AHE Server Deployment and Hosting Applications

Stefan Zasada
University College London

AHE Server Implementation

- WSRF::Lite => services developed in Perl
- WebDAV server
- GridSAM => Globus grid
  => Sun Grid Engine
  => Condor pool
  => Unicore
- MyProxy
- PostgreSQL database
- Apache/Tomcat container

Service Architecture of the AHE

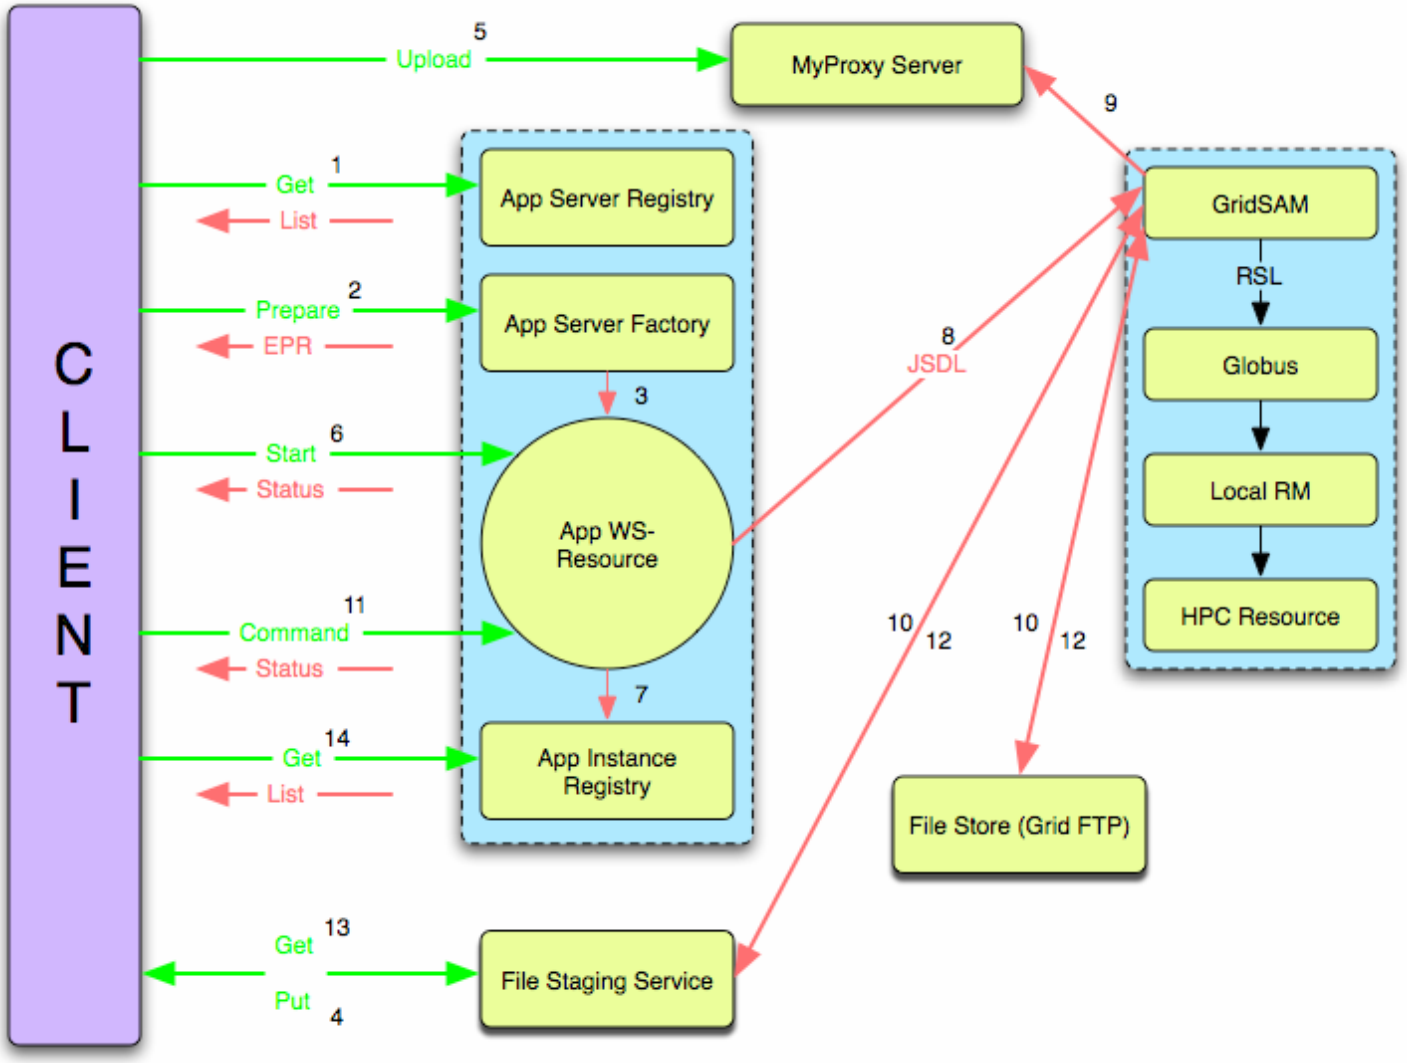

WSRF::Lite

- WSRF::Lite – An Implementation of the Web Services Resource Framework

  http://www.sve.man.ac.uk/Research/AtoZ/ILCT

- We use WSRF::Lite **http://www.sve.man.ac.uk/Research/AtoZ/ILCT** as the AHE middleware.

- Each instance of the running Application on the Grid is represented by a stateful Web Service Resource that conforms to the WSRF specification.

GridSAM

- Produced by Imperial College through the OMII managed programme, distributed with the OMII distribution
- Hosted in the OMII Tomcat Container
- Key to GridSAM is JSDL – Job Submission Description Language. GridSAM provides a uniform abstraction of resource managers – AHE only has to understand JSDL.
- JSDL is a new standard coming out of OGF for replacing things like Globus RSL, etc.

Security

- Client/Server and Server/GridSAM communication secured with mutually authenticated HTTPS
- MyProxy provides a mechanism for giving services GSI Proxy certificates.
- User stores proxy certificate on MyProxy server.
- GridSAM retrieves proxy and uses it to submit to Globus GRAM on NGS machines

File Staging Area

- AHE supports the case where the client has the required input files.
- The File Staging Area, FSA, allows the client to stage files to a place that the application can access them from.
- The client uses HTTP POST to send a file to the FSA and HTTP GET to download a file.
- Output files are placed in FSA by the AHE for the client to download.

FileStore

- Not all files may be stored on the client machine.
- The FileStore is anywhere that a files are stored that are required by the application eg on a GridFTP server.
- The client may not be able to access the FileStore directly - the client may not support the required protocol.

AHE Server Deployment

The expert user must:

- Sets up the OMII container:
  - Installs Tomcat container
  - Deploys WebDav server inside Tomcat
  - Installs and configures PostgreSQL database
  - Installs GridSAM in OMII container
  - Installs WSRF::Lite
  - Installs AHE server
  - Sets up keystores
- Need to install OMII container with HTTPS enabled
- Once deployed, any number of applications can be hosted

OMII Supported Systems

- Recent version of Java (1.4+)
- Perl 5.6 +
- OMII QA Testing on
  - x86_rhas_4 (Redhat Enterprise Linux AS 4)
  - x86_rhes_3 (Redhat Enterprise Linux ES 3)
  - x86_suse_9.0 (SUSE 9)
  - x86_deb_3.1 (Debian)
  - x86_suse_10.0 (OpenSuSE 10.0)
- Limited support for 64-bit systems

Hosting a New Application

Expert user must:

- Install and configure application on all resources on which it is being shared
- Create a JSDL template for the application (easily cloned from exiting template)
- Add the application to the RMInfo.xml file
- Run a script to reread the configuration

Documentation covers whole process of deploying AHE & applications on NGS and TeraGrid

AHE Configuration file

```
<ahe:RM>
  <wsa:EndpointReference xmlns:wsa="http://www.w3.org/2005/03/addressing">
    <wsa:Address>https://ohm.chem.ucl.ac.uk:18443/gridsam_ncsa/services/gridsam</wsa:Address>
  </wsa:EndpointReference>
  <ahe:app>
    <ahe:name>namd</ahe:name>
    <ahe:JSDLTemplate>config/JSDLTemplates/namd.ncsa.jsdl</ahe:JSDLTemplate>
  </ahe:app>
  <ahe:type>TeraGrid</ahe:type>
  <ahe:CPUCount>1262</ahe:CPUCount>
  <ahe:arch>ia64</ahe:arch>
  <ahe:memory>2524</ahe:memory>
  <ahe:virtualMemory>10000</ahe:virtualMemory>
  <ahe:opSys>Linux 2.4.21-SMP</ahe:opSys>
  <ahe:IP>127.0.0.1</ahe:IP>
  <ahe:wallTimeLimit>1440</ahe:wallTimeLimit>
  <ahe:commonName>NCSA-Mercury</ahe:commonName>
</ahe:RM>
```

Sample JSDL Template

```
<JobDefinition xmlns="http://schemas.ggf.org/jsdl/2005/11/jsdl">
  <JobDescription>
    <JobIdentification>
      <JobName>NAMD_NGS_LEEDS</JobName>
      <Description>THIS IS FOR NAMD ON THE NGS</Description>
      <JobAnnotation>NAMD Annotation</JobAnnotation>
    </JobIdentification>
    <Application>
      <mpi:MPIApplication xmlns="http://schemas.ggf.org/jsdl/2005/11/jsdl-posix"
    xmlns:mpi="urn:gridsam:mpi">
        <Executable>/usr/local/Cluster-Apps/namd-2.5-intel/bin/namd2-mpi</Executable>
        <Environment name="NGSMODULES">gm/2.0.8</Environment>
      </mpi:MPIApplication>
    </Application>
  </JobDescription>
</JobDefinition>
```

Develop Client Plugin for App

- Plugins developed in Java
- Plugins implement AHEConfParser interface and follow specific naming convention
- Plugin parses application input files to automatically discover input and output files to stage
- Plugin can also add job arguments, set stdout and stderr etc
- Plugin .class files dropped into plug-in directory and picked up by GUI/command line clients

Current Deployed Applications

- Currently hosting:
  - NAMD
  - LAMMPS
  - DL_POLY
  - LB3D
  - Gromacs
  - CHARMM
- Plan to host:
  - Trubal
  - POLCOMS

Future Plans

- Use to launch RealityGrid steering web service and steered applications
- Clients to run on a PDA (developed at Loughborough)
- Co-allocate resource reservations (HARC)
- Orchestrate complex workflows (using BPEL?)
- Coupled models – host applications which are made up of other application components

Extending Reach

Local UCL resources

GridSAM / SGE

GridSAM / Globus

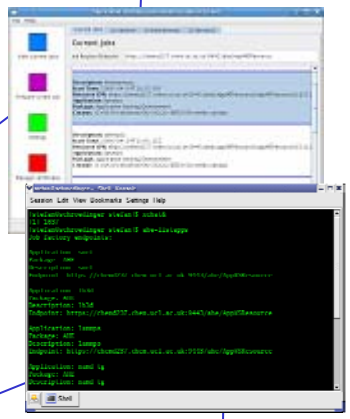

NGS

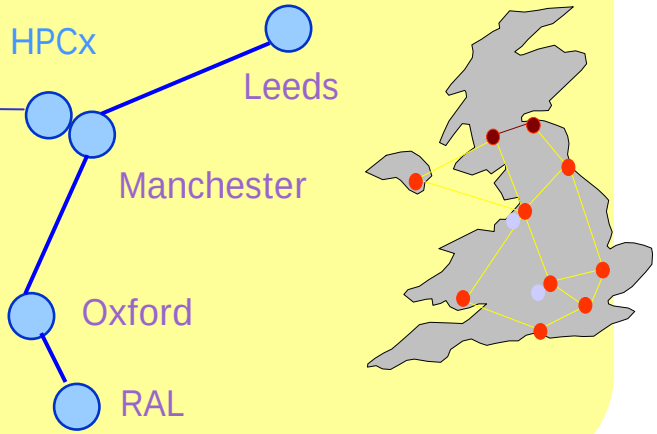

TeraGrid

GridSAM / Globus

GridSAM / UNICORE

DEISA

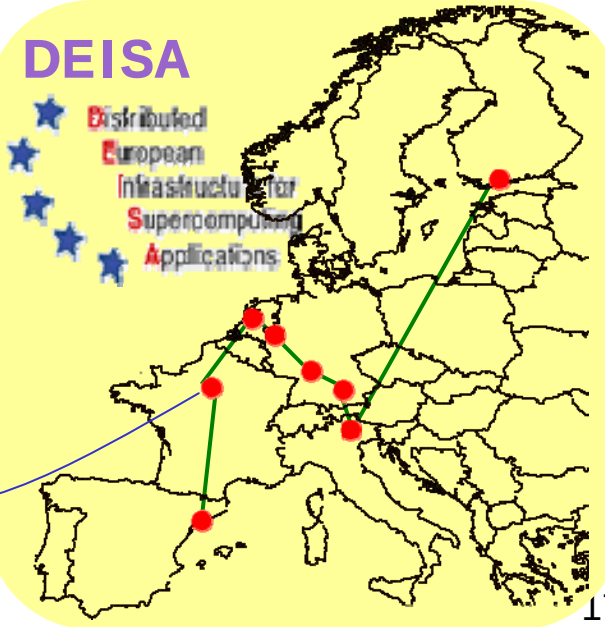

RealityGrid

open middleware infrastructure institute uk
www.omii.ac.uk

EPSRC

Summary

- The AHE provides a lightweight, easily deployable environment for running unmodified scientific applications on the grid and local resources
- The AHE server is designed to be deployed by an expert user who uses it to share applications installed on grid resources
- The client is easily installed by any end user, requiring no intervention by system/network administrators
- We are keen support groups deploying the AHE for their own applications/communities

Acknowledgements

- UCL: Matt Harvey, Laurent Pedesseau, Radhika Saksena, James Suter, Phil Fowler, Kashif Sadiq, Mary-Ann Thyveetil, Giovanni Giupponni, Simon Clifford
- Manchester: Mark Mc Keown, Stephen Pickles, Rob Haines, Andy Porter
- GridSAM Development Team
- RZG: Thomas Soddemann, Hermann Lederer
- NeSC Training Team
- EPSRC
- OMII

Further Information

- Released in OMII 3.2.0

  http://www.omii.ac.uk/downloads/
- RealityGrid web site:

  http://www.realitygrid.org/AHE
- NeSCForge:

  http://forge.nesc.ac.uk/projects/ahe/
- Mailing list:

  http://www.mailinglists.ucl.ac.uk/mailman/listinfo/ahe-discuss