

Name: Trasily Krishnan Shanmugham

Cloud Computing project Report

An online business can use load balancing and auto-scaling to support unexpected usage peaks and save money when usage is lower.

Amazon web services (AWS) provide elastic load balancing to automatically distribute incoming web traffic across multiple Amazon EC2 instances.

If one instance is getting too much traffic, some of it can be rerouted to another instances for processing. With load balancing we can add or remove Amazon EC2 instances as when needed without disrupting the overall flow of information.

The main features of load balancing is distribution

- Load balancing can automatically re-route to other running instances or equally distributes to all the their available instances.
- Node failures - If one instance fails, it reroutes the traffic to other instances.

Managing applications using load balancing has a lots of benefits. It minimizes the risk of overloading one single instance and provides health monitoring for Amazon EC2 instances. Health check monitoring helps monitoring the traffic, handles incoming requests and checks continually to make sure that the load-balancing running between the instances function properly.

We can create and manage security groups with the load balancing to provide additional security and networking options. We can also set up load balancer in single availability zones or multiple availability zones to load balance incoming application traffic across Amazon EC2 instances. An availability zone is like a datacenter, which contains collection of computers in the physical locations, which will protect our application with the single location failure. We can access elastic load balancer through AWS management console, which is simple and intuitive web interface.

Auto scaling is one of the important services to build highly available applications. It eliminates the single point of failure by distributing traffic across availability zones. It can be self-healing by launching new instances, if any of the instances failed. We need to enable load balancer before configuring auto scaling.

We need to launch instances even for auto scaling. When usage is lower, we can switch off the server or instances to save money. Auto scaling can automatically increase the number of instances during high availability and in demand to performance maintenance. Auto Scaling is useful for both to applications that have stable demand patterns or that experience hourly, daily, or weekly variability in usage.

Load Balancing – For example: In a large organization, there are 5 servers and the network is equally distributing between all the servers. If one server fails, it automatically reroutes to the other four servers using dynamic load balancing.

Auto scaling – For example: In a flower shop, the web orders will be high during some occasion. But other normal days, orders will be less. During peak hours, we never know when we will get large number of orders. So we have to be ready with availability of servers using auto scaling. Auto scaling will automatically turn off the servers during less usage to avoid high cost. Also in online shopping, there will be high load on the servers during sale, daytime, weekends, etc. Auto scaling will automatically minimize the number of servers during low traffic (e.g: early morning).

First we need to create an AWS account. We can see how to launch Instance, create load balancer, configure health checks, assign security groups, register, review settings, create load balancer and create auto scaling.

Creating Load balancer:

EC2 Dashboard

Events

Tags

Reports

Limits

INSTANCES

Instances

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Resources

You are using the following Amazon EC2 resources in the US West (N. California) region:

0 Running Instances

0 Volumes

0 Key Pairs

0 Placement Groups

0 Elastic IPs

0 Snapshots

0 Load Balancers

1 Security Groups

Automate application deployments to EC2 with CodeDeploy.

Hide

Create Instance

To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.

Launch Instance

Account Attributes

Supported Platforms

VPC

Default VPC

vpc-34b76b51

Additional Information

Getting Started Guide

Documentation

All EC2 Resources

Forums

Pricing

Contact Us

AWS

Services

Edit

Trasly

Oregon

Support

1. Define Load Balancer

2. Assign Security Groups

3. Configure Security Settings

4. Configure Health Check

5. Add EC2 Instances

6. Add Tags

7. Review

Step 1: Define Load Balancer

Basic Configuration

This wizard will walk you through setting up a new load balancer. Begin by giving your new load balancer a unique name so that you can identify it from other load balancers you might create. You will also need to configure ports and protocols for your load balancer. Traffic from your clients can be routed from any load balancer port to any port on your EC2 instances. By default, we've configured your load balancer with a standard web server on port 80.

Load Balancer name:

Create LB inside:

Create an internal load balancer: ☐ (what's this?)

Enable advanced VPC configuration: ☐

Listener Configuration:

| Load Balancer Protocol | Load Balancer Port | Instance Protocol | Instance Port |
|-----------------------------------|---------------------------------|-----------------------------------|---------------------------------|
| <input type="text" value="HTTP"/> | <input type="text" value="80"/> | <input type="text" value="HTTP"/> | <input type="text" value="80"/> |
| <div>Add</div> | | | |

AWS

Services

Edit

Trasly

Oregon

Support

1. Define Load Balancer

2. Assign Security Groups

3. Configure Security Settings

4. Configure Health Check

5. Add EC2 Instances

6. Add Tags

7. Review

Step 2: Assign Security Groups

You have selected the option of having your Elastic Load Balancer inside of a VPC, which allows you to assign security groups to your load balancer. Please select the security groups to assign to this load balancer. This can be changed at any time.

Assign a security group: ☐ Create a new security group

☒ Select an existing security group

Filter

VPC security groups

| Security Group ID | Name | Description | Actions |
|---|---------|----------------------------|-----------------------------|
| <input checked="" type="checkbox"/> sg-25f01b41 | default | default VPC security group | Copy to new |

AWS

Services

Edit

Trasly

Oregon

Support

1. Define Load Balancer

2. Assign Security Groups

3. Configure Security Settings

4. Configure Health Check

5. Add EC2 Instances

6. Add Tags

7. Review

Step 4: Configure Health Check

Your load balancer will automatically perform health checks on your EC2 instances and only route traffic to instances that pass the health check. If an instance fails the health check, it is automatically removed from the load balancer. Customize the health check to meet your specific needs.

Ping Protocol:

Ping Port:

Ping Path:

Advanced Details

Response Timeout: seconds

Health Check Interval: seconds

Unhealthy Threshold:

Healthy Threshold:

AWS

Services

Edit

Trasly

Oregon

Support

1. Define Load Balancer

2. Assign Security Groups

3. Configure Security Settings

4. Configure Health Check

5. Add EC2 Instances

6. Add Tags

7. Review

Step 7: Review

Please review the load balancer details before continuing

Define Load Balancer

Load Balancer name: My-load-balancer

Scheme: internet-facing

Port Configuration: 80 (HTTP) forwarding to 80 (HTTP)

Edit load balancer definition

Configure Health Check

Ping Target: HTTP:80/index.html

Timeout: 5 seconds

Interval: 30 seconds

Unhealthy Threshold: 2

Healthy Threshold: 10

Edit health check

Add EC2 Instances

Cross-Zone Load Balancing: Enabled

Connection Draining: Enabled, 300 seconds

Instances:

Edit instances

Load Balancer Creation Status

✓ Successfully created load balancer

Load balancer My-load-balancer was successfully created.

Note: It may take a few minutes for your instances to become active in the new load balancer.

Close

Launch Instances and add to load balancer:

EC2 Dashboard

Events

Tags

Reports

Limits

INSTANCES

INSTANCES

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Resources

You are using the following Amazon EC2 resources in the US West (Oregon) region:

0 Running Instances

0 Elastic IPs

0 Volumes

0 Snapshots

0 Key Pairs

1 Load Balancers

0 Placement Groups

1 Security Groups

Automate application deployments to EC2 with CodeDeploy.

Hide

Create Instance

To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.

Launch Instance

Account Attributes

Supported Platforms

VPC

Default VPC

vpc-4d8b1c28

Additional Information

Getting Started Guide

Documentation

All EC2 Resources

Forums

Pricing

Contact Us

1. Choose AMI

2. Choose Instance Type

3. Configure Instance

4. Add Storage

5. Tag Instance

6. Configure Security Group

7. Review

Cancel and Exit

Step 1: Choose an Amazon Machine Image (AMI)

An 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

My AMIs

AWS Marketplace

Community AMIs

Free tier only

Amazon Linux

Free tier eligible

Amazon Linux AMI 2015.03 (HVM), SSD Volume Type - ami-e7527ed7

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

Free tier eligible

Red Hat Enterprise Linux 7.1 (HVM), SSD Volume Type - ami-4dbf9e7d

Red Hat Enterprise Linux version 7.1 (HVM), EBS General Purpose (SSD) Volume Type

Root device type: ebs

Virtualization type: hvm

Select

64-bit

1. Choose AMI

2. Choose Instance Type

3. Configure Instance

4. Add Storage

5. Tag Instance

6. Configure Security Group

7. Review

Step 3: Configure Instance Details

Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot Instances to take advantage of the lower pricing, assign an access management role to the instance, and more.

Number of instances

1

Purchasing option

☐ Request Spot Instances

Network

vpc-4d8b1c28 (172.31.0.0/16) (default)

Create new VPC

Subnet

No preference (default subnet in any Availability Zone)

Create new subnet

Auto-assign Public IP

Use subnet setting (Enable)

IAM role

None

Create new IAM role

Shutdown behavior

Stop

Enable termination protection

☐ Protect against accidental termination

Monitoring

☐ Enable CloudWatch detailed monitoring

Additional charges apply.

Cancel

Previous

Review and Launch

Next: Add Storage

1. Choose AMI

2. Choose Instance Type

3. Configure Instance

4. Add Storage

5. Tag Instance

6. Configure Security Group

7. Review

Step 4: Add Storage

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 also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

| Type | Device | Snapshot | Size (GiB) | Volume Type | IOPS | Delete on Termination | Encrypted |
|------|-----------|---------------|------------|-----------------------|-----------|-------------------------------------|---------------|
| Root | /dev/xvda | snap-bfb086e1 | 8 | General Purpose (SSD) | 24 / 3000 | <input checked="" type="checkbox"/> | Not Encrypted |

Add New Volume

1. Choose AMI

2. Choose Instance Type

3. Configure Instance

4. Add Storage

5. Tag Instance

6. Configure Security Group

7. Review

Step 6: Configure Security Group

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. [Learn more](#) about Amazon EC2 security groups.

Assign a security group:

☒ Create a new security group

☐ Select an existing security group

Security group name:

launch-wizard-1

Description:

launch-wizard-1 created 2015-06-14T13:59:56.515-07:00

| Type | Protocol | Port Range | Source |
|------|----------|------------|--------------------|
| SSH | TCP | 22 | Anywhere 0.0.0.0/0 |

Add Rule

1. Choose AMI

2. Choose Instance Type

3. Configure Instance

4. Add Storage

5. Tag Instance

6. Configure Security Group

7. Review

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

Improve your instances' security. Your security group, launch-wizard-1, is open to the world.

Your instances may be accessible from any IP address. We recommend that you update your security group rules to allow access from known IP addresses only.

You can also open additional ports in your security group to facilitate access to the application or service you're running, e.g., HTTP (80) for web servers. [Edit security groups](#)

AMI Details

Amazon Linux

Free tier eligible

Amazon Linux AMI 2015.03 (HVM), SSD Volume Type - ami-e7527ed7

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

Edit AMI

Instance Type

| Instance Type | ECUs | vCPUs | Memory (GiB) | Instance Storage (GB) | EBS-Optimized Available | Network Performance |
|---------------|----------|-------|--------------|-----------------------|-------------------------|---------------------|
| t2.micro | Variable | 1 | 1 | EBS only | - | Low to Moderate |

Edit instance type

Security Groups

[Edit security groups](#)

AWS

Services

Edit

Trasly

Oregon

Support

Launch Status

Your instances are now launching

The following instance launches have been initiated: [i-a5642852](#) [View launch log](#)

Get notified of estimated charges

[Get notified of estimated charges](#)

Add and Remove Instances

The table below lists all your running EC2 Instances. Check the boxes in the Select column to add those instances to this load balancer.

Warning: Unchecking instances and clicking save will remove these instances from your load balancer.

Add or Remove Instances

| <input type="checkbox"/> | Instance | Name | State | Security Groups | Zone | Subnet ID | Subnet CIDR |
|--------------------------|------------|------|---------|-----------------|------------|-----------------|----------------|
| <input type="checkbox"/> | i-a5642852 | | running | launch-wizard-1 | us-west-2b | subnet-fa8ee29f | 172.31.16.0/20 |

AWS

Services

Edit

Trasly

Oregon

Support

EC2 Dashboard

Events

Tags

Reports

Limits

INSTANCES

Instances

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Snapshots

NETWORK & SECURITY

Security Groups

Elastic IPs

Placement Groups

Load Balancers

Key Pairs

Network Interfaces

Create Load Balancer

Actions

Filter: Search Load Balancers

<< 1 to 1 of 1 >>

| <input type="checkbox"/> | Load Balancer Name | DNS Name | Port Configuration | Availability Zones | Instance Count | Health Check |
|--------------------------|--------------------|------------------------------|----------------------------------|---------------------------|----------------|--------------|
| <input type="checkbox"/> | My-load-balancer | My-load-balancer-10041784... | 80 (HTTP) forwarding to 80 (...) | us-west-2a, us-west-2b... | 1 Instance | HTTP:80/... |

Load balancer: My-load-balancer

Description

Instances

Health Check

Monitoring

Security

Listeners

Tags

Connection Draining: Enabled, 300 seconds (Edit)

Edit Instances

| Instance ID | Name | Availability Zone | Status | Actions |
|-------------|------|-------------------|----------------|---|
| i-a5642852 | | us-west-2b | OutOfService ⓘ | Remove from Load Balancer |

Edit Availability Zones

| Availability | Subnet ID | Subnet CIDR | Instance | Healthy? | Actions |
|--------------|-----------------|----------------|----------|--|---|
| us-west-2a | subnet-4722a030 | 172.31.32.0/20 | 0 | No (Availability Zone contains no healthy instances) | Remove from Load Balancer |

AWS

Services

Edit

Trasly

Oregon

Support

EC2 Dashboard

Events

Tags

Reports

Limits

INSTANCES

Instances

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Create Load Balancer

Actions

Filter: Search Load Balancers

<< 1 to 1 of 1 >>

| <input type="checkbox"/> | Load Balancer Name | DNS Name | Port Configuration | Availability Zones | Instance Count | Health Check |
|--------------------------|--------------------|------------------------------|----------------------------------|---------------------------|----------------|--------------|
| <input type="checkbox"/> | My-load-balancer | My-load-balancer-10041784... | 80 (HTTP) forwarding to 80 (...) | us-west-2a, us-west-2b... | 1 Instance | HTTP:80/... |

Load balancer: My-load-balancer

Description

Instances

Health Check

Monitoring

Security

Listeners

Tags

| Security Group ID | Name | Description |
|-------------------|---------|----------------------------|
| sg-25f01b41 | default | default VPC security group |

AWS

Services

Edit

Trasly

Oregon

Support

EC2 Dashboard

Events

Tags

Reports

Limits

INSTANCES

Instances

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Snapshots

NETWORK & SECURITY

Security Groups

Elastic IPs

Placement Groups

Load Balancers

Key Pairs

Create Load Balancer

Actions

Filter: Search Load Balancers

<< 1 to 1 of 1 >>

| <input type="checkbox"/> | Load Balancer Name | DNS Name | Port Configuration | Availability Zones | Instance Count | Health Check |
|--------------------------|--------------------|------------------------------|----------------------------------|---------------------------|----------------|--------------|
| <input type="checkbox"/> | My-load-balancer | My-load-balancer-10041784... | 80 (HTTP) forwarding to 80 (...) | us-west-2a, us-west-2b... | 1 Instance | HTTP:80/ind |

Load balancer: My-load-balancer

Description

Instances

Health Check

Monitoring

Security

Listeners

Tags

DNS Name:

My-load-balancer-1004178448.us-west-2.elb.amazonaws.com (A Record)

Note: Because the set of IP addresses associated with a LoadBalancer can change over time, you should never create an "A" record with any specific IP address. If you want to use a friendly DNS name for your load balancer instead of the name generated by the Elastic Load Balancing service, you should create a CNAME record for the LoadBalancer DNS name, or use Amazon Route 53 to create a hosted zone. For more information, see [Using Domain Names With Elastic Load Balancing](#).

Scheme:

internet-facing

Status:

0 of 1 Instances in service

Port Configuration:

80 (HTTP) forwarding to 80 (HTTP)

Stickiness: Disabled (Edit)

Availability Zones:

subnet-4722a030 - us-west-2a, subnet-54a00104 - us-west-2b

Create Autoscaling Group and add instances:

The screenshot shows the AWS Management Console interface for the 'Auto Scaling' service. The left-hand navigation pane lists various AWS services, with 'Auto Scaling Groups' highlighted under the 'AUTO SCALING' category. The main content area is titled 'Welcome to Auto Scaling' and includes a 'Create Auto Scaling group' button. Below this, there are three sections: 'Reusable Instance Templates', 'Automated Provisioning', and 'Adjustable Capacity', each with an icon and a brief description. The right-hand side of the console shows 'Additional Information' links such as 'Getting Started Guide', 'Documentation', and 'All EC2 Resources'.

Create Auto Scaling Group

[Cancel and Exit](#)

To create an Auto Scaling group, you will first need to choose a template that your Auto Scaling group will use when it launches instances for you, called a launch configuration. Choose a launch configuration or create a new one, and then apply it to your group.

Later, if you want to use a different template, you can create another launch configuration and apply it to this group, even if you already have instances running in it. Using this method, you can update the software that your group uses when it launches new instances.



Step 1: Create launch configuration

First, define a template that your Auto Scaling group will use to launch instances. You can change your group's launch configuration at any time.



Step 2: Create Auto Scaling group

[Cancel](#) [Create launch configuration](#)

The screenshot shows the 'Create Launch Configuration' wizard in the AWS Management Console. The wizard is divided into six steps: 1. Choose AMI, 2. Choose Instance Type, 3. Configure details, 4. Add Storage, 5. Configure Security Group, and 6. Review. The current step is '3. Configure details'. The form includes fields for 'Name' (myauto), 'Purchasing option' (Request Spot Instances), 'IAM role' (None), and 'Monitoring' (Enable CloudWatch detailed monitoring). Below the form is a section for 'Advanced Details'. At the bottom, there is a modal dialog titled 'Select an existing key pair or create a new key pair' which contains a list of key pairs (Mykey) and a checkbox for acknowledging access to the private key file.

AWS

Services

Edit

Trasly

Oregon

Support

1. Configure Auto Scaling group details

2. Configure scaling policies

3. Configure Notifications

4. Configure Tags

5. Review

Create Auto Scaling Group

Cancel and Exit

Launch Configuration

myauto

Group name

autoscale

Group size

Start with 1 Instances

Network

vpc-4d8b1c28 (172.31.0.0/16) (default)

Create new VPC

Subnet

subnet-fa8ee29f(172.31.16.0/20) | Default in us-west-2b

subnet-4722a030(172.31.32.0/20) | Default in us-west-2a

subnet-54a0010d(172.31.0.0/20) | Default in us-west-2c

Create new subnet

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

AWS

Services

Edit

Trasly

Oregon

Support

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

Scale between 1 and 1 instances. These will be the minimum and maximum size of your group.

Increase Group Size

Name:

Increase Group Size

Execute policy when:

No alarm selected

Add new alarm

Take the action:

Add

0

instances

And then wait:

300

seconds before allowing another scaling activity

AWS

Services

Edit

Trasly

Oregon

Support

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

AWS

Services

Edit

Trasly

Oregon

Support

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

autoscale

Group size

1

Minimum Group Size

1

Maximum Group Size

1

Subnet(s)

subnet-fa8ee29f,subnet-4722a030,subnet-54a0010d

Load Balancers

My-load-balancer

Health Check Type

EC2

Health Check Grace Period

300

Detailed Monitoring

No

Scaling Policies

Edit scaling policies

Increase Group Size

With alarm = None; Add 0 instances and 300 seconds between activities

Decrease Group Size

With alarm = None; Remove 0 instances and 300 seconds between activities

Notifications

Edit notifications

Cancel

Previous

Create Auto Scaling group

Auto Scaling group creation status

✓ **Successfully created Auto Scaling group**
[View creation log](#)

- ▼ View
- [View your Auto Scaling groups](#)
 - [View your launch configurations](#)

► Here are some helpful resources to get you started

Close

INSTANCES

Instances

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Snapshots

NETWORK & SECURITY

Security Groups

Elastic IPs

Placement Groups

Load Balancers

Key Pairs

Network Interfaces

AUTO SCALING

Create Auto Scaling group

Actions

Filter:

< 1 to 1 of 1 Auto Scaling Groups >

| Name | Launch Configuration | Instances | Desired | Min | Max | Availability Zones | Default Cooldown | Health |
|-----------|----------------------|-----------|---------|-----|-----|--------------------------------|------------------|--------|
| autoscale | myauto | 0 | 1 | 1 | 1 | us-west-2c, us-west-2b, us-... | 300 | 300 |

Auto Scaling Group: autoscale

Details

Activity History

Scaling Policies

Instances

Notifications

Tags

Filter: Any Status

< 1 to 1 of 1 History Items >

| Status | Description | Start Time | End Time |
|------------|--|-----------------------------|-----------------------------|
| Successful | Launching a new EC2 instance: i-7fd02eb7 | 2015 June 14 14:28:35 UTC-7 | 2015 June 14 14:27:07 UTC-7 |

INSTANCES

Instances

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Snapshots

NETWORK & SECURITY

Security Groups

Elastic IPs

Placement Groups

Load Balancers

Key Pairs

Network Interfaces

AUTO SCALING

Create Auto Scaling group

Actions

Filter:

< 1 to 1 of 1 Auto Scaling Groups >

| Name | Launch Configuration | Instances | Desired | Min | Max | Availability Zones | Default Cooldown | Health |
|-----------|----------------------|-----------|---------|-----|-----|--------------------------------|------------------|--------|
| autoscale | myauto | 0 | 1 | 1 | 1 | us-west-2c, us-west-2b, us-... | 300 | 300 |

Auto Scaling Group: autoscale

Details

Activity History

Scaling Policies

Instances

Notifications

Tags

Launch Configuration

myauto

Load Balancers

My-load-balancer

Desired

1

Min

1

Max

1

Health Check Type

EC2

Health Check Grace Period

300

Availability Zone(s)

us-west-2c, us-west-2b, us-west-2a

Subnet(s)

subnet-faBee29f,subnet-4722a030,subnet-54a0010d

Default Cooldown

300

Placement Group

Suspended Processes

INSTANCES

Instances

Spot Requests

Reserved Instances

IMAGES

AMIs

Bundle Tasks

ELASTIC BLOCK STORE

Volumes

Snapshots

NETWORK & SECURITY

Security Groups

Elastic IPs

Placement Groups

Load Balancers

Key Pairs

Network Interfaces

AUTO SCALING

Create Auto Scaling group

Actions

Filter:

< 1 to 1 of 1 Auto Scaling Groups >

| Name | Launch Configuration | Instances | Desired | Min | Max | Availability Zones | Default Cooldown | Health |
|-----------|----------------------|-----------|---------|-----|-----|--------------------------------|------------------|--------|
| autoscale | myauto | 1 | 1 | 1 | 1 | us-west-2c, us-west-2b, us-... | 300 | 300 |

Auto Scaling Group: autoscale

Details

Activity History

Scaling Policies

Instances

Notifications

Tags

Actions

Filter: Any Health Status Any Lifecycle State

< 1 to 1 of 1 Instances >

| Instance ID | Lifecycle | Launch Configuration Name | Availability Zone | Health Status |
|-------------|-----------|---------------------------|-------------------|---------------|
| i-7fd02eb7 | InService | myauto | us-west-2c | Healthy |