ICCGI 2009
Cannes, August 23-27, 2009

Computing & Data Service Infrastructures for the Global Information Age

Internet, Web, Grids, Clouds, and Telecom, and how do they come together

Wolfgang Gentzsch
The DEISA Project & Board of Directors of OGF
gentzsch at rzg.mpg.de

Content

- Service Infrastructures
  - Water, Power, Roads, Grids
- Components:
  - HPC Centers, Grids, Clouds, Internet, Web,...
- Example:
  - The DEISA Ecosystem for HPC Applications
- Next-Generation e-Infrastructure:
  - Service Oriented Enterprise
  - Digital City

Service Infrastructures, nothing new

Distributed European Infrastructure for Supercomputing Applications

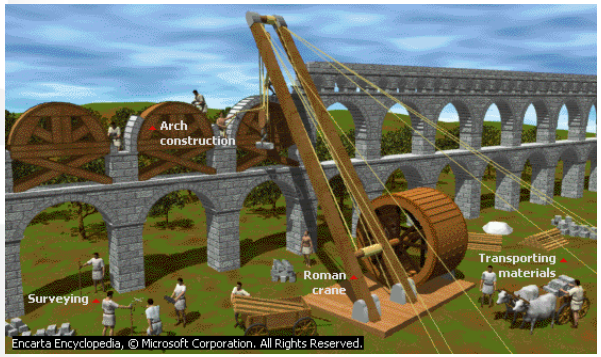

Ancient Rome: 10 aqueducts, 150,000 m^3 water each day

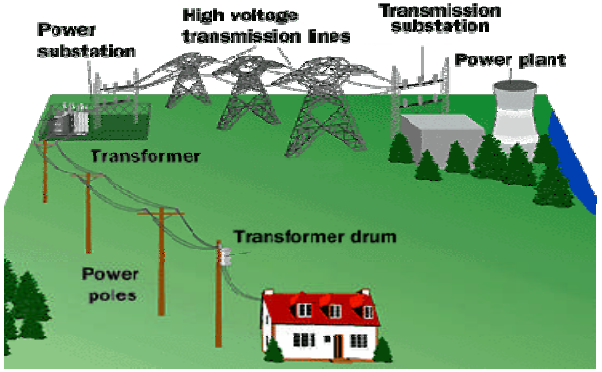

Electrical Power Grid Infrastructure

Transportation Grids

EGEE – Enabling Grid in E-SciencE

Requirements for an e-Infrastructure

Distributed European Infrastructure for Supercomputing Applications

- Transparent
- Secure
- Scalable
- Fast, at your finger tip
- Inexpensive, pay-per-user, ...
- ... with access via Internet !

Distributed
European
Infrastructure for
Supercomputing
Applications

DEISA

Components of an e-Infrastructure:

Servers, Clusters, Grids and Clouds

Terminology

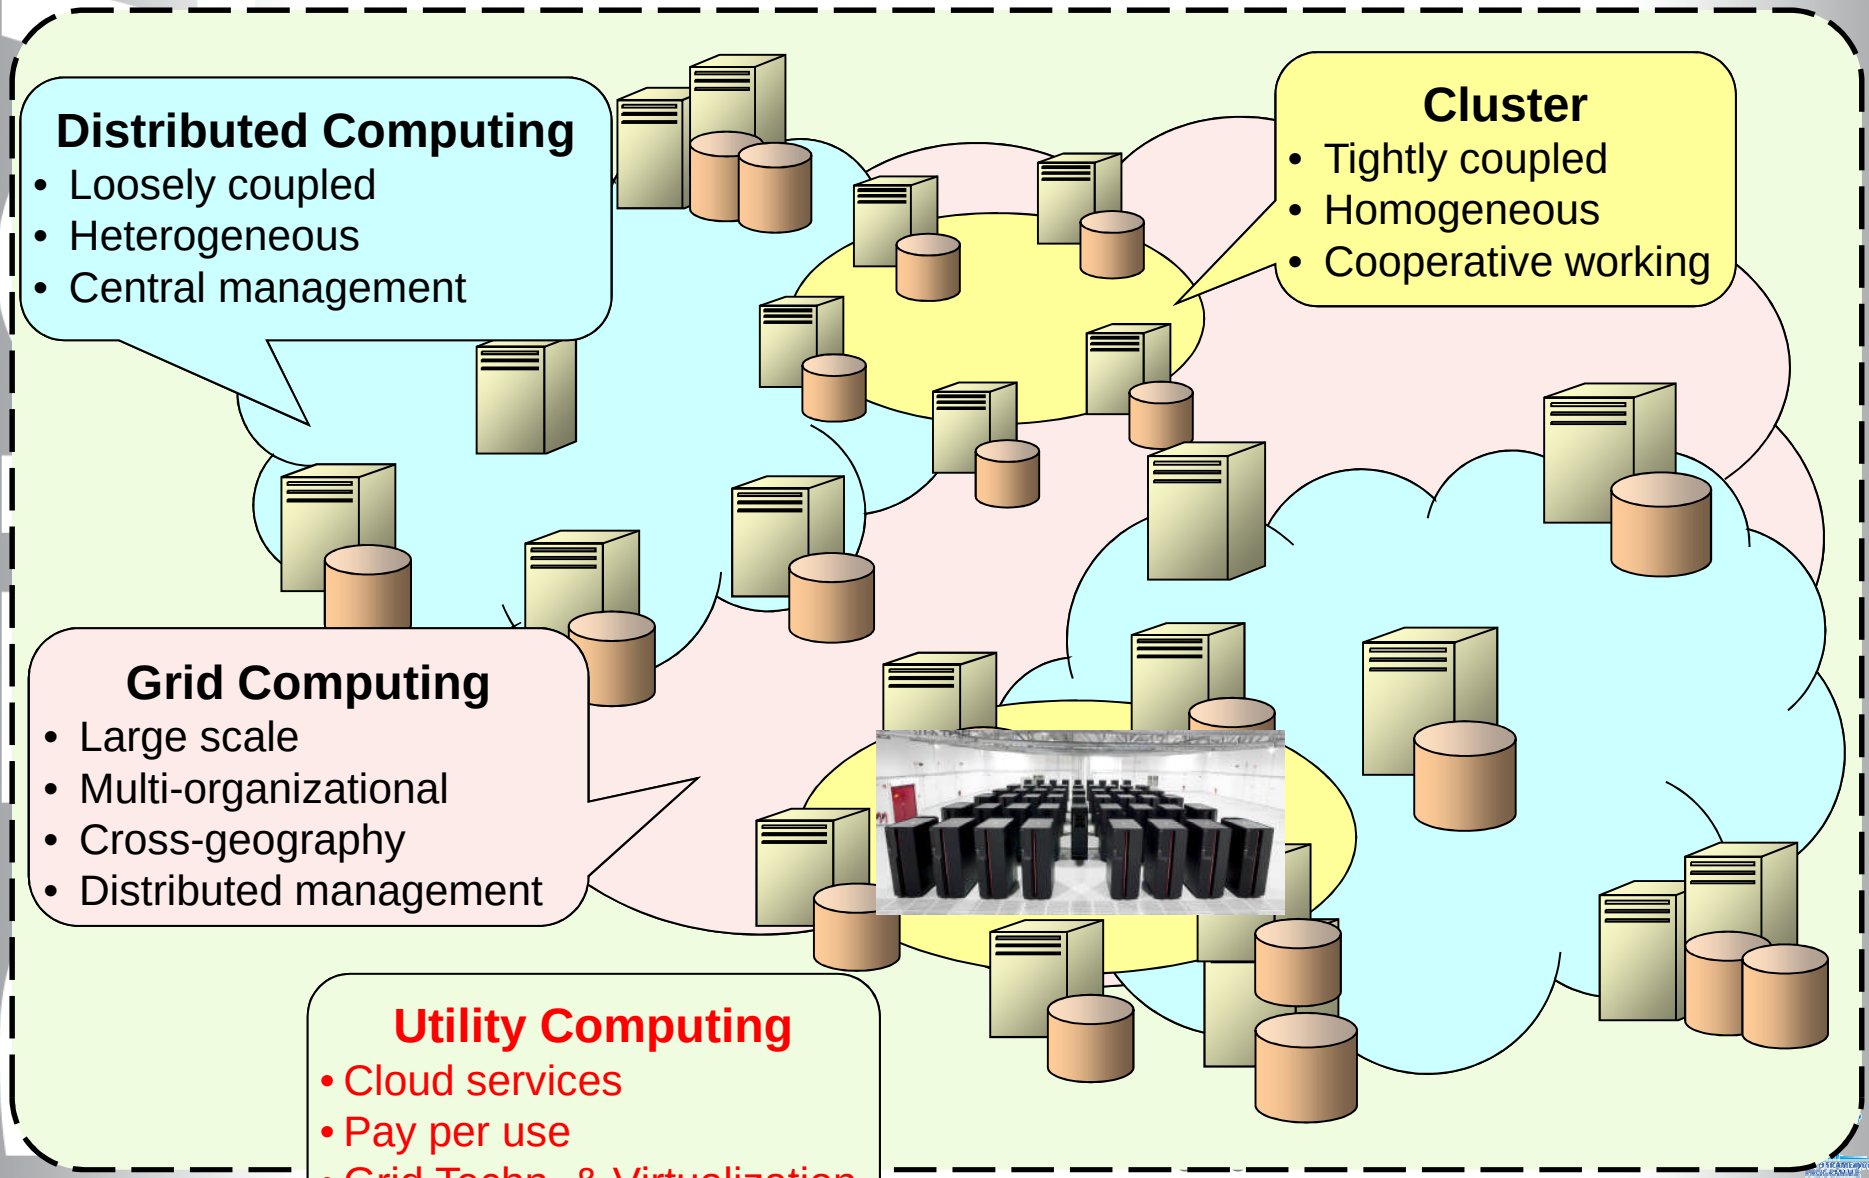

Clusters & HPC Centers



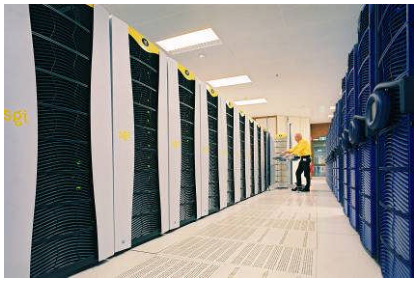

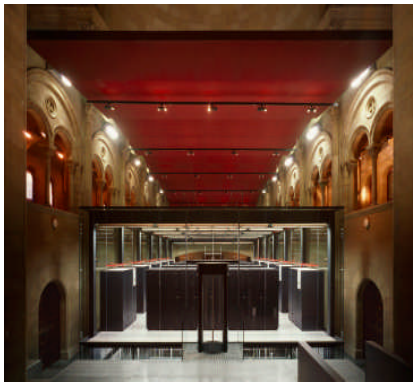


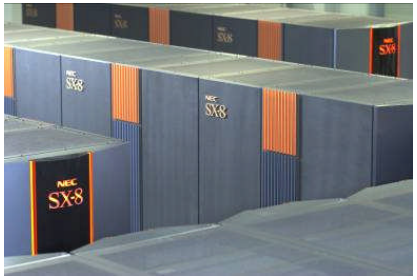


RI-222919

HPC Centers

Distributed European Infrastructure for Supercomputing Applications

- HPC Centers: **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

Grids

Distributed European Infrastructure for Supercomputing Applications

1998: The Grid: Blueprint for a New Computing 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

Grids *)

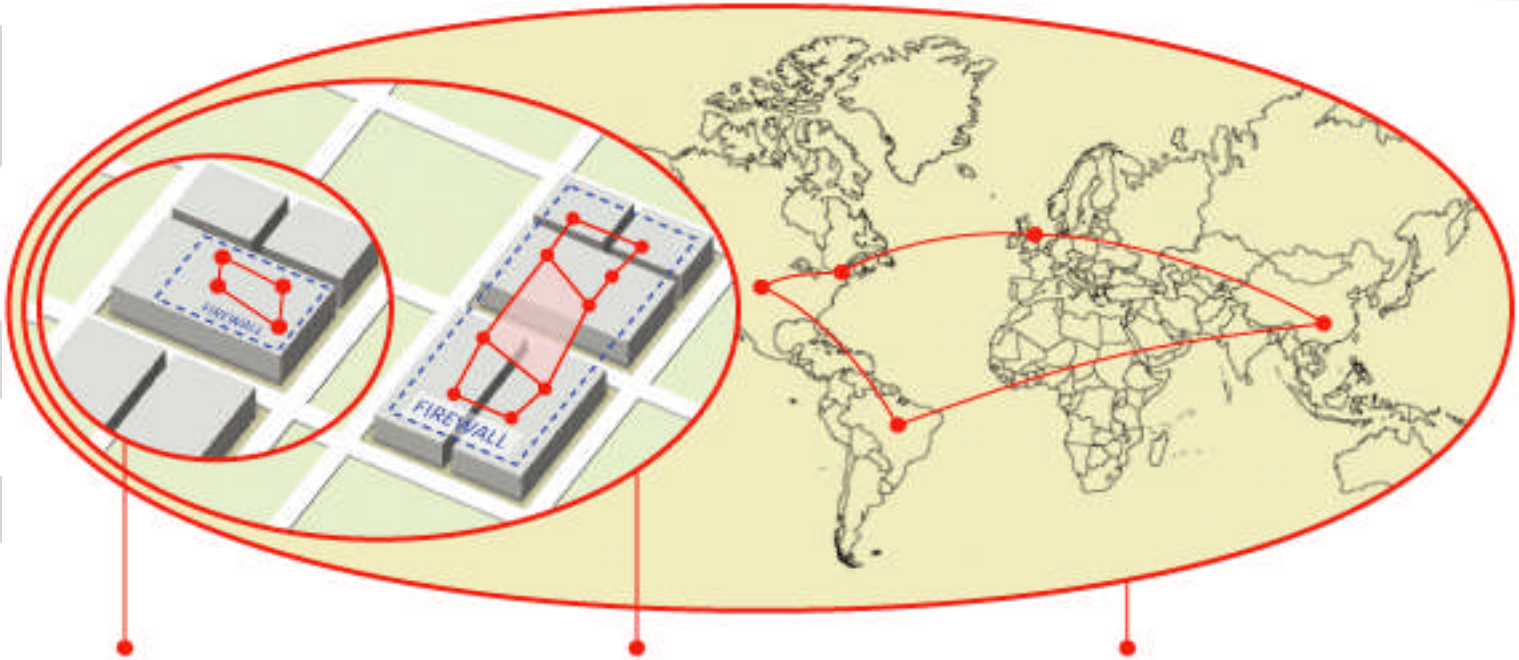

Departmental Grids **Enterprise Grids** **Global Grids**

*) Sun, 2001

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

Example: ANEKA Cloud platform

Distributed European Infrastructure for Supercomputing Applications

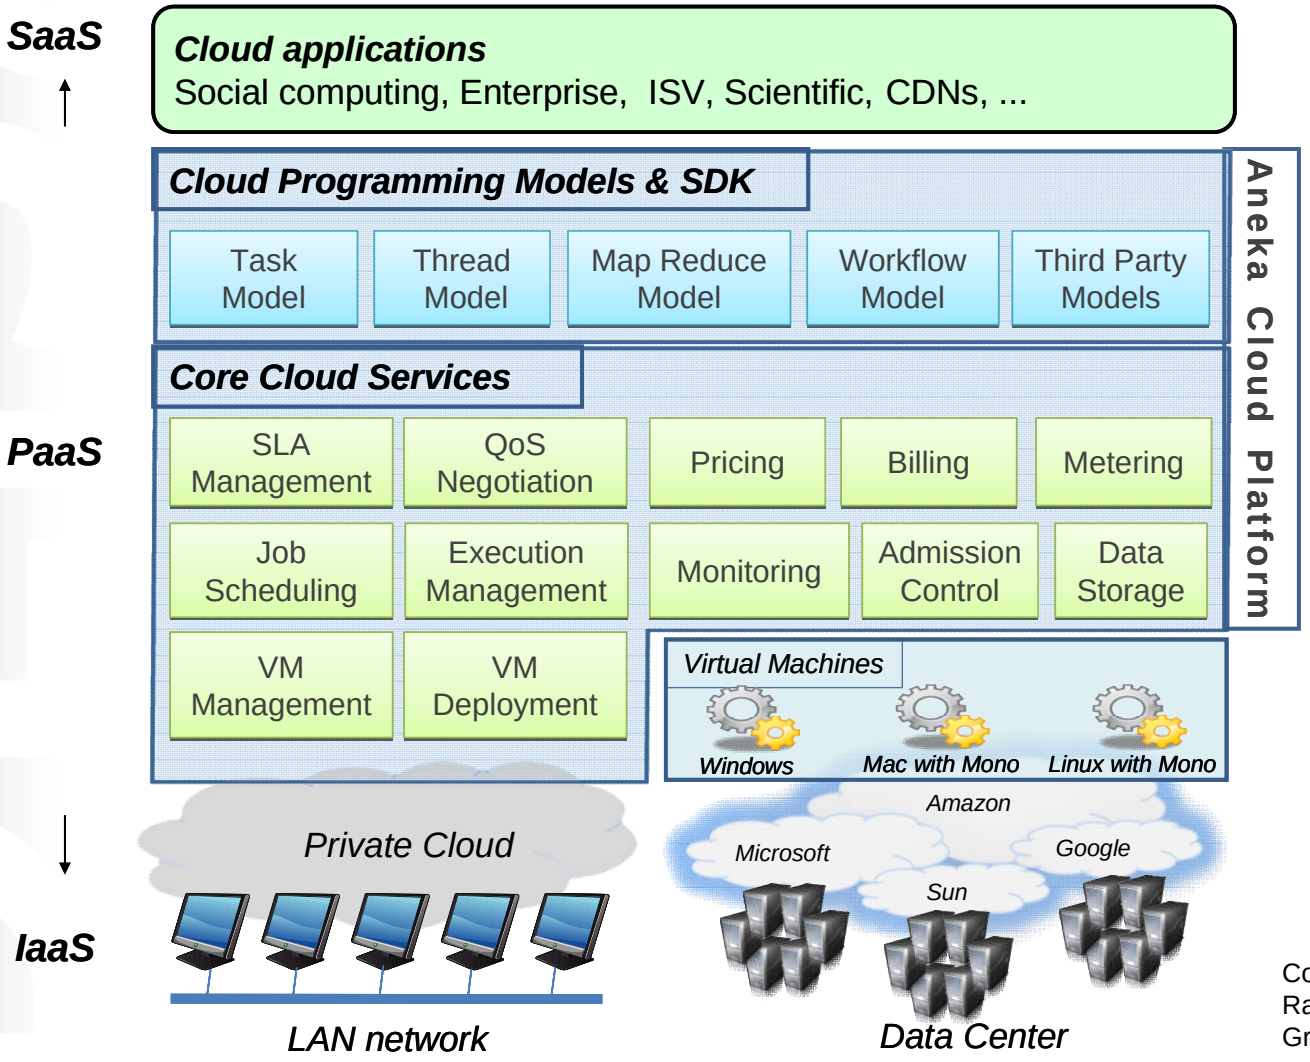

Courtesy:
Raj Buyya
Grids Lab

From Enterprise Grid to Data Cloud: Standard Grid Concepts

(often not scalable)

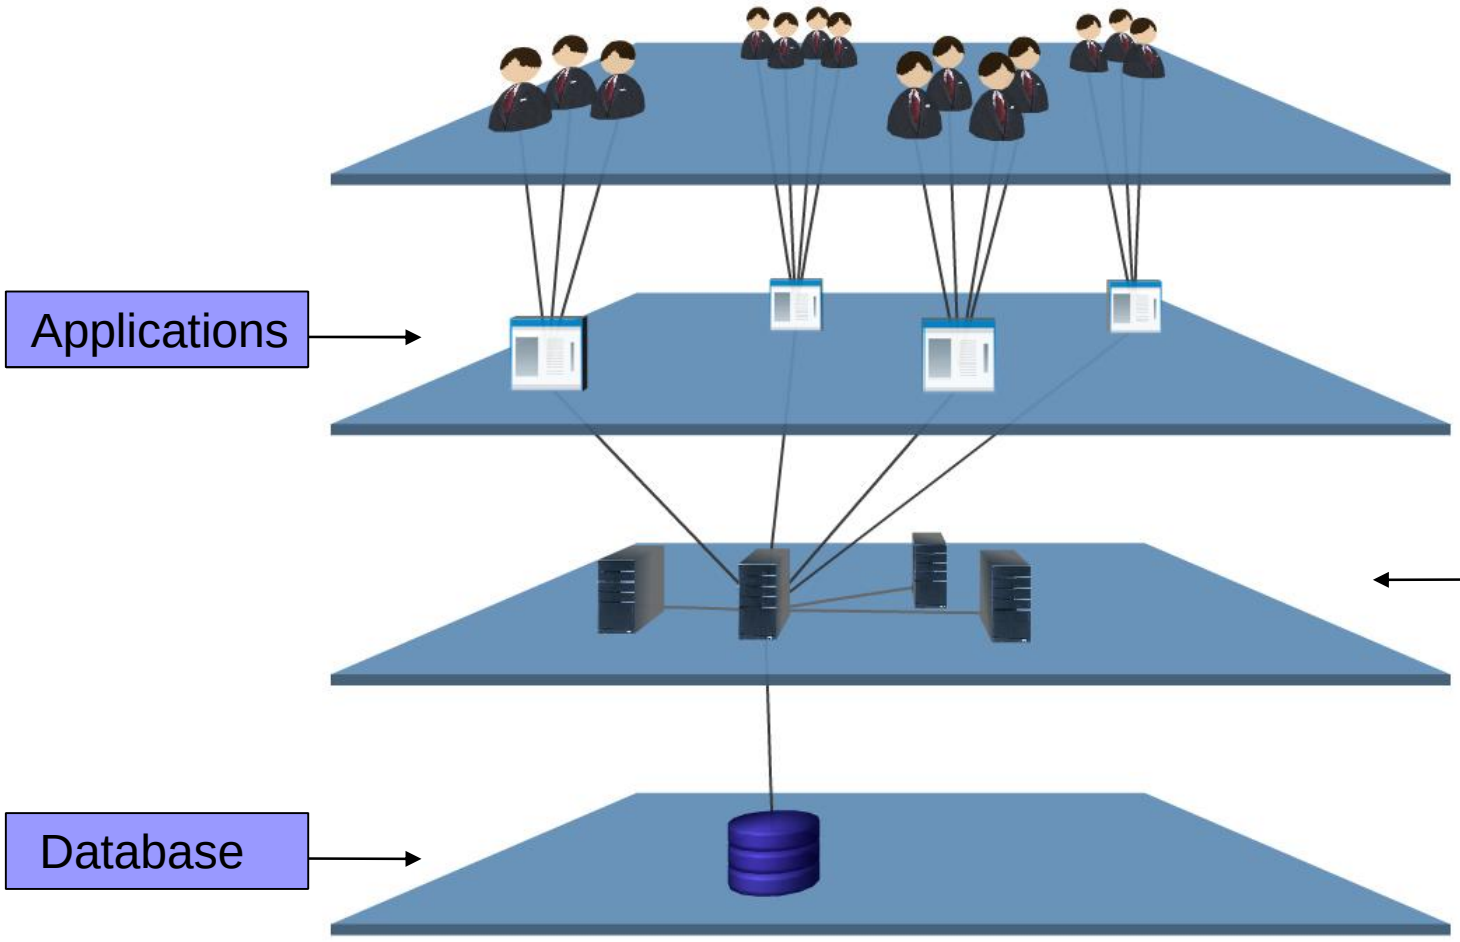

Courtesy: GridwiseTech

From Enterprise Grid to Data Cloud: Cache and Grid concept

Distributed European Infrastructure for Supercomputing Applications

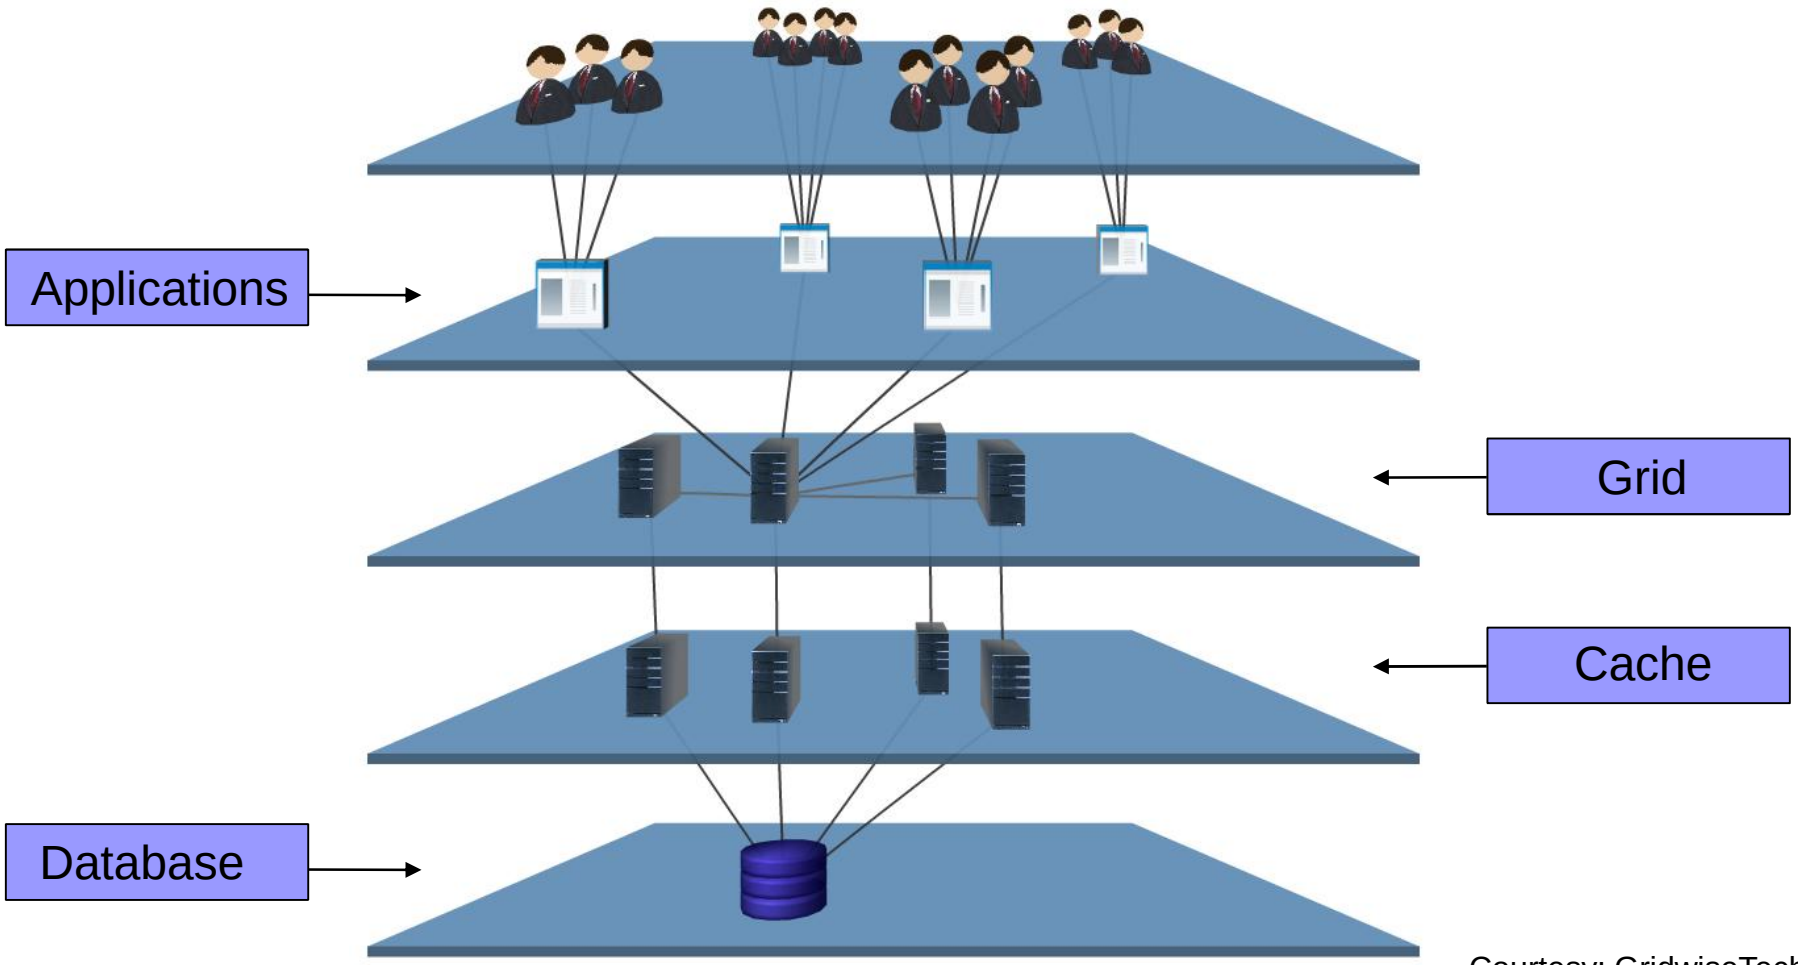

Courtesy: GridwiseTech

From Enterprise Grid to Data Cloud: Architecture

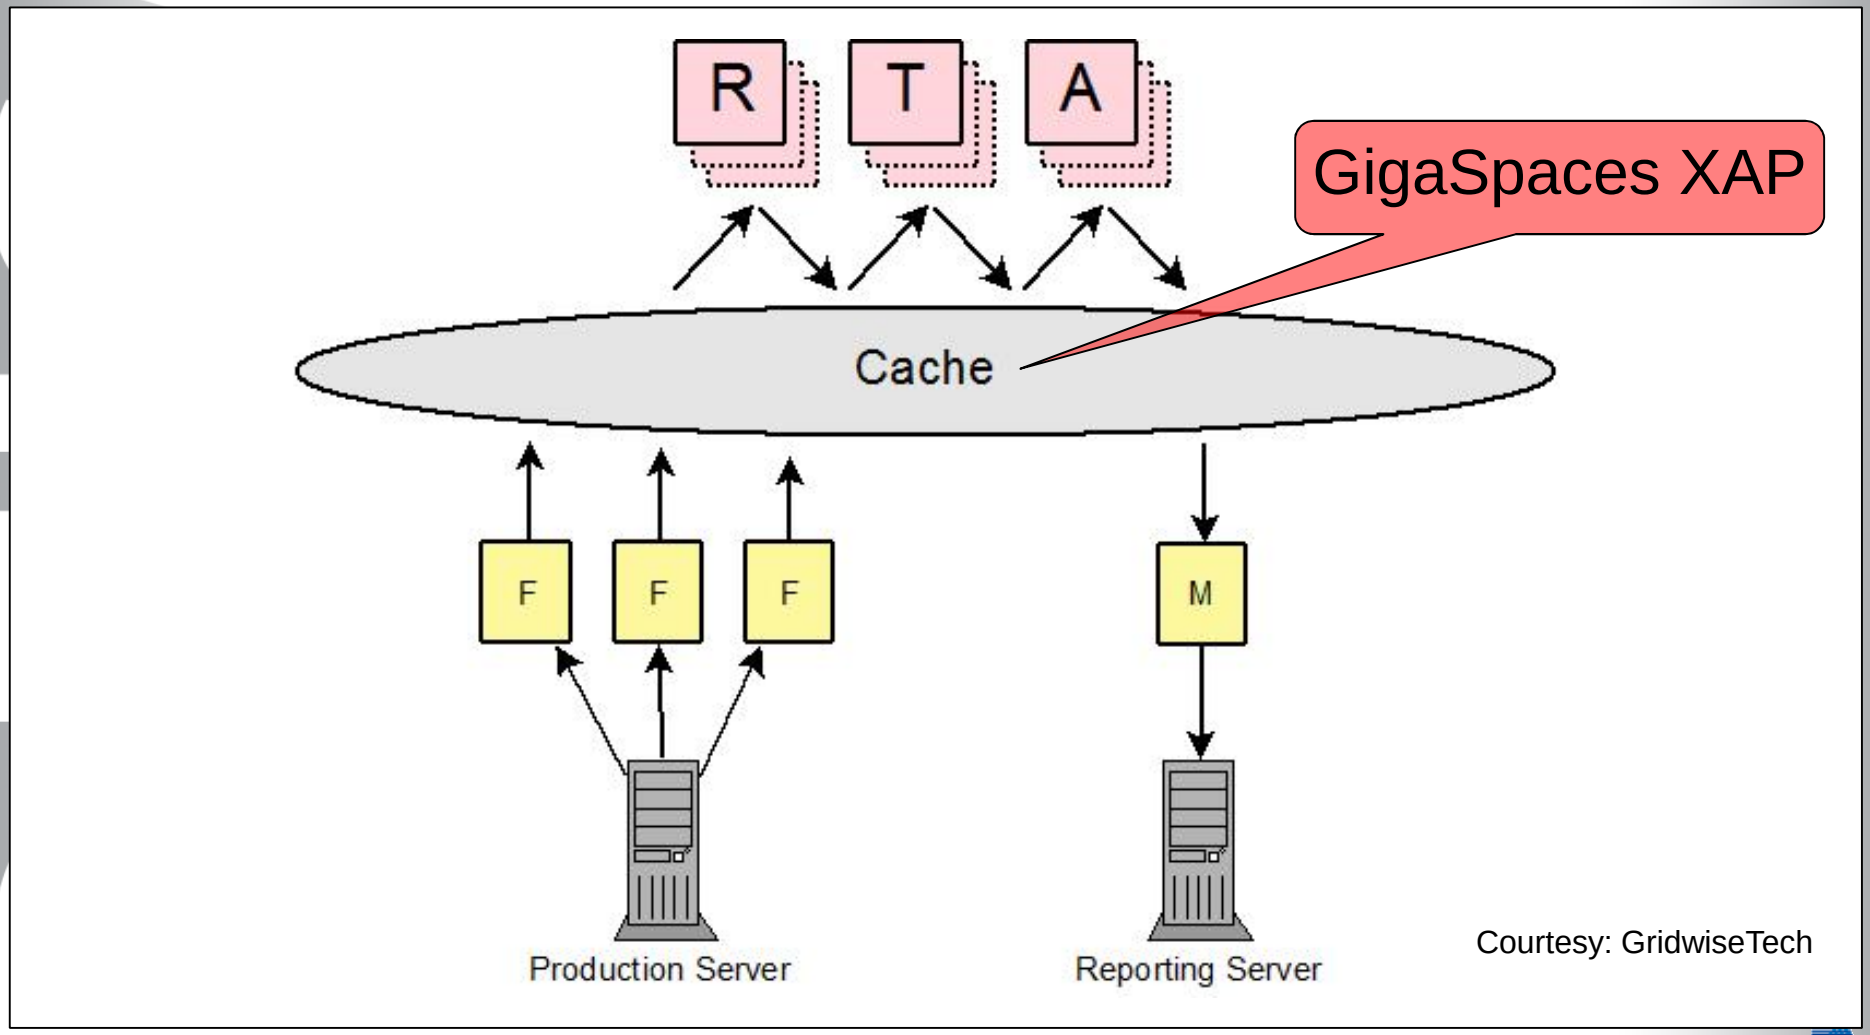

From Enterprise Grid to Data Cloud: a scalable Infrastructure

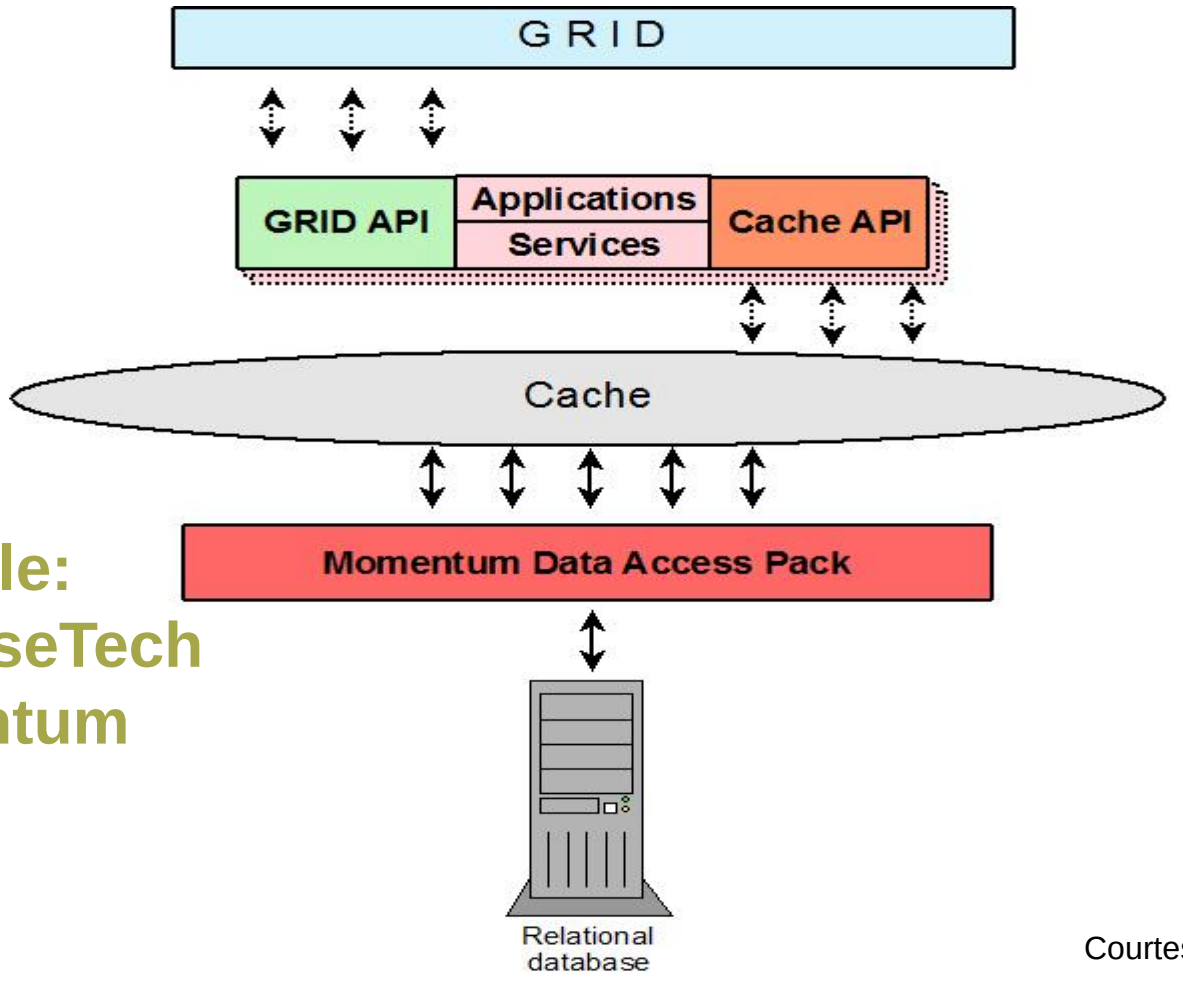

Example: GridwiseTech Momentum

Courtesy: GridwiseTech

Distributed
European
Infrastructure for
Supercomputing
Applications

Example of an e-Infrastructure:

The DEISA Ecosystem for HPC Grand-Challenge Applications

Distributed European Infrastructure for Supercomputing Applications

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.

new "petaflop" supercomputers

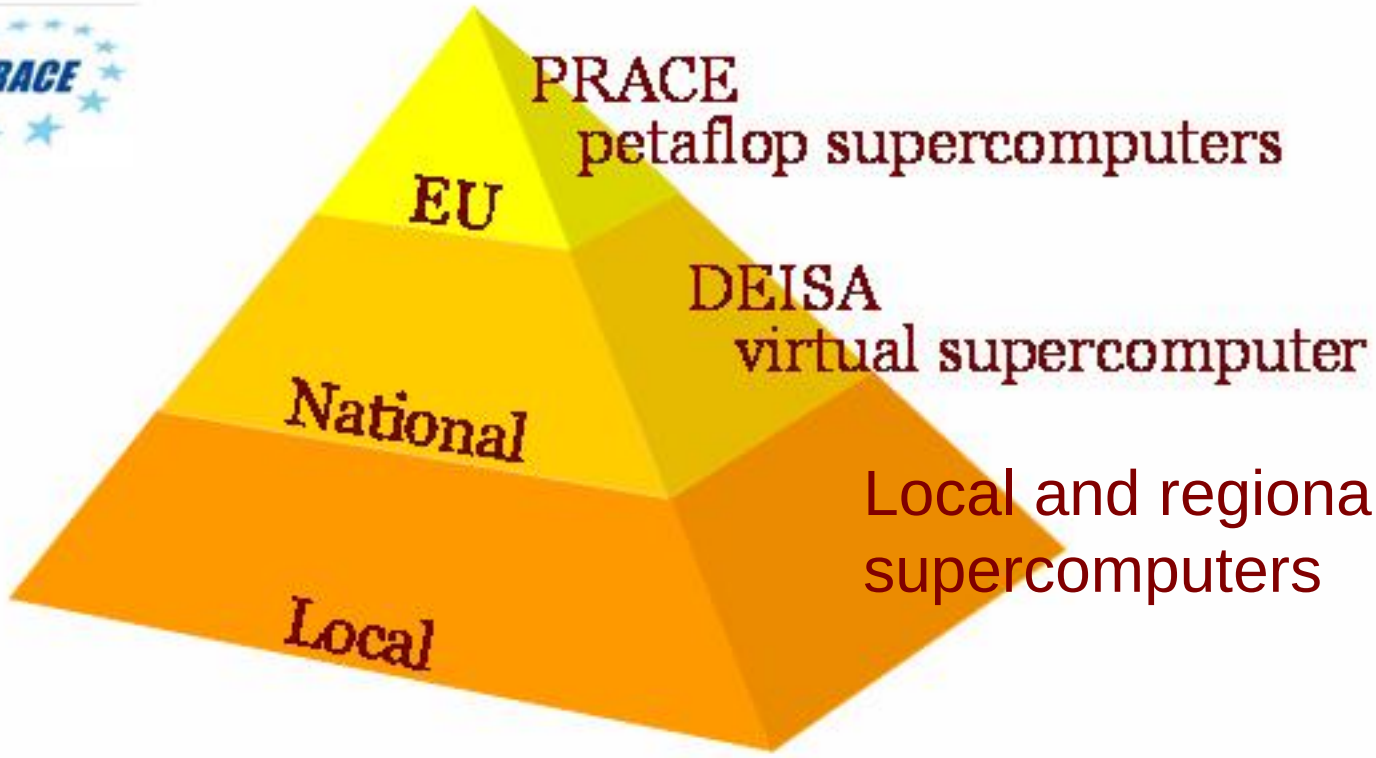

Local and regional supercomputers

Mario Campolargo
European Commission
OGF23, June 2008

DEISA Partners

Distributed European Infrastructure for Supercomputing Applications

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

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

SEVENTH FRAMEWORK PROGRAMME

Categories of DEISA services

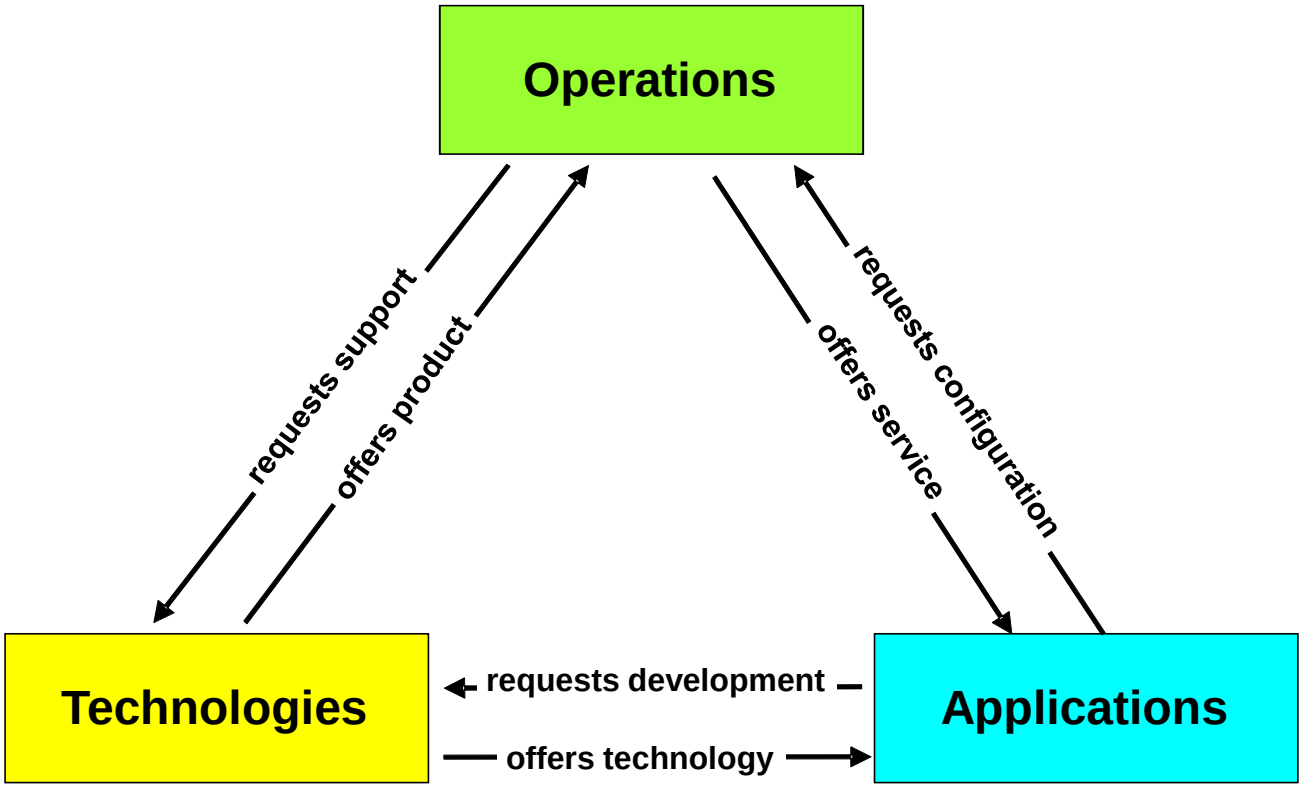

RI-222919

DEISA Service Layers

Distributed European Infrastructure for Supercomputing Applications

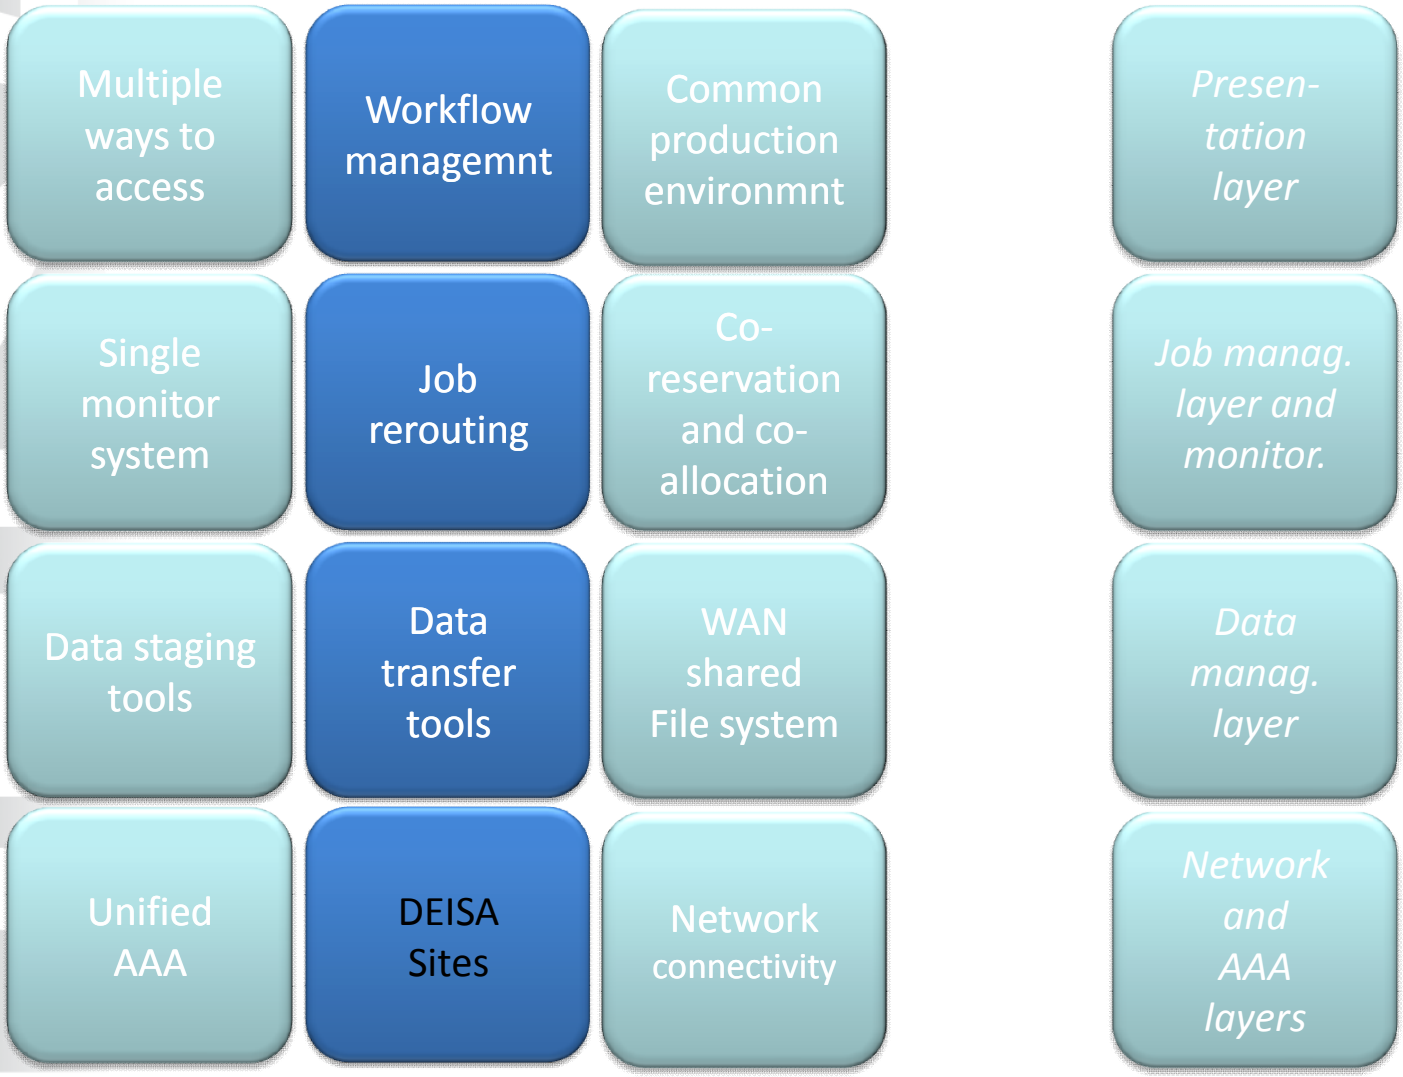

DEISA UNICORE Infrastructure

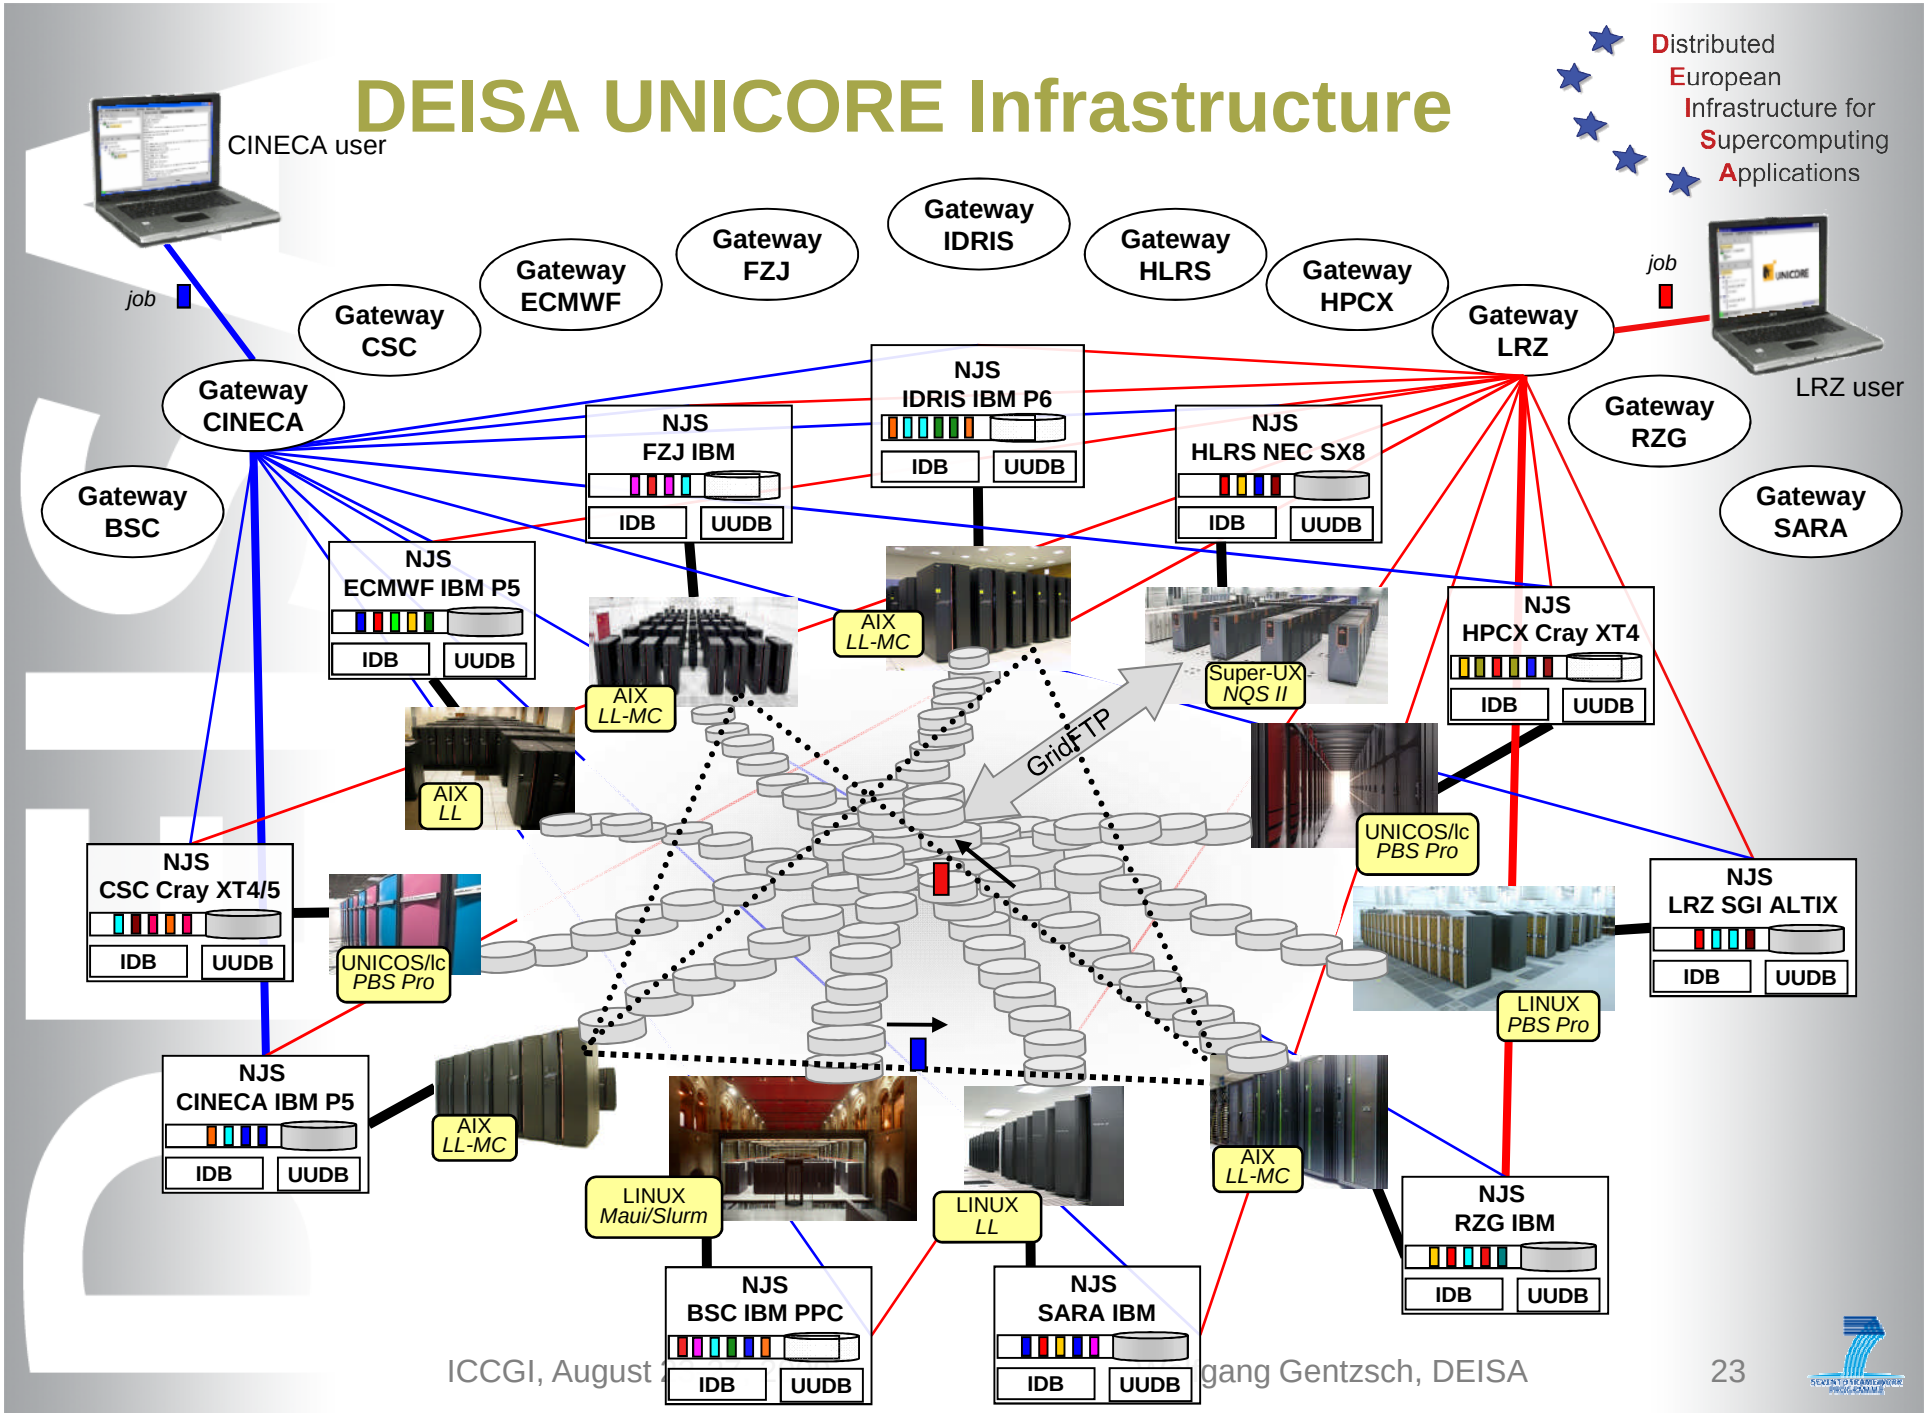

DEISA Global File System

Distributed European Infrastructure for Supercomputing Applications

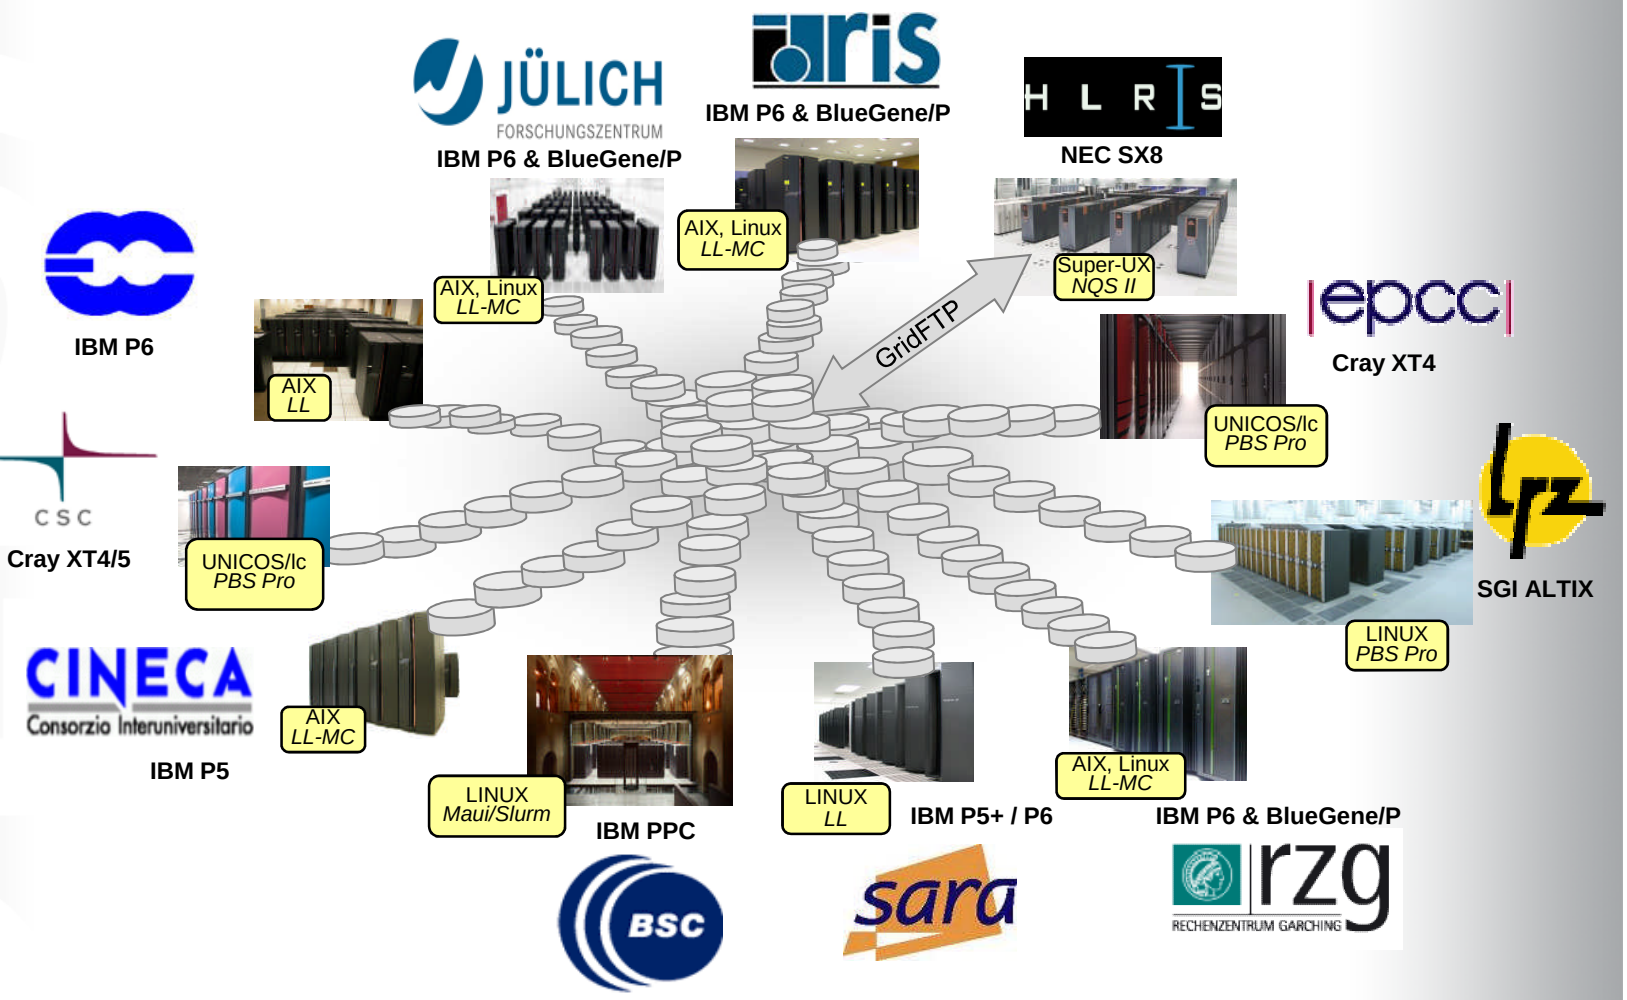

Global transparent file system based on the Multi-Cluster General Parallel File System (MC-GPFS of IBM)

Management of users in DEISA

Distributed European Infrastructure for Supercomputing Applications

- 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

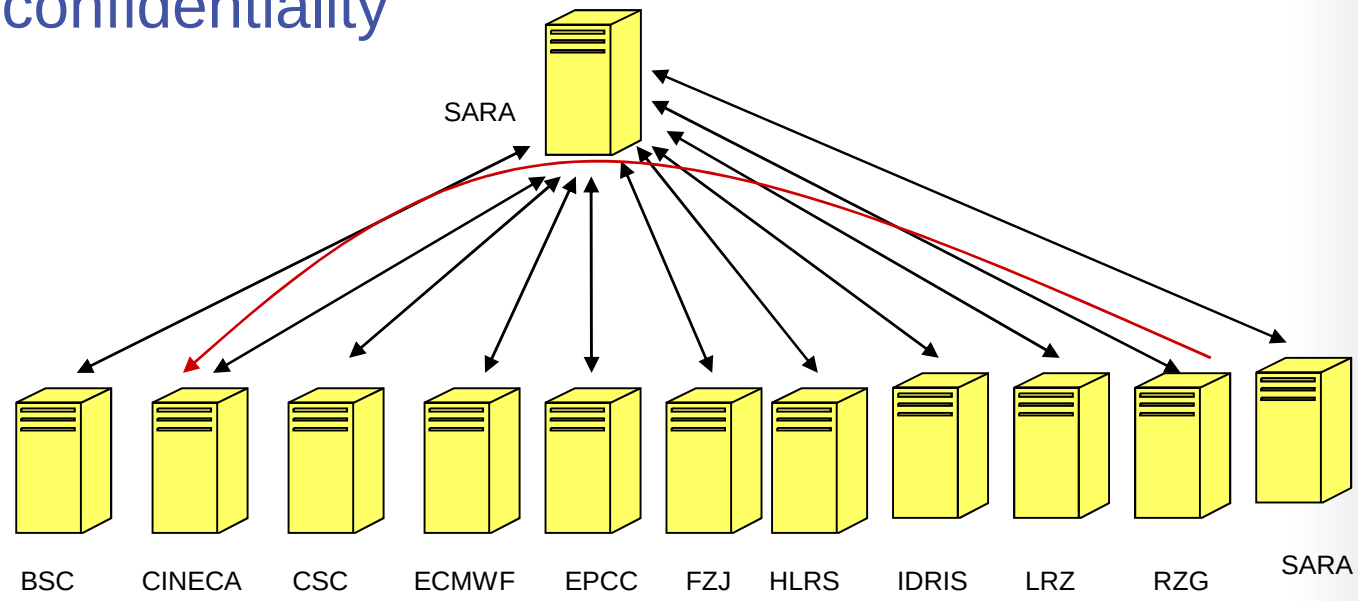

Common User Administration

- 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!

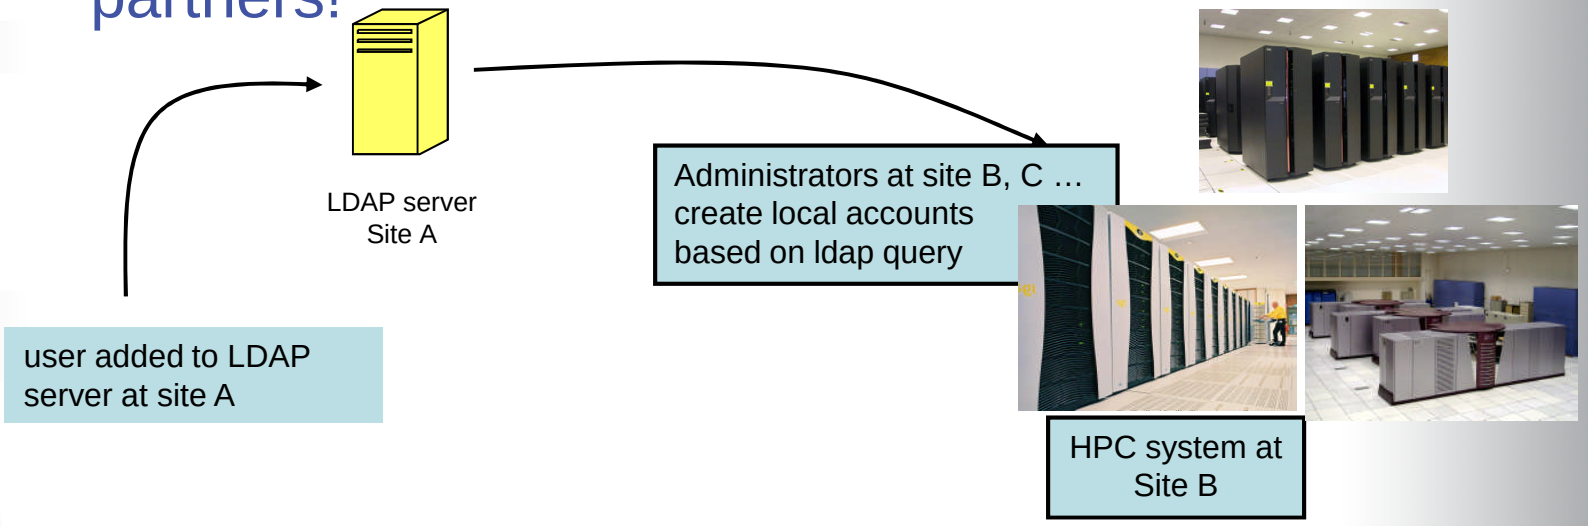

Life Sciences in DEISA

Distributed European Infrastructure for Supercomputing Applications

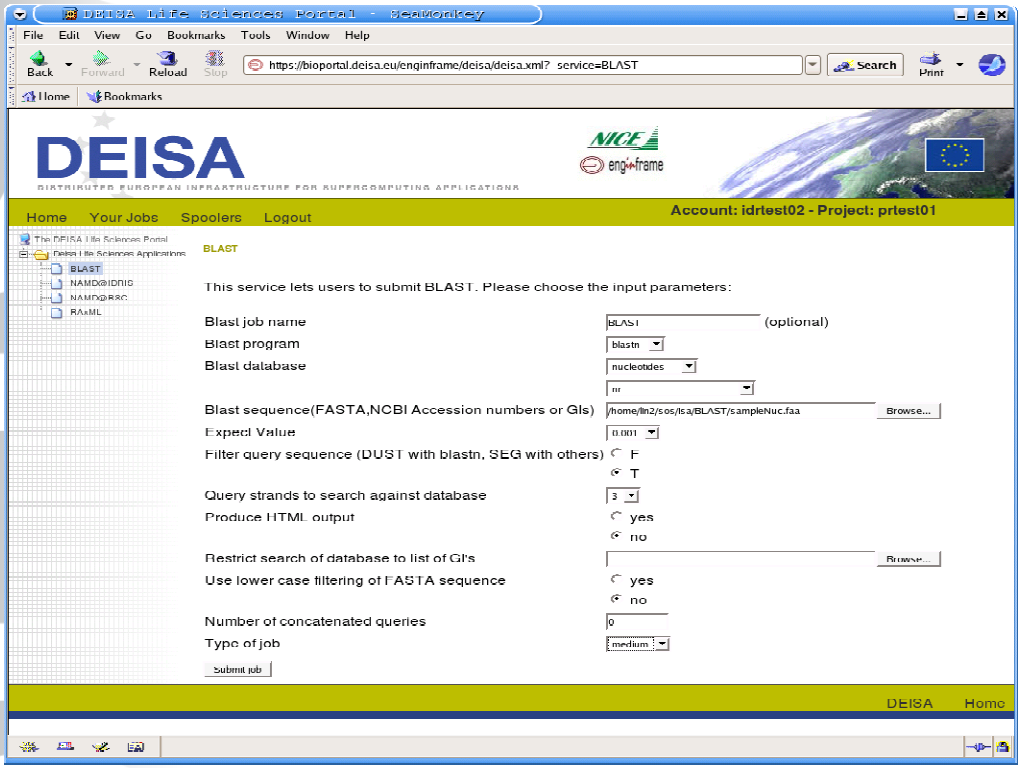

DEISA Life Science Portal based on NICE / EnginFrame

A Virtual Community

Promoting parallel apps in the life science community

Running big simulations on DEISA infrastructure that couldn't be done locally

Providing ease of access to resources

Application support for life science portal

NICE EnginFrame Cluster/Grid/Cloud Portal

Distributed European Infrastructure for Supercomputing Applications

Provides remote, interactive, transparent, and secure access to applications and data on your corporate Intranet or Internet, or in the Cloud.

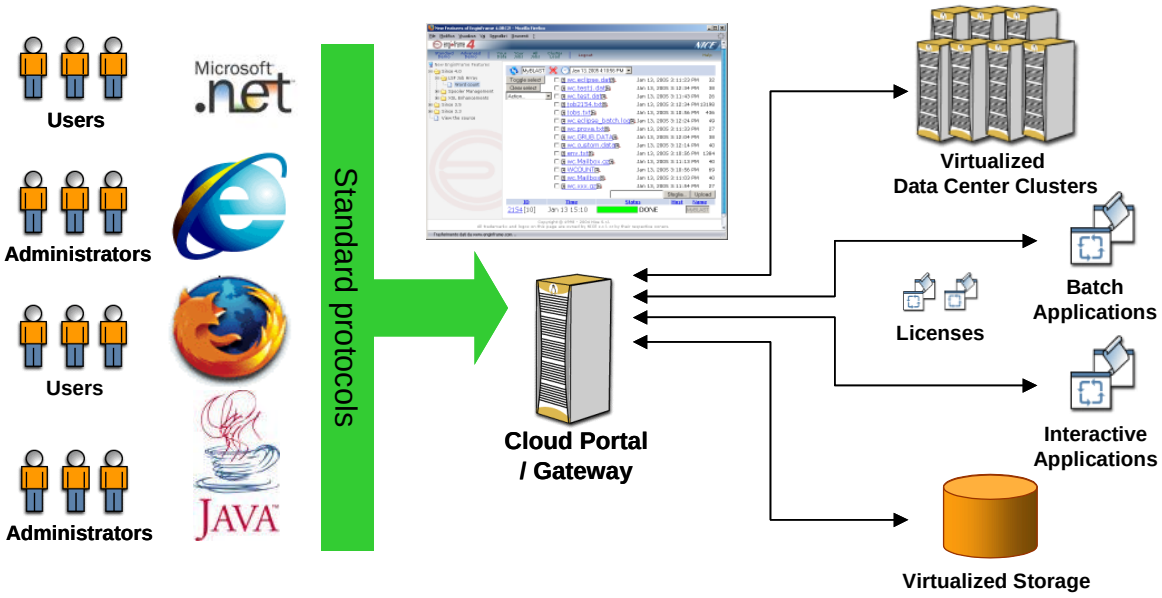

Users and administrators can access and control computing resources via an intuitive and standard Web interface virtually anywhere using a standard Web browser.

DEISA Extreme Computing Initiative (DECI)

Distributed European Infrastructure for Supercomputing Applications

- 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+)

Distributed
European
Infrastructure for
Supercomputing
Applications

Next-Generation e-Infrastructure

Convergence of Bandwidth, Clouds, and Mobile Devices

New Powerful End-User Devices

Distributed European Infrastructure for Supercomputing Applications

Carry-along PCs (CAPS) , Ultra-Mobile PCs (UMPC) , Cellphones

CAPC, from Samsung, South Korea

Fujitsu UMPC

OLPC from MIT

T83 Tablet from Asus, Taiwan, demoed at CeBit 2007

Asus R2H

CAPC from HTC

The Symbian OS

Distributed European Infrastructure for Supercomputing Applications

SYMBIAN FOUNDATION

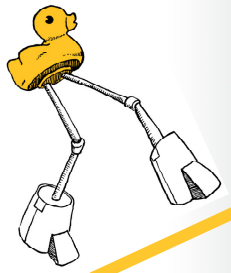

1998 — Symbian Ltd was founded by Ericsson, Motorola, Nokia and Psion to create Symbian OS

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 Community

Distributed European Infrastructure for Supercomputing Applications

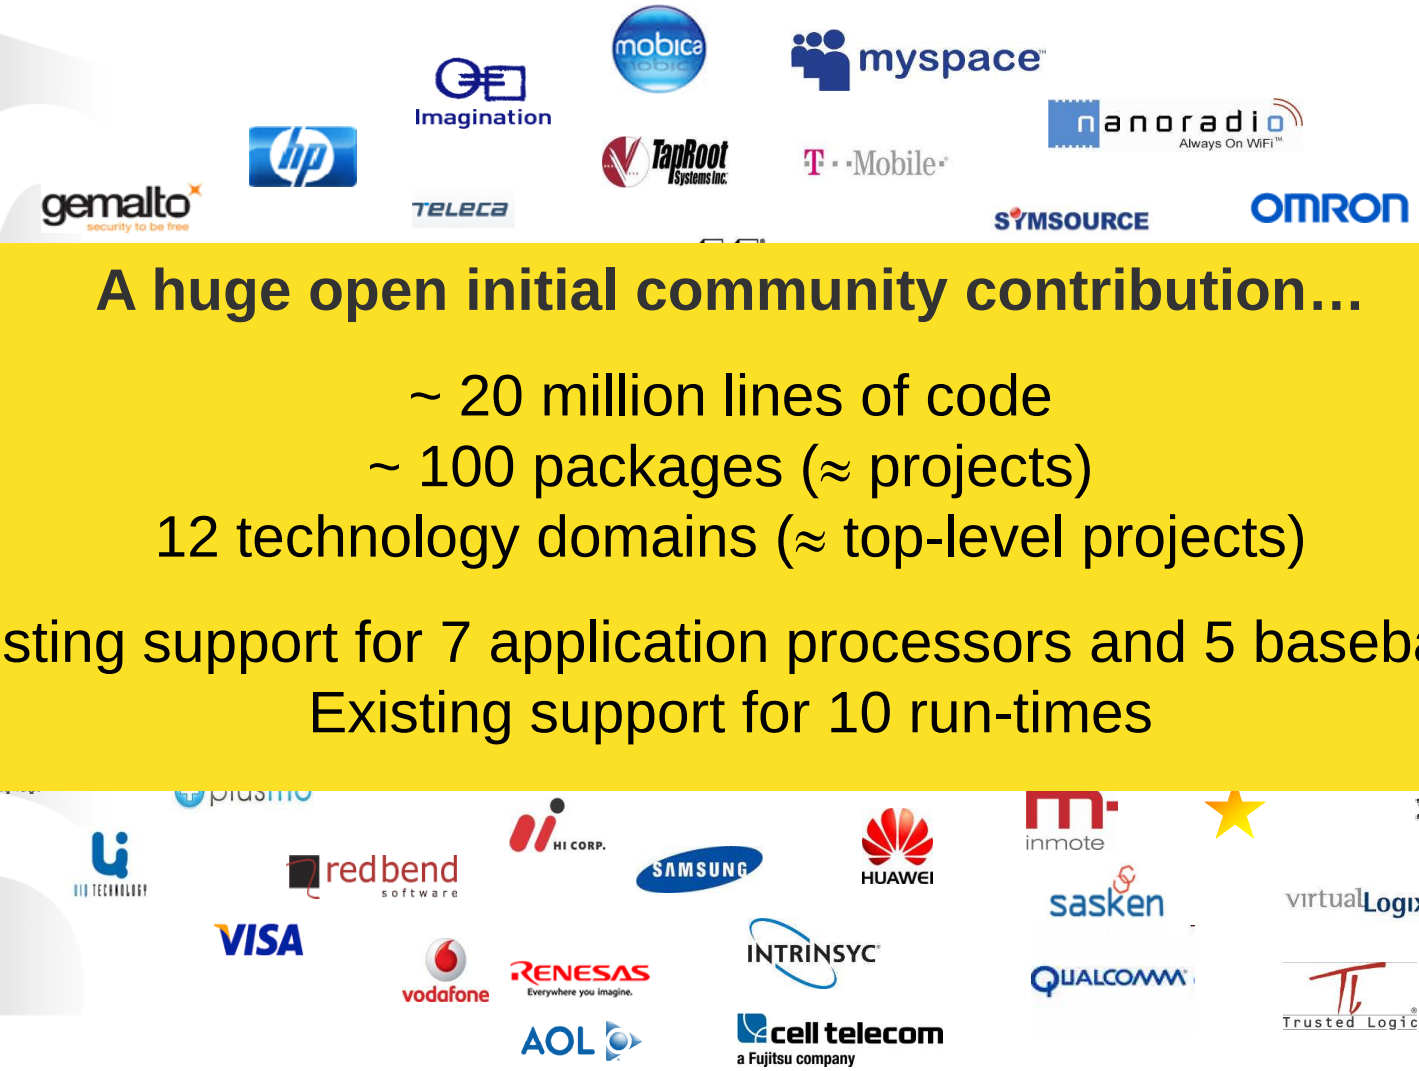

A huge open initial community contribution…

~ 20 million lines of code
~ 100 packages (≈ projects)
12 technology domains (≈ top-level projects)

Existing support for 7 application processors and 5 basebands
Existing support for 10 run-times

The Symbian Platform

(code on the device)

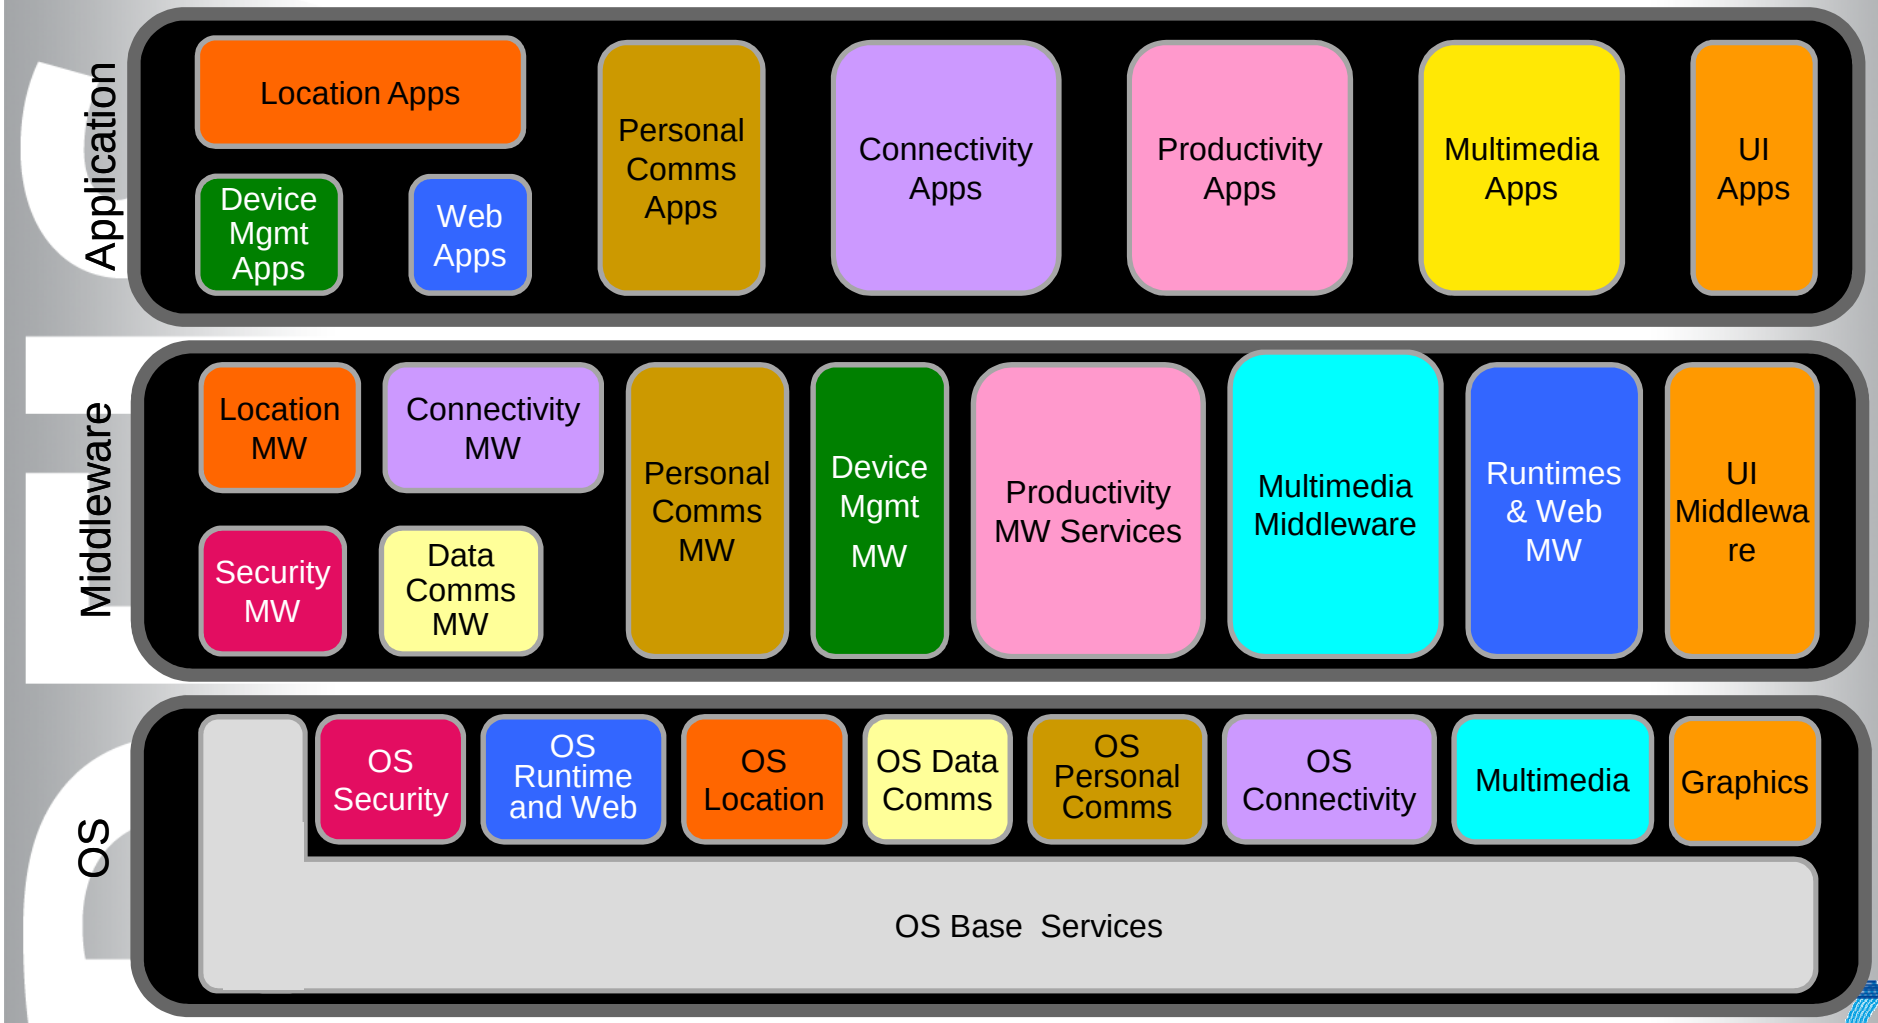

Vision:
The Service Oriented Enterprise

Distributed European Infrastructure for Supercomputing Applications

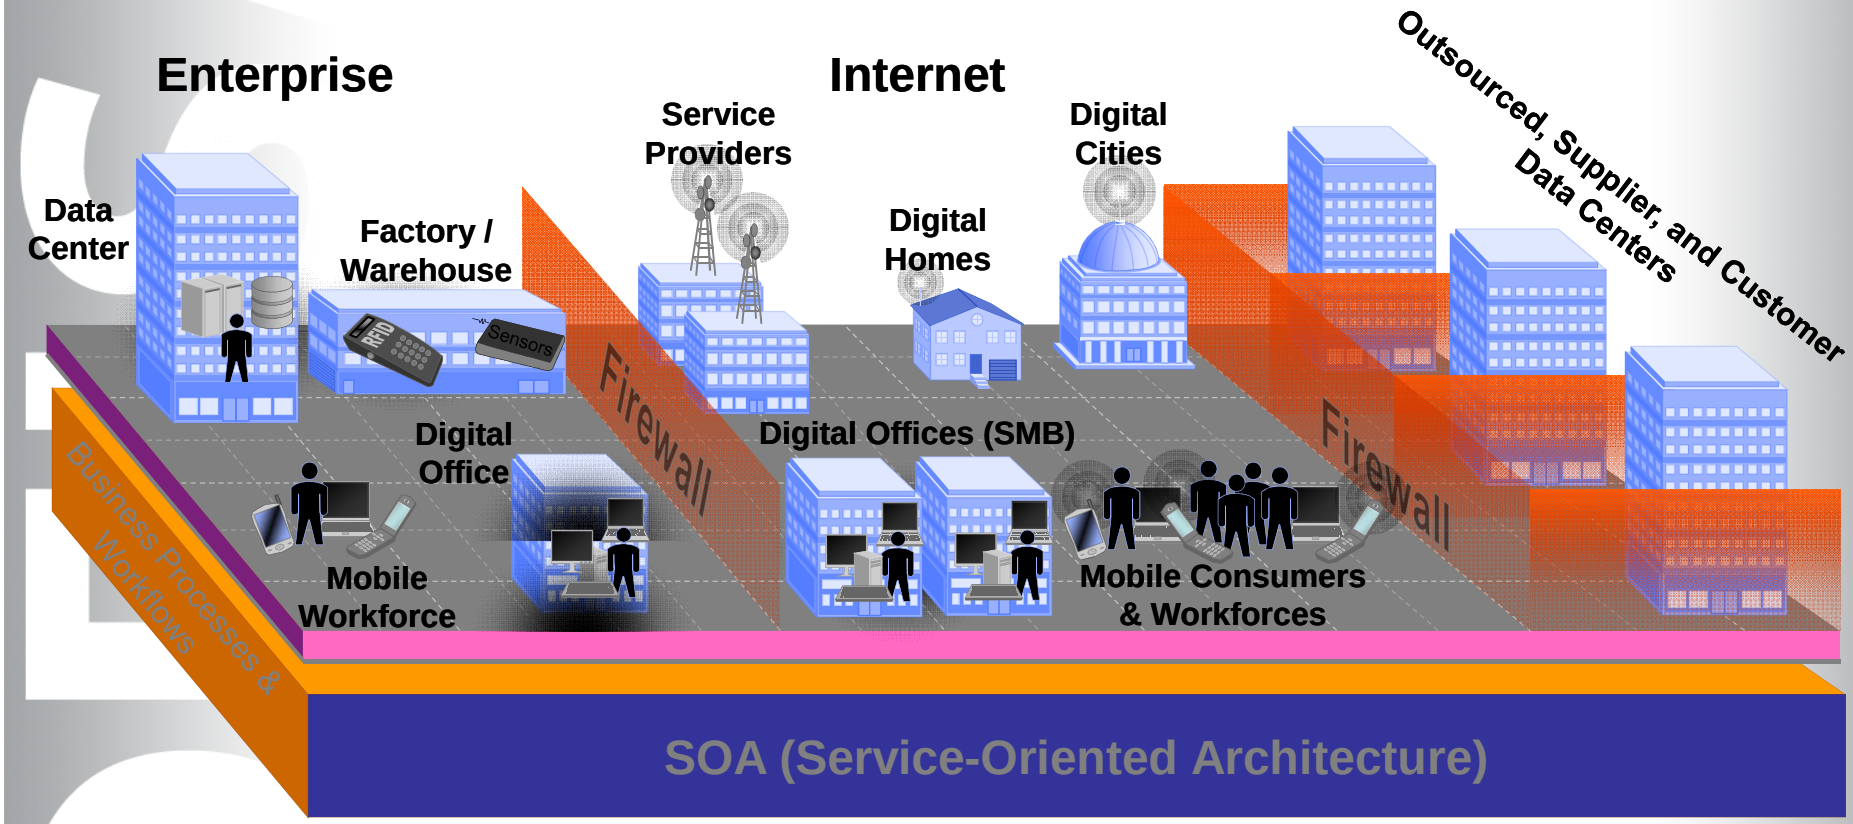

Courtesy: Robert Fogel, Intel

jlm3

Distributed European Infrastructure for Supercomputing Applications

A Peek at Intel's Digital City Vision

Courtesy Robert Fogel, Intel

Slide 37

jlm3 Does Intel own the images? If not, please replace.
jlmayerX, 3/1/2005

Facets of the Digital City

Distributed European Infrastructure for Supercomputing Applications

Serving Citizens

Digital Education

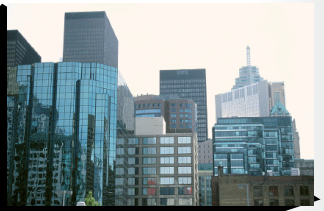

Economic Vitality (Digital Office)

Digital Govt GAPP Programs

Digital City

Digital Healthcare

Bridging the Digital Divide

Digital Home

Safety & Security

The Digital City is the Fabric that Connects the Community

Today's Digital Challenge

Distributed European Infrastructure for Supercomputing Applications

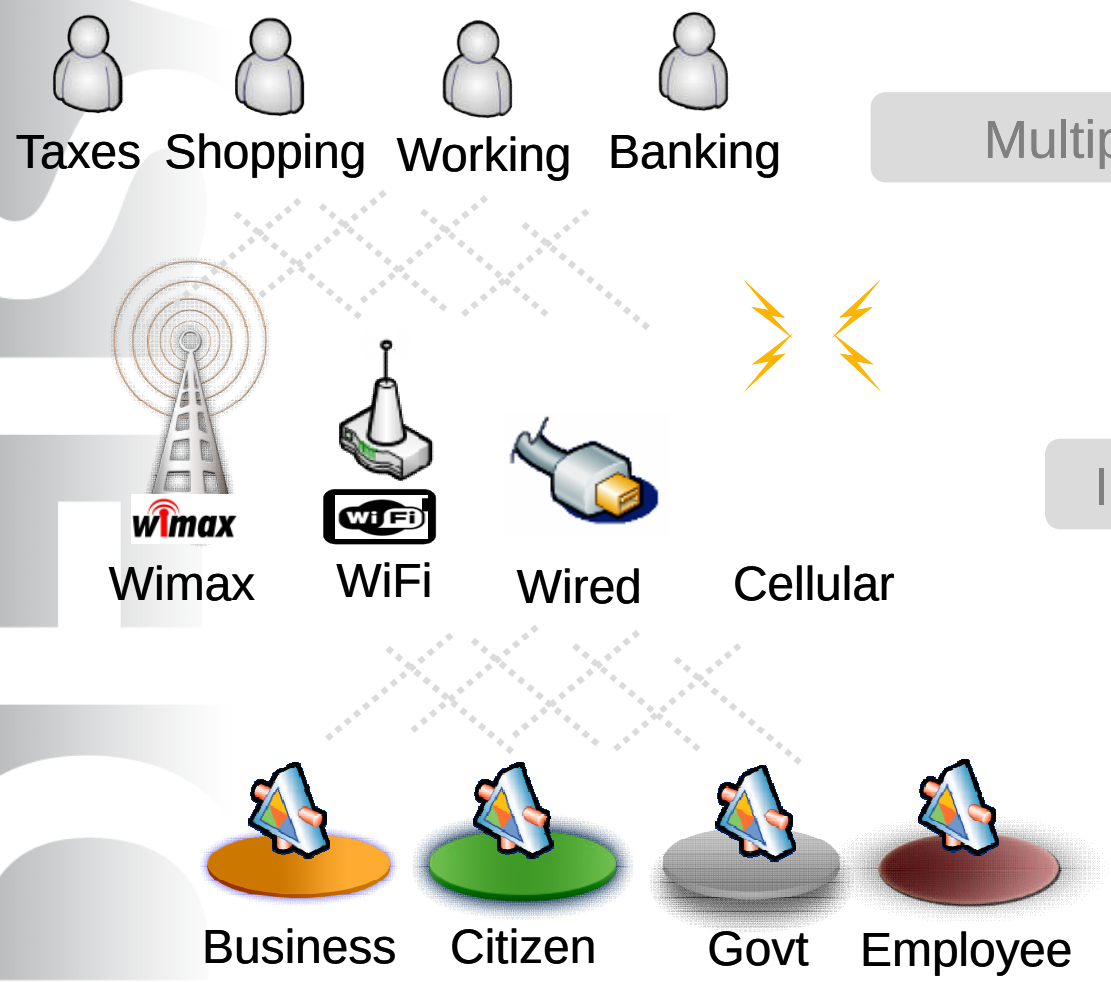

Building the Digital City Today

Distributed European Infrastructure for Supercomputing Applications

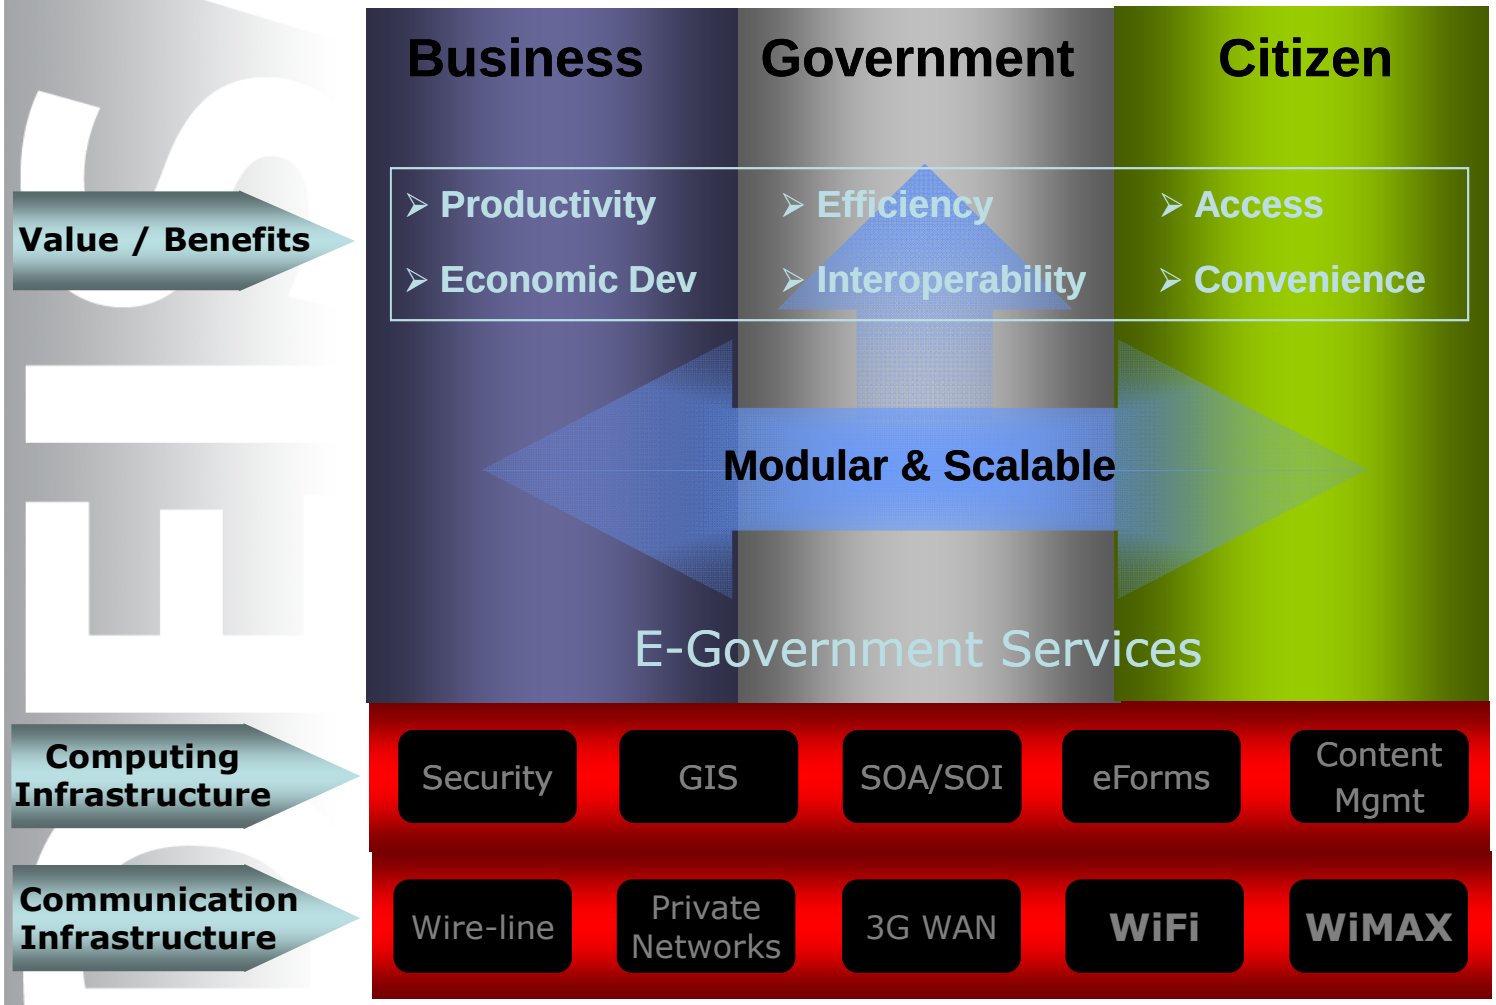

Tomorrow's Integrated Digital City

Distributed European Infrastructure for Supercomputing Applications

A single, portable identity based on strong security

One Identity

An intelligent infrastructure supporting seamless access

One Network

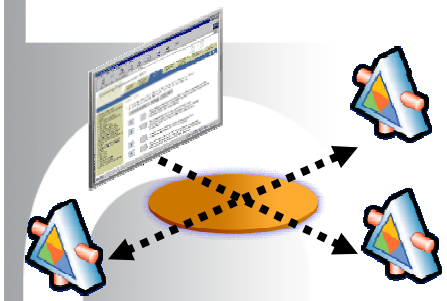

A gateway to integrated e-Services spanning multiple agencies

One Face

Finally: Connected: anyone, anywhere, anytime, any device

- Integrating new devices, data and information sources
- Cells, PDAs, smart sensors, sensor arrays, health monitors
- Embedded in cars, engines, roads, bridges, clothes,...
- Huge amount of data for real-time analysis
- Policies, SLAs, grid & cloud economy
- Support organizational and societal structures, to bridge political and social boundaries . . .

. . . very much like any other vital infrastructure,
e.g. roads, telco, water, electricity, etc…

Distributed European Infrastructure for Supercomputing Applications

Thank You!

Merci!

gentzsch@rzg.mpg.de