

Robert Engel
Alexander Beck-Ratzka
Max-Planck-Institute
for Gravitational Physics




Gravitational Wave Data Analysis

From Deployment to Production
GEO600

7th Astrogrid-D Meeting
TU-München, Jun. 11th 2007

Laser Interferometer Gravitational Wave Observatory



Vue sur www.techno-science.net
Cette image peut être protégée



GEO600 Grid Scenario



Information Service
fetch host list

1



/O=GermanGrid/OU=AEI/CN=Robert Engel



mintaka.aip.de:[svn](#)
mintaka.aip.de:[is](#)



einstein.phys.uwm.edu:[http](#)



buran.aei.mpg.de:[gridftp](#)
buran.aei.mpg.de:[mysql](#)

GEO600 Grid Scenario



Grid Service Monitoring



/O=GermanGrid/OU=AEI/CN=Robert Engel

2



mintaka.aip.de:svn
mintaka.aip.de:is



einstein.phys.uwm.edu:http



buran.aei.mpg.de:gridftp
buran.aei.mpg.de:mysql

GEO600 Grid Scenario



3



/O=GermanGrid/OU=AEI/CN=Robert Engel

Fetch Task from Database



mintaka.aip.de:svn
mintaka.aip.de:is



einstein.phys.uwm.edu:http



buran.aei.mpg.de:gridftp
buran.aei.mpg.de:mysql

GEO600 Grid Scenario



GT4 Job Submission
Fork/PBS/SGE



/O=GermanGrid/OU=AEI/CN=Robert Engel

4



mintaka.aip.de:svn
mintaka.aip.de:is



einstein.phys.uwm.edu:http



buran.aei.mpg.de:gridftp
buran.aei.mpg.de:mysql

GEO600 Grid Scenario



/O=GermanGrid/OU=AEI/CN=Robert Engel

5

Globus Prestage Data



mintaka.aip.de:svn
mintaka.aip.de:is



einstein.phys.uwm.edu:http



buran.aei.mpg.de:gridftp
buran.aei.mpg.de:mysql

GEO600 Grid Scenario



run prestage code
checkout project

6



/O=GermanGrid/OU=AEI/CN=Robert Engel



mintaka.aip.de:[svn](#)
mintaka.aip.de:[is](#)



einstein.phys.uwm.edu:[http](#)




buran.aei.mpg.de:[gridftp](#)
buran.aei.mpg.de:[mysql](#)

GEO600 Grid Scenario



/O=GermanGrid/OU=AEI/CN=Robert Engel



7



mintaka.aip.de:svn
mintaka.aip.de:is



[einstein.phys.uwm.edu:http](http://einstein.phys.uwm.edu)



buran.aei.mpg.de:gridftp
buran.aei.mpg.de:mysql

GEO600 Grid Scenario



/O=GermanGrid/OU=AEI/CN=Robert Engel

8

Globus Poststage Data



mintaka.aip.de:svn
mintaka.aip.de:is




einstein.phys.uwm.edu:http



buran.aei.mpg.de:gridftp
buran.aei.mpg.de:mysql

GEO600 Grid Scenario



/O=GermanGrid/OU=AEI/CN=Robert Engel

- ✓ Automatic Account Management using VORMS
- ✓ Automatic Resource Discovery using MDS/IS
- ✓ Grid Resource Monitoring using gm.pl
- ✓ Automatic GT4 Job Submission
- ✓ Automatic GT4 Job Monitoring
- ✓ Usage grows with the grid growing!



mintaka.aip.de:svn
mintaka.aip.de:is



einstein.phys.uwm.edu:http



buran.aei.mpg.de:gridftp
buran.aei.mpg.de:mysql

Status of AstroGrid – Nov. 2006

	Type	Globus	GSISSH	GSISCP	GRAM
gavo2.aip.de	Server	Yes	Yes	Yes	Yes
gavo3.aip.de	Server	Yes	Yes	Yes	Yes
astar.aip.de	Workstation	Yes	Yes	Yes	Yes
photon.aip.de	Workstation	Yes	Yes	Yes	Yes
dublin.aip.de	Workstation	Yes	No	No	No
supergrid.aei.mpg.de	Server	Yes	Yes	Yes	Yes
altix.lrz-muenchen.de	SMP	Yes	Yes	Yes	Yes
lxsrv1.lrz-muenchen.de	Frontend	Yes	Yes	Yes	Yes
gavo1.aip.de	Workstation	Yes	No	No	No
mintaka.aip.de	Server	Yes	No	No	No
cashmere.aip.de	Server	Yes	Yes	Yes	Yes

Only simple scenarios were well supported!

Status of GEO600 – Nov. 2006

deployment was done using interactive logins

job submission was done by hand, based on resources being available

job monitoring was done by hand

only few grid resources with PBS/SGE were available at the time

Status of AstroGrid / DGrid - June 2007

	Type	Globus	GSISSH	GT4
titan.ari.uni-heidelberg.de	Server	Yes	Yes	Yes
saiph.ari.uni-heidelberg.de	Server	Yes	Yes	Yes
hydra.ari.uni-heidelberg.de	Frontend	Yes	Yes	Yes
srvgrid01.offis.uni-oldenburg.de	Frontend	Yes	Yes	Yes
udo-gt01.grid.uni-dortmund.de	Frontend	Yes	Yes	Yes
lx32ia1.lrz-muenchen.de	Frontend	Yes	No	Yes
lx64ia2.lrz-muenchen.de	Frontend	Yes	Yes	Yes
a01.hlr2.lrz-muenchen.de	Frontend	Yes	Yes	Yes
gramd1.d-grid.uni-hannover.de	Frontend	Yes	Yes	Yes
dgrid-globus.rz.rwth-aachen.de	Frontend	Yes	No	Yes
supergrid.aei.mpg.de	Server	Yes	Yes	Yes
buran.aei.mpg.de	Server	Yes	Yes	Yes

Wide range of resources available!

Status of GEO600 – June 2007

- ✓ we support automatic deployments on all grid resources available
- ✓ we monitor our jobs automatically
- ✓ we automatically submit jobs to Fork/SGE/PBS resources
- ✓ we use a central storage location via gridftp for data
- ✓ we use a MySQL Database to store data and do some accounting
- ✓ we run hundreds of jobs daily each about one hour long
- ✓ we used more than 5000 CPU hours for science last month alone

- ✗ we want to move to the newly purchased Astrogrid storage location
- ✗ we want to integrate with the Information Service to replace MySQL
- ✗ we want to test using a grid scheduler like GridWay!

Deployment Situation

- × All you'll get is a bare account, sometimes less!
- × Only few admins will install software-x.y.z.tar.gz just for you!
- × On most sides you will find gnu but no other fancy compilers!
- × usage policies must be negotiated face to face
- × access policies must be negotiated face to face
- × Fact: the grid is by definition inhomogeneous and dynamic!

Small Steps toward a successful Deployment

- ✓ treat your research code as public
- ✓ test your research code on as many platforms as possible
- ✓ reduce software requirements to a bare minimum
- ✓ have your code in a repository and archives in a web location
- ✓ hand out your code and documentation to external test users

- ✓ try to build your code on interactive grid logins
- ✓ try to script the build process and submit it as a grid job
- ✓ try to submit short test jobs to servers and frontends
- ✓ try to submit test jobs to queuing systems
- ✓ having one solution is good, having two is better!

Progress carefully in small steps!

Globus RSL job description and submission using GT4

- `prestige prestige.sh` from `gridftp://buran.aei.mpg.de/store/GEO600/`
- `prestige job archive <jobid>.tar`

→ `poststage job archive, stdout, stderr and log to gridftp://buran.aei.mpg.de/store/GEO600/`

Globus RSL job description and submission using GT4

- `prestige prestige.sh` from `gridftp://buran.aei.mpg.de/store/GEO600/`
- `prestige job archive <jobid>.tar`

`/bin/sh prestige.sh <arguments>`

- checks and sets environment settings
- checks for svn and checks out `svn://svn.gac-grid.org/software/GEO600` or
- checks for gsiscp and checks out `gsiscp://buran.aei.mpg.de/store/GEO600`
- checks for perl and runs `perl eah.pl <arguments>`

→ exit \$?

- `poststage job archive, stdout, stderr and log` to `gridftp://buran.aei.mpg.de/store/GEO600/`

Globus RSL job description and submission using GT4

- prestige prestige.sh from gridftp://buran.aei.mpg.de/store/GEO600/
- prestige job archive <jobid>.tar

/bin/sh prestige.sh <arguments>

- checks and sets environment settings
- checks for svn and checks out svn://svn.gac-grid.org/software/GEO600 or
- checks for gsiscp and checks out gsiscp://buran.aei.mpg.de/store/GEO600
- checks for perl and runs perl eah.pl <arguments>

perl eah.pl <arguments>

- builds and checks ncurses
- builds and checks MySQL
- builds and checks perl DBI, perl DBD, Time::HiRes
- installs gravitational wave code
- unpacks prestaged data archive
- fetches runtime data from MySQL database on buran.aei.mpg.de:24999
- checks and runs gravitational wave code <arguments>

- packs data archive for poststaging
- sends data back to MySQL database on buran.aei.mpg.de:24999
- exit \$?

- exit \$?

- poststage job archive, stdout, stderr and log to gridftp://buran.aei.mpg.de/store/GEO600/

Globus RSL job description and submission using GT4

- prestige prestige.sh from gridftp://buran.aei.mpg.de/store/GEO600/
- prestige job archive <jobid>.tar

/bin/sh prestige.sh <arguments>

- checks and sets environment settings
- checks for svn and checks out svn://svn.gac-grid.org/software/GEO600 or
- checks for gsiscp and checks out gsiscp://buran.aei.mpg.de/store/GEO600
- checks for perl and runs perl eah.pl <arguments>

perl eah.pl <arguments>

- builds and checks ncurses
- builds and checks MySQL
- builds and checks perl DBI, perl DBD, Time::HiRes
- installs gravitational wave code
- unpacks prestaged data archive
- fetches runtime data from MySQL database on buran.aei.mpg.de:24999
- checks and runs gravitational wave code <arguments>

running the gravitational wave data analysis code

→ exit \$?

- packs data archive for poststaging
- sends data back to MySQL database on buran.aei.mpg.de:24999
- exit \$?

→ exit \$?

- poststage job archive, stdout, stderr and log to gridftp://buran.aei.mpg.de/store/GEO600/

GEO600 / LIGO Production

- ✓ GEO600 sources build on all grid hosts we have access to
- ✓ successful job submission to Fork / PBS and SGE resources
- ✓ ca. 5000 jobs submitted last month alone
- ✓ more than 5000 cpu hours used for gravitational wave data analysis
- ✓ usage rapidly growing ...
- ✓ ... and we have accounting / statistics build in!
- ✓ the infrastructure we developed can be re-used by other use-cases!

Issues to be Resolved

- x access policies vary to a great extent from machine to machine
- x usage policies are not available in machine readable format
- x queue policies are not available in machine readable format
- x account granted, but no \$HOME setup
- x \$HOME setup, but no ssh keys to access internal nodes
- x many nodes, but no queuing system in place
- x small /home partitions that fill up quickly
- x bigger /data partitions available, but that info is not stored
- x the grid looks different from every resource (matrix tests!)
- x The grid-ka certificate runs invalid one day before a presentation!
- x ...

Acknowledgments

Alexander Beck-Ratzka (AEI)
Thomas Radke (AEI)
Helmut Heller (LRZ)
Hamza Mehammed (LRZ)
Stefan Freitag (Dortmund)