



# aws INNOVATE

MODERN APPLICATIONS EDITION

20 October, 2022

# Get started with Karpenter – Just In Time Nodes for your Amazon EKS Clusters

Joshua Samuel

Enterprise Solutions Architect

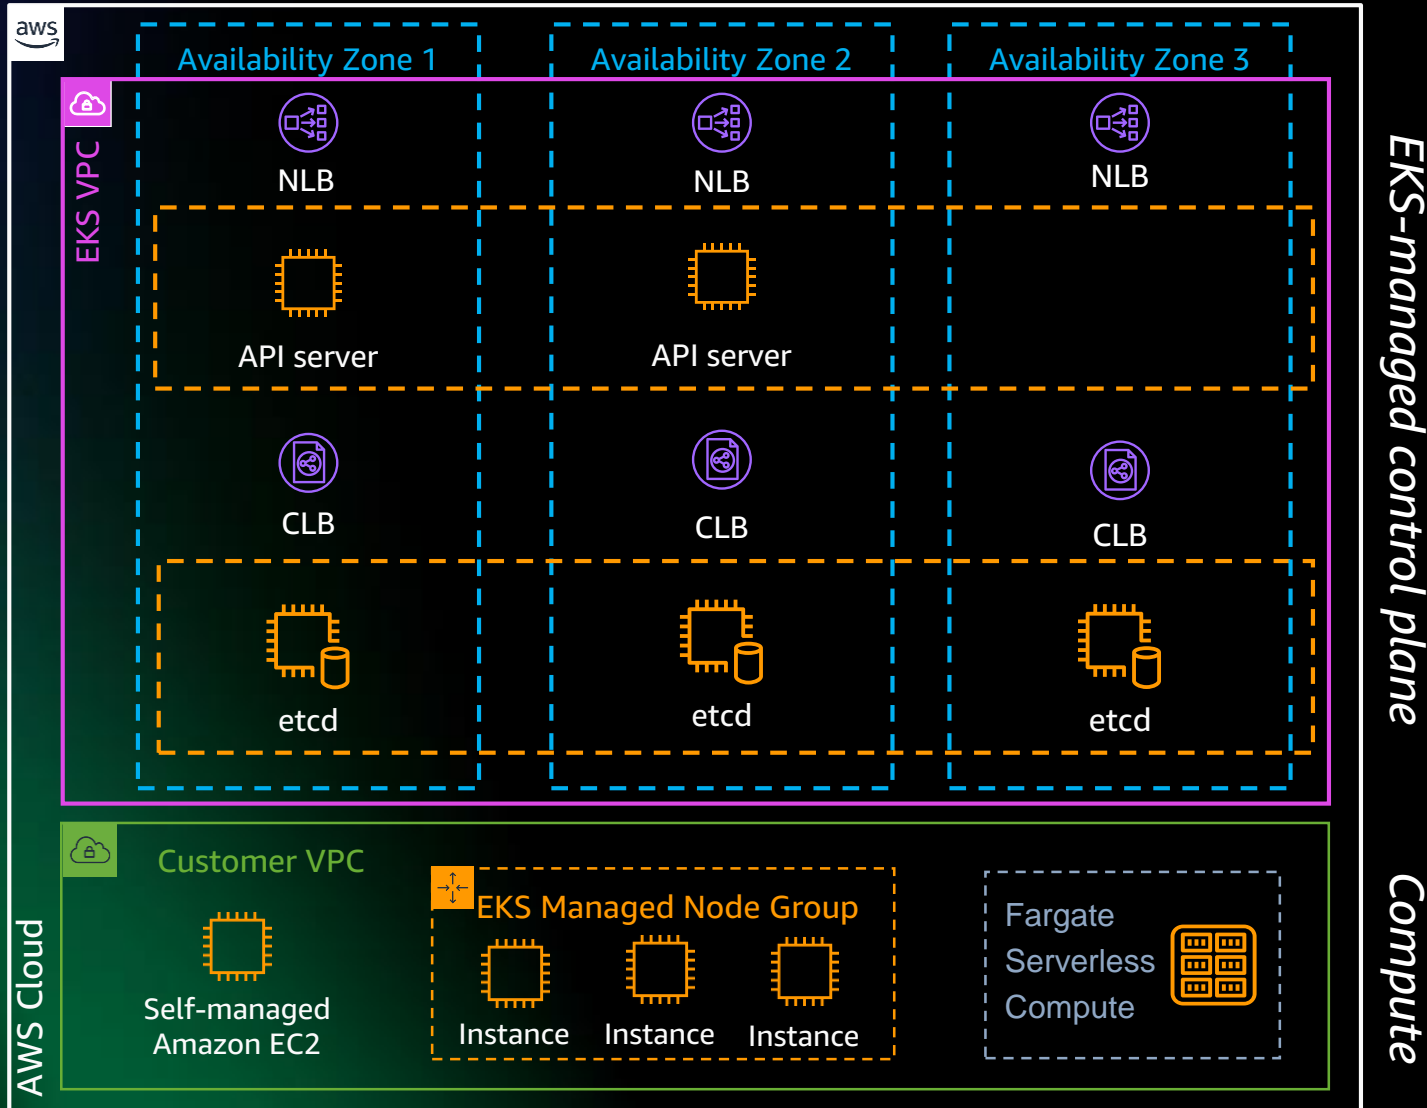
Amazon Web Services



# Agenda

- What is Amazon EKS
- Kubernetes Cluster Autoscaler (CAS)
- Cluster Auto Scaling with Karpenter
- Demo
- Resources to get started

# What is Amazon EKS



## Amazon EKS compute options

### Self-managed Amazon EC2 instances

- Run in your account
- Customer managed
- Highly-flexible

### Amazon EKS managed node group



- Run in your account
- Amazon EKS handles provisioning and instance lifecycle

### AWS Fargate

- Serverless
- AWS managed OS
- pod-based billing


# Kubernetes Cluster Autoscaler (CAS)

Cluster Autoscaler  
(CAS)



Cluster Autoscaler uses EC2's Auto Scaling groups to expand or contract the number of workers in a cluster

ASG instance type RAM/CPU configurations are recommended to be homogenous. which can lead to suboptimal workload placement (or multiple ASGs)




**Compute overhead =  
more costs**

# Cluster Auto Scaling with Karpenter

Karpenter works in tandem with the Kubernetes scheduler and a cluster's compute provider (e.g., Amazon EC2)

- Launches right-sized nodes in response to unschedulable pods
- Terminates underutilized nodes
- Works with Kubernetes workload scheduling constraints



**Karpenter will launch a worker that is best suited to run your workloads**

# Cluster Auto Scaling with Karpenter

Karpenter brings **flexibility** and performance

- Compute is best fit to cluster workloads automatically
- Declarative K8s custom resource definition (CRD) configuration:

```
apiVersion: karpenter.sh/v1alpha5
kind: Provisioner
spec:
  ttlSecondsUntilExpired: 2592000
  ttlSecondsAfterEmpty: 30
  provider:
    cluster:
      name: my-cluster
```



# Karpenter provisioner example

**apiVersion:** karpenter.sh/v1alpha5

**kind:** Provisioner

**metadata:**

**name:** default

**spec:**

**provider:**

**ttlSecondsAfterEmpty:** 30

**ttlSecondsUntilExpired:** 2592000

**requirements:**

- key: karapenter.k8s.aws/instance-family  
  operator: In  
  values: [c5, m5, r5]
- key: karpenter.k8s.aws/instance-size  
  operator: NotIn  
  values: [nano, micro, small, large]
- key: node.kubernetes.io/instance-type  
  operator: NotIn  
  values: [m5.24xlarge]
- key: "topology.kubernetes.io/zone"  
  operator: In  
  values: ["ap-southeast-2a", "ap-southeast-2b"]

How long to wait after node is empty before deletion

Expire nodes older than this age

Instance types to choose from

Instance sizes not to use

Specific instance type(s) not to use

Availability zones to provision within

More configurations available at:

<https://karpenter.sh/v0.16.1/provisioner>



# Demo

# Karpenter roadmap and community

Karpenter v0.16 is **available now**

- Workload consolidation
- Built-in UserData and custom AMIs
- EBS volume and kubelet configuration
- AZ-aware scheduling for stateful workloads

## 2022 Roadmap

- Node termination handling, e.g. Spot
- Instance type settings overrides
- Improved node AMI upgrades



Learn more at:

[github.com/aws/karpenter](https://github.com/aws/karpenter)

# Other resources

- Karpenter:
  - <https://karpenter.sh/>
- Workshop:
  - [https://www.eksworkshop.com/beginner/085\\_scaling\\_karpenter/](https://www.eksworkshop.com/beginner/085_scaling_karpenter/)
  - <https://catalog.us-east-1.prod.workshops.aws/workshops/76a5dd80-3249-4101-8726-9be3eeee09b2/en-US/autoscaling/karpenter>
- Blogs:
  - [Introducing Karpenter – An Open-Source High-Performance Kubernetes Cluster Autoscaler](#)



# Visit the Modern Applications resource hub

Dive deeper with these resources to help you develop an effective plan for your modernization journey.

- Build modern applications on AWS
- Business value of cloud modernization
- An introduction to event-driven architectures
- Accelerate full-stack web and mobile app development
- Determining the total cost of ownership: Comparing serverless and server-based technologies
- Building event-driven architectures with AWS
- Continuous learning, continuous modernization



<https://tinyurl.com/modern-apps-aws>

Visit resource hub




# AWS Training and Certification

Get started with Free Digital Training for you and your team today



Achieve key milestones and plan your next steps with the AWS Modern Application skills training



Access 500+ free digital courses with [AWS Skill Builder](#)



Earn an industry-recognized credential:  
[AWS Certified Developer – Associate](#)  
[AWS Certified DevOps – Professional](#)



Create a self-paced learning roadmap  
[AWS ramp-up guide - Developer](#)  
[AWS ramp-up guide - DevOps](#)

# Thank you for attending AWS Innovate Modern Applications Edition

We hope you found it interesting! A kind reminder to **complete the survey**.  
Let us know what you thought of today's event and how we can improve the event  
experience for you in the future.



[aws-apj-marketing@amazon.com](mailto:aws-apj-marketing@amazon.com)



[twitter.com/AWSCloud](https://twitter.com/AWSCloud)



[facebook.com/AmazonWebServices](https://facebook.com/AmazonWebServices)



[youtube.com/user/AmazonWebServices](https://youtube.com/user/AmazonWebServices)



[slideshare.net/AmazonWebServices](https://slideshare.net/AmazonWebServices)



[twitch.tv/aws](https://twitch.tv/aws)

# Thank you!