

Load Balancing For Clustered Barracuda CloudGen WAF Instances in Amazon Web Services (AWS)

<https://campus.barracuda.com/doc/41098448/>

This guide walks you through the steps to load balance traffic across multiple instances of the Barracuda CloudGen WAF deployed in Amazon Web Services.

In this article

- [Step 1 - Deploy Multiple Barracuda CloudGen WAF Instances in Amazon Web Services](#)
- [Step 2 - Set Up Load Balancing on the Barracuda CloudGen WAF Instances](#)
- [Step 3 - Set Up a High Availability Environment with the Barracuda CloudGen WAF](#)

To set up a High Availability environment with multiple Barracuda CloudGen WAF instances in Amazon Web Services, make sure all services configured on each instance use the WAN IP address of the Barracuda CloudGen WAF.

Step 1 - Deploy Multiple Barracuda CloudGen WAF Instances in Amazon Web Services


Follow the steps in [Deploy the Barracuda Web Application Firewall on Amazon Web Services](#) to deploy multiple Barracuda CloudGen WAF instances. To license and configure your virtual machine, continue with [Barracuda CloudGen WAF Deployment and Quick Start Guide for Amazon Web Services](#). In this example, there are two Barracuda CloudGen WAF instances where Barracuda-WAF1 is the first unit and Barracuda-WAF2 is the second unit.

Step 2 - Set Up Load Balancing on the Barracuda CloudGen WAF Instances

Load Balance the Service on Multiple Barracuda CloudGen WAF Instances Using the Application Load Balancer

1. Log into the [Amazon EC2 Management Console](#).
2. From the EC2 dashboard, select **Load Balancers** under **Load Balancing**.
3. Click **Create Load Balancer**. The **Select load balancer type** page opens.
4. On the **Select load balancer type** page, click **Create under Application Load Balancer**.


Application Load Balancer [Info](#)



Choose an Application Load Balancer when you need a flexible feature set for your applications with HTTP and HTTPS traffic. Operating at the request level, Application Load Balancers provide advanced routing and visibility features targeted at application architectures, including microservices and containers.

[Create](#)

Network Load Balancer [Info](#)



Choose a Network Load Balancer when you need ultra-high performance, TLS offloading at scale, centralized certificate deployment, support for UDP, and static IP addresses for your applications. Operating at the connection level, Network Load Balancers are capable of handling millions of requests per second securely while maintaining ultra-low latencies.

[Create](#)

Gateway Load Balancer [Info](#)



Choose a Gateway Load Balancer when you need to deploy and manage a fleet of third-party virtual appliances that support GENEVE. These appliances enable you to improve security, compliance, and policy controls.

[Create](#)

5. On the **Create Application Load Balancer** page:

1. **Basic Configuration**

1. **Load Balancer name** - Enter a name for the load balancer.
2. **Scheme** - Select the scheme to route the client requests to the server.
 1. **Internet-facing**: Routes client requests over the Internet.
 2. **Internal**: Routes requests using the private IP addresses.
3. **IP address type** - Select the IP address type (IPv4 or dualstack) to communicate with the load balancer.
 1. **IPv4**: Accepts only IPv4 traffic.
 2. **Dualstack**: Accepts IPv4 and IPv6 traffic.

2. **Network mapping**

1. **VPC** - Select the VPC.
2. **Mappings** - Select the availability zones and subnets for the VPC load balancer.

3. **Security groups**

1. **Security groups** - Select an existing security group from the drop-down list and assign it to the load balancer, or choose **Create new security group** to create a new group.

4. **Listeners and routing**

1. **Listener** - Configure the following:
 1. **Protocol** - Select the protocol (HTTP or HTTPS).
 2. **Port** - Specify the port number.
 3. **Default action** - Select an existing target group from the drop-down list, or click **Create target group** to register your targets (services). The load balancer routes the traffic to registered targets (services) using the port and protocol specified for the target group. See the section [Create a Target Group](#).
2. Click **Add Listener**.

5. **Tags (optional)**

1. Specify a key and a value for the tag. Click **Add tag** to add tags.

6. Summary

1. Review your settings before creating the load balancer, and then click **Create load balancer**.

Create a Target Group

You must create a target group and register the Barracuda CloudGen WAF instances or IP addresses for which the traffic needs to be load balanced. The load balancer routes the traffic to the registered targets using the specified port and protocol. For detailed information, refer to the AWS documentation.


1. Log into the [Amazon EC2 Management Console](#).
2. From the **EC2** dashboard, select **Target Groups** under **Load Balancing**.
3. On the **Target groups** page, click **Create target group**.
4. On the **Step 1 Specify group details** page, configure the following:
5. **Basic configuration**
 1. **Choose a target type** - If the Barracuda CloudGen WAF instances have created the service using the system IP address, select **Instances**. If the Barracuda CloudGen WAF instances have services with multiple IP addresses, select **IP addresses**.
 2. **Target group name** - Specify a name for the target group.
 3. **Protocol** - Select the protocol
 4. **Port** - Specify the port number.
 5. **VPC** - Select the VPC with the instances that need to be included in the target group.
 6. **IP address type** - (Available only when the target type is IP addresses) Select the IP address type for the IP addresses. Note: The target group may include only the selected IP address type targets (services).
 7. **Protocol version** - Select the protocol version.
6. **Health checks** - Modify the default settings as needed.
7. **Tags (Optional)** - Expand the tags section, and add the tags with key and value pair.
8. Click **Next**.
9. On the **Register targets** page, do the following:
 1. If the target type is **Instances**:
 1. Select the instances to which the traffic needs to be load balanced and click **Include as pending below**.
 2. If the target type is **IP addresses**:
 1. Select a network **VPC** from the list.
 2. Specify the IP addresses. You can add up to five IP addresses at a time.
 3. Specify the port number to route the traffic to specified IP addresses.
 4. Click **Include as pending below**.
10. Review the targets and click **Create target group**.

Load Balance the Service on the Barracuda CloudGen WAF Instances Using the Classic Load Balancer

1. Log into the [Amazon EC2 Management Console](#).
2. From the EC2 dashboard, select **Load Balancers** under **Load Balancing**.
3. Click **Create Load Balancer**. The **Select load balancer type** page opens.


4. In the **Select load balancer type** page, click **Create** under **Classic Load Balancer**.

Application Load Balancer Info




Choose an Application Load Balancer when you need a flexible feature set for your applications with HTTP and HTTPS traffic. Operating at the request level, Application Load Balancers provide advanced routing and visibility features targeted at application architectures, including microservices and containers.

Network Load Balancer Info



Choose a Network Load Balancer when you need ultra-high performance, TLS offloading at scale, centralized certificate deployment, support for UDP, and static IP addresses for your applications. Operating at the connection level, Network Load Balancers are capable of handling millions of requests per second securely while maintaining ultra-low latencies.


Gateway Load Balancer Info



Choose a Gateway Load Balancer when you need to deploy and manage a fleet of third-party virtual appliances that support GENEVE. These appliances enable you to improve security, compliance, and policy controls.

▼ **Classic Load Balancer - previous generation**

Classic Load Balancer Info



Choose a Classic Load Balancer when you have an existing application running in the EC2-Classical network.

Info AWS will be retiring the EC2-Classical network on August 15, 2022. [Learn more](#)

5. On the **Step 1: Define Load Balancer** page:

1. **Load Balancer Name** - Enter a name for the load balancer.
2. **Create LB Inside** - Select the VPC ID under which the Barracuda CloudGen WAF instances are launched.
3. Leave **Create an internal load balancer** set to the default value.
4. Select the **Enable advanced VPC configuration** check box.
5. **Listener Configuration** - Add the ports where services are created requiring load balancing.
6. Select the subnets for the VPC load balancer.
7. Click **Next: Assign Security Groups**.

6. On the **Step 2: Assign Security Groups** page, choose **Select an existing security group** to select and assign the security group(s) from an existing list, or choose **Create a new security group** to create a new group. Click **Next: Configure Security Settings**.

Ensure the selected group has all ports open, which were configured for the load balancer in Step 5.

7. Note on the **Step 3: Configure Security Settings**: Applicable only if load balancer is listening

to HTTPS traffic.

1. **Select Certificate**
 1. **Certificate Type:** Select the certificate type.
 2. **Certificate:** Select a certificate from the drop-down list.
2. **Select Cipher**
 1. **Predefined Security Policy:** Select a security policy from the drop-down list.
3. Click **Next: Configure Health Check**.
8. In the **Configure Health Check** page:
 1. **Ping Protocol** - Keep the default value, i.e., HTTP.
 2. **Ping Port** - Set to 8000. By default, the Barracuda CloudGen WAF listens on port 8000. If you are using a different port for the Barracuda CloudGen WAF, specify that port number.
 3. **Ping Path** - Enter `/cgi-mod/index.cgi`.
 4. In the **Advanced Details** section, specify required values and click **Next: Add EC2 Instances**.
9. On the **Step 5: Add EC2 Instances** page, select the instances to be added to this load balancer and click **Next: Add Tags**.
10. On the **Step 6. Add Tags** page, specify a key and a value for the tag and click **Review and Create**.
11. On the **Step 7: Review** page, review your settings before creating the load balancer, and then click **Create**.
12. The **Load Balancers** table displays the created load balancer details.

The services configured should be accessed using the DNS name of the created load balancer. For example, in the above example, the DNS name of the load balancer is `waf-lb1-334799786` and the HTTP service created on port 80 can be accessible via `http://waf-lb1-334799786 / http://waf-lb1-334799786:80`.

Step 3 - Set Up a High Availability Environment with the Barracuda CloudGen WAF

Follow these steps to cluster your Barracuda CloudGen WAF virtual machines in Amazon Web Services:

Before clustering your Barracuda CloudGen WAF virtual machines, ensure the following ports are open in the **Security Group** assigned to the Barracuda CloudGen WAF virtual machines:

| Port | Protocol |
|-------|----------|
| 8002 | TCP |
| 32575 | TCP |
| 32576 | UDP |

1. Install each system and ensure that each Barracuda CloudGen WAF is running the same firmware version. Each Barracuda CloudGen WAF in a cluster must have the same model

- number and firmware version.
2. Make a backup of each Barracuda CloudGen WAF configuration.
 3. No processes should be running on any virtual machine when you link them together. To be sure, go to the **ADVANCED > Task Manager** page of each Barracuda CloudGen WAF and verify that no processes are running.
 4. From the **ADVANCED > High Availability** page of Barracuda-WAF1, enter a **Cluster Shared Secret** password, and click **Save Changes**.
 5. From the **ADVANCED > High Availability** page of Barracuda-WAF2, do the following:
 1. Enter the same **Cluster Shared Secret** password, and click **Save Changes**. Both units in a cluster must have the same **Cluster Shared Secret** to communicate with each other.
 2. In the **Clustered Systems** section, enter the WAN IP address of Barracuda-WAF1, and click **Join Cluster**. Never cancel the join cluster task when the join is in progress.

The unit initiating the join cluster inherits the configuration from its peer unit and has its configuration overwritten.
 6. On each Barracuda CloudGen WAF, refresh the **ADVANCED > High Availability** page, and verify the following:
 1. Each system's hostname, serial number, and WAN IP address appears in the **Clustered Systems** list.
 2. The identity of the system (self or peer) displays in the **Type** field.
 3. The **Status** is green for all virtual machines in the cluster.
 7. View the **Cluster Status** from the **BASIC > Dashboard** page, under **Performance Statistics**.

To add more units to the existing cluster, repeat Steps 1 to 5.a., and then do the following:

1. From the **ADVANCED > High Availability** page of the Barracuda CloudGen WAF you are adding to the cluster, enter the WAN IP address of any system in the cluster in the **Peer IP Address** field and click **Join Cluster**. Verify the following:
 1. The configuration of the cluster automatically propagates to the newly added system.
 2. The new unit information propagates to all other units in the cluster.

Figures

1. Application_Load_Balancer.png
2. Load_Balancer_Types.png
3. Ports_Opened.jpg

© Barracuda Networks Inc., 2022 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.