

Analysis Facility at Port d'Informació Científica (PIC)

Francesc Torradeflot

VRE Working Group Meeting
2024-12-04

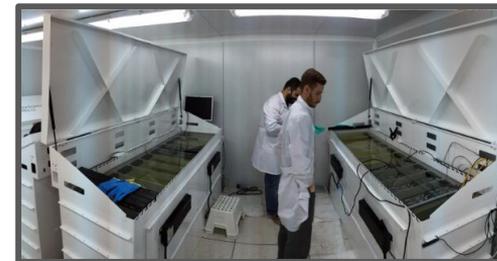
Outline

- Current status
 - Core services
 - JupyterHub
 - CosmoHub
- Future plans

Current status: Core services

- Connectivity
 - 2x100 Gbps to Academic Network
 - 100 PB in+out per year
- Data processing services
 - Disk - dCache: 20 PB (+Ceph 3.5 PB raw)
 - Tape - Enstore: 63 PB
 - Computing - HTCondor: 12700 slots, 16 GPUs
 - Computing - Hadoop: 720 cores, 2.5 PB disk
- Facilities, ~120 kW IT
 - ~80 kW in 150 m² air-cooled room
 - ~40 kW in 25 m² liquid immersion cooling system
- Kubernetes, VMs, etc

IBM TS4500



Current status: JupyterHub

- Launch a jupyter notebook server on PIC's HTC cluster using [jupyterhub](#) and [batchspawner](#)
- User-defined resources
 - CPUs
 - Memory
 - GPUs
- Choose experiment for accounting and POSIX permissions
- Managed with puppet & gitlab CI/CD
- High priority jobs to minimize waiting time

Server Options

Select custom options for your profile

Memory (RSS)

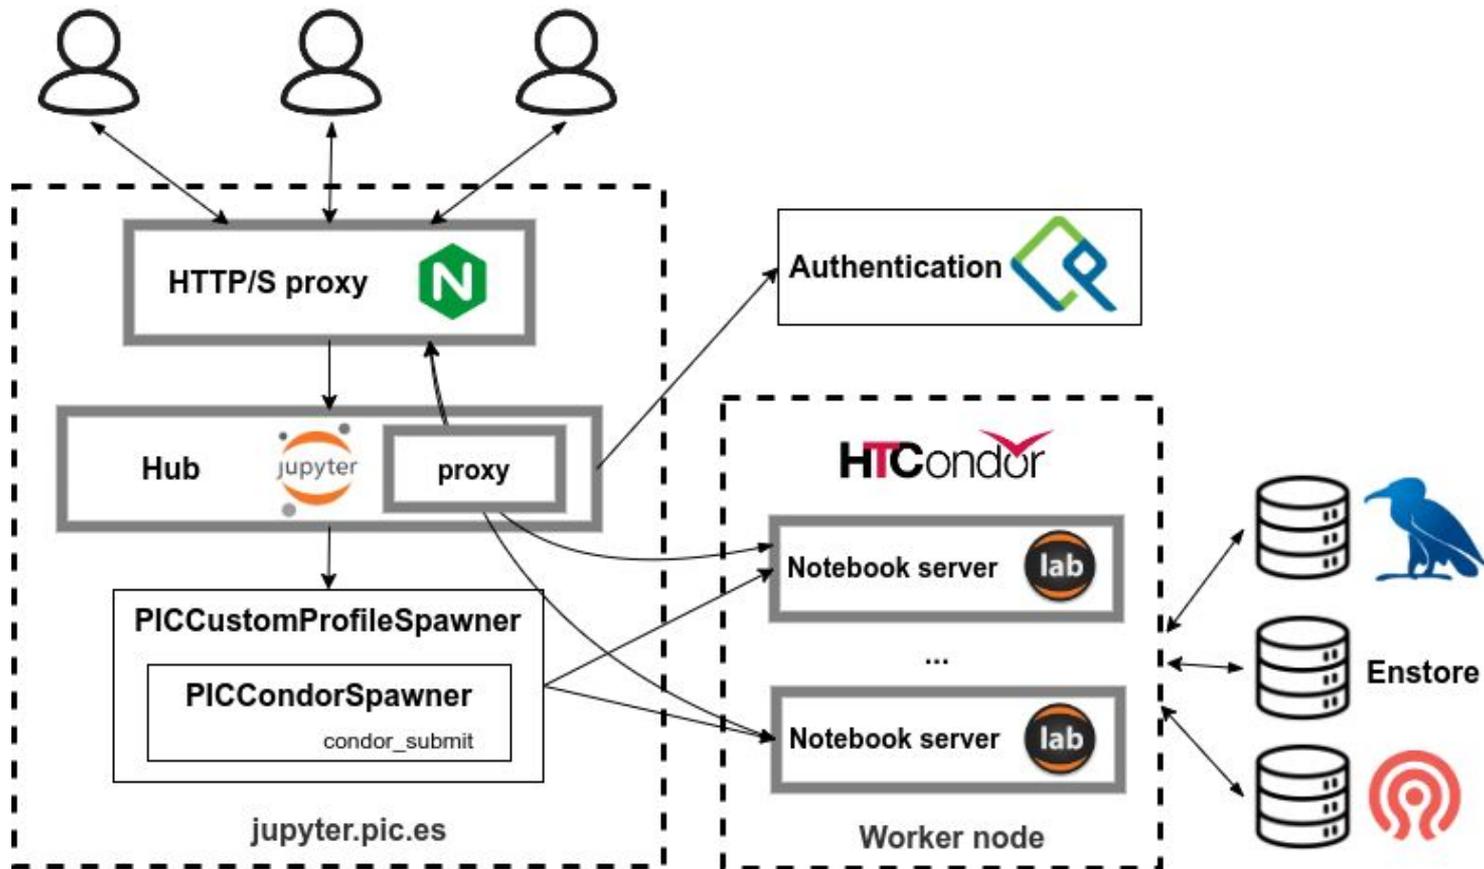
CPUS

GPUS

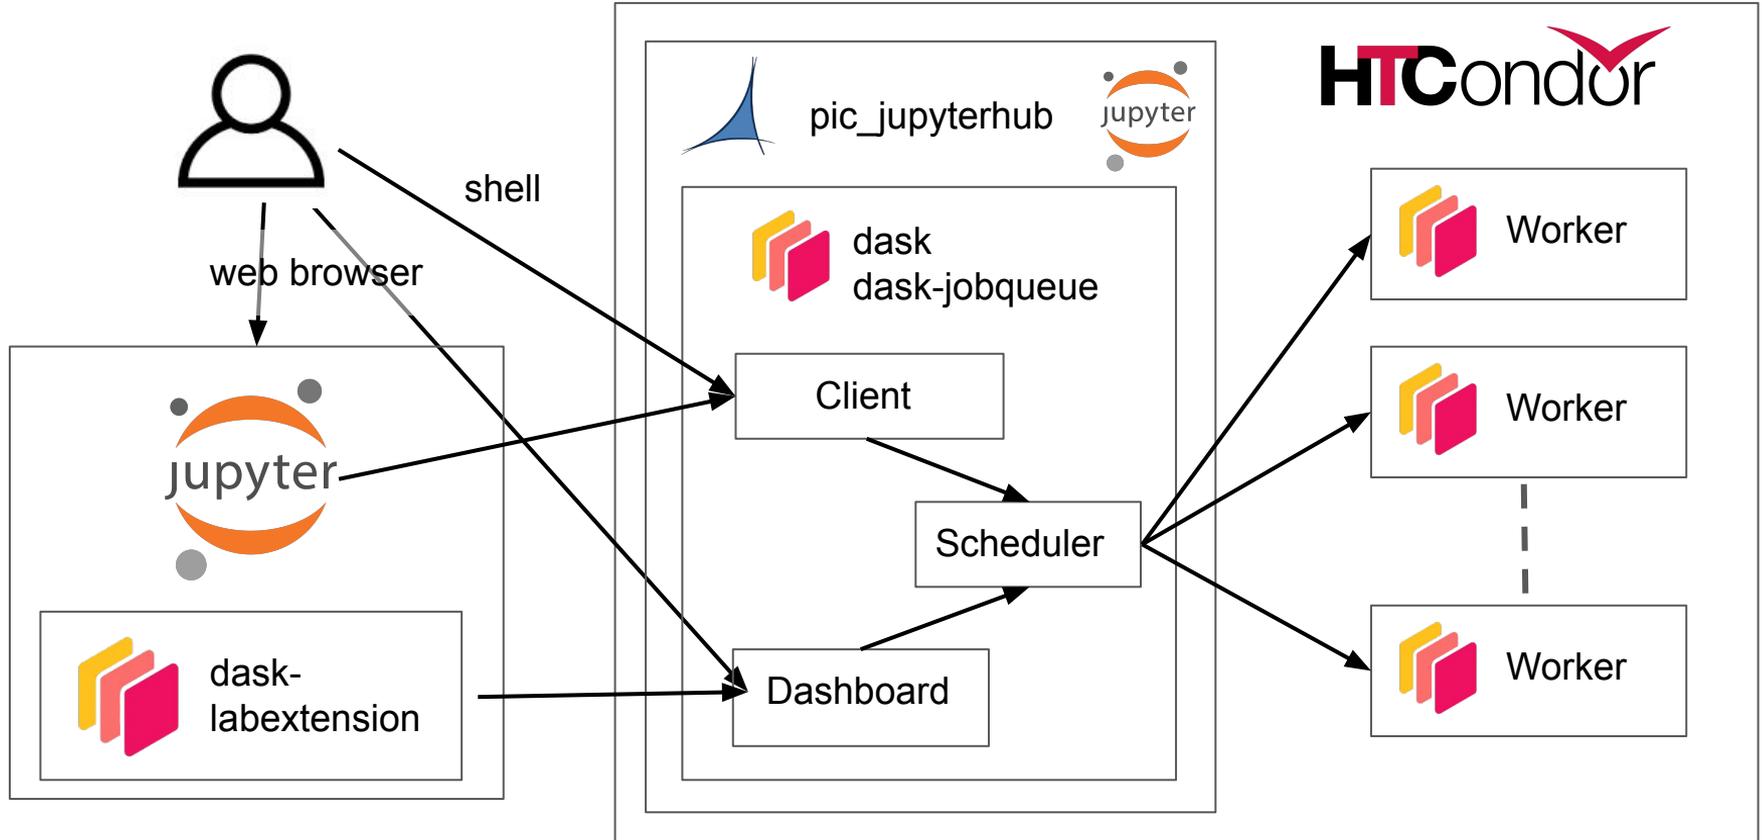
User options

| | |
|------------|------------------------|
| Experiment | Select your experiment |
|------------|------------------------|

Current status: JupyterHub



Current Status: JupyterHub



Current status: Hadoop

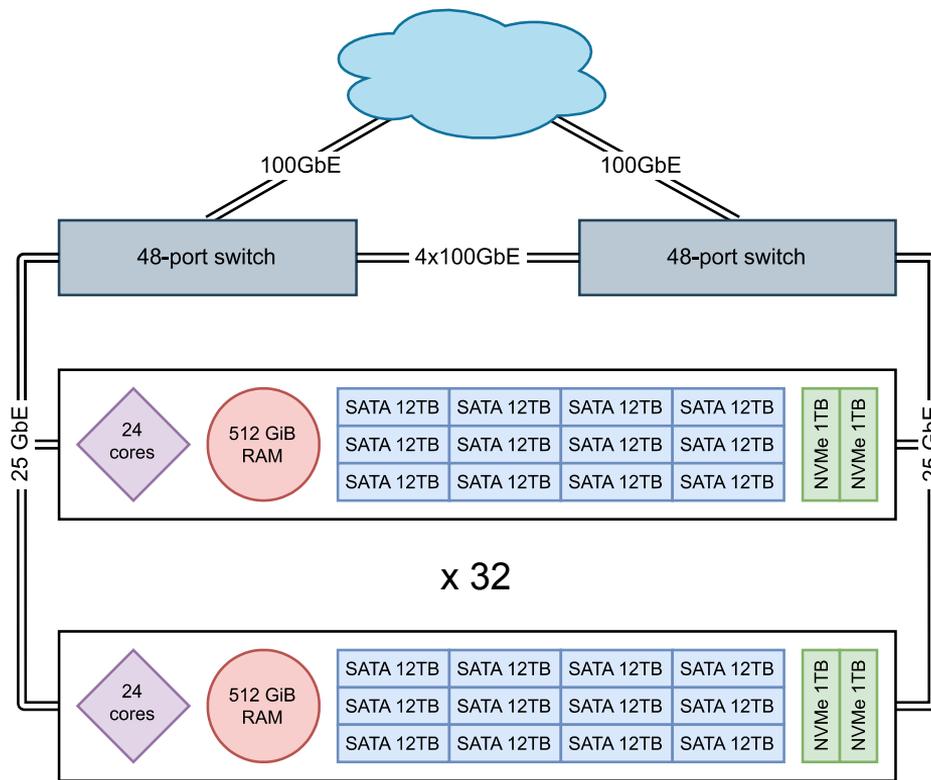
32 nodes:

- 768 cores, 16 TiB RAM
- 60 TB NVMe (for cache)
- 4.3 PiB raw storage (3 PiB usable)

In-house Hadoop distribution: Shepherd

Main use cases:

- **CosmoHub** query processing
- Euclid **mock galaxy catalogs**



Current status: CosmoHub

cosmohub.pic.es



Build your own Universe

Interactive data analysis of massive cosmological data without any SQL knowledge



Billions of observed and simulated galaxies



Superfast queries means superfast results



Features to make you work faster and easier



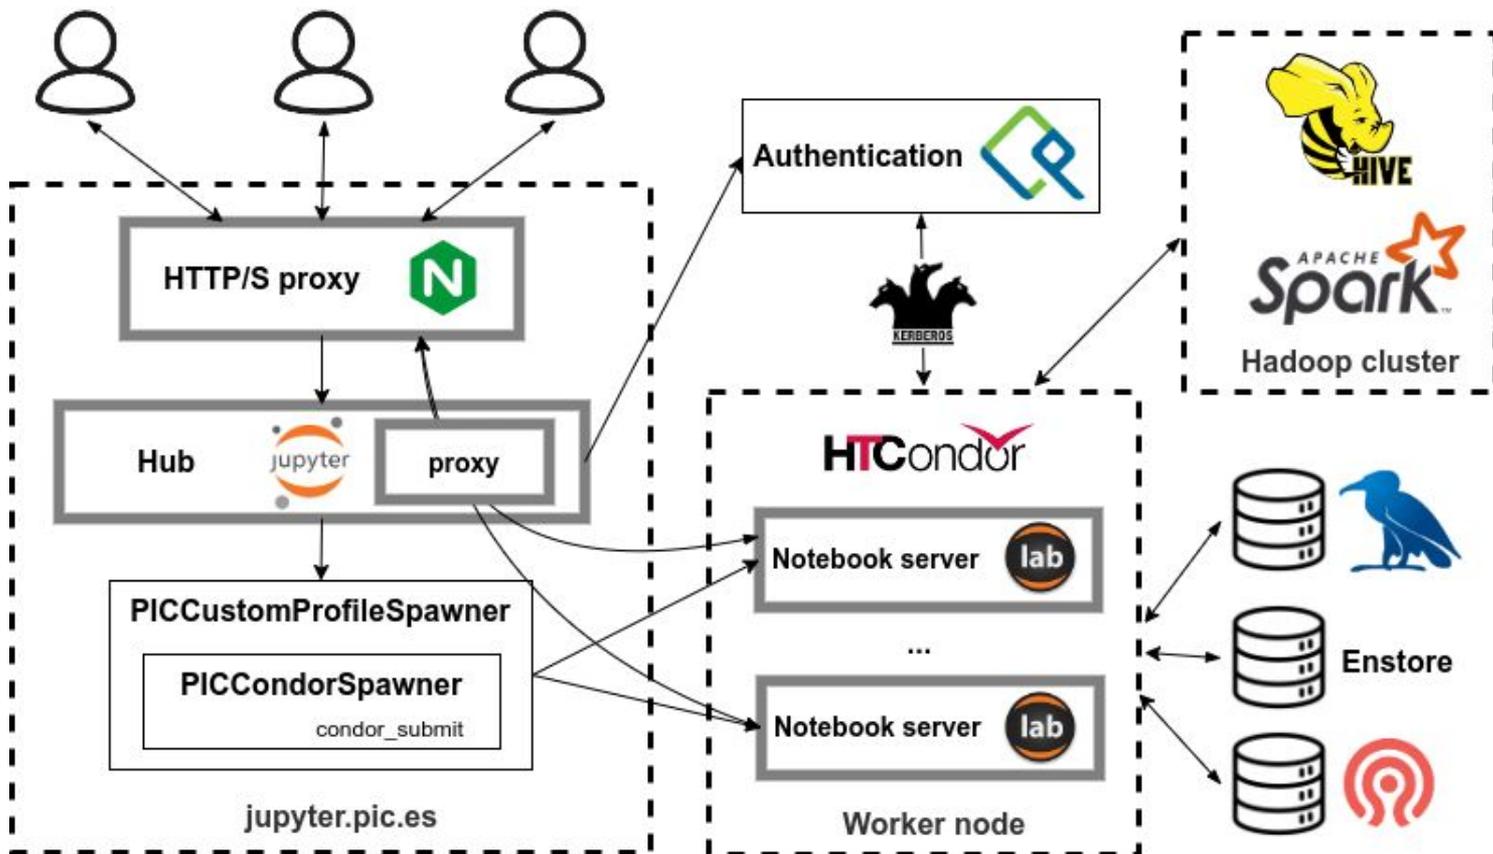
Online plotting preview and data download

Current status: CosmoHub

- Web application for tabulated data exploitation on top of Hive 
- Interactive exploration/visualization
 - **User defined functions**
 - HEALPix, Arrays, spherical geometry
 - **Exploration**
 - Cone search tool
 - Guided process (no SQL knowledge needed)
 - **Visualization**
 - 1D histogram & 2D heatmap
- Distribution
 - Parquet, CSV, FITS, ASDF format
 - Email with a link to download dataset

- Data
 - 90 TiB catalogued data
 - >130 catalogs (simulated and observed)
 - **Supporting multiple projects**
 - DES, PAUS, Euclid, MICE, LST, Gaia, LSST...
- Users
 - >1750 registered users
 - ~150 active users
 - >17K custom catalogs generated
 - >20k interactive queries
- Performance
 - >75% of all queries finish in <3 min
 - Resource queues with reservation
 - Preemption to keep interactive response time

Current status: Jupyter + Hadoop



Future plans

- Core services
 - Move tape system to a new location to avoid humidity issues
 - Migrate from Enstore to CTA
 - Expand and stabilize Ceph cluster
 - Renew immersion cooling system
 - Install and configure newly acquired GPU servers
- Authentication and authorization
 - Integrate with trusted identity providers (LSST, CTA, Alba)
 - Management of 100s of short-lived users (students)
 - Sync users and projects in Alba's EM facility

Future plans

- Hadoop / CosmoHub
 - Migrate CosmoHub API to python3 (still in python2.7!!)
 - Backfilling of Hadoop nodes with HTCondor jobs (in progress)
 - Provide IVOA TAP endpoint to CosmoHub catalogs
 - ADQL compliance
 - Integrate Cosmohub login with PIC's IdP

- JupyterHub
 - Adapt landing page to new users, GPU types, etc
 - Install and integrate rucio-jupyterlab extension
 - Preconfigure Hadoop clients
 - Integration of Dask and Hadoop's backfilling
 - Add KubeSpawner for purely external users (?)

Thank you!