

# XCG: THE VIRGINIA GRID

Karolina Sarnowska  
University of Virginia - Alliance for Computational Science & Engineering

## Outline

- 2
- Grids
- The XCG
- Genesis II
- Use Cases
- Getting Involved

## 3 Grids

What are Grids and Why Researchers Care?

## 4 What are Grids?

- Collections of interconnected resources
- Distributed, heterogeneous and fault-prone
- Administered by different organizations
- Managed via some infrastructure

Examples

- TeraGrid
- Large Hadron Collider Computing Grid
- Northwest Indiana Computation Grid

## 5 Why Grids?

- Computation
  - Faster high-throughput computing
  - More parallel computing
- Data sharing
  - Easier secure access to data
  - Collaboration

## 6 XCG

Current State and Future Plans

### XCG: The Cross Campus Grid

7

- Integrates computing and data resources across Virginia campuses
- Created using Genesis II grid platform


Goals

- ▣ Provide science & engineering research resources
- ▣ Provide secure data sharing
- ▣ Serve as test-bed for Grid standards

### XCG: Current Participants

8

- University of Virginia
- Virginia Tech

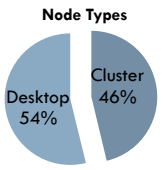


### XCG: Potential Configuration

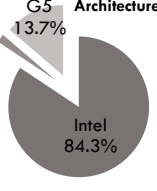
9

100 Teraflops: 13,690 Cores in 4,687 Nodes

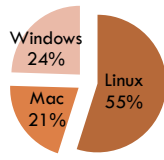
**Node Types**



**Architectures**



**Operating Systems**



### XCG: Current State

10

20% potential cores online

- ▣ University of Virginia ~2,000 cores
- ▣ Virginia Tech ~100 cores

## 11 Genesis II

Building Blocks of the XCG


### 12 Genesis II: The XCG Infrastructure

- Compute & data grid middleware
- Focused on user experience
- Standards based
- Open source
- Secure

## Genesis II: Easy-to-Use

13

- Based on familiar abstraction of a file system



## Genesis II: Easy-to-Use

14

- Everything is a file or directory: **user directories**

```
vcgr:$>ls
/
bes-containers
containers
etc
groups
home
queues
schedulers
uninitialized-containers
users
without-uva-ldap
```

```
vcgr:$>cd home
vcgr:$>ls
home:
CapeKid
admin
bkw1a
burke
cjp8m
crouch
deviney
dgm4d
dk4p
fad5e
franklin
grimshaw
```

## Genesis II: Easy-to-Use

15

- Everything is a file or directory: **queues**

```
vcgr:$>cd queues
vcgr:$>ls
queues:
admin-test-queue
centurion-test-queue
linux-queue
windows-R-queue
windows-short-queue

vcgr:$>cd linux-queue
vcgr:$>ls
linux-queue:
summary
resources
jobs
submission-point
resource-slots
```

```
vcgr:$>cd resources
vcgr:$>ls
resources:
titus
sunfire-pbs
generals-pbs
sulla
camillus-pbs
radio-pbs

vcgr:$>cd titus
vcgr:$>ls
titus:
Gnomad.104
Gnomad.105
Frank's SAS run
reza.alg 21
reza.alg 22
```

## Genesis II: Easy-to-Use

16

- Everything is a file or directory: **host nodes**

```
vcgr:$>cd containers
vcgr:$>ls
containers:
BootstrapContainer
UVA
VT
camillus
centurion001.cs.virginia.edu
pompey.cs.virginia.edu
romulus.cs.virginia.edu
sulla
titus
trajan.cs.virginia.edu

vcgr:$>ls titus/Services
ExportedFilePortType
JNDIAuthPortType
ApplicationDescriptionPortType
BESActivityPortType
EnhancedRMSPortType
XS99AuthPortType
RExportResolverFactoryPortType
IteratorPortType
RandomByteIOPortType
LightweightExportPortType
RExportDirPortType
ExportedDirPortType
QueuePortType
ApplicationDeployerPortType
ExportedRootPortType
GeniBESPortType
RExportResolverPortType
StreamableByteIOPortType
TTYPortType
```

## Mounting the Grid into **Your** File System

17

- See the Grid on local machine as just another directory
  - Genesis II FUSE for Linux
  - G-ICING for Windows

"By default, the user should not have to think."

## Mounting the Grid via FUSE in Linux

18

- Choose local directory
 

```
gbg@titus:~/fuse$ ls
gbg@titus:~/fuse$ pwd
~/uf8/gbg/fuseg
```
- Mount grid into directory
 

```
vcgr:$>fuse --mount /uf8/gbg/fuseg
```
- Access grid via local file system
 

```
gbg@titus:~/fuse$ ls
bes-containers  etc      home    schedulers  users
containers      groups  queues  uninitialized-containers  without-uva-ldap
gbg@titus:~/fuse$
```

19 Use Cases

How to Use the XCG for Computation and Data Sharing

Use Case 1: Computation

- Running large quantities of jobs
- Running same job with many different parameters
- Running MPI jobs

Starting Computations: Describe Job

Create Job Submission Description Language file

- XML based
- Widely adopted standard
- Specify executable, arguments, data staging, etc

```

<JobDefinition> <JobDescription>
<JobIdentification> <JobName>SampleJob</JobName> </JobIdentification>
<Application><POSIXApplication> <Executable>myProgram.exe</Executable>
<Output>stdout</Output> <Error>stderr</Error> <Argument>1</Argument>
</POSIXApplication> </Application>
<DataStaging>
<FileName>input.txt</FileName>
<CreationFlag>overwrite</CreationFlag>
<DeleteOnTermination>true</DeleteOnTermination>
<Source> <URI>rns:/home/feester/sampleJob/input.txt</URI> </Source>
</DataStaging> </JobDescription> </JobDefinition>
    
```

Starting Computations: Describe Job

Automatic generation via GUI tool

Starting Computations: Describe Job

Starting Computations: Describe Job

## Starting Computations: Manage Data

25

- Input files are staged in
- Output files are staged out
- Choose data staging location
  - ▣ Grid space
  - ▣ Local machine

## Starting Computation: Submit Job

26

- Copy to submission point

```
vcgr:$>cp job.xml /queues/linux-queue/submission-point
```

- Qsub to queue

```
vcgr:$>qsub /queues/linux-queue job.xml
```

- Run on execution service

```
vcgr:$>run --jsdl=job.xml /bes-containers/titus
```

## Use Case 2: Data Sharing


27

- Share your files with collaborators
- Access collaborators files
- Read, modify, or view data directly
- Secure access

## Setting Up Data Sharing

28


- Users choose local directories to share on grid
- Access to shared data governed by access-control lists set by owner



## Setting Up Data Sharing

29

- Shared directories can be manipulated locally and on grid



```
vcgr:$>cd /home/tester/MyDocs
vcgr:$>ls
MyDocs:
desktop.ini
SampleJob
Visual Studio 2005
Visual Studio 2008
```

30

## Getting Involved

How to Connect to and Join the XCG

## Connecting to the XCG


31

- Contact UVACSE ([uva-cse@virginia.edu](mailto:uva-cse@virginia.edu))
- Install Genesis II client
  - ▣ Online download
  - ▣ Requires Java 6
  - ▣ Automatically connects to XCG
- Create XCG account
  - ▣ Online request
  - ▣ No special proposal required

## Becoming Part of the XCG

32

- Create unified submission queue to clusters
- Take advantage of idle campus resource cycles
- Personalize resource access policy



## Questions?

33

Contact UVACSE for training and assistance

[uva-cse@virginia.edu](mailto:uva-cse@virginia.edu)

<http://uvacse.virginia.edu>