Tech Talk: (2019_1203-CON)

Fargate Spot

Deepthi Chelupati, Senior Product Manager, EC2
Akshay Ram, Senior Product Manager, Fargate

12/11/2019

aws

Agenda

- What is Fargate Spot?
- Benefits of Fargate Spot
- What applications should I use with Fargate Spot?
- How is it different from Amazon EC2 Spot?
- Demo 1 – using 100% Fargate Spot
- Demo 2 - Mixing Fargate purchase options
- Fargate purchase options
- Best practices
- Resources

aws

What is Fargate Spot?

New pricing option for Fargate – up to 70% off of Fargate prices

Caveat – Fargate tasks can be interrupted with a 2 min notification.

Benefits of Fargate Spot

Up to 70%

Steep discounts for applications that are fault tolerant and time insensitive

Migration friendly

We use the same SIGTERM event which is what customers listen today to perform cleanup

Application first controls

Every service can define a mix of Fargate and Fargate Spot

aws

What applications should I use with Fargate Spot?

Fargate Spot is ideal for:

- ✓Fault-tolerant
- ✓Loosely coupled
- ✓Stateless workloads

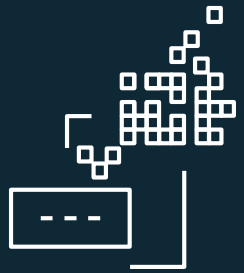

Big data	CI/CD	Web services	HPC

aws

How is it different from Amazon EC2 Spot?

Amazon EC2 Spot

Spare EC2 Capacity

Save up to 90% over On-Demand

Can be *reclaimed* by EC2 (with two minute warning)

Follow EC2 Spot best practices – instance flexibility and use capacity optimized allocation strategy for ASGs.

AWS Fargate Spot

Spare compute Capacity

Save up to 70% over standard Fargate

Can be *reclaimed*
(with two minute warning)

Automatic diversification

Demo 1 – Using 100% Fargate Spot

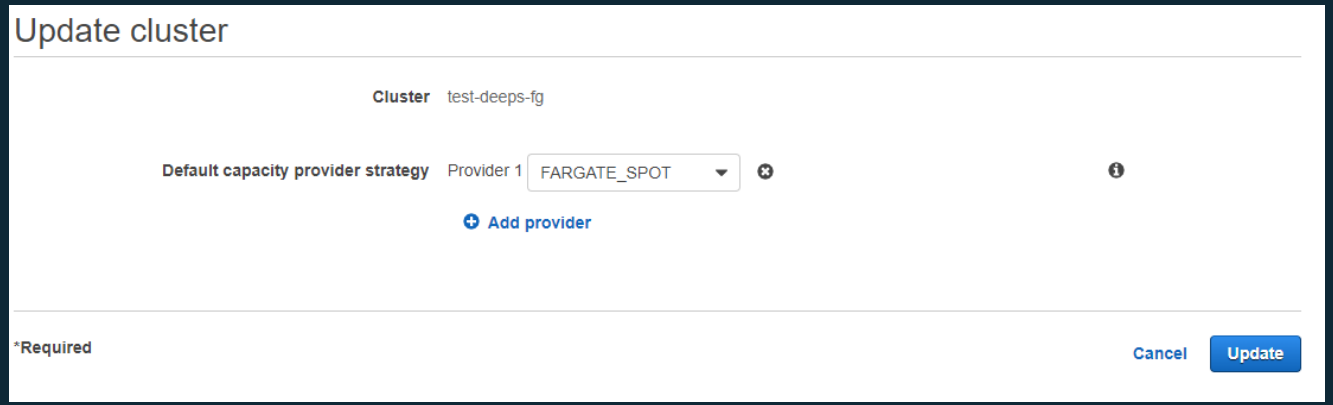

Good fit for 100% Spot:

- ETL
- Batch processing
- Dev/test environments

Demo 2 – Using a mix of Fargate purchase options

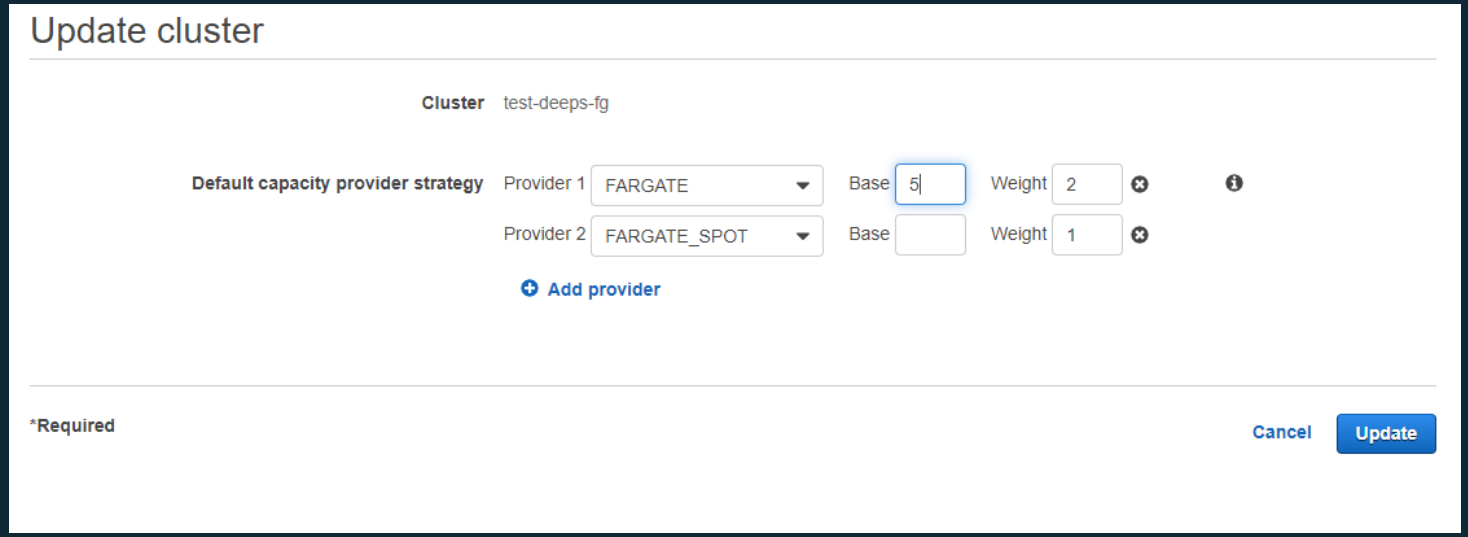

Good fit - Long running services that need baseline compute with intermittent burst compute needs

Fargate Spot Interruption events

```
// ECS Event after an interruption
{
  ...
  'detail-type': 'ECS Task State Change',
  source: 'aws.ecs',
  account: '123456789012',
  ...
  detail: {
    ...
    desiredStatus: 'STOPPED',
    lastStatus: 'RUNNING',
    stoppedReason: 'Your Spot Task was interrupted.',
    stopCode: 'TerminationNotice',
    taskArn: 'arn:aws:ecs:us-west-2:123456789012:task/8c301e27-f5fd-4e1d-b837-221db7285ce4',
    ...
  }
}
```

Fargate Spot Interruption events

```
aws ecs describe-tasks --tasks 8c301e27-f5fd-4e1d-b837-221db7285ce4

{
  "tasks": [
    {
      "taskArn": "arn:aws:ecs:us-west-2:123456789012:task/8c301e27-f5fd-4e1d-b837-
221db7285ce4",
      "clusterArn": "arn:aws:ecs:us-west-2:334552443286:cluster/default",
      "desiredStatus": "STOPPED",
      "lastStatus": "RUNNING",
      ...
      "stopCode": "TerminationNotice"
      "stoppedReason": "Your Spot Task was interrupted.",
      ...
    }
  ]
}
```

Interruption cloud watch logs

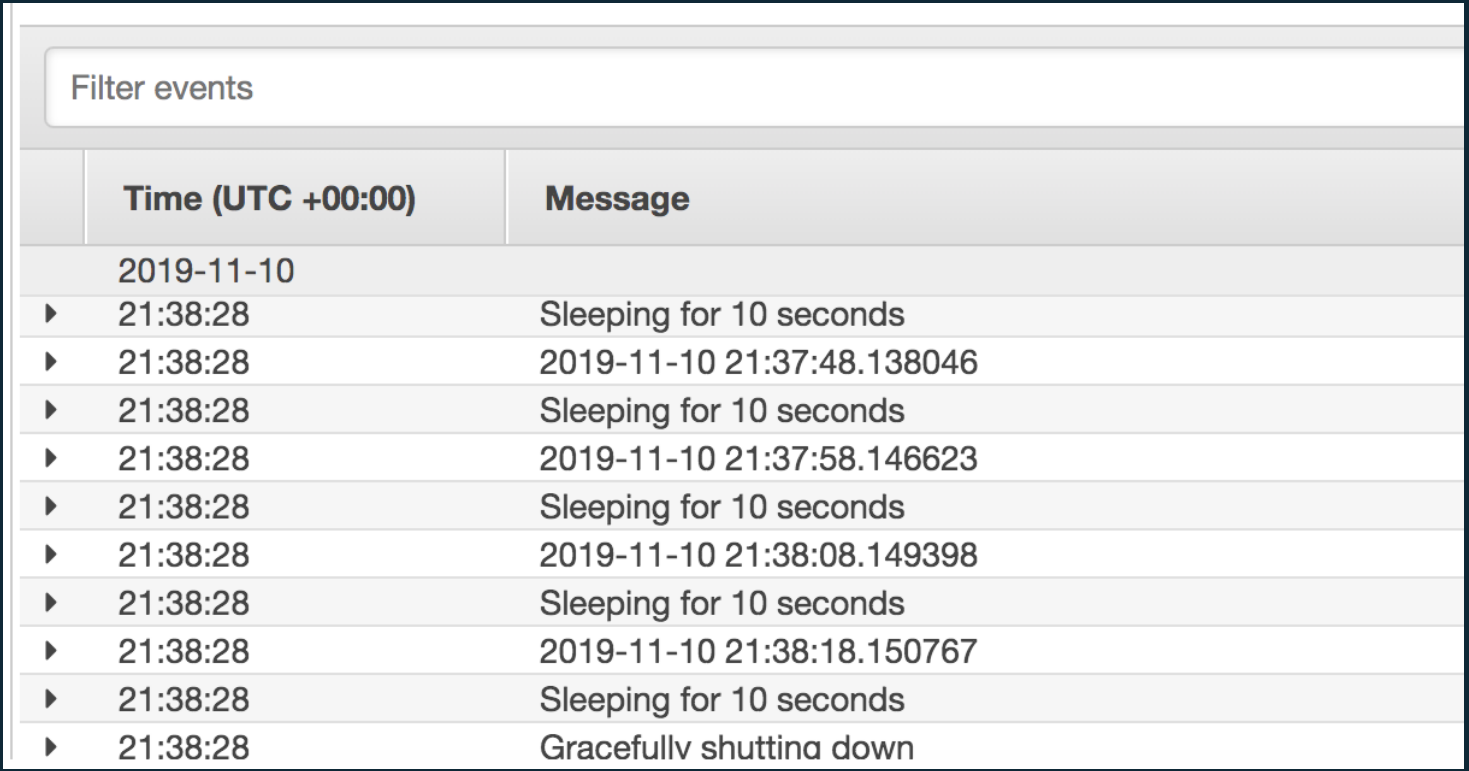

Fargate Purchase options

Fargate

Pay for containers per-second with no long-term commitment

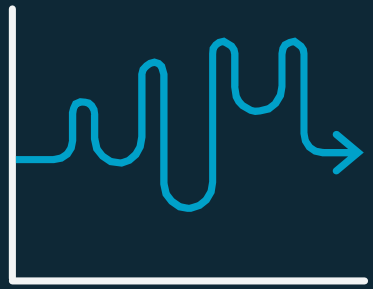

Capacity needs can change rapidly

Compute Savings Plan

Make a 1 or 3-year commitment and receive a significant discount

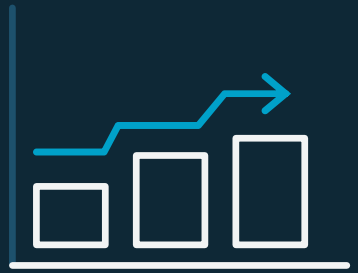

Baseline compute needs known in advance

Fargate Spot

Spare capacity with savings up to 70% off Fargate standard pricing

Fault-tolerant, flexible workloads

aws

Best practices

- Use an ECS Service instead or your own scheduler
- Don’t rely 100% on Fargate Spot for critical workloads with SLA
- Configure a base number of regular Fargate Tasks
- Applications should be fault-tolerant. Handle interruptions gracefully by catching SIGTERM
- Great for: stateless, fault-tolerant workloads

Resources

https://aws.amazon.com/blogs/aws/aws-fargate-spot-now-generally-available/

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS_Fargate.html#fargate-spot

https://www.youtube.com/watch?v=OpNBldEww0Q&t=207s

https://www.twitch.tv/videos/516666118?t=2h59m40s

aws