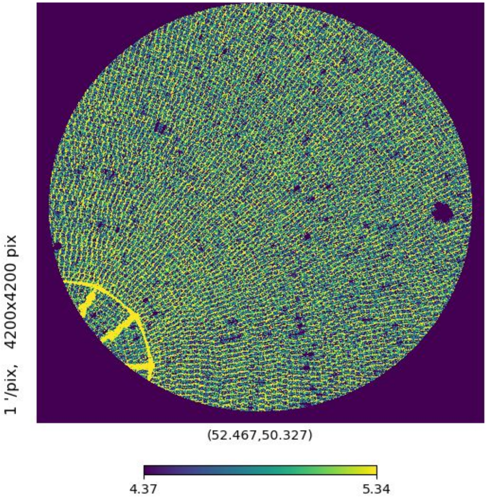

Computing for Euclid Cluster Cosmology

(S. Borgani, T. Castro, M. Costanzi, R. Ingrao, A. Saro)

1. The needed resources for the activities inside the SWG-CL WP-9 for <u>Galaxy Cluster Simulations</u> has been covered in the past by:

- Leonardo@CINECA Early Access Project (~20M core-hr)
- ISCRA-B/CINECA project (~2M core-hr) on cineca
- Couple of small ISCRA-C/CINECA projects for code development (~100k core-hr)
- DRES storage on CINECA (80TB)
- INAF-IA2 Storage (50TB)

⇒ Future activities are covered at the moment by a EuroHPC recently approved project on Leonardo-DCGP (+20M core/h; *PI: T. Castro*)

2. Resources to run chains for <u>cosmological posteriors</u>

- So far based on the use of local clusters (AMONRA@INAF-OATs)
- For DR1 data analysis: expected ~2M core-hr; ISCRA-B@CINECA sufficient in principle, but dedicated and flexible INAF resources would be welcome.

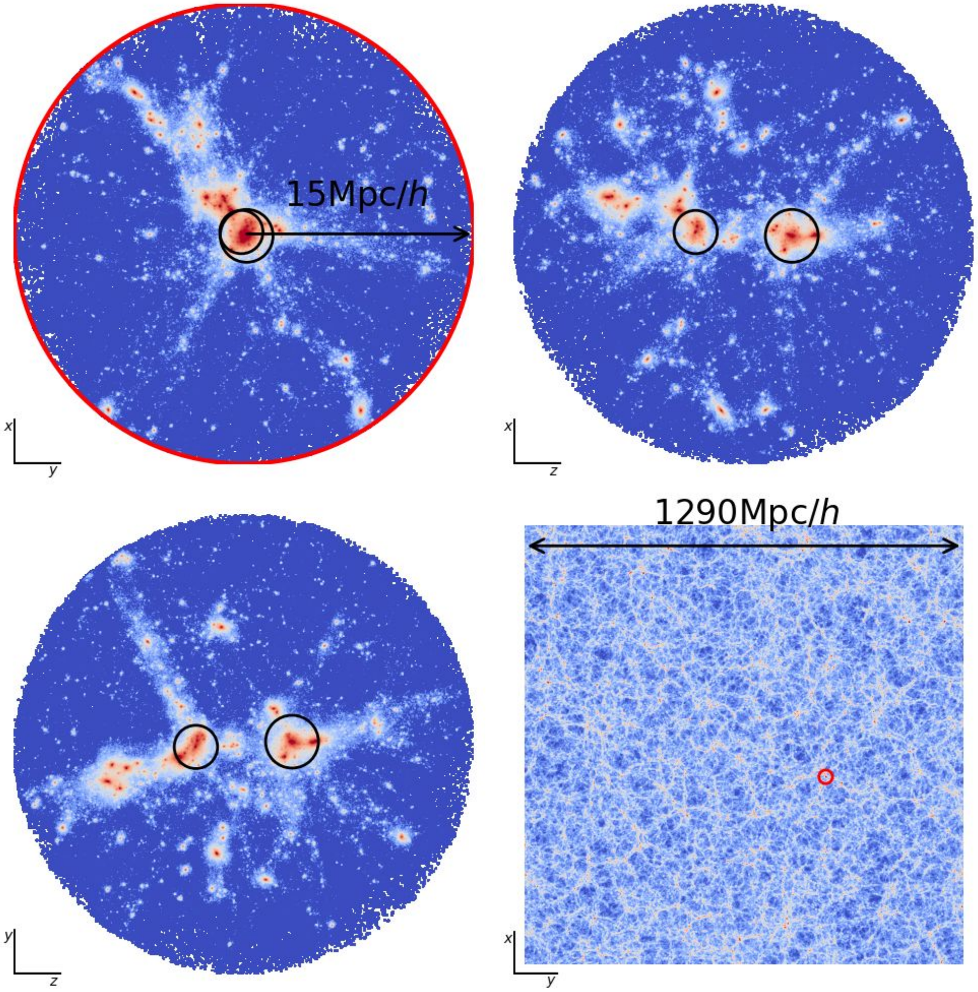

Galaxy clusters and line-of-sight projection effects
credits: Tiago Castro (OATS)

Cosmological simulations for beyond-LCDM models in Euclid

(M. Baldi)

1. **Simulations of cosmologies beyond LCDM** are necessary to achieve Euclid goals on constraining Dark Sector models and have been used in activities of SWG-TH, SWG-CS-WP8, PL-KP-JP6, **using computational resources from:**
   1.1. **PRACE Multi-year Access**: 15 Mio CPU hrs on MareNostrum4 (*DUSTGRAIN* simulations, P.I. M. Baldi)
   1.2. **ISCRA-B projects SIMCODE2 and DZSH**: ~2 Mio CPU hrs on Marconi (*DUSTGRAIN-pathfinder* and *DAKAR2* simulations, P.I. M. Baldi)
   1.3. **ISCRA-C projects DZS, EuNuComp, EuKey, PANDAG4**: ~150k CPU hrs on Marconi, Marconi-100, Galileo-100 (*MassiveNu* Code Comparison simulations, *DUSTGRAIN-pathfinder* extension for HOWLS project)
   1.4. **INFN-InDark** resources at Cineca: ~200k CPU hrs on Marconi-100, Galileo-100 (*CIDER* simulations)
   1.5. **INFN-Euclid** resources at Cineca: ~600k CPU hrs on Galileo-100, Leonardo (*CIDER* simulations)

2. **Storage resources have been provided by:**
   2.1. **INFN-CNAF data center** under Euclid-INFN agreement: ~400TB on disk → storing all raw data and processed data
   2.2. **Local Bologna cluster**: ~20 TB on disk purchased on personal research fundings (M. Baldi) → temporary storage of data for analysis
   2.3. **PIC data center** (Spain): few TB → final products presented in Euclid Publication (for data availability policies)

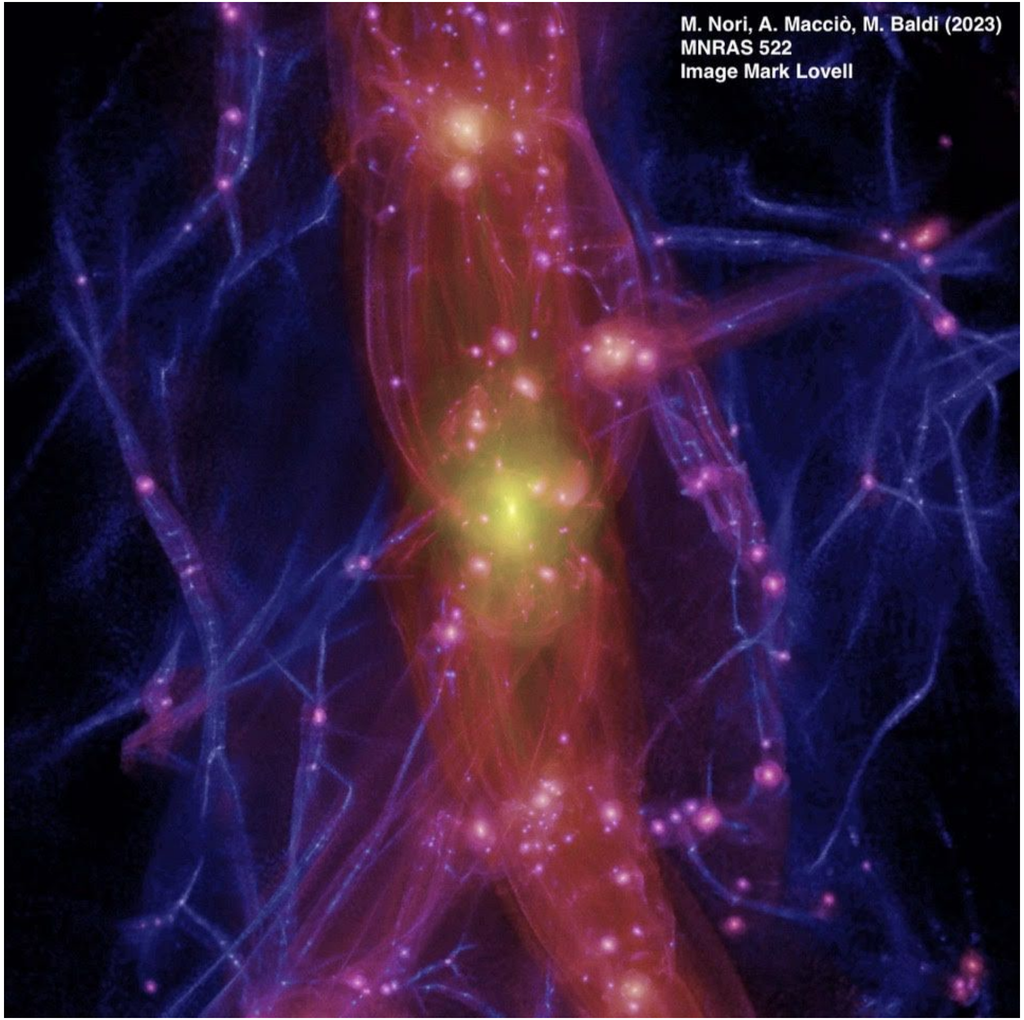

Matter density distribution around an Aquarius halo (Milky-Way like) obtained via AX-Gadget in a Fuzzy dark Matter scenario. Credits: Marco Baldi (UniBO)

Cosmological simulations for beyond-LCDM cosmology in Euclid

Carmelita Carbone: DEMNUni campaign

- **16 DEMNUni XL-simulations:**
  $V=(2\ \mathrm{Gpc}/h)^3$, $N_{part}= 2 \times 2048^3$ (CDM+ν), $M_{cdm} \cong 8 \times 10^{10}\ M_\odot$

baseline Planck cosmology (according to Euclid SWG-coord meeting 11/06/2013)

+

- **M_ν=0, 0.16, 0.32, 0.53 eV (w=-1)** (projected density & particle snaps for all 63 output-times, 5 for Mnu0.32)
- **(M_ν,w_0,w_a)=(0-0.16-0.32 eV, -0.9, ±0.3) & (0-0.16-0.32 eV, -1.1, ±0.3)**
  5 snaps per sim, 34 FoF/Sub, M200_b/c, M500_b/c, M2500_b/c, M_{vir}, galaxy-SHAM, catalogues in 0<z<2 (**240TB OF STORED DATA**)

Implemented Dark Energy Equation of State →

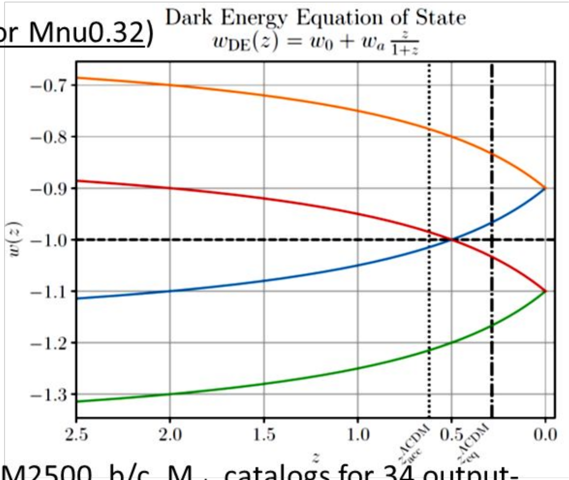

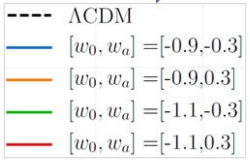

- **50+50 DEMNUni L-simulations (DEMNUni-Cov):**
  $V=(1\ \mathrm{Gpc}/h)^3$, $N_{part}= 2 \times 1024^3$ (CDM+ν), $M_{cdm} \cong 8 \times 10^{10}\ M_\odot$

- **50 sims Planck-LCDM; Mnu=0 eV:** 63 full particle snapshots/sim, FoF/Sub, M200_b/c, M500_b/c, M2500_b/c, M_{vir} catalogs for 34 output-times between z=2 and z=0. Projected densities maps available at all output-times. (**110TB of stored data**)
- **50 sims Planck-LCDM; Mnu=0.16 eV:** 5 full particle snapshots/sim, FoF/Sub, M200_b/c, M500_b/c, M2500_b/c, M_{vir} catalogs for 34 output-times between z=2 and z=0. (**30TB of stored data**)

- **DEMNUni M-simulations (DEMNUni-HigRes): $V=(500\ \mathrm{Mpc}/h)^3$, $N_{part}= 2 \times 2048^3$ (CDM+ν), $M_{cdm} \cong 1.3 \times 10^9\ M_\odot$**
- **2 sims Planck-LCDM; Mnu=0,0.16,0.32 eV:** (resolution enough for Euclid Halpha galaxies)

Total of 26M cpu-hr and 400TB storage @CINECA/INAF/IA2

DEMNUni applications to Euclid

- KP-CL-3 pre-launch paper-4 “*Halo mass function and bias in non-standard models*” of the Galaxy-Clusters SWG.
- KP-CL-3 pre-launch paper-1 “*Calibration of the halo mass function in $\Lambda(\nu)$ cosmologies*” Castro et al. A&A 671, A100 (2023)
- KP-CMBX-2 pre-launch paper2 “*CMBX Mock Simulations*” in CMBX SWG.
- KP-JC-6 pre-launch paper-3 “*Simulations & non-linearities beyond ΛCDM*” joint Theory and Cosmo-SIM SWGs.
- KP-TH-1 pre-launch paper-6 “*Euclid: Nonlinear spectroscopic clustering in beyond-ΛCDM scenarios with Euclid*”, Theory SWG.
- KP-CMBX-2 pre-launch paper-6 “*Numerical covariances for CMB cross Euclid*” (calibration of FLASK lognormal mocks against DEMNUni-CoV mocks)
- KP-CL-2 “*Euclid preparation: Determining the weak lensing mass accuracy and precision for galaxy clusters*”, arXiv:2409.02783

DEMNUni simulations: structure formation in $\nu w_0 w_a$CDM scenarios

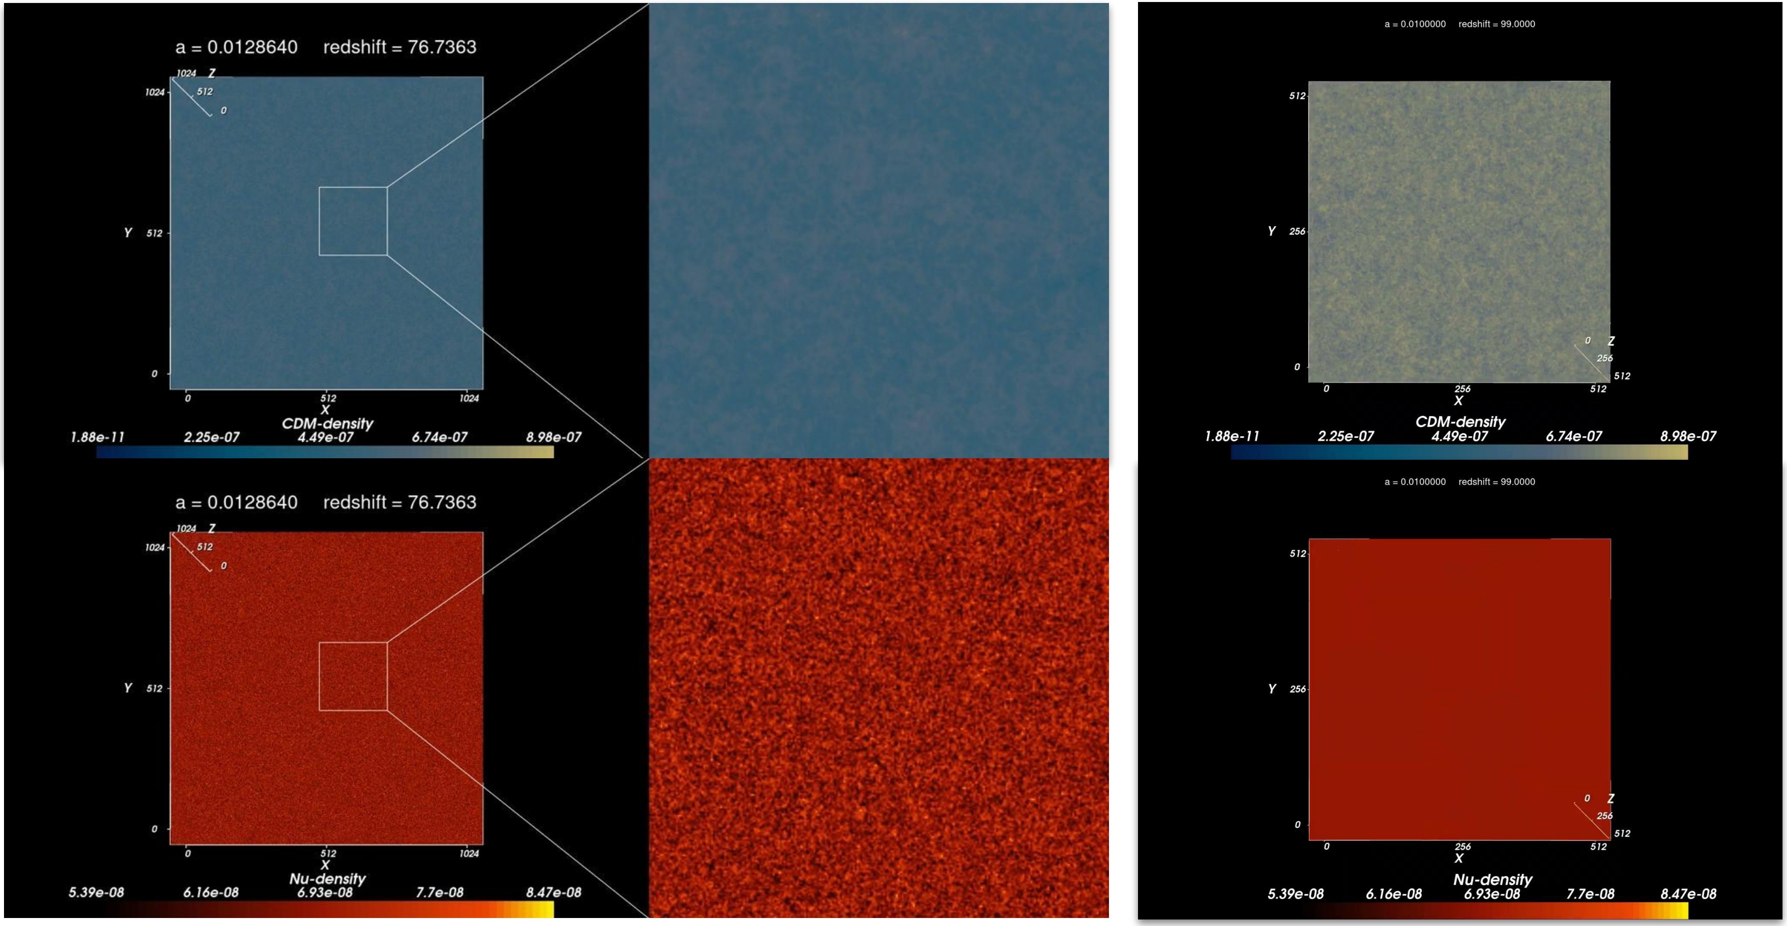

Credits: Tuccari, Sciacca, Vitello (Spoke1&Spoke3)

Analytical Covariances and future likelihood activities

(C. Carbone, S. Camera, M. Baldi, M. Calabrese, C. Giocoli, V. Cardone, D. Sciotti et al.)

1. Resources to run chains for 3x2pt (weak lensing, photo-gc and the cross) cosmological posteriors:

   Runs performed on KiDS infrastructures (NL), france computers (F) and UK facilities

2. Resources to compute the 3x2pt X CMB covariance matrices with FLASK Simulations: Marseille cluster available

3. **Future activities for DR1 KPs** (e.g. DR1-JC1, DR1-JC2, DR1-TH1) not yet covered by sufficient computational resources.
   a. Estimated resources to compute the 3x2pt covariance matrix terms and accounting for survey masks & GLASS Simulations: ~ 2M cpu-hr and up to 1TB of data
   b. Ongoing applications for cosmological parameter inference:
      i. EuroHPC call for Euclid DR1 KPs (coordinated by B. Joachimi)
      ii. DiRAC call for computational time in UK (coordinated by K. Koyama and Pedro Carrilho)

Computing resources in Euclid

- Computing resources for SGS (see table below for OULE3 dated 2020)
- No guaranteed computing resources for science (simulations, covariances, likelihood): only applying to public calls

PF	CPU [core/hours]	Memory [GB]	H/W	Comments
2PCF-GC	2000	150	Xeon E5-2680 v3 with clock speed 2.50 GHz. Each node had 24 physical cores. Nodes have 256GB RAM.	Estimated from Flagship Euclid Wide mock catalog
PK-GC	8	120	Workstation with 1 core with a 3.2GHz Xeon processor.	Estimated from Euclid-size mock catalogs
3PCF-GC	450000	300	2GHz processor.	Extrapolated from smaller catalogs
BK-3PCF	200	1100	Intel Xeon CPU E5-4627 v3 with , 2.60 GHz processor and 250 GB RAM	Extrapolated from smaller catalogs
CM-2PCF-GC	2000x3500	150x3500		Extrapolated from smaller catalog
CM-PK-GC	8x3500	300x3500		Extrapolated from smaller catalogs
VMSP-ID	Deep 1 Wide 6000	Deep 70GB Wide 2GB/Tile	1 core Xeon CPU E5-2630 v2 @ 2.60GHz 128 GB RAM	Full flagship extrapolated from 1 tile
SEL-ID				

Euclid
Science Ground Segment
data processing

INAF USCVIII General Assembly, Galzignano, 14–18 Oct 2024

Daniele Tavagnacco on behalf of SDC-IT
(A.Zacchei, M.Frailis, E.Romelli, S.Galeotta, T.Gasparetto, G.Maggio, D.Maino, F.Rizzo, T.Vassallo, R.Giusteri)

SGS role in Euclid

Define, Organize and Maintain the data analysis infrastructure (ST)

Manage and monitor the instrument and the sky survey progression (IOT)

Process raw data into final science products (OT)

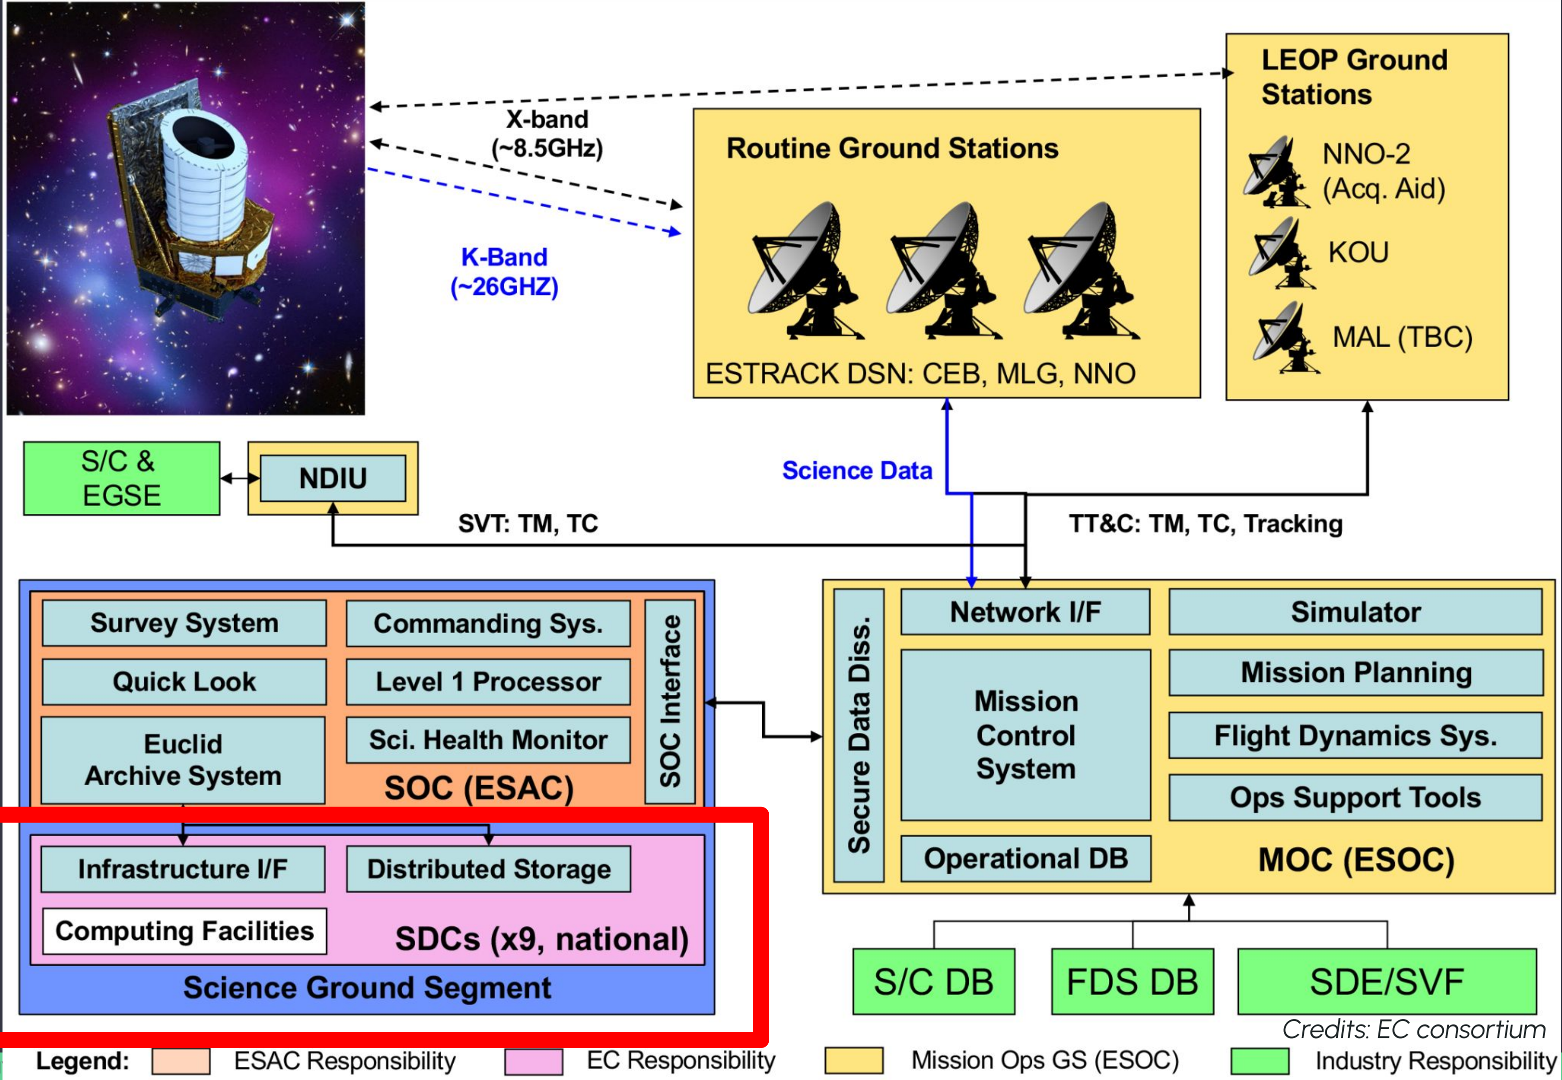

X-band
(~8.5GHz)
K-Band
(~26GHZ)
Routine Ground Stations
ESTRACK DSN: CEB, MLG, NNO
LEOP Ground Stations
NNO-2
(Acq. Aid)
KOU
MAL (TBC)
S/C & EGSE
NDIU
SVT: TM, TC
Science Data
TT&C: TM, TC, Tracking
Survey System
Commanding Sys.
Quick Look
Level 1 Processor
Euclid Archive System
Sci. Health Monitor
SOC (ESAC)
SOC Interface
Infrastructure I/F
Distributed Storage
Computing Facilities
SDCs (x9, national)
Science Ground Segment
Secure Data Diss.
Network I/F
Simulator
Mission Control System
Mission Planning
Flight Dynamics Sys.
Ops Support Tools
Operational DB
MOC (ESOC)
S/C DB
FDS DB
SDE/SVF
Credits: EC consortium
Legend:
ESAC Responsibility
EC Responsibility
Mission Ops GS (ESOC)
Industry Responsibility

- 1 July 2023
- Sun-Earth Lagrange point 2 (L2)
- 6yr nominal mission
- 1.20m primary mirror
- VISible instrument
- Near InfraRead/Spectro instrument

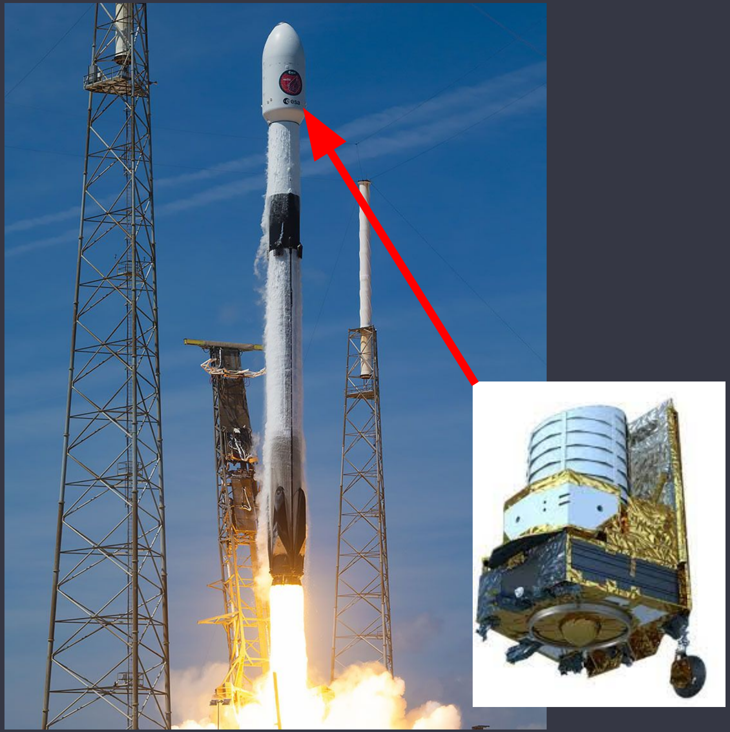

"Euclid is designed to explore the evolution of the dark Universe. It will make a 3D-map of the Universe [...] across more than a third of the sky"

[https://www.esa.int/Science_Exploration/Space_Science/Euclid_overview]

- **15000 deg^2** in 6yr
- FoV ~0.5deg → **+500k images**
- Process: **22 frames/day**

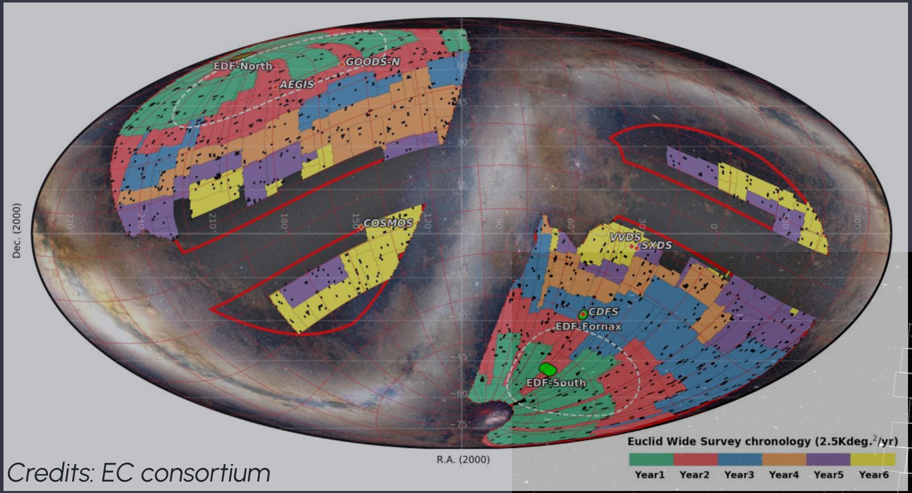

Credits: EC consortium

Euclid FoV

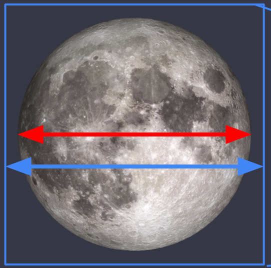

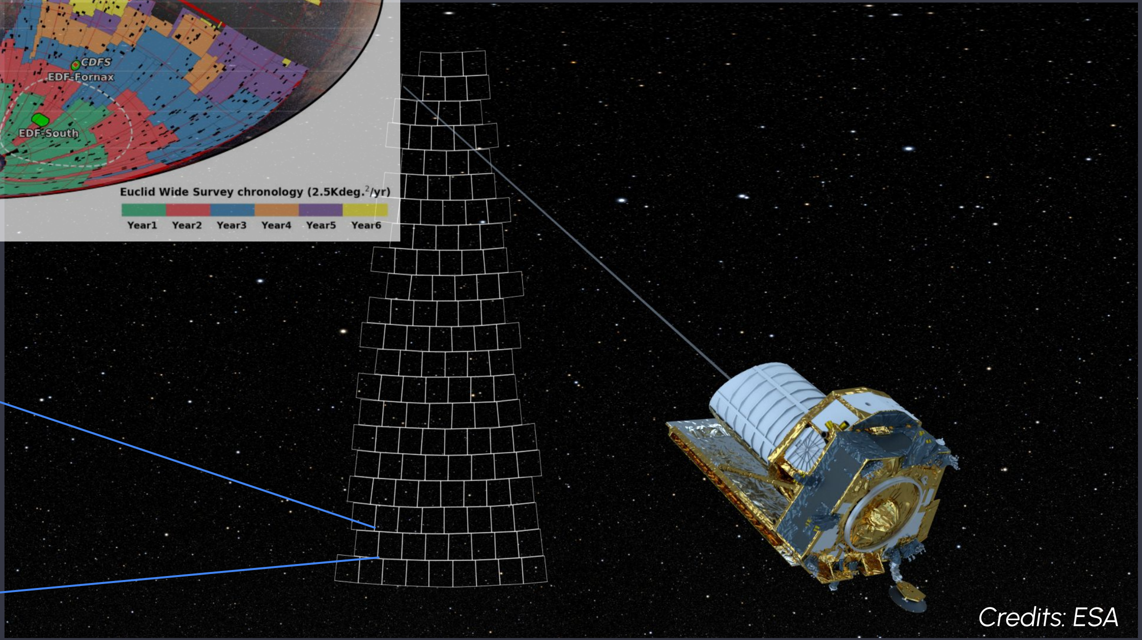

Credits: ESA

raw data images

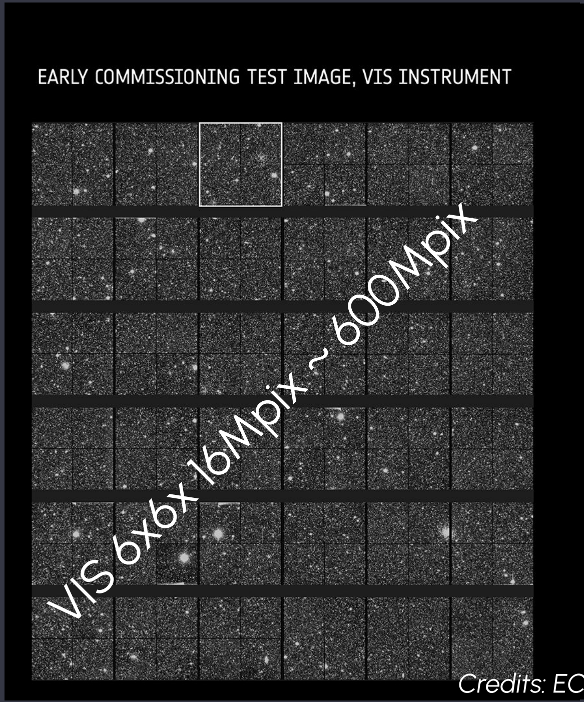

VIS 6x6x 16Mpix ~ 600Mpix

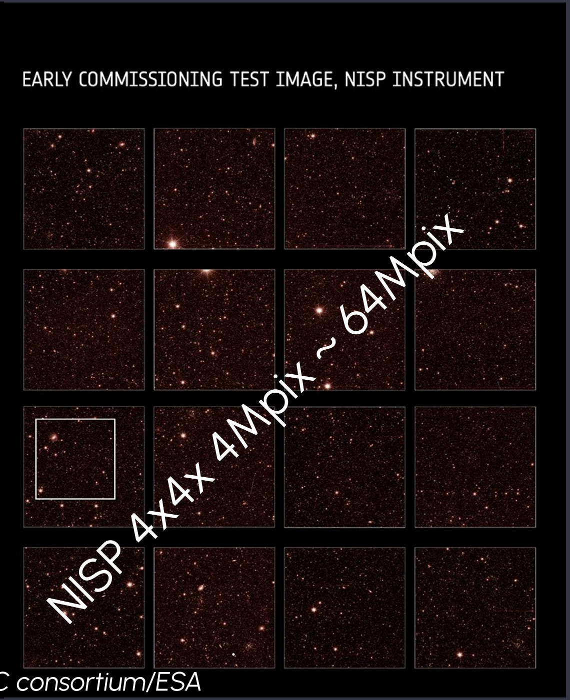

NISP 4x4x 4Mpix ~ 64Mpix

Credits: EC consortium/ESA

+ NISP spectro mode

@download: 100GB/day
@ground: 1TB/day
plus external data
total mission 20-30PB

constraints

Large **Data Volume**
(500+k raw images, external data)

Varied dataset on the processing flow
(images → catalogs → science data)

Continuously **evolving Software**
(data → instrument knowledge → better algorithms)

Euclid SGS as a distributed system

Science Data Centers
national HPC facilities (run)
develop. expertise (integration)

Organization Units
processing definition (algorithms)
science expertise (requirements)

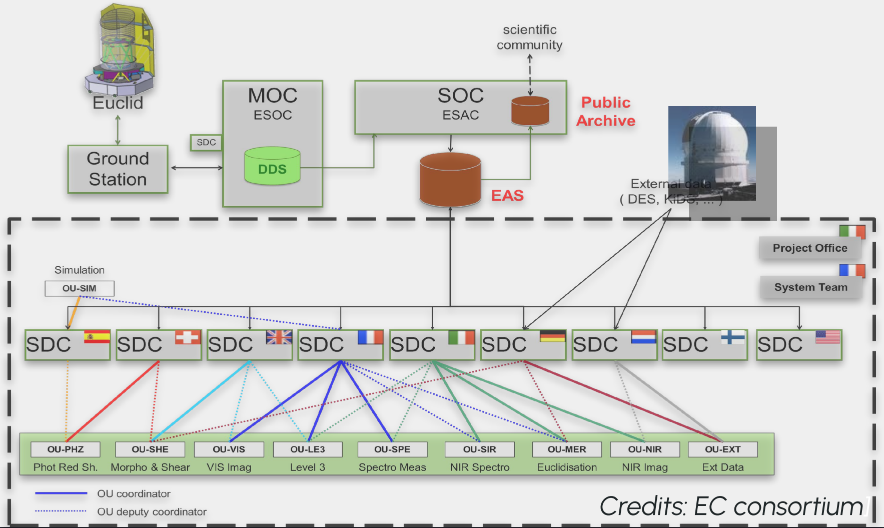

SDC + OU
SW pipelines development/optimization/test/integration/maintenance

… 9 SDCs… 9 OUs… ~1500 persons

SDC(s)

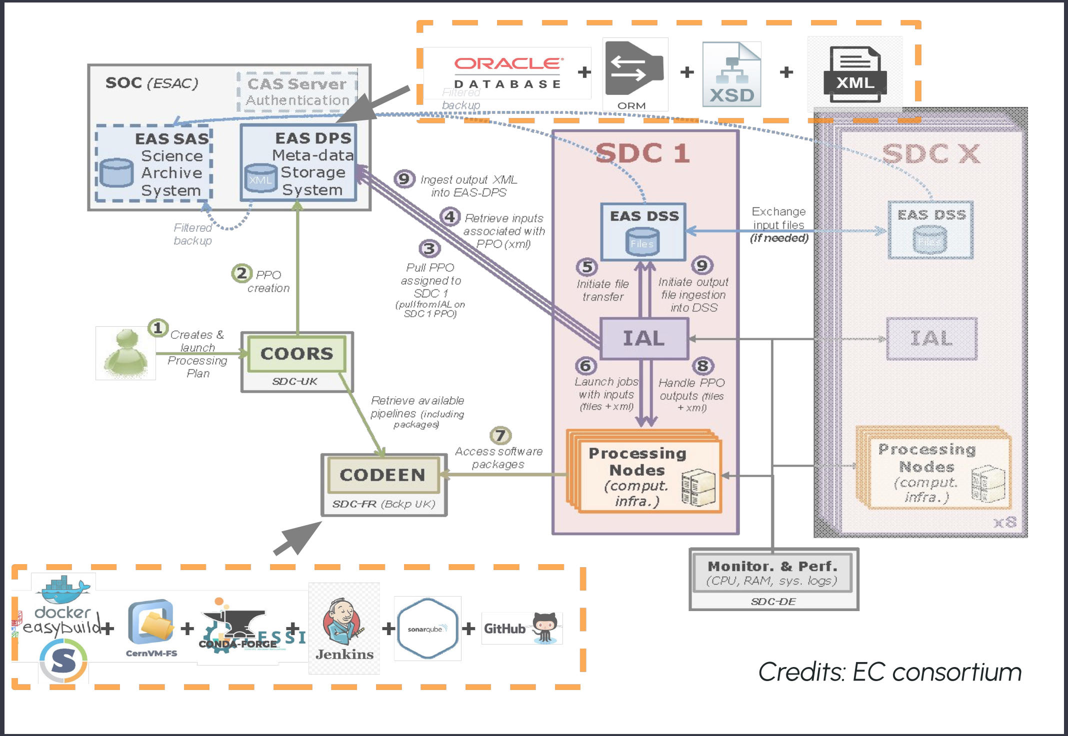

Credits: EC consortium

EAS - Distributed Storage System (DSS)

- Object Storage, SDC
- http data transfer -transparent
- mainly FITS and HDF5 formats

EAS - Data Processing System

- central metadata archive
- Object-to-relational mapping
- Custom query language (REST)
- data product oriented db
- Products as XML files

Euclid common data model

- EAS - Data and Metadata
- XSD to formalize products
- description of metadata *and* data
- evolution via Change Control Board

```xml
    <xs:complexType name="dpdSirCombinedSpectra">
        <xs:sequence>
            <xs:element name="Header" type="sys:genericHeader">
                <xs:annotation> <xs:documentation>Product Generic Header</xs:documentation> </xs:annotation>
            </xs:element>
            <xs:element name="Data"   type="sir:sirCombinedSpectra">
                <xs:annotation> <xs:documentation>Product data container</xs:documentation> </xs:annotation>
            </xs:element>
            <xs:element name="QualityFlags" type="dqc:sqfPlaceHolder" minOccurs="0">
                <xs:annotation> <xs:documentation>Product Quality Flags</xs:documentation> </xs:annotation>
            </xs:element>
            <xs:element name="Parameters"   type="ppr:genericKeyValueParameters" minOccurs="0">
                <xs:annotation> <xs:documentation>Product Optional Parameters</xs:documentation> </xs:annotation>
            </xs:element>
        </xs:sequence>
    </xs:complexType>
</xs:schema>
```

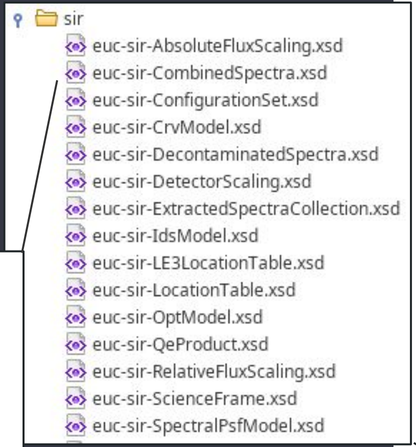

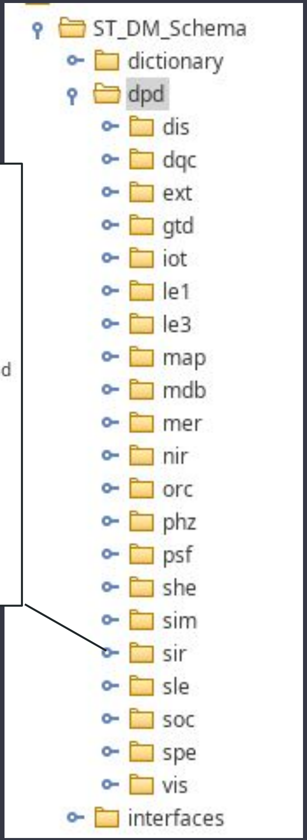

Infrastructure Abstraction Layer

national HPC facilities → possible different HW

- common meta-scheduler
- workflow manager
- common queuing systems

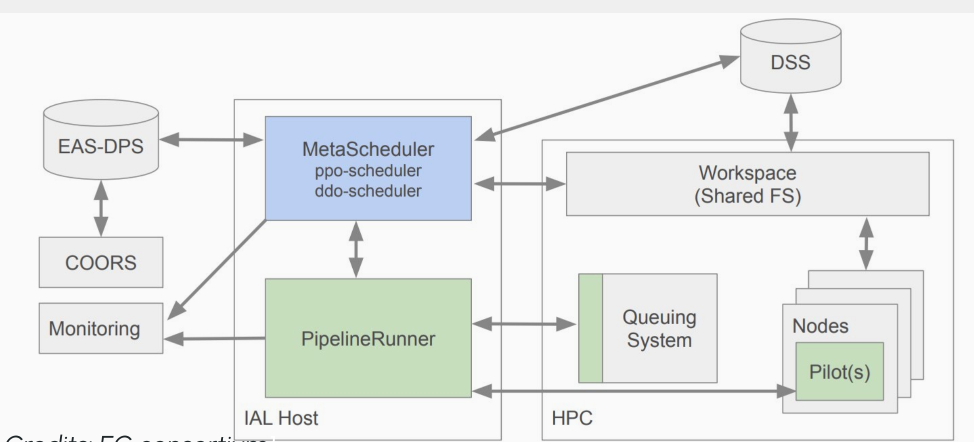

Credits: EC consortium

Workflow (pipelines) definition

- Python scripts specifying relations and requirements of each (PF)
- enforced PF I/O description
- stateless processing

Processing: payload job as common "pilot jobs", profiling of jobs

Infrastructure Abstraction Layer

national HPC facilities → possible different HW

- common meta-scheduler
- workflow manager
- common queuing systems

Workflow (pipelines) definition

- Python scripts specifying relations and requirements of each (PF)
- enforced PF I/O description
- stateless processing

Processing as payload job on common "pilot jobs", profiling of jobs

Euclid processing workflow

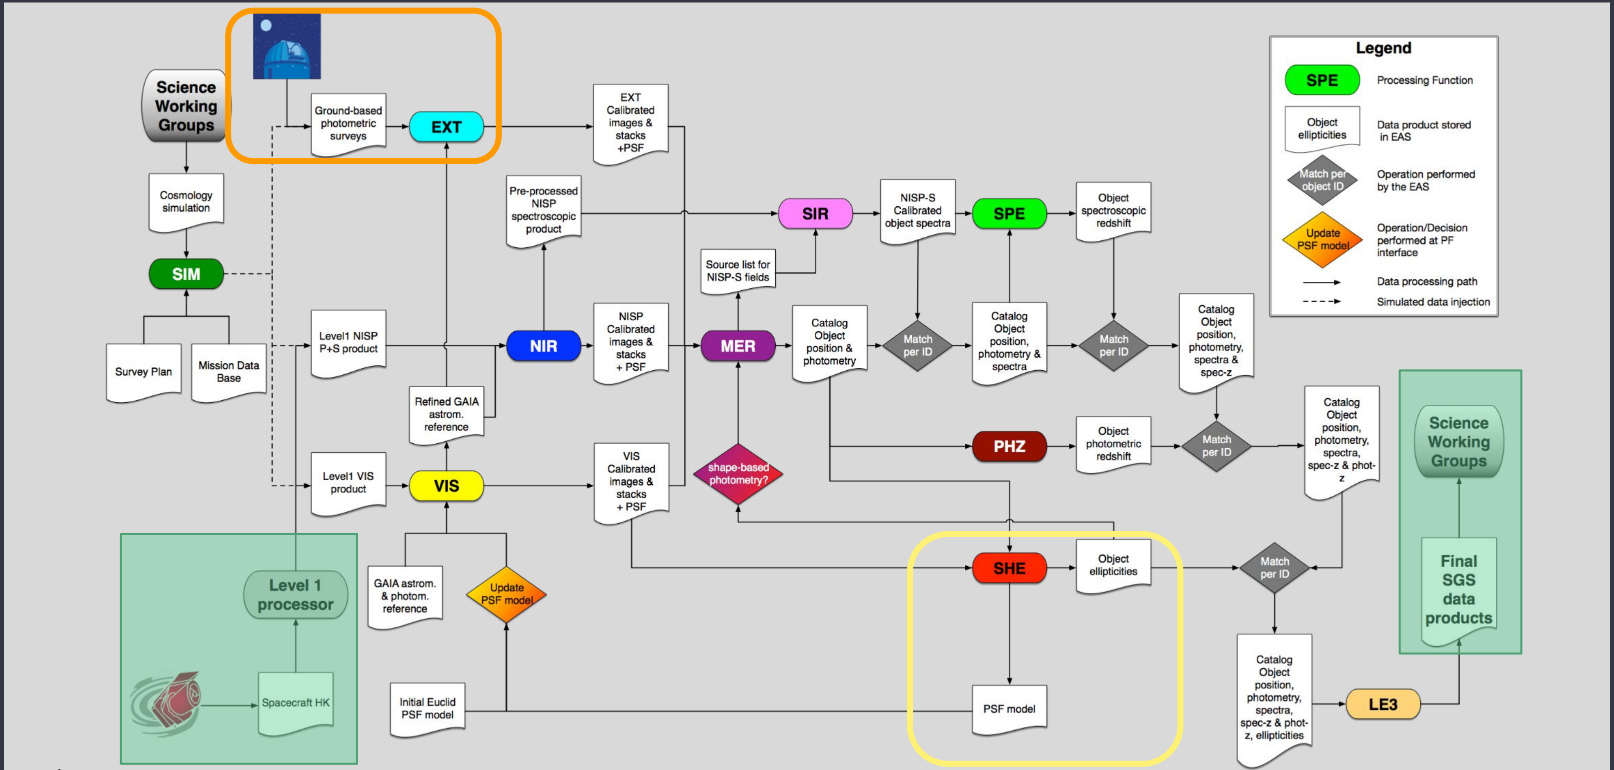

Credits: EC consortium

Software lifecycle in Euclid SGS

Data processing organized in steps (Processing Functions)

- several software projects
- different life cycle timings

Different requirements/granularity

- minimal set of infrastructure req.
- common environment

Multiple datacenters (SDCs)

- synchronization of SW
- different infrastructures

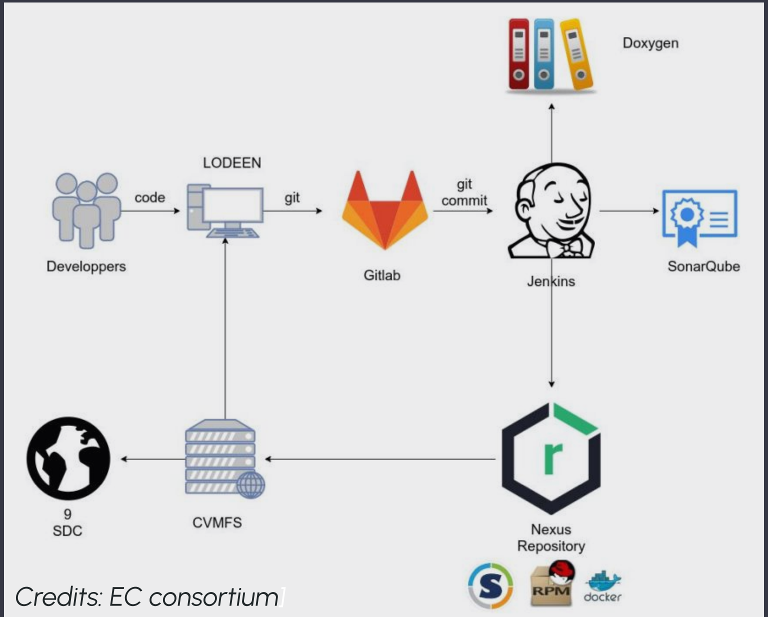

Credits: EC consortium

Gitflow workflow + Jenkins CI/CD + CVMFS packages distribution

SDC-IT* computing infrastructure

development "SDC-IT-DEV" ➔ **OATs**

- 15 nodes: 40 phys cores (no HT), 256GB RAM each node
- 1 node: 32 cores, 180GB RAM, 4x NVIDIA V100 16GB
- 650 TB of parallel storage (BeeGFS)
- InfiniBand

integration/testing "SDC-IT-INT" ➔ **ALTEC** (Turin)

- 8 nodes: 48 cores (24 phys), 256 GB RAM
- 4 nodes: 64 cores (32 phys), 1 TB RAM
- 730TB fast storage (Lustre) and Tier 2 storage shared with OPS

production "SDC-IT-OPS" ➔ **ALTEC** (Turin)

- 12 nodes: 128 cores (64 phys), 1 TB RAM
- 9 nodes: 192 cores (96 phys), 768 GB RAM
- 2 nodes: 192 cores (96 phys); 1.5 TB RAM
- 1.4 PB fast storage (Lustre) + 4.4 PB Tier 2 general storage + 6 PB Tier 3 Tape Lib.
- InfiniBand

...few takeaway notes

within a project, **SGS is one of the longest-lived elements**

technologies evolve during SGS lifecycle

SGS design/implementation and project design/development are entangled

SGS activities require various expertise to cover all operational fields

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