

# VPC Lattice: yet another connectivity option or a game-changer?

12 April 2024 - 4 min. read

[Advanced Networking](#)

[VPC Lattice](#)

"Live as if you were to die tomorrow. Learn as if you were to live forever"

- *Mahatma Gandhi*

I like learning new things, and I am lucky because a significant part of my work involves learning new technologies and evaluating how to use them to solve new and existing problems.

In the past few years, we described many connectivity options to connect our workloads running on AWS ([Transit Gateways](#), Load balancing, [Shared VPCs](#), [PrivateLink and Endpoint Services](#), and VPC peering...)

Last year, **VPC Lattice** reached General Availability, offering a new option for us to explore. You may be wondering why we should learn how to use another service when we already have so many options to choose from... Well, let's find out!

In this article, we'll see that we needed to add another thing... Let's see what this service offers, how it differs from the other connectivity options we already know, and a simple use case.

So, without any further ado, let's dig in!

## Where VPC Lattice stands

As we saw in previous articles, microservices are an architectural approach that allows developers to build and deploy software faster. How

Was this article useful for you? 

single applications now consist of numerous individual services distributed in different AWS accounts and must coordinate. Various technologies can be used together: Lambda, ECS, EKS, etc.

We must face tasks like service discovery, traffic routing, and detailed metrics. There are scenarios where a strict authorization policy is complex to implement using IAM.

Think about allowing communication only between two microservices in different AWS accounts, blocking every other traffic without implementing custom solutions and authentication.

VPC Lattice aims to help organizations overcome these challenges while keeping a secure and manageable approach.

It enables developers to **manage microservices definition, authorization, and configuration, while administrators can manage infrastructure, network, communication policies, and governance**; we'll see more in the following sections.

## VPC Lattice Components

VPC Lattice has some key components:

- **Service:** a VPC Lattice service represents an application and its resources: computational (such as lambdas, EKS pods, EC2 instances), listeners (port and protocols involved), and routing rules to address traffic (weighted routing, path routing, and so on). The application team typically manages these aspects and can be autonomously defined. A service can be shared using AWS Resource Access Manager with other AWS accounts, even outside an organization.
- **Service Network:** a service network is a new concept because it's a logical grouping of services that facilitates connectivity. It can also be shared using AWS Resource Access Manager. This kind of resource is usually defined by infrastructure or network administrators.
- **Service Directory:** the list of service networks that can be used, shared with AWS RAM
- **Auth Policies:** IAM documents that can define access to resources. Auth policies differ from IAM policies because they aren't attached to users, groups, or roles but are used with services and service networks. The

Was this article useful for you?

Yes!

I was expecting something different

Made with Hotjar

Next



or other services access to the attached service.  
conditions, such as

```
vpc-lattice-svcs:RequestMethod
```

to filter allowed methods. For a complete list, you can refer to the [documentation](#).

Let's see which steps are required to set up VPC Lattice.

## Putting all together

Let's briefly look at the roles and steps involved in setting up the communication.

1. The infrastructure admin defines a **Service Network**
2. The infrastructure admin shares the Service Network using **AWS RAM**
3. The service owner defines a **VPC Lattice Service** and its authorization policies
4. The service owner associates the **Service** with the Service Network
5. The infrastructure admin **associates** the Service Network with the **VPCs**

As you can see, there is a certain degree of independence in managing services, so dependencies between teams and operations are reduced.

## A practical example of VPC Lattice usage

For our example, we will use the sample infrastructure and applications from [this URL](#) and analyze the key components deployed on the AWS Console.

You can deploy the reference architecture using the instructions contained in the *ReadMe* file in the repository.

You can also use three different VPCs belonging to a single account to ease the deployment process (as we'll do in this article).

You'll end with this sample infrastructure:

Was this article useful for you?

Yes!

I was expecting something different

Made with Hotjar



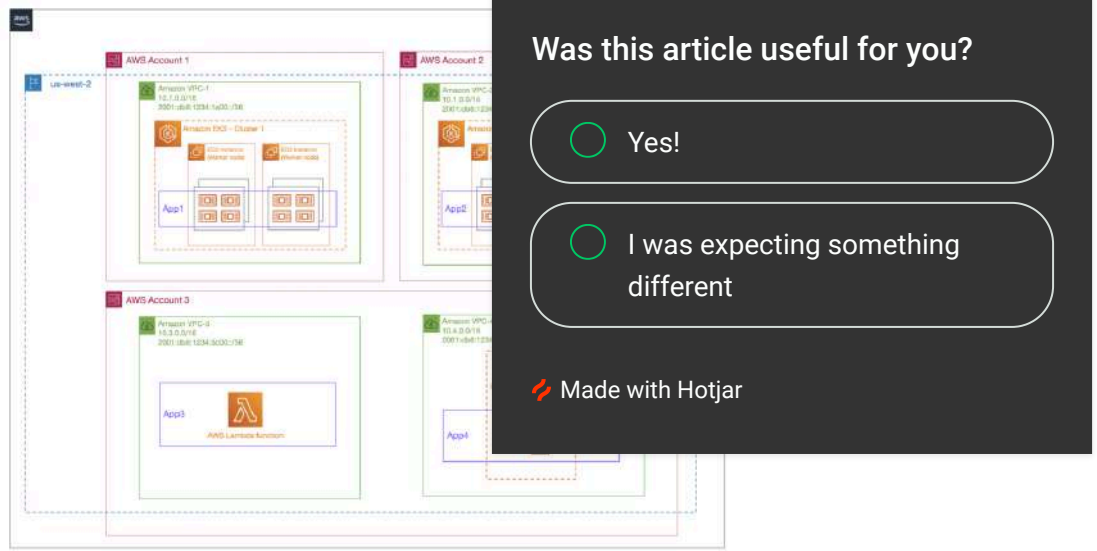


Image source: <https://github.com/aws-samples/build-secure-multi-account-vpc-connectivity-applications-with-amazon-vpc-lattice>

In our example, App1 and App2 will need to communicate bidirectionally, and App3 will be consumed by App1, App2, and App4.

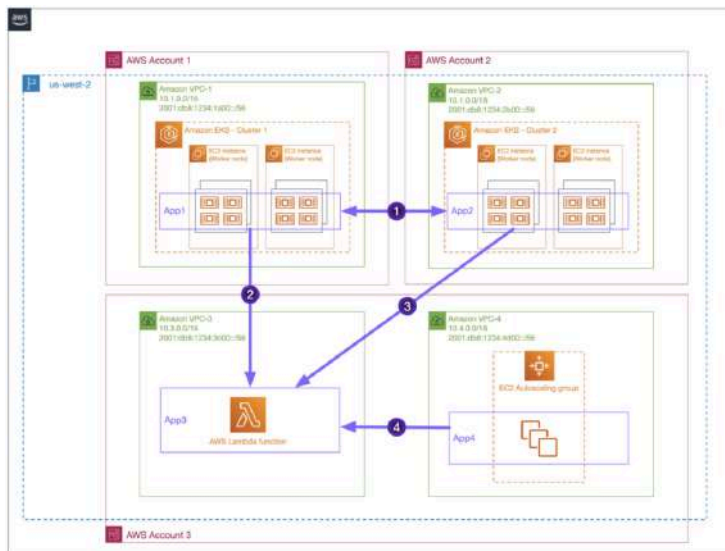


Image source: <https://github.com/aws-samples/build-secure-multi-account-vpc-connectivity-applications-with-amazon-vpc-lattice>

The deployment process will take about 10 - 15 minutes, and you will see that our microservices can communicate.

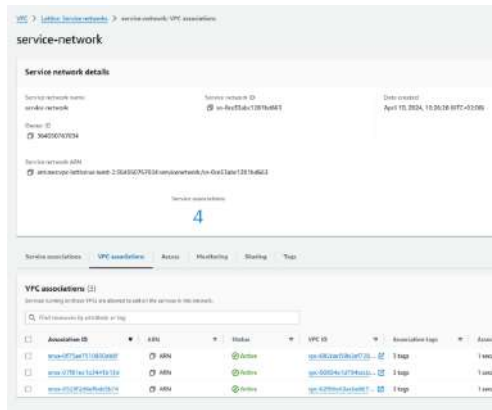
**Bug Alert:** there's a missing step in the deployment instructions; these two lines will make the deployment fail:

```
export TARGETCLUSTER1={TARGET_GROUP_ARN}
export TARGETCLUSTER2={TARGET_GROUP_ARN}
```

You need to change these exports using values from







Was this article useful for you?

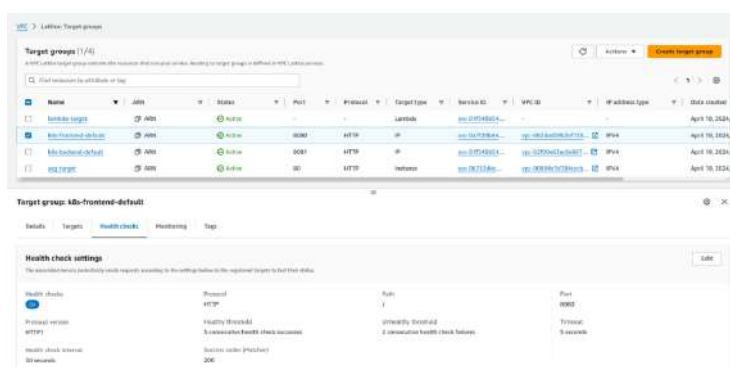
Yes!

I was expecting something different

Made with Hotjar

Now VPC Lattice target groups:

Our Lambda, Autoscaling group, and EKS are added as targets, and like ELB target groups, you can see metrics and define health checks.



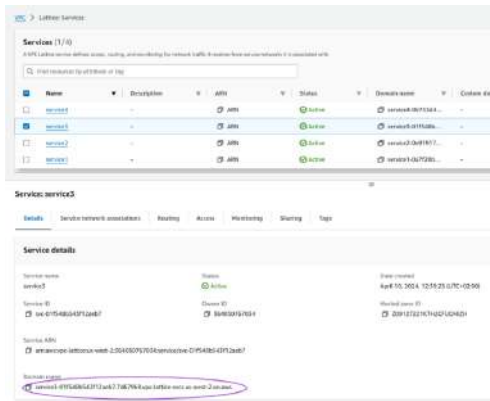
Note that to register our targets managed by EKS clusters, we had to deploy the Gateway API Controller Class into the cluster and register the service using the configuration deployed with the commands:

```
kubectl config use-context cluster1
kubectl apply -f ./vpc-lattice/routes/frontend-export.yaml
```

```
kubectl config use-context cluster2
kubectl apply -f ./vpc-lattice/routes/backend-export.yaml
```

Last but not least, let's see the defined services. You can access them using the DNS domain name defined by VPC Lattice





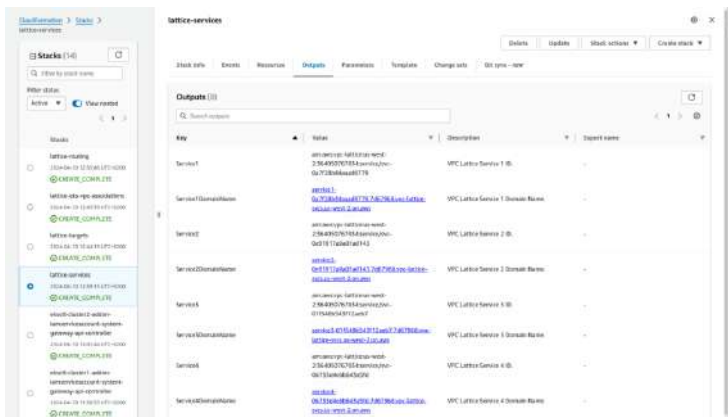
Was this article useful for you?

Yes!

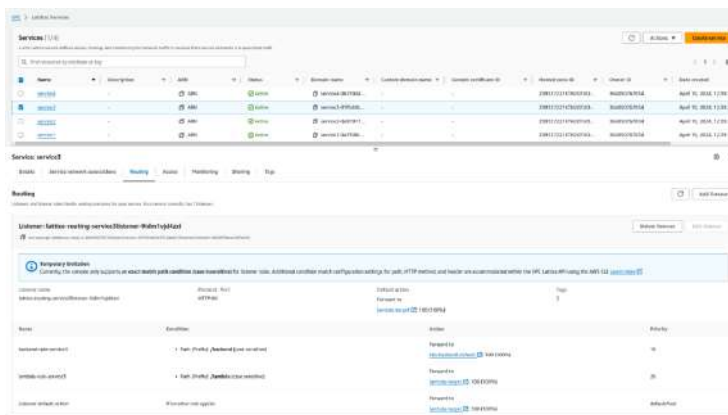
I was expecting something different

Made with Hotjar

You can find the complete list of domain names defined by looking at the Outputs of the CloudFormation stack named "lattice-services":



Each service also has its routing; in this case, you can see that /backend path for service3 forwards to cluster2, while /lambda forwards requests to the lambda application:



Now that you have a deployed infrastructure, feel free to experiment and modify the behaviour of your application and environment (why don't you try IAM authentication between services? ).

This is only a starting point, and you can also find an excellent AWS workshop with a good walkthrough [here](#).

**Last but not least: is VPC Lattice the right t**





The correct answer is, as always: "It depends". Let's see how this will impact on our design.

The first is the mandatory **one-to-one association between services and VPCs**. **Network**. Once you define an association, you cannot change it. You cannot connect other service networks, so you need to plan the impact of each service.

You cannot share services between different regions. This is a **regional construct**: you'll need to use a hybrid approach, resorting to traditional connectivity methods, like Transit Gateway Peering and VPC Peering.

On the other hand, there are great **benefits**: you can free service owners to configure and offer by themselves their components while keeping centralized control in a complex environment.

We didn't have a look at IAM policies, but **implementing IAM Authentication and Authorization is the key to a secure microservice implementation**.

VPC Lattice also offers a complete service mesh solution because it can link EC2, EKS, ECS, and Lambda workloads.

Cost is another important factor: **as your service mesh grows, costs increase proportionally**.

You pay 0.025\$/hour for each service. If you have 10 services you pay 182.5\$/month. A different approach can have a smaller impact on AWS billing for complex architectures with hundreds of microservices.

## That's all for today!

Microservices aren't simple: they can increase the overall complexity of the architecture and require good planning and management over time.

Have you already had the chance to experiment with VPC Lattice? What are your thoughts about this new approach?

Let us know in the comments!

Was this article useful for you?

Yes!

I was expecting something different

Made with Hotjar





**Proud2beCloud** is a blog by **beSharp**, an Italian AP in designing, implementing, and managing complex advanced services on AWS. Before being writers, we worked with AWS services since 2007. We are hungry readers and knowledge seekers. On Proud2beCloud, we regularly share our insights, in-depth news, tips&tricks, how-tos, and more on the discussion!

Was this article useful for you?

Yes!

I was expecting something different

Made with Hotjar



## Damiano Giorgi

Ex on-prem systems engineer, lazy and prone to automating boring tasks. In constant search of technological innovations and new exciting things to experience. And that's why I love Cloud Computing! At this moment, the only "hardware" I regularly dedicate myself to is that my bass; if you can't find me in the office or in the band room try at the pub or at some airport, then!

Copyright © 2011-2024 by beSharp spa - P.IVA IT02415160189

