

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

## Grid Technology

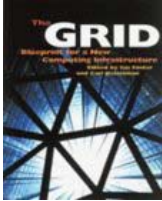
Dr Steven Newhouse, Technical Director  
London e-Science Centre

Department of Computing, Imperial College London

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

## What is the Grid?

“ Grid computing [is] distinguished from conventional distributed computing by its focus on large-scale resource sharing, innovative applications, and, in some cases, high-performance orientation...we review the "Grid problem", which we define as **flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources** - what we refer to as virtual organizations.”



*From "The Anatomy of the Grid: Enabling Scalable Virtual Organizations" by Foster, Kesselman and Tuecke*

2

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

## Why Grids?

- Large-scale science and engineering are done through the **interaction** of people, heterogeneous computing resources, information systems, and instruments, all of which are **geographically and organizationally dispersed**.
- The overall motivation for “Grids” is to facilitate the **routine interactions** of these resources in order to support large-scale science and engineering.

From Bill Johnston 27 July 01

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

## Why Now?

- Moore’s law improvements in computing produce highly functional endsystems
- The Internet and burgeoning wired and wireless provide universal connectivity
- Changing modes of working and problem solving emphasize teamwork, computation
- Network exponentials produce dramatic changes in geometry and geography
  - 9-month doubling: double Moore’s law!
  - 1986-2001: x340,000; 2001-2010: x4000?

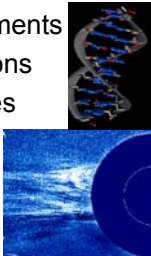
4

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

## What is e-Science?

- Applied Scientists are becoming e-scientists
- Dependent on remote electronic services
- Utilising scarce expensive instruments
- Involvement in global collaborations
- Interaction through mobile devices

**Urgent need for an integrated environment to support this activity**




5

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

## CERN’s Large Hadron Collider

1800 Physicists, 150 Institutes, 32 Countries



100 PB of data by 2010; 50,000 CPUs?

[www.griphyn.org](http://www.griphyn.org)  
 [www.ppdg.net](http://www.ppdg.net)  
 [www.eu-datagrid.org](http://www.eu-datagrid.org)

6

## UK e-Science Programme

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Multi-disciplinary projects
  - Applied Scientists
  - Computer Scientists
- £120M over three years
- Started 2001

7

## London e-Science Centre

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

### 'Enabling the e-Scientist'

- Established applied multi-disciplinary research
- Industrial Collaborations:
  - Sun Centre of Excellence in e-Science
  - Intel Virtual European Centre of Grid Computing
- Cross-campus collaborations:
  - Bioinformatics
  - High Energy Physics
  - Computational Engineering
- Specialisation: Next Generation Grid Middleware
- <http://www.lesc.imperial.ac.uk/>

8

## Classical Grids

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Classical Grids emphasise the sharing of physical resources.
- Existing Grid middleware (e.g. Globus, Condor, Unicore) allows resource discovery, resource allocation, data movement & certification
- Focus on low-level services not the application

9

## The Globus Toolkit™

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Widely adopted Grid Middleware
- Result of 5+ years (100+ man years) work
- Provides protocols to support three services:
  - Job Submission and Monitoring (~HTTPS)
  - Monitoring and Discovery Service (~LDAP)
  - Data Transfer (~FTP)
- Infrastructure to federate distributed resources
  - Support for access control
  - Wide area security mechanisms (PKI)

10

## The Globus Toolkit in One Slide

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

Grid protocols (GSI, GRAM, ...) enable resource sharing within virtual orgs; toolkit provides reference implementation (■ = Globus Toolkit services)

Protocols (and APIs) enable other tools and services for membership, discovery, data mgmt, workflow, ...

11

## GridFTP

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- GSI authentication on data & control channels
- Parallel streams (client to 'clustered' server)
- TCP buffer tuning for maximum performance
- Third party file transfer
- SC2001: 231GB/hour (1.55 GB/s)
  - 8 Striped servers with 4 TCP connections

12

## Personal Condor

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

600 Condor jobs  
personal Condor

- One user & one machine (no root access)
- Provides a persistent execution environment that will run jobs when your policy allows them to run
- Need jobs to run without any user interaction

13

## Layout of the Condor Pool

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

..... = Process Spawned  
→ = ClassAd Communication Pathway

Central Manager  
Cluster Node  
Cluster Node  
Desktop

600 Condor jobs  
Condor Pool

14

## Condor pool within DoC

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

### Imperial College - Department of Computing Condor Pool Machine Statistics for Week

From Thu Nov 21 00:20:23 GMT 2002 to Wed Nov 27 23:49:28 GMT 2002

285.0  
213.7  
142.5  
71.25  
0.0

Total Condor  
Total Idle  
Total Owner

Thu 21 Fri 22 Sat 23 Sun 24 Mon 25 Tue 26 Wed 27

Configure... Zoom In Zoom Out Reset About

[High Filter: The Y-axis is number of machines, the X-axis is time. When graph buffer updating, press "Configure" to view different resolutions or time slices. Also, you can use the mouse to draw a rectangle on the graph and then press "Zoom In" to zoom in to subdivide the data, also Configure or return data zooming. Highlighted down up on graph background in grey.]

15

## Resource Sharing through Condor

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

LeSC Firewall

Submit Machine  
Schedd  
Central Manager  
Collector  
Negotiator  
Pool-Foo  
Collector  
Negotiator  
Pool-Bar  
Collector  
Negotiator

saturn (Solaris E6800) viking-i00 (Linux Xeon) mir (Linux x86)

16

## Mathematicians Solve NUG30

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Looking for the solution to the NUG30 quadratic assignment problem
- An informal collaboration of mathematicians and computer scientists
- Condor-G delivered 3.46E8 CPU seconds in 7 days (peak 1009 processors) in U.S. and Italy (8 sites)

14, 5, 28, 24, 1, 3, 16, 15,  
10, 9, 21, 2, 4, 29, 25, 22,  
13, 26, 17, 30, 6, 20, 19,  
8, 18, 7, 27, 12, 11, 23

MetaNEOS: Argonne, Iowa, Northwestern, Wisconsin

17

## Global Grid Forum - www.gridforum.org

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

Global Grid Forum - Microsoft Internet Explorer

GGF7  
March 4-7 2003  
Hotel Intercontinental Keio Plaza  
Tokyo, JAPAN

18

**Convergence between e-Science and e-Business**

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- A biochemist exploits 10,000 computers to screen 100,000 compounds in an hour;
- A biologist combines a range of diverse and distributed resources (databases, tools, instruments) to answer complex questions;
- 1,000 physicists worldwide pool resources for petaop analyses of petabytes of data
- Civil engineers collaborate to design, execute, & analyze shake table experiments

From Steve Tuecke 12 Oct. 01

**Convergence between e-Science and e-Business (contd.)**

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Climate scientists visualize, annotate, & analyze terabyte simulation datasets
- An emergency response team couples real time data, weather model, population data
- A multidisciplinary analysis in aerospace couples code and data in four companies
- A home user invokes architectural design functions at an application service provider

From Steve Tuecke 12 Oct. 01

**Today's Grid**

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- A Single System Interface
- Transparent wide-area access to large data banks
- Transparent wide-area access to applications on heterogeneous platforms
- Transparent wide-area access to processing resources
- Security, certification, single sign-on authentication
  - Grid Security Infrastructure,
- Data access, Transfer & Replication
  - GridFTP, Giggie
- Computational resource discovery, allocation and process creation
  - GRAM, Unicore, Condor-G

21

**Grid Evolution**

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- 1<sup>st</sup> Generation Grid
  - Computationally intensive, file access/transfer
  - Bag of various heterogeneous protocols & toolkits
  - Recognises internet, Ignores Web
  - Academic teams
- 2<sup>nd</sup> Generation Grid We are here!
  - Data intensive -> knowledge intensive
  - Services-based architecture
  - Recognises Web and Web services
  - Global Grid Forum
  - Industry participation

22

**Convergence on Service Oriented Architectures**

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

```

    graph TD
      SR[Service Registry] -- "Publish Service Description" --> SP[Service Provider]
      SR -- "Find a service" --> SRQ[Service Requestor]
      SRQ -- "Bind to the Service" --> SP
  
```

- Globus: Grid Index Information Service (LDAP)
- Jini: Look-up Server
- CORBA
- Web Services
  - Open Grid Services Architecture in the Grid community



23

**The Future... Web Services**

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Increasingly popular standards-based framework for accessing network applications
  - W3C standardization; Microsoft, IBM, Sun, others
- WSDL: Web Services Description Language
  - Interface Definition Language for Web services
- SOAP: Simple Object Access Protocol
  - XML-based RPC protocol; common WSDL target
- WS-Inspection
  - Conventions for locating service descriptions
- UDDI: Universal Description, Discovery, & Integration
  - Directory for Web services



24

 **Supporting 'Grid Services'** 

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Dynamic Services
  - Created & destroyed on user demand
- Stateful Services
  - Expose internal service state
- Flexible Service Discovery Infrastructures
  - Federate local service registries
  - Build Virtual Organisations



25

 **Open Grid Services Architecture (OGSA)** 

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Utilise standard Web services infrastructures
- Building on current Globus Toolkit:
  - **Grid service**: semantics for service interactions
  - Management of transient instances (& state)
  - **Factory, Registry, Discovery**, other services
  - Reliable and secure transport
- Multiple hosting targets: J2EE, .NET, "C", ...
- Service orientated architecture enables resource virtualisation
- Delivery via open source Globus Toolkit 3.0
  - ⇒ Leverage GT experience, code, mindshare



26

 **Open Grid Services Infrastructure (OGSI)** 

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Global Grid Forum
  - Grid Standards Organisation
  - Modelled on the IETF
  - Mixture of Research & Working Groups
  - Further details [www.gridforum.org](http://www.gridforum.org)
- Grid Services Specification (v1.0),
  - Produced by the OGSI Working Group
  - Entered 60 day comment period March 2003



27

 **Grid Service Specification** 

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Define minimal set of ports & behaviours:
  - GridService
  - HandlerResolver
  - Notification Source, Sink & Subscription ports
  - Factory
  - Registration
- Other issues to be defined elsewhere
  - Exposing an EJB as a Grid Service
  - Security...



28

 **Service Data Elements (SDE)** 

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Describes the public internal service state
- All SDE's have attributes to define:
  - goodFrom
  - goodUntil
  - availableUntil
  - User definable extensibility elements
- Examples
  - Current CPU load, available memory, etc

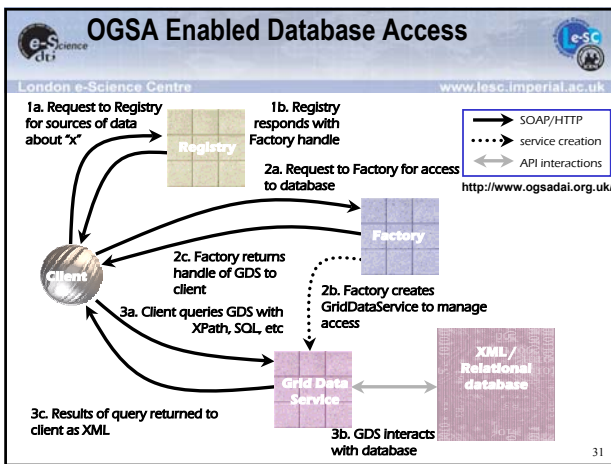
29

 **Factory Interface within the GSS** 

London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)

- Creates a new Grid service instance
  - Reliable creation (once-and-only-once)
- Service-specific creation parameters:
  - Request 16 processors for 4 hours
  - Temporary disk space of 500Mb
- Returns a **Grid Service Handle (GSH)**
  - A globally unique URL
  - Uniquely identifies the instance for all time

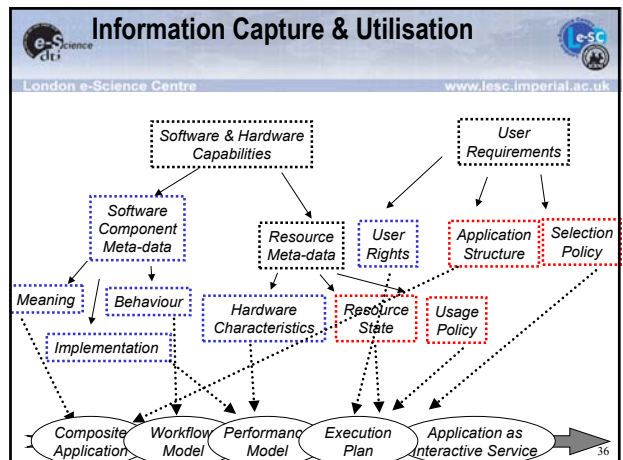
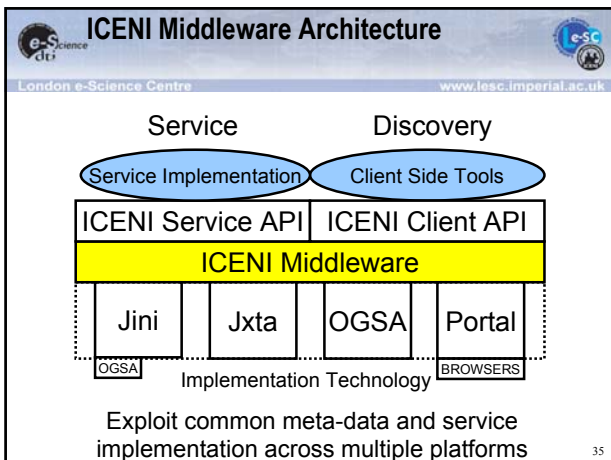
30

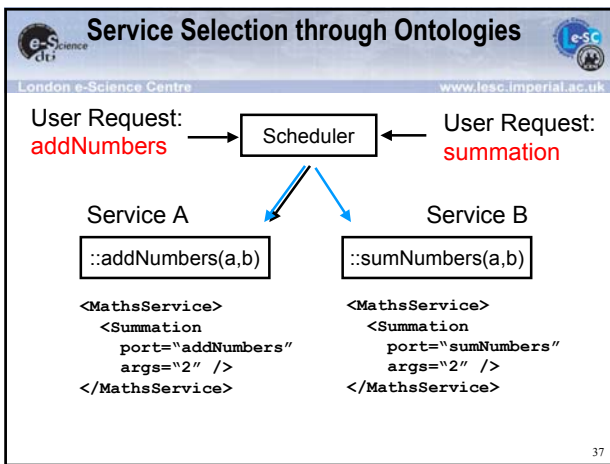


- ### Future Grid Challenges
- London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)
- Need 'power stations' on the Grid
    - Buy (obtain) resources as required
  - Need to understand how applications behave
    - Balance out data transfer vs. compute shipping
  - Need to scalable wide-area service discovery
    - Peer to Peer or centralised servers
    - Meta-data to describe Grid Services
  - Need to exploit distributed services
    - Grid Service Orchestration
    - Optimise service selection and recover from failure
- 32

- ### Classifying Services
- London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)
- Web Service model:
    - Service meta-data in UDDI
    - DAML-S
  - Grid Services (Current)
    - 'Well known' naming of services & ports
  - Grid Services (Future)
    - Domain Specific Service Ontologies
      - MONET, MyGrid, ...
    - Generic Service Ontologies
- 33

- ### ICENI – Imperial College e-Science Networked Infrastructure
- London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)
- The IcenI, under Queen Boudicca, united the tribes of South-East England in a revolt against the occupying Roman forces in AD60.*
- 
- Developed by LeSC Grid Middleware Group
  - Collect and provide relevant Grid meta-data
  - Use to define and develop higher-level services
  - Interaction with other frameworks,
    - e.g. Web Services, Jxta etc.
- 34





- ## Benefits
- London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)
- Decouple service naming and operation
  - Standard tools to classify services
  - Ability to reason about services
  - Providing a persistent service infrastructure
- ## Autonomic Computing
- 38

- ## Market for Computational Services
- London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)
- UK Core e-Science Programme project
  - Interfaces & protocols to trade Grid Services
  - Funded by Department of Trade & Industry
  - Collaborators
    - London e-Science Centre
    - e-Science North West
    - Southampton e-Science Centre
    - UK Grid Support Centre
  - Commercial
    - SME's: Software & service provision
    - Multi-national's: Hardware & Software vendors
- 
- 39

- ## New Economic SDEs
- London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)
- Trust, Reputation & Reliability (0+)
    - Can I rely on this service? Will it deliver?
    - Has this service been audited for performance & integrity?
    - Allow SDE to updated by signed statements?
  - Compensation & Liability (0+)
    - Who can I complain to and can I prosecute?
    - How could the contract be broken?
  - Grid Banking Services (1+)
  - Resource Usage Service (1)
    - Declare which RUS is used by the service
- 40

- ## Grid Economic Service Architecture
- London e-Science Centre [www.lesc.imperial.ac.uk](http://www.lesc.imperial.ac.uk)
- GESA enabled client 'obtains' a GSH
    - Grid Service Handle (GSH) provides unique service ID
    - Represents a service & cost for use
    - Changes in price & status through new service & GSH
  - GSH passed to an existing client
    - Retains existing client interface for economic services
  - Key Issues:
    - Information needed to make the selection (SDE)
    - Two (Multi?) stage commit to support pricing
    - Enable wide ranging exploration of grid economic models
- 41

