




Enabling Grids for
E-science in Europe



Grid-Ireland


Brian Coghlan, John Walsh and Stephen Childs
Trinity College Dublin



Grid-Ireland






- **National Computational Grid**
 - Funded by Enterprise Ireland, HEA, SFI, EU
 - Single Grid layer to support Irish researchers
 - Umbrella for VOs
- **Directors**
 - Brian Coghlan (TCD), Andy Shearer (NUIG)
 - Luke Drury (DIAS), John Morrison (UCC)
- **Grid-Ireland Management**
 - John Walsh (Grid Manager)
 - Stephen Childs (Deputy Grid Manager)



National VOs



	<p>Grid-enabled analysis of natural phenomena DIAS, UCD, NUIG, TCD, UCC, DCU, Armagh, Met Eireann, HEAnet</p> <p>Resources: 200 CPUs at UCD, 16 CPUs at DIAS, 16 shared CPUs at TCD, 40% of 1000 CPUs & 16-way Bull SMP at ICHEC</p>
	<p>non-von-Neumann grid execution architecture UCC, TCD, NUIG</p> <p>Resources: 16 shared CPUs at TCD</p>
	<p>Data mining of Irish National Seabed Survey NUIG, UCD, UCC, GSI, Marine Institute</p> <p>Resources: cluster at NUIG, datafarms at NUIG & UCD</p>
<p>GeneGrid</p>	<p>Exploratory VO for Irish bio-informaticians</p> <p>Resources: shared access to SoloVO resources</p>
<p>TestGrid</p>	<p>Grid administration VO: OpsCentre</p> <p>Resources: 1 virtual test CPU at each site</p>
<p>SoloVO</p>	<p>A catch-all VO for unaffiliated users</p> <p>Resources: 46 (soon 64) CPUs at TCD</p>

International Links



- **UK e-Science:**
 - Gateway in e-Science at Belfast
- **LCG/EGEE:**
 - 184 sites and 24,909 CPUs worldwide
 - TCD OpsCentre is ROC for Ireland
- **CrossGrid/int.eu.grid:**
 - European Interactive Grid

- + SouthEasternEurope
- + SouthWesternEurope
- UKI
 - GridIreland
 - cpDIASie
 - csQUBuk
 - csTCDie
 - csUCCie
 - giAITie
 - giDCUie
 - giDITie
 - giITCie
 - giITTAie
 - giITTRie
 - giNUIMie
 - giRCSie
 - giULie
 - giWITie
 - mpUCDie
 - obsARMuk
 - scgNUIGie
- + LondonT2
- + NorthGrid
- + ScotGrid
- + SouthGrid
- + Tier1A
- + Countries



Grid-Ireland Architecture



- **Grid is a peer-to-peer virtual architecture**
 - Central components + distributed components
- **Grid-Ireland can provide central + some core stuff**
 - National Certification Authority, Brokers, ...
 - Core distributed infrastructure
- **But sites own their compute resources**
 - Heterogeneous mix of architectures
 - Can't tell sites what to do ...
- **Grid-Ireland architecture:**
 - Heterogeneous site resources
 - OpsCentre and national servers in TCD
 - Integrated homogeneous core infrastructure
 - 18 grid gateways: most of Irish 3rd-level sector




Grid-Ireland Gateways



- **Grid gateway:**
 - Firewall
 - Install server
 - CE/SE/UI/NM/WN for grid middleware

- **From the user perspective:**
 - Gateway architecture is the same
 - Hostnames are the same
 - Only UI, CE, SE are of interest
 - UI hosts user's grid account
 - CE is site frontend
 - SE provides storage



- **Gateway queues jobs to site resources**

Grid-Ireland OpsCentre



- **Management of Grid-Ireland infrastructure**
 - Testing, porting, customisation of middleware
 - Deployment to remote sites
 - Remote management of servers
 - Monitoring of Grid services
 - National services (e.g. Certification Authority)
- **Dissemination and training**
 - Grid courses (and online courses in development)
 - Website and Wiki
- **International operations**
 - ROC for Ireland

OpsCentre



Operations Support Team



- **People (4.5 FTE)**
 - Expert and well respected by peers in EGEE
- **Operations Support**
 - grid-ireland-help@cs.tcd.ie → Request Tracker
 - grid-ireland-site-managers@cs.tcd.ie
 - Mon to Fri (9:30am - 6:30pm) + best effort out-of-hours
 - Weekly rotation of **Operator-on-Duty**
 - Weekly operations meeting
 - Liason with HEAnet
 - Security incident, highest priority
 - Understaffed and underfunded ...
 - Evolutionary, adjusting as needs require

International Management




- **Interactions**
 - LCG/EGEE Central Management
 - LCG/EGEE UK/I Federation Management
 - UK e-Science, GridPP DB, BeSC
- **Operations**
 - EGEE ROC for Ireland
 - Weekly operations reports (per registered site)
 - Interact with GGUS (EGEE Global Grid User support)
 - Training interactions with NeSC

Operator on Duty



- **Weekly rotation of member of Operations Team to manage**
 - Request Tracker tickets
 - Takes new ticket, or assigns ticket to particular Operations Staff.
 - Review/follow up of outstanding tickets
 - Review and act upon daily logging e-mail from sites
 - Interact with GGUS (EGEE Global Grid User support).



E-mail support



- **All enquiries should be sent to:**
 - grid-ireland-help@cs.tcd.ie
- **Site Managers mailing list**
 - grid-ireland-site-managers@cs.tcd.ie
 - Low volume, subscription based
 - Discussion about the Grid-Ireland Infrastructure and resources support
 - Announcements of Grid sites downtime (additional HEAnet ticket raised)
- **Grid-Ireland CA**
 - grid-ireland-ca@cs.tcd.ie
 - Brian Coghlan and David O'Callaghan



Middleware Porting Support



- Eamonn Kenny (eamonn.kenny@cs.tcd.ie)
- To support heterogeneous site resources
- Supported Systems:

- CentOS 4.1
- RedHat 9
- Scientific Linux 3.0.4
- Suse 9.3
- RedHat 7.3

- **Work-in-Progress:**

- Fedora Core 4
- Debian Stable
- Solaris 10
- MacOS X 10.X
- AIX,
- IRIX
- Em64t in the works ...



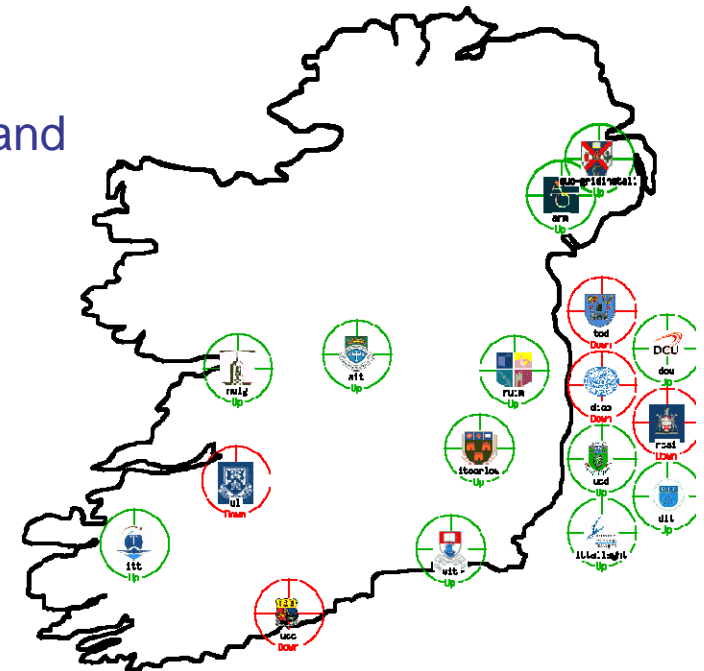
OS Type	Centos 4.1	Redhat 9	SL 3.0.4	SuSE 9.3	Redhat 7.3	FC 4	Debian Stable	Solaris 10	Darwin 7.7	AIX 6.5	Irix 6.5
OS dump	loopback	loopback	loopback	loopback	loopback	loopback	loopback	iso	dump	N/A	N/A
VM build machine	Yes	Yes	Yes	Yes	Yes	No	No	N/A	N/A	N/A	N/A
VDT build	Yes	Yes	Yes	Yes	Yes	No	No	Needs pkg	Yes	No	No
VDT RPM on Web	Yes	Yes	Yes	Yes	Yes	No	No	Needs pkg	tarballs	No	No
VM test WN machine	Yes	Yes	Yes	Yes	Yes	No	No	N/A i386/Sparc	N/A XServer	N/A Server	N/A Itself
Basic WN on Web	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
Basic WN installing	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
Job Submission Tested	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
Gfal built	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
Gfal on web	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
SFTs											

- **Support for installation, systems integration:**
 - YUM/APT repository of Middleware RPMs

Site Health and Monitoring



- **Grid-Ireland Site Functional Tests (SFT)**
 - Tests sites ability to function on the Grid-Ireland
 - Runs twice daily
- **Nagios agents across the infrastructure**
 - Tests that all essential services are working
 - E-mails operator if site service/SFT fails
 - Extensible Plugin based architecture
- **Service monitoring research**
 - Keith Rochford



AIT Infrastructure Servers (ait)

Host	Status	Services	Actions
ait-gridgate	UP	6 OK	[Icons]
ait-gridinstall	UP	2 OK	[Icons]
ait-gridstore	UP	2 OK	[Icons]
ait-gridui	UP	1 OK	[Icons]

ARM Infrastructure Servers (arm)

Host	Status	Services	Actions
arm-gridgate	UP	6 OK	[Icons]
arm-gridinstall	UP	2 OK	[Icons]
arm-gridstore	UP	2 OK	[Icons]
arm-gridui	UP	1 OK	[Icons]

DCU Infrastructure Servers (dcu)

Host	Status	Services	Actions
dcu-gridgate	UP	6 OK	[Icons]
dcu-gridinstall	UP	2 OK	[Icons]
dcu-gridstore	UP	2 OK	[Icons]
dcu-gridui	UP	1 OK	[Icons]

DIAS Infrastructure Servers (dias)

Host	Status	Services	Actions
dias-gridgate	UP	5 OK 1 WARNING	[Icons]

DIT Infrastructure Servers (dit)

Host	Status	Services	Actions
dit-gridgate	UP	6 OK	[Icons]
dit-gridinstall	UP	2 OK	[Icons]
dit-gridstore	UP	2 OK	[Icons]

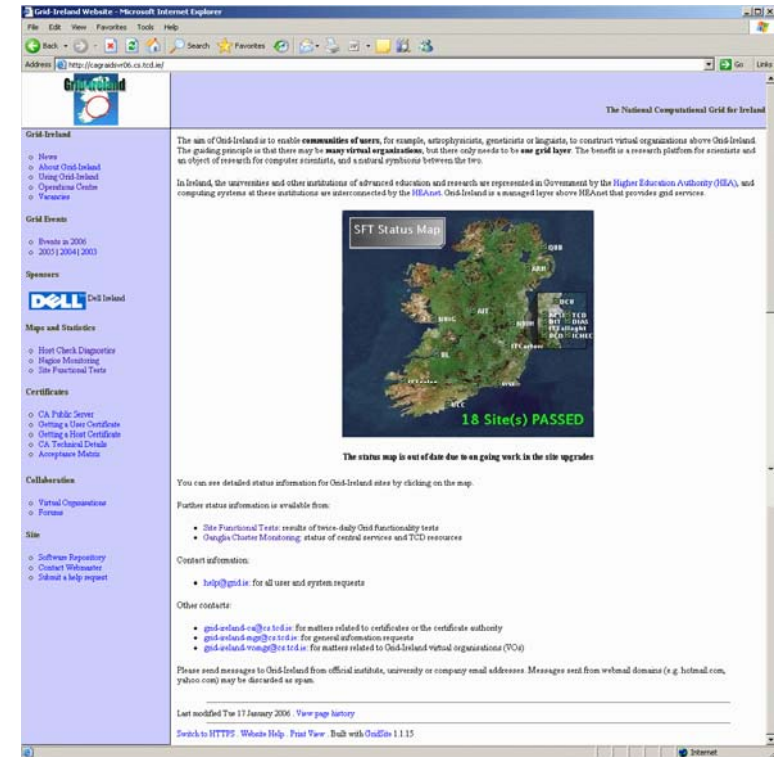
ITCarlow Infrastructure Servers (itcarlow)

Host	Status	Services	Actions
itcarlow-gridgate	UP	6 OK	[Icons]

Grid-Ireland website (I)



- <http://www.grid.ie/>
- News
- Basic Grid-Ireland usage introduction
 - Logging on to UI, then sign-on to Grid
 - Job submission
 - Useful external links
 - Data management
 - Basic Job Description Language (JDL)
- Certificate Application
 - How to, what browsers to use
 - Photographic ID, what Regional Authority to interact with
 - Applicants must read this carefully ...



Grid-Ireland website (II)



- **Porting Status web pages**
- **Grid-certificate-protected web pages**
 - Authorization based on Grid-Ireland User certificate
- **Grid-Ireland Wiki**
 - Public Wiki Pages
 - Grid-certificate-protected Wiki pages, allowing editing
- **CERN CDS agenda-maker system**
 - Course outline, timetables, presenter
 - Supplemental slides & extra materials

Grid-Ireland Courses



- **Training events linked off main page**
- **Training the future Grid experts**
 - After this course, 85 scientists will have undergone some formal introductory grid training
- **eLearning facilities being developed**
 - Adaptive eLearning research
 - Measures knowledge and reacts to suit
 - Learn at your own pace
 - Kathryn Cassidy
- **Training will be on tutorial infrastructure**
 - Won't interfere with production service



Application Development



- **Most MPI based jobs should run on Grid-Ireland unmodified**
 - MPI job submission not yet a smooth process
- **Pure grid applications**
 - Require some knowledge of JDL, workflows and data management
 - Highly parallel independent jobs are trivial
- **The OpsCentre can offer some support in developing applications**

Other Support




- Requests for help for non-Grid-Ireland virtual organisations should interact with Global Grid User Support
 - <http://www.ggus.org>

ID	VO	Date	Info
5586	none	n/a	MISCELLANEOUS (BEIJING-LCG2)
5585	cms	2005-12-02	Q on the job submission
5571	none	2005-12-01	published info inconsistency (RU-Protvin...)
5570	none	n/a	Job submission failed (LCG_KNU)
5569	none	n/a	CPU count erratic (WARSAW-EGEE)
5567	none	n/a	published info inconsistency (UCL-HEP)
5566	none	n/a	published info inconsistency (LIVHEP-LCG...)
5565	none	n/a	published info inconsistency (TW-NCUHEP)
5564	none	n/a	published info inconsistency (Ru-Trotsk...)
5554	none	2005-12-01	MON box wipp-rb.weizmann.ac.il is not in...
5549	atlas	2005-12-01	MON box tbit00.nipne.ro is not in the GO...
5542	none	2005-12-01	MON box se00.inta.es is not in the GOC d...
5523	none	2005-11-30	MON box lcg-lcfnng.its.uiowa.edu not in ...
5520	none	2005-11-30	MON box gridib03.icepp.jp not in the GOC...
5517	none	2005-11-30	MON box gridee01.info.it not in the GOC...

Future Support



- **WebCom-G**
 - Workflow engine
 - But more than just a workflow engine
- **Condensed Graphs**
 - Sequencing tasks for evaluation
 - Multiple strategies
 - Dynamically alterable
 - State captured by Graph Topology
 - Architecture Neutral
 - *Capturing the essence of a solution, not implementation detail*



Future Support



- **The WebCom Abstract Machine**
 - Multi-tiered
- **Pluggable modules**
 - Load Balancing, Fault Tolerance, Comm.Protocols, Engine Modules
 - EJB, Corba, DCOM, Globus, OGSA
- **Dynamically reconfigurable**
 - Via feedback mechanisms

