

Grid and Clouds: a look ahead

Fabrizio Gagliardi

Microsoft Research Europe
External Research, Director

Introduction

- Happy to be here once more to discuss with you latest trends in distributed computing
- Two years ago I anticipated a number of issues with long term sustainability of the Grid computing model, while confirming its validity for a selected number of scientific communities and the emergence of Cloud Computing, where KIT had already a pioneering experience
- Last year we spoke about Cloud computing coming out of age and be widely adopted by the IT industry
- This year I am happy to say that Microsoft has launched in a massive way its Cloud Computing service offering with Azure and MSR with the XCG Cloud Computing initiative is building on this initiatives for supporting science worldwide

eXtreme Computing Group (XCG)

Dan Reed, CVP, MSR, Dennis Gannon, primary interface to Venus-C and other research activity in Europe

XCG Launched in June 2009 as part of MSR (120+ FTEs and growing)

- Incubation, Manycore, Quantum Computing, **Cloud Computing Futures**, Hardware,...

Disruptive ecosystem changes require new hardware/software

- Manycore and parallel computing
- **Cloud computing infrastructure and services**
- Quantum computing models and systems

Intentional R&D

- Working across hardware, software and application boundaries
- Addressing critical Microsoft challenges
- **Engaging industry, government and academia**
- **Building realistic prototypes at scale, both hardware and software**

The Cloud

- A model of computation and data storage based on “pay as you go” access to “unlimited” remote data center capabilities
- A cloud infrastructure provides a framework to manage scalable, reliable, on-demand access to applications
- A cloud is the “invisible” backend to many of our mobile applications
- Historical roots in today’s Internet apps
 - Search, email, social networks
 - File storage



The Cloud is built on massive data centers

Range in size from “edge” facilities to megascale.

Economies of scale

- Approximate costs for a small size center (1K servers) and a larger, 100K server center.



Technology	Cost in small-sized Data Center	Cost in Large Data Center	Ratio
Network	\$95 per Mbps/ Month	\$13 per Mbps/ month	7.1
Storage	\$2.20 per GB/ Month	\$0.40 per GB/ month	5.7
Administration	~140 servers/ Administrator	>1000 Servers/ Administrator	7.1



Each data center is
11.5 times
the size of a football field

Microsoft Advances in DC Deployment

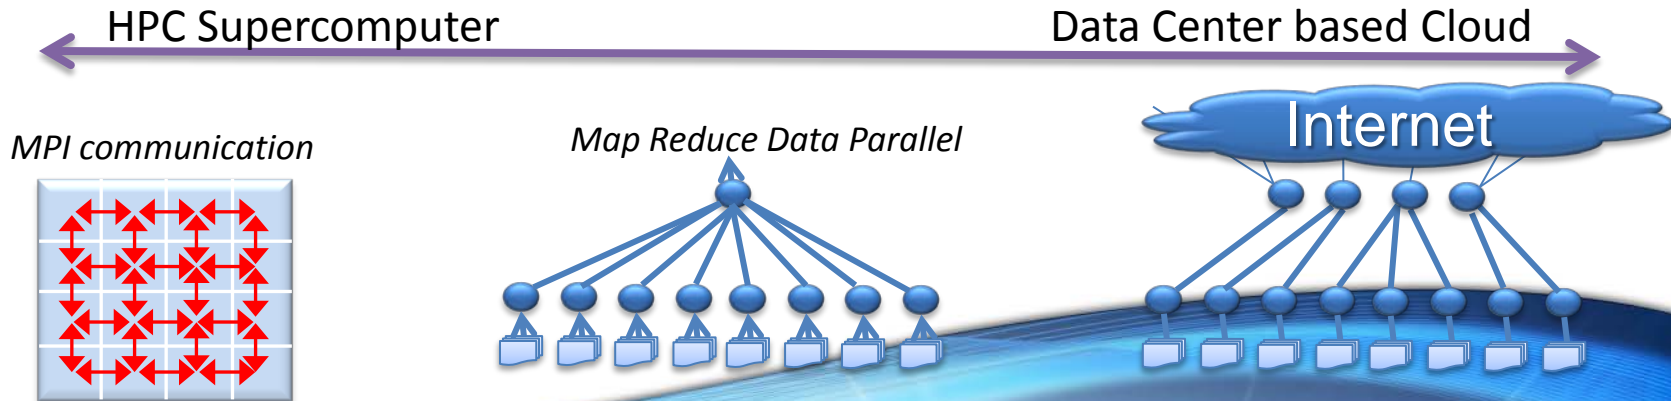
Conquering complexity.

- Building racks of servers & complex cooling systems all separately is not efficient
- Package and deploy into bigger units



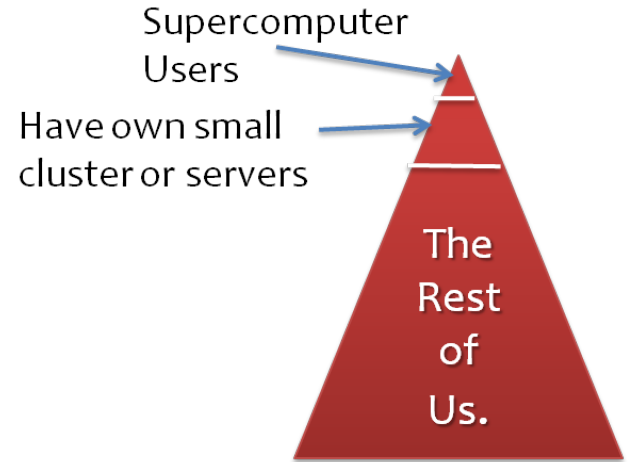
DCs vs Grids, Clusters and Supercomputer Apps

- Supercomputers
 - High parallel, tightly synchronized MPI simulations
- Clusters
 - Gross grain parallelism, single administrative domains
- Grids
 - Job parallelism, throughput computing, heterogeneous administrative domains
- Cloud
 - Scalable, parallel, resilient web services



Our vision: changing the way we do research

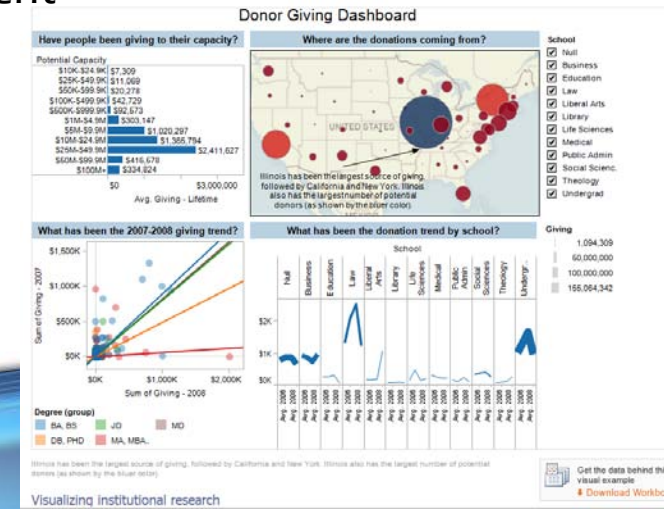
- The Rest of Us
 - Use laptops
 - Got data, now what?
 - And it is really is about data, not the FLOPS...
 - » Our data collections are not as big as we wished
 - » When data collection does grow large, not able to analyze
 - Tools are limited, must dedicate resources to build analysis tools
- Paradigm shifts for research
 - The ability to marshal needed resources on demand
 - » *Without caring or knowing how it gets done...*
 - Funding agencies can request grantees to archive research data
 - The cloud can support very large numbers of users or communities



Focus Client + Cloud for Research

Seamless interaction

- Cloud is the lens that magnifies the power of desktop
- Persist and share data from client in the cloud
- Analyze data initially captured in client tools, such as Excel
 - Analysis as a service (think SQL, Map-Reduce, R/MatLab)
 - Data visualization generated in the cloud, display on client
 - Provenance, collaboration, other 'core' services...



The Clients+Cloud Platform

- At one time the “client” was a PC + browser. Now
 - The Phone
 - The laptop/tablet
 - The TV/Surface/Media wall

And the future

- The instrumented room
- Aware and active surfaces
- Voice and gesture recognition
- Knowledge of where we are
- Knowledge of our health




The Future: an Explosion of Data

Experiments



Simulations




Archives



Literature



Instruments



The Challenge:

Enable Discovery.

Deliver the capability to mine, search and analyze this data in near real time.

**Petabytes
Doubling every
2 years**

Enhance our Lives

Participate in our own health care. Augment experience with deeper understanding.

Changing Nature of Discovery

Complex models

- Multidisciplinary interactions
- Wide temporal and spatial scales

Large multidisciplinary data


- Real-time streams
- Structured and unstructured

Distributed communities

- Virtual organizations
- Socialization and management

Diverse expectations

- Client-centric and infrastructure-centric



The Cloud Landscape

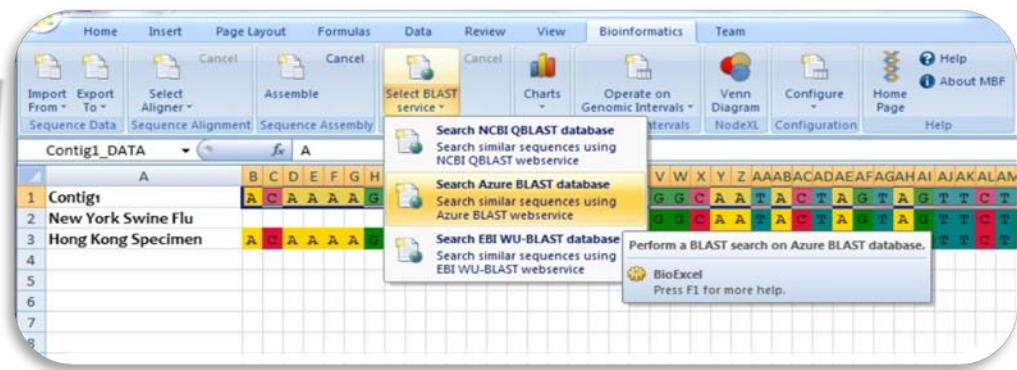
- Infrastructure as a Service (IaaS)
 - Provide a data center and a way to host client VMs and data.
- Platform as a Service (PaaS)
 - Provide a programming environment to build a cloud application
 - The cloud deploys and manages the app for the client
- Software as a Service (SaaS)
 - Delivery of software from the cloud to the desktop



Some examples based on our ongoing
Cloud Research Engagements

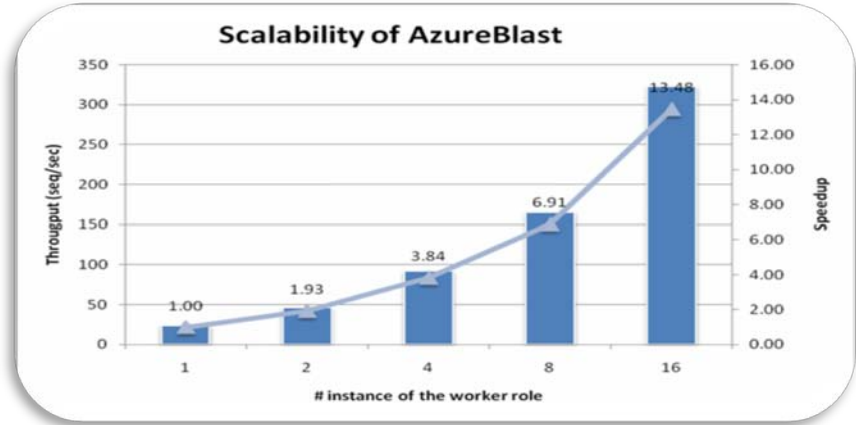
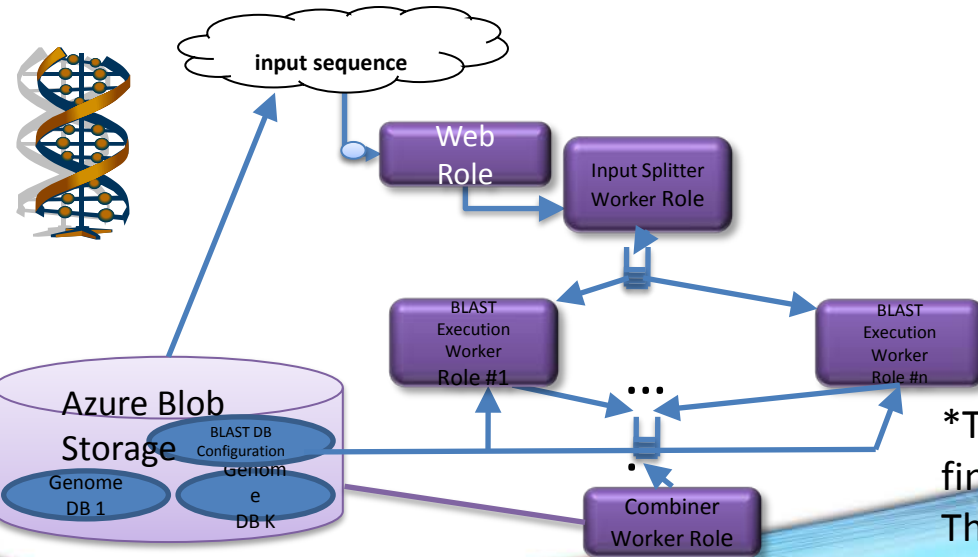


AzureBLAST*



Seamless Experience

- Evaluate data and invoke computational models from Excel.
- Computationally heavy analysis done close to large database of curated data.
- Scalable for large, surge computationally heavy analysis.
- Test local, run on the cloud.



*The **Basic Local Alignment Search Tool (BLAST)** finds regions of local similarity between sequences. The program compares nucleotide or protein sequences to sequence databases

Making Excel the user interface to the cloud

The screenshot shows the Microsoft Excel interface with a custom 'Bioinformatics' ribbon. The ribbon includes the following groups and options:

- Home**: Import From, Export To, Select Aligner, Assemble
- Cancel**: Cancel (multiple instances)
- Select BLAST service**: A dropdown menu with three options:
 - Search NCBI QBLAST database**: Search similar sequences using NCBI QBLAST webservice
 - Search Azure BLAST database**: Search similar sequences using Azure BLAST webservice
 - Search EBI WU-BLAST database**: Search similar sequences using EBI WU-BLAST webservice
- Charts**: Charts
- Operate on Genomic Intervals**: Operate on Genomic Intervals
- Venn Diagram**: Venn Diagram, NodeXL
- Configure**: Configuration
- Help**: Help, About MBF, Home Page

The worksheet contains a table with the following data:


	A	B	C	D	E	F	G	H	V	W	X	Y	Z	AA	AB	AC	AD	AE	AF	AG	AH	AI	AJ	AK	AL	AM
1	Contig1	A	C	A	A	A	A	G	G	G	C	A	A	T	A	C	T	A	G	T	A	G	T	T	C	T
2	New York Swine Flu								G	G	C	A	A	T	A	C	T	A	G	T	A	G	T	T	C	T
3	Hong Kong Specimen	A	C	A	A	A	A	G																		
4																										
5																										
6																										
7																										
8																										

A BioExcel tooltip is visible, stating: "BioExcel Press F1 for more help."

Supporting Smart Sensors and Data Fusion

The NSF Ocean Observing Initiative


- Hundreds of cabled sensors and robots exploring the sea floor
- Data to be collected, curated, mined
- OOI Architecture plan of record, store this data in the cloud




Data collected from:

- Ocean floor sensors, AUV tracks, ship-side cruises, computational models

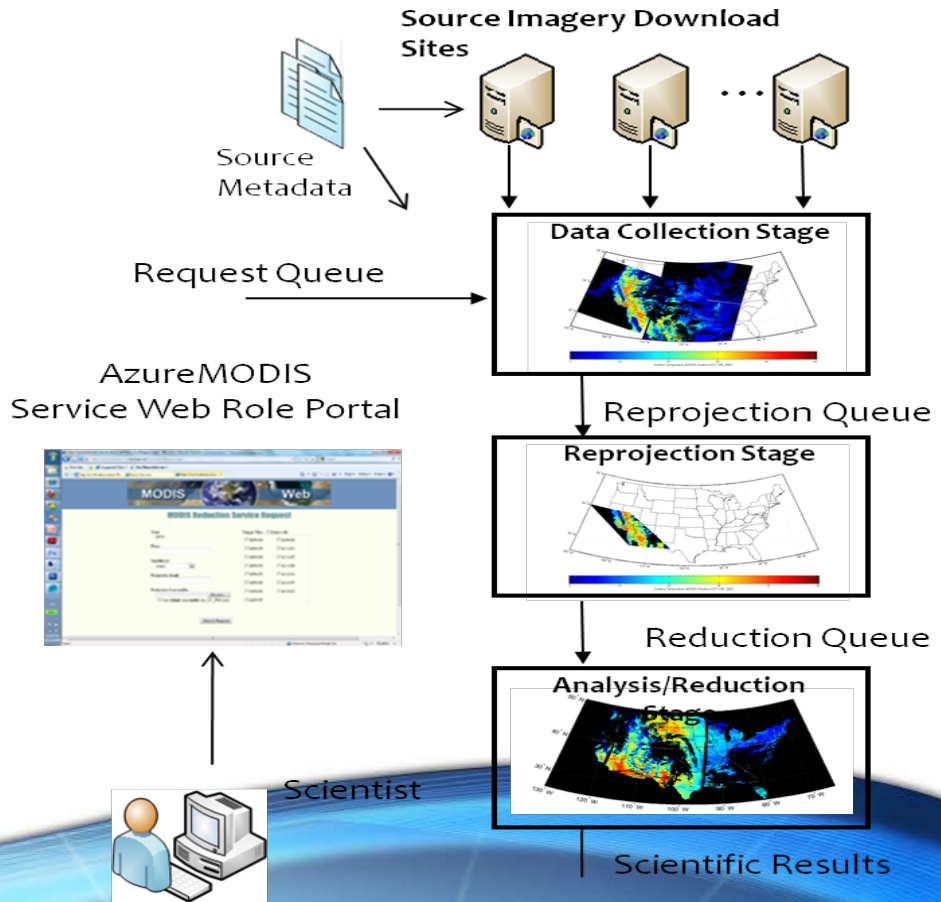
Data moves from **ocean** to shore side **data center** to the **Azure cloud** to your **computer**.



AzureMODIS – Azure Service for Remote Sensing Geoscience



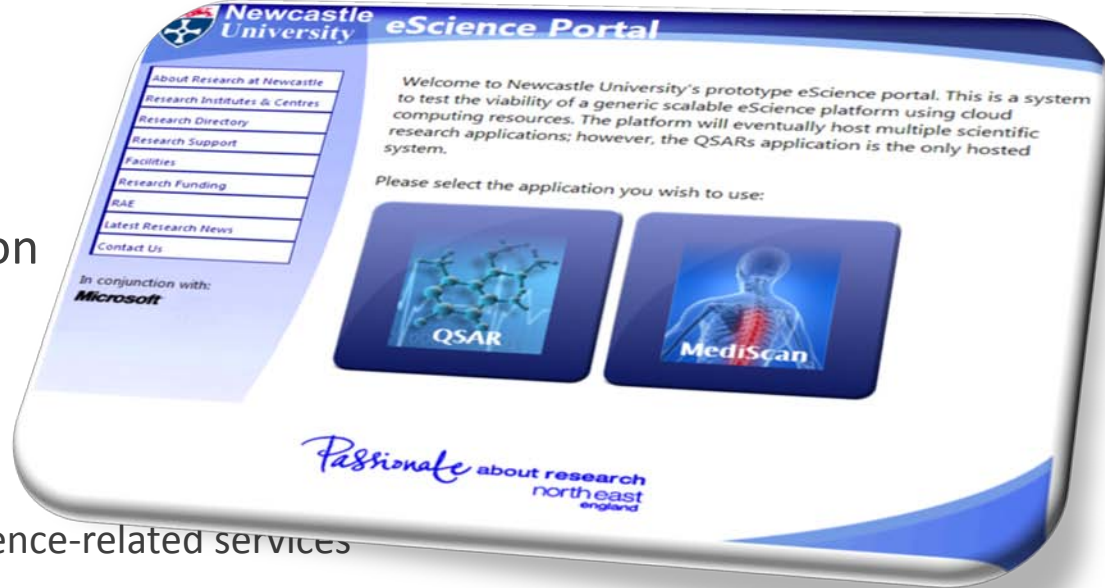
- 5 TB (~600K files) upload of 9 different imagery products from 15 different locations (~6 days of download)
- 4 TB reprojected harmonized imagery ~35000 cpu hours
- 50 GB reduced science variable results ~18000 cpu hours (~14 hour download)
- 50 GB additional reduced science analysis results ~18000 cpu hours (~14 hour download)



Project JUNIOR

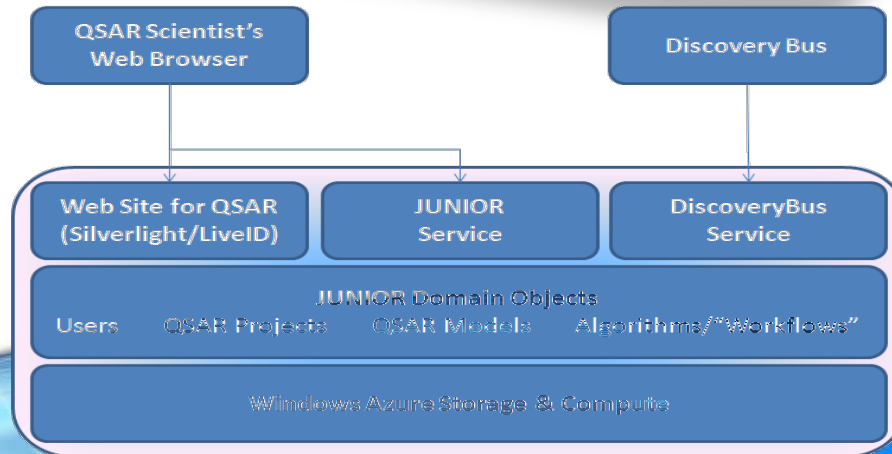
Newcastle University, UK -Paul Watson

- Investigating applicability of commercial clouds for scientific research
- Build a working prototype for use-cases in chemo-informatics
- Uses Microsoft technologies to build science-related services (Windows Azure, Silverlight...)
- Exploits Azure and Amazon Clouds



Built initial proof-of-concept

- Silverlight UI for basic Quantitative Structure-Analysis Relationship (QSAR) modeling
- Demonstrated ability to scale QSAR computations in Windows Azure

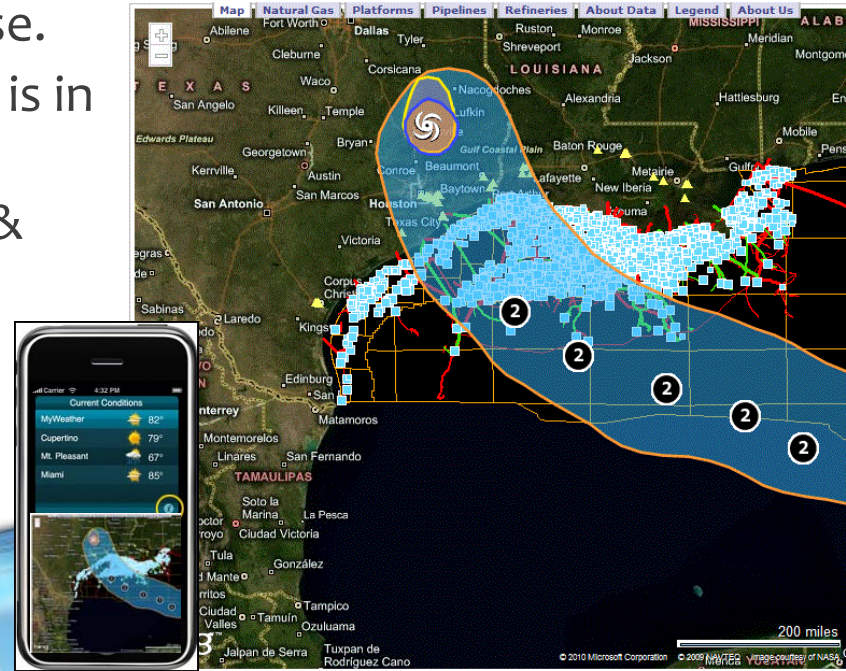


Potential Cloud Applications

- The cloud as a host to data collections
 - *Classical http data* - Web n-grams, photos, blogs, etc.
 - *Geographic & Semi-static* – Maps, satellite, geological records, weather radar, field reports/observations, geo-sensors (seismic, ocean)
 - *Geo-tagged Social record* - census, public government records, local newspapers, health and police records
 - *Scientific* – genome, pathology data, fMRI, astronomy, chemical
 - *Dynamic shared state* - multiplayer games, virtual worlds, social networks
 - *My context* - my devices, my application sessions, my agents.
- 

Geographic & semi-static data

- Maps, satellite, geological records, weather radar, field reports/observations, geo-sensors.
 - Mobile apps for emergency response.
 - Tell me when me or my team/family is in danger.
 - Integrate emergency data streams & field report tweets in the cloud and push to my client devices.



Geo-tagged social record

- Census, public government records, local newspapers, health and police records
 - Build cloud tools to understand societies problems
 - » Links between history of disease, poverty and pollution
 - » Correlate crime rates and bad weather
 - Discover the lost heritage of a people or a place



Dynamic shared state

- multiplayer games, virtual worlds, social networks
 - New modalities of collaboration



Reaching Out: Azure Research Engagement project

In the U.S.

Memorandum of Understanding with the National Science Foundation

- Provide a substantial Azure resource as a donation to NSF
- NSF will provide funding to researchers to use this resource

In Europe

- Direct engagement with the thought leaders in the U.K., France and Germany
- EC engagements where possible (VENUS-C to start off)

In both we provide our engagement team

- We provide workshops, tutorials, best practices and shared services, learn from this community, shape policy...

In Asia

- Exploring possibilities.
- 

VENUS-C



***Virtual multidisciplinary EnvironMents
USing Cloud infrastructures***


***Funding Scheme: Combination of Collaborative Project
and Coordination and Support Action: Integrated
Infrastructure Initiative (I3)***

***Program Topic: INFRA-2010-2 1.2.1. Distributed
Computing Infrastructures***

Consortium

1 (co)	Engineering Ingegneria Informatica S.p.a.	ENG	IT
2	European Microsoft Innovation Centre	EMIC	DE
3	European Charter of Open Grid Forum	OGF.eeig	UK
4	Barcelona Supercomputing Center – Centro Nacional de Supercomputación	BSC-CNS	ES
5	Universidad Politecnica de Valencia	UPV	ES
6	Kungliga Tekniska Hoegskolan	KTH	SE
7	University of the Aegean	AEG	GR
8	Technion	TECH	IL
9	Centre for Computational and Systems Biology	CoSBi	IT
10	University of Newcastle	NCL	UK
11	Consiglio Nazionale delle Ricerche	CNR	IT
12	Collaboratorio	COLB	IT

Goals

1. Create a platform that enables user applications to leverage cloud computing principles and benefits.
 2. Leverage the state of the art to on-board early adopters quickly, incrementally enable interop with existing DCI and push the state of the art where needed to satisfy on-boarding and interop
 3. Create a sustainable infrastructure that enables the cloud computing paradigms for the user communities inside the project, the one from the call for applications, as well as others.
- 

Supporting multiple basic research disciplines

Biomedicine: Integrating widely used tools for Bioinformatics (UPV), System Biology (CosBI) and Drug Discovery (NCL) into the VENUS-C infrastructure

Civil Protection and Emergency: Early fire risk detection (AEG), through an application that will run models on the VENUS-C infrastructure, based on multiple data sources.


Civil Engineering: Support complex computing tasks on Building Information Management for green constructions (provided by COLB) and dynamic building structure analysis (provided by UPV).

D4Science: Integrating computing through VENUS-C on data repositories (CNR). In particular focus will be on Marine Biodiversity through Aquamaps.



Open Call for 20 e-Science Applications

20K€ funding each (in addition to Azure Compute ,
Storage and Network Resources)

- > porting applications to the cloud
 - > education and training
 - > scalability tests
- 

Thanks for your attention

