




Configuring Auto Scale Group as Back-end Servers

Auto Scaling helps you maintain application availability and allows you to scale your Amazon EC2 capacity up or down automatically according to conditions you define. For more information on auto scaling, refer to the Amazon documentation: [Auto Scaling](#) and [Auto Scaling Groups](#) articles.


Before you continue with the steps mentioned below, ensure that you have completed the configuration settings mentioned in the [Clustering the Barracuda Load Balancer ADC Instances in Amazon Web Services](#) article.

To create an auto scaling group, perform the following steps:

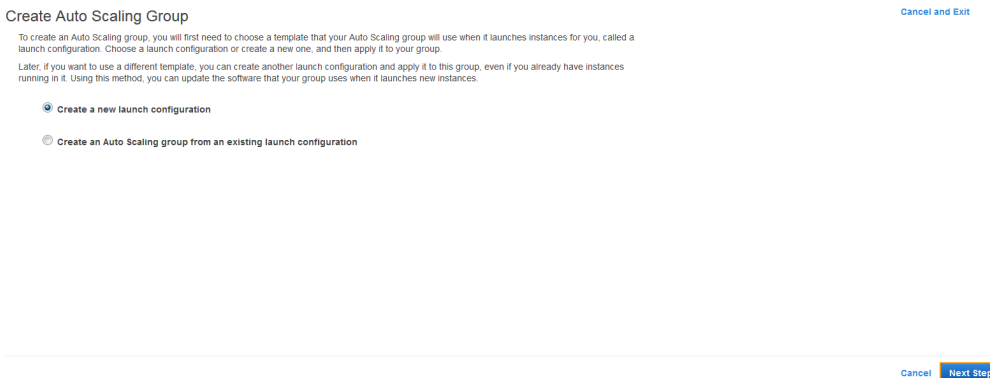
1. Go to the [EC2 Management Console](#).
2. Click **Auto Scaling Groups** under **AUTO SCALING**.



3. Click **Create Auto Scaling group**.



4. On the **Create Auto Scaling Group** page, select **Create a new launch configuration** and click **Next Step**.



5. On the **1. Choose AMI** page, select a server AMI,



1. Choose AMI 2. Choose Instance Type 3. Configure details 4. Add Storage 5. Configure Security Group 6. Review

Create Launch Configuration

All AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.

Quick Start < 1 to 30 of 30 AMIs >

- Amazon Linux AMI 2016.09.0 (HVM), SSD Volume Type - ami-5ec1673e**
The Amazon Linux AMI is an EBS-backed, AWS-supported image. The default image includes AWS command line tools, Python, Ruby, Perl, and Java. The repositories include Docker, PHP, MySQL, PostgreSQL, and other packages.
Root device type: ebs Virtualization type: hvm **Select** 64-bit
- Red Hat Enterprise Linux 7.3 (HVM), SSD Volume Type - ami-8f68c0f**
Red Hat Enterprise Linux version 7.3 (HVM), EBS General Purpose (SSD) Volume Type
Root device type: ebs Virtualization type: hvm **Select** 64-bit
- SUSE Linux Enterprise Server 12 SP2 (HVM), SSD Volume Type - ami-e4a30084**
SUSE Linux Enterprise Server 12 Service Pack 2 (HVM), EBS General Purpose (SSD) Volume Type. Public Cloud, Advanced Systems Management, Web and Scripting, and Legacy modules enabled.
Root device type: ebs Virtualization type: hvm **Select** 64-bit
- Ubuntu Server 16.04 LTS (HVM), SSD Volume Type - ami-a9d276c9**
Ubuntu Server 16.04 LTS (HVM), EBS General Purpose (SSD) Volume Type. Support available from Canonical (http://www.ubuntu.com/cloud/services).
Root device type: ebs Virtualization type: hvm **Select** 64-bit

6. On the **2. Choose Instance Type** page, select an instance type and click **Next: Configure details**.

1. Choose AMI 2. Choose Instance Type 3. Configure details 4. Add Storage 5. Configure Security Group 6. Review

Create Launch Configuration

<input type="checkbox"/>	General purpose	m4.16xlarge	64	256	EBS only	Yes	20 Gigabit
<input checked="" type="checkbox"/>	General purpose	m3.medium	1	3.75	1 x 4 (SSD)	-	Moderate
<input type="checkbox"/>	General purpose	m3.large	2	7.5	1 x 32 (SSD)	-	Moderate
<input type="checkbox"/>	General purpose	m3.xlarge	4	15	2 x 40 (SSD)	Yes	High
<input type="checkbox"/>	General purpose	m3.2xlarge	8	30	2 x 80 (SSD)	Yes	High
<input type="checkbox"/>	Compute optimized	c4.large	2	3.75	EBS only	Yes	Moderate
<input type="checkbox"/>	Compute optimized	c4.xlarge	4	7.5	EBS only	Yes	High
<input type="checkbox"/>	Compute optimized	c4.2xlarge	8	15	EBS only	Yes	High
<input type="checkbox"/>	Compute optimized	c4.4xlarge	16	30	EBS only	Yes	High
<input type="checkbox"/>	Compute optimized	c4.8xlarge	36	60	EBS only	Yes	10 Gigabit
<input type="checkbox"/>	Compute optimized	c3.large	2	3.75	2 x 16 (SSD)	-	Moderate
<input type="checkbox"/>	Compute optimized	c3.xlarge	4	7.5	2 x 40 (SSD)	Yes	Moderate

Cancel Previous **Next: Configure details**

7. On the **3. Configure details** page, do the following:
- Name:** Enter a name for the server instance.
 - IAM role:** Select the IAM role you have created.
 - In the **Advanced Details** section, select **Assign a public IP address to every instance** under **IP Address Type** and click **Next: Add Storage**.

1. Choose AMI 2. Choose Instance Type 3. Configure details 4. Add Storage 5. Configure Security Group 6. Review

Create Launch Configuration

Name

Purchasing option Request Spot Instances

IAM role

Monitoring Enable CloudWatch detailed monitoring [Learn more](#)

Advanced Details

Kernel ID

RAM Disk ID

User data As text As file Input is already base64 encoded
(Optional)

IP Address Type Only assign a public IP address to instances launched in the default VPC and subnet. (default)
 Assign a public IP address to every instance.
 Do not assign a public IP address to any instances.
Note: this option only affects instances launched into an Amazon VPC

Later, if you want to use a different launch configuration, you can create a new one and apply it to any Auto Scaling group. Existing launch configurations cannot be edited.

Cancel Previous **Skip to review** Next: Add Storage

8. On the **4: Add Storage** page, review the storage device settings for the instance. Modify the values if required, and then click **Next: Configure Security Group..**



1. Choose AMI 2. Choose Instance Type 3. Configure details 4. Add Storage 5. Configure Security Group 6. Review

Create Launch Configuration

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can attach additional EBS volumes after launching an instance, but not instance store volumes. <https://docs.aws.amazon.com/console/ec2/launchinstance/storage> about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput	Delete on Termination	Encrypted
Root	/dev/sda1	snap-826344d5	8	General Purpose (SSD)	24 / 3000	N/A	<input checked="" type="checkbox"/>	No

Add New Volume

Free tier eligible customers can get up to 30 GB of EBS storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Cancel Previous **Skip to review** Next: Configure Security Group

9. On the **5: Configure Security Group** page:

1. Choose **Select an existing security group** under **Assign a security group**.
2. Select the security groups that you created and click **Review**.

1. Choose AMI 2. Choose Instance Type 3. Configure details 4. Add Storage 5. Configure Security Group 6. Review

Create Launch Configuration

sg-01e6467 This security group was generated by AWS Marketplace and is based on recommended set

sg-17003d72 Barracuda Load Balancer ADC - BYOL-S-1-009-AutogenByAWSMP-1 This security group was generated by AWS Marketplace and is based on recommended set

Inbound rules for sg-17003d72 Selected security groups: sg-17003d72.

Type	Protocol	Port Range	Source
Custom TCP Rule	TCP	40080	0.0.0.0/0
HTTP	TCP	80	0.0.0.0/0
Custom TCP Rule	TCP	8000	0.0.0.0/0
HTTPS	TCP	443	0.0.0.0/0

Cancel Previous **Review**

10. On the **6: Review** page, review your settings and click **Create launch configuration**.

1. Choose AMI 2. Choose Instance Type 3. Configure details 4. Add Storage 5. Configure Security Group 6. Review

Create Launch Configuration

AMI Details [Edit AMI](#)

Free tier eligible **Ubuntu Server 16.04 LTS (HVM), SSD Volume Type - ami-a9d276c9**
 Ubuntu Server 16.04 LTS (HVM),EBS General Purpose (SSD) Volume Type. Support available from Canonical (<http://www.ubuntu.com/cloud/services>).
 Root device type: ebs Virtualization type: hvm

Instance Type [Edit instance type](#)

Instance Type	ECUs	vCPUs	Memory GiB	Instance Storage (GiB) GiB	EBS-Optimized Available	Network Performance
m3.medium	3	1	3.75	1 x 4	-	Moderate

Launch configuration details [Edit details](#)

Name: multi-AZ-Autoscale
 Purchasing option: On demand
 EBS Optimized: No
 Monitoring: No
 IAM role: byol-autoscale
 Tenancy: Shared tenancy (multi-tenant hardware)

Cancel Previous **Create launch configuration**

11. In the **Select an existing key pair or create a new key pair** pop-up window:

1. Choose an existing key pair or create a new pair.
2. Select the **I acknowledge that I have access to the selected private key file ("key pair name")**, and that **without this file, I won't be able to log into my instance** check box.
3. Click **Create launch configuration**.



Select an existing key pair or create a new key pair ✕

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about [removing existing key pairs from a public AMI](#).

Choose an existing key pair ▼

Select a key pair

az-autoscale-keypair ▼

I acknowledge that I have access to the selected private key file (az-autoscale-keypair.pem), and that without this file, I won't be able to log into my instance.

Cancel
Create launch configuration

12. On the **1. Configure Auto Scaling group details** page, specify values for the following:
 1. **Group name:** Enter a name for the auto scale group.
 2. **Group size:** Enter the number of instances you want to add in the auto scale group.
 3. **Network:** Select the VPC created by you.
 4. **Subnet:** Select the subnets created for the VPC.
 5. Click **Next: Configure scaling policies**.

1. Configure Auto Scaling group details 2. Configure scaling policies 3. Configure Notifications 4. Configure Tags 5. Review

[Cancel and Exit](#)

Create Auto Scaling Group

Launch Configuration ⓘ multi-AZ-Autoscale

Group name ⓘ multi-AZ-Autoscale

Group size ⓘ Start with 2 instances

Network ⓘ adc-vpc C Create new VPC

Subnet ⓘ

us-west-2a	-subnet-1	x
us-west-2b	-subnet-2	x

[Create new subnet](#)

Each instance in this Auto Scaling group will be assigned a public IP address. ⓘ

▶ Advanced Details

[Cancel](#) [Next: Configure scaling policies](#)

13. On the **2. Configure scaling policies** page, select **Keep this group at its initial size** or select **Use scaling policies to adjust the capacity of the group** to configure as per your requirement.

1. Configure Auto Scaling group details 2. Configure scaling policies 3. Configure Notifications 4. Configure Tags 5. Review

Create Auto Scaling Group

You can optionally add scaling policies if you want to adjust the size (number of instances) of your group automatically. A scaling policy is a set of instructions for making such adjustments in response to an Amazon CloudWatch alarm that you assign to it. In each policy, you can choose to add or remove a specific number of instances or a percentage of the existing group size, or you can set the group to an exact size. When the alarm triggers, it will execute the policy and adjust the size of your group accordingly. [Learn more](#) about scaling policies.

Keep this group at its initial size
 Use scaling policies to adjust the capacity of this group

[Cancel](#) [Previous](#) [Review](#) [Next: Configure Notifications](#)

14. On the **3. Configure Notifications** page, click **Add notification** to configure notification settings (if required).



1. Configure Auto Scaling group details 2. Configure scaling policies 3. Configure Notifications 4. Configure Tags 5. Review

Create Auto Scaling Group

Configure your Auto Scaling group to send notifications to a specified endpoint, such as an email address, whenever a specified event takes place, including: successful launch of an instance, failed instance launch, instance termination, and failed instance termination.

If you created a new topic, check your email for a confirmation message and click the included link to confirm your subscription. Notifications can only be sent to confirmed addresses.

Add notification

Cancel Previous Review Next: Configure Tags

15. On the **4: Configure Tags** page, add/remove the tags for the instance (if required) and click **Review**.

1. Configure Auto Scaling group details 2. Configure scaling policies 3. Configure Notifications 4. Configure Tags 5. Review

Create Auto Scaling Group

A tag consists of a case sensitive key-value pair that you can use to identify your group. For example, you could define a tag with Key = Environment and Value = Production. You can optionally choose to apply these tags to instances in the group when they launch. [Learn more](#).

Key	Value	Tag New Instances
Name	multi-AZ-Autoscale	<input checked="" type="checkbox"/>

Add tag 49 remaining

Cancel Previous Review

16. On the **5: Review** page, review your settings before creating the auto scaling group, and click **Create Auto Scaling group**.

1. Configure Auto Scaling group details 2. Configure scaling policies 3. Configure Notifications 4. Configure Tags 5. Review

Create Auto Scaling Group

Please review your Auto Scaling group details. You can go back to edit changes for each section. Click **Create Auto Scaling group** to complete the creation of an Auto Scaling group.

- Auto Scaling Group Details [Edit details](#)
 - Group name: multi-AZ-Autoscale
 - Group size: 2
 - Minimum Group Size: 2
 - Maximum Group Size: 2
 - Subnet(s): subnet-14abf170, subnet-b1c29ec7
 - Health Check Grace Period: 300
 - Detailed Monitoring: No
 - Instance Protection: None
- Scaling Policies [Edit scaling policies](#)
- Notifications [Edit notifications](#)
- Tags [Edit tags](#)
 - Name: multi-AZ-Autoscale tag new instances

Cancel Previous Create Auto Scaling group

Auto Scaling group creation status

✓ Successfully created Auto Scaling group
View creation log

17. The created auto scale group gets displayed in the auto scale group list.

Create Auto Scaling group Actions

Filter: Filter Auto Scaling groups... 1 to 1 of 1 Auto Scaling Groups

Name	Launch Configuration	Instances	Desired	Min	Max	Availability Zones	Default Cooldown	Health Check Gra
multi-AZ-Autos...	multi-AZ-Autoscale	2	2	2	2	us-west-2a, us-west-2b	300	300

To configure auto scale group as your servers, use the Barracuda Load Balancer ADC web interface. Refer to [Clustering the Barracuda Load Balancer ADC Instances in Different Availability Zones](#).



Next Step

Continue with [Clustering the Barracuda Load Balancer ADC Instances in the Same Availability Zone](#) or [Clustering the Barracuda Load Balancer ADC Instances in Different Availability Zones](#).

