From Physics to Industry: EOS outside HEP

Xavier Espinal
on behalf of IT/ST

SLAC
CHEP 2016

t=0

The Big Bang

1 second

10^{-10} seconds

10^{-34} seconds

10^{-43} seconds

10^{32} degrees

10^{27} degrees

10^{15} degrees

10^{10} degrees

10^{9}

radiation

particles

W^+ W^- Z heavy particles carrying the weak force

quark

anti-quark

electron

positron (anti-electron)

proton

neutron

meson

H hydrogen

D deuterium

He helium

Li lithium

DI-9112020_03

t=(yr)2007

The Washington Post

Exabytes: Documenting the 'digital age' and huge growth in computing capacity

THE WORLD'S CAPACITY TO STORE INFORMATION

This chart shows the world's growth in storage capacity for both analog data (books, newspapers, videotapes, etc.) and digital (CDs, DVDs, computer hard drives, smartphone drives, etc.)

In gigabytes or estimated equivalent

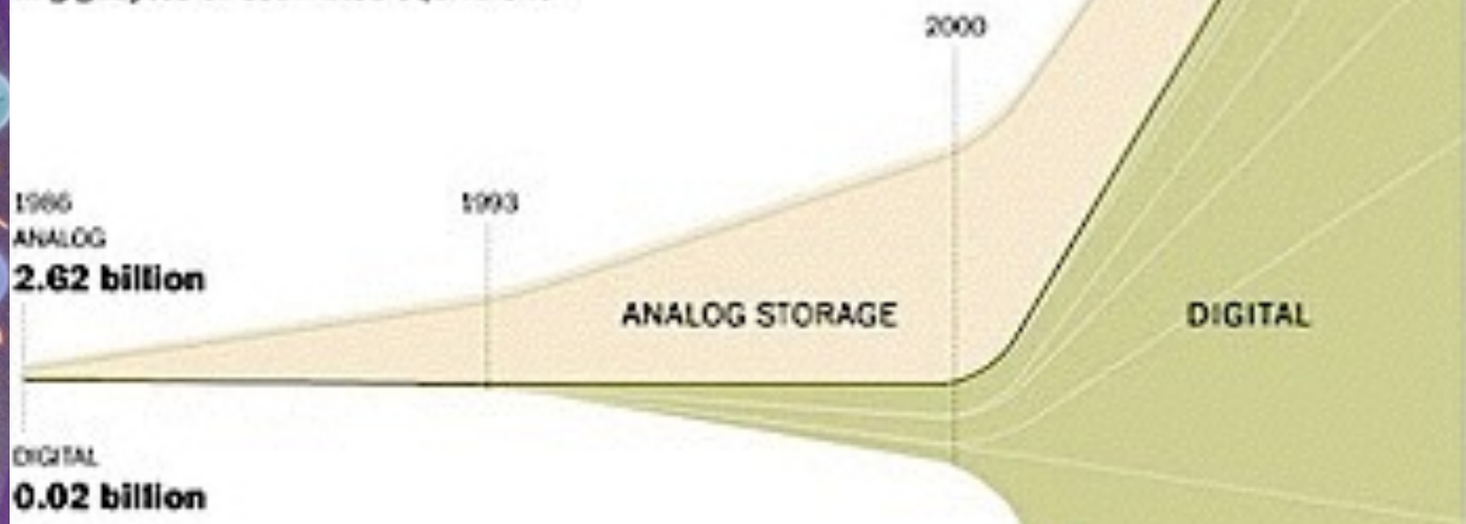

COMPUTING POWER

In 1986, pocket calculators accounted for much of the world's data-processing power.

Percentage of available processing power by device:

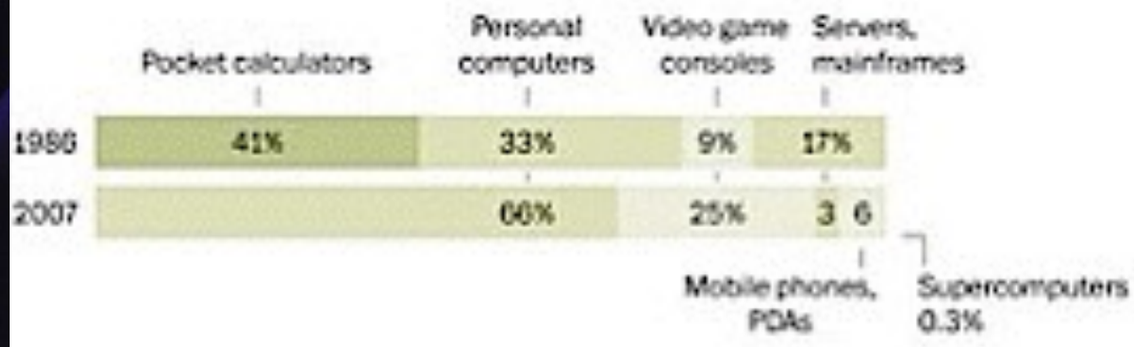

Data production is outdistancing storage solutions

Nowadays' *normal* data volumes are difficult to handle

DIY solutions not scaling: +$$$ ≠ +TBs

'a PetaByte' is not remarkable anymore

Physics pioneered Large Scale Computing

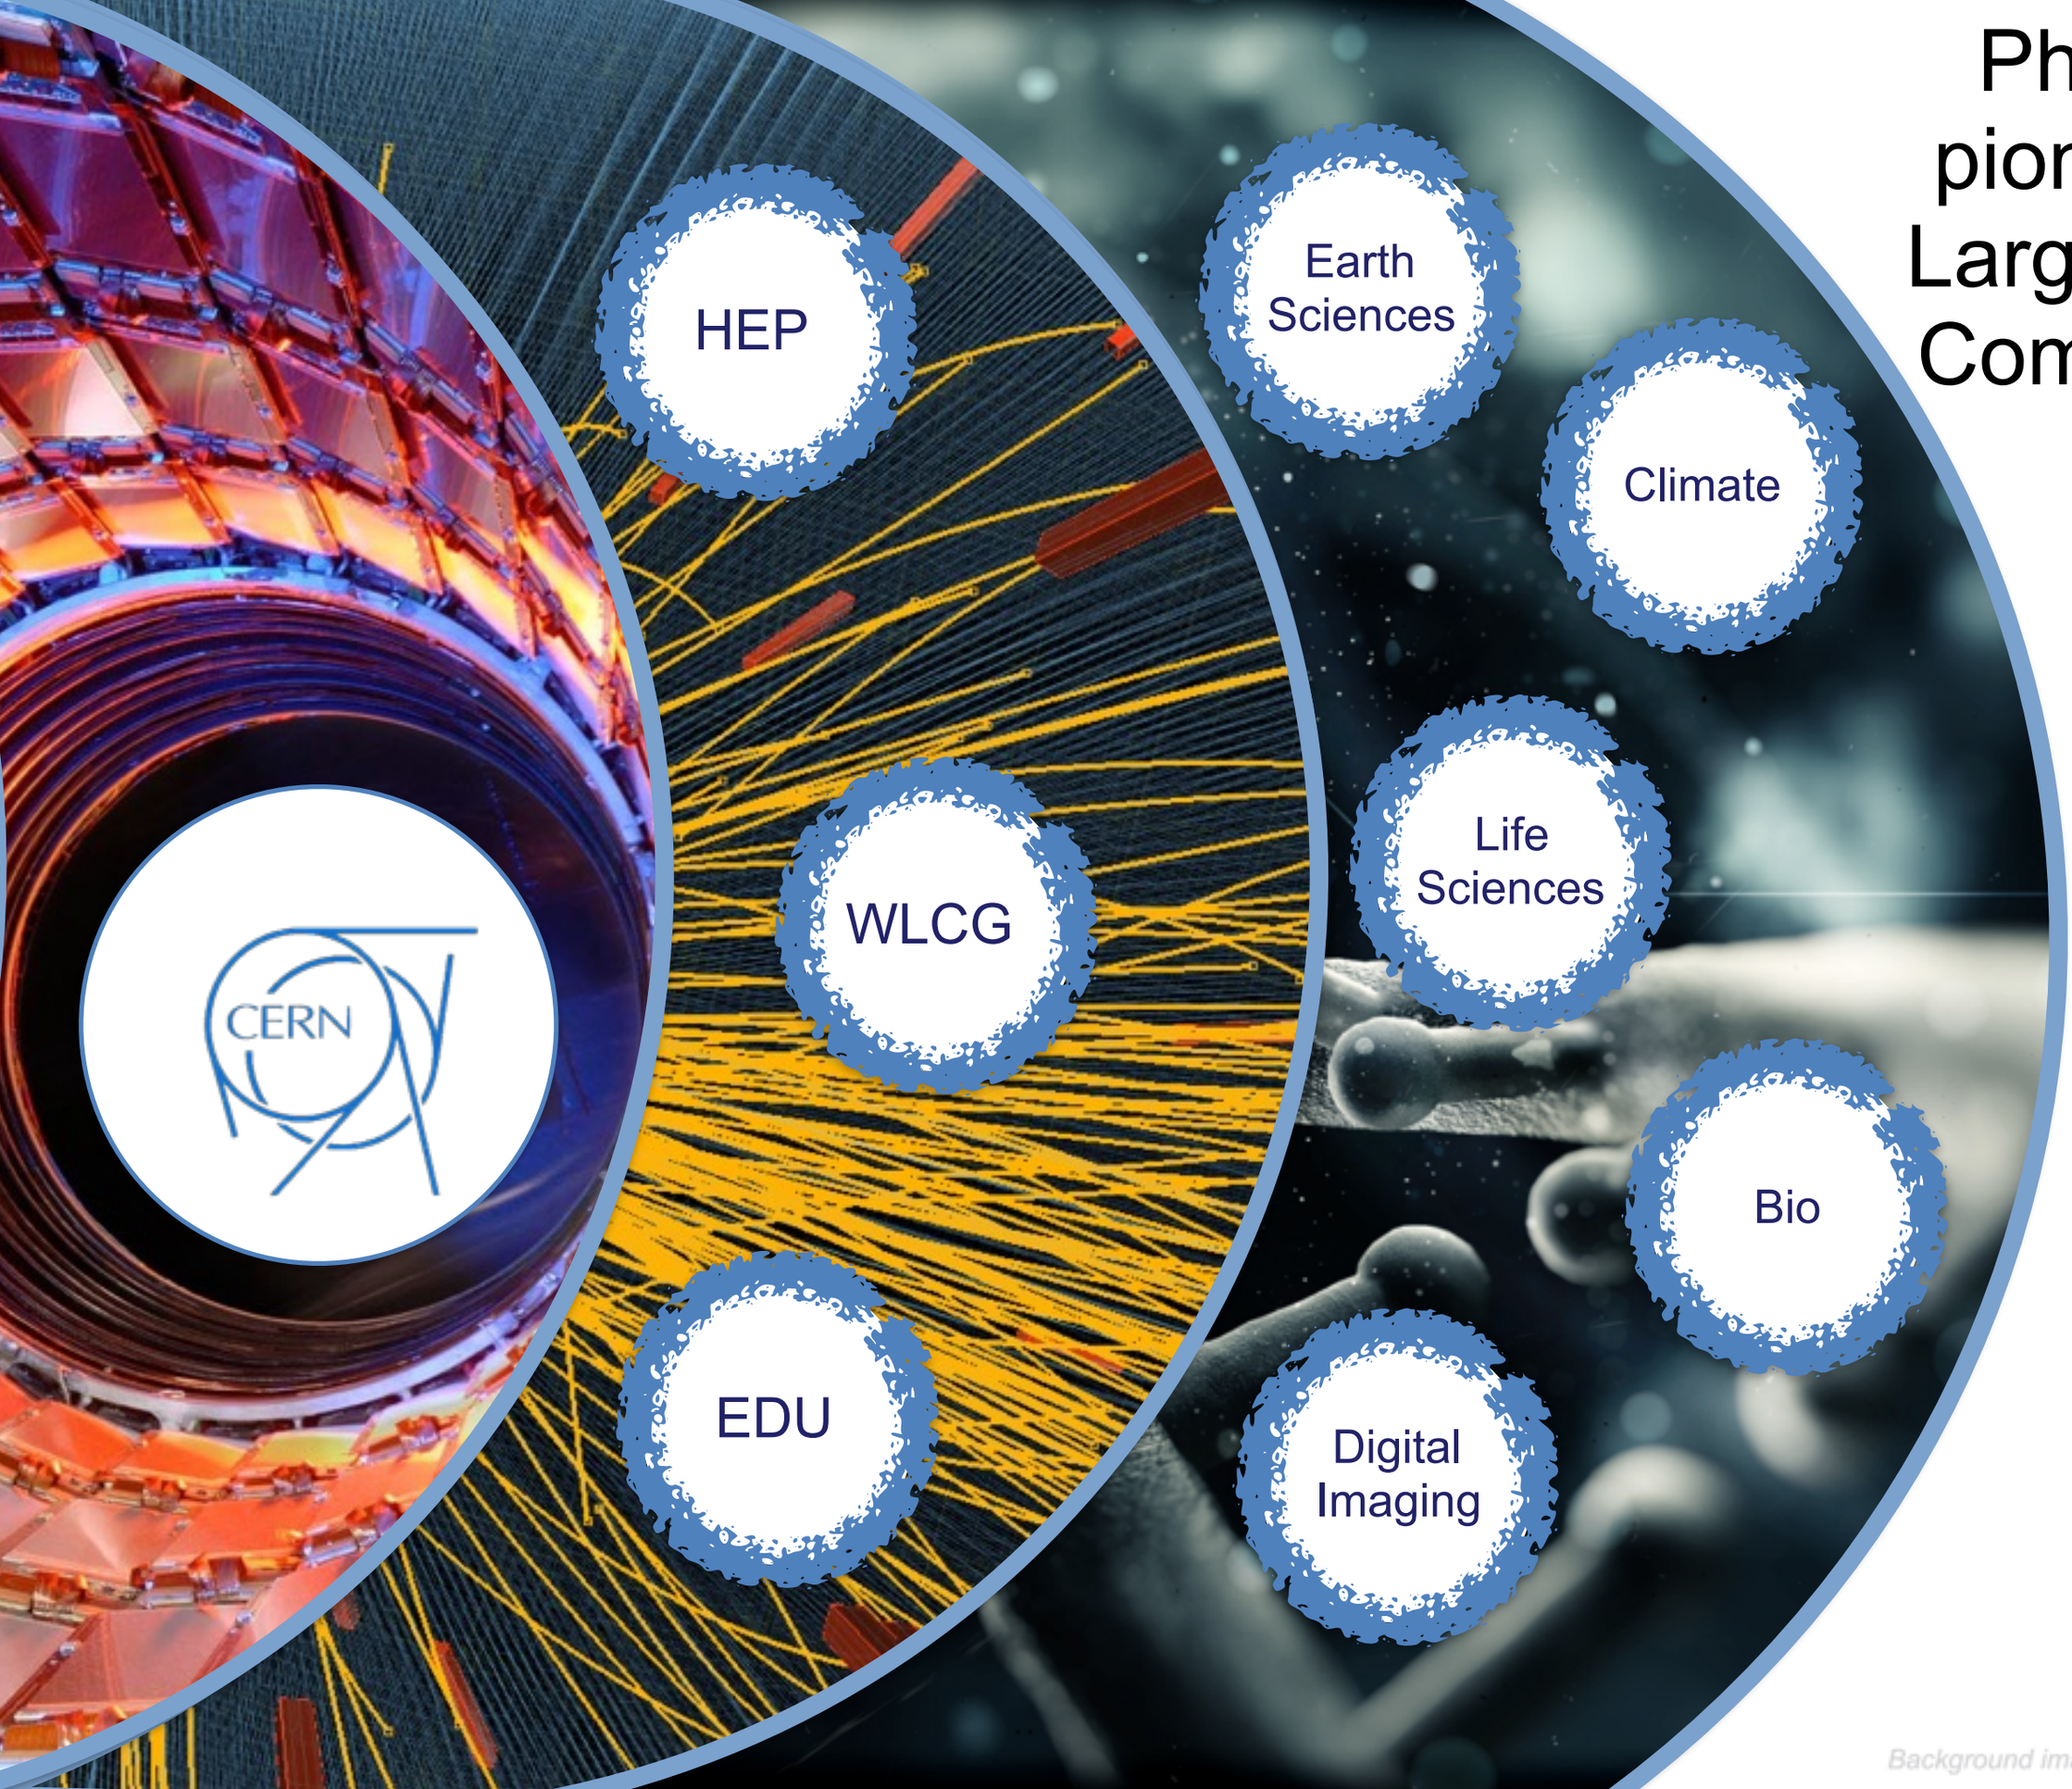

Background image: Shutterstock

CERN

HEP

WLCG

EDU

Earth Sciences

Climate

Life Sciences

Digital Imaging

Bio

Industry

Product

Knowledge

Enterprise

Computing technologies that can percolate from High Energy Physics to Industry is a major success for CHEP community and science

WLCG successfully exposed new technologies. Some still *hidden* but some being *unveiled*: Workload Management Systems, Data Management Frameworks and Storage Technologies

Interfaced **CERN storage services** with **Blue Waters NCSA** using **WLCG's FTS3** to manage the data workflow

Open door for HPC environments to link with our HTC and Distributed Computing expertise

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

SIGN IN NCSA

YOUR BLUE WATERS ABOUT SCIENCE AT BLUE WATERS USING BLUE WATERS EDUCATION & TRAINING NEWS & EVENTS HELP

Mapping Proton Quark Structure in Momentum and Coordinate Space using PetaByte Data-Sets from the COMPASS Experiment at CERN.

Computing technologies that can percolate from High Energy Physics to Industry is a major success for CHEP community and science

WLCG successfully exposed new technologies. Some still *hidden* but some being *unveiled*: Workload Management Systems, Data Management Frameworks and Storage Technologies

HEP community was already data hungry before the *commodity-data explosion* and we developped solutions to cater

EOS as a Large Scale Storage System is opening a new door. HEP-designed but interesting for the *outside world*


Not Your Average
BROWNIES
OUR BROWNIES ARE packed WITH AS
MUCH OF THE FINEST CHOCOLATE AS WE COULD
POSSIBLY FIT IN THE PAN FOR A MOIST,
RICH chocolatesplosion IN EVERY BITE.
DECADENT? YES. WORTH EVERY BITE? Absolutely.
SOMEONE GRAB THE ICE CREAM.
BAKED BY:
Michael

Computing technologies that can percolate from High Energy Physics to Industry is a major success for CHEP community and science

WLCG successfully exposed new technologies. Some still *hidden* but some being *unveiled*: Workload Management Systems, Data Management Frameworks and Storage Technologies

HEP community was already data hungry before the *commodity-data explosion* and we developped solutions to cater

EOS as a Large Scale Storage System is opening a new door. HEP-designed but interesting for the *outside world*

Together with CERN Openlab we started a project with a company (COMTRADE) to start the productisation of EOS

metadata

deployment

consistency

resilience

operational
$$$

maintenance
$$$

redundancy

implementation
$$$

scalability

namespace

QoS
flexibility

book
keeping

"The project scope is the evolution of the EOS system in the direction of simplified usage, installation and maintenance and to extend its utilisation by adding new supported platforms. In the initial phase the emphasis will be in providing a robust installation kit to allow rapid installation of EOS on an agreed set of platforms. The kit will include the necessary installation instructions and tools for operations (admin guide) and for user (user guide). A test suite will exercise the native EOS interface (xroot) and the main access protocols (Fuse, Webdav/HTTP)

evolution of the EOS system

simplified usage, installation and maintenance

providing a robust installation kit

installation instructions and tools for operations

CERN's Disk-only Large Scale Storage System

EB era

Performant and manageable

Easily scalable

CERN main storage platform

Experts in-house

Adapt when required

Re-design if needed

Community storage

Collaboration

CERNBox

Share

Offline work

Sync

Bridge BigData PersonalData

Adaptable
Catering with different uses

Data processing

User Analysis

LHC Data Recording

Sync&Share

EOS now

October 2016

+1200

+45000

850M files

150PB

IT-ST

Phase-I :: provide tools and expand

accomplished

Installation shell script:
- Fully compliant EOS instance
- Headnodes (master and slave)
- Storage nodes

Summary of resources installed and status:
- Servers, available space, etc.
- Functional tests

Uninstallation and rollback capabilities

Documentation
- EOS whitepaper
- Installation script how-to
- Administration
- System description

Acces/authentication:
- Kerberos, LDAP integration, shared secrets

Build platform:
- Gitlab > Jenkins > Docker

Support for different linux distributions

Phase-I :: provide tools and expand

accomplished

aarnet
Australia's Academic and Research Network

European Commission
JOINT RESEARCH CENTRE
The European Commission's in-house science service

universität wien

CERN openlab

UNIVERSITY OF BELGRADE
INSTITUTE OF PHYSICS BELGRADE

中国科学院高能物理研究所
Institute of High Energy Physics
Chinese Academy of Sciences

UNIVERSITÀ
DEGLI STUDI DI TRIESTE

INFN
ISTITUTO NAZIONALE
di FISICA NUCLEARE
sezione di BARI

aarnet
Australia's Academic and Research Network

- Scale-out filesystem underneath the ownCloud app, using the eosd fuse interface for file IO

aspera RSYNC

- Geo-distributed setup: Brisbane, Melbourne, Perth
  - ~1PB (scale to ~20PB next year)

- Australian National University, in Acton Canberra: mirror archives of both genome sequences and open or freely available software distributed among three sites

Australian National University

"This system is presently running 0.3.187, and has been so trouble free that I keep forgetting it's there." David Jericho -AARNetSolutions Architect

Components of the *JRC Earth Observation Data Processing Platform (JEODPP)*

Batch Processing Interface
direct access interface through HTCondor scheduler for experiment/projects use cases

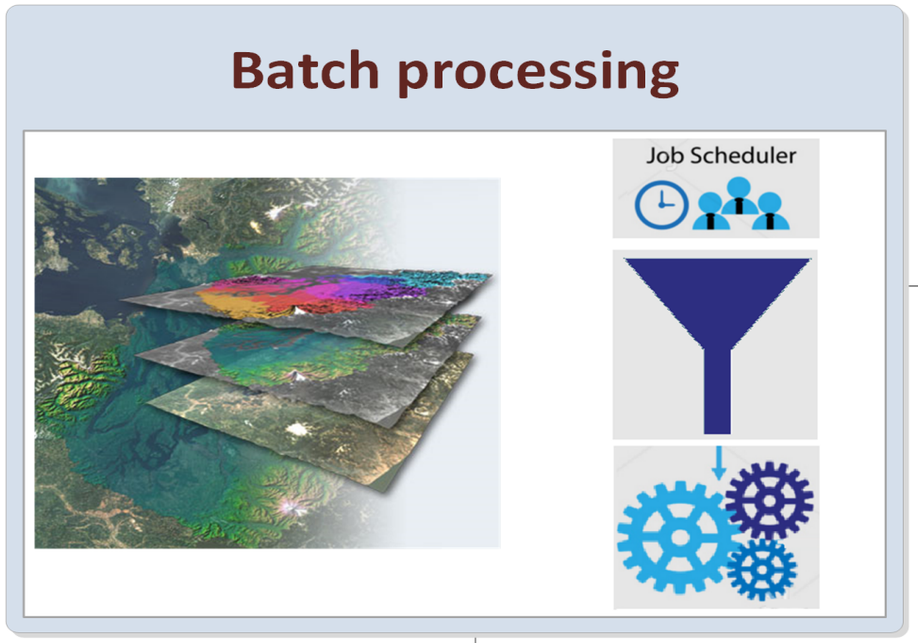

Interactive Processing Interface
web access interface for end users based on Jupyter, Leaflet and custom built image processing libraries

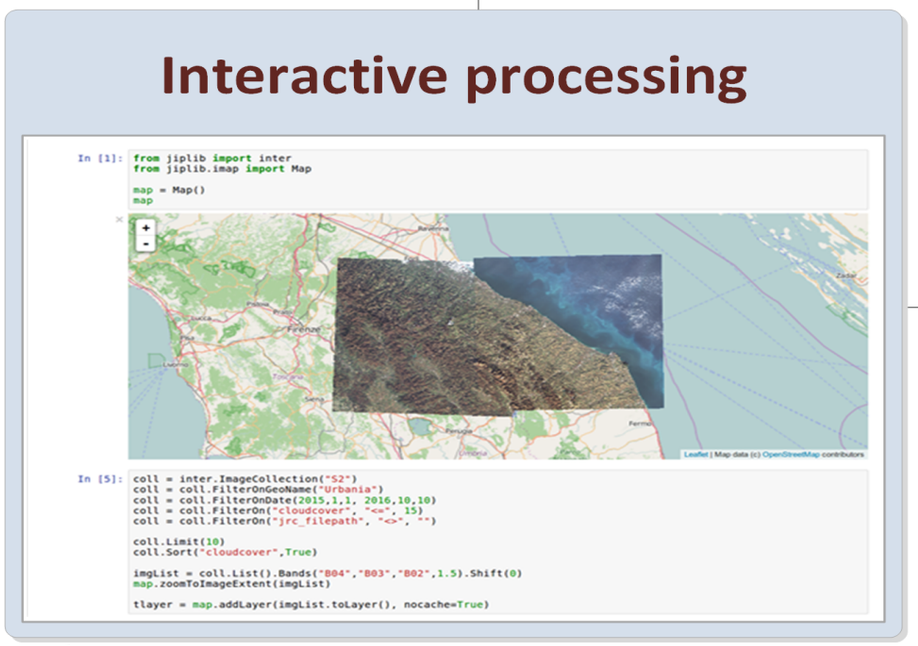

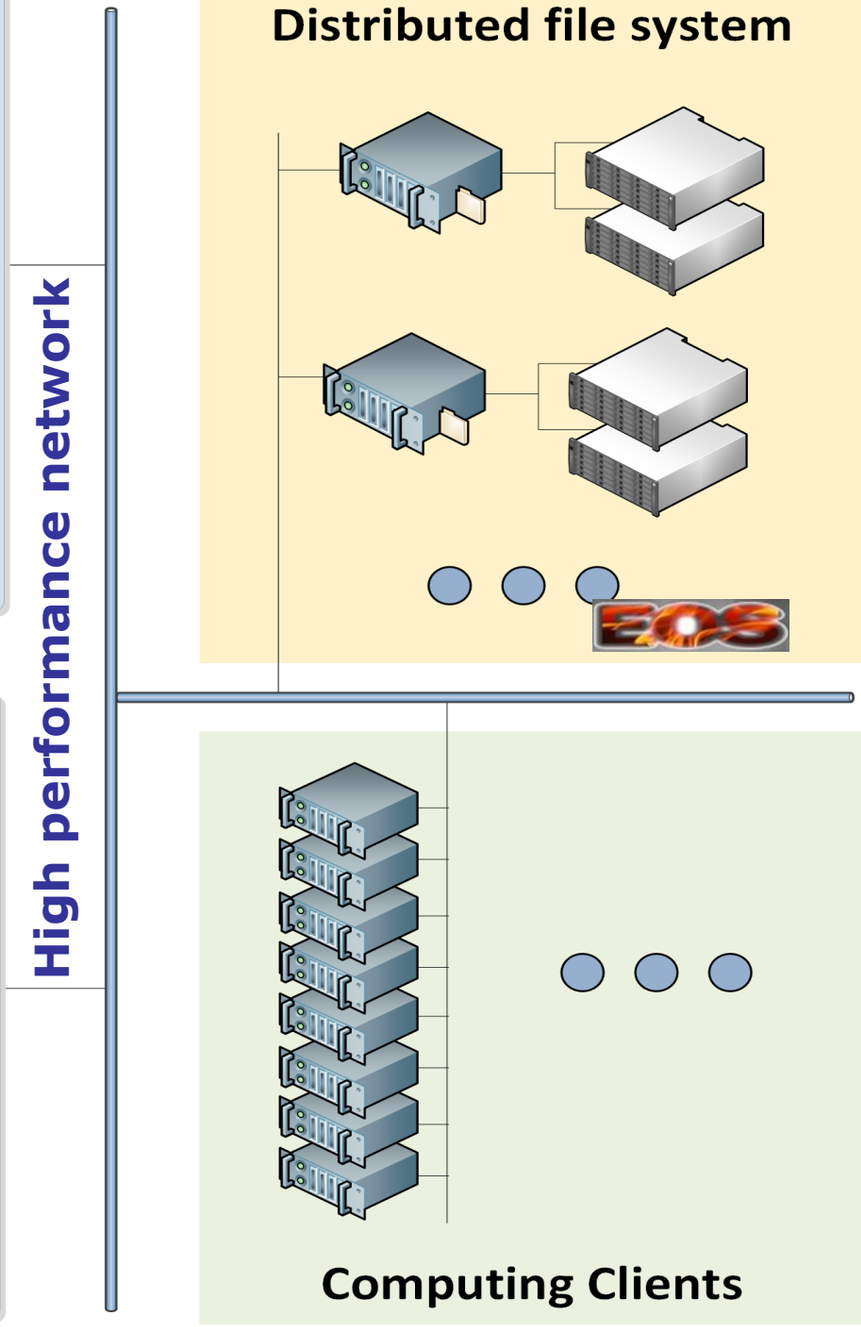

FILESYSTEM
EOS
2 MGM servers
10 FST servers
240 6 TB disks

TOTAL SPACE: 1.44 PB

Current usage:
24.5M files
336TB

PROCESSING
9 2U processing servers
16 1U processing Servers

TOTAL:
600 cores
7+TB memory

European Commission
Joint Research Centre

@V.Vasilev,F.Eyraud (JRC)

Future

Automated testing
- Functional tests
- Benchmarking

Multi-site EOS cluster:
- Geo-scheduling
- RAIN configuration
- Erasure Coding

Continue evolving:
- Installation
- Documentation
- Distro support

Monitoring
- System overview
- Basic metrics

Documentation
- EOS admin guide
- EOS system

Support

Clients
- Native Windows client

Installation/admin
- Docker deployment
- Admin GUI


CERN