

High Throughput Computing with EC2 Spot Instances

HTCondor Week @ The University of Wisconsin, Madison

19 May 2016

Will St. Clair

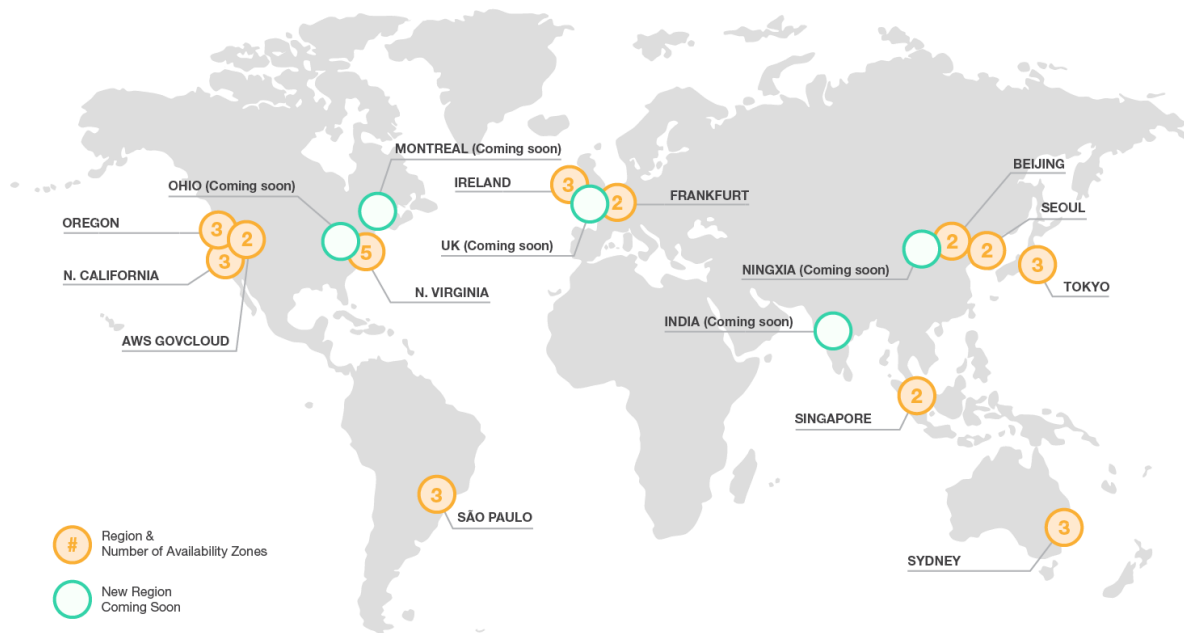
Solutions Architect

wstclair@amazon.com

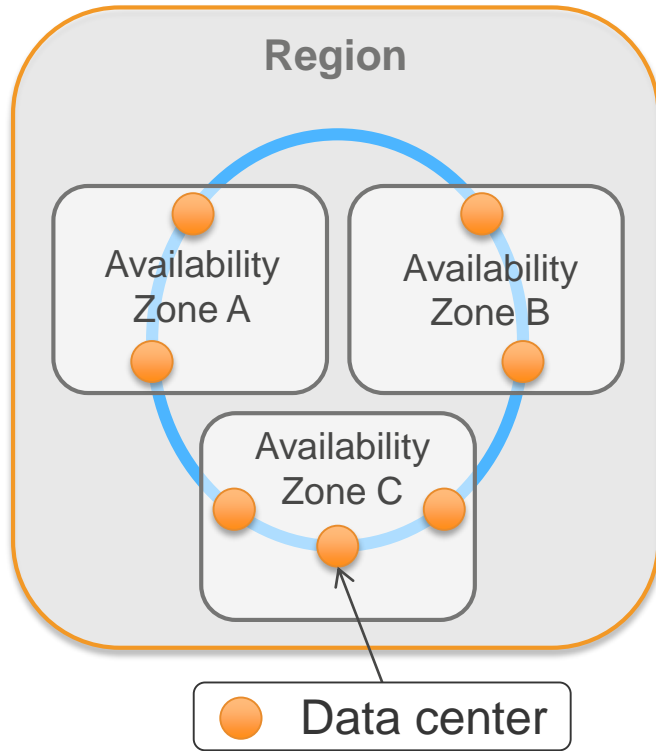
AWS Global Infrastructure

33 Availability Zones in
12 geographic Regions
around the world

11 more Availability Zones
in 5 new Regions coming
online in the next year



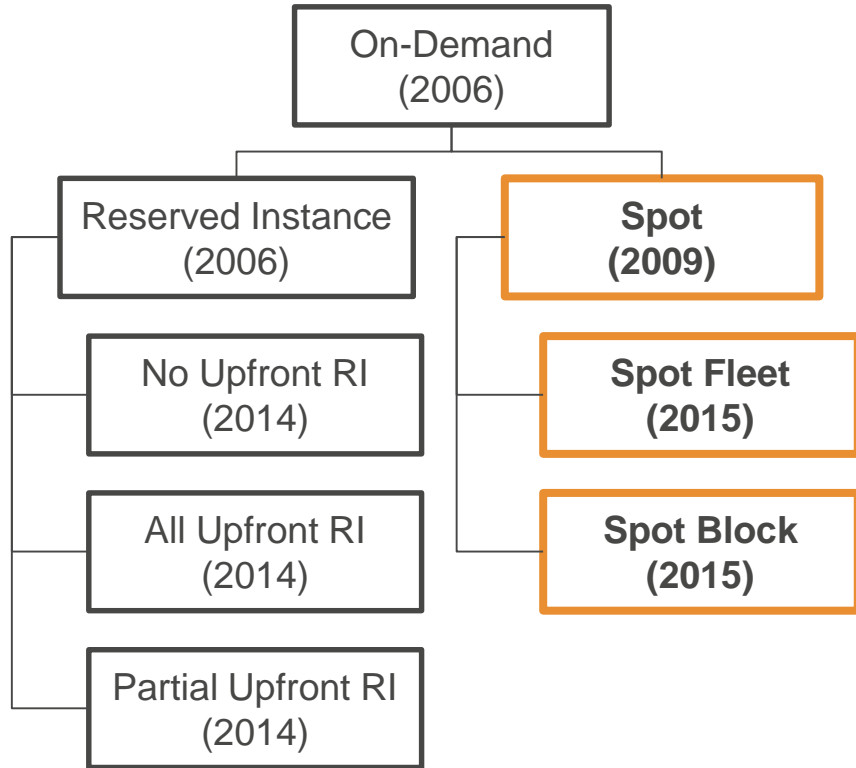
Regions and Availability Zones



A Region is a physical location in the world where we have multiple Availability Zones.



Availability Zones consist of one or more discrete data centers, each with redundant power, networking and connectivity, housed in separate facilities.

EC2 Pricing



- On-Demand
 - Pay a fixed price per hour with no commitments or upfront payments
- Reserved Instances
 - Save up to 70% compared to On-Demand with a 1- or 3-year commitment
 - Your capacity will always be available for the type and zone purchased
- Spot Instances
 - Specify the maximum hourly price you are willing to pay, with the risk of interruption
 - Savings as high as 80-90% compared to On-Demand


What is Amazon EC2 Spot?



Amazon EC2 Spot instances are spare EC2 instances that you can bid on to run your cloud computing applications. Spot instances are available at lower prices than On-Demand, so you can significantly reduce the cost of running your applications, grow your application's compute capacity and throughput for the same budget, and enable new types of cloud computing applications.

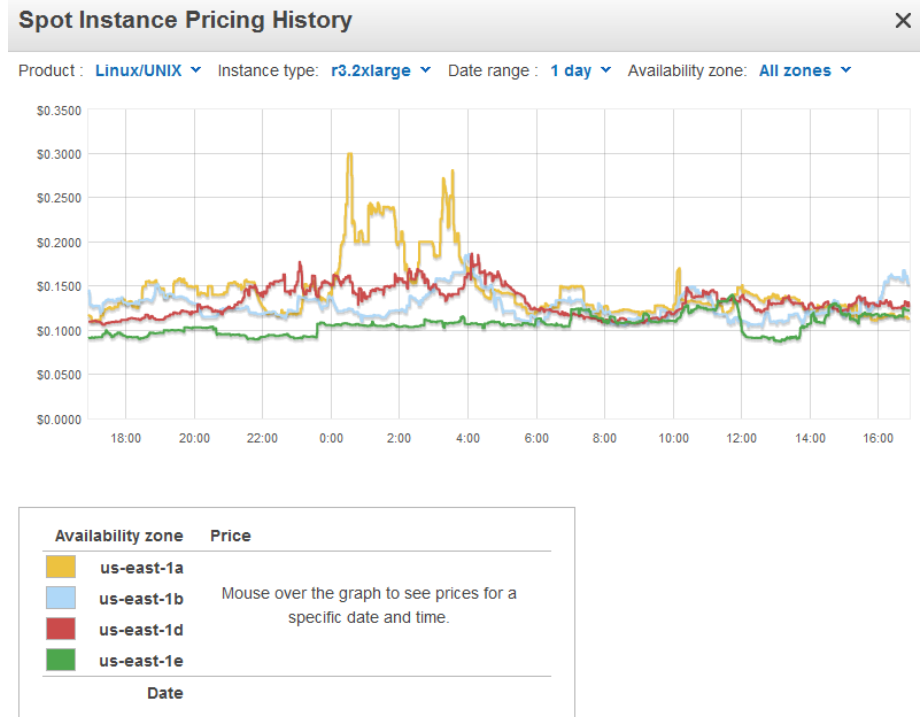
Spare capacity at scale

- AWS has more than a million active customers every month in 190 countries.
- On average, every week, AWS customers are using more compute capacity on Amazon EC2 Spot instances than customers in 2012 were running across all of Amazon EC2.



The Spot market

- Each instance type and Availability Zone combination is its own market
- As long as your bid price exceeds the market price, your instances run at the market (not bid) price
- 2-minute warning prior to shutdown if outbid



View pricing history for Spot Instances in the EC2 Console

Bid vs. Market Price



75% Bid

50% Bid

25% Bid

You pay the market price



Handling the 2-minute warning (example)

```
$ if curl -s http://169.254.169.254/latest/meta-  
data/spot/termination-time | \  
  grep -q .*T.*Z; then instance_id=$(curl -s  
http://169.254.169.254/latest/meta-data/instance-id); \  
/opt/bin/checkpoint.sh; fi
```

1. Check instance metadata service (169.254.169.254) for 2-minute warning
2. If present, run checkpoint.sh

Amazon EC2 Spot Bid Advisor

<https://aws.amazon.com/ec2/spot/bid-advisor/>

Spot Bid Advisor

Region: OS: Bid Price:


Instance Type: vCPU (min): Memory GiB (min): Instance types supported by EMR

Instance Type	vCPU	Memory GiB	Savings over On-Demand	Frequency of being outbid (month)	Frequency of being outbid (week)
m4.xlarge	4	16	91%	Low	Low
m4.2xlarge	8	32	91%	Low	Low
c3.xlarge	4	7.5	82%	Low	Low
c3.4xlarge	16	30	81%	Low	Low
c1.xlarge	8	7	89%	Low	Low
m3.xlarge	8	15	83%	Medium	Medium
m1.xlarge	4	15	90%	Medium	Low
c4.xlarge	4	7.5	83%	Medium	Medium
c4.2xlarge	8	15	87%	Medium	Medium
c4.4xlarge	16	30	84%	Medium	Medium

- Analyzes Spot price history to help you determine a bid price that suits your needs
- The lower your frequency of being outbid, the longer your Spot instances are likely to run without interruption

Optimizing your bidding strategy

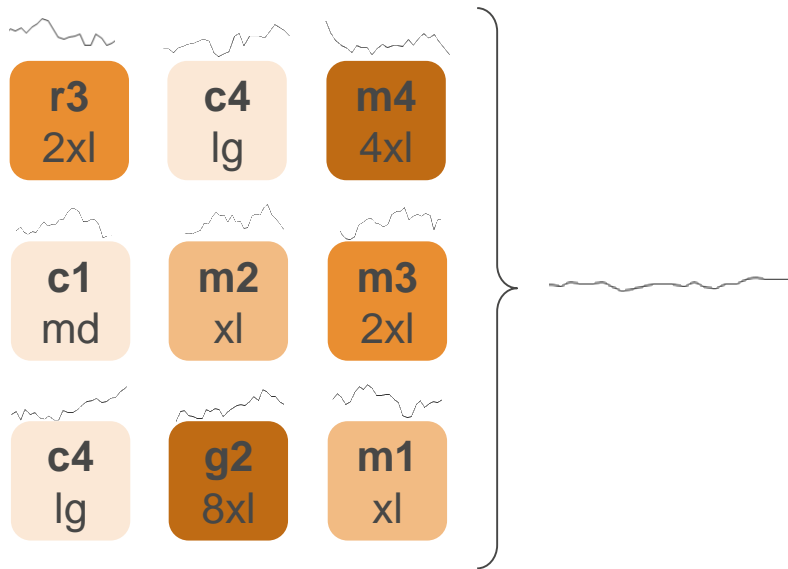
- Relying on a single instance type means more impact from price movements and instance terminations
- Calculating bid strategy for a given \$/core (or \$/GB) is undifferentiated heavy lifting
- Solution? Use **Spot Fleet**



Spot Fleet



- Make a single request for a target amount of resources drawing from multiple pools of resources
- Spot Fleet attempts to maintain its target capacity fleet if your Spot instances are interrupted



Bid based on custom units



- Optimize for the best price or the most instance type diversity to increase availability
- Bid based on your desired \$/vCPU, or \$/vCPU+RAM
- Eliminates the undifferentiated heavy lifting of bid planning

Find instance types

AMI ⓘ Amazon Linux AMI 2015.03.1 (HVM), .

Capacity unit ⓘ Customize your unit...

Minimum unit requirements

vCPUs:

4

Memory:

32 GIB

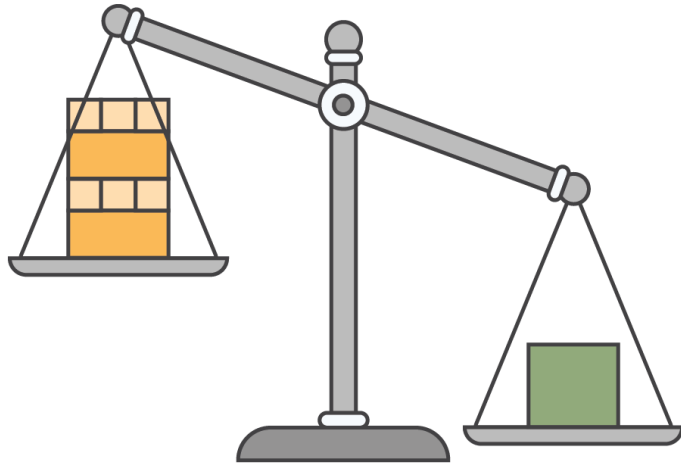
Current generation required

Target capacity ⓘ 100 units

Bid price ⓘ per unit/hour \$ 0.25

% of on-demand

Best practices for HTC



- To maximize goodput, jobs must be broken up into tasks small enough to largely avoid preemption
- Spot market conditions are generally good, but an aggressive bidding strategy can backfire if jobs are interrupted too often

Spot Blocks

- Request guaranteed execution from 1-6 hours
- Up to 50% off On-Demand pricing
 - Regular spot can be as high as 90%
- Use for:
 - Web/app servers
 - Master/coordinator nodes
 - Shared filesystem hosts
 - Stop-loss strategy for outlier jobs



```
$ aws ec2 request-spot-instances \  
  --block-duration-minutes 360 \  
  --instance-count 2 \  
  --spot-price "0.25" ...
```


Recap



Spot Instances

Save on EC2 pricing
by bidding on
unused capacity

+



Spot Fleet

Launch 10s, 100s, or 1000s of
instances across multiple instance
types to meet the target capacity you
define



Spot Block

Launch Spot Instances with
guaranteed execution time of up to 6
hours

A quick word on storage



Amazon Simple Storage Service (S3)

Highly scalable object storage
(data presented as buckets of immutable objects)



Amazon Glacier

Low-cost archival storage
(data presented as vaults of immutable archives)



Amazon Elastic Block Store (EBS)

Block storage for individual hosts
(data presented as disk volumes; analogous to SAN)



Amazon Elastic File System (EFS)

Network file storage
(data presented as a shared file system; analogous to NAS)

PREVIEW

Introducing Elastic File System



- Fully managed shared NFSv4 file system
 - Standard file system semantics
 - Works with standard OS file system APIs
- Highly scalable
 - Reach petabyte scale with 1000s of concurrent connections
 - Throughput and IOPS scale automatically as file systems grow
- Pay only for the storage space you use, with no minimum fee
- Currently in preview in us-west-2 (Oregon)

New magnetic Elastic Block Store (EBS) volumes

	Solid-State Drives (SSD)		Hard disk Drives (HDD)	
Volume Type	General Purpose SSD (gp2)*	Provisioned IOPS SSD (io1)	Throughput Optimized HDD (st1)	Cold HDD (sc1)
Max. IOPS/ Volume	10,000	20,000	500	250
Max. Throughput/ Volume	160 MiB/s	320 MiB/s	500 MiB/s	250 MiB/s
Max. IOPS/ Instance	48,000	48,000	48,000	48,000
Max. Throughput/ Instance	800 MiB/s	800 MiB/s	800 MiB/s	800 MiB/s
Dominant Performance Attribute	IOPS	IOPS	MiB/s	MiB/s

The AWS High Throughput toolbox



Spot Fleet

Automatically manage heterogeneous pools of spare EC2 capacity



Throughput-optimized (st1) EBS volumes

High throughput, low-cost storage for sequential I/O



Spot Block

Launch Spot Instances with guaranteed execution time of up to 6 hours



Elastic File System

Petabyte scale managed NFS service

PREVIEW

Thank you

<https://aws.amazon.com/ec2/spot/>