

Big data and HPC on-demand: Large-scale genome analysis on Helix Nebula – the Science Cloud

Rupert Lueck

Head of IT Services, EMBL

ISC Cloud '12 Mannheim

24 September 2012

EMBL



**HELIX
NEBULA**
THE SCIENCE CLOUD

EMBL: European Molecular Biology Laboratory



- Intergovernmental Research Organization
- Supported by 20 Member States (+1 associated: )
- One of the world's foremost life science institutions
- EIROforum member
- 1500 staff
>70 nationalities

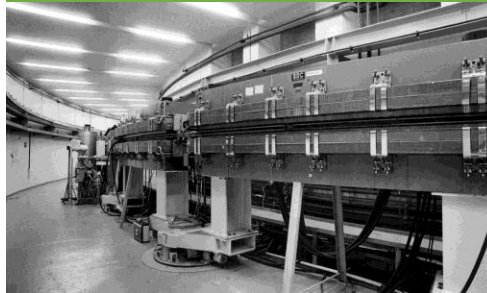
The Five Branches of EMBL

Heidelberg



Basic Molecular Biology
Research
Main Lab / Headquarters

Hamburg



Structural Biology
DESY

Hinxton



European Bioinformatics
Institute (EBI)
Sanger Centre

Grenoble



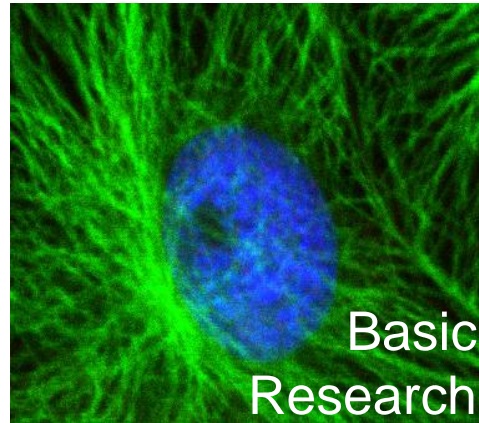
Structural Biology
ILL, ESRF, IBS, UVHCI

Monterotondo

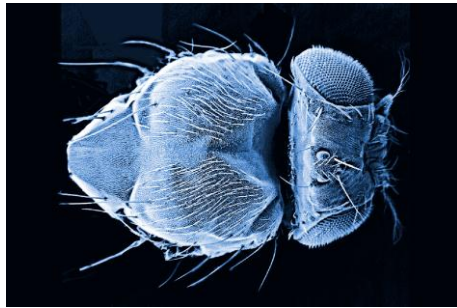
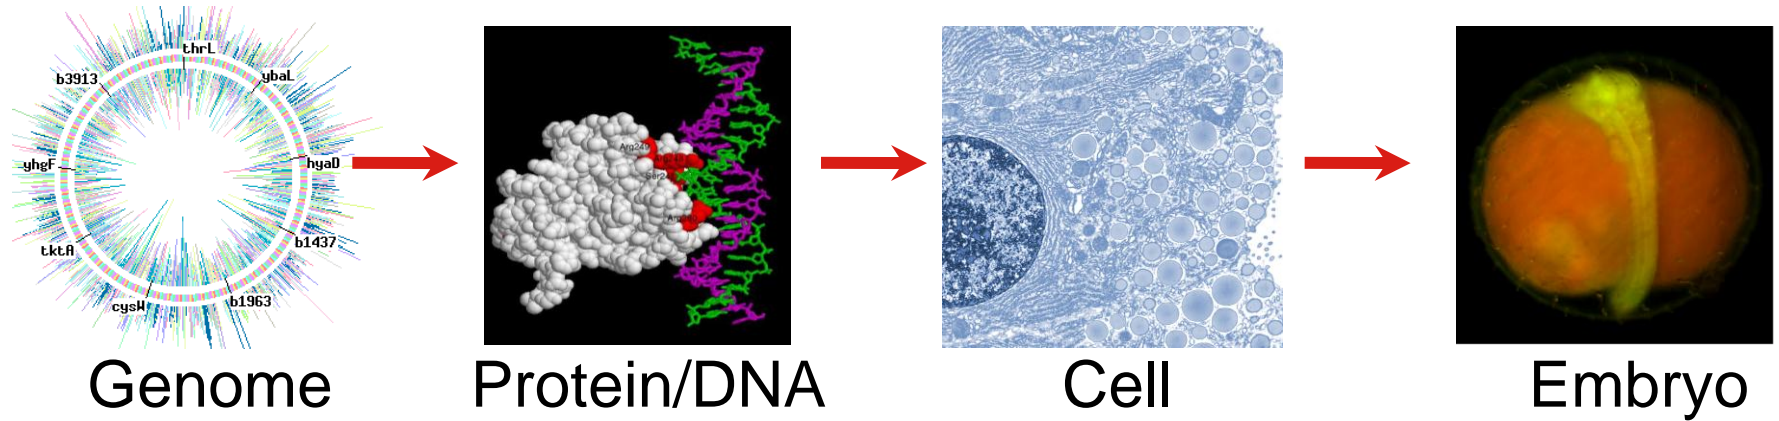


Mousebiology
CNR, EMMA

EMBL's Missions



Systems Biology: From Molecules to Organisms



Development



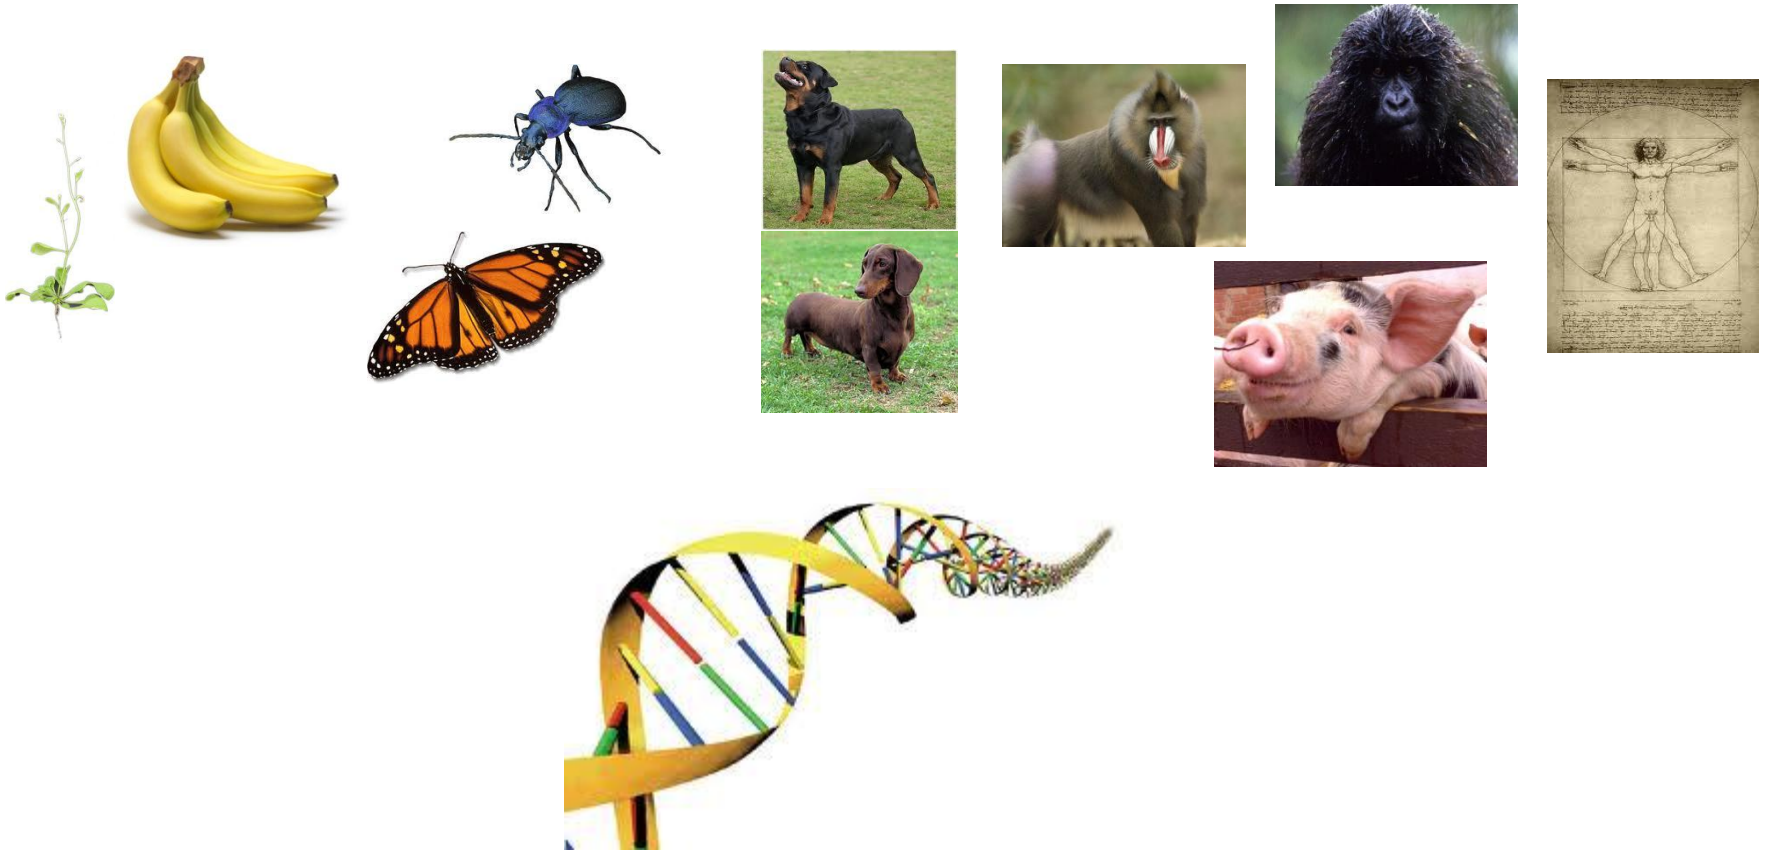
Organisms



Aging

Disease

DNA and Life on Earth

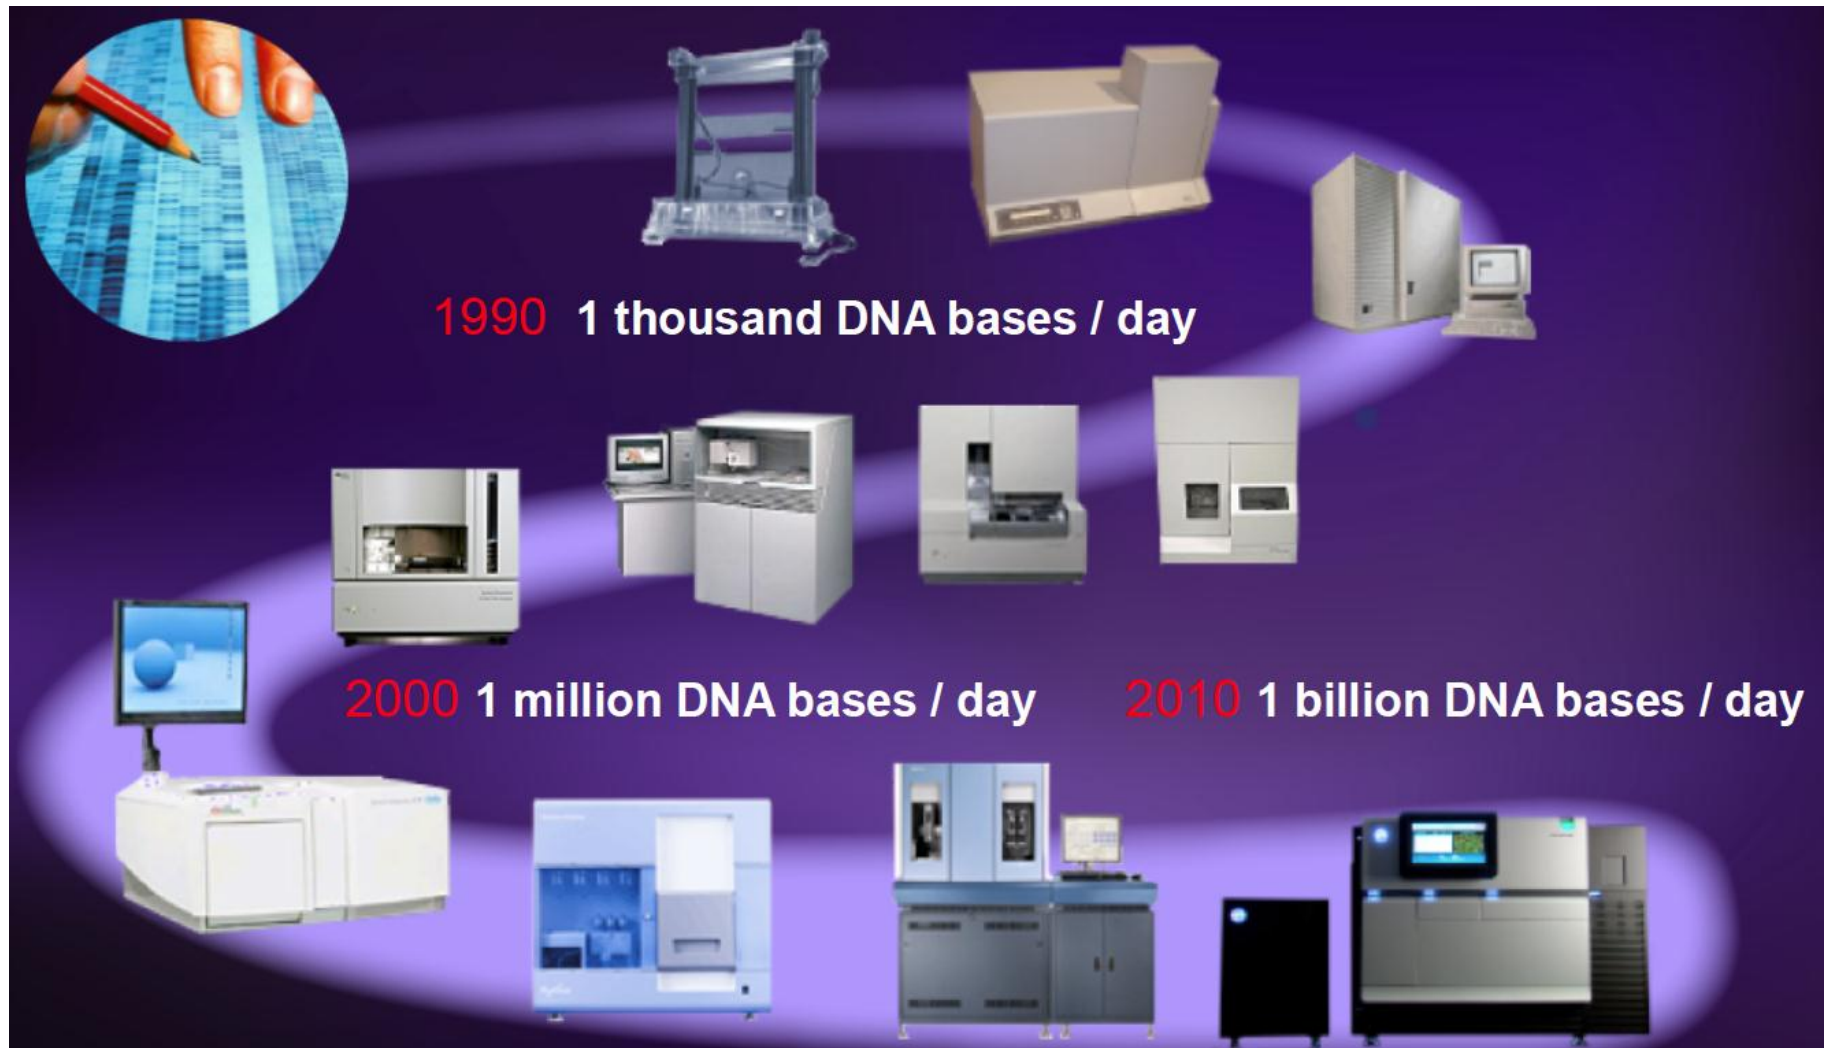


The Sequence Holds the Code for the Organism

Exemplary Big Data Challenge

NEXT GENERATION SEQUENCING (NGS)

Next Generation Sequencing (NGS) Revolution



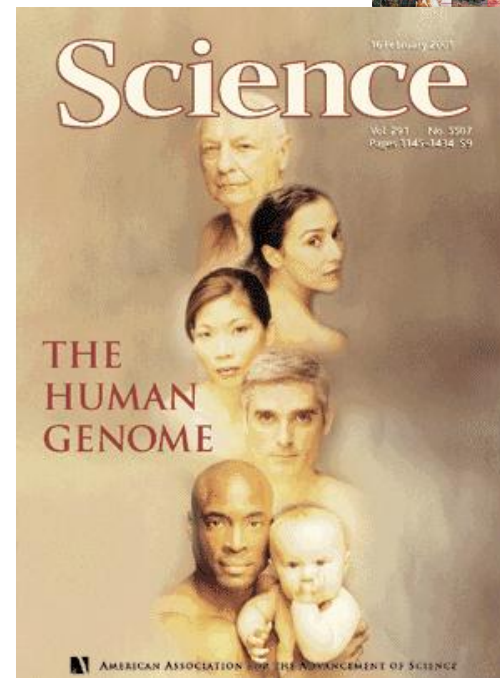
NGS Impact on Human Genome Sequencing

- Human genome project
 - 10 years
 - Large International Consortium
 - Thousands of Sequencers
 - \$3,000,000,000
- Sequencing today
 - < \$10,000
 - A few hours
 - One machine

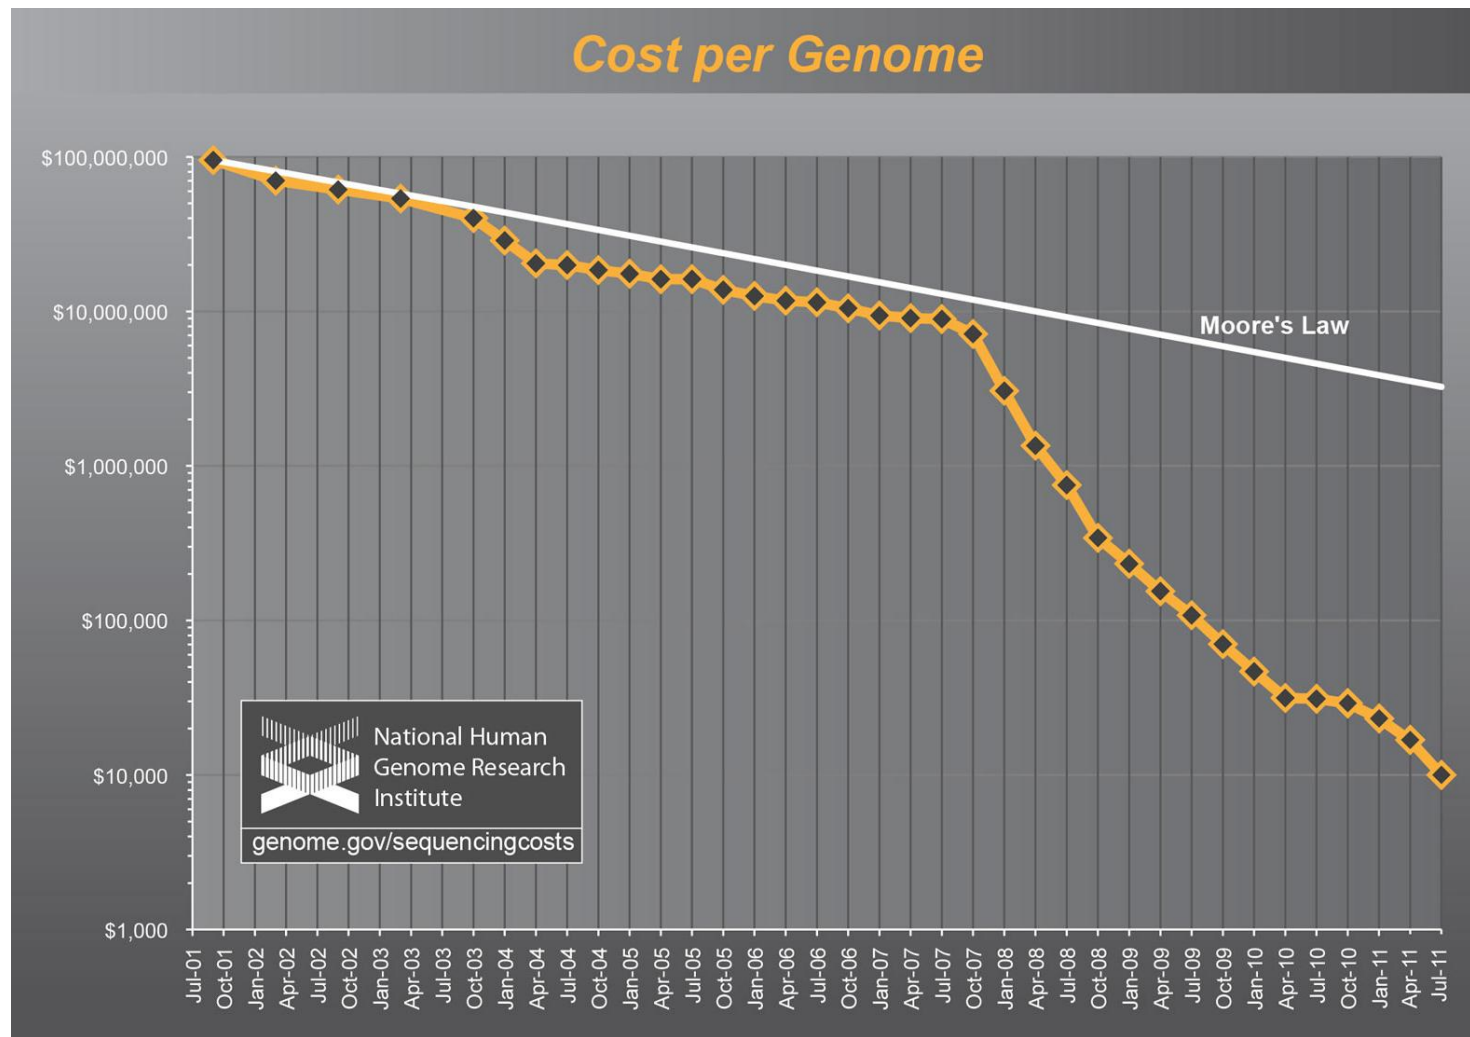
2000



2010



Cost of Sequencing Decreasing Rapidly



Genomic Sequencing is Now an Affordable Solution

Academic
Research
Groups

A screenshot of the 1000 Genomes Project website. The header features the project name in yellow and orange, with a navigation bar below it. A sidebar on the left contains links to 'Home', 'About', and 'Data'. The main content area includes a search bar, a navigation menu with links to 'Database & Species Lists', 'News', 'Events', 'Publications', 'Participants', and 'For G10K Organizers (restricted)', and a large banner for the 'GENOME 10K' project. The banner features a blue DNA double helix and the text 'Unveiling animal diversity'. Below the banner, there is a section titled 'Genome 10K Project' with a description of the project's goals and a 'Join us' button.

1000 Genomes
A Deep Catalog of Human Genetic Variation

Home About Data

Home >

ABOUT THE 1000 GENOMES PROJECT

Project Overview
Project Design
Use of the Project data and samples
Samples included in the project
Publications and project documents

PROJECT OVERVIEW

Recent improvements in sequencing technology have made the 1000 Genomes Project the first project of its kind.

As with other major human genomic projects, the 1000 Genomes Project is a scientific community through free

International Cancer Genome Consortium

Enter keywords Search

Home

Database & Species Lists News Events Publications Participants For G10K Organizers (restricted)

Search: Go

GENOME 10K
Unveiling animal diversity

Genome 10K Project

To understand how complex animal life evolved through changes in DNA and use this knowledge to become better stewards of the planet.

The Genome 10K project aims to assemble a genomic zoo—a collection of DNA sequences representing the genomes of 10,000 vertebrate species, approximately one for every vertebrate genus. The

Join us

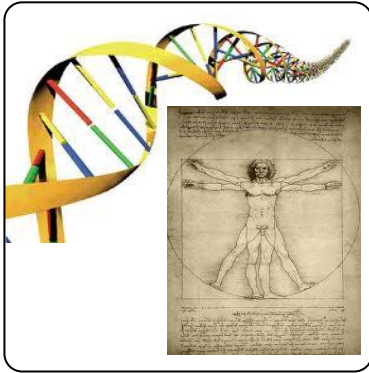
Become a G10K affiliate

Genomic sequencing is
now an affordable solution

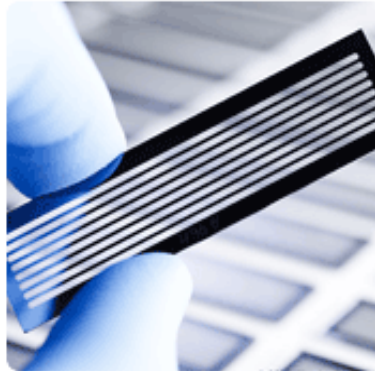
but ...

Read the Sequence to Study the Organism

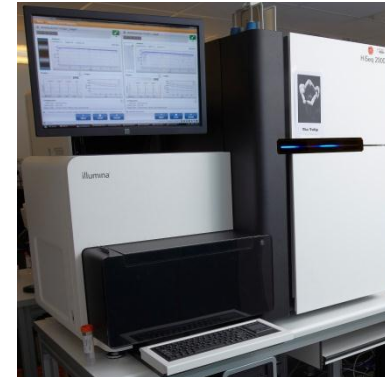
Extract DNA



Prepare



Sequence



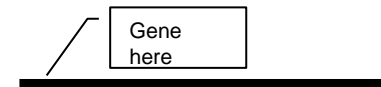
Lab

Assemble



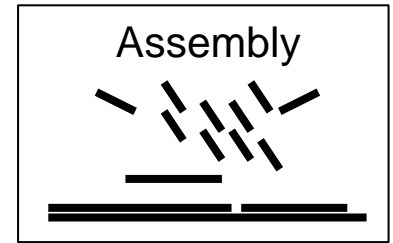
Annotate

In Silico



Requires Computing Infrastructure & Expertise

Problem – 1: Assembly



- NGS output:
millions of very short sequence reads

ATGCATT... 105 ... GTATTCC

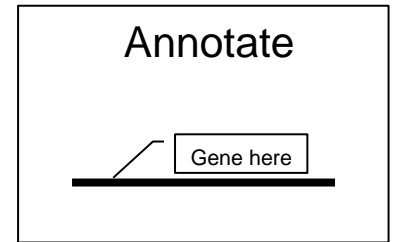
- Genomes contain long strings of bases

ATGCATT... 200,000,000 ...TGCGGATC

- The **short reads** have to be **assembled** into **genomes**
- Up to **1TB RAM** and **many weeks** computation required to solve puzzle

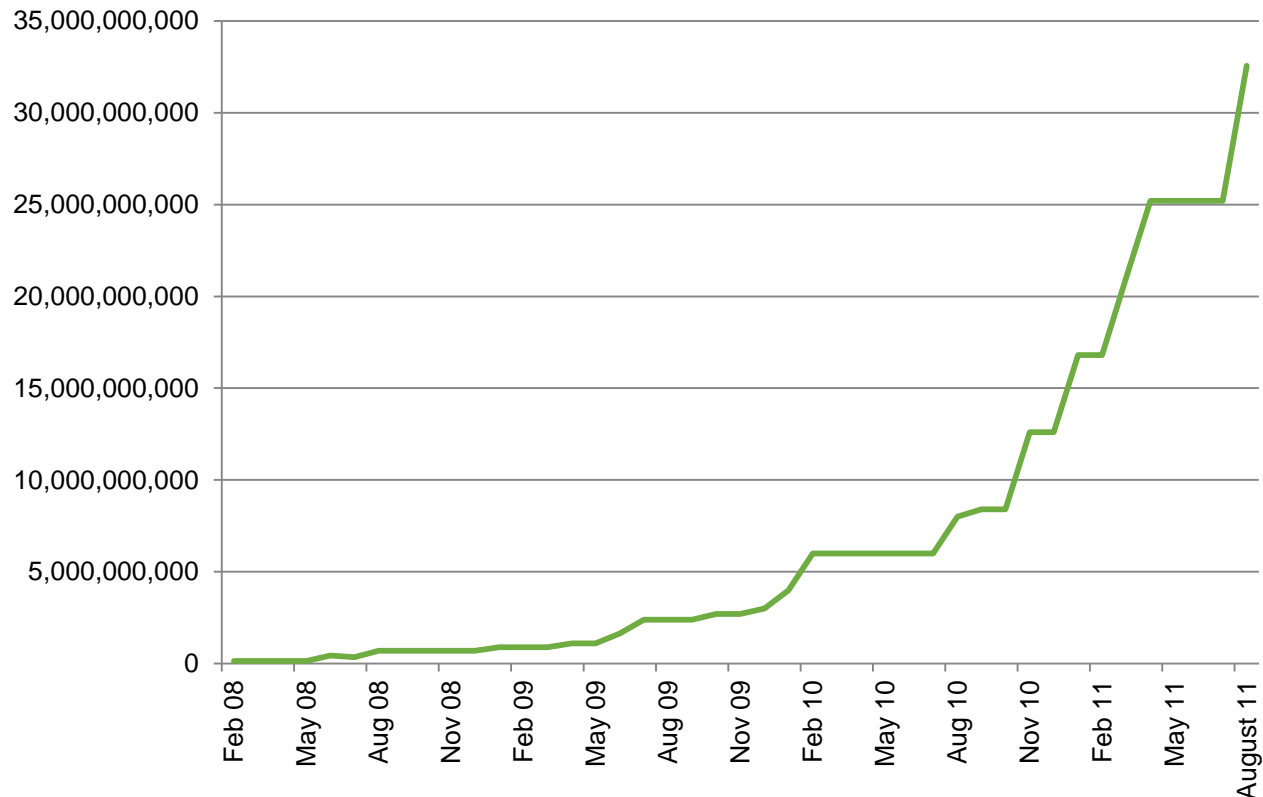
Problem – 2: Annotation

- Strings of assembled bases need to be annotated
- 3 billion bases, ~25k genes
- Looking for genes and regulator elements
- Requires multiple pipelines and databases



Problem - Technology Explosion with NGS

**Bases Sequenced / Sample / Run @ EMBL
(Illumina)**



Sequence Production & IT Infrastructure at EMBL

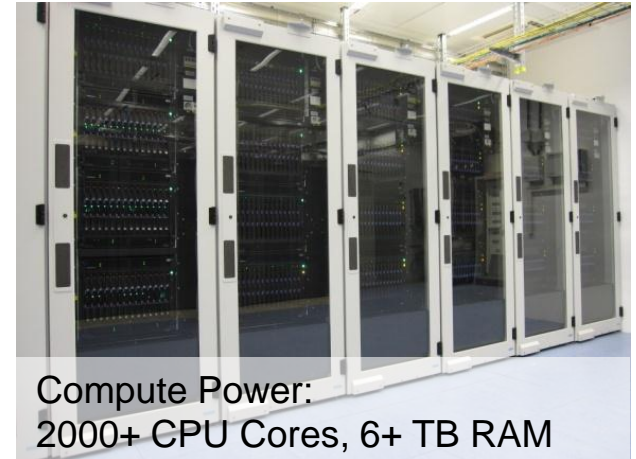
4 x Illumina HiSeq2000



2 x Illumina GAllx



25 TB data
each week



Compute Power:
2000+ CPU Cores, 6+ TB RAM

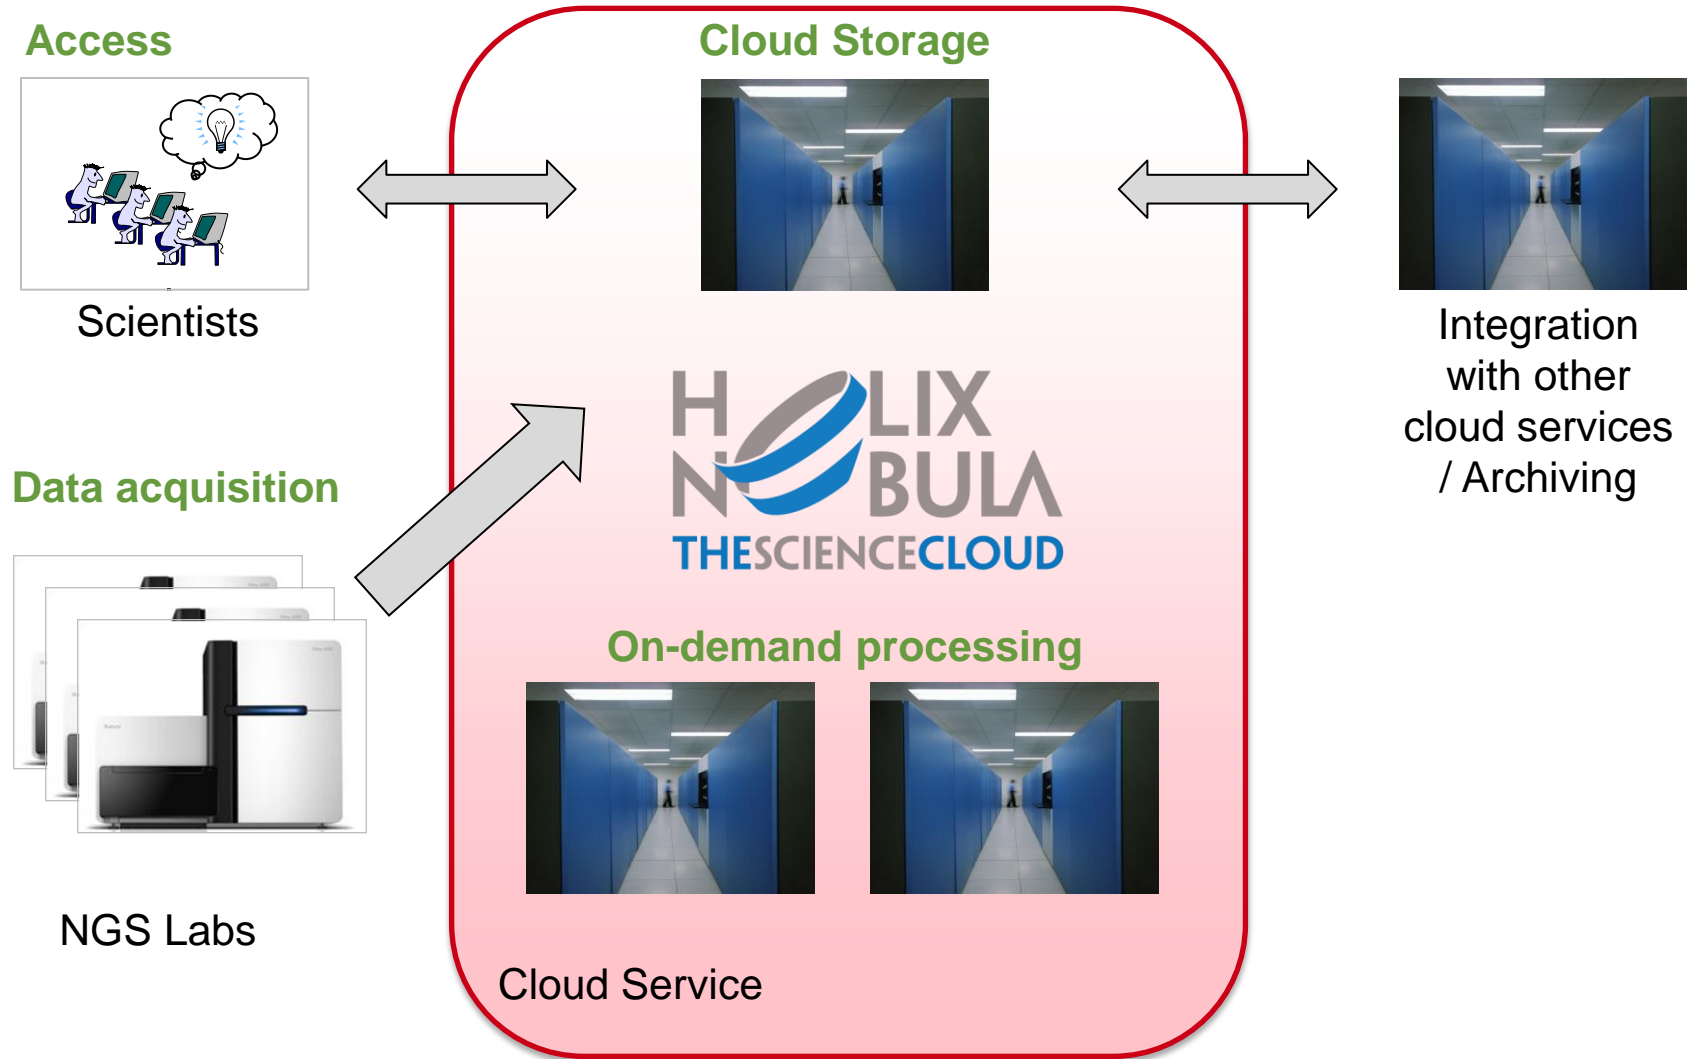


Storage:
1+ PB High Performance Disk

NGS - The Big Picture

- ~ 8.7 million species in the world (estimate)
- ~ 7 billion people
- Sequencers exist in both large centres & small research groups
- 200+ Illumina HiSeq sequencers in Europe alone
 - capacity to sequence 1600 human genomes / month
- Largest centre: Beijing Genomics Institute (BGI)
 - 167 sequencers, 130 HiSeq
 - 2,000 human genomes / day
- 500-1000 Hiseq devices worldwide today
 - 3-6 PB /day
 - 1.1 – 2.2 ExaBbytes / year

EMBL Flagship project: Whole-Genome Assembly



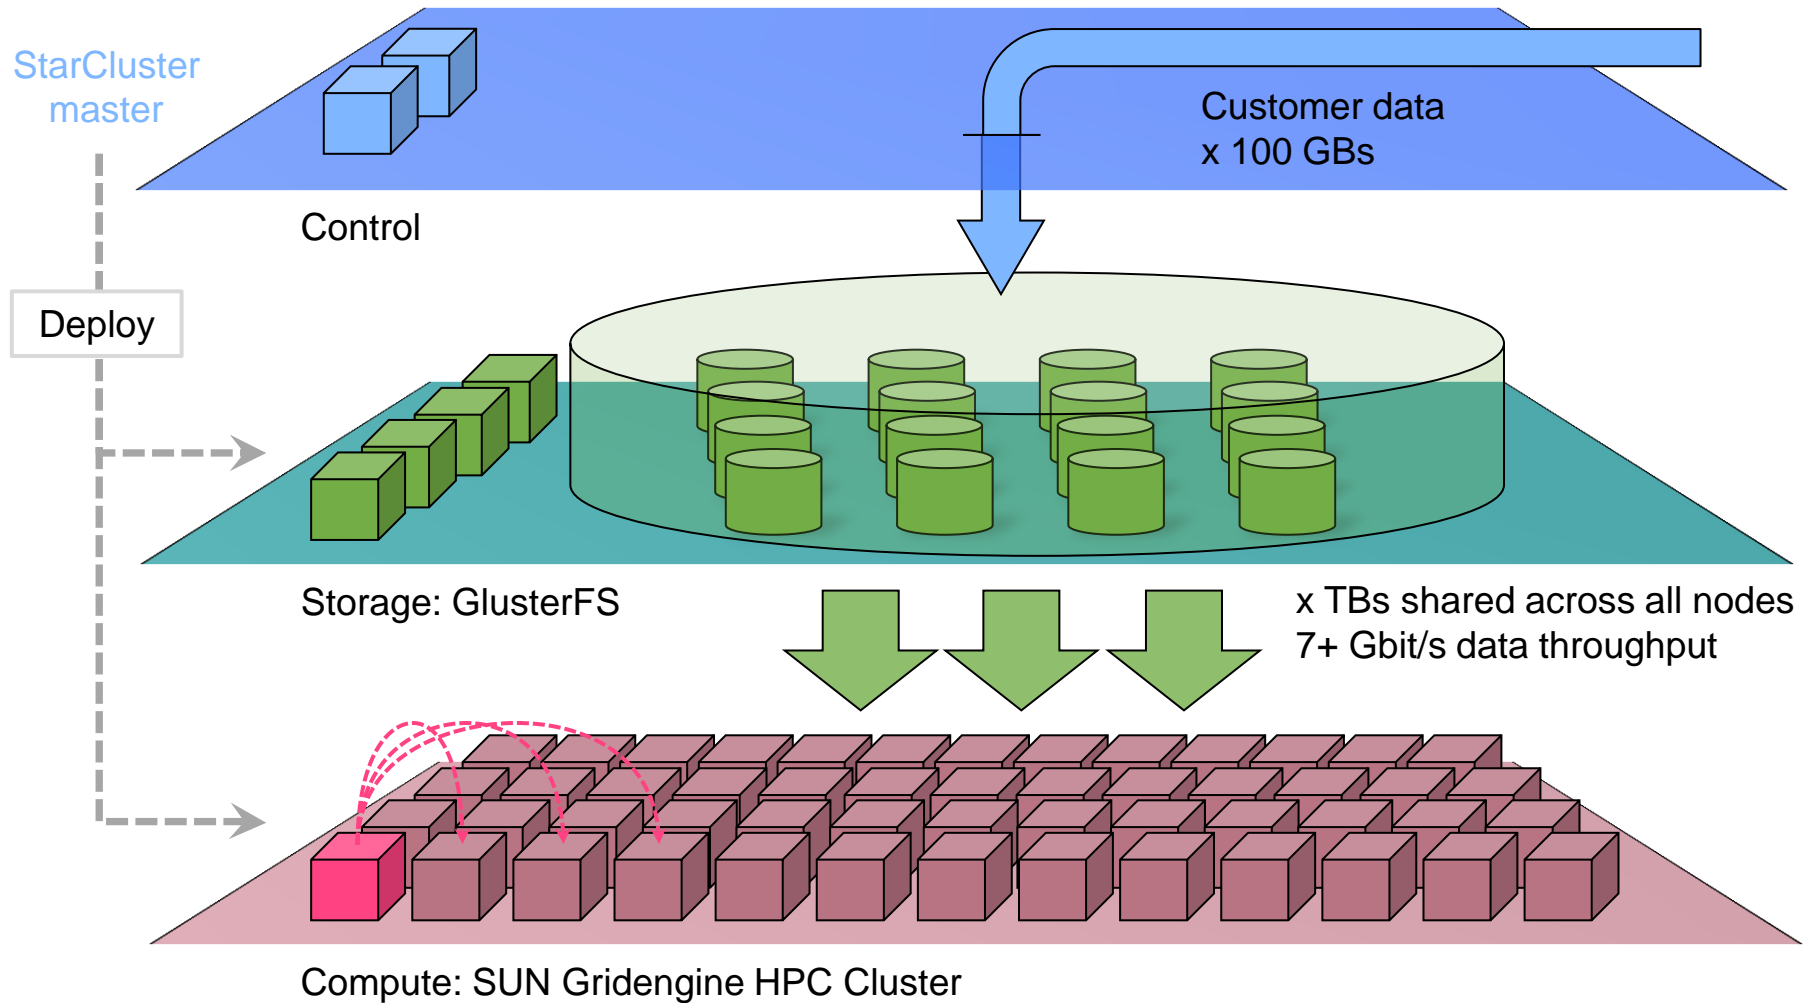
EMBL Flagship Pilot Project

PROOF OF CONCEPT IMPLEMENTATIONS

Proof of Concept Setup

- Multiple Cloud providers
 - ATOS / Sixsq
 - CloudSigma
 - T-Systems
- Each tested 3 major steps with increasing complexity
- Major software components to test
 - Assembly pipeline
 - Annotation pipeline
 - Shared File system
 - StarCluster

EMBL Dynamic Architecture



StarCluster & Sun Grid Engine

Dynamic cluster provisioning

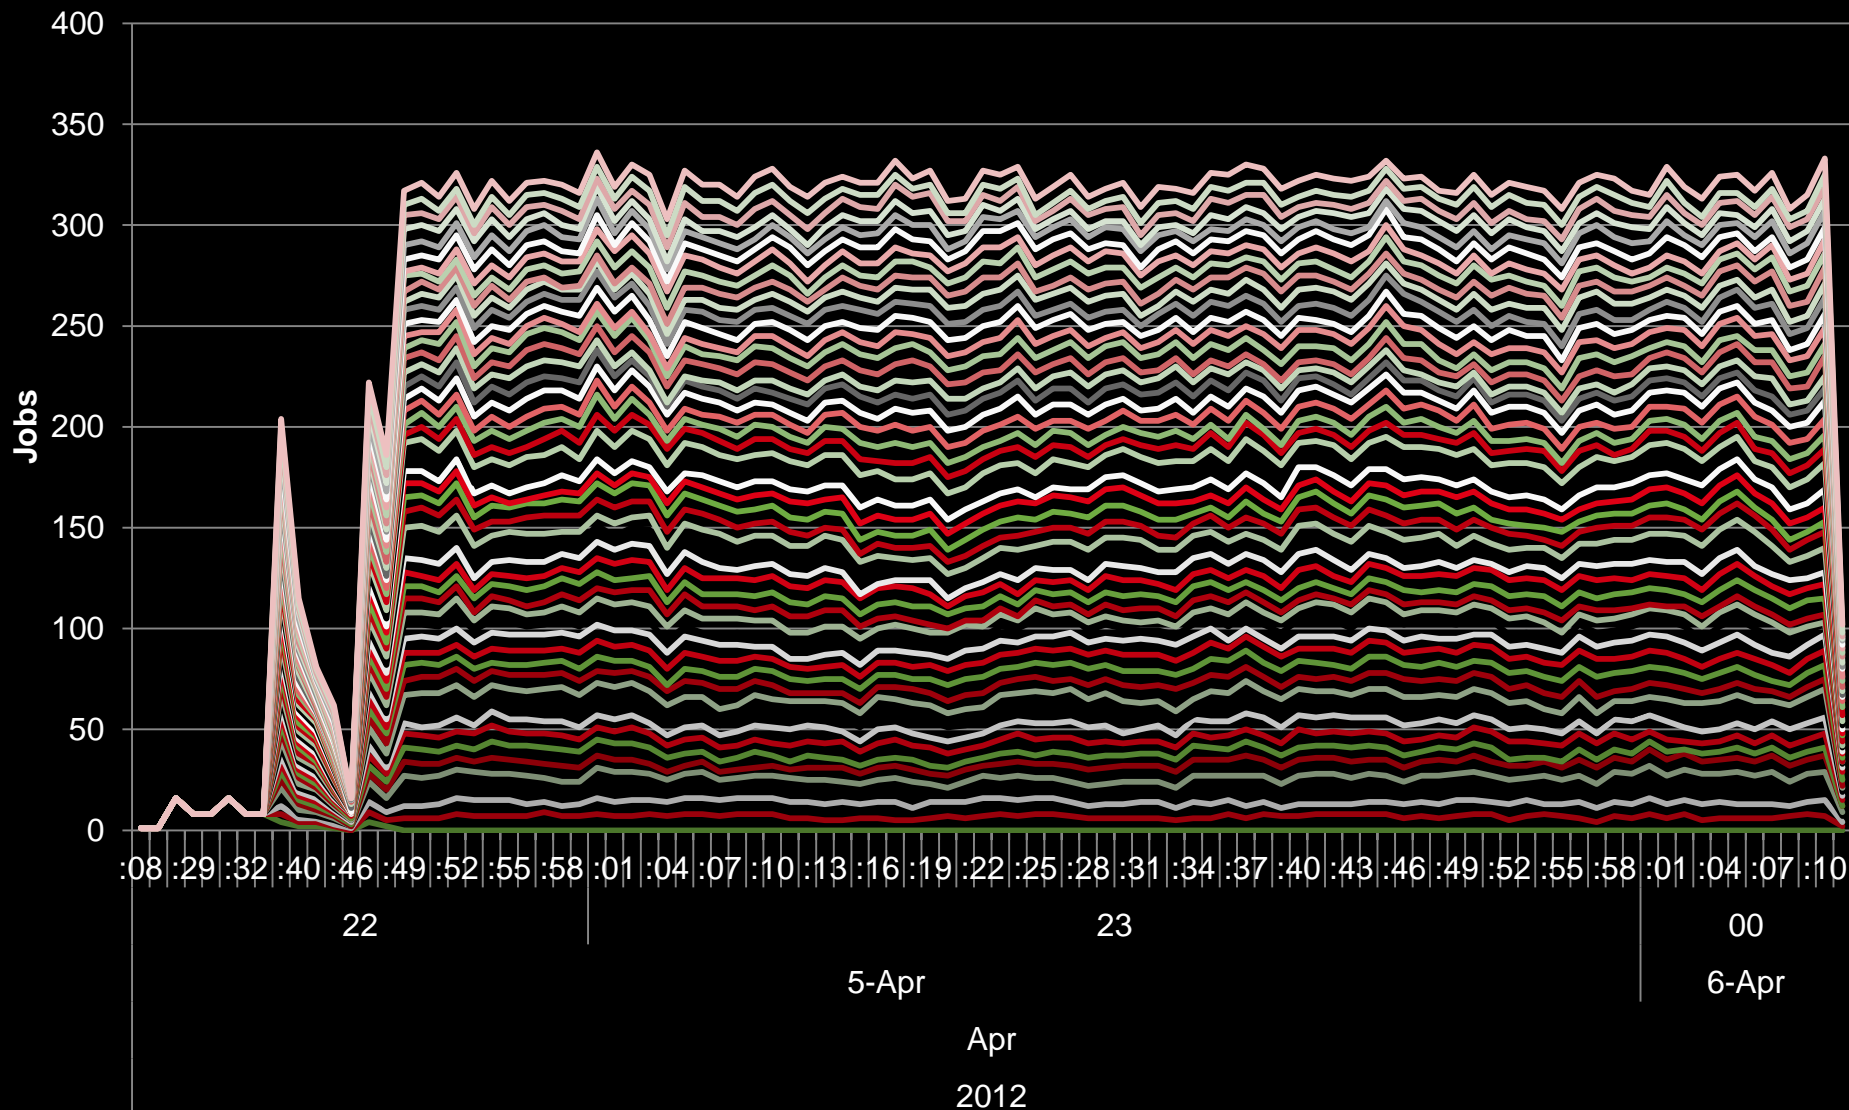
- StarCluster – Dealing with the Fluctuating Workload
 - Manages provisioning of images and setting up of cluster
 - Requires sets of EC2 APIs to work
 - It monitors the number of jobs in the queue and launches more instances
 - Terminates them when no longer required
- Sun Grid Engine
 - Single image running in two modes – master/worker
 - Post-launch configuration

Successful tests of all vendors deployed so far

- StarCluster API integration
- auto-provision 50-node cluster setups
- real world large genome sequencing data
- 100,000s of jobs
- mix of quick parallel jobs and long running serial jobs
- glusterFS stability under high I/O levels
- Initial hurdles (e.g. image deployment, StarCluster integration, network setup) solved

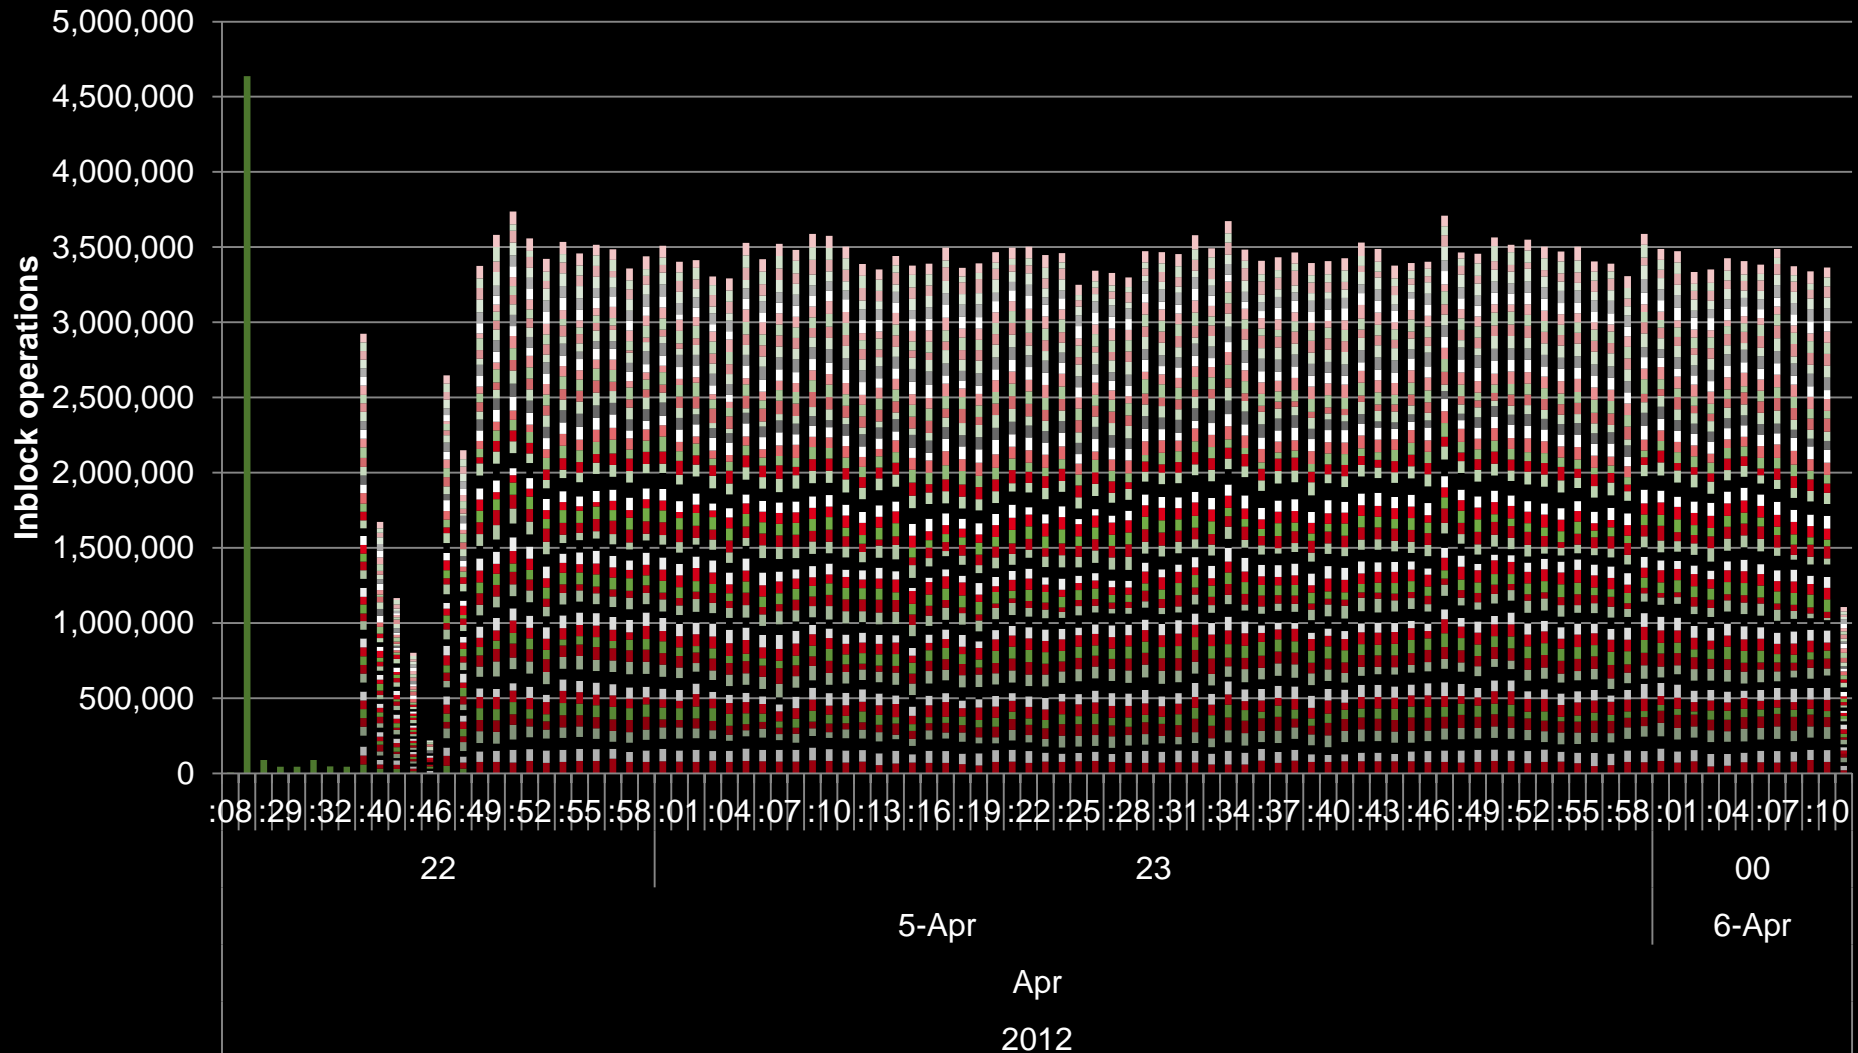
SGE cluster throughput

20.000 annotation jobs / h on 50 nodes



GlusterFS throughput

**60.000 inbound block I/Os / sec
from annotation jobs on 50 nodes**



Next steps

- Identify a suitable model for a future federated Helix Nebula cloud
- Preparations for putting EMBL genome analysis pipeline into production ongoing
- Attract other flagship from within and outside EMBL
 - Through initial success with current genome analysis flagship
 - After implementation of federated cloud model

Helix Nebula PoC Acknowledgements



EMBL

Michael Wahlers
Jonathon Blake
Tobias Rausch
Jürgen Zimmermann
Vladimir Benes
Christian Boulin
Rupert Lueck

EMBL- EBI

Stephen Keenan
Paul Flicek

EMBL



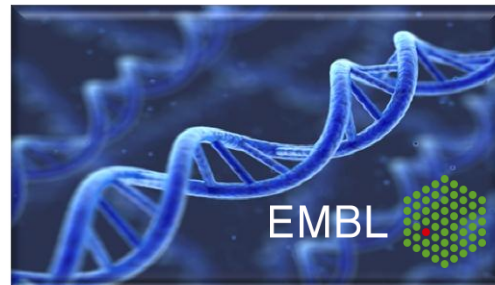
Initial flagships use cases

ATLAS High Energy Physics Cloud Use



To support the computing capacity needs for the ATLAS experiment

Genomic Assembly in the Cloud



A new service to simplify large scale genome analysis; for a deeper insight into evolution and biodiversity

SuperSites Exploitation Platform



To create an Earth Observation platform, focusing on earthquake and volcano research

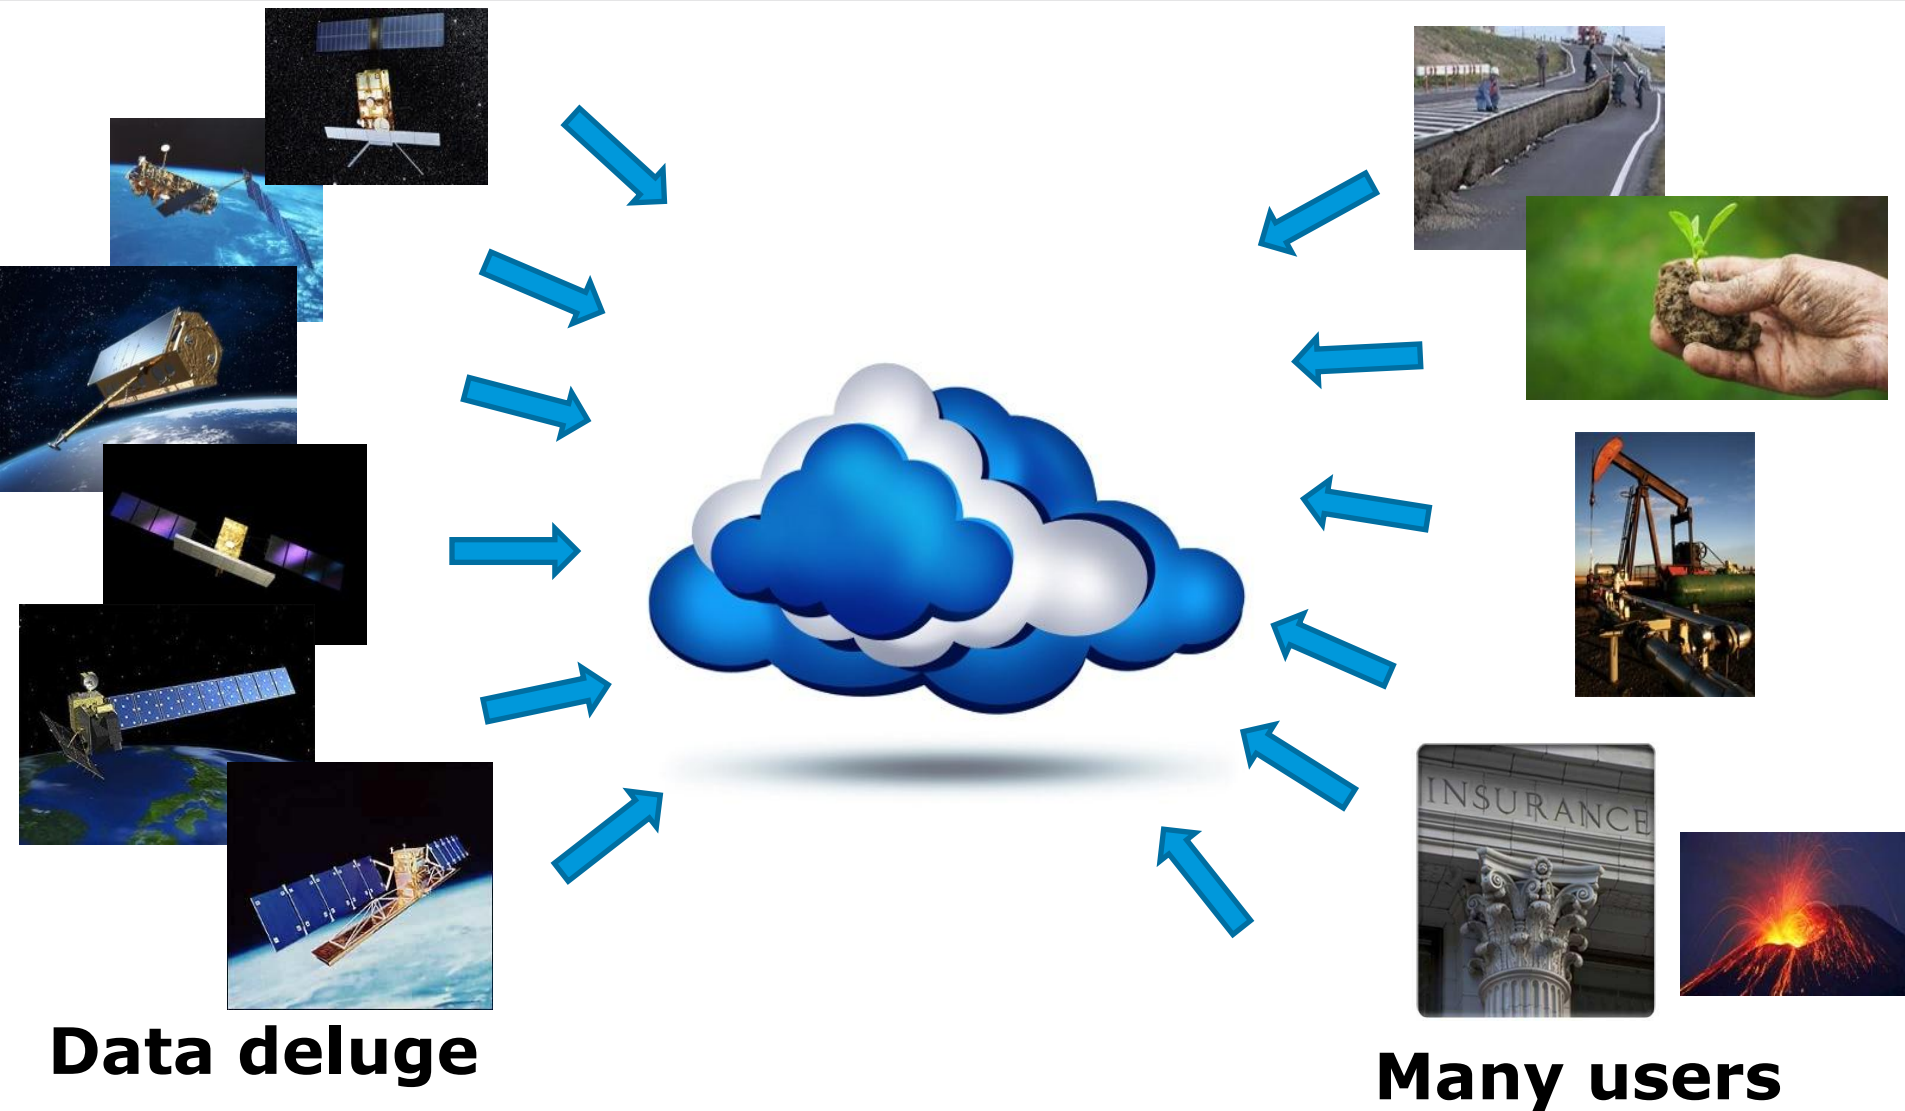
ESA's experience with Helix Nebula and outlook

**Wolfgang Lengert,
ERS and ADM-Aeolus Mission Manager**

presented by Rupert Lueck (EMBL)



Why the Cloud?



1. Can HN cloud computing serve ESA Earth Observation (EO) processing ICT needs?
2. Can ESA deploy an end-to-end platform for Earth Observation exploitation on Helix Nebula?
3. Can an ecosystem of value added service providers develop around such platform?

Approach: SuperSites Exploitation Platform (SSEP)



1. SSEP is vested as an Helix Nebula “flagship”; along sites other flagship at CERN and EMBL.
2. CNES, DLR and CNR agreed to participate to Helix Nebula. The CNR/IREA (Italian Research Council) as a none-space agency contributes with their Radar processor adapted for the cloud.
3. Helix Nebula Prove of Concept participants:
 1. ATOS
 2. CloudSigma
 3. Interoute
 4. T-Systems

1. Can HN cloud computing serve ESA EO processing ICT needs?
2. Can ESA deploy an end-to-end platform for EO exploitation on HN?
3. Can an ecosystem of value added service providers develop around such platform?

Cloud Computing at ESA fingertips integrated in the Grid Processing environment



Browser address bar: gpod.sso.esa.int/admin/ce.aspx

Page title: grid processing on demand

Navigation menu: esa, RSS portal, Home, Services, Workspace, Catalogue, Products, Schedulers, Documentation, My profile, Admin, Help

Search bar: Search, Showing the 13 results found.

User profile: Name: Jordi Farres, Credits: 999, Logout

g-pod logo: grid processing on demand

Computing Element

	Caption	Address
<input type="checkbox"/>	ATOS Cloud CE	212.159.220.77
<input type="checkbox"/>	ATOS Cloud CE (10 WN)	212.159.220.92
<input type="checkbox"/>	CloudSigma Cloud CE	199.180.196.61
<input type="checkbox"/>	Interoute Cloud CE	195.143.228.58
<input type="checkbox"/>	KLIMA CE ESRIN	grid-eo-engine01.esrin.esa.int
<input type="checkbox"/>	KLIMA CE UK	grid-eo-engine06.eo.esa.int
<input type="checkbox"/>	Operational CE 01	grid-eo-engine01.esrin.esa.int
<input type="checkbox"/>	Operational CE 01 64bits	grid-eo-engine01.esrin.esa.int
<input type="checkbox"/>	Operational CE 01 UK	grid-eo-engine06.eo.esa.int

Phase 1: PoC Evaluation Approach



- Performance evaluation (via test scripts)
 - Functions: Data dissemination (upload-cataloguing-download), Data Processing (InSAR, SAR-IPF)
 - Tests: Availability (24x7), Stress, Scalability
- Terms & Conditions evaluation (via questionnaires)
 - Architecture
 - Service Levels
 - Security

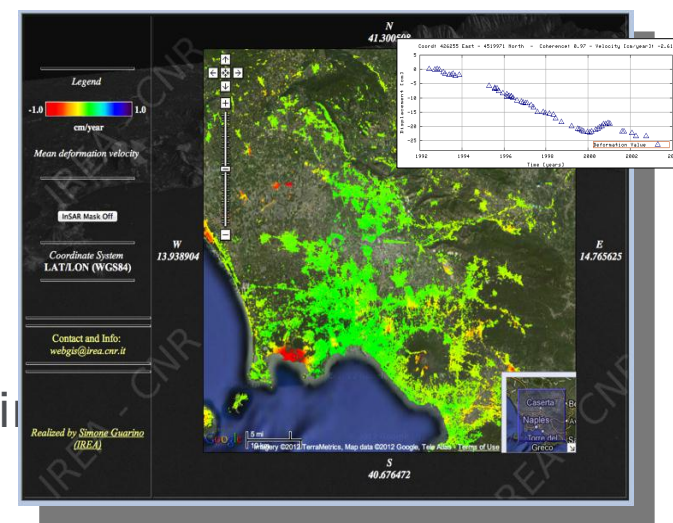
➔ tests have successfully been concluded!

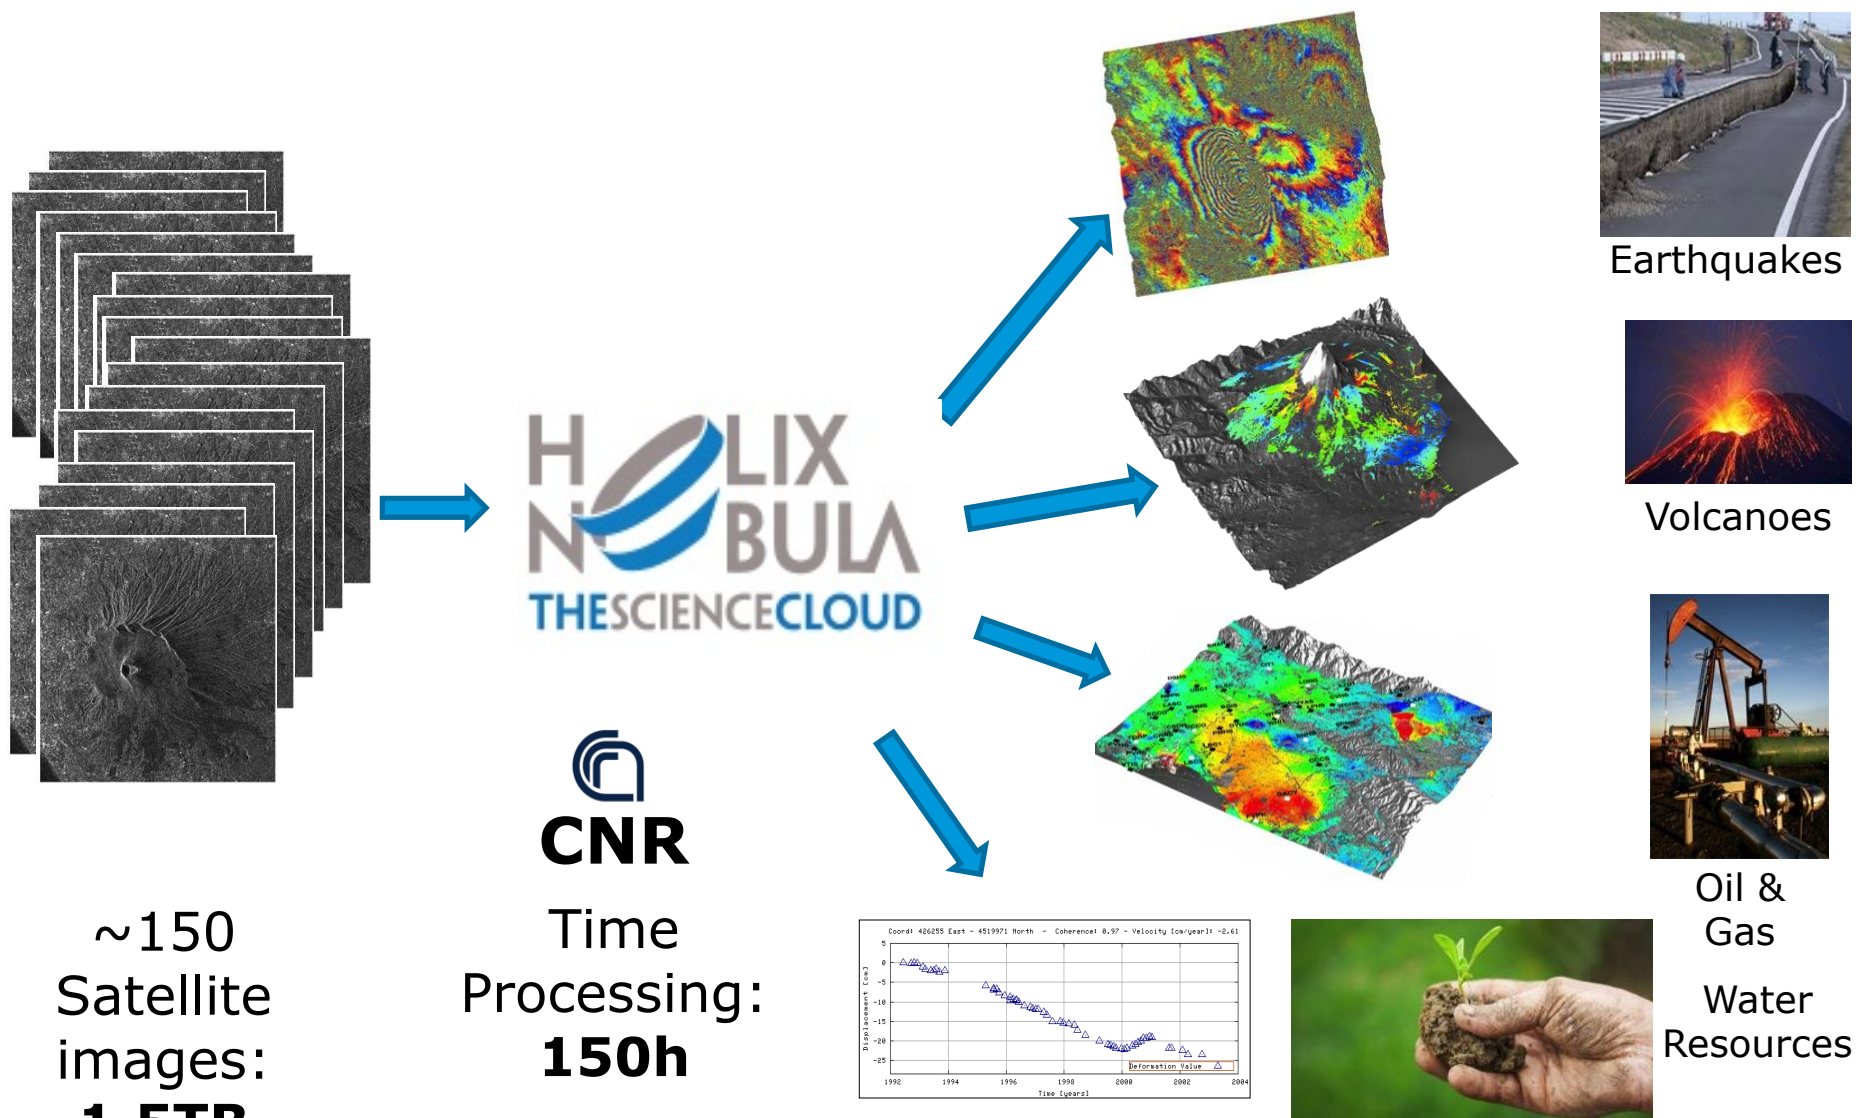
1. Can HN cloud computing serve ESA EO processing ICT needs?
2. Can ESA deploy an end-to-end platform for EO exploitation on HN?
3. Can an ecosystem of value added service providers develop around such platform?

Earth Observation Application Platform exploiting 20 years of satellite data



- EO Application Platform
 - OpenNebula
 - Data Catalogue and Access
 - Map-Reduce computing model
 - Software repository
 - Utilities for sw development and testing
- Cloudification of application
 - CNR / IREA (Italian Research Council in Naples) developed an application (SBAS) measuring the vertical movement of ground in sub cm from space.
 - **SBAS** targets
 - Time series over **20 years** with ESA archive
 - Points of Interest are at **world scale**
 - **TBytes** of data to process

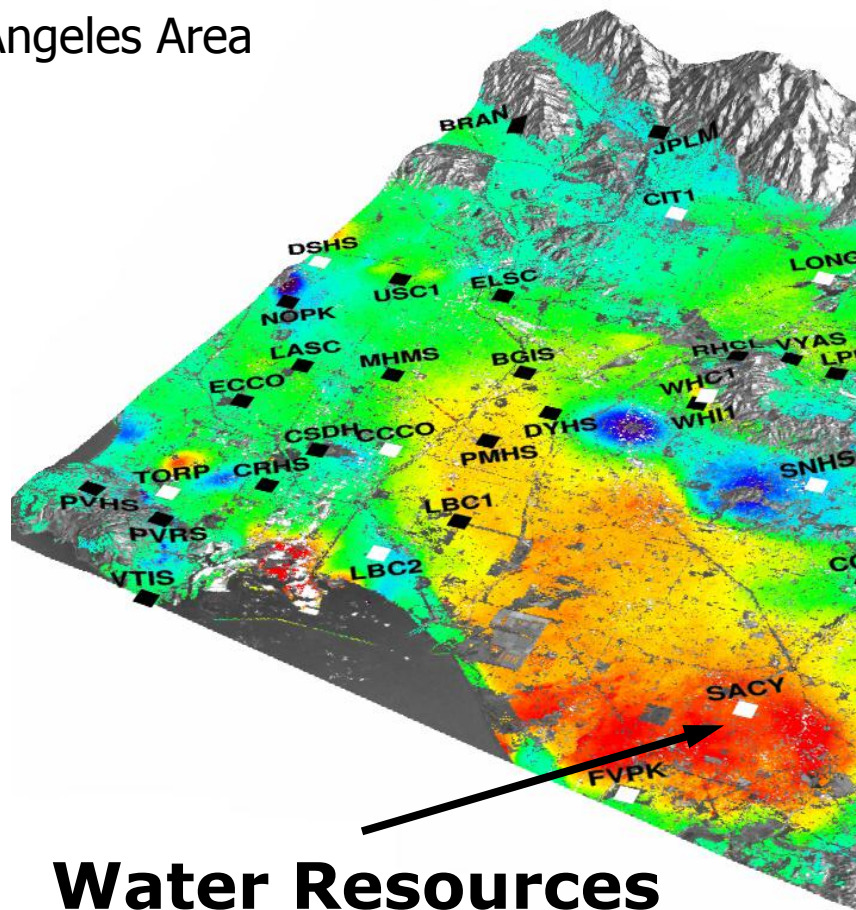




Opportunity: Natural Resources

Los Angeles Area

ERS-1/2 DATA (1995-2002)



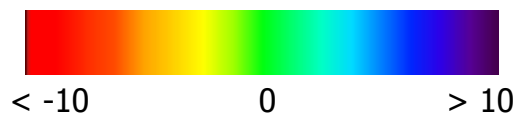
Water Resources



Agriculture

**Sustainable and
social development**

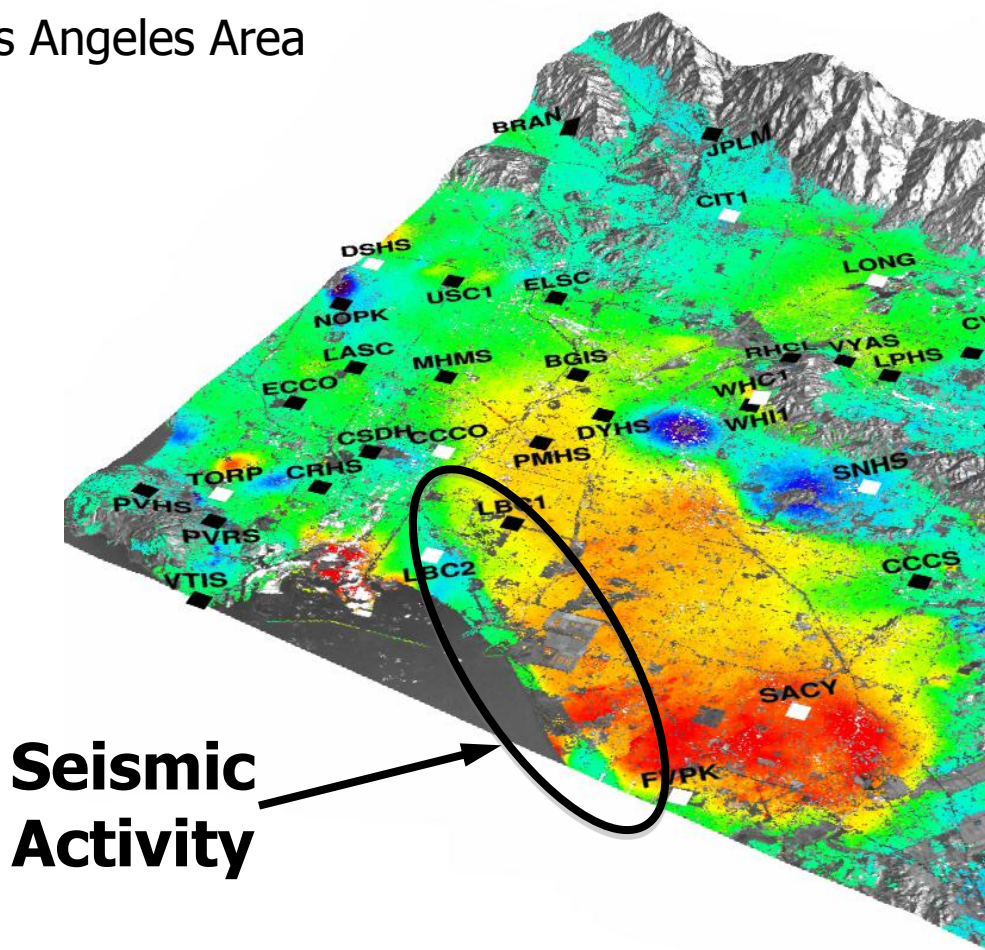
Mean Velocity (mm/a)



Opportunity: Natural Hazards

Los Angeles Area

ERS-1/2 DATA (1995-2002)



Mean Velocity (mm/a)



< -10 0 > 10

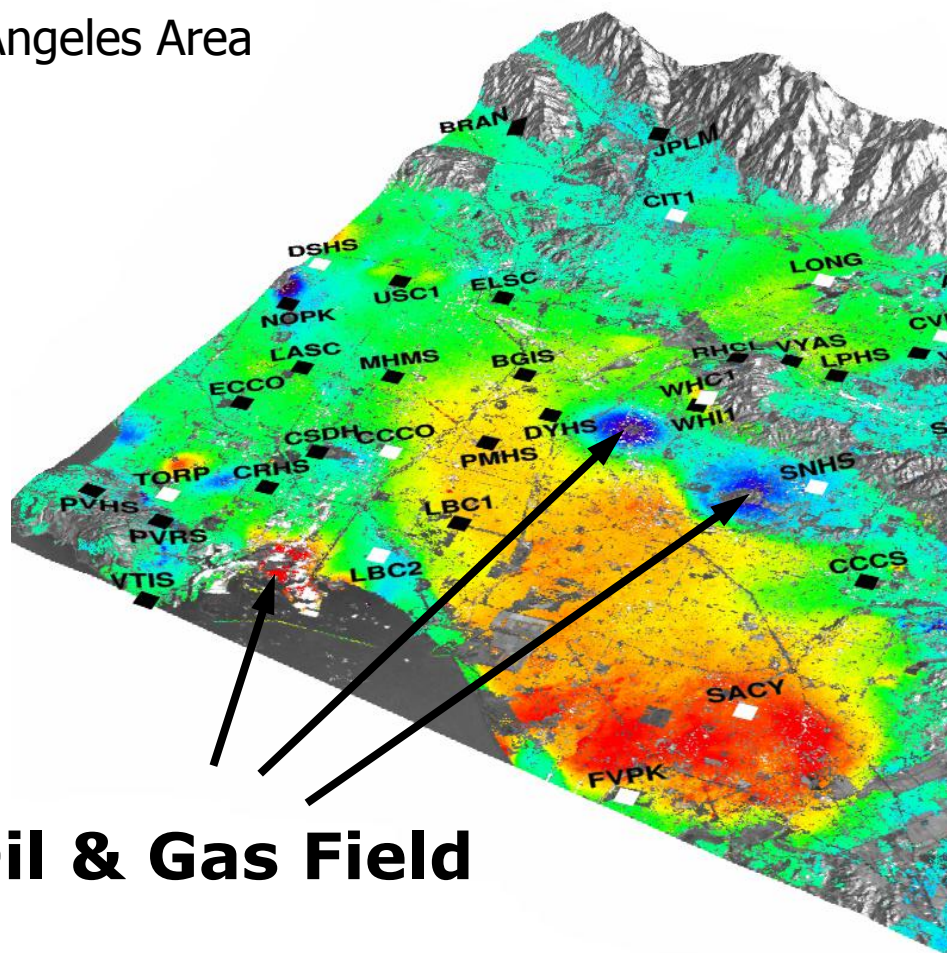


**Civil Protection
Risk Management
Insurances**

Opportunity: Energy Resources

Los Angeles Area

ERS-1/2 DATA (1995-2002)



Oil & Gas Field

Mean Velocity (mm/a)



< -10 0 > 10



1. Can HN cloud computing serve ESA EO processing ICT needs?
2. Can ESA deploy an end-to-end platform for EO exploitation on HN?
3. Can an ecosystem of value added service providers develop around such platform?

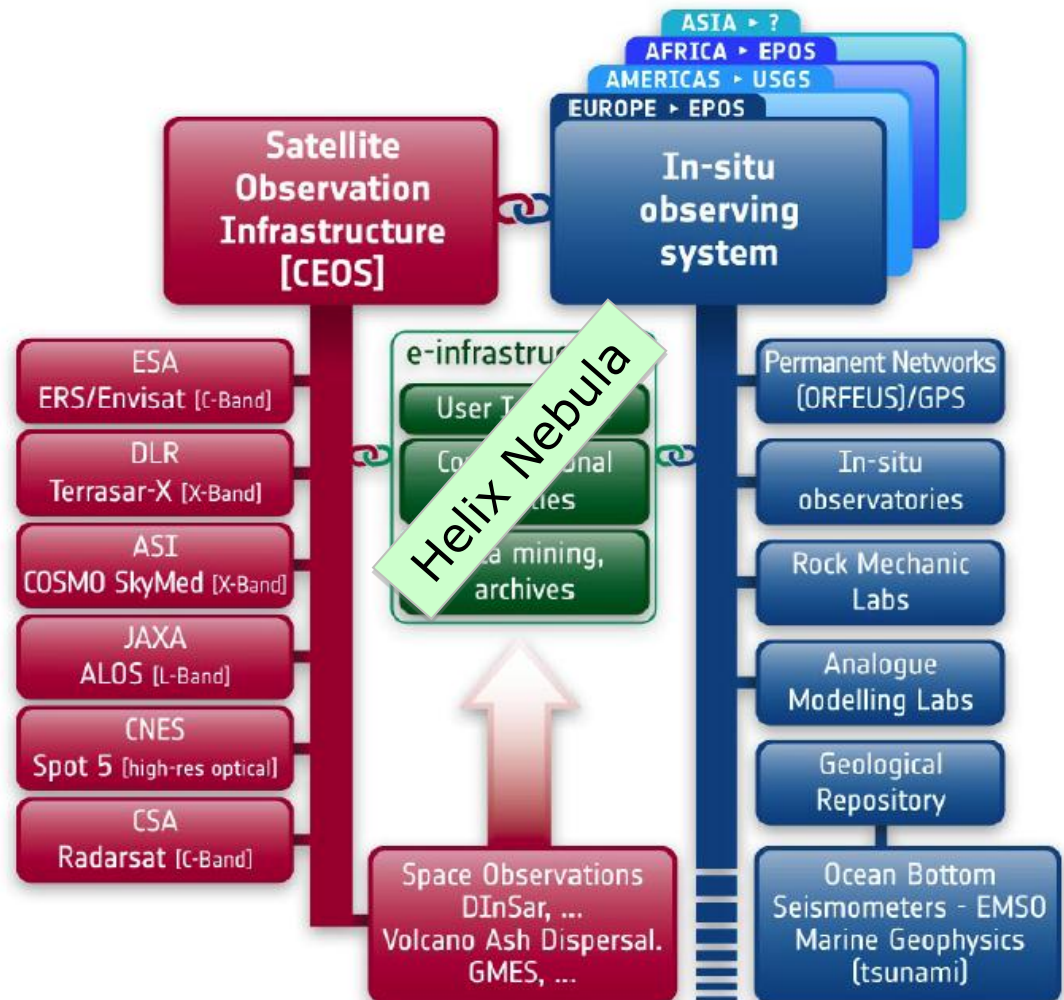
Super Site Exploitation Platform (SSEP)



different actors, different environment helping to understand the Geophysics of Earthquakes and Volcanoes



The Geohazard Supersites partnership pool and coordinate the existing space-based and ground-based observation resources of GEO members to mitigate and to improve the preparedness for geologic disasters



Supersite Exploitation Platform: potential actors benefits

EO data provider benefits:

- Enlarge EO data exploitation (space agencies)
- Increase EO data sales (commercial distributors), in particular EO data archives

End-user benefits:

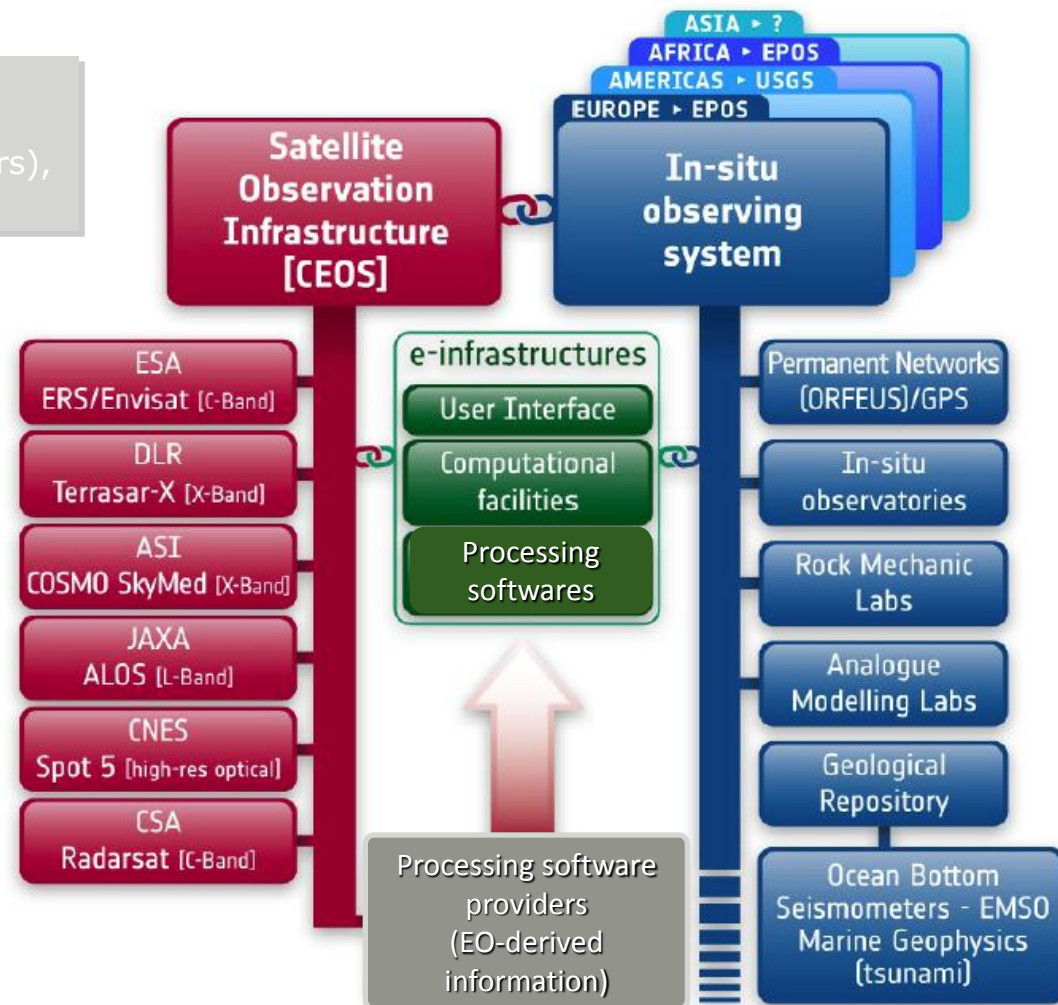
- More data, either free or at low cost
 - Processing capabilities free or at low cost
 - Processing softwares free or at low cost
 - Forum for discussing/exchanging results
- More science

Processing software provider benefits:

- Low investment
- Increase sales
- Increase software visibility

IT companies (computational facilities) benefits:

- New business
- Access to a global user community
- Contribution to science



Example of R&D feedbacks between ESA and EO services industry

- Jun 2003 : **Renewable Energy Industry** (33 companies).
- Oct 2007 : **EO services Industry** (100 companies)
- Sep 2009 : **Insurance** (15 companies).
- May 2008 + 2010 : **World Bank Group**.
- Oct 2009 : **SwissRe** (Flood Risk)
- Jul 2010 : **1st Global Business Biodiversity symposium**.
- Sep 2010 : **Oil & Gas** (104 participants).



- Successful PoC with IaaS Providers
 - Able to perform tests in 3-4 providers
 - Weaknesses could be addressed during following project phases
- Federated HN vs single cloud providers
 - Large differences among providers
 - Multi-sourcing approach recommended for next phase
- Using cloud as a grid vs using native PaaS
 - Evaluation run à la GRID with static provisioning
 - but future use of cloudified applications and dynamic provisioning
- Application cloudification challenges
 - SBAS cloudification required significant effort and deep application expertise

A European cloud computing partnership: big science teams up with big business



Strategic Plan

- ▶ Establish multi-tenant, multi-provider cloud infrastructure
- ▶ Identify and adopt policies for trust, security and privacy
- ▶ Create governance structure
- ▶ Define funding schemes



To support the computing capacity needs for the ATLAS experiment

EMBL



Setting up a new service to simplify analysis of large genomes, for a deeper insight into evolution and biodiversity



To create an Earth Observation platform, focusing on earthquake and volcano research

Atos

Capgemini
CONSULTING. TECHNOLOGY. OUTSOURCING

CloudSigma

eGI

interoute
from the ground to the cloud

logica
be brilliant together

OpenNebula.org
The Open Source Toolkit for Cloud Computing

orange Business Services

SAP

the SERVER LABS
the IT architects

sixsq

Telefonica

terradue 2.0

THALES

...T...Systems



contact@helix-nebula.eu



@HelixNebulaSC



HelixNebula.TheScienceCloud

World Map of High-throughput Sequencers



PoC Steps

- Step 0 – Infrastructure setup and code test
 - Transfer of images and content
 - Set up shared file system: GlusterFS with 4 nodes (1.2TB net)
 - Assembly: SGA assembler tested using small data set
 - Annotation: Manual small batch run against small data set

- Step 2 – Big genome & elastic scalability
 - StarCluster essential in this step
 - automated provisioning of Sun Grid Engine cluster up to 50 nodes
 - Assembly: Large genome sequencing data, mix of quick parallel jobs and long running serial jobs
 - Annotation: Pipeline tested using big data set, 50k - 100k jobs run
 - Validation against run on EMBL infrastructures

- Step 3 (Optional) Large Genome on big box
 - Process a large genome through velvet assembly software
 - Using a high RAM 1TB server



Time
Processing:
150h