Distributed European Infrastructure for Supercomputing Applications

Evolvement of US and European HPC Infrastructures:

DEISA and its contributions

www.deisa.eu

Hermann Lederer

DEISA and RZG, Max Planck Society
lederer@rzg.mpg.de

OGF25 Catania March 2-6, 2009

RI-222919

Distributed European Infrastructure for Supercomputing Applications

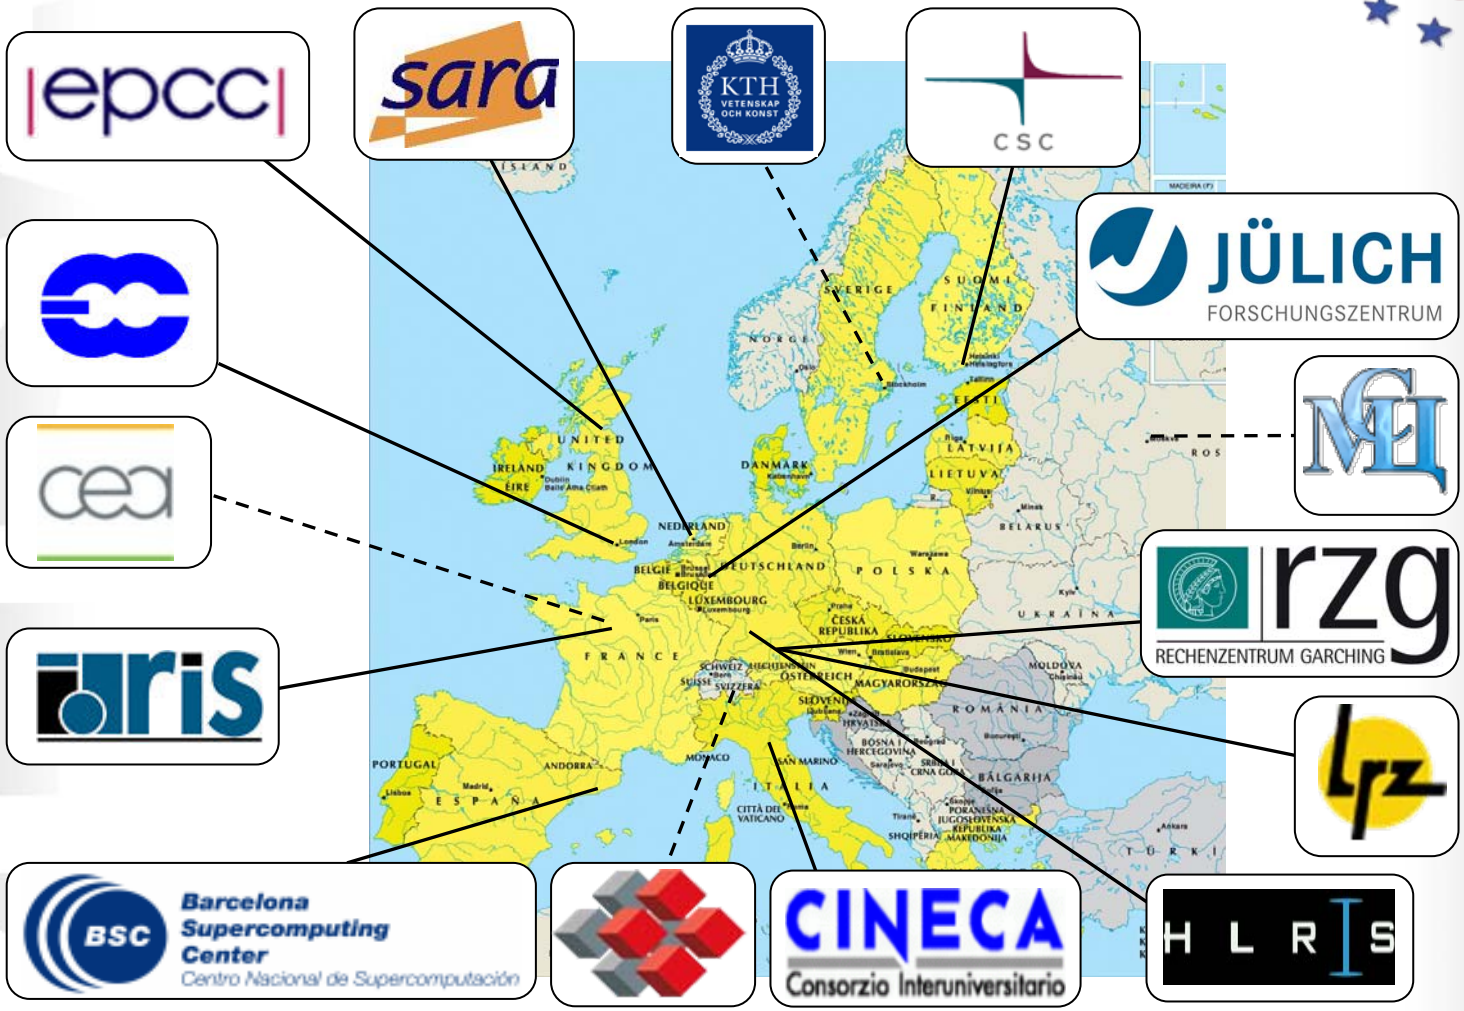

DEISA Consortium and Associate Partners

RI-222919

Distributed European Infrastructure for Supercomputing Applications

Evolvement ???

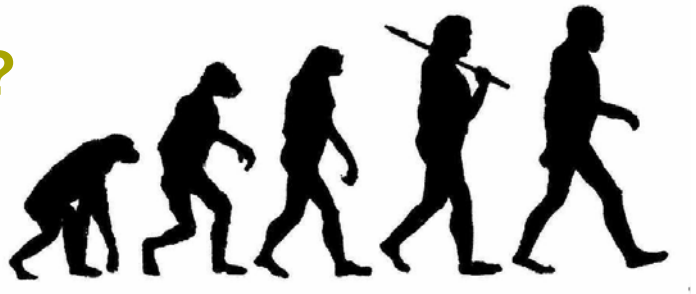

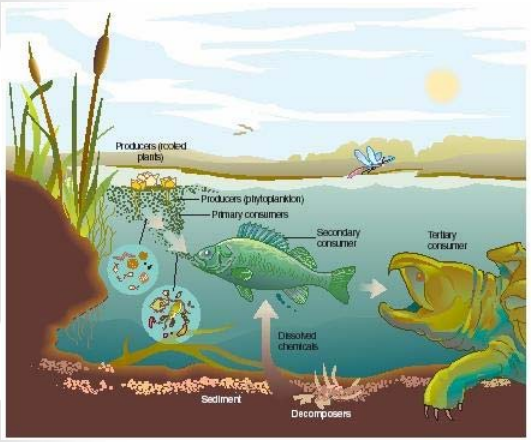

Ecosystem ???

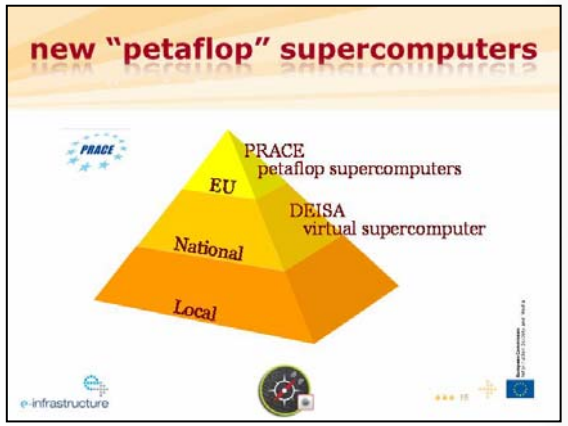

HPC Pyramide ???

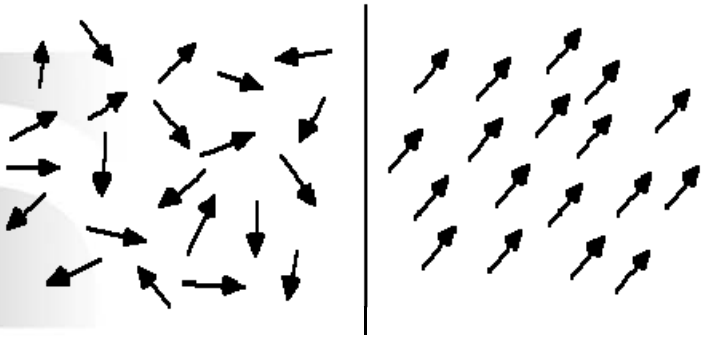

Standardization ???

RI-222919

new "petaflop" supercomputers

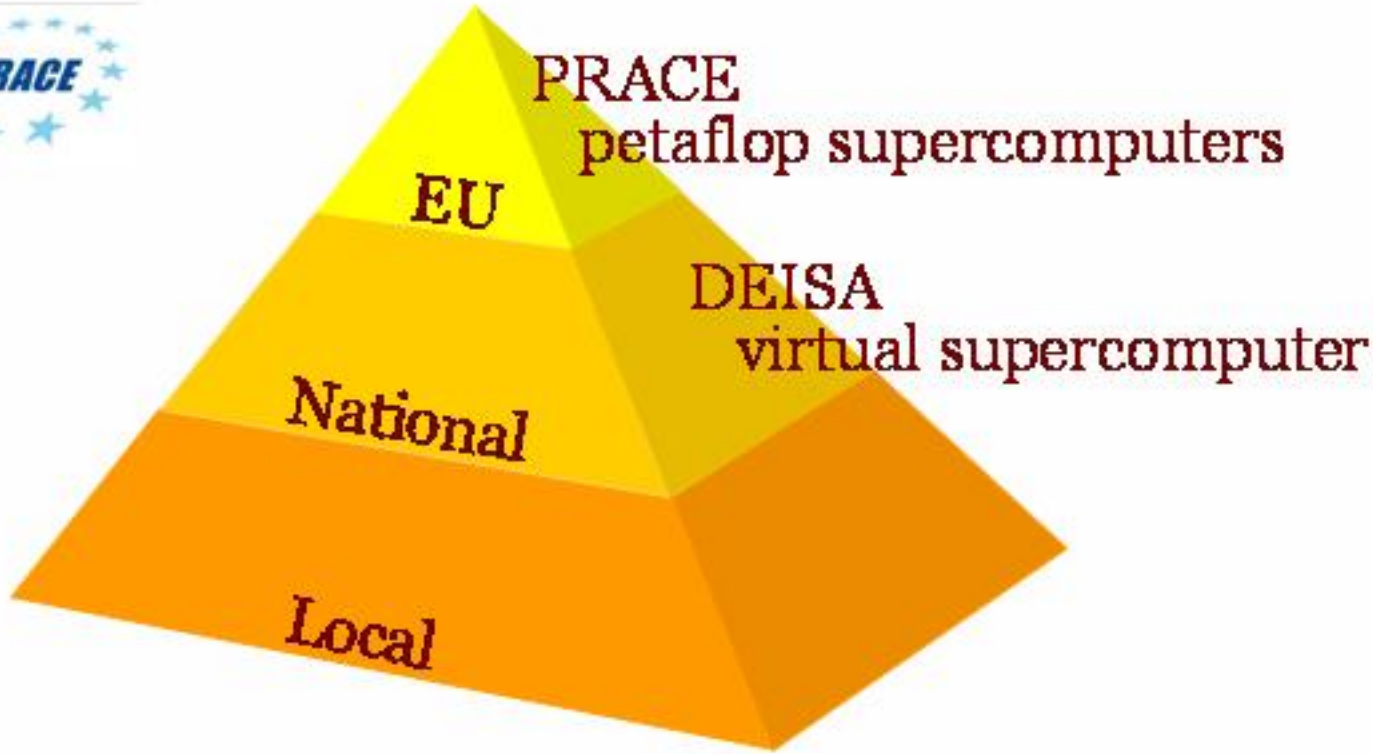

Mario Campolargo
European Commission
OGF23, June 2008

Steps Towards an Integrated European HPC Infrastructure

Distributed European Infrastructure for Supercomputing Applications

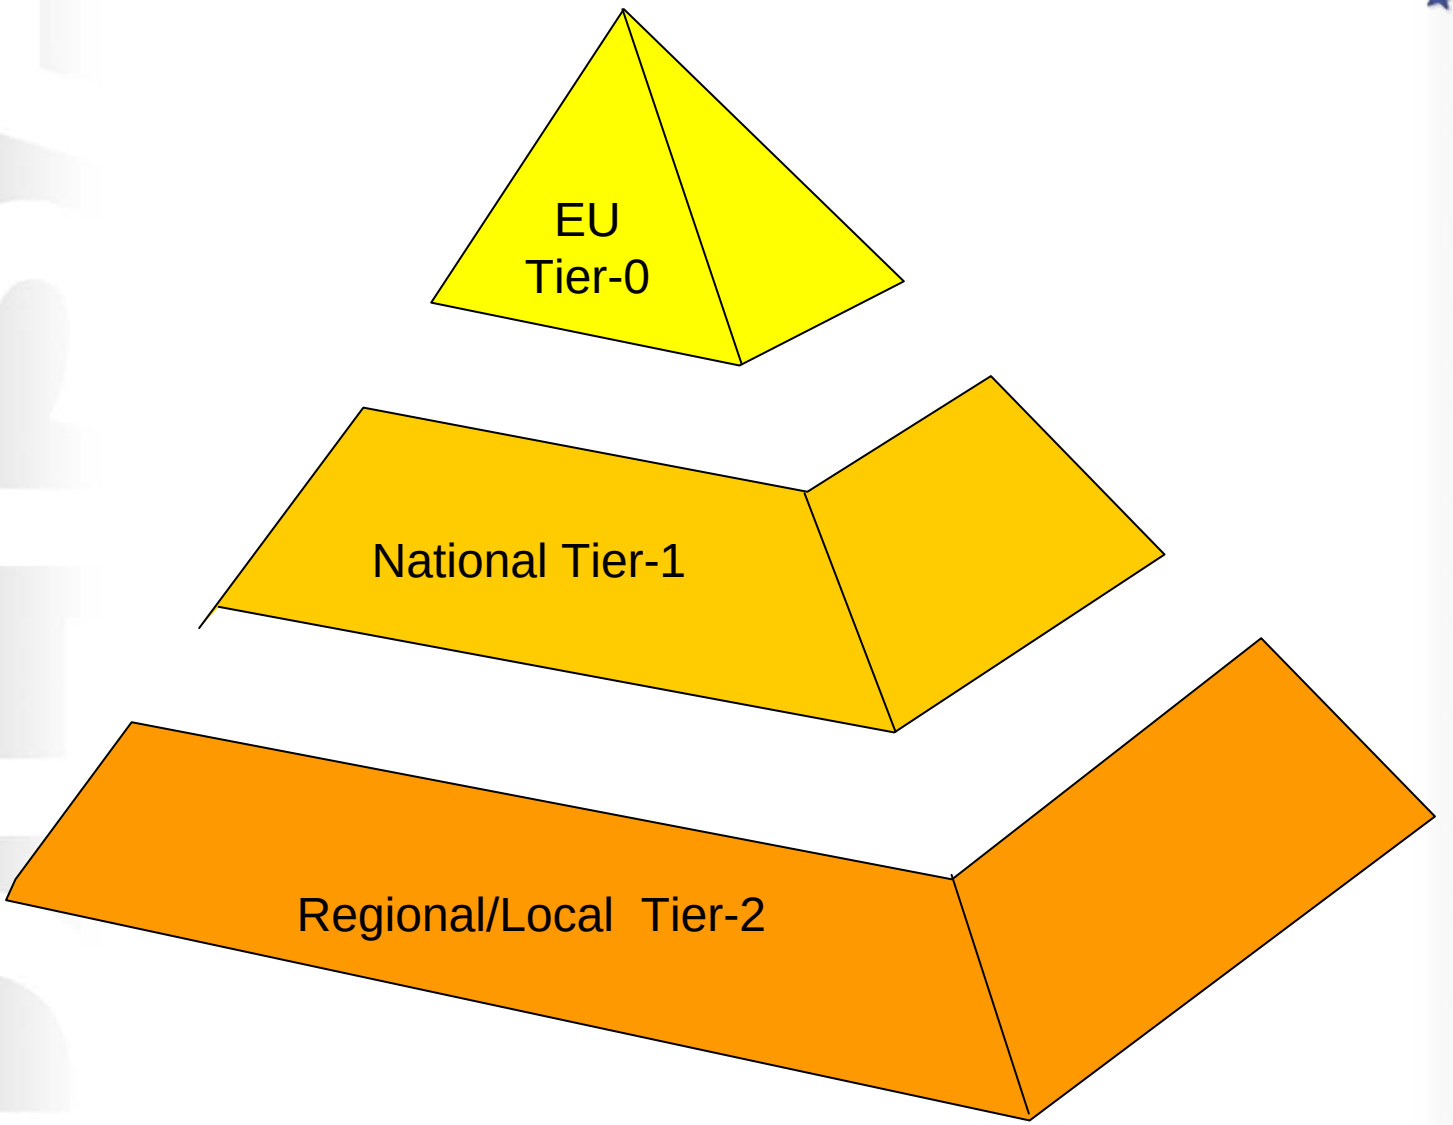

Steps Towards an Integrated European HPC Infrastructure

Distributed European Infrastructure for Supercomputing Applications

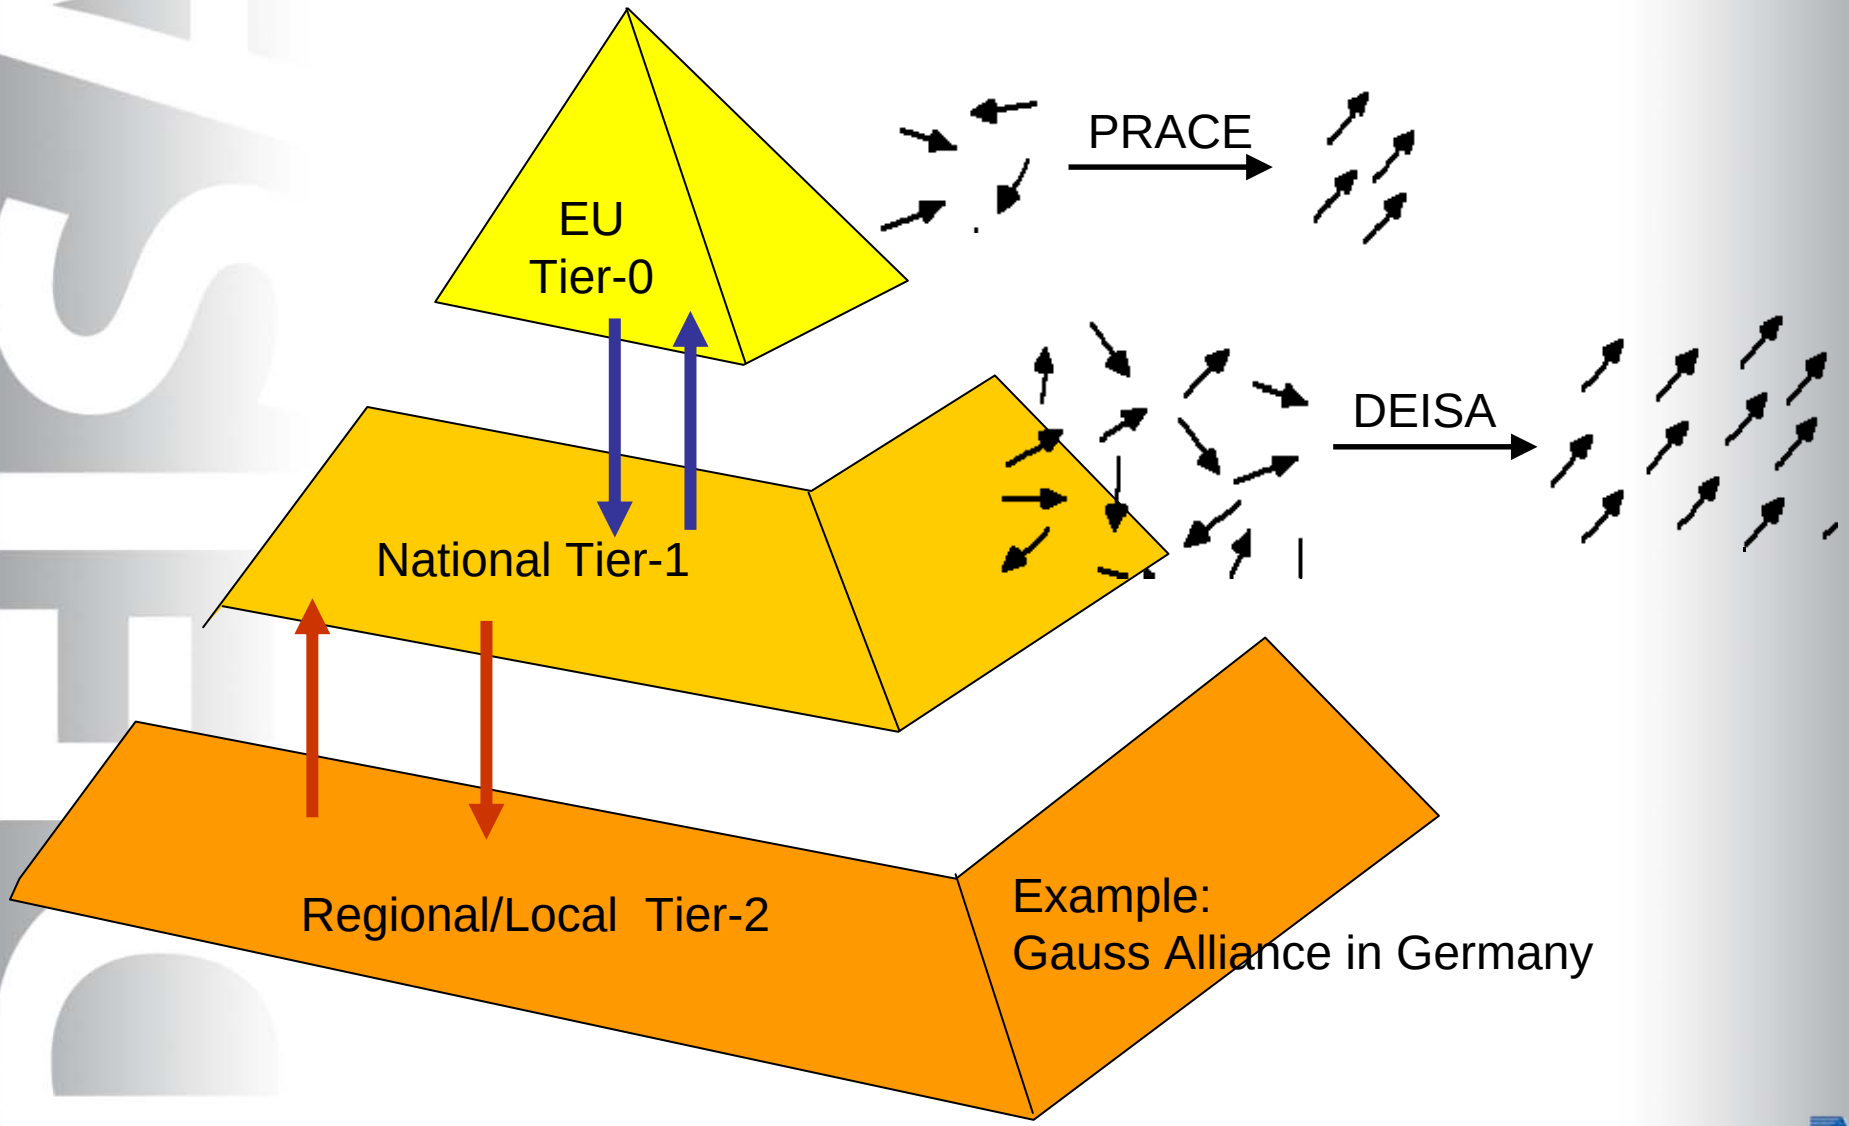

RI-222919

Tier0 / Tier1 Top Layer of the HPC Ecosystem

Distributed European Infrastructure for Supercomputing Applications

T0 : future shared petascale European systems
T1 : leading national supercomputing systems

PRACE
Designing an infrastructure that will enable the operation of shared petascale European systems
Enhancing performance in selected sites and providing wide access to shared systems

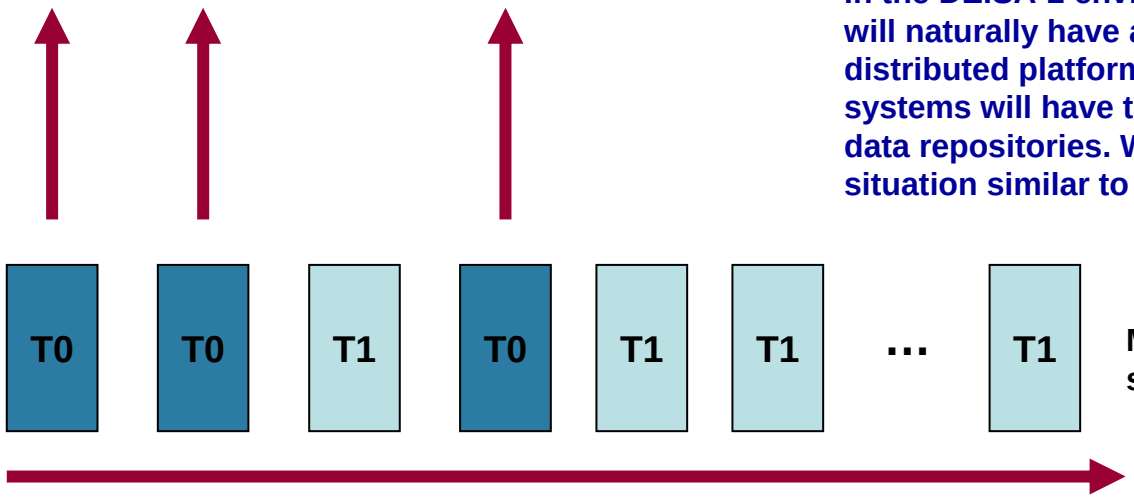

In the DEISA-2 environment, scientists will naturally have access to several distributed platforms, and shared systems will have to access remote data repositories. We will be in a situation similar to TeraGrid.

DEISA-2 : strong integration of T0 and T1 systems (automatically provides wide, seamless and efficient access to shared systems and data repositories)

The DEISA-1 services have been tailored for this mode of operation. There is a positive feedback between the two orthogonal lines of action:
- ***DEISA is paving the way to the efficient operation of T0 systems.***
- ***T0 systems will drive the massive adoption of the DEISA services.***

RI-222919

DEISA Partners & PRACE

Distributed European Infrastructure for Supercomputing Applications

DEISA Center	Country	PRACE
BSC	Spain	x
CINECA	Italy	x
CSC	Finland	x
EPCC	UK	x
ECMWF	UK	
FZJ	Germany	x
HLRS	Germany	x
IDRIS-CNRS	France	x
LRZ	Germany	x
RZG	Germany	
SARA	The Netherlands	x
CEA-CCRT	France	x
CSCS	Switzerland	x
KTH	Sweden	x
JSCC	Russia	

***X** means:*

represented in the organisation being the PRACE Member

RI-222919

Distributed European Infrastructure for Supercomputing Applications

DEISA Supercomputers

State-of-the art supercomputers > 1 PF aggregated peak performance

- Cray XT4/5, Linux
- IBM Power5, Power6, AIX / Linux
- IBM BlueGene/P, Linux (frontend)
- IBM PowerPC, Linux (MareNostrum)
- SGI ALTIX 4700 (Itanium2 Montecito), Linux
- NEC SX8/9 vector systems, Super UX

- Fixed fractions of resources dedicated to DEISA usage
- Systems interconnected with dedicated 10Gb/s network

RI-222919

DEISA dedicated high speed network

Distributed European Infrastructure for Supercomputing Applications

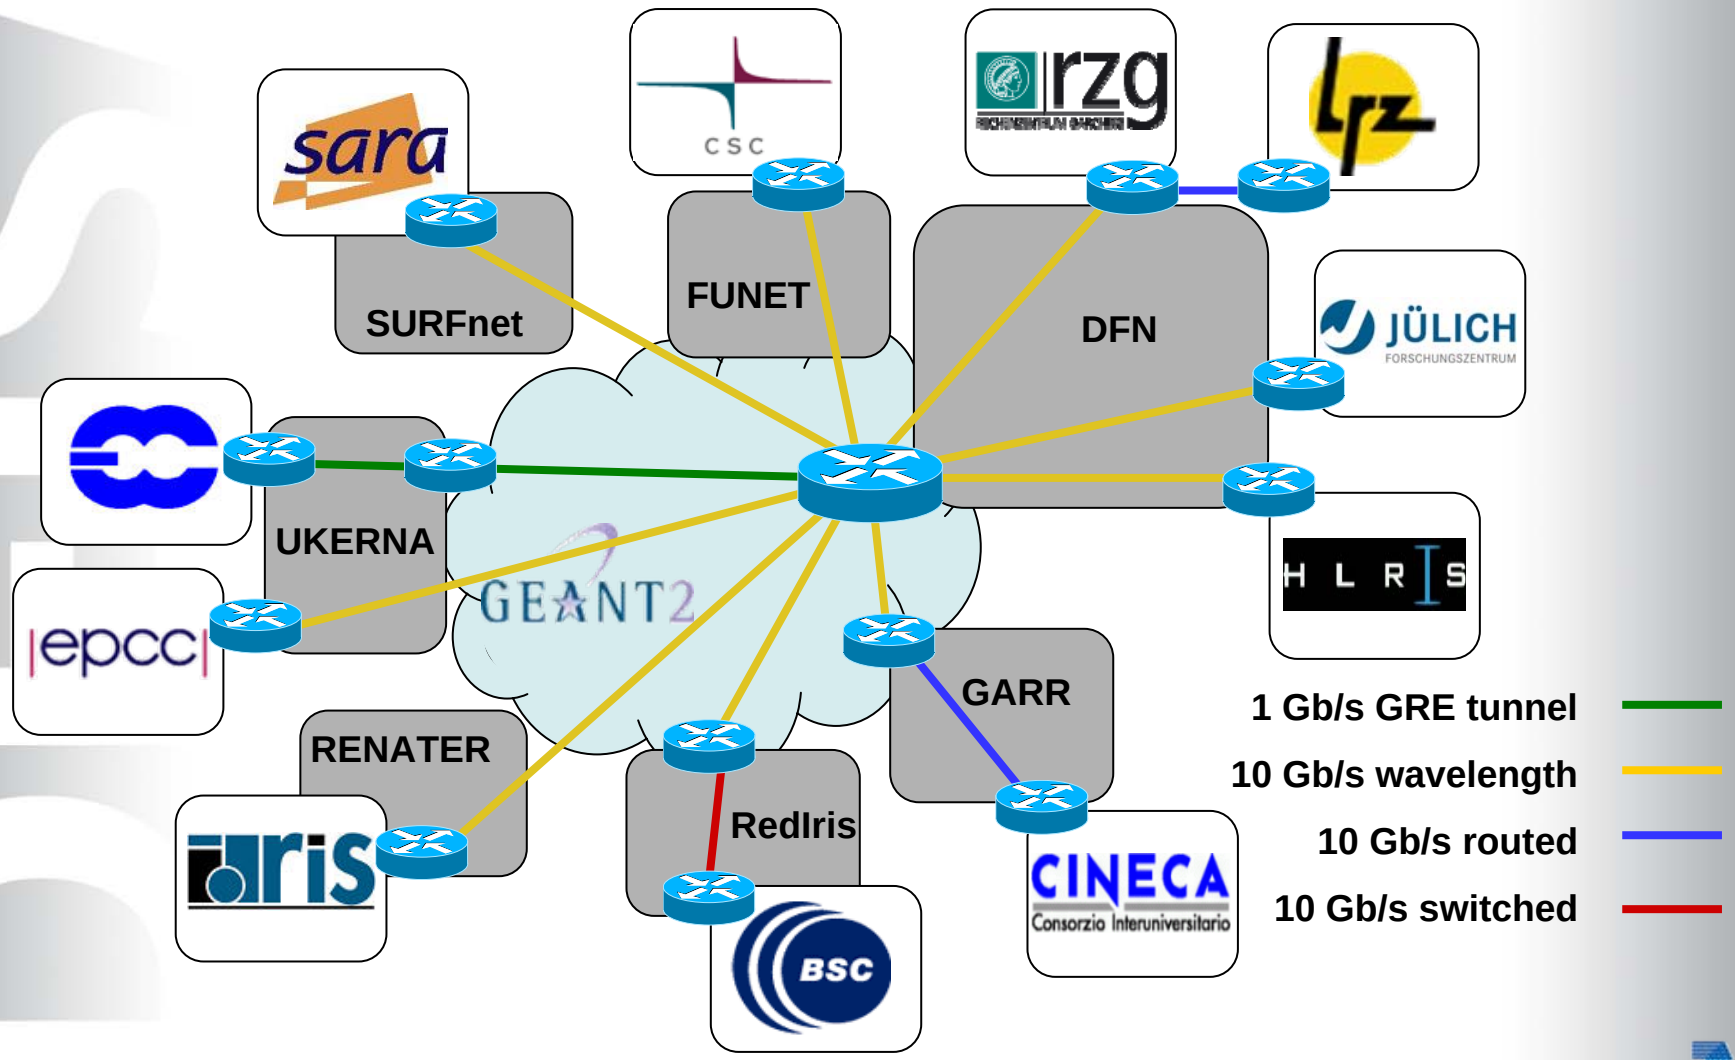

RI-222919

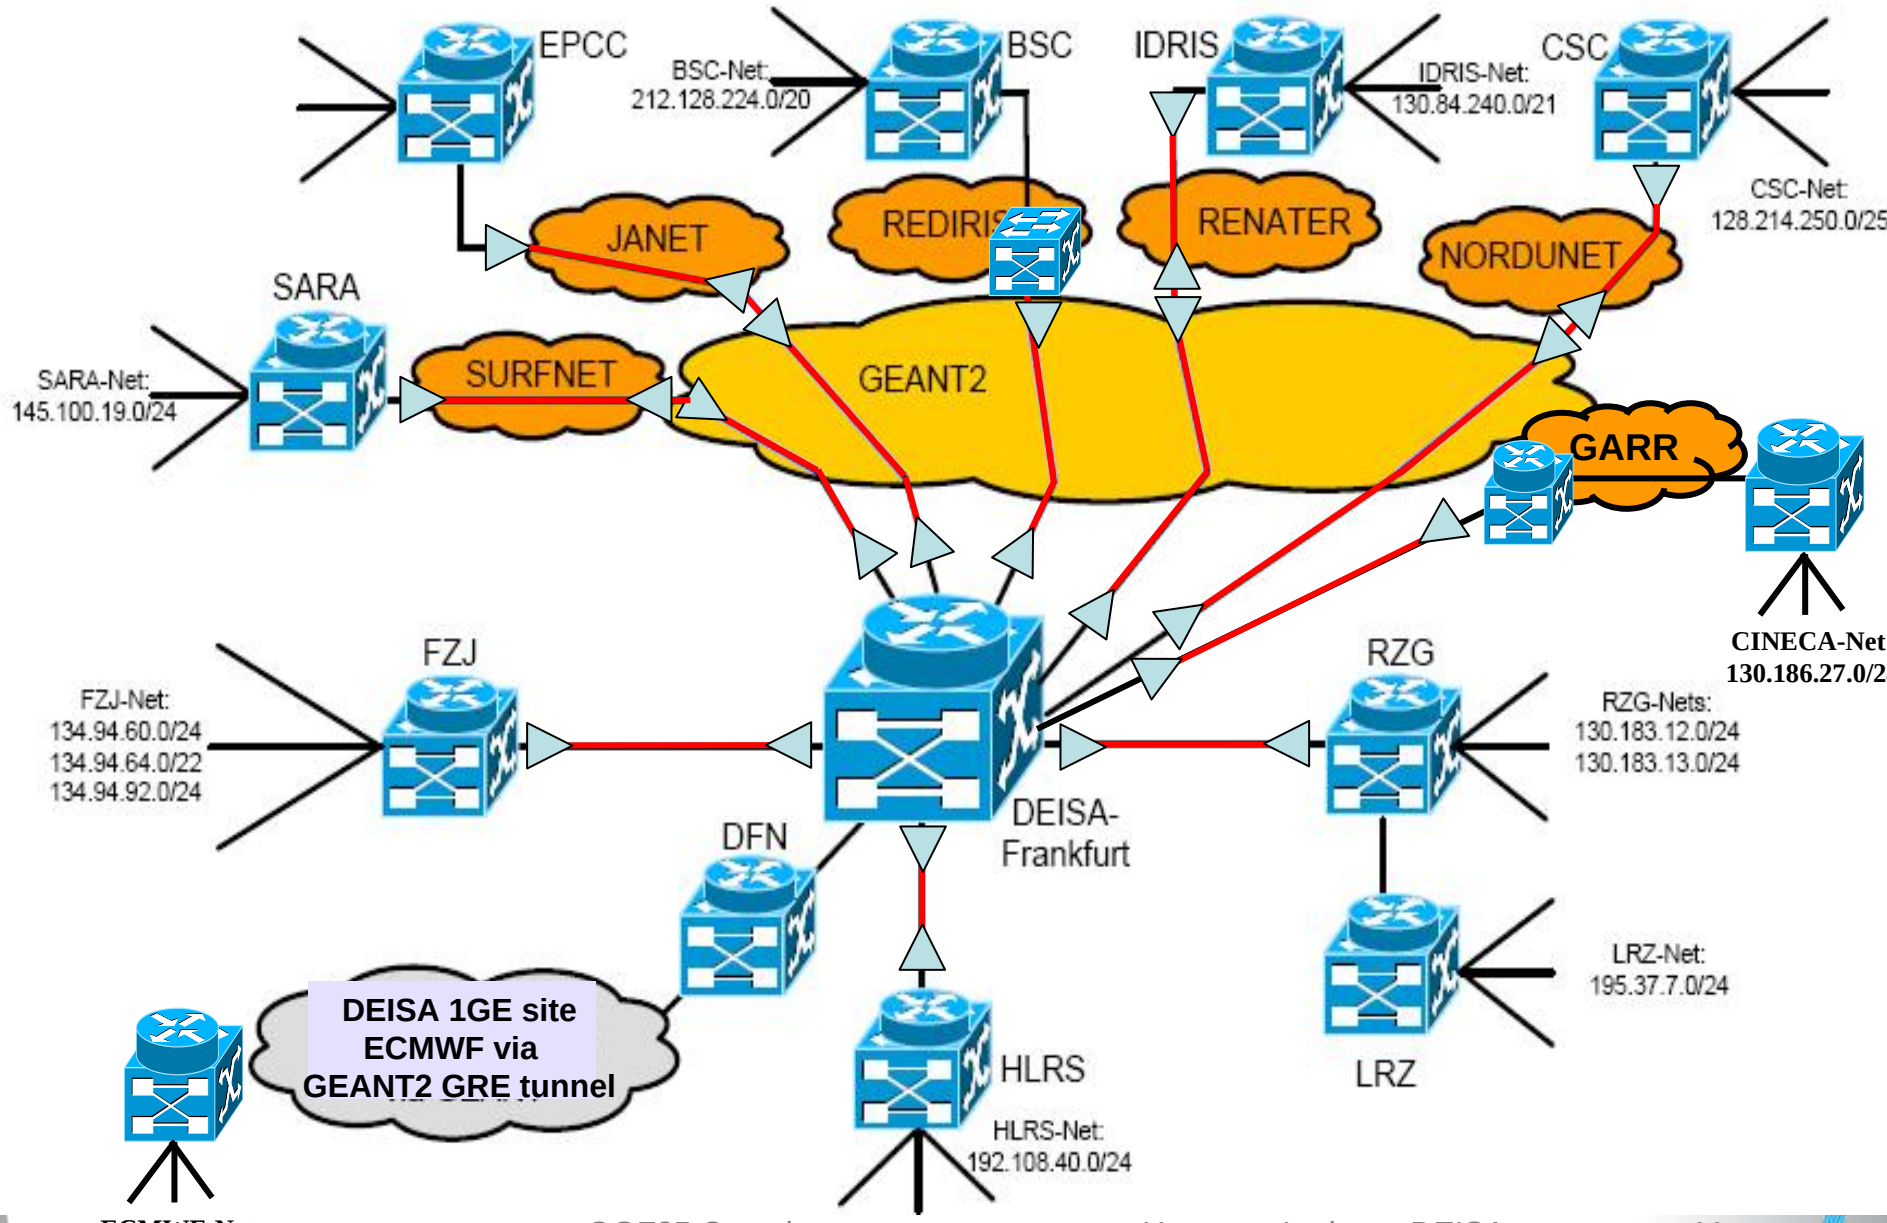
EPCC
BSC
BSC-Net:
212.128.224.0/20
IDRIS
IDRIS-Net:
130.84.240.0/21
CSC
CSC-Net:
128.214.250.0/25
JANET
REDIRIS
RENATER
NORDUNET
SARA
SARA-Net:
145.100.19.0/24
SURFNET
GEANT2
GARR
CINECA-Net:
130.186.27.0/24
FZJ
FZJ-Net:
134.94.60.0/24
134.94.64.0/22
134.94.92.0/24
RZG
RZG-Nets:
130.183.12.0/24
130.183.13.0/24
DEISA-
Frankfurt
DFN
LRZ-Net:
195.37.7.0/24
LRZ
DEISA 1GE site
ECMWF via
GEANT2 GRE tunnel
HLRS
HLRS-Net:
192.108.40.0/24
ECMWF-Net:
136.156.48.0/24

DEISA Global File System at Continental Scale
(based on MC-GPFS)

Distributed European Infrastructure for Supercomputing Applications

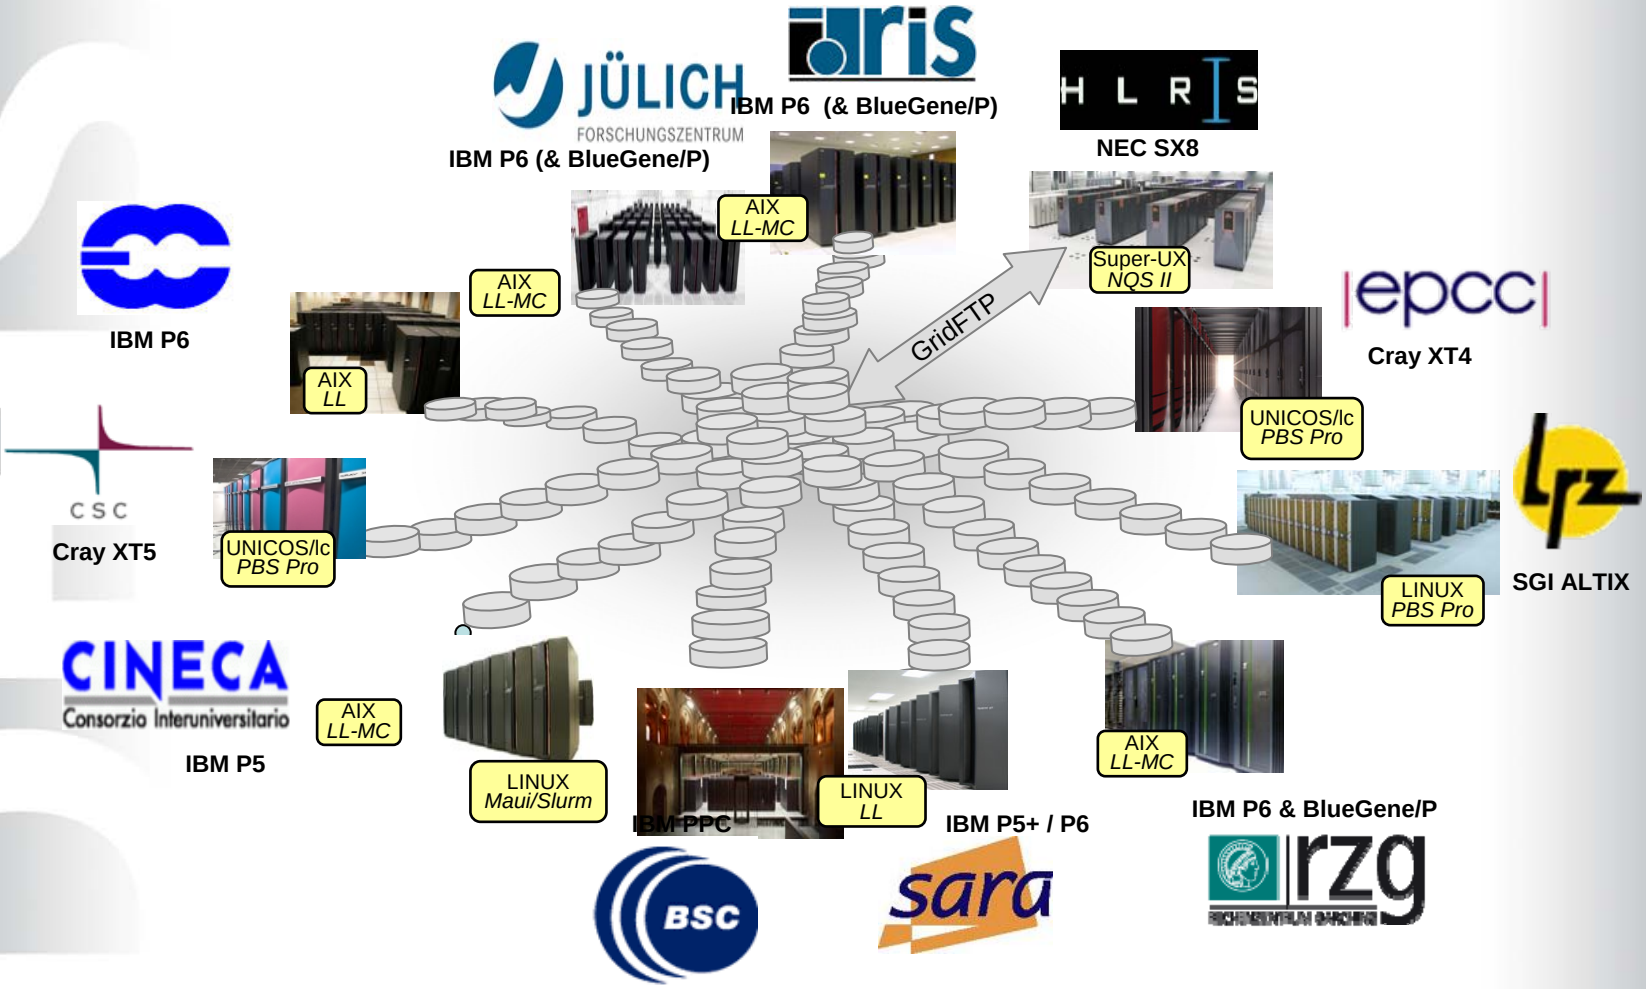

RI-222919

Overall GPFS Configuration

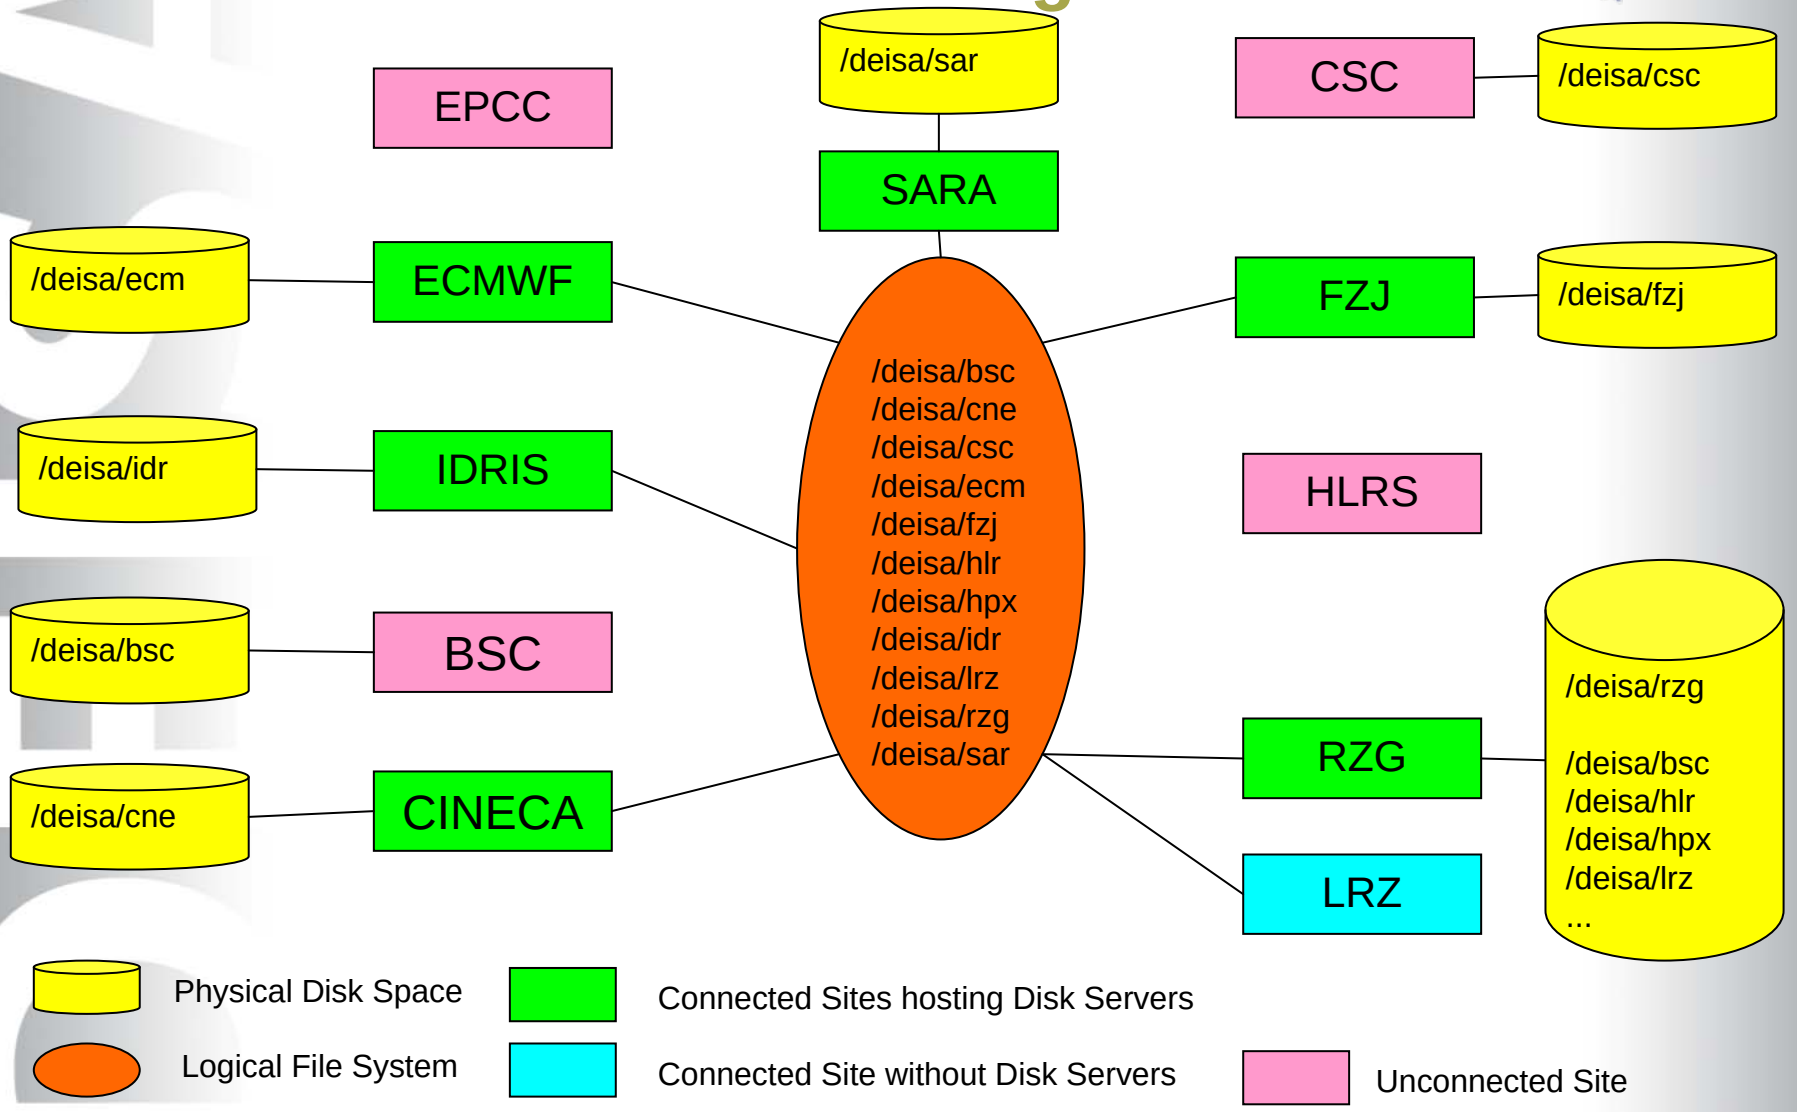

RI-222919

DEISA Extreme Computing Initiative (DECI)

Distributed European Infrastructure for Supercomputing Applications

- DECI launched in early 2005 for complex, demanding, innovative simulations requiring the exceptional capabilities of DEISA
- Multi-national proposals especially 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 in Europe for the most challenging projects
- Most appropriate supercomputer architecture selected for each project
- Mitigation of the rapid performance decay of a single national supercomputer within its short lifetime cycle of typically about 5 years, as implied by Moore's law

RI-222919

DEISA Extreme Computing Initiative

Distributed European Infrastructure for Supercomputing Applications

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

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+)

DECI call 2009 (opened March 1, 2009)

DEISA Extreme Computing Initiative

Distributed European Infrastructure for Supercomputing Applications

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

Involvement of ~ 160 research institutes and universities from 15 European countries

Austria	Finland	France	Germany	Hungary
Italy	Netherlands	Poland	Portugal	Romania
Russia	Spain	Sweden	Switzerland	UK

with collaborators from

four other continents

North America, South America, Asia, Australia

RI-222919

Virtual Community Support

Distributed European Infrastructure for Supercomputing Applications

2008 Started with community support for

Fusion energy research:

EFDA EUROPEAN FUSION DEVELOPMENT AGREEMENT **www.efda.org**

European Fusion Development Agreement

EUFORIA **www.euforia-project.eu**

EU Fusion fOR Iter Applications (EUFORIA)

Life Sciences:

VIROLAB **www.virolab.org**

Virtual laboratory for infectious diseases

2009 Extension of community support for further science areas

Evolution of User Categories in DEISA

Distributed European Infrastructure for Supercomputing Applications

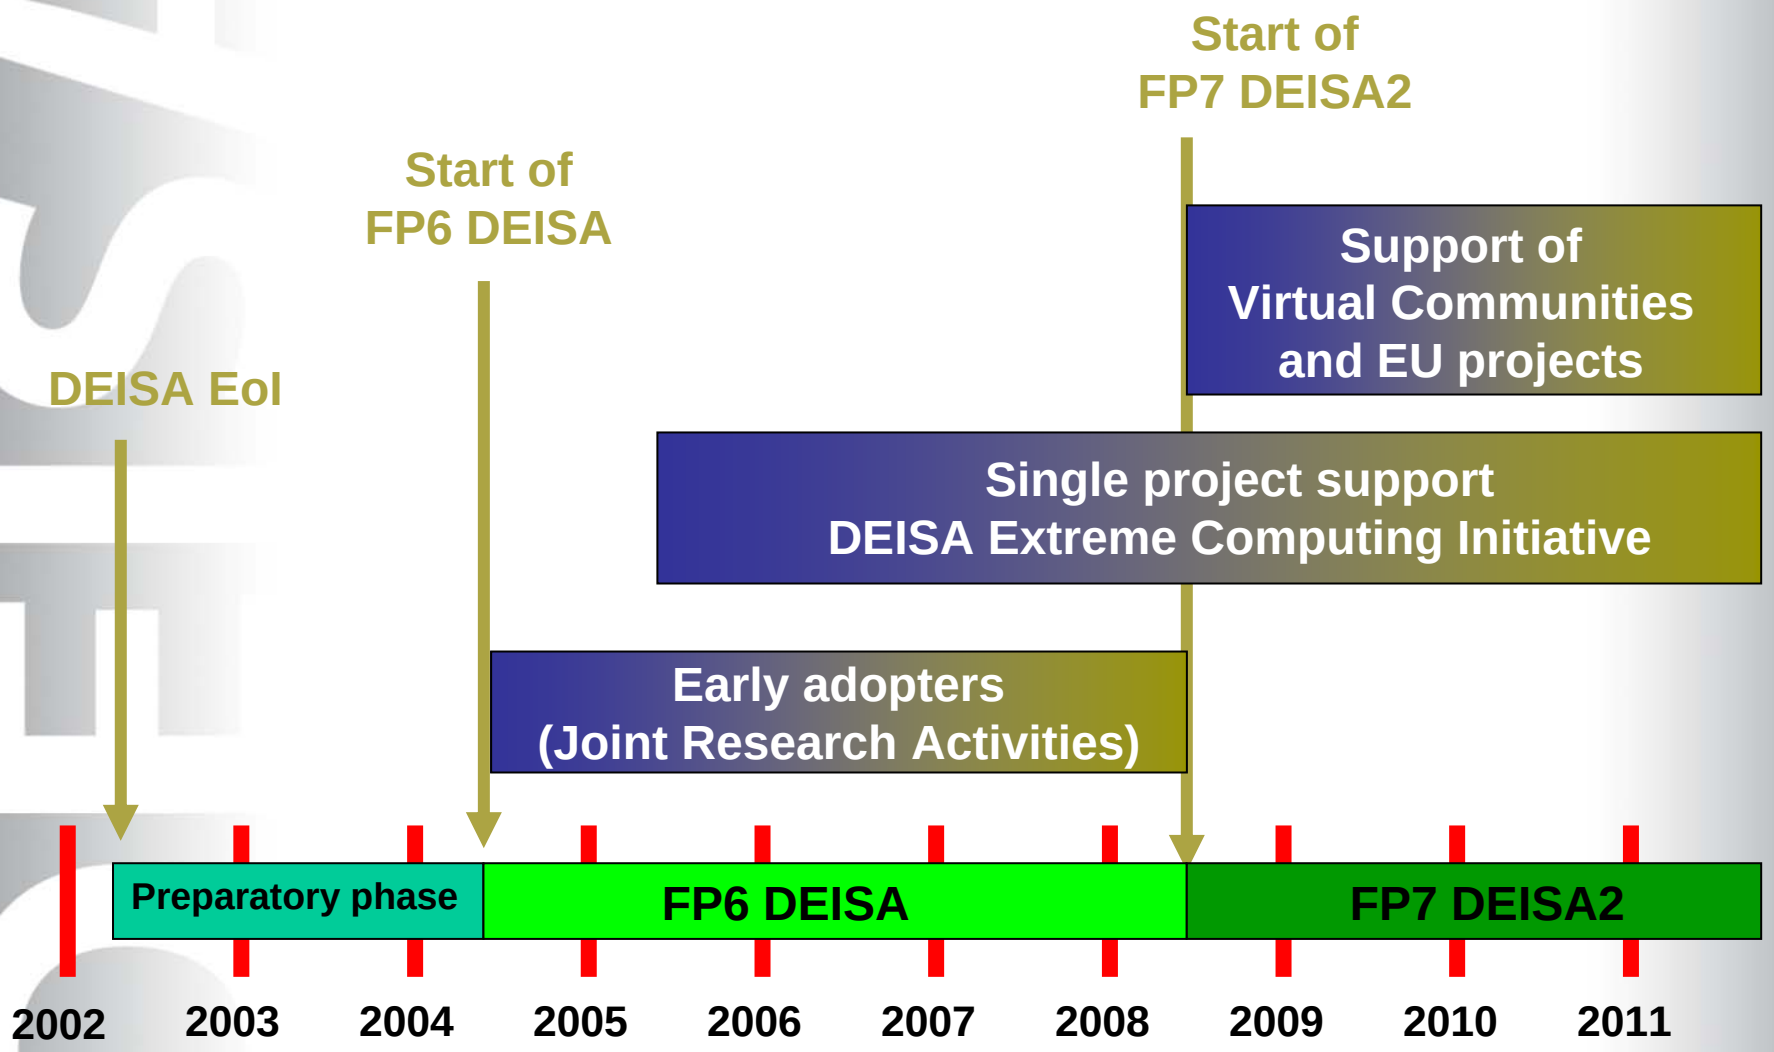

RI-222919

Real needs of HPC users of T0/T1

Distributed European Infrastructure for Supercomputing Applications

HPC users - T0/T1 - are conservative, standard access methods are preferred, no interest in complicated middleware stacks.
Comfortable Data Access among T0/T1/T2 is the key for satisfied users and for the future European HPC Infrastructures!

- Global Login
  - *“HPC users prefer a personal Login in each system”*
- Comfortable Data Access
  - *“HPC users need a global, fast and comfortable access to their data”*
- Common Production Environment
  - *“HPC users do not need an identical but an equivalent HPC software stack”*
- Global Help Desk
  - *“HPC users wish one central point of contact and as local as possible”*
- Application Support
  - "HPC users need help in scalability and adaptation to different architectures"

Distributed European Infrastructure for Supercomputing Applications

Access via Internet

Single sign-on
Secure login
(ssh / gsi-ssh -> D-ssh)
(Unicore, gridFTP)

DEISA Common Production Environment

Different Software Environments

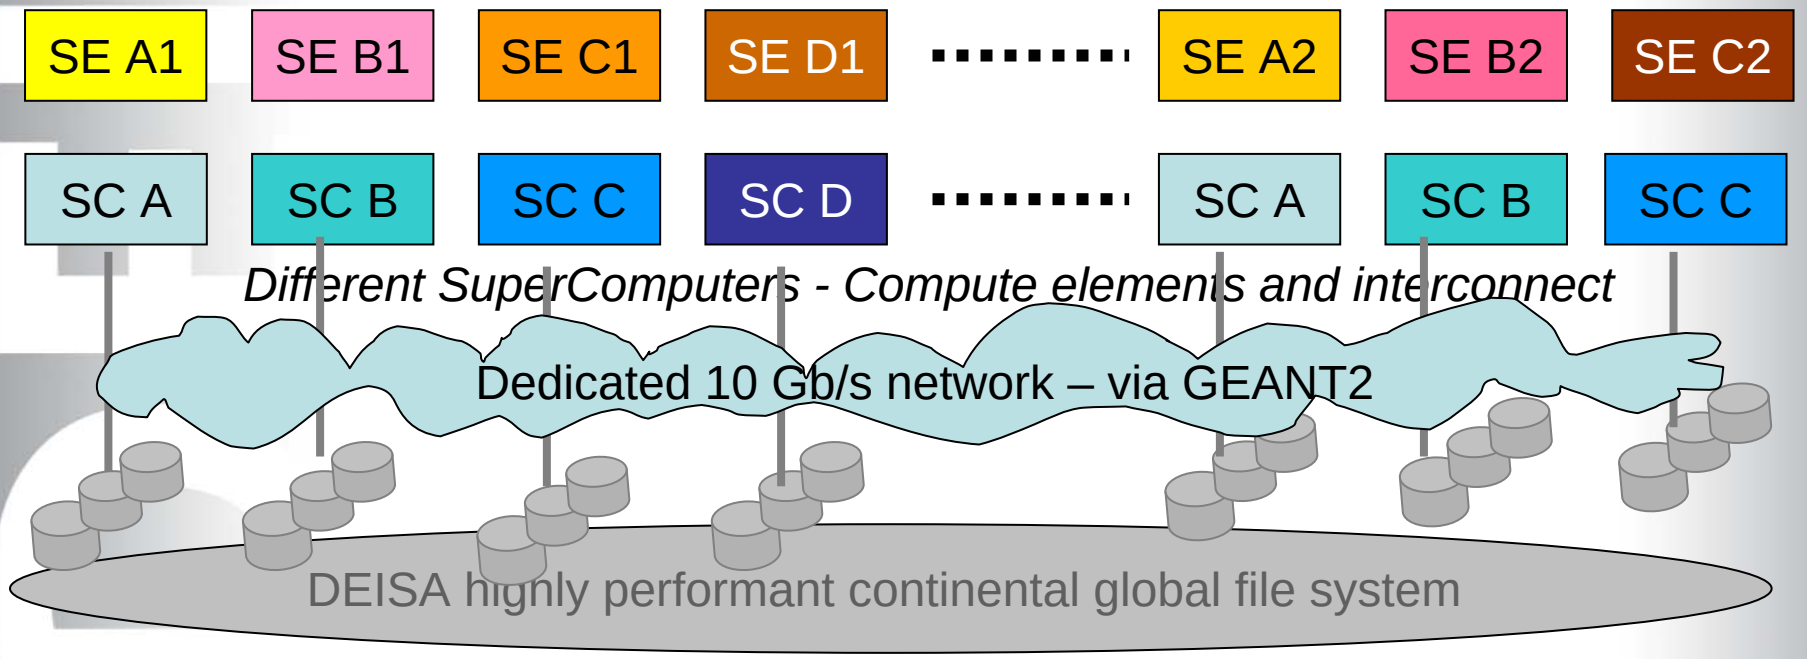

RI-222919

Standardization

Essential to facilitate interoperation of e-infrastructures

- GIN community group within the OGF

- Infrastructure Policy Group

Members: DEISA, EGEE, TeraGrid, OSG, Naregi

Activity: Understanding models of operation to improve interoperation of leading Grid Infrastructures worldwide

Topics: AAAA: Authentication, Authorization, Accounting and Auditing
Resource allocation policies
Portal policies

Meetings: Feb 2008, Boston (during OGF22)
Sep 2008, Singapore (during OGF24)
Mar 2008, Catania (during OGF25)

Evolvements of TeraGrid and DEISA

Distributed European Infrastructure for Supercomputing Applications

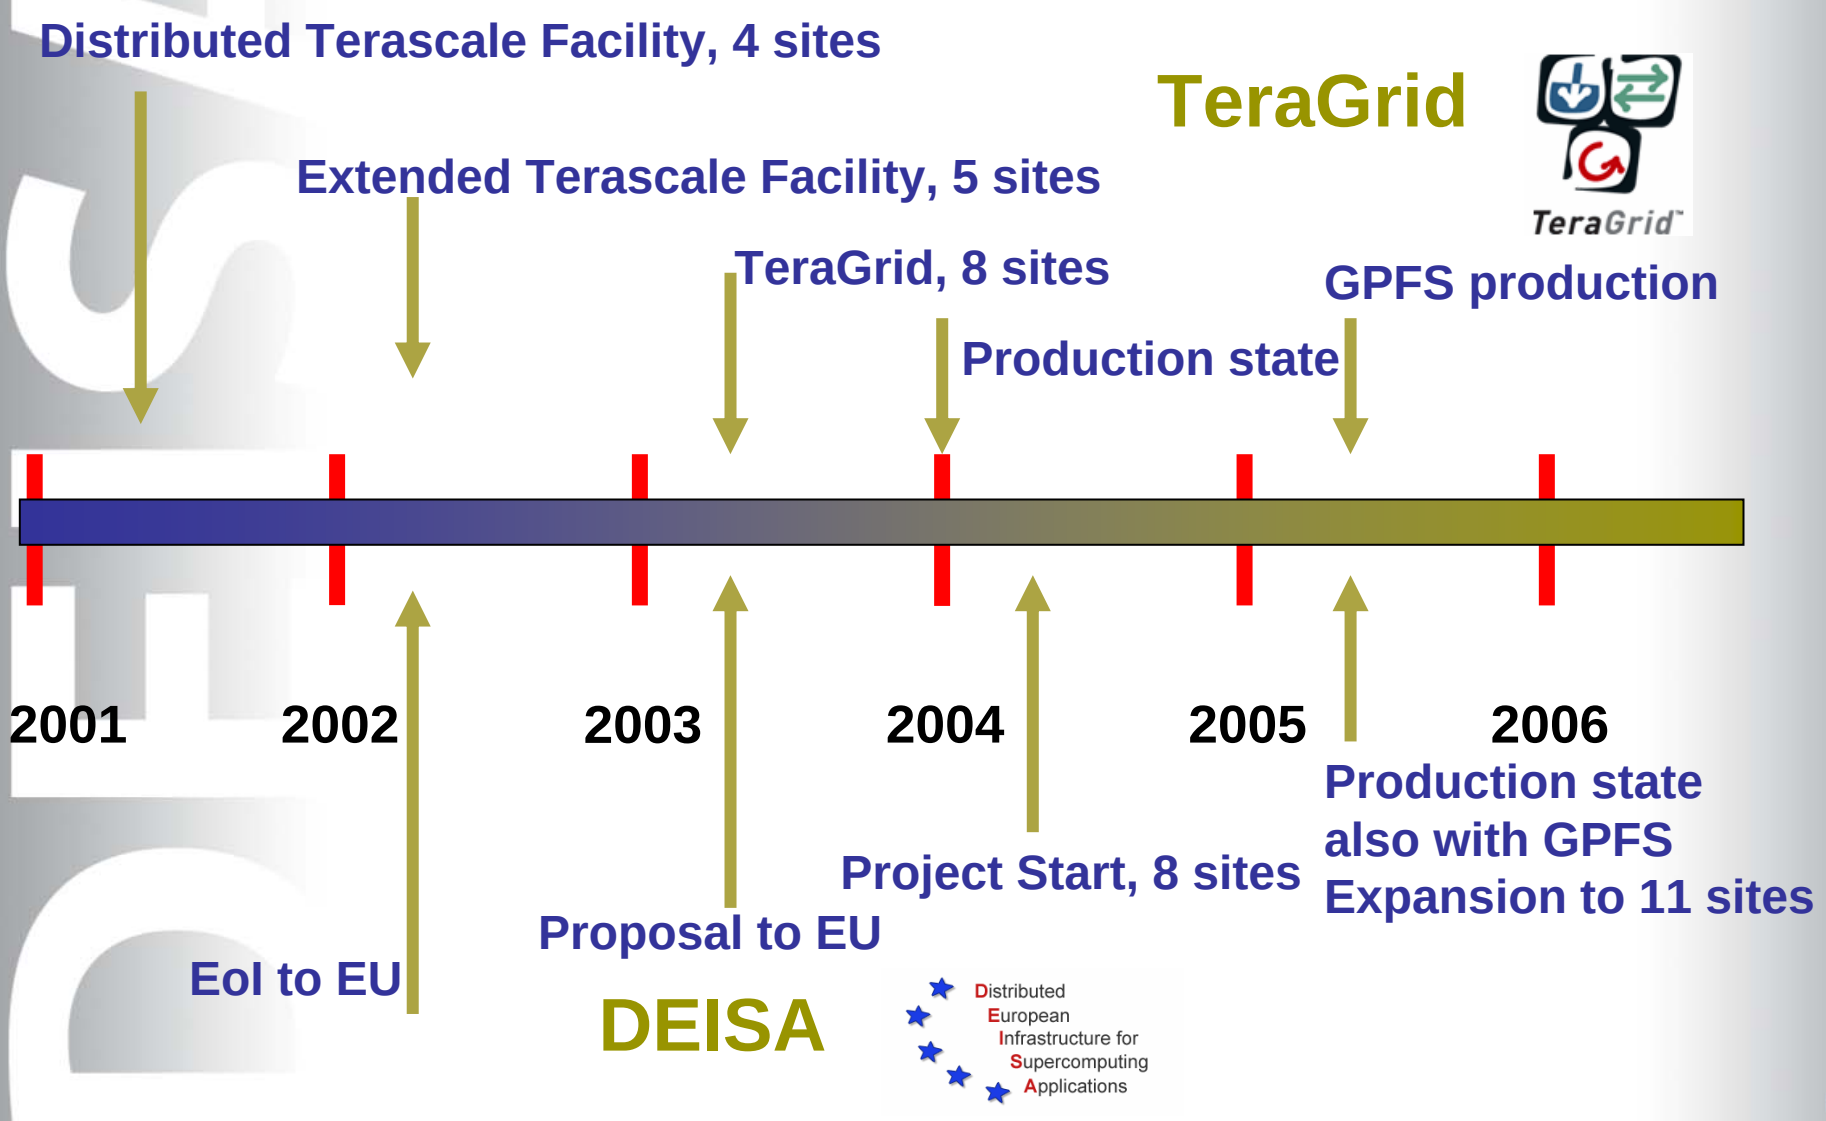

RI-222919

Interoperability Demonstration between DEISA and TeraGrid during SC 2005 in Seattle

Distributed European Infrastructure for Supercomputing Applications

American and European supercomputing infrastructures linked: bridging communities with scalable, wide-area global file systems

Involved TeraGrid Sites

Involved DEISA Sites

Complexity of measures for data access through geographically distributed sites drastically reduced

DEISA PRACE Symposium 2009

Distributed European Infrastructure for Supercomputing Applications

PRACE

HPC Infrastructures for Petascale Applications
11-13 May, 2009, Amsterdam

RI-222919

DEISA Core Infrastructure and Services

Distributed European Infrastructure for Supercomputing Applications

- Dedicated high speed network
- Common AAA
  - Single sign on
  - Accounting/budgeting
- Global data management
  - High performance remote I/O and data sharing with global file systems
  - High performance transfers of large data sets (e.g. GridFTP)
- User Operational infrastructure
  - Distributed Common Production Environment (DCPE)
  - Job management service
  - Common user support and help desk
- System Operational infrastructure
  - Common monitoring and information systems
  - Common system operation
- Global Application Support

DEISA2 Essentials

Distributed European Infrastructure for Supercomputing Applications

- To consolidate the existing DEISA infrastructure and to continue those activities and services that currently contribute to the effective support of world-class computational science in Europe

  A focus is on the provisioning and operation of the specific infrastructure services which allow its users to efficiently work within a distributed high performance computing environment.

- To evolve the DEISA infrastructure towards a robust and persistent European HPC ecosystem, by
  - enhancing the existing services
  - including support for European Virtual Communities
  - collaborating with new European initiatives, especially PRACE that will enable shared European PetaFlop/s supercomputer systems.

- To advance the existing distributed European HPC environment to a turnkey operational solution for a persistent European HPC infrastructure

Vision and Strategy

Distributed European Infrastructure for Supercomputing Applications

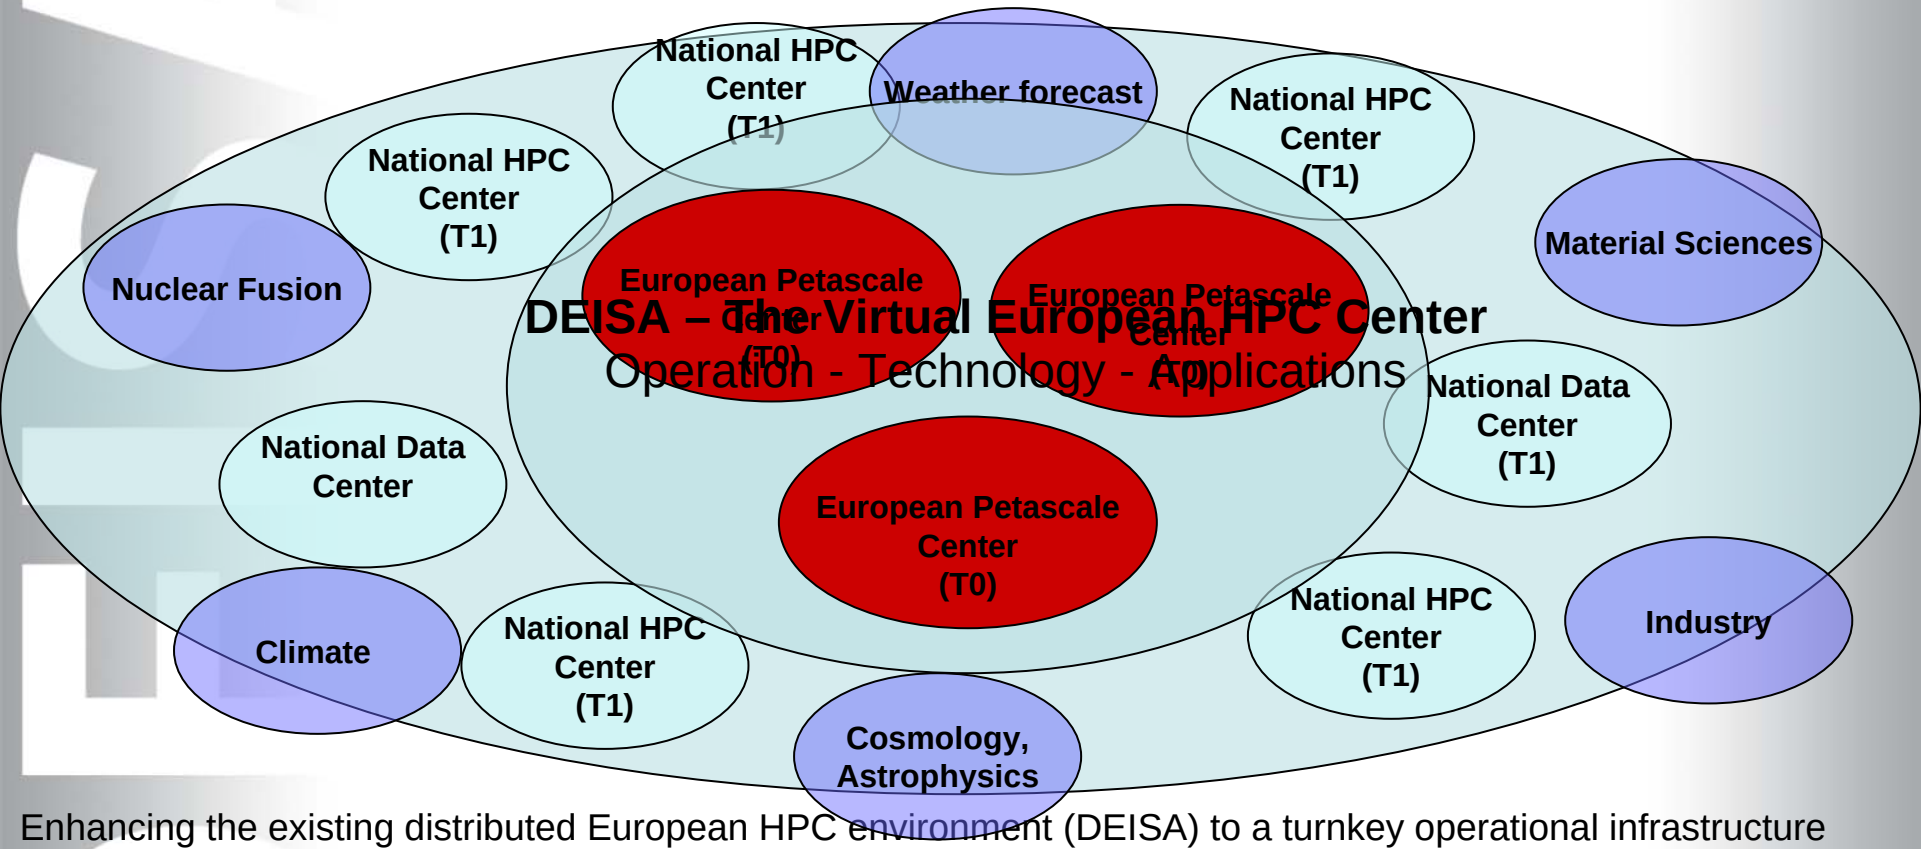

Enhancing the existing distributed European HPC environment (DEISA) to a turnkey operational infrastructure

Advancing the computational sciences in Europe by supporting user communities and extreme computing projects

Enhancing the service provision by offering a complete variety of options of interaction with computational resources

Integration of T-1 and T-0 centres

The Petascale Systems need a transparent access from and into the national data repositories

Bridging worldwide HPC and Grid projects