AWS Wavelength

Developer Guide

AWS Wavelength: Developer Guide

Table of Contents

What is AWS Wavelength?

AWS Wavelength enables developers to build applications that deliver ultra-low latencies to mobile devices and end users. Wavelength deploys standard AWS compute and storage services to the edge of telecommunication carriers' 5G networks. You can extend an Amazon Virtual Private Cloud (VPC) to one or more Wavelength Zones. You can then use AWS resources like Amazon Elastic Compute Cloud (Amazon EC2) instances to run the applications that require ultra-low latency and a connection to AWS services in the Region.

For more information, see AWS Wavelength.

Wavelength concepts

The following are the key concepts for Wavelengths:

- **Wavelength** — A new type of AWS infrastructure designed to run workloads that require ultra-low latency over mobile networks.
- **Wavelength Zone (WZ)** — A zone in the carrier location where the Wavelength infrastructure is deployed. Wavelength Zones are associated with an AWS Region. A Wavelength Zone is a logical extension of the Region, and is managed by the control plane in the Region.
- **VPC** — A customer virtual private cloud (VPC) that spans Availability Zones, Local Zones, and Wavelength Zones, and has deployed resources such as Amazon EC2 instances in the subnets that are associated with the zones.
- **Subnet** — A subnet that you create in a Wavelength Zone. You can create one or more subnets, and then run and manage AWS services, such as Amazon EC2 instances, in the subnet.
- **Carrier gateway** — A carrier gateway serves two purposes. It allows inbound traffic from a carrier network in a specific location, and allows outbound traffic to the carrier network and internet.
- **Network Border Group** — A unique set of Availability Zones, Local Zones, or Wavelength Zones from which AWS advertises IP addresses.
- **Wavelength application** — An application that you run on an AWS resource in a Wavelength Zone.

AWS resources on Wavelength

You can create Amazon EC2 instances, Amazon EBS volumes, and Amazon VPC subnets and carrier gateways in Wavelength Zones. You can also use services that orchestrate or work with EC2, EBS, and VPC, such as:

- Amazon EC2 Auto Scaling
- Amazon EKS clusters
- Amazon ECS clusters
- Amazon EC2 Systems Manager
- Amazon CloudWatch
- AWS CloudTrail
- AWS CloudFormation

The services in Wavelength are part of a VPC that is connected over a reliable connection to an AWS Region for easy access to services running in Regional subnets.

Working with Wavelength

You can create, access, and manage your EC2 resources, Wavelength Zones, and carrier gateways using any of the following interfaces:

- **AWS Management Console**— Provides a web interface that you can use to access your Wavelength resources.
- **AWS Command Line Interface (AWS CLI)** — Provides commands for a broad set of AWS services, including Amazon VPC, and is supported on Windows, macOS, and Linux. The services you use in Wavelength continue to use their own namespace, for example Amazon EC2 uses the "ec2" namespace, and Amazon EBS uses the "ebs" namespace. For more information, see AWS Command Line Interface.
- **AWS SDKs** — Provides language-specific APIs and takes care of many of the connection details, such as calculating signatures, handling request retries, and handling errors. For more information, see AWS SDKs.

When you use any of the interfaces for your Wavelength Zones, use the parent Region.

Pricing

For information about pricing, see the pricing pages for the Region and services that you use with Wavelength, for example Amazon EC2. For more information about pricing, see AWS Pricing.

How AWS Wavelength works

The following diagram demonstrates how you can create a subnet that uses resources in a telecommunication carrier network at a specific location. You create a VPC in the Region. For resources that need to be within the telecommunication carrier network, you opt in to the Wavelength Zone, and then create resources in the Wavelength Zone.

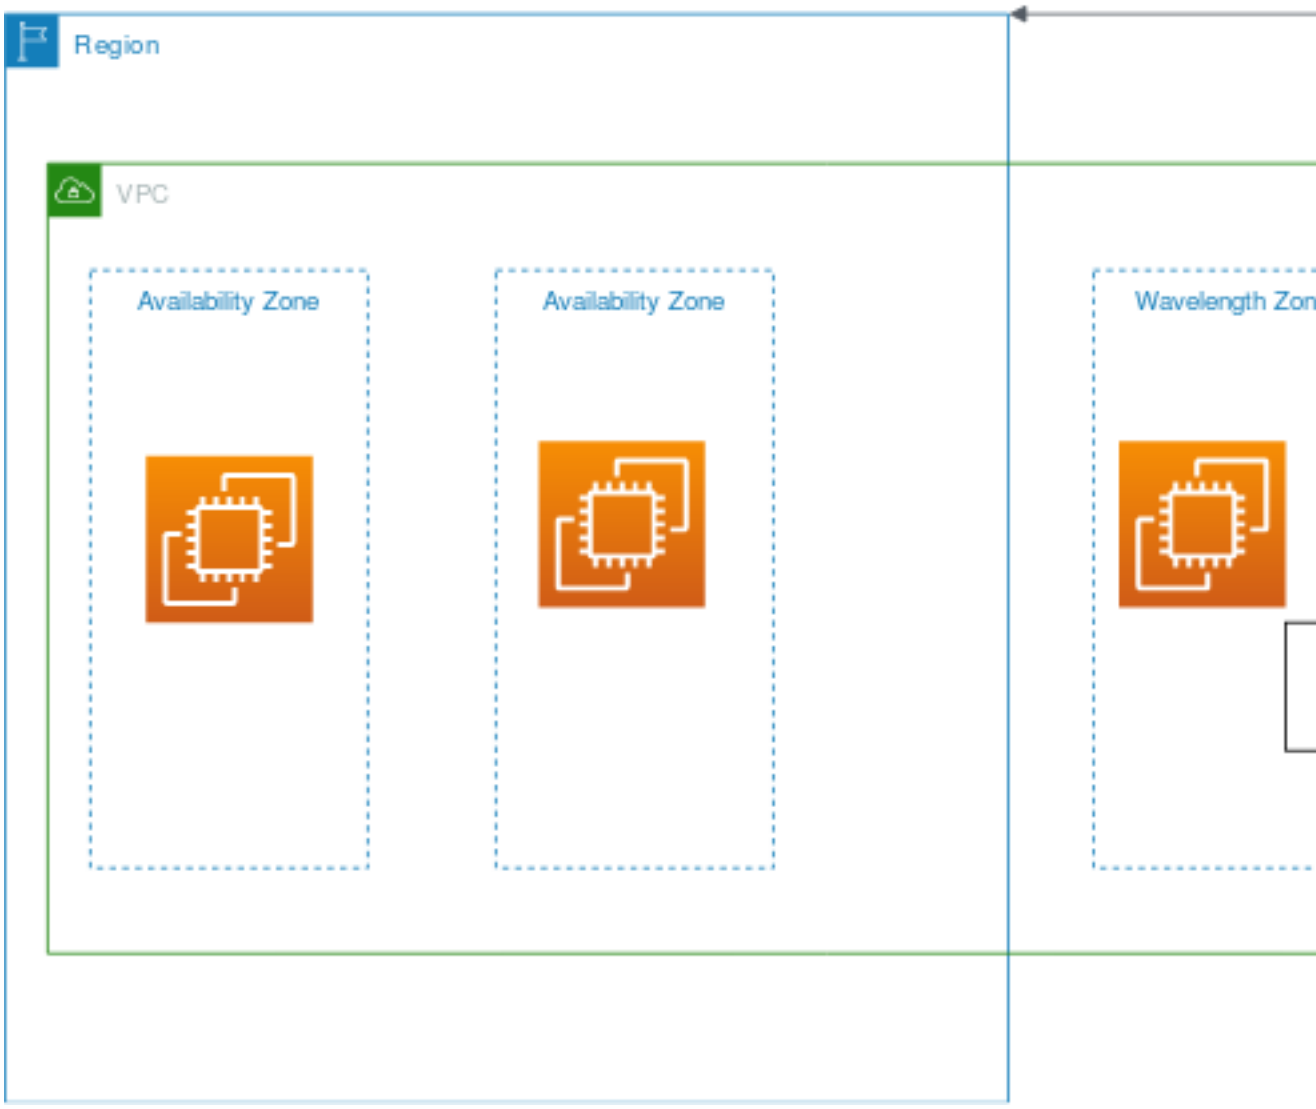

Topics

- Architect apps for Wavelength (p. 6)
- Use cases (p. 7)

VPCs

After you create a VPC in a Region, create a subnet in a Wavelength Zone that is associated with the VPC. In addition to the Wavelength Zone, you can create resources in all of the Availability Zones and Local Zones that are associated with the VPC.

You have control over the VPC networking components, such as IP address assignment, subnets, and route table creation.

VPCs that contain a subnet in a Wavelength Zone can connect to a carrier gateway. A carrier gateway allows you to connect to the following resources:

- 4G/LTE and 5G devices on the telecommunication carrier network
- Outbound traffic to public internet resources

Subnets

Any subnet that you create in a Wavelength Zone inherits the main VPC route table, which includes the local route. The local route enables connectivity between the subnets in the VPC, including the subnets that are in the Wavelength Zone.

AWS recommends that you configure custom route tables for your subnets in Wavelength Zones. The destinations are the same destinations as a subnet in an Availability Zone or Local Zone, with the addition of a carrier gateway. For more information, see the section called “Routing” (p. 5).

Carrier gateways

A carrier gateway serves two purposes. It allows inbound traffic from a carrier network in a specific location, and it allows outbound traffic to the carrier network and internet. There is no inbound connection configuration from the internet to a Wavelength Zone through the carrier gateway.

A carrier gateway supports IPv4 traffic.

Carrier gateways are only available for VPCs that contain subnets in a Wavelength Zone. The carrier gateway provides connectivity between your Wavelength Zone and the telecommunication carrier, and devices on the telecommunication carrier network. The carrier gateway performs NAT of the Wavelength instances' IP addresses to the Carrier IP addresses from a pool that is assigned to the network border group. The carrier gateway NAT function is similar to how an internet gateway functions in a Region.

Carrier IP address

A *Carrier IP address* is the address that you assign to a network interface, which resides in a subnet in a Wavelength Zone (for example an EC2 instance). The carrier gateway uses the address for traffic from the interface to the internet or to mobile devices. The carrier gateway uses NAT to translate the address, and then sends the traffic to the destination. Traffic from the telecommunication carrier network routes through the carrier gateway.

You allocate a Carrier IP address from a network border group, which is a unique set of Availability Zones, Local Zones, or Wavelength Zones from which AWS advertises IP addresses, for example, `us-east-1-wl1-bos-wlz-1`.

Routing

You can set the carrier gateway as a destination in a route table for the following resources:

- VPCs that contain subnets in a Wavelength Zone
- Subnets in Wavelength Zones

Create a custom route table for the subnets in the Wavelength Zones so that the default route goes to the carrier gateway, which then sends traffic to the internet and telecommunication carrier network.

Example: Carrier gateway routing to the public internet

Consider a scenario with the following configuration:

- A VPC with a CIDR block 10.0.0.0/16.
- A subnet in the VPC with a CIDR block 10.0.3.0/24.
- An EC2 instance in the subnet with a private IP address 10.0.3.112.
- A Carrier IP address (198.51.100.130) for the network interface associated with the EC2 instance.
- An IP address association that maps the EC2 instance private IP address (10.0.3.112) to a Carrier IP address (198.51.100.130).

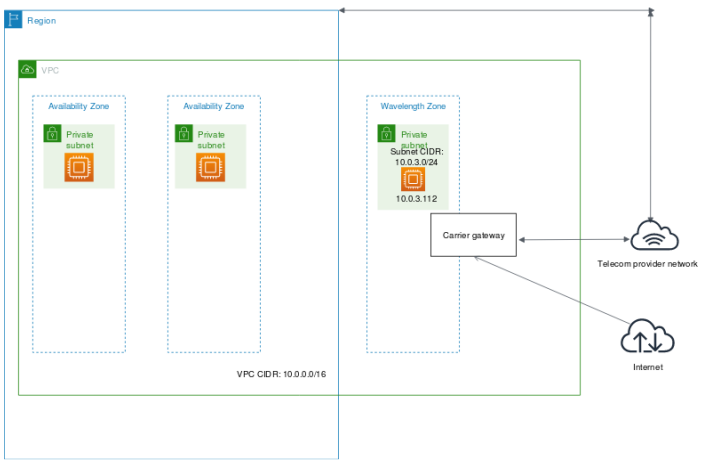

You need the following entries in the Wavelength subnet route table.

Destination	Target	Notes
10.0.0.0/16	Local	This route allows for intra-VPC connectivity, including subnets in the Region.

Destination	Target	Notes
0.0.0.0/0	*carrier-gateway-id*	The Carrier IP address provides internet connectivity through the carrier gateway.

Carrier gateway access to the public internet

The carrier gateway provides access to the internet from your Wavelength subnets. For information about protocol considerations, see the section called “Networking considerations” (p. 19)

Traffic initiated from the EC2 instance for the internet uses the 0.0.0.0/0 route to route traffic to the carrier gateway. The carrier gateway maps the EC2 instance IP address to the Carrier IP address (198.51.100.130), and then sends the traffic to the telecommunication carrier.

DNS

EC2 instances use EC2 DNS to resolve domain names to IP addresses. Route 53 supports DNS features, such as domain registration, and DNS routing. Both public and private hosted Wavelength Zones are supported for routing traffic to specific domains. Route 53 resolvers are hosted in the Region.

You can also use your own DNS services to resolve domain names.

Architect apps for Wavelength

Wavelength Zones are designed for the following workloads:

- Applications that need to connect to compute from 5G mobile devices with ultra-low latency
- Applications that need consistent data rates from mobile devices to compute in a Wavelength Zone

Review *Considerations and quotas* (p. 18), which includes information about available Wavelength Zones, service differences, and Service Quotas.

Consider the following factors when using Wavelength Zones:

- Wavelength Zones are designed for application components that are latency-sensitive.
- AWS recommends that you architect the edge applications in a hub and spoke model with the Region to provide the most scalable, resilient, and cost effective options for components. For more information, see the section called “Best practices” (p. 7)
- Services that run in Wavelength Zones have different compliance than services in an AWS Region. For more information, see the section called “Compliance validation” (p. 16).

Wavelength Zones have network access that is specific to a telecommunication carrier and location. Therefore, you might need to have multiple Wavelength Zones for your latency-sensitive applications to meet your latency requirements. For more information, see the section called “Networking considerations” (p. 19).

Discover the closest Wavelength Zone endpoint

You can use the following procedures to have client devices discover the closest Wavelength Zone endpoint, for example an EC2 instance:

- Register the instance with a discovery service such as AWS Cloud Map. For information about how to register an instance, see Registering Instances in the *AWS Cloud Map Developer Guide*.
- Applications that run on client devices can run latency tests such as `ping` from the client to select the best endpoint that is registered in AWS Cloud Map, or can use the geolocation data from the mobile device.

Best practices

Run the following components in the Region:

- Components that are less latency sensitive
- Components that need to be shared across Zones
- Components that need to persist state, such as databases

Run the application components that need ultra-low latency and higher bandwidth over 5G mobile networks in Wavelength Zones.

For optimal throughput, AWS recommends that you use a public service endpoint when applications in the Wavelength Zone need to connect to AWS services in the parent Region.

Use cases

Using AWS Wavelength Zones can help you accomplish a variety of goals. This section lists a few to give you an idea of the possibilities.

Connected vehicles

Cellular Vehicle-to-Everything (C-V2X) is an increasingly important platform for enabling functionality such as intelligent driving, real-time HD maps, and increased road safety. Low latency access to the compute infrastructure that's needed to run data processing and analytics on AWS Wavelength enables real-time monitoring of data from sensors on the vehicle. This allows for secure connectivity, in-car telematics, and autonomous driving.

Media and entertainment

Wavelength provides the ultra-low latency needed to live stream high-resolution video and high-fidelity audio, and to embed interactive experiences into live video streams. Real-time video analytics provide the ability to generate real-time statistics that enhance the live event experience.

Augmented reality (AR) and virtual reality (VR)

By accessing compute resources on AWS Wavelength, AR/VR applications can reduce the Motion to Photon (MTP) latencies to the benchmark that is needed to offer a realistic customer experience. When you use Wavelength, you can offer AR/VR in locations where it is not possible to run local system servers.

Smart factories

Industrial automation applications use ML inference at the edge to analyze images and videos to detect quality issues on fast moving assembly lines and to trigger actions that address the issues. With Wavelength, these applications can be deployed without having to use expensive, GPU-based servers on the factory floor.

Real-time gaming

Real-time game streaming depends on low latency to preserve the user experience. With Wavelength, you can stream the most demanding games from Wavelength Zones so that they are available on end devices that have limited processing power.

Healthcare

AI/ML driven video analytics and image matching solutions help doctors speed up the diagnosis of observed conditions, such as recognizing polyps. The image or video streams from medical devices are processed in a Wavelength Zone and the response is returned to the surgeon's medical device.

Get started with AWS Wavelength

The following diagram shows the resources that you need to configure to get started using AWS Wavelength.

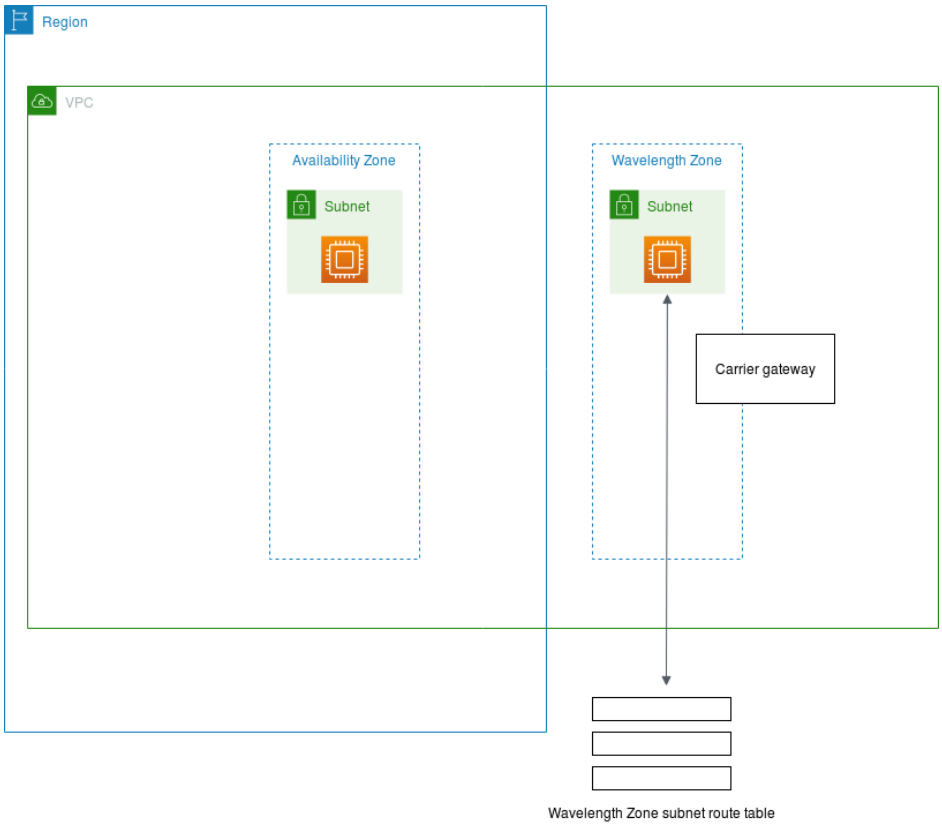

You need to configure the following resources:

- A VPC in your Region
- A carrier gateway

  When you create the carrier gateway, you can also create the subnet. AWS automatically creates the route table, routes, and subnet association.
- A public subnet in an Availability Zone in your Region
- An instance in the public subnet
- An instance in the Wavelength Zone subnet with a Carrier IP address

Before you begin, review *Considerations and quotas* (p. 18), which includes information about available Wavelength Zones, service differences, and Service Quotas.

Tasks

- Step 1: Opt in to Wavelength Zones (p. 10)
- Step 2: Configure your network (p. 10)
- Step 3: Launch an instance in your Availability Zone public subnet (p. 12)
- Step 4: Launch an instance for your Wavelength application (p. 12)
- Step 5: Test the connectivity (p. 14)

Step 1: Opt in to Wavelength Zones

Before you specify a Wavelength Zone for a resource or service, you must opt in to the zone.

You need to request access in order to use Wavelength Zones, before you opt in. For information about how to request Wavelength Zone access, see AWS Wavelength.

Step 2: Configure your network

After you opt in to the Wavelength Zone, create a VPC, a carrier gateway, and a public subnet in the Availability Zone.

Prerequisite

You must have opted in to the Wavelength Zone.

Tasks

- Step 1: Create a VPC (p. 10)
- Step 2: Create a carrier gateway and a subnet associated with the Wavelength Zone (p. 10)
- Step 3: Create a public subnet in an Availability Zone in the Region (p. 11)

Step 1: Create a VPC

Create a VPC using the Amazon VPC Console.

To create a VPC using the console

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **Your VPCs**, **Create VPC**.
3. Specify the following VPC details as necessary:
   - **Name tag**: Optionally provide a name for your VPC. Doing so creates a tag with a key of `Name` and the value that you specify.
   - **IPv4 CIDR block**: Specify an IPv4 CIDR block for the VPC. We recommend that you specify a CIDR block from the private (non-publicly routable) IP address ranges as specified in RFC 1918; for example, `10.0.0.0/16`, or `192.168.0.0/16`.

     Note
     You can specify a range of publicly routable IPv4 addresses. However, we currently do not support direct access to the internet from publicly routable CIDR blocks in a VPC. Windows instances cannot boot correctly if launched into a VPC with ranges from `224.0.0.0` to `255.255.255.255` (Class D and Class E IP address ranges).
4. Choose **Create**.

Step 2: Create a carrier gateway and a subnet associated with the Wavelength Zone

After you create a VPC, create a carrier gateway, and then select the subnets that route traffic to the carrier gateway.

When you choose to automatically route traffic from subnets to the carrier gateway, we create the following resources:

- A carrier gateway
- A subnet. You can optionally assign all carrier gateway tags that do not have a **Key** value of `Name` to the subnet.
- A network ACL with the following resources:
  - A subnet associated with the subnet in the Wavelength Zone
  - Default inbound and outbound rules for all of your traffic.
- A route table with the following resources:
  - A route for all local traffic
  - A route that routes all non-local traffic to the carrier gateway
  - An associated with the subnet

To create a carrier gateway

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **Carrier Gateways**, and then choose **Create carrier gateway**.
3. Optional: For **Name**, enter a name for the carrier gateway.
4. For **VPC**, choose the VPC.
5. Choose **Route subnet traffic to carrier gateway**, and under **Subnets to route** do the following.
   a. Under **Existing subnets in Wavelength Zone**, select the box for each Wavelength subnet to route to the carrier gateway.
   b. To create a subnet in the Wavelength Zone, choose **Add new subnet**, specify the following information, and then choose **Add new subnet**:
      - **Name tag**: Optionally provide a name for your subnet. Doing so creates a tag with a key of `Name` and the value that you specify.
      - **VPC**: Choose the VPC.
      - **Availability Zone**: Choose the Wavelength Zone.
      - **IPv4 CIDR block**: Specify an IPv4 CIDR block for your subnet, for example, `10.0.1.0/24`.
      - To apply the carrier gateway tags to the subnet, select **Apply same tags from this carrier gateway**.
6. (Optional) To add a tag to the carrier gateway, choose **Add tag**, and then do the following:
   - For **Key**, enter the key name.
   - For **Value**, enter the key value.
7. Choose **Create carrier gateway**.

Step 3: Create a public subnet in an Availability Zone in the Region

Create a subnet in an Availability Zone in the Region.

To add a subnet

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **Subnets**, **Create subnet**.
3. Specify the subnet details as necessary and choose **Create**.
   - **Name tag**: Optionally provide a name for your subnet. Doing so creates a tag with a key of `Name` and the value that you specify.

- **VPC**: Choose the VPC for which you're creating the subnet.
- **Availability Zone**: Select an Availability Zone that is in your Region, or select **No Preference** to have AWS choose one for you.
- **IPv4 CIDR block**: Specify a public IPv4 CIDR block for your subnet.

Step 3: Launch an instance in your Availability Zone public subnet

Launch an EC2 instance in the subnet that you created in the Availability Zone. You will use this instance to test the connectivity from the Region to the Wavelength Zone.

You can launch EC2 instances in the public subnet that you created. For information about how to launch an instance in the Amazon EC2 console, see one of the following guides:

- For Linux instances, see Launch an instance in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*.
- For Windows instances, see Launch an instance in the *Amazon Elastic Compute Cloud User Guide for Windows Instances*.

Step 4: Launch an instance for your Wavelength application

After you complete the networking configuration, launch an instance, and then allocate a Carrier IP address for the instance.

Option 1: Launch an instance in the Wavelength Zone subnet and auto assign the Carrier IP address using the AWS CLI

AWS recommends that you use the `run-instances` command, because you can use an option to automatically allocate and associate the Carrier IP address with the network interface.

To launch an instance in a Wavelength Zone

- Use the `run-instances` command to launch an instance in the Wavelength Zone subnet. Use the following options:
  - `subnet:` Set the value to the ID of the subnet in the Wavelength Zone.
  - `AssociateCarrierIpAddress`: Set this value to `true`. This option assigns a Carrier IP address to the network interface for eth0.

  For more information about launching an instance, see run-instances in the AWS CLI Command Reference.

  Example

```
aws ec2 --region us-east-1 run-instances --network-interfaces '[{"DeviceIndex":0,
 "AssociateCarrierIpAddress": true, "SubnetId": "subnet-036aa298f4EXAMPLE"}]' --image-
id ami-04125ecea1EXAMPLE --instance-type t3.medium
```

Option 2: Launch an instance in the Wavelength Zone subnet and allocate and associate a Carrier IP address from the network border group

You can launch EC2 instances in the subnet that you created when you added the carrier gateway. For more information, see the section called "Step 2: Create a carrier gateway and a subnet associated with the Wavelength Zone" (p. 10). Security groups control inbound and outbound traffic for instances in a subnet, just as they do for instances in an Availability Zone subnet. To connect to an EC2 instance in a subnet, specify a key pair when you launch the instance, just as you do for instances in an Availability Zone subnet. For information about how to launch an instance in the Amazon EC2 console, see one of the following guides:

- For Linux instances, see Launch an instance in the *Amazon Elastic Compute Cloud User Guide for Linux Instances*.
- For Windows instances, see Launch an instance in the *Amazon Elastic Compute Cloud User Guide for Windows Instances*.

To allocate and associate a Carrier IP address

1. Use `allocate-address` to allocate a Carrier IP address. For more information, see allocate-address in the *AWS CLI Command Reference*.

   Example

   ```
   aws ec2 allocate-address --region us-east-1 --domain vpc --network-border-group us-
   east-1-wl1-bos-wlz-1
   ```

   Output

   ```
   {
       "AllocationId": "eipalloc-05807b62acEXAMPLE",
       "PublicIpv4Pool": "amazon",
       "NetworkBorderGroup": "us-east-1-wl1-bos-wlz-1",
       "Domain": "vpc",
       "CarrierIp": "155.146.10.111"

   }
   ```

2. Use `associate-address` to associate the Carrier IP address with the EC2 instance. For more information, see associate-address in the *AWS CLI Command Reference*.

   Example

   ```
   aws ec2 associate-address --allocation-id eipalloc-05807b62acEXAMPLE --network-
   interface-id eni-1a2b3c4d
   ```

   Output

```
{
    "AssociationId": "eipassoc-02463d08ceEXAMPLE",
}
```

Step 5: Test the connectivity

Before you test the connectivity, review the section called “Networking considerations” (p. 19).

Test the connectivity from the instance in the Region to the Wavelength Zone instance. Depending on your operating system, use **ssh** or **rdp** to connect to the Carrier IP address of your Region instance. You can use a secure bastion host.

Before you run these tests, configure the VPC security group to allow ICMP traffic. For more information, see Creating a security group in the *Amazon VPC User Guide*.

Run the ping command to the Wavelength Zone instance. In the following example, the IP address of the subnet in the Wavelength Zone is 10.0.3.112.

```
ping 10.0.3.112
Pinging 10.0.3.112
Reply from 10.0.3.112:  bytes=32 time=<1ms TTL=128
Reply from 10.0.3.112:  bytes=32 time=<1ms TTL=128
Reply from 10.0.3.112:  bytes=32 time=<1ms TTL=128

Ping statistics for 10.0.3.112
Packets:  Sent = 3,  Received = 3,  Lost = 0 (0% lost)

Approximate round trip time in milliseconds
Minimum = 0ms,  Maximum = 0ms,  Average = 0ms
```

Test the connectivity from the instance in the Wavelength Zone instance to the carrier network. Depending on your operating system, use **ssh** or **rdp** to connect to the Carrier IP address of your Wavelength Zone instance. You can use a secure bastion host.

You need a device on the carrier network in order to test the connectivity from the Wavelength Zone to the carrier network. In addition, Headspin, which is part of the AWS Partner Network, provides devices on carrier networks for functional testing. For more information, see Headspin.

Run the ping command to an address in the carrier network. In the following example, the carrier network IP address is 198.51.100.130.

```
ping 198.51.100.130
Pinging 198.51.100.130
Reply from 198.51.100.130:  bytes=32 time=<1ms TTL=128
Reply from 198.51.100.130:  bytes=32 time=<1ms TTL=128
Reply from 198.51.100.130:  bytes=32 time=<1ms TTL=128

Ping statistics for 198.51.100.130
Packets:  Sent = 3,  Received = 3,  Lost = 0 (0% lost)

Approximate round trip time in milliseconds
Minimum = 0ms,  Maximum = 0ms,  Average = 0ms
```

Security in AWS Wavelength

Security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The shared responsibility model describes this as security of the cloud and security in the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the AWS Compliance Programs. To learn about the compliance programs that apply to AWS Wavelength, see the section called "Compliance validation" (p. 16).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation describes the differences when services run in a Wavelength Zone. For detailed information about service security, see the following:

- Security in Amazon EC2
- Security in Amazon EC2 Auto Scaling
- Security in Amazon ECS clusters
- Security in Amazon EKS

Contents

Resilience in AWS Wavelength

AWS recommends that you architect edge applications in a hub and spoke model with the Region providing the most scalable, resilient, and cost effective options for components that are less latency sensitive, that need to be shared across Zones, or that have states that need to persist. Then, use Wavelength Zones for the application components that need ultra-low latency, higher bandwidth, or increased quality of service over 5G mobile networks.

If you need to replicate your data or applications in a Wavelength Zone, AWS recommends that you use an Availability Zone in the Region that is not the parent zone as the failover zone. In the following example, you can use the `describe-availability-zones` command to retrieve information about Wavelength Zone Z. The parent Availability Zone is Availability Zone A. In this case, you replicate your resources in Availability Zone B.

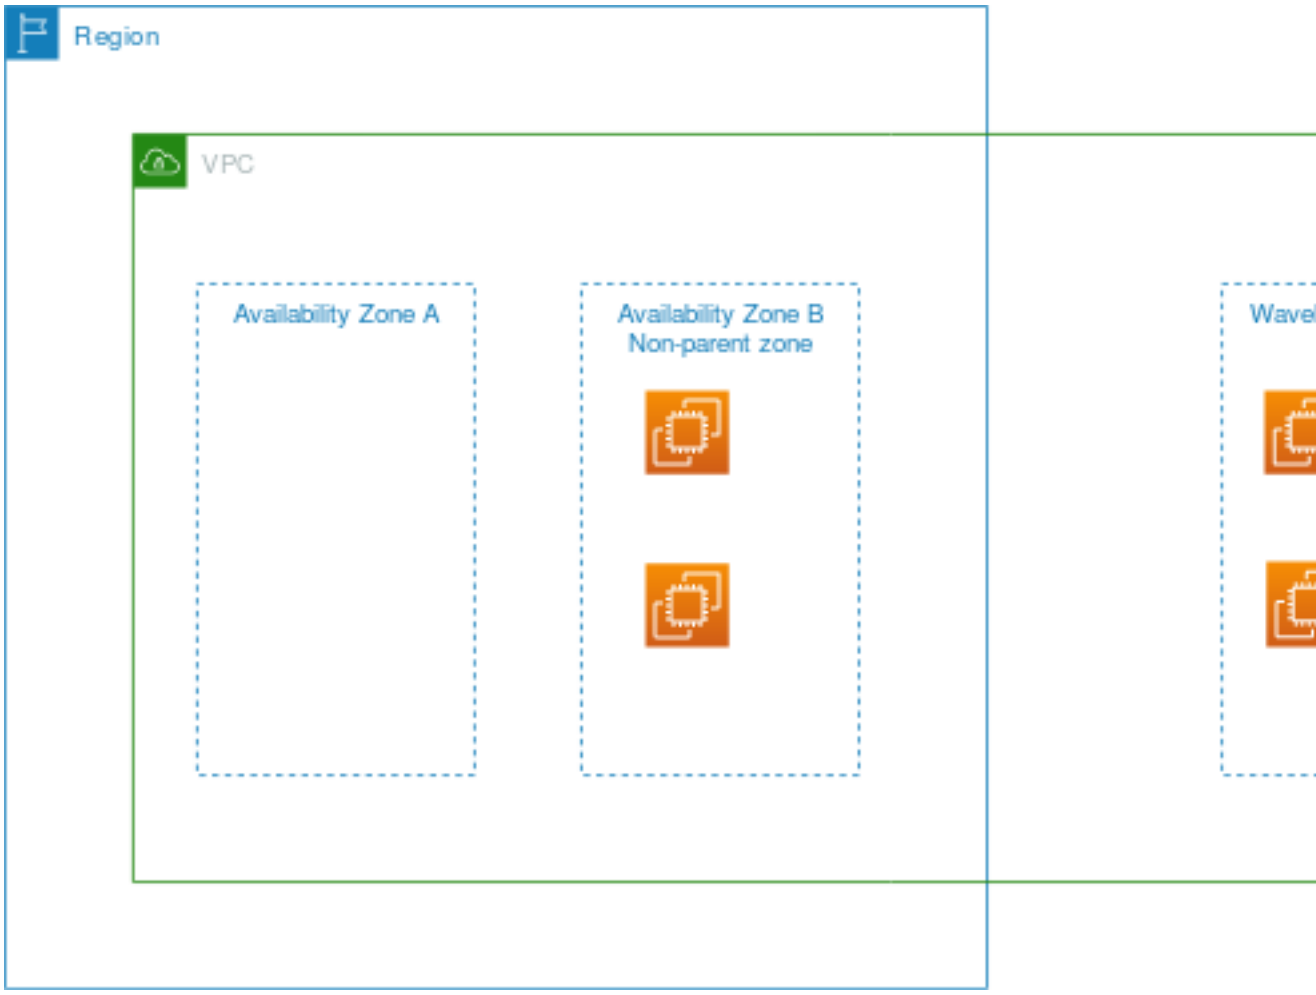

To learn more about resiliency in Amazon EC2 and Amazon EC2 Auto Scaling, see the following:

- Resilience in Amazon EC2 in the *Amazon EC2 User Guide for Linux Instances*
- Resilience in Amazon EC2 Auto Scaling in the *Amazon EC2 Auto Scaling User Guide*.

For more information about AWS Regions, Availability Zones, Local Zones, and Wavelength Zones, see AWS Global Infrastructure.

Compliance validation for AWS Wavelength

The existing compliance certifications for AWS Services apply to services running entirely in an AWS Region. The services running in a Wavelength Zone require a separate evaluation for certifications.

Under the shared responsibility model, AWS is responsible for the hardware and software that run AWS services. This applies to AWS Wavelength, just as it does to an AWS Region. This includes patching the infrastructure software and configuring infrastructure devices. As a customer, you are responsible for implementing best practices for data encryption, patching the operating system and applications, identity and access management, and operating system, network, and firewall configurations.

AWS has responsibility for configuring and maintaining a network connection between the Wavelength Zone and the AWS Region. Communication sent over this connection between the Wavelength Zone and the Region is encrypted by AWS.

Third-party auditors assess the security and compliance of services AWS Wavelength as part of multiple AWS compliance programs.

The certification currently supported is HIPAA.

For more information about your compliance responsibility when using Amazon EC2, see Compliance validation for Amazon EC2 in the *Amazon EC2 User Guide for Linux Instances*.

Wavelength considerations and quotas

Available Wavelength Zones

The following table lists the Wavelength Zones provided by an AWS account.

AWS Region	Location	Carrier	Wavelength Zone ID	Network Border Group
US East (N. Virginia)	Atlanta	Verizon	`us-east-1-wl1-atl-wlz-1`	`us-east-1-wl1-atl-wlz-1`
US East (N. Virginia)	Boston	Verizon	`us-east-1-wl1-bos-wlz-1`	`us-east-1-wl1-bos-wlz-1`
US East (N. Virginia)	Dallas	Verizon	`us-east-1-wl1-dfw-wlz-1`	`us-east-1-wl1-dfw-wlz-1`
US East (N. Virginia)	Miami	Verizon	`us-east-1-wl1-mia-wlz-1`	`us-east-1-wl1-mia-wlz-1`
US East (N. Virginia)	New York City	Verizon	`us-east-1-wl1-nyc-wlz-1`	`us-east-1-wl1-nyc-wlz-1`
US East (N. Virginia)	Washington DC	Verizon	`us-east-1-wl1-was-wlz-1`	`us-east-1-wl1-was-wlz-1`
US West (Oregon)	Denver	Verizon	`us-west-2-wl1-den-wlz-1`	`us-west-2-wl1-den-wlz-1`
US West (Oregon)	Las Vegas	Verizon	`us-west-2-wl1-las-wlz-1`	`us-west-2-wl1-las-wlz-1`
US West (Oregon)	San Francisco Bay area	Verizon	`us-west-2-wl1-sfo-wlz-1`	`us-west-2-wl1-sfo-wlz-1`
US West (Oregon)	Seattle	Verizon	`us-west-2-wl1-sea-wlz-1`	`us-west-2-wl1-sea-wlz-1`
Asia Pacific (Seoul)	Daejeon	SKT	`ap-northeast-2-wl1-cjj-wlz-1`	`ap-northeast-2-wl1-cjj-wlz-1`
Asia Pacific (Tokyo)	Tokyo	KDDI	`ap-northeast-1-wl1-nrt-wlz-1`	`ap-northeast-1-wl1-nrt-wlz-1`

For more information, see AWS Global Infrastructure.

The number and mapping of Wavelength Zones per Region might vary between AWS accounts. To get a list of the Wavelength Zones that are available to your account, you can use the Amazon EC2 console or the command line interface.

Networking considerations

The following controls are enabled by the carrier gateway for Internet flows by default and cannot be removed:

Protocol	Between EC2 instance and the internet	Between EC2 instance and a device on the carrier network
TCP	outbound and the response	allowed
UDP	denied	allowed
ICMP	allowed	allowed

- TCP is allowed for outbound and response
- UDP from the internet is denied

  UDP traffic from a device on the carrier network is allowed to route to an EC2 instance in a Wavelength Zone.
- ICMP is allowed

In addition, inbound routing from the carrier network is optimized for devices in the location of the Wavelength Zone. For example, a Wavelength Zone in the San Francisco Bay area allows low latency access only from devices that are in that metro area and carrier network.

Amazon EC2 considerations

Take the following information into consideration when you launch EC2 instances in Wavelength Zones:

- The following instance types are supported:
  - t3.medium
  - t3.xlarge
  - r5.2xlarge
  - g4dn.2xlarge
- You cannot use dedicated instances or hosts.

You can find the EC2 instance types available by using the `describe-instance-type-offerings` command.

```
aws ec2 describe-instance-type-offerings --location-type availability-zone --filters
 Name=location,Values=us-east-1-wl1-bos-wlz-1 --region us-east-1
```

Output

```
{
```

```
    "InstanceTypeOfferings": [
        {
            "InstanceType": "t3.xlarge",
            "LocationType": "availability-zone",
            "Location": "us-east-1-wl1-bos-wlz-1"
        },
        {
            "InstanceType": "t3.medium",
            "LocationType": "availability-zone",
            "Location": "us-east-1-wl1-bos-wlz-1"
        },
        {
            "InstanceType": "r5.2xlarge",
            "LocationType": "availability-zone",
            "Location": "us-east-1-wl1-bos-wlz-1"
        },
        {
            "InstanceType": "g4dn.2xlarge",
            "LocationType": "availability-zone",
            "Location": "us-east-1-wl1-bos-wlz-1"
        }
    ]
}
```

Amazon EBS considerations

Take the following information into consideration when you use Amazon Elastic Block Store for EC2 instances that are in Wavelength Zones:

- Snapshots of EBS volumes and AMIs are stored in the AWS Region.
- You can only use `gp2` volumes.
- The default limit for `gp2` storage is 30 TB.

  You can request an increase for this value.

Amazon Elastic Kubernetes Service considerations

Take the following information into consideration when you run an Amazon EKS cluster:

- You must run Kubernetes 1.17 or later.
- When you create your Amazon EKS cluster, you must select an Availability Zone in the VPC, and not a Wavelength Zone.
- When you create your Amazon EKS cluster for private subnets only, you need to add VPC endpoints for Amazon ECR and Amazon Simple Storage Service. For more information, see the section called "Amazon VPC considerations" (p. 20).
- To create node groups in Wavelength Zones for your Amazon EKS cluster, see Launching self-managed Amazon Linux 2 nodes in the *Amazon EKS User Guide*.
- To apply the `aws-auth` ConfigMap to your Amazon EKS cluster, see Managing users or IAM roles for your cluster in the *Amazon EKS User Guide*.

Amazon VPC considerations

Take the following information into consideration when you run Amazon VPC:

- If you want to use VPC endpoints, you must create the endpoint in an Availability Zone in the VPC. You cannot create the endpoint in a Wavelength Zone.
- You cannot assign an IPV6 addresses to subnets that are in Wavelength Zones.

Multiple Wavelength Zone considerations

Note
EC2 instances that are in two different Wavelength Zones in the same VPC are not allowed to communicate with each other. If you need Wavelength Zone to Wavelength Zone communication, AWS recommends that you use multiple VPCs, one for each Wavelength Zone. You can use a transit gateway to connect the VPCs. This configuration enables communication between instances in the Wavelength Zones. For information about how to create the configuration for multiple Wavelength Zones, see Extending your VPC resources to Wavelength Zones in the *Amazon VPC User Guide*.

Quotas

Wavelength VPCs and Wavelength subnets count toward the Amazon VPC quotas. For more information about Amazon VPC quotas, see Amazon VPC Quotas in the *Amazon VPC User Guide*.

You can use the Service Quotas console (https://console.aws.amazon.com/servicequotas/) to view your EC2, EBS, and VPC quotas. For more information about how to view your Service Quotas, see Viewing a Service Quota in the *Service Quotas User Guide*.

Document history for AWS Wavelength Developer Guide

The following table describes the documentation for this release of AWS Wavelength. We also update the documentation frequently to address the feedback that you send us.

- **API version: latest** 2016-11-15
- **Latest documentation update:** September 22, 2020

Change	Description	Release Date
Additional Wavelength Zones	This release introduces new Wavelength Zones. For more information, see the section called "Available Wavelength Zones" (p. 18).	September 22, 2020
Initial release	This release introduces AWS Wavelength.	August 6, 2020