Deep Dive on AWS PrivateLink

Colm MacCárthaigh, AWS VPC

January 31st, 2018

Deep Dive on AWS PrivateLink

- Very quick VPC fundamentals
- PrivateLink for AWS Services
- How PrivateLink works: HyperPlane
- PrivateLink for customer and partner services

aws

VPC Fundamentals

What is VPC?

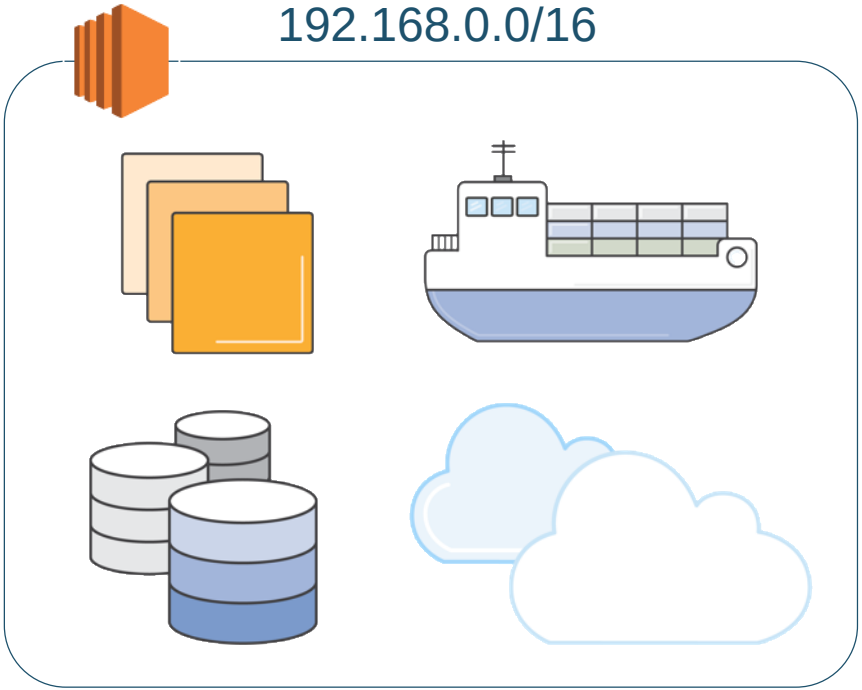

What is VPC?

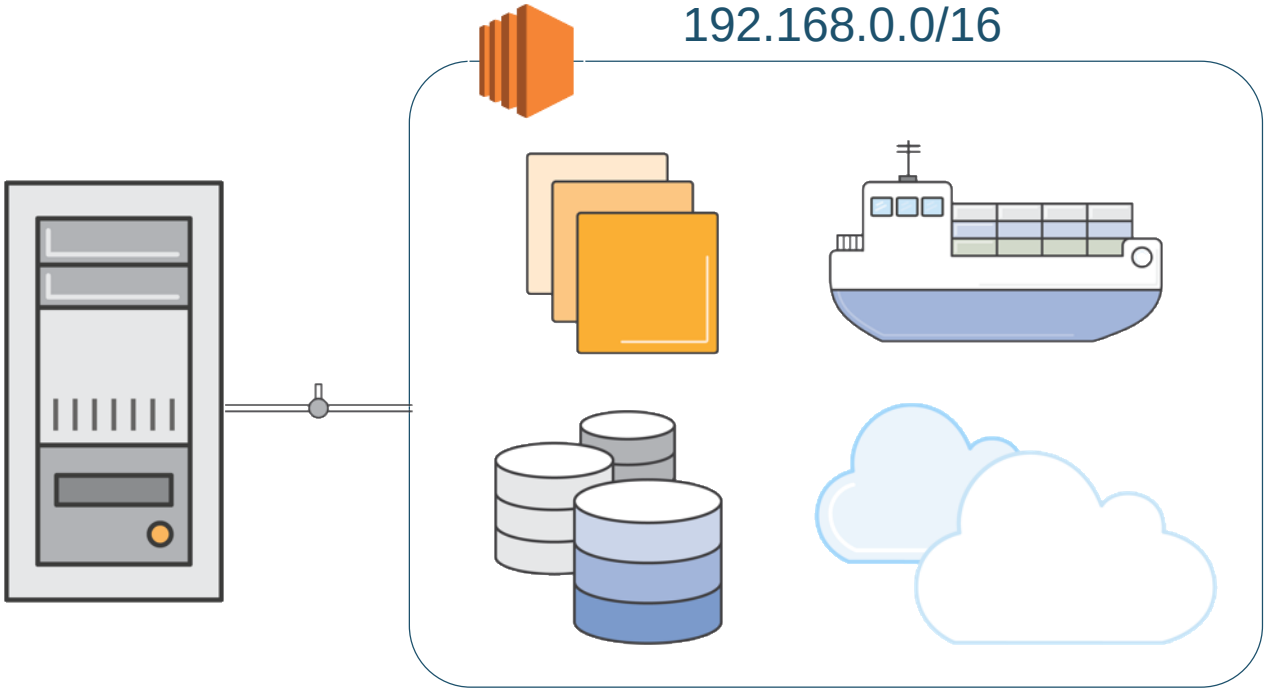

What is VPC?

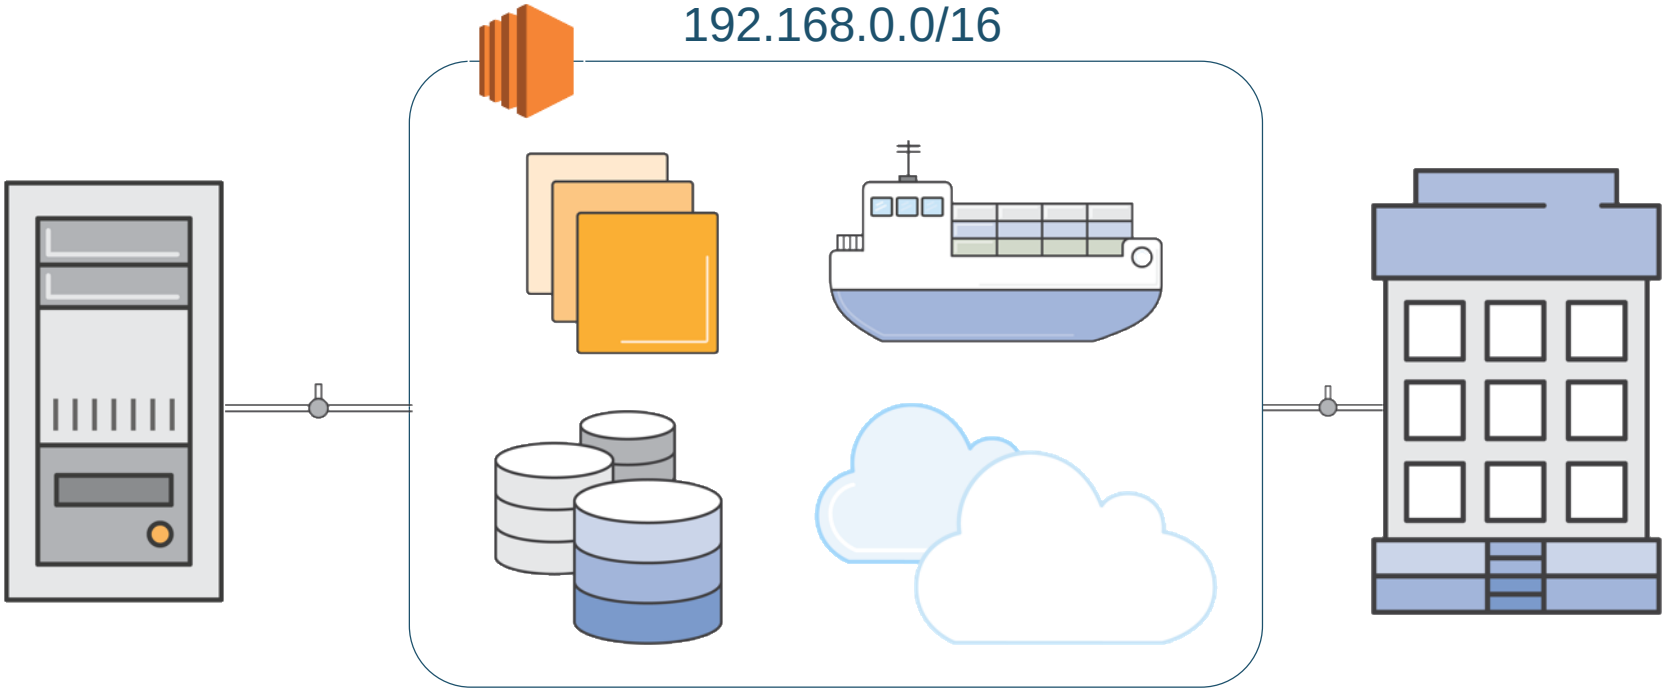

VPC Endpoints

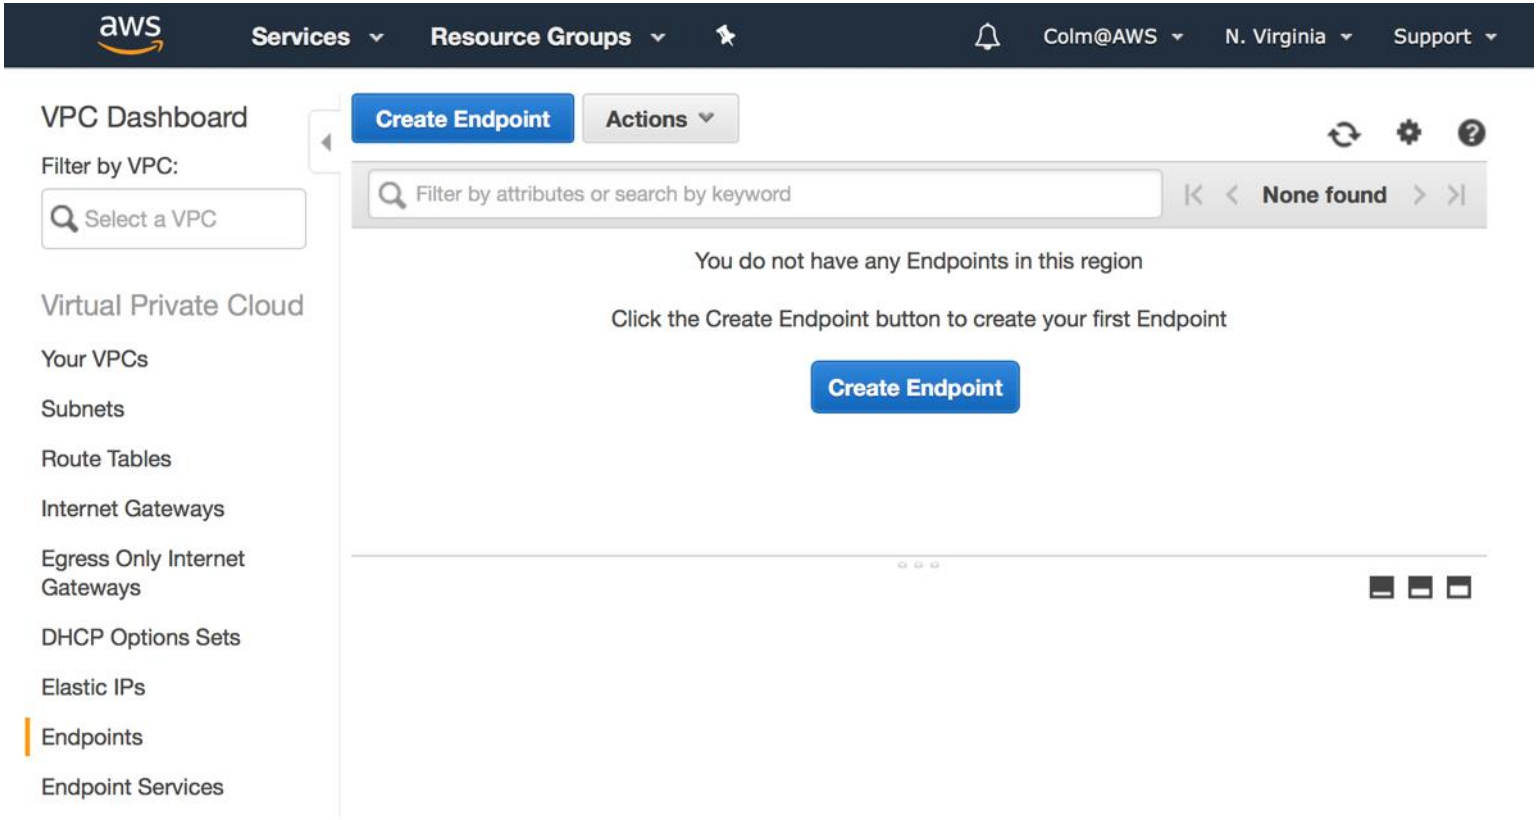

VPC Endpoints

- Launched in 2015
- Support for Amazon S3 and DynamoDB
- Services keep their public IPs, but are connected directly to a VPC.
- Now called Gateway Endpoints

PrivateLink for AWS Services

PrivateLink

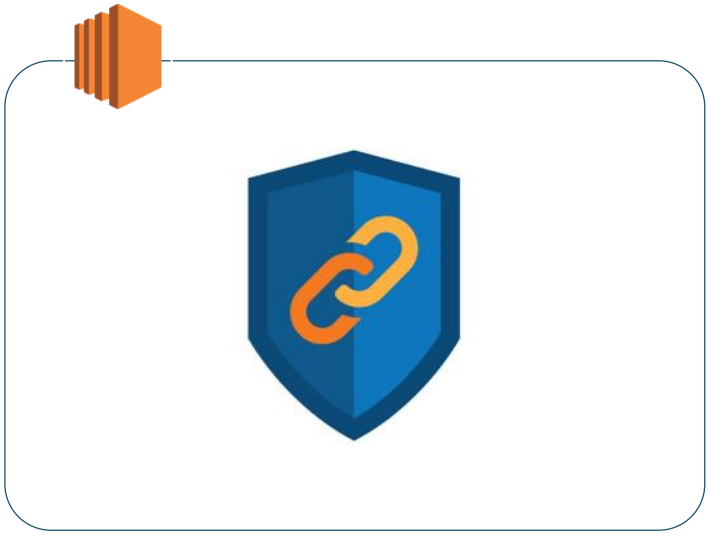

PrivateLink

PrivateLink

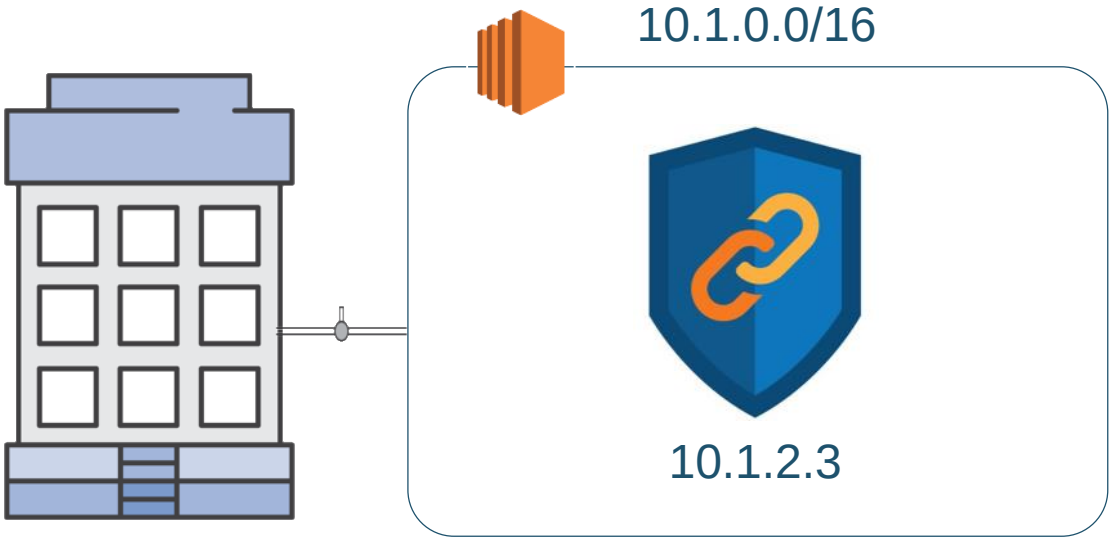

PrivateLink for AWS Services

- Launched on November 8th, 2017
- APIs available as PrivateLink Endpoints: Kinesis, EC2, ELB, EC2 Systems Manager, Service Catalog, KMS
- IP connectivity is fully private
- VPC DNS over-rides the service names

PrivateLink for AWS Services

- PrivateLink Endpoints also have regional and zonal names
- PrivateLink Endpoints work with VPC Security Groups
- PrivateLink Endpoints work with IAM Policies
- PrivateLink Endpoints work via Direct Connect

aws

PrivateLink for AWS Services

Details | Subnets | Security Groups

Endpoint ID	vpce-04a3c569d0096034d
Status	available
Service name	com.amazonaws.us-east-1.kinesis-streams
DNS Names	vpce-04a3c569d0096034d-hxkjzkdh.kinesis.us-east-1.vpce.amazonaws.com (Z7HUB22UULQXV) vpce-04a3c569d0096034d-hxkjzkdh-us-east-1c.kinesis.us-east-1.vpce.amazonaws.com (Z7HUB22UULQXV) vpce-04a3c569d0096034d-hxkjzkdh-us-east-1b.kinesis.us-east-1.vpce.amazonaws.com (Z7HUB22UULQXV) vpce-04a3c569d0096034d-hxkjzkdh-us-east-1a.kinesis.us-east-1.vpce.amazonaws.com (Z7HUB22UULQXV) kinesis.us-east-1.amazonaws.com (ZF38KS9GFNLME)

VPC ID	vpc-29022550 | VPCE Test VPC
Creation Time	August 22, 2017 at 11:32:21 AM UTC-7
Endpoint type	Interface

aws

How PrivateLink works: HyperPlane

HyperPlane

- In production since 2015, based on the S3 Load Balancer
- Supports Elastic Filesystem, VPC NAT Gateway, Network Load Balancer, and PrivateLink
- Massively scalable and fault-tolerant distributed system for managing connections

HyperPlane

HyperPlane nodes make transactional decisions and share state in tens of microseconds.

aws

HyperPlane

For NLB: HyperPlane selects the target instance, IP, or container that should handle a connection

HyperPlane and Shuffle Sharding

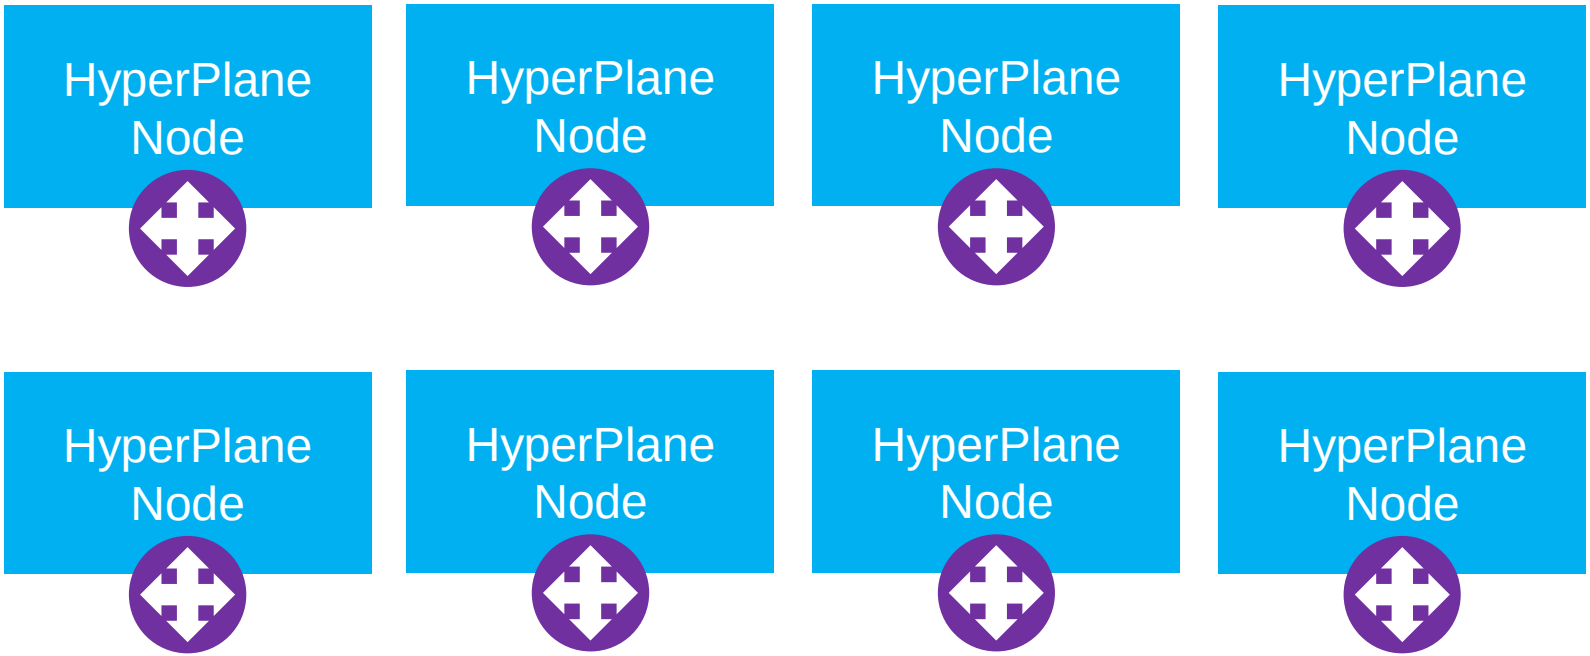

aws

HyperPlane and Shuffle Sharding

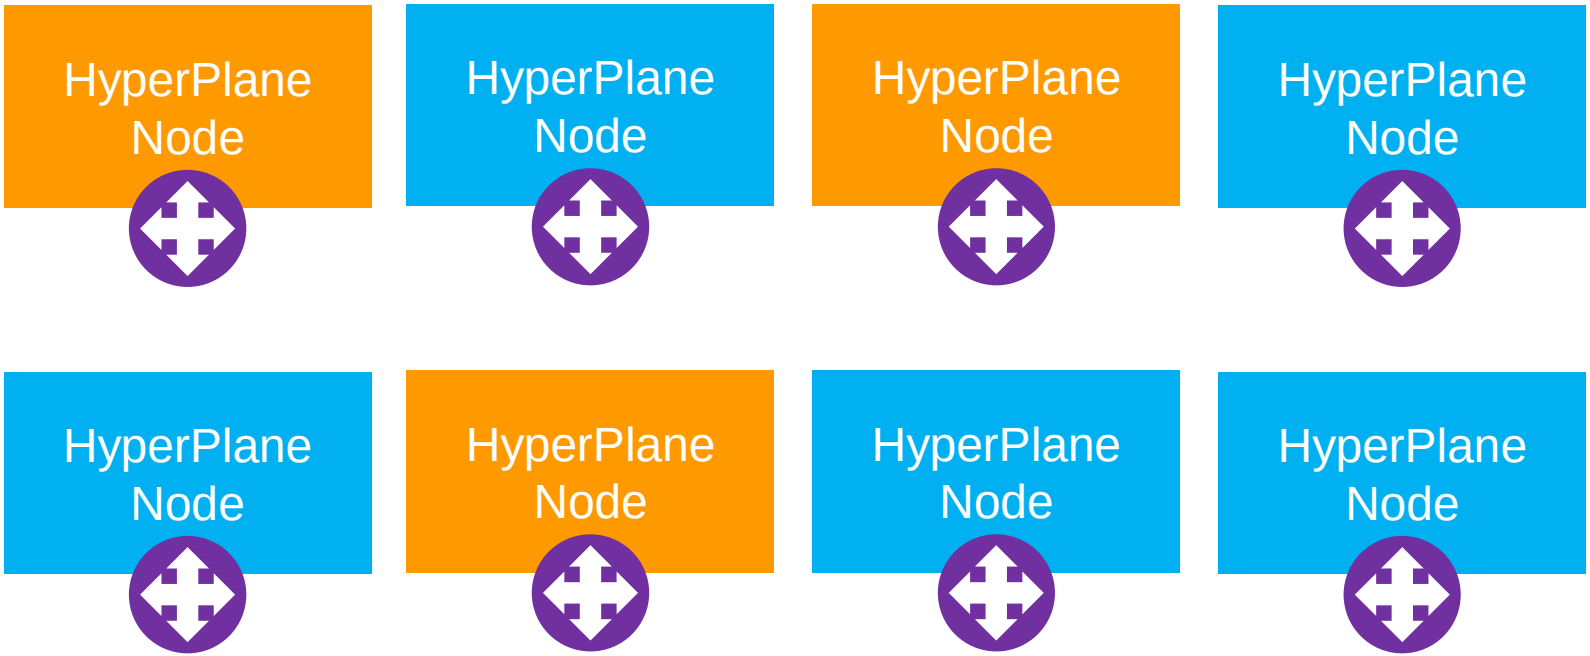

aws

HyperPlane and Shuffle Sharding

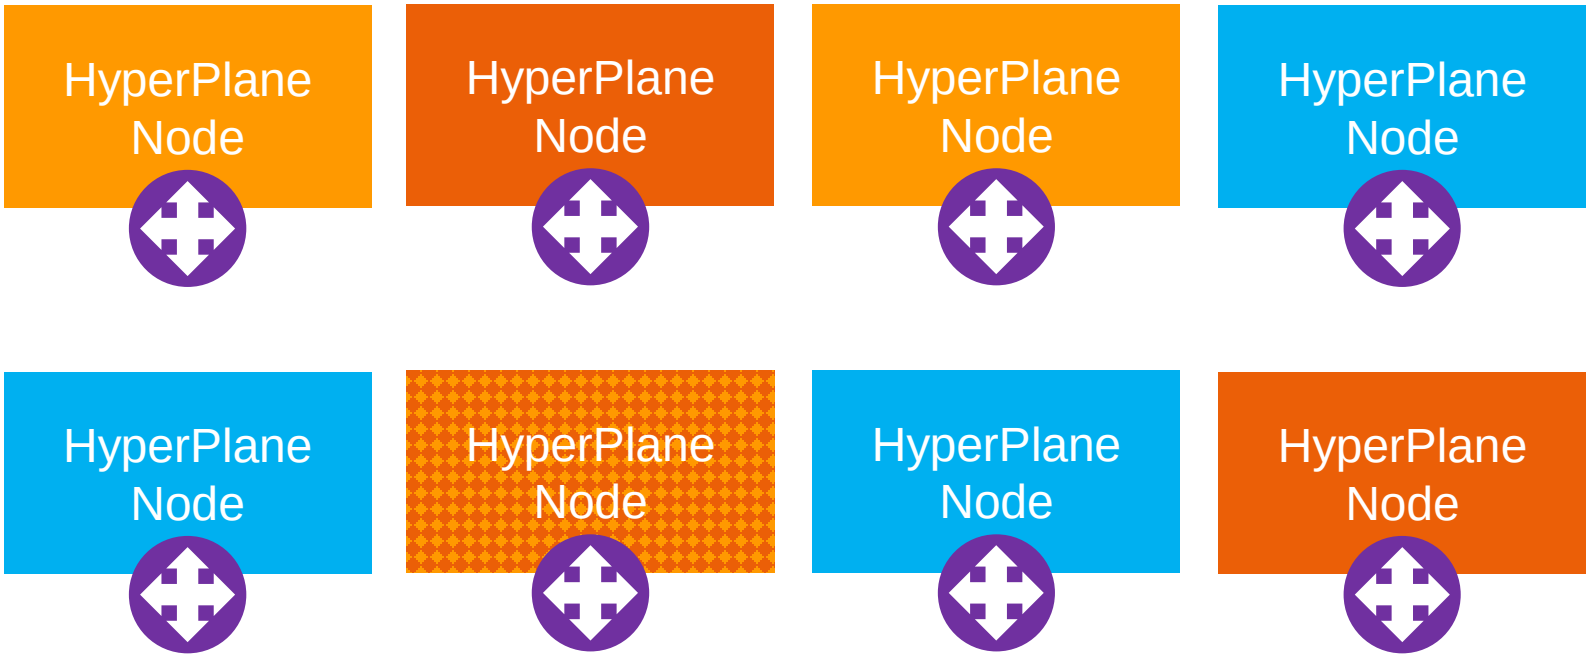

aws

HyperPlane and Shuffle Sharding

Potential Overlap	Percentage chance
0	18%
1	54%
2	26%
3	2%

aws

HyperPlane and Shuffle Sharding

Potential Overlap	Percentage chance
0	77%
1	21%
2	1.8%
3	0.06%
4	0.0006
5	0.00000013

aws

HyperPlane and Shuffle Sharding

Potential Overlap	Percentage chance
0	77%
1	21%
2	1.8%
3	**0%**
4	**0%**
5	**0%**

aws

HyperPlane

- Every HyperPlane resource has 5Gbit/sec of capacity by default, and scales in increments of 5Gbit/sec … to Terabits

- Sub-millisecond latency, hundreds of millions of connections, millions of connections per second

PrivateLink for Customer and Partner Services

PrivateLink for Customers and Partners

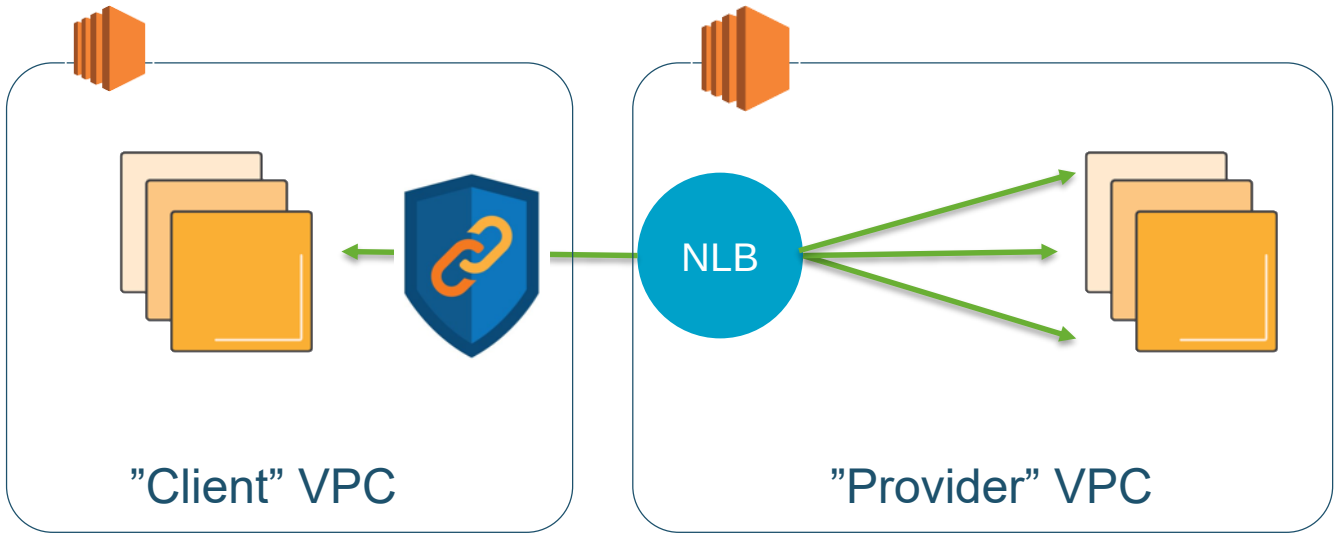

PrivateLink for Customers and Partners

- Service providers: offer services securely and privately directly into Customer VPCs, optionally integrate with AWS Marketplace.
- Micro-service architectures: compartmentalize micro-services into their own VPCs

PrivateLink for Customers and Partners

- On the “Client” or “Consumer” VPC side, not much changes.

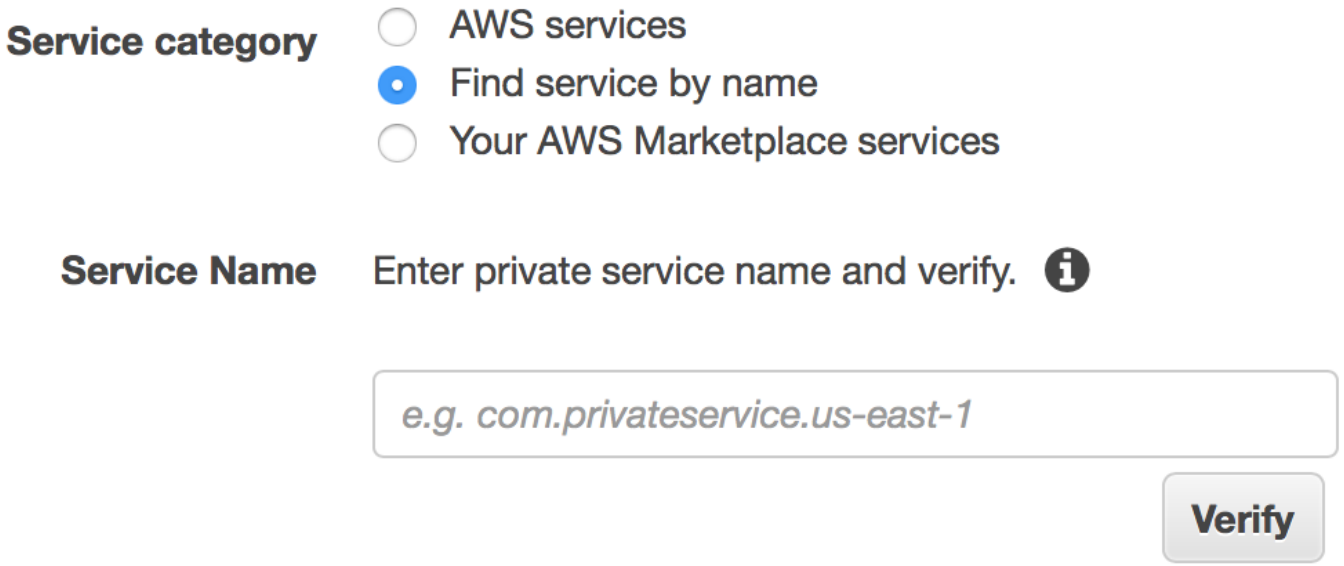

PrivateLink for Customers and Partners

- On the “Service” or “Provider” side, think of it as extending a Network Load Balancer
- Create and Manage an NLB as normal
- Associate it with an Endpoint Service, choose your availability-zones
- Decide if Endpoints require acceptance or not

aws

PrivateLink for Customers and Partners

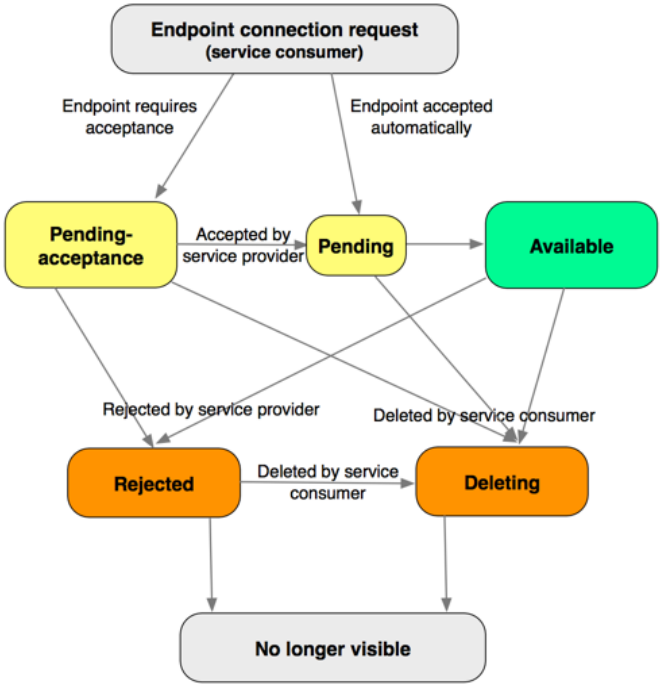

Notification of lifecycle events available via AWS SNS

You can automate sign-ups, leave, events

Invoke Lambda from SNS

PrivateLink for Customers and Partners

- Use the “wildcard pattern” for easy integration with SSL/TLS
- Example: *.supercoolservice.com
- When customers create Endpoints, intercept via SNS notifications and create CNAMEs. E.g. colm.supercoolservice.com
- Give the customer the CNAME name to use.

aws

PrivateLink for Customers and Partners

- If a provider has targets and NLB in each zone, then those zones will be available to customers
- Best latency by being in as many zones as possible, use at least two for availability
- Reminder: zone names vary between accounts!
- Cross-region setups come with availability and data sovereignty risks

aws

PrivateLink for Customers and Partners

- Single-Tenant mode: Create a PrivateLink NLB for every client/customer
- Multi-Tenant mode: allow many customers to use the same PrivateLink NLB
- How do we tell Endpoint traffic from different VPCs apart?

aws

PrivateLink for Customers and Partners

- Method 1: use traditional accounts/passwords/security-tokens at application level.
- Method 2: use separate NLBs and different listener ports on the targets.
- Method 3: Enable the ProxyProtocolV2 pre-amble.

aws

PrivateLink for Customers and Partners

The Proxy Protocol header also includes the ID of the endpoint. This information is encoded using a custom Type-Length-Value (TLV) vector as follows.

Field	Length (in octets)	Description
Type	1	PP2_TYPE_AWS (0xEA)
Length	2	The length of Value
Value	1	PP2_SUBTYPE_AWS_VPCE_ID (0x01)
	variable (Value length minus 1)	The ID of the endpoint

aws

PrivateLink for Customers and Partners

- Provider side pays for the NLB costs as normal
- Client side pays PrivateLink Endpoint costs as normal
- AWS Marketplace Integration available for easy discovery and billing

Takeaways

PrivateLink Takeaways

- PrivateLink Endpoints are “in” your VPC, accessible via Direct Connect too
- Powered by HyperPlane and provisioned for scale
- You can create and manage your own PrivateLink services
- More services coming!

aws

Thank you!