DiCloud user guide

Emmanuel Cecchet
Oct 26, 2010

1. What is DiCloud?

DiCloud is a testbed that allows users to **experiment** with Amazon Web Service cloud resources in the context of the National Science Foundation\'s Global Environment for Network Innovations (GENI) project. DiCloud interfaces with a prototype GENI control framework. More information about the DiCloud project is available here.

1.1 How does DiCloud work?

DiCloud gives the user access to the following resources:

- Elastic Compute Cloud (EC2) instances: servers with local storage that are accessible through the Internet via a public IP address and from other EC2 instances via a local private IP address.
- Simple Storage Service (S3) buckets: a bucket is a storage space varying from 1 byte to 5 GB that can store objects. The objects can only be accessed through a put/get interface.
- Elastic Block Storage (EBS) volumes: EBS is a Network Attach Storage (NAS) that can be mounted and accessed by only one EC2 instance at a time. Any filesystem can be created on an EBS volume.

A limited budget is allocated for the resources that are monitored in real time using Amazon CloudWatch. Once the allocated budget has expired, resources are automatically revoked and experiments are stopped until further funding is available.

1.2 How do I get an account?

Send an email to **dicloud@cs.umass.edu** to request an account.

2. First login

The DiCloud home page is located at http://geni.cs.umass.edu/vise/dicloud.php

Use the login and password that has been provided by the administrator.

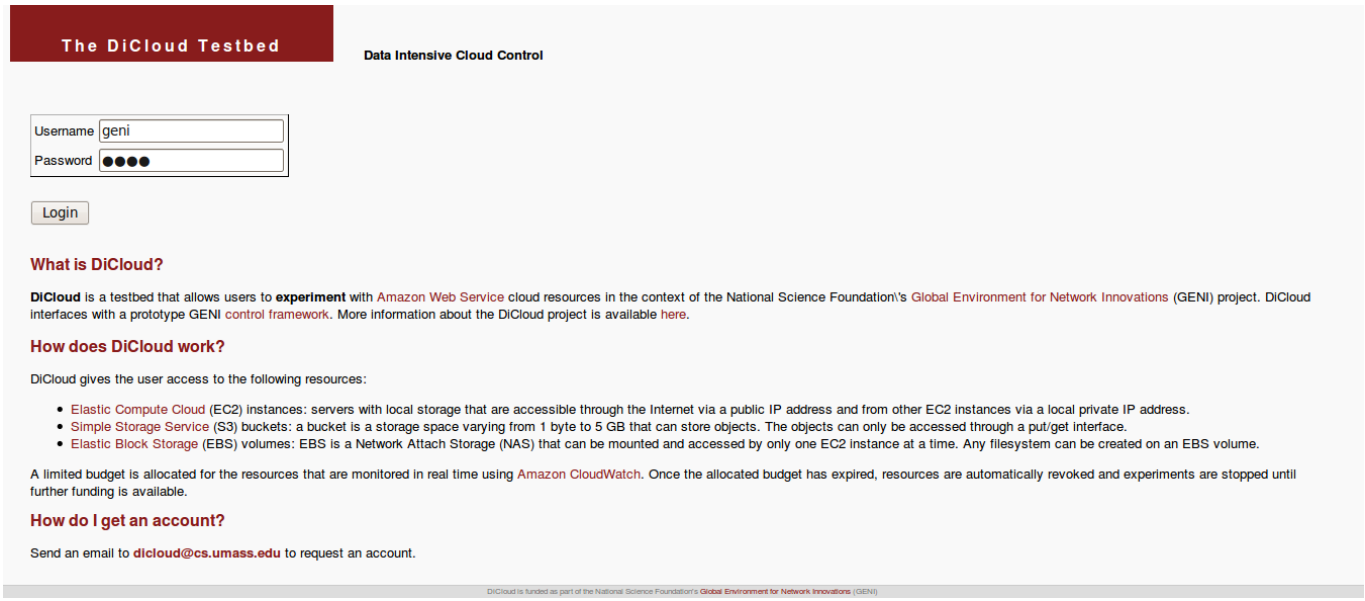

If your login is successful, you will reach your home page that looks like:

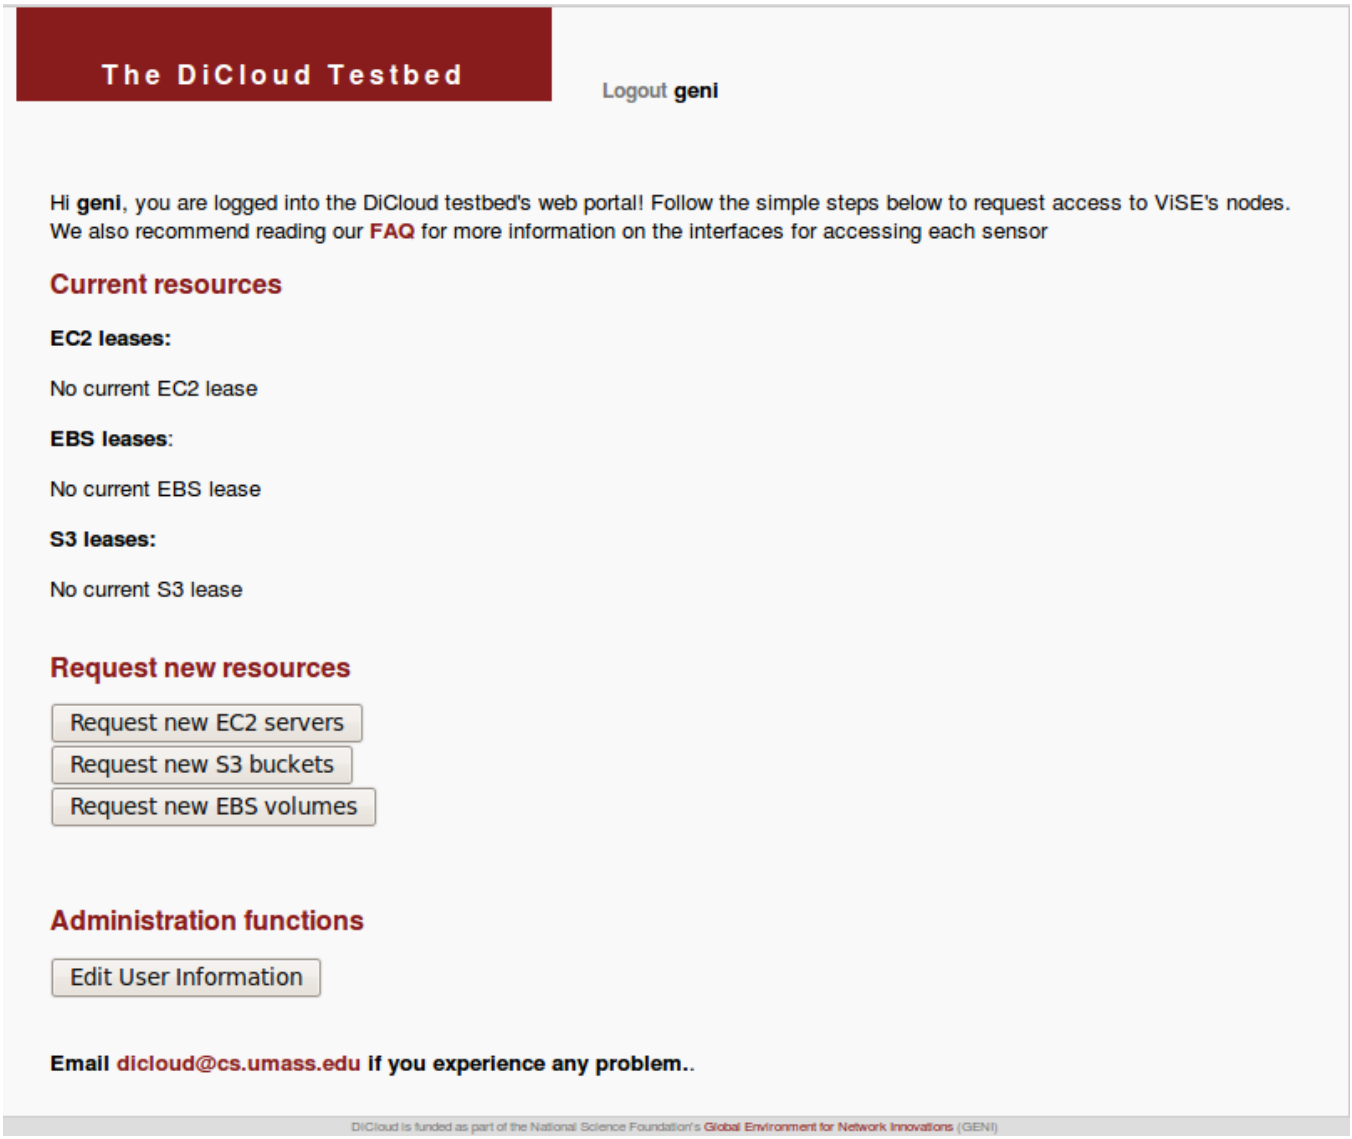

2.1 Update your password

Click on the ‘Edit User Information’ button from your home page.
Then click on the ‘Change password’ link next to your user name to update your password.
If you don’t have administrator rights, you will have to logout and log in again with your new password.

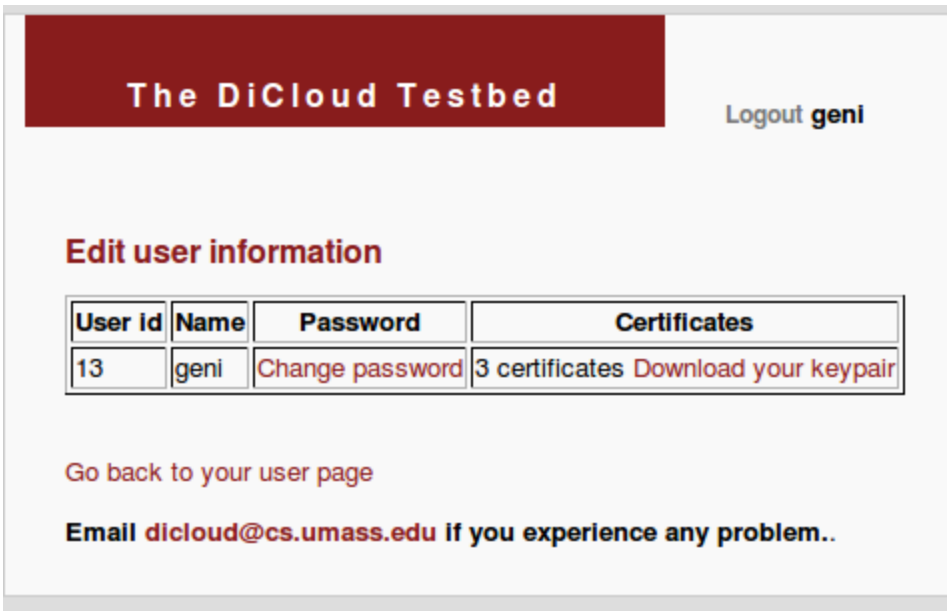

2.2 Download your keypair

The keypair is a key that is need to SSH into EC2 instances. Each instance that will be launched on your behalf will only be accessible using your keypair.
Click on the 'Edit User Information' button from your home page.
Then click on the Download your keypair' link next to your user name. A page similar to the following one will appear:

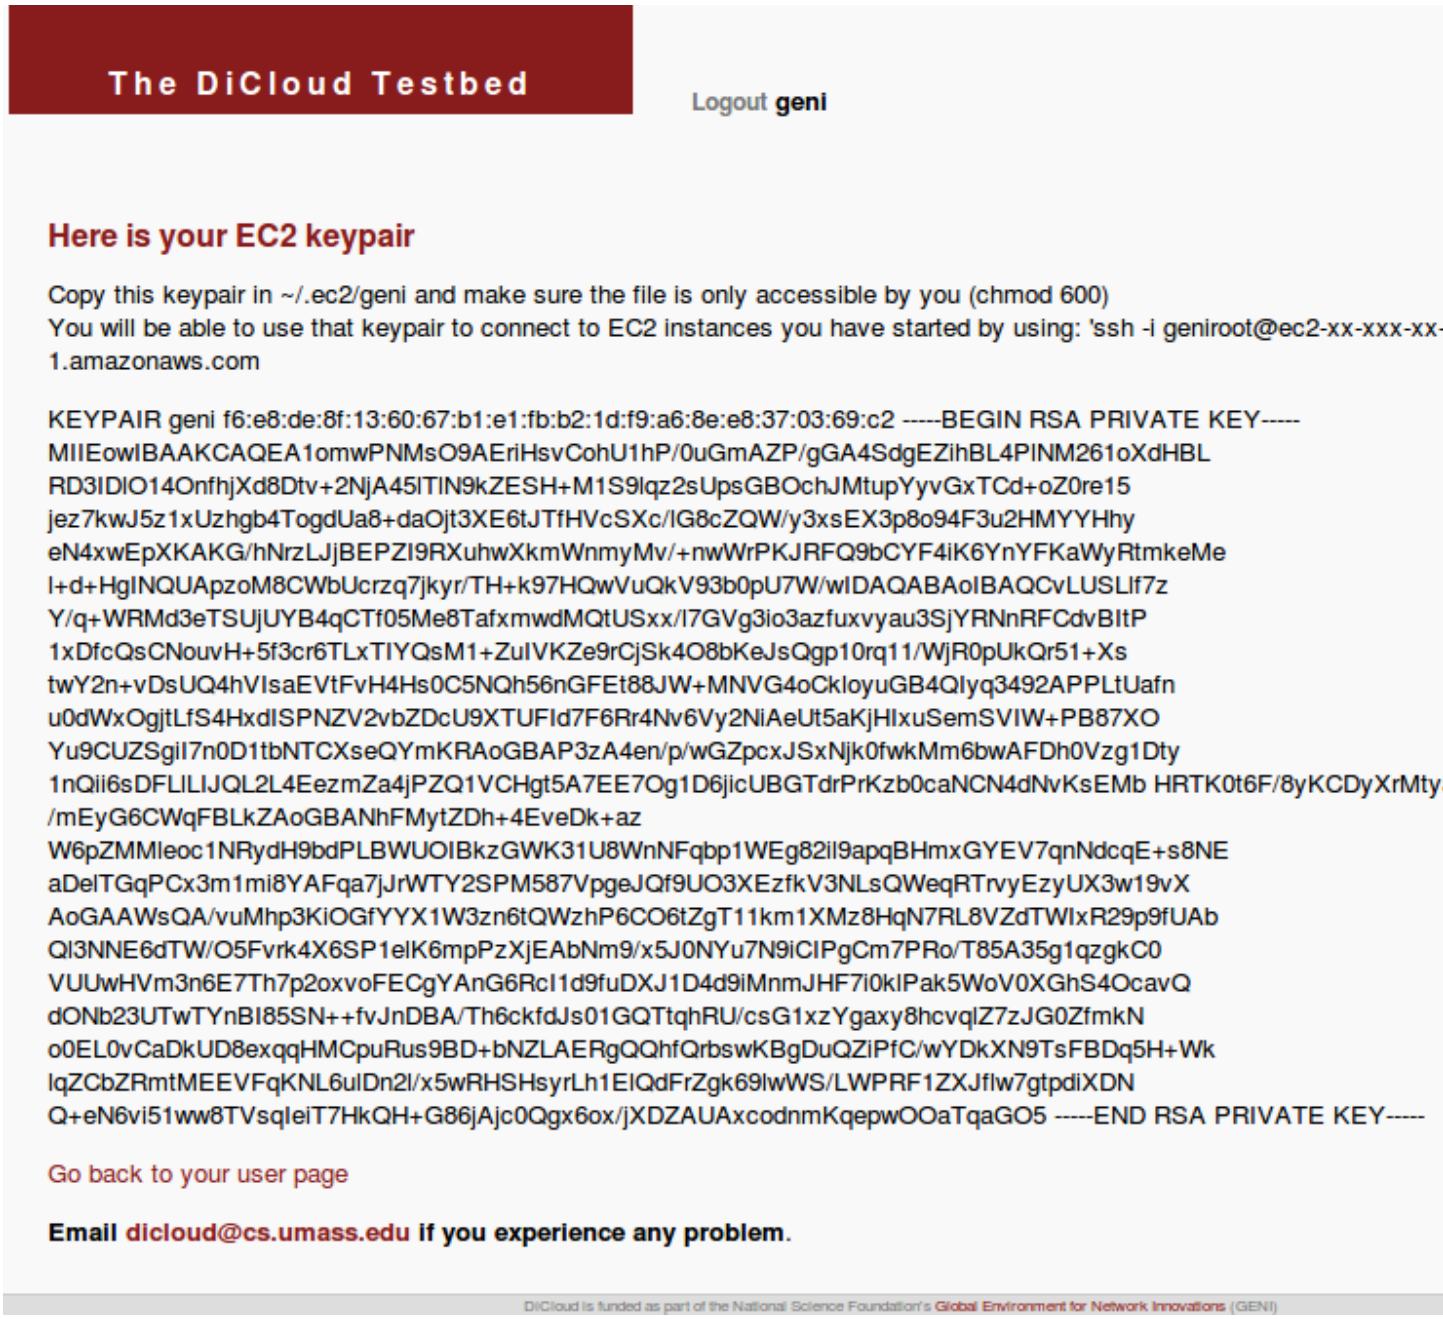

Copy/paste this entire key, including the -----BEGIN RSA PRIVATE KEY----- and -----END RSA PRIVATE KEY----- lines into ~/.ec2/yourlogin. Next, it is important to change the permissions of your keypair file, or else EC2 will not let you connect to it via SSH. To do this, just type the following in your /.ec2 directory:
$ chmod 600 ~/.ec2/yourlogin

You will be able to log into your EC2 instances using a command line like:
ssh -i yourlogin root@ec2-xx-xxx-xx-xx.compute-1.amazonaws.com

Note that the -i parameter is the name of the keypair (which is your login) and not the name of the file. If you try to give the path to the certificate file, ssh will fail.

3. Using EC2 resources

To use EC2 servers, simply click on the ‘Request new EC2 servers’ on your home page. Select the kind and number of instances you need as follows:

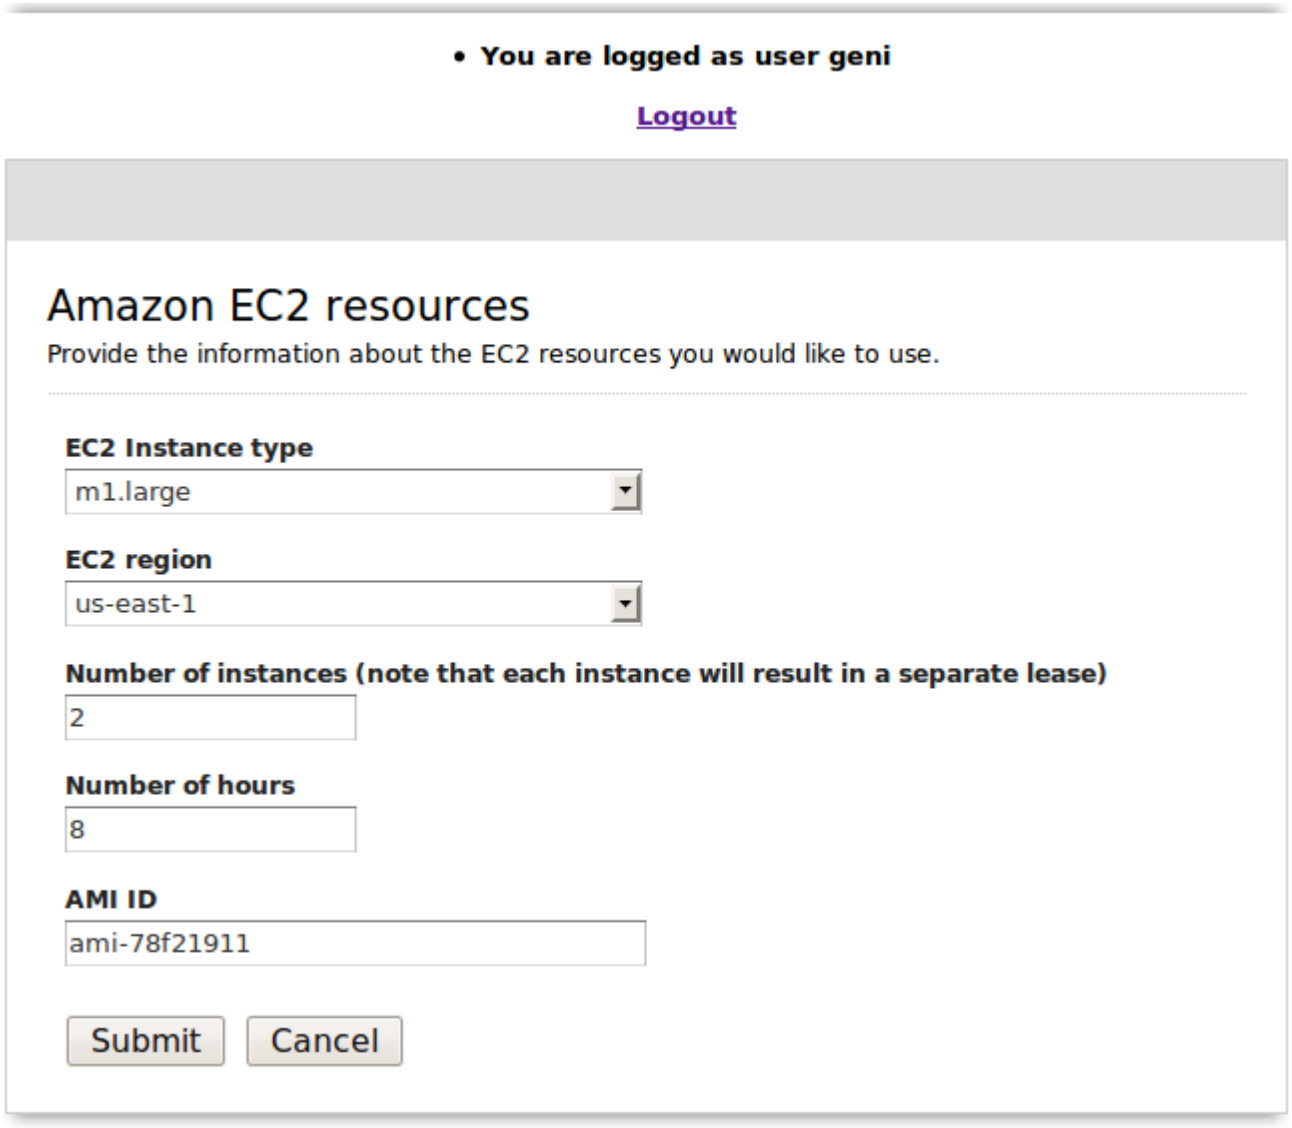

Leases are granted immediately if the budget allows it. Resources are automatically reclaimed at the end of the lease or when the budget has been used, whichever comes first. This means that servers are automatically shutdown when the lease is over or the budget is used.

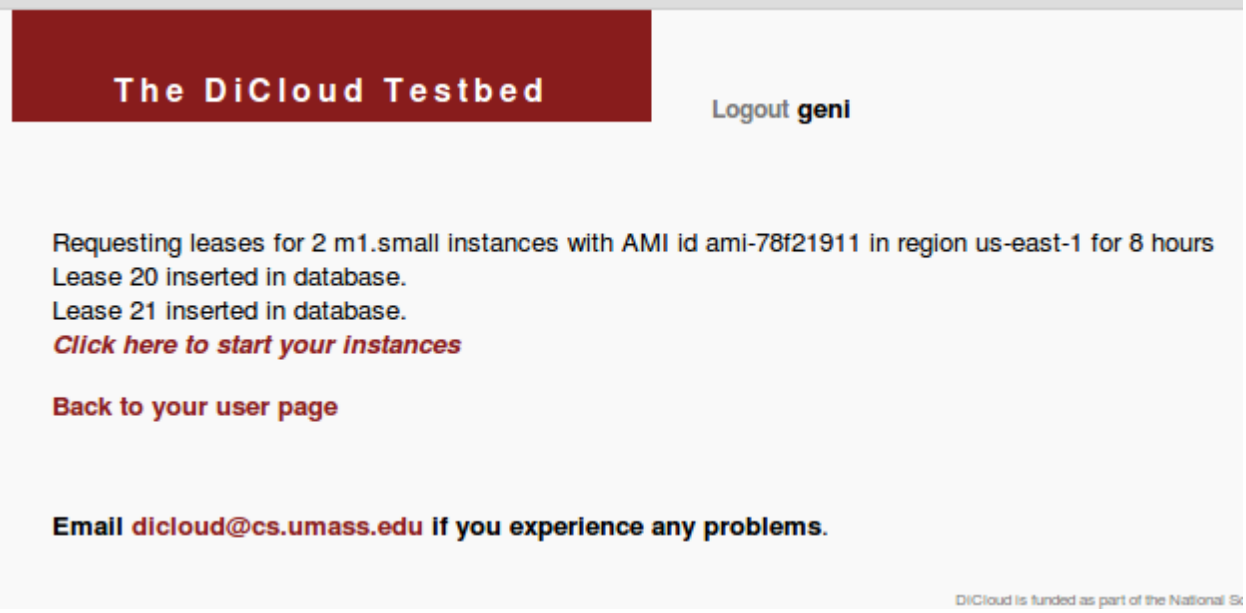

To start your server instances, you need to click on the link **'Click here to start your instances**'. The leases will then be processed and the instances will be started as needed. The EC2 server will the be started and an instance id will be assigned to the server:

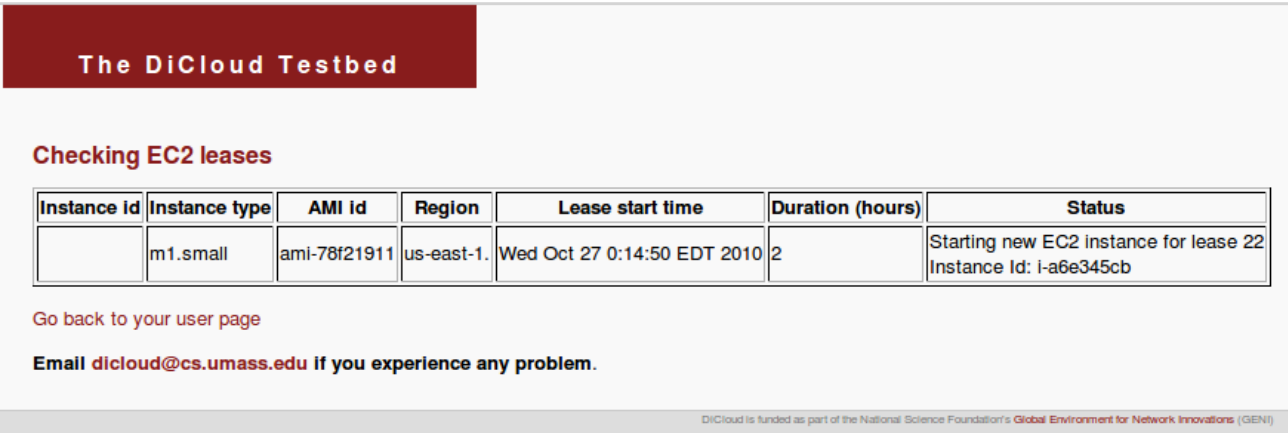

It usually takes a while for EC2 servers to start. You can check the status on the instance from your home page by clicking on the 'Check' link of the EC2 instance.

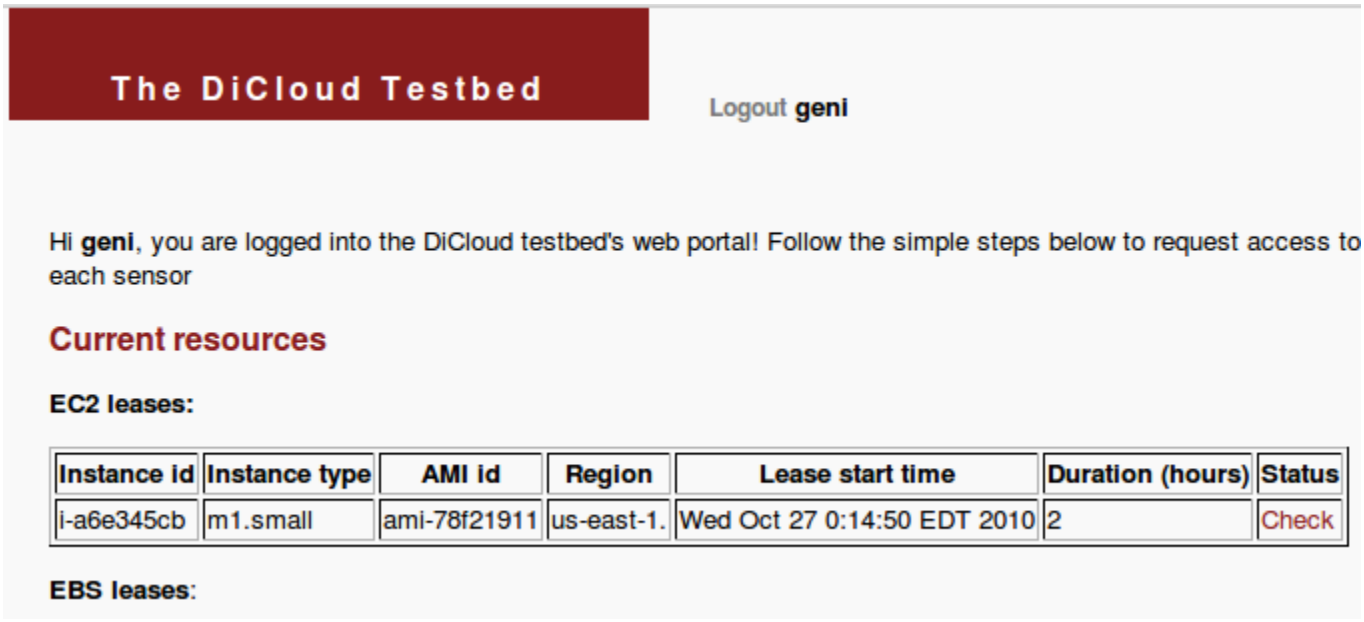

When an instance has been successfully started, you will obtain its IP address in the page generated by following the ‘Check’ link. Here is an example of the resulting page:

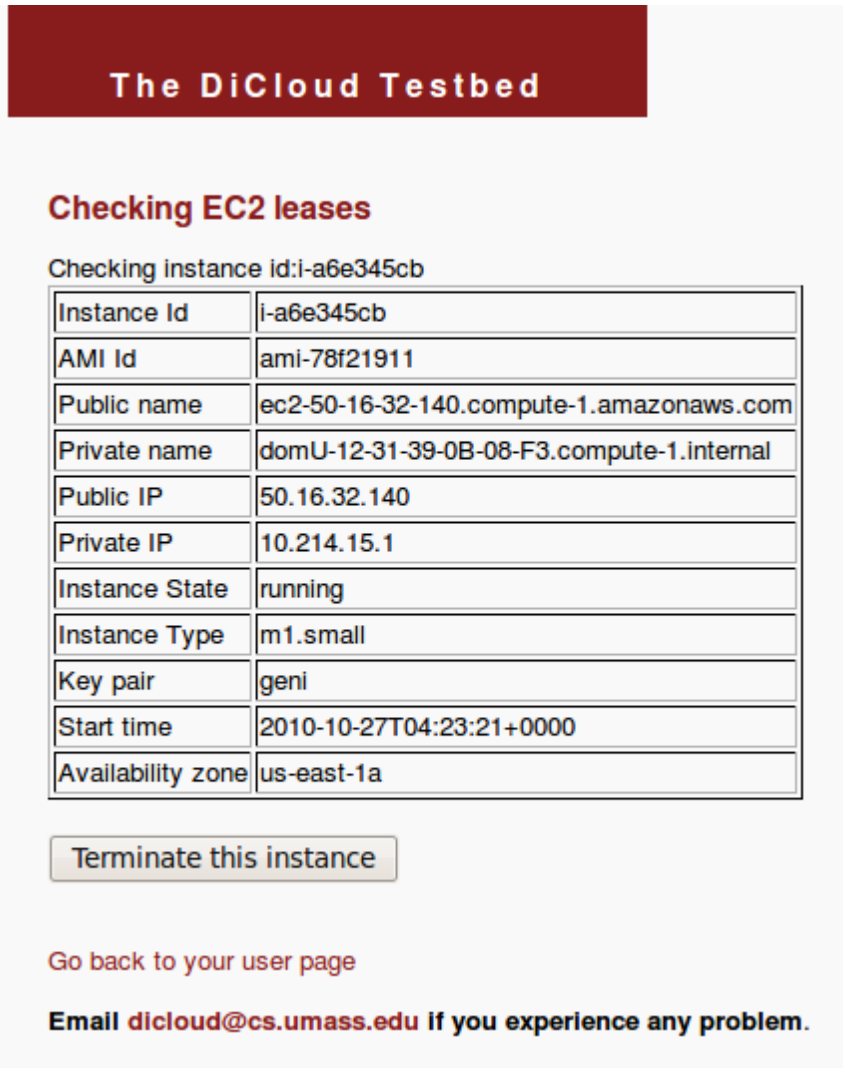

You can connect to the instance above using your keypair (see section 2.2) with a command line like:
ssh -i yourlogin root@ec2-50-16-32-140.compute-1.amazonaws.com
or
ssh -i yourlogin root@50.16.32.140

The private IP can be used to communicate with other EC2 instances.

You can also force to terminate this instance (kill the virtual machine) by clicking the 'Terminate this instance' button. You will be able to restart the instance as long as the lease has not expired by clicking the '**Process your leases**' button from your home page.

4. EBS resources

EBS is a network attached storage. It provides a block device that can be attached to a running

instance. To create a new EBS volume, click on the '**Request new EBS volumes**' from your home page.

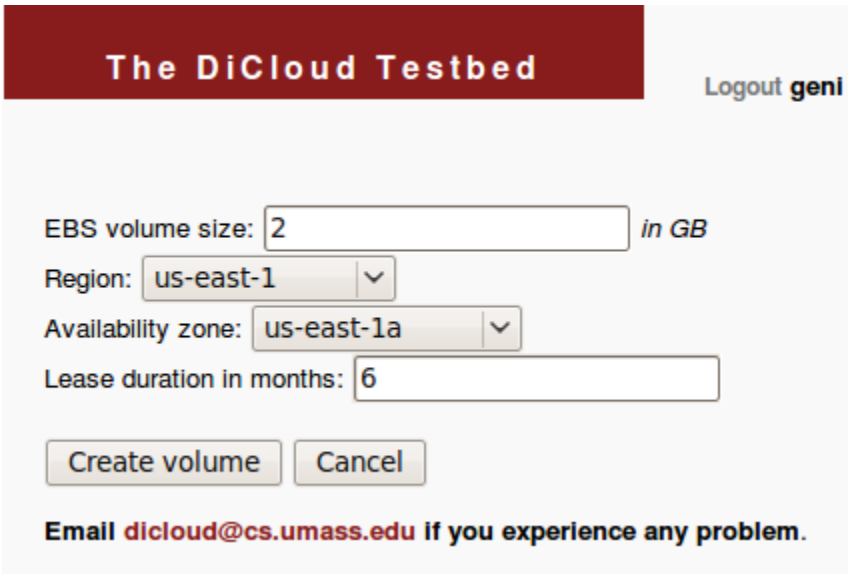

Note that EBS volumes can only be attached to EC2 instances running in the same availability zone of the same region.

Volumes are created as soon as you click the 'Create volume' button.

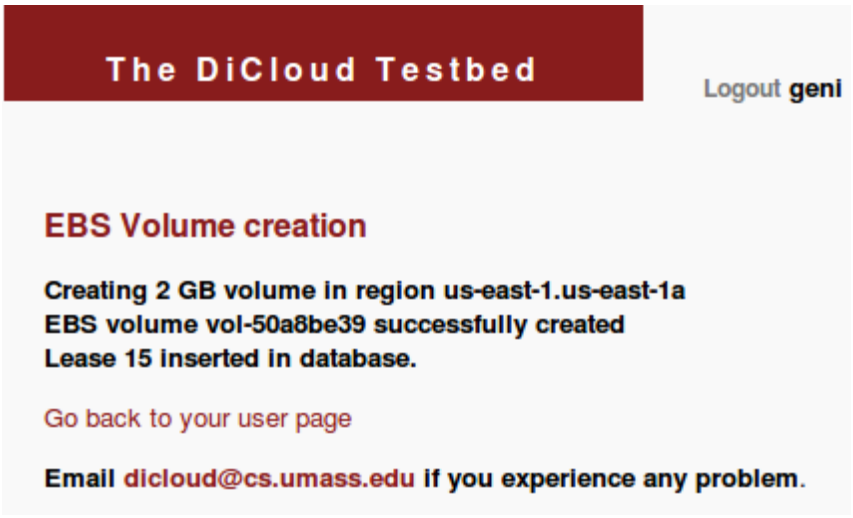

As soon as the volume is created you can attach it to a running instance. You can check the status of an EBS volume from your home page by clicking on the 'Check' link next to the instance.

Current resources

EC2 leases:

Instance id	Instance type	AMI id	Region	Lease start time	Duration (hours)	Status
i-a6e345cb	m1.small	ami-78f21911	us-east-1.	Wed Oct 27 0:14:50 EDT 2010	2	Check

EBS leases:

Volume ID	Size (in GB)	Region	Date Created	Duration (months)	Operations
vol-50a8be39	2	us-east-1.us-east-1a	Wed Oct 27 0:38:46 EDT 2010	6	Check / Attach / Detach / Delete

Click ‘Attach’ to attach the volume to a running instance as follows:

The DiCloud Testbed

Logout **geni**

EBS Attach Volume

EC2 instance id to attach to: i-a6e345cb

Name to give to the device: /dev/hdd

Attach the volume

Current list of EC2 leases:

Instance id	Instance type	AMI id	Region	Lease start time	Duration (hours)	Status
i-a6e345cb	m1.small	ami-78f21911	us-east-1.	Wed Oct 27 0:14:50 EDT 2010	2	Check

Go back to your user page

Email dicloud@cs.umass.edu if you experience any problem.

You can check the status of the EBS volume to make sure it has been attached properly:

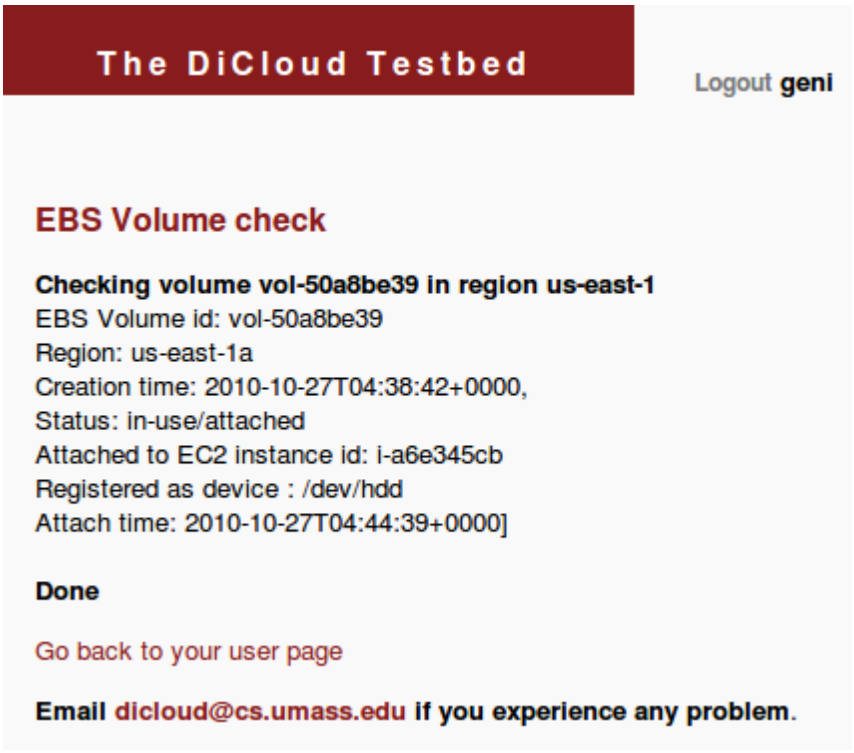

The volume can be detached before the instance it is attached to terminates by clicking the 'Detach' link.

The **'Delete'** link permanently destroys the EBS volume and the associated lease is immediately terminated.

5. S3 usage

S3 buckets can contain a large number of objects (files). The S3 bucket name has to be unique world-wide since it is registered in the Amazon DNS under bucket-name.s3.amazonaws.com.

New buckets can be allocated using the '**Request new S3 buckets**' from the user home page. We recommend to only use 1 bucket per user.

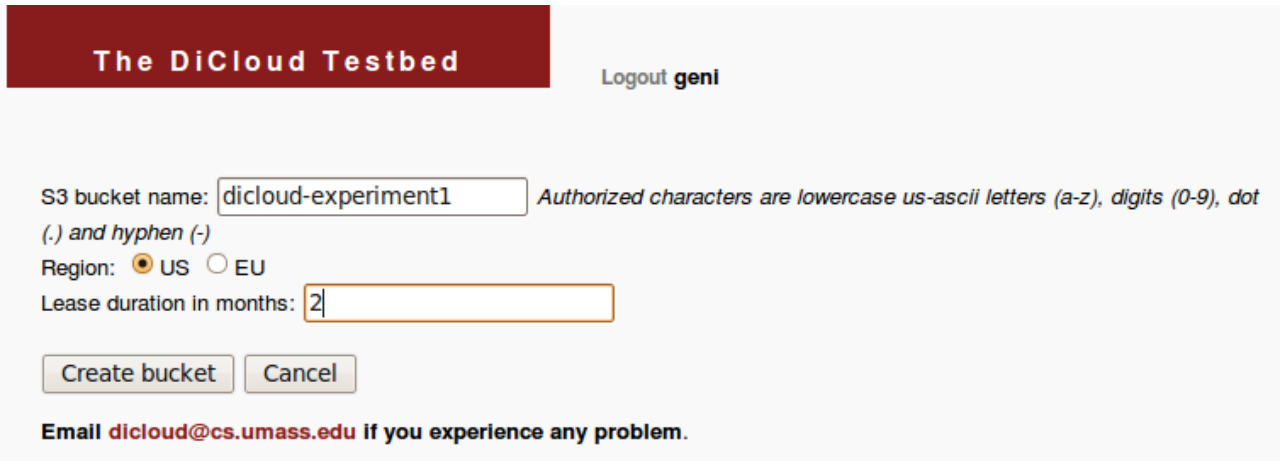

It is then possible to put or get files in the bucket using the appropriate links from the user home page:

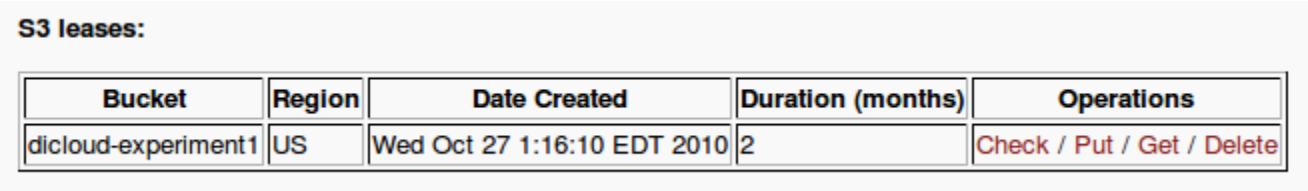

The put operation uploads a file to S3 as follows:

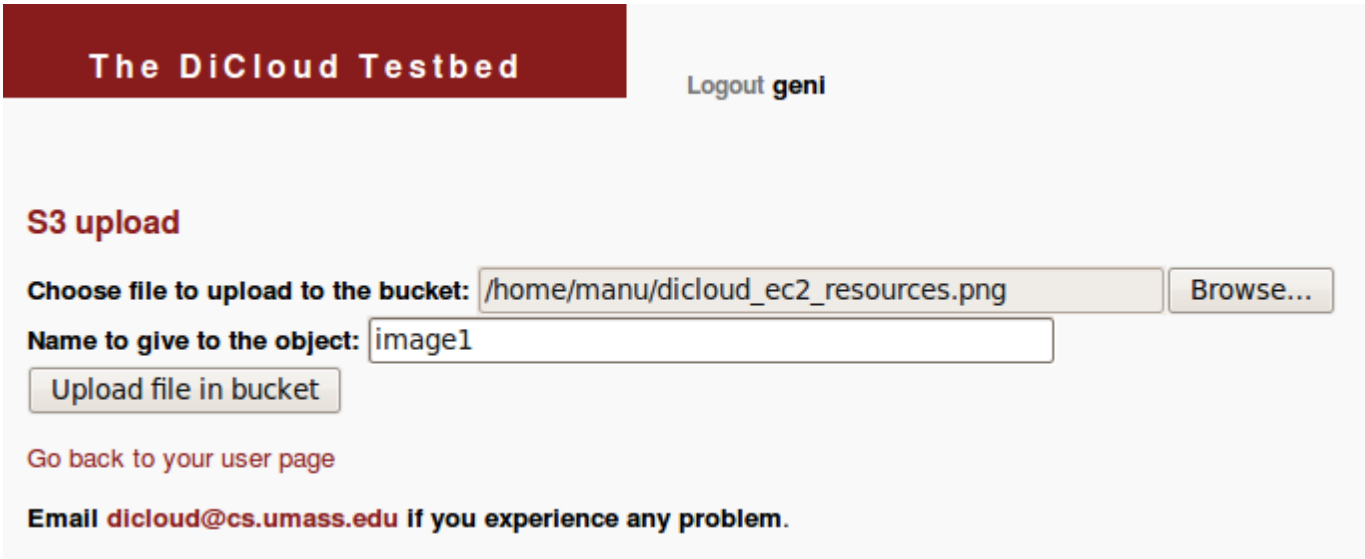

Note that the file is first uploaded to the DiCloud portal and then uploaded to S3. This is why you will see the name of a temporary file during the upload process:

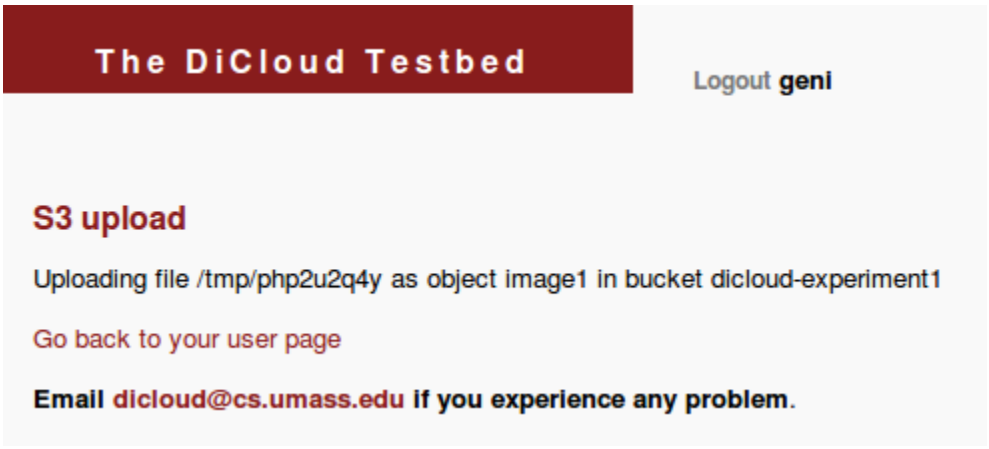

To retrieve a file from an S3 bucket, use the object name (that was given when you did put the file):

6. Pricing

Currently, only the administrator can see the budget expenses. Here are some indications of the cost of the different operations.

6.1. EC2 instance prices

Each instance is started with CloudWatch for monitoring. The current prices per hour (in USD) are as follows:

CLOUDWATCH = -0.015 per instance per hour.

m1.small,us-west-1=-.095
m1.large,us-west-1=-.38
m1.xlarge,us-west-1=-.76
m2.2xlarge,us-west-1=-.57
m2.xlarge,us-west-1=-1.34
m2.4xlarge,us-west-1=-2.68
c1.medium,us-west-1=-.19
c1.xlarge,us-west-1=-.76

m1.small,us-east-1=-.085
m1.large,us-east-1=-.34
m1.xlarge,us-east-1=-.68
m2.2xlarge,us-east-1=-.5
m2.xlarge,us-east-1=-1.2
m2.4xlarge,us-east-1=-2.4
c1.medium,us-east-1=-.17
c1.xlarge,us-east-1=-.68

m1.small,eu-west-1=-.095
m1.large,eu-west-1=-.38
m1.xlarge,eu-west-1=-.76
m2.2xlarge,eu-west-1=-.57
m2.xlarge,eu-west-1=-1.34
m2.4xlarge,eu-west-1=-2.68
c1.medium,eu-west-1=-.19
c1.xlarge,eu-west-1=-.76

m1.small,ap-southeast-1=-.095
m1.large,ap-southeast-1=-.38
m1.xlarge,ap-southeast-1=-.76
m2.2xlarge,ap-southeast-1=-.57
m2.xlarge,ap-southeast-1=-1.34
m2.4xlarge,ap-southeast-1=-2.68

c1.medium,ap-southeast-1=-.19
c1.xlarge,ap-southeast-1=-.76

The network traffic is charged per GB (to and from the Internet traffic have different costs):
EC2_NETWORK_IN_PER_GB = 0; (-0.10 after Nov 1, 2010)
EC2_NETWORK_OUT_PER_GB = -.15;

6.2. Other prices

S3 storage is billed per GB per month plus the network bandwidth at EC2 pricing. Note that downloading data from S3 has to go through the DiCloud proxy which requires in+out network traffic. Each put and get operation also has a nominal cost.
S3_STORAGE_PER_GB_PER_MONTH = -.15;
S3_PUT_OPERATION = -.01 / 1000;
S3_GET_OPERATION = -.01 / 10000;

EBS is charged per million IOs (read or write) on the volume plus the cost of storage per GB per month (same cost as S3)
EBS_IO_PER_MILLION = -.1;