
AWS
re:Invent
DECEMBER 2 – 6, 2024 | LAS VEGAS, NV

KUB204

Automate your entire Kubernetes cluster with Amazon EKS Auto Mode

Alex Kestner
(he/him)
Product Manager
Amazon EKS

Todd Neal
(he/him)
Software Development Engineer
Amazon EKS

Alex Kestner

(he/him)
Product Manager
Amazon EKS

Todd Neal

(he/him)
Software Development Engineer
Amazon EKS

Agenda

01 Amazon EKS until now

02 Introducing Amazon EKS Auto Mode

03 Amazon EKS Auto Mode Demo

04 Amazon EKS Auto Mode Deep Dive

The Cloud, Kubernetes, and Amazon EKS

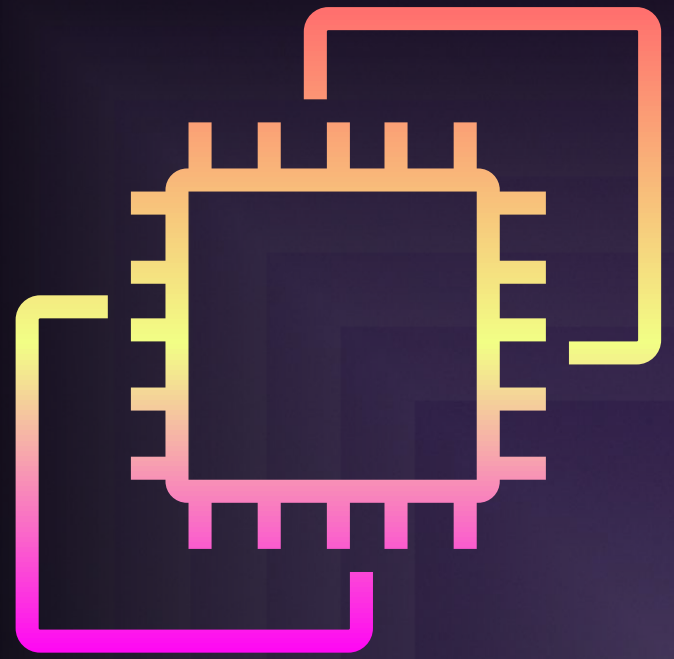

Enter Kubernetes

84%

using in production
or evaluating

CNCF Annual Survey 2023, cncf.io/reports/cncf-annual-survey-2023

Why Kubernetes?

Simplicity

195 CNCF projects

Consistency

100s of compatible tools

Extensibility

Unlimited customization

Amazon EKS

- The most trusted and secure way to run Kubernetes
- Enables you to build reliable, stable, and secure applications in any environment
- Fully upstream and certified conformant Kubernetes

7 years of managed Kubernetes on AWS

2018

EKS Generally Available

Managed Cluster Version Updates

GPU Support

HIPAA eligible

2019

ISO, PCI, and SOC Compliance

Expansion to 15 AWS regions

CSI drivers for EBS, EFS, FSx for Lustre

Pod security policies

Managed Node Groups

EKS Fargate

EKS on AWS Outposts

2020

Price reduced to $0.10 per hour

Secrets Encryption

SLA raised to 99.95%

EKS CIS Benchmark

ACK Project

Load balancer controller

EKS on AWS Local zones

K8s Resources in EKS console

EKS Add-ons

2021

OIDC access authentication

Karpenter project

Cluster creation reduced by 40%

Control plane scaling

EKS Anywhere GA

EKS Connector

FedRamp High Compliance

EFA, P4d instance support

2022

IPv6 clusters available

PrivateLink support

Local clusters on Outposts

Trainium instances
Add-ons from AWS marketplace

Nitro enclaves

Launch time reductions

2023

Expansion to 32 AWS regions

Kubernetes Network Policy enforcement

Upgrade Insights

Extended version support

Version launch acceleration

Pod Identity

S3 Mountpoints CSI Driver

2024

Amazon Linux 2023

Automatic Version Upgrades

Zonal Shift

Full IPv6 support

Multi-cluster dashboard

Auto Mode

Hybrid Nodes

Node health & auto-repair

Launch Counter // 0

AMAZON EKS

Runs tens of millions

of clusters every year

Deploying and operating Kubernetes applications

Deploying and operating Kubernetes applications

Deploying and operating Kubernetes applications

Provision
control plane

Install
plug-ins

Deploying and operating Kubernetes applications

Provision control plane

Install plug-ins

Select and provision compute

Deploying and operating Kubernetes applications

Provision control plane

Install plug-ins

Select and provision compute

Deploy application

Deploying and operating Kubernetes applications

Provision control plane

Install plug-ins

Select and provision compute

Deploy application

Monitor, repair, and upgrade

Deploying and operating Kubernetes applications

Provision control plane

Install plug-ins

Select and provision compute

Deploy application

Monitor, repair, and upgrade

Analyze and optimize

Amazon EKS cluster architecture

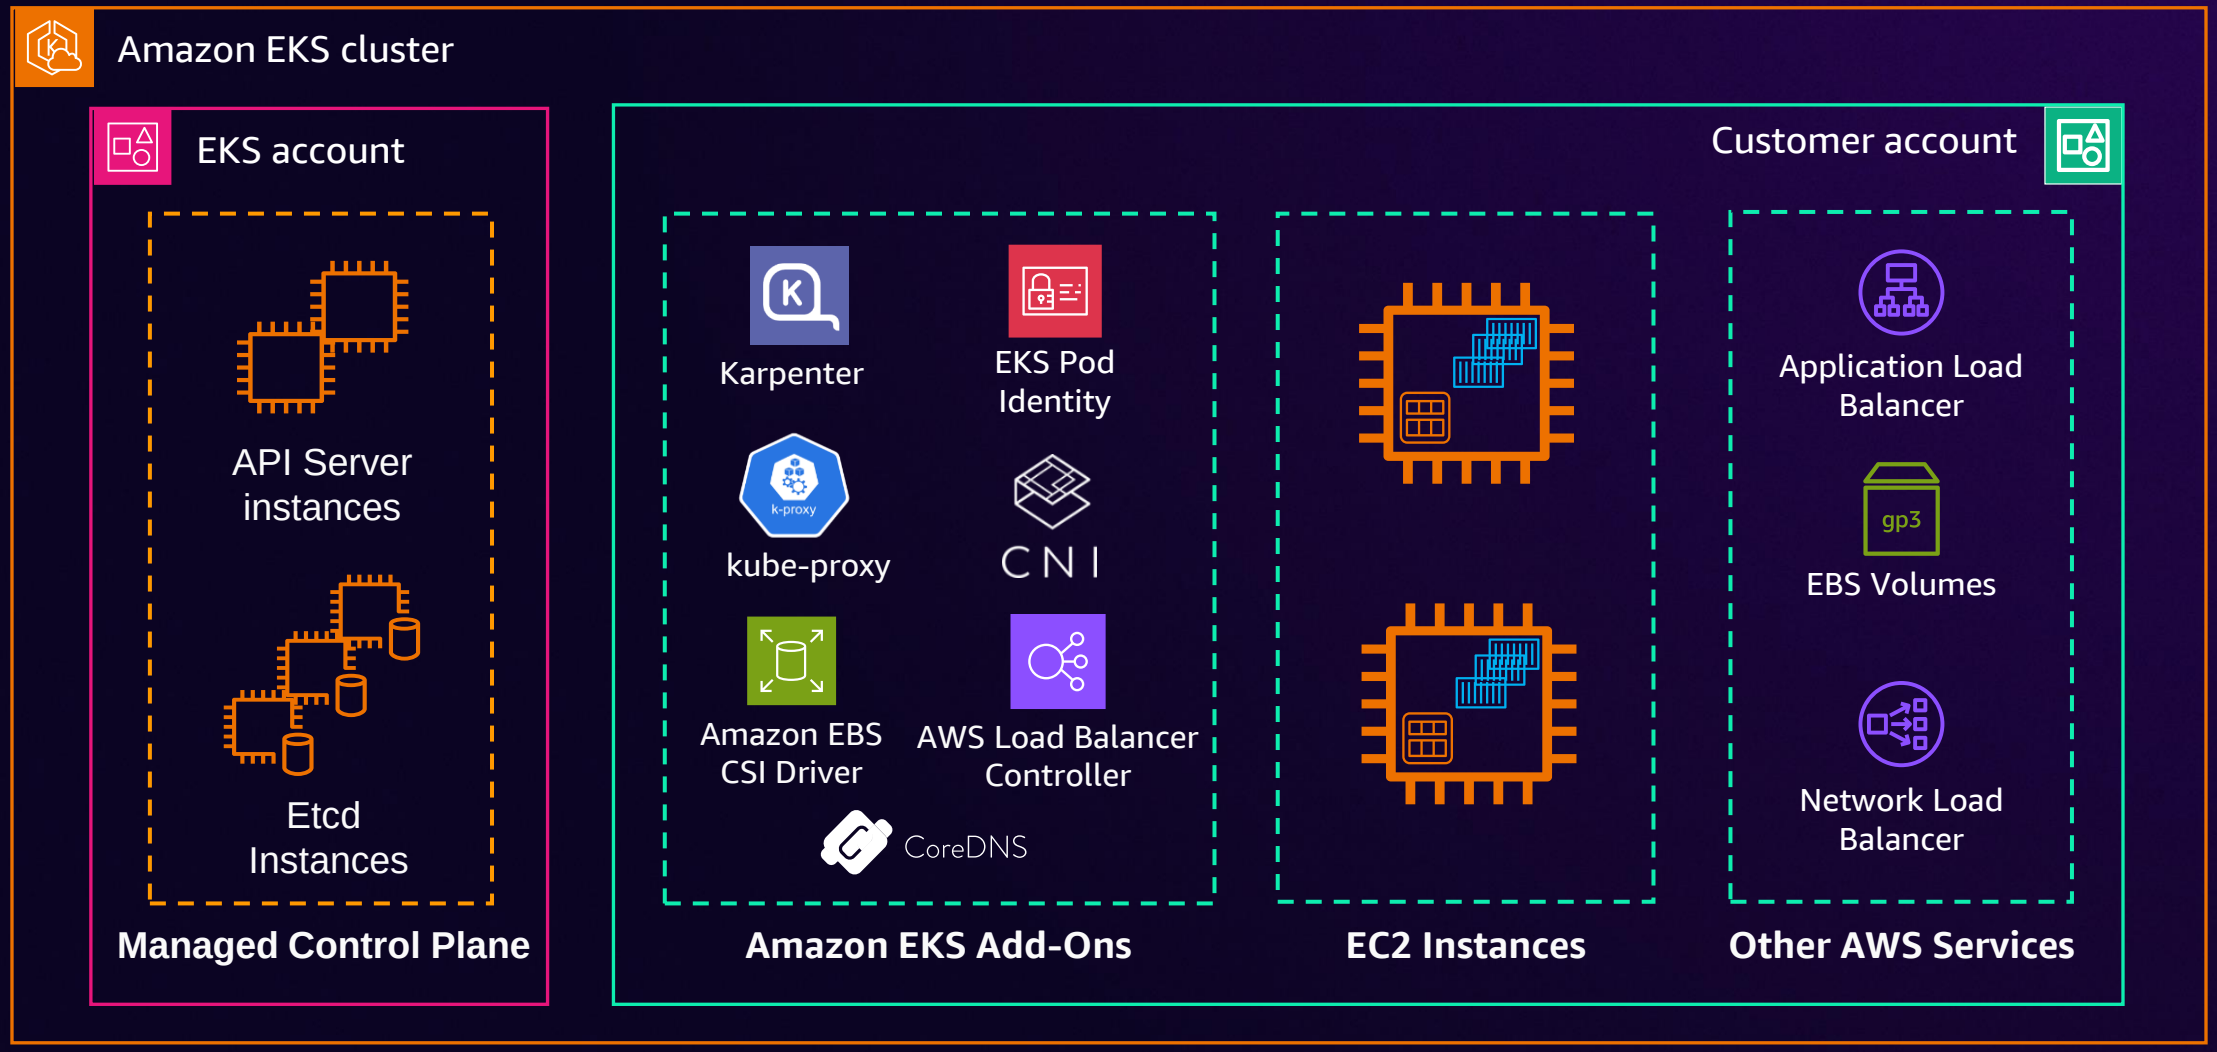

INTRODUCING AMAZON EKS

Auto Mode

Automate your entire Kubernetes cluster infrastructure

WHAT IS AMAZON EKS

Auto Mode

- Application-ready clusters, with essential Kubernetes capabilities
- Dynamic compute scaling, AWS-managed instances, secure and healthy by default
- Automatic compute optimization, simplified cluster upgrades

WHY AMAZON EKS

Auto Mode

Accelerate innovation by offloading cluster operations to AWS

Improve performance, availability, and security of your applications with AWS operational excellence

Optimize compute costs with automatic capacity planning, maintenance and dynamic scaling

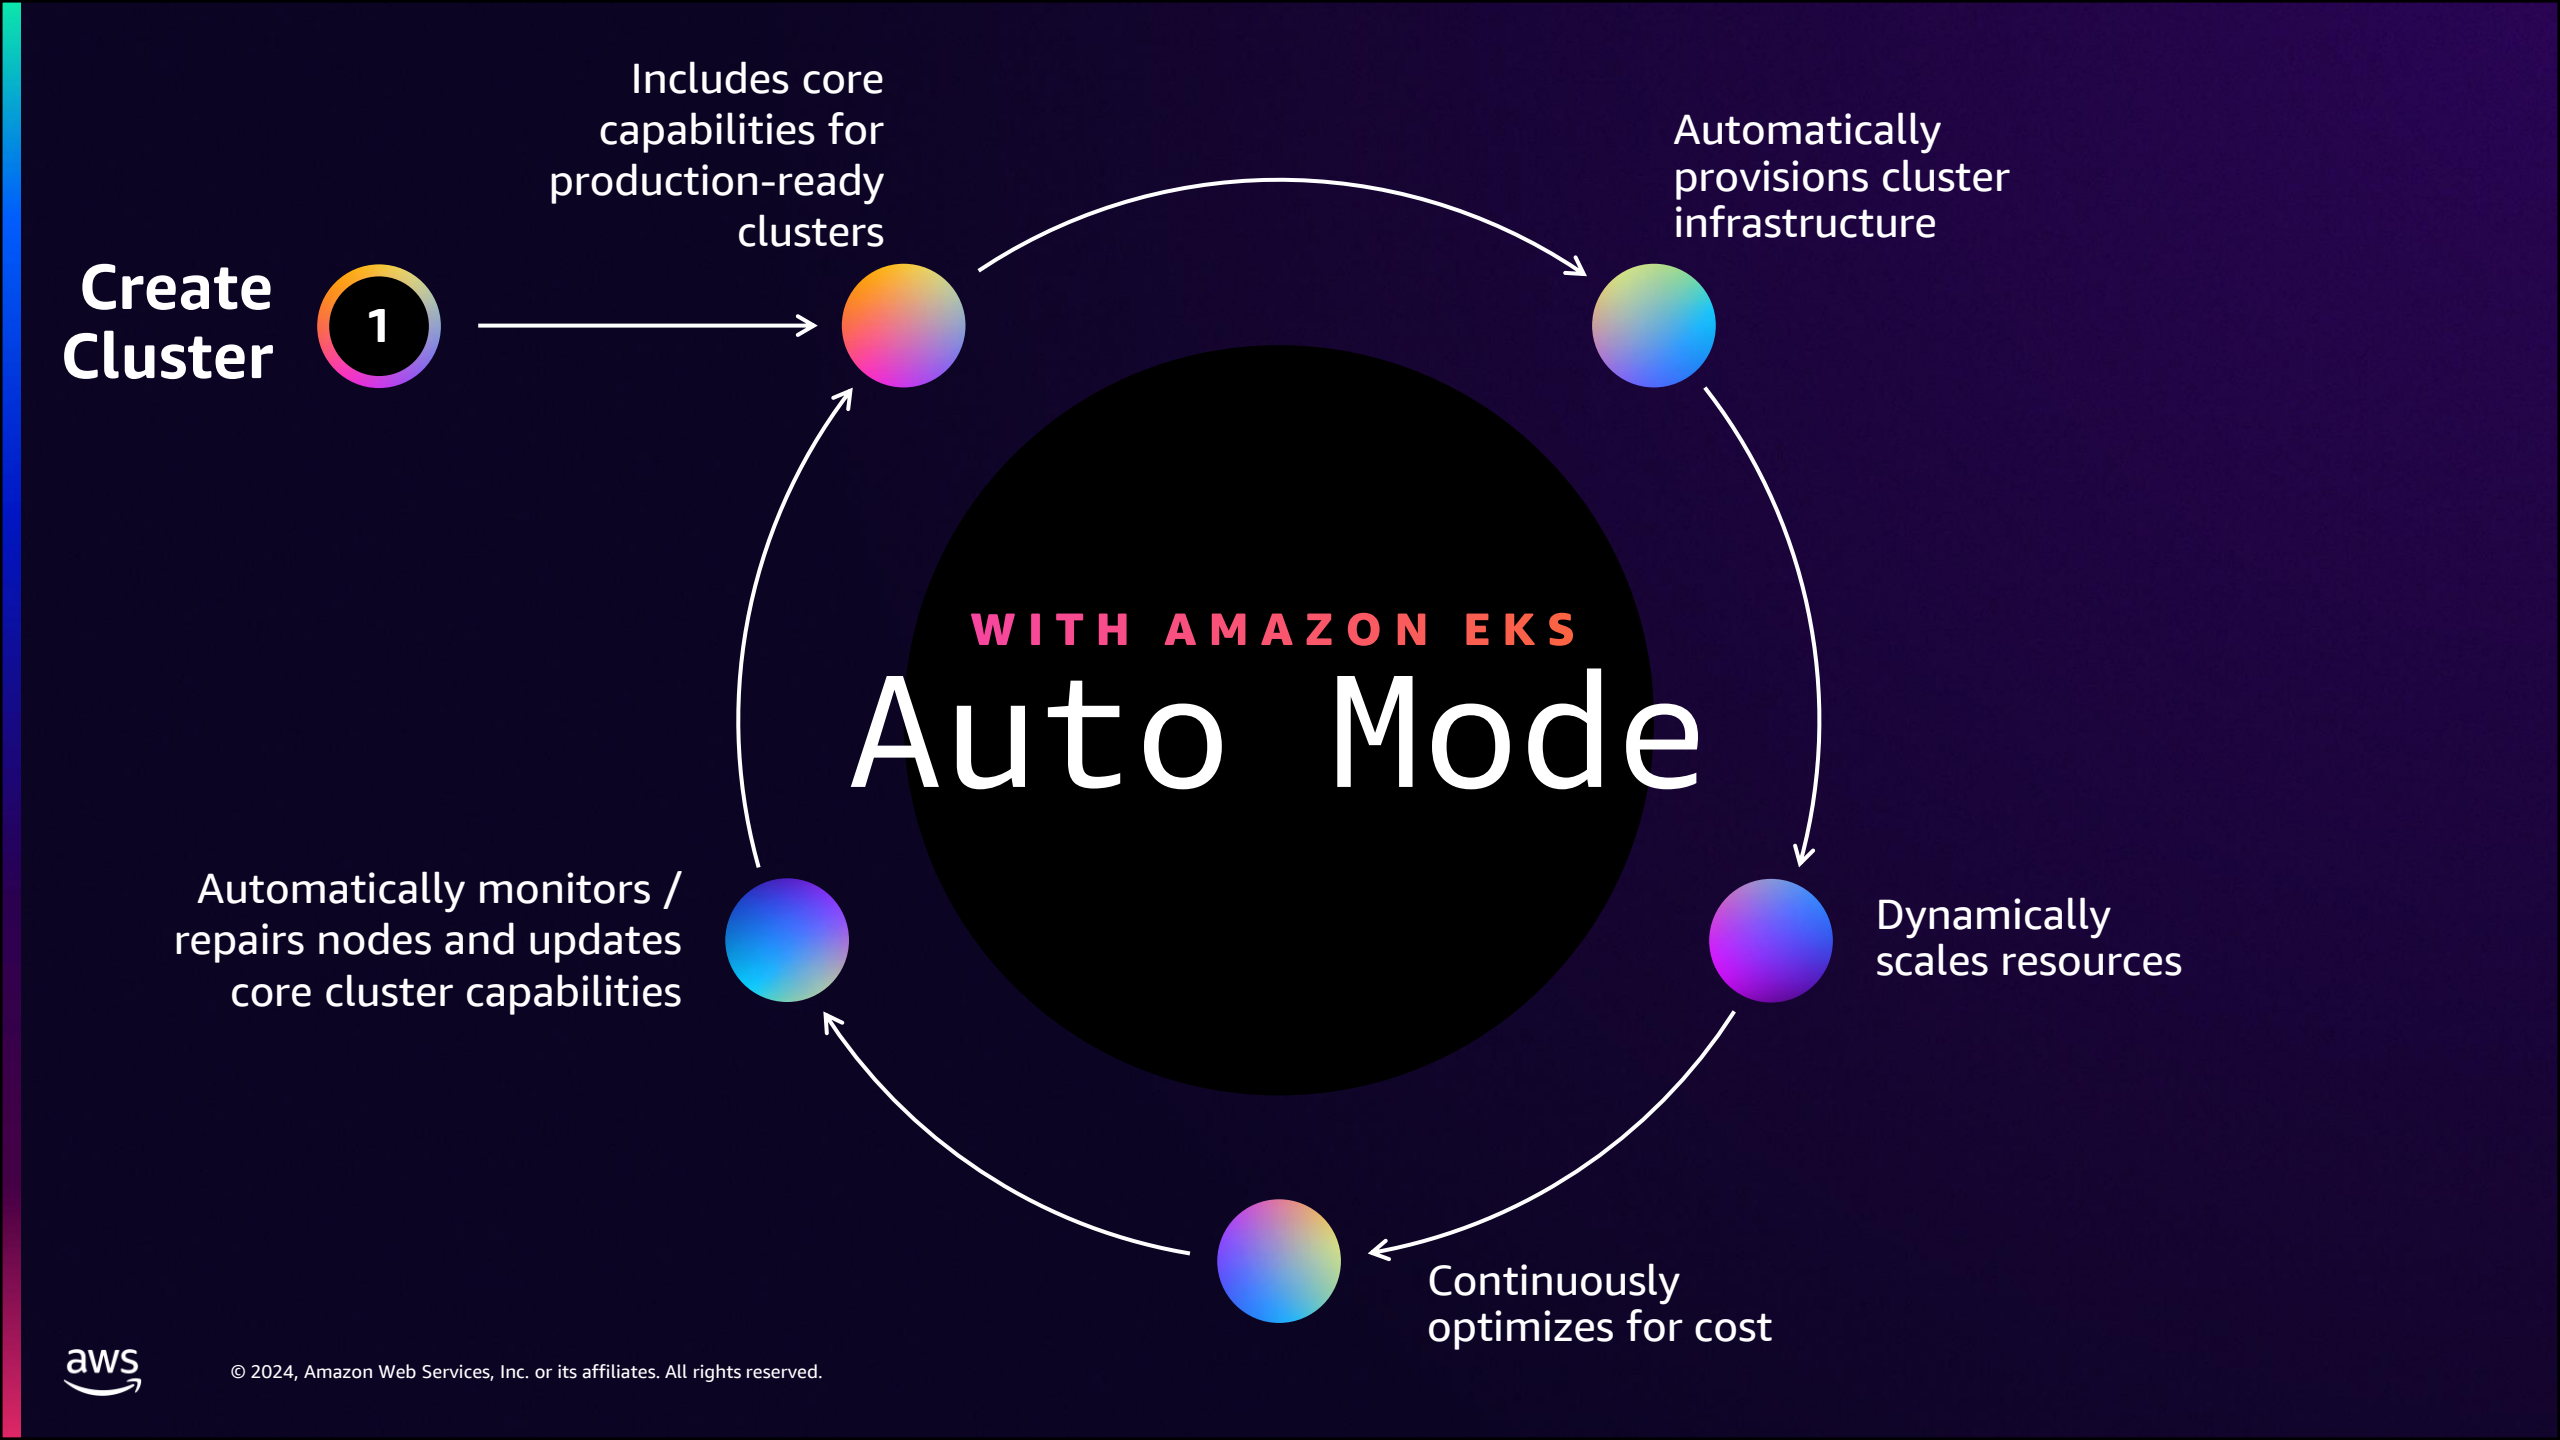

Cluster architecture with EKS Auto Mode

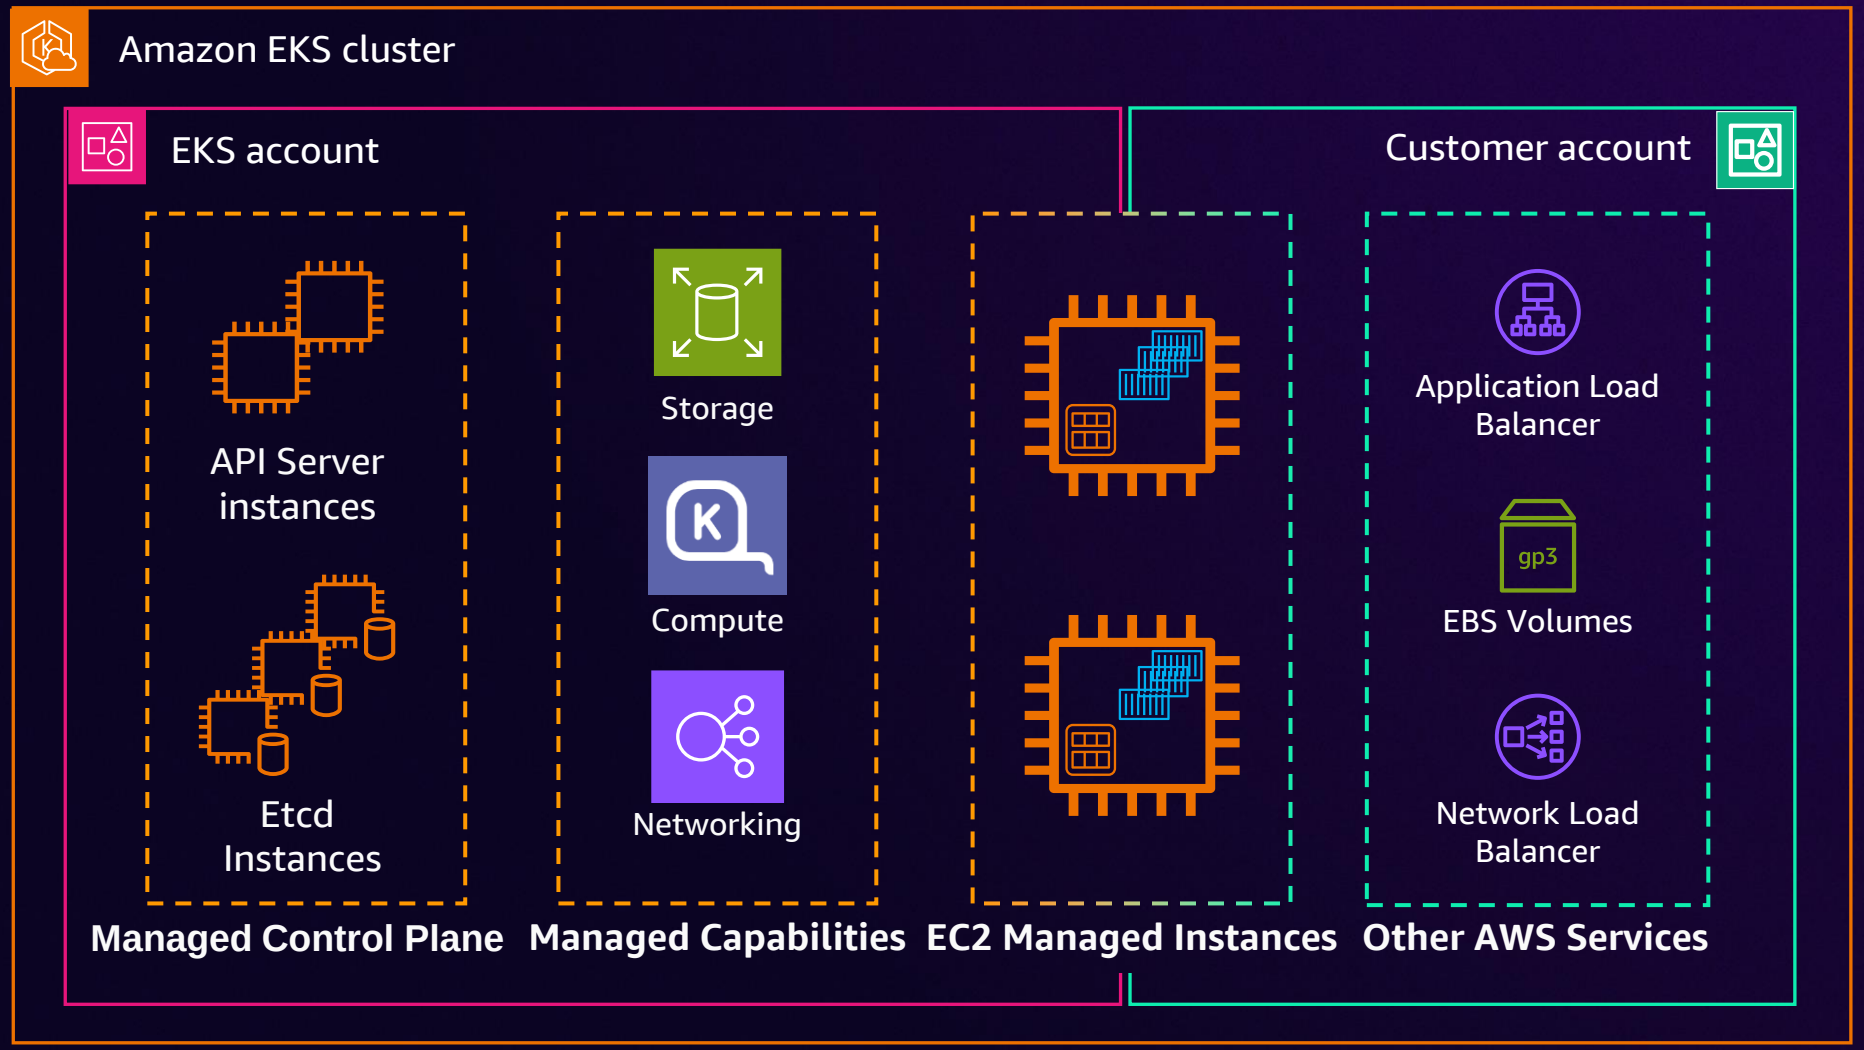

Shared Responsibility Model with EKS previously

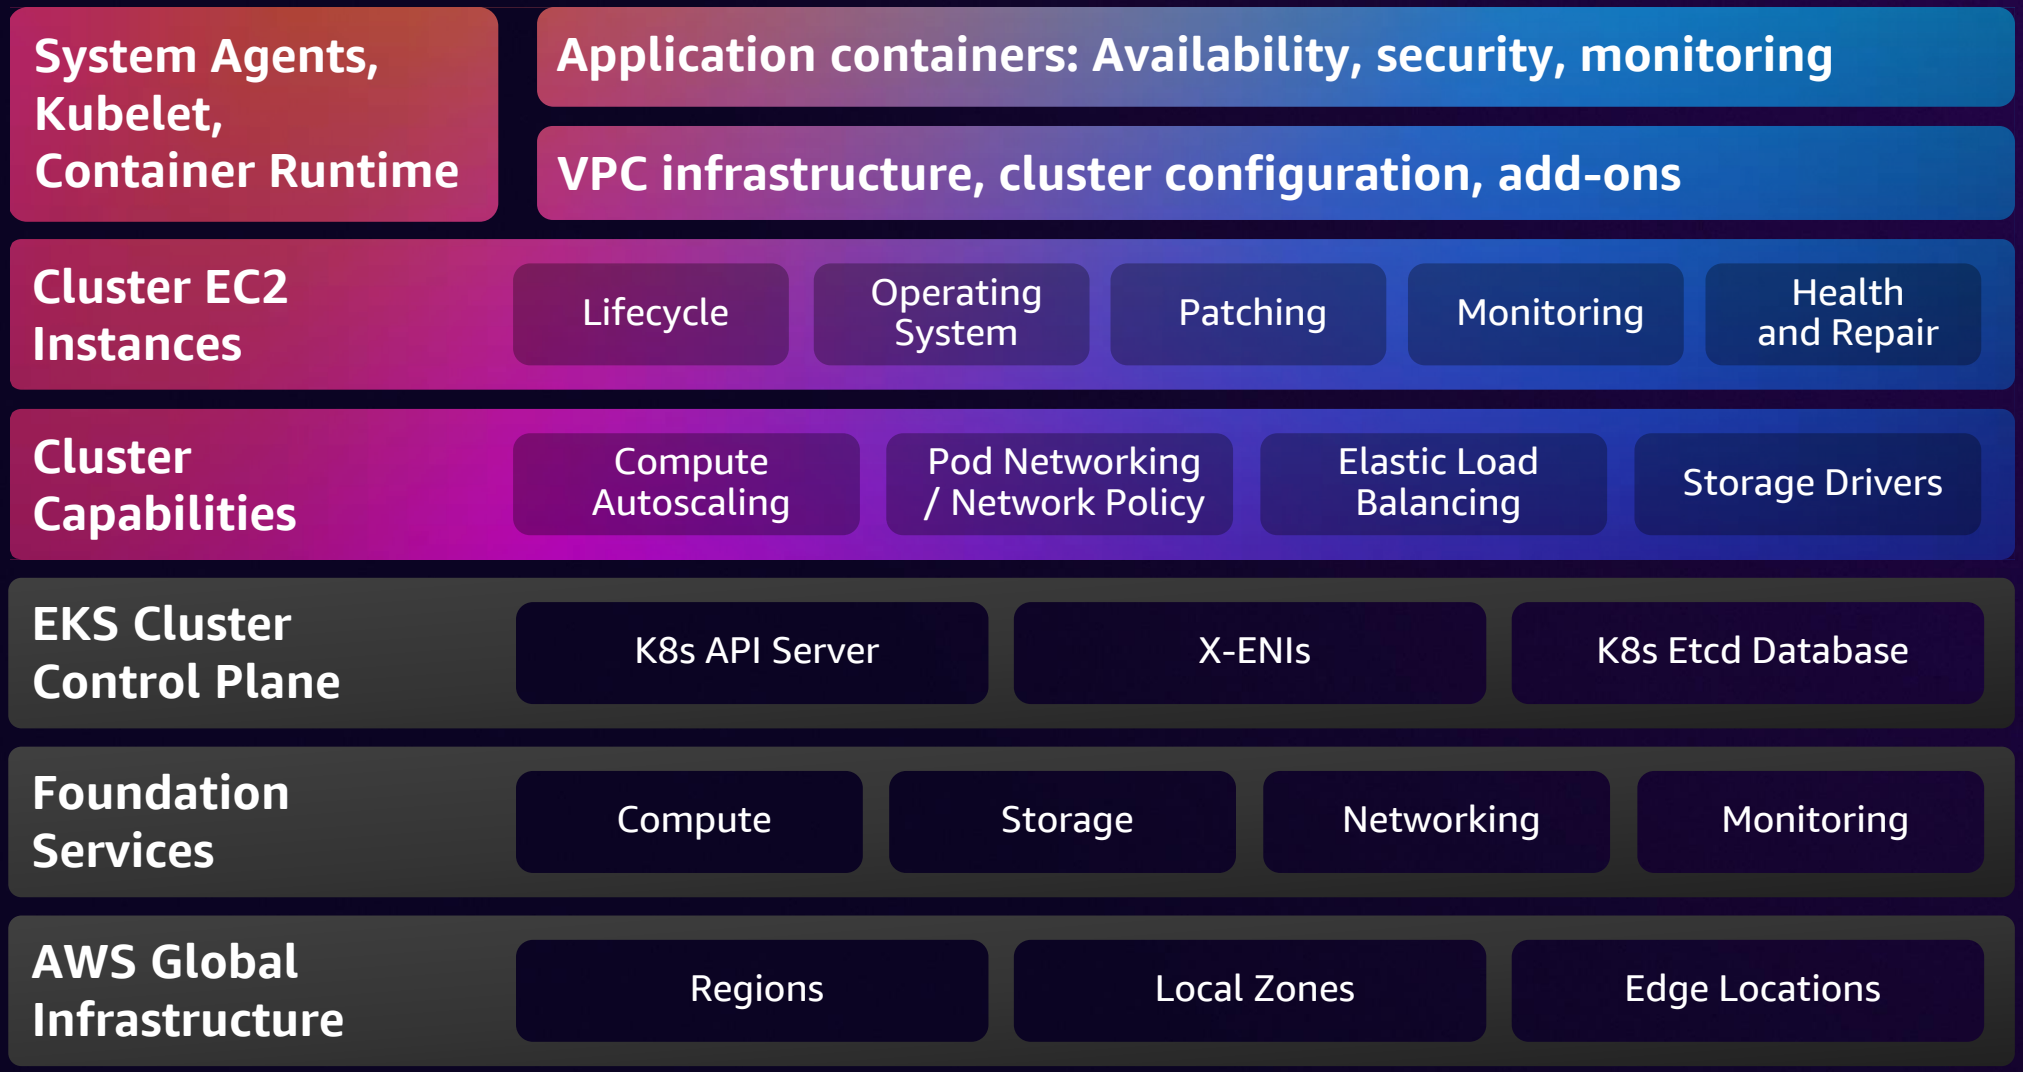

Managed by Customers

Managed by Amazon Web Services

Shared Responsibility Model with EKS Auto Mode

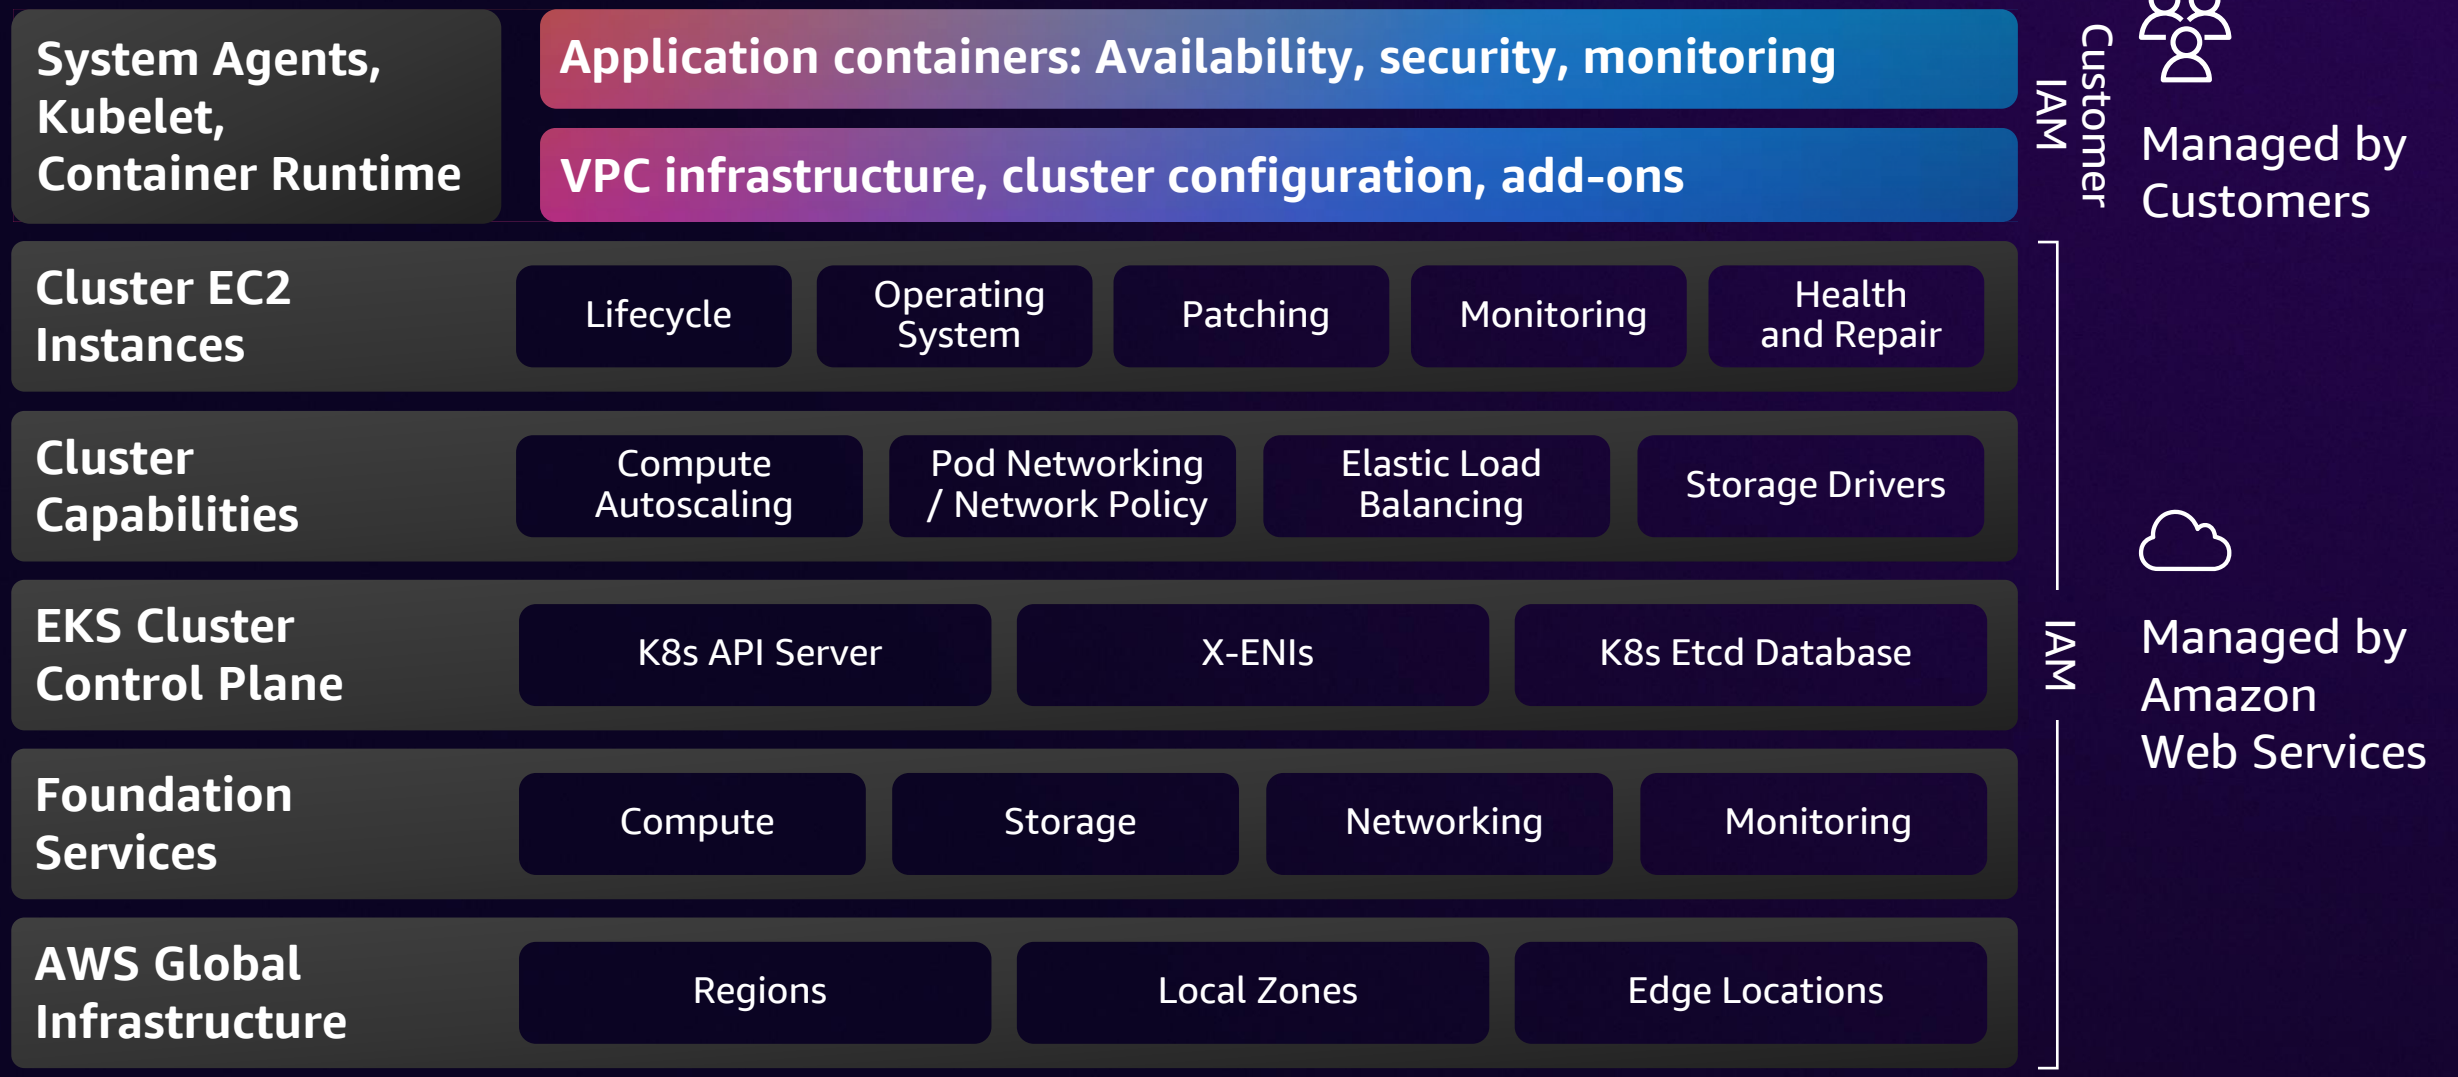

ASTRONOMER

is on a mission to deliver the world's data.

“Amazon EKS Auto Mode reduced cluster management time by 50%, allowing us to support exponentially more clusters per engineer. We've already seen a 20% reduction in compute costs and expect these efficiencies to continue.

Mike Pountney

Director of Software Engineering – Astronomer.io

FICO®

(NYSE: FICO)

is a leading analytics software company, helping businesses in 90+ countries make better decisions that drive higher levels of growth, profitability and customer satisfaction.

“Amazon EKS Auto Mode allows us to focus on our core business instead of spending time managing the underlying infrastructure. It empowers us to quickly get started on Amazon EKS, while at the same time enhancing our security footprint.

Johannes Koch

Sr Engineer Technical Architecture – FICO

Amazon EKS Auto Mode demo

Introducing Amazon EKS Auto Mode

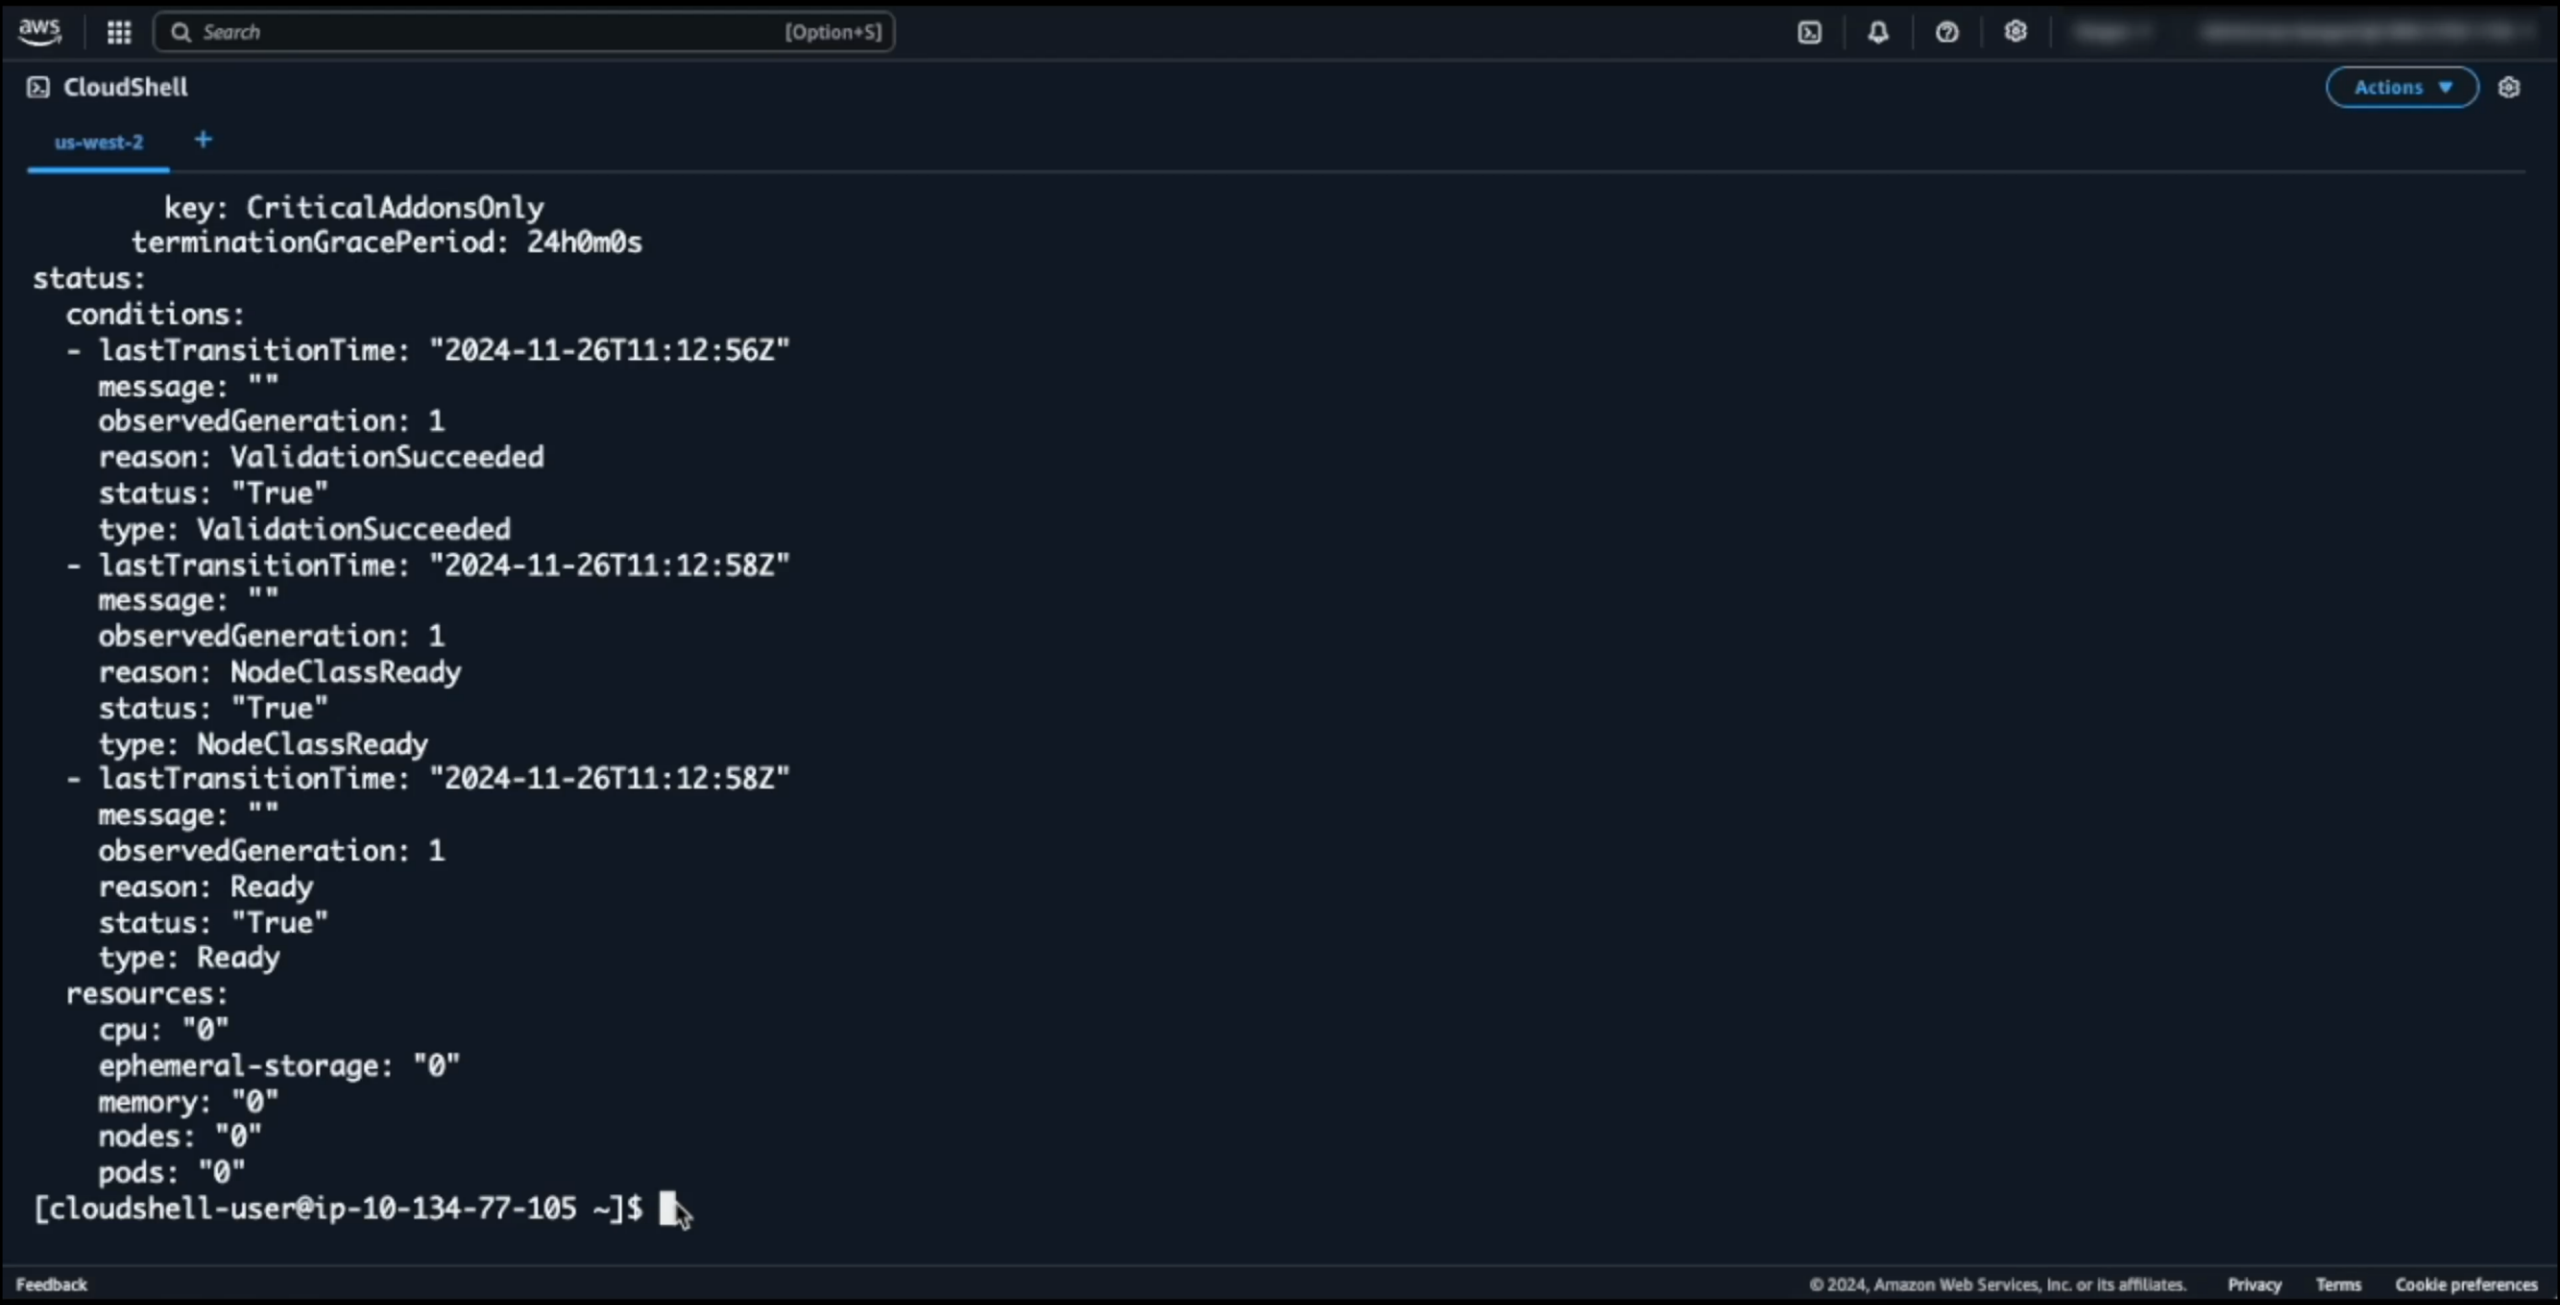
Search
[Option+S]
CloudShell
Actions
us-west-2
key: CriticalAddonsOnly
terminationGracePeriod: 24h0m0s
status:
conditions:
- lastTransitionTime: "2024-11-26T11:12:56Z"
message: ""
observedGeneration: 1
reason: ValidationSucceeded
status: "True"
type: ValidationSucceeded
- lastTransitionTime: "2024-11-26T11:12:58Z"
message: ""
observedGeneration: 1
reason: NodeClassReady
status: "True"
type: NodeClassReady
- lastTransitionTime: "2024-11-26T11:12:58Z"
message: ""
observedGeneration: 1
reason: Ready
status: "True"
type: Ready
resources:
cpu: "0"
ephemeral-storage: "0"
memory: "0"
nodes: "0"
pods: "0"
[cloudshell-user@ip-10-134-77-105 ~]$
Feedback
© 2024, Amazon Web Services, Inc. or its affiliates.
Privacy
Terms
Cookie preferences

Retail Store Sample Home Catalog Cart 1

Home / Checkout

Guest checkout

1 Customer 2 Shipping 3 Payment 4 Review

First name

Last name

Email

you@example.com

Address

1234 Main St

Address 2 (Optional)

Apartment or suite

City

State

Choose...

Zip

Summary

Total price: $385

Total: **$385**

Health Monitoring and Auto Repair

Amazon EKS Auto Mode deep dive

Amazon EKS Auto Mode capabilities

MANAGED CAPABILITIES PROVIDE COMPUTE, STORAGE, AND NETWORKING

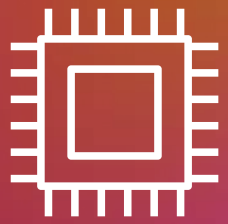

Compute

Compute autoscaling and continuous right-sizing

Container-optimized OS, secure by default

Health monitoring and auto-repair

Networking

Fully-managed, simplified VPC CNI and networking policy

Fully-managed load balancing for services and applications

Cluster DNS on every node for service discovery

Storage

Provision and manage block storage resources / lifecycles automatically as needed by applications

Amazon EKS Auto Mode compute

AUTOMATED PROVISIONING, SCALING, AND MANAGEMENT OF CLUSTER COMPUTE

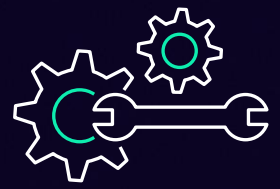

Amazon EC2 Managed Instances

Launches AWS-managed Amazon EC2 Instances in your account to provide fully managed compute for your applications

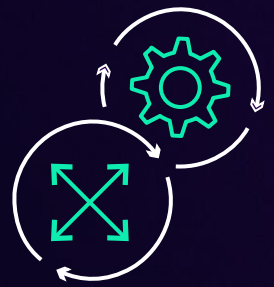

Dynamic selection and scaling

Provisions right-sized compute resources and dynamically scales based on demand for improved application availability

Continuous Cost Optimization

Removes underutilized nodes, replaces expensive nodes with cheaper alternatives, consolidating workloads to improve cost efficiency

Diverse workloads / purchase options

Access to breadth and depth of EC2 instance types and purchase options to provide you flexibility to choose the right compute

What is Bottlerocket?

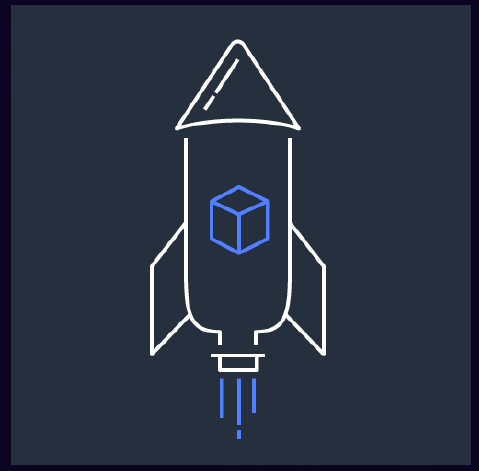

Bottlerocket is purpose-built to run containers. It doesn't include many packages, tools, interpreters, and dependencies in general purpose Linux distributions which aren't needed to run containers. This **improves security posture and reduces operational and security overhead**

Open-source, minimal container OS based on Amazon Linux

`github.com/bottlerocket-os/bottlerocket`

Amazon EKS Auto Mode networking

BUILT-IN POD / SERVICE NETWORKING AND CLUSTER DNS

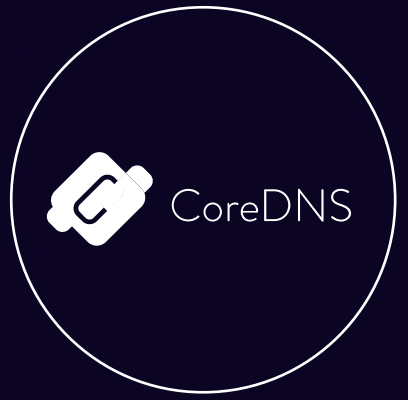

Out-of-the box **managed CoreDNS** on every node

Streamlined and **fully-managed Amazon VPC CNI** and **default network policy** enforcement

Fully-managed network proxy for cluster nodes

Run as systemd services on the node

Amazon EKS Auto Mode load balancing

MANAGED LOAD BALANCER CONTROLLER RUNS ON AWS-OPERATED INFRASTRUCTURE

- Pre-configured networking best-practices, e.g. prefix delegation, secondary IP mode fallback, etc.
- It satisfies Kubernetes Ingress resources by provisioning Application Load Balancers
- It satisfies Kubernetes Service resources by provisioning Network Load Balancers
- Create new ingress class, required for launch

NLB and ALB provisioned when you annotate Kubernetes service/ingress

Amazon EKS Auto Mode storage capability

MANAGED EBS CSI CONTROLLER RUNS ON AWS-OPERATED INFRASTRUCTURE

Amazon EBS CSI Driver

- Provision and manage the lifecycle of block storage resources with Amazon EBS CSI driver and node agent
- New StorageClass for EKS Auto Mode: `ebs.csi.eks.amazonaws.com`
- Create new storage class referencing the new provisioner, required at launch
- Other storage types, e.g. EFS, S3, available through existing EKS add-ons


```yaml
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: eks-auto-ebs-csi-sc
  annotations:
    storageclass.kubernetes.io/is-default-
class: "true"
provisioner: ebs.csi.eks.amazonaws.com
volumeBindingMode: WaitForFirstConsumer
parameters:
  type: gp3
```

Works with Kubernetes, works with EKS Auto Mode

ALWAYS KUBERNETES CONFORMANT AND UPSTREAM COMPATIBLE

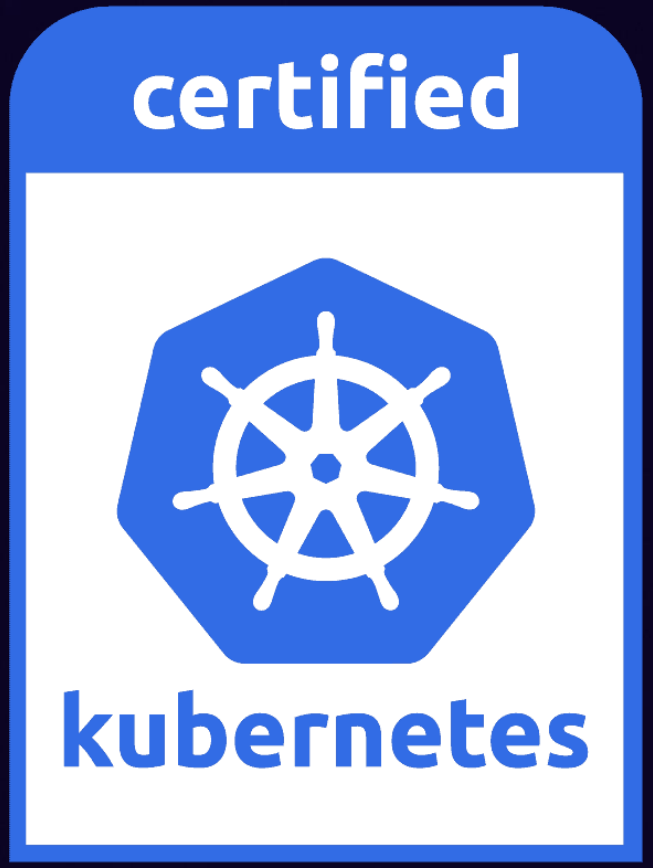

- Kubernetes conformant
- Leverage the whole CNCF Landscape of partners and tools
- Recommended for nearly all use cases

Amazon EKS Auto Mode validated partners


Pulumi

HashiCorp

aqua

kubecost

solo.io

WIZ


DATADOG


CROWDSTRIKE


dynatrace


tetrate


Kong


nirmata


AccuKnox

Amazon EKS Auto Mode summary

ENABLING YOU TO FOCUS ON ACTIVITIES CRITICAL TO YOUR BUSINESS

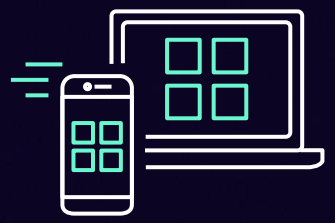

Increase agility and accelerate innovation, by offloading Kubernetes cluster management to AWS

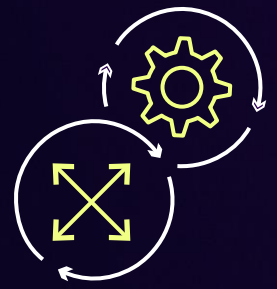

Reduce operational overhead, with streamlined updates and a simpler user experience

Lower costs, increase application availability, and improve performance by dynamically right-sizing compute

Secure by default with automated OS patches, updates and ephemeral compute to limit security risks

Learn more about Amazon EKS Auto Mode

SCAN QR CODES

Learn more about Amazon EKS Auto Mode

Register for EKS Auto Mode webinar

Read the launch blog

Streamline Kubernetes cluster management with new Amazon EKS Auto Mode

To get started refer **EKS Auto Mode User Guide**

Thank you!

Please complete the session survey in the mobile app

Alex Kestner
kestnera@amazon.com

Todd Neal
tnealt@amazon.com