

Status of the APAC Grid Program in Australia.

Marco La Rosa

on behalf of John O'Callaghan
Executive Director, APAC

Australian Partnership for Advanced Computing

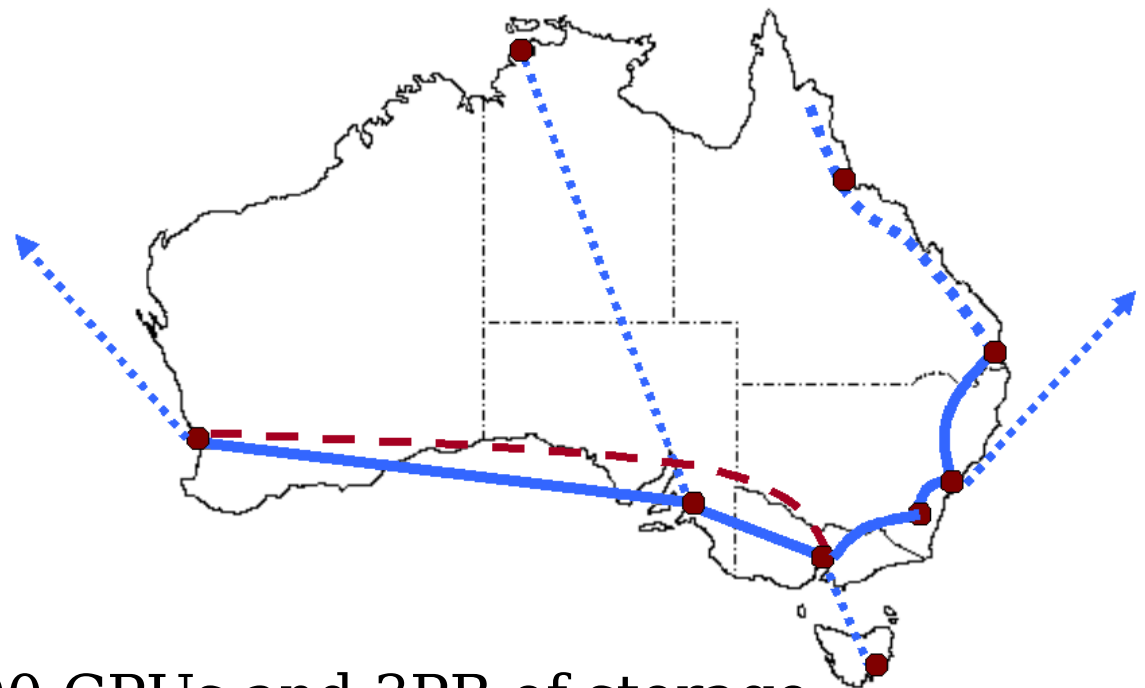
- **provides High Performance Computing, Grid infrastructure and services to Australian researchers.**

APAC National Facility:

- 1928 CPU SGI Altix at ANU in Canberra

Mid range supercomputers at other APAC partner sites:

- iVEC (Western Australia),
- SAPAC (South Australia),
- VPAC (Victoria),
- TPAC (Tasmania),
- ac3 (New South Wales),
- QCIF (Queensland),
- CSIRO



- 25 supercomputers, >4500 CPUs and 3PB of storage.

APAC National Grid Program



Aims to provide Grid infrastructure and services to support a variety of applications


→ started in early 2004.

Defined in terms of projects:

- **Compute** infrastructure,
- **Information** infrastructure,
- **User** interfaces and portals

Applications:

- Geoscience
- Earth Systems Science
- High Energy Physics
- Bioinformatics
- Chemistry
- Astronomy



SEE GRID community website

Solid Earth and Environment GRID



Different communities – different infrastructure requirements

Motivation for Virtualization



- need to support **multiple Grid interfaces** at each site in early 2004: GT2, LCG – expecting GT4, gLite
- frequent reboots during installation / testing
stability issues with some middleware esp. GT4 beta
- Incompatibilities and potential conflicts between different versions of Globus on the same machine

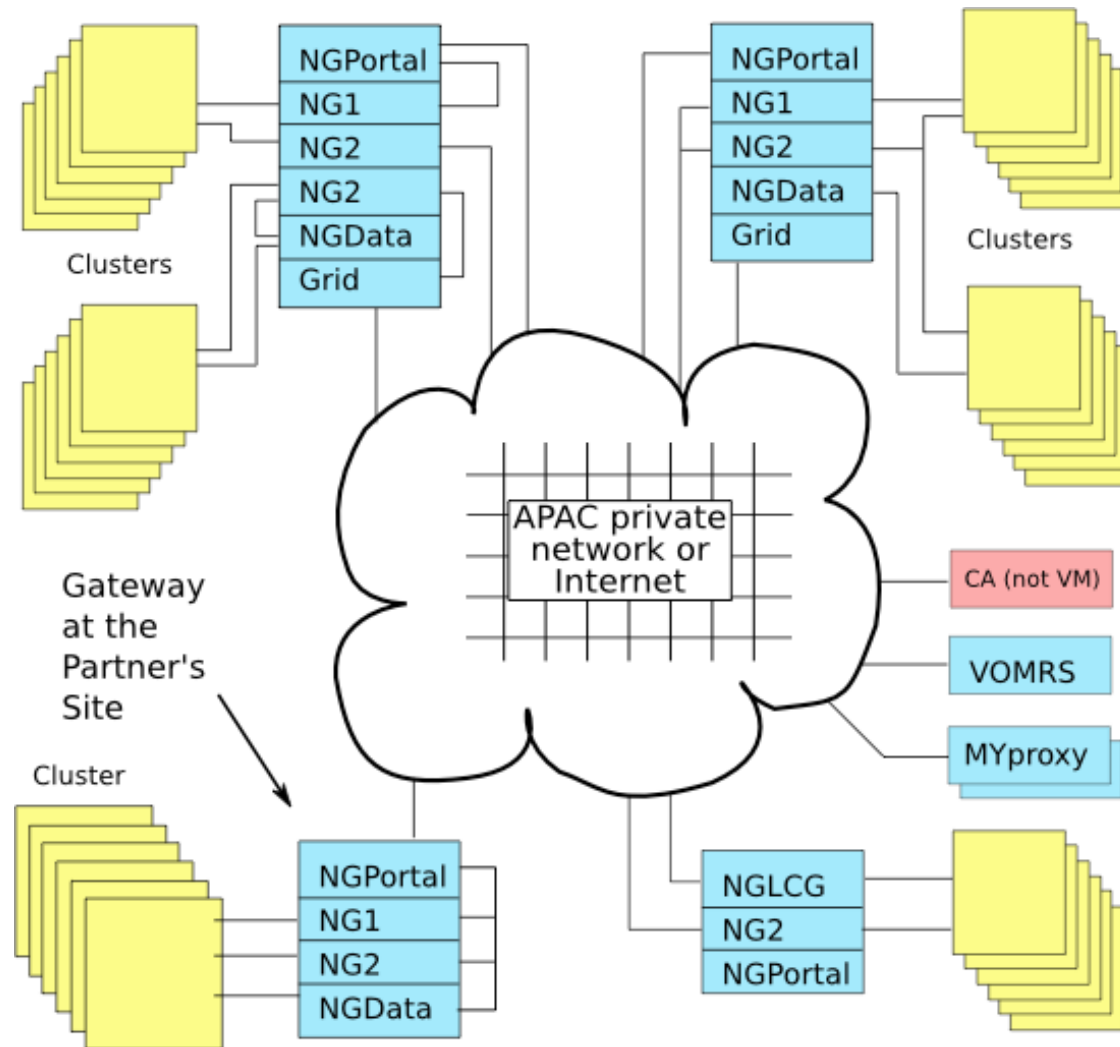
A single

Grid gateway machine hosting

virtual machines providing different middleware stacks

- CentOS, dual 2.8GHz Xeon, 4GB RAM, mirrored 300GB SCSI
- Xen virtualization (Initially v2, now v3)
chosen due to cost and license restrictions of VMWare

Architectural Overview



The road to deployment I



Initially planned for one group to build VM images for deployment at partner facilities.

Issues:

- GT4 knowing the host name at build time
- local admin's not happy with the 'black box' approach
- local admin's wanting to understand the process
- number of local changes after deploying the 'black box' proving to be unmanageable

Solution:

- RPM based installation – manageable
- updates coordinated between sites
- easier to support multiple flavours of middleware and deploy new ones

The road to deployment II



- using a common **VM simplifies** grid middleware **build and deployment**, and easier roll out to partner sites
- makes **grid gateway** model much **easier to implement**
- **modification and testing simplified** - can duplicate a VM in minutes and try new ideas or software
- limits impact of restart / rebuild / crash of grid middleware on the rest of the grid and HPC infrastructure
- savings in cost, power, floor space

The road to deployment III



Xen is new technology

- High I/O load can crash the network stack on a VM
 - much less of a problem in Xen v3 than v2 – but lower I/O throughput
- Xen (< v3.0.3) did not dynamically switch processes to least loaded CPUs
 - needed careful static allocation of CPUs to VMs
- Hypervisor and domain0 kernel recompilations to access full 4GB RAM in 32-bit mode
 - memory split amongst VMs is a limitation
- Rebuilt glibc to enable Native POSIX Thread Library

Current / planned infrastructure



NG2

- CentOS 4.4, VDT 1.6.1a with Globus 4.03, Java 1.5.0.9, Prima Auth Module 0.5

NGDATA

- CentOS 4.4, Globus 4.03 + GridFTP server, GSISSH, kernel configured for high performance TCP transfers, data transfer testing tools


NGPORTAL

- CentOS 4.4, VDT 1.6.1a with Globus 4.03, Java 1.5.0.9, Prima Auth Module 0.5, Tomcat 5.5.20, Ant 1.6.5, Gridsphere 2.2.7, Gridportlets 1.3.2
-

NGGLITE

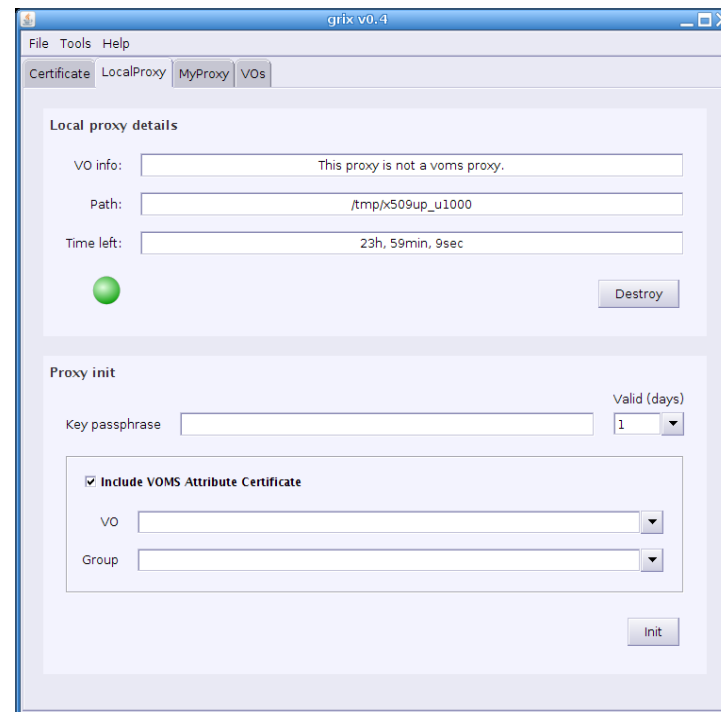
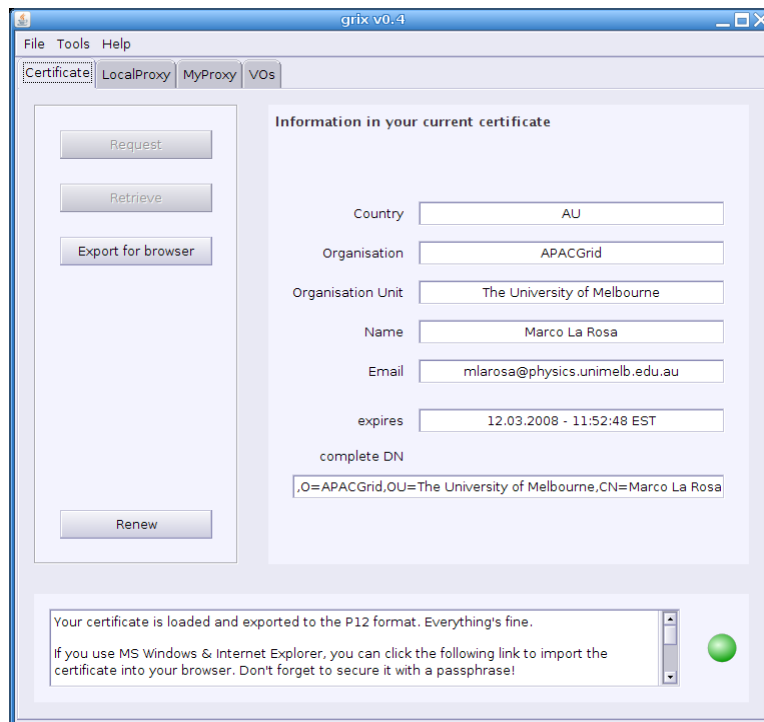
- gLite: some version

The users



- Small early adopter communities – [small VO's](http://grix.vpac.org)
- VOMS / VOMRS mapping to the “merit allocation scheme” (MAS) is still being understood

GRIX : <http://grix.vpac.org>
“help users handle authentication related tasks
in a Grid environment”



User profile: Geoscience



“The [Solid Earth and Environment Grid](#) community ... bring together people in the earth, environmental and computing sciences ... transparent access to data and knowledge about the earth ... enhance our ability to explore for and manage our natural and mineral resources”

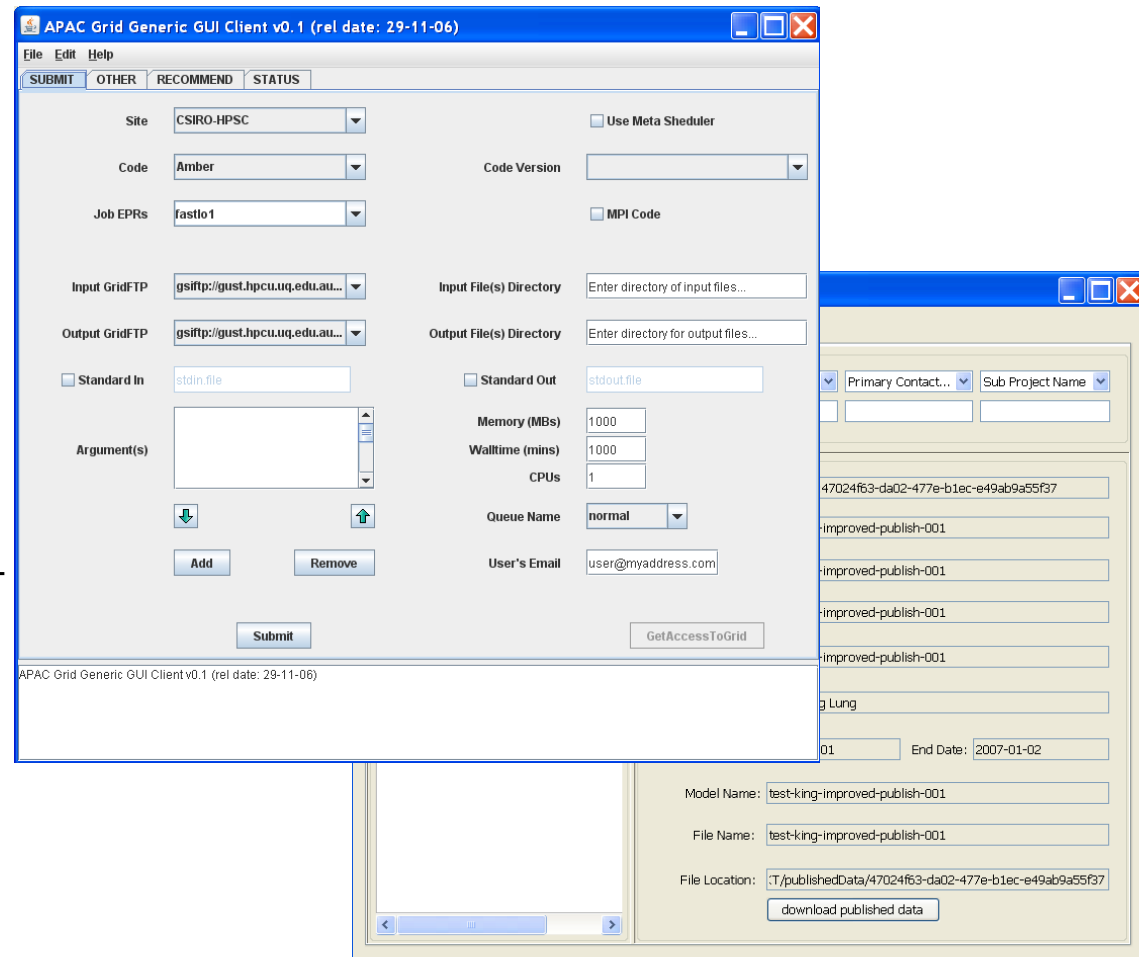
<https://www.seegrid.csiro.au/twiki/bin/view/Main/AboutSEEGrid>

Goals:

- Intuitive and simple to use client applications

Result:

- MDS for Grid resource information
- ebXML dataset registration
- SRB for storage
- Intelligent services
- Data archiving tool
- Generic Grid client



User profile: Earth Systems Science

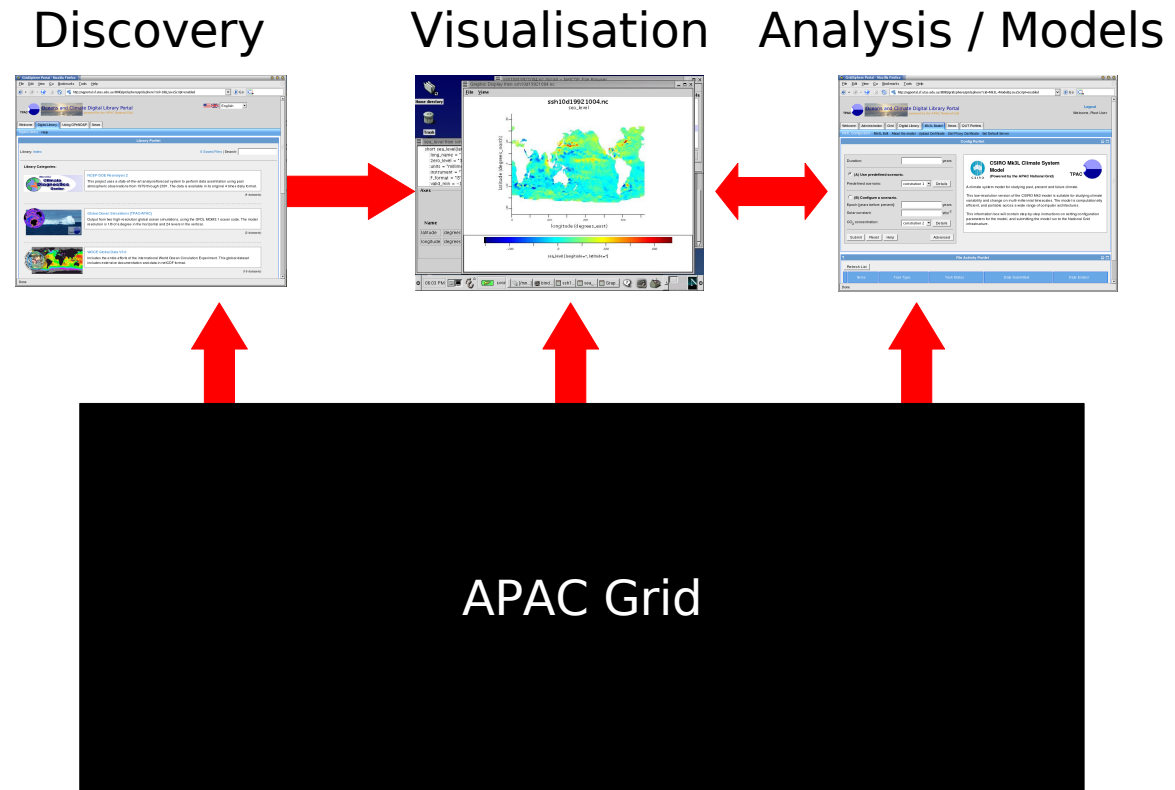


Oceans and Climate Digital Library

<http://digitallibrary.tpac.org.au>

- Climate Diagnosis Centre
- ocean modelling results
- International World Ocean Circulation Experiment (WOCE)
- Australian Antarctic datasets

Uses OPeNDAP framework



User profile: High Energy Physics



users that like getting their hands dirty

2005:

- developed a Grid workflow for Belle MC which used an LCG-2.6 resource and vanilla Globus-2 resources
- home built meta-scheduler GQSched
(<http://epp.ph.unimelb.edu.au/EPPGrid/SoftwareGQSched>)

2006:

- development of expertise in the deployment and usage of the LCG / gLite middleware
- deployment and operation of a pilot Tier 2 facility

39 CPU's, 12TB disk storage

CE, SE, Mon, BDII

various other node types for testing / evaluation

Program goals



The APAC Grid program aims to remove the barrier of access to new users and communities who can benefit from a collaborative research infrastructure.

Advanced users and communities are able to use the infrastructure in the manner which is most suitable for them.

Concluding remarks



- the APAC Grid has reached a critical level of maturity and stability such that early adopters are now able to use the infrastructure to do their research
- development of the higher level interfaces to hide the raw interface to the infrastructure is progressing smoothly
- other (non-early adopter) communities are now being introduced to what they can do with the available resource

The end



Thankyou

Marco La Rosa
mlarosa@physics.unimelb.edu.au

and

'See you at APAC07'

October 8 – 12

Perth, Australia

www.apac.edu.au/apac07