



Enabling Grids for E-scienceE

Supporting Diverse Local Virtual Organisations

Douglas McNab

ScotGrid - University of Glasgow

www.eu-egee.org



- **Case Study 1**
 - Device Modelling Group [vo.nanocmos.ac.uk]
 - Issues and Experiences
 - Outcomes
- **Case Study 2**
 - Photonics [vo.optics.ac.uk]
 - Issues and Experiences
 - Outcomes
- **Case Study 3**
 - Solid State Physics [vo.ssp.ac.uk]
 - Issues and Experiences
 - Outcomes
- **Lessons Learned**

- **Local Device Modelling group**
 - Approx 10 Users in UK based in Glasgow and Manchester
 - 1 Month Usage Trial
 - Originally used to direct batch submission
 - Usage: Python Code, 7 day, serial array jobs
- **Issues**
 - Initial access to the system
 - It was felt this was overly complicated.
 - Job Submission
 - Various methods; command line, Ganga, proxy lifetime issues.
 - Job Monitoring
 - Disparity between command line and Ganga.
 - Reliability
 - Perception that there has a high failure rate.

- **Outcomes**

- Initial access to the system

- Streamlined local documentation for access.

- Job Submission

- Creation of **gqsub**, a qsub like interface to gLite middleware
 - *See poster contribution titled 'Data awareness in gqsub'.*

- Job Monitoring

- The usage of **gqsub** re-introduced consistency with command line tools.

- Reliability

- The perception of unreliable system was unfounded. It seemed to stem from past experiences with early gLite components such as the RB.
- Update of procedures surrounding middleware upgrade notifications.

- **Local Photonics group**

- 3 local users and potential UK users in London, Southampton and some in France. Now installed in Greece.
- Usage: Licensed software by Lumerical Inc, MPI

- **Issues**

- License Management
 - Always a sticking point for Grid, EGEE and gLite.
- Integration for use with gLite middleware
 - In order to support licences management at the batch level custom job parameters need to be able to passed from JDL.
- MPI specification
 - MPI version available but not all PBS/Torque node specifications available.
- Small VO
 - Lack of resources, limited site licenses, potential external users.
 - *SL4/5 & UI install a stumbling block.*

- **Outcomes**

- License Management

- Licences dealt with via consumable resources in the scheduler and the batch system.

- Integration for use with gLite middleware

- Passing custom job parameters with JDL, in this case a software license tag, was made possible with the introduction of CREAM.

- MPI specification

- MPI works but still additional development required. Using custom job parameters with JDL allowed specification of whole nodes and processes per node (:ppn).

- Small VO

- Creation of an off-the-self optics UI Virtual Box image.
 - *Copy your certificate and submit the test job.*

Submission of a licensed Lumerical job

JDL

JDL

```
cerequisites =
" GlueHostApplicationSoftwareRunTimeEnvironment=="FDTD\ "";
```

```
Requirements= other.
GlueHostApplicationSoftwareRunTimeEnvironment=="FDTD";
```

or

```
In [pbs/lsf]_local_submit_attributes.sh

if [ "$GlueHostApplicationSoftwareRunTimeEnvironment" ==
"FDTD" ]; then
echo "#PBS -l software=FDTD"
fi
```

```
In /opt/glite/etc/glite_wms.conf

CeForwardParameters =
"GlueHostApplicationSoftwareRunTimeEnvironment";
```

CREAM CE
(via BLAH)

WMS

FDTD
LICENSE
MANAGER

PBS Batch System
Resource_List.software = FDTD

PBS/MAUI manages licenses through
specification of:
NODECFG[GLOBAL] GRES=FDTD:10

- **Local Solid State Physics group**
 - 2 local users
 - New to batch/grid submission
 - Usage: User code and CASTEP
- **Issues**
 - MPI Integration
 - MPI versions from EGEE not suitable in all cases.
 - MPI Specification
 - MPI over many nodes expected but again propagation of generic parameters from JDL not possible.
 - Small VO
 - Lack of resources, limited experience.

- **Outcomes**

- MPI Integration

- Rebuild of OPENMPI for F90 and Torque Support.

- MPI Specification

- 'cerequisites' JDL attribute for CREAM coupled with `pbs_local_submit_attribute.sh` functionality allowed specification on `:ppn` and `wholenodes`.

JDL

```
cerequisites = "smpgranularity == ${PPN}";
```

CREAM CE
(via BLAH)

```
In [pbs/lsf]_local_submit_attributes.sh
```

```
...
```

```
if [ ! -z "$smpgranularity" ] ; then
```

```
    echo "#PBS -l nodes=$bls_opt_mpinodes:ppn=$smpgranularity"
```

```
....
```

- Small VO

- Customised user scripts as a precursor to `gqsub` usage.

- **‘We just want direct submission to the batch system’**
 - Custom tools can provide the interface users desire
 - They should be shielded from needless complexity
- **As you would expect small VOs require more help**
 - NA4 and SA1 co-location at Glasgow has helped a great deal.
 - Allowing middleware customisation to fit with user needs.
- **Computing Element**
 - Fundamental to support the needs of the smaller VO
 - CREAM has made this possible.