Italian Big Data Open Platform for Research and Innovation

Roberta Turra

Venezia, 16 November 2017

Agenda

- Il Cineca
- Il contesto nazionale e Europeo
- Infrastruttura e modalità di accesso
- La piattaforma big data – i primi passi
- La piattaforma e i servizi
- Le opportunità di sviluppo (progetti e convenzioni)

  EUDAT (data life cycle), Telethon (repository), ISTAT (web scraper e text mining), ImediaCity (deep learning)
- I prossimi passi
- La label di i-Space

Il Cineca

CINECA è un consorzio interuniversitario senza scopo di lucro al servizio della ricerca nazionale.

Nasce nel ’69 con lo scopo di promuovere l’utilizzo dei più avanzati sistemi di elaborazione dell’informazione a favore della ricerca tecnologica e scientifica.

Punto di riferimento per il sistema accademico nazionale … ma anche per le imprese

CINECA con la sua infrastruttura si posiziona:
- MARCONI 14 posto Top500 per sistemi di calcolo dedicati alla ricerca e big data

Top500 List - November 2017

	Performance Computing Japan	M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path Fujitsu				
10	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280	
11	DOE/SC/Argonne National Laboratory United States	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,945
12	Texas Advanced Computing Center/Univ. of Texas United States	**Stampede2** - PowerEdge C6320P/C6420, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path Dell EMC	368,928	8,317.7	18,215.8	
13	GSIC Center, Tokyo Institute of Technology Japan	**TSUBAME3.0** - SGI ICE XA, IP139-SXM2, Xeon E5-2680v4 14C 2.4GHz, Intel Omni-Path, NVIDIA Tesla P100 SXM2 HPE	135,828	8,125.0	12,127.1	792
14	CINECA Italy	**Marconi Intel Xeon Phi** - CINECA Cluster, Lenovo SD530, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path Lenovo	314,384	7,471.1	15,372.0	
15	United Kingdom Meteorological Office United Kingdom	Cray XC40, Xeon E5-2695v4 18C 2.1GHz, Aries interconnect Cray Inc.	241,920	7,038.9	8,128.5	
16	Barcelona Supercomputing Center Spain	**MareNostrum** - Lenovo SD530, Xeon Platinum 8160 24C 2.1GHz, Intel Omni-Path Lenovo	153,216	6,470.8	10,296.1	1,632

Italian HPC systems

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
14	**Marconi Intel Xeon Phi** - CINECA Cluster, Lenovo SD530, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path , Lenovo CINECA Italy	314,384	7,471.1	15,372.0	
37	**HPC2** - iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.8GHz, Infiniband FDR, NVIDIA K20x , IBM Exploration & Production - Eni S.p.A. Italy	72,000	3,188.0	4,605.0	1,227
51	**HPC3** - Lenovo NeXtScale nx360M5, Xeon E5-2697v4 18C 2.3GHz, Infiniband EDR, NVIDIA Tesla K80 , Lenovo Energy Company (A) Italy	66,000	2,592.0	3,798.6	
72	**Marconi Intel Xeon** - Lenovo NeXtScale nx360M5, Xeon E5-2697v4 18C 2.3GHz, Omni-Path , Lenovo CINECA Italy	54,432	1,723.9	2,003.1	1,360.8
396	**GALILEO** - IBM NeXtScale nx360M4, Xeon E5-2630v3 8C 2.4GHz, Infiniband QDR, Intel Xeon Phi 7120P , IBM/Lenovo CINECA Italy	50,232	684.3	1,103.1	2,825.6
440	**DAVIDE** - E4 OP206 cluster, IBM Power8+ 8C 3.26GHz, Infiniband EDR, NVIDIA Tesla P100 SXM2 , E4 Computer Engineering S.p.A. E4 Computer Engineering Italy	10,560	615.4	951.2	78.3

Mission

- Provide ***High Performance Computing and Data*** capabilities and capacity. Largest HPC infrastructure in Italy and one of the largest in Europe
- Support academic users and large scientific communities in maximizing systems load and utilization
- Design and develop applications for massive data processing
- Foster technology transfer and innovation towards the private sector

Presence in Europe

PRACE - Partnership for Advanced Computing in Europe

- The PRACE AISBL is a fully self-financed non–for-profit organization under Belgian law.
- 25 member countries (October 2012): Austria, Belgium, Bulgaria, Cyprus, Czech Republic, Denmark, Finland, **France**, **Germany**, Greece, Hungary, Ireland, Israel, **Italy**, The Netherlands, Norway, Poland, Portugal, Slovakia, Slovenia, **Spain**, Sweden, Switzerland, Turkey and the UK.

ETP4HPC – European Technology Platform for HPC

- Incorporated as a Dutch association
- Steering board with 15 members:
  - Research centers (5): BSC, CEA , **Cineca (Delegate Carlo Cavazzoni)**, Fraunhofer, FZJ Juelich
  - European SMEs (4): Allinea, ClusterVision, Megware, ParTec
  - European controlled corporations (4): ARM, Atos/Bull, Eurotech, Seagate
  - International companies with R&D in Europe (2): IBM, Intel

BDVA – European Technology Platform Big Data Value Association

- The BDVA AISBL is a fully self-financed non–for-profit organization under Belgian law.

European funded R&D projects

e-Infrastructure:

PRACE

EUDAT
European Data Infrastructure

plus ENVRI

elixir

EPOS
EUROPEANPLATEOBSERVINGSYSTEM

Energy efficiency

ANTAR EX

Exascale technology

MONT BLANC

DEEP-ER

European Flagship

MAX

Center of Excellence

Human Brain Project

Fabric of the future

ff
FORTISSIMO

Multimedia/
Cultural Heritage

i media cities

ARROW

FORWARD

Users, suppliers and access channels

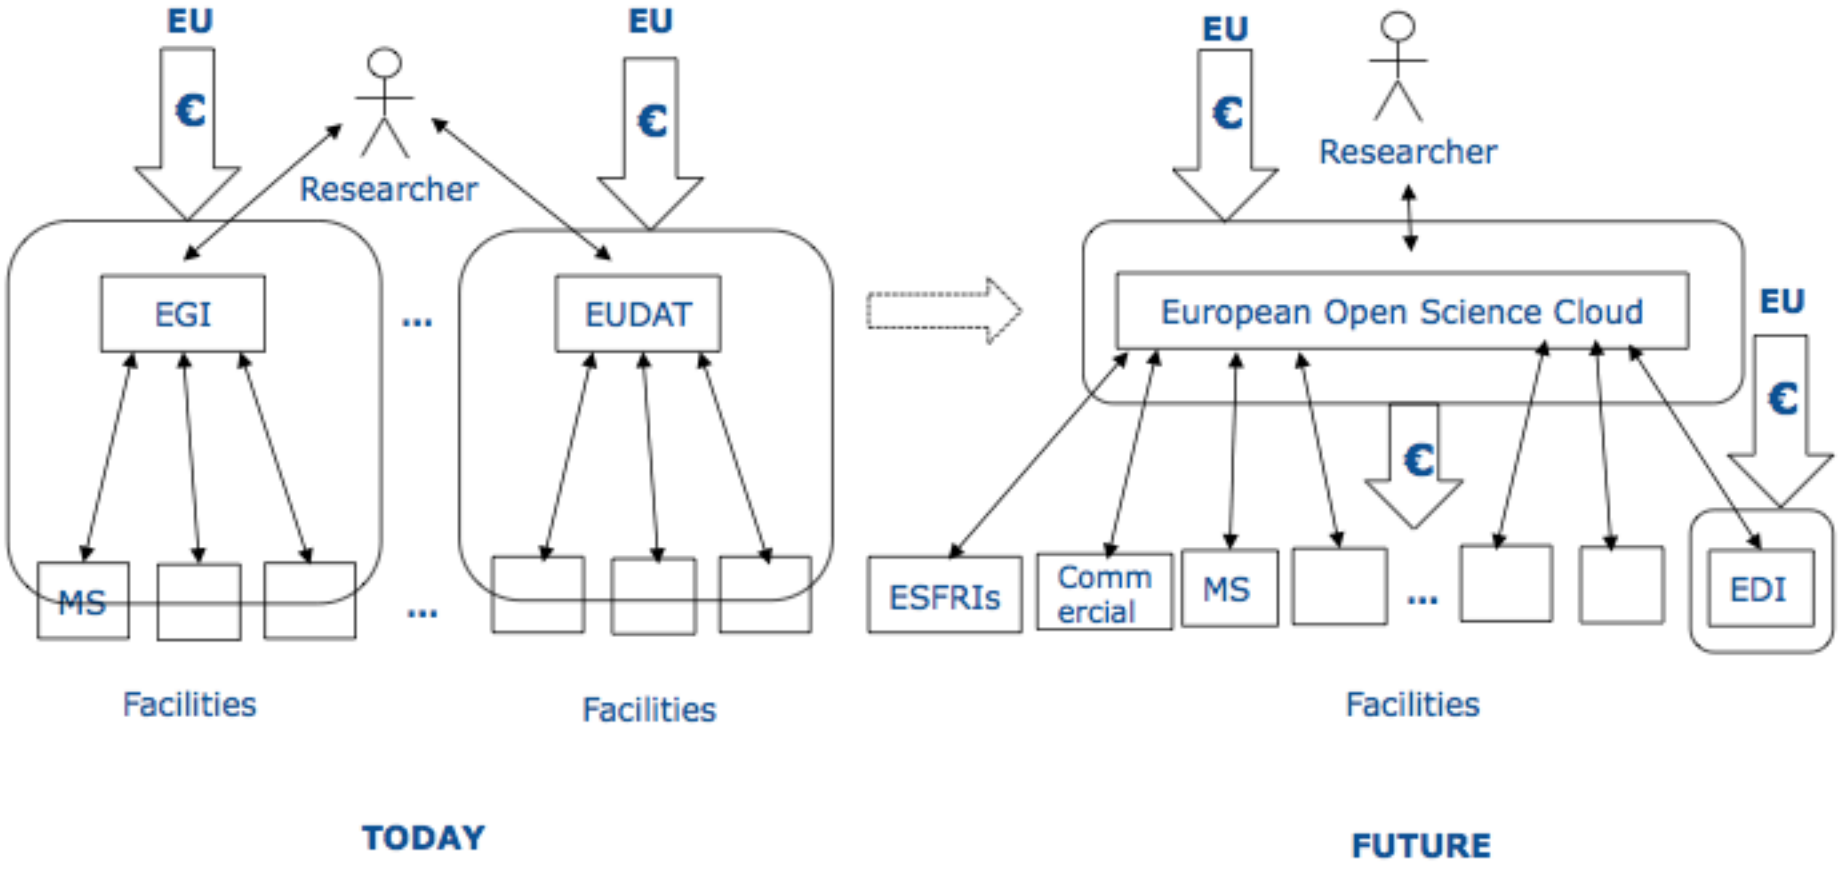

EuroHPC Governance structure

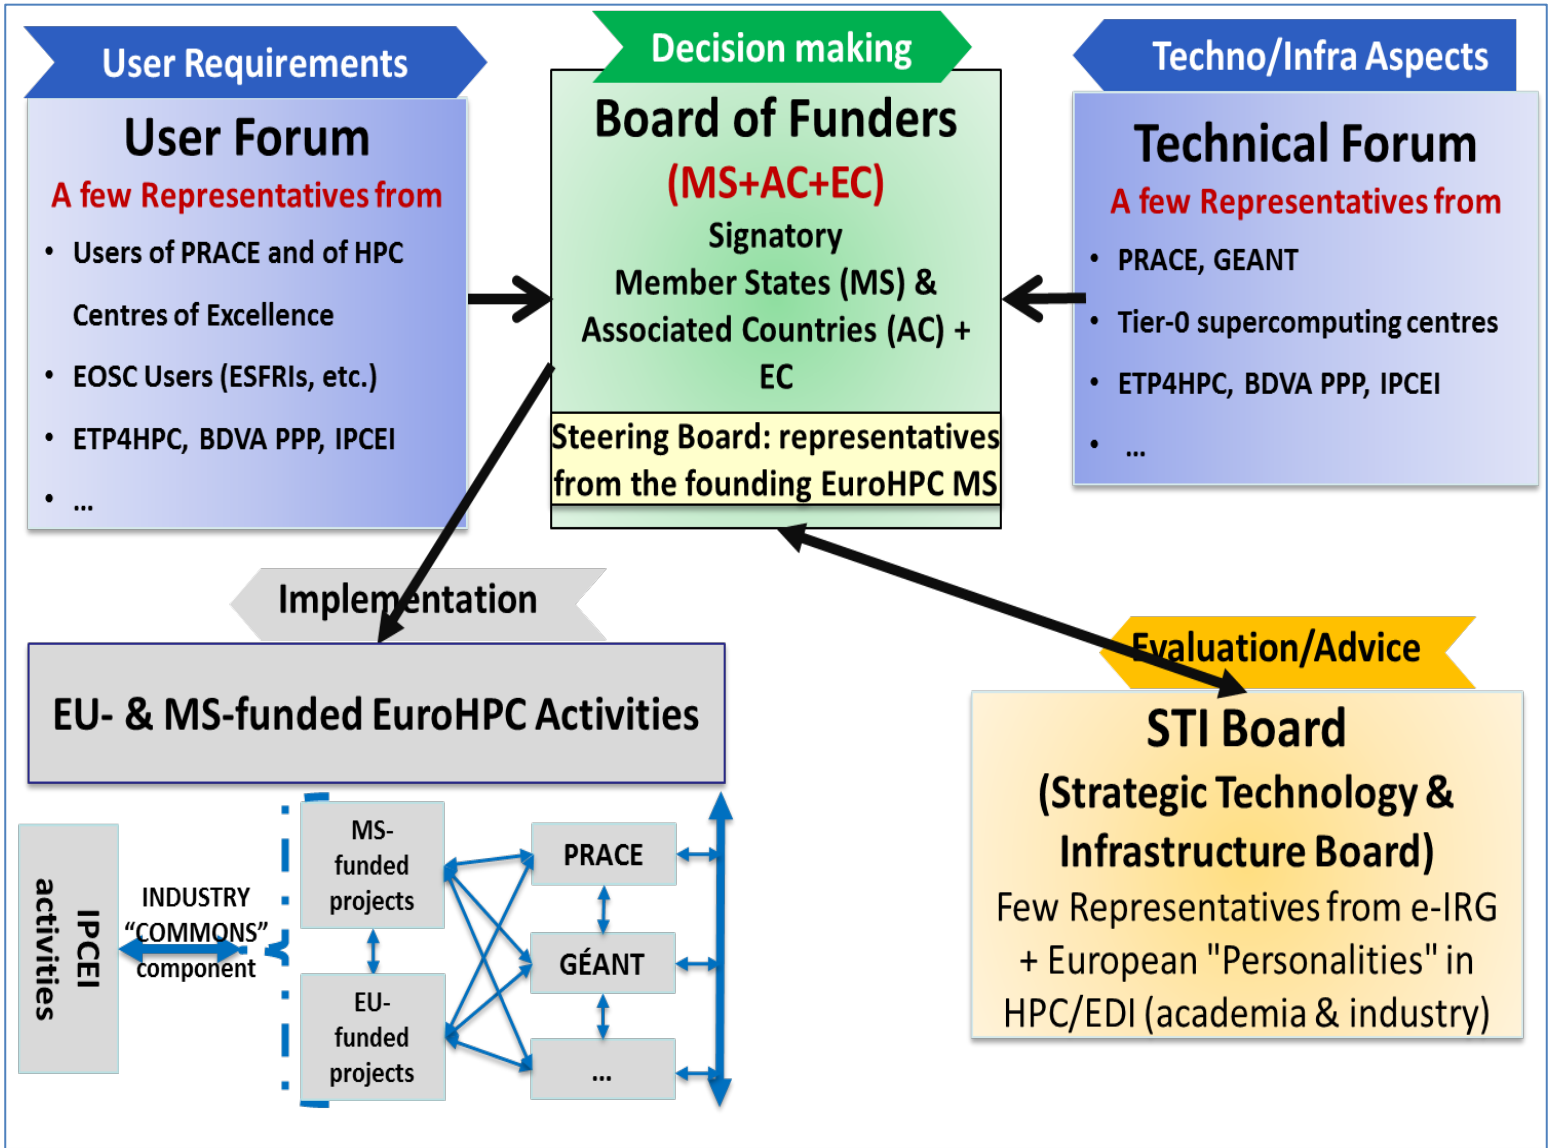

National Big Data Infrastructure

Integration of CINECA-HPC and INFN-HTC computing infrastructure to provide services to:

- Institutional basic and applied research
- Enabling for Public administrations
- Proof of concept and innovation for private organizations and industries

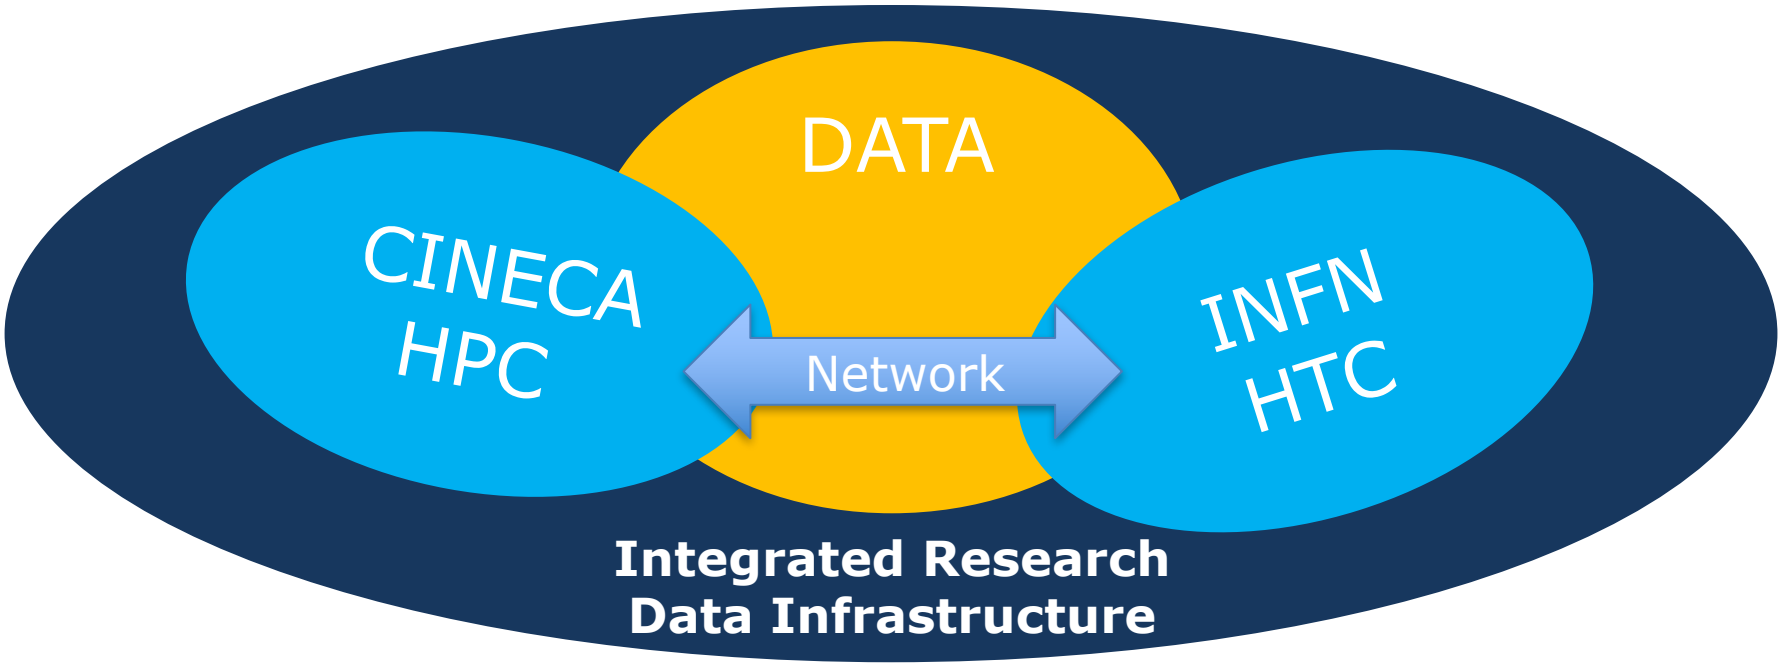

ER Big-data Framework Project

- Goal: leverage Big-Data/HPC HW investment
  - Minimize duplication
    - federated HW Infrastructure Project
  - Deliver enabling value in multiple application domains
    - Vertical Added-Value Booster Projects
    - Test Bed

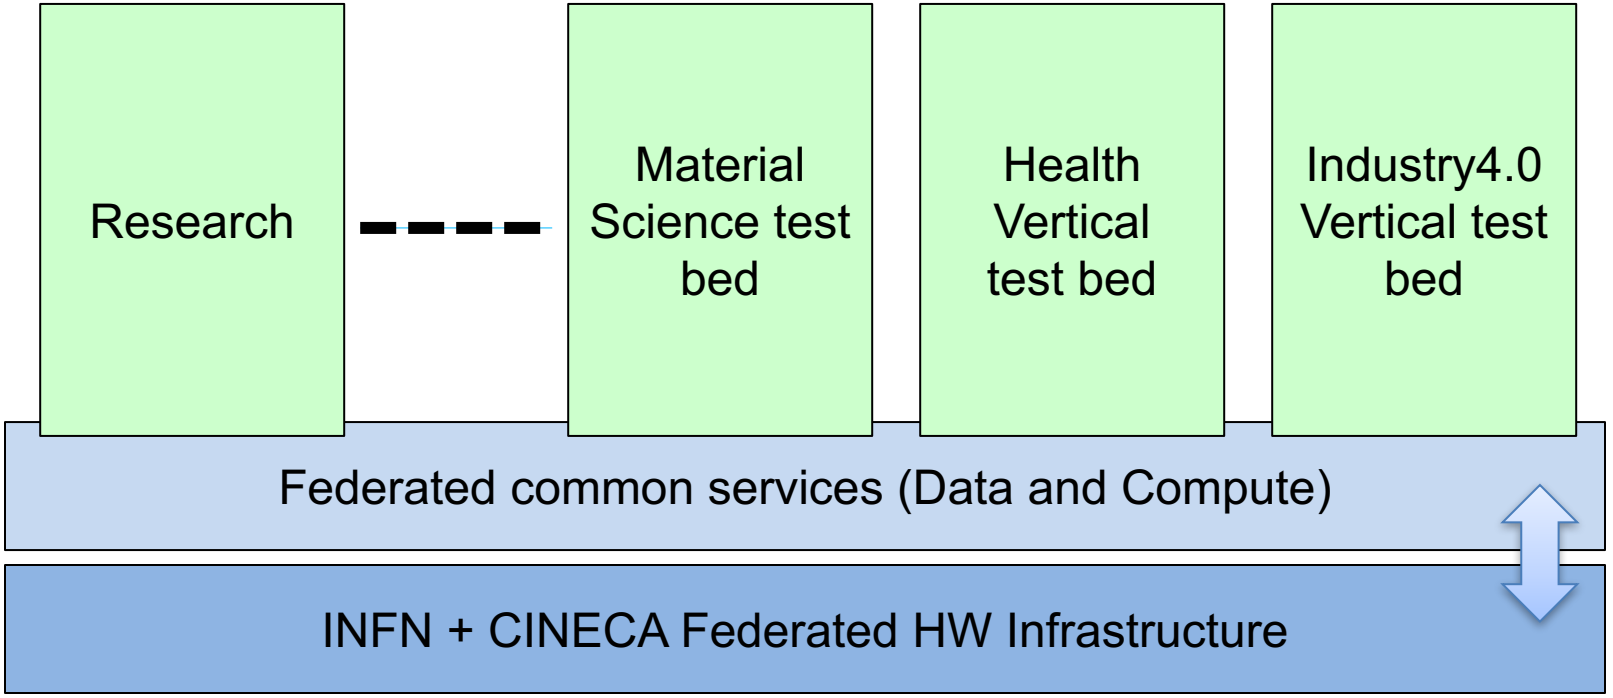

Systems evolution (HPC Cloud)

- Ensure high performance on single node
- Different workloads supported
- Interactive Computer
- Isolated environment for high-security environments

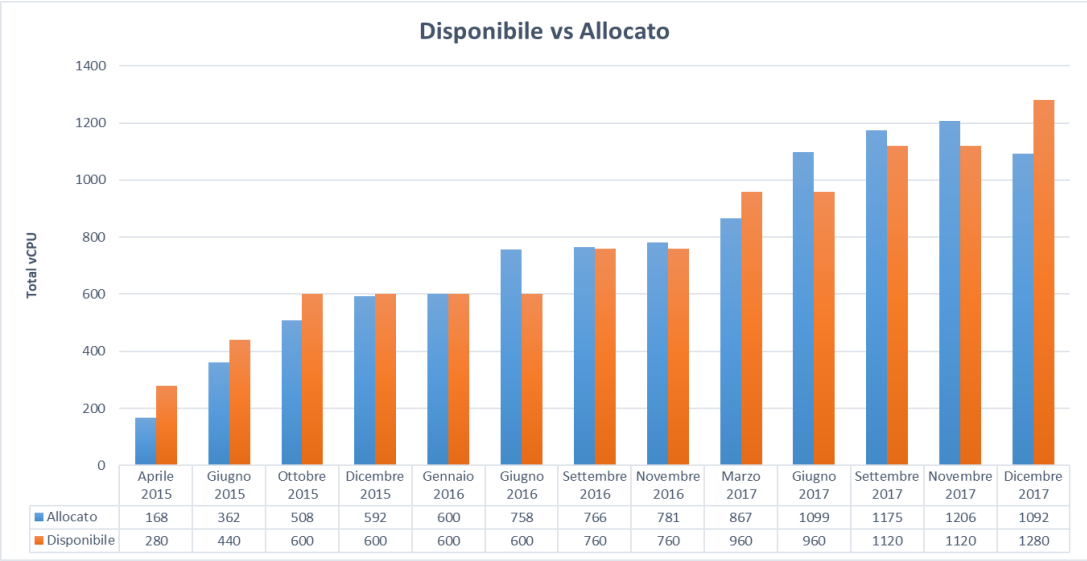

	Aprile 2015	Giugno 2015	Ottobre 2015	Dicembre 2015	Gennaio 2016	Giugno 2016	Settembre 2016	Novembre 2016	Marzo 2017	Giugno 2017	Settembre 2017	Novembre 2017	Dicembre 2017
Allocato	168	362	508	592	600	758	766	781	867	1099	1175	1206	1092
Disponibile	280	440	600	600	600	600	760	760	960	960	1120	1120	1280

Trend di crescita delle risorse assegnate rispetto a quelle disponibili

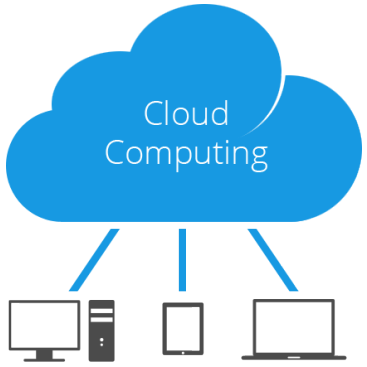

Systems evolution (Data Processing)

The GREEN 500

#18 Nov2017

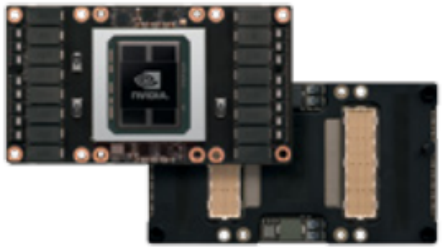

SPECIFICATIONS	
GPU Architecture	**NVIDIA Pascal**
NVIDIA CUDA® Cores	**3584**
Double-Precision Performance	**5.3 TeraFLOPS**
Single-Precision Performance	**10.6 TeraFLOPS**
Half-Precision Performance	**21.2 TeraFLOPS**
GPU Memory	**16 GB CoWoS HBM2**
Memory Bandwidth	**732 GB/s**
Interconnect	**NVIDIA NVLink**
Max Power Consumption	**300 W**
ECC	**Native support with no capacity or performance overhead**
Thermal Solution	**Passive**
Form Factor	**SXM2**
Compute APIs	**NVIDIA CUDA, DirectCompute, OpenCL™, OpenACC**

TeraFLOPS measurements with NVIDIA GPU Boost™ technology

- Result of a PCP (Pre-Commercial Procurement) commissioned by PRACE
- Based on OpenPOWER architecture, using IBM POWER8 processors with NVLink bus and the ultra performing GPGPU NVIDIA® TESLA® P100 SXM2.
  nodes: 45 x (2 Power8 + 4 Tesla P100)
- Low power consumption

PRACE CINECA CSC epcc GENCI JÜLICH FORSCHUNGSZENTRUM

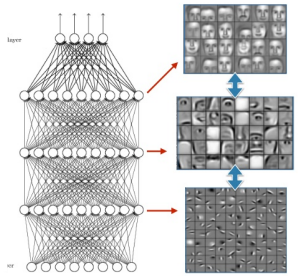

Access to resources

- **Open access peer review**: PRACE, ISCRA, LISA
- **Partnership**: INFN, OGS, IIT, SISSA, INAF, ICTP, Polimi, Bicocca, MI Statale, Unibo,….
- **Open Innovation for Industries and SME**
  - Co-founding Fortissimo, PRACE SHAPE, Lisa for innovation.

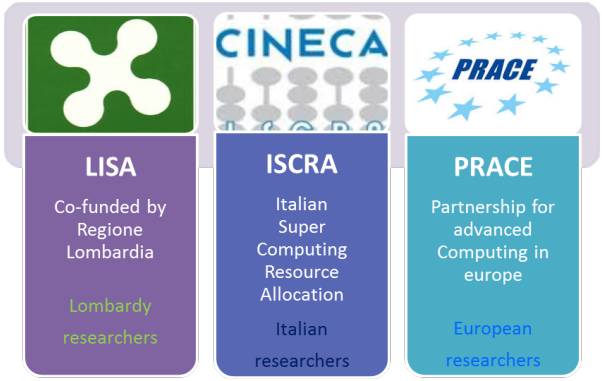

Collaborations with qualified national players

- **Eni E&P Research**
  - HPC system management
  - Production management
  - Applications development
  - Innovation technology
- **Protezione Civile / SMR Regione Emilia Romagna**
  - Operation numerical weather forecast
  - Data Post processing
- **ARPA Piemonte**
  - Environmental numerical forecast
- **ISTAT**
  - Web crawlers
  - Predictive analysis
- **Telethon**
  - Integrated peer review
  - Repository genomic data
- **Human Technopole**

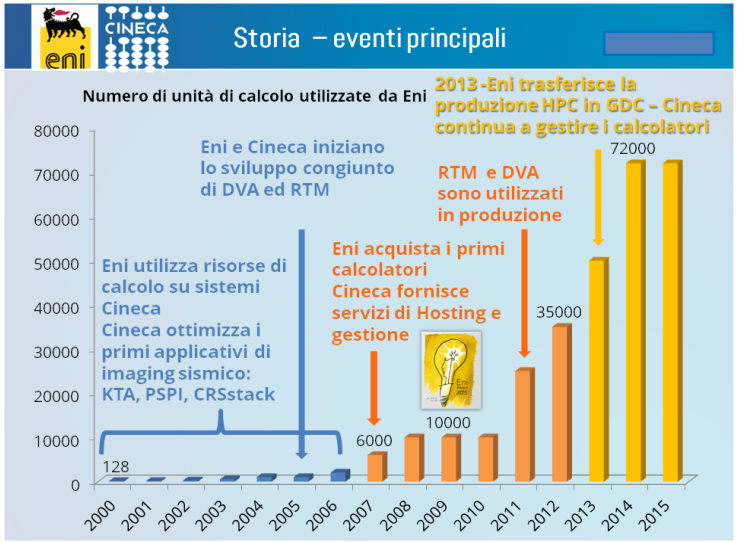

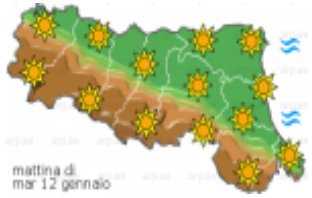

The genesis of the Big Data Platform

Effective actions for enabling Big Data started in 2014 with the adoption of a **data-centric architecture** of the computing center and a new computing system, **PICO**, for the storage and processing of large volumes of data.

A call for **Expression of Interest** supported the first dozen big data projects and new classes of users. Some of these went on under collaboration agreements.

Since then, the **Bioinformatics** community has been growing and several projects have been supported in the **IND4.0** domain and in **other domains** (Digital Humanities, Insurance, e-Government, Media, Energy).

The platform is being developed inside **EU funded projects** and **collaboration agreements**.

The Big Data Platform

Cineca provides:

- A flexible **platform** to accommodate different kinds of users with different requirements and usage models (from GPU intensive workloads on the HPC infrastructure in batch mode, to the more interactive containerized usage of the resources, to the platform as a service available on cloud computing)
- Different options for data storage (from the data lake, unstructured storage, to data base with metadata management, to repositories)
- A vendor independent advice service
- **Expertise** in data science and in a variety of application domains
- Ad hoc data **services** (management, preservation, annotation, analytics, …)

A flexible platform

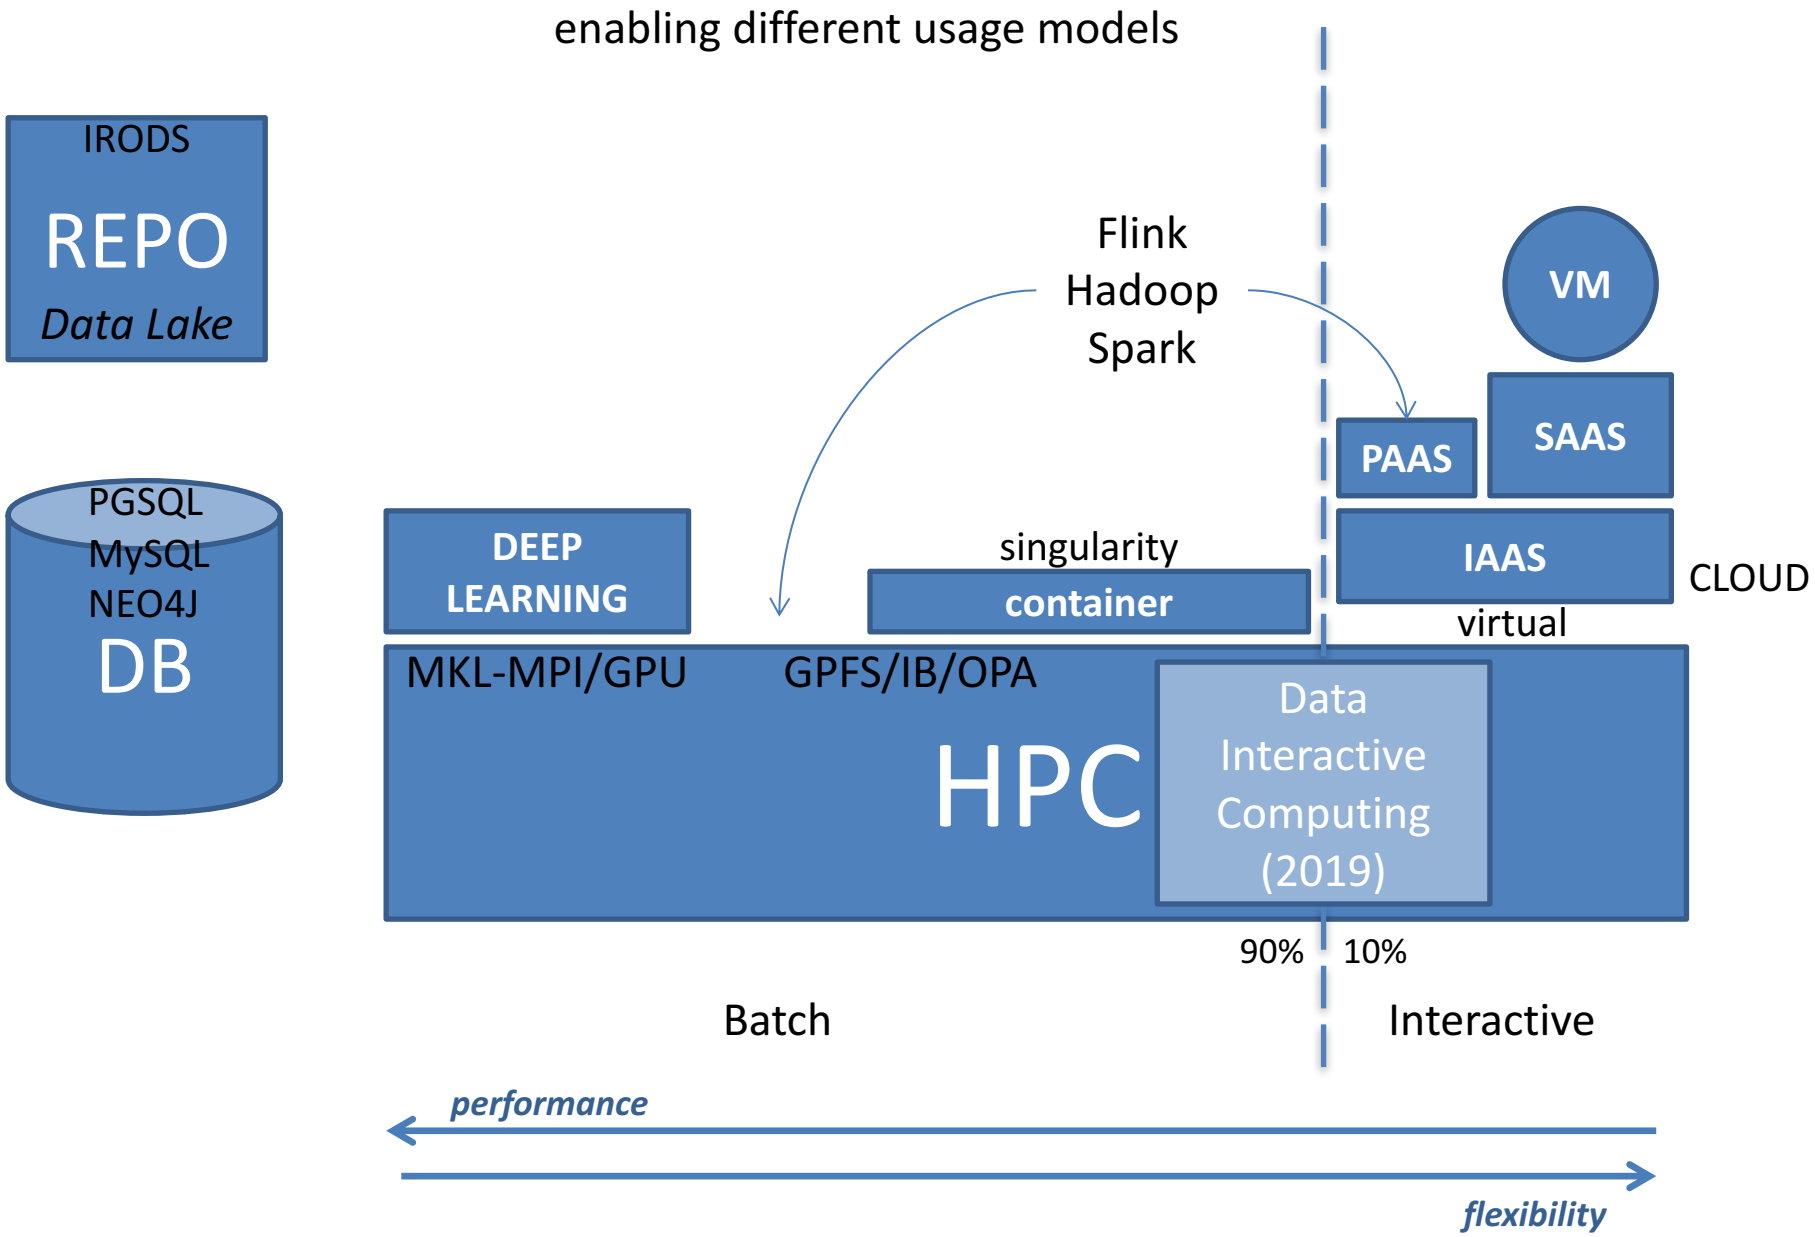

Services

Exploitation of the Infrastructure
open access to HPC/HPDA storage and computing resources, Cloud Computing, computing in batch, interactive and streaming modes

Advanced middleware and software tools

Data management
collection, preparation, annotation, curation, linking, security, access control, long-term preservation, post-processing

Data analytics
Predictive modeling, Supervised and unsupervised learning, Association rules, Sequential patterns, Link analysis, Recommenders, Natural Language Processing, Named Entities Recognition, Information Extraction, Automatic classification, Sentiment Analysis, Semantic metadata generation, Automatic annotation, Speaker segmentation, Automatic Speech Recognition, Video segmentation, Keyframes extraction, Semantic metadata generation from video items, Image recognition

Visualization
Remote visualization, Computer vision and visual computing, Computer Graphics, 3d modeling and rendering, Immersive device programming , Render farm service, Virtual Reality, Augmented Reality, Virtual museum and exhibition design

User support and Specialist support covering different scientific fields, technologies, programming languages, and techniques

Training and Education
Specialized training (workshops on massive data analysis and international summer schools on parallel computing, data analytics and computer graphics), Cooperation with universities (lab activity for master programs, post-doc programs), Knowledge transfer during the projects life cycle

Technology transfer and consulting
Development of proof of concept and innovation projects for businesses to demonstrate the added value and ROIs

Big Data Projects

- EUDAT (data life cycle),
- Telethon and ImediaCity (repository),
- ISTAT (web scraper, text mining),
- Reggia di Caserta (sentiment analysis, virtual exhibitor),
- UNIPOL and ImediaCity (deep learning)
- …

Data explosion

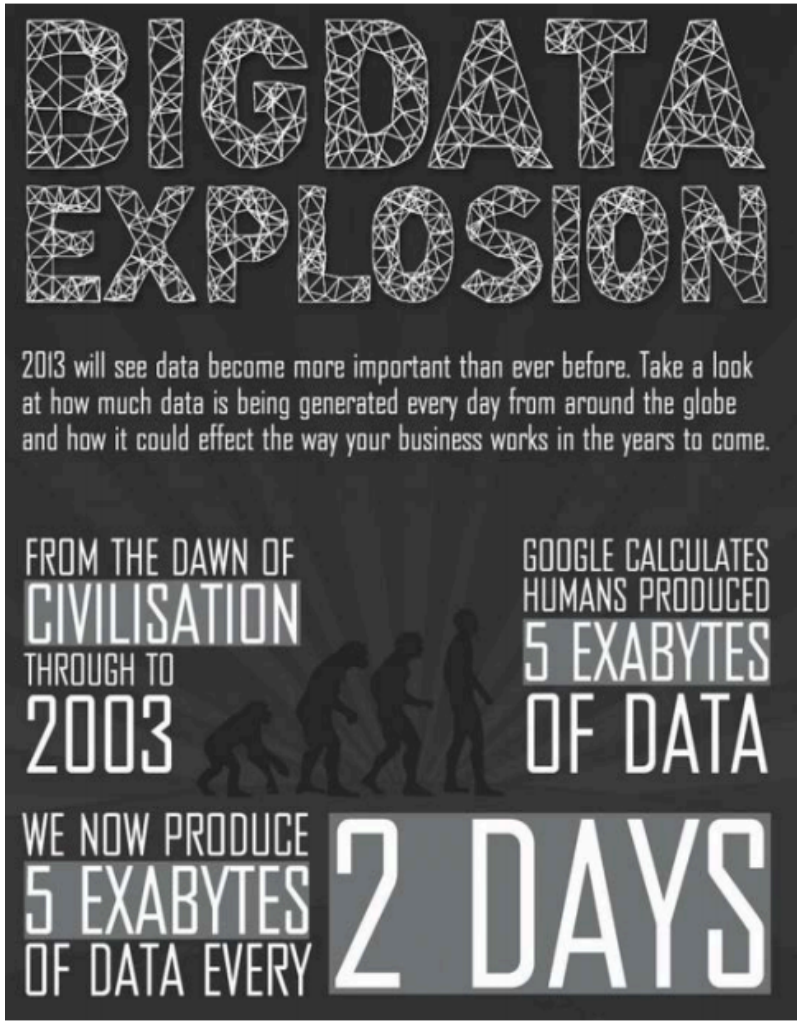

- More and more data is being created
- Issue is not creating data, but being able to navigate and use it
- Data management is critical to make sure data are well-organised, understandable and reusable

Data loss

Digital data are fragile and susceptible to loss for a wide variety of reasons

- Natural **disaster**
- Facilities infrastructure failure
- Storage **failure**
- Server hardware/software failure
- Application software failure
- **Format obsolescence**
- Human error
- Malicious attack
- **Loss of staffing competencies**

Data persistency issues

URL decay in MEDLINE—a 4-year follow-up study

Jonathan D. Wren*
+ Author Affiliations
*To whom correspondence should be addressed.

Received January 22, 2008.
Revision received March 11, 2008.
Accepted April 6, 2008.

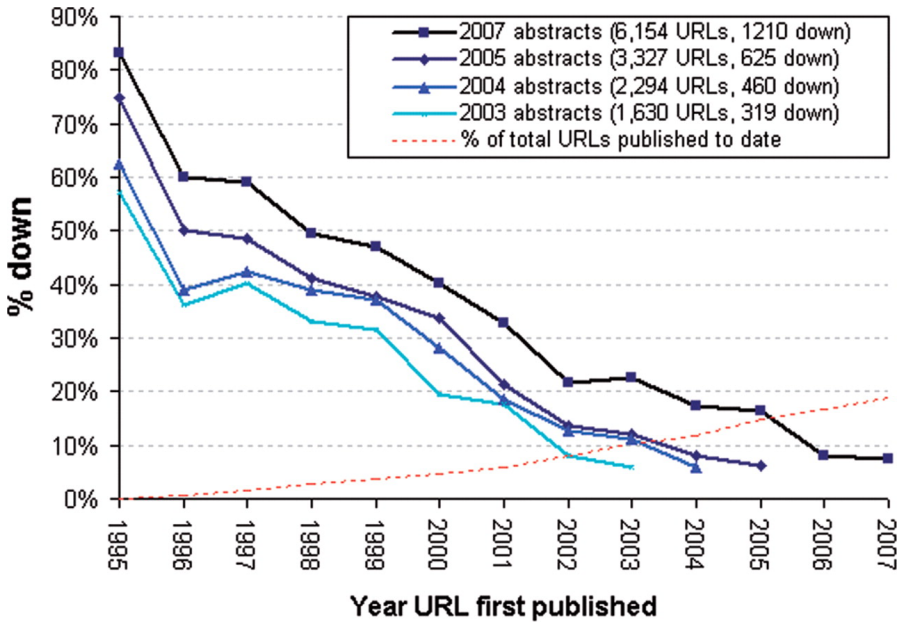

- Link rot – more 404 errors generated over time
- Reference rot* – link rot plus content drift i.e. webpages evolving and no longer reflecting original content cited

* Term coined by Hiberlink http://hiberlink.org

The importance of managing research data

- To make research easier
- To improve data usage and reusability maximizing investments
- To avoid accusations of fraud or bad science
- To share data for others to use and learn from others' results
- To get credit for producing it
- Because funders or your organisation require it, see. H2020 Data Pilot and future programs

Well-managed data opens up opportunities for re-use, integration and new science

RESEARCH DATA – OPEN BY DEFAULT

Projects must have

Provides information on:

the data the research will generate

how to ensure its curation, preservation and sustainability

what parts of that data will be open (and how)

Data Management Activities

DATA COLLECTION
New or existing data are collected or generated

DATA PROCESSING
Preparation of acquired data inputs

DATA QUALITY CONTROL
Measure and monitor data quality

DATA DESCRIPTION
Identify (PID, DOI, …) and document data with extended metadata to allow for understanding, harvesting and consuming the data itself

DATA UPDATE
Change the content of data (correction, add delayed data from sensor, …)

DATA SHARING / DATA PUBLISHING
Make available community data stores through web sites, web services, data catalogues, …

DATA DISCOVERY
Find data based on metadata and /or provenance information

DATA ANALYSIS
Exploration and interpretation of well-managed, processed data for the purpose of knowledge discovery

DATA PROVENANCE
Documenting the various operations that occurred on data to achieve reproducibility and citability

DATA BACKUP
Management of physical risks to the data

DATA LOGN TERM PRESERVATION
Preserving data for long-term use, re-use and accessibility

Data Life Cycle

- Data One definition

*"The data life cycle provides a **high level overview** of the **stages** involved in successful **management and preservation of data** for use and reuse. **Multiple versions of a data life cycle** exist with differences attributable to variation in practices across domains or communities."*

Data Life Cycle

The EUDAT Data Domain

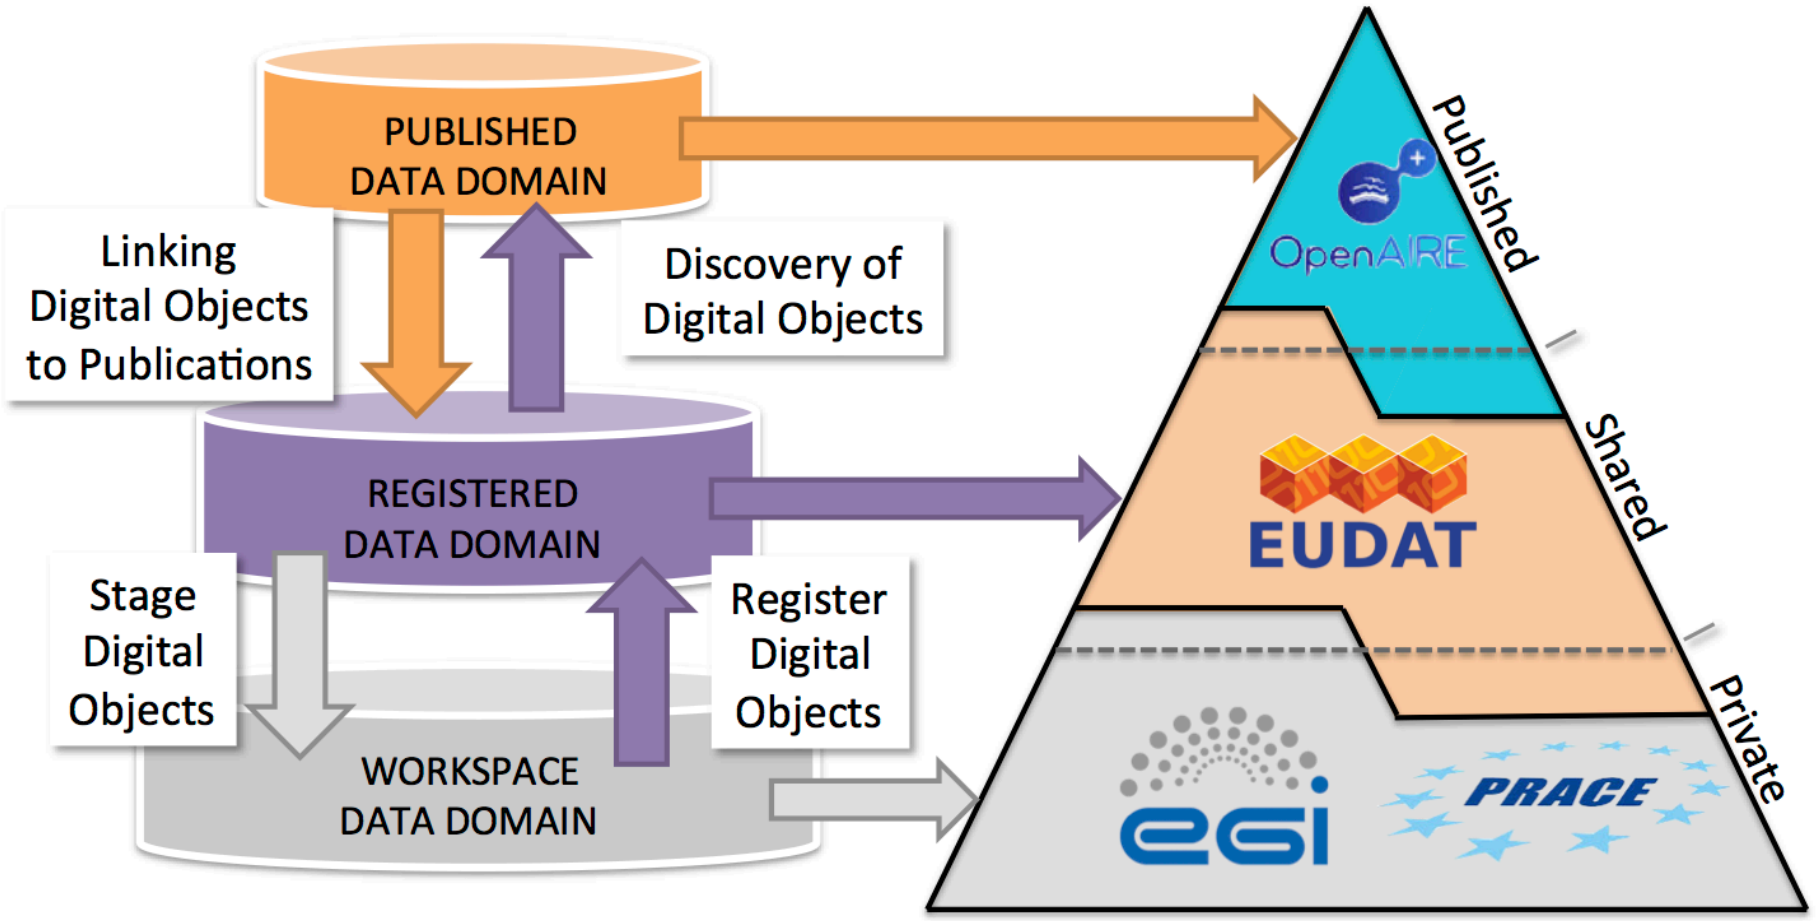

Research Data Life Cycle

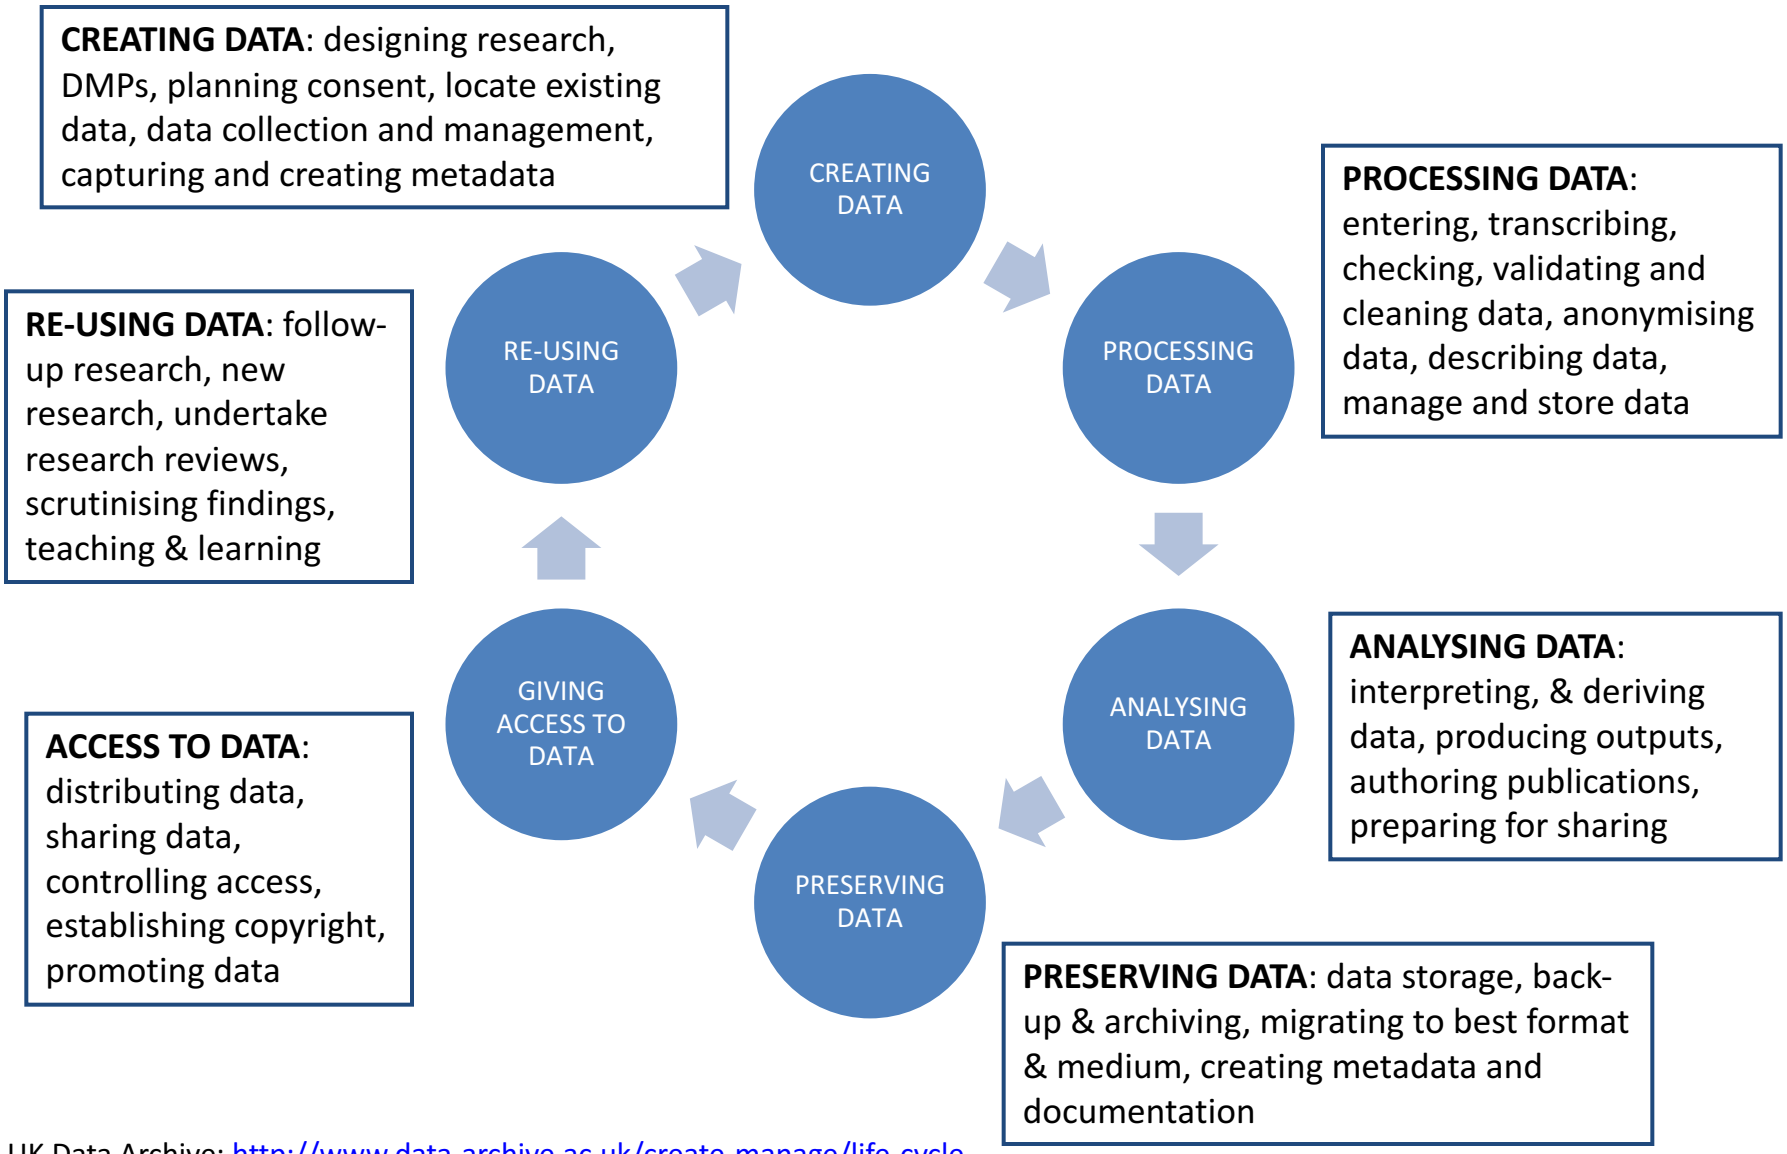

CREATING DATA: designing research, DMPs, planning consent, locate existing data, data collection and management, capturing and creating metadata

PROCESSING DATA: entering, transcribing, checking, validating and cleaning data, anonymising data, describing data, manage and store data

ANALYSING DATA: interpreting, & deriving data, producing outputs, authoring publications, preparing for sharing

PRESERVING DATA: data storage, back-up & archiving, migrating to best format & medium, creating metadata and documentation

ACCESS TO DATA: distributing data, sharing data, controlling access, establishing copyright, promoting data

RE-USING DATA: follow-up research, new research, undertake research reviews, scrutinising findings, teaching & learning

Ref: UK Data Archive: http://www.data-archive.ac.uk/create-manage/life-cycle

What EUDAT does

- Provides and operates an integrated set of data services covering different management aspects
  - ***Data Preservation*** through replication and identification, i.e. PID
  - ***Data Sharing***
  - ***Data Publication***
  - ***Data Discovery***
  - ***Data Access***
  - ***Data Tagging***
- Fosters the definition of a common Data Model in collaboration with other organizations such as RDA
- Trains users and managers in the management of research data including the creation of Data Management Plans
- Connect experts, and representatives from scientific user communities to develop synergic and compound services

Data Life Cycle

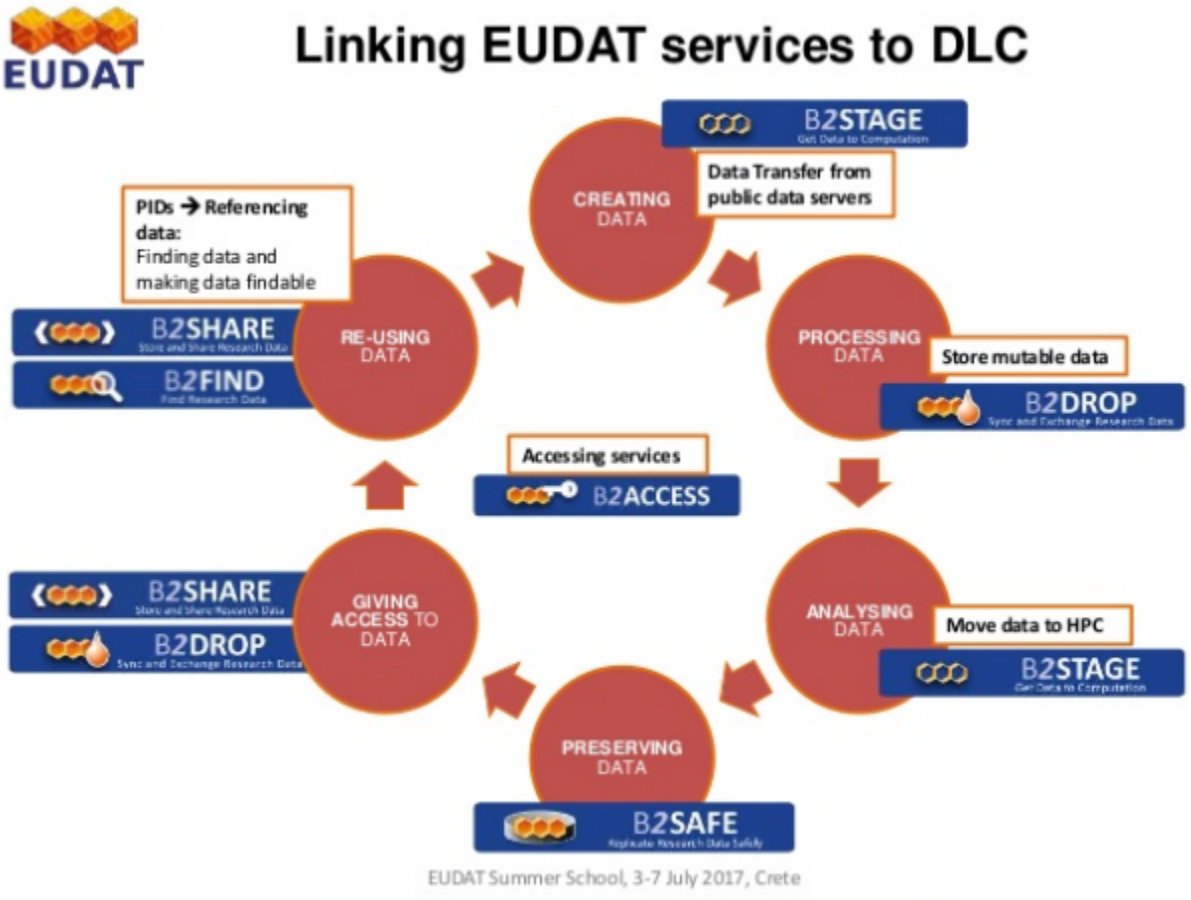

Data analytics software

- Apache Hadoop/Spark related applications (a, b)
  - Spark (Python, Scala and R shells)
  - Spark libraries (MLlib, Graphx, SQL, streaming)
  - Python Machine Learning libraries (Sci-Kit, Pylearn2, …)
  - R (SparkR, BigR)
  - H2O
  - Mahout
- Interfaces/Notebooks (a,b)
  - Jupyter/all-spark-notebook (Scala, R, Python)
  - Spark Notebook
  - Apache Zeppelin
  - Beaker Notebook
  - R Studio server
- High Performance tools
  - Google Tensorflow
  - Intel DAAL – Data Analytics Acceleration Library

Tested software frameworks

- Community based
  - you can clone/fork them from github
- Common approach:
  - NN is built by defining it's computational graph
  - High level language is preferred (Python, protobuf, LUA)
- Intra-node multi-GPU parallelization
- Hardware-independent layer
- Optimized backend engines:
  - MKL and MKL-DNN for Intel based CPUs and many-cores
  - cuBLAS and cuDNN for Nvidia GPUs
- Inter-node scaling:
  - inital support: Tensorflow (gRPC), Caffe (MPI)
  - dedicated communication library (Intel MLSL)

Caffe

theano

The Big Data Platform – Next steps

- The provision of cloud services for big data applications and Application as a service – ongoing
- The provision of Hybrid cloud services (mixing resources of different clouds providers) – planned
- Staffing of the big data and HPDA team – ongoing
- Easy access to resources for industrial users: a call for open innovation – planned
- End-to-end solutions for data processing: managing all phases, from data production or collection to end-user visualization – ongoing
- Provision of innovative platforms for artificial intelligence (ML and DL) – ongoing

Deep Learning at HPC

- Model size is important:
  - winning model of ILSVRC2012 classification task AlexNet: 5 conv. Layers + 3 fullyconn. layers

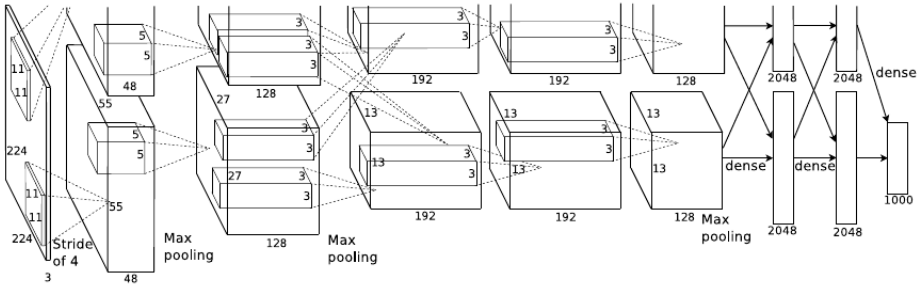

- Data size scaling
  - Larger datasets improve accuracy
- High Arithmetic intensity:
  - Multiple inputs, multiple outputs, batch: GEMM

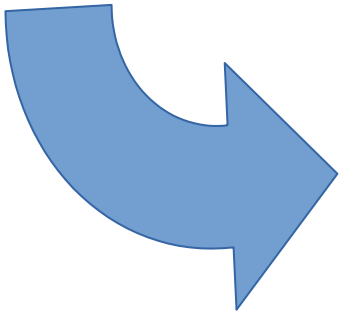

- Dedicated facilities
  - data storage
  - parallel computation hardware (manycores, GPUs)
- Optimized primitives

HPC-BD use case

Hybrid Twin : Wind Turbine use case

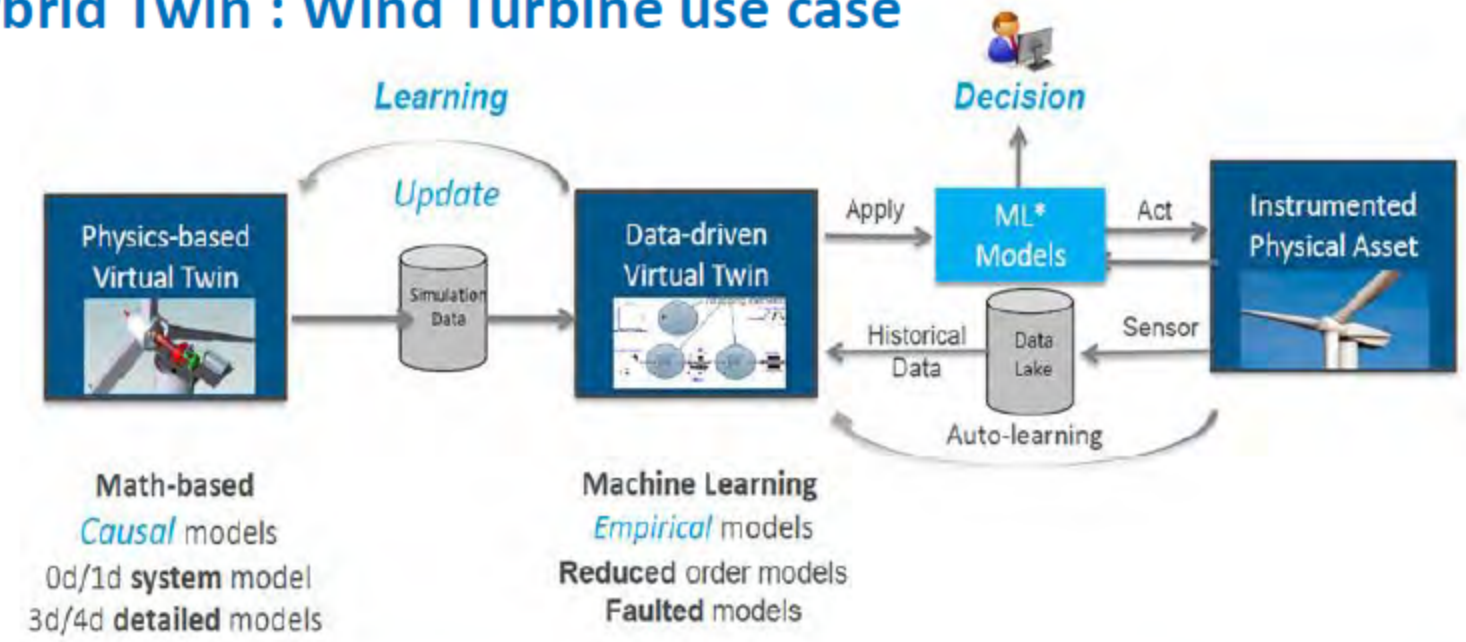

BDVA i-Spaces

Data Innovation Spaces

Accelerate take up of data driven innovation in private sectors like Manufacturing 4.0, Logistics, e Commerce, Media, Aerospace, Automobile, Energy, Agriculture and Agroindustry, Pharma; as well as in non-profit sectors (e-Government, Environment, Pubic Health, Smart Cities).

Provide services to enable and support the **development and validation of new BigData use-cases** (and deliver Economical, Societal and Environmental Value to their local ecosystems)

Data experimenting platforms that **host Closed as well as Open Data** from Business and Public sources (language resources, geospatial data, healthcare data, economic statistics, transport data, weather data…) for the purpose of the execution of **pre-competitive Proof-of-Concept projects**.

i-Spaces offer trusted and secure environment allowing Research, Education and Innovation stakeholder to innovate with data, acting as hubs to **connect different stakeholders**.

With their infrastructure, tools and skill offerings i-Space help the surrounding (industrial and public) ecosystem to overcome high-risk investment burdens for faster adoption of what new BigData technologies can offer to them.

i-Space characteristics

- cross-organizational
- multi-sectorial
  - Opposite to light-house projects
- Access to (valuable) data
  - Focus on private data (not only public data)
- Secure, trusted environment
  - Ownership of data preserved
  - Data in Europe
- Support data sharing
  - For multiple purpose
  - multi-lingual
  - Pan-European
- Support data integration
- Support for legal aspects
- Support for start-ups, web entrepreneurs
  - Access to incubation

- High-performance infrastructure (Hardware and Software Platforms)
- Support for executing Tests, trials or validation on Big Data sources
- Ecosystem
  - Usage of data for Research, Innovation, Education

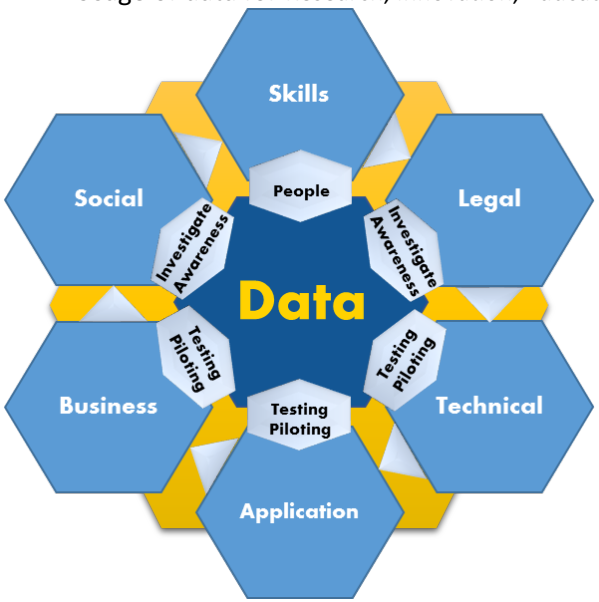

BDVA labelled i-Spaces

BDV i-Space BRONZE

BDV i-Space SILVER

BDV i-Space GOLD

- Excellence of INFRASTRUCTURE
- QUALITY of SERVICES
- PROJECTS & SECTORS
- IMPACT to ECOSYSTEM
- BUSINESS strategy

BDVA labelled i-Spaces

- Excellence of INFRASTRUCTURE
- QUALITY of SERVICES
- PROJECTS & SECTORS
- IMPACT to ECOSYSTEM
- BUSINESS strategy

BDV i-Space BRONZE

BDV i-Space SILVER

BDV i-Space GOLD

Smart Data Innovation Lab

INSTITUT Mines-Télécom

GROUPE-GENES TERALAB

RISE

BIGDATACoE BARCELONA by eurecat

KNOW Center

egi

CINECA

ITI INSTITUTO TECNOLÓGICO DE INFORMÁTICA


thank you!