

# GridWay Overview

John-Paul Robinson

[jpr@uab.edu](mailto:jpr@uab.edu)

University of Alabama at Birmingham

SURAgriid

All-Hands Meeting

Washington, D.C.

March 15, 2007

# What is GridWay?

- Gridway is a Distributed Resource Manager (DRM)
  - meta-scheduler schedules jobs across multiple clusters in a grid
- Clusters support Local Resource Managers (LRM)
  - cluster scheduler schedules jobs across multiple compute nodes in a cluster

# GridWay Features

- “GridWay stands out from other metascheduling systems because it has been specifically designed to work on top of Globus services”
- Java & C application bindings
- Shell level services for familiar job control
- Implements Distributed Resource Manager Application API (DRMAA)
- Job Submission Description Language (JDSL)

# GridWay Requirements

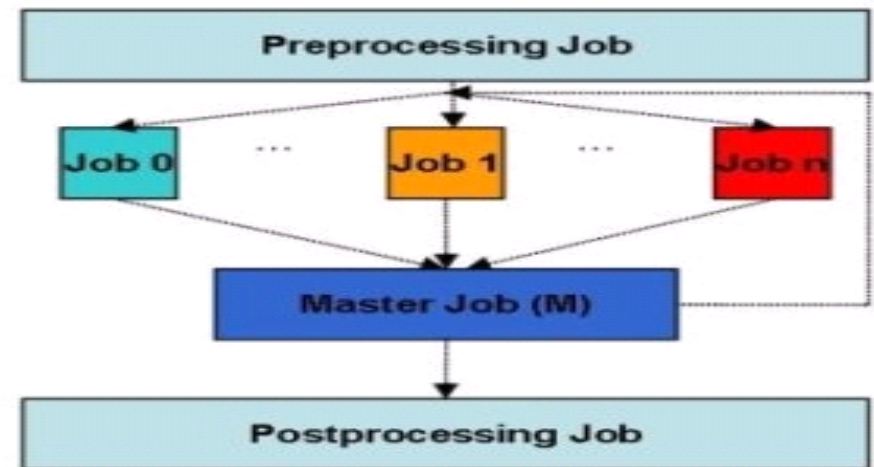

- On Cluster Side
  - Supports both Globus web and pre-web services
  - Simple tests to validate GTK configuration
    - [Verifying Globus Install](#)
- On Submit Side
  - Globus (C and Java)
  - BerkelyDB for accounting (optional)
  - GridWay metascheduler
  - Other standard build dependencies

# How Does GridWay Work?

1. User creates a job file
2. User submits job to GridWay
3. GridWay selects a resource
4. GridWay submits job to resource
5. GridWay monitors and reports job status
6. User receives job output files locally

# What Kind of Jobs are Good?

- Embarrassingly Distributed
  - Many Tasks / Little Communication
  - Lots of Apps fall in this category
- Master-Worker
  - One Controller / Many Parallel Tasks
  - Good for Grid-based workflows
- Custom Schedulers



Images Courtesy of GridWay

# Resource Discovery

- User discovers resources with `gwhost`
- Available resources are list in GridWay configuration
- GridWay discovers current state with MDS

```
[afgane@everest00 drmaa]$ gwhost
```

HID	NICE	OS	ARCH	MHZ	%CPU	MEM(F/T)	DISK(F/T)	N(U/F/T)	LRMS	HOSTNAME
0	500	Linux2.6.9-22.E	x86_6	1005	200	1184/2010	60130/76652	0/4/4	PBS	mileva.hpc.odu.edu
1	500	Linux2.4.21-32.	x86_6	3200	200	2702/3848	57374/63479	0/2/2	Fork	olympus.cis.uab.edu
2	500	NULLNULL	NULL	0	0	0/0	0/0	0/0/2	Fork	wave.cis.uab.edu
3	500	Linux2.4.20-8sm	x86	2657	395	34/2015	10825/69704	0/4/4	Fork	titanic.hpcl.cis.uab.edu
4	500	Linux2.4.21-20.	x86_6	1603	200	11/1977	70759/79604	0/18/29	SGE	everest.cis.uab.edu

```
[afgane@everest00 drmaa]$
```

# Running a Simple Job

- Create a job script

```
[afgane@everest00 drmaa]$ cat testjob
EXECUTABLE=/bin/uname
ARGUMENTS= -a
STDOUT_FILE=stdout.${JOB_ID}
STDERR_FILE=stderr.${JOB_ID}
RESCHEDULE_ON_FAILURE=no
NUMBER_OF_RETRIES=3
[afgane@everest00 drmaa]$ █
```

# Submit Job and Monitor

- gws submit -t testjob
- gwps

```
[afgane@everest00 drmaa]$ gws submit -t testjob
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  pend ---- 000 08:02:35 --:--:-- 0:00:00 0:00:00 --  testjob  --
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  wrap pend 000 08:02:35 --:--:-- 0:00:04 0:00:02 --  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  wrap pend 000 08:02:35 --:--:-- 0:00:07 0:00:02 --  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  wrap pend 000 08:02:35 --:~:~:~ 0:00:11 0:00:02 --  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  wrap pend 000 08:02:35 --:~:~:~ 0:00:19 0:00:02 --  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  wrap pend 000 08:02:35 --:~:~:~ 0:00:28 0:00:02 --  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  wrap pend 000 08:02:35 --:~:~:~ 0:00:36 0:00:02 --  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  epil ---- 000 08:02:35 --:~:~:~ 0:00:41 0:00:08 --  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ gwps
USER          JID AID TID DM  EM  RWS START  END    EXEC  XFER  EXIT TEMPLATE  HOST
afgane        63  --  --  done ---- 000 08:02:35 08:03:41 0:00:41 0:00:12 0  testjob  mileva.hpc.odu.edu/PBS
[afgane@everest00 drmaa]$ █
```

# User Receives Results

- Standard Output and Error are directed to local files with “jobid” extension by default

```
[afgane@everest00 drmaa]$ cat stdout.63
Linux mileva-0-2.local 2.6.9-22.ELsmp #1 SMP Sat Oct 8 21:32:36 BST 2005 x86_64 x86_64 x86
_64 GNU/Linux
[afgane@everest00 drmaa]$ █
```

# More Advanced Example

- dynamicBLAST
  - Example of Master/Worker model
  - Grid Workflow Management
  - Parallel BLAST execution
  - Coordinate job on many grid nodes

# Summary of Output

```
[afgane@everest00 drmaa]$ ./dblast51.sh
[2007/03/15 08:27:08] Started Dynamic BLAST.
[2007/03/15 08:27:08] Using job properties file: props.properties
...
[2007/03/15 08:27:08] Analisis job successfully submitted, job ID: 64
[2007/03/15 08:28:01] Job 64 finished regularly with exit status 0
...
[2007/03/15 08:28:01] Total number of fragments to be created: 10
[2007/03/15 08:28:01] Fragmentation job successfully submitted, ID: 65
[2007/03/15 08:29:03] Job 65 finished regularly with exit status 0
...
[2007/03/15 08:29:03] Gathering resource information...
[2007/03/15 08:29:29] Successfully accessed current resource data.
availableResourceList=[[mileva.hpc.odu.edu, 4], [titanic.hpcl.cis.uab.edu, 4],
[olympus.cis.uab.edu, 2], [everest.cis.uab.edu, 18]]
...
[2007/03/15 08:29:29] Remaining num jobs (frags) to submit: 10
[2007/03/15 08:29:29] Creating job(s) for mileva.hpc.odu.edu with 4 available node(s).
[2007/03/15 08:29:29] Creating job(s) for titanic.hpcl.cis.uab.edu with 4 available node(s).
[2007/03/15 08:29:29] Creating job(s) for olympus.cis.uab.edu with 2 available node(s).
  exit_time=08:30:37
Job usage (for jobs [66, 67, 68, 69]) on titanic.hpcl.cis.uab.edu:
...
*****
All fragments have been processed.
*****
...

```

# Information and Acknowledgements

- GridWay has a very good web site with good documentation
  - <http://www.gridway.org>
  - Install Guide
  - Applications Guide
  - User Guide <-- start here
- Thanks to Enis Afgan ([afgane@uab.edu](mailto:afgane@uab.edu)) and Puri Bangalore ([puri@uab.edu](mailto:puri@uab.edu)) for support in getting up to speed with GridWay and dynamicBLAST