
aws
SUMMIT
LONDON | 27 APRIL 2022

AR-X02

Towards Continuous Resilience

Veliswa Boya
Sr Developer Advocate
Amazon Web Services

Veliswa Boya, Snr Developer Advocate

Based in Johannesburg, South Africa

Past Industries: Banks and insurance companies

Past roles: Software Developer, Solutions Architect

……”because production outages”

The cost of downtime

$5,600 per minute >>> Gartner

$1million per hour >>> IDC

A 101 Startups

What other big headline outages can you remember from the past?

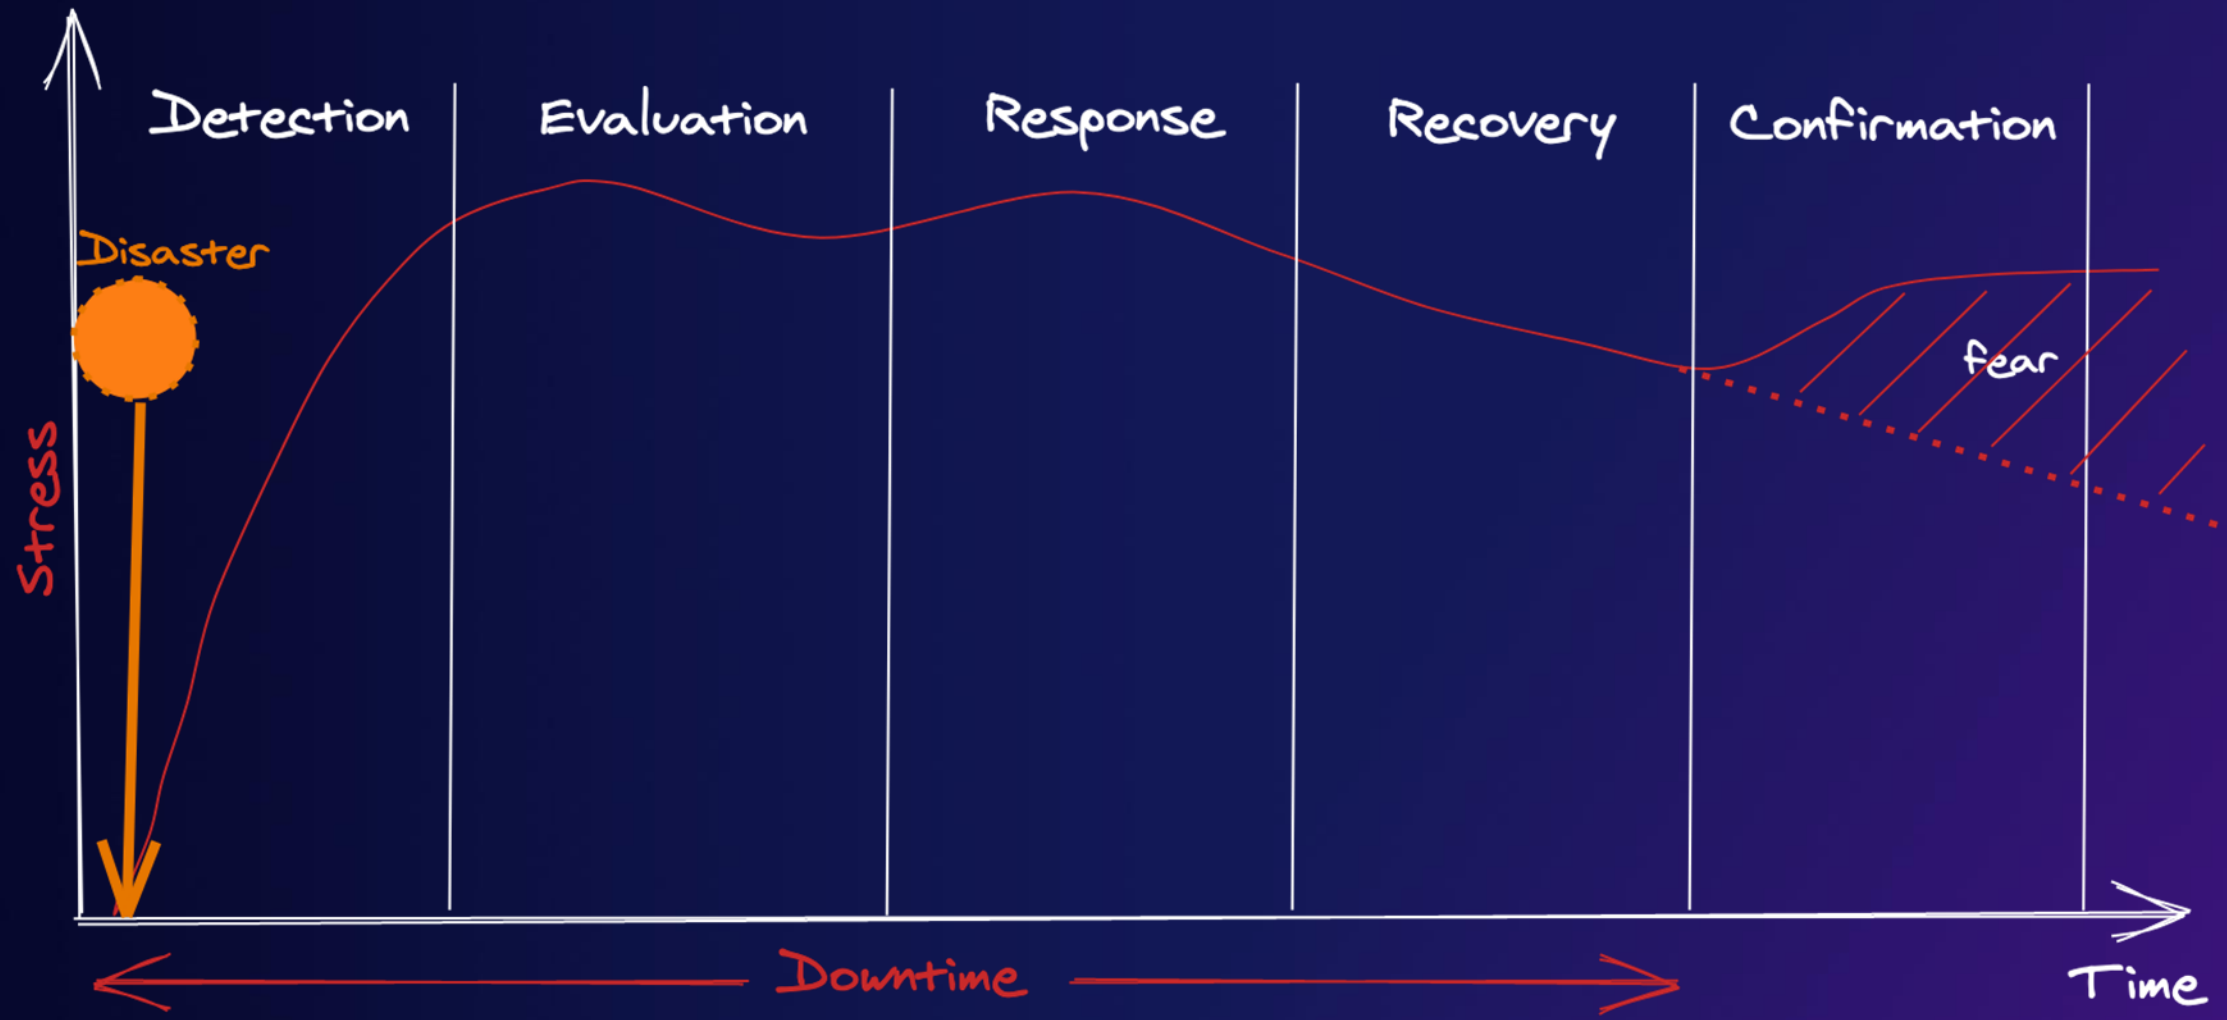

Avoid outages

Processes

Rules

Controls

Change management

Maximize robustness

End-to-end, fuzz, and regression testing

Code freezes

Segregation of duties

Advisory boards for changes

Planned deployments

Outsourcing with strict SLAs and penalties

> **“Failures are a given, and everything will eventually fail over time.”**

Werner Vogels
CTO, Amazon.com

How do you trade fear of failure for a growth opportunity and improvement?

Embracing failure and resilience

Resilience: Ability of a system to respond to, absorb, adapt to, and eventually recover from unexpected conditions

Four essential capabilities in a resilient system

ANTICIPATE

Understanding what to expect; imagining potential failures and mitigating those in advance

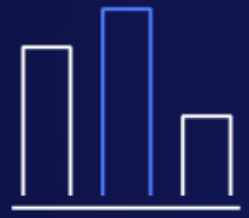

MONITORING

Understanding what to look for, in both internal and external conditions

RESPONDING

Understanding what to do and adjusting responses, if necessary, in a flexible way

LEARNING

Understanding and sharing what has happened to promote learning and changes

"Resilience Engineering In Practice", by Nemeth C., Hollnagel E. and Dekker S.

Four essential capabilities in a resilient system

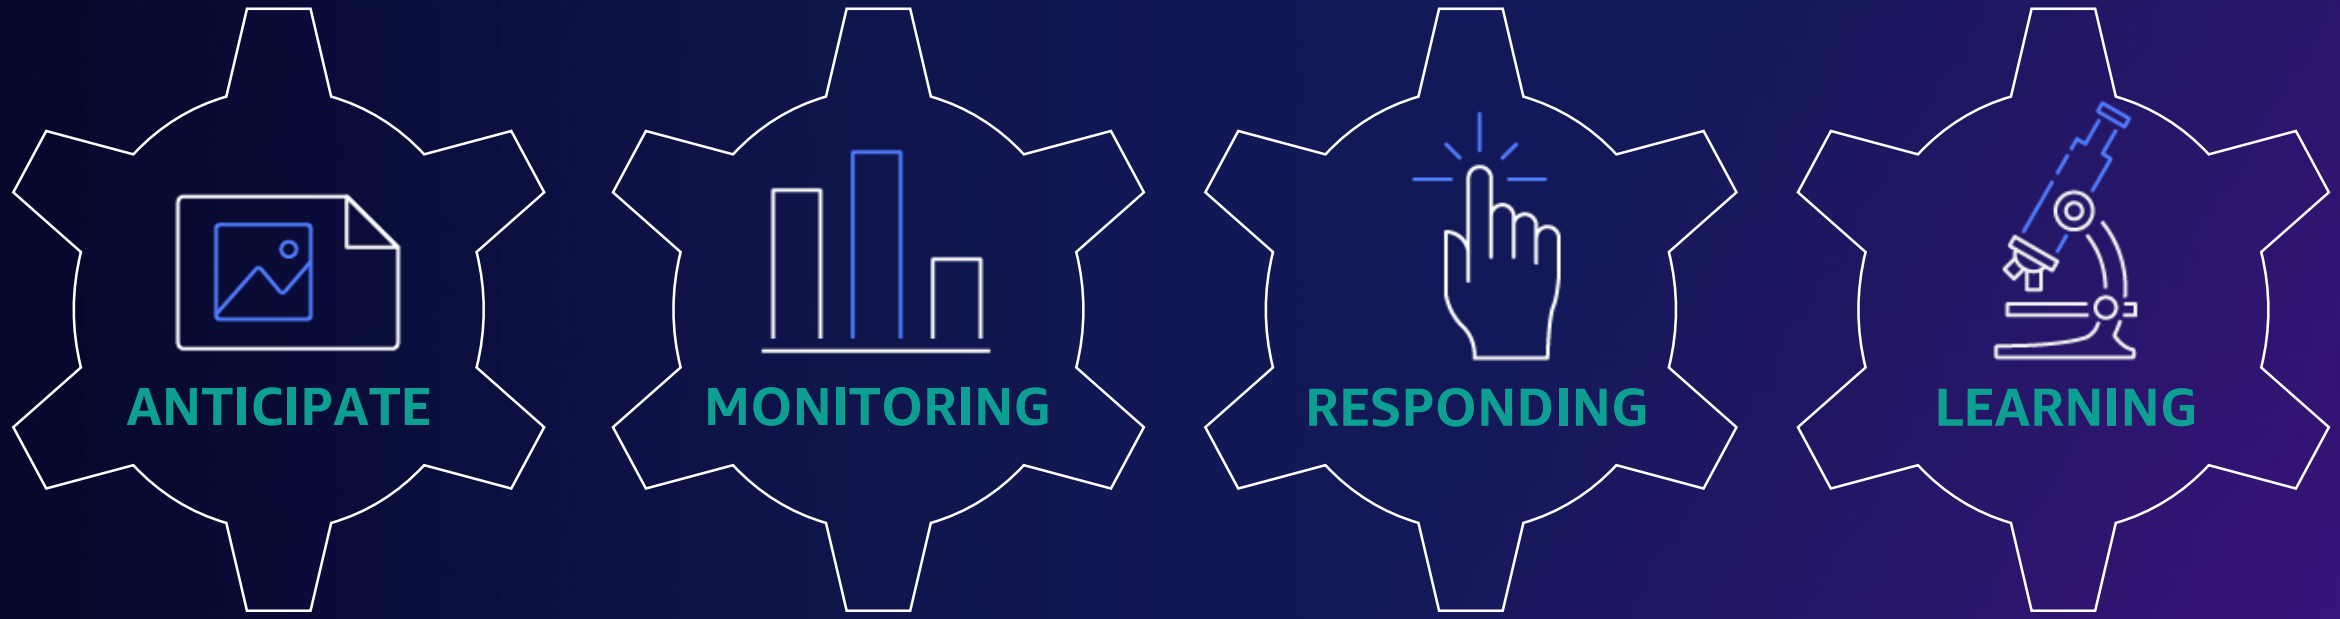

"Resilience Engineering In Practice", by Nemeth C., Hollnagel E. and Dekker S.

Continuous resilience*

*Non-exhaustive list

- Health checks
- Tracing
- Alarming
- Dashboarding

- Code review
- Resilient application patterns
- Immutability
- Simple designs
- Area-of-impact reduction patterns
- ML-powered recommendations

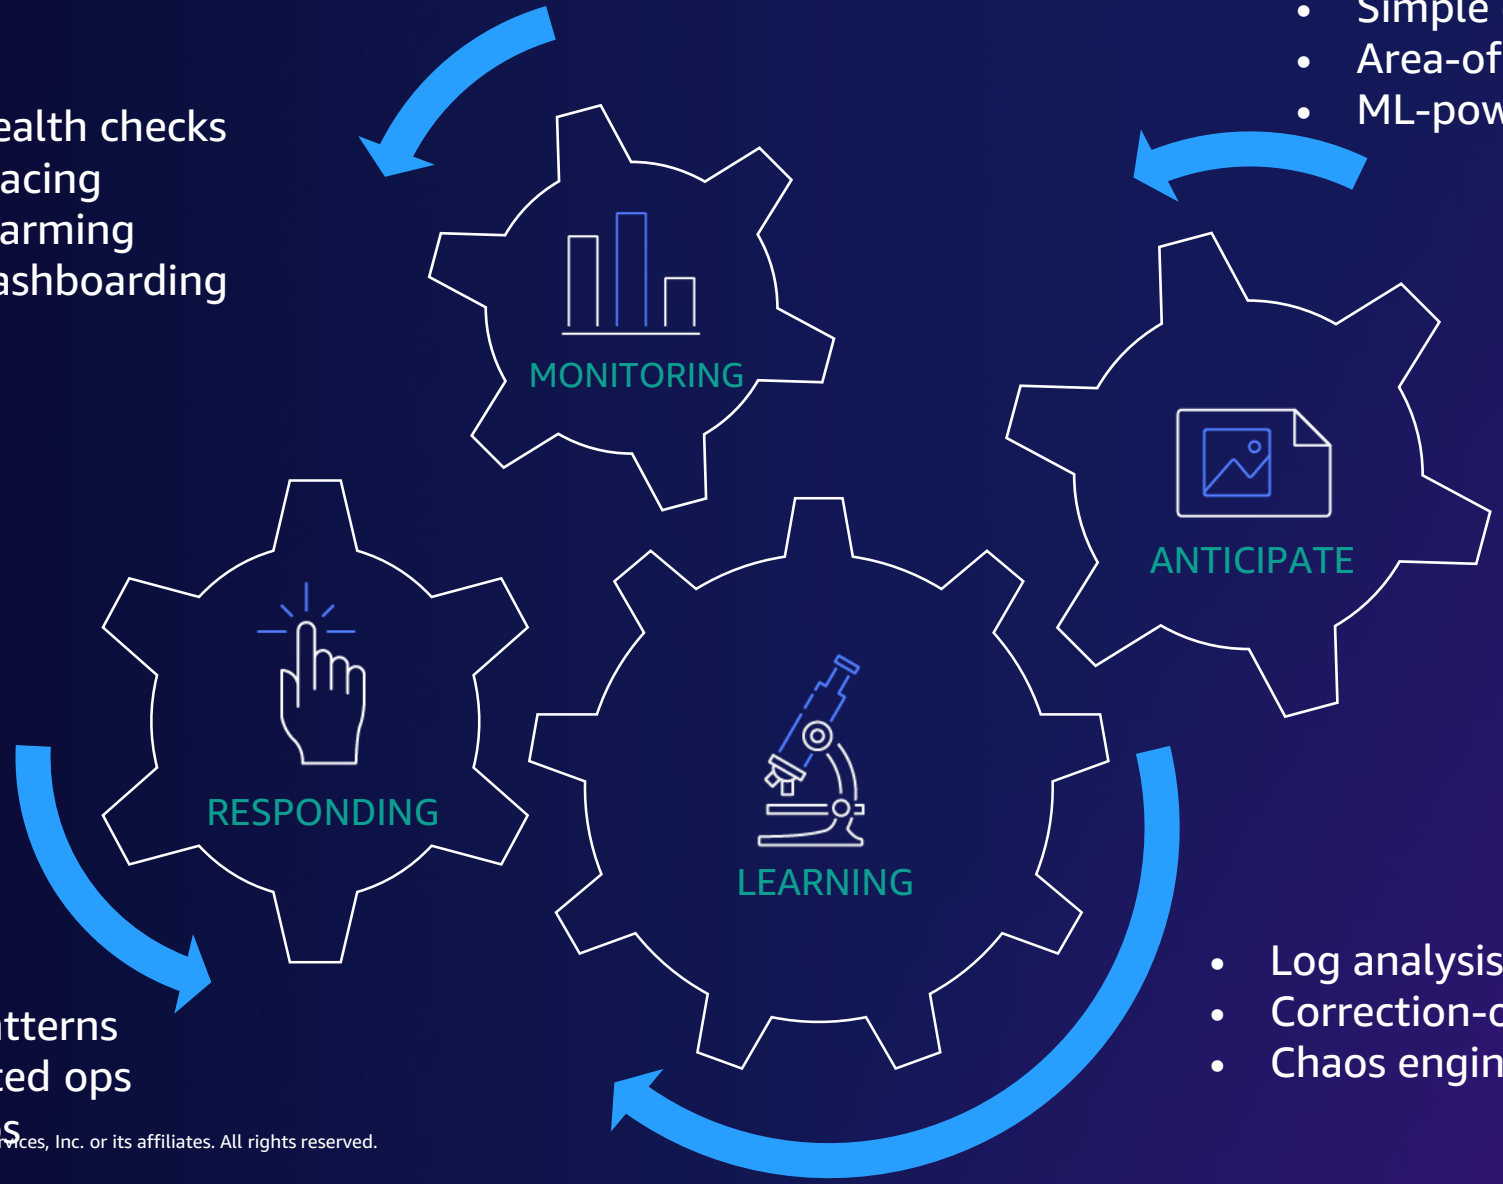

- Event-driven patterns
- Recovery-oriented ops
- ML-powered ops

- Log analysis
- Correction-of-errors (COE)
- Chaos engineering

Resilience:
Anticipate

Continuous resilience

ANTICIPATE

- Code review
- Resilient application patterns
- Immutability
- Simple designs
- Area-of-impact reduction patterns
- ML-powered recommendations

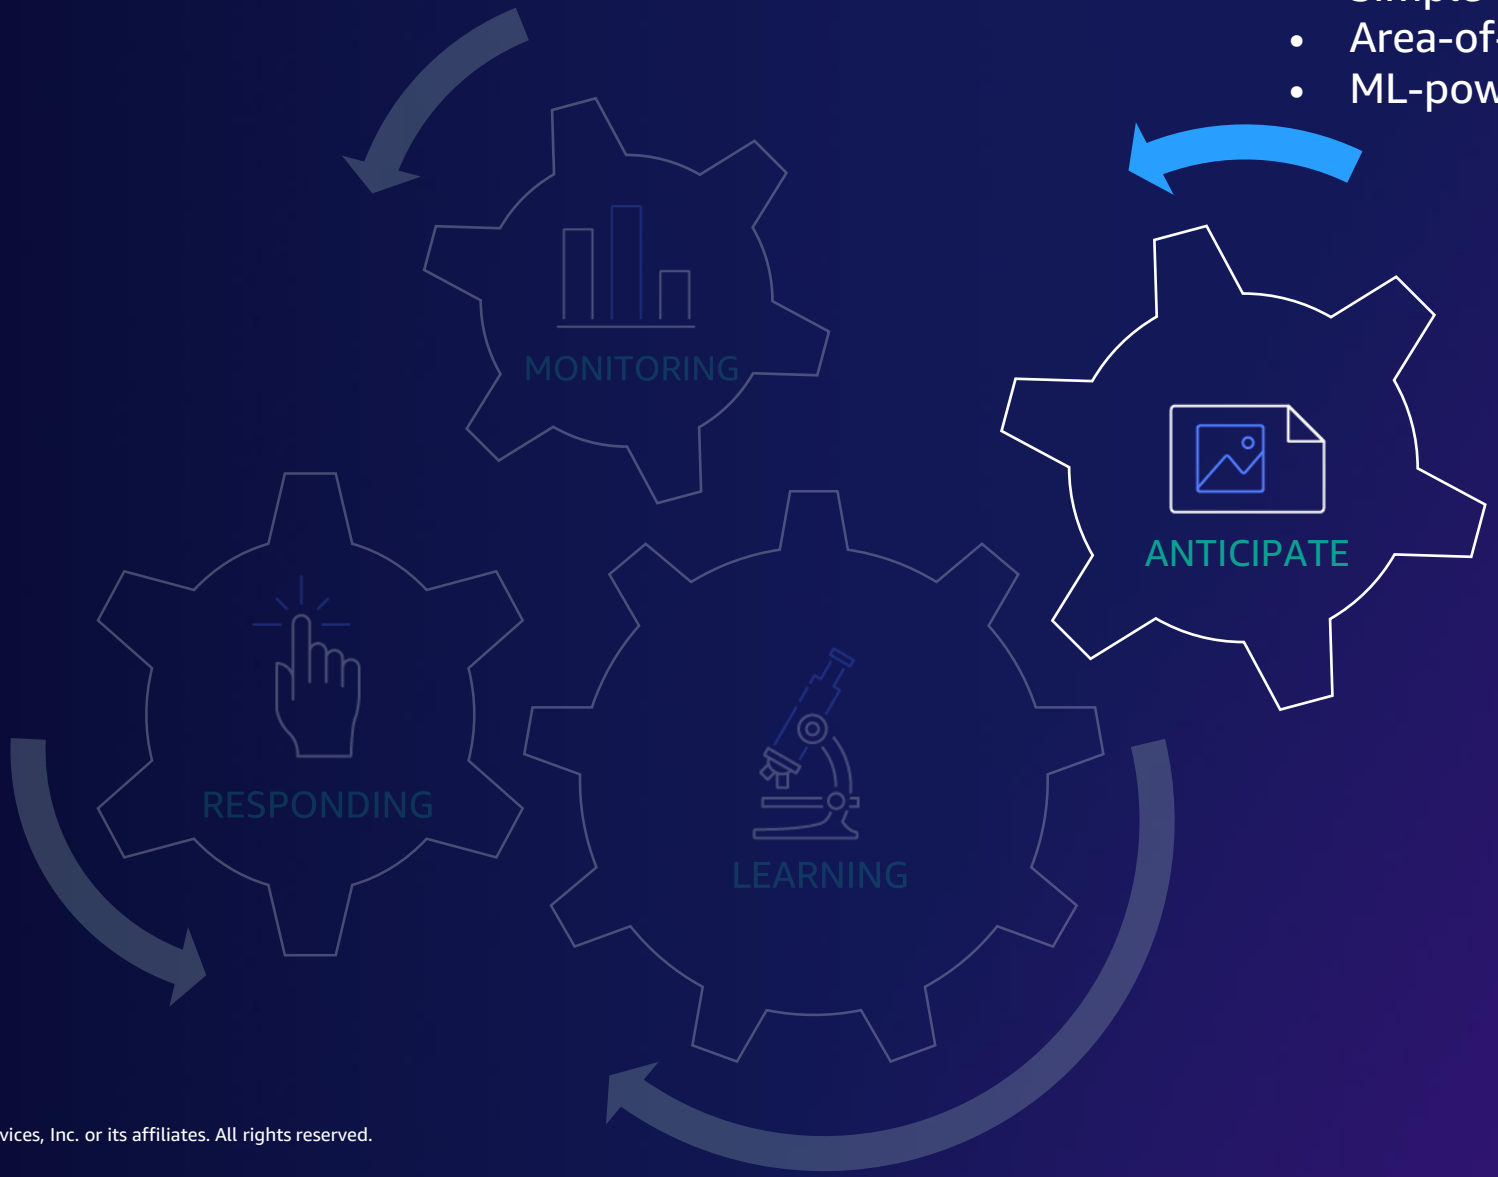

Code review and profilers at Prime Day 2018

WITH AMAZON CODEGURU

Code is considered good if it:

- Does what it should
- Follows a consistent style
- Is easy to understand
- Is well-documented
- Can be tested

+325%

Efficiency increase in CPU utilization YoY

-39%

Lower cost YoY

Code review and profiling with Amazon CodeGuru

USING ML TO CODE REVIEW AND OPTIMIZE HIGH-PERFORMING APPLICATIONS

CodeGuru Reviewer

CodeGuru Profiler

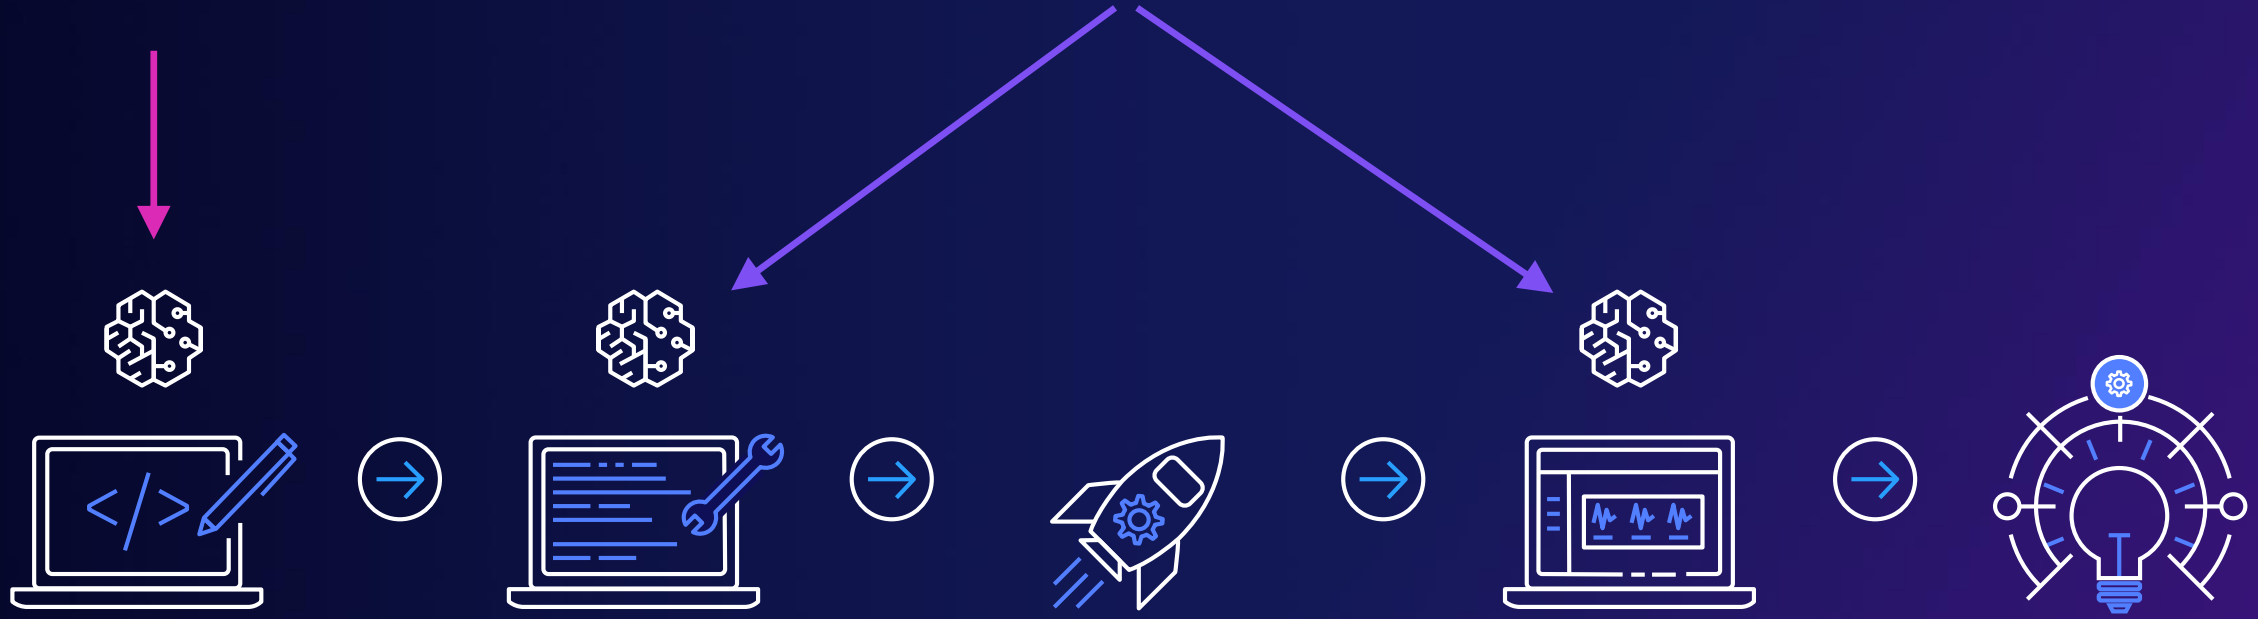

CODING

Built-in code reviews with actionable recommendations

BUILD & TEST

Detect and optimize the expensive lines of code preproduction

DEPLOY

MEASURE

Easily identify performance and cost improvements in production environment

Resilient application patterns

Service (client)

- Set timeouts
- Retries w/back off
  - Add jitter
  - Set maximum retries
- Limit queue sizes

Backend

- Rate limit
- Rejection (load shedding)

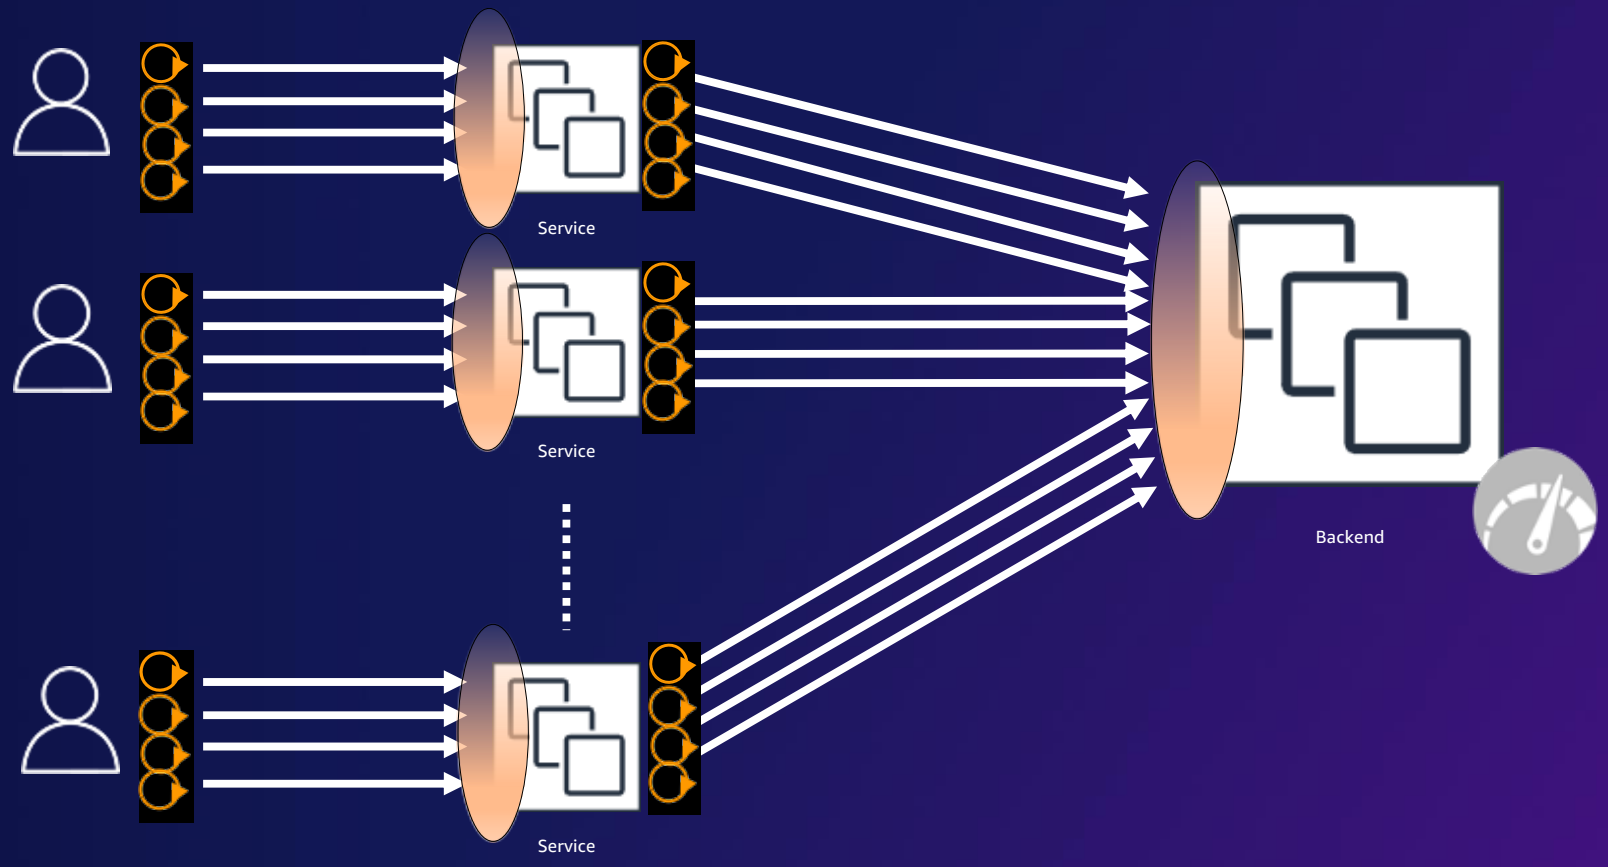

Timeouts, retries, and backoff with jitter

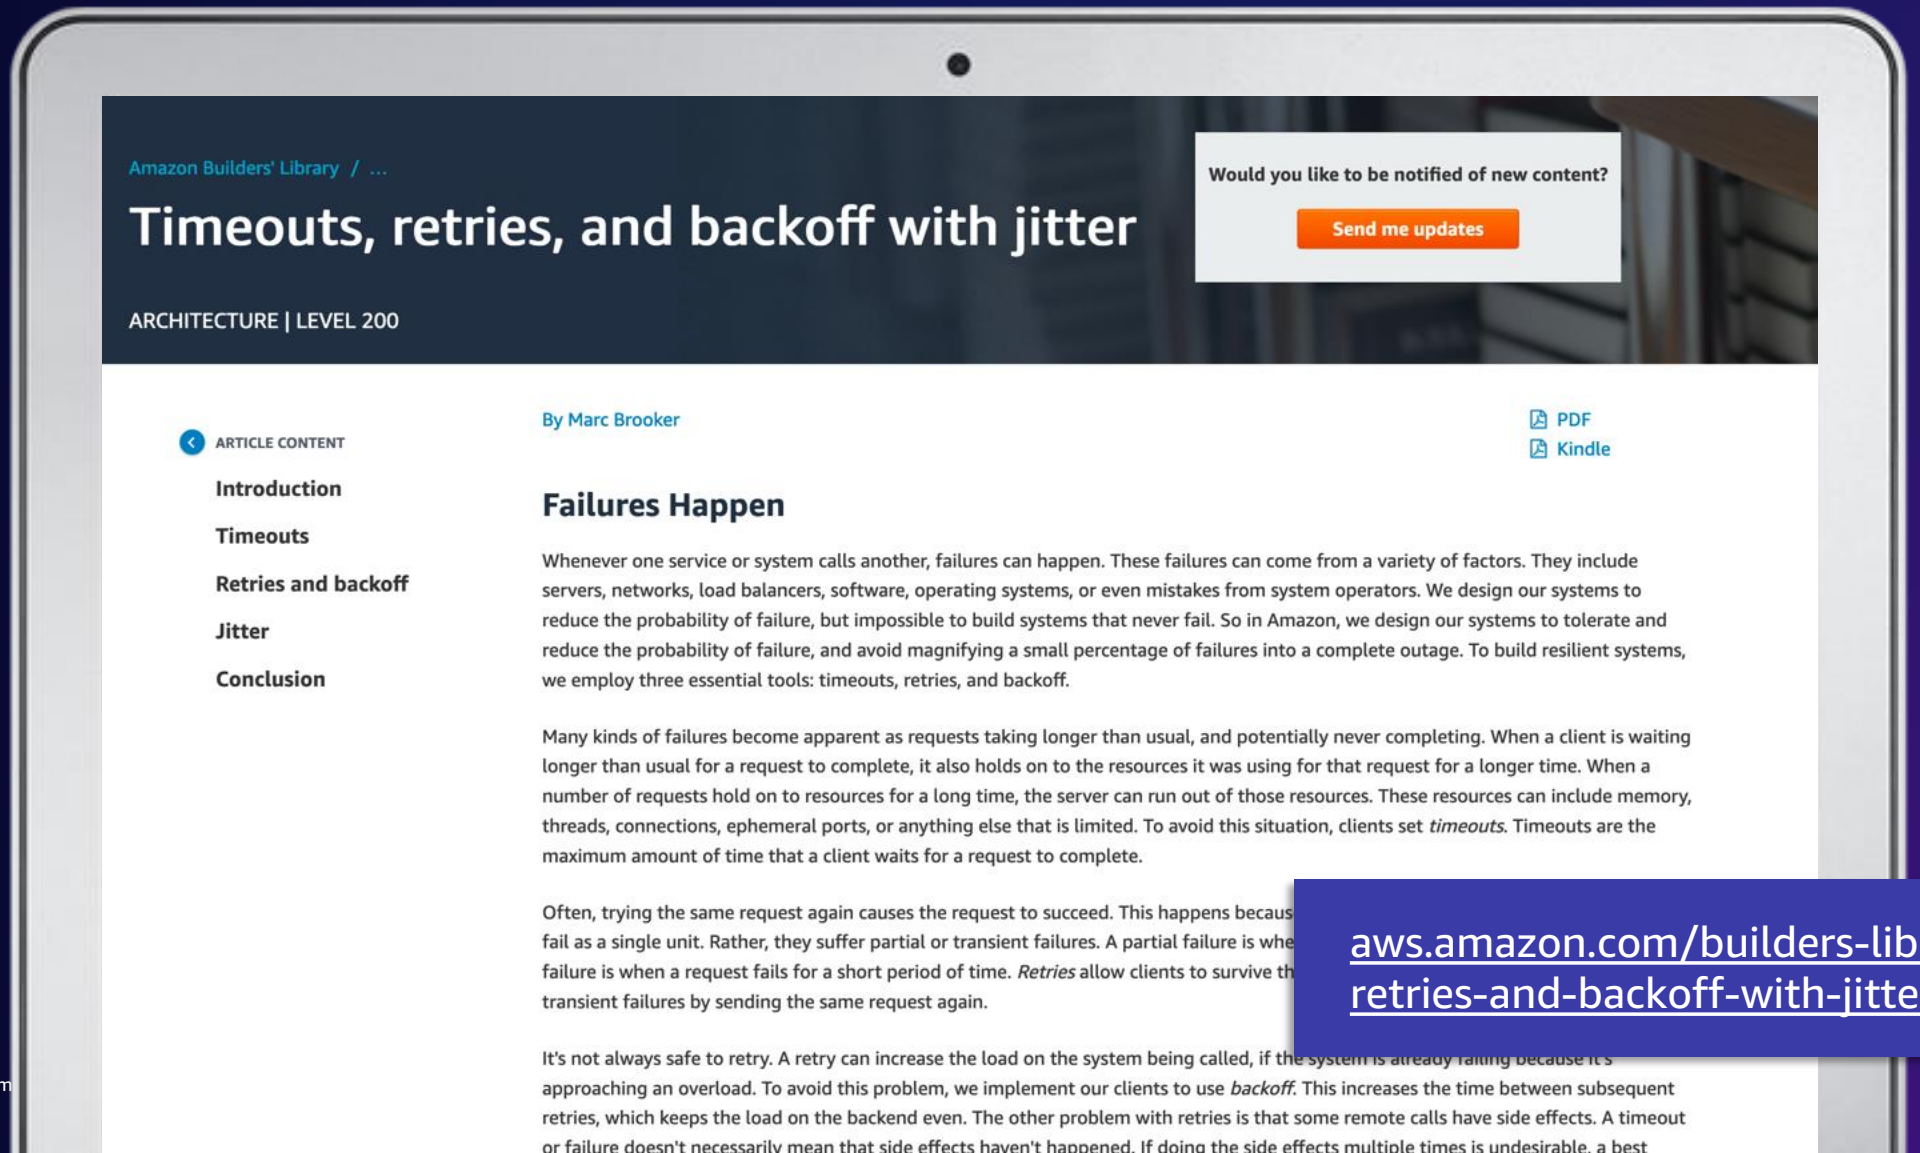

aws.amazon.com/builders-library/timeouts-retries-and-backoff-with-jitter/

Simple designs and constant work

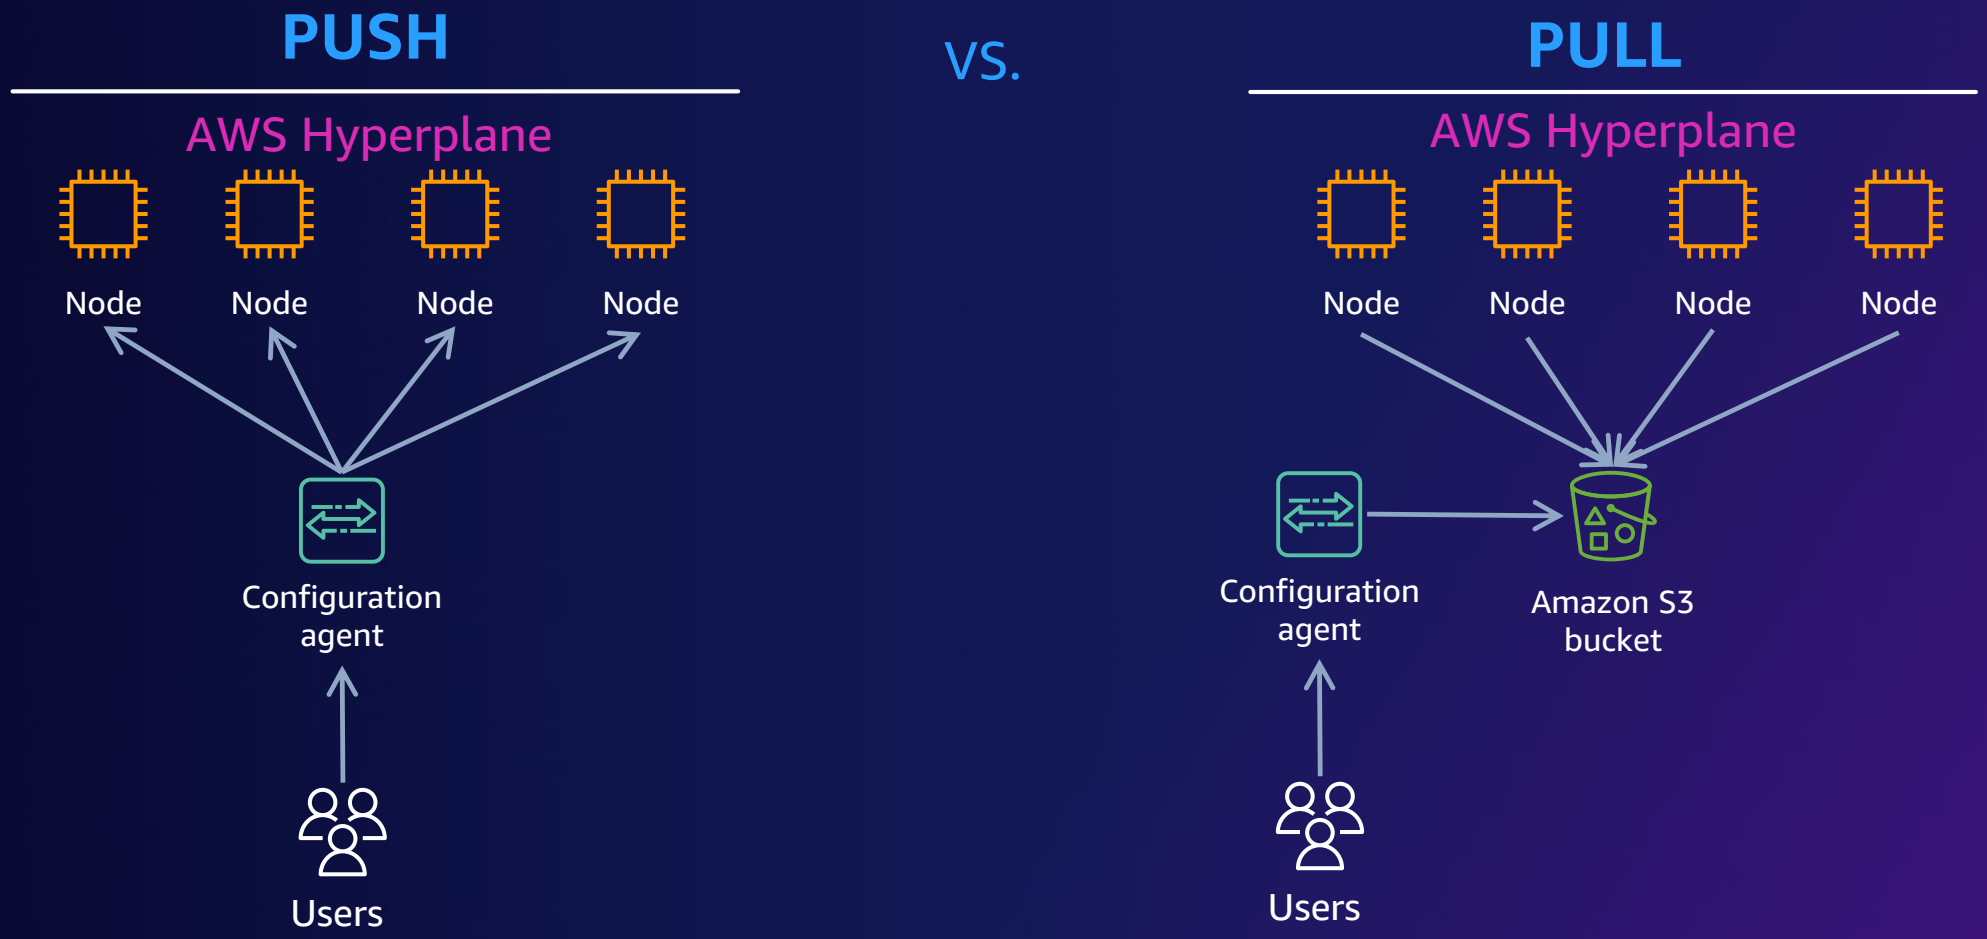

Reliability, constant work, and a good cup of coffee

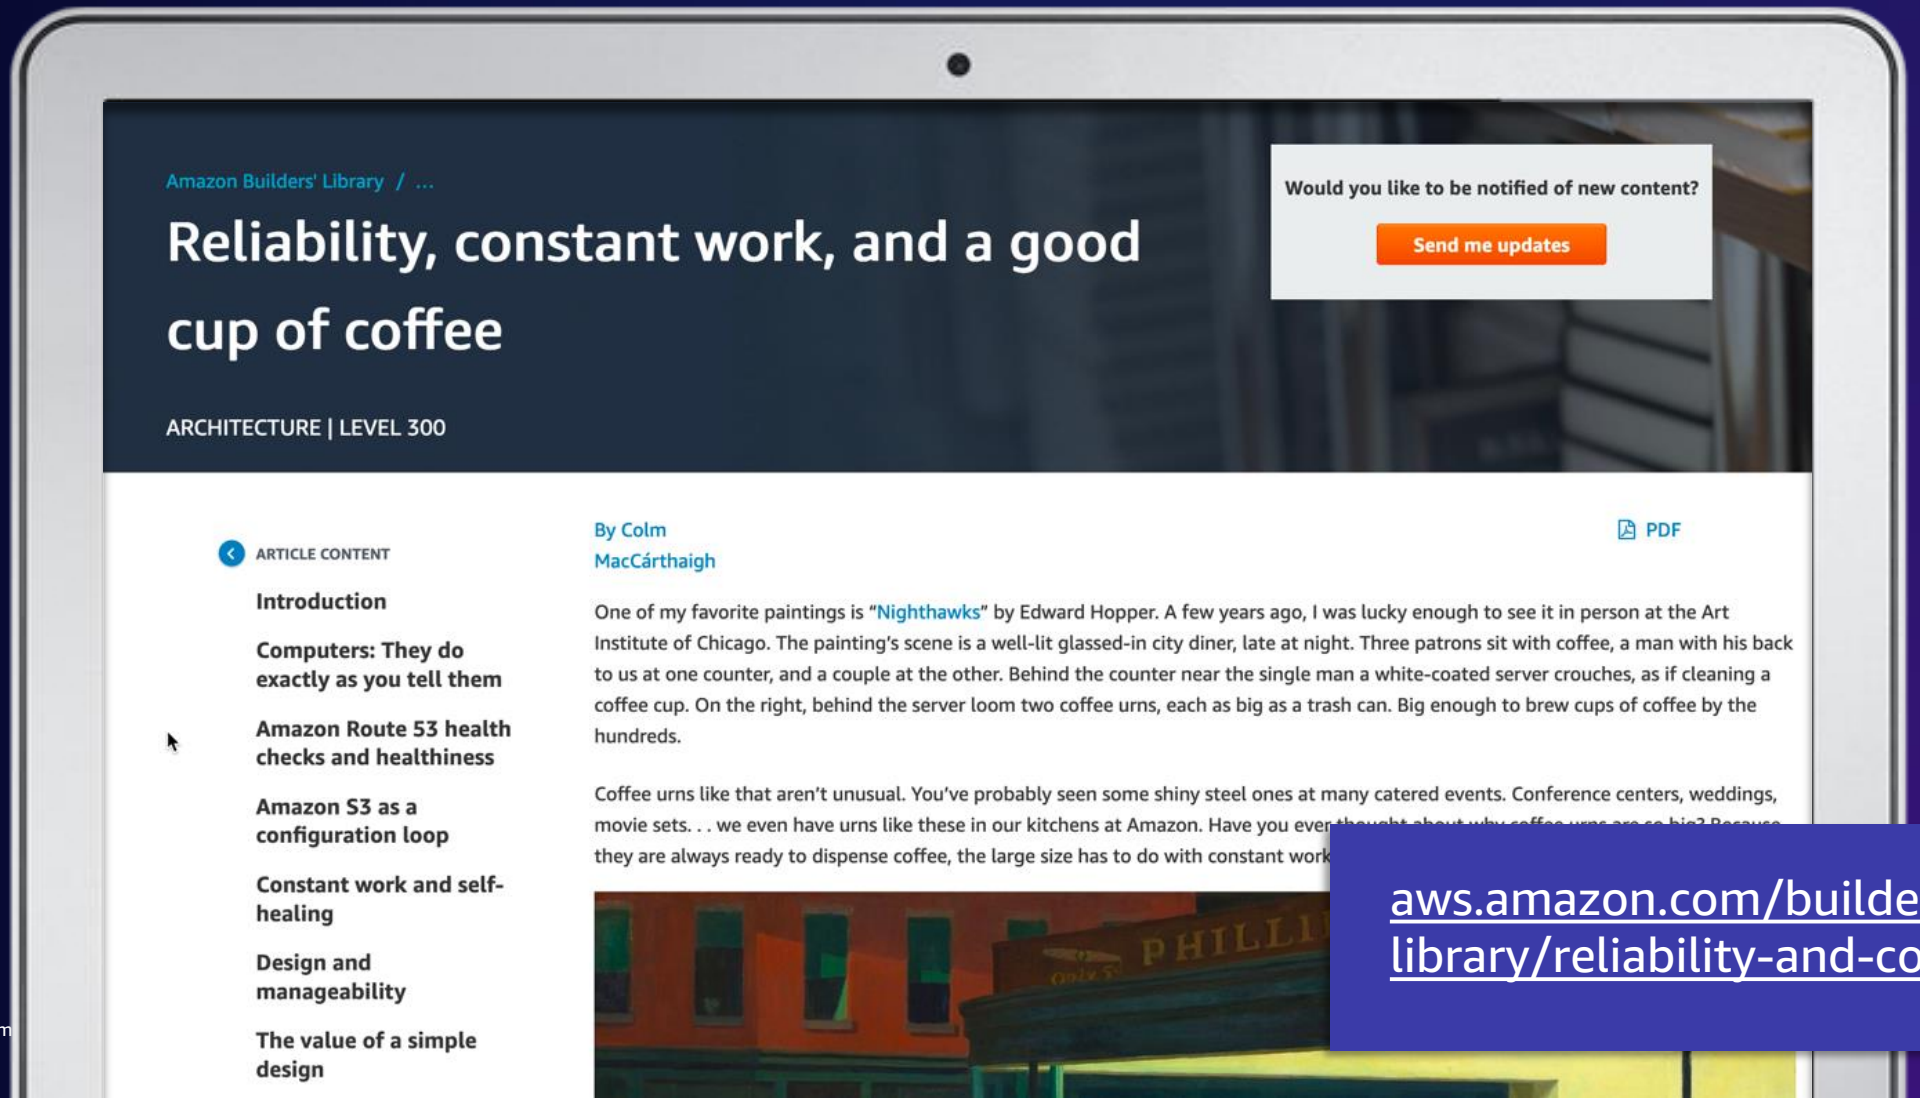

aws.amazon.com/builders-library/reliability-and-constant-work

Limiting impact of failures with cells

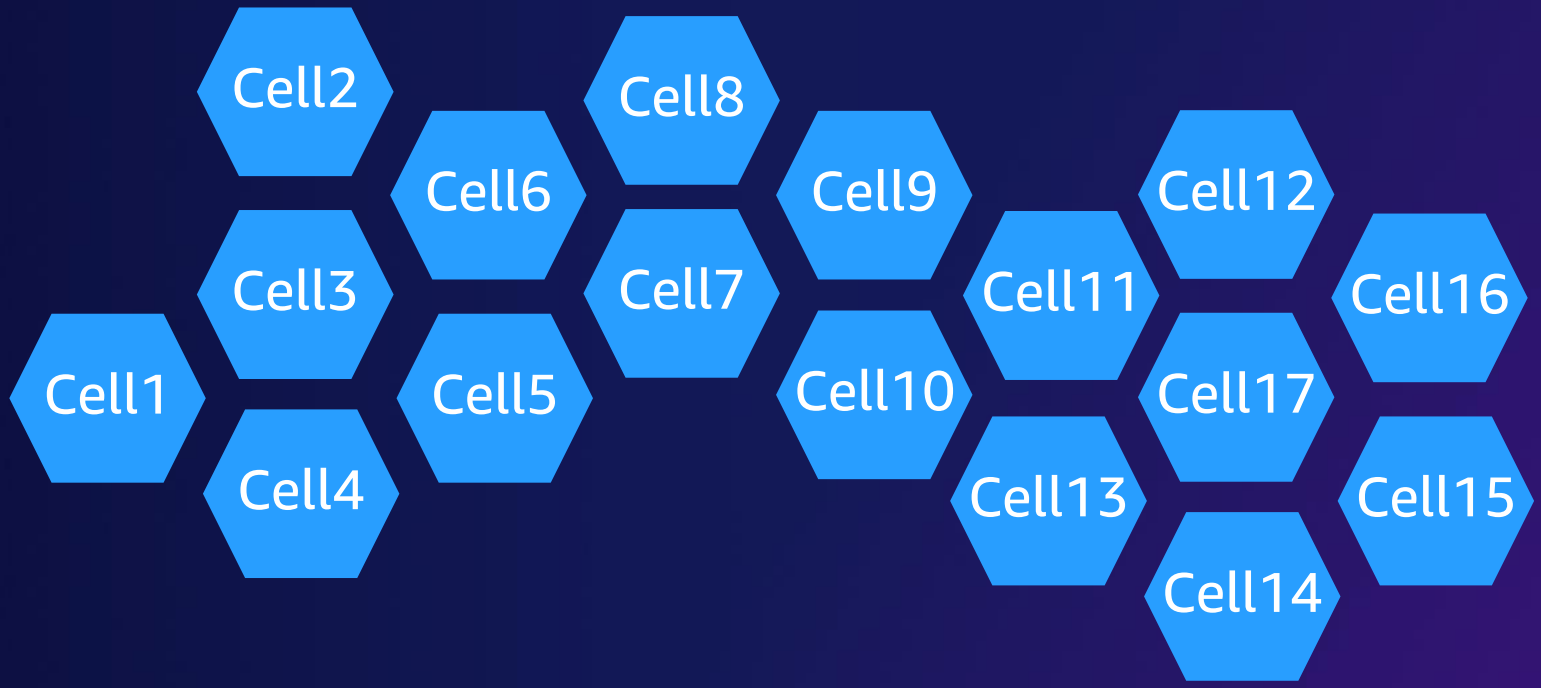

Typical service application

MULTI-AVAILABILITY ZONE

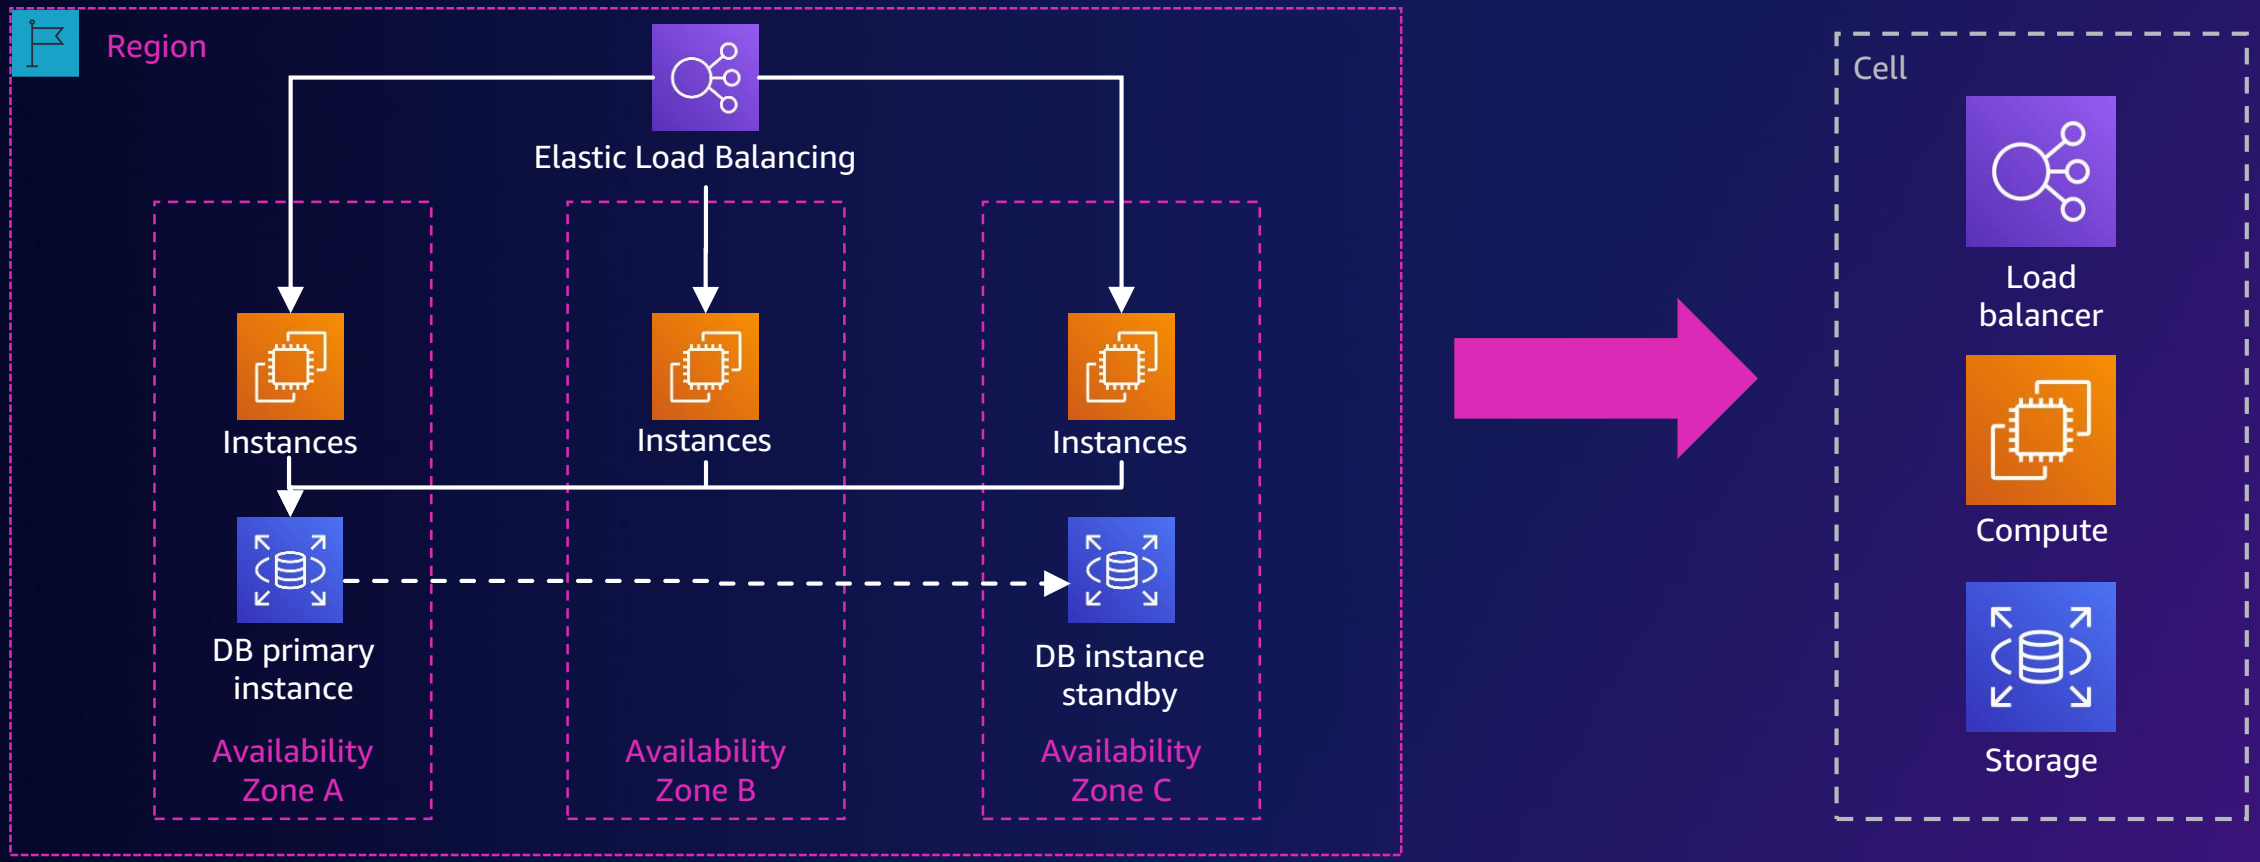

Cell-based architecture

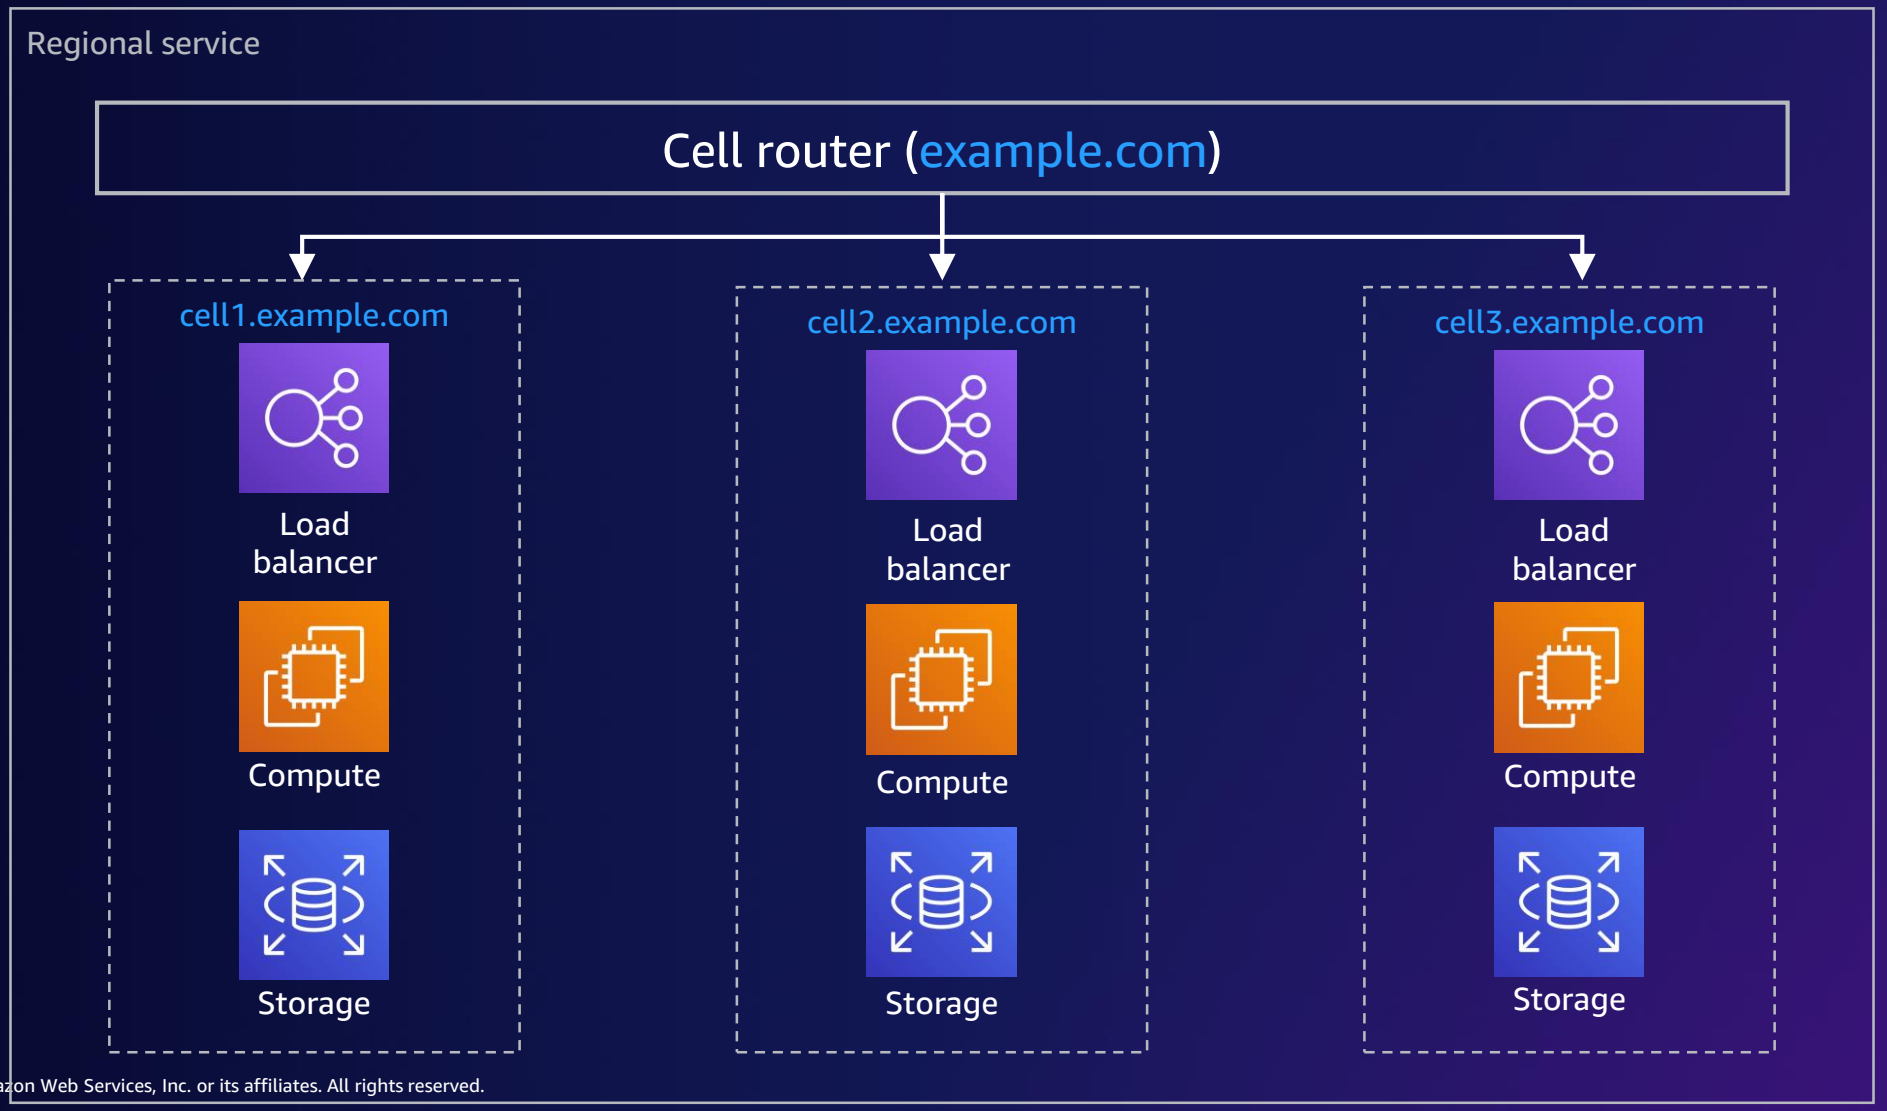

System properties

- Workload isolation
- Failure containment
- Scale out vs. scale up
- Testability
- Manageability

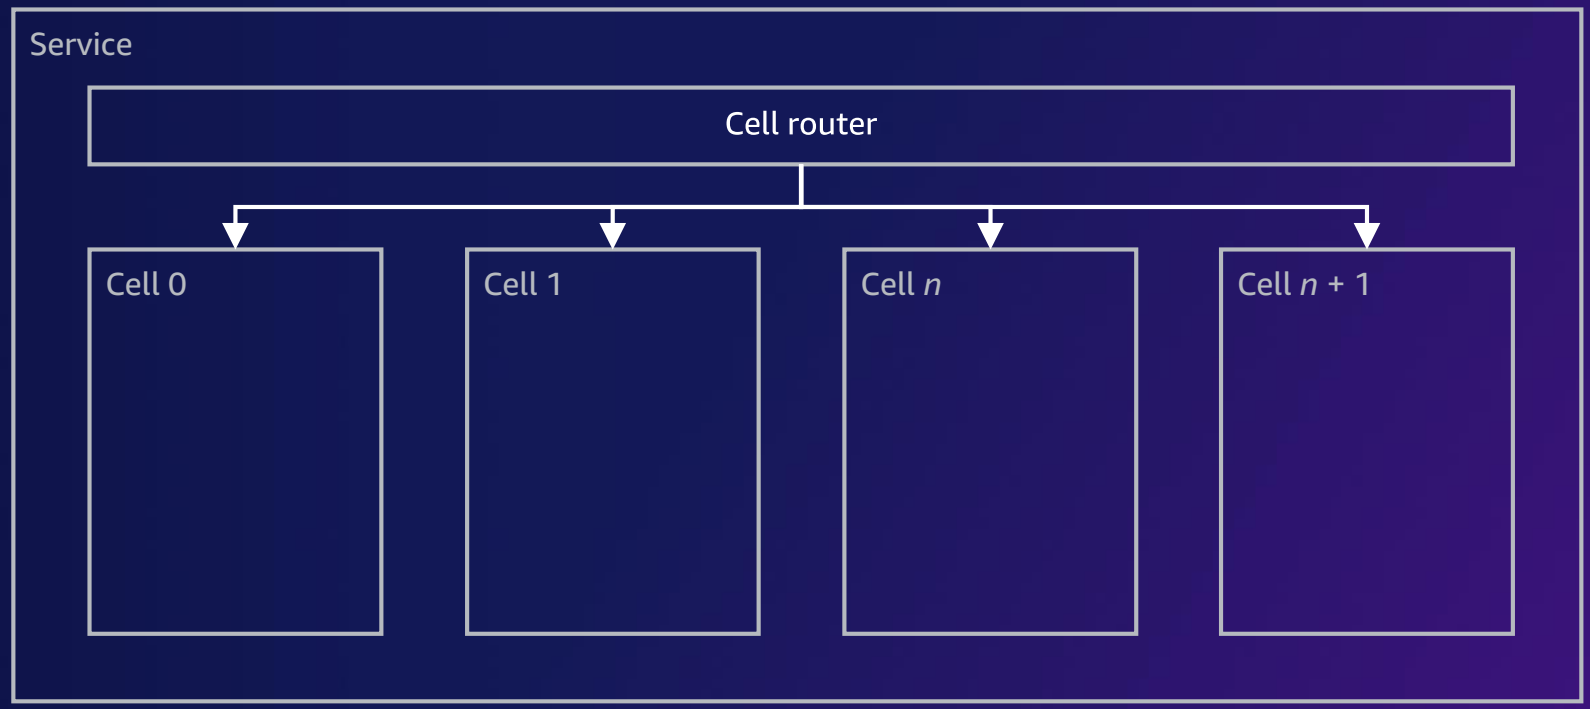

In-place server update

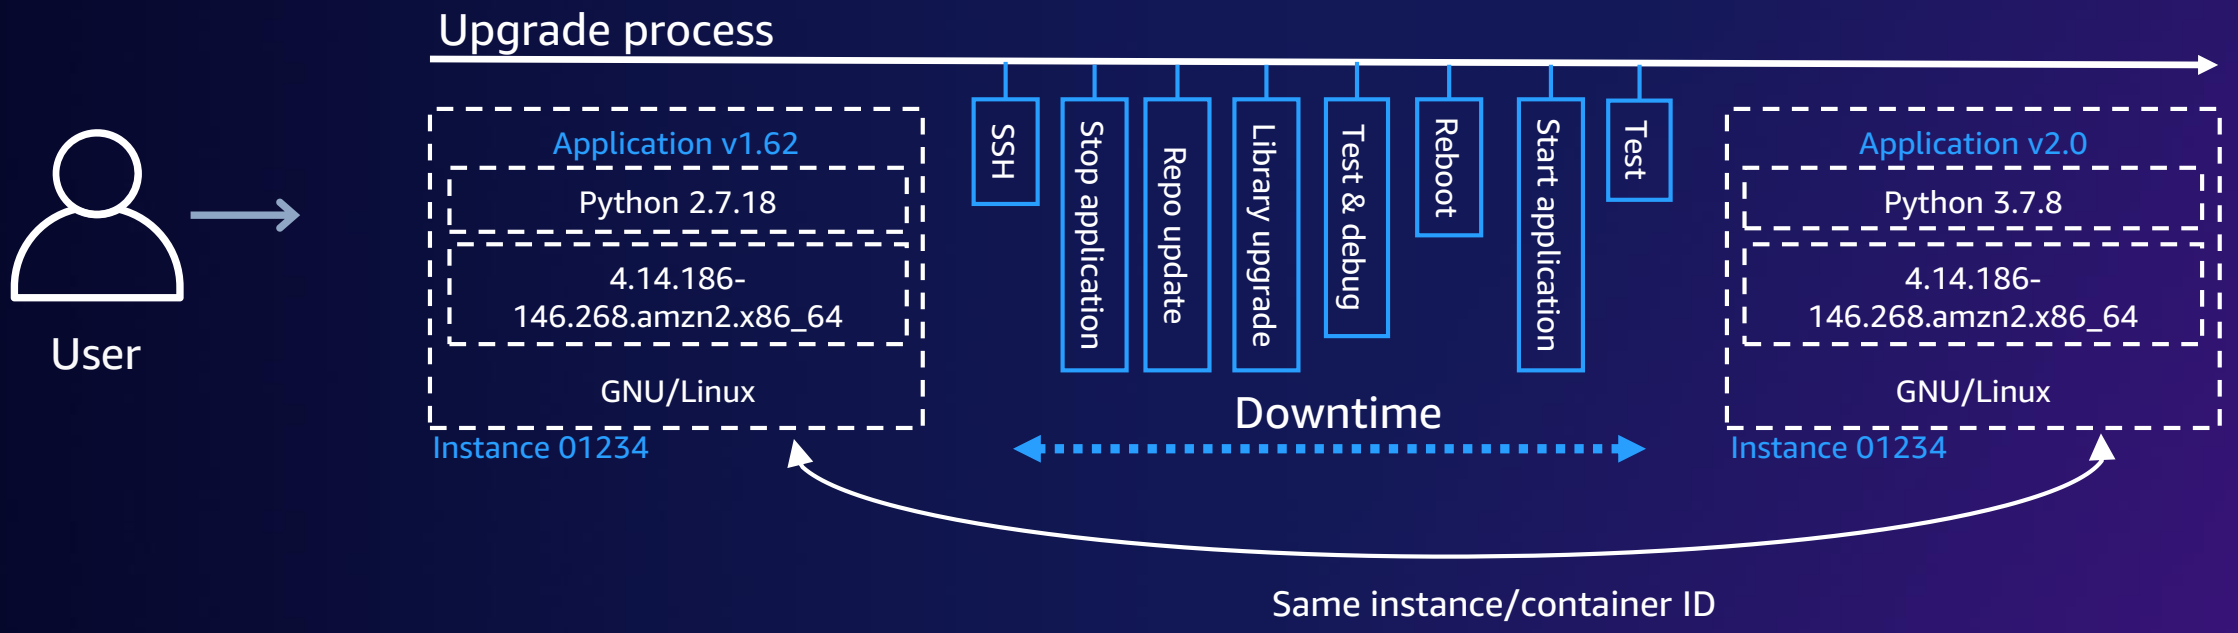

Immutable deployments

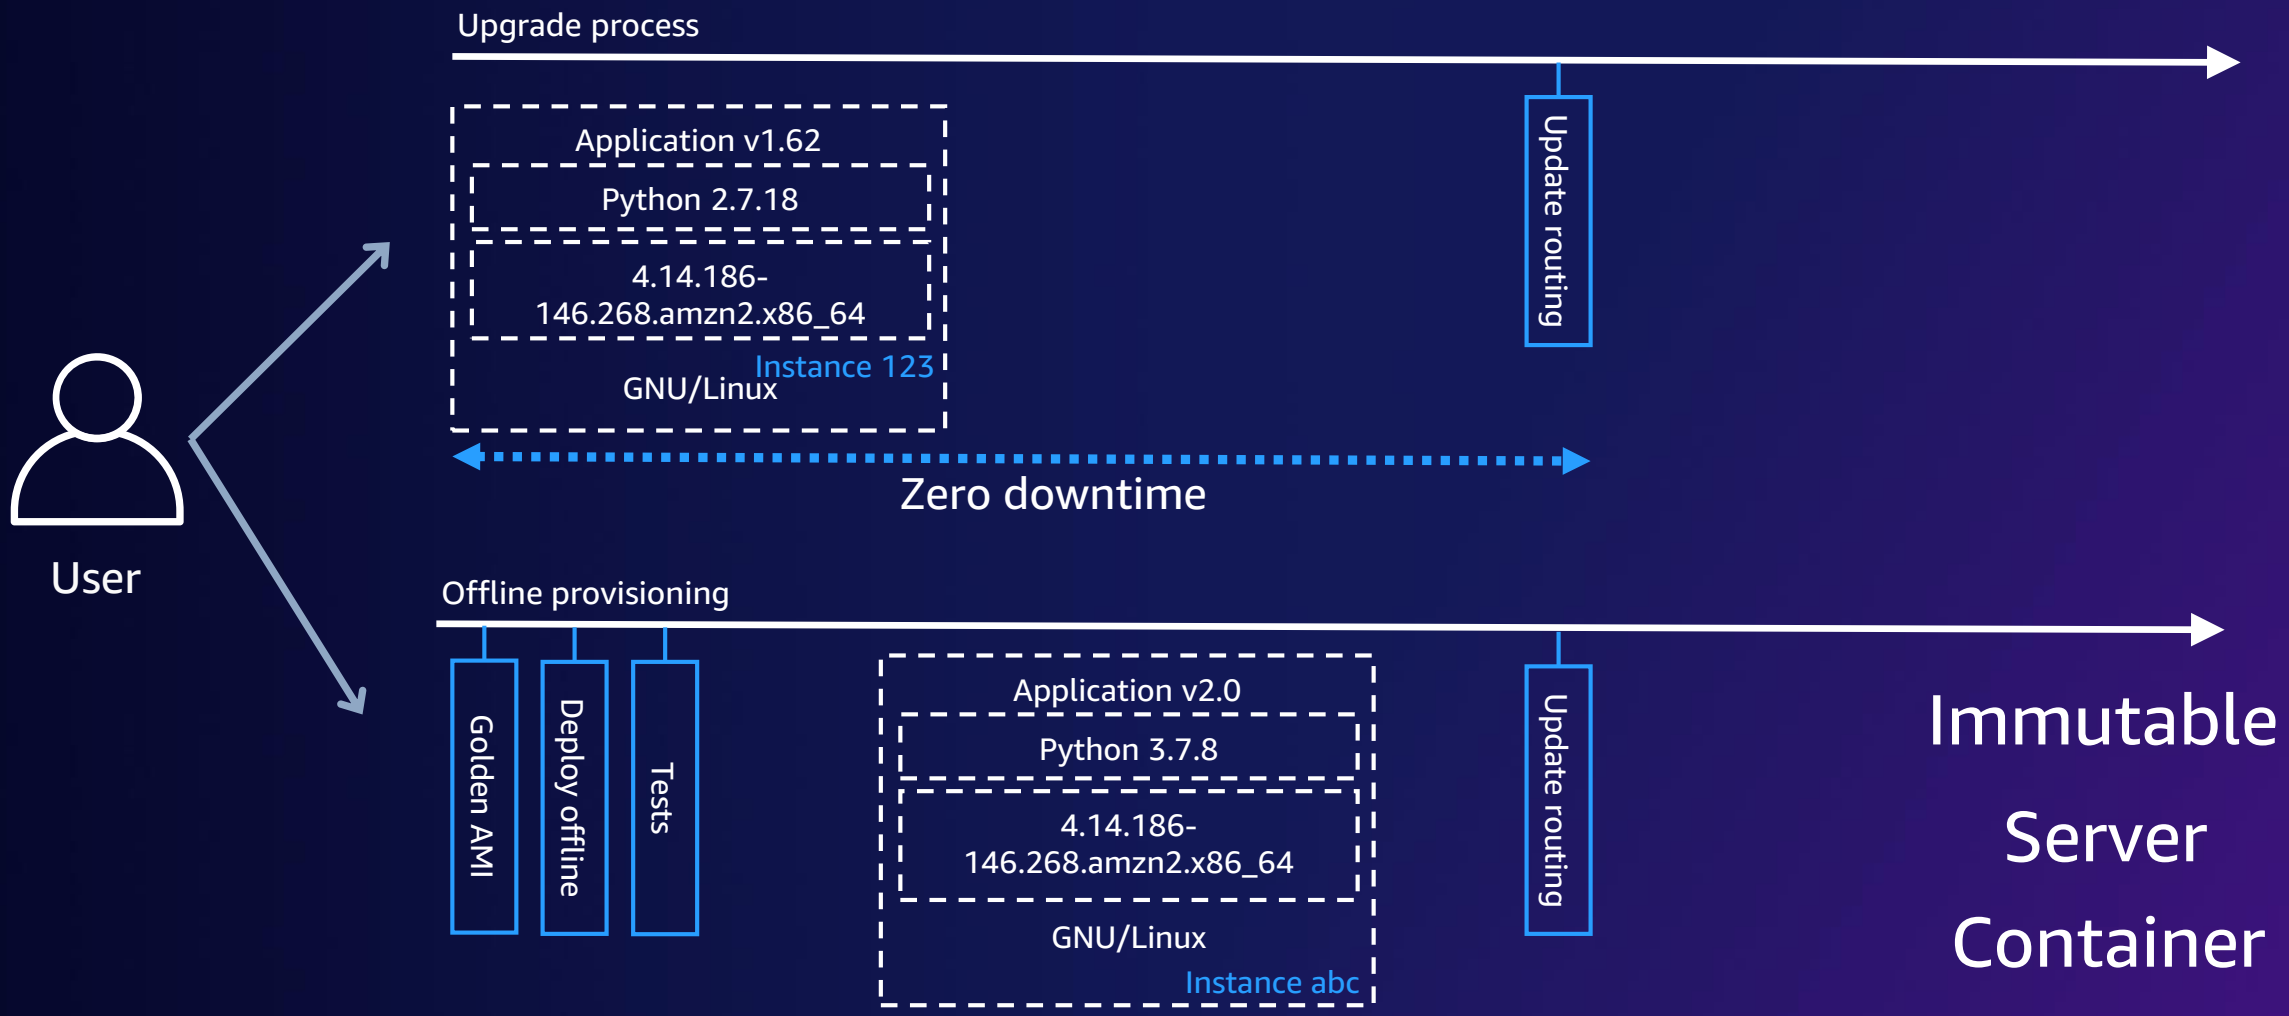

Immutable

Server

Container

Function

Safe deployment with canaries

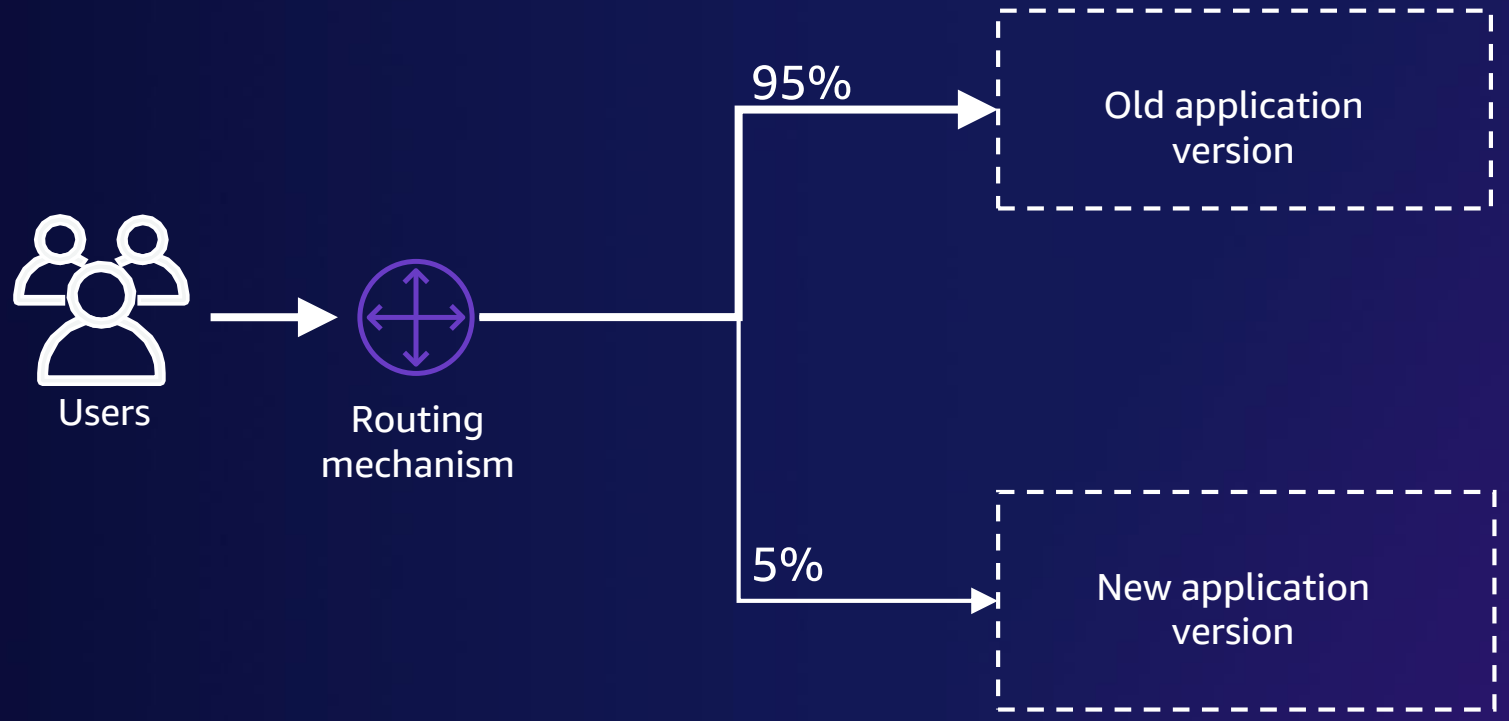

Safe deployment with canaries

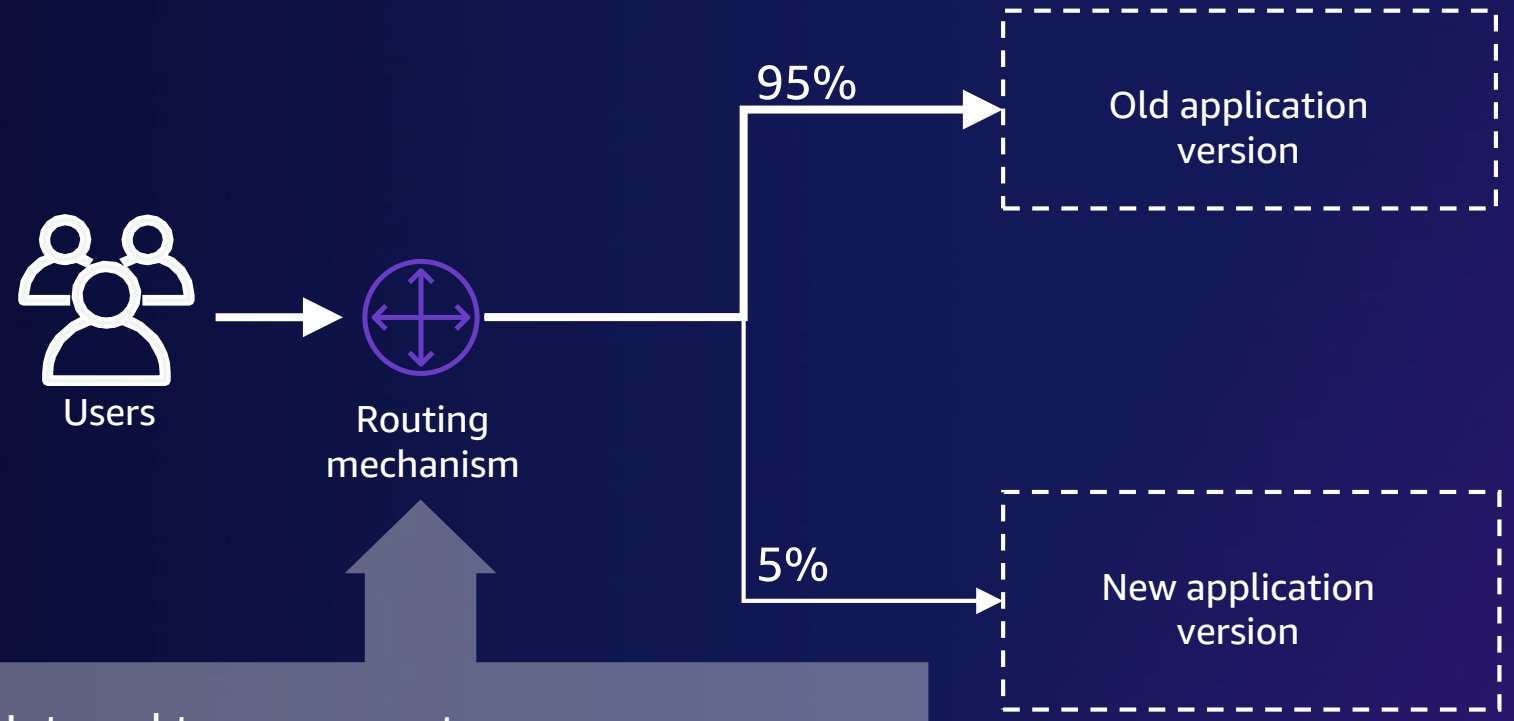

- Internal teams vs. customers
- Paying customers vs. non-paying customers
- Geographic-based routing
- Feature flags (FeatureToggle)
- Random

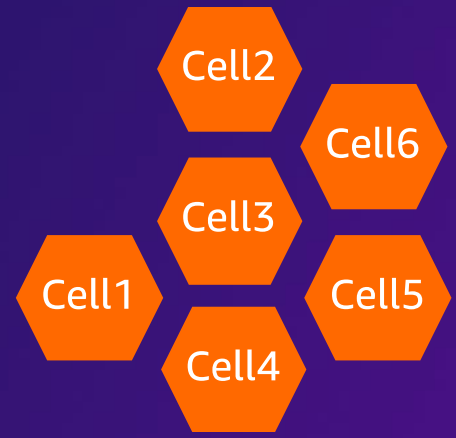

Canary deployment on AWS

Amazon Route 53

Weighted routing policy

Auto Scaling groups

Rolling deployments

Application Load Balancer

Weighted target groups

AWS CodeDeploy

Amazon API Gateway

Release deployments

AWS Lambda

Alias traffic shifting

Amazon ECS

Load balancer target group

Resilience:
Monitoring

Continuous resilience

MONITORING

- Health checks
- Tracing
- Alarming
- Dashboarding

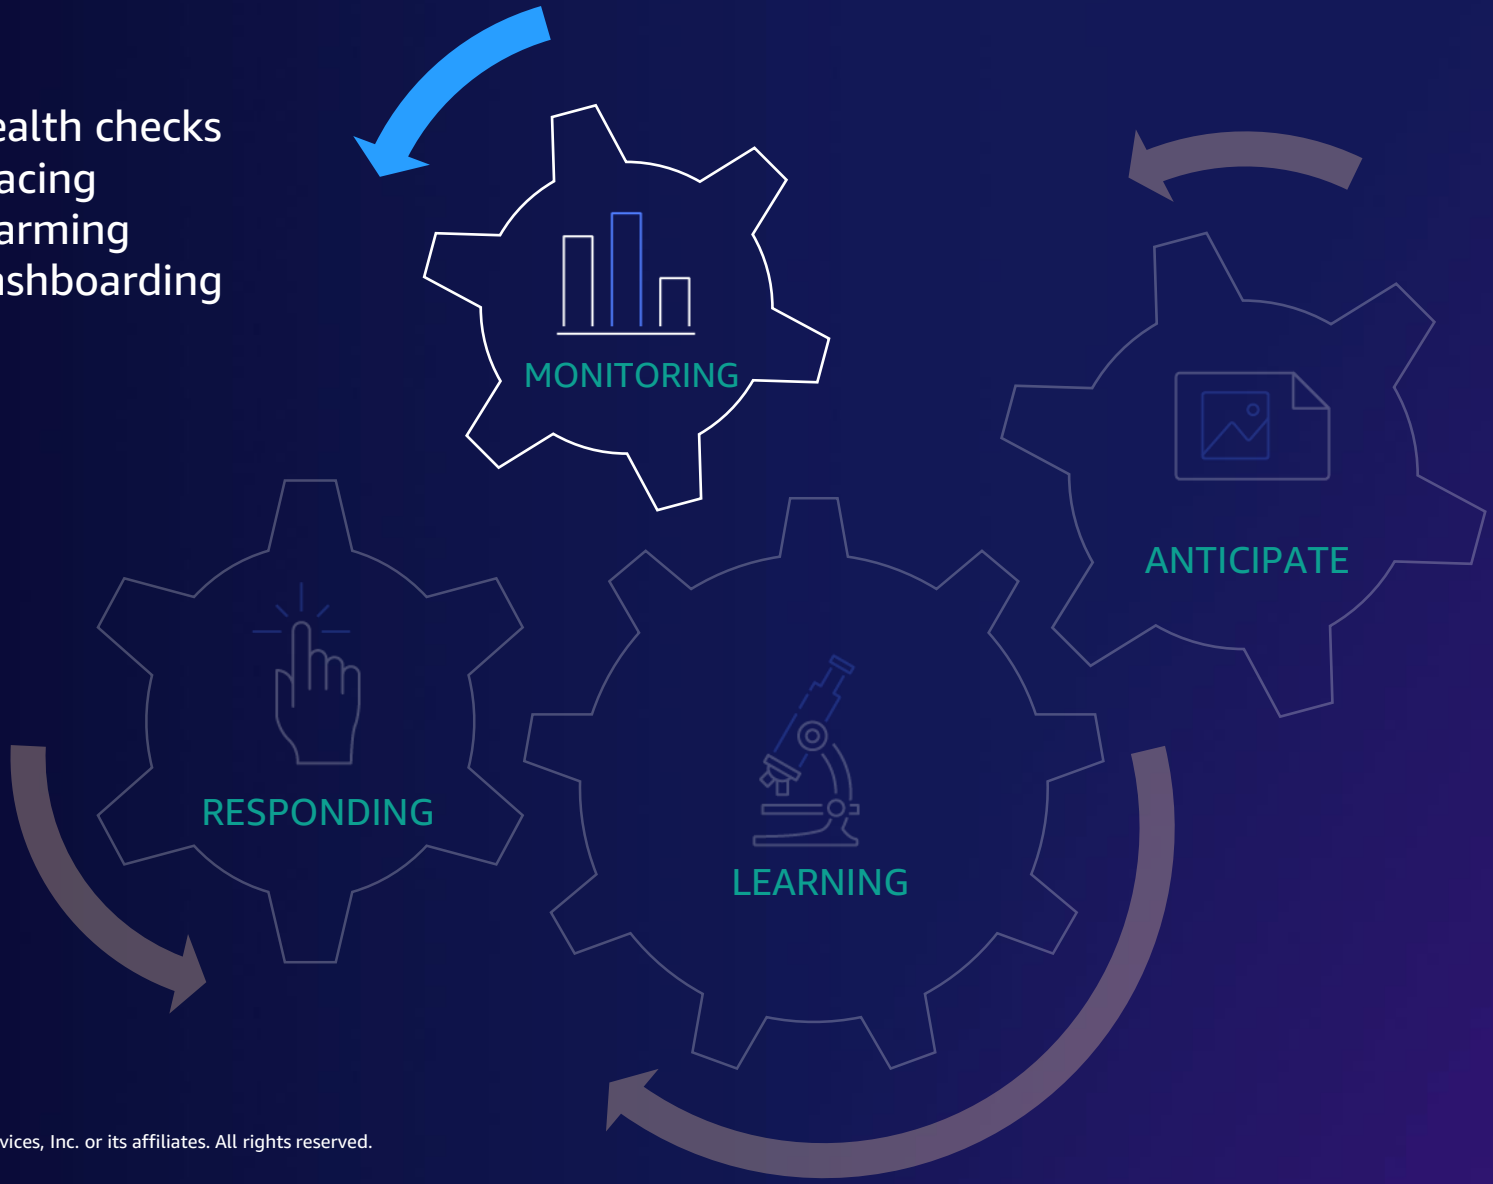

Monitoring more than failures

WITH OBSERVABILITY

High-level overviews

Highly granular insights into the implicit failure modes

Context about its inner workings

"Distributed Systems Observability," by Cindy Sridharan

Monitor the health of your applications

THREE PILLARS OF OBSERVABILITY TOOLING

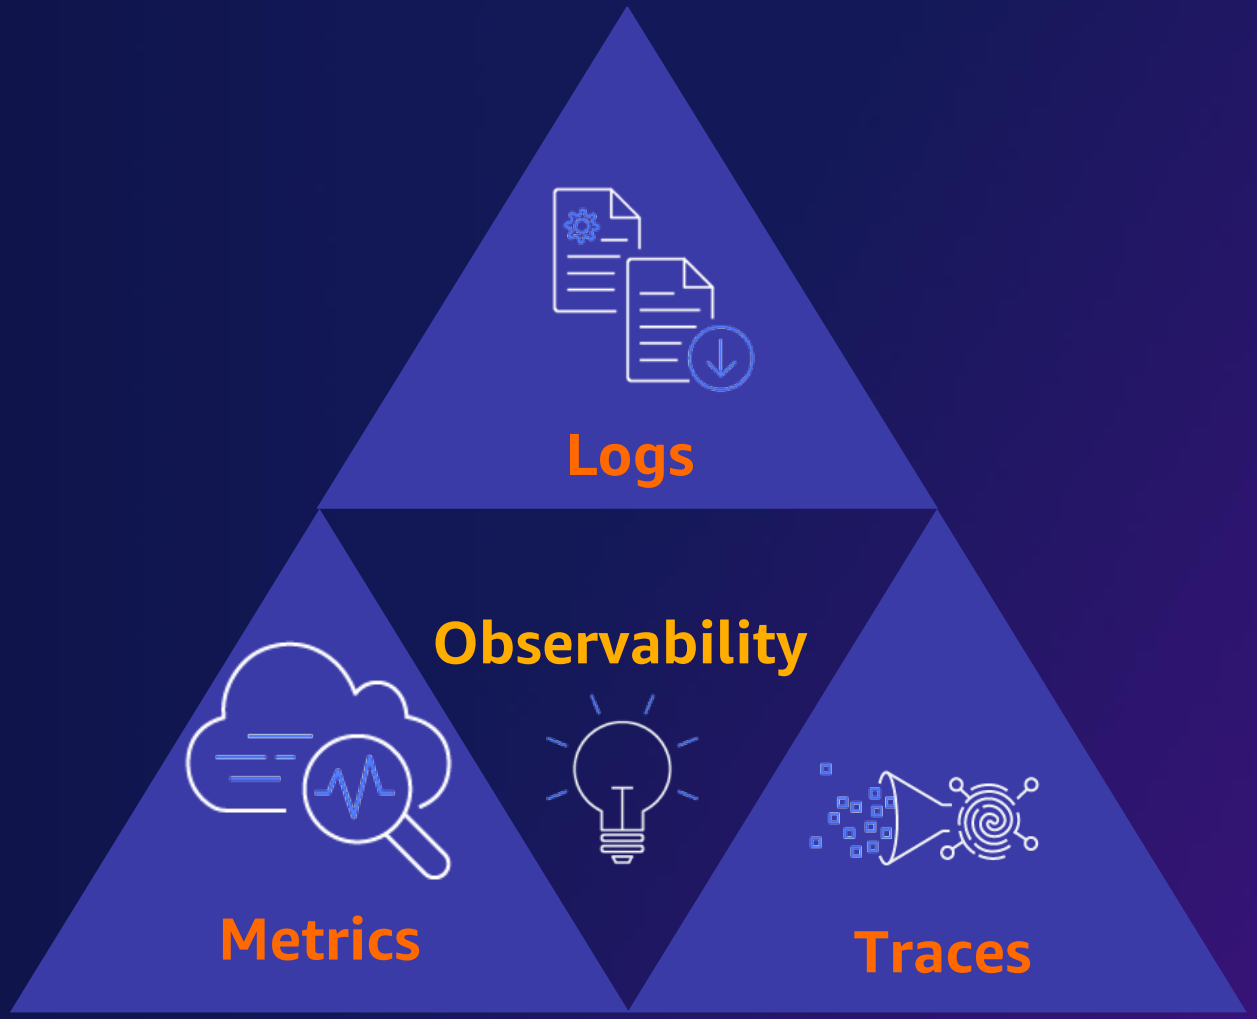

AWS observability tools

- Infrastructure monitoring
- Application monitoring
- Synthetic monitoring

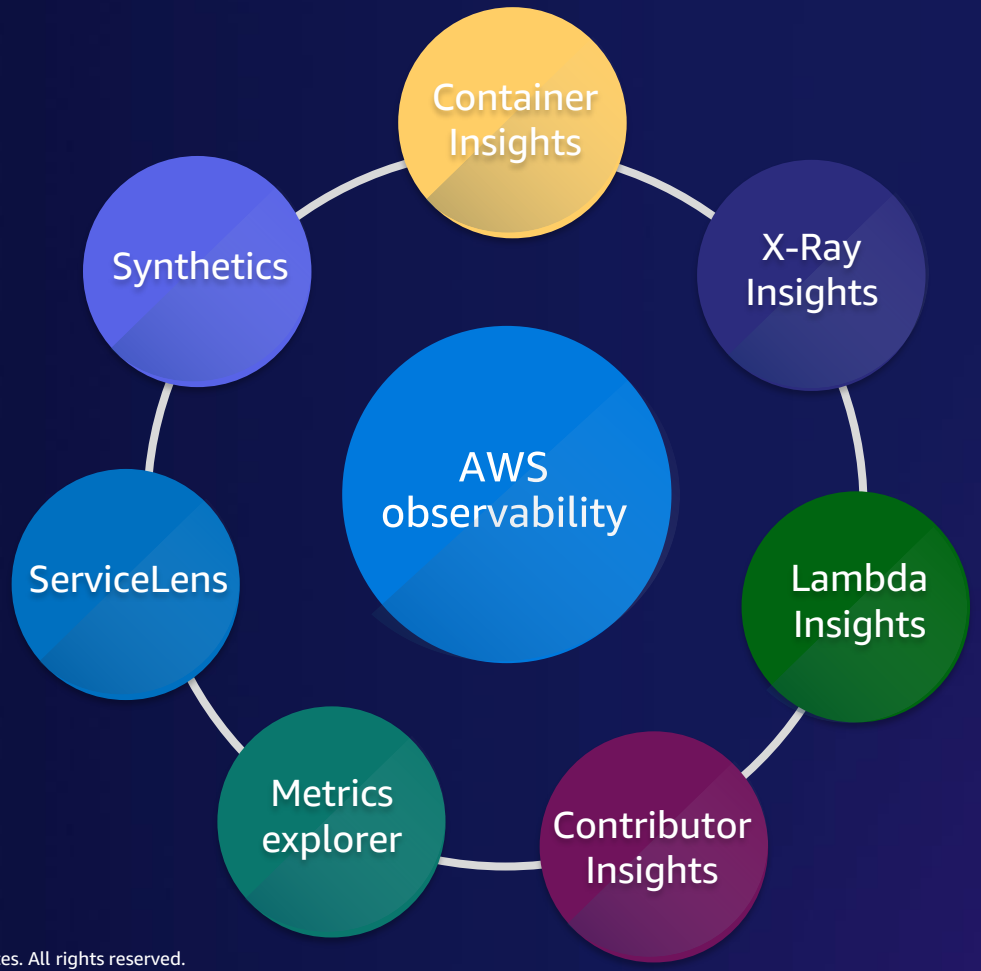

AWS Distro for OpenTelemetry

CloudWatch ServiceLens

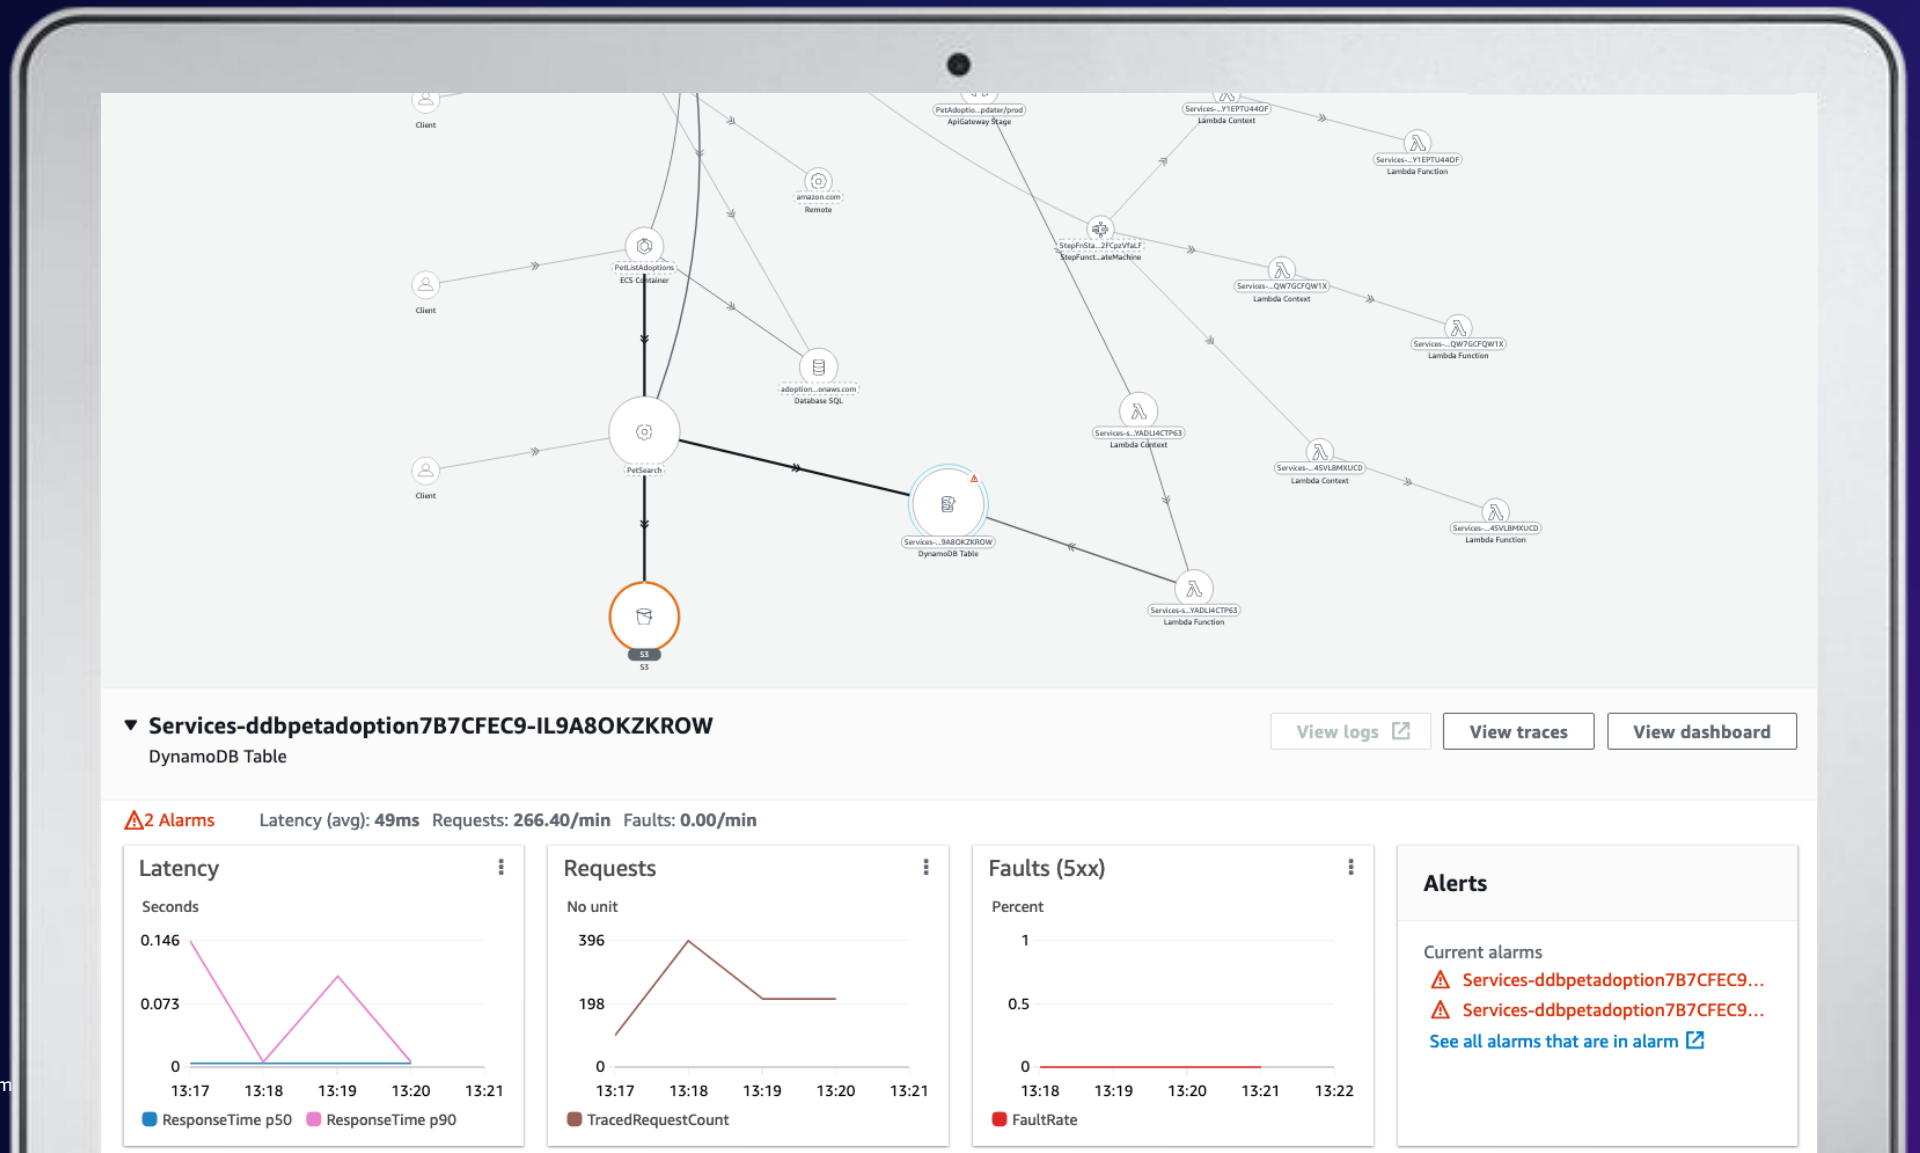

Resilience:
Responding

Continuous resilience

RESPONDING

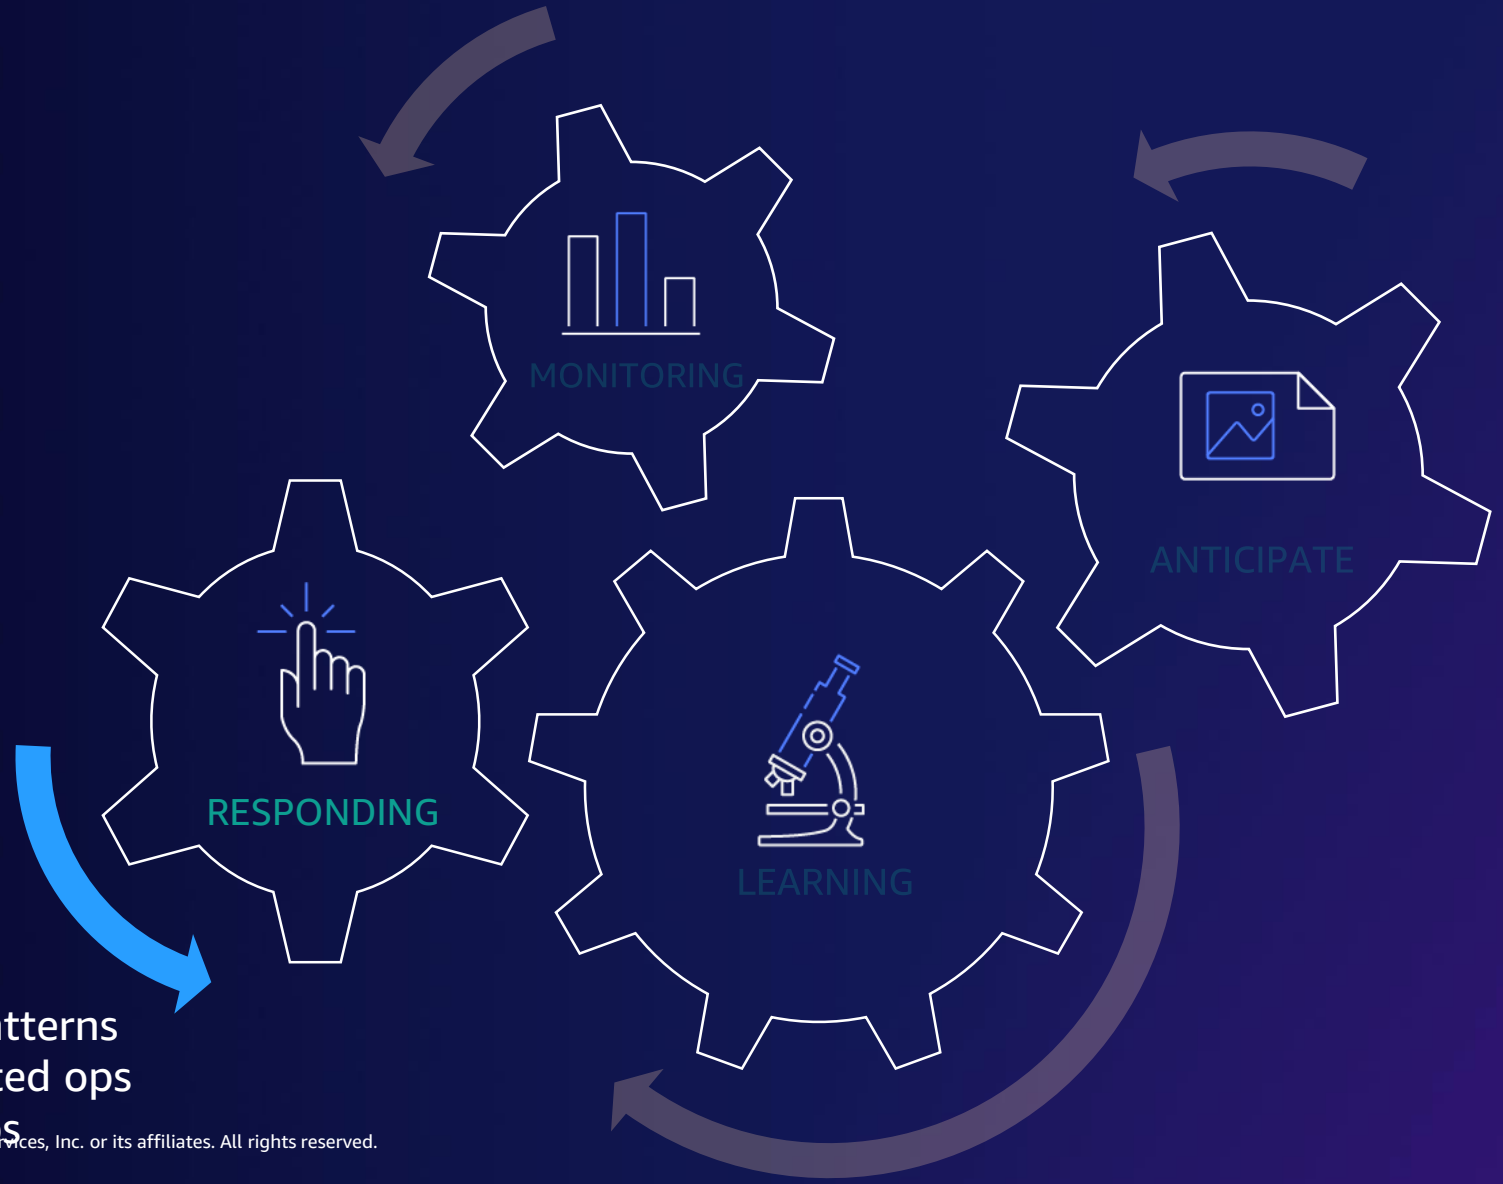

- Event-driven patterns
- Recovery-oriented ops
- ML-powered ops

“Being able to recover quickly from failure is more important than having failures less often.”

John Allspaw
Co-Founder of Adaptive Capacity Labs

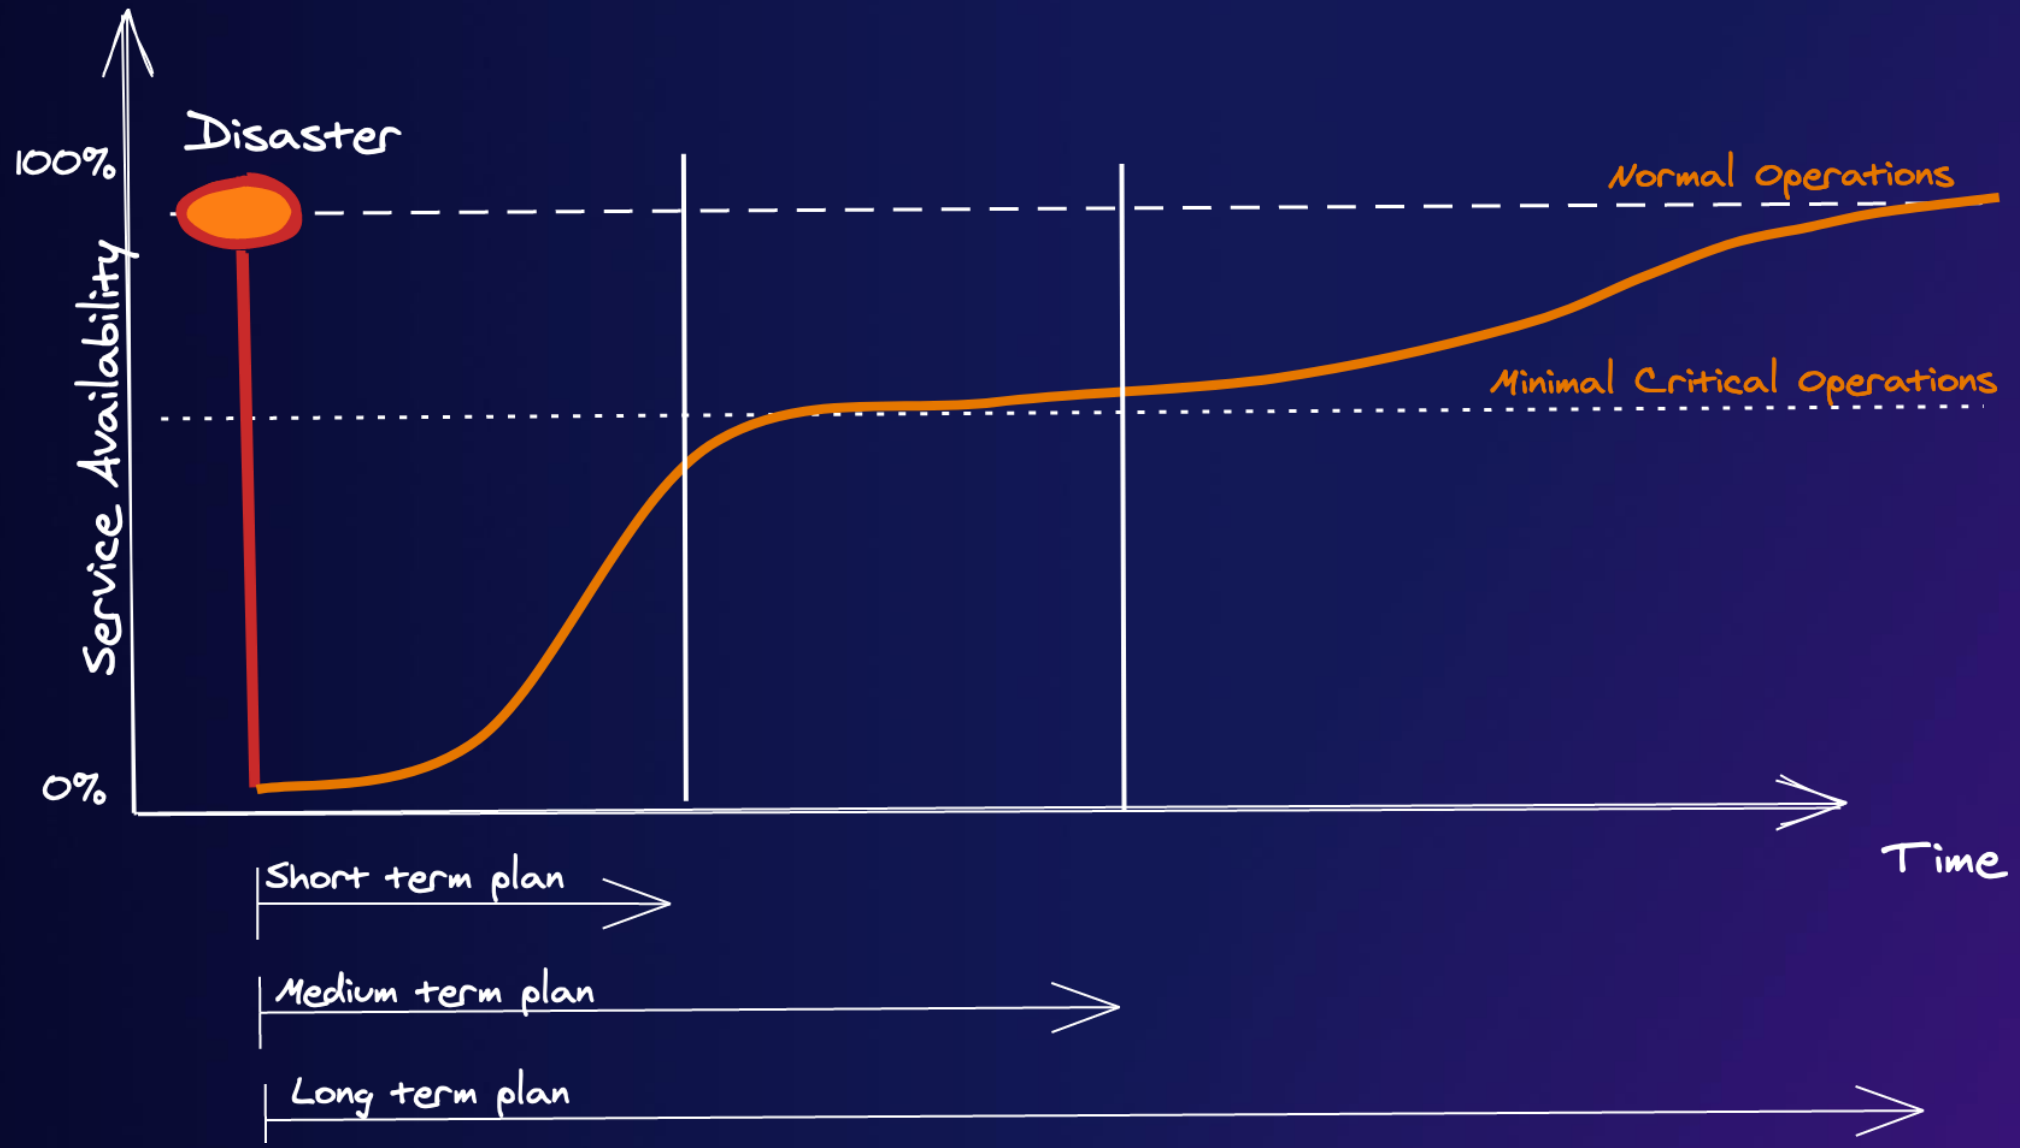

Event-driven patterns

AUTOMATED RESPONSES THAT RUN IN RESPONSE TO EVENTS

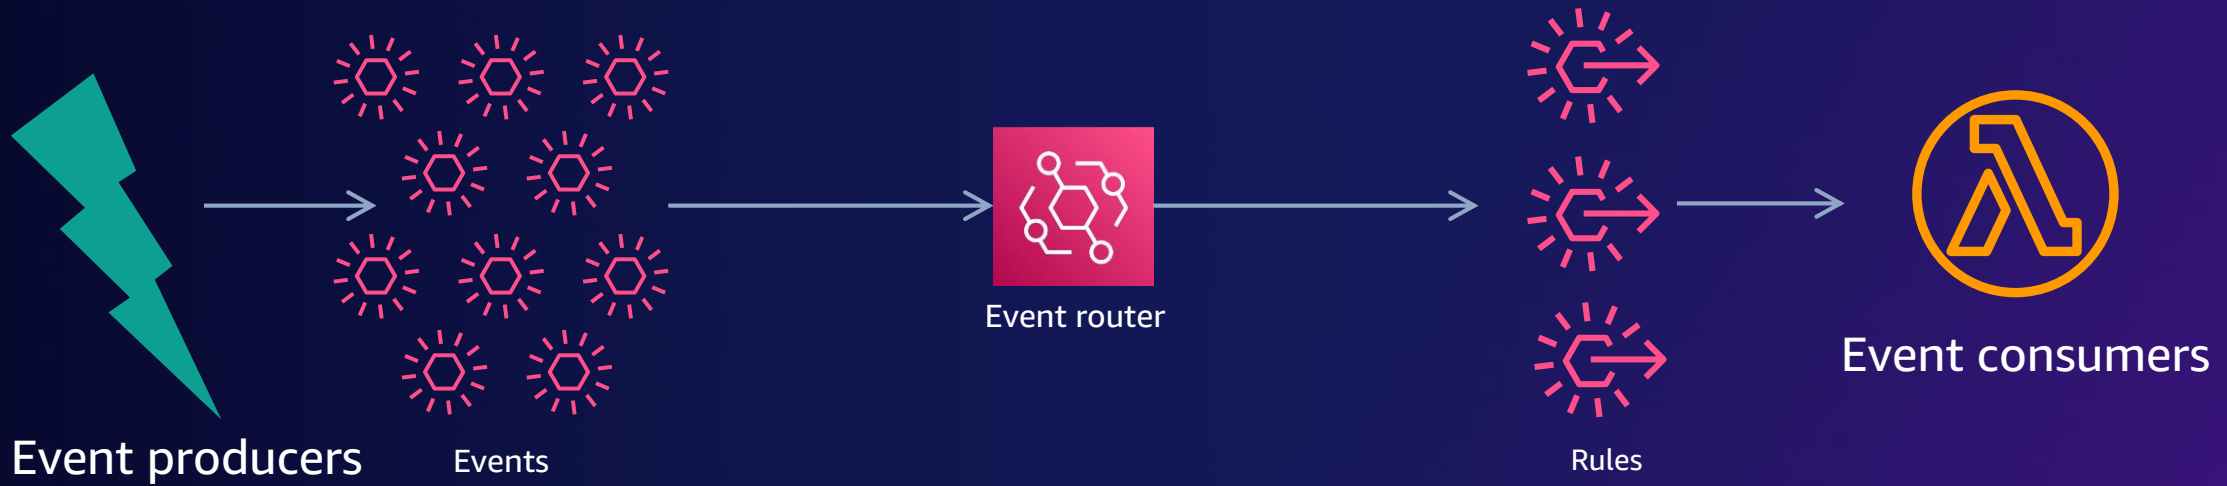

Event-driven patterns

AUTOMATED RESPONSES THAT RUN IN RESPONSE TO EVENTS

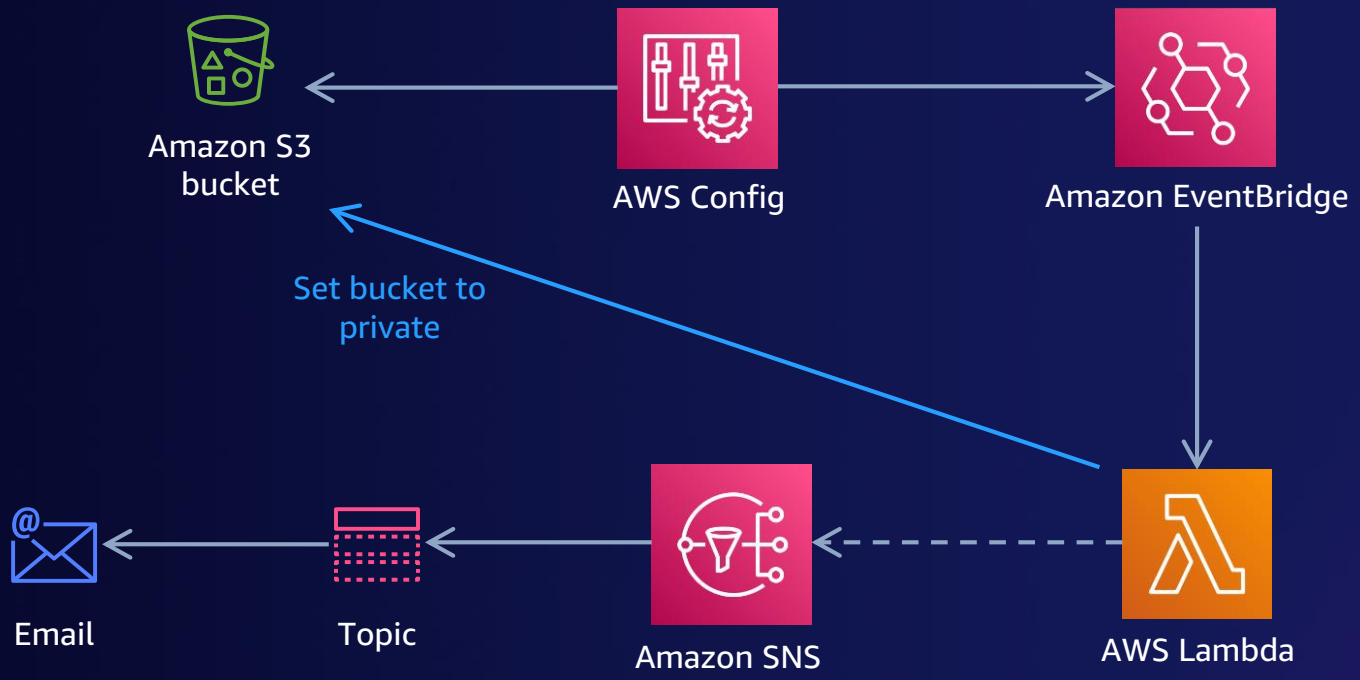

github.com/awslabs/aws-config-rules/tree/master/aws-config-conformance-packs

Security considerations

AUTOMATED RESPONSES THAT RUN IN RESPONSE TO EVENTS

Mutability is one of the most critical attack vectors for cyber crimes

ML-powered ops

WITH AMAZON DEVOPS GURU

Informed by years of Amazon.com and AWS operational excellence

Resilience:
Learning

Continuous resilience

LEARNING

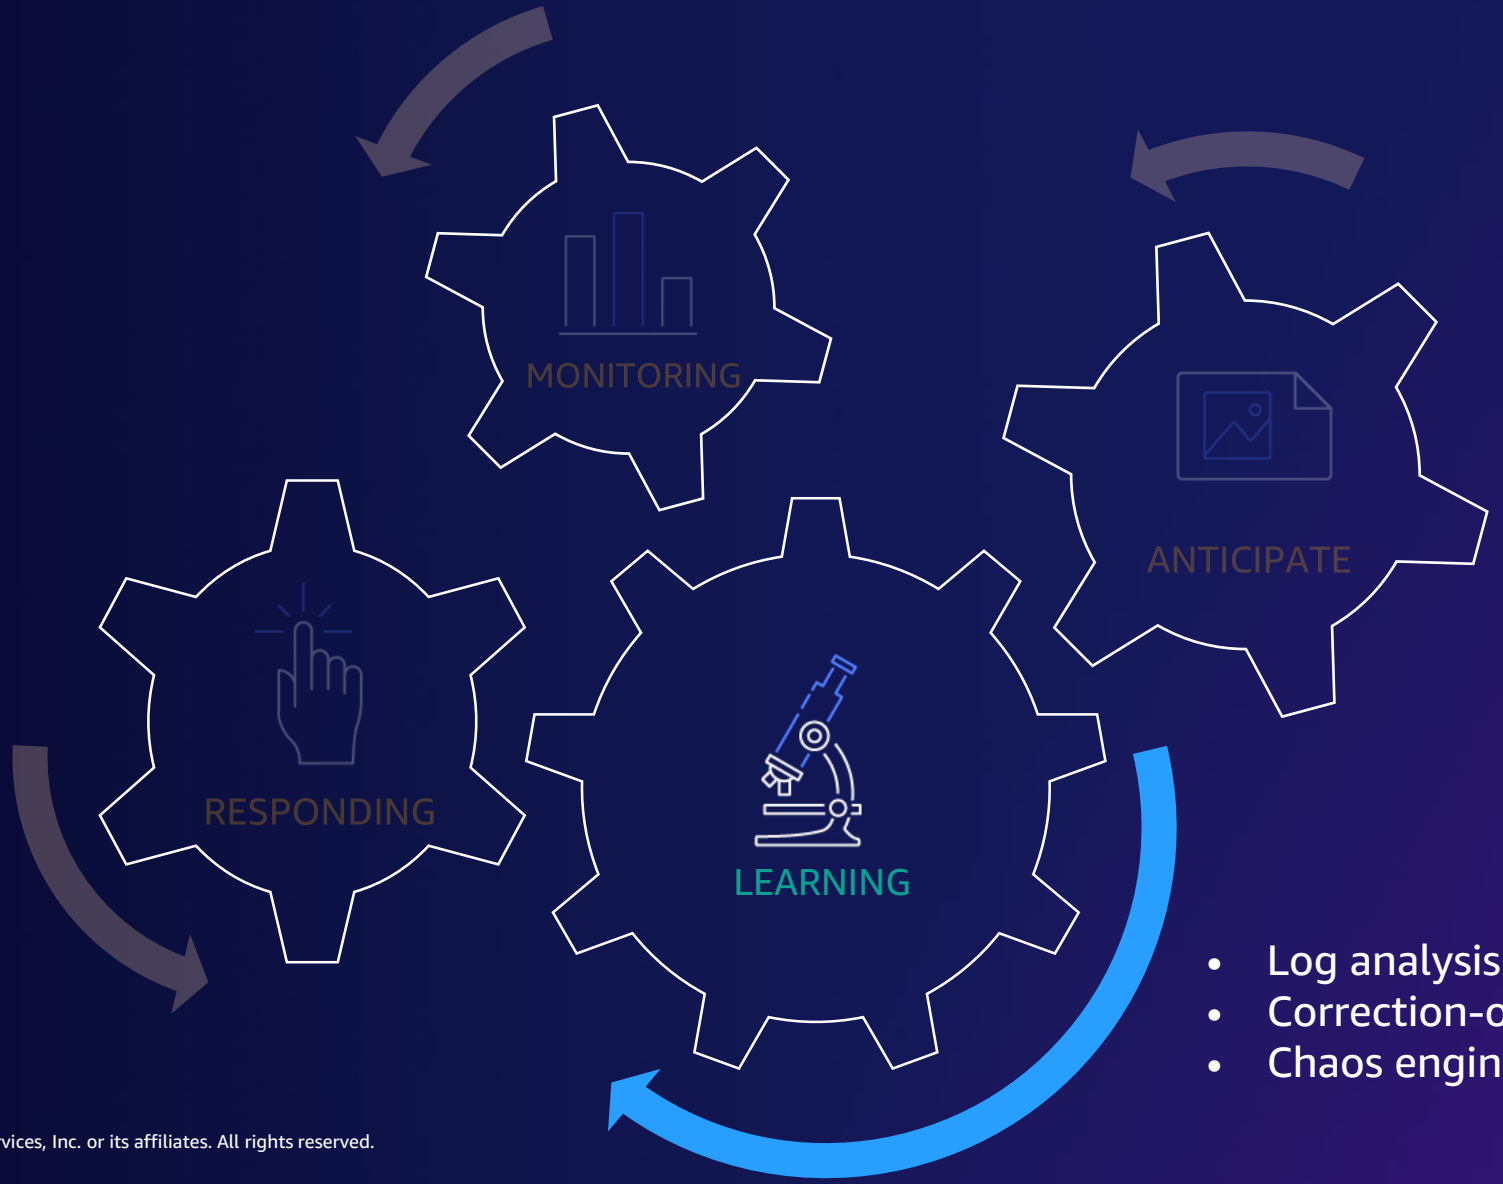

- Log analysis
- Correction-of-errors (COE)
- Chaos engineering

Correction of errors (COE)

POSTMORTEMS

Mechanism to learn from mistakes

- ✓ Technical flaws
- ✓ Process flaws
- ✓ Documentation flaws
- ✓ Organizational flaws
- ✓ Other flaws

Mechanism to identify contributing factors to failures

Mechanism to drive CONTINUOUS IMPROVEMENT

Anatomy of a COE

- What happened?
- What data do you have to support this? – metrics and graphs
- What was the impact on customers and your business?
- What are the contributing factors? – don’t stop at operators
- What lessons did you learn?
  - What corrective actions are you taking?
  - Actions items
  - Related items (trouble tickets, etc.)

https://youtu.be/yQiRli2ZPxU

Chaos engineering

WITH AWS FAULT INJECTION SIMULATOR

- Improve resilience and performance
- Uncover hidden issues
- Expose blind spots
  Monitoring, observability, and alarm
- And more

AWS Fault Injection Simulator

AWS News Blog

AWS Fault Injection Simulator – Use Controlled Experiments to Boost Resilience

by Jeff Barr | on 15 MAR 2021 | in AWS Re:Invent, Launch, News | Permalink | Comments | Share

0:00 / 0:00

Voiced by Amazon Polly

AWS gives you the components that you need to build systems that are highly reliable: multiple Regions (each with multiple Availability Zones), Amazon CloudWatch (metrics, monitoring, and alarms), Auto Scaling, Load Balancing, several forms of cross-region replication, and lots more. When you put them together in line with the guidance provided in the Well-Architected Framework, your systems *should* be able to keep going even if individual components fail.

However, you won't know that this is indeed the case until you perform the right kinds of tests. The relatively new field of Chaos Engineering (based on pioneering work done by "Master of Disaster" Jesse R
Amazon.com, and then taken into high gear by the Netflix Chaos Monkey) focuses
creating disruptive events, observing how the system responds, and implementing
out the areas for improvements, Chaos Engineering helps to discover blind spots th
alarming, uncovers once-hidden implementation issues, and gives you an opportun
with an eye toward improving recovery time. To learn a lot more about this topic, start with Chaos Engineering – Part 1

aws.amazon.com/blogs/aws/aws-fault-injection-simulator-use-controlled-experiments-to-boost-resilience/

“Excellence, then, is not an act, but a habit.”

Will Durant

Writer, historian, philosopher, and teacher

Additional resources

The Amazon Builder's Library
How Amazon builds and operates software
https://amzn.to/2OjslEV

The Resilient Architecture Collection
A list of resiliency-related blog posts
https://bit.ly/30ERxs0

The Chaos Engineering Collection
A list of chaos engineering-related blog posts
https://bit.ly/3rK9dOK

Immutable Infrastructure
Reliability, consistency, and confidence through immutability
https://bit.ly/3lcQxor

Towards Operational Excellence
On culture, tools, and processes
https://bit.ly/3bG5a08

The Cloud Architect
Build resilient, scalable, and highly available cloud architectures
https://bit.ly/3ezpzWI

Four Concepts for Resilience and the Implications for the Future of Resilience Engineering
By David D. Woods, The Ohio State University
https://bit.ly/3qFOgmv

Learn in-demand AWS Cloud skills

AWS Skill Builder

Access **500+ free** digital courses and Learning Plans

Explore resources with a variety of skill levels and **16+** languages to meet your learning needs

Deepen your skills with digital learning on demand

Train now

AWS Certifications

Earn an industry-recognized credential

Receive Foundational, Associate, Professional, and Specialty certifications

Join the **AWS Certified community** and get exclusive benefits

Access **new** exam guides

Thank you!

Veliswa Boya

Twitter: @vel12171

Please complete the session survey