# Using Load Balancers

## Lab overview

Follow these steps to create and configure a load balancer, register a webpage as a target for the load balancer, and test the load balancer.

## Duration

This lab requires approximately 30 minutes to complete.

## AWS service restrictions

In this lab environment, access to AWS services and service actions might be restricted to the ones that are needed to complete the lab instructions. You might encounter errors if you attempt to access other services or perform actions beyond the ones that are described in this lab.

## Access the AWS Management Console

1. To start the lab session, choose Start Lab in the upper-right corner of the page.
    - The lab session starts.
    - A timer displays in the upper-right corner of the page and shows the time remaining in the session.
      Tip: To refresh the session length at any time, choose Start Lab again before the timer reaches 0:00.
2. Before continuing, wait until the lab environment is ready. The environment is ready when the circle icon next to the AWS link in the upper-left corner turns green.
3. To return to these instructions, choose the Readme link in the upper-right corner.
4. To connect to the AWS Management Console, choose the AWS link in the upper-left corner, above the terminal window.
   A new browser tab opens and connects you to the AWS Management Console.
   Tip: If a new browser tab does not open, a banner or icon is usually at the top of your browser with the message that your browser is preventing the site from opening pop-up windows. Choose the banner or icon, and then choose Allow pop-ups.

Note: If you find a dialog prompting you to switch to the new console home, choose Switch to the new Console Home.

## Task 1. Launch an EC2 instance

In this task, you launch an Amazon Elastic Compute Cloud (Amazon EC2) instance as you have in previous labs.

3. In the search box to the right of Services, search for and choose EC2 to open the EC2 console.
4. In the left navigation pane, choose EC2 Dashboard to ensure you are on the dashboard page.
5. Choose the Launch instance button in the middle of the page, then select Launch instance from the dropdown menu.
6. In the Name and tags panel:
    - For Name enter Web Server 1
7. In the Application and OS Images panel:
    - For Quick Start, keep the default Amazon Linux chosen
8. In the Instance type panel:
    - Keep the default instance type, t2.micro.
9. In the Key pair (login) panel:
    - From the Key pair name - required dropdown list, choose vockey.
10. In the Network settings section, choose Edit.
    - From the Subnet dropdown list, choose the existing subnet in Availability Zone us-east-1a.
    - For Security group name - required, enter Web Server security group
    - For Description - required, enter Security group for my web server
    - In the Inbound security groups rules section, Select Remove to remove the default rule.
    - Choose Add security group rule to configure new rule as below
        - Type : HTTP
        - Source type : Anywhere

11. In the Configure storage panel:
    - Keep the default storage configuration.

12. Scroll down, and expand Advanced Details Panel, then configure:
    - Scroll down to the field User data.
    - Copy the following code and paste it into the User data field.
    - #!/bin/bash
      yum update -y
      yum -y install httpd
      systemctl enable httpd
      systemctl start httpd
      echo '<html><h1>Hello World! This is server 1.</h1></html>' >
      /var/www/html/index.html
    - This script does the following:
        - Updates the server
        - Installs an Apache web server (httpd)
        - Configures the web server to automatically start on boot
        - Starts the web server
        - Creates a simple webpage
    -

13. Choose Launch instance .

14. On the next screen, Choose View all instances.

15. Before you continue, wait for your instance to display the following:
    - Instance state: Running
    - Status check: 2/2 checks passed
      Tip: To refresh the instance information,  choose the refresh icon.

## Task 2. Access your EC2 instance’s website

In this task, you will access the content from the web server on the EC2 instance that you just created.

18. Select the Web Server 1 instance that you created earlier in this lab.

19. In the Details tab copy the Public IPv4 address of your instance, then open a new tab in your web browser and paste in and load the address.
    It should display the web server page with the message *Hello World! This is server 1.**
    Note: If web page is not displayed, ensure that you are accessing page using http:// (not https://).

## Task 3. Create a second EC2 instance for load balancing

In this task, you will create a second EC2 instance so that you will later be able to configure load balancing between the two instances.

20. Return to the EC2 Management Console browser tab.

21. Select the Web Server 1 instance.

22. From the Actions menu, choose Images and templates, then choose Launch more like this
    A Launch an instance page opens.

23. In the Name and tags pane, change the Name to Web Server 2.

24. In the Key pair (login) section, from the Key pair name - required dropdown list, choose vockey.

25. From the Subnet dropdown list, choose the existing subnet in Availability Zone us-east-1b.

26. Scroll down, and expand the Advanced Details section, then scroll down to the User data field.

27. Use copy and paste to replace the existing code with the code shown below.

28. #!/bin/bash
    yum update -y
    yum -y install httpd
    systemctl enable httpd
    systemctl start httpd
    echo '<html><h1>Hello World! This is server 2.</h1></html>' > /var/www/html/index.html

29. Note: This script is almost the same as the one that you used for the first instance. However, notice that it says This is server 2. The text that displays when you access Web Server 2 will be different than the text of Web Server 1. When you access the instances through the load balancer, this difference in text is how you will know which instance is displayed.

28. Choose Launch instance .

29. On the next screen, Choose View all instances.

30. Before you continue, wait for your instance to display the following:
    Instance state: Running
    Status check: 2/2 checks passed
    Tip: To refresh the instance information, choose the refresh icon.

## Task 4. Access the website on the second EC2 instance

31. Select the Web Server 2 instance.

32. In the Details tab copy the Public IPv4 address of your instance, then open a new tab in your web browser and paste in and load the address.

33. This will open a new tab in your web browser and display the web server page with the message Hello World! This is server 2.
    Make a note of the Availability Zones where the Web Server 1 and Web Server 2 instances are running. For example, us-east-1a and us-east-1b. You will need this information in the next task.

## Task 5. Create a load balancer

34. Back in the EC2 console, in the left navigation pane, under Load Balancing, choose Load Balancers.

35. Select Create Load Balancer.
    Tip: An Application Load Balancer has many features and can be used for HTTP and HTTPS traffic.

36. Under Application Load Balancer, choose Create.

37. In the Basic Configuration panel:
    - For Name, enter myloadbalancer.

38. In the Network mapping panel:
    - Under Mappings, select the Availability Zones that you created the two instances in. For example, us-east-1a and us-east-1b.
      Note: The Subnet to use in each selected Availability Zone will be automatically populated.

39. In the Security Groups panel:
    - Choose Web Server security group from the drop down menu.
    - After you close the drop down menu, choose the X next to the default security group to remove it.
40. In the Listeners and routing panel:
    - Choose Create target group.
      This will open a new tab in your browser.
41. In the Basic Configuration panel:
    - Keep the target type set to Instances.
    - For Target group name enter myalbTG
42. In the Health checks panel:
    - For Health check path, enter index.html after the forward slash ( / )
    - The path should look like the following: /index.html
43. Choose Next.

44. In the Register targets page, in the Available instances panel, check the boxes next to the Web Server 1 and Web Server 2 instances that you created in this lab.

45. Choose Include as pending below.
    Verify that both instances now appear in the Targets list below.

46. Choose Create target group.
    A banner displays the message that the target group was successfully created.

47. Return to the Load Balancers console tab in the browser.

48. In the Listeners and routing section, under Listener choose the refresh icon.

49. From the dropdown, chose the myalbTG target group you created.

50. Scroll down and chose Create load balancer.
    When the load balancer is created, a Successfully created load balancer message displays.

51. Choose View load balancer.
    Before continuing, ensure that the State of the load balancer that you just created changes to Active.
    It can take a few minutes for it to become active.
    Tip: To refresh the load balancer information, choose the refresh icon.

## Task 6. Test the load balancer

In this task, you will test the load balancer that you just created.

50. Select the load balancer that you just created, and expand the Details section.

51. Under Details, copy the DNS name value to your clipboard.

52. Open a new tab in your web browser, paste the DNS name that you just copied, and press Enter. If your load balancer is working, the Hello World! message displays. Notice whether the message says This is server 1 or This is server 2.

53. Refresh the browser tab a few times.
    Notice when the message changes between This is server 1 and This is server 2. When the message changes, it means that the load balancer has directed you to the web server on the other EC2 instance that you created in this lab.

## Lab complete

Congratulations! You have completed the lab.

54. Log out of the AWS Management Console.
    - In the upper-right corner of the page, choose your user name. Your user name begins with voclabs/user.
    - Choose Sign Out.

55. Choose End Lab at the top of this page, and then select Yes to confirm that you want to end the lab.