

User Management in DEISA

The DEISA VO view

Jules Wolfrat
SARA, wolfrat@sara.nl
HPDC'08 workshop
June 24, 2008


www.deisa.eu



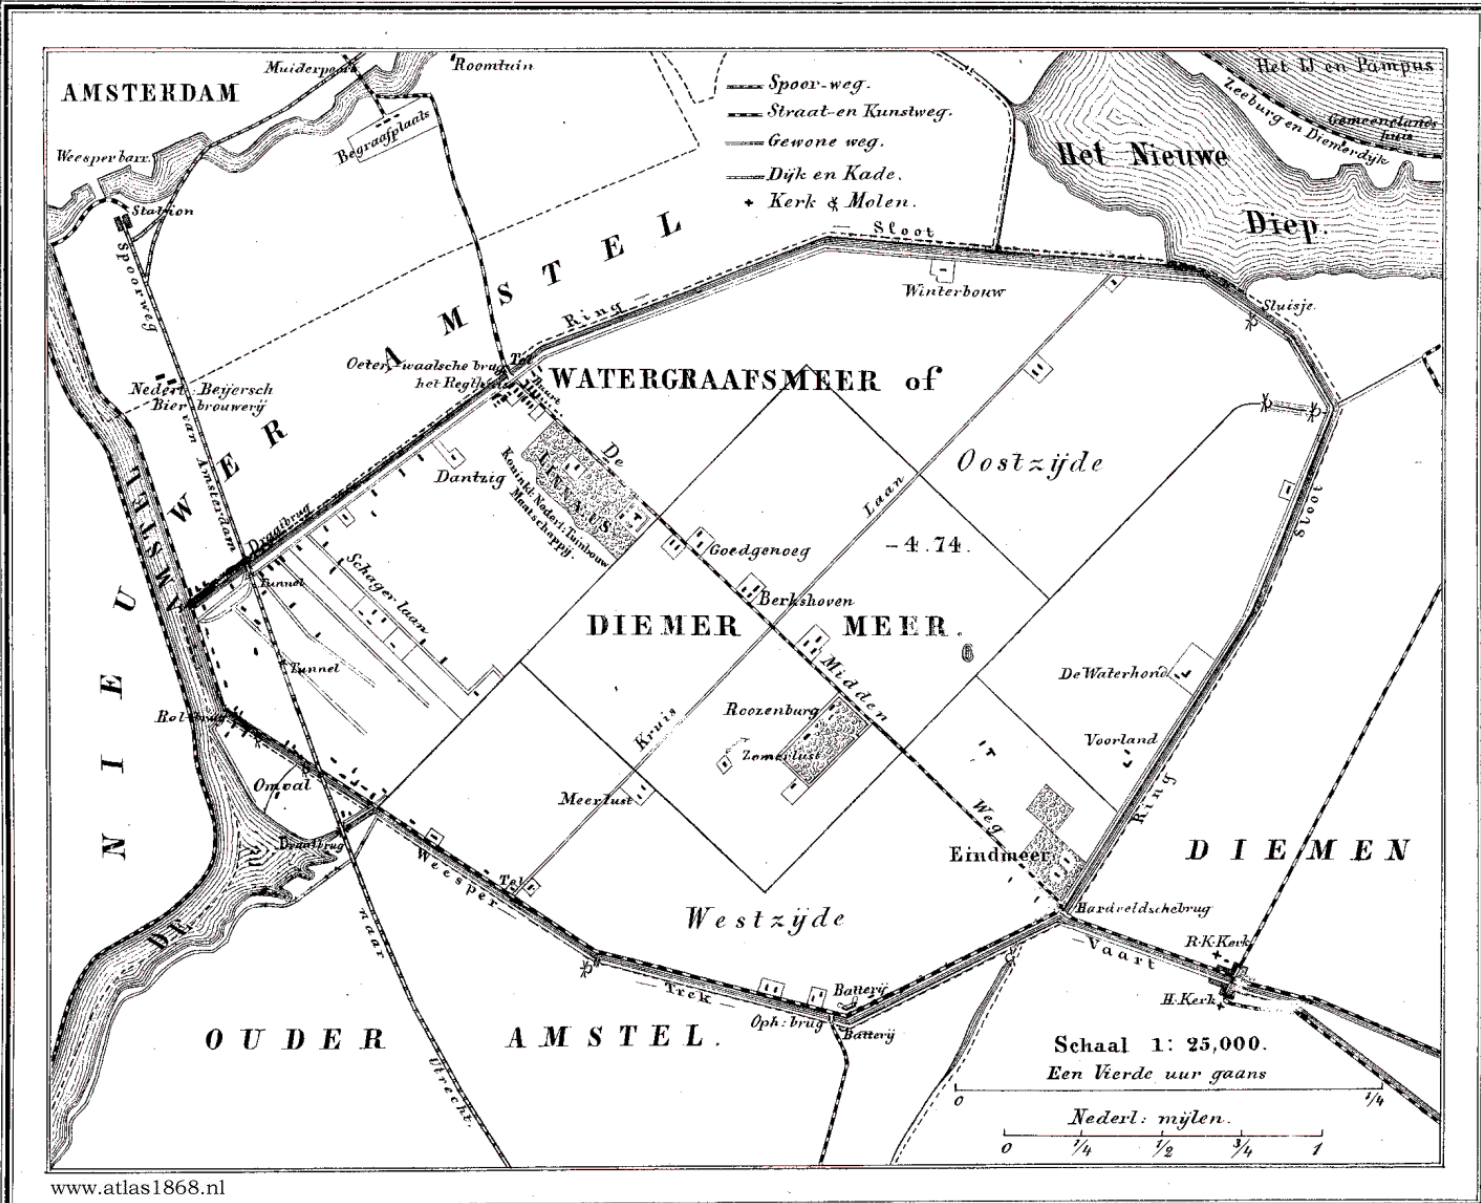
RI-222919



DEISA



★ Distributed
★ European
structure for
ercomputing
lications



ture for
omputing
ations





Outline

- Introduction to DEISA
- Management of users in DEISA
- Interoperability
- VO view DEISA

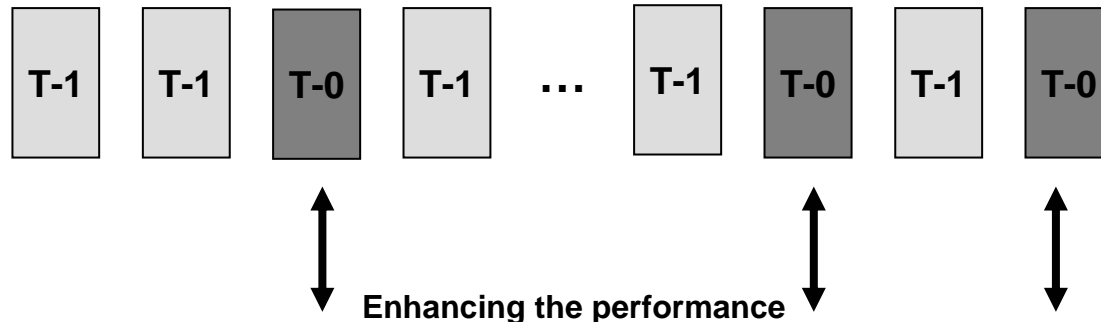
Vision of a European HPC Ecosystem in FP7

DEISA2 - The Infrastructure for the European HPC Ecosystem

Deep operational and technological integration of European HPC (T-0 and T-1) centres and systems providing efficient seamless access to shared HPC resources and large data repositories designing and approving an operational model for a large European Virtual HPC Centre.

Providing scientists access to a large distributed HPC environment via integrated services.

DEISA is paving the way to the efficient operation of the T-0 and T-1 ecosystem

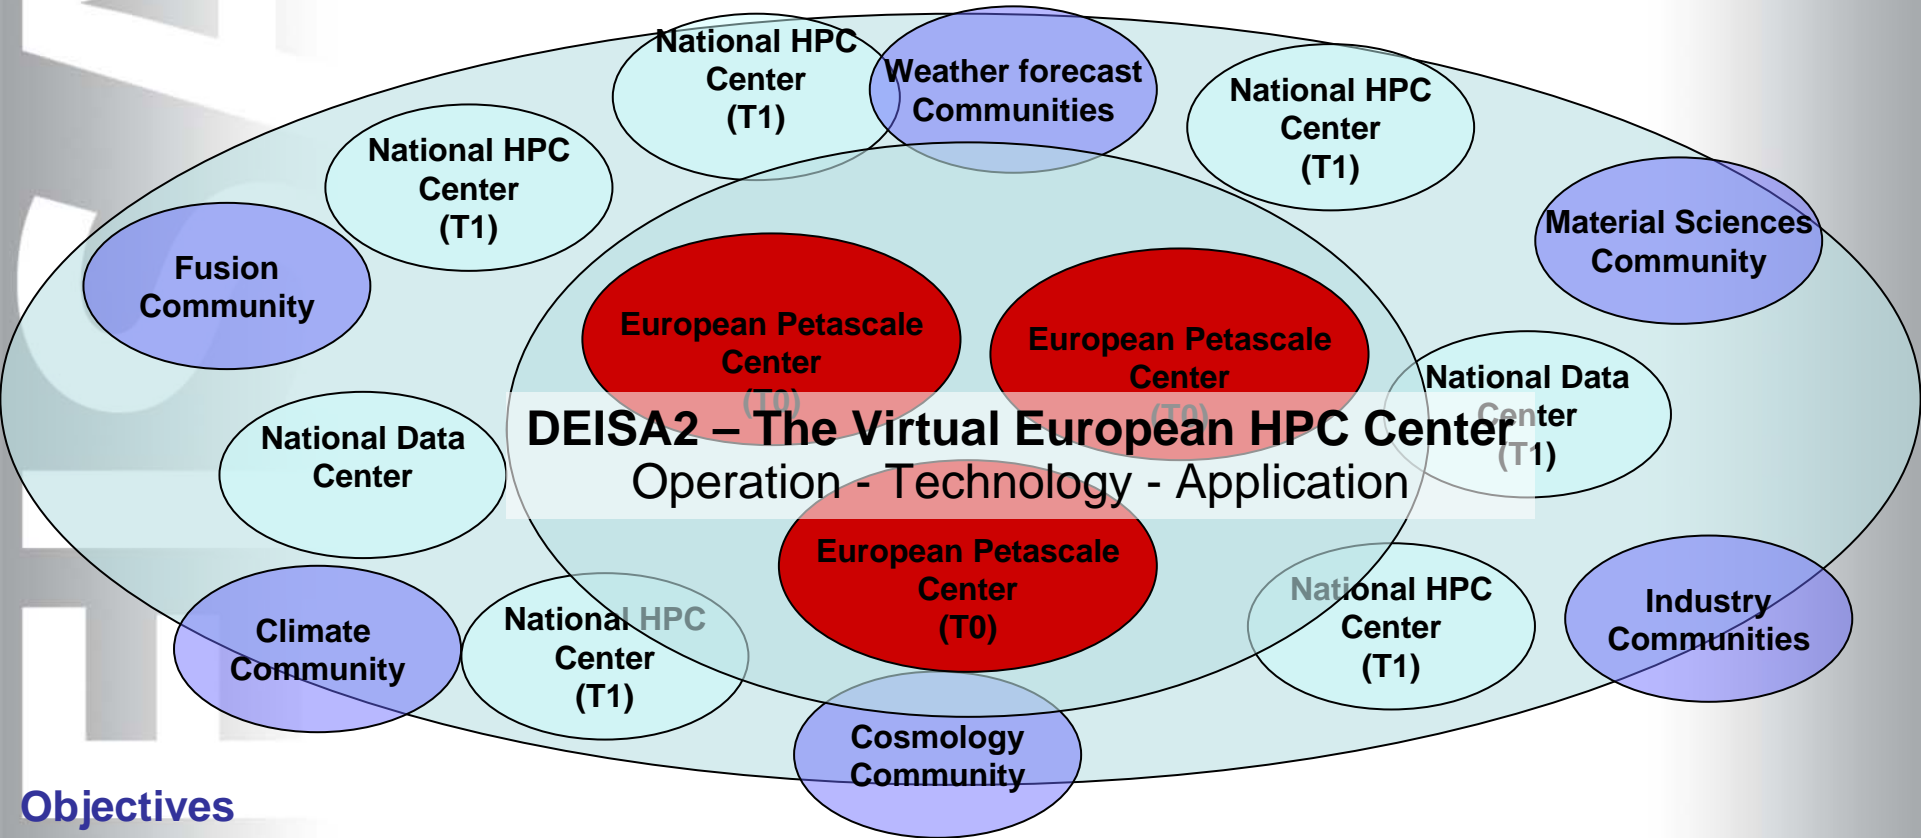


PRACE

Building a world-class pan-European High Performance HPC Ecosystem which is operated under the umbrella of an European Legal Entity adopting operational and technological concepts and services designed and approved by DEISA2.

Towards a European HPC Infrastructure

DEISA2



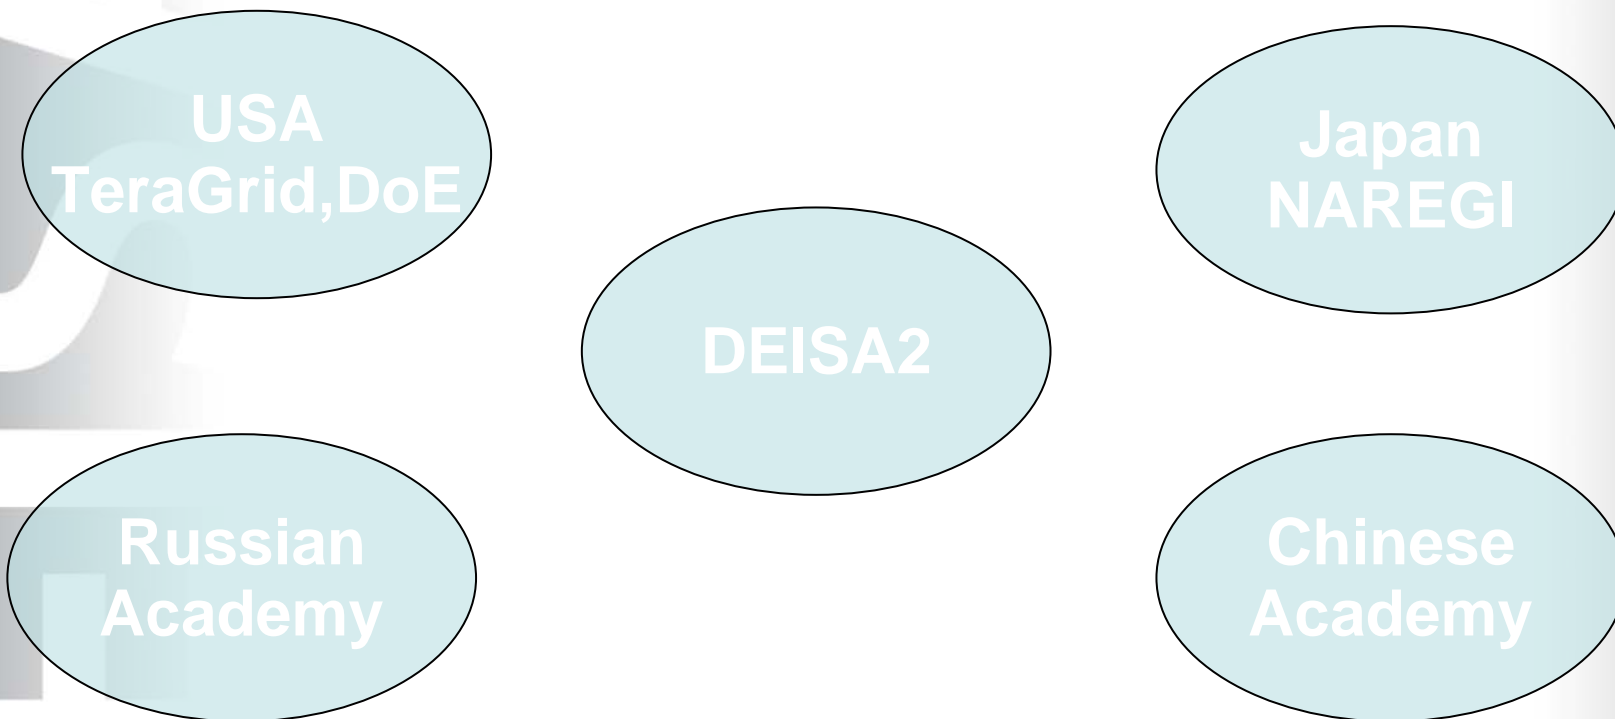
Objectives

- 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 T1 and T0 centres
 - The Petascale Systems need a transparent access from and into the national data repositories
- Bridging worldwide HPC projects



Towards a European HPC Infrastructure

DEISA2



Objectives

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

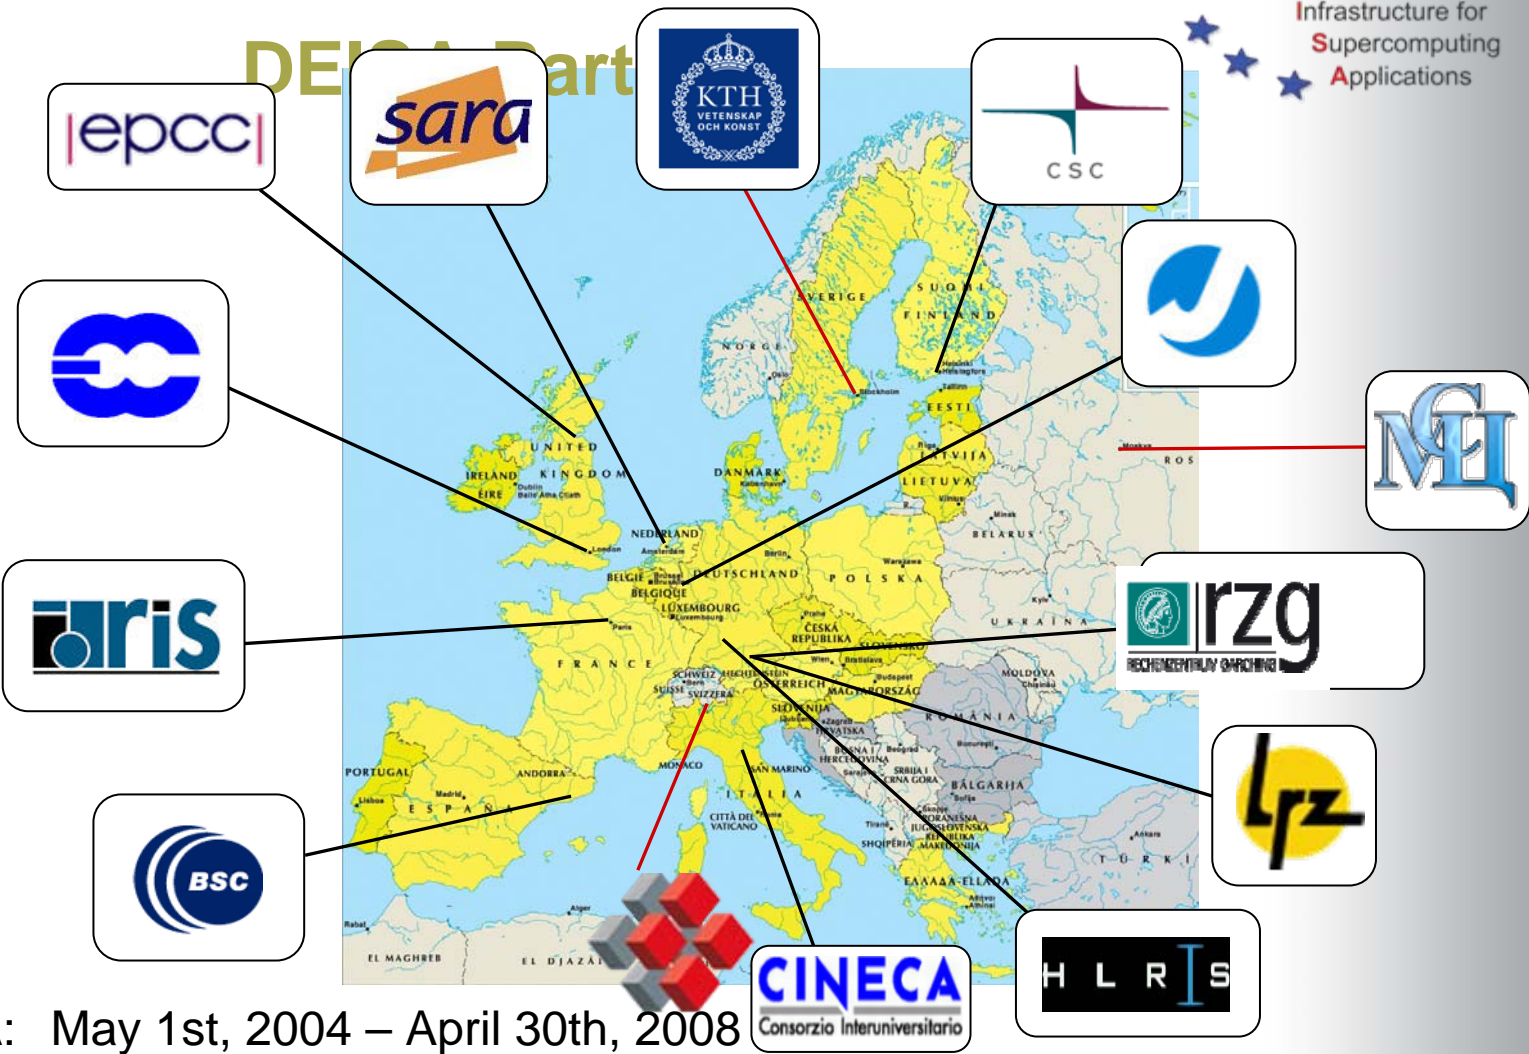
Advancing the computational sciences in Europe by supporting user communities and extreme computing projects

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

Integration of T1 and T0 centres

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

Bridging worldwide HPC projects



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

Three new partners joined June 2005 (eDEISA)

DEISA2: May 1st, 2008 – April 30th, 2011 with additional associated partners



Partners/ Associate Partners

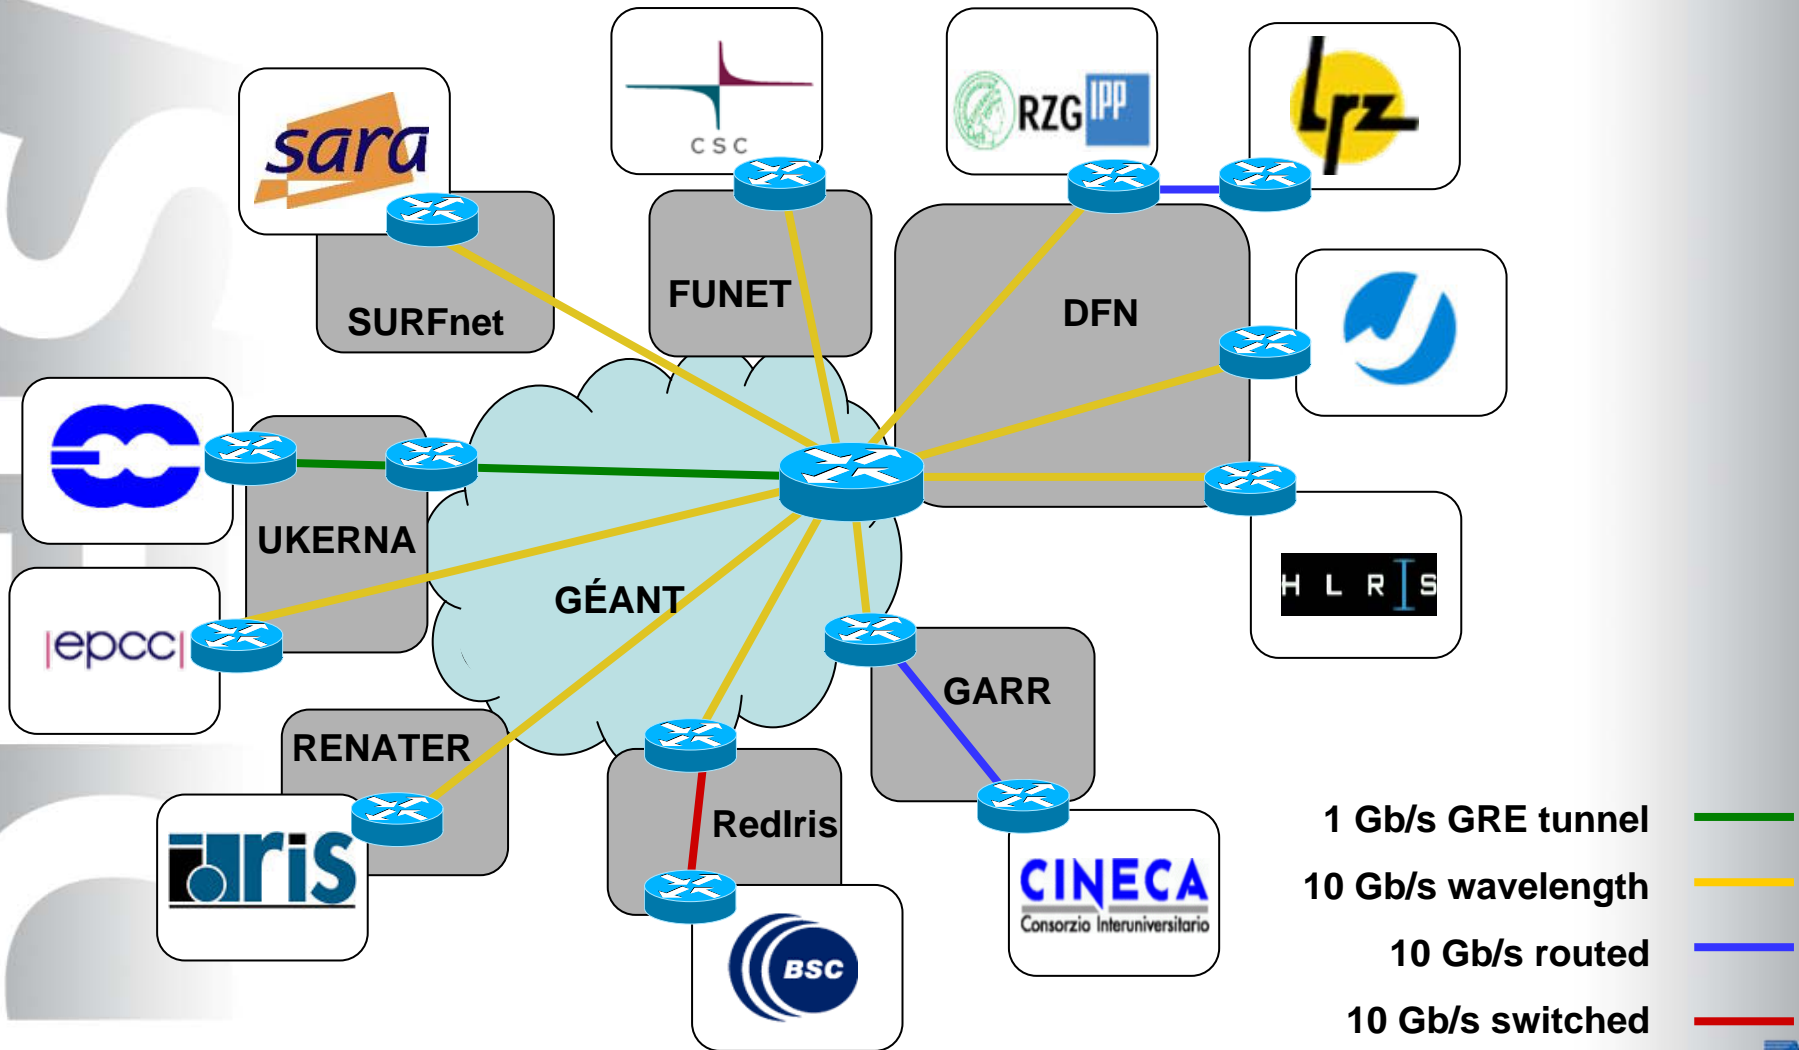


BSC	<i>Barcelona Supercomputing Centre</i>	Spain
CINECA	<i>Consortio Interuniversitario per il Calcolo Automatico</i>	Italy
CSC	<i>Finnish Information Technology Centre for Science</i>	Finland
EPCC	<i>University of Edinburgh and CCLRC</i>	UK
ECMWF	<i>European Centre for Medium-Range Weather Forecast</i>	UK (int)
FZJ	<i>Research Centre Juelich</i>	Germany
HLRS	<i>High Performance Computing Centre Stuttgart</i>	Germany
IDRIS	<i>Institut du Développement et des Ressources en Informatique Scientifique - CNRS</i>	France
LRZ	<i>Leibniz Rechenzentrum Munich</i>	Germany
RZG	<i>Rechenzentrum Garching of the Max Planck Society</i>	Germany
SARA	<i>Dutch National High Performance Computing</i>	Netherlands
KTH	<i>Kungliga Tekniska Högskolan</i>	Sweden
CSCS	<i>Swiss National Supercomputing Centre</i>	Switzerland
JSCC	<i>Joint Supercomputer Center of the Russian Academy of Sciences</i>	Russia

The basic DEISA infrastructures and services

- Dedicated high speed network infrastructure
- Common AAA infrastructure
- Global data management infrastructure
 - Integrating distributed data with distributed computing platforms, including hierarchical storage management and databases. Major highlights are:
 - High performance remote I/O and data sharing with global file systems, using full network bandwidth
 - High performance transfers of large data sets, using full network bandwidth
- DCPE (DEISA Common Production Environment)
 - The job management service
 - The science gateways (portals) to supercomputing resources
- Common Operation Environment
 - Common monitoring and Information systems
 - Common system operation
 - Common help desk
- Global Application Support

DEISA network infrastructure



DEISA Extreme Computing Initiative

DECI call 2005

51 proposals, 12 European countries involved
30 mio cpu-h requested
29 proposals accepted, 12 mio cpu-h
(standardized to P4+ at FZJ)

DECI call 2006


41 proposals, 12 European countries involved
28 mio cpu-h requested
23 proposals accepted, 12 mio cpu-h

DECI call 2007

63 proposals, 14 European countries involved
(US, Canada, Brazil, Israel)
70 mio cpu-h requested
45 proposals accepted, 30 mio cpu-h

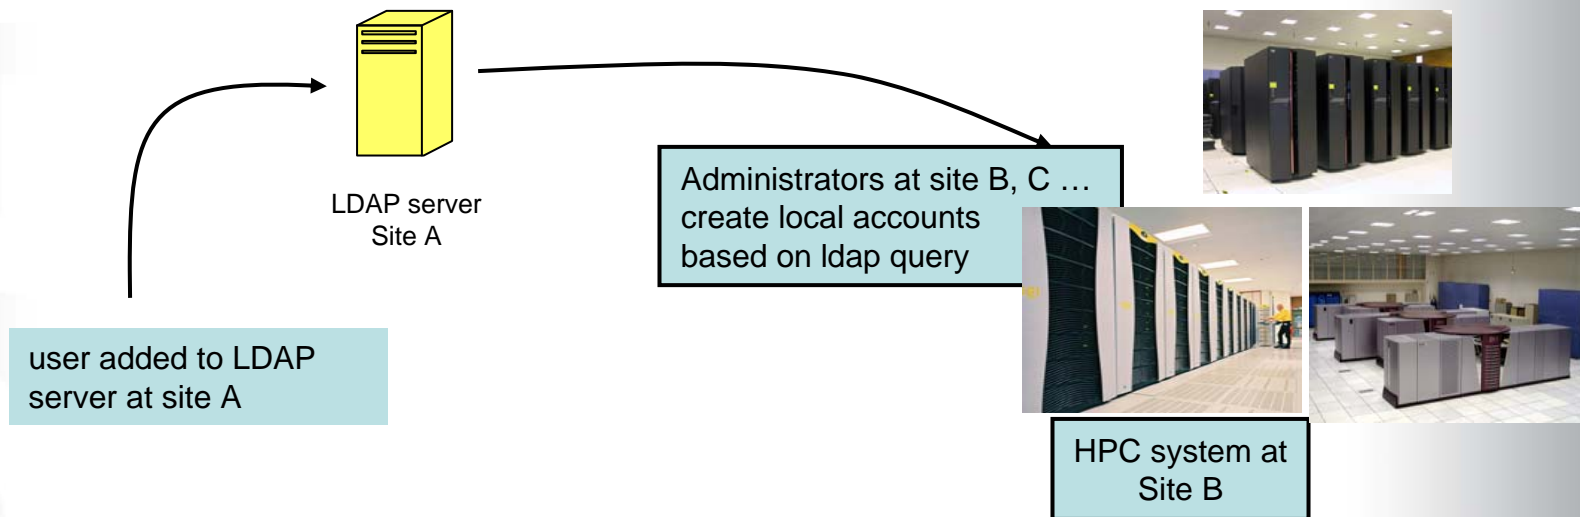
Management of users in DEISA

- For the administration of DEISA users a dedicated distributed repository is built based on LDAP
- Only trusted servers are authorized to access each other (based on X.509 certs) and encrypted communication is used to maintain confidentiality (TLS based)



User Administration System (1)

- 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!



User Administration System (2)

- Some 20 attributes used for the registration of users using existing object classes and a DEISA specific defined schema
- Information is used for
 - Providing authorization information (UNIX accounts, X.509 certs for UNICORE UUDB, grid-mapfile for GT4).
 - Additional information to comply with requirements of partners:
 - Phone number, email address, Science field, Nationality, Status, Project
 - e.g. Nationality because of export regulations for some of the systems in use
- To avoid overlap between DEISA uid numbers and local numbers each site uses reserved ranges
 - GPFS authorization based on uids, so must be the same at each site
- Policies for administrators defined, e.g. what to do if user account has to be deactivated.

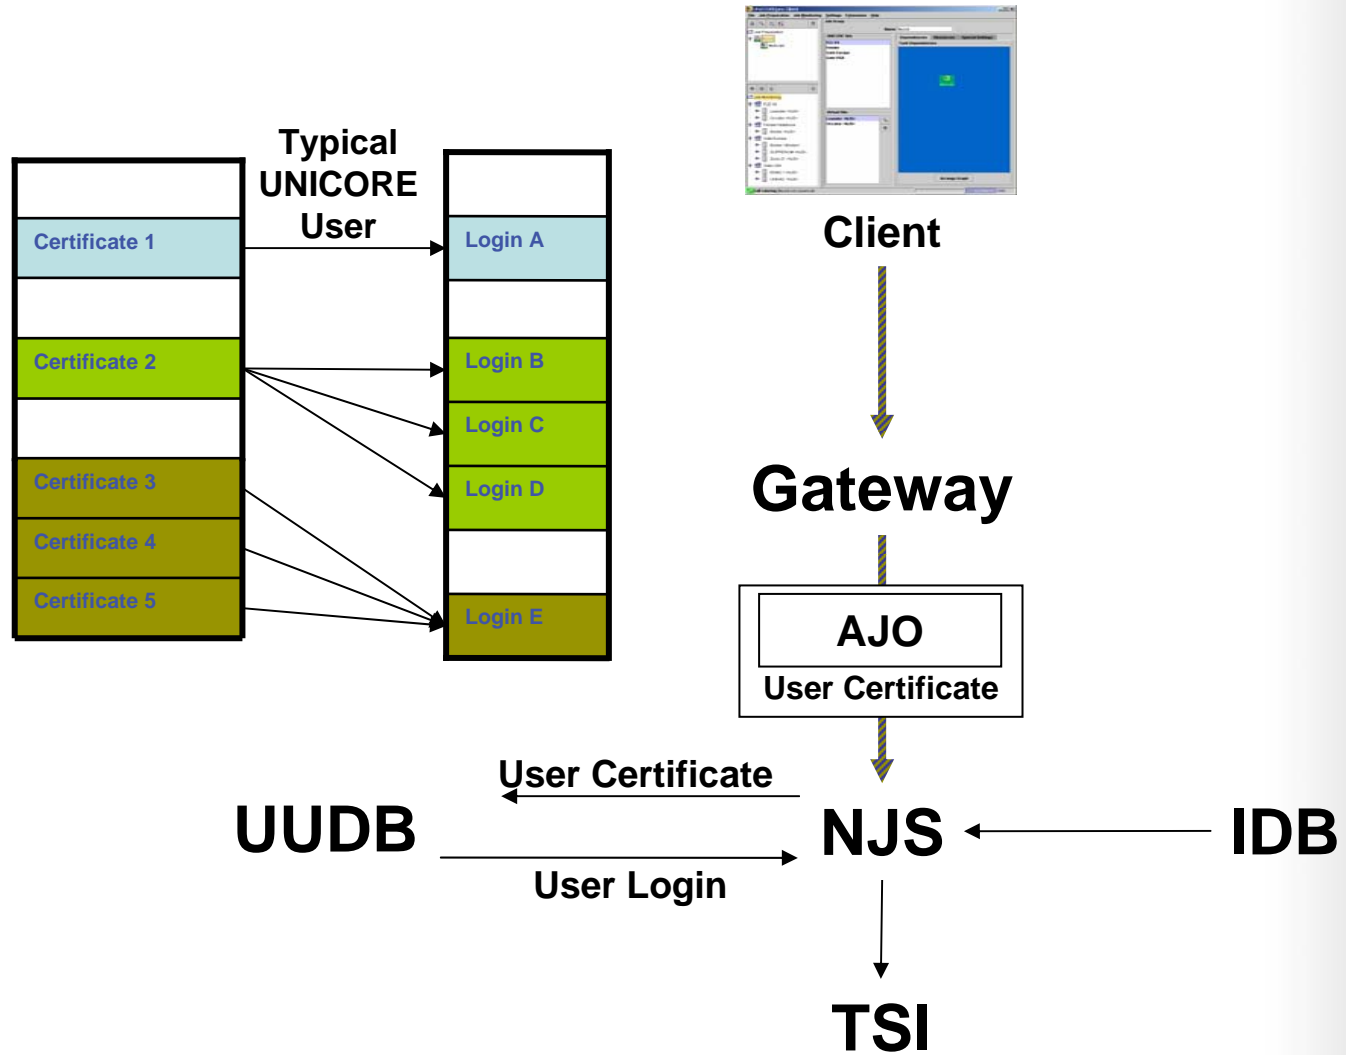
Monitoring the Administration service

- Service is not critical as n/a doesn't impact production service directly
- But important that failing servers are detected in time (order of hours) and that registrations are conform to the policies defined
- Every two hours check of availability of LDAP servers implemented (just see if there is a response on a simple query)
 - Not responding site is notified through a message using the INCA framework
- Audit of all registrations is performed on a daily basis, also using the INCA framework
 - Not all partners ok, but no critical errors (e.g. errors in optional attributes)

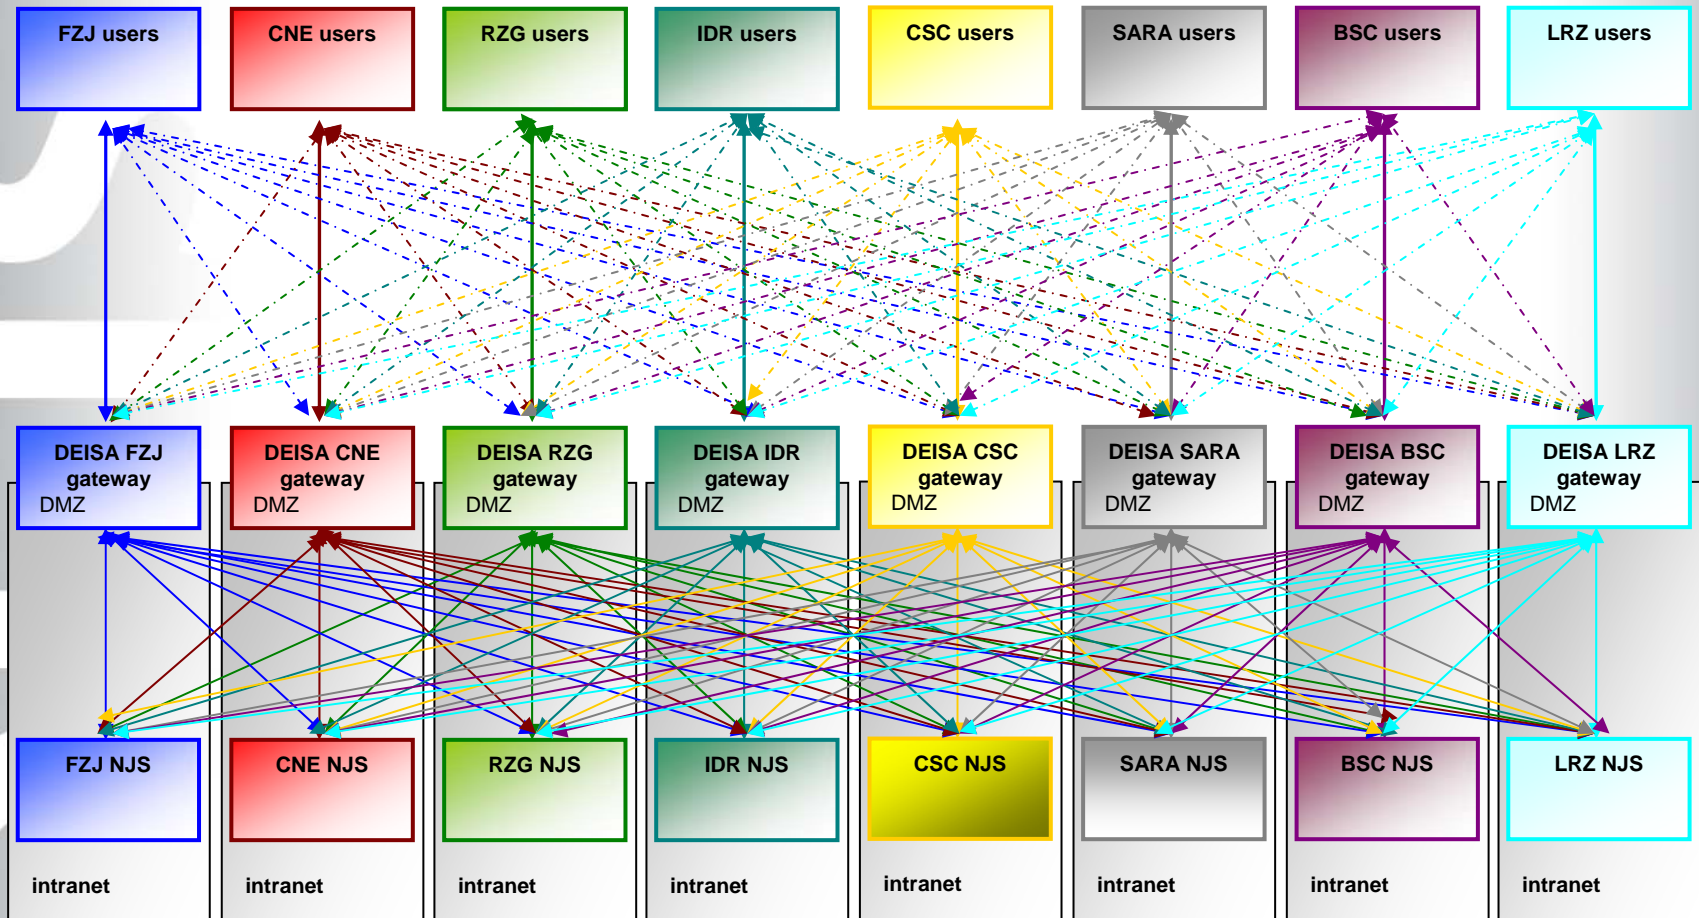
LDAP Availability	LDAP
LDAP_at_BSC	pass
LDAP_at_CINECA	pass
LDAP_at_CSC	pass
LDAP_at_ECMWF	pass
LDAP_at_EPCC	pass
LDAP_at_FZJ	pass
LDAP_at_HLRS	pass
LDAP_at_IDRIS	pass
LDAP_at_LRZ	pass
LDAP_at_RZG	pass
LDAP_at_SARA	pass

LDAP Audit	LDAP
LDAP_audit_at_BSC	error
LDAP_audit_at_CINECA	error
LDAP_audit_at_CSC	pass
LDAP_audit_at_ECMWF	error
LDAP_audit_at_EPCC	pass
LDAP_audit_at_FZJ	pass
LDAP_audit_at_HLRS	error
LDAP_audit_at_IDRIS	pass
LDAP_audit_at_LRZ	pass
LDAP_audit_at_RZG	error
LDAP_audit_at_SARA	error

UNICORE AuthN & AuthZ (1)



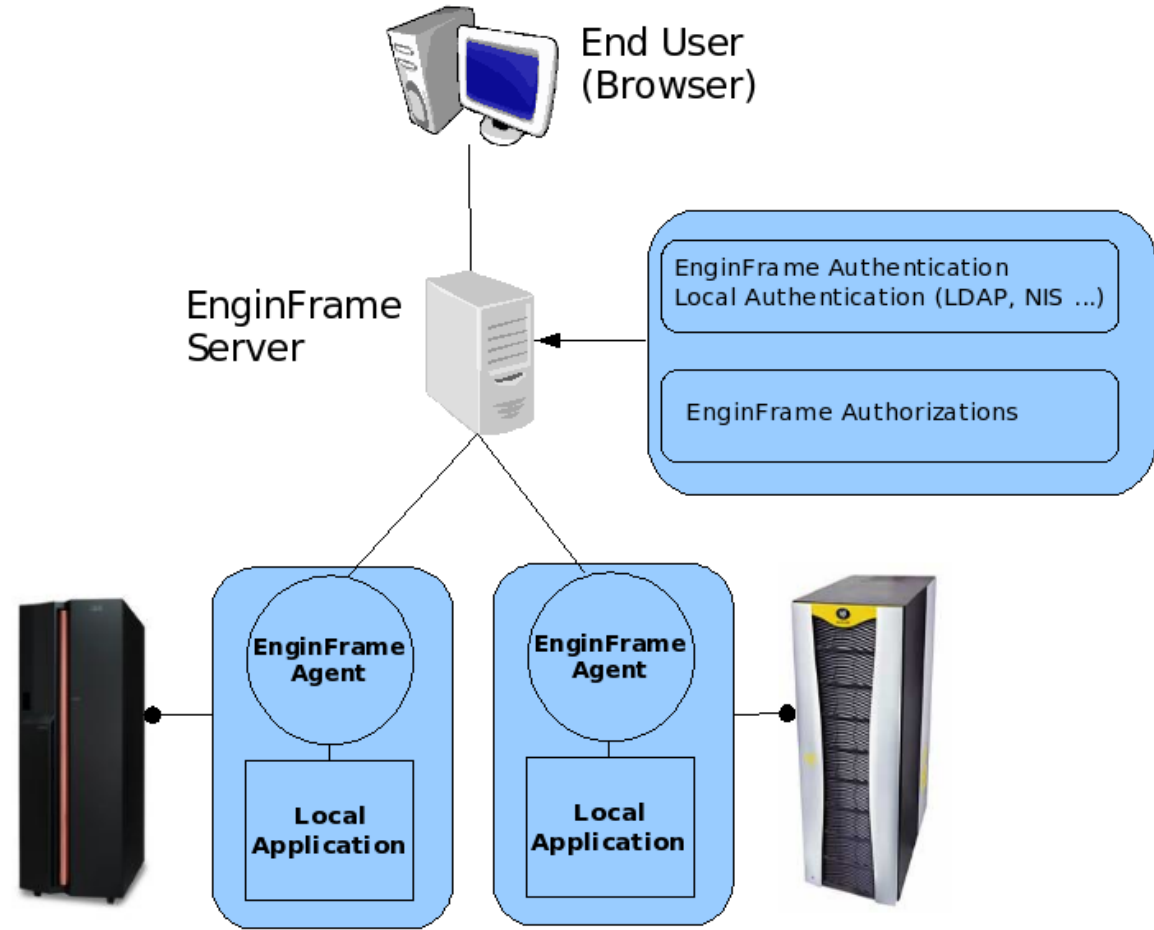
UNICORE config for DEISA



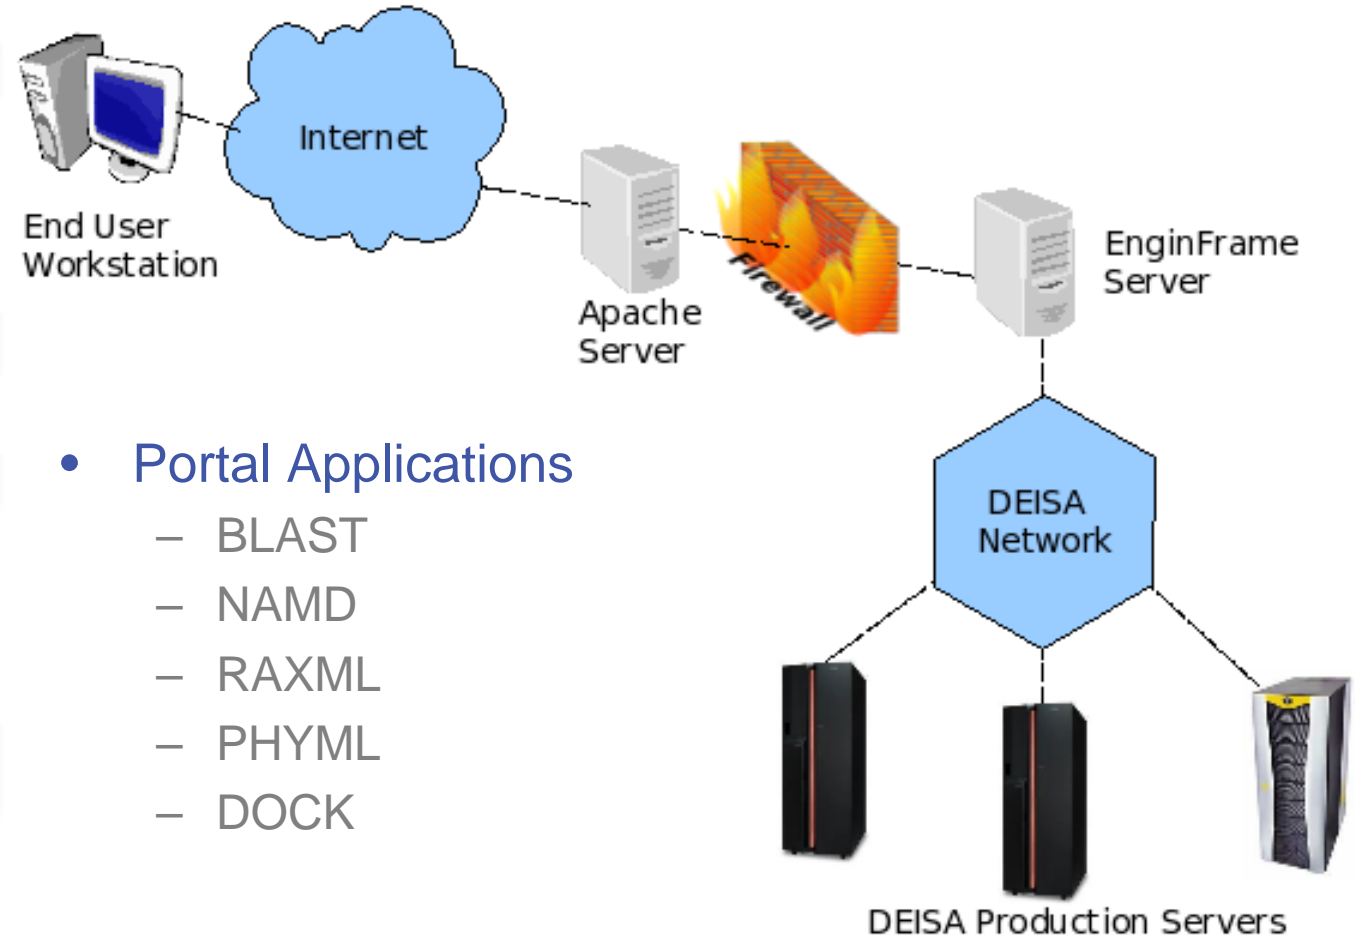
UNICORE AuthN & AuthZ (2)

- UNICORE AuthN and AuthZ is based on X.509 certificates
 - AuthZ based on Subject Name mapping to uids in UUDB (like the gridmapfile)
 - UUDB is maintained at each site. So sites can decide if user can get access through UNICORE, e.g. based on the project the user is working on. Subject names are distributed using the LDAP system.
 - Subject name can be mapped to more than one uid, the user can specify with UNICORE which uid to use

Portal : EnginFrame Architecture



Portal: Deployment on the DEISA grid



- Portal Applications
 - BLAST
 - NAMD
 - RAXML
 - PHYML
 - DOCK

Mgmt of Portal users

- Creation of Accounts
 - Each user must sign a form (that mentions the DEISA AUP)
 - The project manager sends to the site managing the portal server (e.g. IDRIS) the list of its client IPs
 - IDRIS checks that the nationality of the users are accepted by the site(s) that will host the project application.
 - IDRIS registers the project and users in LDAP
 - IDRIS registers the project and the users on the portal

Interoperability

www.deisa.eu



RI-222919



Authentication

- DEISA vision is one global federated namespace, enabling Single Sign-on facilities.
 - Direct (interactive) access to systems required for
 - Testing/debugging of code on specific architectures
 - Submission and checking of production jobs
- DEISA AuthN is X.509 based PKI
 - DEISA trusts IGTF accredited CAs
 - Relying party of EUGridPMA
 - SSH functionality based on X.509 certificates required
 - Handling of X.509 certs cumbersome for users, should be hidden, e.g. using Shibboleth technology
- No interoperability issues for AuthN



Authorization (1)

- Authorization information from LDAP service not directly usable by other projects because most use VOMS based information
 - But both systems give attribute information about users, used differently
 - VOMS information used dynamically to add information to certificates
 - LDAP information used to update authorization DBs (UNICORE UADB and GT4 grid-mapfile)
 - DEISA can operate a VOMS server feeding the information from the LDAP system
 - DEISA can import VOMS information into LDAP system, but currently VOMS based system doesn't supply all needed information

DEISA view on VOs

- Within HPC/DEISA accounts are managed by sites. LDAP repository managed by sites, not by VOs!
 - Partners/sites are the VOs. From the AUP (Acceptable Use Policy): “the Virtual Organization is defined as the DEISA partner that registers you as a DEISA user”. Users are associated through a partner (site).
 - But we also have application communities, where scientists from different regions collaborate and where the Principal Investigator authorizes the members of the project
 - Who to trust and how?
- Important to build a trust basis for authorization information providers, like for the CAs
 - Agreement on policies for management of AuthZ information providers (the LDAP and VOMS servers) needed – at least information on operational model must be available
 - Policies for DEISA are described in User Administration guide
 - Within the IGTF community an AuthZ WG has been formed to investigate this

Accounting

- Usage information published using OGF UR-WG format recommendation
- Developed our own system for publishing usage records
 - Each site publishes data locally in DB (eXist)
 - Access based on role – user, PI (principal investigator), site admin
 - Based on X.509 SubjectName
 - GUI for producing reports (DART)
 - Conversion (comparison) between systems based on CPU performance

Conclusion

- For interoperability technical problems must be solved, but no fundamental barriers
- Middleware must be able to deal with information from Identity Providers (IdPs)
 - E.g. UNICORE 6 can use Shibboleth issued attribute assertions
- Trust building between IdPs important

DEISA

Thank you!

www.deisa.eu



RI-222919



