HPDC 2009
Garching, June 9 - 13, 2009

Distributed European Infrastructure for Supercomputing Applications

Sustainable HPC Infrastructures

HPC Centers, Grids, and Clouds
- The DEISA Experience –

- Your Research and the Big Picture -

Wolfgang Gentzsch
DEISA & OGF
gentzsch at rzg.mpg.de

Distributed European Infrastructure for Supercomputing Applications

Contents

- Service Infrastructures
- Components: HPC Centers, Grids, and Clouds
- Examples: Research e-Infrastructures
- Example, Research: DEISA Ecosystem for HPC
- Example, Industry: Telecoms and the Cloud

Service Infrastructures

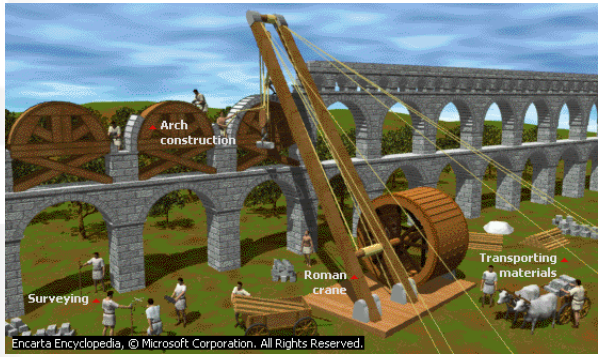

Ancient Rome: 10 aqueducts, 150,000 m^3 water each day

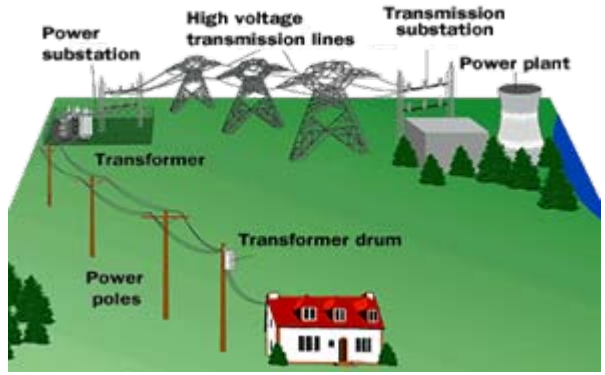

Electrical Power Grid Infrastructure

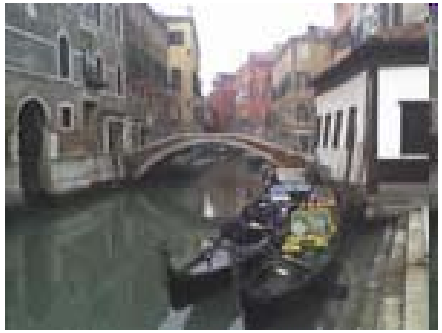

Transportation Grids

EGEE – Enabling Grid in E-SciencE

HPC Centers

- HPC Centers are **service providers**, for past 35 years
- Computing, storage, applications, data, etc IT services
- Serve (local) research, education, and industry
- Very professional: to end-users, they look (almost) like Cloud services (Amazon Cloud definition: easy, secure, flexible, on demand, pay per use, self serve)

RI-222919

RoadRunner, Breaking the Petaflop/s Barrier

Distributed European Infrastructure for Supercomputing Applications

- 1986, Cray-2 breaking the **Giga**flop/s barrier
- 1997, Intel ASCI Red, breaking **Tera**flop/s barrier
- **2008**, IBM RoadRunner, breaking **Peta**flop/s
  - At DOE's Los Alamos National Laboratory
  - 1.026 Linpack Petaflop/s solving 2 Mio equations
  - 6912 dual-core Opteron & 12960 IBM Cell eDP
  - #1 on Top500 of June 2008

RI-222919

Grids

1998: The Grid: Blueprint for a New Computing Infrastructure:

"... hardware and software infrastructure ... **dependable, consistent, pervasive, inexpensive access** to high-end computational capabilities."

2002: The Anatomy of the Grid:

"... coordinated resource **sharing** and problem solving in dynamic multi-institutional **virtual organizations**."

Quotes: Ian Foster, Carl Kesselman, Steve Tuecke

RI-222919

Grids (in 2001)

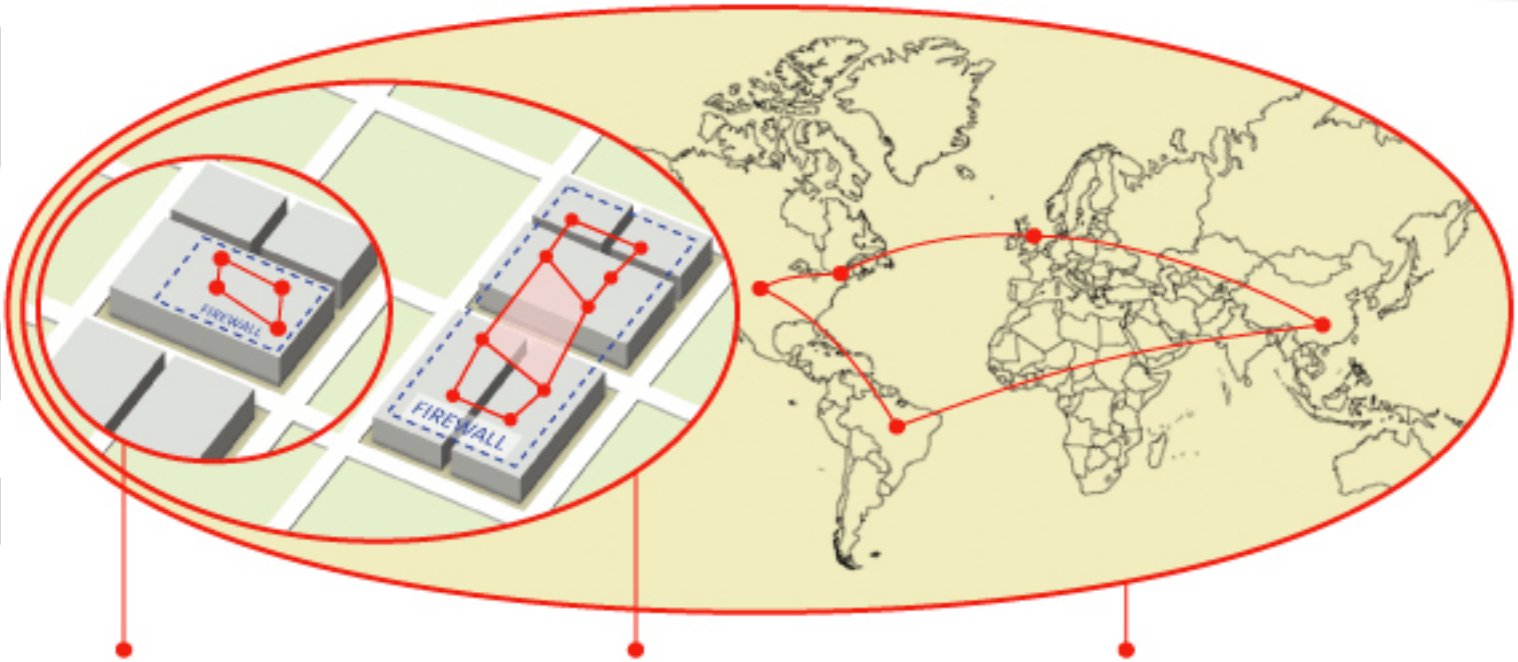

Departmental Grids | **Enterprise Grids** | **Global Grids**

RI-222919

Research Grids work:

Industry Grids work:

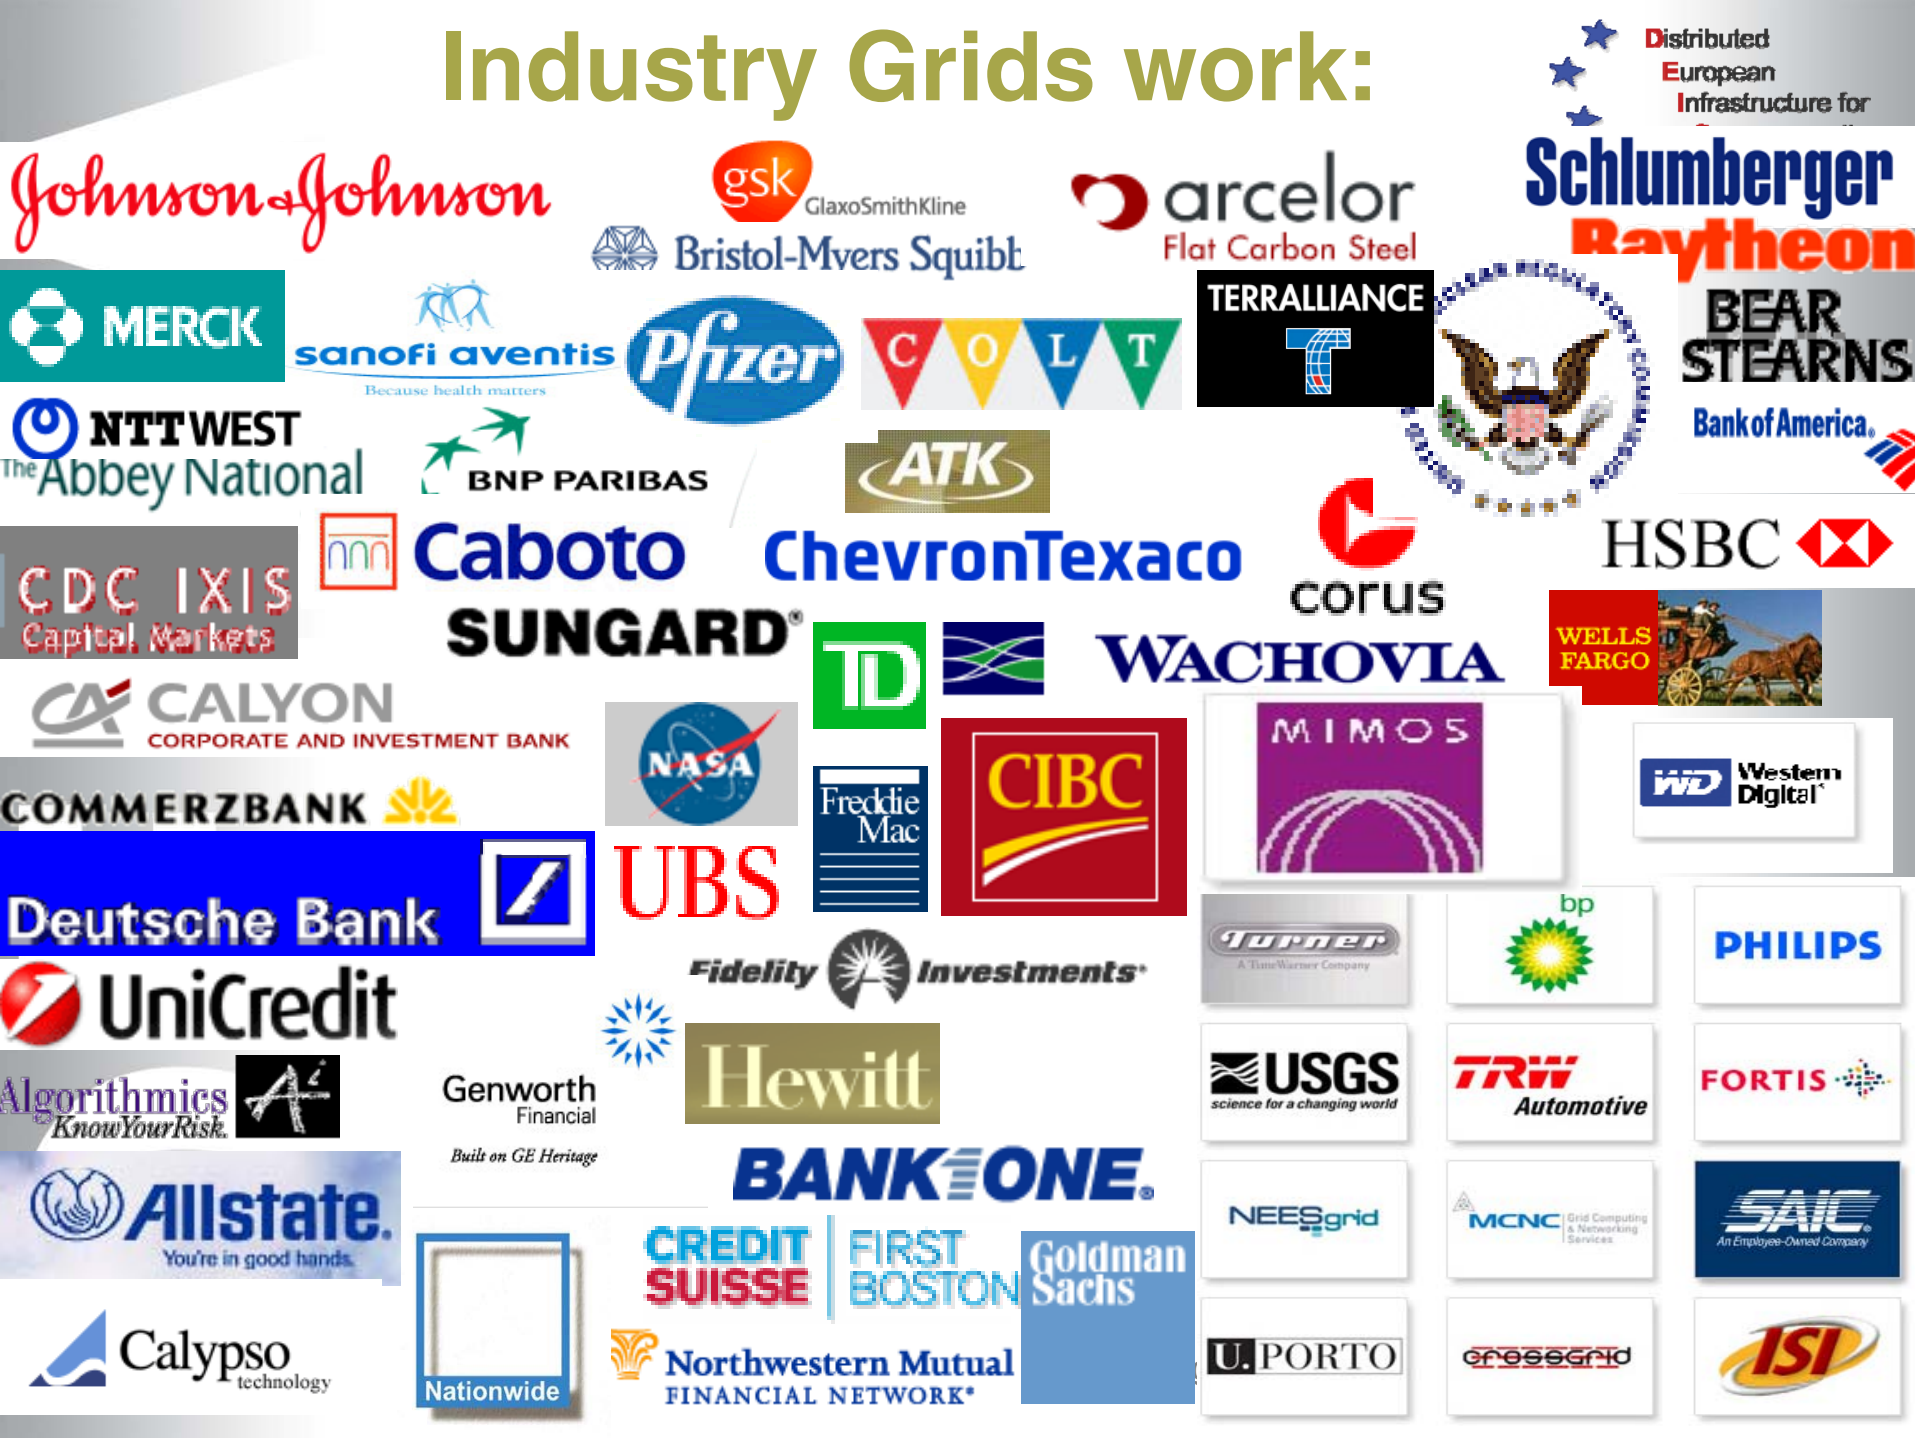

Cloud... X as a Service

Cloud: dynamically scalable and virtualized resources provided as a service over the Internet

Infrastructure (IaaS)

Platform (PaaS)

Software (SaaS)

- Accessible online, anytime, anywhere
- Pay for what you use
- Available on demand
- Service Level Agreements
- Automated:
- Scalability
- Failover
- Concurency management

The Cloud of Cloud Companies

Distributed European Infrastructure for Supercomputing Applications

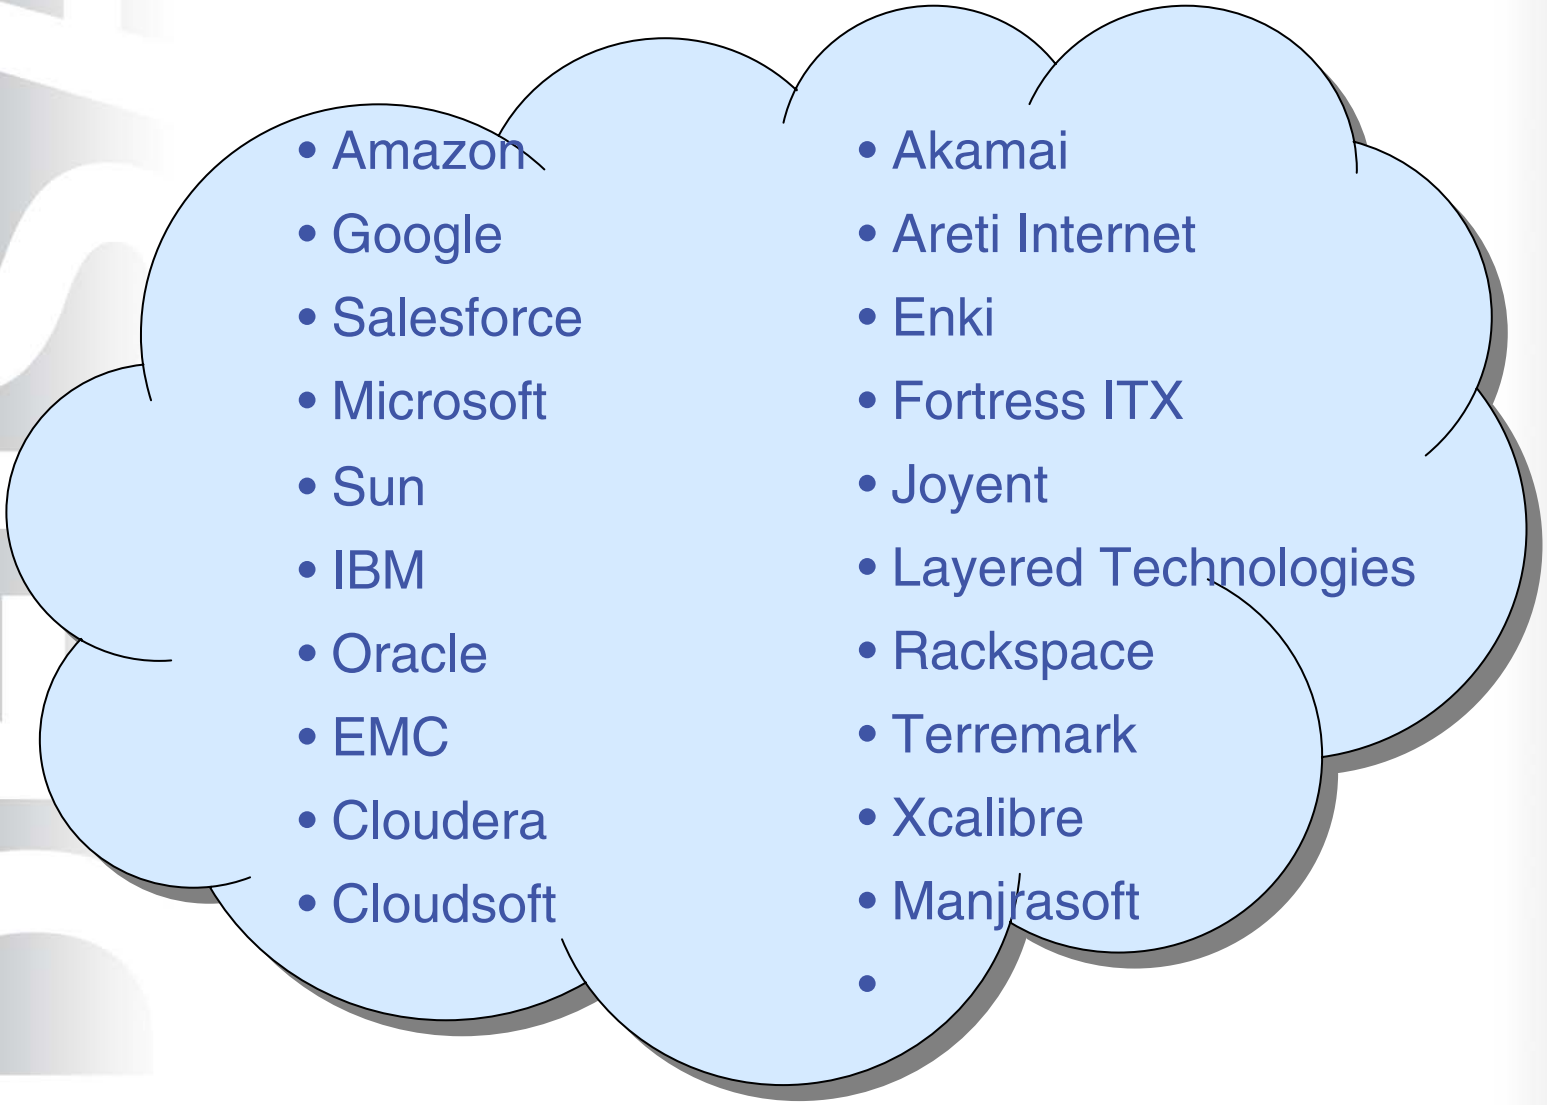

Example: ANEKA Cloud platform

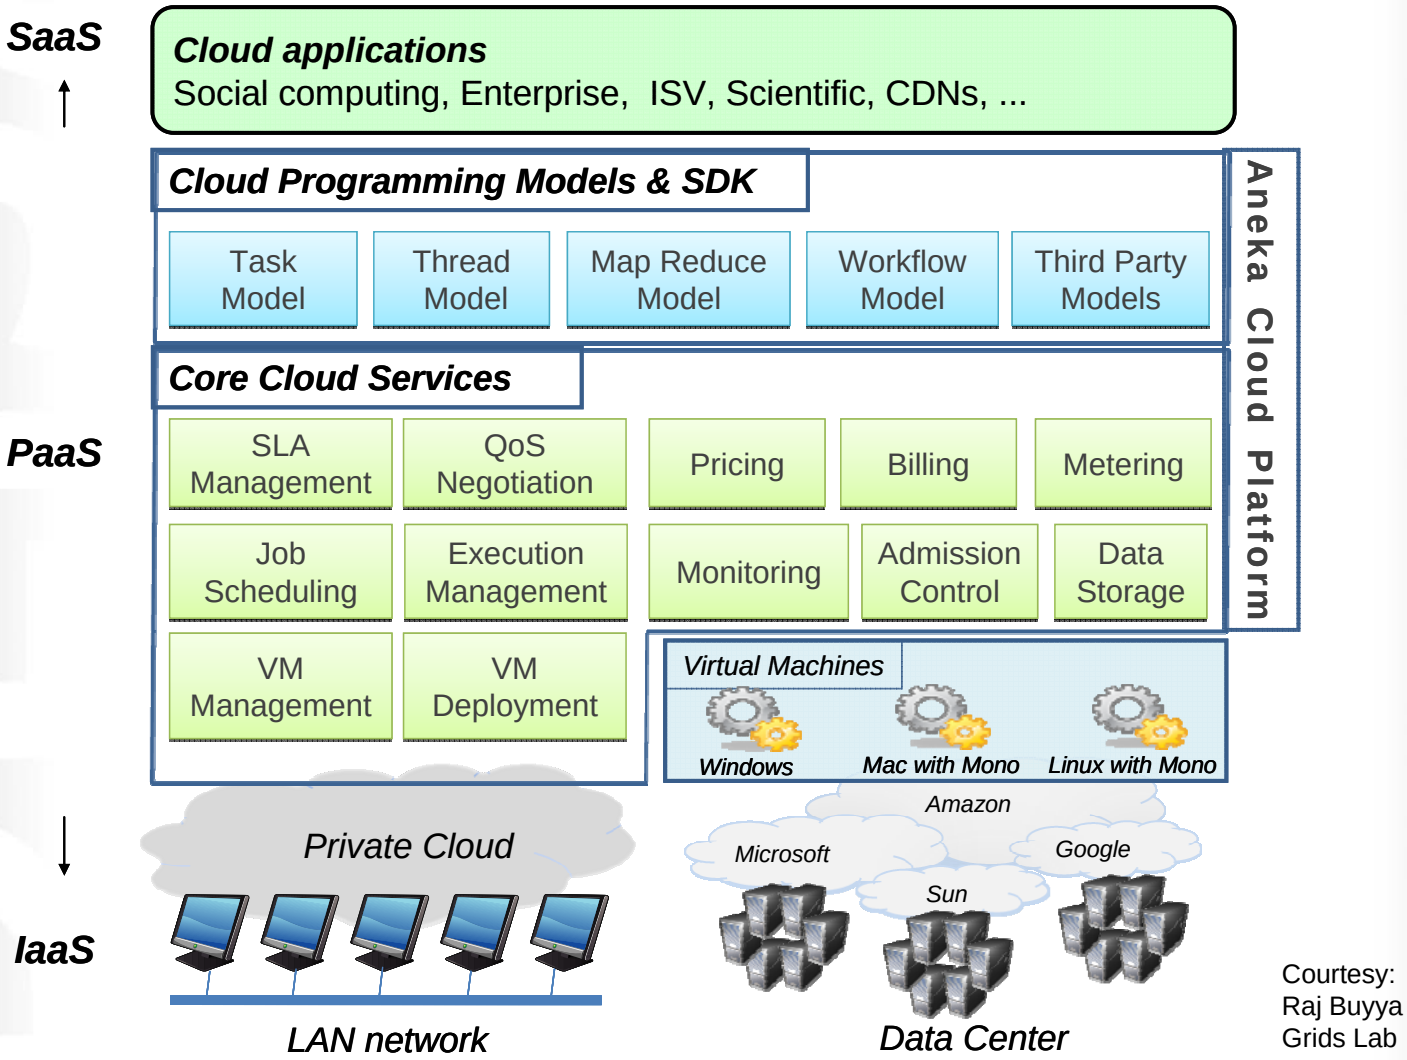

Courtesy:
Raj Buyya
Grids Lab

RI-222919

The Challenges for e-Infrastructures

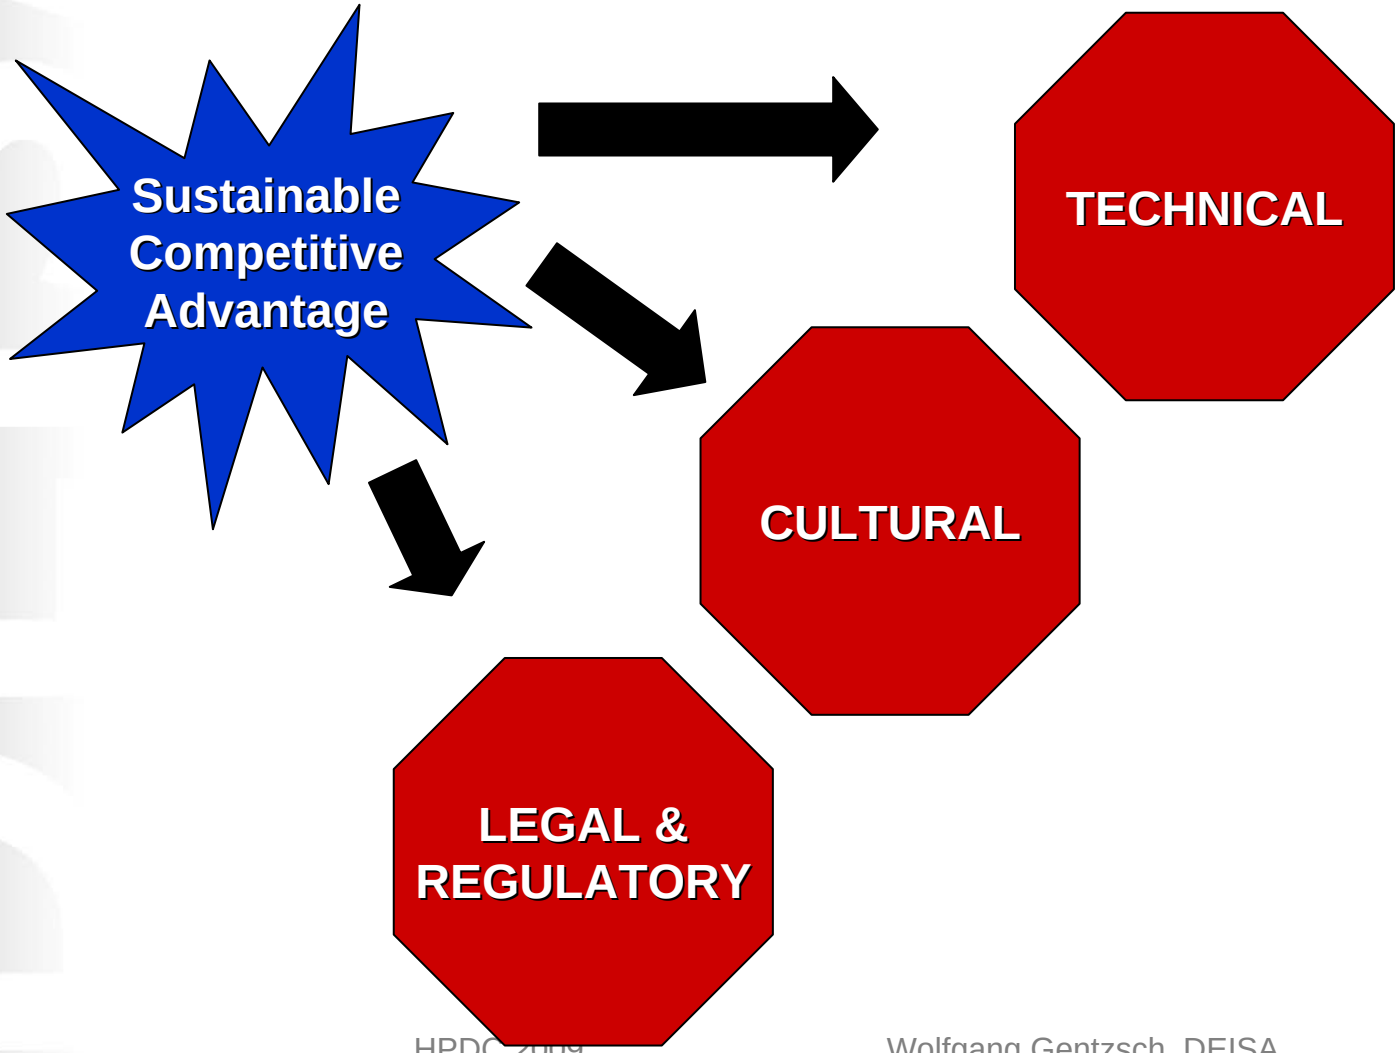

e-Infrastructure Challenges

- **Sensitive data**, sensitive applications (med.patient records)
- Different organizations have different **ROI**
- **Accounting**, who pays for what (sharing!)
- **Security** policies: consistent and enforced across the grid !
- **Interoperability** of components and grids (standards ?)
- Current IT culture is not predisposed to **sharing** resources
- Not all applications are grid-ready or **grid-enabled**
- **Open source** is not equal open source (read the little print)
- SLAs based on open source (**liability**?)
- "Static" **licensing** model don't embrace grid
- Protection of **intellectual property**
- **Legal** issues (FDA, HIPAA, multi-country grids)

Distributed European Infrastructure for Supercomputing Applications

DEISA

Examples of a successful Research e-Infrastructures

NEESGrid

Distributed European Infrastructure for Supercomputing Applications

Realtime access to earthquake
Shake table experiments at remote sites.

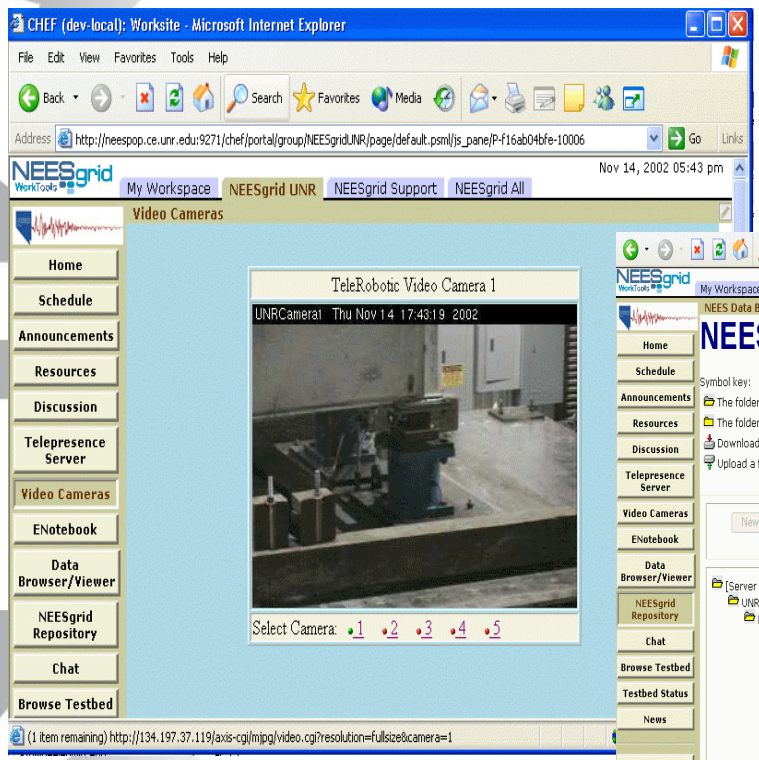

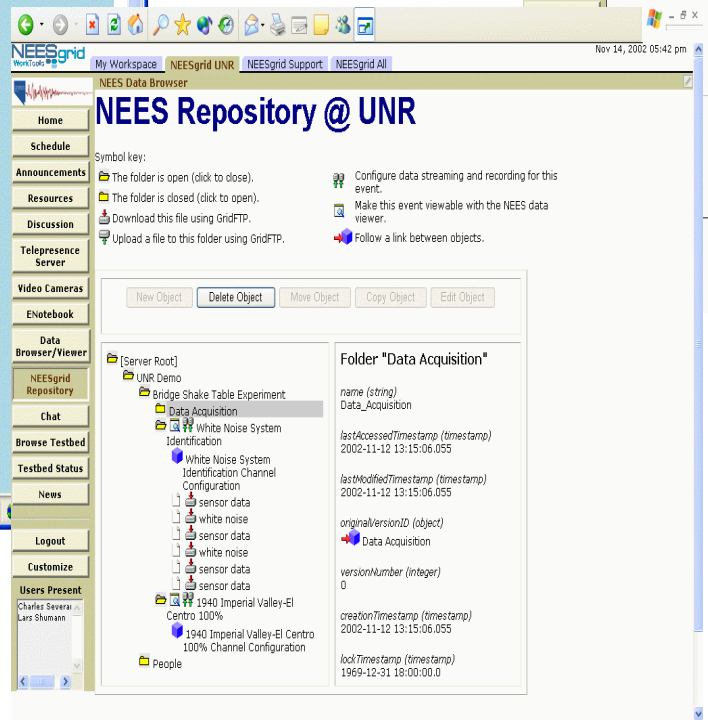

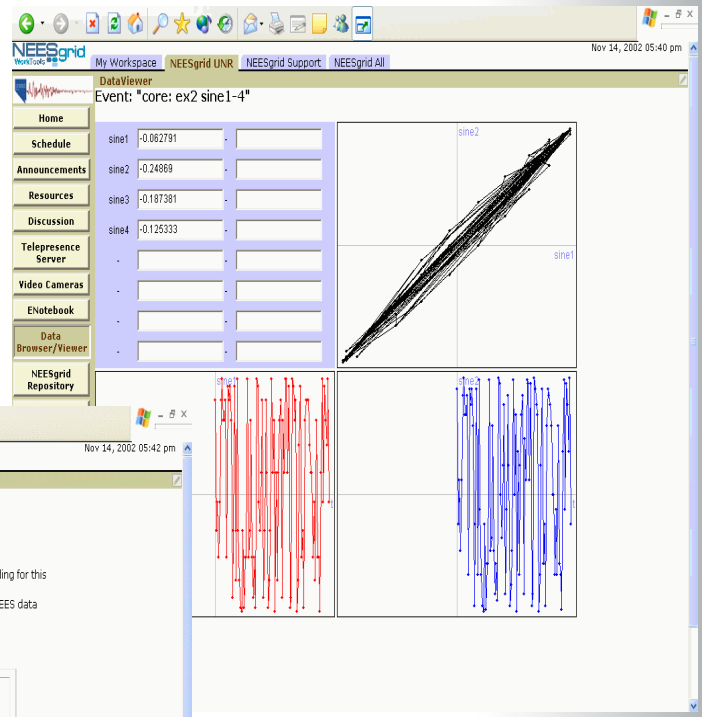

NEESgrid

RI-222919

BIRN – Biomedical Information

RI-222919

Geological Information Grid Portal

Distributed European Infrastructure for Supercomputing Applications

RI-222919

Mesoscale Meteorology

Distributed European Infrastructure for Supercomputing Applications

NSF LEAD project - making the tools that are needed to make accurate predictions of tornados and hurricanes.

- Data exploration and Grid workflow

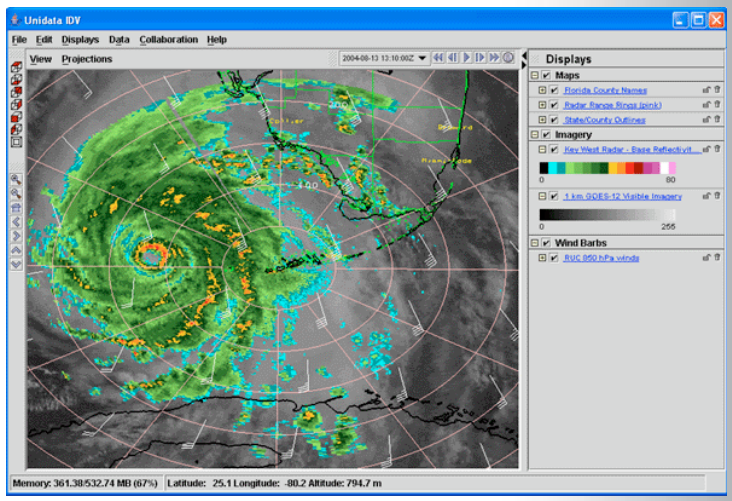

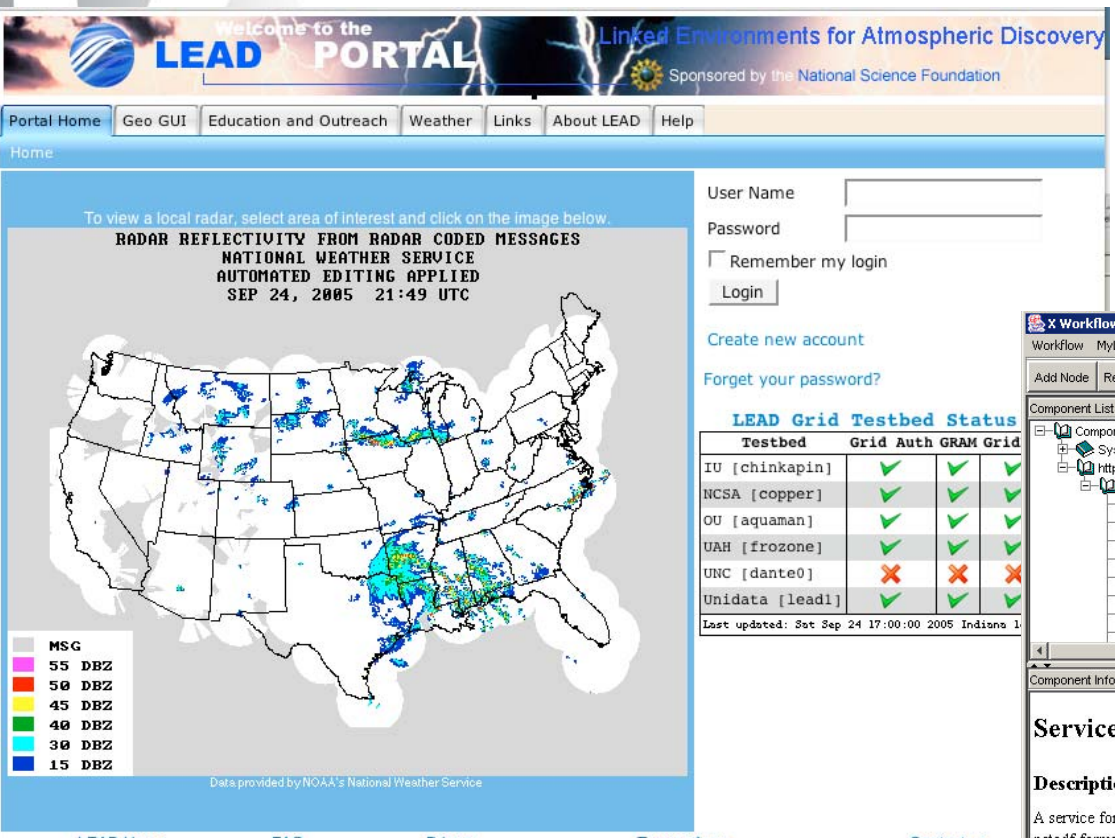

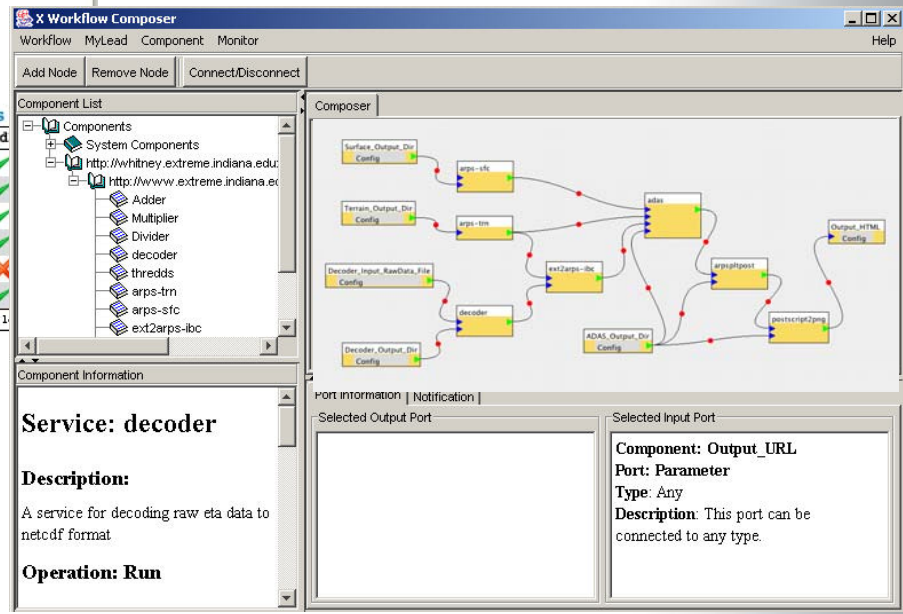

Renci Bio Portal

Distributed European Infrastructure for Supercomputing Applications

Providing access to biotechnology tools running on a back-end Grid.

- leverage state-wide investment in bioinformatics
- undergraduate & graduate education, faculty research
- another portal soon: national evolutionary synthesis center

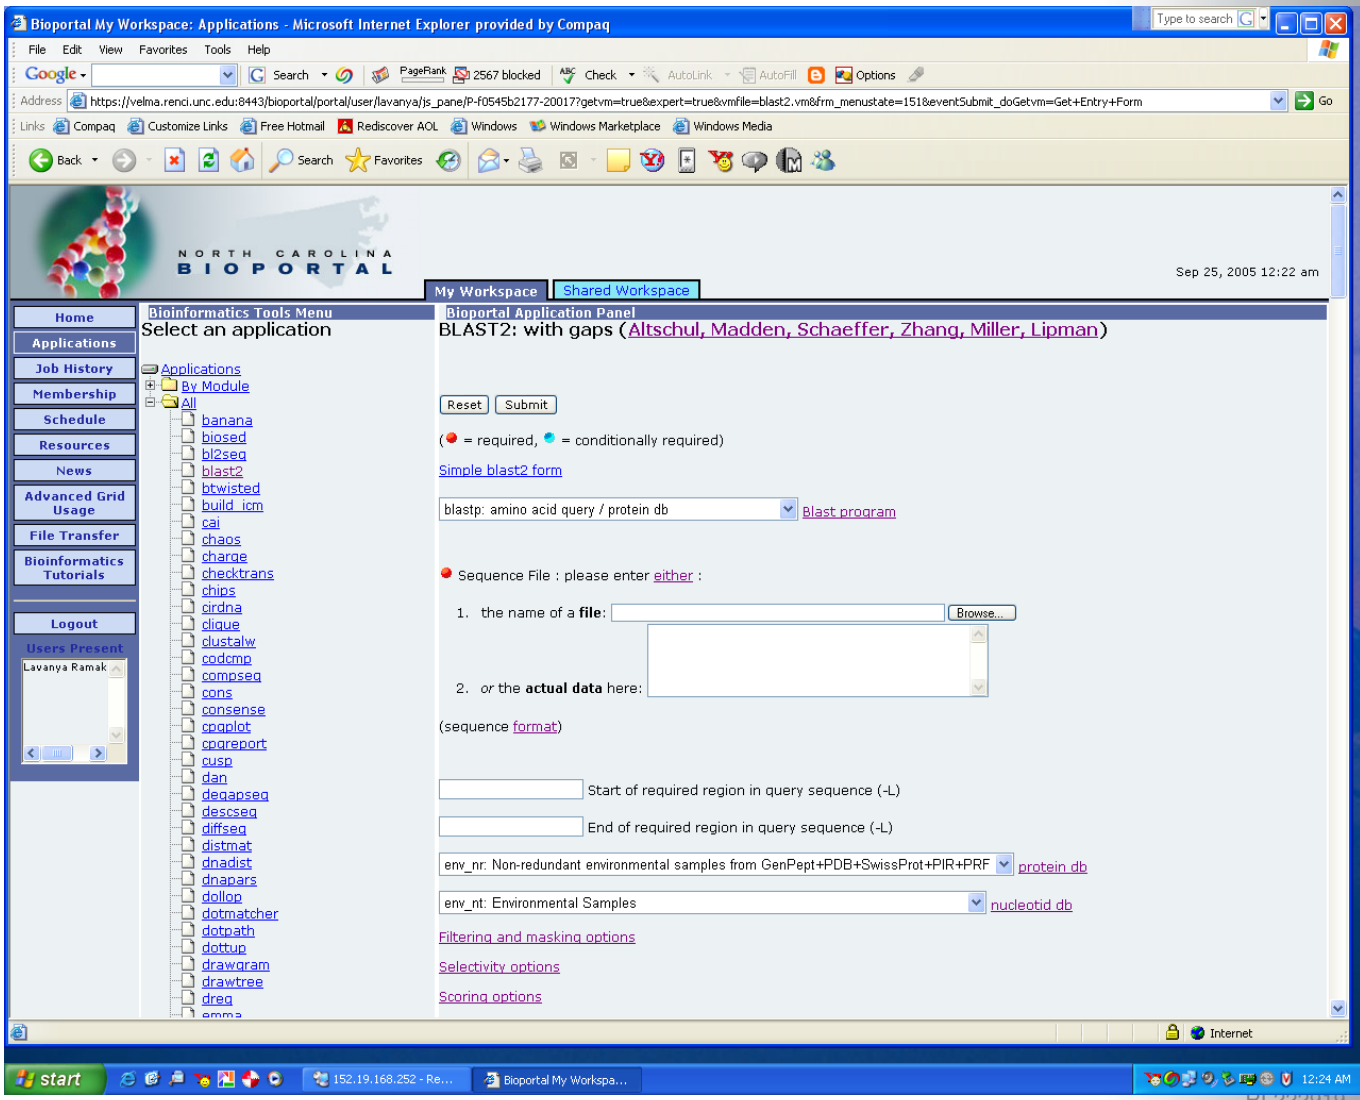

NESCent — National Evolutionary Synthesis Center

Nanohub - nanotechnology

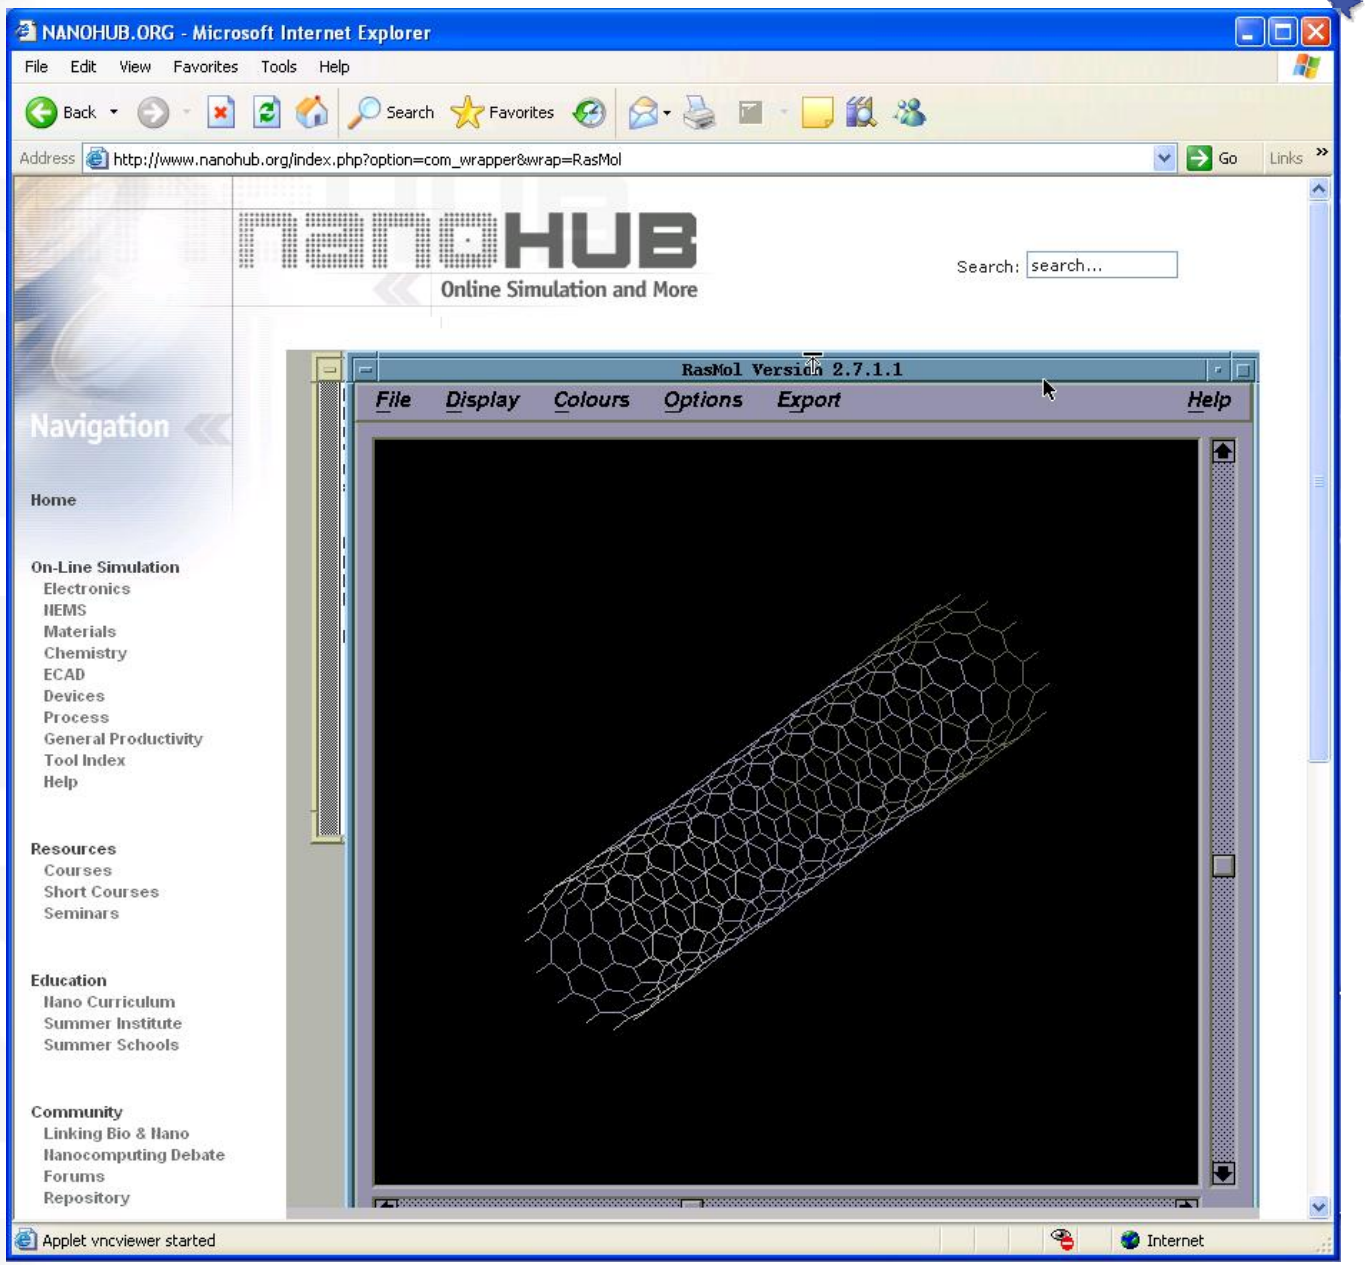

RI-222919

X-Ray Crystallography

Distributed European Infrastructure for Supercomputing Applications

C.I.M.A. Common Instrument Middleware Architecture

Welcome to the Crystallography Portal

Username: Password: Login

Remember me on this computer

Login Help

Home | Current Status | Data Repository | About

- IUB IUMSC
- IUB Myers Hall
- Purdue Crystallography Center
- CSAF Sydney, Australia
- Minnesota X-ray Lab
- ChemMatCARS - Univ. of Chicago at APS
- Other collaborators
- NCS Southampton, UK

The Purdue Chemistry Crystallography Center

The Purdue Chemistry Crystallography Center

Disable your browser's cache to get the live stream!

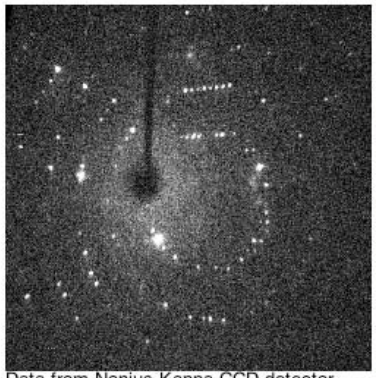

Data from Nonius Kappa CCD detector
(Under development!)
Total Number of jpg: 10
Frame: s01f0010.jpg
All available jpg images
Browse the 20 latest jpg images

<□>

Streaming video from the lab showing the Nonius instrument

<□>

Streaming video from the crystal microscope on the Nonius diffractometer

Local date/time: 2005-09-24 11:36:54

These values are updated approx. every 60 sec. Times in UTC

LabJack U12

Instrument Enclosure Temp. & Humidity:	23.4 C Rel. Humid. 43.1 %	2005-09-24 16:35:59
Chill Water In:	16.4 C	2005-09-24 16:36:25
Chill Water Out:	19.3 C	2005-09-24 16:36:25
Generator Relay Voltage: All previous voltages	3.42 X-ray Generator is: OFF	2005-09-24 16:36:48

ServoGrid Portal

Distributed European Infrastructure for Supercomputing Applications

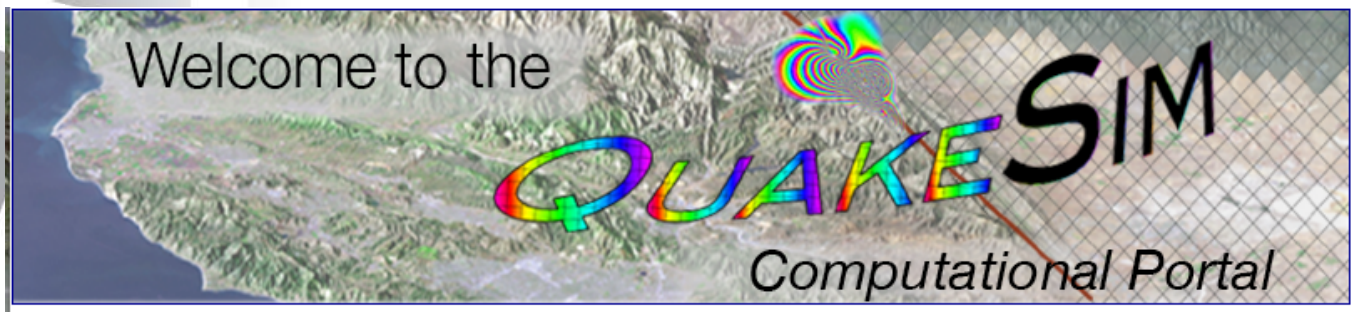

Username:
Password:
Login
Create New Account | Login Help

SERVOGrid FaultDB_Search QuakeTables Portal Search

SERVOGrid

SERVO Grid

Solid Earth Research Virtual Observatory Grid

- QuakeSim home page.
- Old GEM General Earthquake Modeling Web Site
- SLIDE Distributed File System for NASA Computational Technology Project
- Report from the Earth Science Enterprise Computational Technology Requirements Workshop April 30-May 1 2002 where SERVO concept first introduced
- Discover the Grid at the Grid Forum or at this collection of papers
- Other collected papers and presentations on SERVOGrid and related topics are available from the Community Grids Lab publications page.

QuakeSim Web Portal | User Manual | Support | Report Bugs | QuakeSim Web Site

Participating Institutions:

IU CGL | NASA JPL | UC Davis | UC Irvine | USC

Belfast Gene Grid Portal

Distributed European Infrastructure for Supercomputing Applications

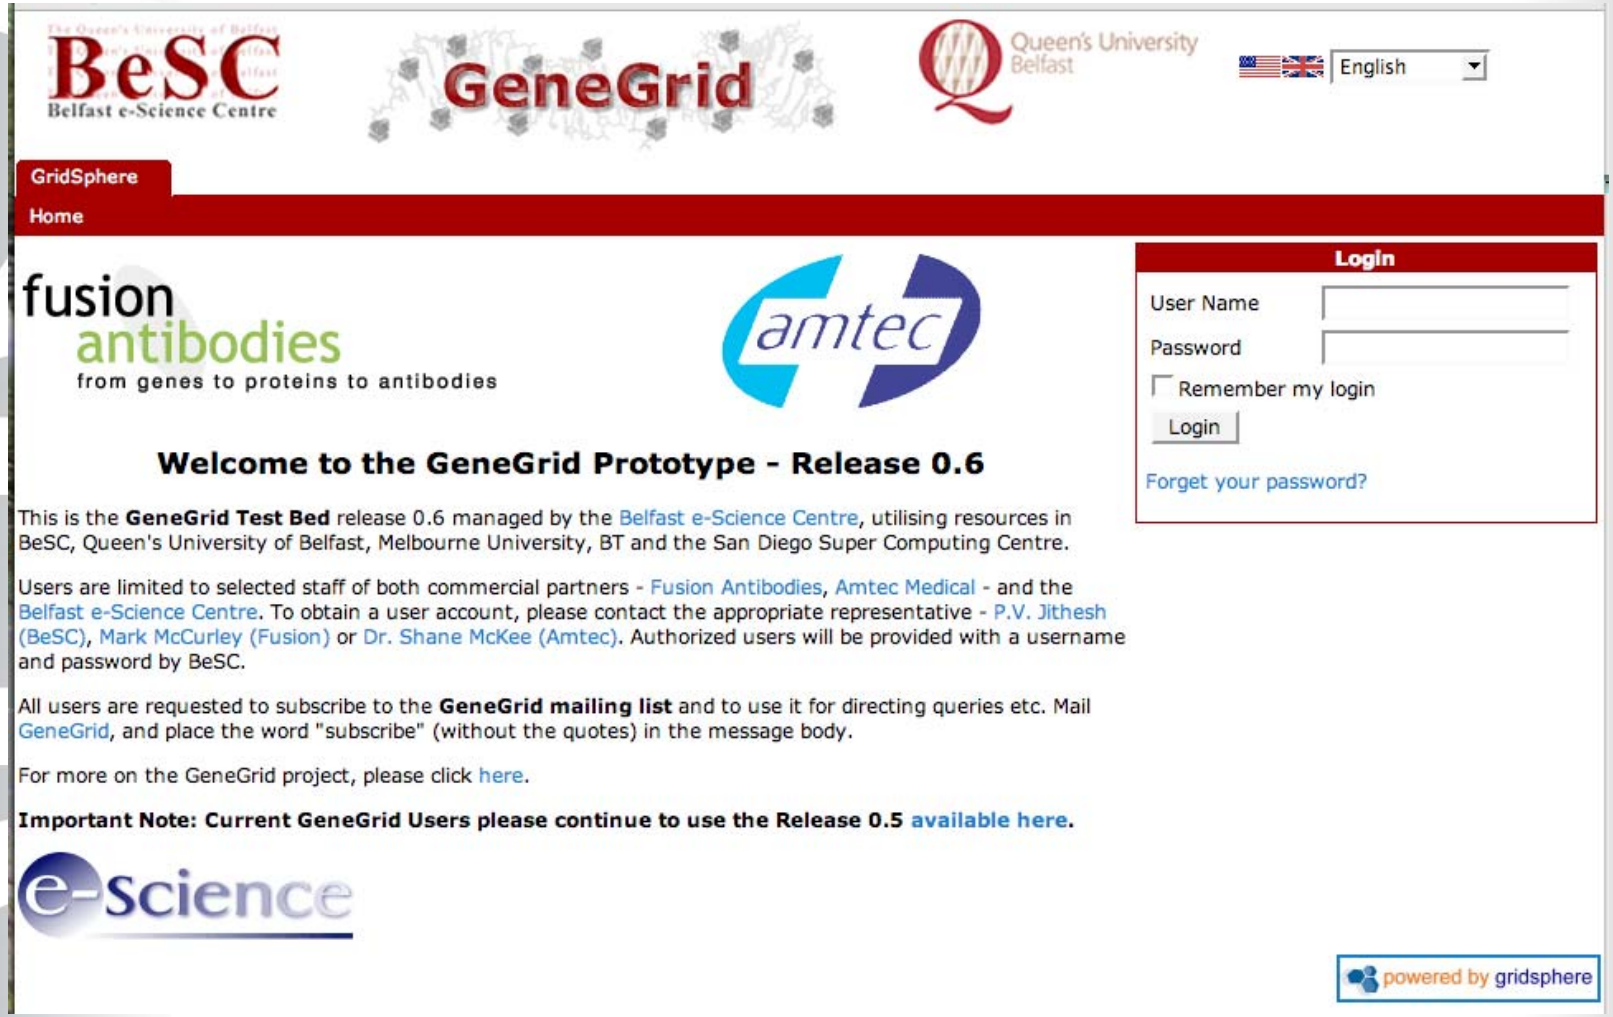

RI-222919

MyGrid - Bioinformatics

Distributed European Infrastructure for Supercomputing Applications

myGrid

Navigate

- Home
- About
- Downloads
- Components
  - Component Overview
  - Research Components
- Using myGrid
- Research Using myGrid
- Links
- Publications
- Contact

Log In

Username

Password

Log In

New Account Signup
Forgot Password

myGrid Architecture

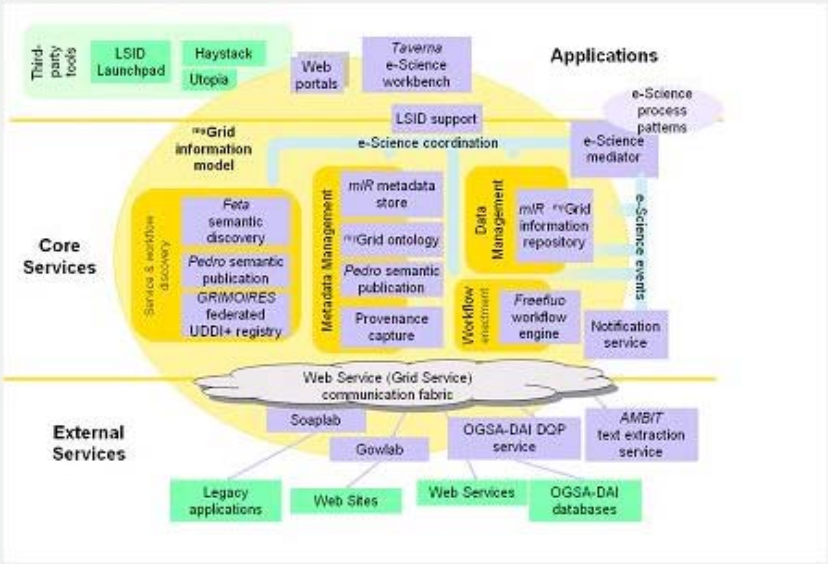

myGrid components - overview

myGrid is a collection of services and components that allows the high level integration of biological applications. The architecture provides the infrastructure necessary, in a web service environment, t e-science workbench that actively supports the scientific lifecycle. Each component or service contri system that allows the e-scientist to perform complex in-silico experiments across distributed bioinfo resources.

RI-222919

Distributed European Infrastructure for Supercomputing Applications

Example of a successful Research e-Infrastructure:

The DEISA Ecosystem for HPC Grand-Challenge Applications

Distributed European Infrastructure for Supercomputing Applications

new "petaflop" supercomputers

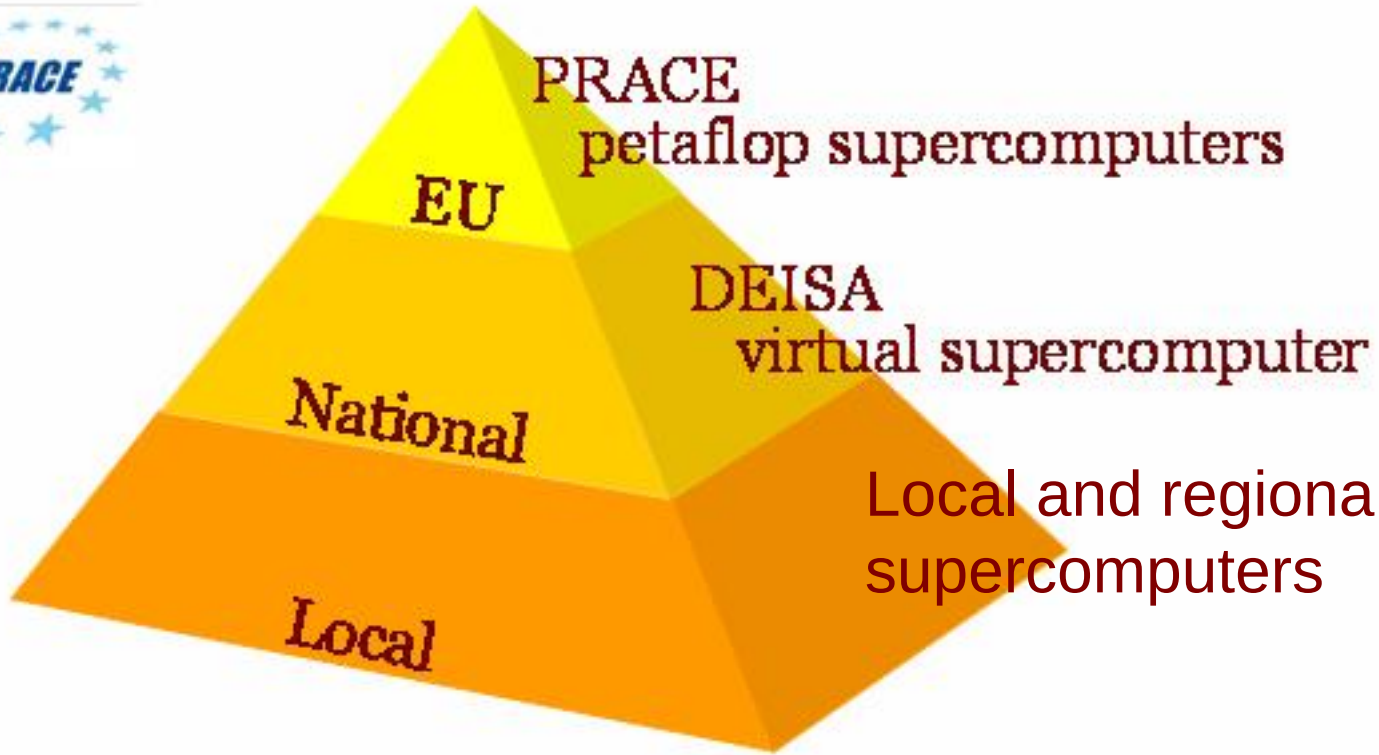

Local and regional supercomputers

Mario Campolargo
European Commission
OGF23, June 2008

DEISA: Vision and Mission

Distributed European Infrastructure for Supercomputing Applications

Vision:

Persistent European HPC ecosystem integrating Tier-1 (Tflop/s) centres and European Tier-0 (Pflop/s) centres.

Mission:

Enhance Europe's capability in computing and science by integrating most powerful supercomputers into a European HPC e-infrastructure.

Built European Supercomputing Service on top of existing national services, based on the deployment and operation of a persistent, production quality, distributed supercomputing environment with continental scope.

DEISA Partners

Distributed European Infrastructure for Supercomputing Applications

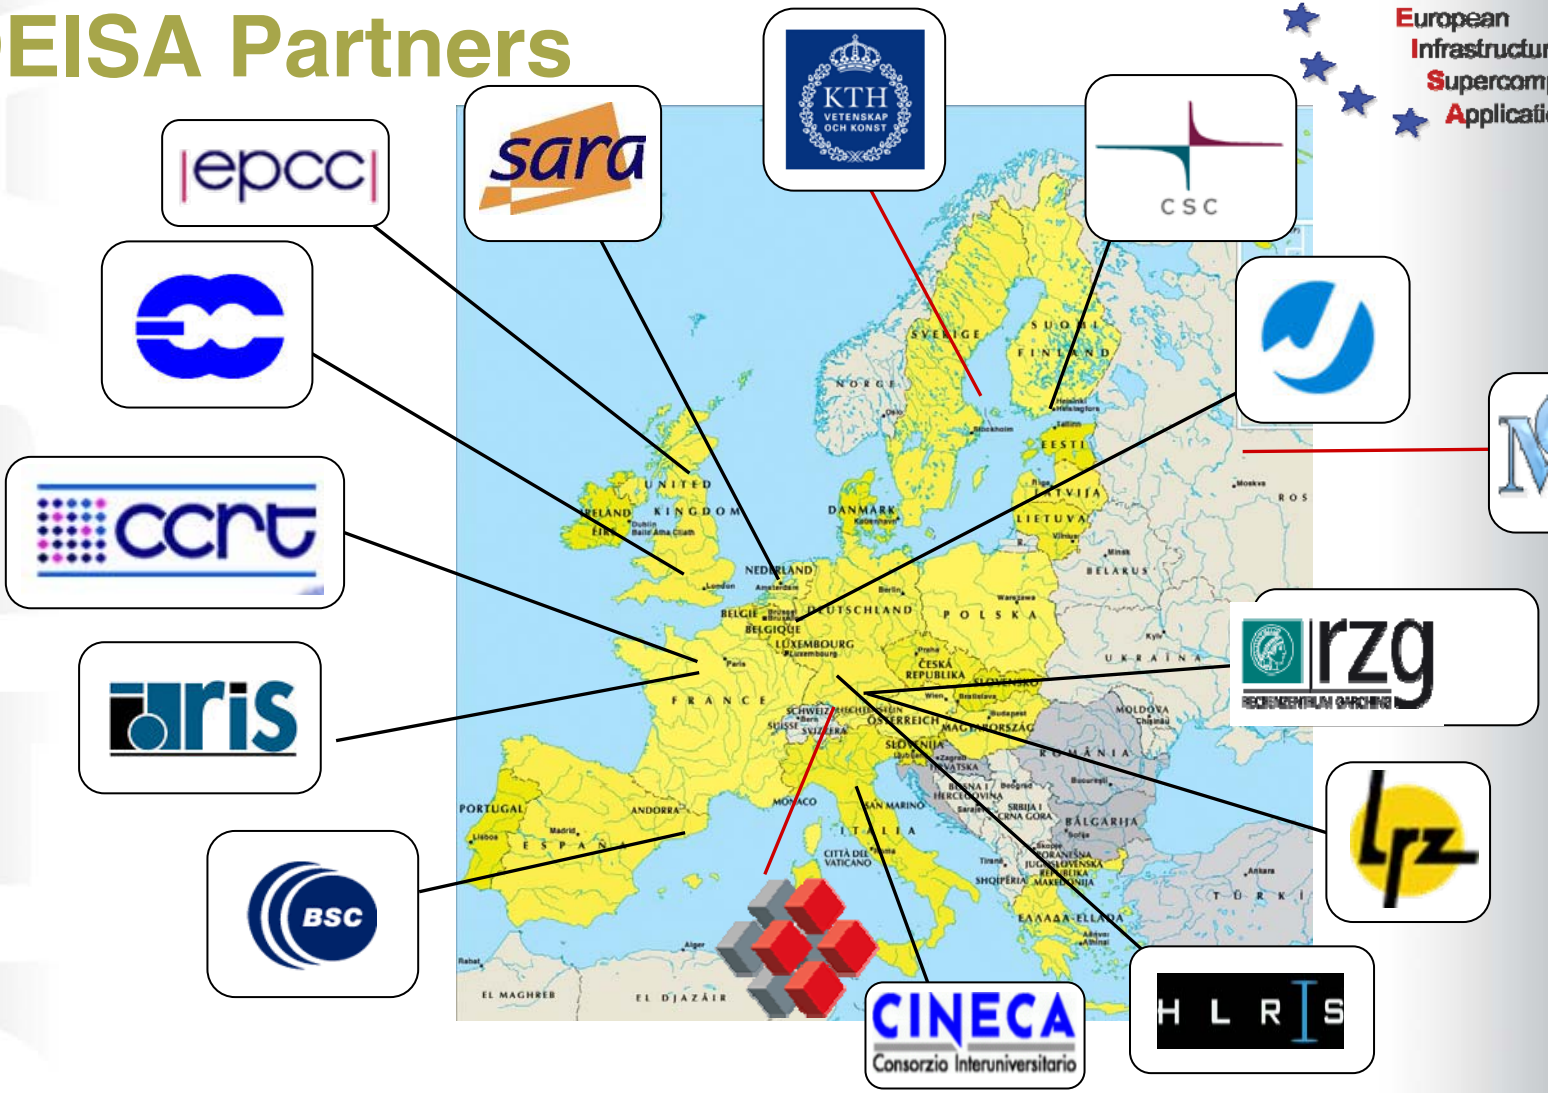

DEISA1: May 1st, 2004 – April 30th, 2008

DEISA2: May 1st, 2008 – April 30th, 2011

RI-222919

DEISA dedicated high speed network on GEANT2 and the NRENs

Distributed European Infrastructure for Supercomputing Applications

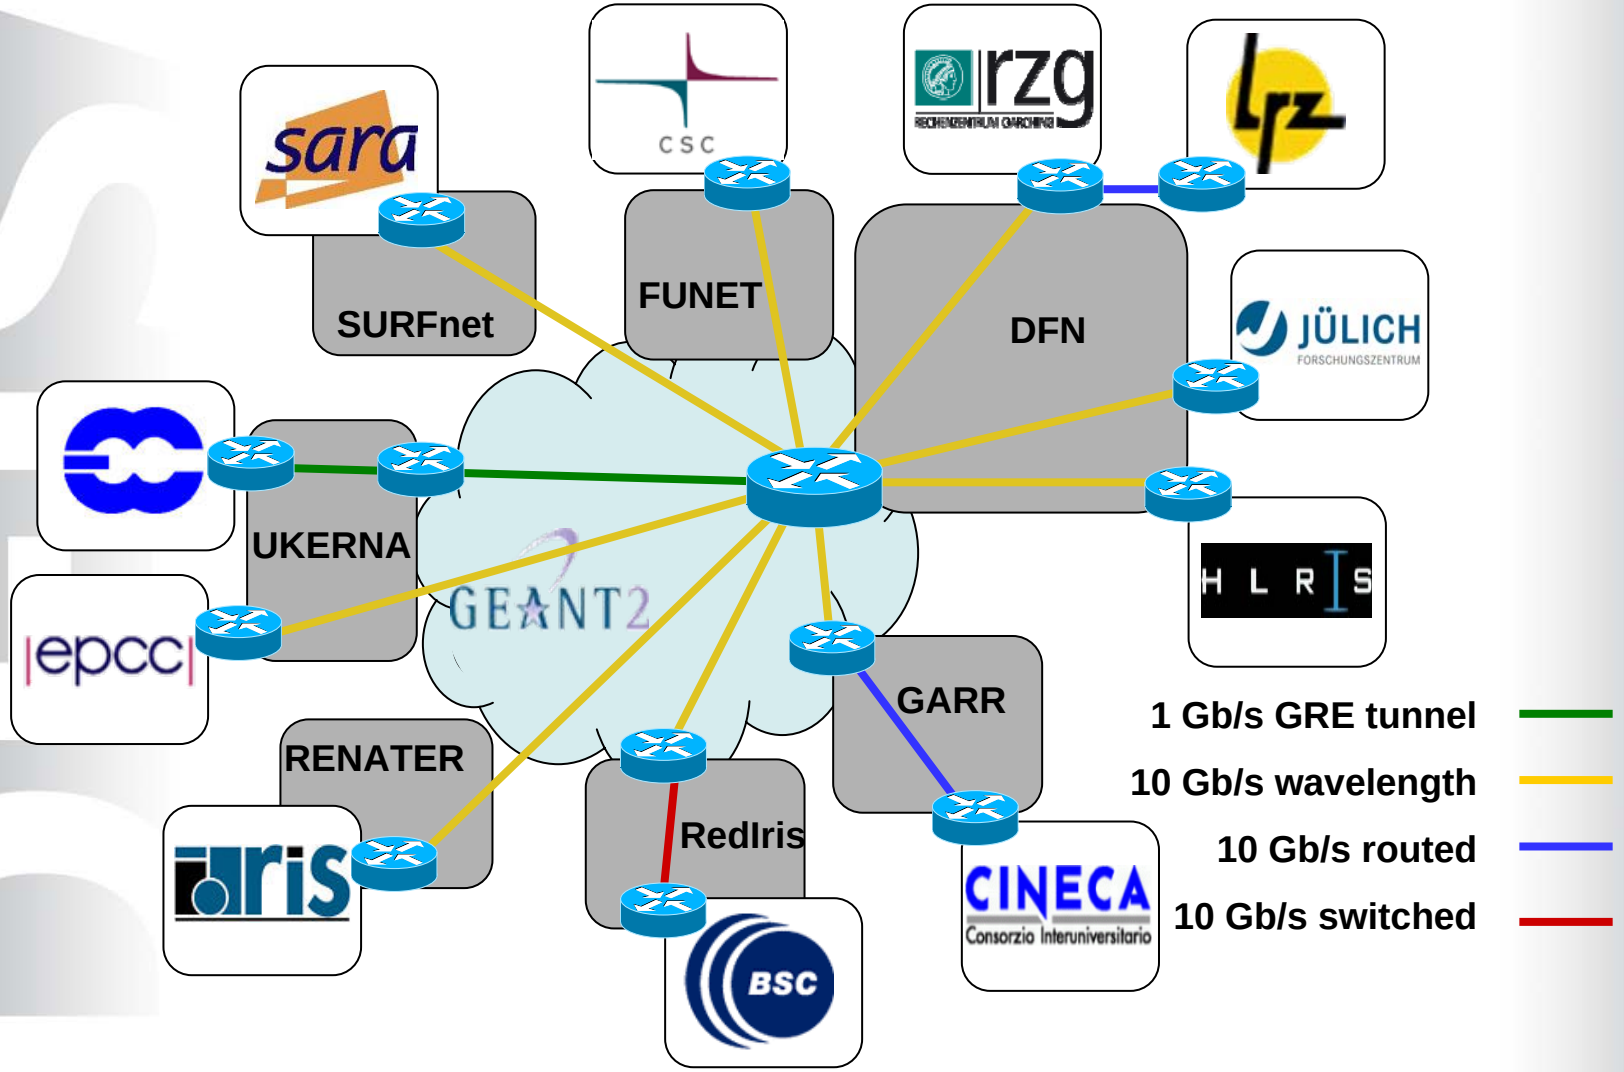

RI-222919

Categories of DEISA services

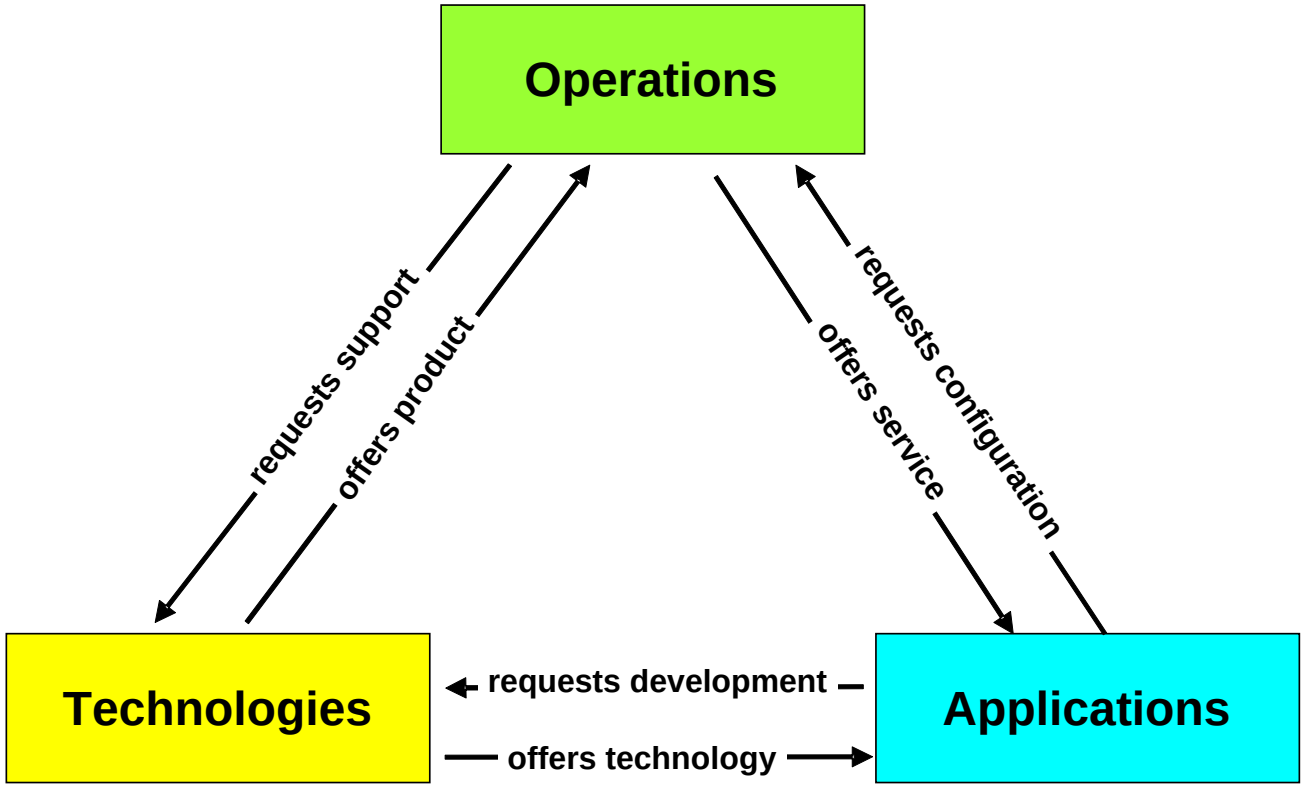

RI-222919

DEISA Service Layers

Distributed European Infrastructure for Supercomputing Applications

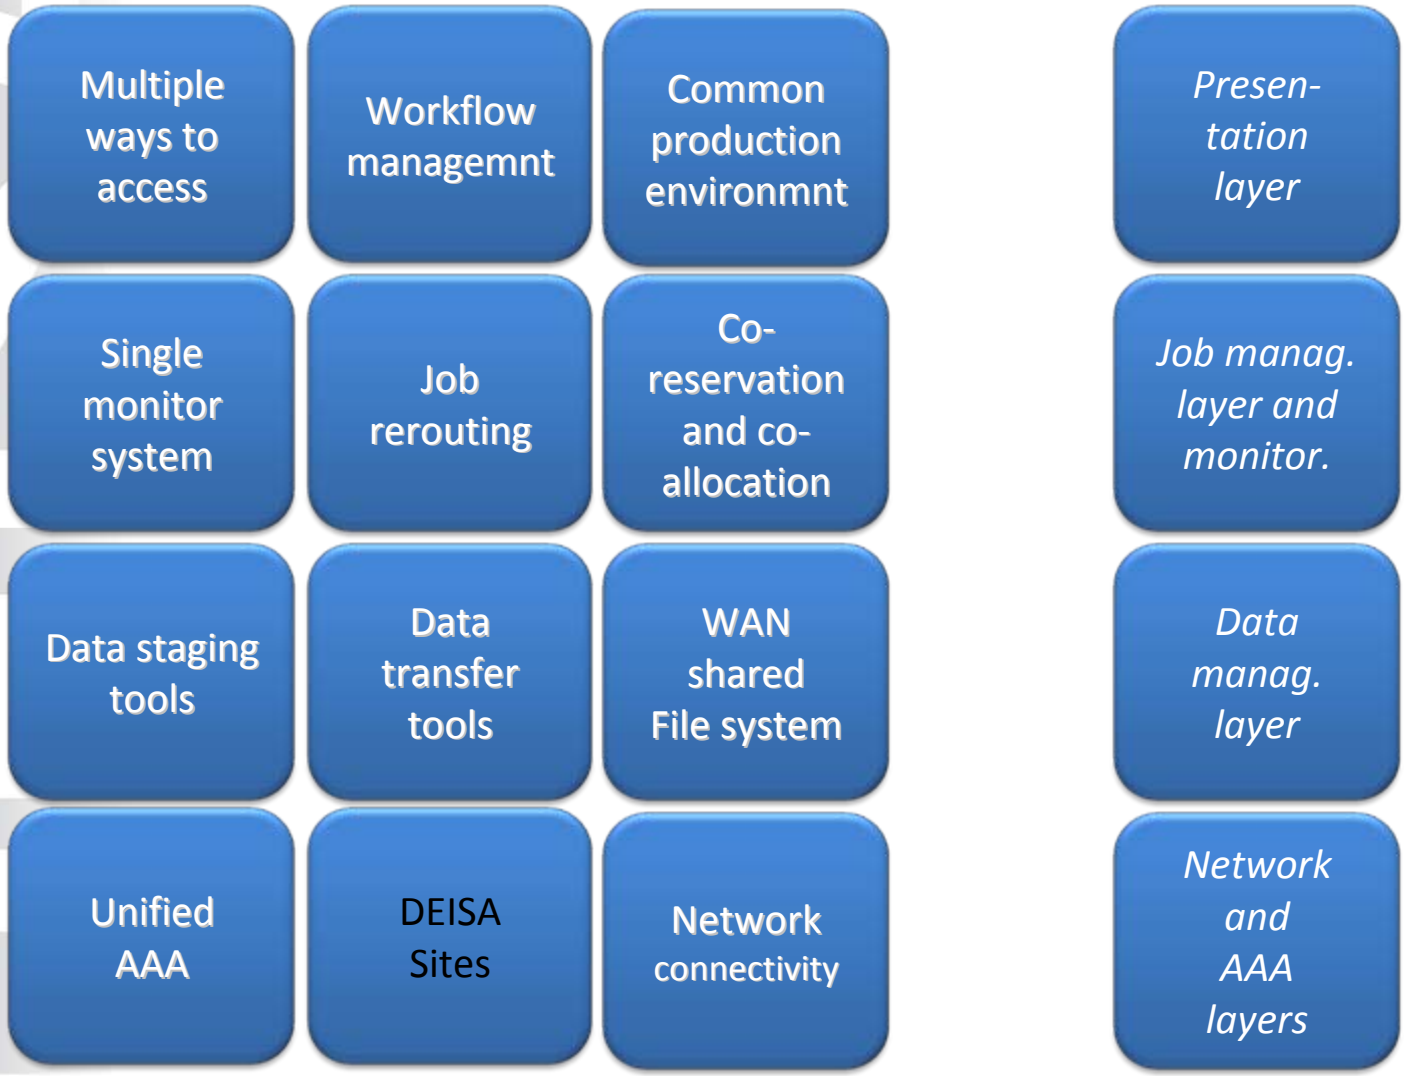

RI-222919

DEISA UNICORE Infrastructure

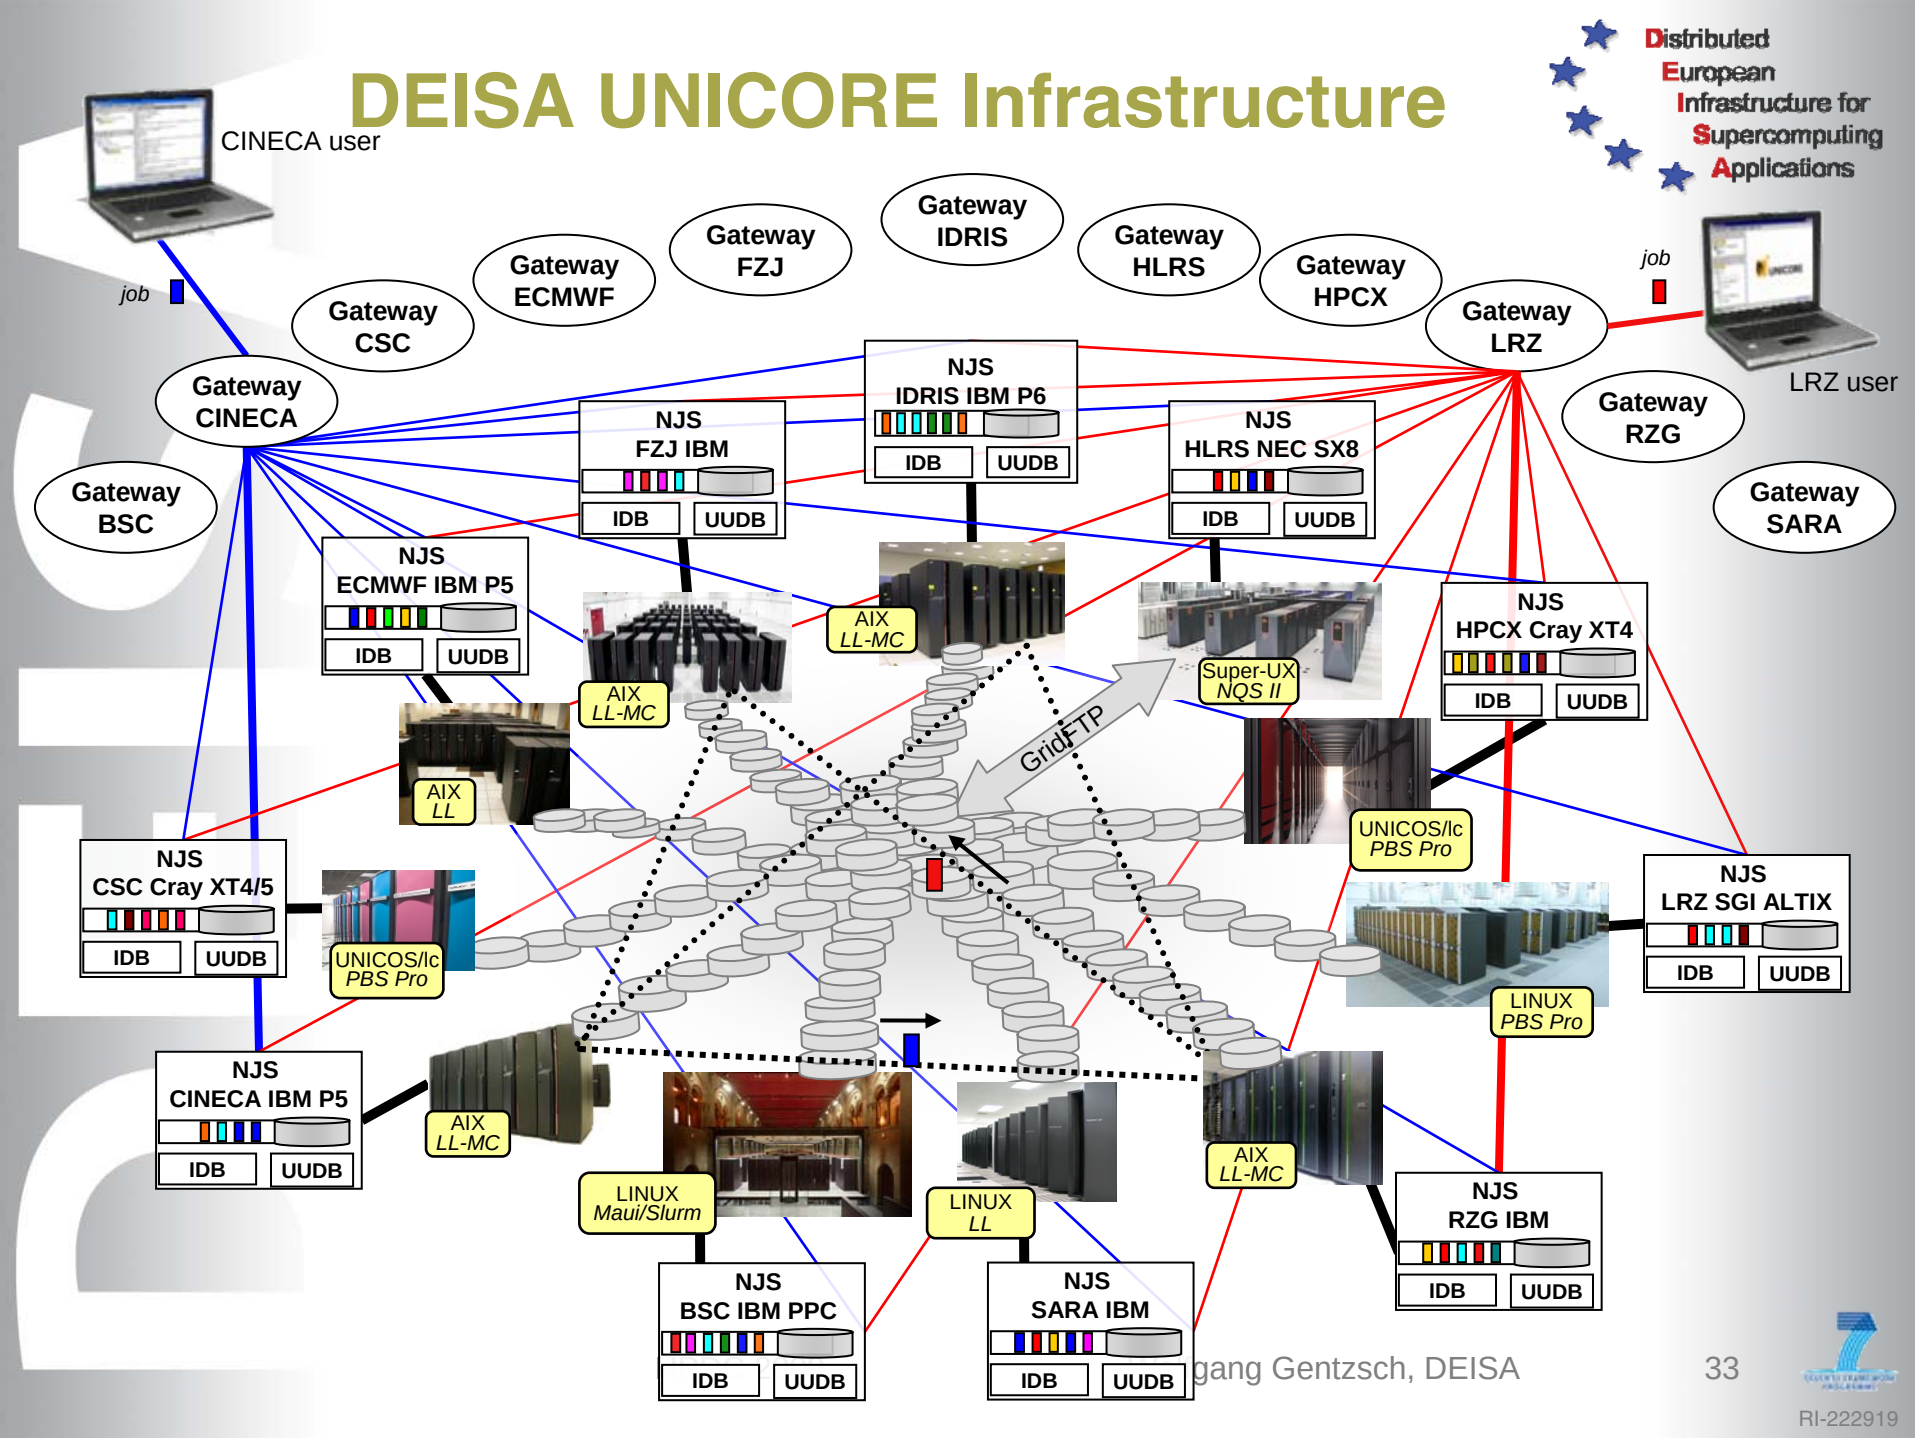

RI-222919

DEISA Global File System

Distributed European Infrastructure for Supercomputing Applications

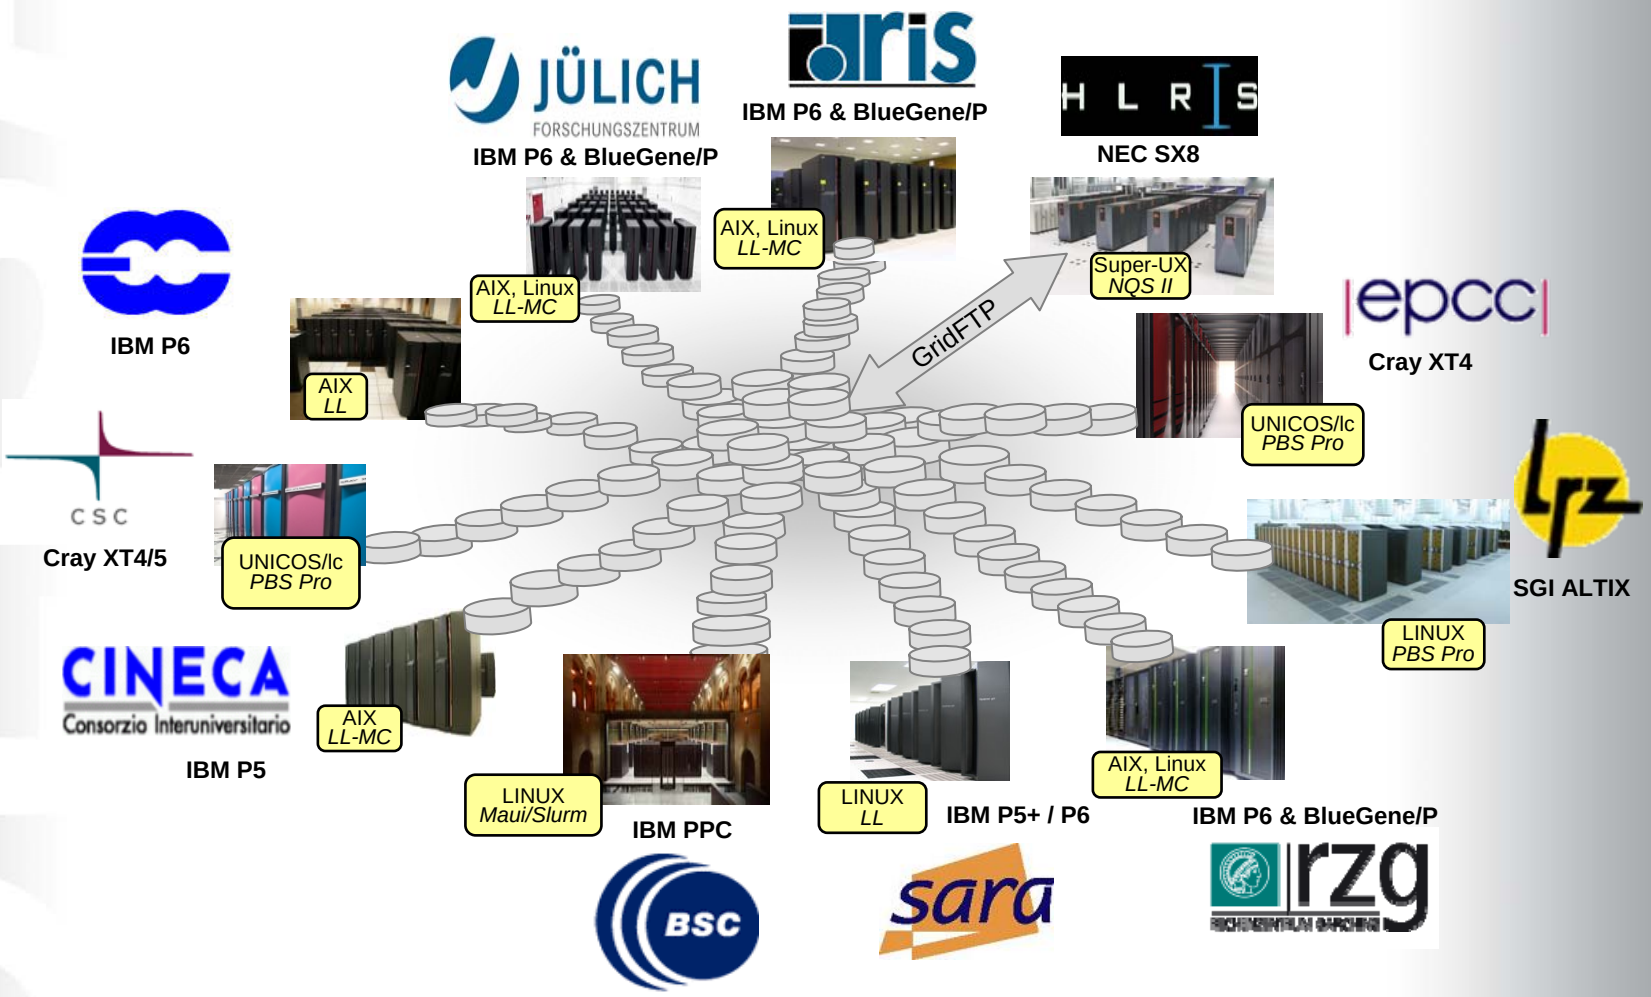

Global transparent file system based on the Multi-Cluster General Parallel File System (MC-GPFS of IBM)

Management of users in DEISA

- A dedicated LDAP-based distributed repository administers DEISA users
- Trusted LDAP servers are authorized to access each other (based on X.509 certificates) and encrypted communication is used to maintain confidentiality

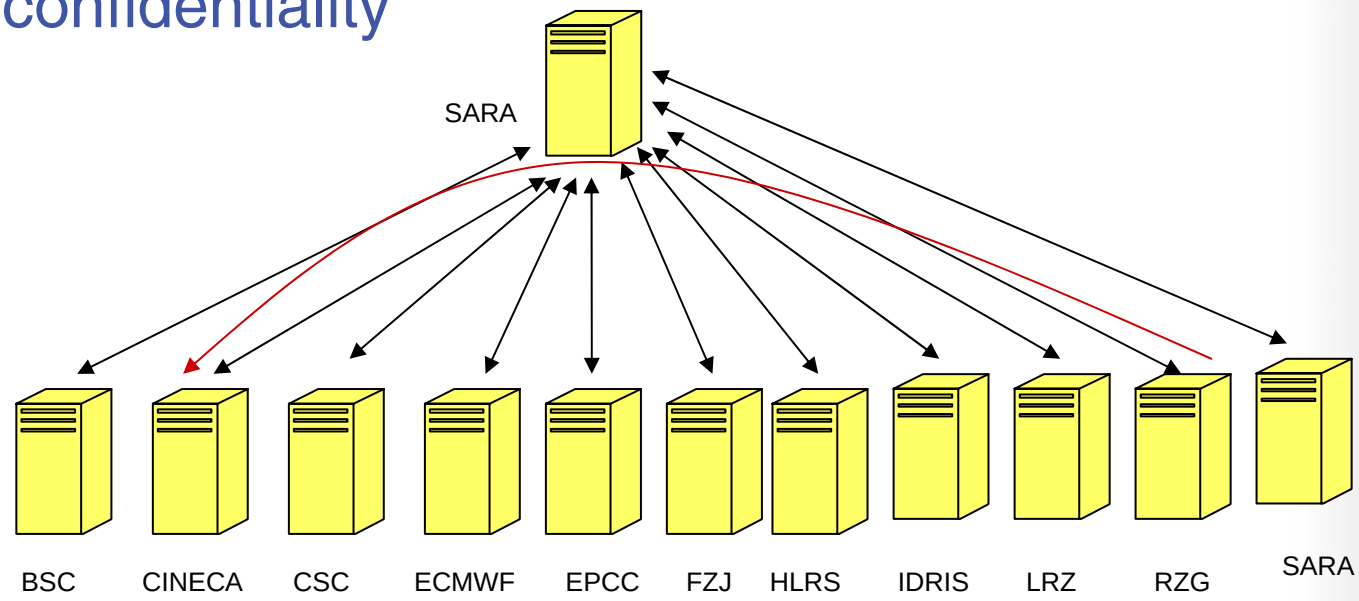

RI-222919

Common User Administration

- Each partner is responsible for the registration of users affiliated to the partner (home organization)
- Other partners update local user administration (LDAP, NIS, /etc/passwd) with data from other sites on a daily basis. Based on trust between partners!

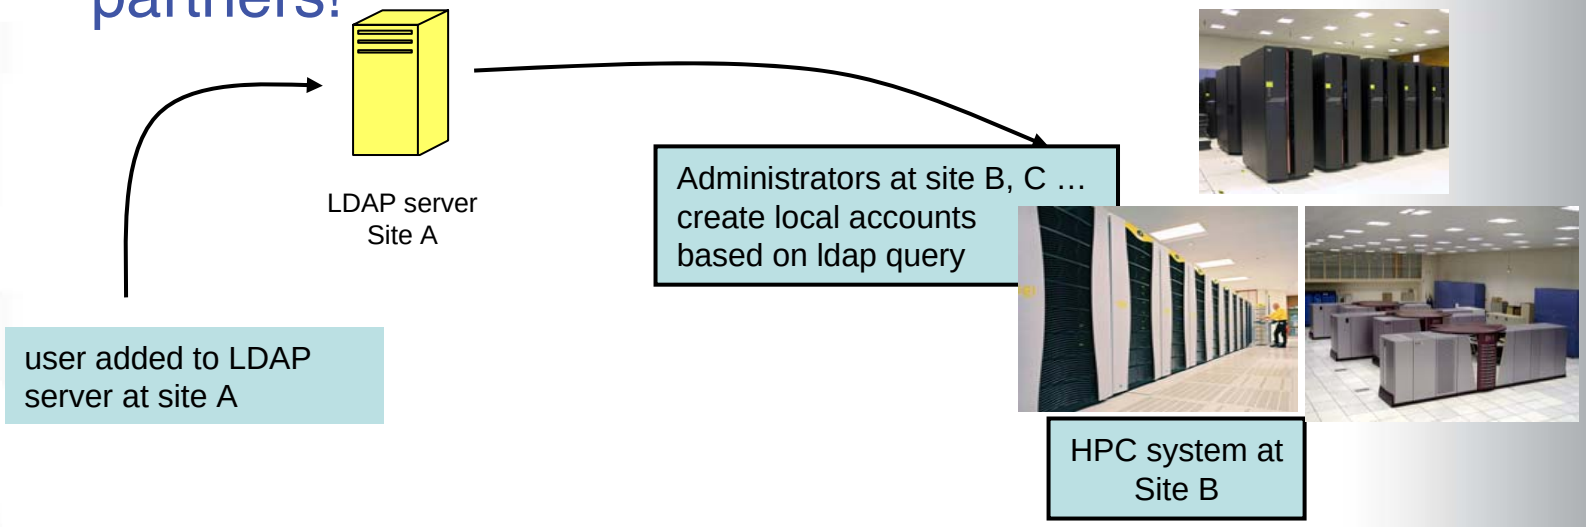

DEISA Extreme Computing Initiative (DECI)

- DECI launched in 2005: complex, demanding, innovative simulations requiring the exceptional capabilities of DEISA
- Multi-national proposals encouraged
- Proposals reviewed by national evaluation committees
- Projects chosen on the basis of innovation potential, scientific excellence, relevance criteria, and national priorities
- Most powerful HPC architectures for most challenging projects
- Most appropriate supercomputer architecture selected

DEISA Extreme Computing Initiative

Distributed European Infrastructure for Supercomputing Applications

Calls for Proposals for challenging supercomputing projects from all areas of science

DECI call 2005

51 proposals, 12 European countries involved, co-investigator from US)
30 mio cpu-h requested
29 proposals accepted, 12 mio cpu-h awarded (normalized to IBM P4+)

DECI call 2006

41 proposals, 12 European countries involved
co-investigators from N + S America, Asia (US, CA, AR, ISRAEL)
28 mio cpu-h requested
23 proposals accepted, 12 mio cpu-h awarded (normalized to IBM P4+)

DECI call 2007

63 proposals, 14 European countries involved, co-investigators from N + S America, Asia, Australia (US, CA, BR, AR, ISRAEL, AUS)
70 mio cpu-h requested
45 proposals accepted, ~30 mio cpu-h awarded (normalized to IBM P4+)

DECI call 2008 (ending June 30, 2008)

66 proposals, 15 European countries involved, co-investigators from N + S America, Asia, Australia
134 mio cpu-h requested (normalized to IBM P4+)
42 proposals accepted, 48 mio cpu-h awarded (normalized to IBM P4+)

RI-222919

Distributed European Infrastructure for Supercomputing Applications

Example of a Next-Generation e-Infrastructure in Industry

Telecommunications

The silent revolution in Telecommunications:

Convergence of smart phones, broadband, and cloud computing

New Powerful End-User Devices

Distributed European Infrastructure for Supercomputing Applications

Carry-along PCs (CAPS), Ultra-Mobile PCs (UMPC), Smart mobile phones

CAPC, from Samsung, South Korea

Fujitsu UMPC

OLPC from MIT

T83 Tablet from Asus, Taiwan, demoed at CeBit 2007

CAPC from HTC

Asus R2H

The Symbian Mobile OS

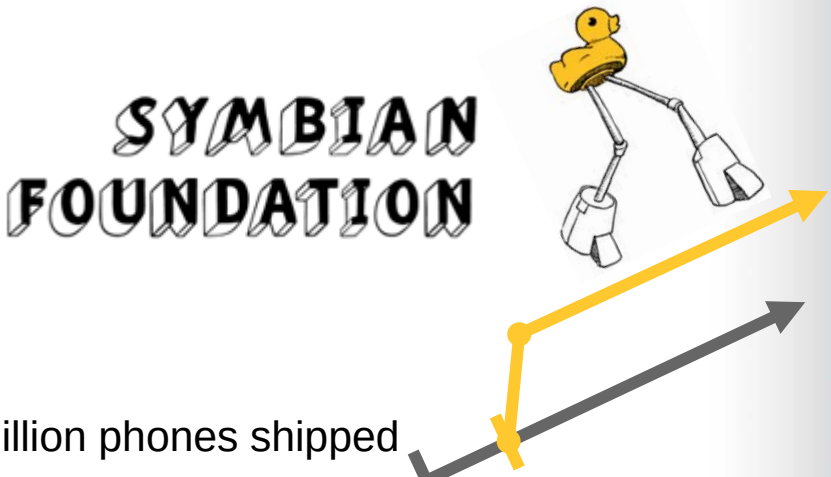

1998: Symbian Ltd was founded by Ericsson, Motorola, Nokia and Psion to create Symbian OS

2006: 100 million phones shipped

2008: 250 different models since Symbian's creation

2009: 250 million phones shipped

- Nokia acquires Symbian Ltd
- The Symbian Foundation is announced by the initial board members.

The Symbian OS

Powerful mobile platform for deep-content services

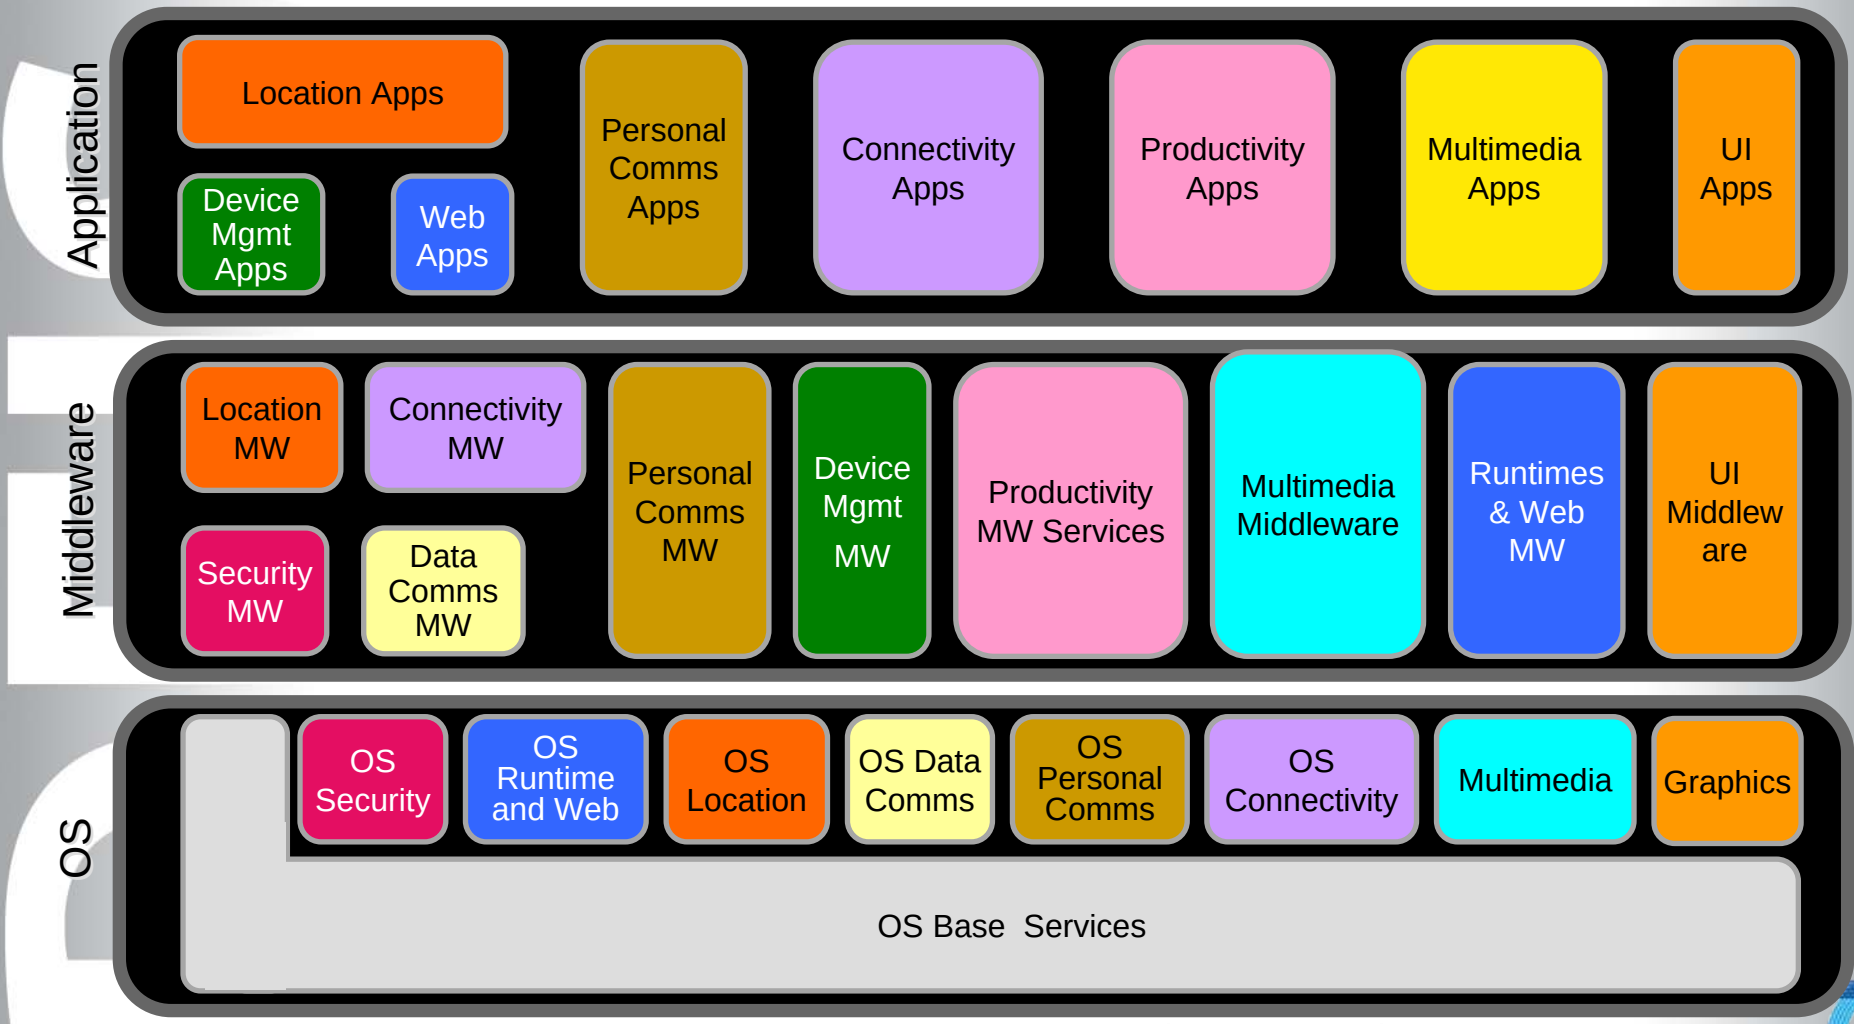

RI-222919

Standard Grid Concepts

are often not scalable for millions of users

Distributed European Infrastructure for Supercomputing Applications

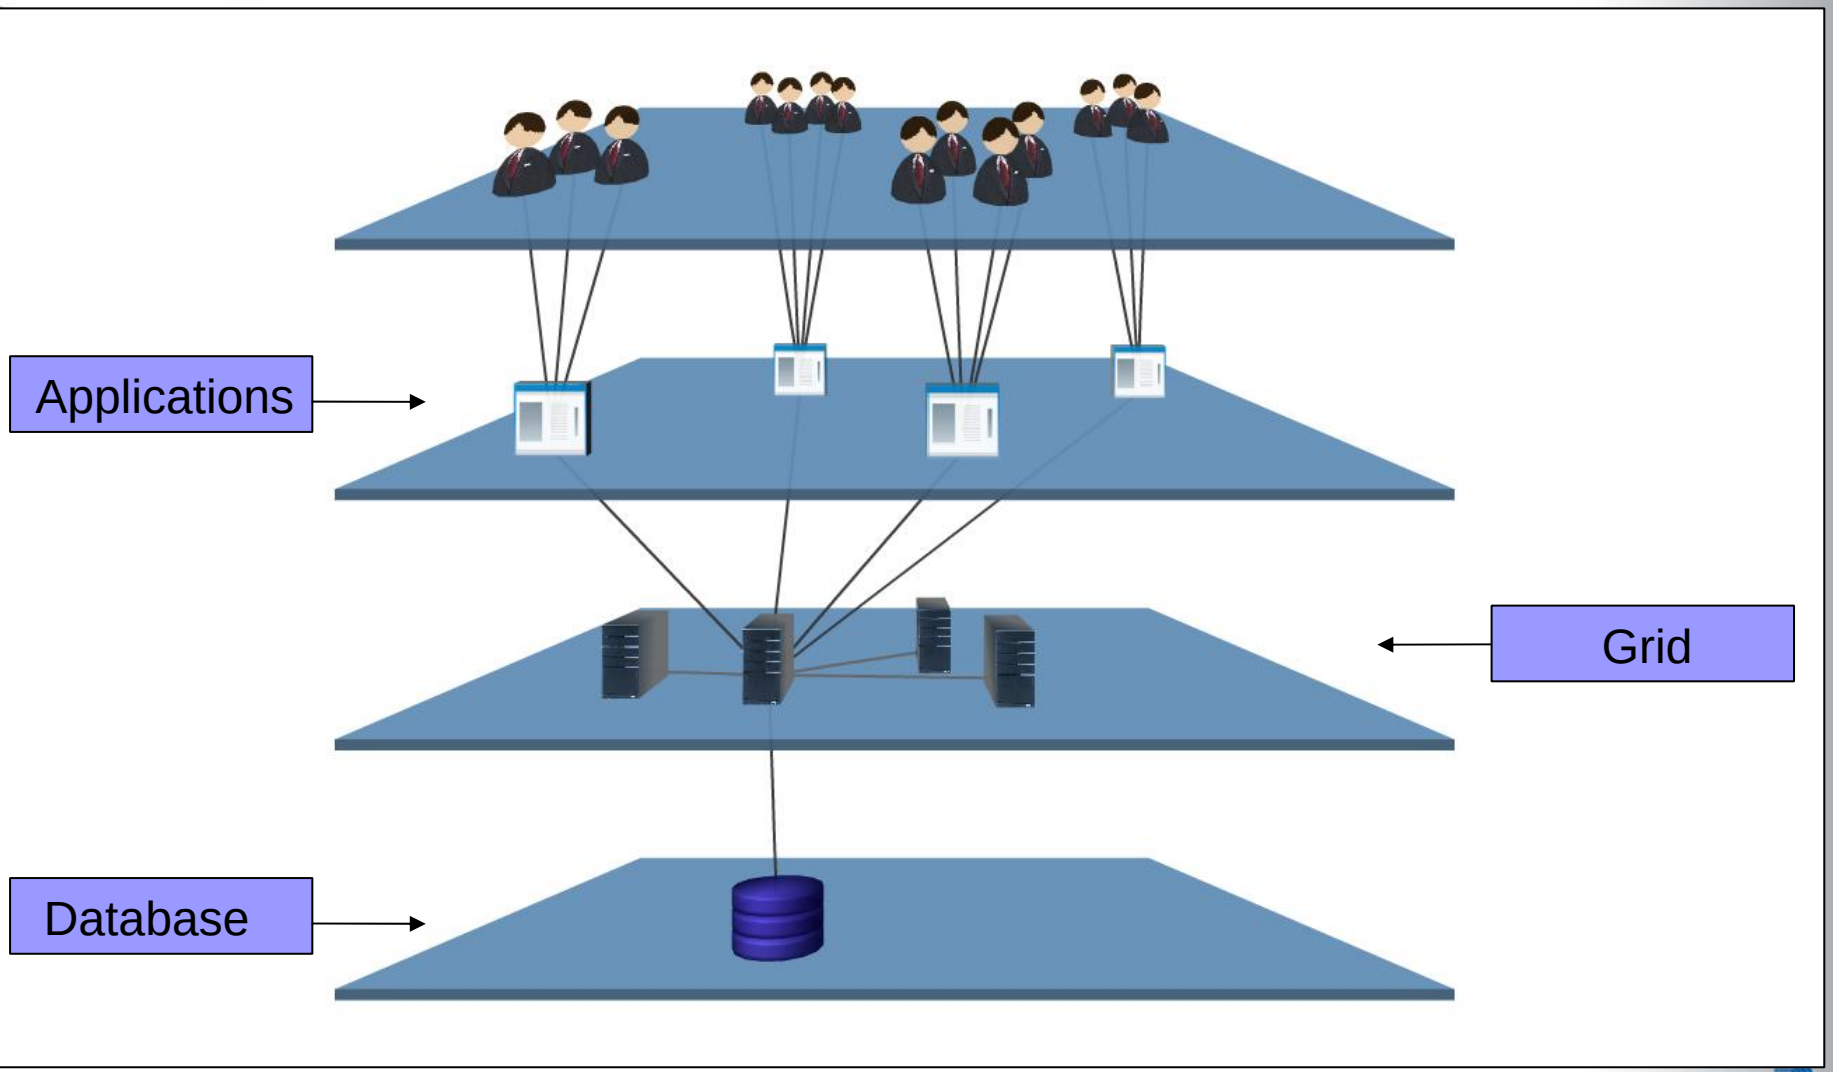

Grid and virtual cache concept

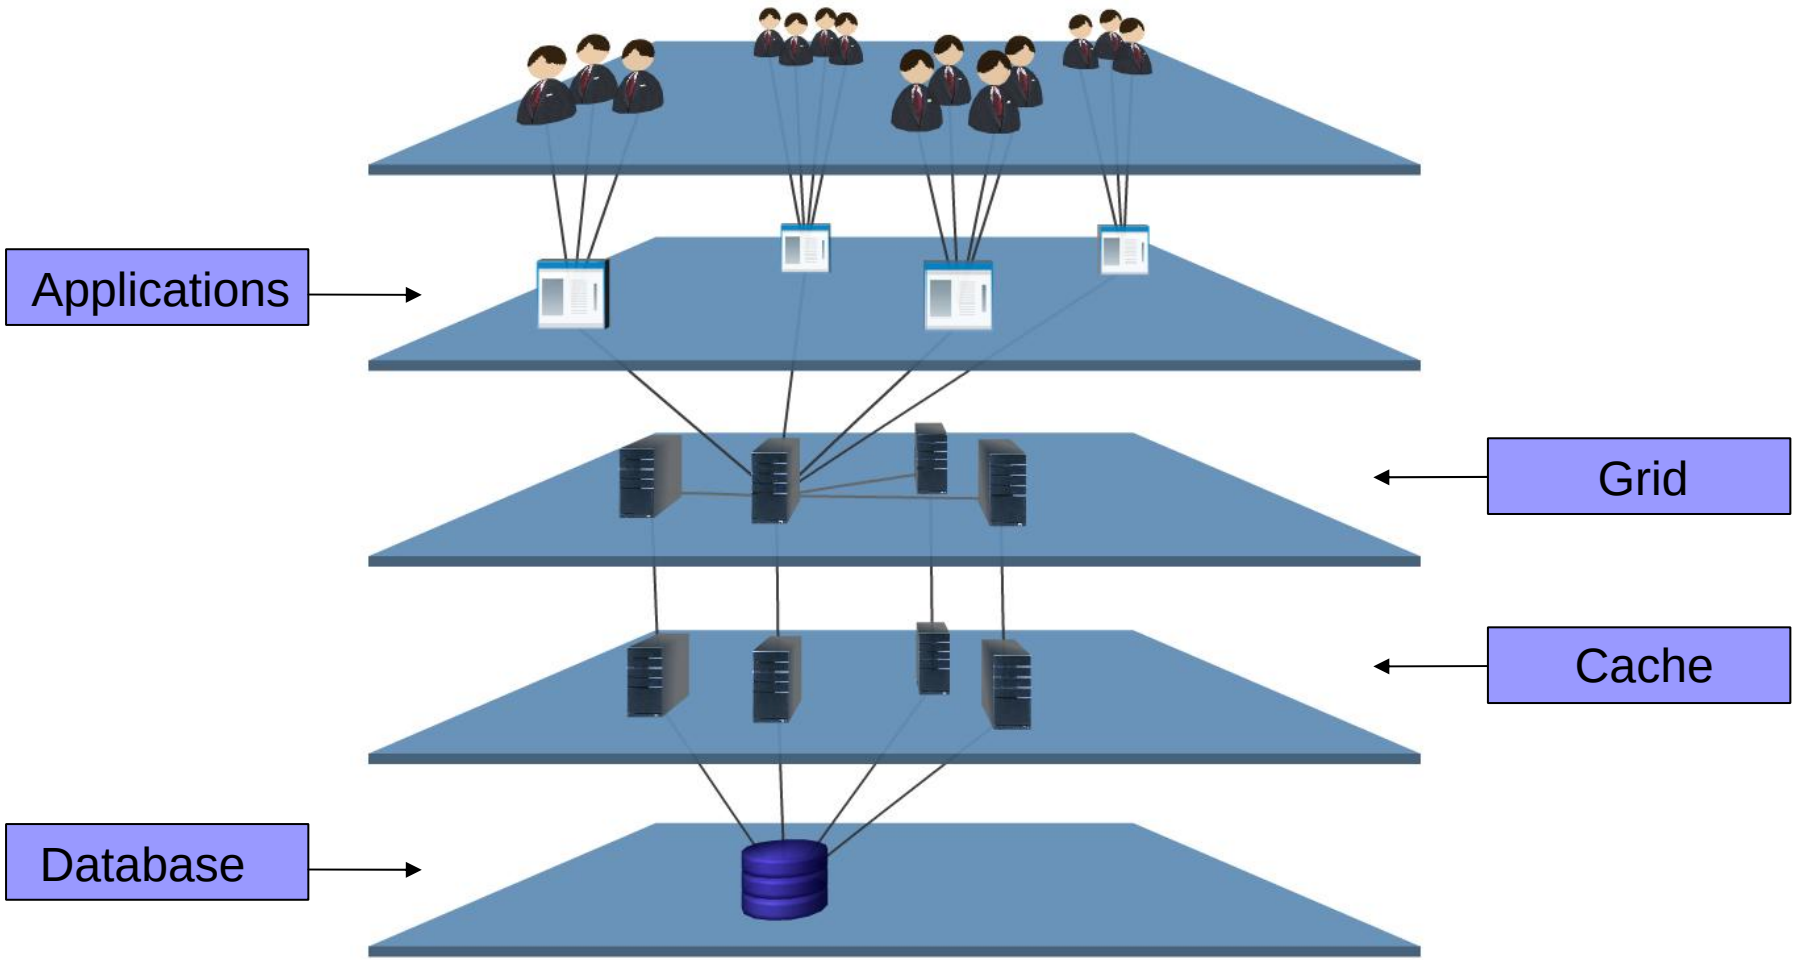

The Core: A scalable Infrastructure

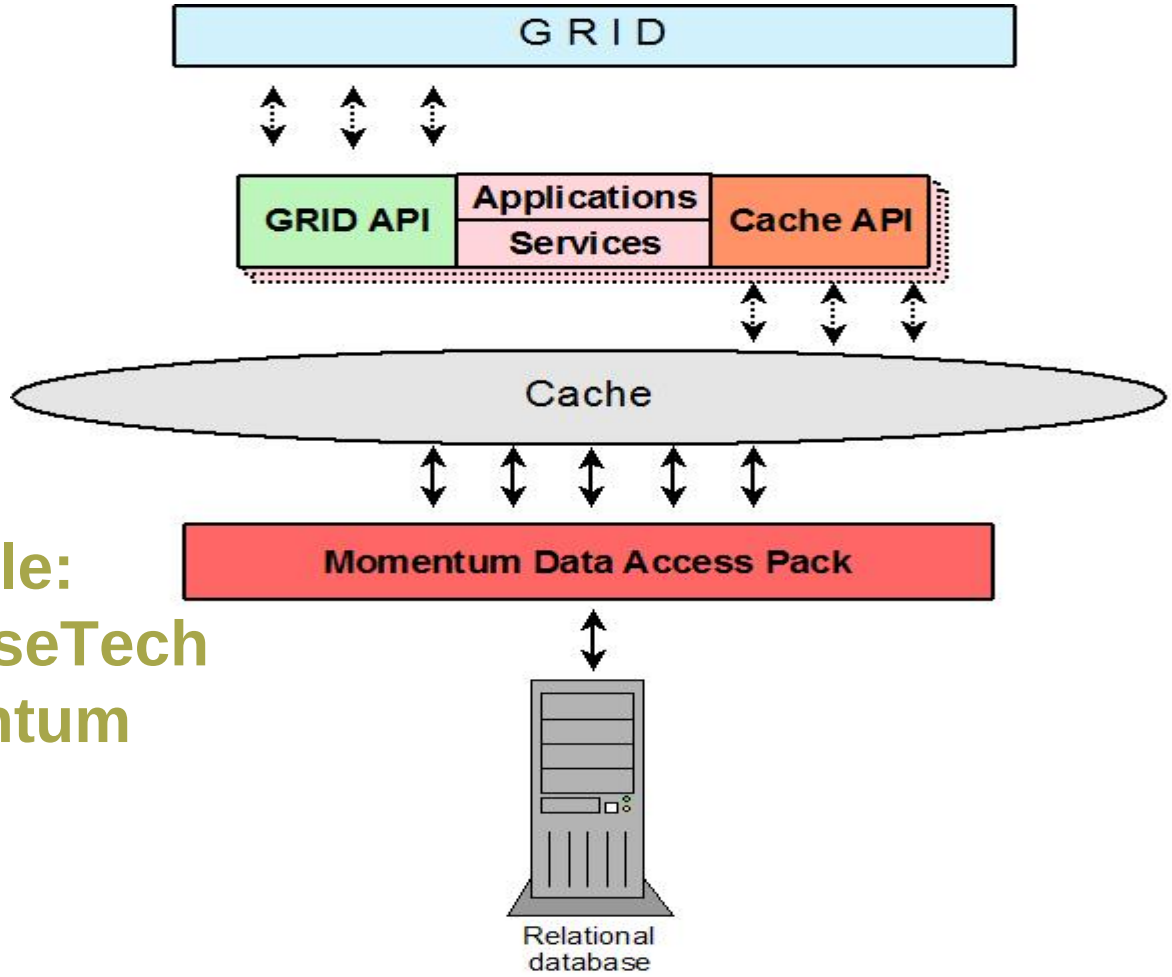

Example: GridwiseTech Momentum

Virtualization Everywhere

- Devices like desktop, laptop, minis, smart phones, and sensors, are islands of capabilities today.
- However, with an image (or clone) in the cloud
- … enabling deep content services not able on the device
- smart mobiles will become the primary (access) device for all communication and computing needs
- ISV's will provide a value-added service by mobilizing their application

Device-centric => Information-centric

RI-222919

ISV's "mobilize" their application

ISV's are providing value-added services
By developing light versions for smart phones

- Salesforce: Mobile Lite, CRM
- Omniture: SiteCatalyst, Web analysis
- Interchange: SalesNow
- Workday: HR and ERP

Business Model:

- attract new customers for their phone services
- attract new customers for their Web services

Your Digital Shadow

All your devices will have a clone
in the cloud
Sitting in a secure container
Accompanying you anywhere
As your Lifelog

or

- Your digital shadow in the cloud -

More about DEISA in:

Thank You!

Gentzsch@rzg.mpg.de

RI-222919