

Building Computational Communities From Federated Resources

Nathalie Furmento, **Steven Newhouse**,
John Darlington

Grid Middleware Group
London - Science Centre
Imperial College, London



Euro-Par2001, Manchester, 28th-31st August 2001

1

Outline

- Computational Grids and e-Science
- Using Java & Jini as a Grid Middleware
- Exploiting Grid Meta-Data
- Current Status
- Summary & Further Work

2

Computational Grids & e-Science

- Connect networked resources (the Grid)
 - Linked instruments/storage/computation
 - Coupled computing (meta-computing - high performance)
- Optimise use of resources
 - Wider access
 - Better utilisation of resources (throughput)
- High-level use (e-Science)
 - Ease of use for domain specialists
 - Support scientific collaboration and knowledge management (connect & build communities)

3

Building Computational Communities

- By federating resources from real organisations (to form *virtual organisations*)
- Share, combine and use computational resources (software and hardware) in an open, transparent and optimal manner
- Need information about:
 - The resources
 - The users
 - The applications (*See presentation 4.00, Thursday, Room D*)

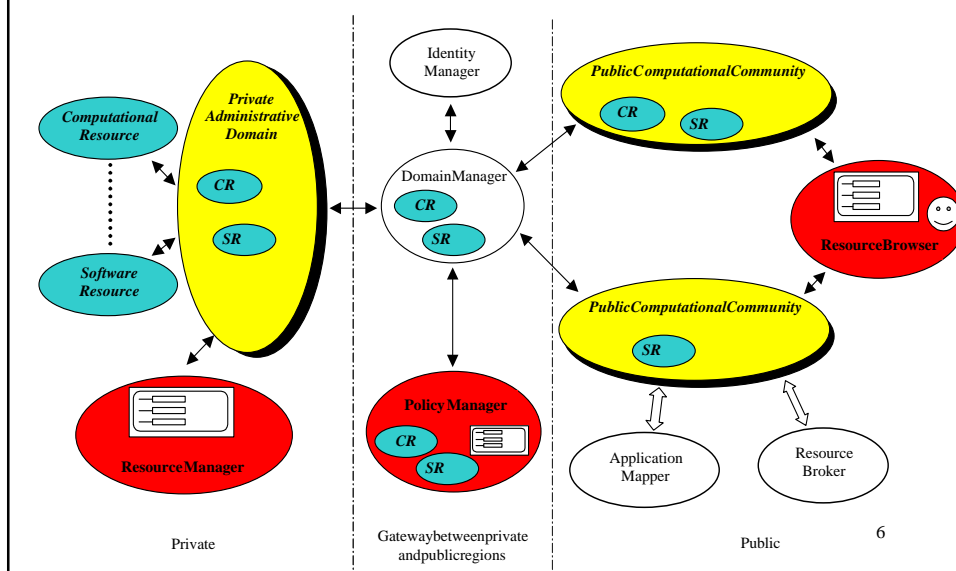
4

Federating Distributed Resources

- Real Organisations
 - Need to express ownership and retain control
 - Grid resources are 'not' free – someone must pay!
 - Manage the resources as a single unit
- Virtual Organisations
 - Hide complexity from the user
 - Optimise resource utilisation for all jobs and users
 - Easy to contribute (and withdraw) resources

5


Java/Jini Grid Architecture



6

Managing Internal Resources

- The resources advertise their capabilities to a private domain
- Three types of resources:
 - Computational resources
 - Storage resources
 - Software resources
- Annotate using XML



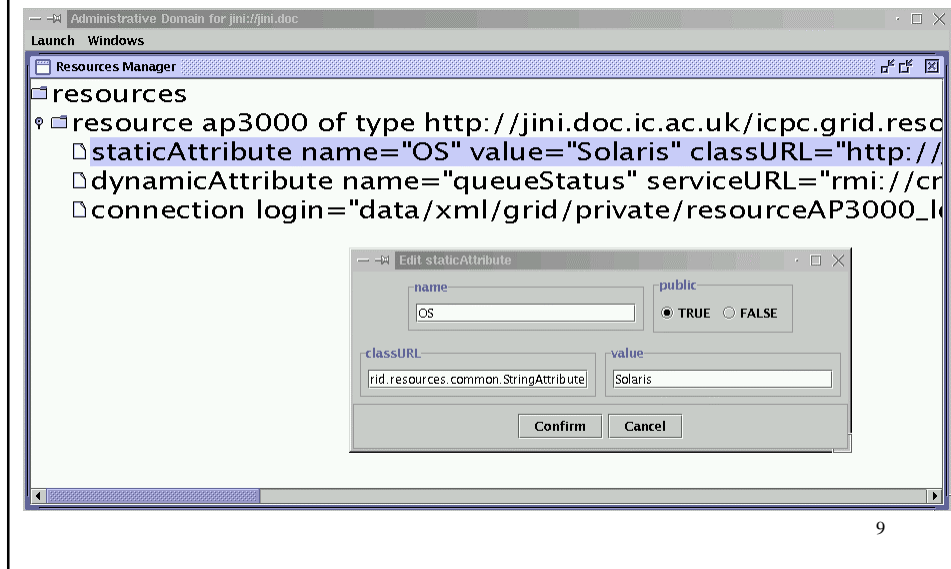
7

A Computational Resource: the AP3000

```
<resource name="ap3000" publish="jini://jini.doc/"
  type="...ComputationalResource">
  <staticAttribute name="OS" value="Solaris"
    public="TRUE" classURL="...StringAttribute"/>
  <dynamicAttribute name="queueStatus"
    serviceURL="rmi:...getAP3000QSStatus"
    classURL="...AP3000" public="TRUE"
    interval="5" property="...queueStatus.prp"/>
  <connection login=" ...AP3000_login.xml.enc"/>
</resource>
```

8

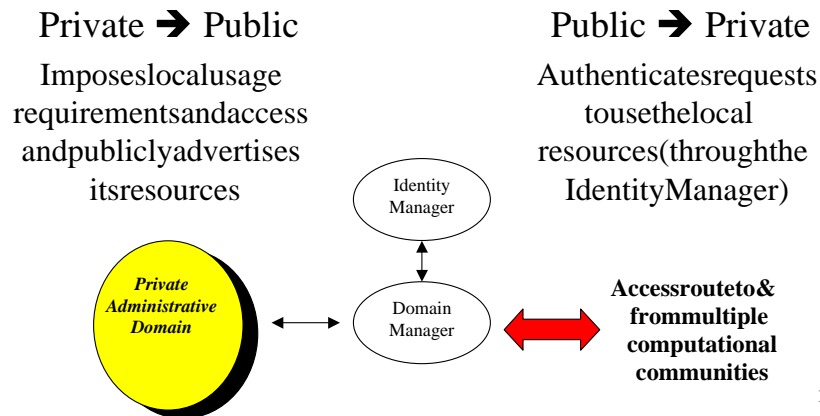
TheResourceManagerGUI



9

TheDomainManager

Sole route between the private and the public areas of the infrastructure



10

Trusting Organisations & Users

- Recognise three entities: individuals, groups and organisations (or domains)
- Entities verified through public key architecture (IdentityManager)
- Locally managed access control list determine which entities have access to local resources
- Non-local users can be mapped to individual, single, or guest accounts (resource dependant policy)

11

Access Control List

```
<domainManager manage="jini://jini.doc/" name="icpc">
  <promote domain="jini://trident.doc">
    <deny startDay="saturday"> <entity type="group"
      name="ra"/> </deny>
    <allow> <entity type="organisation" name="ICPC"/>
    </allow>
    <resource name="ap3000">
      <allow> <entity type="person" name="nfurSigned"/>
      </allow>
    </resource>
    <resource name="atlas"/>
  </promote>
  <!--DomainManager promotes resources in a second JINI -->
  <promote domain="jini://ariane.doc"> ... </promote>
</domainManager>
```

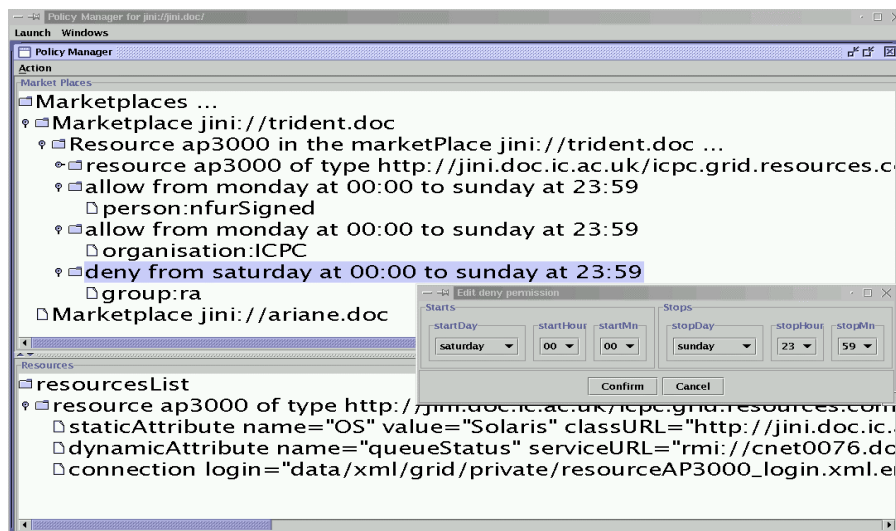
12

Individuals, Groups & Organisations

```
<identificationOrganisation
  publish="jini://jini.doc" name="ICPC"
  keystore="./data/security/keystore/keystore.ICPC">
  <group name="ala">
    <user name="aem3Signed"/>
    <user name="asm100Signed"/>
    <user name="nforSigned"/>
  </group>
  <group name="ra">
    <user name="asm100Signed"/>
    <user name="nforSigned"/>
  </group>
  <knownOrganisation name="inria"/>
</identificationOrganisation>
```

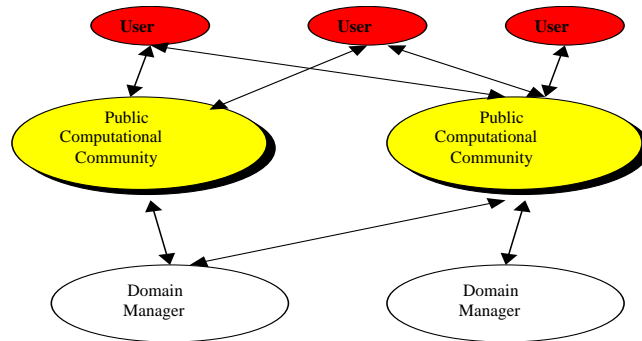
13

The Policy Manager GUI



14

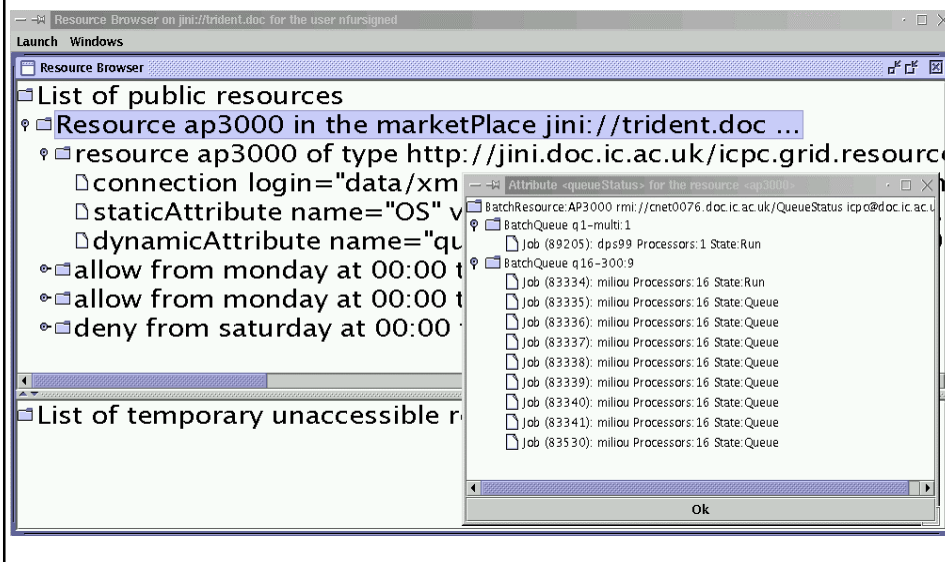
Building Computational Communities



The Computational Community allows Users & Domain Managers to advertise their capabilities and requirements.

15

The Resource Browser GUI



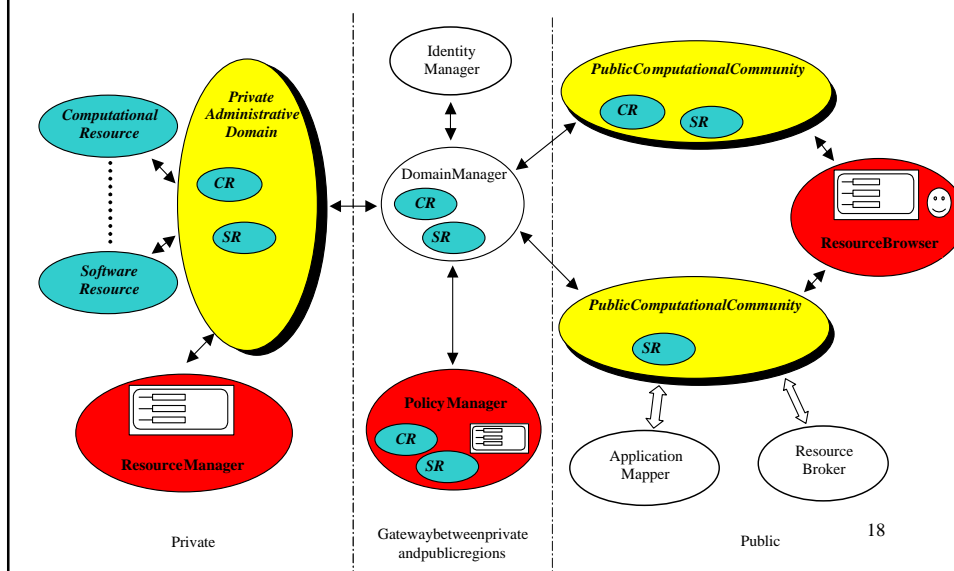
Resources in a Computational Community

A resource might be:

- *Available* → shown with all its public attributes and access policies. It's possible to get the values of the dynamic attribute(s) and to connect to the resource
- *Temporary unavailable* → only the access policies are shown
- *Always unavailable* → no information is shown

17

Java/Jini Grid Architecture



18

Exploiting Grid Meta -data

- The Computational Community contains information:
 - On resource capability, usage and access policy
 - On the user's job and their application
- Use this information to ensure:
 - The user's job is run as specified (e.g. deadline)
 - All resources are fully utilised (if possible)
 - The mapping of jobs to resources is 'optimal'

19

Higher-level Grid Services

- Application Mapper
 - Uses resource and application information to optimise resource selection for an application
(e.g. are 16 PC processors better than 8 Alpha processors?)
- Resource Broker
 - Uses computational economicsto balance priorities
(e.g. 16 underused PC processors may cost less for longer than 8 subscribed Alpha processors) over
 - Factor queue length into resource selection
(e.g. Will a slow unloaded resource provide shorter job turnaround time than a faster but heavily loaded resource?)

20

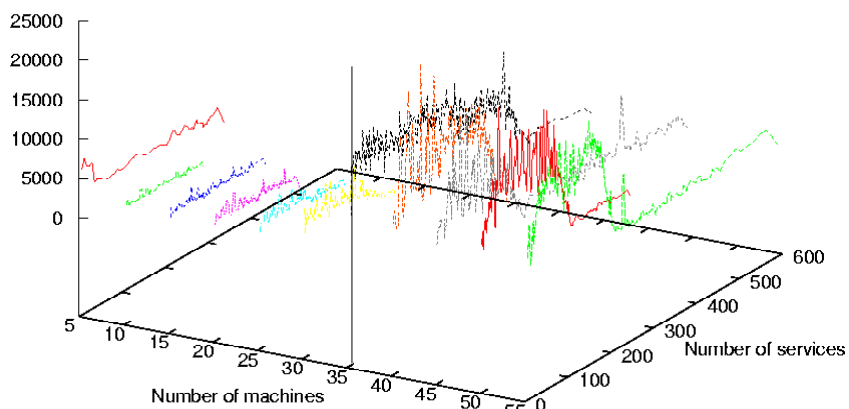
CurrentStatus

- Improving functionality & performance
- Merging Application Mapper and Resource Broker infrastructures into middleware
- Allow higher -level services to 'plugin' to the computational community
- Simulated loading using local Condor pool

21

Response to middleware under load

Response time (milliseconds)



Server: 300MHz PIII 60Mb

22

Summary

- Use Java & Jini to federate Grid resources
- Hold meta-data from the user, their applications and the resources to guide decision making
- Use the meta-data to provide automatic higher-level services to hide complexity from the user
- Application Mapper and Resource Broker enables the scientist to fully exploit the resources within their computational communities

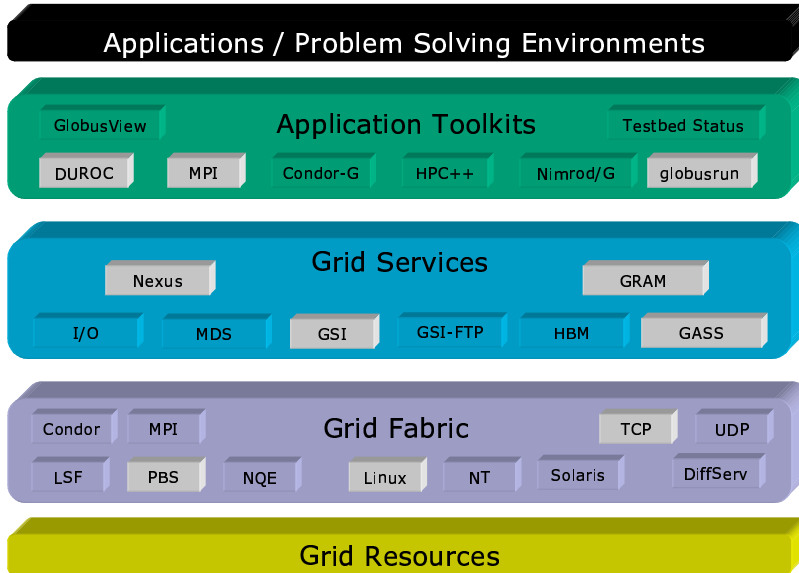
23

Future Work

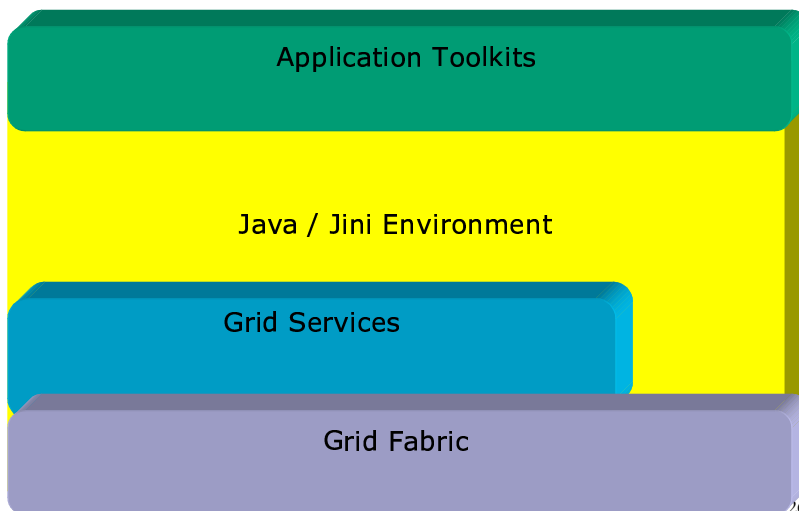
- Integration with other execution environments
- Developing code and user level documentation
- Integration with other Grid Infrastructures

24

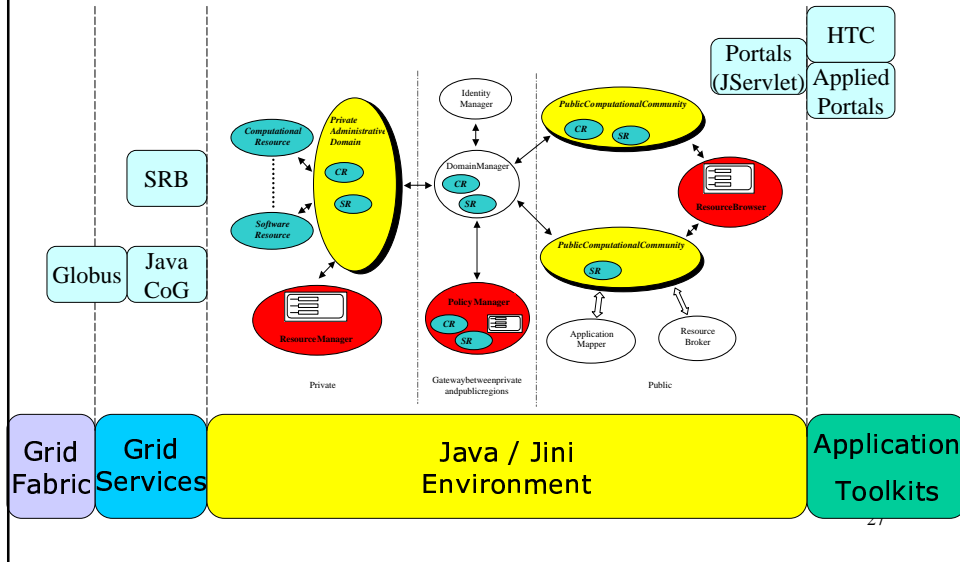
LayeredGridArchitecture



ExploitLayeredArchitecture



Java and Jini Within the Grid



Acknowledgements

- Londone -ScienceCentre, GridMiddlewareGroup:
 - JohnDarlington
 - Steven Newhouse
 - AnthonyMayer
 - Nathalie Furmento
 - Stephen McGough
 - TonyField
- Furtherinformation:
 - <http://www-icpc.doc.ic.ac.uk/components/>
- Contact: icpc-sw@doc.ic.ac.uk
- Funding:EPSRCGR/N13371

TheResourceBrowserGUI(2)



```
ap3000
The local command-shell is not available in this distribution.
Initializing random generator, please wait...done
Connected to server running SSH-1.5-1.2.27
Last login: Thu Aug 23 15:13:40 2001 from polaris.doc.ic.a
System News (updated 23/08/00):  http://www-icpc.doc.ic.ac.uk/services/news/
Email address for ICPC helpdesk: icpc-help@doc.ic.ac.uk
No mail.
System News (updated 23/08/00):  http://www-icpc.doc.ic.ac.uk/services/news/
Email address for ICPC helpdesk: icpc-help@doc.ic.ac.uk
~:> whoami
whoami
grid-00
~:> □
```

LoadingGraph(1)

