Distributed European Infrastructure for Supercomputing Applications

ESAC Workshop, March 11 – 12, 2010

Service Infrastructures for Science: HPC, Grids, and Clouds

The DEISA Experience

Wolfgang Gentzsch
EU DEISA Project, OGF Board of Directors
gentzsch@rzg.mpg.de

Cyberinfrastructure Vision for 21st Century

Distributed European Infrastructure for Supercomputing Applications

Dr. Arden L. Bement, Jr. Director of the National Science Foundation, in March 2007

New cultural community that supports peer collaboration and new modes of education, based upon:

- open access to leadership computing;
- data and information resources;
- online instruments and observatories;
- visualization and collaboration services.

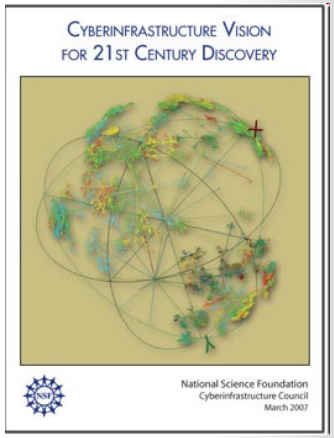

CI enables distributed knowledge communities collaborate across disciplines, distances and cultures.

Research & edu communities becoming virtual organizations that transcend geographic and institutional boundaries.

We built Service Infrastructures

Distributed European Infrastructure for Supercomputing Applications

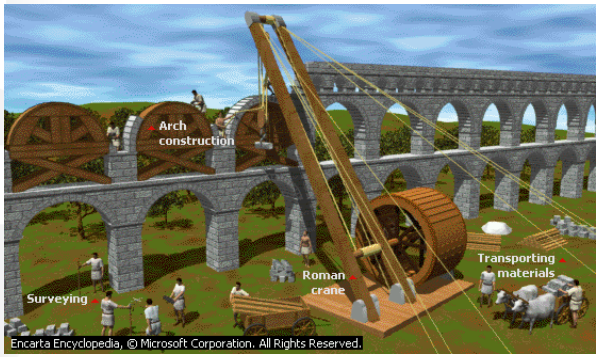

Ancient Rome: 10 aqueducts
150,000 m^3 of water each day

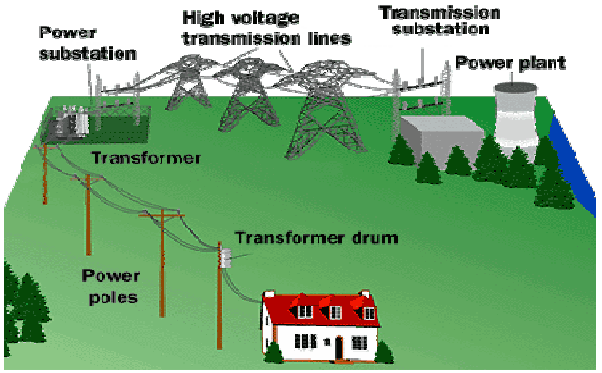

Electrical Power Grid
Infrastructure

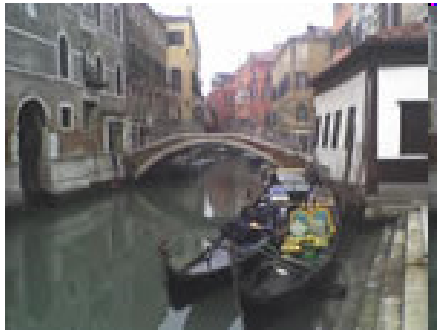

Transportation
Land, water, air

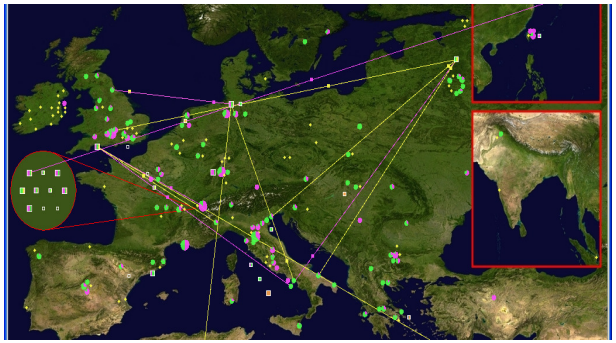

Internet
WWW, Grids, Clouds

Service Infrastructure for Collaboration in Science

eGee
Enabling Grids for E-sciencE

- ~ 280 sites in 54 countries
- > 200 Virtual Organizations
- ~ 110 000 CPUs (March 2009)
- ➢ 20 PB storage
- ➢ 16000 users
- > 250K jobs/day

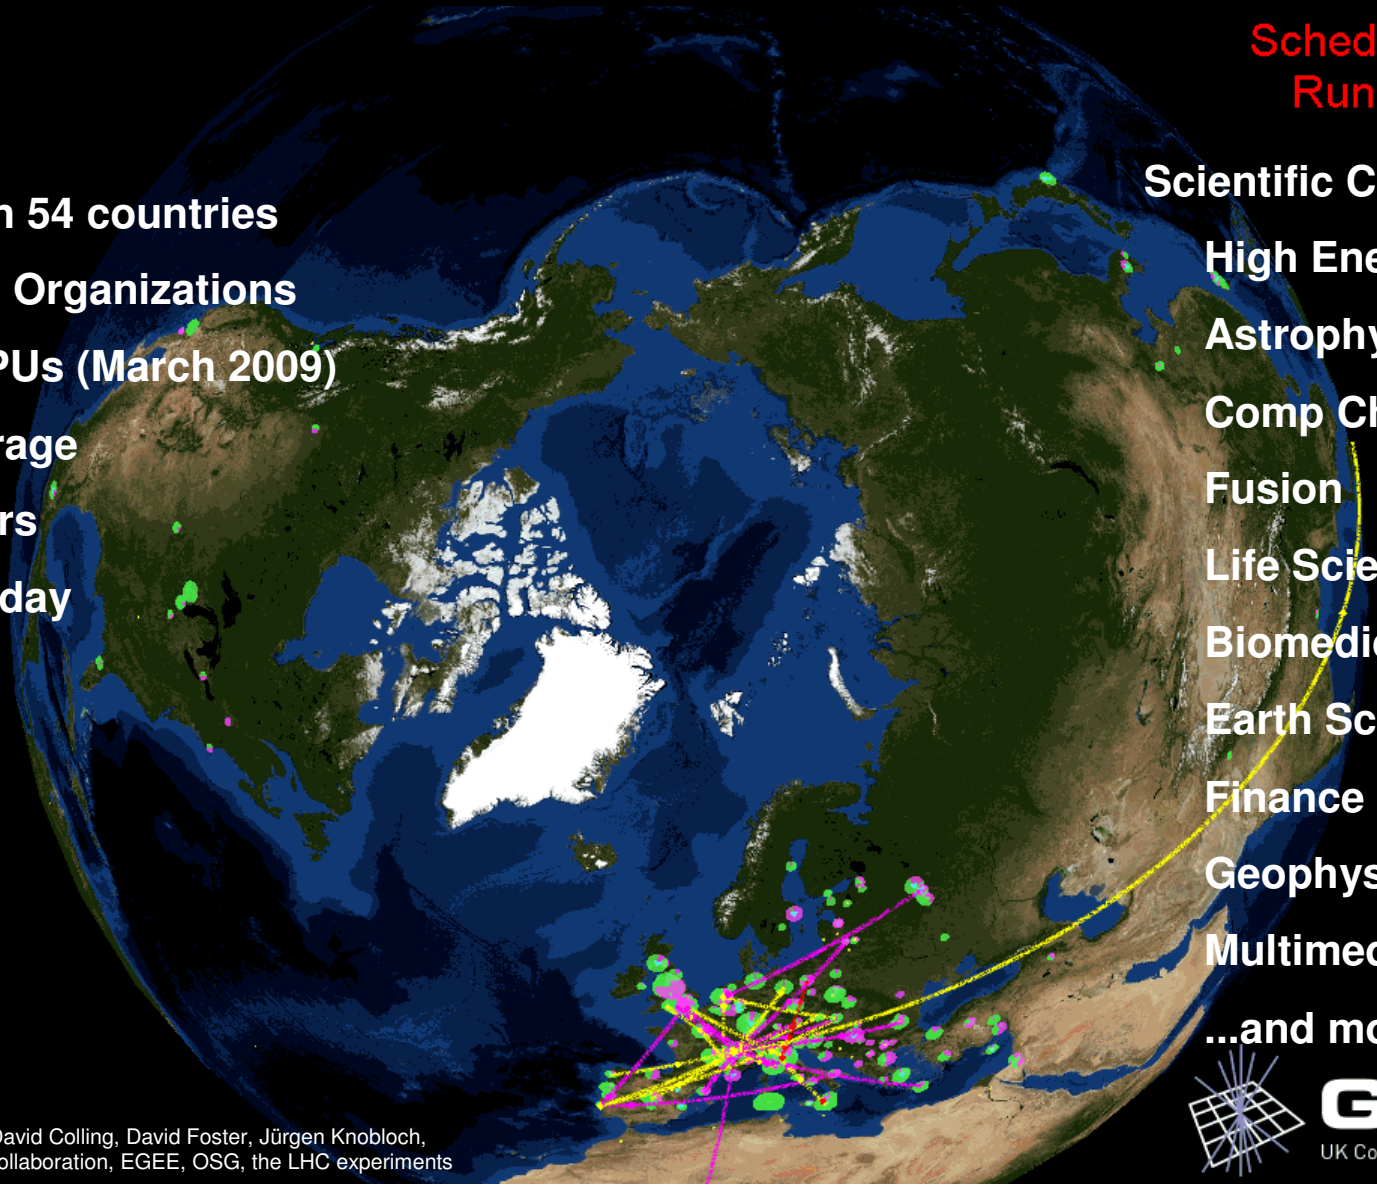

Scheduled = 21539
Running = 25374

Scientific Communities:

- High Energy Physics
- Astrophysics
- Comp Chemistry
- Fusion
- Life Sciences
- Biomedicine
- Earth Sciences
- Finance
- Geophysics
- Multimedia
- ...and more

Acknowledgements:
Julia Andreeva, Ian Bird, David Colling, David Foster, Jürgen Knobloch, Faïrouz Malek, the LCG Collaboration, EGEE, OSG, the LHC experiments

GridPP
UK Computing for Particle Physics

Requirements for an e-Infrastructure

Enabling Grids for E-sciencE

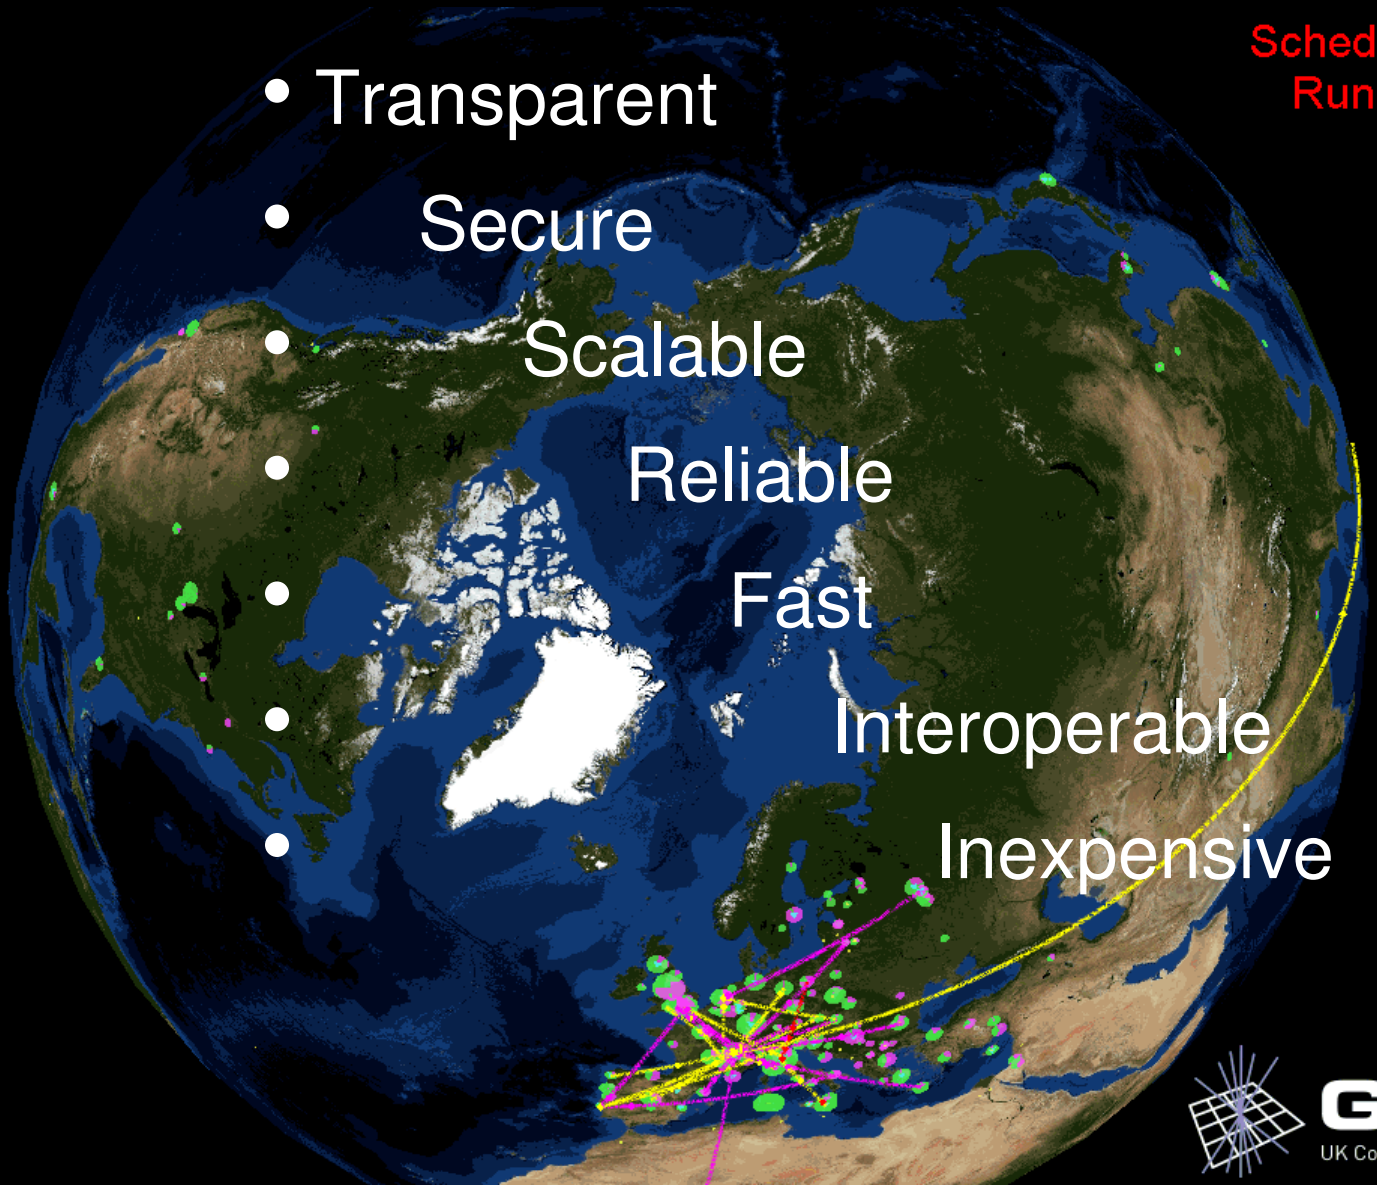

- Transparent
- Secure
- Scalable
- Reliable
- Fast
- Interoperable
- Inexpensive

GridPP
UK Computing for Particle Physics

DEISA

Distributed European Infrastructure for Supercomputing Applications

Components of an e-Infrastructure:

Networks, HPC, Grids and Clouds

Dedicated high speed network (10 Gb/s)

Distributed European Infrastructure for Supercomputing Applications

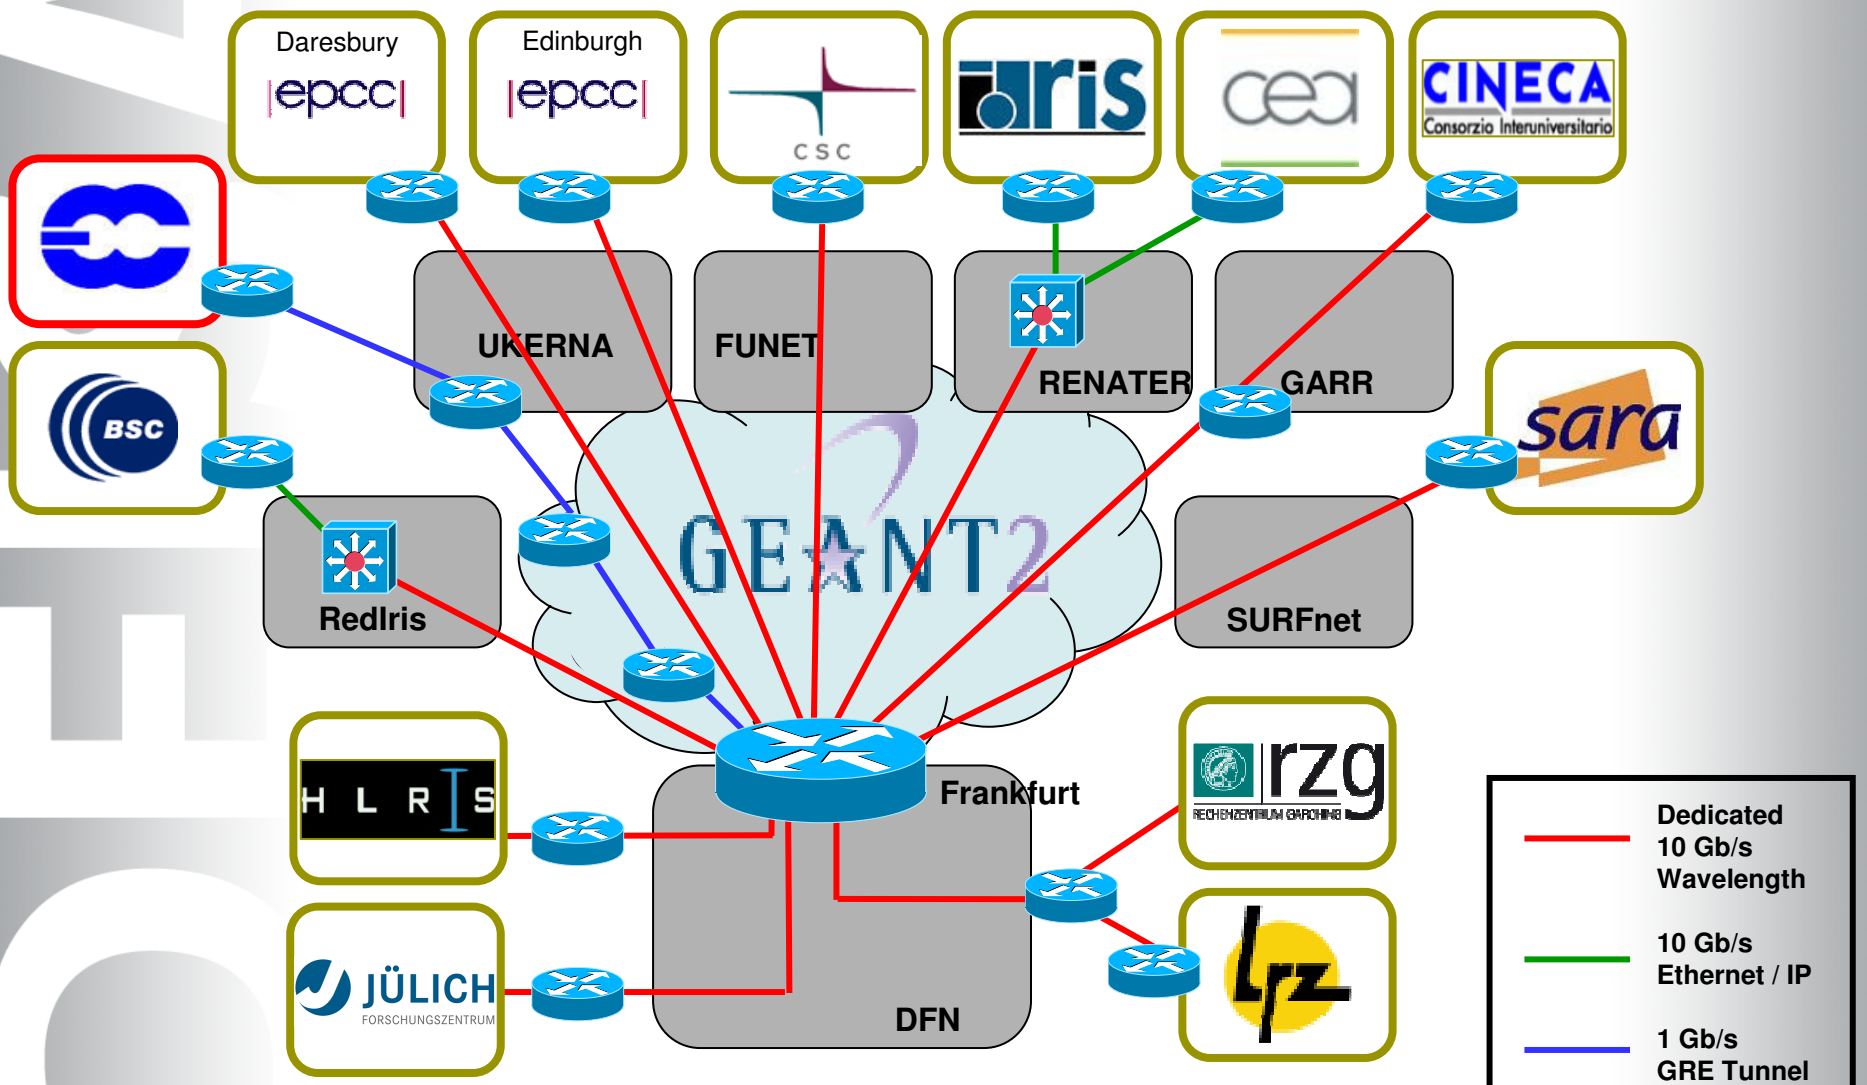

HPC Centers

Distributed European Infrastructure for Supercomputing Applications

- HPC Centers are **service providers**, for past 35 years
- IT Service: Computing, storage, applications, data, etc
- 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)

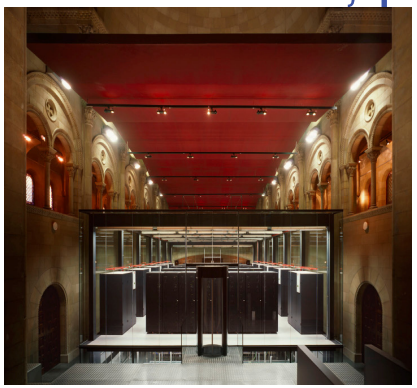

Grids

Distributed European Infrastructure for Supercomputing Applications

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

> ". . . 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

Distributed European Infrastructure for Supercomputing Applications

Grids

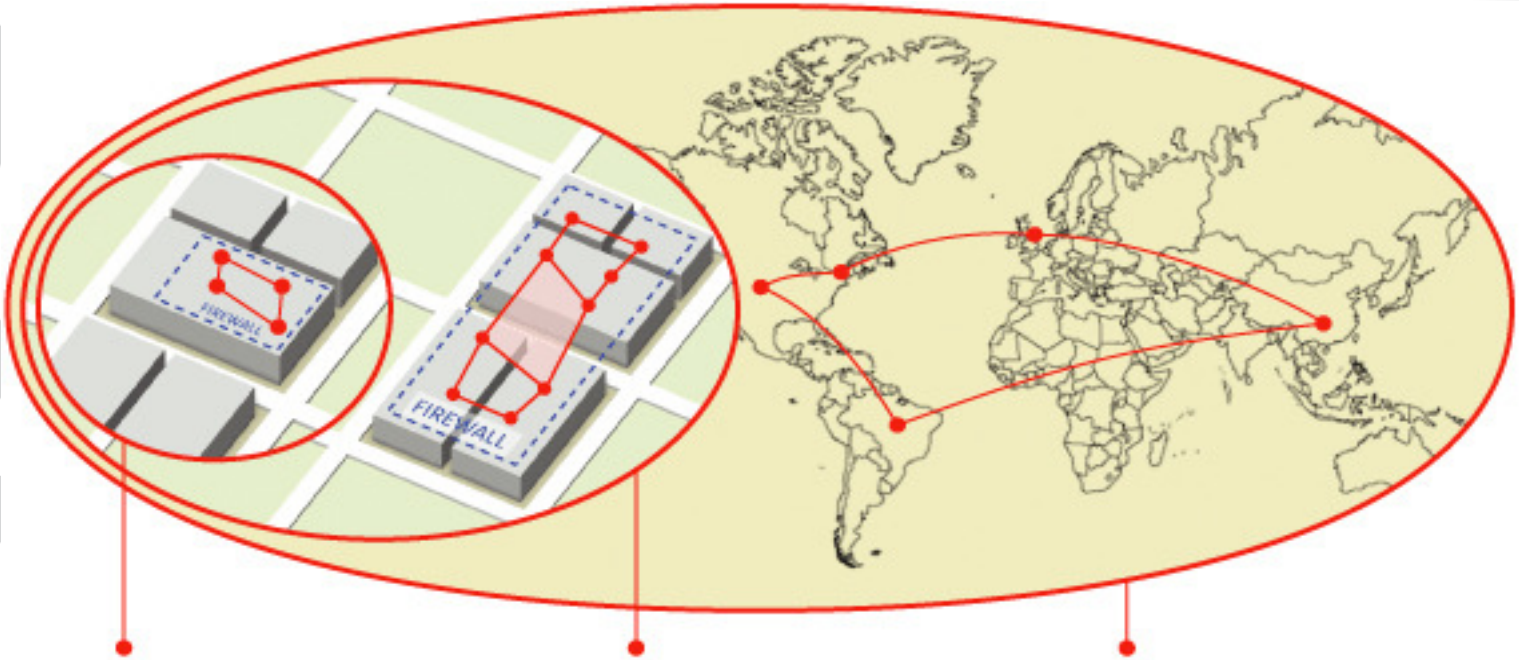

Departmental Grids

Enterprise Grids

Global Grids


Grids
TERAGRID
Distributed European Infrastructure for
UNICORE
Distributed European Infrastructure for Supercomputing Applications
Open Science Grid
INFN GRID
EUROGRID
National Grid
NG
SINGAPORE
Data GRID
Grid3 2003
K*GRID
iVDgL
PPDG
ACCESS GRID
grangenet
GRIP
UK Particle Physics
DataTAG
NEESgrid
Building the National Virtual Collaboratory for Earthquake Engineering.
GriPhyN
Data Intensive Science
CERN
openlab for DataGrid applications
Developing Solutions for the Data-Intensive Science of the Large Hadron Collider
LCG
Grid5000
WestGrid
NORDUGRID
Grid Solution for Wide Area Computing and Data Handling
GridCat
tota la informació sobre el projecte
presentació
descripció tècnica
els serveis oferts
participants
documentació pública
com puc participar?
accés al Portal
Intranet per a investigadors
estat dels serveis actius
dintre del grid
APGrid
Asia-Pacific Grid
Grid Consortium Japan
Ninf
GridLab
eGee
Enabling Grids for E-science in Europe
NAREGI
超高速コンピュータ網形成プロジェクト
National Research Grid Initiative
Grid Applications
Grid Middleware
Networking
国立情報学研究所グリッド研究開発推進拠点 NII -The National Institute of Informatics
grid.it project
Grid-Ireland

Cloud... 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

Why should my App run in the Grid ?

Distributed European Infrastructure for Supercomputing Applications

- Closer collaboration with colleagues (VCs)
- R&D projects University - Industry
- More resources => faster/more/accurate processing
- Different architectures serve different apps
- Failover: move jobs to another system

. . . and why in the Cloud ?

- No upfront cost for additional resources
- CapEx => OpEx, pay-per-use
- Elasticity, scaling up and down
- Hybrid solution (private and public cloud)

Example of a successful e-Infrastructure:

The DEISA Ecosystem for HPC Grand-Challenge Applications

Distributed European Infrastructure for Supercomputing Applications

Status and Requirements
Example: The German D-Grid

Distributed European Infrastructure for Supercomputing Applications

Scientific Field (numbers in TeraFlop/s)	2005-2007	2007-2009	2010
Climate and Earth System Research	20	50-100	>500
Geophysics	1	10-100	>1000
Nanostructure Physics	1	10-50	>200
Solid-State Physics	1	50-100	>1000
Computational Fluid Dynamics	2.5	25-100	>1000
Astrophysics	10	50-100	>500
Elementary Particle Physics and Physics of Hadrons and Nuclei	30	100	>1000
Materials Science	10	50-100	>500
Theoretical Chemistry	3	25-125	>300
Soft Matter	3	30	>200
Biophysics and Bioinformatics	3	15-80	>1000
Plasma Physics	10	50	>500

A. Bode, W. Hillebrandt, and Th. Lippert: German Scientific Case for the BMBF, 8/2005

new "petaflop" supercomputers

The European Union Long-Term Strategy

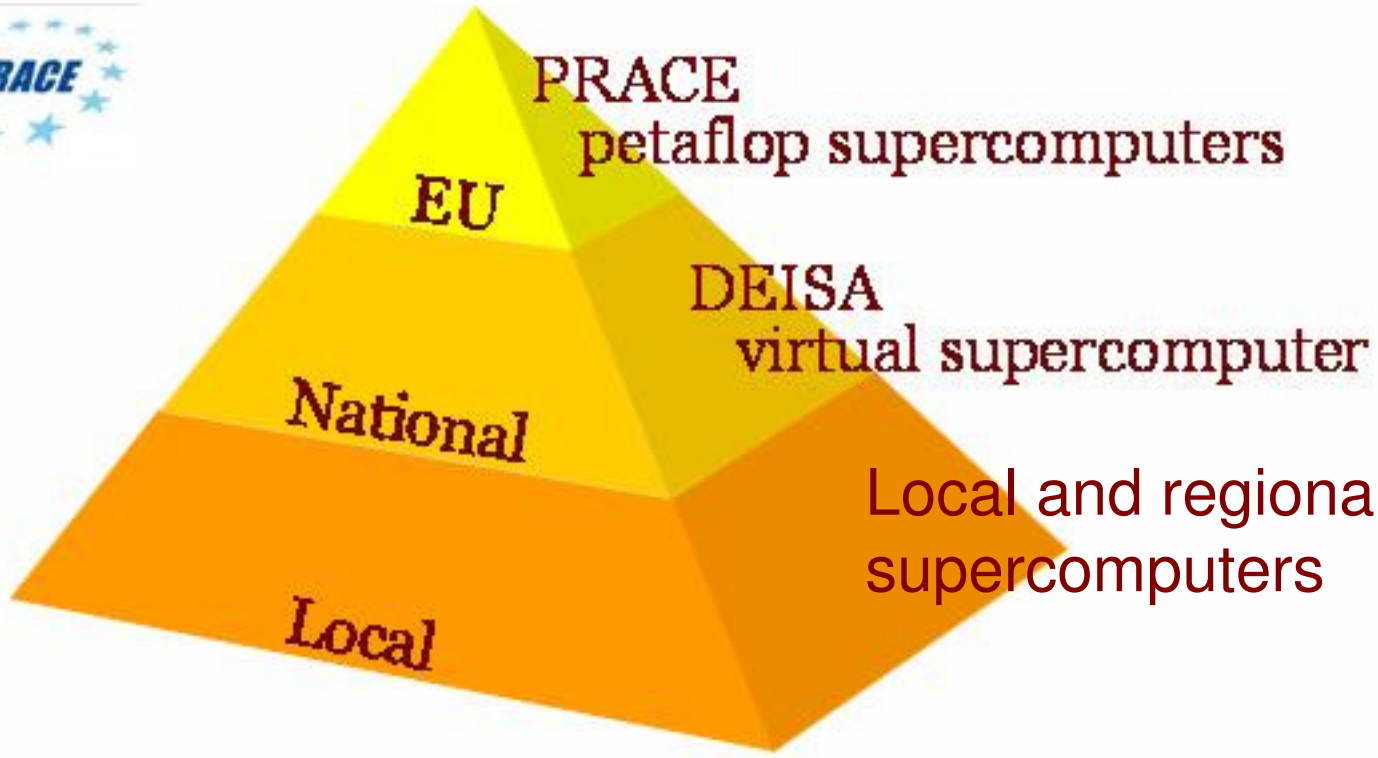

Local and regional supercomputers

Mario Campolargo
European Commission
OGF23, June 2008

e-infrastructure

DEISA Partners

Distributed European Infrastructure for Supercomputing Applications

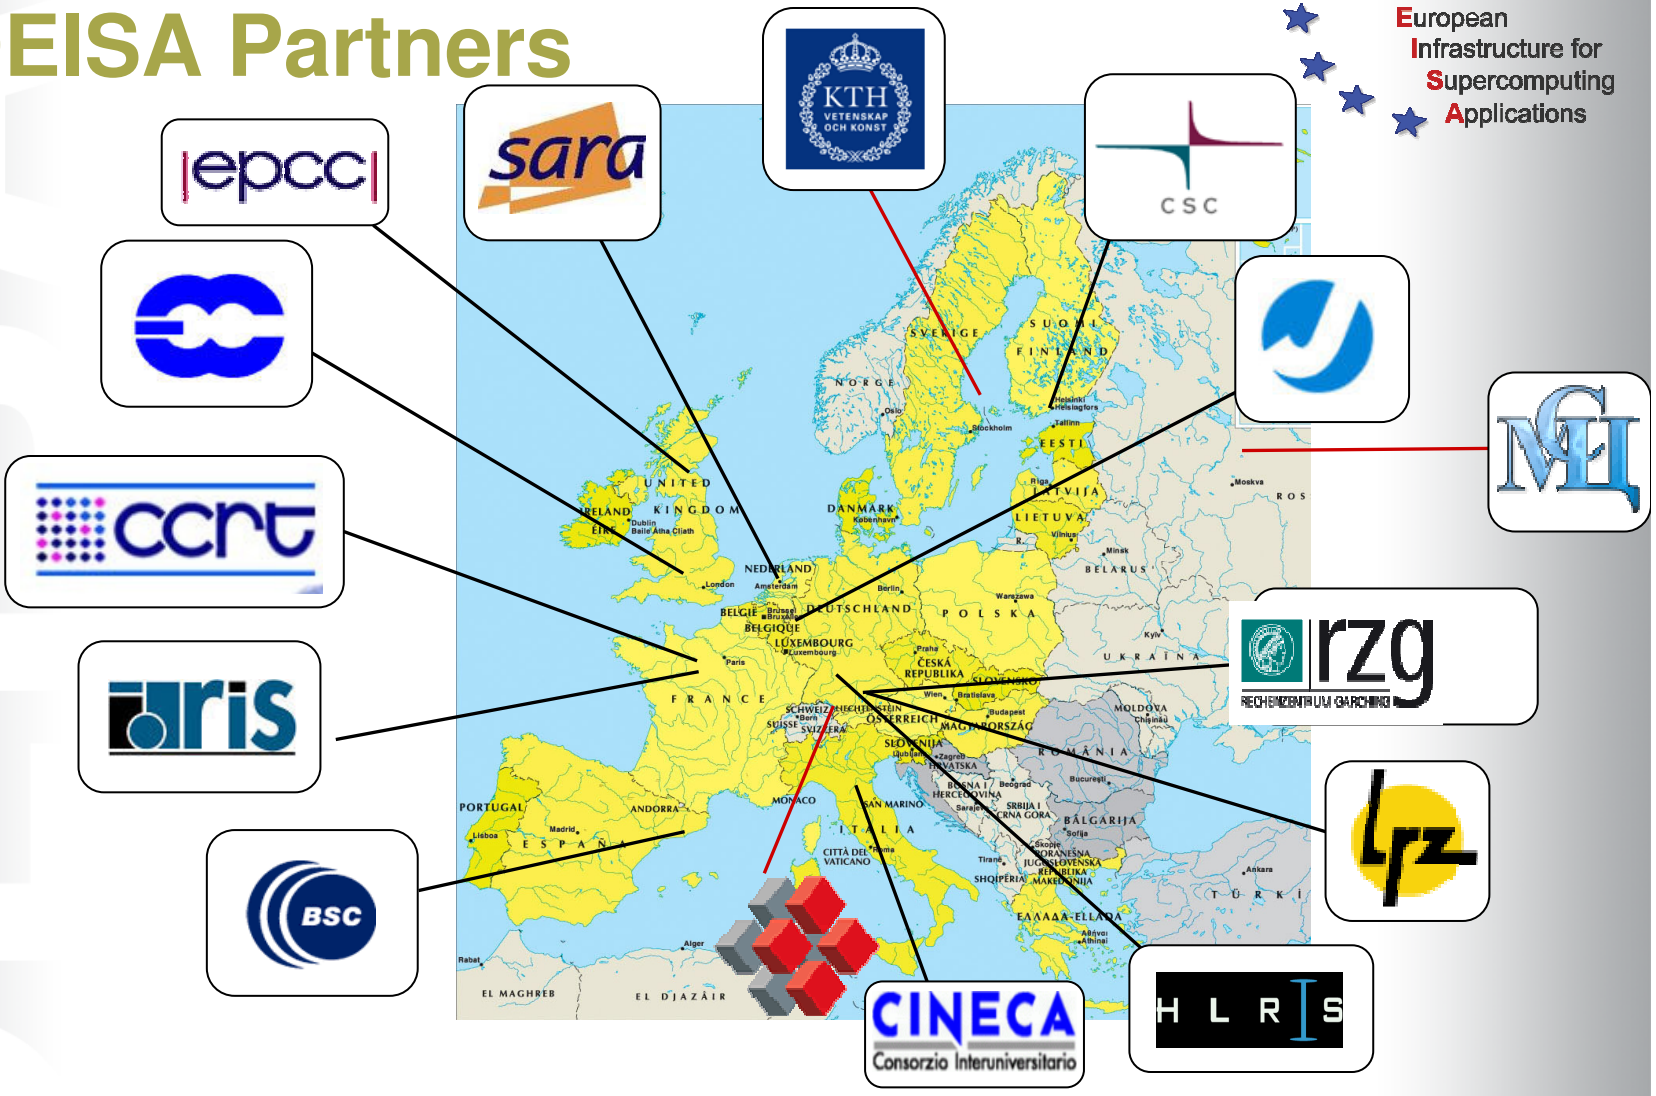

DEISA1: May 1st, 2004 – April 30th, 2008

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

SEVENTH FRAMEWORK PROGRAMME

Dedicated high speed network (10 Gb/s)

Distributed European Infrastructure for Supercomputing Applications

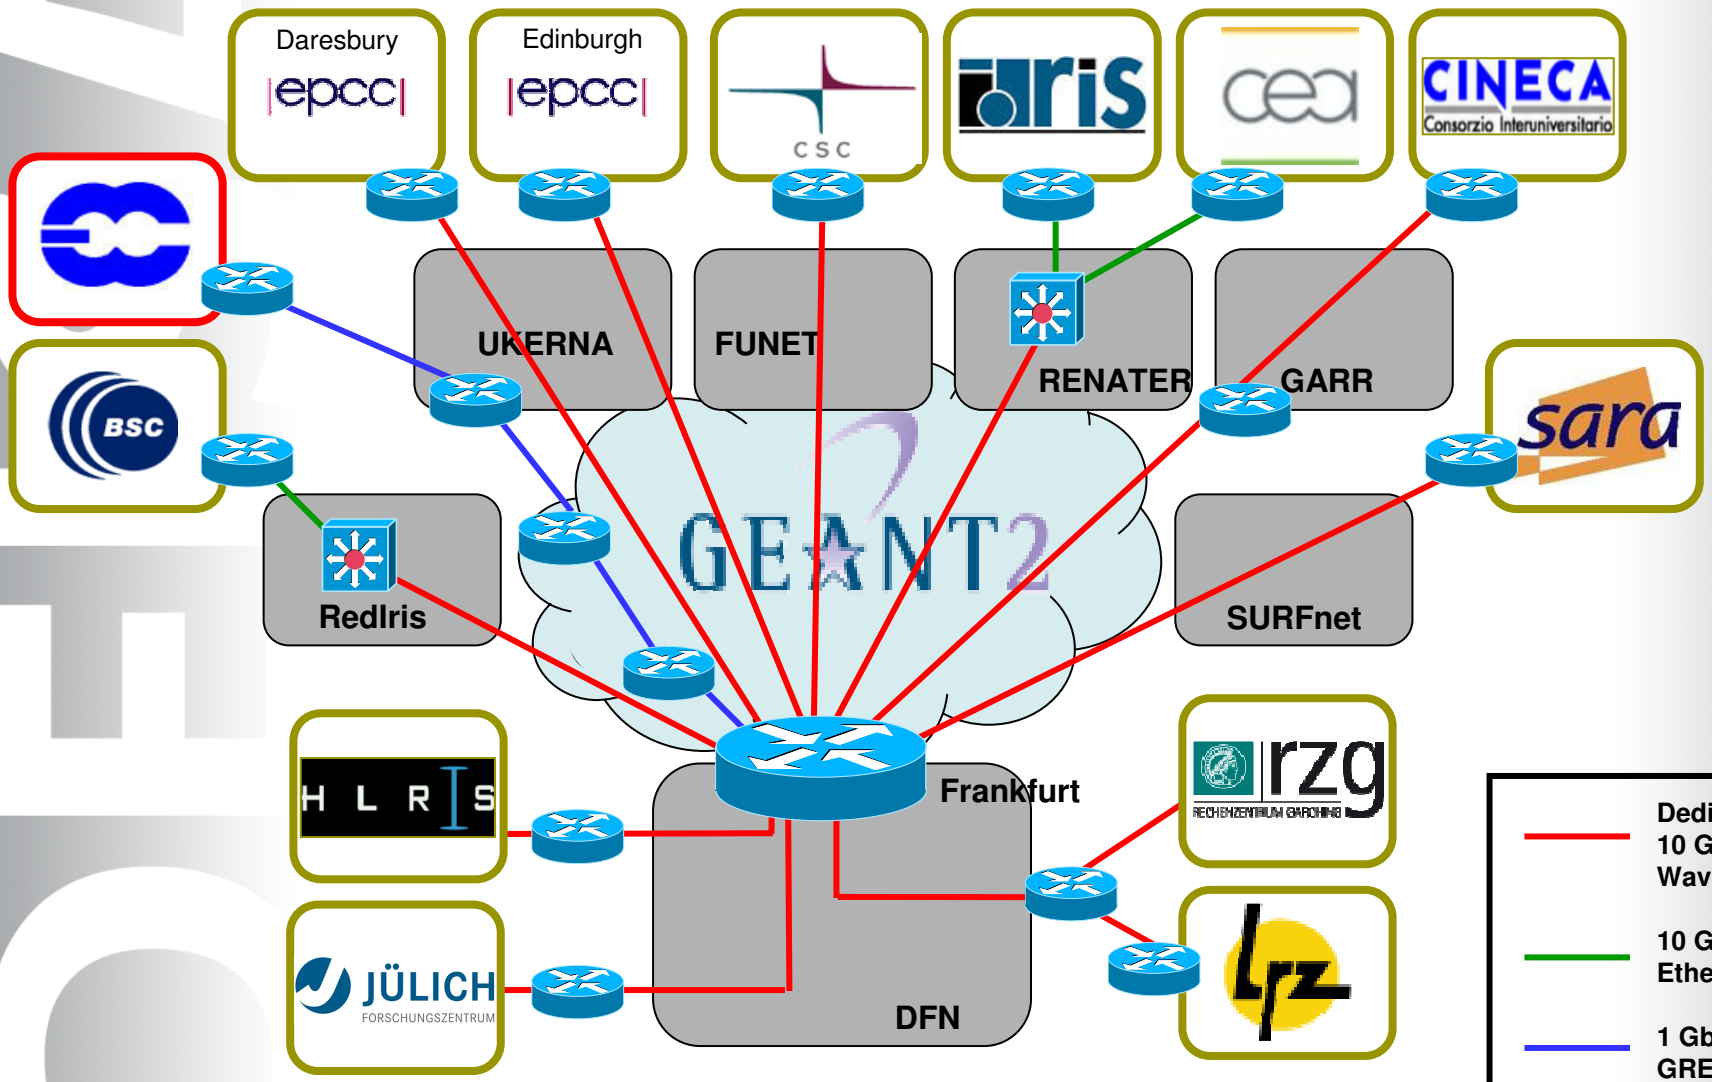

RI-222919

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.

Unified Access and Use of HPC Resources

Distributed European Infrastructure for Supercomputing Applications

Access via Internet

single sign-on (based on X.509 'Grid' certificates)
gsi-ssh -> D-ssh
Unicore, gridFTP

DEISA Common Production Environment

Different Software Environments

SE A1 | SE B1 | SE C1 | SE D1 | ..... | SE E1 | SE B2 | SE C2

Different SuperComputers - Compute elements and interconnect

Dedicated 10 Gb/s network – via GEANT2

DEISA highly performant continental global file system

RI-222919

DEISA Service Layers

Distributed European Infrastructure for Supercomputing Applications

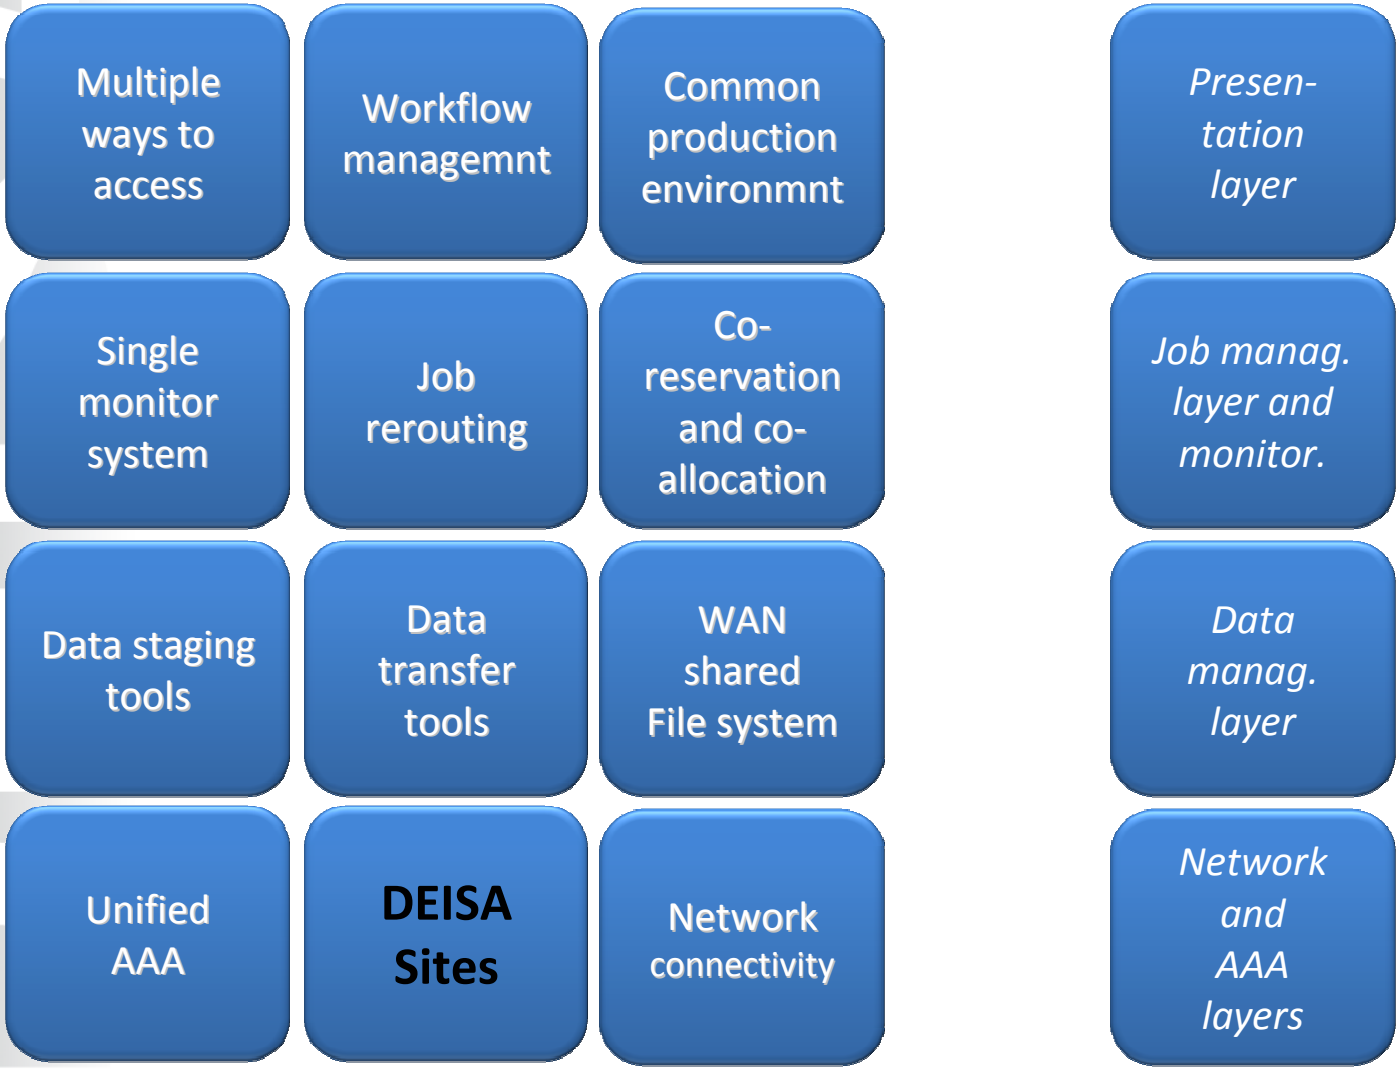

RI-222919

DEISA UNICORE Infrastructure

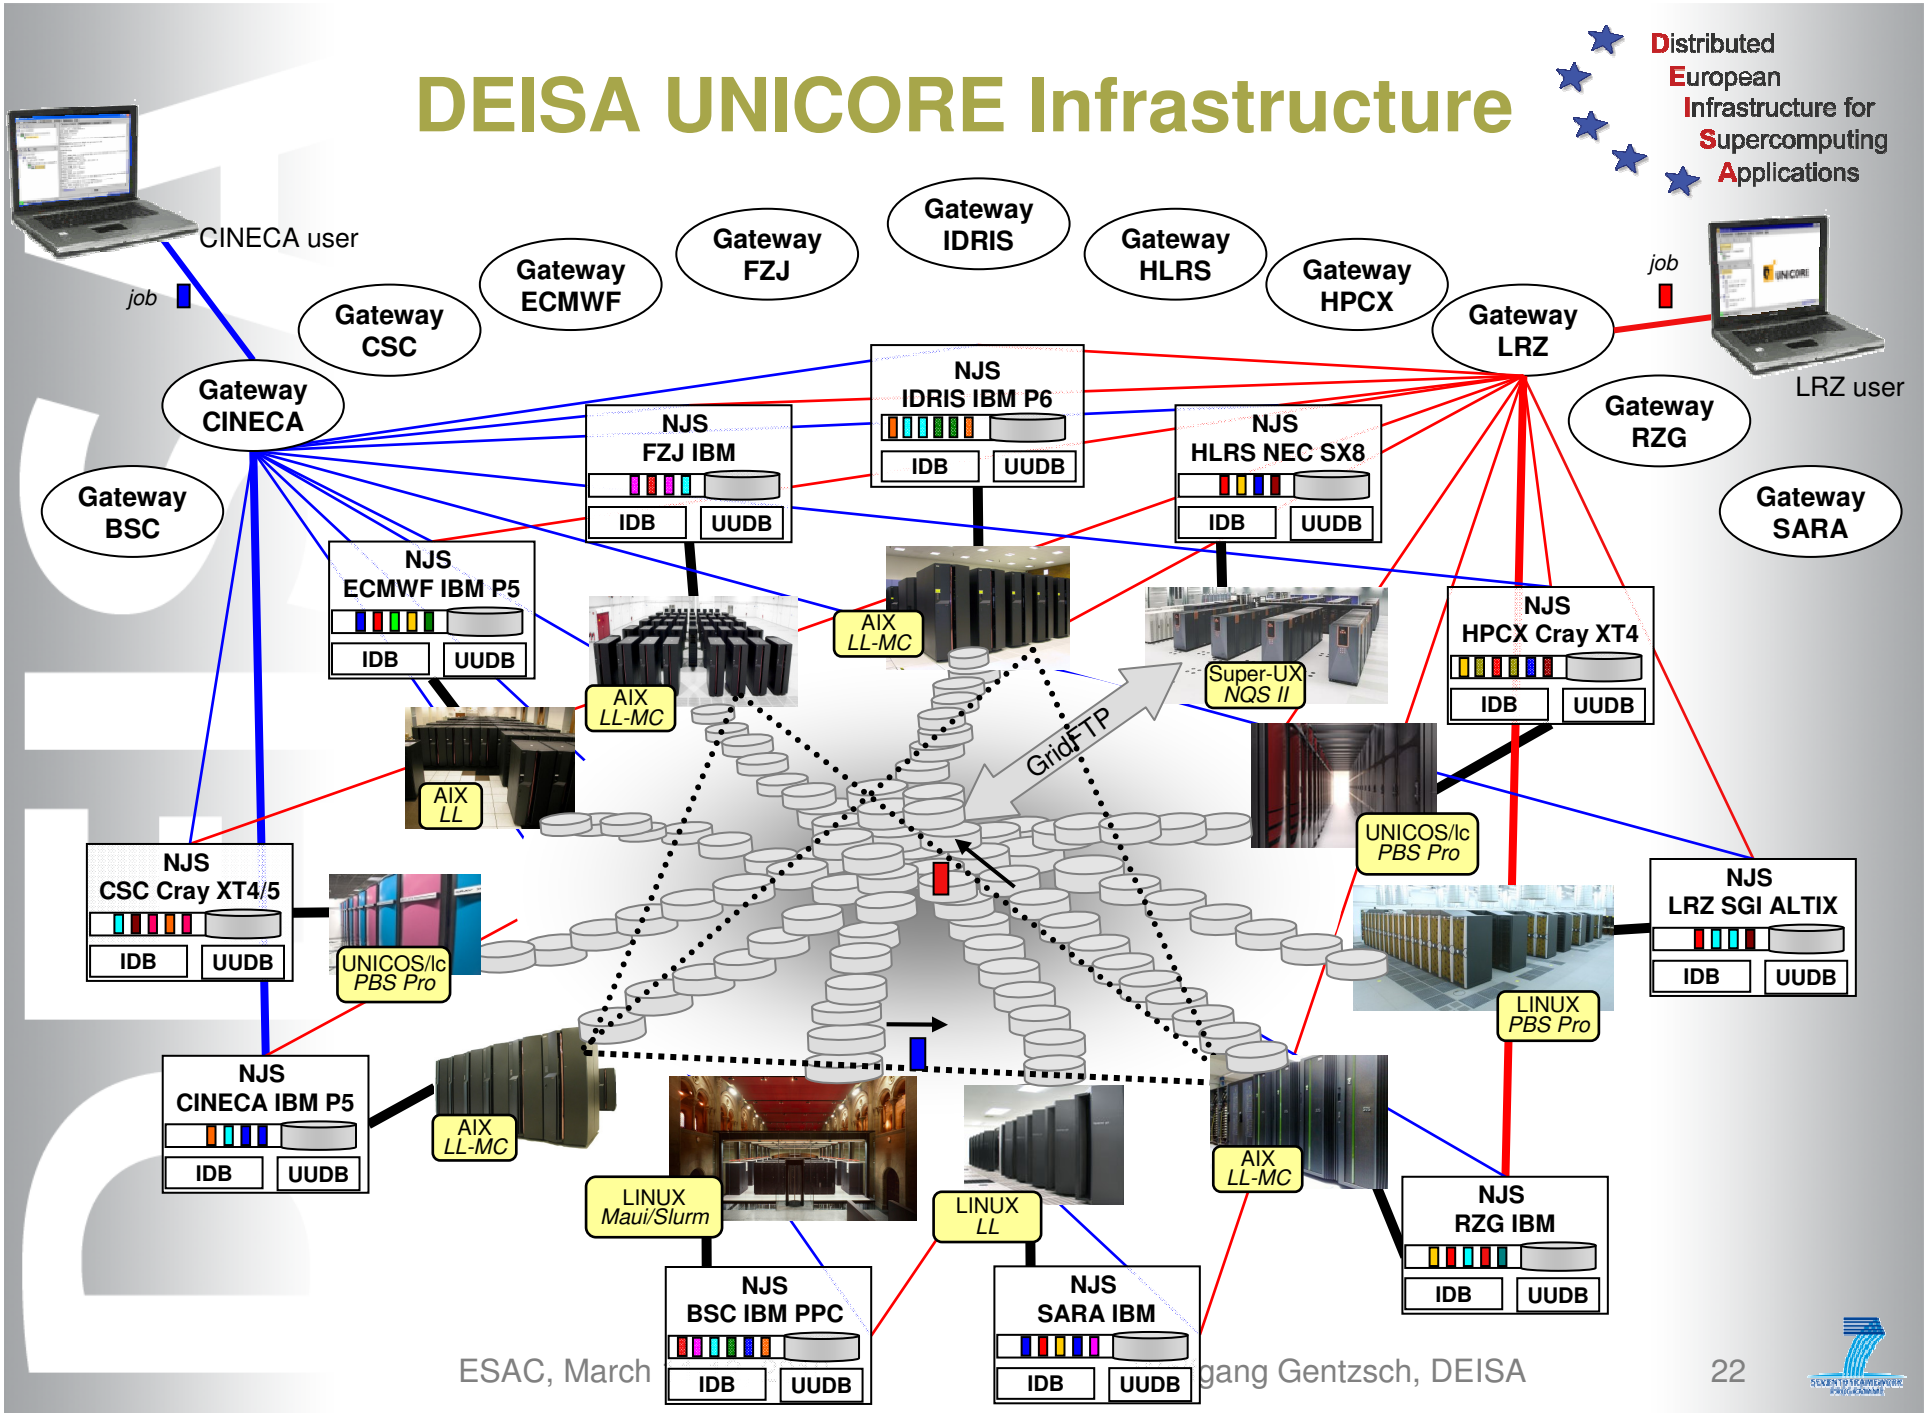

RI-222919

DEISA Global File System

Distributed European Infrastructure for Supercomputing Applications

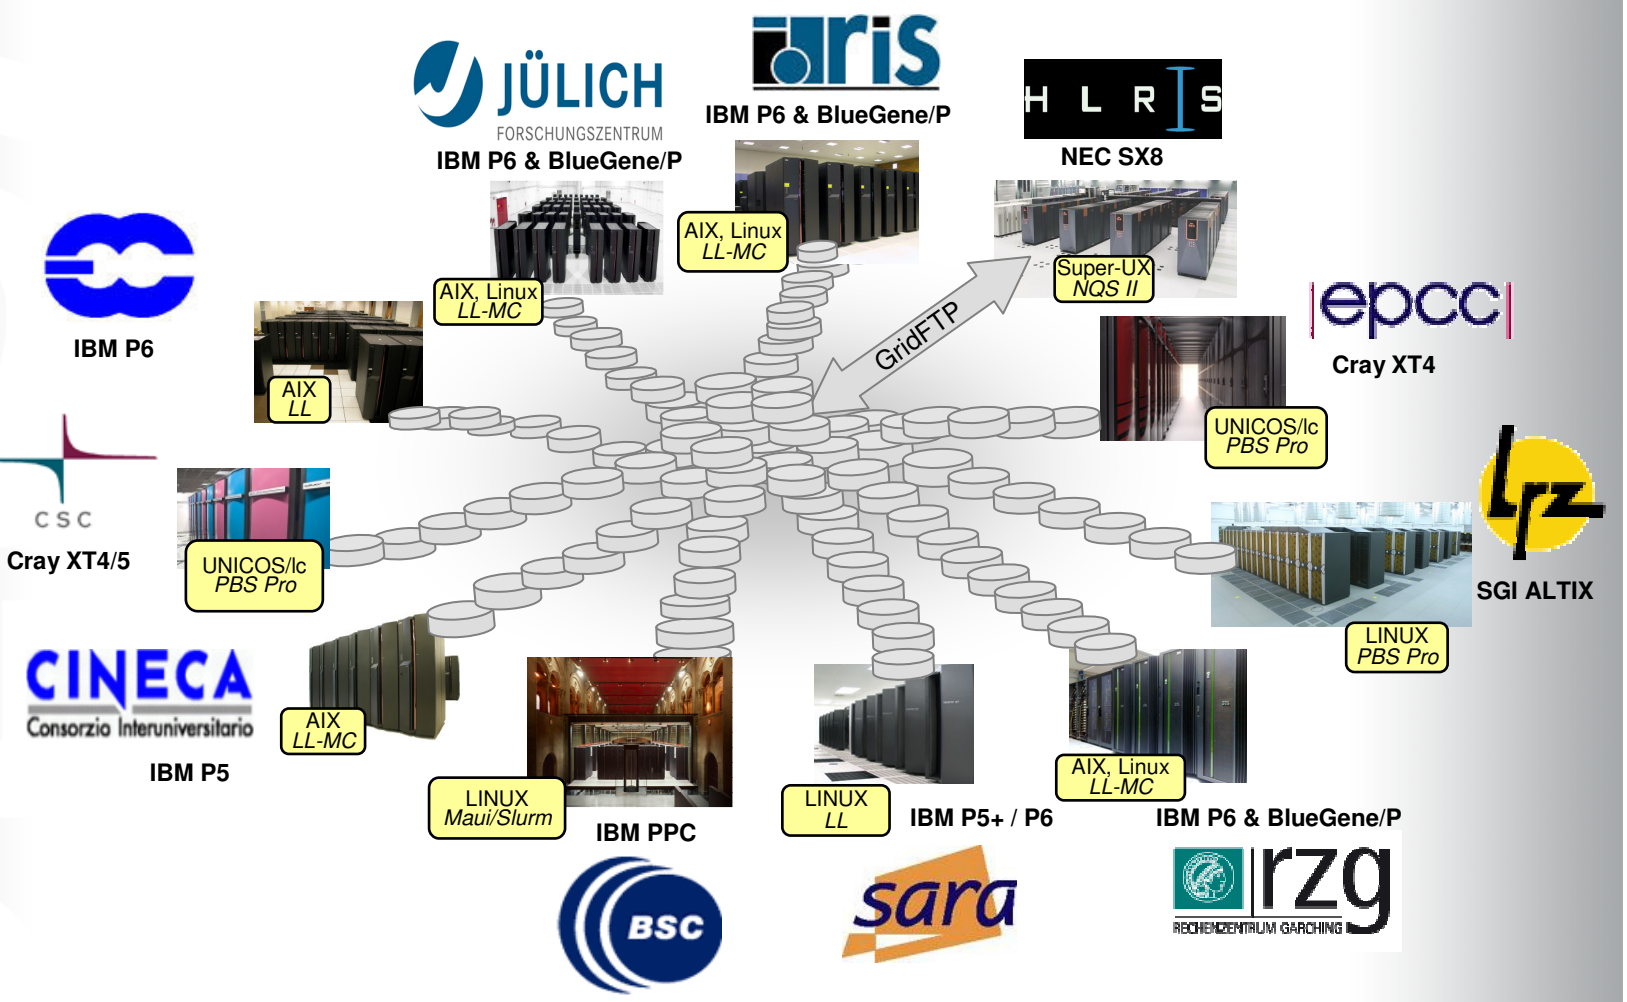

Global transparent file system based on the Multi-Cluster General Parallel File System (MC-GPFS of IBM)

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

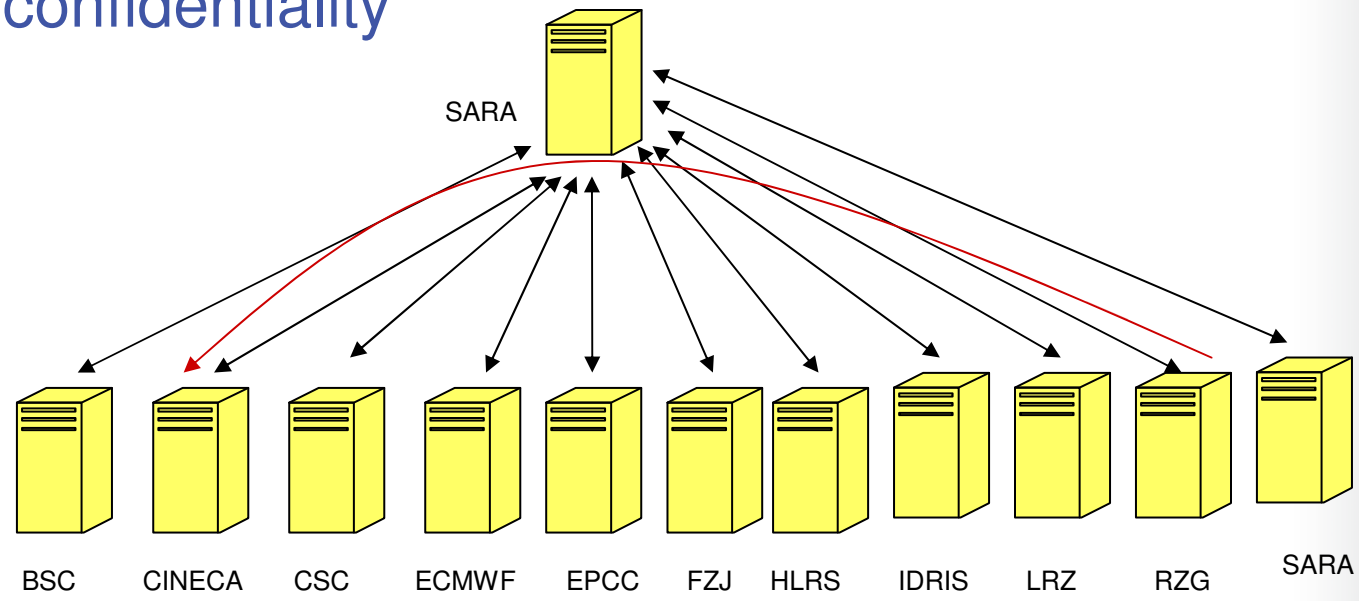

Common User Administration

Distributed European Infrastructure for Supercomputing Applications

- 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!

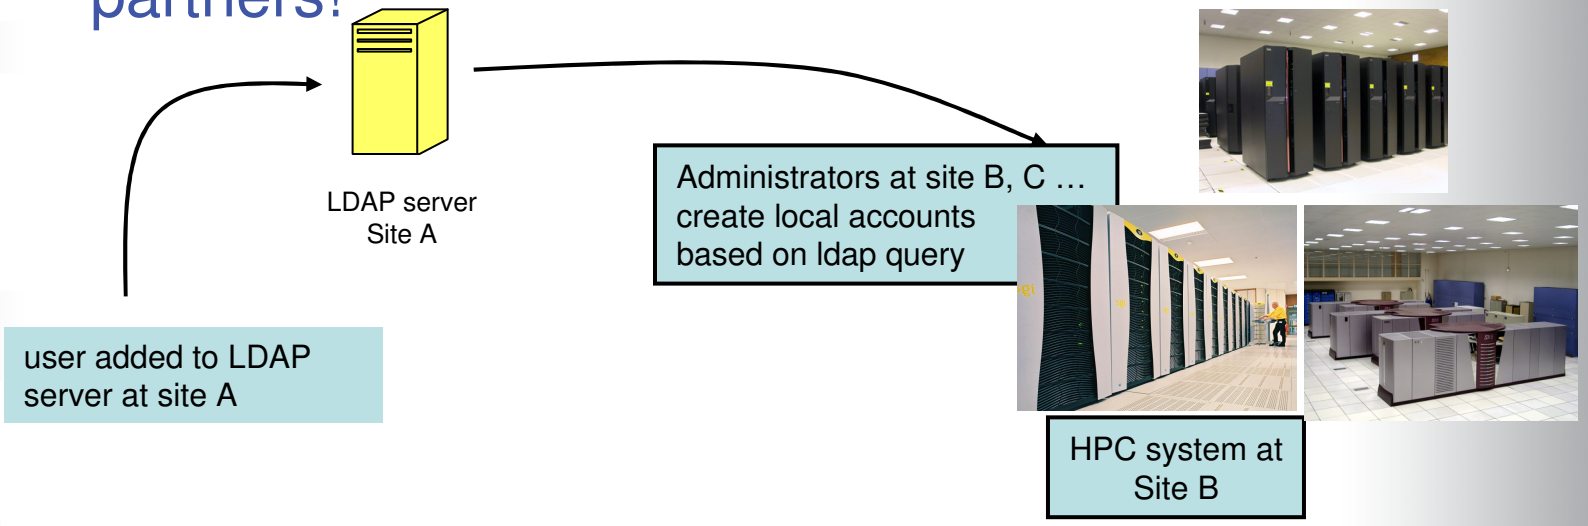

Federated Operation of DEISA

Distributed European Infrastructure for Supercomputing Applications

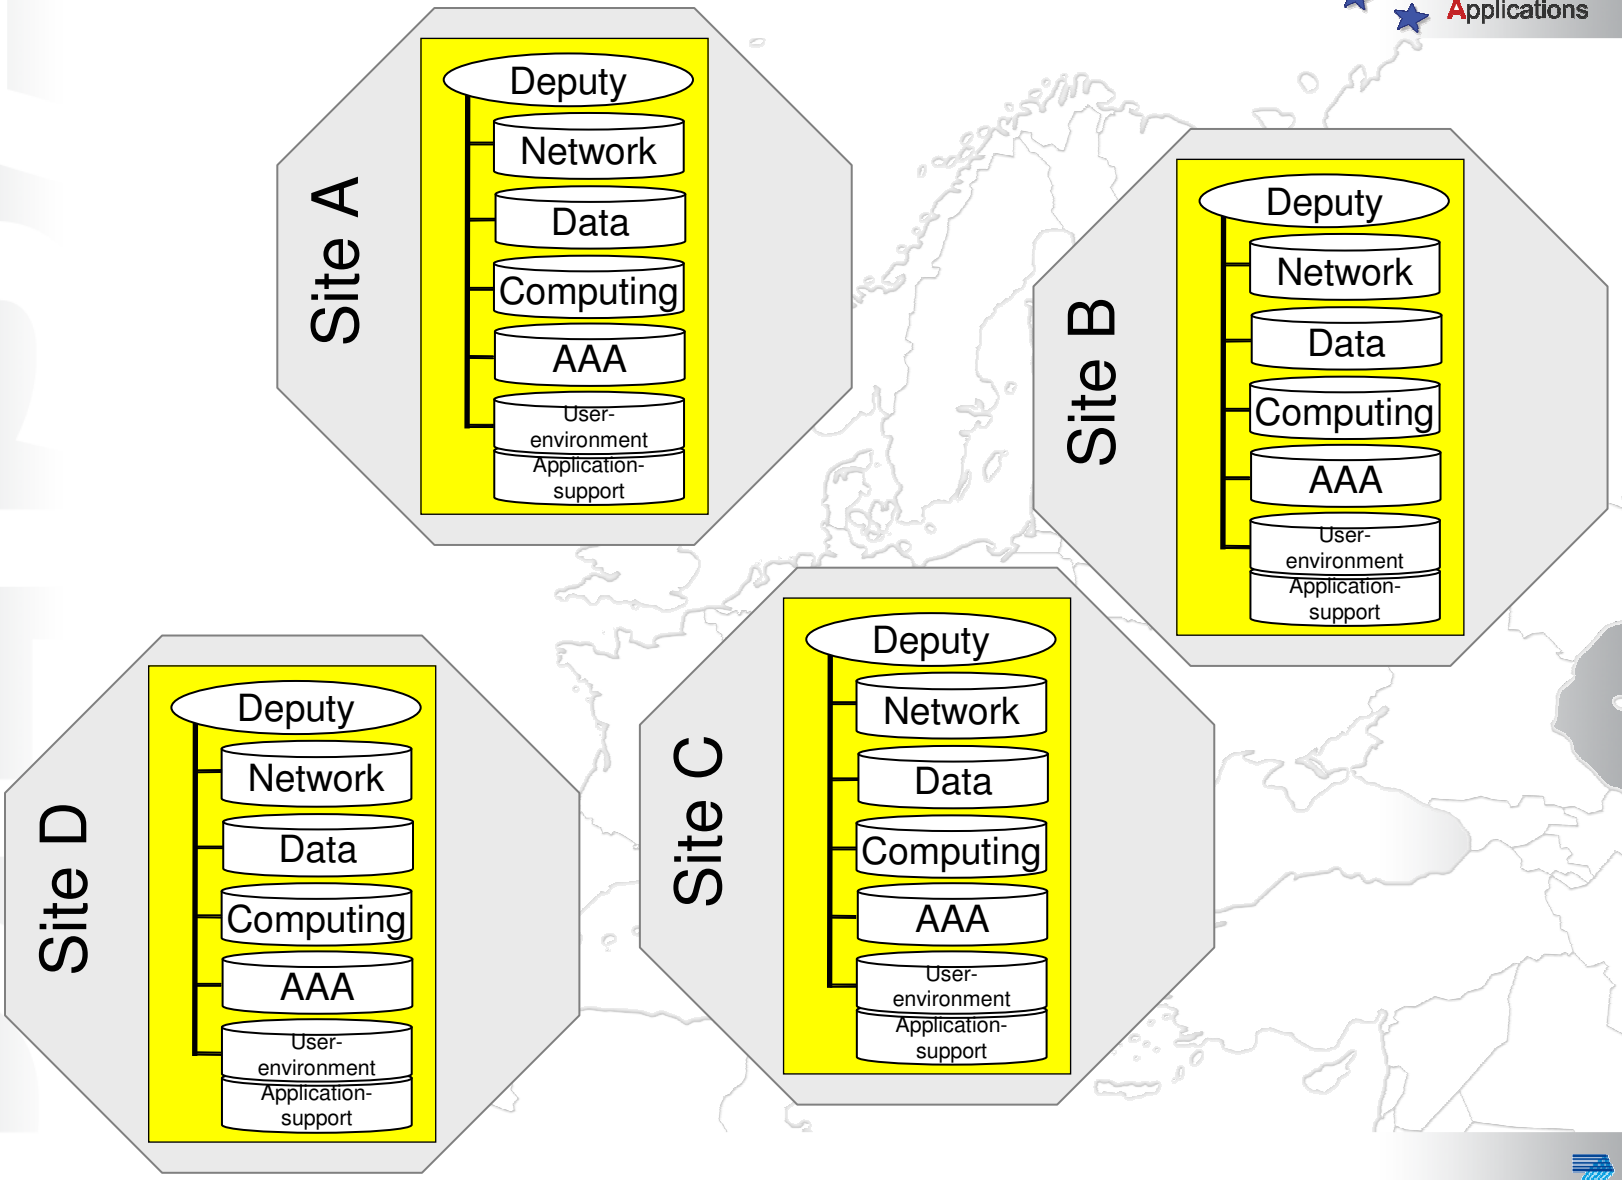

Federated Operation of DEISA

Distributed European Infrastructure for Supercomputing Applications

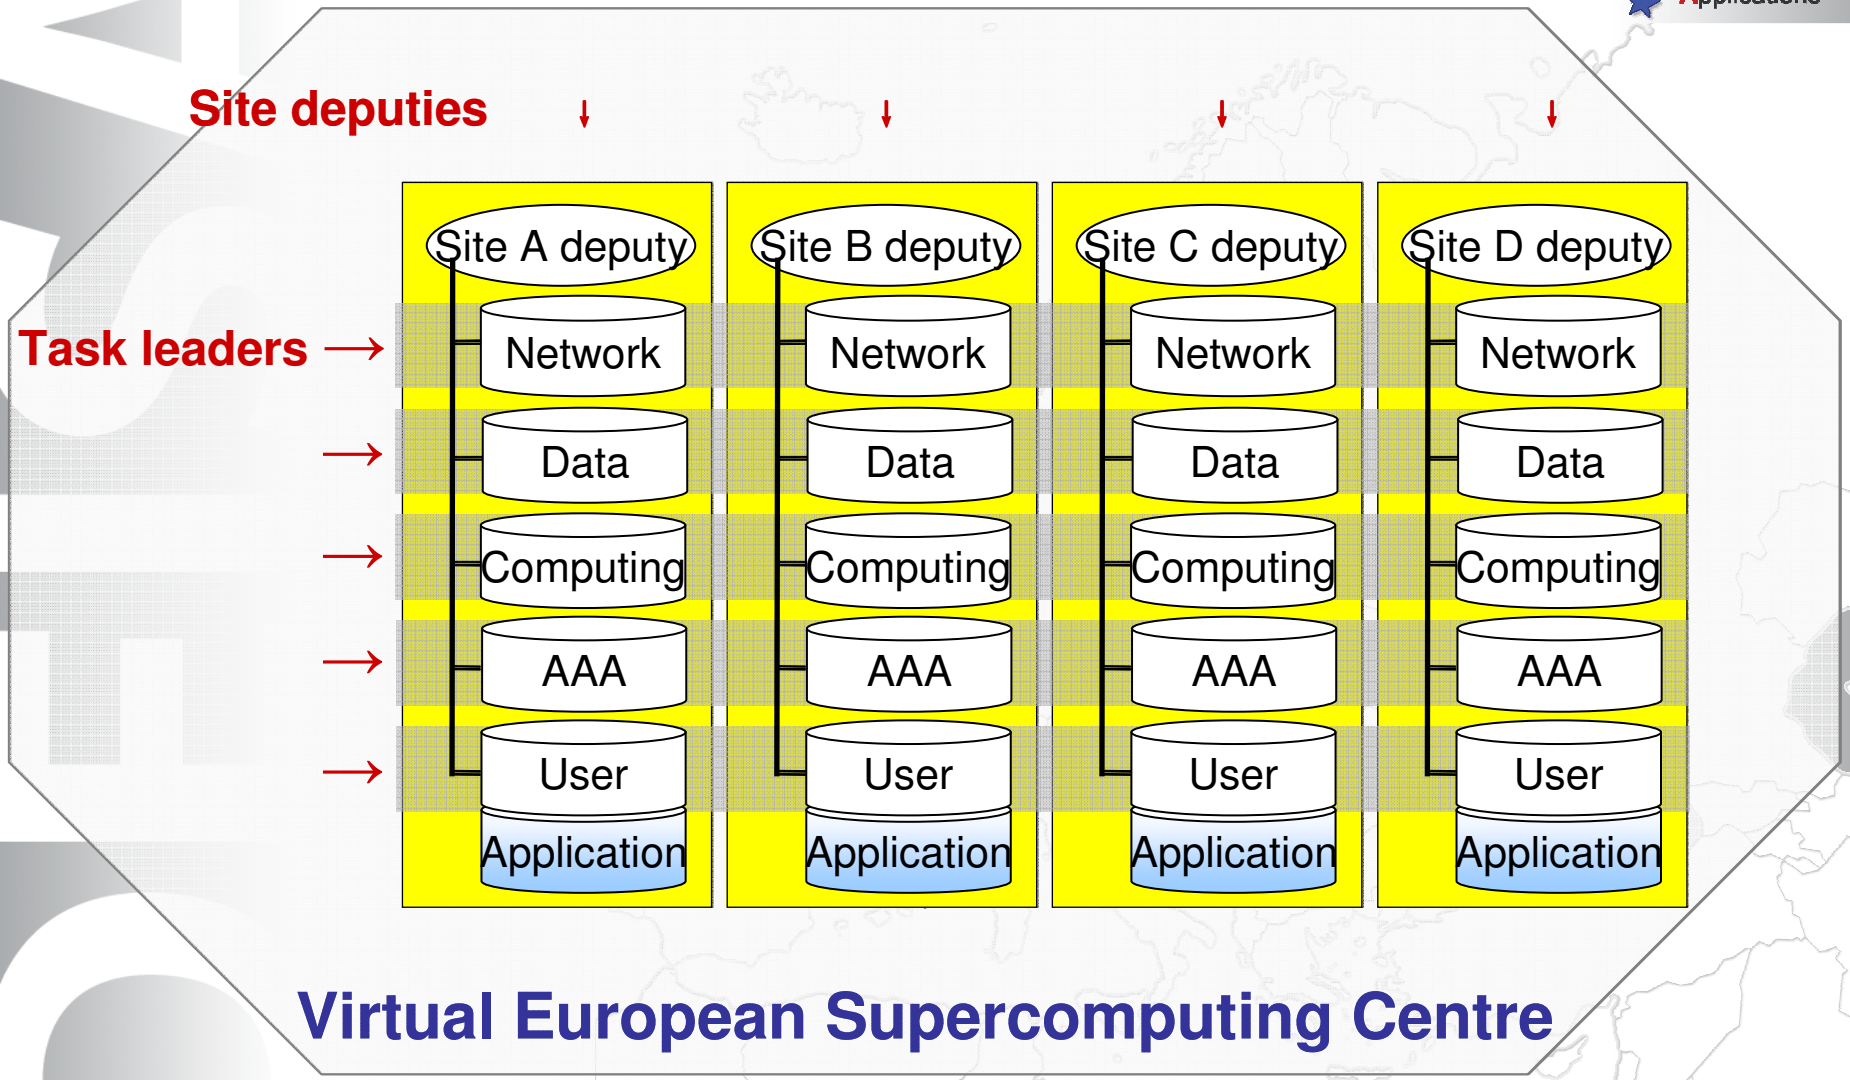

Federated Operation of DEISA

Distributed European Infrastructure for Supercomputing Applications

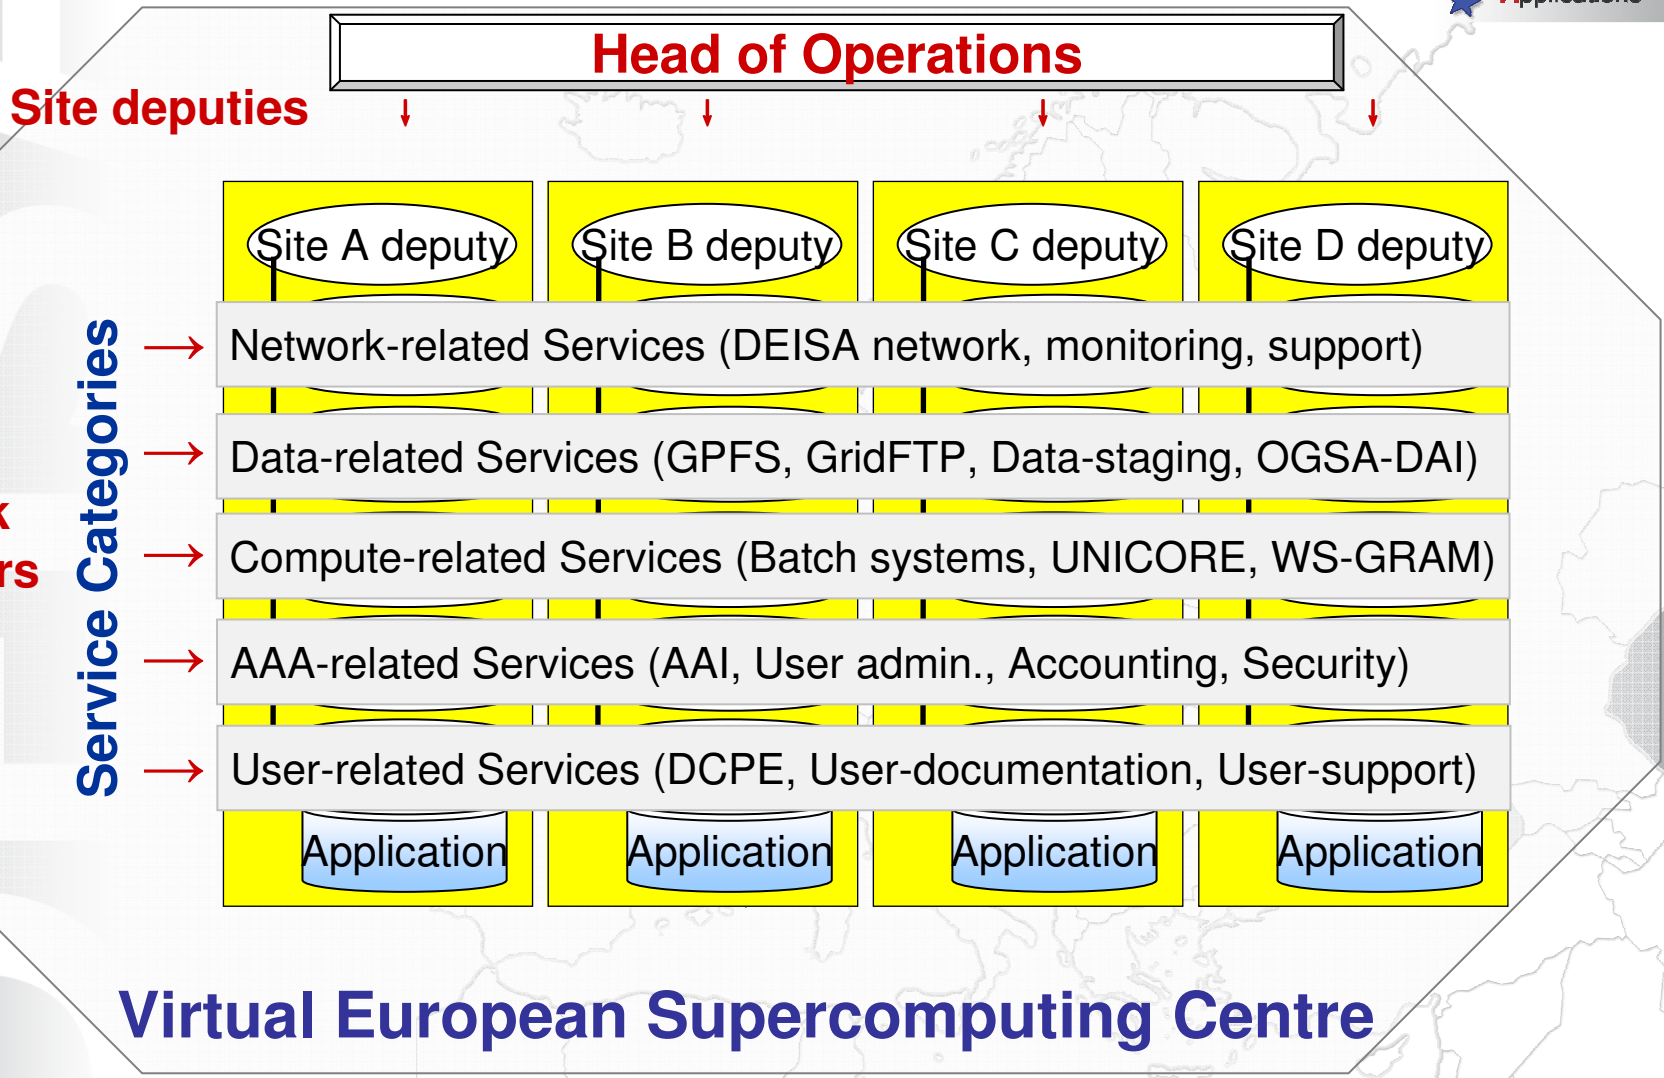

RI-222919

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

Projects from DECI calls 2005, 2006, 2007, 2008, 2009

Involvement of over 180 research institutes and universities from 25 European countries:

Austria	Belgium	Cyprus	Denmark	Finland
France	Germany	Greece	Hungary	Ireland
Italy	Latvia	Norway	Poland	Portugal
Romania	Russia	Slovac Rep.	Spain	Sweden
Switzerland	Netherlands	Turkey	Ukraine	UK

with collaborators from four other continents

North America, South America, Asia, Australia

Projects and Science Communities

Distributed European Infrastructure for Supercomputing Applications

DECI call 2005

29 proposals accepted — 12 mio core-h granted*

DECI call 2006

28 proposals accepted — 12 mio core-h granted*

DECI call 2007

45 proposals accepted — 30 mio core-h granted*

DECI call and Science Communities 2008

42 proposals accepted — 50 mio core-h granted*

3 communities — 5 mio core-h granted*

DECI call and Science Communities 2009

50 proposals accepted — 60 mio core-h granted*

7 communities — 12 mio core-h granted*

*) Core-h normalized to IBM P4+@1.7GHz

DECI: DEISA Extreme Computing Initiative
Yearly call for proposals

Communities: Virtual Scientific Communities

RI-222919

Science Communities Support

Distributed European Infrastructure for Supercomputing Applications

Life Sciences

VIROLAB

Virtual Physiological Human network of excellence

Fusion Energy Research

EFDA EUROPEAN FUSION DEVELOPMENT AGREEMENT

EUFORIA

Space Science / Cosmology

esa planck

VIRGO

Climate Research

enes European Network for Earth System Modelling

2008 3 communities — 5 mio core-h granted*
2009 7 communities — 12 mio core-h granted*

Aerodynamic Shape Optimization

Distributed European Infrastructure for Supercomputing Applications

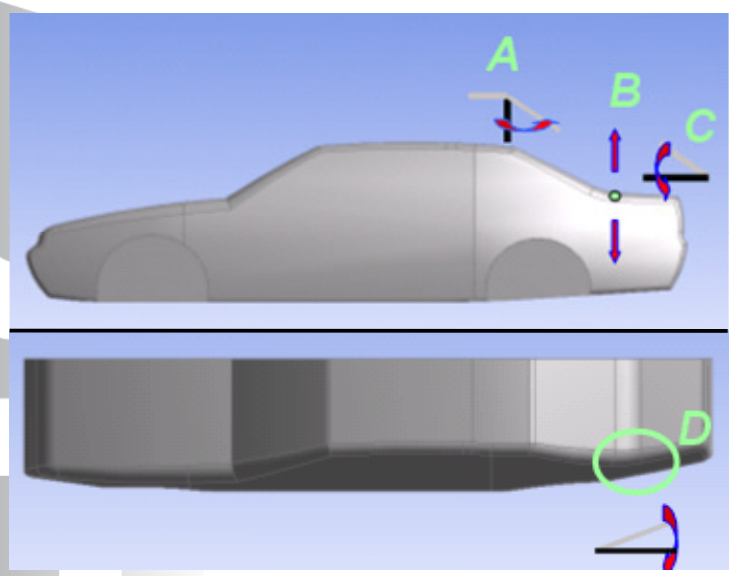

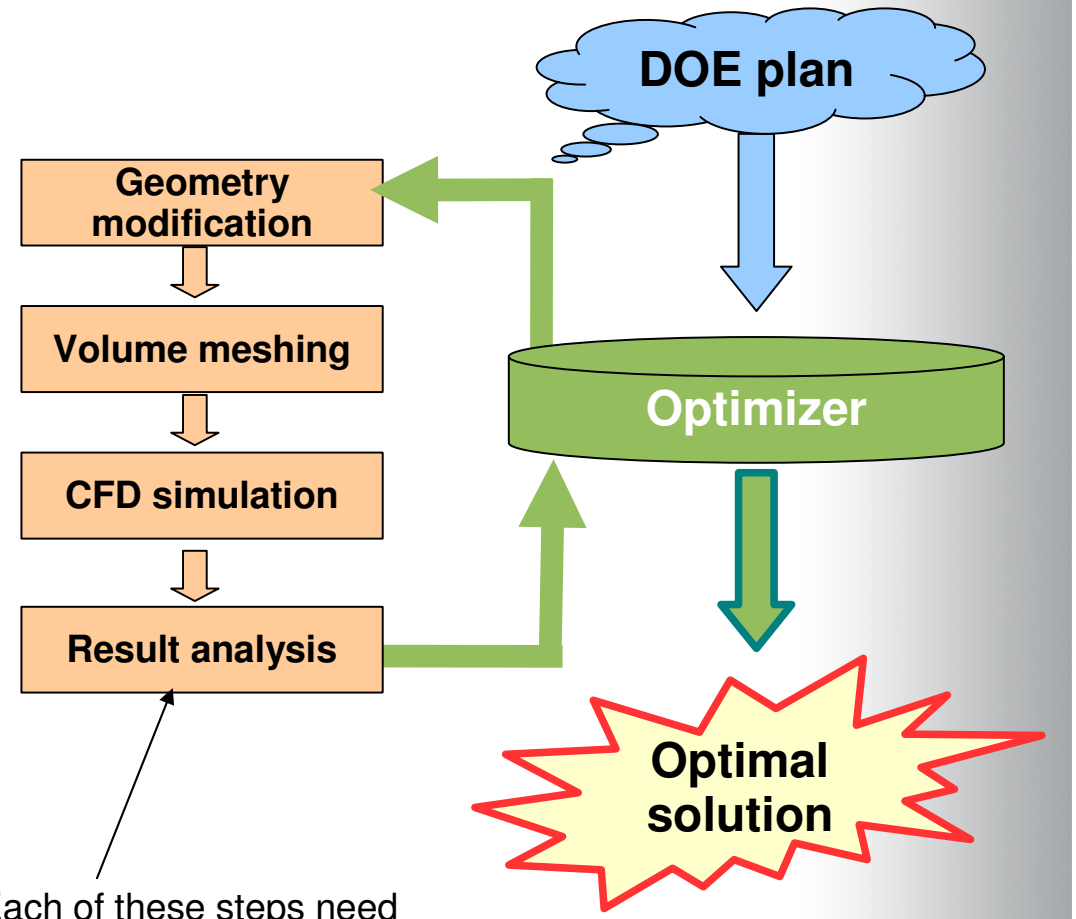

- 4 parameters to be optimized
- cubic face centered DOE
- 25 cases+16 extra cases for error estim.
- polynomial response function
- 70 hours wall clock time on 64 cpus

Each of these steps need to be fully automated and controlled by the optimizer

Combustion / Radiation

- Study the impact of radiative heat transfer (RHT) on the combustion process (2D)
- Couple combustion (AVBP), the RHT (Rayon) codes and the pollutant formation (AVBP)
- Parallelization of the Rayon code and improvement of the coupling part
- Load balancing issue
- 3D extension proposed to DECI and accepted

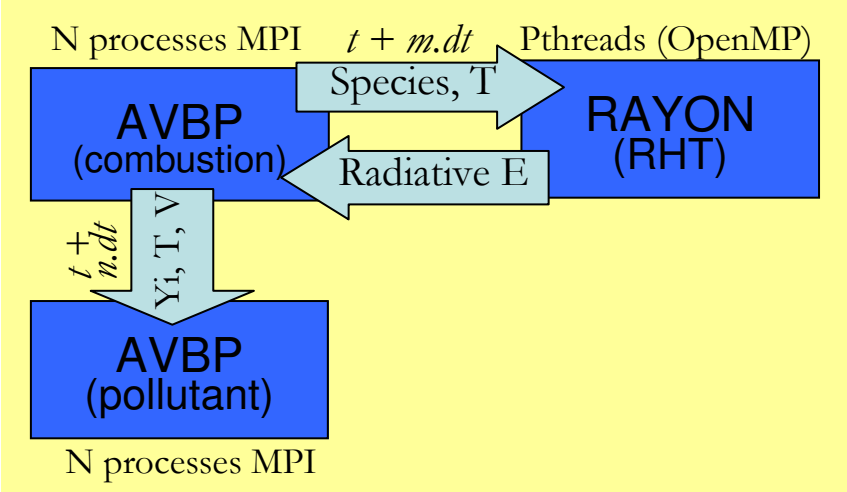

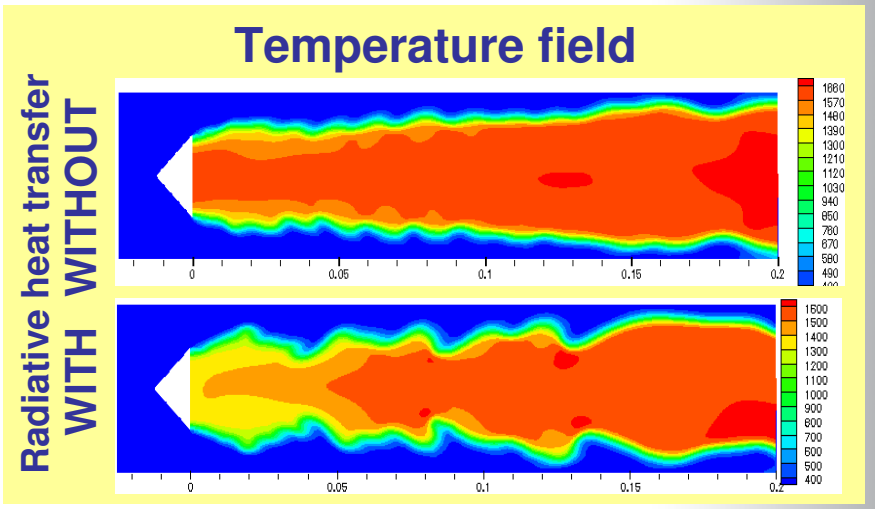

Climate Modelling

Distributed European Infrastructure for Supercomputing Applications

Climate research moves towards new levels of complexity:

Stepping from Climate (=Atmosphere+Ocean) to Earth System Modelling

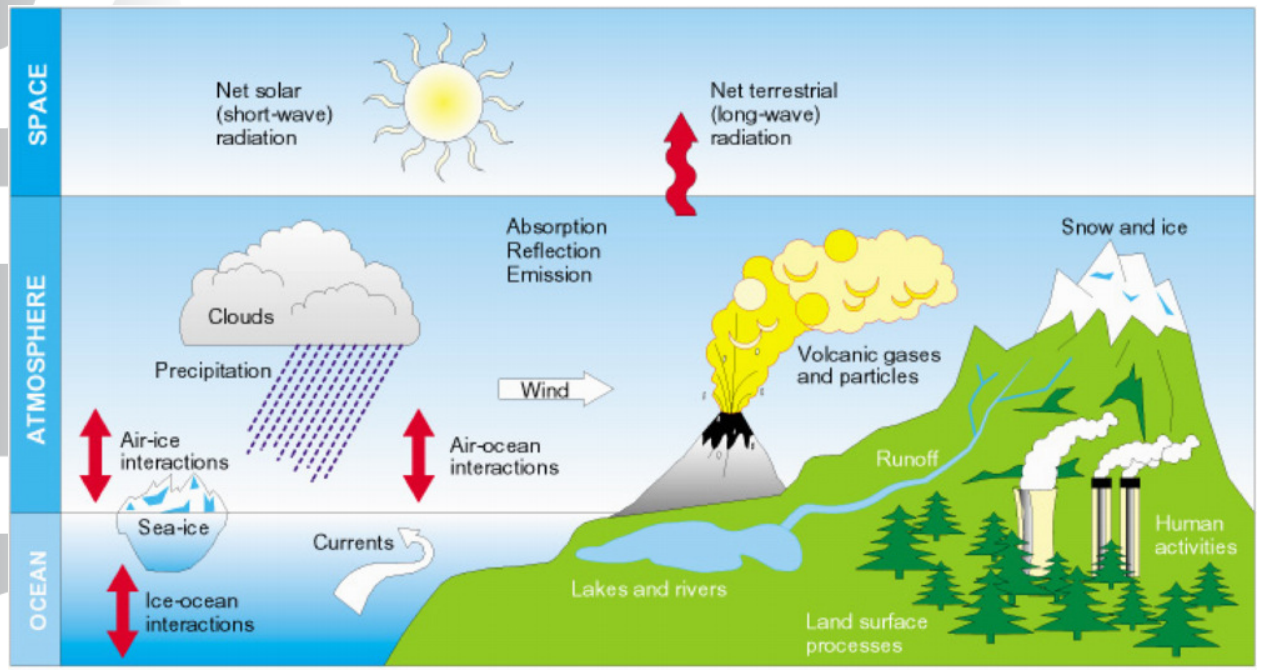

Earth system model wishlist:

Higher spatial and temporal resolution

Quality: Improved subsystem models

Atmospheric chemistry (ozone, sulfates,..)

Bio-geochemistry (Carbon cycle, ecosystem dynamics,..)

Increased Computational demand factor: O(1000 -10000)

AWI

Climate Research
Statistics of Climate Variability

Distributed European Infrastructure for Supercomputing Applications

Project to study climate trends, each 50 TB output data

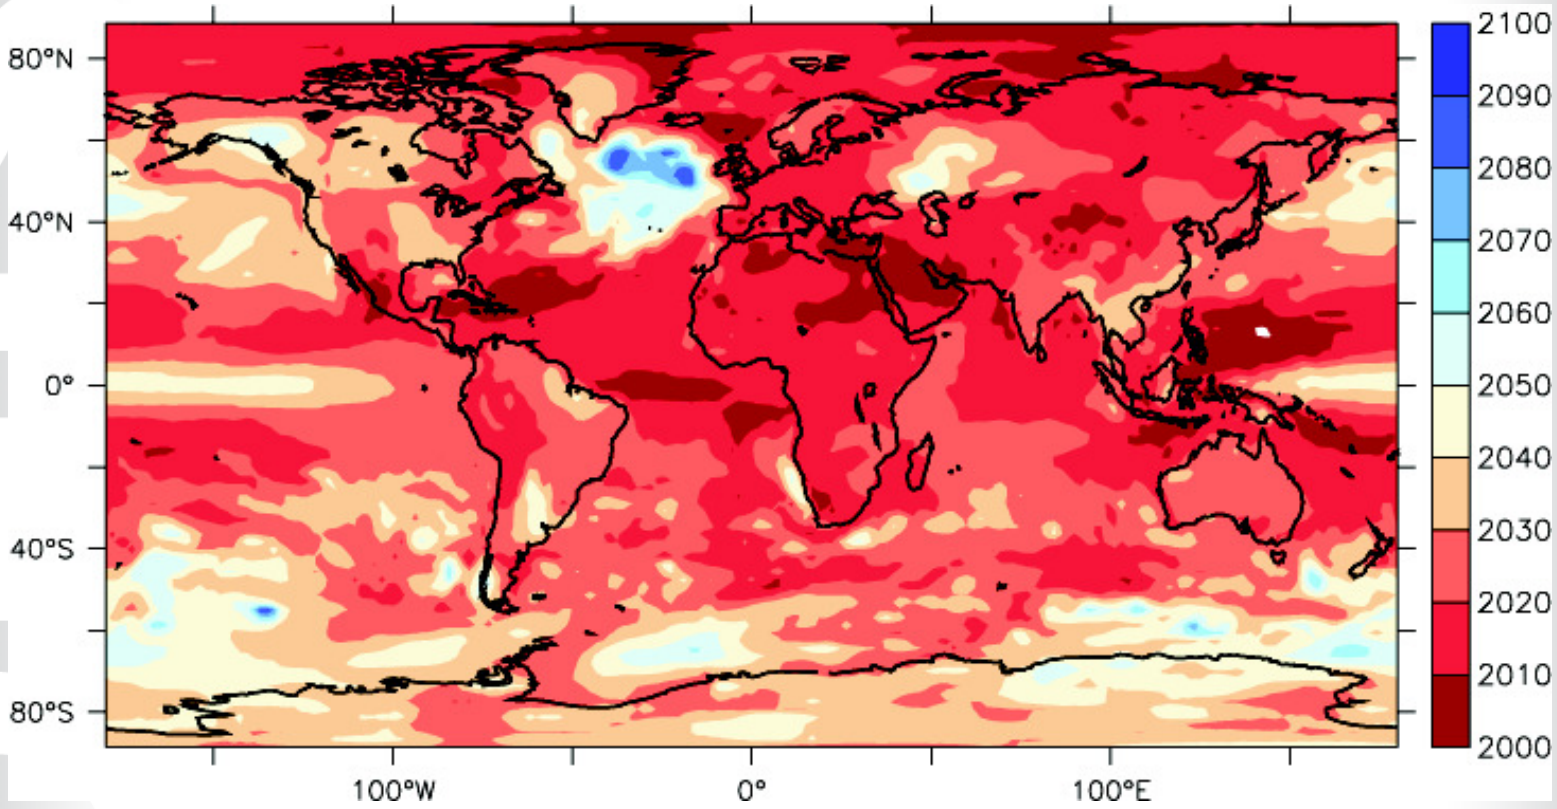

H. Dijkstra, U Utrecht, and W. Hazeleger, KNMI

Environmental Application

- Study the impact of water cycles of the hydrological and vegetation models on climate models
- Coupling area in West Africa
- Best performances with a vector and scalar platform
- Improve extensibility of the architecture and the coupling part
- AMMA project, PhD thesis, 2 publ. and 2 comms.

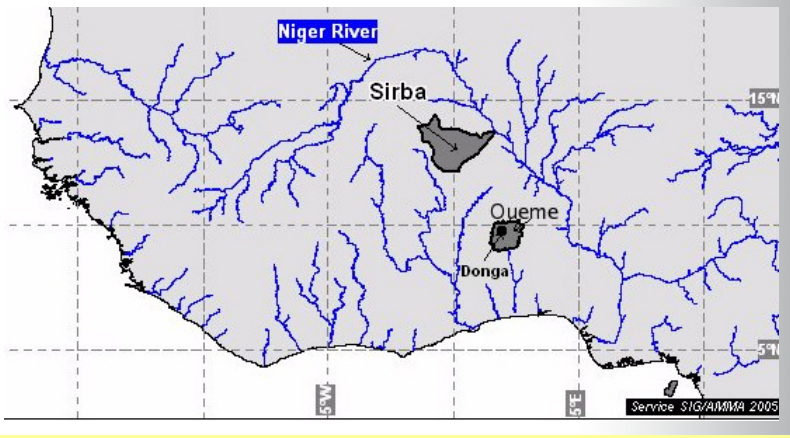

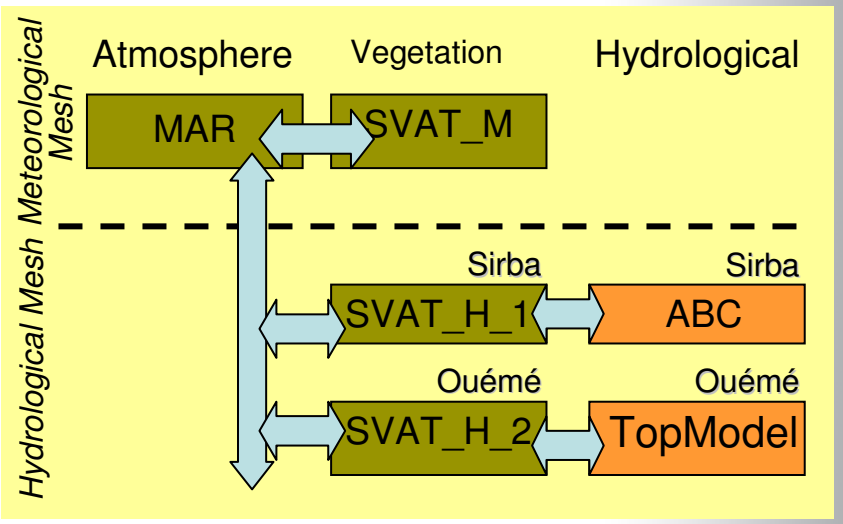

Materials Science

First-principles statistical mechanics for molecular switches at surfaces (MolSwitch)

Distributed European Infrastructure for Supercomputing Applications

Azobenzene on copper, silver and gold surfaces

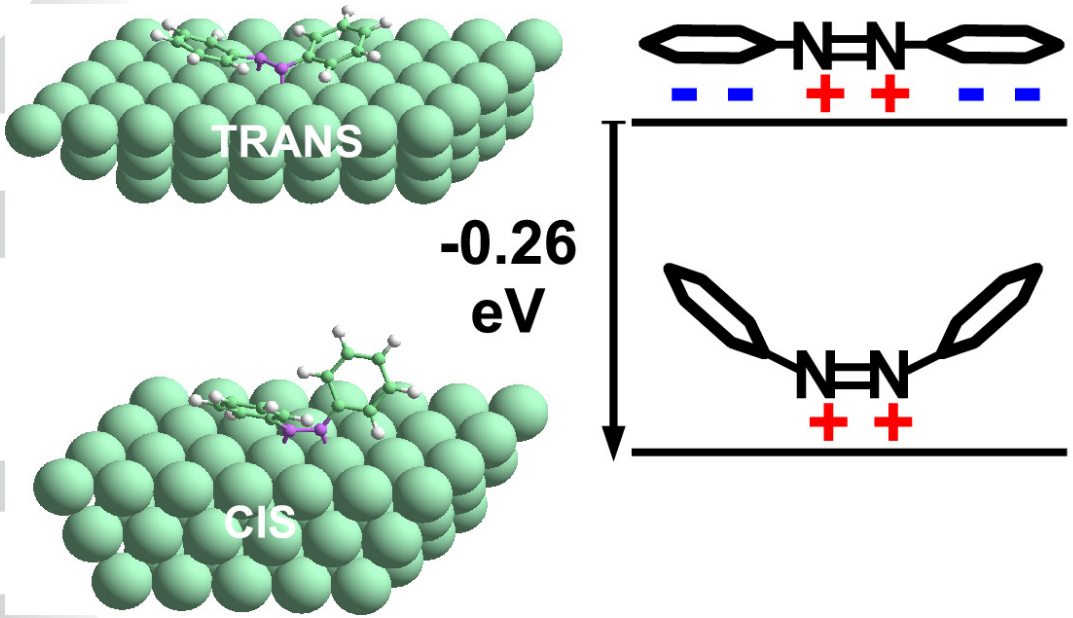

Controlled reversible switching should be possible on Ag surfaces

Courtesy: K. Reuter, FHI

Polymer Research

Distributed European Infrastructure for Supercomputing Applications

Cover story of Nature - May 24, 2007

Curvy membranes make proteins attractive

For almost two decades, physicists have been on the track of membrane mediated interactions. Simulations in DEISA have now revealed that curvy membranes make proteins attractive

Nature 447 (2007), 461-464

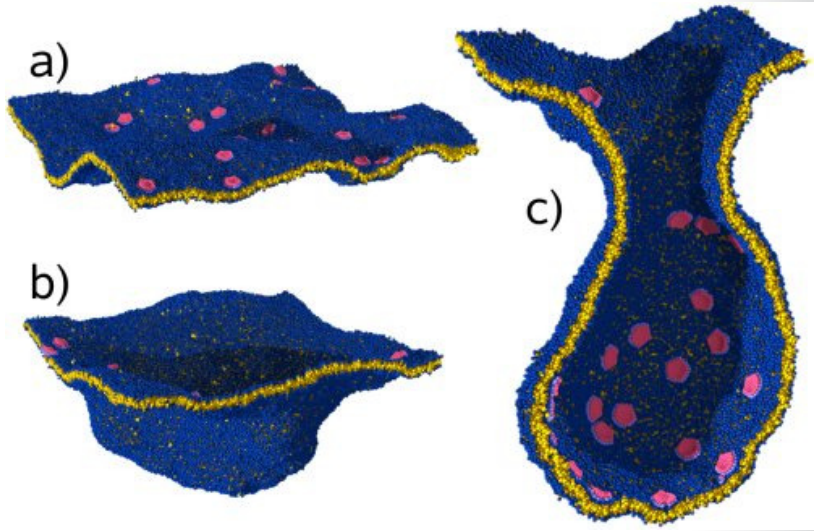

a) proteins (red) adhere on a membrane (blue/yellow) and locally bend it;

b) this triggers a growing invagination.

c) cross-section through an almost complete vesicle

B. J. Reynwar et al.: ***Aggregation and vesiculation of membrane proteins by curvature mediated interactions***, NATURE Vol 447|24 May 2007| doi:10.1038/nature05840

Cosmology Project

- Study galaxy formation in cosmology
- Physics / modules: Gravitation, Hydrodynamics, Chemistry
- Best performance on heterogeneous platforms
- Load balancing issue and improvement of the coupling part
- Proposed to DECI

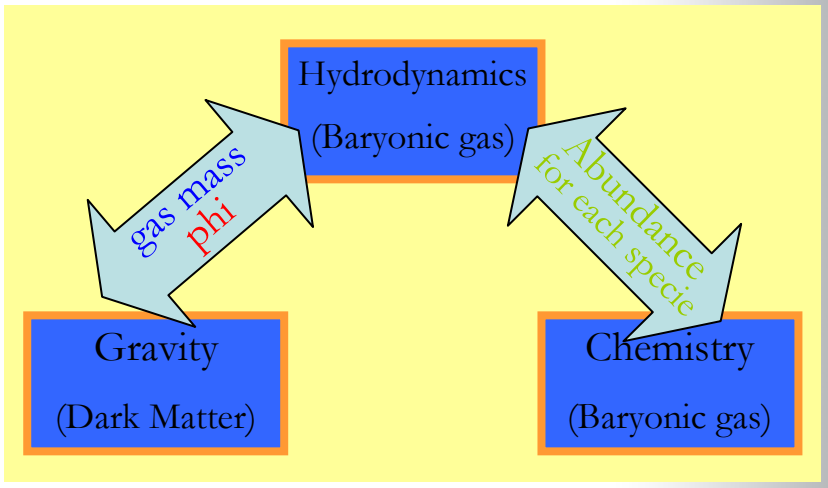

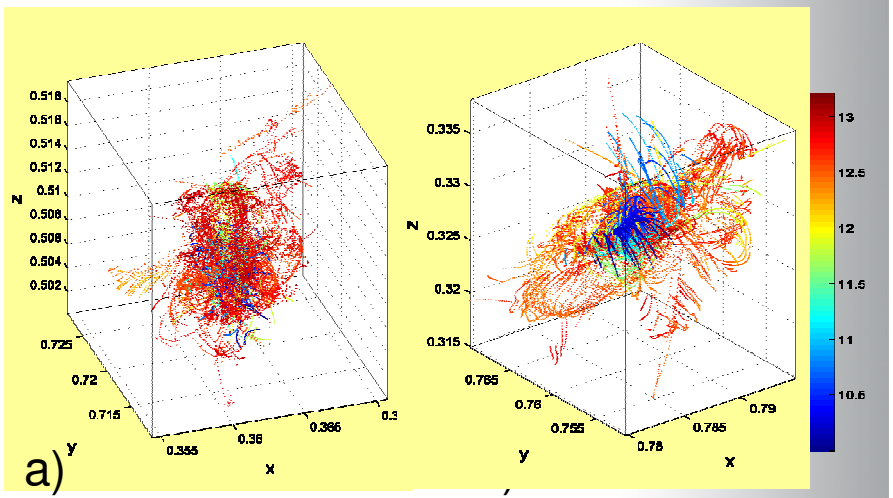

PLASMA Physics
Theory Support of ITER Experiment

Distributed European Infrastructure for Supercomputing Applications

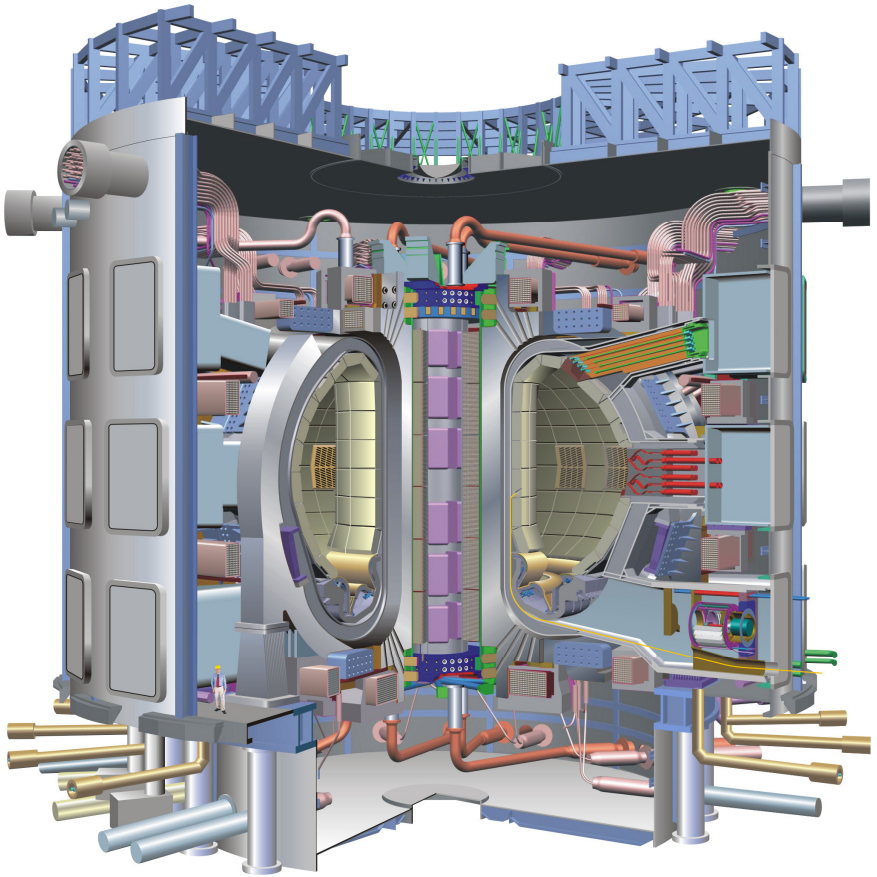

DECI Project GYROKINETICS

Distributed European Infrastructure for Supercomputing Applications

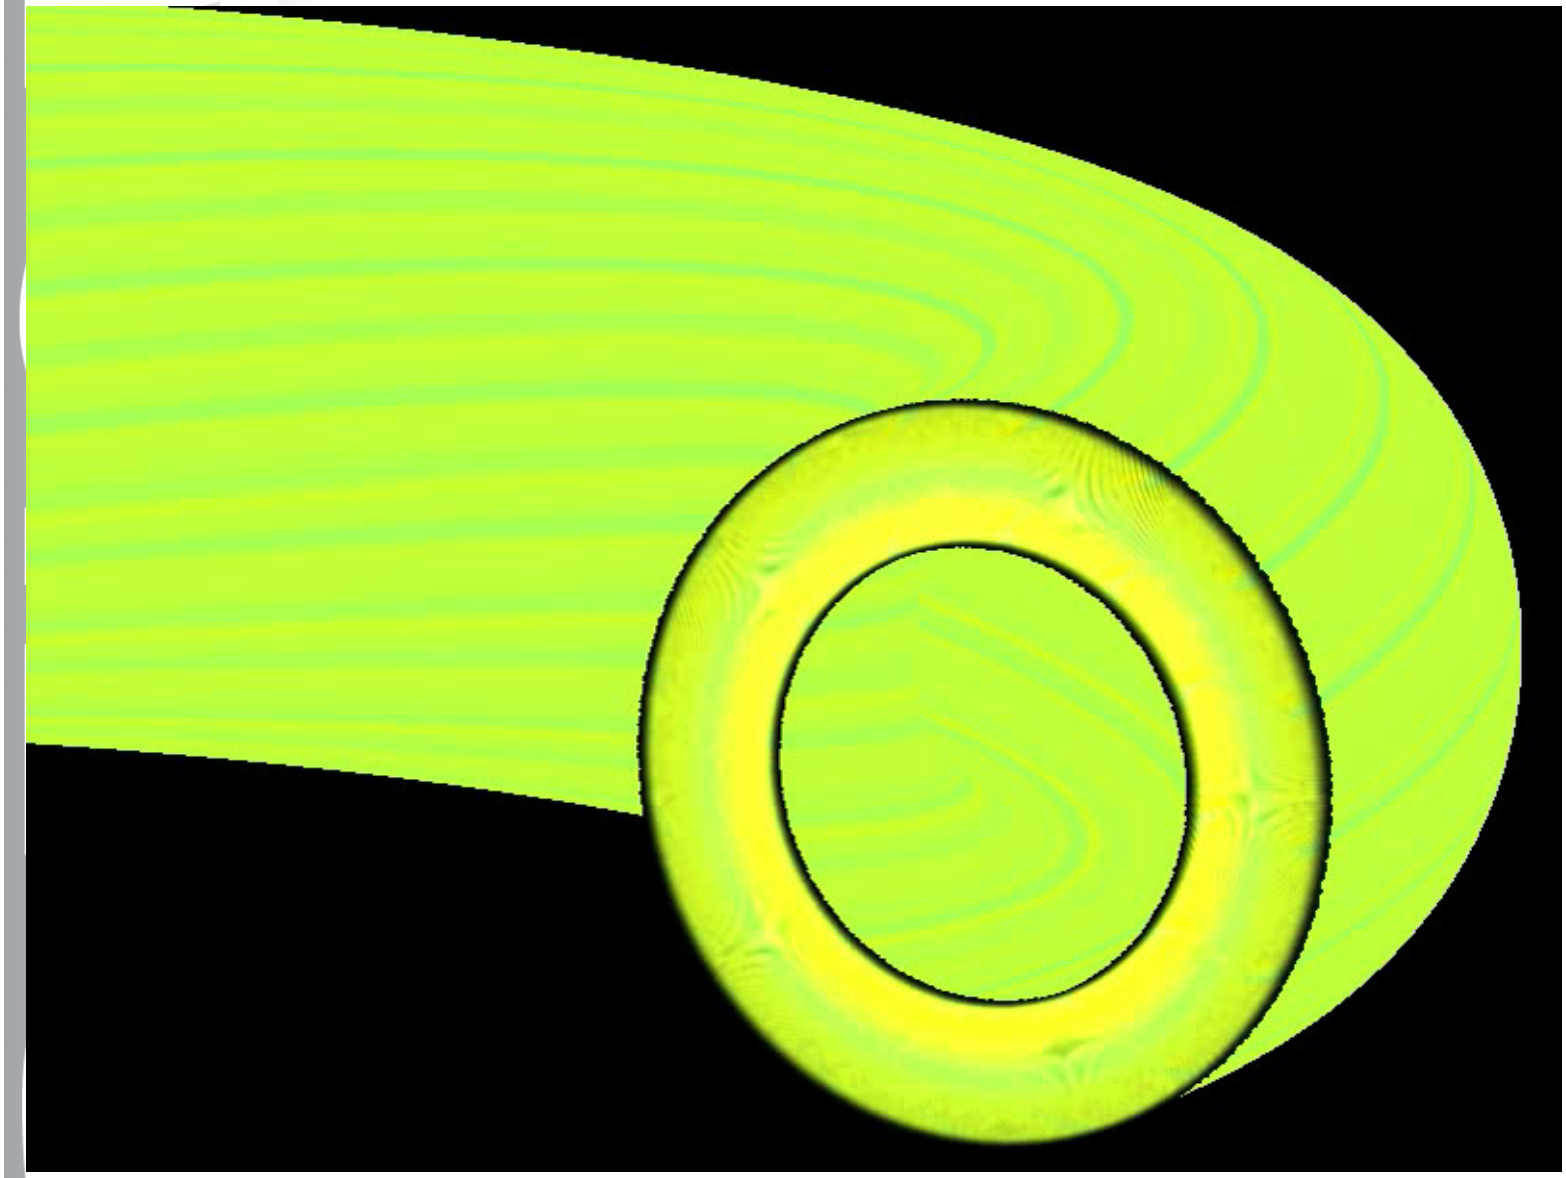

Core-Plasma Turbulence Simulations with GENE

MPI for Plasma Physics

Distributed European Infrastructure for Supercomputing Applications

Thank You for your attention

Gentzsch @ rzg.mpg.de