

Introducing ECS Service Discovery

Communication service for building
containerized microservices on AWS

Nathan Peck

April 23, 2018

Application communication is evolving




Across the room




Across cities / continents 1:1




Dynamic name,
number, and location



Functional calls



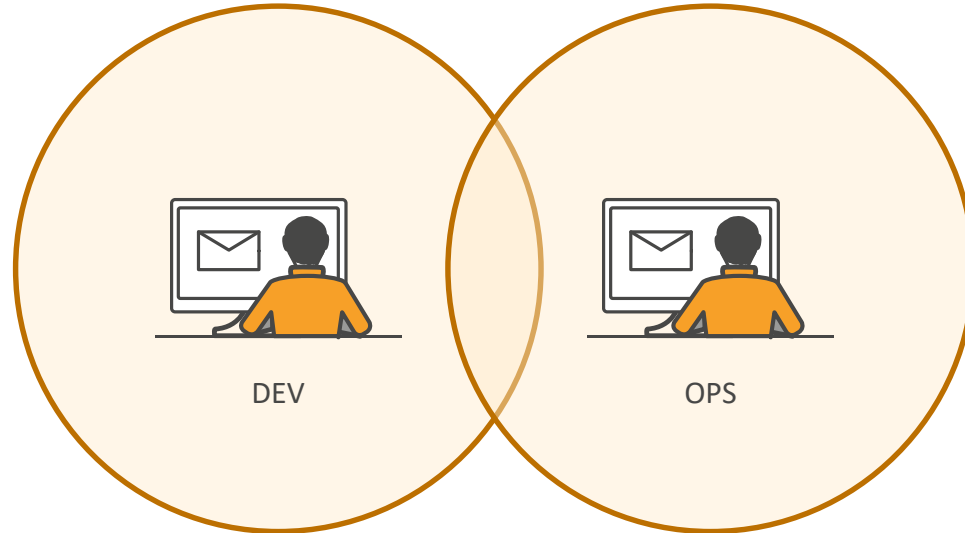
Known endpoints, APIs



Find endpoints,
then connect

Developers need to connect microservices

Build apps
invoking other services
by name



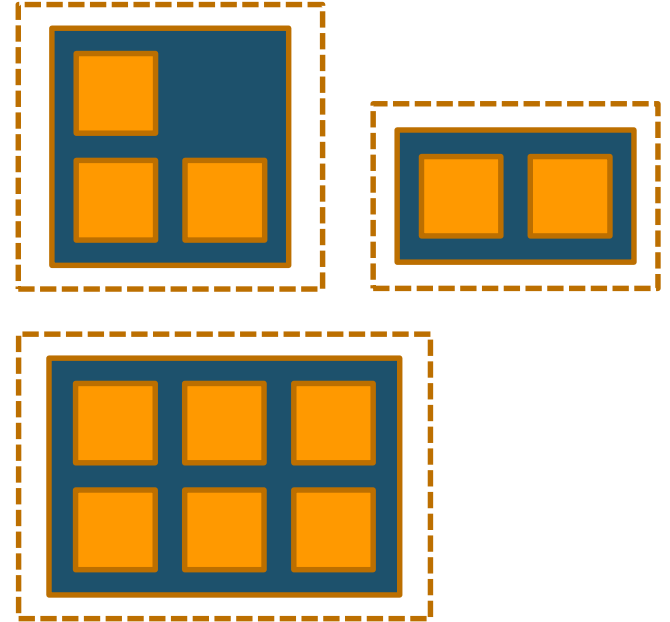
Ensure that service
name resolves to
correct IP/port

What is Service Discovery?

“Where is Service X?”

Friendly name -> IP + port

E.g., app: {10.0.4.5:8080, 10.0.4.6:8080 }



Why is it non trivial?

Dynamic by design:

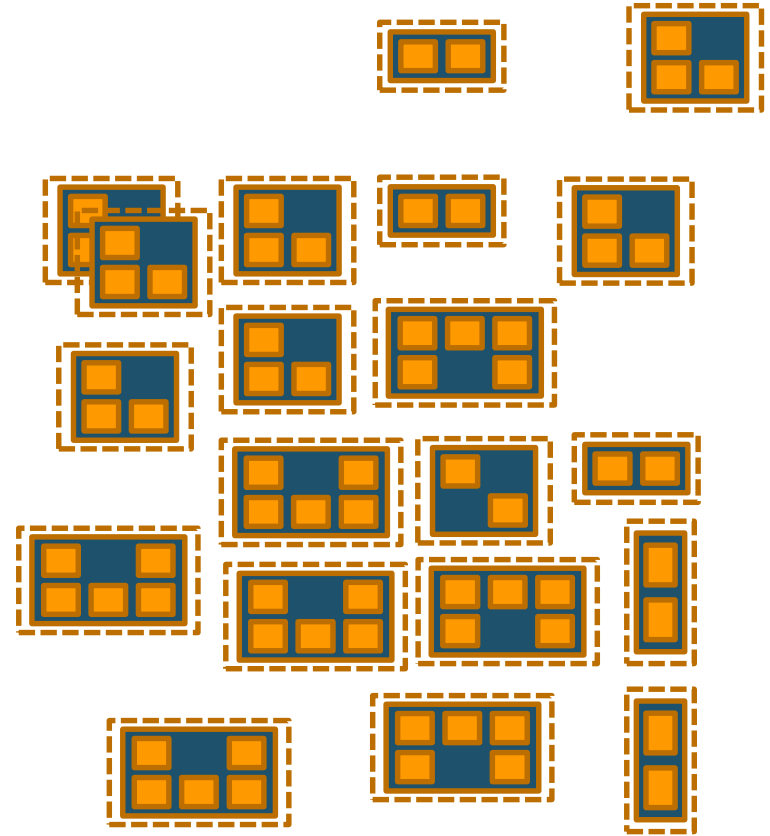
Number of containers & instances

Auto assigned IP addresses & ports

Placement, scheduling, scaling

Deployments and upgrades

Health and connectivity



Decision criteria



Service Registry

Where is info about services stored?


Registration mechanism

How is a service added to the registry?


Discovery mechanism

How does an application send traffic to a service from the registry?


Current patterns require install, setup and management



Load Balancers



Key-value store



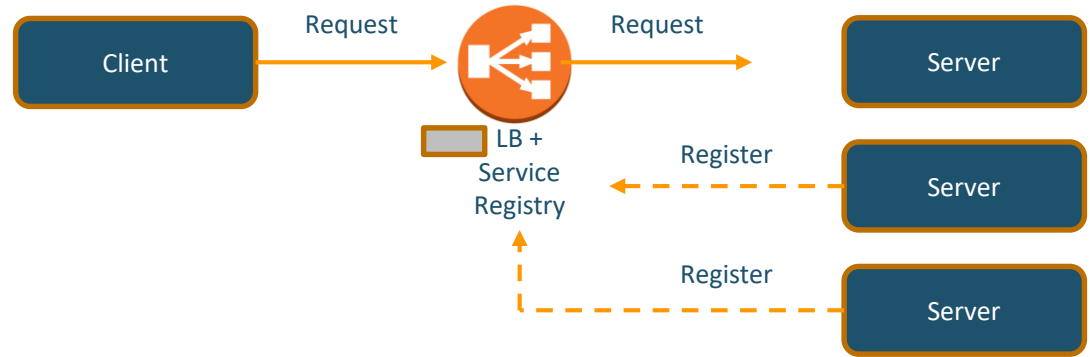
Service Mesh

Load Balancers – Server Side Discovery

E.g.,
Load Balancer

Benefits
Client is simpler

Drawbacks
Install, manage
Availability, capacity
More hops

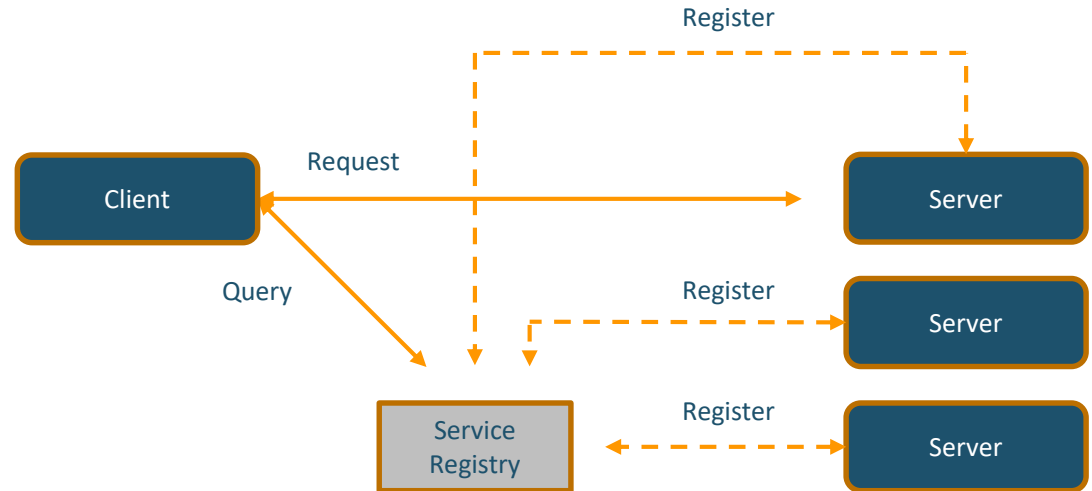


DNS based – Client Side Discovery

E.g.,
Route 53 Based

Benefits
Fewer hops

Drawbacks
Client must be registry aware
Client implements discovery logic



Requires registrations by agents

E.g.,

Consul, Etcd and Zookeeper

Benefits

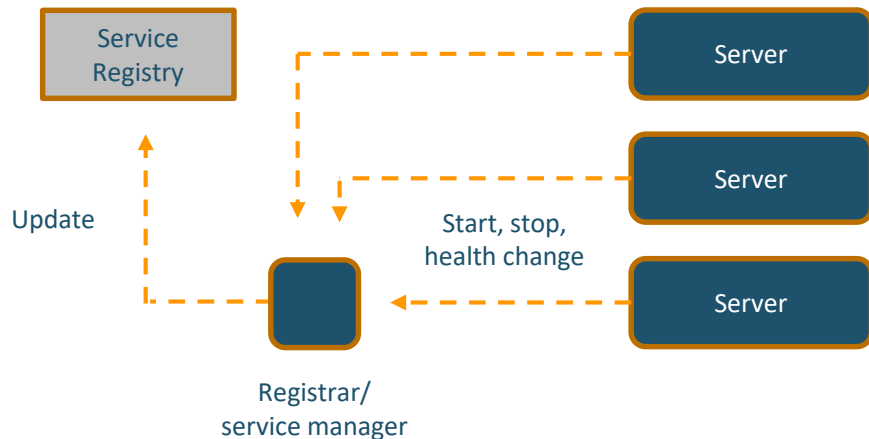
Registrar registers, unregisters and performs health checks

Service less complex

Drawbacks

State may not reflect whether service can handle requests, only active or unavailable

Install, configure and manage other component, unless it is a part of infrastructure



Requires registrations by agents

E.g.,
Consul, Etcd and Zookeeper

Benefits

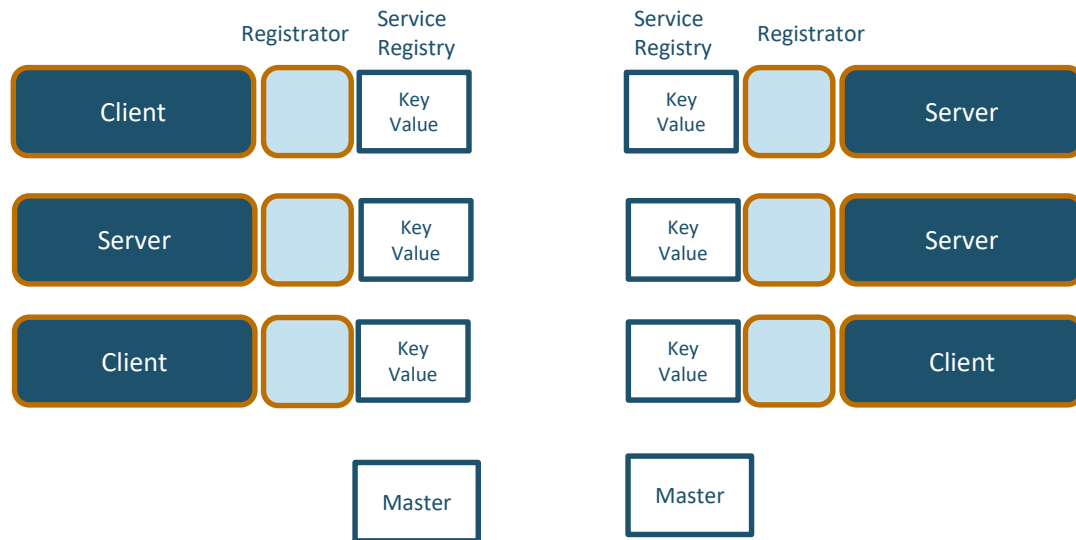
Registrar registers, unregisters and performs health checks

Service less complex

Drawbacks

State may not reflect whether service can handle requests, only active or unavailable

Install, configure and manage other component, unless it is a part of infrastructure




This should be easier!

Predictable Names
for services



Auto updated with
latest, healthy IP,
port



Managed: No
overhead of
installation or
monitoring




High availability,
high scale



Extensible: Flexible
boundaries for
auto discovery



Introducing managed service discovery for ECS




DEV

OPS

You build apps where
services are invoked by
name & name resolves to
IP/port dynamically

You turn on service
discovery during
deployment — service
creation

ECS service discovery is powered by Route 53



ECS updates service registry based on naming convention, task registrations, de-registrations and health

Route 53 provides Service Registry

Demo

<https://servicediscovery.ranman.com/>



Enables these use cases

1

Blue green deployments

- *myapp.staging.local*
- *myapp.prod.local*
- Private IP
- abstract cluster details

2

Internal micro services

- *web.myapp.local*
- *Expose Private IP*

3

External micro services

- *web.myapp.mycompany.com*
- *Expose public IP or ELB EIP*
- *network + container health check*

Enables these use cases

4

Across ECS & Kubernetes

- `Service1.myapp.ecs`
- `Service2.myapp.eks`

5

Across ECS &
AWS & On-Prem

- `Service1.myapp.ecs`
- `Service2.myapp.ec2`
- `Service3.myapp.onprem`

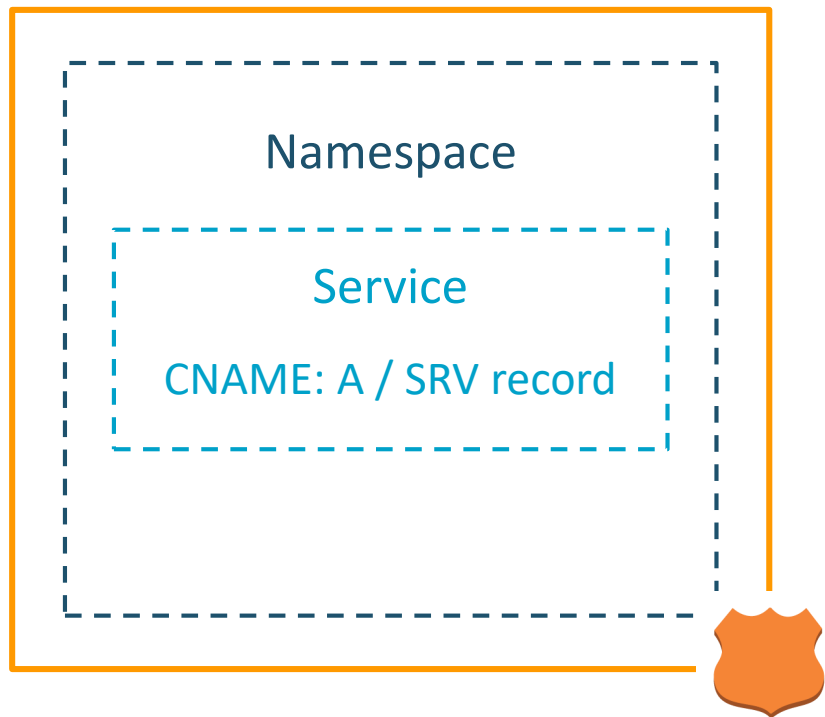
6

Expose to service mesh

- `Service1.myapp.local`
- `Service2.myapp.local`

What's New?

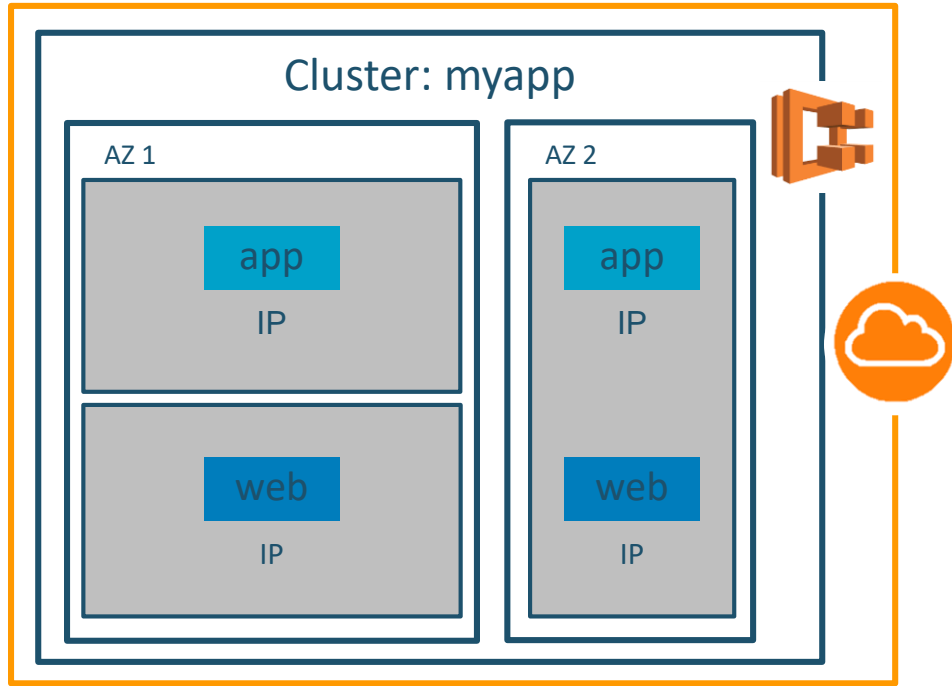
Route 53 provides Service Registry



Route 53 provides APIs to create

- Namespace
- CNAME per service autaname
- A records per task IP
- SRV records per task IP + port

ECS schedules & places service endpoints

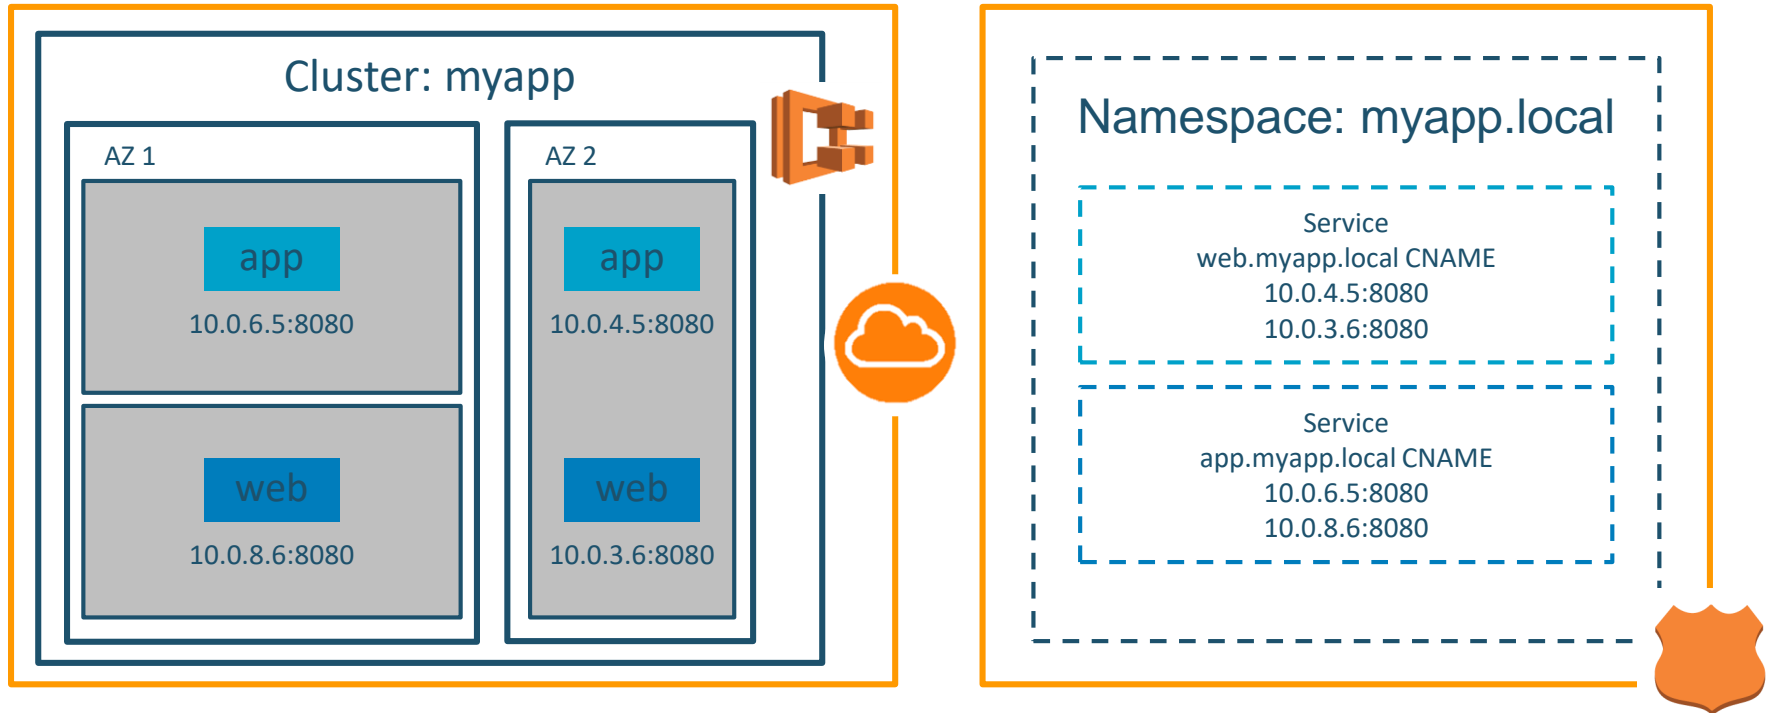


ECS Scheduler updates on:

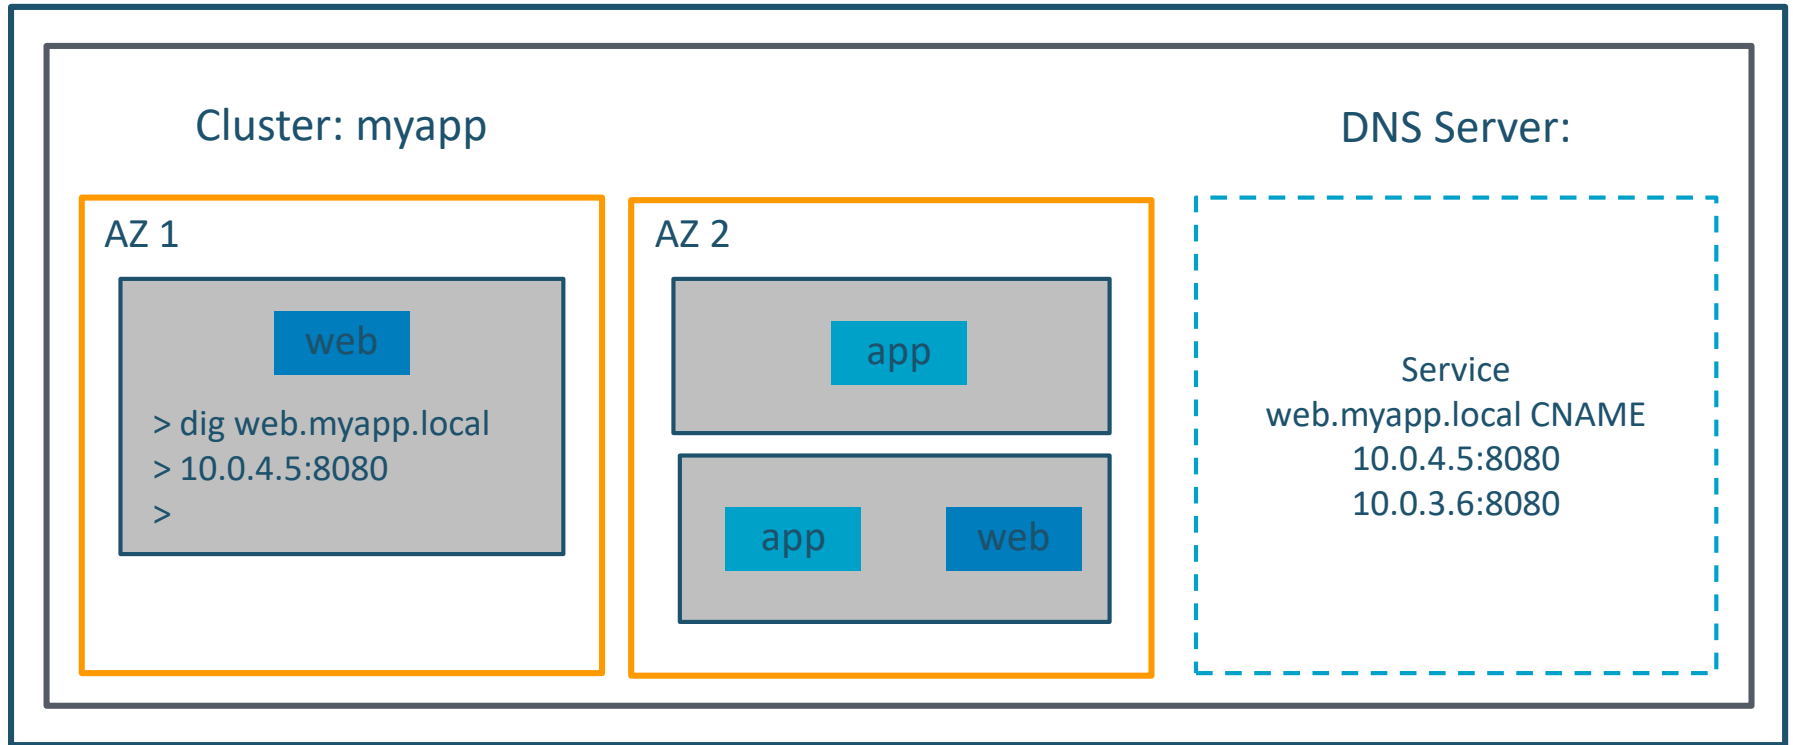
- Service scaling
 - Task registrations
 - Task de-registrations
- Task health
- Scheduling / Placement changes
- ECS instance changes

ECS maintains latest state of the dynamic environment in Service Registry

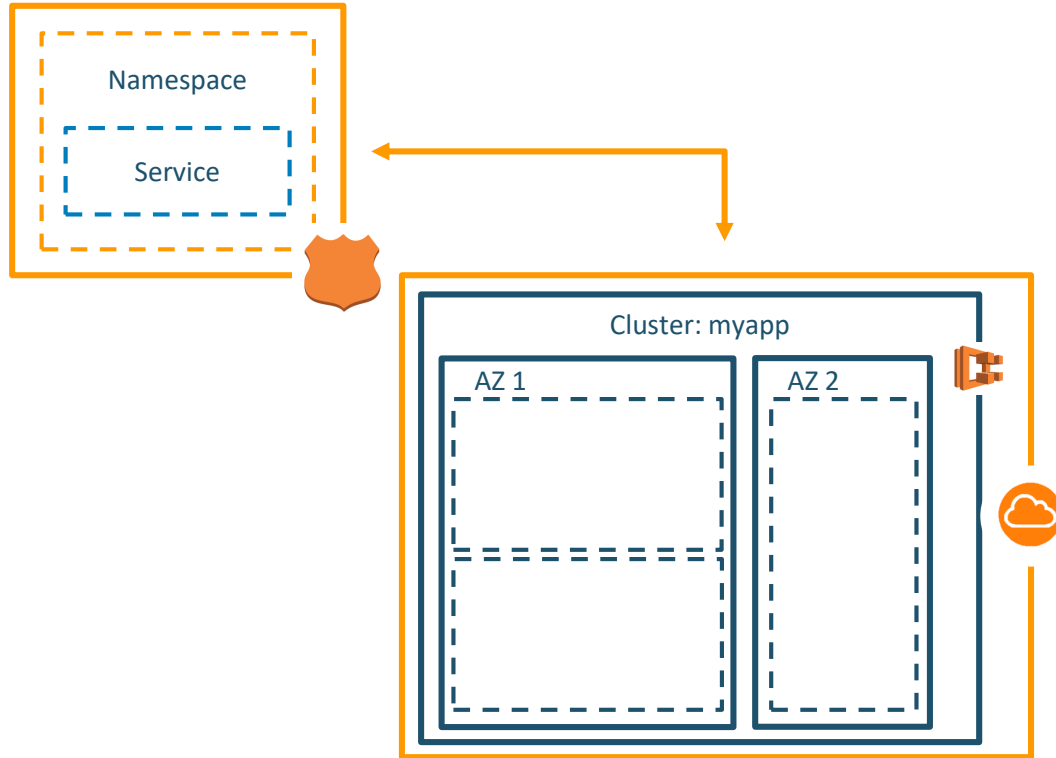
ECS updates service endpoints in Route 53



Services connect to latest endpoints via DNS



Benefits of this approach



Managed

- Setup once, use forever

Highly available

- Same availability and scalability as Route 53


Extensible

- Public APIs that can be used across AWS
- Works across clusters, accounts, AZs
- Works across AWS services

Amazon Route 53 Auto Naming

What is Route 53 Auto Naming?

API that powers ECS
Service Discovery



Service name registration and
management tool



DNS-based service discovery
mechanism




Introducing new abstractions

Namespace

Service

Service Instance



Namespace



Service



Instance
172.10.0.1


Namespace

Logical group of services

Private or public visibility

Service-linked hosted zone in Amazon

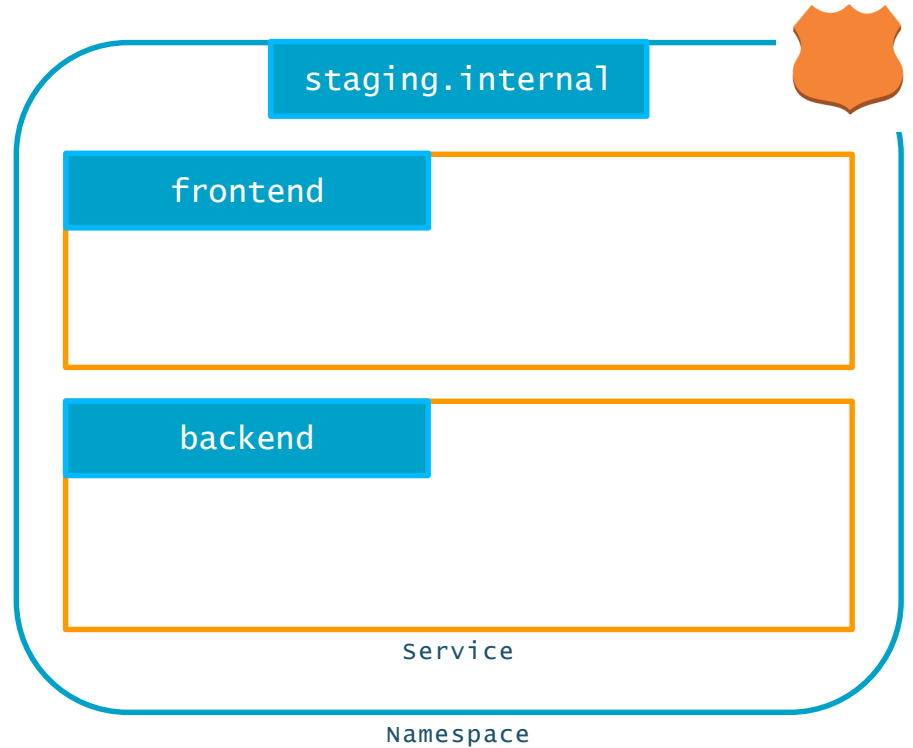
Route 53



Namespace

Namespace management APIs:

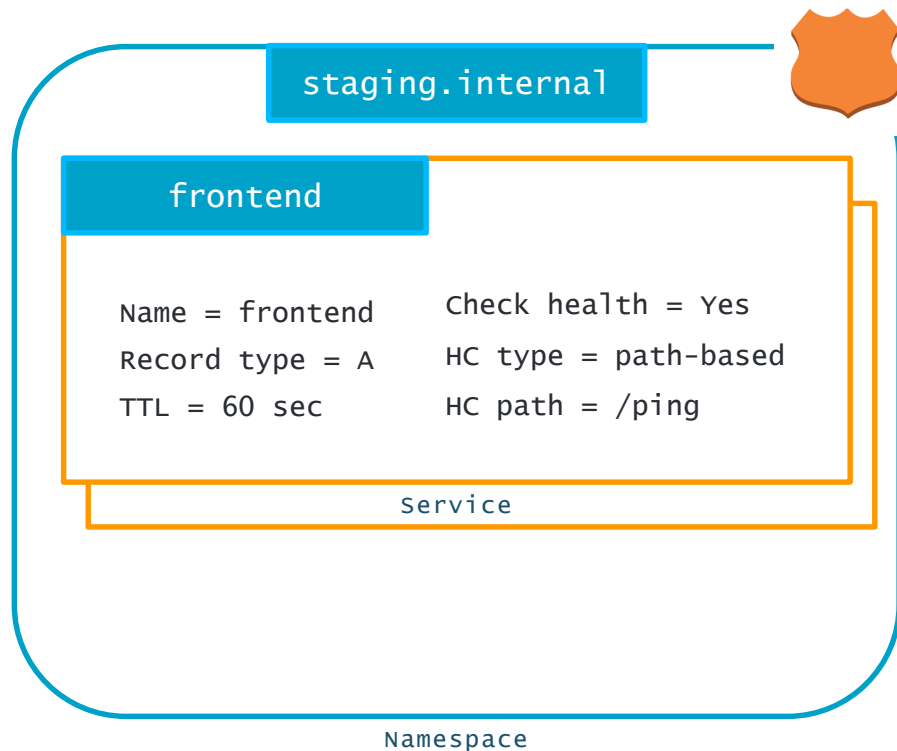
`createPublicDnsNamespace`
`createPrivateDnsNamespace`
`deleteNamespace`
`listNamespaces`
`getNamespace`



Service

ECS: application component run on one or many tasks

Auto Naming: service naming and health checking template



Service

Service management APIs:

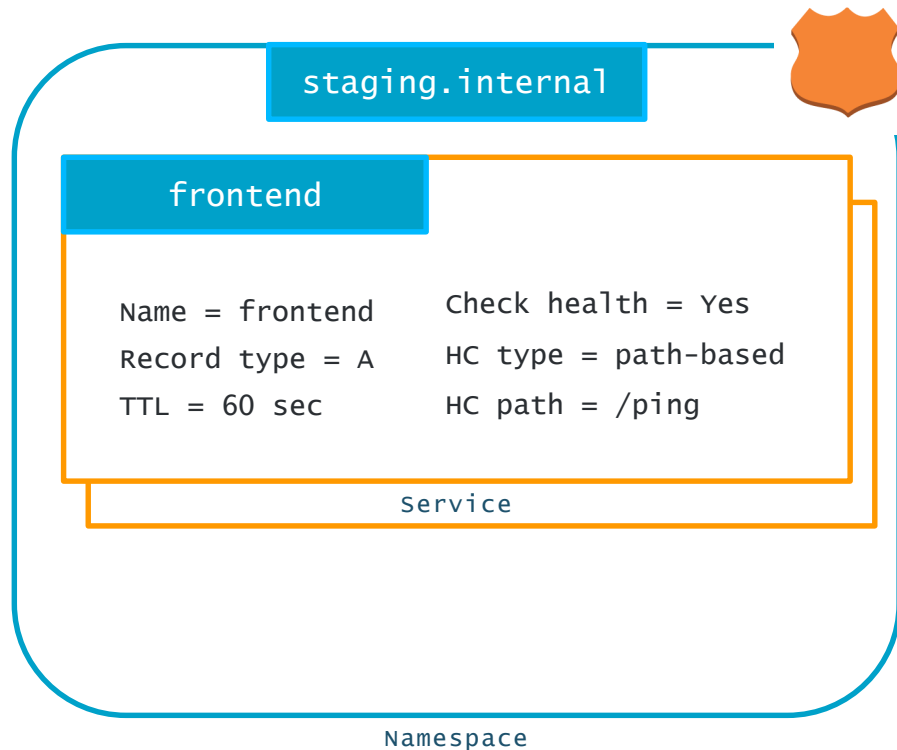
createService

deleteService

updateService

listServices

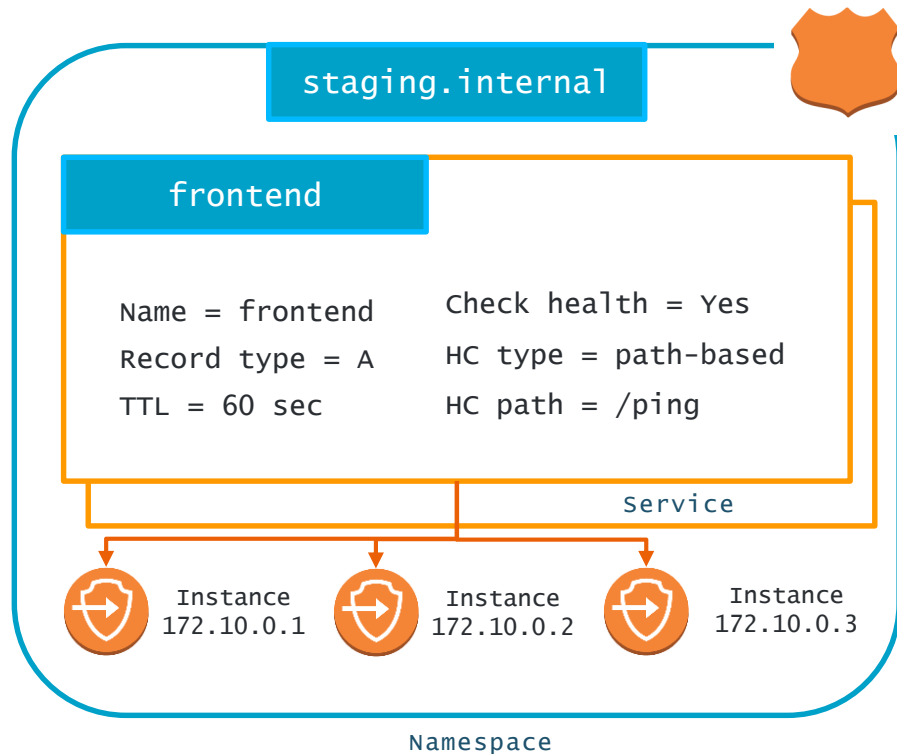
getService



Service Instance

Represents an actual endpoint

Addressable by *ip* or *ip:port*



Service Instance

Service Instance management APIs:

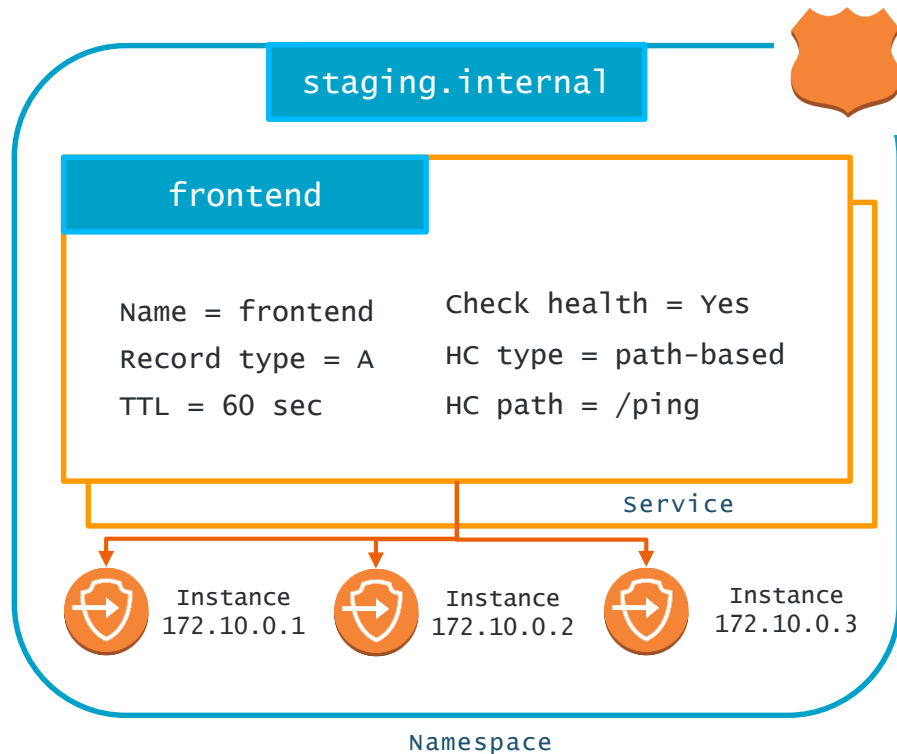
registerInstance

deregisterInstance

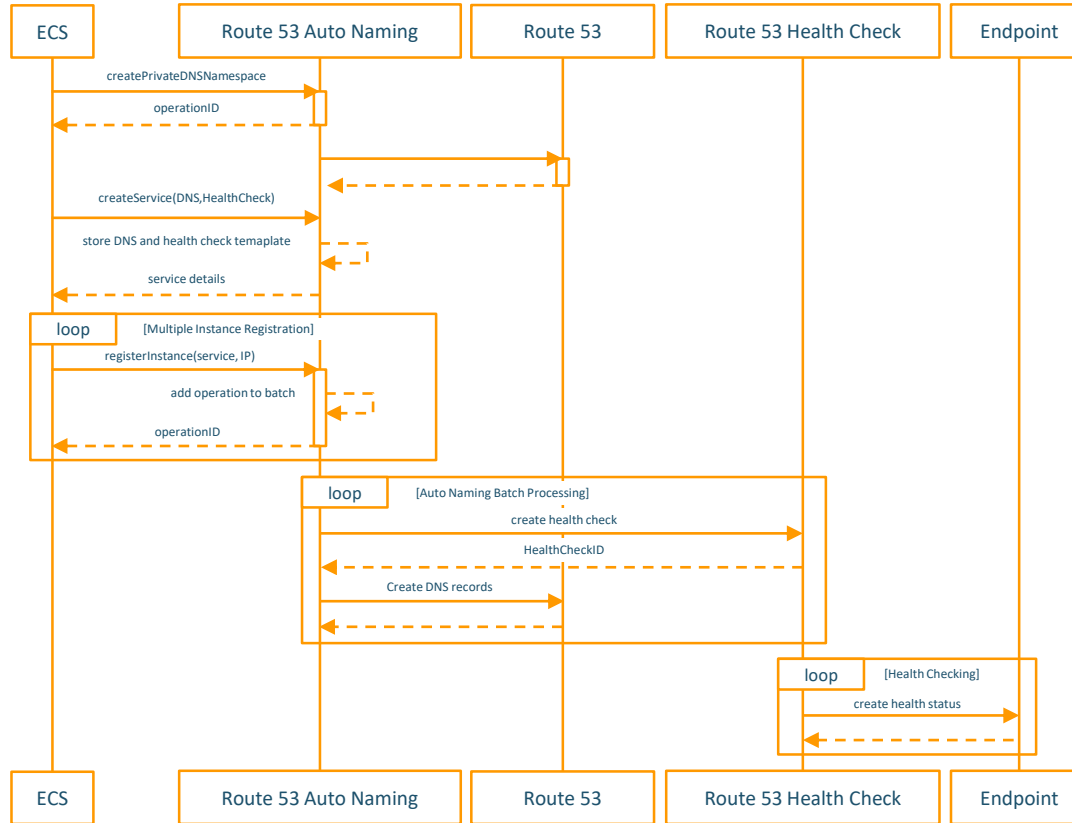
listInstances

getInstance

getInstanceHealthStatus



Route 53 Auto Naming Workflow



Discovery over DNS

Auto Naming uses Route 53
Multivalue Answer Routing



Each DNS query returns up to 8
healthy* endpoints



Client-side load balancing



* If health-checking is enabled for the corresponding Route 53 Auto Naming Service

Amazon Route 53 Auto Naming benefits

Managed API designed for
micro-service architectures



Regional control plane for
better availability



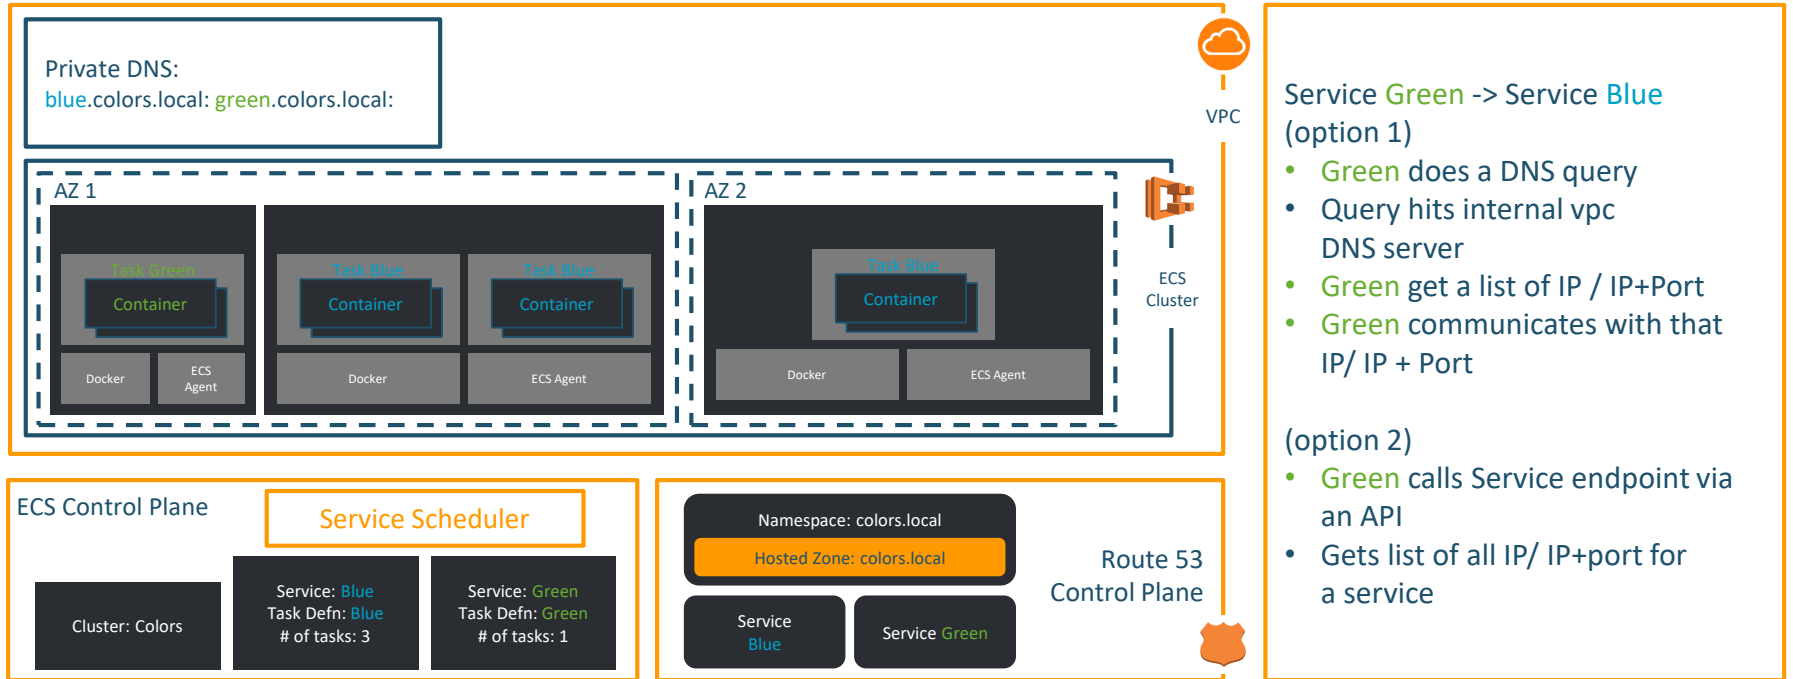
Discovery via Route 53 DNS
with 100% availability SLA



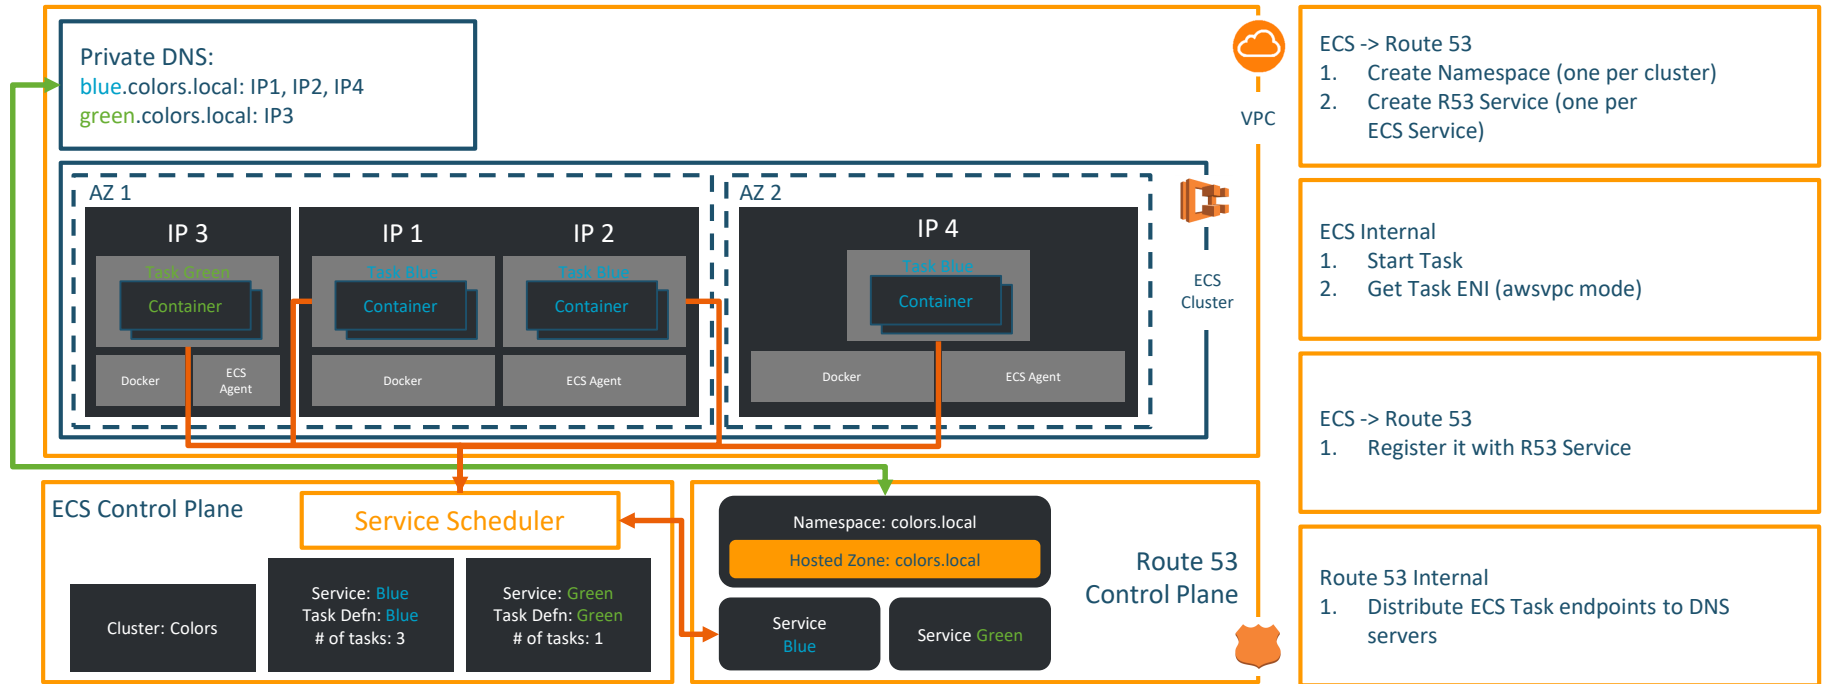
Under the Hood



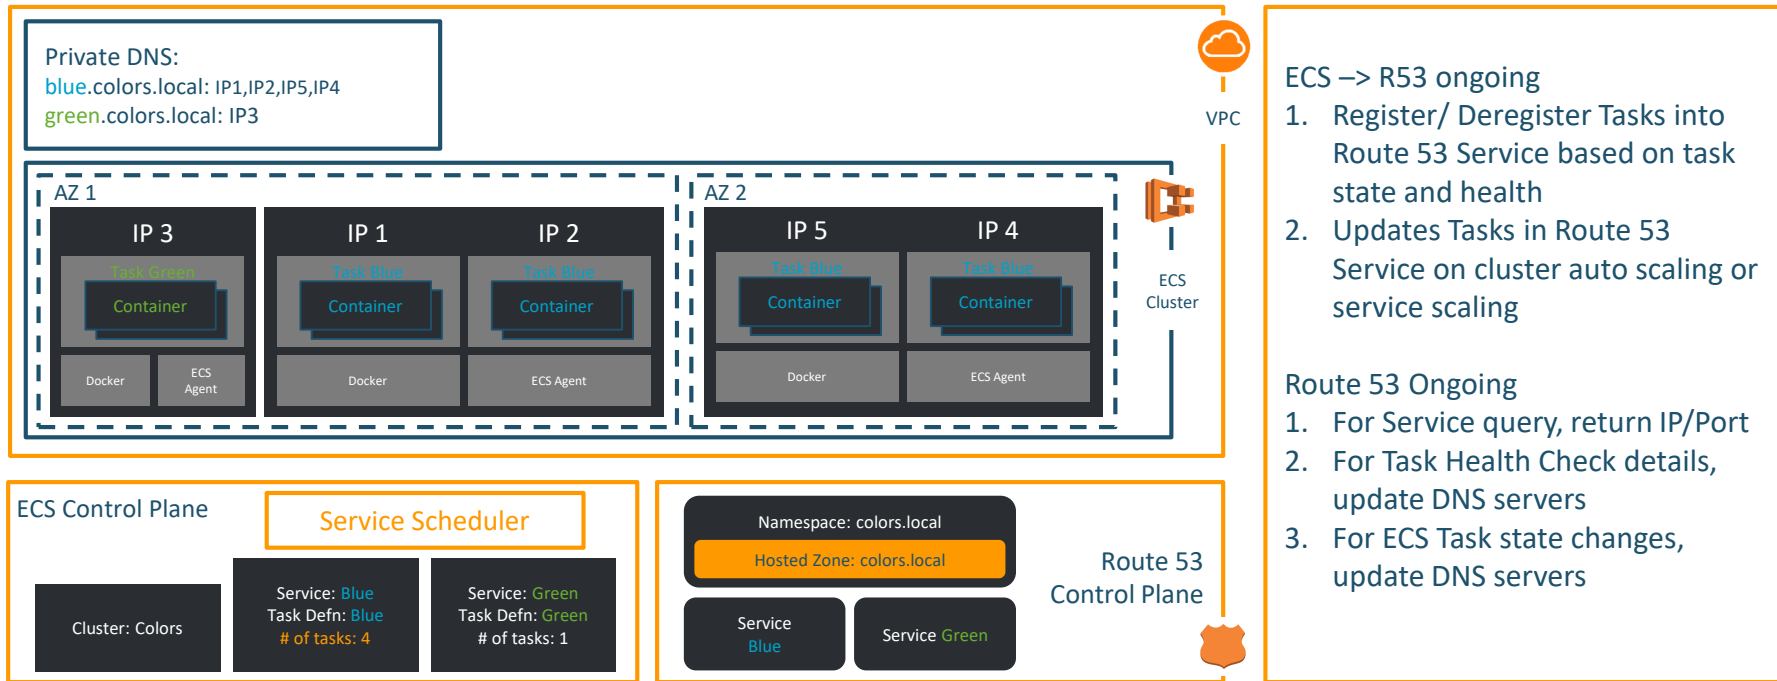
Services can discover other services using DNS



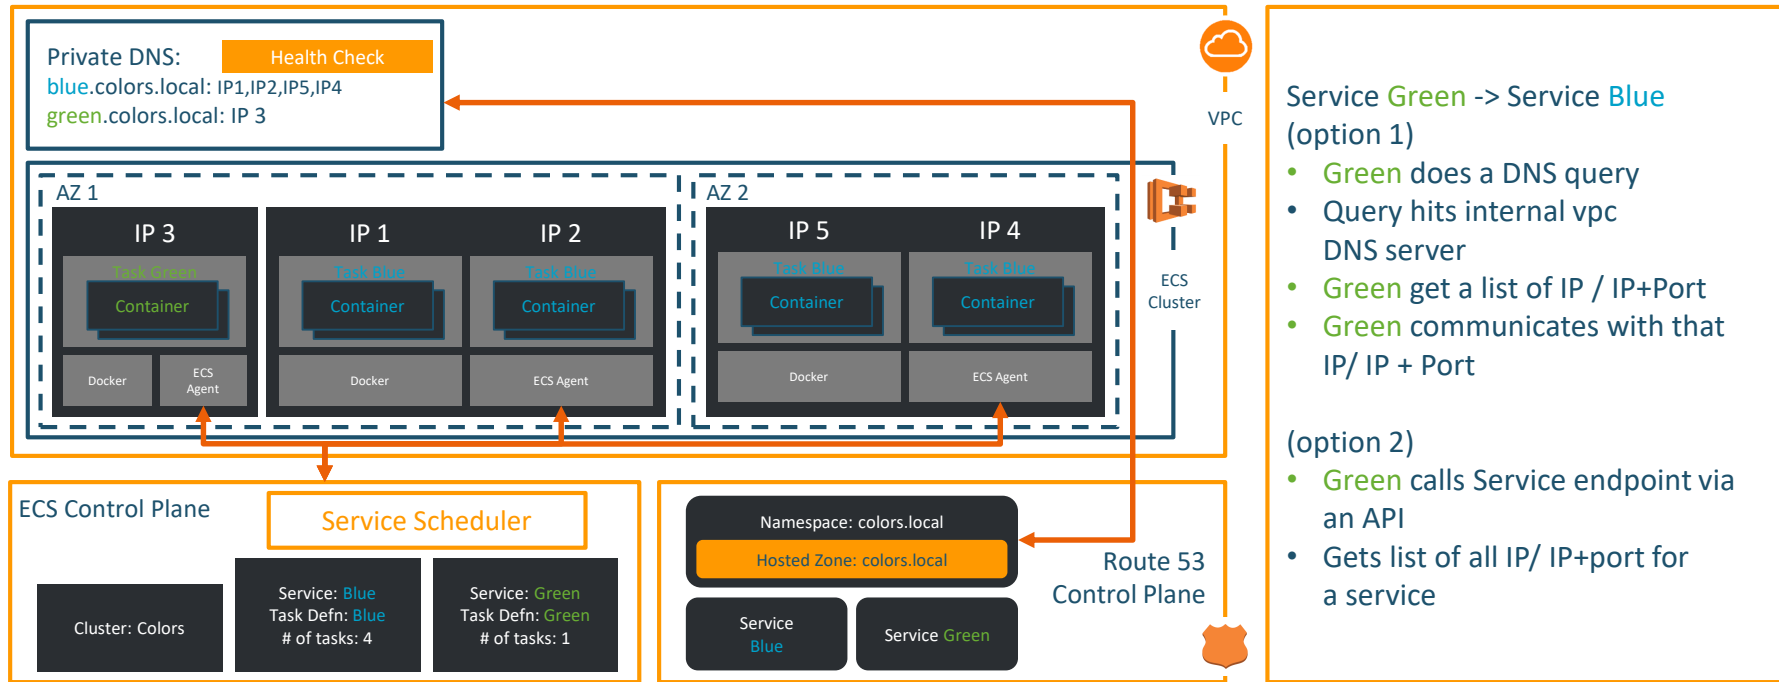
Services can discover other services using DNS



Services can discover other services using DNS



Services can discover other services using DNS



Requirements & Limits

- Tasks must use “awsipc” network mode
- Must specify container health check endpoint for essential containers in task
- Console workflow assumes 1:1:1 mapping between vpc:cluster:app
- Dependent on Route 53 propagation delay for registering new IPs (scaling) or stopping traffic to unhealthy tasks
- Client side logic needed for query, resolution and retries, 8 DNS records fetched in each query
- Private VPC: container health check only, no way to add on network health

How do I use it?

- Native integration with Amazon ECS

[Works with Fargate and EC2 launch types](#)

- Also available for Kubernetes!

<https://github.com/kubernetes-incubator/external-dns>

- Available in all AWS regions where Amazon ECS and Route 53 Auto Naming are available

<https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services>

Thank You!

Learn more at <https://aws.amazon.com/ecs/resources>



peckn@amazon.com



nathankpeck



nathanpeck