

## **Module 4: ELB AND ROUTE 53**

**Assignment Submitted By:-Hitesh Chauhan**

**Course Offered: Advanced Cloud Computing And Devops**

## Problem Statement:

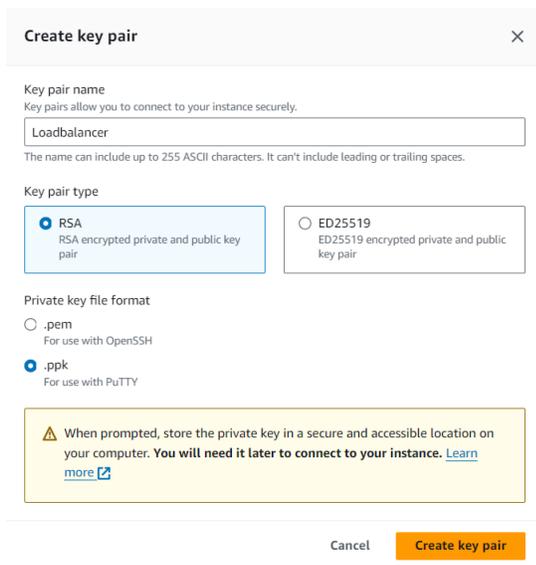
You work for XYZ Corporation which uses on-premise solutions and a limited number of systems. With the increase in requests in their application, the load also increases. So, to handle the load the corporation must buy more systems almost on a regular basis. Realizing the need to cut down the expenses on systems, they decided to move their infrastructure to AWS.

## Tasks To Be Performed:

1. Manage the scaling requirements of the company by:
  - a. Deploying multiple compute resources on the cloud as soon as the load increases and the CPU utilization exceeds 80%
  - b. Removing the resources when the CPU utilization goes under 60%
2. Create a load balancer to distribute the load between compute resources.
3. Route the traffic to the company's domain

## Answer

Login to AWS Console and go to the EC2 dashboard.



The screenshot shows the 'Create key pair' dialog box in the AWS Management Console. It includes a text input field for the 'Key pair name' containing 'Loadbalancer'. Below this, there are two radio button options for 'Key pair type': 'RSA' (selected) and 'ED25519'. Underneath, there are two radio button options for 'Private key file format': '.pem' and '.ppk' (selected). A yellow warning box at the bottom states: 'When prompted, store the private key in a secure and accessible location on your computer. You will need it later to connect to your instance. Learn more'. At the bottom right, there are 'Cancel' and 'Create key pair' buttons.

Create a Key Pair for accessing the EC2 instances.

## Create Security Groups

Go to the Security Groups  Create Security Group (cloudwatch agent) and allow port 80 and 22.

## Compare security group rules Info

Amazon EC2 evaluates all the rules of the selected security groups to control inbound and outbound traffic. You can select more security groups to view their inbound rules to help you to decide how to secure your instance from incoming traffic.

Common security groups

Select security groups

default sg-0087d67d5ce0fb457 X  
VPC: vpc-0ca4bd13727b758cb

Security groups that you add or remove here will be added to or removed from all your network interfaces.

Inbound rules (2)

Security group name	Security group ID	Type	Protocol	Port range	Source	Description
default	sg-0087d67d5ce0fb457	HTTP	tcp	80	0.0.0.0/0	-
default	sg-0087d67d5ce0fb457	ssh	tcp	22	0.0.0.0/0	-

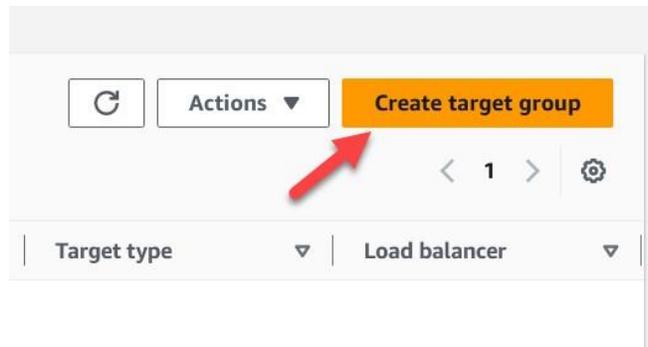
Now Security Group is created successfully.

EC2 > Instances > Launch an instance

Success

Successfully initiated launch of instances ([i-0cb9644581f5ad938](#), [i-0e429874168d1ac12](#))

EC2 Dashboard > Load Balancing > Target Groups > Create Target Group



Choose target Type

## Basic configuration

Settings in this section can't be changed after the target group is created.

### Choose a target type

Instances

- Supports load balancing to instances within a specific VPC.
- Facilitates the use of [Amazon EC2 Auto Scaling](#) to manage and scale your EC2 capacity.

IP addresses

- Supports load balancing to VPC and on-premises resources.
- Facilitates routing to multiple IP addresses and network interfaces on the same instance.
- Offers flexibility with microservice based architectures, simplifying inter-application communication.
- Supports IPv6 targets, enabling end-to-end IPv6 communication, and IPv4-to-IPv6 NAT.

Lambda function

- Facilitates routing to a single Lambda function.
- Accessible to Application Load Balancers only.

Application Load Balancer

- Offers the flexibility for a Network Load Balancer to accept and route TCP requests within a specific VPC.
- Facilitates using static IP addresses and PrivateLink with an Application Load Balancer.

### Target group name

A maximum of 32 alphanumeric characters including hyphens are allowed, but the name must not begin or end with a hyphen.

### Protocol : Port

Choose a protocol for your target group that corresponds to the Load Balancer type that will route traffic to it. Some protocols now include anomaly detection for the targets and you can set mitigation options once your target group is created. This choice cannot be changed after creation

1-65535

### IP address type

Only targets with the indicated IP address type can be registered to this target group.

IPv4

Each instance has a default network interface (eth0) that is assigned the primary private IPv4 address. The instance's primary private IPv4 address is the one that will be applied to the target.

IPv6

Each instance you register must have an assigned primary IPv6 address. This is configured on the instance's default network interface (eth0). [Learn more](#)

### VPC

Select the VPC with the instances that you want to include in the target group. Only VPCs that support the IP address type selected above are available in this list.

  
IPv4 VPC CIDR: 172.31.0.0/16

### Protocol version

HTTP1

Send requests to targets using HTTP/1.1. Supported when the request protocol is HTTP/1.1 or HTTP/2.

HTTP2

Send requests to targets using HTTP/2. Supported when the request protocol is HTTP/2 or gRPC, but gRPC-specific features are not available.

gRPC

Send requests to targets using gRPC. Supported when the request protocol is gRPC.

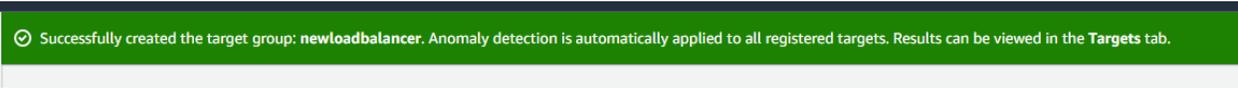
Click Next

## Register targets

This is an optional step to create a target group. However, to ensure that your load balancer routes traffic to this target group you must register your targets.

Available instances (2/2)							
Instance ID	Name	State	Security groups	Zone	Private IPv4 address		
<input checked="" type="checkbox"/>	i-0e429874168d1ac12	my loadbalancer	Running	cloudwatch agent	us-east-1a	172.31.82.143	
<input checked="" type="checkbox"/>	i-0cb9644581f5ad938	my loadbalancer	Running	cloudwatch agent	us-east-1a	172.31.89.127	

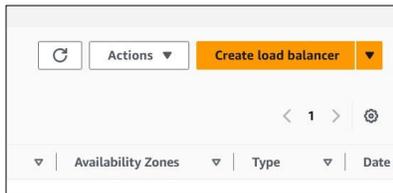
Click Create Target Group



Now Target group has been created successfully

## Load Balancer

EC2 Dashboard > Load Balancing > Load Balancer > Create Load Balancer



### Basic configuration

**Load balancer name**  
Name must be unique within your AWS account and can't be changed after the load balancer is created.  
  
A maximum of 32 alphanumeric characters including hyphens are allowed, but the name must not begin or end with a hyphen.

**Scheme** [Info](#)  
Scheme can't be changed after the load balancer is created.

- Internet-facing**  
An internet-facing load balancer routes requests from clients over the internet to targets. Requires a public subnet. [Learn more](#)
- Internal**  
An internal load balancer routes requests from clients to targets using private IP addresses.

**IP address type** [Info](#)  
Select the type of IP addresses that your subnets use.

- IPv4**  
Recommended for internal load balancers.
- Dualstack**  
Includes IPv4 and IPv6 addresses.

## Network mapping [Info](#)

The load balancer routes traffic to targets in the selected subnets, and in accordance with your IP address settings.

### VPC [Info](#)

The load balancer will exist and scale within the selected VPC. The selected VPC is also where the load balancer targets must be hosted unless routing to Lambda or on-premis using VPC peering. To confirm the VPC for your targets, view [target groups](#). For a new VPC, [create a VPC](#).

-  
vpc-0ca4bd13727b758cb  
IPv4 VPC CIDR: 172.31.0.0/16



### Mappings [Info](#)

Select at least two Availability Zones and one subnet per zone. The load balancer routes traffic to targets in these Availability Zones only. Availability Zones that are not supported by the load balancer or the VPC are not available for selection.

#### Availability Zones

- us-east-1a (use1-az2)
- us-east-1b (use1-az4)
- us-east-1c (use1-az6)
- us-east-1d (use1-az1)
- us-east-1e (use1-az3)
- us-east-1f (use1-az5)

## Security groups [Info](#)

A security group is a set of firewall rules that control the traffic to your load balancer. Select an existing security group, or you can [create a new security group](#).

#### Security groups

Select up to 5 security groups



cloudwatch agent   
sg-0513f7b0108994883 VPC: vpc-0ca4bd13727b758cb

## Listeners and routing [Info](#)

A listener is a process that checks for connection requests using the port and protocol you configure. The rules that you define for a listener determine how the load balancer routes requests to its registered targets.

▼ Listener HTTP:80 Remove

Protocol	Port	Default action	<a href="#">Info</a>
HTTP	80 1-65535	Forward to	newloadbalancer Target type: Instance, IPv4

[Create target group](#)

Listener tags - optional  
Consider adding tags to your listener. Tags enable you to categorize your AWS resources so you can more easily manage them.

Add listener tag  
You can add up to 50 more tags.

Add listener

Select the Target group which we have created

Add-on services [Edit](#) Tags [Edit](#)  
None None

Attributes

ⓘ Certain default attributes will be applied to your load balancer. You can view and edit them after creating the load balancer.

Cancel [Create load balancer](#)

## Create Load Balancer

**Summary**  
Review and confirm your configurations. [Estimate cost](#)

<b>Basic configuration</b> <a href="#">Edit</a> myalb <ul style="list-style-type: none"><li>Internet-facing</li><li>IPv4</li></ul>	<b>Security groups</b> <a href="#">Edit</a> <ul style="list-style-type: none"><li>cloudwatch agent <a href="#">sg-0513f7b0108994883</a></li></ul>	<b>Network mapping</b> <a href="#">Edit</a> VPC <a href="#">vpc-0ca4bd13727b758cb</a> <ul style="list-style-type: none"><li>us-east-1a <a href="#">subnet-0da62f71a2c23d8a8</a></li><li>us-east-1b <a href="#">subnet-07500936c65530e04</a></li><li>us-east-1c <a href="#">subnet-09c8bc6fedafb2bce</a></li><li>us-east-1d <a href="#">subnet-0d828a638aec31575</a></li><li>us-east-1e <a href="#">subnet-084eae892db4de22c</a></li><li>us-east-1f <a href="#">subnet-0c3a227a6d896e22e</a></li></ul>	<b>Listeners and routing</b> <a href="#">Edit</a> <ul style="list-style-type: none"><li>HTTP:80 defaults to <a href="#">newloadbalancer</a></li></ul>
<b>Service integrations</b> <a href="#">Edit</a> AWS WAF: None AWS Global Accelerator: None	<b>Tags</b> <a href="#">Edit</a> None		

**Creation workflow and status**

► **Server-side tasks and status**  
After completing and submitting the above steps, all server-side tasks and their statuses become available for monitoring.

Cancel [Create load balancer](#)

**Load balancers (1/1)** Actions Create load balancer

Elastic Load Balancing scales your load balancer capacity automatically in response to changes in incoming traffic.

Filter load balancers

Name	DNS name	State	VPC ID	Availability Zones	Type	Date created
myalb	myalb-1562156745.us-eas...	Provisioning	vpc-0ca4bd13727b75...	6 Availability Zones	application	September 29, 2024, 13:05 (UTC+05:30)

**Load balancer: myalb**

Details | Listeners and rules | Network mapping | Resource map - new | Security | Monitoring | Integrations | Attributes | Tags

**Details**

Load balancer type Application	Status Provisioning	VPC vpc-0ca4bd13727b758cb	Load balancer IP address type IPv4
Scheme Internet-facing	Hosted zone Z35SXDOTRQ7X7K	Availability Zones <a href="#">subnet-0d828a638aec31575</a> us-east-1d (use1-az1) <a href="#">subnet-09c8bc6fedafb2bce</a> us-east-1c (use1-az6) <a href="#">subnet-0c3a227a6d896e22e</a> us-east-1f (use1-az5)	Date created September 29, 2024, 13:05 (UTC+05:30)

ALB creation will take a few minutes

## Create Launch Template

EC2 Dashboard **Instances** **Launch Templates**

**New launch template**

[Create launch template](#)

**Launch template name and description**

Launch template name - *required*

myalunchtemplate

Must be unique to this account. Max 128 chars. No spaces or special characters like '&', '\*', '@'.

Template version description

v1

Max 255 chars

Auto Scaling guidance [Info](#)

Select this if you intend to use this template with EC2 Auto Scaling

Provide guidance to help me set up a template that I can use with EC2 Auto Scaling

▶ **Template tags**

▶ **Source template**

- Select OS as Amazon Linux
- Instance Type is t2.micro (Free Tier Eligible)
- Select the Key Pair
- Select the Security Group as WEBSG

In Advanced options provide the user data script to make your website ready

✔ **Success**

Successfully created [mylaunchtemplate1 \(lt-069c6cb563ce83e2d\)](#).

## Create Auto Scaling Groups

EC2 Dashboard 7 Auto Scaling 7 Create Auto Scaling Groups

### Create Auto Scaling group

Get started with EC2 Auto Scaling by creating an Auto Scaling group.

[Create Auto Scaling group](#)

Provide the Auto Scaling Group name

### Choose launch template [Info](#)

Specify a launch template that contains settings common to all EC2 instances that are launched by this Auto

#### Name

Auto Scaling group name

Enter a name to identify the group.

Must be unique to this account in the current Region and no more than 255 characters.

## Launch template [Info](#)

**i** For accounts created after May 31, 2023, the EC2 console only supports creating Auto Scaling groups with launch templates. Creating Auto Scaling groups with launch configurations is not recommended but still available via the CLI and API until December 31, 2023.

### Launch template

Choose a launch template that contains the instance-level settings, such as the Amazon Machine Image (AMI), instance type, key pair, and security groups.

mylaunchtemplate1

[Create a launch template](#)

### Version

Default (1)

[Create a launch template version](#)

### Description

v1

### Launch template

mylaunchtemplate1 [↗](#)  
lt-069c6cb563ce83e2d

### Instance type

t2.micro

### AMI ID

ami-0ebfd941bbafe70c6

### Security groups

-

### Request Spot Instances

No

### Key pair name

Loadbalancer

### Security group IDs

sg-0513f7b0108994883 [↗](#)

### Additional details

#### Storage (volumes)

-

#### Date created

Sun Sep 29 2024 13:13:21  
GMT+0530 (India Standard Time)

Click Next.

## Instance type requirements [Info](#)

[Override launch template](#)

You can keep the same instance attributes or instance type from your launch template, or you can choose to override the launch template by specifying different instance attributes or manually adding instance types.

Launch template <a href="#">mylaunchtemplate1</a> lt-069c6cb563ce83e2d	Version Default	Description v1
Instance type t2.micro		

## Network [Info](#)

For most applications, you can use multiple Availability Zones and let EC2 Auto Scaling balance your instances across the zones. The default VPC and default subnets are suitable for getting started quickly.

### VPC

Choose the VPC that defines the virtual network for your Auto Scaling group.

172.31.0.0/16 Default

[Create a VPC](#)

### Availability Zones and subnets

Define which Availability Zones and subnets your Auto Scaling group can use in the chosen VPC.

[Create a subnet](#)

Cancel

Skip to review

Previous

Next

Select the VPC Settings and Subnets

### Network Info

For most applications, you can use multiple Availability Zones and let EC2 Auto Scaling balance your instances across the zones. The default VPC and default subnets are suitable for getting started quickly.

**VPC**  
Choose the VPC that defines the virtual network for your Auto Scaling group.

vpc-0ca4bd13727b758cb  
172.31.0.0/16 Default

[Create a VPC](#)

**Availability Zones and subnets**  
Define which Availability Zones and subnets your Auto Scaling group can use in the chosen VPC.

Select Availability Zones and subnets

- us-east-1a | subnet-0da62f71a2c23d8a8  
172.31.80.0/20 Default
- us-east-1b | subnet-07500936c65530e04  
172.31.16.0/20 Default
- us-east-1c | subnet-09c8bc6fedafb2bce  
172.31.32.0/20 Default
- us-east-1d | subnet-0d828a638aec31575  
172.31.0.0/20 Default
- us-east-1e | subnet-084eae892db4de22c  
172.31.48.0/20 Default
- us-east-1f | subnet-0c3a227a6d896e22e  
172.31.64.0/20 Default

[Create a subnet](#)

Click NEXT

### Configure advanced options - *optional* Info

Integrate your Auto Scaling group with other services to distribute network traffic across multiple servers using a load balancer or to establish service-to-service communications using VPC Lattice. You can also set options that give you more control over health check replacements and monitoring.

**Load balancing Info**

Use the options below to attach your Auto Scaling group to an existing load balancer, or to a new load balancer that you define.

- No load balancer  
Traffic to your Auto Scaling group will not be fronted by a load balancer.
- Attach to an existing load balancer  
Choose from your existing load balancers.
- Attach to a new load balancer  
Quickly create a basic load balancer to attach to your Auto Scaling group.

**Attach to an existing load balancer**  
Select the load balancers that you want to attach to your Auto Scaling group.

- Choose from your load balancer target groups  
This option allows you to attach Application, Network, or Gateway Load Balancers.
- Choose from Classic Load Balancers

**Existing load balancer target groups**  
Only instance target groups that belong to the same VPC as your Auto Scaling group are available for selection.

Select target groups

- newloadbalancer | HTTP  
Application Load Balancer: myalb

Load Balancing options Attach to an existing load balancer.

Choose from your load balancer target groups (Since we already created one)

### Health checks

Health checks increase availability by replacing unhealthy instances. When you use multiple health checks, all are evaluated, and if at least one fails, instance replacement occurs.

EC2 health checks  
[Always enabled](#)

Additional health check types - optional [Info](#)

Turn on Elastic Load Balancing health checks **Recommended**  
Elastic Load Balancing monitors whether instances are available to handle requests. When it reports an unhealthy instance, EC2 Auto Scaling can replace it on its next periodic check.

[EC2 Auto Scaling will start to detect and act on health checks performed by Elastic Load Balancing. To avoid unexpected terminations, first verify the settings of these health checks in the \[Load Balancer console\]\(#\)](#)

Health check grace period [Info](#)  
This time period delays the first health check until your instances finish initializing. It doesn't prevent an instance from terminating when placed into a non-running state.

10 seconds

### Additional settings

Monitoring [Info](#)

Enable group metrics collection within CloudWatch

Default instance warmup [Info](#)  
The amount of time that CloudWatch metrics for new instances do not contribute to the group's aggregated instance metrics, as their usage data is not reliable yet.

Enable default instance warmup

Cancel Skip to review Previous **Next**

Click **Next**

### Group size - optional [Info](#)

Specify the size of the Auto Scaling group by changing the desired capacity. You can also specify minimum and maximum capacity limits. Your desired capacity must be within the limit range.

Desired capacity  
1

Minimum capacity  
1

Maximum capacity  
2

### Scaling policies - optional

Choose whether to use a scaling policy to dynamically resize your Auto Scaling group to meet changes in demand. [Info](#)

**Target tracking scaling policy**

Choose a desired outcome and leave it to the scaling policy to add and remove capacity as needed to achieve that outcome.

None

Scaling policy name

Target Tracking Policy

Metric type [Info](#)

Monitored metric that determines if resource utilization is too low or high. If using EC2 metrics, consider enabling detailed monitoring for better scaling performance.

Average CPU utilization

Target value

80

Instance warmup [Info](#)

5 seconds

Disable scale in to create only a scale-out policy

Configure Scaling Policies as described in questions, when CPU 80% utilized it should create another EC2 instance

### Instance scale-in protection - optional

Instance scale-in protection

If protect from scale in is enabled, newly launched instances will be protected from scale in by default.

Enable instance scale-in protection

Cancel

Skip to review

Previous

Next

Click **Next**

### Add notifications - optional [Info](#)

Send notifications to SNS topics whenever Amazon EC2 Auto Scaling launches or terminates the EC2 instances in your Auto Scaling group.

Add notification

Cancel

Skip to review

Previous

Next

Click **Next**

Step 6: Add tags

Edit

Tags (0)

Key	Value	Tag new instances
No tags		

Cancel

Previous

Create Auto Scaling group

Create ASG

Access the Website using ALB URL.



The ASG is working fine as expected

Instances (4) info

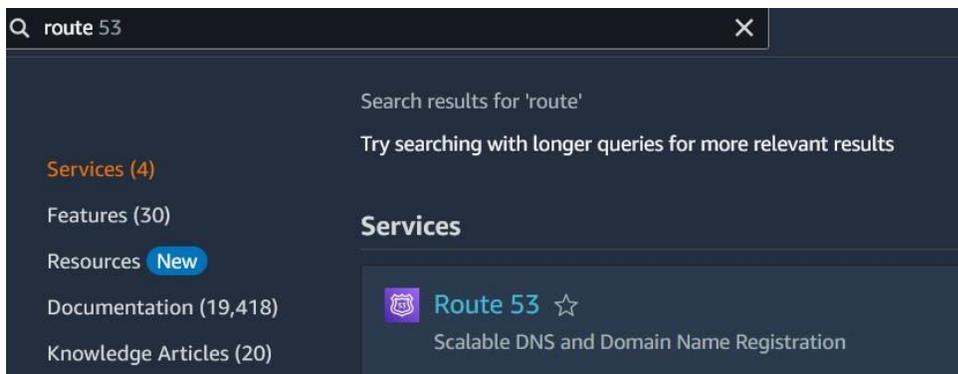
Last updated less than a minute ago

Connect Instance state Actions Launch instances

Find Instance by attribute or tag (case-sensitive) All states

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Public IPv4 ...	Elastic IP
	i-0a0054f9f211b0457	Stopped	t2.micro	-	View alarms +	us-east-1a	-	-	-
my loadbalancer	i-0e429874168d1ac12	Running	t2.micro	2/2 checks passec	View alarms +	us-east-1a	ec2-34-239-133-111.co...	34.239.133.111	-
my loadbalancer	i-0cb9644581f5ad938	Running	t2.micro	2/2 checks passec	View alarms +	us-east-1a	ec2-3-89-111-52.comp...	3.89.111.52	-
	i-0beb9efa4853dd46e	Running	t2.micro	2/2 checks passec	View alarms +	us-east-1a	ec2-3-82-245-32.comp...	3.82.245.32	-

## Map the Company Domain



Go to Route 53

Create a hosted zone ex: (hiteshchauhan.xyz)

# Create hosted zone [Info](#)

## Hosted zone configuration

A hosted zone is a container that holds information about how you want to route traffic for a domain, such as example.com, and its subdomains.

### Domain name [Info](#)

This is the name of the domain that you want to route traffic for.

Valid characters: a-z, 0-9, ! " # \$ % & ' ( ) \* + , - / : ; < = > ? @ [ \ ] ^ \_ ` { } . ~

### Description - optional [Info](#)

This value lets you distinguish hosted zones that have the same name.

The description can have up to 256 characters. 0/256

### Type [Info](#)

The type indicates whether you want to route traffic on the internet or in an Amazon VPC.

**Public hosted zone**  
A public hosted zone determines how traffic is routed on the internet.

**Private hosted zone**  
A private hosted zone determines how traffic is routed within an Amazon VPC.

[Route 53](#) > [Hosted zones](#) > [hiteshchauhan.xyz](#)

[Public](#) **hiteshchauhan.xyz** [Info](#)

Delete zone

Test record

Configure query logging

▶ Hosted zone details

Edit hosted zone

[Records \(2\)](#)

[DNSSEC signing](#)

[Hosted zone tags \(0\)](#)

**Records (2)** [Info](#)

Automatic mode is the current search behavior optimized for best filter results. To change modes go to settings.



Delete record

Import zone file

Create record

🔍 Filter records by property or value

Type ▼

Routing pol... ▼

Alias ▼

< 1 > ⌘

<input type="checkbox"/>	Record name	Type	Routin...	Differ...	Alias	Value/Route traffic to	TTL (s...)	Health ...	Evalu...	R..
<input type="checkbox"/>	hiteshchauhan.xyz	NS	Simple	-	No	ns-435.awsdns-54.com. ns-1013.awsdns-62.net. ns-1086.awsdns-07.org. ns-1981.awsdns-55.co.uk.	172800	-	-	-
<input type="checkbox"/>	hiteshchauhan.xyz	SOA	Simple	-	No	ns-435.awsdns-54.com. awsd...	900	-	-	-

# Create record [Info](#)

## Quick create record

[Switch to wizard](#)

### ▼ Record 1

Delete

Record name [Info](#)

myweb .hiteshchauhan.xyz

Keep blank to create a record for the root domain.

Record type [Info](#)

CNAME – Routes traffic to another domain name and to some AWS reso... ▼

Alias

Value [Info](#)

www.example.com

Enter multiple values on separate lines.

TTL (seconds) [Info](#)

300

1m

1h

1d

Routing policy [Info](#)

Simple routing ▼

Recommended values: 60 to 172800 (two days)

Add another record

Cancel

Create records

## Quick create record

[Switch to wizard](#)

### ▼ Record 1

Delete

Record name [Info](#)

myweb .hiteshchauhan.xyz

Keep blank to create a record for the root domain.

Record type [Info](#)

CNAME – Routes traffic to another domain name and to some AWS reso... ▼

Alias

Value [Info](#)

myalb-1562156745.us-east-1.elb.amazonaws.com

Enter multiple values on separate lines.

TTL (seconds) [Info](#)

300

1m

1h

1d

Routing policy [Info](#)

Simple routing ▼

Recommended values: 60 to 172800 (two days)

Add another record

Cancel

Create records

Hosted zone details

Edit hosted zone

Records (3) DNSSEC signing Hosted zone tags (0)

Records (3) Info

Delete record Import zone file Create record

Automatic mode is the current search behavior optimized for best filter results. To change modes go to settings.

Filter records by property or value Type Routing pol... Alias

< 1 > ⌘

<input type="checkbox"/>	Record name	Type	Routin...	Differ...	Alias	Value/Route traffic to	TTL (s...)	Health ...	Evalu...	R..
<input type="checkbox"/>	hiteshchauhan.xyz	NS	Simple	-	No	ns-435.awsdns-54.com. ns-1013.awsdns-62.net. ns-1086.awsdns-07.org. ns-1981.awsdns-55.co.uk.	172800	-	-	-
<input type="checkbox"/>	hiteshchauhan.xyz	SOA	Simple	-	No	ns-435.awsdns-54.com. awsd...	900	-	-	-
<input type="checkbox"/>	myweb.hiteshchauhan.xyz	CNAME	Simple	-	No	myalb-1562156745.us-east-...	300	-	-	-