



---

CycleCloud Tutorial

Overview	3
Account Registration	3
Cluster Configuration	8
Purchase CycleCloud Credits	10
Starting Your Cluster	13
Cluster Management using CycleServer	17
Uploading Data to the Cluster	19
Submitting Jobs	20
Condor	21
SGE	24
Torque	27
Downloading Results	30
Stopping Your Cluster	30
Billing	33
Cluster Configuration Options	34

## CycleCloud Tutorial

This tutorial will get you up to speed with using the CycleCloud<sup>SM</sup> webapp. After this tutorial you will be able to login to CycleCloud, create and start a cluster, copy data to your cluster's shared filesystem, submit your job and get your results.

### Overview

CycleCloud allows you to start high-performance, high-throughput computing clusters using cloud infrastructure built upon Amazon Web Services. CycleCloud currently supports three scheduling options: Condor, SGE and Torque. Each cluster is managed by the CycleServer<sup>TM</sup> Management Tool and provides a shared file system. Clusters can be configured to dynamically expand to deal with incoming jobs and then terminate executing nodes once jobs are complete.

Additional [CycleCloud information](#).

### Account Registration

The first step to using CycleCloud is to [register for a developer account](#).

Enter your relevant information and click on 'View Terms of Service'.

## CREATE ACCOUNT

**Company/Organization:**

**First Name:**

**Last Name:**

**Email Address:**


**Password:**

**Password (again):**

**Phone Number:**

**Validation:**  
What is 0 + 9?

Terms and conditions have not been accepted.



After reading through the terms of service click on the 'accept' option.

**Accept or decline agreement terms**

### Terms of Service

**4.1.** The Service shall be subject to the privacy policy for the Service which is available at [/privacy/](#). You agree to the use of your data in accordance with CycleCloud's privacy policies.

**4.2.** CycleCloud shall not be liable to you or to any third party for any modification, price change, suspension or discontinuance of the Service.

**4.3.** Since CycleCloud uses Amazon Web Services your Content, Data and Applications are subject to Amazon's privacy policy as well. That policy is located at <http://aws.amazon.com/privacy/>.

### 5. Payment


**5.1.** CycleCloud uses a prepayment system for use of the Service. You must have a positive account balance to continue use of the Service. If your account balance becomes negative you agree to pay any applicable fees. A negative account balance may cause termination of Service and may result in loss of Content, Data and Applications.

**5.2.** In consideration of your use of any of the Service, you agree to pay applicable fees for the Service in the amounts set forth on the Service detail page. Fees for any new Service or new Service feature will be effective upon posting by us on the CycleCloud Website for the applicable Service. We may increase or add new fees by giving you 30 days' advance notice. Such notice will be posted on the CycleCloud Website on the Service detail page for the affected Service. You agree that you are responsible for checking the CycleCloud Website each month to confirm whether there are any new fees and their effective date(s).

**5.3.** In the case of account cancellation or termination, the remaining account balance after all applicable fees have been removed will be refunded to you. Refunds may take over a month to apply due to the account billing process.

I accept the terms of this agreement

I decline



Once you click 'Create Account' you will receive an email that requires you to activate your account.

## CREATE ACCOUNT

**Company/Organization:**

**First Name:**

**Last Name:**

**Email Address:**


**Password:**

**Password (again):**

**Phone Number:**

**Validation:**  
What is 0 + 9?

Terms and conditions have been accepted.

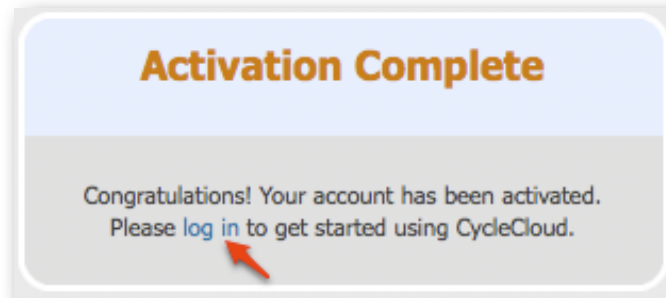


### Registration Complete

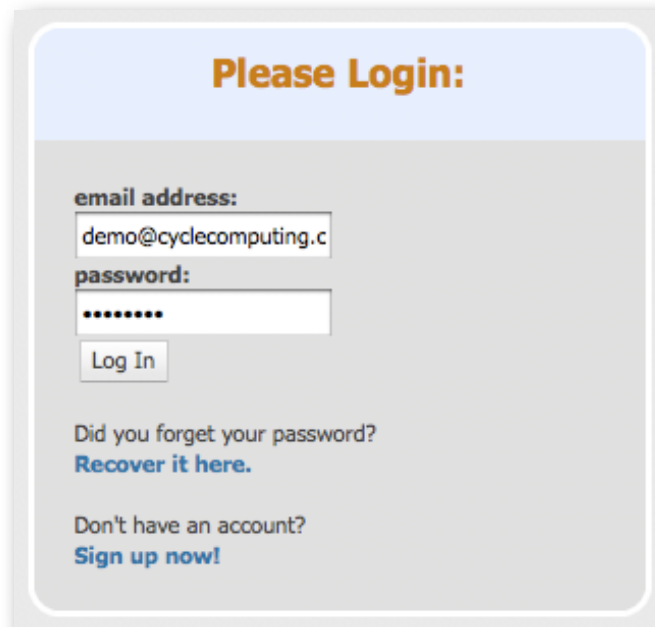
**Thank you for registering for an account on  
CycleCloud!**

An activation email has been sent to you. Please  
check your email and follow the instructions activate  
your account and get started using CycleCloud.

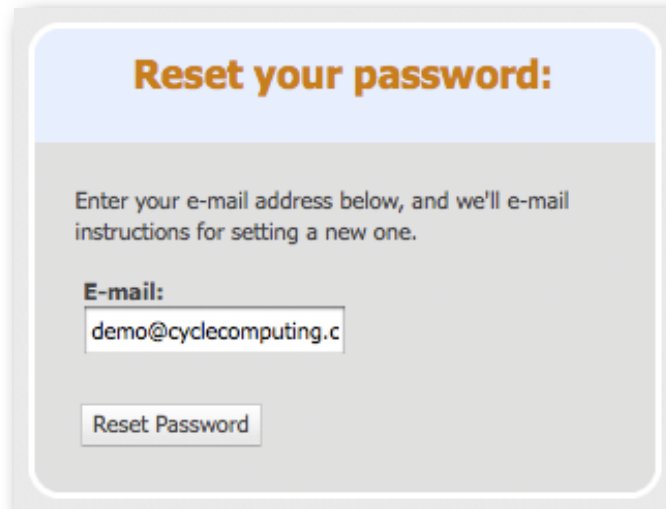
Clicking on the activation link in the email will validate your email account and activate your account.



You can now [log in](#) to your account.



If you ever forget your password you can [recover it](#).



**Reset your password:**

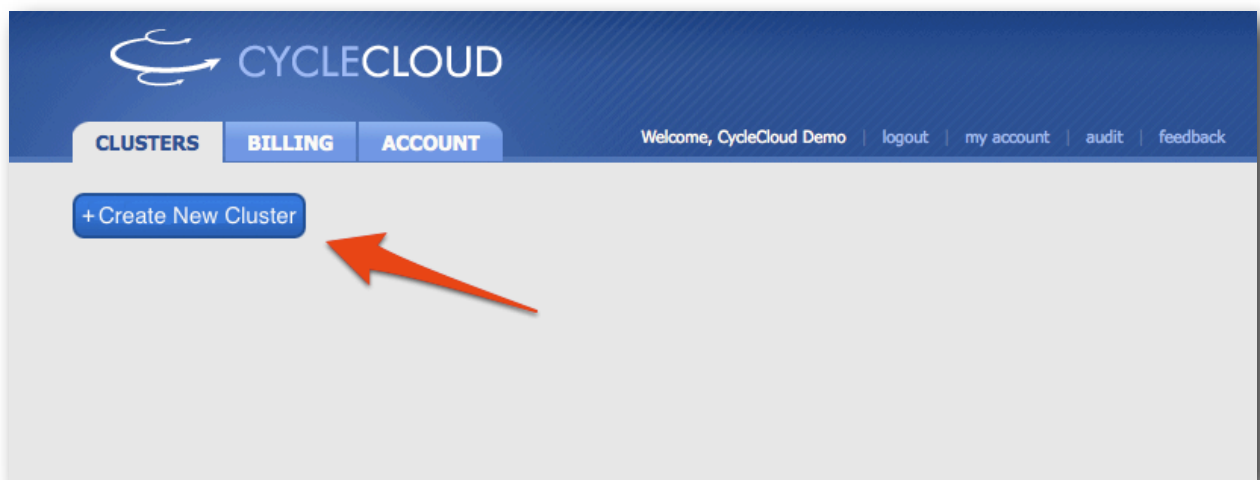
Enter your e-mail address below, and we'll e-mail instructions for setting a new one.

**E-mail:**

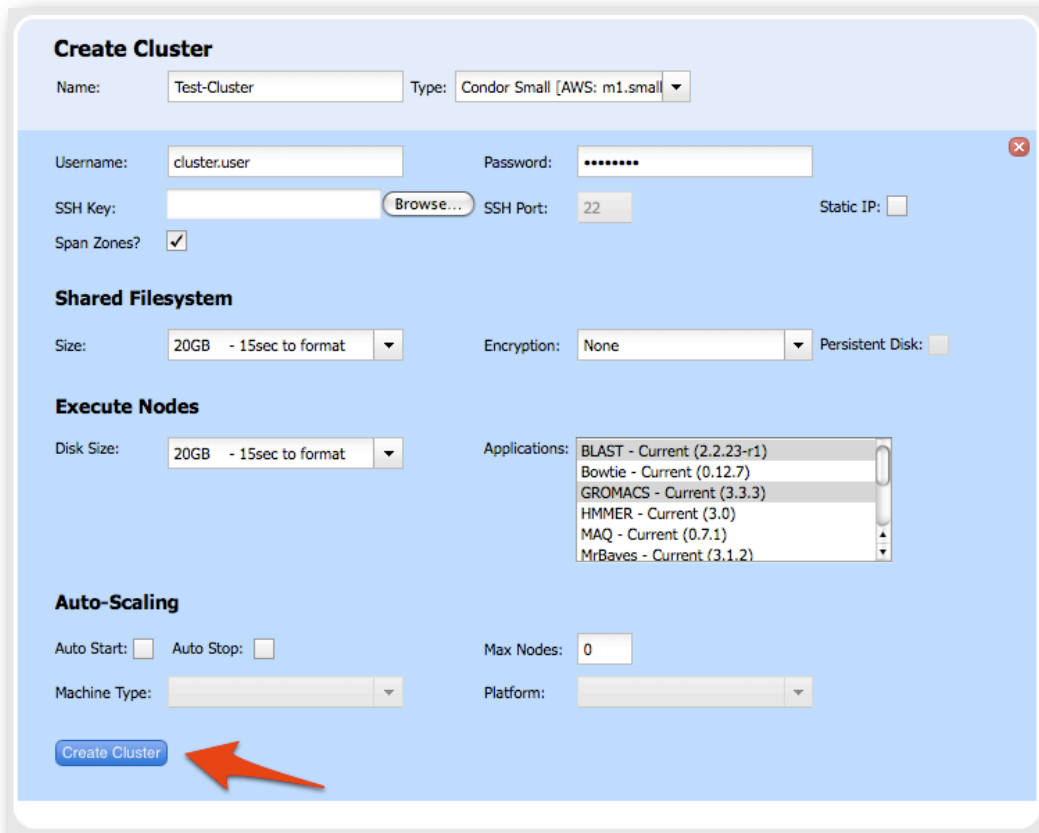
If you require capabilities beyond what the Developer Plan can provide you may [contact us](#) to upgrade your plan.

## Cluster Configuration

Once you have logged in you can now create your first cluster by clicking on 'Create Cluster'.



This will open a form with a number of cluster configuration options. We will go over the various options in detail below but for now you will need to create a cluster name, choose a cluster type, add a cluster username and password.



The screenshot shows a 'Create Cluster' form with the following fields and options:

- Name:** Test-Cluster
- Type:** Condor Small [AWS: m1.small]
- Username:** cluster.user
- Password:** [masked]
- SSH Key:** [empty] **Browse...**
- SSH Port:** 22
- Static IP:** [checkbox]
- Span Zones?**
- Shared Filesystem:**
  - Size:** 20GB - 15sec to format
  - Encryption:** None
  - Persistent Disk:** [checkbox]
- Execute Nodes:**
  - Disk Size:** 20GB - 15sec to format
  - Applications:**
    - BLAST - Current (2.2.23-r1)
    - Bowtie - Current (0.12.7)
    - GROMACS - Current (3.3.3)
    - HMMER - Current (3.0)
    - MAQ - Current (0.7.1)
    - MrBaves - Current (3.1.2)
- Auto-Scaling:**
  - Auto Start:** [checkbox]
  - Auto Stop:** [checkbox]
  - Max Nodes:** 0
  - Machine Type:** [dropdown]
  - Platform:** [dropdown]

A red arrow points to the **Create Cluster** button at the bottom left of the form.

Once you click on 'Create Cluster' your cluster will be generated for you. The cluster is not running and you will need to purchase credits before you can start it.

## Purchase CycleCloud Credits

CycleCloud requires prepayment before you can start clusters. If you try to start a cluster or nodes with a zero or negative balance you will get an error message that will direct you to the [payments page](#).

### Purchase CycleCloud Credits

CycleCloud uses a prepay billing model. CycleCloud requires a customer to retain a positive account balance for continued use of the service. If an account's balance is emptied the account runs the risk of service interruption which may result in the termination of running clusters and the possible loss of data due to the non-persistent nature of the cloud.


Funds can be added to an account in \$10.00 increments. We use Paypal for payment processing. You do not need a Paypal account to use the Paypal system since it accepts all major credit cards. Once a deposit has been made to your account via Paypal it may take up to an hour before it is cleared and registered by CycleCloud. A refund of the remaining balance will be issued through Paypal in the case of account cancellation.

If you have any questions or concerns please contact us.

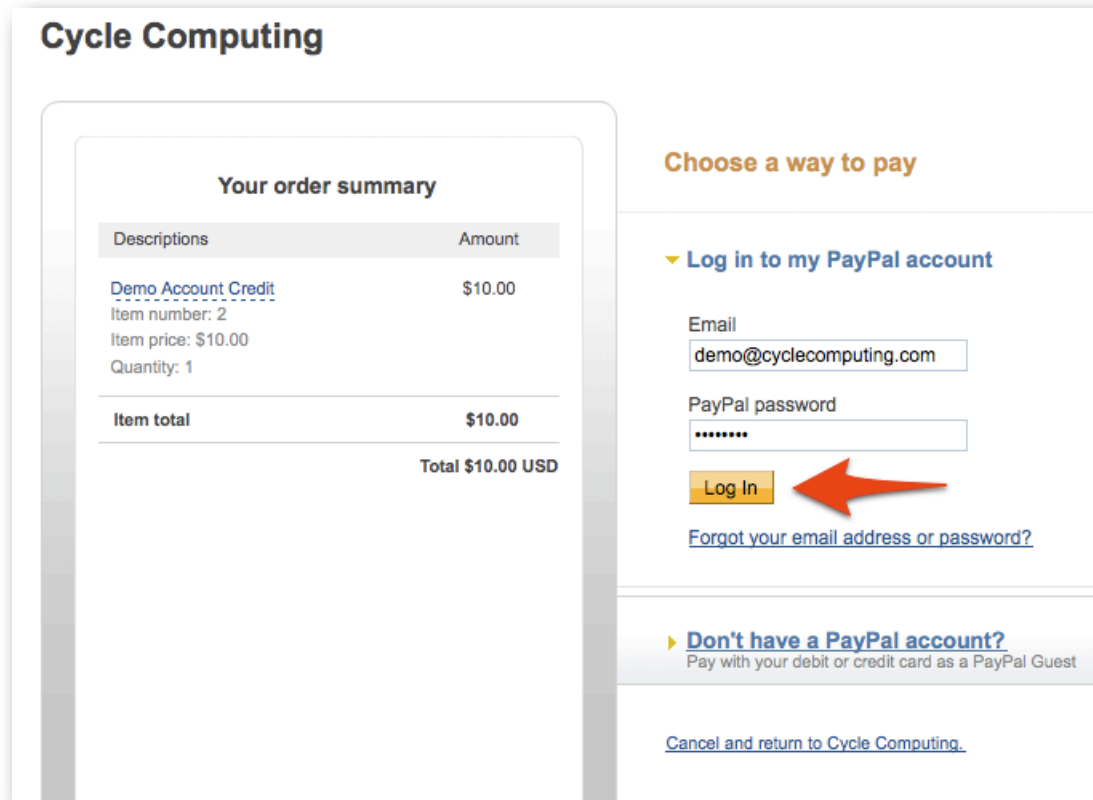
**Make a deposit to Demo's account balance:**

\$10.00 CycleCloud Credit:

Total Purchase: \$10.00 [Buy Credits](#)



We use PayPal™ for handling all payments so when you click on ‘Buy Credits’ you will be brought to a PayPal payment page. You will need to log on to the PayPal page using your PayPal user credentials. If you don’t have a Paypal account you will need to create one.



The screenshot shows the Cycle Computing checkout page. On the left, under 'Your order summary', there is a table with two columns: 'Descriptions' and 'Amount'. The table lists 'Demo Account Credit' for \$10.00, with item number 2, price \$10.00, and quantity 1. Below this, the 'Item total' is \$10.00, and the overall 'Total' is \$10.00 USD. On the right, under 'Choose a way to pay', there are two main options. The first is 'Log in to my PayPal account', which includes input fields for 'Email' (demo@cyclecomputing.com) and 'PayPal password' (masked with dots), a 'Log In' button (highlighted with a red arrow), and a link for 'Forgot your email address or password?'. The second option is 'Don't have a PayPal account?', which allows payment as a guest. At the bottom, there is a link to 'Cancel and return to Cycle Computing.'

Descriptions	Amount
Demo Account Credit Item number: 2 Item price: \$10.00 Quantity: 1	\$10.00
<b>Item total</b>	<b>\$10.00</b>
<b>Total \$10.00 USD</b>	

**Choose a way to pay**

▼ **Log in to my PayPal account**

Email  
demo@cyclecomputing.com

PayPal password  
\*\*\*\*\*

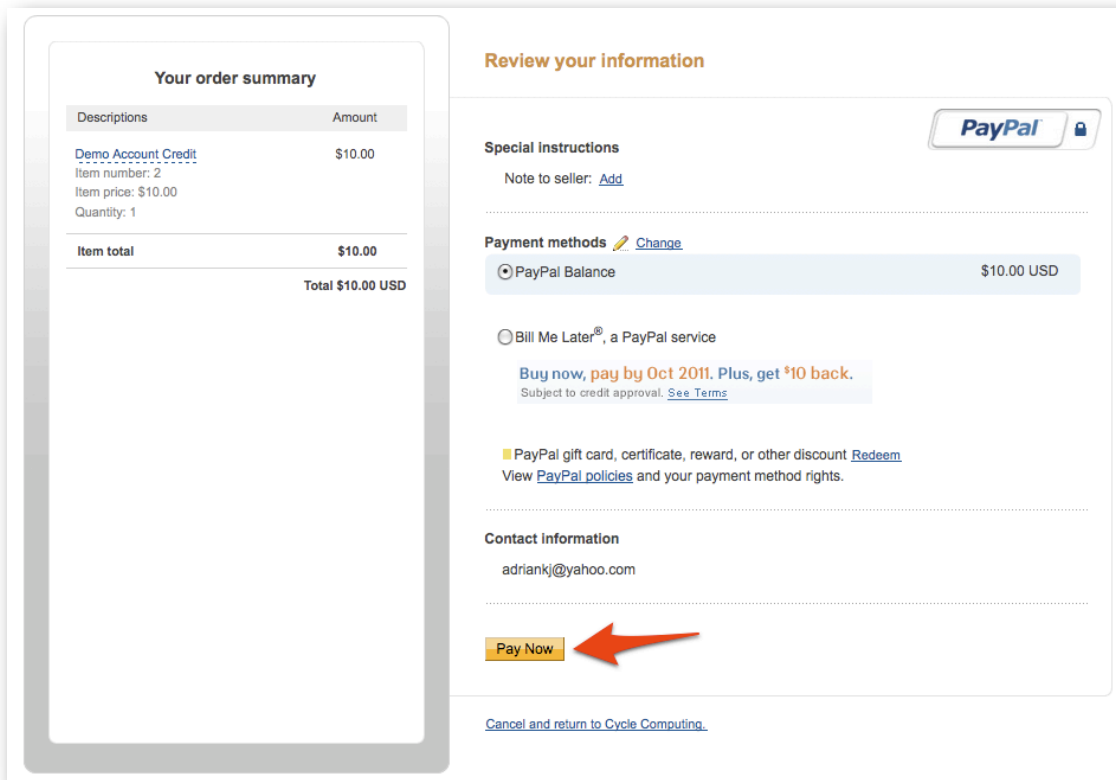
**Log In** ←

[Forgot your email address or password?](#)

► **Don't have a PayPal account?**  
Pay with your debit or credit card as a PayPal Guest

[Cancel and return to Cycle Computing.](#)


Once you have logged into PayPal you can verify the purchase amount and then purchase the CycleCloud credits.



**Your order summary**

Descriptions	Amount
Demo Account Credit Item number: 2 Item price: \$10.00 Quantity: 1	\$10.00
<b>Item total</b>	<b>\$10.00</b>
<b>Total \$10.00 USD</b>	

**Review your information**

**Special instructions** 

Note to seller: [Add](#)

---

**Payment methods** [Change](#)

PayPal Balance \$10.00 USD

Bill Me Later®, a PayPal service


Buy now, pay by Oct 2011. Plus, get \$10 back.  
Subject to credit approval. [See Terms](#)

PayPal gift card, certificate, reward, or other discount [Redeem](#)  
View [PayPal policies](#) and your payment method rights.

---

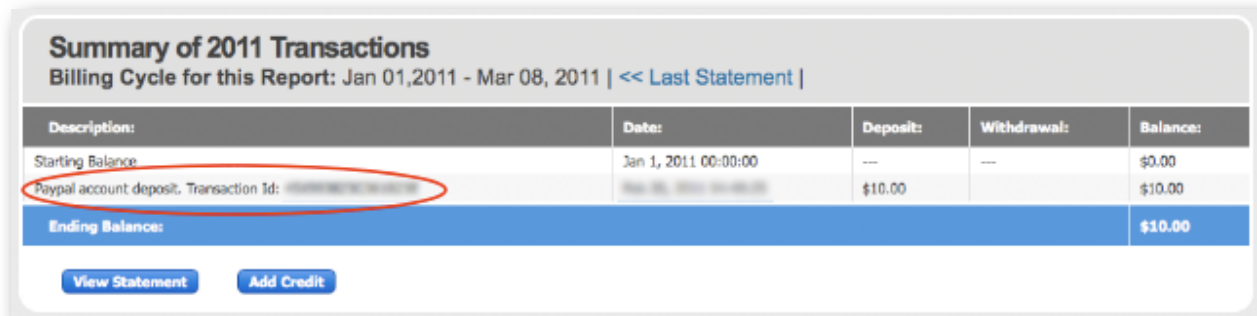
**Contact information**

adriankj@yahoo.com

**Pay Now** 

[Cancel and return to Cycle Computing.](#)

Your payment will now be processed by PayPal. You should receive a receipt from PayPal and CycleCloud will receive a notice about your payment. Once CycleCloud has notice of your payment it will credit your account and you will be able to start a cluster. This can take up to an hour. You can view your transactions in the [ledger](#).



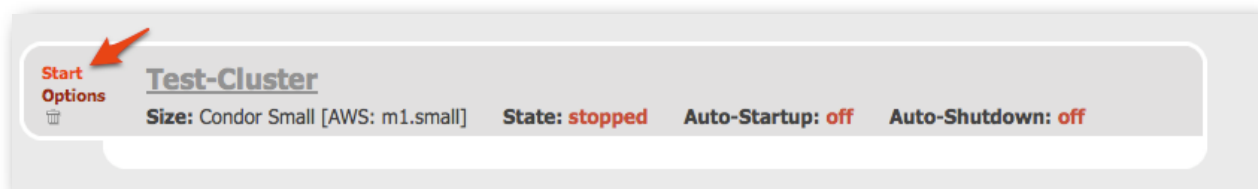
**Summary of 2011 Transactions**  
Billing Cycle for this Report: Jan 01, 2011 - Mar 08, 2011 | [<< Last Statement](#) |

Description:	Date:	Deposit:	Withdrawal:	Balance:
Starting Balance	Jan 1, 2011 00:00:00	---	---	\$0.00
Paypal account deposit. Transaction Id: <a href="#">XXXXXXXXXXXX</a>	Feb 28, 2011 00:00:00	\$10.00		\$10.00
<b>Ending Balance:</b>				<b>\$10.00</b>

[View Statement](#) [Add Credit](#)

## Starting Your Cluster

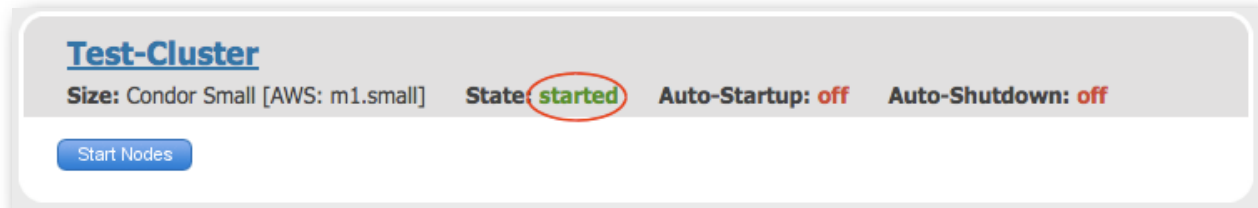
Now that CycleCloud credits have been deposited you can start the cluster by clicking the 'Start' link that will appear when you hover over your cluster.



A cluster has an underlying recipe that describes which instances must be run in which order. In the case of the Condor Small cluster that recipe requires that two instances start. The first is the Central Manager which contains the CycleServer Management Tool and the second is the Scheduler which hosts the shared filesystem and scheduling tools. When the cluster is started both of the instances must start before the cluster is ready for use. We can see the cluster's pending requests.

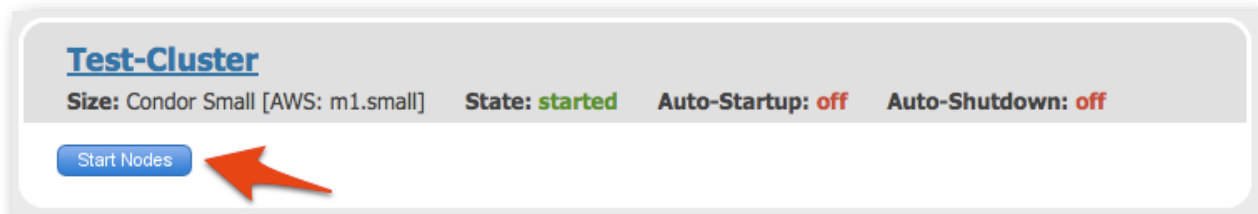


The time it takes for the instances to start-up is dependent upon a number of variables including the applications, the scheduler size, the size of the shared filesystem and a number of others. The cluster's state will change to 'started' once the cluster is ready for use.



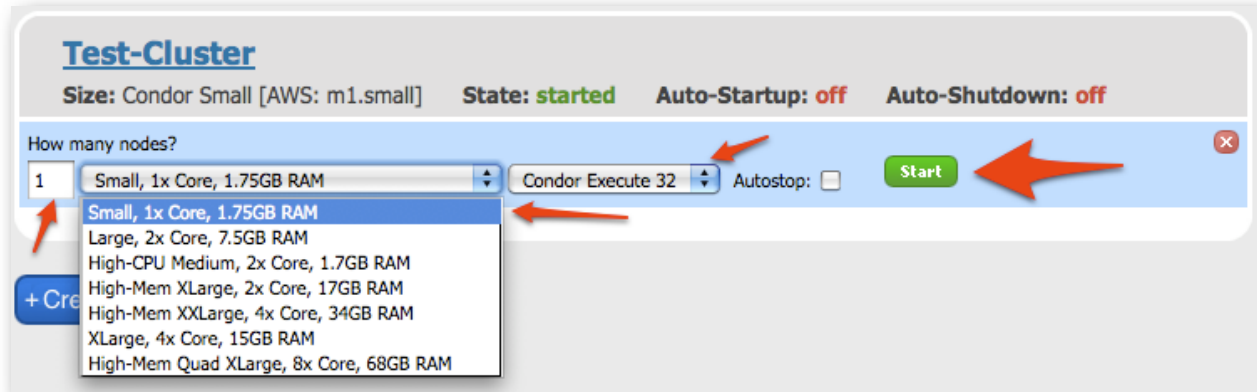
**Test-Cluster**  
Size: Condor Small [AWS: m1.small]    State: **started**    Auto-Startup: off    Auto-Shutdown: off  
[Start Nodes](#)

Execute nodes can be added to the cluster once it is started. The execute nodes are responsible for running any submitted jobs. The cluster can be configured to automatically size the cluster based upon pending jobs but for this tutorial we'll just manually start a single execute node by clicking the 'Start Nodes' button.

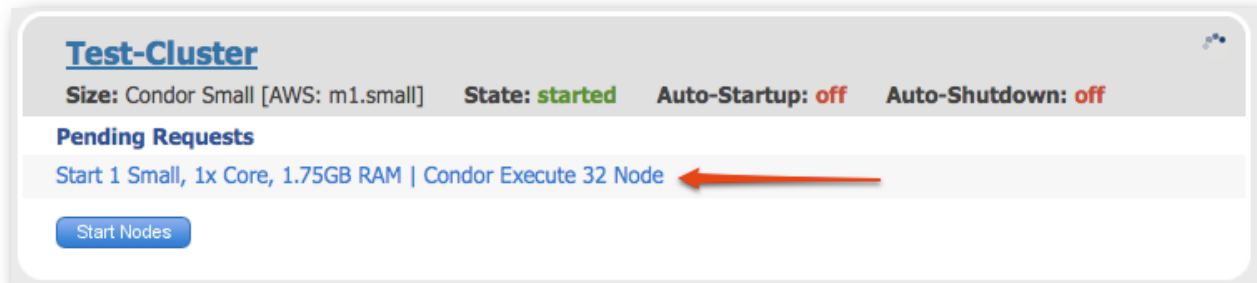


**Test-Cluster**  
Size: Condor Small [AWS: m1.small]    State: **started**    Auto-Startup: off    Auto-Shutdown: off  
[Start Nodes](#)

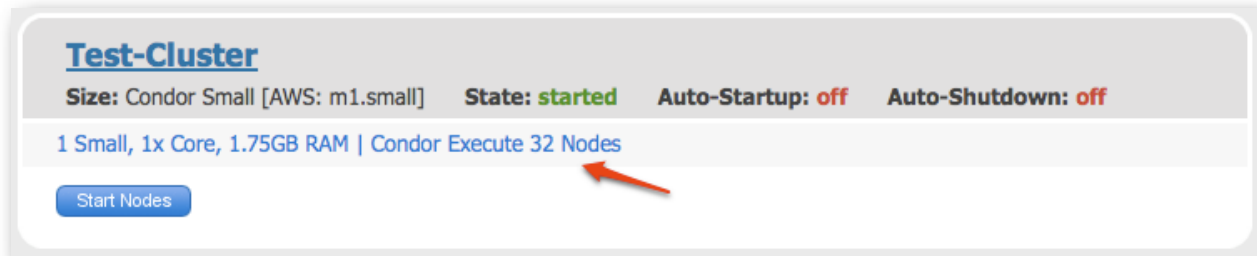
You will need to determine what type of machine you want started, what image will run on the machine, and how many you wish to start. In this case we've chosen to start a single small 32-bit instance. The node is requested once the 'Start' button is clicked.



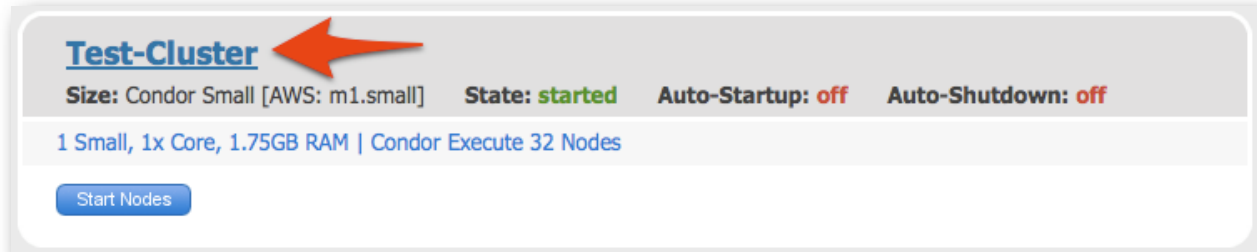
We can now see that a single small instance has been requested.




Once the request has been fulfilled we can see the running instance.



You can view details of your running cluster by clicking on the cluster's name.



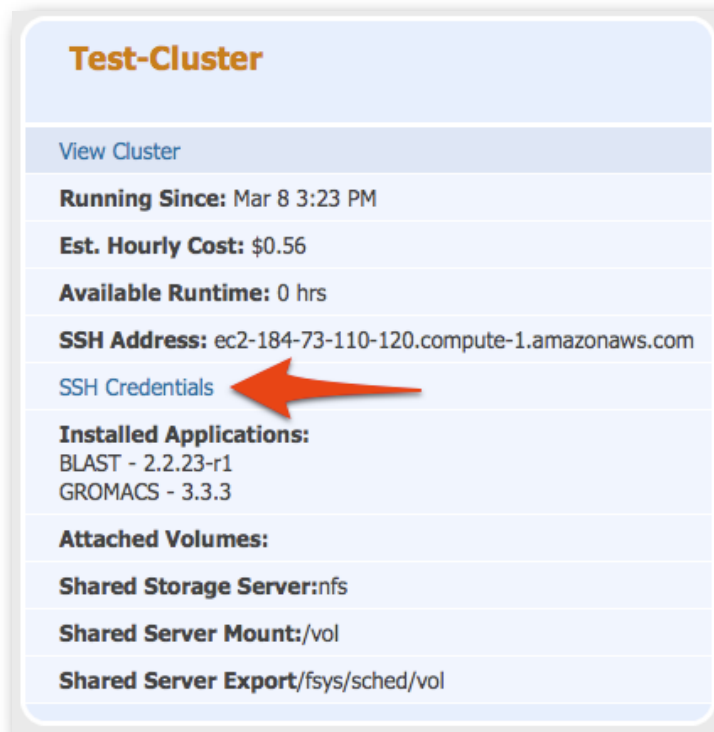
**Test-Cluster** 

**Size:** Condor Small [AWS: m1.small]    **State:** started    **Auto-Startup:** off    **Auto-Shutdown:** off

1 Small, 1x Core, 1.75GB RAM | Condor Execute 32 Nodes

[Start Nodes](#)

The detail sidebar contains useful data on the pool including a rough estimate on the hourly cost, the SSH address of the scheduler and a list of all installed applications. You can also see the cluster credentials which are used for accessing the CycleServer management tool and SSHing into the cluster's nodes.



**Test-Cluster**


[View Cluster](#)

**Running Since:** Mar 8 3:23 PM

**Est. Hourly Cost:** \$0.56

**Available Runtime:** 0 hrs

**SSH Address:** ec2-184-73-110-120.compute-1.amazonaws.com

[SSH Credentials](#) 

**Installed Applications:**  
BLAST - 2.2.23-r1  
GROMACS - 3.3.3

**Attached Volumes:**


**Shared Storage Server:** nfs

**Shared Server Mount:** /vol

**Shared Server Export:** /fsys/sched/vol

## Cluster Management using CycleServer

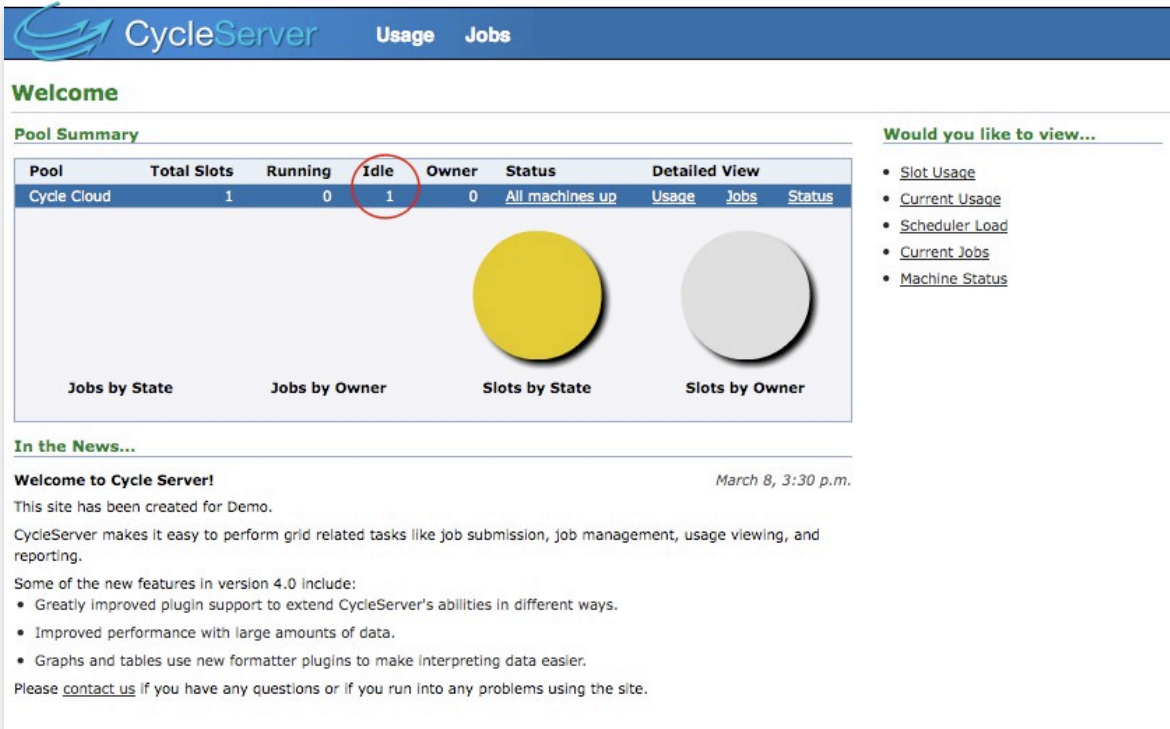
CycleServer can provide a comprehensive view into your cluster. You can access CycleServer by clicking, 'View'.



You can log into CycleServer with the cluster credentials you set when you first created your cluster.



Once logged into CycleServer you can view your resources and jobs. In this case we can see that we have one idle core that is waiting for a job.



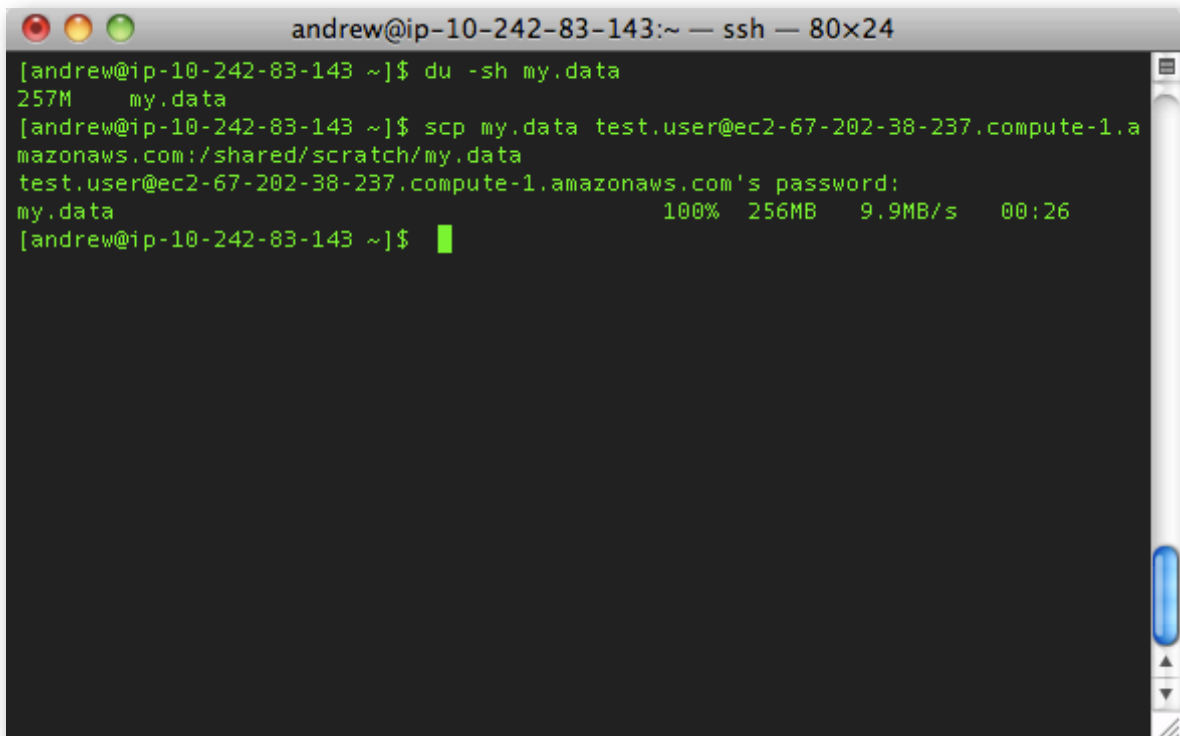
The screenshot shows the CycleServer interface with the 'Usage' tab selected. A table displays the 'Pool Summary' for 'Cycle Cloud'. The 'Idle' column is circled in red, showing a value of 1. Below the table are four circular charts: 'Jobs by State' (yellow), 'Jobs by Owner' (grey), 'Slots by State' (yellow), and 'Slots by Owner' (grey). A sidebar on the right offers links to 'Slot Usage', 'Current Usage', 'Scheduler Load', 'Current Jobs', and 'Machine Status'. A 'Welcome to Cycle Server!' message is visible at the bottom, dated March 8, 3:30 p.m.

Pool	Total Slots	Running	Idle	Owner	Status	Detailed View		
Cycle Cloud	1	0	1	0	All machines up	Usage	Jobs	Status

You can find the full CycleServer documentation in the [CycleServer user guide](#).

## Uploading Data to the Cluster

Before you can run your jobs on the cluster you will need to upload your data. Use [SCP](#) to upload the data to the shared file system on the scheduler. You can get the scheduler address from the cluster's details page. You can also look up the credentials that are used for SSH from the cluster's detail page as well.



```
andrew@ip-10-242-83-143:~ — ssh — 80x24
[andrew@ip-10-242-83-143 ~]$ du -sh my.data
257M    my.data
[andrew@ip-10-242-83-143 ~]$ scp my.data test.user@ec2-67-202-38-237.compute-1.amazonaws.com:/shared/scratch/my.data
test.user@ec2-67-202-38-237.compute-1.amazonaws.com's password:
my.data                                100% 256MB  9.9MB/s   00:26
[andrew@ip-10-242-83-143 ~]$
```

## Submitting Jobs

Use a standard SSH connection to connect to the scheduler. Use the 'SSH credentials' input during the creation during the cluster to connect to the 'SSH Address' that is listed within the cluster information pane. For the example below the ssh address for our Test-Cluster scheduler is `ec2-184-73-110-120.compute-1.amazonaws.com` and the username is `cluster.user`.



```
cluster.user@ip-10-122-163-53:~ — ssh — 80x24
admiral:~ andrew$ ssh cluster.user@ec2-184-73-110-120.compute-1.amazonaws.com
cluster.user@ec2-184-73-110-120.compute-1.amazonaws.com's password:

  cyclecloud

Role: Cyclecloud Condor Scheduler
Pool created for: Demo
Pool: Test-Cluster

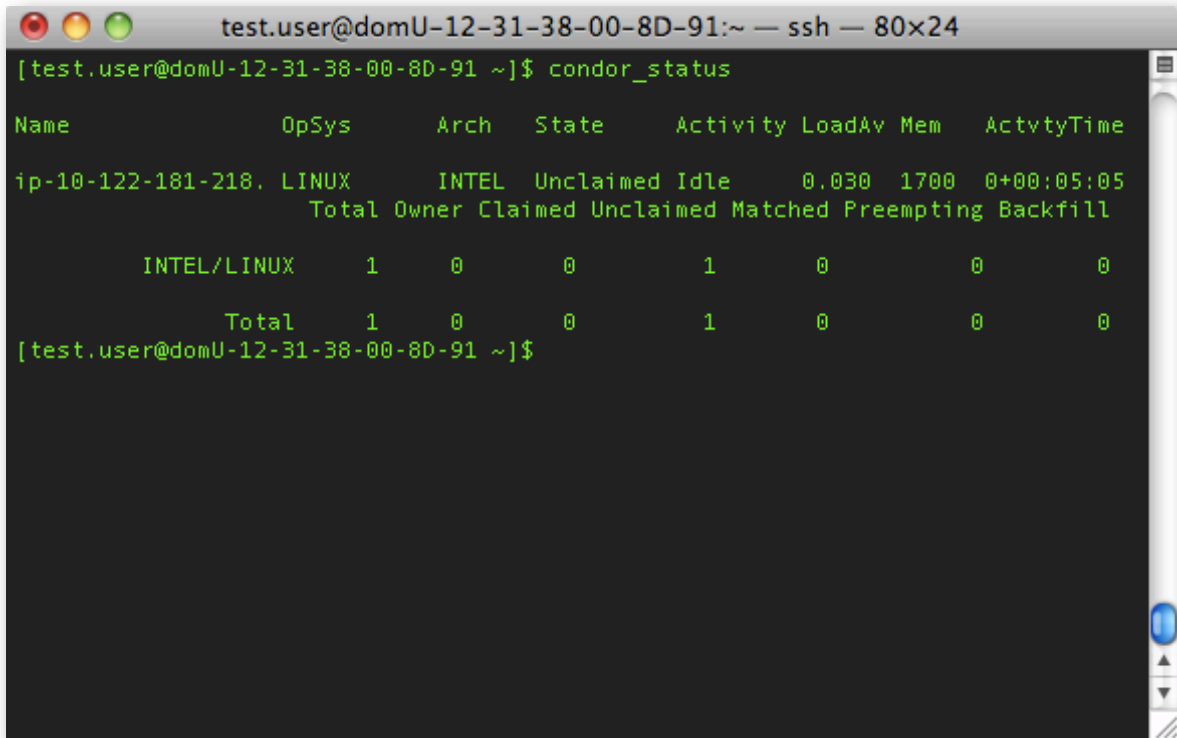
[cluster.user@ip-10-122-163-53 ~]$
```

Optionally, an ssh public key can be uploaded when creating a cluster. Using this feature, it is possible to access the scheduler without having to enter a password. A keypair can be created using the `ssh-keygen` command on your workstation.

CycleCloud currently supports three job schedulers. Which method you use to submit and monitor your jobs will depend on which type of scheduler you picked when you created your cluster.

## Condor

Examine the execute node that is available for jobs:



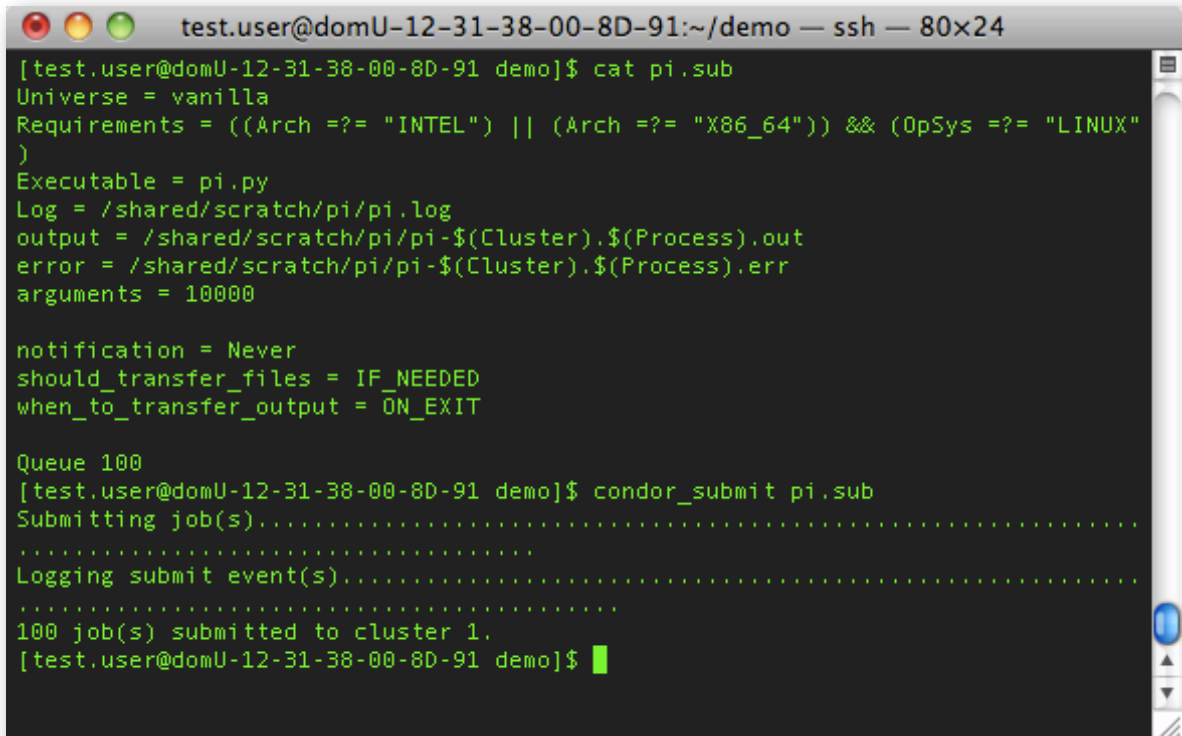
```

test.user@domU-12-31-38-00-8D-91:~ — ssh — 80x24
[test.user@domU-12-31-38-00-8D-91 ~]$ condor_status

Name                OpSys      Arch      State      Activity LoadAv Mem      ActvtyTime
ip-10-122-181-218. LINUX      INTEL     Unclaimed Idle      0.030  1700  0+00:05:05
                    Total Owner Claimed Unclaimed Matched Preempting Backfill
                    INTEL/LINUX    1    0    0    1    0    0    0
                    Total      1    0    0    1    0    0    0
[test.user@domU-12-31-38-00-8D-91 ~]$

```

Submit 100 jobs:



```
test.user@domU-12-31-38-00-8D-91:~/demo — ssh — 80x24
[test.user@domU-12-31-38-00-8D-91 demo]$ cat pi.sub
Universe = vanilla
Requirements = ((Arch == "INTEL") || (Arch == "X86_64")) && (OpSys == "LINUX"
)
Executable = pi.py
Log = /shared/scratch/pi/pi.log
output = /shared/scratch/pi/pi-$(Cluster).$(Process).out
error = /shared/scratch/pi/pi-$(Cluster).$(Process).err
arguments = 10000

notification = Never
should_transfer_files = IF_NEEDED
when_to_transfer_output = ON_EXIT

Queue 100
[test.user@domU-12-31-38-00-8D-91 demo]$ condor_submit pi.sub
Submitting job(s).....
.....
Logging submit event(s).....
.....
100 job(s) submitted to cluster 1.
[test.user@domU-12-31-38-00-8D-91 demo]$
```

Monitor the running jobs on the execute node:

```

test.user@domU-12-31-38-00-8D-91:~/demo -- ssh -- 80x24
[test.user@domU-12-31-38-00-8D-91 demo]$ condor_q|more

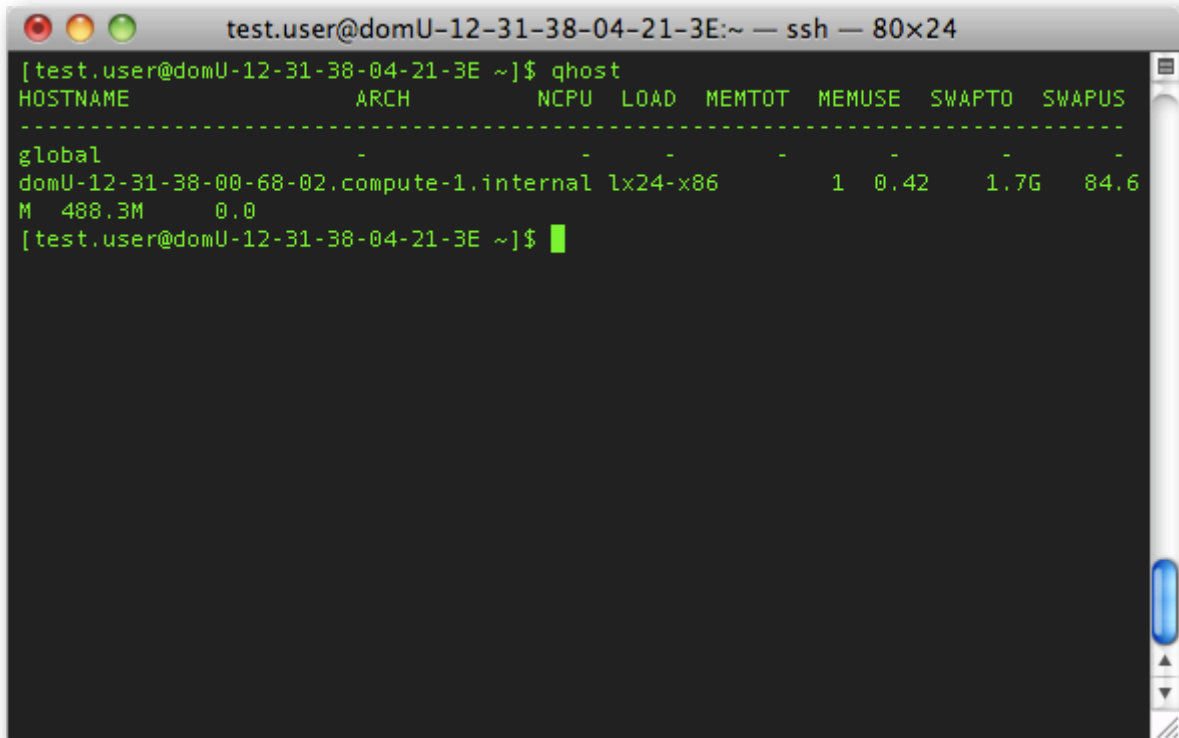
-- Submitter: domU-12-31-38-00-8D-91.compute-1.internal : <10.252.146.95:57546>
: domU-12-31-38-00-8D-91.compute-1.internal
ID      OWNER      SUBMITTED  RUN_TIME ST PRI  SIZE CMD
1.0     test.user  3/8 14:40  0+00:01:21 R 0   0.0 pi.py 10000
1.1     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.2     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.3     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.4     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.5     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.6     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.7     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.8     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.9     test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.10    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.11    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.12    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.13    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.14    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.15    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.16    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000
1.17    test.user  3/8 14:40  0+00:00:00 I 0   0.0 pi.py 10000

```

Further information at [http://www.cs.wisc.edu/condor/manual/v7.4/2\\_Users\\_Manual.html](http://www.cs.wisc.edu/condor/manual/v7.4/2_Users_Manual.html)

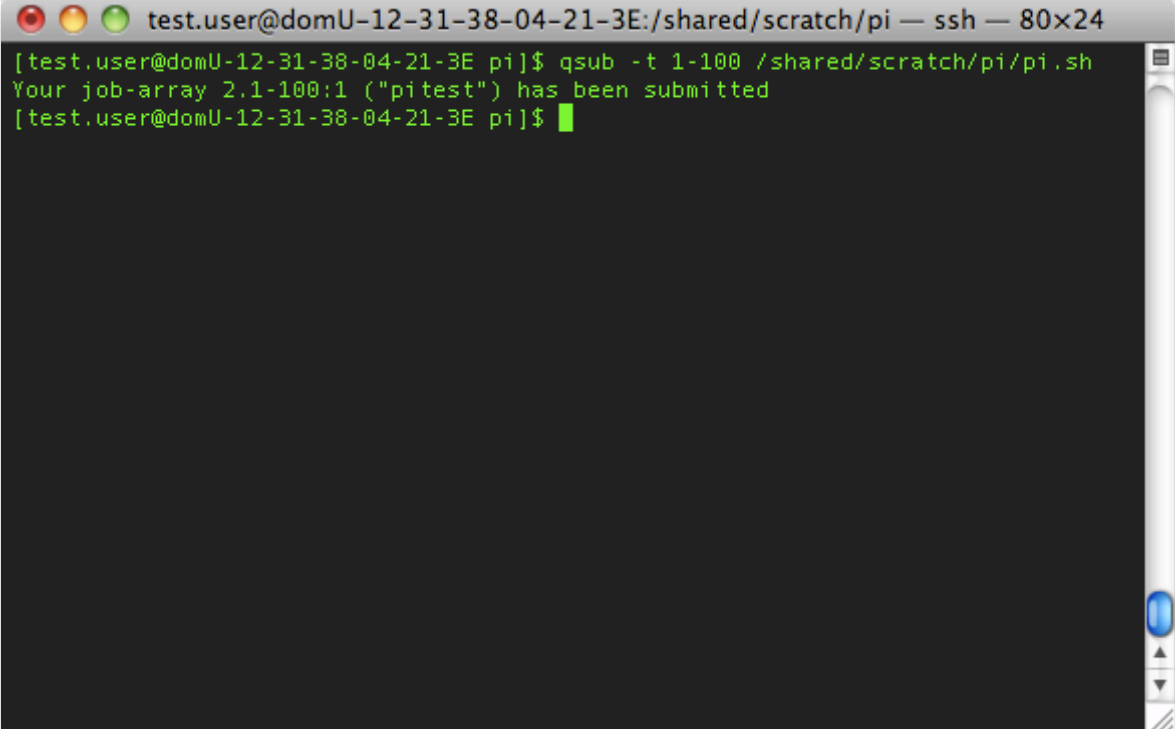
## SGE

Examine the execute node available for jobs:



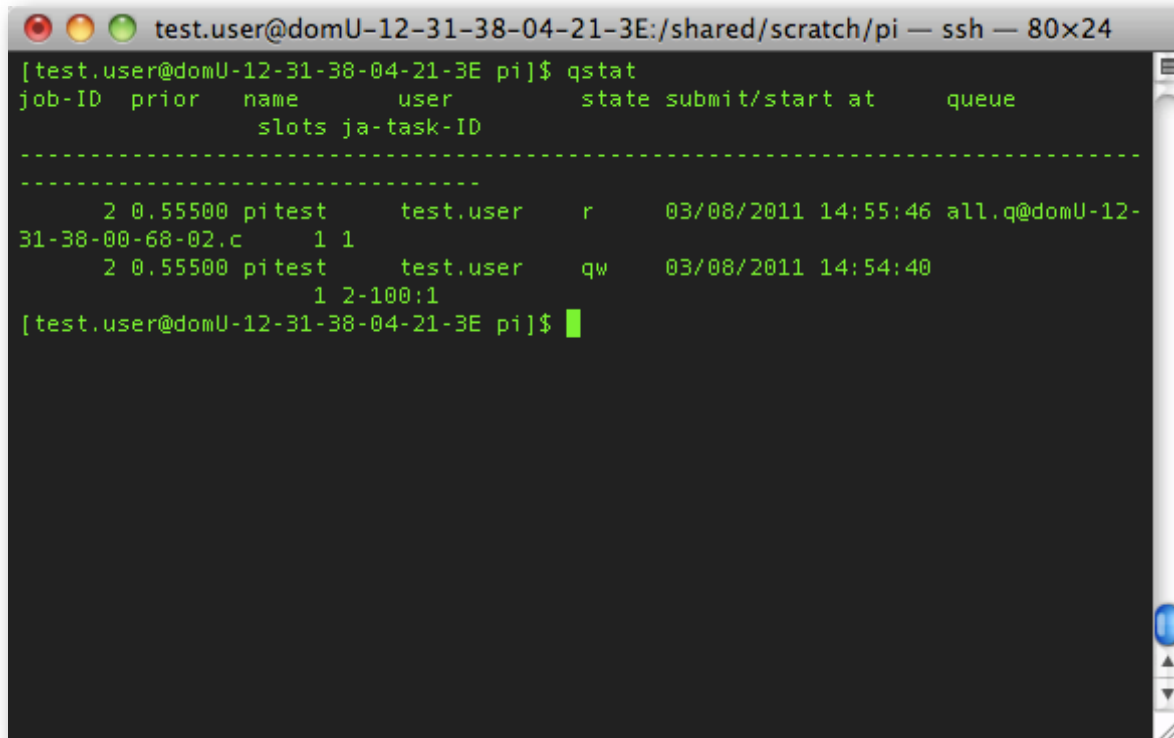
```
test.user@domU-12-31-38-04-21-3E:~ — ssh — 80x24
[test.user@domU-12-31-38-04-21-3E ~]$ qhost
HOSTNAME          ARCH          NCPU  LOAD  MEMTOT  MEMUSE  SWAPTO  SWAPUS
-----
global            -             -      -     -        -        -        -
domU-12-31-38-00-68-02.compute-1.internal 1x24-x86      1  0.42   1.7G   84.6
M  488.3M    0.0
[test.user@domU-12-31-38-04-21-3E ~]$
```

Submit 100 jobs:



```
test.user@domU-12-31-38-04-21-3E:/shared/scratch/pi — ssh — 80x24
[test.user@domU-12-31-38-04-21-3E pi]$ qsub -t 1-100 /shared/scratch/pi/pi.sh
Your job-array 2.1-100:1 ("pitest") has been submitted
[test.user@domU-12-31-38-04-21-3E pi]$
```

Monitor the running jobs on the execute node:

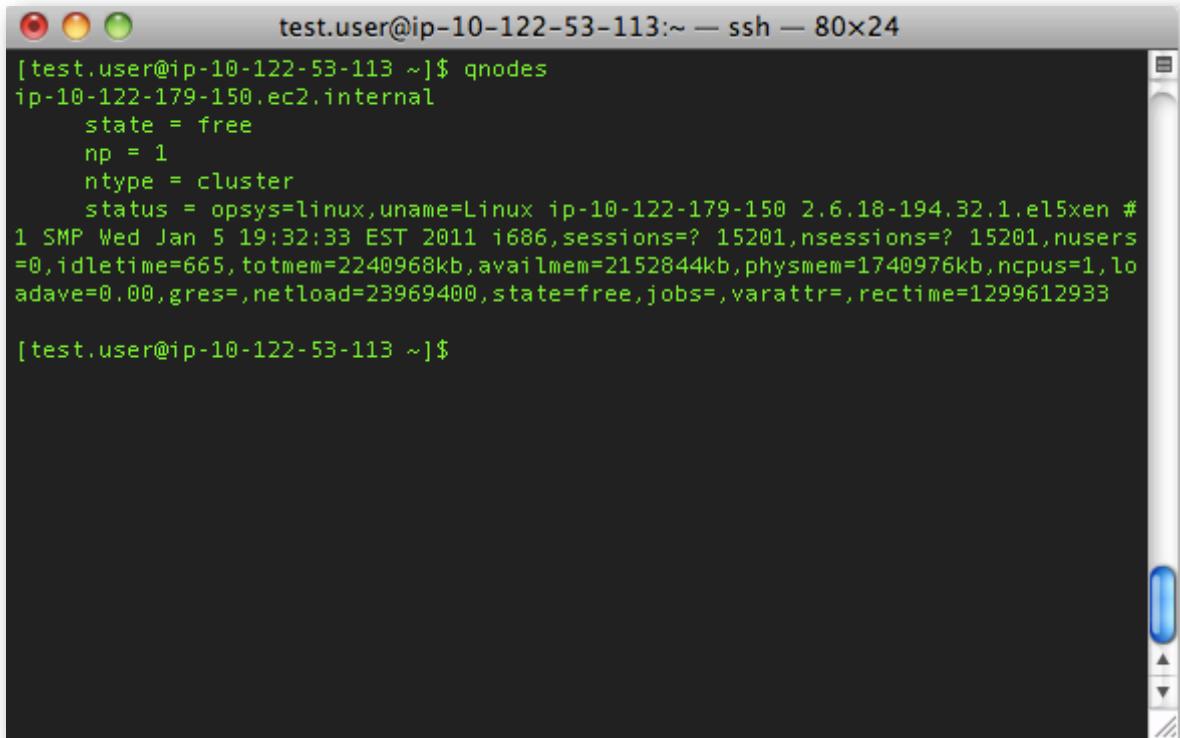


```
test.user@domU-12-31-38-04-21-3E:/shared/scratch/pi — ssh — 80x24
[test.user@domU-12-31-38-04-21-3E pi]$ qstat
job-ID  prior  name          user          state submit/start at    queue
      slots ja-task-ID
-----
      2 0.55500 pittest      test.user     r      03/08/2011 14:55:46 all.q@domU-12-
31-38-00-68-02.c  1 1
      2 0.55500 pittest      test.user     qw     03/08/2011 14:54:40
      1 2-100:1
[test.user@domU-12-31-38-04-21-3E pi]$
```

Further information at <http://wikis.sun.com/display/gridengine62u5/Using>

## Torque

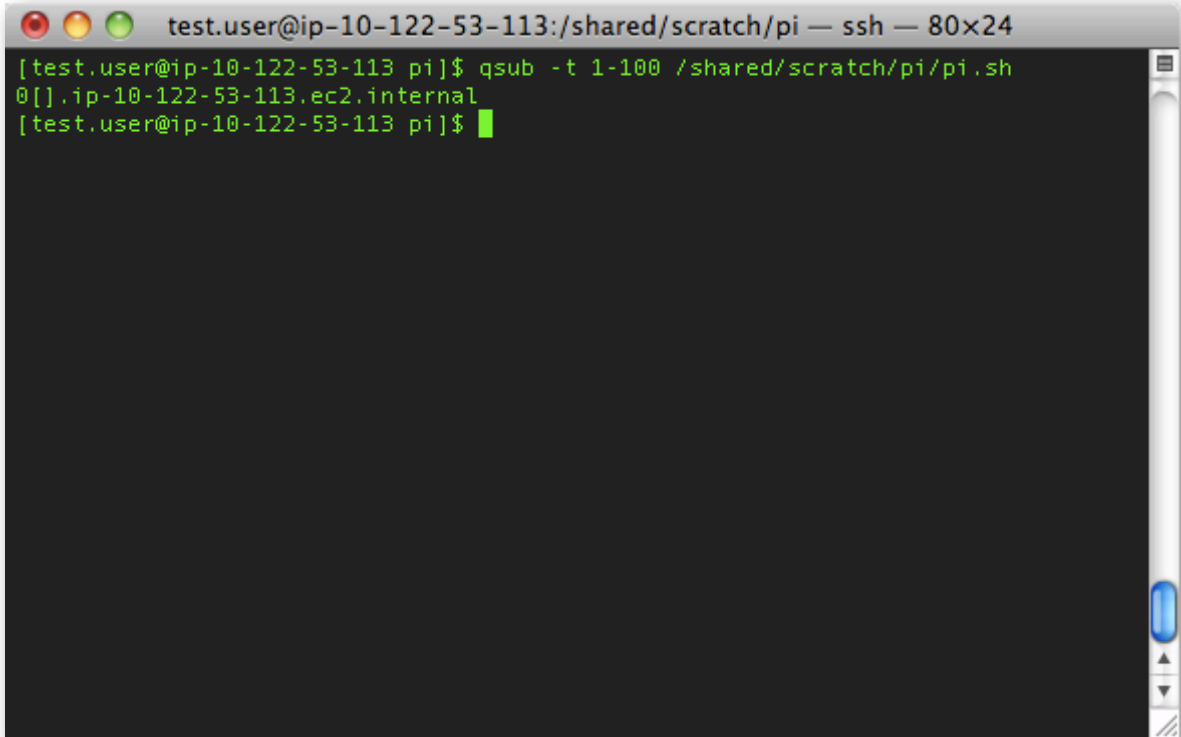
Examine the execute node available for jobs:



```
test.user@ip-10-122-53-113:~ — ssh — 80x24
[test.user@ip-10-122-53-113 ~]$ qnodes
ip-10-122-179-150.ec2.internal
  state = free
  np = 1
  ntype = cluster
  status = opsys=linux,uname=Linux ip-10-122-179-150 2.6.18-194.32.1.el5xen #
1 SMP Wed Jan 5 19:32:33 EST 2011 i686,sessions=? 15201,nsessions=? 15201,nusers
=0,idletime=665,totmem=2240968kb,availmem=2152844kb,physmem=1740976kb,ncpus=1,lo
adave=0.00,gres=,netload=23969400,state=free,jobs=,varattr=,rectime=1299612933

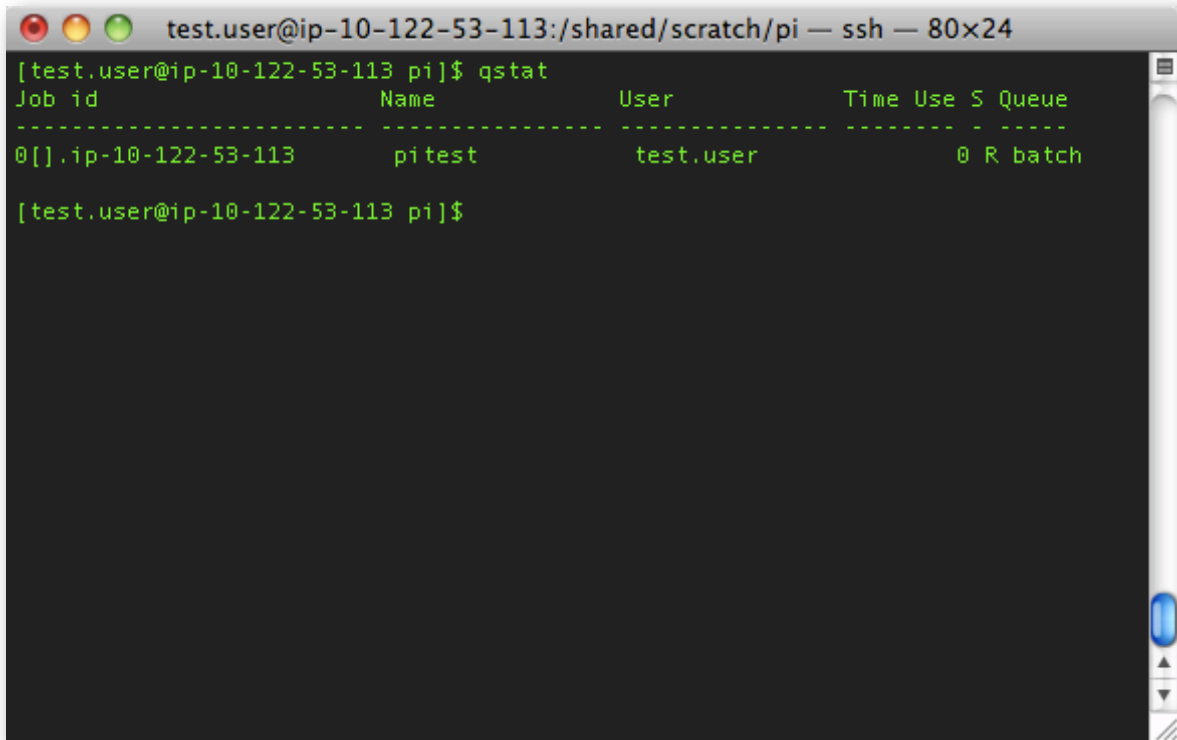
[test.user@ip-10-122-53-113 ~]$
```

Submit 100 jobs:

A terminal window with a dark background and light green text. The window title is "test.user@ip-10-122-53-113:/shared/scratch/pi — ssh — 80x24". The terminal shows the command "qsub -t 1-100 /shared/scratch/pi/pi.sh" being executed, followed by the output "0[] .ip-10-122-53-113.ec2.internal" and a new prompt line. The window has standard macOS-style window controls (red, yellow, green buttons) at the top left and a vertical scrollbar on the right side.

```
test.user@ip-10-122-53-113:/shared/scratch/pi — ssh — 80x24
[test.user@ip-10-122-53-113 pi]$ qsub -t 1-100 /shared/scratch/pi/pi.sh
0[] .ip-10-122-53-113.ec2.internal
[test.user@ip-10-122-53-113 pi]$
```

Monitor the running jobs on the execute node:

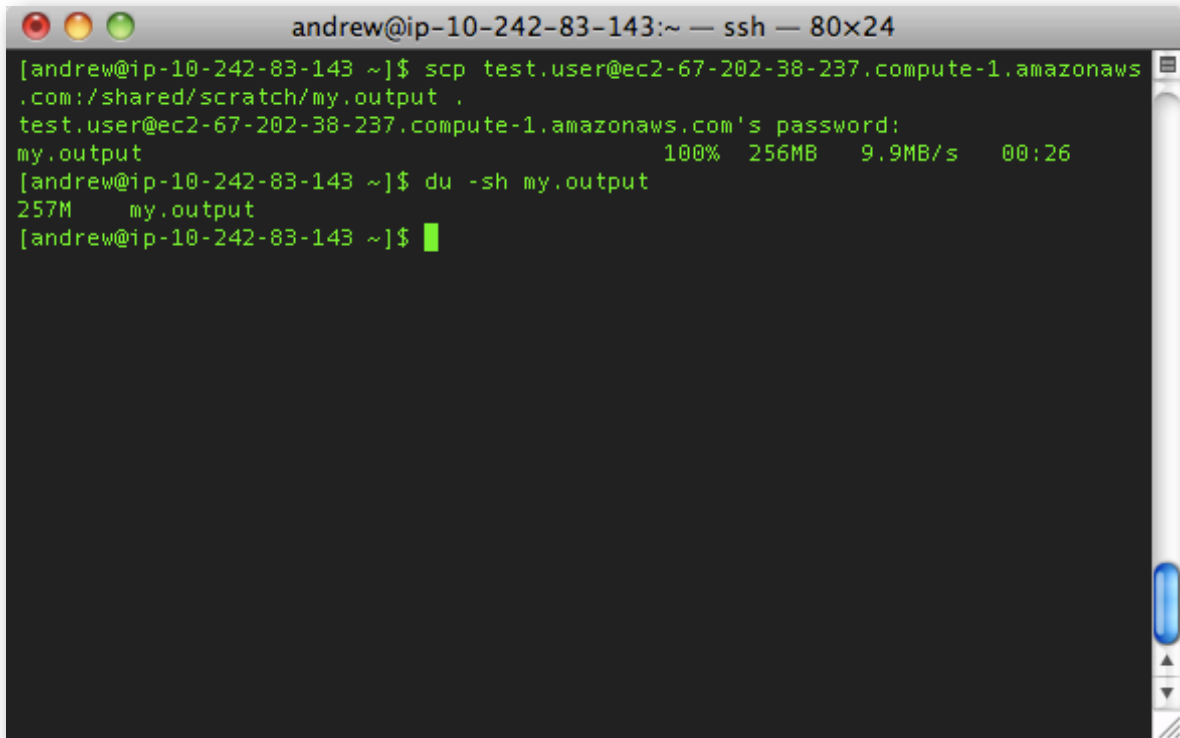


```
test.user@ip-10-122-53-113:/shared/scratch/pi — ssh — 80x24
[test.user@ip-10-122-53-113 pi]$ qstat
Job id          Name          User          Time Use S Queue
-----
0[] .ip-10-122-53-113  pitest       test.user          0 R batch
[test.user@ip-10-122-53-113 pi]$
```

Further information at <http://www.clusterresources.com/torquedocs21/users/2.1jobsubmission.shtml>

## Downloading Results


Use SCP to download the results from the shared filesystem on the scheduler:



```
andrew@ip-10-242-83-143:~ — ssh — 80x24
[andrew@ip-10-242-83-143 ~]$ scp test.user@ec2-67-202-38-237.compute-1.amazonaws.com:/shared/scratch/my.output .
test.user@ec2-67-202-38-237.compute-1.amazonaws.com's password:
my.output 100% 256MB 9.9MB/s 00:26
[andrew@ip-10-242-83-143 ~]$ du -sh my.output
257M    my.output
[andrew@ip-10-242-83-143 ~]$
```

## Stopping Your Cluster

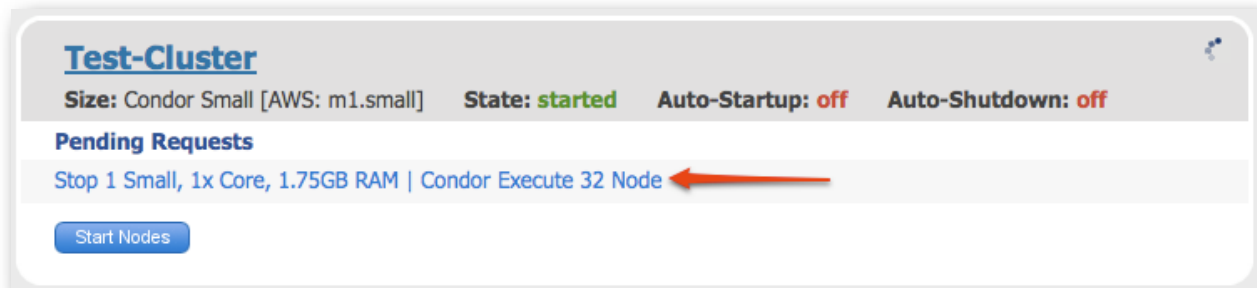
Once you have finished running jobs and have downloaded your data you can terminate the cluster. You can shut down individual nodes by clicking on the 'Stop' link that appears when you hover over the node you would like to stop.




Choose the number of nodes you want to stop and click the 'Stop' button.



This will show all the pending stop requests.



To terminate the cluster you can click on the 'Stop' link that appears when you hover over the cluster's name.



Once you have confirmed your intentions to stop the cluster the remaining instances in the cluster will be terminated.

**Test-Cluster** 

**Size:** Condor Small [AWS: m1.small]    **State:** **stopping**    **Auto-Startup:** **off**    **Auto-Shutdown:** **off**

**Pending Requests**

Stop 1 Small, 1x Core, 1.75GB RAM | Condor Manager 32 Node


Stop 1 Small, 1x Core, 1.75GB RAM | Condor Scheduler 32 Node

Once the instances have stopped the cluster will be moved into a 'stopped' state. A stopped cluster will not have any instance charges but it will accrue charges for persistent storage attached to the cluster and any static IP addresses that may have been used.

**Test-Cluster**

**Size:** Condor Small [AWS: m1.small]    **State:** **stopped**    **Auto-Startup:** **off**    **Auto-Shutdown:** **off**

You can delete the cluster by clicking the garbage can icon. A deleted cluster will destroy any persistent storage volumes so please be certain that you have downloaded all your results. It will also release any static IP addresses associated with the cluster. A deleted cluster is not charged for anything.

**Start Options** 

**Test-Cluster**

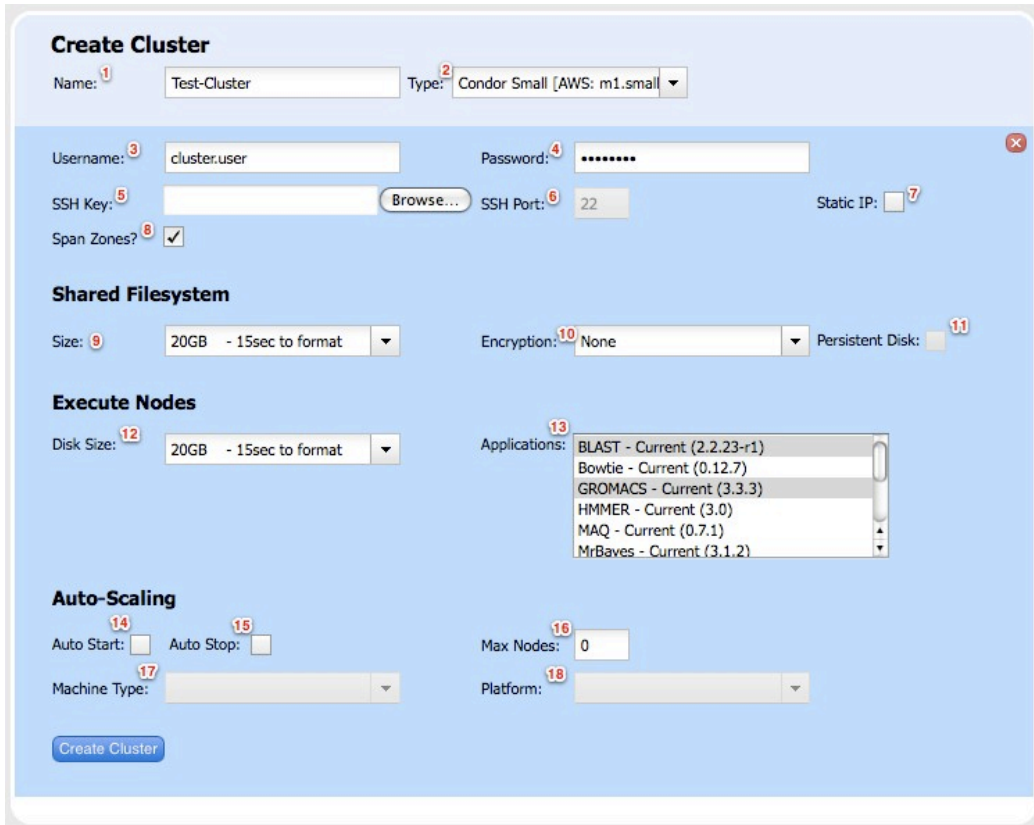
**Size:** Condor Small [AWS: m1.small]    **State:** **stopped**    **Auto-Startup:** **off**    **Auto-Shutdown:** **off**

## Billing

You can access the [billing page](#) by clicking on the billing tab. The billing page has a breakdown of your usage for each month itemized by cluster. You can access the [pricing page](#) to see the costs of the different cluster types and the individual node types.

## Cluster Configuration Options

There are quite a few options that can be set during cluster creation. We'll go over the current options below.



The screenshot shows a 'Create Cluster' form with the following fields and options:

- Name:** 1. Text input field containing 'Test-Cluster'.
- Type:** 2. Dropdown menu showing 'Condor Small [AWS: m1.small]'.
- Username:** 3. Text input field containing 'cluster.user'.
- Password:** 4. Password input field with masked characters.
- SSH Key:** 5. Text input field with a 'Browse...' button.
- SSH Port:** 6. Text input field containing '22'.
- Static IP:** 7. Text input field.
- Span Zones?** 8. Checked checkbox.
- Shared Filesystem:**
  - Size:** 9. Dropdown menu showing '20GB - 15sec to format'.
  - Encryption:** 10. Dropdown menu showing 'None'.
  - Persistent Disk:** 11. Checked checkbox.
- Execute Nodes:**
  - Disk Size:** 12. Dropdown menu showing '20GB - 15sec to format'.
  - Applications:** 13. List box containing: BLAST - Current (2.2.23-r1), Bowtie - Current (0.12.7), GROMACS - Current (3.3.3), HMMER - Current (3.0), MAQ - Current (0.7.1), MrBaves - Current (3.1.2).
- Auto-Scaling:**
  - Auto Start:** 14. Unchecked checkbox.
  - Auto Stop:** 15. Unchecked checkbox.
  - Max Nodes:** 16. Text input field containing '0'.
  - Machine Type:** 17. Dropdown menu.
  - Platform:** 18. Dropdown menu.

A 'Create Cluster' button is located at the bottom left of the form.

1. **Name:** The name of the cluster. Spaces are not allowed.
2. **Type:** The type of the cluster. This sets the type and size of scheduler. The scheduler size is displayed in the square brackets.
3. **Cluster username:** The username that is used for SSH and CycleServer.
4. **Cluster password:** The password used for SSH and CycleServer.
5. **SSH Key:** Optionally upload an SSH key for accessing the scheduler without a password.

6. **SSH Port:** Can customize the ssh port. Not all accounts have this option enabled.
7. **Static IP:** If checked, a static IP address will be reserved and attached to the scheduler when started. There is a charge associated with checking this. Please check your pricing page to see how much this option may cost.
8. **Span Zones:** If checked, the nodes are allowed to span availability zones if necessary. If unchecked all instances will be forced into the same availability zone.
9. **Shared Filesystem Size:** The size of the shared filesystem. A larger filesystem will take longer to format and thus the cluster will take longer to boot.
10. **Encryption:** Type of encryption to use on the filesystem. Using encryption can increase boot time.
11. **Persistent Disk:** If checked, the shared file system will be installed on a persistent volume that persists even when the cluster is stopped. There is an additional charge for using persistent storage and not all accounts will have access to this option.
12. **Execute Node Disk Size:** The size of the ephemeral disk to format on the execute nodes. Larger disks take longer to boot.
13. **Applications:** The applications that will be installed on the scheduler and execute nodes. Can select multiple applications by holding down control/ command.
14. **Auto-start:** If selected, the default nodes will be automatically started to deal with jobs.
15. **Auto-stop:** If selected, nodes will be stopped when there is no longer jobs.
16. **Max-nodes:** the maximum number of auto-start nodes that should be started to deal with idle jobs.
17. **Machine type:** the type of machine to use for auto-start nodes.
18. **Platform:** the platform type to be used for auto-start nodes.

