



EC2 Cloud Spot Instances for ATLAS Jobs



Val Hendrix: vchendrix@lbl.gov

Lawrence Berkeley National Lab

Introduction

- Cost of ATLAS jobs on EC2 Cloud
- Amazon EC2 Spot Instances
- Reduce costs using Spots Instances
- Spot Instances for ATLAS Jobs

Cost Estimate (Cost per 1K Evt)



Assume we run 8 concurrent jobs for all cases, the cost per 1k event is calculated.

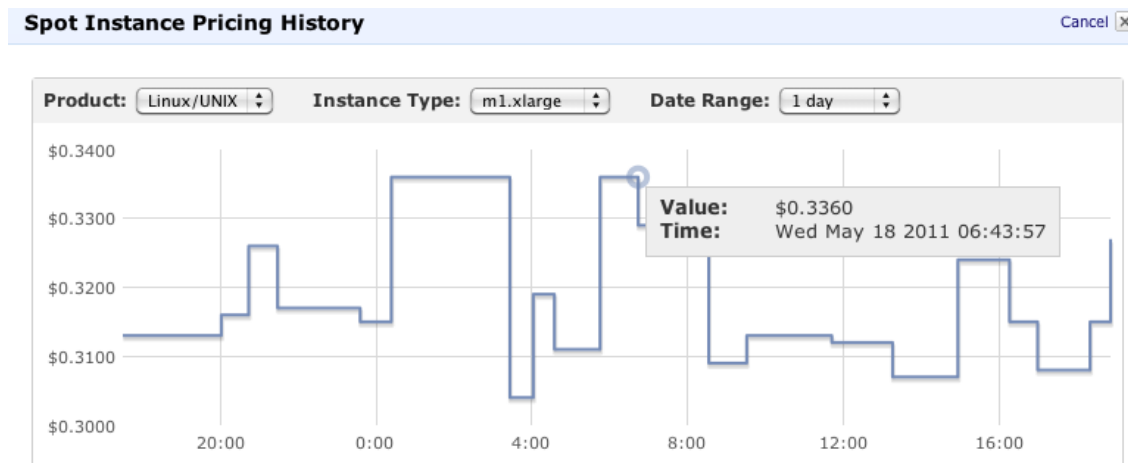
Cost / 1K Event (USD)	EC2		Tier3 Size Center	Large Center (hundreds K cores)
	m1.xlarge	cc1.4xlarge		
Simulation	37	26	11	5
Reconstruction	1.88+Storage	1.24+Storage	0.48	0.24

EC2 cost doesn't include storage (in/out/store), which is very significant as well!!!!

Yushu Yao, Performance of ATLAS Jobs in the EC2 Cloud, April 4, 2011.

EC2 Spot Instances

- Gives the ability to bid for unused Amazon EC2 capacity
- The price fluctuates based on supply and demand (range below: \$0.304 – \$0.336)




Thursday, May 19, 2011

Spot vs. On-Demand Instances

	On-Demand	Spot Instance
<i>Pricing</i>	Flat Hourly Rate	Variable Hourly Rate
<i>Start-up</i>	Immediate	May be delayed *
<i>Termination</i>	Only by user	May be terminated by EC2

* if the spot price is equal to your max bid price your instance may or may not be started based on the number of requests at that price

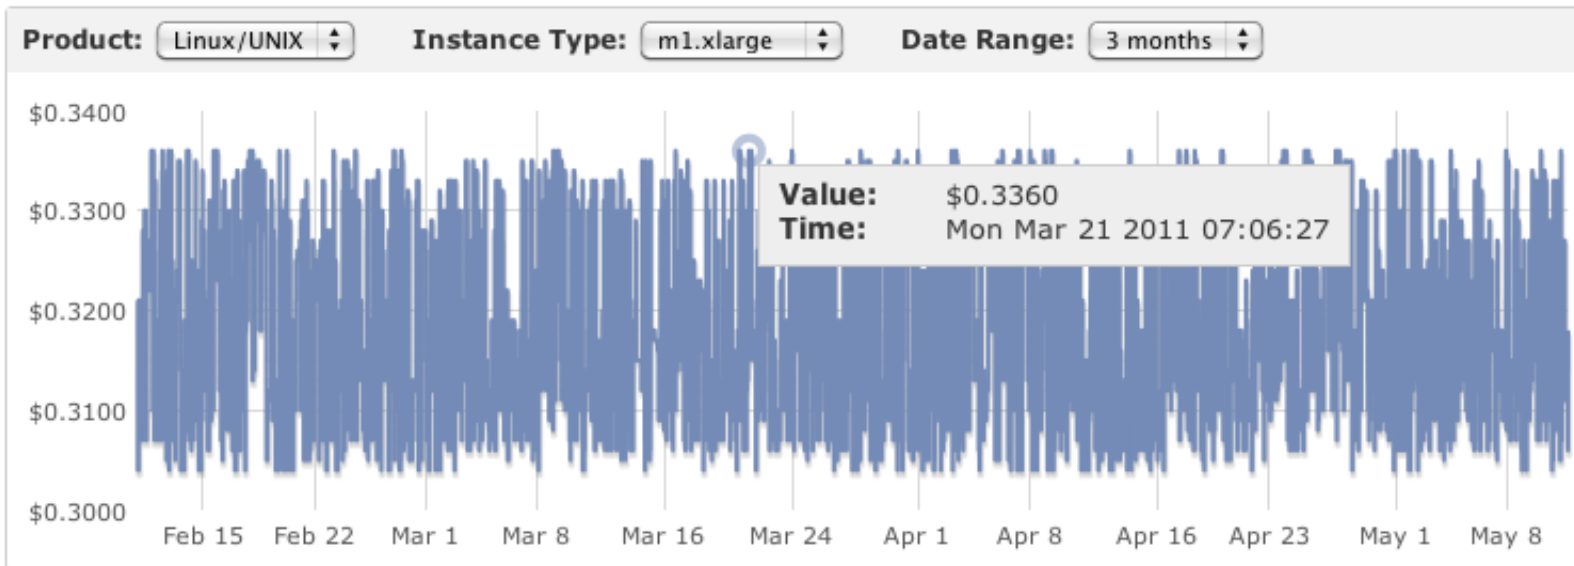
Thursday, May 19, 2011



Review the price history

```
$ ec2-describe-spot-price-history -t m1.xlarge
```

```
SPOTINSTANCEPRICE    0.330000 2011-03-12T12:11:24-0800  m1.xlarge  Linux/UNIX
SPOTINSTANCEPRICE    0.306000 2011-03-12T14:59:06-0800  m1.xlarge  Linux/UNIX
SPOTINSTANCEPRICE    0.307000 2011-03-12T16:02:26-0800  m1.xlarge  Linux/UNIX
...
```



* On-Demand \$0.76 per instance hour

Consider a max bid price

- You will often pay less per instance hour than your max bid price
 - everyone pays same spot price no matter what their max bid is
- Spot instances will be terminated when the spot price goes above your max bid price
 - if the spot price is equal to your max bid price your request may or may not get terminated
 - make sure workloads and applications can handle this
 - you are not charged for partial hours unless you terminate your own instance
- The durability of your request can either be one-time or persistent

Choose a request type

- Use a **one-time request** with a **high spot price** expecting that your instances get enough compute time to complete the job without interruption.
- Use a **persistent request** with a **lower spot price** combining a number of instances run over time whose aggregate time would be enough to complete the job.

Hourly Cost Per Core					
m1.xlarge On-Demand	m1.xlarge High Spot Price	cc1.4xlarge On-Demand	cc1.4xlarge High Spot Price	Tier 3 Center	Large Data Centers (Tier 1)
\$0.19	\$0.08325	\$0.20	N/A*	~\$0.08	~\$0.04

Initial Questions

- *What does it mean to terminate an instance in terms of the ATLAS jobs?*

We can either :

- Use ATLAS bookkeeping in order to resume event processing on restart.
- Use Amazon Simple Queue Service (SQS) for bookkeeping

And store results with either:

- Amazon EBS back AMIs which creates a new EBS for each instance.
(Amazon EBS volumes are created in a particular Availability Zone and can be from 1 GB to 1 TB in size.)
- Amazon Simple Storage Service (S3)
(unlimited data, data transfer between EC2 and S3 is free via COPY request within an Amazon Region.)

- *What do we do with partially processed input files? Luminosity block alignment?*
- *Perhaps the simplest solution is to overbid by 10% above highest spot price.*

Conclusion

- Running ATLAS jobs using EC2 On-Demand Instances is too expensive
- Amazon EC2 Spot Instances can reduce the costs of running ATLAS jobs possibly making it competitive with the cost of Tier 3 sized centers.
- Therefore, there is much more hope that EC2 and Tier3 costs are compatible which gives us a strong reason to pursue this.
- I have run some initial tests over the last week and found that there seems to be a time difference in the delayed start-up depending on the time of day. Also, my average spot price for m1.xlarge was \$0.318 per instance hour with a max bid price higher than the highest spot price of \$0.336.

Resources

- Yushu Yao, *Performance of ATLAS jobs on EC2 Cloud*
 - <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/AtlasPublicCloudMagellan>
- Amazon Documentation:
 - <http://aws.amazon.com/ec2/instance-types/>
 - <http://aws.amazon.com/ec2/spot-instances/>
- <http://cloudexchange.org/>: visualizations of the spot prices by region and instance type.

Spot Instance Features

- **Durability**
 - *Persistent*: remains in consideration after each instance termination and is resubmitted with initial parameters
 - *One-time*: a request is only satisfied once
- **Cluster** (may reduce the chance that your spot request is fulfilled)
 - *Launch Groups*: ensures that cluster instances start and terminate simultaneously
 - *Availability Zones*: ensure that cluster instances are launched in the same Availability Zone
- **Pricing history** is provided by Amazon to help you determine an appropriate spot price.