



Enabling Grids for E-scienceE

Grid Middleware

Mike Mineter

mjm@nesc.ac.uk

www.eu-egee.org



- **This talk was prepared by Mike Mineter of NeSC and includes slides on LCG-2 from previous tutorials and talks delivered by:**
 - the EDG training team
 - EGEE colleagues
- OGSA-DAI team

- **Has the goal of mapping the grid middleware territory**
 - Identifying what is commonly used in production grids
 - Noting some current gaps and a couple of the landmarks that might be important for Virtual Digital Libraries

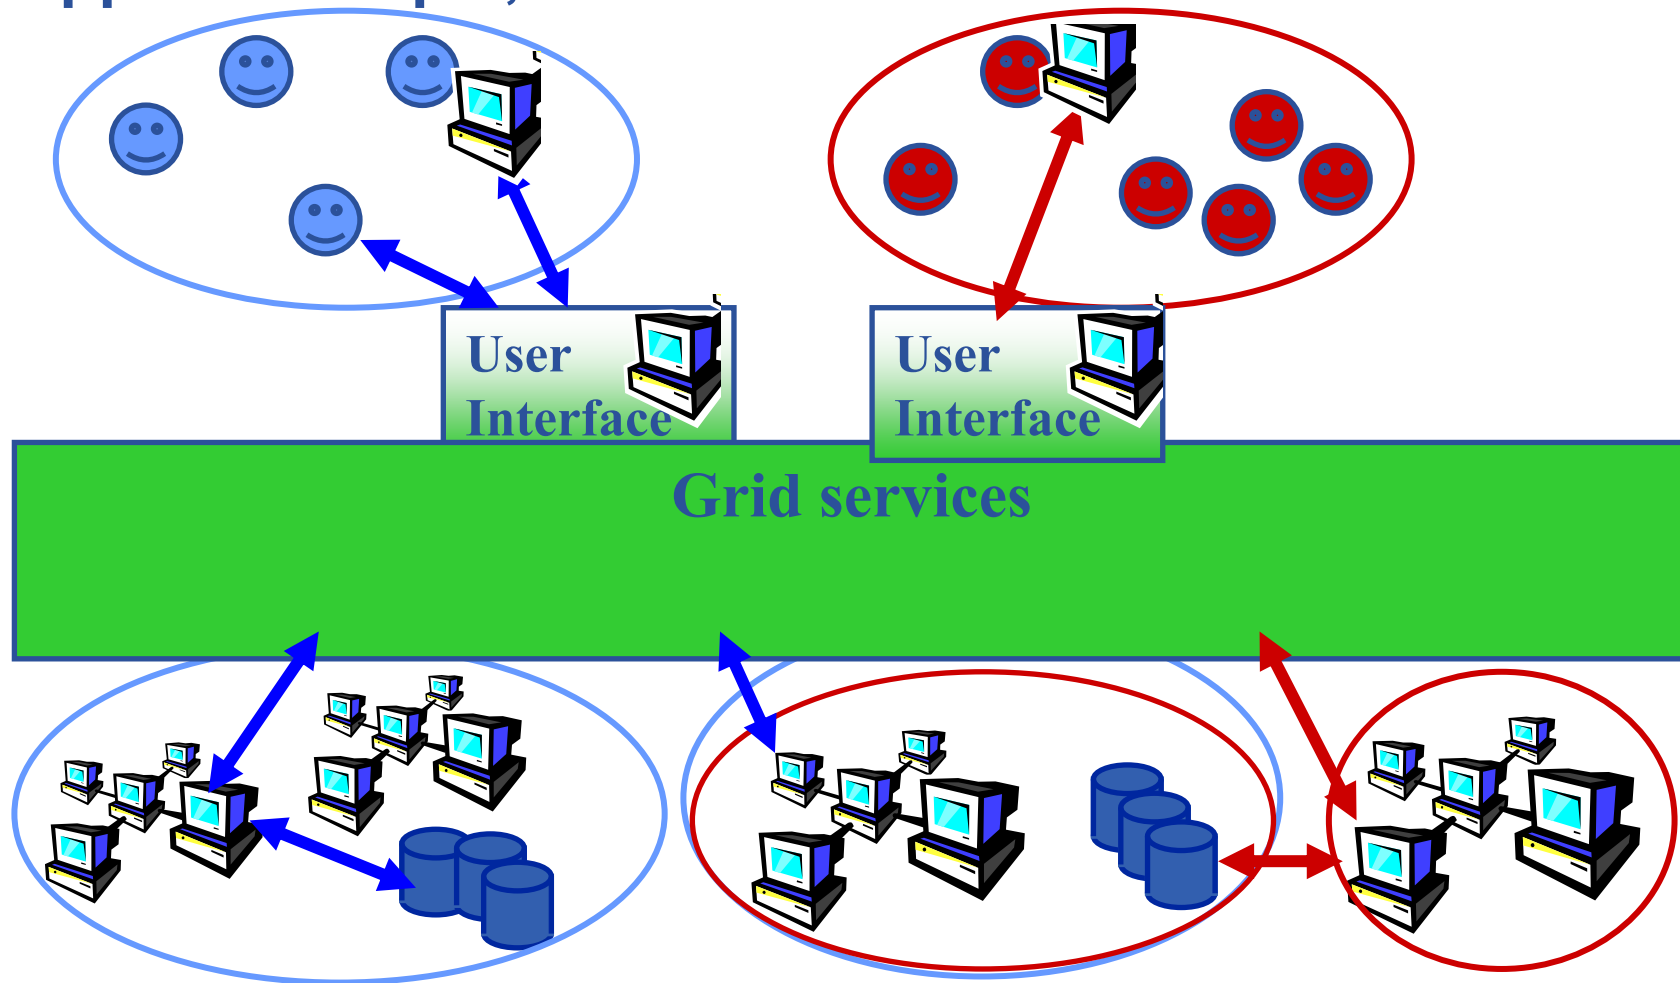
- **Preface: EGEE and DILIGENT**
- **The basic tools and services**
- **Building on the basics: from tools to a grid**
- **More on selected services**
- **Emerging standards**

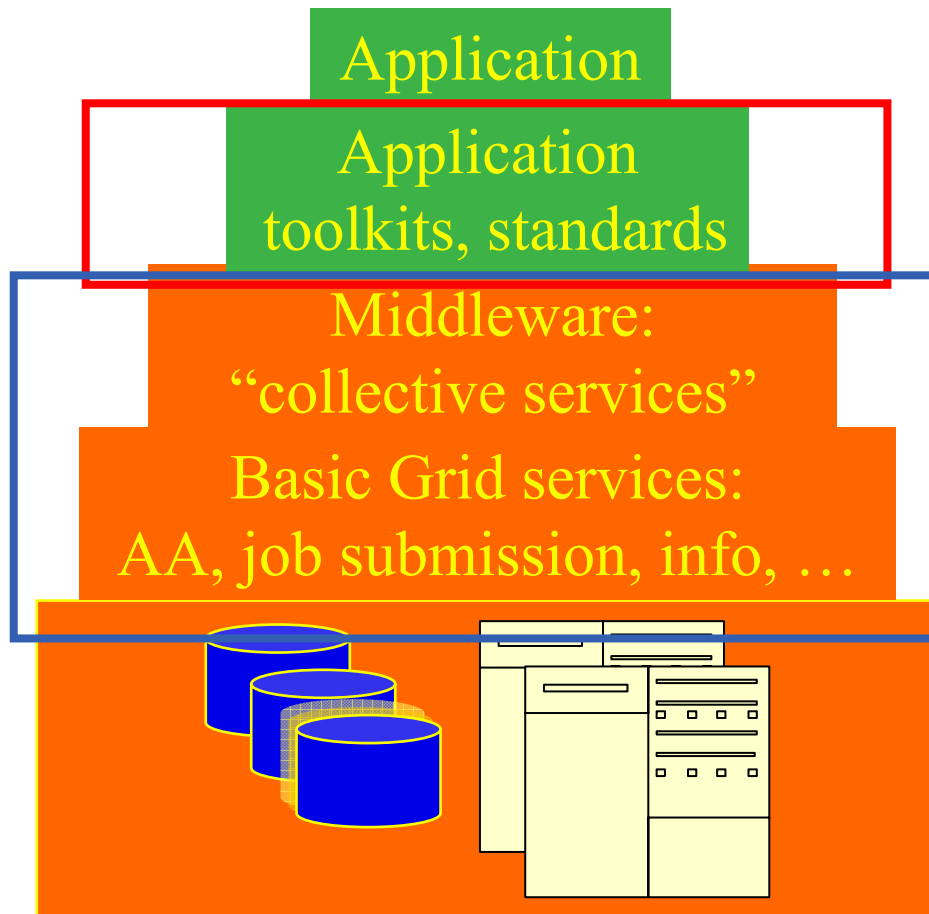


- **Preface: EGEE and DILIGENT**
- The basic tools and services
- Building on the basics: from tools to a grid
- More on selected services
- Emerging standards



- EGEE is establishing a production grid service to support multiple, diverse VO's





- Application development environment, portals,
- Semantics, ontologies


Focus of this talk: Generic grid middleware

- **EGEE is establishing a production grid service to support multiple, diverse Virtual Organisations (VOs)**
- **Yet the hardened middleware on which it is based still has evident roots in research in “big science”**
- **Hence EGEE...**
 - Seeks and embraces new VOs
 - Selects new VOs to set new challenges
 - And looks forward to Virtual Digital Libraries setting quite a few more....

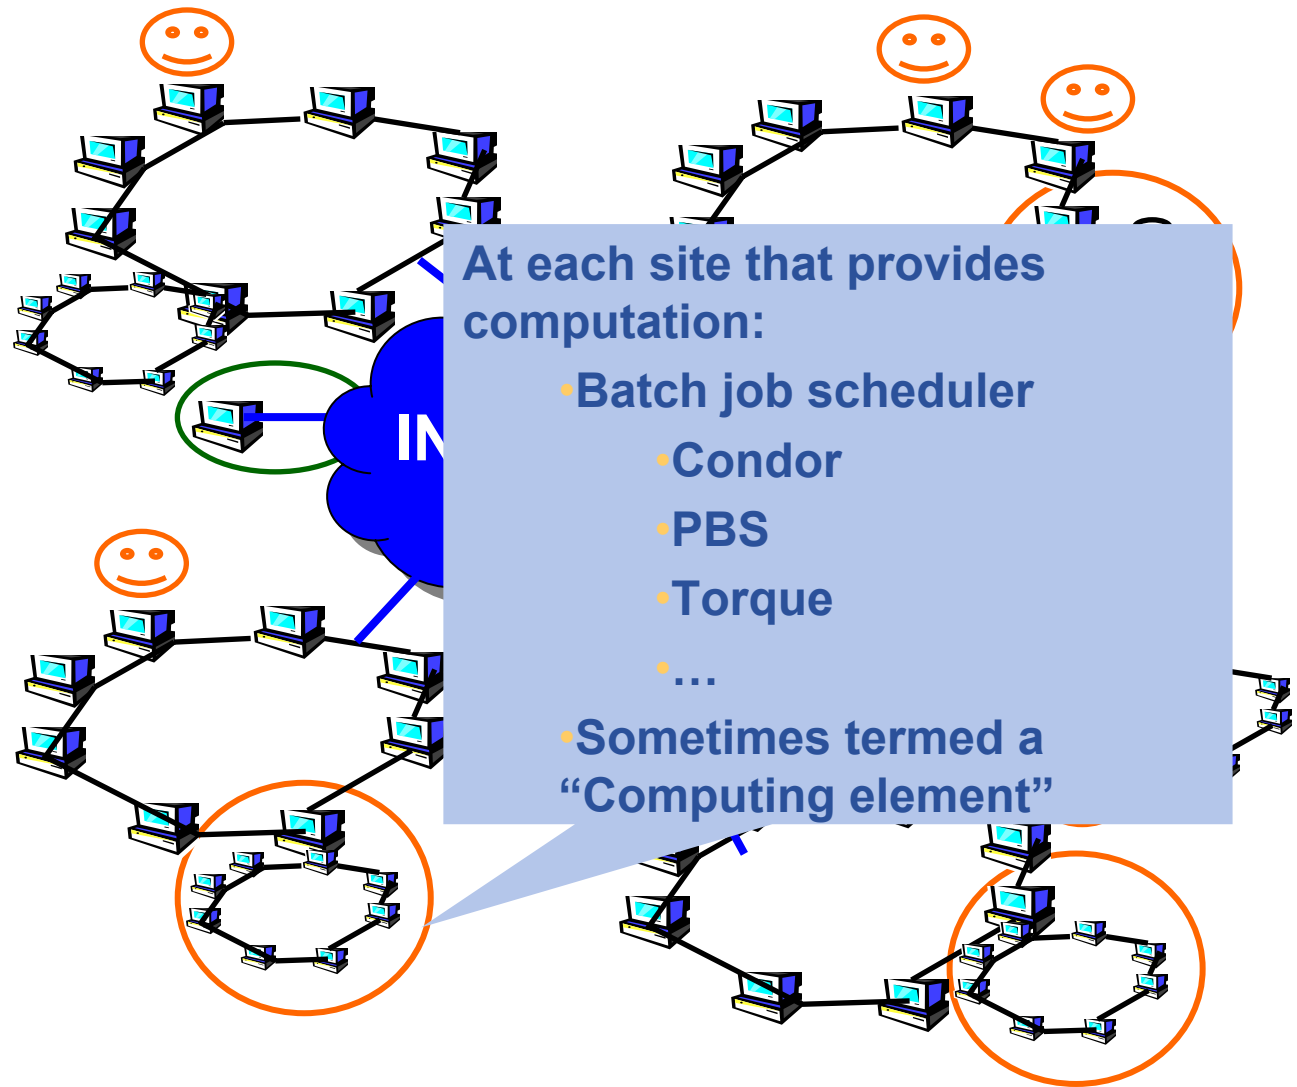
- Preface: EGEE and DILIGENT
- **The basic tools and services**
- Building on the basics: from tools to a grid
- More on selected services
- Emerging standards



- **Grid middleware runs on each shared resource**
 - Data storage
 - (Usually) batch queues on pools of processors
- **Users join VO's**
- **Virtual organisation negotiates with sites to agree access to resources**
- **Distributed services (both people and middleware) enable the grid, allow single sign-on**

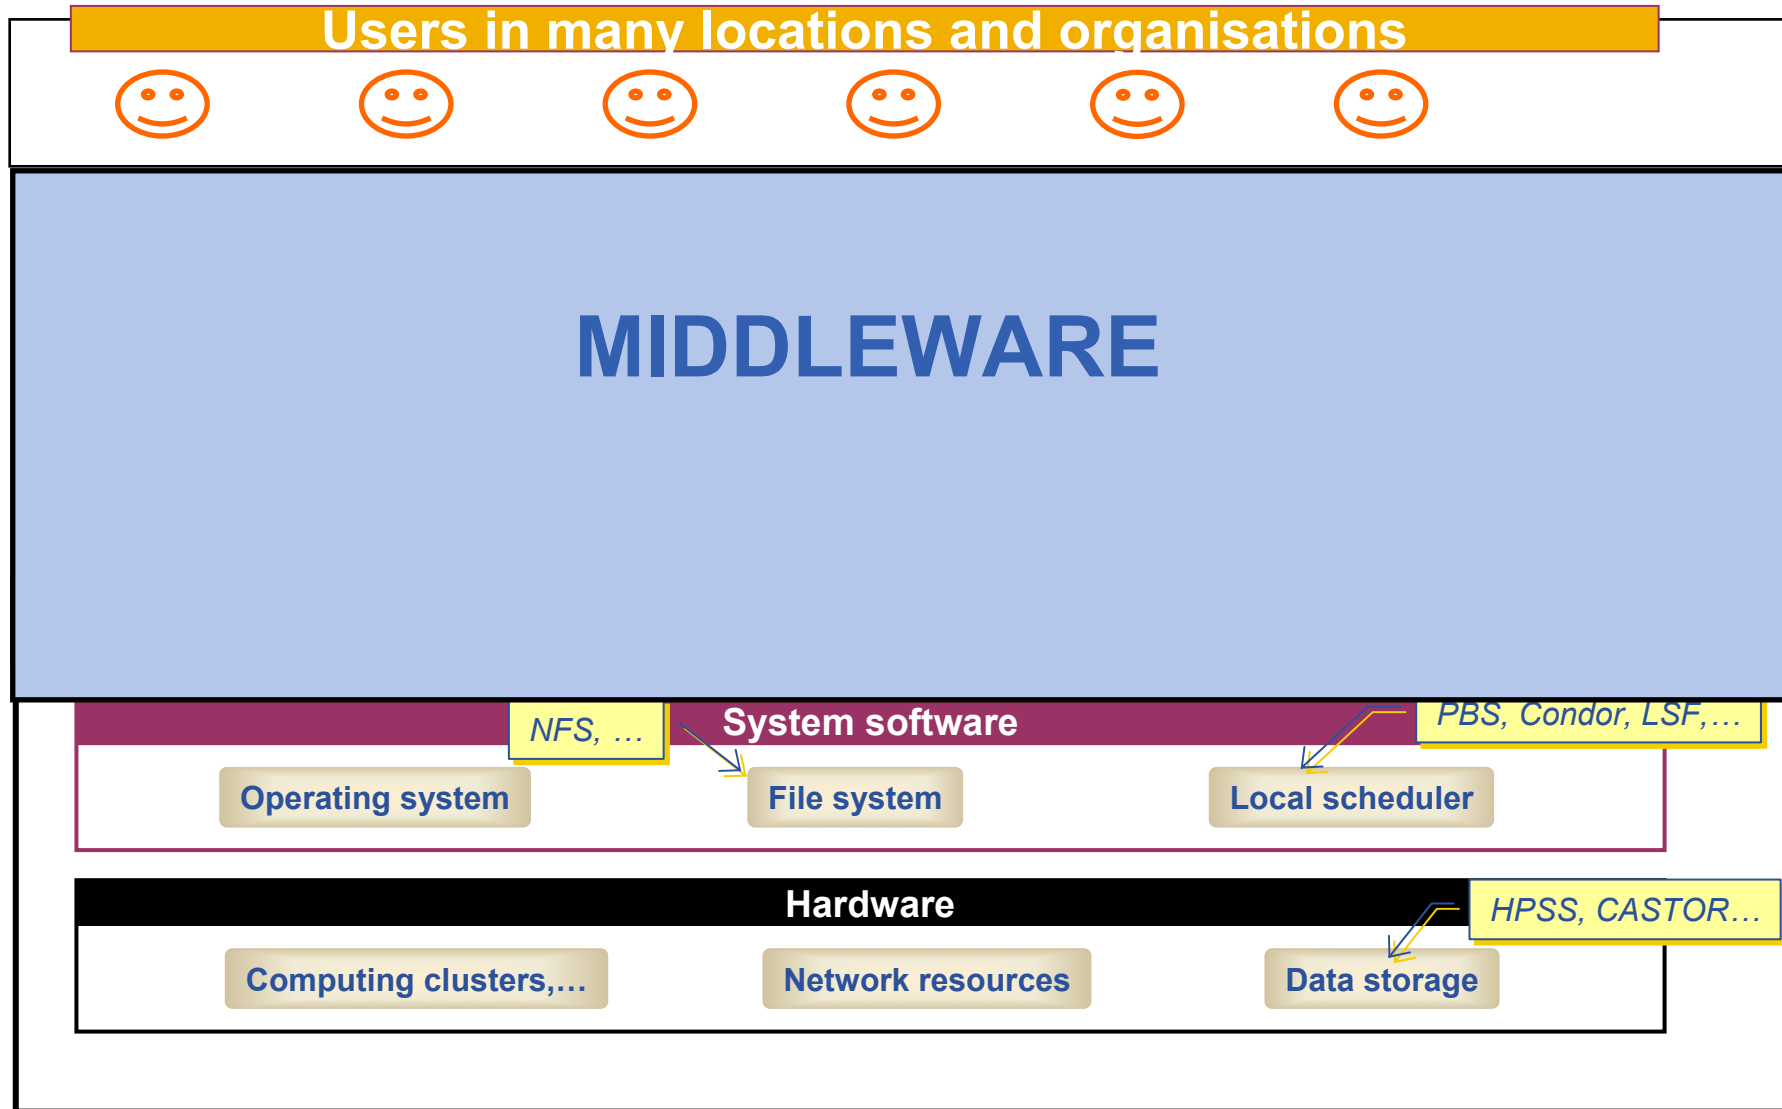


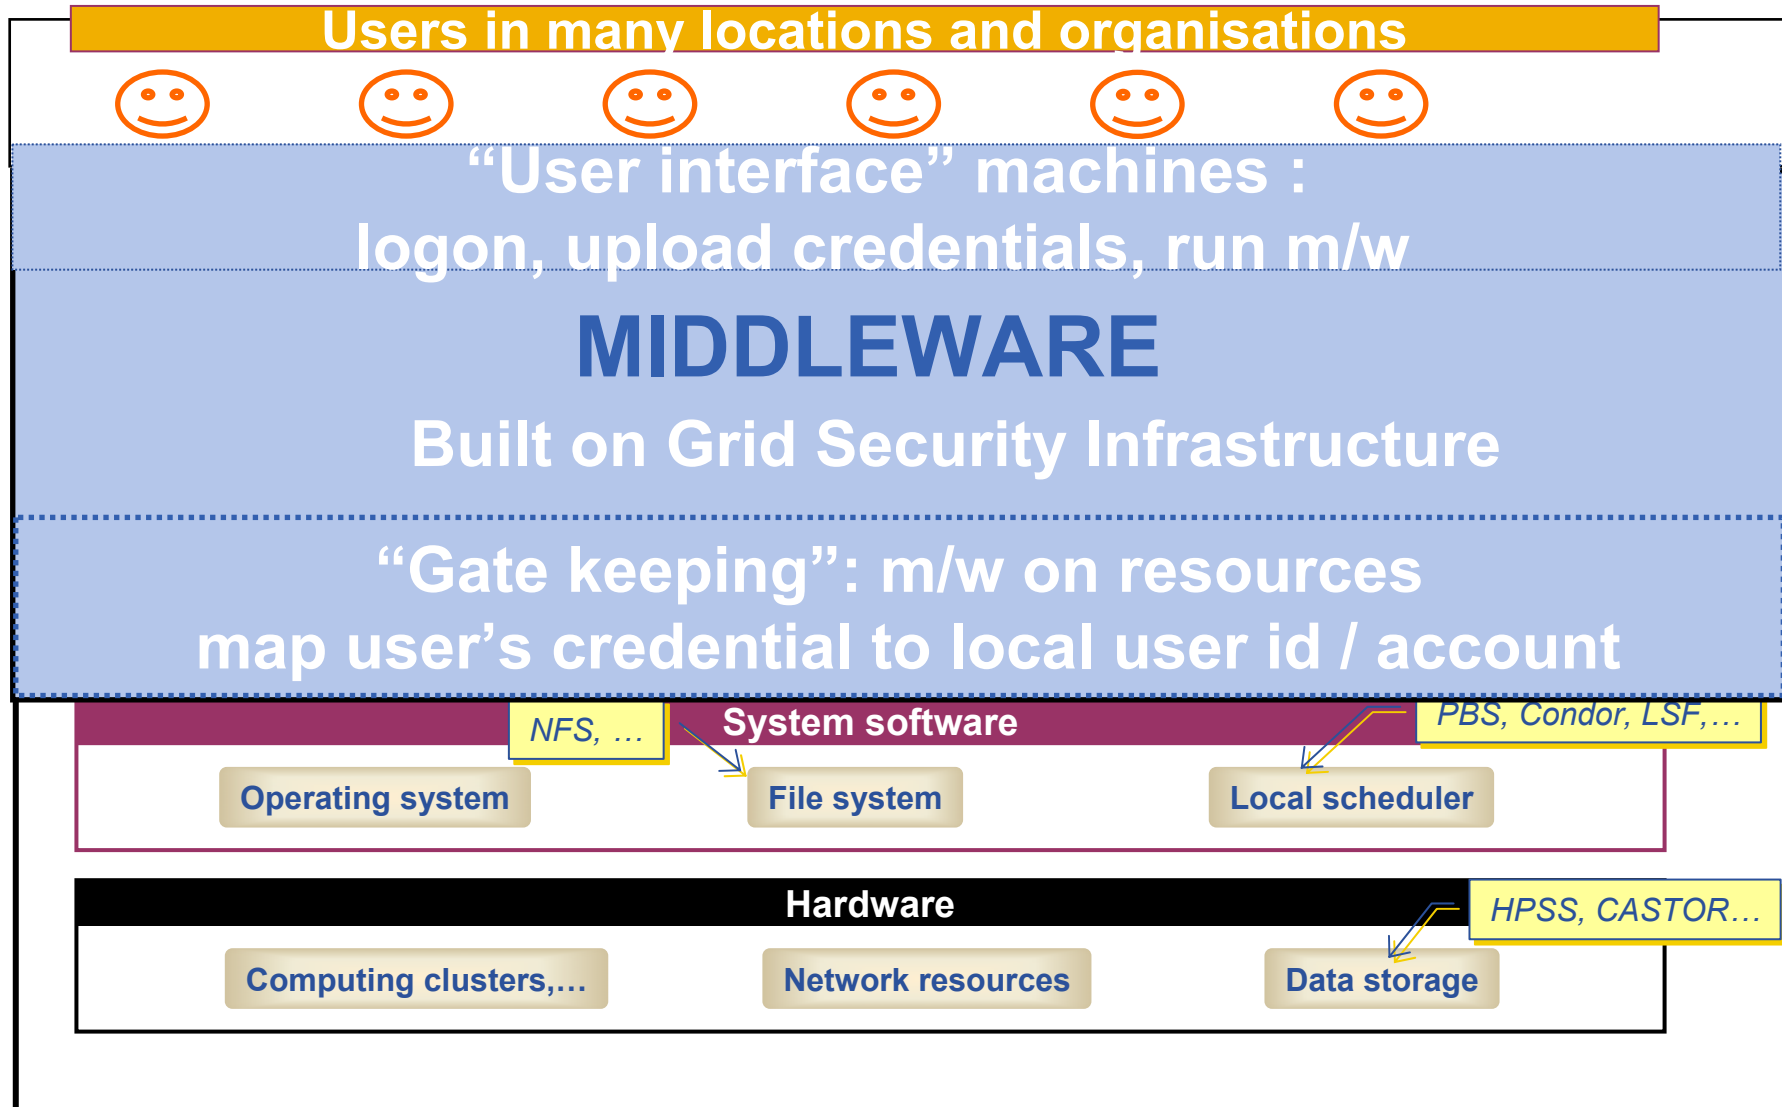
- **Grid middleware runs on each shared resource**
 - Data storage
 - (Usually) batch queues on pools of processors
- **Users join VO's**
- **Virtual organisation negotiates with sites to agree access to resources**
- **Distributed services (both people and middleware) enable the grid, allow single sign-on**

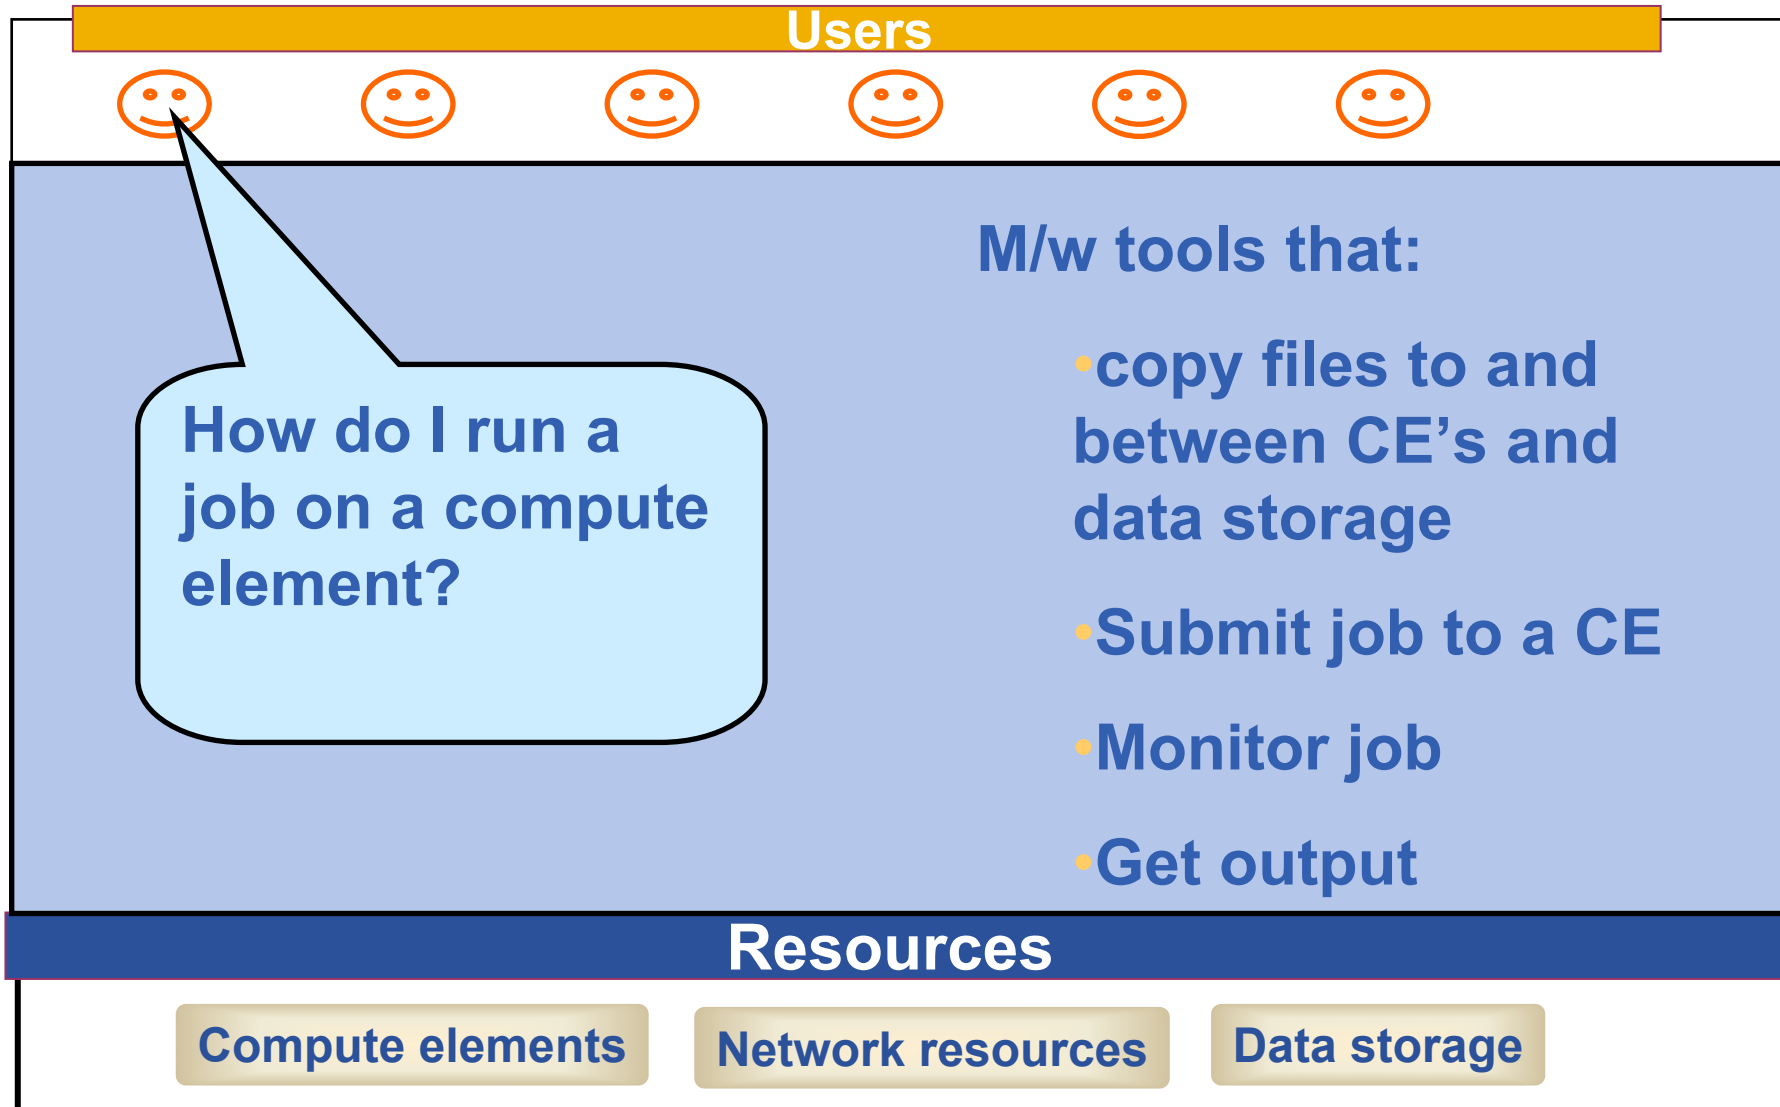


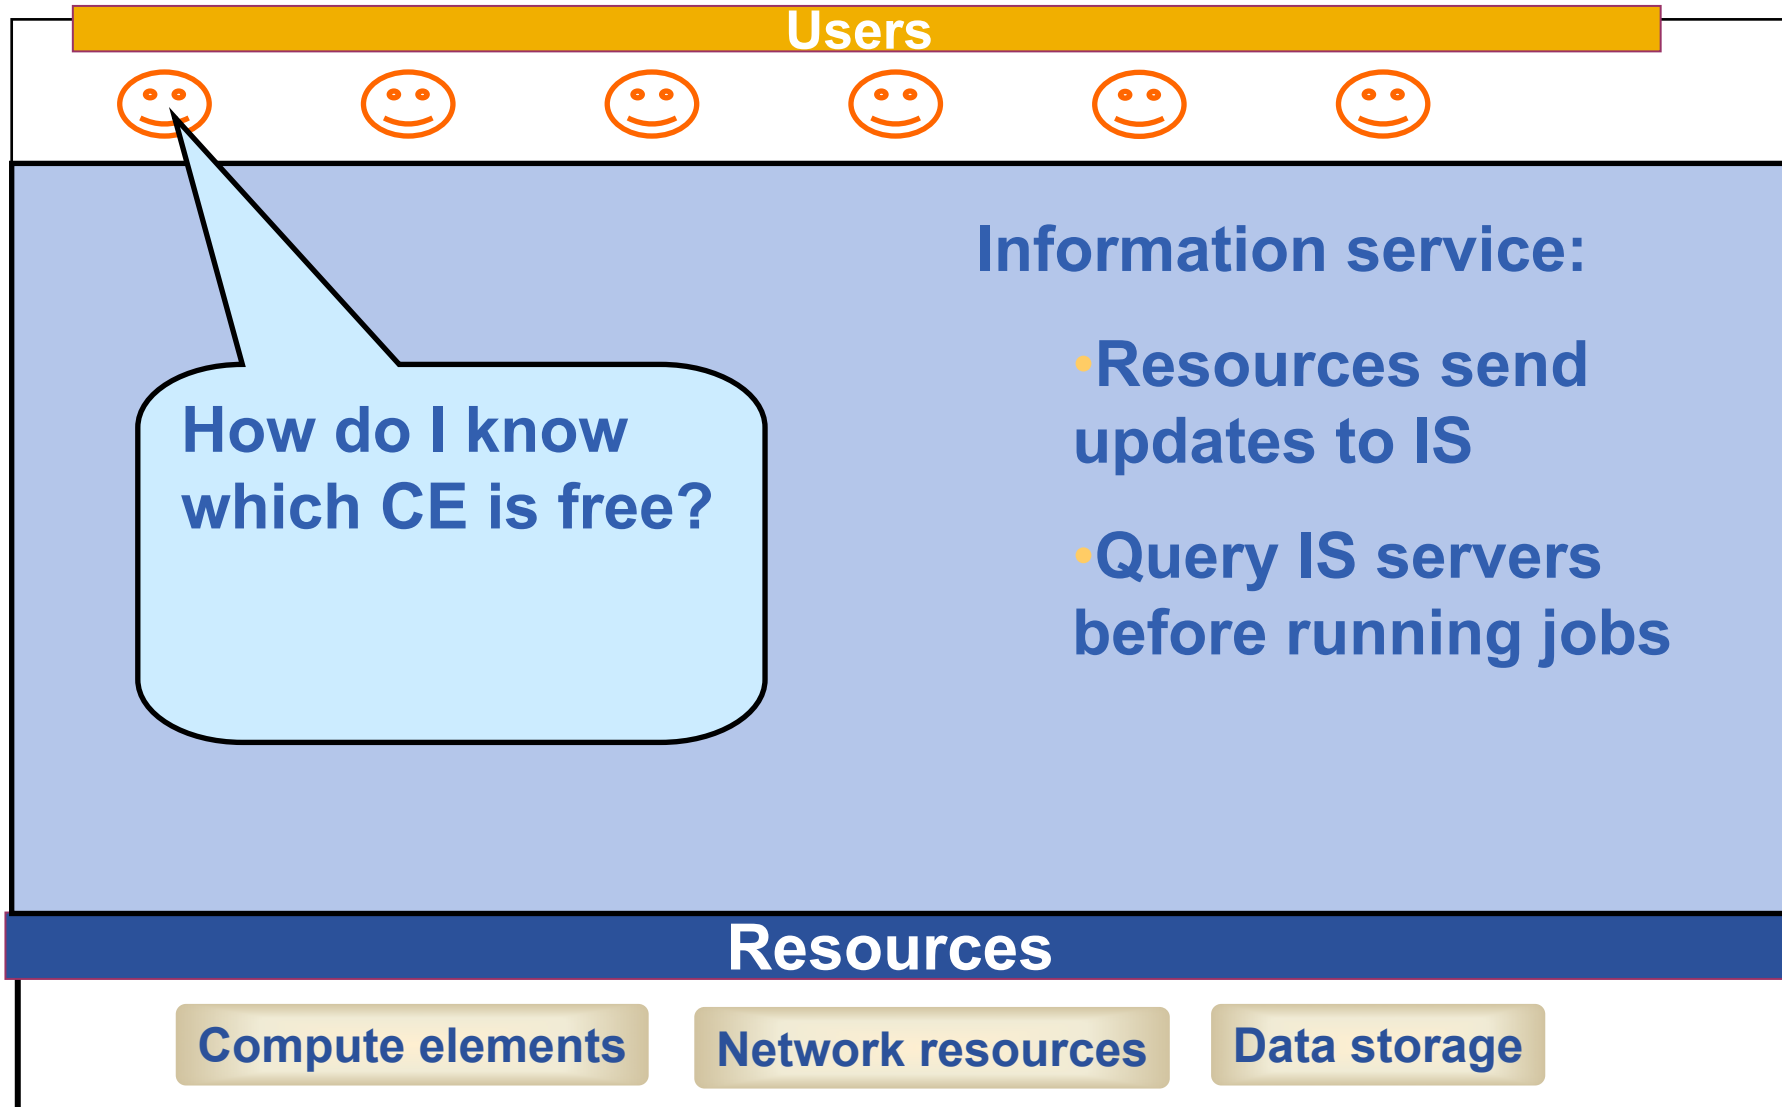
At each site that provides computation:

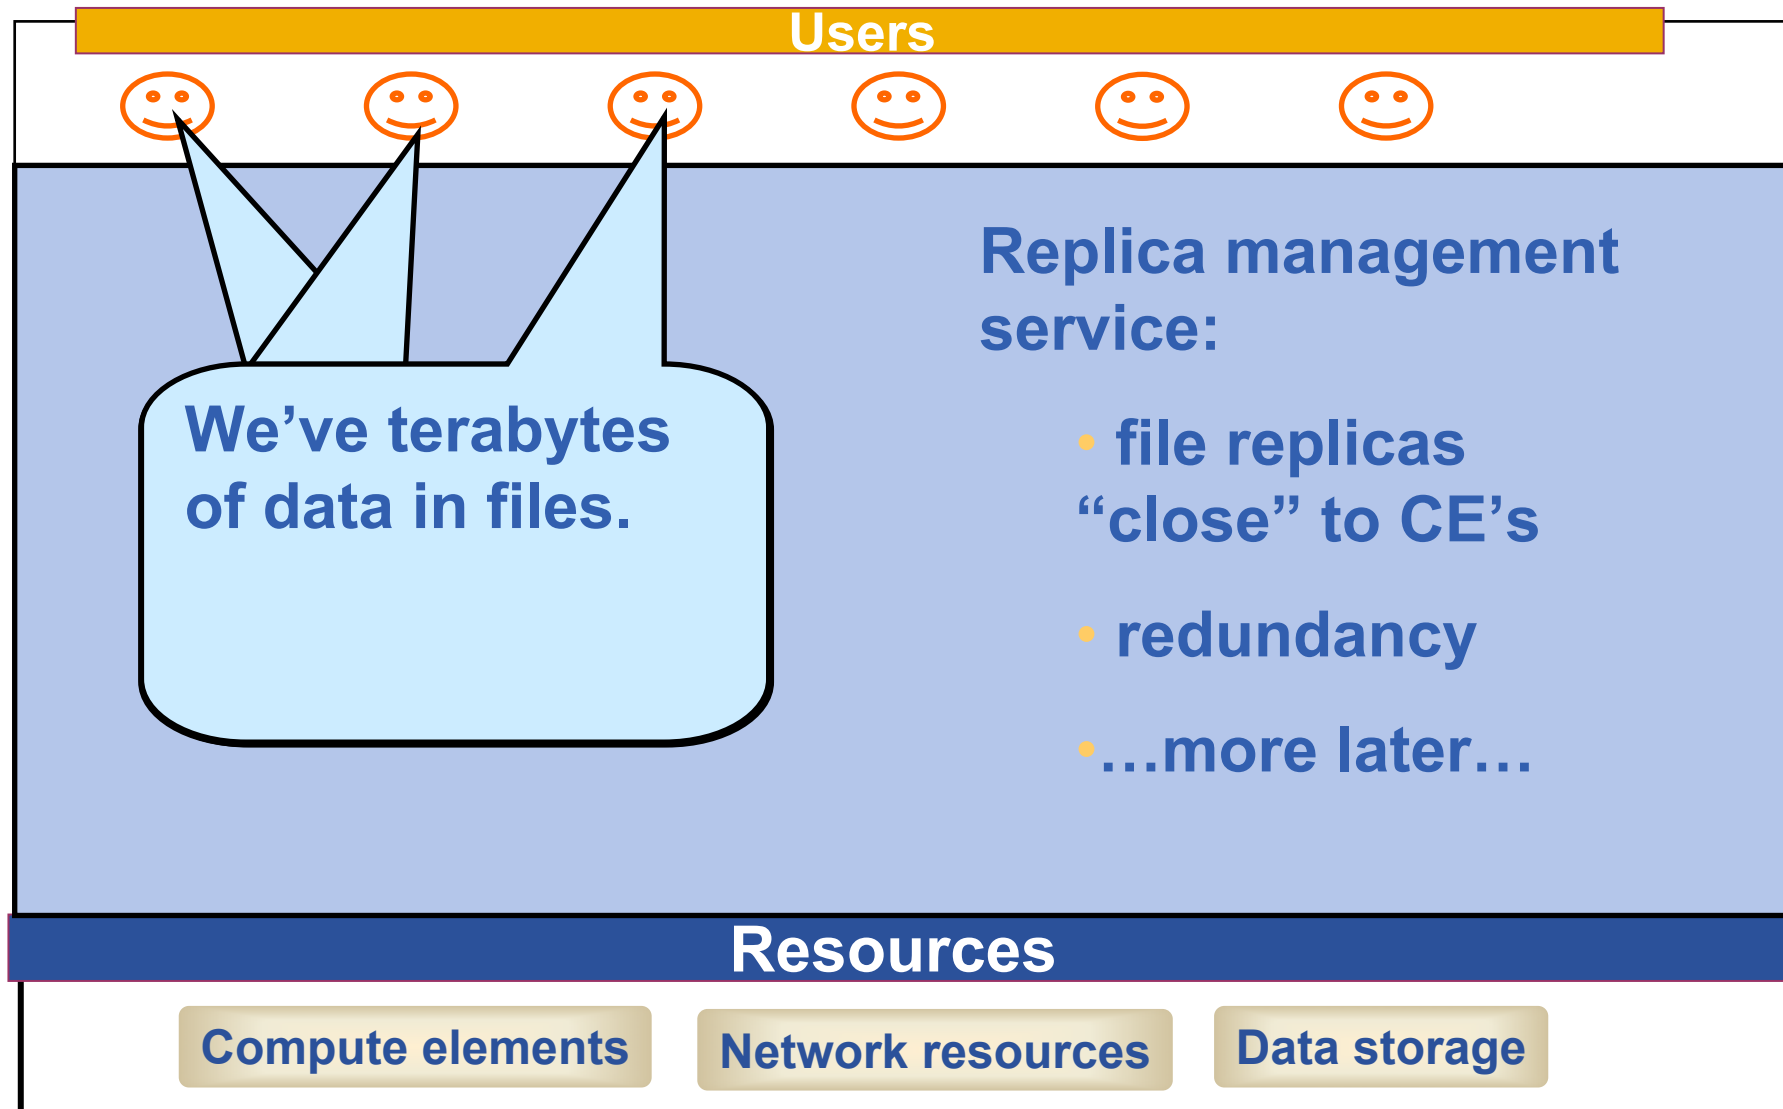
- Batch job scheduler
 - Condor
 - PBS
 - Torque
 - ...
- Sometimes termed a "Computing element"











- **A software toolkit: a modular “bag of technologies”**
 - Made available under liberal open source license
- **Not turnkey solutions, but *building blocks* and *tools* for application developers and system integrators**
- **Tools built on Grid Security Infrastructure to include:**
 - Job submission: run a job on a remote computer
 - Information services: So I know which computer to use
 - File transfer: so large data files can be transferred
 - Replica management: so I can have multiple versions of a file “close” to the computers where I want to run jobs
- **Production grids are (currently) based on the Globus Toolkit release 2**
- **Globus Alliance: <http://www.globus.org/>**

- **GT2 Toolkit**
- **An example of the command line interface:**
 - Job submission – need to know name of a CE to use

```
globus-job-submit grid-data.rl.ac.uk/jobmanager-pbs /bin/hostname -f
```

```
https://grid-data.rl.ac.uk:64001/1415/1110129853/
```

```
globus-job-status https://grid-data.rl.ac.uk:64001/1415/1110129853/
```

```
DONE
```

```
globus-job-get-output https://grid-data.rl.ac.uk:64001/1415/1110129853/
```

```
grid-data12.rl.ac.uk
```

- Preface: EGEE and DILIGENT
- The basic tools and services
- **Building on the basics: from tools to a grid**
- More on selected services
- Emerging standards

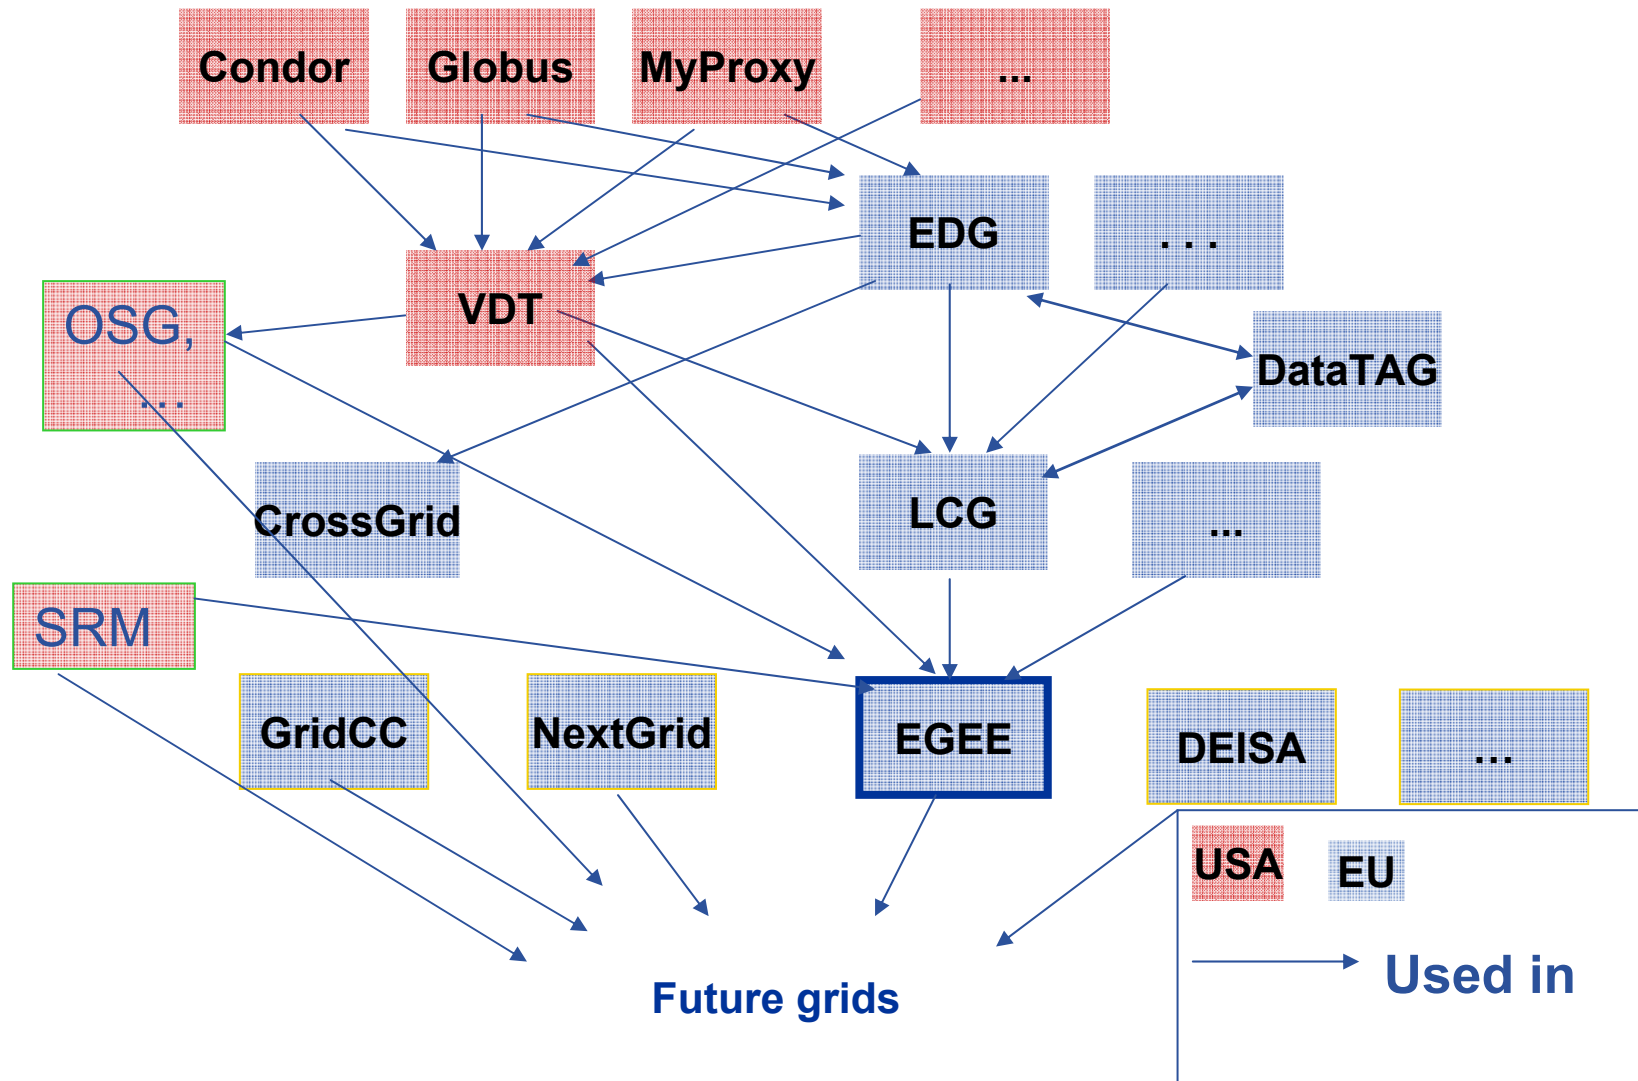


- **GT2: a toolkit – not a turnkey solution**
- **Need higher level tools**
 - E.g. so submit a job to “a grid” not a CE
- **...including services for**
 - **Logging** who’s done what, statistics about jobs,...
 - **Monitoring** whats happening on the grid
- **Illustrate this using LCG middleware (Large Hadron Collider Compute Grid)**
 - But first note the ecosystem of projects and middleware...
 - (with apologies to other projects I could have included)


Parts of the Grid “ecosystem”

2001

2004



- **“The Virtual Data Toolkit (VDT) is an ensemble of grid middleware that can be easily installed and configured. In our experience, installing grid software is challenging and time consuming. The goal of the VDT is to make it as easy as possible for users to deploy, maintain and use grid middleware.”**
<http://www.cs.wisc.edu/vdt/>



- Submit job to grid via the “resource broker”, in LCG...
`edg_job_submit my.jdl`

Application code can build on this using API's

Example JDL file

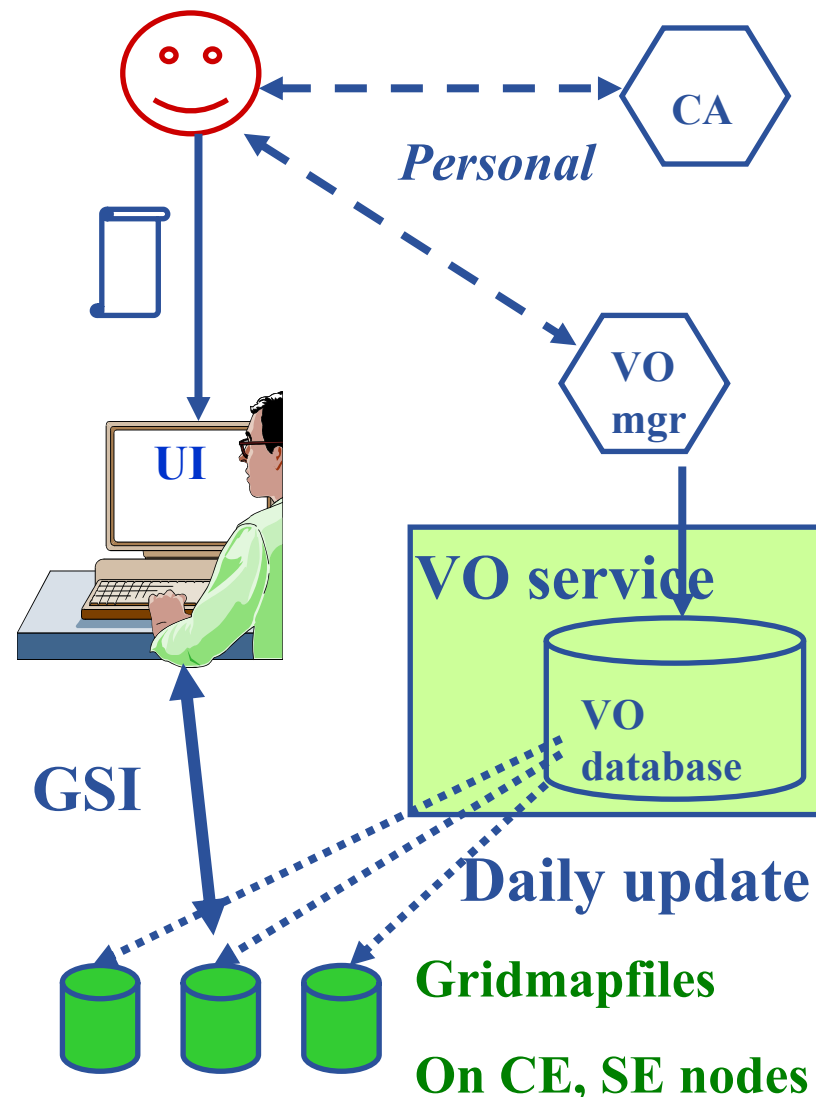
```
Executable = "gridTest";
StdError = "stderr.log";
StdOutput = "stdout.log";
InputSandbox = {"/home/joda/test/gridTest"};
OutputSandbox = {"stderr.log", "stdout.log"};
InputData = "lfn:testbed0-00019";
DataAccessProtocol = "gridftp";
Requirements = other.Architecture=="INTEL" && \
               other.OpSys=="LINUX" && other.FreeCpus >=4;
Rank = "other.GlueHostBenchmarkSF00";
```

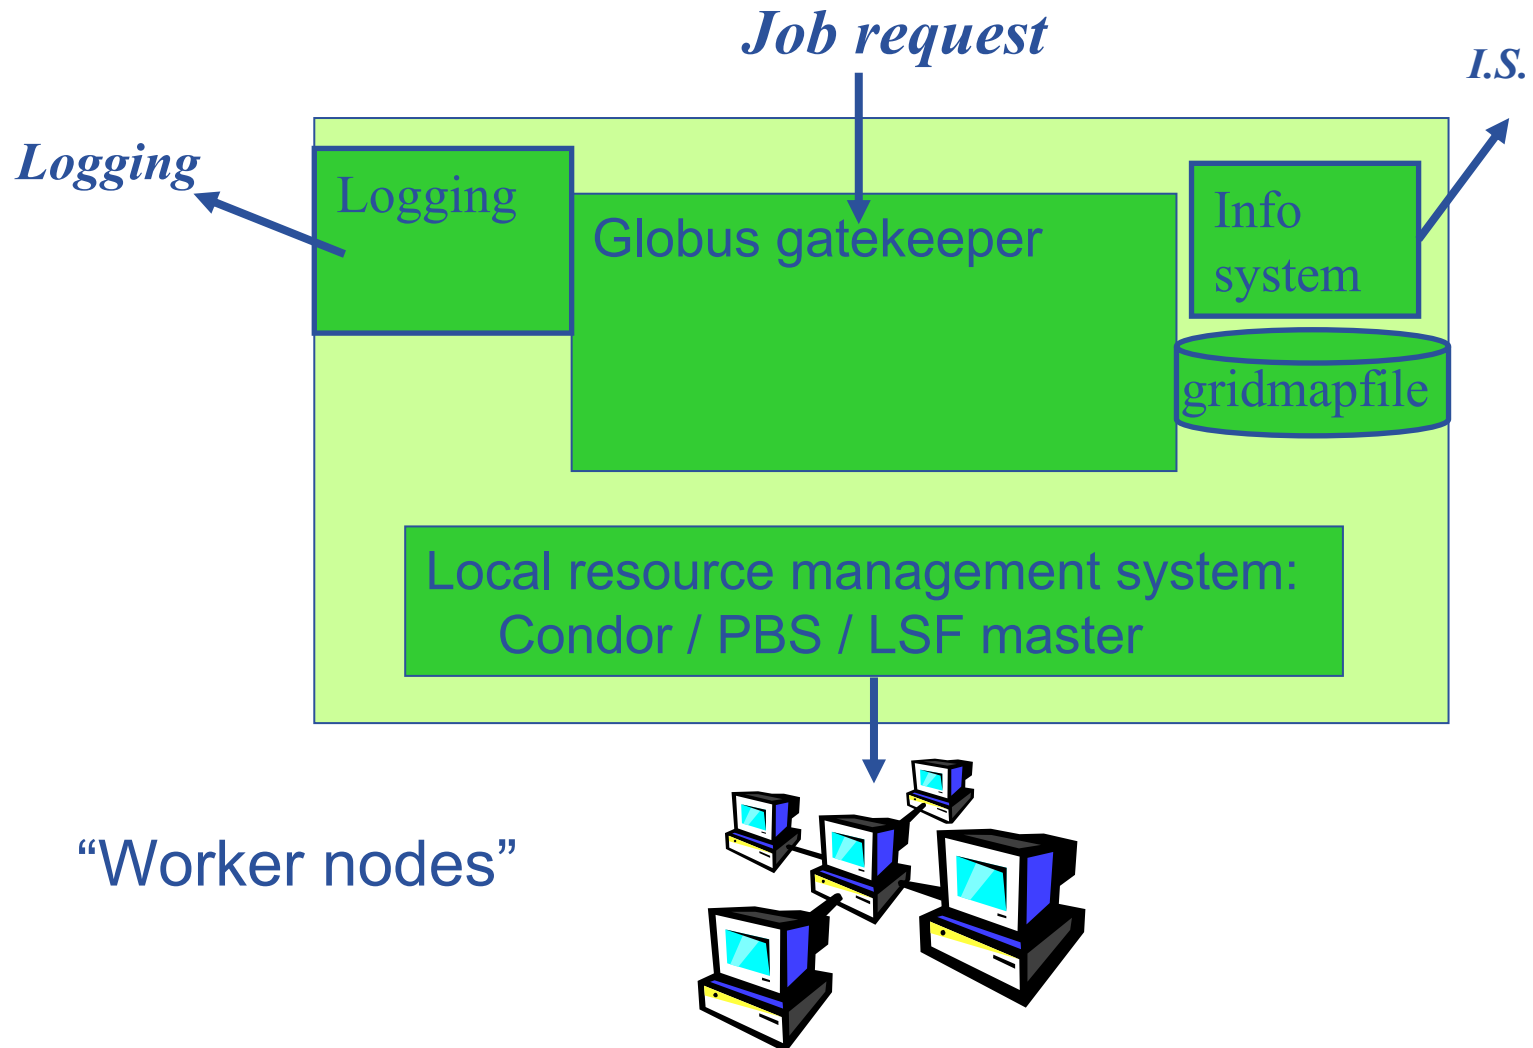

- **Authentication**

- User obtains certificate from CA
- Connects to UI by ssh
- Downloads certificate
- Invokes Proxy server
- **Single logon** – to UI - then **Grid Security Infrastructure identifies user to other machines**

- **Authorisation - currently**

- User joins Virtual Organisation
- VO negotiates access to Grid nodes and resources (CE, SE)
- Authorisation tested by CE, SE:
gridmapfile maps user to local account



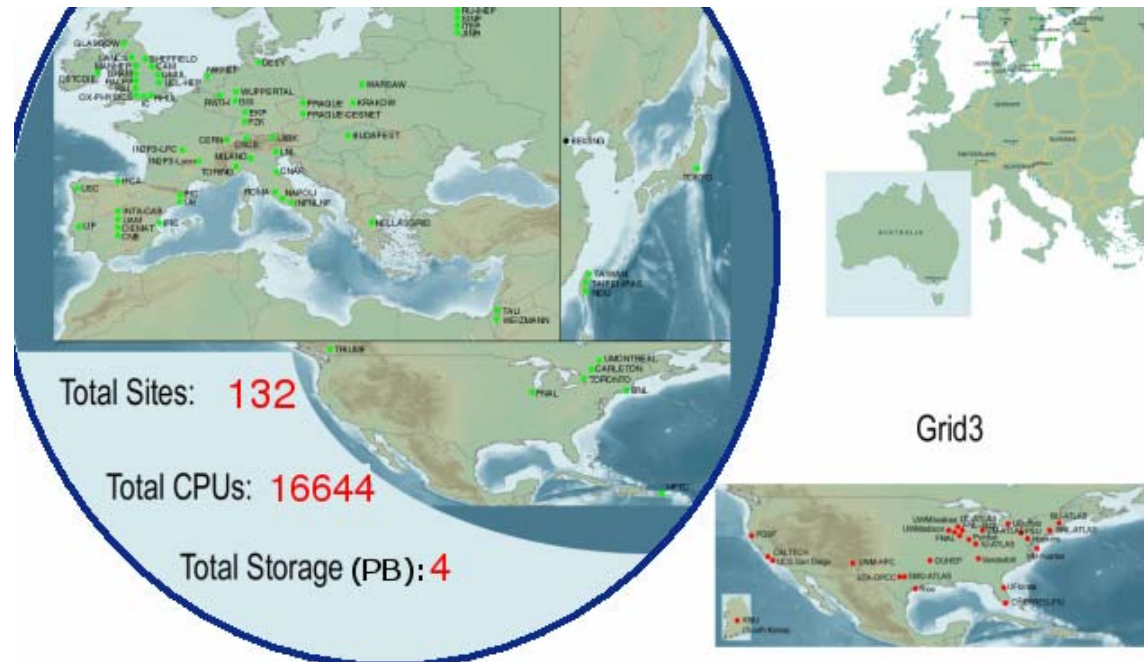


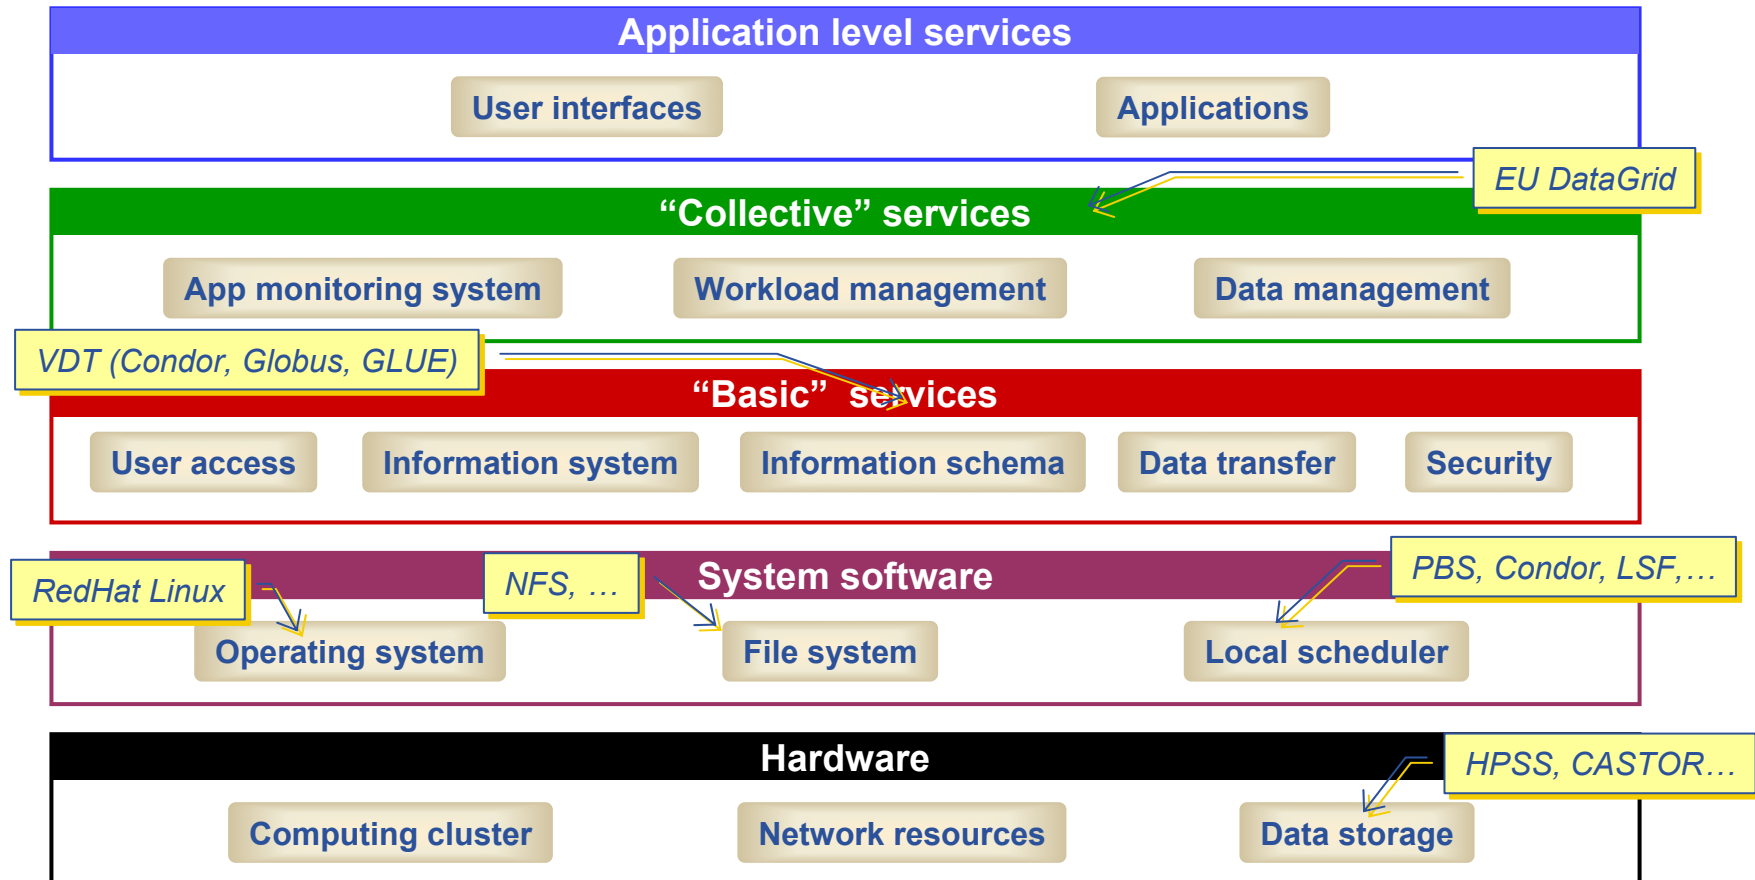
- **Real-time monitor**

- <http://www.hep.ph.ic.ac.uk/e-science/projects/demo/index.html>

- **Current status**

- <http://goc.grid-support.ac.uk/gridsite/monitoring/>





- Preface: EGEE and DILIGENT
- The basic tools and services
- Building on the basics: from tools to a grid
- **More on selected services:**
 - Data services
 - AA
 - Portals
- Emerging standards



Simple data files

- **Middleware supporting**
 - **Replica files**
 - *move data to computation*
 - Persistency
 - *Can add new storage technology*
 - Redundancy
 - “Close to Computing Element”
 - **Logical filenames**
 - **Catalogue**: maps logical name to physical storage device/file
 - **Virtual filesystems**
POSIX-like I/O
- **“SRM”**:
included in gLite

• Structured data

- RDBMS, XML databases
- Require ***extendable*** middleware tools to support
 - *Move computation near to data*
 - *easy access, controlled by AA*
 - *integration and federation*
- Hence **OGSA-DAI**
DAI: Data Access and Integration
- Now part of latest GT toolkit
- A few slides on OGSA-DAI follow
(from Malcolm Atkinson and from Mario Antonioletti)

- Motivation
 - **No business or research team is satisfied with a single resource**
- Data Curation Expertise Human Centred
- Integration Human centred
- Domain

Federation or Virtualisation preceding integration

Discipline insight needed here

Integration function

Revised request minutes later
Revised request after months of thought

- Sources inevitably heterogeneous
- Time-varying content, structure & policies
- Robust, stable steerable integration services
 - **Higher-level services over multiple resources**
 - **Fundamental requirements for (re)negotiation**

or kit of integration tools to be interwoven with an application?

Federation or

This is the motivation for & home of OGSA-DAI

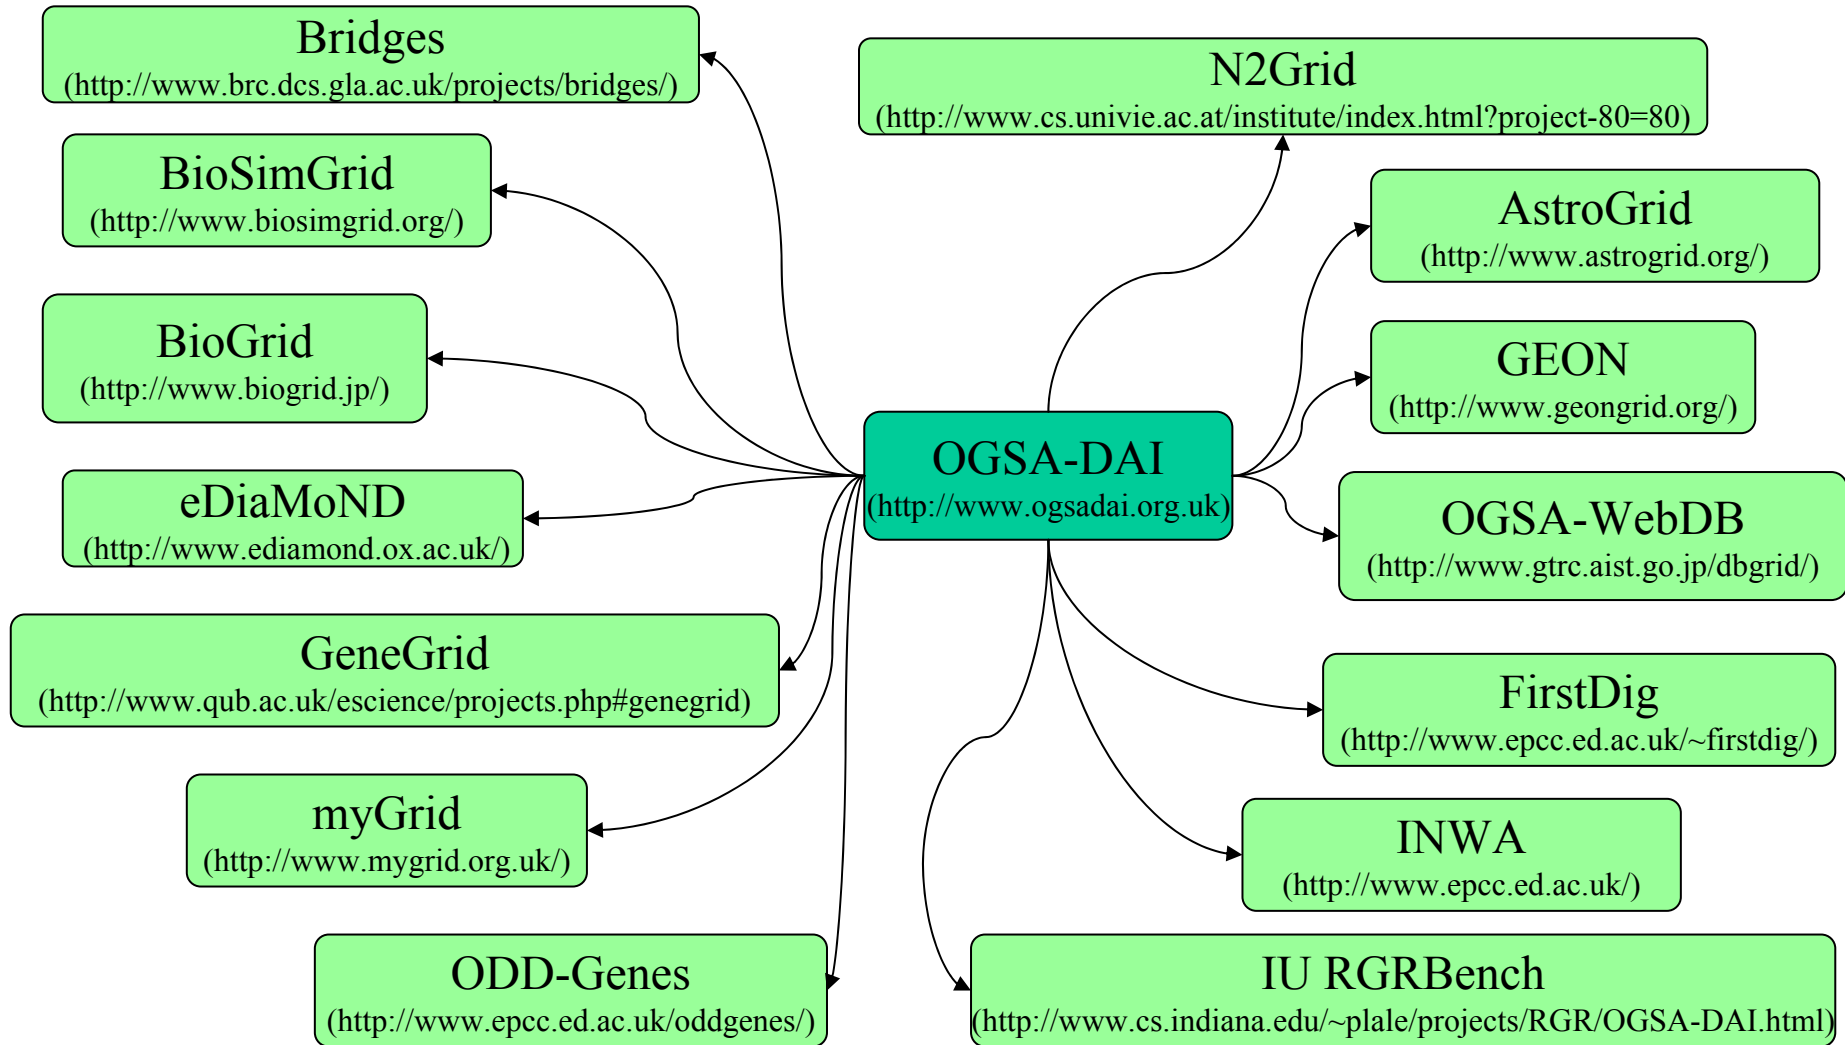
- Identify the recurrent requirements
- Provide one infrastructure that meets them
- Wide use enables a robust, reliable and supported set of facilities
- Steadily increase power of facilities
- Steadily raise the level of abstraction
- Standardise & achieve multi-national investment

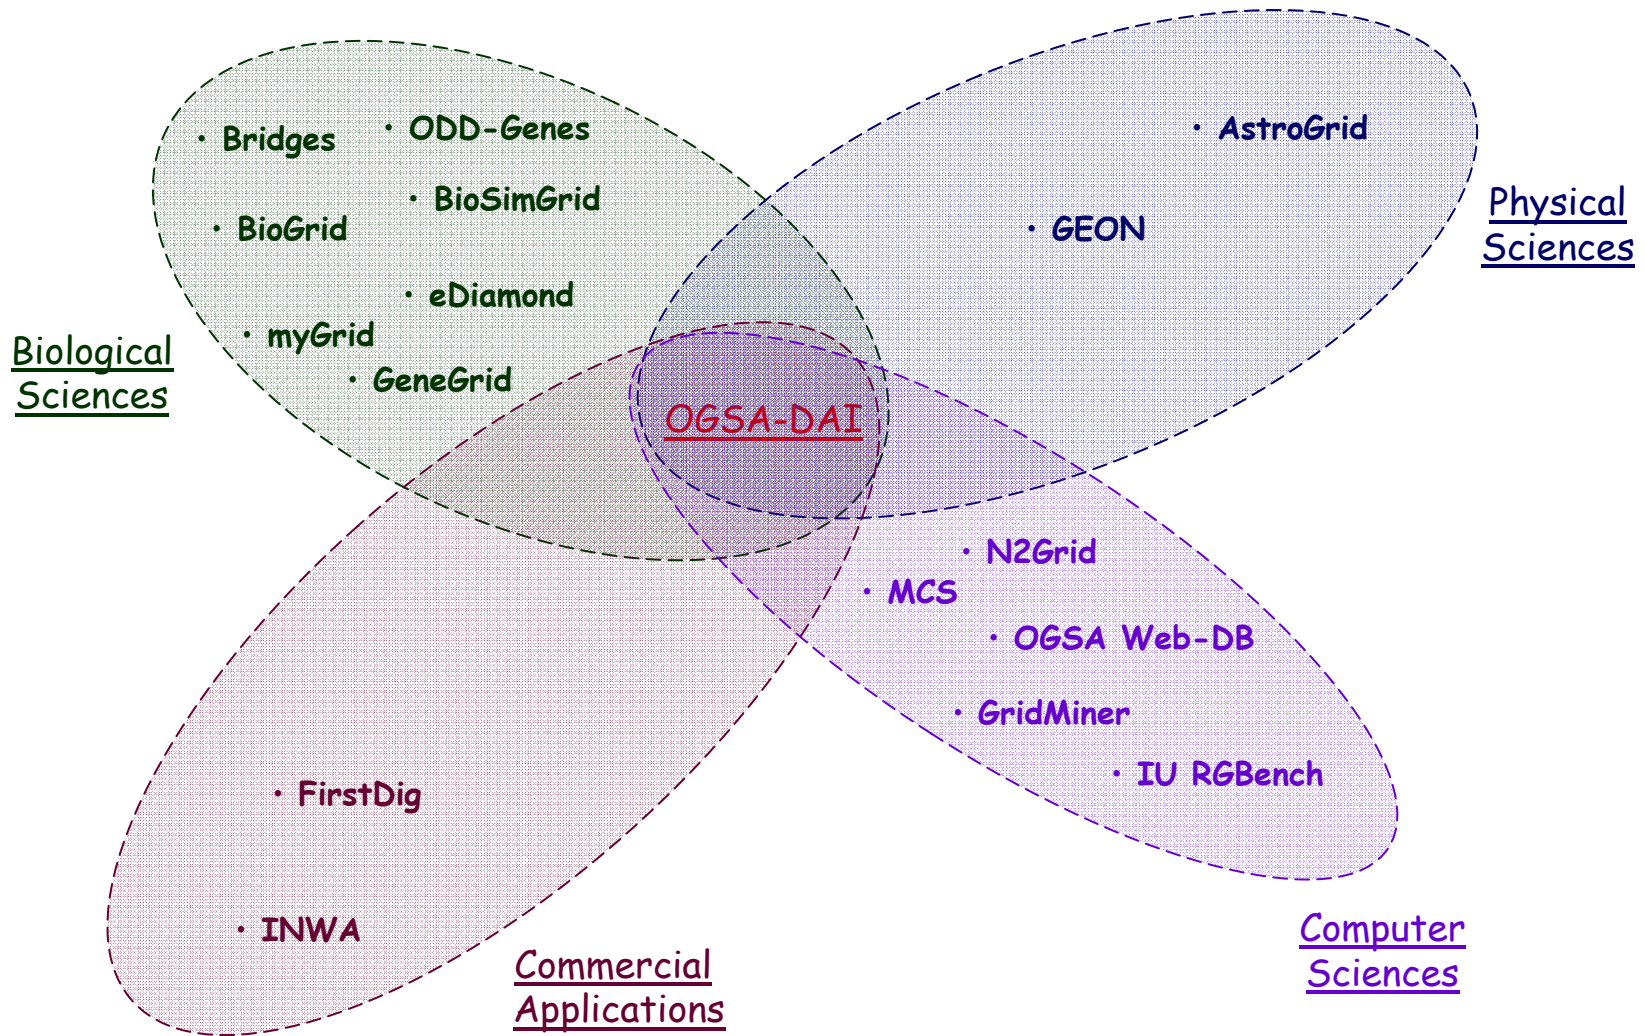
- Higher-level services over multiple resources
- Fundamental requirements for (re)negotiation

with an application?

- **OGSA-DAI is collaboration between:**
 - EPCC
 - IBM (+ Oracle in phase 1)
 - National e-Science Centre
 - Manchester University
 - Newcastle University
- **Project funding:**
 - OGSA-DAI, 2002-03,
 - £3.3 million from the UK Core e-Science fund
 - DAIT (DAI Two), 2003-06
 - £1.3 million from the UK e-Science Core Programme II
- **"OGSA-DAI" is a trade mark**
- **690 downloads between May 04 and 13 Dec 04**
 - **Actual user downloads not search engine crawlers**
 - **Does not include downloads as part of GT3.2 releases**
- **Total of 966 registered users**
- **<http://www.ogsadai.org.uk>**

**Funded by UK's
Department of Trade &
Industry + Engineering
& Physical Sciences
Research Council as
part of the e-Science
Core Programme**





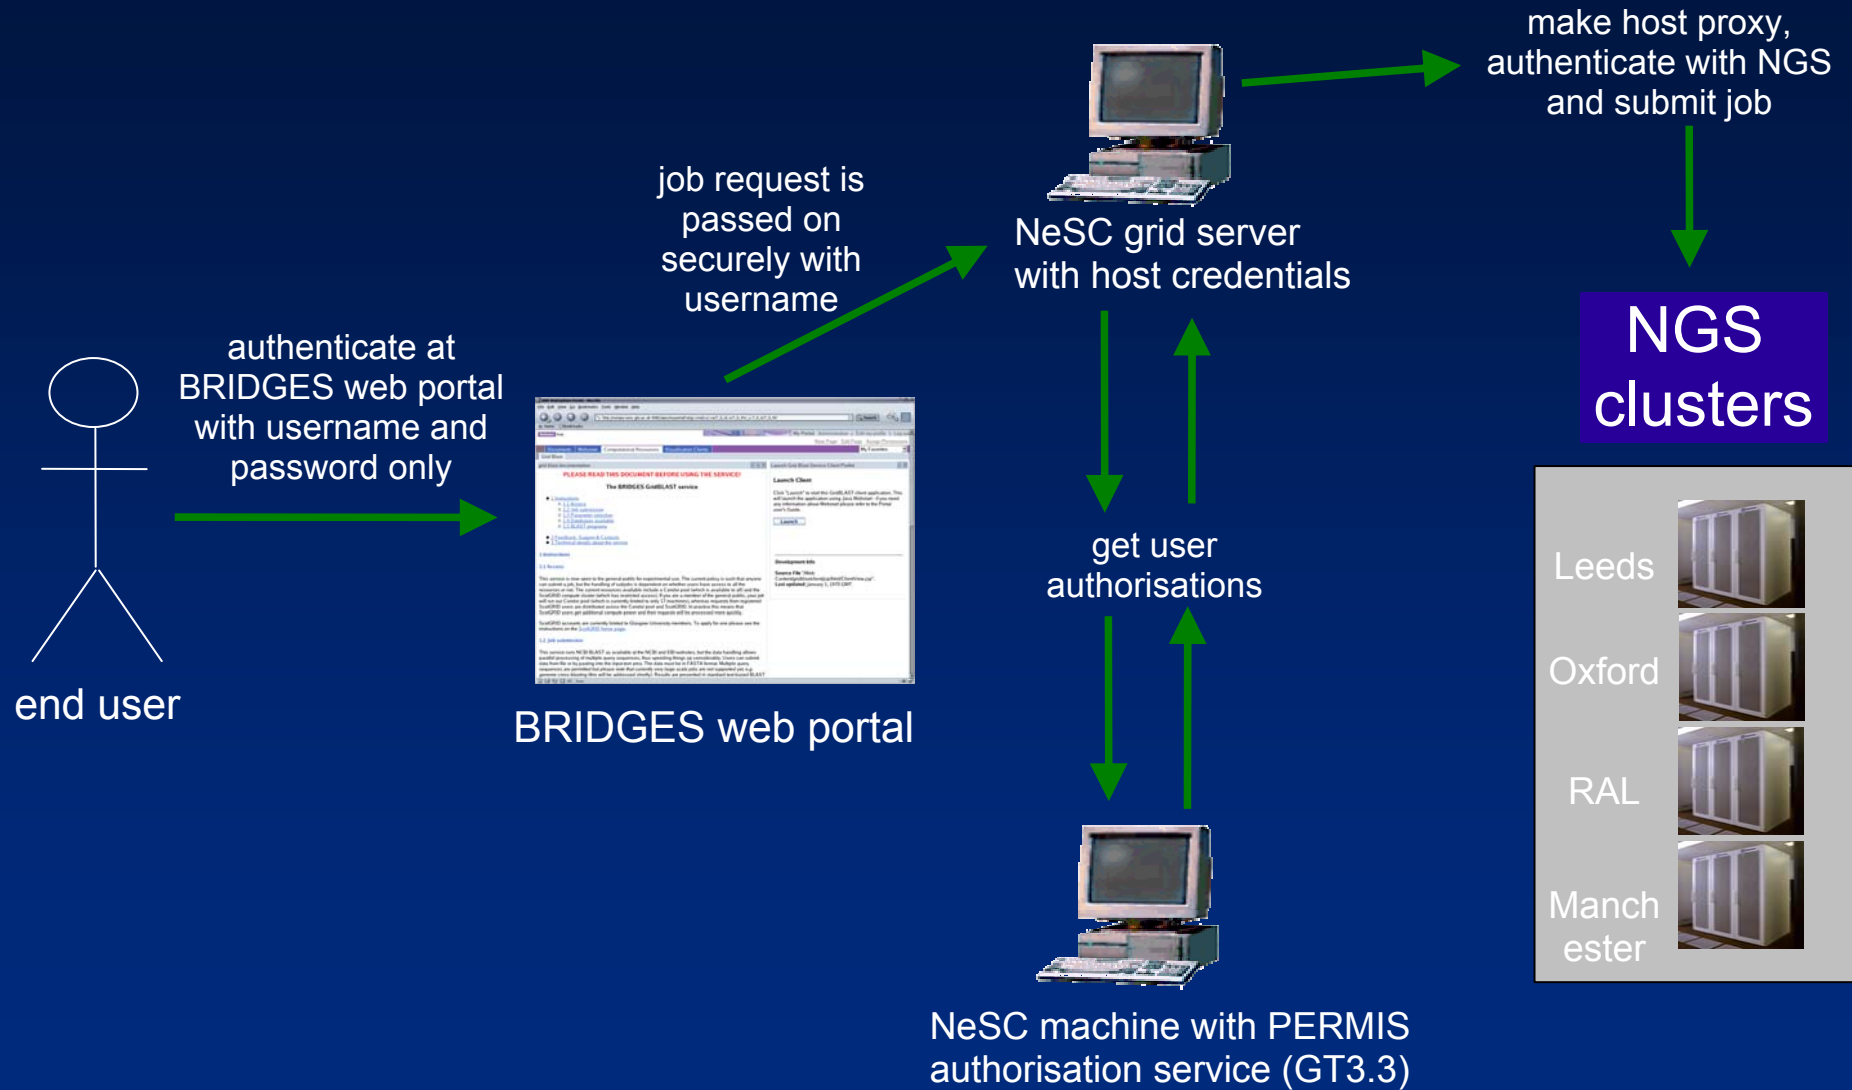
- Preface: EGEE and DILIGENT
- The basic tools and services
- Building on the basics: from tools to a grid
- **More on selected services:**
 - Data services
 - **Authentication, Authorisation**
 - Portals
- Emerging standards



- **Heterogeneity in Virtual Organisation membership**
 - In LCG-2, VOs are homogeneous
 - In EGEE (gLite) “VOMS” VO Management Service is used and will support roles
 - So different authorisations for different people in a VO
- **Users’ willingness to manage digital certificates.**
 - need avoid obstacles especially for early users
 - not all are as LINUX-literate as physicists
 - Effect: some communities resist certificates ..
 - Example from UK in next slide, using PERMIS in BRIDGES
 - Virtual Digital Libraries will entail building with existing DL communities
 - AA systems exist outside grids (ATHENS in the UK)
 - *Shibboleth is joining the party... see GridShib, ESP-GRID, DYVOSE*
- **Note that EGEE has a security activity – in next talk**
- **New developments needed for Digital Rights Management??**
 - E.g. as VDL federate with new data sources

- **PERMIS** <http://www.permis.org/>
- **BRIDGES**
<http://www.brc.dcs.gla.ac.uk/projects/bridges/>
- **DYVOSE** <http://labserv.nesc.gla.ac.uk/projects/dyvose/>
- **ESP-GRID**
http://e-science.ox.ac.uk/oesc/projects/index.xml.ID=body.1_div.20
- **GridShib** <http://grid.ncsa.uiuc.edu/GridShib/>
- references to ESP-GRID, gridShib – thanks to Mark Baker, Amy Apon, C. Ferner, J. Brown, “Emerging Grid Standards”, April 2005, <http://csdl.computer.org/comp/mags/co/2005/04/r4043abs.htm>

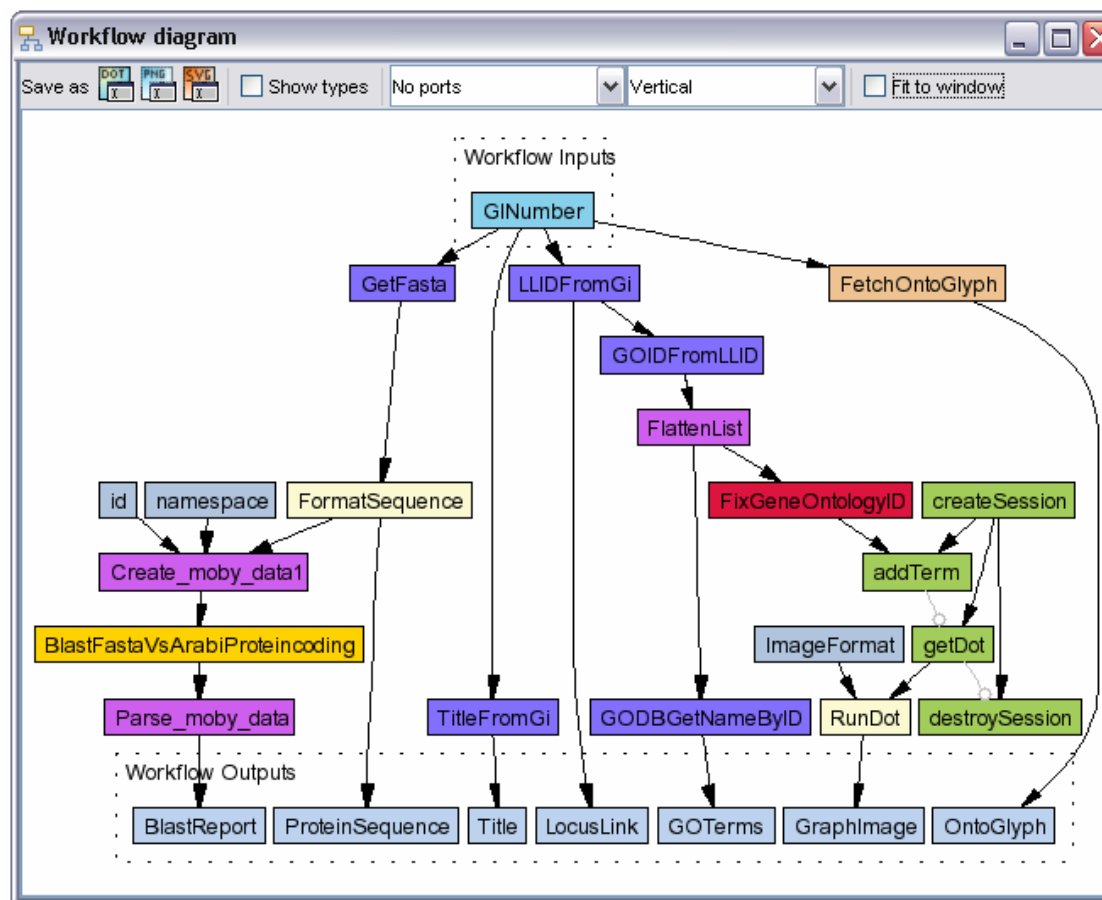
Security in BRIDGES – summary



- Preface: EGEE and DILIGENT
- The basic tools and services
- Building on the basics: from tools to a grid
- **More on selected services:**
 - Data services
 - AA
 - **Workflow**
- Emerging standards

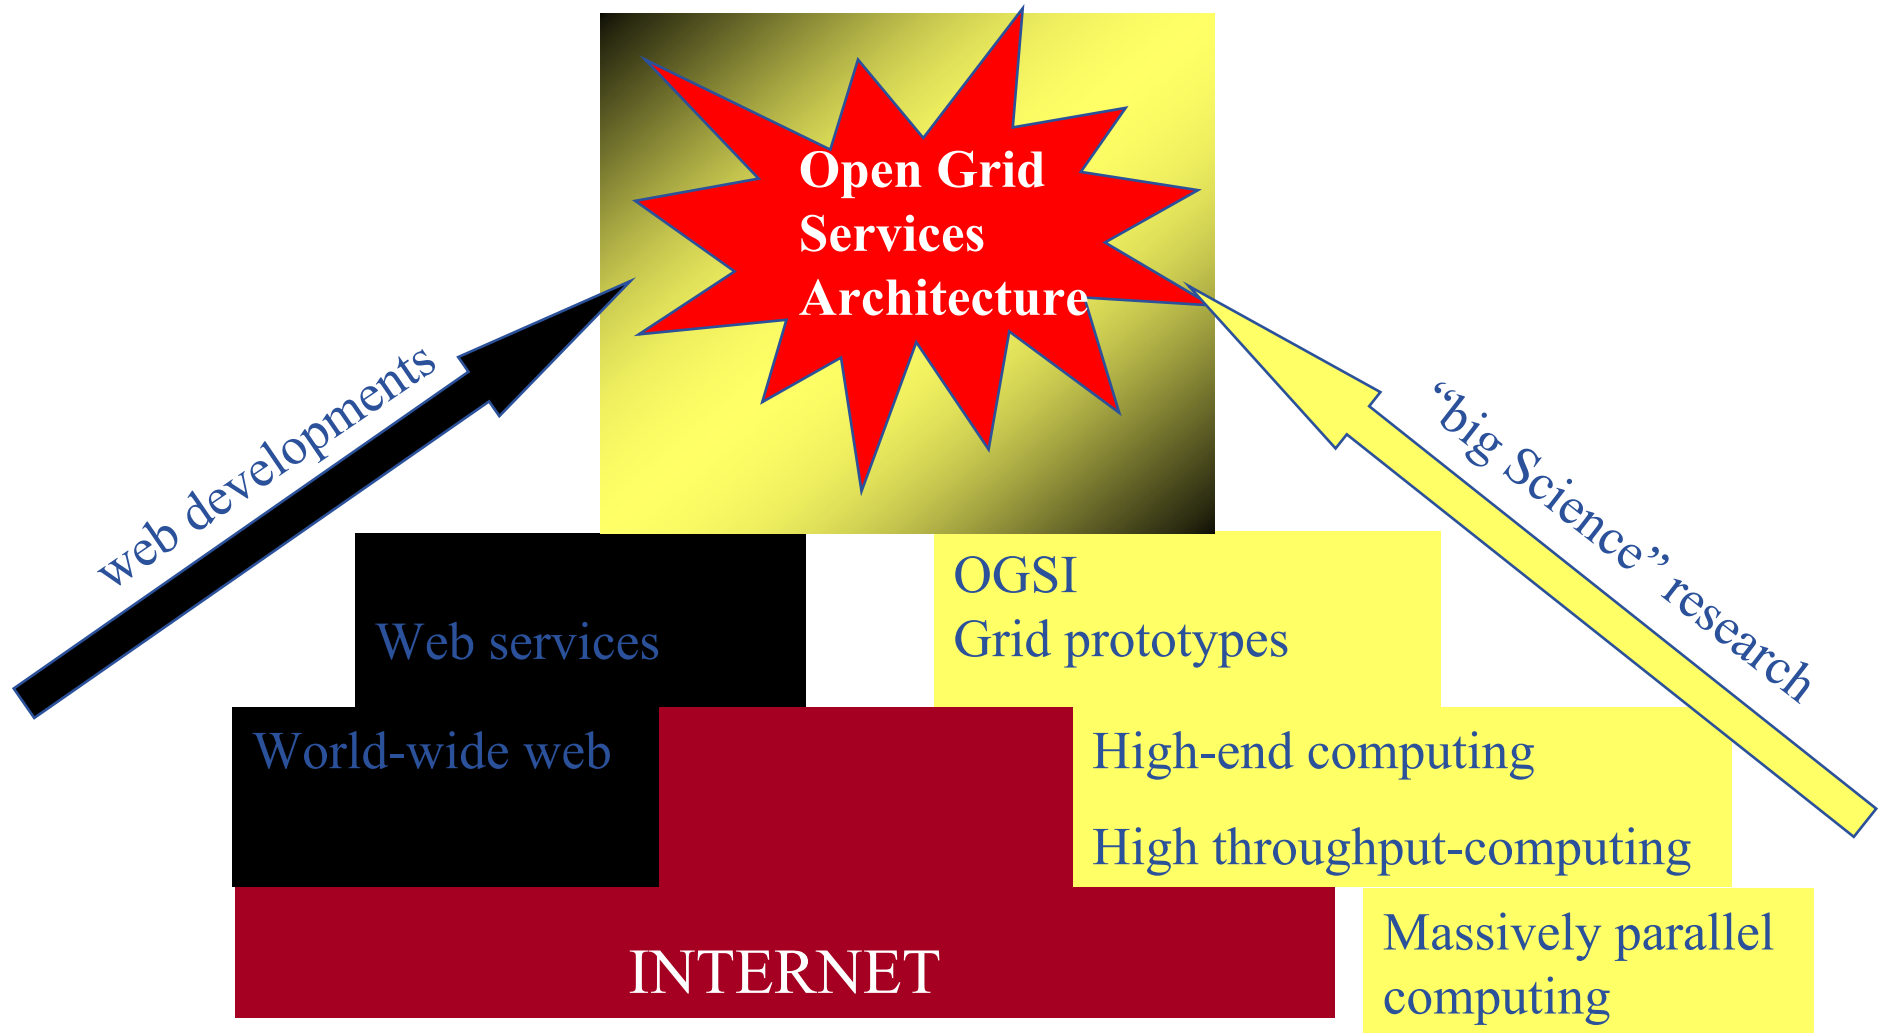


- Taverna in MyGrid <http://www.mygrid.org.uk/>
- “allows the e-Scientist to describe and enact their experimental processes in a structured, repeatable and verifiable way”
- GUI
- Workflow language
- enactment engine



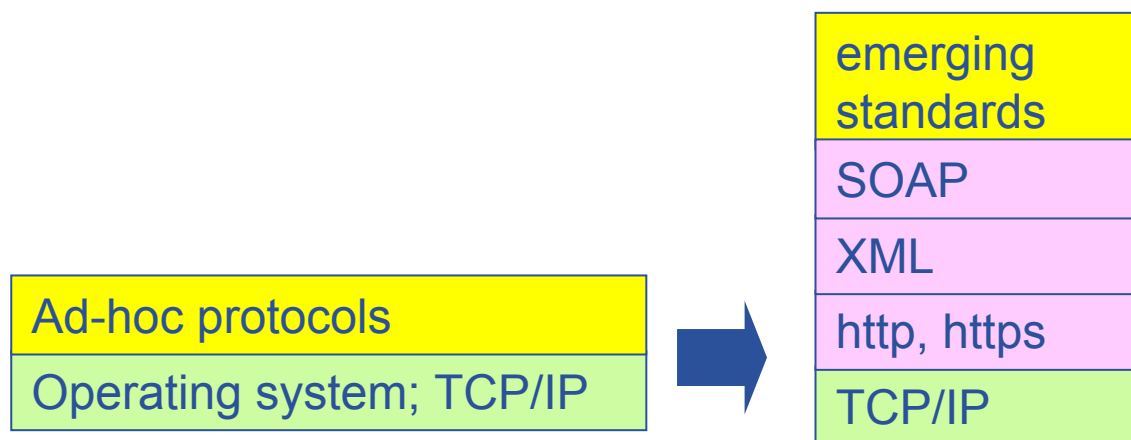
- Preface: EGEE and DILIGENT
- The basic tools and services
- Building on the basics: from tools to a grid
- More on selected services:
 - Data services
 - AA
 - Workflow
- **Emerging standards**





Why?

- Need service orientation for grids
- Opens integration of grids and WS worlds
- Leverage WS hosting environments



- **From the rich grid ecosystem ...**
- **Production services are**
 - **Built on tools and services for**
 - Authorisation and authentication
 - Job submission (direct to a Computing Element)
 - File replication
 - **...with higher level services**
 - Job submission to “a grid” (via resource broker)
 - Monitoring
 - Logging
 - **..and upon these, toolkits and services for creating new applications**
 - Workflow
 - Portals (demo this afternoon)
 - ...
- **Authorisation and authentication underpin it all**
 - resource-sharing across organisations, without centralised control

- Global Grid Forum <http://www.gridforum.org/>
- Globus Alliance <http://www.globus.org/>
- SRM see http://www.nesc.ac.uk/esi/visitors/reports/115_report.pdf
and URLs in this report
- OGSA-DAI <http://www.ogsadai.org.uk>
- Condor <http://www.cs.wisc.edu/condor/>
- VDT <http://www.cs.wisc.edu/vdt/>
- Open Science Grid <http://www.opensciencegrid.org/>
- Grid Center <http://www.grid-center.org/>
- LCG <http://lcg.web.cern.ch/LCG/>