

Grid on Demand

Willem Toorop and Alain van Hoof

{wtoorop,ahoof}@os3.nl

June 30, 2010

Research Question

Can Grid computing be offered as a Cloud service and what type of applications benefit from this?


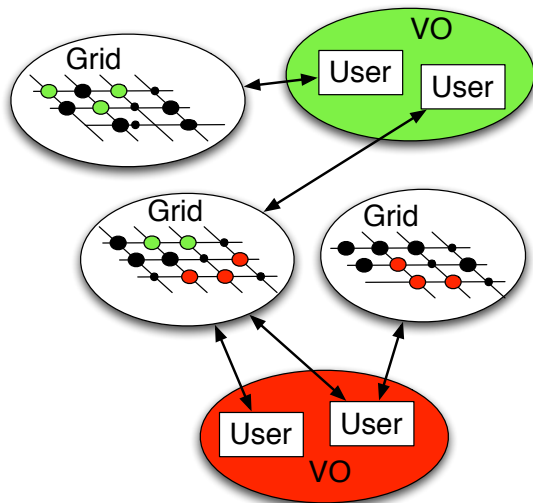
The Grid and The Cloud: Sharing Resources

Sharing Resources to minimizing idle resources

Resource	Capacity	Goods
Processor	Flops	Compute
Hard-disk memory	Bytes	Storage
Network	Bytes/sec	Communication Bandwidth

	Grid	Cloud (IaaS)
Distribution	By Collaboration	By Cloud Provider
Resource Type	Heterogeneous	Homogeneous
Abstraction	Middleware (open standards)	Virtualization
Sharing Guarantee	Security infrastructure	Financial infra

The Grid and The Cloud: Dividing the Resources



A Grid in The Cloud: Urgent Computing

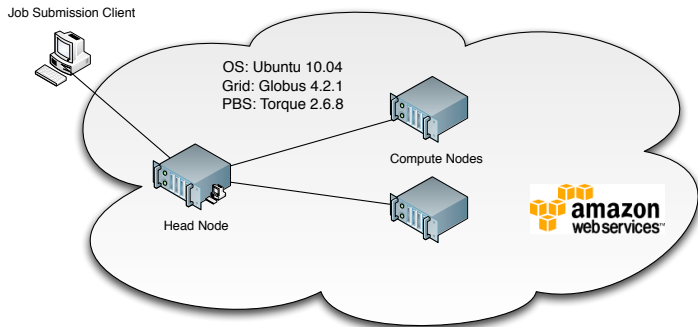


Using the elasticity and scalability of the Cloud and providing a Grid interface, current e-Science (computationally intensive science) applications can use a **Grid on Demand** when there is a high demand for resources for a short period of time.

Example

UrbanFlood (<http://urbanflood.eu>) wants to calculate flooding predictions when actual flooding happens, Grid on Demand can provide this.

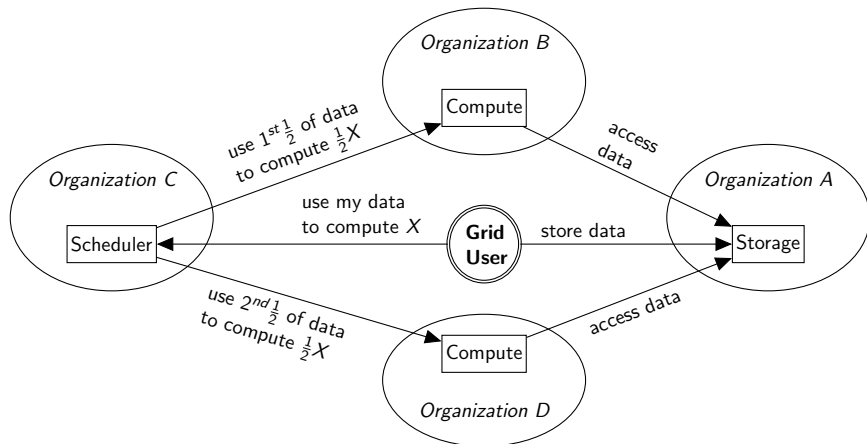
Grid on Demand



Grid Interface: Globus Toolkit 4.2 (open standard, widely used)
Queue Manager (Distribution of jobs to nodes): Torque 2.6.2
Amazon EC2 Cloud Provider (de-facto standard with open API)

What makes the Grid, the Grid

Typical Grid usage example spanning multiple organizations

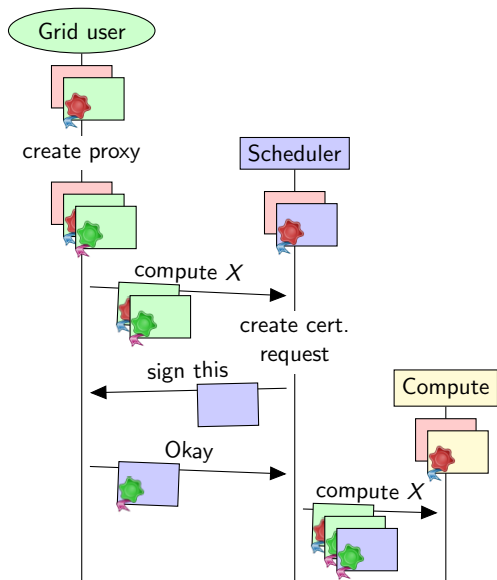



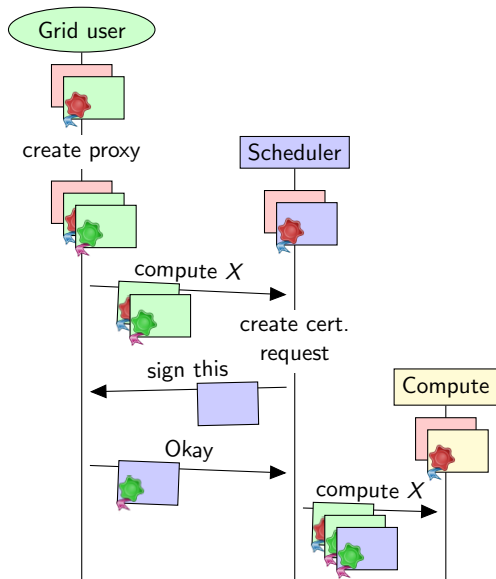
Proxy certificates for delegation

The Grid is a collaboration between multiple organizations.

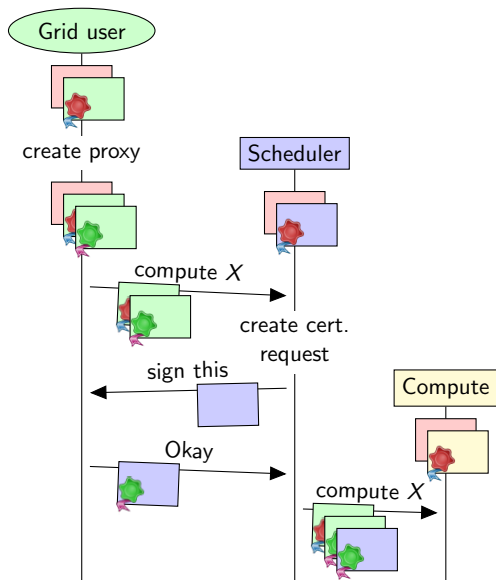
- Grid users need to be authorized to use remote resources
- Those resources should be authorized to perform actions *on behalf of* the Grid user

Proxy certificates enable delegation of authority






Virtual Organizations are based on delegation by proxy certificates.



Proxy certificates allow for dynamic resource allocation,

But the hosts still need a CA signed certificate.

This is a problem when hostname's are assigned dynamically!




Proxy certificates allow for dynamic resource allocation,

But the hosts still need a CA signed certificate.

This is a problem when hostname's are assigned dynamically!

Fixed CA in image?



Proxy certificates allow for dynamic resource allocation,

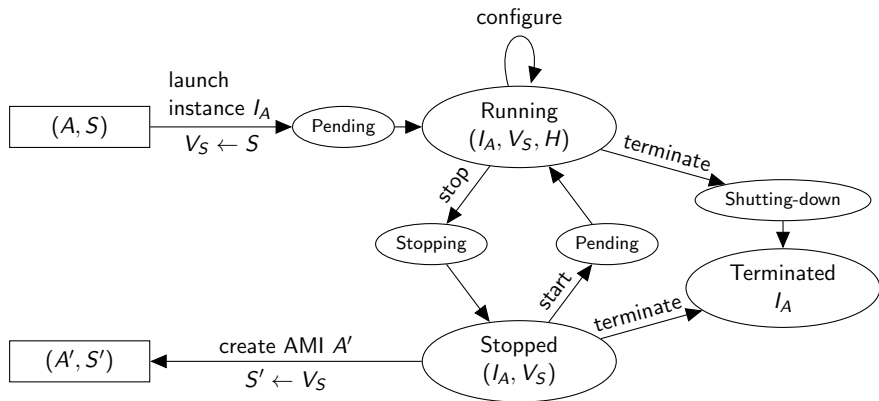
But the hosts still need a CA signed certificate.

This is a problem when hostname's are assigned dynamically!

Fixed CA in image?

Dynamic CA on every Launch?

Easy Image management with Amazon Elastic Block Store

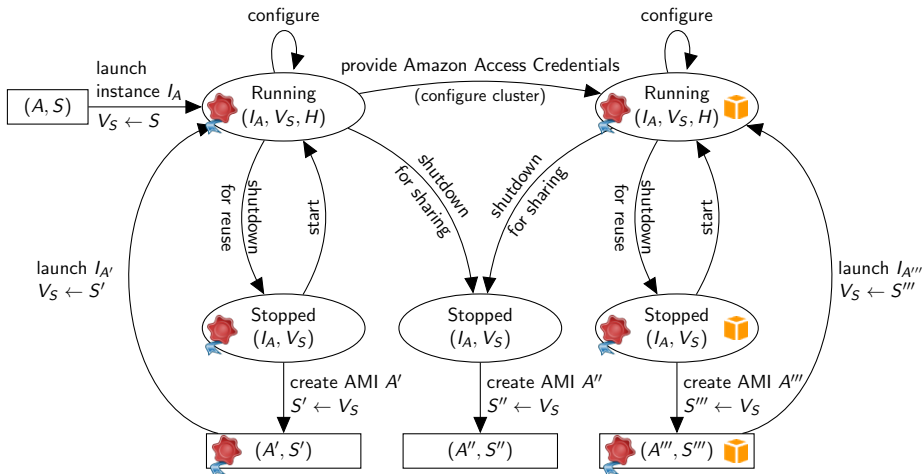


Life cycle of Amazon Machine Images using Elastic Block Storage

A = Amazon Machine Image, S = Snapshot,

I = Instance, V = Volume, H = Hostname

EBS for configuration stages



AWS Management Co... x

https://console.aws.amazon.com/ec2/home?region=eu-west-1#?c=EC2&s=Hc

aws.amazon.com | AWS | Products | Developers | Community | Support | Account | Welcome, Jaap van Ginkel | Settings | Sign out

Amazon S3 | **Amazon EC2** | Amazon Elastic MapReduce | Amazon CloudFront | Amazon RDS

Navigation

Region: EU West

- > EC2 Dashboard
- INSTANCES
 - > Instances
 - > Spot Requests
- IMAGES
 - > AMIs
 - > Bundle Tasks
- ELASTIC BLOCK STORE
 - > Volumes
 - > Snapshots
- NETWORKING & SECURITY
 - > Elastic IPs
 - > Security Groups
 - > Key Pairs
 - > Load Balancers

Amazon EC2 Console Dashboard

Getting Started

To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.

[Launch Instance](#)

Note: Your instances will launch in the EU West (Ireland) region.

Service Health

Current Status	Details
Amazon EC2 (EU - Ireland)	Service is operating normally

[View complete service health details](#)

My Resources

You are using the following Amazon EC2 resources in the EU West (Ireland) region: [Refresh](#)

- 0 Running Instances
- 0 Elastic IPs
- 2 EBS Volumes
- 4 EBS Snapshots
- 2 Key Pairs
- 1 Security Group
- 0 Load Balancers

Related Links

- > Documentation
- > All EC2 Resources
- > Forums

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | [Feedback](#) | [Support](#) | [Privacy Policy](#) | [Terms of Use](#)

An amazon.com company

The screenshot shows the AWS Management Console interface for the 'Request Instances Wizard'. The wizard is in the 'CHOOSE AN AMI' step. The 'Community AMIs' tab is selected, and the search filter is set to 'grid on demand'. A table displays three AMIs, each with a 'Select' button. A mouse cursor is hovering over the 'Select' button for the second AMI.

Request Instances Wizard Cancel X

CHOOSE AN AMI INSTANCE DETAILS CREATE KEY PAIR CONFIGURE FIREWALL REVIEW

Choose an Amazon Machine Image (AMI) from one of the tabbed lists below by clicking its **Select** button.

[Quick Start](#) [My AMIs](#) [Community AMIs](#)

Viewing: All Images 1 to 3 of 3 Items

AMI ID	Root Device	Manifest	Platform	
ami-53133927	ebs	021450451553/Grid on Demand v1.11	Other Linux	Select
ami-57ebc123	ebs	021450451553/Grid on Demand v1.13	Other Linux	Select
ami-f3eac087	ebs	021450451553/Grid on Demand v1.12	Other Linux	Select

Navigation

Region: EU West

Amazon S3 | **Amazon EC2** | Amazon Elastic MapReduce | Amazon CloudFront | Amazon RDS

My Instances

Launch Instance | Instance Actions | Reserved Instances | Show/Hide | Refresh | Help

Viewing: All Instances | All Instance Types

Instance	AMI ID	Root	Type	Status	Lifecycle	Public DNS
<input checked="" type="checkbox"/> i-64d34c13	ami-57ebc123	ebs	m1.large	running	normal	ec2-79-125-48-213.eu-west-1.comput
<input type="checkbox"/> i-ee861499	ami-5313392	ebs	m1.large	sto		
<input type="checkbox"/> i-ba3ca0cd	ami-f3eac087	ebs	m1.large	ter		

1 EC2 Instance selected

EC2 Instance: i-64d34c13

Description | Monitoring

AMI ID: ami-57ebc123 | Zone

Security Groups: default | Type

Status: running | Own

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | Feedback | Support | Privacy Policy | Terms of Use |

https://console.aws.amazon.com/ec2/home?... | An amazon.com company


```
willem@willem-laptop: ~ 84x24
willem@willem-laptop:~$ ssh -i default.pem root@ec2-79-125-48-213.eu-west-1.compute.
amazonaws.com
Warning: Permanently added 'ec2-79-125-48-213.eu-west-1.compute.amazonaws.com,79.125
.48.213' (RSA) to the list of known hosts.
Please login in as griduser to immediatly use the Grid on Demand:

ssh -i default.pem -L 8080:localhost:80 \
    griduser@ec2-79-125-48-213.eu-west-1.compute.amazonaws.com

or login as ubuntu for maintenance:

ssh -i default.pem -L 8080:localhost:80 \
    ubuntu@ec2-79-125-48-213.eu-west-1.compute.amazonaws.com

Then Grid on Demand can be configured and monitored via:

http://localhost:8080/

Connection to ec2-79-125-48-213.eu-west-1.compute.amazonaws.com closed.
willem@willem-laptop:~$
```

```
griduser@ec2-79-125-48-213: ~ 84x24
willem@willem-laptop:~$ ssh -i default.pem root@ec2-79-125-48-213.eu-west-1.compute.
amazonaws.com
Warning: Permanently added 'ec2-79-125-48-213.eu-west-1.compute.amazonaws.com,79.125
.48.213' (RSA) to the list of known hosts.
Please login in as griduser to immediatly use the Grid on Demand:

ssh -i default.pem -L 8080:localhost:80 \
    griduser@ec2-79-125-48-213.eu-west-1.compute.amazonaws.com

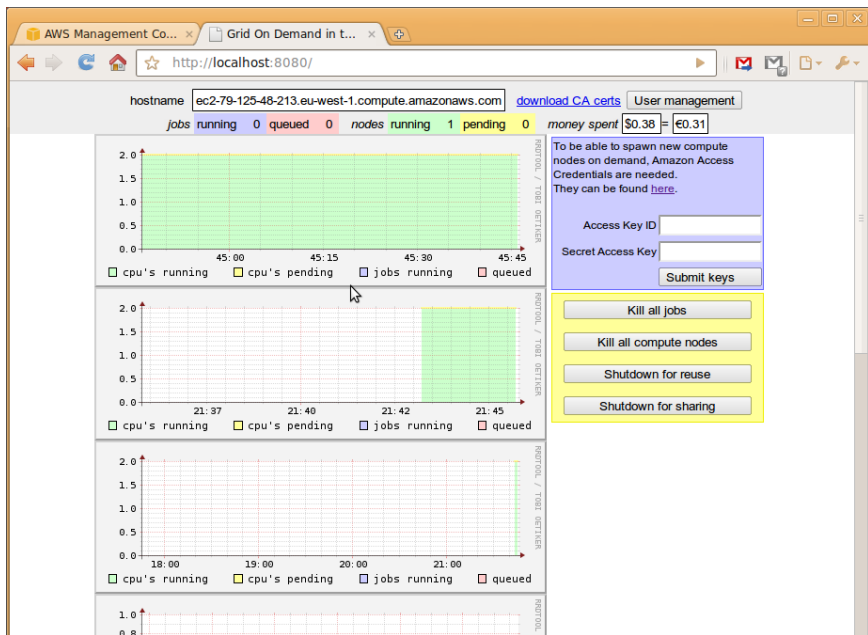
or login as ubuntu for maintenance:

ssh -i default.pem -L 8080:localhost:80 \
    ubuntu@ec2-79-125-48-213.eu-west-1.compute.amazonaws.com

Then Grid on Demand can be configured and monitored via:

http://localhost:8080/

Connection to ec2-79-125-48-213.eu-west-1.compute.amazonaws.com closed.
willem@willem-laptop:~$ ssh -i default.pem -L 8080:localhost:80 \
>    griduser@ec2-79-125-48-213.eu-west-1.compute.amazonaws.com
Last login: Mon Jun 28 20:16:08 2010 from happusxs4all.nl
griduser@ec2-79-125-48-213:~$
```



```
griduser@ec2-79-125-48-213: ~ 84x24
or login as ubuntu for maintenance:

ssh -i default.pem -L 8080:localhost:80 \
  ubuntu@ec2-79-125-48-213.eu-west-1.compute.amazonaws.com

Then Grid on Demand can be configured and monitored via:

http://localhost:8080/

Connection to ec2-79-125-48-213.eu-west-1.compute.amazonaws.com closed.
willem@willem-laptop:~$ ssh -i default.pem -L 8080:localhost:80 \
> griduser@ec2-79-125-48-213.eu-west-1.compute.amazonaws.com
Last login: Mon Jun 28 20:16:08 2010 from happus.xs4all.nl
griduser@ec2-79-125-48-213:~$ grid-proxy-info
subject : /O=Grid/OU=GlobusTest/OU=GridOnDemand/CN=Grid User/CN=1418860555
issuer   : /O=Grid/OU=GlobusTest/OU=GridOnDemand/CN=Grid User
identity : /O=Grid/OU=GlobusTest/OU=GridOnDemand/CN=Grid User
type     : RFC 3820 compliant impersonation proxy
strength : 512 bits
path     : /tmp/x509up_u1003
timeleft : 11:57:02
griduser@ec2-79-125-48-213:~$ globus-job-run localhost /usr/bin/whoami
torqueuser
griduser@ec2-79-125-48-213:~$ █
```

AWS Management Co... x Grid On Demand in t... x

http://localhost:8080/

hostname [download CA certs](#) [User management](#)

jobs running 0 queued 0 nodes running 1 pending 0 money spent \$0.38 = €0.31

cpu's running (green) cpu's pending (yellow) jobs running (blue) queued (red)

cpu's running (green) cpu's pending (yellow) jobs running (blue) queued (red)

cpu's running (green) cpu's pending (yellow) jobs running (blue) queued (red)

cpu's running (green) cpu's pending (yellow) jobs running (blue) queued (red)

aws-portal.amazon.com/gp/aws/.../index.ht...

To be able to spawn new compute nodes on demand, Amazon Access Credentials are needed. They can be found [here](#).

Access Key ID

Secret Access Key

Access Credentials

There are three types of access credentials used to authenticate your requests to AWS services: (a) access key X.509 certificates, and (c) key pairs. Each access credential type is explained below.

Access Keys | X.509 Certificates | Key Pairs

Use access keys to make secure REST or Query protocol requests to any AWS service API. We create one for you when your account is created — see your access key below.

Your Access Keys

Created	Access Key ID	Secret Access Key	Status
January 31, 2010	AKIAJPB6XI2ZSJI6RLA	Show	Active (Make Inactive)
February 15, 2010	AKIAJIIQSOYC3ZWISH3Q	Show	Active (Make Inactive)

For your protection, you should never share your secret access keys with anyone. In addition, industry best practice recommends frequent key rotation.

[Learn more about Access Keys](#)

Sign-In Credentials

To sign in to AWS web sites and applications, AWS requires your Amazon e-mail address and password. Additionally, AWS supports the AWS Multi-Factor Authentication option. Each sign-in credential is explained below.

Amazon E-mail Address and Password

To sign in to secure pages on the AWS web site, the AWS Management Console, the AWS Discussion Forums, the AWS Premium Support site, you need to provide your Amazon e-mail address and password.

javascript:showAccessKey('AKIAJPB6XI2ZSJI6RLA')

AWS Management Co... x Grid On Demand in t... x Amazon Web Services x

http://localhost:8080/

hostname [download CA certs](#)

jobs running 0 queued 0 nodes running 1 pending 0 money spent \$0.38 = €0.31

cpu's running cpu's pending jobs running queued

cpu's running cpu's pending jobs running queued

cpu's running cpu's pending jobs running queued

ROOTJOB / TOBI GETTIER

ROOTJOB / TOBI GETTIER

ROOTJOB / TOBI GETTIER

ROOTJOB / TOBI GETTIER

To be able to spawn new compute nodes on demand, Amazon Access Credentials are needed. They can be found [here](#).

Access Key ID

Secret Access Key

AWS Management Co... x Grid On Demand in t... x Amazon Web Services x

http://localhost:8080/

hostname [download CA certs](#) [User management](#)

jobs running 0 queued 0 nodes running 1 pending 0 money spent \$0.38 = €0.31

incl. head node

minimum

maximum

startup window

startup type

	mem	cores	CU's	price
<input checked="" type="radio"/>	7.5	2	4	\$0.38
<input type="radio"/>	15	4	8	\$0.76
<input type="radio"/>	17.1	2	6.5	\$0.57
<input type="radio"/>	34.2	4	13	\$1.34
<input type="radio"/>	68.4	8	26	\$2.68
<input type="radio"/>	7	8	20	\$1.16

```
griduser@ec2-79-125-48-213: ~
willem@diglett: ~ 84x22
willem@diglett:~$ grid-proxy-init
Your identity: /O=dutchgrid/O=users/O=uva/OU=wins/CN=Willem Toorop
Creating proxy ..... Done
Your proxy is valid until: Tue Jun 29 11:50:33 2010
willem@diglett:~$
```

hostname [download CA certs](#)

jobs running 0 queued 0 nodes running 1 pending 0 money spent \$0.38 = €0.31

[certificates](#)

localhost:8080/gridondemand-ca-certs.tgz

```
griduser@ec2-79-125-48-213: ~
willem@diglett: ~
willem@diglett:~$ grid-proxy-init
Your identity: /O=dutchgrid/O=users/O=uva/OU=wins/CN=Willem Toorop
Creating proxy ..... Done
Your proxy is valid until: Tue Jun 29 11:55:08 2010
willem@diglett:~$ tar xzvf gridondemand-ca-certs.tgz
.globus/certificates/
.globus/certificates/globus-host-ssl.conf.5c703074
.globus/certificates/grid-security.conf.5c703074
.globus/certificates/globus-user-ssl.conf.5c703074
.globus/certificates/5c703074.signing_policy
.globus/certificates/5c703074.0
willem@diglett:~$ for i in {1..10}
> do
>   globus-job-submit ec2-79-125-48-213.eu-west-1.compute.amazonaws.com \
>     /bin/sleep `expr $i \* 5` &
>   sleep $i
> done
```

AWS Management Co... x Grid On Demand in t... x Amazon Web Services x

http://localhost:8080/

hostname [download CA certs](#) [User management](#)

jobs running 2 queued 3 nodes running 1 pending 1 money spent \$0.76 = €0.61

incl. head node
 minimum
 maximum
 startup window

startup type				
	mem	cores	CU's	price
<input checked="" type="radio"/>	7.3	2	4	\$0.38
<input type="radio"/>	15	4	8	\$0.76
<input type="radio"/>	17.1	2	6.5	\$0.57
<input type="radio"/>	34.2	4	13	\$1.34
<input type="radio"/>	68.4	8	26	\$2.68
<input type="radio"/>	7	8	20	\$1.16

AWS Management Co... x Grid On Demand in t... x Amazon Web Services x

http://localhost:8080/

hostname [download CA certs](#)

jobs running 0 queued 0 nodes running 2 pending 0 money spent \$0.76 = €0.61

incl. head node

minimum

maximum

startup window

startup type

	mem	cores	CU's	price
<input checked="" type="radio"/>	7.5	2	4	\$0.38
<input type="radio"/>	15	4	8	\$0.76
<input type="radio"/>	17.1	2	6.5	\$0.57
<input type="radio"/>	34.2	4	13	\$1.34
<input type="radio"/>	68.4	8	26	\$2.68
<input type="radio"/>	7	8	20	\$1.16

cpu's running cpu's pending jobs running queued

cpu's running cpu's pending jobs running queued

cpu's running cpu's pending jobs running queued

AWS Management Co... x Grid On Demand in t... x Amazon Web Services x

http://localhost:8080/

hostname [download CA certs](#)

jobs running 0 queued 0 nodes running 1 pending 0 money spent \$1.14 = €0.92

cpu's running (green) | cpu's pending (yellow) | jobs running (blue) | queued (red)

incl. head node
 minimum
 maximum
 startup window
 startup type

	mem	cores	CU's	price
<input checked="" type="radio"/>	7.5	2	4	\$0.38
<input type="radio"/>	15	4	8	\$0.76
<input type="radio"/>	17.1	2	6.5	\$0.57
<input type="radio"/>	34.2	4	13	\$1.34
<input type="radio"/>	68.4	8	26	\$2.68
<input type="radio"/>	7	8	20	\$1.16

AWS Management Co... x Grid On Demand in t... x Amazon Web Services x

http://localhost:8080/

hostname [download CA certs](#)

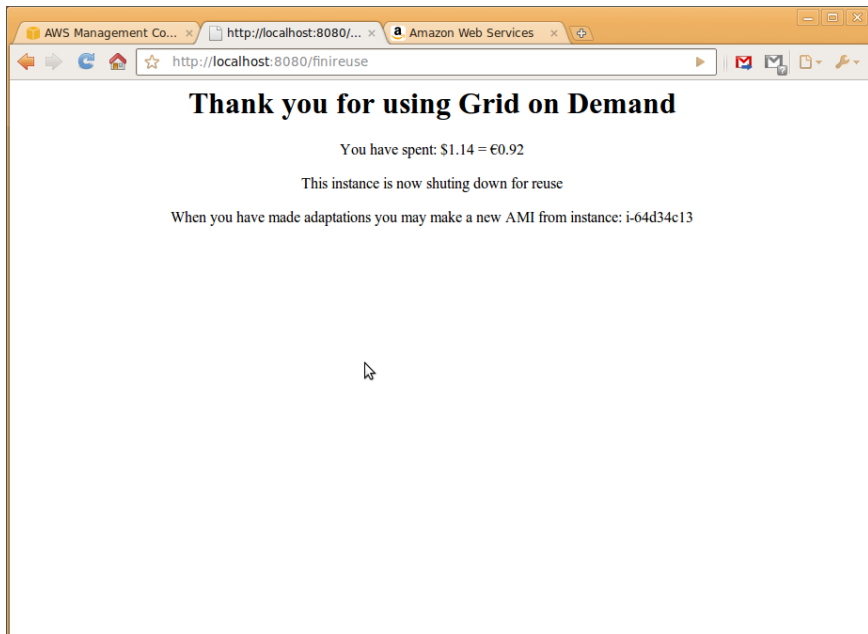
jobs running 0 queued 0 nodes running 1 pending 0 money spent \$1.14 = €0.92

cpu's running (green) | cpu's pending (yellow) | jobs running (blue) | queued (red)

incl. head node
 minimum
 maximum
 startup window

startup type

	mem	cores	CU's	price
<input checked="" type="radio"/>	7.5	2	4	\$0.38
<input type="radio"/>	15	4	8	\$0.76
<input type="radio"/>	17.1	2	6.5	\$0.57
<input type="radio"/>	34.2	4	13	\$1.34
<input type="radio"/>	68.4	8	26	\$2.68
<input type="radio"/>	7	8	20	\$1.16



A screenshot of a web browser window. The browser has two tabs: "AWS Management Co..." and "http://localhost:8080/...". The address bar shows "http://localhost:8080/finireuse". The main content area displays the following text:

Thank you for using Grid on Demand

You have spent: \$1.14 = €0.92

This instance is now shutting down for reuse

When you have made adaptations you may make a new AMI from instance: i-64d34c13

A mouse cursor is visible in the center of the page.

The screenshot displays the AWS Management Console for Amazon EC2. The main content area shows a table of instances under the heading "My Instances". A context menu is open over the instance with ID "i-64d34c13".

Instance	AMI ID	Root	Type	Status	Lifecycle	Public DNS
i-ee861499	ami-5313392	ebs	m1.large	stopped	normal	
<input checked="" type="checkbox"/>	i-64d34c13	ami-57ebc12	ebs	m1.large	stopped	normal
i-40d24d37	ami-57ebc12	ebs	m1.large	terminating		

The context menu for instance "i-64d34c13" includes the following options:

- Instance Management
 - Connect
 - Get System Log
 - Create Image (EBS AMI)**
 - Launch More Like This
 - Disassociate IP Address
- Instance Lifecycle
 - Terminate
 - Reboot
 - Stop
 - Start
- CloudWatch Monitoring
 - Enable CloudWatch
 - Disable CloudWatch

Below the table, the console shows details for the selected instance "EC2 Instance: i-64d34c13".

Description		Monitoring	
AMI ID:	ami-57ebc123	Zone:	
Security Groups:	default	Type:	
Status:	stopped	Owner:	021450

"Real use" and Performance test

To test the possibilities and performance of Grid on Demand a **Biomedical Application** and a **workflow manager** are used

"Real use" and Performance test

To test the possibilities and performance of Grid on Demand a **Biomedical Application** and a **workflow manager** are used

Biomedical Application: **WAVE**

Parallel (MonteCarlo simulation) application

"Real use" and Performance test

To test the possibilities and performance of Grid on Demand a **Biomedical Application** and a **workflow manager** are used

Biomedical Application: **WAVE**

Parallel (MonteCarlo simulation) application

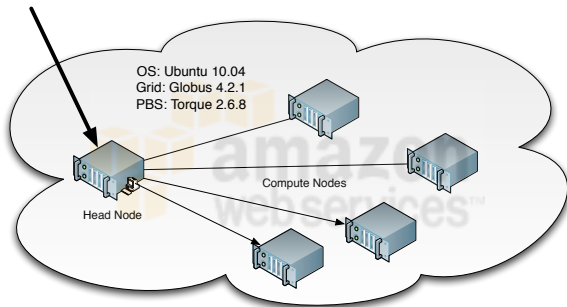
Workflow Manager

- A (graphical) tool to assist complex e-Science application creation
- Creates a series of jobsubmissions
- WS-VLAM created by UvA SNE Group
- Connects to Globus Grid Interface

Connecting the Grid on Demand to a Workflow Manager



WAVE
Biomedical
Compute
Application





Preliminary performance results: total execution time



Grid on demand compared to a 32 node physical UvA cluster (DAS3) with Globus Grid interface

Grid	Jobs	Total execution time
DAS3	28	00:42:58
Grid on Demand	28	00:46:22
DAS3	98	02:46:36
Grid on Demand	98	00:52:26

Preliminary performance results of job pending and running time: 30 jobs



Preliminary performance results of job pending and running time: 100 jobs



Conclusion

- Grid on Demand can offer Grid as a Cloud service
- Parallel applications run/scale very well on Grid on Demand
- Grid on Demand can be a solution for Urgent Computing demands

Questions