

The background features a dark blue gradient with abstract geometric shapes. On the left, a large triangle is formed by a vertical orange line and a diagonal orange line. On the right, a large curved shape in shades of orange and red sweeps across the frame. The text is centered in the upper right area.

AWS re:Invent

NOV. 29 – DEC. 3, 2021 | LAS VEGAS, NV

NFX301

How Netflix is using IPv6 to enable hyperscale networking

Donavan Fritz (he/him)

Sr Network SRE

Netflix



Agenda


Why IPv6


Co-innovation between Netflix and AWS

Netflix progress

Lessons learned and best practices

How to get started and how to show IPv6 is worthwhile







AWS Cloud



us-west-2



Instances



Instances



Instances



us-east-1



Instances



Instances



Instances



eu-west-1



Instances

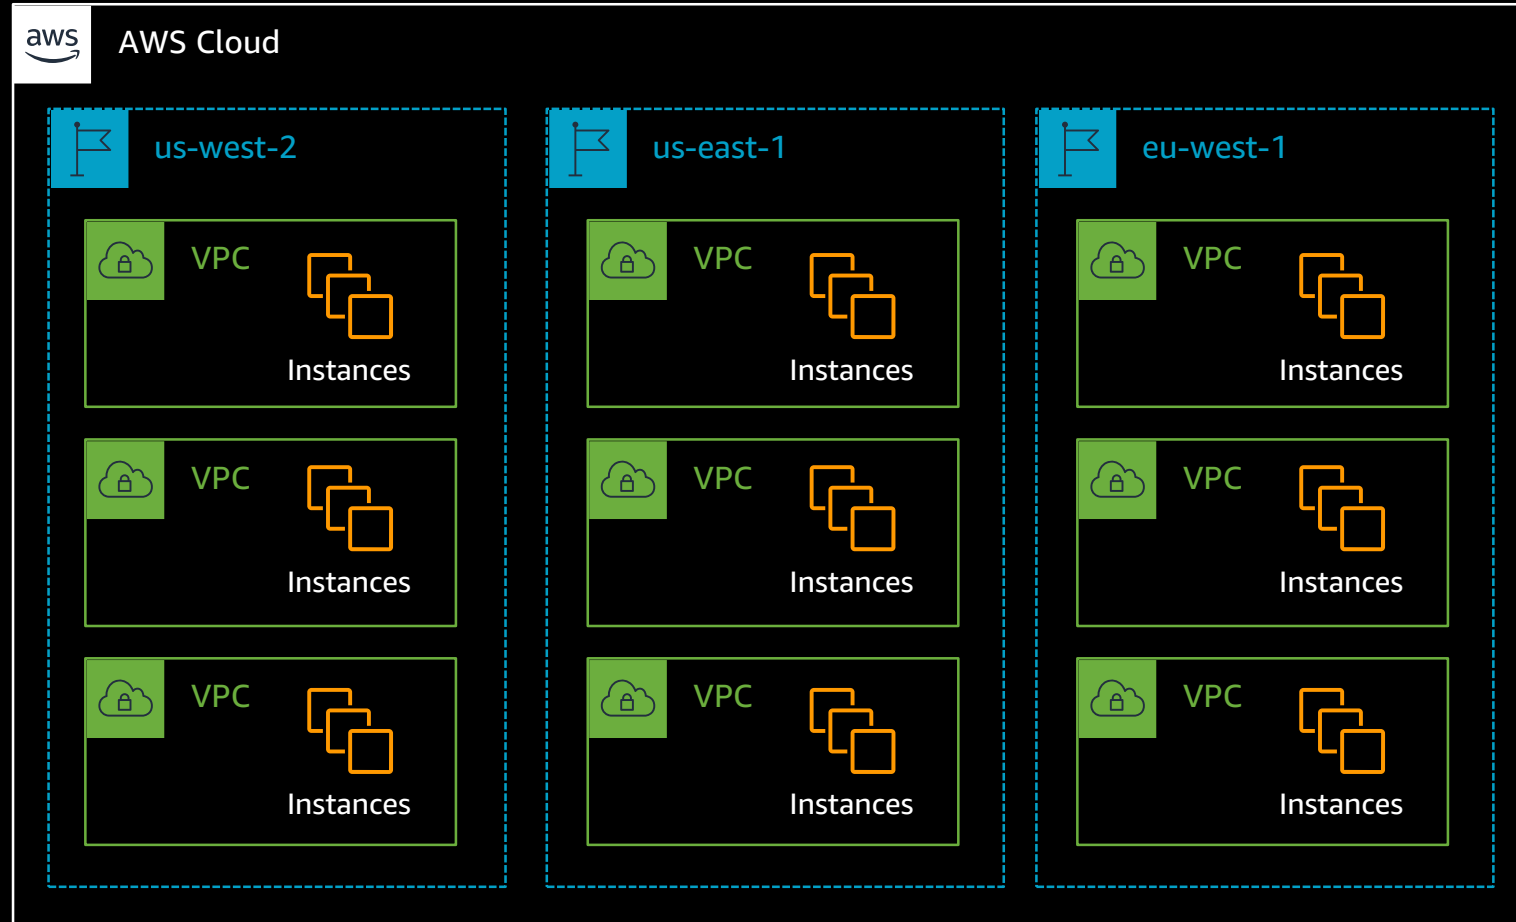


Instances

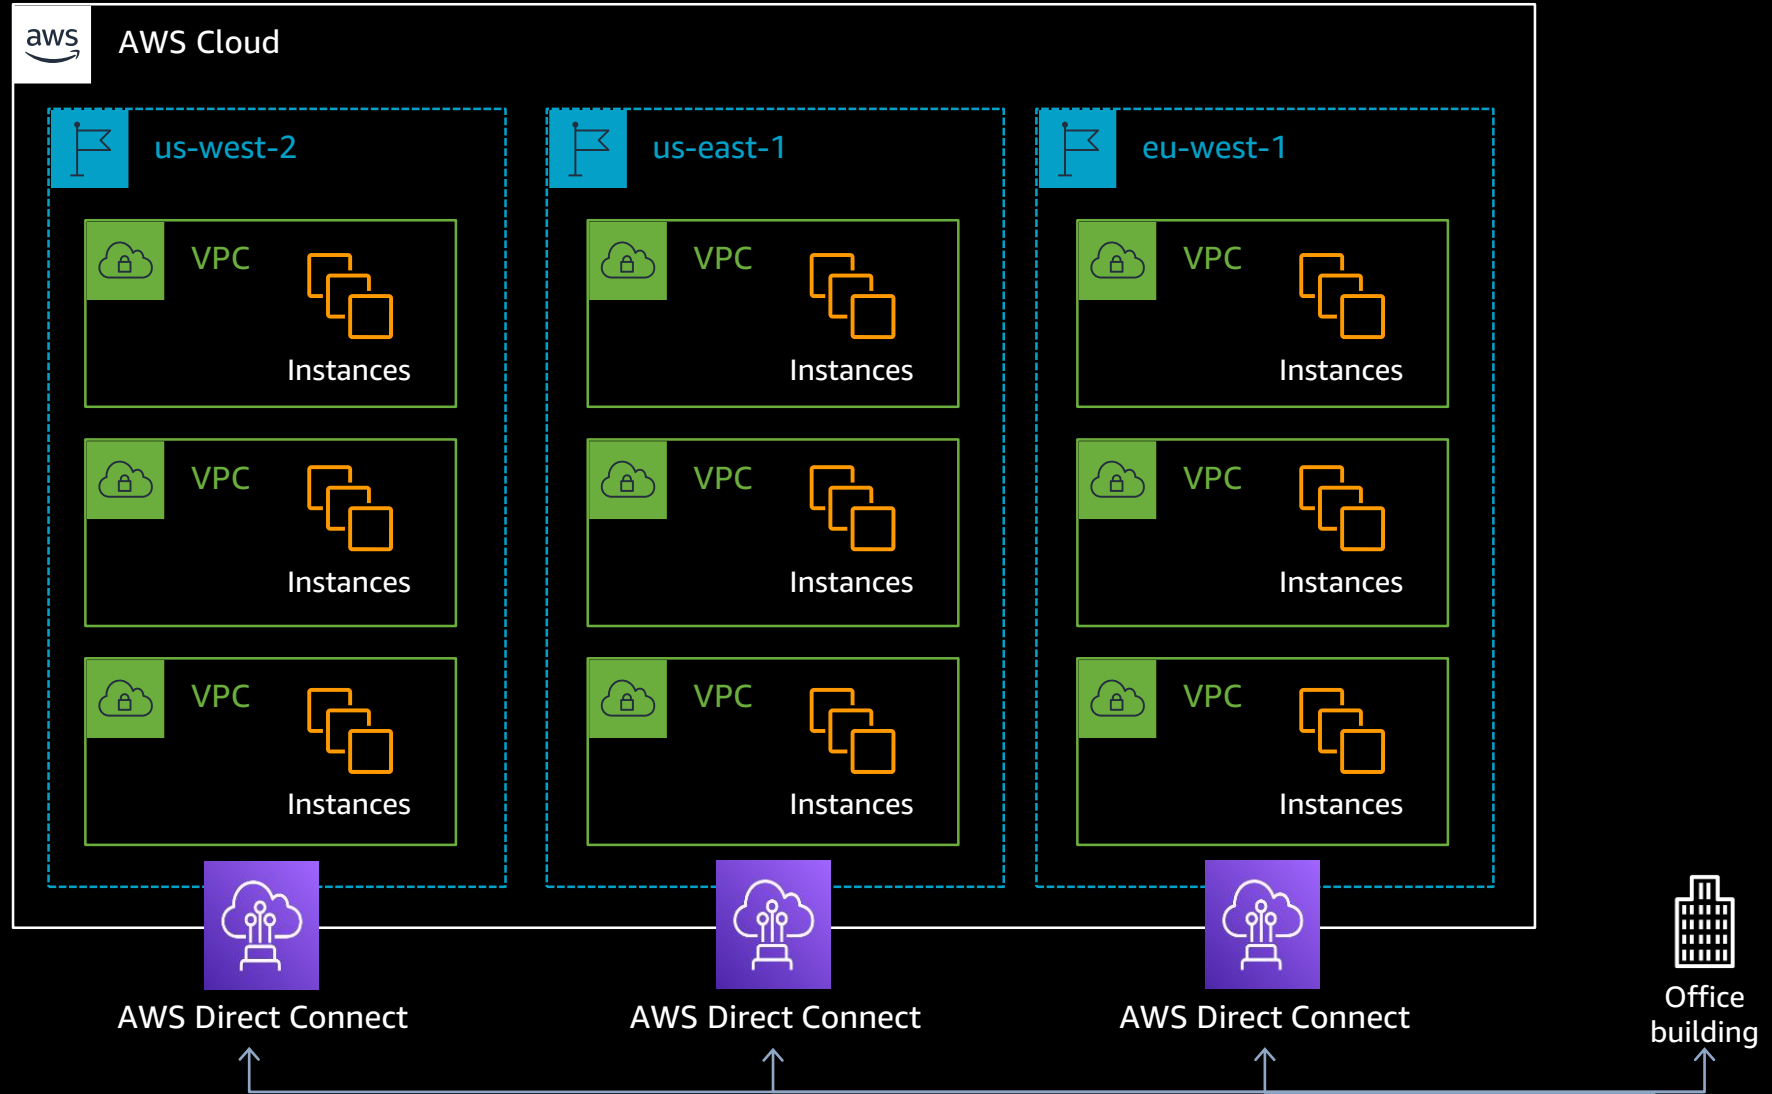


Instances

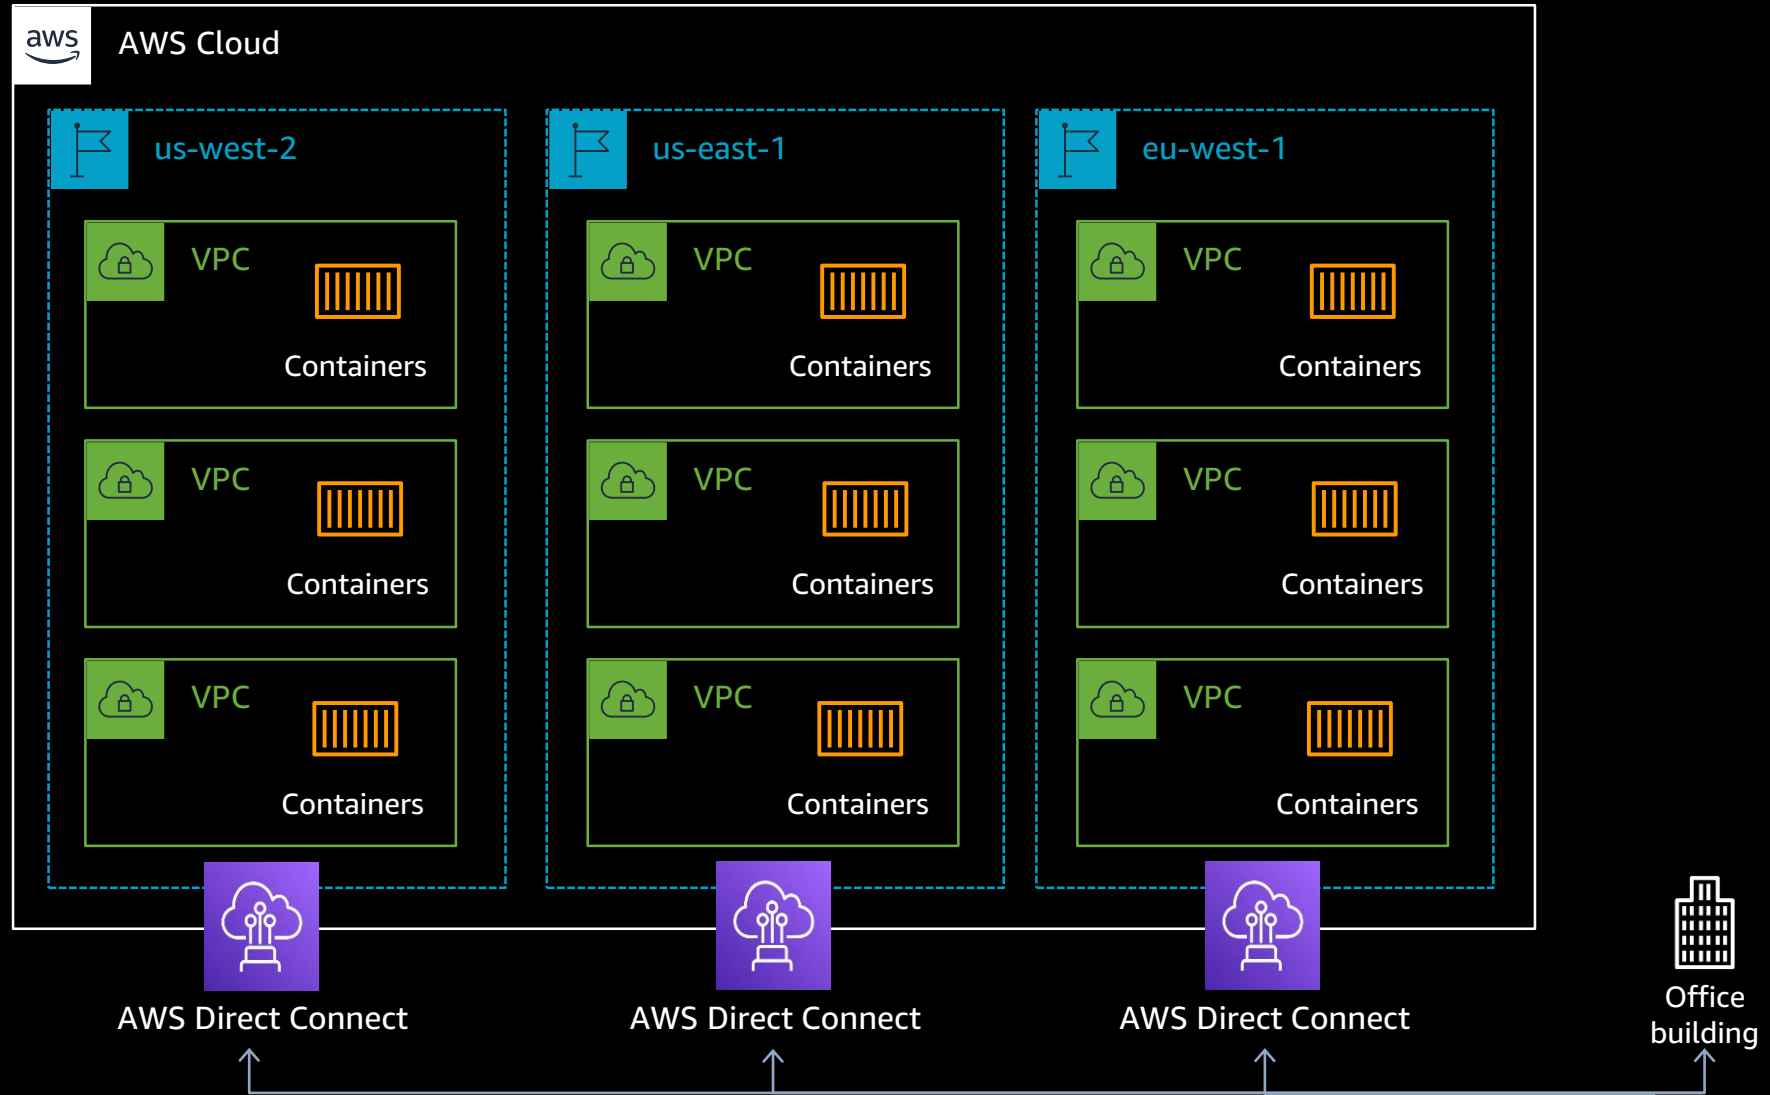
Flat network



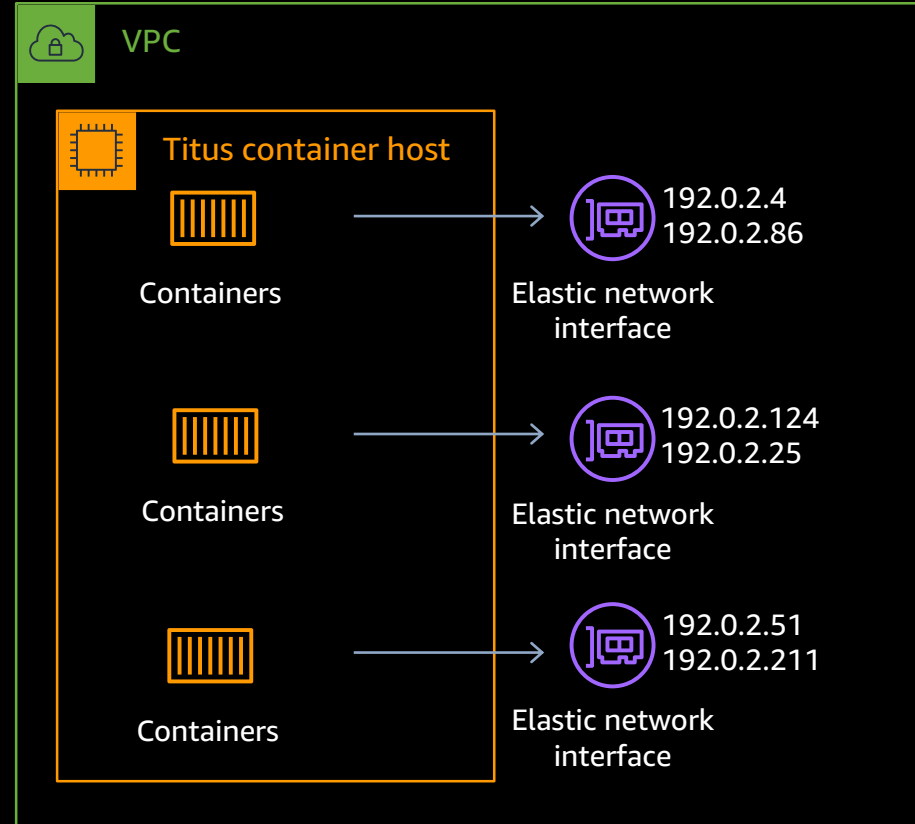

Flat network



Flat network and now, containers



Flat network and now, containers



Where we have been

Flat network

- Carry over from Amazon EC2 classic
- Lots of client-side load balancing
- Business logic

Containers

- IP per container
- Same network posture as EC2 instances

Where we have been

Flat network

Carry over from Amazon EC2 classic
Lots of client-side load balancing
Business logic

Containers

IP per container
Same network posture as EC2 instances

Where we are going

Continued growth

More accounts, more VPCs

On premises

Studio
Gaming (?)

Technical requirements

Flat network

Containers

Continued growth

On premises

100+ VPCs

Full IP reachability

Technical requirements

Flat network

Containers

Continued growth

On premises

“N” IPs per ENI
Short-lived IPs

Technical requirements

Flat network

ENI density

Continued growth

On premises

“N” IPs per ENI
Short-lived IPs

Technical requirements

Flat network

ENI density

Continued growth

On premises

1000+ (?) VPCs
Full IP reachability

Technical requirements

Flat network

ENI density

Continued growth

On premises

Full IP reachability on premises

Flaws with current IPv4 approach

Flat network

Not really “flat” (!)
Public vs. private IPv4

Continued growth

ENI density

On premises

Flaws with current IPv4 approach

Flat network

Not really “flat” (!)
Public vs. private IPv4

Continued growth

ENI density

AWS limits
IPv4 address reuse
EIP update lag

On premises

Flaws with current IPv4 approach

Flat network

Not really “flat” (!)
Public vs. private IPv4

Continued growth

AWS routing limits (VPC peering, etc.)
Private IPv4 address exhaustion

ENI density

AWS limits
IPv4 address reuse
EIP update lag

On premises

Flaws with current IPv4 approach

Flat network

Not really “flat” (!)
Public vs. private IPv4

Continued growth

AWS routing limits (VPC peering, etc.)
Private IPv4 address exhaustion

ENI density

AWS limits
IPv4 address reuse
EIP update lag

On premises

Private IPv4 address exhaustion
AWS routing limits (AWS Direct Connect, etc.)

A group of young people are shown in a crowd, likely at a festival or fair. In the foreground, two young women are smiling and looking towards the right. The woman on the left has curly brown hair and is wearing a colorful, abstract patterned shirt. She is holding a waffle cone with pink ice cream. The woman on the right has long, wavy reddish-brown hair and is wearing a white shirt with horizontal stripes in red, yellow, and blue. She is also holding a waffle cone with ice cream. In the background, other people are visible, some looking in the same direction. The overall atmosphere is happy and social.

What are we going to do?

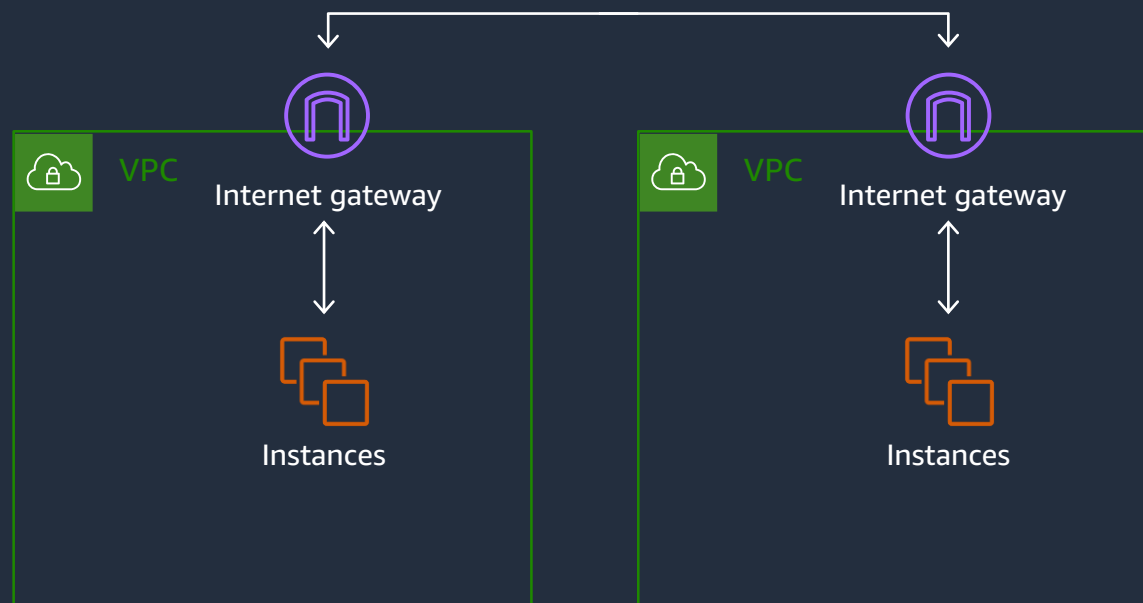
What are we going to do?

Partner with AWS

Ideas considered

“Amazon EC2 classic”

- ✓ We’ve done this before!
- ✗ Does not address ENI density



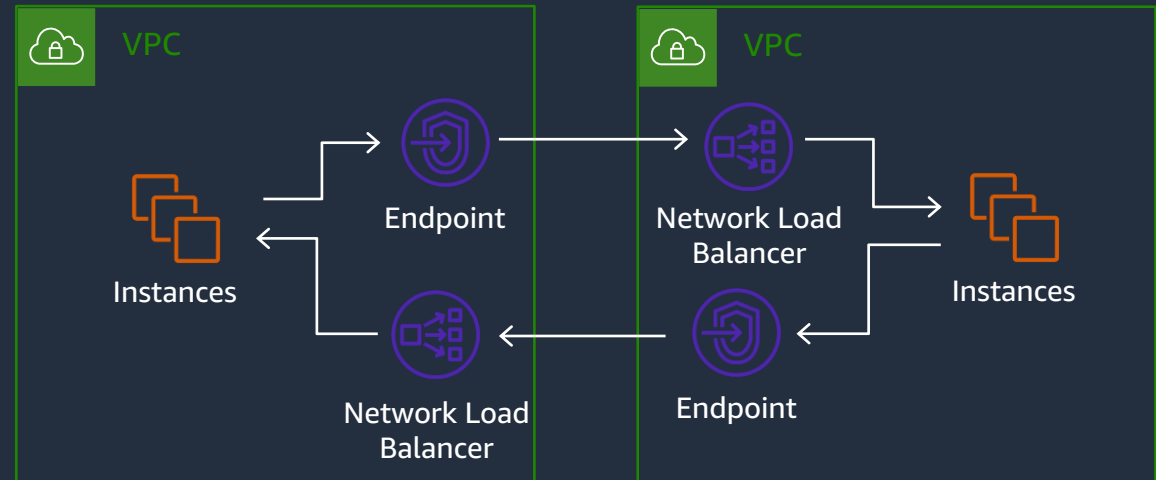
Ideas considered

“Amazon EC2 classic”

- ✓ We’ve done this before!
- ✗ Does not address ENI density

“Tiny bubbles”

- ✓ Well defined pattern
- ✗ Client-side load balancing
- ✗ Does not address ENI density




Ideas considered

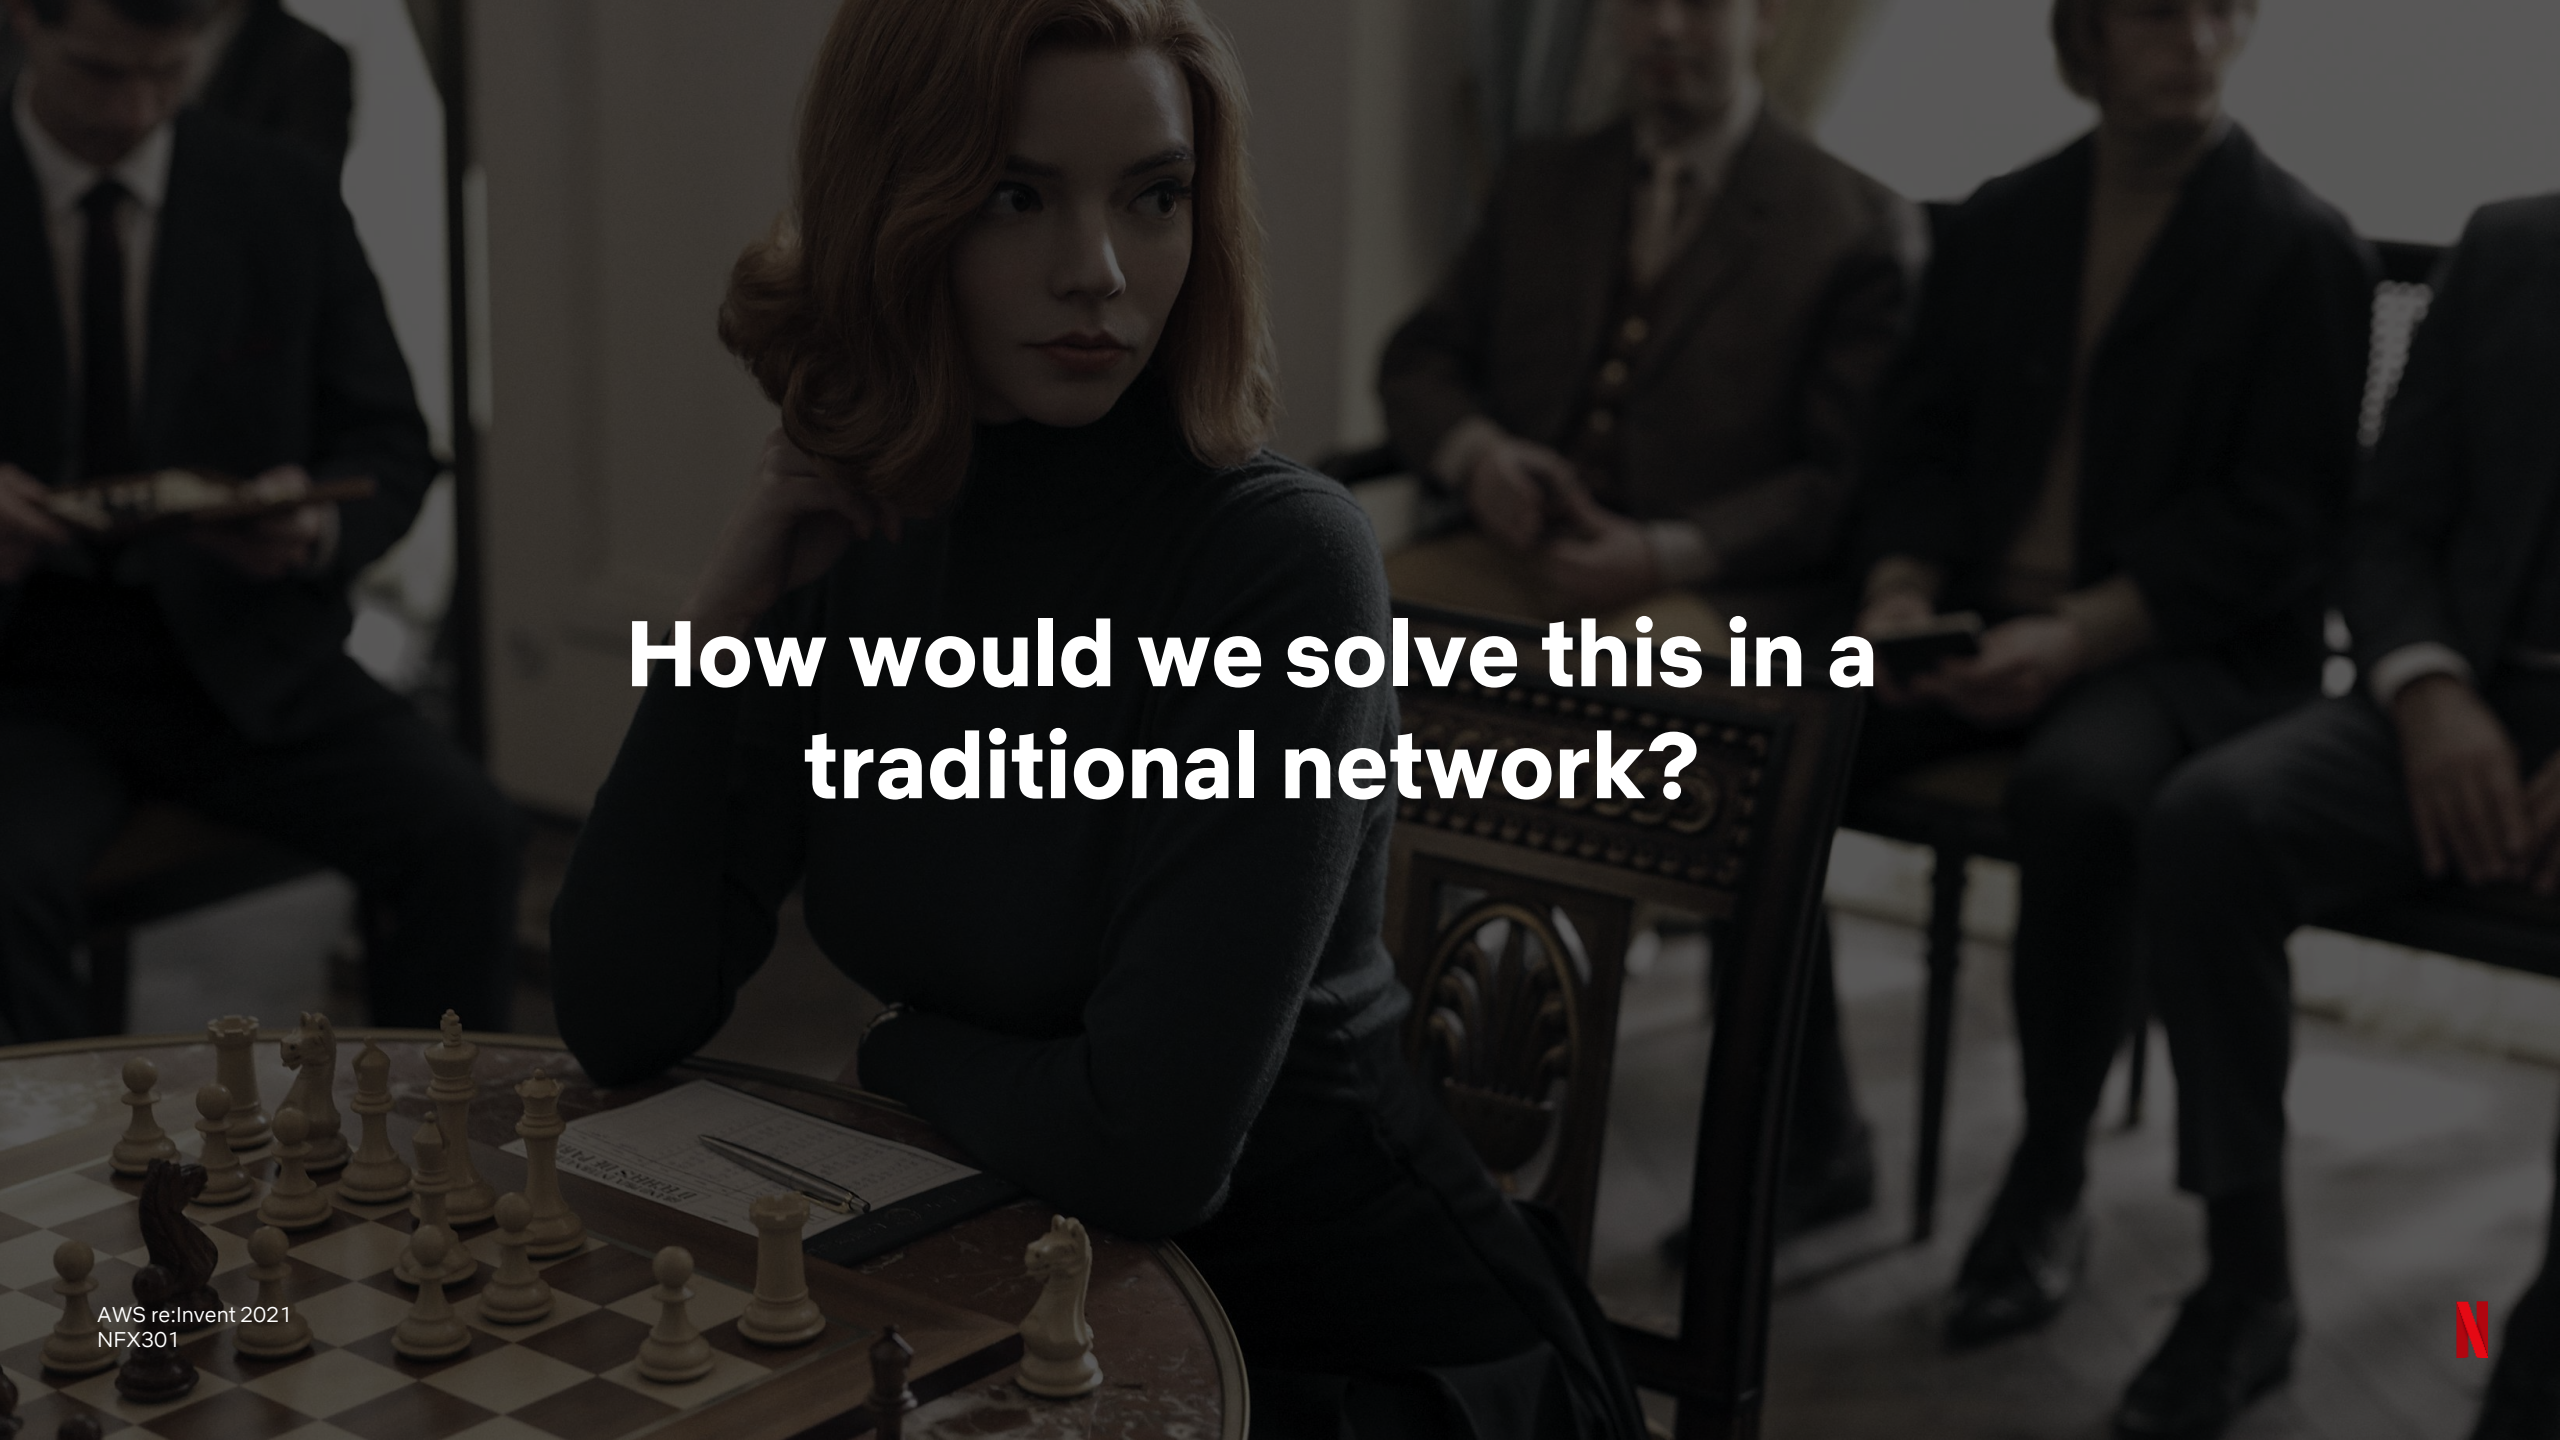
“Amazon EC2 classic”

- ✓ We’ve done this before!
- ✗ Does not address ENI density

“Tiny bubbles”

- ✓ Well defined pattern
- ✗ Client-side load balancing
- ✗ Does not address ENI density



A woman with blonde hair is sitting at a round chess table, looking thoughtfully at the camera. She is wearing a dark turtleneck. In front of her is a chessboard with several pieces. To her left, a man in a suit is looking down at a book. In the background, other people are seated at tables, some looking at their phones. The scene is dimly lit, creating a serious and intellectual atmosphere.

**How would we solve this in a
traditional network?**

How would we solve this in a
traditional network?

**Route
a network
block to the
host**

A dimly lit office scene with three men. On the left, a man with dark curly hair and a brown jacket looks intently at a screen. In the center, a man with blonde hair and glasses, wearing a red shirt and large headphones, also looks at the screen. On the right, a man with a beard and glasses, wearing a patterned blazer, leans in to look at the screen. The background is blurred, showing office cubicles and another person in the distance.

Can we do something similar in AWS?


Focus on ENI density

192.0.2.96 →
192.0.2.15 →
...
192.0.2.99 →
192.0.2.43 →



Elastic network interface

192.0.2.96/28 →



Elastic network interface

Prefix delegation

Focus on ENI density




Source: https://twitter.com/_joemag_/status/1418345704964063232


A person in a red shirt stands in the foreground, looking at a massive wall of digital screens. The screens display various data, including network maps, charts, and advertisements. The background is a dark, futuristic cityscape with glowing lights and signs. The overall scene suggests a high-tech environment, possibly a data center or a control room.

How much IP space do we need?

How much IP space do we need?



10.0.0.0/8




172.16.0.0/12




192.168.0.0/16


How much IP space do we need?



10.0.0.0/8



100.64.0.0/10




172.16.0.0/12




192.168.0.0/16


How much IP space do we need?




How much IP space do we need?



10.0.0.0/8





100.64.0.0/10





172.16.0.0/12





192.168.0.0/16

 Used on premises
 Used on AWS

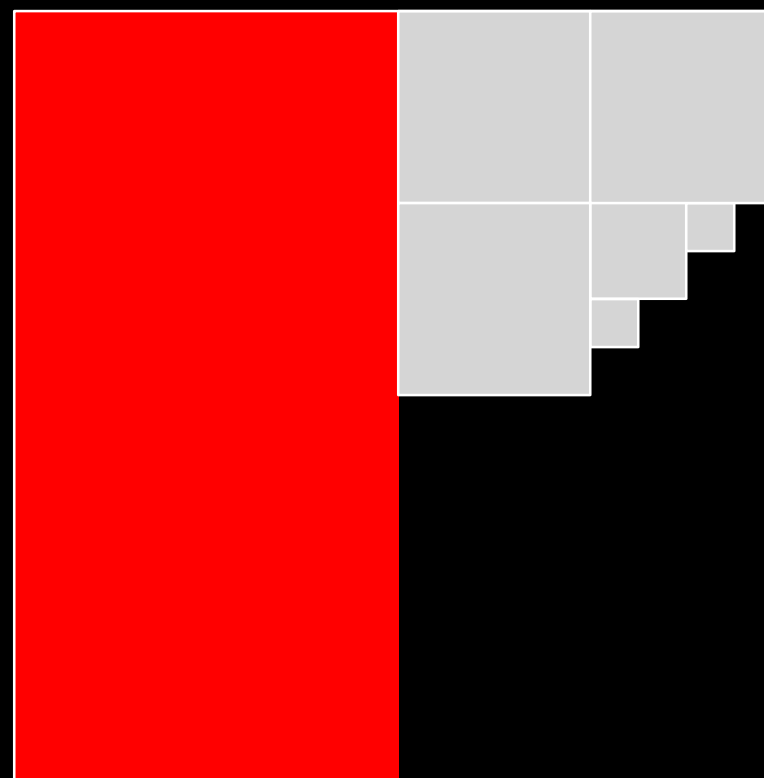
How much IP space do we need?




 Used on premises
 Used on AWS

How much IP space do we need?

/28 per prefix	16 IPs	/28
----------------	--------	-----



10.0.0.0/8




100.64.0.0/10


- Used on premises
- Used on AWS
- Prefix delegation

How much IP space do we need?

/28 per prefix	16 IPs	/28
4 prefixes per ENI	64 IPs	/26



10.0.0.0/8




100.64.0.0/10


- Used on premises
- Used on AWS
- Prefix delegation

How much IP space do we need?

/28 per prefix	16 IPs	/28
4 prefixes per ENI	64 IPs	/26
8K ENIs per zone	524k IPs	/13



10.0.0.0/8




100.64.0.0/10


- Used on premises
- Used on AWS
- Prefix delegation

How much IP space do we need?

/28 per prefix	16 IPs	/28
4 prefixes per ENI	64 IPs	/26
8K ENIs per zone	524k IPs	/13
3 zones per Region	1.5m IPs	/12 & /13



10.0.0.0/8




100.64.0.0/10


- Used on premises
- Used on AWS
- Prefix delegation

How much IP space do we need?

/28 per prefix	16 IPs	/28
4 prefixes per ENI	64 IPs	/26
8K ENIs per zone	524k IPs	/13
3 zones per Region	1.5m IPs	/12 & /13
3 Regions	4.5m IPs	/10 & /13



10.0.0.0/8




100.64.0.0/10

- Used on premises
- Used on AWS
- Prefix delegation

How much IP space do we need?

/28 per prefix	16 IPs	/28
4 prefixes per ENI	64 IPs	/26
8K ENIs per zone	524k IPs	/13
3 zones per Region	1.5m IPs	/12 & /13
3 Regions	4.5m IPs	/10 & /13
Double it!	9m IPs	/9 & /12



A person in a red shirt stands in the foreground, looking at a massive digital display that simulates a city at night. The display is covered in a grid of data, with various advertisements and signs visible, including 'LYBUKED', 'REST', 'XERO', and 'GIRLS GIRLS GIRLS'. The overall scene is dark and futuristic, with the city lights and data grid creating a complex, layered visual.

How much IP space do we need?

A person in a red shirt stands on a balcony, looking out over a sprawling, futuristic city at night. The city is densely packed with skyscrapers, each covered in glowing neon signs and advertisements in various colors like blue, green, red, and yellow. Some signs are in English, others in Japanese or Chinese. The scene is illuminated by the vibrant city lights, creating a cyberpunk atmosphere. The person's silhouette is visible against the bright background of the city.

How much IP space do we need?

A lot

ALTERED CARBON

IPv6


Focus on ENI density

192.0.2.96 →
192.0.2.15 →
...
192.0.2.99 →
192.0.2.43 →



Elastic network interface

192.0.2.96/28 →



Elastic network interface

Focus on ENI density

192.0.2.96 →
192.0.2.15 →
...
192.0.2.99 →
192.0.2.43 →



Elastic network interface

192.0.2.96/28 →



Elastic network interface

2001:db8::96 →
2001:db8::15 →
...
2001:db8::99 →
2001:db8::43 →



Elastic network interface

2001:db8:96::/80 →



Elastic network interface

Does IPv6 solve business network requirements?

Flat network



No NAT

Continued growth

ENI density

On premises

Does IPv6 solve business network requirements?

Flat network



No NAT

Continued growth

ENI density




Prefix delegation


On premises

Does IPv6 solve business network requirements?

Flat network

 No NAT

Continued growth

 No IPv4 address exhaustion
AWS limits (routing, etc.)

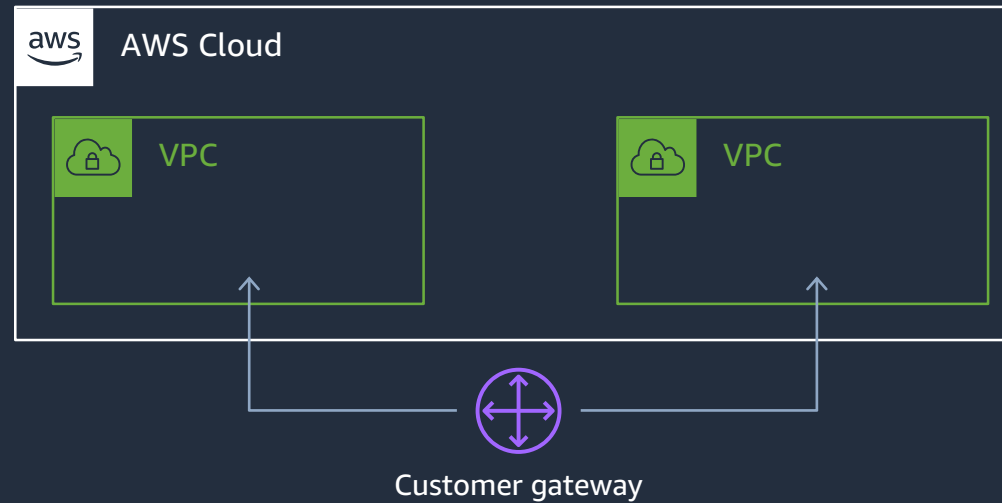
ENI density

 Prefix delegation

On premises

Inter-VPC private IPv4 reachability

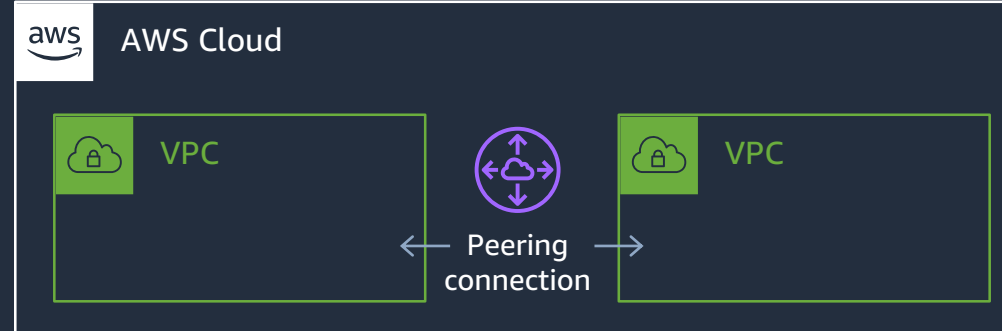
1 – Customer gateway



Inter-VPC private IPv4 reachability

1 – Customer gateway

2 – VPC peering




Inter-VPC private IPv4 reachability

1 – Customer gateway

2 – VPC peering

3 – AWS Transit Gateway




Inter-VPC private IPv4 reachability

1 – Customer gateway

2 – VPC peering

3 – AWS Transit Gateway

4 – Internet gateway



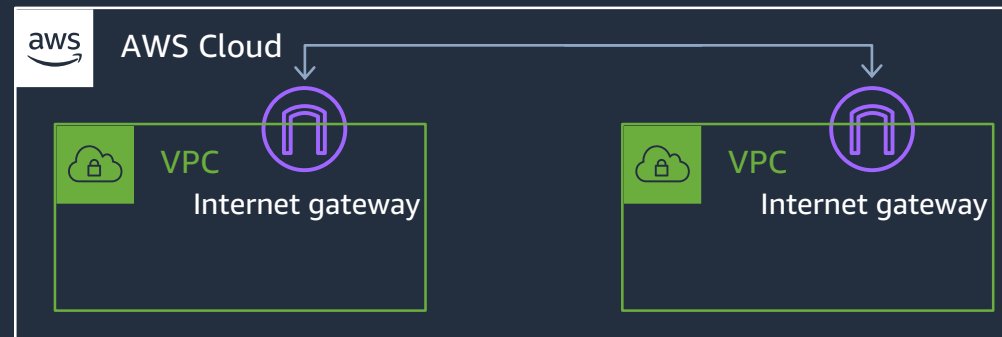
Inter-VPC IPv6 reachability

1 – Customer gateway

2 – VPC peering

3 – AWS Transit Gateway

4 – Internet gateway



Does IPv6 solve business network requirements?

Flat network



No NAT

Continued growth



No IPv4 address exhaustion

AWS limits (routing, etc.)

ENI density




Prefix delegation


On premises

Does IPv6 solve business network requirements?

Flat network

 No NAT

Continued growth

 No IPv4 address exhaustion
Network setup not explicitly required

ENI density

 Prefix delegation

On premises

Does IPv6 solve business network requirements?

Flat network

- ✓ No NAT

Continued growth

- ✓ No IPv4 address exhaustion
Network setup not explicitly required

ENI density

- ✓ Prefix delegation

On premises

- ✓ Network setup not explicitly required



Should we really do this?

Should we really do this?

Yes!



IPv6 co-innovation with AWS

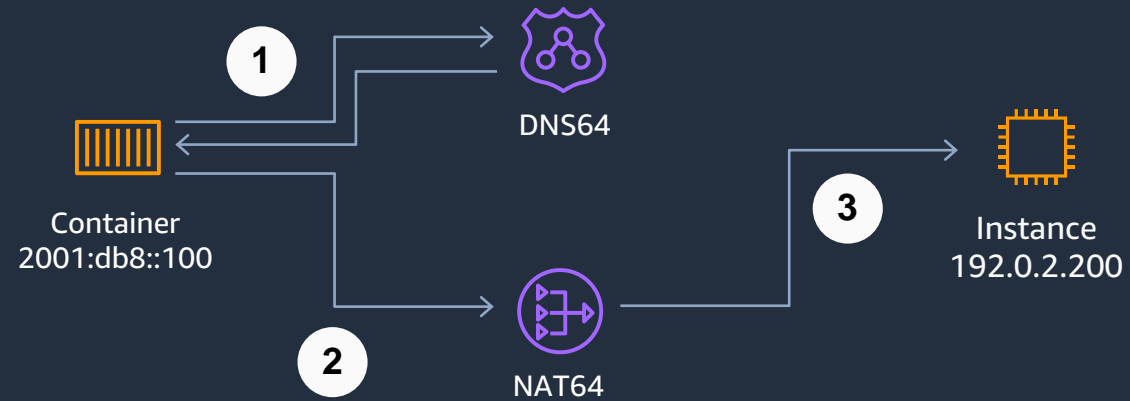
Prefix delegation



IPv6 open issues with AWS

Prefix delegation

IPv6 transition mechanism

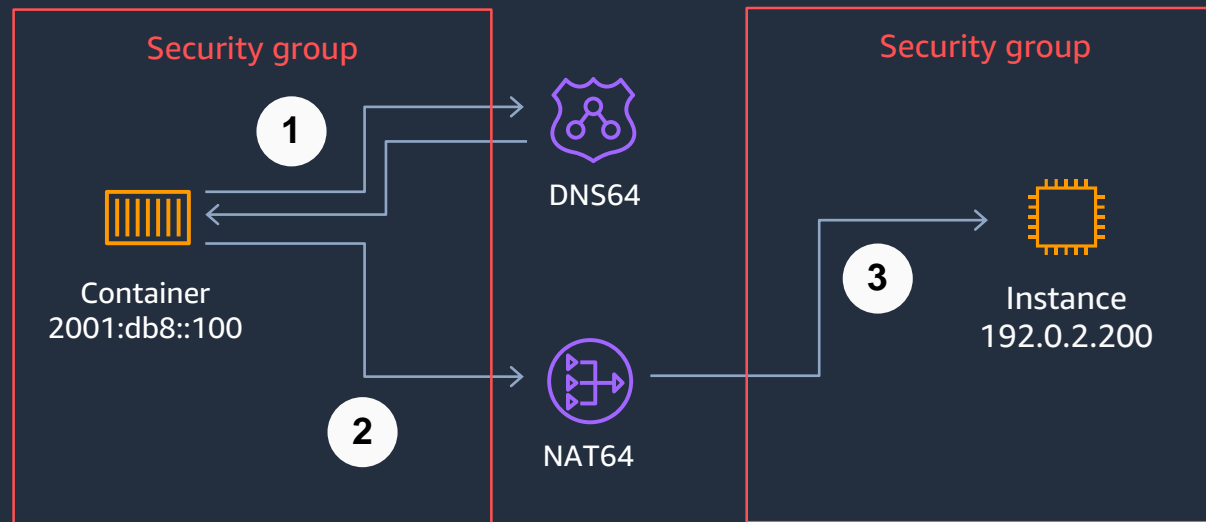


Service discovery **must** be done via DNS!

IPv6 open issues with AWS

Prefix delegation

IPv6 transition mechanism

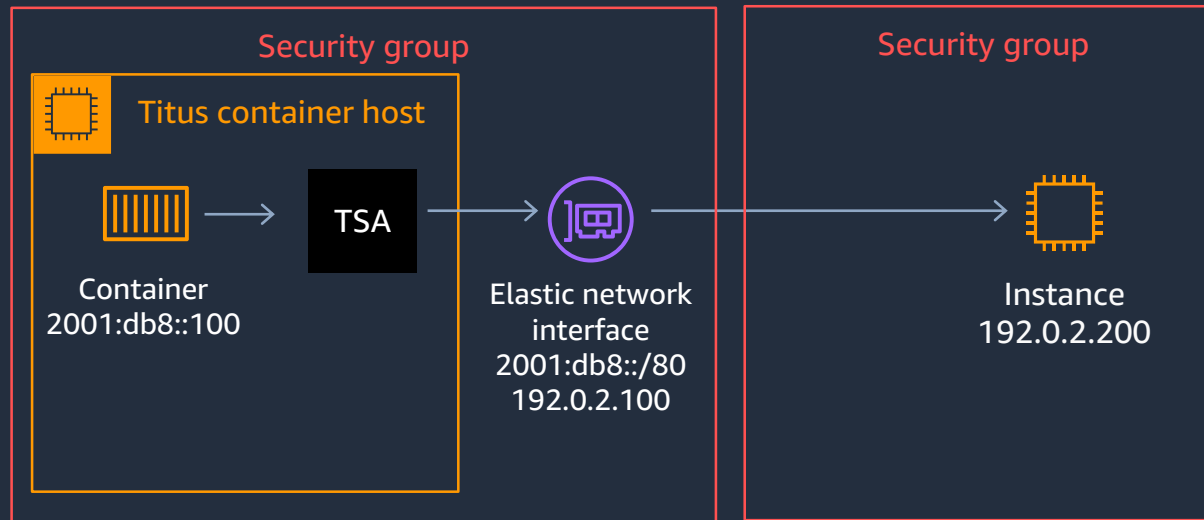


Security group references **do not work!**

Netflix innovation within AWS

Prefix delegation

IPv6 transition mechanism



Security group references **do work!**

More info at the Netflix Linux Plumbers Conference presentation
<https://bit.ly/nflx-tsa>

Netflix AWS VPC
IPv6 adoption progress


IPv6

Prefix delegation

In production



Netflix AWS IPv6 adoption progress 2021



Started 2021 with IPv6 being <1% of all interservice flows in VPC

Lessons learned

- Old code is not fun, especially Java

Lessons learned

Old code is not fun, especially Java

- **Assigning IPv6 to a node does not mean IPv6 is used**

Lessons learned

Old code is not fun, especially Java

Assigning IPv6 to a node does not mean IPv6 is used

- **Happy Eyeballs masks IPv6 problems**

Lessons learned

Old code is not fun, especially Java

Assigning IPv6 to a node does not mean IPv6 is used

Happy Eyeballs masks IPv6 problems

- **Little IPv6 support for AWS Managed Services**

Netflix best practices

- Communication

Netflix best practices

Communication

- **Use BYOIP for IPv6 to have summarizable address space, similar to private IPv4**

Netflix best practices

Communication

Use BYOIP for IPv6 to have summarizable address space, similar to private IPv4

- **Overlay IPv6 with IPv4**

Netflix best practices

Communication

Use BYOIP for IPv6 to have summarizable address space, similar to private IPv4

Overlay IPv6 with IPv4

- **Match IPv4 IP range rules with corresponding IPv6 IP range rules**

How do I get started?

How do I get started?

- **Dual stack workstations and dev machines**

How do I get started?

Dual stack workstations and dev machines

- **Enable IPv6 on the edge**
Update security groups
Update DNS records

How do I get started?

Dual stack workstations and dev machines

Enable IPv6 on the edge

Update security groups

Update DNS records

- **Focus on workloads in VPC**
 - Overlay IPv6 with IPv4**
 - Edge inwards**

How do I show IPv6 is worthwhile to the business?

- **Economics**

IPv6 is faster

Expensive boundary: public IPv4 vs. private IPv4

Middle boxes and operational risk



We have started the journey to IPv6 in VPC and so should you

STRANGER THINGS

Thank you!

Donavan Fritz

dfritz@netflix.com

[@DonavanFritz](https://twitter.com/DonavanFritz)

[linkedin.com/in/DonavanFritz](https://www.linkedin.com/in/DonavanFritz)

