The IT Challenges For Research Infrastructures In Physics

CERN 2013

Laurence.Field@cern.ch

Overview

- Data generation
- Tour of the Research Infrastructures
- The activity in the CRISP project
- Cluster project collaboration

Research Infrastructure

Beamlines

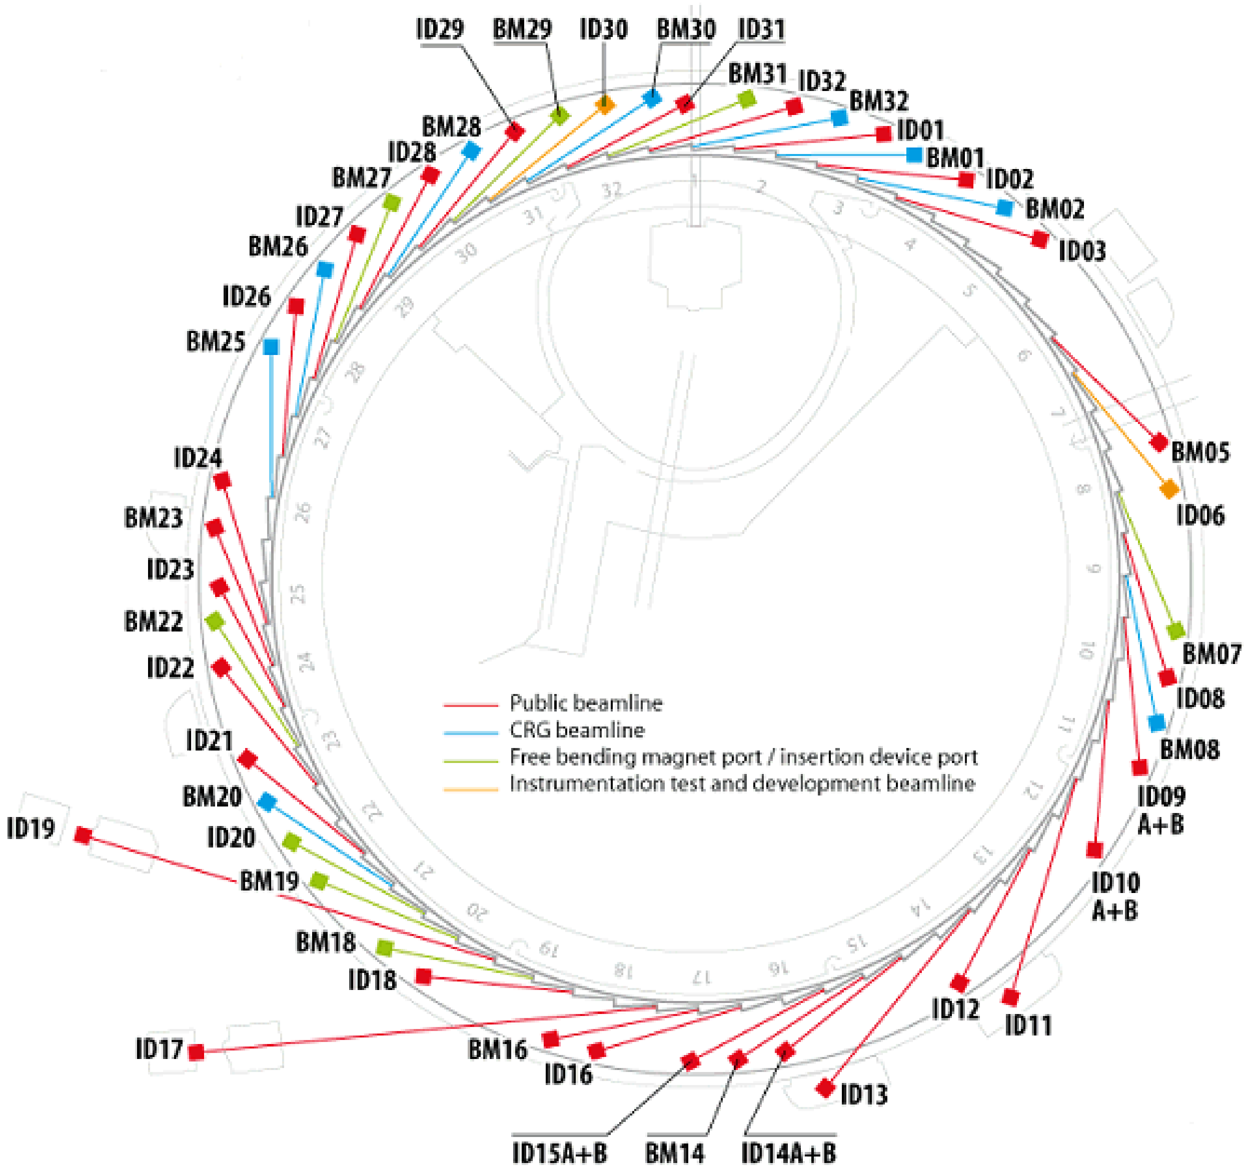
ID29
BM29
ID30
BM30
ID31
BM31
ID32
BM32
ID01
BM01
ID02
BM02
ID03
BM05
ID06
BM07
ID08
BM08
ID09 A+B
ID10 A+B
ID11
ID12
ID13
ID14A+B
BM14
ID15A+B
ID16
BM16
ID17
ID18
BM18
BM19
ID19
ID20
BM20
ID21
ID22
BM22
ID23
BM23
ID24
BM25
ID26
BM26
ID27
BM27
ID28
BM28
Public beamline
CRG beamline
Free bending magnet port / insertion device port
Instrumentation test and development beamline

Beamline

Detector

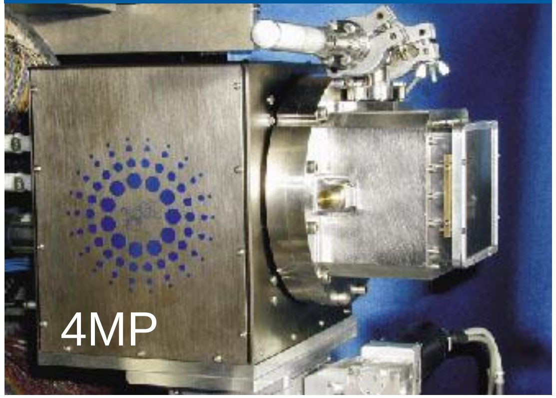

Instrument

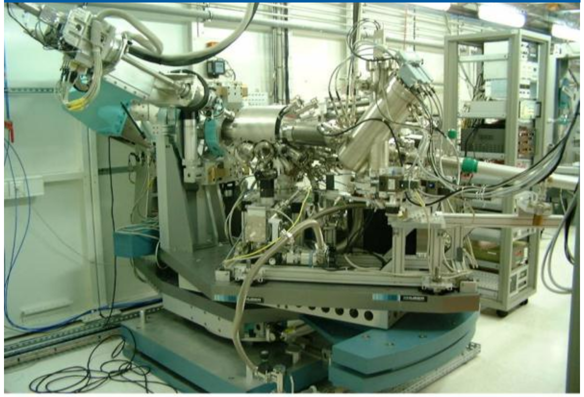

Beam Optics

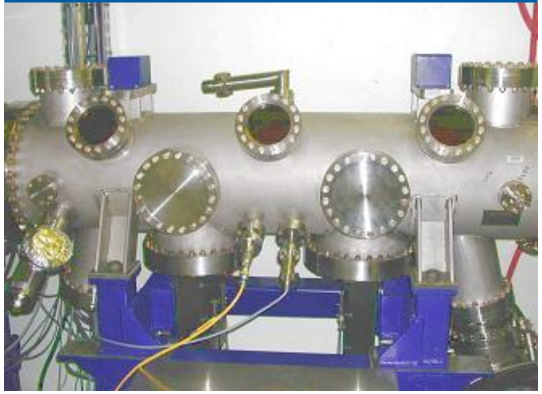

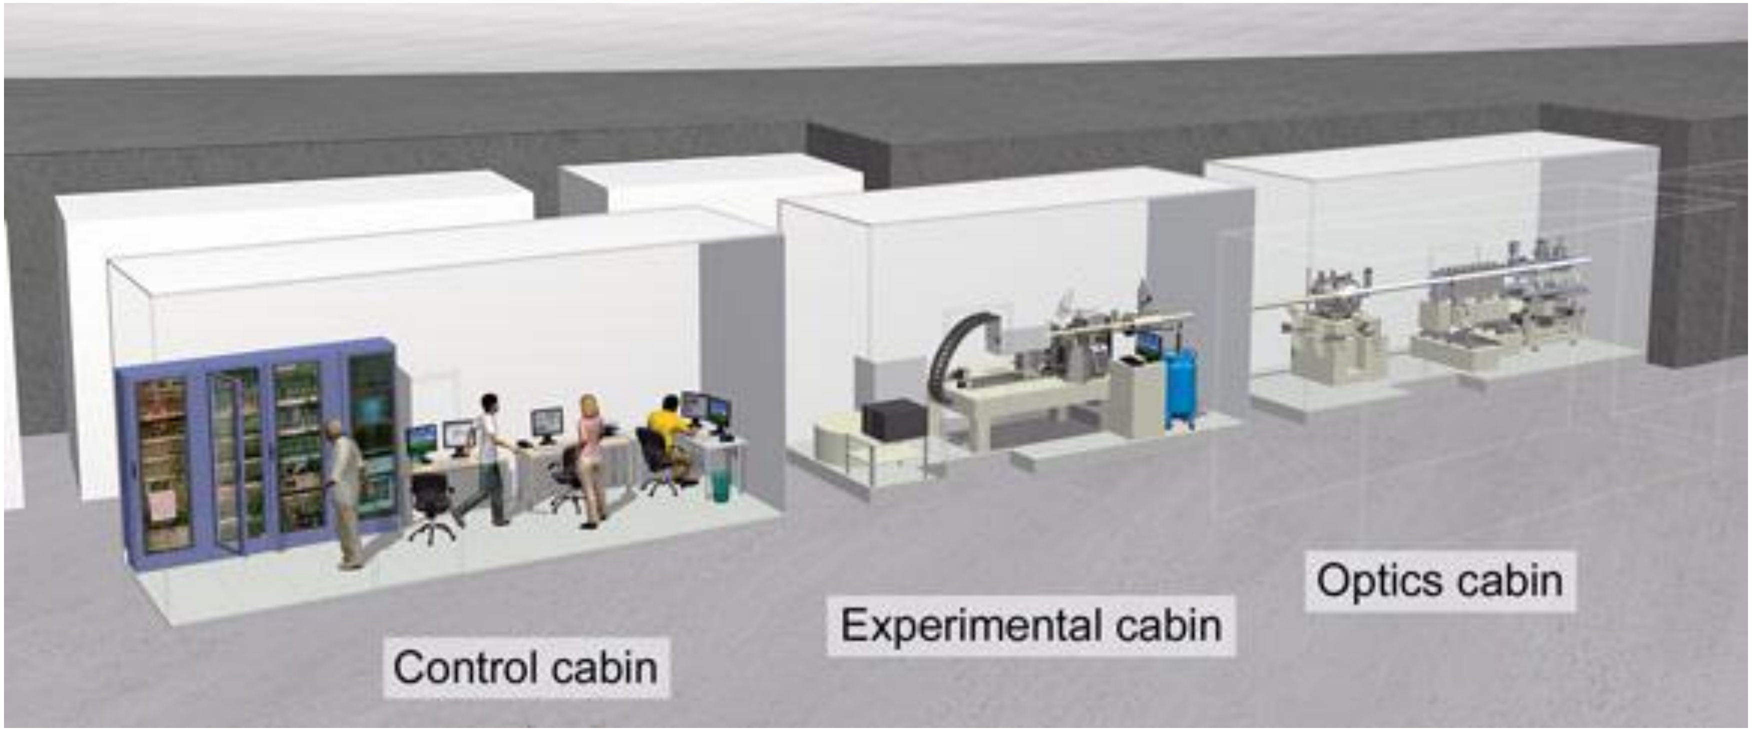

Data

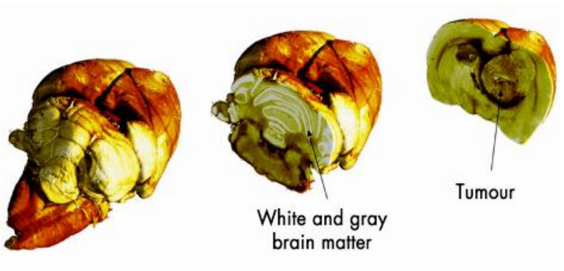

- Pixel-based detectors
  - A series of 2D images
    - 3D volume or a time series
- Data rates and volume are depend on the instrument
  - Rates range between a few kB/s to a few GB/s
- Data locally buffered at the instrument
  - Before being moved to a central storage facility
- MetaData
  - Experiment parameters
- Digital microscopes
  - Providing molecular, atomic and sub-atomic resolutions

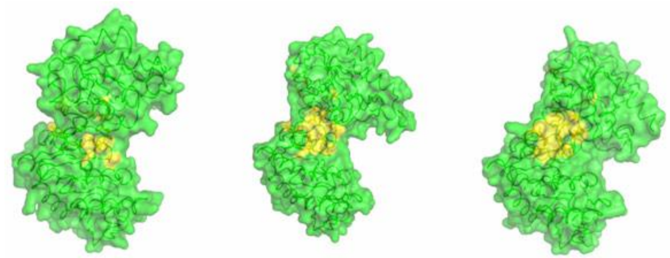

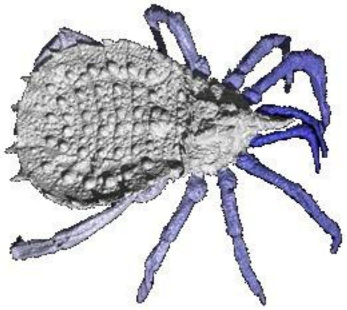

ESRF

- Proposal for *beamtime* submitted using an online application form
- Data moved to the central storage facility using NFS mounts
  - Data deleted from disk 50 days after the end of the experiment
  - Backup kept on tape for 6 months
- SAN storage currently 1.3 PB with plans to expand by another 1PB
- LTO currently a few thousand tapes, ~4-5 PB
- Archiving is implemented on a case by case basis
  - De facto for in-house research
- Local Data Processing and Analysis Facility
  - Multiple individual data analysis clusters
- Plans for new and upgraded experimental stations
  - Resulting in higher data rates

EuroFEL

- Consortium of national FEL and SPS facilities
- Each operates a number of different beamlines
- Aim to aggregate a complete dataset into a single HDF5/NeXus file
- Each facility acts as a custodian for the experimental data
  - Operate their own storage and archival infrastructure
  - Data is remotely accessible through a data portal
  - Online access initially
    - Near-line (tape) after 3-6 months
- Data processing demands are not very demanding
  - A smaller cluster or a single system with GPU-acceleration is sufficient
- Data processing and analysis infrastructure
  - Dedicated per instrument for real-time processing
  - A common compute cluster at the facility

EuroFEL
FREE ELECTRON LASERS OF EUROPE

European XFEL

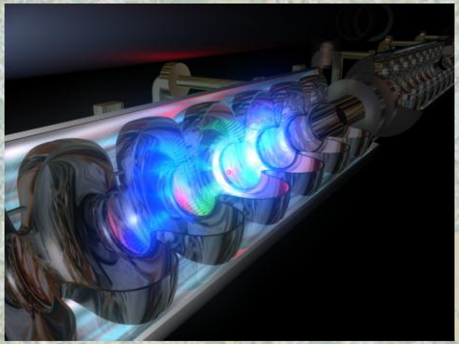

- Will initially operate three beamlines
  - With six different instruments
    - Max. data rate 10 GB/s from a single detector
    - 1PB/year accumulated data volume
- First usage envisaged for 2016
  - To be expanded further at a later stage
- It is expected that all data will be hosted at the facility
  - Online and offline compute infrastructure will also be provided
- Most users will not have sufficient resources for analysis at home
- Data infrastructure must permit remote access to;
  - Scientific data
  - Compute resources
  - Related services
    - Metadata catalogues etc.
    - Data import/export

ILL

- Intense neutron source
  - Application for beamtime is organised via the Visitors Club website
- Provides 40 instruments
  - Data rates vary from 300KB/day to 8.2GB/day
- Data is collected on a local buffer
  - Asynchronously transferred to the central archive
- Currently using a NetApp file server
  - All data (since 1973) can be accessed on-line
- Local support for analysis
  - 1K node compute cluster based on Sun Grid Engine

SPIRAL2

- Linear particle accelerator
  - Commissioning scheduled between mid-2013 to mid-2014
    - Beams scheduled for 2016
- Five different beamlines
- Data rate of up to 300 MB/s
- Data is stored in the GANIL datacentre
  - On-line access for one year
    - Data is then deleted
- Infrastructure requirements;
  - LAN: 5-19Gb/s
  - WAN: 1 Gb/s
  - Storage: 100s TB (extensible and highly available)

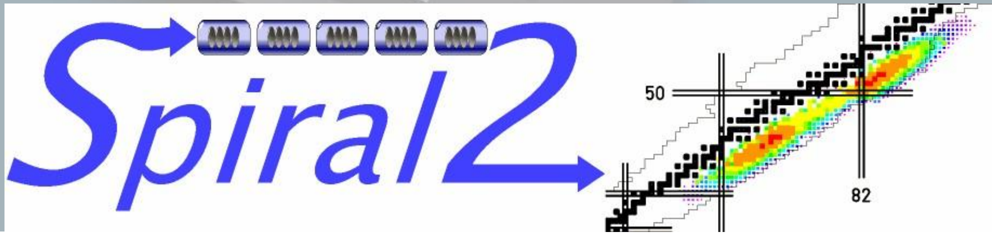

ELI

- A distributed research infrastructure comprised of four facilities
  - Czech Republic, Hungary and Romania.

  - Location of the fourth facility still to be defined
- Provision of ultra-intense and ultra-short laser pulses
- Each facility will have up to 10 instruments
- Peak data rate is 300 MB/shot sustained average of 3GB/s
- Early stage of development
  - Computing model has not been designed
  - Data archiving is needed
    - Federation of those archives is required
- Users proposal to be reviewed by a single selection c

ESS

- Pulsed neutron source similar to the ILL
- Operational in 2019 with 7 instruments
  - Expanding to 22 instruments by 2025
- Data Management, Computing and Software Centre (DMSC)
  - Instrument control
  - Data acquisition and storage
  - Data analysis
  - Instrument simulation
- Peak data rate for a single source will be 400 MB/s

FAIR

- First beam is expected in 2018
- 14 Experiments
  - Both collaborations and user visits
- Computing strategy not finalised
  - Computing models developed using experience from HEP
- Estimate 200K cores and 30 PB storage for the first year
- Two experiments similar to other High Energy Physics (HEP)
  - Clear separation between stages of processed and raw data
- A novel trigger-less detector read-out
  - Data stream exceeding 1TB/sec.

SLHC

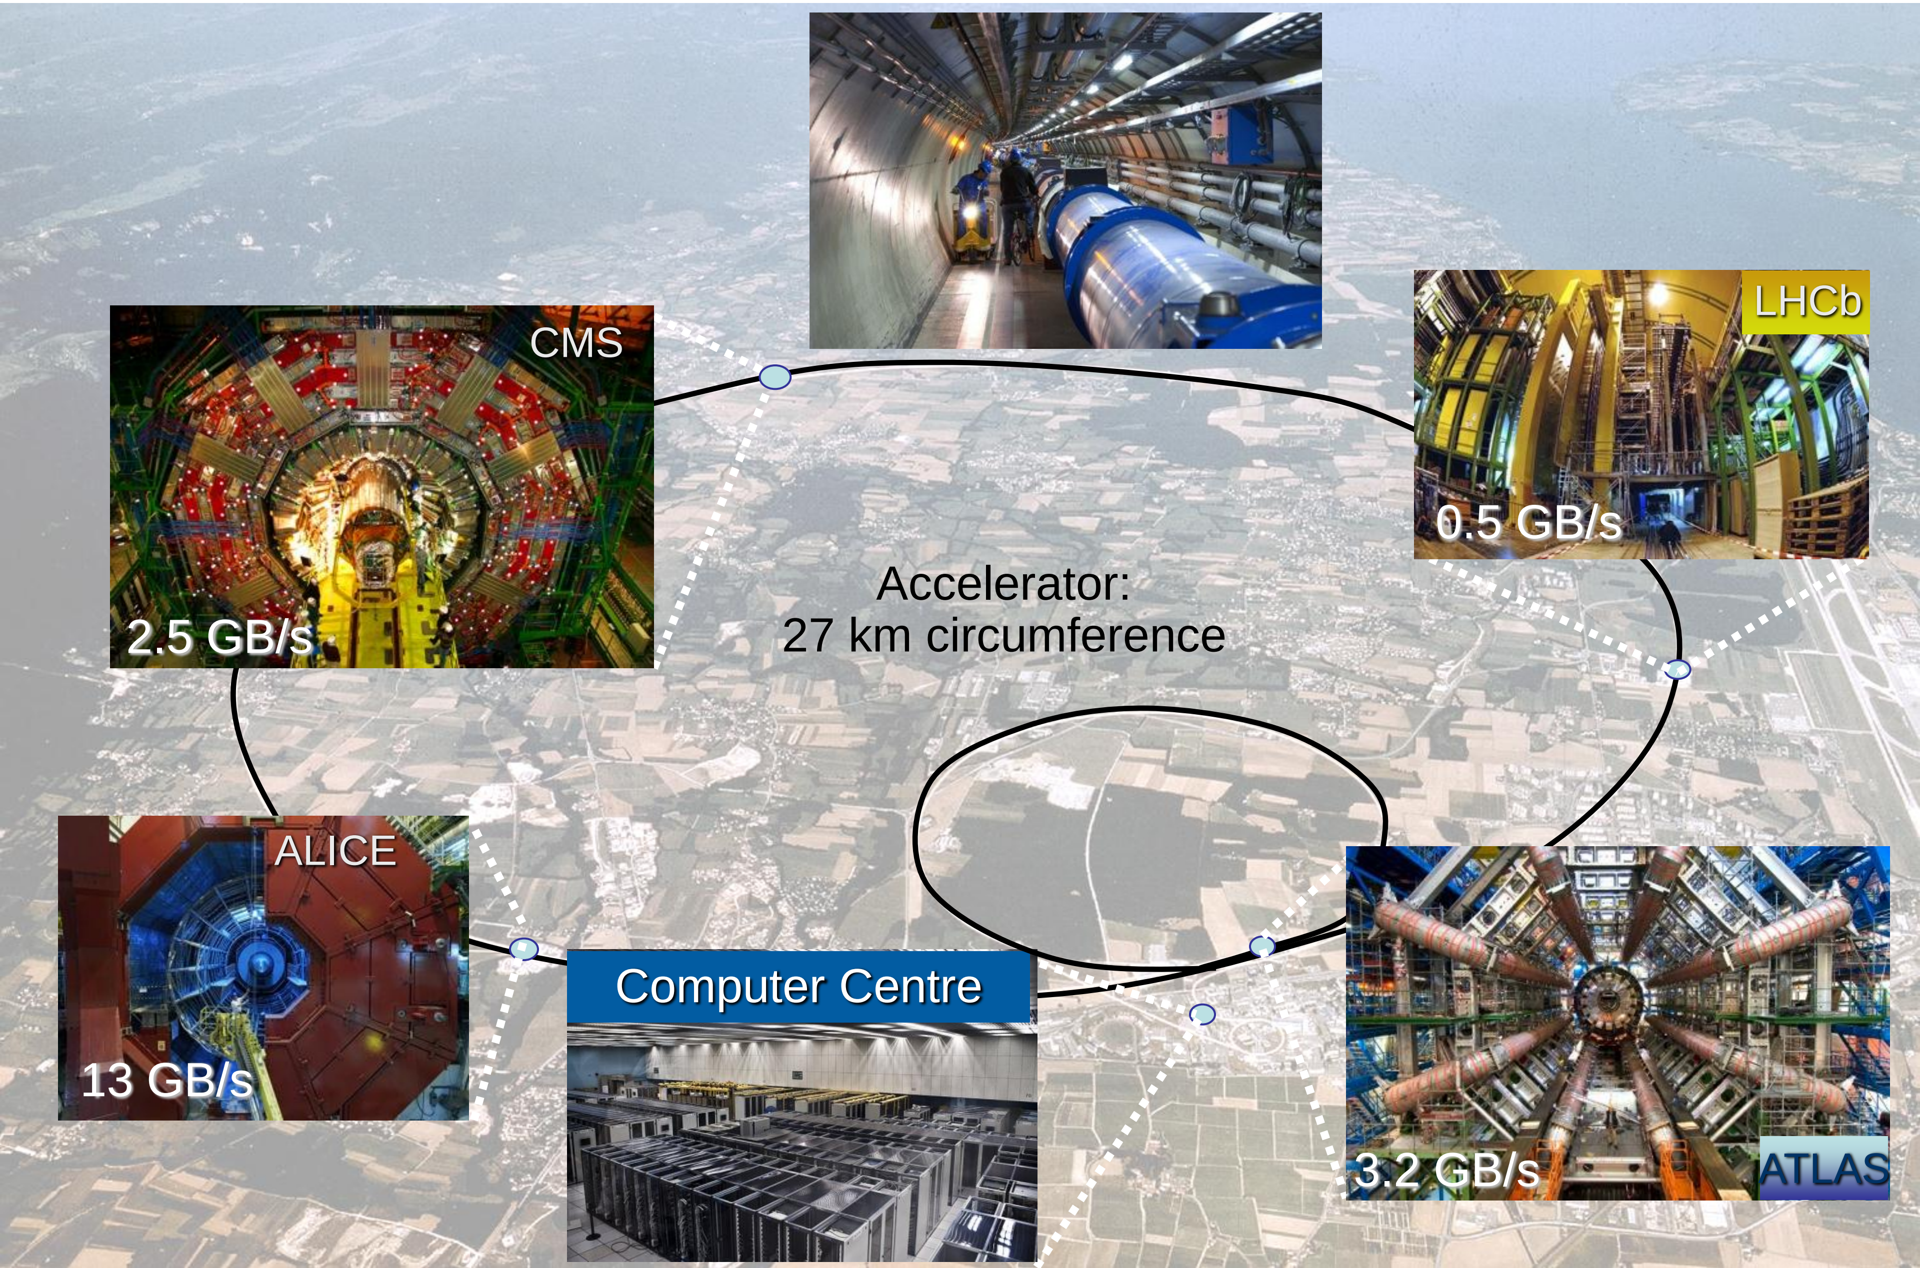
CMS
2.5 GB/s
LHCb
0.5 GB/s
Accelerator:
27 km circumference
ALICE
13 GB/s
Computer Centre
3.2 GB/s
ATLAS

WLCG

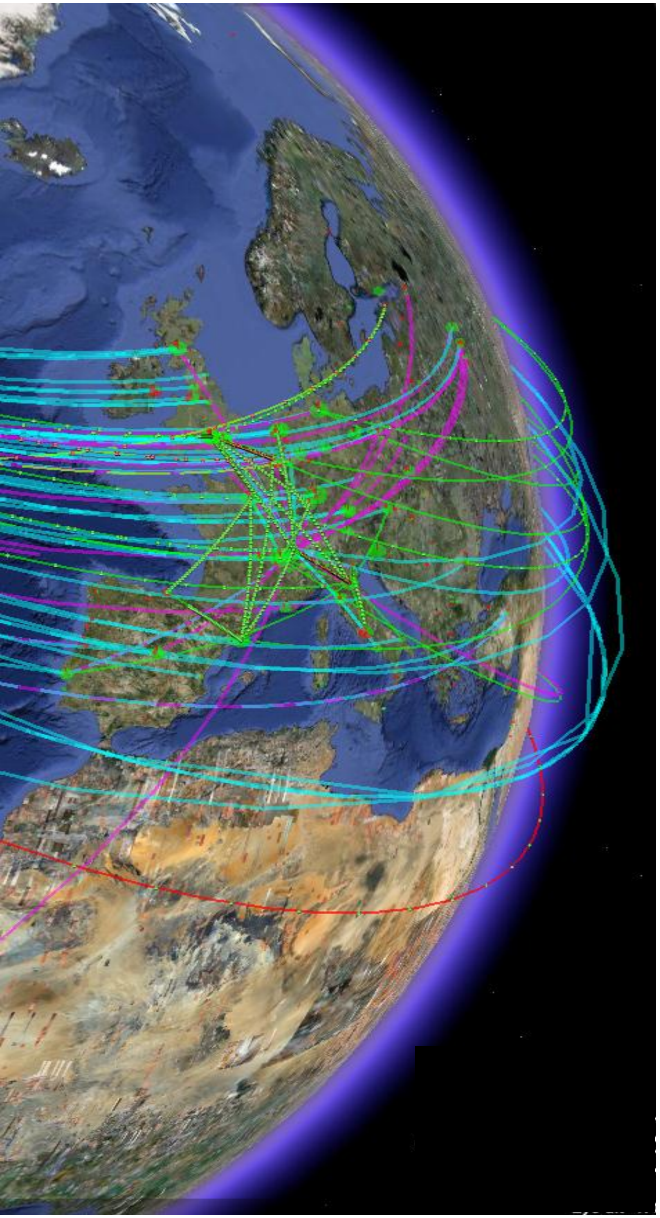

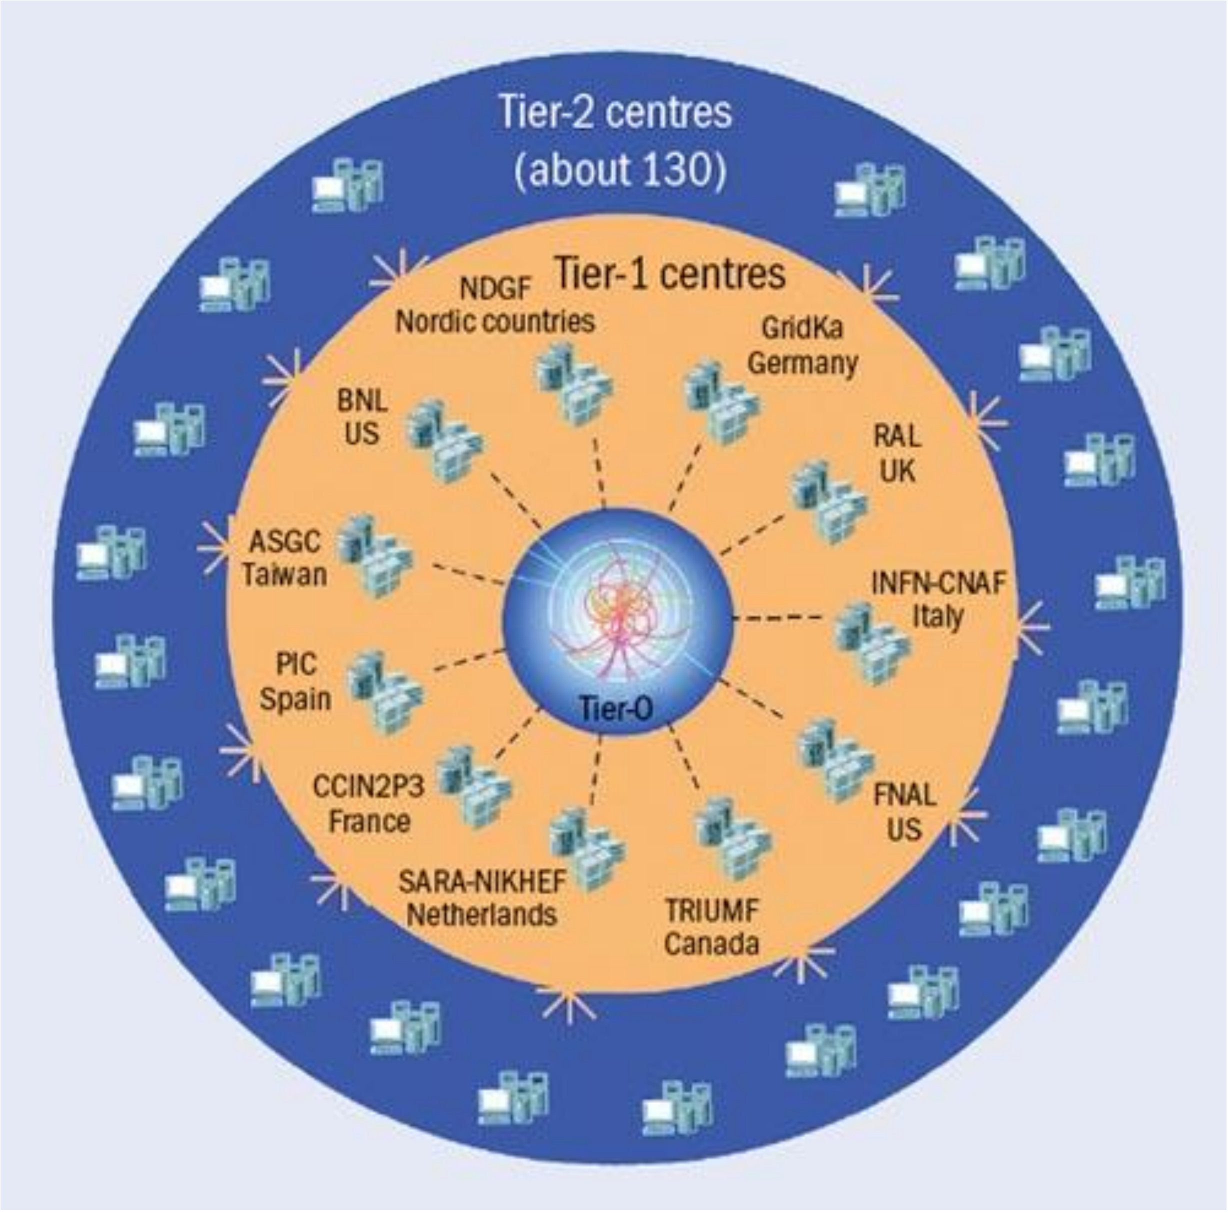
Tier-2 centres
(about 130)
Tier-1 centres
NDGF
Nordic countries
GridKa
Germany
BNL
US
RAL
UK
ASGC
Taiwan
INFN-CNAF
Italy
PIC
Spain
Tier-0
CCIN2P3
France
FNAL
US
SARA-NIKHEF
Netherlands
TRIUMF
Canada

ILC

- Next major high-energy physics facility
  - Collide electrons and positrons at energies of 500 GeV
    - Upgradeable to 1GeV
  - Search for funding and a host country is under way
  - Estimate 8 years for construction with first beams in the 2020s
- Two detectors
  - International Large Detector (ILD)
  - Silicon Detector (SiD)
- Too early to discuss the computing strategy
  - Requirements similar to LHC

SKA

- A global project to build the world’s largest radio telescope
  - Supporting survey or single observer operation
- Led by the SKA Organisation, headquarters in Manchester, UK
- Will be physically built on two sites; Australian and South African
  - Envisaged to operate separately with a limited network connectivity
- Science Computing Facility
  - Will operate core processing including the main data archive
  - Processing Data Store requirements
    - Maximum write data rate: 330 GB/s
    - Maximum read data rate:1650 GB/s
    - Storage duration: 5 hours
    - Storage capacity: 6 PB
  - Data rate of 114GB/obs
- The two data archives will operate independently
  - Replicate data to Regional Science Centres

IT&DM Work Packages

User Identity

Metadata Catalogues

High-speed Data Recording

Distributed Data Infrastructure

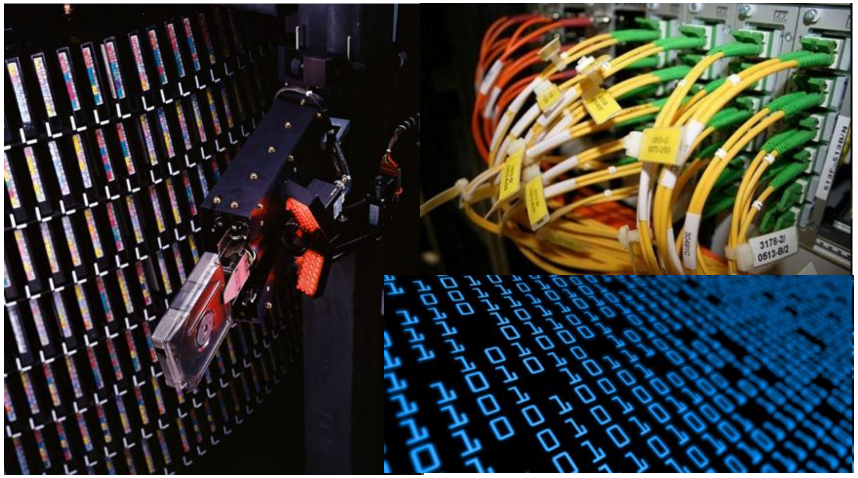

EuroFEL Example

- Users can perform experiments at various facilities
  - Need to combine data stored at those facilities
- Challenges
  - Combination of different data sources
  - Different data formats and policies
- Require transparent access to data and compute resources
  - Based on a cross-facility authentication and authorization infrastructure
- Metadata-model needs to be standardized
  - Data catalogues need to be federated
  - Combine data from different places seamlessly

Umbrella

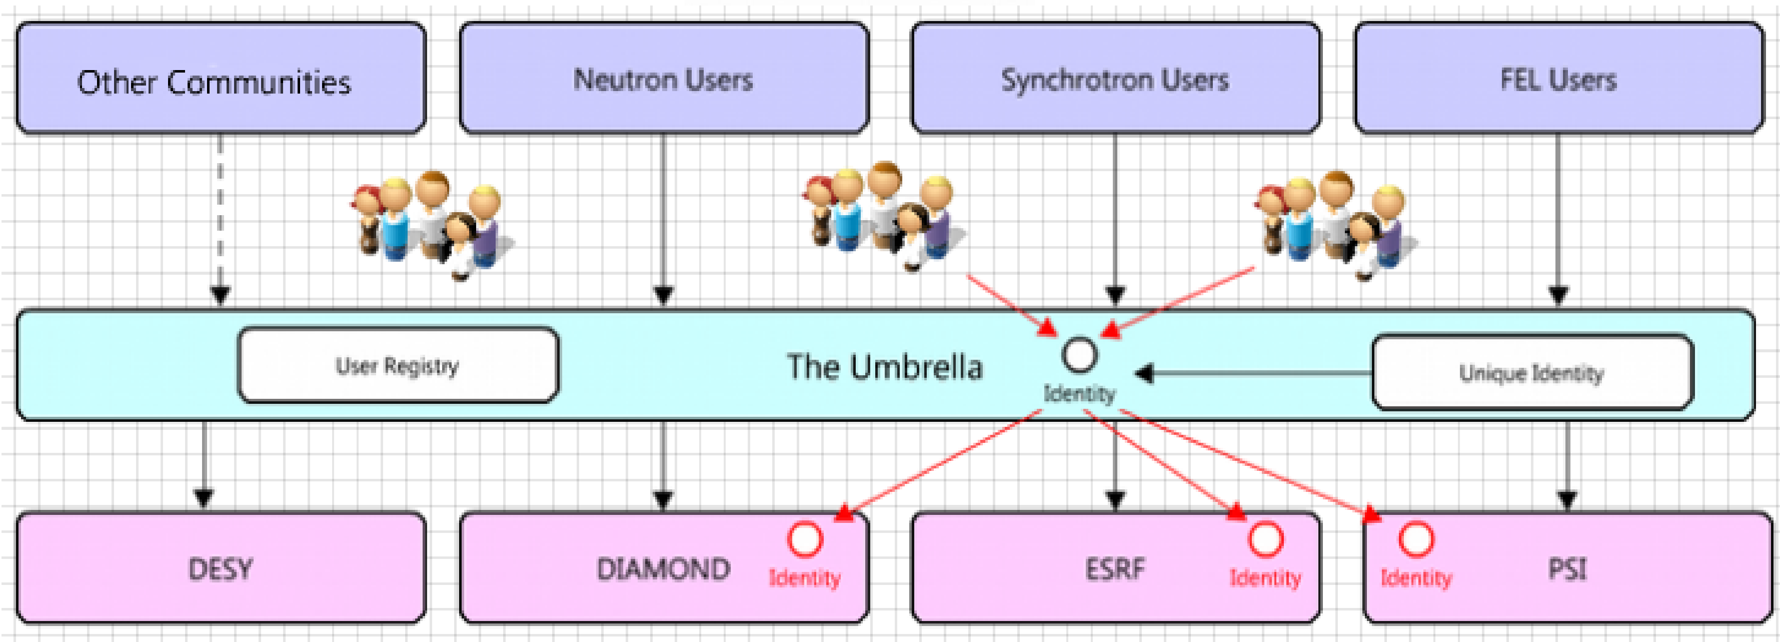
Other Communities
Neutron Users
Synchrotron Users
FEL Users
User Registry
The Umbrella
Identity
Unique Identity
DESY
DIAMOND
Identity
ESRF
Identity
Identity
PSI

ICAT

ICAT instances in the world

Photon and Neutron Data Infrastructure

A user can use the TopCAT frontend client installed at any one of the facilities to access data across any exposed ICAT instance in the world.

TopCAT frontend client at ILL, Grenoble, France

Jamie Hall, Institut Laue-Langevin, Grenoble, France.

Recording and Archiving

Date	Collaboration	Data Volume	Data Rates	Technology
1950's	2-3	Kilobits	W/M	notebooks
1960's	10-15	kB	KB/M	punchcards
1970's	~35	MB	B/s	tape
1980's	~100	GB	KB/s	tape, disk
1990's	700-800	TB	MB/s	tape, disk
2010's	~3000	PB	GB/s	tape, disk

Large Electron–Positron Collider
Operational bewteen1989-2000
Complete data set is a few TB

Large Hadron Collider (LHC)
1 year of LHC data is 25 PB

Distributed Data Infrastructure

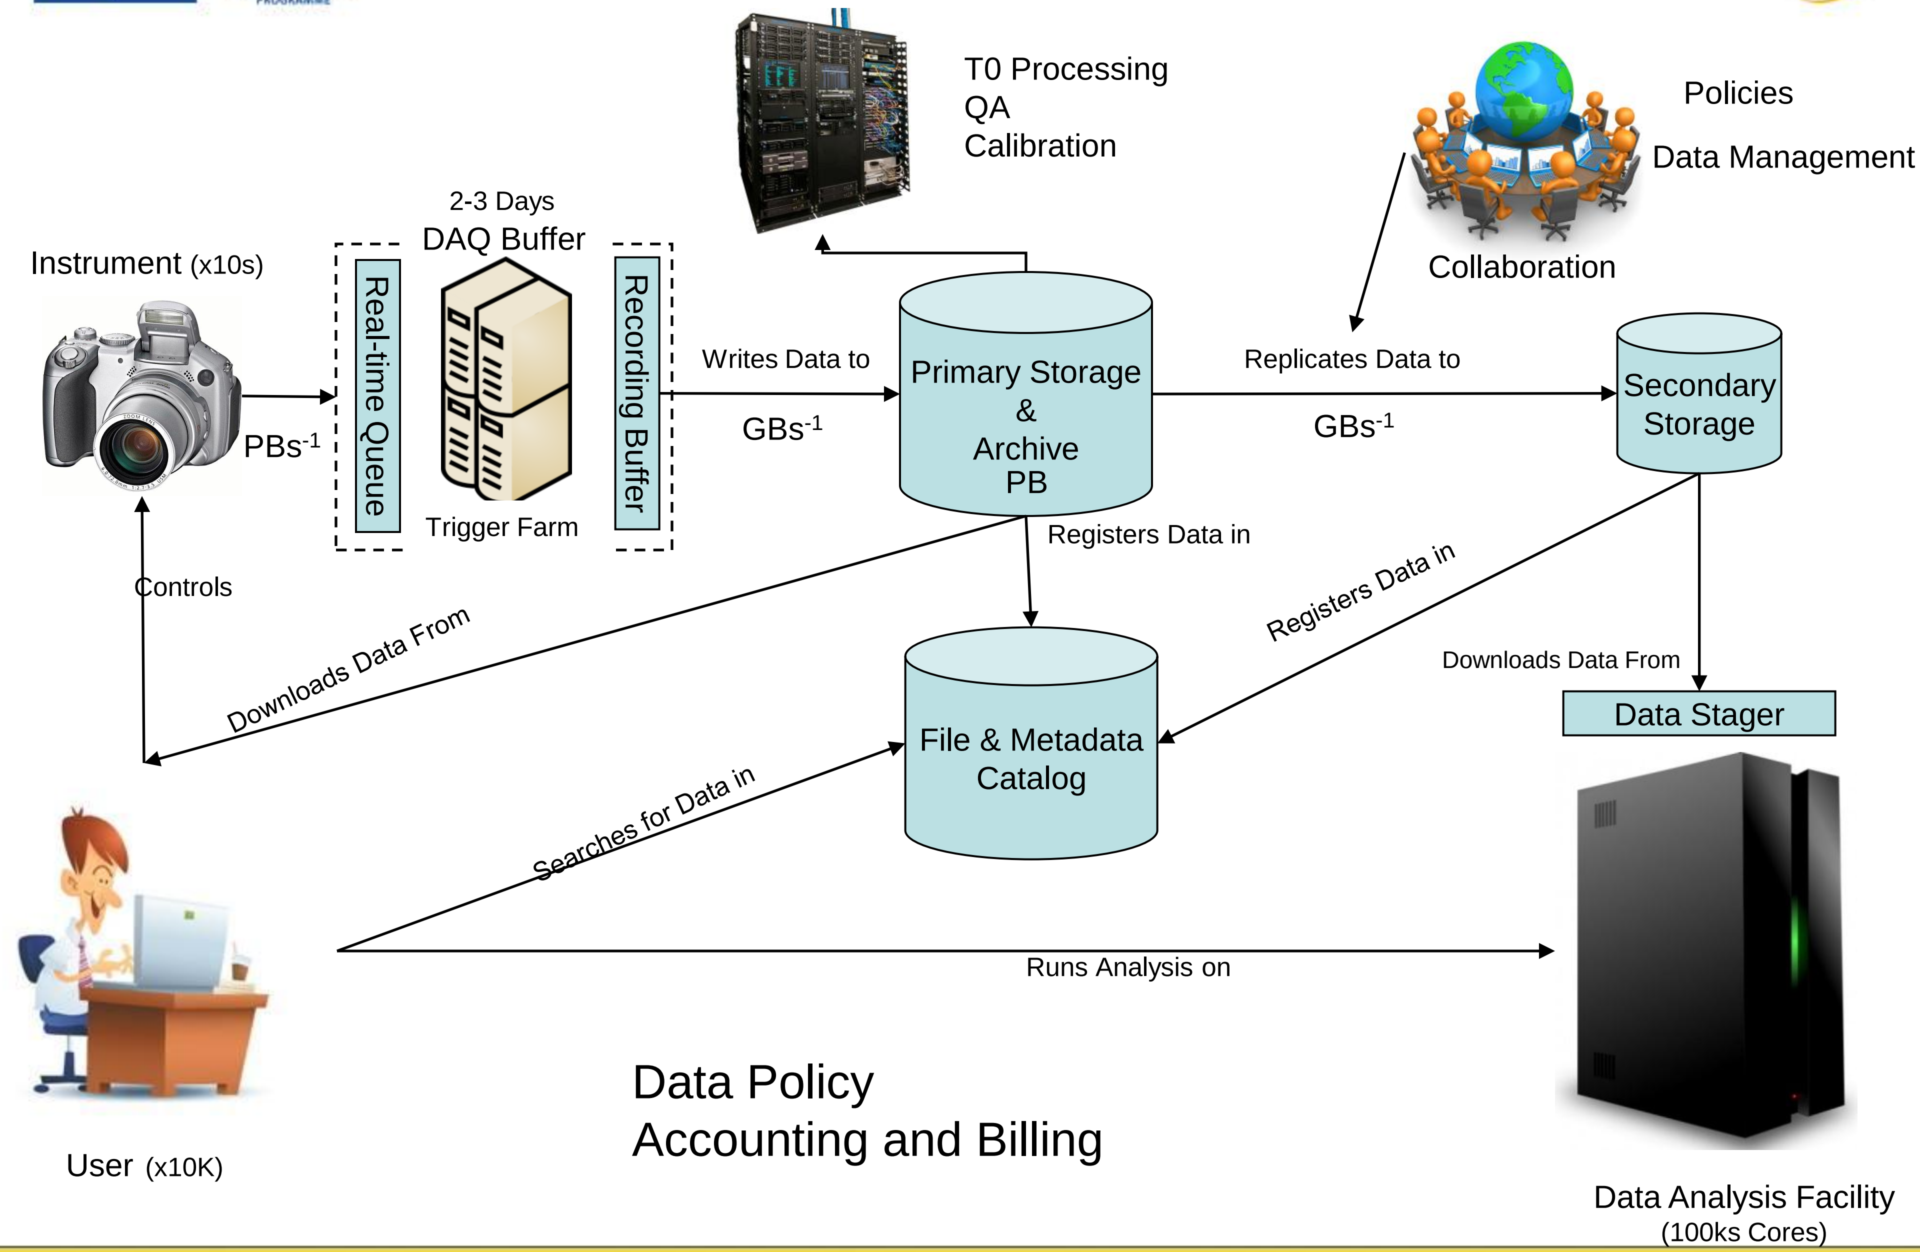

ESFRI Projects

- European Strategy Forum on Research Infrastructures
  - Aims to improve the integration of RIs
    - Providing a coherent approach to development
  - CRISP
    - Digital images
  - ENVRI
    - Environmental Science
      - Satellite data and sensor networks
  - BioMedBridges
    - Biological and biomedical domains
      - Medical imaging and Genome data
  - DASISH
    - Social sciences and humanities
      - Cultural heritage and survey data

Areas of Collaboration

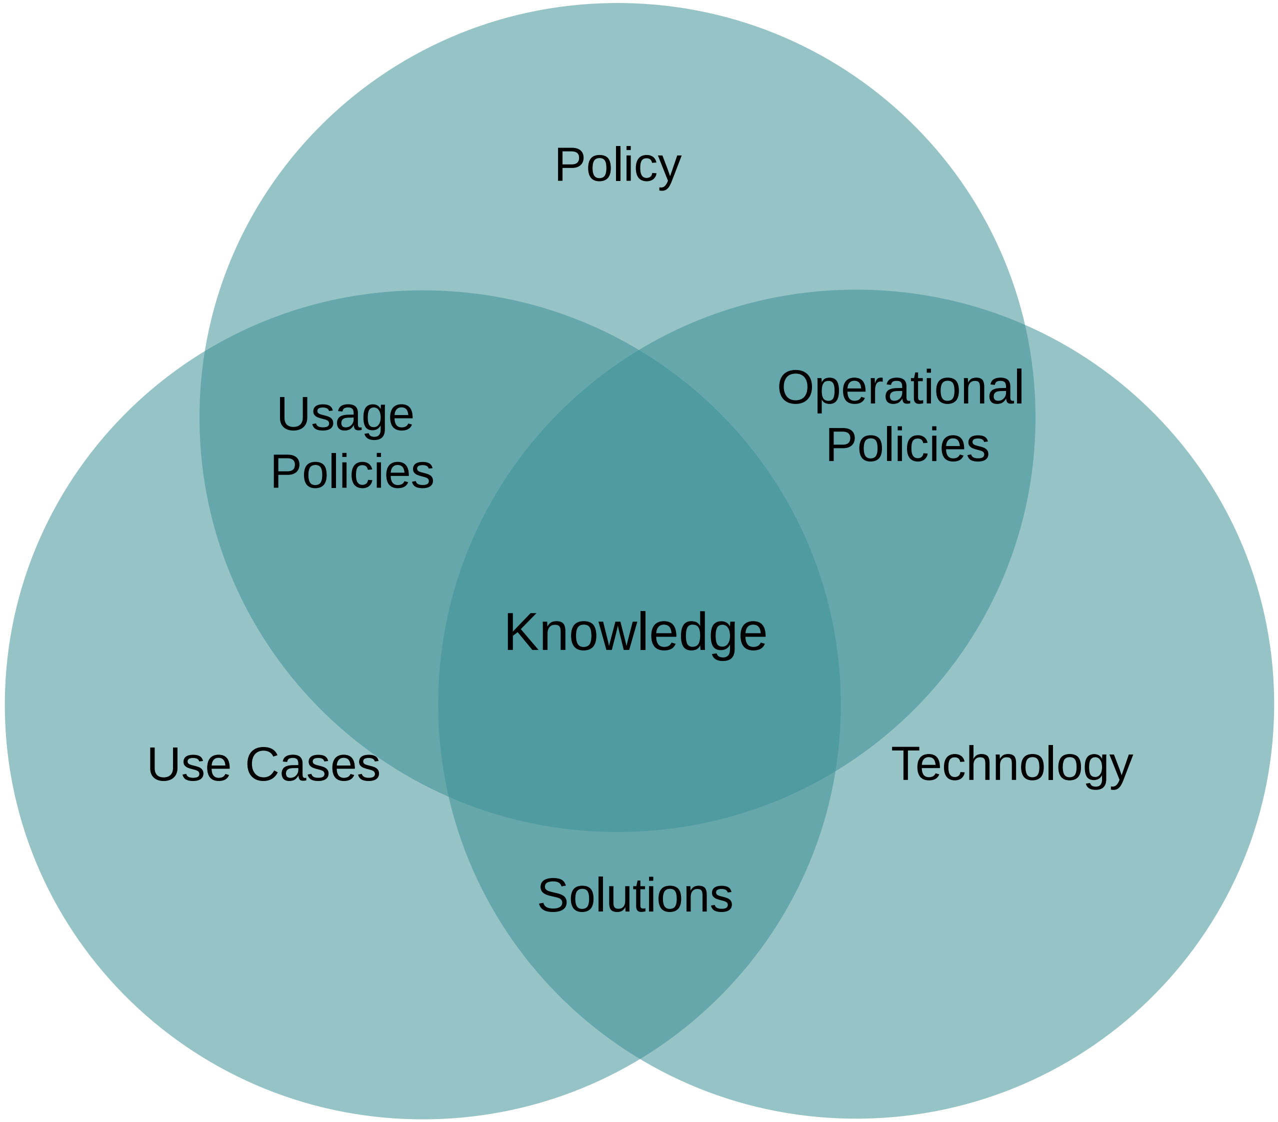

Shared Challenges

- Identity Management
  - Federated Identify Management Initiative
- Data Archiving and Preservation
  - Persistent Identifiers
  - Data Formats
  - Common Solutions
  - Ensuring Viability and Relevance
- Data Discovery
  - Catalogs and Registries
- Data Access
  - Standard Protocols
- Policies
  - Data
  - Infrastructure
  - Collaboration

Summary

- 11 Physics Research Infrastructures
  - Digital images of the very big and very small
    - With associated metadata
- Common computing models
  - User visits and collaborations
- CRISP Investigation four areas
  - Identity management
  - Metadata management
  - High-speed data recording
  - Distributed Data Infrastructure
- Potential collaboration with other ESFRI projects
  - Use Cases
  - Policy
    - Identity management
  - Technology
    - Data archiving and preservation