



Enabling Grids for
E-science in Europe

www.eu-egee.org

EGEE Induction, 17-19 May 2004

Infrastructure and Fabric – EGEE Operations



Ian Bird, CERN
EGEE Operations Manager

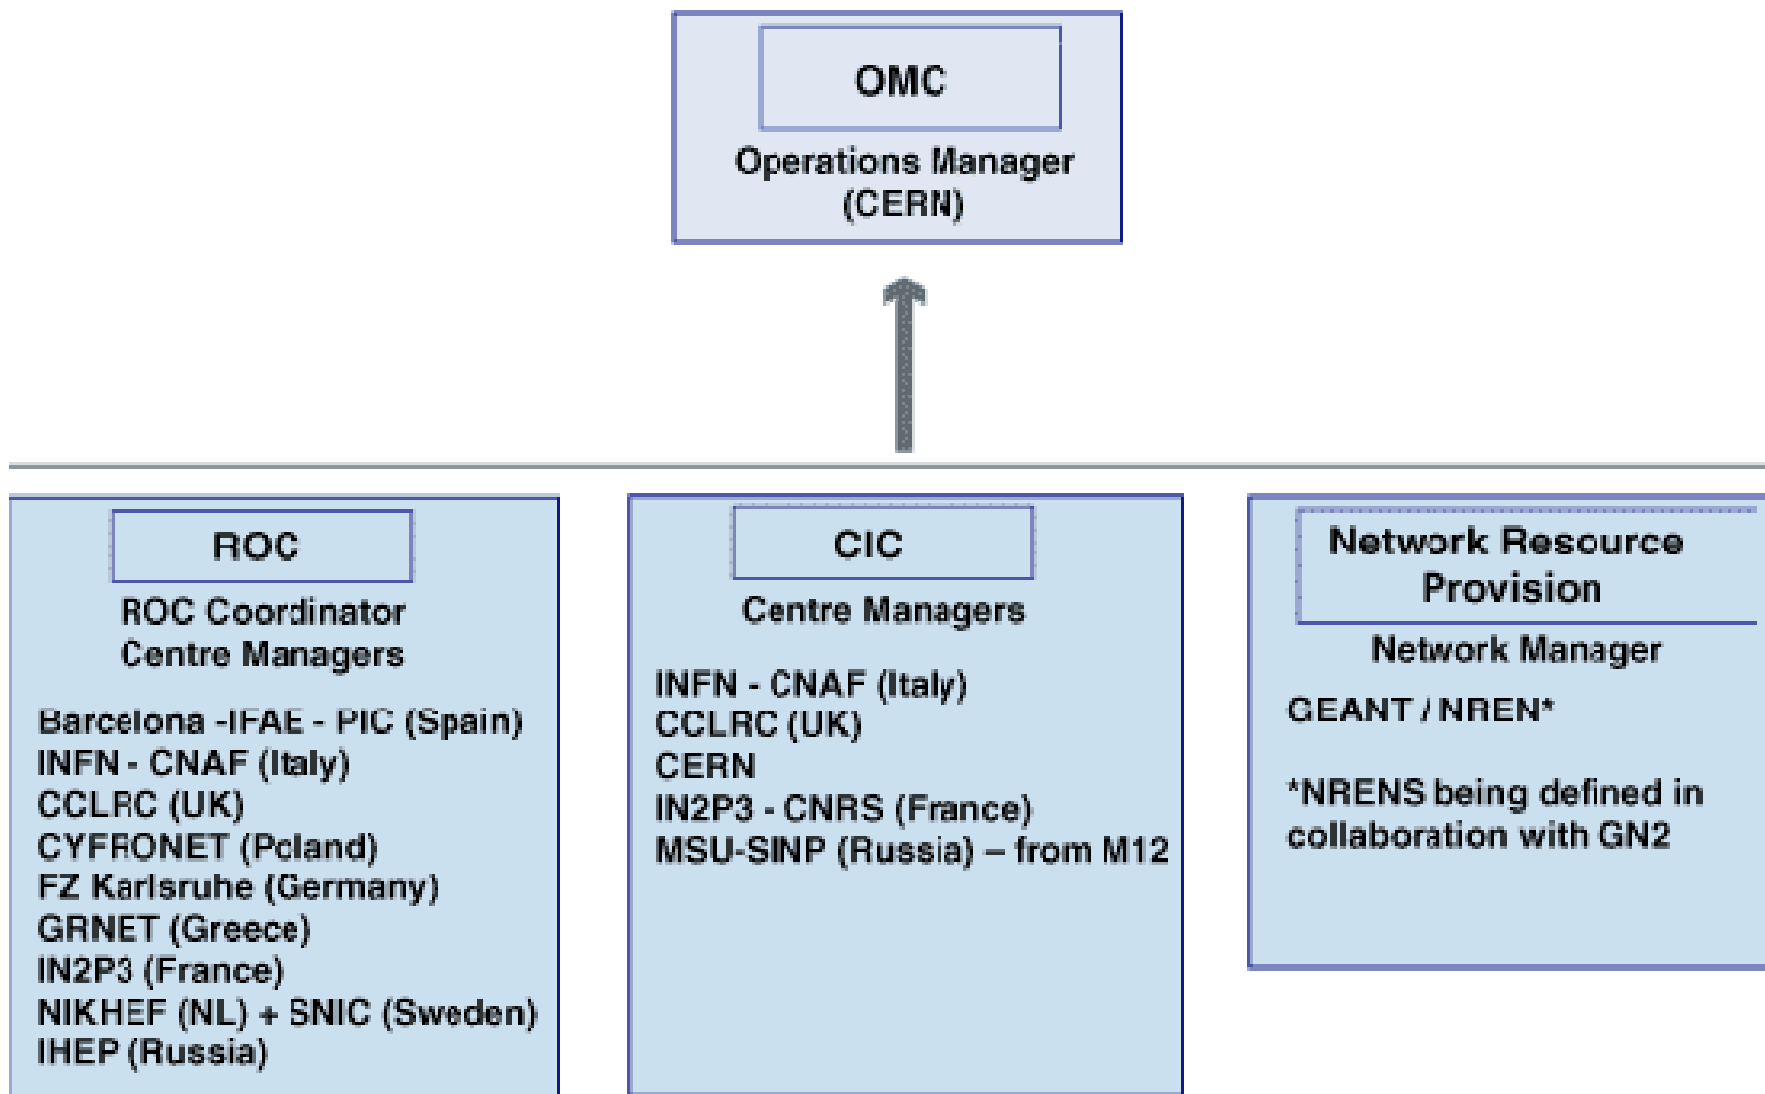
EGEE is a project funded by the European Union under contract IST-2003-508833

Contents

- Introduction – Operational activities
- Organisation – managing the infrastructure
- Infrastructure services
- Deployment process
 - Adding sites and VO's
- Network services – SA2
- Summary



Operations (SA1, SA2) Management



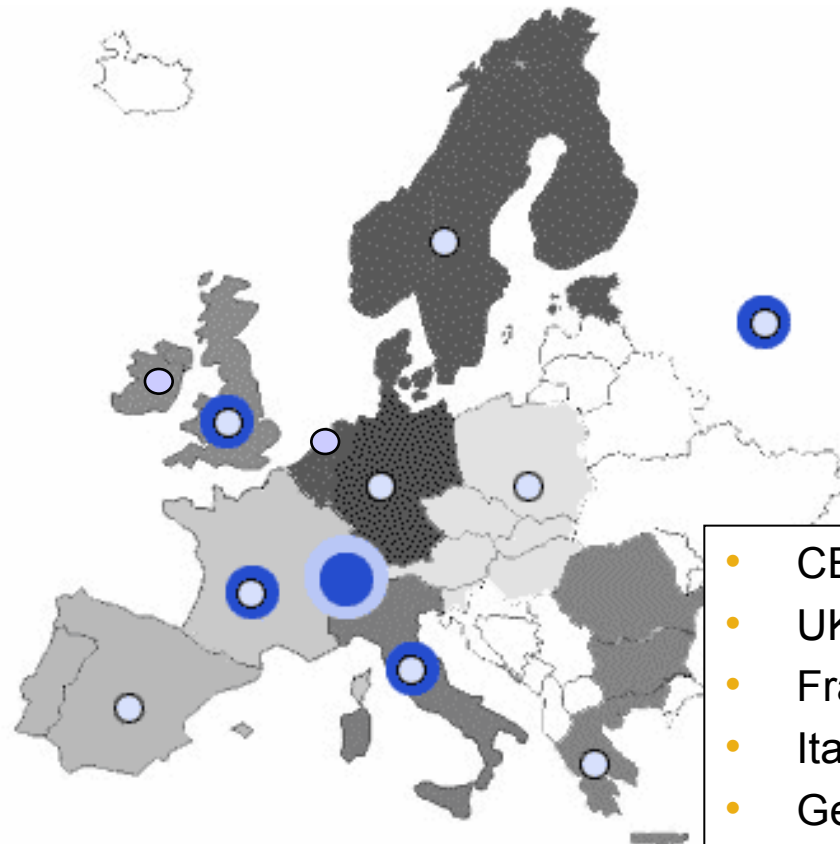
SA1 Objectives

- **Core Infrastructure services:**
 - Operate essential grid services
- **Grid monitoring and control:**
 - Proactively monitor the operational state and performance,
 - Initiate corrective action
- **Middleware deployment and resource induction:**
 - Validate and deploy middleware releases
 - Set up operational procedures for new resources
- **Resource provider and user support:**
 - Coordinate the resolution of problems from both Resource Centres and users
 - Filter and aggregate problems, providing or obtaining solutions
- **Grid management:**
 - Coordinate Regional Operations Centres (ROC) and Core Infrastructure Centres (CIC)
 - Manage the relationships with resource providers via service-level agreements.
- **International collaboration:**
 - Drive collaboration with peer organisations in the U.S. and in Asia-Pacific
 - Ensure interoperability of grid infrastructures and services for cross-domain VO's
 - Participate in liaison and standards bodies in wider grid community

SA2 Objectives

- Ensure that EGEE has access to appropriate networking services provided by GEANT and the NRENs. This includes:
 - Definition of requirements,
 - Specification of services technically and operationally
 - Monitoring of service-level provision
- Define policies for Grid access to the network

Operations Infrastructure



- Operations Management Centre
- Core Infrastructure Centre
- Regional Operations Centre

- CERN (OMC, CIC)
- UK+Ireland (CIC, ROC)
- France (CIC, ROC)
- Italy (CIC, ROC)
- Germany+Switzerland (ROC)
- Northern Europe (ROC)
- South West Europe (ROC)
- South East Europe (ROC)
- Central Europe (ROC)
- Russia (CIC – M12, ROC)
- 48 Partners involved in SA1
- ROC's in several regions are distributed across many sites

~80 funded and
~100 unfunded FTE

The Regional Operations Centres

- The ROC organisation is the focus of SA1 activities:
 - Coordinate and support deployment
 - Coordinate and support operations
 - Coordinate Resource Centre management
 - Negotiate and monitor SLA's within the region
 - Negotiate app access to resources within region
 - Coordinate reporting of SA1 partners within region
 - Coordinate planning for the regional activities
 - Teams:
 - Deployment team
 - 24hour support team (answers user and rc problems)
 - Operations training at RC's
 - Organise tutorials for users
- The ROC is the first point of contact for all:
 - New sites joining the grid and support for them
 - New users and user support

Core Infrastructure Centres

- “Grid Operations Centres” – behaving as a single organisation
- Operate infrastructure services
 - VO services:
 - VO servers, VO registration service
 - RBs, UIs
 - RLS and other database services
 - BDIIs
 - Ensure recovery procedures and fail-over (between CICs)
- Act as Grid Operations Centre
 - Monitoring, proactive troubleshooting
 - Performance monitoring
 - Control sites’ participation in production service
 - Use work done at RAL for LCG GOC as starting point
- Support to ROCs for operational problems
- Operational configuration management and change control
- Accounting and resource usage/availability monitoring

Operations Management Centre

- Located at CERN
- Coordinate operations and management
 - Via ROC managers, CIC managers, policy body
 - Provide security oversight and coordination
 - Coordinate SLA's between regions
- Coordinates with International grid projects
 - Negotiate interoperation policies and frameworks
 - Set up joint projects to address common issues
- Activity coordination
 - Edit execution and implementation plans
 - Coordinate reporting
 - Edit release notes
 - Edit planning guide (cookbooks)

Coordination bodies

- ROC Managers
 - Coordinator – Cristina Vistoli (INFN)
- CIC Managers
 - Need coordinator – need to agree how they work together
- Operations Management
 - OMC, ROC managers, CIC managers, SA2, reps from NA4
 - Resource negotiation policy body – as a subgroup
 - Security group – relationship with JRA3
 - This group is the OAG explained in the TA
- Forum for RC system admins/managers
 - Start a series of workshops (with NA3)



Enabling Grids for
E-science in Europe

The services and test-beds



Production service

- Main production service for production applications
- MUST run reliably, runs only proven stable, debugged middleware and services
 - May be 2 levels – level 1 : certified production; level 2 : awaiting certification (new, or recovering from problems) – controlled by CIC operations centre
- Full support – 24x7 as soon as possible
 - Start with 16x(5-7?) – rotation of coverage between CICs
- Initial service is in place – LCG-2
- Want to add new sites in EGEE federations
 - They join via their ROCs who help deploy middleware

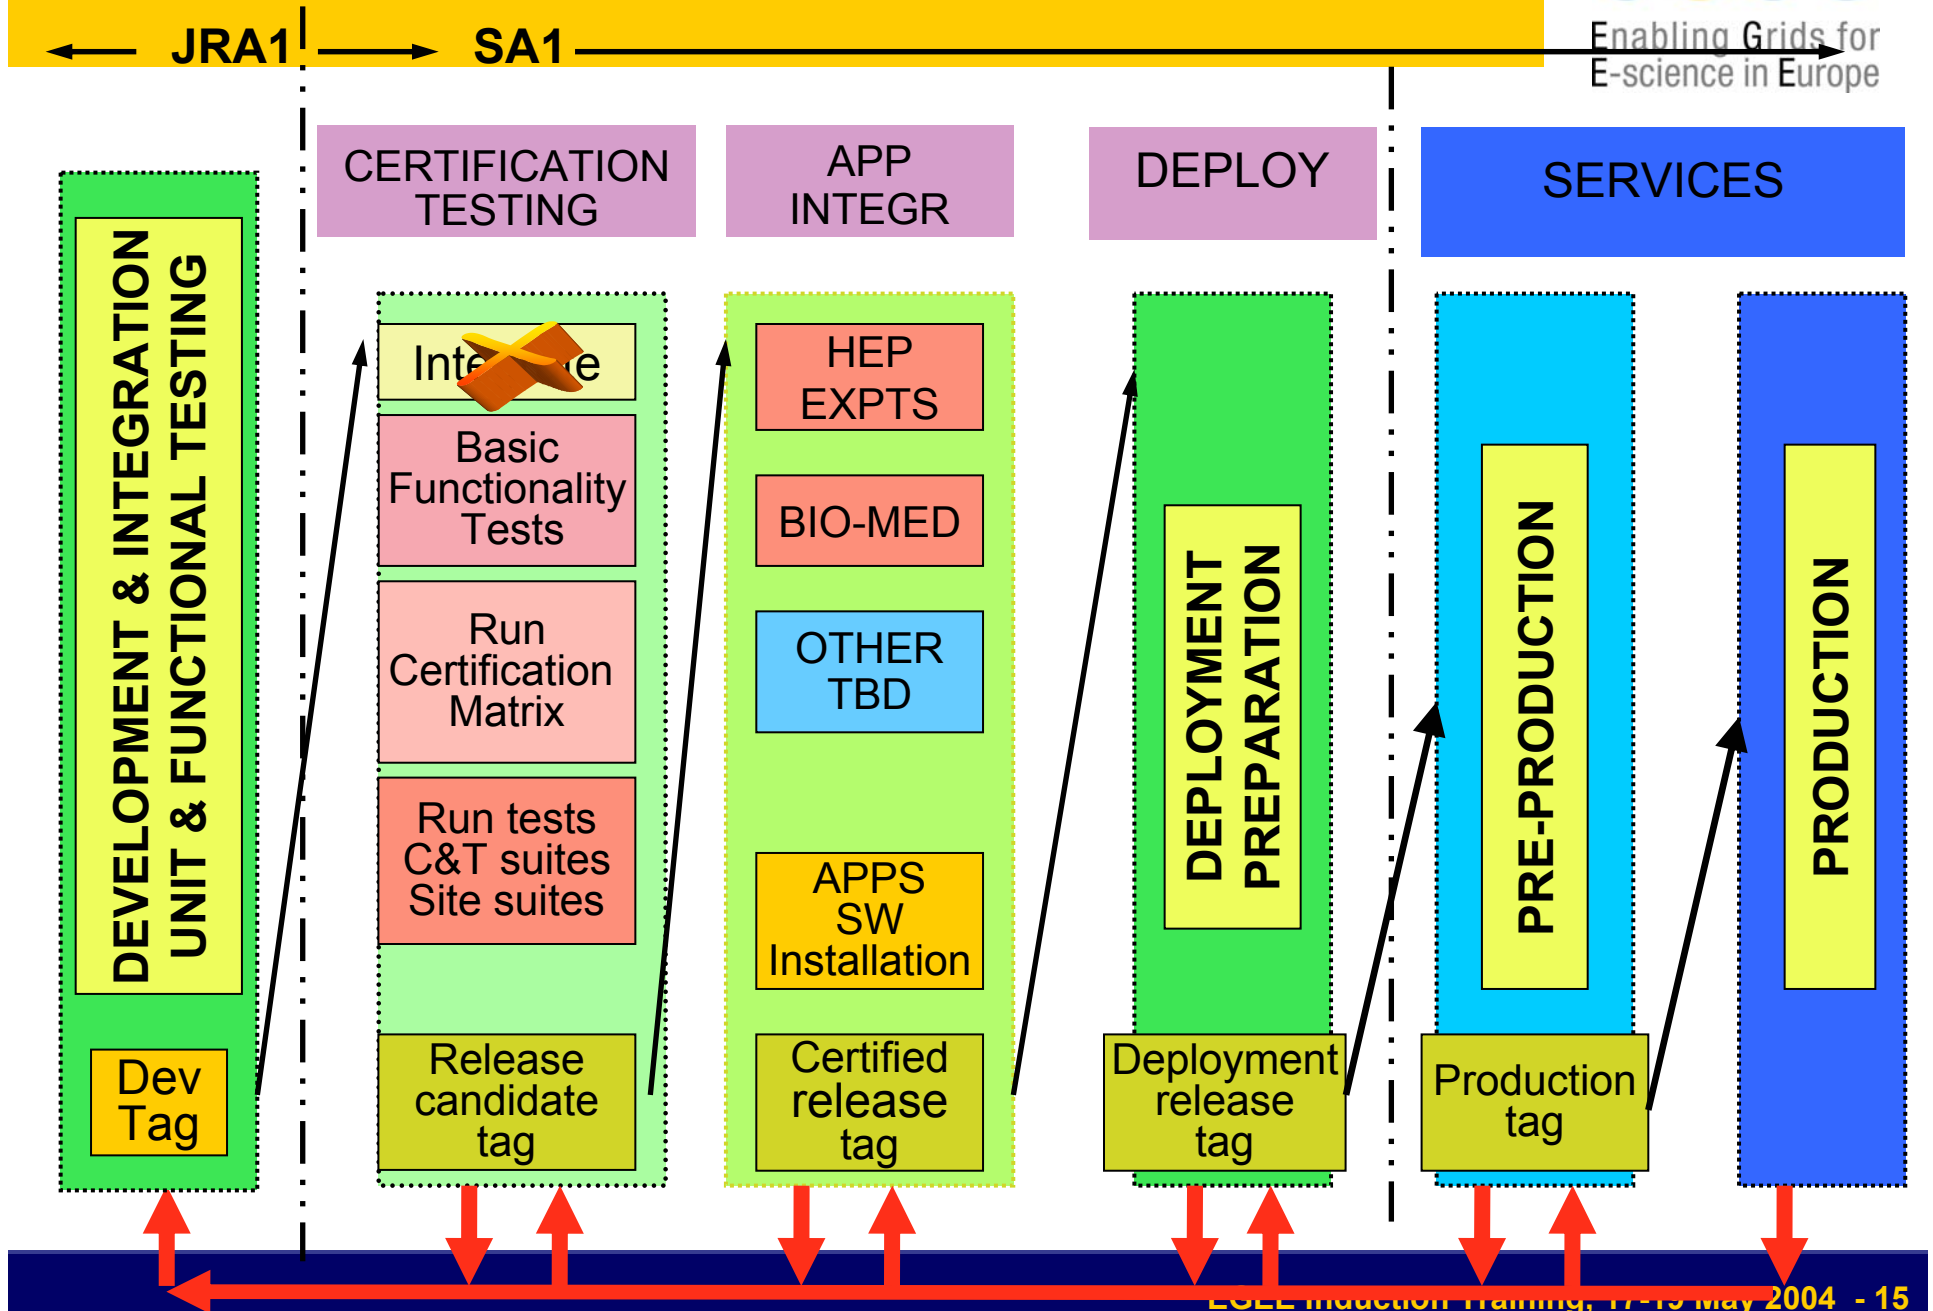
Pre-production service

- For next version middleware
- Initially – start with EGEE middleware as soon as there is a basic release
 - For year 1 pre-prod will run EGEE mw, prod will run LCG-2
 - When EGEE mw ready – move to production and then the pre-prod service will be next EGEE candidate release
 - Even incremental component changes – get away from big-bang changes
 - Expect to updated services on pre-prod even 1 by 1
- Feedback from users, ROC's, CIC's, RC's is essential – this service must be widely deployed but does not need huge resources
- Initial resources – come from EDG app testbed sites, perhaps also some of the new smaller sites
- While waiting for first EGEE release – will deploy LCG-2 to get pre-prod system up
- Need a pre-production service coordinator
- Support is 8x5

Certification test-beds

- Certification test-bed at CERN
- Need some validation with ROC's before going to pre-production
 - ROCs should provide these resources
- Need resources for porting
 - E.g. If a region has particular need – port to their favourite OS and certify middleware

Certification, Testing and Release Cycle



Training/demo service

- Permanent need for tutorials, demonstrations etc.
- Cannot disturb production system, or guarantee pre-production
- Ideally need dedicated (small) service that can be booked for tutorials etc
 - But must be kept in an operational state
 - Need sufficient resources to be available (another testbed!)
- Perhaps can do this via info system and dedicated queues, VO's, etc within the production system
 - Needs some thought to set this up
- This may now be partly addressed by GILDA service (see NA3)

Some remarks

- Existing LCG-2 sites already support many VOs
 - Not only LCG
 - Front-line support for all VOs is via the ROCs
- Process to introduce a new VO
 - Well defined
 - Some tools needed to make the mechanics simpler
- Evaluation of new middleware by applications, and preparation for deployment in EGEE-1
 - This is what the pre-production service is for
- Resource allocation/negotiation
 - OMC/ROC managers/NA4 – negotiate with RC's and apps

Joining EGEE – Overview of process

- Application nominates VO manager
- Find (CIC) to operate VO server
- VO is added to registration procedure
- Determine access policy:
 - Propose discussion (body) NA4 + ROC manager group
 - Which sites will accept to run app (funding, political constraints)
 - Need for a test VO?
- Modify site configs to allow the VO access
- Negotiate CICs to run VO-specific services:
 - VO server (see above)
 - RLS service if required
 - Resource Brokers (can be some general at CIC and others owned by apps), UIs – general at CIC/ROC – or on apps machines etc
 - Potentially (if needed) BDII to define apps view of resources
- Application software installation
 - Understand application environment, and how installed at sites
- Many of these issues can be negotiated by NA4/SA1 in a short discussion with the new apps community

Resource Negotiation Policy

- The EGEE infrastructure is intended to support and provide resources to many virtual organisations
 - Initially HEP (4 LHC experiments) + Biomedical
 - Each RC supports many VOs and several application domains – situation now for centres in LCG, EDG, EDT
- Initially must balance resources contributed by the application domains and those that they consume
 - Maybe specifically funded for one application
 - In 1st 6 months sufficient resources are committed to cover requirements
- Allocation across multiple sites will be made at the VO level.
 - EGEE will establish inter-VO allocation guidelines
 - E.g. High Energy Physics experiments have agreed to make no restrictions on resource usage by physicists from different institutions
- Resource centres may have specific allocation policies
 - E.g. due to funding agency attribution by science or by project
 - Expect a level of peer review within application domains to inform the allocation process

Resource allocation – 2

- New VOs and Resource centres will be required to satisfy minimum requirements
 - Commit to bring a level of additional resources consistent with their requirements
 - The project must demonstrate that on balance this level of commitment is less than that required for the user community to perform the same work outside the grid
 - The difference will come from the access to idle resources of other VOs and resource centres
 - This is the essence of a grid infrastructure
- All compute resources made available to EGEE will be connected to the grid infrastructure.
 - Significant potential for sites to have additional resources
 - A small number of nodes at each site will be dedicated to operating the grid infrastructure services
- Requirement on JRA1 to provide mechanisms to implement/enforce quotas, etc
- Selection of new VO/RC via NA4
 - In accordance with policies designed and proposed by e-IRG (NA5)

New Resource Centres

- Procedure for new sites to join LCG2/EGEE is well defined and documented
- Sites can join now
- Coordination for this is via the ROCs
 - Who will support the installations, set-up, and operation

Security Issues

- SA1 and JRA3 both have security responsibility
 - SA1 – operational security
- CA's – JRA3
 - Procedures for accepting new CA's
 - Operation of Catch-All CA (CNRS)
 - SA1 runs CERN CA
- Operational security
 - Security group based on LCG group and its work
- VO-Management and policies
- Incident Response
- Security Audit
- Accounting
 - Integrity, access and privacy (policies needed)
- Rules of Conduct
- Service Level Agreements

Expected Computing Resources

Region	CPU nodes Month 1	Disk (TB) Month 1	CPU Nodes Month 15	Disk (TB) Month 15
CERN	900	140	1800	310
UK + Ireland	100	25	2200	300
France	400	15	895	50
Italy	553	60.6	679	67.2
North	200	20	2000	50
South West	250	10	250	10
Germany + Switzerland	100	2	400	67
South East	146	7	322	14
Central Europe	385	15	730	32
Russia	50	7	152	36
Totals	3084	302	8768	936

Month 24

resource centres

10

20

50

Sites in LCG-2/EGEE-0

- No information
- Scheduled Maintenance
- Timeout
- Globus OK
- RB OK



Status for Resource Broker CERN_lxn1188: Fri May 14 16:45:26 BST 2004

Regional Centres Connected to the LCG Grid

07-May-04

<i>country</i>	<i>centre</i>	<i>country</i>	<i>centre</i>
Austria	UIBK	Portugal	LIP, Lisbon
Canada	TRIUMF, Vancouver	Russia	SINP, Moscow
	Univ. Montreal	Spain	PIC, Barcelona
	Univ. Alberta		IFIC, Valencia
Czech Republic	CESNET, Prague		IFCA, Santander
	University of Prague		University of Barcelona
France	IN2P3, Lyon**		Uni. Santiago de Compostela
Germany	FZK, Karlsruhe		CIEMAT, Madrid
	DESY		UAM, Madrid
	University of Aachen	Switzerland	CERN
	University of Wuppertal		CSCS, Manno**
Greece	GRNET, Athens	Taiwan	Academia Sinica, Taipei
Holland	NIKHEF, Amsterdam		NCU, Taipei
Hungary	KFKI, Budapest	UK	RAL
Israel	Tel Aviv University**		Cavendish, Cambridge
	Weizmann Institute		Imperial, London
Italy	CNAF, Bologna		Lancaster University
	INFN, Torino		Manchester University
	INFN, Milano		Sheffield University
	INFN, Roma		QMUL, London
	INFN, Legnaro	USA	FNAL
Japan	ICEPP, Tokyo**		BNL**
Poland	Cyfronet, Krakow		

** not yet in LCG-2



Enabling Grids for
E-science in Europe



> 40 sites

> 3,100 CPUs

Centres in process of being connected

<i>country</i>	<i>centre</i>
China	IHEP, Beijing
India	TIFR, Mumbai
Pakistan	NCP, Islamabad

Hewlett Packard to provide "Tier 2-like" services for
LCG, initially in Puerto Rico

EGEE Induction Training, 17-19 May

SA2: Network Resource Provision Goals, Objectives and Approach

- Ensures EGEE access to network services provided by GEANT and the NRENs to link users, resources and operational management

Do this by managing the relationship between EGEE and GEANT

- Tasks

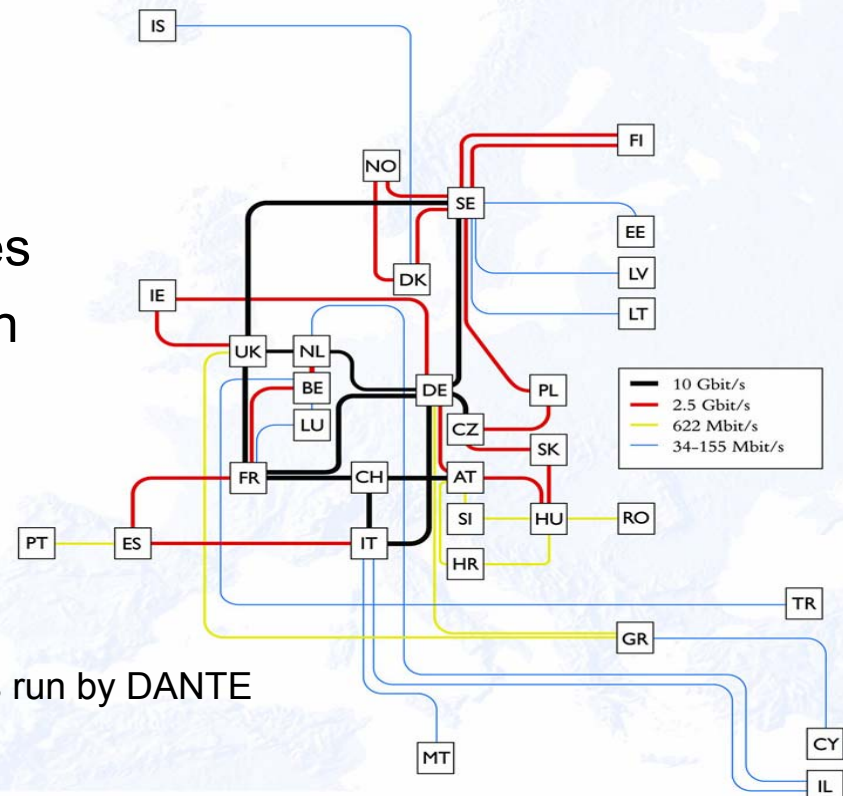
- Definition of requirements
- Specification of services
- Definition of network access policies
- Monitoring of service level provision



GEANT High-speed pan-European backbone linking NRENs run by DANTE

NRENs National Research and Educational Networks

DANTE Not-for-profit company that manages GEANT



SA2 Approach : Network services

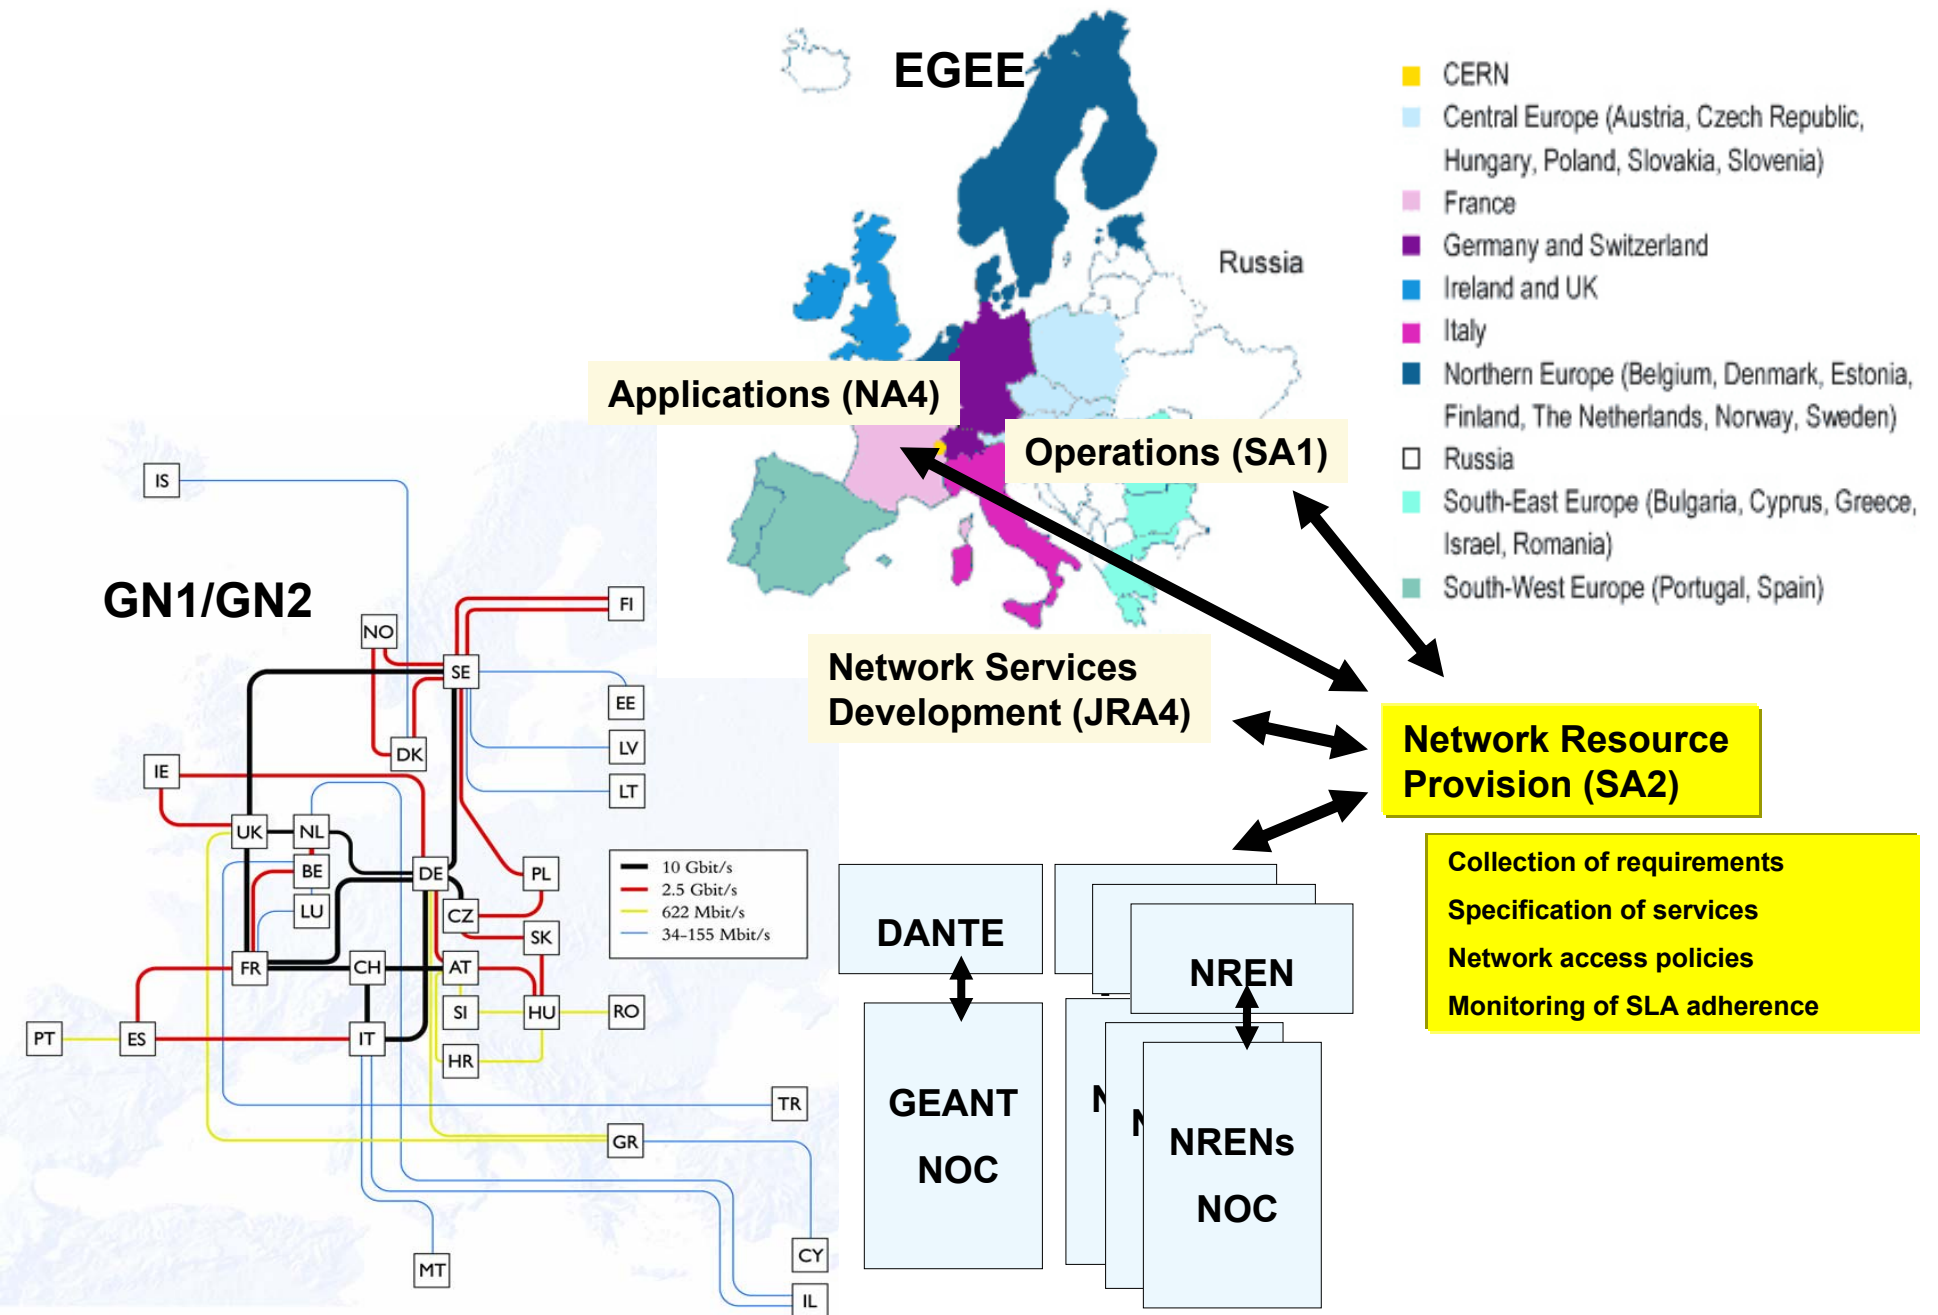
- Definition through standard modelling process :
 - Filling of SLRs (Service Level Request) by end users and applications
 - Ex : Bandwidth allocation
 - Flow classification, MPLS VPN (L2, L3), GMPLS, Lightpath.
 - Definition of SLSs (Service Level Specification) by SA2, to be implemented by GEANT and the NRENs, in conjunction with JRA4 activity
 - Signature of SLAs (Service Level Agreement)
 - Client : Operations, Applications, Virtual Organizations ?
 - GEANT/NRENs

SA2 Approach : Operational Interface

- Network Operation Centre (NOC) operational procedure study on GEANT and NRENS
 - To select NRENS
 - In EGEE : GARR, DFN, GRNET, CESNET,
 - Outside EGEE : SURFNET, RENATER, UKERNA ...
- Incremental integration with EGEE GOCs
 - Trouble Ticket systems study.
 - Define interfaces.
 - Homogenous system at the EGEE level.

SA2 Approach

- Outside the scope of EGEE to provide connections for any user or resource site
 - Sites must have adequate bandwidth & performance to join the production grid facility.
 - EGEE can help a particular site to improve its connectivity.
- Go beyond existing best effort IP service to meet the needs of a production level grid network.
- Network provision can itself be view as a class of Grid resource.



SA2 Team

UREC will manage SA2 and oversee both SA2 and JRA4 activities, and will be responsible for DANTE and the NRENs liaison

Participant	Description of Role	FTE (EU funded + unfunded)
CNRS/UREC	Network Co-ordinator overseeing both service (SA2) and research activities (JRA4); responsible for DANTE and the NRENs liaison. Network resource provision requirements SLR/SLS/SLA definitions Operational model	Jean-Paul Gautier Mathieu Goutelle
RCC KI	Network resource provision requirements SLR/SLS/SLA definitions Operational interface between RDIG, Russian network providers and EGEE.	Sergei Teryaev X
GRNET	Network resource provision requirements SLR/SLS/SLA definitions	Afrodite Sevasti

SA2 Milestones and deliverables

PM	Deliverable or Milestone	Item
M3	Milestone MSA2.1	First meeting of EGEE-GEANT/NRENS Liaison Board
M6	Deliverable DSA2.1	Survey of pilot application requirements on networks, initial SLRs and service classes.
M9	Milestone MSA2.2	Initial requirements aggregation model, specification of services as SLSs on the networks,
M12	Milestone MSA2.3	Operational interface between EGEE and GEANT/NRENS.
M12	Deliverable DSA2.2	Institution of SLAs and appropriate policies.
M24	Deliverable DSA2.3	Revised SLAs and policies.

Summary

- ~50% of project funding is for Operations (SA1 + SA2)
- 48 partners participate in SA1
- Management – distributed
 - Regional responsibilities in the ROCs
 - Coordinated at CERN
- Production service is operational – based on LCG
 - New VOs and new sites are joining now
- Services on track to meet first 2 milestones
 - Set up ROCs and CICs
- Operations activity has big influence on other areas of the project
 - middleware, security, etc.