Use Cases from Science Projects: CTA

Matthias Fuessling, Nadine Neyroud, Frederic Gillardo

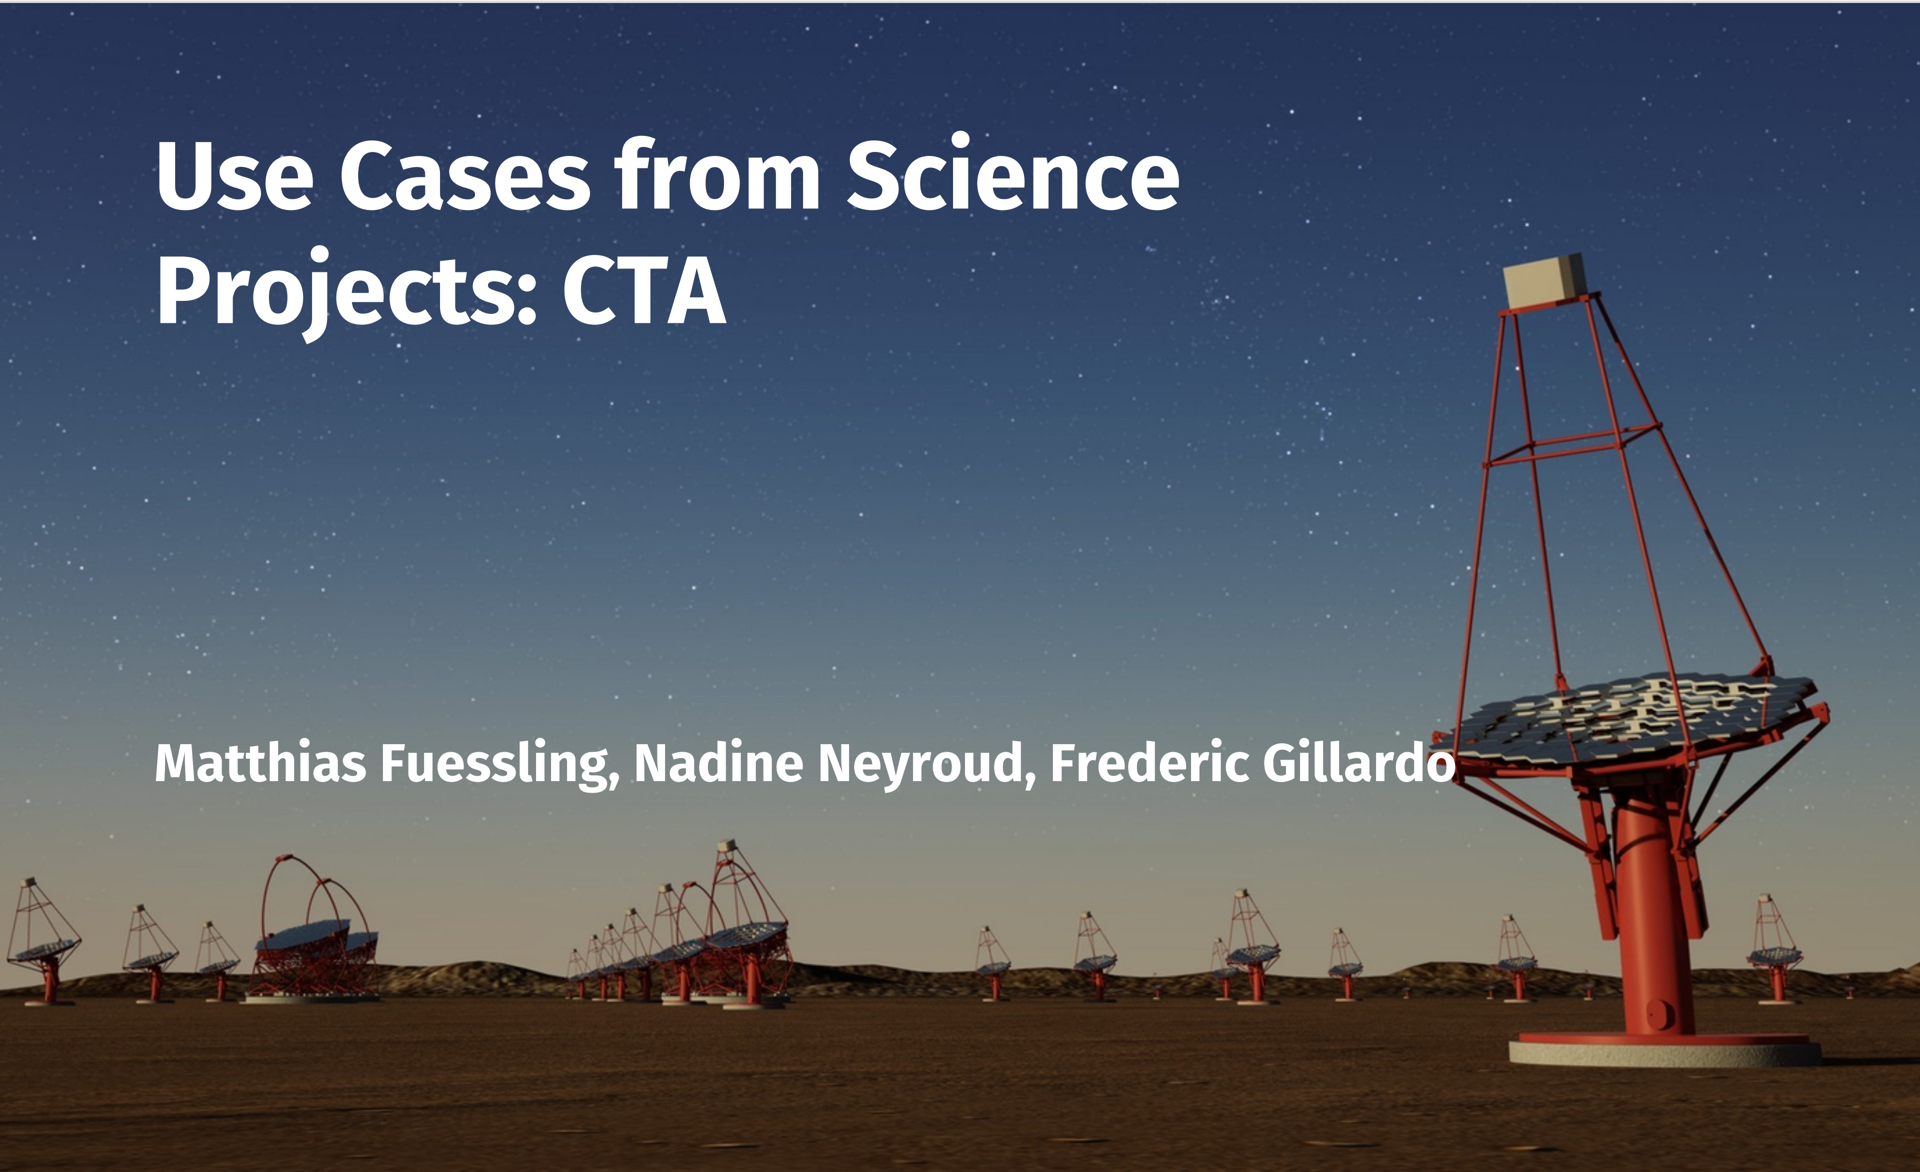

Project Overview

The Cherenkov Telescope Array

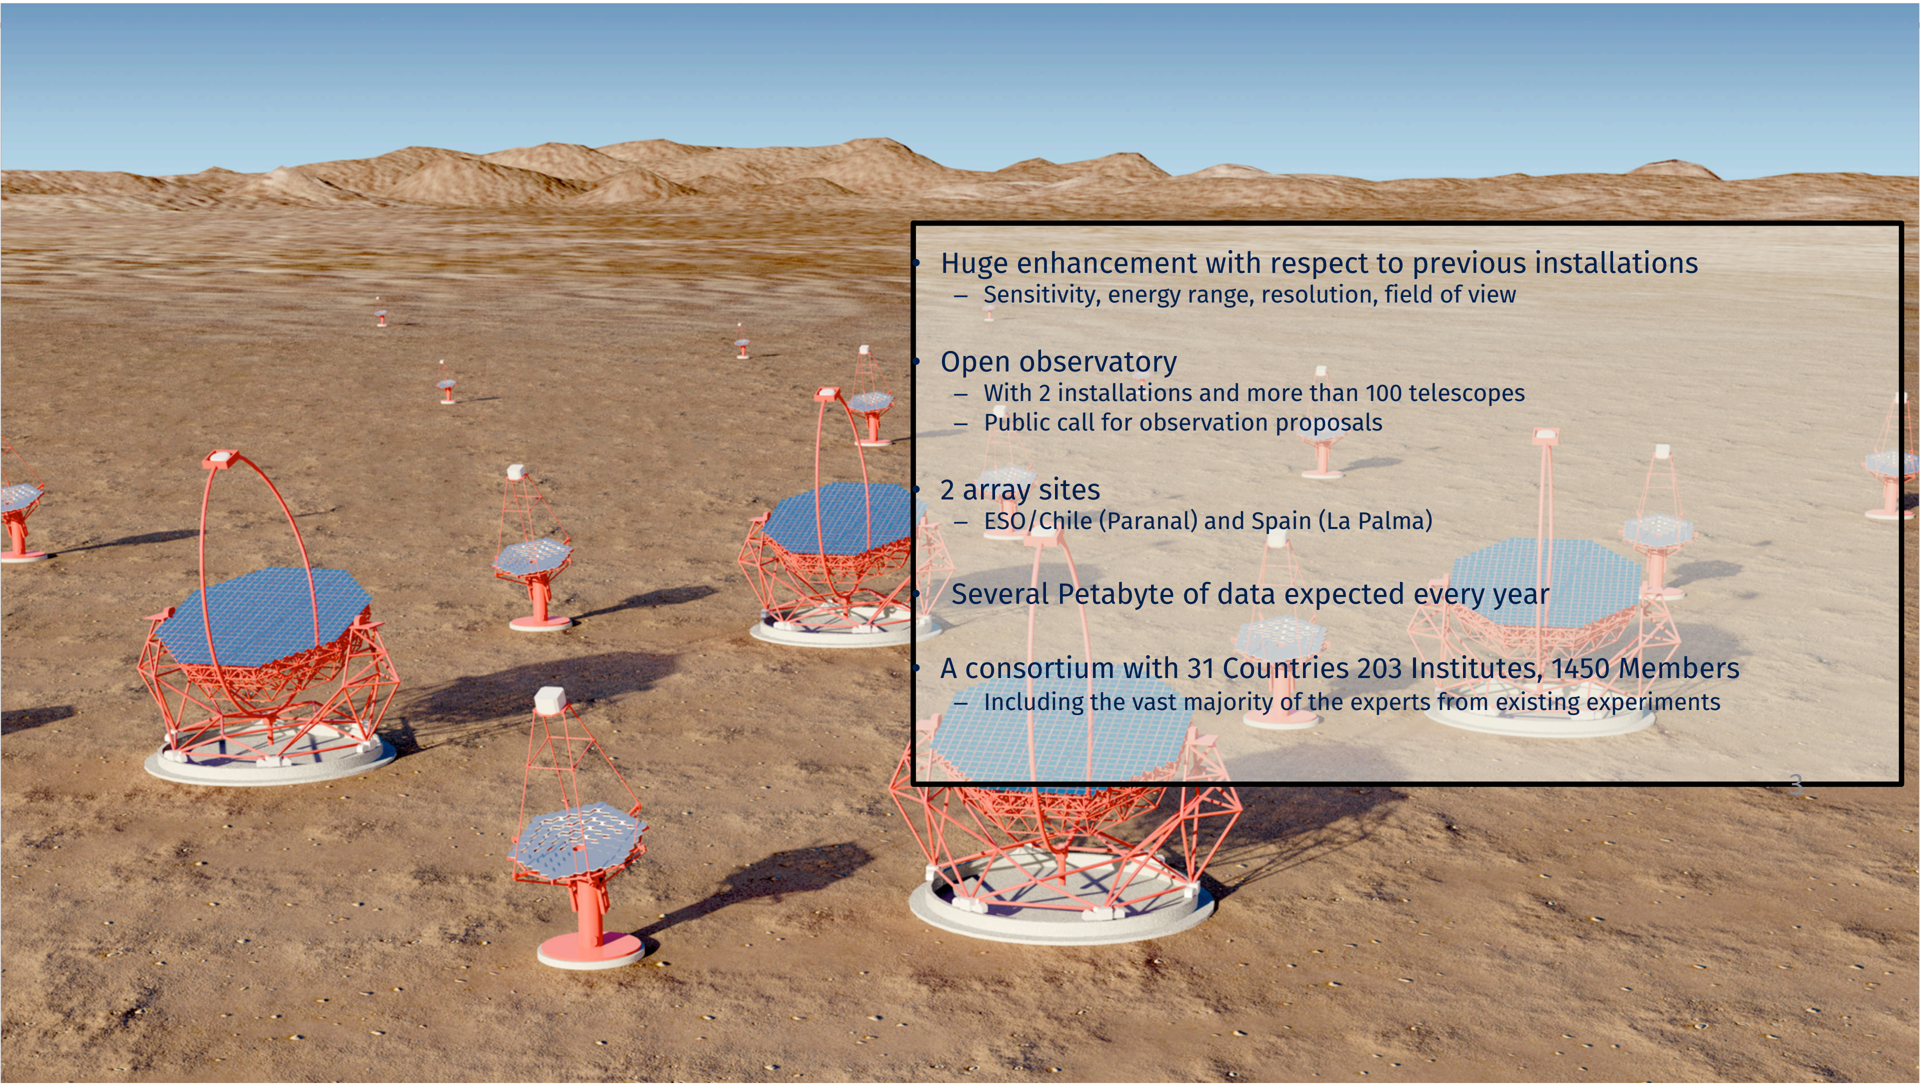

- Huge enhancement with respect to previous installations
  - Sensitivity, energy range, resolution, field of view
- Open observatory
  - With 2 installations and more than 100 telescopes
  - Public call for observation proposals
- 2 array sites
  - ESO/Chile (Paranal) and Spain (La Palma)
- Several Petabyte of data expected every year
- A consortium with 31 Countries 203 Institutes, 1450 Members
  - Including the vast majority of the experts from existing experiments

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Detection principle

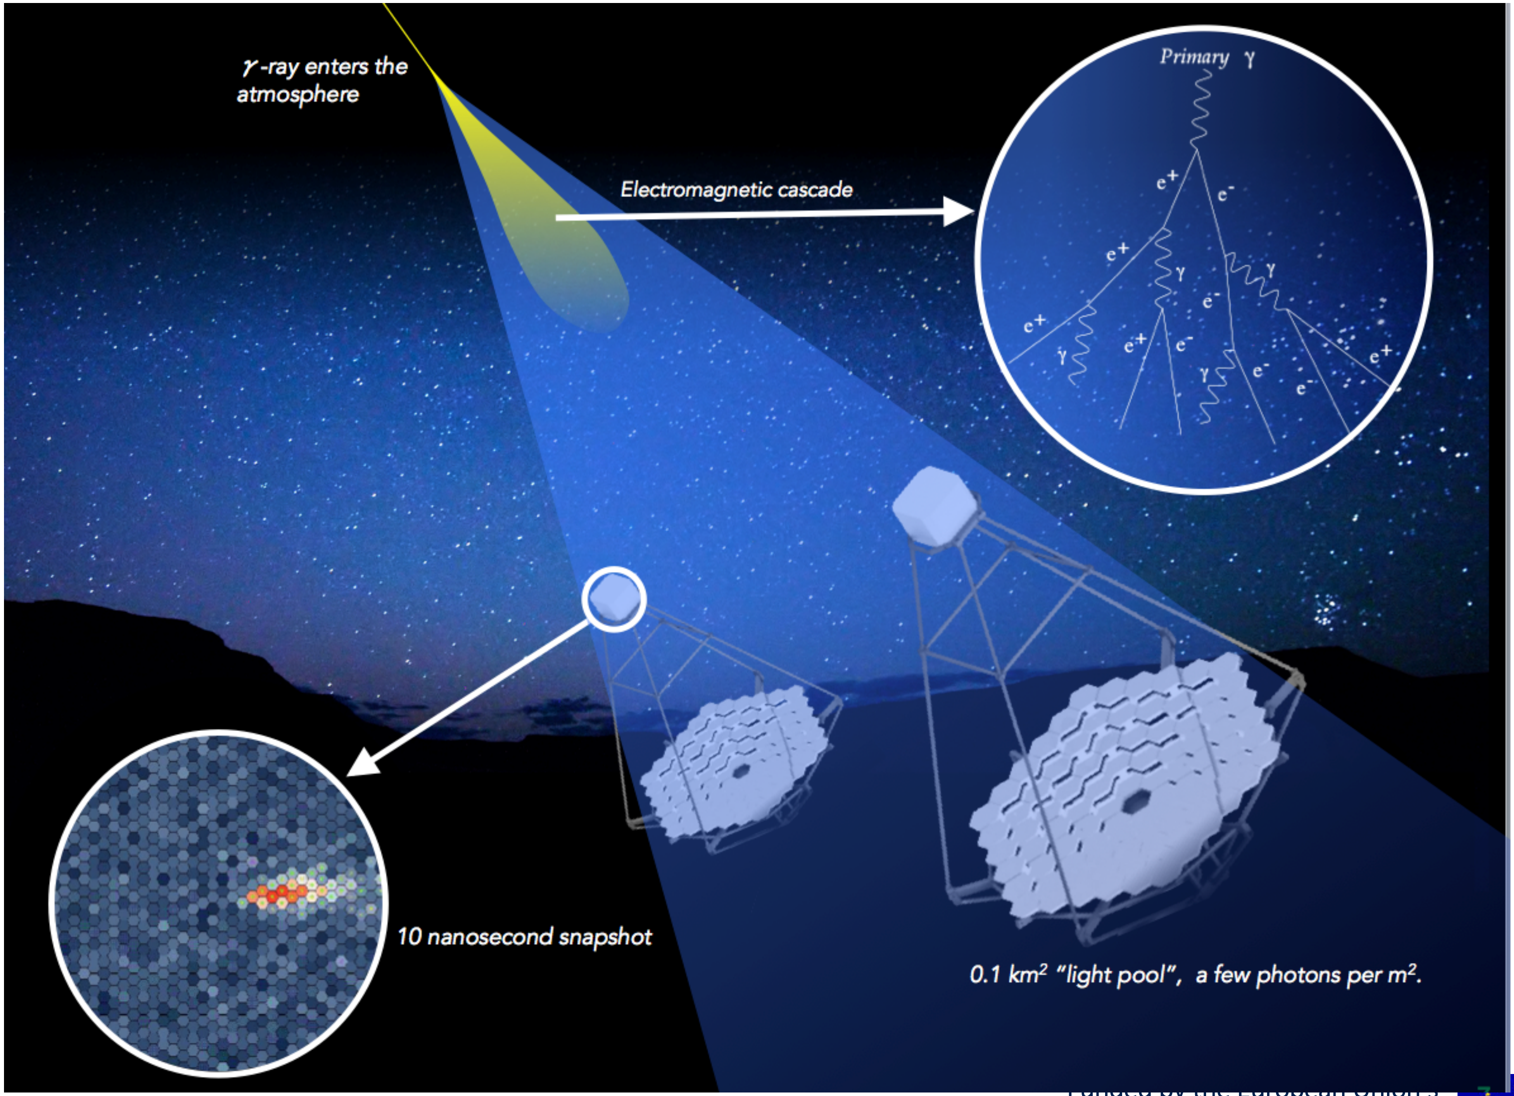

Funded by the European Union's
Horizon 2020 - Grant N° 824064

CTA telescope prototypes

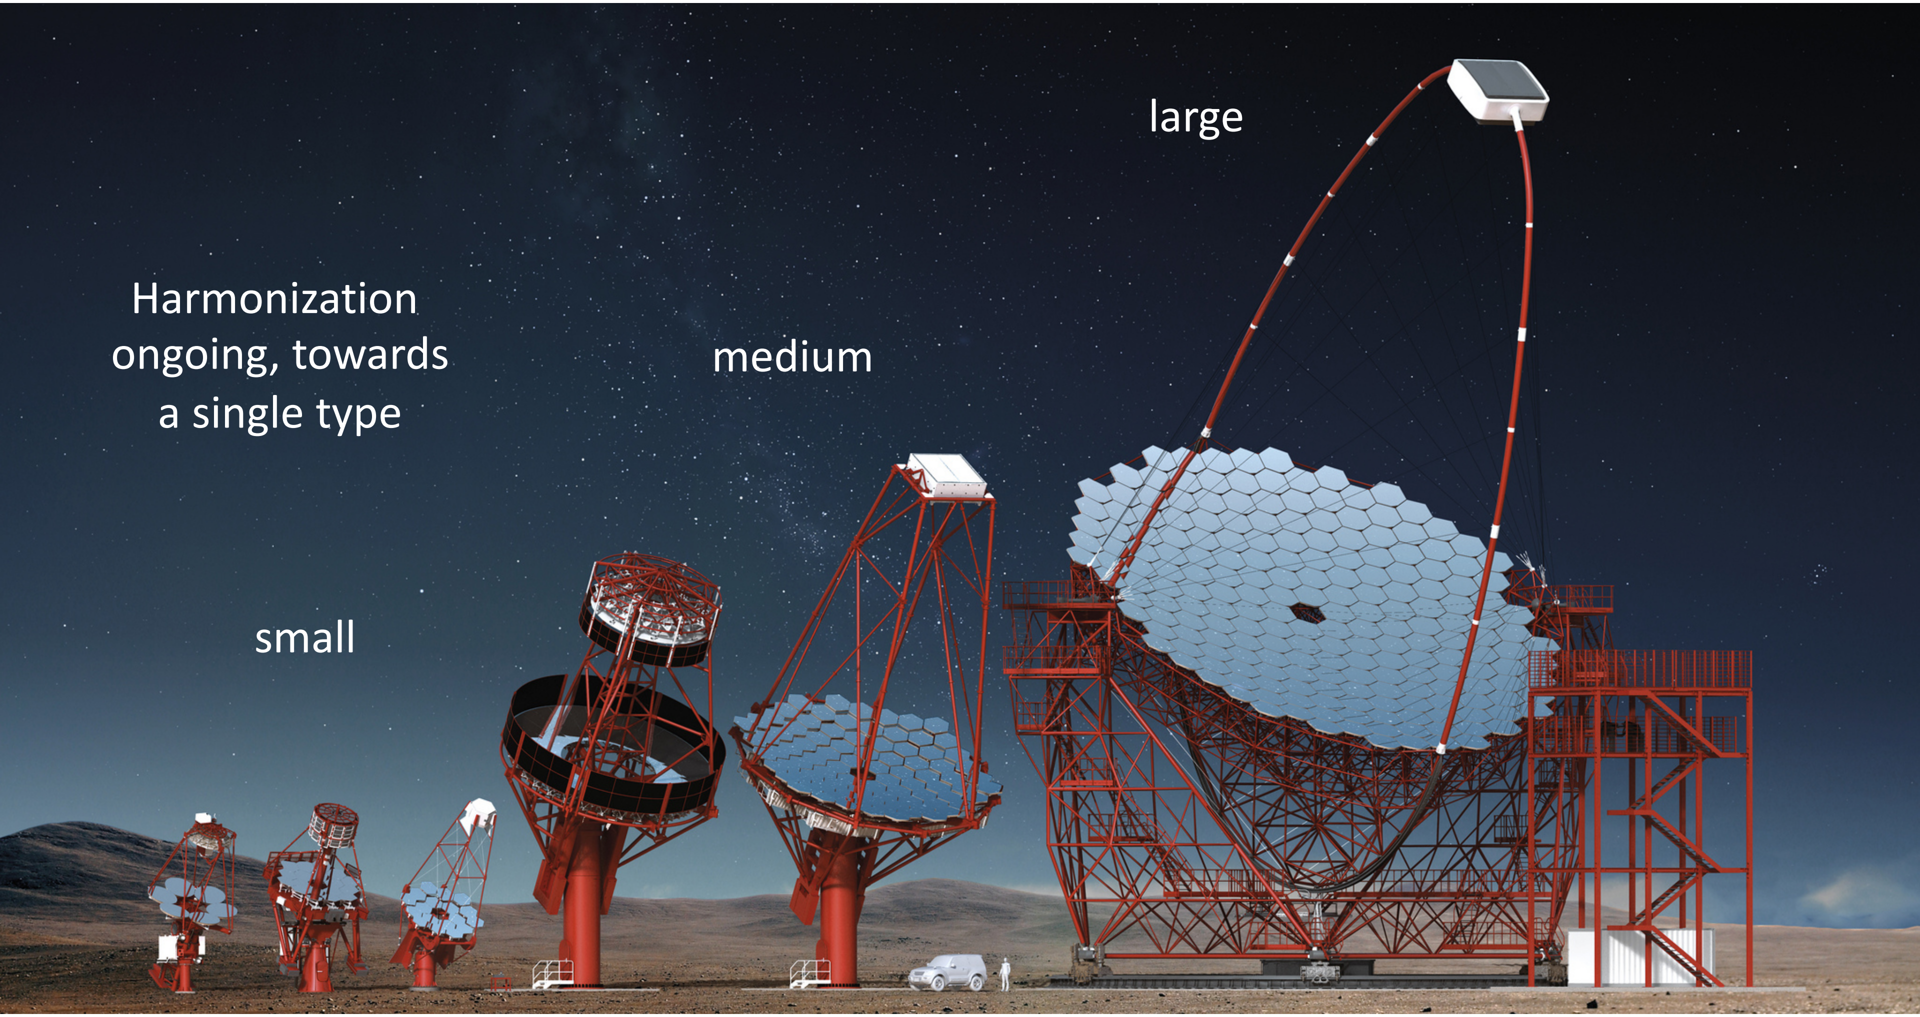

Funded by the European Union's
Horizon 2020 - Grant N° 824064

- **2005** Scientists propose CTA as Next Generation Facility for Ground-based Gamma-Ray Astronomy
- **2007** First Community Meeting on CTA
- **2007** Definition of the CTA Design Study
- **2008** CTA Consortium Established
- **2008** CTA included in the Roadmap of the European Strategy Forum on Research Infrastructures (ESFRI)
- **2010 – 2014** CTA Preparatory Phase Supported by FP7 Grant from the European Commission
- **2010** Project Office Established in Heidelberg
- **2011** Publication of the "Design Concepts for the Cherenkov Telescope Array"
- **2012** Funding Agencies Sign Declaration of Intent – Beginning of CTA Pre-Construction Phase
- **2015** Consortium Grows to more than 1,500 Members from 32 Countries
- **06 2016** HQ and Data Management Centre Site Decisions
- **2016** North Site Hosting Agreements Finalized
- **11 2016** Memorandum of Understanding for Construction
- **03 2017** CTA Offices Open in Bologna
- **2018** South Site Hosting Agreement Finalized
- **2020** European Research Infrastructure Consortium (ERIC) Established
- **2020** First Pre-Production Telescopes on Site
- **2022** Beginning of Observatory Operations
- **2025** Construction Project Completion

CTA sites: arrays, headquarters, data center

CTA at Paranal & La Palma

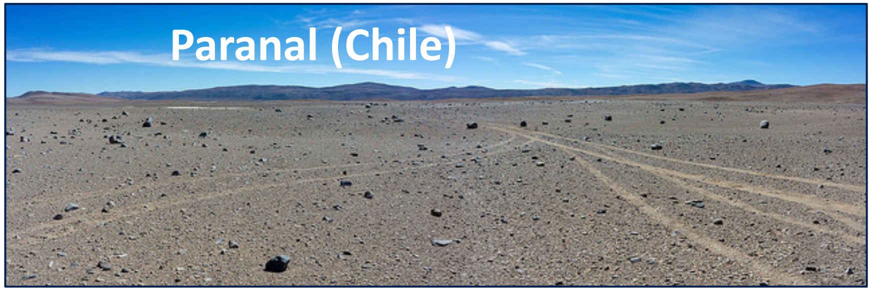

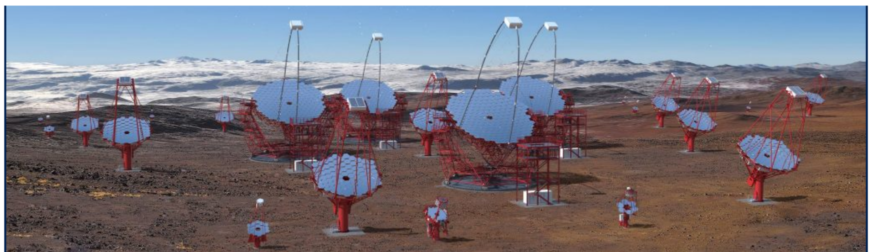

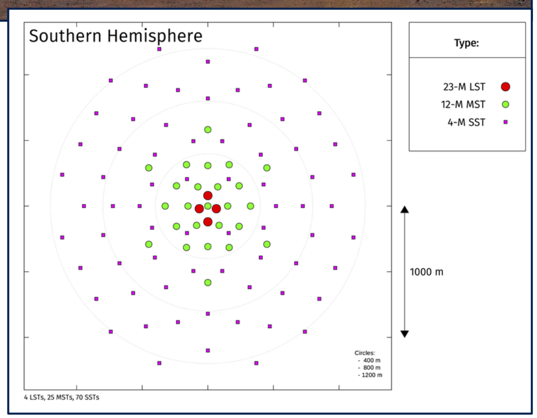

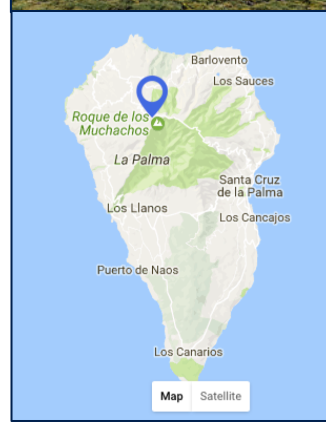

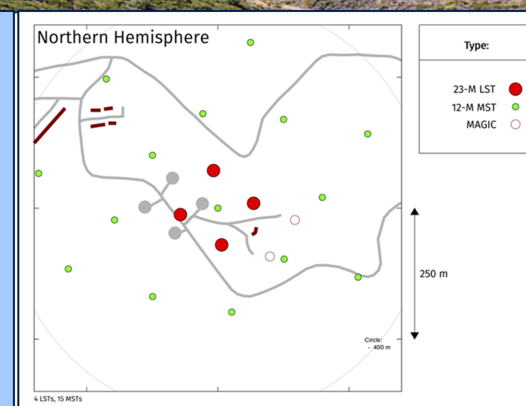

The CTA Observatory (CTAO)

- In 2014, the CTA Observatory GmbH was founded as interim legal entity, located at Heidelberg, under German law
  - To prepare the CTA implementation (select and prepare two array sites + Science Data Management Centre)
- The final legal entity for full construction, a *European Research Infrastructure Consortium* (ERIC), is being set up under European Union law (early 2020?)
- Bologna (Italy) selected as HQ; Project Office is moving to Bologna and is steadily growing
- The *Science Data Management Centre* (SDMC) will be built up at DESY in Berlin-Zeuthen (Germany), in a new building

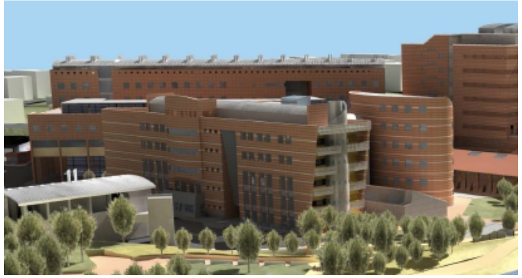

Artist's conception

Horizon 2020 - Grant N° 824064

CTA as an Observatory

- A Guest Observer Facility
  - For the **first time** in this waveband
    - Existing instruments are run as experiments
  - Annual cycles, TAC ranking, long-term schedule
  - Proposal preparation support, tracking, helpdesk +
  - Public science data archive
    - After proprietary period
- Two Telescope Arrays – one Observatory
  - Inter-site coordination
  - Uniform approach to science operations

Funded by the European Union's Horizon 2020 - Grant N° 824064

CTA Observatory

- Users will receive their data fully calibrated in FITS format, and be provided analysis tools
- After a one-year proprietary period, data are open
- During 1st decade, available observation time split roughly evenly between Key Science Projects (in particular surveys) and open time

- CTA is a Software instrument
  - Software plays a critical role in all steps of the Observatory
- Main Challenges
  - Sub-array operation, wide field of view, instrument response generation, background modelling, rapid alert generation and response, data volume, science operations during construction

Funded by the European Union's
Horizon 2020 - Grant N° 824064

CTA Science Operations Overview

Science User Support System (SUSS)

Array Control and Data Acquistion (ACADA)

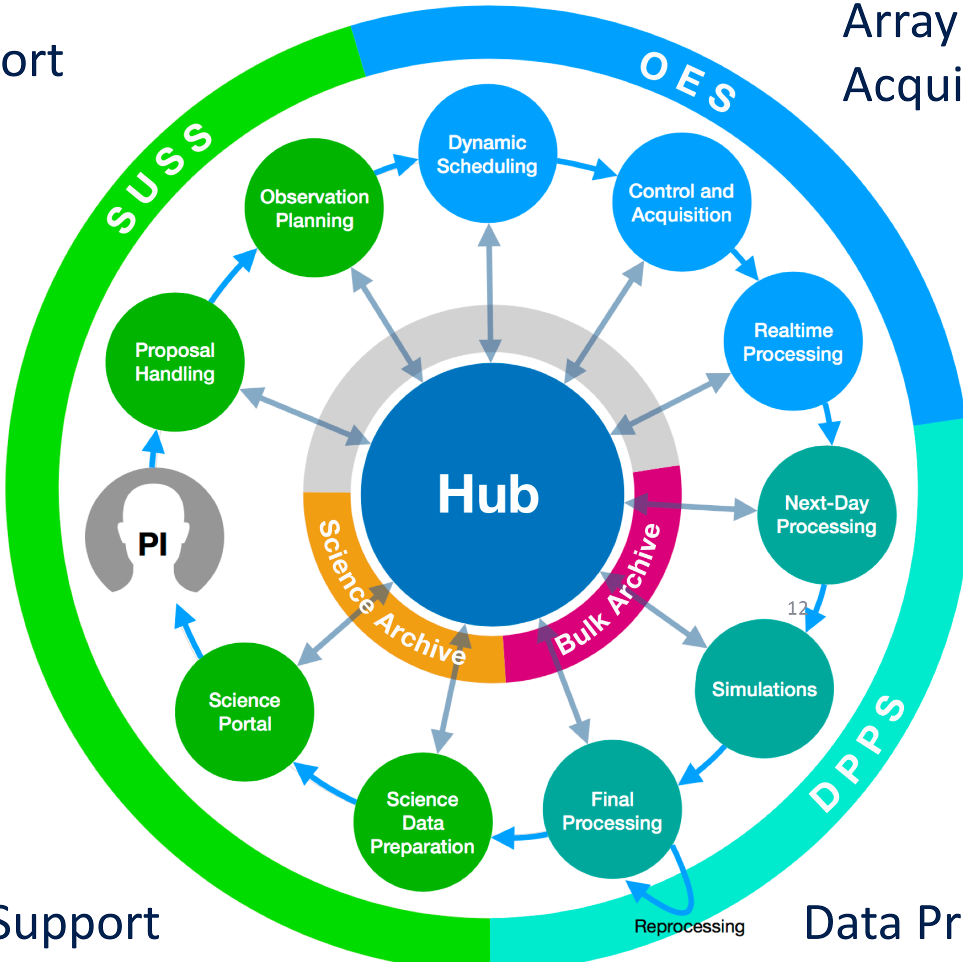

Science Operations Support (Hub)

Data Processing and Preservation System (DPPS)

Funded by the European Union's Horizon 2020 - Grant N° 824064

Overall Data Flow

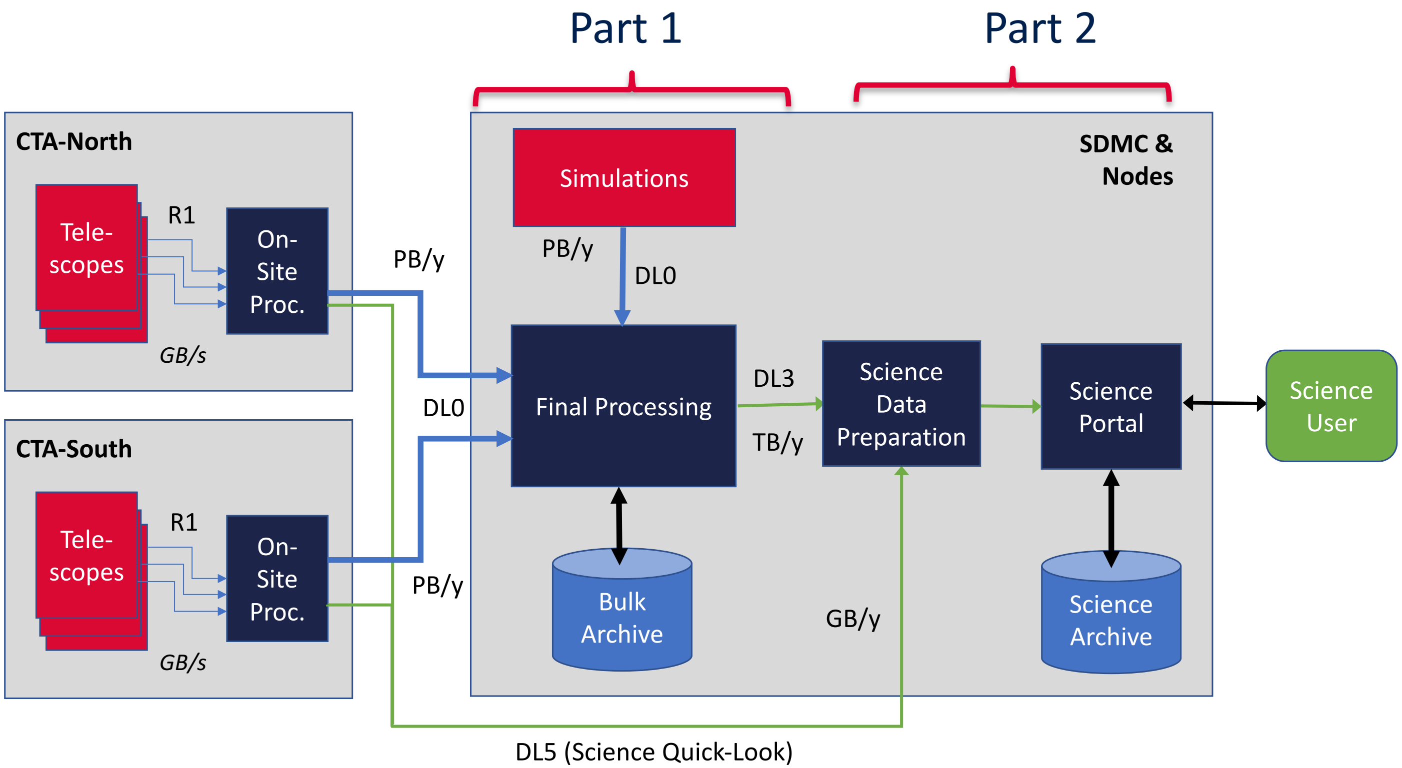

- On-site computing
  - Next-day data processing for quick-look and science alerts
  - On-site buffer and data transfer
- Off-site computing
  - Simulations and final processing (DL0 -> DL3)
  - Bulk Archive (DL0-DL3)
  - Science data preparation (DL3 -> DL5)
  - Science Archive (DL3-DL5)
  - Open access through Science Portal

Data Levels

Data Level	Short Name	Definition
R0	Raw Internal	On-site streamed raw data, not normally preserved long-term in this form. R0 content and format is internal to each device / controllable system, such as raw data transmitted from the physical device / system to its respective server in the on-site Data Centre.
R1	Raw Common	On-site stream raw data meeting common standards, transmitted on-site from a Camera or other on-site system to the OES. This is the first level of data seen by the OES, that will typically need some pre-processing from the R0 data format. Exceptionally, some R1 data may be stored for engineering purposes.
DL0	Raw Archived	All archival data from the data acquisition hardware/software, transmitted from the OES to the DPPS. This is the lowest level of data that are intended for long-term storage in the bulk archive. This includes both camera event data and technical data from other sub-systems, such as non-camera devices or software.
DL1	Processed	Processed DL0 data that may include telescope-level (TEL) data and parameters derived from them. Typical contents include calibrated image charge, Hillas parameters, and a usable telescope pattern. DL1 data is not normally stored long-term.
DL2	Recon-structed	Reconstructed shower parameters such as energy, direction, particle ID, and related signal discrimination parameters. Does not include telescope-level (TEL) information. For each event this information may be repeated for multiple reconstruction and discrimination methods. DL2 data is not normally stored long-term.
DL3	Reduced	Sets of selected events with a single final set of reconstruction and discrimination parameters, along with associated instrumental response characterizations and any technical data needed for science analysis.
DL4	Binned	Data product produced by binning of DL3 data, including data cubes and maps which are suitable for combination/summation to produce DL5 products.
DL5	Science	Data product produced by combination of DL4 products an extraction target specific region(s) of interest. Includes for example light-curves and spectra, along with associated data such as source models and fit results.
DL6	High-Level	High-level or legacy observatory data, such as survey maps and source catalogues.

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Data Levels

Data Level	Short Name	Definition
R0	Raw Internal	On-site to each respect nternal n to its
R1	Raw Common	On-site system proces -site rposes.
DL0	Raw Archived	All arch the low camera . This is h ware.
DL1	Processed	Proces Typical data is n. n. DL1
DL2	Recon-structed	Recons parame repeate ination be ng-term.
DL3	Reduced	Sets of associa with lysis.
DL4	Binned	Data pr combin for
DL5	Science	Data pr Include results. terest. and fit
DL6	High-Level	High-level or legacy observatory data, such as survey maps and source catalogues.

Image/Data Cube
Per Telescope

Mixed Tel/Event (Temporary)

Event List

Binned

Funded by the European Union's
Horizon 2020 - Grant N° 824064

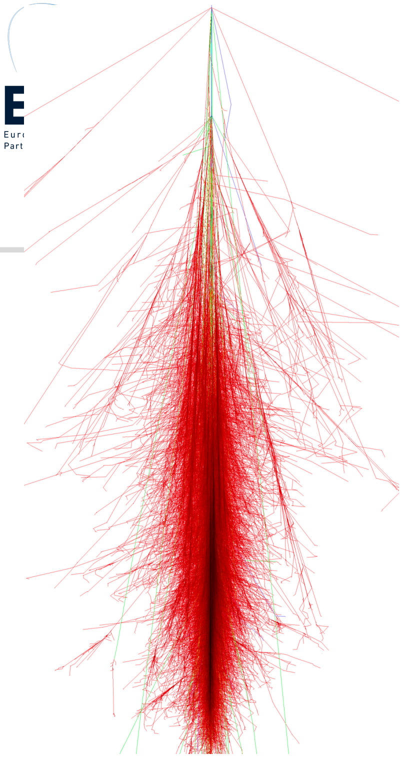

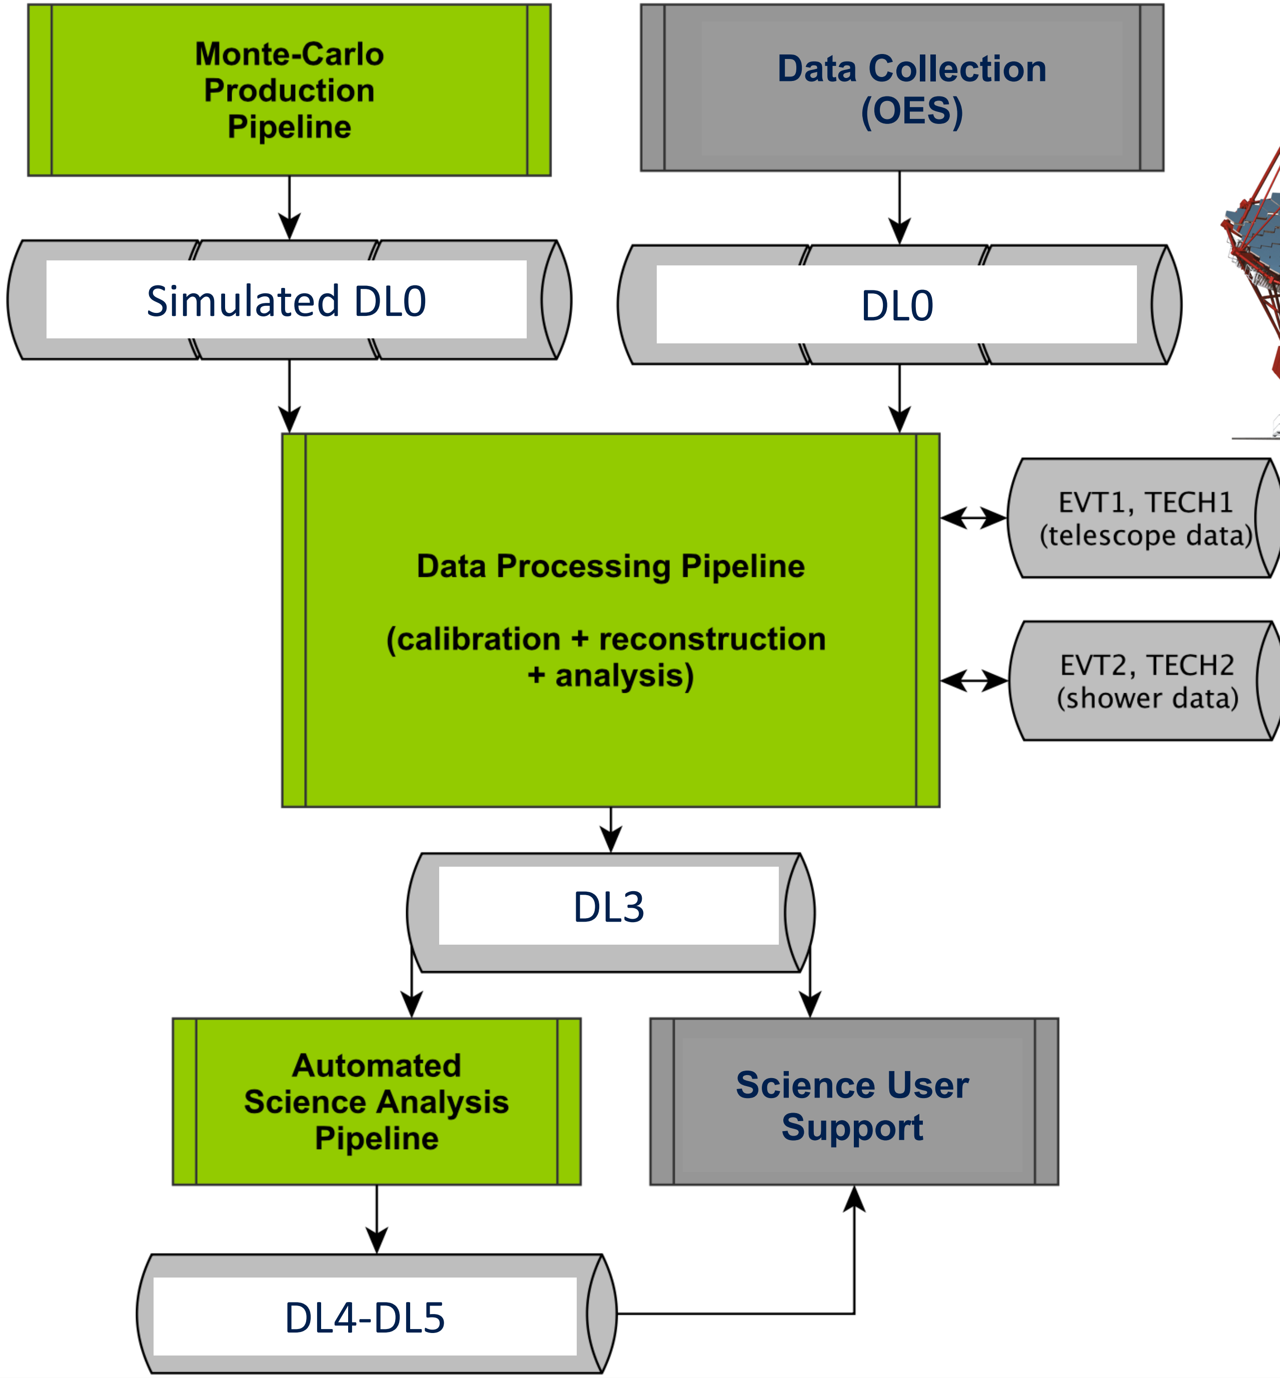

Monte-Carlo Production Pipeline
Data Collection (OES)
Simulated DL0
DL0
Data Processing Pipeline
(calibration + reconstruction + analysis)
EVT1, TECH1 (telescope data)
EVT2, TECH2 (shower data)
DL3
Automated Science Analysis Pipeline
Science User Support
DL4-DL5

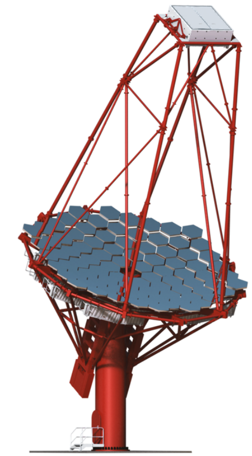

ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

Data Categories

cta

- **A** (real-time) Data products produced and distributed rapidly for science alert generation and rapid data quality evaluation. These generally have the lowest precision and highest systematic uncertainties arising from basic calibration and simplified analysis techniques. — minute
- **B** (next-day) Data products produced and distributed after some off-line processing on-site, by the next observation day. These have somewhat better precision and lower systematics than category-A products, but still may use simplified analysis or calibration techniques appropriate for relatively fast science alerts and proposal monitoring. — day
- **C** (final) Data products produced by the full high-quality data processing chain, off-site in CTA data centres, with a delay of ~weeks from data taking. These use the best calibration and algorithms, providing precision and systematics meeting or exceeding CTA requirements, and thus are the products intended for final analysis and publication of results. — month

DL1-DL5 produced for all categories

Part 1: From Observations to Reduced Data

From On-site Data to Science-Ready Data

- Transfer raw data (DL0) from site
  - Bandwidth sufficient after pixel selection
- Preservation *in bulk data archive*
  - Bulk data archive ‘ingestion’ – backup made
- Pipeline analysis
  - Calibration and final image/cube processing
  - Combination of telescope information
  - Event reconstruction
- Enrich data with appropriate metadata (incl. provenance)
- Derivation of matching IRFs
  - Selection from IRF grid from sims, or tailored sim for this observation
- Delivery science-ready data products (DL3) to science archive
  - of Event List + IRFs + binned products for this observation

- is independent on the specific proposal or observation and should be as automated as possible
- is (almost) continuous process (process data when observations come in)
- is organized by CTAO staff only, not accessible by science users

Funded by the European Union’s
Horizon 2020 - Grant N° 824064

Example Workflow

DL0:
- one folder per telescope per observation

Simulations / Monitoring
- used for reconstruction and identification of good time intervals

DL3:
- one file per good time interval in an observation

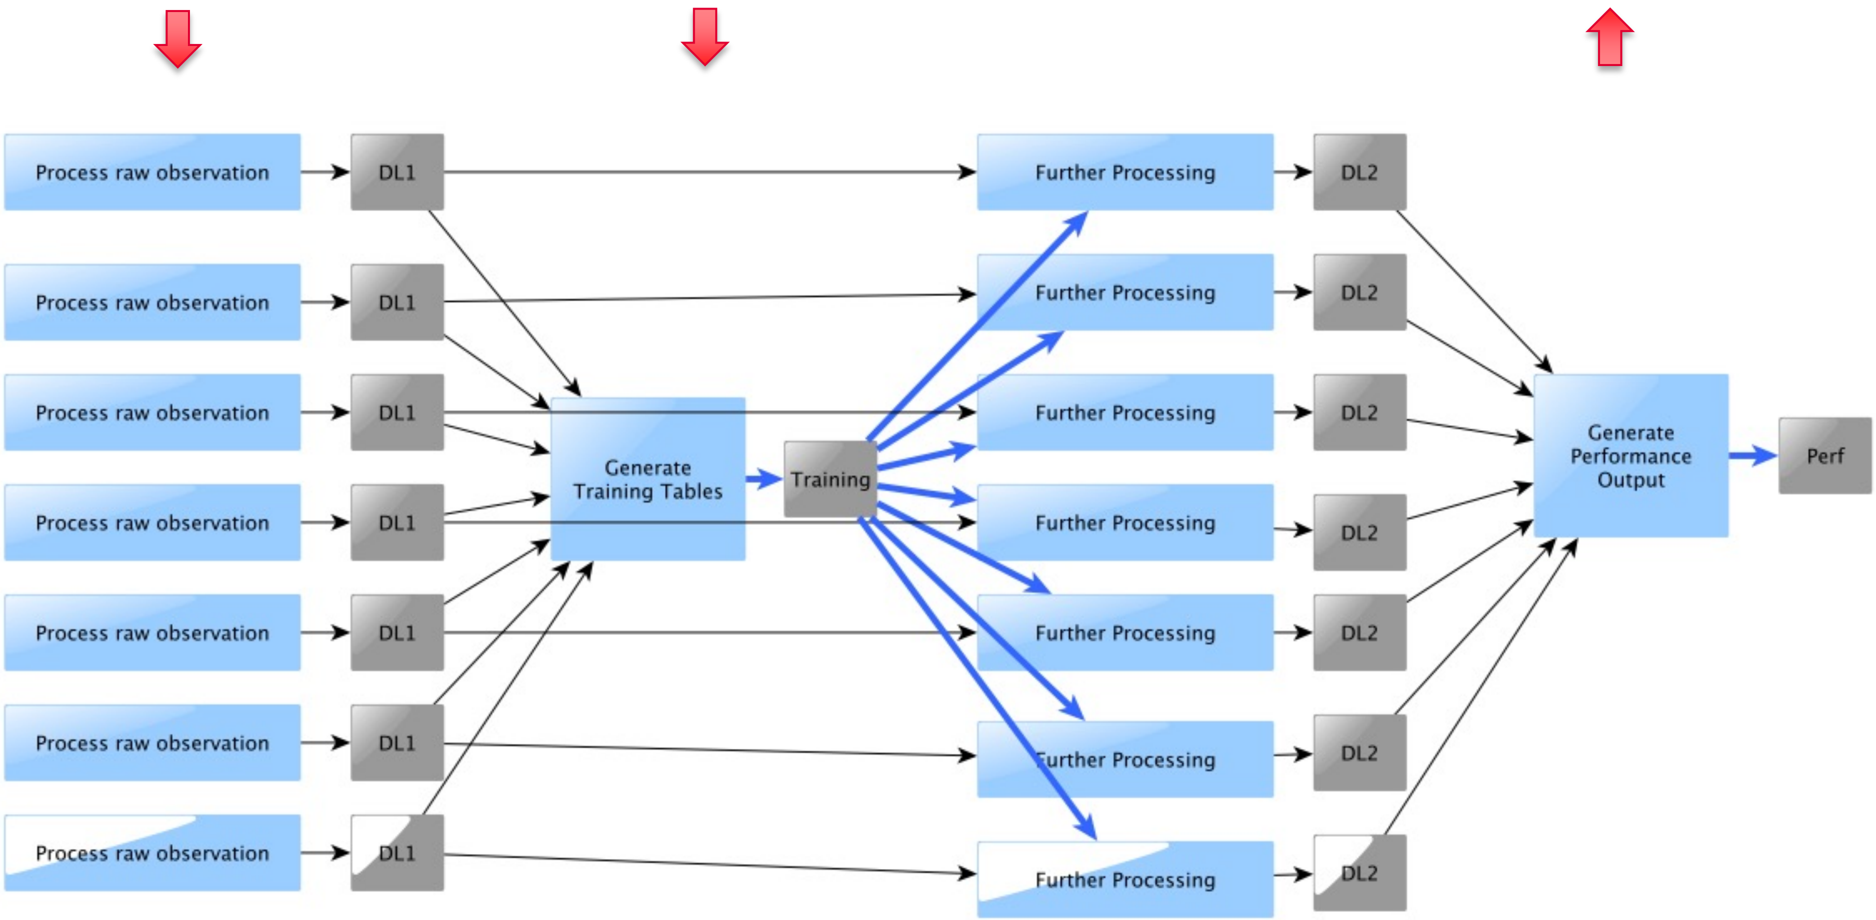

Raw Data Production

- CTA raw data size (as powers of 10 only):
  - 100 telescopes (CTA-south)
  - 10,000 array triggers per second
  - 10 telescopes on average per trigger
  - 10-100 image frames per telescope camera
  - 1,000 to 10,000 pixels per camera
  - ÷ 100 lossy and lossless compression
  - O (10) PB/yr of raw data (DL0)

- Number of raw data files:
  - One data/file stream per telescope and observation
  - Files probably split (file size not defined)
  - O (10^6) files/yr of raw data, equal amount for simulations

- Monte-Carlo Simulations
  - (basically continuously, similar data volume and number of files)

- In addition
  - Monitoring data, configurations, …

- Yearly reprocessing of *all* data with new calibration and reconstruction (30 year lifetime of CTA…)

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Data Levels and Storage Requirements

- Strong data reduction along the processing steps
  - From PB/y (DL0) to GB/y (high-level science data, DL3-DL5)
- Total storage requirements
  - +6 PB/y on Disk
  - +21 PB/y on Tape
- For simulations:
  - GRID-based approach using DIRAC
  - already now: 130M HS06 CPU-Hours used and 10 PB transferred
- Model with distributed resources explored for operations phase

Data Level	Short Name	Definition
R0	Raw Internal	On-site streamed raw data, not normally preserved long-term in this form. R0 content and format is internal to each device / controllable system, such as raw data transmitted from the physical device / system to its respective server in the on-site Data Centre.
R1	Raw Common	On-site stream raw data meeting common standards, transmitted on-site from a Camera or other on-site system to the OES. This is the first level of data seen by the OES, that will typically need some pre-processing from the R0 data format. Exceptionally, some R1 data may be stored for engineering purposes.
DL0	Raw Archived	All archival data from the data acquisition hardware/software, transmitted from the OES to the DPPS. This is the lowest level of data that are intended for long-term storage in the bulk archive. This includes both camera event data and technical data from other sub-systems, such as non-camera devices or software.
DL1	Processed	Processed DL0 data that may include telescope-level (TEL) data and parameters derived from them. Typical contents include calibrated image charge, Hillas parameters, and a usable telescope pattern. DL1 data is not normally stored long-term.
DL2	Recon-structed	Reconstructed shower parameters such as energy, direction, particle ID, and related signal discrimination parameters. Does not include telescope-level (TEL) information. For each event this information may be repeated for multiple reconstruction and discrimination methods. DL2 data is not normally stored long-term.
DL3	Reduced	Sets of selected events with a single final set of reconstruction and discrimination parameters, along with associated instrumental response characterizations and any technical data needed for science analysis.
DL4	Binned	Data product produced by binning of DL3 data, including data cubes and maps which are suitable for combination/summation to produce DL5 products.
DL5	Science	Data product produced by combination of DL4 products an extraction target specific region(s) of interest. Includes for example light-curves and spectra, along with associated data such as source models and fit results.
DL6	High-Level	High-level or legacy observatory data, such as survey maps and source catalogues.

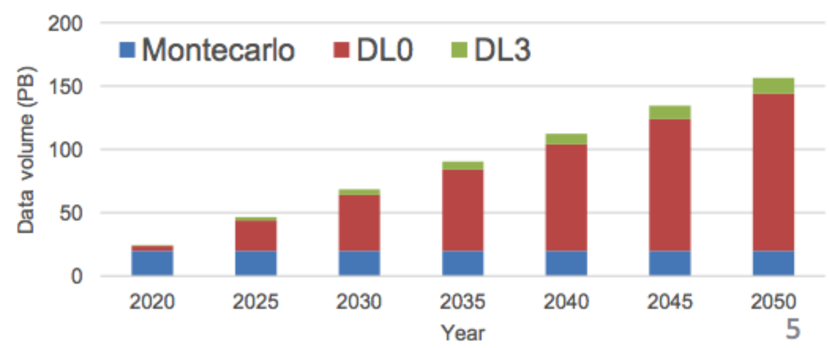

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Intercontinental link

- « Relatively » small capacity data centers planned on-site (Chile, La Palma – Spain), science data management center in Europe
- Network link quality/stability obviously not so good from Chile
  - Requirement and current baseline is 1 Gbps for both sites, higher bandwidth under investigation
- On-site buffering and transfer to Europe
  - No archiving on-site, only buffering of data for limited period

Funded by the European Union's Horizon 2020 - Grant N° 824064

CTA Bulk Archive

- Bulk Archive
  - Archive bulk raw-data (DL0-DL3) and associated metadata information over the lifetime of CTA
  - Preserve DL0 simulation data for at least 3 years after production
  - No public access, only CTAO staff
- Performance of bulk archive
  - Located at least at 2 sites with 300 km distance
  - Handle increasing data volume of ~10 PB/yr
  - Allow fast (re-)processing of data (annual reprocessing of all data within 1 month)
  - No data loss over the full lifetime of CTA

Funded by the European Union's
Horizon 2020 - Grant N° 824064

CTA Bulk Archive requirements 1/2

- Bulk Archive will follow (as much as possible) OAIS standard as reference architecture
- OAIS is an ISO standard, defined by Committee for Space Data Systems

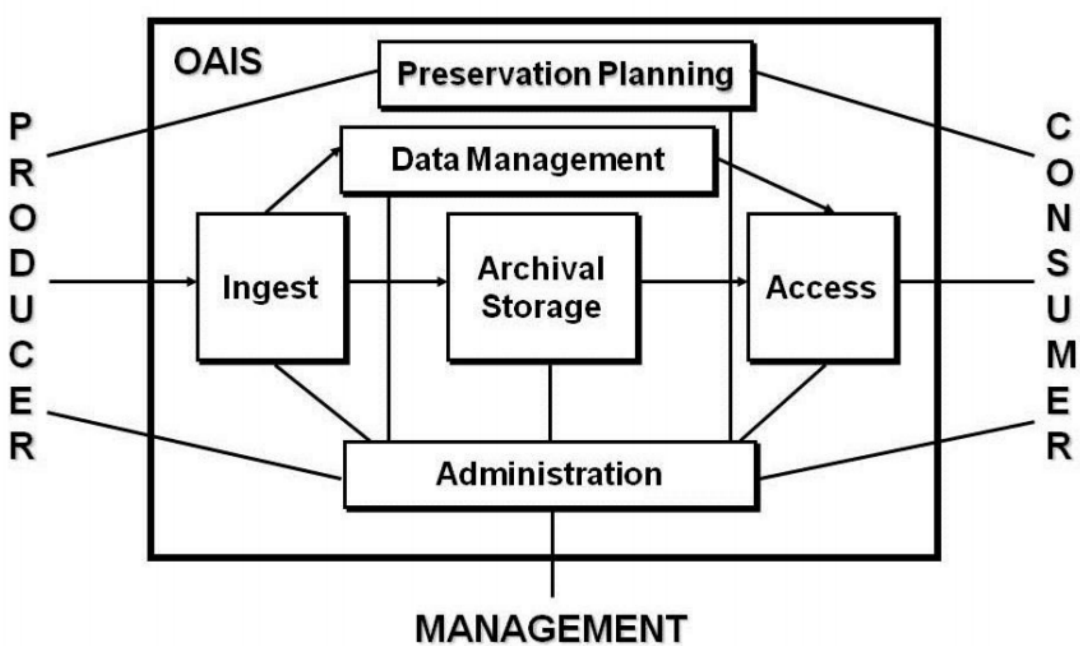

Diagram taken from the OAIS Reference Model by the Digital Preservation Coalition

d by the European Union's
Horizon 2020 - Grant N° 824064

CTA Bulk Archive requirements 2/2

- Storage-related requirements
  - Management, operation of and access to Bulk Archive only by CTA Observatory staff (no public access)
  - Access rights management supporting specific roles for CTA Observatory staff (Archive Manager, Data Processing Manager)
  - Unique identifiers for all data products that are independent of the storage location or number of copies
  - Versioning of data products (limited numbers of versions to be archived)
  - Placement, Replacement, Duplication, Migration of data products and metadata
  - Multiple locations and different QoS with at least 2 sites with 300 km distance
  - Archive bulk raw-data (DL0-DL3) and associated metadata information over the lifetime of CTA
  - Preserve DL0 simulation data for at least 3 years after production (max size 20 PB)
  - Handle increasing data volume of ~10 PB/yr
  - Embargo of 1 year (at the archive level?)

- Ingest/Access-related requirements
  - Metadata extraction (nice to be transactional) and browsing
  - Update of metadata and regeneration of metadata from data products
  - Confirmation of the availability of requested data products and estimation of the retrieval time within 1s of the search request on average
  - Raw data to be processed within 1 month
  - 1 full re-processing per year

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Open Questions for WP2

- all data will be organized in datasets

  possible datasets at different processing levels:

  - dataset per telescope per observation (DL0)
  - dataset per observation (DL3)
  - dataset per science target (DL5)
- Colocation between CPU & Data

  For Raw data we expect to have colocation between CPU & Data & Software
- Additional requirements
  - Command “Staging” (to move from tape to disc)
  - Retrieve the physical location of a file (to support CPU & Data colocation)
- Workflow and Resource management likely DIRAC-based

Funded by the European Union’s Horizon 2020 - Grant N° 824064

CTA Software Stack

- All software is prototype status and will evolve (might even switch technology)
  - File Transfer System (via intercontinental link)
    - So far looked into FTS
  - Workflow and Computing Resource Management System
    - Prototype based on DIRAC
  - Data Processing Pipelines
    - Prototype based on python
  - Archive Management Software
    - Prototypes based on OneData, …

- DL0 data model and data format
  - Data model currently under discussion
  - Data format will be chosen following benchmark tests

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Topics related to Data lake

- “Data lake of the core CTA processing engine”
  - Caveat: only simulated CTA data
- Goal: large-scale automatic processing O(10PB/y) with minimal operations costs
- Limited number of users (CTAO staff)
- Flexibility of computing model
  - Distributed vs. centralized computing
  - Add/remove data centers
  - Migrate storage resources and technology
  - Usage of HPC resources when applicable
  - …
- Flexibility of workflows
  - Evolution of processing workflows
  - …
- Queries and metadata expected to be rather stable
  - will evolve if new algorithms (e.g. ML, deep learning) become available

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Part 2: From Reduced Data to Science User

From Science-Ready Data to Science User

- Science-ready data products (DL3) available *in science archive*
  - Event List + IRFs + binned products for the observations
- Automatic quick-look data products (DL4, DL5) dependent on the science users proposal
  - Skymaps, spectra, …
- Enrich data with appropriate metadata (incl. provenance, science-related metdata)
- Data protected according to proprietary period and PI access rights
- Notification of science user on availability of data products

- Reprocessing and dissemination of data products when new DL3 is available from data processing / bulk archive, announced release after verification and notification

Funded by the European Union's
Horizon 2020 - Grant N° 824064

CTA Science Archive

- Science Archive
  - Long-term Preservation of Science Data (DL3-DL6) with associated metadata and provenance information
  - Small amount of data, Large number of users
  - public access, access rights management
  - allow for a user-contributed part (e.g. catalogues)
  - Preservation beyond lifetime CTA
- Access to Science Archive
  - Via CTA Science Data Management Center
  - Link to IVOA
- Computing model tbd
  - Distribution may not be strictly needed as volume is limited

Funded by the European Union's Horizon 2020 - Grant N° 824064

Science Archive Requirements

Requirements for the science archive

- public access with data rights management (proprietary period)
- Preserve science data products (DL3-DL5, DL6) beyond the lifetime of CTA
- Preserve associated metadata and provenance information
- Versioning of data products and software
- Flexible queries for users
- Fast access (product searches within 1 min)
- Support automatic data processing (DL3→DL5) and data products verification
- Support data browsing and interactive exploration
- Support standard interfaces (VO compliance)
- Unique identifiers for at least DL5 and DL6 products
- Highly available, high quality of service and products

Further considered:

- User-contributed part of the science archive for end-users
- Link of data products identifiers (e.g. DOIs) with publications and usage

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Science User Perspective

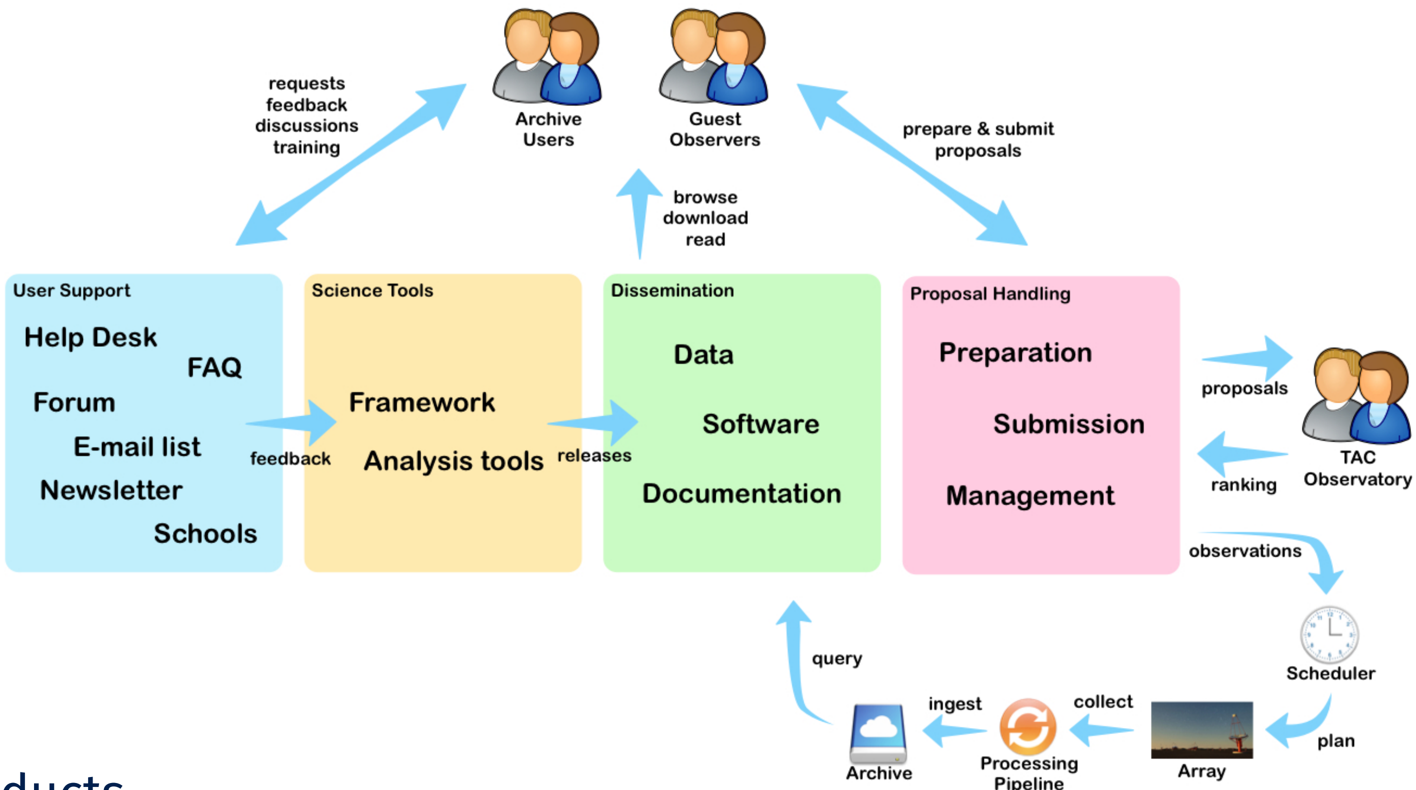

- Products
  - Photon (candidate) event list data (FITS)
  - Instrument response functions, background model
  - Science analysis tool suite, supporting docs.

Funded by the European Union's
Horizon 2020 - Grant N° 824064

CTA Observatory as a User Facility

- Science Archive contains science-ready Data Products low in volume, but rich in information
  - Volume is O(TB/yr), incl. all metadata and provenance information, incl. science alerts, enables users to build catalogues covering the full sky

- User of the CTA Observatory is the world-wide community (>1000)
  - Users will have wide range of backgrounds and expertise in terms of gamma-ray data analysis

- Products and services offered to the community include
  - Data products (Event lists, Instrument Response Function, Quicklook products (skymaps, spectra, lightcurves)
  - Tools for proposal preparation
  - Tools for Science Analysis
  - User support services
  - Science Portal, FAQ, Newsletter
  - PI interactions, incl. science alerts
  - Interactive workflows
  - Visualisation and Interactive data exploration
  - Establish link to VO and publish publicly accessible data
  - ...

- Products and services are in the core of the CTA Science Portal

Funded by the European Union's Horizon 2020 - Grant N° 824064

CTA Software Stack

- All software is prototype status and will evolve (might even switch technology)
  - Science Analysis Tools
    - Prototypes SW packages based on python and C++ (gammapy, ctools)
  - Science Gateway (not portal)
    - Prototype framework for integration of services based on http, rabbitmq, protobuf
  - Science Portal
    - Investigating interactive data exploration via Jupyter notebooks
- DL3 data model and data format
  - Data model currently under discussion for both event lists and instrument response function
  - Data format will be fits

Funded by the European Union's Horizon 2020 - Grant N° 824064

Topics related to Science Platform

- Queries and metadata need to be flexible
  - Queries are science-driven and user-based
- Provenance metadata (link to WP4)
  - Establish provenance information in Science Platform metadata
  - Based on experiences in ASTERICS
- Multi-messenger and metadata (incl. provenance) (link to WP4)
  - CTA key science includes multi-messenger multi-wavelength projects
  - Enrich a Science Platform with multi-messenger information and related services
    - Where does the archive of all science alerts live?
  - Add provenance information related to multi-messenger information
- Open questions
  - How to deal with data under embargo (in proprietary period)?

Funded by the European Union's
Horizon 2020 - Grant N° 824064

That’s it.