

# Save up to 90% on CI/CD Workloads with Amazon EC2 Spot Instances

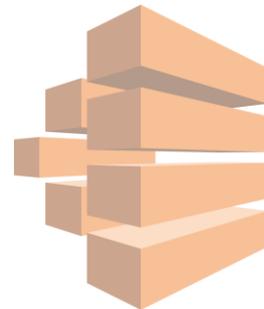
Chad Schmutzer, Solutions Architect – EC2 Spot  
schmutze@amazon.com

November 2018



# Agenda

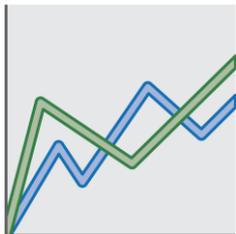
- Amazon EC2 Spot Instances
- Use cases
- Let's take it for a spin!
- Q&A



# Amazon EC2 Purchasing Options

## On-Demand

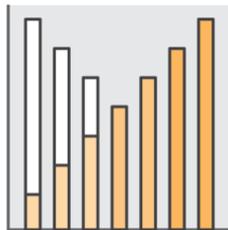
Pay for compute capacity by **the second** with no long-term commitments



Spiky workloads, to define needs

## Reserved Instances

Make a 1- or 3-year commitment and receive a **significant discount** off On-Demand prices



Committed, steady-state usage

## Spot Instances

Spare EC2 capacity at **savings of up to 90%** off On-Demand prices



Fault-tolerant, flexible, stateless workloads

**WHAT IT FEELS LIKE TO**



**REDUCE OPERATING COSTS BY 90%**

# What are EC2 Spot Instances?

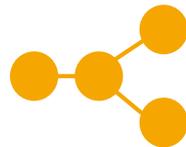
Spare EC2 capacity with some great advantages



**Low Cost**



**Faster Results**



**Easy Access**



**Resource  
Flexibility**



**The price changes  
infrequently based on  
supply and demand of  
spare capacity**



**Just request capacity and  
pay the current rate - when  
we need the capacity back  
you'll get a 2 minute  
warning**

# Spot is EASY to use

Simplified access model, predictable prices, pause & resume



## Simplified access

Via existing EC2 API calls such as RunInstances



## Simplified Pricing Set it & Forget it

Low predictable prices without bidding



## Pause & Resume

With hibernation and stop-start

**EC2 Spot Fleet:** Simply tell EC2 Spot Fleet how much capacity you need and Fleet does the rest

# Spot Instances: Spare compute capacity at scale



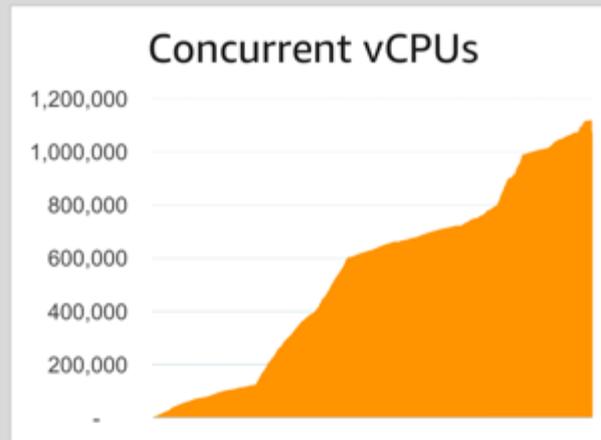
Supports all EC2 Instance Types and Available in all Public Regions and Availability Zones  
*GovCloud Coming Soon!*

# Spot Instances: Spare compute capacity at scale



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**.

CLEMSON UNIVERSITY RAN THE LARGEST HIGH PERFORMANCE CLUSTER ON SPOT WITH **1,100,000 CONCURRENT vCPUs**



<http://amzn.to/2wmYkWm>

# Spot Instances: Spare compute capacity at scale

Creating a 1.3 Million vCPU Grid on AWS using EC2 Spot Instances and TIBCO GridServer

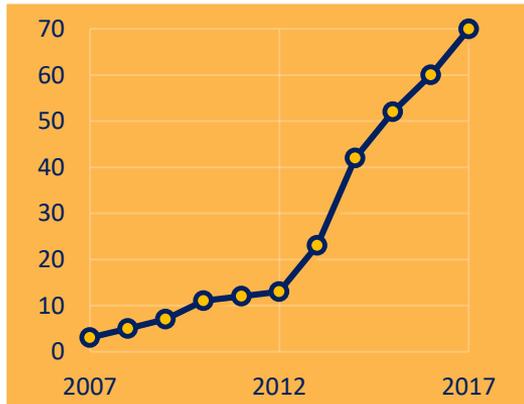
<https://amzn.to/2zvM8Zn>

Univa Demonstrates Extreme Scale Automation by Deploying More Than One Million Cores in a Single Univa Grid Engine Cluster using AWS

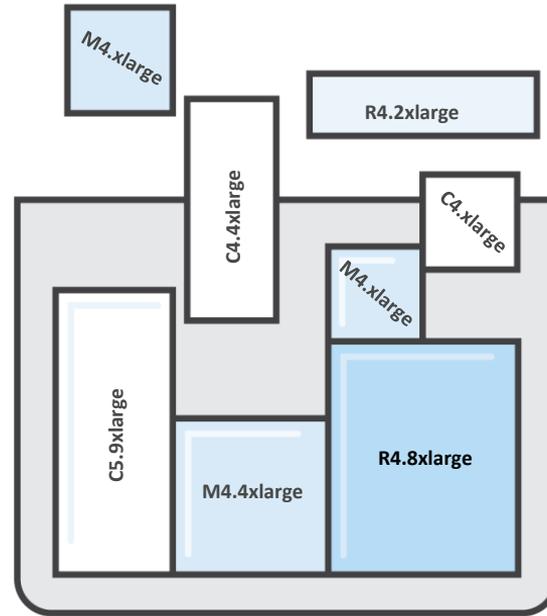
<http://bit.ly/2NHbK8B>

# Understanding how to use EC2 Spot pools

Each instance family, each instance size, in each Availability Zone, in every Region is a separate Spot pool



100s of Instance Options



Instance Flexibility and Diversification  
are crucial

# What does it mean to be instance flexible?

## Instance Size

- m4.xlarge capacity is different from m4.large. Use as many sizes as you can efficiently!

## Instance Family

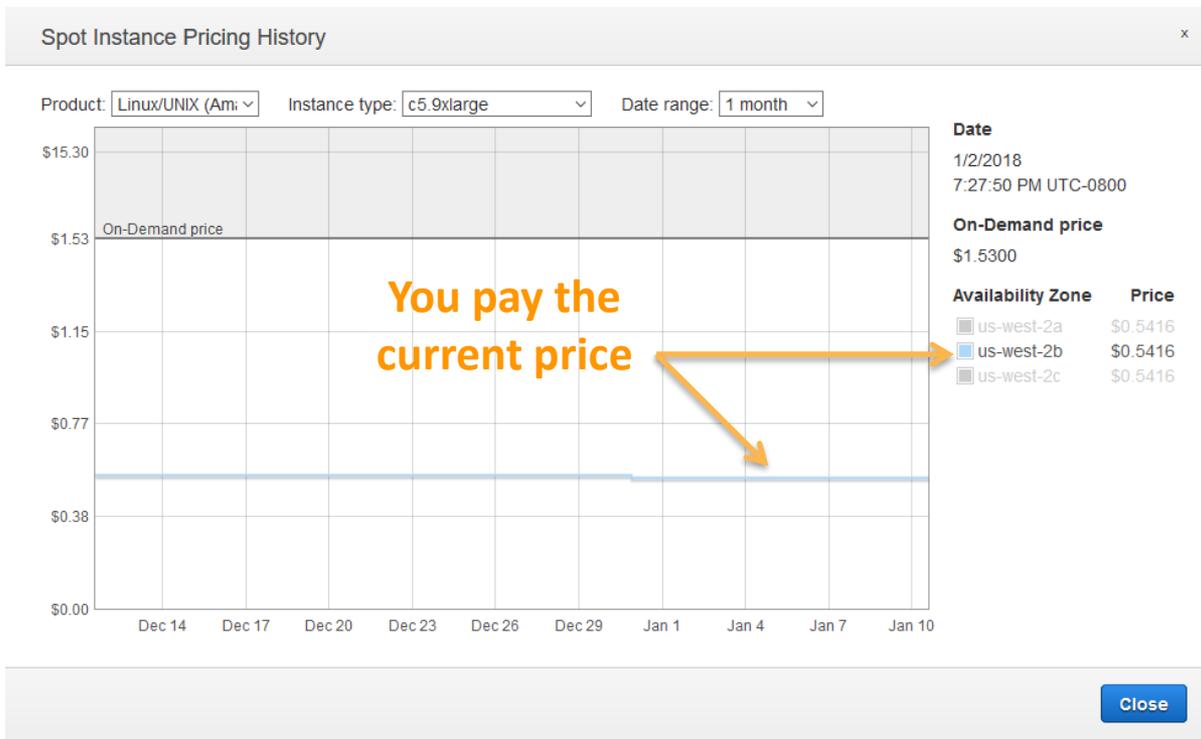
- If you're using the c4.large, you can almost certainly use m4.large and r4.large. They have the same number of vCPUs, just with some extra memory!

## Availability Zone

- Availability Zones consist of one or more discrete data centers. us-east-2a and us-east-2b capacity is different! If your application can use multiple AZs, do it!

# EC2 Spot Instance Pricing

- You will never be charged more than the maximum price you specified
- While your instance runs, you are charged the Spot price that is in effect for that period



# Worried about interruptions? Don't be.

*Over the last 3 months, >95% of Spot Instance interruptions were from a customer terminating the instance because the application had completed its work.*

# EC2 Spot Instance Advisor

Spot Instance Advisor

Region: **US East (N. Virginia)** OS: **Linux/UNIX**

Instance type filter:  
vCPU (min): **1** Memory GiB (min): **0**  Instance types supported by EMR

Instance Type	vCPU	Memory GiB	Savings over On-Demand*	Frequency of interruption ▾
m3.large	2	7.5	78%	<5%
t2.large	2	8	70%	<5%
d2.xlarge	16	122	70%	<5%
d2.8xlarge	36	244	70%	<5%
r4.xlarge	4	30.5	76%	<5%
i2.xlarge	4	30.5	70%	<5%
m3.medium	1	3.75	90%	<5%
r3.xlarge	4	30.5	82%	<5%
m4.2xlarge	8	32	69%	<5%
g3.16xlarge	64	488	70%	<5%

Display all 90 instance types

✓ Percentage of interruptions published to the Spot Instance Advisor

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

# EC2 Spot Fleet

Automates the management of Spot instances – Simply tell Spot Fleet how much capacity you need and Fleet does the rest.

## Fast Access to Capacity



- Launch Thousands of Spot Instances with single API call
- Maintain capacity or scale up or down based on metrics and thresholds you set

## Optimize Price vs. Availability



- Find the lowest priced horsepower that works for you or diversify your fleet to increase availability
- Auto-attach to a Load Balancer

## Customize Based on Application



- Create your own capacity unit based on your application requirements

# Spot Instances Use Cases

## Big data



**FINRA** has **saved up to 50%** from its on-premises solution, increased elasticity/scalability, and accelerated reprocessing requests from months to days with EC2 Spot Instances

## Containers & CI/CD



**Yelp** runs **millions of tests** every day with EC2 Spot Instances. Yelp improved test result response time from **2 days to 30 minutes** and has also delivered a large reduction in execution costs with Spot.

## HPC & batch



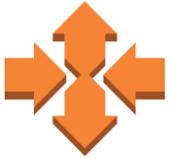
**TLG Aerospace** saw a **75% reduction in the cost** per CFD simulation with Amazon EC2 Spot Instances. They were able to pass those savings along to their customers and be more competitive.

## Stateless web services



**AdRoll** have been able to seamlessly scale their infrastructure, better serve customers across the globe, and **reduce our fixed costs by 75% and operational costs by 83%**.with AWS solution, including EC2 Spot Instances

# Amazon EC2 Spot works with...



Auto Scaling



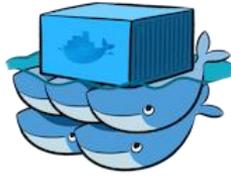
AWS Batch



Amazon EMR



Amazon ECS



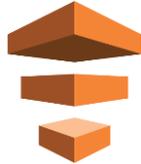
Apache MESOS™



AWS OpsWorks



AWS CloudFormation



AWS Data Pipeline



Bamboo



bole

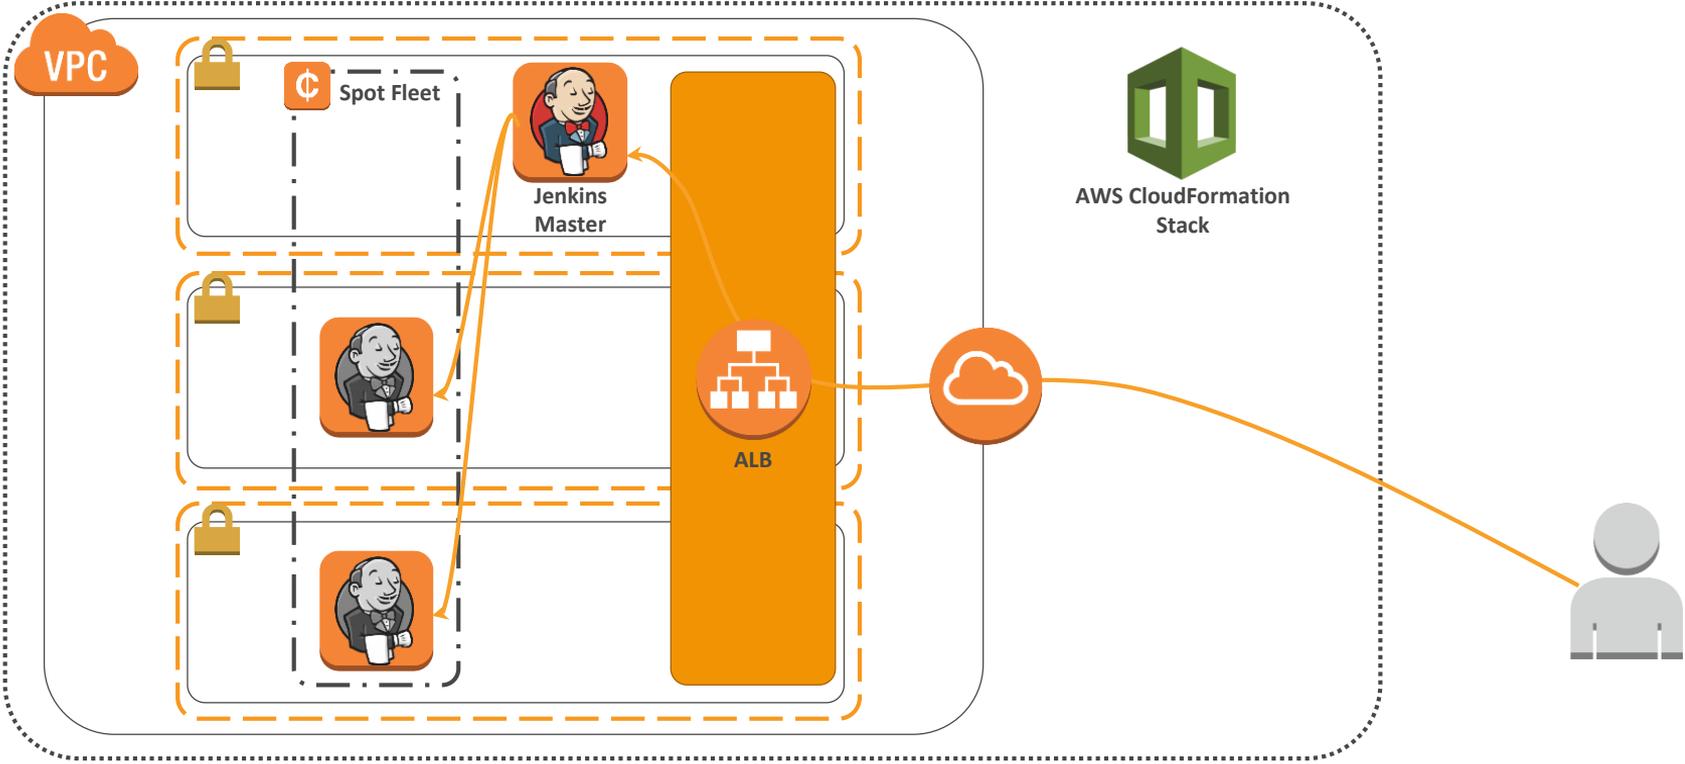
cloudera



Would you like to see a demo?



# Demo architecture



Any questions?



# Thank You!

- Get started
  - [Amazon EC2 Spot Instances](#)
  - <https://plugins.jenkins.io/ec2-fleet>
  - <https://marketplace.atlassian.com/apps/1218167/awsspotfleetbambooplugin>
- Demos
  - <https://github.com/awslabs/ec2-spot-labs/tree/master/ec2-spot-fleet-jenkins-plugin>