



Elettra Sincrotrone Trieste



Elettra
Sincrotrone
Trieste

Computing and Storage Architecture

At

Elettra-Sincrotrone Trieste



USC-C General Assembly, Trieste

Iztok Gregori, 11th March 2026



Elettra
Sincrotrone
Trieste

What will we talk about

- Introduction – Who, where and what?
- Infrastructure Overview
- Scientific Data Acquisition
- HPC Processing



Elettra
Sincrotrone
Trieste

Who, where and what?



USC-C General Assembly, Trieste

Iztok Gregori, 11th March 2026



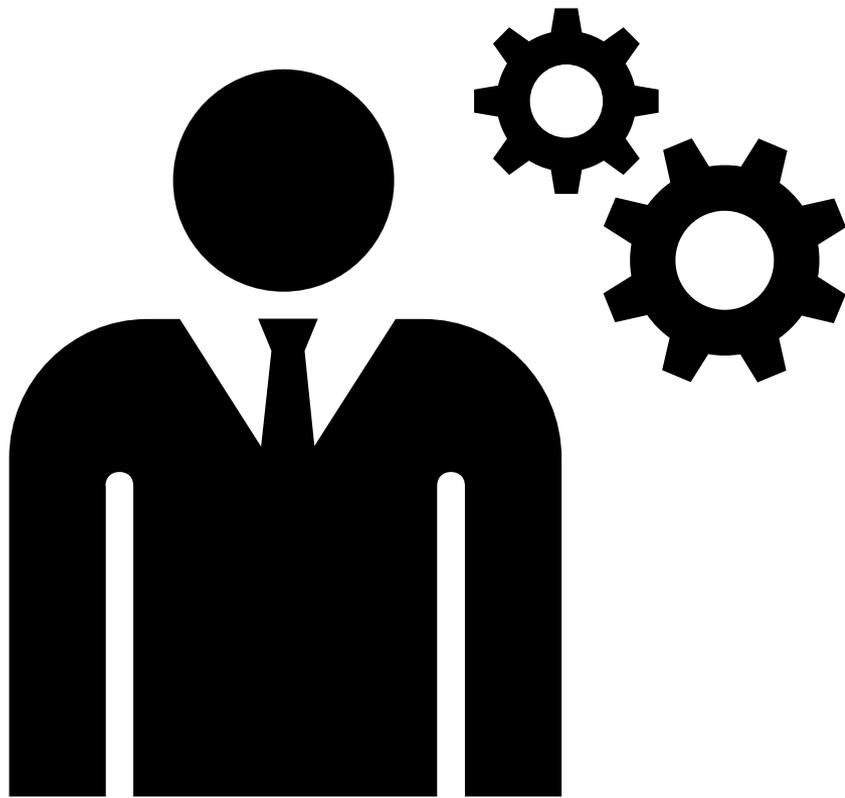
Elettra
Sincrotrone
Trieste

What is Elettra?

- Multidisciplinary research center located in Trieste, Italy
- Running 2 particle accelerators producing synchrotron light
- 30+ beam-lines/laboratories
- 350 people + over 1000 visiting scientists every year



Who am I?



About myself

- Linux system administrator
- Part of the ICT System and Service group in the IT department since 2017
- Working with Virtualization and Storage infrastructure, HPC and much more



Elettra
Sincrotrone
Trieste

Infrastructure Overview



USC-C General Assembly, Trieste

Iztok Gregori, 11th March 2026



Elettra
Sincrotrone
Trieste

Build on Opensource Software

- Around 99% of infrastructure and 90% of data acquisition runs on opensource software.
- Key elements:
 - Proxmox for Virtualization
 - Ceph for Storage
 - Slurm and Kubernetes for HPC



ceph



Elettra
Sincrotrone
Trieste

Elettra – IT Numbers



- 4 Datacenter.
- 150+ physical servers,
~1000 VM.
- 10 PB raw storage.
- Core Network at 10
Gb/s.
(updating to 100 Gb/s)



Elettra - Clusters



- 2 Hyper-converged clusters (Proxmox+Ceph) for machine controls.
- 1 Hyper-converged cluster (Proxmox+Ceph) for support services.
- 1 cluster (Proxmox and Ceph) scientific data handling.
- 1 cluster (Proxmox and Ceph) for backup.
- 1 cluster (Ceph) for S3 storage.



Elettra
Sincrotrone
Trieste

Scientific Data Acquisition



USC-C General Assembly, Trieste

Iztok Gregori, 11th March 2026

12



Scientific Data Acquisition

- Different beam-lines, different storage requirements, common acquisition pipeline.
- Pipeline managed through in-house build application (VUO).
- Data retention is 10 years.



Data Acquisition Pipeline

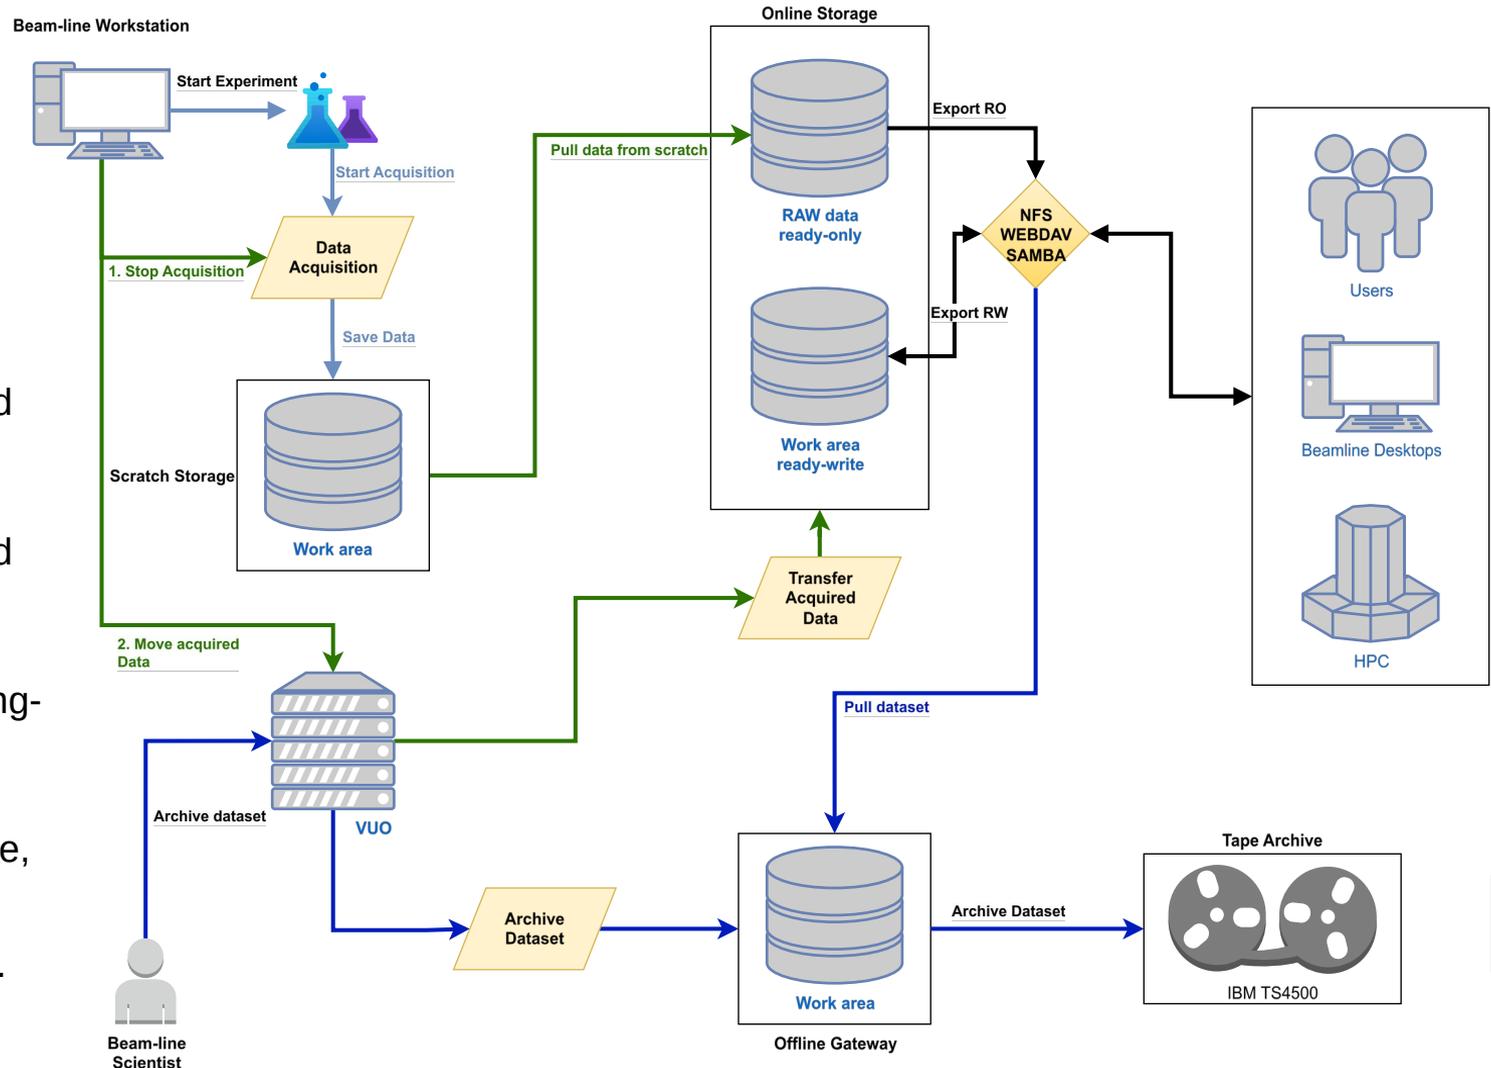
1. Acquired experiment data is saved on the **scratch storage**.
2. Once the acquisition is done, data is transferred to the **online storage** for later consumption.
3. Automatically and on request data is copied (or moved) to the **offline storage** (LTO tapes).



Elettra
Sincrotrone
Trieste

GLOSSARY

- **Scratch**
Storage space used during acquisition
- **Online**
Storage space used after acquisition
- **Offline**
Space (tape) for long-term storage.
- **VUO**
Virtual Unified Office, core application for managing data acquisition pipeline.





Elettra
Sincrotrone
Trieste

High Performance Computing



USC-C General Assembly, Trieste

Iztok Gregori, 11th March 2026

16



Online processing

- Data Processing in-real-time (during acquisition).
- CPU & GPU processing pipeline, in-house development.
 - 3 x Nvidia A100
 - 3 x Nvidia L40S
 - 4 x Nvidia L4
 - 6 x Nvidia T4
- On-premise Kubernetes cluster as computing infrastructure.



Batch processing

- “Offline” (not in real time) processing of data.
- “Kalculus” cluster
 - Slurm batch scheduler.
 - 56 nodes, 3 partitions, 1500+ cpu cores.



Elettra
Sincrotrone
Trieste

QUESTIONS?

www.elettra.eu