Public Clouds for science: a new sustainable business model for distributed computing infrastructures

Fabrizio Gagliardi
Microsoft Research

Microsoft Research Connections

Outline

- The Challenge of the “long tail” of science
- Is there a sustainable financial model for scientific data?
- Large Data Centers powering commercial clouds
- Clouds vs Grids
- The heritage of VENUS-C
- VENUS-C Communities

The data explosion is transforming science

Petabytes
Doubling & Doubling

- Every area of science is now engaged in data-intensive research
- Researchers need
  - Technology to publish and share data in the cloud
  - Data analytics tools to explore massive data collections
  - A sustainable economic model for scientific analysis, collaboration and data curation

Microsoft Research Connections

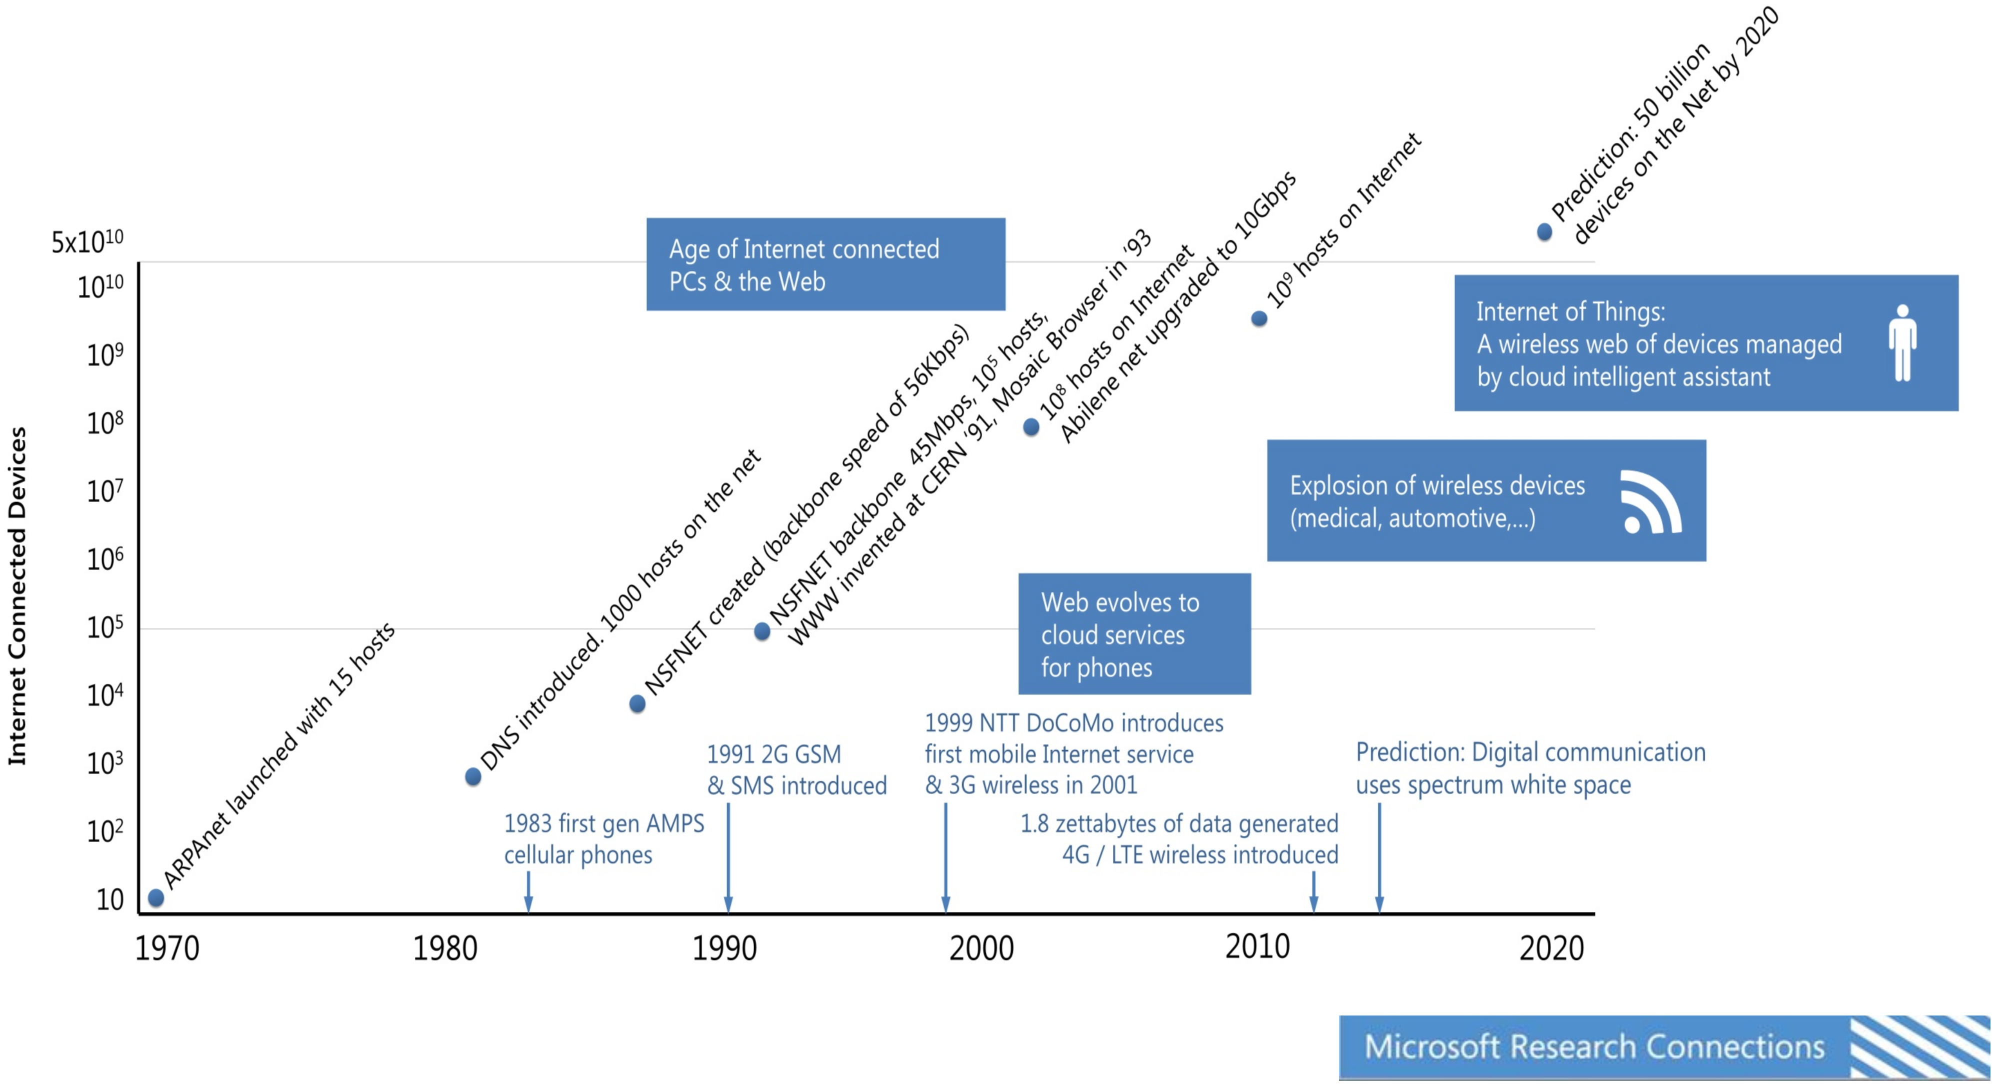

Internet Connected Devices
5x10^10
10^10
10^9
10^8
10^7
10^6
10^5
10^4
10^3
10^2
10
1970
1980
1990
2000
2010
2020
ARPAnet launched with 15 hosts
DNS introduced. 1000 hosts on the net
NSFNET created (backbone speed of 56Kbps)
NSFNET backbone 45Mbps, 10^5 hosts,
WWW invented at CERN '91, Mosaic Browser in '93
10^8 hosts on Internet
Abilene net upgraded to 10Gbps
10^9 hosts on Internet
Prediction: 50 billion
devices on the Net by 2020
Age of Internet connected
PCs & the Web
Internet of Things:
A wireless web of devices managed
by cloud intelligent assistant
Explosion of wireless devices
(medical, automotive,...)
Web evolves to
cloud services
for phones
1983 first gen AMPS
cellular phones
1991 2G GSM
& SMS introduced
1999 NTT DoCoMo introduces
first mobile Internet service
& 3G wireless in 2001
1.8 zettabytes of data generated
4G / LTE wireless introduced
Prediction: Digital communication
uses spectrum white space
Microsoft Research Connections

Where is the Big Data Now?

Big Data Collections (these are approximations, items from talks and rumors)

- Facebook
  - 140B photos @100K per = 14 PBytes
  - Growing at 70B photos/year. 30,000 servers
  - 25TB of log data per day
  - Social Graph 500M nodes.
- Ebay - 16PB (rumored)
- Human Genome
  - 1 individual = 30TB raw format from Illumina IG sequencer
  - Final form 1.5GB (or 20MB with a good reference genome)
  - Library of 1B people = 20PB
- Crawling the web (estimate)
  - 1T unique URLs @ 50K each = 50PB
- Email (estimate)
  - Hotmail: 350M users, Gmail 260M users = 600M users @1GB each = 600PB
  - (In 2006, the email traffic, excluding spam, accounted for 6 exabytes)
- **Nearly 70% of the digital universe will be generated by individuals by 2020**
  - **Cross 1 zettabyte by end 2012.**

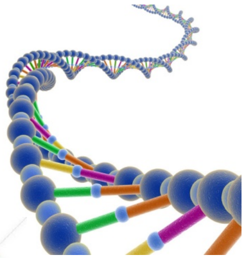

Data Streaming to/from major sources

- Twitter “fire hose”
  - 50M tweets/day*(140+64)B/tw = 10GB/day = 1Mb/sec
- Google search (estimate)
  - 2.36 Mb/sec input queries. 100 Mb/sec out.
- LHC
  - 15 PB/year = 41TB/day = 1.712 TB/hr = 3.8Gb/sec
- Email (non-spam)
  - Gmail 1B emails/day = 75TB/day = 7Gb/sec,
- SKA
  - Raw data = 960PB/day, Final processed data = 80Gb/sec
- Zynga
  - 1PB/day = 92Gb/sec
- Facebook
  - 600K photos delivered per second @25K. = 120 Gb/sec

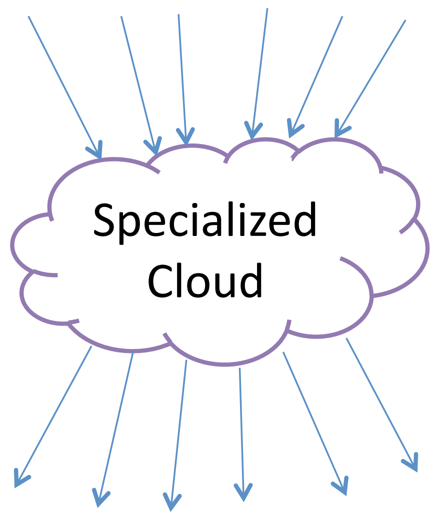

The Long Tail of Science

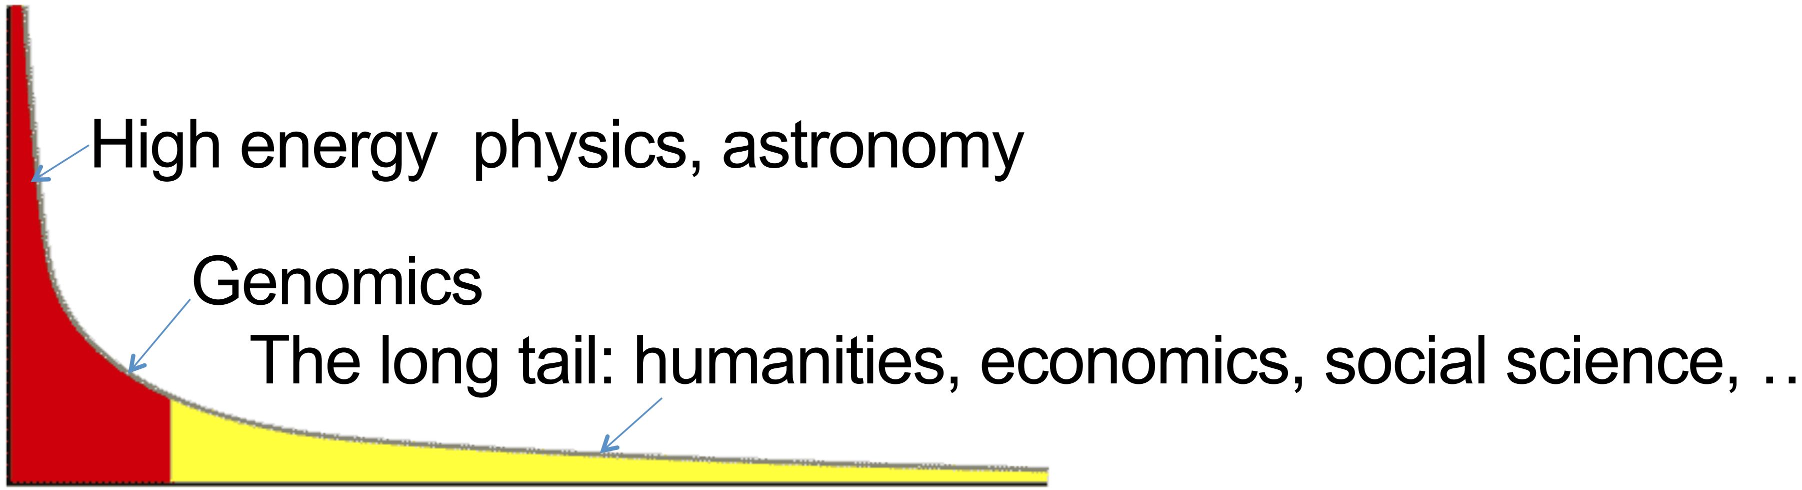

Collectively "long tail" science is generating a lot of data
- Estimated at several PBs per year and growing fast

The EC and US NSF require all data produced by the publicly funded projects to be made openly accessible:
- Universities are struggling with this new load
- Data must be preserved
- Data must be sharable, searchable, and analyzable

Microsoft Research Connections

Data Centers and the Microsoft Cloud

Microsoft Research Connections

The Microsoft Cloud is Built on Data Centers

~100 Globally Distributed Data Centers

Range in size from "edge" facilities to megascale (100K to 1M servers)

The Microsoft Cloud

- Purpose-built data centers to host containers at large scale
  - Cost $500 million, 100,000 square foot facility (10 football fields)
- 40 foot-long shipping containers can house as many as 2,500 servers each
  - Density of 10 times amount of compute in equivalent space in traditional data center
- Deliver an average PUE of 1.22
  - Power Usage Effectiveness benchmark from The Green Grid™ consortium on energy efficiency

Data Center Infrastructure

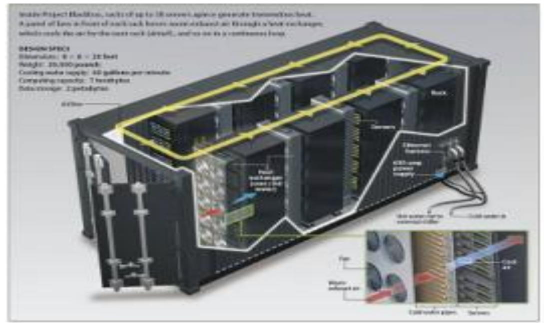

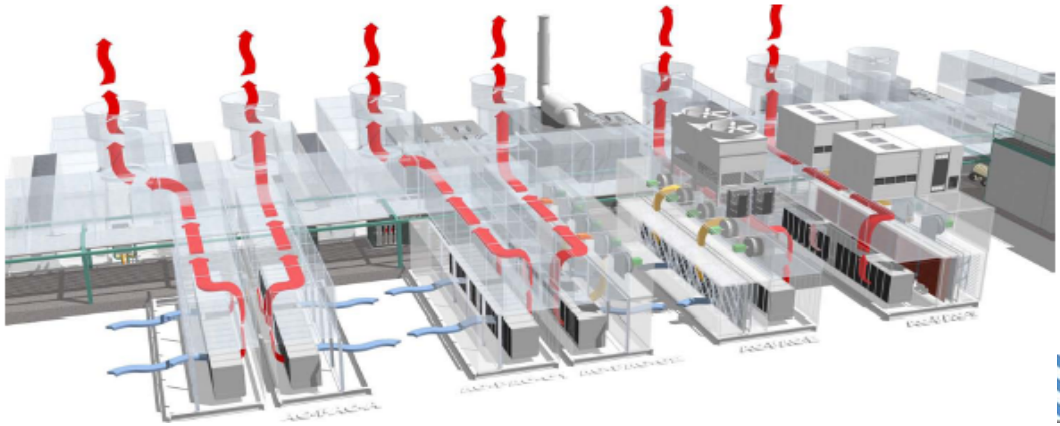

Cloud Properties

- **Designed to Provide Information and Computation to Many Users**

- Automatic Deployment and Management of Virtual Machine Instances
  - tens to thousands & dynamic scalability
- Dynamic Fault Recovery of Failed Resources
  - Cloud Services must run 24x7
- Automatic Data Replication
  - Geo-replication if needed
- Two levels of parallelism
  - Thousands of concurrent users
  - Thousands of servers for a single task

Grids vs Clouds

- Computing and data servers
  - EGEE (in 2010): **150 K** CPU cores in 300 sites in 50 countries around the world for 15,000 users (10,000 from LHC). The majority of users have now access via National Grid Initiatives (NGIs) with **EGI**
  - Cost so far for the EU: > 150 M € (excluding CAPEX)
  - MS Chicago Data Center = 50 containers = **100 K 8-core servers** – cost (estimated) > 500 M €
- Architecture
  - Grids: Heterogeneous systems in different administrative domains – limited scalability
  - Clouds: Homogeneous systems in single administrative domains – high scalability
- Cost models
  - GRIDS: contribute resources in kind to VOs
  - Clouds: pay per use with fine granularity

EMI Services Deployment

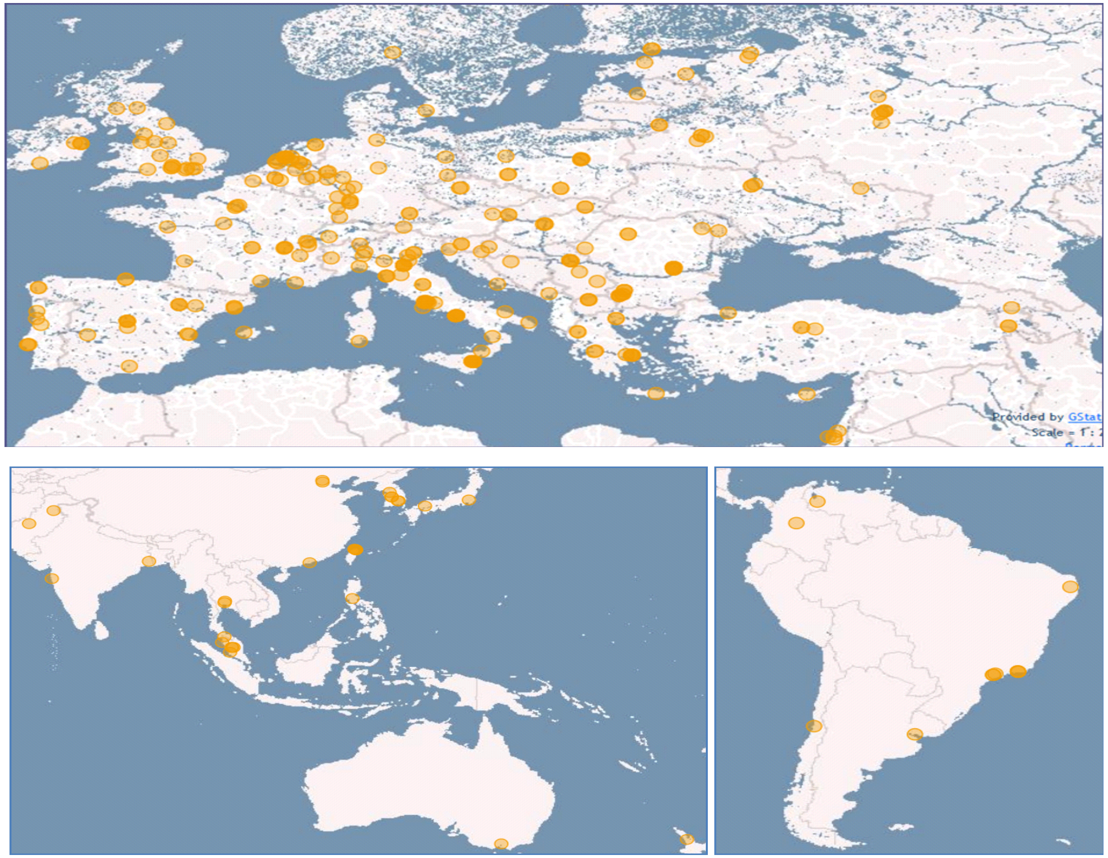

As of May 2012 the EMI services are deployed on all EGI sites

352 EGI sites
299 from 42 Euro/CERN
27 from Asia-Pacific
26 from Canada and LA

A cumulative total of **1095** service instances are deployed

Estimated base of around **20000** end users of which around **2000** are infrastructure operators

Courtesy of A. Di Meglio

The Heritage from VENUS-C

Microsoft Research Connections

Industry contribution to the European Cloud Strategy

Building an industry-quality, highly scalable & flexibale Cloud infrastructure

Microsoft
EMIC – MICGR - MRL

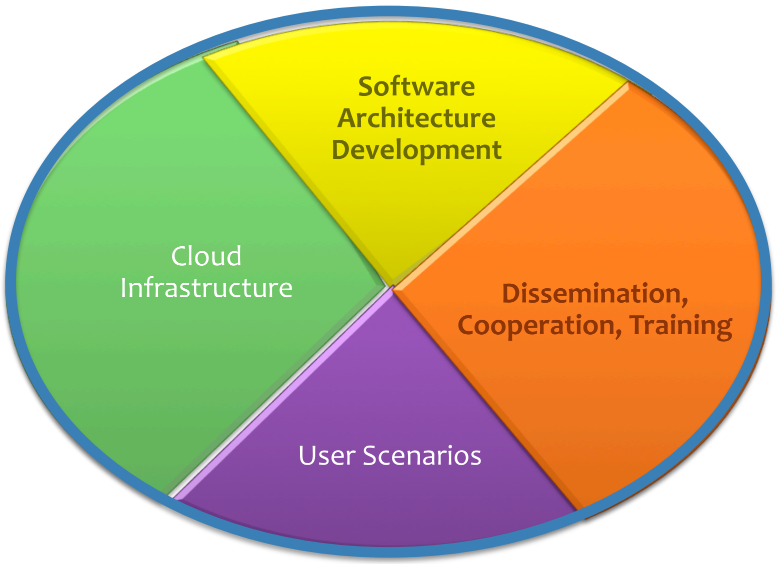

Coordinated by Engineering – Investment in infrastructure provision & software development.

Microsoft invests in Azure resources & manpower through Redmond & its European data centres

A user-centric Approach

Building a Cloud Infrastructure with user needs interwoven
Bringing about fundamental changes in scientific discovery & innovation

Impossibile visualizzare l'immagine. La memoria del computer potrebbe essere insufficiente per aprire l'immagine oppure l'immagine potrebbe essere danneggiata. Riavviare il computer e aprire di nuovo il file. Se viene visualizzata di nuovo la x rossa, potrebbe essere necessario eliminare l'immagine e inserirla di nuovo.

Structural Analysis of buildings

Impossibile visualizzare l'immagine. La memoria del computer potrebbe essere insufficiente per aprire l'immagine oppure l'immagine potrebbe essere danneggiata. Riavviare il computer e aprire di nuovo il file. Se viene visualizzata di nuovo la x rossa, potrebbe essere necessario eliminare l'immagine e inserirla di nuovo.

Building Information Management

AquaMaps – Marine Biodiversity data

Civil Protection and Emergencies

Bioinformatics

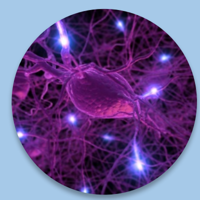

Systems Biology

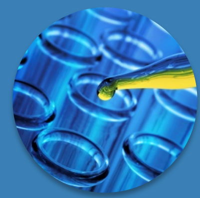

Drug Discovery

User Community in VENUS-C

- The VENUS-C project involved 7 scenarios, 15 pilots and 5 experiments, from 10 European countries
- 30% of them were not Universities

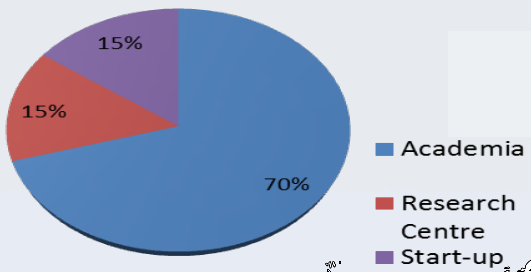

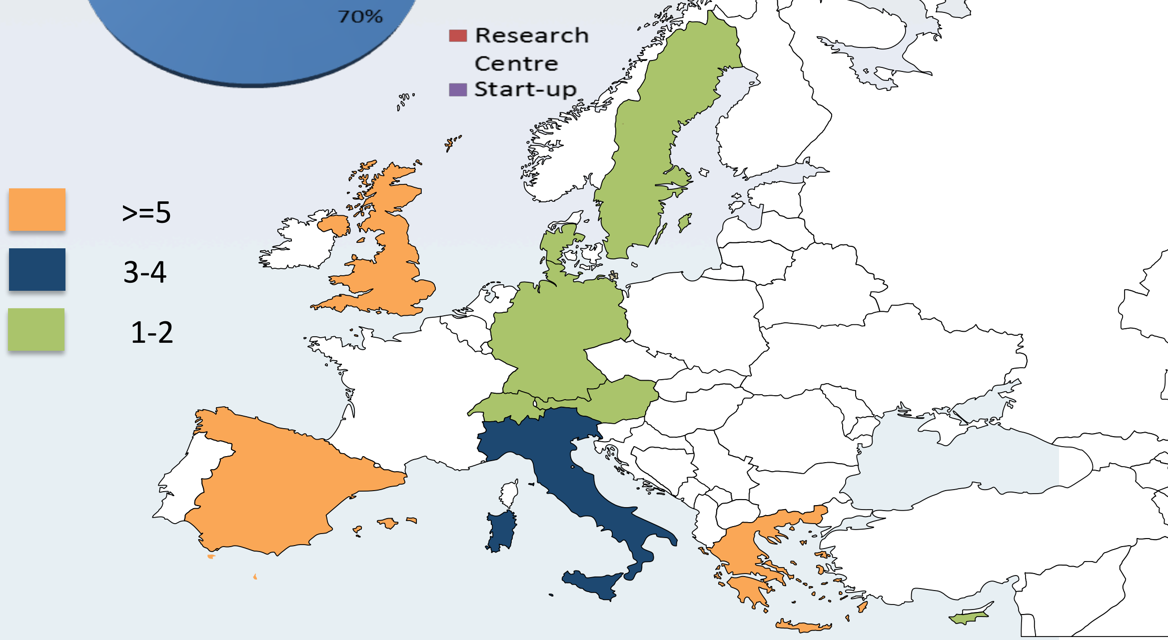

VENUS-C

Virtual multidisciplinary EnviroNments USing Cloud infrastructures

User Community

Civil Protection (1)
- Fire Risk estimation and fire propagation

Chemistry (3)
- Lead Optimization in Drug Discovery
- Molecular Docking

Biodiversity & Biology (2)
- Biodiversity maps in marine species
- Gait simulation

Civil Eng. and Arch. (4)
- Structural Analysis
- Building information Management
- Energy Efficiency in Buildings
- Soil structure simulation

Physics (1)
- Simulation of Galaxies configuration

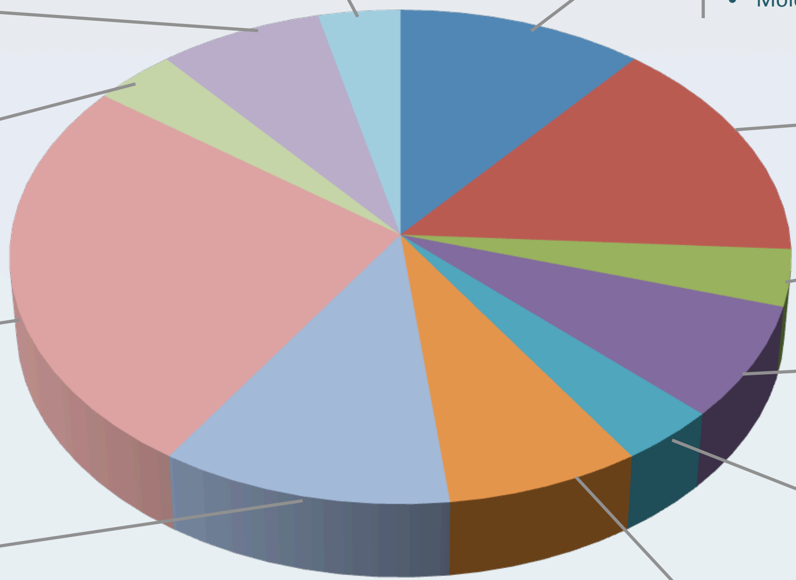

Earth Sciences (1)
- Seismic propagation

Mol, Cell. & Gen. Bio. (7)
- Genomic sequence analysis
- RNA prediction and analysis
- System Biology
- Loci Mapping
- Micro-arrays quality.

ICT (2)
- Logistics and vehicle routing
- Social networks analysis

Medicine (3) (*)
- Intensive Care Units decision support.
- IM Radiotherapy planning.
- Brain Imaging

Mathematics (1)
- Computational Algebra

Mech, Naval & Aero. Eng. (2)
- Vessels monitoring
- Bevel gear manufacturing simulation

(*) Applications from Medicine use only anonymised data released for research in the Cloud.

VENUS-C Validation - by application developers -

Addressing review comment #6 – “Verify that VENUS-C eases the adaptation”.

- Performance
- Error management
- Scalability
- Completeness
- Interoperability
- Learning curve
- Convenience
- Robustness

- Based on Speed-up comparing with local resources
- Clarity of error messages and easy availability of logs
- Performance penalty with the increased number of resources
- Full implementation of the requirements
- Flexibility to switch between different platforms and support of standards
- Easy to learn how to adapt the applications
- Qualitative evaluation of the failure ratios.

1-5 scale: 1-poor, 2- low, 3- medium, 4- good, 5- excellent.

VENUS-C
Virtual multidisciplinary EnviroNments USing Cloud infrastructures

Validation Scores

- Average evaluation was in the interval 4-5 in 81% of the cases
- Interop, Completeness and Learning Curve got the highest scores

	GW	COMPSs	CDMI	Accounting	Avg
Performance	4.50	3.75	3.67	4.40	4.08
Error management	4.00	3.20	4.00	4.00	3.80
Scalability	4.40	4.67	2.00	4.33	3.85
Completeness	4.17	4.60	4.67	4.40	4.46
Interoperability	4.00	4.60	5,00	4.60	4.55
Learning curve	4.17	4.33	4.67	4.60	4.44
Convenience	4.33	4.17	4.00	4.25	4.19
Robustness	4.33	3.67	3.33	4.40	3.93
Average	4.24	4.12	3.92	4.37	4.14

VENUS-C Success Stories

- Interactive computation of fire risk and fire propagation estimation
- Access to burst-scalable cloud compute and storage
- Web-based GIS based on Bing Map

[Wild Fire Demo]

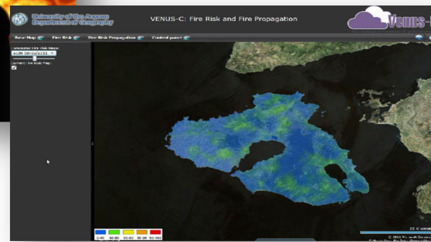

collaboratorio
innovators in architecture

- Collaboratorio & its new start-up Green Prefab
- Collaborative platform for the design of ecofriendly & affordable buildings
- Selected by INTESA SAN PAOLO Start-up intitiative; expanding to US

"We feel like pioneers in the right direction to the still untouched gold mine," Furio Barzon

Extending VENUS-C with Pilots & Experiments

Engineering & Science


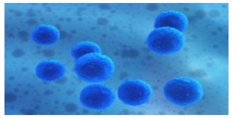

Architecture & Civil Engineering
Biology

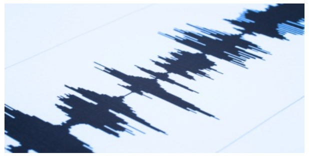
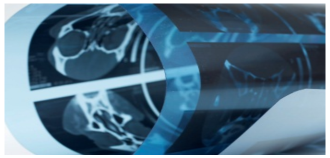
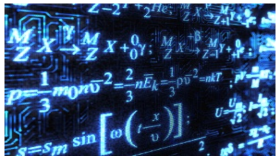



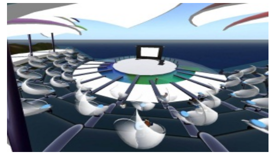

NEW DISCIPLINES

Earth Sciences, Healthcare, Maths, Mechanical Engineering, Physics, Social Media, Education

Start-ups

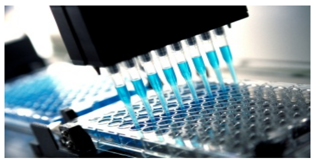

Computer resources can be scaled as required without committing to large capital purchases, which is critical to the success of our small business. **Molplex UK**

PERSEUS

DFRC is part of the EU Flagship project PERSEUS on maritime security. Scaling our platform with VENUS-C will enable us to support future growth in terms of vessels monitored in real time & usability by operators.

Conclusions and Next Steps

An opportunity to create a new model for data-intensive science

- Cloud data services from commercial providers open the door for a new paradigm for research
- A Research Data Services cloud
  - Open and extensible
  - Easily accessed by simple desktop/web analysis applications
  - Encourages scientific collaboration
  - Allows scientific analysis of massive data collections without requiring each researcher to acquire a private supercomputer
- With an ecosystem that supports a **marketplace** of research tools and domain expertise
  - Providing an economic sustainability model for data preservation and use
  - Allowing researchers to outsource special tasks to expert service providers

Microsoft Research Connections

The Data for Science Sustainability Challenge

- Can we create a sustainable *economic* model for the long tail of science?
  - The funding agencies will not **directly** support an exponentially growing data collection now will be able to continue to fund dedicated computing and data resources to each project they fund

- Our hypothesis
  - We can create an ecosystem that supports a **marketplace** of research tools and domain expertise
    - Allowing researchers to outsource special tasks to expert service providers
    - Funding will come from subscriptions from individual researchers, academic institutions and private sectors

European Cloud Computing Strategy

Vice-President Neelie Kroes, responsible for the Digital Agenda

Three Pillars for Cloud

- *Legal frameworks*
- *Technical and commercial fundamental elements*
- *Development of the cloud market by supporting pilot projects of cloud deployments*

Official opening of the Microsoft Cloud & Interoperability Center, March 2011

Cloud Power

Neelie Kroes on international standardisation & open specifications
"I count here on the further support and commitment of Microsoft and all the other participants."

More science for $$$

- Public funding agencies tend to allocate a non negligible part of their grants to provisioning of compute services:
  - Typical HPC users are well served by Surper Computer Centres, Private Grids (HEP LHC) and dedicated computing solutions
  - Everybody else (the long tail of the computational scientific community) ends up buying local clusters and storing data results in Silos
- Faster to deploy than conventional HPC in emerging scientific and business communities
- Distributing, managing and curating data is better served by a virtual, scalable and elastic Cloud infrastructure
- Economy of scale, energy costs and environmental impact are better addressed by Cloud computing
- Virtualisation of computing infrastructures can support funding agencies in developing new funding models:
  - Moving from CAPEX to OPEX
- Leading to more science per tax payer €

Credits: Dennis Gannon/MSR, Ignacio Blanquer/UPV, Alberto Di Meglio/CERN

Thanks you!