RETHINK big

Roadmap for European Technologies
in Hardware and Networking for Big Data

RETHINK big Project

Gina ALIOTO, Adrián CRISTAL, Osman UNSAL

BDEC for Europe Workshop, Barcelona

28 Jan 2015

www.rethinkbig-project.eu

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no 619788.

Project Overview

- **The Partners**

INDUSTRY-DRIVEN

- **The Project**: 2-year, Coordination and Support Action (CSA), Coordinated by BSC, Start: 1 Mar 2014

CO-DESIGN OPPORTUNITIES

- **The Mission**: To deliver a strategic roadmap for how technology advancements in *hardware*, *networking* and *algorithms* can be exploited for Big Data analytics.

Why Hardware matters?

- **All Software runs on hardware….**
  - MapReduce
  - CUDA
  - OpenCL
  - MPI
  - OpenMP
  - …
- **You might not want to admit it but…**
  - You are constrained by hardware and the network

- **In the next 10 years, the HW will change more dramatically than it has in the past 10 years**

- **HW will influence the products and services that you provide**

Challenges

- **Work with different areas**
  - Applications and end users
  - Software Tools
  - Systems
  - Network
  - Hardware
- **Work with different requirements**
  - Velocity
  - Volume
  - Variety
  - Real Time
  - Sensors
  - Power consumption…

U.S. Big Data Scene

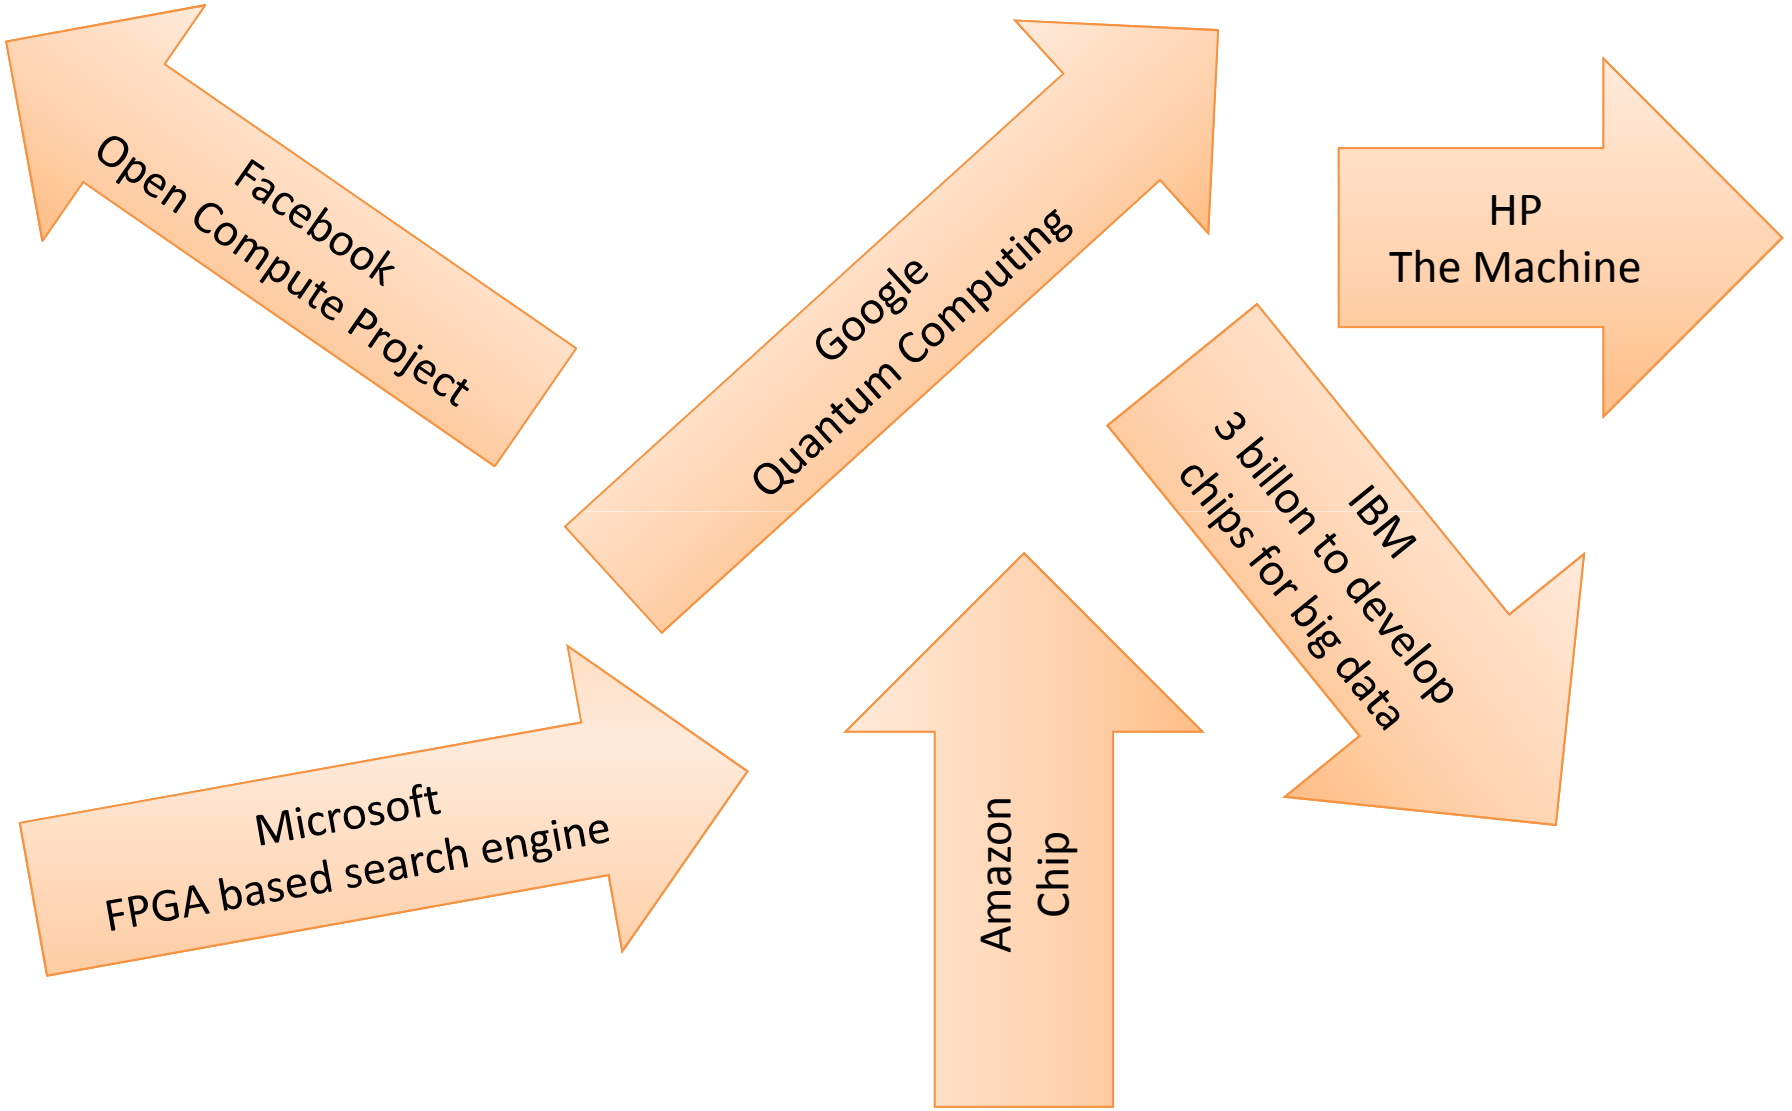

EU State-of-the-art

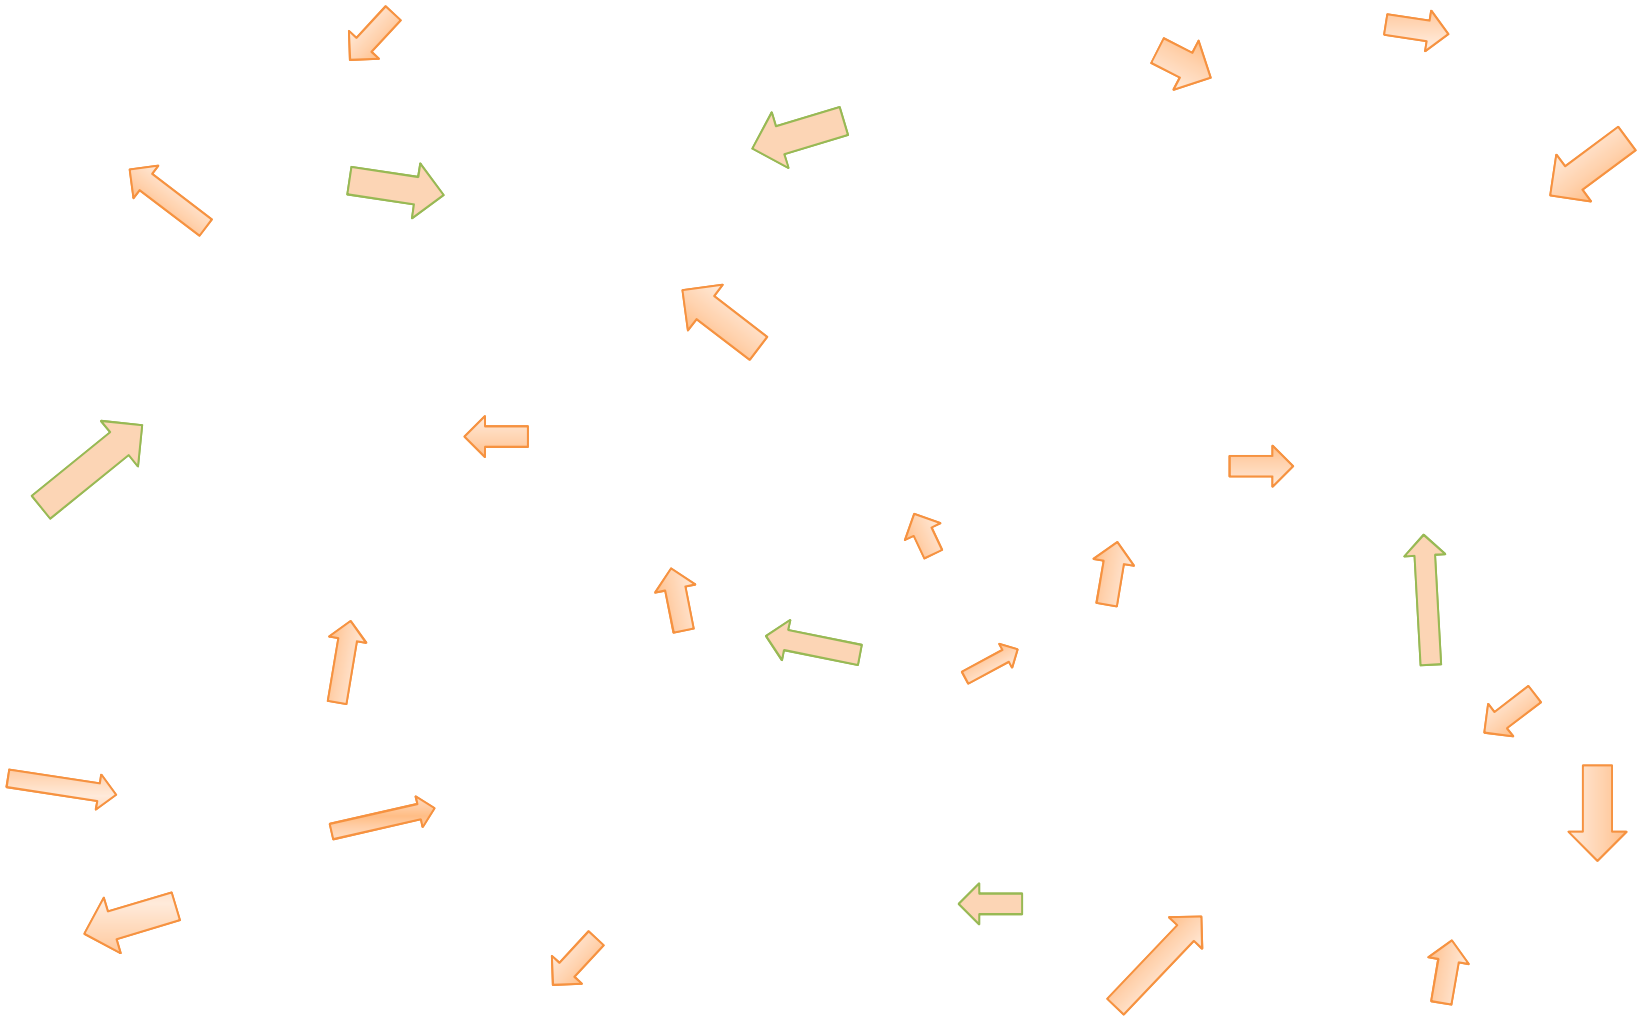

EU Possible Approach

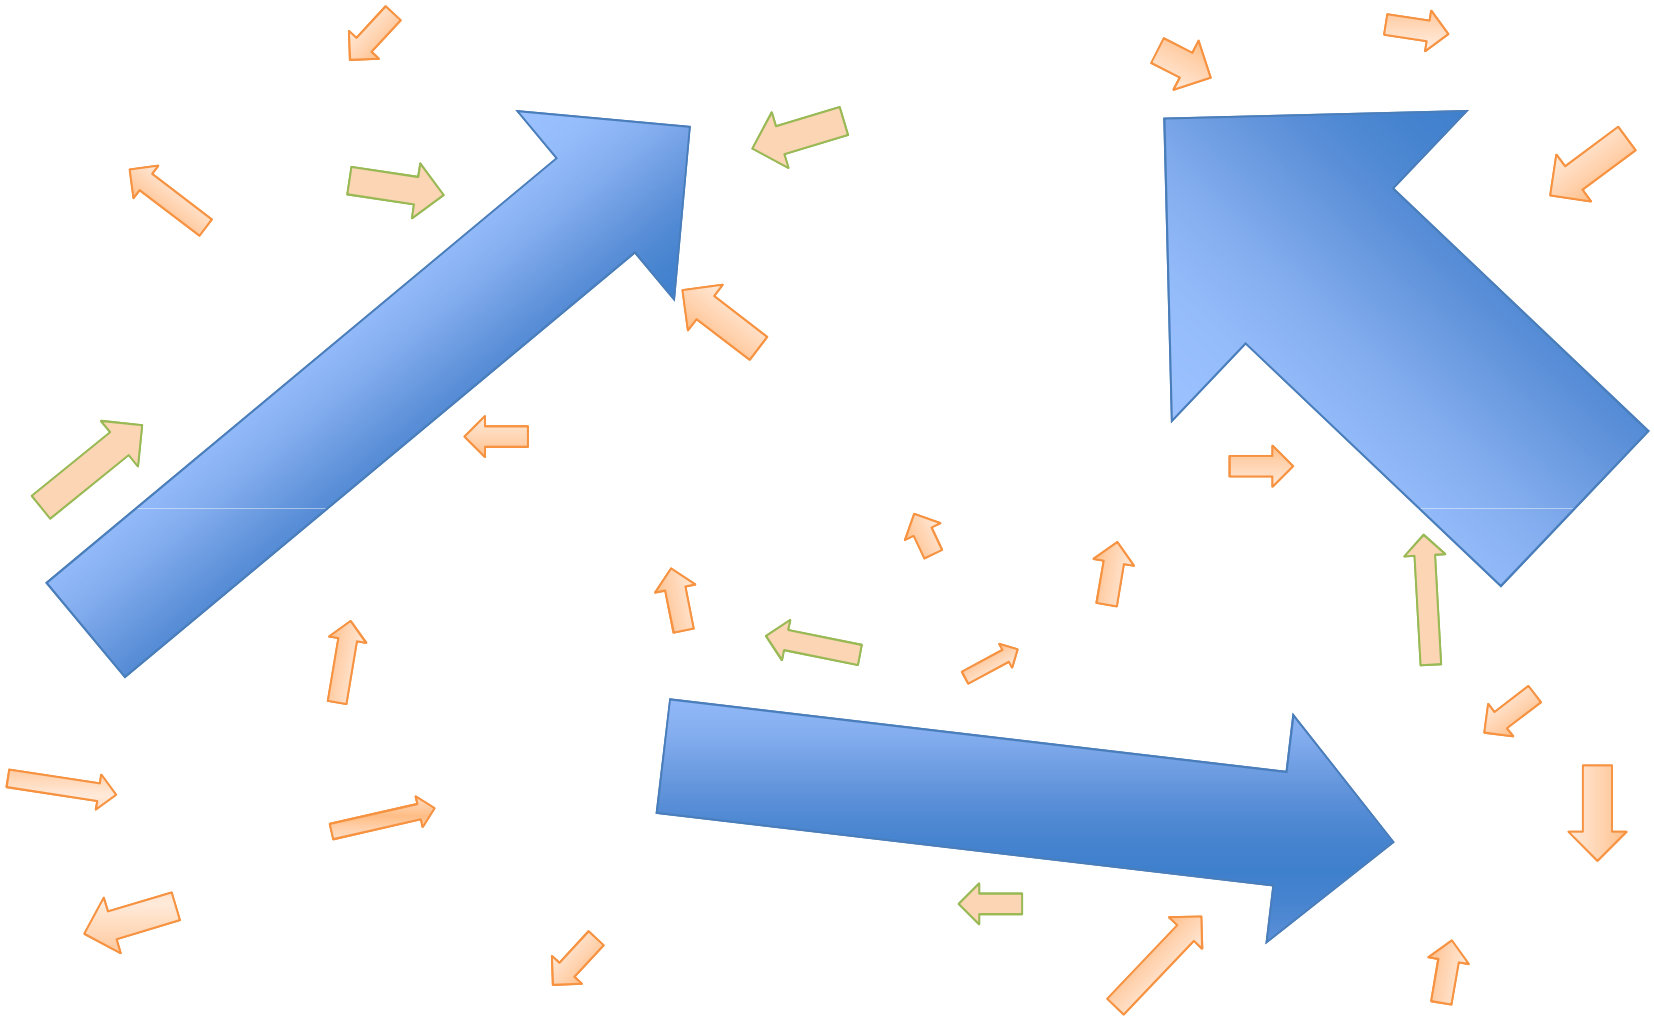

Key: Hardware/Software Holistic Design

- Hardware needs to be software-aware
- Software needs to be hardware-aware

What happens if HW does not consider SW

- **Many (supposedly great) changes in HW architecture do not survive**
  - Cell processor (Playstation 3 processor)
    - Master-Slave processor model programmed using DMAs
      - -> Extremely difficult for programmers
  - Itanium processor (VLIW)
    - Very Long instruction word explicitly harnesses instruction level parallelism through Compiler
      - -> Compilers could not extract required parallelism

What happens if SW does not consider HW

- **Terasort contest: sorting 100TB data**
  - Number 1: Vanilla Hadoop
    - 2100 nodes, 12 cores per node, 64 Gb per node
      - 24.000 cores
      - 134 Tb memory
    - 52 nodes, [illegible] cores per node, 24 Gb
      - 416 cores
      - 1,2 Tb memory
    - Time: 8300 secs and 6400 secs
    - Cost in Amazon: $ 294 and 226

Vanilla Hadoop is [illegible] program, but needs 57X more [illegible], 100X more memory, and [illegible]ts 2X performance

As of March 2014

Application Challenges

- Science and Engineering Applications
- Life Sciences
- Future Internet and Social Networking
- Business, Finance, Information Marketplaces

Enabling Technologies

- Conventional / Unconventional HW and processing technology
- Distributed Architectures, Devices and Sensors, Memory and Storage
- Networks
- Frameworks, SW Models, Algorithms, Data Stuctures and Visualization

First Working Group Meeting: 18,19 Sep 2014

- **Objectives:** Identify challenges across European Big Data sectors, Develop a shared language, Engage key strategists
- **Attendees:** 70 Experts from 49 Organizations, 38 External

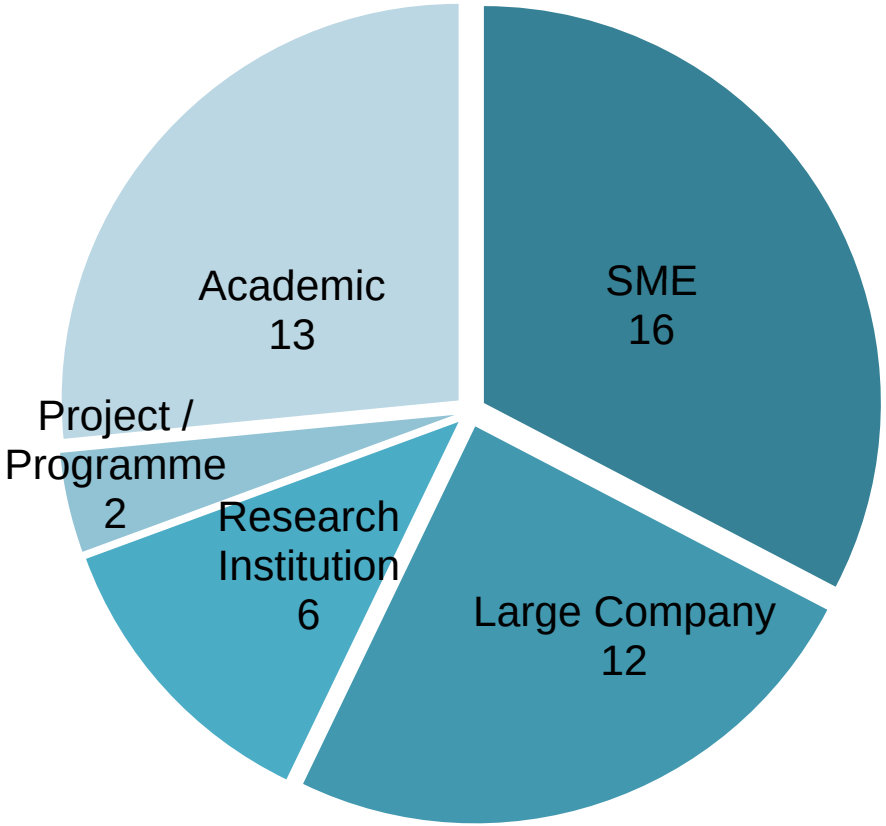

INDUSTRY PARTICIPANTS INCLUDED:	
Users	***THALES***, AIRBUS, Boehringer-Ingleheim, AGT International, Capgemini, Cloud&Heat, The Unbelievable Machine Company, NextWorks
Providers	***ARM, THALES, Alcatel Lucent Bell Labs,*** Telefonica, T-Systems, Bull, TT Tech, Lacie (Seagate), Kalray, Okkam

First Working Group Meeting: Results

- We are concerned with BIG DATA and…
  - The economics of memory from SSD and Nand Drives to PCM, FeRAM, MRAM, PRAM…
  - Memristors
  - Analog Computing
  - Neuromorphic computing
  - 3-D stacking
  - Optical computing
  - Neural network algorithms
  - Accelerators
  - …and beyond

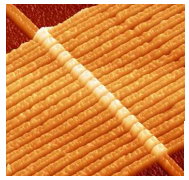

HP Labs: An image of a circuit with 17 memristors captured by an atomic force microscope.

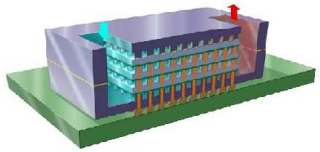

Figure from EPFL
http://esl.epfl.ch/page-58161-en.html

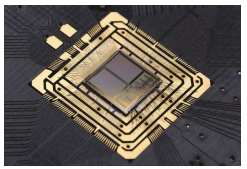

Neuromorphic chip developed at Bernstein center

The world in 3D (3D Stacking)

- Very large bandwidth
- Very low latency
- Large amounts of memory on a chip
- Locality will be extremely important
- Thermal problems

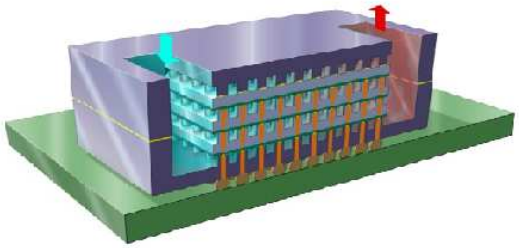

Figure from EPFL
http://esl.epfl.ch/page-58161-en.html

Non-volatile memory

- **New technologies**

  (STT-RAM, CB-RAM, RRAM, …)
- **More density**
- **Replacement for DRAMs**
  - Endurance problem
- **Large influence on software**
  - Data base systems
  - File systems

Dark Silicon Era

- **Thermal problems**
  - Not all cores will be able to be on all the time
- **Extensive use of Accelerators**
- **Reconfigurable computing**

EU SWOT

- **Strengths for Big Data HW**
  - Embedded and real-time systems -> IoT
  - Low-power -> Datacenter
  - Health Care, security, trust, privacy
  - Machine Learning
  - HPC expertise (PRACE, CERN)

- **Weaknesses**
  - Culture of distributed collaboration
  - Device/circuit, Algorithms

Many potentially disruptive SMEs, research

- Machine Learning IC (Artificial Learning)
- Servers as heaters (Qarnot, Cloud&Heat)

- HLS for FPGAs (Synflow)
- Debugging for FPGAs (Yugo Systems)

- Graph Database Acceleration (Neo Tech.)

- DNA memories (European Bioinformatics)
- Optical computing (Optalysis)

RETHINK big Project

Big Data Value cPPP

www.rethinkbig-project.eu

This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no 619788.

BDV cPPP: Structure

- EU and Industry agree on a cPPP to conduct strategic Big Data Value research and innovation projects
- Requires setting up of a Big Data Value association
- Must involve broad stakeholder community

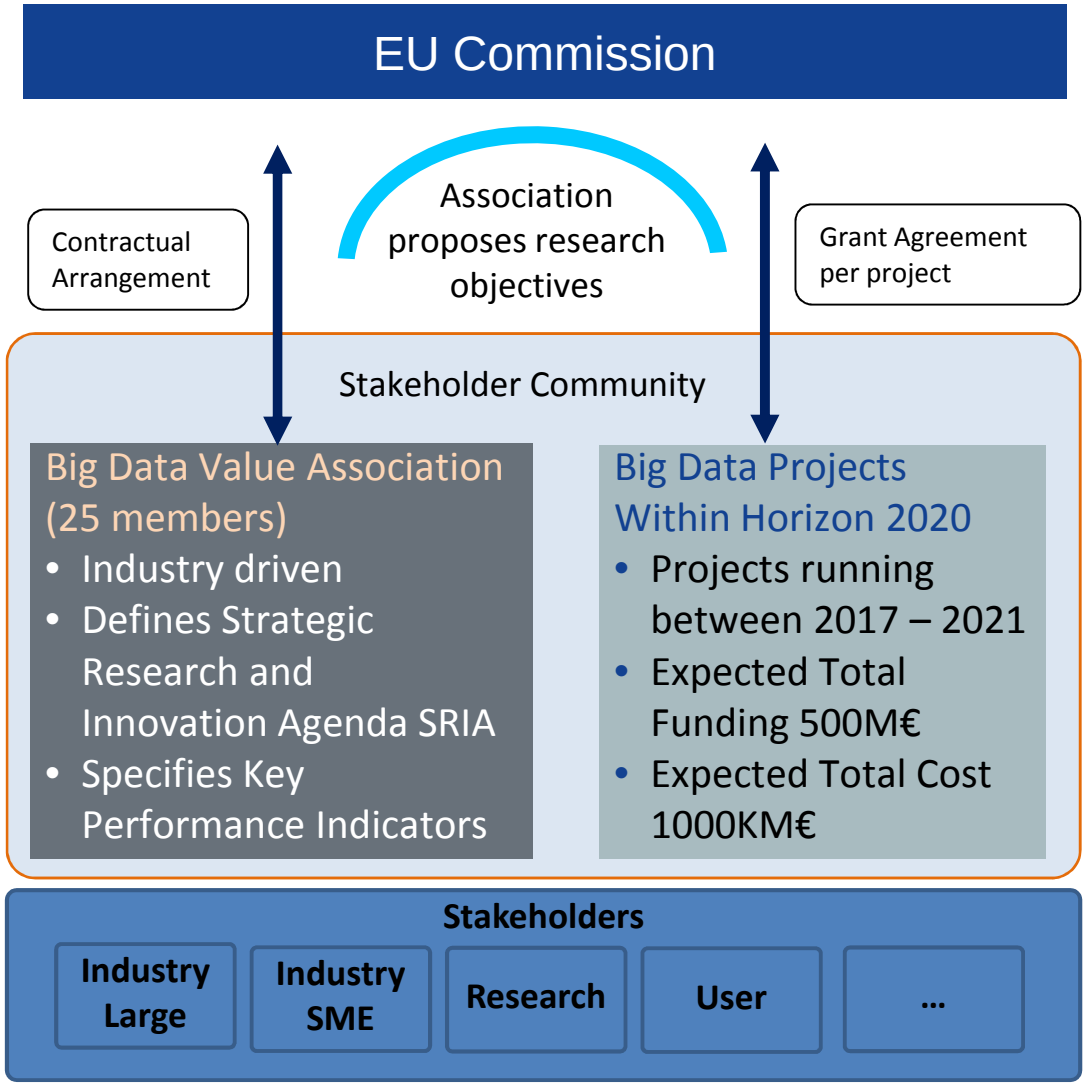

RETHINK big Involvement in BDV cPPP

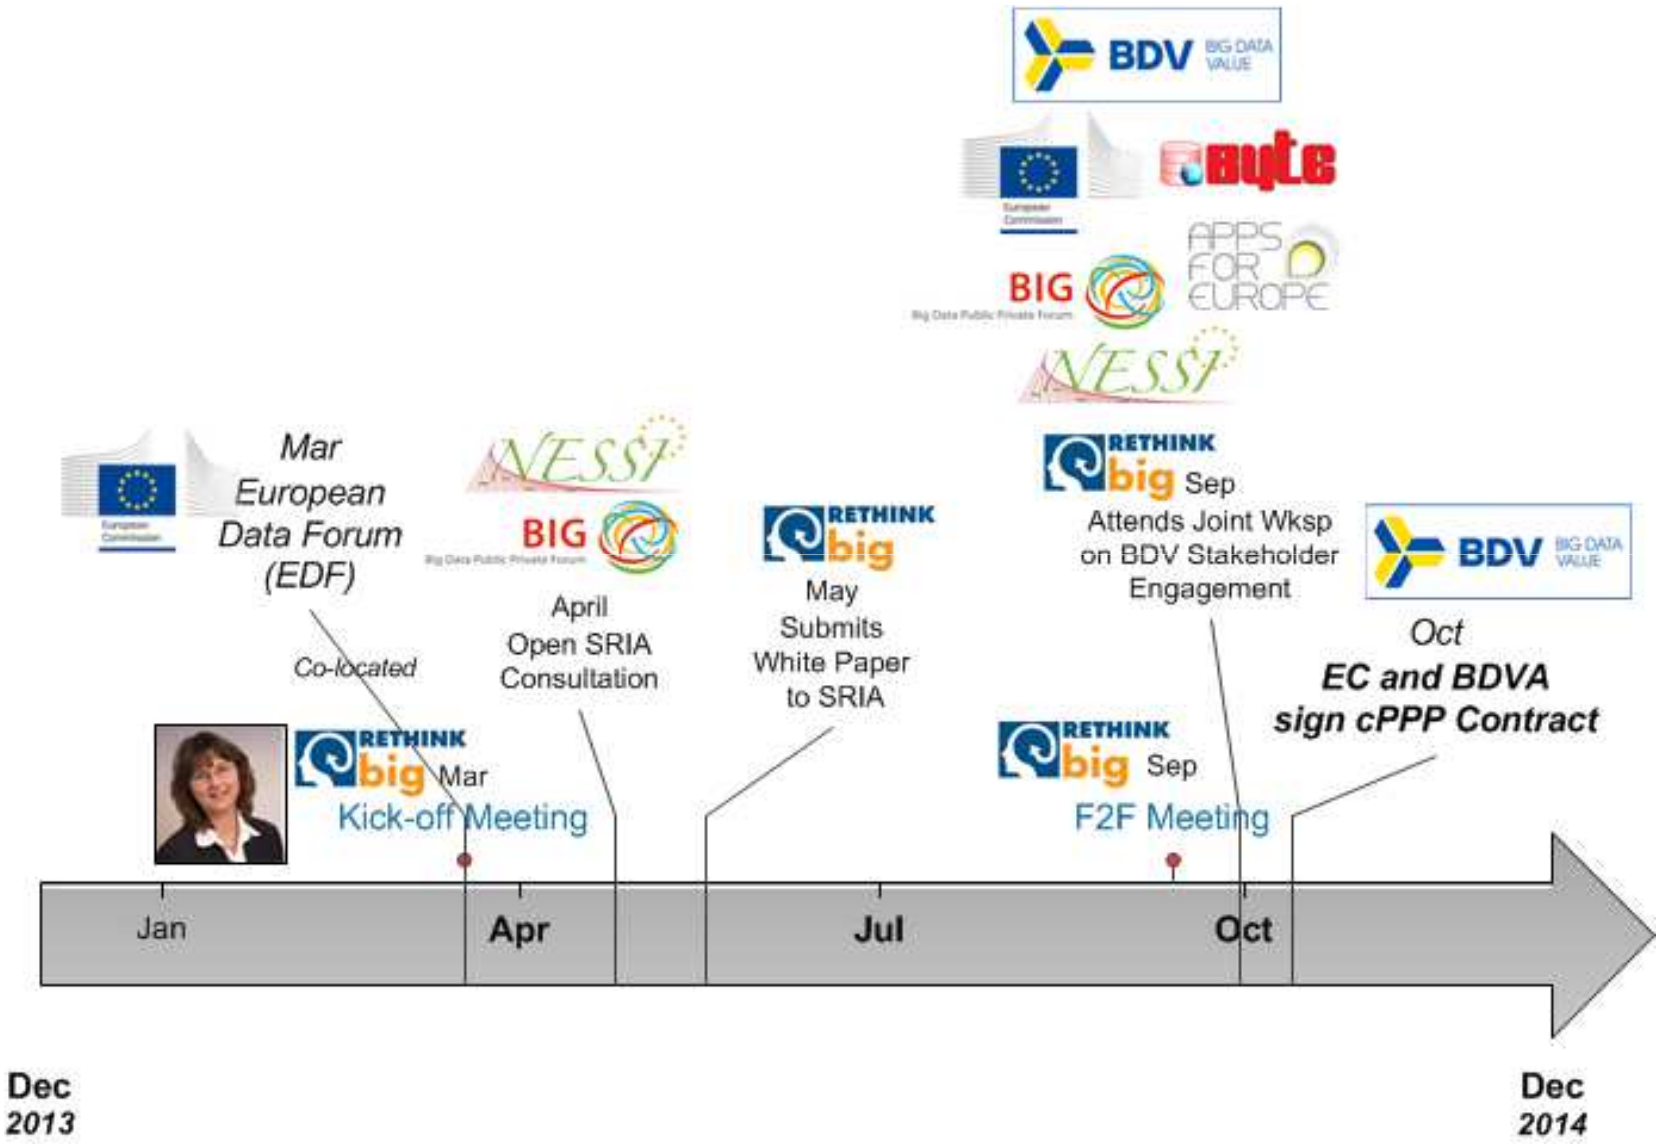

BDV cPPP Content: Multidisciplinary approach

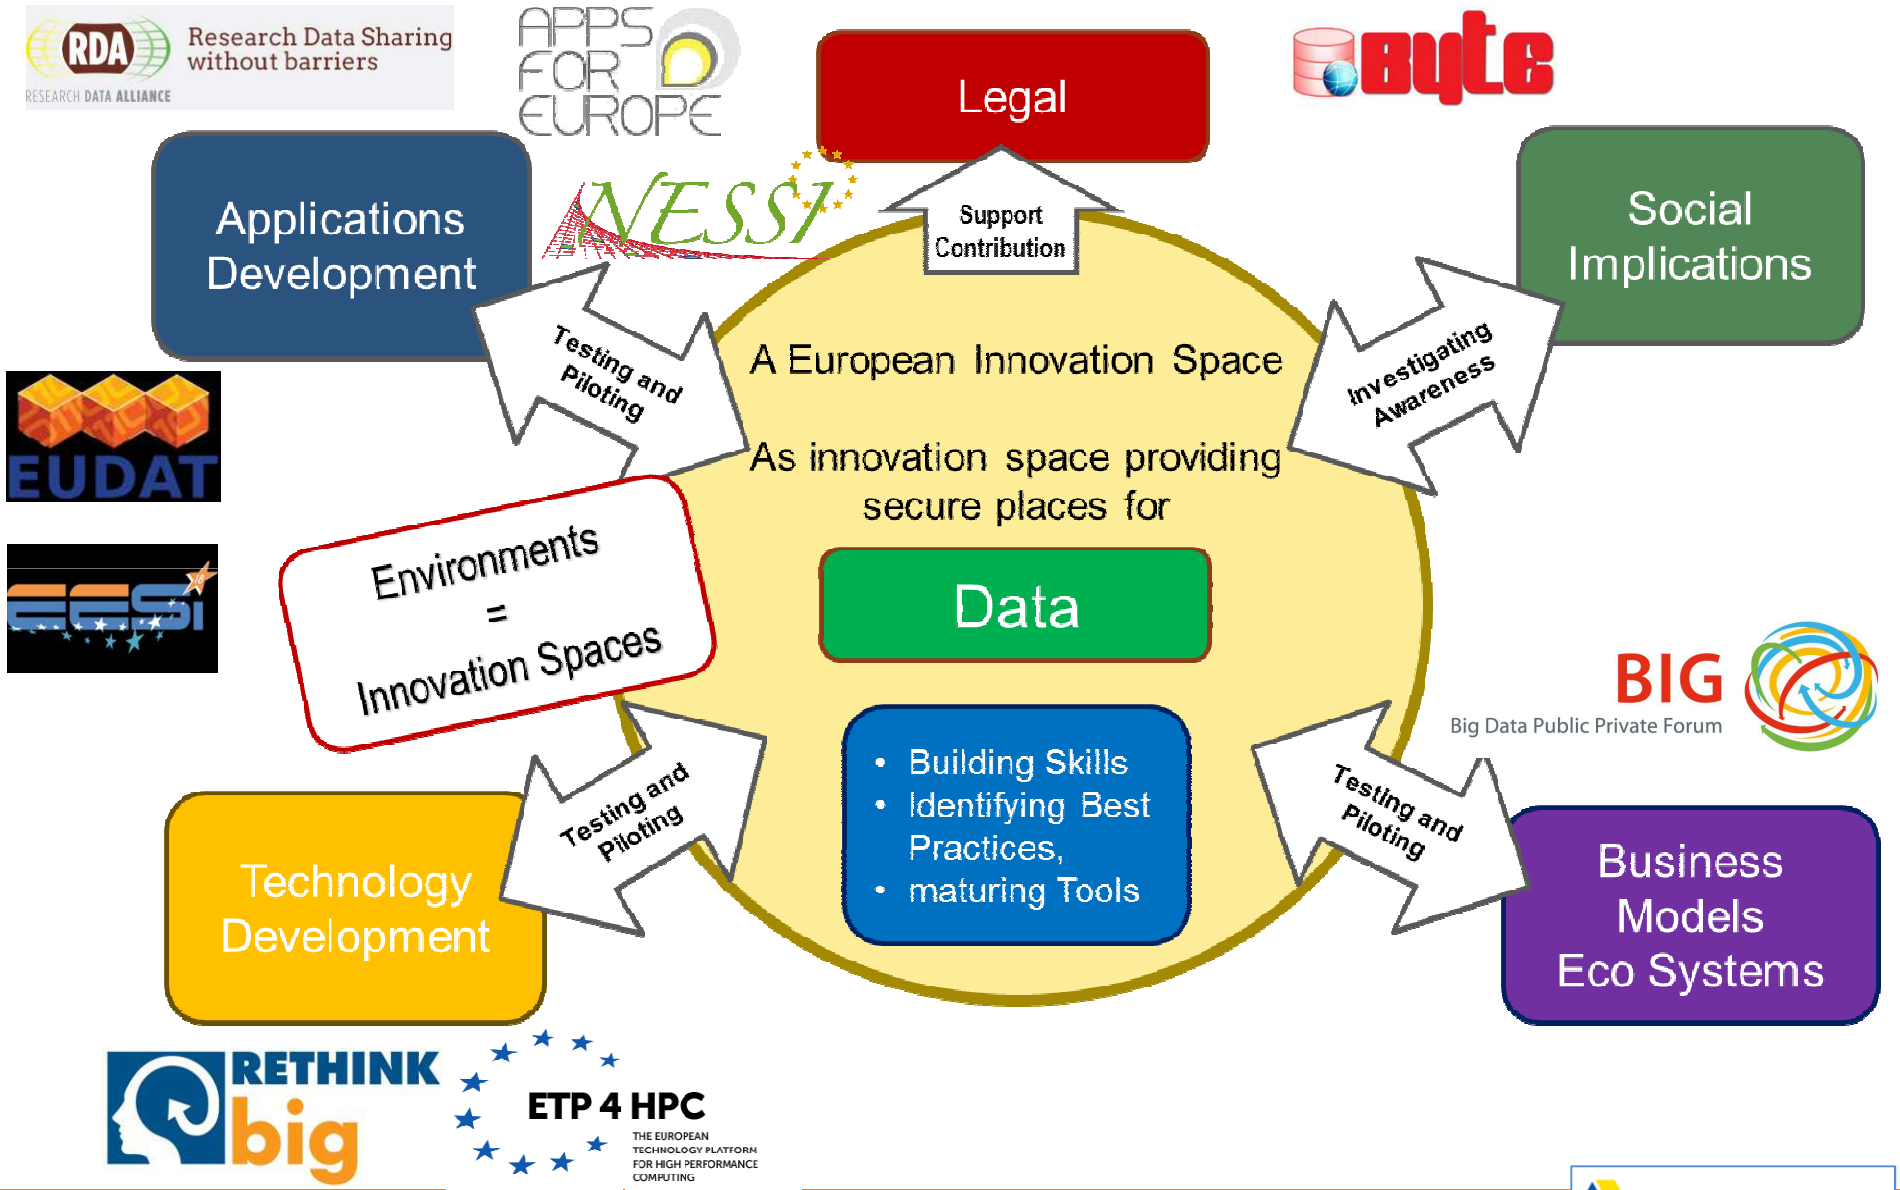

BDV BIG DATA VALUE

Thank you

www.RETHINKbig-project.eu

RETHINK big Project and BDV cPPP

- Building an ecosystem:

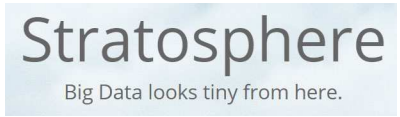

Human Brain Project