



GridWay interoperability through BES

EGI Technical Forum 2012
Prague, Czech Republic
September 17-21, 2012

Dr Ismael Marín Carrión¹, Dr Eduardo Huedo¹,
Dr Stephen Crouch², Dr Ignacio M. Llorente¹



¹ Universidad Complutense de Madrid
² University of Southampton

UNIVERSITY OF
Southampton

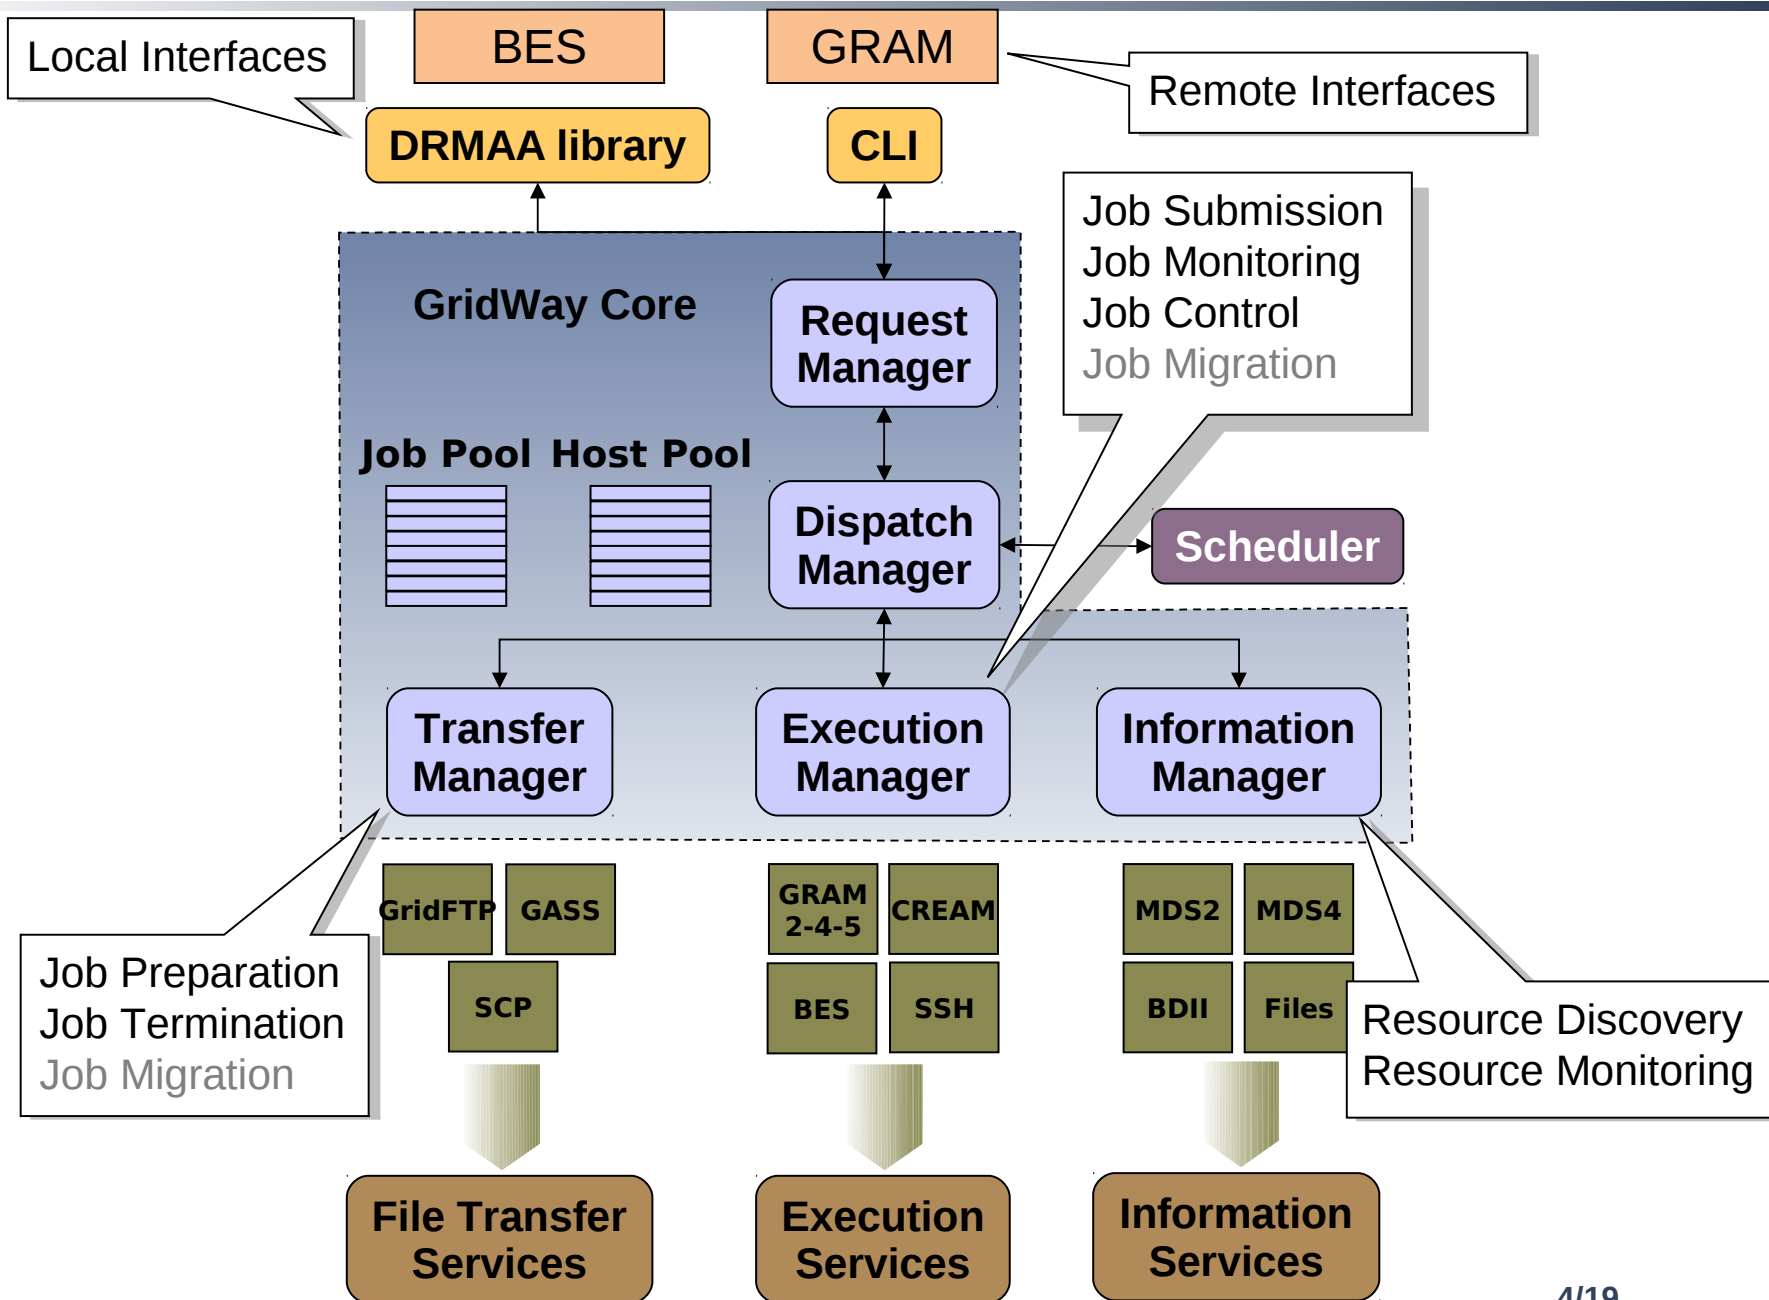


Initiative for Globus in Europe



- 1. Introduction**
2. Initiative for Globus in Europe
3. BES support
4. Demo

- The GridWay metascheduler enables large-scale, reliable and efficient sharing of computing resources over different grid middleware, providing a single point of access for them
- GridWay provides a LRM-like CLI for submitting, monitoring, synchronizing and controlling jobs
- GridWay implements the OGF standard DRMAA API, assuring compatibility of applications with LRM systems that implement the standard
- GridWay provides BES- and GRAM-compliant interfaces, enabling the submission and monitoring of jobs through a standard interface



1. Introduction
- 2. Initiative for Globus in Europe**
3. BES Support
4. Demo



INITIATIVE FOR GLOBUS IN EUROPE

- The IGE project is a EU FP7 project to coordinate European Globus activities
- Main objectives:
 - Support the European computing infrastructures and their users
 - Provide a central point of contact in Europe for Globus
 - Strengthen the influence of European developers and users on the development of Globus
- IGE serves as a comprehensive service provider for the European e-infrastructures regarding the development, customization, provisioning, support, and maintenance of components of the Globus Toolkit

IGE Testbed

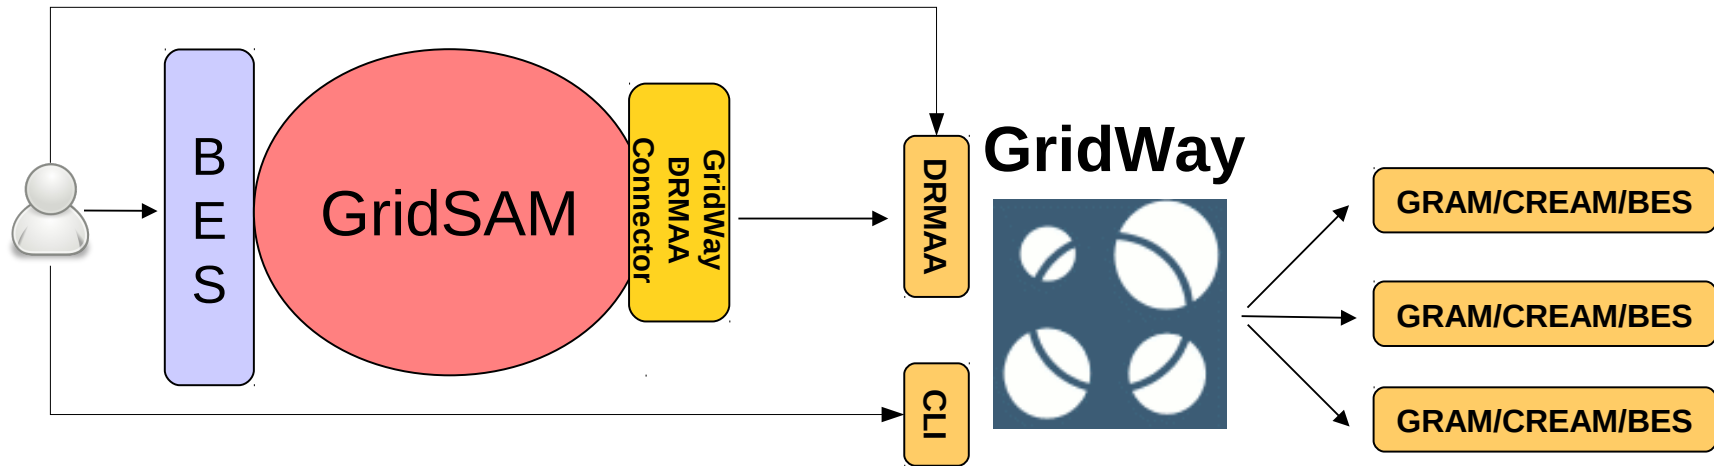
- GridWay instance provided by:
 - Universidad Complutense de Madrid (UCM), Spain
- Globus 5 resources provided by:
 - Technische Universität Dortmund (TUDO), Germany
 - Poznan Supercomputing and Networking Center (PSNC), Poland
 - Leibniz-Rechenzentrum, Bayerische Akademie der Wissenschaften (BADW-LRZ), Germany
 - Universitatea Technica Cluj-Napoca (UTCN), Romania
 - Stichting voor Fundamenteel Onderzoek der Materie – Institute for Subatomic Physics (FOM-NIKHEF), Netherlands
 - University of Edinburgh – Edinburgh Parallel Computing Centre (UEDIN-EPCC), United Kingdom
- Other resources:
 - University of Southampton (SOTON), United Kingdom



1. Introduction
2. Initiative for Globus in Europe
- 3. BES Support**
4. Demo

- The Basic Execution Service (BES) specification defines Web Services interfaces for creating, monitoring, and controlling computational entities
- JSDL is used to describe jobs
- GridWay provides support for the BES standard both as a client and as a server
 - BES Execution driver
 - BES Interface for GridWay
- BES driver provides an abstraction layer that enables users to submit jobs to BES interfaces, and control and monitor the execution of jobs
- BES interface for GridWay enables the remote access to GridWay's metascheduling capabilities through a standard interface, allowing users to access services provided by different grid middlewares
 - Based on GridSAM
 - Communication is based on the DRMAA API

Interoperability Scenario



- Interfaces: BES Interface / GridWay CLI
- GridSAM: Job submission interface that provides a BES interface. It not provides metascheduling capabilities.
- GridWay connector: It enables the interoperation of GridSAM and GridWay. It allows end-users to submit jobs to GridWay through a BES interface.
- GridWay metascheduler
- GridWay drivers that interfaces with Grid services: GRAM, BES, CREAM, GridFTP, BDII, etc.

1. Introduction
2. Initiative for Globus in Europe
3. BES Support
- 4. Demo**

BES Driver

- Configuration

```
$cat /usr/etc/gwd.conf  
...  
IM_MAD = static_2:gw_im_mad_static:-l etc/BES_hosts.list:dummy:bes  
EM_MAD = bes:GW_em_mad_bes::jsdl  
TM_MAD = dummy:gw_tm_mad_dummy:-u gsiftp\://gridway.fdi.ucm.es
```

BES Driver

- Create your proxy
- Use *gwhost* command to see available resources:

HID	PRIO	OS	ARCH	MHZ	%CPU	MEM(F/T)	DISK(F/T)	N(U/F/T)	LRMS	HOSTNAME
0	1	Linux2.6.32.27-	x86_64	2533	400	2007/2007	71G/71G	0/2/2	jobmanager-sge	gt5-ige.drg.lrz.de
1	1	Linux2.6.18-238	x86_64	1995	600	2007/2007	927G/927G	0/44/44	jobmanager-pbs	udo-gt01.grid.tu-dortmund.de
2	1	Linux2.6.18-194	x86_64	2993	800	2011/2011	3670M/3670M	0/8/8	jobmanager-fork	ve.nikhef.nl
3	1	Linux2.6.18-194	x86_64	1600	100	1024/1024	40G/40G	0/1/1	jobmanager-fork	gt1.epcc.ed.ac.uk
4	1	Linux2.6.18-194	x86_64	2328	400	512/512	95G/95G	0/1/1	jobmanager-pbs	gt01.ige.psnc.pl
5	1	Linux2.6.32-36	x86_64	2000	400	2009/2009	4250G/4250G	0/4/4	jobmanager-fork	gt-ige.utcluj.ro
6	1	Linux2.6.35-22	x86_64	1995	200	2010/2010	8683M/8683M	0/2/2	bes-pbs	147.96.25.29

- and get more detailed information specifying a Host ID:

```
$ gwhost 0
```

HID	PRIO	OS	ARCH	MHZ	%CPU	MEM(F/T)	DISK(F/T)	N(U/F/T)	LRMS	HOSTNAME
0	1	Linux2.6.32.27-	x86_6	2533	400	2007/2007	71G/71G	0/2/2	jobmanager-sge	gt5-ige.drg.lrz.de

QUEUE	NAME	SL(F/T)	WALLT	CPUT	COUNT	MAXR	MAXQ	STATUS	DISPATCH	PRIORITY
all.q		2/2	0	0	0	0	0	1	Batch	Inte

- Describe the job. Create a job template:

```
EXECUTABLE = primes.exe
ARGUMENTS = 0 1000
STDOUT_FILE = stdout_file.${JOB_ID}
STDERR_FILE = stderr_file.${JOB_ID}
RANK = CPU_MHZ
REQUIREMENTS = LRMS_NAME= "bes-pbs"
```

- Use *gwsbmit* command to submit the job:

```
$ gwsbmit -t primes.jt
```

BES Driver

- Check the resources that match job requirements with `gwhost -m <job_id>`:

```
$ gwhost -m 0
HID QNAME      RANK  PRIO  SLOTS HOSTNAME
6   pbs         1995  1     2     147.96.25.29
```

- Follow the evolution of the job with `gwps` command:

```
$ gwps
USER      JID DM   EM   START      END        EXEC       XFER      EXIT NAME      HOST
imarin:0  0   done  ---- 22:48:32 22:52:18 0:03:40 0:00:00 0   primes.jt 147.96.25.29/bes-pbs
imarin:0  1   done  ---- 22:53:10 23:00:48 0:07:25 0:00:00 0   primes.jt 147.96.25.29/bes-pbs
imarin:0  2   done  ---- 09:44:47 09:45:04 0:00:02 0:00:11 0   jt        gt1.epcc.ed.ac.uk/jobma
imarin:0  3   done  ---- 11:56:16 11:56:47 0:00:02 0:00:23 0   jt        gt-ige.utcluj.ro/jobman
....
imarin:0  26  wrap  pend 12:29:06 ---:---:-- 0:00:07 0:00:00 --  primes.jt 147.96.25.29/bes-pbs
```

- HINT: Use `gwps -c <seconds>` for continuous output.
- See the job history with `gwhistory` command:

```
$ gwhistory 0
HID START      END        PROLOG  WRAPPER EPILOG  MIGR      REASON  QUEUE  HOST
6   22:48:38 22:52:18 0:00:00 0:03:40 0:00:00 0:00:00  ----  pbs    147.96.25.29/bes-pbs
```

- Check the results!

BES Interface

- Delegate your credentials to a MyProxy server:

```
$ myproxy-init -s gridway.fdi.ucm.es
```

- Create a JSDL file for describing the job. Specify how to retrieve your credentials from a MyProxy server, requirements, ...

```
<jSDL:JobDefinition xmlns:jSDL="http://schemas.ggf.org/jSDL/2005/11/jSDL"
  xmlns:jSDL-posix="http://schemas.ggf.org/jSDL/2005/11/jSDL-posix"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <jSDL:JobDescription>
    ...
    <jSDL:Application>
      <jSDL-posix:POSIXApplication>
        <jSDL-posix:Executable>primes.exe</jSDL-posix:Executable>
        <jSDL-posix:Argument>0 1000</jSDL-posix:Argument>
        <jSDL-posix:Output>primes_stdout</jSDL-posix:Output>
      </jSDL-posix:POSIXApplication>
      <MyProxy xmlns="urn:gridsam:myproxy">
        <ProxyServer>gridway.fdi.ucm.es</ProxyServer>
        <ProxyServerPort>7512</ProxyServerPort>
        <ProxyServerUserName>imarin</ProxyServerUserName>
        <ProxyServerPassPhrase>****</ProxyServerPassPhrase>
        <ProxyServerLifetime>10000</ProxyServerLifetime>
      </MyProxy>
    ...
  </jSDL:Application>
  <jSDL:Resources>
    ...
  </jSDL:Resources>
```

BES Interface

- Describe the job. Specify datastaging items:

```

...
<jsdsl:DataStaging>
  <jsdsl:FileName>primes.exe</jsdsl:FileName>
  ...
  <jsdsl:Source>
    <jsdsl:URI>gsiftp://gridway.fdi.ucm.es//tmp/primes.exe</jsdsl:URI>
  </jsdsl:Source>
</jsdsl:DataStaging>
<jsdsl:DataStaging>
  <jsdsl:FileName>primes_stdout</jsdsl:FileName>
  ...
  <jsdsl:Target>
    <jsdsl:URI>gsiftp://gridway.fdi.ucm.es//tmp/primes_stdout</jsdsl:URI>
  </jsdsl:Target>
</jsdsl:DataStaging>
</job:JobDescription>
<MyProxy xmlns="urn:gridsam:myproxy">
  <ProxyServer>gridway.fdi.ucm.es</ProxyServer>
  <ProxyServerPort>7512</ProxyServerPort>
  <ProxyServerUserName>imarin</ProxyServerUserName>
  <ProxyServerPassPhrase>****</ProxyServerPassPhrase>
  <ProxyServerLifetime>10000</ProxyServerLifetime>
</MyProxy>
</jsddl:JobDefinition>

```


BES Interface

- Use a BES client to submit and monitor your jobs.
- Submit a single job using the GridSAM client:

```
$ ./gridsam.sh BESCreateActivity -s
"https://gridway.fdi.ucm.es:8443/gridsam/services/bes?wsdl" -j bes-prim.es.jsdl >
job_id
```

```
$ cat job_id
<?xml version="1.0" encoding="UTF-8"?>
<EndpointReference xmlns="http://www.w3.org/2005/08/addressing">
...<ID>urn:gridsam:13e0999c3848481f0138484838b80001</ID>...</EndpointReference>
```

- Follow the evolution of your BES jobs:

```
$ watch -n 5 ./gridsam.sh BESGetActivityStatuses -s
"https://gridway.fdi.ucm.es:8443/gridsam/services/bes?wsdl" -file job_id
...
<?xml version="1.0" encoding="UTF-8"?>
<GetActivityStatusesResponse xmlns="http://schemas.ggf.org/bes/2006/08/bes-factory">
...<ActivityStatus state="Running"/>...</GetActivityStatusesResponse>
```

- Check how your BES job is scheduled by GridWay
- Check the results at /tmp/prim.es_stdout

- **The GridWay Project:** <http://gridway.org/>
- **GridWay Development:** <http://dev.gridway.org/>
- **Documentation:** <http://gridway.org/doku.php?id=documentation>
- **Support:** <http://gridway.org/doku.php?id=support>
 - User discussion: gridway-user@globus.org
 - Announcements: gridway-announce@globus.org
- **The IGE Project:** <http://www.ige-project.eu/>
- **GridSAM:** <http://www.omii.ac.uk/wiki/GridSAM>
- **BES Interface for GridWay:** <http://www.gridway.org/doku.php?id=ecosystem:gridwaybes>

**Thank you
for your attention!**