Cloud Day Prague

O2 UNIVERSUM| 23 OCTOBER 2024

APP306

Resilient architectures at scale: Real-life use cases from Amazon

Dragoș Mădărășan
Solutions Architect Team Lead
AWS

Resilient architectures

Resilience

The ability of an application to resist or recover from faults or load spikes, and remain functional.

Amazon got big

AMAZON PRIME DAY 2023 STATISTICS

Prime members globally purchased over
375M items

Global Amazon Prime Day sales
$12.7B

Amazon SQS

86M peak requests per second

Amazon Aurora

318 billion transactions
2,140 terabytes stored

Amazon DynamoDB

126M peak requests per second

Amazon EBS

15.4T requests per day

https://bit.ly/prime-day-2023-stats
https://bit.ly/prime-day-2023-record
https://bit.ly/prime-day-2023-on-aws

aws

Architecture 2024

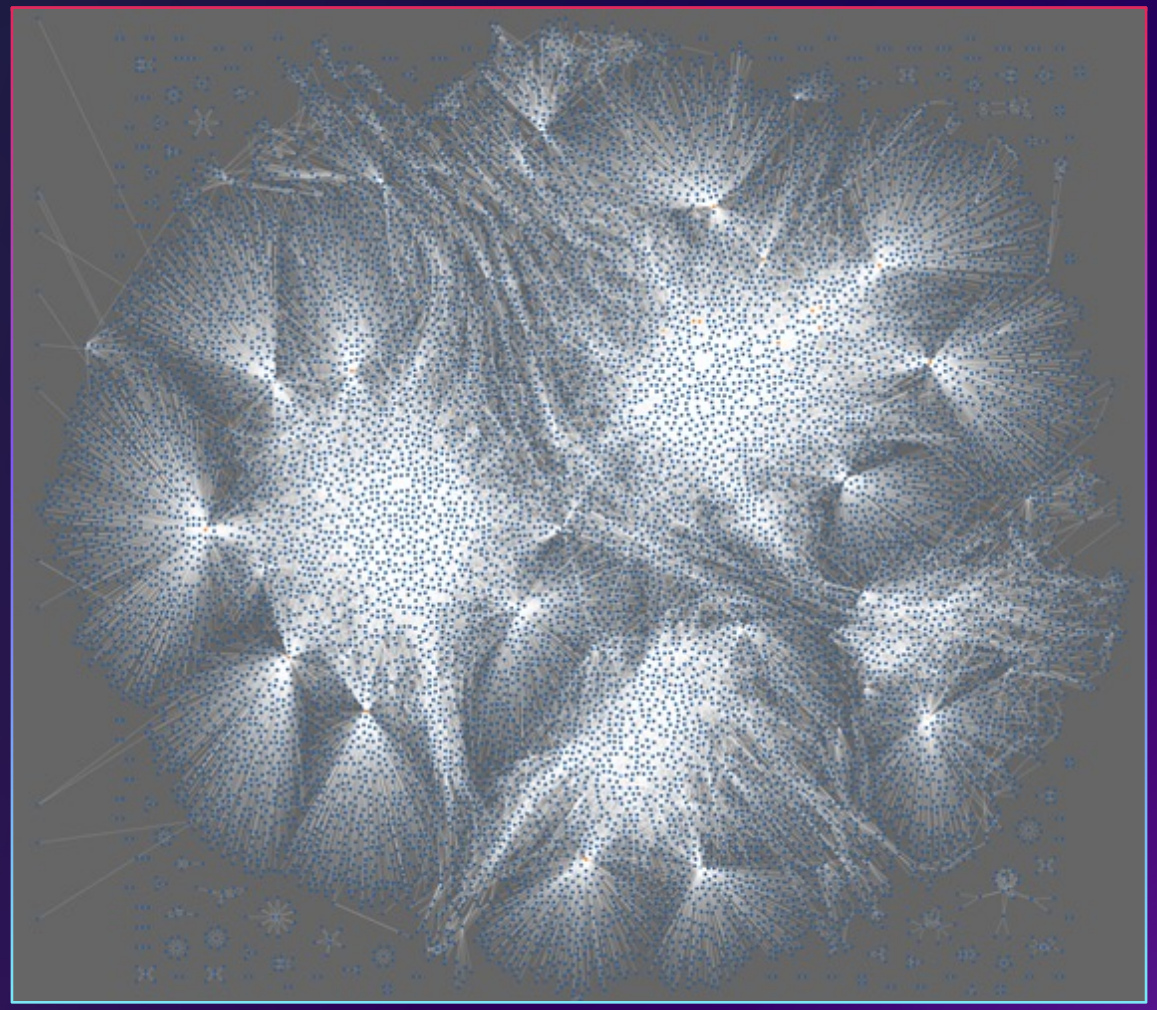

Amazon product “detail page”

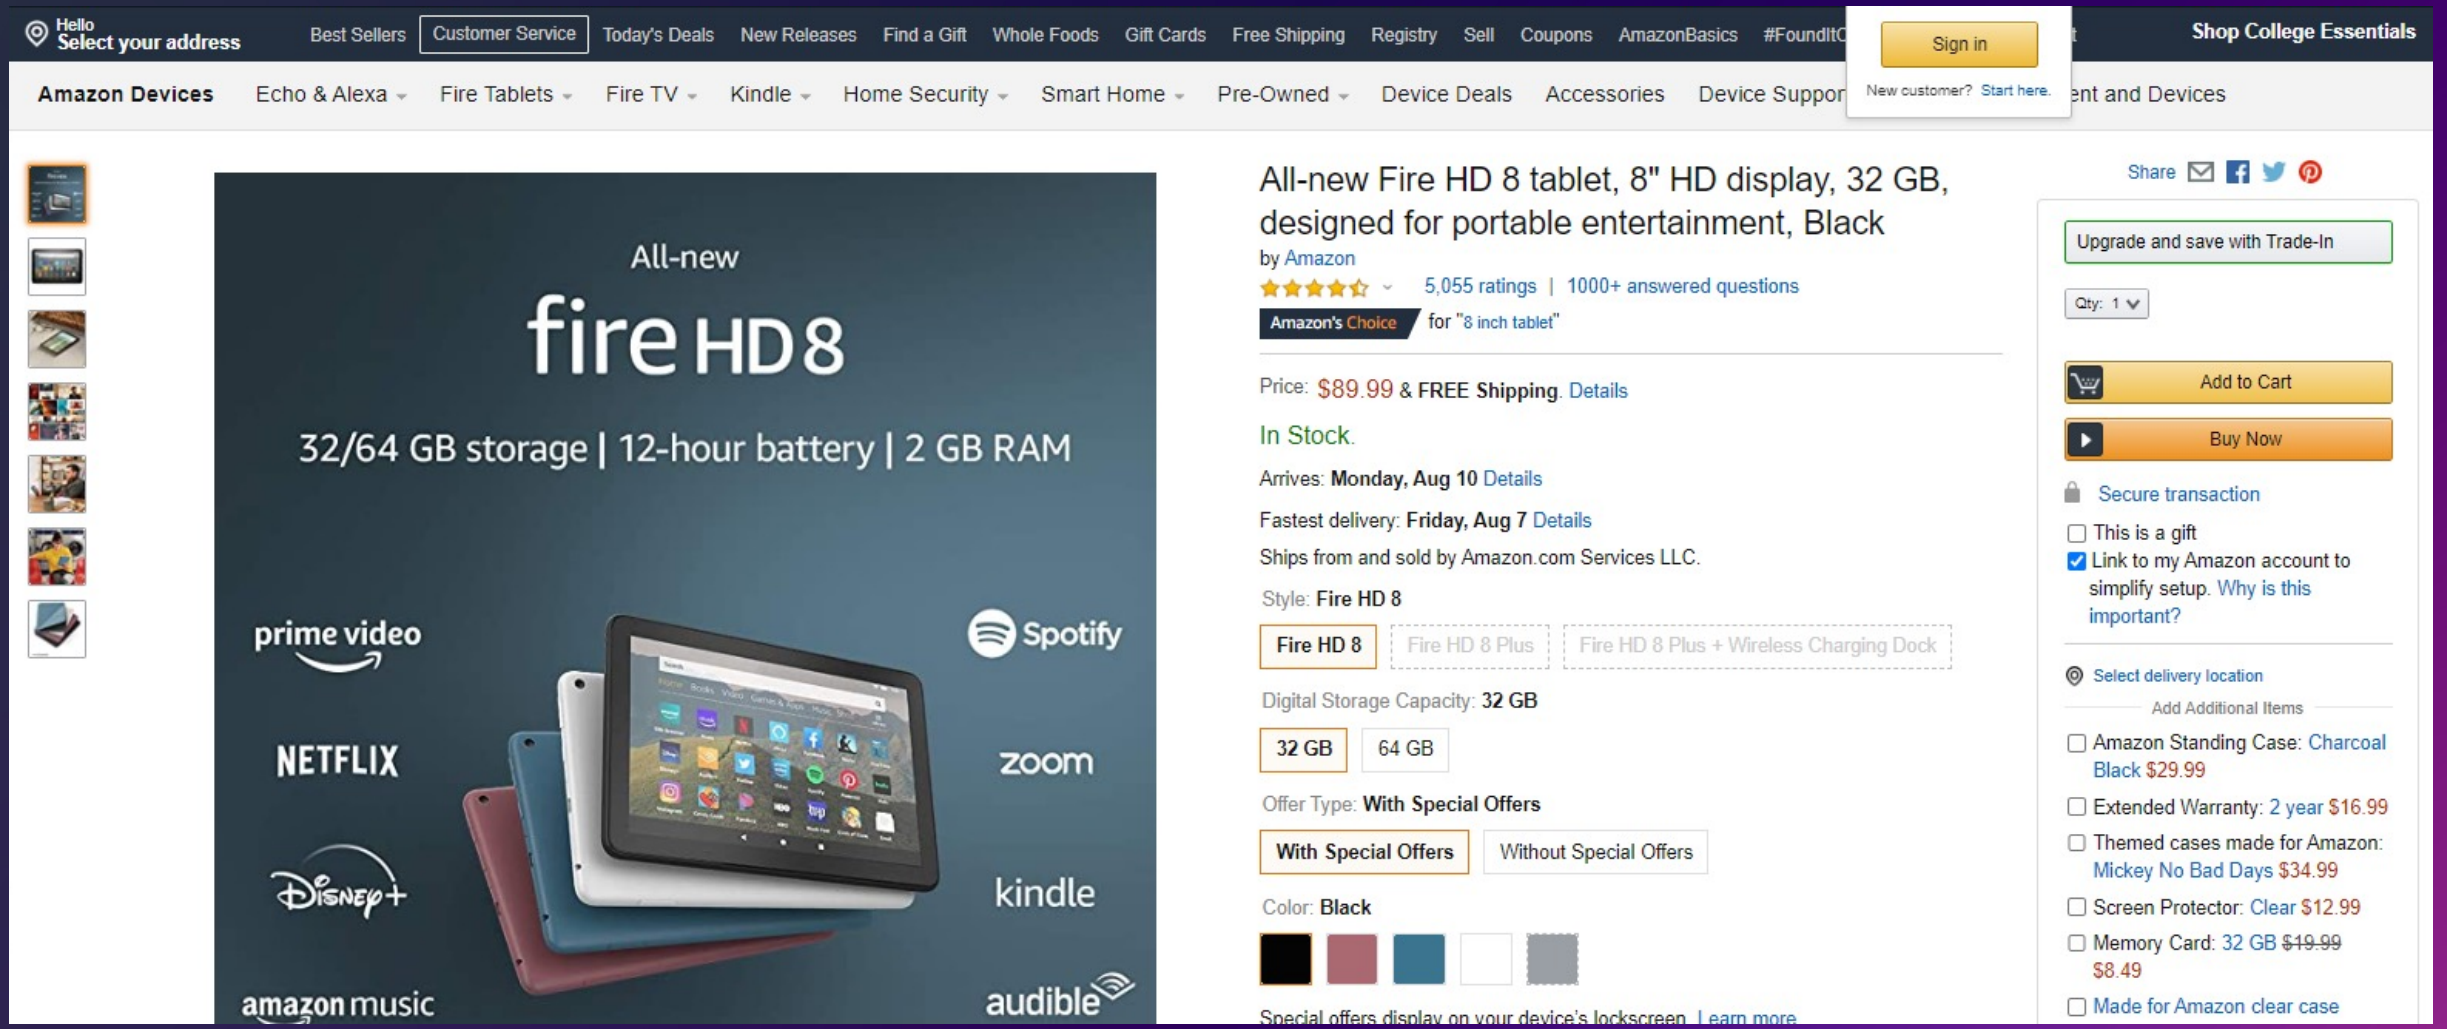

Microservices can enable resilience and scale

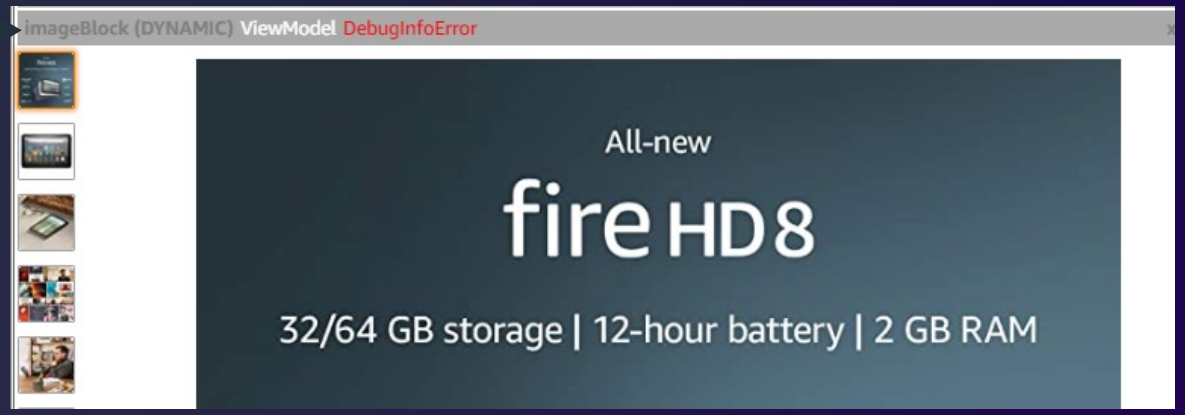

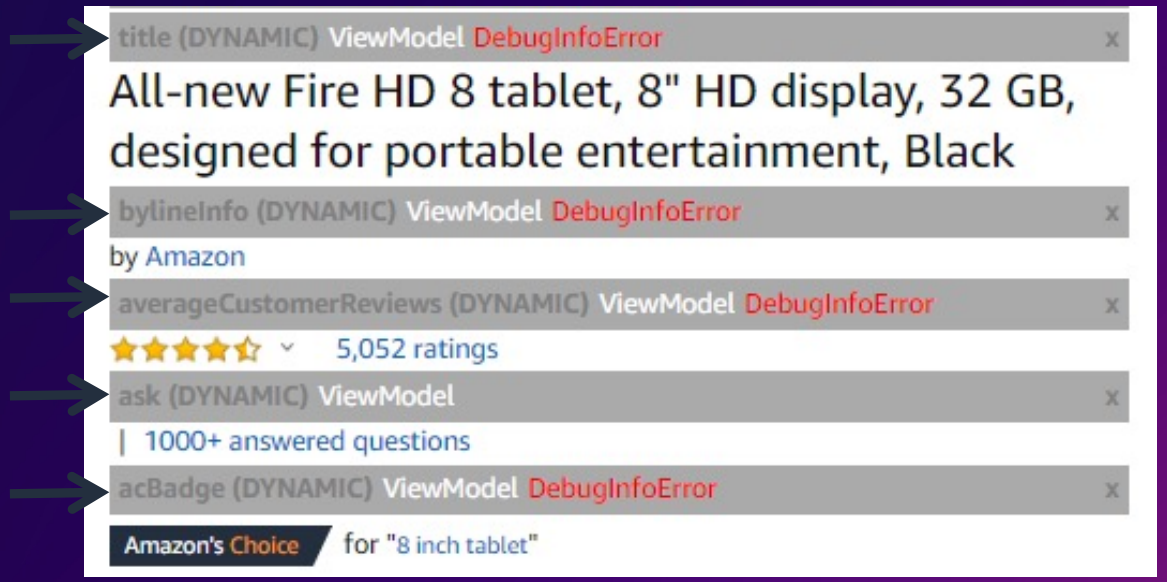

Prime Video and Amazon Music improve availability

Using fault isolation with cell-based architectures

Traditional scaling has gaps

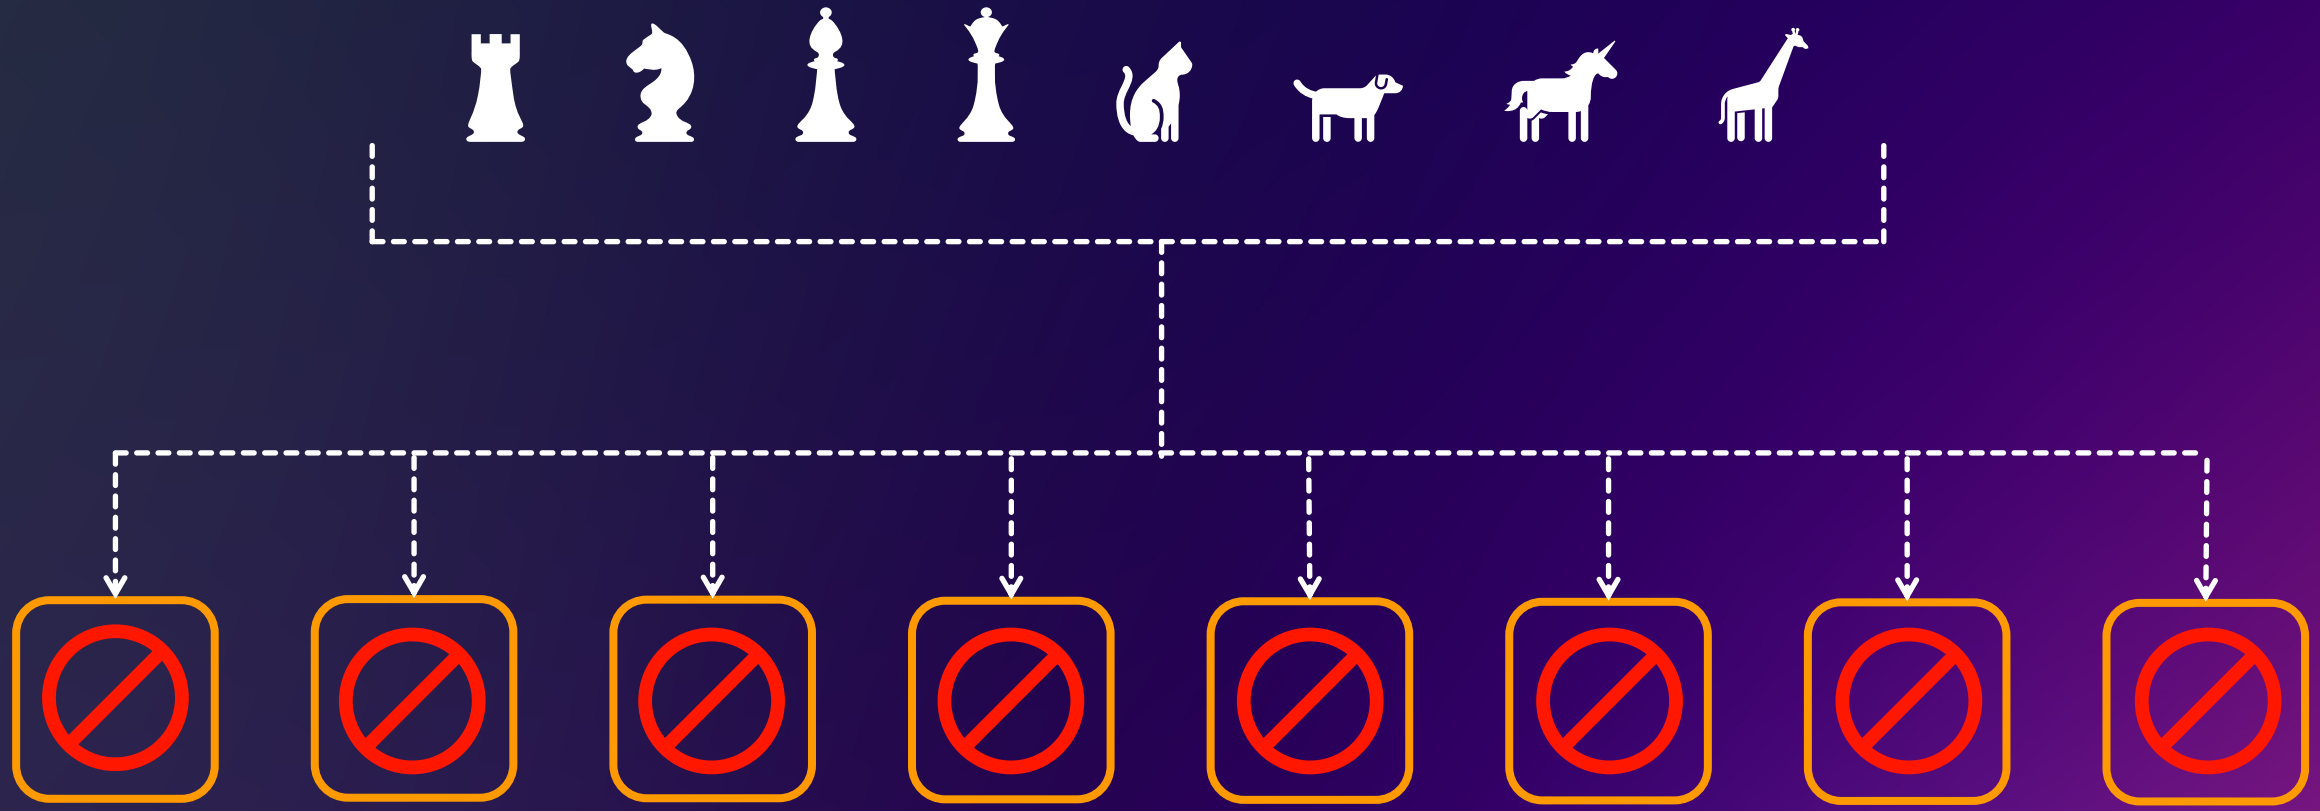

Blast radius = |*customers*|

Cell-based scaling and fault isolation

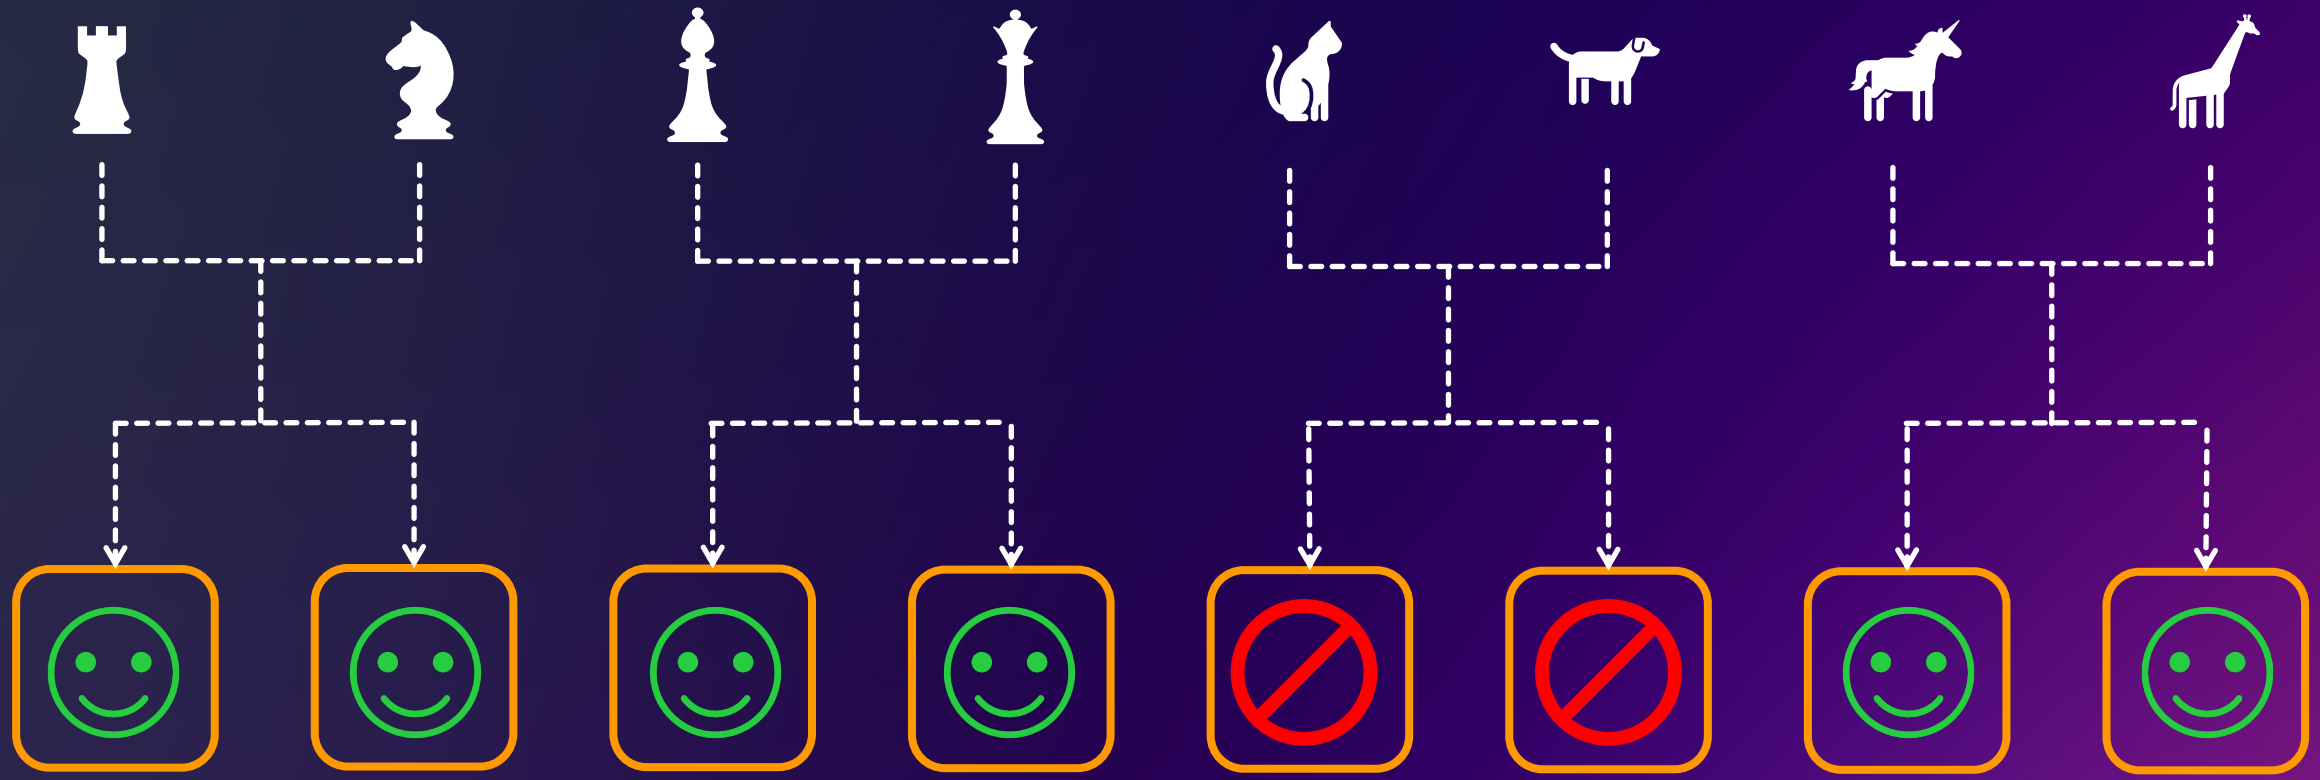

$$\text{Blast radius} = \frac{|customers|}{|cells|}$$

Cell-based architecture

Design pattern where a service is split into multiple deployment stacks, called “cells”

- Independent instance of the service
- Independently service the full workload of customers
- Cells share nothing

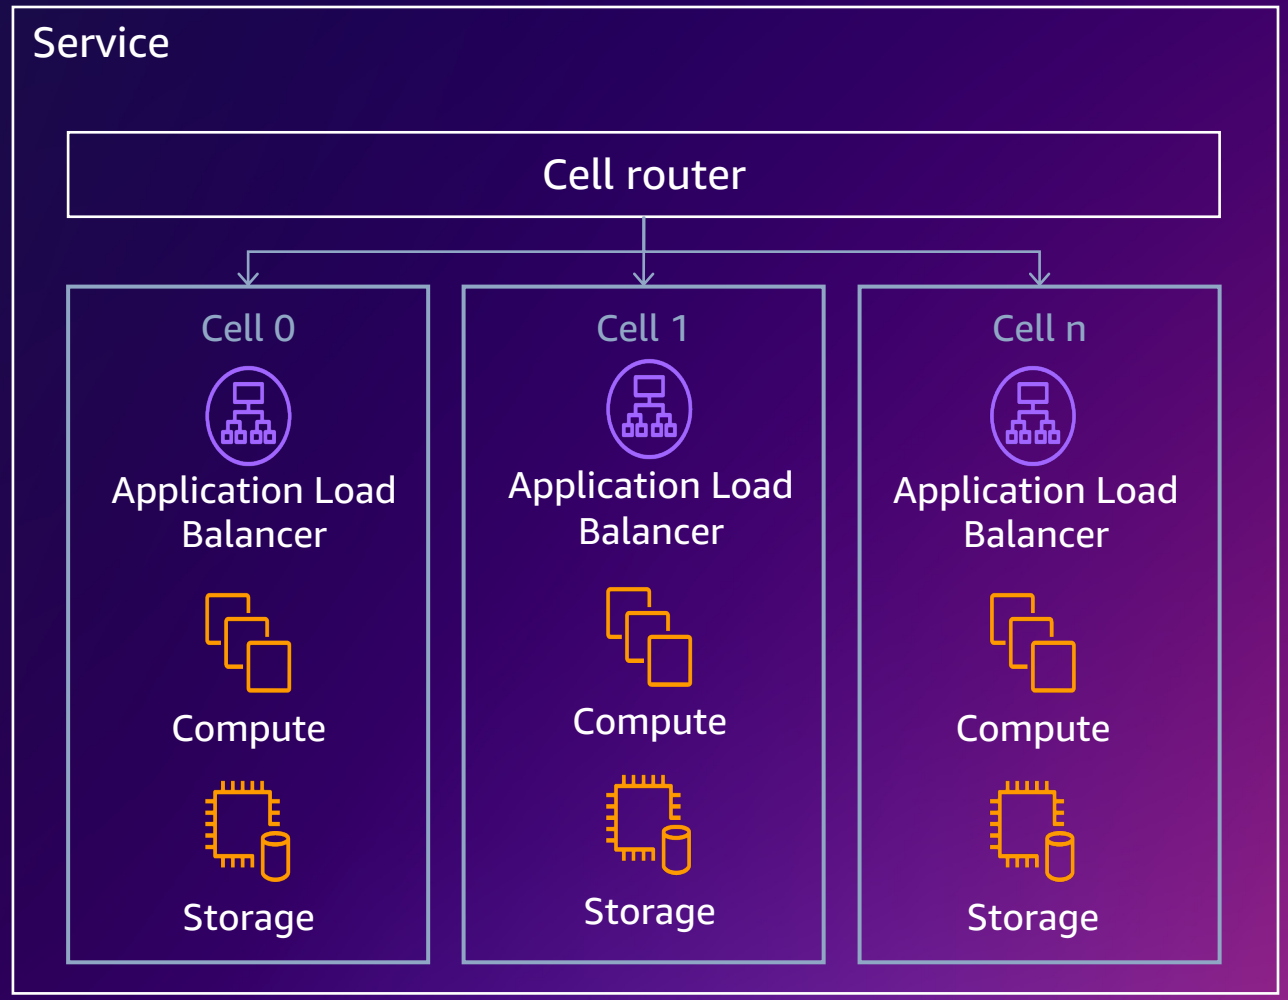

Use cases - goals

prime video

Simplifying global setups

amazon music

Fault isolation

Cells in each Region

KEY DECISIONS

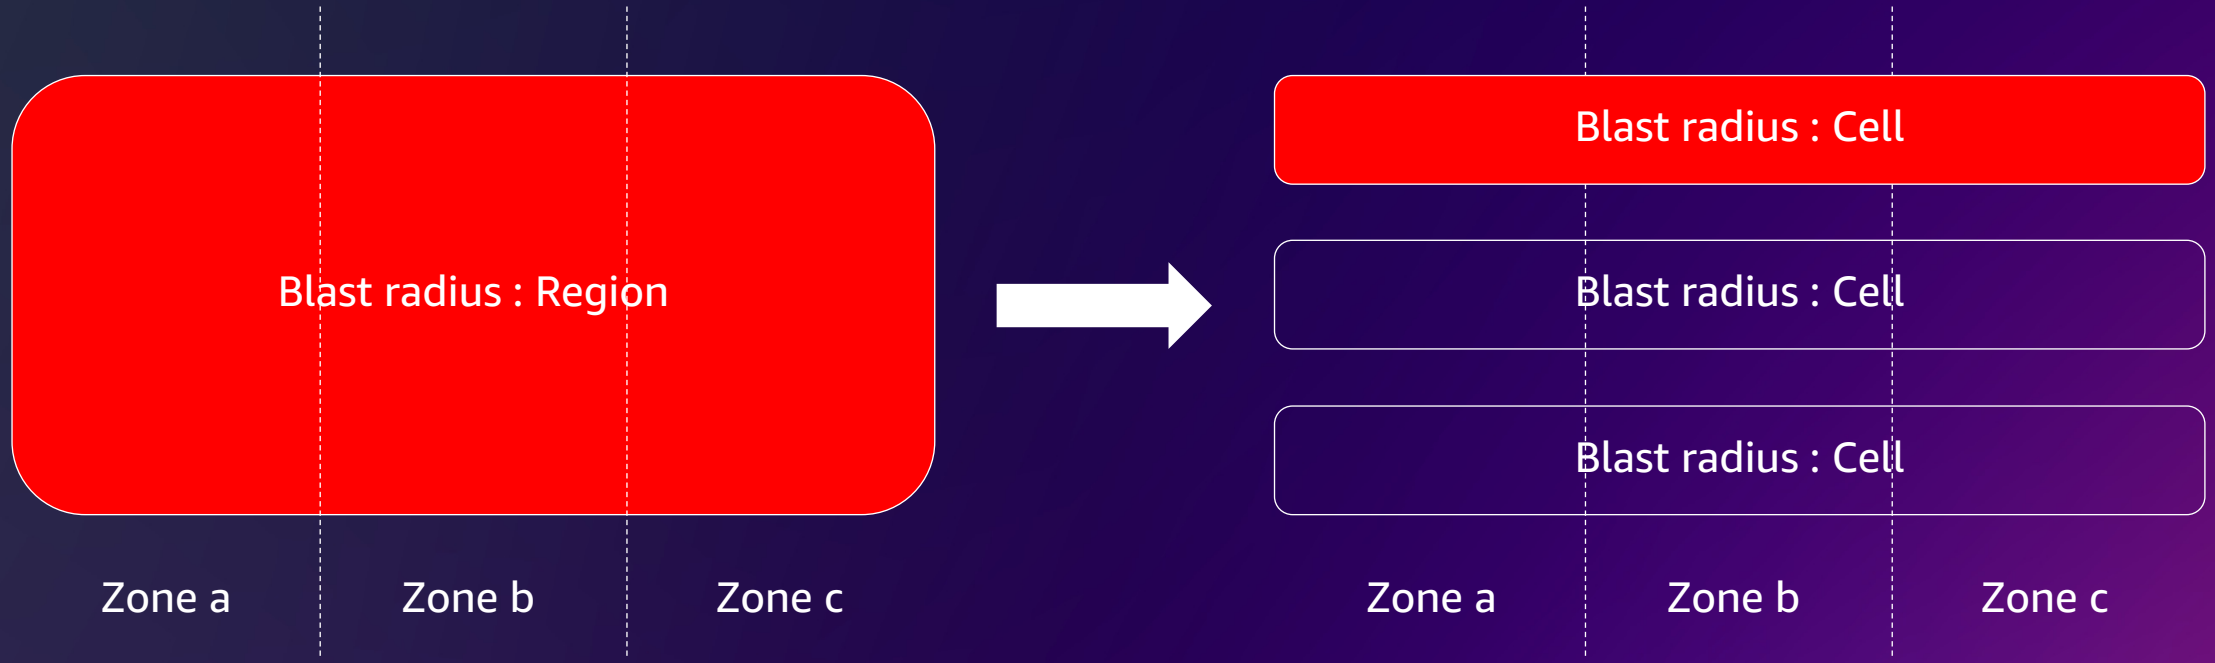

AWS Fault Isolation Boundaries Whitepaper:

Cellular traffic policy to isolate failures

KEY DECISIONS

First round robin routing to a cell DNS

Second geo-proximity routing to the closest Region

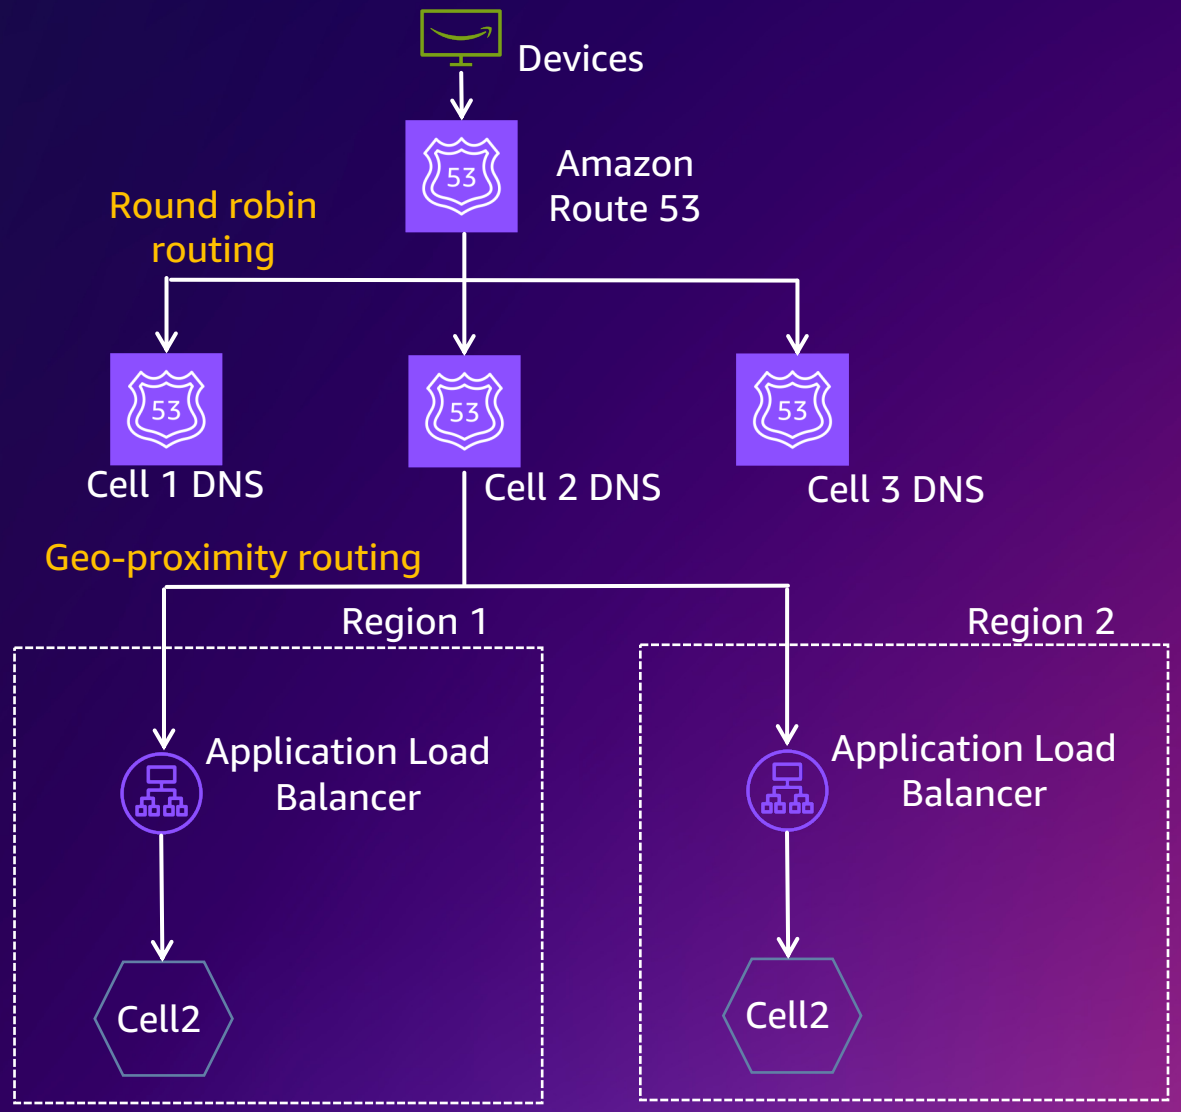

Calculated health check

KEY DECISIONS

- Route 53 health check that calls bootstrap API
- CloudWatch-alarm health check based on ELB 5xx errors (>100 errors for 1 minute)

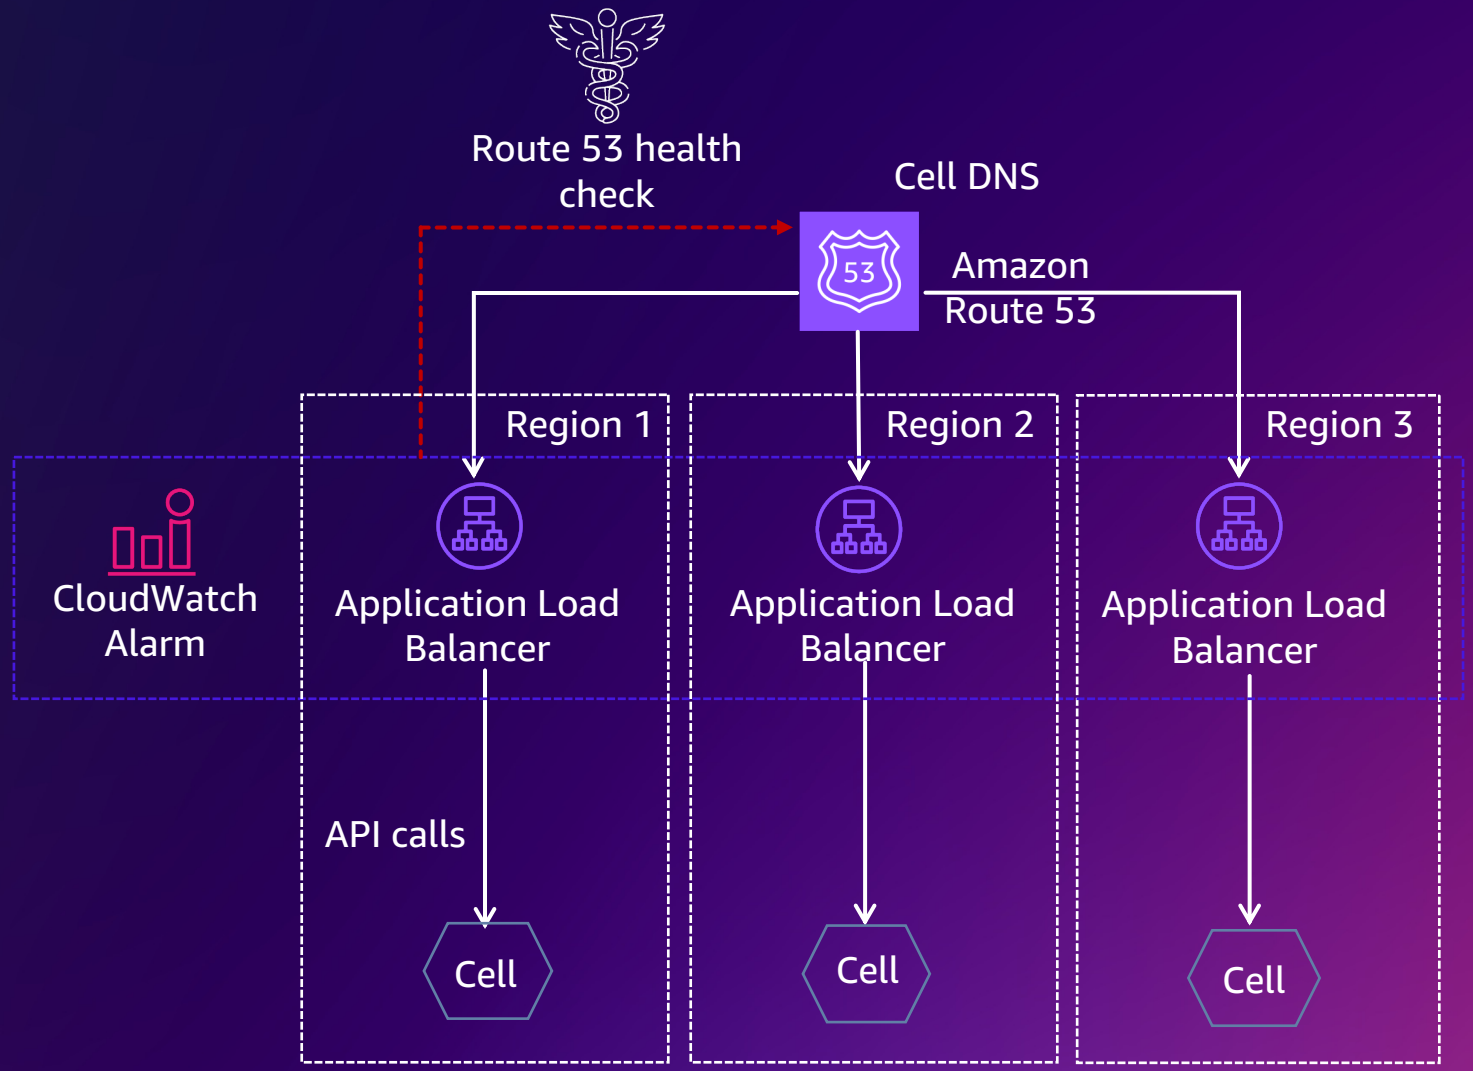

Outcomes - Increased availability

prime video

99.9996% availability

Improved ability to failover that comes with cellularisation

$$\text{Availability} = \frac{(\textit{Total Requests} - \textit{Errors})}{\textit{Total Requests}}$$

Global edge service availability (4 weeks)

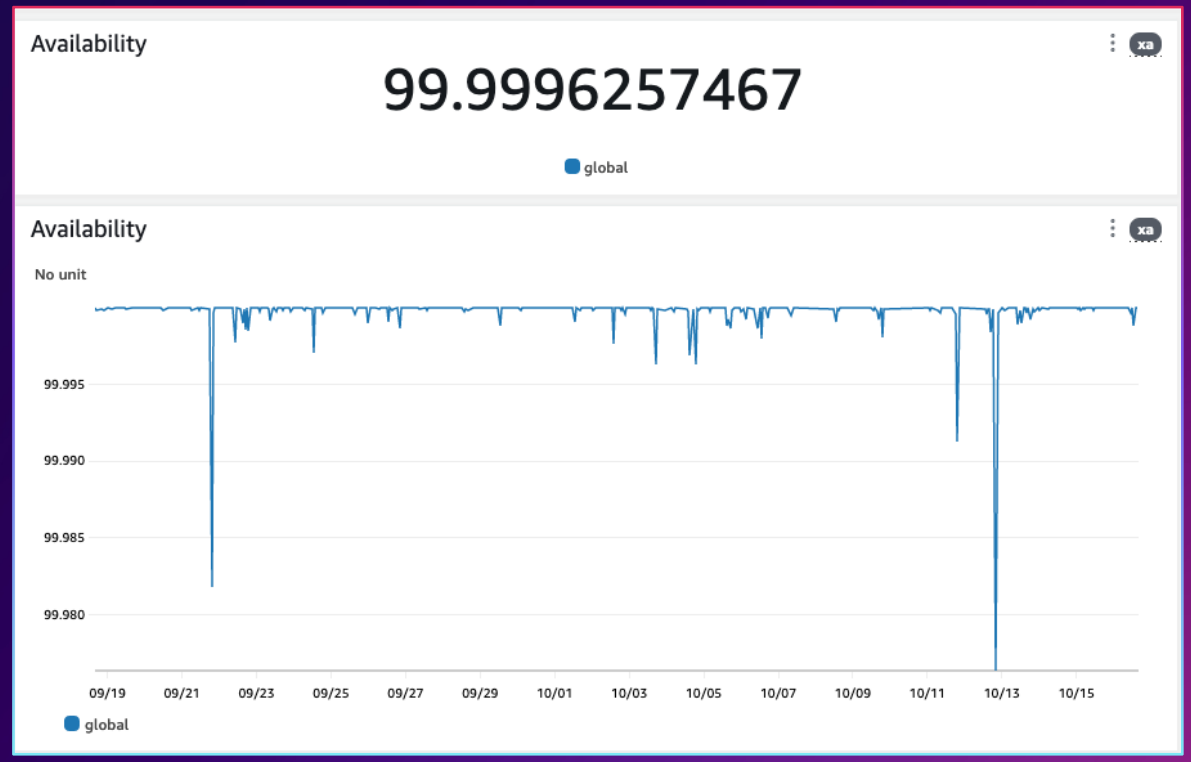

Cell routing and authentication logic

amazon music

KEY DECISIONS

AWS Fargate service hosts both cell-routing logic and authentication logic

Static routing configuration stored as config as code

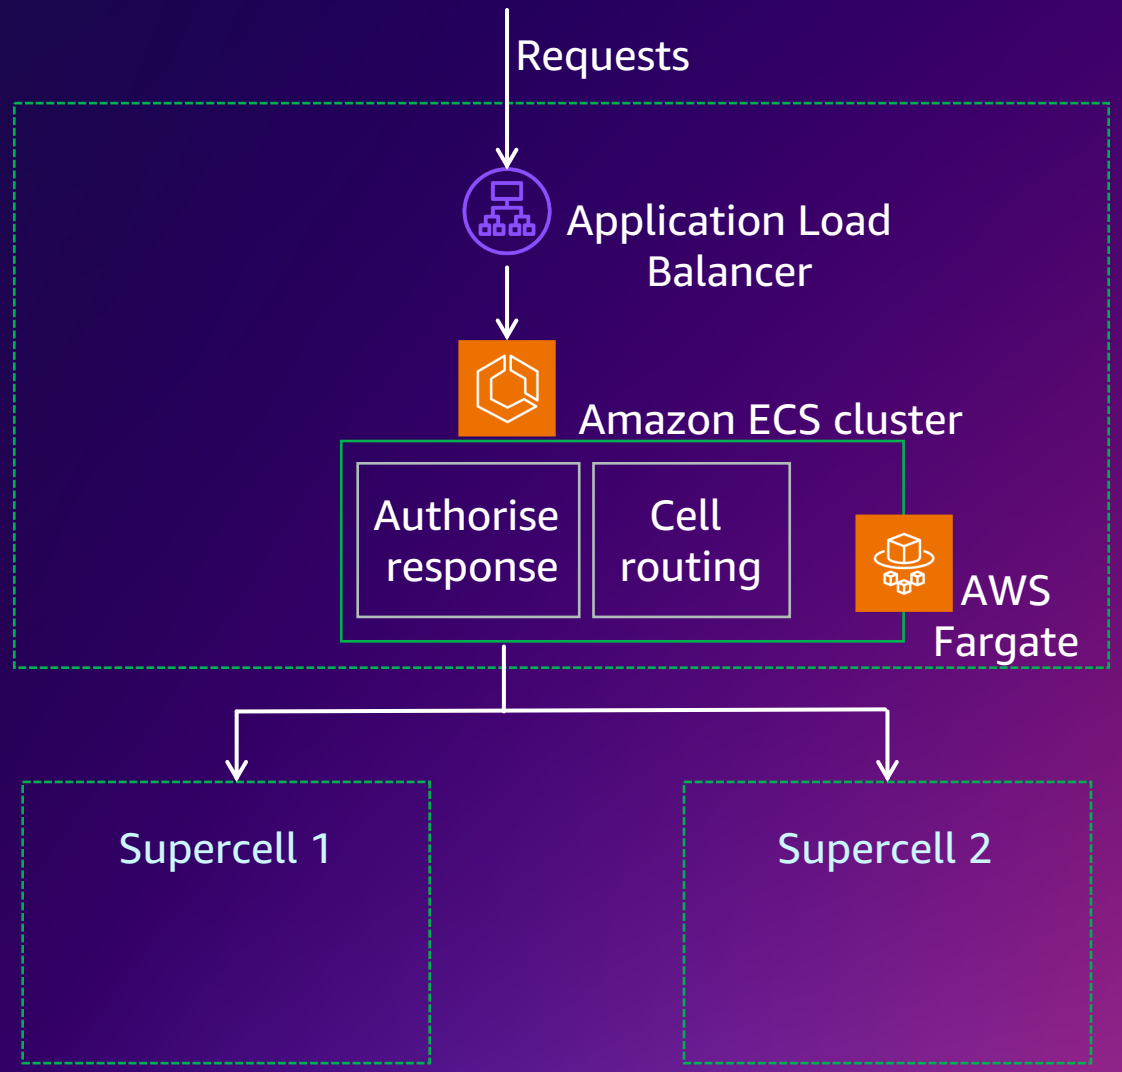

Cell mapping strategy

amazon music

KEY DECISIONS

First layer: manual assignment using device type

Second layer: route to cells base on event tier

- Event based mapping - event mapped to different tiers
- Partition key of device type

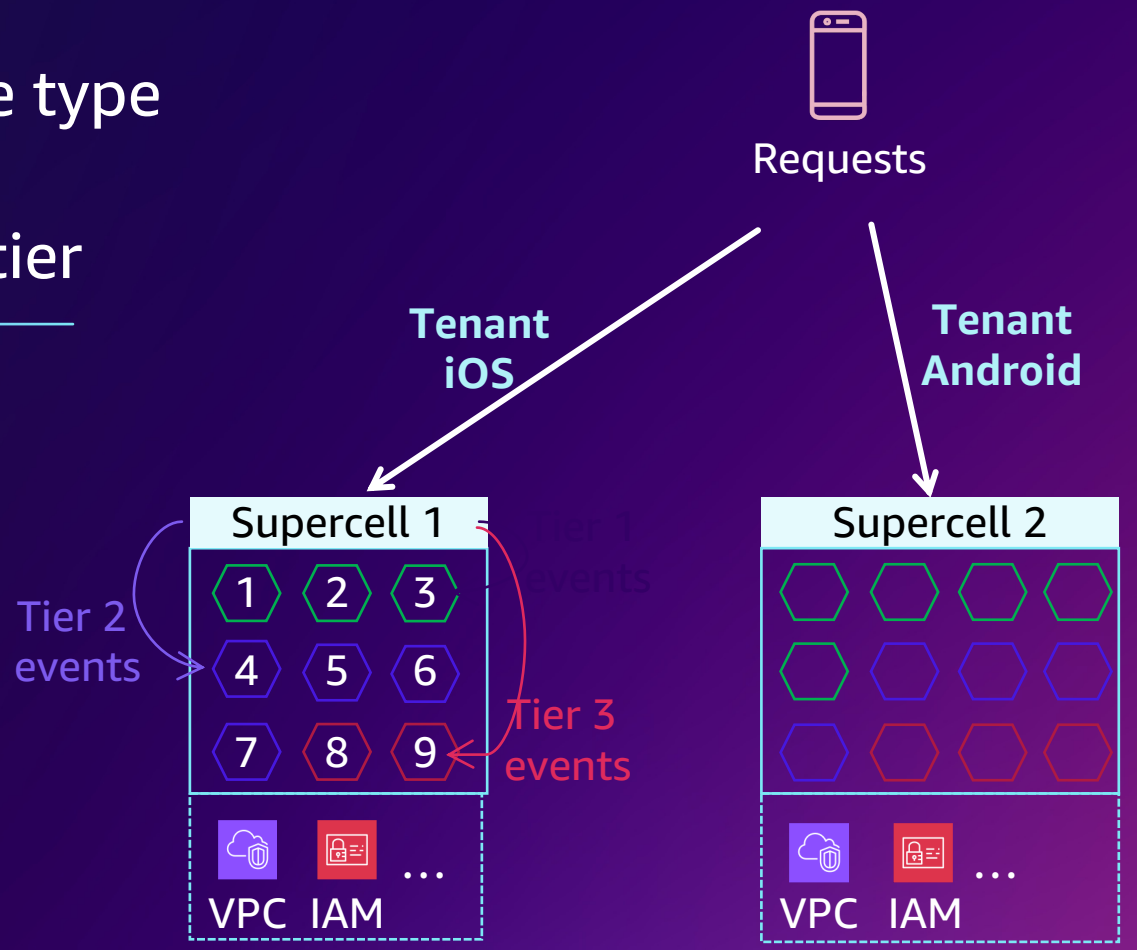

Cell size management model

KEY DECISIONS

amazon music

Atomic Cells = add more identical cells

Non-atomic Cells = scale existing cells up

Cellularisation as a requirement

amazon music

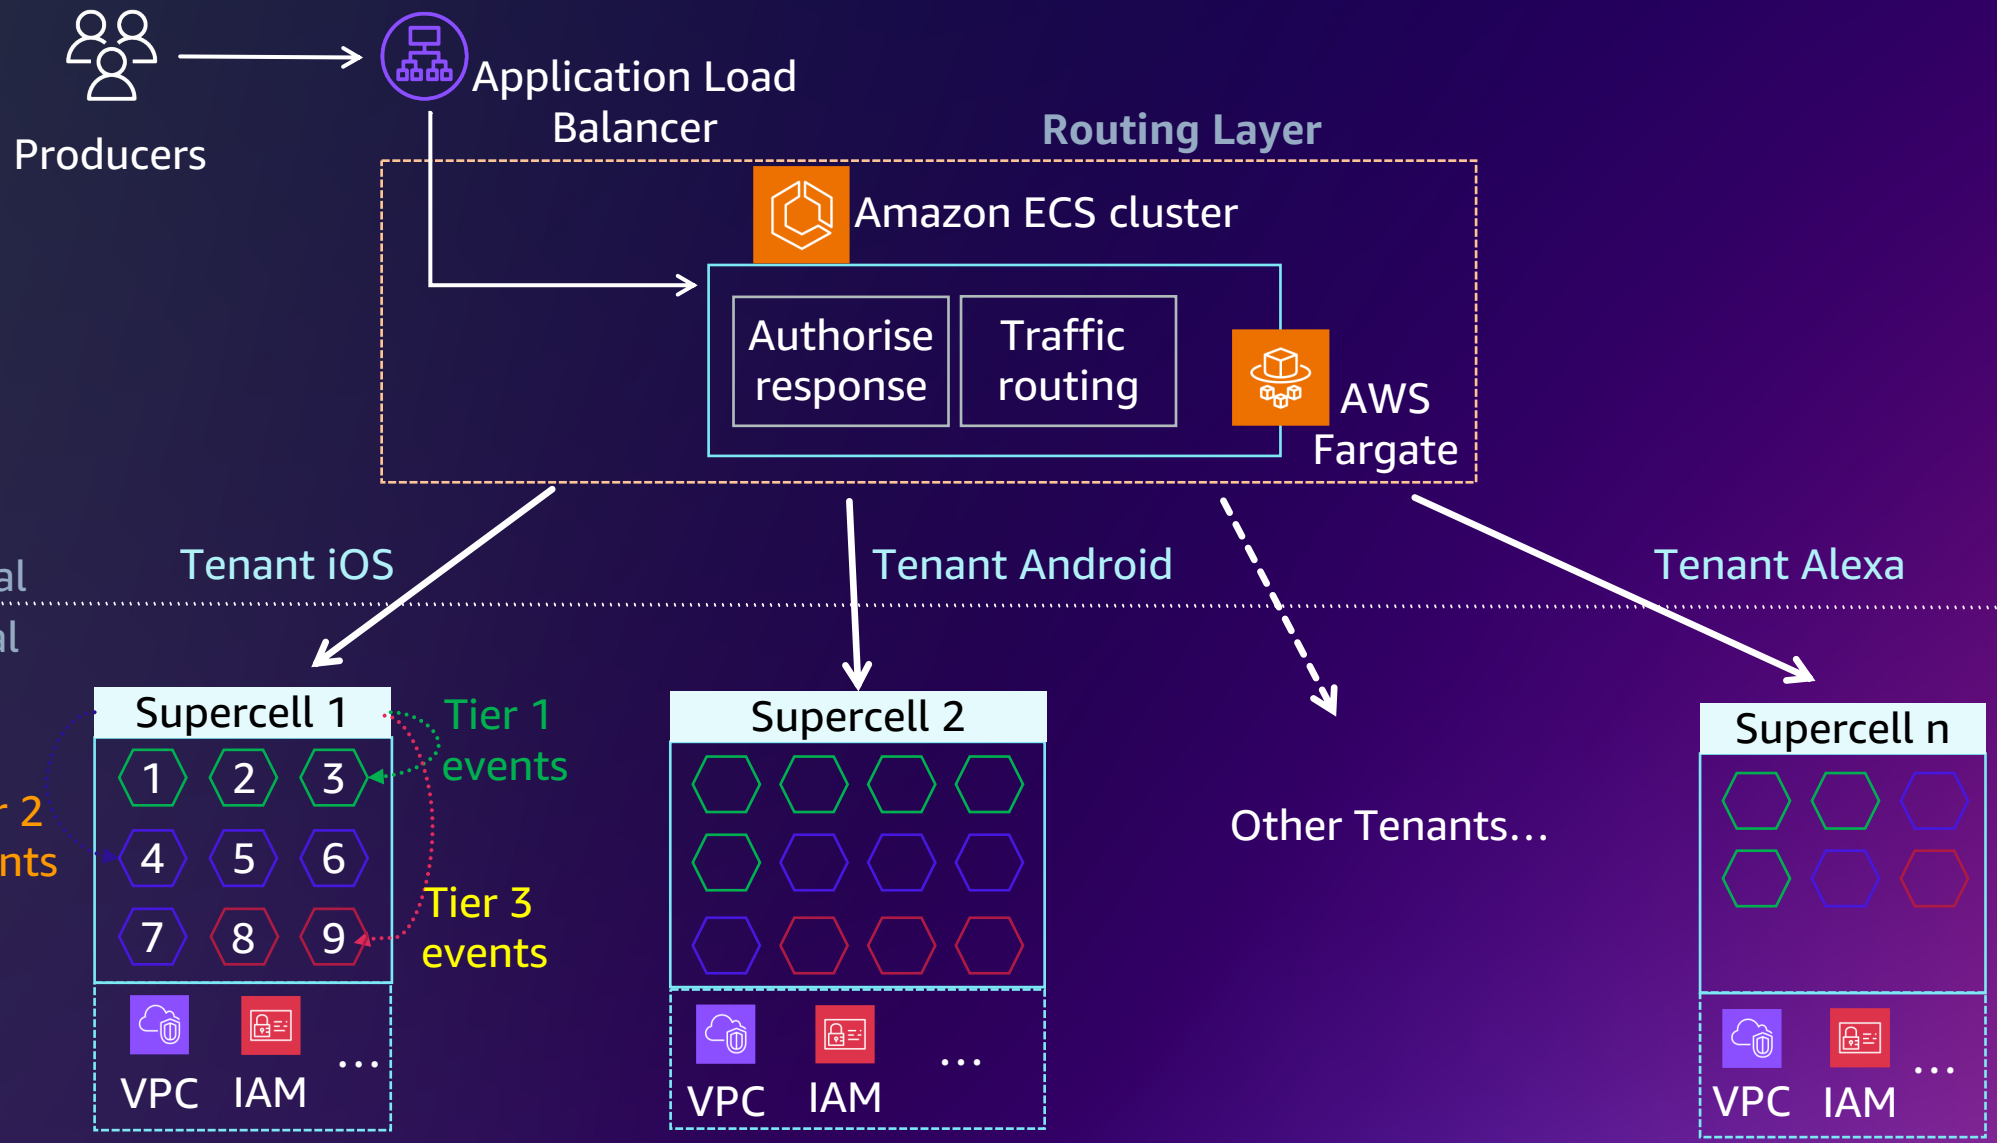

Summary

prime video

Stateless systems –

- ALB + Lambda + SQS

Routing policy –

- round robin and geo proximity

amazon music

Stateless systems –

- ALB + Lambda + Kinesis

Routing policy –

- device type and event based

Outcome – Increased Availability and Resilience

aws

Ring’s massive-scale event-driven architecture

Over 129K RPS with 99.9999% availability

Ring doorbells, camera, home security

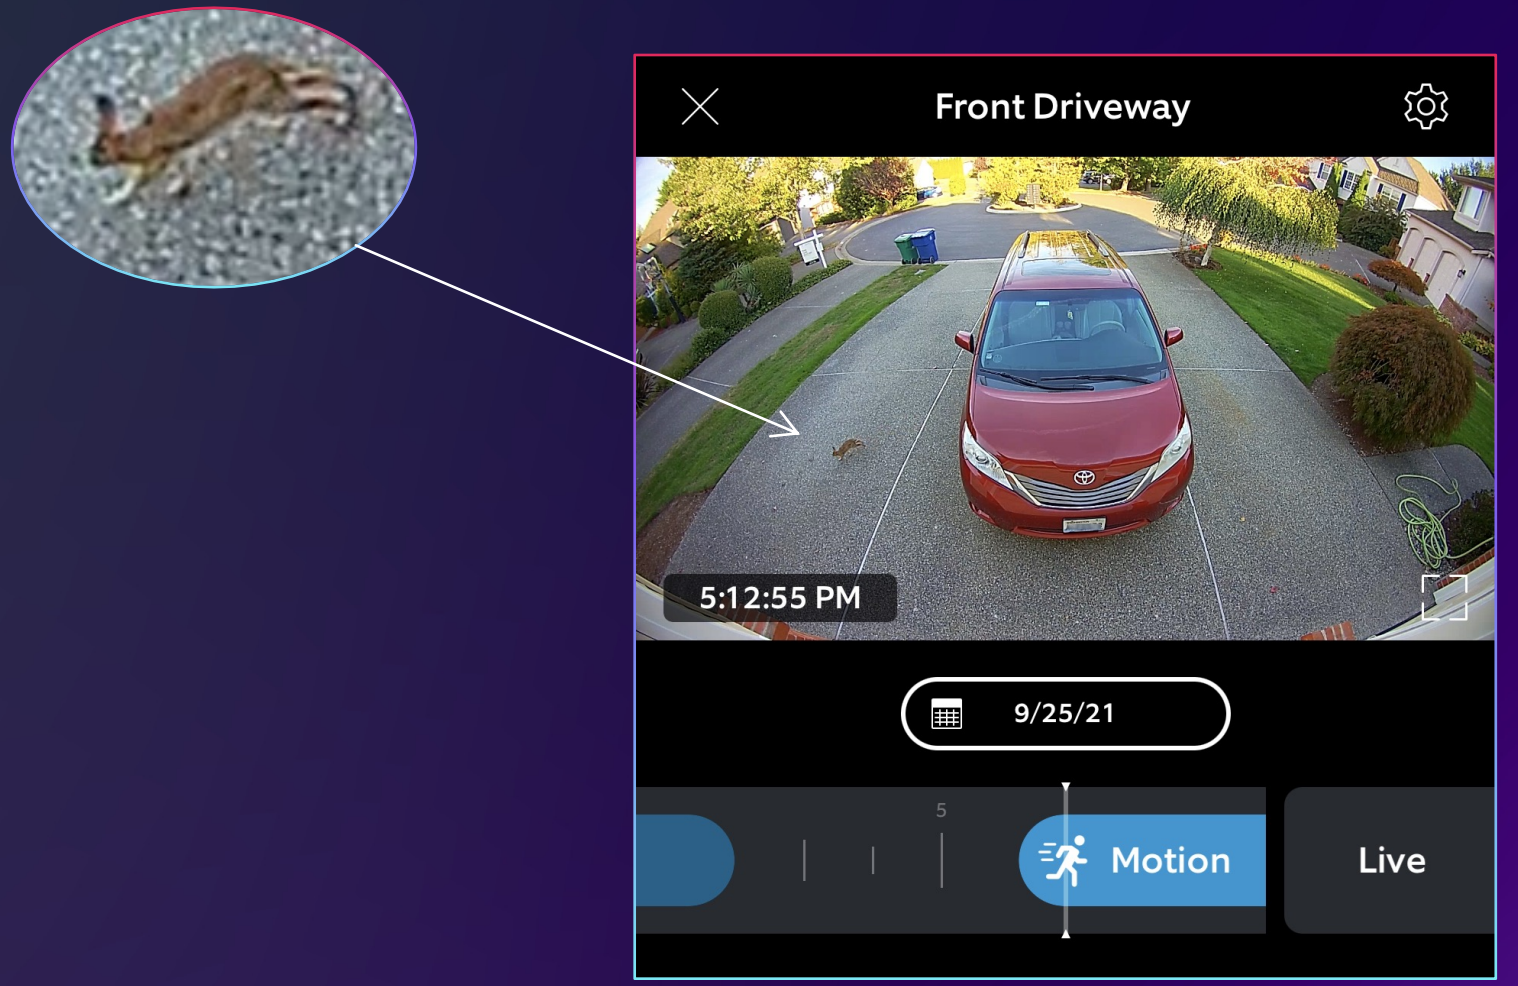

Ring video processing

Ring's "Prime Day": Halloween!

PREPARING FOR HIGH DEMAND

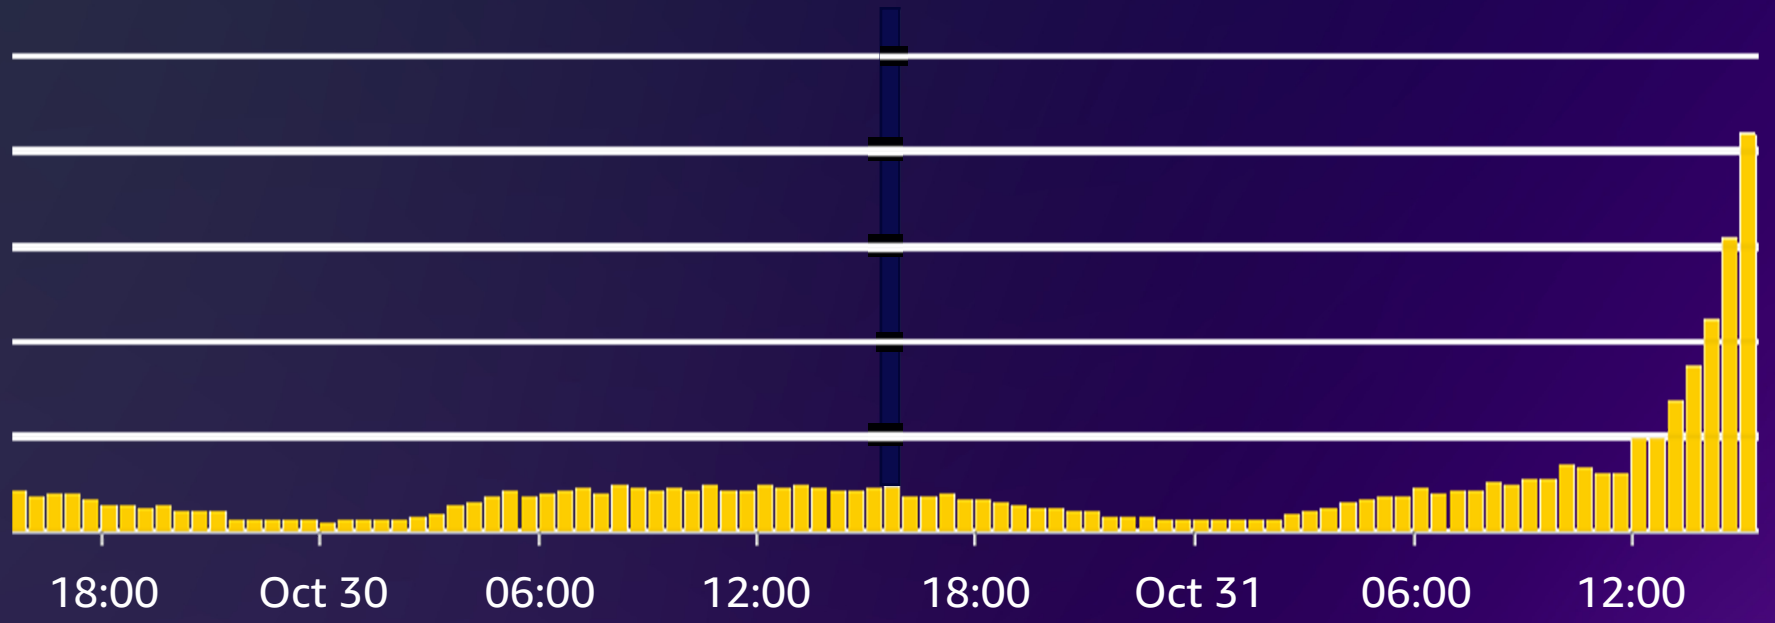

Auto scaling for Ring

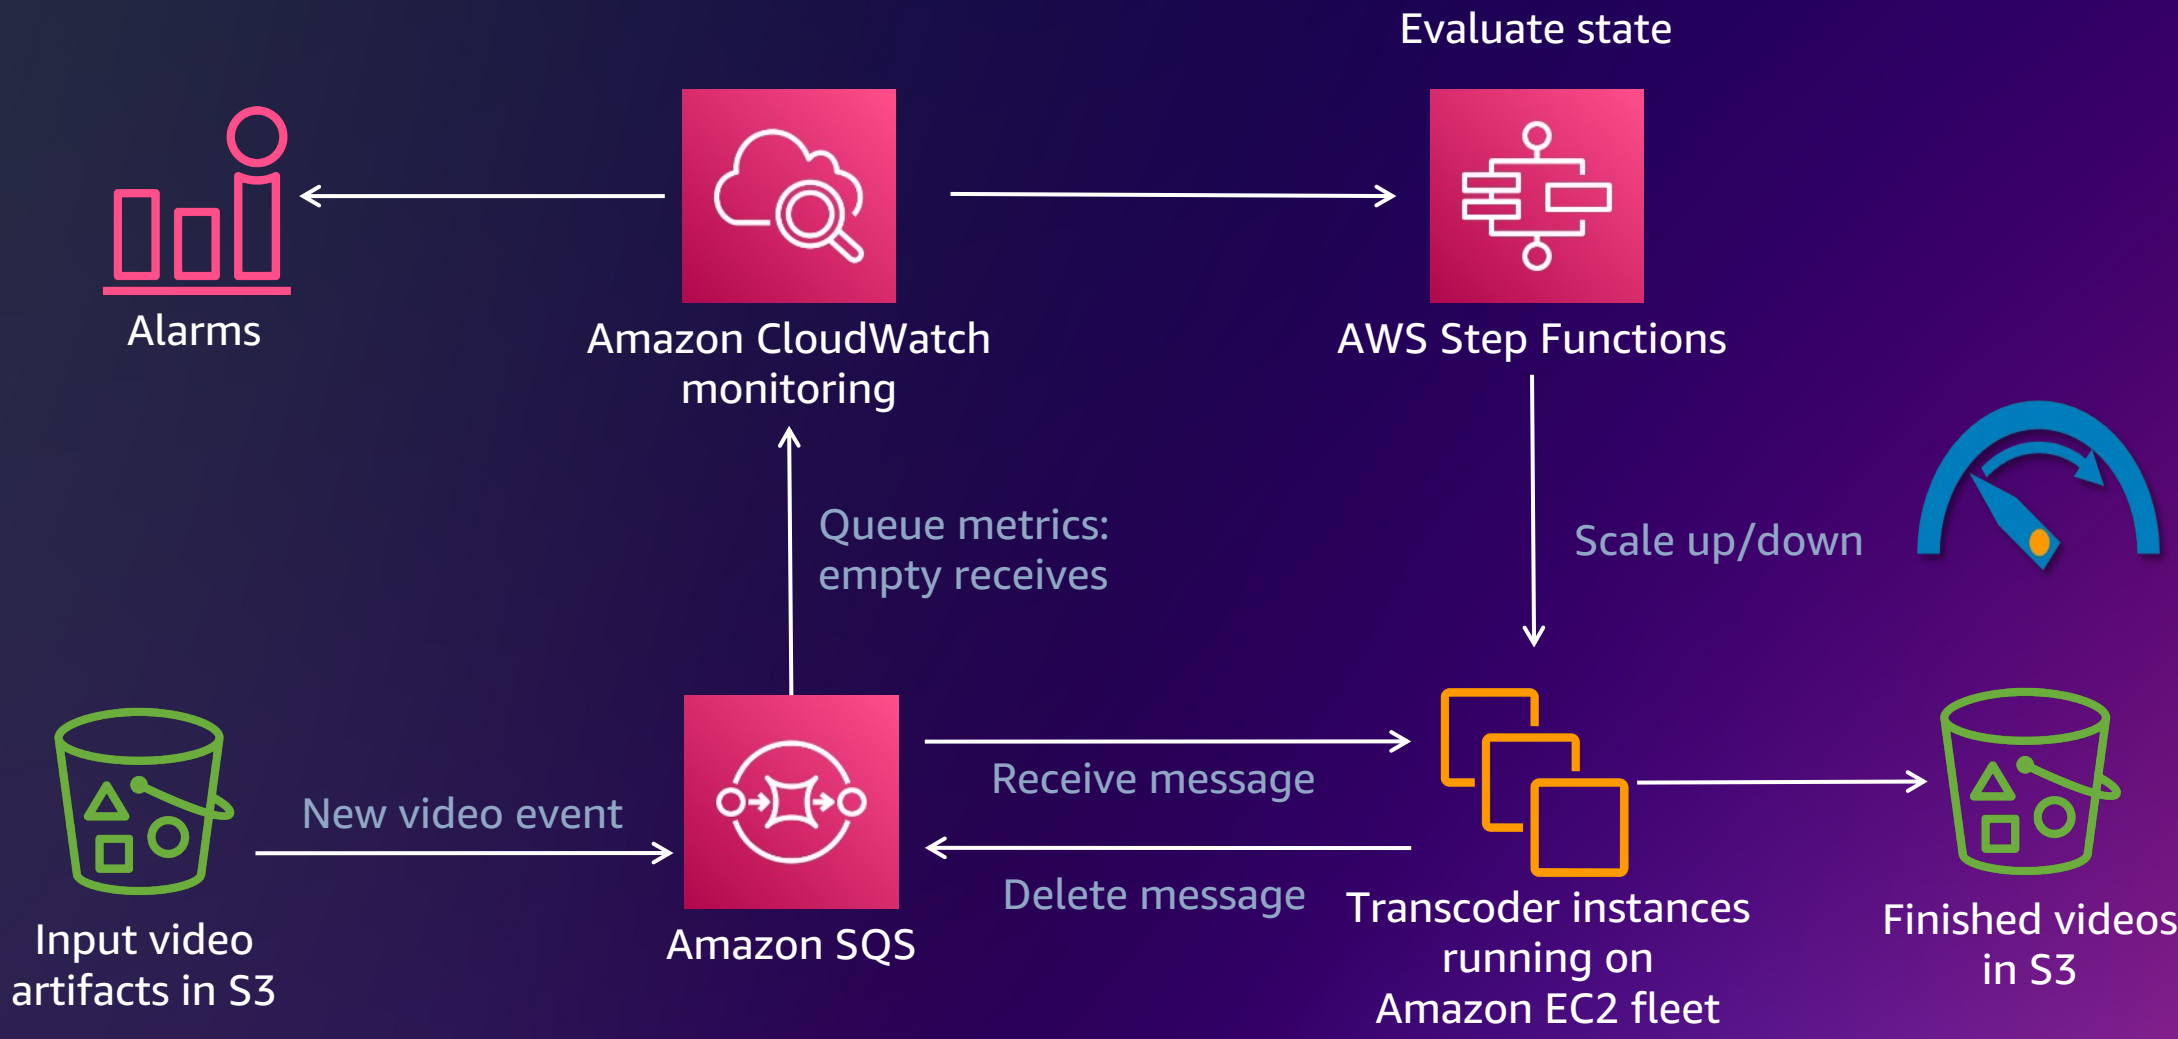

The quest for low latency

Previous story was about reducing latency to see video

Next is a story about reducing latency between all events and actions

Event		Consumer
Stream start	→	Unified notification service
Device detects motion		Sends push message to user device

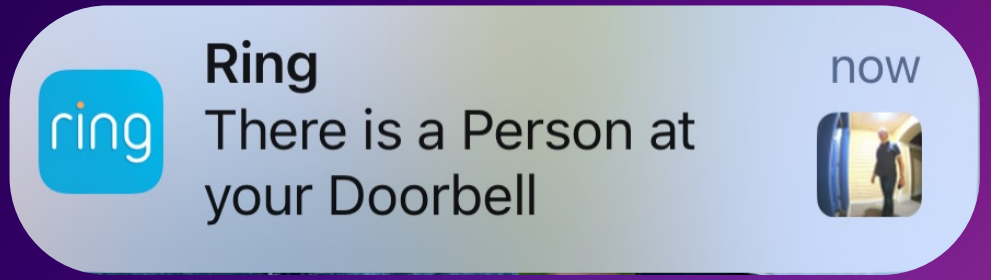

SEB

- Tiers:

  Service: API/Router

  Data Plane

  Consumer proxy

- Multi-cell

- Separate AWS Accounts

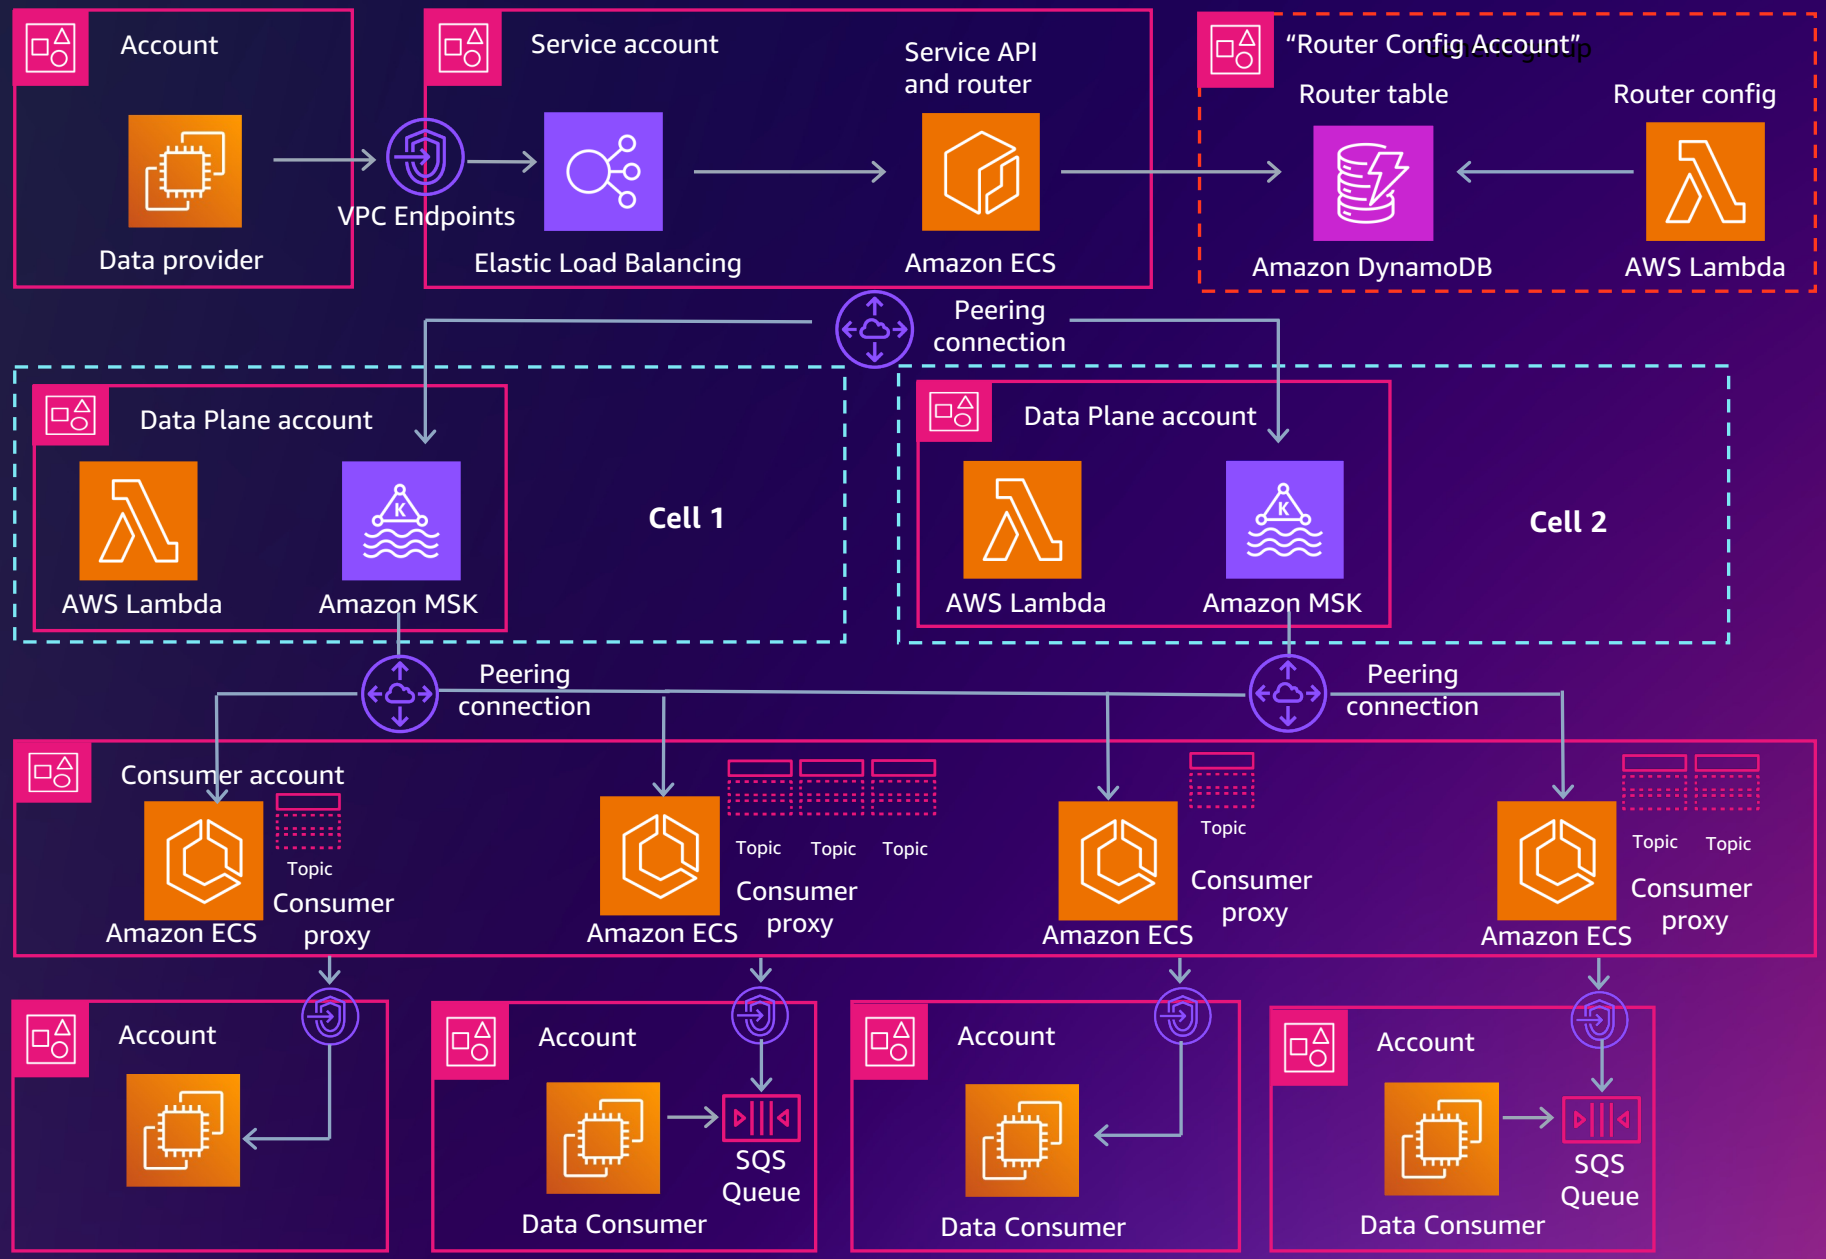

SEB

- Tiers:

  Service:
  API/Router

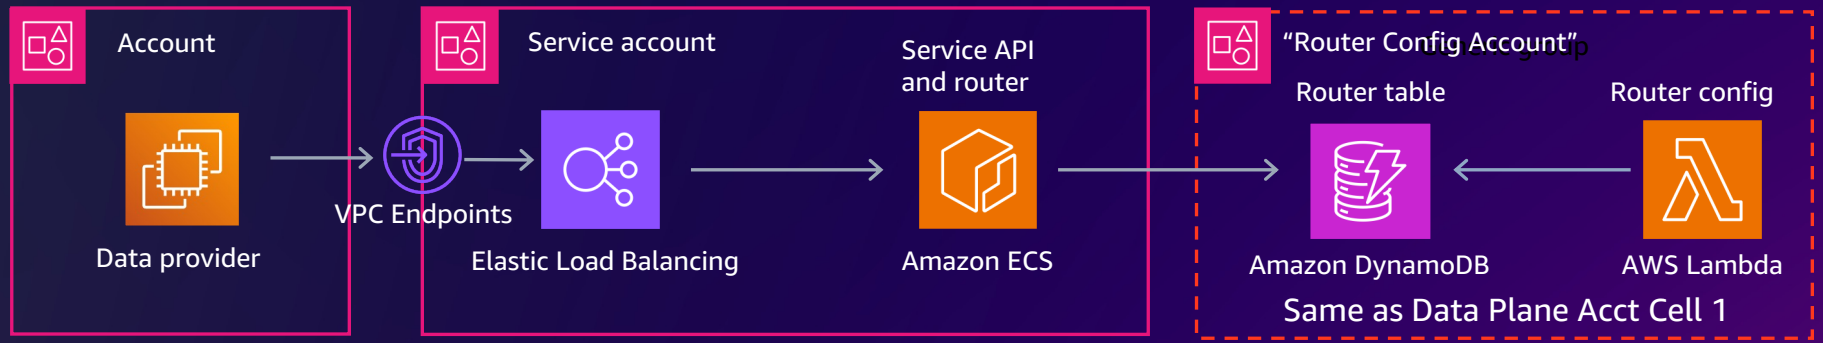

Streaming Event Bus (SEB)

- Tiers:

Data Plane

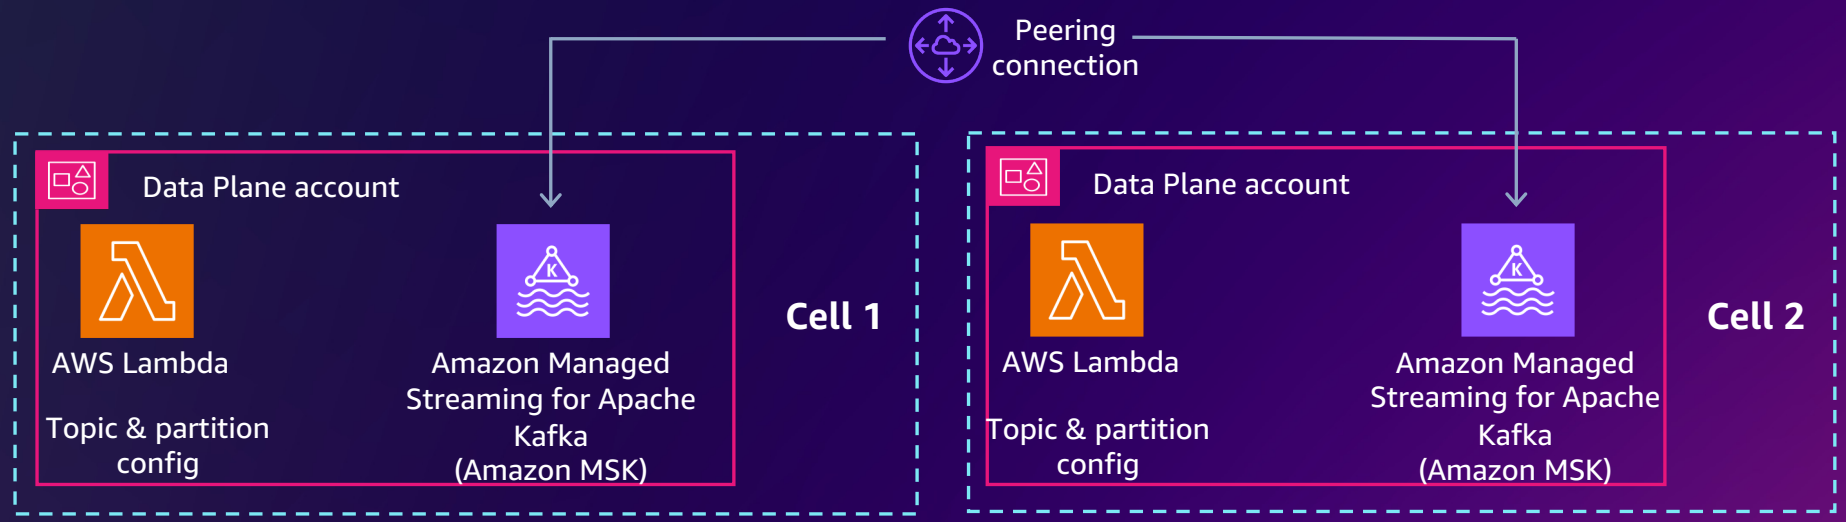

Streaming Event Bus (SEB)

- Tiers:

  Consumer proxy

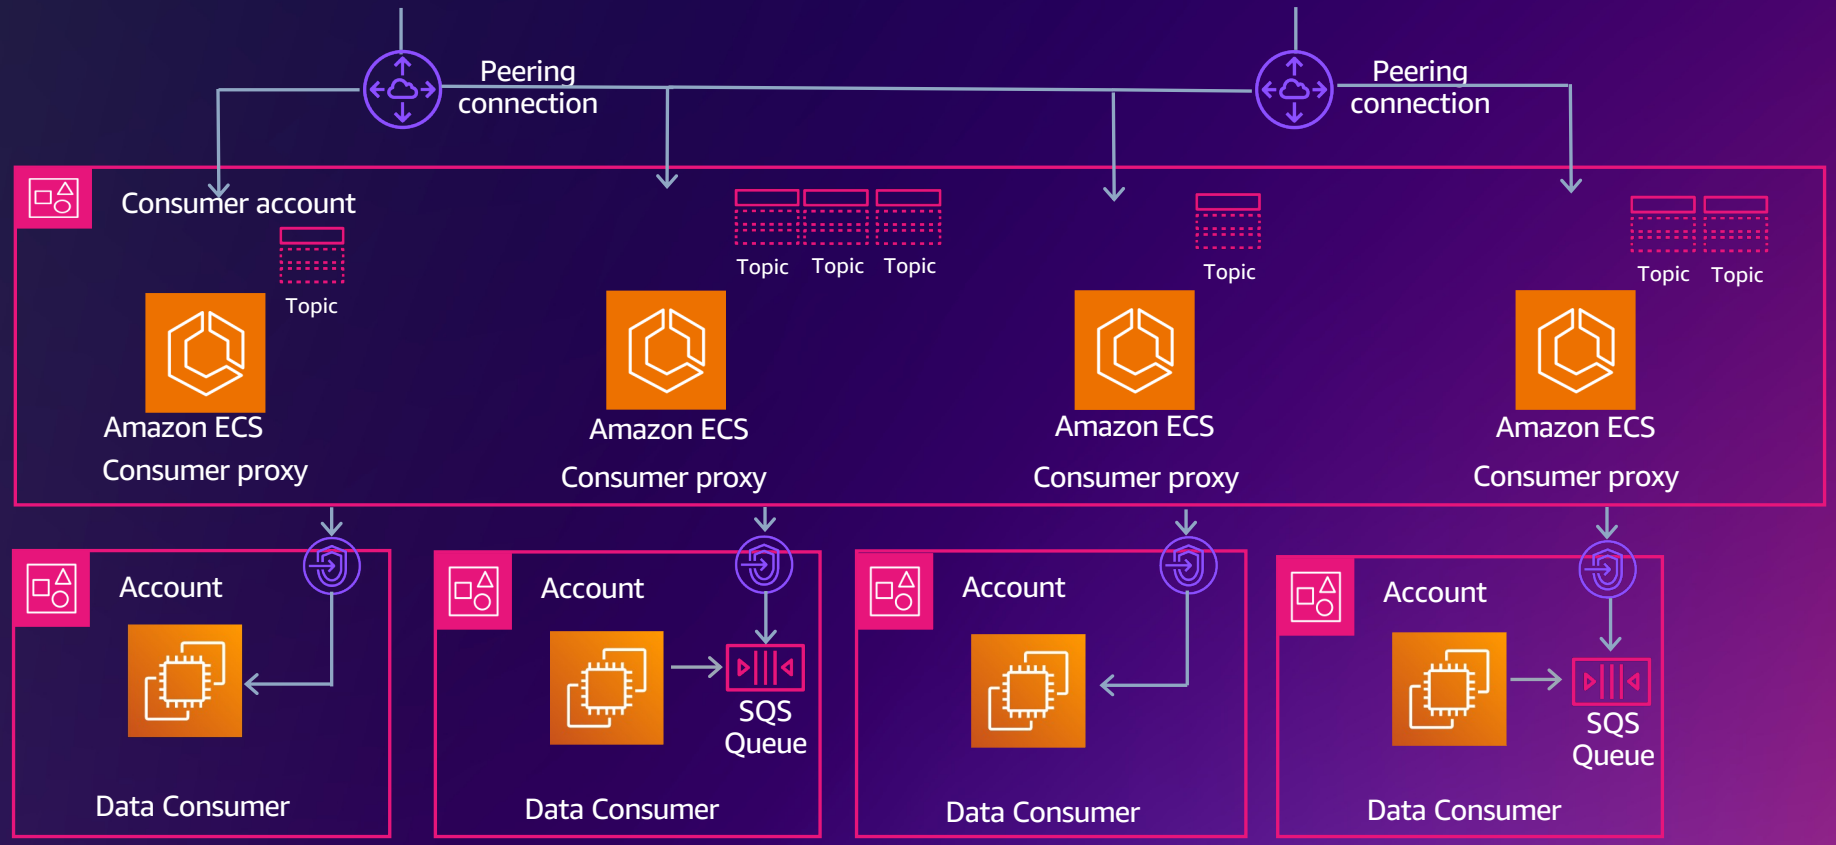

SEB

- Tiers:

  Service: API/Router

  Data Plane

  Consumer proxy

- Multi-cell
- Separate AWS Accounts

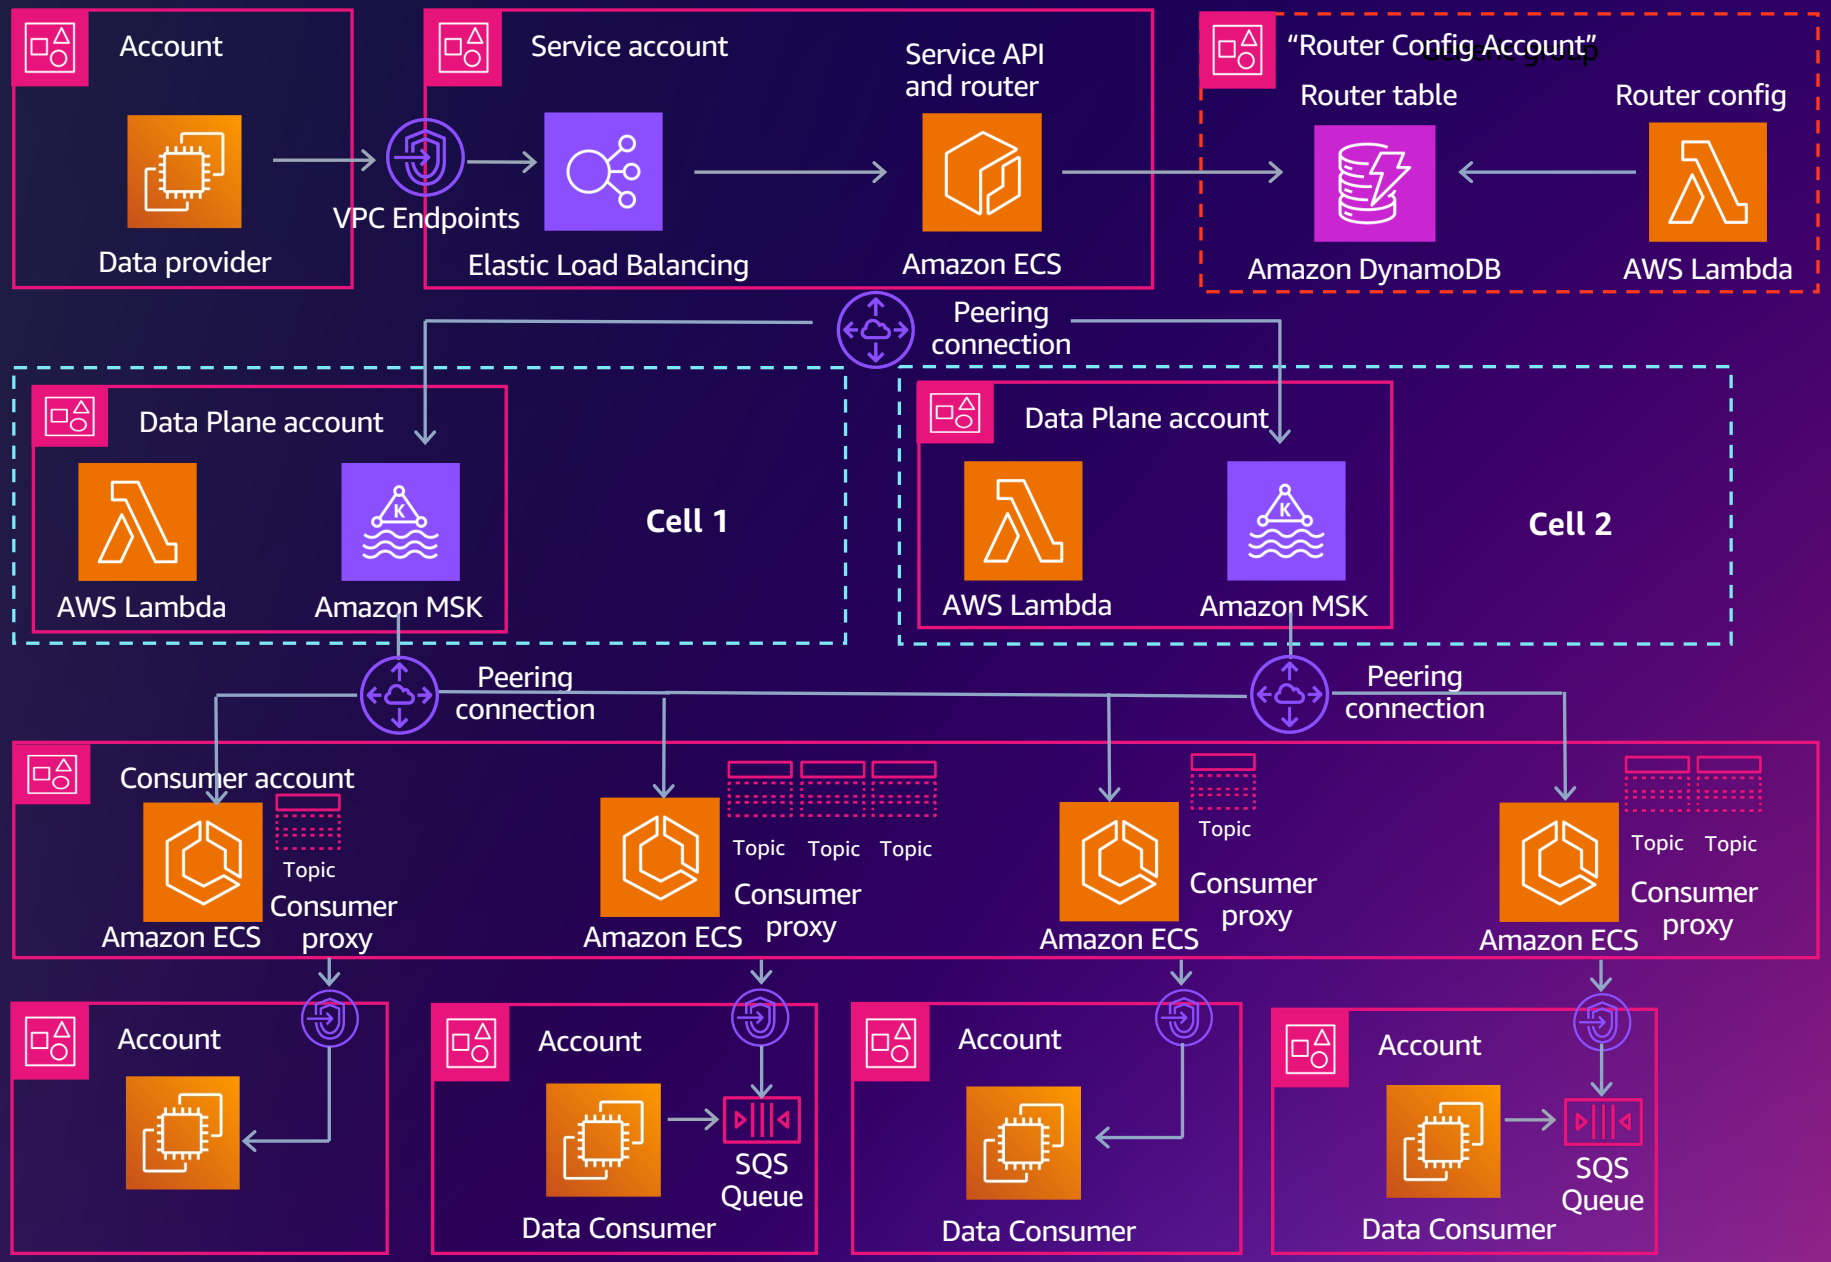

SEB – Active/active architecture

Cell-based architecture

Scaling: Topics mapped to a cell

Resilience: Impact: limited (to blast radius)

Active/active architecture

Any cell can handle all topics

Can failover

Resilience: Impact: none

Scaling: Still map topics to cell

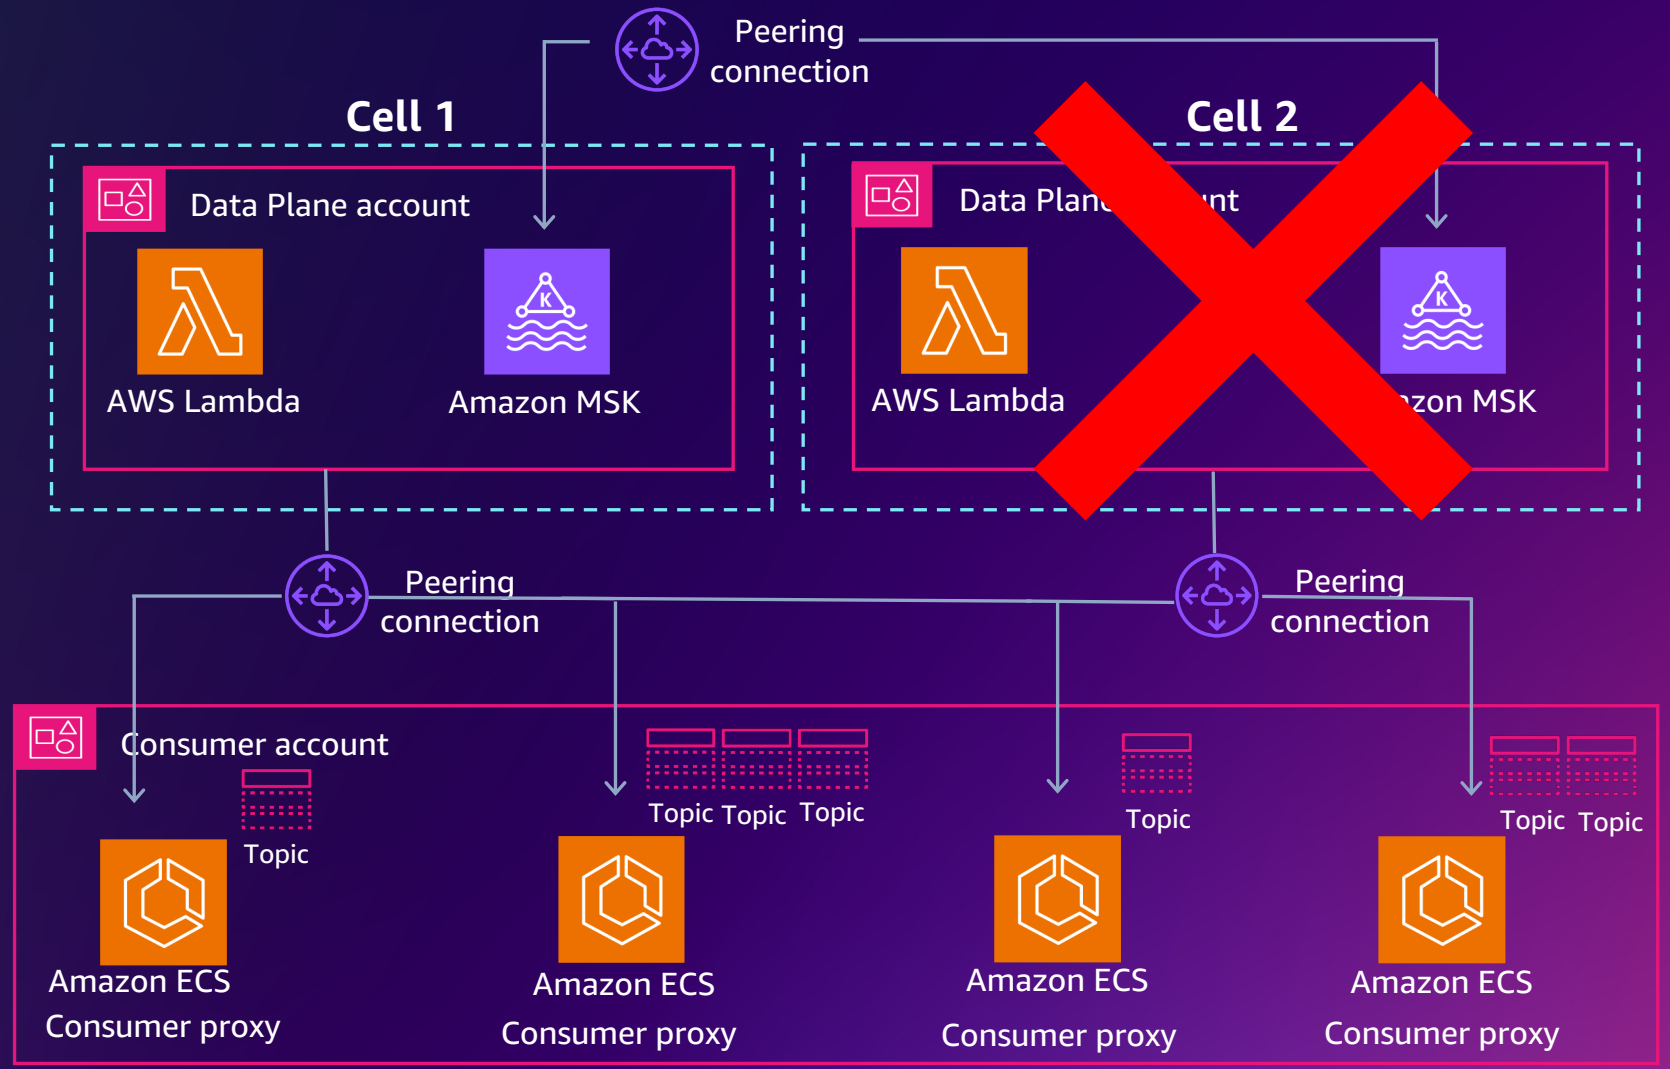

SEB – Add an additional layer of resilience

FALLBACK CELL

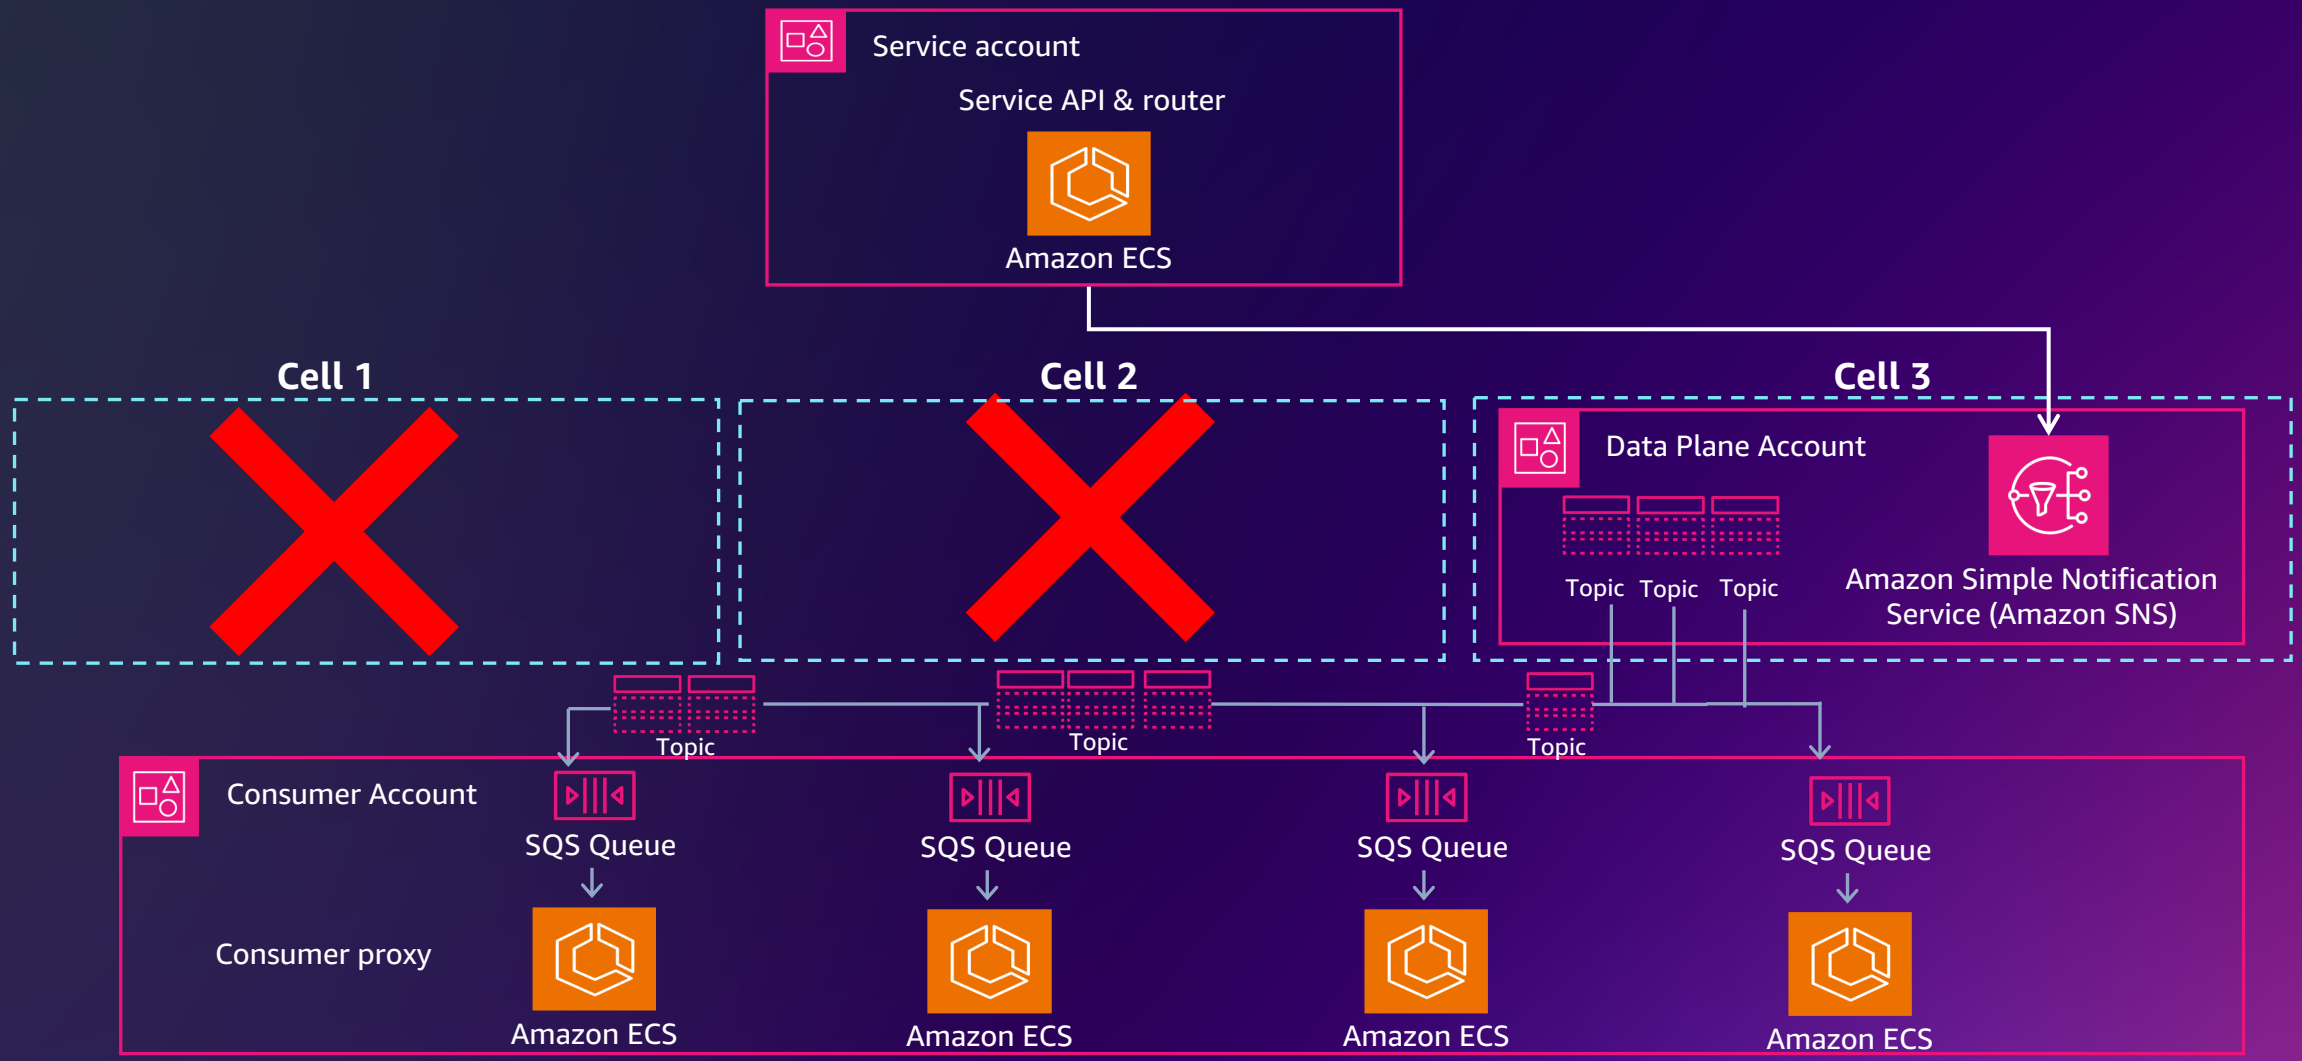

Circuit breaker pattern used by SEB

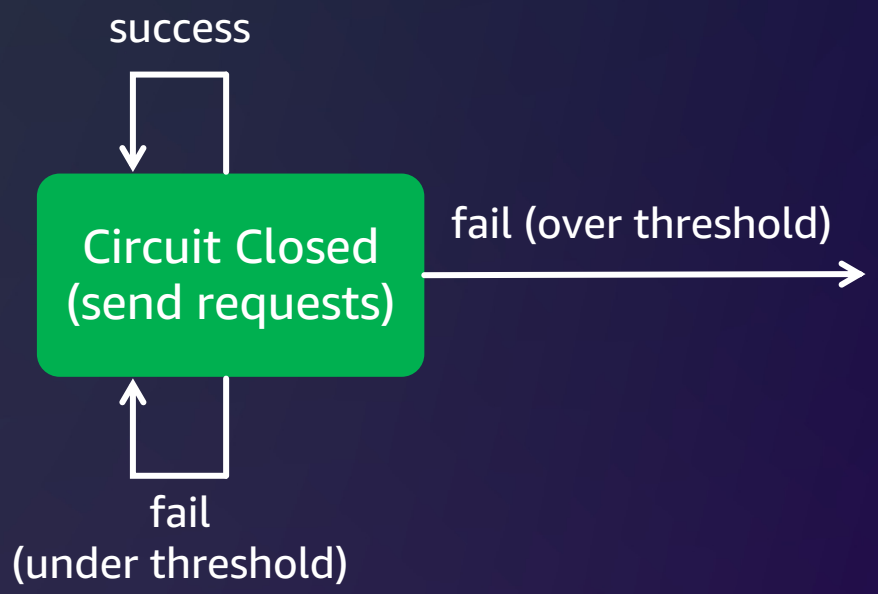

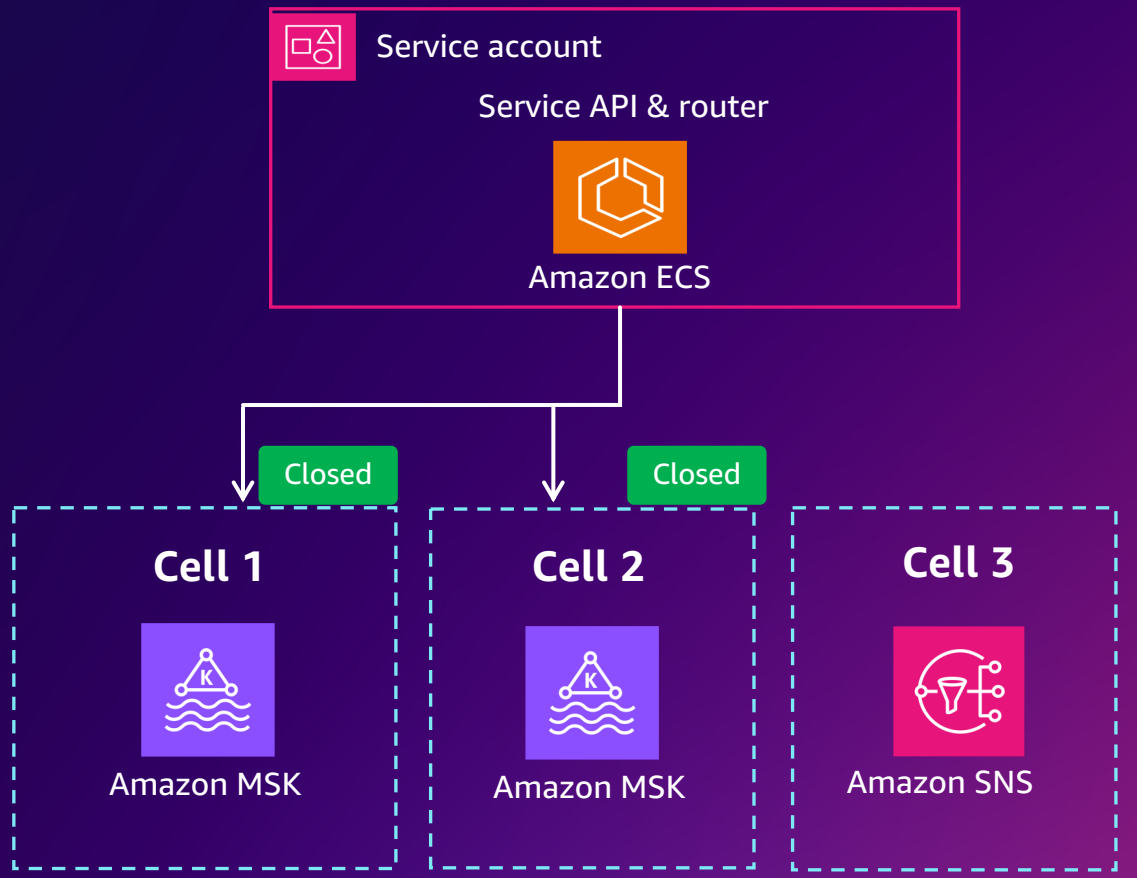

Circuit breaker pattern used by SEB

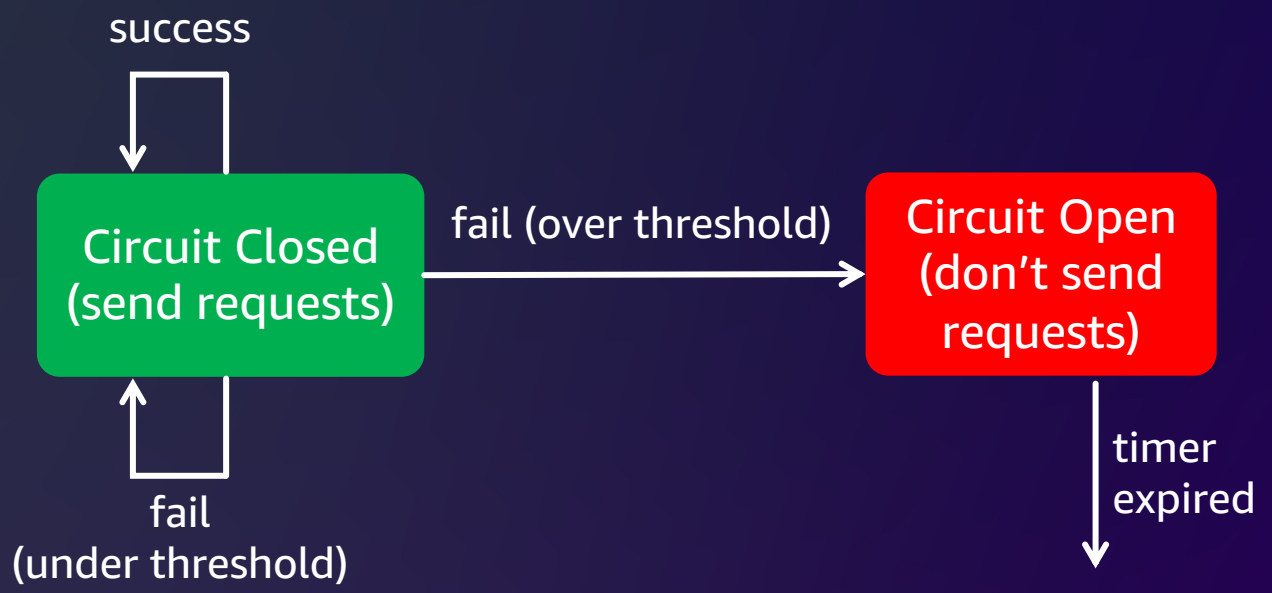

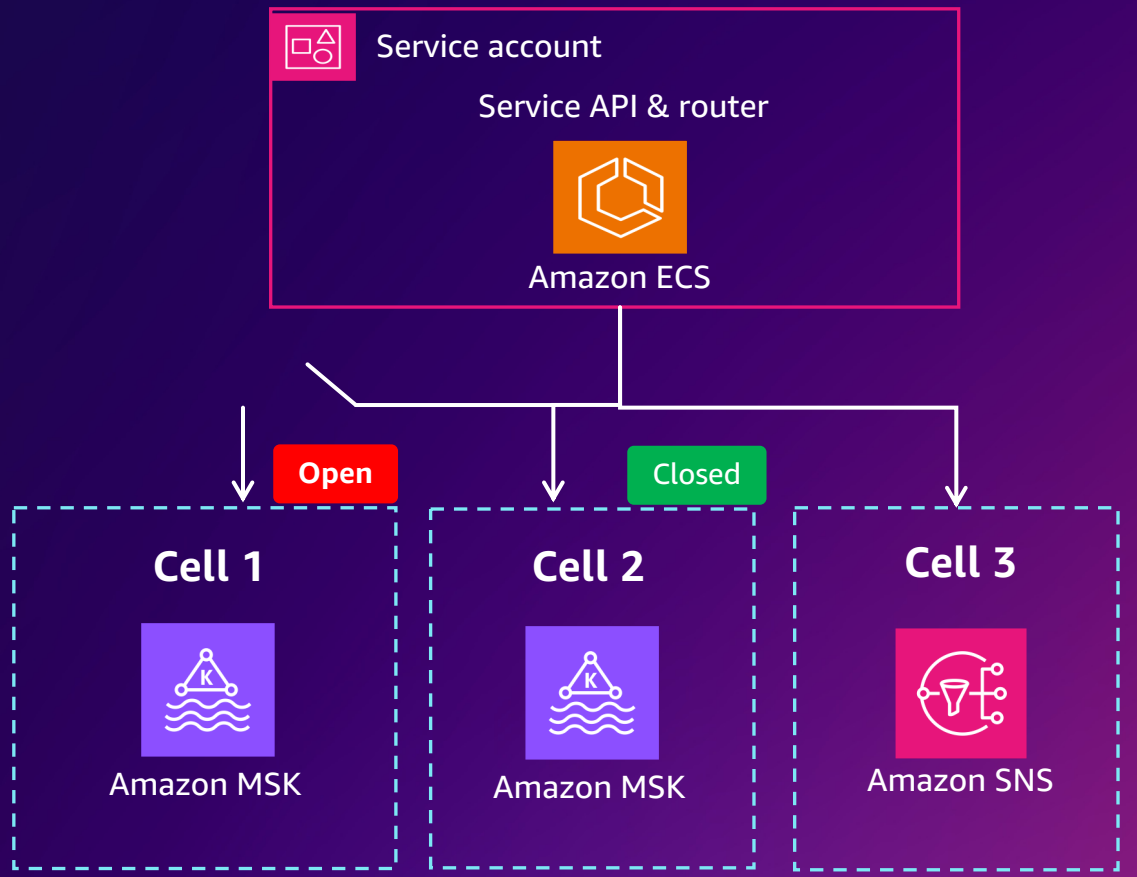

Circuit breaker pattern used by SEB

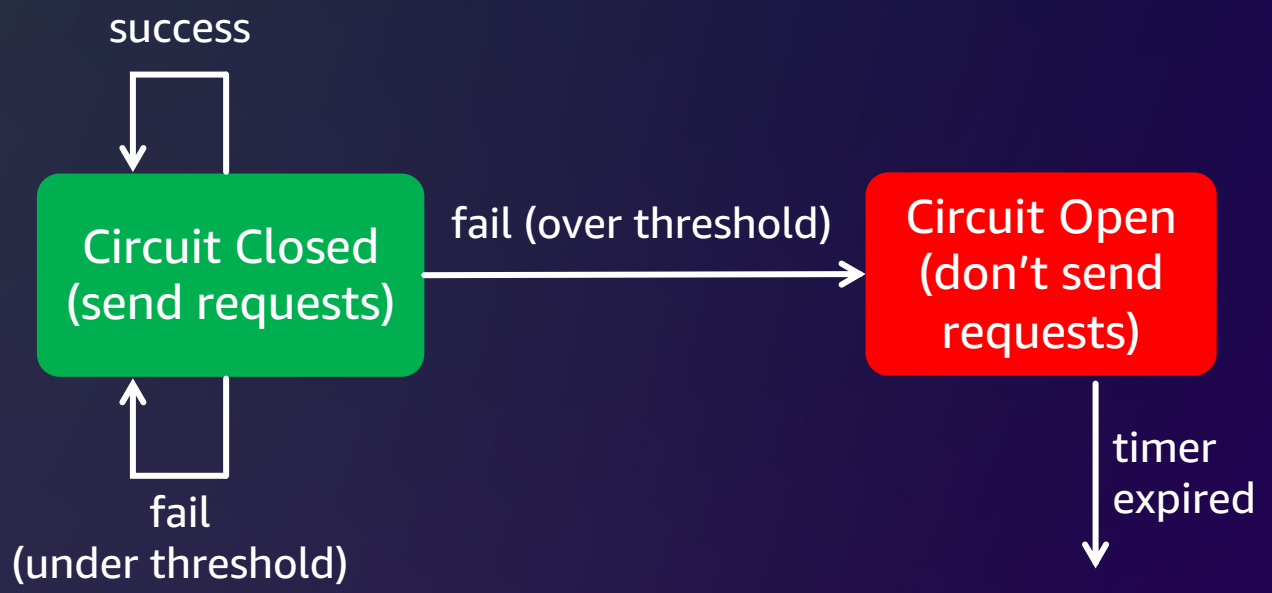

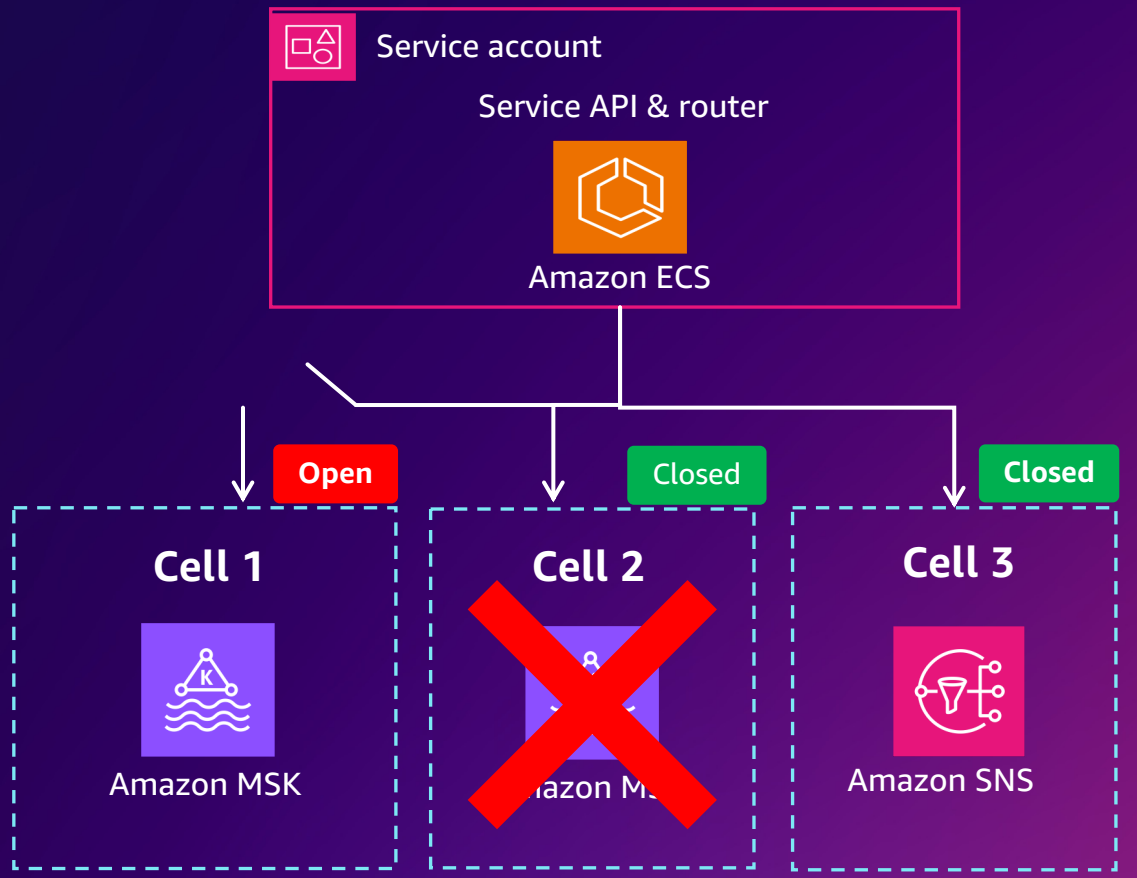

Circuit breaker pattern used by SEB

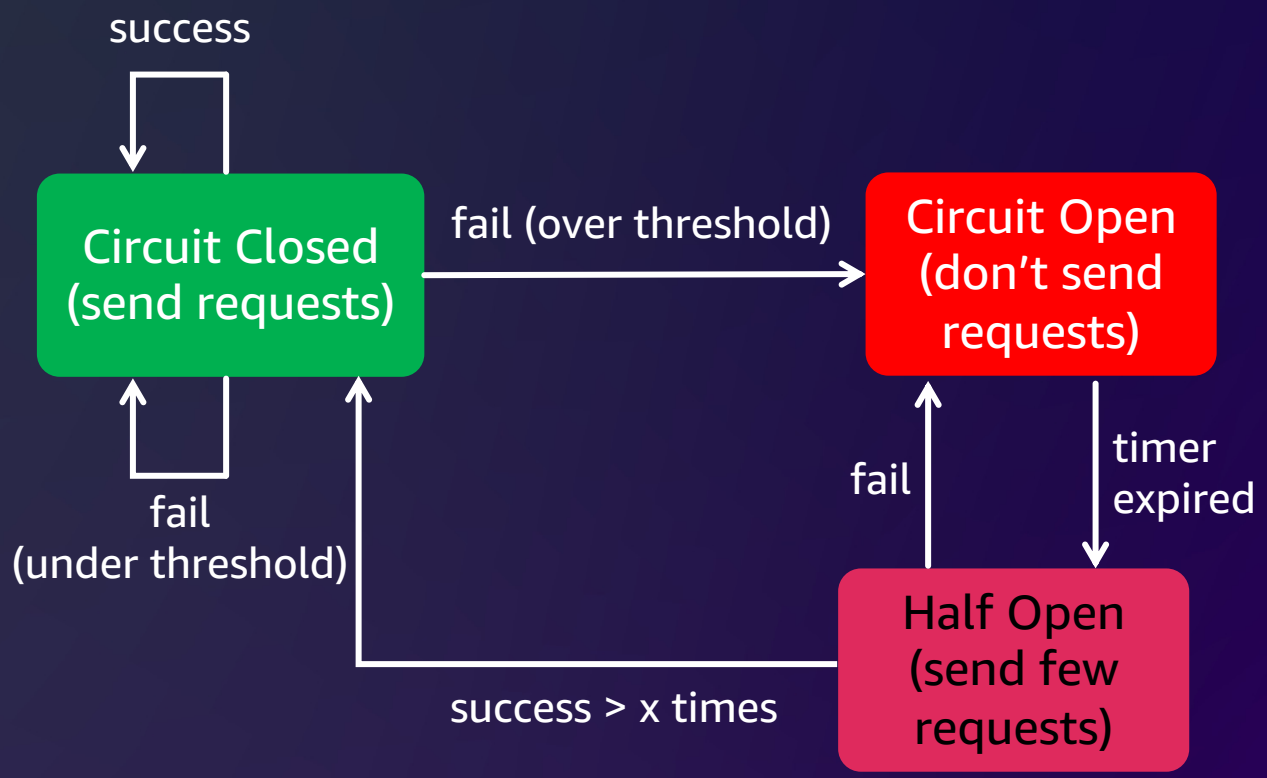

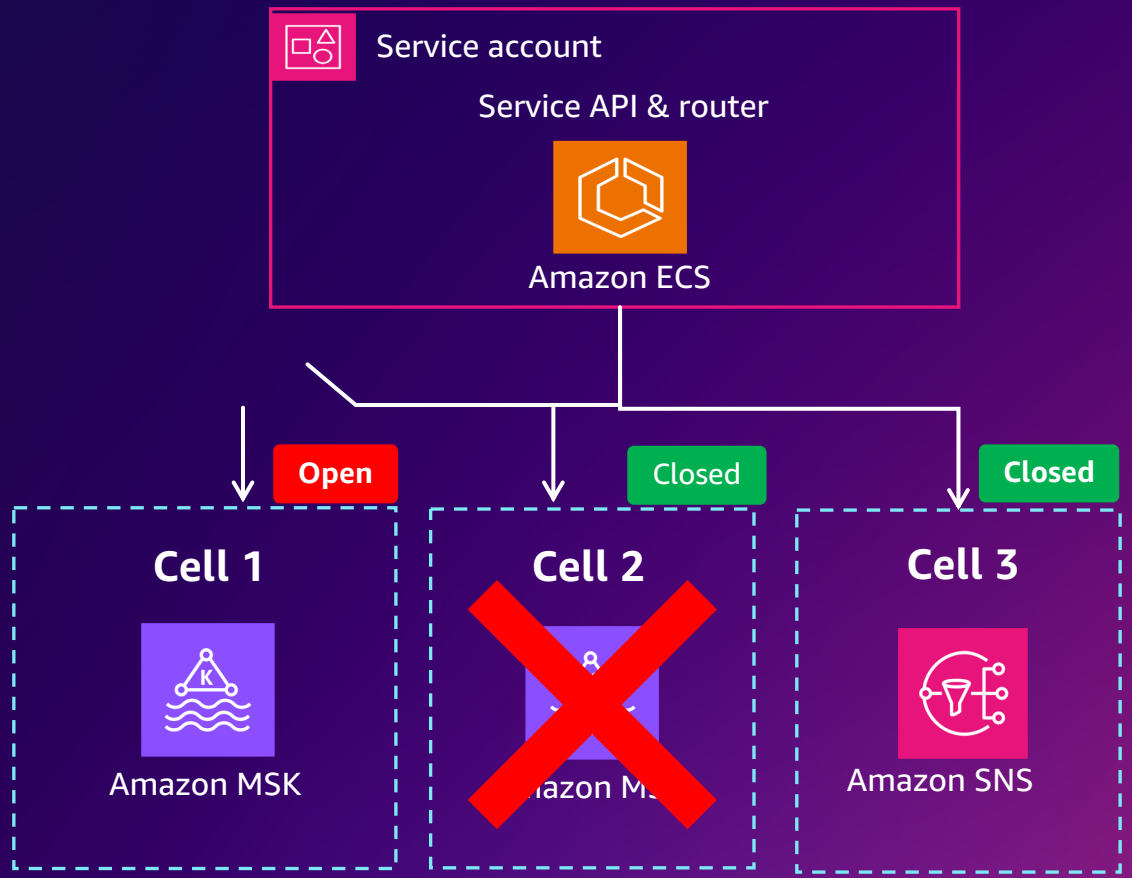

Actual availability, as measured

ACTIVE/ACTIVE STRATEGY WITH MULTI-LAYER FALLBACK

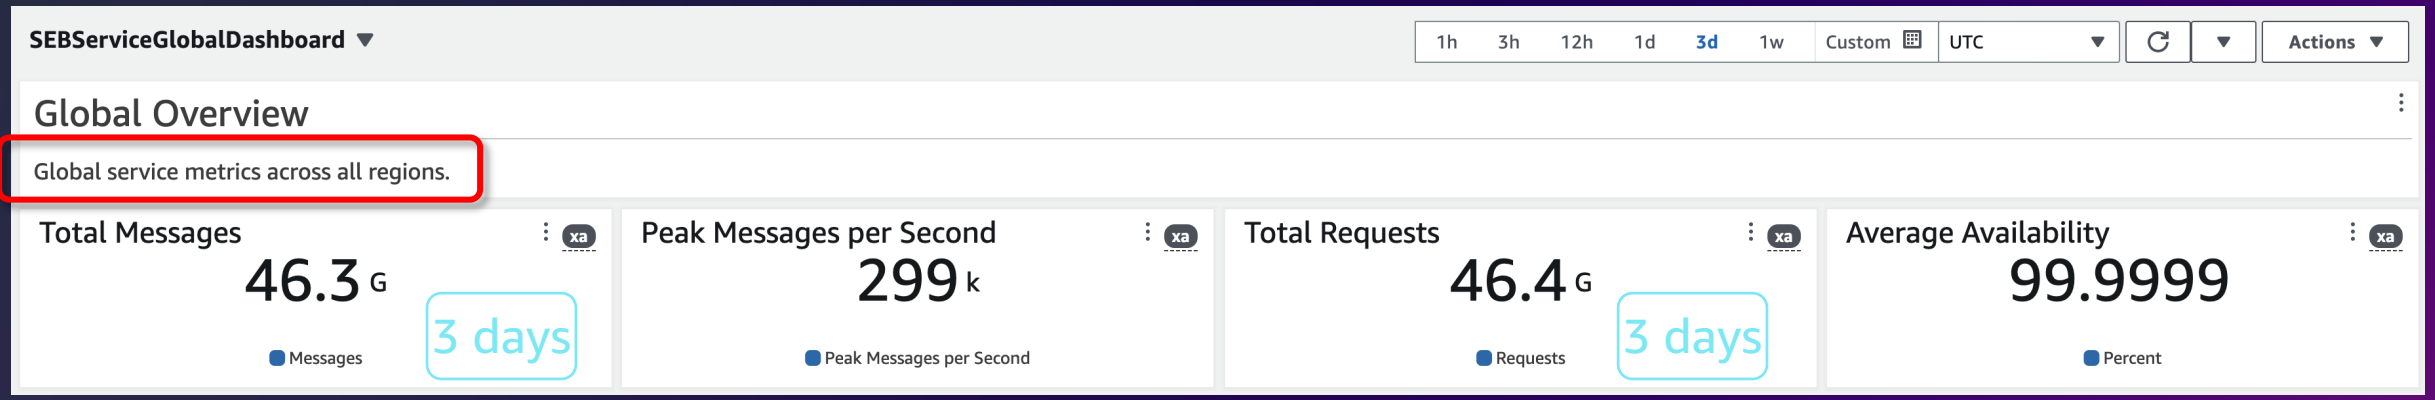

Actual availability, as measured

ACTIVE/ACTIVE STRATEGY WITH MULTI-LAYER FALLBACK

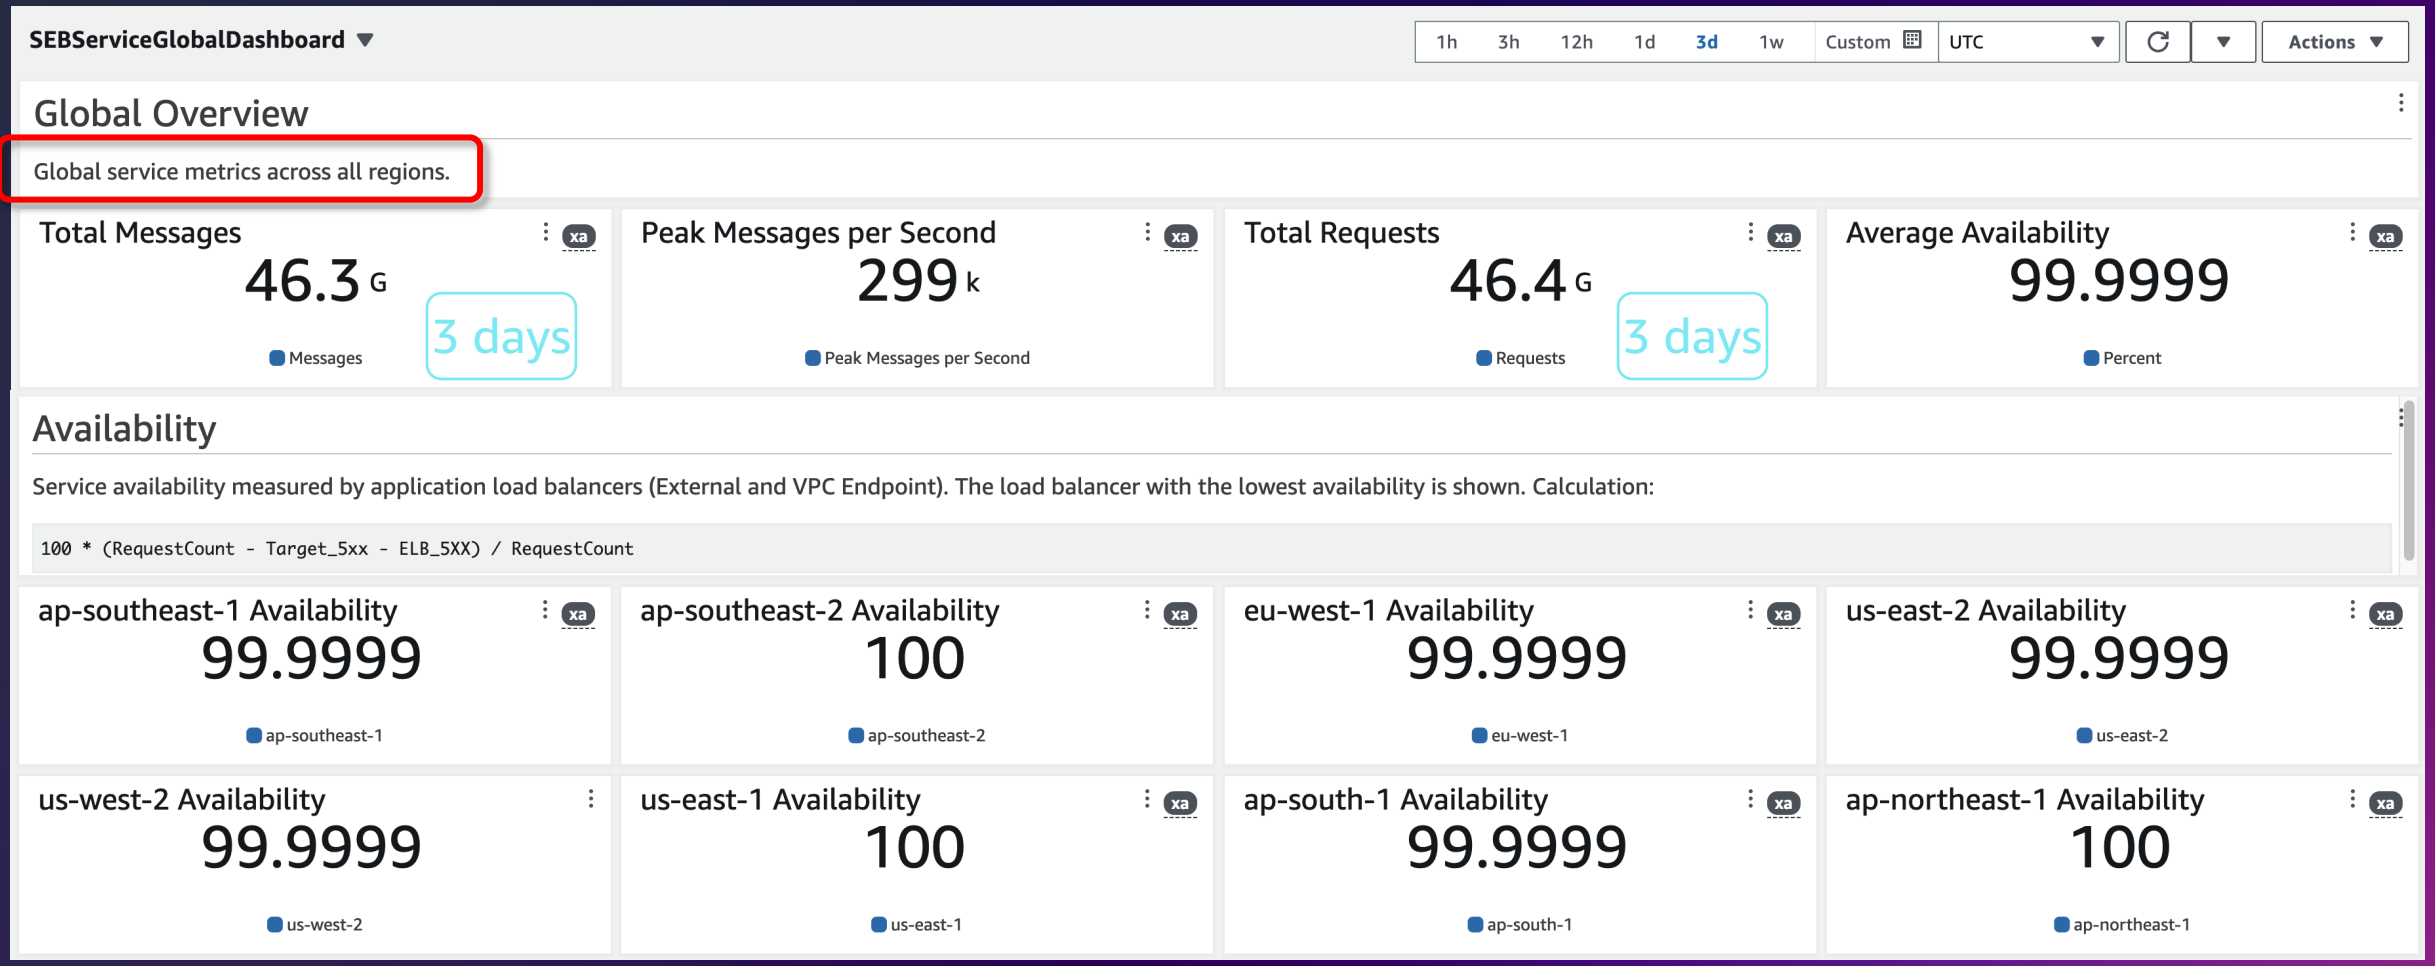

Conclusion

- Resilience at scale on the cloud is a design decision
- Real-life examples…
  - Are a great way to learn
  - Can inspire better architectures

Purpose-built AWS resilience offerings

BUILD RESILIENT, HIGHLY AVAILABLE APPLICATIONS IN THE AWS CLOUD

AWS Resilience Hub

Analyse the components of your application to uncover potential resilience weaknesses

AWS Fault Injection Service

Improve application performance, observability, and resilience through controlled fault injection experiments

aws.com/fis

AWS Elastic Disaster Recovery

Minimise downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications

AWS Backup

Protect data at scale using this cost-effective, fully managed, policy-based service

Amazon Route53 Application Recovery Controller

Automate management and coordination of recovery for your applications across AWS AZs or Regions

Resilience lifecycle framework

Strategies, services & mechanisms you can weave into your existing processes

Thank you!

Dragoș Mădărășan

Please complete the session survey in the Events App

App Store

Google Play Store