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

NET403

Planet-scale networking: How AWS powers the world's largest networks

Stephen Callaghan
Senior Principal Network Engineer
Amazon Infrastructure

Jorge Vasquez
Senior Principal Software Engineer
Amazon Infrastructure

Agenda

- Role of the AWS network
- Design goals and principles
- Deep ownership
- What’s next?

Continuing the series

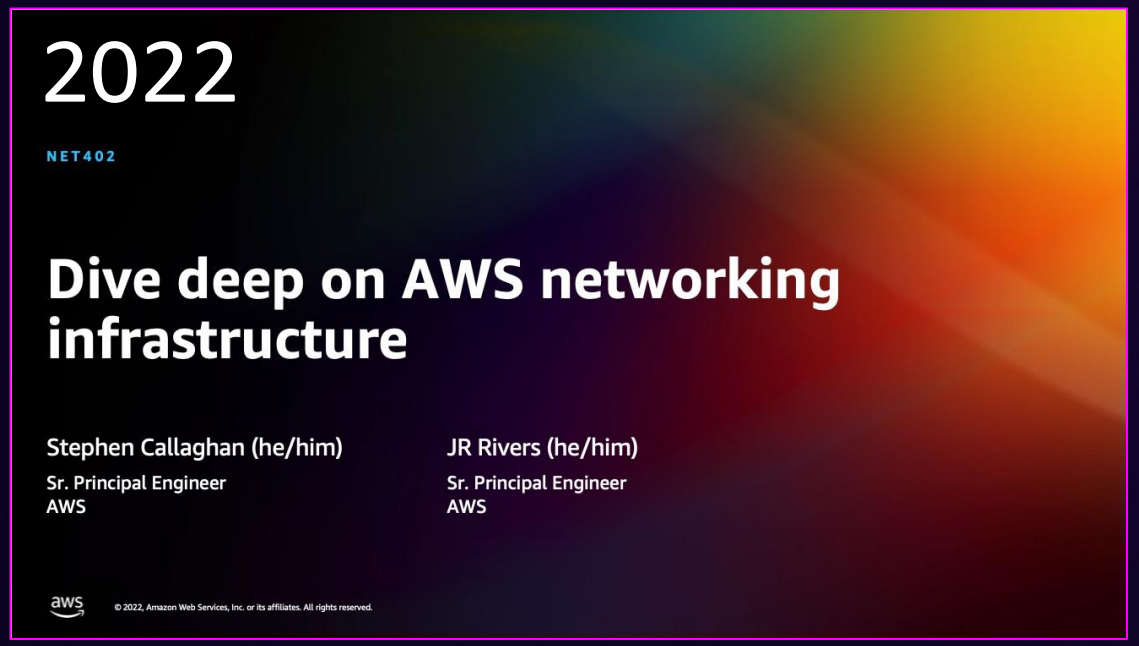

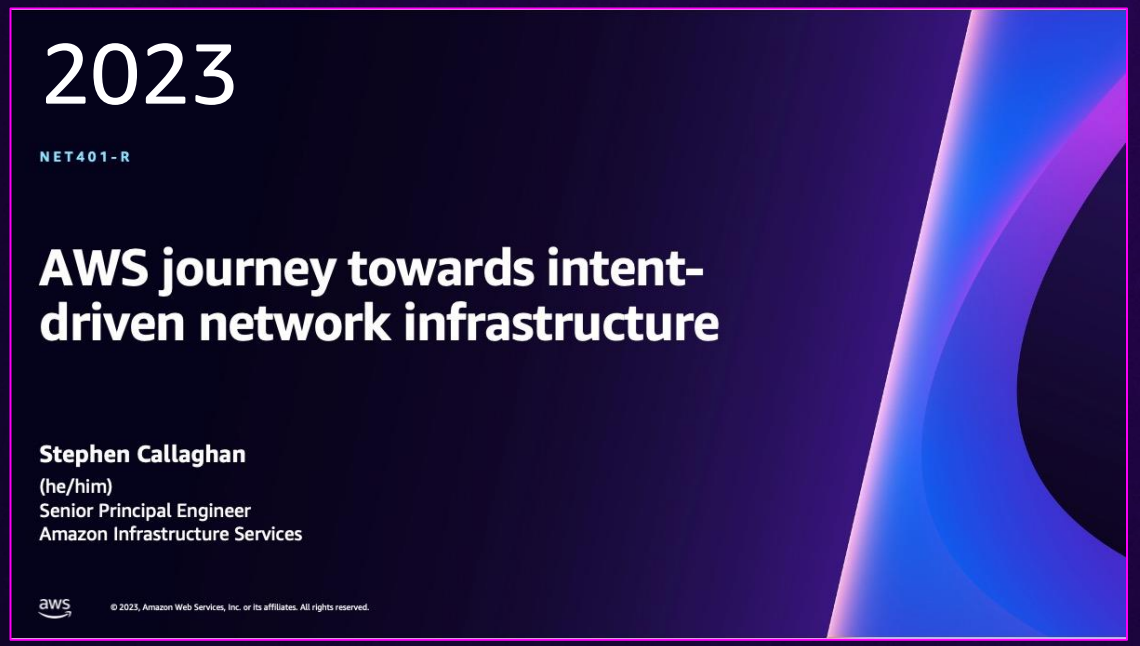

AWS networking

Infrastructure networking	Amazon EC2 networking	Edge networking
Routers/Switches	Amazon Virtual Private Cloud (Amazon VPC)	Amazon Route 53
Copper/Optical cables	Elastic network interface	AWS Global Accelerator
Datacenters	AWS NAT gateway	Amazon CloudFront
Inter-Region backbone	Elastic Fabric Adapter (EFA)	AWS Direct Connect
Internet peering/transit	Placement groups	AWS Cloud WAN

Planet scale networking

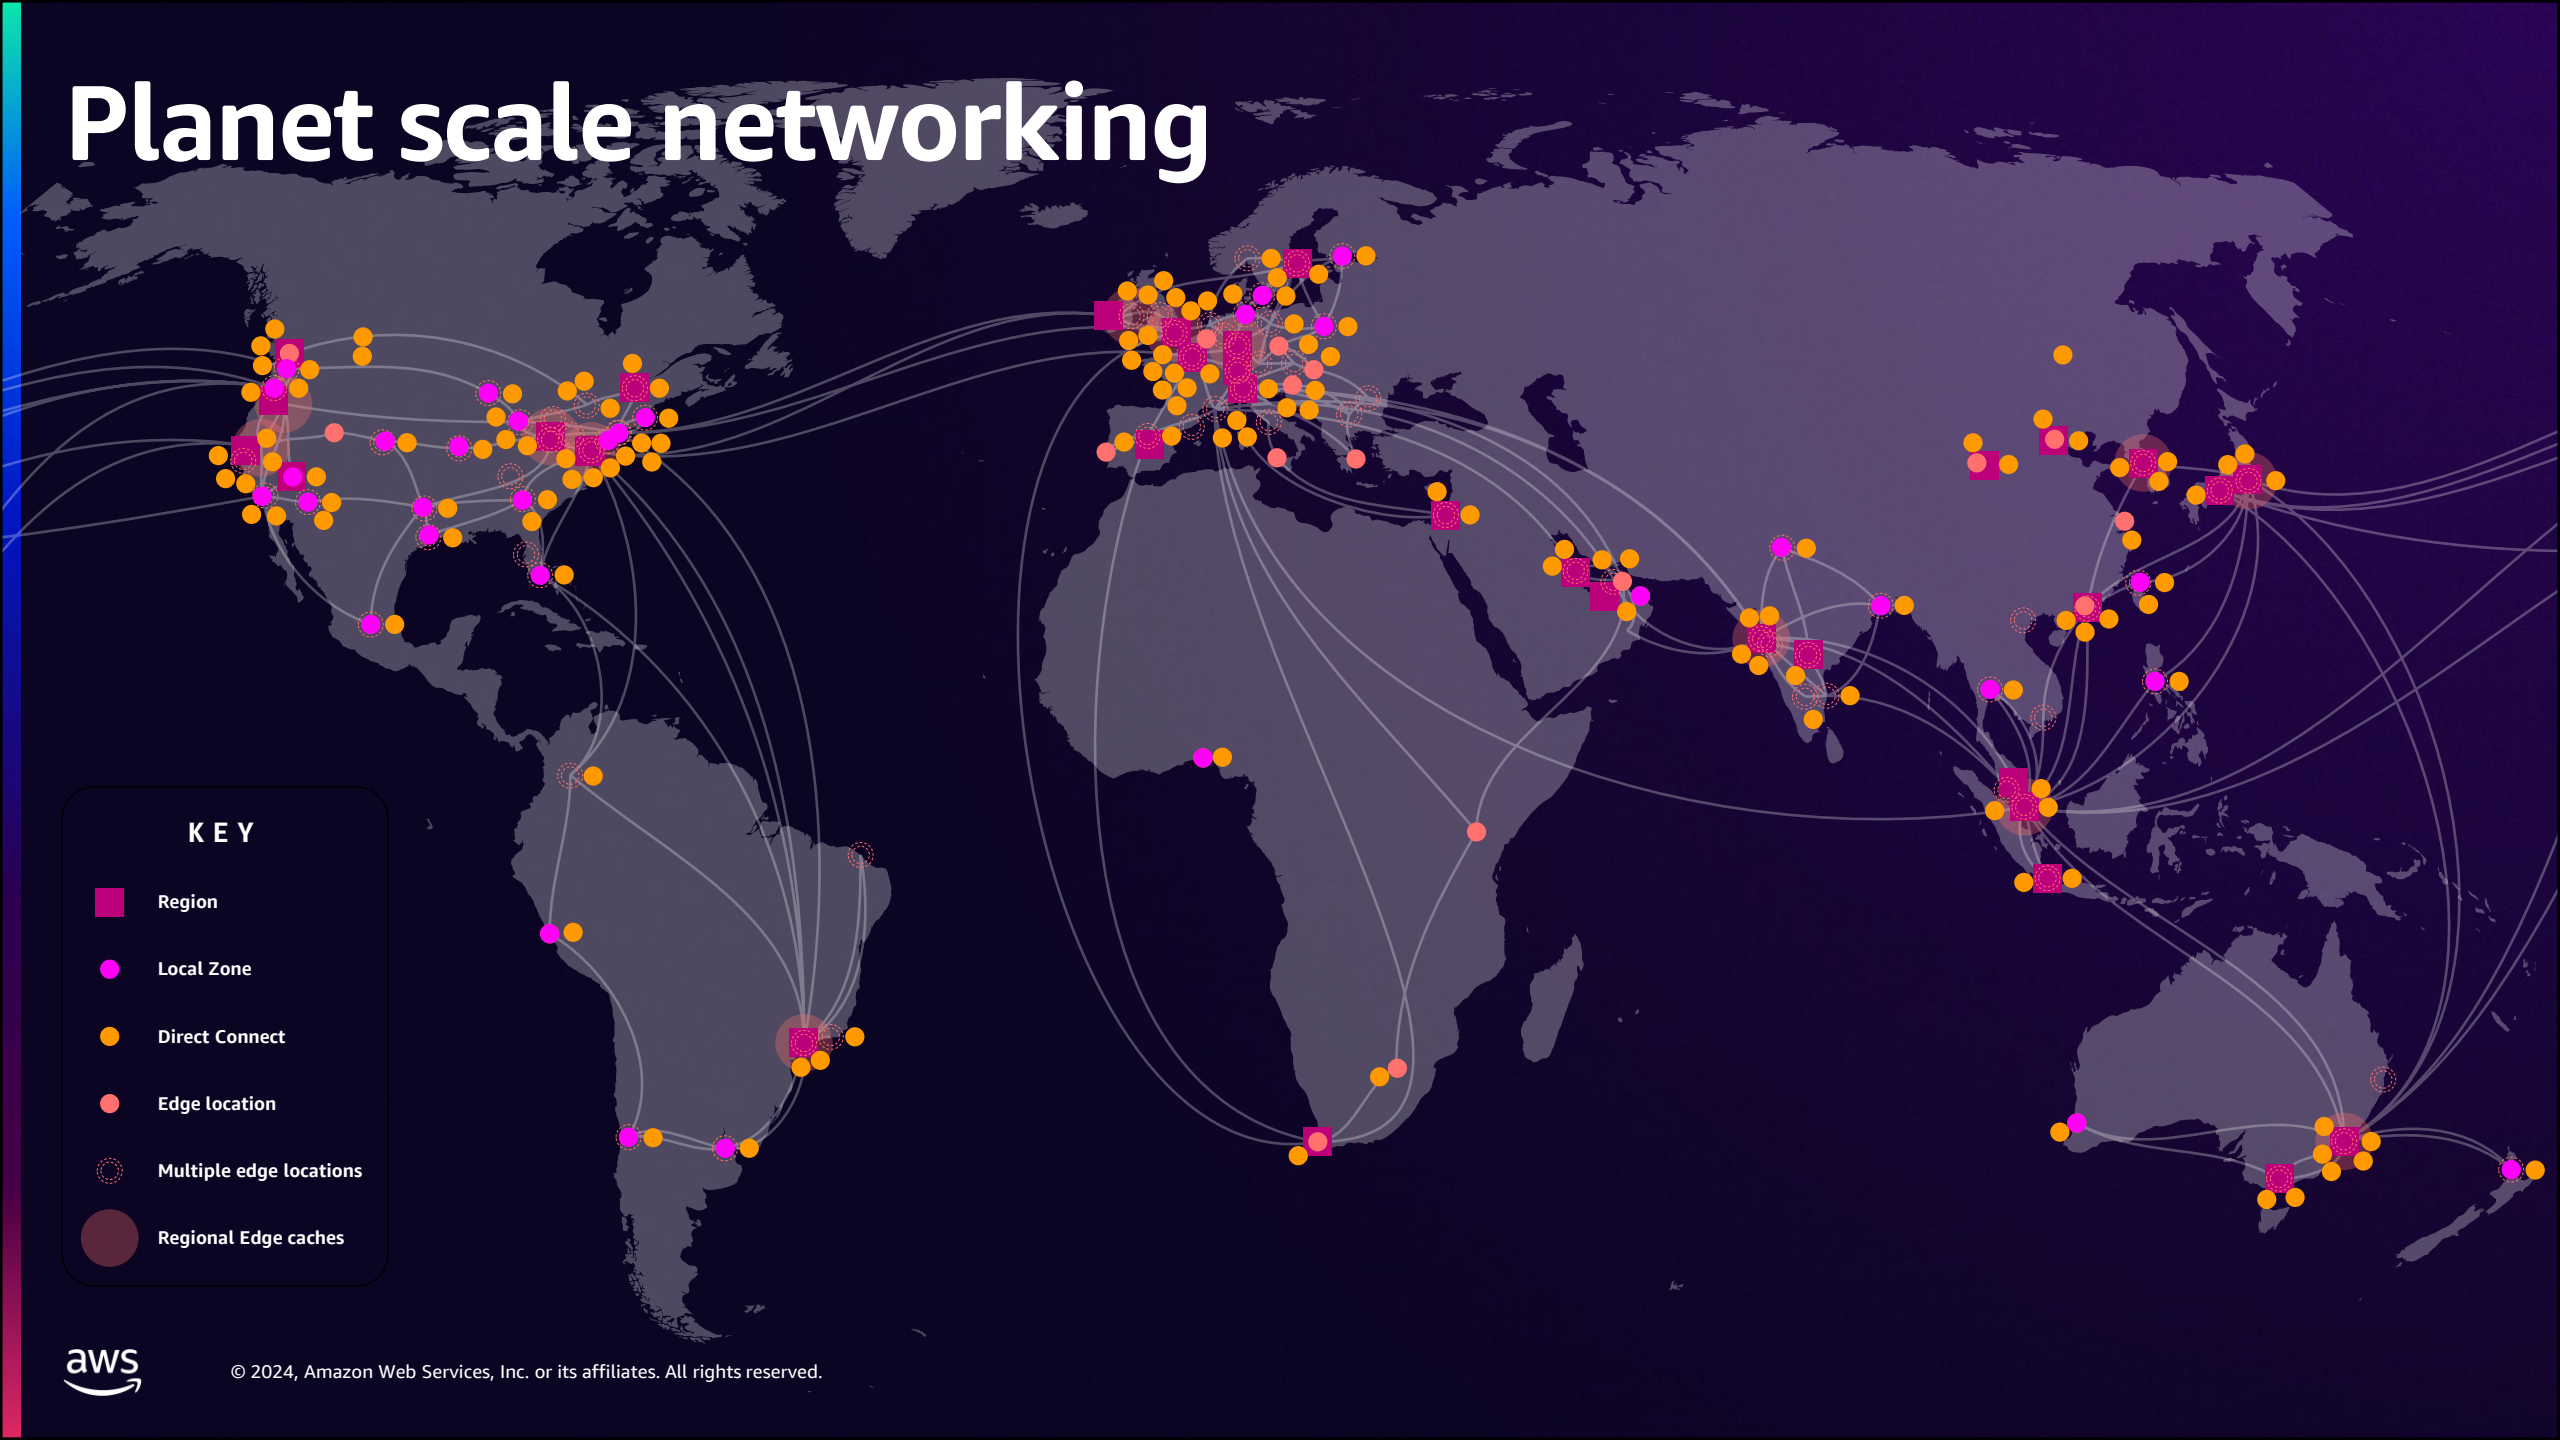

AWS custom network hardware

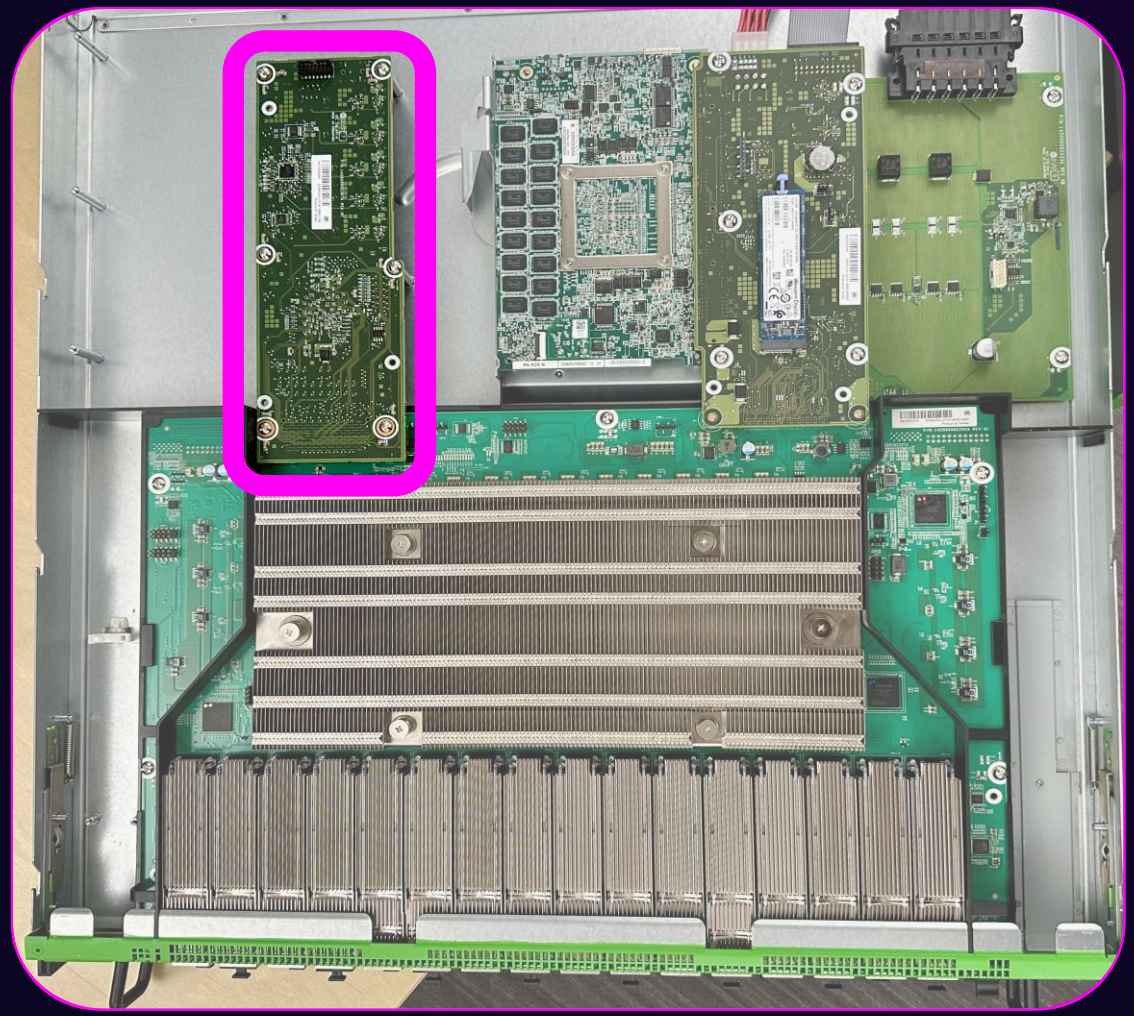

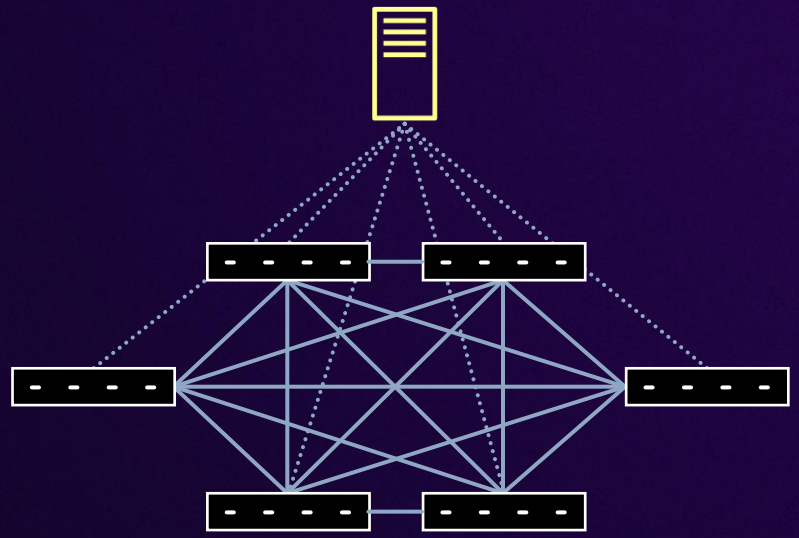

Statically stable

Deterministic

Highly visible

Low scope of impact

Hybrid control plane

Role of the AWS network

Be so performant and reliable that we're out of the way of your workloads

Design goals

Flexible

amazon Search Who We Are What We Do Our Workplace Our Impact Our Planet Follow Us Subscribe EN

News / AWS

6 min

June 26, 2024

4 ways AWS is engineering infrastructure to power generative AI

Written by Prasad Kalyanaraman, VP of Infrastructure Services at AWS

Flexible

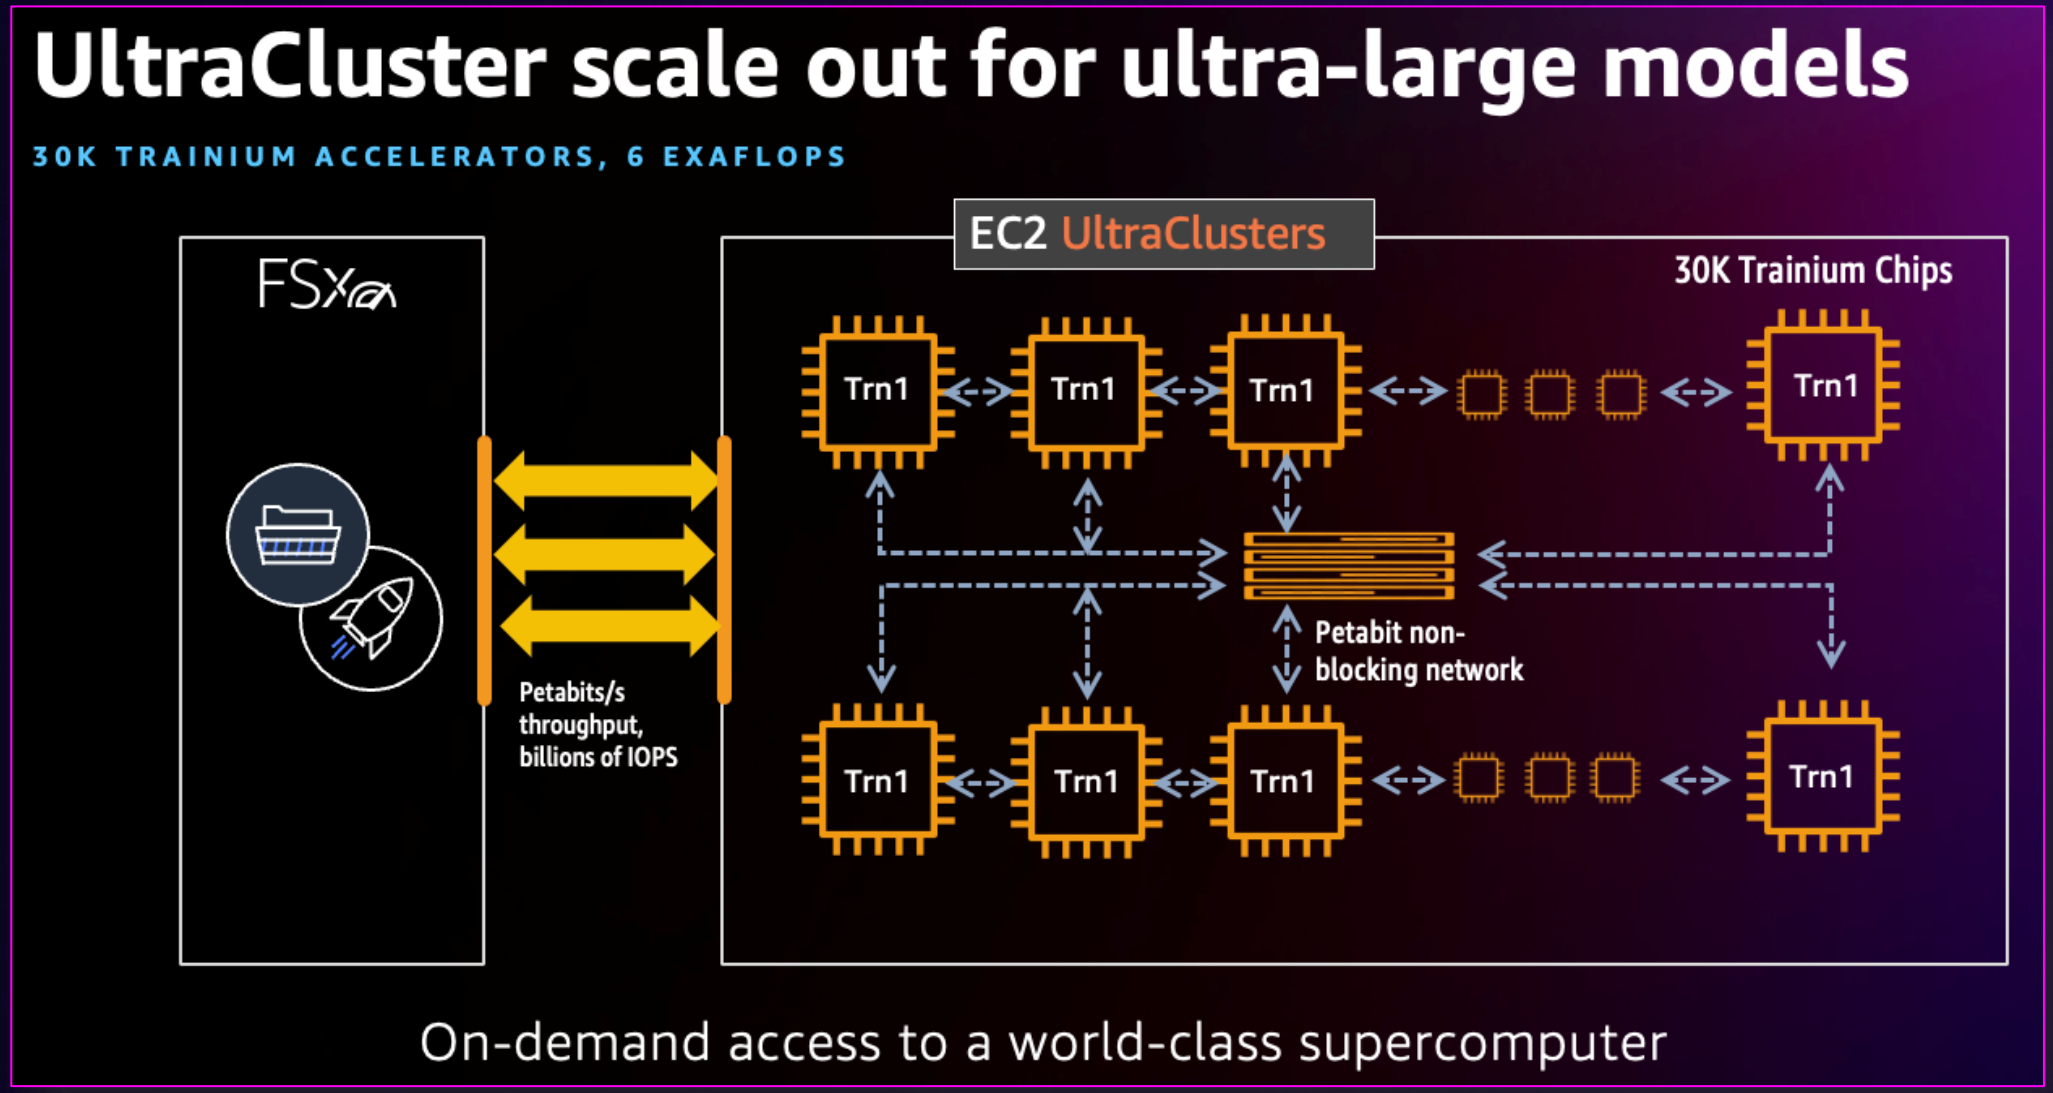

AI training and inference

amazon | ANTHROP\C

COMPANY NEWS

Amazon and Anthropic deepen their shared commitment to advancing generative AI

AWS

AWS can help reduce the carbon footprint of AI workloads by up to 99%. Here's how.

aws | NVIDIA

AWS

AWS and NVIDIA extend their collaboration to advance generative AI

AI training and inference

AWS

Amazon Bedrock introduces Anthropic's Claude 3.5 Sonnet to customers, their most powerful AI model to date

AWS

Llama 3.2 models from Meta are now available on AWS, offering more options for building generative AI applications

AWS

Amazon Bedrock customers have more choice in AI models with Mistral Large now available

Local zones

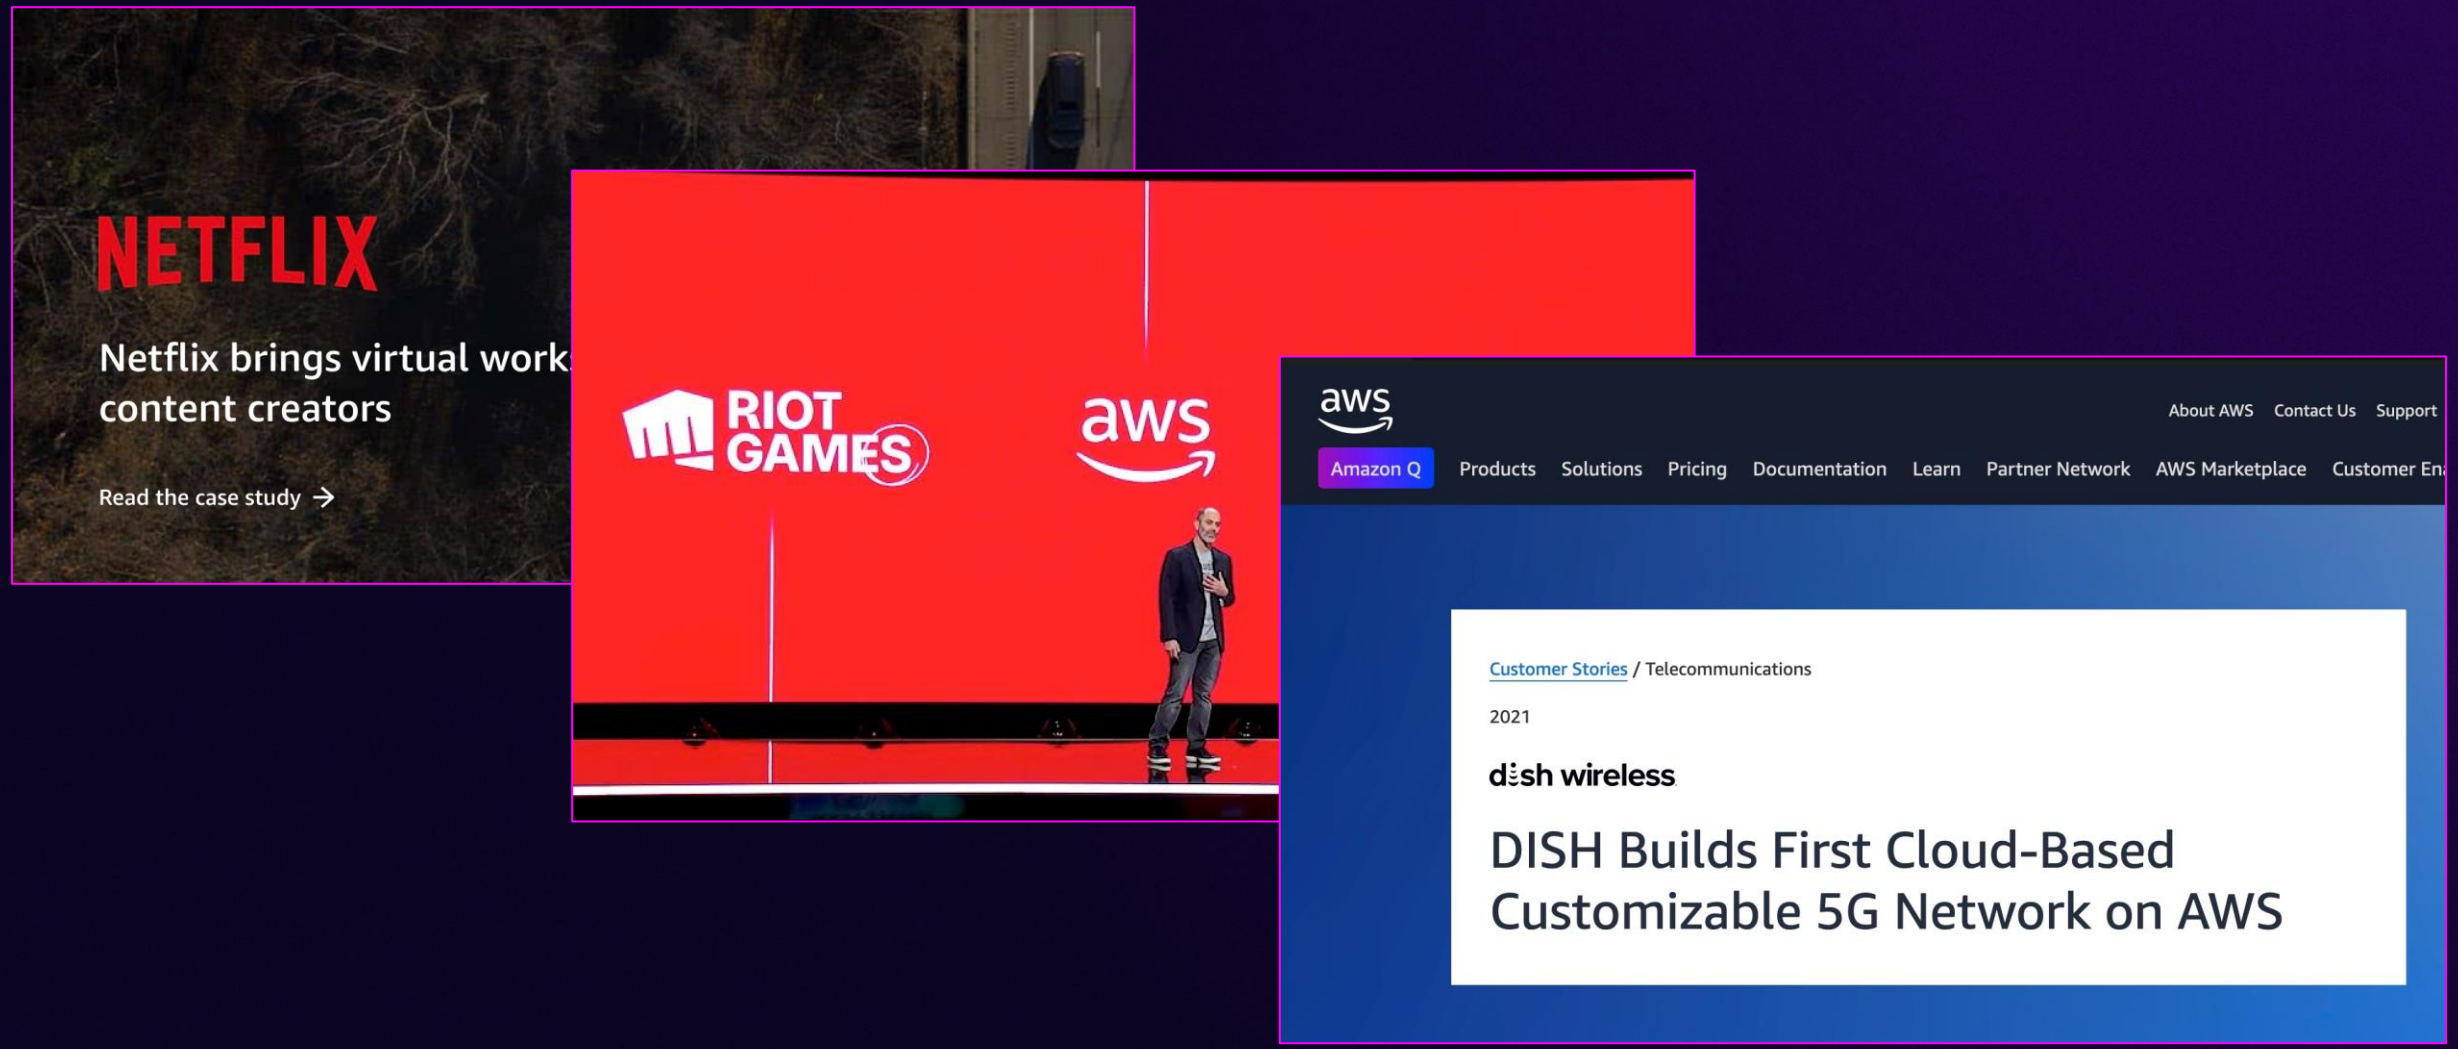

Video streaming

ENTERTAINMENT

'Thursday Night Football' behind-the-scenes: How Prime's NFL coverage comes to life

AWS

How AWS is using AI to bring Formula 1 fans closer to the race

ENTERTAINMENT

Prime Video reaches a landmark 11-year streaming deal with the NBA and the WNBA

“The stack”

Subset of design goals

- Security
- Availability
- Features

Security

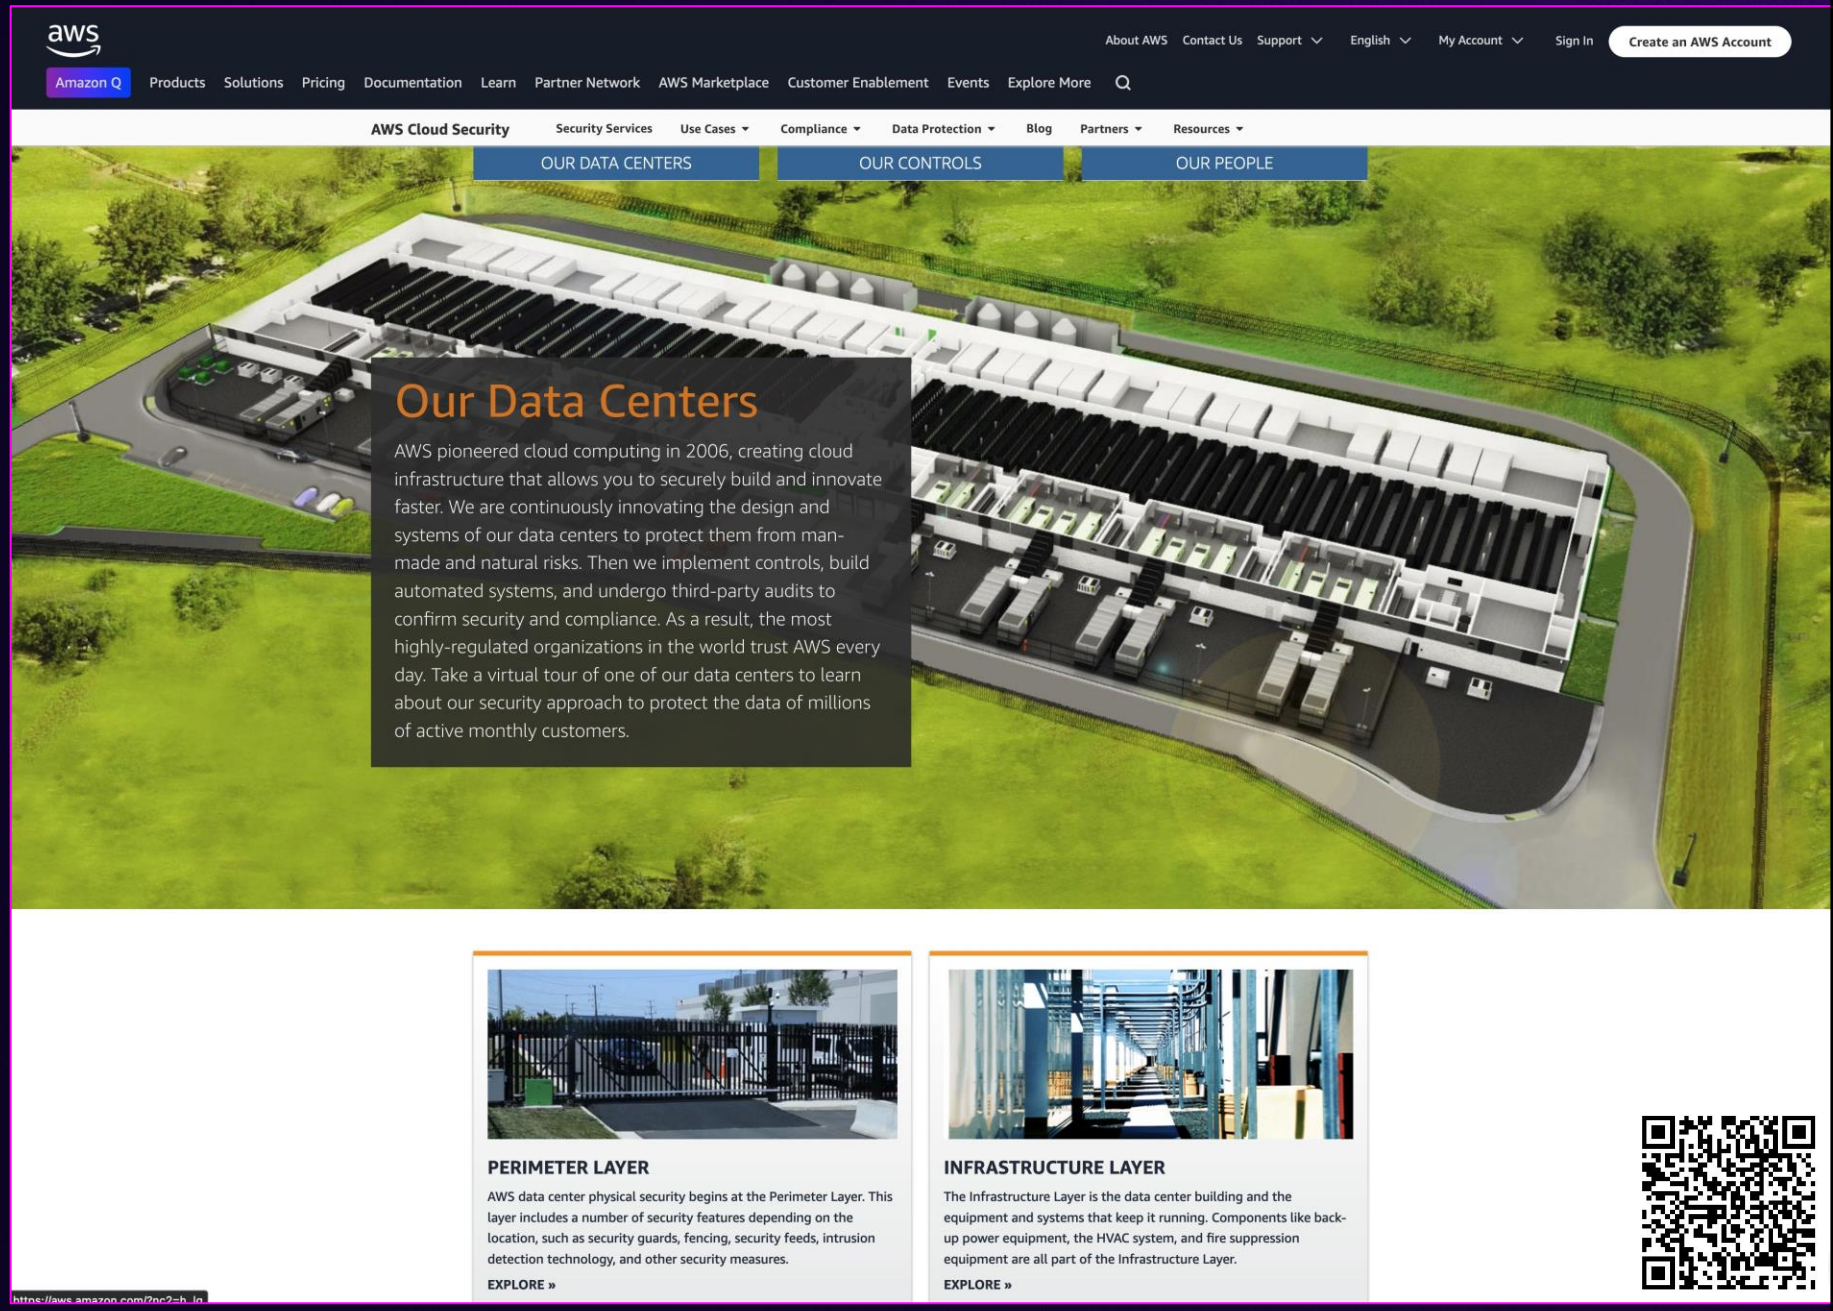

Internet security

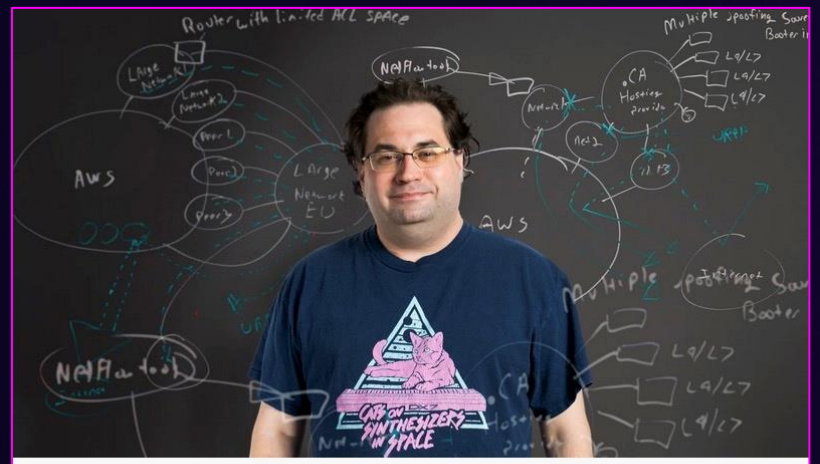

AWS

How one Amazon engineer is making the internet a safer place for all of us (and cats, too)

AWS

Amazon helps the US Department of Justice thwart international cybercriminal group Anonymous Sudan

Availability

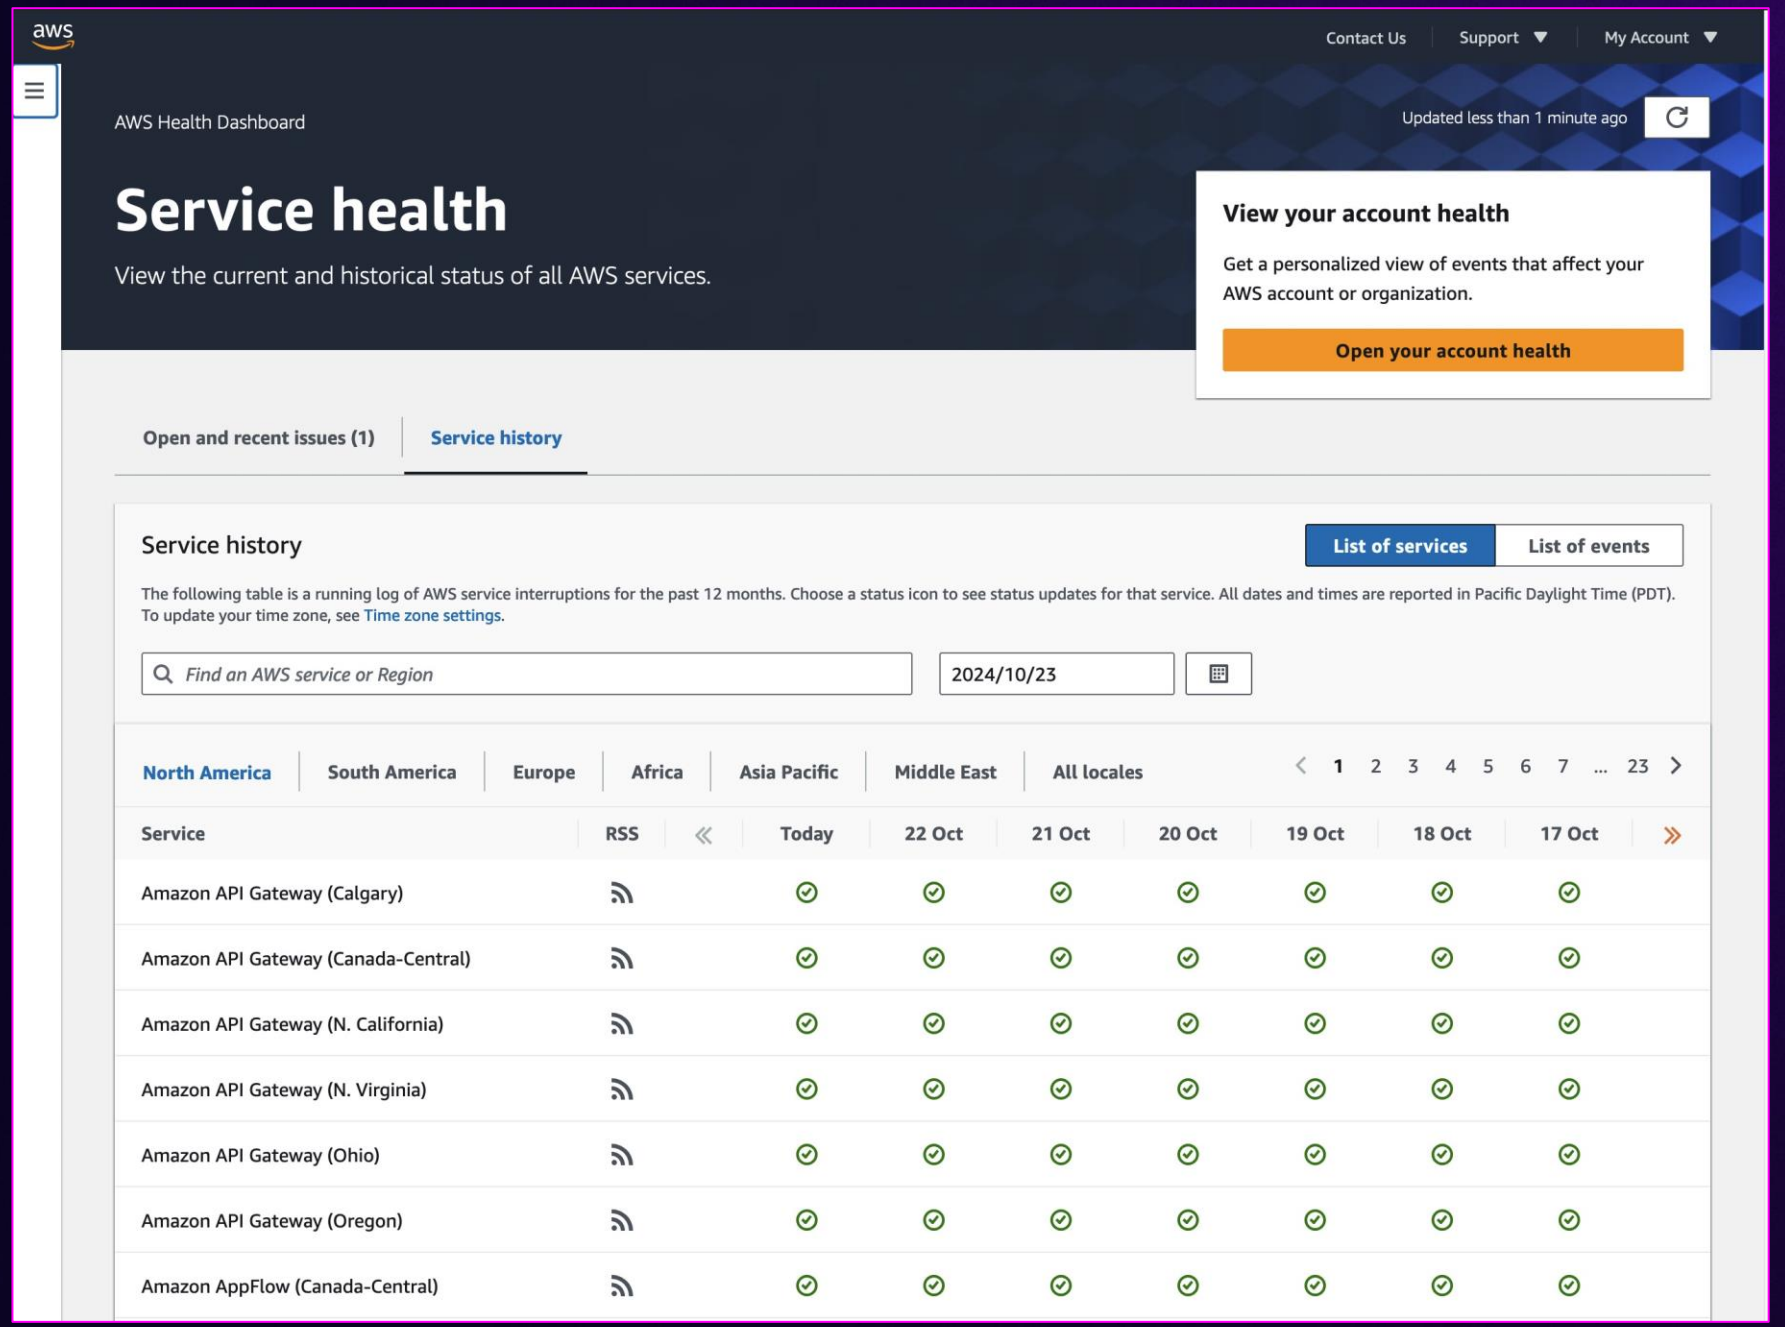

Consistency

SMOKEPING

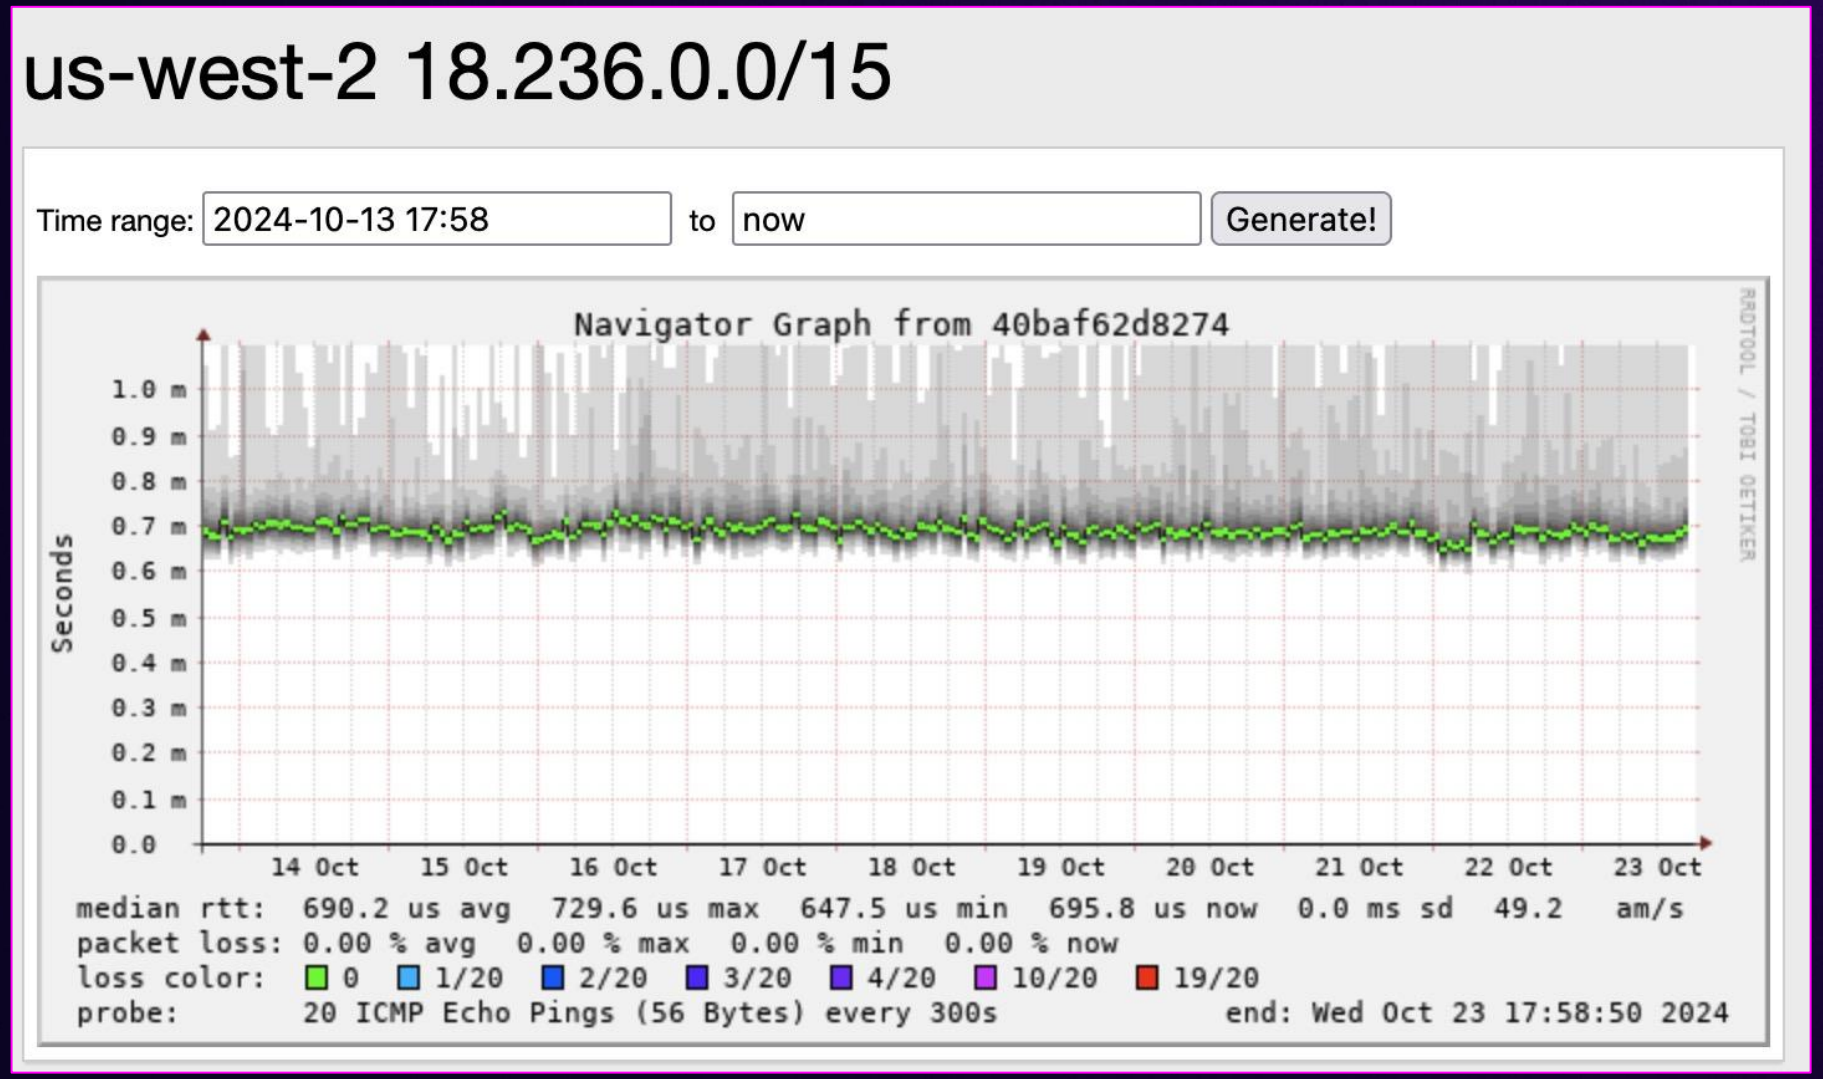

Resiliency

- Plan for failure
- Create multiple options
- Choose the best place

Resiliency example

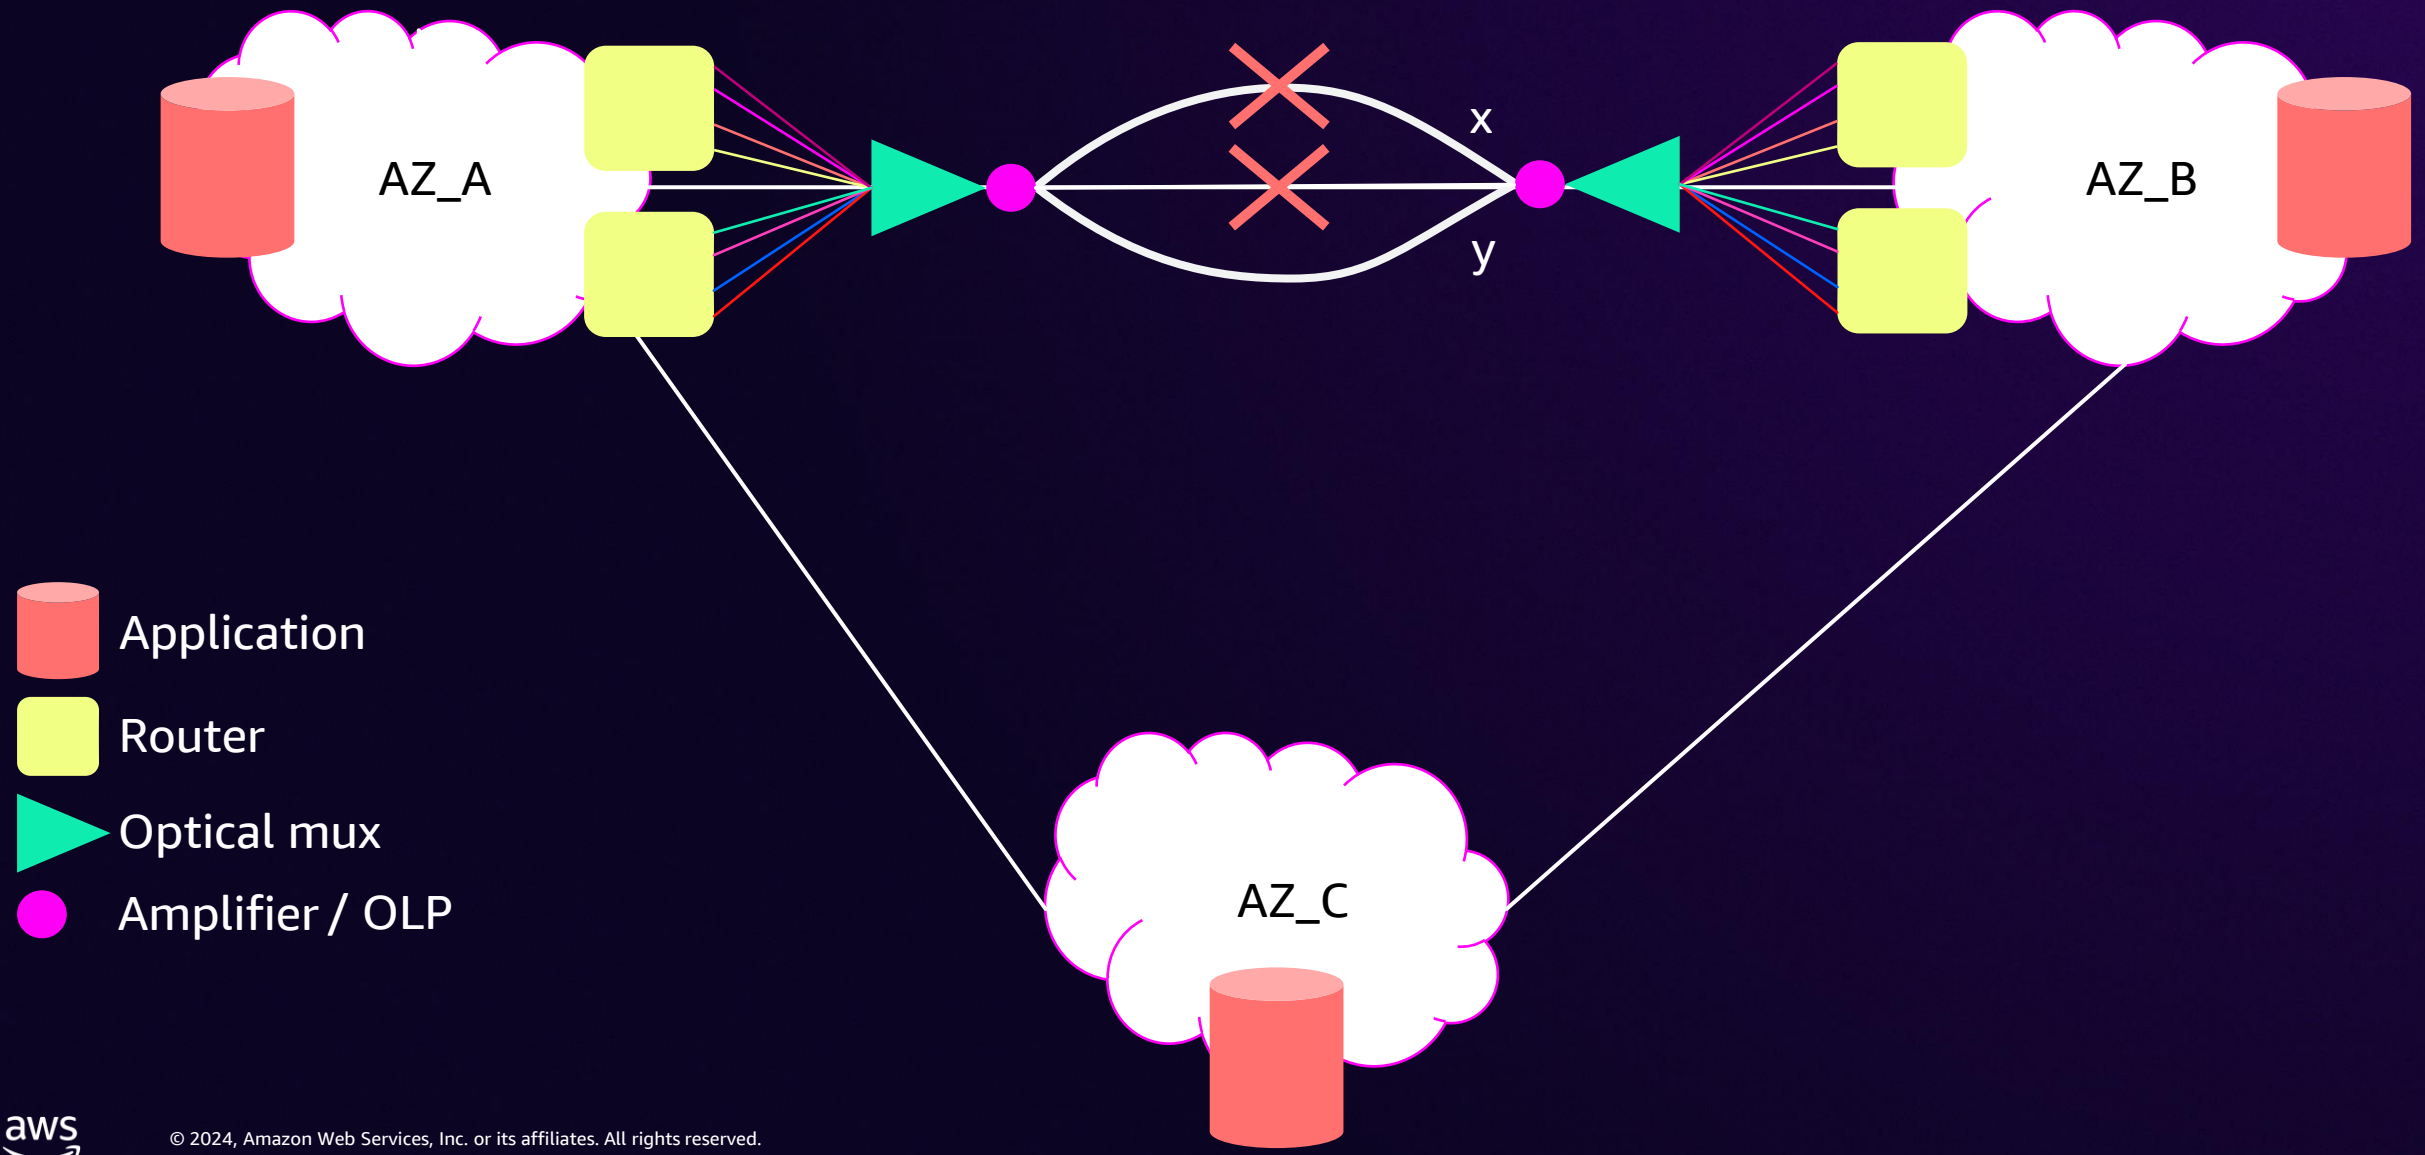

Resiliency

Robustness

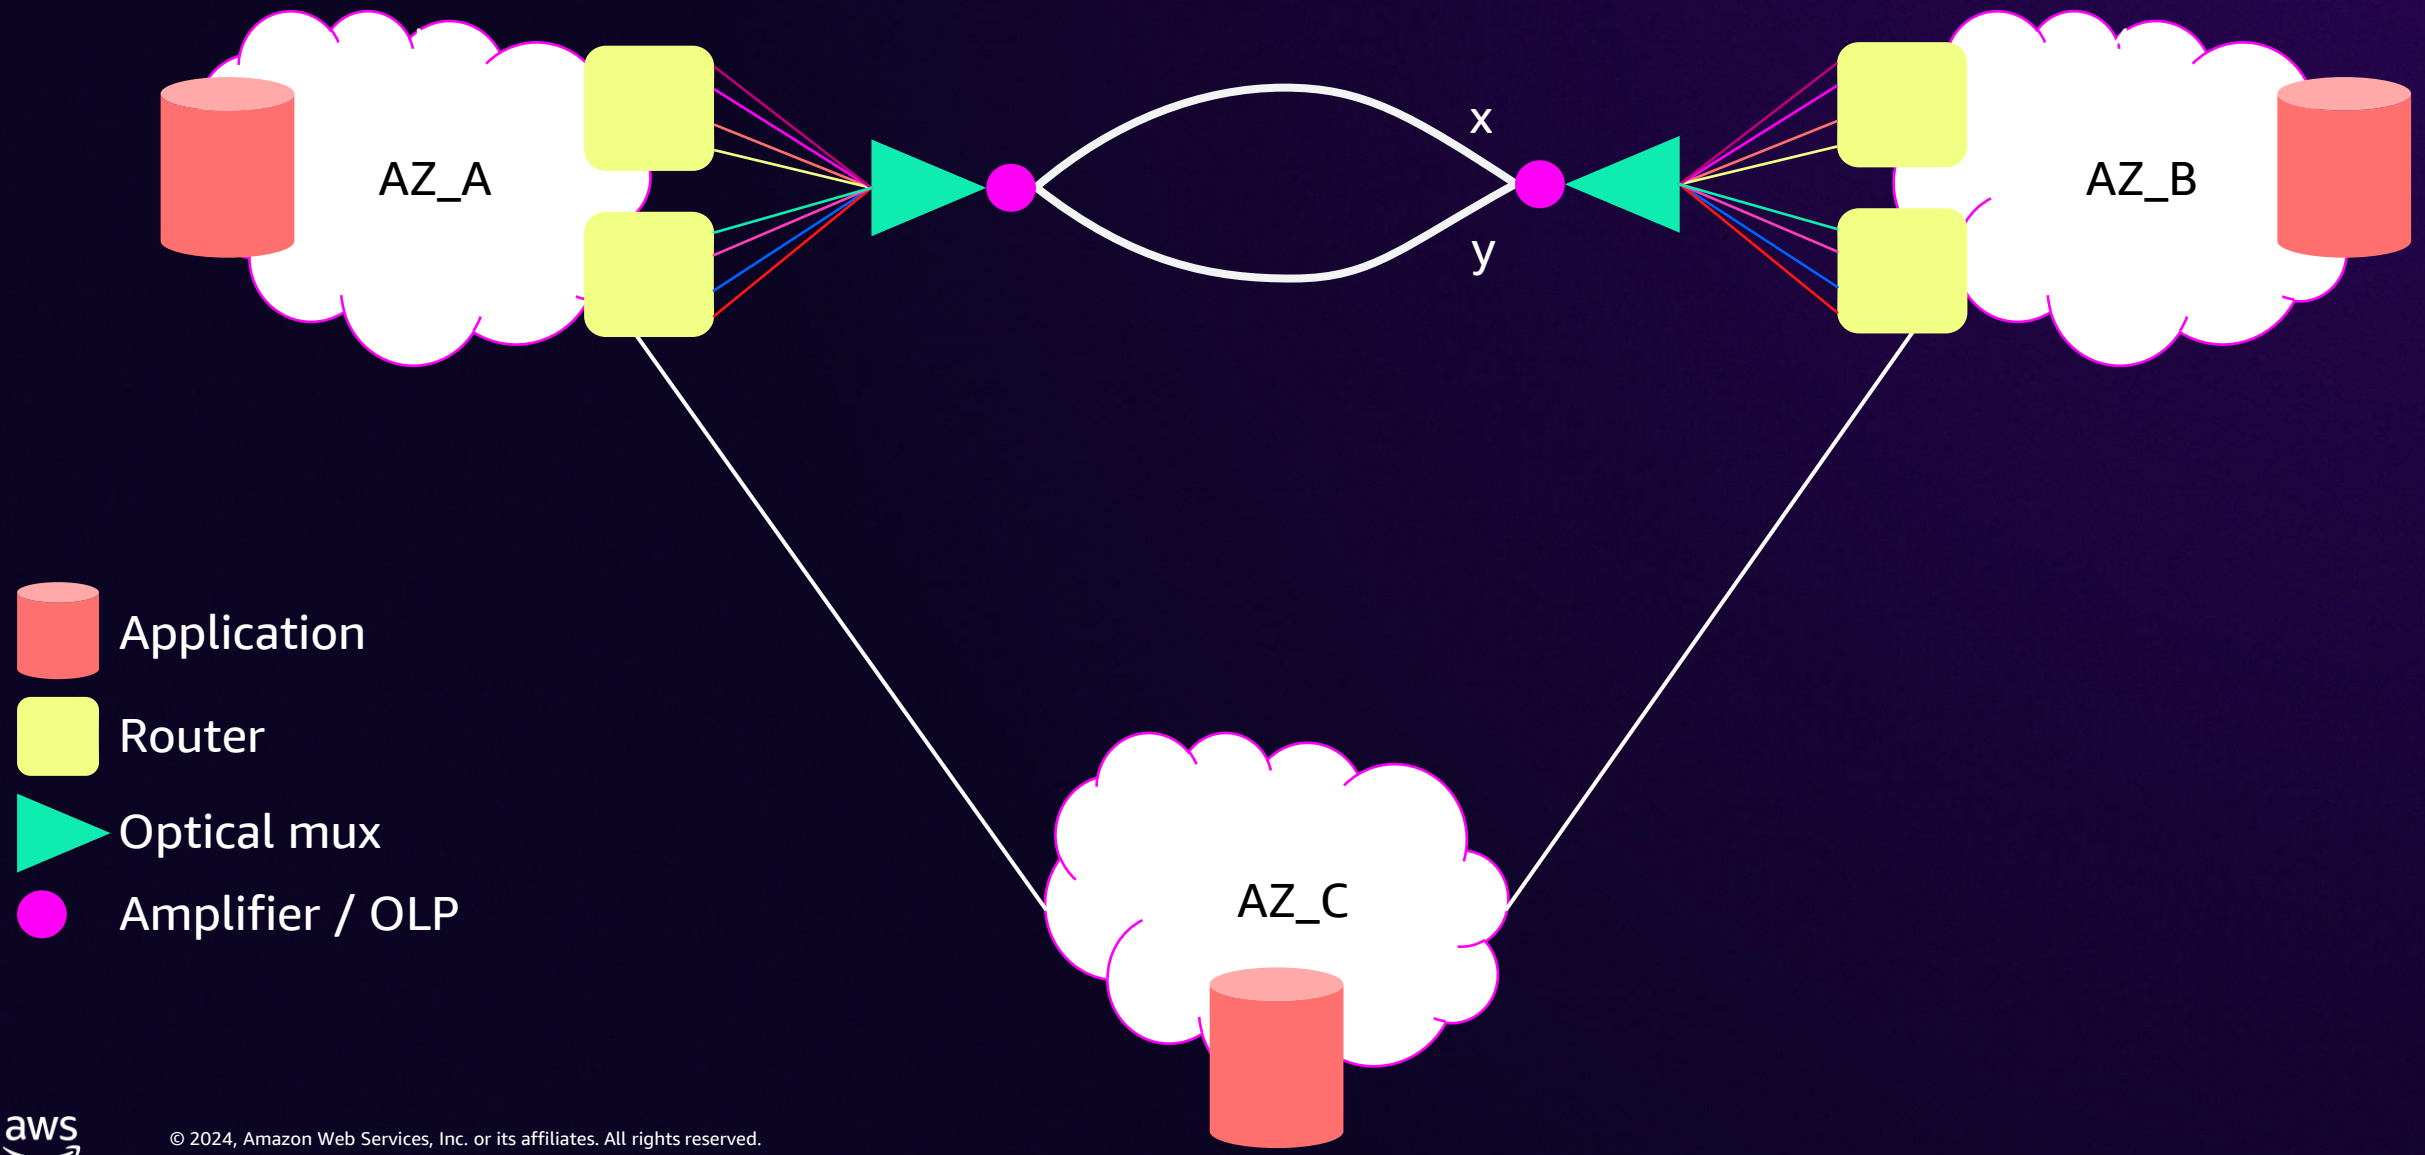

Capacity management

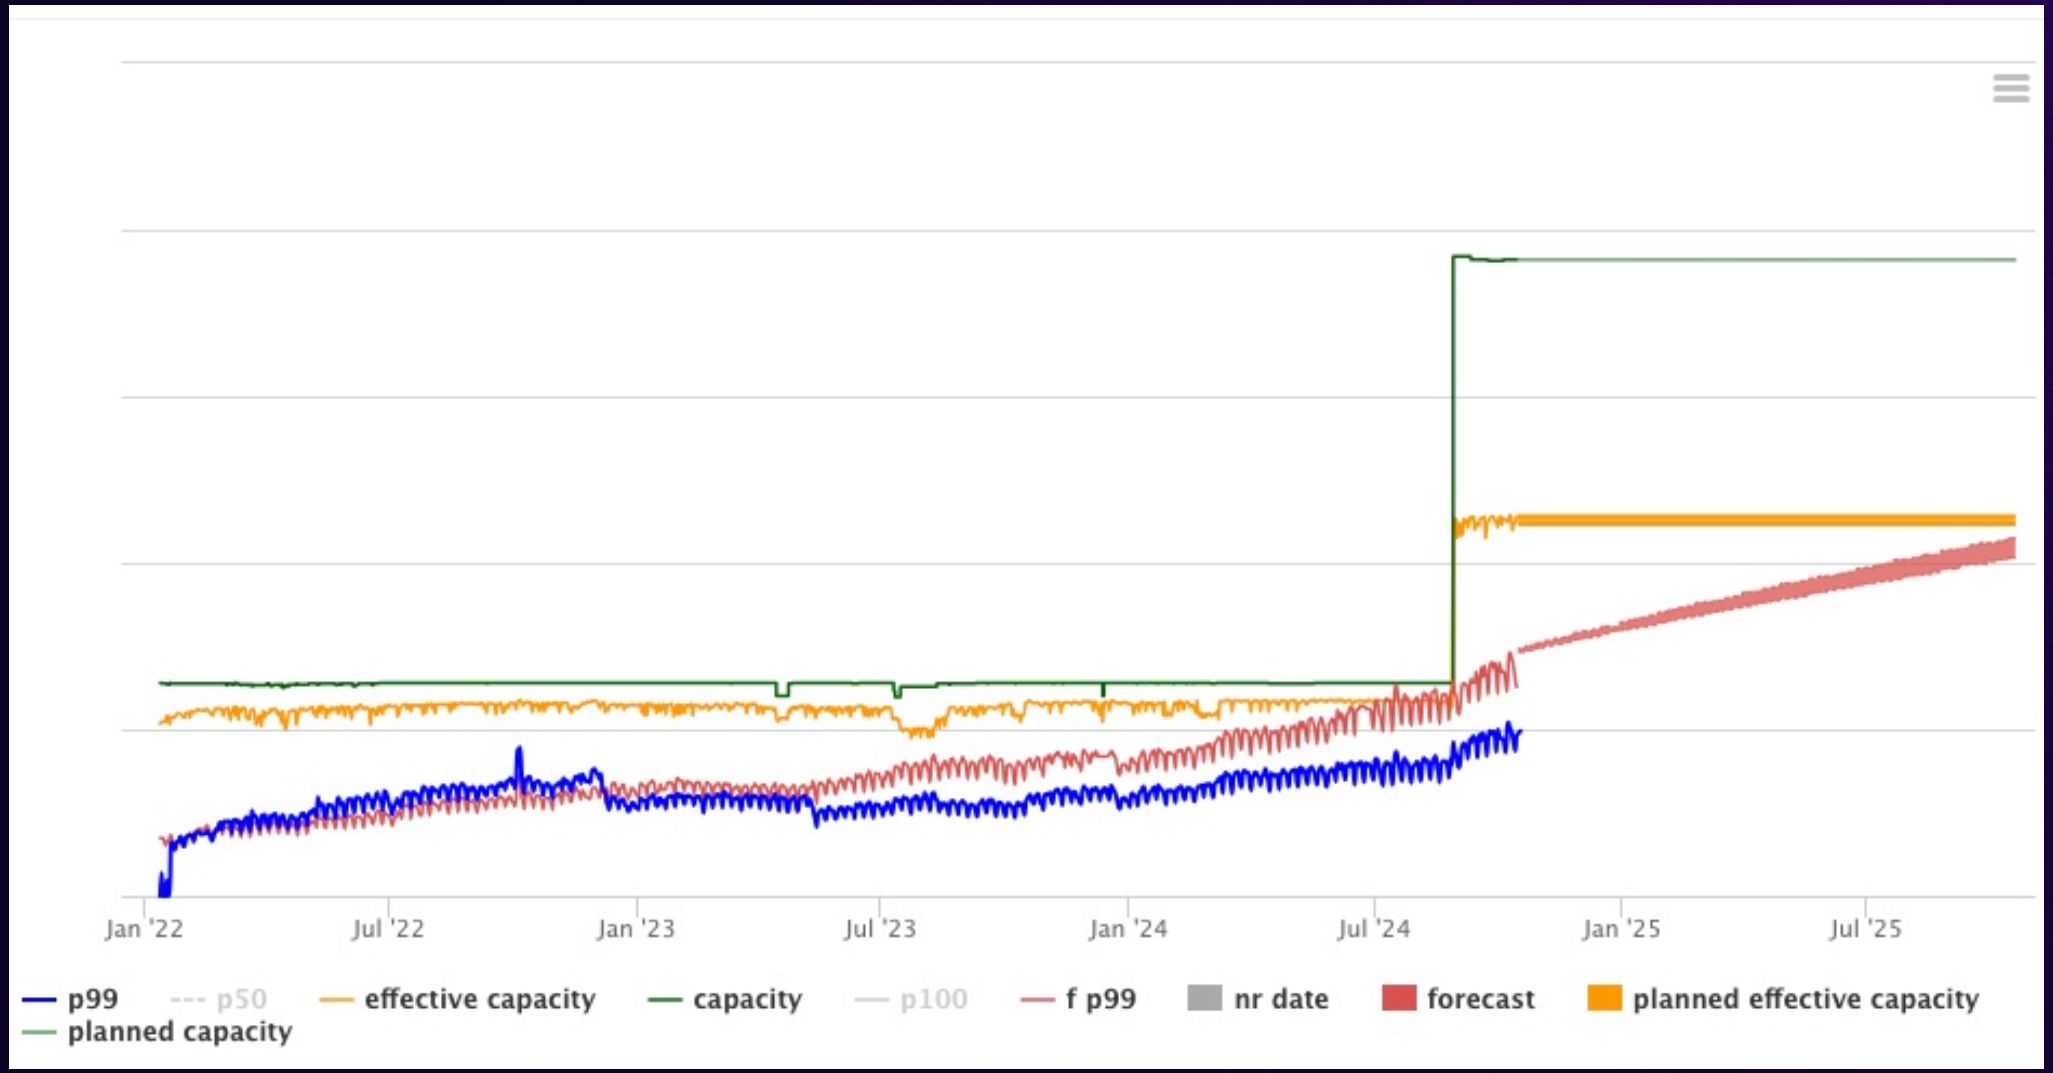

Capacity management

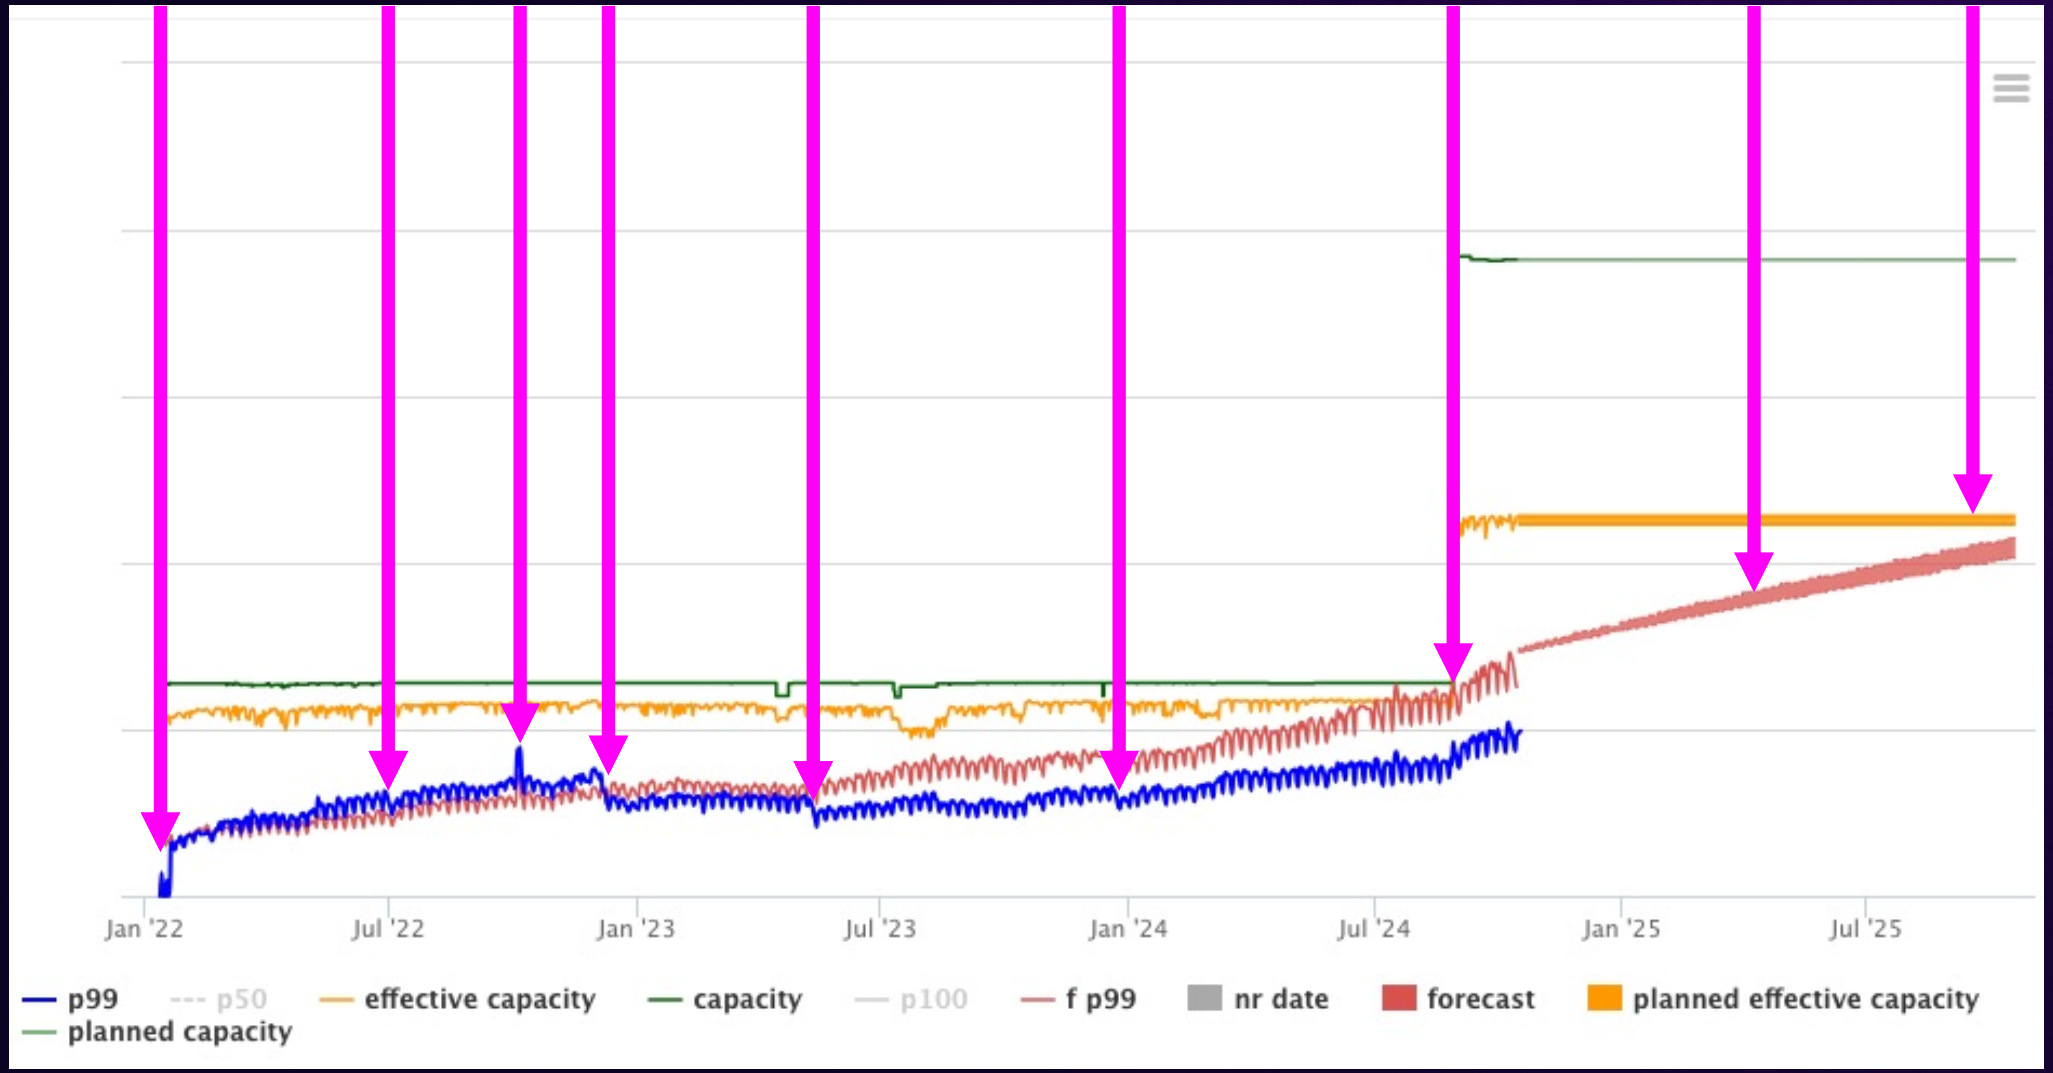

Features

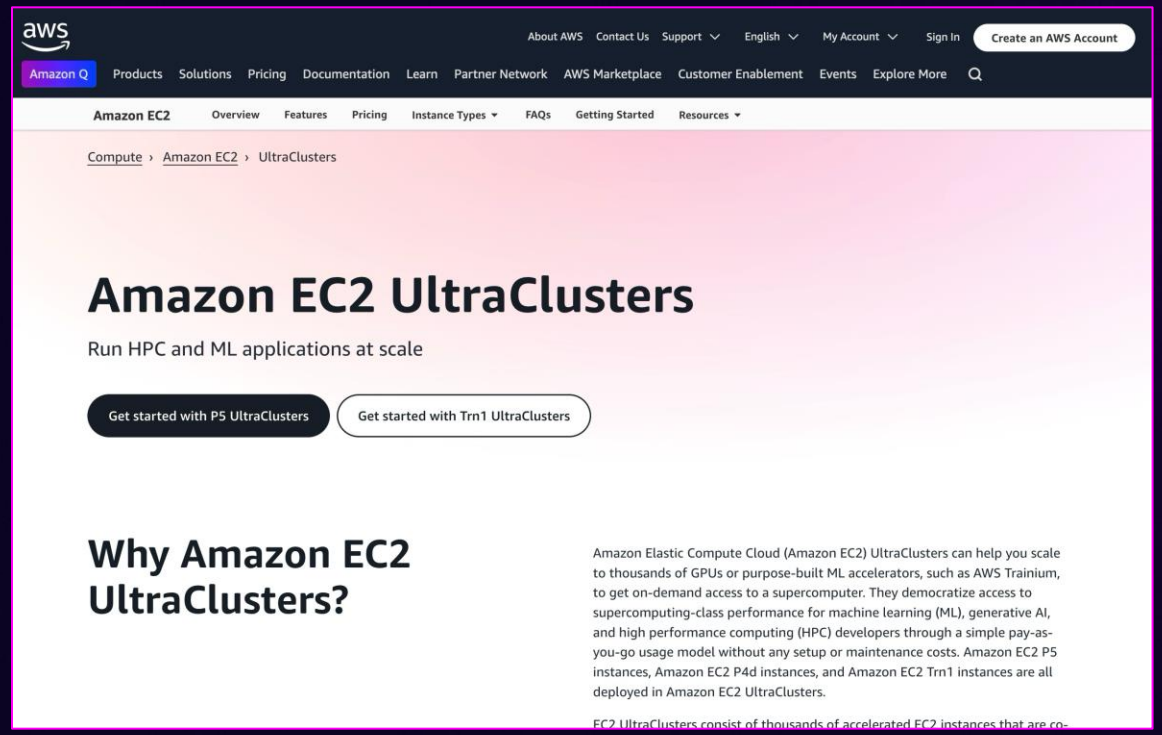

Network design principles

1. Automate as much as possible

1. Automate as much as possible

- Configuration
- Telemetry
- Traffic engineering

1. Automate as much as possible

850 raw events per second

2.4 human engagements per hour

2. We should fail seldomly and in a predictable way

2. We should fail seldomly and in a predictable way

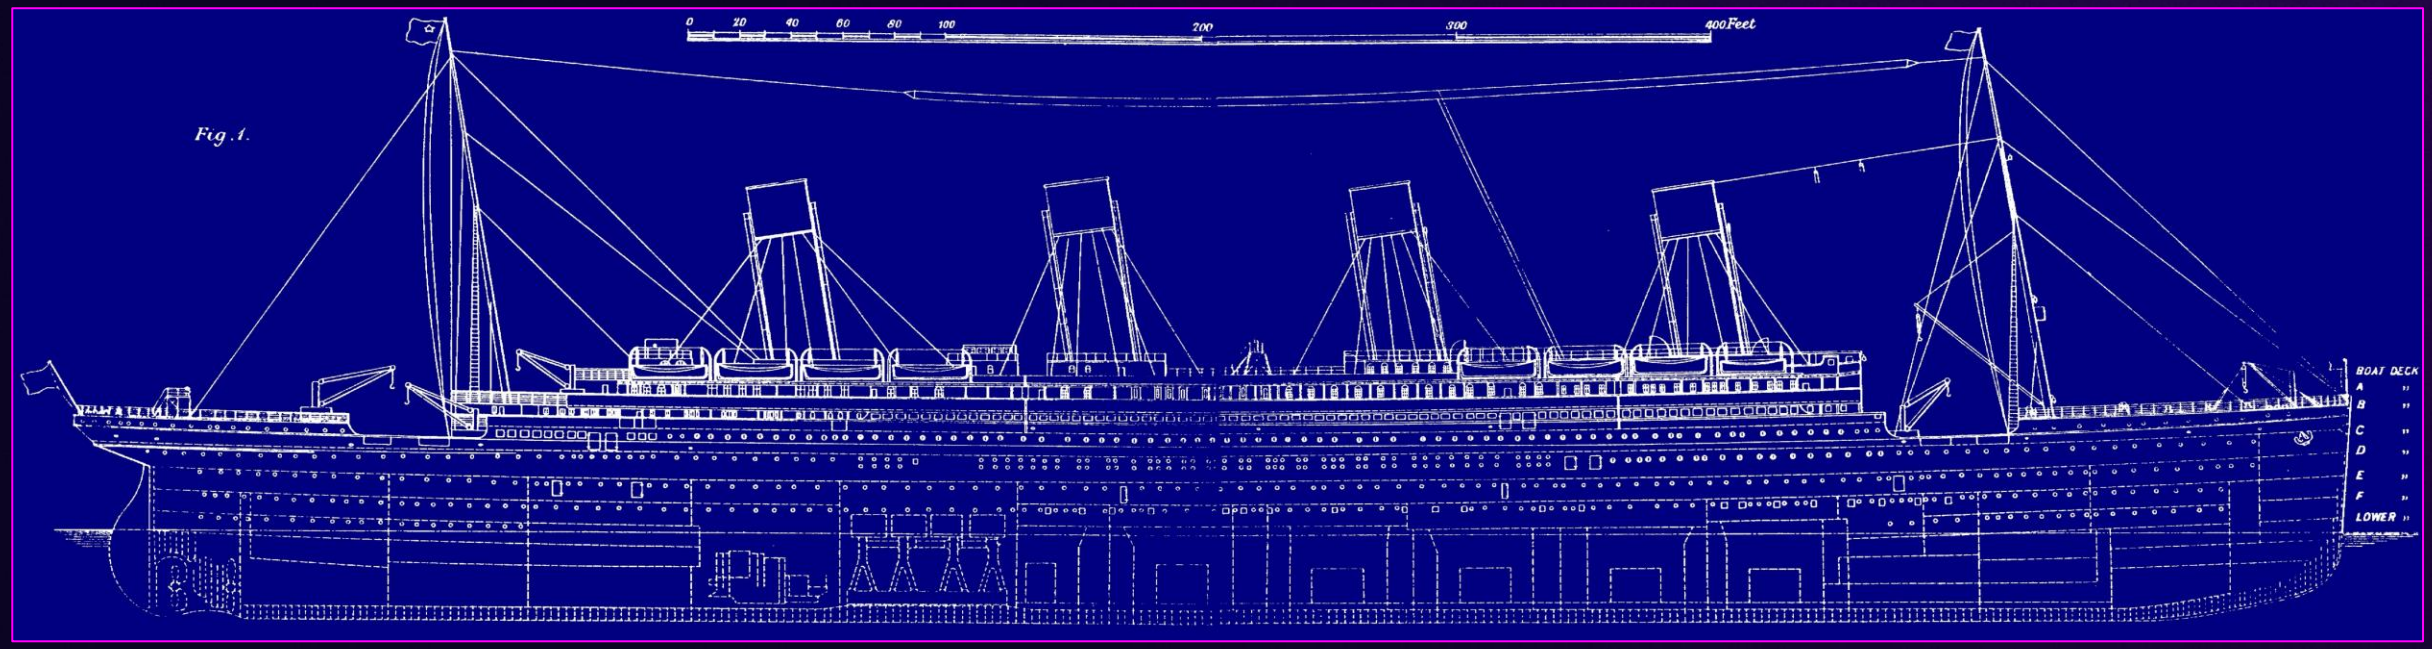

Image © by Mr Aquitania

3. Don’t reach unprecedented scale

Photo © by Björn Strey

3. Don’t reach unprecedented scale

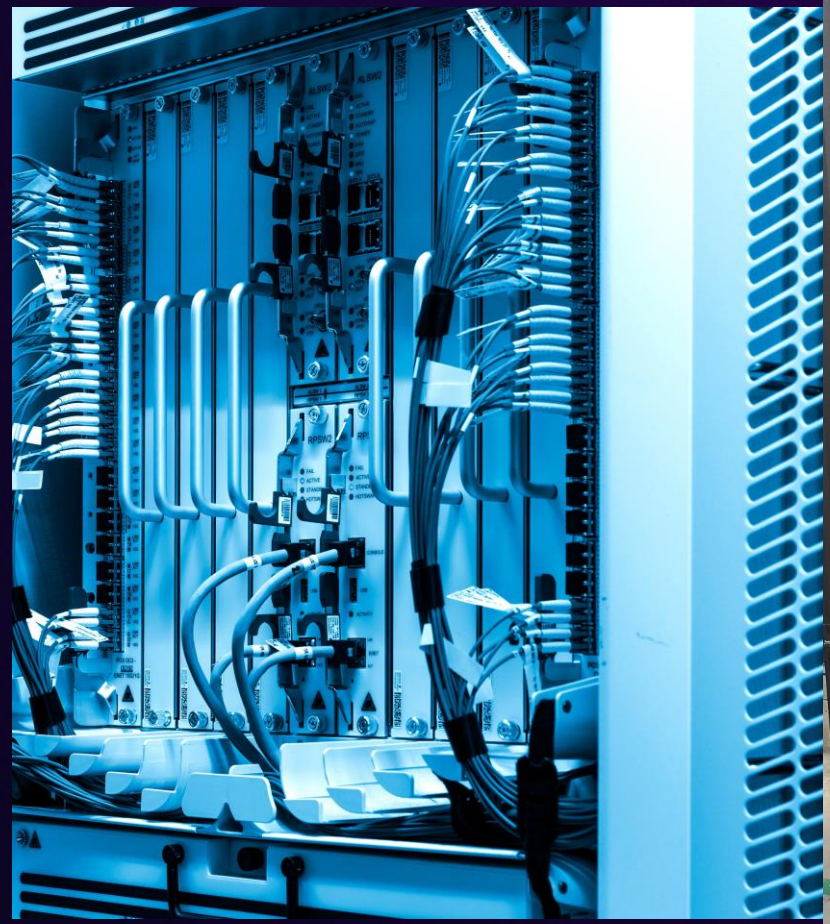

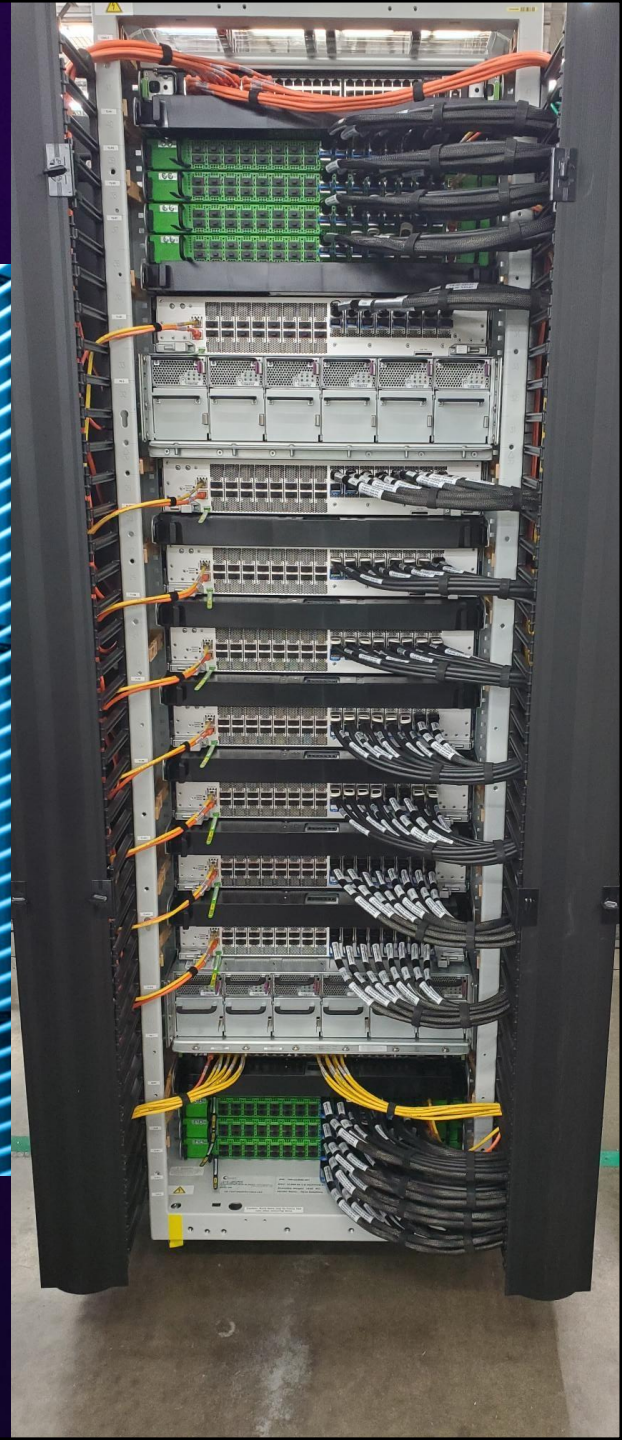

3. Don’t reach unprecedented scale

Network designs

=

Distributed systems designs

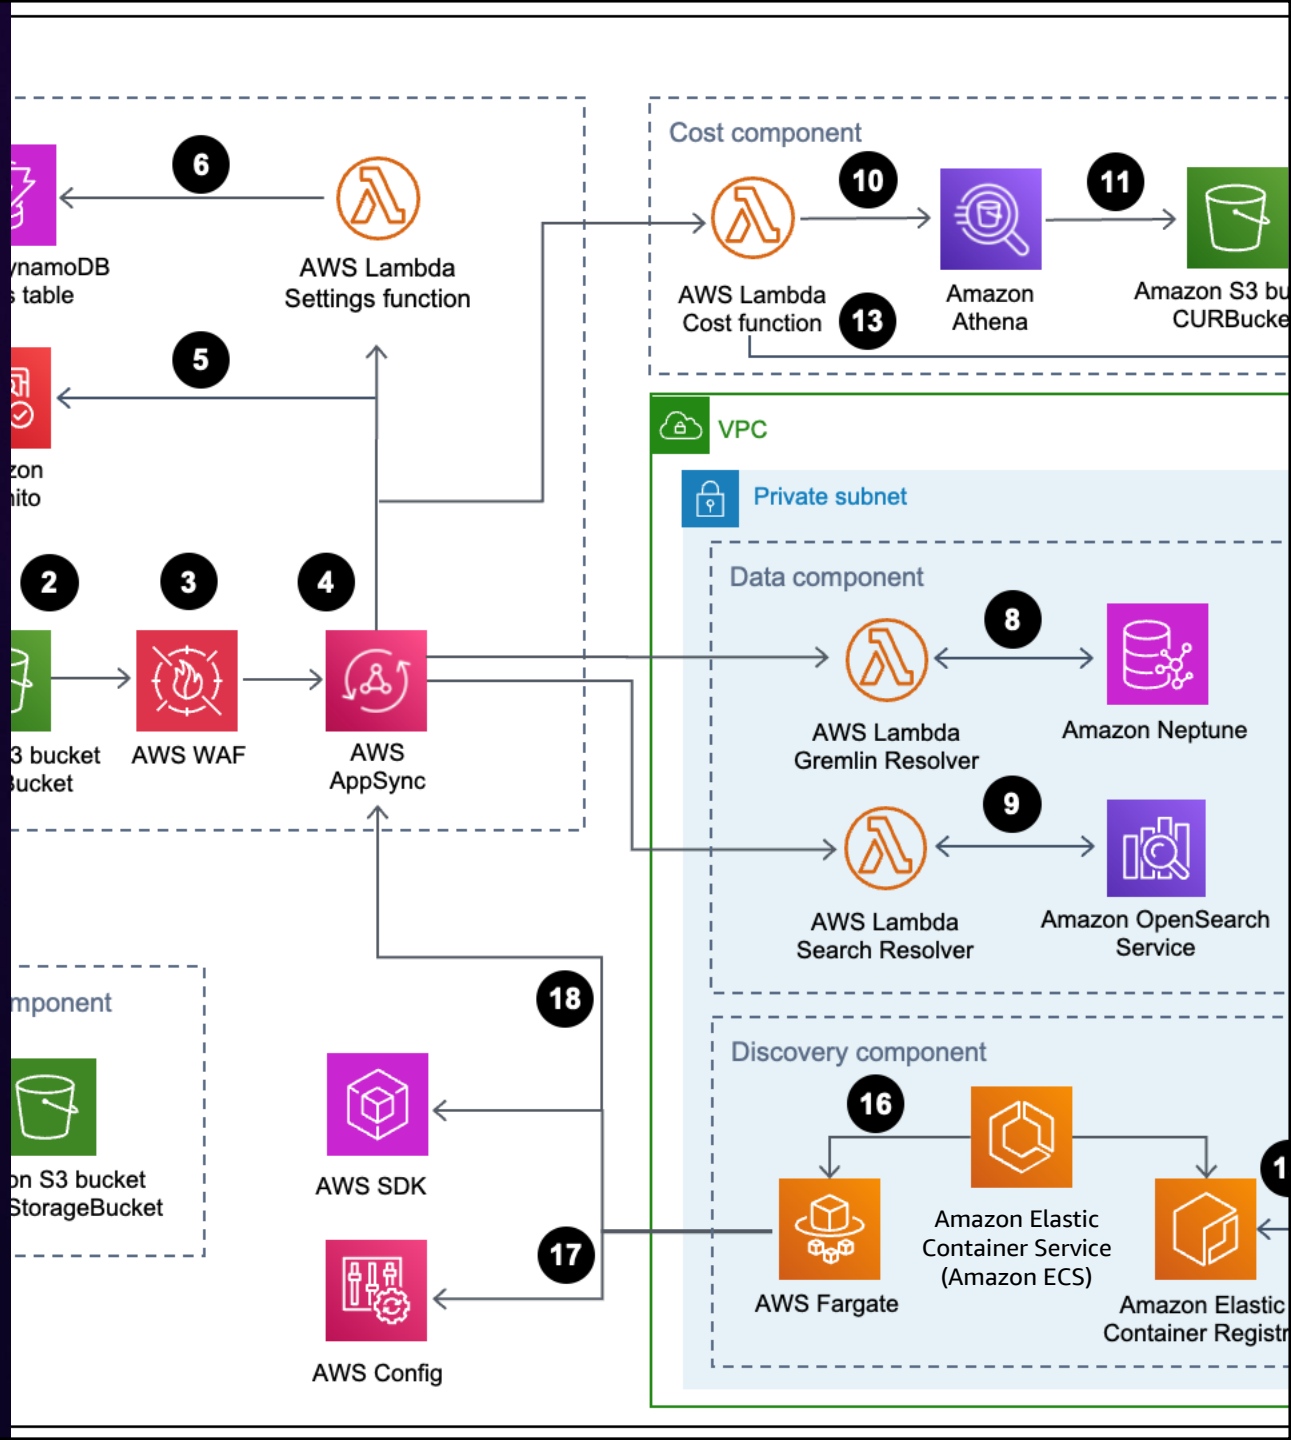

Design for fallible humans by using Intents (2023)

Design for fallible humans by using Intents

EC2 to internet

40X improvement

SDN journey

1. Monitoring, alarming, triangulation

2. Automation

3. Closed-loop controllers

Outbound traffic engineering

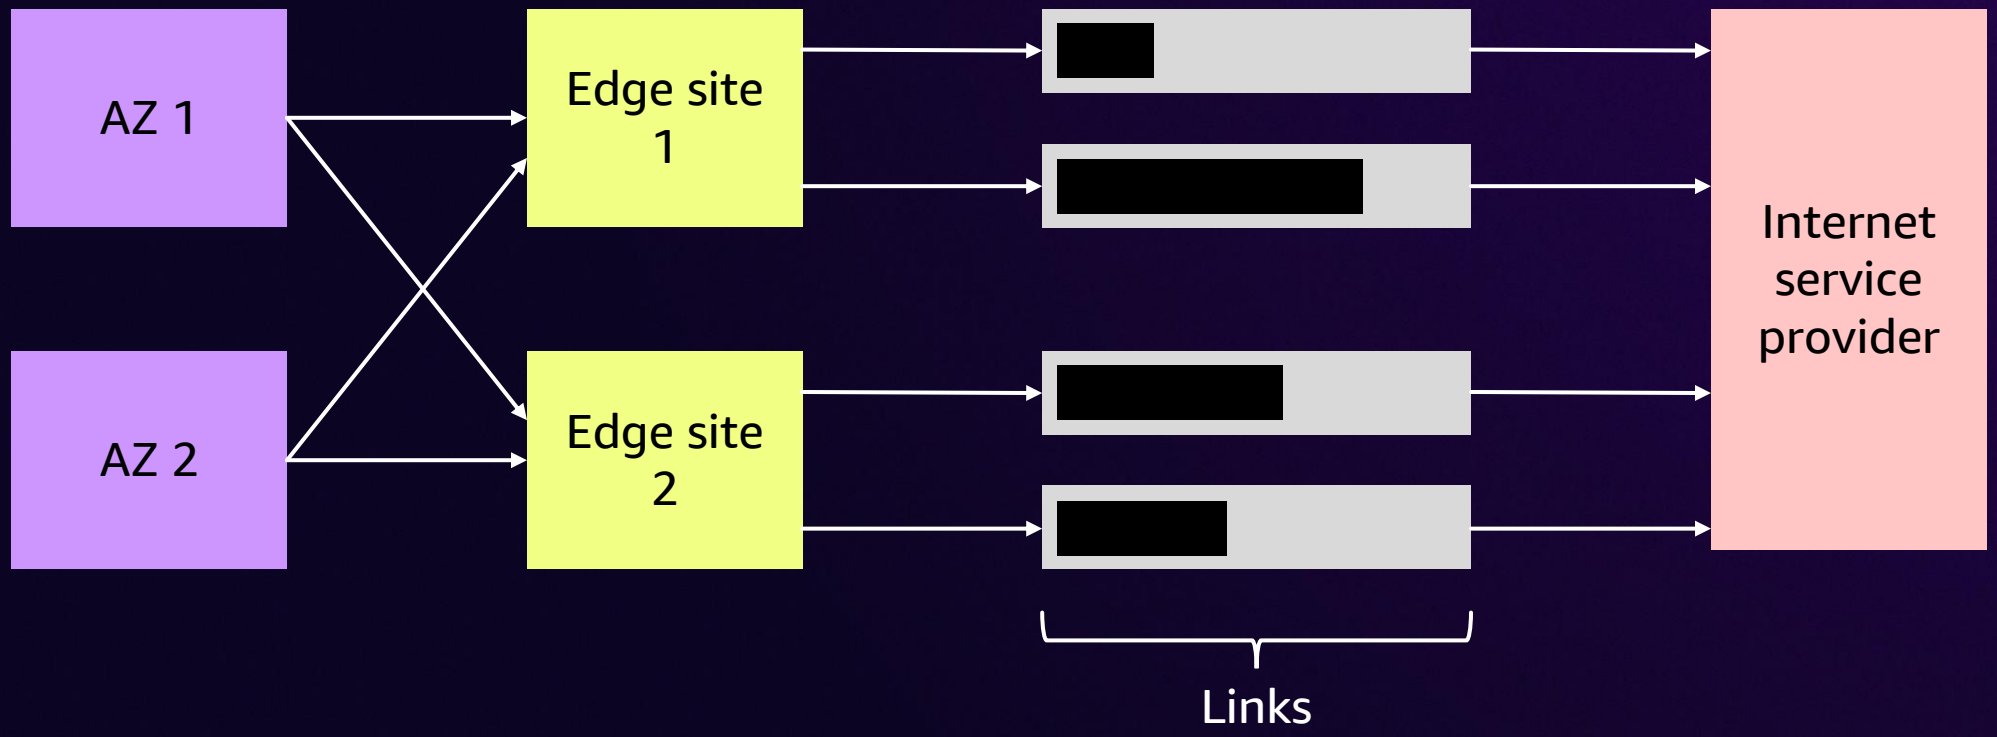

Congestion mitigation

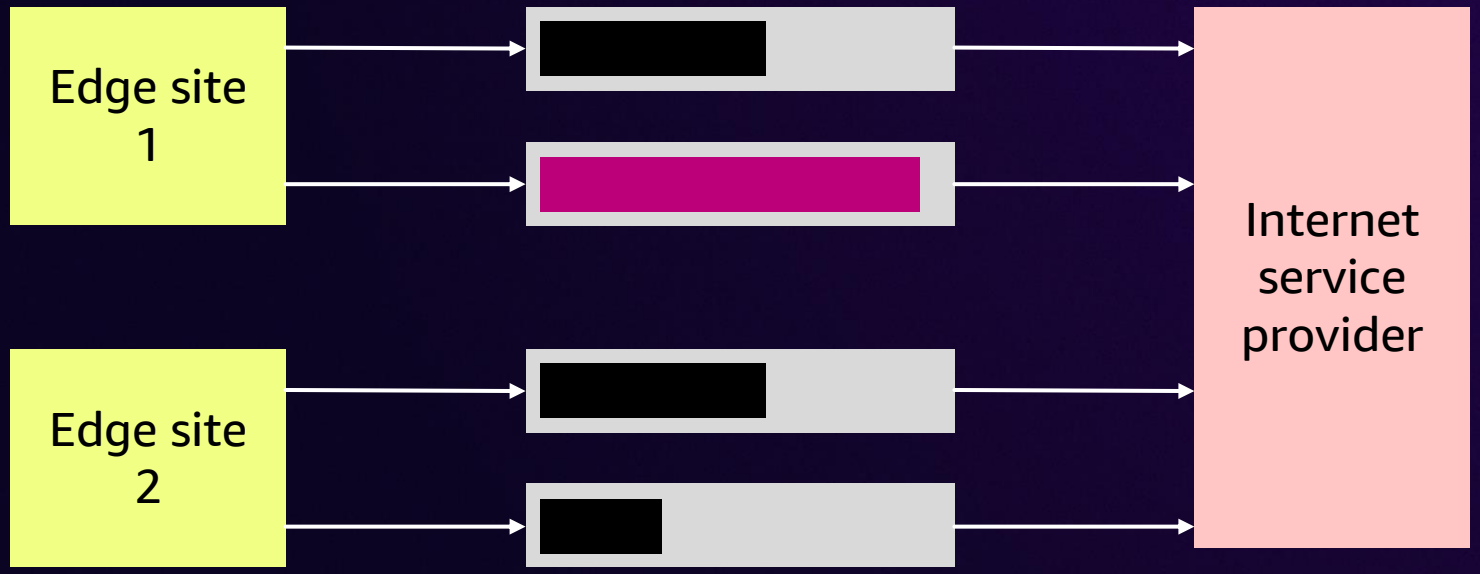

Congestion mitigation

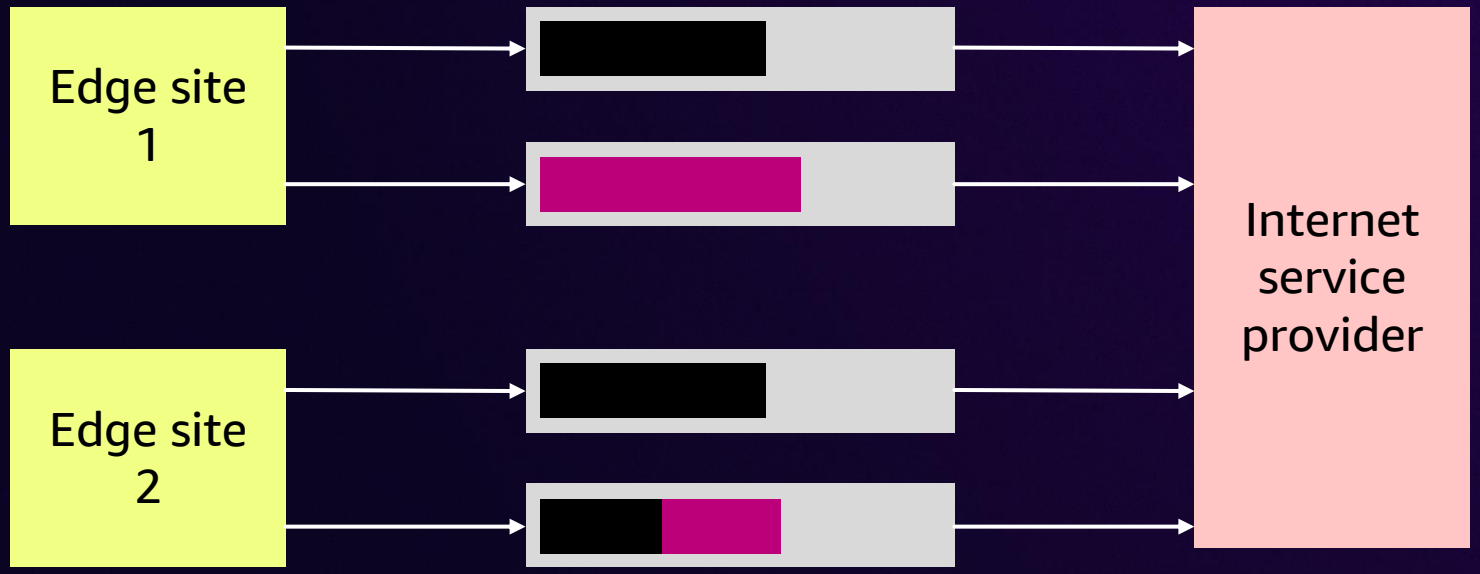

Link balancing

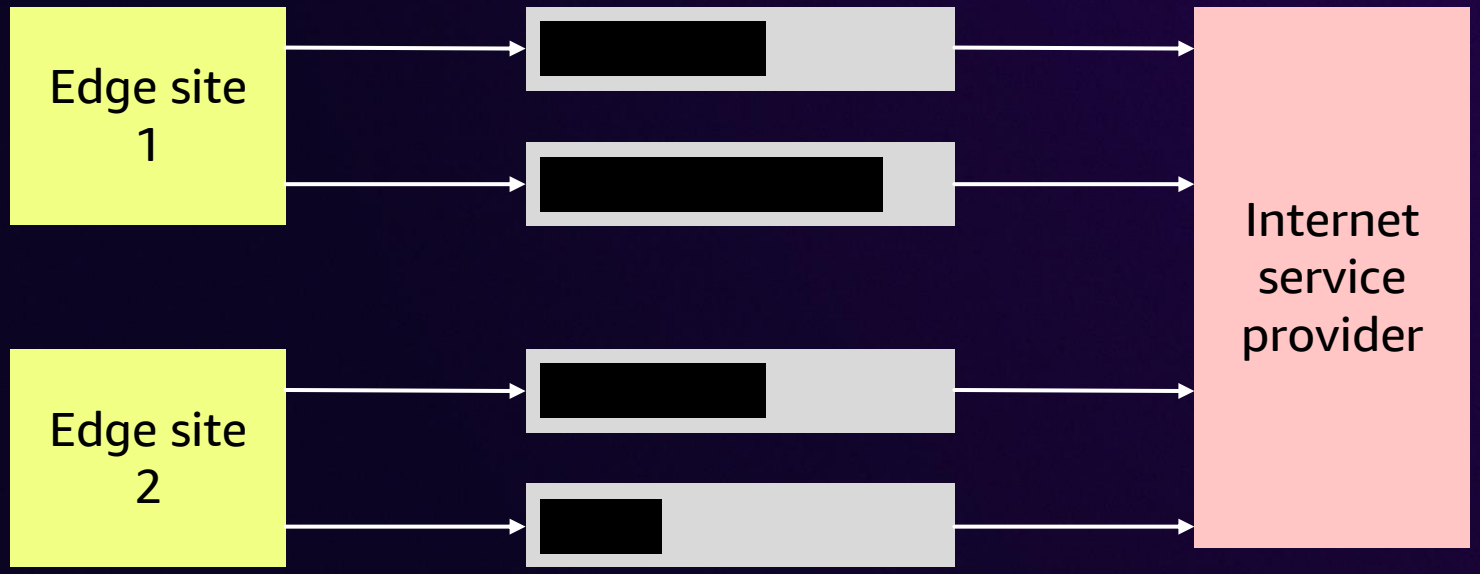

Link balancing

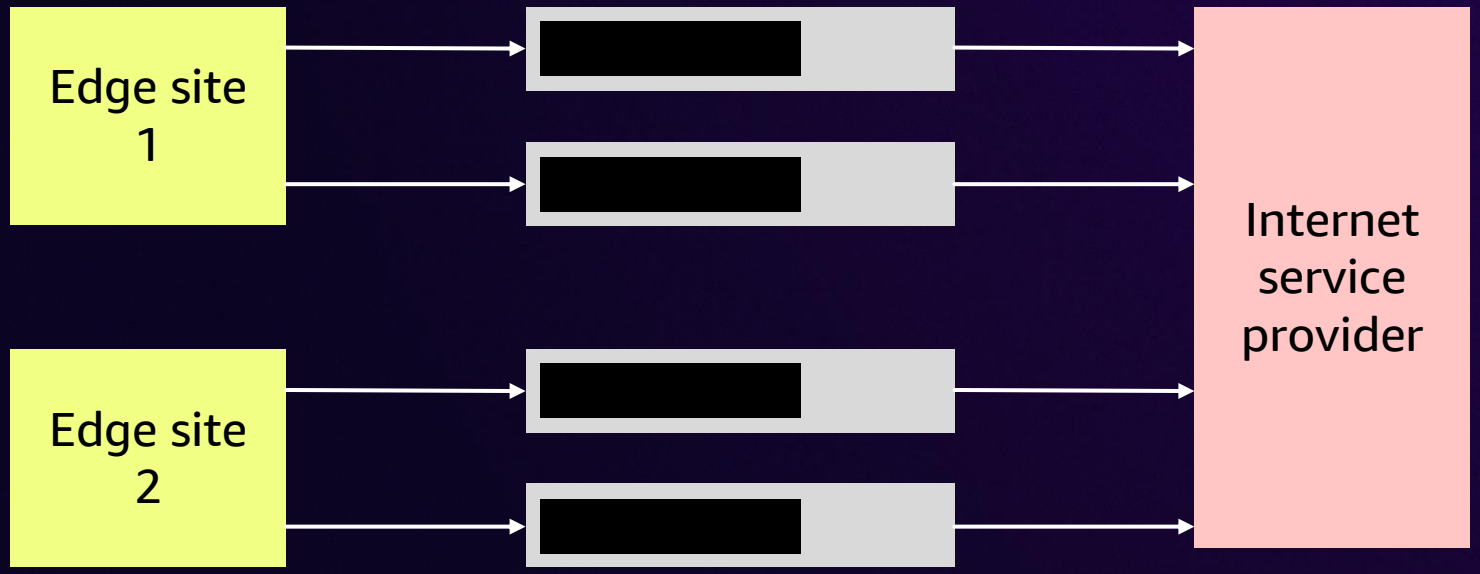

Performance optimization

Performance optimization

AZ independence

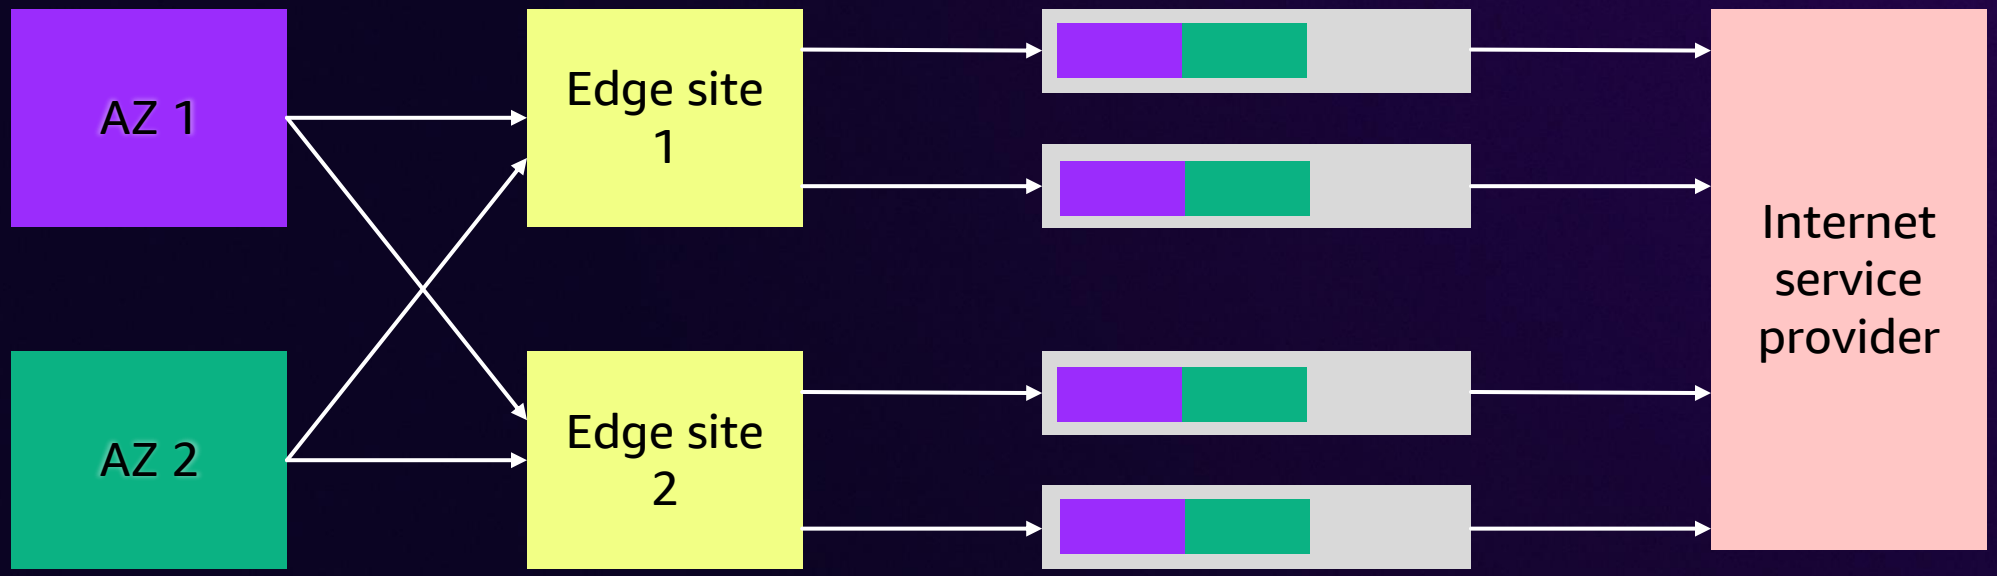

AZ independence

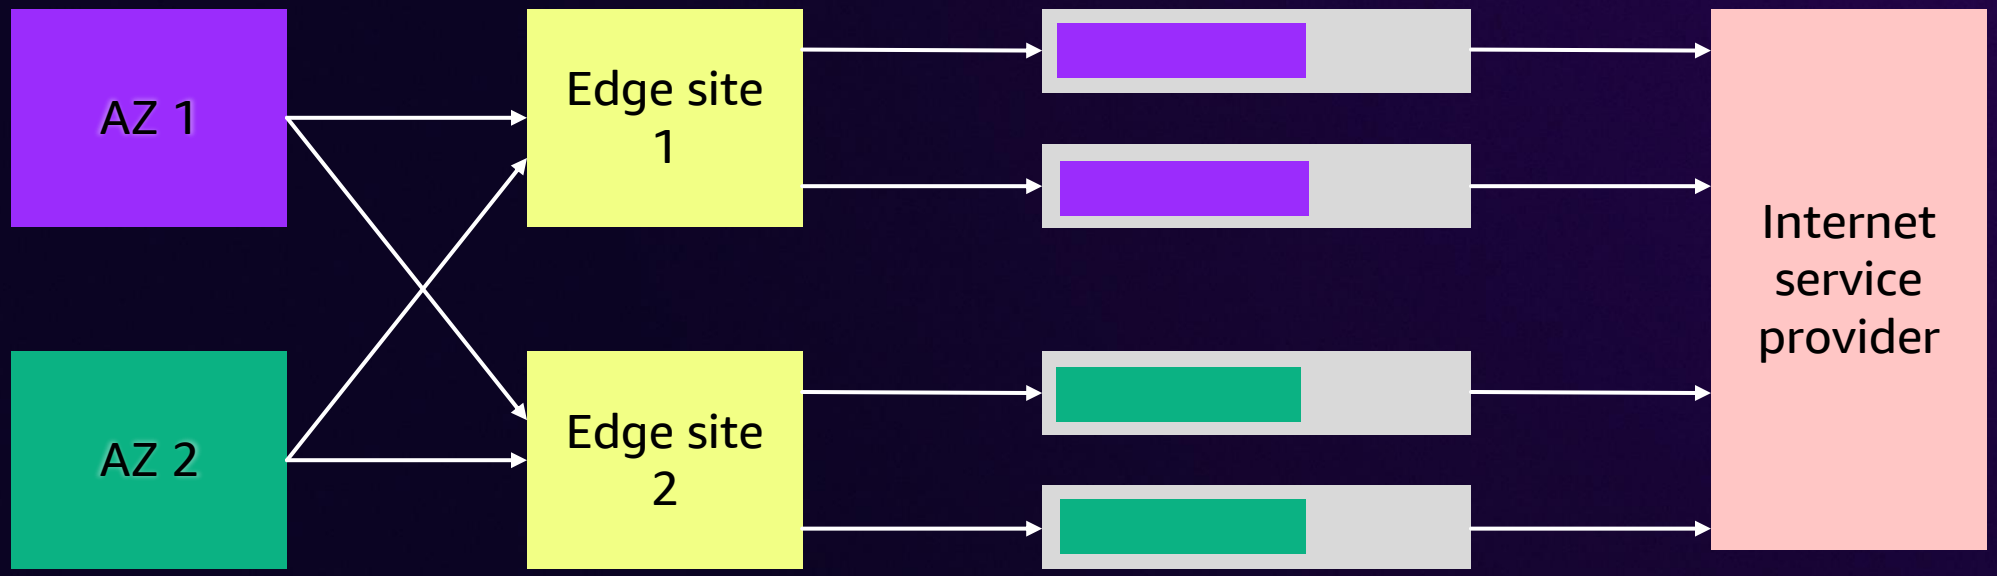

Inbound traffic engineering

Inbound traffic engineering

7X faster

mitigations for inbound congestion

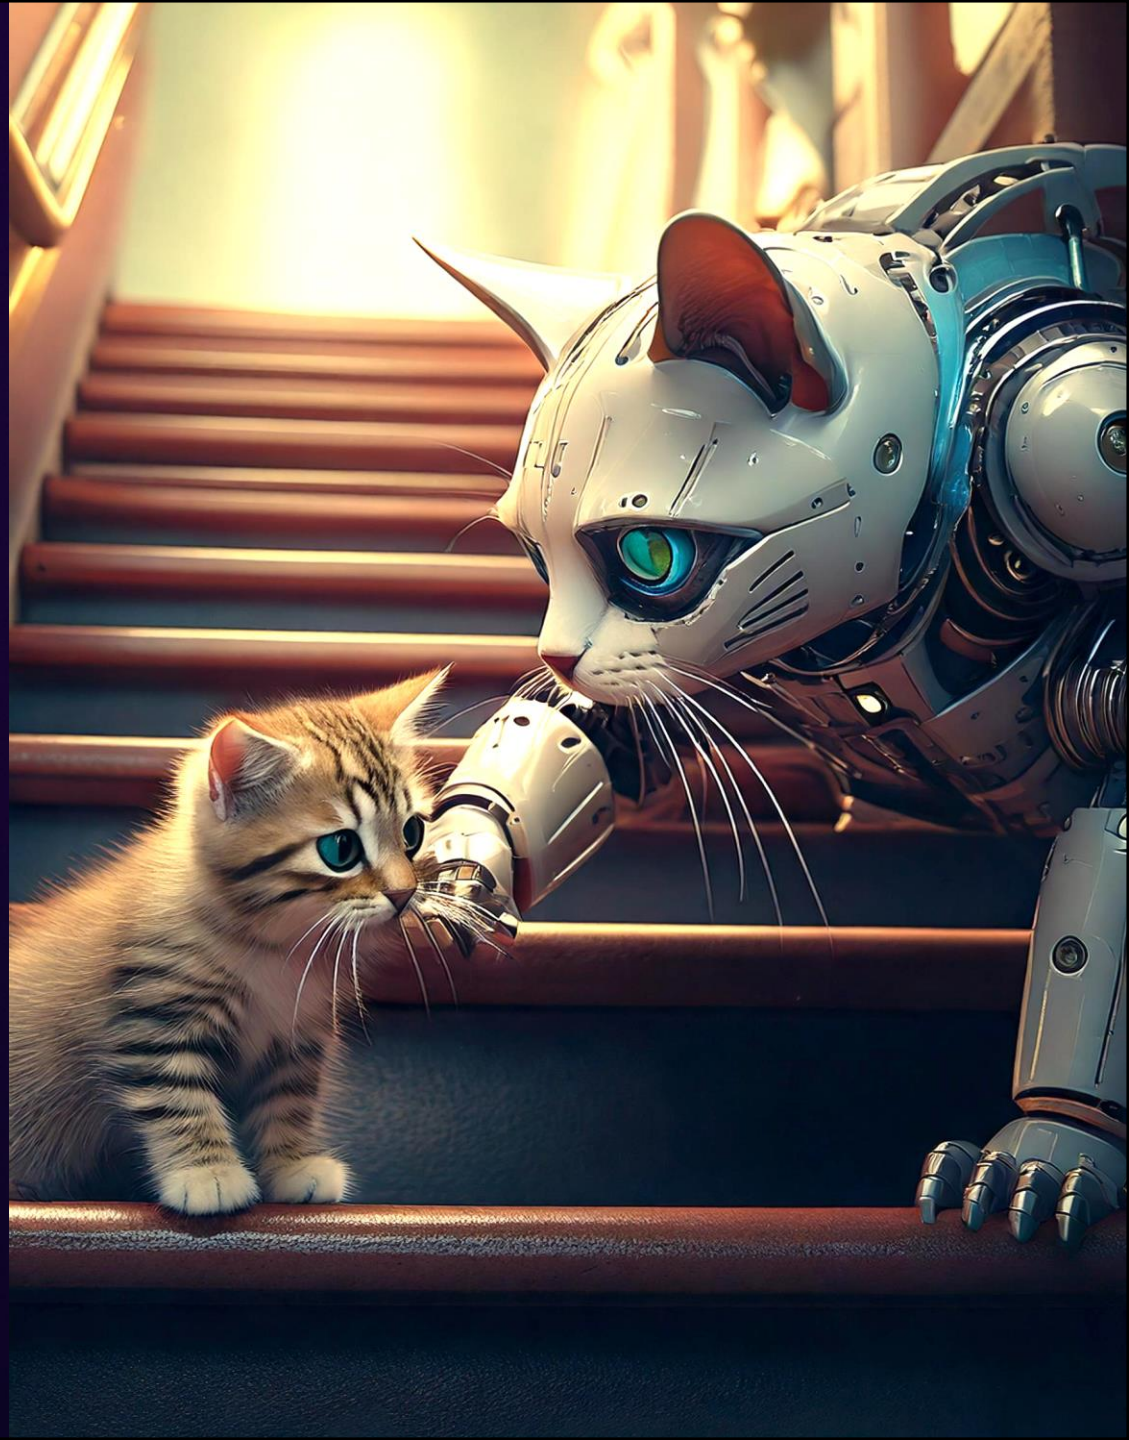

Traffic engineering on the WAN

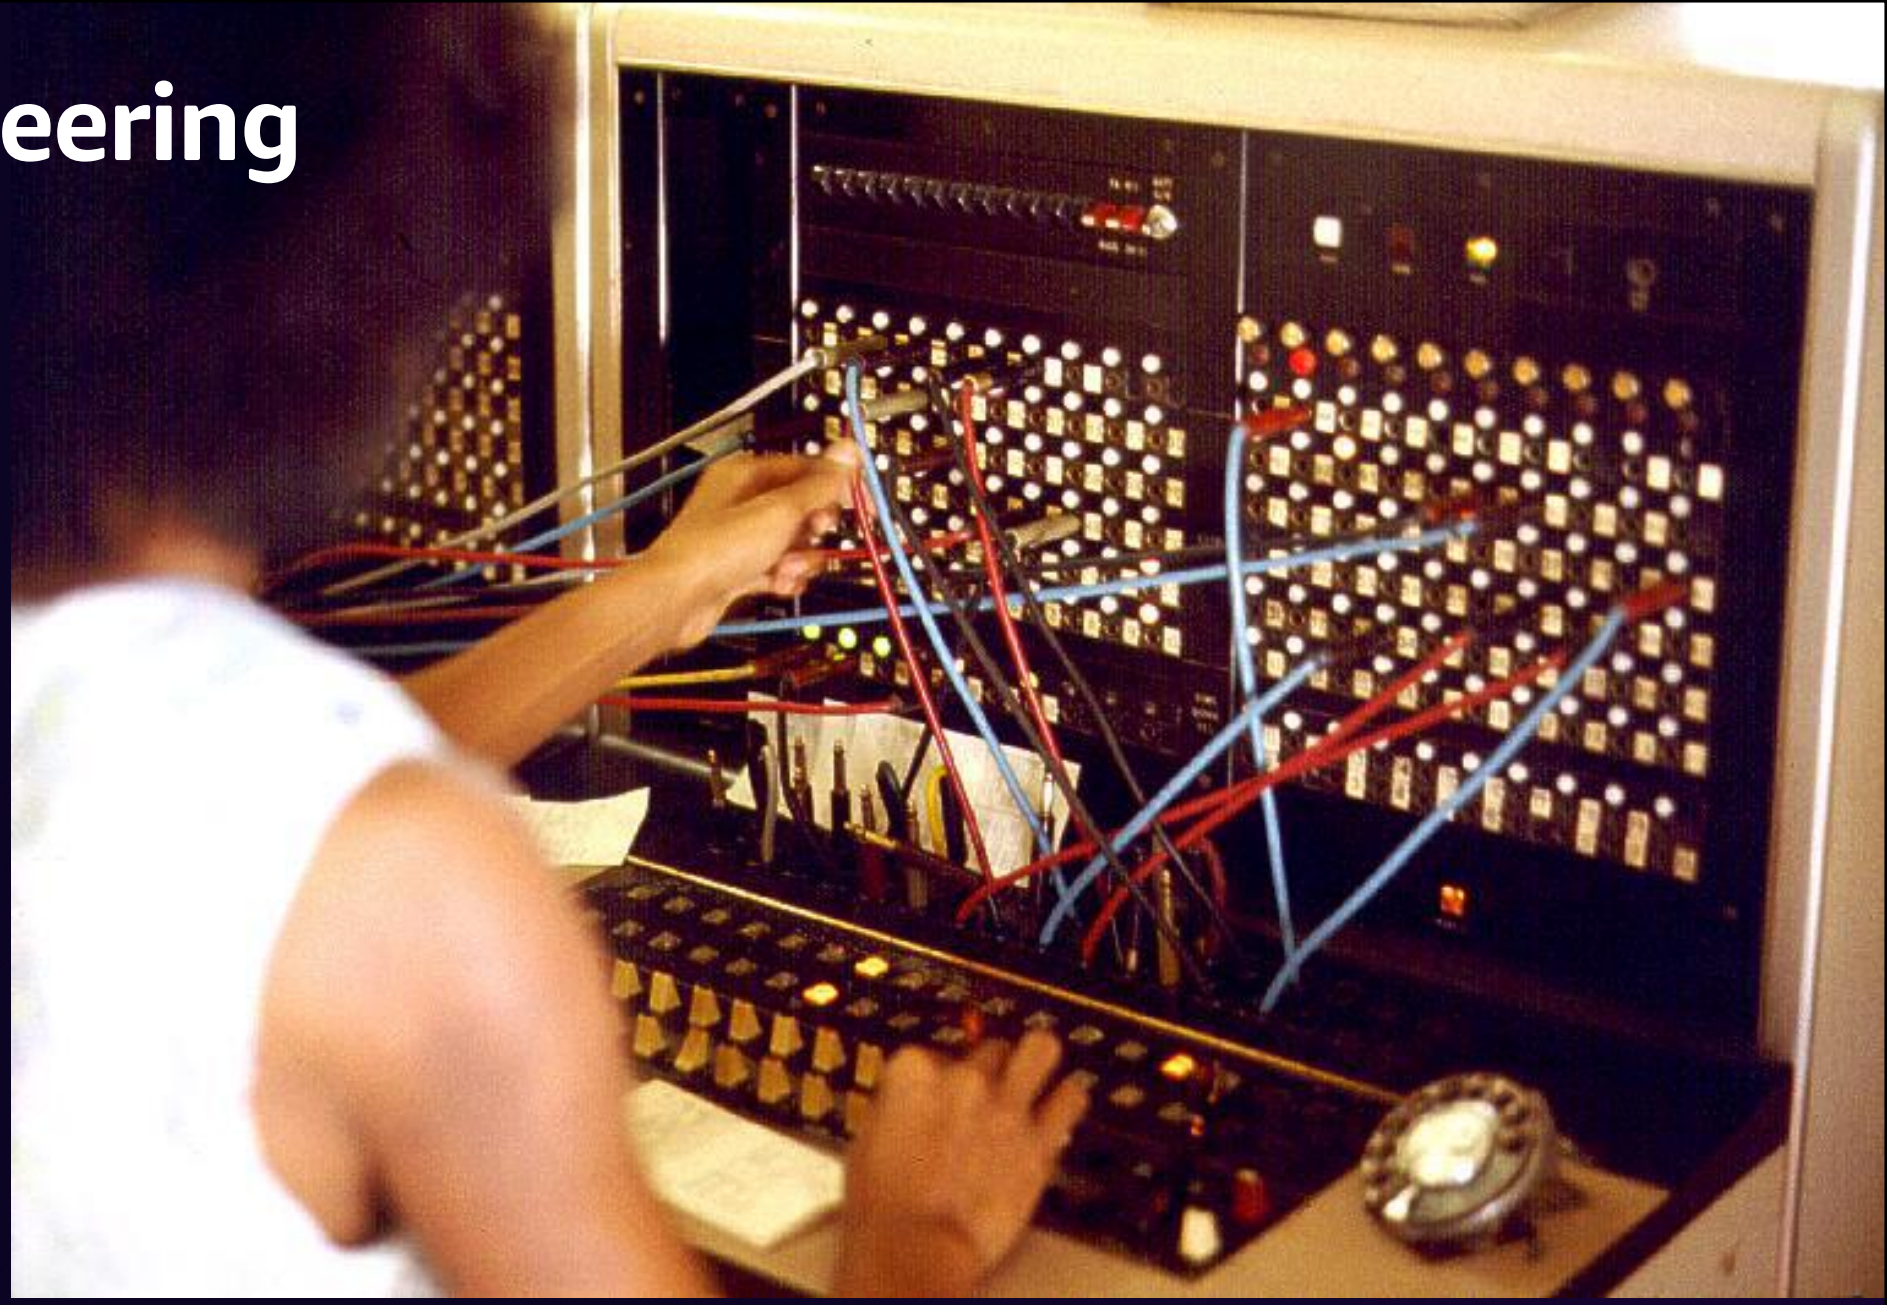

Photo © by Joseph A. Carr

SDN with global brains

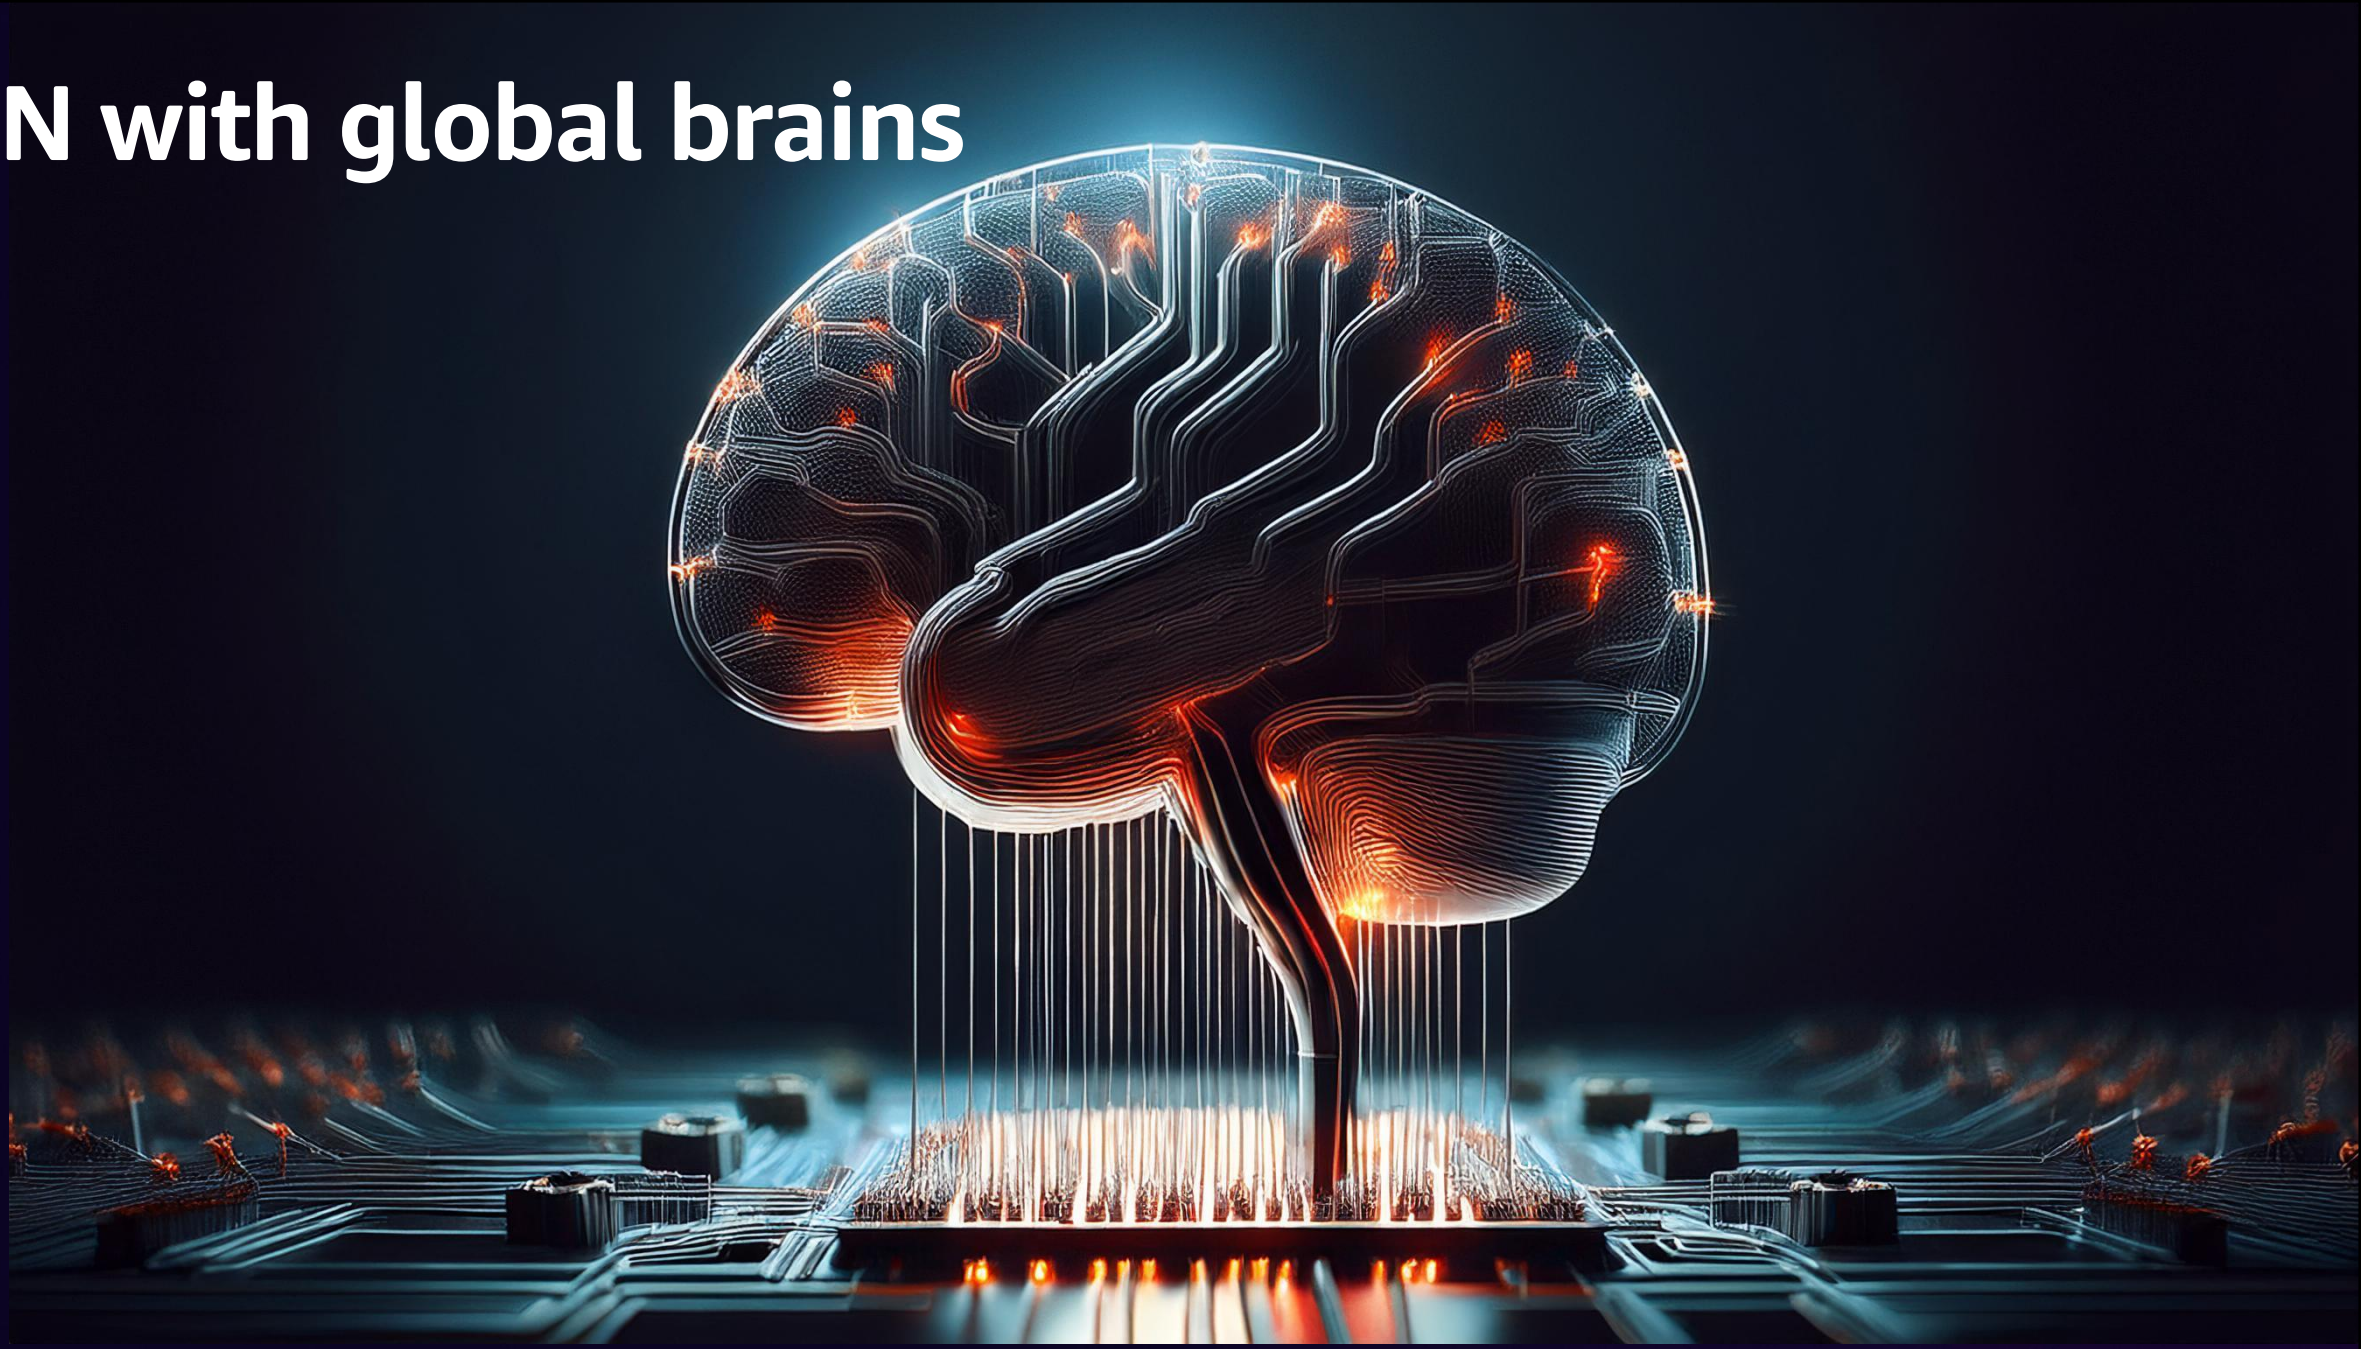

Distributed and optimal SDN

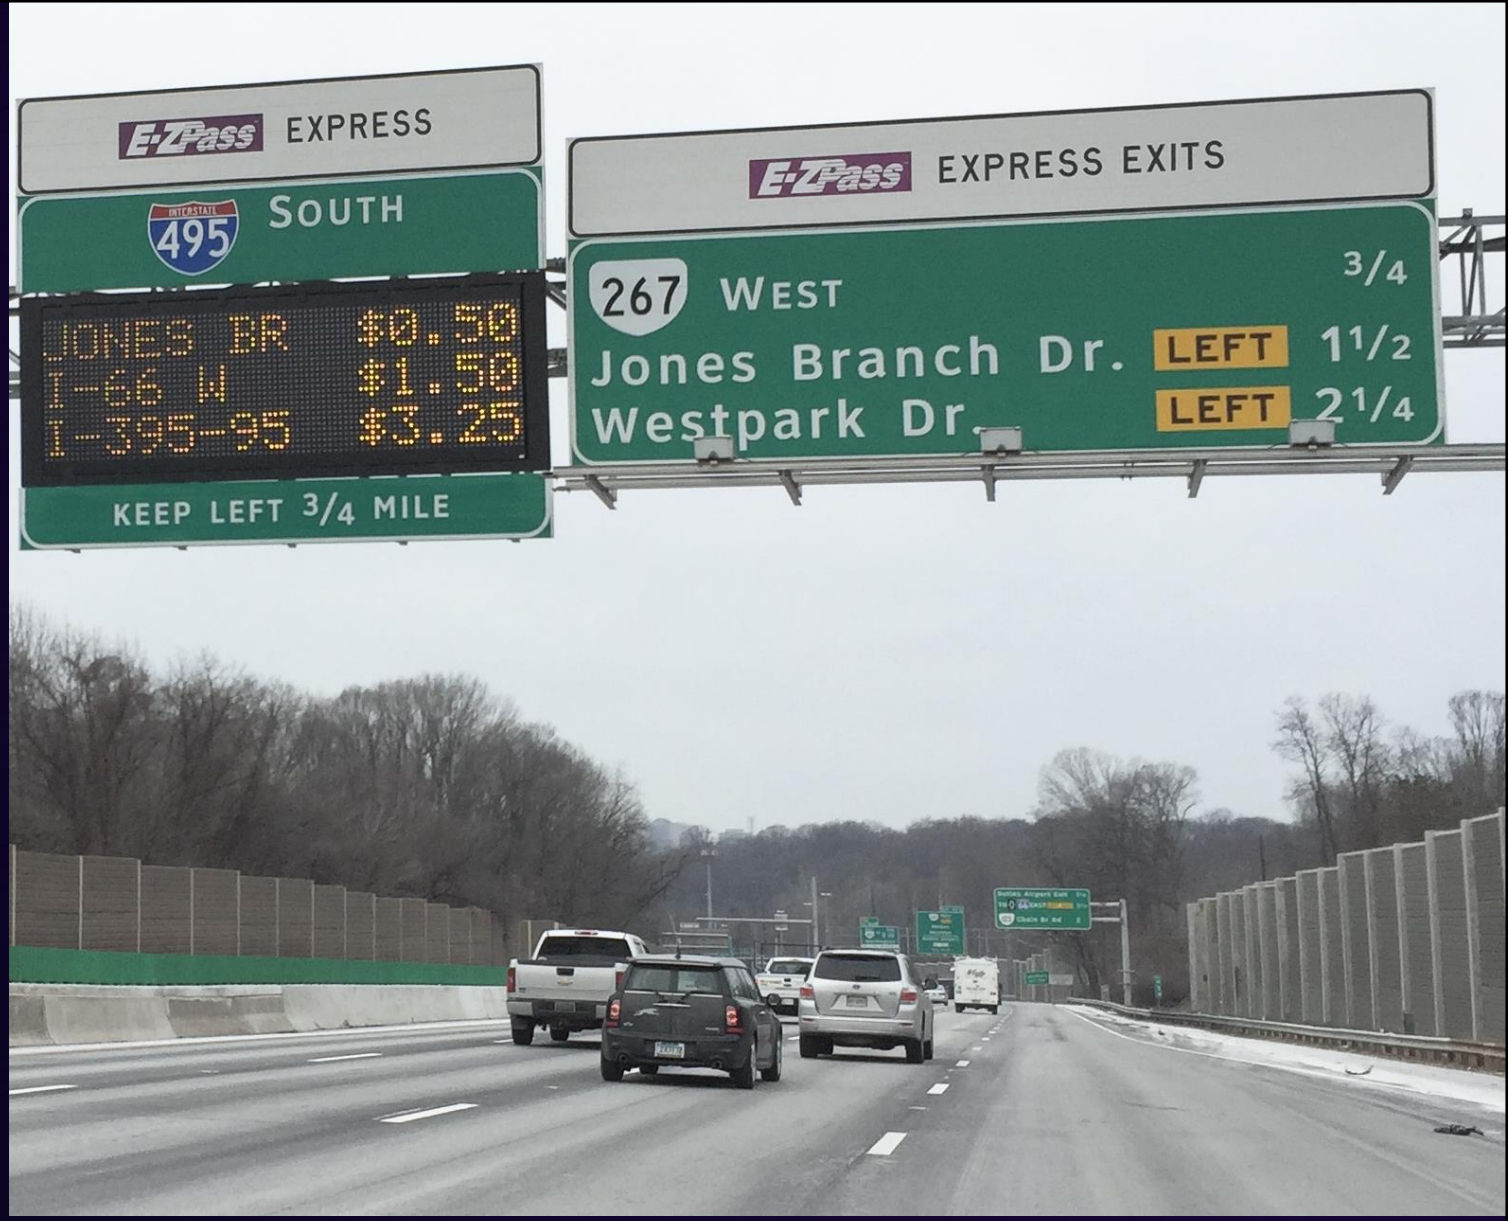

Photo © by Famartin

Distributed and optimal SDN

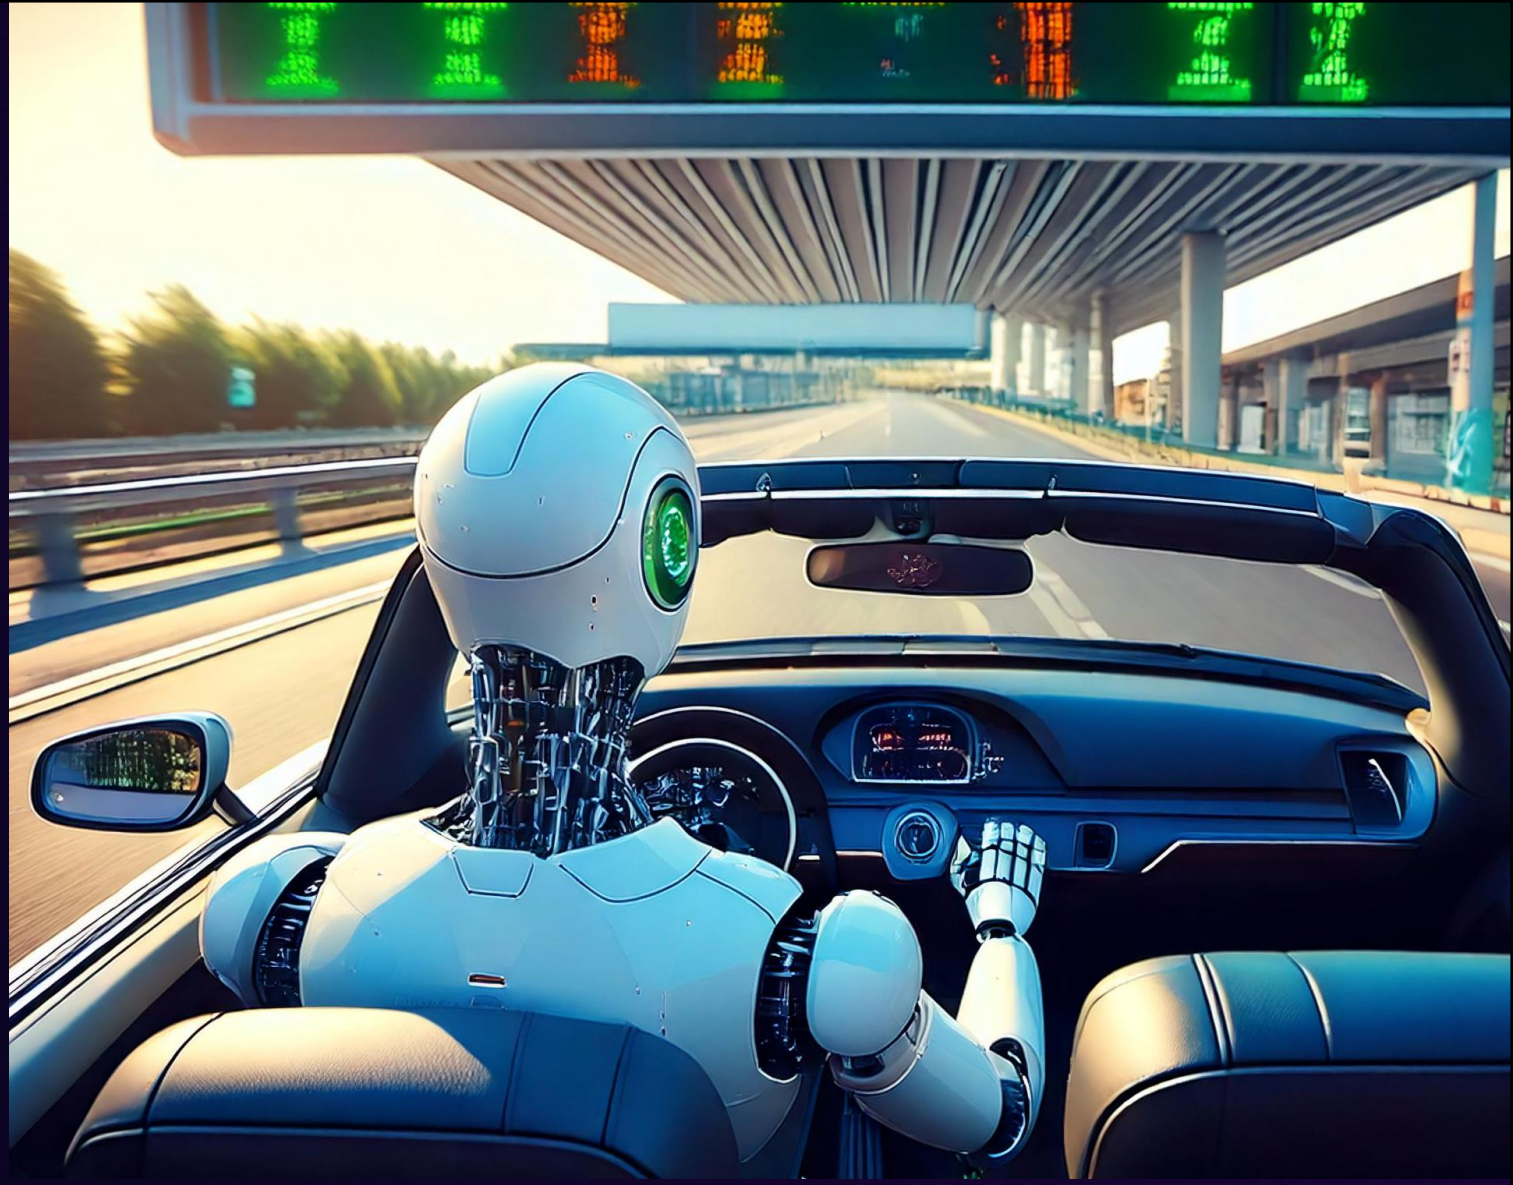

Example

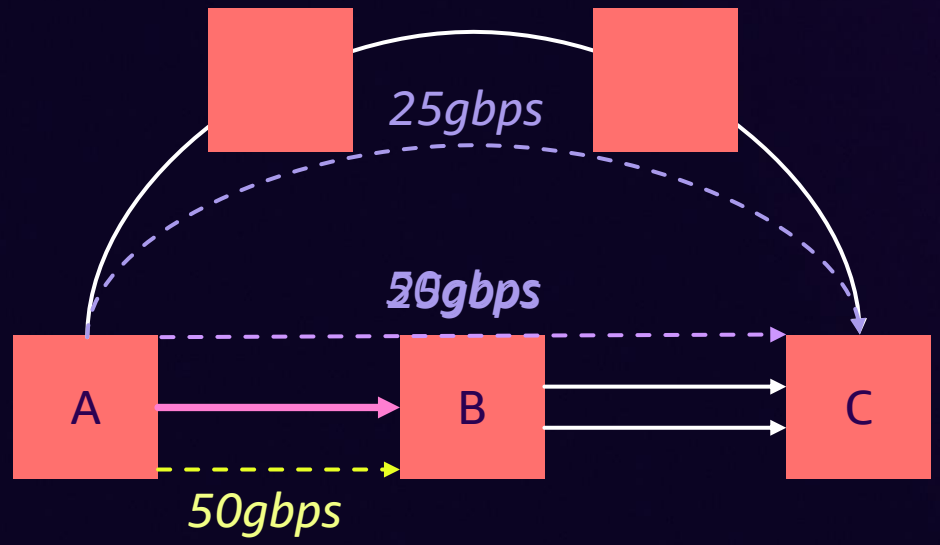

Each link: 10ms, 75gbps

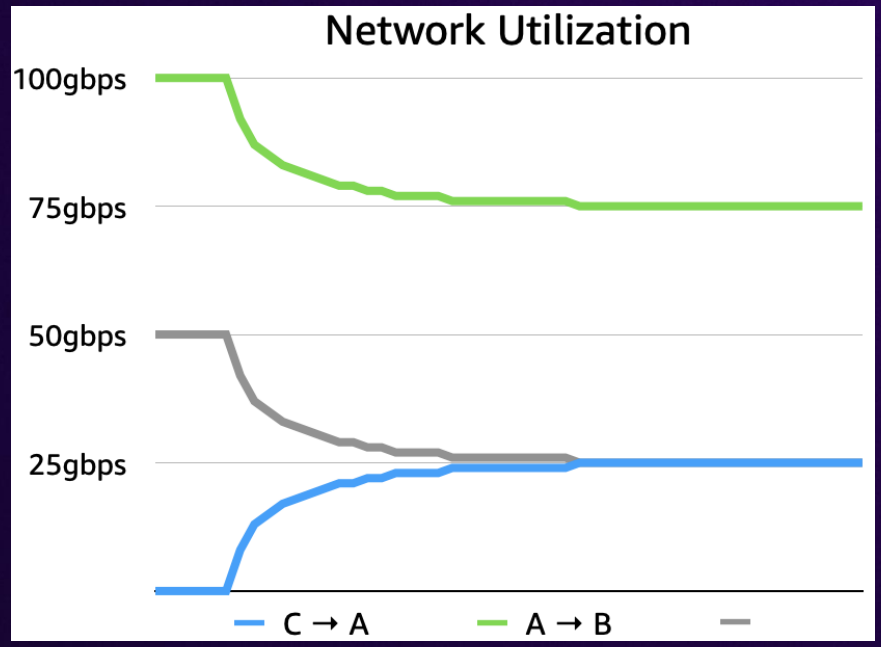

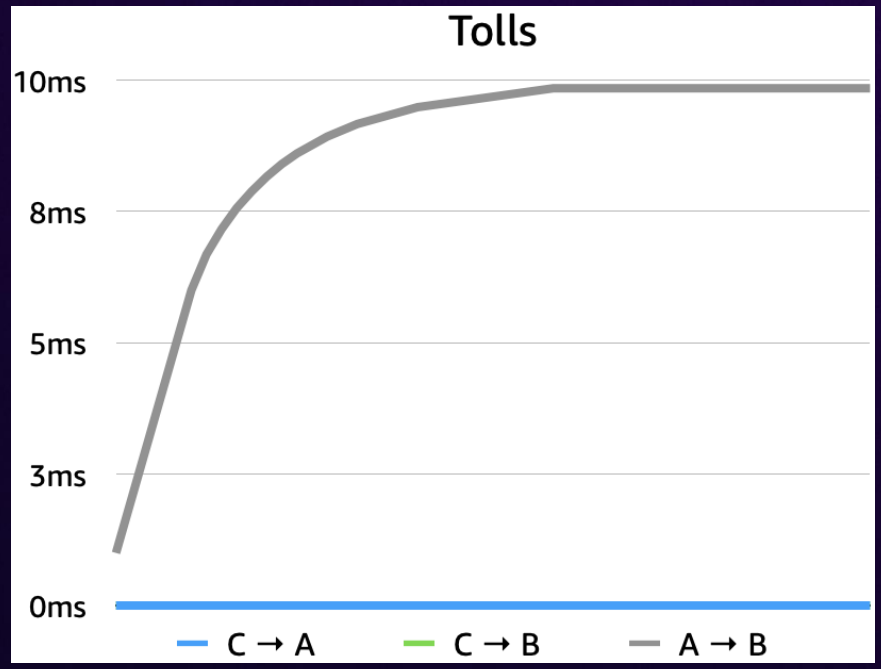

Demand-driven innovation

Purpose-built

Incremental

Ownership

(And tiny plastic connectors)

The full hardware package

Racks

- Amazon rack 1
- Similar to outpost - AWS Expo
- Standardized building blocks

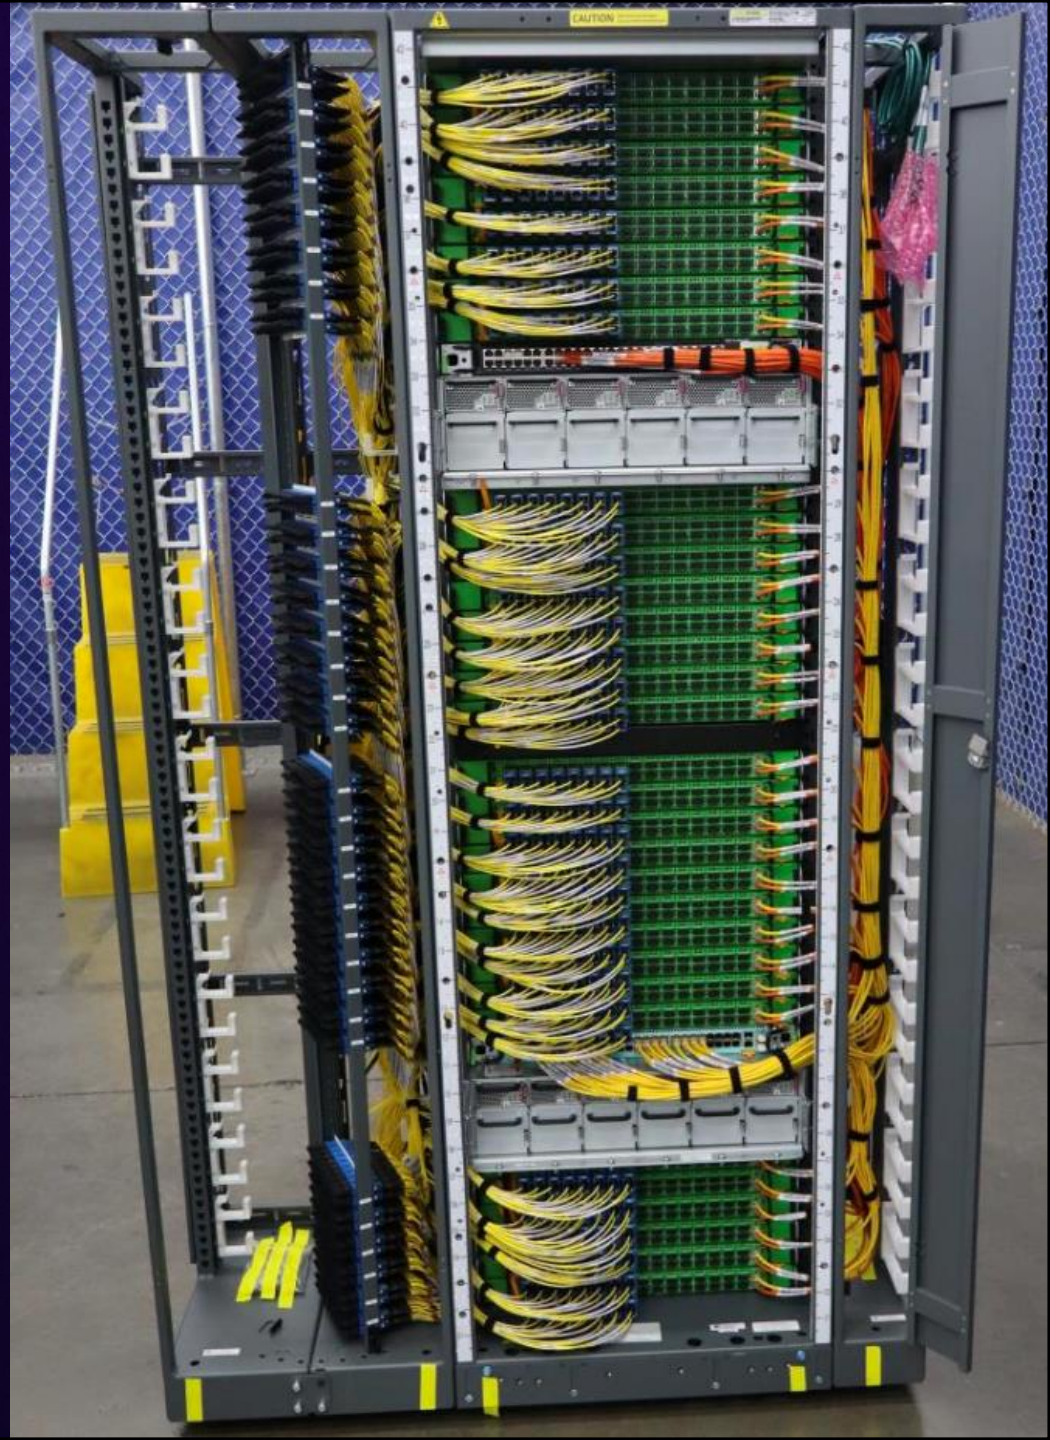

Cabling

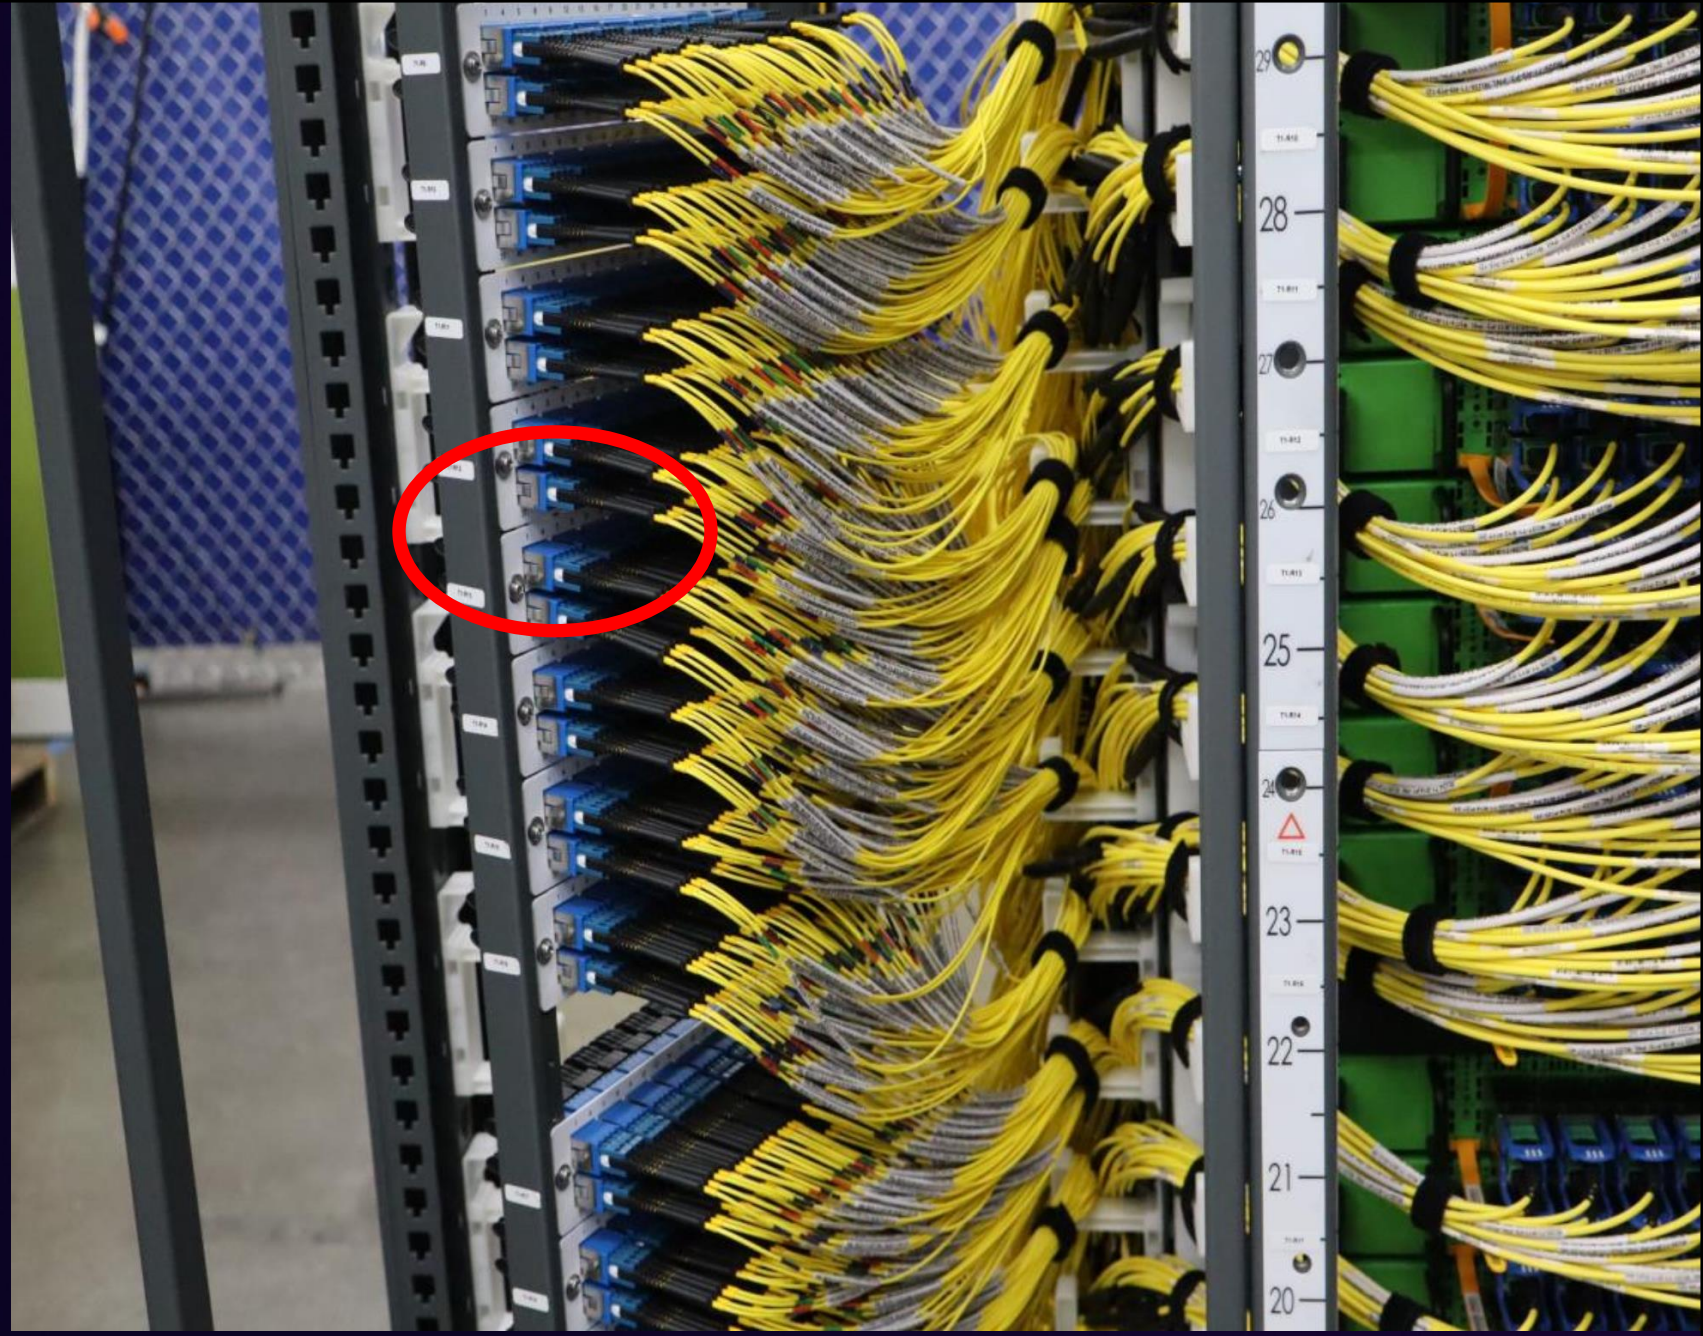

Cabling

- Pre-wired positions
- 144-fiber trunk cables
- SN to MPO connectors
- Network <> Host racks

Connectors

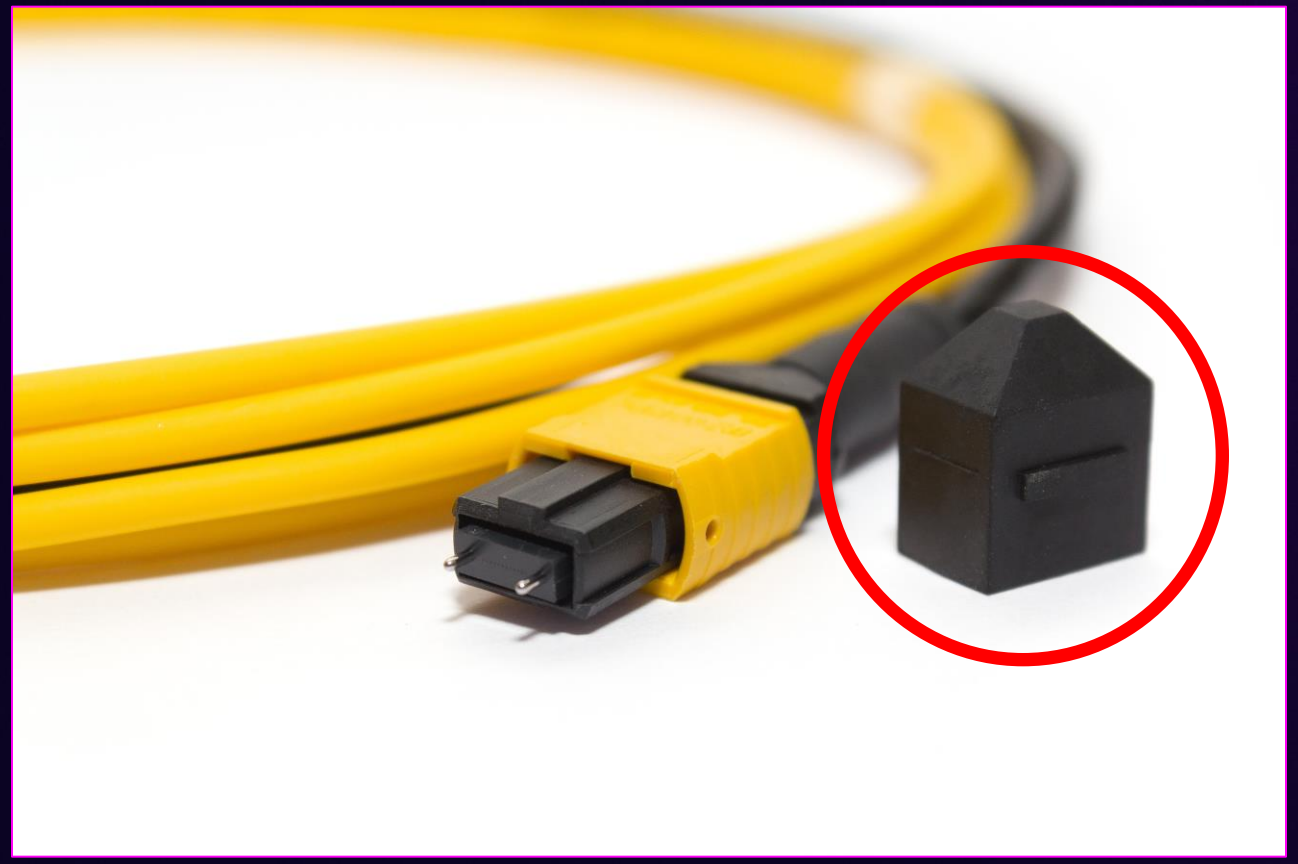

MPO

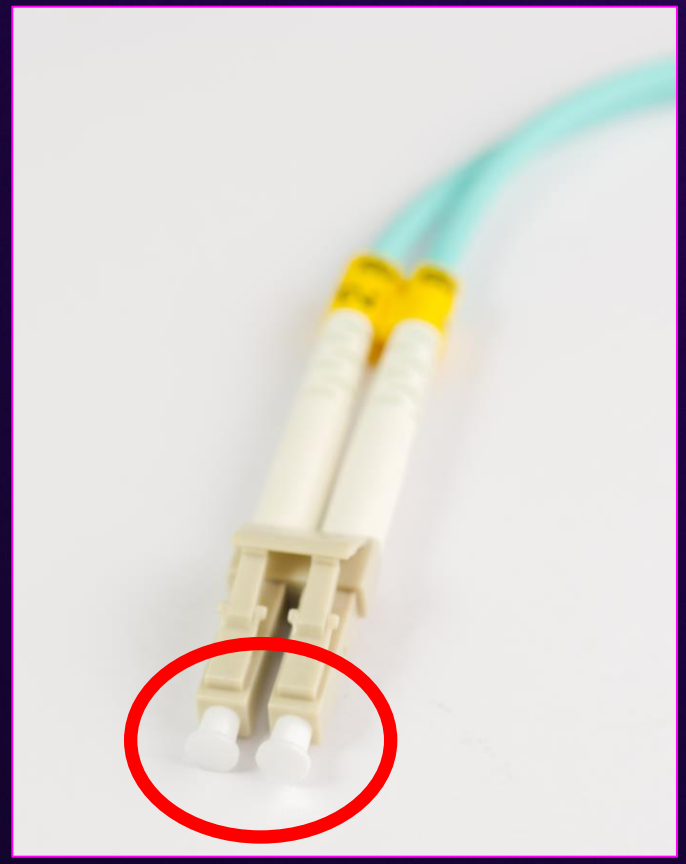

LC

Connectors

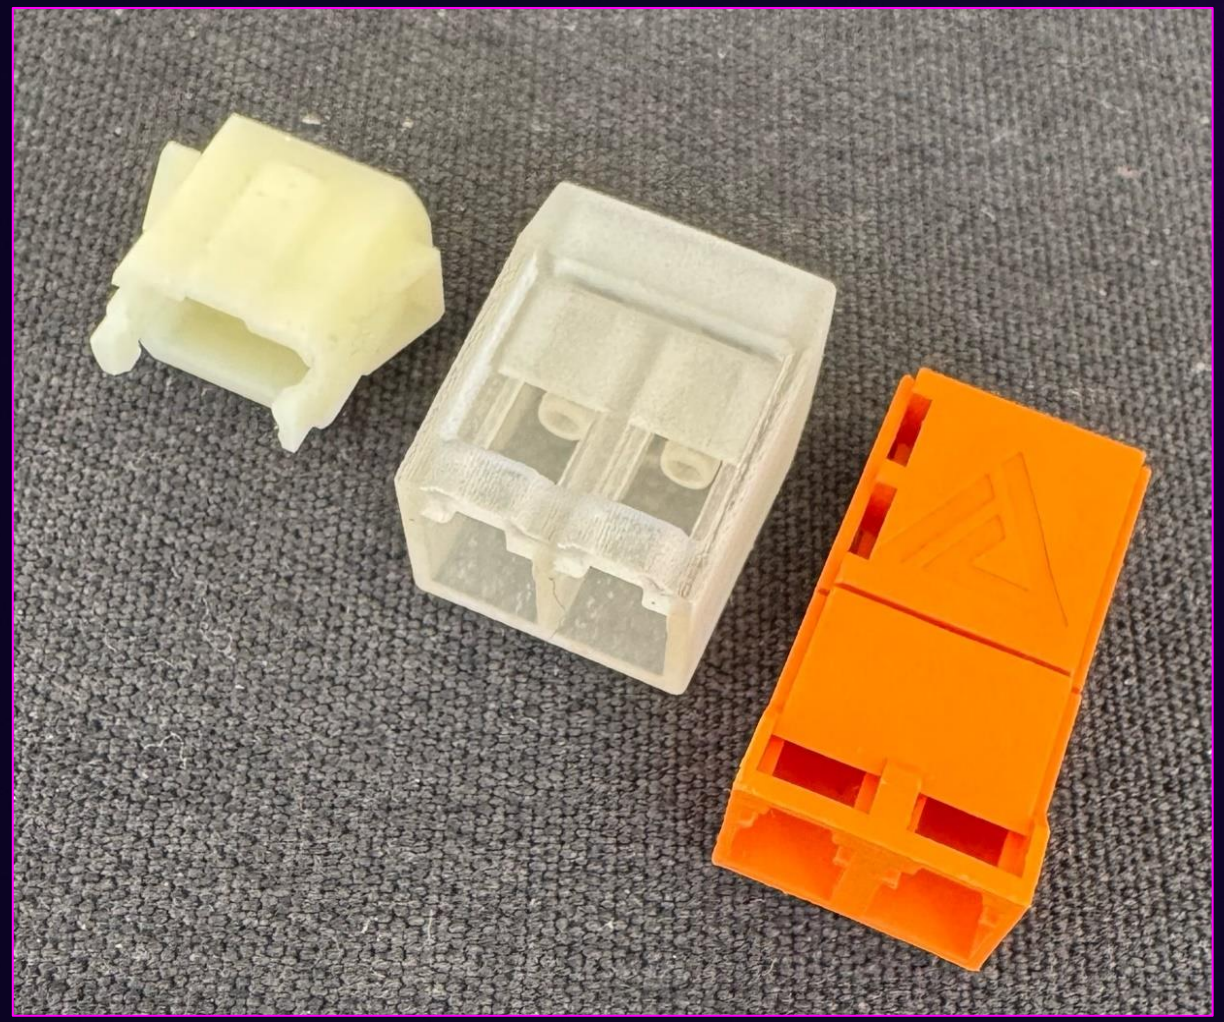

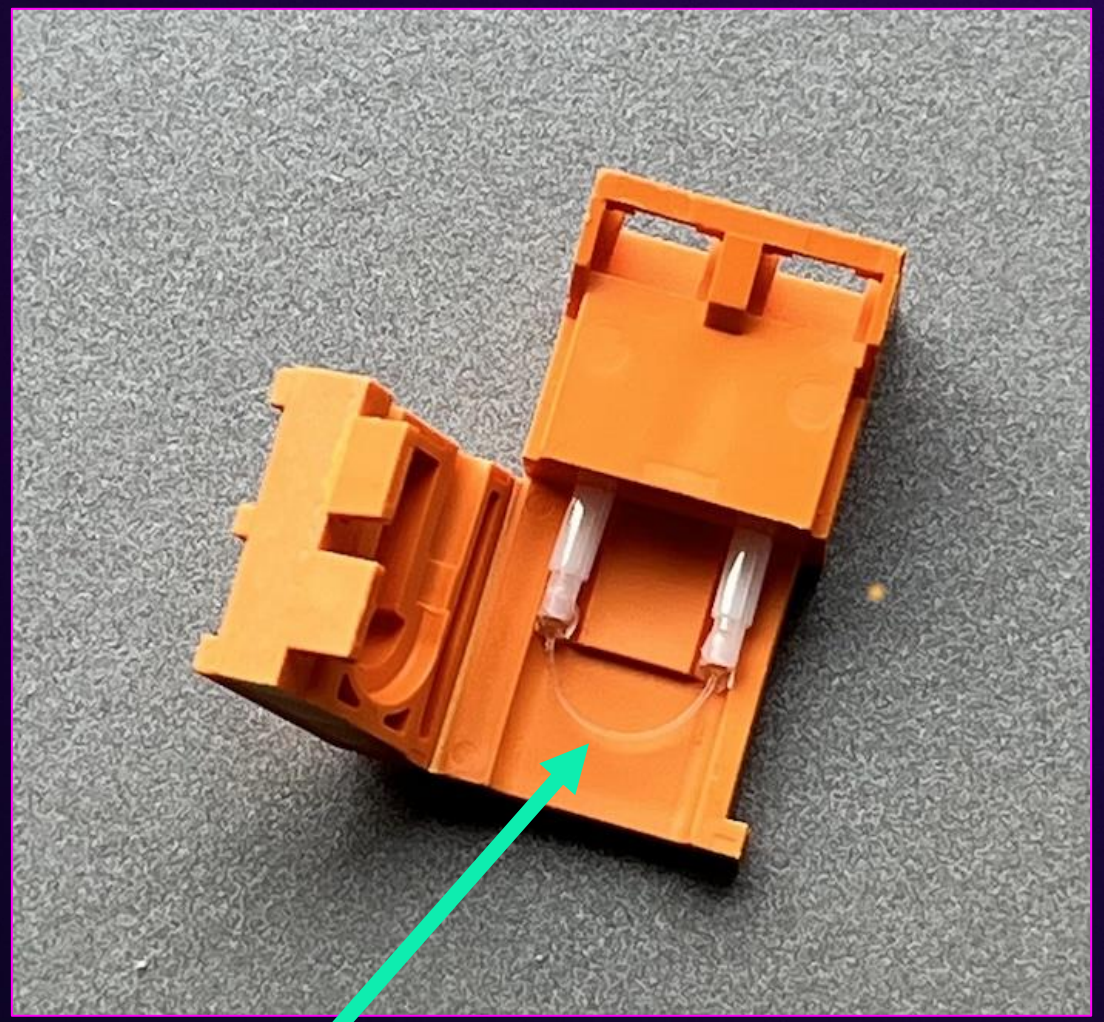

Monitoring

Telemetry

Active probing

Correlation

Remediation

Monitoring

PASSIVE TELEMETRY

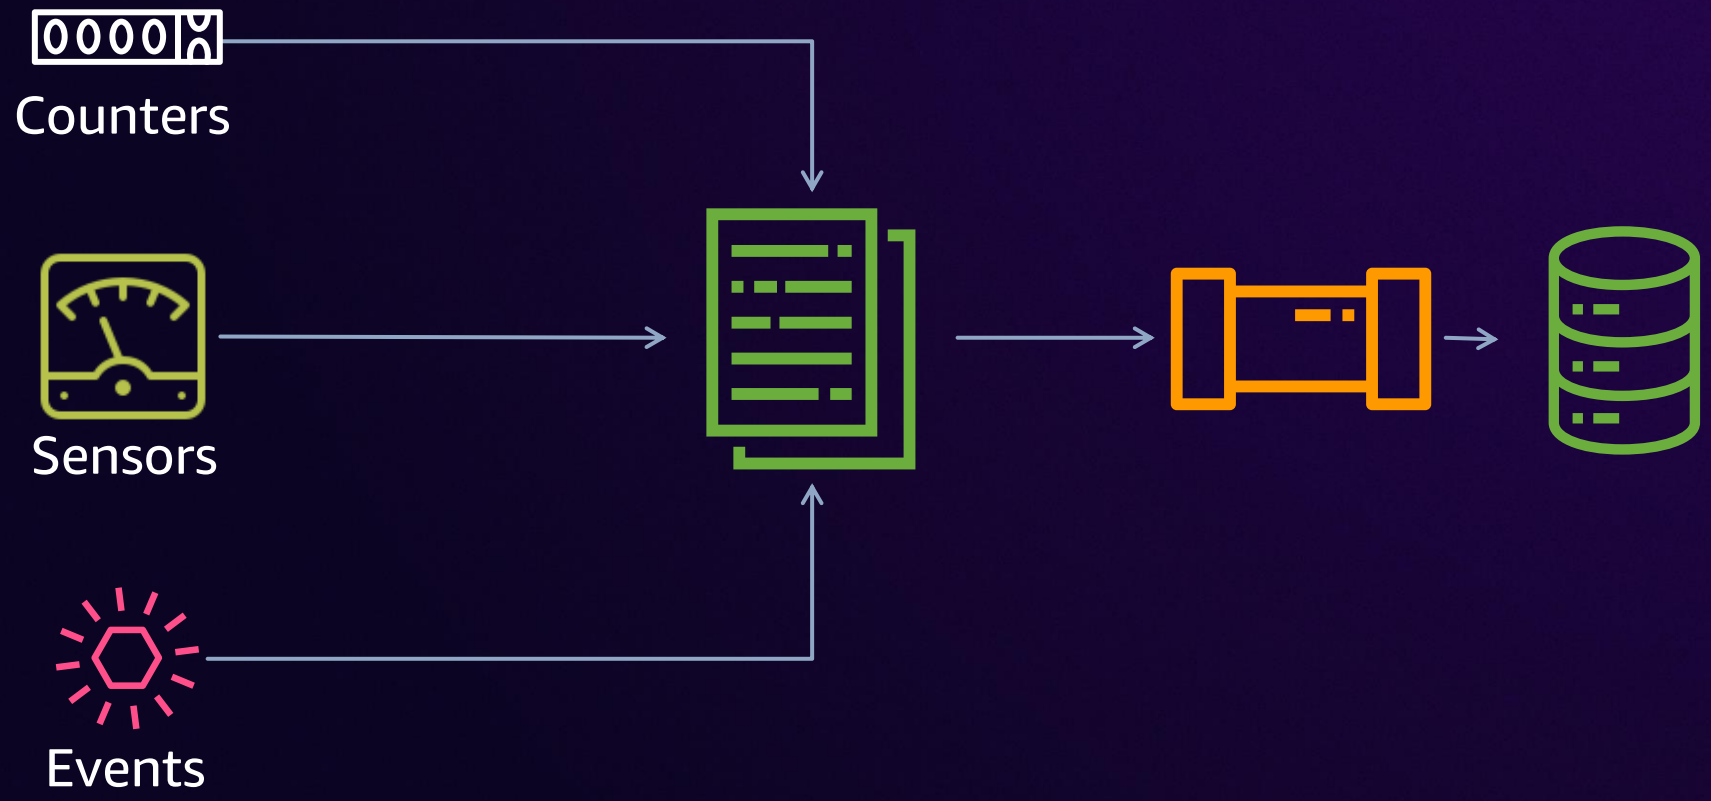

7 billion observations per minute

Monitoring

ACTIVE MONITORING

>1.5B probes per minute

Monitoring

ACTIVE MONITORING

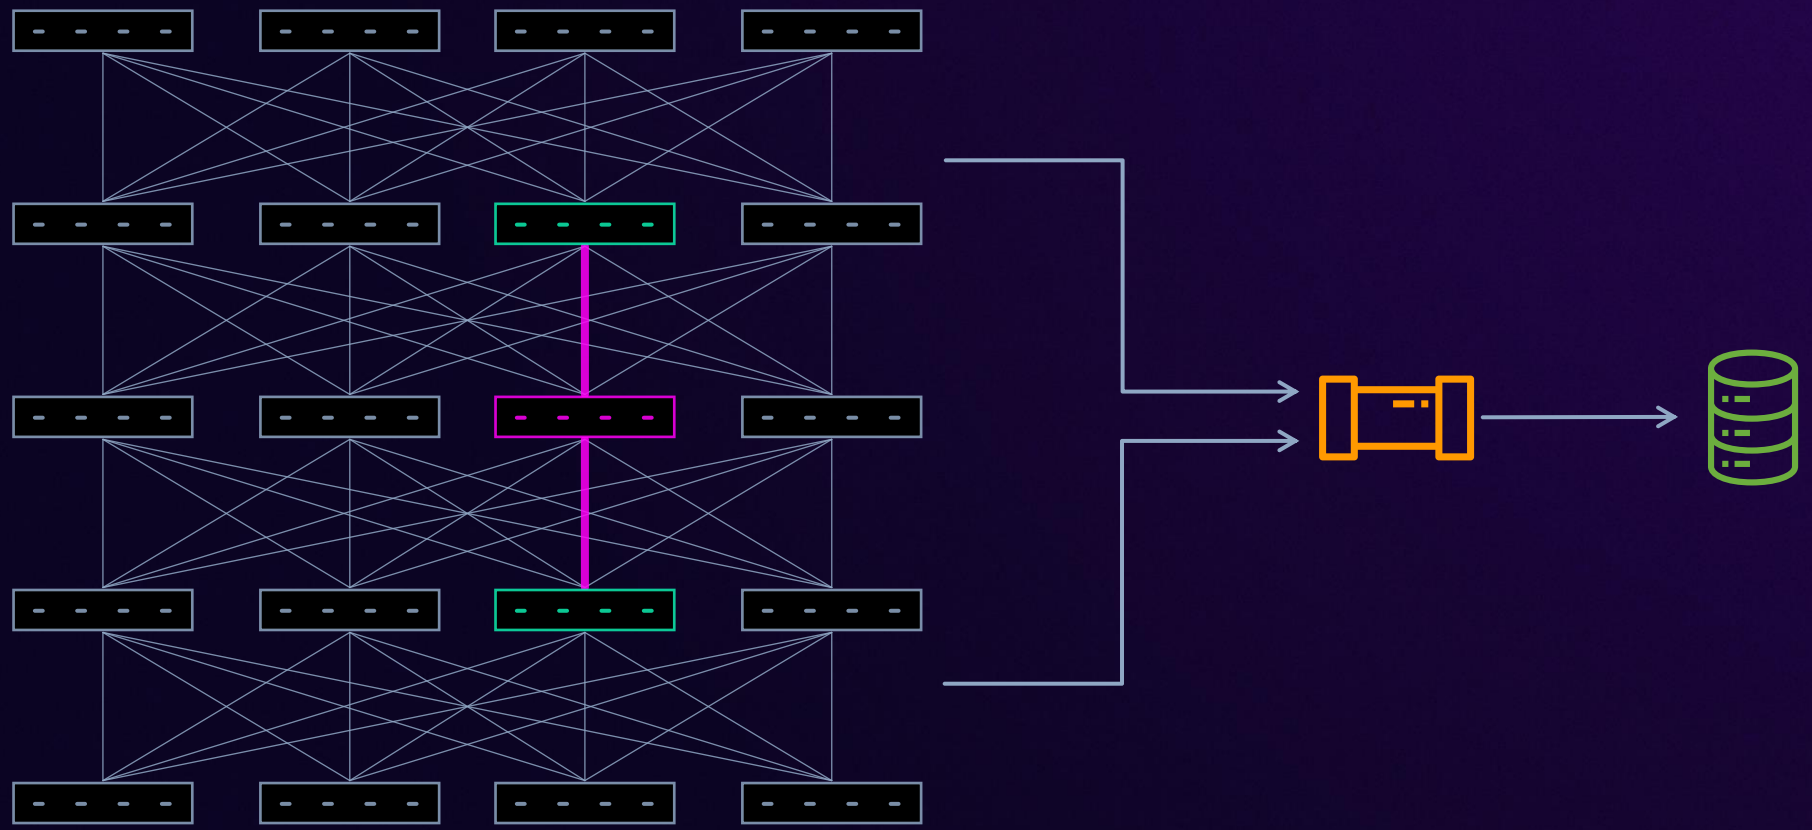

Monitoring

ACTIVE MONITORING

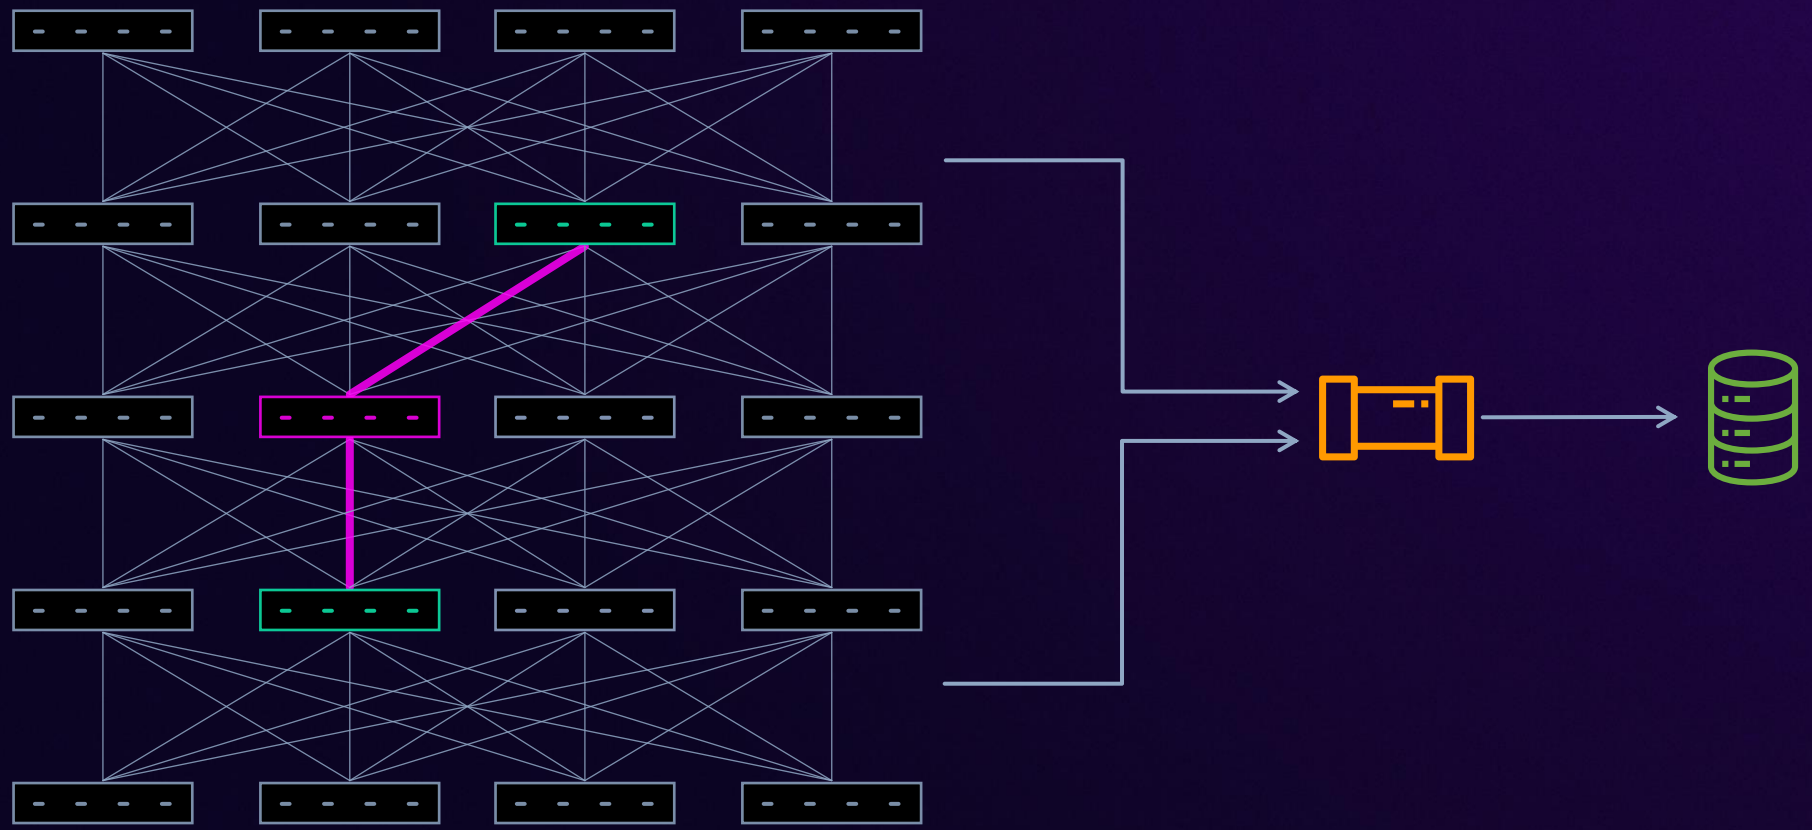

Monitoring

25B alerts per year

>1B actionable

>3M distinct events

98% fully-automated actions

2% escalated

What’s next?

Power and connectivity

Space and power

AWS

AWS plans to invest $
Mississippi, the larges
investment in the stat

Read more

SUSTAINABILITY

How Amazon is supporting farmers through solar and wind farm investments

Read more

September 18, 2024

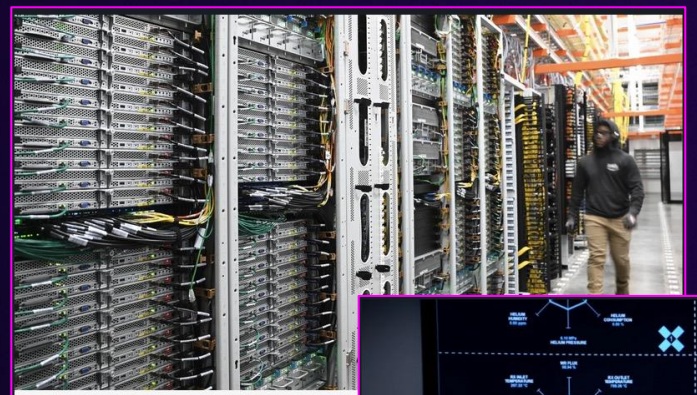

AWS

AWS plans to invest
Indiana, the largest
in the state's history

Read more

SUSTAINABILITY

Amazon signs agreements for innovative nuclear energy projects to address growing energy demands

Read more

October 16, 2024

Space and time

300,000 km per second

1 light year
= 9 quadrillion km

Latency

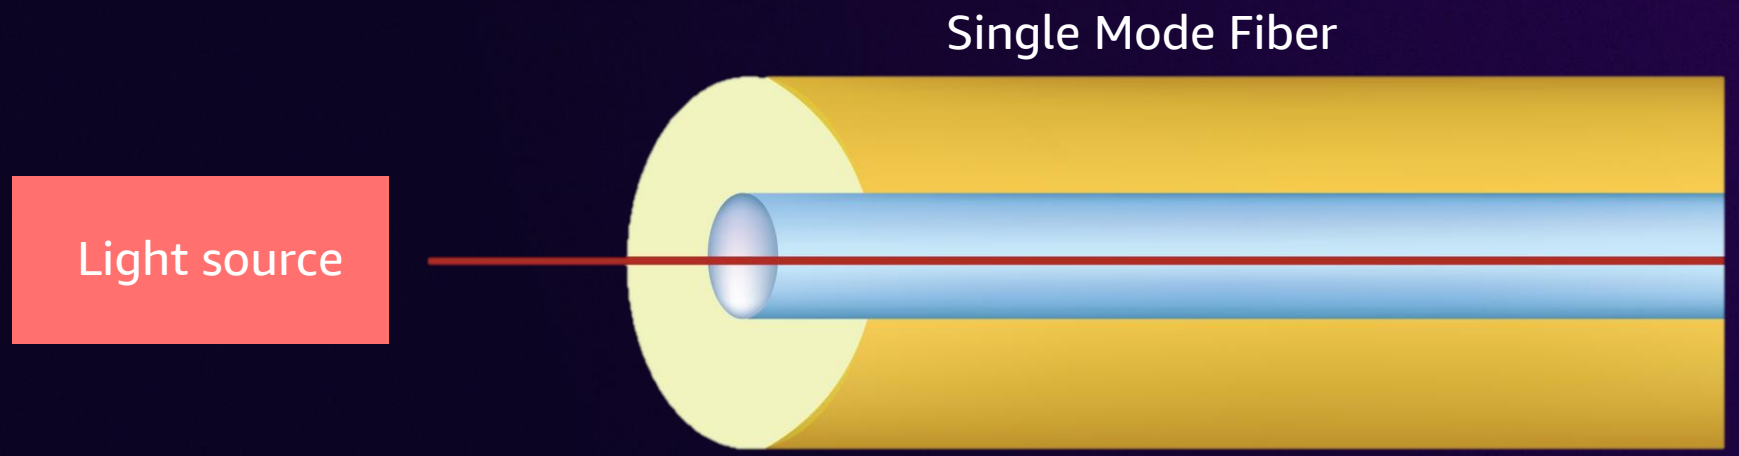

Index of refraction: 1.44

Speed of light in fiber = 3x10^8 / 1.44 = 2.08x10^8

1m = 5ns

1km = 5µs

Region growth

IntraAZ latency limit

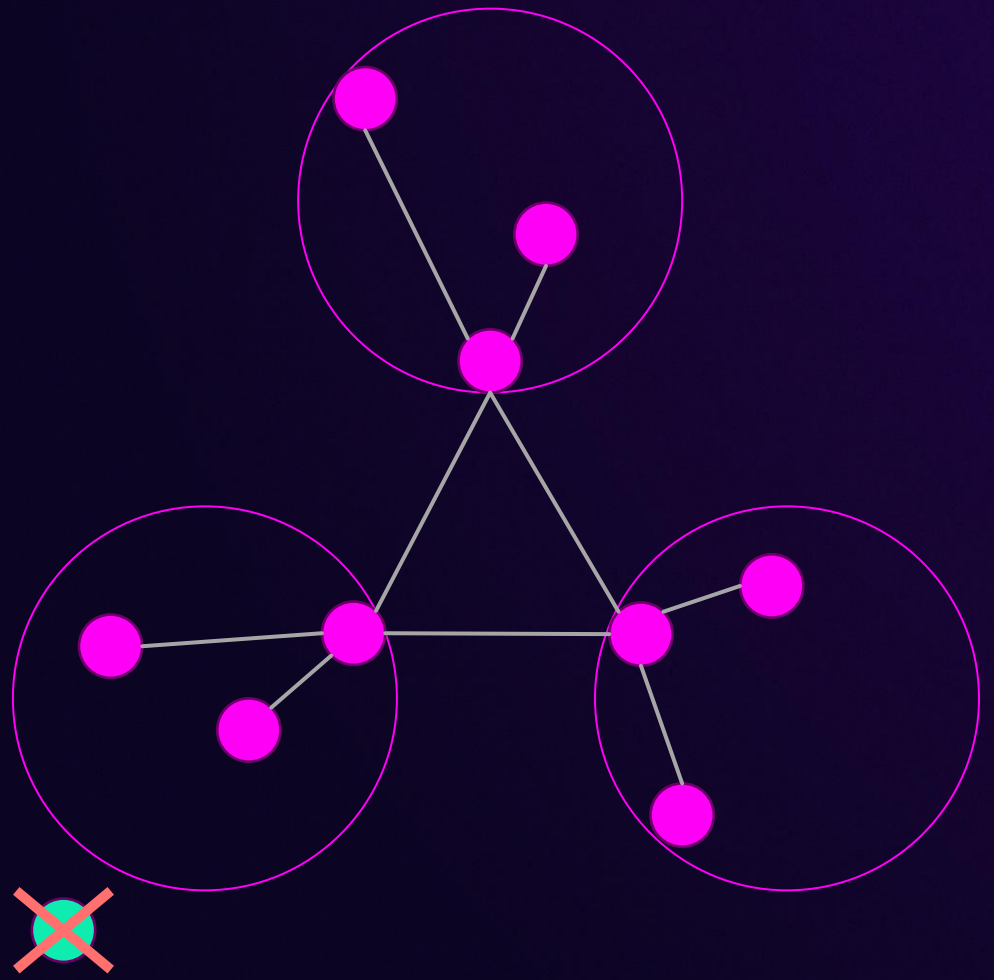

Hollow core fiber

- Replace silica core with air
- 25+ years
- DARPA

Source: DARPA

Hollow core fiber

BENEFITS

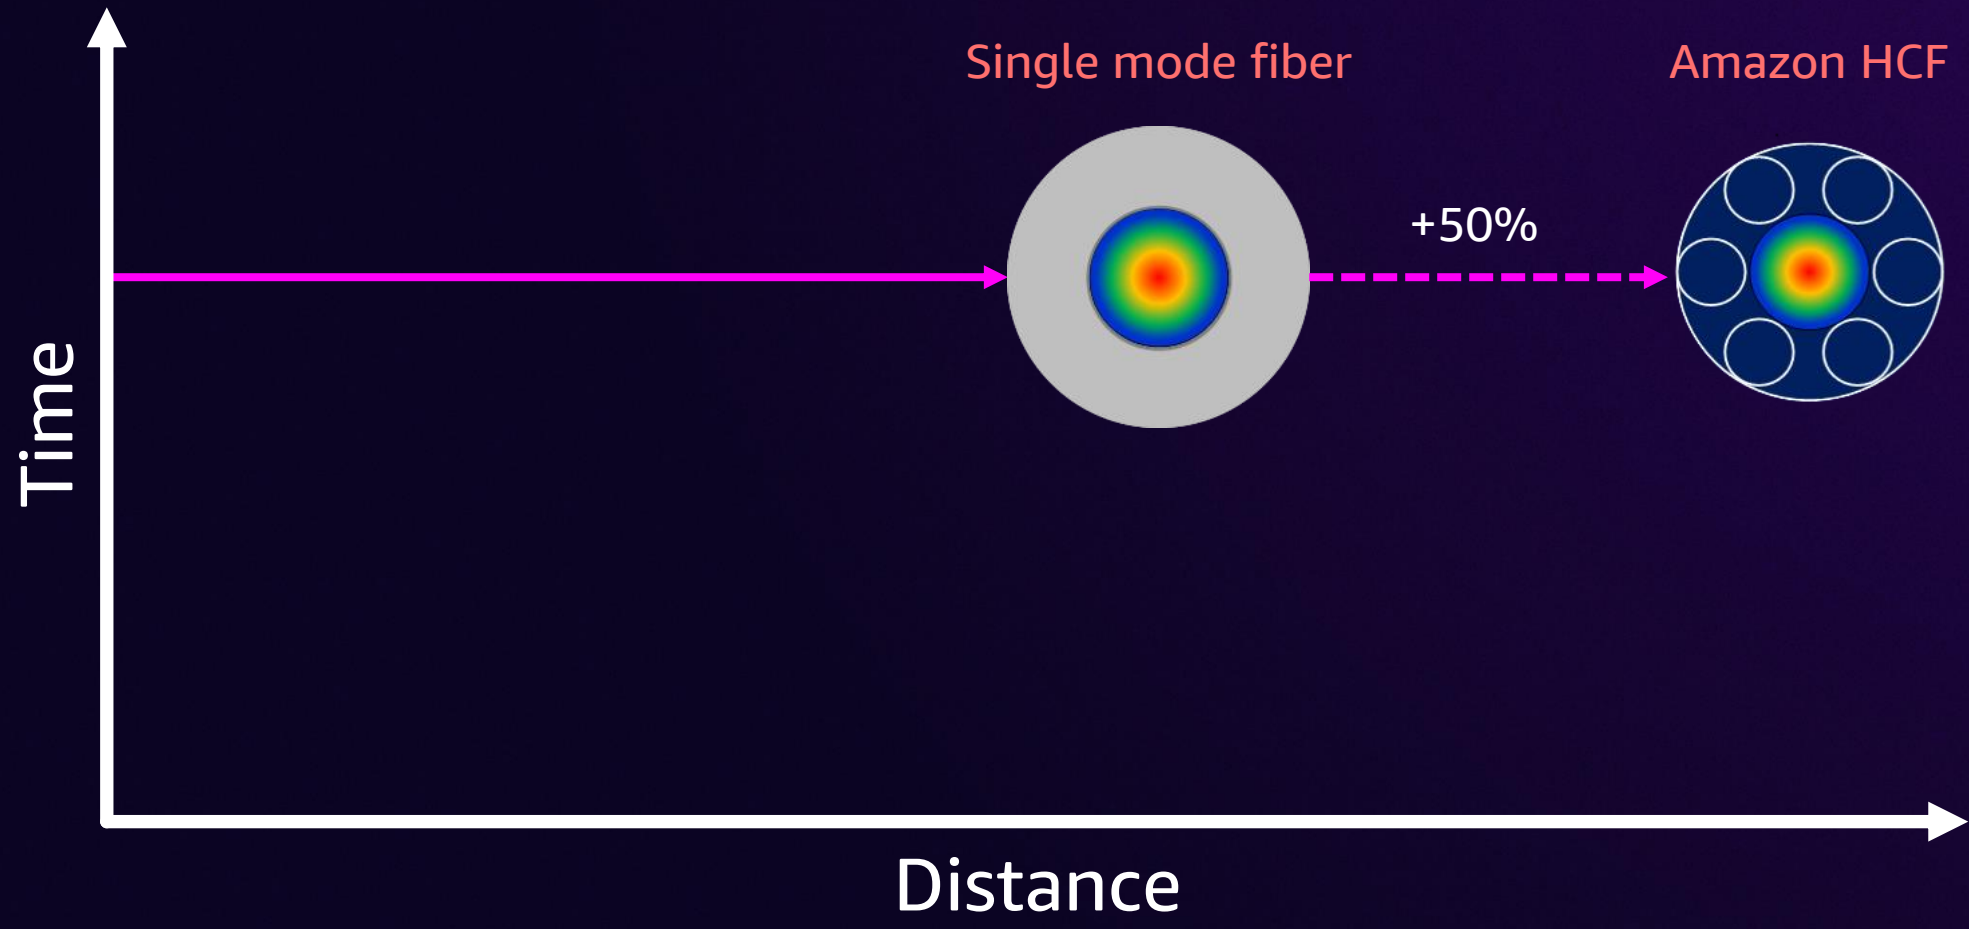

Hollow core fiber

THE AMAZON WAY

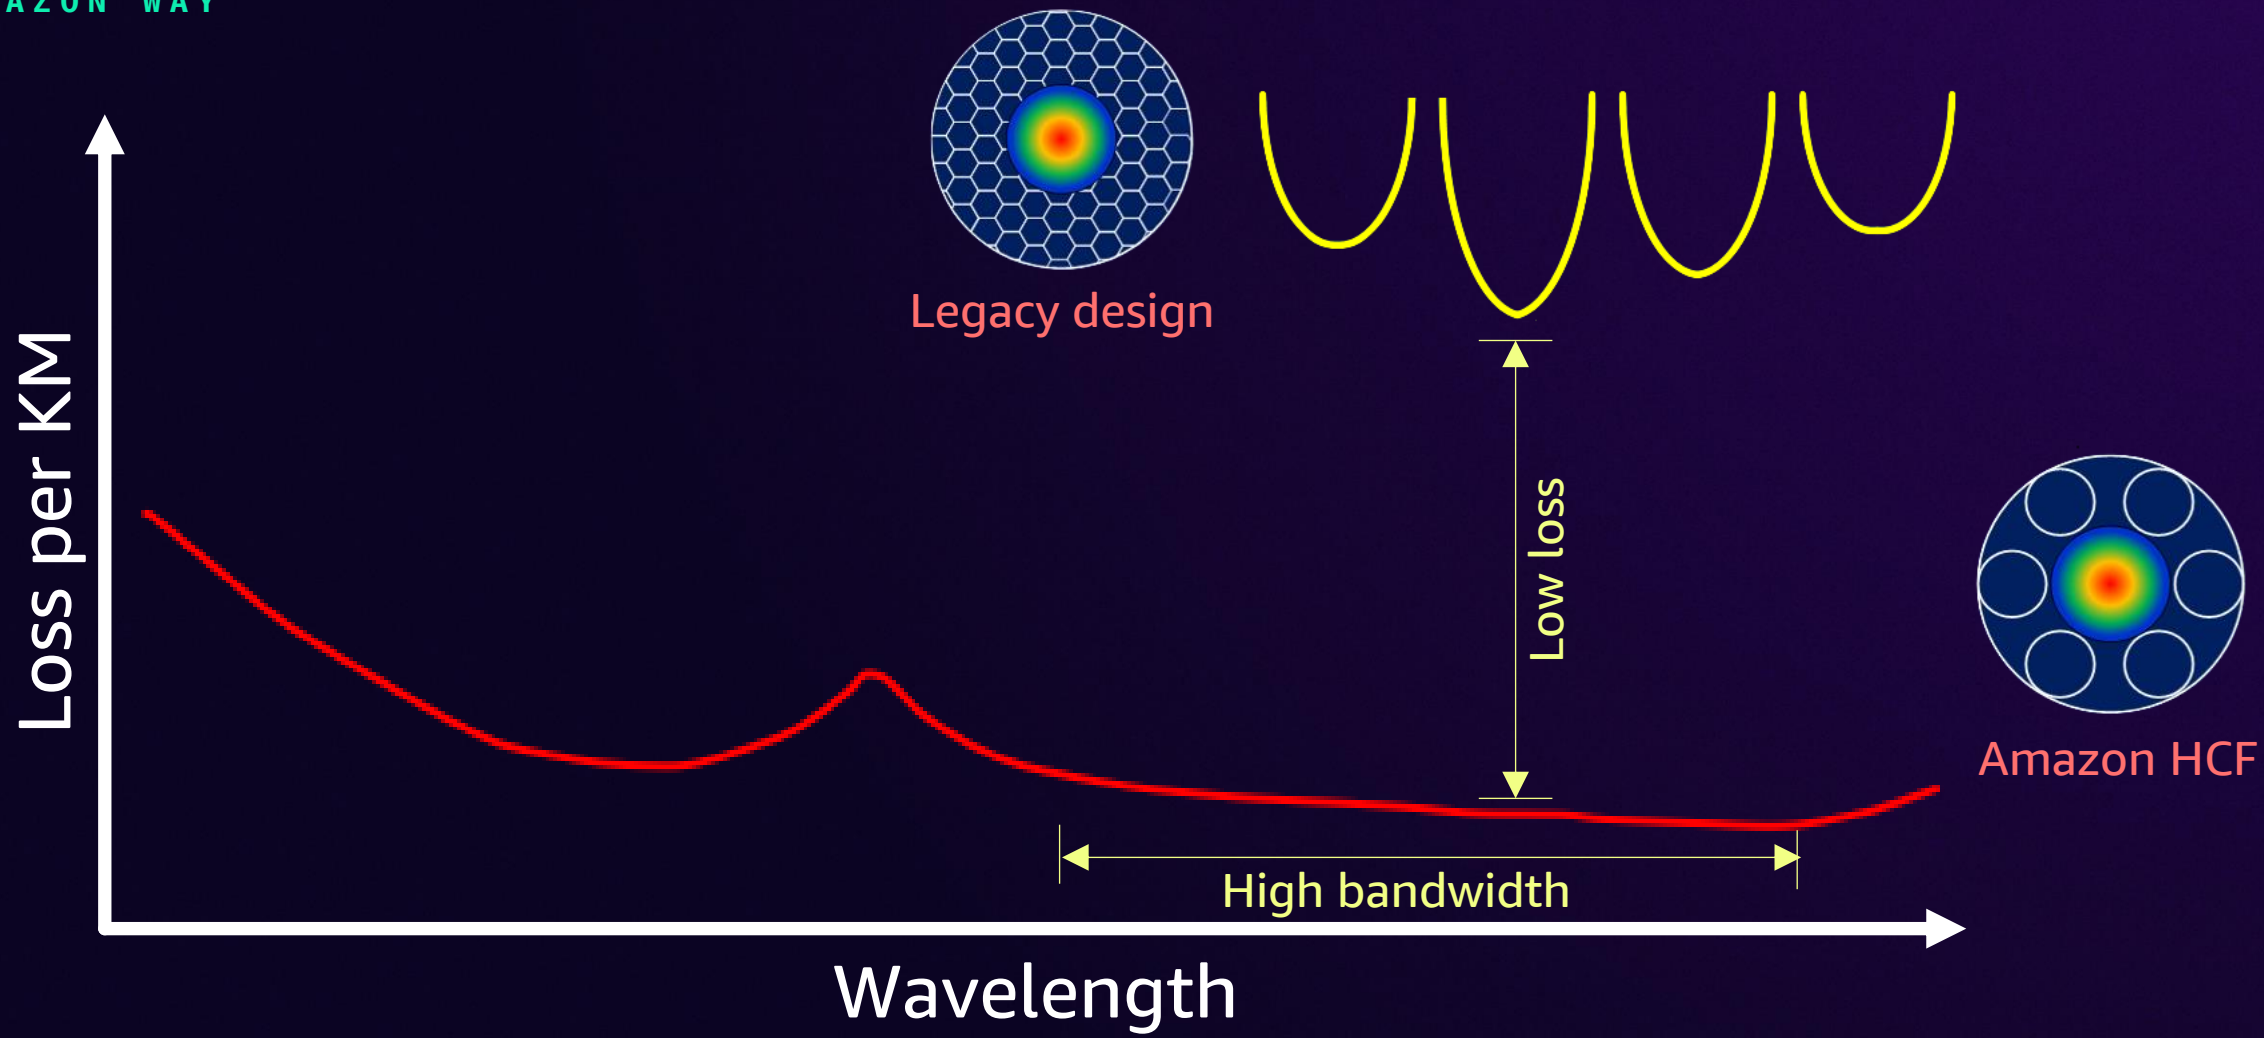

Region growth

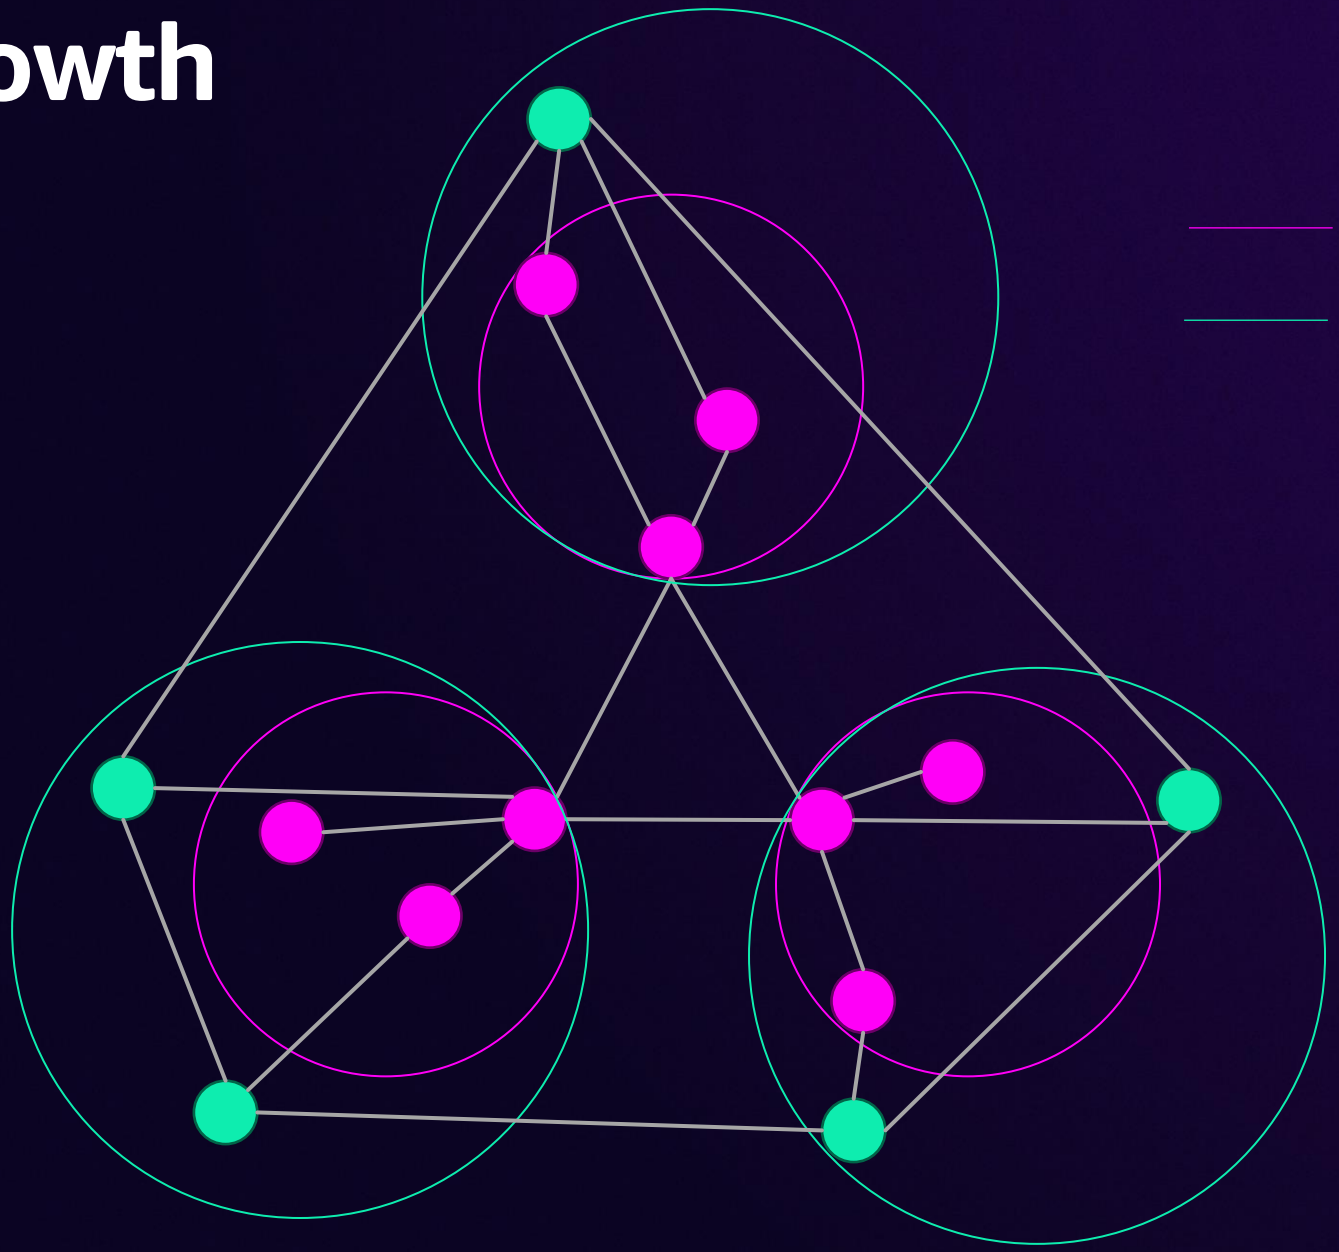

IntraAZ latency limit

IntraAZ latency HCF

1.5x radius
= 2.25x area

Fiber innovations

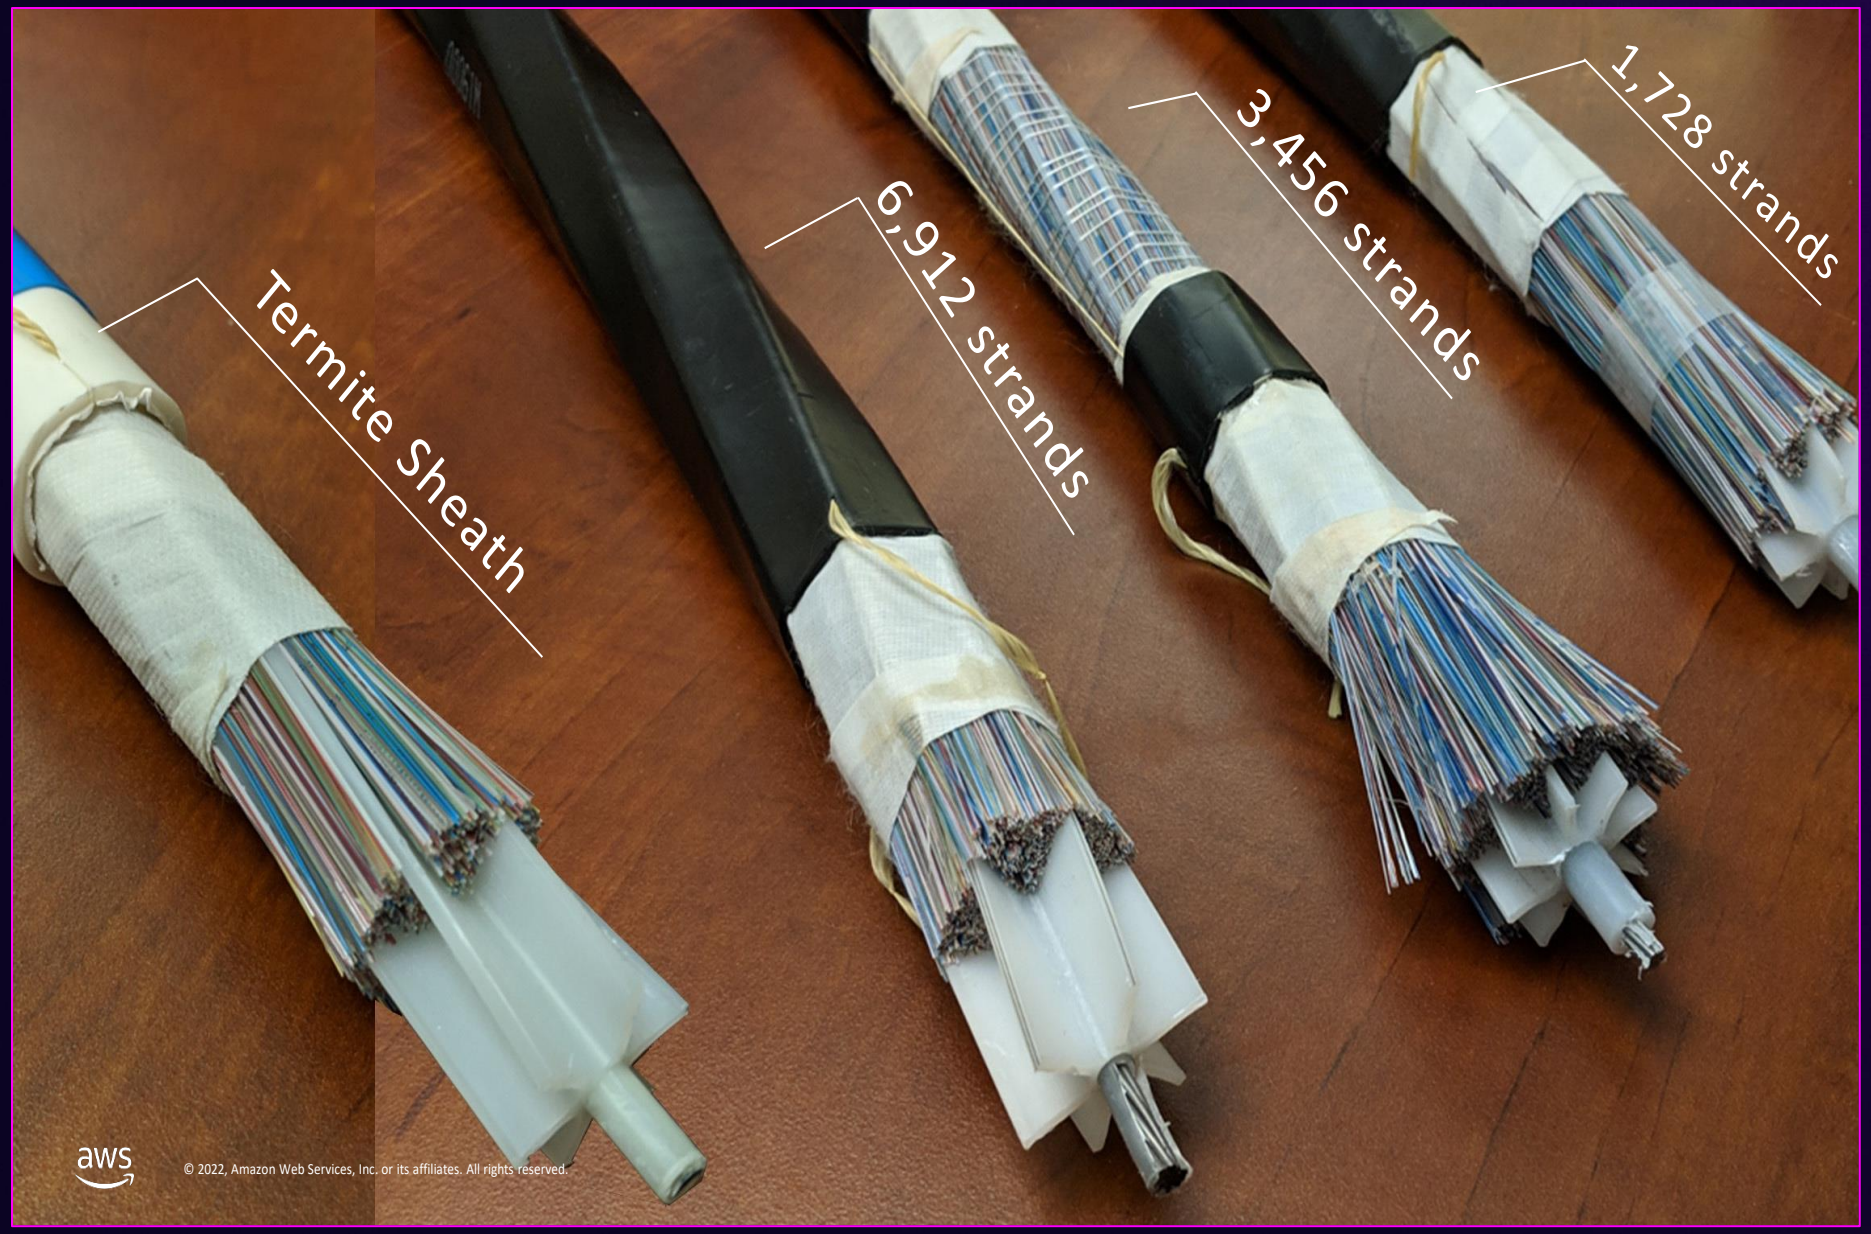

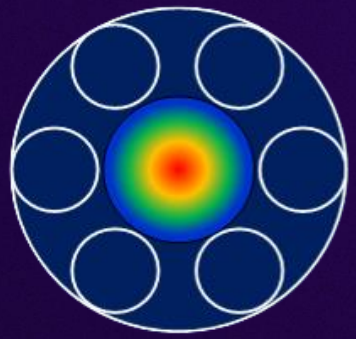

Hollow core

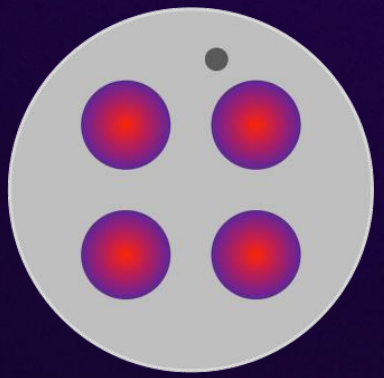

Multi core

Wrap up

Wrap up

- Design
- Ownership
- Constant Improvements
- Tiny plastic connectors
- Future challenges

Survey

Thank you!

Please complete the session survey in the mobile app

Stephen Callaghan
Senior Principal Technologist
Amazon Infrastructure

Jorge Vasquez
Senior Principal Technologist
Amazon Infrastructure