

---

# Amazon Elastic Compute Cloud

## User Guide

API Version 2013-10-01



# Amazon Web Services

## Amazon Elastic Compute Cloud: User Guide

Amazon Web Services

Copyright © 2013 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

The following are trademarks of Amazon Web Services, Inc.: Amazon, Amazon Web Services Design, AWS, Amazon CloudFront, Cloudfront, Amazon DevPay, DynamoDB, ElastiCache, Amazon EC2, Amazon Elastic Compute Cloud, Amazon Glacier, Kindle, Kindle Fire, AWS Marketplace Design, Mechanical Turk, Amazon Redshift, Amazon Route 53, Amazon S3, Amazon VPC. In addition, Amazon.com graphics, logos, page headers, button icons, scripts, and service names are trademarks, or trade dress of Amazon in the U.S. and/or other countries. Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon.

All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

---

What is Amazon EC2? .....	1
Instances and AMIs .....	4
Regions and Availability Zones .....	7
Root Device Volume .....	12
Get Set Up .....	18
Getting Started .....	22
Step 1: Launch an Instance .....	23
Step 2: Connect to Your Instance .....	24
Step 3: Add a Volume .....	27
Step 4: Clean Up .....	29
Tutorial: Installing a LAMP Web Server .....	30
Tutorial: Hosting a WordPress Blog .....	36
Amazon Machine Images .....	43
AMI Types .....	44
Finding a Suitable AMI .....	47
Shared AMIs .....	49
Finding Shared AMIs .....	49
Making an AMI Public .....	50
Sharing an AMI with Specific AWS Accounts .....	51
Using Bookmarks .....	52
Guidelines for Shared Linux AMIs .....	53
Paid AMIs .....	57
Creating Amazon EBS-Backed AMIs Using the Console .....	60
Creating Your Own AMIs .....	62
Creating Amazon EBS-Backed Linux AMIs .....	63
Creating Instance Store-Backed Linux/UNIX AMIs .....	66
Tools You Need .....	67
From an Existing AMI .....	68
From a Loopback .....	72
Creating and Launching an AMI from a Snapshot .....	80
Using Your Own Linux Kernels .....	80
Copying AMIs .....	84
Amazon Linux .....	87
Instances .....	94
Instance Types .....	94
Micro Instances .....	95
H1 Instances .....	102
HS1 Instances .....	104
GPU Instances .....	105
EBS-Optimized Instances .....	107
Placement Groups .....	108
Resizing Instances .....	111
Spot Instances .....	115
Getting Started with Spot Instances .....	116
Viewing Spot Instance Pricing History .....	118
Creating a Spot Instance Request .....	122
Finding Running Spot Instances .....	126
Canceling Spot Instance Requests .....	129
Fundamentals of Spot Instances .....	132
Placing Spot Requests .....	132
Spot Instance Limits .....	133
Customizing Your Spot Requests .....	134
Tracking Spot Requests with Bid Status Codes .....	135
Tagging Spot Instance Requests .....	141
Protecting Your Spot Instance Data Against Interruptions .....	142
Planning for Interruptions .....	142
Persisting Your Root EBS Partition .....	142
Walkthroughs: Using Spot Instances with AWS Services .....	143



Managing Spot Instances with Auto Scaling .....	143
Tools for Managing Auto Scaling with Spot Instances .....	144
Launching Spot Instances with Auto Scaling .....	146
Obtaining Information About the Instances Launched by Auto Scaling .....	149
Updating the Bid Price for the Spot Instances .....	153
Scheduling Spot Bid Requests .....	155
Using Auto Scaling to Get Notifications for Spot Instances .....	156
Using CloudFormation Templates to Launch Spot Instances .....	158
Launching Amazon Elastic MapReduce Job Flows with Spot Instances .....	159
Launching Spot Instances in Amazon Virtual Private Cloud .....	160
Advanced Tasks .....	161
Subscribe to Your Spot Instance Data Feed .....	161
Programming Spot with AWS Java SDK .....	164
Tutorial: Amazon EC2 Spot Instances .....	166
Tutorial: Advanced Amazon EC2 Spot Request Management .....	175
Starting Clusters on Spot Instances .....	191
Reserved Instances .....	193
Getting Started with Reserved Instances .....	194
Tools for Working with Reserved Instances .....	197
Reserved Instance Fundamentals .....	199
Choosing Reserved Instances Based on Your Usage Plans .....	199
Understanding Reserved Instance Pricing Tiers .....	200
Understanding the Pricing Benefit of Reserved Instances .....	208
Reserved Instances and Consolidated Billing .....	209
Reserved Instance Marketplace .....	209
Buying Reserved Instances .....	212
Becoming a Buyer .....	213
Purchasing Reserved Instances .....	214
Reading Your Statement (Invoice) .....	221
Obtaining Information About Your Reserved Instances .....	222
Modifying Your Reserved Instances .....	227
Changing the Instance Type of Your Reservations .....	230
Submitting Modification Requests .....	232
Selling in the Reserved Instance Marketplace .....	236
Registering as a Seller .....	237
Selling Your Reserved Instances .....	240
After Your Reserved Instance Is Sold .....	260
Requirements Checklist for Reserved Instances .....	262
Instance Lifecycle .....	263
Launch .....	266
Launching an Instance .....	266
Launching an Instance from a Backup .....	271
Launching an AWS Marketplace Instance .....	271
Connect .....	273
Connect Using MindTerm .....	273
Connect Using PuTTY .....	274
Connect Using SSH .....	279
Connect Using RDP .....	282
Stop and Start .....	284
Reboot .....	286
Terminate .....	287
Instance Metadata and User Data .....	290
Importing and Exporting Instances .....	301
Importing EC2 Instances .....	301
Before You Get Started .....	302
Using the Amazon EC2 VM Import Connector to Import Your Virtual Machine to Amazon EC2 .....	303
Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2 .....	318

Troubleshooting Instance Importation .....	330
Exporting EC2 Instances .....	332
Monitoring Your Instances .....	335
Monitoring Your Instances with CloudWatch .....	336
Monitoring the Status of Your Instances .....	346
Monitoring Instances with Status Checks .....	346
Monitoring Events for Your Instances .....	349
Troubleshooting .....	353
Launching Your Instance .....	353
Connecting to Your Instance .....	355
Stopping Your Instance .....	358
Terminating Your Instance .....	359
Failed Status Checks .....	360
Instance Capacity .....	383
General .....	383
Network and Security .....	385
Key Pairs .....	385
Security Groups .....	392
Controlling Access .....	399
IAM Policies .....	401
IAM Roles .....	408
Network Access .....	412
Amazon VPC .....	414
Supported Platforms .....	417
Instance IP Addressing .....	419
Elastic IP Addresses .....	428
Elastic Network Interfaces .....	431
Storage .....	444
Amazon EBS .....	446
EBS Volumes .....	447
EBS Volume Types .....	448
Creating or Restoring a Volume .....	449
Using Public Data Sets .....	453
Attaching a Volume to an Instance .....	455
Making a Volume Available for Use .....	459
Describing Volumes .....	462
Monitoring the Status of Your Volumes .....	464
Detaching a Volume from an Instance .....	475
Deleting a Volume .....	477
Expanding a Volume .....	479
EBS Snapshots .....	484
Creating a Snapshot .....	485
Deleting a Snapshot .....	487
Copying a Snapshot .....	488
Describing Snapshots .....	491
Sharing Snapshots .....	492
EBS Performance .....	494
EC2 Configuration .....	495
I/O Characteristics .....	497
Workload Demand .....	497
Pre-Warm Volumes .....	498
RAID Configuration .....	501
Benchmark Volumes .....	503
API and Command Overview .....	506
Instance Store .....	508
Amazon S3 .....	517
Block Device Mapping .....	517
Resources and Tags .....	528

Resource Locations .....	528
Listing and Filtering Your Resources .....	529
Tagging Your Resources .....	532
Setting Up the CLI Tools .....	541
Verify the Signature .....	547
Making API Requests .....	552
Query Requests .....	553
Troubleshooting API Request Errors .....	555
Ensuring Idempotency .....	557
SOAP Requests .....	559
Document History .....	560

# What is Amazon EC2?

---

Amazon Elastic Compute Cloud (Amazon EC2) provides resizable computing capacity in the Amazon Web Services (AWS) cloud. Using Amazon EC2 eliminates your need to invest in hardware up front, so you can develop and deploy applications faster. You can use Amazon EC2 to launch as many or as few virtual servers as you need, configure security and networking, and manage storage. Amazon EC2 enables you to scale up or down to handle changes in requirements or spikes in popularity, reducing your need to forecast traffic.

## Features of Amazon EC2

Amazon EC2 provides the following features:

- Virtual computing environments, known as *instances*
- Pre-configured templates for your instances, known as *Amazon Machine Images (AMIs)*, that package the bits you need for your server (including the operating system and additional software)
- Various configurations of CPU, memory, storage, and networking capacity for your instances, known as *instance types*
- Secure login information for your instances using *key pairs* (AWS stores the public key, and you store the private key in a secure place)
- Storage volumes for temporary data that's deleted when you stop or terminate your instance, known as *instance store volumes*
- Persistent storage volumes for your data using Amazon Elastic Block Store (Amazon EBS), known as *Amazon EBS volumes*
- Multiple physical locations for your resources, such as instances and Amazon EBS volumes, known as *regions* and *Availability Zones*
- A firewall that enables you to specify the protocols, ports, and source IP ranges that can reach your instances using *security groups*
- Static IP addresses for dynamic cloud computing, known as *Elastic IP addresses*
- Metadata, known as *tags*, that you can create and assign to your Amazon EC2 resources
- Virtual networks you can create that are logically isolated from the rest of the AWS cloud, and that you can optionally connect to your own network, known as *virtual private clouds (VPCs)*

For more information about the features of Amazon EC2, see the [Amazon EC2 product page](#).

# How to Get Started with Amazon EC2

The first thing you need to do is get set up to use Amazon EC2. After you are set up, you are ready to complete the Getting Started tutorial for Amazon EC2. Whenever you need more information about a feature of Amazon EC2, you can read the technical documentation.

## Getting Started

- [Get Set Up for Amazon EC2 \(p. 18\)](#)
- [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#)
- [Getting Started with Amazon EC2 Windows Instances](#)

## Basics

- [Instances and AMIs \(p. 4\)](#)
- [Instance Types \(p. 94\)](#)
- [Regions and Availability Zones \(p. 7\)](#)
- [Tags \(p. 532\)](#)

## Networking and Security

- [Amazon EC2 Key Pairs \(p. 385\)](#)
- [Security Groups \(p. 392\)](#)
- [Elastic IP Addresses \(EIP\) \(p. 428\)](#)
- [Amazon EC2 and Amazon VPC \(p. 414\)](#)

## Storage

- [Amazon EBS \(p. 446\)](#)
- [Instance Store \(p. 508\)](#)

If you have questions about whether AWS is right for you, [Contact AWS Sales](#). If you have technical questions about Amazon EC2, use the [Amazon EC2 forum](#).

# Related Services in AWS

You can provision Amazon EC2 resources, such as instances and volumes, directly using Amazon EC2. You can also provision Amazon EC2 resources using other services in AWS. For more information, see the following documentation:

- [Auto Scaling Developer Guide](#)
- [AWS CloudFormation User Guide](#)
- [AWS Elastic Beanstalk Developer Guide](#)
- [AWS OpsWorks User Guide](#)

To automatically distribute incoming application traffic across multiple instances, use Elastic Load Balancing. For more information, see [Elastic Load Balancing Developer Guide](#).

To monitor basic statistics for your instances and Amazon EBS volumes, use Amazon CloudWatch. For more information, see [Monitoring Your Instances with CloudWatch \(p. 336\)](#).

To get a managed relational database in the cloud, use Amazon Relational Database Service (Amazon RDS) to launch a database instance. Although you can set up a database on an EC2 instance, Amazon RDS offers the advantage of handling your database management tasks, such as patching the software, backing up, and storing the backups. For more information, see [Amazon Relational Database Service Developer Guide](#).

## Accessing Amazon EC2

Amazon EC2 provides a web-based user interface, the Amazon EC2 console. If you've signed up for an AWS account, you can access the Amazon EC2 console by signing into the AWS Management Console and selecting **EC2** from the console home page.

If you prefer to use a command line interface, there are several options:

### Amazon EC2 Command Line Interface (CLI) Tools

Provide commands for Amazon EC2, Amazon EBS, and Amazon VPC, and is supported on Windows, Mac, and Linux/UNIX. To get started, see [Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX \(p. 541\)](#) or [Installing the Amazon EC2 Command Line Interface Tools on Windows](#). For more information about the commands, see [Commands \(CLI Tools\)](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

### AWS Command Line Interface (CLI)

Provides commands for a broad set of AWS products, and is supported on Windows, Mac, and Linux/UNIX. To get started, see [AWS Command Line Interface User Guide](#). For more information about the commands for Amazon EC2, see [ec2](#).

### AWS Tools for Windows PowerShell

Provides commands for a broad set of AWS products for those who script in the PowerShell environment. To get started, see [AWS Tools for Windows PowerShell User Guide](#).

Amazon EC2 provides a Query API. These requests are HTTP or HTTPS requests that use the HTTP verbs GET or POST and a Query parameter named `Action`. For more information about the API actions for Amazon EC2, see [Actions](#) in the *Amazon Elastic Compute Cloud API Reference*.

If you prefer to build applications using language-specific APIs instead of submitting a request over HTTP or HTTPS, AWS provides libraries, sample code, tutorials, and other resources for software developers. These libraries provide basic functions that automatically take care of tasks such as cryptographically signing your requests, retrying requests, and handling error responses, so that it is easier for you to get started. For more information, see [AWS SDKs and Tools](#).

## Pricing for Amazon EC2

When you sign up for AWS, you can get started with Amazon EC2 for free using the [AWS Free Usage Tier](#).

Amazon EC2 provides the following purchasing options for instances:

### On-Demand Instances

Pay for the instances that you use by the hour, with no long-term commitments or upfront payments.

### Reserved Instances

Make a low, one-time, upfront payment for an instance, reserve it for a one- or three-year term, and pay a significantly lower hourly rate for these instances.

### Spot Instances

Specify the maximum hourly price that you are willing to pay to run a particular instance type. The Spot Price fluctuates based on supply and demand, but you never pay more than the maximum price you specified. If the Spot Price moves higher than your maximum price, Amazon EC2 shuts down your Spot Instances.

For a complete list of charges and specific prices for Amazon EC2, see [Amazon EC2 Pricing](#).

To calculate the cost of a sample provisioned environment, see [AWS Economics Center](#).


To see your bill, go to your [AWS Account Activity page](#). Your bill contains links to usage reports that provide details about your bill. To learn more about AWS account billing, see [AWS Account Billing](#).

If you have questions concerning AWS billing, accounts, and events, [Contact AWS Support](#).

For an overview of Trusted Advisor, a service that helps you optimize the costs, security, and performance of your AWS environment, see [AWS Trusted Advisor](#). You can access Trusted Advisor from the Amazon EC2 console by clicking the **AWS Trusted Advisor** link underneath **Resources**.

## Instances and AMIs

An *Amazon Machine Image (AMI)* is a template that contains a software configuration (for example, an operating system, an application server, and applications). From an AMI, you launch an *instance*, which is a copy of the AMI running as a virtual server in the cloud. You can launch multiple instances of an AMI, as shown in the following figure.



Your instances keep running until you stop or terminate them, or until they fail. If an instance fails, you can launch a new one from the AMI.

## Instances

You can launch different types of instances from a single AMI. An *instance type* essentially determines the hardware of the host computer used for your instance. Each instance type offers different compute and memory capabilities. Select an instance type based on the amount of memory and computing power that you need for the application or software that you plan to run on the instance. For more information about the specifications for each Amazon EC2 instance type, see [Instance Type Details](#).

After you launch an instance, it looks like a traditional host, and you can interact with it as you would any computer. You have complete control of your instances; you can use **sudo** to run commands that require root privileges.

Your AWS account has a limit on the number of instances that you can have running. For more information about this limit, and how to request an increase, see [How many instances can I run in Amazon EC2](#) in the Amazon EC2 General FAQ.

In addition to the limit on running instances, there is a limit on the overall number of instances that you can have (whether running, stopped, or in any other state except for terminated). This overall instance limit is two times your running instance limit.

## Storage for Your Instance

The root device for your instance contains the image used to boot the instance. For more information, see [Amazon EC2 Root Device Volume \(p. 12\)](#).

Your instance may include local storage volumes, known as instance store volumes, which you can configure at launch time with block device mapping. For more information, see [Block Device Mapping \(p. 517\)](#). After these volumes have been added to and mapped on your instance, they are available for you to mount and use. If your instance fails, or if your instance is stopped or terminated, the data on these volumes is lost; therefore, these volumes are best used for temporary data. For important data, you should use a replication strategy across multiple instances in order to keep your data safe, or store your persistent data in Amazon S3 or Amazon EBS volumes. For more information, see [Storage \(p. 444\)](#).

## Security Best Practices

- Use AWS Identity and Access Management (IAM) to control access to your AWS resources, including your instances. You can create IAM users and groups under your AWS account, assign security credentials to each, and control the access that each has to resources and services in AWS. For more information, see [Controlling Access to Amazon EC2 Resources \(p. 399\)](#).
- Restrict access by only allowing trusted hosts or networks to access ports on your instance. For example, you can restrict SSH access by restricting incoming traffic on port 22. For more information, see [Amazon EC2 Security Groups \(p. 392\)](#).
- Review the rules in your security groups regularly, and ensure that you apply the principle of *least privilege*—only open up permissions that you require. You can also create different security groups to deal with instances that have different security requirements. Consider creating a bastion security group that allows external logins, and keep the remainder of your instances in a group that does not allow external logins.
- Disable password-based logins for instances launched from your AMI. Passwords can be found or cracked, and are a security risk. For more information, see [Disable Password-Based Logins for Root \(p. 54\)](#). For more information about sharing AMIs safely, see [Shared AMIs \(p. 49\)](#).

## Stopping, Starting, and Terminating Instances

### Stopping an instance

When an instance is stopped, the instance performs a normal shutdown, and then transitions to a `stopped` state. All of its Amazon EBS volumes remain attached, and you can start the instance again at a later time.

You are not charged for additional instance hours while the instance is in a stopped state. A full instance hour will be charged for every transition from a stopped state to a running state, even if this happens multiple times within a single hour. If the instance type was changed while the instance was stopped, you will be charged the rate for the new instance type after the instance is started. All of the associated Amazon EBS usage of your instance, including root device usage, is billed using typical Amazon EBS prices.



When an instance is in a stopped state, you can attach or detach Amazon EBS volumes. You can also create an AMI from the instance, and you can change the kernel, RAM disk, and instance type.

### Terminating an instance

When an instance is terminated, the instance performs a normal shutdown, then the attached Amazon EBS volumes are deleted unless the volume's `deleteOnTermination` attribute is set to `false`. The instance itself is also deleted, and you can't start the instance again at a later time.

To prevent accidental termination, you can disable instance termination. If you do so, ensure that the `disableApiTermination` attribute is set to `true` for the instance. To control the behavior of an instance shutdown, such as `shutdown -h` in Linux or `shutdown` in Windows, set the `instanceInitiatedShutdownBehavior` instance attribute to `stop` or `terminate` as desired. Instances with Amazon EBS volumes for the root device default to `stop`, and instances with instance-store root devices are always terminated as the result of an instance shutdown.

For more information, see [Instance Lifecycle](#) (p. 263).

## AMIs

Amazon Web Services (AWS) publishes many Amazon Machine Images (AMIs) that contain common software configurations for public use. In addition, members of the AWS developer community have published their own custom AMIs. You can also create your own custom AMI or AMIs; doing so enables you to quickly and easily start new instances that have everything you need. For example, if your application is a website or a web service, your AMI could include a web server, the associated static content, and the code for the dynamic pages. As a result, after you launch an instance from this AMI, your web server starts, and your application is ready to accept requests.

All AMIs are categorized as either *backed by Amazon EBS*, which means that the root device for an instance launched from the AMI is an Amazon EBS volume, or *backed by instance store*, which means that the root device for an instance launched from the AMI is an instance store volume created from a template stored in Amazon S3.

The description of an AMI indicates the type of root device (either `ebs` or `instance store`). This is important because there are significant differences in what you can do with each type of AMI. For more information about these differences, see [Storage for the Root Device](#) (p. 45).

# Regions and Availability Zones

Amazon EC2 is hosted in multiple locations world-wide. These locations are composed of regions and Availability Zones. Each *region* is a separate geographic area. Each region has multiple, isolated locations known as *Availability Zones*. Amazon EC2 provides you the ability to place resources, such as instances, and data in multiple locations. Resources aren't replicated across regions unless you do so specifically.

Amazon operates state-of-the-art, highly-available data centers. Although rare, failures can occur that affect the availability of instances that are in the same location. If you host all your instances in a single location that is affected by such a failure, none of your instances would be available.

## Note


Some AWS resources might not be available in all regions and Availability Zones. Ensure that you can create the resources you need in the desired regions or Availability Zone before deploying your applications.

## Topics

- [Region and Availability Zone Concepts \(p. 7\)](#)
- [Describing Your Regions and Availability Zones \(p. 9\)](#)
- [Specifying the Region for a Resource \(p. 10\)](#)
- [Launching Instances in an Availability Zone \(p. 11\)](#)
- [API and Command Overview \(p. 11\)](#)

## Region and Availability Zone Concepts

Each region is completely independent. Each Availability Zone is isolated, but the Availability Zones in a region are connected through low-latency links. The following diagram illustrates the relationship between regions and Availability Zones.



Amazon EC2 resources are either global, tied to a region, or tied to an Availability Zone. For more information, see [Resource Locations \(p. 528\)](#).

## Regions

Each Amazon EC2 region is designed to be completely isolated from the other Amazon EC2 regions. This achieves the greatest possible fault tolerance and stability.

Amazon EC2 provides multiple regions so that you can launch Amazon EC2 instances in locations that meet your requirements. For example, you might want to launch instances in Europe to be closer to your European customers or to meet legal requirements. The following table lists the regions that provide support for Amazon EC2.

Code	Name
ap-northeast-1	Asia Pacific (Tokyo) Region
ap-southeast-1	Asia Pacific (Singapore) Region
ap-southeast-2	Asia Pacific (Sydney) Region
eu-west-1	EU (Ireland) Region
sa-east-1	South America (Sao Paulo) Region
us-east-1	US East (Northern Virginia) Region
us-west-1	US West (Northern California) Region
us-west-2	US West (Oregon) Region

When you view your resources, you'll only see the resources tied to the region you've specified. This is because regions are isolated from each other, and we don't replicate resources across regions automatically.

When you work with an instance using the command line interface or API actions, you must specify its regional endpoint. For more information about the regions and endpoints for Amazon EC2, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

When you launch an instance, you must select an AMI that's in the same region. If the AMI is in another region, you can copy the AMI to the region you're using. For more information, see [Copying AMIs \(p. 84\)](#).

All communications between regions is across the public Internet. Therefore, you should use the appropriate encryption methods to protect your data. Data transfer between regions is charged at the Internet data transfer rate for both the sending and the receiving instance. For more information, see [Amazon EC2 Pricing - Data Transfer](#).

## Availability Zones

You can list the Availability Zones that are available to your account. For more information, see [Describing Your Regions and Availability Zones \(p. 9\)](#).

When you launch an instance, you can select an Availability Zone or let us choose one for you. If you distribute your instances across multiple Availability Zones and one instance fails, you can design your application so that an instance in another Availability Zone can handle requests.

You can also use Elastic IP addresses to mask the failure of an instance in one Availability Zone by rapidly remapping the address to an instance in another Availability Zone. For more information, see [Elastic IP Addresses \(EIP\) \(p. 428\)](#).

To ensure that resources are distributed across the Availability Zones for a region, we independently map Availability Zones to identifiers for each account. For example, your Availability Zone `us-east-1a` might not be the same location as `us-east-1a` for another account. Note that there's no way for you to coordinate Availability Zones between accounts.

As Availability Zones grow over time, our ability to expand them can become constrained. If this happens, we might restrict you from launching an instance in a constrained Availability Zone unless you already have an instance in that Availability Zone. Eventually, we might also remove the constrained Availability Zone from the list of Availability Zones for new customers. Therefore, your account might have a different number of available Availability Zones in a region than another account.

## Describing Your Regions and Availability Zones

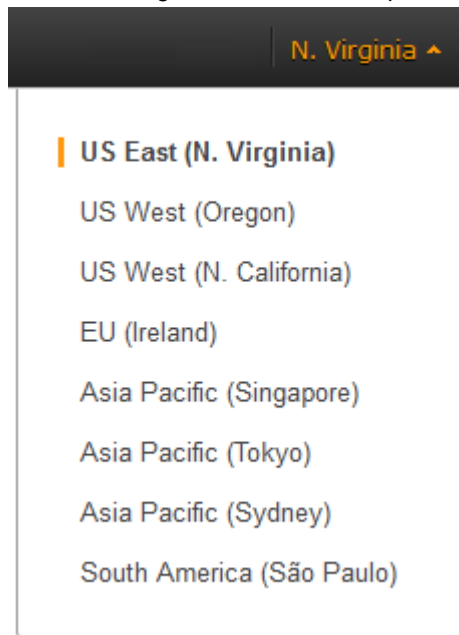
You can use the AWS Management Console or the Amazon EC2 command line interface to determine which regions and Availability Zones are available for your use.

```
PROMPT> ec2-describe-availability-zones --region us-east-1
AVAILABILITYZONE us-east-1a available us-east-1
AVAILABILITYZONE us-east-1b available us-east-1
AVAILABILITYZONE us-east-1c available us-east-1
AVAILABILITYZONE us-east-1d available us-east-1
```

## AWS Management Console

### To find your regions and Availability Zones

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, view the options in the region selector.



3. After you select a region, you can view your Availability Zones within that region when you launch an instance or create an Amazon EBS volume.
  - a. In the navigation pane, click **Volumes**.
  - b. View the options in the **Availability Zones** list.
  - c. When you are finished, click **Cancel**.

## Command Line Interface

Use the following command to describe your regions.

```
PROMPT> ec2-describe-regions
REGION us-east-1 ec2.us-east-1.amazonaws.com
REGION ap-northeast-1 ec2.ap-northeast-1.amazonaws.com
REGION ap-southeast-1 ec2.ap-southeast-1.amazonaws.com
..
```

Use the following command to describe the Availability Zones within the `us-east-1` region.

```
PROMPT> ec2-describe-availability-zones --region us-east-1
AVAILABILITYZONE us-east-1a available us-east-1
AVAILABILITYZONE us-east-1b available us-east-1
AVAILABILITYZONE us-east-1c available us-east-1
AVAILABILITYZONE us-east-1d available us-east-1
```

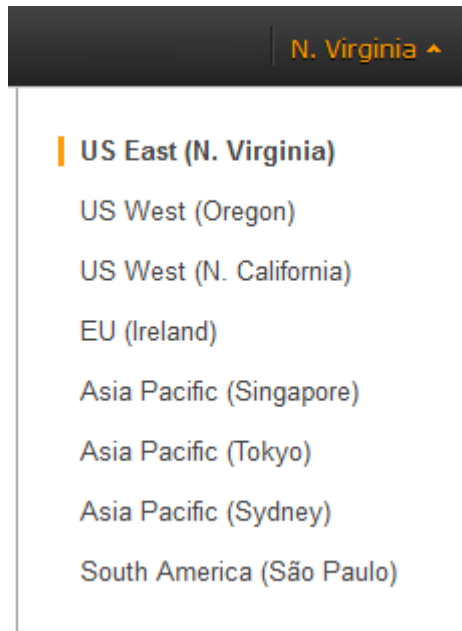
## Specifying the Region for a Resource

Every time you create an Amazon EC2 resource, you can specify the region for the resource. This section explains how to specify the region for a resource.

### AWS Management Console

#### To specify the region for a resource

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Use the region selector in the navigation bar.



### Command Line Interface

To specify the region to use for all commands, set the value of the `EC2_URL` environment variable to the regional endpoint. For example, `https://ec2.us-west-1.amazonaws.com`.

Alternatively, you can use the `--region` or `-U` command line option with each individual command. For example, `--region us-west-1` or `-U https://ec2.us-west-1.amazonaws.com`.

For more information about the endpoints for Amazon EC2, see [Amazon Elastic Compute Cloud Endpoints](#).

## Launching Instances in an Availability Zone

When you launch an instance, select a region that puts your instances closer to specific customers, or meets the legal or other requirements you have. By launching your instances in separate Availability Zones, you can protect your applications from the failure of a single location.

When you launch an instance, you can optionally specify an Availability Zone in the region that you are using. If you do not specify an Availability Zone, we select one for you. When you launch your initial instances, we recommend that you accept the default Availability Zone, because this enables us to select the best Availability Zone for you based on system health and available capacity. If you launch additional instances, only specify an Availability Zone if your new instances must be close to, or separated from, your running instances.

## AWS Management Console

### To specify an Availability Zone for your instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. On the dashboard, click **Launch Instance**.
3. Follow the directions for the wizard to launch the instance. On the **Configure Instance Details** page, you can select one of the Availability Zone options from the list, or select **No Preference** to enable us to select the best Availability Zone for you.

## Command Line Interface

To specify an Availability Zone for your instance, use the `--availability-zone` option with the `ec2-run-instances` command.

```
PROMPT> ec2-run-instances ami_id --availability-zone zone
```

## API and Command Overview

The following table summarizes the available commands and corresponding API actions for regions and Availability Zones.

Description	Command and API Action
Describes the Availability Zones that are available to you.	<code>ec2-describe-availability-zones</code> <code>DescribeAvailabilityZones</code>
Describes the regions that are available to you.	<code>ec2-describe-regions</code> <code>DescribeRegions</code>

# Amazon EC2 Root Device Volume

When you launch an Amazon EC2 instance, the *root device volume* contains the image used to boot the instance. When we introduced Amazon EC2, all AMIs were backed by Amazon EC2 instance store, which means the root device for an instance launched from the AMI is an instance store volume created from a template stored in Amazon S3. After we introduced Amazon EBS, we introduced AMIs that are backed by Amazon EBS. This means that the root device for an instance launched from the AMI is an Amazon EBS volume created from an Amazon EBS snapshot. You can choose between AMIs based by Amazon EC2 instance store and AMIs backed by Amazon EBS. We recommend that you use AMIs backed by Amazon EBS, because they launch faster and use persistent storage.

## Topics

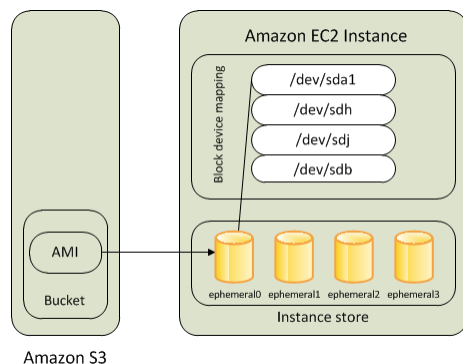
- [Root Device Storage Concepts](#) (p. 12)
- [Choosing an AMI by Root Device Type](#) (p. 14)
- [Displaying the Root Device Type of Your Instance](#) (p. 14)
- [Changing the Root Device Volume to Persist](#) (p. 14)
- [Root Device Storage Usage Scenarios](#) (p. 15)

## Root Device Storage Concepts

You can launch an instance from one of two types of AMIs: an Amazon EC2 instance store-backed AMI or an Amazon EBS-backed AMI. The description of an AMI includes which type of AMI it is; you'll see the root device referred to in some places as either *ebs* (for Amazon EBS-backed) or *instance store* (for Amazon EC2 instance store-backed). This is important because there are significant differences between what you can do with each type of AMI. For more information about these differences, see [Storage for the Root Device](#) (p. 45).

### Instance Store-backed Instances

Instances that use instance stores for the root device automatically have instance store volumes available, with one serving as the root device volume. When an instance is launched, the image that is used to boot the instance is copied to the root volume (typically *sda1*). Any data on the instance store volumes persists as long as the instance is running, but this data is deleted when the instance is terminated (instance store-backed instances do not support the **Stop** action) or if it fails (such as if an underlying drive has issues).




After an instance store-backed instance fails or terminates, it cannot be restored. If you plan to use Amazon EC2 instance store-backed instances, we highly recommend that you distribute the data on your instance stores across multiple Availability Zones. You should also back up the data on your instance store volumes to persistent storage on a regular basis.

For more information, see [Amazon EC2 Instance Store](#) (p. 508).

### Amazon EBS-backed Instances

Instances that use Amazon EBS for the root device automatically have an Amazon EBS volume attached. When you launch an Amazon EBS-backed instance, we create an Amazon EBS volume for each EBS snapshot referenced by the AMI you use. You can optionally use other Amazon EBS volumes or instance store volumes.



An Amazon EBS-backed instance can be stopped and later restarted without affecting data stored in the attached volumes. There are various instance- and volume-related tasks you can do when an Amazon EBS-backed instance is in a stopped state. For example, you can modify the properties of the instance, you can change the size of your instance or update the kernel it is using, or you can attach your root volume to a different running instance for debugging or any other purpose.

By default, the root device volume and the other Amazon EBS volumes attached when you launch an Amazon EBS-backed instance are automatically deleted when the instance terminates. For information about how to change this behavior when you launch an instance, see [Changing the Root Device Volume to Persist](#) (p. 14).

By default, any Amazon EBS volumes that you attach to a running instance are detached with their data intact when the instance terminates. You can attach a detached volume to any running instance.

If an Amazon EBS-backed instance fails, you can restore your session by following one of these methods:

- Stop and then start again.
- Automatically snapshot all relevant volumes and create a new AMI. For more information, see [Creating Amazon EBS-Backed Linux AMIs](#) (p. 63).
- Attach the volume to the new instance by following these steps:
  1. Create a snapshot of the root volume.
  2. Register a new AMI using the snapshot.
  3. Launch a new instance from the new AMI.
  4. Detach the remaining Amazon EBS volumes from the old instance.
  5. Reattach the Amazon EBS volumes to the new instance.

We recommend using either the first or the second method for failed instances with normal volume size and the third method for failed instances with large volumes.



## Choosing an AMI by Root Device Type

The AMI that you specify when you launch your instance determines the type of root device volume that your instance has.

### To choose an EBS-backed AMI

1. Open the EC2 console.
2. In the navigation pane, click **AMIs**.
3. From the filter lists, select the image type (such as **Public images**), the operating system (such as **Amazon Linux**), and **EBS images**.
4. (Optional) To get additional information to help you make your choice, click the **Show/Hide Columns** icon, update the columns to display, and click **Close**.
5. Choose an AMI and write down its AMI ID.

### To choose an instance store-backed AMI

1. Open the EC2 console.
2. In the navigation pane, click **AMIs**.
3. From the filter lists, select the image type (such as **Public images**), the operating system (such as **Amazon Linux**), and **Instance store images**.
4. (Optional) To get additional information to help you make your choice, click the **Show/Hide Columns** icon, update the columns to display, and click **Close**.
5. Choose an AMI and write down its AMI ID.

## Displaying the Root Device Type of Your Instance

### To display the root device type of an instance

1. Open the EC2 console.
2. In the navigation pane, click **Instances**, and select the instance.
3. Check the value of **Root device type** in the details pane as follows:
  - If the value is `ebs`, this is an EBS-backed instance.
  - If the value is `instance store`, this is an instance store-backed instance.

## Changing the Root Device Volume to Persist

By default, the root device volume for an AMI backed by Amazon EBS is deleted when the instance terminates. To change the default behavior, set the `DeleteOnTermination` flag to `false` in the instance's block device mapping.

## AWS Management Console

### To change the root device volume to persist when you launch an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the Amazon EC2 console dashboard, click **Launch Instance**.

3. On the **Choose an Amazon Machine Image (AMI)** page, choose the AMI to use and click **Select**.
4. Follow the wizard to complete the **Choose an Instance Type** and **Configure Instance Details** pages.
5. On the **Add Storage** page, deselect the **Delete On Termination** check box for the Root volume.
6. Complete the remaining wizard pages, and then click **Launch**.

You can verify the setting by viewing details for the root device volume on the instance's details pane in the AWS Management Console. Next to **Block devices**, click the entry for the root device volume. By default, **Delete on termination** is `True`. If you change the default behavior, **Delete on termination** is `False`.

## Command Line Interface

Use the `ec2-run-instances` command, and include a block device mapping that sets the `deleteOnTermination` flag for the root device to `false`. Include the `-v` option to run the command in verbose mode.

```
PROMPT> ec2-run-instances ami_id -b root_device_name::false other parameters... -v
```

The root device is typically `/dev/sda1`, or `xvda` (for Windows). The following is an example.

```
PROMPT> ec2-run-instances ami-1a2b3c4d -b /dev/sda1::false other parameters... -v
```

If you're using the command line tools on a Windows system, you must put quotation marks around the block device mapping value.

```
PROMPT> ec2-run-instances ami-1a2b3c4d -b "xvda::false" other parameters... -v
```

By running the command in verbose mode, you can see the underlying request and response, and confirm that the `deleteOnTermination` value is set to `false`, as shown here.

```
...  
<blockDeviceMapping>  
  <item>  
    <deviceName>/dev/sda1</deviceName>  
    <ebs>  
      <deleteOnTermination>>false</deleteOnTermination>  
    </ebs>  
  </item>  
</blockDeviceMapping>  
...
```

For more information, see [ec2-run-instances](#).

## Root Device Storage Usage Scenarios

You can implement EBS-backed AMIs by creating a set of snapshots and registering an AMI that uses those snapshots. The AMI publisher controls the default size of the root device through the size of the

snapshot. The default size can be increased up to 1TiB to accommodate the requirements of the application either at the time you register the EBS-backed AMI or while you launch the EBS-backed instance.

You cannot decrease the size of your root device to less than the size of the AMI. To decrease the size of your root device, create your own AMI with the desired size for the root device and then launch an instance from that AMI.

### **To launch an EBS-backed instance with increased root device storage disk size**

1. Select an EBS-backed AMI to launch your instance from.
2. Check the root device size and note the AMI ID.
3. Using the command line interface, launch the instance by specifying the AMI ID and the mapping of the root device with the increased size.
4. Connect to the instance.
5. Check the size of the root device on the instance. The increased size of the root device is not apparent yet. This is because the file system does not recognize the increased size on the root device.
6. Resize the file system.
7. Check the size of the root device.
8. The root device of the newly launched instance now shows the increased size.

### **To increase the size of the root device for a running EBS-backed instance**

1. Get the ID of the Amazon EBS volume and the Availability Zone of a running instance for which you want to increase the root storage size.
2. Stop the instance.
3. Detach the original volume from the instance.
4. Create a snapshot of the detached volume.
5. Create a new volume from the snapshot by specifying a larger size.
6. Attach the new volume to the stopped instance.
7. Start the instance and get the new IP address/hostname.
8. Connect to the instance using the new IP address/hostname.
9. Resize the root file system to the extent of the new Amazon EBS volume.
10. Check the size of the root device. The root device now shows the increased size.
11. (Optional) Delete the old Amazon EBS volume, if you no longer need it.

The following are the tasks for creating a snapshot of the root device of an instance store-backed instance. The snapshot is created using an Amazon EBS volume. You can use this snapshot to create a new EBS-backed AMI or to launch another instance.

### **To create a snapshot of the root device of an instance store-backed instance**

1. Launch an instance from an instance store-backed AMI.
2. Create a 10GiB Amazon EBS volume in the same Availability Zone as that of your newly launched instance.

#### **Note**

Use this volume to create a snapshot of the root partition of an instance store-backed AMI. The resulting snapshot is the same size as the root partition; the maximum size of the root partition in an instance store-backed AMI is 10GiB.

3. Attach the volume to the running instance using either the AWS Management Console or the command line tools.
4. Format the volume with a file system.

5. [Linux] Create a directory and then mount the volume on the newly-created directory.
6. Copy the data on the root storage device to the newly-attached volume.
7. Unmount and detach the volume from the instance.
8. Create a snapshot of the volume.

Instance store-backed AMIs are limited to 10GiB storage for the root device. If you require additional storage on your root device, you must first convert the instance store-backed AMI to an EBS-backed AMI and then launch an EBS-backed instance with increased root storage.

**Note**

This conversion procedure works with a Linux AMI, but step 6 fails with a Windows AMI.

**To convert an instance store-backed AMI to an EBS-backed AMI (Linux only)**

1. Launch an instance from an instance store-backed AMI.
2. Create a 10GiB Amazon EBS volume in the same Availability Zone as that of your newly-launched instance.

Use this volume to create a snapshot of the root partition of the instance store-backed AMI. The resulting snapshot is the same size as the root partition; the maximum size of the root partition in an instance store-backed AMI is 10GiB.

3. Attach the volume to the running instance using either the AWS Management Console or the command line interface.
4. Format the volume with a file system.
5. Create a directory and then mount the volume on the newly-created directory.
6. Copy the data on the root storage device to the newly-attached volume.
7. Unmount and detach the volume from the instance.
8. Create a snapshot of the volume.
9. Register the snapshot of the volume as an AMI.

# Get Set Up for Amazon EC2

---

Before you use Amazon EC2 for the first time, complete the following tasks:

1. [Sign Up for AWS](#) (p. 18)
2. [Create a Key Pair](#) (p. 18)
3. [Create a Security Group](#) (p. 20)

Note that if you plan to launch instances in multiple regions, you'll need to create a key pair and a security group in each region. For more information about regions, see [Regions and Availability Zones](#) (p. 7).

## Sign Up for AWS

When you sign up for Amazon Web Services (AWS), your AWS account is automatically signed up for all services in AWS, including Amazon EC2. You are charged only for the services that you use.

With Amazon EC2, you pay only for what you use. If you are a new AWS customer, you can get started with Amazon EC2 for free. For more information, see [AWS Free Usage Tier](#).

If you have an AWS account already, skip to the next task. If you don't have an AWS account, use the following procedure to create one.

### To create an AWS account

1. Go to <http://aws.amazon.com>, and then click **Sign Up**.
2. Follow the on-screen instructions.

Part of the sign-up procedure involves receiving a phone call and entering a PIN using the phone keypad.

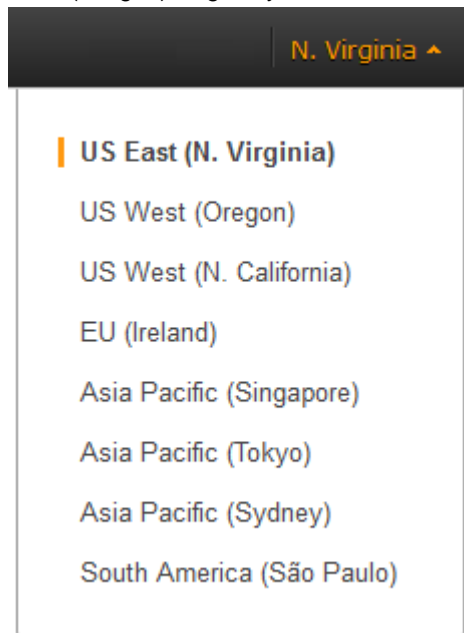
## Create a Key Pair

AWS uses public-key cryptography to secure the login information for your instance. A Linux instance has no password; you use a key pair to log in to your instance securely. You specify the name of the key pair when you launch your instance, then provide the private key when you log in.

If you haven't created a key pair already, you can create one using the Amazon EC2 console.

### To create a key pair

1. Open the Amazon EC2 console.
2. From the navigation bar, select a region for the key pair. You can select any region that's available to you, regardless of your location. This choice is important because some Amazon EC2 resources can be shared between regions, but key pairs can't. For example, if you create a key pair in the US West (Oregon) Region, you can't see or use the key pair in another region.



3. Click **Key Pairs** in the navigation pane.
4. Click **Create Key Pair**.
5. Enter a name for the new key pair in the **Key pair name** field of the **Create Key Pair** dialog box, and then click **Create**. Choose a name that is easy for you to remember, such as your user name, followed by `-key-pair`, plus the region name. For example, `your_user_name-key-pair-region_name`.
6. The private key file is automatically downloaded by your browser. The base file name is the name you specified as the name of your key pair, and the file name extension is `.pem`. Save the private key file in a safe place.

#### Important

This is the only chance for you to save the private key file. You'll need to provide the name of your key pair when you launch an instance and the corresponding private key each time you connect to the instance.

7. If you will use an SSH client on a Linux computer to connect to your Linux instance, use the following command to set the permissions of your private key file so that only you can read it.

```
$ chmod 400 your_user_name-key-pair-region_name.pem
```


For more information, see [Amazon EC2 Key Pairs \(p. 385\)](#).

If you'll connect to your Linux instance from a computer running Linux, you'll specify the `.pem` file to your SSH client. If you'll connect to your Linux instance from a computer running Windows, you can use either


MindTerm or PuTTY. If you plan to use PuTTY, you'll need to install it and use the following procedure to convert the `.pem` file to a `.ppk` file.

**(Optional) To prepare to connect to a Linux instance from Windows using PuTTY**

1. Download and install PuTTY from <http://www.chiark.greenend.org.uk/~sgtatham/putty/>. Be sure to install the entire suite.
2. Start PuTTYgen (for example, from the **Start** menu, click **All Programs > PuTTY > PuTTYgen**).
3. Under **Type of key to generate**, select **SSH-2 RSA**.



4. Click **Load**. By default, PuTTYgen displays only files with the extension `.ppk`. To locate your `.pem` file, select the option to display files of all types.



5. Select the private key file that you created in the previous procedure and click **Open**. Click **OK** to dismiss the confirmation dialog box.
6. Click **Save private key**. PuTTYgen displays a warning about saving the key without a passphrase. Click **Yes**.
7. Specify the same name for the key that you used for the key pair (for example, `your_user_name-key-pair-region_name`). PuTTY automatically adds the `.ppk` file extension.

## Create a Security Group

Security groups act as a firewall for associated instances, controlling both inbound and outbound traffic at the instance level. You must add rules to a security group that enable you to connect to your instance from your IP address using SSH. You can also add rules that allow inbound and outbound HTTP and HTTPS access from anywhere.

**Tip**

You'll need the public IP address of your local computer, which you can get using a service. For example, we provide the following service: <http://checkip.amazonaws.com/>. To locate another service that provides your IP address, use the search phrase "what is my IP address." If you are connecting through an Internet service provider (ISP) or from behind a firewall without a static IP address, you need to find out the range of IP addresses used by client computers.

**To create a security group with least privilege**

1. Open the Amazon EC2 console.
2. Click **Security Groups** in the navigation pane.
3. Click **Create Security Group**.
4. Specify `your_user_name_SG_region_name` as the name of the security group, and provide a description. Click **Yes, Create**.
5. Select the security group that you just created.
6. On the **Inbound** tab, create the following rules, and then click **Apply Rule Changes**:

- Select `HTTP` from the **Create a new rule** list, make sure that **Source** is `0.0.0.0/0`, and then click **Add Rule**.
- Select `HTTPS` from the **Create a new rule** list, make sure that **Source** is `0.0.0.0/0`, and then click **Add Rule**.
- Select `SSH` from the **Create a new rule** list. In the **Source** box, specify the public IP address of your computer or network in CIDR notation, and then click **Add Rule**. To specify an individual IP address in CIDR notation, add the prefix `/32`. For example, if your IP address is `203.0.113.25`, specify `203.0.113.25/32`. If your company allocates addresses from a range, specify the entire range, such as `203.0.113.0/24`.

**Caution**


For security reasons, we don't recommend that you allow SSH access from all IP addresses (`0.0.0.0/0`) to your instance, except for testing purposes and only for a short time.

For more information, see [Amazon EC2 Security Groups \(p. 392\)](#).



# Getting Started with Amazon EC2 Linux Instances

Let's get started with Amazon Elastic Compute Cloud (Amazon EC2) by launching, connecting to, and using a Linux instance. We'll use the AWS Management Console, a point-and-click web-based interface, to complete the example architecture shown in the following diagram:



The instance is an Amazon EBS-backed instance (meaning that the root volume is an Amazon EBS volume). We'll also create and attach an additional Amazon EBS volume. You can either specify the Availability Zone in which your instance runs, or let us select an Availability Zone for you. When you launch your instance, you secure it by specifying a key pair and security group. (This exercise assumes that you created a key pair and a security group when getting set up; see [Get Set Up for Amazon EC2](#).) When you connect to your instance, you must specify the private key of the key pair that you specified when launching your instance.

To complete this exercise, perform the following tasks:

1. [Launch an Amazon EC2 Instance \(p. 23\)](#)
2. [Connect to Your Instance \(p. 24\)](#)
3. [Add a Volume to Your Instance \(p. 27\)](#)
4. [Clean Up Your Instance and Volume \(p. 29\)](#)

## Related Topics

If you'd prefer to launch and connect to a Windows instance, see this tutorial: [Getting Started with Amazon EC2 Windows Instances](#).

For tutorials that show you how to use additional AWS products and services with Amazon EC2, see [Getting Started with AWS](#).

# Launch an Amazon EC2 Instance

You can launch a Linux instance using the AWS Management Console as described in this topic. Before you begin, be sure that you've completed the steps in [Get Set Up for Amazon EC2](#).

## Note

If you'd prefer to launch a Windows instance, see [Getting Started with Amazon EC2 Windows Instances](#).

## Important

When you sign up for AWS, you can get started with Amazon EC2 for free using the [AWS Free Usage Tier](#). If you created your AWS account less than 12 months ago, and have not already exceeded the Free Usage Tier benefits for Amazon EC2 and Amazon EBS, it will not cost you anything to complete this tutorial, because we help you select options that are within the Free Usage Tier benefits. Otherwise, you'll incur the standard Amazon EC2 usage fees from the time that you launch the instance until you terminate the instance (which is the final task of this tutorial), even if it remains idle. The total charges to complete this tutorial outside the Free Usage Tier are minimal (typically only a few dollars).

## To launch an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the console dashboard, click **Launch Instance**.
3. The **Select an Amazon Machine Image (AMI)** page displays a list of basic configurations called Amazon Machine Images (AMIs) that serve as templates for your instance. Select the 64-bit Amazon Linux AMI. Notice that this configuration is marked "Free tier eligible."
4. On the **Select an Instance Type** page, you can select the hardware configuration of your instance. The **t1.micro** instance is selected by default. Click **Review and Launch** to let the wizard complete other configuration settings for you, so you can get started quickly.
5. On the **Review Instance Launch** page, you can review the settings for your instance.

Under **Security Groups**, you'll see that the wizard created and selected a security group for you. Instead, select the security group that you created when getting set up using the following steps:

- a. Click **Edit security groups**.
  - b. On the **Configure Security Group** page, ensure the **Select an existing security group** option is selected.
  - c. Select your security group from the list of existing security groups, and click **Review and Launch**.
6. On the **Review Instance Launch** page, click **Launch**.
  7. In the **Select an existing key pair or create a new key pair** dialog box, select **Choose an existing key pair**, then select the key pair you created when getting set up.

Alternatively, you can create a new key pair. Select **Create a new key pair**, enter a name for the key pair, and then click **Download Key Pair**. This is the only chance for you to save the private key file, so be sure to download it. Save the private key file in a safe place. You'll need to provide the

name of your key pair when you launch an instance and the corresponding private key each time you connect to the instance.

A key pair enables you to connect to a Linux instance through SSH. Therefore, don't select the **Proceed without a key pair** option. If you launch your instance without a key pair, then you can't connect to it.

When you are ready, select the acknowledgment check box, and then click **Launch Instances**.

8. A confirmation page lets you know that your instance is launching. Click **View Instances** to close the confirmation page and return to the console.
9. On the **Instances** screen, you can view the status of your instance. It takes a short time for an instance to launch. When you launch an instance, its initial state is `pending`. After the instance starts, its state changes to `running`, and it receives a public DNS name. (If the **Public DNS** column is hidden, click the **Show/Hide** icon and select **Public DNS**.)

### Next Step

Now that you've launched your instance, you can connect to it and use it. For more information, see [Connect to Your Instance \(p. 24\)](#).

## Connect to Your Instance

After you launch your instance, you can connect to it and use it the way that you'd use a computer sitting in front of you.

If you receive an error while attempting to connect to your instance, see [Troubleshooting Connecting to Your Instance](#).

Before you try to connect to your instance, be sure that you've completed the following tasks:

- **Get the public DNS name of the instance**  
You can get the public DNS for your instance using the Amazon EC2 console (check the **Public DNS** column; if this column is hidden, click the **Show/Hide** icon and select **Public DNS**). If you prefer, you can use the [ec2-describe-instances](#) command.
- **Locate the private key**  
You'll need the fully-qualified path of the `.pem` file for the key pair that you specified when you launched the instance.
- **Enable inbound SSH traffic from your IP address to your instance**  
Ensure that the security group associated with your instance allows incoming SSH traffic from your IP address. For more information, see [Authorizing Network Access to Your Instances](#).

There are several ways to connect to a Linux instance. Choose the method that meets your needs:

- [Option 1: Connect Using Your Browser \(p. 25\)](#)
- [Option 2: Connect from Windows Using PuTTY \(p. 25\)](#)
- [Option 3: Connect from Linux Using an SSH Client \(p. 27\)](#)

### Next Step

After you've successfully launched and connected to your instance, you can do any of the following:

- Continue to the next step in this tutorial, [Add a Volume to Your Instance \(p. 27\)](#).

- Continue using this instance with a different tutorial, such as [Installing a LAMP Web Server](#) or [Hosting a WordPress Blog](#).
- Skip to the last step in this tutorial, [Clean Up Your Instance and Volume \(p. 29\)](#), to terminate the instance so that you don't continue to incur charges.

## Option 1: Connect Using Your Browser

You must have Java installed and enabled in the browser. If you don't have Java already, you can contact your system administrator to get it installed, or follow the steps outlined in the following pages: [Install Java](#) and [Enable Java in your web browser](#).

### To connect to your Linux instance using a web browser

1. From the Amazon EC2 console, click **Instances** in the navigation pane.
2. Select the instance, and then click **Connect**.
3. Click **A Java SSH client directly from my browser (Java required)**.
4. Amazon EC2 automatically detects the public DNS name of your instance and populates **Public DNS** for you. It also detects the key pair that you specified when you launched the instance. Complete the following, and then click **Launch SSH Client**.
  - a. In **User name**, enter `ec2-user`.

**Tip**  
For Amazon Linux, the user name is `ec2-user`. For RHEL5, the user name is often `root` but might be `ec2-user`. For an Ubuntu, AMI the user name is `ubuntu`. Otherwise, check with your AMI provider.
  - b. In **Private key path**, enter the fully qualified path to your private key (`.pem`) file.
  - c. Click **Store in browser cache** to store the location of the private key in your browser cache. This enables Amazon EC2 to detect the location of the private key in subsequent browser sessions, until you clear your browser's cache.
5. When prompted to add the host to your set of known hosts, click **No**.
6. If necessary, click **Yes** to trust the certificate.
7. Click **Run** to run the MindTerm client.
8. If you accept the license agreement, click **Accept**.
9. If this is your first time running MindTerm, a series of dialog boxes asks you to confirm setup for your home directory and other settings. Confirm these settings. A window opens and you are connected to your instance.

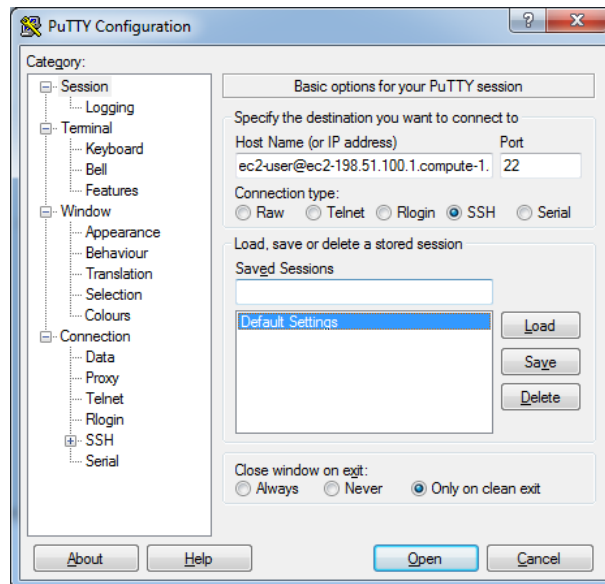
## Option 2: Connect from Windows Using PuTTY

PuTTY doesn't use `.pem` files, it uses `.ppk` files. If you haven't already generated a `.ppk` file, do so now. For more information, see [To prepare to connect to a Linux instance from Windows using PuTTY](#).

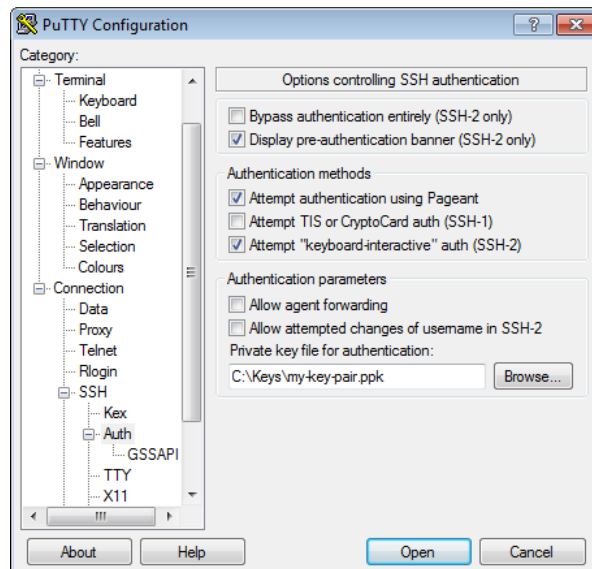
### To connect to your Linux instance using PuTTY

1. Start PuTTY (from the **Start** menu, click **All Programs > PuTTY > PuTTY**).
2. In the Category pane, select **Session** and complete the following fields:
  - a. In the **Host Name** box, enter `ec2-user@public_dns_name`.
  - b. Under **Connection type**, select **SSH**.

- c. Ensure that **Port** is 22.



3. In the **Category** pane, expand **Connection**, expand **SSH**, and then select **Auth**. Complete the following:
  - a. Click **Browse**.
  - b. Select the `.ppk` file that you generated for your key pair, and then click **Open**.
  - c. Click **Open** to start the PuTTY session.



4. If this is the first time you have connected to this instance, PuTTY displays a security alert dialog box that asks whether you trust the host you are connecting to. Click **Yes**. A window opens and you are connected to your instance.

## Option 3: Connect from Linux Using an SSH Client

Your Linux computer most likely includes an SSH client by default. You can check for an SSH client by typing **ssh** at the command line. If your computer doesn't recognize the command, the OpenSSH project provides a free implementation of the full suite of SSH tools. For more information, see <http://www.openssh.org>.

Open your command shell and run the following command:

```
ssh -i /path/key_pair.pem ec2-user@public_dns_name
```

### Tip

For Amazon Linux, the user name is `ec2-user`. For RHEL5, the user name is often `root` but might be `ec2-user`. For an Ubuntu, AMI the user name is `ubuntu`. Otherwise, check with your AMI provider.

## Add a Volume to Your Instance

Now that you've launched and connected to your Linux instance, you can run the following command on your instance to view its mounted volumes.

```
$ df -h
```

For a micro instance, your output should look something like this.

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/xvda1	8.0G	1.1G	6.9G	14%	/
tmpfs	298M	0	298M	0%	/dev/shm

The `/dev/xvda1` volume is the root device volume. It contains the image used to boot the instance. Notice that there's some room to install additional software on your instance. For example, you can use the **yum** command to download and install packages.

If you need additional storage for your data, a simple solution is to add Amazon EBS volumes to your instance. An Amazon EBS volume serves as network-attached storage for your instance. Let's add a volume to the Linux instance that you've launched. First we'll use the EC2 console to create the volume and attach it to the instance, and then we'll mount the volume to make it available.

### To create and attach an Amazon EBS volume

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation bar, select the region in which you created your instance, and then click **Instances** in the navigation pane.

The console displays the list of current instances in that region. Select your Linux instance. In the **Description** tab in the bottom pane note the **Availability Zone** for the instance.

3. In the navigation pane, under **Elastic Block Store**, click **Snapshots**. Select **Public Snapshots** from the **Viewing** list. Select a snapshot from the list and note its snapshot ID. The Free Usage Tier provides up to 30 GB of Amazon Elastic Block Storage; therefore, to avoid being charged for this tutorial, choose a snapshot that is smaller than 30 GB.

Note that this tutorial assumes that you create the volume using a snapshot as described in this step. If you create an empty volume instead, we'll ask you to perform an additional step in the next procedure.

4. Click **Create Volume**.
5. The **Create Volume** dialog box is preconfigured with the snapshot ID and volume size of the snapshot you selected. Configure the following, and then click **Yes, Create**:
  - Select the `Standard` volume type to create a standard EBS volume.
  - Select the same **Availability Zone** that you used when you created your instance. Otherwise, you can't attach the volume to your instance.
6. In the navigation pane, under **Elastic Block Store**, click **Volumes**. Notice that your newly created volume appears there and the state of the volume is `available`, so it's ready to be attached to an instance.
7. Right-click the newly created volume and select **Attach Volume**.
8. In the **Attach Volume** dialog box, configure the following, and then click **Yes, Attach**:
  - Select your Linux instance from the list.
  - Specify an unused device name for that instance. We'll use `/dev/sdf` in this tutorial. If you select a different device name, be sure to note it as you'll need this information in the next procedure.

You'll notice that in the **Details** pane for your volume, the state of the volume is `in-use`, and the volume is attached to your instance with the device name `/dev/sdf`. However, if you return to your instance and run the `df -h` command again, you won't see the volume yet. That's because we need to mount the volume to make it available.

### To make a volume available

1. To mount the volume as `/mnt/my-data`, run the following commands.

```
$ sudo mkdir /mnt/my-data
$ sudo mount /dev/sdf /mnt/my-data
```

If you attached the volume using a device name other than `/dev/sdf`, be sure to specify that device name. Otherwise, you might receive the following error when you run this mount command: "mount: you must specify the filesystem type".

Note that if you created an empty volume instead of creating a volume from a snapshot in the previous procedure, you'll need to format the volume using `mkfs` before you can mount it. Do not use `mkfs` if you created the volume from a snapshot, as this will delete the public data set. For more information, see [Making the Volume Available on Linux](#).

2. Now when you run the `df -h` command, you'll see output like the following.

```
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      8.6G  1.2G  7.4G  14% /
tmpfs           313M    0  313M   0% /dev/shm
/dev/xvdf       5.0G  4.3G  442M  91% /mnt/my-data
```

3. To view the contents of the new volume, run the following command.

```
$ ls /mnt/my-data
```

At this point, you have completed the example architecture for this tutorial. You can continue to customize and use your instance for as long as you wish.

### Important

Remember, if you launched an instance in the Free Usage Tier, there are no charges. Otherwise, as soon as your instance starts to boot, you're billed for each hour or partial hour that you keep the instance running, even if the instance is idle. You'll stop incurring charges for a regular instance as soon as the instance status changes to `shutting down` or `terminated`.

When you're finished with your instance, don't forget to clean up any resources you've used and terminate the instance, as shown in the next step, [Clean Up Your Instance and Volume](#) (p. 29).

## Clean Up Your Instance and Volume

After you've finished with the instance and the Amazon EBS volume that you created for this tutorial, you should clean up. First, terminate the instance, which detaches the volume from the instance, and then delete the volume.

Terminating an instance effectively deletes it because you can't reconnect to an instance after you've terminated it. This differs from stopping the instance; when you stop an instance, it is shut down and you are not billed for hourly usage or data transfer (but you are billed for any Amazon EBS volume storage). Also, you can restart a stopped instance at any time. For more information about the differences between stopping and terminating an instance, see [Stopping Instances](#).

### To terminate the instance

1. Locate your instance in the list of instances on the **Instances** page. If you can't find your instance, verify that you have selected the correct region.
2. Right-click the instance, and then click **Terminate**.
3. Click **Yes, Terminate** when prompted for confirmation.

EBS volumes can persist even after your instance is terminated. If you created and attached an EBS volume in the previous step, it was detached when you terminated the instance. However, you must delete the volume, or you'll be charged for volume storage if the storage amount exceeds the limit of the Free Usage Tier. After you delete a volume, its data is gone and the volume can't be attached to any instance.

### To delete the volume

1. Locate the volume that you created in the list of volumes on the **Volumes** page. If you can't find your volume, verify that you have selected the correct region.
2. Right-click the volume, and then click **Delete**.
3. Click **Yes, Delete** when prompted for confirmation. Amazon EC2 begins deleting the volume.



# Tutorial: Installing a LAMP Web Server

---

The following procedures help you install the Apache web server with PHP and MySQL support on your Amazon EC2 instance (sometimes called a LAMP web server or LAMP stack). You can use this server to host a static website or deploy a dynamic PHP application that reads and writes information to a database. These instructions are intended for use with the Amazon Linux AMI, but the commands and file locations are similar for Red Hat and CentOS AMIs. For more information about other distributions, see their specific documentation.

## Prerequisites

This tutorial assumes that you have already launched an instance with a public DNS name that is reachable from the Internet. For more information, see [Launch an Amazon EC2 Instance \(p. 23\)](#). You must also have configured your security group to allow SSH (port 22), HTTP (port 80), and HTTPS (port 443) connections. For more information about these prerequisites, see [Get Set Up for Amazon EC2 \(p. 18\)](#).

## To install and start the LAMP web server

1. [Connect to your instance \(p. 24\)](#).
2. To ensure that all of your software packages are up to date, perform a quick software update on your instance. This process may take a few minutes, but it is important to make sure you have the latest security updates and bug fixes.

### Note

The `-y` option installs the updates without asking for confirmation. If you would like to examine the updates before installing, you can omit this option.

```
[ec2-user ~]$ sudo yum update -y
```

3. Now that your instance is current, you can install the Apache web server, MySQL, and PHP software packages. Use the `yum groupinstall` command to install multiple software packages and all related dependencies at the same time.

```
[ec2-user ~]$ sudo yum groupinstall -y "Web Server" "MySQL Database" "PHP Support"
```

**Note**

Non-Amazon Linux AMIs may have subtle differences in their group names. If the above command fails because of an invalid group name, use the **yum grouplist** command and scan the output for similar groups, such as "MySQL Database server" instead of "MySQL Database", and use the appropriate group name for your distribution.

4. Install the `php-mysql` package.

```
[ec2-user ~]$ sudo yum install -y php-mysql
```

5. Start the Apache web server.

```
[ec2-user ~]$ sudo service httpd start
Starting httpd: [ OK ]
```

6. Use the **chkconfig** command to configure the Apache web server to start at each system boot.

```
[ec2-user ~]$ sudo chkconfig httpd on
```

**Tip**

The **chkconfig** command does not provide any confirmation message when you successfully enable a service. You can verify that **httpd** is on by running the following command.

```
[ec2-user ~]$ chkconfig --list httpd
httpd          0:off  1:off  2:on   3:on   4:on   5:on   6:off
```

Here, **httpd** is on in runlevels 2, 3, 4, and 5 (which is what you want to see).

7. Test your web server. In a web browser, enter the public DNS address (or the public IP address) of your instance; you should see the Apache test page. You can get the public DNS for your instance using the Amazon EC2 console (check the **Public DNS** column; if this column is hidden, click the **Show/Hide** icon and select **Public DNS**).

**Tip**

If you are unable to see the Apache test page, check that the security group you are using contains a rule to allow HTTP (port 80) traffic. For information about adding an HTTP rule to your security group, see [Adding Rules to a Security Group \(p. 396\)](#).

**Important**

If you are not using an Amazon Linux AMI, you may also need to configure the firewall on your instance to allow these connections. For more information about how to configure the firewall, see the documentation for your specific distribution.

## Amazon Linux AMI Test Page

This page is used to test the proper operation of the Apache HTTP server after it has been installed. If you can read this page, it means that the Apache HTTP server installed at this site is working properly.

**If you are a member of the general public:**

The fact that you are seeing this page indicates that the website you just visited is either experiencing problems, or is undergoing routine maintenance.


If you would like to let the administrators of this website know that you've seen this page instead of the page you expected, you should send them e-mail. In general, mail sent to the name "webmaster" and directed to the website's domain should reach the appropriate person.

For example, if you experienced problems while visiting `www.example.com`, you should send e-mail to "webmaster@example.com".

**If you are the website administrator:**

You may now add content to the directory `/var/www/html/`. Note that until you do so, people visiting your website will see this page, and not your content. To prevent this page from ever being used, follow the instructions in the file `/etc/httpd/conf.d/welcome.conf`.

You are free to use the image below on web sites powered by the Apache HTTP Server:



### Note

This test page only appears when there is no content in `/var/www/html`. When you add content to the document root, your content appears at the public DNS address of your instance instead of this test page.

Apache **httpd** serves files that are kept in a directory called the Apache document root. The Amazon Linux AMI Apache document root is `/var/www/html`, which is owned by `root` by default.

```
[ec2-user ~]$ ls -l /var/www
total 16
drwxr-xr-x 2 root root 4096 Jul 12 01:00 cgi-bin
drwxr-xr-x 3 root root 4096 Aug 7 00:02 error
drwxr-xr-x 2 root root 4096 Jan 6 2012 html
drwxr-xr-x 3 root root 4096 Aug 7 00:02 icons
```

To allow `ec2-user` to manipulate files in this directory, you need to modify the ownership and permissions of the directory. There are many ways to accomplish this task; in this tutorial, you add a `www` group to your instance, and you give that group ownership of the `/var/www` directory and add write permissions for the group. Any members of that group will then be able to add, delete, and modify files for the web server.

### To set file permissions

1. Add the `www` group to your instance.

```
[ec2-user ~]$ sudo groupadd www
```

2. Add your user (in this case, `ec2-user`) to the `www` group.

```
[ec2-user ~]$ sudo usermod -a -G www ec2-user
```

**Important**

You need to log out and log back in to pick up the new group. You can use the **exit** command, or close the terminal window.

3. Log out and then log back in again and verify your membership in the `www` group.
  - a. Log out.

```
[ec2-user ~]$ exit
```

- b. Reconnect to your instance and then run the following command to verify your membership in the `www` group.

```
[ec2-user ~]$ groups
ec2-user wheel www
```

4. Change the group ownership of `/var/www` and its contents to the `www` group.

```
[ec2-user ~]$ sudo chown -R root:www /var/www
```

5. Change the directory permissions of `/var/www` and its subdirectories to add group write permissions and to set the group ID on future subdirectories.

```
[ec2-user ~]$ sudo chmod 2775 /var/www
[ec2-user ~]$ find /var/www -type d -exec sudo chmod 2775 {} +
```

6. Recursively change the file permissions of `/var/www` and its subdirectories to add group write permissions.

```
[ec2-user ~]$ find /var/www -type f -exec sudo chmod 0664 {} +
```

Now `ec2_user` (and any future members of the `www` group) can add, delete, and edit files in the Apache document root. Now you are ready to add content, such as a static website or a PHP application.

**To test your LAMP web server**

If your server is installed and running, and your file permissions are set correctly, your `ec2-user` account should be able to create a simple PHP file in the `/var/www/html` directory that will be available from the Internet.

1. Create a simple PHP file in the Apache document root.

```
[ec2-user ~]$ echo "<?php phpinfo(); ?>" > /var/www/html/phpinfo.php
```

**Tip**

If you get a "Permission denied" error when trying to run this command, try logging out and logging back in again to pick up the proper group permissions that you configured in [To set file permissions \(p. 32\)](#).

- In a web browser, enter the URL of the file you just created. This URL is the public DNS address of your instance followed by a forward slash and the filename. For example:

```
http://my.public.dns.amazonaws.com/phpinfo.php
```

You should see the PHP information page.



- Delete the `phpinfo.php` file. Although this can be very useful information to you, it should not be broadcast to the Internet for security reasons.

```
[ec2-user ~]$ rm /var/www/html/phpinfo.php
```

### To secure the MySQL server

The default installation of the MySQL server has several features that are great for testing and development, but they should be disabled or removed for production servers. The `mysql_secure_installation` command walks you through the process of setting a root password and removing the insecure features from your installation. Even if you are not planning on using the MySQL server, performing this procedure is a good idea.

- Start the MySQL server so that you can run `mysql_secure_installation`.

```
[ec2-user ~]$ sudo service mysqld start
Initializing MySQL database: Installing MySQL system tables...
OK
Filling help tables...
OK

To start mysqld at boot time you have to copy
support-files/mysql.server to the right place for your system

PLEASE REMEMBER TO SET A PASSWORD FOR THE MySQL root USER !
...

Starting mysqld: [ OK ]
```

2. Run `mysql_secure_installation`.

```
[ec2-user ~]$ sudo mysql_secure_installation
```

- a. When prompted, enter a password for the `root` account.
  - i. Enter the current `root` password. By default, the `root` account does not have a password set, so press **Enter**.
  - ii. Type **Y** to set a password, and enter a secure password twice. For more information about creating a secure password, go to <http://www.pctools.com/guides/password/>. Make sure to store this password in a safe place.
- b. Type **Y** to remove the anonymous user accounts.
- c. Type **Y** to disable remote `root` login.
- d. Type **Y** to remove the test database.
- e. Type **Y** to reload the privilege tables and save your changes.

3. (Optional) Stop the MySQL server if you do not plan to use it right away. You can restart the server when you need it again.

```
[ec2-user ~]$ sudo service mysqld stop  
Stopping mysqld: [ OK ]
```

You should now have a fully functional LAMP web server. If you add content to the Apache document root at `/var/www/html`, you should be able to view that content at the public DNS address for your instance.

### Related Topics

For more information on transferring files to your instance or installing a WordPress blog on your web server, see the following topics:

- [Transferring Files to Your Instance with WinSCP \(p. 278\)](#)
- [Transferring Files to Linux/UNIX Instances from Linux/UNIX with SCP \(p. 280\)](#)
- [Tutorial: Hosting a WordPress Blog with Amazon EC2 \(p. 36\)](#)

For more information about the Apache web server, go to <http://httpd.apache.org/>. For more information about the MySQL database server, go to <http://www.mysql.com/>. For more information about the PHP programming language, go to <http://php.net/>.

# Tutorial: Hosting a WordPress Blog with Amazon EC2

---

The following procedures will help you install, configure, and secure a WordPress blog on your Amazon EC2 instance. This tutorial is intended for use with the Amazon Linux AMI, but the commands and file locations are similar for Red Hat and CentOS AMIs. For more information about other distributions, see their specific documentation.

This tutorial is a good introduction to using Amazon EC2 in that you have full control over a web server that hosts your WordPress blog, which is not typical with a traditional hosting service. Of course, that means that you are responsible for updating the software packages and maintaining security patches for your server as well. For a more automated WordPress installation that does not require direct interaction with the web server configuration, the AWS CloudFormation service provides a WordPress template that can also get you started quickly. For more information, see [Get Started](#) in the *AWS CloudFormation User Guide*. If you'd prefer to host your WordPress blog on a Windows instance, see [Deploying a WordPress Blog on Your Amazon EC2 Windows Instance](#) in the *Amazon Elastic Compute Cloud Microsoft Windows Guide*.

## Prerequisites

This tutorial assumes that you have launched an instance with a functional web server with PHP and MySQL support. Your Amazon EC2 security group should also allow HTTP and HTTPS traffic. If you do not already have a functional web server, see [Tutorial: Installing a LAMP Web Server \(p. 30\)](#) to create one and then return to this tutorial to install WordPress. For information about adding rules to your security group, see [Adding Rules to a Security Group \(p. 396\)](#).

## To download and unzip the WordPress installation package

1. Download the latest WordPress installation package with the `wget` command. The command below should always download the latest release.

```
[ec2-user ~]$ wget https://wordpress.org/latest.tar.gz
--2013-08-09 17:19:01-- https://wordpress.org/latest.tar.gz
Resolving wordpress.org (wordpress.org)... 66.155.40.249, 66.155.40.250
Connecting to wordpress.org (wordpress.org)|66.155.40.249|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 4028740 (3.8M) [application/x-gzip]
Saving to: latest.tar.gz
```

```
100%[=====>] 4,028,740 20.1MB/s in 0.2s
2013-08-09 17:19:02 (20.1 MB/s) - latest.tar.gz saved [4028740/4028740]
```

2. Unzip and unarchive the installation package. The installation folder is unzipped to a folder called `wordpress`.

```
[ec2-user ~]$ tar -xzf latest.tar.gz
[ec2-user ~]$ ls
latest.tar.gz  wordpress
```

### To create a MySQL user and database for your WordPress installation

Your WordPress installation needs to store information, such as blog post entries and user comments, in a database. This procedure will help you create a database for your blog and a user that is authorized to read and save information to that database.

1. Log in to the MySQL client as the `root` user. Enter your MySQL `root` password when prompted; this may be different than your `root` UNIX password or it may even be empty if you have not secured your MySQL server.

#### Important

If you have not secured your MySQL server yet, it is very important that you do so. For more information, see [To secure the MySQL server \(p. 34\)](#).

```
[ec2-user ~]$ mysql -u root -p
Enter password:
```

2. Create a user and password for your MySQL database. Your WordPress installation uses these values to communicate with your MySQL database. Enter the following command, substituting a unique user name and password.

```
mysql> CREATE USER 'wordpress-user'@'localhost' IDENTIFIED BY
'your_strong_password';
Query OK, 0 rows affected (0.00 sec)
```

Make sure that you create a strong password for your user. Do not use the single quote character ( `'` ) in your password, because this will break the above command. For more information about creating a secure password, go to <http://www.pctools.com/guides/password/>. Do not reuse an existing password, and make sure to store this password in a safe place.

3. Create your database. Give your database a descriptive, meaningful name, such as `wordpress-db`.

#### Note

The punctuation marks surrounding the database name in the command below are called backticks. The backtick ( ``` ) key is usually located above the **Tab** key on a standard keyboard. Backticks are not always required, but they allow you to use otherwise illegal characters, such as hyphens, in database names.

```
mysql> CREATE DATABASE `wordpress-db`;
Query OK, 1 row affected (0.01 sec)
```



- Grant full privileges for your database to the WordPress user you created earlier.

```
mysql> GRANT ALL PRIVILEGES ON `wordpress-db`.* TO "wordpress-user"@"local
host";
Query OK, 0 rows affected (0.00 sec)
```

- Flush the MySQL privileges to pick up all of your changes.

```
mysql> FLUSH PRIVILEGES;
Query OK, 0 rows affected (0.01 sec)
```

- Exit the `mysql` client.

```
mysql> exit
Bye
```

### To create and edit the `wp-config.php` file

The WordPress installation folder contains a sample configuration file called `wp-config-sample.php`. In this procedure, you copy this file and edit it to fit your specific configuration.

- Copy the `wp-config-sample.php` file to a file called `wp-config.php`. This creates a new configuration file and keeps the original sample file intact as a backup.

```
[ec2-user ~]$ cd wordpress/
[ec2-user wordpress]$ cp wp-config-sample.php wp-config.php
```

- Edit the `wp-config.php` file with your favorite text editor (such as **nano** or **vim**) and enter values for your installation. If you do not have a favorite text editor, `nano` is much easier for beginners to use.

```
[ec2-user wordpress]$ nano wp-config.php
```

- Find the line that defines `DB_NAME` and change `database_name_here` to the database name you created in [Step 3 \(p. 37\)](#) of [To create a MySQL user and database for your WordPress installation \(p. 37\)](#).

```
define('DB_NAME', 'wordpress-db');
```

- Find the line that defines `DB_USER` and change `username_here` to the database user you created in [Step 2 \(p. 37\)](#) of [To create a MySQL user and database for your WordPress installation \(p. 37\)](#).

```
define('DB_USER', 'wordpress-user');
```

- Find the line that defines `DB_PASSWORD` and change `password_here` to the strong password you created in [Step 2 \(p. 37\)](#) of [To create a MySQL user and database for your WordPress installation \(p. 37\)](#).

```
define('DB_PASSWORD', 'your_strong_password');
```

- d. Find the section called Authentication Unique Keys and Salts. These KEY and SALT values provide a layer of encryption to the browser cookies that WordPress users store on their local machines. Basically, adding long, random values here makes your site more secure. Visit <https://api.wordpress.org/secret-key/1.1/salt/> to randomly generate a set of key values that you can copy and paste into your `wp-config.php` file. To paste text into a PuTTY terminal, place the cursor where you want to paste the text and right-click your mouse inside the PuTTY terminal.

For more information about security keys, go to [http://codex.wordpress.org/Editing\\_wp-config.php#Security\\_Keys](http://codex.wordpress.org/Editing_wp-config.php#Security_Keys).

#### Note

The values below are for example purposes only; do not use these values for your installation.

```
define('AUTH_KEY',          ' #U$$+[RXN8:b^-L 0(WU_+ c+WFkI~c]o]-
bHw+)/Aj[wTwSiZ<Qb[mghEXcRh-');
define('SECURE_AUTH_KEY',  ' Zsz._P=1/|y.Lq)XjlkWslY5NJ76E6EJ.AV0pCK
ZZB,*~*r ?6OP$eJT@;+(ndLg');
define('LOGGED_IN_KEY',    ' ju}qwre3V*+8f_zOWF?{LlGsQ]Ye@2Jh^,8x>)Y
|;(^[Iw]Pi+LG#A4R?7N`YB3');
define('NONCE_KEY',       ' P(g62HeZxEes|LnI^i=H,[XwK9I&[2s|: ?ON}VJM%?;v2v]v+;+^9eXUahg@::Cj');
define('AUTH_SALT',       ' C$DpB4Hj[JK: ?{q1`SRVa: { :7yShy(9A@5wg+`JJVb1fk%_-Bx*M4(qc[Qg%JT!h');
define('SECURE_AUTH_SALT', ' d!uRu#}+q#{f$Z?Z9uFPG.${+S{n~1M&%@~gL>U>NV<zpD-@2-Es7Q10-bp28EKv');
define('LOGGED_IN_SALT',  ' ;j{00P*owZf)kVD+FVLn~~
>./Y%Ug4#I^*LVd9QeZ^&XmK|e(76miC+&W&+^0P/');
define('NONCE_SALT',     ' -97r*v/cgxLmp?Zy4zUU4r99QQ_rGs2LTd%P;|_e1tS)8_B/, .6[=UK<J_y9?JWG');
```

- e. Save the file and exit your text editor.

## To move your WordPress installation to the Apache document root

Now that you've unzipped the installation folder, created a MySQL database and user, and customized the WordPress configuration file, you are ready to move your installation files to your web server document root so you can run the installation script that completes your installation. The location of these files depends on whether you want your WordPress blog to be available at the root of your web server (for example, [my.public.dns.amazonaws.com](http://my.public.dns.amazonaws.com)) or in a subdirectory or folder (for example, [my.public.dns.amazonaws.com/blog](http://my.public.dns.amazonaws.com/blog)).

### Important

Choose the location where you want your blog to be available and only run the `mv` associated with that location. If you run both sets of commands below, you will get an error message on the second `mv` command because the files you are trying to move are no longer there.

1. To make your blog available at [my.public.dns.amazonaws.com](http://my.public.dns.amazonaws.com), move the files in the `wordpress` folder (but not the folder itself) to the Apache document root (`/var/www/html` on Amazon Linux AMIs).

```
[ec2-user wordpress]$ mv * /var/www/html/
```

2. OR, to make your blog available at `my.public.dns.amazonaws.com/blog` instead, create a new folder called `blog` inside the Apache document root and move the files in the `wordpress` folder (but not the folder itself) to the new `blog` folder.

```
[ec2-user wordpress]$ mkdir /var/www/html/blog
[ec2-user wordpress]$ mv * /var/www/html/blog
```

### Important

If you are not moving on to the next procedure immediately, stop the Apache web server (`httpd`) now for security purposes. After you move your installation to the Apache document root, the WordPress installation script is unprotected and an attacker could gain access to your blog if the Apache web server were running. To stop the Apache web server, enter the **`sudo service httpd stop`** command. If you are moving on to the next procedure, you do not need to stop the Apache web server.

### To run the WordPress installation script

1. Use the **`chkconfig`** command to ensure that the `httpd` and `mysqld` services start at every system boot.

```
[ec2-user wordpress]$ sudo chkconfig httpd on
[ec2-user wordpress]$ sudo chkconfig mysqld on
```

2. Verify that the MySQL server (`mysqld`) is running.

```
[ec2-user wordpress]$ sudo service mysqld status
mysqld (pid 4746) is running...
```

If the `mysqld` service is not running, start it.

```
[ec2-user wordpress]$ sudo service mysqld start
Starting mysqld: [ OK ]
```

3. Verify that your Apache web server (`httpd`) is running.

```
[ec2-user wordpress]$ sudo service httpd status
httpd (pid 502) is running...
```

If the `httpd` service is not running, start it.

```
[ec2-user wordpress]$ sudo service httpd start
Starting httpd: [ OK ]
```

4. Verify that the `php` and `php-mysql` packages are installed. Your output may look slightly different, but look for the `Installed Packages` section.

```
[ec2-user wordpress]$ yum list installed php php-mysql
Loaded plugins: priorities, security, update-motd, upgrade-helper
amzn-main | 2.1 kB | 00:00
amzn-updates | 2.3 kB | 00:00
```

```

Installed Packages
php.x86_64                    5.3.27-1.0.amzn1      @amzn-
updates
php-mysql.x86_64            5.3.27-1.0.amzn1      @amzn-
updates
    
```

**Note**

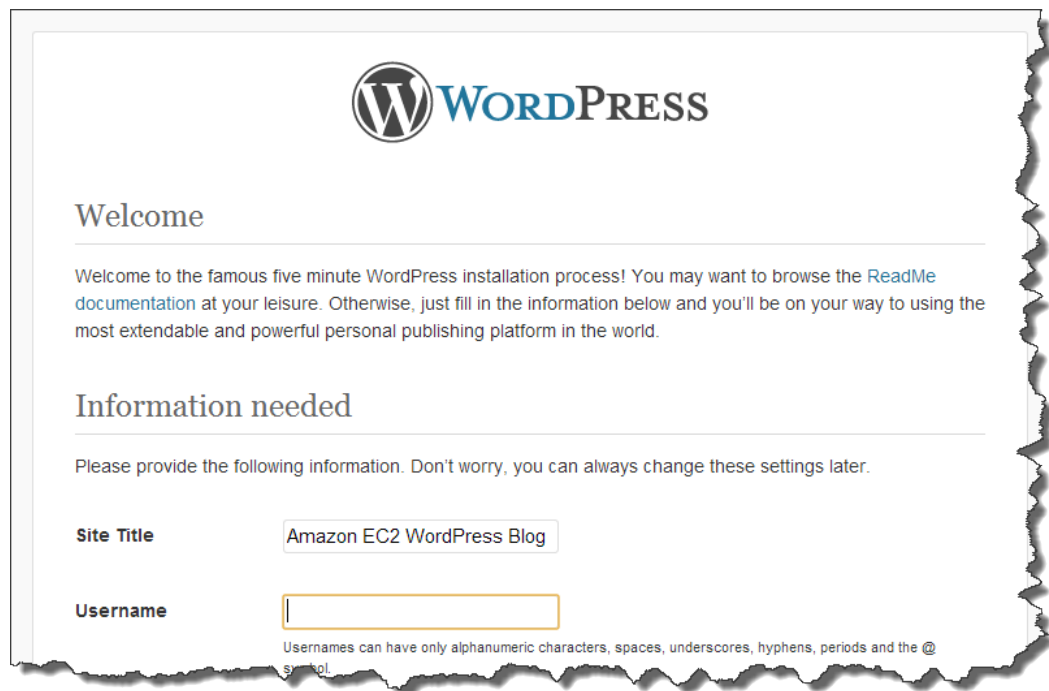
If either of these packages are not listed as installed, install them with the following command and then restart the **httpd** service.

```

[ec2-user wordpress]$ sudo yum install -y php php-mysql
[ec2-user wordpress]$ sudo service httpd restart
    
```

5. In a web browser, enter the URL of your WordPress blog (either the public DNS address for your instance, or that address followed by the `blog` folder). You should see the WordPress installation screen.

```
http://my.public.dns.amazonaws.com
```



6. Enter the remaining installation information into the WordPress installation wizard.

Field	Value
Site Title	Enter a name for your WordPress site.
Username	Enter a name for your WordPress administrator. For security purposes you should choose a unique name for this user, since this will be more difficult to exploit than the default user name, admin.

Field	Value
<b>Password</b>	Enter a strong password and then enter it again to confirm. Do not reuse an existing password, and make sure to store this password in a safe place.
<b>Your E-mail</b>	Enter the email address you want to use for notifications.

7. Click **Install WordPress** to complete the installation.

Congratulations, you should now be able to log into your WordPress blog and start posting entries.

If your WordPress blog becomes popular and you need more compute power, you might consider migrating to a larger instance type; for more information, see [Resizing Your Instance \(p. 111\)](#). If your blog requires more storage space than you originally accounted for, you could expand the storage space on your instance; for more information, see [Expanding the Storage Space of a Volume \(p. 479\)](#). If your MySQL database needs to grow, you could consider moving your database to [Amazon RDS](#) to take advantage of the service's autoscaling abilities.

For information about WordPress, see the WordPress Codex help documentation at <http://codex.wordpress.org/>. For more information about troubleshooting your installation, go to [http://codex.wordpress.org/Installing\\_WordPress#Common\\_Installation\\_Problems](http://codex.wordpress.org/Installing_WordPress#Common_Installation_Problems). For information about making your WordPress blog more secure, go to [http://codex.wordpress.org/Hardening\\_WordPress](http://codex.wordpress.org/Hardening_WordPress). For information about keeping your WordPress blog up-to-date, go to [http://codex.wordpress.org/Updating\\_WordPress](http://codex.wordpress.org/Updating_WordPress).

# Amazon Machine Images (AMI)

---

An Amazon Machine Image (AMI) is a template that contains a software configuration for your server (for example, an operating system, an application server, and applications). You specify an AMI when you launch an instance, which is a virtual server in the cloud. The AMI provides the software for the root volume of the instance. You can launch as many instances from your AMI as you need.

## Using an AMI

When you are connected to an instance, you can use it just like you use any other server. For information about launching, connecting, and using your instance, see [Amazon EC2 Instances \(p. 94\)](#).

You can search for an AMI that meets the criteria for your instance. You can search for AMIs provided by AWS or AMIs provided by the community. For more information, see [AMI Types \(p. 44\)](#) and [Finding a Suitable AMI \(p. 47\)](#).

## Creating Your Own AMI

You can customize the instance that you launch from a public AMI and then save that configuration as a custom AMI for your own use. Instances that you launch from your AMI use all the customizations that you've made.

The root storage device of the instance determines the process you follow to create an AMI. The root volume of an instance is either an Amazon EBS volume or an instance store volume. For information, see [Amazon EC2 Root Device Volume \(p. 12\)](#).

To create an Amazon EBS-backed AMI, see [Creating Amazon EBS-Backed Linux AMIs \(p. 63\)](#). To create an instance store-backed AMI, see [Creating Instance Store-Backed Linux/UNIX AMIs \(p. 66\)](#).

To help categorize and manage your AMIs, you can assign custom *tags* to them. For more information, see [Tagging Your Amazon EC2 Resources \(p. 532\)](#).

## Buying, Sharing, and Selling AMIs

After you create an AMI, you can keep it private so that only you can use it, or you can share it with a specified list of AWS accounts. You can also make your custom AMI public so that the community can use it. Building a safe, secure, usable AMI for public consumption is a fairly straightforward process, if you follow a few simple guidelines. For information about how to create and use shared AMIs, see [Shared AMIs \(p. 49\)](#).

You can purchase an AMIs from a third party, including AMIs that come with service contracts from organizations such as Red Hat. You can also create an AMI and sell it to other Amazon EC2 users. For more information about buying or selling AMIs, see [Paid AMIs \(p. 57\)](#).

### Amazon Linux

The Amazon Linux AMI is a supported and maintained Linux image provided by AWS. The following are some of the features of Amazon Linux.

- Amazon Linux is a stable, secure, and high-performance execution environment for applications running on Amazon EC2.
- Amazon Linux is provided at no additional charge to Amazon EC2 users.
- The Amazon Linux AMI is an Amazon EBS-backed, PV-GRUB image that includes Linux 3.4, AWS tools, and repository access to multiple versions of MySQL, PostgreSQL, Python, Ruby, and Tomcat.
- Amazon Linux is updated on a regular basis to include the latest components, and these updates are also made available in the **yum** repositories for installation on running instances.
- Amazon Linux includes packages that enable easy integration with AWS services, such as the Amazon EC2 API and AMI tools, the Boto library for Python, the Elastic Load Balancing tools.

For more information, see [Amazon Linux \(p. 87\)](#).

## AMI Types

You can select an AMI to use based on the following characteristics:

- Region (see [Regions and Availability Zones \(p. 7\)](#))
- Operating system
- Architecture (32-bit or 64-bit)
- [Launch Permissions \(p. 44\)](#)
- [Storage for the Root Device \(p. 45\)](#)

## Launch Permissions

The owner of an AMI determines its availability by specifying launch permissions. Launch permissions fall into the following categories.

Launch Permission	Description
public	The owner grants launch permissions to all AWS accounts.
explicit	The owner grants launch permissions to specific AWS accounts.
implicit	The owner has implicit launch permissions for an AMI.

Amazon and the Amazon EC2 community provide a large selection of public AMIs. For more information, see [Shared AMIs \(p. 49\)](#). Developers can charge for their AMIs. For more information, see [Paid AMIs \(p. 57\)](#).

## Storage for the Root Device

All AMIs are categorized as either *backed by Amazon EBS* or *backed by instance store*. The former means that the root device for an instance launched from the AMI is an Amazon EBS volume created from an Amazon EBS snapshot. The latter means that the root device for an instance launched from the AMI is an instance store volume created from a template stored in Amazon S3. For more information, see [Amazon EC2 Root Device Volume \(p. 12\)](#).

This section summarizes the important differences between the two types of AMIs. The following table provides a quick summary of these differences.

Characteristic	Amazon EBS-Backed	Amazon Instance Store-Backed
Boot time	Usually less than 1 minute	Usually less than 5 minutes
Size limit	1 TiB	10 GiB
Root device volume	Amazon EBS volume	Instance store volume
Data persistence	Data on Amazon EBS volumes persists after instance termination*; you can also attach instance store volumes that don't persist after instance termination.	Data on instance store volumes persists only during the life of the instance; you can also attach Amazon EBS volumes that persist after instance termination.
Upgrading	The instance type, kernel, RAM disk, and user data can be changed while the instance is stopped.	Instance attributes are fixed for the life of an instance.
Charges	You're charged for instance usage, Amazon EBS volume usage, and storing your AMI as an Amazon EBS snapshot.	You're charged for instance usage and storing your AMI in Amazon S3.
AMI creation/bundling	Uses a single command/call	Requires installation and use of AMI tools
Stopped state	Can be placed in stopped state where instance is not running, but the root volume is persisted in Amazon EBS	Cannot be in stopped state; instances are running or terminated

\* By default, Amazon EBS-backed instance root volumes have the `DeleteOnTermination` flag set to `true`, which causes the volume to be deleted upon instance termination. For information about how to change this so that the volume persists following termination, see [Changing the Root Device Volume to Persist \(p. 14\)](#).

### Size Limit

Amazon EC2 instance store-backed AMIs are limited to 10 GiB storage for the root device, whereas Amazon EBS-backed AMIs are limited to 1 TiB. Many Windows AMIs come close to the 10 GiB limit, so you'll find that Windows AMIs are often backed by an Amazon EBS volume.

#### Note

All Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 AMIs are backed by an Amazon EBS volume by default because of their larger size.



## Stopped State

You can *stop* an Amazon EBS-backed instance, but not an Amazon EC2 instance store-backed instance. Stopping causes the instance to stop running (its status goes from `running` to `stopping` to `stopped`). A stopped instance persists in Amazon EBS, which allows it to be restarted. Stopping is different from terminating; you can't restart a terminated instance. Because Amazon EC2 instance store-backed AMIs can't be stopped, they're either running or terminated. For more information about what happens and what you can do while an instance is stopped, see [Stop and Start Your Instance \(p. 284\)](#).

## Default Data Storage and Persistence

Instances that use an instance store volume for the root device automatically have instance store available (the root volume contains the root partition and you can store additional data). Any data on an instance store volume is deleted when the instance fails or terminates (except for data on the root device). You can add persistent storage to your instance by attaching one or more Amazon EBS volumes.

Instances that use Amazon EBS for the root device automatically have an Amazon EBS volume attached. The volume appears in your list of volumes like any other. The instances don't use any available instance store volumes by default. You can add instance storage or additional Amazon EBS volumes using a block device mapping. For more information, see [Block Device Mapping \(p. 517\)](#). For information about what happens to the instance store volumes when you stop an instance, see [Stop and Start Your Instance \(p. 284\)](#).

## Boot Times

Amazon EBS-backed AMIs launch faster than Amazon EC2 instance store-backed AMIs. When you launch an Amazon EC2 instance store-backed AMI, all the parts have to be retrieved from Amazon S3 before the instance is available. With an Amazon EBS-backed AMI, only the parts required to boot the instance need to be retrieved from the snapshot before the instance is available. However, the performance of an instance that uses an Amazon EBS volume for its root device is slower for a short time while the remaining parts are retrieved from the snapshot and loaded into the volume. When you stop and restart the instance, it launches quickly, because the state is stored in an Amazon EBS volume.

## AMI Creation

To create Linux/UNIX AMIs backed by instance store, you must create an image of your instance on the instance itself, but there aren't any API actions to help you. To create a Windows AMI backed by instance store, there's an API action that creates an image and another API action that registers the AMI.

AMI creation is much easier for AMIs backed by Amazon EBS. The `CreateImage` API action creates the AMI on both Linux/UNIX and Windows. This API action creates your Amazon EBS-backed AMI and registers it. There's also a button in the AWS Management Console that lets you create an image from a running instance. For more information, see [Creating Amazon EBS-Backed Linux AMIs \(p. 63\)](#).

## How You're Charged

With AMIs backed by instance store, you're charged for AMI storage and instance usage. With AMIs backed by Amazon EBS, you're charged for volume storage and usage in addition to the AMI and instance usage charges.

With Amazon EC2 instance store-backed AMIs, each time you customize an AMI and create a new one, all of the parts are stored in Amazon S3 for each AMI. So, the storage footprint for each customized AMI is the full size of the AMI. For Amazon EBS-backed AMIs, each time you customize an AMI and create a new one, only the changes are stored. So the storage footprint for subsequent AMIs you customize after the first is much smaller, resulting in lower AMI storage charges.

When an Amazon EBS-backed instance is stopped, you're not charged for instance usage; however, you're still charged for volume storage. We charge a full instance hour for every transition from a stopped state to a running state, even if you transition the instance multiple times within a single hour. For example, let's say the hourly instance charge for your instance is \$0.10. If you were to run that instance for one hour without stopping it, you would be charged \$0.10. If you stopped and restarted that instance twice during that hour, you would be charged \$0.30 for that hour of usage (the initial \$0.10, plus 2 x \$0.10 for each restart).

## Finding a Suitable AMI

Before you select an AMI, consider the following requirements you might have for the instances you'll launch:

- The region
- The operating system
- The architecture: 32-bit (`i386`) or 64-bit (`x86_64`)
- The root device type: Amazon EBS or instance store
- The provider: Amazon Web Services, Oracle, IBM, Microsoft, or the community

## Finding an AMI Using the Amazon EC2 Console

### To find a suitable AMI using the console

1. Open the Amazon EC2 console.
2. From the navigation bar, select a region. You can select any region that's available to you, regardless of your location. This is the region in which you'll launch your instance.
3. In the navigation pane, click **AMIs**.
4. (Optional) Use the **Filter** options to scope the list of displayed AMIs to the AMIs that interest you. For example, to list all AMIs provided by AWS, select **Public images** and then **Amazon images**.
5. (Optional) Click the **Show/Hide Columns** icon to select which image attributes to display, such as the root device type. Alternatively, you can select an AMI from the list and view its properties in the **Details** tab.
6. Before you select an AMI, it's important that you check whether it's backed by instance store or by Amazon EBS and that you are aware of the effects of this difference. For more information, see [Storage for the Root Device \(p. 45\)](#).
7. To launch an instance from this AMI, select it and then click **Launch**. For more information, see [Launch Your Instance \(p. 266\)](#). If you're not ready to launch the instance now, write down the AMI ID (`ami-xxxxxxx`) for later.

## Finding an AMI Using the Amazon EC2 CLI

Use the `ec2-describe-images` command to list your AMIs and Amazon's public AMIs.

```
PROMPT> ec2-describe-images -o self -o amazon
```

The following example shows only part of the resulting output from the command (information for 10 AMIs).

```
IMAGE ami-d8699bb1 amazon/ami-vpc-nat-1.0.0-beta.i386-ebs amazon available
public i386 machine aki-407d9529 ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-33d88c5f 8
IMAGE ami-c6699baf amazon/ami-vpc-nat-1.0.0-beta.x86_64-ebs amazon available
public x86_64 machine aki-427d952b ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-57d88c3b 8
IMAGE ami-30f30659 amazon/amzn-ami-0.9-beta.i386-ebs amazon available public
i386 machine aki-407d9529 ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-d895cdb3 10
IMAGE ami-0af30663 amazon/amzn-ami-0.9.7-beta.x86_64-ebs amazon available public
x86_64 machine aki-427d952b ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-f295cd99 10
IMAGE ami-3ac33653 amazon/amzn-ami-0.9.8-beta.i386-ebs amazon available public
i386 machine aki-407d9529 ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-14ba967f 10
IMAGE ami-38c33651 amazon/amzn-ami-0.9.8-beta.x86_64-ebs amazon available public
x86_64 machine aki-427d952b ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-10b9957b 10
IMAGE ami-08728661 amazon/amzn-ami-0.9.9-beta.i386-ebs amazon available public
i386 machine aki-407d9529 ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-674a930d 10
IMAGE ami-2272864b amazon/amzn-ami-0.9.9-beta.x86_64-ebs amazon available public
x86_64 machine aki-427d952b ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-8926ffe3 10
IMAGE ami-76f0061f amazon/amzn-ami-2010.11.1-beta.i386-ebs amazon available
public i386 machine aki-407d9529 ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-cba692a1 8
IMAGE ami-74f0061d amazon/amzn-ami-2010.11.1-beta.x86_64-ebs amazon available
public x86_64 machine aki-427d952b ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-ffa69295 8
IMAGE ami-8c1fece5 amazon/amzn-ami-2011.02.1.i386-ebs amazon available public
i386 machine aki-407d9529 ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-22fc264e 8
IMAGE ami-8elfece7 amazon/amzn-ami-2011.02.1.x86_64-ebs amazon available public
x86_64 machine aki-427d952b ebs paravirtual xen
BLOCKDEVICEMAPPING /dev/sda1 snap-a6fc26ca 8
```

To reduce the number of displayed AMIs, use a filter to list only the types of AMIs that interest you. For example, use `--filter "platform=windows"` to display only Windows-based AMIs.

After locating an AMI that meets your needs, write down its ID (ami-xxxxxxx). You can use this AMI to launch an instances. For more information, see [Launching an Instance Using the Amazon EC2 CLI](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

# Shared AMIs

A *shared AMI* is an AMI that a developer created and made available for other developers to use. One of the easiest ways to get started with Amazon EC2 is to use a shared AMI that has the components you need and then add custom content.

You use a shared AMI at your own risk. Amazon can't vouch for the integrity or security of AMIs shared by other Amazon EC2 users. Therefore, you should treat shared AMIs as you would any foreign code that you might consider deploying in your own data center and perform the appropriate due diligence.

We recommend that you get an AMI from a trusted source. If you have questions or observations about a shared AMI, use the [AWS forums](#).

Amazon's public images have an aliased owner, which appears as `amazon` in the account field. This enables you to find AMIs from Amazon easily. Other users can't alias their AMIs.

## Topics

- [Finding Shared AMIs \(p. 49\)](#)
- [Making an AMI Public \(p. 50\)](#)
- [Sharing an AMI with Specific AWS Accounts \(p. 51\)](#)
- [Using Bookmarks \(p. 52\)](#)
- [Guidelines for Shared Linux AMIs \(p. 53\)](#)

## Finding Shared AMIs

You can use the Amazon EC2 console, Amazon EC2 CLI, or the Amazon EC2 API to shared shared AMIs.

### Finding a Shared AMI Using the Console

#### To find a shared private AMI using the console

1. Open the Amazon EC2 console.
2. In the navigation pane, click **AMIs**.
3. In the first filter, select **Private images**. All AMIs that have been shared with you are listed.

#### To find a shared public AMI using the console

1. Open the Amazon EC2 console.
2. In the navigation pane, click **AMIs**.
3. To find shared AMIs, select **Public images** from the **Filter** drop-down list.
4. Use filters to list only the types of AMIs that interest you. For example, select **Amazon images** to display only Amazon's public images.

### Finding a Shared AMI Using the CLI

#### To find a shared public AMI using the command line tools

Use the `ec2-describe-images` command to list AMIs. You can scope the list to the types of AMIs that interest you, as shown in the following examples.

The following command lists all public AMIs using the `-x all` option. This list includes any public AMIs that you own.

```
PROMPT> ec2-describe-images -x all
```

The following command lists the AMIs for which you have explicit launch permissions. This list excludes any such AMIs that you own.

```
PROMPT> ec2-describe-images -x self
```

The following command lists the AMIs owned by Amazon. Amazon's public AMIs have an aliased owner, which appears as `amazon` in the account field. This enables you to find AMIs from Amazon easily. Other users can't alias their AMIs.

```
PROMPT> ec2-describe-images -o amazon
```

The following command lists the AMIs owned by the specified AWS account.

```
PROMPT> ec2-describe-images -o <target_uid>
```

The `<target_uid>` is the account ID that owns the AMIs for which you are looking.

To reduce the number of displayed AMIs, use a filter to list only the types of AMIs that interest you. For example, use `--filter "platform=windows"` to display only Windows-based AMIs.

## Using Shared AMIs

Before you use a shared AMI, take the following steps to confirm the instance is not doing anything malicious.

1. Check the SSH authorized keys file. The only key in the file should be the key you used to launch the AMI.
2. Check open ports and running services.
3. Check if SSH allows root password logins. If so, disable them; for more information, see [Disable Password-Based Logins for Root \(p. 54\)](#).
4. Check whether there are any other user accounts that might allow back-door entry to your instance. Accounts with superuser privileges are particularly dangerous.
5. Verify that all cron jobs are legitimate.

## Making an AMI Public

Amazon EC2 enables you to share your AMIs with other AWS accounts. You can allow all AWS accounts to launch the AMI (make the AMI public), or only allow a few specific accounts to launch the AMI. You are not billed when your AMI is launched by other AWS accounts; only the accounts launching the AMI are billed.

Before you share an AMI, make sure to read the security considerations in [Guidelines for Shared Linux AMIs \(p. 53\)](#).

### Note

If an AMI has a product code, you can't make it public. You must share the AMI with specific AWS accounts.

## Sharing a Public AMI Using the Console

### To share a public AMI using the console

1. Open the Amazon EC2 console.
2. In the navigation pane, click **AMIs**.
3. Select your AMI in the list, click the **Permissions** tab, and select the **Public** radio button.

## Sharing a Public AMI Using the CLI

Each AMI has a `launchPermission` property that controls which AWS accounts, besides the owner's, are allowed to use that AMI to launch instances. By modifying the `launchPermission` property of an AMI, you can make the AMI public (which grants launch permissions to all AWS accounts or share it with only the AWS accounts that you specify).

You can add or remove account IDs from the list of accounts that have launch permissions for an AMI. To make the AMI public, specify the `all` group. You can specify both public and explicit launch permissions.

### To make an AMI public

Use the `ec2-modify-image-attribute` command as follows to add the `all` group to the `launchPermission` list for the specified AMI.

```
PROMPT> ec2-modify-image-attribute ami-2bb65342 --launch-permission -a all
launchPermission      ami-2bb65342      ADD      group      all
```

To verify the launch permissions of the AMI, use the following command.

```
PROMPT> ec2-describe-image-attribute ami-2bb65342 -l
launchPermission      ami-2bb65342      group      all
```

To make the AMI private again, remove the `all` group from its launch permissions. Note that the owner of the AMI always has launch permissions and is therefore unaffected by this command.

```
PROMPT> ec2-modify-image-attribute ami-2bb65342 -l -r all
launchPermission      ami-2bb65342      REMOVE    group      all
```

## Sharing an AMI with Specific AWS Accounts

You can share an AMI with specific AWS accounts without making the AMI public. All you need are the AWS account IDs.

## Sharing an AMI Using the Console

### To grant explicit launch permissions using the console

1. Open the Amazon EC2 console.
2. In the navigation pane, click **AMIs**.
3. Select your AMI in the list, then click the **Permissions** tab.
4. Specify the AWS account number of the user with whom you want to share the AMI in the **AWS Account Number** field, then click **Save**.

To share this AMI with multiple users, click **Add Permission** and repeat the above step until you have added all the required users.

5. To allow create volume permissions for snapshots, check **Add "create volume" permissions to the following associated snapshots when creating permissions.**

**Note**

You do not need to share the Amazon EBS snapshots than an AMI references in order to share the AMI. Only the AMI itself needs to be shared; the system automatically provides the instance access to the referenced Amazon EBS snapshots for the launch.

## Sharing an AMI Using the CLI

Use the `ec2-modify-image-attribute` command to share an AMI as shown in the following examples.

### To grant explicit launch permissions using the CLI

The following command grants launch permissions for the specified AMI to the specified AWS account.

```
PROMPT> ec2-modify-image-attribute ami-2bb65342 -l -a 111122223333
launchPermission      ami-2bb65342      ADD      userId  111122223333
```

### To remove launch permissions for an account

The following command removes launch permissions for the specified AMI from the specified AWS account:

```
PROMPT> ec2-modify-image-attribute ami-2bb65342 -l -r 111122223333
launchPermission      ami-2bb65342      REMOVE  userId  111122223333
```

### To remove all launch permissions

The following command removes all public and explicit launch permissions from the specified AMI. Note that the owner of the AMI always has launch permissions and is therefore unaffected by this command.

```
PROMPT> ec2-reset-image-attribute ami-2bb65342 -l
launchPermission      ami-2bb65342      RESET
```

## Using Bookmarks

If you have created a public AMI, or shared an AMI with another AWS user, you can create a *bookmark* that allows a user to access your AMI and launch an instance in their own account immediately. This is an easy way to share AMI references, so users don't have to spend time finding your AMI in order to use it.

Note that your AMI must be public, or you must have shared it with the user to whom you want to send the bookmark.

### To create a bookmark for your AMI

1. Type a URL with the following information, where `<region>` is the region in which your AMI resides, and `<ami_id>` is the ID of the AMI:

```
https://console.aws.amazon.com/ec2/home?region=<region>#LaunchInstanceWizard:ami=<ami_id>
```

For example, this URL launches an instance from the ami-2bb65342 AMI in the us-east-1 region:

```
https://console.aws.amazon.com/ec2/v2/home?region=us-east-1#LaunchInstanceWizard:ami=ami-2bb65342
```

2. Distribute the link to users who want to use your AMI.
3. To use a bookmark, click the link or copy and paste it into your browser. The launch wizard opens, with the AMI already selected.

## Guidelines for Shared Linux AMIs

If you follow these guidelines, you'll provide a better user experience and make your users' instances less vulnerable to security issues.

If you are building AMIs for AWS Marketplace, see [Building AMIs for AWS Marketplace](#) for guidelines, policies and best practices.

For additional information about sharing AMIs safely, see the following articles on the AWS Developer Resources website:

- [How To Share and Use Public AMIs in A Secure Manner](#)
- [Public AMI Publishing: Hardening and Clean-up Requirements](#)

### Shared AMI Guidelines

1	<a href="#">Update the AMI Tools at Boot Time (p. 53)</a>
2	<a href="#">Disable Password-Based Logins for Root (p. 54)</a>
3	<a href="#">Disable Root Access (p. 54)</a>
4	<a href="#">Install Public Key Credentials (p. 55)</a>
5	<a href="#">Disabling sshd DNS Checks (Optional) (p. 56)</a>
6	<a href="#">Identify Yourself (p. 56)</a>
7	<a href="#">Protect Yourself (p. 56)</a>

## Update the AMI Tools at Boot Time

For AMIs backed by instance store, we recommend that your AMIs download and upgrade the Amazon EC2 AMI creation tools during startup. This ensures that new AMIs based on your shared AMIs have the latest AMI tools.

For [Amazon Linux](#), add the following to `rc.local`:



```
# Update the Amazon EC2 AMI tools
echo " + Updating EC2 AMI tools"
yum update -y aws-amitools-ec2
echo " + Updated EC2 AMI tools"
```

Use this method to automatically update other software on your image.

**Note**

When deciding which software to automatically update, consider the amount of WAN traffic that the update will generate (your users will be charged for it) and the risk of the update breaking other software on the AMI.

For other distributions, make sure you have the latest AMI tools.

## Disable Password-Based Logins for Root

Using a fixed root password for a public AMI is a security risk that can quickly become known. Even relying on users to change the password after the first login opens a small window of opportunity for potential abuse.

To solve this problem, disable password-based logins for the root user. Additionally, we recommend you disable root access.

### To disable password-based logins for root

1. Open the `/etc/ssh/sshd_config` file with a text editor and locate the following line:

```
#PermitRootLogin yes
```

2. Change the line to:

```
PermitRootLogin without-password
```

The location of this configuration file might differ for your distribution, or if you are not running OpenSSH. If this is the case, consult the relevant documentation.

## Disable Root Access

When you work with shared AMIs, it is a known best practice to have a secure environment; one of the elements associated with a secure environment is ensuring the root password is not empty. To do this, log into your running instance and issue the following command to disable root access:

```
sudo passwd -l root
```

**Note**

This does not impact the use of `sudo`.

## Remove SSH Host Key Pairs

If you plan to share an AMI derived from a public AMI, remove the existing SSH host key pairs located in `/etc/ssh`. This forces SSH to generate new unique SSH key pairs when someone launches an instance using your AMI, improving security and reducing the likelihood of "man-in-the-middle" attacks.

The following list shows the SSH files to remove.

- `ssh_host_dsa_key`
- `ssh_host_dsa_key.pub`
- `ssh_host_key`
- `ssh_host_key.pub`
- `ssh_host_rsa_key`
- `ssh_host_rsa_key.pub`

### Important

If you forget to remove the existing SSH host key pairs from your public AMI, our routine auditing process notifies you and all customers running instances of your AMI of the potential security risk. After a short grace period, we mark the AMI private.

## Install Public Key Credentials

After configuring the AMI to prevent logging in using a password, you must make sure users can log in using another mechanism.

Amazon EC2 allows users to specify a public-private key pair name when launching an instance. When a valid key pair name is provided to the `RunInstances` API call (or through the command line API tools), the public key (the portion of the key pair that Amazon EC2 retains on the server after a call to `CreateKeyPair` or `ImportKeyPair`) is made available to the instance through an HTTP query against the instance metadata.

To log in through SSH, your AMI must retrieve the key value at boot and append it to `/root/.ssh/authorized_keys` (or the equivalent for any other user account on the AMI). Users can launch instances of your AMI with a key pair and log in without requiring a root password.

```
if [ ! -d /root/.ssh ] ; then
    mkdir -p /root/.ssh
    chmod 700 /root/.ssh
fi
# Fetch public key using HTTP
curl http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key > /tmp/my-key
if [ $? -eq 0 ] ; then
    cat /tmp/my-key >> /root/.ssh/authorized_keys
    chmod 700 /root/.ssh/authorized_keys
    rm /tmp/my-key
fi
```

This can be applied to any user account; you do not need to restrict it to root.

### Note

Rebundling an instance based on this image includes the key with which it was launched. To prevent the key's inclusion, you must clear out (or delete) the `authorized_keys` file or exclude this file from rebundling.

## Disabling sshd DNS Checks (Optional)

Disabling sshd DNS checks slightly weakens your sshd security. However, if DNS resolution fails, SSH logins still work. If you do not disable sshd checks, DNS resolution failures prevent all logins.

### To disable sshd DNS checks

1. Open the `/etc/ssh/sshd_config` file with a text editor and locate the following line:

```
#UseDNS yes
```

2. Change the line to:

```
UseDNS no
```

#### Note

The location of this configuration file can differ for your distribution or if you are not running OpenSSH. If this is the case, consult the relevant documentation.

## Identify Yourself

Currently, there is no easy way to know who provided a shared AMI, because each AMI is represented by an account ID.

We recommend that you post a description of your AMI, and the AMI ID, in the [Amazon EC2 forum](#). This provides a convenient central location for users who are interested in trying new shared AMIs. You can also post the AMI to the [Amazon Machine Images \(AMIs\)](#) page.

## Protect Yourself

The previous sections described how to make your shared AMIs safe, secure, and usable for the users who launch them. This section describes guidelines to protect yourself from the users of your AMI.

We recommend against storing sensitive data or software on any AMI that you share. Users who launch a shared AMI might be able to rebundle it and register it as their own. Follow these guidelines to help you to avoid some easily overlooked security risks:

- Always delete the shell history before bundling. If you attempt more than one bundle upload in the same image, the shell history contains your secret access key. The following example should be the last command executed before bundling from within the instance.

```
rm ~/.bash_history ~/.zsh_history
```

For the following two commands, AWS recommends using the `--e` (`--exclude`) option on `ec2-bundle-vol` to instruct it to skip the directories and subdirectories listed within the parameter from the bundle operation. List the directories and subdirectories containing secret information within the parameter. For more information, see `ec2-bundle-vol` in the *Amazon Elastic Compute Cloud Command Line Reference*.

- Bundling a running instance requires your private key and X.509 certificate. Put these and other credentials in a location that is not bundled (such as the instance store).
- Exclude the ssh authorized keys when bundling the image. The Amazon public images store the public key used to launch an instance with its ssh authorized keys file.

**Note**

Unfortunately, it is not possible for this list of guidelines to be exhaustive. Build your shared AMIs carefully and take time to consider where you might expose sensitive data.

## Paid AMIs

A *paid AMI* is an AMI that you can purchase from a developer.

Amazon EC2 integrates with Amazon DevPay and AWS Marketplace, enabling developers to charge other Amazon EC2 users for the use of their AMIs or to provide support for instances. For more information about Amazon DevPay, see the [Amazon DevPay](#) site.

The AWS Marketplace is an online store where you can buy software that runs on AWS; including AMIs that you can use to launch your EC2 instance. The AWS Marketplace AMIs are organized into categories, such as Developer Tools, to enable you to find products to suit your requirements. For more information about AWS Marketplace, see the [AWS Marketplace](#) site.

Launching an instance from a paid AMI is the same as launching an instance from any other AMI. No additional parameters are required. The instance is charged according to the rates set by the owner of the AMI, as well as the standard usage fees for the related web services; for example, the hourly rate for running a m1.small instance type in Amazon EC2. The owner of the paid AMI can confirm whether a specific instance was launched using that paid AMI.

**Important**

All paid AMIs from Amazon DevPay are backed by instance store. AWS Marketplace supports AMIs backed by Amazon EBS.

**Topics**

- [Selling Your AMI \(p. 57\)](#)
- [Finding a Paid AMI \(p. 58\)](#)
- [Purchase a Paid AMI \(p. 58\)](#)
- [Getting the Product Code for Your Instance \(p. 59\)](#)
- [Using Paid Support \(p. 59\)](#)
- [Bills for Paid and Supported AMIs \(p. 60\)](#)
- [Managing Your AWS Marketplace Subscriptions \(p. 60\)](#)

## Selling Your AMI

You can sell your AMI using either AWS Marketplace or Amazon DevPay. Both help customers buy software that runs on AWS, but AWS Marketplace offers a better shopping experience, making it easier for customers to find your AMI. AWS Marketplace also supports AWS features that Amazon DevPay doesn't support, such as Amazon EBS-backed AMIs, Reserved Instances, and Spot Instances.

For information about how to sell your AMI on AWS Marketplace, see [Selling on AWS Marketplace](#).

For information about how to sell your AMI on Amazon DevPay, see [Using DevPay with Your Amazon EC2 AMI](#).

## Finding a Paid AMI

There are several ways that you can find AMIs that are available for you to purchase. For example, you can use [AWS Marketplace](#), the Amazon EC2 console, or the Amazon EC2 CLI. Alternatively, a developer might let you know about a paid AMI themselves.

### Finding a Paid AMI Using the Amazon EC2 Console

#### To find a paid AMI using the AMIs page

1. Open the Amazon EC2 console.
2. In the navigation pane, click **AMIs**.
3. Select **Public images** from the first **Filter** list, **Marketplace images** from the second **Filter** list, and the operating system from the third **Filter** list.

### Finding a Paid AMI Using AWS Marketplace

#### To find a paid AMI using AWS Marketplace

1. Open [AWS Marketplace](#).
2. Enter the name of the operating system in the search box, and click **Go**.
3. To scope the results further, use one of the categories or filters.
4. Each product is labeled with its product type: either AMI or Software as a Service.

### Finding a Paid Windows AMI Using the Amazon EC2 CLI

#### To find a paid Windows AMI using the Amazon EC2 CLI

You can also find a paid Windows AMI using the [ec2-describe-images](#) command as follows.

```
PROMPT> ec2-describe-images
```

This command returns numerous fields that describe each AMI. If the output for an AMI contains a product code, it is a paid AMI. The following example output from `ec2-describe-images` for a paid AMI. The product code is `ACD42B6F`.

```
IMAGE    ami-a5bf59cc    image_source    123456789012    available public
ACD42B6F    x86_64    machine    instance-store
```

## Purchase a Paid AMI

You must sign up for (purchase) a paid AMI before you can launch an instance using the AMI.

Typically a seller of a paid AMI presents you with information about the AMI, including its price and a link where you can buy it. When you click the link, you're first asked to log into AWS, and then you can purchase the AMI.

#### **Important**

You don't get the discount from Reserved Instances if you use a paid AMI from Amazon DevPay. That is, if you purchase Reserved Instances, you don't get the lower price associated with them

when you launch a paid AMI. You always pay the price that's specified by the seller of the paid AMI.

## Purchasing a Paid AMI Using the Console

You can purchase a paid AMI by using the Amazon EC2 launch wizard. For more information, see [Launching an AWS Marketplace Instance \(p. 271\)](#).

## Subscribing to a Product Using AWS Marketplace

To use the AWS Marketplace, you must have an AWS account. To launch instances from AWS Marketplace products, you must be signed up to use the Amazon EC2 service, and you must be subscribed to the product from which to launch the instance. There are two ways to subscribe to products in the AWS Marketplace:

- **The AWS Marketplace website:** You can launch preconfigured software quickly with the 1-Click deployment feature.
- **The EC2 launch wizard:** You can search for an AMI and launch an instance directly from the wizard. For more information, see [Launching an AWS Marketplace Instance \(p. 271\)](#).

## Purchasing a Paid AMI From a Developer

The developer of a paid AMI can enable you to purchase a paid AMI that isn't listed in AWS Marketplace. The developer provides you with a link that enables you to purchase the product through Amazon. You can sign in with your Amazon.com credentials and select a credit card that's stored in your Amazon.com account to use when purchasing the AMI.

## Getting the Product Code for Your Instance

You can determine whether your instance has an Amazon DevPay or AWS Marketplace product code using its instance metadata. For more information about retrieving metadata, see [Instance Metadata and User Data \(p. 290\)](#).

To retrieve a product code, use the following query:

```
GET http://169.254.169.254/2007-03-01/meta-data/product-codes
```

If the instance has a product code, Amazon EC2 returns it. For example:

```
774F4FF8
```

## Using Paid Support

Amazon EC2 also enables developers to offer support for software (or derived AMIs). Developers can create support products that you can sign up to use. During sign-up for the support product, the developer gives you a product code, which you must then associate with your own AMI. This enables the developer to confirm that your instance is eligible for support. It also ensures that when you run instances of the product, you are charged according to the terms for the product specified by the developer.

### Important

You can't use a support product with Reserved Instances. You always pay the price that's specified by the seller of the support product.

To associate a product code with your AMI, use the `ec2-modify-image-attribute` command as follows, where `ami_id` is the ID of the AMI and `product_code` is the product code:

```
PROMPT> ec2-modify-image-attribute ami_id --product-code product_code
```

After you set the product code attribute, it cannot be changed or removed.

## Bills for Paid and Supported AMIs

At the end of each month, you receive an email with the amount your credit card has been charged for using any paid or supported AMIs during the month. This bill is separate from your regular Amazon EC2 bill. For more information, see [Paying For AWS Marketplace Products](#).

## Managing Your AWS Marketplace Subscriptions

On the AWS Marketplace website, you can check your subscription details, view the vendor's usage instructions, manage your subscriptions, and more.

### To check your subscription details

1. Log in to the [AWS Marketplace](#).
2. Click **Your Account**.
3. Click **Manage Your Software Subscriptions**.
4. All your current subscriptions are listed. Click **Usage Instructions** to view specific instructions for using the product, for example, a user name for connecting to your running instance.

### To cancel an AWS Marketplace subscription

1. Ensure that you have terminated any instances running from the subscription.
  - a. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
  - b. In the navigation pane, click **Instances**.
  - c. Select the instance, and select **Terminate** from the **Actions** menu. When prompted, click **Yes, Terminate**.
2. Log in to the [AWS Marketplace](#), and click **Your Account**, then **Manage Your Software Subscriptions**.
3. Click **Cancel subscription**. You are prompted to confirm your cancellation.

#### Note

After you've canceled your subscription, you are no longer able to launch any instances from that AMI. To use that AMI again, you need to resubscribe to it, either on the AWS Marketplace website, or through the launch wizard in the Amazon EC2 console.

# Creating Amazon EBS-Backed AMIs Using the Console

### Topics

- [Create an AMI from an Instance \(p. 61\)](#)

- [Delete an AMI and a Snapshot \(p. 62\)](#)

This section walks you through the process of creating an Amazon EBS-backed AMI from a running Amazon EBS-backed instance. For more information about Amazon EBS-backed AMIs and instance store-backed AMIs, see [Storage for the Root Device \(p. 45\)](#). For instructions that use the command line tools or API, see [Creating Amazon EBS-Backed Linux AMIs \(p. 63\)](#).

## Create an AMI from an Instance

### To create an AMI from a running Amazon EBS-backed instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. If you don't have a running instance that uses an Amazon EBS volume for the root device, you must launch one. For instructions, see [Launching an Instance \(p. 266\)](#).
3. [Optional] Connect to the instance and customize it however you want. For example, you can install software and applications, copy data, or attach additional EBS volumes. For more information about connecting to an instance, see [Connect to Your Amazon EC2 Instance \(p. 273\)](#).
4. In the navigation pane, click **Instances** to view a list of your instances. Right-click your running instance and select **Create Image**.

#### Tip

If this option is disabled, your instance isn't an Amazon EBS-backed instance.

The **Create Image** dialog box appears.

5. Fill in the requested information as follows, and click **Create Image**.
  - a. A unique name for the image.
  - b. [Optional] A description of the image (up to 255 characters).
  - c. By default, Amazon EC2 shuts down the instance, takes snapshots of any attached volumes, creates and registers the AMI, and then reboots the instance. Select **No reboot** if you don't want your instance to be shut down.

#### Warning

If you select the **No Reboot** option, the file system integrity of the created image can't be guaranteed.

- d. [Optional] You can modify the root volume, EBS volumes, and instance store volumes as follows:
  - To change the size of the root volume, locate the **Root** volume in the **Type** column, and fill in the **Size** field.
  - To suppress an EBS volume specified by the block device mapping of the AMI used to launch the instance, locate the EBS volume in the list and click **Delete**.
  - To add an EBS volume, click **Add New Volume**, select **EBS** from the **Type** list, and fill in the fields. When you then launch an instance from your new AMI, these additional volumes are automatically attached to the instance. Empty volumes must be formatted and mounted. Volumes based on a snapshot must be mounted.
  - To suppress an instance store volume specified by the block device mapping of the AMI used to launch the instance, locate the volume in the list and click **Delete**.
  - To add an instance store volume, click **Add New Volume**, select **Instance Store** from the **Type** list, and select a device name from the **Device** list. When you launch an instance from your new AMI, these additional volumes are automatically initialized and mounted. These volumes don't contain data from the instance store volumes of the running instance from which you based your AMI.

For more information, see [Amazon EC2 Root Device Volume \(p. 12\)](#), [Amazon EC2 Instance Store \(p. 508\)](#), and [Block Device Mapping \(p. 517\)](#)



6. Click **AMIs** in the navigation pane to view the AMI's status. While the new AMI is being created, its status is `pending`.

It takes a few minutes for the whole process to finish.

7. After your new AMI's status is `available`, go to the **Snapshots** page and view the new snapshot that was created for the new AMI. Any instance you launch from the new AMI uses this snapshot for its root device volume.
8. Go back to the **AMIs** page, select the image, and click **Launch**. The launch wizard opens.
9. Walk through the wizard to launch an instance of your new AMI.
10. After your instance's status is `running`, connect to the instance and verify that any changes you made to the original AMI have persisted.

You now have a new AMI and snapshot that you just created. Both continue to incur charges to your account until you stop or delete them.

## Delete an AMI and a Snapshot

### To delete an AMI and a snapshot

1. Go to the **AMIs** page. Select the AMI, click **Actions**, then select **Deregister**. When asked for confirmation, click **Continue**.  
The image is deregistered, which means it is deleted and can no longer be launched.
2. Go to the **Snapshots** page. Right-click the snapshot and select **Delete Snapshot**. When asked for confirmation, click **Yes, Delete**.  
The snapshot is deleted.

## Creating Your Own AMIs

There are many public AMIs available to you. To see the available AMIs, go to [Amazon Machine Images \(AMIs\)](#). If the available public AMIs don't provide everything that you're looking for, you can create an AMI that meets your needs.

Creating your own AMI helps you make the most of Amazon EC2. Your AMI becomes the basic unit of deployment; it enables you to rapidly boot new custom instances as you need them. This section gives an overview of your AMI creation options, identifies the tools you need, and walks you through the process.

Before you begin this section, you should be familiar with AMI and instance concepts. For more information, see the following sections:

- [Amazon Machine Images \(AMI\)](#) (p. 43)
- [Amazon EC2 Instances](#) (p. 94)
- [Amazon Elastic Block Store](#) (p. 446)

## Overview of the AMI Creation Process

There are a few different ways to create an AMI. The process you must follow to create an AMI depends on whether you are creating an Amazon EBS-backed AMI or an Amazon EC2 instance store-backed AMI. There are significant differences between Amazon EBS-backed and Amazon EC2 instance store-backed AMIs, such as data persistence. For information on the differences between these choices, see [Storage for the Root Device](#) (p. 45).

First, decide which operating system and root device volume you want, and then you'll know which of these processes to use for creating the AMI:

- **Amazon EBS-backed AMI**  
The same general process applies to Linux/UNIX and Windows.
  - **Linux/UNIX**—[Creating Amazon EBS-Backed Linux AMIs \(p. 63\)](#)
  - **Windows**—[Creating an Amazon EBS-Backed Windows AMI \(Amazon Elastic Compute Cloud Microsoft Windows Guide\)](#)
- **Amazon EC2 instance store-backed AMI**  
The process depends on the operating system.
  - **Linux/UNIX**—[Creating Instance Store-Backed Linux/UNIX AMIs \(p. 66\)](#)
  - **Windows**—[Creating an Instance Store-Backed Windows AMI \(Amazon Elastic Compute Cloud Microsoft Windows Guide\)](#)
- **Creating an AMI from a Resized Instance**

If you have [changed the size of your instance \(p. 111\)](#) and the root device for your instance is an instance store volume, you must create an AMI from your current instance, launch a new instance from this AMI with the instance type you need, then terminate the instance you no longer need. For detailed instructions on how to do this, click [here \(p. 112\)](#).

## Creating Amazon EBS-Backed Linux AMIs

To create an Amazon EBS-backed Linux AMI, start with an Amazon EBS-backed AMI (for example, one of the public AMIs that Amazon provides), and modify it to suit your particular needs (note that as Amazon EBS-backed instances are stored as Amazon EBS data, standard storage rates apply). If you start with an Amazon instance store-backed instance, you cannot create an Amazon EBS-backed AMI using these instructions. For more information about Amazon EBS-backed AMIs and instance store-backed AMIs, see [Storage for the Root Device \(p. 45\)](#).

### Note

This topic describes the process for creating an Amazon EBS-backed Linux AMI. For information about Amazon EBS-backed Windows AMIs, see [Creating an Amazon EBS-Backed Windows AMI](#). For instructions using an instance store-backed AMI, see [Creating Instance Store-Backed Linux/UNIX AMIs \(p. 66\)](#).

### Topics

- [Creating an Amazon EBS-Backed Linux AMI \(p. 63\)](#)
- [Special Cases \(p. 64\)](#)
- [How to Create Amazon EBS-Backed AMIs \(p. 64\)](#)
- [Converting Amazon EC2 instance store-backed AMIs to EBS-Backed AMIs \(p. 66\)](#)

## Creating an Amazon EBS-Backed Linux AMI

### To create an Amazon EBS-backed Linux AMI

1. Launch an instance of an Amazon EBS-backed AMI that is similar to the AMI that you want to create. For example, you might take a public AMI that uses the operating system you want to use for your AMI.  
  
The instance must be launched from an Amazon EBS-backed AMI; you cannot use an Amazon EC2 instance store-backed AMI.
2. When the instance is running, customize it as desired. For example, you could attach additional Amazon EBS volumes, load applications, or copy data.

### Important

If you customize your instance with instance-store volumes or additional EBS volumes besides the root device, the new AMI contains block device mapping information for those volumes and new instances automatically launch with the additional volumes. However, instance-store volumes on the new instance won't contain any customized data. For more information, see [Create an AMI from an Instance](#). If you want your data to persist, you must use an EBS volume instead of an instance-store (ephemeral) volume. For more information, see [Block Device Mapping \(p. 517\)](#).

3. When the instance is set up the way you want it, it is best to stop the instance before you create the AMI to ensure data integrity. Follow these steps to stop the instance:
  - a. Right-click your running instance and select **Stop**.
  - b. When prompted to confirm this, click **Yes, Stop**.
  
4. Create an AMI from that instance.

It takes several minutes for the entire process to complete. If you customized the instance with instance store volumes or additional EBS volumes besides the root device, the new AMI contains block device mapping information for those volumes. When you launch an instance from your new AMI, the instance automatically launches with the additional volumes. The instance store volumes are new and don't contain any data from the instance store volumes of the original instance used to create the AMI.

Amazon EC2 powers down the instance before creating the AMI to ensure that everything on the instance is stopped and in a consistent state during the creation process. If you're confident that your instance is in a consistent state appropriate for AMI creation, you can add the `--no-reboot` flag to `ec2-create-image` or `CreateImage` that tells Amazon EC2 not to power down and reboot the instance. With this flag, the instance remains running throughout the AMI creation process. Some file systems, such as xfs, can freeze and unfreeze activity, making it safe to create the image without rebooting the instance.

## Special Cases

In some cases, the general tasks in creating Amazon EBS-backed AMIs don't apply:

- You don't have the original AMI from which to launch instances.

In this case, you can create an Amazon EBS-backed AMI by registering a snapshot of a root device. You must own the snapshot and it must be a Linux/UNIX system (this process is not available for Windows instances). For more information about creating an AMI this way, see [Launching an Instance from a Backup \(p. 271\)](#).

- You have an Amazon EC2 instance store-backed Linux/UNIX AMI.

In this case, you can convert the AMI to be backed by Amazon EBS. You cannot convert a Windows AMI backed by instance store. For more information about converting a Linux/UNIX AMI, see [Converting Amazon EC2 instance store-backed AMIs to EBS-Backed AMIs \(p. 66\)](#).

## How to Create Amazon EBS-Backed AMIs

You can create an Amazon EBS-backed AMI by using the AWS Management Console, the command line tools, or the API. The following section describes the steps using each tool or interface.

## AWS Management Console

For instructions that use the AWS Management Console, see [Creating Amazon EBS-Backed AMIs Using the Console \(p. 60\)](#).

## Command Line Interface

### To create an Amazon EBS-backed AMI

1. Use the `ec2-create-image` command to create an image.

```
PROMPT> ec2-create-image -n your_image_name instance_id
```

For example:

```
PROMPT> ec2-create-image -n "My AMI" i-eb977f82
```

Amazon EC2 creates an image and returns an AMI ID.

```
IMAGE ami-8675309
```

2. If you want to check whether the AMI is ready, use the `ec2-describe-images` command as follows:

```
$ ec2-describe-images -o self
```

Amazon EC2 returns information about the AMI.

If the AMI you start with doesn't already have the storage devices you want attached, you can add them by creating EBS volumes or using block device mapping. To create EBS volumes, use `ec2-create-volume` and `ec2-attach-volume`.

You also can call `ec2-run-instances` with block device mapping information for the devices you want to add. For more information about block device mapping, see [Block Device Mapping \(p. 517\)](#).

## API

To create an Amazon EBS-backed AMI, construct the following query request to create an image:

```
https://ec2.amazonaws.com/  
?Action=CreateImage  
&InstanceId=instance_id  
&Name=My_Ami  
&AUTHPARAMS
```

In the following example response, Amazon EC2 creates the image and returns its AMI ID.

```
<CreateImageResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <imageId>ami-8675309</imageId>  
</CreateImageResponse>
```

AMI creation can take time. You can check whether the AMI is ready using `DescribeImages`.

If the AMI you start with doesn't already have the storage devices you want attached, you can add them by creating EBS volumes or using block device mapping. To create EBS volumes, use `CreateVolume` and `AttachVolume`.

You also can call `RunInstances` with block device mapping information for the devices you want to add. For more information about block device mapping, see [Block Device Mapping](#) (p. 517).

## Converting Amazon EC2 instance store-backed AMIs to EBS-Backed AMIs

There's no simple API or button in the AWS Management Console that converts an existing Amazon EC2 instance store-backed AMI to an Amazon EBS-backed AMI. However, you can convert Amazon EC2 instance store-backed Linux/UNIX AMIs to EBS-backed systems manually.

### Important

You can't convert an instance store-backed Windows AMI to an EBS-backed Windows AMI. You must start with a public EBS-backed Windows AMI, modify it to meet your specifications, and then create an image from it. For information, see [Creating Amazon EBS-Backed AMIs Using the Console](#) (p. 60).

The following table describes the conversion process.

### How to convert a Linux/UNIX Amazon EC2 instance store-backed AMI to an EBS-backed AMI

1	Copy the AMI's root device information to an Amazon EBS volume. For more information, see the related task list in <a href="#">Root Device Storage Usage Scenarios</a> (p. 15)
2	Create a snapshot of that volume. For more information, see <a href="#">Creating an Amazon EBS Snapshot</a> (p. 485).
3	Register the image with a block device mapping that maps the root device name of your choice to the snapshot you just created. For an example, see <a href="#">Block Device Mapping</a> (p. 517).

You might find it useful to refer to available blog posts that discuss conversion. The following are two example blogs; AWS, however, takes no responsibility for the completeness or accuracy of the content:

- <http://www.elastician.com/2009/12/creating-ebs-backed-ami-from-s3-backed.html>
- <http://coderslike.us/2009/12/07/amazon-ec2-boot-from-ebs-and-ami-conversion/>

## Creating Instance Store-Backed Linux/UNIX AMIs

### Topics

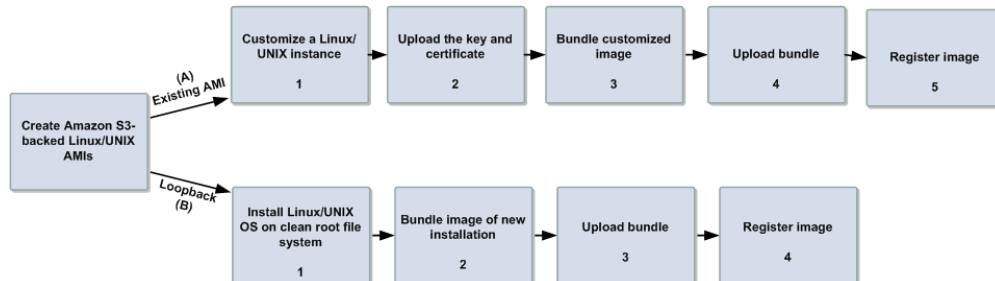
- [Tools You Need](#) (p. 67)
- [Creating an Instance Store-Backed AMI From an Existing AMI](#) (p. 68)
- [Creating an Instance Store-Backed AMI From a Loopback](#) (p. 72)

For Linux/UNIX systems, you have two common ways to prepare Amazon EC2 instance store-backed AMIs. The easiest method (A) involves launching an existing public AMI and modifying it according to your requirements. For more information, see [Creating an Instance Store-Backed AMI From an Existing AMI](#) (p. 68).

Another approach (B) is to build a fresh installation either on a stand-alone machine or on an empty file system mounted by loopback. The process entails building an operating system installation from scratch.

After you've built the installation package to your satisfaction, you must bundle it using the AMI tool for bundling volumes and register it using the command line tool for registering images. For information, see [Creating an Instance Store-Backed AMI From a Loopback \(p. 72\)](#).

The following diagram shows the general tasks in creating Amazon EC2 instance store-backed Linux/UNIX AMIs.



This section discusses the steps for creating AMIs from an existing file and from a loopback, and some basics about the AMI tools.

## Tools You Need

Amazon created the Amazon EC2 AMI tools to help you perform specific tasks for Amazon EC2 instance store-backed Linux/UNIX AMIs. You use these AMI tools, which are a set of command line utilities, for bundling and uploading Amazon EC2 instance store-backed Linux/UNIX AMIs. You also use these AMI tools for managing these bundled images. For information about the specific AMI tools, see [AMI Tools Reference](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

When you bundle an Amazon EC2 instance store-backed Linux/UNIX AMI and you start with an instance, you use the AMI tools for bundling and uploading the bundle, and then you use the API tools to register the image. If you are creating an Amazon EC2 instance store-backed AMI from a loopback, you first prepare the instance, then use the AMI tools to bundle before you use the API tools to register the image you created.

If you are starting with an instance of an Amazon public AMI, it might already have the AMI tools installed. Try running the command `ec2-bundle-vol` to check if the instance already has the AMI tools.

If the tools are already installed, you can jump to the section that discusses the bundling process you want to complete:

- [Creating an Instance Store-Backed AMI From an Existing AMI \(p. 68\)](#)
- [Creating an Instance Store-Backed AMI From a Loopback \(p. 72\)](#)

If the tools aren't installed, read on. This section describes installation and usage information when using AMI tools.

## Install the AMI Tools

The AMI tools are available in both a zip file and as an RPM suitable for running on Fedora Core with Ruby 1.8.2 (or greater) installed. You need root privileges to install the software.

For information about installing the AMI tools, see [Amazon EC2 AMI Tools](#).

### To install the AMI tools

1. Install Ruby using the `yum` package manager.

```
# yum install ruby
```

2. Install the AMI tools RPM.

```
# rpm -i ec2-ami-tools-x.x-xxxx.i386.rpm
```

## View the AMI Tools Documentation

This section describes how to view Linux/UNIX documentation.

### To view the manual for each tool

- Append **--manual** to the command that invokes the tool.

```
$ ec2-bundle-image --manual
```

### To view help for each tool

- Append **--help** to the command that invokes the tool.

```
$ ec2-bundle-image --help
```

## Creating an Instance Store-Backed AMI From an Existing AMI

To quickly and easily get a new working AMI, start with an existing public AMI or one of your own. You can then modify it and create a new AMI.

### Before You Get Started

1. If the AMI tools are not already installed, [install them \(p. 67\)](#)
2. Before you select an AMI, determine whether the instance types you plan to launch are 32-bit or 64-bit. For more information, see [Instance Types \(p. 94\)](#).
3. Make sure you are using GNU Tar 1.15 or later.
4. Install the Amazon EC2 API tools. Go to [Amazon EC2 API Tools](#) for more information and to download the tools from Amazon S3.

#### Note

To ensure you have the latest and most reliable version, we recommend that you install the Amazon EC2 API tools only from Amazon S3.

### Tasks to Use an Existing AMI to Create a New AMI

1. [Customize an Instance \(p. 69\)](#)
2. [Upload the Key and Certificate \(p. 69\)](#)
3. [Bundle a Customized Image \(Requires Root Privileges\) \(p. 70\)](#)

4. [Upload a Bundled AMI \(p. 71\)](#)
5. [Register the AMI \(p. 71\)](#)

## Customize an Instance

Customizing an instance involves the following series of steps:

1. Selecting an AMI from available AMIs.
2. Launching an instance from the AMI you selected.
3. Making changes to (thus, *customizing*) the instance, such as altering the Linux configuration, adding software, and configuring web applications.

For more information, see [Launch Your Instance \(p. 266\)](#).

After you've launched an instance according to your specifications, proceed to the next steps to create a new AMI using the customized instance.

## Upload the Key and Certificate

Your new AMI must be encrypted and signed to ensure that only you and Amazon EC2 can access it. To accomplish this, you must upload your Amazon EC2 private key and X.509 certificate to an instance store directory on your running instance. The private key and the certificate will be used in the AMI bundling process.

The private key and certificate files must not be bundled with the image. To prevent this, create a separate directory for these files. This directory will be specifically excluded from the bundle. In these examples, the private key and certificate files will be stored in the `/tmp/cert` directory. You can either use a secure file transfer program such as WinSCP to copy the files from your computer to your instance, or you can upload them directly. You must grant write permissions to the directory where these files will be uploaded. The following command grants write permission to `/tmp/cert`.

```
$ sudo chmod 777 /tmp/cert
```

### To upload your Amazon EC2 private key and X.509 certificate

Copy your Amazon EC2 private key and X.509 certificate to the instance using a secure copy function such as `scp`.

The following shows the syntax to use with the `scp` command.

```
$ scp -i <keypair_name> <private_keyfile> <certificate_file> <username>@<dns_location>:<instance_store_directory>
```

Where,

- `<keypair_name>` is the `.pem` file.
- `<private_keyfile>` is the file that contains the private key.
- `<certificate_file>` is the file that contains the certificate.
- `<username>` is the login name you use to log in to your instance; for example, `ec2-user`.
- `<dns_location>` is the DNS location of the instance within Amazon EC2.
- `<instance_store_directory>` is the directory where your instance store is mounted.

`scp` displays the names of the files copied and some performance statistics.



The following is an example of a fully specified `scp` command using the Amazon Linux AMI.

```
$ scp -i gsg-keypair pk-HKZYKTAIG2ECMXIYIBH3HXV4ZBEXAMPLE.pem cert-HKZYKTAIG2ECMXIYIBH3HXV4ZBEXAMPLE.pem ec2-user@ec2-203-0-113-25.compute-1.amazonaws.com:/tmp/cert
pk-HKZYKTAIG2ECMXIYIBH3HXV4ZBEXAMPLE.pem 100% 717 0.7KB/s 00:00
cert-HKZYKTAIG2ECMXIYIBH3HXV4ZBEXAMPLE.pem 100% 685 0.7KB/s 00:00
```

## Bundle a Customized Image (Requires Root Privileges)

When you have the image that meets your specifications, you need to bundle it for uploading to Amazon S3. The bundling process requires both the AMI Tools (to install the AMI Tools, click [here \(p. 67\)](#)) and root privileges. Root privileges can be obtained with the `sudo su` or `sudo su-` command:

```
sudo su -[root password]
```

### To bundle a customized image

Use the `ec2-bundle-vol` command. Make sure to exclude the directory where the private key and certificate files are stored with the `-e` option. This option excludes files that may contain sensitive information, such as your AWS credentials. By default, the bundle process excludes files that might contain sensitive information. These files include `*.sw`, `*.swo`, `*.swp`, `*.pem`, `*.priv`, `*id_rsa*`, `*id_dsa*`, `*.gpg`, `*.jks`, `*/.ssh/authorized_keys`, and `*/.bash_history`. To include all of these files, use the `--no-filter` option. To include some of these files, use the `--include` option.

```
# ec2-bundle-vol -k <private_keyfile> -c <certificate_file> -u <your_aws_ac
count_id> -e <cert_location>
```

- `<private_keyfile>` is the file that contains the private key.
- `<certificate_file>` is the file that contains the certificate.
- `<user_id>` is the ID associated with your AWS account. This is your AWS account ID without dashes. It consists of 12 to 15 characters, and it's *not* the same as your access key ID.
- `<cert_location>` is the directory that contains the private key and certificate files, which must be excluded from the bundle.

#### Note

If SELinux is enabled on your system, be sure to disable it before running `ec2-bundle-vol`.

Specify your architecture. In the following example, `x86_64` is used:

```
# -r x86_64
```

Execute the following command to bundle the local machine root file system.

```
# ec2-bundle-vol -e /tmp/cert -k pk-HKZYKTAIG2ECMXIYIBH3HXV4ZBEXAMPLE.pem -c
cert-HKZYKTAIG2ECMXIYIBH3HXV4ZBEXAMPLE.pem -u 111122223333
```

```
Please specify a value for arch [x86_64]:
Copying / into the image file /tmp/image...
```

```
Excluding:
/sys
...
/tmp/cert
...
1+0 records in
1+0 records out
1048576 bytes (1.0 MB) copied, 0.00172 s, 610 MB/s
mke2fs 1.42.3 (14-May-2012)
Bundling image file...
Splitting /tmp/image.tar.gz.enc...
Created image.part.00
Created image.part.01
...
Created image.part.NN
Generating digests for each part...
Digests generated.
Creating bundle manifest...
ec2-bundle-vol complete.
```

## Upload a Bundled AMI

You must upload the bundled AMI to Amazon S3 before Amazon EC2 can access it. This task is necessary when you create Amazon EC2 instance store-backed AMIs from an existing instance or from a loopback. Use the `ec2-upload-bundle` command to upload the bundled AMI that you created earlier. Amazon S3 stores data objects in buckets, which are similar to directories. All buckets must have globally unique names. The `ec2-upload-bundle` command uploads the bundled AMI to a specified bucket. If the specified bucket exists and belongs to another AWS account, the `ec2-upload-bundle` command will fail.

### Important

The specified Amazon S3 bucket must exist, and it must have been created in the same region as the instance being uploaded.

### To upload the bundled AMI

Use the `ec2-upload-bundle` command as follows:

```
$ ec2-upload-bundle -b <your-s3-bucket> -m <manifest_path> -a <access_key> -s <secret_key>
```

- `<your-s3-bucket>` is the Amazon S3 bucket that the bundle will be uploaded to. You can also upload the bundle to a subfolder of the bucket, such as `my-awsbucket/uploaded-images/image-1`. If the subfolder does not exist, it will be created.
- `<manifest_path>` is the full path to the manifest file (for example, `/tmp/image.manifest.xml`). The manifest file will reside in the destination directory that was specified in the `ec2-bundle-vol` command.
- `<access_key>` is your AWS access key ID.
- `<secret_key>` is your AWS secret key.

The AMI manifest file and all image parts are uploaded to Amazon S3. The manifest file is encrypted with the Amazon EC2 public key before being uploaded.

## Register the AMI

You must register your image with Amazon EC2, so that Amazon EC2 can locate it and run instances based on it. This task is necessary when you create Amazon EC2 instance store-backed AMIs from an

existing file or from a loopback. If you make any changes to the source image stored in Amazon S3, you must reregister the image.

### To register the AMI that you created and uploaded to Amazon S3

Use the `ec2-register` command (which is part of the EC2 CLI tools, not the AMI tools) as follows:

```
$ ec2-register <your-s3-bucket>/<path>/image.manifest.xml -n <image_name> -O  
<your_access_key> -W <your_secret_key>
```

#### Important

The capitalization of the bucket name and path in `<your-s3-bucket>/<path>` must match exactly what was passed in the `ec2-upload-bundle` command.

This command registers the AMI in the default region. To specify a different region, set the `EC2_URL` environment variable, or use the `--region` option with the `ec2-register` command.

Amazon EC2 returns an AMI identifier, the value next to the `IMAGE` tag, that you can use to run instances.

## Creating an Instance Store-Backed AMI From a Loopback

Creating AMIs through a loopback involves doing a full operating system installation on a clean root file system, but avoids having to create a new root disk partition and file system on a physical disk. After you have installed your operating system, you can bundle the resulting image as an AMI with the `ec2-bundle-image` command, which is part of the AMI tools (and not an API action). For more information about the `ec2-bundle-image` command and the AMI tools, go to the [Amazon Elastic Compute Cloud Command Line Reference](#).

#### Note

This method works only with AMIs that use instance stores for their root devices. This method is not applicable for AMIs backed by Amazon EBS.

### Before You Get Started

1	Before you select an AMI, determine whether the instance types you plan to launch are 32-bit or 64-bit. For more information, see <a href="#">Instance Types (p. 94)</a> .
2	Make sure you are using GNU Tar 1.15 or later.
3	This topic uses Fedora Core 4. Please make any adjustments for your distribution.

### Tasks to Create a New AMI Through a Loopback

1	Install Linux/UNIX and Prepare the System <ol style="list-style-type: none"><li><a href="#">Create a File to Host the AMI (p. 73)</a></li><li><a href="#">Create a Root File System Inside the File (p. 73)</a></li><li><a href="#">Mount the File through Loopback (p. 74)</a></li><li><a href="#">Prepare for the Installation (p. 75)</a></li><li><a href="#">Install the Operating System (p. 76)</a></li><li><a href="#">Configure the Operating System (p. 77)</a></li></ol>
2	<a href="#">Bundle the Loopback File Image (p. 78)</a>

3	<a href="#">Upload a Bundled AMI (p. 79)</a>
4	<a href="#">Register the AMI (p. 80)</a>

## Create a File to Host the AMI

The `dd` utility can create files of arbitrary sizes. Make sure to create a file large enough to host the operating system, tools, and applications that you will install. For example, a baseline Linux/UNIX installation requires about 700 MB, so your file should be at least 1 GB.

### To create a file to host the AMI

- Enter the following command:

```
# dd if=/dev/zero of=image_name bs=1M count=size
```

The `<image_name>` is the name of the image file you are creating and `<size>` is the size of the file in megabytes.

### Example

The following example creates a 1 GB file (1024\*1 MB).

```
# dd if=/dev/zero of=my-image.fs bs=1M count=1024  
1024+0 records in  
1024+0 records out
```

## Create a Root File System Inside the File

The `mkfs` utility has several variations that can create a file system inside the image file you are creating. Typical Linux/UNIX installations default to `ext2` or `ext3` file systems.

### To create an `ext3` file system

- Enter the following command:

```
# mke2fs -F -j <image_name>
```

The `<image_name>` is the name of the image file.

## Example

The following example creates an ext3 file system.

```
# mke2fs -F -j my-image.fs
mke2fs 1.38 (30-Jun-2005)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
131072 inodes, 262144 blocks
13107 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=268435456
8 block groups
32768 blocks per group, 32768 fragments per group
16384 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376

Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 24 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.
```

## Mount the File through Loopback

The loopback module enables you to use a normal file as if it were a raw device, which gives you a file system within a file. Mounting a file system image file through loopback presents it as part of the normal file system. You can then modify it using your favorite file management tools and utilities.

### To mount the file through loopback

1. Enter the following command to create a mount point in the file system where the image will be attached:

```
# mkdir <image_mountpoint>
```

The `<image_mountpoint>` is the location where the image will be mounted.

2. Mount the file system image:

```
# mount -o loop <image_name> <image_mountpoint>
```

The `<image_name>` is the name of the image file and `<image_mountpoint>` is the mount location.

## Example

The following commands create and mount the my-image.fs image file.

```
# mkdir /mnt/ec2-fs
# mount -o loop my-image.fs /mnt/ec2-fs
```

## Prepare for the Installation

Before the operating system installation can proceed, you must create and prepare the newly created root file system.

### To prepare for the installation

1. Create a `/dev` directory and populate it with a minimal set of devices. You can ignore the errors in the output.

```
# mkdir /mnt/ec2-fs/dev
# /sbin/MAKEDEV -d <image_mountpoint>/dev -x console
# /sbin/MAKEDEV -d <image_mountpoint>/dev -x null
# /sbin/MAKEDEV -d <image_mountpoint>/dev -x zero
```

The `<image_mountpoint>` is the mount location.

2. Create the `fstab` file within the `/etc` directory and add the following:

```
/dev/sda1 / ext3 defaults 1 1
none /dev/pts devpts gid=5,mode=620 0 0
none /dev/shm tmpfs defaults 0 0
none /proc proc defaults 0 0
none /sys sysfs defaults 0 0
```

3. Create a temporary YUM configuration file (e.g., `yum-xen.conf`) and add the following content.

```
[fedora]
name=Fedora $releasever - $basearch
failovermethod=priority
#baseurl=http://download.fedoraproject.org/pub/fedora/linux/releases/$relea
sever/Everything/$basearch/os/
mirrorlist=https://mirrors.fedoraproject.org/metalink?repo=fedora-$relea
sever&arch=$basearch
enabled=1
#metadata_expire=7d
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-fedora-$basearch

[updates]
name=Fedora $releasever - $basearch - Updates failovermethod=priority
#baseurl=http://download.fedoraproject.org/pub/fedora/linux/updates/$relea
sever/$basearch/
mirrorlist=https://mirrors.fedoraproject.org/metalink?repo=updates-released-
f$releasever&arch=$basearch
enabled=1
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-fedora-$basearch
```

This step ensures that all the required basic packages and utilities are installed. You can locate this file anywhere on your main file system (not on your loopback file system) and is used only during installation.

4. Enter the following:

```
# mkdir <image_mountpoint>/proc
# mount -t proc none <image_mountpoint>/proc
```

The `<image_mountpoint>` is the mount location. A `groupadd` utility bug in the `shadow-utils` package (versions prior to 4.0.7-7) requires you to mount the new `proc` file system manually with the preceding command.

## Example

These commands create the `/dev` directory and populate it with a minimal set of devices:

```
# mkdir /mnt/ec2-fs/dev
# /sbin/MAKEDEV -d /mnt/ec2-fs/dev -x console
MAKEDEV: mkdir: File exists
MAKEDEV: mkdir: File exists
MAKEDEV: mkdir: File exists
# /sbin/MAKEDEV -d /mnt/ec2-fs/dev -x null
MAKEDEV: mkdir: File exists
MAKEDEV: mkdir: File exists
MAKEDEV: mkdir: File exists
# /sbin/MAKEDEV -d /mnt/ec2-fs/dev -x zero
MAKEDEV: mkdir: File exists
MAKEDEV: mkdir: File exists
MAKEDEV: mkdir: File exists
```

This example creates and mounts the `/mnt/ec2-fs/proc` directory.

```
# mkdir /mnt/ec2-fs/proc
# mount -t proc none /mnt/ec2-fs/proc
```

## Install the Operating System

At this stage, the basic directories and files are created and you are ready to install the operating system. Depending on the speed of the host and network link to the repository, this process might take a while.

### To install the operating system

- Enter the following command:

```
# yum -c <yum_configuration_file> --installroot=<image_mountpoint> -y groupin
stall Base
```

The `<yum_configuration_file>` is the name of the YUM configuration file and `<image_mountpoint>` is the mount location.

You now have a base installation, which you can configure for operation inside Amazon EC2 and customize for your use.

## Example

This example installs the operating system at the `/mnt/ec2-fs` mount point using the `yum-xen.conf` YUM configuration file.

```
# yum -c yum-xen.conf --installroot=/mnt/ec2-fs -y groupinstall Base
Setting up Group Process
Setting up repositories
base                100% |=====| 1.1 kB    00:00
updates-released   100% |=====| 1.1 kB    00:00
comps.xml           100% |=====| 693 kB    00:00
comps.xml           100% |=====| 693 kB    00:00
Setting up repositories
Reading repository metadata in from local files
primary.xml.gz      100% |=====| 824 kB    00:00
base                : ##### 2772/2772
Added 2772 new packages, deleted 0 old in 15.32 seconds
primary.xml.gz      100% |=====| 824 kB    00:00
updates-re: ##### 2772/2772
Added 2772 new packages, deleted 0 old in 10.74 seconds
...
Complete!
```

## Configure the Operating System

After successfully installing the base operating system, you must configure your networking and hard drives to work in the Amazon EC2 environment.

### To configure the operating system

1. Edit (or create) `/mnt/ec2-fs/etc/sysconfig/network-scripts/ifcfg-eth0` and make sure it contains at least the following information:

```
DEVICE=eth0
BOOTPROTO=dhcp
ONBOOT=yes
TYPE=Ethernet
USERCTL=yes
PEERDNS=yes
IPV6INIT=no
```

#### Note

The Amazon EC2 DHCP server ignores hostname requests. If you set `DHCP_HOSTNAME`, the local hostname will be set on the instance but not externally. Additionally, the local hostname will be the same for all instances of the AMI, which might be confusing.

2. Verify that the following line appears in the `/mnt/ec2-fs/etc/sysconfig/network` file so that networking starts:

```
NETWORKING=yes
```

3. Add the following lines to `/mnt/ec2-fs/etc/fstab` so that local disk storage on `/dev/sda2` and swap space on `/dev/sda3` are mounted at system startup:



```
/dev/sda2 /mnt ext3 defaults 0 0
/dev/sda3 swap swap defaults 0 0
```

#### Note

The `/dev/sda2` and `/dev/sda3` storage locations only apply to small instances. For more information on instance storage, see [the section called “Instance Store” \(p. 508\)](#).

4. Allocate appropriate system run levels so that all your required services start at system startup. For example, to enable a service on multiuser and networked run levels, use the following commands:

```
# chroot /mnt/ec2-fs /bin/sh
# chkconfig --level 345 my-service on
# exit
```

Your new installation is successfully installed and configured to operate in the Amazon EC2 environment.

5. Enter the following commands to unmount the image:

```
# umount <image_mountpoint>/proc
# umount -d <image_mountpoint>
```

The `<image_mountpoint>` is the mount location.

#### Example

The following example unmounts the installation from the `/mnt/ec2-fs` mount point.

```
# umount /mnt/ec2-fs/proc
# umount -d /mnt/ec2-fs
```

## Bundle the Loopback File Image

### To bundle the loopback file image

- Enter the following command:

```
# ec2-bundle-image -i <image_name>.img -k <private_keyfile> -c <certificate_file> -u <user_id>
```

The `<image_name>` is the name of the image file, `<private_keyfile>` is the file that contains the private key, `<certificate_file>` is the file that contains the certificate, and `<user_id>` is the ID associated with your AWS account.

#### Note

The user ID is your AWS account ID without dashes. It consists of 12 to 15 characters, and it's *not* the same as your Access Key ID.

## Example

The `ec2-bundle-image` command bundles an image created in a loopback file.

```
# ec2-bundle-image -k pk-HKZYKTAIG2ECMXIYBH3HXV4ZBEXAMPLE.pem -c cert-HKZYK
TAIG2ECMXIYBH3HXV4ZBEXAMPLE.pem -u 111122223333 -i image.img -d bundled/ -p
fred -r x86_64
Please specify a value for arch [i386]:
Bundling image file...
Splitting bundled/fred.gz.crypt...
Created fred.part.00
Created fred.part.01
Created fred.part.02
Created fred.part.03
Created fred.part.04
Created fred.part.05
Created fred.part.06
Created fred.part.07
Created fred.part.08
Created fred.part.09
Created fred.part.10
Created fred.part.11
Created fred.part.12
Created fred.part.13
Created fred.part.14
Generating digests for each part...
Digests generated.
Creating bundle manifest...
ec2-bundle-image complete.
```

## Upload a Bundled AMI

You must upload the bundled AMI to Amazon S3 before Amazon EC2 can access it. This task is necessary when you create Amazon EC2 instance store-backed AMIs from an existing instance or from a loopback. Use the `ec2-upload-bundle` command to upload the bundled AMI that you created earlier. Amazon S3 stores data objects in buckets, which are similar to directories. All buckets must have globally unique names. The `ec2-upload-bundle` command uploads the bundled AMI to a specified bucket. If the specified bucket exists and belongs to another AWS account, the `ec2-upload-bundle` command will fail.

### Important

The specified Amazon S3 bucket must exist, and it must have been created in the same region as the instance being uploaded.

## To upload the bundled AMI

Use the `ec2-upload-bundle` command as follows:

```
$ ec2-upload-bundle -b <your-s3-bucket> -m <manifest_path> -a <access_key> -s
<secret_key>
```

- `<your-s3-bucket>` is the Amazon S3 bucket that the bundle will be uploaded to. You can also upload the bundle to a subfolder of the bucket, such as `my-awsbucket/uploaded-images/image-1`. If the subfolder does not exist, it will be created.
- `<manifest_path>` is the full path to the manifest file (for example, `/tmp/image.manifest.xml`). The manifest file will reside in the destination directory that was specified in the `ec2-bundle-vol` command.
- `<access_key>` is your AWS access key ID.

- `<secret_key>` is your AWS secret key.

The AMI manifest file and all image parts are uploaded to Amazon S3. The manifest file is encrypted with the Amazon EC2 public key before being uploaded.

## Register the AMI

You must register your image with Amazon EC2, so that Amazon EC2 can locate it and run instances based on it. This task is necessary when you create Amazon EC2 instance store-backed AMIs from an existing file or from a loopback. If you make any changes to the source image stored in Amazon S3, you must reregister the image.

### To register the AMI that you created and uploaded to Amazon S3

Use the `ec2-register` command (which is part of the EC2 CLI tools, not the AMI tools) as follows:

```
$ ec2-register <your-s3-bucket>/<path>/image.manifest.xml -n <image_name> -O  
<your_access_key> -W <your_secret_key>
```

#### Important

The capitalization of the bucket name and path in `<your-s3-bucket>/<path>` must match exactly what was passed in the `ec2-upload-bundle` command.

This command registers the AMI in the default region. To specify a different region, set the `EC2_URL` environment variable, or use the `--region` option with the `ec2-register` command.

Amazon EC2 returns an AMI identifier, the value next to the `IMAGE` tag, that you can use to run instances.

## Creating and Launching an AMI from a Snapshot

If you have a snapshot of the root device volume of an instance, you can terminate that instance and then later launch a new instance from the snapshot. You must first register the snapshot, then create and launch the resulting AMI, as explained in [Launching an Instance from a Backup \(p. 271\)](#).

## Using Your Own Linux Kernels

To enable user-provided kernels on EC2 instances, Amazon has published Amazon Kernel Images (AKIs) that use a system called *PV-GRUB*. PV-GRUB is a paravirtual boot loader that runs a patched version of GNU GRUB 0.97. When you start an instance, PV-GRUB loads the kernel specified by your image's `menu.lst` file.

PV-GRUB understands standard `grub.conf` or `menu.lst` commands, which allows it to work with all currently supported Linux distributions. Older distributions such as Ubuntu 10.04 LTS, Oracle Enterprise Linux or CentOS 5.x require a special "ec2" or "xen" kernel package, while newer distributions include the required drivers in the default kernel package.

Most modern paravirtual AMIs use a PV-GRUB AKI by default (including all of the paravirtual Linux AMIs available in the Amazon EC2 Launch Wizard Quick Start menu), so there are no additional steps that you need to take to use a different kernel on your instance, provided that the kernel you want to use is compatible with your distribution. You can verify that the kernel image for an AMI is a PV-GRUB AKI by executing the following command with the Amazon EC2 command line tools (substituting the kernel image ID you want to check):

```
$ ec2-describe-images -a -F image-id=aki-880531cd  
IMAGE    aki-880531cd    amazon/pv-grub-hd0_1.04-x86_64.gz ...
```

The `name` field of the output should contain `pv-grub`.

### Topics

- [Limitations of PV-GRUB \(p. 81\)](#)
- [Configuring GRUB \(p. 81\)](#)
- [Amazon PV-GRUB Kernel Image IDs \(p. 82\)](#)

## Limitations of PV-GRUB

PV-GRUB has the following limitations:

- You can't use the 64-bit version of PV-GRUB to start a 32-bit kernel or vice versa.
- You can't specify an Amazon ramdisk image (ARI) when using a PV-GRUB AKI.
- AWS has tested and verified that PV-GRUB works with these file system formats: EXT2, EXT3, EXT4, JFS, XFS, and ReiserFS. Other file system formats might not work.
- PV-GRUB can boot kernels compressed using the `gzip`, `bzip2`, `lzo`, and `xz` compression formats.
- Cluster AMIs don't support or need PV-GRUB, because they use full hardware virtualization (HVM). While paravirtual instances use PV-GRUB to boot, HVM instance volumes are treated like actual disks, and the boot process is similar to the boot process of a bare metal operating system with a partitioned disk and bootloader.
- PV-GRUB versions 1.03 and earlier don't support GPT partitioning, they support MBR partitioning only.
- If you plan to use a logical volume manager (LVM) with Amazon EBS volumes, you need a separate boot partition outside of the LVM. Then you can create logical volumes with the LVM.

## Configuring GRUB

To boot PV-GRUB, a `GRUB menu.lst` file must exist in the image; the most common location for this file is `/boot/grub/menu.lst`.

The following is an example of a `menu.lst` configuration file for booting an AMI with a PV-GRUB AKI. In this example, there are two kernel entries to choose from: `Amazon Linux 2013.09` (the original kernel for this AMI), and `Vanilla Linux 3.11.6` (a newer version of the Vanilla Linux kernel from <https://www.kernel.org/>). The `Vanilla` entry was copied from the original entry for this AMI, and the `kernel` and `initrd` paths were updated to the new locations. The `default 0` parameter points the boot loader to the first entry it sees (in this case, the `Vanilla` entry), and the `fallback 1` parameter points the bootloader to the next entry if there is a problem booting the first.

```
default 0
fallback 1
timeout 0
hiddenmenu

title Vanilla Linux 3.11.6
root (hd0)
kernel /boot/vmlinuz-3.11.6 root=LABEL=/ console=hvc0
initrd /boot/initrd.img-3.11.6

title Amazon Linux 2013.09 (3.4.62-53.42.amzn1.x86_64)
root (hd0)
kernel /boot/vmlinuz-3.4.62-53.42.amzn1.x86_64 root=LABEL=/ console=hvc0
initrd /boot/initramfs-3.4.62-53.42.amzn1.x86_64.img
```

You don't need to specify a fallback kernel in your `menu.lst` file, but we recommend that you have a fallback when you test a new kernel. PV-GRUB can fall back to another kernel in the event that the new kernel fails. Having a fallback kernel allows the instance to boot even if the new kernel isn't found.

PV-GRUB checks the following locations for `menu.lst`, using the first one it finds:

- `(hd0)/boot/grub`
- `(hd0,0)/boot/grub`
- `(hd0,0)/grub`
- `(hd0,1)/boot/grub`
- `(hd0,1)/grub`
- `(hd0,2)/boot/grub`
- `(hd0,2)/grub`
- `(hd0,3)/boot/grub`
- `(hd0,3)/grub`

Note that PV-GRUB 1.03 and earlier only check one of the first two locations in this list.

## Amazon PV-GRUB Kernel Image IDs

PV-GRUB AKIs are available in all Amazon EC2 regions. There are AKIs for both 32-bit and 64-bit architecture types. Most modern AMIs use a PV-GRUB AKI by default.

We recommend that you always use the latest version of the PV-GRUB AKI, as not all versions of the PV-GRUB AKI are compatible with all instance types. Use the following command to get a list of the PV-GRUB AKIs for the current region:

```
$ ec2-describe-images -o amazon --filter "name=pv-grub-*.gz"
```

Note that PV-GRUB is the only AKI available in the `ap-southeast-2` region. You should verify that any AMI you want to copy to this region is using a version of PV-GRUB that is available in this region.

The following are the current AKI IDs for each region. There is no longer any difference between the `hd0` and `hd00` AKIs, though we continue to provide both naming schemes for backward compatibility. You should register new AMIs using an `hd0` AKI.

### ap-northeast-1, Asia Pacific (Tokyo) Region

Image ID	Image Name
<code>aki-136bf512</code>	<code>pv-grub-hd0_1.04-i386.gz</code>
<code>aki-176bf516</code>	<code>pv-grub-hd0_1.04-x86_64.gz</code>
<code>aki-196bf518</code>	<code>pv-grub-hd00_1.04-i386.gz</code>
<code>aki-1f6bf51e</code>	<code>pv-grub-hd00_1.04-x86_64.gz</code>

### ap-southeast-1, Asia Pacific (Singapore) Region

Image ID	Image Name
<code>aki-ae3973fc</code>	<code>pv-grub-hd0_1.04-i386.gz</code>

Image ID	Image Name
aki-503e7402	pv-grub-hd0_1.04-x86_64.gz
aki-563e7404	pv-grub-hd00_1.04-i386.gz
aki-5e3e740c	pv-grub-hd00_1.04-x86_64.gz

### ap-southeast-2, Asia Pacific (Sydney) Region

Image ID	Image Name
aki-cd62fff7	pv-grub-hd0_1.04-i386.gz
aki-c362fff9	pv-grub-hd0_1.04-x86_64.gz
aki-c162fffb	pv-grub-hd00_1.04-i386.gz
aki-3b1d8001	pv-grub-hd00_1.04-x86_64.gz

### eu-west-1, EU (Ireland) Region

Image ID	Image Name
aki-68a3451f	pv-grub-hd0_1.04-i386.gz
aki-52a34525	pv-grub-hd0_1.04-x86_64.gz
aki-5ea34529	pv-grub-hd00_1.04-i386.gz
aki-58a3452f	pv-grub-hd00_1.04-x86_64.gz

### sa-east-1, South America (Sao Paulo) Region

Image ID	Image Name
aki-5b53f446	pv-grub-hd0_1.04-i386.gz
aki-5553f448	pv-grub-hd0_1.04-x86_64.gz
aki-5753f44a	pv-grub-hd00_1.04-i386.gz
aki-5153f44c	pv-grub-hd00_1.04-x86_64.gz

### us-east-1, US East (Northern Virginia) Region

Image ID	Image Name
aki-8f9dcae6	pv-grub-hd0_1.04-i386.gz
aki-919dcaf8	pv-grub-hd0_1.04-x86_64.gz
aki-659ccb0c	pv-grub-hd00_1.04-i386.gz
aki-499ccb20	pv-grub-hd00_1.04-x86_64.gz

### us-gov-west-1, AWS GovCloud (US)

Image ID	Image Name
aki-1fe98d3c	pv-grub-hd0_1.04-i386.gz
aki-1de98d3e	pv-grub-hd0_1.04-x86_64.gz
aki-63e98d40	pv-grub-hd00_1.04-i386.gz
aki-61e98d42	pv-grub-hd00_1.04-x86_64.gz

### us-west-1, US West (Northern California) Region

Image ID	Image Name
aki-8e0531cb	pv-grub-hd0_1.04-i386.gz
aki-880531cd	pv-grub-hd0_1.04-x86_64.gz
aki-960531d3	pv-grub-hd00_1.04-i386.gz
aki-920531d7	pv-grub-hd00_1.04-x86_64.gz

### us-west-2, US West (Oregon) Region

Image ID	Image Name
aki-f08f11c0	pv-grub-hd0_1.04-i386.gz
aki-fc8f11cc	pv-grub-hd0_1.04-x86_64.gz
aki-e28f11d2	pv-grub-hd00_1.04-i386.gz
aki-e68f11d6	pv-grub-hd00_1.04-x86_64.gz

## Copying AMIs

You can easily copy the Amazon Machine Images (AMIs) that you own to other AWS regions and scale your applications to take advantage of AWS's geographically diverse regions.

Copying your AMIs provides the following benefits:

- **Consistent global deployment:** You can copy an AMI from one region to another, enabling you to launch consistent instances based from the same AMI into different regions.
- **Scalability:** You can more easily design and build world-scale applications that meet the needs of your users, regardless of their location.
- **Performance:** You can increase performance by distributing your application, as well as locating critical components of your application in closer proximity to your users. You can also take advantage of region-specific features, such as instance types or other AWS services.
- **High availability:** You can design and deploy applications across AWS regions, to increase availability.

## AMI Copy

You can copy an AMI to as many regions as you like, using the AWS Management Console, the Amazon EC2 CLI, or the Amazon EC2 API. You can copy an AMI to the same region. You can copy both Amazon EBS-backed AMIs and instance-store-backed AMIs.

There are no charges for copying an AMI. However, standard storage and data transfer rates apply.

Each copy of an AMI results in a new AMI with its own unique AMI ID. The new AMI is fully independent of the source AMI; there is no link to the original (source) AMI. You can modify the new AMI without affecting the source AMI. The reverse is also true: you can modify the source AMI without affecting the new AMI. Therefore, if you make changes to the source AMI and want those changes to be reflected in the AMI in the destination region, you must recopy the source AMI to the destination region.

We don't copy launch permissions, user-defined tags, or Amazon S3 bucket permissions from the source AMI to the new AMI. After the copy operation is complete, you can apply launch permissions, user-defined tags, and Amazon S3 bucket permissions to the new AMI.

We try to find matching AKIs and ARIs for the new AMI in the destination region. If we can't find a matching AKI or ARI, then we don't copy the AMI. If you are using the AKIs and ARIs that we recommend, the copy operation registers the AMI with the appropriate AKI and ARI in the destination region. If you get an error message "Failed to find matching AKI/ARI", it means that the destination region doesn't contain an AKI or ARI that matches those specified in the source AMI. If your AMI uses a PV-GRUB AKI, then you can update the AMI to leverage the latest version of PV-GRUB. For more information on PV-GRUB and AKIs, see [Using Your Own Linux Kernels \(p. 80\)](#).

## Copying an Amazon EC2 AMI

Prior to copying an AMI, you must ensure that the contents of the source AMI are updated to support running in a different region. For example, you should update any database connection strings or similar application configuration data to point to the appropriate resources. Otherwise, instances launched from the new AMI in the destination region may still use the resources from the source region, which can impact performance and cost.

## AWS Management Console

### To copy an AMI using the console

1. Open the Amazon EC2 console.
2. From the navigation bar, select the region that contains the AMI to copy.
3. In the navigation pane, click **AMIs**.
4. Select the AMI to copy, click **Actions**, and then click **Copy AMI**.
5. In the **AMI Copy** page, set the following fields, and then click **Copy AMI**:
  - **Destination region:** Select the region to which you want to copy the AMI.
  - **Name:** Specify a name for the new AMI.
  - **Description:** By default, the description includes information about the source AMI so that you can identify a copy from the original. You can change this description as necessary.
6. We display a confirmation page to let you know that the copy operation has been initiated and provide you with the ID of the new AMI.



To check on the progress of the copy operation immediately, click the provided link to switch to the destination region. To check on the progress later, click **Done**, and then when you are ready, use the navigation pane to switch to the destination region.

The initial status of the destination AMI is `pending` and the operation is complete when the status is `available`.

## Command Line Interface

### To copy an AMI using the CLI

You can copy an AMI using the `ec2-copy-image` command. This command initiates the copy operation and registers the new AMI in the destination region.

This command is submitted to and initiated from the destination region endpoint.

## API

### To copy an AMI using the API

You can copy an AMI using `CopyImage`.

This call is submitted to and initiated from the destination region endpoint.

## Stopping a Pending AMI Copy Operation

### AWS Management Console

#### To stop an AMI copy operation using the console

1. Open the Amazon EC2 console.
2. From the navigation bar, select the destination region from the region selector.
3. In the navigation pane, click **AMIs**.
4. Select the AMI you want to stop copying, click **Actions**, and then click **Deregister**.
5. When asked for confirmation, click **Continue**.

## Command Line Interface

### To stop the AMI copy operation using the CLI

You can stop an AMI copy operation using the `ec2-deregister` command.

This command stops the copy operation and deregisters the new AMI in the destination region.

## API

### To stop the AMI copy operation using the API

You can stop an AMI copy operation using `DeregisterImage`.

This action stops the copy operation and deregisters the new AMI in the destination region.

# Amazon Linux

Amazon Linux is provided by Amazon Web Services (AWS). It is designed to provide a stable, secure, and high-performance execution environment for applications running on Amazon EC2. It also includes packages that enable easy integration with AWS, including launch configuration tools and many popular AWS libraries and tools. AWS provides ongoing security and maintenance updates to all instances running Amazon Linux.

To launch an Amazon Linux instance, use an Amazon Linux AMI. AWS provides Amazon Linux AMIs to Amazon EC2 users at no additional cost.

## Topics

- [Finding the Amazon Linux AMI \(p. 87\)](#)
- [Launching and Connecting to an Amazon Linux Instance \(p. 87\)](#)
- [Identifying Amazon Linux AMI Images \(p. 88\)](#)
- [Included AWS Command Line Tools \(p. 88\)](#)
- [cloud-init \(p. 89\)](#)
- [Repository Configuration \(p. 90\)](#)
- [Adding Packages \(p. 91\)](#)
- [Accessing Source Packages for Reference \(p. 91\)](#)
- [Developing Applications \(p. 92\)](#)
- [Instance Store Access \(p. 92\)](#)
- [Product Life Cycle \(p. 92\)](#)
- [Security Updates \(p. 92\)](#)
- [Support \(p. 93\)](#)

## Finding the Amazon Linux AMI

For a list of the latest Amazon Linux AMIs, see [Amazon Linux AMIs](#).

## Launching and Connecting to an Amazon Linux Instance

After locating your desired AMI, note the AMI ID. You can use the AMI ID to launch and then connect to your instance.

Amazon Linux does not allow remote root SSH by default. Also, password authentication is disabled to prevent brute-force password attacks. To enable SSH logins to an Amazon Linux instance, you must provide your key pair to the instance at launch. You must also set the security group used to launch your instance to allow SSH access. By default, the only account that can log in remotely using SSH is `ec2-user`; this account also has `sudo` privileges. If you want to enable remote root log in, please be aware that it is less secure than relying on key pairs and a secondary user.

For information on launching and using your Amazon Linux instance, see [Launch Your Instance \(p. 266\)](#). For information on connecting to your Amazon Linux instance, see [Connecting to Your Linux/UNIX Instance \(p. 279\)](#).

## Identifying Amazon Linux AMI Images

Each image contains a unique `/etc/image-id` that identifies the AMI. This file contains information about the image.

The following is an example of the `/etc/image-id` file:

```
[ec2-user@ip-10-159-3-200 ~]$ cat /etc/image-id
image_name="amzn-ami-pv"
image_version="2013.09"
image_arch="x86_64"
image_file="amzn-ami-pv-2013.09.0.x86_64.ext4"
image_stamp="3373-c521"
image_date="20130925011245"
recipe_name="amzn ami"
recipe_id="55ae74f2-622a-de5d-2288-c450-6402-3ea3-a8b3b196"
```

The `image_name`, `image_version`, and `image_arch` items come from the build recipe that Amazon used to construct the image. The `image_stamp` is simply a unique random hex value generated during image creation. The `image_date` item is in `YYYYMMDDhhmmss` format, and is the UTC time of image creation. The `recipe_name` and `recipe_id` refer to the name and ID of the build recipe Amazon used to construct the image, which identifies the current running version of Amazon Linux. This file will not change as you install updates from the yum repository.

Amazon Linux contain a `/etc/system-release` file that specifies the current release that is installed. This file is updated through yum and is part of the system-release rpm.

The following is an example of a `/etc/system-release` file:

```
# cat /etc/system-release
Amazon Linux AMI release 2013.09
```

Amazon Linux also contains a machine readable version of the `/etc/system-release` file found in `/etc/system-release-cpe` and follows the CPE specification from MITRE ([CPE](#)).

## Included AWS Command Line Tools

The following popular command line tools for AWS integration and usage have been included in Amazon Linux or in the repositories:

- `aws-amitools-ec2`
- `aws-apitools-as`
- `aws-apitools-cfn`
- `aws-apitools-common`
- `aws-apitools-ec2`
- `aws-apitools-elb`
- `aws-apitools-iam`
- `aws-apitools-mon`
- `aws-apitools-rds`
- `aws-cfn-bootstrap`
- `aws-cli`
- `aws-scripts-ses`

To simplify the configuration of these tools, a simple script has been included to prepare `AWS_CREDENTIAL_FILE`, `JAVA_HOME`, `AWS_PATH`, `PATH`, and product-specific environment variables after a credential file has been installed.

Also, to allow the installation of multiple versions of the API and AMI tools, we have placed symlinks to the desired versions of these tools in `/opt/aws`, as described here:

#### **`/opt/aws/bin`**

Symlink farm to `/bin` directories in each of the installed tools directories.

#### **`/opt/aws/{apitools|amitools}`**

Products are installed in directories of the form `[name]-version` and symlink `[name]` attached to the most recently installed version.

#### **`/opt/aws/{apitools|amitools}/[name]/environment.sh`**

Used by `/etc/profile.d/aws-apitools-common.sh` to set product-specific environment variables, such as `EC2_HOME`.

## cloud-init

The cloud-init package is an open source application built by Canonical that is used to bootstrap Linux images in a cloud computing environment, such as Amazon EC2. Amazon Linux contains a customized version of cloud-init. It enables you to specify actions that should happen to your instance at boot time. You can pass desired actions to cloud-init through the user data fields when launching an instance. This means you can use common AMIs for many use cases and configure them dynamically at startup. Amazon Linux also uses cloud-init to perform initial configuration of the `ec2-user` account.

For more information about cloud-init, see <https://help.ubuntu.com/community/CloudInit>.

Amazon Linux uses the following cloud-init actions (configurable in `/etc/sysconfig/cloudinit`):

- action: INIT (always runs)
  - Setting a default locale.
  - Setting the hostname.
  - Parsing and handling user data.
- action: CONFIG\_SSH
  - Generating host private SSHkeys.
  - Adding user's public SSHkeys to `.ssh/authorized_keys` for easy login and administration.
- action: PACKAGE\_SETUP
  - Preparing yum repo.
  - Handles package actions defined in user data.
- action: RUNCMD
  - Runs a shell command.
- action: RUN\_USER\_SCRIPTS
  - Executes user scripts found in user data.
- action: CONFIG\_MOUNTS
  - Mounts ephemeral drives.
- action: CONFIG\_LOCALE
  - Sets the locale in the locale config file according to user data.

## Supported User-Data Formats

The cloud-init package supports user-data handling of a variety of formats:

- Gzip
  - If user-data is gzip compressed, cloud-init will decompress the data and handle as appropriate.
- MIME multipart
  - Using a MIME multipart file, you can specify more than one type of data. For example, you could specify both a user-data script and a cloud-config type. Each part of the multipart file can be handled by cloud-init if it is one of the supported formats.
- Base64 decoding
  - If user-data is base64 encoded, cloud-init determines if it can understand the decoded data as one of the supported types. If it understands the decoded data, it will decode the data and handle as appropriate. If not, it returns the base64 data intact.
- User-Data script
  - Begins with "#!" or "Content-Type: text/x-shellsript".
  - The script will be executed by "/etc/init.d/cloud-init-user-scripts" level during first boot. This occurs late in the boot process (after the initial configuration actions were performed).
- Include file
  - Begins with "#include" or "Content-Type: text/x-include-url".
  - This content is an include file. The file contains a list of URLs, one per line. Each of the URLs will be read, and their content will be passed through this same set of rules. The content read from the URL can be gzipped, MIME-multi-part, or plain text.
- Cloud Config Data
  - Begins with "#cloud-config" or "Content-Type: text/cloud-config".
  - This content is cloud-config data. See the examples for a commented example of supported config formats.
- Cloud Boothook
  - Begins with "#cloud-boothook" or "Content-Type: text/cloud-boothook".
  - This content is boothook data. It is stored in a file under /var/lib/cloud and then executed immediately.
  - This is the earliest "hook" available. Note that there is no mechanism provided for running it only once. The boothook must take care of this itself. It is provided with the instance ID in the environment variable INSTANCE\_ID. Use this variable to provide a once-per-instance set of boothook data.

## Repository Configuration

Beginning with the 2011.09 release of Amazon Linux, Amazon Linux AMIs are treated as snapshots in time, with a repository and update structure that always gives you the latest packages when you run `yum update -y`.

The repository structure is configured to deliver a continuous flow of updates that allow you to roll from one version of Amazon Linux to the next. For example, if you launch an instance from an older version of the Amazon Linux AMI (such as 2013.03 or earlier) and run `yum update -y`, you end up with the latest packages.

You can disable rolling updates for Amazon Linux by enabling the *lock-on-launch* feature. The lock-on-launch feature locks your newly launched instance to receive updates only from the specified release of the AMI. For example, you can launch a 2012.03 AMI and have it receive only the updates that were released prior to the 2013.09 AMI, until you are ready to migrate to the 2013.09 AMI. To enable lock-on-launch in new instances, launch it with the following user data passed to cloud-init, using either the Amazon EC2 console or the `ec2-run-instances` command with the `-f` flag.

```
#cloud-config
repo_releasever:2013.09
```

### To lock existing instances to their current AMI release version

1. Edit `/etc/yum.conf`.
2. Comment out `releasever=latest`.
3. Run **yum clean all** to clear the cache.

## Adding Packages

Amazon Linux is designed to be used with online package repositories hosted in each Amazon EC2 region. These repositories provide ongoing updates to packages in the Amazon Linux AMI, as well as access to hundreds of additional common open source server applications. The repositories are available in all regions and are accessed using yum update tools, as well as on the [Linux AMI packages site](#). Hosting repositories in each region enables us to deploy updates quickly and without any data transfer charges. The packages can be installed by issuing yum commands, such as the following example:

```
# sudo yum install httpd
```

Access to the Extra Packages for Enterprise Linux (EPEL) repository is configured, but it is not enabled by default. EPEL provides third-party packages in addition to those that are in the Amazon Linux repositories. The third-party packages are not supported by AWS.

If you find that Amazon Linux does not contain an application you need, you can simply install the application directly on your Amazon Linux instance. Amazon Linux uses RPM and yum for package management, and that is likely the simplest way to install new applications. You should always check to see if an application is available in our central Amazon Linux repository first, because many applications are available there. These applications can easily be added to your Amazon Linux instance.

To upload your applications onto a running Amazon Linux instance, use `scp` or `sftp` and then configure the application by logging on to your instance. Your applications can also be uploaded during the instance launch by using the `PACKAGE_SETUP` action from built-in the cloud-init package. For more information, see [cloud-init](#) (p. 89).

#### Important

If your instance is running in a virtual private cloud (VPC), you must attach an Internet Gateway to the VPC in order to contact the yum repository. For more information, see [Internet Gateways](#) in the *Amazon Virtual Private Cloud User Guide*.

## Accessing Source Packages for Reference

You can view the source of packages you have installed on your instance for reference purposes by using tools provided in Amazon Linux. Source packages are available for all of the packages included in Amazon Linux and the online package repository. Simply determine the package name for the source package you want to install and use the `get_reference_source` command to view source within your running instance. For example:

```
# get_reference_source -p httpd
```

The following is a sample response:

```
# get_reference_source -p httpd
# get_reference_source -p httpd

Requested package: httpd
```

```
Found package from local RPM database: httpd-2.2.25-1.0.amzn1.x86_64
Corresponding source RPM to found package: httpd-2.2.25-1.0.amzn1.src.rpm

Are these parameters correct? Please type 'yes' to continue: yes
Source RPM downloaded to: /usr/src/srpm/debug/ httpd-2.2.25-1.0.amzn1.src.rpm
```

The source RPM will be placed in the `/usr/src/srpm/debug` directory of your instance. From there it can be unpacked, and, for reference, you can view the source tree using standard RPM tools. After you finish debugging, the package will be available for use.

### Important

If your instance is running in a virtual private cloud (VPC), you must attach an Internet Gateway to the VPC in order to contact the yum repository. For more information, see [Internet Gateways](#) in the *Amazon Virtual Private Cloud User Guide*.

## Developing Applications

A full set of Linux development tools is provided in the yum repository for Amazon Linux. To develop applications on Amazon Linux, simply select the development tools you need with yum. Alternatively, many applications developed on CentOS and other similar distributions should run on Amazon Linux.

## Instance Store Access

The instance store drive `ephemeral0` is mounted in `/media/ephemeral0` only on Amazon instance store-backed AMIs. This is different than many other images that mount the instance store drive under `/mnt`.

## Product Life Cycle

The Amazon Linux AMI is updated regularly with security and feature enhancements. If you do not need to preserve data or customizations on your Amazon Linux instances, you can simply relaunch new instances with the latest Amazon Linux AMI. If you need to preserve data or customizations for your Amazon Linux instances, you can maintain those instances through the Amazon Linux yum repositories. The yum repositories contain all the updated packages. You can choose to apply these updates to your running instances.

Older versions of the AMI and update packages will continue to be available for use, even as new versions are released. However, in some cases, if you're seeking support for an older version of Amazon Linux; through AWS Support, we might ask you to move to newer versions as part of the support process.

## Security Updates

Security updates are provided via the Amazon Linux AMI yum repositories as well as via updated Amazon Linux AMIs. Security alerts will be published in the [Amazon Linux AMI Security Center](#). For more information on AWS security policies or to report a security problem, visit the [AWS Security Center](#).

Amazon Linux AMIs are configured to download and install security updates at launch time. This is controlled via a cloud-init setting called `repo_upgrade`. The following snippet of cloud-init configuration shows how you can change the settings in the user data text you pass to your instance initialization:

```
#cloud-config
repo_upgrade:security
```

The possible values for the `repo_upgrade` setting are as follows:

**security**

Apply outstanding updates that Amazon marks as security updates.

**bugfix**

Apply updates that Amazon marks as bug fixes. Bug fixes are a larger set of updates, which include security updates and fixes for various other minor bugs.

**all**

Apply all applicable available updates, regardless of their classification.

**none**

Do not apply any updates to the instance on startup.

The default setting for `repo_upgrade` is `security`. That is, if you don't specify a different value in your user data, by default the Amazon Linux AMI will perform the security upgrades at launch for any packages installed at that time. Amazon Linux AMI will also notify you of any updates to the installed packages by listing the number of available updates upon login using the `motd`. To install these updates, you will need to run `sudo yum upgrade` on the instance.

**Important**

If your instance is running in a virtual private cloud (VPC), you must attach an Internet Gateway to the VPC in order to contact the yum repository. For more information, see [Internet Gateways](#) in the *Amazon Virtual Private Cloud User Guide*.

## Support

Support for installation and use of the base Amazon Linux AMI is included through subscriptions to AWS Support. For more information, see [AWS Support](#).

We encourage you to post any questions you have about Amazon Linux to the [Amazon EC2 forum](#).



# Amazon EC2 Instances

---

If you're new to Amazon EC2, see the following topics to get started:

- [What is Amazon EC2? \(p. 1\)](#)
- [Get Set Up for Amazon EC2 \(p. 18\)](#)
- [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#)
- [Getting Started with Amazon EC2 Windows Instances](#)
- [Instance Lifecycle \(p. 263\)](#)

Before you launch a production environment, you need to answer the following questions.

**Q. What purchasing option best meets my needs?**

Amazon EC2 supports On-Demand Instances (the default), [Spot Instances \(p. 115\)](#), and [Reserved Instances \(p. 193\)](#). For more information, see [Amazon EC2 Pricing](#).

**Q. What instance type best meets my needs?**

Amazon EC2 provides different instance types to enable you to choose the CPU, memory, storage, and networking capacity that you need to run your applications. For more information, see [Instance Types \(p. 94\)](#).

**Q. Which type of root volume meets my needs?**

Each instance is backed by Amazon EBS or backed by instance store. Select an AMI based on which type of root volume you need. For more information, see [Storage for the Root Device \(p. 45\)](#).

**Q. Would I benefit from using a virtual private cloud?**

If you can launch instances in either EC2-Classic or EC2-VPC, you'll need to decide which platform meets your needs. For more information, see [Supported Platforms \(p. 417\)](#) and [Amazon EC2 and Amazon Virtual Private Cloud \(VPC\) \(p. 414\)](#).

## Instance Types

When you launch an instance, the *instance type* that you specify determines the hardware of the host computer used for your instance. Each instance type offers different compute, memory, and storage capabilities. Select an instance type based on the requirements of the application or software that you plan to run on your instance.

Amazon EC2 provides each instance with a consistent and predictable amount of CPU capacity, regardless of its underlying hardware.

Amazon EC2 dedicates some resources of the host computer, such as CPU, memory, and instance storage, to a particular instance. Amazon EC2 shares other resources of the host computer, such as the network and the disk subsystem, among instances. If each instance on a host computer tries to use as much of one of these shared resources as possible, each receives an equal share of that resource. However, when a resource is under-utilized, an instance can consume a higher share of that resource while it's available.

Each instance type provides higher or lower minimum performance from a shared resource, depending on their size. For example, instance types with high I/O performance have a larger allocation of shared resources. Allocating a larger share of shared resources also reduces the variance of I/O performance. For most applications, moderate I/O performance is more than enough. However, for applications that require greater or more consistent I/O performance, consider an instance type with higher I/O performance.

To obtain additional, dedicated capacity for EBS I/O, you can launch some instance types as EBS-optimized instances. For more information, see [EBS-Optimized Instances \(p. 107\)](#).

To optimize your instances for high performance computing (HPC) applications, you can launch some instance types in a placement group. For more information, see [Placement Groups \(p. 108\)](#).

## Available Instance Types

Amazon EC2 provides the instance types listed in the following table. For more information about the specifications for each Amazon EC2 instance type, see [Instance Type Details](#).

Instance Family	Instance Types
General purpose	m1.small   m1.medium   m1.large   m1.xlarge   m3.xlarge   m3.2xlarge
Compute optimized	c1.medium   c1.xlarge   cc2.8xlarge
Memory optimized	m2.xlarge   m2.2xlarge   m2.4xlarge   cr1.8xlarge
Storage optimized	h1.4xlarge   hs1.8xlarge
Micro instances	t1.micro
GPU instances	cg1.4xlarge   g2.2xlarge

To determine which instance type best meets your needs, we recommend that you launch an instance and use your own benchmark application. Because you pay by the instance hour, it's convenient and inexpensive to test multiple instance types before making a decision. If your needs change, you can resize your instance later on. For more information, see [Resizing Your Instance \(p. 111\)](#).

The following topics provide additional information about some of the instance types:

- [Micro Instances \(p. 95\)](#)
- [H1 Instances \(p. 102\)](#)
- [HS1 Instances \(p. 104\)](#)
- [GPU Instances \(p. 105\)](#)

## Micro Instances

Micro instances (`t1.micro`) provide a small amount of consistent CPU resources and allow you to increase CPU capacity in short bursts when additional cycles are available. They are well suited for lower throughput applications and websites that require additional compute cycles periodically.

The `t1.micro` instance is available as an Amazon EBS-backed instance only.

This documentation describes how `t1.micro` instances work so that you can understand how to apply them. It's not our intent to specify exact behavior, but to give you visibility into the instance's behavior so you can understand its performance (to a greater degree than you typically get with, for example, a multi-tenant shared web hosting system).

#### Topics


- [Hardware Specifications](#) (p. 96)
- [Optimal Application of Micro Instances](#) (p. 96)
- [Available CPU Resources During Spikes](#) (p. 98)
- [When the Instance Uses Its Allotted Resources](#) (p. 98)
- [Comparison with the `m1.small` Instance Type](#) (p. 100)
- [AMI Optimization for Micro Instances](#) (p. 102)

## Hardware Specifications

For more information about the specifications for each Amazon EC2 instance type, see [Instance Type Details](#).


## Optimal Application of Micro Instances

A `t1.micro` instance provides spiky CPU resources for workloads that have a CPU usage profile similar to what is shown in the following figure.




The instance is designed to operate with its CPU usage at essentially only two levels: the normal low background level, and then at brief spiked levels much higher than the background level. We allow the instance to operate at up to 2 EC2 compute units (ECUs) (one ECU provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor). The ratio between the maximum level and the background level is designed to be large. We designed `t1.micro` instances to support tens of requests per minute on your application. However, actual performance can vary significantly depending on the amount of CPU resources required for each request on your application.


Your application might have a different CPU usage profile than that described in the preceding section. The next figure shows the profile for an application that isn't appropriate for a `t1.micro` instance. The application requires continuous data-crunching CPU resources for each request, resulting in plateaus of CPU usage that the `t1.micro` instance isn't designed to handle.



The next figure shows another profile that isn't appropriate for a `t1.micro` instance. Here the spikes in CPU use are brief, but they occur too frequently to be serviced by a micro instance.



The next figure shows another profile that isn't appropriate for a `t1.micro` instance. Here the spikes aren't too frequent, but the background level between spikes is too high to be serviced by a `t1.micro` instance.



In each of the preceding cases of workloads not appropriate for a `t1.micro` instance, we recommend that you consider using a different instance type. For more information about instance types, see [Instance Types](#) (p. 94).

## Available CPU Resources During Spikes

When your instance *bursts* to accommodate a spike in demand for compute resources, it uses unused resources on the host. The amount available depends on how much contention there is when the spike occurs. The instance is never left with zero CPU resources, whether other instances on the host are spiking or not.


## When the Instance Uses Its Allotted Resources

We expect your application to consume only a certain amount of CPU resources in a period of time. If the application consumes more than your instance's allotted CPU resources, we temporarily limit the instance so it operates at a low CPU level. If your instance continues to use all of its allotted resources, its performance will degrade. We will increase the time that we limit its CPU level, thus increasing the time before the instance is allowed to burst again.


If you enable Amazon CloudWatch monitoring for your `t1.micro` instance, you can use the "Avg CPU Utilization" graph in the AWS Management Console to determine whether your instance is regularly using all its allotted CPU resources. We recommend that you look at the maximum value reached during each given period. If the maximum value is 100%, we recommend that you use Auto Scaling to scale out (with additional `t1.micro` instances and a load balancer), or move to a larger instance type. For more information about Auto Scaling, see the [Auto Scaling Developer Guide](#).

The following figures show the three suboptimal profiles from the preceding section and what it might look like when the instance consumes its allotted resources and we have to limit its CPU level. If the instance consumes its allotted resources, we restrict it to the low background level.


The next figure shows the situation with the long plateaus of data-crunching CPU usage. The CPU hits the maximum allowed level and stays there until the instance's allotted resources are consumed for the period. At that point, we limit the instance to operate at the low background level, and it operates there until we allow it to burst above that level again. The instance again stays there until the allotted resources are consumed and we limit it again (not seen on the graph).



The next figure shows the situation where the requests are too frequent. The instance uses its allotted resources after only a few requests and so we limit it. After we lift the restriction, the instance maxes out its CPU usage trying to keep up with the requests, and we limit it again.




The next figure shows the situation where the background level is too high. Notice that the instance doesn't have to be operating at the maximum CPU level for us to limit it. We limit the instance when it's operating above the normal background level and has consumed its allotted resources for the given period. In this case (as in the preceding one), the instance can't keep up with the work, and we limit it again.




## Comparison with the m1.small Instance Type

The `t1.micro` instance provides different levels of CPU resources at different times (up to 2 ECUs). By comparison, the `m1.small` instance type provides 1 ECU at all times. The following figure illustrates the difference.




The following figures compare the CPU usage of a `t1.micro` instance with an `m1.small` instance for the various scenarios we've discussed in the preceding sections.


The first figure that follows shows an optimal scenario for a `t1.micro` instance (the left graph) and how it might look for an `m1.small` instance (the right graph). In this case, we don't need to limit the `t1.micro` instance. The processing time on the `m1.small` instance would be longer for each spike in CPU demand compared to the `t1.micro` instance.



The next figure shows the scenario with the data-crunching requests that used up the allotted resources on the `t1.micro` instance, and how they might look with the `m1.small` instance.




The next figure shows the frequent requests that used up the allotted resources on the `t1.micro` instance, and how they might look on the `m1.small` instance.



The next figure shows the situation where the background level used up the allotted resources on the `t1.micro` instance, and how it might look on the `m1.small` instance.





## AMI Optimization for Micro Instances

We recommend that you follow these best practices when optimizing an AMI for the `t1.micro` instance type:

- Design the AMI to run on 600 MB of RAM
- Limit the number of recurring processes that use CPU time (for example, cron jobs, daemons)

In Linux, you can optimize performance using swap space and virtual memory (for example, by setting up swap space in a separate partition from the root file system).

In Windows, when you perform significant AMI or instance configuration changes (for example, enable server roles or install large applications), you might see limited instance performance, because these changes can be memory intensive and require long-running CPU resources. We recommend that you first use a larger instance type when performing these changes to the AMI, and then run the AMI on a `t1.micro` instance for normal operations.

## H1 Instances

H1 instances (`hi1.4xlarge`) can deliver tens of thousands of low-latency, random I/O operations per second (IOPS) to applications. They are well suited for the following scenarios:

- NoSQL databases (for example, Cassandra and MongoDB)
- Clustered databases
- Online transaction processing (OLTP) systems

You can cluster H1 instances in a placement group. For more information, see [Placement Groups \(p. 108\)](#).

By default, you can run up to two `hi1.4xlarge` instances. If you need more than two `hi1.4xlarge` instances, you can request more using the [Amazon EC2 Instance Request Form](#).

### Topics

- [Hardware Specifications \(p. 102\)](#)
- [Disk I/O Performance \(p. 103\)](#)
- [SSD Storage \(p. 103\)](#)

## Hardware Specifications

The `hi1.4xlarge` instance type is based on solid-state drive (SSD) technology.

For more information about the specifications for each Amazon EC2 instance type, see [Instance Type Details](#).

## Disk I/O Performance

Using Linux paravirtual (PV) AMIs, H1 instances can deliver more than 120,000 4 KB random read IOPS and between 10,000 and 85,000 4 KB random write IOPS (depending on active logical block addressing span) to applications across two SSD data volumes. For hardware virtual machines (HVM) and Windows AMIs, performance is approximately 90,000 4 KB random read IOPS and between 9,000 and 75,000 4 KB random write IOPS.

The maximum sequential throughput on all AMI types (Linux PV, Linux HVM, and Windows) per second is approximately 2 GB read and 1.1 GB write.

## SSD Storage

This section contains important information you need to know about SSD storage. With SSD storage:

- The primary data source is an instance store with SSD storage.
- Read performance is consistent and write performance can vary.
- Write amplification can occur.
- The TRIM command is not currently supported.

## Instance Store with SSD Storage

The `hi1.4xlarge` instances use an EBS-backed root device. However, their primary data storage is provided by the SSD volumes in the instance store. Like other instance store volumes, these instance store volumes persist only for the life of the instance. Because the root device of the `hi1.4xlarge` instance is EBS-backed, you can still start and stop your instance. When you stop an instance, your application persists, but your production data in the instance store does not persist. For more information about instance store volumes, see [Amazon EC2 Instance Store \(p. 508\)](#).

## Variable Write Performance

Write performance depends on how your applications utilize logical block addressing (LBA) space. If your applications use the total LBA space, write performance can degrade by about 90 percent. Benchmark your applications and monitor the queue depth (the number of pending I/O requests for a volume) and I/O size.

## Write Amplification

Write amplification refers to an undesirable condition associated with flash memory and SSDs, where the actual amount of physical information written is a multiple of the logical amount intended to be written. Because flash memory must be erased before it can be rewritten, the process to perform these operations results in moving (or rewriting) user data and metadata more than once. This multiplying effect increases the number of writes required over the life of the SSD, which shortens the time that it can reliably operate. The `hi1.4xlarge` instances are designed with a provisioning model intended to minimize write amplification.

Random writes have a much more severe impact on write amplification than serial writes. If you are concerned about write amplification, allocate less than the full tebibyte of storage for your application (also known as over provisioning).

## The TRIM Command

The TRIM command enables the operating system to notify an SSD that blocks of previously saved data are considered no longer in use. TRIM limits the impact of write amplification.

TRIM support is not available for `hi1.4xlarge` instances at this time. We hope to add TRIM support in the future.

## HS1 Instances

HS1 instances (`hs1.8xlarge`) provide very high storage density and high sequential read and write performance per instance. They are well suited for the following scenarios:

- Data warehousing
- Hadoop/MapReduce
- Parallel file systems

You can cluster HS1 instances in a placement group. For more information, see [Placement Groups \(p. 108\)](#).

By default, you can run up to two HS1 instances. If you need more than two HS1 instances, you can request more using the [Amazon EC2 Instance Request Form](#).

### Topics

- [Hardware Specifications \(p. 104\)](#)
- [Disk I/O Performance \(p. 104\)](#)
- [Storage Information \(p. 104\)](#)

## Hardware Specifications

HS1 instances support both EBS-backed and instance store-backed Amazon Machine Images (AMIs). HS1 instances support both paravirtual (PV) and hardware virtual machine (HVM) AMIs. HS1 instances are capable of delivering 2.4 GB per second of sequential read and 2.6 GB per second of sequential write performance when using a block size of 2 MiB.

For customers using Microsoft Windows Server, HS1 instances are only supported with the Microsoft Windows Server AMIs for Cluster Instance Type. HS1 instances do not currently support EBS optimization, but provide high bandwidth networking and can also be used with Amazon Elastic Block Store (Amazon EBS) provisioned I/O operations per second (IOPS) volumes for improved consistency and performance.

For more information about the specifications for each Amazon EC2 instance type, see [Instance Type Details](#).

## Disk I/O Performance

HS1 instances are capable of delivering 2.6 GB per second of sequential read and write performance when using a block size of 2 MiB.

## Storage Information

This section contains important information you need to know about the storage used with HS1 instances.

## Instance Store with HS1 Instances

HS1 instances support both instance store and Amazon EBS root device volumes. However, even when using an EBS-backed instance, primary data storage is provided by the hard disk drives in the instance store. Like other instance store volumes, these instance store volumes persist only for the life of the instance. Therefore, when you stop an instance (when using an EBS-backed root volume), your application persists, but your production data in the instance store does not persist. For more information about instance store volumes, see [Amazon EC2 Instance Store \(p. 508\)](#).

## Disk Initialization

If you plan to run an HS1 instance in a steady state for long periods of time, we recommend that you zero the hard disks first for improved performance. This process can take as long as six hours to complete.

## Setting the Memory Limit

Many Linux-based Amazon Machine Images (AMIs) come with `CONFIG_XEN_MAX_DOMAIN_MEMORY` set to 70. We recommend that you set this as appropriate for 117 GiB of memory.

## Setting the User Limit (ulimit)

Many Linux-based Amazon Machine Images (AMIs) come with a default ulimit of 1024. We recommend that you increase the ulimit to 2048.

# GPU Instances

If you require high parallel processing capability, you'll benefit from using GPU instances, which provide access to NVIDIA GPUs with up to 1,536 CUDA cores and 4 GB of video memory. You can use GPU instances to accelerate many scientific, engineering, and rendering applications by leveraging the Compute Unified Device Architecture (CUDA) or OpenCL parallel computing frameworks. You can also use them for graphics applications, including game streaming, 3-D application streaming, and other graphics workloads.

GPU instances run as HVM-based instances. Hardware virtual machine (HVM) virtualization uses hardware-assist technology provided by the AWS platform. With HVM virtualization, the guest VM runs as if it were on a native hardware platform, except that it still uses paravirtual (PV) network and storage drivers for improved performance. This enables Amazon EC2 to provide dedicated access to one or more discrete GPUs in each GPU instance.

You can cluster GPU instances into a cluster placement group. Cluster placement groups provide low latency and high-bandwidth connectivity between the instances within a single Availability Zone. For more information, see [Placement Groups \(p. 108\)](#).

### Topics

- [Hardware Specifications \(p. 105\)](#)
- [GPU Instance Limitations \(p. 106\)](#)
- [AMIs for GPU Instances \(p. 106\)](#)
- [Installing the NVIDIA Driver on Linux \(p. 106\)](#)
- [Installing the NVIDIA Driver on Windows \(p. 107\)](#)

## Hardware Specifications

For more information about the specifications for each Amazon EC2 instance type, see [Instance Type Details](#).

## GPU Instance Limitations

GPU instances currently have the following limitations:

- They aren't available in every region.
- They must be launched from HVM AMIs.
- They can't access the GPU unless the NVIDIA drivers are installed.
- They aren't available for use with Amazon DevPay.
- We limit the number of instances that you can run. For more information, see [How many instances can I run in Amazon EC2?](#) in the Amazon EC2 FAQ. To request an increase in these limits, use the following form: [Request to Increase Amazon EC2 Instance Limit](#).

## AMIs for GPU Instances

You can launch a G2 instance using Windows Server 2012 and Windows Server 2008 R2 AMIs. In addition, you can launch Linux HVM AMIs with the following operating systems: Amazon Linux, SUSE Enterprise Linux, and Ubuntu. If you encounter the following error and want to launch AMIs with operating systems not listed here, contact [Customer Service](#) or reach out through [EC2 forums](#).

```
Client.UnsupportedOperation: Instances of type 'g2.2xlarge' may not be launched from AMI <ami-id>.
```

You can launch a CG1 instance using any HVM AMI.

To help you get started, NVIDIA provides AMIs for GPU instances for Amazon Linux and Windows. These reference AMIs include the NVIDIA driver, which enables full functionality and performance of the NVIDIA GPUs. For a list of AMIs with the NVIDIA driver, see [AWS Marketplace \(NVIDIA GRID\)](#).

## Installing the NVIDIA Driver on Linux

A GPU instance must have the appropriate NVIDIA driver. The NVIDIA driver you install must be compiled against the kernel that you intend to run on your instance.

Amazon provides updated and compatible builds of the NVIDIA kernel drivers for each official kernel upgrade. If you decide to use a different NVIDIA driver version than the one Amazon provides, or decide to use a kernel that's not an official Amazon build, you must uninstall the Amazon-provided NVIDIA packages from your system to avoid conflicts with the versions of the drivers you are trying to install.

Use this command to uninstall Amazon-provided NVIDIA packages:

```
$ sudo yum erase nvidia cudatoolkit
```

The Amazon-provided CUDA toolkit package has dependencies on the NVIDIA drivers. Uninstalling the NVIDIA packages erases the CUDA toolkit. You must reinstall the CUDA toolkit after installing the NVIDIA driver.

You can download NVIDIA drivers from <http://www.nvidia.com/Download/Find.aspx>. Select a driver for the NVIDIA GRID K520 (G2 instances) or Tesla M-Class M2050 (CG1 instances) for Linux 64-bit systems. For more information about installing and configuring the driver, open the **ADDITIONAL INFORMATION** tab on the download page for the driver on the NVIDIA website and click the README link.

## Manually Install the NVIDIA Driver

### To install the driver for an Amazon Linux AMI

1. Make sure the `kernel-devel` package is installed and matches the version of the kernel you are currently running.

```
$ yum install kernel-devel-`uname -r`
```

2. Run the self-install script to install the NVIDIA driver. For example:

```
$ /root/NVIDIA-Linux-x86_64_319.60.run
```

3. Reboot the instance. For more information, see [Reboot Your Instance \(p. 286\)](#).
4. Confirm that the driver is functional. The response for the following command lists the installed NVIDIA driver version and details about the GPUs.

```
$ /usr/bin/nvidia-smi -q -a
```

## Installing the NVIDIA Driver on Windows

To install the NVIDIA driver on your Windows instance, log on to your instance as Administrator using Remote Desktop. You can download NVIDIA drivers from <http://www.nvidia.com/Download/Find.aspx>. Select a driver for the NVIDIA GRID K520 (G2 instances) or Tesla M-Class M2050 (CG1 instances) for your version of Windows Server. Open the folder where you downloaded the driver and double-click the installation file to launch it. Follow the instructions to install the driver and reboot your instance as required. To verify that the GPU is working properly, check device manager.

When using Remote Desktop, GPUs that use the WDDM driver model are replaced with a non-accelerated Remote Desktop display driver. In order to access your GPU hardware, you must use a different remote access tool, such as VNC. You can also use one of the GPU AMIs from the AWS Marketplace because they provide remote access tools that support 3-D acceleration.

## EBS-Optimized Instances

An EBS-optimized instance uses an optimized configuration stack and provides additional, dedicated capacity for EBS I/O. This optimization provides the best performance for your EBS volumes by minimizing contention between EBS I/O and other traffic from your Amazon EC2 instance.

When you use an EBS-optimized instance, you pay an additional low, hourly fee for the dedicated capacity.

EBS optimization enables instances to fully utilize the IOPS provisioned on an EBS volume. EBS-optimized instances deliver dedicated throughput to EBS, with options between 500 Mbps and 1,000 Mbps, depending on the instance type you use. When attached to an EBS-optimized instance, Provisioned IOPS volumes are designed to deliver within 10 percent of their provisioned performance 99.9 percent of the time in a given year. For more information, see [Provisioned IOPS Volumes \(p. 449\)](#).

The table below describes the instance types that can be launched as EBS-optimized instances, the dedicated EBS throughput they provide, and the maximum number of input/output operations per second (IOPS) that you can drive over the EBS connection.

Instance Type	EBS-optimized Available	Maximum Throughput (MB/s)*	Max 16K IOPS
t1.micro	No	32 MB/s	2,000
m1.small	No	64 MB/s	4,000
m1.medium	No	64 MB/s	4,000
m1.large	Yes	64 MB/s	4,000
m1.xlarge	Yes	128 MB/s	8,000
m3.xlarge	Yes	64 MB/s	4,000
m3.2xlarge	Yes	128 MB/s	8,000
c1.medium	No	32 MB/s	2,000
c1.xlarge	Yes	128 MB/s	8,000
cc2.8xlarge	N/A**	800 MB/s	48,000
m2.xlarge	No	64 MB/s	4,000
m2.2xlarge	Yes	64 MB/s	4,000
m2.4xlarge	Yes	128 MB/s	8,000
cr1.8xlarge	N/A**	800 MB/s	48,000
hi1.4xlarge	N/A**	800 MB/s	48,000
hs1.8xlarge	N/A**	800 MB/s	48,000
g2.2xlarge	Yes	128 MB/s	8,000
cg1.4xlarge	N/A**	800 MB/s	48,000

To launch an EBS-optimized instance using the AWS Management Console, select the **Launch as EBS-optimized instance** option in the launch wizard.

If the instance type that you've selected can't be launched as EBS-optimized instance, or if the instance type is EBS-optimized by default, this option is not available.

To launch an EBS-optimized instance using the `ec2-run-instances` command, specify the `--ebs-optimized` option.

## Placement Groups

A *placement group* is a logical grouping of instances within a single Availability Zone. Using placement groups enables applications to get the full-bisection bandwidth and low-latency network performance required for tightly coupled, node-to-node communication typical of HPC applications. For more information, see [High Performance Computing \(HPC\) on AWS](#).

First, you create a placement group and then you launch multiple instances into the placement group. We recommend that you launch the number of instances that you need in the placement group in a single launch request. If you try to add more instances to the placement group later, you increase your chances of getting an insufficient capacity error.

If you stop an instance in a placement group and then start it again, it still runs in the placement group. However, the start fails if there isn't enough capacity for the instance.

If you receive a capacity error when launching an instance in a placement group, stop and restart the instances in the placement group, and then try the launch again.

### Topics

- [Placement Group Limitations \(p. 109\)](#)
- [Launching Instances into a Placement Group \(p. 109\)](#)
- [Deleting a Placement Group \(p. 110\)](#)

## Placement Group Limitations

Placement groups have the following limitations:

- A placement group can't span multiple Availability Zones.
- The name you specify for a placement group a name must be unique within your AWS account.
- Instances that you launch into a placement group must be based on an HVM AMI with an Amazon EBS root device volume.
- The following are the only instance types that you can use when you launch an instance into a placement group:
  - `cc2.8xlarge`
  - `cg1.4xlarge`
  - `cr1.8xlarge`
  - `hi1.4xlarge`
  - `hs1.8xlarge`
- A placement group can't contain instances of different instance types. After you launch one or more instances of a particular type into a placement group, you can only add instances of that type to the placement group. You can't change the instance type supported by a placement group after you've launched an instance into it.
- You can't merge placement groups. Instead, you must terminate the instances in one placement group, and then relaunch those instances into the other placement group.
- Amazon EC2 Reserved Instances are not guaranteed within a specific requested placement group.

## Launching Instances into a Placement Group

We suggest that you create an AMI specifically for the instances that you'll launch into a placement group.

### To launch an instance into a placement group

1. Open the Amazon EC2 console.
2. Create an AMI for your instances.
  - a. From the Amazon EC2 dashboard, click **Launch Instance**. Be sure to select an HVM AMI for cluster instances. After you complete the wizard, click **Launch**.
  - b. Connect to your instance. (For more information, see [Connect to Your Amazon EC2 Instance \(p. 273\)](#).)
  - c. Install software and applications on the instance, copy data, or attach additional Amazon EBS volumes.
  - d. In the navigation pane, click **Instances**, select your instance, click **Actions**, and then click **Create Image**. Provide the information requested by the **Create Image** dialog box, and then click **Create Image**.



- e. (Optional) You can terminate this instance if you have no further use for it.
3. Create a placement group.
    - a. In the navigation pane, click **Placement Groups**.
    - b. Click **Create Placement Group**.
    - c. In the **Create Placement group** dialog box, provide a name for the placement group that is unique in the AWS account you're using, and then click **Yes, Create**.

When the status of the placement group is `available`, you can launch instances into the placement group.
  4. Launch your instances into the placement group.
    - a. In the navigation pane, click **Instances**.
    - b. Click **Launch Instance**. Complete the wizard as directed, taking care to select the following:
      - The HVM AMI that you created
      - The number of instances that you'll need
      - The placement group that you created

If you prefer, you can use the [ec2-create-image](#) command to create your AMI, the [ec2-create-placement-group](#) command to create your placement group, and use the [ec2-run-instances](#) command to launch an instance into the placement group.

## Deleting a Placement Group

You can delete a placement group if you need to replace it or no longer need a placement group. Before you can delete your placement group, you must terminate all instances that you launched into the placement group.

### To delete a placement group

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Instances**.
3. Select and terminate all instances in the placement group. (You can verify that the instance is in a placement group before you terminate it by checking the value of **Placement Group** in the details pane.)
4. In the navigation pane, click **Placement Groups**.
5. Select the placement group, and then click **Delete**.
6. When prompted for confirmation, click **Yes, Delete**.

If you prefer, you can use the [ec2-terminate-instances](#) command to terminate the instances and the [ec2-delete-placement-group](#) command to delete the placement group.

## Resizing Your Instance

As your needs change, you might find that your instance is over-utilized (the instance type is too small) or under-utilized (the instance type is too large). If this is the case, you can change the size of your instance. For example, if your `t1.micro` instance is too small for its workload, you can change it to an `m1.small` instance.

The process for resizing an instance varies depends on the type of its root device volume, as follows:

- If the root device for your instance is an Amazon EBS volume, you can easily resize your instance by changing its instance type.
- If the root device for your instance is an instance store volume, you must migrate to a new instance.

To determine the root device type of your instance, open the Amazon EC2 console, click **Instances**, select the instance, and check the value of **Root device type** in the details pane. The value is either `ebs` or `instance store`.

For more information about root device volumes, see [Storage for the Root Device \(p. 45\)](#).

### Topics

- [Resizing an EBS-backed Instance \(p. 111\)](#)
- [Resizing an Instance Store-backed Instance \(p. 112\)](#)

## Resizing an EBS-backed Instance

You must stop your EBS-backed instance before you can change its instance type. When you stop and start an instance, we move it to new hardware. If the instance is running in EC2-Classic, we give it new public and private IP addresses, and disassociate any Elastic IP address that's associated with the instance. Therefore, to ensure that your users can continue to use the applications that you're hosting on your instance uninterrupted, you must re-associate any Elastic IP address after you restart your instance. For more information, see [Stop and Start Your Instance \(p. 284\)](#).

Use the following procedure to resize an Amazon EBS-backed instance using the AWS Management Console.

### To resize an EBS-backed instance

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Instances**, and select the instance.
3. [EC2-Classic] If the instance has an associated Elastic IP address, write down the Elastic IP address and the instance ID shown in the details pane.
4. Click **Actions**, and then click **Stop**.
5. In the confirmation dialog box, click **Yes, Stop**. It can take a few minutes for the instance to stop.  
  
[EC2-Classic] When the instance state becomes `stopped`, the **Elastic IP**, **Public DNS**, **Private DNS**, and **Private IPs** fields in the details pane are blank to indicate that the old values are no longer associated with the instance.
6. With the instance still selected, click **Actions**, and then click **Change Instance Type**. Note that this action is disabled if the instance state is not `stopped`.
7. In the **Change Instance Type** dialog box, in the **Instance Type** list, select the type of instance that you need, and then click **Apply**.
8. To restart the stopped instance, select the instance, click **Actions**, and then click **Start**.

9. In the confirmation dialog box, click **Yes, Start**. It can take a few minutes for the instance to enter the `running` state.

[EC2-Classic] When the instance state is `running`, the **Public DNS**, **Private DNS**, and **Private IPs** fields in the details pane contain the new values that we assigned to the instance.

10. [EC2-Classic] If your instance had an associated Elastic IP address, you must reassociate it as follows:
  - a. In the navigation pane, click **Elastic IPs**.
  - b. Select the Elastic IP address that you wrote down before you stopped the instance.
  - c. Click **Associate Address**.
  - d. Select the instance ID that you wrote down before you stopped the instance, and then click **Associate**.

## Resizing an Instance Store-backed Instance

You can create an image from your current instance, launch a new instance from this image with the instance type you need, and then terminate the instance you no longer need. To ensure that your users can continue to use the applications that you're hosting on your instance uninterrupted, you must take any Elastic IP address that you've associated with your current instance and associate it with the new instance.

### Prerequisites

Before you begin, ensure that you have the following:

- Your AWS account ID. To retrieve your account ID, go to [Your Security Credentials](#) and expand **Account Identifiers**.
- Your access key ID. If you need to retrieve or create an access key ID, go to [Your Security Credentials](#), and expand **Access Keys**.
- Your secret access key. You can't retrieve your secret access key. Therefore, if you can't find your secret access key, you'll need to create a new one. To create a secret access key, go to [Your Security Credentials](#), expand **Access Keys**, and click **Create New Access Key**.
- An X.509 certificate. If you need to create an X.509 certificate, go to [Your Security Credentials](#), expand **X.509 Certificates**, and click **Create New Certificate**.
- A private key (.pem) file. If you need to create a key pair, follow the directions in [Creating Your Key Pair Using Amazon EC2 \(p. 386\)](#).
- An Amazon S3 bucket. To create an Amazon S3 bucket, open the Amazon S3 console and click **Create Bucket**.
- A connection to your instance store-backed instance. For more information, see [Connect to Your Amazon EC2 Instance \(p. 273\)](#).

### To resize an instance store-backed instance

1. Write down the ID of the instance store-backed instance that you're starting with. If the instance has an associated Elastic IP address, write down the Elastic IP address as well.
2. From the instance, create a temporary directory for your credentials using the following **mkdir** command:

```
mkdir /tmp/cert
```

3. From your computer, use [scp \(p. 280\)](#) (Linux/UNIX), [pscp \(p. 277\)](#) (Windows), or [winscp \(p. 278\)](#) (Windows) to copy your private key file (for example, `my-key-pair.pem`) and your certificate file (`cert-mycert.pem`) from your computer to the `/tmp/cert` directory of your instance.
4. From the instance, obtain root privileges using the following **sudo** command:

```
sudo su
```

5. From the instance, verify that the Amazon EC2 AMI tools are installed. For example, run the following command: **ec2-bundle-vol --help**. If this command is not available, download the [Amazon EC2 AMI Tools](#).

Create an image using the following [ec2-bundle-vol](#) command:

```
ec2-bundle-vol -k /tmp/cert/my-key-pair.pem -c /tmp/cert/cert-mycert.pem -  
u your_aws_account_id -r x86_64 -e /tmp/cert
```

It can take a few minutes to create the image. When this command completes, your `tmp` directory contains the image files (`image.manifest.xml`, plus multiple `image.part.xx` files).

6. From the instance, upload the image to your Amazon S3 bucket (for example, `my-s3-bucket`) using the following [ec2-upload-bundle](#) command. Note that if the folders don't exist in the bucket, this command creates them.

```
ec2-upload-bundle -b my-s3-bucket/images/image_name -m /tmp/image.manifest.xml  
-a your_access_key_id -s your_secret_access_key
```

7. (Optional) After the image is uploaded to Amazon S3, you can remove the image files from the `tmp` directory on the instance using the following **rm** command:

```
rm -rf /tmp/im*
```

8. From either the instance or your computer, verify that the CLI tools are installed using the following command: **ec2-version**. If this command is not available, follow the directions in [Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX \(p. 541\)](#).

Register the image with Amazon EC2 using the following [ec2-register](#) command. Note that you don't need to specify the `-O` and `-W` options if you've set the `AWS_ACCESS_KEY` and `AWS_SECRET_KEY` environment variables.

```
ec2-register my-s3-bucket/images/image_name/image.manifest.xml -n image_name  
-O your_access_key_id -W your_secret_access_key
```

Write down the ID of the image (in the form `ami-xxxxxxx`) that's returned by the command.

9. Open the Amazon EC2 console and in the navigation pane, select **AMIs**. From the filter lists, select **Owned by me**, and select the image whose ID you wrote down when you registered it. Notice that **AMI Name** is the name that you specified when you registered the image and **Source** is your Amazon S3 bucket.
10. Click **Launch**. When you specify options in the launch wizard, be sure to specify the new instance type that you need. It can take a few minutes for the instance to enter the `running` state.
11. If the instance that you started with had an associated Elastic IP address, you must associate it with the new instance as follows:
  - a. In the navigation pane, click **Elastic IPs**.

- b. Select the Elastic IP address that you wrote down at the beginning of this procedure.
  - c. Click **Associate Address**.
  - d. Select the ID of the new instance, and then click **Associate**.
12. (Optional) You can terminate the instance that you started with, if it's no longer needed. Select the instance and check its instance ID against the instance ID that you wrote down at the beginning of this procedure to verify that you are terminating the correct instance. Click **Actions**, and then click **Terminate**.

## Spot Instances

If you have flexibility on when your application will run, you can bid on unused Amazon EC2 compute capacity, called Spot Instances, and lower your costs significantly. Set by Amazon EC2, the Spot price for these instances fluctuates periodically depending on the supply of and demand for Spot Instance capacity.

To use Spot Instances, you place a Spot Instance request (your bid) specifying the maximum price you are willing to pay per hour per instance. If the maximum price of your bid is greater than the current Spot price, your request is fulfilled and your instances run until you terminate them or the Spot price increases above your maximum price. Your instance can also be terminated when your bid price equals the market price, even when there is no increase in the market price. This can happen when demand for capacity rises, or when supply fluctuates.

### Note

You can run Amazon EC2 Spot Instances on the following platforms: Amazon Linux/UNIX, SUSE Linux Enterprise Server, and Windows (without SQL Server).

You will often pay less per hour than your maximum bid price. The Spot price is adjusted periodically as requests come in and the available supply of instances changes. Everyone pays that same Spot price for that period regardless of whether their maximum bid price was higher, and you will never pay more than your hourly maximum bid price.

## Quick Look: Getting Started with Spot Instances Video

The following video shows how to launch your first Spot Instance using the AWS Management Console. This video includes instructions on placing a bid, determining when the instance is fulfilled, and canceling the instance. [Getting Started with Spot Instances](#)

## Checklist for Getting Started with Spot Instances

If you want to get started working with Spot Instances, here are the resources you need to get going.

- [Getting Started with Spot Instances](#) (p. 116)
  - [Viewing Spot Instance Pricing History](#) (p. 118)
  - [Creating a Spot Instance Request](#) (p. 122)
  - [Finding Running Spot Instances](#) (p. 126)
  - [Canceling Spot Instance Requests](#) (p. 129)
- [Fundamentals of Spot Instances](#) (p. 132)
  - [Placing Spot Requests](#) (p. 132)
  - [Tagging Spot Instance Requests](#) (p. 141)
  - [Protecting Your Spot Instance Data Against Interruptions](#) (p. 142)
- [Walkthroughs: Using Spot Instances with AWS Services](#) (p. 143)
  - [Managing Spot Instances with Auto Scaling](#) (p. 143)
  - [Using CloudFormation Templates to Launch Spot Instances](#) (p. 158)
  - [Launching Amazon Elastic MapReduce Job Flows with Spot Instances](#) (p. 159)
  - [Launching Spot Instances in Amazon Virtual Private Cloud](#) (p. 160)
- [Advanced Tasks](#) (p. 161)
  - [Subscribe to Your Spot Instance Data Feed](#) (p. 161)
  - [Programming Spot with AWS Java SDK](#) (p. 164)

- [Starting Clusters on Spot Instances \(p. 191\)](#)

## What's New in Spot Instances

Here's a quick look at what's new in Spot Instances:

- [Customizing Your Spot Requests \(p. 134\)](#)
- [Tracking Spot Requests with Bid Status Codes \(p. 135\)](#)

## Getting Started with Spot Instances

In this section, we will step you through how to get started using Spot Instances. First, we'll take you through what you need to know before you begin and the prerequisites you need to get started.

We will walk you through the following activities:

- [Viewing Spot Instance Pricing History \(p. 118\)](#)
- [Creating a Spot Instance Request \(p. 122\)](#)
- [Finding Running Spot Instances \(p. 126\)](#)
- [Canceling Spot Instance Requests \(p. 129\)](#)

## Before You Begin

If you are new to Spot Instances, take a look at [Prerequisites for Using Spot Instances \(p. 116\)](#) to make sure you can take full advantage of the benefits of this Amazon EC2 product. If you have been using Amazon EC2 and you're ready to proceed, click one of the steps in the preceding list to get going.

Before requesting a Spot Instance, consider configuring your Amazon Machine Image (AMI) so that your application does the following:

- Automatically performs the tasks you want at start-up because the instance will start asynchronously without notification. For example, if you're using batch processes, you could set up your AMI to pull jobs from an Amazon Simple Queue Service (Amazon SQS) queue.
- Stores important data regularly in a place that won't be affected by instance termination. For example, you could store your data using Amazon Simple Storage Service (Amazon S3), Amazon SimpleDB, or Amazon Elastic Block Store (Amazon EBS).

### Important

Although Spot Instances can use Amazon EBS-backed AMIs, be aware that Spot Instances do not support the Stop/Start feature. In other words, you can't stop and start Spot Instances launched from an AMI that has an Amazon EBS root device. For information about Amazon EBS, see [Amazon Elastic Block Store \(Amazon EBS\)](#).

- Handles termination gracefully.

For information about creating AMIs, see [Creating Your Own AMIs \(p. 62\)](#).

## Prerequisites for Using Spot Instances


To work with Amazon EC2 Spot Instances, we assume you have read and completed the instructions described in [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#), which provides information on creating your Amazon EC2 account and credentials.

In addition, whichever you choose to use—the AWS Management Console, the Amazon EC2 command line interface (CLI), or the Amazon EC2 application programming interface (API)—Amazon EC2 provides tools for Spot Instances that you can use to assess Spot price history, submit Spot Instance requests (also called *bids*), and manage your Spot requests and instances. You can also use, develop, and manage your applications using the AWS SDKs. For more information, see [Tools for Amazon Web Services](#).

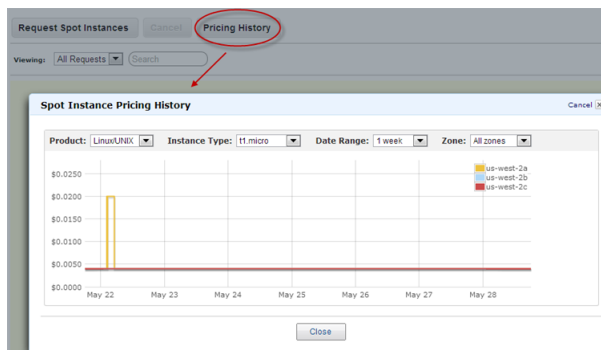
## AWS Management Console

In the AWS Management Console, the EC2 console has tools specifically designed for Spot Instance request tasks. The EC2 console also has general tools that you can use to manage the instances launched when your requests are fulfilled. The following list identifies the tools you can use in the EC2 console:

- The **Spot Requests** page is the main way you interact with your Spot Instance requests.



- **Spot Instance Pricing History** gives you an insight in the pricing patterns for specific Spot Instance types in Availability Zones over a defined period.



- Use the **Request Spot Instances** page to submit a Spot Instance request and specify the details of the instance to be launched when your request succeeds.
- Use the **Instances** page to manage the instances launched when your Spot request succeeds.

## Amazon EC2 Command Line Interface (CLI) Tools

You can use Amazon EC2 CLI tools specifically designed for managing Spot requests. To manage the instances launched when your Spot request is fulfilled, use the same commands that you use for On-Demand EC2 instances.

### Note

If you use the Amazon EC2 command line interface (CLI) tools, we assume that you have read and completed the instructions described in the [Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX \(p. 541\)](#). It walks you through setting up your environment for use with the CLI tools.

The following table lists the commands you use for Spot request tasks.



Task	CLI
<a href="#">Viewing Spot Instance Pricing History (p. 118).</a>	<code>ec2-describe-spot-price-history</code>
<a href="#">Finding Running Spot Instances (p. 126).</a>	<code>ec2-describe-spot-instance-requests</code>
<a href="#">Creating a Spot Instance Request (p. 122).</a>	<code>ec2-request-spot-instances</code>
<a href="#">Subscribe to Your Spot Instance Data Feed (p. 163).</a>	<code>ec2-create-spot-datafeed-subscription</code>
<a href="#">Data Feed Filename and Format (p. 162).</a>	<code>ec2-describe-spot-datafeed-subscription</code>
<a href="#">Delete a Spot Instance Data Feed (p. 164).</a>	<code>ec2-delete-spot-datafeed-subscription</code>
<a href="#">Canceling Spot Instance Requests (p. 129).</a>	<code>ec2-cancel-spot-instance-requests</code>

For information about CLI commands, go to the [Amazon Elastic Compute Cloud Command Line Reference](#).

## API

You can use Amazon EC2 API tools specifically designed to manage your Spot requests. To manage the instances launched when your Spot request is fulfilled, use the same API actions that you use for On-Demand EC2 instances.

The following table lists the API actions you use for Spot request tasks.

Task	API
<a href="#">Viewing Spot Instance Pricing History (p. 118).</a>	<code>DescribeSpotPriceHistory</code>
<a href="#">Finding Running Spot Instances (p. 126).</a>	<code>DescribeSpotInstanceRequests</code>
<a href="#">Creating a Spot Instance Request (p. 122).</a>	<code>RequestSpotInstances</code>
<a href="#">Subscribe to Your Spot Instance Data Feed (p. 163).</a>	<code>CreateSpotDatafeedSubscription</code>
<a href="#">Data Feed Filename and Format (p. 162).</a>	<code>DescribeSpotDatafeedSubscription</code>
<a href="#">Delete a Spot Instance Data Feed (p. 164).</a>	<code>DeleteSpotDatafeedSubscription</code>
<a href="#">Canceling Spot Instance Requests (p. 129).</a>	<code>CancelSpotInstanceRequests</code>

For information about API actions, go to the [Amazon Elastic Compute Cloud API Reference](#).

## AWS Java SDK

Java developers can go to the *AWS SDK for Java* to consult the Java tutorials on Spot Instances:

- [Tutorial: Amazon EC2 Spot Instances \(p. 166\)](#)
- [Tutorial: Advanced Amazon EC2 Spot Request Management \(p. 175\)](#)

## Viewing Spot Instance Pricing History

The Spot price represents the price above which you have to bid to guarantee that a single Spot request is fulfilled. When your bid is above the Spot price, your Spot Instance is launched, and if the Spot price rises above your bid price, your Spot Instance is terminated. You might choose to bid above the current Spot price so that your Spot request is fulfilled quickly. However, before specifying a price with which you

want to bid for your Spot Instance, we recommend that you view the Spot price history. You can view the Spot price history for the last 90 days for any *pool* of Spot Instances sharing the same instance type, operating system, and Availability Zone.

For example, let's say you want to bid on a Linux/UNIX t1.micro instance to be launched in the us-east-1 region. To view past prices in this Spot pool, specify these values using the **Spot Instance Pricing History** page of the AWS Management Console, the `DescribeSpotPriceHistory` API action, or `ec2-describe-spot-price-history` CLI command. If you need to launch your Spot Instance in a specific Availability Zone, you can specify that Availability Zone when retrieving the Spot price history.

After you review the Spot price history, you might choose to bid at a price that would have given you 75 percent Spot Instance uptime in the past. Or, you might choose to bid two times the current Spot price because doing so would have given you 99 percent uptime in the past. However you frame your bid, keep in mind that past performance of Spot prices is not a guarantee of future results. Spot prices vary based on real-time supply and demand conditions, and the conditions that generated certain Spot prices or pricing patterns in the past may not repeat themselves in the future.

#### Note

Make sure you have set up the prerequisites for working with Amazon EC2. If you haven't, see [Prerequisites for Using Spot Instances \(p. 116\)](#).

If you are using an API version earlier than 2011-05-15, the `DescribeSpotPriceHistory` action or the `ec2-describe-spot-price-history` command will return the lowest price across the region for the given time period and the prices will be returned in chronological order.

## AWS Management Console

### To view Spot Price history

1. From the [Amazon EC2 console](#), click **Spot Requests** in the navigation pane.

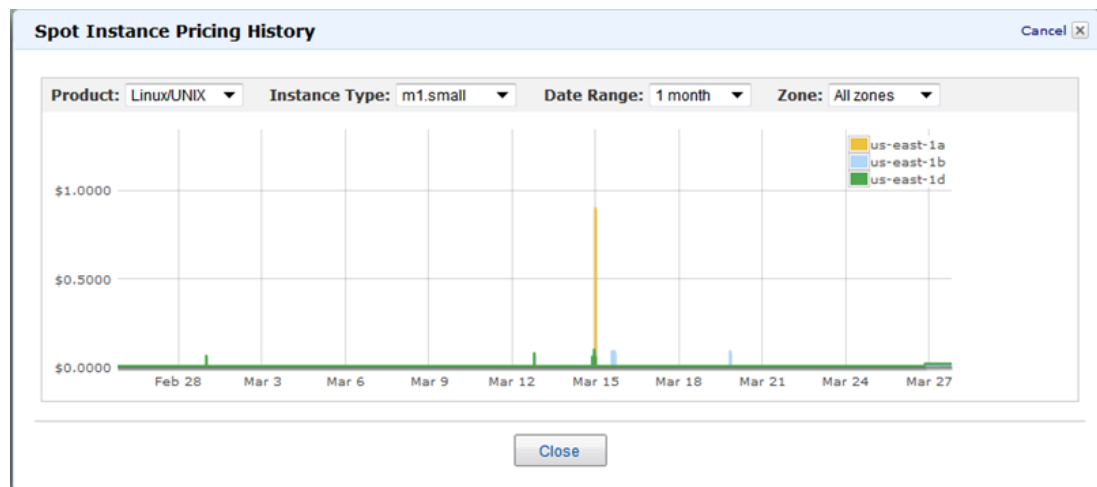
The **Spot Requests** pane opens on the right. It will list your Spot requests if you have any.

2. At the top of the pane, click the **Pricing History** button.

The console displays the **Spot Instance Pricing History** page.

3. If you want to view the Spot price history for specific Availability Zones, click the **Zone** drop-down list and select an Availability Zone.

The **Spot Instance Pricing History** page displays the Spot Instance pricing history for all zones or the zone you selected.



- Using the price history as a guide, select a price that you think would likely keep your instances running for the period of time you need.

## Amazon EC2 Command Line Interface (CLI) Tools

### To view Spot price history

- Enter the following command:

```
PROMPT> ec2-describe-spot-price-history -H --instance-type m1.xlarge
```

Amazon EC2 returns output similar to the following:

```
SPOTINSTANCEPRICE 0.384000 2011-05-25T11:37:48-0800 m1.xlarge
Windows us-east-1b
SPOTINSTANCEPRICE 0.384000 2011-05-25T11:37:48-0800 m1.xlarge
Windows us-east-1d
...
SPOTINSTANCEPRICE 0.242000 2011-04-18T14:39:14-0800 m1.xlarge SUSE
Linux us-east-1d
SPOTINSTANCEPRICE 0.242000 2011-04-18T14:39:14-0800 m1.xlarge SUSE
Linux us-east-1a
```

In this example, the price for the `m1.xlarge` instance type ranges between \$0.242 and \$0.384.

- Using the price history as a guide, select a price that you think would likely keep your instances running for the period of time that you need.

### Tip

You can filter the Spot history data so it includes only instance types or dates of interest to you. For more information about how to filter the results, go to [ec2-describe-spot-price-history](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

## API

### To view Spot price history

- Construct the following Query request.

```
https://ec2.amazonaws.com/
?Action=DescribeSpotPriceHistory
&InstanceType=instance_type
&...auth parameters...
```

Following is an example response.

```
<DescribeSpotPriceHistoryResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <spotPriceHistorySet>
    <item>
      <instanceType>m1.small</instanceType>
      <productDescription>Linux/UNIX</productDescription>
      <spotPrice>.28</spotPrice>
      <timestamp>2009-12-01T11:51:50.000Z</timestamp>
```

```
    <availabilityZone>&region_iad;&az1;</availabilityZone>
  </item>
  <item>
    <instanceType>m1.small</instanceType>
    <productDescription>Linux/UNIX</productDescription>
    <spotPrice>.28</spotPrice>
    <timestamp>2009-12-01T11:51:50.000Z</timestamp>
    <availabilityZone>&region_iad;&az1;</availabilityZone>
  </item>
  <item>
    <instanceType>m1.small</instanceType>
    <productDescription>Linux/UNIX</productDescription>
    <spotPrice>.31</spotPrice>
    <timestamp>2009-12-01T11:51:50.000Z</timestamp>
    <availabilityZone>&region_iad;&az2;</availabilityZone>
  </item>
  <item>
    <instanceType>m1.small</instanceType>
    <productDescription>Linux/UNIX</productDescription>
    <spotPrice>.30</spotPrice>
    <timestamp>2009-12-01T11:51:50.000Z</timestamp>
    <availabilityZone>&region_iad;&az2;</availabilityZone>
  </item>
  <item>
    <instanceType>m1.small</instanceType>
    <productDescription>Linux/UNIX</productDescription>
    <spotPrice>.25</spotPrice>
    <timestamp>2009-12-01T11:51:50.000Z</timestamp>
    <availabilityZone>&region_iad;&az3;</availabilityZone>
  </item>
  <item>
    <instanceType>m1.small</instanceType>
    <productDescription>Linux/UNIX</productDescription>
    <spotPrice>.28</spotPrice>
    <timestamp>2009-12-01T11:51:50.000Z</timestamp>
    <availabilityZone>&region_iad;&az3;</availabilityZone>
  </item>
  <item>
    <instanceType>m1.small</instanceType>
    <productDescription>Linux/UNIX</productDescription>
    <spotPrice>.35</spotPrice>
    <timestamp>2009-12-01T11:51:50.000Z</timestamp>
    <availabilityZone>&region_iad;&az3;</availabilityZone>
  </item>
</spotPriceHistorySet>
<nextToken/>
</DescribeSpotPriceHistoryResponse>
```

- Using the price history as a guide, select a price that you think would likely keep your instances running for the period of time you need.

**Tip**

You can filter the Spot history data so it includes only instance types or dates of interest to you. For more information about how to filter the results, go to [DescribeSpotPriceHistory](#) in the *Amazon Elastic Compute Cloud API Reference*.

**What do you want to do next?**

- [Creating a Spot Instance Request \(p. 122\)](#)
- [Finding Running Spot Instances \(p. 126\)](#)
- [Launching Spot Instances in Amazon Virtual Private Cloud \(p. 160\)](#)

## Creating a Spot Instance Request

After deciding on your bid price, you are ready to request a Spot Instance. Using the AWS Management Console, the API, or the CLI, you can specify how many Spot Instances you'd like, what instance type to use, and the bid price you are willing to pay.

If you request multiple Spot Instances at once, Amazon EC2 creates separate Spot Instance Requests with unique IDs (e.g., sir-1a2b3c4d) so that you can track the status of each request separately. For information about tracking Spot requests, see [Tracking Spot Requests with Bid Status Codes \(p. 135\)](#).

In this section, we discuss how to create requests for Amazon EC2 Spot Instances using the AWS Management Console, the API, or the CLI. For more information about requesting Spot Instances, see the following:

- [Spot Instance Limits \(p. 133\)](#)
- [Types of Spot Instance Requests \(p. 134\)](#)
- [Launching Spot Instances in Launch Groups and Availability Zones \(p. 135\)](#)
- [Spot Instance Request States \(p. 134\)](#)
- [Launching Spot Instances in Amazon Virtual Private Cloud \(p. 160\)](#)

### Note

Make sure you have set up the prerequisites for working with Amazon EC2. If you haven't, see [Prerequisites for Using Spot Instances \(p. 116\)](#).

## AWS Management Console

### To create a Spot Instance request

1. From the [Amazon EC2 console](#), click **Spot Requests** in the navigation pane.
2. Click the **Request Spot Instances** button.

Except for the Spot Instance options that you specify to configure your Spot Instances, the process for requesting Spot Instances is the same as the process for launching On-Demand instances. You go through the same steps when you click the **Launch Instance** button from the **Instances** page: Choose an Amazon Machine Image (AMI) and an instance type, and configure the details of the instance. For information, see [Launch Your Instance \(p. 266\)](#).

3. On the **Configure Instance Details** screen, **Request Spot Instances** is checked as the **Purchasing Option** by default if you get to this configure screen from the **Spot Requests** page. It is only when the **Request Spot Instances** option is selected that the options related to Spot Instance requests are displayed.

Specify the **Maximum price** you are willing to pay for your Spot Instance. Notice that the current Spot Price for each of the Availability Zones in your region is listed. Use this information as a guide on the price to specify so that your request is fulfilled. Your instance launches and runs if your bid price exceeds the Spot Price.

To explore Spot Price trends, click **Pricing History** in the **Spot Requests** page. For information, see [Viewing Spot Instance Pricing History \(p. 118\)](#).

**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

**Number of instances** ⓘ

---

**Purchasing option** ⓘ  Request Spot Instances

**Current price** ⓘ

us-east-1a	0.013
us-east-1b	0.013
us-east-1d	0.013
us-east-1e	0.013

**Maximum price** ⓘ \$

**Launch group** ⓘ

**Availability Zone group** ⓘ

**Request valid from** ⓘ Any time [Edit](#)

**Request valid to** ⓘ Any time [Edit](#)

**Persistent request** ⓘ  Persistent request

- Select the options you want to define your Spot request and to configure your Spot Instance settings.

Option	Description
<b>Max Price</b>	Specifies the maximum price you are willing to pay per instance hour.
<b>Request valid from and Request valid to</b>	<p>Defines the validity period of a Spot request. The validity period, beginning with <b>Request valid from</b> through <b>Request valid to</b>, specifies the length of time that your request will remain valid. By default, a Spot request will be considered for fulfillment from the time it is created until it is either fulfilled or canceled by you. However, you can constrain the validity period using these options.</p> <p><b>Note</b> The end time you specify doesn't apply to the Spot Instances launched by this request. This end time only applies to the Spot Instance request.</p>
<b>Persistent request</b>	Determines whether your request is one-time or persistent. By default, it is one-time. A one-time request can only be fulfilled once. A persistent request is considered for fulfillment whenever there is no Spot Instance running for the same request.
<b>Launch group</b>	Groups a set of requests together in the same launch group. All requests in a launch group have their instances started and terminated together. For more information about using Launch group in your Spot request, see <a href="#">Launching Spot Instances in Launch Groups and Availability Zones</a> (p. 135).

Option	Description
<b>Availability Zone group</b>	Groups a set of requests together in the same Availability Zone. All requests that share an Availability Zone group will launch Spot Instances in the same Availability Zone. For more information about using Availability Zone group in your Spot request, see <a href="#">Launching Spot Instances in Launch Groups and Availability Zones</a> (p. 135).

- Continue with the wizard as you normally would when launching instances. For information, see [Launch Your Instance](#) (p. 266).

If your request specified multiple instances, Amazon EC2 creates a separate Spot Instance request for each instance. The requests are displayed on the **Spot Requests** page.

## Amazon EC2 Command Line Interface (CLI) Tools

### To create a Spot price request

- Enter the following command:

```
PROMPT> ec2-request-spot-instances
--price price
--user-data data
--instance-count count
--type one-time / persistent
[--valid-from timestamp]
[--valid-until timestamp]
[--launch-group launchgroup]
(run-instances-arguments)
```

For example:

```
PROMPT> ec2-request-spot-instances --price 0.32 ami-1234abcd --key MyKeypair
--group webserv --instance-type m1.small --instance-count 10 --type one-time
```

Amazon EC2 returns output similar to the following:

```
SPOTINSTANCEREQUEST    sir-09fb0a04    0.32    one-time    Linux/UNIX
    open    2010-04-06T10:03:09+0200    ami-1234abc    m1.small
MyKeypair    webserv    monitoring-disabled
...
```

If your request specified multiple instances, Amazon EC2 creates a separate Spot Instance request for each instance; each instance has a different ID (e.g., sir-09fb0a04).

## API

### To create a Spot price request

- Construct a Query request similar to the following.

```
https://ec2.amazonaws.com/  
?Action=RequestSpotInstances  
&SpotPrice.1=0.50  
&InstanceCount.1=10  
&Type.1=one-time  
&AvailabilityZoneGroup.1=MyAzGroup  
&LaunchSpecification.ImageId.1=ami-43a4412a  
&LaunchSpecification.KeyName.1=MyKeypair  
&LaunchSpecification.GroupSet.1=websrv  
&LaunchSpecification.InstanceType.1=m1.small  
&...auth parameters...
```

Following is an example response.

```
<RequestSpotInstancesResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <spotInstanceRequestSet>  
    <item>  
      <spotInstanceRequestId>sir-876aff12</spotInstanceRequestId>  
      <spotPrice>0.32</spotPrice>  
      <type>one-time</type>  
      <state>open</state>  
      <fault/>  
      <validFrom/>  
      <validUntil/>  
      <launchGroup/>  
      <availabilityZoneGroup>MyAzGroup</availabilityZoneGroup>  
      <launchSpecification>  
        <imageId>ami-43a4412a</imageId>  
        <keyName>MyKeypair</keyName>  
        <groupSet>  
          <item>  
            <groupId>websrv</groupId>  
          </item>  
        </groupSet>  
        <instanceType>m1.small</instanceType>  
        ...  
      </launchSpecification>  
      <instanceId/>  
      <createTime>2009-10-19T00:00:00+0000</createTime>  
      <productDescription>Linux/UNIX</productDescription>  
      <tagSet/>  
    </item>  
    ...  
  </spotInstanceRequestSet>  
</RequestSpotInstancesResponse>
```

If your request specified multiple instances, Amazon EC2 creates a separate Spot Instance request for each instance; each instance has a different ID (e.g., sir-876aff12).

#### What do you want to do next?

- [Finding Running Spot Instances \(p. 126\)](#)
- [Canceling Spot Instance Requests \(p. 129\)](#)
- [Planning for Interruptions \(p. 142\)](#)



- [Launching Spot Instances in Amazon Virtual Private Cloud \(p. 160\)](#)
- [Tagging Spot Instance Requests \(p. 141\)](#)
- [Subscribe to Your Spot Instance Data Feed \(p. 161\)](#)

## Finding Running Spot Instances

Your Spot request will launch a Spot Instance when the Spot price is below your bid price. A Spot Instance will run until either its maximum bid price is no longer higher than the Spot price or you terminate the Spot Instance yourself. (If your bid price is exactly equal to the Spot price, there is a chance your Spot Instance will remain running.) When your Spot request is fulfilled and has running instances, your Spot request will be *active* (as opposed to *open*, *closed*, or *anceled*). In this section, we describe how to find running Spot Instances associated with your active Spot requests.

### Note

Make sure you have set up the prerequisites for working with Amazon EC2. If you haven't, go to [Prerequisites for Using Spot Instances \(p. 116\)](#).

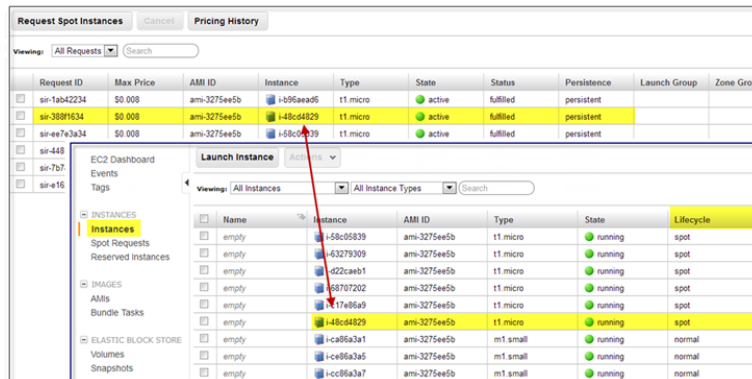
## AWS Management Console

### To find running Spot Instances

1. From the [Amazon EC2 console](#), click **Instances** in the navigation pane.

The console displays a list of running instances.

2. To identify which of the instances were launched as a result of Spot requests, match the Instance IDs in the **Instances** page— where the **Lifecycle** column lists the instance as *spot*—with Instance IDs in the **Spot Requests** page. (You might need to turn on the display of the column by clicking **Show/Hide** in the top right corner.)



## Amazon EC2 Command Line Interface (CLI) Tools

### To find running Tasks Spot Instances

1. Use `ec2-describe-spot-instance-requests`. If your Spot Instance request has been fulfilled (an instance has been launched), the instance ID appears in the response.

```
PROMPT> ec2-describe-spot-instance-requests
SPOTINSTANCEREQUEST      sir-e1471206      0.09      one-time      Linux/UNIX
      active 2010-09-13T16:50:44-0800
      i-992cf7dd      ami-813968c4      m1.small      MyKey      default
      monitoring-disabled
```

- Alternatively, you can use `ec2-describe-instances` with the following filter: `--filter instance-lifecycle=spot`. If you're using the command line interface tools on a Windows system, you might need to use quotation marks (`--filter "instance-lifecycle=spot"`). For more information about filters, see [Listing and Filtering Your Resources \(p. 529\)](#).

```
PROMPT> ec2-describe-instances --filter instance-lifecycle=spot
```

Amazon EC2 returns output similar to the following:

```
RESERVATION      r-b58651f1      111122223333      default
INSTANCE         i-992cf7dd      ami-813968c4      ec2-184-72-8-111.us-west-
1.compute.amazonaws.com  ip-10-166-105-139.us-west-1.compute.internal
  running  MyKey  0      m1.small      2010-09-13T23:54:40+0000
      us-west-1a      aki-a13667e4      ari-a33667e6
  monitoring-disabled  184.72.8.111      10.166.105.139      ebs
      spot      sir-e1471206      paravirtual
```

## API

### To find running Spot Instances

- Construct a `DescribeInstances` Query request and use a filter to look for instances where `instance-lifecycle=spot`. For more information about filters, see [Listing and Filtering Your Resources \(p. 529\)](#).

```
https://ec2.amazonaws.com/
?Action=DescribeInstances
&Filter.1.Name=instance-lifecycle
&Filter.1.Value.1=spot
&...auth parameters...
```

Following is an example response. It includes an `instanceLifecycle` element with `spot` as the value.

```
<DescribeInstancesResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  ...
  <instancesSet>
    <item>
      <instanceId>i-992cf7dd</instanceId>
      <imageId>ami-813968c4</imageId>
      <instanceState>
        <code>16</code>
        <name>running</name>
      </instanceState>
      <privateDnsName>ip-10-166-105-139.&region_sfo;.compute.intern
al</privateDnsName>
      <dnsName>ec2-184-72-8-111.&region_sfo;.compute.&api-domain;</dnsName>
      <reason/>
      <keyName>MyKey</keyName>
      <amiLaunchIndex>0</amiLaunchIndex>
      <productCodes/>
      <instanceType>m1.small</instanceType>
```

```
<launchTime>2010-09-13T23:54:40.000Z</launchTime>
<placement>
  <availabilityZone>&region_sfo;&az1</availabilityZone>
  <groupName/>
</placement>
<kernelId>aki-a13667e4</kernelId>
<ramdiskId>ari-a33667e6</ramdiskId>
<monitoring>
  <state>disabled</state>
</monitoring>
<privateIpAddress>10.166.105.139</privateIpAddress>
<ipAddress>184.72.8.111</ipAddress>
<architecture>i386</architecture>
<rootDeviceType>ebs</rootDeviceType>
<rootDeviceName>/dev/sda1</rootDeviceName>
<blockDeviceMapping>
  <item>
    <deviceName>/dev/sda1</deviceName>
    <ebs>
      <volumeId>vol-61088f0a</volumeId>
      <status>attached</status>
      <attachTime>2010-09-13T23:54:42.000Z</attachTime>
      <deleteOnTermination>>true</deleteOnTermination>
    </ebs>
  </item>
</blockDeviceMapping>
<i>instanceLifecycle</i><instanceLifecycle>spot</instanceLifecycle>
<spotInstanceRequestId>sir-e1471206</spotInstanceRequestId>
<virtualizationType>paravirtual</virtualizationType>
</item>
</instancesSet>
</DescribeInstancesResponse>
```

2. Alternatively, you can use `DescribeSpotInstanceRequests`. If your Spot Instance request has been fulfilled (an instance has been launched), the instance ID appears in the response. Following is an excerpt from a response.

```
...
<spotInstanceRequestSet>
  <item>
    <spotInstanceRequestId>sir-e1471206</spotInstanceRequestId>
    <spotPrice>0.09</spotPrice>
    <type>one-time</type>
    <state>active</state>
    <launchSpecification>
      <imageId>ami-813968c4</imageId>
      <keyName>MyKey</keyName>
      <groupSet>
        <item>
          <groupId>default</groupId>
        </item>
      </groupSet>
      <instanceType>m1.small</instanceType>
      <blockDeviceMapping/>
      <monitoring>
        <enabled>>false</enabled>
```

```
    </monitoring>
  </launchSpecification>
  <instanceId>i-992cf7dd</instanceId>
  <createTime>2010-09-13T23:50:44.000Z</createTime>
  <productDescription>Linux/UNIX</productDescription>
  <launchedAvailabilityZone>&region_iad;&az3;</launchedAvailabilityZone>

</item>
<spotInstanceRequestSet/>
...
```

### What do you want to do next?

- [Canceling Spot Instance Requests](#) (p. 129)
- [Launching Spot Instances in Amazon Virtual Private Cloud](#) (p. 160)
- [Persisting Your Root EBS Partition](#) (p. 142)
- [Tagging Spot Instance Requests](#) (p. 141)
- [Subscribe to Your Spot Instance Data Feed](#) (p. 161)

## Canceling Spot Instance Requests

If you no longer want your Spot request, you can cancel it. You can only cancel Spot Instance requests that are open or active. Your Spot request is *open* when your request has not yet been fulfilled and no instances have been launched. Your Spot request is *active* when your request has been fulfilled, and Spot Instances have launched as a result. If your Spot request is active and has an associated running Spot Instance, canceling the request does not automatically terminate the instance: You must terminate the running Spot Instance manually.

### Note

Make sure you have set up the prerequisites for working with Amazon EC2. If you haven't, go to [Prerequisites for Using Spot Instances](#) (p. 116).

Your Spot request can be in other states in addition to open and active. For more information, see [Spot Instance Request States](#) (p. 134). You can then track the status of each of the requests separately. For information about tracking Spot requests, see [Tracking Spot Requests with Bid Status Codes](#) (p. 135).

A *persistent* Spot request will try to relaunch a Spot Instance whenever it is interrupted or terminated. If you want to end both the Spot request and its Spot Instance, make sure you cancel the Spot request before you terminate the Spot Instance to ensure that the request doesn't launch a new instance.

## AWS Management Console

### To cancel Spot Instance requests

1. From the [Amazon EC2 console](#), click **Spot Requests** in the navigation pane.

The console displays a list of Spot Instance requests.

	Request ID	Max Price	AMI ID	Instance	Type	State	Status
<input type="checkbox"/>	sir-280a7a34	\$0.050	ami-e565ba8c	i-8a2b30e6	m1.small	active	fulfilled
<input type="checkbox"/>	sir-b0277635	\$0.050	ami-e565ba8c	i-a31bb2ce	m1.small	active	fulfilled
<input type="checkbox"/>	sir-95335234	\$0.070	ami-e565ba8c	i-b6b021da	m1.small	active	fulfilled

2. Select the Spot requests you want to cancel and click the **Cancel** button.
3. If you don't want the associated Spot Instances to continue running, you can cancel them. Select them in the **Instances** page, then right-click and select **Terminate**.

## Amazon EC2 Command Line Interface (CLI) Tools

### To cancel Spot Instance requests

1. Enter the following command to see your Spot Instance requests:

```
PROMPT> ec2-describe-spot-instance-requests
```

Amazon EC2 returns output similar to the following:

```
SPOTINSTANCEREQUEST   sir-09fb0a04   0.04   one-time   Linux/UNIX   closed
2010-04-06T10:03:09+0200   ami-b232d0db   m1.small   gsg-keypair   default
monitoring-disabled
```

2. To cancel the request, enter the following:

```
PROMPT> ec2-cancel-spot-instance-requests sir-09fb0a04
```

Amazon EC2 returns output similar to the following:

```
SPOTINSTANCEREQUEST   sir-09fb0a04   canceled
```

3. To terminate the instances associated with the Spot request you just canceled, enter the following:

```
PROMPT> ec2-terminate-instances i-48cd4829
```

Amazon EC2 returns output similar to the following:

```
INSTANCE   i-48cd4829   running shutting-down
```

### Tip

You can filter the list of Spot Instance requests to return only certain EC2 instance types. For more information about how to filter the results, go to [ec2-describe-spot-instance-requests](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

## API

### To cancel Spot Instance requests

1. Construct the following Query request to see your Spot Instance requests.

```
https://ec2.amazonaws.com/  
?Action=DescribeSpotInstanceRequests  
&...auth parameters...
```

Following is an example response.

```
<DescribeSpotInstanceRequestsResponse xmlns="http://ec2.amazon  
aws.com/doc/2013-10-01/">  
  <spotInstanceRequestSet>  
    <item>  
      <spotInstanceRequestId>sir-8675309a</spotInstanceRequestId>  
      <spotPrice>0.32</spotPrice>  
      <type>one-time</type>  
      <state>open</state>  
      <fault/>  
      <validFrom/>  
      <validUntil/>  
      <launchGroup/>  
      <availabilityZoneGroup>MyAzGroup</availabilityZoneGroup>  
      <launchSpecification>  
        <imageId> i-43a4412a</imageId>  
        <keyName>MyKeypair</keyName>  
        <groupSet>webserv</groupSet>  
        <instanceType>m1.small</instanceType>  
      </launchSpecification>  
      <instanceId>i-123345678</instanceId>  
      <createTime>2009-10-19T00:00:00+0000</createTime>  
      <productDescription>Linux/UNIX</productDescription>  
      <launchedAvailabilityZone>&region_iad;&az3;</launchedAvailabilityZone>  
  
    </item>  
  </spotInstanceRequestSet>  
</DescribeSpotInstanceRequestsResponse>
```

2. Construct a Query request to cancel the Spot Instance requests.

```
https://ec2.amazonaws.com/  
?Action=CancelSpotInstanceRequests  
&SpotInstanceRequestId.1=sir-8675309a  
&...auth parameters...
```

Following is an example response.

```
<CancelSpotInstanceRequestsResponse xmlns="http://ec2.amazonaws.com/doc/2013-  
10-01/">  
  <spotInstanceRequestId>sir-8675309a</spotInstanceRequestId>  
</CancelSpotInstanceRequestsResponse>
```

### Tip

You can filter the list of Spot Instance requests to return only certain EC2 instance types. For more information about how to filter the results, go to [DescribeSpotInstanceRequests](#) in the *Amazon Elastic Compute Cloud API Reference*.

### What do you want to do next?

- [Creating a Spot Instance Request](#) (p. 122)
- [Viewing Spot Instance Pricing History](#) (p. 118)
- [Finding Running Spot Instances](#) (p. 126)
- [Launching Spot Instances in Amazon Virtual Private Cloud](#) (p. 160)

## Fundamentals of Spot Instances

If your application can tolerate interruptions and can handle dynamically changing compute capacity and prices, Amazon EC2 Spot Instances can be a powerful tool for optimizing your application's costs and increasing its computational scale. Because Spot Instances are often offered at a large savings compared to regular On-Demand instances, you might be able to take advantage of Spot Instances to both accelerate your computational job and lower its overall cost.

This section presents the basics you need to know so you can create successful Spot requests, manage Spot Instances effectively, and plan for contingencies such as interruptions and changes in capacity.

- A Spot Instance request (also called a *bid*) defines the kind of Spot Instances you want and the maximum price you're willing to pay for them. How you specify your bid determines the fulfillment of the request and the successful launch of Spot Instances. For information, see [Placing Spot Requests](#) (p. 132).
- You can use tags to manage your Spot Instances. For information, see [Tagging Spot Instance Requests](#) (p. 141).
- The availability of Spot Instances varies significantly depending on Amazon EC2 capacity. When you use Spot Instances, you should expect and plan for interruptions and mitigate their effect on your applications. For information, see [Protecting Your Spot Instance Data Against Interruptions](#) (p. 142).

## Placing Spot Requests

The Spot Instance request (or *bid*) is a core element in Amazon EC2 Spot Instances. Your specifications for a Spot request—price, Availability Zone, launch group, and schedule—affect the fulfillment of the request and the successful launch of Spot Instances.

This section discusses the bid fundamentals you need to know to help you manage your Spot requests.

- [Spot Instance Limits](#) (p. 133)

Discusses the number of Spot Instances per instance type that you can launch for each region.

- [Customizing Your Spot Requests](#) (p. 134)

Discusses the options you can specify to customize your Spot request.

- [Tracking Spot Requests with Bid Status Codes](#) (p. 135)

Discusses Spot bid status information and how you can use it to understand what is going on with your Spot requests.

Depending on how quickly you want your Spot request fulfilled and how long you want your Spot Instances to continue running, there are a number of factors you need to consider when you bid on Spot Instances.

The requirements you include in your Spot request can affect the chances that it will be fulfilled. If you have a number of requirements that must be met, your request can take longer to fulfill.

The maximum Spot Instance price you're willing to pay—your *bid price*—may also affect how long your Spot Instances will run. If you want your instances to run as long as possible and have a low probability of interruption, you might consider submitting a higher bid price. If you are less concerned about Spot interruptions and simply want to take advantage of Spot cost savings, you should submit a lower bid price.

Your bid price is not necessarily the price that you pay. For example, if you bid \$0.500, and the Spot price is \$0.300 per hour for the period, you only pay \$0.300 per hour. If the Spot price drops, you pay less. If the Spot price increases, you pay the new price up to your bid price. If the Spot price rises above your bid price (or another capacity constraint is encountered), your Spot Instance is interrupted.

**Note**

Bidding high *does not guarantee* that your Spot Instance won't be interrupted. Also, do not assume that the Spot price can never remain high for long periods of time. If Spot supply is tight for an extended period of time, the Spot price can be elevated for a sustained period of time, pegged to the bid price of the highest bidders.

## Quick Look: Bidding Strategies Video

The following video describes strategies to use when bidding for Spot Instances. [Deciding on Your Spot Bidding Strategies](#)

After you submit your bid for Spot Instances, use bid status codes to track your request. For information, see [Tracking Spot Requests with Bid Status Codes](#) (p. 135).

## Spot Instance Limits

Generally, you are limited to running a total of 100 Spot Instances per region. Certain instance types have a different limit amount per region as shown in the following table. In addition, keep in mind that your Spot Instances can be terminated any time that the Spot market price goes above your bid price, or any time that there is reduced available capacity for the instance type in the region. So, although you might be within your Spot limit level for a particular instance type in the region, you can still get evicted as a result of price or capacity changes.

Instance Type	Spot Limit
g2.2xlarge	10
cg1.4xlarge	10
m2.2xlarge	20
m2.4xlarge	20
cc1.4xlarge	20
cc2.8xlarge	20
hi1.4xlarge	Not Offered
hs1.8xlarge	Not Offered
All Other Instance Types	100

**Note**

The instance types *cc2.8xlarge*, *cg1.4xlarge*, and *cr1.8xlarge* are not available in all regions.



If you need more Spot Instances, complete the [Amazon EC2 instance request form](#) with your use case and your instance increase will be considered. Limit increases are tied to the region specified in the Spot Instance request.

## Customizing Your Spot Requests

You can use a number of options to customize your Spot requests so you get instances that address your needs. These options are discussed in the following topics:

- [Types of Spot Instance Requests \(p. 134\)](#)
- [Launching Spot Instances in Launch Groups and Availability Zones \(p. 135\)](#)

The options you choose can affect the success of your Spot requests as well as the lifespan of the Spot Instances that these requests launch.

## Spot Instance Request States

A Spot Instance request can be in one of the following states:

- **Open**—The request is not fulfilled.
- **Active**—The request is currently active (fulfilled) and has an associated Spot Instance.
- **Failed**—The request failed because bad parameters were specified.
- **Closed**—The request either completed (a Spot Instance was launched and subsequently was interrupted or terminated), or was not fulfilled within the period specified.
- **Canceled**—The request is canceled because one of two events took place: You canceled the request, or the bid request went past its expiration date.

Depending on conditions, an instance might still be running even if the request is closed or canceled.

The **valid until** option that you specify when you submit a Spot Instance request applies only to the *request*; this deadline doesn't apply to Spot Instances that are *launched* by the request. This means that when your Spot Instance request is canceled, either because the request went beyond its expiration date or because you manually canceled it, the Spot Instances that were launched previously through the now-canceled request don't automatically get terminated. The Spot Instance service will terminate a running instance only when the Spot price equals or exceeds your price. However, you can always terminate your instances manually.

You can track the status of your Spot Instance requests as well as the status of the instances launched through bid status codes. For information about tracking Spot requests, see [Tracking Spot Requests with Bid Status Codes \(p. 135\)](#).

## Types of Spot Instance Requests

You can make two types of Spot Instance requests—a one-time request or a persistent request. A one-time request remains active until one of the following conditions is met: All of the requested instances launch, the request expires, or you cancel the request. For example, if you create a one-time request for three instances, the request is considered complete after all three instances launch.

Persistent Spot Instance requests remain active until they expire or you cancel them, even if the requests were previously satisfied. For example, if you create a persistent Spot Instance request for one instance at \$0.300, Amazon EC2 launches or keeps your instance running if your maximum bid price is above \$0.300 and terminates your instance if your maximum bid price is below \$0.300.

With both one-time and persistent requests, instances continue to run until they no longer exceed the Spot price, you terminate them, or the instances terminate on their own. If the maximum price is exactly equal to the Spot price, an instance might or might not continue running (depending on available capacity).

## Launching Spot Instances in Launch Groups and Availability Zones

You can opt to have your Spot Instances launch at the same time or in the same Availability Zone. To tell Amazon EC2 to launch your instances only if all the instances in the request can be fulfilled, specify a launch group in your Spot Instance request.

To specify a launch group using the Amazon EC2 console, you select **Launch Group** in the Request Instances Wizard. If you use the CLI or API tools, you specify the `launch_group` option when you use the `ec2-request-spot-instances` command or call the `RequestSpotInstances` action.

If you want all your instances to be launched together in a single Availability Zone, specify an Availability Zone group. Do this by selecting **Availability Zone Group** when using the Request Instances Wizard in the Amazon EC2 console, or by specifying the `availability-zone-group` option when you use the `ec2-request-spot-instances` CLI command or call the `RequestSpotInstances` API action.

Although the launch group and Availability Zone group specifications can be advantageous, use them only when needed. Avoiding these requirements increases the chances that your Spot Instance request can be fulfilled. Also, specifying a launch group increases the chance that your Spot Instances will be terminated: When one instance in the launch group is terminated, all of the instances in the group are terminated.

To identify a bid price for a Spot Instance in a specific Availability Zone, see [Viewing Spot Instance Pricing History \(p. 118\)](#) for the price over time in the Availability Zone where you want to make your bid. If you use the Amazon EC2 console, go to the **Spot Instance Pricing History** page and select the Availability Zone from the **Zone** drop-down menu. You can also call `DescribeSpotPriceHistory` API action or use the `ec2-describe-spot-price-history` CLI command. You will see the price history for the specified Availability Zone with the most recent set of prices listed first. If you don't specify an Availability Zone, you will get the lowest prices across all Availability Zones for the time period, starting with the most recent set of prices.

### Note

When you use the `DescribeSpotPriceHistory` action or the `ec2-describe-spot-price-history` command before the 2011-05-15 API version, you will get the lowest price across the region for the given time period and the prices will be returned in chronological order—that is, from the oldest to the most recent.

In addition, keep in mind when reviewing Spot prices that past performance is not a guarantee of future results. Spot prices vary based on real-time supply and demand conditions, and the conditions that generated certain Spot prices or pricing patterns in the past may not repeat themselves in the future.

For more information about Availability Zones, see [Region and Availability Zone Concepts \(p. 7\)](#).

## Tracking Spot Requests with Bid Status Codes

### Topics

- [The Life Cycle of a Spot Request \(p. 136\)](#)
- [Spot Bid Status Code Reference \(p. 139\)](#)

Bid statuses can help you track your Amazon Elastic Compute Cloud (EC2) Spot Instance requests, plan your use of Spot Instances, and bid strategically. Spot bid status information is composed of a status code, the update time, and a status message. Together, they help you determine the disposition of your Spot request. For example, a bid status can tell you the reason why your Spot request isn't fulfilled yet or list the constraints that are impeding the fulfillment of your Spot request.

### Note


A *bid* is a *request* that is conditionally tied to a price. The terms are used interchangeably in this documentation.

You can use any of the Amazon EC2 Spot Instance tools to obtain status information about your Spot requests. Spot bid status information will be available for new Spot requests as well as bid requests that you already have open. However, if you're using either the Amazon EC2 CLI or API tools to obtain bid status information, you need to use version 2012-10-01 or later of the API tools. The following table summarizes how you can get bid status information using the Amazon EC2 tools.

Tool	Use	To get
Spot Requests page in the Amazon EC2 console	Status column	Status code
	Bottom pane, tooltip	Status message
Amazon EC2 CLI	<code>ec2-request-spot-instances</code> or <code>ec2-describe-spot-instance-requests</code>	Status code, status message, and update time
	Amazon EC2 API	<code>RequestSpotInstances</code> or <code>DescribeSpotInstanceRequests</code>

### The Life Cycle of a Spot Request

This diagram shows you the paths your Spot request can follow throughout its life cycle. Following the diagram, we explain each step of the life cycle—from making a request to the terminal state.



When you request a Spot Instance, your request goes through an evaluation process. At each step of the process—also called the Spot request *life cycle*—specific events dictate successive request states. In the life cycle diagram, each step is depicted as a node, and each node's status code describes the status of your Spot request and Spot Instance.

1. **Pending evaluation** – As soon as you make your Spot Instance request, it goes into the `pending-evaluation` state.

Status Code	Request State	Instance State
<code>pending-evaluation</code>	<code>open</code>	<code>(none)</code>

The only reason your Spot bid request does not proceed to this `pending-evaluation` state is that at least one of the parameters is invalid (`bad-parameters`).

Status Code	Request State	Instance State
<code>bad-parameters</code>	<code>closed</code>	<code>(none)</code>

2. **Pending evaluation holding state** – Your Spot Instance request will report a status code in the *holding* state if one or more options you specified are valid but not met, or if we encounter a Spot capacity constraint. The options defined by your request affect the likelihood of it being fulfilled. For example, if you specified a max bid price below the current Spot price, or requested that your Spot Instances be launched together in one Availability Zone (AZ), your Spot request will stay in a *holding* state until the Spot price goes below your bid or the Availability Zone constraint is met.

The following table lists the status codes reported when your bid is in a *holding state* waiting for a constraint to be met.

Status Code	Request State	Instance State
<code>capacity-not-available</code>	<code>open</code>	<code>(none)</code>
<code>capacity-oversubscribed</code>	<code>open</code>	<code>(none)</code>
<code>price-too-low</code>	<code>open</code>	<code>(none)</code>
<code>not-scheduled-yet</code>	<code>open</code>	<code>(none)</code>
<code>launch-group-constraint</code>	<code>open</code>	<code>(none)</code>
<code>az-group-constraint</code>	<code>open</code>	<code>(none)</code>
<code>placement-group-constraint</code>	<code>open</code>	<code>(none)</code>
<code>constraint-not-fulfillable</code>	<code>open</code>	<code>(none)</code>

3. **Pending evaluation/fulfillment terminal state** – Before your request reaches the pending fulfillment phase, your Spot Instance request may end up in a *terminal* state if the valid dates you specified for the bid expired, you canceled the bid yourself, or a system error occurred.

**Note**

A *canceled* Spot request does not necessarily mean that the Spot Instance is terminated. In contrast, a *closed* Spot request means the request will no longer be considered and the Spot Instance is terminated.

The following table lists the status codes reported when your bid ends up in the *terminal* state after pending evaluation or pending fulfillment.

Status Code	Request State	Instance State
schedule-expired	closed	(none)
canceled-before-fulfillment*	canceled	(none)
bad-parameters	failed	(none)
system-error	closed	(none)

\* This status results from a cancellation initiated by a user.

4. **Pending fulfillment** – When the constraints you specified (if any) are met, and your bid price is equal to or higher than the current Spot market price, your Spot request goes into the `pending-fulfillment` state.

Status Code	Request State	Instance State
pending-fulfillment	open	(none)

At this point, the Amazon EC2 Spot service is getting ready to provision the instances that you requested. If the process stops at this point, it would likely be due to cancellation by the user before a Spot Instance was launched, or an unexpected system error (see *Pending evaluation/fulfillment terminal state*).

5. **Fulfilled** – When all the specifications for your Spot Instance are met, your Spot request is fulfilled. The Amazon EC2 Spot service launches the Spot Instance, which may take a few minutes.

Status Code	Request State	Instance State
fulfilled	active	starting up → running

At this point, your instance launches, and your bid request is `active`. Your instance will continue to run as long as your bid price is at or above the Spot market price, EC2 has spare Spot capacity for your instance type, and you don't terminate the instance.

**Note**

If your bid price is equal to the market price, your request can stay `fulfilled`, but you run the risk of your Spot Instance getting terminated at any time that the market price goes up. In some cases, your Spot Instance could be evicted even if your bid equaled the Spot price because Spot Instances were oversubscribed at that price. In such cases, you will get the `instance-terminated-capacity-oversubscribed` status code. See *Fulfilled terminal state*.

6. **Fulfilled terminal state** – Your Spot request will report a status code in the *terminal* state if a change in Spot price or capacity requires the Spot service to terminate your Spot Instance. You will also get a terminal state status code if you cancel the Spot request or terminate the Spot Instance.

The following table lists the status codes reported when your instance is in the *terminal* state after the bid was fulfilled.

Status Code	Request State	Instance State
request-canceled-and-instance-running	canceled[c]	running

Status Code	Request State	Instance State
marked-for-termination	closed	running
instance-terminated-by-price	closed[a]	(Spot) terminated
instance-terminated-by-user (or spot-instance-terminated-by-user, which will be deprecated)	closed[b]	terminated[c]
instance-terminated-no-capacity	closed[a]	(Spot) terminated
<del>instance-terminated-capacity-oversubscribed</del>	closed[a]	(Spot) terminated
<del>instance-terminated-launch-gpu-constraint</del>	closed[a]	(Spot) terminated

[a] If the Spot bid request is persistent, the state of the request becomes *open*. See the next item, Persistent requests.

[b] If you terminate your instance first, there may be a delay before the Spot service realizes that your Spot Instance was terminated. For persistent requests, this means that for user-terminated instances it could take some time before your Spot request re-enters the `pending-evaluation` state.

[c] Actions are initiated by the user.

7. **Persistent requests** – When your Spot Instances are terminated (either by you or the Amazon EC2 Spot service), if you specified a *persistent* Spot request, then the request will be considered again. This request returns to the `pending-evaluation` state. Your Spot bid request will undergo the same evaluation process that it originally went through, and an entirely new Spot Instance will be launched for the persistent Spot request.

### Spot Bid Status Code Reference

The following list contains Spot bid status codes and explanations about what the code means for your Spot request:

- **az-group-constraint**

The Spot service cannot launch all the instances you requested in the same Availability Zone.

- **bad-parameters**

One or more of the parameters of your Spot request was invalid or malformed (for example, the AMI you specified may not exist). The bid status message will tell you which parameter is invalid.

- **canceled-before-fulfillment**

The user canceled the Spot request before it was fulfilled.

- **capacity-not-available**

There is no capacity available for the instances you requested.

- **capacity-oversubscribed**

The number of Spot requests with bid prices equal to or higher than your bid price exceeds the available capacity in this pool.

- **constraint-not-fulfillable**

The Spot request cannot be fulfilled because one or more constraints are invalid. For example, you provided an invalid *from* or *to* or an invalid Availability Zone. The bid status message will tell you which constraint was invalid.

- **fulfilled**

Your bid request was fulfilled, so your Spot request is active, and the Amazon EC2 Spot service is launching (or has launched) your Spot Instance.

- **instance-terminated-by-user (or spot-instance-terminated-by-user)**

You terminated a Spot instance that had been fulfilled, so the bid state becomes closed (unless it's a persistent bid), and the instance state is terminated.

**Note**

The status codes `instance-terminated-by-user` and `spot-instance-terminated-by-user` represent the same status. The code `spot-instance-terminated-by-user` will be deprecated.

- **instance-terminated-by-price**

The Spot price rose above your bid price. If your request is a persistent bid, the process—or *life cycle*—restarts and your bid will again be pending evaluation.

- **instance-terminated-capacity-oversubscribed**

Your instance was terminated because the number of Spot requests with bid prices equal to or higher than your bid price has exceeded the available capacity in this pool. This means that your instance was interrupted even though the Spot price may not have changed because your bid was at the Spot price.

**Note**

When there are several requests at a certain bid price and only some of them need to be terminated, the Spot service randomly selects the requests to be terminated.

- **instance-terminated-launch-group-constraint**

One of the instances in your launch group was terminated, so the launch group constraint is no longer fulfilled.

- **instance-terminated-no-capacity**

There is no longer any Spot capacity available for the instance.

- **launch-group-constraint**

The Spot service cannot launch all the instances you requested at the same time.

**Note**

All requests in a launch group have their instances started and terminated together. For more information about using Launch Group in your Spot request, see [Launching Spot Instances in Launch Groups and Availability Zones \(p. 135\)](#).

- **marked-for-termination**

Your Spot Instance is marked for termination.

- **not-scheduled-yet**

Your Spot request will not be evaluated until the scheduled date.

- **pending-evaluation**

As soon as you make your Spot Instance request, it goes into the pending-evaluation state while the system evaluates the parameters of your request.

- **pending-fulfillment**

Your Spot request is pending fulfillment while the system tries to provision your Spot Instance.

- **placement-group-constraint**

Your Spot request cannot be fulfilled yet because a new Spot Instance cannot currently be added to the placement group you specified.

- **price-too-low**

You have a bid request that cannot be fulfilled because the bid price is below the Spot price. In this case, no instance is launched, and your bid remains open.

- **request-canceled-instance-running**

You canceled your Spot request while the Spot instance was still running, so the request is canceled, but the instance remains running.

- **schedule-expired**

Your Spot request has expired because it was not fulfilled before the valid *to* date.

- **system-error**

There was an unexpected system error. If this is a recurring issue, please contact customer support for assistance.

## Tagging Spot Instance Requests

To help categorize and manage your Spot Instance requests, you can tag them with metadata of your choice. You tag your Spot Instance requests in the same way that you tag other Amazon EC2 resources. Create a tag by specifying a key and value, and assigning the pair to the Spot Instance request. You use the [CreateTags](#) action or the [ec2-create-tags](#) command. For information about how to use tags, see [Tagging Your Amazon EC2 Resources \(p. 532\)](#).

The tags you create for your Spot Instance requests only apply to the requests. These tags don't propagate to the instances that the requests launch. To categorize the Spot Instances that are launched from your tagged request, you must create a separate set of tags for the instances using the same commands.

For example, you have Spot Instance request *sir-69047e11* and you want to label it with the tag `Spot=Test`. To do this, use the following command:

```
PROMPT> ec2-create-tags sir-69047e11 --tag "Spot=Test"
```

Amazon EC2 returns this information:

```
TAG      spot-instance-request    sir-69047e11    Spot    Test
```

You can also confirm the tag information by using the `ec2-describe-tags` command.

When your request is fulfilled and a Spot Instance launches, you will see that the tag you used for your Spot Instance request is not applied to the Spot Instance. In the following example command, we are obtaining information about the Spot Instance *i-b8ca48d8* that was launched as a result of your Spot Instance request *sir-69047e11*, tagged `Spot=Test`.

```
PROMPT> ec2-describe-instances i-b8ca48d8
```

The call returns details about the instance with no tag information, showing that the tag for your Spot Instance request does not apply to the Spot Instance that the request launched. To tag your Spot Instance, use the following command:



```
PROMPT> ec2-create-tags i-b8ca48d8 --tag "SpotI=Test1"
```

Amazon EC2 returns this information:

TAG	instance	i-b8ca48d8	SpotI	Test1
-----	----------	------------	-------	-------

You can create similar tags using the API. For more information, see the API section of [Tagging Your Amazon EC2 Resources](#) (p. 532).

## Protecting Your Spot Instance Data Against Interruptions

Customer demand for Amazon EC2 Spot Instances can vary significantly from moment to moment, and the availability of Spot Instances can also vary significantly depending on how much Amazon EC2 capacity is available. Therefore, it is always possible that your Spot instance will be interrupted if another user bids higher than you or if we have to reclaim spare capacity.

### Note

No matter how high you bid, there is always a risk that your Spot Instance will be interrupted. We strongly recommend against bidding above the On-Demand price or using Spot for applications that cannot tolerate interruptions.

The following sections discuss how to plan for and mitigate the impact of interruptions.

- [Planning for Interruptions](#) (p. 142)
- [Persisting Your Root EBS Partition](#) (p. 142)

## Planning for Interruptions

### Topics

- [Quick Look: Managing Interruptions Video](#) (p. 142)

Because Spot Instances can terminate at any time, applications that run on Spot Instances must terminate cleanly. Although we attempt to cleanly terminate your instances, your application should be prepared to deal with an immediate shutdown.

To test your application, launch and terminate it as an On-Demand instance.

### Quick Look: Managing Interruptions Video

The following video shows how some customers manage the interruption of their Spot instances. [How to Manage Spot Instance Interruptions](#)

For more information about strategies to manage interruptions, go to the following sections:

- [Launching Amazon Elastic MapReduce Job Flows with Spot Instances](#) (p. 159)
- [Using Auto Scaling to Get Notifications for Spot Instances](#) (p. 156)
- [Starting Clusters on Spot Instances](#) (p. 191)

## Persisting Your Root EBS Partition

Amazon Elastic Block Store (Amazon EBS) can be an effective way to store data that you otherwise might lose when your Spot Instance terminates.

### Important

Although Spot Instances can use Amazon EBS-backed AMIs, be aware that Spot Instances do not support the Stop/Start feature. In other words, you can't stop and start Spot Instances launched from an AMI that has an Amazon EBS root device. For information about Amazon EBS, see [Amazon Elastic Block Store \(Amazon EBS\)](#).

To set up the persistence of Spot Instance data, you map the Spot Instances that will be launched to an existing Amazon Elastic Block Store (Amazon EBS) snapshot. Set the `delete-on-termination` flag to `false`; this indicates that Amazon EC2 shouldn't delete the Amazon EBS volume when the spot instance terminates.

Let's walk through making an example Spot request with the following specifications:

- Bid price of \$0.500
- One instance of the m1.xlarge instance type
- Block device mapping to a snapshot that shouldn't be deleted when the Spot Instance is terminated

You can do this example using either the CLI or API tools. Using the CLI, your example request should look like this:

```
PROMPT> ec2-request-spot-instances -p 0.5 -t m1.xlarge -n 1 -b '/dev/sdb=snap-a123bcde:20:false' ami-8e1fece7
```

For more information, see:

- [Block Device Mapping \(p. 517\)](#)
- [ec2-request-spot-instances](#)
- [RequestSpotInstances](#)

## Walkthroughs: Using Spot Instances with AWS Services

### Topics

- [Managing Spot Instances with Auto Scaling \(p. 143\)](#)
- [Using CloudFormation Templates to Launch Spot Instances \(p. 158\)](#)
- [Launching Amazon Elastic MapReduce Job Flows with Spot Instances \(p. 159\)](#)
- [Launching Spot Instances in Amazon Virtual Private Cloud \(p. 160\)](#)

You can use AWS services with Spot Instances. In this section, we will show you how Amazon EC2 Spot Instances works with services, such as Auto Scaling, Elastic MapReduce, and Amazon Virtual Private Cloud (Amazon VPC).

## Managing Spot Instances with Auto Scaling

### Topics

- [Tools for Managing Auto Scaling with Spot Instances \(p. 144\)](#)
- [Launching Spot Instances with Auto Scaling \(p. 146\)](#)
- [Obtaining Information About the Instances Launched by Auto Scaling \(p. 149\)](#)
- [Updating the Bid Price for the Spot Instances \(p. 153\)](#)
- [Scheduling Spot Bid Requests \(p. 155\)](#)

- [Using Auto Scaling to Get Notifications for Spot Instances \(p. 156\)](#)

You can take advantage of Auto Scaling features to manage your Amazon EC2 Spot Instances. With Auto Scaling, you can scale up or down your capacity based on demand by setting up Auto Scaling to make Spot bids on your behalf.

In addition, you can use Auto Scaling's scheduling functionality for more granular control over when to bid and launch your Spot Instances. You also can use an Amazon Simple Notification Service (Amazon SNS)-backed Auto Scaling feature that sends notifications each time specified events—such as the launch or termination of instances—take place. Using Auto Scaling's scheduled actions, you can set up bids to expire at a set time.

When you use Auto Scaling with Spot Instances, there are some Auto Scaling tools that you have to use instead of the Spot Instance tools that might be familiar. There are also a number of Spot Instance options that you cannot use. Here's a summary of differences:

- **Setting your bid price.** When you use Auto Scaling to launch Spot Instances, you set your bid price in an Auto Scaling launch configuration.
- **Spot market price and your bid price.** This is the same as current behavior. If the market price for Spot Instances rises above your bid price for a running instance, Amazon EC2 will terminate your instance.
- **Changing your bid price.** If you want to change your bid price, you have to create a new launch configuration and associate it with your Auto Scaling group. You cannot update the existing launch configuration.
- **New bid price and running instances.** When you change your bid price by creating a new launch configuration, instances already launched will continue to run as long as the bid price for those running instances is higher than the current market price for Spot Instances.
- **Spot and Auto Scaling instance termination.** Amazon EC2 terminates a Spot Instance when the bid price for that instance falls below the Spot market price. Auto Scaling terminates instances based on a combination of criteria, including the launch configuration it is associated with, length of time in a billing hour the instance has been running, and the Availability Zone in which it is launched. For more information about instance termination in Auto Scaling, see [Auto Scaling Instance Termination](#).
- **Maintaining your Spot Instances.** When your instance is terminated, Auto Scaling will try to launch another instance to replace it in order to maintain your specified desired capacity. However, whether or not Auto Scaling successfully launches an instance depends on the bid price as compared to the Spot market price: If the bid price is higher than the Spot market price, then an instance will be launched; if the Spot market price is higher than the bid price, then no instance will be launched at that point.
- **Persistent bids.** Auto Scaling will continue submitting bids for you as long as you keep your Auto Scaling group and launch configuration. The Auto Scaling group specifies the desired number of instances you want maintained and the launch configuration specifies the bid price for your Spot Instances.

For more information about Auto Scaling, see [Auto Scaling Developer Guide](#). To get started quickly with Auto Scaling, see [Auto Scaling Getting Started Guide](#).

## Tools for Managing Auto Scaling with Spot Instances

You will use the Auto Scaling command line interface (CLI) tools to use the Auto Scaling features with Spot Instances. Currently, the AWS Management Console does not have support to create and manage Auto Scaling objects. However, you can view most details about bids and Amazon EC2 instances launched by Auto Scaling on the AWS Management Console.

To use the Auto Scaling CLI tools, download and unzip the package, and set up the tools on your computer just as you set up your Amazon EC2 CLI tools. The Auto Scaling CLI tools and the Amazon EC2 CLI tools are shipped as different tools packages.

The following table lists resources that can help you get started with the Auto Scaling CLI tools. For more information, go to [Auto Scaling Tools](#) in the *Auto Scaling Developer Guide*.

### Resources for Using Auto Scaling CLI Tools

Task	Resource
Download the Auto Scaling command line interface tools	<a href="#">Auto Scaling Command Line Tool</a> on the Amazon Web Services site.
Setup and Install Auto Scaling command line interface tools	<a href="#">Install the Command Line Interface</a> in the <i>Auto Scaling Developer Guide</i> .
Read the readme installation instructions	<i>Install Drive\AutoScaling-n.n.nn.n\README</i>

### To get help about Auto Scaling CLI commands

After you have installed the Auto Scaling CLI tools, you can get more information about the commands by querying the command line.

1. Open the Command Line.
2. Navigate to *Install Drive\AutoScaling-n.n.nn.n\bin*.
3. Type `as-cmd` and press **Enter**.

The command line returns a list of Auto Scaling CLI commands and a short description of what each command does.

For example, here are some of the Auto Scaling CLI commands we will use in this section.

Command	Description
<code>as-create-launch-config</code>	Creates a new launch configuration with specified attributes.
<code>as-create-auto-scaling-group</code>	Creates a new Auto Scaling group with specified name and other attributes.
<code>as-describe-auto-scaling-groups</code>	Describes the Auto Scaling groups, if the groups exist.
<code>as-describe-auto-scaling-instances</code>	Describes the Auto Scaling instances, if the instances exist.
<code>as-describe-scaling-activities</code>	Describes a set of activities or all activities belonging to a group.
<code>as-delete-auto-scaling-group</code>	Deletes the specified auto scaling group, if the group has no instances and no scaling activities are in progress.

For common conventions the documentation uses to represent command symbols, see [Document Conventions](#).

4. Type `as-command-name --help`.

The command line returns a description, syntax information, and examples of how to use the specified Auto Scaling CLI command.

For Amazon EC2 CLI tools, see the following:

- **Amazon EC2 CLI tools.** [Amazon Elastic Compute Cloud CLI Reference](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.
- **Amazon EC2 CLI tools specific to Spot Instances.** [Amazon EC2 Command Line Interface \(CLI\) Tools \(p. 117\)](#) in the *Amazon Elastic Compute Cloud User Guide*.
- **Other Amazon EC2 tools you can use with Spot instances.** [Getting Started with Spot Instances \(p. 116\)](#) in the *Amazon Elastic Compute Cloud User Guide*.

## Launching Spot Instances with Auto Scaling

In this section, we will create an Auto Scaling launch configuration and an Auto Scaling group that launch Spot Instances. We will use the Auto Scaling command line interface (CLI) commands to create the launch configuration and the Auto Scaling group, and to verify and obtain information about them and the Amazon EC2 instances that they launch.

### Prerequisites

If you're not familiar with how to create a launch configuration or an Auto Scaling group, we recommend that you go through the steps in the [Basic Auto Scaling Configuration](#) in the *Auto Scaling Developer Guide*. Use the basic scenario to get started with the infrastructure that you need in most Auto Scaling scenarios.

If you don't have the Auto Scaling CLI tools installed on your computer, you must install them to do this walkthrough. For information, see [Tools for Managing Auto Scaling with Spot Instances](#) in the *Amazon Elastic Compute Cloud User Guide*. You can also use the information in [Install the Command Line Interface](#) in the *Auto Scaling Developer Guide*.

In this scenario, we will perform the following tasks:

- [Step 1: Create a Launch Configuration \(p. 146\)](#)
- [Step 2: Create an Auto Scaling Group \(p. 147\)](#)

If you already have a launch configuration and Auto Scaling group, here are other related Spot Instance tasks that you can do using Auto Scaling:

- [Updating the Bid Price for the Spot Instances \(p. 153\)](#)
- [Scheduling Spot Bid Requests \(p. 155\)](#)
- [Using Auto Scaling to Get Notifications for Spot Instances \(p. 156\)](#)
- [Advanced Tasks \(p. 161\)](#)

For more information about Auto Scaling, see [What is Auto Scaling?](#) in the *Auto Scaling Developer Guide*. For information about scenarios using Auto Scaling, see [Using Auto Scaling](#) also in the *Auto Scaling Developer Guide*.

### Step 1: Create a Launch Configuration

An Auto Scaling launch configuration is a template that contains the parameters necessary to launch new Amazon EC2 instances. You can attach only one launch configuration to an Auto Scaling group at a time. You have to create a launch configuration and then an Auto Scaling group in order to launch instances with Auto Scaling.

To place bids for Spot Instances, use the `as-create-launch-config` Auto Scaling CLI command with the `--spot-price` option. Specify the maximum price you want to pay for an instance. This price will be used by Auto Scaling to bid for your instances, but this price is not necessarily what you pay for the

instances when they are launched. You will pay the Spot price. For example, you bid \$0.05 for m1.small instances. Your bid gets fulfilled if the current market price for m1.small Spot Instance is \$0.03, or any other price below \$0.05, and you will be charged the current price of \$0.03 per hour. In fact, as long as your current bid is higher than the market price, your bid will be fulfilled and a Spot Instance will be launched for you.

You can get guidance on the price to bid by checking Spot price history. You can do this by using the AWS Management Console, CLI, or API. For more information, go to [Viewing Spot Instance Pricing History](#) (p. 118).

The `as-create-launch-config` command takes the following arguments:

```
as-create-launch-config LaunchConfigurationName --image-id value --instance-type
value [--spot-price value] [--iam-instance-profile value] [--block-device-mapping
"key1=value1,key2=value2..." ] [--monitoring-enabled|--monitoring-disabled]
[--kernel value ] [--key value ] [--ramdisk value] [--group value[,value...]
] [--user-data value] [--user-data-file value] [General Options]
```

The only required options are the launch configuration name, image ID, and instance type. For this walkthrough, specify:

- Launch configuration name = `spot1c-5cents`

**Note**

When Auto Scaling launches instances, it does not distinguish the Spot Instances from the On-Demand instances. To help you identify which of your instances are Spot Instances, consider assigning a launch configuration name that includes *spot* and the bid price.

- Image ID = `ami-e565ba8c`

The Image ID identifies the Amazon Machine Image (AMI) that will be used to launch the instances. If you don't have an AMI, and you want to find a suitable one, see [Amazon Machine Images \(AMIs\)](#).

- Instance type = `m1.small`
- Spot price = `$0.05`

This parameter is optional. If you want to launch Spot Instances, you must specify the Spot bid price that will be used to bid for Spot Instances. Spot Instance bid prices must be specified in US dollars.

Your command should look similar to the following example:

```
as-create-launch-config spot1c-5cents --image-id ami-e565ba8c --instance-type
m1.small --spot-price "0.05"
```

You should get a confirmation like the following example:

```
OK-Created launch config
```

## Step 2: Create an Auto Scaling Group

An Auto Scaling group is a collection of Amazon EC2 instances that shares similar characteristics and to which you want to apply certain scaling actions. You can use the Auto Scaling group to automatically scale the number of instances or maintain a fixed number of instances. You can attach only one launch configuration to an Auto Scaling group at a time. You have to create a launch configuration and then an Auto Scaling group in order to launch Amazon EC2 instances with Auto Scaling.

The `as-create-auto-scaling-group` command takes the following arguments:

```
as-create-auto-scaling-group AutoScalingGroupName --availability-zones
value[,value...] --launch-configuration value --max-size value --min-size value
[--default-cooldown value] [--desired-capacity value] [--grace-period value]
[--health-check-type value] [--load-balancers value[, value]] [--placement-group
value] [--vpc-zone-identifier value] [General Options]
```

This command requires that you specify a name for your Auto Scaling group, a launch configuration, one or more Availability Zones, a minimum group size, and a maximum group size. The Availability Zones you choose determine the physical location of your Auto Scaling instances. The minimum and maximum group size tells Auto Scaling the minimum and maximum number of instances the Auto Scaling group should have.

Desired capacity is an important component of the `as-create-auto-scaling-group` command. Although it is an optional parameter, desired capacity tells Auto Scaling the number of instances you want to run initially. To adjust the number of instances you want running in your Auto Scaling group, you change the value of `--desired-capacity`. If you don't specify `--desired-capacity`, its value is the same as minimum group size.

For more detailed information on the syntax of the `as-create-auto-scaling-group` command, see [Basic Auto Scaling Configuration](#) in the *Auto Scaling Developer Guide*. You can also get help information from the command line: Run the `as-create-auto-scaling-group --help`. For more information, go to [Resources for Using Auto Scaling CLI Tools](#) (p. 145).

For this walkthrough, specify these values for the command:

- Auto Scaling Group name = `spotasg`

**Note**

When Auto Scaling launches instances, it does not distinguish the Spot Instances from the On-Demand instances. To help you identify which of your instances are Spot Instances, consider assigning spot-specific names to the Auto Scaling group that launches Spot Instances.

- Launch configuration name = `spotlc-5cents`
- Availability Zone = `us-east-1a,us-east-1b`
- Max size = 5
- Min size = 1
- Desired capacity = 3

Your command should look like the following example:

```
as-create-auto-scaling-group spotasg --launch-configuration spotlc-5cents --
availability-zones "us-east-1a,us-east-1b" --max-size 5 --min-size 1 --desired-
capacity 3
```

You should get confirmation like the following example:

```
OK-Created AutoScalingGroup
```

**What do you want to do next?**

- [Obtaining Information About the Instances Launched by Auto Scaling](#) (p. 149)
- [Updating the Bid Price for the Spot Instances](#) (p. 153)
- [Scheduling Spot Bid Requests](#) (p. 155)
- [Using Auto Scaling to Get Notifications for Spot Instances](#) (p. 156)



- [Advanced Tasks \(p. 161\)](#)

For more information about Auto Scaling, see [What is Auto Scaling?](#) in the *Auto Scaling Developer Guide*. For information about scenarios using Auto Scaling, see [Using Auto Scaling](#) also in the *Auto Scaling Developer Guide*.

## Obtaining Information About the Instances Launched by Auto Scaling

You can use the Auto Scaling CLI tools to obtain information about your launch configuration, Auto Scaling groups and Amazon EC2 instances launched by Auto Scaling.

In this section, we will use the following Auto Scaling CLI commands to get information about the Spot price bids and Spot Instances that Auto scaling makes and launches for you.

- `as-describe-scaling-activities`—You can use the information about Auto Scaling activities that this command returns to check the status of the bids submitted for you by Auto Scaling.
- `as-describe-auto-scaling-groups`—You can use the information about Auto Scaling groups that this command returns to confirm that Auto Scaling is launching your Spot Instances according to your specifications.

### To check the status of the bids that Auto Scaling is making for you

The `as-describe-scaling-activities` command lists the activities that Auto Scaling performed for a specified Auto Scaling group.

This is the basic syntax:

```
as-describe-scaling-activities [ActivityIds [ActivityIds ...]]  
[--auto-scaling-group value] [--max-records value] [General Options]
```

Specifying the Auto Scaling group and the Activity ID are optional. If you don't specify the Auto Scaling group, the command will return all activities for all Auto Scaling groups associated with your account. If you specify the Auto Scaling group, only the activities for that Auto Scaling group will be listed.

In this scenario, we are using the `as-describe-scaling-activities` command to see state of your bid. Assume that there is only one Auto Scaling group (spotasg) and you want to list all activities.

1. Open a command line and navigate to the bin folder of your Auto Scaling CLI tools directory.
2. Type the command: `as-describe-scaling-activities --auto-scaling-group spotasg --headers`

The information you get should be similar to the following example.

```
ACTIVITY  ACTIVITY-ID          END-TIME          GROUP-  
NAME  CODE          MESSAGE  
ACTIVITY  31bcbb67-7f50-4b88-ae7e-e564a8c80a90          spotasg  
    WaitingForSpotInstanceId  Placed Spot instance request: sir-fc8a3014.  
    Waiting for instance(s)  
ACTIVITY  770bbeb5-407c-404c-a826-856f65db1c57          spotasg  
    WaitingForSpotInstanceId  Placed Spot instance request: sir-69101014.  
    Waiting for instance(s)  
ACTIVITY  597e4ebd-220e-42bc-8ac9-2bae4d20b8d7  2012-05-23T17:40:22Z  spotasg  
    Successful
```



In this response, you know that your bids were placed, one of the bids is successful, and Auto Scaling is waiting for the other two bids.

**Note**

If the `as-describe-scaling-activities` command returns a list that includes *Failed* activities, check the data you specified in the launch configuration. For example:

- The Amazon Machine Image (AMI) might not be valid anymore.
- The bid price specified in the launch configuration is lower than the Spot market price.

3. If you run the `as-describe-scaling-activities` command again later, you can be getting information that is similar to the following example:

ACTIVITY NAME	ACTIVITY-ID CODE	END-TIME	GROUP-
ACTIVITY	90630906-b40f-41a6-967a-cd6534b2dfca	2012-06-01T02:32:15Z	spotasg
	Successful		
ACTIVITY	a1139948-ad0c-4600-9efe-9dab8ce23615	2012-06-01T00:48:02Z	spotasg
	Successful		
ACTIVITY	33001e70-6659-4494-a817-674d1b7a2f58	2012-06-01T02:31:11Z	spotasg
	Successful		

The output shows that the listed activities were successful. Because we know that *spotasg* is an Auto Scaling group that uses a launch configuration with a spot bid price, you can assume that the activities represent the bids placed by Auto Scaling.

4. If you want to get more details about the activities and instances, use the `--show-xml` option of `as-describe-scaling-activities`. Enter the following command  
`as-describe-scaling-activities --auto-scaling-group spotasg --show-xml`.

The information you get should be similar to the following example.

```
<DescribeScalingActivitiesResponse xmlns="http://autoscaling.&api-do
main:/doc/2011-01-01/">
  <DescribeScalingActivitiesResult>
    <NextToken>b5a3b43e-10c6-4b61-8e41-2756db1fb8f5</NextToken>
    <Activities>
      <member>
        <StatusCode>Successful</StatusCode>
        <Progress>0</Progress>
        <ActivityId>90630906-b40f-41a6-967a-cd6534b2dfca</ActivityId>
        <StartTime>2012-06-01T00:48:21.942Z</StartTime>
        <AutoScalingGroupName>spotasg</AutoScalingGroupName>
        <Cause>At 2012-06-01T00:48:21Z a difference between desired and ac
tual capacity changing the desired capacity, increasing the capacity from
2 to 3.</Cause>
        <Details>{}</Details>
        <Description>Launching a new EC2 instance: i-fe30d187</Description>

        <EndTime>2012-06-01T02:32:15Z</EndTime>
      </member>
      <member>
        <StatusCode>Successful</StatusCode>
        <Progress>0</Progress>
        <ActivityId>a1139948-ad0c-4600-9efe-9dab8ce23615</ActivityId>
        <StartTime>2012-06-01T00:47:51.293Z</StartTime>
        <AutoScalingGroupName>spotasg</AutoScalingGroupName>
```

```
<Cause>At 2012-06-01T00:47:51Z an instance was taken out of service
in response to a system health-check.</Cause>
<Details>{}</Details>
<Description>Terminating EC2 instance: i-88ce28f1</Description>
<EndTime>2012-06-01T00:48:02Z</EndTime>
</member>
<member>
  <StatusCode>Successful</StatusCode>
  <Progress>0</Progress>
  <ActivityId>33001e70-6659-4494-a817-674dlb7a2f58</ActivityId>
  <StartTime>2012-06-01T00:46:19.723Z</StartTime>
  <AutoScalingGroupName>spotasg</AutoScalingGroupName>
  <Cause>At 2012-06-01T00:46:19Z a difference between desired and ac
tual capacity changing the desired capacity, increasing the capacity from
2 to 3.</Cause>
  <Details>{}</Details>
  <Description>Launching a new EC2 instance: i-2c30d155</Description>

  <EndTime>2012-06-01T02:31:11Z</EndTime>
</member>
...
</Activities>
</DescribeScalingActivitiesResult>
<ResponseMetadata>
  <RequestId>d02af4bc-ad8f-11e1-85db-83e1968c7d8d</RequestId>
</ResponseMetadata>
</DescribeScalingActivitiesResponse>
```

The XML output shows more detail about the Spot and Auto Scaling activity.

- Cause: At 2012-06-01T00:48:21Z a difference between desired and actual capacity changing the desired capacity, increasing the capacity from 2 to 3. Description: Launching a new EC2 instance: i-fe30d187

If an instance is terminated and the number of instances falls below the desired capacity, Auto Scaling will launch a new instance so that the total number of your running instances rises back to the level specified for desired capacity.

- Cause: At 2012-06-01T00:47:51Z an instance was taken out of service in response to a system health-check. Description: Terminating EC2 instance: i-88ce28f1

Auto Scaling maintains the desired number of instances by monitoring the health status of the instances in the Auto Scaling group. When Auto Scaling receives notification that an instance is *unhealthy* or terminated, Auto Scaling launches another instance to take the place of the unhealthy instance. For information, see [Configure Health Checks for Your Auto Scaling group](#) in the *Auto Scaling Developer Guide*.

**Note**

Auto Scaling provides the cause of instance termination that is not the result of a scaling activity. This includes instances that have been terminated because the Spot market price exceeded their bid price.

When Auto Scaling attempts to replace terminated instances resulting from the Spot market price rising above the instances' bid price, Auto Scaling will place the bid specified in the current launch configuration and attempt to launch another instance to maintain the desired capacity.

## To confirm that Auto Scaling is launching your Spot Instances according to your specifications

Use `as-describe-auto-scaling-groups`. The command will show details about the group and instances launched. For information about the `as-describe-auto-scaling-groups` command, see [Verify Auto Scaling Group Creation](#) in the *Auto Scaling Developer Guide*.

1. Open a command line and navigate to the bin folder of your Auto Scaling CLI tools directory.
2. Type the command:`as-describe-auto-scaling-groups spotasg --headers`

### Note

The `--headers` option supplies the column name so you know what data is being returned.

The information you get should be similar to the following example.

AUTO-SCALING-GROUP	GROUP-NAME	LAUNCH-CONFIG	AVAILABILITY-ZONES		
MIN-SIZE	MAX-SIZE	DESIRED-CAPACITY			
AUTO-SCALING-GROUP	spotasg	spotlc-5cents	us-east-1b,us-east-1a		
1	5	3			
INSTANCE	INSTANCE-ID	AVAILABILITY-ZONE	STATE	STATUS	LAUNCH-CONFIG
INSTANCE	i-2c30d155	us-east-1a	InService	Healthy	spotlc-5cents
INSTANCE	i-fe30d187	us-east-1a	InService	Healthy	spotlc-5cents
INSTANCE	i-c630d1bf	us-east-1a	InService	Healthy	spotlc-5cents

You can see that Auto Scaling launched 3 instances in us-east-1a, as you specified, and they are all running.

3. Additionally, you can find out details about the Spot Instances launched for you by Auto Scaling, by using the `as-describe-auto-scaling-instances` command.

This is the basic syntax:

```
as-describe-auto-scaling-instances [InstanceIds [InstanceIds ...]]  
[--max-records value] [General Options]
```

Specifying `InstanceIds` is optional. If you specify it, the command will return information about the instance, if it exists. If you don't specify `InstanceIds`, the command returns information about all instances associated with your Auto Scaling account.

In this walkthrough, we are assuming that you created one launch configuration and Auto Scaling group, and you want to find out details about all your Spot Instances.

Your command should look like the following example:

```
as-describe-auto-scaling-instances --headers
```

The information you get should be similar to the following example:

INSTANCE	INSTANCE-ID	GROUP-NAME	AVAILABILITY-ZONE	STATE	STATUS
		LAUNCH-CONFIG			
INSTANCE	i-2c30d155	spotasg	us-east-1a	InService	HEALTHY
		spotlc-5cents			
INSTANCE	i-c630d1bf	spotasg	us-east-1a	InService	HEALTHY
		spotlc-5cents			
INSTANCE	i-fe30d187	spotasg	us-east-1a	InService	HEALTHY
		spotlc-5cents			

### What do you want to do next?

- [Updating the Bid Price for the Spot Instances](#) (p. 153)
- [Scheduling Spot Bid Requests](#) (p. 155)
- [Using Auto Scaling to Get Notifications for Spot Instances](#) (p. 156)
- [Advanced Tasks](#) (p. 161)

## Updating the Bid Price for the Spot Instances

Auto Scaling launch configurations cannot be changed. If you want to modify your bid price for Spot Instances, you must create a new launch configuration.

If, for example, you want to launch a set of Spot Instances that have a higher likelihood of running uninterrupted for a long time, you can use a higher bid price. To do this, you must create a new launch configuration, using the same procedure that you followed earlier in this walkthrough. (For more information, go to [Step 1: Create a Launch Configuration](#) (p. 146).)

Specify the following values:

- Launch configuration name = spotlc-7cents
- Image ID = ami-e565ba8c

#### Note

If you don't have an AMI, and you want to find a suitable one, see [Amazon Machine Images \(AMIs\)](#).

- Instance type = m1.small
- Spot price = \$0.07

Your command should look similar to the following example:

```
as-create-launch-config spotlc-7cents --image-id ami-e565ba8c --instance-type m1.small --spot-price "0.07"
```

You should get a confirmation like the following example:

```
OK-Created launch config
```

After you have created the new launch configuration successfully, create a new Auto Scaling group specifying the new launch configuration.

Your command should look similar to the following example:

```
as-create-auto-scaling-group spotasg-7cents --launch-configuration spotlc-7cents --availability-zones "us-east-1a,us-east-1b" --max-size 5 --min-size 10 --desired-capacity 3
```

You should get a confirmation like the following example:

```
OK-Created AutoScalingGroup
```

You can view the status of your Spot bid and a list of the bids that Auto Scaling placed for you by running `as-describe-scaling-activities` soon after you create your Auto Scaling group.

Your command should look similar to the following example:

```
as-describe-scaling-activities --headers
```

If not all your bids are fulfilled, you will get information that looks like the following example:

```
ACTIVITY  ACTIVITY-ID  END-TIME  GROUP-
NAME      CODE          MESSAGE
ACTIVITY  5879cc50-1e40-4539-a754-1cb084f1aecd  spotasg-
7cents  WaitingForSpotInstanceId  Placed Spot instance request: sir-93828812.
Waiting for instance(s)
ACTIVITY  777fbe1b-7a24-4aaf-b7a9-d368d0511878  spotasg-
7cents  WaitingForSpotInstanceId  Placed Spot instance request: sir-016cf812.
Waiting for instance(s)
ACTIVITY  f4b00f81-eaea-4421-80b4-a2e3a35cc782  spotasg-
7cents  WaitingForSpotInstanceId  Placed Spot instance request: sir-cf60ea12.
Waiting for instance(s)
ACTIVITY  31bcbb67-7f50-4b88-ae7e-e564a8c80a90  spotasg
WaitingForSpotInstanceId  Placed Spot instance request: sir-fc8a3014.
Waiting for instance(s)
ACTIVITY  770bbeb5-407c-404c-a826-856f65db1c57  spotasg
WaitingForSpotInstanceId  Placed Spot instance request: sir-69101014.
Waiting for instance(s)
ACTIVITY  597e4ebd-220e-42bc-8ac9-2bae4d20b8d7  2012-05-23T17:40:22Z  spotasg
Successful
ACTIVITY  eca158b4-a6f9-4ec5-a813-78d42c1738e2  2012-05-23T17:40:22Z  spotasg
Successful
ACTIVITY  1a5bd6c6-0b0a-4917-8cf0-eee1044a179f  2012-05-23T17:22:19Z  spotasg
Successful
ACTIVITY  c285bf16-d2c4-4ae8-acad-7450655facb5  2012-05-23T17:22:19Z  spotasg
Successful
ACTIVITY  127e3608-5911-4111-906e-31fb16d43f2e  2012-05-23T15:38:06Z  spotasg
Successful
ACTIVITY  bfb548ad-8bc7-4a78-a7db-3b41f73501fc  2012-05-23T15:38:07Z  spotasg
Successful
ACTIVITY  82d2b9bb-3d64-46d9-99b6-054a9ecf5ac2  2012-05-23T15:30:28Z  spotasg
Successful
ACTIVITY  95b7589b-f8ac-49bc-8c83-514bf664b4ee  2012-05-23T15:30:28Z  spotasg
Successful
ACTIVITY  57bbf77a-99d6-4d94-a6db-76b2307fb9de  2012-05-23T15:16:34Z  spotasg
Successful
ACTIVITY  cdef758e-0be2-416e-b402-0ef521861039  2012-05-23T15:16:17Z  spotasg
Successful
ACTIVITY  d7e0a3ed-7067-4583-8a87-1561b3de2aed  2012-05-23T14:51:46Z  spotasg
```

Successful			
ACTIVITY	da5471ab-482c-4680-b430-99e4173d2bd7	2012-05-23T14:52:48Z	spotasg
Successful			
ACTIVITY	78701f3d-a747-46e1-8b0f-8aff22834f46	2012-05-23T14:38:17Z	spotasg
Successful			
ACTIVITY	274d4772-3614-4f5c-8858-026b33635be3	2012-05-23T14:38:16Z	spotasg
Successful			
ACTIVITY	1024abb2-bf84-4fae-b717-a398bac91c4f	2012-05-23T14:22:39Z	spotasg
Successful			

Bids are represented by values such as `sir-93828812` and `sir-016cf812`.

When you create a new launch configuration that sets a new bid price for Spot Instances, and you have Spot Instances already running based on a different price, these instances will continue running and will only be terminated if the Spot market price goes above the bid price on which it was based.

#### What do you want to do next?

- [Scheduling Spot Bid Requests \(p. 155\)](#)
- [Using Auto Scaling to Get Notifications for Spot Instances \(p. 156\)](#)
- [Advanced Tasks \(p. 161\)](#)

## Scheduling Spot Bid Requests

You can set up Auto Scaling to launch a certain number of instances at a specific time. This capability is useful if, for example, you want to take advantage of a window of time when prices historically are lower, or you want to terminate Spot Instances at a specific time.

We will use the Auto Scaling CLI command `as-put-scheduled-update-group-action` to set up a schedule. This is the basic syntax:

```
as-put-scheduled-update-group-action ScheduledActionName --auto-scaling-group value [--desired-capacity value] [--end-time value][--max-size value][--min-size value] [--recurrence value][--start-time value][--time value][General Options]
```

In this scenario, use the following values:

- Scheduled action name: `as-spotbid-schedule`
- Auto Scaling group: `spotasg`
- Start time: `2012-05-15T19:10:00Z`
- End time: `2012-05-15T19:15:00Z`
- Desired capacity: `20`

Your command should look similar to the following example:

```
as-put-scheduled-update-group-action as-spotbid-schedule --auto-scaling-group spotasg --desired-capacity 20 --start-time 2012-05-15T19:10:00Z --end-time 2012-05-15T19:15:00Z
```

You should get a confirmation like the following example:

```
OK-Put Scheduled Update Group Action
```

To check your scheduled action, run `as-describe-scheduled-actions`.

You will get information similar to the following example:

```
UPDATE-GROUP-ACTION spotasg as-spotbid-schedule 2012-05-15T19:10:00Z 20
```

#### What do you want to do next?

- [Obtaining Information About the Instances Launched by Auto Scaling \(p. 149\)](#)
- [Updating the Bid Price for the Spot Instances \(p. 153\)](#)
- [Using Auto Scaling to Get Notifications for Spot Instances \(p. 156\)](#)
- [Advanced Tasks \(p. 161\)](#)

## Using Auto Scaling to Get Notifications for Spot Instances

You can set up Auto Scaling to send notifications to you as instances are terminated and launched. When the Spot market price goes up and your bid price falls below it, Amazon EC2 terminates the instances that were launched based on that bid price. If your Spot Instances are terminated, Auto Scaling will try to submit your bid and launch replacement instances and to ensure the capacity you specified for your Auto Scaling group. You can set up Auto Scaling to notify you when instance launch or termination events occur.

There are two ways to get notifications for Spot Instances:

- Auto Scaling
- AWS SDK Sample Code Tool

#### Spot Notifications Sample in AWS Java SDK

You can also use the AWS Java SDK to develop applications that monitor your Spot Instance usage in ways that are customized to your use case. The Spot Notifications sample application is a Java application that demonstrates techniques for monitoring Amazon EC2 instance status, Spot Instance requests, and Spot price fluctuations. The application is documented and freely available for download at [How to Track Spot Instance Activity with the Spot-Notifications Sample Application](#). You are free to adapt the application for your own purposes, or use it as a guide in developing your own application for monitoring Spot Instances. For more information, go to the [AWS SDK for Java](#).

#### Configuring Auto Scaling groups to send notifications about your Spot Instances

In this portion of the walkthrough, you learn how to set up Amazon SNS to send email notifications.

To do this, you need the following:

*An Amazon Resource Name (ARN).* You generate an ARN when you create an Amazon Simple Notification Service (Amazon SNS) topic. A topic is a communication channel to send messages and subscribe to notifications. It provides an access point for publishers and subscribers to communicate with each other. An endpoint, such as an email address, must be subscribed to the topic so the endpoint can receive messages published to the topic. To create a topic, go to [Create a Topic](#) in the *Amazon Simple Notification Service Getting Started Guide*.

*An Auto Scaling Group.* Use the Auto Scaling group that you created earlier in this walkthrough.

*A notification configuration.* You configure an Auto Scaling group to send notifications when specified events take place by calling the `as-put-notification-configuration` CLI command or the `PutNotificationConfiguration` API action. We will discuss the steps in setting up a notification configuration later in this section. For more information about the command, go to [PutNotificationConfiguration](#) in the *Auto Scaling API Reference*.

*A list of Auto Scaling notification types.* These notification types are the events that will cause the notification to be sent.

### Set Up Notifications Using Auto Scaling

1. After you've created your Amazon SNS topic and you have the ARN, you are ready to set up the notification configuration. (To create a topic, go to [Create a Topic](#) in the *Amazon Simple Notification Service Getting Started Guide*.)

To configure your Auto Scaling group to send notifications when specified events take place, use the `as-put-notification-configuration` CLI command.

The `as-put-notification-configuration` command takes the following arguments:

```
as-put-notification-configuration AutoScalingGroupName --notification-types value --topic-arn topic-ARN [General Options]
```

You need to specify the Auto Scaling group name, the ARN, and the notification types.

For this example, specify:

- Auto Scaling group name: `spotasg`

Specify the Auto Scaling group that you want to get notifications about. In this walkthrough, you want Auto Scaling to notify you when instances are launched and terminated for the `spotasg` Auto Scaling group.

- ARN: `arn:placeholder:MyTopic`

#### Note

ARNs are unique identifiers for Amazon Web Services (AWS) resources. Replace the ARN placeholder with your ARN.

- Notification types: `autoscaling:EC2_Instance_Launch`,  
`autoscaling:EC2_Instance_Terminate`

The notification types are the events you want Auto Scaling to send you e-mail about.

Open a command prompt and enter the `as-put-notification-configuration` command.

```
as-put-notification-configuration spotasg--topic-arn arn:placeholder:MyTopic
--notification-types autoscaling:EC2_INSTANCE_LAUNCH, autoscaling:EC2_IN
STANCE_TERMINATE
```

Auto Scaling returns the following:

```
OK-Put Notification Configuration
```

You now have a notification configuration that sends a notification to the endpoint subscribed in the `arn:placeholder:MyTopic` ARN whenever instances are launched or terminated.

2. Verify the notification configuration.

To verify the notification actions associated with your Auto Scaling group when specified events take place, use `as-describe-notification-configurations`.



The `as-describe-notification-configurations` command takes the following arguments:

```
as-describe-notification-configurations [--auto-scaling-groups  
value[,value...]] [--maxrecords value] [General Options]
```

If you specify the Auto Scaling group, this command returns a full list of all notification configuration for the Auto Scaling group listed. If you don't provide an Auto Scaling group name, the service returns the full details of all Auto Scaling groups. The command also returns a token if there are more pages to retrieve. To get the next page, call this action again with the returned token as the `next-token` argument. For this walkthrough, specify:

- Auto Scaling group name: `spotasg`

Open a command prompt and enter the `as-describe-notification-configurations` command.

```
as-describe-notification-configurations --auto-scaling-groups spotasg -headers
```

Auto Scaling returns the following:

```
NOTIFICATION-CONFIG  GROUP-NAME  TOPIC-ARN  NOTIFICATION-TYPE-NAME  
NOTIFICATION-CONFIG  spotasg    arn:placeholder:spotasg  autoscaling:EC2_IN  
STANCE_LAUNCH  
NOTIFICATION-CONFIG  spotasg    arn:placeholder:spotasg  autoscaling:EC2_IN  
STANCE_TERMINATE
```

You have confirmed that you have a notification configuration set up for the `spotasg` Auto Scaling group.

When Auto Scaling launches instances to reach or maintain desired capacity, as specified in your Auto Scaling group, SNS sends a notification to your email address with information about the instances. You can check your email Inbox for this information, then run `as-describe-auto-scaling-group` to get information about current instances in the Auto Scaling group and confirm that the instances listed in your email actually exist.

### What do you want to do next?

- [Obtaining Information About the Instances Launched by Auto Scaling \(p. 149\)](#)
- [Updating the Bid Price for the Spot Instances \(p. 153\)](#)
- [Scheduling Spot Bid Requests \(p. 155\)](#)
- [Advanced Tasks \(p. 161\)](#)

## Using CloudFormation Templates to Launch Spot Instances

You can use AWS CloudFormation templates to launch Spot Instances and the AWS resources you need for an application. Templates are text files describing a collection, called *stack*, of AWS resources that you want to deploy as a group. For more information about AWS CloudFormation templates, see [Learn Template Basics](#). For more information about using AWS CloudFormation, see the [AWS CloudFormation User Guide](#).

You can write your own template from scratch, or you can use one of the example templates from the [AWS CloudFormation Sample Template Library](#). The following templates in AWS CloudFormation utilize Spot Instances:

- Asynchronous Queue-Based Processing

This template creates an auto-scaled worker that monitors work (messages) in a queue. The application is auto-scaled based on the amount of work in the queue. When there is work, Auto Scaling scales up; when there is no work, Auto Scaling scales down. Each message contains a command/script to run, an input file location, and an output location for the results.

To download this template, go to [Asynchronous Queue-Based Processing](#).

- Bees with Machine Guns

This template creates a load balancer, a controller, and the instances behind the load balancer; fires up and triggers an attack; and stores logs on Amazon Simple Storage Service (Amazon S3) and then shuts down (if enabled). For this template, instances use the Amazon Linux AMI, and the setup requires that you already have an AWS Identity and Access Management (IAM) SSL certification. You can modify this template to remove the SSL dependencies.

**Note**

To launch this template, you need a new SSH private key that you create specifically for this task. You'll have to log in to the instance created by the template and provide this private key information.

To download this template, go to [Bees with Machine Guns](#).

- Grid Computing for Spot Instances

This template uses Star Cluster to launch and bootstrap a cluster of Amazon EC2 instances for high performance computational tasks using Spot pricing. You only need to provide the number of instances, instance size, and Spot price that you want to use.

To download this template, go to [Grid Computing for Spot Instances](#).

You can create stacks from these templates by using the AWS Management Console, the AWS CloudFormation command line interface tools, or the AWS CloudFormation API.

Before you use these templates, you must:

1. Have an AWS account and sign up for AWS CloudFormation.
2. Decide on the template to use.
3. Enter the parameters required for the stack.
4. Create the stack.

To get started with AWS CloudFormation, [Getting Started with AWS CloudFormation](#).

For more information about using AWS CloudFormation, see the [AWS CloudFormation User Guide](#).

## Launching Amazon Elastic MapReduce Job Flows with Spot Instances

You can launch an Amazon Elastic MapReduce job flow in Spot Instances. Amazon Elastic MapReduce is a data analysis tool that simplifies the set up and management of a computer cluster, the source data, and the computational tools that help you implement sophisticated data processing jobs quickly. For more information, see [What is Amazon EMR?](#)

In addition, you can use Spot Instances with Amazon EMR clusters and AWS Data Pipeline. For information, see [Amazon EC2 Spot Instances with Amazon EMR Clusters and AWS Data Pipeline](#).

## Quick Look: Using Spot Instances with Amazon Elastic MapReduce Video

The following video describes how Spot Instances work in Amazon EMR and walks you through the process of launching a job flow on Spot Instances from the AWS Management Console: [Using Spot Instances with Amazon ElasticMapReduce](#)

## Launching Spot Instances in Amazon Virtual Private Cloud

Amazon Virtual Private Cloud (Amazon VPC) lets you define a virtual networking environment in a private, isolated section of the Amazon Web Services (AWS) cloud. You have complete control of this virtual network and you can use advanced security features and network access control at the instance level and subnet level. For more information about Amazon VPC, see the [Amazon Virtual Private Cloud User Guide](#).

If you want to take advantage of the features of Amazon VPC when you use Spot Instances, specify in your Spot request that your instances are to be launched in Amazon VPC. To ensure that your Spot Instances are launched in Amazon VPC:

- Determine if your AWS account supports a default VPC. For more information, see [Detecting Your Supported Platforms and Whether You Have a Default VPC](#).
- If your AWS account supports a default VPC, and
  - You want to launch your Spot Instances in that VPC

You don't have to perform any additional steps when you request Spot Instances. The **Request Spot Instances** wizard selects the subnet ID of your default VPC in which to launch the instance.

- You want to use a different VPC for the Spot Instances

Create a new Amazon VPC and specify the subnet ID of the VPC that you just created.

### Note

Support for default VPC is available in select regions, and will be available in all regions soon. For more information, see [Default VPC Basics](#).

- If your AWS account does not support a default VPC, create an Amazon VPC and specify the subnet of the VPC that you just created.

This topic discusses what you need to know if you want your Spot Instances to launch in an Amazon VPC and your account does not support default VPCs, or your account supports default VPCs but you want your Spot Instances to launch in a VPC that is not the default.

For more information about setting up an Amazon VPC, see [Getting Started with Amazon VPC](#) in the *Amazon Virtual Private Cloud Getting Started Guide*.

## Quick Look: Launching Spot Instances in Amazon VPC Video

The following video shows how to launch your first Spot Instance in the Amazon Virtual Private Cloud (Amazon VPC) using the AWS Management Console. This video includes instructions for creating your Amazon VPC subnet, placing a bid, determining when the instance is fulfilled, and canceling the instance. [Launching Spot Instances in Amazon VPC Video](#)

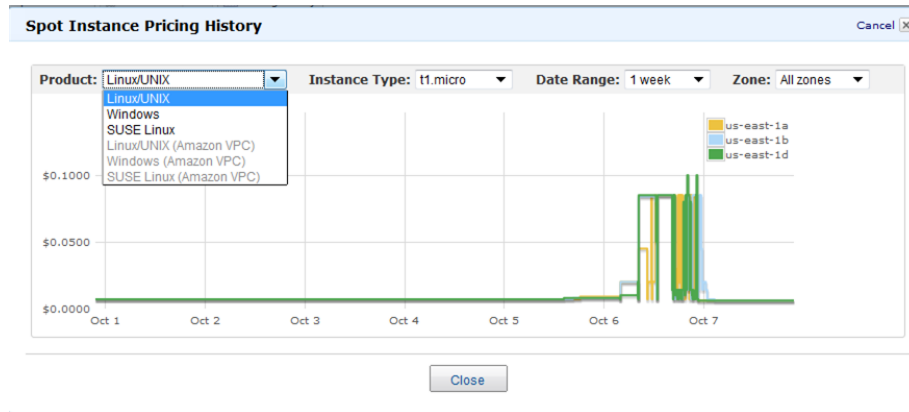
The process for making a Spot Instance request that launches in Amazon VPC is the same as the process you follow when you make a Spot Instance request in a non-VPC portion of Amazon EC2. The main differences are that you:

- Base your Spot price bid on Spot price history of Spot Instances in VPCs.

When you use the `DescribeSpotPriceHistory` action or the `ec2-describe-spot-price-history` command, add *Amazon VPC* to the `product-description` filter. For example:

```
PROMPT> ec2-describe-spot-price-history -s '2011-09-08T00:00:00Z' -t m1.xlarge  
-d "Linux/UNIX (Amazon VPC)"
```

Using the AWS Management Console, check the **Spot Instance Pricing History** page to see the Spot pricing history for Amazon EC2 instances running in both Amazon EC2 and Amazon VPC.



- Specify the VPC subnet in which you want to launch your Spot Instance.

When you use the `RequestSpotInstances` action or the `ec2-request-spot-instances` command, specify the ID of the Amazon VPC subnet in which you want to launch the Spot Instance.

```
PROMPT> ec2-request-spot-instances ami-8e1fece7 -t m1.xlarge -p '$0.01' -n 5  
-r 'one-time' -s 'subnet-baab943d3'
```

When you launch the **Request Instances Wizard** from the Spot Instance page of the AWS Management Console, select a subnet after specifying that you're launching the Spot Instance into a VPC.

For more information about using Amazon VPC, see the [Amazon Virtual Private Cloud User Guide](#).

## Advanced Tasks

Now that you have created Spot Instance requests and worked with Spot Instances, this section discusses some advanced tasks.

- [Subscribe to Your Spot Instance Data Feed \(p. 161\)](#)
- [Programming Spot with AWS Java SDK \(p. 164\)](#)
- [Starting Clusters on Spot Instances \(p. 191\)](#)

If you think we should add other advanced tasks to this section, let us know at [spot-instance-feedback@amazon.com](mailto:spot-instance-feedback@amazon.com).

## Subscribe to Your Spot Instance Data Feed

If you want to monitor your Spot Instance usage, you can subscribe to your Spot Instance data feed, which stores usage logs in Amazon Simple Storage Service (Amazon S3).

This section describes the data feed content and how to create the data feed for Spot Instances. You can create one data feed per account.

## Spot Instance Data Feed Overview

To help you understand the charges for your Spot Instances, Amazon EC2 provides access to a data feed that details your Spot Instance usage and pricing. This data feed is sent to the Amazon S3 bucket of your choice.

To have a data feed delivered to an Amazon S3 bucket, you need to create a Spot Instances data feed subscription using the Amazon EC2 API. When you create this subscription, you can specify an Amazon S3 bucket to deliver the data feed files to, and a filename prefix to use to avoid collisions.

### Data Feed Filename and Format

The Spot Instance data feed filename uses the following format (with the date and hour in UTC):

```
{Bucket}.s3.amazonaws.com/{Optional Prefix}/{AWS Account ID}.{YYYY}-{MM}-{DD}-  
{HH}.{n}.{Unique ID}.gz
```

For example, if your bucket name is `myawsbucket`, and you name your prefix `myprefix`, your filenames look similar to this:

```
myawsbucket.s3.amazonaws.com/myprefix/111122223333.2010-03-17-20.001.pwBdGTJG.gz
```

Data feed files arrive in your bucket typically once an hour and each hour of usage is typically covered in a single data file. These files are compressed (gzip) before delivery into your bucket. We can write multiple files for a given hour of usage where files are very large (for example, when file contents for the hour exceed 50 MB before compression).

#### Note

If you don't have a Spot Instance running during a certain hour, you won't receive a data feed file for that hour.

The Spot Instance data feed files are tab-delimited. Each line in the data file corresponds to one instance-hour. Each line contains the fields listed in the following table.

Field	Description
Timestamp	The timestamp used to determine the price charged for this instance-hour.
UsageType	Indicates the type of usage and instance type being charged for. For m1.small Spot Instances, this field is set to "SpotUsage." For all other instance types, this field is set to "SpotUsage:{instance-type}," for example, "SpotUsage:c1.medium."
Operation	Indicates the product being charged for. For Linux/UNIX Spot Instances, this field is set to "RunInstances." For Microsoft Windows, this field is set to "RunInstances:0002." Spot usage is grouped according to Availability Zone.
InstanceID	The instance ID for the Spot Instance that generated this instance-hour.
MyBidID	The Spot Instance request ID for the request that generated this instance-hour.
MyMaxPrice	The maximum price specified for this Spot Instance request.
MarketPrice	The Spot price at the time specified in the Timestamp field.
Charge	The price charged for this instance-hour.
Version	The version included in the data feed filename for this record.

## Preparing Amazon S3 for Data Feeds

When you subscribe to data feeds, you tell Amazon EC2 which bucket you want to store the data feed file in. Before you subscribe to the data feed, consider the following when choosing your S3 bucket:

- You must have Amazon S3 FULL\_CONTROL permission on the bucket you provide.

If you're the bucket owner, you have this permission by default. If you're not the bucket owner, the bucket owner must grant your AWS account FULL\_CONTROL permission.

- When you create your data feed subscription, Amazon EC2 updates the designated bucket's ACL to allow read and write permissions for the AWS data feeds account.
- Each data feed file has its own ACL (separate from the bucket's ACL).

The bucket owner has FULL\_CONTROL permission for the data files. The data feed account has read and write permission.

- Removing the permissions for the data feed account does not disable the data feed.

If you remove those permissions but don't disable the data feed (which you do with the control API), we reinstate those permissions the next time the data feeds account needs to write a data file to your bucket.

- If you delete your data feed subscription, Amazon EC2 doesn't remove the read/write permissions for the data feed account on either the bucket or the data files.

You must perform remove the read/write permissions yourself.

## Subscribe to Your Spot Instance Data Feed

### Amazon EC2 Command Line Interface (CLI) Tools

#### To subscribe to your Spot Instance data feed

- Enter the following command:

```
PROMPT> ec2-create-spot-datafeed-subscription --bucket myawsbucket [--prefix  
prefix ]
```

Amazon EC2 returns output similar to the following:

```
SPOTDATAFEEDSUBSCRIPTION      111122223333      myawsbucket      prefix  
Active
```

### API

#### To subscribe to your Spot Instance data feed

- Construct the following Query request.

```
https://ec2.amazonaws.com/  
?Action=CreateSpotDatafeedSubscription  
&Bucket=myawsbucket  
&Prefix=my-spot-subscription  
&...auth parameters...
```

Following is an example response.

```
<CreateSpotDatafeedSubscriptionResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <requestId>59dbff89-35bd-4eac-99ed-be587EXAMPLE</requestId>
  <spotDatafeedSubscription>
    <ownerId>999988887777</ownerId>
    <bucket>myawsbucket</bucket>
    <prefix>my-spot-subscription</prefix>
    <state>Active</state>
    <fault>/</fault>
  </spotDatafeedSubscription>
</CreateSpotDatafeedSubscriptionResponse>
```

## Delete a Spot Instance Data Feed

### Command Line Tools

#### To delete a Spot Instance data feed

- To delete a data feed, enter the following command:

```
PROMPT> ec2-delete-spot-datafeed-subscription
```

If the request is successful, the output is empty.

### API

#### To delete a Spot Instance data feed

- Construct the following Query request.

```
https://ec2.amazonaws.com/
?Action=DeleteSpotDatafeedSubscription
&...auth parameters...
```

Following is an example response. It confirms that the subscription was deleted.

```
<DeleteSpotDatafeedSubscriptionResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <requestId>59dbff89-35bd-4eac-99ed-be587EXAMPLE</requestId>
  <return>>true</return>
</DeleteSpotDatafeedSubscriptionResponse>
```

## Programming Spot with AWS Java SDK

This section will hold the two SDK tutorials

You can use the AWS Java SDK to program Spot Instances:

- [Tutorial: Amazon EC2 Spot Instances \(p. 166\)](#)

- [Tutorial: Advanced Amazon EC2 Spot Request Management \(p. 175\)](#)



## Tutorial: Amazon EC2 Spot Instances

### Overview

Spot Instances allow you to bid on unused Amazon Elastic Compute Cloud (Amazon EC2) capacity and run the acquired instances for as long as your bid exceeds the current *Spot Price*. Amazon EC2 changes the Spot Price periodically based on supply and demand, and customers whose bids meet or exceed it gain access to the available Spot Instances. Like On-Demand Instances and Reserved Instances, Spot Instances provide you another option for obtaining more compute capacity.

Spot Instances can significantly lower your Amazon EC2 costs for batch processing, scientific research, image processing, video encoding, data and web crawling, financial analysis, and testing. Additionally, Spot Instances give you access to large amounts of additional capacity in situations where the need for that capacity is not urgent.

To use Spot Instances, place a Spot Instance request specifying the maximum price you are willing to pay per instance hour; this is your bid. If your bid exceeds the current Spot Price, your request is fulfilled and your instances will run until either you choose to terminate them or the Spot Price increases above your bid (whichever is sooner).

It's important to note two points:

- You will often pay less per hour than your bid. Amazon EC2 adjusts the Spot Price periodically as requests come in and available supply changes. Everyone pays the same Spot Price for that period regardless of whether their bid was higher. Therefore, you might pay less than your bid, but you will never pay more than your bid.
- If you're running Spot Instances and your bid no longer meets or exceeds the current Spot Price, your instances will be terminated. This means that you will want to make sure that your workloads and applications are flexible enough to take advantage of this opportunistic capacity.

Spot Instances perform exactly like other Amazon EC2 instances while running, and like other Amazon EC2 instances, Spot Instances can be terminated when you no longer need them. If you terminate your instance, you pay for any partial hour used (as you would for On-Demand or Reserved Instances). However, if the Spot Price goes above your bid and your instance is terminated by Amazon EC2, you will not be charged for any partial hour of usage.

This tutorial provides a quick overview of how to use the Java programming language to do the following.

- Submit a Spot Request
- Determine when the Spot Request becomes fulfilled
- Cancel the Spot Request
- Terminate associated instances

### Prerequisites

To use this tutorial you need to be signed up for Amazon Web Services (AWS). If you have not yet signed up for AWS, go to the [Amazon Web Services website](#), and click **Create an AWS Account** in the upper right corner of the page. In addition, you also need to install the [AWS Java SDK](#).

If you are using the Eclipse development environment, we recommend that you install the [AWS Toolkit for Eclipse](#). Note that the AWS Toolkit for Eclipse includes the latest version of the AWS SDK for Java.

### Step 1: Setting Up Your Credentials

To begin using this code sample, you need to populate your credentials in the `AwsCredentials.properties` file. Specifically, you need to populate your secret key and access key.

Copy and paste your Access Key and Secret Access Key into the `AwsCredentials.properties` file.

Now that you have configured your settings, you can get started using the code in the example.

## Step 2: Setting Up a Security Group

A *security group* acts as a firewall that controls the traffic allowed in and out of a group of instances. By default, an instance is started without any security group, which means that all incoming IP traffic, on any TCP port will be denied. So, before submitting our Spot Request, we will set up a security group that allows the necessary network traffic. For the purposes of this tutorial, we will create a new security group called "GettingStarted" that allows Secure Shell (SSH) traffic from the IP address where you are running your application from. To set up a new security group, you need to include or run the following code sample that sets up the security group programmatically.

After we create an `AmazonEC2` client object, we create a `CreateSecurityGroupRequest` object with the name, "GettingStarted" and a description for the security group. Then we call the `ec2.createSecurityGroup` API to create the group.

To enable access to the group, we create an `ipPermission` object with the IP address range set to the CIDR representation of the subnet for the local computer; the "/10" suffix on the IP address indicates the subnet for the specified IP address. We also configure the `ipPermission` object with the TCP protocol and port 22 (SSH). The final step is to call `ec2.authorizeSecurityGroupIngress` with the name of our security group and the `ipPermission` object.

```
1
  // Retrieves the credentials from an AwsCredentials.properties file.
  AwsCredentials credentials = null;
  try {
5     credentials = new PropertiesCredentials(
        GettingStartedApp.class.getResourceAsStream("AwsCredentials.properties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCredentials.properties.");
        System.out.println(e1.getMessage());
10    System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Create a new security group.
    try {
        CreateSecurityGroupRequest securityGroupRequest = new CreateSecurityGroupRequest("GettingStartedGroup", "Getting Started Security Group");
        ec2.createSecurityGroup(securityGroupRequest);
20 } catch (AmazonServiceException ase) {
    // Likely this means that the group is already created, so ignore.
    System.out.println(ase.getMessage());
    }

25 String ipAddr = "0.0.0.0/0";

    // Get the IP of the current host, so that we can limit the Security
    // Group by default to the ip range associated with your subnet.
    try {
30    InetAddress addr = InetAddress.getLocalHost();
```

```
        // Get IP Address
        ipAddr = addr.getHostAddress()+"/10";
    } catch (UnknownHostException e) {
35 }

    // Create a range that you would like to populate.
    ArrayList<String> ipRanges = new ArrayList<String>();
    ipRanges.add(ipAddr);
40

    // Open up port 22 for TCP traffic to the associated IP
    // from above (e.g. ssh traffic).
    ArrayList<IpPermission> ipPermissions = new ArrayList<IpPermission> ();
    IpPermission ipPermission = new IpPermission();
45 ipPermission.setIpProtocol("tcp");
    ipPermission.setFromPort(new Integer(22));
    ipPermission.setToPort(new Integer(22));
    ipPermission.setIpRanges(ipRanges);
    ipPermissions.add(ipPermission);
50

    try {
        // Authorize the ports to the used.
        AuthorizeSecurityGroupIngressRequest ingressRequest =
            new AuthorizeSecurityGroupIngressRequest("GettingStartedGroup", ip
Permissions);
55     ec2.authorizeSecurityGroupIngress(ingressRequest);
    } catch (AmazonServiceException ase) {
        // Ignore because this likely means the zone has
        // already been authorized.
        System.out.println(ase.getMessage());
60 }
}
```

You can view this entire code sample in the `CreateSecurityGroupApp.java` code sample. Note you only need to run this application once to create a new security group.

You can also create the security group using the AWS Toolkit for Eclipse. Go to the [toolkit documentation](#) for more information.

### Step 3: Submitting Your Spot Request

To submit a Spot request, you first need to determine the instance type, Amazon Machine Image (AMI), and maximum bid price you want to use. You must also include the security group we configured previously, so that you can log into the instance if desired.

There are several instance types to choose from; see [Instance Types \(p. 94\)](#) for a complete list. For this tutorial, we will use `t1.micro`, the cheapest instance type available. Next, we will determine the type of AMI we would like to use. We'll use `ami-8c1fece5`, the most up-to-date Amazon Linux AMI available when we wrote this tutorial. The latest AMI may change over time, but you can always determine the latest version AMI by:

1. Logging into the AWS Management Console, clicking the EC2 tab, and, from the EC2 Console Dashboard, attempting to launch an instance.
2. In the window that displays AMIs, just use the AMI ID as shown in the wizard. Alternatively, you can use the `DescribeImages` API, but leveraging that command is outside the scope of this tutorial.

There are many ways to approach bidding for Spot instances; to get a broad overview of the various approaches you should view the [Bidding for Spot Instances](#) video. However, to get started, we'll describe three common strategies: bid to ensure cost is less than on-demand pricing; bid based on the value of the resulting computation; bid so as to acquire computing capacity as quickly as possible.

- **Reduce Cost below On-Demand** You have a batch processing job that will take a number of hours or days to run. However, you are flexible with respect to when it starts and when it completes. You want to see if you can complete it for less cost than with On-Demand Instances. You examine the Spot Price history for instance types using either the AWS Management Console or the Amazon EC2 API. For more information, see [Viewing Spot Instance Pricing History \(p. 118\)](#). After you've analyzed the price history for your desired instance type in a given Availability Zone, you have two alternative approaches for your bid:
  - You could bid at the upper end of the range of Spot Prices (which are still below the On-Demand price), anticipating that your one-time Spot request would most likely be fulfilled and run for enough consecutive compute time to complete the job.
  - Or, you could bid at the lower end of the price range, and plan to combine many instances launched over time through a persistent request. The instances would run long enough--in aggregate--to complete the job at an even lower total cost. (We will explain how to automate this task later in this tutorial.)
- **Pay No More than the Value of the Result** You have a data processing job to run. You understand the value of the job's results well enough to know how much they are worth in terms of computing costs. After you've analyzed the Spot Price history for your instance type, you choose a bid price at which the cost of the computing time is no more than the value of the job's results. You create a persistent bid and allow it to run intermittently as the Spot Price fluctuates at or below your bid.
- **Acquire Computing Capacity Quickly** You have an unanticipated, short-term need for additional capacity that is not available through On-Demand Instances. After you've analyzed the Spot Price history for your instance type, you bid above the highest historical price to provide a high likelihood that your request will be fulfilled quickly and continue computing until it completes.

After you choose your bid price, you are ready to request a Spot Instance. For the purposes of this tutorial, we will bid the On-Demand price (\$0.03) to maximize the chances that the bid will be fulfilled. You can determine the types of available instances and the On-Demand prices for instances by going to Amazon EC2 Pricing page. To request a Spot Instance, you simply need to build your request with the parameters you chose earlier. We start by creating a `RequestSpotInstanceRequest` object. The request object requires the number of instances you want to start and the bid price. Additionally, you need to set the `LaunchSpecification` for the request, which includes the instance type, AMI ID, and security group you want to use. Once the request is populated, you call the `requestSpotInstances` method on the `AmazonEC2Client` object. The following example shows how to request a Spot Instance.

```
1
  // Retrieves the credentials from a AWScredentials.properties file.
  AWScredentials credentials = null;
  try {
5     credentials = new PropertiesCredentials(
        GettingStartedApp.class.getResourceAsStream("AwsCredentials.properties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCredentials.properties.");
        System.out.println(e1.getMessage());
10    System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
```

```
15 // Initializes a Spot Instance Request
    RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

    // Request 1 x t1.micro instance with a bid price of $0.03.
20 requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(1));

    // Setup the specifications of the launch. This includes the
    // instance type (e.g. t1.micro) and the latest Amazon Linux
25 // AMI id available. Note, you should always use the latest
    // Amazon Linux AMI id or another of your choosing.
    LaunchSpecification launchSpecification = new LaunchSpecification();
    launchSpecification.setImageId("ami-8c1fece5");
    launchSpecification.setInstanceType("t1.micro");
30

    // Add the security group to the request.
    ArrayList<String> securityGroups = new ArrayList<String>();
    securityGroups.add("GettingStartedGroup");
    launchSpecification.setSecurityGroups(securityGroups);
35

    // Add the launch specifications to the request.
    requestRequest.setLaunchSpecification(launchSpecification);

    // Call the RequestSpotInstance API.
40 RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

Running this code will launch a new Spot Instance Request. There are other options you can use to configure your Spot Requests. To learn more, see [Tutorial: Advanced Amazon EC2 Spot Request Management \(p. 175\)](#) or the [RequestSpotInstances](#) API in the Java SDK.

**Note**

You will be charged for any Spot Instances that are actually launched, so make sure that you cancel any requests and terminate any instances you launch to reduce any associated fees.

**Step 4: Determining the State of Your Spot Request**

Next, we want to create code to wait until the Spot request reaches the "active" state before proceeding to the last step. To determine the state of our Spot request, we poll the [describeSpotInstanceRequests](#) method for the state of the Spot request ID we want to monitor.

The request ID created in Step 2 is embedded in the response to our `requestSpotInstances` request. The following example code shows how to gather request IDs from the `requestSpotInstances` response and use them to populate an `ArrayList`.

```
1 // Call the RequestSpotInstance API.
    RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
    List<SpotInstanceRequest> requestResponses = requestResult.getSpotInstance
Requests();
5 // Setup an arraylist to collect all of the request ids we want to
    // watch hit the running state.
```

```
ArrayList<String> spotInstanceRequestIds = new ArrayList<String>();

10 // Add all of the request ids to the hashset, so we can determine when they
    // hit the
    // active state.
    for (SpotInstanceRequest requestResponse : requestResponses) {
        System.out.println("Created Spot Request: "+requestResponse.getSpotIn
            stanceRequestId());
        spotInstanceRequestIds.add(requestResponse.getSpotInstanceRequestId());
15 }
```

To monitor your request ID, call the `describeSpotInstanceRequests` method to determine the state of the request. Then loop until the request is not in the "open" state. Note that we monitor for a state of not "open", rather a state of, say, "active", because the request can go straight to "closed" if there is a problem with your request arguments. The following code example provides the details of how to accomplish this task.

```
1
    // Create a variable that will track whether there are any
    // requests still in the open state.
    boolean anyOpen;
5
    do {
        // Create the describeRequest object with all of the request ids
        // to monitor (e.g. that we started).
        DescribeSpotInstanceRequestsRequest describeRequest = new DescribeSpot
            InstanceRequestsRequest();
10        describeRequest.setSpotInstanceRequestIds(spotInstanceRequestIds);

        // Initialize the anyOpen variable to false - which assumes there
        // are no requests open unless we find one that is still open.
        anyOpen=false;
15
        try {
            // Retrieve all of the requests we want to monitor.
            DescribeSpotInstanceRequestsResult describeResult = ec2.describeS
                potInstanceRequests(describeRequest);
            List<SpotInstanceRequest> describeResponses = describeResult.get
                SpotInstanceRequests();
20
                // Look through each request and determine if they are all in
                // the active state.
                for (SpotInstanceRequest describeResponse : describeResponses) {
                    // If the state is open, it hasn't changed since we attempted
25                    // to request it. There is the potential for it to transition
                    // almost immediately to closed or canceled so we compare
                    // against open instead of active.
                    if (describeResponse.getState().equals("open")) {
                        anyOpen = true;
30                    break;
                    }
                }
        } catch (AmazonServiceException e) {
            // If we have an exception, ensure we don't break out of
35            // the loop. This prevents the scenario where there was
            // blip on the wire.
```

```
        anyOpen = true;
    }
40    try {
        // Sleep for 60 seconds.
        Thread.sleep(60*1000);
    } catch (Exception e) {
        // Do nothing because it woke up early.
45    }
    } while (anyOpen);
```

After running this code, your Spot Instance Request will have completed or will have failed with an error that will be output to the screen. In either case, we can proceed to the next step to clean up any active requests and terminate any running instances.

### Step 5: Cleaning Up Your Spot Requests and Instances

Lastly, we need to clean up our requests and instances. It is important to both cancel any outstanding requests *and* terminate any instances. Just canceling your requests will not terminate your instances, which means that you will continue to pay for them. If you terminate your instances, your Spot requests may be canceled, but there are some scenarios—such as if you use persistent bids—where terminating your instances is not sufficient to stop your request from being re-fulfilled. Therefore, it is a best practice to both cancel any active bids and terminate any running instances.

The following code demonstrates how to cancel your requests.

```
1    try {
        // Cancel requests.
        CancelSpotInstanceRequestsRequest cancelRequest = new CancelSpotInstanceRequestsRequest(spotInstanceRequestIds);
5    ec2.cancelSpotInstanceRequests(cancelRequest);
    } catch (AmazonServiceException e) {
        // Write out any exceptions that may have occurred.
        System.out.println("Error canceling instances");
        System.out.println("Caught Exception: " + e.getMessage());
10    System.out.println("Response Status Code: " + e.getStatusCode());
        System.out.println("Error Code: " + e.getErrorCode());
        System.out.println("Request ID: " + e.getRequestId());
    }
```

To terminate any outstanding instances, you will need the instance ID associated with the request that started them. The following code example takes our original code for monitoring the instances and adds an `ArrayList` in which we store the instance ID associated with the `describeInstance` response.

```
1    // Create a variable that will track whether there are any requests
    // still in the open state.
    boolean anyOpen;
5
    // Initialize variables.
    ArrayList<String> instanceIds = new ArrayList<String>();

    do {
```

```
10    // Create the describeRequest with all of the request ids to
    // monitor (e.g. that we started).
    DescribeSpotInstanceRequestsRequest describeRequest = new DescribeSpot
InstanceRequestsRequest();
    describeRequest.setSpotInstanceRequestIds(spotInstanceRequestIds);

15    // Initialize the anyOpen variable to false, which assumes there
    // are no requests open unless we find one that is still open.
    anyOpen = false;

    try {
20        // Retrieve all of the requests we want to monitor.
        DescribeSpotInstanceRequestsResult describeResult = ec2.describeS
potInstanceRequests(describeRequest);
        List<SpotInstanceRequest> describeResponses = describeResult.get
SpotInstanceRequests();

        // Look through each request and determine if they are all
25        // in the active state.
        for (SpotInstanceRequest describeResponse : describeResponses) {
            // If the state is open, it hasn't changed since we
            // attempted to request it. There is the potential for
            // it to transition almost immediately to closed or
30            // canceled so we compare against open instead of active.
            if (describeResponse.getState().equals("open")) {
                anyOpen = true;
                break;
            }
35
            // Add the instance id to the list we will
            // eventually terminate.
            instanceIds.add(describeResponse.getInstanceId());
        }
40    } catch (AmazonServiceException e) {
        // If we have an exception, ensure we don't break out
        // of the loop. This prevents the scenario where there
        // was blip on the wire.
        anyOpen = true;
45    }

    try {
        // Sleep for 60 seconds.
        Thread.sleep(60*1000);
50    } catch (Exception e) {
        // Do nothing because it woke up early.
    }
} while (anyOpen);
```

Using the instance IDs, stored in the `ArrayList`, terminate any running instances using the following code snippet.

```
1    try {
        // Terminate instances.
        TerminateInstancesRequest terminateRequest = new TerminateInstances
Request(instanceIds);
```



```
5     ec2.terminateInstances(terminateRequest);
    } catch (AmazonServiceException e) {
        // Write out any exceptions that may have occurred.
        System.out.println("Error terminating instances");
        System.out.println("Caught Exception: " + e.getMessage());
10    System.out.println("Reponse Status Code: " + e.getStatusCode());
        System.out.println("Error Code: " + e.getErrorCode());
        System.out.println("Request ID: " + e.getRequestId());
    }
```

### Bringing It All Together

To bring this all together, we provide a more object-oriented approach that combines the preceding steps we showed: initializing the EC2 Client, submitting the Spot Request, determining when the Spot Requests are no longer in the open state, and cleaning up any lingering Spot request and associated instances. We create a class called `Requests` that performs these actions.

We also create a `GettingStartedApp` class, which has a main method where we perform the high level function calls. Specifically, we initialize the `Requests` object described previously. We submit the Spot Instance request. Then we wait for the Spot request to reach the "Active" state. Finally, we clean up the requests and instances.

The complete source code is available for download at [GitHub](#).

Congratulations! You have just completed the getting started tutorial for developing Spot Instance software with the AWS Java SDK.

### Next Steps

We recommend that you take the Java Developers: [Tutorial: Advanced Amazon EC2 Spot Request Management \(p. 175\)](#).

## Tutorial: Advanced Amazon EC2 Spot Request Management

### Overview

Spot Instances allow you to bid on unused Amazon Elastic Compute Cloud (Amazon EC2) capacity and run those instances for as long as your bid exceeds the current Spot Price. Amazon EC2 changes the Spot Price periodically based on supply and demand. Customers whose bids meet or exceed the Spot Price gain access to the available Spot Instances. Like On-Demand Instances and Reserved Instances, Spot Instances provide you an additional option for obtaining more compute capacity.

Spot Instances can significantly lower your Amazon EC2 costs for batch processing, scientific research, image processing, video encoding, data and web crawling, financial analysis, and testing. Additionally, Spot Instances can provide access to large amounts of additional compute capacity when your need for the capacity is not urgent.

This tutorial provides a quick overview of some advanced Spot Request features, such as detailed options to create Spot requests, alternative methods for launching Spot Instances, and methods to manage your instances. This tutorial is not meant to be a complete list of all advanced topics associated with Spot Instances. Instead, it gives you a quick reference of code samples for some of the commonly used methods for managing Spot Requests and Spot Instances.

### Prerequisites

To use this tutorial you need to be signed up for Amazon Web Services (AWS). If you have not yet signed up for AWS, go to the [Amazon Web Services website](#), and click **Create an AWS Account** in the upper right corner of the page. In addition, you also need to install the [AWS Java SDK](#).

If you are using the Eclipse development environment, we recommend that you install the [AWS Toolkit for Eclipse](#). Note that the AWS Toolkit for Eclipse includes the latest version of the AWS SDK for Java.

### Step 1: Setting Up Your Credentials

To begin using this code sample, you need to populate your credentials in the `AwsCredentials.properties` file. Specifically, you need to populate your `secretKey` and `accessKey`.

Copy and paste your access key ID and secret access key into the `AwsCredentials.properties` file.

### Step 2: Setting Up a Security Group

Additionally, you need to configure your *security group*. A security group acts as a firewall that controls the traffic allowed in and out of a group of instances. By default, an instance is started without any security group, which means that all incoming IP traffic, on any TCP port will be denied. So, before submitting our Spot Request, we will set up a security group that allows the necessary network traffic. For the purposes of this tutorial, we will create a new security group called "GettingStarted" that allows Secure Shell (SSH) traffic from the IP address where you are running your application from. To set up a new security group, you need to include or run the following code sample that sets up the security group programmatically.

After we create an `AmazonEC2` client object, we create a `CreateSecurityGroupRequest` object with the name, "GettingStarted" and a description for the security group. Then we call the `ec2.createSecurityGroup` API to create the group.

To enable access to the group, we create an `ipPermission` object with the IP address range set to the CIDR representation of the subnet for the local computer; the "/10" suffix on the IP address indicates the subnet for the specified IP address. We also configure the `ipPermission` object with the TCP protocol and port 22 (SSH). The final step is to call `ec2.authorizeSecurityGroupIngress` with the name of our security group and the `ipPermission` object.

(The following code is the same as what we used in the first tutorial.)

```
1
  // Retrieves the credentials from an AWSCredentials.properties file.
  AWSCredentials credentials = null;
  try {
5     credentials = new PropertiesCredentials(
        GettingStartedApp.class.getResourceAsStream("AwsCredentials.properties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCredentials.properties.");
        System.out.println(e1.getMessage());
10    System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Create a new security group.
    try {
        CreateSecurityGroupRequest securityGroupRequest =
            new CreateSecurityGroupRequest("GettingStartedGroup", "Getting
Started Security Group");
20    ec2.createSecurityGroup(securityGroupRequest);
    } catch (AmazonServiceException ase) {
        // Likely this means that the group is already created, so ignore.
        System.out.println(ase.getMessage());
    }
25
    String ipAddr = "0.0.0.0/0";

    // Get the IP of the current host, so that we can limit the Security Group
    // by default to the ip range associated with your subnet.
30 try {
        InetAddress addr = InetAddress.getLocalHost();

        // Get IP Address
        ipAddr = addr.getHostAddress()+"/10";
35 } catch (UnknownHostException e) {
    }

    // Create a range that you would like to populate.
    ArrayList<String> ipRanges = new ArrayList<String>();
40 ipRanges.add(ipAddr);

    // Open up port 22 for TCP traffic to the associated IP from
    // above (e.g. ssh traffic).
    ArrayList<IpPermission> ipPermissions = new ArrayList<IpPermission> ();
45 IpPermission ipPermission = new IpPermission();
    ipPermission.setIpProtocol("tcp");
    ipPermission.setFromPort(new Integer(22));
    ipPermission.setToPort(new Integer(22));
    ipPermission.setIpRanges(ipRanges);
50 ipPermissions.add(ipPermission);

    try {
        // Authorize the ports to be used.
        AuthorizeSecurityGroupIngressRequest ingressRequest =
65        new AuthorizeSecurityGroupIngressRequest("GettingStartedGroup", ip
```

```
Permissions);
    ec2.authorizeSecurityGroupIngress(ingressRequest);
} catch (AmazonServiceException ase) {
    // Ignore because this likely means the zone has already
    // been authorized.
60    System.out.println(ase.getMessage());
}
```

You can view this entire code sample in the `advanced.CreateSecurityGroupApp.java` code sample. Note you only need to run this application once to create a new security group.

You can also create the security group using the AWS Toolkit for Eclipse. Go to the [toolkit documentation](#) for more information.

### Detailed Spot Instance Request Creation Options

As we explained in [Tutorial: Amazon EC2 Spot Instances \(p. 166\)](#), you need to build your request with an instance type, an Amazon Machine Image (AMI), and maximum bid price.

Let's start by creating a `RequestSpotInstanceRequest` object. The request object requires the number of instances you want and the bid price. Additionally, we need to set the `LaunchSpecification` for the request, which includes the instance type, AMI ID, and security group you want to use. After the request is populated, we call the `requestSpotInstances` method on the `AmazonEC2Client` object. An example of how to request a Spot instance follows.

(The following code is the same as what we used in the first tutorial.)

```
1
    // Retrieves the credentials from an AWSCredentials.properties file.
    AWSCredentials credentials = null;
    try {
5        credentials = new PropertiesCredentials(
            GettingStartedApp.class.getResourceAsStream("AwsCredentials.properties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCredentials.properties.");
        System.out.println(e1.getMessage());
10       System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Initializes a Spot Instance Request
    RequestSpotInstancesRequest requestRequest = new RequestSpotInstancesRequest();

    // Request 1 x t1.micro instance with a bid price of $0.03.
20    requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(1));

    // Set up the specifications of the launch. This includes the
    // instance type (e.g. t1.micro) and the latest Amazon Linux
25 // AMI id available. Note, you should always use the latest
    // Amazon Linux AMI id or another of your choosing.
```

```
LaunchSpecification launchSpecification = new LaunchSpecification();
launchSpecification.setImageId("ami-8c1fece5");
launchSpecification.setInstanceType("t1.micro");
30
// Add the security group to the request.
ArrayList<String> securityGroups = new ArrayList<String>();
securityGroups.add("GettingStartedGroup");
launchSpecification.setSecurityGroups(securityGroups);
35
// Add the launch specification.
requestRequest.setLaunchSpecification(launchSpecification);

// Call the RequestSpotInstance API.
40 RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

### Persistent vs. One-Time Requests

When building a Spot request, you can specify several optional parameters. The first is whether your request is one-time only or persistent. By default, it is a one-time request. A one-time request can be fulfilled only once, and after the requested instances are terminated, the request will be closed. A persistent request is considered for fulfillment whenever there is no Spot Instance running for the same request. To specify the type of request, you simply need to set the Type on the Spot request. This can be done with the following code.

```
1
// Retrieves the credentials from an
// AWSCredentials.properties file.
AWSCredentials credentials = null;
5 try {
    credentials = new PropertiesCredentials(
        GettingStartedApp.class.getResourceAsStream("AwsCredentials.proper
ties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCre
dentials.properties.");
10    System.out.println(e1.getMessage());
        System.exit(-1);
    }

// Create the AmazonEC2Client object so we can call various APIs.
15 AmazonEC2 ec2 = new AmazonEC2Client(credentials);

// Initializes a Spot Instance Request
RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

20 // Request 1 x t1.micro instance with a bid price of $0.03.
    requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(1));

// Set the type of the bid to persistent.
25 requestRequest.setType("persistent");

// Set up the specifications of the launch. This includes the
// instance type (e.g. t1.micro) and the latest Amazon Linux
```

```
    // AMI id available. Note, you should always use the latest
30 // Amazon Linux AMI id or another of your choosing.
    LaunchSpecification launchSpecification = new LaunchSpecification();
    launchSpecification.setImageId("ami-8c1fece5");
    launchSpecification.setInstanceType("t1.micro");

35 // Add the security group to the request.
    ArrayList<String> securityGroups = new ArrayList<String>();
    securityGroups.add("GettingStartedGroup");
    launchSpecification.setSecurityGroups(securityGroups);

40 // Add the launch specification.
    requestRequest.setLaunchSpecification(launchSpecification);

    // Call the RequestSpotInstance API.
    RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
45
```

### Limiting the Duration of a Request

You can also optionally specify the length of time that your request will remain valid. You can specify both a starting and ending time for this period. By default, a Spot request will be considered for fulfillment from the moment it is created until it is either fulfilled or canceled by you. However you can constrain the validity period if you need to. An example of how to specify this period is shown in the following code.

```
1
    // Retrieves the credentials from an AWSCredentials.properties file.
    AWSCredentials credentials = null;
    try {
5        credentials = new PropertiesCredentials(
            GettingStartedApp.class.getResourceAsStream("AwsCredentials.proper
ties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCre
dentials.properties.");
        System.out.println(e1.getMessage());
10        System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Initializes a Spot Instance Request
    RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

    // Request 1 x t1.micro instance with a bid price of $0.03.
20 requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(1));

    // Set the valid start time to be two minutes from now.
    Calendar cal = Calendar.getInstance();
25 cal.add(Calendar.MINUTE, 2);
    requestRequest.setValidFrom(cal.getTime());

    // Set the valid end time to be two minutes and two hours from now.
```

```
    cal.add(Calendar.HOUR, 2);
30 requestRequest.setValidUntil(cal.getTime());

    // Set up the specifications of the launch. This includes
    // the instance type (e.g. t1.micro)

35 // and the latest Amazon Linux AMI id available.
    // Note, you should always use the latest Amazon
    // Linux AMI id or another of your choosing.
    LaunchSpecification launchSpecification = new LaunchSpecification();
    launchSpecification.setImageId("ami-8c1fece5");
40 launchSpecification.setInstanceType("t1.micro");

    // Add the security group to the request.
    ArrayList<String> securityGroups = new ArrayList<String>();
    securityGroups.add("GettingStartedGroup");
45 launchSpecification.setSecurityGroups(securityGroups);

    // Add the launch specification.
    requestRequest.setLaunchSpecification(launchSpecification);

50 // Call the RequestSpotInstance API.
    RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

## Grouping Your Amazon EC2 Spot Instance Requests

You have the option of grouping your Spot instance requests in several different ways. We'll look at the benefits of using launch groups, Availability Zone groups, and placement groups.

If you want to ensure your Spot instances are all launched and terminated together, then you have the option to leverage a launch group. A launch group is a label that groups a set of bids together. All instances in a launch group are started and terminated together. Note, if instances in a launch group have already been fulfilled, there is no guarantee that new instances launched with the same launch group will also be fulfilled. An example of how to set a Launch Group is shown in the following code example.

```
1
    // Retrieves the credentials from an AWSCredentials.properties file.
    AWSCredentials credentials = null;
    try {
5        credentials = new PropertiesCredentials(
            GettingStartedApp.class.getResourceAsStream("AwsCredentials.proper
ties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCre
dentials.properties.");
        System.out.println(e1.getMessage());
10        System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Initializes a Spot Instance Request
    RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();
```

```
// Request 5 x t1.micro instance with a bid price of $0.03.
20 requestRequest.setSpotPrice("0.03");
   requestRequest.setInstanceCount(Integer.valueOf(5));

   // Set the launch group.
   requestRequest.setLaunchGroup("ADVANCED-DEMO-LAUNCH-GROUP");
25

   // Set up the specifications of the launch. This includes
   // the instance type (e.g. t1.micro) and the latest Amazon Linux
   // AMI id available. Note, you should always use the latest
   // Amazon Linux AMI id or another of your choosing.
30 LaunchSpecification launchSpecification = new LaunchSpecification();
   launchSpecification.setImageId("ami-8c1fece5");
   launchSpecification.setInstanceType("t1.micro");

   // Add the security group to the request.
35 ArrayList<String> securityGroups = new ArrayList<String>();
   securityGroups.add("GettingStartedGroup");
   launchSpecification.setSecurityGroups(securityGroups);

   // Add the launch specification.
40 requestRequest.setLaunchSpecification(launchSpecification);

   // Call the RequestSpotInstance API.
   RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

If you want to ensure that all instances within a request are launched in the same Availability Zone, and you don't care which one, you can leverage Availability Zone groups. An Availability Zone group is a label that groups a set of instances together in the same Availability Zone. All instances that share an Availability Zone group and are fulfilled at the same time will start in the same Availability Zone. An example of how to set an Availability Zone group follows.

```
1
   // Retrieves the credentials from an AWSCredentials.properties file.
   AWSCredentials credentials = null;
   try {
5     credentials = new PropertiesCredentials(
       GettingStartedApp.class.getResourceAsStream("AwsCredentials.proper
ties"));
   } catch (IOException e1) {
       System.out.println("Credentials were not properly entered into AwsCre
dentials.properties.");
       System.out.println(e1.getMessage());
10    System.exit(-1);
   }

   // Create the AmazonEC2Client object so we can call various APIs.
   AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15

   // Initializes a Spot Instance Request
   RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

   // Request 5 x t1.micro instance with a bid price of $0.03.
```



```
20 requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(5));

    // Set the availability zone group.
    requestRequest.setAvailabilityZoneGroup("ADVANCED-DEMO-AZ-GROUP");
25
    // Set up the specifications of the launch. This includes the instance
    // type (e.g. t1.micro) and the latest Amazon Linux AMI id available.
    // Note, you should always use the latest Amazon Linux AMI id or another
    // of your choosing.
30 LaunchSpecification launchSpecification = new LaunchSpecification();
    launchSpecification.setImageId("ami-8c1fece5");
    launchSpecification.setInstanceType("t1.micro");

    // Add the security group to the request.
35 ArrayList<String> securityGroups = new ArrayList<String>();
    securityGroups.add("GettingStartedGroup");
    launchSpecification.setSecurityGroups(securityGroups);

    // Add the launch specification.
40 requestRequest.setLaunchSpecification(launchSpecification);

    // Call the RequestSpotInstance API.
    RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

You can specify an Availability Zone that you want for your Spot Instances. The following code example shows you how to set an Availability Zone.

```
1
    // Retrieves the credentials from an AWSCredentials.properties file.
    AWSCredentials credentials = null;
    try {
        5 credentials = new PropertiesCredentials(
            GettingStartedApp.class.getResourceAsStream("AwsCredentials.proper
ties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCre
dentials.properties.");
        System.out.println(e1.getMessage());
10    System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Initializes a Spot Instance Request
    RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

    // Request 1 x t1.micro instance with a bid price of $0.03.
20 requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(1));

    // Set up the specifications of the launch. This includes the instance
    // type (e.g. t1.micro) and the latest Amazon Linux AMI id available.
```

```
25 // Note, you should always use the latest Amazon Linux AMI id or another
    // of your choosing.
    LaunchSpecification launchSpecification = new LaunchSpecification();
    launchSpecification.setImageId("ami-8c1fece5");
    launchSpecification.setInstanceType("t1.micro");
30
    // Add the security group to the request.
    ArrayList<String> securityGroups = new ArrayList<String>();
    securityGroups.add("GettingStartedGroup");
    launchSpecification.setSecurityGroups(securityGroups);
35
    // Set up the availability zone to use. Note we could retrieve the
    // availability zones using the ec2.describeAvailabilityZones() API. For
    // this demo we will just use us-east-1a.
    SpotPlacement placement = new SpotPlacement("us-east-1b");
40
    launchSpecification.setPlacement(placement);

    // Add the launch specification.
    requestRequest.setLaunchSpecification(launchSpecification);
45
    // Call the RequestSpotInstance API.
    RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

Lastly, you can specify a *placement group* if you are using High Performance Computing (HPC) Spot instances, such as cluster compute instances or cluster GPU instances. Placement groups provide you with lower latency and high-bandwidth connectivity between the instances. An example of how to set a placement group follows.

```
1
    // Retrieves the credentials from an AWSCredentials.properties file.
    AWSCredentials credentials = null;
    try {
2     credentials = new PropertiesCredentials(
        GettingStartedApp.class.getResourceAsStream("AwsCredentials.proper
ties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCre
dentials.properties.");
        System.out.println(e1.getMessage());
10     System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Initializes a Spot Instance Request
    RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

    // Request 1 x t1.micro instance with a bid price of $0.03.
20 requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(1));

    // Set up the specifications of the launch. This includes the instance
```

```
    // type (e.g. t1.micro) and the latest Amazon Linux AMI id available.
25 // Note, you should always use the latest Amazon Linux AMI id or another
    // of your choosing.
    LaunchSpecification launchSpecification = new LaunchSpecification();
    launchSpecification.setImageId("ami-8c1fece5");
    launchSpecification.setInstanceType("t1.micro");
30
    // Add the security group to the request.
    ArrayList<String> securityGroups = new ArrayList<String>();
    securityGroups.add("GettingStartedGroup");
    launchSpecification.setSecurityGroups(securityGroups);
35
    // Set up the placement group to use with whatever name you desire.
    // For this demo we will just use "ADVANCED-DEMO-PLACEMENT-GROUP".
    SpotPlacement placement = new SpotPlacement();
    placement.setGroupName("ADVANCED-DEMO-PLACEMENT-GROUP");
40 launchSpecification.setPlacement(placement);

    // Add the launch specification.
    requestRequest.setLaunchSpecification(launchSpecification);

45 // Call the RequestSpotInstance API.
    RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

All of the parameters shown in this section are optional. It is also important to realize that most of these parameters—with the exception of whether your bid is one-time or persistent—can reduce the likelihood of bid fulfillment. So, it is important to leverage these options only if you need them. All of the preceding code examples are combined into one long code sample, which can be found in the `com.amazonaws.codesamples.advanced.InlineGettingStartedCodeSampleApp.java` class.

### How to Persist a Root Partition After Interruption or Termination

One of the easiest ways to manage interruption of your Spot instances is to ensure that your data is checkpointed to an Amazon Elastic Block Store (Amazon EBS) volume on a regular cadence. By checkpointing periodically, if there is an interruption you will lose only the data created since the last checkpoint (assuming no other non-idempotent actions are performed in between). To make this process easier, you can configure your Spot Request to ensure that your root partition will not be deleted on interruption or termination. We've inserted new code in the following example that shows how to enable this scenario.

In the added code, we create a `BlockDeviceMapping` object and set its associated Elastic Block Storage (EBS) to an EBS object that we've configured to not be deleted if the Spot Instance is terminated. We then add this `BlockDeviceMapping` to the `ArrayList` of mappings that we include in the launch specification.

```
1
    // Retrieves the credentials from an AWSCredentials.properties file.
    AWSCredentials credentials = null;
    try {
2     credentials = new PropertiesCredentials(
        GettingStartedApp.class.getResourceAsStream("AwsCredentials.properties"));
    } catch (IOException e1) {
        System.out.println("Credentials were not properly entered into AwsCre
dentials.properties.");
```

```
        System.out.println(e1.getMessage());
10     System.exit(-1);
    }

    // Create the AmazonEC2Client object so we can call various APIs.
    AmazonEC2 ec2 = new AmazonEC2Client(credentials);
15
    // Initializes a Spot Instance Request
    RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

    // Request 1 x t1.micro instance with a bid price of $0.03.
20 requestRequest.setSpotPrice("0.03");
    requestRequest.setInstanceCount(Integer.valueOf(1));

    // Set up the specifications of the launch. This includes the instance
    // type (e.g. t1.micro) and the latest Amazon Linux AMI id available.
25 // Note, you should always use the latest Amazon Linux AMI id or another
    // of your choosing.
    LaunchSpecification launchSpecification = new LaunchSpecification();
    launchSpecification.setImageId("ami-8c1fece5");
    launchSpecification.setInstanceType("t1.micro");
30
    // Add the security group to the request.
    ArrayList<String> securityGroups = new ArrayList<String>();
    securityGroups.add("GettingStartedGroup");
    launchSpecification.setSecurityGroups(securityGroups);
35
    // Create the block device mapping to describe the root partition.
    BlockDeviceMapping blockDeviceMapping = new BlockDeviceMapping();
    blockDeviceMapping.setDeviceName("/dev/sda1");

40 // Set the delete on termination flag to false.
    EbsBlockDevice ebs = new EbsBlockDevice();
    ebs.setDeleteOnTermination(Boolean.FALSE);
    blockDeviceMapping.setEbs(ebs);

45 // Add the block device mapping to the block list.
    ArrayList<BlockDeviceMapping> blockList = new ArrayList<BlockDeviceMap
ping>();
    blockList.add(blockDeviceMapping);

    // Set the block device mapping configuration in the launch specifications.
50 launchSpecification.setBlockDeviceMappings(blockList);

    // Add the launch specification.
    requestRequest.setLaunchSpecification(launchSpecification);

55 // Call the RequestSpotInstance API.
    RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
```

Assuming you wanted to re-attach this volume to your instance on startup, you can also use the block device mapping settings. Alternatively, if you attached a non-root partition, you can specify the Amazon EBS volumes you want to attach to your Spot instance after it resumes. You do this simply by specifying a snapshot ID in your `EbsBlockDevice` and alternative device name in your `BlockDeviceMapping` objects. By leveraging block device mappings, it can be easier to bootstrap your instance.

Using the root partition to checkpoint your critical data is a great way to manage the potential for interruption of your instances. For more methods on managing the potential of interruption, please visit the [Managing Interruption](#) video.

## How to Tag Your Spot Requests and Instances

Adding [tags to EC2 resources](#) can simplify the administration of your cloud infrastructure. A form of metadata, tags can be used to create user-friendly names, enhance searchability, and improve coordination between multiple users. You can also use tags to automate scripts and portions of your processes.

To add tags to your resources, you need to tag them *after* they have been requested. Specifically, you must add a tag after a Spot request has been submitted or after the `RunInstances` call has been performed. The following code example illustrates adding tags.

```
1
  /*
   * Copyright 2010-2011 Amazon.com, Inc. or its affiliates. All Rights Re
served.
   *
   5  * Licensed under the Apache License, Version 2.0 (the "License").
   * You may not use this file except in compliance with the License.
   * A copy of the License is located at
   *
   * http://aws.amazon.com/apache2.0
  10 *
   * or in the "license" file accompanying this file. This file is distributed
   * on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either
   * express or implied. See the License for the specific language governing
   * permissions and limitations under the License.
  15 */
package com.amazonaws.codesamples.advanced;

import java.io.IOException;
import java.util.ArrayList;
  20 import java.util.List;

import com.amazonaws.AmazonServiceException;
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.PropertiesCredentials;
  25 import com.amazonaws.codesamples.getting_started.GettingStartedApp;
import com.amazonaws.services.ec2.AmazonEC2;
import com.amazonaws.services.ec2.AmazonEC2Client;
import com.amazonaws.services.ec2.model.CancelSpotInstanceRequestsRequest;
import com.amazonaws.services.ec2.model.CreateTagsRequest;
  30 import com.amazonaws.services.ec2.model.DescribeSpotInstanceRequestsRequest;
import com.amazonaws.services.ec2.model.DescribeSpotInstanceRequestsResult;
import com.amazonaws.services.ec2.model.LaunchSpecification;
import com.amazonaws.services.ec2.model.RequestSpotInstancesRequest;
import com.amazonaws.services.ec2.model.RequestSpotInstancesResult;
  35 import com.amazonaws.services.ec2.model.SpotInstanceRequest;
import com.amazonaws.services.ec2.model.Tag;
import com.amazonaws.services.ec2.model.TerminateInstancesRequest;

/**
  40 * Welcome to your new AWS Java SDK based project!
   *
   * This class is meant as a starting point for your console-based application
that
```

```

    * makes one or more calls to the AWS services supported by the Java SDK,
such as EC2,
    * SimpleDB, and S3.
45  *
    * In order to use the services in this sample, you need:
    *
    * - A valid Amazon Web Services account. You can register for AWS at:
    *   https://aws-portal.amazon.com/gp/aws/developer/registration/in
dex.html
50  *
    * - Your account's Access Key ID and Secret Access Key:
    *   http://aws.amazon.com/security-credentials
    *
    * - A subscription to Amazon EC2. You can sign up for EC2 at:
55  *   http://aws.amazon.com/ec2/
    *
    */

public class InlineTaggingCodeSampleApp {
60
    /**
    * @param args
    */
    public static void main(String[] args) {
65        //=====
        //===== Submitting a Request =====
        //=====

        // Retrieves the credentials from an AWSCredentials.properties file.
70        AWSCredentials credentials = null;
        try {
            credentials = new PropertiesCredentials(
                GettingStartedApp.class.getResourceAsStream("AwsCredentials.proper
ties"));
        } catch (IOException e1) {
75            System.out.println("Credentials were not properly entered into
AwsCredentials.properties.");
            System.out.println(e1.getMessage());
            System.exit(-1);
        }

80        // Create the AmazonEC2Client object so we can
        // call various APIs.
        AmazonEC2 ec2 = new AmazonEC2Client(credentials);

        // Initializes a Spot Instance Request
85        RequestSpotInstancesRequest requestRequest = new RequestSpotInstances
Request();

        // Request 1 x t1.micro instance with a bid price of $0.03.
        requestRequest.setSpotPrice("0.03");
        requestRequest.setInstanceCount(Integer.valueOf(1));
90

        // Set up the specifications of the launch. This includes
        // the instance type (e.g. t1.micro) and the latest Amazon
        // Linux AMI id available. Note, you should always use the
        // latest Amazon Linux AMI id or another of your choosing.
95        LaunchSpecification launchSpecification = new LaunchSpecification();

```

```
        launchSpecification.setImageId("ami-8c1fece5");
        launchSpecification.setInstanceType("t1.micro");

        // Add the security group to the request.
100    ArrayList<String> securityGroups = new ArrayList<String>();
        securityGroups.add("GettingStartedGroup");
        launchSpecification.setSecurityGroups(securityGroups);

        // Add the launch specifications to the request.
105    requestRequest.setLaunchSpecification(launchSpecification);

        //=====//
        //===== Getting the Request ID from the Request =====//
        //=====//

110    // Call the RequestSpotInstance API.
        RequestSpotInstancesResult requestResult = ec2.requestSpotInstances(re
questRequest);
        List<SpotInstanceRequest> requestResponses = requestResult.getSpotIn
stanceRequests();

115    // Set up an arraylist to collect all of the request ids we want to
        // watch hit the running state.
        ArrayList<String> spotInstanceRequestIds = new ArrayList<String>();

        // Add all of the request ids to the hashset, so we can
120    // determine when they hit the active state.
        for (SpotInstanceRequest requestResponse : requestResponses) {
            System.out.println("Created Spot Request: "+requestResponse.getSpot
InstanceRequestId());
            spotInstanceRequestIds.add(requestResponse.getSpotInstanceRequest
Id());
        }

125    //=====//
        //===== Tag the Spot Requests =====//
        //=====//

130    // Create the list of tags we want to create
        ArrayList<Tag> requestTags = new ArrayList<Tag>();
        requestTags.add(new Tag("keyname1", "value1"));

        // Create a tag request for the requests.
135    CreateTagsRequest createTagsRequest_requests = new CreateTagsRequest();
        createTagsRequest_requests.setResources(spotInstanceRequestIds);
        createTagsRequest_requests.setTags(requestTags);

        // Try to tag the Spot request submitted.
140    try {
            ec2.createTags(createTagsRequest_requests);
        } catch (AmazonServiceException e) {
            // Write out any exceptions that may have occurred.
            System.out.println("Error terminating instances");
145    System.out.println("Caught Exception: " + e.getMessage());
            System.out.println("Reponse Status Code: " + e.getStatusCode());
            System.out.println("Error Code: " + e.getErrorCode());
            System.out.println("Request ID: " + e.getRequestId());
        }

150
```





```

205     }
        } while (anyOpen);

        //=====//
        //===== Tag the Spot Instances =====//
210     //=====//

        // Create the list of tags we want to create
        ArrayList<Tag> instanceTags = new ArrayList<Tag>();
        instanceTags.add(new Tag("keyname1","value1"));
215

        // Create a tag request for instances.
        CreateTagsRequest createTagsRequest_instances = new CreateTagsRequest();
        createTagsRequest_instances.setResources(instanceIds);
        createTagsRequest_instances.setTags(instanceTags);

220

        // Try to tag the Spot instance started.
        try {
            ec2.createTags(createTagsRequest_instances);
        } catch (AmazonServiceException e) {
225            // Write out any exceptions that may have occurred.
            System.out.println("Error terminating instances");
            System.out.println("Caught Exception: " + e.getMessage());
            System.out.println("Reponse Status Code: " + e.getStatusCode());
            System.out.println("Error Code: " + e.getErrorCode());
230            System.out.println("Request ID: " + e.getRequestId());
        }

        //=====//
        //===== Canceling the Request =====//
235     //=====//

        try {
            // Cancel requests.
            CancelSpotInstanceRequestsRequest cancelRequest = new CancelSpotIn
instanceRequestsRequest(spotInstanceRequestIds);
240            ec2.cancelSpotInstanceRequests(cancelRequest);
        } catch (AmazonServiceException e) {
            // Write out any exceptions that may have occurred.
            System.out.println("Error canceling instances");
            System.out.println("Caught Exception: " + e.getMessage());
245            System.out.println("Reponse Status Code: " + e.getStatusCode());
            System.out.println("Error Code: " + e.getErrorCode());
            System.out.println("Request ID: " + e.getRequestId());
        }

250     //=====//
        //===== Terminating any Instances =====//
        //=====//
        try {
            // Terminate instances.
255            TerminateInstancesRequest terminateRequest = new TerminateInstances
Request(instanceIds);
            ec2.terminateInstances(terminateRequest);
        } catch (AmazonServiceException e) {
            // Write out any exceptions that may have occurred.
            System.out.println("Error terminating instances");
260            System.out.println("Caught Exception: " + e.getMessage());

```

```
        System.out.println("Reponse Status Code: " + e.getStatusCode());
        System.out.println("Error Code: " + e.getErrorCode());
        System.out.println("Request ID: " + e.getRequestId());
    }
265 } // main
    }
```

Tags are a simple first step toward making it easier to manage your own cluster of instances. To read more about tagging Amazon EC2 resources, go to [Using Tags](#) in the Amazon Elastic Compute Cloud User Guide.

### Bringing It All Together

To bring this all together, we provide a more object-oriented approach that combines the steps we showed in this tutorial into one easy to use class. We instantiate a class called `Requests` that performs these actions. We also create a `GettingStartedApp` class, which has a main method where we perform the high level function calls.

The complete source code is available for download at [GitHub](#).

Congratulations! You've completed the Advanced Request Features tutorial for developing Spot Instance software with the AWS SDK for Java.

## Starting Clusters on Spot Instances

Grids are a form of distributed computing that enable a user to leverage multiple instances to perform parallel computations. Customers—such as [Numerate](#), [Scribd](#), and the [University of Barcelona/University of Melbourne](#)—use Grid Computing with Spot Instances because this type of architecture can take advantage of Spot Instance's built-in elasticity and low prices to get work done faster at a more cost-effective price.

To get started, a user will break down the work into discrete units called jobs, and then submit that work to a "master node." These jobs will be queued up, and a process called a "scheduler" will distribute that work out to other instances in the grid, called "worker nodes." After the result is computed by the worker node, the master node is notified, and the worker node can take the next operation from the queue. If the job fails or the instance is interrupted, the job will automatically be re-queued by the scheduler process.

As you work to architect your application, it is important to choose the appropriate amount of work to be included in your job. We recommend breaking your jobs down into a logical grouping based on the time it would take to process. Typically, you will want to create a workload size less than an hour, so that if you have to process the workload again, it doesn't cost you additional money (you don't pay for the hour if we interrupt your instance).

Many customers use a Grid scheduler, such as Oracle Grid Engine or UniCloud, to set up a cluster. If you have long-running workloads, the best practice is to run the master node on On-Demand or Reserved Instances, and run the worker nodes on Spot or a mixture of On-Demand, Reserved, and Spot Instances. Alternatively, if you have a workload that is less than an hour or you are running a test environment, you may want to run all of your instances on Spot. No matter the setup, we recommend that you create a script to automatically re-add instances that may be interrupted. Some existing tools—StarCluster, for example— can help you manage this process.

### Quick Look: How to Launch a Cluster on Spot Video

Chris Dagdigan, from AWS Solution Provider [BioTeam](#), provides a quick overview of how to start a cluster from scratch in about 10 to 15 minutes on Amazon EC2 Spot Instances using StarCluster. StarCluster is an open source tool created by a lab at MIT that makes it easy to set up a new Oracle Grid Engine cluster.

In this video, Chris walks through the process of installing, setting up, and running simple jobs on a cluster. Chris also leverages Spot Instances, so that you can potentially get work done faster and potentially save between 50 percent to 66 percent. [How to Launch a Cluster on Spot](#)

## Reserved Instances

Amazon Elastic Compute Cloud (Amazon EC2) Reserved Instances is a pricing model that enables you to reserve capacity for your EC2 instances and lowers your average instance cost. With Reserved Instances, you pay a low, one-time fee for the capacity reservation and then receive a significant discount on the hourly charge for your instances. When you want to use your reserved capacity, you launch an EC2 instance with the same configuration as the reserved capacity that you purchased. Amazon Web Services (AWS) will automatically apply the discounted hourly rate that is associated with your capacity reservation. You are charged the discounted hourly rate for your EC2 instance for as long as you own the Reserved Instance. When the term of your Reserved Instance ends, you can continue using the EC2 instance without interruption. However, you will now be charged at the On-Demand rate.

Reserved Instances can provide substantial savings over owning your own hardware or running only On-Demand instances, as well as help assure that the capacity you need is available to you when you require it.

To purchase an Amazon EC2 Reserved Instance, you must select an instance type (such as m1.small), platform (Linux/UNIX, Windows, or Windows with SQL Server), location (region and Availability Zone), and term (either one year or three years). If you want your Reserved Instance to run on a specific Linux/UNIX platform, you must identify that platform when you purchase the reserved capacity. Then, when you're ready to use the Reserved Instance that you purchased, you must choose an Amazon Machine Image (AMI) that runs that specific Linux/UNIX platform, along with any other specifications you identified during the purchase.

For example, if you require a one-year, *SUSE Linux*, m1.medium Reserved Instance in the Singapore region, purchasing a one-year, *Linux/UNIX*, m1.medium Reserved Instance in the Singapore region will not give you the capacity guarantee and pricing benefit. Both your Reserved Instance purchase and the instance you launch must specify the same *SUSE Linux* product platform.

For product pricing information, see the following pages:

- [AWS Service Pricing Overview](#)
- [Amazon EC2 On-Demand Instances Pricing](#)
- [Amazon EC2 Reserved Instance Pricing](#)

## Reserved Instance Overview

The following information will help you get started working with Amazon EC2 Reserved Instances:

- Complete all the prerequisite tasks first—such as registration, signing up, and installing the tools—so you can start working with Reserved Instances. For more information, see [Getting Started with Reserved Instances](#) (p. 194).
- Before you buy and sell Reserved Instances, you can learn more about them by reading [Steps for Using Reserved Instances](#) (p. 195).
- Standard one- and three-year terms for Reserved Instances are available for purchase from AWS, and non-standard terms are available for purchase from third-party resellers through the Reserved Instance Marketplace.
- Optimize your Reserved Instance costs by selecting the pricing model that best matches how often you plan to use your instances. For more information, see [Choosing Reserved Instances Based on Your Usage Plans](#) (p. 199).
- Learn more about the pricing benefit of Reserved Instances. For more information, see [Understanding the Pricing Benefit of Reserved Instances](#) (p. 208).
- Understand Reserved Instance pricing tiers and how to take advantage of discount pricing. For more information, see [Understanding Reserved Instance Pricing Tiers](#) (p. 200).

- You can sell your unused Reserved Instances in the Reserved Instance Marketplace. The Reserved Instance Marketplace makes it possible for sellers who have Reserved Instances that they no longer need to find buyers who are looking to purchase additional capacity. Reserved Instances bought and sold through the Reserved Instance Marketplace work like any other Reserved Instances. For more information, see [Reserved Instance Marketplace](#) (p. 209).

For a checklist that summarizes requirements for working with Reserved Instances and the Reserved Instance Marketplace, see [Requirements Checklist for Reserved Instances](#) (p. 262).

## What Do You Want to Do Next?

- Learn:
  - [Getting Started with Reserved Instances](#) (p. 194)
  - [Buying Reserved Instances](#) (p. 212)
  - [Selling in the Reserved Instance Marketplace](#) (p. 236)
  - [Using Reserved Instances in Amazon VPC](#) (p. 196)
- Start:
  - [Becoming a Buyer](#) (p. 213)
  - [Purchasing Reserved Instances](#) (p. 214)
  - [Obtaining Information About Your Reserved Instances](#) (p. 222)
  - [Modifying Your Reserved Instances](#) (p. 227)
  - [Registering as a Seller](#) (p. 237)
  - [Listing Your Reserved Instance](#) (p. 242)

## Getting Started with Reserved Instances

### Topics

- [Get Set Up](#) (p. 194)
- [Steps for Using Reserved Instances](#) (p. 195)
- [Using Reserved Instances in Amazon VPC](#) (p. 196)
- [Tools for Working with Reserved Instances](#) (p. 197)

You can use Amazon Elastic Compute Cloud (Amazon EC2) Reserved Instances to reserve capacity for your instances and get the benefits of lower-cost computing. With Reserved Instances you pay a low, one-time fee and in turn receive a significant discount on the hourly charge for your instance. Reserved Instances can provide substantial savings over owning your own hardware or running only On-Demand instances, as well as help assure that the capacity you need is available to you when you require it. This topic takes you through the basic information you need to get started with Reserved Instances

### Get Set Up

Before you get started working with Reserved Instances, you should complete the following tasks:

- Sign up.
  - To work with Reserved Instances, read and complete the instructions described in [Getting Started with Amazon EC2 Linux Instances](#) (p. 22), which provides information on signing up for your Amazon EC2 account and credentials.
- Install the tools.

You can use the Amazon EC2 tools—AWS Management Console, Amazon EC2 Command Line Interface (CLI) tools, or the Amazon EC2 API—to work with EC2 Reserved Instances and search for offerings. For more information, see [Tools for Working with Reserved Instances \(p. 197\)](#).

If you want to start working with Reserved Instances using specific tools, see [AWS Management Console \(p. 197\)](#), the [Command Line Interface Tools \(p. 197\)](#), or the [API \(p. 198\)](#).

## Steps for Using Reserved Instances

There are five sets of steps to follow when you use Reserved Instances. You can *purchase* Amazon EC2 Reserved Instances, and then you can *launch* them. You can *view* the Reserved Instances you have, *modify* them, and you can *sell* unused Reserved Instances in the Reserved Instance Marketplace. (Restrictions apply. For information, see [Requirements Checklist for Reserved Instances \(p. 262\)](#).) This section describes purchasing, launching, viewing, modifying, and selling Reserved Instances.

You can use the AWS Management Console, the Amazon EC2 CLI tools, or the Amazon EC2 API to perform any of these tasks. Before you get started, you need to set up the prerequisite accounts and tools. For more information, see [Get Set Up \(p. 194\)](#).

### 1. Purchase.

- a. Determine how much capacity you want to reserve. Specify the following criteria for your instance reservation.
  - Platform (for example, Linux/UNIX).

#### Note

When you want your Reserved Instance to run on a specific Linux/UNIX platform, you must identify the specific platform when you purchase the reserved capacity. Then, when you launch your instance with the intention of using the reserved capacity you purchased, you must choose the Amazon Machine Image (AMI) that runs that specific Linux/UNIX platform, along with any other specifications you identified during the purchase.

- Instance type (for example, m1.small).
  - Term (time period) over which you want to reserve capacity (one or three years).
  - Tenancy specification, if you want to reserve capacity for your instance to run in single-tenant hardware (*dedicated* tenancy, as opposed to *shared*).
  - Region and Availability Zone where you want to run the instance.
- b. Choose the offering type that best addresses how much you want to pay and how often you plan to run your instances.
    - Heavy Utilization
    - Medium Utilization
    - Light Utilization

For more information about these offering types, see [Choosing Reserved Instances Based on Your Usage Plans \(p. 199\)](#).

- c. Search for offerings that meet the criteria you specified.
- d. Purchase offerings that fulfill your requirements.

For more information, see [Purchasing Reserved Instances \(p. 214\)](#).

### 2. Launch.

To use your Reserved Instance, launch an On-Demand EC2 instance with the same criteria as your Reserved Instance (the region, Availability Zone, instance type, and platform specified when you purchased your Reserved Instance). See step 1.

**Note**

Reserved Instance pricing benefits and capacity guarantees automatically apply to any running EC2 instances you have that aren't already covered by a reservation.

For more information, see [Launch Your Instance \(p. 266\)](#).

3. View.

You can view the Reserved Instances that you own or that are available to your account, and confirm that your instances are running as specified.

For more information, see [Obtaining Information About Your Reserved Instances \(p. 222\)](#).

4. Modify.

You can modify your Reserved Instances by moving them between Availability Zones within the same region, changing their instance type to another instance type in the same instance family, or modifying the network platform of the Reserved Instances between EC2-VPC and EC2-Classic.

For more information, see [Modifying Your Reserved Instances \(p. 227\)](#).

5. Sell Reserved Instance capacity that you no longer need.

a. Register as a seller in the Reserved Instance Marketplace using the Seller Registration wizard. For more information, see [Registering as a Seller \(p. 237\)](#).

**Note**

Not every customer can sell in the Reserved Instance Marketplace and not all Reserved Instances can be sold in the Reserved Instance Marketplace. For information, see [Requirements Checklist for Reserved Instances \(p. 262\)](#).

b. Decide on a price for the Reserved Instances that you want to sell. For more information, see [Pricing Your Reserved Instances \(p. 241\)](#).

c. List your Reserved Instances. For more information, see [Listing Your Reserved Instance \(p. 242\)](#).

d. Find out how you get paid when your Reserved Instances are sold. For more information, see [Getting Paid \(p. 261\)](#).

## Using Reserved Instances in Amazon VPC

To launch Reserved Instances in Amazon Virtual Private Cloud (Amazon VPC), you must either have an account that supports a default VPC or you must purchase an Amazon VPC Reserved Instance.

If your account does not support a default VPC, you must purchase an Amazon VPC Reserved Instance by selecting a platform that includes *Amazon VPC* in its name. For more information, see [Detecting Your Supported Platforms and Whether You Have a Default VPC](#). For information about Amazon VPC, see [What is Amazon VPC?](#) in the *Amazon Virtual Private Cloud User Guide*.

If your account supports a default VPC, the list of platforms available does not include *Amazon VPC* in its name because all platforms have default subnets. In this case, if you launch an instance with the same configuration as the capacity you reserved and paid for, that instance is launched in your default VPC and the capacity guarantees and billing benefits are applied to your instance. For information about default VPCs, see [Your Default VPC and Subnets](#) in the *Amazon Virtual Private Cloud User Guide*.

You can also choose to purchase Reserved Instances that are physically isolated at the host hardware level by specifying *dedicated* as the instance tenancy. For more information about Dedicated Instances, see [Using EC2 Dedicated Instances Within Your VPC](#) in the *Amazon Virtual Private Cloud User Guide*.

## Tools for Working with Reserved Instances

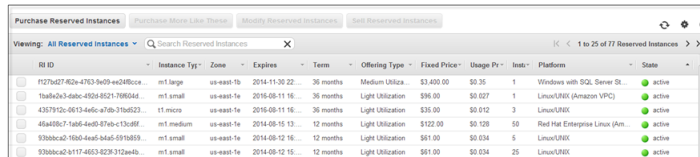
You can use the [AWS Management Console \(p. 197\)](#), the [Command Line Interface Tools \(p. 197\)](#), or the [API \(p. 198\)](#) to list or search for available Amazon EC2 Reserved Instances, purchase reserved capacity, manage your Reserved Instance, and sell your unused Reserved Instances.

If you use the CLI tools, first you should read and complete the instructions described in [Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX \(p. 541\)](#). The Getting Started topic walks you through setting up your environment for use with the CLI tools.

### AWS Management Console

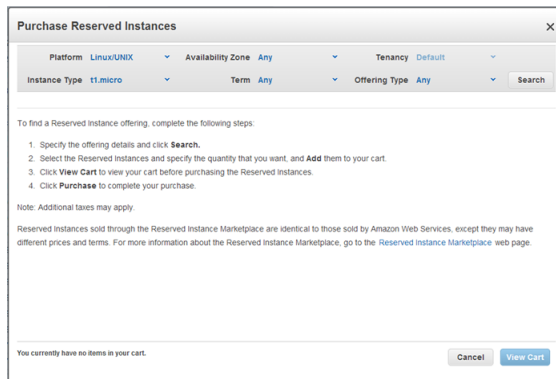
The AWS Management Console has tools specifically designed for Reserved Instances tasks. You can find them in the Amazon EC2 console. You will also find general tools that you can use to manage the instances launched when you use your Reserved Instances.

- The **Reserved Instances** page is where you work with your Reserved Instances.



RI ID	Instance Type	Zone	Expires	Term	Offering Type	Fixed Price	Usage Pr	Inst.	Platform	State
f127b2d7-62a-4763-9e09-e2d80ca...	m1.large	us-east-1b	2014-11-30 22:00	36 months	Medium Utiliza...	\$3,400.00	\$0.35	1	Windows with SQL Server S...	active
15a8e2a3-dabc-452a-8521-7980d4...	m1.small	us-east-1a	2015-08-11 16:00	36 months	Light Utilization	\$96.00	\$0.027	1	Linux/UNIX (Amazon VPC)	active
431791e3-0613-446c-a7b-37e0523...	t1.micro	us-east-1a	2015-09-11 16:00	36 months	Light Utilization	\$35.00	\$0.012	3	Linux/UNIX	active
46a0d071-7a6d-4a0b-0784-c10a9f...	m1.medium	us-east-1a	2014-09-12 16:00	12 months	Light Utilization	\$122.00	\$0.109	50	RHEL for Enterprise Linux (Am...	active
93bba65-1684-4a4d-844d-593a855...	m1.small	us-east-1a	2014-09-12 16:00	12 months	Light Utilization	\$61.00	\$0.034	5	Linux/UNIX	active
93bba65-8117-4653-823f-312a94b...	m1.small	us-east-1a	2014-09-12 16:00	12 months	Light Utilization	\$61.00	\$0.034	25	Linux/UNIX	active

- Use the **Purchase Reserved Instances** page to specify the details of the Reserved Instances you want to purchase.



Purchase Reserved Instances

Platform: Linux/UNIX | Availability Zone: Any | Tenancy: Default

Instance Type: t1.micro | Term: Any | Offering Type: Any | Search

To find a Reserved Instance offering, complete the following steps:

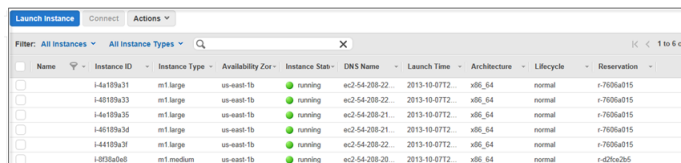
1. Specify the offering details and click **Search**.
2. Select the Reserved Instances and specify the quantity that you want, and **Add** them to your cart.
3. Click **View Cart** to view your cart before purchasing the Reserved Instances.
4. Click **Purchase** to complete your purchase.

Note: Additional taxes may apply.

Reserved Instances sold through the Reserved Instance Marketplace are identical to those sold by Amazon Web Services, except they may have different prices and terms. For more information about the Reserved Instance Marketplace, go to the [Reserved Instance Marketplace](#) web page.

You currently have no items in your cart. Cancel View Cart

- Use the **Instances** page to launch your Reserved Instances and manage them.



Name	Instance ID	Instance Type	Availability Zone	Instance Status	DNS Name	Launch Time	Architecture	Lifecycle	Reservation
i-4a189a31	i-4a189a31	m1.large	us-east-1b	running	ec2-54-209.22...	2013-10-07T2...	x86_64	normal	r-7006a015
i-4a189a23	i-4a189a23	m1.large	us-east-1b	running	ec2-54-209.22...	2013-10-07T2...	x86_64	normal	r-7006a015
i-4a189a25	i-4a189a25	m1.large	us-east-1b	running	ec2-54-209.21...	2013-10-07T2...	x86_64	normal	r-7006a015
i-4a189a3d	i-4a189a3d	m1.large	us-east-1b	running	ec2-54-209.21...	2013-10-07T2...	x86_64	normal	r-7006a015
i-4a189a3f	i-4a189a3f	m1.large	us-east-1b	running	ec2-54-209.22...	2013-10-07T2...	x86_64	normal	r-7006a015
i-031a0a68	i-031a0a68	m1.medium	us-east-1b	running	ec2-54-209.20...	2013-10-07T2...	x86_64	normal	r-df0ca2b6

### Command Line Interface Tools

To purchase Reserved Instances, or sell them in the Reserved Instance Marketplace, you can use Amazon EC2 command line interface (CLI) tools specifically designed for these tasks. To manage the instances



when your Reserved Instances are launched, use the same commands in the CLI tools that you use for any other Amazon EC2 instances.

The following table lists the commands in the CLI tools that you use specifically for Reserved Instances tasks.

Task	CLI
List Reserved Instances that you have purchased.	<a href="#">ec2-describe-reserved-instances</a>
Modify the Reserved Instances you own.	<a href="#">ec2-modify-reserved-instances</a>
View the modifications made to your Reserved Instances.	<a href="#">ec2-describe-reserved-instances-modifications</a>
View the Reserved Instances offerings that are available for purchase.	<a href="#">ec2-describe-reserved-instances-offerings</a>
Create a listing of the Reserved Instances you want to sell in the Reserved Instance Marketplace.	<a href="#">ec2-create-reserved-instances-listing</a>
View the details of your Reserved Instance listings in the Reserved Instance Marketplace.	<a href="#">ec2-describe-reserved-instances-listings</a>
Purchase a Reserved Instance.	<a href="#">ec2-purchase-reserved-instances-offering</a>
Cancel your active Reserved Instances listing in the Reserved Instance Marketplace.	<a href="#">ec2-cancel-reserved-instances-listing</a>

For information about CLI commands, see the [Amazon Elastic Compute Cloud Command Line Reference](#).

## API

To purchase Reserved Instances, you use API calls specifically designed for these tasks. To manage the instances when your Reserved Instances are launched, use the same API calls that you use for any other Amazon EC2 instances.

The following table lists the API calls you use for Reserved Instances tasks.

Task	API
List Reserved Instances that you have purchased.	<a href="#">DescribeReservedInstances</a>
Modify the Reserved Instances you own.	<a href="#">ModifyReservedInstances</a>
View the modifications made to your Reserved Instances.	<a href="#">DescribeReservedInstancesModifications</a>
View the Reserved Instances offerings that are available for purchase.	<a href="#">DescribeReservedInstancesOfferings</a>
Create a listing of the Reserved Instances you want to sell in the Reserved Instance Marketplace.	<a href="#">CreateReservedInstancesListing</a>
View the details of your Reserved Instance listings in the Reserved Instance Marketplace.	<a href="#">DescribeReservedInstancesListings</a>
Purchase a Reserved Instance.	<a href="#">PurchaseReservedInstancesOffering</a>

Task	API
Cancel your active Reserved Instances listing in the Reserved Instance Marketplace.	<code>CancelReservedInstancesListing</code>

For information about API actions, see the [Amazon Elastic Compute Cloud API Reference](#).

## Reserved Instance Fundamentals

This section discusses fundamental concepts that can help you optimize the benefits of Reserved Instances, and use and manage them effectively.

- [Choosing Reserved Instances Based on Your Usage Plans \(p. 199\)](#)—Select the pricing model that best matches how often you plan to use your instances.
- [Understanding Reserved Instance Pricing Tiers \(p. 200\)](#)—Take advantage of the Reserved Instances pricing tier discounts in a region, when the total upfront list price of your Reserved Instances in the region is \$250,000 USD or more.
- [Understanding the Pricing Benefit of Reserved Instances \(p. 208\)](#)—Learn how the Reserved Instances pricing benefits are applied.
- [Reserved Instance Marketplace \(p. 209\)](#)—Understand the flexibility provided by the Reserved Instance Marketplace to AWS customers who can sell the remainder of their Reserved Instances when their needs change, or buy Reserved Instances with less than the full standard terms from other AWS customers.

## Choosing Reserved Instances Based on Your Usage Plans

You can select a Reserved Instance fee structure based on how often you plan to use your instance. We offer three Reserved Instance types to address your projected utilization of the instance: *Heavy Utilization*, *Medium Utilization*, or *Light Utilization*. (To use these Reserved Instance types, make sure you have API version 2011-11-01 or later.)

*Heavy Utilization* Reserved Instances can be used to enable workloads that have a consistent baseline of capacity, or they can run steady-state workloads. Heavy Utilization Reserved Instances require the highest upfront commitment. However, if you plan to run your Reserved Instances around 35 percent of the time or more for a three-year term, you may be able to earn the largest savings (up to around 65 percent off of the On-Demand price if your instance utilization is 100 percent) of any of the offering types. Unlike other Reserved Instances offering types, with Heavy Utilization Reserved Instances, you pay a one-time fee, followed by a lower hourly fee for the duration of the term *regardless* of whether or not your instance is running.

*Medium Utilization* Reserved Instances are the best option if you plan to use your Reserved Instances a substantial amount of the time, but want either a lower one-time fee or the flexibility to stop paying for your instance when you shut it off. Medium Utilization is a cost-effective option when you plan to run your Reserved Instances approximately between 19 percent and 35 percent of the time over the Reserved Instance term. This option can save you up to 59 percent off of the On-Demand price. With Medium Utilization Reserved Instances, you pay a slightly higher one-time fee than with Light Utilization Reserved Instances, but you receive lower hourly usage rates when you run an instance. (This offering type is equivalent to the Reserved Instance offering available before API version 2011-11-01.)

*Light Utilization* Reserved Instances are ideal for periodic workloads that run only a couple of hours a day or a few days per week. Some use cases, such as disaster recovery, also require reserved capacity to meet potential demand without notice. Using Light Utilization Reserved Instances, you pay a one-time fee followed by a discounted hourly usage fee when your instance is running. You start saving when your instance is running approximately between 11 percent and 19 percent of the time over the Reserved

Instance term, and you can save up to 49 percent off of the On-Demand rates over the entire term of your Reserved Instance.

**Note**

With *Light Utilization* and *Medium Utilization* Reserved Instances, you pay a one-time upfront fee and then only pay the hourly price when you use the instance. With *Heavy Utilization* Reserved Instances, you pay a one-time upfront fee and commit to paying an hourly rate for every hour of the Reserved Instance's term *whether or not you use it*.

Remember that discounted usage fees for Reserved Instance purchases are tied to your specifications of type and Availability Zone of your instance. If you shut down a running EC2 instance on which you have been getting a discounted rate as a result of a Reserved Instance purchase, and the term of the Reserved Instance has not yet expired, you will continue to get the discounted rate if you launch another instance with the same specifications during the remainder of the term.

The following table summarizes the differences between the Reserved Instances offering types.

**Reserved Instance Offerings**

Offering	Upfront Cost	Usage Fee	Advantage
Heavy Utilization	Highest	Lowest hourly fee. Applied to the whole term whether or not you're using the Reserved Instance.	Lowest overall cost if you plan to utilize your Reserved Instances more than approximately 35 percent of the time over a 3-year term.
Medium Utilization	Average	Hourly usage fee charged for each hour you use the instance. We encourage you to turn off your instances when you aren't using them so you won't be charged for them.	Suitable for elastic workloads or when you expect moderate usage, approximately between 19 percent and 35 percent of the time over a 3-year term.
Light Utilization	Lowest	Hourly usage fee charged. Highest fees of all the offering types, but they apply only when you're using the Reserved Instance. We strongly encourage you to turn off your instances when you aren't using them so you won't be charged for them.	Highest overall cost if you plan to run all of the time; however it's the lowest overall cost if you anticipate you will use your Reserved Instances infrequently, approximately between 11 percent and 19 percent of the time over a 3-year term.

## Understanding Reserved Instance Pricing Tiers

To qualify for Amazon Elastic Compute Cloud (Amazon EC2) Reserved Instances pricing tier discounts in a region, the total upfront list price of your Reserved Instances in the region must be \$250,000 USD or more. When your account qualifies for a discount pricing tier, it will automatically receive discounts on upfront and usage fees for all Reserved Instance purchases that you make within that tier level from that point on.

This section introduces you to Reserved Instances pricing tiers and how to take advantage of the pricing tier discounts.

- [What Are the Reserved Instance Pricing Tiers?](#) (p. 201)
- [Current Limitations](#) (p. 201)
- [Determining Your Pricing Tier Level](#) (p. 201)
- [How Do Pricing Tier Discounts Get Applied?](#) (p. 203)

## What Are the Reserved Instance Pricing Tiers?

The following table lists the qualifications for each pricing tier and the discount that is applied when the total upfront list price of your active Reserved Instances in a region crosses into the range of that pricing tier.

Tier level	Total upfront list price of active Reserved Instances in the region	Discount applied to upfront and usage fees for Reserved Instances purchased in the tier
Tier 0	\$0 - \$250,000	Standard Reserved Instance upfront and usage fees. No discount.
Tier 1	\$250,000 - \$2,000,000	10 percent discount
Tier 2	\$2,000,000 - \$5,000,000	20 percent discount
Tier 3	Over \$5,000,000	<a href="#">Contact Us</a>

## Current Limitations

The following limitations currently apply to Reserved Instances pricing tiers:

- Amazon EC2 Reserved Instance purchases are the only purchases that will apply toward your Amazon EC2 Reserved Instance pricing tier discounts. And the Amazon EC2 Reserved Instance pricing tiers and related discounts apply only to purchases of Amazon EC2 Reserved Instances.
- Amazon EC2 Reserved Instance pricing tiers do not apply to Reserved Instances for Windows with SQL Server Standard or Windows with SQL Server Web.
- Amazon EC2 Reserved Instances purchased as part of a tiered discount cannot be sold in the Reserved Instance Marketplace. For more information about the Reserved Instance Marketplace, see [Reserved Instance Marketplace](#) (p. 209).

For a checklist that summarizes requirements for working with Reserved Instances and the Reserved Instance Marketplace, see [Requirements Checklist for Reserved Instances](#) (p. 262).

## Determining Your Pricing Tier Level

To determine which Reserved Instance pricing tier applies to you in a particular region, compare the sum of your Reserved Instances' *upfront* list prices to the total upfront list price required for the pricing tier. *List price* is the undiscounted Reserved Instance price that you see in the AWS Management Console or the AWS marketing website. Keep in mind that if the price of Reserved Instances drops after you buy Reserved Instances, that price drop might not be reflected in the undiscounted Reserved Instance price of your Reserved Instances.

## Amazon Elastic Compute Cloud User Guide Reserved Instance Fundamentals

Seller	Term	Effective Rate	Upfront Price	Hourly Rate	Availability Zone	Offering Type	Quantity Available	Desired Quantity
3rd Party	2 months	\$0.085	\$53.32	\$0.048	us-east-1a	Medium Utilization	1	1
AWS	12 months	\$0.082	\$122.00	\$0.068	us-east-1d	Light Utilization	Unlimited	1
AWS	12 months	\$0.082	\$122.00	\$0.068	us-east-1b	Light Utilization	Unlimited	1
AWS	12 months	\$0.082	\$122.00	\$0.068	us-east-1a	Light Utilization	Unlimited	1
AWS	36 months	\$0.061	\$192.00	\$0.054	us-east-1d	Light Utilization	Unlimited	1

List price is not the same as paid price. *Paid price* is the actual amount that you paid for the Reserved Instances. (The term *paid price* is the same as the term *fixed price* that you see when you use Reserved Instance tools in the AWS Management Console, command line interface (CLI), or the API.) If the Reserved Instance was purchased at a discount, the paid price will be lower than the list price. If the Reserved Instance was purchased without a discount, the paid price will equal the list price. So, if you purchased your Reserved Instances without discounts, you can use the fixed price value in the console, the CLI, or the API to determine which pricing tier your purchase falls under.

### Note

In the EC2 console, you might need to turn on the display of the **Fixed Price** column by clicking **Show/Hide** in the top right corner.

RI ID	Instance Type	Zone	Start	Expires	Term	Offering Type	Fixed Price	Instance Count	Recurring Chrg	Platform
46405c7-042...	m1 large	us-east-1b	2013-09-09 09...	2014-08-30 14...	11.7 months	Medium Utiliza...	\$0.00	1		Linux/UNIX
b847693-169...	m1 large	us-east-1a	2013-02-28 09...	2013-02-28 15...	12 months	Heavy Utilization	\$780.00	1	Hourly: \$0.054	Linux/UNIX

Assuming you purchased your Reserved Instances without discounts, you can determine the pricing tier for your account by calculating the total fixed price for all your Reserved Instances in a region. To do this using the console, calculate the sum of the amounts in the **Fixed Price** column.

RI ID	Instance Type	Zone	Start	Expires	Term	Offering Type	Fixed Price	Instance Count
46405c7-042...	m1 large	us-east-1b	2013-09-09 09...	2014-08-30 14...	11.7 months	Medium Utiliza...	\$0.00	1
b847693-169...	m1 large	us-east-1a	2013-02-28 09...	2013-02-28 15...	12 months	Heavy Utilization	\$780.00	1
b847693-169...	m1 large	us-east-1a	2013-02-28 09...	2013-02-28 15...	12 months	Heavy Utilization	\$2,030.00	1
b8c09749-0208...	m1 large	us-east-1a	2013-09-09 09...	2013-09-09 09...	11.7 months	Medium Utiliza...	\$0.00	1
1ba8e243-33f...	m1 large	us-east-1d	2013-02-28 09...	2013-03-02 16...	12 months	Medium Utiliza...	\$640.00	1
938bbc2-47f6...	t1 micro	us-east-1d	2013-02-26 12...	2013-03-02 16...	12 months	Light Utilization	\$23.00	1
938bbc2-422...	m1 large	us-east-1d	2013-02-28 09...	2013-03-02 16...	12 months	Light Utilization	\$276.00	1
945c4137-5ab...	m1 large	us-east-1d	2013-08-30 14...	2013-08-30 14...	12 months	Medium Utiliza...	\$554.00	1
b8c09749-bda...	m1 large	us-east-1d	2013-02-28 09...	2013-03-02 16...	12 months	Medium Utiliza...	\$1,670.00	1
b8c09749-c56...	m1 large	us-east-1d	2013-02-28 09...	2013-03-02 16...	12 months	Light Utilization	\$1,370.00	1
f127bd27-80fc...	m1 large	us-east-1d	2013-08-30 14...	2013-09-09 09...	12 months	Medium Utiliza...	\$0.00	1

Using the Amazon EC2 CLI or the API, you can determine the pricing tier of your account by calculating the sum of the `FixedPrice` (CLI) or `fixedPrice` (API) values returned by the `ec2-describe-reserved-instances` command or the `DescribeReservedInstances` action, respectively.

Your `ec2-describe-reserved-instances` command should look like the following example:

```
PROMPT> ec2-describe-reserved-instances --headers
```

Amazon EC2 returns output similar to the following example:

```
PROMPT> ec2-describe-reserved-instances
Type ReservedInstancesId AvailabilityZone InstanceType ProductDescription Dura
tion FixedPrice UsagePrice InstanceCount Start State Currency InstanceTenancy
OfferingType
RESERVEDINSTANCES f127bd27-f0e9-43bb-89f5-1b8c030bc8f9 us-east-1b m1.small
```

```
Linux/UNIX 1y 227.5 0.03 1 2011-11-28T16:17:12+0000 default active USD Medium
Utilization
RESERVEDINSTANCES f127bd27-f62e-4763-9e09-ee24f8ccef5d us-east-1b m1.large
Windows with SQL Server 3y 3400.0 0.35 1 2011-12-02T06:13:25+0000 default
active USD Medium Utilization
RESERVEDINSTANCES 1ba8e2e3-e8b2-4637-a124-ec9c11495ac9 us-east-1d m1.small
Linux/UNIX (Amazon VPC) 1y 280.0 0.035 1 2011-12-02T06:05:28+0000 dedicated
active USD Medium Utilization
RESERVEDINSTANCES 46a408c7-89fe-4b02-bb5e-ecb9bbc510eb us-east-1d m2.xlarge
Linux/UNIX 1y 1000.0 0.5 1 2011-12-02T06:05:27+0000 default active USD Lower
Utilization
RESERVEDINSTANCES af9f760e-96ae-4d12-8abe-8e8e1e7a2bbdb us-east-1d t1.micro
Linux/UNIX 1y 54.0 0.0070 10 2011-12-02T06:12:09+0000 default active USD Medium
Utilization
RESERVEDINSTANCES 46a408c7-ae07-4611-9d1c-f6d3a947f8d3 us-east-1a m1.small
Linux/UNIX (Amazon VPC) 1y 227.5 0.03 1 2011-11-08T18:00:02+0000 default active
USD Medium Utilization
RESERVEDINSTANCES bbcd9749-1211-4134-a7e7-0cdfec1caca5 us-east-1a t1.micro
Linux/UNIX 1y 54.0 0.0070 1 2011-11-08T18:03:20+0000 default active USD Medium
Utilization
RESERVEDINSTANCES d16f7a91-556f-4db5-afc9-4dd0673334c6 us-east-1a m1.large
Windows with SQL Server 3y 3400.0 0.35 1 2011-12-02T06:22:03+0000 default
active USD Medium Utilization
REQUEST ID d9121e8b-e7c1-49f2-88cd-b478c999751
```

For an example using the API action, see [DescribeReservedInstances](#).

## How Do Pricing Tier Discounts Get Applied?

If a single purchase of Reserved Instances in a region takes you over the threshold of a discount tier, then the portion of that purchase that is above the price threshold will be charged at the discounted rate. Refer to the table in the previous section for tiers and discount price points.

### Note

Amazon EC2 Reserved Instance purchases are the only purchases that determine your Amazon EC2 Reserved Instance pricing tiers, and the Amazon EC2 Reserved Instance pricing tiers apply only to Amazon EC2 Reserved Instance purchases.

Here's an example that shows you the discount effect of a purchase of Reserved Instances that crosses the discount tier threshold. Let's assume you currently have \$200,000 worth of active Reserved Instances in the us-east-1 Region. You purchase 75 Reserved Instances at the list price of \$1,000 each. That's a total of \$75,000, which brings the total amount you have paid for active Reserved Instances to \$275,000. Since the discount pricing threshold is \$250,000, the first \$50,000 of your new purchase would not receive a discount. The remaining \$25,000, which exceeds the discount pricing threshold, would be discounted by 10 percent (\$2,500). This means you will only be charged \$22,500 for the remainder of your purchase (25 instances), and you will be charged discounted usage rates for those 25 Reserved Instances. (However, keep in mind that your total upfront— *undiscounted* —list price is still \$275,000.)

After the total upfront list price of your active Reserved Instances in a region crosses into the discount pricing tier, any future purchase of Reserved Instances in that region will be charged at a discounted rate. As long as your total list price stays above the price point for the discount tier, all future purchases of Reserved Instances in that region will be discounted. If your total list price falls below that price point for the discount tier—for example, if some of your Reserved Instances expire—succeeding purchases of Reserved Instances in the region will not be discounted. However, you will continue to get the discount against the already purchased Reserved Instances that originally went within the discount pricing tier.



If your account is part of a consolidated billing account, you can benefit from the Reserved Instance pricing tiers. A consolidated billing account aggregates into a single list price the list prices of all of the active Reserved Instances accounts in a region that are part of the consolidated billing account. When the total list price of active Reserved Instances for the consolidated billing accounts reaches the discounted tier level, any Reserved Instances purchased after this point by any member of the consolidated account will be charged at the discounted rate (as long as the total list price for that consolidated account stays above the discount tier price point).

Here's how Reserved Instance purchases work with consolidated billing: Let's assume that two accounts—A and B—are part of a consolidated billing account. All the active Reserved Instances in the consolidated billing account are in one region. Account A has Reserved Instances worth \$135,000; Account B has Reserved Instances worth \$115,000. The total upfront cost of the consolidated bill of accounts A and B is \$250,000. Remember, \$250,000 is the discount pricing threshold. This means that when either or both of the A and B accounts purchase additional Reserved Instances, the cost of the new purchases will be discounted by 10 percent. So, when account B purchases Reserved Instances at a list price of \$15,000, the consolidated account will only be charged \$13,500 for the new Reserved Instances (\$15,000 minus the 10 percent discount of \$1,500 equals \$13,500), and account B will be charged discounted usage rates for those new Reserved Instances.

For more information about how the benefits of Reserved Instances apply to consolidated billing accounts, see [Reserved Instances and Consolidated Billing \(p. 209\)](#).

### Purchasing at a Discount Tier Price

When you purchase Reserved Instances, Amazon EC2 will automatically apply any discounts to the part of your Reserved Instance purchase that falls within a discount tier. You don't need to do anything differently, and you can purchase using any of the Amazon EC2 tools.

- AWS Management Console: Click the **Purchase Reserved Instances** button on the **Reserved Instances** page of the [Amazon EC2 console](#).
- Amazon EC2 CLI: Use the `ec2-purchase-reserved-instances-offering` command.
- Amazon EC2 API: Call the `PurchaseReservedInstancesOffering` action.

If your purchase crosses into a discounted pricing tier, the console, the `ec2-describe-reserved-instances` command, or the `DescribeReservedInstances` action will show multiple entries for that purchase. You will see an entry for that part of the purchase that will be charged the regular Reserved Instance price, and another entry or entries for that part of the purchase that will be charged the applicable discounted rate.

Consequently, the Reserved Instance ID returned by your purchase CLI command or API action will be different from the actual ID of the new Reserved Instances. For more information, see [Reserved Instance IDs \(p. 205\)](#).

For example, let's say that your account has a list price total of \$245,000 in Reserved Instances in a region. You purchase more Reserved Instances at a total list price of \$10,000. Remember that the threshold for the 10 percent discount tier is \$250,000, so \$5,000 of your purchase crosses into the discount tier. When the transaction is complete, the list you'll get from the console, the CLI, or the API will show two reservations—a reservation for \$5,000 at the undiscounted rate, and another for \$4,500, which is the discounted price. The discount is 10 percent of \$5,000 or \$500.

The pricing discount applies to future purchases after the total list price cost of your active Reserved Instances reaches the discounted pricing tier. However, if some of your Reserved Instances expire and the total list price for your active Reserved Instances falls below the discounted pricing tier level, then the next purchase you make will be at the retail rates for Reserved Instances. To illustrate, let's use the previous example: You currently have a list price total of \$254,500 in Reserved Instances. \$250,000 is at the undiscounted rate, and \$4,500 is at the discounted rate. In two months, Reserved Instances worth \$20,000 expire, bringing your total down to \$234,500, which is below the threshold for the 10 percent discount tier. When you subsequently purchase \$15,000 in Reserved Instances, you will be charged the

retail rate. However, the \$4,500 in Reserved Instances that you purchased earlier at the discounted rate will continue to be charged at the discounted rate.

The main point to keep in mind is that *list price* is the undiscounted price of the Reserved Instance at the time of purchase, while the *fixed price* value in the console, the CLI, and the API is the *paid price* of all Reserved Instances purchased. The discount tier is based on *list price*.

Remember that when you purchase Reserved Instances, one of four possible scenarios will occur:

- Your purchase of Reserved Instances within a region is still below the discount threshold. This means that you don't get a discount.
- Your purchase of Reserved Instances within a region crosses the threshold of the first discount tier. This means that the newly purchased Reserved Instances will go into the pending state while the purchase is processed. For this purchase, the Reserved Instances service will purchase Reserved Instances for you at two different rates: An amount at the undiscounted rate, and an amount at the discounted rate. It is important to understand that the Reserved Instance IDs you get back from the CLI or API when you use the purchase command will be different from the new Reserved Instance IDs that will actually be created at the completion of the purchase. These IDs are returned when you use the describe command. For an explanation of this difference, see the next section, [Reserved Instance IDs \(p. 205\)](#).
- Your entire purchase of Reserved Instances within a region is completely within one discount tier. This means that the newly purchased Reserved Instances will go into the pending state while the purchase is processed. For this purchase, the Reserved Instances service will purchase Reserved Instances at the appropriate discount level for you. It is important to understand that, as with the previous scenario, the Reserved Instance IDs you get back will be different from the new Reserved Instance IDs that will actually be created at the completion of the purchase. These IDs are returned when you use the describe command. For an explanation of this difference, see the next section, [Reserved Instance IDs \(p. 205\)](#).
- Your purchase of Reserved Instances within a region crosses from a lower discount tier to a higher discount tier. As with the previous scenario, this means that the newly purchased Reserved Instances will go into the pending state while the purchase is processed. For this purchase, the Reserved Instances service will purchase Reserved Instances for you at two different rates: An amount at the first or lower discounted rate, and an amount at the higher discounted rate. It is important to understand that the Reserved Instance IDs you get back will be different from the new Reserved Instance IDs that will actually be created at the completion of the purchase. For an explanation of this difference, see the next section, [Reserved Instance IDs \(p. 205\)](#).

## Reserved Instance IDs

When your total purchase crosses one or more Reserved Instance discount pricing tiers, the Reserved Instance IDs returned by your purchase command can be different from the Reserved Instance IDs returned by the describe command that you call after the purchase is complete. What happens is that the Reserved Instance service generates several Reserved Instance IDs because your purchase crossed from an undiscounted tier to a discounted tier, or from one discounted tier to another. There will be a Reserved Instance ID for each set of Reserved Instances in a tier.

Using the CLI, here's an example output showing the `ec2-describe-reserved-instances` command not recognizing the Reserved Instance ID `1ba8e2e3-edf1-43c3-b587-7742bc77b9ba`, which was returned by `ec2-purchase-reserved-instances-offering`.

```
$ ec2-describe-reserved-instances -H --region sa-east-1 1ba8e2e3-edf1-43c3-b587-7742bc77b9ba
Type ReservedInstancesId AvailabilityZone InstanceType ProductDescription Duration FixedPrice UsagePrice InstanceCount Start State Currency InstanceTenancy OfferingType
```



In this example, the Reserved Instance ID generated by the purchase command is like an intermediate ID that is used while the purchase is being processed. Because your purchase crossed from the undiscounted tier (tier 0) to the first discounted tier (tier 1), the Reserved Instance service actually generates several Reserved Instance IDs. After the purchase is complete, when you use the describe command, the service returns Reserved Instance ID `bbcd9749-05f0-4ada-96c8-812f5f0ab9b3`, the ID for the Reserved Instances that you purchased at the undiscounted list price of \$20,000 each. The service also returns Reserved Instance IDs `1ba8e2e3-346e-4e5b-a2e2-b559243f2325` and `af9f760e-868c-48f4-87e2-44576dbf05ef`, the IDs for the Reserved Instances that you purchased at the 10 percent discount rate (\$20,000 minus \$2,000).

Your entire purchase would look like the following example:

```
$ ec2-describe-reserved-instances -H --region sa-east-1 bbcd9749-05f0-4ada-96c8-812f5f0ab9b3
Type ReservedInstancesId AvailabilityZone InstanceType ProductDescription Duration FixedPrice UsagePrice InstanceCount Start State Currency InstanceTenancy OfferingType
RESERVEDINSTANCES bbcd9749-05f0-4ada-96c8-812f5f0ab9b3 sa-east-1a t1.micro Linux/UNIX 3y 20000.0 0.0090 2 2012-03-02T23:20:16+0000 default payment-pending USD Medium Utilization
$ ec2-describe-reserved-instances -H --region sa-east-1 1ba8e2e3-346e-4e5b-a2e2-b559243f2325
Type ReservedInstancesId AvailabilityZone InstanceType ProductDescription Duration FixedPrice UsagePrice InstanceCount Start State Currency InstanceTenancy OfferingType
RESERVEDINSTANCES 1ba8e2e3-346e-4e5b-a2e2-b559243f2325 sa-east-1a t1.micro Linux/UNIX 3y 18000.0 0.0080 3 2012-03-02T23:20:17+0000 default payment-pending USD Medium Utilization
$ ec2-describe-reserved-instances -H --region sa-east-1 af9f760e-868c-48f4-87e2-44576dbf05ef
Type ReservedInstancesId AvailabilityZone InstanceType ProductDescription Duration FixedPrice UsagePrice InstanceCount Start State Currency InstanceTenancy OfferingType
RESERVEDINSTANCES af9f760e-868c-48f4-87e2-44576dbf05ef sa-east-1a t1.micro Linux/UNIX 3y 18000.0 0.0080 5 2012-03-02T23:20:18+0000 default payment-pending USD Medium Utilization
```

### Scenario Showing Purchases that Cross Pricing Tiers

Let's walk through an example scenario in which your purchases of Reserved Instances cross the various pricing tiers.

Two months ago, you purchased 100 Reserved Instances in the us-east-1a region at \$2000 each. That purchase totaled \$200,000. The list price for this purchase is \$2000. The amount you paid was \$2000 per Reserved Instance, so it is the paid price, and the paid price is the value that will be shown under fixed price. Your purchases are still within the first, undiscounted tier. (For information about tier thresholds, see the [What Are the Reserved Instance Pricing Tiers?](#) (p. 201) table.)

The following table illustrates this example.

Purchase Number	List Price	Amount Paid	Fixed Price	Total RI Purchased	Total List Price Cost	Total Paid Amount
Purchase 1	\$2000	\$2000	\$2000	100	\$200,000	\$200,000

Later you want to purchase more Reserved Instances. Let's say that AWS lowered prices and the same type of Reserved Instance now is available at \$1000 each. You purchase 75 of these Reserved Instances.

The list price for this purchase is \$1000. The purchase of 75 Reserved Instances at \$1000 each totals \$75,000. This raises the total cost of your active Reserved Instances to \$275,000. The threshold for the discount tier is \$250,000. This purchase crosses into the first discount tier, tier 1 in the [What Are the Reserved Instance Pricing Tiers? \(p. 201\)](#) table.

In this discount tier, you get a 10 percent discount on all your purchases in the same region of Reserved Instances above \$250,000. So, you will pay the new list price of \$1000 each for the first 50 Reserved Instances (total amount paid of \$50,000). And you will pay \$900 each—the \$1000 list price, minus the 10 percent discount—for the remaining 25 Reserved Instances (total amount paid of \$22,500). Your *fixed price* and the *amount paid* for the discounted Reserved Instances will both show \$900.

The following table illustrates this example.

Purchase Number	List Price	Amount Paid	Fixed Price	Total RI Purchased	Total List Price Cost	Total Paid Amount
Purchase 1	\$2000	\$2000	\$2000	100	\$200,000	\$200,000
Purchase 2	\$1000	\$1000	\$1000	50	\$50,000	\$50,000
	\$1000	\$900	\$900	25	\$25,000	\$22,500
Totals					\$275,000	\$272,500

Six months later, let's assume that your business has experienced tremendous growth, and you need to purchase 1800 additional Reserved Instances in the same region. At a list price of \$1000, this purchase will total \$1,800,000. When you add this new purchase to your previous purchases of already active Reserved Instances in the amount of \$272,500, your new total will be \$2,072,500. This new total crosses the threshold for the next discount tier (tier 2 in the [What Are the Reserved Instance Pricing Tiers? \(p. 201\)](#) table). In this tier, discounts of 20 percent apply to purchases of \$2,000,000 and above.

Your new purchase will be charged two different discount rates: The 1725 Reserved Instances that fall within tier 1 will be discounted at 10 percent. The remaining 75 Reserved Instances that put the total list price cost above \$2,000,000, and thus are in tier 2, will be discounted at 20 percent.

Purchase Number	List Price	Amount Paid	Fixed Price	Total RI Purchased	Total List Price Cost	Total Paid Amount
Purchase 1	\$2000	\$2000	\$2000	100	\$200,000	\$200,000
Purchase 2	\$1000	\$1000	\$1000	50	\$50,000	\$50,000
	\$1000	\$900	\$900	25	\$25,000	\$22,500
Purchase 3	\$1000	\$900	\$900	1725	\$1,725,000	\$1,552,500
	\$1000	\$800	\$800	75	\$75,000	\$60,000
				Total	\$2,075,000	\$1,885,000

### Reading Your Bill

A bill for a Reserved Instances purchase that qualifies for a discount shows the split purchase if your purchase crosses a pricing tier. Your bill will reflect a breakdown of costs similar to the tiered price scenario discussed in the previous section. In the previous example, you save \$390,000 (the difference between the total list price cost and the total paid amount) using Reserved Instances pricing tiers.

## Understanding the Pricing Benefit of Reserved Instances

When you purchase Reserved Instances, you get two benefits: a capacity reservation for a number of EC2 instances you will launch at some future time, and a discounted hourly fee. The rest of the experience of launching and working with Reserved Instances is the same as working with On-Demand EC2 instances.

This section discusses how Amazon EC2 applies the pricing benefit of Reserved Instances.

### Applying the Pricing Benefit of Reserved Instances

With Reserved Instances, you pay an upfront fee for the capacity reservation on available Amazon EC2 instances based on the specifications (such as product platform, instance type, Availability Zone, etc.) that you defined during the purchase. After you purchase Reserved Instances, you cannot change these specifications. For example, if you purchased a c1.medium instance, you cannot change the capacity reservation to a c1.xlarge instance. (However, you can sell your unused Reserved Instances in the Reserved Instance Marketplace.)

In addition to the capacity reservation, you also get a discounted rate on the hourly fee for running On-Demand EC2 instances that are associated with the same account that purchased the Reserved Instances. For the discount to apply, the On-Demand instances must match the specifications for the Reserved Instances.

For example, let's say user A is running the following ten On-Demand EC2 instances:

- (4) m1.small instances in Availability Zone us-east-1a
- (4) c1.medium instances in Availability Zone us-east-1b
- (2) c1.xlarge instances in Availability Zone us-east-1b

Then user A purchases the following six Medium Utilization Reserved Instances:

- (2) m1.small instances in Availability Zone us-east-1a
- (3) c1.medium instances in Availability Zone us-east-1a
- (1) c1.xlarge instance in Availability Zone us-east-1b

Then user A purchases the following six Medium Utilization Reserved Instances:

When he purchases the Reserved Instances, user A pays an upfront fee for the capacity reservation so he can launch the six instances to his specifications when he needs them. In addition, he gets a discount on the hourly usage fees for the equivalent of six instances each month. Since he already has instances running when he purchases the Reserved Instances, Amazon EC2 will automatically apply the discounted hourly rates to the already running On-Demand instances that match the specifications for the Reserved Instances he purchased.

This is what happens:

- Amazon EC2 applies the discounted usage fee rate for two m1.small Reserved Instances that user A purchased to two of the four running m1.small Amazon EC2 instances in Availability Zone us-east-1a.

The other two EC2 instances in Availability Zone us-east-1a will be charged at the current On-Demand rate.

- Amazon EC2 doesn't apply discounted rates from the three c1.medium Reserved Instances that user A purchased because these c1.medium Reserved Instances are specified to run in a different Availability Zone from the zone currently running c1.medium Amazon EC2 instances.

The four running c1.medium Amazon EC2 instances will be charged at the current On-Demand rate.

If user A launches a c1.medium EC2 instance in Availability Zone us-east-1a, then Amazon EC2 will apply the Reserved Instance discounted usage fee rate to that instance.

- Amazon EC2 applies the discounted usage fee rate for one c1.xlarge Reserved Instance that user A purchased to one of the two running c1.xlarge Amazon EC2 instances in Availability Zone us-east-1b.

The other c1.xlarge EC2 instance in Availability Zone us-east-1b will be charged at the current On-Demand rate.

In this example scenario, by purchasing the six Reserved Instances, user A saves on the hourly fee charged against two m1.small and one c1.xlarge On-Demand EC2 instances he had already running. At the same time, he is assured of the capacity to run the six Reserved Instances when he needs them.

## Reserved Instances and Consolidated Billing

The pricing benefits of Reserved Instances are shared when the purchasing account is part of a set of accounts billed under one consolidated billing (CB) payer account. Consolidated billing allows you to pay all of your charges using one account, the CB payer account. The amount of hourly usage for each month across all sub-accounts is also aggregated in the CB payer account. This billing is typically useful for companies in which there are different functional teams or groups. For more information on consolidated billing, see *Consolidated Billing* in [About AWS Account Billing](#).

For Reserved Instances, the amount of usage across all linked accounts is aggregated in the CB payer account, by the hour for each month. Then the normal Reserved Instance logic is applied to calculate the bill.

For example, your account is part of a consolidated billing account, and using your account you purchase Reserved Instances. The upfront cost of the Reserved Instances is paid by the CB payer account, and the discount is spread across the sub-accounts. The allocation of the total cost is determined by the ratio of each sub-account's usage divided by the total usage of the CB payer account. However, the capacity reservation remains with the sub-account that purchased the reservation—in this example, your account. Keep in mind that capacity reservation only applies to the product platform, instance type, and Availability Zone specified in the purchase.

For more information about how the discounts of the Reserved Instance pricing tiers apply to consolidated billing accounts, see [Understanding Reserved Instance Pricing Tiers \(p. 200\)](#).

## Reserved Instance Marketplace

The Reserved Instance Marketplace is an online marketplace that provides AWS customers the flexibility to sell their unused Amazon Elastic Compute Cloud (Amazon EC2) Reserved Instances to other businesses and organizations. Customers can also browse the Reserved Instance Marketplace to find a wide selection of Reserved Instance term lengths and pricing options sold by other AWS customers (listed as *3rd-Party* sellers). The Reserved Instance Marketplace gives customers the flexibility to sell the remainder of their Reserved Instances as their needs change. For example, a customer may want to move instances to a new AWS region, change to a new instance type, or sell capacity for projects that end before the term expires. Amazon EC2 Instances purchased on the Reserved Instance Marketplace offer the same capacity reservations as Reserved Instances purchased directly from AWS.

### Note

Some restrictions—such as what is required to become a seller and when you can sell your reserved capacity—apply. For information about restrictions and requirements for Reserved Instances and the Reserved Instance Marketplace, see [Requirements Checklist for Reserved Instances \(p. 262\)](#).

For a buyer, there are a few differences between these Reserved Instances and Reserved Instances purchased directly from Amazon Web Services (AWS):

- **Term.** The Reserved Instances that you purchase from third-party sellers on the Reserved Instance Marketplace will have less than a full standard term remaining. Full standard terms for Reserved Instances available from AWS run for one year or three years.
- **Upfront price.** Third-party Reserved Instances can be sold at different upfront prices. The usage or recurring fees will remain the same as the fees set when the Reserved Instances were originally purchased from AWS.
- **Tiered discounts.** Amazon EC2 Reserved Instances purchased at a reduced cost because a discount tier threshold had been crossed cannot be sold in the Reserved Instance Marketplace. For information about the Reserved Instance pricing tiers, see [Understanding Reserved Instance Pricing Tiers \(p. 200\)](#)

As a seller, you can choose to list some or all of your Reserved Instances, and you can choose the upfront price you want to receive. Your Reserved Instances are then listed in the Reserved Instance Marketplace and are available for purchase. You will be charged a service fee of 12 percent of the total upfront price for each Reserved Instance you sell in the Reserved Instance Marketplace. You can use your Reserved Instance until it's sold. When you sell, you are giving up the capacity reservation and the accompanying discounted fees. This means that you can continue to use your instance after you have sold your capacity reservation. You will just have to pay the On-Demand price for the instance, starting from the time that the reserved capacity on the instance was sold.

**Note**

Only Amazon EC2 Reserved Instances can be sold in the Reserved Instance Marketplace. Other AWS Reserved Instances, such as Amazon Relational Database Service (Amazon RDS) and Amazon ElastiCache Reserved Instances cannot be sold on the Reserved Instance Marketplace.

For information about Reserved Instances, see [Reserved Instances \(p. 193\)](#).

## Buyer Overview

The Reserved Instance Marketplace is useful if you want to buy Reserved Instances with terms that are different from the terms offered by AWS. You can search the marketplace for Reserved Instances with configurations that address your specific business needs.

### Quick Start: Buying in the Reserved Instance Marketplace Video

The following video shows you how to buy Reserved Instances in the Reserved Instance Marketplace using the AWS Management Console. [Getting Started Buying Reserved Instances in the Reserved Instance Marketplace](#)

### Requirements

To purchase Reserved Instances in the Reserved Instance Marketplace, you must have a valid Amazon Web Services (AWS) account. For information on setting up an AWS account, see [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#).

If you have purchased Amazon EC2 Reserved Instances in the past, you will find that the process and tools for purchasing Reserved Instances in the Reserved Instance Marketplace are very familiar.

### Steps to buying Reserved Instances

The steps to purchasing Reserved Instances—whether they are standard AWS Reserved Instances or instances in the Reserved Instance Marketplace—are the same.

1. Specify the details of the Reserved Instance you want to purchase.
2. Select the Reserved Instance you want from the list identified by the Reserved Instance Marketplace based on your specifications.
3. Confirm your choice and purchase.

For more information, see [Purchasing Reserved Instances \(p. 214\)](#).

## Seller Overview

The Reserved Instance Marketplace will be useful to you if you own Reserved Instances and want to sell the remainder of the term of your reserved capacity, or if your business is looking for Reserved Instances with configurations that are different from the ones you currently own. The marketplace gives you the opportunity to sell your instances to businesses with needs for short-term workloads and they want to purchase Reserved Instances outside the standard one-year and three-year term lengths.

Listing on the Reserved Instance Marketplace provides you the flexibility to move to new Reserved Instance configurations when your business needs change. For example, say you currently own several three-year, m1.xlarge Reserved Instances in the EU (Ireland) Region. This year, your customer base expanded to Asia, so you need an m1.large Reserved Instance that you can use in the Asia Pacific (Tokyo) Region. You can use the Reserved Instance Marketplace to sell the remainder of the term on some of your m1.xlarge Reserved Instances purchased in the EU (Ireland) Region and purchase capacity in the Asia Pacific (Tokyo) Region.

### Note

AWS will charge you a service fee of 12 percent of the total upfront price of each Reserved Instance you sell in the marketplace.

### Quick Start: Selling in the Reserved Instance Marketplace Video

The following video shows you how to sell Reserved Instances in the Reserved Instance Marketplace using the AWS Management Console. This video includes instructions on registering as a seller and listing your instances. [Getting Started Selling Reserved Instances in the Reserved Instance Marketplace](#)

### Requirements

- **Register as a seller.** Any US legal entity or non-US legal entity can sell in the Reserved Instance Marketplace by first registering as a seller. For information, see [Registering as a Seller \(p. 237\)](#).
- **Complete the tax registration.** Sellers who have 200 or more transactions or who plan to sell \$20,000 or more in Reserved Instances will have to provide additional information about their business for tax reasons. For information, see [Tax Information \(p. 239\)](#).
- **Provide a US bank.** AWS must have your bank information in order to disburse funds collected when you sell your Reserved Instance. The bank you specify must have a US address. For more information, see [Your Bank \(p. 238\)](#).

### Steps to selling Reserved Instances

After you have registered as a seller and have provided all required information, you are ready to sell your Reserved Instances in the Reserved Instance Marketplace.

1. Select the Reserved Instances you want to sell.
2. Choose the price at which you want your Reserved Instances to sell.
3. List your Reserved Instances.

For more information, see [Selling in the Reserved Instance Marketplace \(p. 236\)](#).

## What Do You Want to Do Next?

- Learn about:
  - [Getting Started with Reserved Instances \(p. 194\)](#)
  - [Understanding the Pricing Benefit of Reserved Instances \(p. 208\)](#)
  - [Understanding Reserved Instance Pricing Tiers \(p. 200\)](#)



- Start:
  - [Obtaining Information About Your Reserved Instances \(p. 222\)](#)
  - [Purchasing Reserved Instances \(p. 214\)](#)
  - [Registering as a Seller \(p. 237\)](#)
  - [Listing Your Reserved Instance \(p. 242\)](#)

## Buying Reserved Instances

You can purchase Amazon EC2 Reserved Instances with one- or three-year terms from Amazon Web Services (AWS) or you can purchase EC2 Reserved Instances from third-party sellers who own EC2 Reserved Instances that they no longer need. Reserved Instances bought from third parties and sold through the Reserved Instance Marketplace work like Reserved Instances purchased from AWS, and the purchase process is the same. The only differences are that Reserved Instances purchased from third parties will have less than a full term remaining, and they can be sold at different upfront prices.

For a buyer, the Reserved Instance Marketplace provides increased selection and flexibility by allowing you to search for Reserved Instances that most closely match your preferred combination of instance type, region, and duration.

It is important to note that once you have purchased a Reserved Instance (either from a third-party seller in the Reserved Instance Marketplace or from AWS), you cannot cancel your purchase. However, you can modify your Reserved Instances and you can sell them if your needs change. For information about modifying your Reserved Instances, see [Modifying Your Reserved Instances \(p. 227\)](#). For information about selling your Reserved Instances in the Reserved Instance Marketplace, see [Selling in the Reserved Instance Marketplace \(p. 236\)](#).

This buyer's guide contains the following sections:

- [How Buying Works \(p. 212\)](#)—Provides an overview of what you need to get started buying in the Reserved Instance Marketplace.
- [Becoming a Buyer \(p. 213\)](#)—Discusses what you need to do to become a buyer in the Reserved Instance Marketplace, the information you have to disclose, and the reasons why certain information is necessary.
- [Purchasing Reserved Instances \(p. 214\)](#)—Walks you through the purchase process, which involves tasks that you likely will be repeating if you decide to use the Reserved Instance Marketplace.
- [Reading Your Statement \(Invoice\) \(p. 221\)](#)—Helps you understand your bill.

For general information about the Reserved Instance Marketplace, see [Reserved Instance Marketplace \(p. 209\)](#). For information about selling Reserved Instances in the Reserved Instance Marketplace, see [Selling in the Reserved Instance Marketplace \(p. 236\)](#). For basic information about Reserved Instances, see [Reserved Instances \(p. 193\)](#).

For product pricing information, see the following pages:

- [Amazon EC2 Reserved Instance Pricing](#)
- [Amazon EC2 On-Demand Instances Pricing](#)
- [AWS Service Pricing Overview](#)

## How Buying Works

Before you start using the Reserved Instance Marketplace, you must first create an account with AWS. To make a purchase in the Reserved Instance Marketplace, you specify the details of the Reserved Instances that you want to purchase in the AWS Management Console search wizard. You then are

presented with options that match your request. Any Reserved Instance you select must be purchased just as it is listed—that is, you cannot change the term, the instance type, and so on.

After you select Reserved Instances to purchase, AWS will provide you a quote on the total cost of your selections. When you decide to proceed with the purchase, AWS will automatically place a *limit price* on the purchase price, so the total cost of your Reserved Instances will not exceed the amount you were quoted. If the price rises for any reason, AWS will automatically return you to the previous screen and let you know that your purchase did not complete because the price had changed. In addition, if at the time of purchase, there are offerings similar to your choice but at a lower price, AWS will sell you the offerings at the lower price instead of your higher-priced choice.

The Reserved Instance pricing tier discounts only apply to purchases made from AWS. These discounts do not apply to purchases of third-party Reserved Instances. For example, if you purchase \$250,000 worth of Reserved Instances from a third party in the Reserved Instance Marketplace, and then you purchase another set of third-party Reserved Instances that puts your total list price above the first discount pricing tier threshold, you will not get the 10 percent discount associated with that tier. However, if after your purchases from the marketplace cross a discount pricing threshold, you then purchase a Reserved Instance from AWS, your purchase from AWS will be discounted. For more information, see [Understanding Reserved Instance Pricing Tiers \(p. 200\)](#).

You can determine if your Reserved Instance transaction has completed by looking at the **Reserved Instances** page in the Amazon EC2 console, or the results of the `ec2-describe-reserved-instances` command, or the `DescribeReservedInstances` action. If the purchase is successful, your Reserved Instance will transition from the *pending-payment* state to the *active* state.

When you buy Reserved Instances and your payment fails, the console, the `ec2-describe-reserved-instances` command, and the `DescribeReservedInstances` action will display this failed transaction by showing the Reserved Instance that you attempted to purchase to be *payment-failed*, changing from the previous *payment-pending* state.

## Becoming a Buyer

Becoming a buyer is simple and easy. If you already have an Amazon Web Services (AWS) account, you are ready to start purchasing. For more information, see the following sections:

- [Understanding the Information a Buyer Discloses \(p. 213\)](#)
- [Purchasing Reserved Instances \(p. 214\)](#)

If you don't have an AWS account, you first have to sign up and create an account with AWS. For more information, see [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#). If you are new to Reserved Instances, see [Getting Started with Reserved Instances \(p. 194\)](#).

## Understanding the Information a Buyer Discloses

Some basic information about the buyer will be shared with the seller. If you are the buyer, your ZIP code and country information will be provided to the seller in the disbursement report. This information will enable sellers to calculate any necessary transaction taxes that they have to remit to the government (such as sales tax or value-added tax). In rare circumstances, AWS might have to provide the seller with your email address, so that the seller can contact you regarding questions related to the sale (for example, tax questions).

For similar reasons, AWS will share the legal entity name of the seller on the buyer's purchase invoice. In addition, if you need additional information about the seller for tax or related reasons, you can call AWS Customer Service.



## Next Steps

After you have signed up with AWS, you can begin buying Reserved Instances in the Reserved Instance Marketplace.

- To find the Reserved Instances that address your specific business needs, see [Purchasing Reserved Instances \(p. 214\)](#).
- To understand your invoice, see [Reading Your Statement \(Invoice\) \(p. 221\)](#).
- To sell your unused Reserved Instances, see [Selling in the Reserved Instance Marketplace \(p. 236\)](#).

For information about Reserved Instances, see [Reserved Instances \(p. 193\)](#).


## Purchasing Reserved Instances

The procedure for buying Amazon EC2 Reserved Instances from third parties in the Reserved Instance Marketplace is essentially the same as the procedure for purchasing Reserved Instances from Amazon Web Services (AWS). You can purchase Reserved Instances in the Reserved Instance Marketplace using the AWS Management Console, the EC2 command line interface (CLI) tools, or the EC2 API.

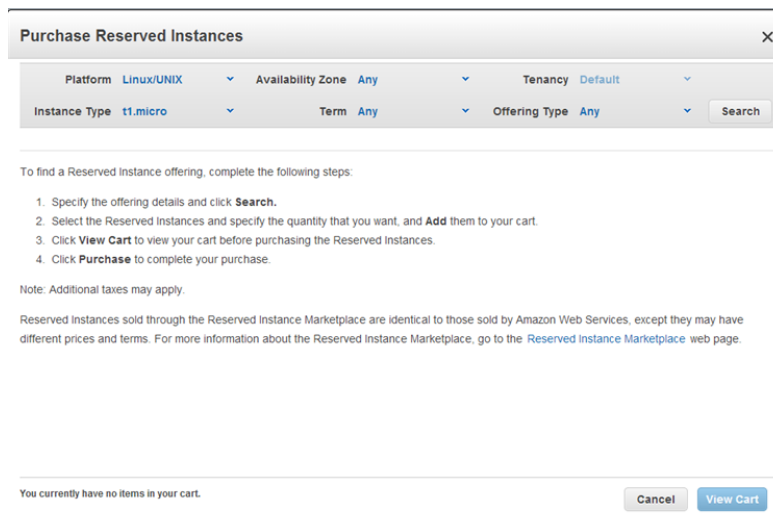
### AWS Management Console

#### To find and purchase a Reserved Instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Reserved Instances** in the **Navigation** pane.
3. In the **Reserved Instances** page, click **Purchase Reserved Instances**.



4. In the **Purchase Reserved Instances** page, specify the details of the Reserved Instances you want to purchase (use familiar filters like **Platform**, **Instance Type**, **Availability Zone**, **Term**, and **Tenancy**), and click **Search**.



## Amazon Elastic Compute Cloud User Guide Buying Reserved Instances

The **Purchase Reserved Instances** wizard will display a list of Reserved Instances that meet your search criteria.

5. Select the Reserved Instances that you want, enter the quantity that you want to purchase, and click **Add to Cart**. You can continue to select more Reserved Instances and add them to your cart.

The **Seller** column indicates whether the seller is a **3rd Party** seller or **AWS**. Notice that the **Term** column gives you non-standard terms if the seller is a third-party seller. At the bottom of the page, the **Purchase Reserved Instances** wizard keeps a running tally of the total in your cart.

**Purchase Reserved Instances**

Platform: Linux/UNIX | Availability Zone: Any | Tenancy: Default  
Instance Type: m1.medium | Term: Any | Offering Type: Any

Seller	Term	Effective Rate	Upfront Price	Hourly Rate	Availability Zone	Offering Type	Quantity Available	Desired Quantity	
AWS	12 months	\$0.082	\$122.00	\$0.068	us-west-2c	Light Utilization	Unlimited	1	Add to Cart
AWS	12 months	\$0.082	\$122.00	\$0.068	us-west-2a	Light Utilization	Unlimited	1	Add to Cart
AWS	12 months	\$0.082	\$122.00	\$0.068	us-west-2b	Light Utilization	Unlimited	1	Add to Cart
3rd Party	7 months	\$0.064	\$180.00	\$0.028	us-west-2a	Heavy Utilization	1	1	Add to Cart
AWS	36 months	\$0.061	\$192.00	\$0.054	us-west-2c	Light Utilization	Unlimited	1	Add to Cart
AWS	36 months	\$0.061	\$192.00	\$0.054	us-west-2a	Light Utilization	Unlimited	1	Add to Cart
AWS	36 months	\$0.061	\$192.00	\$0.054	us-west-2b	Light Utilization	Unlimited	1	Add to Cart
AWS	12 months	\$0.074	\$277.00	\$0.042	us-west-2c	Medium Utilization	Unlimited	1	Add to Cart
AWS	12 months	\$0.074	\$277.00	\$0.042	us-west-2a	Medium Utilization	Unlimited	1	Add to Cart
AWS	12 months	\$0.074	\$277.00	\$0.042	us-west-2b	Medium Utilization	Unlimited	1	Add to Cart
3rd Party	23 months	\$0.057	\$319.24	\$0.038	us-west-2a	Medium Utilization	1	1	Add to Cart

Your cart: 2 Reserved Instances, Total Due Now: **\$314.00**  
Additional taxes may apply.

Cancel View Cart

6. Click **View Cart** to see a summary of the Reserved Instances that you have selected.

If you want to add more Reserved Instances to your cart, click **Add More To Cart**. If you want to remove an item from your cart, click **Delete**.

Or, click **Cancel** if you want to start over or search for a different set of Reserved Instances.

**Purchase Reserved Instances**

**Shopping Cart**

Items to buy now	Upfront Price	Hourly Price	Quantity
Light Utilization Linux/UNIX m1.medium in us-west-2c for 12 months Payment Terms: Upfront + Every Hour Sold By: Amazon <a href="#">Delete</a>	\$122.00	\$0.068	1
Light Utilization Linux/UNIX m1.medium in us-west-2c for 36 months Payment Terms: Upfront + Every Hour Sold By: Amazon <a href="#">Delete</a>	\$192.00	\$0.054	1

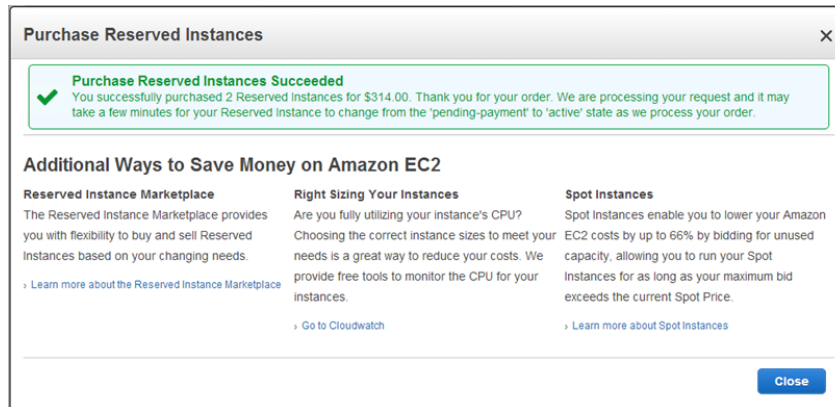
Your cart: 2 Reserved Instances, Total Due Now: **\$314.00**  
Additional taxes may apply.

Add More To Cart Purchase

- Click **Purchase** when you have all the Reserved Instances you want to purchase, and you want to check out.

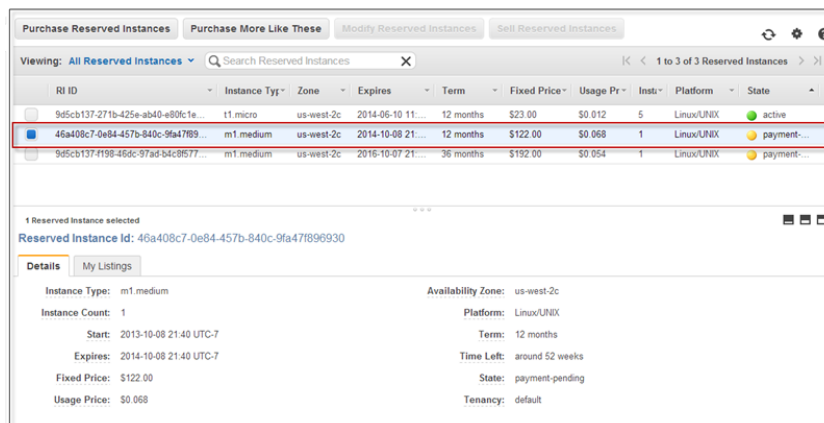
**Note**

If at the time of purchase, there are offerings similar to your choice but with a lower price, AWS will sell you the offerings at the lower price.



Your purchase is complete.

- To verify your order, go to the **Reserved Instances** page in the EC2 console.



The **Reserved Instances** page displays a list of Reserved Instances that belong to your account, including the new Reserved Instance that you just purchased.

You can use your Reserved Instance any time after your purchase is complete. This means that the **State** of your Reserved Instance has changed from *payment-pending* to *active*. To use your Reserved Instance, you launch an instance in the same way you launch an On-Demand instance. Just make sure to specify the same criteria that you specified for your Reserved Instance. AWS will automatically charge you the lower hourly rate. You do not have to restart your instance.

## Amazon EC2 CLI

### To find and purchase a Reserved Instance

- Use `ec2-describe-reserved-instances-offerings` to get a list of Reserved Instance offerings that match your specifications. In this example, we'll check to see what m1.small, Linux/UNIX Reserved Instances are available in the sa-east-1b Availability Zone.

```
PROMPT> ec2-describe-reserved-instances-offerings -t m1.small -z sa-east-1b  
-d Linux/UNIX --headers
```

Amazon EC2 returns output similar to the following example:

```
PROMPT> ec2-describe-reserved-instances-offerings  
Type Source ReservedInstancesOfferingId AvailabilityZone InstanceType Duration  
FixedPrice UsagePrice ProductDescription Currency InstanceTenancy Offering  
Type  
OFFERING AWS 4b2293b4-3236-49f5-978d-a74c3example sa-east-1b m1.small 3y  
574.0 0.0 Linux/UNIX USD default Heavy Utilization  
Type Frequency Amount  
RECURRING-CHARGE Hourly 0.021  
OFFERING AWS 3a98bf7d-07e1-4b33-8e11-e5314example sa-east-1b m1.small 3y  
473.0 0.031 Linux/UNIX USD default Medium Utilization  
OFFERING AWS 438012d3-5fc5-4e49-a88e-273edexample sa-east-1b m1.small 3y  
203.0 0.055 Linux/UNIX USD default Light Utilization  
OFFERING AWS d586503b-bb92-41fa-9065-e5b90example sa-east-1b m1.small 1y  
372.94 0.0 Linux/UNIX USD default Heavy Utilization  
Type Frequency Amount  
RECURRING-CHARGE Hourly 0.03  
OFFERING AWS ceb6a579-b235-41e2-9aad-15a23example sa-east-1b m1.small 1y  
307.13 0.04 Linux/UNIX USD default Medium Utilization  
OFFERING AWS 649fd0c8-4ffb-443d-824d-ee3fexample sa-east-1b m1.small 1y  
131.63 0.07 Linux/UNIX USD default Light Utilization  
OFFERING 3rd Party b6121943-9faf-4350-8047-bc6d4example sa-east-1b m1.small  
10m - 0.032 Linux/UNIX USD default Medium Utilization  
Type Count Price  
PRICING_DETAIL 2 $1.2  
OFFERING 3rd Party 08edcff2-8143-4c1d-b23c-e4c11example sa-east-1b m1.small  
5m - 0.032 Linux/UNIX USD default Medium Utilization  
Type Count Price  
PRICING_DETAIL 19 $1.2  
PRICING_DETAIL 4 $1.23
```

The preceding output shows a part of the overall offerings that are available.

#### Tip

You can filter this list to return only certain types of Reserved Instances offerings of interest to you. For more information about how to filter the results, see [ec2-describe-reserved-instances-offerings](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

2. From the list of available Reserved Instances, purchase the Reserved Instances that meet your requirements. To purchase a Reserved Instance, use the following command.

```
PROMPT> ec2-purchase-reserved-instances-offering --offering offering --in  
stance-count count
```

Amazon EC2 returns output similar to the following:

```
PURCHASE af9f760e-c1c1-449b-8128-1342dexample 438012d3-80c7-42c6-9396-  
a209cexample
```

The response includes the offering ID and a reservation ID.

3. Write down and save the reservation ID for future reference.

4. Verify the purchase.

```
PROMPT> ec2-describe-reserved-instances
```

Amazon EC2 returns output similar to the following:

```
RESERVEDINSTANCE af9f760e-c1c1-449b-8128-1342dexample sa-east-1b  
m1.small          1y          227.5    0.03    Linux/UNIX Active
```

You can run your Reserved Instance any time after your purchase is complete. To run your Reserved Instance, you launch it in the same way you launch an On-Demand instance. Make sure to specify the same criteria that you specified for your Reserved Instance. AWS will automatically charge you the lower hourly rate.

## Amazon EC2 API

### To find and purchase a Reserved Instance

1. Use `DescribeReservedInstancesOfferings` to get a list of Reserved Instance offerings that match your specifications. In this example, we'll check to see the available Linux/UNIX, Heavy Utilization Reserved Instances.

```
https://ec2.amazonaws.com/?Action=DescribeReservedInstancesOfferings  
&MaxResults=50  
&ProductDescription=Linux%2FUNIX  
&OfferingType=Heavy+Utilization  
&AUTHPARAMS
```

#### Note

When using the Query API the “/” is denoted as “%2F”.

Following is an example response.

```
<DescribeReservedInstancesOfferingsResponse xmlns='http://ec2.amazon  
aws.com/doc/2012-08-15/'>  
  <requestId>768e52ac-20f5-42b1-8559-e70e9example</requestId>  
  <reservedInstancesOfferingsSet>  
    <item>  
      <reservedInstancesOfferingId>d0280f9e-afcl-47f3-9899-  
c3a2cexample</reservedInstancesOfferingId>  
      <instanceType>m1.xlarge</instanceType>  
      <availabilityZone>us-east-1a</availabilityZone>  
      <duration>25920000</duration>  
      <fixedPrice>195.0</fixedPrice>  
      <usagePrice>0.0</usagePrice>  
      <productDescription>Linux/UNIX</productDescription>  
      <instanceTenancy>dedicated</instanceTenancy>  
      <currencyCode>USD</currencyCode>  
      <offeringType>Heavy Utilization</offeringType>  
      <recurringCharges>  
        <item>  
          <frequency>Hourly</frequency>  
          <amount>0.2</amount>
```

```
        </item>
      </recurringCharges>
      <marketplace>true</marketplace>
      <pricingDetailsSet>
        <item>
          <price>195.0</price>
          <count>1</count>
        </item>
        <item>
          <price>310.0</price>
          <count>1</count>
        </item>
        <item>
          <price>377.0</price>
          <count>1</count>
        </item>
        <item>
          <price>380.0</price>
          <count>1</count>
        </item>
      </pricingDetailsSet>
    </item>
  <item>
    <reservedInstancesOfferingId>649fd0c8-7846-46b8-8f84-a6400example</reservedInstancesOfferingId>
    <instanceType>m1.large</instanceType>
    <availabilityZone>us-east-1a</availabilityZone>
    <duration>94608000</duration>
    <fixedPrice>1200.0</fixedPrice>
    <usagePrice>0.0</usagePrice>
    <productDescription>Linux/UNIX</productDescription>
    <instanceTenancy>default</instanceTenancy>
    <currencyCode>USD</currencyCode>
    <offeringType>Heavy Utilization</offeringType>
    <recurringCharges>
      <item>
        <frequency>Hourly</frequency>
        <amount>0.052</amount>
      </item>
    </recurringCharges>
    <marketplace>>false</marketplace>
    <pricingDetailsSet/>
  </item>
</reservedInstancesOfferingsSet>
<nextToken>QUUVo/0S3X6nEBjSQZR/pRRlCPP/5Lrx79Wyxexample</nextToken>
</DescribeReservedInstancesOfferingsResponse>
```

2. From the list of available Reserved Instances in the previous example, select the Reserved Instance you want and specify a limit price.

```
https://ec2.amazonaws.com/?Action=PurchaseReservedInstancesOffering
&ReservedInstancesOfferingId=d0280f9e-afc1-47f3-9899-c3a2cexample
&InstanceCount=1
&LimitPrice.Amount=200
&AUTHPARAMS
```

Following is an example response.

```
<PurchaseReservedInstancesOfferingResponse xmlns="http://ec2.amazonaws.com/doc/2012-08-15/">
  <requestId>59dbff89-35bd-4eac-99ed-be587example</requestId>
  <reservedInstancesId>e5a2ff3b-7d14-494f-90af-0b5d0example</reservedInstancesId>
</PurchaseReservedInstancesOfferingResponse>
```

3. To verify the purchase, check for your new Reserved Instance.

```
http://ec2.amazonaws.com/?Action=DescribeReservedInstances
&AUTHPARAMS
```

Following is an example response:


```
<DescribeReservedInstancesResponse xmlns='http://ec2.amazonaws.com/doc/2012-08-15/'>
  <requestId>ebe3410a-8f37-441d-ae11-2e78eexample</requestId>
  <reservedInstancesSet>
    <item>
      <reservedInstancesId>e5a2ff3b-7d14-494f-90af-0b5d0example</reservedInstancesId>
      <instanceType>m1.xlarge</instanceType>
      <availabilityZone>us-east-1a</availabilityZone>
      <start>2012-08-23T15:19:31.071Z</start>
      <duration>2592000</duration>
      <fixedPrice>195.0</fixedPrice>
      <usagePrice>0.0</usagePrice>
      <instanceCount>1</instanceCount>
      <productDescription>Linux/UNIX</productDescription>
      <state>active</state>
      <instanceTenancy>dedicated</instanceTenancy>
      <currencyCode>USD</currencyCode>
      <offeringType>Heavy Utilization</offeringType>
      <recurringCharges>
        <item>
          <frequency>Hourly</frequency>
          <amount>0.2</amount>
        </item>
      </recurringCharges>
    </item>
  </reservedInstancesSet>
</DescribeReservedInstancesResponse>
```

You can run your Reserved Instance any time after your purchase is complete. To run your Reserved Instance, you launch it in the same way you launch an On-Demand EC2 instance. Make sure to specify the same criteria that you specified for your Reserved Instance. AWS will automatically charge you the lower hourly rate.

## Reading Your Statement (Invoice)

### Account Activity

You can find out about the charges and fees to your account by viewing your **Account Activity** page in the AWS Management Console. You can access it by clicking the drop-down arrow beside your account name.



The **Account Activity** page will show all charges against your account—such as upfront and one-time fees and recurring charges. You can get both a summary of all your charges and a detailed list of your charges.

The upfront charges from your purchase of third-party Reserved Instances in the Reserved Instance Marketplace will be listed in the **AWS Marketplace Charges** section, with the name of the seller displayed beside it. However, all recurring or usage charges for these Reserved Instances will be listed in the **AWS Service Charges** section.

Account Activity Welcome [Rick Ingram](#) | [Sign Out](#)  
Account Number: [XXXXXXXXXX](#)

**You are eligible for the AWS Free Usage Tier.** See the [Getting Started Guide: AWS Free Usage Tier](#) to learn how to get started with the free usage tier.

**Monitor your estimated charges.** Enable [Now](#) to begin setting billing alerts that automatically e-mail you when charges reach a threshold you define. [Learn More](#)

**This Month's Activity as of August 29, 2012**  
The statement period for this report is August 1 - August 31, 2012. The charges on this page currently show activity through approximately 08/29/2012 02:01 GMT.

Summary	
<b>AWS Service Charges</b>	<b>\$69.60</b>
New Purchases and Adjustments <small>(<a href="#">View Info</a>)</small>	\$69.60
<a href="#">View charges and download PDFs</a>	
<b>AWS Marketplace Charges</b>	<b>\$7.91</b>
New Purchases and Adjustments <small>(<a href="#">View Info</a>)</small>	\$7.91
<a href="#">View charges and download PDFs</a>	
<b>Total new charges for this statement</b>	<b>\$77.51</b>
<a href="#">Payment Summary</a>	-\$77.51
<b>Outstanding balance for this statement</b>	<b>\$0.00</b>

Details	
<a href="#">Expand All Services</a>   <a href="#">Collapse All Services</a> <span style="float: right;"><a href="#">Printer Friendly Version</a></span>	
<b>AWS Service Charges</b>	<b>\$69.60</b>
<a href="#">Amazon Elastic Compute Cloud</a>	\$69.60
<a href="#">Download Usage Report</a>	
<a href="#">Amazon Simple Notification Service</a>	\$0.00
<a href="#">Download Usage Report</a>	
<a href="#">VAT to be collected</a>	\$0.00
<b>AWS Marketplace Charges</b>	<b>\$7.91</b>
<a href="#">Amazon Elastic Compute Cloud (sold by ThoughtPac)</a>	\$1.00
<a href="#">RI Marketplace: ml.small Linux/UNIX RI in sa-east-1-0 (sold by ThoughtPac)</a>	\$6.91
<a href="#">VAT to be collected</a>	\$0.00

\* Usage and recurring charges for this statement period will be charged on your next billing date, September 1, 2012. Estimated charges shown on this page, or shown on any notifications that we send to you, may differ from your actual charges for this statement period. This is because estimated charges presented on this page do not include usage charges accrued during this statement period after the date you view this page. Similarly, information about admitted charges sent to you in a notification do not include usage charges accrued during this statement period after the date we send you the notification. One-time fees and subscription charges are assessed separately from usage and recurring charges, on the date that they occur.

You can view the charges online, and you can also download a PDF rendering of the charge information.



**Summary**

AWS Service Charges			\$69.60
<b>New Purchases and Adjustments</b> (More Info)			\$69.60
<input type="checkbox"/> View charges and download PDFs			
AWS Services: Subscription charge	August 24, 2012		\$1.20
AWS Services: Subscription charge	August 24, 2012		\$2.40
AWS Services: Subscription charge	August 22, 2012		\$6.00
AWS Services: Subscription charge	August 27, 2012		\$34.00
AWS Services: Subscription charge			
AWS Services: Subscription charge			
AWS Services: Subscription charge			
AWS Services: Subscription charge			
AWS Services: Subscription charge			
AWS Services: Subscription charge			
<b>Total New Purchases and Adjustment</b>			
AWS Marketplace Charges			
<b>New Purchases and Adjustments</b> (More Info)			
<input type="checkbox"/> View charges and download PDFs			
AWS Marketplace: Subscription charge			
AWS Marketplace: Subscription charge			
AWS Marketplace: Subscription charge			
AWS Marketplace: Subscription charge			
AWS Marketplace: Subscription charge			
AWS Marketplace: Subscription charge			
AWS Marketplace: Subscription charge			
<b>Total New Purchases and Adjustment</b>			
<b>Total new charges for this statement</b>			

**Amazon Web Services Statement**

Email or talk to us about your AWS account or bill, visit [aws.amazon.com/contact-us/](http://aws.amazon.com/contact-us/)

---

**Statement Summary**

Statement Number: \_\_\_\_\_  
 Statement Date: August 20, 2012

**TOTAL AMOUNT DUE ON August 20, 2012 \$1.20**

---

This statement is for the billing period August 1 - August 31, 2012

Greetings from Amazon Web Services, we're writing to provide you with an electronic statement of your transactions on the AWS Marketplace. Additional information regarding your bill, individual service charge details, and your account history are available on the Account Activity Page.

Summary	
<b>AWS Marketplace Charges</b>	\$1.20
1 x Amazon Elastic Compute Cloud (one time fee)	\$1.20
Charges	\$0.00
Credits	\$0.00
Tax *	\$0.00
<b>Total for this statement</b>	<b>\$1.20</b>

\* Details of services from Japan on which consumption tax is included are provided on the Account Activity Page, visit [aws.amazon.com/](http://aws.amazon.com/)

\* This statement is not an invoice. It is a statement of transactions on the AWS Marketplace. If you need an invoice, please contact the seller listed in the description of the transaction.

Detail	
Amazon Elastic Compute Cloud sold by ThoughtPot	\$1.20
1 x Amazon Elastic Compute Cloud (one time fee)	\$1.20
Availability Zone - GRUZ	
m1.small	
Charges	\$0.00
Estimated US sales tax to be collected	\$0.00

The **Detail** section contains information about the Reserved Instance—such as the Availability Zone, instance type, cost, and number of instances. It also includes the name of the seller of the Reserved Instances that you purchased.

## Obtaining Information About Your Reserved Instances

Information about your Reserved Instances, such as state, instance type, Availability Zone, and term is useful when you decide to use the capacity reservation. You can check information about the Reserved Instances that are available to your account using any of the Amazon EC2 tools that you've used either for purchasing or selling.

### Reserved Instance States

Reserved Instances can be in one of the following states:

- **Active**—The Reserved Instance is available for use.
- **Payment-Pending**—Amazon Web Services (AWS) is processing your payment for the Reserved Instance. You will be able to use the Reserved Instance when the state becomes Active.
- **Retired**—The Reserved Instance has been terminated. It could have reached this state because of any of the following reasons:
  - AWS did not receive your payment. For example, the credit card transaction did not go through.
  - The Reserved Instance term expired.
  - The Reserved Instance was canceled.

Status information displayed in the **State** column in the **Reserved Instance** page is different from the status information displayed in the **Listing State** in the **My Listings** tab. The State column displays status of the *Reserved Instance*; the Listing State displays status of the Reserved Instance *Listing*. The **My Listings** tab shows information only if you are a seller in the Reserved Instance Marketplace. For more information, see [Reserved Instance Listing States](#) (p. 260).

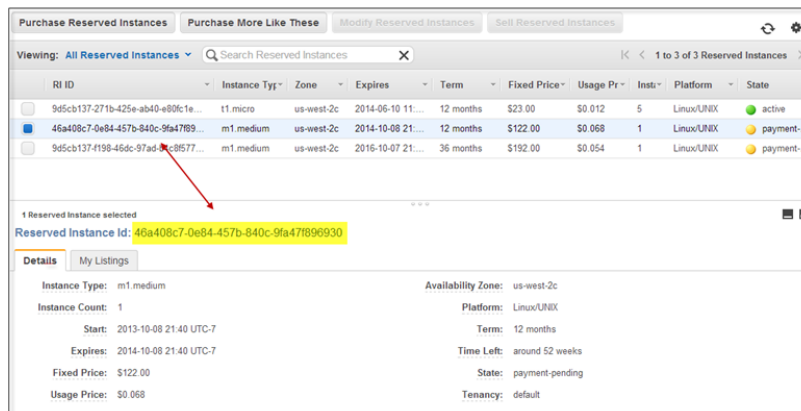
## AWS Management Console

### To view your listing

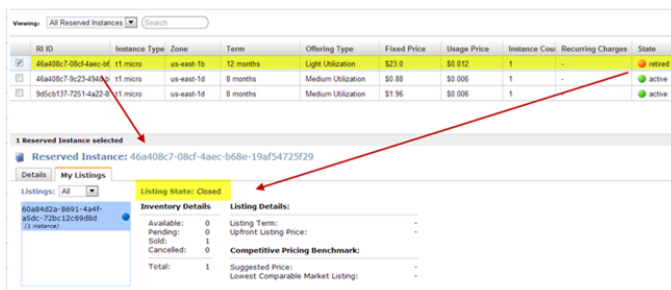
1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Reserved Instances** in the **Navigation** pane.

The **Reserved Instances** page displays a list of your account's instances.

3. Select a Reserved Instance. The **Details** tab displays details about the instance you selected.



4. If you are a seller in the Reserved Instance Marketplace and you want information about your Reserved Instance listing, click the **My Listings** tab. You will see details about the Reserved Instance listing you selected.



## Amazon EC2 CLI

### To view your listing

- Run `ec2-describe-reserved-instances-listings` to get details about your listing.

```
PROMPT> ec2-describe-reserved-instances-listings
```

Amazon EC2 returns output similar to the following:

```
PROMPT> ec2-describe-reserved-instances-listings
Type ReservedInstancesListingId ReservedInstancesId CreateDate UpdateDate
Status StatusMessage
LISTING 615d8a10-8224-4c19-ba7d-b9aa0example 1ba8e2e3-d20d-44ec-b202-
fcb6aexample Wed Aug 22 09:02:58 PDT 2012 Wed Aug 22 14:24:26 PDT 2012 can
celled cancelled
INSTANCE-COUNT available 0
INSTANCE-COUNT sold 0
INSTANCE-COUNT cancelled 1
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 10 $1.2
PRICE-SCHEDULE 9 $1.08
PRICE-SCHEDULE 8 $0.96
PRICE-SCHEDULE 7 $0.84
PRICE-SCHEDULE 6 $0.72
PRICE-SCHEDULE 5 $0.6
PRICE-SCHEDULE 4 $0.48
PRICE-SCHEDULE 3 $0.36
PRICE-SCHEDULE 2 $0.24
PRICE-SCHEDULE 1 $0.12
LISTING d5fa5166-83c3-40e4-abb2-b7298example 1ba8e2e3-d20d-44ec-b202-
fcb6aexample Wed Aug 22 14:31:55 PDT 2012 Wed Aug 22 14:42:40 PDT 2012 closed
closed
INSTANCE-COUNT available 0
INSTANCE-COUNT sold 1
INSTANCE-COUNT cancelled 0
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 10 $0.9
PRICE-SCHEDULE 9 $0.81
PRICE-SCHEDULE 8 $0.72
PRICE-SCHEDULE 7 $0.63
PRICE-SCHEDULE 6 $0.54
PRICE-SCHEDULE 5 $0.45
PRICE-SCHEDULE 4 $0.36
PRICE-SCHEDULE 3 $0.27
PRICE-SCHEDULE 2 $0.18
PRICE-SCHEDULE 1 $0.09
....
LISTING 095c0e18-c9e6-4692-97e5-653e0example b847fa93-c736-4eae-bca1-
e3147example Tue Aug 28 18:21:07 PDT 2012 Tue Aug 28 18:21:07 PDT 2012 active
active
INSTANCE-COUNT available 1
INSTANCE-COUNT sold 0
INSTANCE-COUNT cancelled 0
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 5 $1.2
PRICE-SCHEDULE 4 $1.2
PRICE-SCHEDULE 3 $1.2
PRICE-SCHEDULE 2 $1.2
PRICE-SCHEDULE 1 $1.2
```

## Amazon EC2 API

### To view your listing

- Call `DescribeReservedInstancesListings` to get details about your listing.

The call should look like this example:

```
http://ec2.amazonaws.com/?Action=DescribeReservedInstancesListings
&AUTHPARAMS
```

Following is an example response.

```
<DescribeReservedInstancesListingsResponse>
  <requestId>cec5c904-8f3a-4de5-8f5a-ff7f9example</requestId>
  <reservedInstancesListingsSet>
    <item>
      <reservedInstancesListingId>5ec28771-05ff-4b9b-aa31-
9e57dexample</reservedInstancesListingId>
      <reservedInstancesId>f127bd27-cee4-443a-a76b-a5af9example</re
servedInstancesId>
      <createDate>2012-08-30T17:11:09.449Z</createDate>
      <updateDate>2012-08-30T21:00:42.300Z</updateDate>
      <status>active</status>
      <statusMessage>active</statusMessage>
      <instanceCounts>
        <item>
          <state>Available</state>
          <instanceCount>2</instanceCount>
        </item>
        <item>
          <state>Sold</state>
          <instanceCount>1</instanceCount>
        </item>
        <item>
          <state>Cancelled</state>
          <instanceCount>0</instanceCount>
        </item>
        <item>
          <state>Pending</state>
          <instanceCount>0</instanceCount>
        </item>
      </instanceCounts>
      <priceSchedules>
        <item>
          <term>11</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>true</active>
        </item>
        <item>
          <term>10</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>false</active>
        </item>
      </priceSchedules>
    </item>
  </reservedInstancesListingsSet>
</DescribeReservedInstancesListingsResponse>
```

```
<item>
  <term>9</term>
  <price>2.5</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>8</term>
  <price>2.0</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>7</term>
  <price>2.0</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>6</term>
  <price>2.0</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>5</term>
  <price>1.5</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>4</term>
  <price>1.5</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>3</term>
  <price>0.7</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>2</term>
  <price>0.7</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
<item>
  <term>1</term>
  <price>0.1</price>
  <currencyCode>USD</currencyCode>
  <active>>false</active>
</item>
</priceSchedules>
<tagSet/>
<clientToken>listRI1</clientToken>
</item>
```

```
</reservedInstancesListingsSet>  
</DescribeReservedInstancesListingsResponse>
```

## Modifying Your Reserved Instances

When your computing needs change, you can modify your Reserved Instances and continue to benefit from your capacity reservation. You can move your Reserved Instances between Availability Zones within the same region, and you can change the network platform between EC2-Classic and EC2-VPC. In addition, starting with Amazon EC2 API version 2013-10-01, you can change the instance type of your reservation to a larger or smaller instance type in the same family (e.g., the M1 instance type) if capacity is available and if your reservation is for the Linux/UNIX or Windows product platform.

Modification does not change the remaining term of your Reserved Instances. Their end dates remain the same. There is no fee for modifying your Reserved Instances, and you do not receive any new bills or invoices.

Instances covered by Reserved Instances continue to run even if you modify your reservation. However, after modification, the pricing benefit of the Reserved Instances starts applying to instances that match the new parameters of your Reserved Instances. You are charged at the On-Demand rate for the EC2 instances no longer receiving the benefits of the modified Reserved Instances, unless your account has other applicable reservations, in which case you will be charged at the rate of these other Reserved Instances.

Modification does not affect how you use, purchase, or sell Reserved Instances. When you purchase Reserved Instances, you still need to specify the product platform, instance type, network platform, tenancy, term length, offering type, and Availability Zone.

You can view the status of your request in the AWS Management Console, or when you call the `DescribeReservedInstancesModifications` API action or `ec2-describe-reserved-instances-modifications` CLI command.

This section discusses the modification process:

- [Understanding the Modification Process \(p. 227\)](#)—What happens when I submit a modification request?
- [Determining the Status of Your Modification \(p. 229\)](#)—How do I track my modification requests?
- [Requirements for Modification \(p. 230\)](#)—Which Reserved Instances can I modify?
- [Changing the Instance Type of Your Reservations \(p. 230\)](#)—When can I change the instance size of my Reserved Instances?
- [Submitting Modification Requests \(p. 232\)](#)—How do I modify my Reserved Instances?

## Understanding the Modification Process

You can modify your Reserved Instances in one or more of the following ways:

- Move them between Availability Zones within the same region.

If you modify the Availability Zone of your Reserved Instances, the capacity reservation and pricing benefit stop applying to instance usage in the original zone, and start applying to usage in the new Availability Zone.

- Change the network platform of the Reserved Instances between EC2-VPC and EC2-Classic.

If you modify the network platform of your Reserved Instances, the capacity reservation stops applying to instance usage with the original network platform, and starts applying to usage with the new network

platform. However, the pricing benefit continues to apply to both EC2-Classic and EC2-VPC instance usage matching the remaining Reserved Instances parameters.

- Upgrade or downgrade the instance type of your Reserved Instances within the same instance family.

If you modify the instance type of your Reserved Instances, you must ensure that the instance family of your reservation has a larger or smaller instance type available and that you have enough applicable Reserved Instances to make the change.

You can modify your whole reservation or a subset of your reservation. When you modify a subset of your reservation, Amazon EC2 splits your original Reserved Instances into two or more new Reserved Instances. For example, if you have Reserved Instances for 10 instances in us-east-1a, and decide to move 5 instances to us-east-1b, the modification request results in two new Reserved Instances—one for 5 instances in us-east-1a (the original zone), and the other for 5 instances in us-east-1b.

Amazon EC2 fulfills your modification request as soon as possible, depending on available capacity. Until your modification request completes, the capacity reservation and pricing benefit associated with your Reserved Instances continue to be based on the original parameters of your reservation.

**Note**

You cannot cancel or change a pending modification request after you submit it. While your modification is being processed, the status of the Reserved Instances that you're modifying is *active (pending modification)*. After the modification has completed—successfully or otherwise—you can submit another modification request to roll back any changes you made.

If your Reserved Instances modification request succeeds:

- The modified reservation becomes effective immediately and the pricing benefit of the Reserved Instances is applied to the new instances beginning at the hour of the modification request. For example, if you successfully modify your Reserved Instances at 9:15PM, the pricing benefit transfers to your new instance at 9:00PM. (You can get the *effective date* of the modified Reserved Instances by using the `DescribeReservedInstances` API action or the `ec2-describe-reserved-instances` CLI command.)
- The end date of the modified Reserved Instances is the same as the original end date of the reservation. If you successfully modify a three-year reservation that had 16 months left in its term, the resulting modified reservation is a 16-month Reserved Instance with the same end date as the original Reserved Instances.

For example, in the following table, Modification #1 shows RI af9f760... was modified successfully and retired on 2013-08-30 16:00 UTC-7 and a new RI 46a408c... was created as a result of the modification on the same date-time that RI af9f760... was retired. Modification #2 shows RI 46a408c... was modified successfully and retired on 2013-09-03 14:00 UTC-7 and a new RI b847fa9... was created as a result of the modification on the same date-time that RI 46a408c... was retired.

RI ID	Start	Expires	Term	State
93bbca2-46a...	2013-08-28 10:01 UTC-7	2016-08-27 10:01 UTC-7	36 months	active
93bbca2-6a7...	2013-09-04 14:00 UTC-7	2014-08-30 15:45 UTC-7	12 months	active
93bbca2-c7a...	2013-08-30 14:00 UTC-7	2014-08-28 10:01 UTC-7	12 months	active
b847fa93-0b3...	2013-09-03 14:00 UTC-7	2014-08-30 15:45 UTC-7	12 months	active
Modification #2	2013-08-28 10:01 UTC-7	2013-08-30 14:00 UTC-7	12 months	retired
46a408c7-61d...	2013-08-30 16:00 UTC-7	2013-09-03 14:00 UTC-7	12 months	retired
Modification #1	2013-08-30 15:45 UTC-7	2013-08-30 16:00 UTC-7	12 months	retired
af9f760e-92bf...	2013-08-30 15:45 UTC-7	2013-08-30 16:00 UTC-7	12 months	retired
d16f7a91-b9ef...	2013-08-30 16:00 UTC-7	2013-09-04 14:00 UTC-7	12 months	retired

- The original reservation is retired. Its end date is the start date of the new reservation, and the end date of the new reservation is the same as the end date of the original Reserved Instance when it was active.

- The modified Reserved Instances shows \$0 fixed price and not the fixed price of the original Reserved Instances.

**Note**

The fixed price of the modified reservation does not affect the discount tier calculations applied to your account, which are based on the fixed price of the original reservation.

If your modification request fails:

- Your Reserved Instances maintain the original properties that they had prior to your request.
- Your Reserved Instances are available for another modification request.

You can determine the status of your request by looking at the *state* of the Reserved Instances that you are modifying. For information, see [Determining the Status of Your Modification \(p. 229\)](#).

## Determining the Status of Your Modification

The state of your modification request is displayed in the **State** field in the AWS Management Console. You can also use the `DescribeReservedInstancesModifications` API action to get detailed information about your modification request: The *state* returned shows your request as *in-progress*, *fulfilled*, or *failed*.

You can only modify your Reserved Instances if they are *active*. You are not able to modify Reserved Instances that are in any other state. You also cannot modify your Reserved Instances if they are listed in the Reserved Instance Marketplace. For more information, see [Requirements for Modification \(p. 230\)](#). If your Reserved Instances are not in the active state or cannot be modified, the **Modify Reserved Instances** button in the AWS Management Console is not enabled. If you use the API to modify Reserved Instances that are not active, you will get an error.

In cases when you select Reserved Instances whose Availability Zone and network platform can be modified, but whose instance types cannot be changed, the **Modify Reserved Instances** button is enabled and you can proceed to the **Modify Reserved Instances** page. However, you do not have the option to modify the instance type. For more information, see [Changing the Instance Type of Your Reservations \(p. 230\)](#).

### Modification States

**The modification request is being processed.** While the modification request is being processed, the status of the Reserved Instances being modified show as *active (pending modification)*. The Reserved Instances are in this state only for a short period. The state reverts to *active* or becomes *retired*, depending on the success of the modification. (If you use the `DescribeReservedInstancesModifications` API action, the status of your modification request should show *processing*.)

**The modification succeeded.** If the modification is successful, the Reserved Instances being modified are *retired*, and new Reserved Instances are created with the modification configuration that you requested. The status of these new Reserved Instances is *active*. (If you use the `DescribeReservedInstancesModifications` API action, the status of your modification request should show *fulfilled*.)

**Note**

For the brief period that the new Reserved Instances are being activated, the original Reserved Instances show as *retired (pending modification)*.

**The modification failed.** If the modification did not complete, the Reserved Instances being modified return to the *active* state. (If you use the `DescribeReservedInstancesModifications` API action, the status of your modification request should show *failed*.) For information about why some Reserved Instances cannot be modified, see [Requirements for Modification \(p. 230\)](#).



## Requirements for Modification

Amazon EC2 processes your modification request if we have sufficient Reserved Instances capacity for your target configuration, and the following conditions are met:

- You own the Reserved Instances that you are modifying.
- The Reserved Instances you are modifying are active.
- The Reserved Instances are not pending another modification request.

You may modify your Reserved Instances as frequently as you like. However, you cannot submit a modification request for Reserved Instances that are still pending a previous modification request; that is, the previous modification request is in the *active (pending modification)* state.

- The Reserved Instances are not listed in the Reserved Instance Marketplace.

To modify your Reserved Instances that are listed in the Reserved Instance Marketplace, cancel the listing, modify the Reserved Instances, and then list them again.

In addition, you cannot modify a Reserved Instance Marketplace offering before or at the same time that you purchase it. However, you can submit a modification request after you purchase a Marketplace offering. For more information, see [Reserved Instance Marketplace \(p. 209\)](#).

- The new configuration settings that you request for your Reserved Instances must be a unique combination of Availability Zone, instance type, and network platform.
- The Availability Zone of the Reserved Instances is online.
- The Availability Zone and network platform attribute that you request must be available to your account.
- For instance type modifications, the Reserved Instances you are modifying must be for the Amazon Linux/UNIX or Windows (without SQL Server) product platforms.
- The new instance type of the Reserved Instances must be in the same instance family as the original reservation.
- If you select multiple Reserved Instances for modification and one or more of these Reserved Instances are for a product that does not allow instance type modification, the **Modify Reserved Instances** page will not give you the option of changing the instance type of any of the selected Reserved Instances. To change the instance type of your Reserved Instances, make sure you only select the Reserved Instances with a product that allows instance type modification (currently Linux/UNIX and Windows).
- If you are upgrading the instance type of your Reserved Instances, your account must have enough smaller instance types in the same reservation to consolidate into a larger instance type of Reserved Instances.

For information about modifying instance types, see [Changing the Instance Type of Your Reservations \(p. 230\)](#). For information about how to modify Reserved Instances, see [Submitting Modification Requests \(p. 232\)](#).

## Changing the Instance Type of Your Reservations

Under certain conditions, you can adjust the instance type of your Reserved Instances. If you have capacity reservations for Amazon Linux/UNIX or Windows (without SQL Server) in instance families with multiple instance sizes, you can request a modification of your Reserved Instances to different instance types within the same family. Your request proceeds successfully if the capacity exists and the modification does not change the instance size footprint of your Reserved Instances.

For example, you can divide a reservation for one m1.large instance into four m1.small instances, or you can combine a reservation for four m1.small instances into one m1.large instance. In either case, the instance size footprint of the reservation does not change. On the other hand, you cannot change your reservation for two m1.small instances into one m1.large instance because the existing instance size footprint of your reservation is smaller than the proposed footprint.

Instance size footprints are determined by the normalization factor of the instance type and the number of instances in the reservation. This section discusses the two ways you can change the instance type of your Reserved Instances, and how you use the instance type's normalization factor to figure out what instance type modifications you can make.

- [Understanding Instance Normalization Factor \(p. 231\)](#)
- [Upgrading Your Instance Type \(p. 232\)](#)
- [Downgrading Your Instance Type \(p. 232\)](#)

The following instance types cannot be modified because there are no other sizes in their families:

- t1.micro
- cc1.4xlarge
- cc2.8xlarge
- cg1.8xlarge
- cr1.8xlarge
- hi1.4xlarge
- hs1.8xlarge

When you try to modify Reserved Instances that have instance types or product platforms that cannot be modified, you can proceed to the **Modify Reserved Instances** page, but the Amazon EC2 console limits the modification options to **Network**, **Availability Zone**, and **Count**.

Keep in mind that instance type modifications are allowed only if other Reserved Instances specification details match—such as region, utilization type, tenancy, product, end date—and capacity is available. In addition, such modifications do not change the discounted hourly usage fee that you are billed, or the end date of your reservation.

## Understanding Instance Normalization Factor

Each Reserved Instance has an instance size footprint. When you modify the instance type of Reserved Instances, the footprint is maintained even if the instance type is downsized or upsized. A modification request is not processed if the footprint of the target configuration does not match the original configuration.

The size of an instance type's footprint can be calculated by using its normalization factor, which is based on the type's size within the instance family. (Normalization factors are only meaningful within the same instance family; instance types cannot be modified from one family to another.)

The following table illustrates the normalization factor that applies within an instance family. For example, an m1.small instance has a normalization factor of 1, an m1.medium instance has a factor of 2, and an m1.large instance has a factor of 4.

Instance Size	Normalization Factor
small	1
medium	2
large	4
xlarge	8
2xlarge	16
4xlarge	32

Instance Size	Normalization Factor
8xlarge	64

Each Reserved Instance has a total number of *normalized instance units*, which is equal to the instance count multiplied by the normalization factor of the instance type. For example, an m1.medium has a normalization factor of 2 so a Reserved Instance for four m1.medium instances is worth eight normalized instance units. You arrive at the value this way:

$$4 \text{ [count]} \times 2 \text{ [normalization factor]}$$

You can allocate your Reserved Instances into different instance sizes across the same instance family as long as the total normalized instance units of your Reserved Instances remain the same. If you have Reserved Instances for four m1.medium instances, you can turn it into a reservation for eight m1.small instances (m1.small instances have a normalization factor of 1, thus  $8 \times 1$ ). The resulting Reserved Instances would have the same normalized instance units or instance size footprint.

## Upgrading Your Instance Type

You can consolidate a reservation of multiple smaller instance types that belong to one instance family into a reservation for a larger instance type in the same family. To upgrade or *upsized* your reservation, you must have enough of the smaller instance types in the same reservation to consolidate into larger instance types, and your reservation's overall instance size footprint does not change.

For example, you can convert four m1.small instances, which is equivalent to four normalized instance units ( $4 \times 1$ ), into one m1.large instance, which is also equivalent to four normalized units ( $1 \times 4$ ). However, you cannot convert a reservation for a single m1.small instance ( $1 \times 1$ ) into a reservation for an m1.large instance ( $1 \times 4$ ). The two reservations are not equal.

Only the smaller instance types from the M1, M2, M3, and C1 instance families can be changed or upgraded to larger instance type sizes.

## Downgrading Your Instance Type

You can change Reserved Instances for a large instance type to several Reserved Instances of smaller instance types of the same family. When you downgrade or *downsize* your Reserved Instances, you are actually dividing a large reservation into multiple smaller reservations. Just as with upsizing instance types, you can downsize successfully only if your reservation's overall instance size footprint does not change.

For example, a reservation for two m1.large instances equals eight normalized instance units ( $2 \times 4$ ). Assuming there is capacity, they can be converted to a Reserved Instance for four m1.medium instances, which is also equivalent to eight normalized instance units ( $4 \times 2$ ).

Only the larger instance types from the M1, M2, M3, and C1 families can be changed or downgraded to smaller instance type sizes.

For information about Amazon EC2 instance types, see [Instance Types \(p. 94\)](#). For information about Reserved Instances modification requests, see [Submitting Modification Requests \(p. 232\)](#).

## Submitting Modification Requests

You can modify your Reserved Instances using the AWS Management Console, the `ec2-modify-reserved-instances` CLI command, or the `ModifyReservedInstances` API action. This section shows you how to modify your Reserved Instances using the AWS Management Console.

To modify your Reserved Instances programmatically, see the following:


- [ec2-modify-reserved-instances](#) in the *Amazon Elastic Compute Cloud Command Line Reference*
- [ModifyReservedInstances](#) in the *Amazon Elastic Compute Cloud API Reference*
- [AWS SDK for Java](#)

## Modifying Reserved Instances Using the AWS Management Console

You can modify your Reserved Instances in a few steps using the **Reserved Instances** pages and screens in AWS Management Console.

1. Identify the Reserved Instances to modify.

On the **Reserved Instances** page, select the Reserved Instances to modify, and click **Modify Reserved Instances**. You can select one or more Reserved Instances to modify.



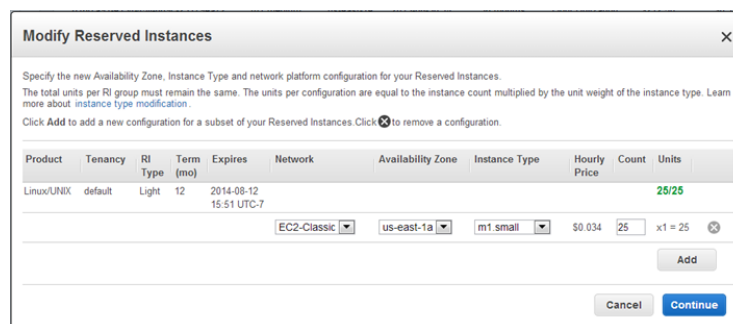
2. Decide how to modify your Reserved Instances.

On the **Modify Reserved Instances** page, specify the new configuration for your reservation—the Availability Zone, and the network platform (EC2-VPC or EC2-Classical)—and how many instances in your reservation to modify.

You also can modify the instance type of your Reserved Instances if they belong to instance families that have multiple instance types. For information, see [Changing the Instance Type of Your Reservations](#) (p. 230).

### Note

On the **Modify Reserved Instances** page, identical Reserved Instances (aside from the Availability Zone, network platform, and instance type) are grouped together. Also, if you select several Reserved Instances to modify, you get a message informing you if one or more of the selected Reserved Instances cannot be modified. For example, the status of the Reserved Instances might not be active, or the instance type cannot be modified. For information, see [Requirements for Modification](#) (p. 230).



On this page, you can also do the following:

- Add another configuration for a subset of your reservation.

To make two or more separate changes to subsets of your Reserved Instances so that one change is contained in one row, and another change is contained in another row, and so on, click **Add**.

For example, if you have 25 m1.small EC2-Classic Reserved Instances in us-east-1a and you want to modify 5 to be in us-east-1b, 5 to be in us-east-1c, and the remaining 10 in us-east-1d, do the following:

- Select us-east-1b in the **Availability Zone** list, and type 5 in the **Count** field.
- Click **Add** to add another row. Then select us-east-1c in the **Availability Zone** list, and type 5 in the **Count** field. Because you don't want to change the network platform or the instance type in this case, select EC2-Classic and m1.small in the **Network** and **Instance Type** fields, respectively.
- Click **Add** to add another row. This time, enter the specifications to keep the remaining Reserved Instances in their original configuration—EC2-Classic (**Network**), us-east-1a (**Availability Zone**), and 10 (**Count**)—but change them to be m1.medium (**Instance Type**).
- Monitor the number of Reserved Instances that you are modifying. Each target configuration row for your Reserved Instances has a count and units value. Count is the literal number of instances for the current instance type. Units represent the total instance size of your reservation relative to its instance family. The value of the units equals the count multiplied by the normalization factor of the instance type. For information about instance type modification, see [Changing the Instance Type of Your Reservations \(p. 230\)](#).

The **Modify Reserved Instances** page keeps track of the total count and units of the Reserved Instances that you have allocated for the target (or new) configurations versus the total number of Reserved Instances available for you to modify. You cannot proceed with your modification if the total units you have specified for your target configurations do not equal the total number of units of Reserved Instances available to modify.

For the modification example in which you started with 25 EC2-Classic Reserved Instances for m1.small instance type, here's how the Count-Units tally appears if you specify the following changes:

- Specify 5 instances in us-east-1b: the tally is *5/25*.
- Add another row specifying 5 instances us-east-1c: the tally is *10/25*.
- Add another row specifying 10 instances us-east-1d, but change the instance type to m1.medium: the tally is *30/25*.

The allocated total is displayed in red if you have specified either more or fewer Reserved Instances than are available for modification, but you are able to continue adding and deleting configurations. However, you cannot click **Continue** if the allocated total does not match the available total.

- Delete a configuration row by clicking the **X** button, if you decide that you do not want to allocate part of your reservation to that configuration.

**Note**

If the **Modify Reserved Instances** page contains only one row for configuration changes, you are not able to delete that row.

In the modification example, on the **Modify Reserved Instances** page, there is a set of three configurations being considered for your Reserved Instances: 5 each in us-east-1b and us-east-1c, and the remaining 10 in us-east-1d. The total units specified in the target configuration is more than the Reserved Instances available to modify. To delete 5 m1.small instances in us-east-1b, click **X** for the row with **Availability Zone** us-east-1b. With this deletion, note that the total units tally is *25/25* and displayed in green.

3. Confirm your modification choices.

When you finish specifying your target configurations, and click **Submit Modifications**, Amazon EC2 checks your modification request. If the configuration settings you requested were unique, the instance count that you specified matches the total number of Reserved Instances available for you to modify, you will be informed that Amazon EC2 is processing your request.

**We have received your Modify Reserved Instances request.**

✓ The Reserved Instances in your request will be **active (pending modification)** until the modification completes. If the modification succeeds, your Reserved Instances will be **retired**, and new **active** ones will reflect the modifications; otherwise, your original Reserved Instances will return to the **active** state.

**Note**

At this point, Amazon EC2 has only determined that the parameters of your modification request are valid, and Amazon EC2 can now process the request. Only during the processing of the modification request, after you get this screen, is the capacity check made. Your modification request can still fail, but at a later point, due to capacity not available.

4. Understand the messages about your modification request.

In some situations, you might get a message indicating incomplete or failed modification requests instead of a confirmation. Use the information in such messages as a starting point for resubmitting another modification request.

- Not all selected Reserved Instances can be processed for modification.

In the following message, Amazon EC2 identifies and lists the Reserved Instances that cannot be modified. If you receive a message like this, go to the **Reserved Instances** page and check the information details about these capacity reservations.

**⚠ We have received your Modify Reserved Instances request, but some of your requests could not be processed.**

The Reserved Instances in your request will be **active (pending modification)** until the modification completes. If the modification succeeds, your Reserved Instances will be **retired**, and new **active** ones will reflect the modifications; otherwise, your original Reserved Instances will return to the **active** state.

However, there was an error processing your request to modify the following **Reserved Instances**. Please try submitting your modification request for these Reserved Instances again.

- b847fa93-4da4-4d22-9401-6b791aaef686
- f127bd27-2d9c-4df5-ac3b-4fe42624f3d3

If the problem persists, contact AWS Customer Support.

- Error in processing your modification request.

You submitted one or more Reserved Instances for modification and none of your requests can be processed. Depending on the number of Reserved Instances you are modifying, you can get different versions of the message.

**⚠ Error**

The error(s) below occurred while we were processing your Modify Reserved Instances request. If you cannot resolve them, please contact AWS Customer Support.

Invalid value for 'targetReservedInstancesConfigurations': Configuration LeaseConfiguration(availabilityZone=us-east-1a, targetPlatform=EC2-Classical, instanceCount=1) is a duplicate.

**⚠ Error**

The error(s) below occurred while we were processing your Modify Reserved Instances request. If you cannot resolve them, please contact AWS Customer Support.

Invalid value for 'targetReservedInstancesConfigurations': Configuration LeaseConfiguration(availabilityZone=us-east-1a, targetPlatform=EC2-VPC, instanceCount=1) is a duplicate.

Invalid value for 'targetReservedInstancesConfigurations': Configuration LeaseConfiguration(availabilityZone=us-east-1a, targetPlatform=EC2-VPC, instanceCount=1) is a duplicate.

Invalid value for 'targetReservedInstancesConfigurations': Configuration LeaseConfiguration(availabilityZone=us-east-1a, targetPlatform=EC2-Classical, instanceCount=1) is a duplicate.

In these messages, Amazon EC2 displays the reasons why your request cannot be processed. For example, you might have specified the same target configuration—a combination of Availability Zone and platform—for one or more subsets of the Reserved Instances you are modifying. Try submitting these modification requests again, but ensure that instance details of the Reserved Instances match, and that the target configurations for all subsets of the Reserved Instances being modified are unique.

What do you want to do next?

- [Obtaining Information About Your Reserved Instances \(p. 222\)](#)
- [Buying Reserved Instances \(p. 212\)](#)



- [Selling in the Reserved Instance Marketplace \(p. 236\)](#)

## Selling in the Reserved Instance Marketplace

The Reserved Instance Marketplace gives you the flexibility to sell the remainder of your Reserved Instances as your needs change—for example, if you want to move instances to a different Amazon Web Services (AWS) region, change to another instance type, or sell capacity for projects that end before the Reserved Instance term expires. (Some restrictions—such as what is required to become a seller and when you can sell your reserved capacity apply. For information about restrictions and requirements for Reserved Instances and the Reserved Instance Marketplace, see [Requirements Checklist for Reserved Instances \(p. 262\)](#).)

As soon as you list your Reserved Instances, they will be included in a list that other AWS customers can view. AWS groups Reserved Instances based on the type of instance being listed, the duration of the term remaining, and the hourly price. This grouping makes it easier for buyers to find the Reserved Instances they want to purchase. From the list, customers choose to purchase the instance that best matches their criteria and decide on the best tradeoff between quoted upfront price and hourly price.

To fulfill a buyer's request, AWS first sells the Reserved Instance with the lowest upfront price in the specified grouping; then it sells the Reserved Instance with the next lowest price, until the buyer's entire purchase order is fulfilled. AWS processes the transaction and transfers ownership of the Reserved Instance to the buyer. Sellers will receive a cash disbursement for their Reserved Instances through a wire transfer directly into their bank account.

When you sell in the Reserved Instance Marketplace, the buyer's ZIP code and country information will be provided to you through a disbursement report. With this information, you will be able to calculate any necessary tax you need to remit to the government. Your business name (as the seller) will also be provided on the purchase invoice of the buyer. AWS charges an administrative fee (12 percent of the total upfront price) for selling in the Reserved Instance Marketplace. For more information, see *Reserved Instance Marketplace* in the [Amazon EC2 Reserved Instances](#) product page.

You retain control of your Reserved Instance until it's sold. When you sell, what you are giving up is the capacity reservation and the discounted recurring fees. You can continue to use your instance after you have sold the reserved capacity, but AWS will now charge you the On-Demand price. The On-Demand price will start from the time that your Reserved Instance was sold. Thus, if you don't want to be charged On-Demand prices for instances that you use, purchase more reserved capacity or terminate your instances when your capacity reservation is sold (or expires).

This topic walks you through the steps to selling in the Reserved Instance Marketplace:

- [Registering as a Seller \(p. 237\)](#)—Register as a seller and specify a bank that has a US address. If you plan on 200 or more transactions or if you plan to sell \$20,000 or more worth of Reserved Instances over the course of a year, you also have to provide tax information.
- [Selling Your Reserved Instances \(p. 240\)](#)—List an active Reserved Instance that has more than one month left in its term. You also must own the Reserved Instance for longer than a month.
- [After Your Reserved Instance Is Sold \(p. 260\)](#)—Find out when your Reserved Instance is sold, and how you get paid.
- [Quick Start: Selling in the Reserved Instance Marketplace Video \(p. 211\)](#)—Pick up the information you need to quickly get started selling in the Reserved Instance Marketplace.

For information about buying Reserved Instances in the Reserved Instance Marketplace, see [Buying Reserved Instances \(p. 212\)](#). For basic information about Reserved Instances, see [Reserved Instances \(p. 193\)](#).

## Important Notes About Selling in the Reserved Instance Marketplace

- To become a seller in the Reserved Instance Marketplace, you must register as a seller, and specify a bank that has a US address. For information, see [Registering as a Seller \(p. 237\)](#).
- Reserved Instances can be sold after they have been active for at least 30 days and when AWS has received the upfront payment.
- You can sell up to \$50,000 in Reserved Instances per year. If you need to sell more Reserved Instances, complete the [Request to Raise Sales Limit on Amazon EC2 Reserved Instances](#) form.
- Amazon EC2 Reserved Instances purchased at a reduced cost resulting from tiered discount cannot be sold in the Reserved Instance Marketplace. For more information about discount pricing tiers, see [Understanding Reserved Instance Pricing Tiers \(p. 200\)](#).


For a checklist that summarizes requirements for working with Reserved Instances and the Reserved Instance Marketplace, see [Requirements Checklist for Reserved Instances \(p. 262\)](#).

## Registering as a Seller

### Topics

- [Your Bank \(p. 238\)](#)
- [Tax Information \(p. 239\)](#)
- [Seller Registration Confirmation \(p. 240\)](#)
- [Sharing Information with the Buyer \(p. 240\)](#)
- [Next Steps \(p. 240\)](#)

To be able to sell in the Reserved Instance Marketplace, your first task is to register as a seller.



If you haven't created an AWS account yet, you need to do this first before you register for Reserved Instance Marketplace. Complete the instructions described in [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#), which provides information about creating your Amazon EC2 account and credentials.

You can access the registration process through the [Reserved Instance Marketplace Seller Registration](#) web page. If you try to create a listing and you have not registered, AWS will direct you to this seller registration page.

Registering means providing the name of your business, information about your bank, and your business's tax identification number. Usually, you only have to provide your information once. However, you can update personal and banking information through the [Reserved Instance Marketplace Seller Registration](#)



## Amazon Elastic Compute Cloud User Guide Selling in the Reserved Instance Marketplace

web page. Log in to AWS using the account you used when you first registered and the page will send you directly to the personal and banking information pages.

amazon web services

Welcome [Sign Out](#)

Account Information — Manage Bank Account — Confirmation

### Reserved Instance Marketplace

Account Information

Enter your business information below so that Amazon.com may contact you regarding your Selling account. Information about our customers is an important part of our business, and we are not in the business of selling it to others. We share customer information only as described in the [Privacy Notice](#).

**Business Information**

**Business Name:**   
Name of the Seller of Record that will appear on any legal documentation Amazon.com sends you.

[Continue](#)

© 2012 Amazon Web Services LLC or its affiliates. All rights reserved.  
An [amazon.com](#) company

## Your Bank

AWS must have your bank information in order to disburse funds collected when you sell your Reserved Instance. The bank you specify must have a US address.

On the **Manage Bank Account** page, provide the following information about the bank through which you will receive payment:

- Bank account holder name
- Routing number
- Account number
- Bank account type

### Note

If you are using a corporate bank account, you will be prompted to send via fax (1-206-765-3424) the information about the bank account.

amazon web services

Welcome [Sign Out](#)

Account Information — **Manage Bank Account** — Confirmation

### Add Bank Account

We register your bank account for the deposit of your sales proceeds. Note that it must be a U.S. bank account. Bank account information is subject to the [Privacy Notice](#).

John Customer  
John Customer  
Customer ID: 1234

12-123456789 1234  
0123456789

FOR DEPOSIT ONLY

FOR

123456789 1234567891234

Bank Account Holder Name: \*

Routing Number: \*   
(9 digits)

Account Number: \*   
(3-17 digits)

\* Required

[Back](#) [Continue](#)

© 2012 Amazon Web Services LLC or its affiliates. All rights reserved.  
An [amazon.com](#) company

You can change the default bank account through which you receive disbursements. Just go to the [Reserved Instance Marketplace Seller Registration](#) web page using the account you used when you first registered and the page will send you directly to the personal and banking information pages.

After you have completed the registration, AWS verifies your bank account and sets it as the default bank. You will not be able to receive disbursements until AWS has verified your account with the bank. Verification with the bank can take up to two weeks, so if your account is a new one, you will not get the disbursement as a result of a sale for up to two weeks. For an established account, it will usually take about two days for disbursements to complete.

## Tax Information

Your sale of Reserved Instances on the Reserved Instance Marketplace might be subject to a transactional tax, such as sales tax or value-added tax. You should check with your business's tax, legal, finance, or accounting department to determine if transaction-based taxes are applicable. You are responsible for collecting and sending the transaction-based taxes to the appropriate tax authority.

As part of the seller registration process, you have the option of completing a tax interview. We encourage you to complete this process if any of the following apply:

- You want AWS to generate a Form 1099-K.
- You anticipate having either 200 or more transactions or \$20,000 or more in sales of Reserved Instances in a calendar year. A transaction can involve one or more Reserved Instances. If you choose to skip this step during registration, and later you reach transaction 199, you will get a message saying, "You have reached the transaction limit for pre-tax. Please complete the tax interview at [http://portal.aws.amazon.com/ec2/ri/seller\\_registration?action=taxInterview](http://portal.aws.amazon.com/ec2/ri/seller_registration?action=taxInterview)."
- You are a non-US seller. In this case, you must electronically complete Form W-8BEN.

If you complete the tax interview, the tax information you enter will differ depending on whether your business is a US or non-US legal entity. If you are a US seller, you must provide AWS with your tax identification number along with your business contact information.

The screenshot shows the Amazon Web Services interface for a 'Tax Interview'. At the top left is the Amazon Web Services logo. On the right, it says 'Welcome Nina Ingram, Sign Out'. The main heading is 'Tax Interview'. Below this is a progress bar showing 'Progress: 10'. Underneath the progress bar is a 'Getting started' section with the following text: 'If you are unsure about how to answer a question, refer to **Help** by either selecting the ? next to the entry field for pop-up Help, or by reviewing the application help available at any point in the interview process. We will now begin by collecting your tax information.' At the bottom of the page, there are 'Back' and 'Continue' buttons. The footer contains the copyright notice: '© 2012 Amazon Web Services LLC or its affiliates. All rights reserved. An amazon.com company'.

After you complete the tax registration process, AWS will file Form 1099-K, and you will receive a copy of it through the US mail on or before January 31 in the year after the year that your tax account reaches the threshold levels.

As you fill out the tax interview, keep in mind the following:

- Information provided by AWS, including the information in this topic, does not constitute tax, legal, or other professional advice. To find out how the IRS reporting requirements might affect your business, or if you have other questions, please contact your tax, legal, or other professional advisor.
- In order to fulfill the IRS reporting requirements as efficiently as possible, answer all questions and enter all information requested during the interview.

- Check your answers. Avoid misspellings or entering incorrect tax identification numbers. They can result in an invalidated tax form.

For more information about IRS requirements and the Form 1099-K, go to the [IRS website](#).

## Seller Registration Confirmation

After we receive your completed seller registration, you will get an email confirming your registration and informing you that you can get started selling in the Reserved Instance Marketplace.

## Sharing Information with the Buyer

When you sell your Reserved Instances in the Reserved Instance Marketplace, AWS will share your company's legal name on the buyer's statement to comply with US regulations. In addition, if the buyer calls AWS customer service because the buyer needs to contact you for an invoice or for some other tax-related reason, AWS may need to provide the buyer with your email address so that the buyer can contact you directly.

For similar reasons, the buyer's ZIP code and country information will be provided to the seller in the disbursement report. As a seller you might need this information to accompany any necessary transaction taxes that you remit to the government (such as sales tax and value-added tax).

AWS cannot offer tax advice, but if your tax specialist determines that you need specific additional information, contact AWS Customer Support.

## Next Steps

After you have successfully registered as a seller in the Reserved Instance Marketplace, you can begin selling your Reserved Instances. Selling in the Reserved Instance Marketplace requires the following steps:

1. Deciding which Reserved Instances you want to sell.

Identify the active Reserved Instances that you want to sell and select the upfront price at which you want to sell them. For more information, see [Pricing Your Reserved Instances \(p. 241\)](#).

2. Listing your Reserved Instances.

Include your Reserved Instances in the Reserved Instance Marketplace listings. For more information, see [Listing Your Reserved Instance \(p. 242\)](#).

3. Viewing your listings.

You can monitor your Reserved Instances and view your listings. For more information, see [Obtaining Information About Your Reserved Instances \(p. 222\)](#).

4. Canceling and changing your listings.

You can change your listings by first canceling and then relisting. For more information, see [Canceling and Changing Your Listings \(p. 250\)](#).

For information about selling your Reserved Instances, see [Selling Your Reserved Instances \(p. 240\)](#).

## Selling Your Reserved Instances

### Topics

- [Pricing Your Reserved Instances \(p. 241\)](#)
- [Listing Your Reserved Instance \(p. 242\)](#)
- [Canceling and Changing Your Listings \(p. 250\)](#)

This section walks you through how to list and sell your Reserved Instances in the Reserved Instance Marketplace. If you haven't registered as a seller yet, you must do so first. For information, see [Selling Your Reserved Instances \(p. 240\)](#).

As a registered seller, you can choose to sell one or more of your Reserved Instances, and you can choose to sell all of them in one listing. In addition, you can list any type of Reserved Instance—including any configuration of instance type, platform, region, and Availability Zone—as long as the following requirements are met:

- You have already paid the upfront cost for the Reserved Instance you are listing.

This means that you can list your Reserved Instance when it is in the *Active* state. However, keep in mind that a Reserved Instance can be in the *Active* state before AWS actually receives your payment. If this is the case, the Reserved Instance Marketplace will not allow you to list your Reserved Instance until the payment for the upfront fee is collected.

- You have owned the Reserved Instance for at least a month.
- There is at least a month remaining in the term of the Reserved Instance you are listing.

You can list the remainder of the Reserved Instance term rounded down to the nearest month. For example, if you have 9 months and 13 days remaining on your Reserved Instance, you can list 9 months for your Reserved Instance.

- The Reserved Instance you are selling is not a discounted or private Reserved Instance. You cannot list these types of Reserved Instances in the Reserved Instance Marketplace.

For a checklist that summarizes requirements for working with Reserved Instances and the Reserved Instance Marketplace, see [Requirements Checklist for Reserved Instances \(p. 262\)](#).

To get details about your existing Reserved Instances, you can use any of the following tools:

- The **Reserved Instances** page in the **EC2** console of the [AWS Management Console \(p. 197\)](#).
- The `ec2-describe-reserved-instances` CLI command.
- The `DescribeReservedInstances` API action.

For more information about Reserved Instances, see [Reserved Instances \(p. 193\)](#).

## Pricing Your Reserved Instances

When you're a seller on the Reserved Instance Marketplace, the upfront fee is the only fee that you can specify for the Reserved Instance. The upfront fee is the one-time fee that the buyer pays when the buyer purchases a Reserved Instance. The seller cannot specify the usage fee or the recurring fee. The buyer of your Reserved Instance will be paying the same usage or recurring fees that were set when the Reserved Instances were originally purchased. This usage fee applies to Light Utilization and Medium Utilization Reserved Instances, and the fee differs depending on which Reserved Instance offering you're using. The recurring fee is the hourly fee that you pay for Heavy Utilization, whether or not you're using the Reserved Instance. For more information about the availability of Reserved Instances offerings based on utilization, see [Choosing Reserved Instances Based on Your Usage Plans \(p. 199\)](#).

### Setting a Pricing Schedule

You can set different upfront fees (prices) that are based on when your Reserved Instance sells. You can specify one price if the Reserved Instance sells immediately, and you can specify another price if the Reserved Instance were to sell in any subsequent month. Since the value of Reserved Instances decreases over time, by default, AWS will set prices to decrease linearly—that is, the price drops in equal increments month over month. You can choose to set the prices differently.

For example, if your Reserved Instance has nine months of its term remaining, you can specify the amount you would accept if a customer were to purchase that Reserved Instance with nine months remaining, and you could set another price with five months remaining, and yet another price with one month remaining.

## Listing Your Reserved Instance

You can list the Reserved Instances you want to sell in the Reserved Instance Marketplace by using the AWS Management Console, the Amazon EC2 CLI, or the Amazon EC2 API.

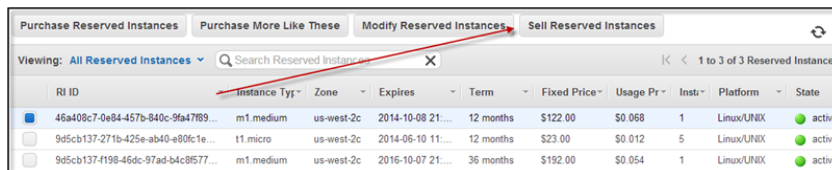
### AWS Management Console

#### To list a Reserved Instance in the Reserved Instance Marketplace

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Reserved Instances** in the **Navigation** pane.

The **Reserved Instances** page displays a list of your account's instances.

3. Select the Reserved Instances you want to list on the marketplace, and click **Sell Reserved Instances**.

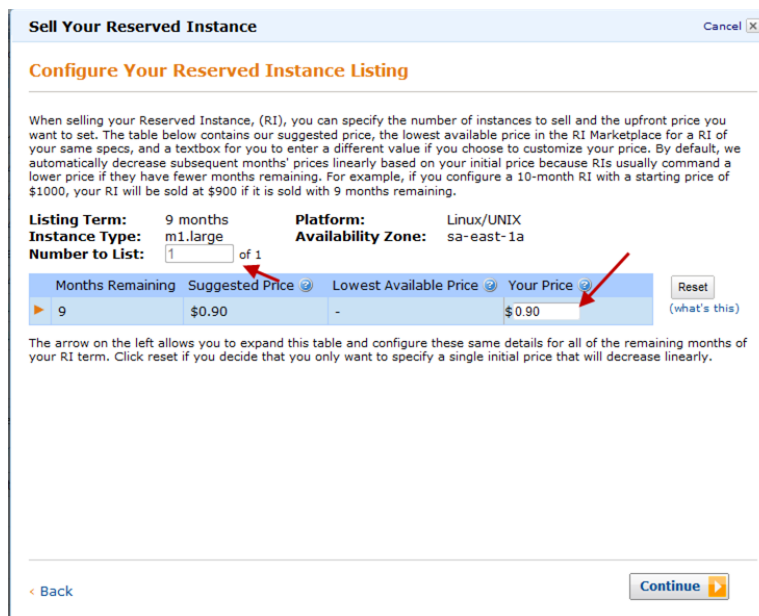


RI ID	Instance Type	Zone	Expires	Term	Fixed Price	Usage Pr	Instr	Platform	State
46a408c7-0e84-457b-840c-9fa4789...	m1.medium	us-west-2c	2014-10-08 21:...	12 months	\$122.00	\$0.068	1	Linux/UNIX	active
9d5cb137-271b-425e-ab40-e80c1e...	t1.micro	us-west-2c	2014-06-10 11:...	12 months	\$23.00	\$0.012	5	Linux/UNIX	active
9d5cb137-4198-46dc-97ad-b4c8577...	m1.medium	us-west-2c	2016-10-07 21:...	36 months	\$192.00	\$0.054	1	Linux/UNIX	active

#### Note

If you have not completed the Reserved Instance Marketplace seller registration process, you will be prompted to complete this process now. For information about the seller registration process, see [Registering as a Seller \(p. 237\)](#).

4. In the **Configure Your Reserved Instance Listing** page, for **Number to List**, set the number of instances to sell and, for **Your Price**, set the upfront price for the remaining time period.



**Sell Your Reserved Instance** Cancel

### Configure Your Reserved Instance Listing

When selling your Reserved Instance, (RI), you can specify the number of instances to sell and the upfront price you want to set. The table below contains our suggested price, the lowest available price in the RI Marketplace for a RI of your same specs, and a textbox for you to enter a different value if you choose to customize your price. By default, we automatically decrease subsequent months' prices linearly based on your initial price because RIs usually command a lower price if they have fewer months remaining. For example, if you configure a 10-month RI with a starting price of \$1000, your RI will be sold at \$900 if it is sold with 9 months remaining.

**Listing Term:** 9 months      **Platform:** Linux/UNIX  
**Instance Type:** m1.large      **Availability Zone:** sa-east-1a  
**Number to List:** 1 of 1

Months Remaining	Suggested Price	Lowest Available Price	Your Price
9	\$0.90	-	\$0.90

The arrow on the left allows you to expand this table and configure these same details for all of the remaining months of your RI term. Click reset if you decide that you only want to specify a single initial price that will decrease linearly.

< Back Continue >

You can see how the value of your Reserved Instance will change over the remainder of the term by clicking the arrow on the left of the **Months Remaining** column. By default, AWS sets the price to decrease linearly. This means the price drops by equal increments each month.

5. If you are an advanced user and you want to customize the pricing, you can enter different values for the subsequent months. To return to the default linear price drop, click the **Reset** button. Click **Continue** when you are finished configuring your listing.

**Sell Your Reserved Instance** Cancel

### Configure Your Reserved Instance Listing

When selling your Reserved Instance, (RI), you can specify the number of instances to sell and the upfront price you want to set. The table below contains our suggested price, the lowest available price in the RI Marketplace for a RI of your same specs, and a textbox for you to enter a different value if you choose to customize your price. By default, we automatically decrease subsequent months' prices linearly based on your initial price because RIs usually command a lower price if they have fewer months remaining. For example, if you configure a 10-month RI with a starting price of \$1000, your RI will be sold at \$900 if it is sold with 9 months remaining.

**Listing Term:** 9 months      **Platform:** Linux/UNIX  
**Instance Type:** m1.large      **Availability Zone:** sa-east-1a  
**Number to List:** 1 of 1

Months Remaining	Suggested Price	Lowest Available Price	Your Price
9 months	0.90	-	\$ 0.90
8 months	0.80	-	\$ 0.80
7 months	0.70	-	\$ 0.70
6 months	0.60	-	\$ 0.60
5 months	0.50	-	\$ 0.50
4 months	0.40	-	\$ 0.40
3 months	0.30	-	\$ 0.30

Reset (what's this)

The arrow on the left allows you to expand this table and configure these same details for all of the remaining months of your RI term. Click reset if you decide that you only want to specify a single initial price that will decrease linearly.

< Back Continue

6. When you are satisfied with the details of your listing as displayed by the **Confirm Your Reserved Instance Listing** page, click **List Reserved Instance**.

You will get a confirmation that your listing is being processed.

**Sell Your Reserved Instance** Cancel

It may take a few minutes before your Reserved Instance listing changes from pending to active. When active, your Reserved Instance will appear in the purchase search results grouped by its unique attributes, such as Availability Zone. You can use your Reserved Instance until it is sold. When your Reserved Instance listing sells, any running instance will now be charged at the On-Demand rate.

### Interested in Purchasing New Reserved Instances?

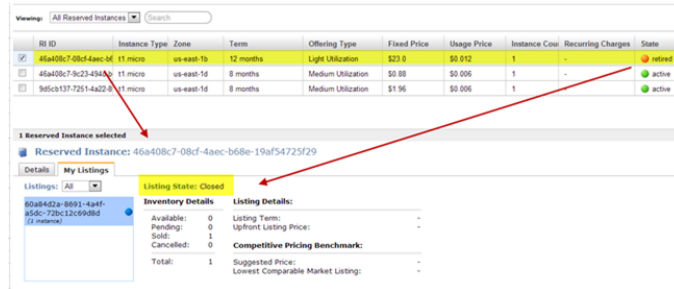
**Find the Optimal Reserved Instance**  
Visit the Reserved Instance Marketplace to purchase a new Reserved Instance to replace the one you are selling. As usual, you can simply make a low, one-time payment to reserve compute capacity, and in turn, receive a significant discount on the hourly charge for that instance.

**Right Sizing Your Instances**  
Are you fully utilizing your instance's CPU? Choosing the correct instance sizes to meet your needs is a great way to reduce your costs. We provide free tools to monitor the CPU for your instances.

[> Learn more about Reserved Instances](#)      [> Go to CloudWatch to see your utilization](#)

Close

7. To view the details of your Reserved Instance listing, on the **Reserved Instances** page, select the Reserved Instance you want to view, and click the **My Listings** tab.



## Amazon EC2 CLI

### To list a Reserved Instance in the Reserved Instance Marketplace

1. Get a list of your Reserved Instances by calling `ec2-describe-reserved-instances`.

```
PROMPT> ec2-describe-reserved-instances --headers
```

Amazon EC2 returns output similar to the following:

```
PROMPT> ec2-describe-reserved-instances --headers
Type ReservedInstancesId AvailabilityZone InstanceType ProductDescription
Duration FixedPrice UsagePrice InstanceCount Start State Currency Instan
ceTenancy OfferingType
RESERVEDINSTANCES f127bd27-9f30-41d3-bf45-9af45example sa-east-1a m1.large
Linux/UNIX 10m 1.0 0.0 1 2012-08-22T21:41:51+0000 active USD default Medium
Utilization
RESERVEDINSTANCES 1ba8e2e3-d20d-44ec-b202-fcb6aexample sa-east-1b m1.small
Linux/UNIX 10m 1.2 0.032 3 2012-08-21T14:02:00+0000 retired USD default
Medium Utilization
RESERVEDINSTANCES 4357912c-6f69-4966-a407-6f0cbexample sa-east-1b m1.small
Linux/UNIX 10m 1.2 0.032 3 2012-08-21T14:02:00+0000 active USD default
Medium Utilization
RESERVEDINSTANCES 4357912c-d032-4a97-9b49-5eb3aexample sa-east-1b m1.small
Linux/UNIX 10m 1.2 0.032 1 2012-08-21T14:02:00+0000 retired USD default
Medium Utilization
...
```

Select the Reserved Instance ID of the Reserved Instance you want to list in the Reserved Instance Marketplace.

2. Specify the Reserved Instance ID of the Reserved Instance you want to list and call `ec2-create-reserved-instances-listing`. You have to specify the following required parameters:
  - Reserved Instance ID
  - Instance count
  - MONTH:PRICE

The command should look like this example:



```
PROMPT> ec2-create-reserved-instances-listing --reserved-instance b847fa93-c736-4eae-bca1-3147example --instance-count 1 05:01.20 04:01.00 01:00.75 --headers
```

Amazon EC2 returns output similar to the following:

```
PROMPT>LISTING 2a0ff720-f62e-4824-8ed1-7dd0aexample b847fa93-c736-4eae-bca1-e3147example Wed Aug 29 13:59:11 PDT 2012 Wed Aug 29 13:59:11 PDT 2012 active
active
INSTANCE-COUNT available 1
INSTANCE-COUNT sold 0
INSTANCE-COUNT cancelled 0
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 5 $1.2
PRICE-SCHEDULE 4 $1.0
PRICE-SCHEDULE 3 $1.0
PRICE-SCHEDULE 2 $1.0
PRICE-SCHEDULE 1 $0.75
```

3. To view the details of your Reserved Instance listing, run `ec2-describe-reserved-instances-listings` with the listing ID `095c0e18-c9e6-4692-97e5-653e0example`.

```
PROMPT> ec2-describe-reserved-instances-listings 095c0e18-c9e6-4692-97e5-653e0example
```

Amazon EC2 returns output similar to the following:

```
PROMPT> ec2-describe-reserved-instances-listings 095c0e18-c9e6-4692-97e5-653e0example
Type ReservedInstancesListingId ReservedInstancesId CreateDate UpdateDate
Status StatusMessage
LISTING 095c0e18-c9e6-4692-97e5-653e0example b847fa93-c736-4eae-bca1-e3147example Tue Aug 28 18:21:07 PDT 2012 Tue Aug 28 18:21:07 PDT 2012 active
active
INSTANCE-COUNT available 1
INSTANCE-COUNT sold 0
INSTANCE-COUNT cancelled 0
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 5 $1.2
PRICE-SCHEDULE 4 $1.2
PRICE-SCHEDULE 3 $1.2
PRICE-SCHEDULE 2 $1.2
PRICE-SCHEDULE 1 $1.2
```

## Amazon EC2 API

### To list a Reserved Instance in the Reserved Instance Marketplace

1. Get a list of your Reserved Instances by calling `DescribeReservedInstances`.



```
https://ec2.amazonaws.com/
?Action=DescribeReservedInstances
&AUTHPARAMS
```

Following is an example response.

```
<DescribeReservedInstancesResponse xmlns='http://ec2.amazonaws.com/doc/2012-08-15/'>
  <requestId>ebe3410a-8f37-441d-ae11-2e78eexample</requestId>
  <reservedInstancesSet>
    <item>
      <reservedInstancesId>f127bd27-cee4-443a-a76b-a5af9example</reservedInstancesId>
      <instanceType>m1.large</instanceType>
      <availabilityZone>us-east-1a</availabilityZone>
      <start>2012-08-07T15:19:31.071Z</start>
      <duration>31536000</duration>
      <fixedPrice>276.0</fixedPrice>
      <usagePrice>0.156</usagePrice>
      <instanceCount>5</instanceCount>
      <productDescription>Linux/UNIX</productDescription>
      <state>active</state>
      <instanceTenancy>default</instanceTenancy>
      <currencyCode>USD</currencyCode>
      <offeringType>Light Utilization</offeringType>
      <recurringCharges/>
    </item>
  </reservedInstancesSet>
</DescribeReservedInstancesResponse>
```

Note the Reserved Instance ID of the Reserved Instance that you want to list in the Reserved Instance Marketplace.

2. Create a listing for three Reserved Instances from Reserved Instance ID *f127bd27-cee4-443a-a76b-a5af9example* and specify the following pricing schedule.

Term (remaining months)	11	10	9	8	7	6	5	4	3	2	1
Price specified for period	2.5			2.0			1.5		0.7		0.1
Price	2.5	2.5	2.5	2.0	2.0	2.0	1.5	1.5	0.7	0.7	0.1

The call should look like this example:

```
https://ec2.amazonaws.com/?Action=CreateReservedInstancesListing
&ClientToken=myIdempToken1
&ReservedInstancesId=f127bd27-cee4-443a-a76b-a5af9example
&InstanceCount=3
&PriceSchedules.0.Price=2.5&PriceSchedules.0.Term=11
&PriceSchedules.1.Price=2.0&PriceSchedules.1.Term=8
&PriceSchedules.2.Price=1.5&PriceSchedules.2.Term=5
&PriceSchedules.3.Price=0.7&PriceSchedules.3.Term=3
&PriceSchedules.4.Price=0.1&PriceSchedules.4.Term=1
```

&AUTHPARAMS

Following is an example response.

```
<CreateReservedInstancesListingResponse>
<requestId>a42481af-335a-4e9e-b291-bd18dexample</requestId>
<reservedInstancesListingsSet>
  <item>
    <reservedInstancesListingId>5ec28771-05ff-4b9b-aa31-
9e57dexample</reservedInstancesListingId>
    <reservedInstancesId>f127bd27-cee4-443a-a76b-a5af9example</reserved
InstancesId>
    <createDate>2012-08-30T17:11:09.449Z</createDate>
    <updateDate>2012-08-30T17:11:09.468Z</updateDate>
    <status>active</status>
    <statusMessage>active</statusMessage>
    <instanceCounts>
      <item>
        <state>Available</state>
        <instanceCount>3</instanceCount>
      </item>
      <item>
        <state>Sold</state>
        <instanceCount>0</instanceCount>
      </item>
      <item>
        <state>Cancelled</state>
        <instanceCount>0</instanceCount>
      </item>
      <item>
        <state>Pending</state>
        <instanceCount>0</instanceCount>
      </item>
    </instanceCounts>
    <priceSchedules>
      <item>
        <term>11</term>
        <price>2.5</price>
        <currencyCode>USD</currencyCode>
        <active>true</active>
      </item>
      <item>
        <term>10</term>
        <price>2.5</price>
        <currencyCode>USD</currencyCode>
        <active>false</active>
      </item>
      <item>
        <term>9</term>
        <price>2.5</price>
        <currencyCode>USD</currencyCode>
        <active>false</active>
      </item>
      <item>
        <term>8</term>
```

```
        <price>2.00</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>7</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>6</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>5</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>4</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>3</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>2</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>1</term>
        <price>0.1</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
</priceSchedules>
<tagSet/>
<clientToken>listRI1</clientToken>
</item>
</reservedInstancesListingsSet>
</CreateReservedInstancesListingResponse>
```

3. To view the details of your Reserved Instance listing, run `DescribeReservedInstancesListings`.

The command should look like this example:

```
http://ec2.amazonaws.com/?Action=DescribeReservedInstancesListings
&AUTHPARAMS
```

Following is an example response.

```
<DescribeReservedInstancesListingsResponse>
  <requestId>cec5c904-8f3a-4de5-8f5a-ff7f9example</requestId>
  <reservedInstancesListingsSet>
    <item>
      <reservedInstancesListingId>5ec28771-05ff-4b9b-aa31-
9e57dexample</reservedInstancesListingId>
      <reservedInstancesId>f127bd27-cee4-443a-a76b-a5af9example</re
servedInstancesId>
      <createDate>2012-08-30T17:11:09.449Z</createDate>
      <updateDate>2012-08-30T17:11:09.468Z</updateDate>
      <status>active</status>
      <statusMessage>active</statusMessage>
      <instanceCounts>
        <item>
          <state>Available</state>
          <instanceCount>3</instanceCount>
        </item>
        <item>
          <state>Sold</state>
          <instanceCount>0</instanceCount>
        </item>
        <item>
          <state>Cancelled</state>
          <instanceCount>0</instanceCount>
        </item>
        <item>
          <state>Pending</state>
          <instanceCount>0</instanceCount>
        </item>
      </instanceCounts>
      <priceSchedules>
        <item>
          <term>11</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>true</active>
        </item>
        <item>
          <term>10</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>>false</active>
        </item>
        <item>
          <term>9</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>>false</active>
        </item>
        <item>
          <term>8</term>
          <price>2.0</price>
        </item>
      </priceSchedules>
    </item>
  </reservedInstancesListingsSet>
</DescribeReservedInstancesListingsResponse>
```

```
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>7</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>6</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>5</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>4</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>3</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>2</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>1</term>
        <price>0.1</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
</priceSchedules>
<tagSet/>
<clientToken>listRI1</clientToken>
</item>
</reservedInstancesListingsSet>
</DescribeReservedInstancesListingsResponse>
```

## Canceling and Changing Your Listings

After listing your Reserved Instances in the Reserved Instance Marketplace, you can manage your listing using the [AWS Management Console \(p. 251\)](#), the [Amazon EC2 CLI \(p. 252\)](#), or the [Amazon EC2 API \(p. 253\)](#).

You can view details of your listing. You also can cancel your listing, as long as it hasn't been purchased yet.

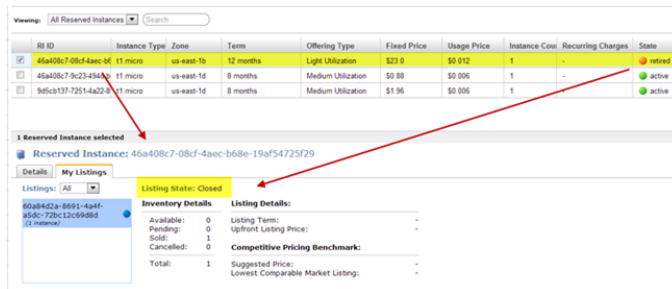
Currently, you cannot modify your listing directly. However, you can change your listing by first canceling it and then creating another listing with new parameters.

In this section, we will show you how to perform the tasks of viewing, canceling and changing your Reserved Instance Marketplace listings, using the console, the CLI, and the API.

## AWS Management Console

### To view your listing

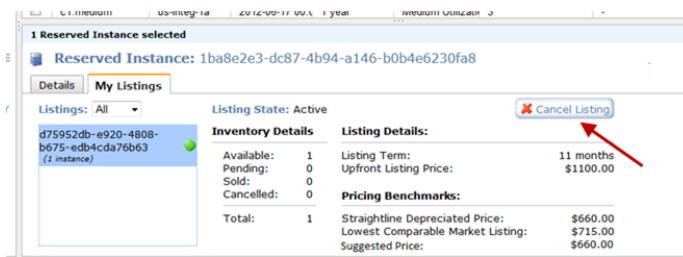
1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Reserved Instances** in the **Navigation** pane.
3. Right-click your Reserved Instance.
4. Click the **My Listings** tab. You will see details about the Reserved Instance listing you selected.



### To cancel your listing

You can cancel your listing at any time, as long as it's in the *active* state. You cannot cancel the listing if it's already matched or being processed for a sale. If some of the instances in your listing are matched and you cancel the listing, only the remaining unmatched instances will be removed from the listing.

1. Go to the **Reserved Instances** page in the Amazon EC2 console, and right-click your Reserved Instance.
2. On the **My Listings** tab, click **Cancel Listing**.



### To change your listing

Currently, you cannot modify your listing directly. However, you can change your listing by first canceling it and then creating another listing with new parameters.

1. Cancel your active Reserved Instance listing. For more information, see the previous procedure.
2. Create a new listing. For more information, see [Listing Your Reserved Instance \(p. 242\)](#).

## Amazon EC2 CLI

### To view your listing

- Run `ec2-describe-reserved-instances-listings` to get details about your listing.

```
PROMPT> ec2-describe-reserved-instances-listings
```

Amazon EC2 returns output similar to the following:

```
PROMPT> ec2-describe-reserved-instances-listings
Type ReservedInstancesListingId ReservedInstancesId CreateDate UpdateDate
Status StatusMessage
LISTING 615d8a10-8224-4c19-ba7d-b9aa0example 1ba8e2e3-d20d-44ec-b202-
fcb6aexample Wed Aug 22 09:02:58 PDT 2012 Wed Aug 22 14:24:26 PDT 2012 can
cancelled cancelled
INSTANCE-COUNT available 0
INSTANCE-COUNT sold 0
INSTANCE-COUNT cancelled 1
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 10 $1.2
PRICE-SCHEDULE 9 $1.08
PRICE-SCHEDULE 8 $0.96
PRICE-SCHEDULE 7 $0.84
PRICE-SCHEDULE 6 $0.72
PRICE-SCHEDULE 5 $0.6
PRICE-SCHEDULE 4 $0.48
PRICE-SCHEDULE 3 $0.36
PRICE-SCHEDULE 2 $0.24
PRICE-SCHEDULE 1 $0.12
LISTING d5fa5166-83c3-40e4-abb2-b7298example 1ba8e2e3-d20d-44ec-b202-
fcb6aexample Wed Aug 22 14:31:55 PDT 2012 Wed Aug 22 14:42:40 PDT 2012 closed
closed
INSTANCE-COUNT available 0
INSTANCE-COUNT sold 1
INSTANCE-COUNT cancelled 0
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 10 $0.9
PRICE-SCHEDULE 9 $0.81
PRICE-SCHEDULE 8 $0.72
PRICE-SCHEDULE 7 $0.63
PRICE-SCHEDULE 6 $0.54
PRICE-SCHEDULE 5 $0.45
PRICE-SCHEDULE 4 $0.36
PRICE-SCHEDULE 3 $0.27
PRICE-SCHEDULE 2 $0.18
PRICE-SCHEDULE 1 $0.09
....
LISTING 095c0e18-c9e6-4692-97e5-653e0example b847fa93-c736-4eae-bca1-
e3147example Tue Aug 28 18:21:07 PDT 2012 Tue Aug 28 18:21:07 PDT 2012 active
active
INSTANCE-COUNT available 1
INSTANCE-COUNT sold 0
```

```
INSTANCE-COUNT cancelled 0
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 5 $1.2
PRICE-SCHEDULE 4 $1.2
PRICE-SCHEDULE 3 $1.2
PRICE-SCHEDULE 2 $1.2
PRICE-SCHEDULE 1 $1.2
```

### To cancel your listing

You can cancel your listing at any time, as long as it's in the *active* state. You cannot cancel the listing if it's already matched or being processed for a sale.

- Run `ec2-cancel-reserved-instances-listing` to cancel your listing.

```
PROMPT> ec2-cancel-reserved-instances-listing 095c0e18-c9e6-4692-97e5-653e0example
```

Amazon EC2 returns output similar to the following:

```
PROMPT> ec2-cancel-reserved-instances-listing
Type ReservedInstancesListingId ReservedInstancesId CreateDate UpdateDate
Status StatusMessage
LISTING 095c0e18-c9e6-4692-97e5-653e0example b847fa93-c736-4eae-bca1-
e3147example Tue Aug 28 18:21:07 PDT 2012 Tue Aug 28 18:21:07 PDT 2012 can
celled cancelled
INSTANCE-COUNT available 0
INSTANCE-COUNT sold 0
INSTANCE-COUNT cancelled 1
INSTANCE-COUNT pending 0
PRICE-SCHEDULE 5 $1.2
PRICE-SCHEDULE 4 $1.2
PRICE-SCHEDULE 3 $1.2
PRICE-SCHEDULE 2 $1.2
PRICE-SCHEDULE 1 $1.2
```

### To change your listing

Currently, you cannot modify your listing directly. However, you can change your listing by first canceling it and then creating another listing with new parameters.

1. Cancel your active Reserved Instance listing. For more information, see the previous procedure.
2. Create a new listing. For more information, see [Listing Your Reserved Instance \(p. 242\)](#).

## Amazon EC2 API

### To view your listing

- Call `DescribeReservedInstancesListings` to get details about your listing.

The command should look like this example:



```
http://ec2.amazonaws.com/?Action=DescribeReservedInstancesListings
&AUTHPARAMS
```

Following is an example response.

```
<DescribeReservedInstancesListingsResponse>
  <requestId>cec5c904-8f3a-4de5-8f5a-ff7f9example</requestId>
  <reservedInstancesListingsSet>
    <item>
      <reservedInstancesListingId>5ec28771-05ff-4b9b-aa31-
9e57dexample</reservedInstancesListingId>
      <reservedInstancesId>f127bd27-cee4-443a-a76b-a5af9example</re
servedInstancesId>
      <createDate>2012-08-30T17:11:09.449Z</createDate>
      <updateDate>2012-08-30T21:00:42.300Z</updateDate>
      <status>active</status>
      <statusMessage>active</statusMessage>
      <instanceCounts>
        <item>
          <state>Available</state>
          <instanceCount>2</instanceCount>
        </item>
        <item>
          <state>Sold</state>
          <instanceCount>1</instanceCount>
        </item>
        <item>
          <state>Cancelled</state>
          <instanceCount>0</instanceCount>
        </item>
        <item>
          <state>Pending</state>
          <instanceCount>0</instanceCount>
        </item>
      </instanceCounts>
      <priceSchedules>
        <item>
          <term>11</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>true</active>
        </item>
        <item>
          <term>10</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>false</active>
        </item>
        <item>
          <term>9</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>false</active>
        </item>
        <item>
          <term>8</term>
          <price>2.0</price>
        </item>
      </priceSchedules>
    </item>
  </reservedInstancesListingsSet>
</DescribeReservedInstancesListingsResponse>
```

```
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>7</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>6</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>5</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>4</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>3</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>2</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>1</term>
        <price>0.1</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
</priceSchedules>
<tagSet/>
<clientToken>listRI1</clientToken>
</item>
</reservedInstancesListingsSet>
</DescribeReservedInstancesListingsResponse>
```

### To cancel your listing

You can cancel your listing at any time, as long as it's in the *active* state. You cannot cancel the listing if it's already matched or being processed for a sale.

- Run `CancelReservedInstancesListing` to cancel your listing `5ec28771-05ff-4b9b-aa31-9e57dexample` and remove it from the Reserved Instance Marketplace.

The command should look like this example:

```
https://ec2.amazonaws.com/?Action=CancelReservedInstancesListing
&ReservedInstancesListingId.0=5ec28771-05ff-4b9b-aa31-9e57dexample
&AUTHPARAMS
```

Following is an example response.

```
<CancelReservedInstancesListingResponse>
  <requestId>bec2cf62-98ef-434a-8a15-886fcexample</requestId>
  <reservedInstancesListingsSet>
    <item>
      <reservedInstancesListingId>5ec28771-05ff-4b9b-aa31-
9e57dexample</reservedInstancesListingId>
      <reservedInstancesId>f127bd27-cee4-443a-a76b-a5af9example</re
servedInstancesId>
      <createDate>2012-08-30T17:11:09.449Z</createDate>
      <updateDate>2012-08-31T14:12:23.468Z</updateDate>
      <status>cancelled</status>
      <statusMessage>cancelled</statusMessage>
      <instanceCounts>
        <item>
          <state>Available</state>
          <instanceCount>0</instanceCount>
        </item>
        <item>
          <state>Sold</state>
          <instanceCount>1</instanceCount>
        </item>
        <item>
          <state>Cancelled</state>
          <instanceCount>2</instanceCount>
        </item>
        <item>
          <state>Pending</state>
          <instanceCount>0</instanceCount>
        </item>
      </instanceCounts>
      <priceSchedules>
        <item>
          <term>1</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>>false</active>
        </item>
        <item>
          <term>10</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>>false</active>
        </item>
        <item>
          <term>9</term>
          <price>2.5</price>
        </item>
      </priceSchedules>
    </item>
  </reservedInstancesListingsSet>
</CancelReservedInstancesListingResponse>
```

```
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>8</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>7</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>6</term>
        <price>2.0</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>5</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>4</term>
        <price>1.5</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>3</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>2</term>
        <price>0.7</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
    <item>
        <term>1</term>
        <price>0.1</price>
        <currencyCode>USD</currencyCode>
        <active>>false</active>
    </item>
</priceSchedules>
<tagSet/>
<clientToken>listRI1</clientToken>
</item>
</reservedInstancesListingsSet>
</CancelReservedInstancesListingResponse>
```

In the response to the Cancel request, you see the canceled listing. Another way to see that canceled listing is by calling `DescribeReservedInstancesListings`. The request will look like this example:

```
http://ec2.amazonaws.com/?Action=DescribeReservedInstancesListings
&AUTHPARAMS
```

The response will look like the following example:

```
<DescribeReservedInstancesListingsResponse>
  <requestId>bec2cf62-98ef-434a-8a15-367c0example</requestId>
  <reservedInstancesListingsSet>
    <item>
      <reservedInstancesListingId>5ec28771-05ff-4b9b-aa31-
9e57dexample</reservedInstancesListingId>
      <reservedInstancesId>f127bd27-cee4-443a-a76b-a5af9example</re
servedInstancesId>
      <createDate>2012-08-30T17:11:09.449Z</createDate>
      <updateDate>2012-08-31T14:12:23.468Z</updateDate>
      <status>cancelled</status>
      <statusMessage>cancelled</statusMessage>
      <instanceCounts>
        <item>
          <state>Available</state>
          <instanceCount>0</instanceCount>
        </item>
        <item>
          <state>Sold</state>
          <instanceCount>1</instanceCount>
        </item>
        <item>
          <state>Cancelled</state>
          <instanceCount>2</instanceCount>
        </item>
        <item>
          <state>Pending</state>
          <instanceCount>0</instanceCount>
        </item>
      </instanceCounts>
      <priceSchedules>
        <item>
          <term>11</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>>false</active>
        </item>
        <item>
          <term>10</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>>false</active>
        </item>
        <item>
          <term>9</term>
          <price>2.5</price>
          <currencyCode>USD</currencyCode>
          <active>>false</active>
        </item>
      </priceSchedules>
    </item>
  </reservedInstancesListingsSet>
</DescribeReservedInstancesListingsResponse>
```

```
<item>
  <term>8</term>
  <price>2.0</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
<item>
  <term>7</term>
  <price>2.0</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
<item>
  <term>6</term>
  <price>2.0</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
<item>
  <term>5</term>
  <price>1.5</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
<item>
  <term>4</term>
  <price>1.5</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
<item>
  <term>3</term>
  <price>0.7</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
<item>
  <term>2</term>
  <price>0.7</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
<item>
  <term>1</term>
  <price>0.1</price>
  <currencyCode>USD</currencyCode>
  <active>false</active>
</item>
</priceSchedules>
<tagSet/>
<clientToken>listRI1</clientToken>
</item>
</reservedInstancesListingsSet>
</DescribeReservedInstancesListingsResponse>
```

## To change your listing

Currently, you cannot modify your listing directly. However, you can change your listing by first canceling it and then creating another listing with new parameters.

1. Cancel your active Reserved Instance listing. For more information, see the previous procedure.
2. Create a new listing. For more information, see [Listing Your Reserved Instance \(p. 242\)](#).

## After Your Reserved Instance Is Sold

### Topics

- [Reserved Instance Listing States \(p. 260\)](#)
- [Lifecycle of a Listing \(p. 260\)](#)
- [Getting Paid \(p. 261\)](#)
- [Notifications \(p. 262\)](#)

When your Reserved Instance is sold, AWS will send you an email notification. Each day that there is any kind of activity (for example, you create a listing; you sell a listing; or AWS sends funds to your account), you will get one email notification capturing all the activities of the day. For more information, see [Notifications \(p. 262\)](#).

You can track the status of your Reserved Instance listings by looking at the **My Listings** tab of the selected Reserved Instance on the **Reserved Instance** page in the Amazon EC2 console. The tab contains the **Listing State** as well as information about the term, listing price, and a breakdown of how many instances in the listing are available, pending, sold, and cancelled. You can also use the `ec2-describe-reserved-instances-listings` CLI command or the `DescribeReservedInstancesListings` API call, with the appropriate filter to obtain information about your Reserved Instance listings.

## Reserved Instance Listing States

**Listing State** displays the current status of your Reserved Instance listings:

- **Active**—The listing is available for purchase.
- **Cancelled**—The listing is canceled and won't be available for purchase in the marketplace.
- **Closed**—The Reserved Instance is not listed. A Reserved Instance might be *Closed* because the sale of the listing was completed.

The information displayed by **Listing State** is about the status of your listing in the Reserved Instance Marketplace. It is different from the status information that is displayed by the **State** column in the Reserved Instance page. This **State** information is about your Reserved Instance. For more information, see [Reserved Instance States \(p. 222\)](#).

## Lifecycle of a Listing

Now that you have created a listing, let's walk through what happens when your listing sells.

When *all* the instances in your listing are matched and sold, the **My Listings** tab shows that your **Total instance count** matches the count listed under **Sold**, there are no **Available** instances left for your listing, and its **Status** is *closed*.

When only *a portion* of your listing is sold, AWS retires the Reserved Instances in the listing and creates the number of Reserved Instances equal to the Reserved Instances remaining in the count. So, the

Reserved Instances listing ID and the listing that it represents, which now has an instance count of fewer instances for sale, is still active.

Any future sales of Reserved Instances in this listing are processed this way. When all the Reserved Instances in the listing are sold, AWS marks the listing as *closed*.

For example, let's say you created a listing *Reserved Instances listing ID 5ec28771-05ff-4b9b-aa31-9e57dexample* with an instance count of 5.

Your **My Listings** tab in the **Reserved Instance** page of the Amazon EC2 console will display the listing this way:

*Reserved Instance listing ID 5ec28771-05ff-4b9b-aa31-9e57dexample*

- Total instance count = 5
- Sold = 0
- Available = 5
- Status = active

Let's say that a buyer purchases two of the instances, which leaves a count of three instances still available for sale. As a result of this partial sale, AWS creates a new Reserved Instance with an instance count of three to represent the remaining three that are still for sale.

This is how your listing will look in your **My Listings** tab:

*Reserved Instance listing ID 5ec28771-05ff-4b9b-aa31-9e57dexample*

- Total instance count = 5
- Sold = 2
- Available = 3
- Status = active

If you decide to cancel your listing and a portion of that listing has already sold, the cancellation is not effective on the portion that has been sold. Only the portion of the listing not yet sold will no longer be available in the Reserved Instance Marketplace.

## Getting Paid

As soon as AWS receives funds from the buyer of your Reserved Instance, AWS sends a message to your email address—that is, the email address associated with the account that is registered as owner of the Reserved Instance that was sold.

AWS sends an Automated Clearing House (ACH) wire transfer to the bank account that you specified when you registered for the Reserved Instance Marketplace. Typically, this transfer occurs between one to three days after your Reserved Instance has been matched. You can view the state of this disbursement by viewing your Reserved Instance disbursement report. Disbursements take place once a day. Keep in mind that you will not be able to receive disbursements until AWS has received verification from your bank. This period can take up to two weeks.

The Reserved Instance you sold will continue to appear in the results of `DescribeReservedInstances` calls you make for another 60 days before disappearing from the list. You will receive a pro-rated refund for the portion of the upfront fee you paid for the Reserved Instance that you did not use.



## Notifications

As a seller in the Reserved Instance Marketplace, you will receive an email digest of the Reserved Instance Marketplace activities pertaining to your account. On any given day, you will receive one email digest, and you will only receive this email if one or a combination of activities occurred that day:

- You created a new listing in the Reserved Instance Marketplace.
- You sold one or more of the Reserved Instances you listed.
- AWS posted a disbursement to your bank account as a result of a sale of part or all of your listing in the Reserved Instance Marketplace.

Your email digest will look similar to the following:

Dear John Doe,

Your account's Reserved Instance Marketplace activity for 2012/08/27 is shown below:

**New Listings**

Platform	Instance Type	Offering Type	Term	Availability Zone	Price	Count	Total Potential Earnings
Linux/UNIX	m1.large	Medium Utilization	2 months	us-east-1a	\$24.44	2	\$99.11

**Recent Sales**

Platform	Instance Type	Offering Type	Term	Availability Zone	Price	Count	Total Earnings
Linux/UNIX (Amazon VPC)	m1.large	Heavy Utilization	35 months	us-east-1a	\$1,480.50	1	\$1,317.64
Linux/UNIX (Amazon VPC)	m1.large	Medium Utilization	35 months	us-east-1a	\$1,800.00	1	\$1,420.50

**Disbursements**

Zip Code	Country	Disbursement Amount	Status
12345	United States	\$1,056.11	Success
<b>Total</b>		<b>\$1,056.11</b>	

We initiated a transfer to your checking account (ending with 2634) of \$1,056.11 on 2012/07/19. Funds usually arrive within 3 - 5 banking days, but times vary by bank.

If you have questions about your Reserved Instance Marketplace activity above, please visit the Reserved Instance Marketplace web page at <http://aws.amazon.com/reserved-instance-marketplace>.

Sincerely,  
The Amazon EC2 Team

## Requirements Checklist for Reserved Instances

Amazon EC2 Reserved Instances and the Reserved Instance Marketplace can be a powerful and cost-saving strategy for running your business. However, before you use Reserved Instances or the Reserved Instance Marketplace, ensure that you meet the requirements for purchase and sale. You also must understand the details and restrictions on certain elements of Reserved Instances and the Reserved Instance Marketplace—including seller registration, banking, using the AWS Free tier, dealing with cancelled instances, and so on. Use this topic as a checklist for buying and selling Reserved Instances, and for buying and selling in the Reserved Instance Marketplace.

### Reserved Instances

- **AWS account**—You need to have an AWS account in order to purchase Reserved Instances. If you don't have an AWS account, you should read and complete the instructions described in [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#), which provides information on signing up for your Amazon EC2 account and credentials.
- **AWS free tier**—The AWS free usage tier is available for new AWS accounts. If you are using the AWS free usage tier to run Amazon EC2 instances, and then you purchase a Reserved Instance, you will be charged for the Reserved Instance under standard pricing guidelines. For information about the free tier and the applicable services and usage amounts, see [AWS Free Usage Tier](#).

### Buying Reserved Instances

- **Usage fee for Heavy Utilization**—With Light Utilization and Medium Utilization Reserved Instances, you pay a one-time upfront fee and then only pay the hourly price when you use the instance. With Heavy Utilization Reserved Instances, you pay a low, one-time upfront fee and commit to paying an hourly rate for every hour of the Reserved Instance's term *whether or not you use it*. For information, see [Choosing Reserved Instances Based on Your Usage Plans \(p. 199\)](#).

- **Tiered discounts on purchases**—The Reserved Instance pricing tier discounts only apply to purchases made from AWS. These discounts do not apply to purchases of third-party Reserved Instances. For information, see [Understanding Reserved Instance Pricing Tiers \(p. 200\)](#).
- **Cancellation of purchase**—Before you confirm your purchase, review the details of the Reserved Instances that you plan to buy, and make sure that all the parameters are accurate. After you purchase a Reserved Instance (either from a third-party seller in the Reserved Instance Marketplace or from AWS), you cannot cancel your purchase. However, you can sell the Reserved Instance if your needs change. For information about selling your Reserved Instance in the Reserved Instance Marketplace, see [Selling Your Reserved Instances \(p. 240\)](#).

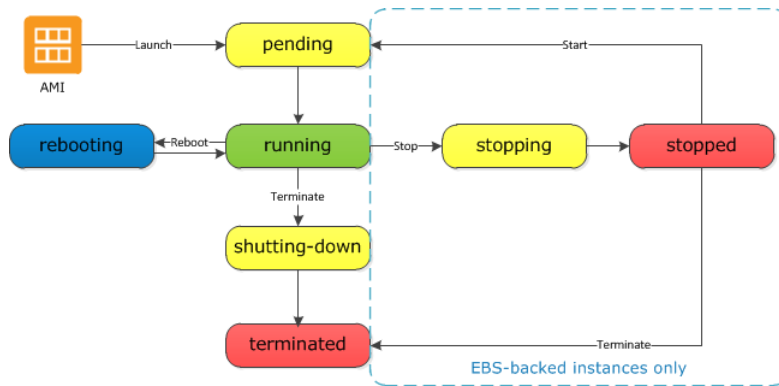
### Selling Reserved Instances and the Reserved Instance Marketplace

- **Seller requirement**—To become a seller in the Reserved Instance Marketplace, you must register as a seller. For information, see [Selling Your Reserved Instances \(p. 240\)](#).
- **Bank requirement**—AWS must have your bank information in order to disburse funds collected when you sell your Reserved Instance. The bank you specify must have a US address. For more information, see [Your Bank \(p. 238\)](#).
- **Tax requirement**—Sellers who have 200 or more transactions or who plan to sell \$20,000 or more in Reserved Instances will have to provide additional information about their business for tax reasons. For information, see [Tax Information \(p. 239\)](#).
- **Minimum selling**—The minimum price allowed in the Reserved Instance Marketplace is \$1.01.
- **When Reserved Instances can be sold**—Reserved Instances can be sold only after AWS has received the upfront payment and the Reserved Instance has been active (you've owned it) for at least 30 days. In addition, there must be at least a month remaining in the term of the Reserved Instance you are listing.
- **Modifying your listing**—Currently, you cannot modify your listing in the Reserved Instance Marketplace directly. However, you can change your listing by first canceling it and then creating another listing with new parameters. For information, see [Canceling and Changing Your Listings \(p. 250\)](#). You also can change your Reserved Instances before listing them. For information, see [Modifying Your Reserved Instances \(p. 227\)](#).
- **Selling discounted Reserved Instances**—Amazon EC2 Reserved Instances purchased at a reduced cost resulting from a tiering discount cannot be sold in the Reserved Instance Marketplace. For more information about the Reserved Instance Marketplace, see [Reserved Instance Marketplace \(p. 209\)](#).
- **Service fee**—AWS will charge you a service fee of 12 percent of the total upfront price of each Reserved Instance you sell in the marketplace. (The upfront price is the price the seller is charging for his Reserved Instance.)
- **Other AWS Reserved Instances**—Only Amazon EC2 Reserved Instances can be sold in the Reserved Instance Marketplace. Other AWS Reserved Instances, such as Amazon Relational Database Service (Amazon RDS) and Amazon ElastiCache Reserved Instances cannot be sold on the Reserved Instance Marketplace.

## Instance Lifecycle

This topic describes the lifecycle of an Amazon EC2 instance, from the moment you launch it through its termination. By working with Amazon EC2 to manage your instance, you ensure that your customers have the best possible experience with the applications or sites that you host on your instance.

The following illustration represents the transitions between instance states.



## Instance Launch

When you launch an instance, it enters the `pending` state. The instance type that you specified at launch determines the hardware of the host computer for your instance. We use the Amazon Machine Image (AMI) you specified at launch to boot the instance. After the instance is ready for you, it enters the `running` state. You can connect to your running instance and use it the way that you'd use a computer sitting in front of you.

As soon as your instance starts to boot, you're billed for each hour or partial hour that you keep the instance running (even if the instance remains idle and you don't connect to it).

For more information, see [Launch Your Instance \(p. 266\)](#) and [Connect to Your Amazon EC2 Instance \(p. 273\)](#).

## Instance Stop and Start

If your instance fails a status check or is not running your applications as expected, and if the root volume of your instance is an Amazon EBS volume, you can stop and start your instance to try to fix the problem.

When you stop your instance, it enters the `stopping` state, and then the `stopped` state. We don't charge hourly usage or data transfer fees for your instance after you stop it, but we do charge for the storage for any Amazon EBS volumes. While your instance is in the `stopped` state, you can modify certain attributes of the instance, including the instance type.

When you start your instance, it enters the `pending` state, and we move the instance to a new host computer. Therefore, when you stop and start your instance, you'll lose any data on the instance store volumes on the previous host computer.

If your instance is running in EC2-Classic, it receives a new private IP address, which means that an Elastic IP address associated with the private IP address is no longer associated with your instance. If your instance is running in EC2-VPC, it retains its private IP address, which means that an Elastic IP address associated with the private IP address or network interface is still associated with your instance.

Each time you transition an instance from `stopped` to `running`, we charge a full instance hour, even if these transitions happen multiple times within a single hour.

For more information, see [Stop and Start Your Instance \(p. 284\)](#).

## Instance Reboot

You can reboot your instance using the Amazon EC2 console, the Amazon EC2 CLI, and the Amazon EC2 API. We recommend that you use Amazon EC2 to reboot your instance instead of running the operating system reboot command from your instance.

Rebooting an instance is equivalent to rebooting an operating system; the instance remains on the same host computer and maintains its public DNS name, private IP address, and any data on its instance store volumes. It typically takes a few minutes for the reboot to complete, but the time it takes to reboot depends on the instance configuration.

Rebooting an instance doesn't start a new instance billing hour.

For more information, see [Reboot Your Instance \(p. 286\)](#).

## Instance Termination

When you've decided that you no longer need an instance, you can terminate it. As soon as the status of an instance changes to `shutting-down` or `terminated`, you stop incurring charges for that instance.

After you terminate an instance, it remains visible in the console for a short while, and then the entry is deleted. You can also describe a terminated instance using the CLI and API. You can't connect to or recover a terminated instance.

Each Amazon EBS-backed instance supports the `InstanceInitiatedShutdownBehavior` attribute, which controls whether the instance stops or terminates when you initiate a shutdown from within the instance itself (for example, by using the **shutdown** command on Linux). The default behavior is to stop the instance. You can modify the setting of this attribute while the instance is running or stopped.

For more information, see [Terminate Your Instance \(p. 287\)](#).

## Differences Between Reboot, Stop, and Terminate

The following table summarizes the key differences between rebooting, stopping, and terminating your instance.

Characteristic	Reboot	Stop/start	Terminate
Host computer	The instance stays on the same host computer	The instance runs on a new host computer	None
Private and public IP addresses	These addresses stay the same	The instance gets new private and public IP addresses	None
Elastic IP addresses	The Elastic IP address remains associated with the instance	EC2-Classic: The Elastic IP address is disassociated from the instance EC2-VPC: The Elastic IP address remains associated with the instance	The Elastic IP address is disassociated from the instance
Instance store volumes	The data is preserved	The data is lost	The data is lost

# Launch Your Instance

## Topics

- [Launching an Instance \(p. 266\)](#)
- [Launching an Instance from a Backup \(p. 271\)](#)
- [Launching an AWS Marketplace Instance \(p. 271\)](#)

An instance is a virtual server in the AWS cloud. You launch an instance from an Amazon Machine Image (AMI). The AMI provides the operating system, application server, and applications for your instance.

To begin, the instance state is `pending`. When the instance state is `running`, the instance has started booting. There might be a short time before you can connect to the instance. The instance receives a public DNS name that you can use to contact the instance from the Internet. The instance also receives a private DNS name that other instances within the same Amazon EC2 network (EC2-Classic or EC2-VPC) can use to contact the instance.

When you sign up for AWS, you can get started with Amazon EC2 for free using the [AWS Free Usage Tier](#). You can either leverage the Free Usage Tier to launch and use a Micro instance for free for 12 months. If you launch an instance that is not within the Free Usage Tier, you incur the standard Amazon EC2 usage fees for the instance. For more information, see the [Amazon EC2 Pricing](#).

### Important

When you launch an instance that is not within the Free Usage Tier, you are charged hourly for the time that the instance is running, even if it remains idle.


When you are finished with an instance, be sure to terminate it. For more information, see [Terminate Your Instance \(p. 287\)](#).

## Launching an Instance

Before you launch your instance, be sure that you are set up. For more information, see [Get Set Up for Amazon EC2 \(p. 18\)](#).

### To launch an instance

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation bar at the top of the screen, the current region is displayed. Click the region's name to select the region for the instance. This choice is important because some Amazon EC2 resources can be shared between regions, while others can't. Select the region that meets your needs. For more information, see [Resource Locations \(p. 528\)](#).



3. From the Amazon EC2 console dashboard, click **Launch Instance**.
4. The **Choose an Amazon Machine Image (AMI)** page displays a list of basic configurations called Amazon Machine Images (AMIs). An AMI contains all the information needed to create a new instance. For example, an AMI might contain the software required to act as a web server: for example, Linux, Apache, and your web site. You can select the type of AMI to use by using the categories on the left pane.
  - The **Quick Start** category displays a selection of popular AMIs to help you get started quickly. To keep things simple, AWS marks the AMIs that are available in the free usage tier with **Free tier eligible**.
  - The **My AMIs** category displays private AMIs that you own, or private AMIs that have been shared with you.
  - The **AWS Marketplace** is an online store where you can buy software that runs on AWS, including AMIs that you can use to launch your instance. For more information about launching an instance from the AWS Marketplace, see [Launching an AWS Marketplace Instance \(p. 271\)](#).
  - The **Community AMIs** option displays public AMIs that AWS community members have made available for others to use.

To filter the list of available AMIs, use the filter options on the left. For example, to view Amazon EBS-backed community AMIs, select the **Community AMIs** category, select the **Amazon Linux** check box under **Operating system**, then select **EBS** under **Root device type**.

**Note**

As you choose an AMI, it's important to note whether the AMI is backed by an instance store or by Amazon EBS. For more information, see [Storage for the Root Device \(p. 45\)](#).

Choose the AMI to use and click **Select**.

5. On the **Choose an Instance Type** page, select the hardware configuration and size of the instance to launch. Larger instance types have more CPU and memory. For more information about instance families, see [Instance Types \(p. 94\)](#).

To stay within the free tier, select the **Micro instances** category, select the **t1.micro** instance, and then click **Next: Configure Instance Details**.

### Note

If you are new to AWS and would like to set up an instance quickly and easily, you can click **Review and Launch** at this point to accept default configuration settings, and launch your instance. For more information, see [Getting Started with Amazon EC2 Linux Instances \(p. 22\)](#).

6. On the **Configure Instance Details** page, change the following settings as necessary (expand **Advanced Details** to see all the settings), and then click **Next: Add Storage**:
  - **Number of instances:** Enter the number of instances to launch.
  - **Purchasing option:** Select **Request Spot Instances** to launch a Spot Instance. For more information, see [Spot Instances \(p. 115\)](#).
  - **Network:** Your account may support the EC2-Classic and EC2-VPC platforms, or EC2-VPC only. For more information, see [Supported Platforms \(p. 417\)](#). To launch into EC2-Classic, select **Launch into EC2-Classic**. If your account supports EC2-VPC only, you can launch your instance into your default VPC, which has already been created for you. Regardless of which platform your account supports, you can also launch the instance into a VPC that you have already created, also known as a nondefault VPC.

If you are launching into EC2-Classic:

- **Availability Zone:** Select the Availability Zone to use. To let AWS choose an Availability Zone for you, select **No preference**.

If you are launching into a VPC:

- **Network:** Select a VPC, or to create a new VPC, click **Create new VPC** to go the VPC console. When you are done, return to the wizard and click the **Refresh** button to load your VPC in the list.
- **Subnet:** Select the subnet into which to launch your instance. If your account is EC2-VPC only, select **No preference** to let AWS choose a default subnet in any Availability Zone. To create a new subnet, click **Create new subnet** to go to the VPC console. When you are done, return to the wizard and click the **Refresh** button to load your subnet in the list.
- **Public IP:** Select this check box to request that your instance receives a public IP address. For more information about public IP addressing, see [Amazon EC2 Instance IP Addressing \(p. 419\)](#).
- **IAM role:** If applicable, select an AWS Identity and Access Management (IAM) role to associate with the instance. For more information, see [Controlling Access to Amazon EC2 Resources \(p. 399\)](#).
- **Shutdown behavior:** Select whether the instance should stop or terminate when shut down.
- **Enable termination protection:** Select this check box to prevent accidental termination. For more information, see [Enabling Termination Protection for an Instance \(p. 288\)](#).
- **Monitoring:** Select this check box to enable detailed monitoring of your instance using Amazon CloudWatch. Additional charges apply. For more information, see [Introduction to Amazon CloudWatch](#) in the *Amazon CloudWatch Developers Guide*.
- **EBS-Optimized instance:** An Amazon EBS-optimized instance uses an optimized configuration stack and provides additional, dedicated capacity for Amazon EBS I/O. If the instance type supports this feature, select this check box to enable it. Additional charges apply. For more information, see [EBS-Optimized Instances \(p. 107\)](#).
- **Tenancy:** If you are launching your instance into a VPC, you can select **Dedicated tenancy** to run your instance on isolated, dedicated hardware. Additional charges apply.
- **Network interfaces:** If you are launching an instance into VPC and you did not select **No Preference** for your subnet, you can specify up to two network interfaces in the wizard. Click **Add IP** to assign more than one IP address to the selected interface. For more information about network interfaces, see [Elastic Network Interfaces \(ENI\) \(p. 431\)](#). If you selected the **Public IP** check box above, you can only assign a public IP address to a single, new network interface with the device index of eth0. For more information, see [Assigning a Public IP Address \(p. 423\)](#).
- **Kernel ID:** Select **Use default** unless you want to use a specific kernel.
- **RAM disk ID:** Select **Use default** unless you want to use a specific RAM disk. If you have selected a kernel, you may need to select a specific RAM disk with the drivers to support it.



- **Placement group:** A placement group is a logical grouping for your cluster instances. Select an existing placement group, or create a new one.
  - **User data:** You can specify user data to configure an instance during launch, or to run a configuration script. To attach a file, select the **As file** option and browse for the file to attach.
7. On the **Add Storage** page, you can specify volumes to attach to the instance besides the volumes specified by the AMI (such as the root device volume). You can change the following options, then click **Next: Tag Instance** when done:
- **Type:** Select instance store or Amazon EBS volumes to associate with your instance. The type of volume available in the list depends on the instance type you've chosen.
  - **Device:** Select from the list of available device names for the volume.
  - **Snapshot:** Enter the name or ID of the snapshot from which to restore a volume. You can also search for public snapshots by typing text into the **Snapshot** field. Snapshot descriptions are case-sensitive.
  - **Size:** For Amazon EBS-backed volumes, you can specify a storage size. Note that even if you have selected an AMI and instance that are eligible for the free usage tier, you need to keep under 30 GB of total storage to stay within the free usage tier.
  - **Volume Type:** For Amazon EBS volumes, select either a Standard or Provisioned IOPS volume. For more information, see [Amazon Elastic Block Store \(Amazon EBS\)](#) (p. 446).
  - **IOPS:** If you have selected a Provisioned IOPS volume type, then you can enter the number of I/O operations per second (IOPS) that the volume can support.
  - **Delete on Termination:** For Amazon EBS volumes, select this check box to delete the volume when the instance is terminated.
8. On the **Tag Instance** page, specify tags for the instance by providing key and value combinations. Click **Create Tag** to add more than one tag to your resource. Click **Next: Configure Security Group** when you are done.

For more information about tags, see [Tagging Your Amazon EC2 Resources](#) (p. 532).

9. On the **Configure Security Group** page, the wizard automatically defines the launch-wizard-x security group to allow you to connect to your instance.

A security group defines firewall rules for your instances. These rules specify which incoming network traffic is delivered to your instance. All other traffic is ignored. The launch-wizard-x security group automatically allows traffic on either SSH (port 22) for Linux instances, or RDP (port 3389) for Windows instances.

### Caution

The launch-wizard-x security group enables all IP addresses (0.0.0.0/0) to access your instance over the specified ports. This is acceptable for this short exercise, but it's unsafe for production environments. You should authorize only a specific IP address or range of addresses to access your instance.

You can modify the default rule in the launch-wizard-x security group to suit your needs. For example, if you want to use your instance as a web server, open ports 80 (HTTP) and 443 (HTTPS) to allow Internet traffic to reach your instance.

Select **My IP** from the **Source** list to let the wizard add your computer's public IP address. However, if you are connecting through an ISP or from behind your firewall without a static IP address, you need to find out the range of IP addresses used by client computers. Click **Add Rule** to add more rules to the group.



You can also select the **Select an existing security group** option to use one of your existing security groups, for example, the one you created when getting set up. You can't edit an existing group's rules, but you can copy its rules into a new group clicking the **Copy to new** link for that group.

For more information about security groups, see [Amazon EC2 Security Groups \(p. 392\)](#).

When you are done, click **Review and Launch**.

10. On the **Review Instance Launch** page, check the details of your instance, and make any necessary changes by clicking the appropriate **Edit** link.

When you are ready, click **Launch**.

11. In the **Select an existing key pair or create a new key pair** dialog box, you can choose an existing key pair, or create a new one. For example, select **Choose an existing key pair**, then select the key pair you created when getting set up.

To launch your instance, select the acknowledgment check box, then click **Launch Instances**.

For more information about creating key pairs, see [Amazon EC2 Key Pairs \(p. 385\)](#).

#### **Important**

We recommend against selecting the **Proceed without key pair** option. If you launch an instance without a key pair, you won't be able to connect to it. This option is used only when you are creating your own AMI and don't need to connect to the instance.

12. If the instance state immediately goes to `terminated` instead of `running`, you can get information about why the instance didn't launch. For more information, see [What To Do If An Instance Immediately Terminates \(p. 353\)](#).

## Launching More Instances From a Template

The Amazon EC2 console provides a **Launch More Like This** wizard option that enables you to use a current instance as a template for launching other instances. This option automatically populates the Amazon EC2 launch wizard with certain configuration details from the selected instance.

#### **Note**

The **Launch More Like This** wizard option does not clone your selected instance; it only replicates some configuration details. To create a copy of your instance, first create an AMI from it, then launch more instances from the AMI. For more information about creating your own AMI, see [Creating Your Own AMIs \(p. 62\)](#).

The following configuration details are copied from the selected instance into the launch wizard:

- AMI ID
- Instance type
- Availability Zone, or the VPC and subnet in which the selected instance is located
- Tags associated with the instance, if applicable
- Kernel ID and RAM disk ID, if applicable
- IAM role associated with the instance, if applicable
- Security group associated with the instance
- Tenancy setting, if launching into a VPC (shared or dedicated)
- EBS-optimization setting (true or false)

The following configuration details are not copied from your selected instance; instead, the wizard applies their default settings or behavior:

- Storage: The default storage configuration is determined by the AMI and the instance type.

- Public IP address: The option to assign a public IP address to your instance is enabled by default when launching into a default subnet.
- Termination protection: Disabled by default.
- Shutdown behavior: Set to 'stop' by default.
- User data: None by default.

### To use your current instance as a template

1. On the Instances page, select the instance you want to use.
2. Click **Actions**, and select **Launch More Like This**.
3. The launch wizard opens on the **Review Instance Launch** page. You can check the details of your instance, and make any necessary changes by clicking the appropriate **Edit** link.

When you are ready, click **Launch** to select a key pair and launch your instance.

## Launching an Instance from a Backup

With an Amazon EBS-backed instance, you can back up the root device volume of the instance by creating a snapshot. When you have a snapshot of the root device volume of an instance, you can terminate that instance and then later launch a new instance from the snapshot. This can be useful if you don't have the original AMI that you launched an instance from, but you need to be able to launch an instance using the same image.

### Important

At this time, although you can create a Windows AMI from a snapshot, you can't launch an instance from the AMI.

Use the following procedure to create an AMI from the root volume of your instance. If you prefer, you can use the [ec2-register](#) command instead.

### To create an AMI from your root volume

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, under **Elastic Block Store**, click **Snapshots**.
3. Click **Create Snapshot**.
4. Select the root volume, and then click **Create**.
5. Select the snapshot that you just created, and then click **Create Image**.
6. In the **Create Image from EBS Snapshot** dialog box, complete the fields to create your AMI, then click **Yes, Create**. Be sure to do the following:
  - Select the architecture from the **Architecture** list (**i386** for 32-bit or **x86\_64** for 64-bit).
  - Select the AKI from the **Kernel ID** list. If you select the default AKI or don't select an AKI, you'll be required to specify an AKI every time you launch an instance. In addition, your instance may fail the health check because the default AKI is incompatible with the instance.
7. In the navigation pane, select **AMIs**.
8. Select the AMI that you just created, and then click **Launch**. Follow the wizard to launch your instance.

## Launching an AWS Marketplace Instance

You can subscribe to an AWS Marketplace product and launch an instance from the product's AMI using the Amazon EC2 launch wizard. For more information about paid AMIs, see [Paid AMIs \(p. 57\)](#). To cancel

your subscription after launch, you first have to terminate all instances running from it. For more information, see [Managing Your AWS Marketplace Subscriptions \(p. 60\)](#).

### To launch an instance from the AWS Marketplace using the launch wizard

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the Amazon EC2 dashboard, click **Launch Instance**.
3. On the **Choose an Amazon Machine Image (AMI)** page, select the **AWS Marketplace** category on the left. Find a suitable AMI by browsing the categories, or using the search functionality. Click **Select** to choose your product.
4. A dialog displays an overview of the product you've selected. You can view the pricing information, as well as any other information that the vendor has provided. When you're ready, click **Continue**.

#### Note

You are not charged for using the product until you have launched an instance with the AMI. Take note of the pricing for each supported instance type, as you will be prompted to select an instance type on the next page of the wizard.

5. On the **Choose an Instance Type** page, select the hardware configuration and size of the instance to launch. When you're done, click **Next: Configure Instance Details**.
6. On the next pages of the wizard, you can configure your instance, add storage, and add tags. For more information about the different options you can configure, see [Launching an Instance \(p. 266\)](#). Click **Next** until you reach the **Configure Security Group** page.

The wizard creates a new security group according to the vendor's specifications for the product. The security group may include rules that allow all IP addresses (0.0.0.0/0) access on SSH (port 22) or RDP (port 3389). We recommend that you adjust these rules to allow only a specific IP address or range of addresses to access your instance over those specific ports.

When you are ready, click **Review and Launch**.

7. On the **Review Instance Launch** page, check the details of the AMI from which you're about to launch the instance, as well as the other configuration details you set up in the wizard. When you're ready, click **Launch** to choose or create a key pair, and launch your instance.
8. Depending on the product you've subscribed to, the instance may take a few minutes or more to launch. You are first subscribed to the product before your instance can launch. If there are any problems with your credit card details, you will be asked to update your account details. When the launch confirmation page displays, click **View Instances** to go to the Instances page.

#### Note

You are charged the subscription price as long as your instance is running, even if it is idle. If your instance is stopped, you may still be charged for storage.

9. When your instance is in the **running** state, you can connect to it. To do this, select your instance in the list and click **Connect**. Follow the instructions in the dialog. For more information about connecting to your instance, see [Connect to Your Amazon EC2 Instance \(p. 273\)](#).

#### Important

Check the vendor's usage instructions carefully, as you may need to use a specific user name to log in to the instance. For more information about accessing your subscription details, see [Managing Your AWS Marketplace Subscriptions \(p. 60\)](#).

## Launching an AWS Marketplace AMI Instance Using the API and CLI

To launch instances from AWS Marketplace products using the API or command line tools, first ensure that you are subscribed to the product. You can then use the `RunInstances` request or the `ec2-run-instances` command to launch an instance using the product's AMI ID.

## Connect to Your Amazon EC2 Instance

This section describes how to connect to instances that you launched and how to transfer files between your local computer and your Amazon EC2 instance.

Your Computer	Your Instance	Topic
All	Linux/UNIX	<a href="#">Connecting to Your Instance from Your Web Browser Using MindTerm (p. 273)</a>
Linux/UNIX	Linux/UNIX	<a href="#">Connecting to Your Linux/UNIX Instances Using SSH (p. 279)</a>
Windows	Linux/UNIX	<a href="#">Connecting to Linux/UNIX Instances from Windows Using PuTTY (p. 274)</a>
All	Windows	<a href="#">Connecting to Windows Instances Using RDP (p. 282)</a>

After you connect to your instance you can try one of our tutorials, such as [Tutorial: Installing a LAMP Web Server \(p. 30\)](#) or [Tutorial: Hosting a WordPress Blog with Amazon EC2 \(p. 36\)](#).

## Connecting to Your Instance from Your Web Browser Using MindTerm

After you launch your instance, you can connect to it and use it the way that you'd use a computer sitting in front of you.

If you receive an error while attempting to connect to your instance, see [Troubleshooting Connecting to Your Instance](#).

The following instructions explain how to connect to your instance using MindTerm through the Amazon EC2 console.

### Prerequisites

- **Install Java**  
Your Linux computer most likely includes Java. If not, see [How do I enable Java in my web browser?](#) On a Windows or Mac client, you must run your browser using administrator credentials. For Linux, additional steps may be required if you are not logged in as `root`.
- **Enable Java in your browser**  
For instructions, see [http://java.com/en/download/help/enable\\_browser.xml](http://java.com/en/download/help/enable_browser.xml).
- **Locate the private key**  
You'll need the fully-qualified path of the `.pem` file for the key pair that you specified when you launched the instance.
- **Enable inbound SSH traffic from your IP address to your instance**  
Ensure that the security group associated with your instance allows incoming SSH traffic from your IP address. For more information, see [Authorizing Network Access to Your Instances](#).

## Starting MindTerm

### To connect to your instance using a web browser with MindTerm

1. In the Amazon EC2 console, click **Instances** in the navigation pane.
2. Select the instance, and then click **Connect**.
3. Click **A Java SSH client directly from my browser (Java required)**.
4. Amazon EC2 automatically detects the public DNS name of your instance and the name of the populates **Public DNS** for you. It also detects name of the key pair that you specified when you launched the instance. Complete the following, and then click **Launch SSH Client**.
  - a. In **User name**, enter the user name to log in to your instance.

**Tip**  
For Amazon Linux, the user name is `ec2-user`. For RHEL5, the user name is often `root` but might be `ec2-user`. For an Ubuntu, AMI the user name is `ubuntu`. Otherwise, check with your AMI provider.
  - b. In **Private key path**, enter the fully-qualified path to your private key (`.pem`) file.
  - c. (Optional) Click **Store in browser cache** to store the location of the private key in your browser cache. This enables Amazon EC2 to detect the location of the private key in subsequent browser sessions, until you clear your browser's cache.
5. If necessary, click **Yes** to trust the certificate.
6. Click **Run** to run the MindTerm client.
7. If you accept the license agreement, click **Accept**.
8. If this is your first time running MindTerm, a series of dialog boxes asks you to confirm setup for your home directory and other settings. Confirm these settings. A window opens and you are connected to your instance.

## Connecting to Linux/UNIX Instances from Windows Using PuTTY

After you launch your instance, you can connect to it and use it the way that you'd use a computer sitting in front of you.

If you receive an error while attempting to connect to your instance, see [Troubleshooting Connecting to Your Instance](#).

The following instructions explain how to connect to your instance using PuTTY, a free SSH client for Windows.

### Topics

- [Prerequisites](#) (p. 274)
- [Converting Your Private Key Using PuTTYgen](#) (p. 275)
- [Starting a PuTTY Session](#) (p. 276)
- [Transferring Files to Your Instance with the PuTTY Secure Copy Client](#) (p. 277)
- [Transferring Files to Your Instance with WinSCP](#) (p. 278)

### Prerequisites

- **Install PuTTY**  
Download and install PuTTY from the [PuTTY download page](#). Be sure to install the entire suite.


- **Get the ID of the instance**  
You can get the ID of your instance using the Amazon EC2 console (from the **Instance ID** column). If you prefer, you can use the [ec2-describe-instances](#) command.
- **Get the public DNS name of the instance**  
You can get the public DNS for your instance using the Amazon EC2 console (check the **Public DNS** column; if this column is hidden, click the **Show/Hide** icon and select **Public DNS**). If you prefer, you can use the [ec2-describe-instances](#) command.
- **Locate the private key**  
You'll need the fully-qualified path of the `.pem` file for the key pair that you specified when you launched the instance.
- **Enable inbound SSH traffic from your IP address to your instance**  
Ensure that the security group associated with your instance allows incoming SSH traffic from your IP address. For more information, see [Authorizing Network Access to Your Instances](#).

## Converting Your Private Key Using PuTTYgen


PuTTY does not natively support the private key format (`.pem`) generated by Amazon EC2. PuTTY has a tool named PuTTYgen, which can convert keys to the required PuTTY format (`.ppk`). You must convert your private key into this format (`.ppk`) before attempting to connect to your instance using PuTTY.

### To convert your private key

1. Start PuTTYgen (for example, from the **Start** menu, click **All Programs > PuTTY > PuTTYgen**).
2. Under **Type of key to generate**, select **SSH-2 RSA**.



3. Click **Load**. By default, PuTTYgen displays only files with the extension `.ppk`. To locate your `.pem` file, select the option to display files of all types.



4. Select your `.pem` file and click **Open**. Click **OK** to dismiss the confirmation dialog box.
5. Click **Save private key** to save the key in the format that PuTTY can use. PuTTYgen displays a warning about saving the key without a passphrase. Click **Yes**.

#### Note

A passphrase on a private key is an extra layer of protection, so even if your private key is discovered, it can't be used without the passphrase. The downside to using a passphrase is that it makes automation harder because human intervention is needed to log on to an instance, or copy files to an instance.

6. Specify the same name for the key that you used for the key pair (for example, `my-key-pair`). PuTTY automatically adds the `.ppk` file extension.

Your private key is now in the correct format for use with PuTTY. You can now connect to your instance using PuTTY's SSH client.

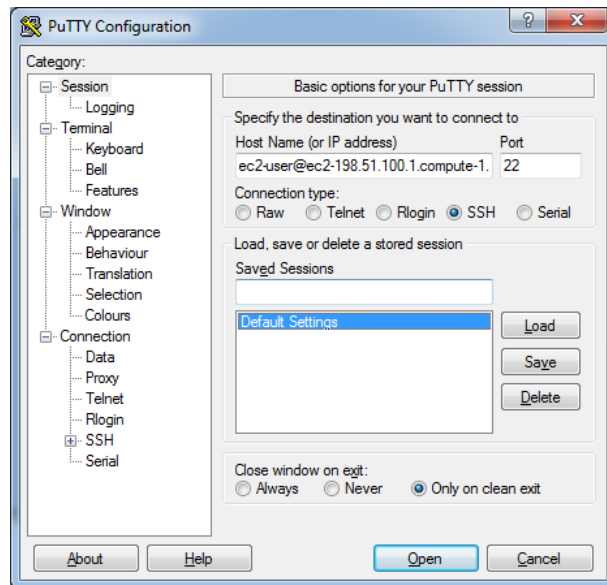
## Starting a PuTTY Session

### To start a PuTTY session

- (Optional) If you've launched a public AMI from a third party, run the `ec2-get-console-output` command on your local system (not on the instance), and locate the `SSH HOST KEY FINGERPRINTS` section. Note the fingerprints (for example, `1f:51:ae:28:bf:89:e9:d8:1f:25:5d:37:2d:7d:b8:ca:9f:f5:f1:6f`) so that you can compare them to the fingerprints of the instance.

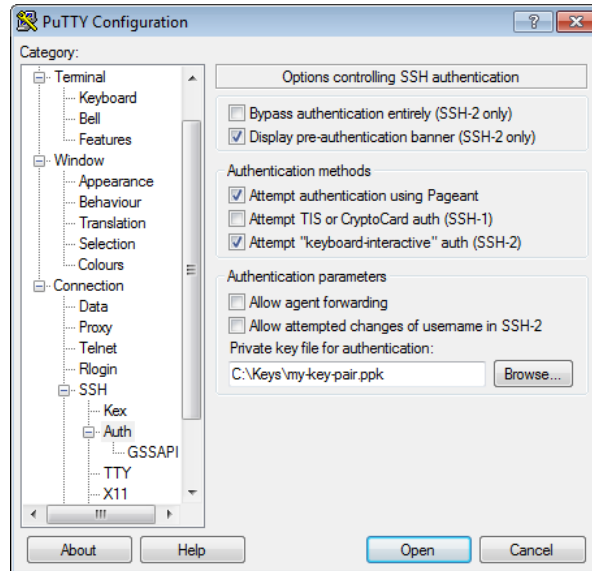
```
C:\> ec2-get-console-output instance_id
```

- Start PuTTY (from the **Start** menu, click **All Programs > PuTTY > PuTTY**).
- In the **Category** pane, select **Session** and complete the following fields:
  - In the **Host Name** box, enter `user_name@public_dns_name`. Be sure to specify the appropriate user name for your AMI. For example:
    - For an Amazon Linux AMI, the user name is `ec2-user`.
    - For a RHEL5 AMI, the user name is often `root` but might be `ec2-user`.
    - For an Ubuntu AMI, the user name is `ubuntu`.
    - Otherwise, check with your AMI provider.
  - Under **Connection type**, select **SSH**.
  - Ensure that **Port** is 22.



- In the **Category** pane, expand **Connection**, expand **SSH**, and then select **Auth**. Complete the following:
  - Click **Browse**.
  - Select the `.ppk` file that you generated for your key pair, and then click **Open**.

- c. (Optional) If you plan to start this session again later, you can save the session information for future use. Select **Session** in the **Category** tree, enter a name for the session in **Saved Sessions**, and then click **Save**.
- d. Click **Open** to start the PuTTY session.



5. If this is the first time you have connected to this instance, PuTTY displays a security alert dialog box that asks whether you trust the host you are connecting to.
6. (Optional) If you've launched a public AMI, verify that the fingerprint in the security alert matches the fingerprint that you obtained in step 1. If these fingerprints don't match, someone might be attempting a "man-in-the-middle" attack. If they match, continue to the next step.
7. Click **Yes**. A window opens and you are connected to your instance.

**Note**

If you specified a passphrase when you converted your private key to PuTTY's format, you must provide that passphrase when you log in to the instance.

## Transferring Files to Your Instance with the PuTTY Secure Copy Client

The PuTTY Secure Copy client (PSCP) is a command-line tool that you can use to transfer files between your Windows computer and your Linux/UNIX instance. If you prefer a graphical user interface (GUI), you can use an open source GUI tool named WinSCP. For more information, see [Transferring Files to Your Instance with WinSCP \(p. 278\)](#).

To use PSCP, you'll need the private key you generated in [Converting Your Private Key Using PuTTYgen \(p. 275\)](#). You'll also need the public DNS address of your Linux/UNIX instance.

The following example transfers the file `Sample_file.txt` from a Windows computer to the `/usr/local` directory on a Linux/UNIX instance:

```
C:\> pscp -i C:\Keys\my-key-pair.ppk C:\Sample_file.txt user_name@pub  
lic_dns:/usr/local/Sample_file.txt
```



## Transferring Files to Your Instance with WinSCP

WinSCP is a GUI-based file manager for Windows that allows you to upload and transfer files to a remote computer using the SFTP, SCP, FTP, and FTPS protocols. WinSCP allows you to drag and drop files from your Windows machine to your Linux instance or synchronize entire directory structures between the two systems.

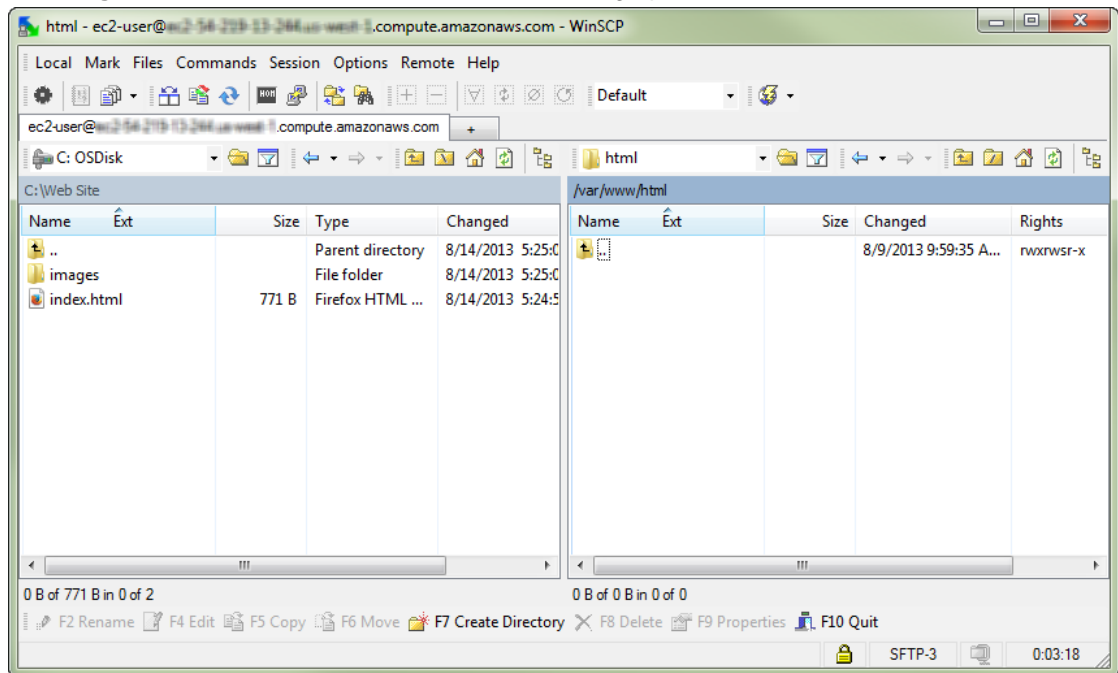
To use WinSCP, you'll need the private key you generated in [Converting Your Private Key Using PuTTYgen](#) (p. 275). You'll also need the public DNS address of your Linux/UNIX instance.

1. Download and install WinSCP from <http://winscp.net/eng/download.php>. For most users the default installation options are OK.
2. Start WinSCP.
3. At the **WinSCP login** screen, for **Host name**, enter the public DNS address for your instance.
4. For **User name**, enter the default user name for your AMI. For Amazon Linux AMIs, the user name is `ec2-user`. For Red Hat AMIs the user name is `root`, and for Ubuntu AMIs the user name is `ubuntu`.
5. For **Private key**, enter the path to your private key, or click the `...` button to browse for the file.

### Note

WinSCP requires a PuTTY private key file (`.ppk`). You can convert a `.pem` security key file to the `.ppk` format using PuTTYgen. For more information, see [Converting Your Private Key Using PuTTYgen](#) (p. 275).

6. (Optional) In the left panel, click **Directories**, and then, for **Remote directory**, enter the path for the directory you want to add files to.
7. Click **Login** to connect, and click **Yes** to add the host fingerprint to the host cache.



8. After the connection is established, in the connection window your Linux instance is on the right and your local machine is on the left. You can drag and drop files directly into the remote file system from your local machine. For more information on WinSCP, see the project documentation at <http://winscp.net/eng/docs/start>.

## Connecting to Your Linux/UNIX Instances Using SSH

After you launch your instance, you can connect to it and use it the way that you'd use a computer sitting in front of you.

If you receive an error while attempting to connect to your instance, see [Troubleshooting Connecting to Your Instance](#).

The following instructions explain how to connect to your instance using an SSH client.

### Topics

- [Prerequisites \(p. 279\)](#)
- [Connecting to Your Linux/UNIX Instance \(p. 279\)](#)
- [Transferring Files to Linux/UNIX Instances from Linux/UNIX with SCP \(p. 280\)](#)

### Prerequisites

- **Install an SSH client**  
Your Linux computer mostly likely includes an SSH client by default. You can check for an SSH client by typing `ssh` at the command line. If your computer doesn't recognize the command, the OpenSSH project provides a free implementation of the full suite of SSH tools. For more information, see <http://www.openssh.org>.
- **Install the Amazon EC2 CLI Tools**  
(Optional) If you're using a public AMI from a third party, use the `ec2-get-console-output` command to verify the fingerprint. For more information, see [Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX \(p. 541\)](#).
- **Get the ID of the instance**  
You can get the ID of your instance using the Amazon EC2 console (from the **Instance ID** column). If you prefer, you can use the `ec2-describe-instances` command.
- **Get the public DNS name of the instance**  
You can get the public DNS for your instance using the Amazon EC2 console (check the **Public DNS** column; if this column is hidden, click the **Show/Hide** icon and select **Public DNS**). If you prefer, you can use the `ec2-describe-instances` command.
- **Locate the private key**  
You'll need the fully-qualified path of the `.pem` file for the key pair that you specified when you launched the instance.
- **Enable inbound SSH traffic from your IP address to your instance**  
Ensure that the security group associated with your instance allows incoming SSH traffic from your IP address. For more information, see [Authorizing Network Access to Your Instances](#).

## Connecting to Your Linux/UNIX Instance

### To connect to your instance using SSH

1. (Optional) If you've launched a public AMI from a third party, run the `ec2-get-console-output` command on your local system (not on the instance), and locate the `SSH HOST KEY FINGERPRINTS` section. Note the fingerprints (for example, `1f:51:ae:28:bf:89:e9:d8:1f:25:5d:37:2d:7d:b8:ca:9f:f5:f1:6f`) so that you can compare them to the fingerprints of the instance.

```
ec2-get-console-output instance_id
```

2. In a command line shell, change directories to the location of the private key file that you created when you launched the instance.
3. Use the **chmod** command to make sure your private key file isn't publicly viewable. For example, if the name of your private key file is `my-key-pair.pem`, you would use the following command:

```
chmod 400 my-key-pair.pem
```

4. Use the **ssh** command to connect to the instance. You'll specify the private key (`.pem`) file and `user_name@public_dns_name`. For Amazon Linux, the user name is `ec2-user`. For RHEL5, the user name is often `root` but might be `ec2-user`. For an Ubuntu, AMI the user name is `ubuntu`. Otherwise, check with your AMI provider.

```
ssh -i my-key-pair.pem ec2-user@ec2-198-51-100-1.compute-1.amazonaws.com
```

You'll see a response like the following.

```
The authenticity of host 'ec2-198-51-100-1.compute-1.amazonaws.com
(10.254.142.33)'
can't be established.
RSA key fingerprint is
1f:51:ae:28:bf:89:e9:d8:1f:25:5d:37:2d:7d:b8:ca:9f:f5:f1:6f.
Are you sure you want to continue connecting (yes/no)?
```

5. (Optional) If you've launched a public AMI, verify that the fingerprint in the security alert matches the fingerprint that you obtained in step 1. If these fingerprints don't match, someone might be attempting a "man-in-the-middle" attack. If they match, continue to the next step.
6. Enter `yes`.

You'll see a response like the following.

```
Warning: Permanently added 'ec2-198-51-100-1.compute-1.amazonaws.com' (RSA)
to the list of known hosts.
```

## Transferring Files to Linux/UNIX Instances from Linux/UNIX with SCP

One way to transfer files between your local computer and a Linux/UNIX instance is to use Secure Copy (SCP). This section describes how to transfer files with SCP. The procedure is very similar to the procedure for connecting to an instance with SSH.

### Prerequisites

- **Install an SCP client**

Most Linux, UNIX, and Apple computers include an SCP client by default. If yours doesn't, the OpenSSH project provides a free implementation of the full suite of SSH tools, including an SCP client. For more information, go to <http://www.openssh.org>.

- **Get the ID of the instance**

You can get the ID of your instance using the Amazon EC2 console (from the **Instance ID** column). If you prefer, you can use the `ec2-describe-instances` command.

- **Get the public DNS name of the instance**

You can get the public DNS for your instance using the Amazon EC2 console (check the **Public DNS** column; if this column is hidden, click the **Show/Hide** icon and select **Public DNS**). If you prefer, you can use the [ec2-describe-instances](#) command.

- **Locate the private key**

You'll need the fully-qualified path of the `.pem` file for the key pair that you specified when you launched the instance.

- **Enable inbound SSH traffic from your IP address to your instance**

Ensure that the security group associated with your instance allows incoming SSH traffic from your IP address. For more information, see [Authorizing Network Access to Your Instances](#).

The following procedure steps you through using SCP to transfer a file. If you've already connected to the instance with SSH and have verified its fingerprints, you can start with the step that contains the SCP command (step 4).

### To use SCP to transfer a file

1. (Optional) If you've launched a public AMI from a third party, run the [ec2-get-console-output](#) command on your local system (not on the instance), and locate the `SSH HOST KEY FINGERPRINTS` section. Note the fingerprints (for example, `1f:51:ae:28:bf:89:e9:d8:1f:25:5d:37:2d:7d:b8:ca:9f:f5:f1:6f`) so that you can compare them to the fingerprints of the instance.

```
ec2-get-console-output instance_id
```

2. In a command shell, change directories to the location of the private key file that you specified when you launched the instance.
3. Use the `chmod` command to make sure your private key file isn't publicly viewable. For example, if the name of your private key file is `my-key-pair.pem`, you would use the following command:

```
chmod 400 my-key-pair.pem
```

4. Transfer a file to your instance using the instance's public DNS name. For example, if the name of the private key file is `my-key-pair`, the file to transfer is `SampleFile.txt`, and the public DNS name of the instance is `ec2-198-51-100-1.compute-1.amazonaws.com`, use the following command to copy the file to the `ec2-user` home directory.

```
scp -i my-key-pair.pem SampleFile.txt ec2-user@ec2-198-51-100-1.compute-1.amazonaws.com:~
```

#### Tip

For Amazon Linux, the user name is `ec2-user`. For RHEL5, the user name is often `root` but might be `ec2-user`. For an Ubuntu, AMI the user name is `ubuntu`. Otherwise, check with your AMI provider.

You'll see a response like the following.

```
The authenticity of host 'ec2-198-51-100-1.compute-1.amazonaws.com
(10.254.142.33)'
can't be established.
RSA key fingerprint is
1f:51:ae:28:bf:89:e9:d8:1f:25:5d:37:2d:7d:b8:ca:9f:f5:f1:6f.
Are you sure you want to continue connecting (yes/no)?
```

- (Optional) If you've launched a public AMI, verify that the fingerprint in the security alert matches the fingerprint that you obtained in step 1. If these fingerprints don't match, someone might be attempting a "man-in-the-middle" attack. If they match, continue to the next step.
- Enter **yes**.

You'll see a response like the following.

```
Warning: Permanently added 'ec2-198-51-100-1.compute-1.amazonaws.com' (RSA)
to the list of known hosts.
Sending file modes: C0644 20 SampleFile.txt
Sink: C0644 20 SampleFile.txt
SampleFile.txt                               100%   20    0.0KB/s   00:00
```

To transfer files in the other direction (from your Amazon EC2 instance to your local computer), simply reverse the order of the host parameters. For example, to transfer the SampleFile.txt file from your Amazon EC2 instance back to the home directory on your local computer as SampleFile2.txt, use the following command on your local computer.

```
scp -i my-key-pair.pem ec2-user@ec2-198-51-100-1.compute-1.amazonaws.com:~/Sample
File.txt ~/SampleFile2.txt
```

## Connecting to Windows Instances Using RDP

After you launch your instance, you can connect to it and use it the way that you'd use a computer sitting in front of you.

If you receive an error while attempting to connect to your instance, see [Troubleshooting Connecting to Your Instance](#).

The following instructions explain how to connect to your instance using an RDP client.

### Topics

- [Prerequisites \(p. 282\)](#)
- [Connecting to Your Windows Instance \(p. 283\)](#)
- [Transfer Files to Windows Server Instances from Windows \(p. 284\)](#)

### Prerequisites

- Install an RDP client**

Your Windows computer includes an RDP client by default. You can check for an RDP client by typing **mstsc** at a Command Prompt window. If your computer doesn't recognize this command, see the [Microsoft Windows home page](#) and search for the download for Remote Desktop Connection. For Mac OS X, you can use [Microsoft's Remote Desktop Client](#). For Linux/UNIX, you can use [rdesktop](#).

- Get the ID of the instance**

You can get the ID of your instance using the Amazon EC2 console (from the **Instance ID** column). If you prefer, you can use the [ec2-describe-instances](#) command.

- Get the public DNS name of the instance**

You can get the public DNS for your instance using the Amazon EC2 console (check the **Public DNS** column; if this column is hidden, click the **Show/Hide** icon and select **Public DNS**). If you prefer, you can use the [ec2-describe-instances](#) command.

- **Locate the private key**  
You'll need the fully-qualified path of the `.pem` file for the key pair that you specified when you launched the instance.
- **Enable inbound RDP traffic from your IP address to your instance**  
Ensure that the security group associated with your instance allows incoming RDP traffic from your IP address. For more information, see [Authorizing Network Access to Your Instances](#).
- For the best experience using Internet Explorer, run the latest version.

## Connecting to Your Windows Instance

To connect to a Windows instance, you must retrieve the initial administrator password, and then specify this password when you connect to your instance using Remote Desktop.

### Note

Windows instances are limited to two simultaneous remote connections at one time. If you attempt a third connection, an error will occur. For more information, see [Configure the Number of Simultaneous Remote Connections Allowed for a Connection](#).

### To connect to your Windows instance

1. In the Amazon EC2 console, select the instance, and then click **Connect**.
2. In the **Connect To Your Instance** dialog box, click **Get Password** (it will take a few minutes after the instance is launched before the password is available).
3. Click **Browse** and navigate to the private key file you created when you launched the instance. Select the file and click **Open** to copy the entire contents of the file into contents box.
4. Click **Decrypt Password**. The console displays the default administrator password for the instance in the **Connect To Your Instance** dialog box, replacing the link to **Get Password** shown previously with the actual password.
5. Record the default administrator password, or copy it to the clipboard. You need this password to connect to the instance.
6. Click **Download Remote Desktop File**. Your browser prompts you to either open or save the `.rdp` file. Either option is fine. When you have finished, you can click **Close** to dismiss the **Connect To Your Instance** dialog box.
7. If you opened the `.rdp` file, you'll see the **Remote Desktop Connection** dialog box. If you saved the `.rdp` file, navigate to your downloads directory, and double-click the `.rdp` file to display the dialog box. You may get a warning that the publisher of the remote connection is unknown. Click **Connect** to connect to your instance. You may get a warning that the security certificate could not be authenticated. Click **Yes** to continue.
8. Log in to the instance as prompted, using `Administrator` as the user name and the default administrator password that you recorded or copied previously. If the user name includes a domain (`domain\Administrator`), delete the text before `Administrator`.

After you connect, we recommend that you do the following:

- Change the Administrator password from the default value. You change the password while logged on to the instance itself, just as you would on any other Windows Server.
- Create another user account with administrator privileges on the instance. Another account with administrator privileges is a safeguard if you forget the Administrator password or have a problem with the Administrator account.

## Transfer Files to Windows Server Instances from Windows

You can work with your instance the same way you would work with any Windows server. For example, you can transfer files between an Amazon EC2 Windows instance and your local Windows computer using the local file sharing feature of Windows Remote Desktop. If you enable this option in your Windows Remote Desktop Connection software, you can access your local files from your Amazon EC2 Windows instances. You can access local files on hard disk drives, DVD drives, portable media drives, and mapped network drives. For information about this feature, go to the [Microsoft Support website](#) or go to [The most useful feature of Remote Desktop I never knew about](#) on the MSDN Blogs website.

## Stop and Start Your Instance

You can stop and restart your instance if it has an Amazon EBS volume as its root device. The instance retains its instance ID, but can change as described in the Overview section.

When you stop an instance, we shut it down. We don't charge hourly usage for a stopped instance, or data transfer fees, but we do charge for the storage for any Amazon EBS volumes. Each time you start a stopped instance we charge a full instance hour, even if you make this transition multiple times within a single hour. If you have stopped your Amazon EBS-backed instance and it appears "stuck" in the `stopping` state, you can forcibly stop it. For more information, see [Troubleshooting Stopping Your Instance](#) (p. 358).

While the instance is stopped, you can treat its root volume like any other volume, and modify it (for example, repair file system problems or update software). You just detach the volume from the stopped instance, attach it to a running instance, make your changes, detach it from the running instance, and then reattach it to the stopped instance. Make sure that you reattach it using the storage device name that's specified as the root device in the block device mapping for the instance.

If you decide that you no longer need an instance, you can terminate it. As soon as the state of an instance changes to `shutting-down` or `terminated`, we stop charging for that instance. For more information, see [Terminate Your Instance](#) (p. 287).

## Overview

When you stop a running instance, the following happens:

- The instance performs a normal shutdown and stops running; its status changes to `stopping` and then `stopped`.
- Any Amazon EBS volumes remain attached to the instance, and their data persists.
- Any data stored in the RAM of the host computer or the instance store volumes of the host computer is gone.
- EC2-Classic: We release the public and private IP addresses for the instance when you stop the instance, and assign new ones when you restart it.

EC2-VPC: The instance retains its private IP addresses when stopped and restarted. We release the public IP address and assign a new one when you restart it.

- EC2-Classic: We disassociate any Elastic IP address that's associated with the instance. You're charged for Elastic IP addresses that aren't associated with an instance. When you restart the instance, you must associate the Elastic IP address with the instance; we don't do this automatically.

EC2-VPC: The instance retains its associated Elastic IP addresses. You're charged for any Elastic IP addresses associated with a stopped instance.

- When you stop and restart a Windows instance, by default, we change the instance host name to match the new IP address and initiate a reboot. By default, we also change the drive letters for any attached Amazon EBS volumes. For more information about these defaults and how you can change them, see



[Configuring a Windows Instance Using the EC2Config Service](#) in the *Amazon Elastic Compute Cloud Microsoft Windows Guide*.

- If you've registered the instance with a load balancer, it's likely that the load balancer won't be able to route traffic to your instance after you've stopped and restarted it. You must de-register the instance from the load balancer after stopping the instance, and then re-register after starting the instance. For more information, see [De-Registering and Registering Amazon EC2 Instances](#) in the *Elastic Load Balancing Developer Guide*.

## Stopping and Starting Your Instances

You can only stop an Amazon EBS-backed instance.

### To verify the root device type of your instance

1. In the navigation pane, click **Instances**, and select the instance.
2. Check the value of **Root device type** in the details pane as follows:
  - If the value is `ebs`, this is an Amazon EBS-backed instance.
  - If the value is `instance store`, this is an instance store-backed instance. You can't stop an instance store-backed instance.

You can start and stop your Amazon EBS-backed instance using the AWS Management Console as follows. (If you prefer, you can use the [ec2-stop-instances](#) and [ec2-start-instances](#) commands.)

### To stop and start an Amazon EBS-backed instance

1. In the navigation pane, click **Instances**, and select the instance.
2. [EC2-Classic] If the instance has an associated Elastic IP address, write down the Elastic IP address and the instance ID shown in the details pane.
3. Click **Actions**, and then click **Stop**. If **Stop** is disabled, either the instance is already stopped or its root device is an instance store volume.
4. In the confirmation dialog box, click **Yes, Stop**. It can take a few minutes for the instance to stop.

[EC2-Classic] When the instance state becomes `stopped`, the **Elastic IP**, **Public DNS**, **Private DNS**, and **Private IPs** fields in the details pane are blank to indicate that the old values are no longer associated with the instance.

5. While your instance is stopped, you can modify certain instance attributes. For more information, see [Modifying a Stopped Instance \(p. 286\)](#).
6. To restart the stopped instance, select the instance, click **Actions**, and then click **Start**.
7. In the confirmation dialog box, click **Yes, Start**. It can take a few minutes for the instance to enter the `running` state.

[EC2-Classic] When the instance state becomes `running`, the **Public DNS**, **Private DNS**, and **Private IPs** fields in the details pane contain the new values that we assigned to the instance.

8. [EC2-Classic] If your instance had an associated Elastic IP address, you must reassociate it as follows:
  - a. In the navigation pane, click **Elastic IPs**.
  - b. Select the Elastic IP address that you wrote down before you stopped the instance.
  - c. Click **Associate Address**.
  - d. Select the instance ID that you wrote down before you stopped the instance, and then click **Associate**.



## Modifying a Stopped Instance

You can modify the following attributes of an instance only when it is stopped:

- Instance type
- User data
- Kernel
- RAM disk

If you try to modify these attributes using the CLI or API while the instance is running, Amazon EC2 returns the `IncorrectInstanceState` error.

You can change the instance type and user data attributes using the AWS Management Console as shown here. You can't use the AWS Management Console to modify the kernel or RAM disk attributes.

### To change the instance type for a stopped instance using the console

1. In the navigation pane, click **Instances**.
2. Select the stopped instance, click **Actions**, and then click **Change Instance Type**.
3. In the **Change Instance Type** dialog box, in the **Instance Type** list, select the type of instance you need, and then click **Apply**.

For more information, see [Resizing Your Instance \(p. 111\)](#).

### To change the user data for a stopped instance using the console

1. In the navigation pane, click **Instances**.
2. Select the stopped instance, click **Actions**, and then click **View/Change User Data**.
3. In the **View/Change User Data** dialog box, update the user data, and then click **Save**. Note that you can't change the user data if the instance is running, but you can view it.

If you prefer, you can use the `ec2-modify-instance-attribute` command. This command can change the instance type, user data, kernel, and RAM disk.

## Reboot Your Instance

An instance reboot is equivalent to an operating system reboot. In most cases, it takes only a few minutes to reboot your instance. When you reboot an instance, it remains on the same physical host, so your instance keeps its public DNS name, private IP address, and any data on its instance store volumes.

Rebooting an instance doesn't start a new instance billing hour, unlike stopping and restarting your instance.

We might schedule your instance for a reboot for necessary maintenance, such as to apply updates that require a reboot. No action is required on your part; we recommend that you wait for the reboot to occur within its scheduled window. For more information, see [Monitoring Events for Your Instances \(p. 349\)](#).

We recommend that you use Amazon EC2 to reboot your instance instead of running the operating system reboot command from your instance.

Use the following procedure to reboot an instance using the console. If you prefer, you can use the `ec2-reboot-instances` command instead.

### To reboot an instance

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Instances**.
3. Select the instance, click **Actions**, and then click **Reboot**.
4. Click **Yes, Reboot** when prompted for confirmation.

## Terminate Your Instance

When you've decided that you no longer need an instance, you can terminate it. As soon as the state of an instance changes to `shutting-down` or `terminated`, you stop incurring charges for that instance.

You can't connect to or restart an instance after you've terminated it. However, you can launch additional instances using the same AMI. If you'd rather stop and restart your instance, see [Stop and Start Your Instance \(p. 284\)](#).

If your instance is in the `shutting-down` state for longer than usual, it will eventually be cleaned up (terminated) by automated processes within the Amazon EC2 service. For more information, see [Troubleshooting Terminating \(Shutting Down\) Your Instance \(p. 359\)](#).

### Topics

- [Instance Termination \(p. 287\)](#)
- [Terminating an Instance \(p. 288\)](#)
- [Enabling Termination Protection for an Instance \(p. 288\)](#)
- [Changing the Instance Initiated Shutdown Behavior \(p. 289\)](#)
- [Preserving Amazon EBS Volumes on Instance Termination \(p. 289\)](#)

## Instance Termination

After you terminate an instance, it remains visible in the console for a short while, and then the entry is deleted.

When an instance terminates, the data on any instance store volumes associated with that instance is deleted.

By default, any Amazon EBS volumes that you attach as you launch the instance are automatically deleted when the instance terminates. However, by default, any volumes that you attach to a running instance persist even after the instance terminates. This behavior is controlled by the volume's `DeleteOnTermination` attribute, which you can modify. For more information, see [Preserving Amazon EBS Volumes on Instance Termination \(p. 289\)](#).

You can prevent an instance from being terminated accidentally by someone using the AWS Management Console, the CLI, and the API. This feature is available for both Amazon EC2 instance store-backed and Amazon EBS-backed instances. Each instance has a `DisableApiTermination` attribute with the default value of `false` (the instance can be terminated through Amazon EC2). You can modify this instance attribute while the instance is running or stopped (in the case of Amazon EBS-backed instances). For more information, see [Enabling Termination Protection for an Instance \(p. 288\)](#).

You can control whether an instance should stop or terminate when shutdown is initiated from the instance using the operating system command for system shutdown. For more information, see [Changing the Instance Initiated Shutdown Behavior \(p. 289\)](#).

If you run a script on instance termination, your instance might have an abnormal termination, because we have no way to ensure that shutdown scripts run. Amazon EC2 attempts to shut an instance down

cleanly and run any system shutdown scripts; however, certain events (such as hardware failure) may prevent these system shutdown scripts from running.

## Terminating an Instance

Use the following procedure to terminate an instance using the console. If you prefer, you can use the [ec2-terminate-instances](#) command instead.

### To terminate an instance

1. Before you terminate the instance, verify that you won't lose any data by checking that your Amazon EBS volumes won't be deleted on termination and that you've copied any data that you need from your instance store volumes to Amazon EBS or Amazon S3.
2. Open the Amazon EC2 console.
3. In the navigation pane, click **Instances**.
4. Select the instance, click **Actions**, and then click **Terminate**.
5. Click **Yes, Terminate** when prompted for confirmation.

## Enabling Termination Protection for an Instance

By default, you can terminate your instance using the Amazon EC2 console, command line interface, or API. If you want to prevent your instance from being accidentally terminated using Amazon EC2, you can enable *termination protection* for the instance. The `DisableApiTermination` attribute controls whether the instance can be terminated using the console, CLI, or API. By default, termination protection is disabled for your instance. You can set the value of this attribute when you launch the instance, while the instance is running, or while the instance is stopped (for Amazon EBS-backed instances).

The `DisableApiTermination` attribute does not prevent you from terminating an instance by initiating shutdown from the instance (using the operating system command for system shutdown) when the `InstanceInitiatedShutdownBehavior` attribute is set. For more information, see [Changing the Instance Initiated Shutdown Behavior](#) (p. 289).

Instances that are part of an Auto Scaling group are not covered by termination protection. For more information, see [Instance Termination Policy for Your Auto Scaling Group](#) in the *Auto Scaling Developer Guide*.

Use the following procedures to enable or disable termination protection using the console. If you prefer, you can use the `ec2-modify-instance-attribute` command instead.

### To enable termination protection for an instance at launch time

1. On the dashboard of the Amazon EC2 console, click **Launch Instance** and follow the directions in the wizard.
2. On the **Configure Instance Details** page, select the **Enable termination protection** check box.

### To enable termination protection for a running or stopped instance

1. Select the instance, click **Actions**, and then click **Change Termination Protection**.
2. Click **Yes, Enable**.

### To disable termination protection for a running or stopped instance

1. Select the instance, click **Actions**, and then click **Change Termination Protection**.
2. Click **Yes, Disable**.

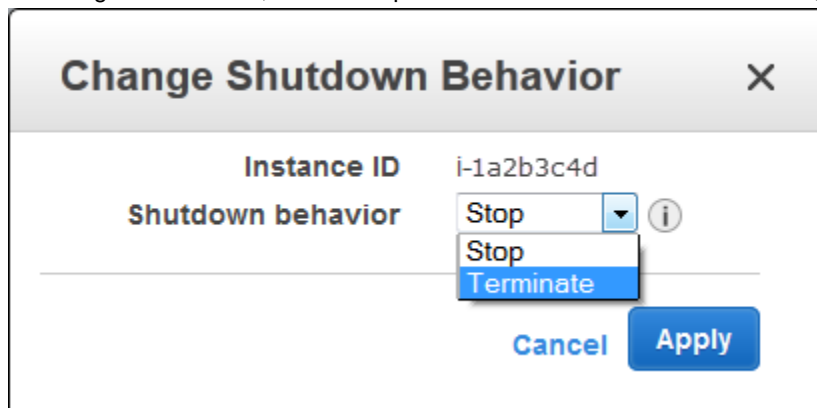
## Changing the Instance Initiated Shutdown Behavior

By default, when you initiate a shutdown from an Amazon EBS-backed instance using the `shutdown` command, the instance stops. You can change this behavior using the `InstanceInitiatedShutdownBehavior` attribute for the instance so that it terminates instead. You can update this attribute while the instance is running or stopped.

You can view and update the `InstanceInitiatedShutdownBehavior` attribute using the AWS Management Console as follows. If you prefer, you can use the `ec2-describe-instance-attribute` and `ec2-modify-instance-attribute` commands instead.

### To change the shutdown behavior of an instance

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Instances**.
3. Select the instance, click **Actions**, and then click **Change Shutdown Behavior**. The current behavior is already selected.
4. To change the behavior, select an option from the **Shutdown behavior** list, and then click **Apply**.



## Preserving Amazon EBS Volumes on Instance Termination

By default, we do the following:

- Preserve any volumes that you attach to a running instance even after the instance terminates
- Preserve any volumes that you attach to your instance at launch when you stop and restart an instance
- Delete the volumes that you attach to your instance at launch when you terminate the instance

You can change this behavior using the `DeleteOnTermination` attribute for the volume. If the value of this attribute is `true`, we delete the volume after the instance terminates; otherwise, we preserve the volume. If the `DeleteOnTermination` attribute of a volume is `false`, the volume persists in its current state. You can take a snapshot of the volume, and you can attach it to another instance.

If you detach a volume that you attached to your instance at launch, and then reattach it, we preserve it even after the instance terminates. In other words, its `DeleteOnTermination` attribute is set to `false`.

You can see the value for the `DeleteOnTermination` attribute on the volumes attached to an instance by looking at the instance's block device mapping. For more information, see [Viewing the EBS Volumes in an Instance Block Device Mapping \(p. 525\)](#).

You can change the value of a volume's `DeleteOnTermination` attribute while the instance is running using the `ec2-modify-instance-attribute` command. The following example preserves the specified volumes by setting their `DeleteOnTermination` attributes to `false`.

```
ec2-modify-instance-attribute -b "/dev/sda1::false" -b "/dev/sdh=snap  
shot_id::false"
```

Note that you can also use the `-b` option to set the `DeleteOnTermination` attribute when you launch an instance using the `ec2-run-instances` command, as shown in the following example.

```
ec2-run-instances ami-1a2b3c4d -k my-key-pair -b "/dev/sda1::false"
```

For information about how to set this attribute while launching an instance using the console, see [Changing the Root Device Volume to Persist \(p. 14\)](#).

## Instance Metadata and User Data

*Instance metadata* is data about your EC2 instance that you can use to configure or manage the running instance. Instance metadata is divided into categories. For more information, see [Instance Metadata Categories \(p. 297\)](#).

EC2 instances can also include *dynamic data*, such as an instance identity document that is generated when the instance is launched. For more information, see [Dynamic Data Categories \(p. 300\)](#).

You can also access the *user data* that you supplied when launching your EC2 instance. For example, you can specify parameters for configuring your instance, or attach a simple script. You can also use this data to build more generic AMIs that can be modified by configuration files supplied at launch time. For example, if you run web servers for various small businesses, they can all use the same AMI and retrieve their content from the Amazon S3 bucket you specify in the user data at launch. To add a new customer at any time, simply create a bucket for the customer, add their content, and launch your AMI. If you launch more than one instance at the same time, the user data is available to all instances in that reservation.

Because you can access instance metadata and user data from within your running instance, you do not need to use the Amazon EC2 console or the CLI tools. This can be helpful when you're writing scripts to run from within your instance. For example, you can access your instance's local IP address from within the running instance to manage a connection to an external application.

### Important

Although you can only access instance metadata and user data from within the instance itself, the data is not protected by cryptographic methods. Anyone who can access the instance can view its metadata. Therefore, you should take suitable precautions to protect sensitive data (such as long-lived encryption keys). You should not store sensitive data, such as passwords, as user data.

For more information about adding user data when you launch an instance, see [Launching an Instance \(p. 266\)](#). You can add or modify user data on Amazon EBS-backed instances when they're stopped. For more information about adding user data to a stopped instance, see [Modifying a Stopped Instance \(p. 286\)](#).

When you are adding user data, take note of the following:

- User data is treated as opaque data: what you give is what you get back. It is up to the instance to be able to interpret it.
- User data is limited to 16 KB. This limit applies to the data in raw form, not base64-encoded form.

- User data must be base64-encoded before being submitted to the API. The API command line tools perform the base64 encoding for you. The data is decoded before being presented to the instance. For more information about base64 encodings, go to <http://tools.ietf.org/html/rfc4648>.

#### Topics

- [Retrieving Instance Metadata \(p. 291\)](#)
- [Retrieving User Data \(p. 293\)](#)
- [Retrieving Dynamic Data \(p. 294\)](#)
- [Example: AMI Launch Index Value \(p. 294\)](#)
- [Instance Metadata Categories \(p. 297\)](#)

## Retrieving Instance Metadata

To view all categories of instance metadata from within a running instance, use the following URI:

```
http://169.254.169.254/latest/meta-data/
```

Note that you are not billed for HTTP requests used to retrieve instance metadata and user data.

On a Linux instance, you can use a tool such as cURL, or use the GET command, for example:

```
$ GET http://169.254.169.254/latest/meta-data/
```

You can also download the Instance Metadata Query tool, which allows you to query the instance metadata without having to type out the full URI or category names:

<http://aws.amazon.com/code/1825>

On a Windows instance, you can install a tool such as GNU Wget or cURL to retrieve instance metadata at the command line, or you can copy and paste the URI into a browser. If you do not want to install any third-party tools, you can use PowerShell cmdlets to retrieve the URI. For example, if you are running version 3.0 or later of PowerShell, use the following cmdlet:

```
C:\> invoke-restmethod -uri http://169.254.169.254/latest/meta-data/
```

#### Important

If you do install a third-party tool on a Windows instance, ensure that you read the accompanying documentation carefully, as the method of calling the HTTP and the output format might be different from what is documented here.

All metadata is returned as text (content type text/plain). A request for a specific metadata resource returns the appropriate value, or a 404 - Not Found HTTP error code if the resource is not available.

A request for a general metadata resource (the URI ends with a /) returns a list of available resources, or a 404 - Not Found HTTP error code if there is no such resource. The list items are on separate lines, terminated by line feeds (ASCII 10).

## Examples of Retrieving Instance Metadata

The following are examples of requests and responses on a Linux instance.

This example gets the available versions of the instance metadata. These versions do not necessarily correlate with an Amazon EC2 API version. The earlier versions are available to you in case you have scripts that rely on the structure and information present in a previous version.

```
$ GET http://169.254.169.254/  
1.0  
2007-01-19  
2007-03-01  
2007-08-29  
2007-10-10  
2007-12-15  
2008-02-01  
2008-09-01  
2009-04-04  
2011-01-01  
2011-05-01  
2012-01-12  
latest
```

This example gets the top-level metadata items. Some items are only available for instances in a VPC. For more information about each of these items, see [Instance Metadata Categories](#) (p. 297).

```
$ GET http://169.254.169.254/latest/meta-data/  
ami-id  
ami-launch-index  
ami-manifest-path  
block-device-mapping/  
hostname  
instance-action  
instance-id  
instance-type  
kernel-id  
local-hostname  
local-ipv4  
mac  
network/  
placement/  
public-hostname  
public-ipv4  
public-keys/  
reservation-id  
security-groups
```

These examples get the value of some of the metadata items from the preceding example.

```
$ GET http://169.254.169.254/latest/meta-data/ami-id  
ami-2bb65342
```

```
$ GET http://169.254.169.254/latest/meta-data/reservation-id  
r-fea54097
```

```
$ GET http://169.254.169.254/latest/meta-data/hostname  
ec2-203-0-113-25.compute-1.amazonaws.com
```

This example gets the list of available public keys.

```
$ GET http://169.254.169.254/latest/meta-data/public-keys/  
0=my-public-key
```

This example shows the formats in which public key 0 is available.

```
$ GET http://169.254.169.254/latest/meta-data/public-keys/0  
openssh-key
```

This example gets public key 0 (in the OpenSSH key format).

```
$ GET http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key  
ssh-rsa MIICiTCcAfICCQD6m7oRw0uXOjANBgkqhkiG9w0BAQUFADCBiDELMAKGA1UEBhMC  
VVMxCzAJBgNVBAGTAldBMRAdG9YDVQQHEwdTZWF0dGx1MQ8wDQYDVQQKEwZBbWF6  
b24xFDASBgNVBA5TC0lBTsBDb25zb2xlMRlWZAYDVQQDEw1UZXR0Q2lsYWMxH3Ad  
BgkqhkiG9w0BCQEWEG5vb25lQGftYXpvi5jb20wHhcNMTEwNDI1MjA0NTIxWhcN  
MTIwNDI0MjA0NTIxWjCBiDELMAKGA1UEBhMCMVVMxCzAJBgNVBAGTAldBMRAdG9YD  
VQQHEwdTZWF0dGx1MQ8wDQYDVQQKEwZBbWF6b24xFDASBgNVBA5TC0lBTsBDb25z  
b2xlMRlWZAYDVQQDEw1UZXR0Q2lsYWMxH3AdBgkqhkiG9w0BCQEWEG5vb25lQGft  
YXpvi5jb20wgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBAMaK0dn+a4GmWIWJ  
21uUSfwfEvySWtC2XADZ4nB+BLYgVIk60CpiwsZ3G93vUEIO3IyNoH/f0wYK8m9T  
rDHudUZg3qX4waLG5M43q7Wgc/MbQITxOUSQv7c7ugFFDzQGBzZswY6786m86gpE  
Ibb3OhjZnzcVQAaRHhd1QWIMm2nrAgMBAAEwDQYJKoZIhvcNAQEFBQADgYEAtCu4  
nUhVVxYUntned9+h8Mg9q6q+auNKyExzyLwaxlAoo7TJHidbtS4J5iNmZgXL0Fkb  
FFBjvSfpJiLJ00zbnNYS5f6GuoEDmFJl0ZxBHjJnyp378OD8uTs7fLvJx79LjSTb  
NYiytVbZPQUQ5Yaxu2jXnimvw3rrszlaEXAMPLE my-public-key
```

This example shows the information available for a specific network interface (indicated by the MAC address) on an NAT instance in the EC2-Classic platform.

```
$ GET http://169.254.169.254/latest/meta-data/network/inter  
faces/macs/02:29:96:8f:6a:2d/  
device-number  
local-hostname  
local-ipv4s  
mac  
owner-id  
public-hostname  
public-ipv4s
```

This example gets the subnet ID for an EC2 instance launched into a VPC.

```
$ GET http://169.254.169.254/latest/meta-data/network/inter  
faces/macs/02:29:96:8f:6a:2d/subnet-id  
subnet-be9b61d7
```

## Retrieving User Data

To retrieve user data, use the following URI:

```
http://169.254.169.254/latest/user-data
```



Requests for user data returns the data as it is (content type application/x-octetstream).

This shows an example of returning comma-separated user data.

```
$ GET http://169.254.169.254/latest/user-data
1234, john, reboot, true | 4512, richard, | 173, , ,
```

This shows an example of returning line-separated, user data.

```
$ GET http://169.254.169.254/latest/user-data
[general]
instances: 4

[instance-0]
s3-bucket: <user_name>

[instance-1]
reboot-on-error: yes
```

## Retrieving Dynamic Data

To retrieve dynamic data from within a running instance, use the following URI:

```
http://169.254.169.254/latest/dynamic/
```

This example shows how to retrieve the high-level instance identity categories:

```
GET http://169.254.169.254//latest/dynamic/instance-identity/
pkcs7
signature
document
```

## Example: AMI Launch Index Value

This example demonstrates how you can use both user data and instance metadata to configure your instances.

Alice wants to launch four instances of her favorite Linux database AMI, with the first acting as master and the remaining three acting as replicas. When she launches them, she wants to add user data about the replication strategy for each replicant. She is aware that this data will be available to all four instances, so she needs to structure the user data in a way that allows each instance to recognize which parts are applicable to it. She can do this using the `ami-launch-index` instance metadata value, which will be unique for each instance.

Here is the user data that Alice has constructed:

```
replicate-every=1min | replicate-every=5min | replicate-every=10min
```

The `replicate-every=1min` data defines the first replicant's configuration, `replicate-every=5min` defines the second replicant's configuration, and so on. Alice decides to provide this data as an ASCII string with a pipe symbol (|) delimiting the data for the separate instances.

Alice launches four instances, specifying the user data:

**Amazon Elastic Compute Cloud User Guide**  
**Example: AMI Launch Index Value**

---

```
PROMPT> ec2-run-instances ami-2bb65342 -n 4 -d "replicate-every=1min | replicate-  
every=5min | replicate-every=10min"
```

```
RESERVATION      r-fea54097          598916040194    default  
INSTANCE i-10a64379  ami-2bb65342 pending 0 m1.small 2010-03-19T13:59:03+0000  
us-east-1a aki-94c527fd ari-96c527ff monitoring-disabled ebs  
INSTANCE i-10a64380  ami-2bb65342 pending 0 m1.small 2010-03-19T13:59:03+0000  
us-east-1a aki-94c527fd ari-96c527ff monitoring-disabled ebs  
INSTANCE i-10a64381  ami-2bb65342 pending 0 m1.small 2010-03-19T13:59:03+0000  
us-east-1a aki-94c527fd ari-96c527ff monitoring-disabled ebs  
INSTANCE i-10a64382  ami-2bb65342 pending 0 m1.small 2010-03-19T13:59:03+0000  
us-east-1a aki-94c527fd ari-96c527ff monitoring-disabled ebs
```

After they're launched, all instances have a copy of the user data and the common metadata shown here:

- AMI id: ami-2bb65342
- Reservation ID: r-fea54097
- Public keys: none
- Security group name: default
- Instance type: m1.small

However, each instance has certain unique metadata.

*Instance 1*

Metadata	Value
instance-id	i-10a64379
ami-launch-index	0
public-hostname	ec2-203-0-113-25.compute-1.amazonaws.com
public-ipv4	67.202.51.223
local-hostname	ip-10-251-50-12.ec2.internal
local-ipv4	10.251.50.35

*Instance 2*

Metadata	Value
instance-id	i-10a64380
ami-launch-index	1
public-hostname	ec2-67-202-51-224.compute-1.amazonaws.com
public-ipv4	67.202.51.224
local-hostname	ip-10-251-50-36.ec2.internal
local-ipv4	10.251.50.36

*Instance 3*

Metadata	Value
instance-id	i-10a64381
ami-launch-index	2
public-hostname	ec2-67-202-51-225.compute-1.amazonaws.com
public-ipv4	67.202.51.225
local-hostname	ip-10-251-50-37.ec2.internal
local-ipv4	10.251.50.37

*Instance 4*

Metadata	Value
instance-id	i-10a64382
ami-launch-index	3
public-hostname	ec2-67-202-51-226.compute-1.amazonaws.com
public-ipv4	67.202.51.226
local-hostname	ip-10-251-50-38.ec2.internal
local-ipv4	10.251.50.38

Alice can use the ami-launch-index value to determine which portion of the user data is applicable to a particular instance.

1. She connects to one of the instances, and retrieves the ami-launch-index for that instance to ensure it is one of the replicants:

```
$ GET http://169.254.169.254/latest/meta-data/ami-launch-index
2
```

2. She saves the ami-launch-index as a variable:

```
$ ami_launch_index=`GET http://169.254.169.254/latest/meta-data/ami-launch-index`
```

3. She saves the user data as a variable:

```
$ user_data=`GET http://169.254.169.254/latest/user-data/`
```

4. Finally, Alice runs a Linux cut command to extract the portion of the user data that is applicable to that instance:

```
$ echo $user_data | cut -d"|" -f"$ami_launch_index"
replicate-every=5min
```

## Instance Metadata Categories

The following table lists the categories of instance metadata.

Data	Description	Version Introduced
ami-id	The AMI ID used to launch the instance.	1.0
ami-launch-index	If you started more than one instance at the same time, this value indicates the order in which the instance was launched. The value of the first instance launched is 0.	1.0
ami-manifest-path	The path to the AMI's manifest file in Amazon S3. If you used an EBS-backed AMI to launch the instance, the returned result is unknown.	1.0
ancestor-ami-ids	The AMI IDs of any instances that were rebundled to create this AMI. This value will only exist if the AMI manifest file contained an <code>ancestor-amis</code> key.	2007-10-10
block-device-mapping/ami	The virtual device that contains the root/boot file system.	2007-12-15
block-device-mapping/ebs <i>N</i>	The virtual devices associated with Amazon EBS volumes, if any are present. This value is only available in metadata if it is present at launch time. The <i>N</i> indicates the index of the Amazon EBS volume (such as <code>ebs1</code> or <code>ebs2</code> ).	2007-12-15
block-device-mapping/eph emeral <i>N</i>	The virtual devices associated with ephemeral devices, if any are present. The <i>N</i> indicates the index of the ephemeral volume.	2007-12-15
block-device-mapping/root	The virtual devices or partitions associated with the root devices, or partitions on the virtual device, where the root (/ or C:) file system is associated with the given instance.	2007-12-15
block-device-mapping/swap	The virtual devices associated with swap. Not always present.	2007-12-15
hostname	The private hostname of the instance. In cases where multiple network interfaces are present, this refers to the <code>eth0</code> device (the device for which the device number is 0).	1.0

**Amazon Elastic Compute Cloud User Guide**  
**Instance Metadata Categories**

Data	Description	Version Introduced
iam/info	Returns information about the last time the instance profile was updated, including the instance's LastUpdated date, InstanceProfileArn, and InstanceProfileId.	2012-06-01
iam/security-credentials/ role-name	Where <i>role-name</i> is the name of the IAM role associated with the instance. Returns the temporary security credentials (AccessKeyId, SecretAccessKey, SessionToken, and Expiration) associated with the IAM role.	2012-06-01
instance-action	Notifies the instance that it should reboot in preparation for bundling. Valid values: none   shutdown   bundle-pending.	2008-09-01
instance-id	The ID of this instance.	1.0
instance-type	The type of instance. For more information, see <a href="#">Instance Types (p. 94)</a> .	2007-08-29
kernel-id	The ID of the kernel launched with this instance, if applicable.	2008-02-01
local-hostname	The private DNS hostname of the instance. In cases where multiple network interfaces are present, this refers to the eth0 device (the device for which the device number is 0).	2007-01-19
local-ipv4	The private IP address of the instance. In cases where multiple network interfaces are present, this refers to the eth0 device (the device for which the device number is 0).	1.0
mac	The instance's media access control (MAC) address. In cases where multiple network interfaces are present, this refers to the eth0 device (the device for which the device number is 0).	2011-01-01
network/interfaces/macs/ mac/device-number	The device number associated with that interface. Each interface must have a unique device number. The device number serves as a hint to device naming in the instance; for example, device-number is 2 for the eth2 device.	2011-01-01
network/interfaces/macs/ mac/ipv4-associations/ public-ip	The private IPv4 addresses that are associated with each public-ip address and assigned to that interface.	2011-01-01

**Amazon Elastic Compute Cloud User Guide**  
**Instance Metadata Categories**

Data	Description	Version Introduced
network/interfaces/macs/mac/local-hostname	The interface's local hostname.	2011-01-01
network/interfaces/macs/mac/local-ipv4s	The private IP addresses associated with the interface.	2011-01-01
network/interfaces/macs/mac/mac	The instance's media access control (MAC) address.	2011-01-01
network/interfaces/macs/mac/owner-id	The ID of the owner of the network interface. In multiple-interface environments, an interface can be attached by a third party, such as Elastic Load Balancing. Traffic on an interface is always billed to the interface owner.	2011-01-01
network/interfaces/macs/mac/public-hostname	The interface's public DNS. If the instance is in a VPC, this category is only returned if the <code>enableDnsHostnames</code> attribute is set to <code>true</code> . For more information, see <a href="#">Using DNS with Your VPC</a> .	2011-01-01
network/interfaces/macs/mac/public-ipv4s	The elastic IP addresses associated with the interface. There may be multiple IP addresses on an instance.	2011-01-01
network/interfaces/macs/mac/security-groups	Security groups to which the network interface belongs. Returned only for EC2 instances launched into a VPC.	2011-01-01
network/interfaces/macs/mac/security-group-ids	IDs of the security groups to which the network interface belongs. Returned only for EC2 instances launched into a VPC. For more information on security groups in the EC2-VPC platform, see <a href="#">Security Groups for Your VPC</a> .	2011-01-01
network/interfaces/macs/mac/subnet-id	The ID of the subnet in which the interface resides. Returned only for EC2 instances launched into a VPC.	2011-01-01
network/interfaces/macs/mac/subnet-ipv4-cidr-block	The CIDR block of the subnet in which the interface resides. Returned only for EC2 instances launched into a VPC.	2011-01-01
network/interfaces/macs/mac/vpc-id	The ID of the VPC in which the interface resides. Returned only for EC2 instances launched into a VPC.	2011-01-01
network/interfaces/macs/mac/vpc-ipv4-cidr-block	The CIDR block of the VPC in which the interface resides. Returned only for EC2 instances launched into a VPC.	2011-01-01
placement/availability-zone	The Availability Zone in which the instance launched.	2008-02-01

**Amazon Elastic Compute Cloud User Guide**  
**Instance Metadata Categories**

Data	Description	Version Introduced
product-codes	Product codes associated with the instance, if any.	2007-03-01
public-hostname	The instance's public DNS. If the instance is in a VPC, this category is only returned if the <code>enableDnsHostnames</code> attribute is set to <code>true</code> . For more information, see <a href="#">Using DNS with Your VPC</a> .	2007-01-19
public-ipv4	The public IP address. If an elastic IP address is associated with the instance, the value returned is the elastic IP address.	2007-01-19
public-keys/0/openssh-key	Public key. Only available if supplied at instance launch time.	1.0
ramdisk-id	The ID of the RAM disk specified at launch time, if applicable.	2007-10-10
reservation-id	ID of the reservation.	1.0
security-groups	The names of the security groups applied to the instance.  <b>Note</b> Only EC2 instances launched into a VPC can change security groups after launch. These changes will be reflected here and in <code>network/interfaces/network/interfaces/security-groups</code> .	1.0

## Dynamic Data Categories

The following table lists the categories of dynamic data.

Data	Description	Version introduced
fws/instance-monitoring	Value showing whether the customer has enabled detailed one-minute monitoring in CloudWatch. Valid values: <code>enabled</code>   <code>disabled</code>	2009-04-04
instance-identity/document	JSON containing instance attributes, such as instance-id, private IP address, etc.	2009-04-04
instance-identity/pkcs7	Used to verify the document's authenticity and content against the signature.	2009-04-04
instance-identity/signature	Data that can be used by other parties to verify its origin and authenticity.	2009-04-04

# Importing and Exporting Instances

You can import a virtual machine (VM) from a Citrix Xen, Microsoft Hyper-V, or VMware vSphere virtualization platform and then launch it in Amazon EC2. Later, you can export that Amazon EC2 instance back to Citrix Xen, Microsoft Hyper-V, or VMware vSphere.

## Topics

- [Importing EC2 Instances \(p. 301\)](#)
- [Exporting EC2 Instances \(p. 332\)](#)

## Importing EC2 Instances

### Topics

- [Components in Your VM Environment \(p. 302\)](#)
- [Before You Get Started \(p. 302\)](#)
- [Using the Amazon EC2 VM Import Connector to Import Your Virtual Machine to Amazon EC2 \(p. 303\)](#)
- [Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2 \(p. 318\)](#)
- [Troubleshooting Instance Importation \(p. 330\)](#)

There are two ways you can launch an instance in the Amazon Elastic Compute Cloud (Amazon EC2). You can launch an instance from an AMI that you created or selected from a catalog. Or, you can launch an instance from a virtual machine (VM) that you imported from a Citrix Xen, Microsoft Hyper-V, VMware Workstation, or VMware vSphere virtualization environment. This section covers using VMs from Citrix, Microsoft, or VMware to launch instances.

To use your virtual machine as an instance in Amazon EC2, you must first export it from the virtualization environment using its tools. Then you import it to Amazon EC2 using the Amazon EC2 command line or API tools.

If you are importing a VMware vSphere VM, you can also use the Amazon EC2 VM Import Connector for VMware (Connector), a plug-in that integrates with VMware vSphere Client, to perform the task.

Whether you use the command line tools, the API, or the Connector, you will follow the same general process for importing VMs or volumes to Amazon EC2. You need to complete these tasks, which are all discussed in this section:

1. Prepare the virtual machine for import to Amazon EC2. For more information, go to [Before You Get Started \(p. 302\)](#)
2. Export the virtual machine from the virtualization environment.
  - [Exporting from Citrix \(p. 319\)](#)
  - [Exporting from Microsoft Hyper-V \(p. 322\)](#)
  - [Exporting from VMware \(p. 323\)](#)
3. Import the virtual machine to Amazon EC2.  
For information about using the command line tools to import your VM, see [Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2 \(p. 318\)](#). For information about using the Connector, see [Using the Amazon EC2 VM Import Connector to Import Your Virtual Machine to Amazon EC2 \(p. 303\)](#).
4. Upload the instance to Amazon EC2.
5. Launch the instance in Amazon EC2.



## Components in Your VM Environment

The table in this section describes the typical components in your VM environment.

Component	Description	Product Name
Virtualization product	Virtualization service for managing virtual computing infrastructure	Citrix Xen Microsoft Hyper-V VMware Workstation VMware vSphere (vSphere)
Client	The software you need on your computer to access and manage your virtualization environment	Citrix Xen Center Microsoft Hyper-V Manager VMware Workstation VMware vSphere Client
Server	The management platform for the virtualization environment	Citrix XenServer Microsoft Hyper-V VMware vCenter Server
Amazon EC2 VM Import Connector for VMware (Connector)	The virtual appliance, a plug-in to the management platform of the virtualization Server, that enables the import of virtual machines into Amazon EC2 using the Client interface	Citrix Xen (not applicable) Microsoft Hyper-V (not applicable) VMware vCenter—Amazon EC2 VM Import Connector for VMware vCenter (Connector)

## Before You Get Started

This section discusses the things you need to know and what you must have before you begin the process of importing your virtual machine.

**Operating Systems**—The following operating systems can be imported to Amazon EC2:

- Microsoft Windows Server 2003 (Standard, Datacenter, Enterprise).
- Microsoft Windows Server 2003 R2 (Standard, Datacenter, Enterprise).
- Microsoft Windows Server 2008 (Standard, Datacenter, Enterprise).
- Microsoft Windows Server 2008 R2 (Standard, Datacenter, Enterprise).

**Image Formats**—We support import of the following image formats for importing both volumes and instances to Amazon Web Services:

- RAW format for importing volumes and instances.
- Virtual Hard Disk (VHD) image formats, which are compatible with Microsoft Hyper-V and Citrix Xen virtualization products.
- ESX Virtual Machine Disk (VMDK) image formats, which are compatible with VMware ESX and VMware vSphere virtualization products.

**Known Limitations**—The importing of instances and volumes is subject to the following limitations:

- You can have up to five conversions and tasks in progress at the same time per Region.
- Typically, you import a compressed version of a disk image; the expanded image cannot exceed 1TB.

- Tasks must complete within 7 days of the start date.
- Importing virtual machines with more than one virtual disk is not supported. We suggest that you import the VM with only the boot volume, and import any additional disks using `ImportVolume` (`ec2-import-volume`) in the command line. After the `ImportInstance` task is complete, use `AttachVolume` (`ec2-attach-volume`) to associate the additional volumes with your instance.
- Multiple network interfaces are not currently supported. When converted and imported, your instance will have a single virtual NIC using DHCP for address assignment.
- For vCenter 4.0 and vSphere 4.0 users, remove any attached CD-ROM images or ISOs from the virtual machine.
- Internet Protocol version 6 (IPv6) IP addresses are not supported.

## Preparing Your Virtual Machine

Use the following guidelines to configure your virtual machine before exporting it from the virtualization environment.

### Important

If you are importing a virtual machine from Citrix Xen, you must uninstall the Citrix Tools for Virtual Machines from the VM. If you don't uninstall the tools, your import will fail. For more information, see [Exporting from Citrix \(p. 319\)](#).

- Enable Remote Desktop (RDP) for remote access.
- Make sure your host firewall (Windows firewall), if configured, allows access to RDP. Otherwise, you will not be able to access your instance after the conversion is complete.
- Make sure all user accounts use secure passwords. This includes the administrator account.

### Note

All accounts must have passwords. Otherwise, the import might fail.

- Disable any antivirus or intrusion detection software on your virtual machine. These services can be re-enabled after the import process is complete.
- Disconnect any CD-ROM drives (virtual or physical).
- Do not Sysprep your virtual machine images. We recommend that you import the image and then use the Amazon EC2 Config service to Sysprep it.
- Set your network to DHCP. If you are using a private IP address, be sure to use a non-reserved private IP address in your VPC subnet. Amazon Virtual Private Cloud (Amazon VPC) reserves the first four private IP address in a VPC subnet.
- Shut down your virtual machine before exporting it.
- Make sure you've installed .NET Framework 3.5, as required by [Amazon Windows EC2Config Service](#).

## Using the Amazon EC2 VM Import Connector to Import Your Virtual Machine to Amazon EC2

### Topics

- [Before You Install the Connector \(p. 304\)](#)
- [Installing the Connector for VMware vCenter \(p. 305\)](#)
- [Configuring the Connector for VMware vCenter \(p. 309\)](#)
- [Using the Connector for VMware vCenter \(p. 314\)](#)
- [Get Diagnostic Information from the Connector for VMware vCenter \(p. 316\)](#)
- [Uninstalling the Connector for VMware vCenter \(p. 317\)](#)

You can use the Amazon EC2 VM Import Connector virtual appliance (vApp), a plug-in for VMware vCenter, to import virtual machines from your VMware vSphere infrastructure to Amazon EC2. The Connector is a virtual appliance that works with VMware vCenter Server only. It provides an easy-to-use interface, enhancing your existing management tools to work with the Amazon EC2 VM Import Connector.

**Note**

You cannot use the Connector to import Citrix Xen or Microsoft Hyper-V virtual machines to Amazon EC2. Instead, use the command line tools to import your Citrix and Hyper-V virtual machines to Amazon EC2. You can also choose to use the command line tools to import your VMware VMs. For more information, see [Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2 \(p. 318\)](#).

By comparing the following procedures, you can see that using the Connector simplifies the process of importing your VMware VMs.

**To import VMware VMs using the Connector**

1. Import the VM to Amazon EC2
2. Launch the instance

**To import VMware VMs using Amazon EC2**

1. Export the VM from the virtualization environment
2. Import your VM to Amazon EC2
3. Upload the instance to Amazon EC2
4. Launch the instance in Amazon EC2

**Before You Install the Connector**

To use the VM Import Connector, you first need to install the Connector virtual appliance. Before you install, read through the general prerequisites listed in the [Before You Get Started \(p. 302\)](#) section and make sure that your virtualization environment meets the following requirements:

**VMware infrastructure**

- vSphere 4.0, 4.1, 5.0, or 5.1
- vCenter 4.0, 4.1, 5.0, or 5.1

**Amazon EC2 VM Import Connector vApp**

- 256MB RAM
- Minimum 250GB of disk space

**Note**

Although the Connector virtual appliance is small, it temporarily stores the VM images that you import to Amazon EC2. The data store must be large enough to accommodate these images, plus the Connector. We recommend a data store size of 250GB or larger.

To estimate the disk space you need, multiply the maximum number of parallel import tasks you want to run with the Connector by the average size you expect your VMs to be, then add about 10GB for the virtual appliance. For example, if you project that you will run a maximum of 5 imported VMs averaging 75GB in size, then you'll need about 385GB of disk space.

**Internet access**

- Outbound Internet access, either direct or via a proxy, from the Connector appliance on TCP port 443 (SSL) to Amazon EC2 and Amazon S3.

**Note**

If you are adding a firewall rule to allow this access and want to further restrict this access to Amazon EC2 and Amazon S3, you can add the hosts at the published endpoints as the destinations on TCP port 443. For more information about the current endpoints, see [Regions and Endpoints](#).

- DHCP server
- A static IP address or an IP reservation via DHCP for the virtual appliance

**Note**

You must set up a static DHCP lease or configure a static IP before you begin the configuration process.

### Connector local network access

- Inbound TCP 443 (SSL) from vCenter Server and vSphere Client
- Inbound TCP 80 (HTTP) from the LAN for Connector Web Console

### Administrative rights

- To VMware vSphere and VMware vCenter for installation. For information on how to allow a user without administrative rights to use the Connector to import VMs to Amazon EC2, see [To grant permission to non-administrative users to import VMs to Amazon EC2 \(p. 311\)](#).

### AWS access credentials

- The Connector stores AWS credentials for each VMware vCenter user. In this way, multiple users with separate AWS credentials can use the same Connector. Alternatively, you can use the same AWS account and credentials with more than one user. Each VMware vSphere user will need to fill out the information in the **Enter AWS Credentials** dialog box the first time he or she uses the Connector.
- Your AWS account must be subscribed to Amazon EC2 before you start the import process.

**Note**

The Connector virtual appliance stores your AWS credentials. To protect these credentials from unauthorized use, grant access to the virtual appliance's console and configuration only to administrators.

## Installing the Connector for VMware vCenter


After you have confirmed that you have all the prerequisites and your virtualization environment meets the minimum requirements, you are now ready to install the Connector virtual appliance (vApp) for VMware vCenter.

The Connector virtual appliance is an Open Virtualization Format (OVF) package that is distributed in an Open Virtualization Application (OVA) file. Installing the Connector involves downloading the OVA file and deploying the OVF template.

### To Install the Connector for the VMware vCenter


1. Download the OVA package for the Connector virtual appliance from [Amazon Web Services Developer Tools](#) and save it to your Downloads folder.
2. Start the vSphere Client and connect to your vCenter Server.

- Using the vSphere Client, deploy the OVF template contained in the OVA file that you downloaded. On the **File** menu, select **Deploy OVF Template** and point to the location where you downloaded the OVA package.




A series of **Deploy OVF Template** screens walks you through the deployment process.


- Confirm that the **Deploy OVF Template** screen is displaying correct information about the VM Import Connector, and click **Next**.



- Specify the name and location of the Connector or accept the default, then click **Next**.



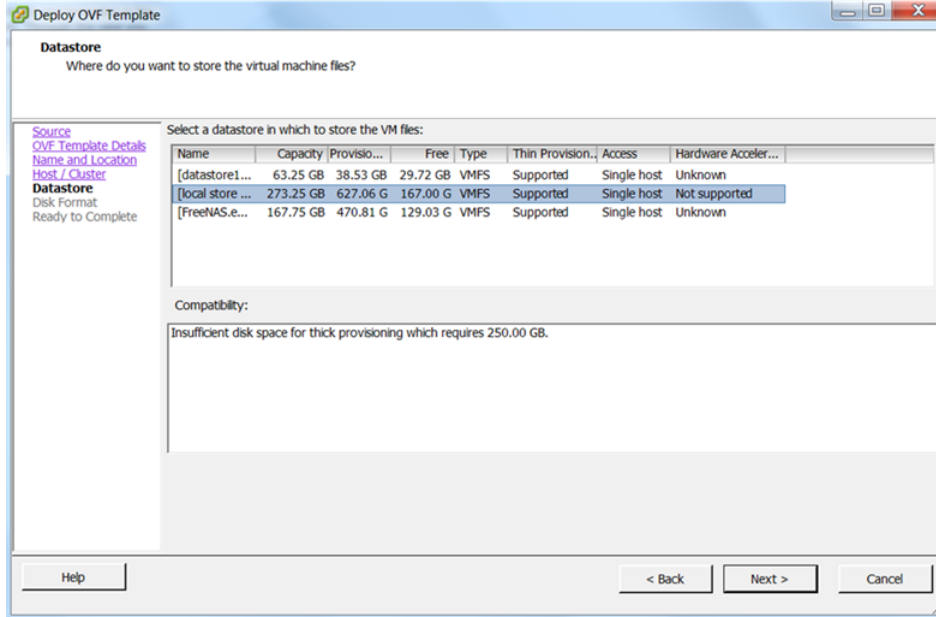
6. Select the host or cluster in which you want the Connector to run, then click **Next**.



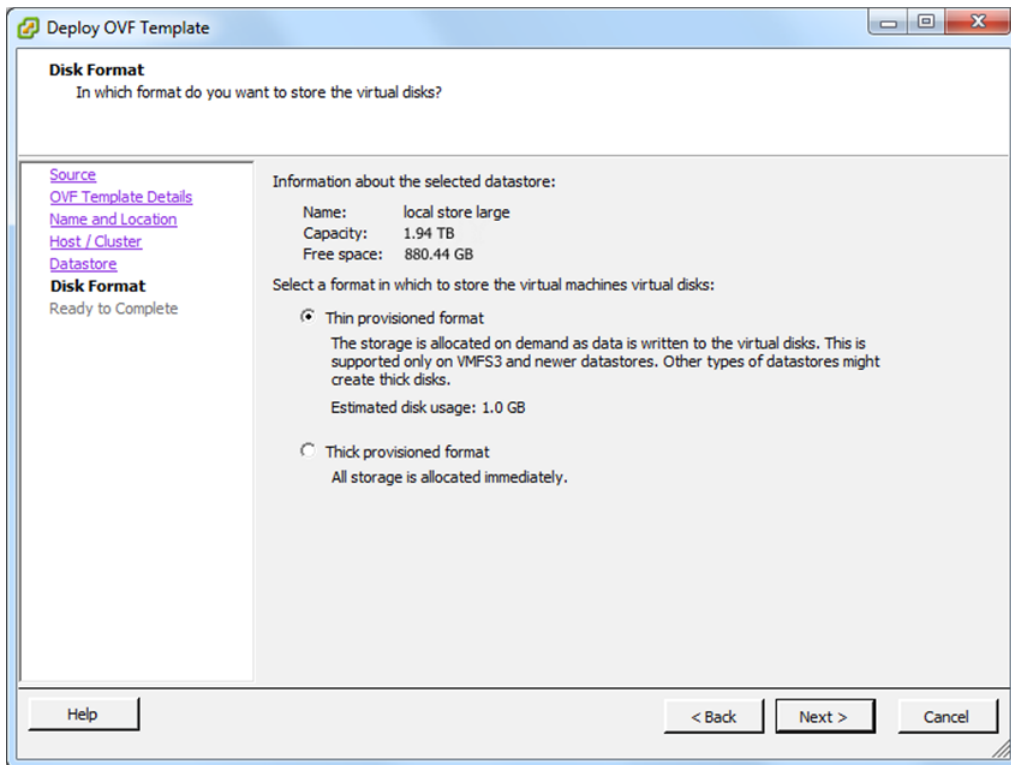
7. Select the data store where the Connector will be stored, then click **Next**.

**Note**

Although the Connector virtual appliance is small, it temporarily stores the images of the virtual machines that import to Amazon EC2. Therefore, the data store must be large enough to accommodate these images, as well as the Connector. We recommend a data store size of 250GB or larger.



- 8. Select **Thin provisioned** format, then click **Next**.



9. Confirm the details you selected and click **Finish**.

The Connector virtual appliance will now install.

## Configuring the Connector for VMware vCenter

After you install the EC2 VM Import Connector, you must obtain its IP address and password and register it with the vCenter Server. To obtain the Connector's IP address and password, you first start the Connector appliance in vCenter, then go to the vCenter **Console** tab. The tab displays the IP address and password when the Connector is running. In a web browser, use this information to log in to the Connector and register it with the vCenter Server.

If the VMware vSphere Client was running when you installed the Connector virtual application, close and restart the vSphere Client, then follow these procedures.

### To start the Connector for VMware vCenter

1. On the vSphere Client, right-click the Connector you just installed, select **Power** and then **Power On**.
2. To open the console, right-click the Connector you just installed and click **Open Console**.

When the Connector is running, you will find its IP address and password displayed in the vCenter **Console**. Your Connector information will be similar to the following example:

```
Amazon EC2 VM Import Connector for VMware vCenter
Management Website -> 192.0.2.99
Management Password -> tURILx3K
login: █
```

### To register the Connector with the VMware vCenter

1. Open a web browser and, in the address bar, type the Connector IP address that you obtained from the **Console** in the vCenter, and log in with the password.




2. When you are logged in, the browser displays Connector status information. Note that the Connector is not yet registered with the vCenter Server. Confirm that everything else in the **Connector Status** list has a status of **OK**.





3. If the version of the Connector that you are registering is not an upgrade, click **Register Now**. If you are upgrading the Connector, you must take the next two steps before registering.
  - Confirm that all import tasks are complete.
  - Deregister the old Connector from vCenter. To do this, go to the **vCenter** tab and click **Deregister**.



The **vCenter Connector Registration** page appears.

4. Provide the IP address of the vCenter with which you want to register the Connector. The user name and password you use must have administrative rights. Click **Register Connector with vCenter**.

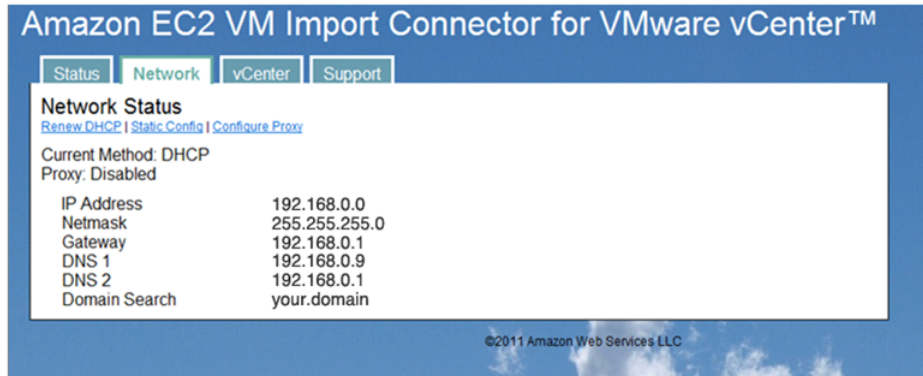
#### Note

If an error occurs, check that the VMware vCenter IP address or name that you provided is correct. Also, confirm that the Connector virtual appliance has network access to TCP port 443 on your vCenter Server. If the user name or password you provided is incorrect, you will see a description of this error.

### To configure a proxy

If you need to configure a proxy to allow the Connector to reach the Internet, you can do it using the Connector's web interface.

1. Go to the **Network** tab, click **Configure Proxy**.



The **Network Configuration:Proxy** page appears.


2. Provide the information required and click **Configure Proxy**.



After using the web browser to register the Connector with the vCenter Server, as an administrator you can import virtual machines to Amazon EC2. If you're going to import VMs using only an administrator account, skip the following section and go to [Using the Connector for VMware vCenter \(p. 314\)](#). If you want non-administrative users to import VMs to Amazon EC2, you must grant them permission using the vCenter.

### To grant permission to non-administrative users to import VMs to Amazon EC2

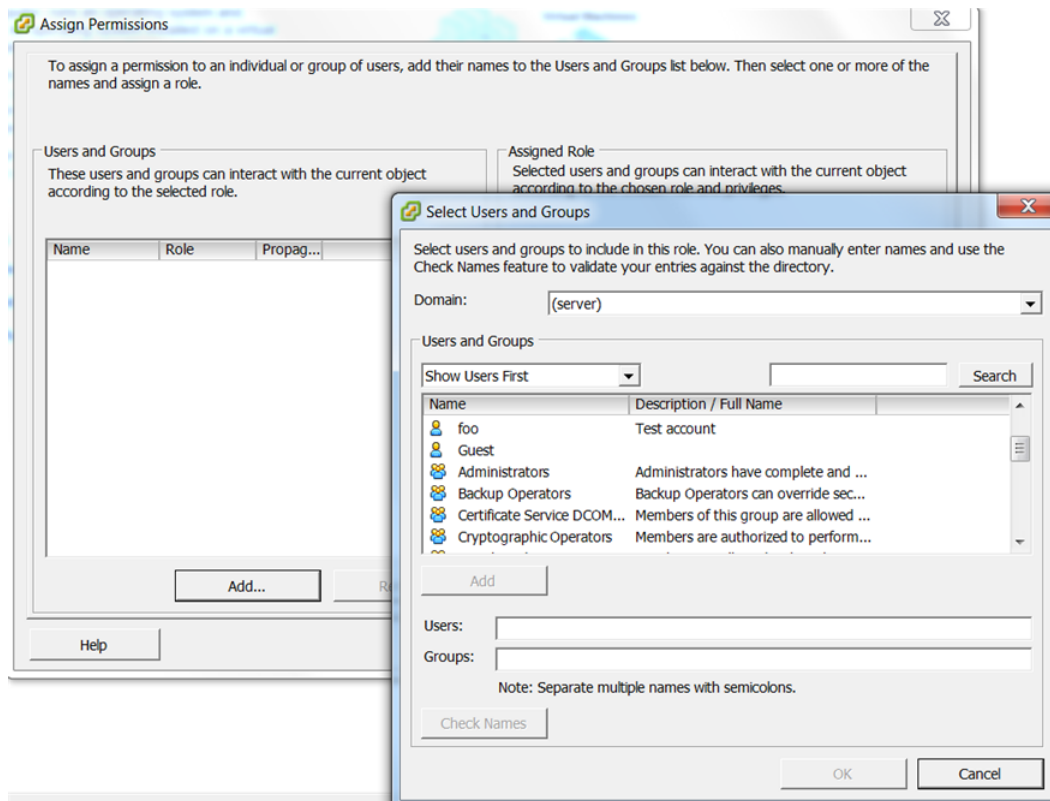
1. Log in to vCenter as Administrator and from **Home**, navigate to **Roles**, and click **Add Role**.



2. In the **Add New Role** dialog box, type the name for the new role and specify the following permissions.
  - Under **Global**, select **Cancel task**.
  - Under **Tasks**, select **Create task** and **Update task**.
  - Under **vApp**, select **Export** and **View OVF Environment**.

Click **OK**.


3. To grant the new role permission to vCenter users or groups who will be importing virtual machines to Amazon EC2, right-click your specific vCenter in the tree-view pane and select **Add permission**. The **Assign Permissions** dialog box opens.
4. In the **Users and Groups** box on the left, select the users or groups you want to add to the new role you created.



If you don't have users defined yet, click **Add**. The **Select Users and Groups** dialog box opens.

5. Select and add the users you want to add to the new role. When you have identified all the users for the role, click **OK**.
6. In the **Assigned Role** box on the right of the **Assign Permissions** dialog box, select the role that you previously created.
7. Clear the check box for **Propagate to Child Objects** and click **OK**.

Repeat the same process to assign the role of virtual machine power user to all users and groups that you want to allow to import VMs to Amazon EC2. You can do this at the VM object level or at a higher level in the hierarchy. If you assign the role at a higher level, you must select the check box for **Propagate to Child Objects**.



## Using the Connector for VMware vCenter

This section shows you how to use the Connector to import a virtual machine to Amazon EC2 for the first time using an account with administrative rights.

Confirm that you have prepared the virtual machine according to the guidelines in [Preparing Your Virtual Machine \(p. 303\)](#).

### Important

If you don't enable RDP and disable the Windows-based firewall, your VM will import successfully to Amazon EC2, but you will not be able to log in.

These requirements must be satisfied:

- The virtual machine must be turned off.
- The virtual machine must only use a single virtual hard drive (multiple partitions are OK). Any additional disks must be detached or import will fail.
- The virtual hard drive cannot be larger than one terabyte (1TB).
- The Connector virtual appliance must have sufficient free hard drive space to temporarily store the compressed VMDK image while it is being imported to Amazon EC2.

### To use the Connector to import a VM for the first time using an account with administrative rights

1. Log in to VMware vCenter using the VMware vSphere Client. If you had a session open while you were installing the Connector, notice that the **Import to EC2** tab becomes visible.

### Note

The first time you log in, you will see an SSL certificate warning. This warning indicates that the SSL certificate being used by the Connector cannot be verified by an external source. This is expected behavior. Your session will continue to use SSL for encryption. Check the **Install this certificate and do not display a security warning** option at the bottom of the screen and click **Ignore**.

You are now logged in to the vCenter Server.

2. On the left pane, navigate the tree view to the virtual machine you want to import. Select it.
3. On the right pane, select the **Import to EC2** tab.

### Note


You might see another security warning about SSL certificates. Click **Yes** to continue.

The Connector initializes.

4. In the **Enter AWS Credentials** dialog box, provide your **Access Key ID** and **Secret Access Key**.

The main Connector page appears when your AWS credentials are verified.

5. Back in vCenter, select the virtual machine you want to import and go to the **Import to EC2** tab.



6. In the **Instance Parameters** dialog box, specify the values for the following options, then click the **Import to EC2** button.

- **Region**—(Required) The AWS Region in which you want to launch your imported instance. Select one of the options from the drop-down list.
- **Availability Zone**—(Optional) Your choice of Availability Zone in the Region that you have selected. If not specified, Amazon will select one for you.
- **Operating System**—(Required) The operating system that your virtual machine is using. Selecting an incorrect value can cause the import task to fail.

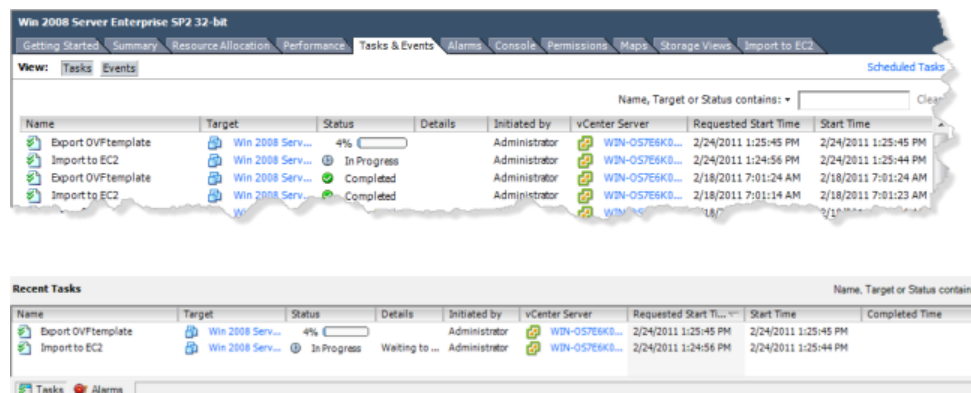
For virtual machines using Windows Server 2008 R2, always select **Microsoft Windows Server 2008 (64-bit)**. If the virtual machine you're importing runs on Windows Server 2008 SP2, determine first whether the operating system is a 32-bit or 64-bit **System Type** then select the **Operating System** accordingly. For more information about 32-bit and 64-bit Windows, go to [32-bit and 64-bit Windows: frequently asked questions](#). For more information about how to determine System Type for Windows Server 2003, go to [Microsoft Support](#).

- **Instance Type**—(Required) The instance type to use when your virtual machine becomes an instance in EC2. Instance types allow you to choose the CPU and memory requirements. For more information on instance types and pricing, go to [Amazon EC2 Pricing](#).
- **Launch in VPC**—(Optional) Check this box if you want your virtual machine to become an instance within Amazon Virtual Private Cloud (Amazon VPC). For information, go to [Amazon Virtual Private Cloud \(Amazon VPC\)](#).

- **Subnet**—(Required only if launching in VPC) Select the subnet that you want your instance placed within in a VPC.
- **Private IP address**—(Optional, applies only if launching in VPC) Specify the private IP address of your instance within VPC.
- **Security Group**—(Required) Select the security group to use with your instance. Defaults to the default security group.

The values you specified are listed in the **Confirm Import Options** box. Check the information and click **Import**.

7. Monitor the progress of the import task in the **Tasks & Events** or **Recent Tasks** tab of the VMware vSphere Client.



- **Export OVF template**—Creates a stream-optimized VMDK image. This process consolidates your virtual machine to a single image. In addition, stream-optimized VMDKs are compressed and are well-suited for transfer over a WAN connection. The stream-optimized VMDK will be temporarily stored on your Connector virtual appliance.
- **Import to EC2**—Transfers the stream-optimized VMDK that was created in the first task to Amazon EC2, and converts your virtual machine to an Amazon EC2 instance.

The **Import to EC2** task can take up to a few hours to complete. In addition, you might notice that the task progress will pause for up to 10 minutes at times. This is expected behavior.

### Important

Keep your session in VMware vCenter open until all tasks complete. If you quit the vSphere Client or log off of your vCenter, the import task will not complete successfully, and the progress indicator will not be updated. If this occurs, you can use the command line tools to check the status of your import task. For more information, see [Checking on the Status of Your Import in Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2 \(p. 318\)](#). When you have verified that the task is complete, you can safely cancel the import task by right-clicking the task in the vSphere Client and selecting **Cancel**.

Although you can see your instance in the AWS Management Console when the import process begins, do not launch your instance until the import process completes. For information about launching instances, see [Launching an Instance \(p. 266\)](#).


## Get Diagnostic Information from the Connector for VMware vCenter

If you are having trouble with the Connector for VMware vCenter, you can download diagnostic information to help you determine a cause.



### To get diagnostic information from the Connector for VMware vCenter

1. Open a web browser and, in the address bar, type the Connector IP address that you obtained from the **Console** in the vCenter, and log in with the password.



2. When you are logged in, the browser displays Connector status information. Note that the Connector is not yet registered with the vCenter Server. Confirm that everything else in the **Connector Status** list has a status of **OK**.



3. Click the **Support** tab, then click **Download Debugging File**.

The Connector will generate the connector-debug.tar.gz file that you can use to troubleshoot the Connector.

### Uninstalling the Connector for VMware vCenter

If you no longer want to use the Connector for VMware vCenter and you want to uninstall the virtual appliance, you will follow a two-part process:

- Deregister the Connector from the vCenter Server.
- Shut down the virtual appliance.

#### To uninstall the Connector from the VMware vCenter

1. Open a web connection to the Connector's IP address and log in.
2. Click the **vCenter** tab, then click **Deregister**.



3. In the **vCenter Connector Deregistration** page, enter the vCenter IP information and user name and password, then click **Deregister Connector with vCenter**.
4. When the Connector is no longer registered with the vCenter Server, shut down the virtual appliance and remove it from vCenter.

## Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2

### Topics

- [Exporting Your Virtual Machines from Their Virtual Environment \(p. 319\)](#)
- [Importing Your Virtual Machine into Amazon EC2 \(p. 324\)](#)

In this section, you'll learn how to use the Amazon EC2 command line tools to import your Citrix, Microsoft Hyper-V, or VMware virtual machine to Amazon EC2. If you haven't already installed the Amazon EC2 command line tools, see [Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX \(p. 541\)](#).

Importing VMs into Amazon EC2 is a two-step process. First, you export your virtual machine from the virtualization environment. Next, you create an import task and upload your virtual machine into Amazon EC2.

Use the following commands when you perform import tasks using the Amazon EC2 command line tools:

Command	Description
<code>ec2-import-instance</code>	Creates a new import instance task using metadata from the specified disk image and imports the instance to Amazon EC2.
<code>ec2-import-volume</code>	Creates a new import volume task using metadata from the specified disk image and imports the volume to Amazon EC2.
<code>ec2-resume-import</code>	Resumes the upload of a disk image associated with an import instance or import volume task ID.
<code>ec2-describe-conversion-tasks</code>	Lists and describes your conversion tasks.
<code>ec2-cancel-conversion-task</code>	Cancels the active conversion task. The task can be the import of an instance or volume.
<code>ec2-delete-disk-image</code>	Deletes a partially or fully uploaded disk image for conversion from Amazon S3.

For information about these commands and other EC2 commands, go to the [Amazon Elastic Compute Cloud Command Line Reference](#).

### Note

You can use a graphical user interface provided through the Amazon EC2 VM Import Connector (Connector) to import your VMware virtual machines to Amazon EC2. For more information, see [Using the Amazon EC2 VM Import Connector to Import Your Virtual Machine to Amazon EC2 \(p. 303\)](#). You cannot use the Connector to import Citrix or Microsoft Hyper-V virtual machines.

## Exporting Your Virtual Machines from Their Virtual Environment

The process of exporting virtual machines depends on the source virtualization environment. In this section, we will show you the basic steps to export virtual machines from Citrix, Microsoft Hyper-V, and VMware virtualization products. For in-depth information, consult the documentation for these products.

- [Exporting from Citrix \(p. 319\)](#)
- [Exporting from Microsoft Hyper-V \(p. 322\)](#)
- [Exporting from VMware \(p. 323\)](#)

Before starting the export process, prepare the Windows Server environment of the virtual machine you are exporting so that:

- Remote desktop is enabled.
- Windows firewall allows public RDP traffic.
- Autologon is disabled.
- There are no pending Microsoft updates and the computer is not set to install software when it reboots.

### Exporting from Citrix


Before you export your virtual machines from Citrix XenCenter, you must perform the following tasks:

- Remove the **Citrix Tools for Virtual Machines** from the VM.
- Use the Citrix XenCenter console to remove the tools.


If you have multiple virtual disks that you want to export from Citrix Xen, we recommend that you export the disks one at a time. The export of multiple virtual disks at the same time results in a list of randomly named VHD files.

### To remove Citrix tools for virtual machines

1. In **XenCenter**, select the virtual machine you want to export, and click the **Console** tab, which shows the Windows desktop of the virtual machine.



- Using the **Console**, access the Control Panel of the virtual machine's Windows operating system and uninstall the **Citrix Tools for Virtual Machines**.




- After the tools are removed, reboot the virtual machine when prompted, log in again and then shut down using Windows.  
You can now proceed to export the VM.

## To export an image from Citrix


1. In the Citrix XenCenter, select the virtual machine you want to export, and then go to the **Tools** menu, click **Virtual Appliance Tools**, and then **Export Appliance**.

### Note

Do not export the image by right-clicking a stopped instance and selecting "Export to File."



2. When the **XenServer OVF Appliance Export Wizard** starts, specify the destination of the VM files, click **Next**, accept the defaults, and proceed through the screens until you click **Finish**.




3. When the export completes, you can proceed and import the VM files to Amazon EC2 by following the steps in [Importing Your Virtual Machine into Amazon EC2 \(p. 324\)](#) and specifying `vhd` as the file format.

## Exporting from Microsoft Hyper-V


To export Hyper-V virtual disks from Microsoft, you use the Hyper-V Manager.

### To export a Hyper-V image from Microsoft

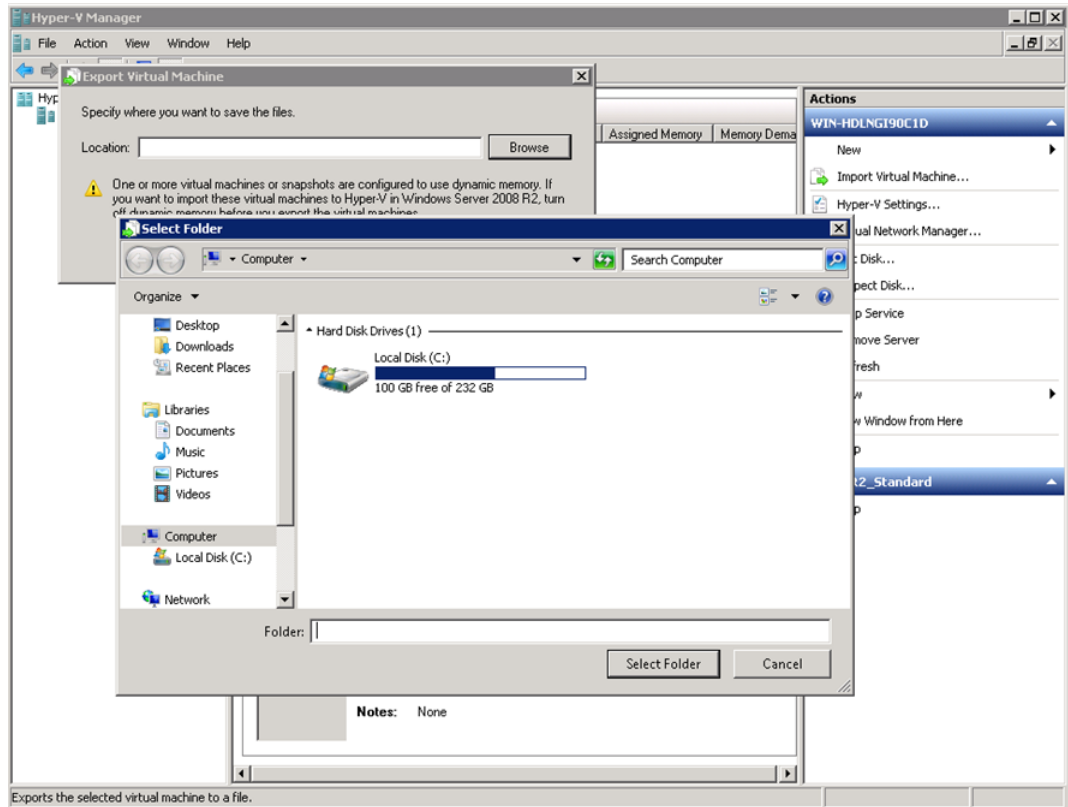
1. In the **Hyper-V Manager**, shut down the virtual machine you want to export.



2. In the **Actions** pane for the virtual machine, select **Export**.



3. In the **Export Virtual Machine** dialog box, for **Location**, click **Browse**, navigate to a destination location that has plenty of space, and click **Export**.



4. Track the export progress through the **Status** of your VM in the Hyper-V Manager. Wait for the export to complete.


When the export completes, you can proceed and import the VM files to Amazon EC2 by following the steps in [Importing Your Virtual Machine into Amazon EC2 \(p. 324\)](#) and specifying VHD as the file format. The VHD file will be located in the folder you specified in the **Export Virtual Machine** dialog box.

## Exporting from VMware


Before you can import a VMware vSphere VM or volume into Amazon EC2, you must export the VMDK disk image file. The following procedure shows you how to use the VMware vSphere Client to export a VM (and VMDK file). For more detailed information, consult your VMware documentation.

### To export a disk image from VMware


1. In the VMware vSphere Client, select the virtual machine to export.
2. Shut down the machine.



- From the **File** menu, select **Export**, and then **Export OVF Template**.



- In the **Export OVF Template** dialog box, enter a name for the disk image file and a directory to save it in.



You will see a box displaying the progress of the export. Because of the size and number of files associated with the VM, and the network connection, the export might take some time. If the connection times out, the export process will fail. If this happens, restart the export. When the export completes, vSphere saves the disk image file in the directory you specified. Use the name of the VMDK file as an argument in `ec2-import-instance` to import a virtual machine into Amazon EC2, or in `ec2-import-volume` to import a volume into Amazon Elastic Block Store (Amazon EBS).

## Importing Your Virtual Machine into Amazon EC2

After exporting your virtual machine from the third-party virtualization environment, you can import it into Amazon EC2. The import process is the same regardless of the origin of the virtual machine.

Here are some important things to know about your VM instance, as well as some security and storage recommendations:

- Amazon EC2 automatically assigns a DHCP IP address to your instance. The DNS name and IP address are available via the `ec2-describe-instances` command when the instance starts running.
- Your instance has only one Ethernet network interface.

- To specify an Amazon Virtual Private Cloud (Amazon VPC) subnet to use when you create the conversion task, use the `--subnet subnet_id` option with the `ec2-import-instance` command. Otherwise your instance will use a public IP address. We recommend you use a restrictive security group to control access to your instance. For more information about the `--subnet subnet_id` option, see [ec2-import-instance](#).
- We recommend that your instance contain strong passwords for all user accounts.
- We recommend that you install the [Amazon Windows EC2Config Service](#) after you import your virtual machine into Amazon EC2.

## To import a virtual machine

You can import a virtual machine into Amazon EC2. If the import of the virtual machine is interrupted, you can use the `ec2-resume-import` command to resume the import from where it stopped. For more information, see [Resuming an Upload \(p. 328\)](#).

- Use `ec2-import-instance` to create a new import instance task.  
The syntax of the command is:

```
ec2-import-instance DISK_IMAGE_FILENAME -t INSTANCETYPE -f FORMAT -a ARCHITECTURE-SYSTEM -b S3_BUCKET_NAME -o OWNER -w SECRETKEY
```

The following command creates an import instance task that imports a Windows Server 2008 SP2 (32-bit) VM.

### Note

The example uses the VMDK format. You can also use VHD or RAW.

```
ec2-import-instance ./WinSvr8-2-32-disk1.vmdk -f VMDK -t ml.small -a i386 -b myawsbucket -o AKIAIOSFODNN7EXAMPLE -w wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
```

This request uses the VMDK file, `winSvr8-2-32-disk1.vmdk`, to create the import task. The output is similar to the following example.

```
Requesting volume size: 25 GB
Disk image format: Stream-optimized VMDK
Converted volume size: 26843545600 bytes (25.00 GiB)
Requested EBS volume size: 26843545600 bytes (25.00 GiB)
TaskType          IMPORTINSTANCE TaskId import-i-fhbx6hua ExpirationTime
 2011-09-09T15:03:38+00:00 Status active StatusMessage Pending
InstanceID        i-6ced060c
DISKIMAGE         DiskImageFormat VMDK DiskImageSize 5070303744
VolumeSize        25 AvailabilityZone us-east-1c Approximate
BytesConverted    0 Status active StatusMessage Pending
Creating new manifest at testImport/9cba4345-b73e-4469-8106-
2756a9f5a077/Win_2008_R1_EE_64.vmdkmanifest.xml
Uploading the manifest file
Uploading 5070303744 bytes across 484 parts
0% |-----| 100%
   |=====|
Done
```



## Checking on the Status of Your Import

The `ec2-describe-conversion-tasks` command returns the status of an import. Status values include:

- **active**—Your instance or volume is still importing.
- **cancelling**—Your instance or volume is still being canceled.
- **cancelled**—Your instance or volume is canceled.
- **completed**—Your instance or volume is ready to use.

The imported instance is in the stopped state. You use `ec2-start-instance` to start it. For more information, go to [ec2-start-instances](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

## To check the status of your import

- Use `ec2-describe-conversion-task` to return the status of the task. The syntax of the command is:

```
ec2-describe-conversion-tasks TASKID
```

The following example enables you to see the status of your import instance task.

```
ec2-describe-conversion-tasks import-i-ffvko9js
```

The following response shows that the `IMPORTINSTANCE` status is active, and 73747456 bytes out of 893968896 have been converted.

```
TaskType      IMPORTINSTANCE TaskId import-i-ffvko9js ExpirationTime
2011-06-07T13:30:50+00:00 Status active StatusMessage Pending
InstanceID    i-17912579
DISKIMAGE     DiskImageFormat VMDK   DiskImageSize 893968896 VolumeSize
12            AvailabilityZone us-east-1a   ApproximateBytesCon
verted        73747456     Status active StatusMessage Pending
```

The following response shows that the `IMPORTINSTANCE` status is active, and at 7% progress and that the `DISKIMAGE` is completed.

```
TaskType      IMPORTINSTANCE TaskId import-i-ffvko9js ExpirationTime
2011-06-07T13:30:50+00:00 Status active StatusMessage Progress:
7% InstanceID  i-17912579
DISKIMAGE     DiskImageFormat VMDK   DiskImageSize 893968896 VolumeId
vol-9b59daf0  VolumeSize      12     AvailabilityZone us-
east-1a       ApproximateBytesConverted 893968896 Status completed
```

The following response shows that the `IMPORTINSTANCE` status is completed.

```
TaskType      IMPORTINSTANCE TaskId import-i-ffvko9js ExpirationTime
2011-06-07T13:30:50+00:00 Status completed InstanceID i-
17912579
DISKIMAGE     DiskImageFormat VMDK   DiskImageSize 893968896 VolumeId
vol-9b59daf0  VolumeSize      12     AvailabilityZone us-
east-1a       ApproximateBytesConverted 893968896 Status completed
```

### Note

The `IMPORTINSTANCE` status is what you use to determine the final status. The `DISKIMAGE` status will be completed for a period of time before the `IMPORTINSTANCE` status is completed.

You can now use commands such as `ec2-stop-instance`, `ec2-start-instance`, `ec2-reboot-instance`, and `ec2-terminate-instance` to manage your instance.

### Note

By default, when you terminate an instance, Amazon EC2 does not delete the associated Amazon EBS volume. You can optionally use the `ec2-modify-instance-attribute` command to change this behavior.

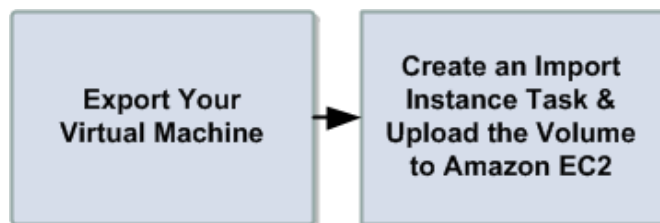
## Importing Your Volumes into Amazon EBS

This section describes how to import your data storage into Amazon EBS, and then attach it to one of your existing Amazon EC2 instances. Amazon EC2 supports importing RAW and VMDK disk formats.

### Important

We recommend utilizing Amazon EC2 security groups to limit network access to your imported instance. Configure a security group to allow only trusted Amazon EC2 instances and remote hosts to connect to RDP and other service ports. For more information about security groups, see [Amazon EC2 Security Groups \(p. 392\)](#).

After you have exported your virtual machine from the virtualization environment, importing the volume to Amazon EBS is a single-step process. You create an import task and upload the volume in one step, as illustrated in the following diagram.



## To import a volume into Amazon EBS

1. Use `ec2-import-volume` to create a task that allows you to upload your volume into Amazon EBS. The syntax of the command is:

```
$ ec2-import-volume DISK_IMAGE_FILENAME -f FORMAT -s SIZE-IN-GB -z AVAILABILITY_ZONE -b S3_BUCKET_NAME -o OWNER -w SECRETKEY
```

The following example creates an import volume task for importing a volume to the us-east-1 region.

### Note

The example uses the VMDK format. You can also use VHD or RAW.

```
Requesting volume size: 25 GB
Disk image format: Stream-optimized VMDK
Converted volume size: 26843545600 bytes (25.00 GiB)
Requested EBS volume size: 26843545600 bytes (25.00 GiB)
TaskType          IMPORTVOLUME   TaskId  import-vol-ffut5xv4   ExpirationTime
2011-09-09T15:22:30+00:00      Status  active  StatusMessage  Pending
DISKIMAGE          DiskImageFormat  VMDK    DiskImageSize  5070303744
```

```
VolumeSize      25      AvailabilityZone    us-east-1d      Approximate
BytesConverted   0
Creating new manifest at tesdfgting/0fd8fcf5-04d8-44ae-981f-
3c9f56d04520/Win_2008_R1_EE_64.vmdkmanifest.xml
Uploading the manifest file
Uploading 5070303744 bytes across 484 parts
0% |-----| 100%
   |=====|
Done
```

Amazon EC2 returns a task ID that you use in the next step. In this example, the ID is `import-vol-ffut5xv4`.

2. Use `ec2-describe-conversion-tasks` to confirm that your volume imported successfully.

```
$ ec2-describe-conversion-tasks import-vol-ffut5xv4
TaskType      IMPORTVOLUME      TaskId  import-vol-ffut5xv4      ExpirationTime
2011-09-09T15:22:30+00:00      Status  completed
DISKIMAGE     DiskImageFormat  VMDK    DiskImageSize  5070303744
VolumeId      vol-365a385c     VolumeSize  25      AvailabilityZone
us-east-1d    ApproximateBytesConverted  5070303744
```

The status in this example is `completed`, which means the import succeeded.

3. Use `ec2-attach-volume` to attach the Amazon EBS volume to one of your existing Amazon EC2 instances. The following example attaches the volume, `vol-2540994c`, to the `i-a149ec4a` instance on the device, `/dev/sde`.

```
$ ec2-attach-volume vol-2540994c -i i-a149ec4a -d /dev/sde
ATTACHMENT vol-2540994c i-a149ec4a /dev/sde attaching 2010-03-
23T15:43:46+00:00
```

## Resuming an Upload

Connectivity problems can interrupt an upload. When you resume an upload, Amazon EC2 automatically starts the upload from where it stopped. The following procedure steps you through determining how much of an upload succeeded and how to resume it.

### To resume an upload

- Use the task ID with `ec2-resume-import` to continue the upload. The command uses the HTTP HEAD action to determine where to resume.

```
ec2-resume-import DISK_IMAGE_FILENAME -t TASK_ID -o OWNER -w SECRETKEY
```

The following example resumes an import instance task.

```
Disk image size: 5070303744 bytes (4.72 GiB)
Disk image format: Stream-optimized VMDK
Converted volume size: 26843545600 bytes (25.00 GiB)
Requested EBS volume size: 26843545600 bytes (25.00 GiB)
Uploading 5070303744 bytes across 484 parts
0% |-----| 100%
   |=====|
```

```
Done
Average speed was 10.316 MBps
The disk image for import-i-ffni8aei has been uploaded to Amazon S3
where it is being converted into an EC2 instance. You may monitor the
progress of this task by running ec2-describe-conversion-tasks. When
the task is completed, you may use ec2-delete-disk-image to remove the
image from S3.
```

## Canceling an Upload

Use `ec2-cancel-conversion-task` to cancel an active conversion task. The task can be the upload of an instance or a volume. The command removes all artifacts of the conversion, including uploaded volumes or instances.

If the conversion is complete or still transferring the final disk image, the command fails and returns an exception similar to the following:

```
Client.CancelConversionTask Error: Failed to cancel conversion task import-i-
fh95npoc
```

### To cancel a conversion task

- Use the task ID of the upload you want to delete with `ec2-cancel-conversion-task`. The following example cancels the upload associated with the task ID `import-i-fh95npoc`.

```
PROMPT> ec2-cancel-conversion-task import-i-fh95npoc
```

The output for a successful cancellation is similar to the following:

```
CONVERSION-TASK import-i-fh95npoc
```

You can use the `ec2-describe-conversion-tasks` command to check the status of the cancellation. For example:

```
$ ./ec2-describe-conversion-tasks import-i-fh95npoc
TaskType      IMPORTINSTANCE  TaskId  import-i-fh95npoc      ExpirationTime
2010-12-20T18:36:39+00:00      Status  cancelled  InstanceID      i-825063ef
DISKIMAGE     DiskImageFormat VMDK    DiskImageSize  2671981568
VolumeSize    40              AvailabilityZone      us-east-1c  ApproximateBytesCon
verted        0              Status  cancelled
```

In the above example, the status is *cancelled*. If it were still in process, the status would be *cancelling*.

## Cleaning up After an Upload

You can use `ec2-delete-disk-image` to remove the image file after it is uploaded. If you do not delete it, you will be charged for its storage in Amazon S3.

### To delete a disk image

- Use the task ID of the disk image you want to delete with `ec2-delete-disk-image`.

The following example deletes the disk image associated with the task ID, `import-i-fh95npoc`.

```
PROMPT> ec2-delete-disk-image import-i-fh95npoc
```

The output for a successful cancellation is similar to the following:

```
DELETE-TASK import-i-fh95npoc
```

## Troubleshooting Instance Importation

When you import a virtual machine using the `ec2-import-instance` command, the import task might stop at 56 percent completion, and then fail. To investigate what went wrong, you can use the `ec2-describe-conversion-tasks` command to describe the instance. Then, you should see the following message:

*FirstBootFailure: This import request failed because the Windows instance failed to boot and establish network connectivity.*

When you receive this message it means that your virtual disk image was unable to perform one of the following steps:

- Boot up and start Windows.
- Install Amazon EC2 networking and disk drivers.
- Use a DHCP-configured network interface to retrieve an IP address.
- Activate Windows using the Amazon EC2 Windows volume license.

## Best Practices for Avoiding First Boot Failures

The following best practices can help you to avoid Windows first boot failures.

### Disable anti-virus and anti-spyware software and firewalls.

These types of software can prevent installing new Windows services or drivers or prevent unknown binaries from running. Software and firewalls can be re-enabled after importing.

### Do not harden your operating system.

Security configurations, sometimes called hardening, can prevent unattended installation of EC2 drivers. There are numerous Windows configuration settings that can prevent import. These settings can be reapplied once imported.

### Disable or delete multiple bootable partitions.

If your virtual machine boots and requires you to choose which boot partition to use, the import may fail.

## Possible Causes of First Boot Failures

This inability of the virtual disk image to boot up and establish network connectivity could be due to any of the following causes.

### Topics

- [The installation of Windows is not valid on the virtual machine. \(p. 331\)](#)
- [TCP/IP networking and DHCP are enabled. \(p. 331\)](#)

- [A volume that Windows requires is missing from the virtual machine.](#) (p. 331)
- [Windows always boots into System Recovery Options.](#) (p. 331)
- [The virtual machine was created using a physical-to-virtual \(P2V\) conversion process.](#) (p. 332)
- [Windows Activation fails.](#) (p. 332)

### **The installation of Windows is not valid on the virtual machine.**

#### **Cause**

The installation of Windows must be valid before you can successfully import the virtual machine.

#### **Resolution**

Do not run SYSPREP before shutting down the Amazon EC2 instance. After the instance is imported, you can run SYSPREP from the instance before you create an Amazon Machine Image (AMI). Importing creates a single instance, so running SYSPREP is not necessary.

Ensure that the installation process is fully complete and that Windows boots (without user intervention) to a login prompt.

### **TCP/IP networking and DHCP are enabled.**

#### **Cause**

For any Amazon EC2 instance, including those in Amazon VPC, TCP/IP networking and DHCP must be enabled. Within a VPC you can define an IP address for the instance either before or after importing the instance. Do not set a static IP address before exporting the instance.

#### **Resolution**

Ensure that TCP/IP networking is enabled. For more information, see [Setting up TCP/IP \(Windows Server 2003\)](#) or [Configuring TCP/IP \(Windows Server 2008\)](#) at the *Microsoft TechNet website*.

Ensure that DHCP is enabled. For more information, see [What is DHCP](#) at the *Microsoft TechNet web site*.

### **A volume that Windows requires is missing from the virtual machine.**

#### **Cause**

Importing a virtual machine into Amazon EC2 only imports the boot disk, all other disks must be detached and Windows must be able to boot before importing the virtual machine. For example, Active Directory often stores the Active Directory database on the D: \ drive. A domain controller cannot boot if the Active Directory database is missing or inaccessible.

#### **Resolution**

Detach any secondary and network disks attached to the Windows virtual machine before exporting.

Move any Active Directory databases from secondary drives or partitions onto the primary Windows partition. For more information, see ["Directory Services cannot start" error message when you start your Windows-based or SBS-based domain controller](#) at the *Microsoft Support website*.

### **Windows always boots into System Recovery Options.**

#### **Cause**

Windows can boot into System Recovery Options for a variety of reasons, including when Windows is pulled into a virtualized environment from a physical machine, also known as P2V.

### Resolution

Ensure that Windows boots to a login prompt before exporting and preparing for import.

Do not import virtualized Windows instances that have come from a physical machine.

**The virtual machine was created using a physical-to-virtual (P2V) conversion process.**

### Cause

A P2V conversion occurs when a disk image is created by performing the Windows installation process on a physical machine and then importing a copy of that Windows installation into a virtual machine. Virtual machines, which are created as the result of a P2V conversion are not supported by Amazon EC2 VM import. Amazon EC2 VM import only supports Windows images that were natively installed inside the source virtual machine.

### Resolution

Install Windows in a virtualized environment and migrate your installed software to that new virtual machine.

**Windows Activation fails.**

### Cause

During boot, Windows will detect a change of hardware and attempt activation. During the import process we attempt to switch the licensing mechanism in Windows to a volume license provided by Amazon Web Services. However, if the Windows activation process does not succeed, then the import will not succeed.

### Resolution

Ensure that the version of Windows you are importing supports volume licensing. Beta or preview versions of Windows may not.

## Exporting EC2 Instances

### Topics

- [Before You Get Started \(p. 332\)](#)
- [Export an Instance \(p. 333\)](#)
- [Cancel or Stop the Export of an Instance \(p. 334\)](#)

If you have previously imported an instance running Microsoft Windows Server into Amazon Elastic Compute Cloud (Amazon EC2), you can use the command line tools to export that Microsoft Windows Server instance to Citrix Xen, Microsoft Hyper-V, or VMware vSphere. Exporting an instance can be useful when you want to deploy a copy of your EC2 instance in your on-site virtualization environment.

### Before You Get Started

Before you begin the process of exporting an instance, you need to be aware of the operating systems and image formats we support, and understand the limitations on exporting instances and volumes. You will also need to download and install the EC2 command line tools and sign up for your private key and X.509 certificate before you use the command line interface (CLI) or the API to export your instance. For more information, see [Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX \(p. 541\)](#).

### Operating Systems

The following operating systems can be exported from Amazon EC2:

- Windows Server 2003 R2 (Standard, Enterprise, and Datacenter)
- Windows Server 2008 (Standard, Enterprise, and Datacenter)
- Windows Server 2008 R2 (Standard, Enterprise, and Datacenter)

## Image Formats

We support the following image formats for exporting both volumes and instances from Amazon Web Services (AWS):

- Stream-optimized ESX Virtual Machine Disk (VMDK) image format, which is compatible with VMware ESX and VMware vSphere versions 4 and 5 virtualization products.
- Open Virtual Appliance (OVA) image format, which is compatible with VMware vSphere versions 4 and 5.
- Virtual Hard Disk (VHD) image format, which is compatible with Citrix Xen and Microsoft Hyper-V virtualization products.

## Known Limitations

The exporting of instances and volumes is subject to the following limitations:

- You cannot export Amazon Elastic Block Store (Amazon EBS) data volumes.
- You cannot export an instance that has more than one virtual disk.
- You cannot export an instance that has more than one network interface.

## Export an Instance

You can use the Amazon EC2 command line interface (CLI) to export an instance. The [ec2-create-instance-export-task](#) command gathers all of the information necessary (e.g., instance ID; name of the S3 bucket that will hold the exported image; name of the exported image; VMDK, OVA, or VHD format) to properly export the instance to the selected virtualization format. The exported file is saved in the Amazon Simple Storage Service (Amazon S3) bucket that you designate.

### Note

When you export an instance, you are charged the standard Amazon S3 rates for the bucket where the exported VM is stored. In addition, a small charge reflecting temporary use of an EBS snapshot might appear on your bill. For more information about Amazon S3 pricing, see [Amazon Simple Storage Service \(S3\) Pricing](#).

### To export an instance

1. Create an Amazon S3 bucket where exported instances will be stored. The S3 bucket must grant **Upload/Delete** and **View Permissions** access to the **vm-import-export@amazon.com** account. For more information, see [Creating a Bucket](#) and [Editing Bucket Permissions](#) in the Amazon Simple Storage Service Console User Guide.
2. At a command prompt, type the following command: `ec2-create-instance-export-task INSTANCE_ID -e TARGET_ENVIRONMENT -f DISK_IMAGE_FORMAT -c CONTAINER_FORMAT -b S3_BUCKET`

Where:

*INSTANCE\_ID* is the ID of the instance you want to export.

*TARGET\_ENVIRONMENT* is VMware, Citrix, or Microsoft.

*DISK\_IMAGE\_FORMAT* is VMDK for VMware or VHD for Microsoft Hyper-V and Citrix Xen.



*CONTAINER\_FORMAT* may be optionally set to OVA when exporting to VMware.

*S3\_BUCKET* is the name of the Amazon S3 bucket to which you want to export the instance.

3. To monitor the export of your instance, at the command prompt, type the following command:  
`ec2-describe-export-tasks TASK_ID`

Where:

*TASK\_ID* is the ID of the export task.

## Cancel or Stop the Export of an Instance

You can use the Amazon EC2 command line interface (CLI) to cancel or stop the export of an instance up to the point of completion. The `ec2-cancel-export-task` command removes all artifacts of the export, including any partially created Amazon S3 objects. If the export task is complete or is in the process of transferring the final disk image, the command fails and returns an error.

### To cancel or stop the export of an instance

- At the command prompt, type: `ec2-cancel-export-task TASK_ID`

Where

*TASK\_ID* is the ID of the export task you want to cancel.

# Monitoring Your Instances

## Topics

- [Monitoring Your Instances with CloudWatch \(p. 336\)](#)
- [Monitoring the Status of Your Instances \(p. 346\)](#)

Amazon Web Services (AWS) automatically provides data, such as Amazon CloudWatch metrics and instance status, that you can use to monitor your Amazon EC2 instances:

- CloudWatch metrics are statistical data you can use to view, analyze, and set alarms on the operational behavior of your instances. These metrics include CPU utilization, network traffic, I/O, and latency.
- Instance status provides two types of information:
  - Instance status checks that summarize results of automated tests that you can use to determine whether your instances are affected by specific, detectable problems.
  - Events that provide information about certain activities that are scheduled for your instances, including operational maintenance that AWS may perform such as rebooting and retirement.

For information about monitoring your Amazon EBS volumes, see [Monitoring the Status of Your Volumes \(p. 464\)](#).

# Monitoring Your Instances with CloudWatch

## Topics

- [Monitoring Instances \(p. 336\)](#)
- [Creating and Editing Status Check Alarms \(p. 344\)](#)

Amazon CloudWatch is a service that collects raw data from partnered AWS products such as Amazon EC2 and then processes the information into readable, near real-time metrics. These statistics are recorded for a period of two weeks, allowing you access to historical information and providing you with a better perspective on how your web application or service is performing. For detailed information about Amazon CloudWatch, see the [Amazon CloudWatch Developer Guide](#).

## Monitoring Instances

The following table describes the types of monitoring data available for your Amazon EC2 instances.

Resource	Type	Description
Instances	Basic	Data is available automatically in 5-minute periods at no charge.
Detailed	Data is available in 1-minute periods at an additional cost. To get this level of data, you must specifically enable it for the instance. For the instances where you've enabled detailed monitoring, you can also get aggregated data across groups of similar instances. For information about pricing, go to the <a href="#">Amazon CloudWatch product page</a> .	

You can get monitoring data for your Amazon EC2 instances using either the Amazon CloudWatch API or the AWS Management Console. The console displays a series of graphs based on the raw data from the Amazon CloudWatch API. Depending on your needs, you might prefer to use either the data from the API or the graphs in the console.

### Data from the Amazon Cloudwatch API

You can use the Amazon CloudWatch `GetMetricStatistics` API action to get any of the instance metrics listed in the following table. The *period* refers to how often the system reports a data point for each metric for an instance. If you've enabled detailed monitoring, each data point covers the instance's previous 1 minute of activity. Otherwise, each data point covers the instance's previous 5 minutes of activity.

Metric	Description
CPUUtilization	<p>The percentage of allocated EC2 compute units that are currently in use on the instance. This metric identifies the processing power required to run an application upon a selected instance.</p> <p>Units: <i>Percent</i></p>
DiskReadOps	<p>Completed read operations from all ephemeral disks available to the instance (if your instance uses Amazon EBS, see <a href="#">Amazon EBS Metrics (p. 464)</a>.)</p> <p>This metric identifies the rate at which an application reads a disk. This can be used to determine the speed in which an application reads data from a hard disk.</p> <p>Units: <i>Count</i></p>
DiskWriteOps	<p>Completed write operations to all ephemeral disks available to the instance (if your instance uses Amazon EBS, see <a href="#">Amazon EBS Metrics (p. 464)</a>.)</p> <p>This metric identifies the rate at which an application writes to a hard disk. This can be used to determine the speed in which an application saves data to a hard disk.</p> <p>Units: <i>Count</i></p>
DiskReadBytes	<p>Bytes read from all ephemeral disks available to the instance (if your instance uses Amazon EBS, see <a href="#">Amazon EBS Metrics (p. 464)</a>.)</p> <p>This metric is used to determine the volume of the data the application reads from the hard disk of the instance. This can be used to determine the speed of the application.</p> <p>Units: <i>Bytes</i></p>
DiskWriteBytes	<p>Bytes written to all ephemeral disks available to the instance (if your instance uses Amazon EBS, see <a href="#">Amazon EBS Metrics (p. 464)</a>.)</p> <p>This metric is used to determine the volume of the data the application writes onto the hard disk of the instance. This can be used to determine the speed of the application.</p> <p>Units: <i>Bytes</i></p>
NetworkIn	<p>The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to an application on a single instance.</p> <p>Units: <i>Bytes</i></p>
NetworkOut	<p>The number of bytes sent out on all network interfaces by the instance. This metric identifies the volume of outgoing network traffic to an application on a single instance.</p> <p>Units: <i>Bytes</i></p>

Metric	Description
StatusCheckFailed	<p>A combination of StatusCheckFailed_Instance and StatusCheckFailed_System that reports if either of the status checks has failed. Values for this metric are either 0 (zero) or 1 (one.) A zero indicates that the status check passed. A one indicates a status check failure.</p> <p><b>Note</b> Status check metrics are available at 5 minute frequency and are not available in Detailed Monitoring. For a newly launched instance, status check metric data will only be available after the instance has completed the initialization state. Status check metrics will become available within a few minutes of being in the running state.</p> <p>Units: <i>Count</i></p>
StatusCheckFailed_Instance	<p>Reports whether the instance has passed the EC2 instance status check in the last 5 minutes. Values for this metric are either 0 (zero) or 1 (one.) A zero indicates that the status check passed. A one indicates a status check failure.</p> <p><b>Note</b> Status check metrics are available at 5 minute frequency and are not available in Detailed Monitoring. For a newly launched instance, status check metric data will only be available after the instance has completed the initialization state. Status check metrics will become available within a few minutes of being in the running state.</p> <p>Units: <i>Count</i></p>
StatusCheckFailed_System	<p>Reports whether the instance has passed the EC2 system status check in the last 5 minutes. Values for this metric are either 0 (zero) or 1 (one.) A zero indicates that the status check passed. A one indicates a status check failure.</p> <p><b>Note</b> Status check metrics are available at 5 minute frequency and are not available in Detailed Monitoring. For a newly launched instance, status check metric data will only be available after the instance has completed the initialization state. Status check metrics will become available within a few minutes of being in the running state.</p> <p>Units: <i>Count</i></p>

**Note**

When you get data from Amazon CloudWatch, you can include a `Period` request parameter to specify the granularity of the returned data. This is different than the period we use when we collect the data (either 1-minute periods for detailed monitoring, or 5-minute periods for basic monitoring). We recommend that you specify a period in your request that is equal to or larger than the collection period to ensure that the returned data is valid.

You can use the dimensions in the following table to refine the metrics returned for your instances.

Dimension	Description
AutoScalingGroupName	This dimension filters the data you request for all instances in a specified capacity group. An <i>AutoScalingGroup</i> is a collection of instances you define if you're using the Auto Scaling service. This dimension is available only for EC2 metrics when the instances are in such an AutoScalingGroup. Available for instances with Detailed or Basic Monitoring enabled.
ImageId	This dimension filters the data you request for all instances running this EC2 Amazon Machine Image (AMI). Available for instances with Detailed Monitoring enabled.
InstanceId	This dimension filters the data you request for the identified instance only. This helps you pinpoint an exact instance from which to monitor data. Available for instances with Detailed Monitoring enabled.
InstanceType	This dimension filters the data you request for all instances running with this specified instance type. This helps you categorize your data by the type of instance running. For example, you might compare data from an m1.small instance and an m1.large instance to determine which has the better business value for your application. Available for instances with Detailed Monitoring enabled.

For more information about using the `GetMetricStatistics` action, see [GetMetricStatistics](#) in the *Amazon CloudWatch API Reference*.

## Graphs in the AWS Management Console

After you launch an instance, you can go to the AWS Management Console and view the instance's monitoring graphs. They're displayed when you select the instance on the **Instances** page in the EC2 Dashboard. A **Monitoring** tab is displayed next to the instance's **Description** tab. The following graphs are available:

- Average CPU Utilization (Percent)
- Average Disk Reads (Bytes)
- Average Disk Writes (Bytes)
- Maximum Network In (Bytes)
- Maximum Network Out (Bytes)
- Summary Disk Read Operations (Count)
- Summary Disk Write Operations (Count)
- Summary Status (Any)
- Summary Status Instance (Count)
- Summary Status System (Count)

The AWS Management Console contains a console for Amazon CloudWatch. In the Amazon CloudWatch console you can search and browse all your AWS resource metrics, view graphs to troubleshoot issues and discover trends, create and edit alarms to be notified of problems, and see at-a-glance overviews of your alarms and AWS resources. For more information, see [AWS Management Console](#) in the *Amazon CloudWatch Developer Guide*.

## Enabling or Disabling Detailed Monitoring on an Amazon EC2 Instance

This section describes how to enable or disable detailed monitoring on either a new instance (as you launch it) or on a running or stopped instance. After you enable detailed monitoring, the Amazon EC2 console in the AWS Management Console displays monitoring graphs with a 1-minute period for the instance. You can enable or disable detailed monitoring using the console or the command line interface (CLI).

### AWS Management Console

#### To enable detailed monitoring of an existing EC2 instance

You can enable detailed monitoring of your EC2 instances, which provides data about your instance in 1-minute periods. (There is an additional charge for 1-minute monitoring.) Detailed data is then available for the instance in the AWS Management Console graphs or through the API. To get this level of data, you must specifically enable it for the instance. For the instances on which you've enabled detailed monitoring, you can also get aggregated data across groups of similar instances. An instance must be running or stopped to enable detailed monitoring.

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the **Navigation** pane, click **Instances**.
3. In the list of instances, select a running or stopped instance, click **Actions**, and then click **Enable Detailed Monitoring**.
4. In the **Enable Detailed Monitoring** dialog box, click **Yes, Enable**.
5. In the **Enable Detailed Monitoring** confirmation dialog box, click **Close**.

Detailed data (collected with a 1-minute period) is then available for the instance in the AWS Management Console graphs or through the API.

#### To enable detailed monitoring when launching an EC2 instance

- When launching an instance with the AWS Management Console, select the **Monitoring** check box on the **Configure Instance Details** page of the launch wizard.

After the instance is launched, you can select the instance in the console and view its monitoring graphs on the instance's **Monitoring** tab in the lower pane.

#### To disable detailed monitoring of an EC2 instance

When you no longer want to monitor your instances at 1-minute intervals, you can disable detailed monitoring and use basic monitoring instead. Basic monitoring provides data in 5-minute periods at no charge.

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the **Navigation** pane, click **Instances**.
3. In the list of instances, select a running or stopped instance, click **Actions**, and then click **Disable Detailed Monitoring**.
4. In the **Disable Detailed Monitoring** dialog box, click **Yes, Disable**.
5. In the **Disable Detailed Monitoring** confirmation dialog box, click **Close**.

For information about launching instances, see [Launch Your Instance \(p. 266\)](#).

## Command Line Interface

### To enable detailed monitoring on an existing instance

- Use the `ec2-monitor-instances` command with one or more instance IDs.

```
PROMPT> ec2-monitor-instances i-1a2b3c4d
i-1a2b3c4d    monitoring-pending
```

Detailed data (collected with a 1-minute period) is then available for the instance in the AWS Management Console graphs or through the API.

### To enable detailed monitoring when launching an instance

- Use the `ec2-run-instances` command with the `--monitor` flag.

```
PROMPT> ec2-run-instances ami-2bb65342 -k gsg-keypair --monitor
```

Amazon EC2 returns output similar to the following example. The status of monitoring is listed as *monitoring-enabled*.

```
RESERVATION    r-7430c31d    111122223333    default
INSTANCE       i-ae0bf0c7    ami-2bb65342    pending gsg-keypair    0
  m1.small     2008-03-21T16:19:25+0000    us-east-1a    monitoring-enabled
```

After the instance is running, detailed data (collected with a 1-minute period) is then available for the instance in the AWS Management Console graphs or through the API.

### To disable detailed monitoring of an instance

- Use the `ec2-unmonitor-instances` command with one or more instance IDs.

```
PROMPT> ec2-unmonitor-instances i-1a2b3c4d
```

## API

### To enable detailed monitoring when launching an instance

- Set the `Monitoring.Enabled` parameter to `true` in the `RunInstances` request.

Following is an example Query request.

```
https://ec2.amazonaws.com/
?Action=RunInstances
&ImageId=ami-id
&MaxCount=1
&MinCount=1
&KeyName=keypair-name
&Monitoring.Enabled=true
&...auth parameters...
```



Following is an example response.

```
<RunInstancesResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <reservationId>r-47a5402e</reservationId>
  <ownerId>111122223333</ownerId>
  <groupSet>
    <item>
      <groupId>default</groupId>
    </item>
  </groupSet>
  <instancesSet>
    <item>
      <instanceId>i-2ba64342</instanceId>
      <imageId>ami-60a54009</imageId>
      <instanceState>
        <code>0</code>
        <name>pending</name>
      </instanceState>
      ...
      <monitoring>
        <state>pending</state>
      </monitoring>
      ...
    </item>
  </instancesSet>
</RunInstancesResponse>
```

After the instance is launched, detailed data (collected with a one-minute period) is then available for the instance in the AWS Management Console graphs or through the API.

### To enable detailed monitoring on an existing instance

- Use the `MonitorInstances` action.

Following is a sample Query request.

```
https://ec2.amazonaws.com/
?Action=MonitorInstances
&InstanceId.1=i-1a2b3c4d
&...auth parameters...
```

Following is an example response.

```
<MonitorInstancesResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <instancesSet>
    <item>
      <instanceId>i-1a2b3c4d</instanceId>
      <monitoring>
        <state>pending</state>
      </monitoring>
    </item>
  </instancesSet>
</MonitorInstancesResponse>
```

Detailed data (collected with a one-minute period) is then available for the instance in the AWS Management Console graphs or through the API.

## Create a CloudWatch Alarm

You can create an Amazon CloudWatch alarm that monitors any one of your Amazon EC2 instance's CloudWatch metrics. CloudWatch will automatically send you a notification when the metric reaches a threshold you specify. You can create a CloudWatch alarm on the Amazon EC2 console of the AWS Management Console, or you can use the CloudWatch console and configure more advanced options.

### To create a CloudWatch alarm for high CPU

1. On the [Amazon EC2 console](#), select the instance for which you want to create an alarm.
2. On the instance's **Monitoring** tab in the lower pane, click **Create Alarm**.
3. In the **Create Alarm** dialog box, set the criteria for your alarm. In this example, we'll set an alarm if the instance's average CPU utilization is above 70 percent.
4. The check box next to Send a notification is selected by default. Select an existing topic, or click **Create topic** and enter a name. (Notifications use Amazon Simple Notification Service (Amazon SNS)).
5. In the **With these recipients** box, enter the email addresses of the recipients you want to notify. You can enter up to 10 email addresses, each separated by a comma.
6. Configure the threshold for your alarm.
  - a. In the **Whenever** boxes, select **Average** and **CPU Utilization**.
  - b. In the **Is** boxes, define the threshold for the alarm by selecting **>** and entering **70**.
  - c. In the **For at least** boxes, specify the sampling period and number of samples evaluated by the alarm. You can leave the defaults or define your own. For our example, we'll monitor for 1 period of 15 minutes.

#### Note

A shorter period creates a more sensitive alarm. A longer period can mitigate brief spikes in a metric.

- d. In **Name of alarm**, a name is automatically generated for you. You can type in the field to change the name.

#### Important

You cannot modify the name after you create the alarm.

**Create Alarm** ✕

You can use CloudWatch alarms to be notified automatically whenever metric data reaches a level you define.  
To edit an alarm, first choose whom to notify and then define when the notification should be sent.

**Send a notification to:**  [cancel](#)

**With these recipients:**

**Take the action:**  Stop  Terminate this instance.

**Whenever:**  of

**Is:**   Percent

**For at least:**  consecutive period(s) of

**Name of alarm:**

[Cancel](#) [Create](#)

**CPU Utilization Percent**

60  
40  
20  
0

10/7 10/7 10/7  
08:00 10:00 12:00

i-1a2b3c4d

7. Click **Create**.

After you create the alarm, you can use the **Monitoring** tab in the Amazon EC2 console to view a summary of alarms that have been set for that instance. From there, you can also edit the alarm.

#### Note

If you created a new Amazon SNS topic for this alarm or added new email addresses to an existing topic, each email address added will receive a subscription confirmation email from Amazon SNS. The person who receives the email must confirm it by clicking the included link in order to receive notifications.

## Creating and Editing Status Check Alarms

You can create instance status and system status alarms to notify you when an instance has a failed status check. To create or change these alarms you can use either the AWS Management Console or the command line interface (CLI).

### AWS Management Console

#### To create a status check alarm

You can create status check alarms for an existing instance to monitor instance status or system status. You can configure the alarm to send you a notification by email when an instance fails an instance check or system status check.

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the **Navigation** pane, click **Instances**.
3. Select an instance, and then on the **Status Checks** tab, click **Create Status Check Alarm**.
4. In the **Create Alarm** dialog box, select the **Send a notification to** check box, and then choose an existing Amazon Simple Notification Service (SNS) topic or create a new SNS topic to use for this alarm.
5. In the **With these recipients** box, type your email address (e.g., john.stiles@example.com) and the addresses of any additional recipients, separated by commas.
6. In the **Whenever** drop-down list, select the status check you want to be notified about (e.g., Status Check Failed (Any), Status Check Failed (Instance), or Status Check Failed (System)).

7. In the **For at least** box, set the number of periods you want to evaluate (for example, 2) and in the **consecutive periods** drop-down menu, select the evaluation period duration (for example, 5 minutes) before triggering the alarm and sending an email
8. To change the default name for the alarm, in the **Name of alarm** box, type a friendly name for the alarm (for example, StatusCheckFailed), and then click **Create**.

#### Important

If you added an email address to the list of recipients or created a new topic, Amazon SNS will send a subscription confirmation email message to each new address shortly after you create an alarm. Remember to click the link contained in that message, which confirms your subscription. Alert notifications are only sent to confirmed addresses.

### To edit a status check alarm

If you need to make changes to an instance status alarm, you can edit it.

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the **Navigation** pane, click **Instances**.
3. Select an instance, click **Actions**, and then click **Add/Edit Alarms**.
4. In the **Alarm Details** dialog box, click the name of the alarm.
5. In the **Edit Alarm** dialog box, make the desired changes, and then click **Save**.

## Command Line Interface

### To create a status check alarm using the CLI

You can create a status check alarm using the Amazon CloudWatch CLI. In the following example, the alarm publishes a notification to a specific SNS topic that has the ARN `arn:aws:sns:us-east-1:111111111111:StatusCheckNotifications` when instance `i-ab12345` fails either the instance check or system status check for at least two periods. (The metric is `StatusCheckFailed`.) For more information, see the [mon-put-metric-alarm](#) command in the *Amazon CloudWatch Developer Guide*.

1. At a command prompt, type **mon-list-metrics --headers** to view the list of all available Amazon CloudWatch metrics for the services in AWS that you're using.
2. In the list of metrics, look in the **Namespace** column (second column), and review the Status Check metrics that have the `AWS/EC2` namespace. These are the status check metrics that you can use to create a status check alarm.
3. At the command prompt, enter the following command:

```
mon-put-metric-alarm --alarm-name StatusCheckFailed-Alarm-for-i-ab12345
  --alarm-description "Alarm when StatusCheckFailed metric has a value of
  one for two
  periods" --metric-name StatusCheckFailed --namespace AWS/EC2 --statistic
  Maximum
  --period 300 --threshold 1 --comparison-operator GreaterThanThreshold
  --dimensions
  "InstanceId=i-ab12345" --evaluation-periods 2 --alarm-actions
  arn:aws:sns:us-east-1:111111111111:StatusCheckNotifications --unit Count
```

Where:

The **--alarm-name** is the name of the alarm. This is required.

The **--alarm-description** is a friendly description of the alarm.

The **--metric-name** is one of the available status metrics (e.g., `StatusCheckFailed`, `StatusCheckFailed_Instance`, or `StatusCheckFailed_System`). This is required.

The **--statistic** is one of the following values: `Average`, `Sum`, `Minimum`, or `Maximum`. This is required.

The **--period** is the time frame (in seconds) in which Amazon CloudWatch metrics are collected. In this example, you would enter 300, which is 60 seconds multiplied by 5 minutes. This is required.

The **--threshold** is the value to which the metric will be compared (e.g., 1). This is required.

The **--namespace** is the metric's namespace (e.g., `AWS/EC2`). This is required.

The **--dimensions** are associated with the metric (e.g., `InstanceId=i-ab12345`).

The **--evaluation-periods** is the number of consecutive periods for which the value of the metric must be compared to the threshold. This is required.

The **--alarm-actions** is the list of actions to perform when this alarm is triggered. Each action is specified as an Amazon Resource Number (ARN). In this example, we want the alarm to send us an email using Amazon SNS.

#### Note

You can find the ARN for the Amazon SNS topic that the alarm will use in the Amazon SNS console:

1. Open the Amazon SNS console at <https://console.aws.amazon.com/sns/>.
2. On the **Navigation** pane, under **My Topics**, select the topic you want the alarm to send mail to.
3. The ARN is located in the **Topic ARN** field on the **Topic Details** pane.

The **--unit** is the unit of the metric on which to alarm (e.g., `Count`).

## Monitoring the Status of Your Instances

You can monitor the status of your instances by viewing status checks and scheduled events for your instances. A status check gives you the information that results from automated checks performed by Amazon EC2. These automated checks detect whether specific issues are affecting your instances. The status check information, together with the data provided by Amazon CloudWatch, gives you detailed operational visibility into each of your instances.

You can also see status on specific events scheduled for your instances. Events provide information about upcoming activities such as rebooting or retirement that are planned for your instances, along with the scheduled start and end time of each event.

#### Topics

- [Monitoring Instances with Status Checks \(p. 346\)](#)
- [Monitoring Events for Your Instances \(p. 349\)](#)

## Monitoring Instances with Status Checks

With instance status monitoring you can quickly determine whether Amazon EC2 has detected any problems that may prevent your instances from running applications. Amazon EC2 performs automated checks on every running Amazon EC2 instance to identify hardware and software issues. You can view the results of these status checks to identify specific and detectable problems. This data augments the information that Amazon EC2 already provides about the intended state of each instance (pending,

running, stopping, etc.) as well as the utilization metrics that Amazon CloudWatch monitors (CPU utilization, network traffic, and disk activity).

Status checks are performed every five minutes and each returns a pass or a fail status. If all checks pass, the overall status of the instance is **OK**. If one or more checks fail, the overall status is **impaired**. Status checks are built into Amazon EC2, so they cannot be disabled or deleted. You can, however create or delete alarms that are triggered based on the result of the status checks. For example, you can create an alarm to warn you if status checks fail on a specific instance. For more information, see [Creating and Editing Status Check Alarms \(p. 344\)](#).

There are two types of status checks: system status checks and instance status checks.

**System status checks** monitor the AWS systems required to use your instance to ensure they are working properly. These checks detect problems with your instance that require AWS involvement to repair. When a system status check fails, you can choose to wait for AWS to fix the issue or you can resolve it yourself (for example, by stopping and restarting or terminating and replacing an instance). Examples of problems that cause system status checks to fail include:

- Loss of network connectivity
- Loss of system power
- Software issues on the physical host
- Hardware issues on the physical host

**Instance status checks** monitor the software and network configuration of your individual instance. These checks detect problems that require your involvement to repair. When an instance status check fails, typically you will need to address the problem yourself (for example by rebooting the instance or by making modifications in your operating system). Examples of problems that may cause instance status checks to fail include:

- Failed system status checks
- Misconfigured networking or startup configuration
- Exhausted memory
- Corrupted file system
- Incompatible kernel

#### Note

Status checks that occur during instance reboot or while a Windows instance store-backed instance is being bundled will report an instance status check failure until the instance becomes available again.

## Viewing Status

AWS provides you with several ways to view and work with status checks: You can use the AWS Management Console, interact directly with the API, or use the command line interface.

### AWS Management Console

#### To view status checks

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the Navigation pane, click **Instances**.
3. On the **Instances** page, the **Status Checks** column lists the operational status of each instance.
4. To view an individual instance's status, select the instance, and then click the **Status Checks** tab.

## Amazon Elastic Compute Cloud User Guide Monitoring the Status of Your Instances

Description **Status Checks** Monitoring Tags

Status checks detect problems that may impair this instance from running your applications. [Learn more](#) about status checks.

[Create Status Check Alarm](#)

**System Status Checks** ⓘ

These checks monitor the AWS systems required to use this instance and ensure they are functioning properly.

System reachability check passed

**Instance Status Checks** ⓘ

These checks monitor your software and network configuration for this instance.

Instance reachability check failed at October 7, 2013 11:52:11 AM UTC+2 (16 minutes ago)

[Learn more about this issue](#)

**Additional Resources**

[Submit feedback](#) if our checks do not reflect your experience with this instance or if they do not detect the issues you are having. Please note that we will not respond to customer support issues reported via this form. Please post your issue on the [Developer Forums](#) or contact [AWS Support](#) if you need technical assistance with this instance.

### Note

If you have an instance with a failed status check and the instance has been unreachable for over 20 minutes, you can click **Contact AWS Support** to submit a request for assistance. To try and troubleshoot system or instance status check failures yourself, see [Troubleshooting Instances with Failed Status Checks](#) (p. 360).

### Command Line Tools

To do this	Run this command
Get the status of all instances	<code>ec2-describe-instance-status</code>
Get the status of all instances with a instance status of impaired	<code>ec2-describe-instance-status --filter "instance-status.status=impaired"</code>
Get the status of all instances with a single instance with instance ID i-15a4417c	<code>ec2-describe-instance-status status -i-15a4417c</code>

For more information about using the **ec2-describe-instance-status** command, see [ec2-describe-instance-status](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

### Note

If you have an instance with a failed status check, see [Troubleshooting Instances with Failed Status Checks](#) (p. 360)

### API

You can use the `DescribeInstanceState` action to retrieve the status of your instances. For more information, see [DescribeInstanceState](#) in the *Amazon Elastic Compute Cloud API Reference*.

### Reporting Status

You can provide feedback about your instances if you are having problems with an instance whose status is not shown as impaired, or to send AWS additional details about the problems you are experiencing with an impaired instance.

We use reported feedback to identify issues impacting multiple customers, but do not respond to individual account issues reported via this form. Providing feedback will not change the status check results that you currently see for this instance.

If you are in need of technical assistance specific to your account, please post your question to the Developer Forums or contact Premium Support.

## AWS Management Console

### To report status feedback using the management console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the Navigation pane, click **Instances**.
3. On the **Instances** page, click on the instance on which you want to report status.
4. Click the **Status Checks** tab and then click **Submit feedback**.
5. Complete the information on the **Report Instance Status** page.

## Command Line Tools

Use the `ec2-report-instance-status` command to send status feedback using the command line tools. The command uses the following syntax:

```
ec2-report-instance-status [instance_id ...] [--status ...] [--reason] ..]
```

For more information about using the `ec2-report-instance-status` command, see [ec2-report-instance-status](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

## API

You can use the `ReportInstanceStatus` action to submit feedback about a running instance's status. If your experience with the instance differs from the instance status returned by the `DescribeInstanceStatus` action, use `ReportInstanceStatus` to report your experience with the instance. Amazon EC2 collects this information to improve the accuracy of status checks. For more information, see [ReportInstanceStatus](#) in the *Amazon Elastic Compute Cloud API Reference*.

## Monitoring Events for Your Instances

### Topics

- [Monitoring Events with Instance Status \(p. 350\)](#)
- [Working with an Instance that Has a Scheduled Event \(p. 351\)](#)

Instance status describes specific events that Amazon Web Services (AWS) may schedule for your instances. Events provide information about upcoming activities, such as rebooting or retirement, that are planned for your instances, along with the scheduled start and end time of each event. These scheduled events are not frequent. There are several types of scheduled events:

- **Instance reboot:** AWS may schedule an instance for a reboot for necessary maintenance, such as to apply patch upgrades to an underlying host that *do not* require the host to be rebooted.
- **System reboot:** AWS may schedule an instance for a reboot for necessary maintenance, such as to apply patch upgrades that *do* require the host to be rebooted.
- **Network maintenance:** AWS may schedule network maintenance that includes a scheduled start and end time, during which your instances will not have network connectivity. You will be notified by email if one of your instances is set for network maintenance. The email message indicates when your instance will not have network connectivity.
- **Power maintenance:** AWS may schedule power maintenance that includes a scheduled start and end time, during which your instances may be offline for an extended period and then be rebooted. You will be notified by email if one of your instances is set for power maintenance. The email message indicates when your instance will be rebooted.
- **Instance retirement:** AWS may schedule instances for retirement in cases where there is an unrecoverable issue with the hardware on an underlying host. You will also be notified by email if one



of your instances is set to retiring. The email message indicates when your instance will be permanently retired.

- **Instance stop:** AWS may schedule instances to stop in cases where there is an unrecoverable issue with the hardware on an underlying host. You will also be notified by email if one of your instances is set to stop. The email message indicates when your instance will be stopped.

### Important

No action is required on your part if one of your instances is scheduled for reboot. We recommend that you wait for the reboot to occur within its scheduled maintenance window.

For instances scheduled for network maintenance, power maintenance, stop, or retirement, we recommend that you take the actions detailed later in this section.

## Monitoring Events with Instance Status

You can view scheduled events for your instances using the AWS Management Console, the command line interface (CLI), or the API.

### AWS Management Console

#### To view scheduled events for your instances

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the **Navigation** pane, click **Events**. You can see a list of all resources with events associated with them. You can view only instances list of instances that have upcoming events scheduled.
3. Alternatively, you can do the following to view upcoming scheduled events:
  - a. In the **Navigation** pane, click the **EC2 Dashboard**.
  - b. Under Events, you can see the events associated with your Amazon EC2 instances and volumes.
  - c. On the **Events** page, in **Viewing**, select **Instances** to view only instances. You can also filter on specific status types.

### Command Line Interface

#### To view scheduled events for your instances

- Enter the following command:

```
ec2-describe-instance-status
```

Amazon EC2 returns output similar to the following:

```
INSTANCE i-1a2b3c4d us-east-1d running 16 ok ok active
SYSTEMSTATUS reachability passed
INSTANCESTATUS reachability passed
INSTANCE i-2a2b3c4d us-east-1d running 16 ok ok active
SYSTEMSTATUS reachability passed
INSTANCESTATUS reachability passed
INSTANCE i-3a2b3c4d us-east-1d running 16 ok ok active
SYSTEMSTATUS reachability passed
INSTANCESTATUS reachability passed
INSTANCE i-4a2b3c4d us-east-1d running 16 ok ok retiring YYYY-MM-
DDTHH:MM:SS+0000
```

```
SYSTEMSTATUS reachability passed
INSTANCESTATUS reachability passed
EVENT instance-stop YYYY-MM-DDTHH:MM:SS+0000 The instance is running on de
graded hardware
INSTANCE i-5a2b3c4d us-east-1d running 16 ok ok retiring YYYY-MM-
DDTHH:MM:SS+0000
SYSTEMSTATUS reachability passed
INSTANCESTATUS reachability passed
EVENT instance-retiring YYYY-MM-DDTHH:MM:SS+0000 The instance is running on
degraded hardware
INSTANCE i-6a2b3c4d us-east-1d running 16 ok ok retiring YYYY-MM-
DDTHH:MM:SS+0000
SYSTEMSTATUS reachability passed
INSTANCESTATUS reachability passed
EVENT instance-stop YYYY-MM-DDTHH:MM:SS+0000 The instance is running on de
graded hardware
```

For more information about using the **ec2-describe-instance-status** command, see [ec2-describe-instance-status](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

## API

You can use the `DescribeInstanceStatus` action to retrieve the status of your instances. For more information, see [DescribeInstanceStatus](#) in the *Amazon Elastic Compute Cloud API Reference*.

## Working with an Instance that Has a Scheduled Event

This section discusses tasks you can perform if your instance has one of the following scheduled events.

- Instance or system reboot
- Network maintenance
- Power maintenance
- Instance retirement

### Recommended Tasks for Instances Scheduled for a Reboot

AWS may schedule instances for a reboot to perform maintenance tasks such as patch upgrades to the software on an underlying host. These scheduled reboots are not frequent. There are two types of reboot events: system reboot and instance reboot. In either case, your instance will be rebooted. During a system reboot, the hardware supporting your instance will also be rebooted. During an instance reboot, your instance will be rebooted but the hardware supporting the instance will not be rebooted.

#### Important

No action is required on your part if one of your instances is scheduled for reboot. We recommend that you wait for the reboot to occur automatically within its scheduled maintenance window.

Scheduled reboot events start within their scheduled maintenance window. After initiation, both system and instance reboots typically complete in a matter of minutes. After a reboot completes, your instance is available to use; it is not necessary to wait until the scheduled end time.

To verify that the reboot has occurred, check your scheduled events and verify that the instance no longer shows a scheduled event. We recommend that you check your instance after it is rebooted to ensure that your application is functioning as you expected.

## Optional Alternative Tasks for Instances Scheduled for Reboot

We recommend that you wait for the reboot to occur automatically within its scheduled window. If you choose, you may perform the instance reboot yourself to control the timing of the event.

### Instance Reboot

For instance reboot events, you can perform the reboot yourself (be careful to not shut down or terminate your instance) at any time before the scheduled reboot event begins. The reboot can be initiated using the AWS Management Console, a `RebootInstances` API call, or from within the instance (e.g., at the command prompt). After you reboot, any pending maintenance to the underlying host is performed automatically, and you can begin using your instance again after the instance has fully booted.

#### Note

After you perform the reboot, the scheduled event for the instance reboot is canceled immediately. The event's description is updated in the AWS console to reflect this.

### System Reboot

If you choose to perform a system reboot, your course of action will differ depending on whether your instance's root device volume is an EBS volume or an instance store volume. You can determine the root device type for an instance using either the [DescribeInstances API](#) or the AWS Management Console. In the console, you select an instance and view the root device type listed in the **Description** tab.

## Instances Scheduled for Network Maintenance, Power Maintenance, Stop, or Retirement

The following section outlines steps you can take to migrate or replace instances that are scheduled for network maintenance, power maintenance, or retirement. If you migrate your instances scheduled for network maintenance, power maintenance, stop, or retirement before the maintenance start time, the maintenance event will automatically be cancelled.

## Instances Backed by Amazon EBS

If your instance's root device is an EBS volume, you can stop and restart it (be careful not to shut down or terminate the instance). If you choose to do so, some configuration settings will change. The following changes occur:

- Data in instance store volumes will no longer be available. Before you stop the instance, back up any data you may need. For more information about instance store volumes, see [Amazon EC2 Instance Store \(p. 508\)](#).
- The public DNS name and the private IP address of the instance will change.
- If you associated an Elastic IP address with this instance, stopping this instance also disassociates the Elastic IP address from it (charges still apply). After you start the instance, re-associate it with the Elastic IP address if the address is still required.
- For instances in Amazon Virtual Private Cloud (Amazon VPC), the Elastic IP address and the private IP address remain unchanged.

Before you stop and restart an instance, perform the following tasks:

1. Retrieve any data from instance store that you will need later. This data will not be available after you stop and restart an instance.
2. Take a snapshot of your existing volume (storage charges will apply).
3. Note the necessary configuration data in case you will need it later, including the DNS public name and the private IP address.
4. Stop and restart your instance. For more information, see [Stopping and Starting Your Instances \(p. 285\)](#).
5. Re-associate an Elastic IP address, if this address is necessary.

6. If other applications or instances rely on the public DNS name or the private IP address of this instance, update the information with the new configuration data.

**Note**

After you stop and restart the instance, the scheduled event is canceled immediately. The event's description is updated in the AWS console to reflect this.

### Instances Backed by Instance Store

If your instance's root device is an instance store volume and your instance is scheduled for network maintenance, power maintenance, stop, or retirement, you can opt to launch a replacement instance. After you have launched a replacement instance, you can terminate the original instance. If you choose to take this action, be aware that:

- The public DNS name and the private IP address of your replacement instance will differ from your original instance.
- You must first ensure that you have an AMI with the configuration that you want to use when launching the new instance. For more information, see:
  - [Creating an Instance Store-backed Windows AMI](#)
  - [Creating Instance Store-Backed Linux/UNIX AMIs](#) (p. 66)

## Troubleshooting Instances

The following documentation can help you troubleshoot problems that you might have with your instance.

### Topics

- [What To Do If An Instance Immediately Terminates](#) (p. 353)
- [Troubleshooting Connecting to Your Instance](#) (p. 355)
- [Troubleshooting Stopping Your Instance](#) (p. 358)
- [Troubleshooting Terminating \(Shutting Down\) Your Instance](#) (p. 359)
- [Troubleshooting Instances with Failed Status Checks](#) (p. 360)
- [Troubleshooting Instance Capacity](#) (p. 383)
- [Getting Console Output and Rebooting Instances](#) (p. 383)

You can also search for answers and post questions on the [Amazon EC2 forum](#).

## What To Do If An Instance Immediately Terminates

After you launch an instance, we recommend that you check its status to confirm that it goes from the `pending` status to the `running` status, the `not_terminated` status.

The following are a few reasons why an Amazon EBS-backed instance might immediately terminate:

- You've reached your volume limit. For information about the volume limit, and to submit a request to increase your volume limit, see [Request to Increase the Amazon EBS Volume Limit](#).
- The AMI is missing a required part.
- The snapshot is corrupt.

You can use the Amazon EC2 console, CLI, or API to get information about the reason that the instance terminated.

## AWS Management Console

### To get the reason that an instance terminated

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Instances** to display the instance details.
3. Select your instance.
4. In the **Description** tab, locate the reason next to the label **State transition reason**. If the instance is still running, there's typically no reason listed. If you've explicitly stopped or terminated the instance, the reason is `User initiated shutdown`.

## Command Line Tools

### To get the reason that an instance terminated

1. Use the `ec2-describe-instances` command in verbose mode:

```
PROMPT> ec2-describe-instances instance_id -v
```

2. In the XML response that's displayed, locate the `stateReason` element. It looks similar to the following example.

```
<stateReason>
  <code>Client.UserInitiatedShutdown</code>
  <message>Client.UserInitiatedShutdown: User initiated shutdown</message>
</stateReason>
```

This example response shows the reason code that you'll see after you stop or terminate a running instance. If the instance terminated immediately, you'll see `code` and `message` elements that describe the reason that the instance terminated (for example, `VolumeLimitExceeded`).

## API

### To get information about why the instance terminated

1. Construct the following Query request. For information about AUTHPARAMS, see [Common Query Parameters](#) in the *Amazon Elastic Compute Cloud API Reference*.

```
https://ec2.amazonaws.com/
?Action=DescribeInstances
&AUTHPARAMS
```

2. In the XML response that's displayed, locate the `stateReason` element. It looks similar to the following example.

```
<stateReason>
  <code>Client.UserInitiatedShutdown</code>
  <message>Client.UserInitiatedShutdown: User initiated shutdown</message>
</stateReason>
```

This example response shows the reason code that you'll see after you stop or terminate a running instance. If the instance terminated immediately, you'll see `code` and `message` elements that describe the reason that the instance terminated (for example, `VolumeLimitExceeded`).

## Troubleshooting Connecting to Your Instance

The following are possible problems you may have and error messages you may see while trying to connect to your instance.

### Topics

- [Error connecting to your instance: Connection timed out \(p. 355\)](#)
- [Error: User key not recognized by server \(p. 356\)](#)
- [Error: Host key not found or Authentication failed, permission denied \(p. 357\)](#)
- [Error: Server refused our key or No supported authentication methods available \(p. 358\)](#)
- [Error using MindTerm on Safari Browser \(p. 358\)](#)


## Error connecting to your instance: Connection timed out

If you try to connect to your instance and get an error message `Network error: Connection timed out` or `Error connecting to [instance], reason: -> Connection timed out: connect`, try the following:

- Check your security group rules. You need a security group rule that allows inbound traffic to the proper port(s) from your public IP address. When you use SSH to connect to your Linux instance, this is usually port 22. When you use RDP to connect to your Windows instance, this is usually port 3389. For more information, see [Authorizing Network Access to Your Instances \(p. 412\)](#).
- Check the CPU load on your instance; the server may be overloaded. AWS automatically provides data such as Amazon CloudWatch metrics and instance status, which you can use to see how much CPU load is on your instance and, if necessary, adjust how your loads are handled. For more information, see [Monitoring Your Instances with CloudWatch \(p. 336\)](#).
  - If your load is variable, you can automatically scale your instances up or down using [Auto Scaling](#) and [Elastic Load Balancing](#).
  - If your load is steadily growing, you can move to a larger instance type. For more information, see [Resizing Your Instance \(p. 111\)](#).
- Verify that you are using the private key file that corresponds to the key pair that was selected when the instance was launched. For more information, see [Amazon EC2 Key Pairs \(p. 385\)](#).
- Verify that you are connecting with the appropriate user name for your AMI.
  - For an Amazon Linux AMI, the user name is `ec2-user`.
  - For a RHEL5 AMI, the user name is often `root` but might be `ec2-user`.
  - For an Ubuntu AMI, the user name is `ubuntu`.
  - Otherwise, check with your AMI provider.

If you are using MindTerm to connect, enter the user name in the **User name** box in the **Connect To Your Instance** window.

If you are using PuTTY to connect, enter the user name in the **Host name** box in the **PuTTY Configuration** window.



## Error: User key not recognized by server

### If you use SSH to connect to your instance

- Use `ssh -vvv` to get triple verbose debugging information while connecting:

```
#ssh -vvv -i [your key name].pem ec2-user@[public DNS address of your instance].compute-1.amazonaws.com
```

The following sample output demonstrates what you might see if you were trying to connect to your instance with a key that was not recognized by the server:

```
open/ANT/myusername/.ssh/known_hosts).
debug2: bits set: 504/1024
debug1: ssh_rsa_verify: signature correct
debug2: kex_derive_keys
debug2: set_newkeys: mode 1
debug1: SSH2_MSG_NEWKEYS sent
debug1: expecting SSH2_MSG_NEWKEYS
debug2: set_newkeys: mode 0
debug1: SSH2_MSG_NEWKEYS received
debug1: Roaming not allowed by server
debug1: SSH2_MSG_SERVICE_REQUEST sent
debug2: service_accept: ssh-userauth
debug1: SSH2_MSG_SERVICE_ACCEPT received
debug2: key: boguspem.pem ((nil))
debug1: Authentications that can continue: publickey
debug3: start over, passed a different list publickey
debug3: preferred gssapi-keyex,gssapi-with-mic,publickey,keyboard-interactive,password
debug3: authmethod_lookup publickey
```

```
debug3: remaining preferred: keyboard-interactive,password
debug3: authmethod_is_enabled publickey
debug1: Next authentication method: publickey
debug1: Trying private key: boguspem.pem
debug1: read PEM private key done: type RSA
debug3: sign_and_send_pubkey: RSA
9c:4c:bc:0c:d0:5c:c7:92:6c:8e:9b:16:e4:43:d8:b2
debug2: we sent a publickey packet, wait for reply
debug1: Authentications that can continue: publickey
debug2: we did not send a packet, disable method
debug1: No more authentication methods to try.
Permission denied (publickey).
```

### If you use SSH (MindTerm) to connect to your instance

- If Java is not enabled, the server does not recognize the user key. To enable Java, see [How do I enable Java in my web browser?](#) in the Java documentation.

### If you use PuTTY to connect to your instance

- Verify that your private key (.pem) file has been converted to the format recognized by PuTTY (.ppk). For more information about converting your private key, see [Connecting to Linux/UNIX Instances from Windows Using PuTTY \(p. 274\)](#).

#### Note

In PuTTYgen, load your private key file and select **Save Private Key** rather than **Generate**.

- Verify that you are connecting with the appropriate user name for your AMI. Enter the user name in the **Host name** box in the **PuTTY Configuration** window.
  - For an Amazon Linux AMI, the user name is `ec2-user`.
  - For a RHEL5 AMI, the user name is often `root` but might be `ec2-user`.
  - For an Ubuntu AMI, the user name is `ubuntu`.
  - Otherwise, check with your AMI provider.
- Verify that you have an inbound security group rule to allow inbound traffic to the appropriate port. For more information, see [Authorizing Network Access to Your Instances \(p. 412\)](#).

## Error: Host key not found or Authentication failed, permission denied

If you use MindTerm to connect to your instance and get either of the following errors, `Host key not found in [directory]` or `Authentication failed, permission denied`, verify that you are connecting with the appropriate user name for your AMI. Enter the user name in the **User name** box in the **Connect To Your Instance** window.

The appropriate user names are as follows:

- For an Amazon Linux AMI, the user name is `ec2-user`.
- For a RHEL5 AMI, the user name is often `root` but might be `ec2-user`.
- For an Ubuntu AMI, the user name is `ubuntu`.
- Otherwise, check with your AMI provider.



## Error: Server refused our key or No supported authentication methods available

If you use PuTTY to connect to your instance and get either of the following errors, `Error: Server refused our key` or `Error: No supported authentication methods available`, verify that you are connecting with the appropriate user name for your AMI. Enter the user name in the **User name** box in the **PuTTY Configuration** window.

The appropriate user names are as follows:

- For an Amazon Linux AMI, the user name is `ec2-user`.
- For a RHEL5 AMI, the user name is often `root` but might be `ec2-user`.
- For an Ubuntu AMI, the user name is `ubuntu`.
- Otherwise, check with your AMI provider.

## Error using MindTerm on Safari Browser

If you use MindTerm to connect to your instance, and are using the Safari 6.1 or 7 web browser, you may get the following error:

```
Error connecting to your_instance_ip, reason:  
-> Key exchange failed: Host authentication failed
```

You need to update the browser's security settings to allow the AWS Management Console to run the Java plugin in unsafe mode.

### To enable the Java plugin to run in unsafe mode

1. In Safari, keep the Amazon EC2 console open, and select **Safari**, then **Preferences**, then **Security**.
2. Click **Manage Website Settings**.
3. Select the **Java** plugin on the left, then locate the AWS Management Console URL in the **Currently Open Websites** list. Select **Run in Unsafe Mode** from its associated list.
4. When prompted, click **Trust** in the warning dialog. Click **Done** to return the browser.

## Troubleshooting Stopping Your Instance

If you have stopped your Amazon EBS-backed instance and it appears "stuck" in the `stopping` state, there may be an issue with the underlying host computer.

First, try stopping the instance again. If you are using the Amazon EC2 command line interface tools, be sure to use the `--force` option, as shown in the following example:

```
ec2-stop-instances your_instance_id --force
```

If you can't force the instance to stop, you can create an AMI from the instance and launch a replacement instance.

### To create a replacement instance

1. Open the Amazon EC2 console.

2. In the navigation pane, click **Instances** and select the instance.
3. Click **Actions** and then click **Create Image**.
4. In the **Create Image** dialog box, fill in the following fields and then click **Create Image**:
  - a. Specify a name and description for the AMI.
  - b. Select **No reboot**.
5. Launch an instance from the AMI and verify that the instance is working.
6. Select the stuck instance, click **Actions**, and then click **Terminate**. If the instance also gets stuck terminating, Amazon EC2 automatically forces it to terminate within a few hours.

If you are unable to create an AMI from the instance as described in the previous procedure, you can set up a replacement instance as follows:

#### To create a replacement instance (if the previous procedure fails)

1. Select the instance, open the **Description** tab, and view the **Block devices** list. Select each volume and write down its volume ID. Be sure to note which volume is the root volume.
2. In the navigation pane, click **Volumes**. Select each volume for the instance, click **Actions**, and then click **Create Snapshot**.
3. In the navigation pane, click **Snapshots**. Select each snapshot that you just created and click **Create Volume**.
4. Launch an instance of the same type as the stuck instance (Amazon Linux, Windows, and so on). Note the volume ID and device name of its root volume.
5. In the navigation pane, click **Instances**, select the instance that you just launched, and then click **Stop**.
6. In the navigation pane, click **Volumes**, select the root volume of the stopped instance, click **Actions**, and then click **Detach Volume**.
7. Select the root volume that you created from the stuck instance, click **Attach Volume**, and attach it to the new instance as its root volume (using the device name that you wrote down). Attach any additional non-root volumes to the instance.
8. In the navigation pane, click **Instances** and select the replacement instance. Click **Actions** and then click **Start**. Verify that the instance is working.
9. Select the stuck instance, click **Actions**, and then click **Terminate**. If the instance also gets stuck terminating, Amazon EC2 automatically forces it to terminate within a few hours.

If you're unable to complete these procedures, you can post a request for help to the [Amazon EC2 forum](#). To help expedite a resolution, include the instance ID and describe the steps that you've already taken.

## Troubleshooting Terminating (Shutting Down) Your Instance

If your instance remains in the `shutting-down` state for more than a few minutes, it might be stuck in this state. A common cause is a problem on the underlying host computer. After a few hours, Amazon EC2 terminates your stuck instance. You are not billed for any instance hours while an instance is not in the `running` state. In other words, as soon as the state of an instance changes to `shutting-down`, you stop incurring charges for that instance.

If you are unable to wait for your instance to terminate, you can post a request for help to the [Amazon EC2 forum](#). To help expedite a resolution, include the instance ID and describe the steps that you've already taken.

## Automatically Launch or Terminate Instances

If you terminate all your instances, you may see that we launch a new instance for you. If you launch an instance, you may see that we terminate one of your instances. Generally, these behaviors mean that you've used Auto Scaling or AWS Elastic Beanstalk to automatically scale your computing resources based on criteria that you've defined.

For more information, see the [Auto Scaling Developer Guide](#) or the [AWS Elastic Beanstalk Developer Guide](#).

## Running Scripts on Instance Termination

If you run a script on instance termination, you may experience an abnormal termination because there is no way to ensure that shutdown scripts run. Amazon EC2 tries to shut an instance down cleanly and run system shutdown scripts, but certain events (such as hardware failure) may prevent system shutdown scripts from running.

## Troubleshooting Instances with Failed Status Checks

### Topics

- [Initial Steps You Can Take](#) (p. 360)
- [Troubleshooting Instance Status Checks for Linux-Based Instances](#) (p. 361)

## Initial Steps You Can Take

If your instance fails a status check, first determine whether your applications are exhibiting any problems. If you verify that the instance is not running your applications as expected, follow these steps:

### To investigate impaired instances using the AWS Management Console

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Instances**, and then select your instance.
3. Click the **Status Checks** tab in the details pane to see the individual results for all **System Status Checks** and **Instance Status Checks**.

If a system status check has failed, you can try one of the following options:

- For an instance using an Amazon EBS-backed AMI, stop and re-start the instance.
- For an instance using an instance-store backed AMI, terminate the instance and launch a replacement.
- Wait for Amazon EC2 to resolve the issue.
- Post your issue to the [Amazon EC2 forum](#).

If an instance status check fails, follow these steps:

1. Right-click your instance, and then click **Reboot**. It may take a few minutes for your system to restart.
2. Verify that the problem still exists; in some cases, rebooting may resolve the problem.
3. Wait until the instance shows a `running` state.
4. Right-click the instance and then click **Get System Log**. You can use this information to help identify the problem. Be sure that you rebooted recently to clear unnecessary information from the log.

5. Review the log that appears on the screen.
6. Use the list of known system log error statements below to troubleshoot your issue.
7. If your experience differs from the our check results, or if you are having an issue with your instance that our checks did not detect, click **Submit feedback** at the bottom of the **Status Checks** tab to help us improve our detection tests.
8. If your issue is not resolved, you can post your issue to the [Amazon EC2 forum](#).

## Troubleshooting Instance Status Checks for Linux-Based Instances

For Linux-based instances that have failed an instance status check, such as the instance reachability check, verify that you followed the steps discussed earlier to retrieve the system log. The following list contains some common system log errors and suggested actions you can take to resolve the issue for each error .

### Memory Errors

- [Out of memory: kill process](#) (p. 362)
- [ERROR: mmu\\_update failed \(Memory management update failed\)](#) (p. 362)

### Device Errors

- [I/O error \(Block device failure\)](#) (p. 363)
- [IO ERROR: neither local nor remote disk \(Broken distributed block device\)](#) (p. 364)

### Kernel Errors

- [request\\_module: runaway loop modprobe \(Looping legacy kernel modprobe on older Linux versions\)](#) (p. 365)
- ["FATAL: kernel too old" and "fsck: No such file or directory while trying to open /dev" \(Kernel and AMI mismatch\)](#) (p. 366)
- ["FATAL: Could not load /lib/modules" or "BusyBox" \(Missing kernel modules\)](#) (p. 367)
- [ERROR Invalid kernel \(EC2 incompatible kernel\)](#) (p. 368)

### File System Errors

- [request\\_module: runaway loop modprobe \(Looping legacy kernel modprobe on older Linux versions\)](#) (p. 369)
- [fsck: No such file or directory while trying to open... \(File system not found\)](#) (p. 370)
- [General error mounting filesystems \(Failed mount\)](#) (p. 372)
- [VFS: Unable to mount root fs on unknown-block \(Root filesystem mismatch\)](#) (p. 374)
- [Error: Unable to determine major/minor number of root device... \(Root file system/device mismatch\)](#) (p. 375)
- [XENBUS: Device with no driver...](#) (p. 376)
- [... days without being checked, check forced \(File system check required\)](#) (p. 377)
- [fsck died with exit status... \(Missing device\)](#) (p. 378)

### Operating System Errors

- [GRUB prompt \(grubdom>\)](#) (p. 379)

- Bringing up interface eth0: Device eth0 has different MAC address than expected, ignoring. (Hard-coded MAC address) (p. 380)
- Unable to load SELinux Policy. Machine is in enforcing mode. Halting now. (SELinux misconfiguration) (p. 381)
- XENBUS: Timeout connecting to devices (Xenbus timeout) (p. 382)

## Out of memory: kill process

An out of memory error is indicated by a system log entry similar to the one shown below:

```
[115879.769795] Out of memory: kill process 20273 (httpd) score 1285879  
or a child  
[115879.769795] Killed process 1917 (php-cgi) vsz:467184kB, anon-  
rss:101196kB, file-rss:204kB
```

### Potential Cause

Exhausted memory

### Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Do one of the following: <ul style="list-style-type: none"><li>• Stop the instance, and modify the instance to use a different instance type, and start the instance again. For example, a larger or a memory-optimized instance type.</li><li>• Reboot the instance to return it to a non-impaired status. The problem will probably occur again unless you change the instance type.</li></ul>
Instance store-backed	Do one of the following: <ul style="list-style-type: none"><li>• Terminate the instance and launch a new instance, specifying a different instance type. For example, a larger or a memory-optimized instance type.</li><li>• Reboot the instance to return it to an unimpaired status. The problem will probably occur again unless you change the instance type.</li></ul>

## ERROR: mmu\_update failed (Memory management update failed)

Memory management update failures are indicated by a system log entry similar to the following:

```
...
Press `ESC' to enter the menu... 0 [H[J Booting 'Amazon Linux 2011.09
(2.6.35.14-95.38.amzn1.i686) '

root (hd0)

Filesystem type is ext2fs, using whole disk

kernel /boot/vmlinuz-2.6.35.14-95.38.amzn1.i686 root=LABEL=/ console=hvc0 LANG=
en_US.UTF-8 KEYTABLE=us

initrd /boot/initramfs-2.6.35.14-95.38.amzn1.i686.img

ERROR: mmu_update failed with rc=-22
```

## Potential Cause

Issue with Amazon Linux.

## Suggested Action

Seek assistance by posting your issue to the [Developer Forums](#) or contacting [AWS Support](#).

## I/O error (Block device failure)

An input/output error is indicated by a system log entry similar to the following example:

```
[9943662.053217] end_request: I/O error, dev sde, sector 52428288
[9943664.191262] end_request: I/O error, dev sde, sector 52428168
[9943664.191285] Buffer I/O error on device md0, logical block 209713024
[9943664.191297] Buffer I/O error on device md0, logical block 209713025
[9943664.191304] Buffer I/O error on device md0, logical block 209713026
[9943664.191310] Buffer I/O error on device md0, logical block 209713027
[9943664.191317] Buffer I/O error on device md0, logical block 209713028
[9943664.191324] Buffer I/O error on device md0, logical block 209713029
[9943664.191332] Buffer I/O error on device md0, logical block 209713030
[9943664.191339] Buffer I/O error on device md0, logical block 209713031
[9943664.191581] end_request: I/O error, dev sde, sector 52428280
[9943664.191590] Buffer I/O error on device md0, logical block 209713136
[9943664.191597] Buffer I/O error on device md0, logical block 209713137
[9943664.191767] end_request: I/O error, dev sde, sector 52428288
[9943664.191970] end_request: I/O error, dev sde, sector 52428288
[9943664.192143] end_request: I/O error, dev sde, sector 52428288
[9943664.192949] end_request: I/O error, dev sde, sector 52428288
[9943664.193112] end_request: I/O error, dev sde, sector 52428288
[9943664.193266] end_request: I/O error, dev sde, sector 52428288
...
```

## Potential Causes

Instance type	Potential cause
Amazon EBS-backed	A failed Amazon EBS volume

Instance type	Potential cause
Instance store-backed	A failed physical drive

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	<p>Use the following procedure:</p> <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Detach the volume.</li><li>3. Attempt to recover the volume.</li></ol> <p><b>Tip</b> It's good practice to snapshot your Amazon EBS volumes often. This dramatically decreases the risk of data loss as a result of failure.</p> <ol style="list-style-type: none"><li>4. Re-attach the volume to the instance.</li><li>5. Detach the volume.</li></ol>
Instance store-backed	<p>Terminate the instance and launch a new instance.</p> <p><b>Note</b> Data cannot be recovered. Recover from backups</p> <p><b>Tip</b> It's a good practice to use either Amazon S3 or Amazon EBS for backups. Instance store volumes are directly tied to single host and single disk failures.</p>

## IO ERROR: neither local nor remote disk (Broken distributed block device)

An input/output error on the device is indicated by a system log entry similar to the following example:

```
...
block drbd1: Local IO failed in request_timer_fn. Detaching...

Aborting journal on device drbd1-8.

block drbd1: IO ERROR: neither local nor remote disk

Buffer I/O error on device drbd1, logical block 557056

lost page write due to I/O error on drbd1
```

```
JBD2: I/O error detected when updating journal superblock for drbd1-8.
```

## Potential Causes

Instance type	Potential cause
Amazon EBS-backed	A failed Amazon EBS volume
Instance store-backed	A failed physical drive

## Suggested Action

Terminate the instance and launch a new instance.

For an Amazon EBS-backed instance you can recover data from a recent snapshot by creating an image from it. Any data added after the snapshot cannot be recovered.

## request\_module: runaway loop modprobe (Looping legacy kernel modprobe on older Linux versions)

This condition is indicated by a system log similar to the one shown below. Using an unstable or old Linux kernel (for example, 2.6.16-xenU) can cause an interminable loop condition at startup.

```
Linux version 2.6.16-xenU (builder@xenbat.amazonsa) (gcc version 4.0.1  
20050727 (Red Hat 4.0.1-5)) #1 SMP Mon May 28 03:41:49 SAST 2007
```

```
BIOS-provided physical RAM map:
```

```
Xen: 0000000000000000 - 0000000026700000 (usable)
```

```
OMB HIGHMEM available.
```

```
...
```

```
request_module: runaway loop modprobe binfmt-464c
```

```
request_module: runaway loop modprobe binfmt-464c
```

```
request_module: runaway loop modprobe binfmt-464c
```

```
request_module: runaway loop modprobe binfmt-464c
```

```
request_module: runaway loop modprobe binfmt-464c
```



## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	<p>Use a newer kernel, either Grub-based or static, using one of the following options.</p> <p>Option 1: Terminate the instance and launch a new instance, specifying the <code>-kernel</code> and <code>-ramdisk</code> parameters.</p> <p>Option 2:</p> <ol style="list-style-type: none"> <li>1. Stop the instance.</li> <li>2. Modify the kernel and ramdisk attributes to use a newer kernel.</li> <li>3. Start the instance.</li> </ol>
Instance store-backed	<p>Terminate the instance and launch a new instance, specifying the <code>-kernel</code> and <code>-ramdisk</code> parameters.</p>

## "FATAL: kernel too old" and "fsck: No such file or directory while trying to open /dev" (Kernel and AMI mismatch)

This condition is indicated by a system log similar to the one shown below:

```
Linux version 2.6.16.33-xenU (root@dom0-0-50-45-1-a4-ee.z-2.aes0.internal)
(gcc version 4.1.1 20070105 (Red Hat 4.1.1-52)) #2 SMP Wed Aug 15 17:27:36 SAST
2007
...
FATAL: kernel too old
Kernel panic - not syncing: Attempted to kill init!
```

## Potential Causes

Incompatible kernel and userland.

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	<p>Use the following procedure:</p> <ol style="list-style-type: none"> <li>1. Stop the instance.</li> <li>2. Modify the configuration to use a newer kernel.</li> <li>3. Start the instance.</li> </ol>

For this instance type	Do this
Instance store-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Create an AMI that uses a newer kernel.</li><li>2. Terminate the instance.</li><li>3. Start a new instance from the AMI you created.</li></ol>

## "FATAL: Could not load /lib/modules" or "BusyBox" (Missing kernel modules)

This condition is indicated by a system log similar to the one shown below.

```
[    0.370415] Freeing unused kernel memory: 1716k freed
Loading, please wait...
WARNING: Couldn't open directory /lib/modules/2.6.34-4-virtual: No such file
or directory
FATAL: Could not open /lib/modules/2.6.34-4-virtual/modules.dep.temp for writing:
  No such file or directory
FATAL: Could not load /lib/modules/2.6.34-4-virtual/modules.dep: No such file
or directory
Couldn't get a file descriptor referring to the console
Begin: Loading essential drivers... ...
FATAL: Could not load /lib/modules/2.6.34-4-virtual/modules.dep: No such file
or directory
FATAL: Could not load /lib/modules/2.6.34-4-virtual/modules.dep: No such file
or directory
Done.
Begin: Running /scripts/init-premount ...
Done.
Begin: Mounting root file system... ...
Begin: Running /scripts/local-top ...
Done.
Begin: Waiting for root file system... ...
Done.
Gave up waiting for root device. Common problems:
- Boot args (cat /proc/cmdline)
  - Check rootdelay= (did the system wait long enough?)
  - Check root= (did the system wait for the right device?)
- Missing modules (cat /proc/modules; ls /dev)
FATAL: Could not load /lib/modules/2.6.34-4-virtual/modules.dep: No such file
or directory
FATAL: Could not load /lib/modules/2.6.34-4-virtual/modules.dep: No such file
or directory
ALERT! /dev/sda1 does not exist. Dropping to a shell!

BusyBox v1.13.3 (Ubuntu 1:1.13.3-lubuntu5) built-in shell (ash)
Enter 'help' for a list of built-in commands.

(initramfs)
```

## Potential Causes

One or more of the following conditions can cause this problem:

- Missing ramdisk
- Missing correct modules from ramdisk
- EBS root volume not correctly attached as /dev/sda1

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Select corrected ramdisk for the EBS volume.</li><li>2. Stop the instance.</li><li>3. Detach the volume and repair it.</li><li>4. Attach the volume to the instance.</li><li>5. Start the instance.</li><li>6. Modify the AMI to use the corrected ramdisk.</li></ol>
Instance store-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Terminate the instance and launch a new instance with the correct ramdisk.</li><li>2. Create a new AMI with the correct ramdisk.</li></ol>

## ERROR Invalid kernel (EC2 incompatible kernel)

This condition is indicated by a system log similar to the one shown below.

```
...
root (hd0)

Filesystem type is ext2fs, using whole disk

kernel /vmlinuz root=/dev/sda1 ro

initrd /initrd.img

ERROR Invalid kernel: elf_xen_note_check: ERROR: Will only load images
built for the generic loader or Linux images
xc_dom_parse_image returned -1

Error 9: Unknown boot failure

Booting 'Fallback'
```

```
root (hd0)

Filesystem type is ext2fs, using whole disk

kernel /vmlinuz.old root=/dev/sda1 ro

Error 15: File not found
```

## Potential Causes

One or both of the following conditions can cause this problem:

- Supplied kernel is not supported by Grub.
- Fallback kernel does not exist.

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Replace with working kernel.</li><li>3. Install a fallback kernel.</li><li>4. Modify the AMI by correcting the kernel.</li></ol>
Instance store-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Terminate the instance and launch a new instance with the correct kernel.</li><li>2. Create an AMI with the correct kernel.</li><li>3. (Optional) Seek technical assistance for data recovery using <a href="#">AWS Support</a>.</li></ol>

## request\_module: runaway loop modprobe (Looping legacy kernel modprobe on older Linux versions)

This condition is indicated by a system log similar to the one shown below. Using an unstable or old Linux kernel (for example, 2.6.16-xenU) can cause an interminable loop condition at startup.

```
Linux version 2.6.16-xenU (builder@xenbat.amazonsa) (gcc version 4.0.1
20050727 (Red Hat 4.0.1-5)) #1 SMP Mon May 28 03:41:49 SAST 2007

BIOS-provided physical RAM map:

Xen: 0000000000000000 - 0000000026700000 (usable)
```

```
OMB HIGHMEM available.  
...  
request_module: runaway loop modprobe binfmt-464c  
request_module: runaway loop modprobe binfmt-464c  
request_module: runaway loop modprobe binfmt-464c  
request_module: runaway loop modprobe binfmt-464c  
request_module: runaway loop modprobe binfmt-464c
```

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Use a newer kernel, either Grub-based or static, using one of the following options.  Option 1: Terminate the instance and launch a new instance, specifying the <code>-kernel</code> and <code>-ramdisk</code> parameters.  Option 2: <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Modify the kernel and ramdisk attributes to use a newer kernel.</li><li>3. Start the instance.</li></ol>
Instance store-backed	Terminate the instance and launch a new instance, specifying the <code>-kernel</code> and <code>-ramdisk</code> parameters.

## fsck: No such file or directory while trying to open... (File system not found)

This condition is indicated by a system log similar to the one shown below:

```
Welcome to Fedora  
Press 'I' to enter interactive startup.  
Setting clock : Wed Oct 26 05:52:05 EDT 2011 [ OK ]  
  
Starting udev: [ OK ]  
  
Setting hostname localhost: [ OK ]  
  
No devices found  
Setting up Logical Volume Management: File descriptor 7 left open  
No volume groups found  
[ OK ]
```

```
Checking filesystems
Checking all file systems.
[/sbin/fsck.ext3 (1) -- /] fsck.ext3 -a /dev/sda1
/dev/sda1: clean, 82081/1310720 files, 2141116/2621440 blocks
[/sbin/fsck.ext3 (1) -- /mnt/dbbackups] fsck.ext3 -a /dev/sdh
fsck.ext3: No such file or directory while trying to open /dev/sdh

/dev/sdh:
The superblock could not be read or does not describe a correct ext2
filesystem.  If the device is valid and it really contains an ext2
filesystem (and not swap or ufs or something else), then the superblock
is corrupt, and you might try running e2fsck with an alternate superblock:
    e2fsck -b 8193 <device>

[FAILED]

*** An error occurred during the file system check.
*** Dropping you to a shell; the system will reboot
*** when you leave the shell.
Give root password for maintenance
(or type Control-D to continue):
```

## Potential Causes

- A bug exists in ramdisk filesystem definitions /etc/fstab
- Misconfigured filesystem definitions in /etc/fstab
- Missing/failed drive

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	<p>Use the following procedure:</p> <ol style="list-style-type: none"><li>1. Stop the instance, detach the root volume, repair/modify <code>/etc/fstab</code> the volume, attach the volume to the instance, and start the instance.</li><li>2. Fix ramdisk to include modified <code>/etc/fstab</code> (if applicable).</li><li>3. Modify the AMI to use a newer ramdisk.</li></ol> <p><b>Tip</b> The sixth field in the <code>fstab</code> defines availability requirements of the mount – a nonzero value implies that an <code>fsck</code> will be done on that volume and <i>must</i> succeed. Using this field can be problematic in Amazon EC2 because a failure typically results in an interactive console prompt which is not currently available in Amazon EC2. Use care with this feature and read the Linux man page for <code>fstab</code>.</p>
Instance store-backed	<p>Use the following procedure:</p> <ol style="list-style-type: none"><li>1. Terminate the instance and launch a new instance.</li><li>2. Detach any errant EBS volumes and the reboot instance.</li><li>3. (Optional) Seek technical assistance for data recovery using <a href="#">AWS Support</a>.</li></ol>

## General error mounting filesystems (Failed mount)

This condition is indicated by a system log similar to the one shown below.

```
Loading xenblk.ko module
xen-vbd: registered block device major 8

Loading ehci-hcd.ko module
Loading ohci-hcd.ko module
Loading uhci-hcd.ko module
USB Universal Host Controller Interface driver v3.0

Loading mbcache.ko module
Loading jbd.ko module
Loading ext3.ko module
Creating root device.
Mounting root filesystem.
```

```
kjournald starting. Commit interval 5 seconds

EXT3-fs: mounted filesystem with ordered data mode.

Setting up other filesystems.
Setting up new root fs
no fstab.sys, mounting internal defaults
Switching to new root and running init.
unmounting old /dev
unmounting old /proc
unmounting old /sys
mountall:/proc: unable to mount: Device or resource busy
mountall:/proc/self/mountinfo: No such file or directory
mountall: root filesystem isn't mounted
init: mountall main process (221) terminated with status 1

General error mounting filesystems.
A maintenance shell will now be started.
CONTROL-D will terminate this shell and re-try.
Press enter for maintenance
(or type Control-D to continue):
```

## Potential Causes

Instance type	Potential cause
Amazon EBS-backed	<ul style="list-style-type: none"><li>• Detached or failed EBS volume.</li><li>• Corrupted filesystem.</li><li>• Mismatched ramdisk and AMI combination (e.g., Debian ramdisk with a SUSE AMI).</li></ul>
Instance store-backed	<ul style="list-style-type: none"><li>• A failed drive.</li><li>• A corrupted file system.</li><li>• A mismatched ramdisk and combination (for example, a Debian ramdisk with a SUSE AMI).</li></ul>



## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Detach the root volume.</li><li>3. Attach the root volume to a known working instance.</li><li>4. Run filesystem check (<code>fsck -a /dev/...</code>).</li><li>5. Fix any errors.</li><li>6. Detach the volume from the known working instance.</li><li>7. Attach the volume to the stopped instance.</li><li>8. Start the instance.</li><li>9. Recheck the instance status.</li></ol>
Instance store-backed	Try one of the following: <ul style="list-style-type: none"><li>• Restart a new instance.</li><li>• (Optional) Seek technical assistance for data recovery using <a href="#">AWS Support</a>.</li></ul>

## VFS: Unable to mount root fs on unknown-block (Root filesystem mismatch)

This condition is indicated by a system log similar to the one shown below.

```
Linux version 2.6.16-xenU (builder@xenbat.amazonsa) (gcc version 4.0.1
20050727 (Red Hat 4.0.1-5)) #1 SMP Mon May 28 03:41:49 SAST 2007
...
Kernel command line: root=/dev/sda1 ro 4
...
Registering block device major 8
...
Kernel panic - not syncing: VFS: Unable to mount root fs on unknown-block(8,1)
```

## Potential Causes

Instance type	Potential cause
Amazon EBS-backed	<ul style="list-style-type: none"><li>• Device not attached correctly.</li><li>• Root device not attached at correct device point.</li><li>• Filesystem not expected format.</li><li>• Use of legacy kernel (e.g., 2.6.16-XenU).</li></ul>
Instance store-backed	Hardware device failure.

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Do one of the following: <ul style="list-style-type: none"><li>• Stop and then restart the instance.</li><li>• Modify root volume to attach at the correct device point, possible /dev/sda1 instead of /dev/sda.</li><li>• Stop and modify to use modern kernel.</li></ul>
Instance store-backed	Terminate the instance and launch a new instance using a modern kernel.

## Error: Unable to determine major/minor number of root device... (Root file system/device mismatch)

This condition is indicated by a system log similar to the one shown below.

```
...
XENBUS: Device with no driver: device/vif/0
XENBUS: Device with no driver: device/vbd/2048
drivers/rtc/hctosys.c: unable to open rtc device (rtc0)
Initializing network drop monitor service
Freeing unused kernel memory: 508k freed
:: Starting udevd...
done.
:: Running Hook [udev]
:: Triggering uevents...<30>udev[65]: starting version 173
done.
Waiting 10 seconds for device /dev/xvda1 ...
Root device '/dev/xvda1' doesn't exist. Attempting to create it.
ERROR: Unable to determine major/minor number of root device '/dev/xvda1'.
You are being dropped to a recovery shell
Type 'exit' to try and continue booting
sh: can't access tty; job control turned off
[ramfs /]#
```

## Potential Causes

- Missing or incorrectly configured virtual block device driver
- Device enumeration clash (sda versus xvda or sda instead of sda1)
- Incorrect choice of DomU kernel

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Detach the volume.</li><li>3. Fix the device mapping problem.</li><li>4. Start the instance.</li><li>5. Modify the AMI to address device mapping issues.</li></ol>
Instance store-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Create a new AMI with the appropriate fix (map block device correctly).</li><li>2. Terminate the instance and launch a new instance from the AMI you created.</li></ol>

## XENBUS: Device with no driver...

This condition is indicated by a system log similar to the one shown below.

```
XENBUS: Device with no driver: device/vbd/2048
drivers rtc/hctosys.c: unable to open rtc device (rtc0)
Initializing network drop monitor service
Freeing unused kernel memory: 508k freed
:: Starting udevd...
done.
:: Running Hook [udev]
:: Triggering uevents...<30>udev[65]: starting version 173
done.
Waiting 10 seconds for device /dev/xvda1 ...
Root device '/dev/xvda1' doesn't exist. Attempting to create it.
ERROR: Unable to determine major/minor number of root device '/dev/xvda1'.
You are being dropped to a recovery shell
Type 'exit' to try and continue booting
sh: can't access tty; job control turned off
[ramfs /]#
```

## Potential Causes

- Missing or incorrectly configured virtual block device driver
- Device enumeration clash (sda versus xvda)
- Incorrect choice of DomU kernel

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Detach the volume.</li><li>3. Fix the device mapping problem.</li><li>4. Start the instance.</li><li>5. Modify the AMI to address device mapping issues.</li></ol>
Instance store-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Create an AMI with the appropriate fix (map block device correctly).</li><li>2. Terminate the instance and launch a new instance using the AMI you created.</li></ol>

## ... days without being checked, check forced (File system check required)

This condition is indicated by a system log similar to the one shown below.

```
...
Checking filesystems
Checking all file systems.
[/sbin/fsck.ext3 (1) -- /] fsck.ext3 -a /dev/sda1
/dev/sda1 has gone 361 days without being checked, check forced
```

## Potential Causes

Filesystem check time passed; a filesystem check is being forced.

## Suggested Actions

- Wait until the filesystem check completes. Note that a filesystem check can take a long time depending on the size of the root filesystem.

- Modify your filesystems to remove the filesystem check (fsck) enforcement using tune2fs or tools appropriate for your filesystem.

## fsck died with exit status... (Missing device)

This condition is indicated by a system log similar to the one shown below.

```
Cleaning up ifupdown...
Loading kernel modules...done.
...
Activating lvm and md swap...done.
Checking file systems...fsck from util-linux-ng 2.16.2
/sbin/fsck.xfs: /dev/sdh does not exist
fsck died with exit status 8
[31mfailed (code 8).[39;49m
```

### Potential Causes

- Ramdisk looking for missing drive
- Filesystem consistency check forced
- Drive failed or detached

### Suggested Actions

For this instance type	Do this
Amazon EBS-backed	<p>Try one or more of the following to resolve the issue:</p> <ul style="list-style-type: none"><li>• Stop the instance, attach volume to existing running instance.</li><li>• Manually run consistency checks.</li><li>• Fix ramdisk to include relevant utilities.</li><li>• Modify filesystem tuning parameters to remove consistency requirements (not recommended).</li></ul>
Instance store-backed	<p>Try one or more of the following to resolve the issue:</p> <ul style="list-style-type: none"><li>• Rebundle ramdisk with correct tooling.</li><li>• Modify file system tuning parameters to remove consistency requirements (not recommended).</li><li>• Terminate the instance and launch a new instance.</li><li>• (Optional) Seek technical assistance for data recovery using <a href="#">AWS Support</a>.</li></ul>

## GRUB prompt (grubdom>)

This condition is indicated by a system log similar to the one shown below.

```
GNU GRUB version 0.97 (629760K lower / 0K upper memory)

[ Minimal BASH-like line editing is supported.  For
  the first word, TAB lists possible command
  completions.  Anywhere else TAB lists the possible
  completions of a device/filename. ]

grubdom>
```

### Potential Causes

Instance type	Potential causes
Amazon EBS-backed	<ul style="list-style-type: none"><li>• Missing grub.conf file.</li><li>• Incorrect Grub image used expecting grub.conf at different location.</li><li>• Unsupported filesystem used to store grub.conf.</li></ul>
Instance store-backed	<ul style="list-style-type: none"><li>• Missing grub.conf file.</li><li>• Incorrect Grub image used expecting grub.conf at different location.</li></ul>

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	<p>Option 1: Modify the AMI and relaunch the instance</p> <ol style="list-style-type: none"> <li>1. Modify the source AMI to create a grub.conf at the standard location (/boot/grub/menu.lst).</li> <li>2. Pick the appropriate Grub image, (hd0-1st drive or hd00 – 1st drive, 1st partition).</li> <li>3. Terminate the instance and launch a new one using the AMI the you created.</li> </ol> <p>Option 2: Fix the existing instance</p> <ol style="list-style-type: none"> <li>1. Stop the instance.</li> <li>2. Detach the root filesystem.</li> <li>3. Attach the root filesystem to a known working instance.</li> <li>4. Mount filesystem.</li> <li>5. Create grub.conf.</li> <li>6. Detach filesystem.</li> <li>7. Attach to the original instance.</li> <li>8. Modify kernel attribute to use appropriate Grub (1st disk or 1st partition on 1st disk).</li> <li>9. Start the instance.</li> </ol>
Instance store-backed	<p>Option 1: Modify the AMI and relaunch the instance</p> <ol style="list-style-type: none"> <li>1. Create the new AMI with a grub.conf at the standard location (/boot/grub/menu.lst).</li> <li>2. Pick the appropriate Grub image, (hd0-1st drive or hd00 – 1st drive, 1st partition).</li> <li>3. Terminate the instance and launch a new instance using the AMI you created.</li> </ol> <p>Option 2: Terminate the instance and launch a new instance, specifying the correct kernel.</p> <p><b>Note</b> To recover data from the existing instance, contact AWS Support.</p>

## Bringing up interface eth0: Device eth0 has different MAC address than expected, ignoring. (Hard-coded MAC address)

This condition is indicated by a system log similar to the one shown below:

```
...  
Bringing up loopback interface: [ OK ]  
  
Bringing up interface eth0: Device eth0 has different MAC address than expected,  
ignoring.  
[FAILED]  
  
Starting auditd: [ OK ]
```

## Potential Causes

There is a hardcoded interface MAC in the AMI configuration.

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	<p>Do one of the following:</p> <ul style="list-style-type: none"><li>• Modify the AMI to remove the hard coding and relaunch the instance.</li><li>• Modify the instance to remove the hard coded MAC address.</li></ul> <p>OR</p> <p>Use the following procedure:</p> <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Detach the root volume.</li><li>3. Attach the volume to another instance and modify the volume to remove the hard coded MAC address.</li><li>4. Attach the volume to original instance.</li><li>5. Start the instance.</li></ol>
Instance store-backed	<p>Do one of the following:</p> <ul style="list-style-type: none"><li>• Modify the instance to remove the hard-coded MAC address.</li><li>• Terminate the instance and launch a new instance.</li></ul>

## Unable to load SELinux Policy. Machine is in enforcing mode. Halting now. (SELinux misconfiguration)

This condition is indicated by a system log similar to the one shown below.



```
audit(1313445102.626:2): enforcing=1 old_enforcing=0 auid=4294967295
Unable to load SELinux Policy. Machine is in enforcing mode. Halting now.
Kernel panic - not syncing: Attempted to kill init!
```

## Potential Causes

SELinux has been enabled in error:

- Supplied kernel is not supported by Grub.
- Fallback kernel does not exist.

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Stop the instance.</li><li>2. Detach the volume.</li><li>3. Disable SELinux.</li><li>4. Start the instance.</li></ol>
Instance store-backed	Use the following procedure: <ol style="list-style-type: none"><li>1. Terminate the instance and launch a new instance.</li><li>2. (Optional) Seek technical assistance for data recovery using <a href="#">AWS Support</a>.</li></ol>

## XENBUS: Timeout connecting to devices (Xenbus timeout)

This condition is indicated by a system log similar to the one shown below.

```
Linux version 2.6.16-xenU (builder@xenbat.amazonsa) (gcc version 4.0.1
20050727 (Red Hat 4.0.1-5)) #1 SMP Mon May 28 03:41:49 SAST 2007
...
XENBUS: Timeout connecting to devices!
...
Kernel panic - not syncing: No init found. Try passing init= option to kernel.
```

## Potential Causes

- The block device not is connected to the instance.
- This instance is using a very old DomU kernel.

## Suggested Actions

For this instance type	Do this
Amazon EBS-backed	Do one of the following: <ul style="list-style-type: none"><li>• Modify AMI and instance to use a modern kernel and relaunch the instance.</li><li>• Reboot the instance.</li></ul>
Instance store-backed	Do one of the following: <ul style="list-style-type: none"><li>• Terminate the instance.</li><li>• Modify AMI to use a modern kernel and launch a new instance using this AMI.</li></ul>

## Troubleshooting Instance Capacity

The following errors are related to instance capacity.

### Error: `InsufficientInstanceCapacity`

If you get an `InsufficientInstanceCapacity` error when you try to launch an instance, AWS does not currently have enough available capacity to service your request. If you are requesting a large number of instances, there might not be enough server capacity to host them. You can try again later or specify a smaller number of instances.

### Error: `InstanceLimitExceeded`

If you get an `InstanceLimitExceeded` error when you try to launch an instance, you have reached your concurrent running instance limit. For new AWS accounts, the default limit is 20. If you need additional running instances, complete the form at [Request to Increase Amazon EC2 Instance Limit](#).

## Getting Console Output and Rebooting Instances

Console output is a valuable tool for problem diagnosis. It is especially useful for troubleshooting kernel problems and service configuration issues that could cause an instance to terminate or become unreachable before its SSH daemon can be started.

Similarly, the ability to reboot instances that are otherwise unreachable is valuable for both troubleshooting and general instance management.

Amazon EC2 instances do not have a physical monitor through which you can view their console output. They also lack physical controls that allow you to power up, reboot, or shut them down. To allow these actions, we support these tasks through the Amazon EC2 API and the command line interface tools (CLI).

### Instance Reboot

Just as you can reset a computer by pressing the reset button, you can reset Amazon EC2 instances using `RebootInstances`. For more information, see [Reboot Your Instance \(p. 286\)](#).

**Caution**

For Windows instances, this operation performs a hard reboot that might result in data corruption.

## Instance Console Output

For Linux/UNIX instances, the instance console output displays the exact console output that would normally be displayed on a physical monitor attached to a computer. This output is buffered because the instance produces it and then posts it to a store where the instance's owner can retrieve it.

For Windows instances, the instance console output displays the last three system event log errors.

The posted output is not continuously updated; only when it is likely to be of the most value. This includes shortly after instance boot, after reboot, and when the instance terminates.

**Note**

Only the most recent 64 KB of posted output is stored, which is available for at least 1 hour after the last posting.

Only the instance owner can access the console output. You can retrieve the console output for your instances using [ec2-get-console-output](#) or [GetConsoleOutput](#).

## Instance Recovery When its Host Computer Fails

If there is an unrecoverable issue with the hardware of an underlying host computer, AWS may schedule an instance stop event. You'll be notified of such an event ahead of time by email.

If you have an Amazon EBS-backed instance running on a host computer that fails, you can do the following to recover:

1. Back up any important data on your instance store volumes to Amazon EBS or Amazon S3.
2. Stop the instance.
3. Start the instance.
4. Restore any important data.
5. [EC2-Classic] If the instance had an associated Elastic IP address, you must reassociate it with the instance.

If you have an instance store-backed instance running on a host computer that fails, you can do the following to recover:

1. Create an AMI from the instance using the [ec2-bundle-vol](#) command.
2. Upload the image to Amazon S3 using the [ec2-upload-bundle](#) command.
3. Back up important data to Amazon EBS or Amazon S3.
4. Terminate the instance.
5. Launch a new instance from the AMI.
6. Restore any important data to the new instance.
7. [EC2-Classic] If the original instance had an associated Elastic IP address, you must associate it with the new instance.

For more information, see [Stop and Start Your Instance](#) (p. 284).

# Network and Security

---

This section describes key network and security features related to Amazon EC2.

## Topics

- [Amazon EC2 Key Pairs](#) (p. 385)
- [Amazon EC2 Security Groups](#) (p. 392)
- [Controlling Access to Amazon EC2 Resources](#) (p. 399)
- [Amazon EC2 and Amazon Virtual Private Cloud \(VPC\)](#) (p. 414)
- [Amazon EC2 Instance IP Addressing](#) (p. 419)
- [Elastic IP Addresses \(EIP\)](#) (p. 428)
- [Elastic Network Interfaces \(ENI\)](#) (p. 431)

You can launch an instance into one of two platforms: EC2-Classic or EC2-VPC. An instance that's launched into EC2-Classic or a default VPC is automatically assigned a public IP address. An instance that's launched into a nondefault VPC can be assigned a public IP address on launch. For more information about EC2-Classic and EC2-VPC, see [Supported Platforms](#).

Instances can fail or terminate for reasons outside of your control. If one fails and you launch a replacement instance, the replacement has a different public IP address than the original. However, if your application needs a static IP address, you can use an *Elastic IP address*.

You can use *security groups* to control who can access your instances. These are analogous to an inbound network firewall that enables you to specify the protocols, ports, and source IP ranges that are allowed to reach your instances. You can create multiple security groups and assign different rules to each group. You can then assign each instance to one or more security groups, and we use the rules to determine which traffic is allowed to reach the instance. You can configure a security group so that only specific IP addresses or specific security groups have access to the instance.

## Amazon EC2 Key Pairs

Amazon EC2 uses public-key cryptography to encrypt and decrypt login information. Public-key cryptography uses a public key to encrypt a piece of data, such as a password, then the recipient uses the private key to decrypt the data. The public and private keys are known as a *key pair*.

To log in to your instance, you must create a key pair, specify the name of the key pair when you launch the instance, and provide the private key when you connect to the instance. Linux/UNIX instances have

no password, and you use a key pair to log in using SSH. With Windows instances, you use a key pair to obtain the administrator password and then log in using RDP.

### Creating a Key Pair

You can use Amazon EC2 to create your key pair. For more information, see [Creating Your Key Pair Using Amazon EC2 \(p. 386\)](#).

Alternatively, you could use a third-party tool and then import the public key to Amazon EC2. For more information, see [Importing Your Own Key Pair to Amazon EC2 \(p. 387\)](#).

Each key pair requires a name. Be sure to choose a name that is easy to remember. Amazon EC2 associates the public key with the name that you specify as the key name.

Amazon EC2 stores the public key only, and you store the private key. Anyone who possesses your private key can decrypt your login information, so it's important that you store your private keys in a secure place.

The keys that Amazon EC2 uses are 1024-bit SSH-2 RSA keys. You can have up to five thousand key pairs per region.

### Launching and Connecting to Your Instance

When you launch an instance, you should specify the name of the key pair you plan to use to connect to the instance. If you don't specify the name of an existing key pair when you launch an instance, you won't be able to connect to the instance. When you connect to the instance, you must specify the private key that corresponds to the key pair you specified when you launched the instance. Amazon EC2 doesn't keep a copy of your private key; therefore, if you lose your private key, there is no way to recover it. If you lose the private key for an instance store-backed instance, you can't access the instance; you should terminate the instance and launch another instance using a new key pair. If you lose the private key for an EBS-backed instance, you can regain access to your instance. For more information, see [Connecting to Your Instance if You Lose Your Private Key \(p. 390\)](#).

#### Topics


- [Creating Your Key Pair Using Amazon EC2 \(p. 386\)](#)
- [Importing Your Own Key Pair to Amazon EC2 \(p. 387\)](#)
- [Retrieving the Public Key for Your Key Pair \(p. 389\)](#)
- [Connecting to Your Instance if You Lose Your Private Key \(p. 390\)](#)

## Creating Your Key Pair Using Amazon EC2

Use the following steps to create a key pair using the Amazon EC2 console.

### To have Amazon EC2 create your key pair

1. Open the Amazon EC2 console.
2. From the navigation bar, select a region for the key pair. You can select any region that's available to you, regardless of your location. This choice is important because some Amazon EC2 resources can be shared between regions, but key pairs can't. For example, if you create a key pair in the US West (Oregon) Region, you can't see or use the key pair in another region.



3. Click **Key Pairs** in the navigation pane.
4. Click **Create Key Pair**.
5. Enter a name for the new key pair in the **Key pair name** field of the **Create Key Pair** dialog box, and then click **Create**.
6. The private key file is automatically downloaded by your browser. The base file name is the name you specified as the name of your key pair, and the file name extension is `.pem`. Save the private key file in a safe place.

#### Important

This is the only chance for you to save the private key file. You'll need to provide the name of your key pair when you launch an instance and the corresponding private key each time you connect to the instance.

7. If you will use an SSH client on a Linux computer to connect to your Linux instance, use the following command to set the permissions of your private key file so that only you can read it.

```
$ chmod 400 my-key-pair.pem
```

Alternatively, you can have Amazon EC2 create your key pair using the [ec2-create-keypair](#) command or the [CreateKeyPair](#) action.

## Importing Your Own Key Pair to Amazon EC2

If you used Amazon EC2 to create your key pair, as described in the previous section, you are ready to launch an instance. Otherwise, instead of using Amazon EC2 to create your key pair, you can create an RSA key pair using a third-party tool and then import the public key to Amazon EC2. For example, you can use **ssh-keygen** (a tool provided with the standard OpenSSH installation) to create a key pair. Alternatively, Java, Ruby, Python, and many other programming languages provide standard libraries that you can use to create an RSA key pair.

Amazon EC2 accepts the following formats:

- OpenSSH public key format (the format in `~/.ssh/authorized_keys`)

- Base64 encoded DER format
- SSH public key file format as specified in [RFC4716](#)

Amazon EC2 does not accept DSA keys. Make sure your key generator is set up to create RSA keys.

Supported lengths: 1024, 2048, and 4096.

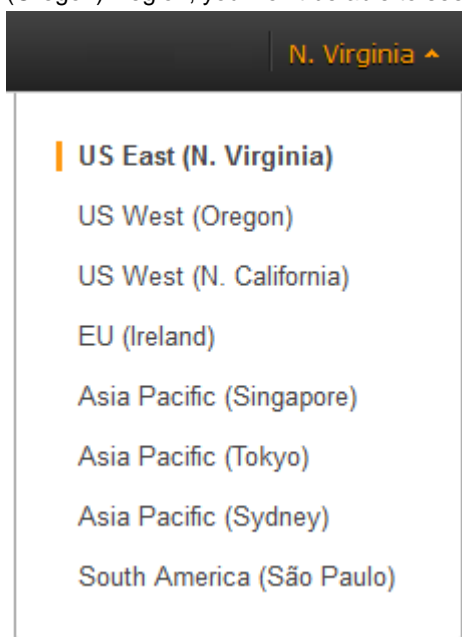
### To create a key pair using a third-party tool

1. Generate a key pair with a third-party tool of your choice.
2. Save the public key to a local file. For example, `C:\keys\my-key-pair.pub`. The file name extension for this file is not important.
3. Save the private key to a different local file that has the `.pem` extension. For example, `C:\keys\my-key-pair.pem`. Save the private key file in a safe place. You'll need to provide the name of your key pair when you launch an instance and the corresponding private key each time you connect to the instance.

Use the following steps to import your key pair using the Amazon EC2 console. (If you prefer, you can use the `ec2-import-keypair` command or the `ImportKeyPair` action to import the public key.)

### To import the public key

1. Open the Amazon EC2 console.
2. From the navigation bar, select the region for the key pair. This choice is important because key pair resources cannot be shared between regions. For example, if you import a key pair into the US West (Oregon) Region, you won't be able to see or use the key pair in another region.



3. Click **Key Pairs** in the navigation pane.
4. Click **Import Key Pair**.
5. In the **Import Key Pair** dialog box, click **Browse**, and select the public key file that you saved previously. Enter a name for the key pair in the **Key pair name** field, and click **Import**.

After the public key file is imported, you can verify that the key pair was imported successfully using the Amazon EC2 console as follows. (If you prefer, you can use the [ec2-describe-keypairs](#) command or the [DescribeKeyPairs](#) action to list your key pairs.)

### To verify that your key pair was imported

1. From the navigation bar, select the region in which you created the key pair.
2. Click **Key Pairs** in the navigation pane.
3. Verify that the key pair that you imported is in the displayed list of key pairs.

## Retrieving the Public Key for Your Key Pair

On a Linux/UNIX instance, the public key content is placed in an entry within `~/.ssh/authorized_keys`. This is done at boot time and enables you to securely access your instance without passwords. You can open this file in an editor to view the public key for your key pair. The following is an example entry for the key pair named `my-key-pair`. It consists of the public key followed by the name of the key pair. For example:

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAClKsfkNkuSevGj3eYhCe53pcjqP3maAhDFcvBS706Vhz2ItxCih+PnDSUaw+WNQn/mZphTk/a/gU8jEzoOWbkM4yxyb/wB96xbiFveSFJuOp/d6RJhJOI0iBXRlsLnBItnctckij7FbtXJMXLvvwJryDUilBMTjYtwB+QhYXUM0zce5Pjz5/i8SeJtjnV3iAoG/cQk+0FzZqaeJAAHco+CY/5WrUBkrHmFJr6HcXkvJdWPkYQS3xqC0+FmUZofz221CBt5IMucxXPkX4rWi+z7wB3RbBQoQzd8v7yeb7Oz1PnWOyN0qFU0XA246RA8QFYiCNYwi3f05p6KLxEXAMPLE my-key-pair
```

On Linux/UNIX, you can use **ssh-keygen** to get the public key for your key pair. Run the following command on a computer to which you've downloaded your private key:

```
$ ssh-keygen -y
```

When prompted by **ssh-keygen**, specify the path to your `.pem` file. The command returns the public key. For example:

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAClKsfkNkuSevGj3eYhCe53pcjqP3maAhDFcvBS706Vhz2ItxCih+PnDSUaw+WNQn/mZphTk/a/gU8jEzoOWbkM4yxyb/wB96xbiFveSFJuOp/d6RJhJOI0iBXRlsLnBItnctckij7FbtXJMXLvvwJryDUilBMTjYtwB+QhYXUM0zce5Pjz5/i8SeJtjnV3iAoG/cQk+0FzZqaeJAAHco+CY/5WrUBkrHmFJr6HcXkvJdWPkYQS3xqC0+FmUZofz221CBt5IMucxXPkX4rWi+z7wB3RbBQoQzd8v7yeb7Oz1PnWOyN0qFU0XA246RA8QFYiCNYwi3f05p6KLxEXAMPLE
```

On Windows, you can use PuTTYgen to get the public key for your key pair. Start PuTTYgen, click **Load**, and select the `.ppk` or `.pem` file. PuTTYgen displays the public key.

The public key that you specified when you launched an instance is also available to you through its instance metadata. To view the public key that you specified when launching the instance, use the following command from your instance:

```
$ GET http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQAClKsfkNkuSevGj3eYhCe53pcjqP3maAhDFcvBS706Vhz2ItxCih+PnDSUaw+WNQn/mZphTk/a/gU8jEzoOWbkM4yxyb/wB96xbiFveSFJuOp/d6RJhJOI0iBXRlsLnBItnctckij7FbtXJMXLvvwJryDUilBMTjYtwB+QhYXUM0zce5Pjz5/i8SeJtjnV3iAoG/cQk+0FzZqaeJAAHco+CY/5WrUBkrHmFJr6HcXkvJdWPkYQS3xqC0+FmUZofz221CBt5IMucxXPkX4rWi+z7wB3RbBQoQzd8v7yeb7Oz1PnWOyN0qFU0XA246RA8QFYiCNYwi3f05p6KLxEXAMPLE my-key-pair
```



Note that if you change the key pair that you use to connect to the instance, as shown in the next section on this page, we don't update the instance metadata to show the new public key; you'll continue to see the public key for the key pair you specified when you launched the instance in the instance metadata. For more information about instance metadata, see [Retrieving Instance Metadata \(p. 291\)](#).

## Connecting to Your Instance if You Lose Your Private Key

If you lose the private key for an EBS-backed instance, you can regain access to your instance. You must stop the instance, detach its root volume and attach it to another instance as a data volume, modify the `authorized_keys` file, move the volume back to the original instance, and restart the instance. This procedure isn't supported for instance store-backed instances or instances whose root volume has an AWS Marketplace product code. For more information about launching, connecting to, and stopping instances, see [Instance Lifecycle \(p. 263\)](#).

To determine the root device type of your instance, open the Amazon EC2 console, click **Instances**, select the instance, and check the value of **Root device type** in the details pane. The value is either `ebs` or `instance store`. If the root device is an instance store volume, you must have the private key in order to connect to the instance. You can also check the value of **Product codes** in the details pane to determine whether you're using an AWS Marketplace product code.

### Prerequisites

Create a new key pair using either the Amazon EC2 console or a third-party tool.

### To connect to an EBS-backed instance with a different key pair

1. From the Amazon EC2 console, click **Instances** in the navigation pane, and select the instance that you'd like to connect to. (We'll refer to this as the original instance.)
2. Save the following information that you'll see to complete this procedure.
  - Write down the instance ID (i-xxxxxxx), AMI ID (ami-xxxxxxx), and Availability Zone of the original instance.
  - Click the entry for `sda1` (the root volume) under **Block devices** in the details pane and write down the volume ID (vol-xxxxxxx).
  - [EC2-Classic] If the original instance has an associated Elastic IP address, write down the Elastic IP address shown under **Elastic IP** in the details pane.
3. Click **Actions**, and then click **Stop**. If **Stop** is disabled, either the instance is already stopped or its root device is an instance store volume.
4. Launch a temporary `t1.micro` instance as follows:
  - Use the same AMI that you used to launch the original instance. If that AMI is unavailable, you can create an AMI that you can use from the stopped instance. For more information, see [Creating Amazon EBS-Backed AMIs Using the Console \(p. 60\)](#).
  - Specify the new key pair that you created.
  - Specify the same Availability Zone as the instance you'd like to connect to.
  - Add the tag `Name=Temporary` to the instance to indicate that this is a temporary instance.
5. In the navigation pane, click **Volumes** and select the root device volume for the original instance (you wrote down its volume ID in a previous step). Click **Actions**, and then click **Detach Volume**. Wait for the state of the volume to become `available`. (You might need to click the **Refresh** icon.)

## Amazon Elastic Compute Cloud User Guide

### Connecting to Your Instance if You Lose Your Private Key

---

- With the volume still selected, click **Actions**, and then click **Attach Volume**. Select the instance ID of the temporary instance, write down the device name specified under **Device** (for example, `/dev/sdf`), and then click **Yes, Attach**.
- Connect to the temporary instance.
- From the temporary instance, mount the volume that you attached to the instance so that you can access its file system. For example, if the device name is `/dev/sdf`, use the following commands to mount the volume as `/mnt/tempvol`:

```
$ sudo mkdir /mnt/tempvol
$ sudo mount /dev/sdf /mnt/tempvol
```

- From the temporary instance, use the following command to update `authorized_keys` on `/dev/sdf` with the new public key from the `authorized_keys` for the temporary instance:

```
$ cp .ssh/authorized_keys /mnt/tempvol/home/ec2-user/.ssh/authorized_keys
```

If this copy succeeded, you can go to the next step.

(Optional) Otherwise, if you don't have permission to edit files in `/mnt/tempvol`, you'll need to update the file using **sudo** and then check the permissions on the file to verify that you'll be able to log into the original instance. Use the following command to check the permissions on the file:

```
$ sudo ls -l /mnt/home/ec2-user/.ssh
total 4
-rw----- 1 222 500 398 Sep 13 22:54 authorized_keys
```

In this example output, `222` is the user ID and `500` is the group ID. Next, use **sudo** to re-run the copy command that failed:

```
$ sudo cp .ssh/authorized_keys /mnt/tempvol/home/ec2-user/.ssh/authorized_keys
```

Run the following command again to determine whether the permissions changed:

```
$ sudo ls -l /mnt/home/ec2-user/.ssh
```

If the user ID and group ID have changed, use the following command to restore them:

```
$ sudo chown 222:500 /mnt/home/ec2-user/.ssh/authorized_keys
```

- From the temporary instance, unmount the volume that you attached so that you can reattach it to the original instance. For example, use the following command to unmount `/dev/sdf`:

```
$ sudo umount -d /dev/sdf
```

- From the Amazon EC2 console, select the volume with the volume ID that you wrote down, click **Actions**, and then click **Detach Volume**. Wait for the state of the volume to become `available`. (You might need to click the **Refresh** icon.)
- With the volume still selected, click **Actions**, and then click **Attach Volume**. Select the instance ID of the original instance, specify the device name `/dev/sda1`, and then click **Yes, Attach**.

### Warning

If you don't specify `sda1` as the device name, you'll be unable to start the original instance. This is because Amazon EC2 expects the root device volume at `sda1`.

13. Select the original instance, click **Actions**, and then click **Start**. After the instance enters the `running` state, you can connect to it using the private key file for your new key pair.
14. [EC2-Classic] If the original instance had an associated Elastic IP address before you stopped it, you must re-associate it with the instance as follows:
  - a. In the navigation pane, click **Elastic IPs**.
  - b. Select the Elastic IP address that you wrote down at the beginning of this procedure.
  - c. Click **Associate Address**.
  - d. Select the ID of the original instance, and then click **Associate**.
15. (Optional) You can terminate the temporary instance if you have no further use for it. Select the temporary instance, click **Actions**, and then click **Terminate**.

## Amazon EC2 Security Groups

A *security group* acts as a firewall that controls the traffic for one or more instances. When you launch an instance, you associate one or more security groups with the instance. You add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances that are associated with the security group. When we decide whether to allow traffic to reach an instance, we evaluate all the rules from all the security groups that are associated with the instance.

This topic provides information about security groups and security group rules.

### Topics

- [Security Groups for EC2-Classic \(p. 392\)](#)
- [Security Groups for EC2-VPC \(p. 393\)](#)
- [Security Group Rules \(p. 393\)](#)
- [Default Security Groups \(p. 394\)](#)
- [Custom Security Groups \(p. 394\)](#)
- [Creating a Security Group \(p. 395\)](#)
- [Describing Your Security Groups \(p. 396\)](#)
- [Adding Rules to a Security Group \(p. 396\)](#)
- [Deleting Rules from a Security Group \(p. 398\)](#)
- [Deleting a Security Group \(p. 398\)](#)
- [API and Command Overview \(p. 399\)](#)

If you have requirements that aren't met by security groups, you can maintain your own firewall on any of your instances in addition to using security groups.

## Security Groups for EC2-Classic

If you're using EC2-Classic, you must use security groups created specifically for EC2-Classic. When you launch an instance in EC2-Classic, you must specify a security group in the same region as the instance. You can't specify a security group that you created for a VPC when you launch an instance in EC2-Classic.

After you launch an instance in EC2-Classic, you can't change its security groups. However, you can add rules to or remove rules from a security group, and those changes are automatically applied to all instances that are associated with the security group.

**Note**

In EC2-Classic, you can associate an instance with up to 500 security groups and add up to 100 rules to a security group.

## Security Groups for EC2-VPC

If you're using EC2-VPC, you must use security groups created specifically for your VPC. When you launch an instance in a VPC, you must specify a security group for that VPC. You can't specify a security group that you created for EC2-Classic when you launch an instance in a VPC.

After you launch an instance in a VPC, you can change its security groups. You can also change the rules of a security group, and those changes are automatically applied to all instances that are associated with the security group.

**Note**

In EC2-VPC, you can associate a network interface with up to 5 security groups and add up to 50 rules to a security group.

When you specify a security group for a nondefault VPC to the CLI or the API actions, you must use the security group ID and not the security group name to identify the security group.

Security groups for EC2-VPC have additional capabilities that aren't supported by security groups for EC2-Classic. For more information about security groups for EC2-VPC, see [Security Groups for Your VPC](#) in the *Amazon Virtual Private Cloud User Guide*.

## Security Group Rules

The rules of a security group control the inbound traffic that's allowed to reach the instances that are associated with the security group and the outbound traffic that's allowed to leave them. By default, security groups allow all outbound traffic.

You can add and remove rules at any time. Your changes are automatically applied to the instances associated with the security group after a short period. You can't modify an existing rule in a security group; you must delete the rule and add a new rule. You can't change the outbound rules for EC2-Classic. Security group rules are always permissive; you can't create rules that deny access.

For each rule, you specify the following:

- The protocol to allow (for example, TCP, UDP, or ICMP).
- TCP and UDP: The range of ports to allow
- ICMP: The ICMP type and code
- One or the following options for the source (inbound rules) or destination (outbound rules):
  - An individual IP address, in CIDR notation. Be sure to use the `/32` prefix after the IP address; if you use the `/0` prefix after the IP address, this opens the port to everyone. For example, specify the IP address `203.0.113.1` as `203.0.113.1/32`.
  - An IP address range, in CIDR notation (for example, `203.0.113.0/24`).
  - The name (EC2-Classic) or ID (EC2-Classic or EC2-VPC) of one of the following security groups:
    - The current security group.
    - EC2-Classic: A different security group for EC2-Classic in the same region
    - EC2-VPC: A different security group for the same VPC
    - EC2-Classic: A security group for another AWS account in the same region (add the AWS account ID as a prefix; for example, `111122223333/sg-edcd9784`)

When you specify a security group as the source or destination for a rule, the rule affects all instances associated with the security group. For example, the incoming traffic is allowed based on the private IP addresses of the instances that are associated with the source security group.

If there is more than one rule for a specific port, we apply the most permissive rule. For example, if you have a rule that allows access to TCP port 22 (SSH) from IP address 203.0.113.1 and another rule that allows access to TCP port 22 from everyone, everyone has access to TCP port 22.

When you associate multiple security groups with an instance, the rules from each security group are effectively aggregated to create one set of rules. We use this set of rules to determine whether to allow access.

### Caution

Because you can assign multiple security groups to an instance, an instance can have hundreds of rules that apply. This might cause problems when you access the instance. Therefore, we recommend that you condense your rules as much as possible.

For more information about IP addresses, see [Amazon EC2 Instance IP Addressing \(p. 419\)](#).

## Default Security Groups

Your AWS account automatically has a *default security group* per region for EC2-Classic. When you create a VPC, we automatically create a default security group for the VPC. If you don't specify a different security group when you launch an instance, the instance is automatically associated with the appropriate default security group.

A default security group is named `default`, and it has an ID assigned by AWS. The following are the initial settings for each default security group:

- Allow inbound traffic only from other instances associated with the default security group
- Allow all outbound traffic from the instance

The default security group specifies itself as a source security group in its inbound rules. This is what allows instances associated with the default security group to communicate with other instances associated with the default security group.

You can change the rules for a default security group. For example, you can add an inbound rule to allow SSH or Remote Desktop connections so that specific hosts can manage the instance.

You can't delete a default security group.

## Custom Security Groups

If you don't want all your instances to use the default security group, you can create your own security groups and specify them when you launch your instances. You can create multiple security groups to reflect the different roles that your instances play; for example, a web server or a database server. For instructions that help you create security groups for web servers and database servers, see [Recommended Security Groups](#) in the *Amazon Virtual Private Cloud User Guide*.

### Note

In EC2-Classic, you can create up to 500 security groups in each region for each account. In EC2-VPC, you can create up to 100 security groups per VPC. The security groups for EC2-Classic do not count against the security group limit for EC2-VPC.

When you create a security group, you must provide it with a name and a description. Security group names and descriptions can be up to 255 characters in length, and are limited to the following characters:

- EC2-Classic: ASCII characters

- EC2-VPC: a-z, A-Z, 0-9, spaces, and `._-:/()#,@[]+=&:{}!$*`

AWS assigns each security group a unique ID in the form `sg-xxxxxxx`. The following are the initial settings for a security group that you create:

- Allow no inbound traffic
- Allow all outbound traffic

After you've created a security group, you can change its inbound rules to reflect the type of inbound traffic that you want to reach the associated instances. In EC2-VPC, you can also change its outbound rules.

To allow instances that have the same security group to communicate, you must explicitly add rules for this. The following table describes the rules that you must add to your security group to enable instances in EC2-Classic to communicate.

Inbound			
Source	Protocol	Port Range	Comments
The ID of the security group	ICMP	All	Allow inbound ICMP access from other instances associated with this security group
The ID of the security group	TCP	0 - 65535	Allow inbound TCP access from other instances associated with this security group
The ID of the security group	UDP	0 - 65535	Allow inbound UDP access from other instances associated with this security group

The following table describes the rules that you must add to your security group to enable instances in a VPC to communicate.

Inbound			
Source	Protocol	Port Range	Comments
The ID of the security group	All	All	Allow inbound traffic from other instances associated with this security group

## Creating a Security Group

### To create a security group for EC2-Classic

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Security Groups**.
3. Click **Create Security Group**.
4. Specify a name and description for the security group. Select `No VPC` for **VPC**, and then click **Yes, Create**.

### To create a security group for EC2-VPC

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Security Groups**.
3. Click **Create Security Group**.
4. Specify a name and description for the security group. Select the ID of your VPC for **VPC**, and then click **Yes, Create**.

You can assign a security group to an instance when you launch the instance. When you add or remove rules, those changes are automatically applied to all instances to which you've assigned the security group.

After you launch an instance in EC2-Classic, you can't change its security groups. After you launch an instance in a VPC, you can change its security groups. For more information, see [Changing an Instance's Security Groups](#) in the *Amazon Virtual Private Cloud User Guide*.

## Describing Your Security Groups

### To describe your security groups for EC2-Classic

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Security Groups**.
3. Select **EC2 Security Groups** from the **Viewing** list.
4. Select a security group. We display general information in the **Details** tab and inbound rules on the **Inbound** tab.

### To describe your security groups for EC2-VPC

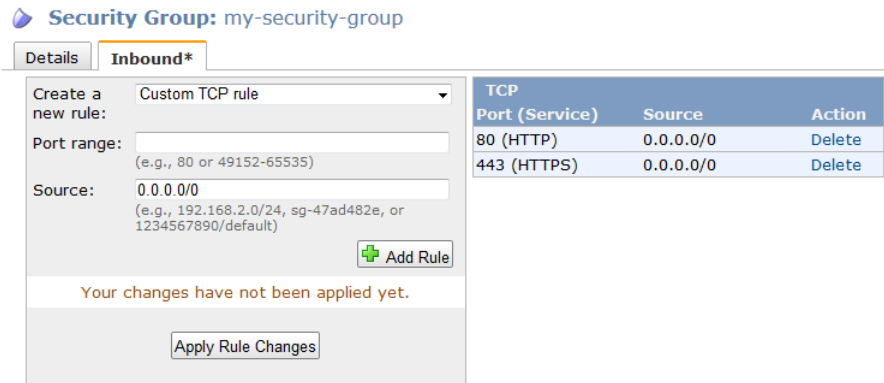
1. Open the Amazon EC2 console.
2. In the navigation pane, click **Security Groups**.
3. Select **VPC Security Groups** from the **Viewing** list.
4. Select a security group. We display general information in the **Details** tab and inbound rules on the **Inbound** tab.

## Adding Rules to a Security Group

When you add a rule to a security group, the new rule is automatically applied to any instances associated with the security group.

### To add rules to a security group

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Security Groups**.
3. Select the security group.
4. You can allow web servers to receive all inbound HTTP and HTTPS traffic. On the **Inbound** tab, select **HTTP** from **Create a new rule**, leave **Source** as `0.0.0.0/0`, and then click **Add Rule**. Notice that the **Apply Rule Changes** button is now enabled, and the text "Your changes have not been applied yet" appears above the button. Add a similar rule for **HTTPS**, and then click **Apply Rule Changes** to add both rules.



5. To connect to a Linux instance, you'll need to allow SSH traffic.
  - a. On the **Inbound** tab, select **SSH** from the **Create a new rule** list.
  - b. In the **Source** box, specify the public IP address of your computer, in CIDR notation. For example, if your IP address is 203.0.113.25, specify 203.0.113.25/32 to list this single IP address in CIDR notation. If your company allocates addresses from a range, specify the entire range, such as 203.0.113.0/24.


**Tip**

You can get the public IP address of your local computer using a service. For example, we provide the following service: <http://checkip.amazonaws.com/>. To locate another service that provides your IP address, use the search phrase "what is my IP address". If you are connecting through an ISP or from behind your firewall without a static IP address, you need to find out the range of IP addresses used by client computers.

**Caution**

If you use 0.0.0.0/0, you enable all IP addresses to access your instance using SSH. This is acceptable for a short time in a test environment, but it's unsafe for production environments. In production, you'll authorize only a specific IP address or range of addresses to access your instance.

- c. Click **Add Rule**, and then click **Apply Rule Changes** to add the rule.



6. You can allow communication between all instances associated with this security group. Start typing the ID of the security group in **Source**; this provides you with a list of security groups. Select the security group from the list, click **Add Rule**, and then click **Apply Rule Changes**.
  - a. Select **All ICMP** from **Create a new rule**. In the **Source** box, specify the ID of the current security group. Click **Add Rule**.
  - b. Select **All TCP** from **Create a new rule**. In the **Source** box, specify the ID of the current security group. Click **Add Rule**.



- c. Select `ALL UDP` from **Create a new rule**. In the **Source** box, specify the ID of the current security group. Click **Add Rule**.
- d. Click **Apply Rule Changes** to add all three rules.

**Security Group: my-security-group**

Details **Inbound\***

Create a new rule: Custom TCP rule

Port range:   
(e.g., 80 or 49152-65535)

Source:   
(e.g., 192.168.2.0/24, sg-47ad482e, or 1234567890/default)

Your changes have not been applied yet.

ICMP		
Port (Service)	Source	Action
ALL	sg-164ff27d	Delete

TCP		
Port (Service)	Source	Action
80 (HTTP)	0.0.0.0/0	Delete
443 (HTTPS)	0.0.0.0/0	Delete
22 (SSH)	203.0.113.0/24	Delete
0 - 65535	sg-164ff27d	Delete

UDP		
Port (Service)	Source	Action
0 - 65535	sg-164ff27d	Delete

7. If you are creating a security group for a VPC, you can also specify outbound rules. For an example, see [Adding and Removing Rules](#) in the *Amazon Virtual Private Cloud User Guide*.

## Deleting Rules from a Security Group

When you delete a rule from a security group, the change is automatically applied to any instances associated with the security group.

### To delete a security group rule

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Security Groups**.
3. Select a security group.
4. Click **Delete** next to each rule that you need to delete.

An asterisk appears on the **Inbound** tab to indicate that there are changes that have not been applied.

5. Click **Apply Rule Changes**.

## Deleting a Security Group

You can't delete a security group that associated with an instance. You can't delete the default security group.

### To delete a security group

1. Open the Amazon EC2 console.
2. In the navigation pane, click **Security Groups**.
3. Select a security group and click **Delete**.
4. Click **Yes, Delete**.

## API and Command Overview

The following table summarizes the available commands and corresponding API actions for security groups.

Description	Command	API Action
Creates a security group.	<a href="#">ec2-create-group</a>	<a href="#">CreateSecurityGroup</a>
Adds one or more rules to a security group.	<a href="#">ec2-authorize</a>	<a href="#">AuthorizeSecurityGroupIngress</a> <a href="#">AuthorizeSecurityGroupEgress</a> (EC2-VPC only)
Describes one or more of your security groups.	<a href="#">ec2-describe-group</a>	<a href="#">DescribeSecurityGroups</a>
[EC2-VPC only] Modifies the security groups an instance is associated with.	<a href="#">ec2-modify-instance-attribute</a>	<a href="#">ModifyInstanceAttribute</a>
Removes one or more rules from a security group.	<a href="#">ec2-revoke</a>	<a href="#">RevokeSecurityGroupIngress</a>
Deletes a security group.	<a href="#">ec2-delete-group</a>	<a href="#">DeleteSecurityGroup</a>

## Controlling Access to Amazon EC2 Resources

Your security credentials identify you to services in AWS and grant you unlimited use of your AWS resources, such as your Amazon EC2 resources. You can use features of Amazon EC2 and AWS Identity and Access Management (IAM) to allow other users, services, and applications to use your Amazon EC2 resources without sharing your security credentials. You can choose to allow full use or limited use of your Amazon EC2 resources.

### Topics

- [Network Access to Your Instance](#) (p. 399)
- [Amazon EC2 Permission Attributes](#) (p. 399)
- [Introduction to IAM and Amazon EC2](#) (p. 400)
- [IAM Policies for Amazon EC2](#) (p. 401)
- [IAM Roles for Amazon EC2](#) (p. 408)
- [Authorizing Inbound Traffic for Your Instances](#) (p. 412)

## Network Access to Your Instance

A security group acts as a firewall that controls the traffic allowed to reach one or more instances. When you launch an instance, you assign it one or more security groups. You add rules to each security group that control traffic for the instance. You can modify the rules for a security group at any time; the new rules are automatically applied to all instances to which the security group is assigned.

For more information, see [Authorizing Inbound Traffic for Your Instances](#) (p. 412).

## Amazon EC2 Permission Attributes

Your organization might have multiple AWS accounts. Amazon EC2 enables you to specify additional AWS accounts that can use your Amazon Machine Images (AMIs) and Amazon EBS snapshots. These permissions work at the AWS account level only; you can't restrict permissions for specific users within

the specified AWS account. All users in the AWS account that you've specified can use the AMI or snapshot.

Each AMI has a `LaunchPermission` attribute that controls which AWS accounts can access the AMI. For more information, see [Making an AMI Public \(p. 50\)](#).

Each Amazon EBS snapshot has a `createVolumePermission` attribute that controls which AWS accounts can use the snapshot. For more information, see [Sharing Snapshots \(p. 492\)](#).

## Introduction to IAM and Amazon EC2

IAM enables you to do the following:

- Create users and groups under your AWS account
- Assign unique security credentials to each user under your AWS account
- Control each user's permissions to perform tasks using AWS resources
- Allow the users in another AWS account to share your AWS resources
- Create roles for your AWS account and define the users or services that can assume them
- Use existing identities for your enterprise to grant permissions to perform tasks using AWS resources

By using IAM with Amazon EC2, you can control whether users in your organization can perform a task using specific Amazon EC2 API actions and whether they can use specific AWS resources.

## Creating an IAM Group and Users

### To create an IAM group and users

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. From the dashboard, click **Create a New Group of Users**.
3. On the **GROUP NAME** page, specify the name of the group.
4. On the **PERMISSIONS** page, specify the policies for the group. You can select a policy template or create custom policies. For example, for Amazon EC2, one of the following policy templates might meet your needs:

- Power User Access
- Read Only Access
- Amazon EC2 Full Access
- Amazon EC2 Read Only Access

For more information about creating custom policies, see [IAM Policies for Amazon EC2 \(p. 401\)](#).

5. On the **USERS** page, enter one or more user names. If the users will use the CLI or API, select **Generate an access key for each User**. Click **Continue**.
6. If you had IAM generate access keys, click **Download Credentials** or **Show User Security Credentials** and save the access keys. As indicated on the dialog box, this is your only chance to retrieve and save your secret access key.
7. If the users will use the console, click **Users** in the navigation pane and do the following for each user:
  - a. Select the user.
  - b. Click the **Security Credentials** tab in the details pane.
  - c. Under **Sign-In Credentials**, click **Manage Password**.
  - d. In the **Manage Password** dialog box, select an option and click **Apply**.

- e. Click **Download Credentials** or **Show User Security Credentials** and save the password.
8. Give each user his or her credentials (access keys and password); this enables them to use services based on the permissions you specified for the IAM group.

## Related Topics

For more information about IAM, see the following:

- [IAM Policies for Amazon EC2 \(p. 401\)](#)
- [IAM Roles for Amazon EC2 \(p. 408\)](#)
- [Identity and Access Management \(IAM\)](#)
- [Using IAM](#)

## IAM Policies for Amazon EC2

By default, IAM users don't have permission to create or modify Amazon EC2 resources, or perform tasks using the Amazon EC2 API. (This means that they also can't do so using the Amazon EC2 console or CLI.) To allow IAM users to create or modify resources and perform tasks, you must create IAM policies that grant IAM users permission to use the specific resources and API actions they'll need, and then attach those policies to the IAM users or groups that require those permissions.

When you attach a policy to a user or group of users, it allows or denies the users permission to perform the specified tasks on the specified resources.

### Topics

- [Policy Syntax \(p. 401\)](#)
- [Actions for Amazon EC2 \(p. 402\)](#)
- [Amazon Resource Names \(ARNs\) for Amazon EC2 \(p. 402\)](#)
- [Condition Keys for Amazon EC2 \(p. 404\)](#)
- [Example Policy Statements for Amazon EC2 \(p. 405\)](#)
- [Checking that Users Have the Required Permissions \(p. 408\)](#)

## Policy Syntax

An IAM policy is a JSON document that consists of one or more statements. Each statement is structured as follows:

```
{
  "Statement": [ {
    "Effect": "effect",
    "Action": "action",
    "Resource": "arn"
    "Condition": {
      "condition": {
        "key": "value"
      }
    }
  }
}
```

```
]
}
```

The *effect* can be `Allow` or `Deny`. By default, IAM users don't have permission to use resources and API actions, so all requests are denied. An explicit allow overrides the default. An explicit deny overrides any allows.

To learn about specifying *action*, see [Actions for Amazon EC2 \(p. 402\)](#). To learn about specifying *arn*, see [Amazon Resource Names \(ARNs\) for Amazon EC2 \(p. 402\)](#).

Conditions are optional. To learn about specifying conditions for Amazon EC2, see [Condition Keys for Amazon EC2 \(p. 404\)](#).

For example IAM policy statements for Amazon EC2, see [Example Policy Statements for Amazon EC2 \(p. 405\)](#). For more information about IAM policies, see [Permissions and Policies](#) in the *Using IAM* guide.

## Actions for Amazon EC2

In an IAM policy statement, you can specify any API action from any service that supports IAM. For Amazon EC2, use the following prefix with the name of the API action: `ec2:`. For example: `ec2:RunInstances` and `ec2:CreateImage`.

To specify multiple actions in a single statement, separate them with commas as follows:

```
"Action": ["ec2:action1", "ec2:action2"]
```

You can also specify multiple actions using wildcards. For example, you can specify all actions whose name begins with the word "Describe" as follows:

```
"Action": "ec2:Describe*"
```

To specify all Amazon EC2 API actions, use a wildcard as follows:

```
"Action": "ec2:*"
```

For a list of Amazon EC2 actions, see [Actions](#) in the *Amazon Elastic Compute Cloud API Reference*.

## Amazon Resource Names (ARNs) for Amazon EC2

Each IAM policy statement applies to the resources that you specify using their ARNs. An ARN has the following general syntax:

```
arn:aws:[service]:[region]:[account]:resourceType/resourcePath
```

**service**

The service (for example, `ec2`).

**region**

The region for the resource (for example, `us-east-1`).

**account**

The AWS account ID, with no hyphens (for example, `123456789012`).

**resourceType**

The type of resource (for example, `instance`).

**resourcePath**

A path that identifies the resource. You can use the \* wildcard in your paths.

The following table describes the ARNs for each type of resource used by the Amazon EC2 API actions. (We'll add ARNs for additional Amazon EC2 resources later in 2013.)

Resource Type	ARN
Customer gateway	arn:aws:ec2:region:account:customer-gateway/cgw-id Where <i>cgw-id</i> is cgw-xxxxxxx
DHCP options set	arn:aws:ec2:region:account:dhcp-options/dhcp-options-id Where <i>dhcp-options-id</i> is dopt-xxxxxxx
Instance	arn:aws:ec2:region:account:instance/instance-id Where <i>instance-id</i> is i-xxxxxxx
Instance profile	arn:aws:iam::account:instance-profile/instance-profile-name Where <i>instance-profile-name</i> is the name of the instance profile, and <i>region</i> isn't used
Internet gateway	arn:aws:ec2:region:account:internet-gateway/igw-id Where <i>igw-id</i> is igw-xxxxxxx
Network ACL	arn:aws:ec2:region:account:network-acl/nacl-id Where <i>nacl-id</i> is acl-xxxxxxx
Placement group	arn:aws:ec2:region:account:placement-group/placement-group-name Where <i>placement-group-name</i> is the placement group name (for example, <i>my-cluster</i> )
Route table	arn:aws:ec2:region:account:route-table/route-table-id Where <i>route-table-id</i> is rtb-xxxxxxx
Security group	arn:aws:ec2:region:account:security-group/security-group-id Where <i>security-group-id</i> is sg-xxxxxxx
Snapshot	arn:aws:ec2:region::snapshot/snapshot-id Where <i>snapshot-id</i> is snap-xxxxxxx, and <i>account</i> isn't used
Volume	arn:aws:ec2:region:account:volume/volume-id Where <i>volume-id</i> is vol-xxxxxxx
VPC	arn:aws:ec2:region:account:vpc/vpc-id Where <i>vpc-id</i> is vpc-xxxxxxx
All Amazon EC2 resources owned by the specified account in the specified region	arn:aws:ec2:region:account:*

Resource Type	ARN
All Amazon EC2 resources	arn:aws:ec2:*

Many Amazon EC2 API actions involve multiple resources. For example, `AttachVolume` attaches an Amazon EBS volume to an instance, so an IAM user must have permission to use the volume and the instance. To specify multiple resources in a single statement, separate their ARNs with commas, as follows:

```
"Resource": [ "arn1", "arn2" ]
```

For information about which ARNs you can use with which Amazon EC2 API actions, see [Granting IAM Users Required Permissions for Amazon EC2 Resources](#) in the *Amazon Elastic Compute Cloud API Reference*.

For more information about ARNs, see [Amazon Resource Names \(ARN\) and AWS Service Namespaces](#) in the *Amazon Web Services General Reference*.

## Condition Keys for Amazon EC2

In a policy statement, you can optionally specify conditions that control when it is in effect. Each condition contains one or more key-value pairs. Condition keys are not case sensitive. We've defined AWS-wide condition keys, plus additional service-specific condition keys.

If you specify multiple conditions, or multiple keys in a single condition, we evaluate them using a logical AND operation. If you specify a single condition with multiple values for one key, we evaluate the condition using a logical OR operation. For permission to be granted, all conditions must be met.

You can also use placeholders when you specify conditions. For example, you can grant an IAM user permission to use resources with a tag that specifies his or her IAM user name. For more information, see [Policy Variables](#) in the *Using IAM* guide.

Amazon EC2 implements the AWS-wide condition keys (see [Available Keys](#)), plus the following service-specific condition keys. (We'll add support for additional service-specific condition keys for Amazon EC2 later in 2013.)

Condition Key	Key/Value Pair	Evaluation Types
ec2:AvailabilityZone	"ec2:AvailabilityZone": " <i>az-api-name</i> " Where <i>az-api-name</i> is the name of the Availability Zone (for example, <i>us-west-2a</i> ). To list your Availability Zones, use <a href="#">ec2-describe-availability-zones</a> .	String, Null
ec2:EbsOptimized	"ec2:EbsOptimized": " <i>optimized-flag</i> " Where <i>optimized-flag</i> is <code>true</code>   <code>false</code>	Boolean, Null
ec2:InstanceProfile	"ec2:InstanceProfile": " <i>instance-profile-arn</i> " Where <i>instance-profile-arn</i> is the instance profile ARN	ARN, Null

Condition Key	Key/Value Pair	Evaluation Types
ec2:InstanceType	"ec2:InstanceType": <i>"instance-type-api-name"</i> Where <i>instance-type-api-name</i> is the name of the instance type ( <i>m1.small</i>   <i>m1.medium</i>   <i>m1.large</i>   <i>m1.xlarge</i>   <i>m3.xlarge</i>   <i>m3.2xlarge</i>   <i>c1.medium</i>   <i>c1.xlarge</i>   <i>cc2.8xlarge</i>   <i>m2.xlarge</i>   <i>m2.2xlarge</i>   <i>m2.4xlarge</i>   <i>cr1.8xlarge</i>   <i>hi1.4xlarge</i>   <i>hs1.8xlarge</i>   <i>t1.micro</i>   <i>cg1.4xlarge</i>   <i>g2.2xlarge</i> ).	String, Null
ec2:ParentSnapshot	"ec2:ParentSnapshot": <i>"snapshot-arn"</i> Where <i>snapshot-arn</i> is the snapshot ARN	ARN, Null
ec2:PlacementGroup	"ec2:PlacementGroup": <i>"placement-group-arn"</i> Where <i>placement-group-arn</i> is the placement group ARN	ARN, Null
ec2:Region	"ec2:Region": <i>"region-name"</i> Where <i>region-name</i> is the name of the region (for example, <i>us-west-2</i> ). To list your regions, use <a href="#">ec2-describe-regions</a> .	String, Null
ec2:ResourceTag/ <i>tag-key</i>	"ec2:ResourceTag/ <i>tag-key</i> ": <i>"tag-value"</i> Where <i>tag-key</i> and <i>tag-value</i> are the tag-key pair	String, Null
ec2:RootDeviceType	"ec2:RootDeviceType": <i>"root-device-type-name"</i> Where <i>root-device-type-name</i> is <i>ebs</i>   <i>instance-store</i>	String, Null
ec2:Tenancy	"ec2:Tenancy": <i>"tenancy-attribute"</i> Where <i>tenancy-attribute</i> is <i>default</i>   <i>dedicated</i>	String, Null
ec2:Volumelops	"ec2:Volumelops": <i>"volume-iops"</i> Where <i>volume-iops</i> is the input/output operations per second (IOPS); the range is 100 to 4000	Numeric, Null
ec2:VolumeSize	"ec2:VolumeSize": <i>"volume-size"</i> Where <i>volume-size</i> is the size of the volume, in GiB	Numeric, Null
ec2:VolumeType	"ec2:VolumeType": <i>"volume-type-name"</i> Where <i>volume-type-name</i> is <i>io1</i>   <i>standard</i>	String, Null
ec2:Vpc	"ec2:Vpc": <i>"vpc-arn"</i> Where <i>vpc-arn</i> is the VPC ARN	ARN, Null

For information about which condition keys you can use with which Amazon EC2 resources, on an action-by-action basis, see [Granting IAM Users Required Permissions for Amazon EC2 Resources](#) in the *Amazon Elastic Compute Cloud API Reference*. For example policy statements for Amazon EC2, see [Example Policy Statements for Amazon EC2](#) (p. 405).

For more information about IAM policies, see [Permissions and Policies](#) in the *Using IAM* guide.

## Example Policy Statements for Amazon EC2

The following examples show policy statements that you could use to control the permissions that IAM users have to Amazon EC2.



- [1: Allow users to list the Amazon EC2 resources that belong to the AWS account \(p. 406\)](#)
- [2: Allow users to describe, launch, stop, start, and terminate all instances \(p. 406\)](#)
- [3: Allow users to stop and start only particular instances \(p. 406\)](#)
- [4. Allow users to manage particular volumes for particular instances \(p. 407\)](#)

### Example 1: Allow users to list the Amazon EC2 resources that belong to the AWS account

The following policy grants users permission to use all Amazon EC2 API actions whose names begin with `Describe`. The `Resource` element uses a wildcard to indicate that users can specify all resources with these API actions. The users don't have permission to use any other API actions (unless another statement grants them permission to do so) because users are denied permission to use API actions by default.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "ec2:Describe*",
    "Resource": "*"
  }]
}
```

### Example 2: Allow users to describe, launch, stop, start, and terminate all instances

The following policy grants users permission to use the API actions specified in the `Action` element. The `Resource` element uses a wildcard to indicate that users can specify all resources with these API actions. The users don't have permission to use any other API actions (unless another statement grants them permission to do so) because users are denied permission to use API actions by default.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances", "ec2:DescribeImages",
      "ec2:DescribeKeyPairs", "ec2:DescribeSecurityGroups",
      "ec2:DescribeAvailabilityZones",
      "ec2:RunInstances", "ec2:TerminateInstances",
      "ec2:StopInstances", "ec2:StartInstances"
    ],
    "Resource": "*"
  }]
}
```

### Example 3: Allow users to stop and start only particular instances

The following policy allows users to start and stop only the instances that have the tag `"department=dev"`, associated with them. Although the ARN in this statement specifies all instances in the US East (Northern Virginia) Region (`us-east-1`) that belong to the specified AWS account, the `Condition` element qualifies when the policy statement is in effect.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:StopInstances",
      "ec2:StartInstances"
    ],
    "Resource": "arn:aws:ec2:us-east-1:123456789012:instance/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/department": "dev"
      }
    }
  ]
}
```

#### Example 4. Allow users to manage particular volumes for particular instances

When an API action requires a caller to specify multiple resources, you must create a policy statement that allows users to access all required resources. If you need to use a `Condition` element with one or more of these resources, you must create multiple statements as shown in this example.

The following policy allows users to attach volumes with the tag `volume_user=iam-user-name` to instances with the tag `department=dev`, and to detach those volumes from those instances. If you attach this policy to an IAM group, the `aws:username` policy variable gives each IAM user in the group permission to attach or detach volumes from the instances with a tag named `volume_user` that has his or her IAM user name as a value.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:AttachVolume",
      "ec2:DetachVolume"
    ],
    "Resource": "arn:aws:ec2:us-east-1:123456789012:instance/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/department": "dev"
      }
    }
  }],
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:AttachVolume",
      "ec2:DetachVolume"
    ],
    "Resource": "arn:aws:ec2:us-east-1:123456789012:volume/*",
    "Condition": {
      "StringEquals": {
        "ec2:ResourceTag/volume_user": "${aws:username}"
      }
    }
  ]
}
```

```
}  
  }  
] }  
}
```

## Checking that Users Have the Required Permissions

After you've created an IAM policy, we recommend that you check whether it grants users the permissions to use the particular API actions and resources they'll need before you put the policy into production.

First, create an IAM user for testing purposes, and then attach the IAM policy that you created to the test user. Then, make a request as the test user.

If the action that you are testing creates or modifies a resource, you should make the request using the `DryRun` parameter (or run the CLI command with the `--auth-dry-run` option). In this case, the call completes the authorization check, but does not complete the operation. For example, you can check whether the user can terminate a particular instance without actually terminating it. If the test user has the required permissions, the request returns `DryRunOperation`; otherwise, it returns `UnauthorizedOperation`.

If the policy doesn't grant the user the permissions that you expected, or is overly permissive, you can adjust the policy as needed and retest until you get the desired results.

### Important

It can take several minutes for policy changes to propagate before they take effect. Therefore, we recommend that you allow five minutes to pass before you test your policy updates.

If an authorization check fails, the request returns an encoded message with diagnostic information. You can decode the message using the `DecodeAuthorizationMessage` action. For more information, see [DecodeAuthorizationMessage](#) in the *AWS Security Token Service API Reference*.

## IAM Roles for Amazon EC2

Applications must sign their API requests with AWS credentials. Therefore, if you are an application developer, you need a strategy for managing credentials for your applications that run on EC2 instances. For example, you can securely distribute your AWS credentials to the instances, enabling the applications on those instances to use your credentials to sign requests, while protecting them from other users. However, it's challenging to securely distribute credentials to each instance, especially those that AWS creates on your behalf, such as Spot Instances or instances in Auto Scaling groups. You must also be able to update the credentials on each instance when you rotate your AWS credentials.

We designed IAM roles so that your applications can securely make API requests from your instances, without requiring you to manage the security credentials that the applications use. Instead of creating and distributing your AWS credentials, you can delegate permission to make API requests using IAM roles as follows:

1. Create an IAM role.
2. Define which accounts or AWS services can assume the role.
3. Define which API actions and resources the application can use after assuming the role.
4. Specify the role when you launch your instances.
5. Have the application retrieve a set of temporary credentials and use them.

For example, you can use IAM roles to grant permissions to applications running on your instances that needs to use a bucket in Amazon S3.

You can specify permissions for IAM roles by creating a policy in JSON format. These are similar to the policies that you create for IAM users. If you make a change to a role, the change is propagated to all instances, simplifying credential management. For more information about creating and using IAM roles, see [Roles](#) in the *Using IAM* guide.

## Retrieving Security Credentials from Instance Metadata

An application on the instance retrieves the security credentials provided by the role from the instance metadata item `iam/security-credentials/role-name`. The application is granted the permissions for the actions and resources that you've defined for the role through the security credentials associated with the role. These security credentials are temporary and we rotate them automatically. We make new credentials available at least five minutes prior to the expiration of the old credentials.

### Warning

If you use services that use instance metadata with IAM roles, ensure that you don't expose your credentials when the services make HTTP calls on your behalf. The types of services that could expose your credentials include HTTP proxies, HTML/CSS validator services, and XML processors that support XML inclusion.

The following command retrieves the security credentials for an IAM role named `s3access`.

```
GET http://169.254.169.254/latest/meta-data/iam/security-credentials/s3access
```

The following is example output.

```
{
  "Code" : "Success",
  "LastUpdated" : "2012-04-26T16:39:16Z",
  "Type" : "AWS-HMAC",
  "AccessKeyId" : "AKIAIOSFODNN7EXAMPLE",
  "SecretAccessKey" : "wJalrXUtnFEMI/K7MDENG/bPxrRfiCYEXAMPLEKEY",
  "Token" : "token",
  "Expiration" : "2012-04-27T22:39:16Z"
}
```

For more information about instance metadata, see [Instance Metadata and User Data \(p. 290\)](#). For more information about temporary credentials, see [Using Temporary Security Credentials](#) in the IAM documentation.

## Granting an IAM User Permission to Launch an Instance with an IAM Role

To enable an IAM user to launch an instance with an IAM role, you must grant the user permission to pass the role to the instance.

For example, the following IAM policy grants users permission to launch an instance with the IAM role named `s3access`.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": "arn:aws:iam::123456789012:role/s3access"
  }]
}
```

```
} ]  
}
```

Alternatively, you could grant IAM users access to all your roles by specifying the resource as "\*" in this policy. However, consider whether users who launch instances with your roles (ones that exist or that you'll create later on) might be granted permissions that they don't need or shouldn't have.

For more information, see [Permissions Required for Using Roles with Amazon EC2](#) in the *Using IAM* guide.

## Launching an Instance with an IAM Role using the Console

You must create an IAM role before you can launch an instance with that role.

### Important

After you create an IAM role, it may take several seconds for the permissions to propagate. If your first attempt to launch an instance with a role fails, wait a few seconds before trying again. For more information, see [Troubleshooting Working with Roles](#) in the *Using IAM* guide.

### To launch an instance with an IAM role using the AWS Management Console

1. Create an IAM role.
  - a. Open the IAM console at <https://console.aws.amazon.com/iam/>.
  - b. In the navigation pane, click **Roles**, and then click **Create New Role**.
  - c. On the first **CONFIGURE ROLE** page, enter a name for the role and click **Continue**.
  - d. On the second **CONFIGURE ROLE** page, click **Select** next to **Amazon EC2**.
  - e. On the **PERMISSIONS** page, specify the policies for the group. You can select a policy template or create custom policies. For example, for Amazon EC2, one of the following policy templates might meet your needs:
    - Power User Access
    - Read Only Access
    - Amazon EC2 Full Access
    - Amazon EC2 Read Only Access
  - f. For more information about creating custom policies, see [IAM Policies for Amazon EC2 \(p. 401\)](#).
2. Launch an instance with the IAM role.
  - a. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
  - b. On the dashboard, click **Launch Instance**.
  - c. Select an AMI, then select an instance type and click **Next: Configure Instance Details**.
  - d. On the **Configure Instance Details** page, select the IAM role you created from the **IAM role** list.
  - e. Configure any other details, then follow the instructions through the rest of the wizard, or click **Review and Launch** to accept default settings and go directly to the **Review Instance Launch** page.
  - f. Review your settings, then click **Launch** to choose a key pair and launch your instance.

3. If you are using the Amazon EC2 API actions in your application, retrieve the AWS security credentials made available on the instance and use them to sign the requests. Note that the AWS SDK does this for you.

```
GET http://169.254.169.254/latest/meta-data/iam/security-credentials/role_name
```

## Launching an Instance with an IAM Role using the CLI

You must create an IAM role before you can launch an instance with that role.

### Important

After you create an IAM role, it may take several seconds for the permissions to propagate. If your first attempt to launch an instance with a role fails, wait a few seconds before trying again. For more information, see [Troubleshooting Working with Roles](#) in the *Using IAM* guide.

### To launch an instance with an IAM role using the IAM and Amazon EC2 CLIs

1. The following example creates an IAM role named `s3access` with a policy that allows the role to use an Amazon S3 bucket, and an instance profile named `s3access`.

```
iam-rolecreate -r s3access -s ec2.amazonaws.com

iam-roleaddpolicy -r s3access -e Allow -a s3:* -c \* -p s3star -o
{"Version": "2008-10-17", "Statement": [{"Effect": "Allow", "Action": ["s3:*"], "Resource": ["*"]}]}

iam-instanceprofilecreate -s s3access -r s3access
arn:aws:iam::111111111111:instance-profile/s3access
```

For more information, see [iam-rolecreate](#), [iam-roleaddpolicy](#), and [iam-instanceprofilecreate](#) in the *IAM Command Line Reference*.

2. Launch an instance using the instance profile. The following example shows a `t1.micro` instance being launched with the instance profile created in step 1.

```
ec2-run-instances -t t1.micro -p arn:aws:iam::111111111111:instance-profile/s3access -k key-pair -g 'Web Server' ami-e565ba8c
RESERVATION    r-11c62773    111111111111    sg-e7ddc68e
INSTANCE       i-9a6843fd    ami-e565ba8c
key-pair       0             us-east-1e     aki-88aa75e1    disabled

ec2-describe-instances
RESERVATION    r-11c62773    111111111111    sg-e7ddc68e
INSTANCE       i-9a6843fd    ami-e565ba8c
ec2-50-19-200-155.compute-1.amazonaws.com    ip-10-28-28-186.ec2.internal
running key-pair       0             t1.micro        2012-04-26T16:29:25.000Z
us-east-1e
aki-88aa75e1    disabled
AIPAJ6OQOSP4IRHXCI6E4
```

For more information, see [ec2-run-instances](#) and [ec2-describe-instances](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

3. If you are using the Amazon EC2 API actions in your application, retrieve the AWS security credentials made available on the instance and use them to sign the requests. Note that the AWS SDK does this for you.

```
GET http://169.254.169.254/latest/meta-data/iam/security-credentials/role_name
```

## Launching an Instance with an IAM Role using the AWS SDK

If you use an AWS SDK to write your application, you automatically get temporary security credentials from the role associated with the current instance. The AWS SDK documentation includes walkthroughs that show how an application can use security credentials from a IAM role to read an Amazon S3 bucket. For more information, see the following topics in the SDK documentation:

- [Using IAM Roles for EC2 Instances with the SDK for Java](#)
- [Using IAM Roles for EC2 Instances with the SDK for .NET](#)
- [Using IAM Roles for EC2 Instances with the SDK for PHP](#)
- [Using IAM Roles for EC2 Instances with the SDK for Ruby](#)

## Authorizing Inbound Traffic for Your Instances

To enable network access to your instance, you must allow inbound traffic to your instance on port 22 (SSH) or 3389 (RDP). To open a port for inbound traffic, add a rule to a security group that you associated with your instance when you launched it.

The following instructions authorize inbound SSH or RDP traffic, but only from your computer's public IP address. To allow traffic from additional IP address ranges, add a new security group rule for each range using these instructions.

### Before You Start

Decide who requires access to your instance; for example, a single host or a specific network that you trust. In this case, we use your local system's public IP address. You can get the public IP address of your local computer using a service. For example, we provide the following service: <http://checkip.amazonaws.com/>. To locate another service that provides your IP address, use the search phrase "what is my IP address". If you are connecting through an ISP or from behind your firewall without a static IP address, you need to find out the range of IP addresses used by client computers.

#### Caution

If you use `0.0.0.0/0`, you enable all IP addresses to access your instance using SSH or RDP. This is acceptable for a short time in a test environment, but it's unsafe for production environments. In production, you'll authorize only a specific IP address or range of addresses to access your instance.

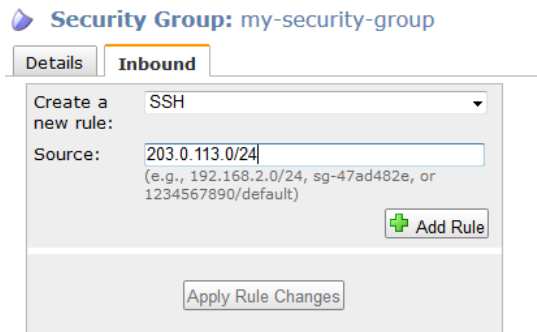
For more information about security groups, see [Amazon EC2 Security Groups \(p. 392\)](#).

## Adding a Rule for Inbound SSH Traffic

### To add a rule to a security group for inbound SSH traffic

1. In the navigation pane of the Amazon EC2 console, click **Instances**. Select your instance and look at the **Description** tab; **Security groups** lists the security groups that are associated with the instance. Click the **view rules** link to display a list of the rules that are in effect for the instance.

2. In the navigation pane, click **Security Groups**. Select one of the security groups associated with your instance.
3. In the details pane, on the **Inbound** tab, select **SSH** from the **Create a new rule** list.



4. In the **Source** box, specify the public IP address of your computer, in CIDR notation. For example, if your IP address is 203.0.113.25, specify 203.0.113.25/32 to list this single IP address in CIDR notation. If your company allocates addresses from a range, specify the entire range, such as 203.0.113.0/24.
5. Click **Add Rule**. An asterisk appears on the **Inbound** tab, indicating that the rule hasn't been applied.
6. When you're finished adding rules, click **Apply Rule Changes**. The new rules are applied to all instances that are associated with the security group.

If you prefer, you can use the `ec2-authorize` command as follows. Be sure to run this command on your local system, not on the instance itself.

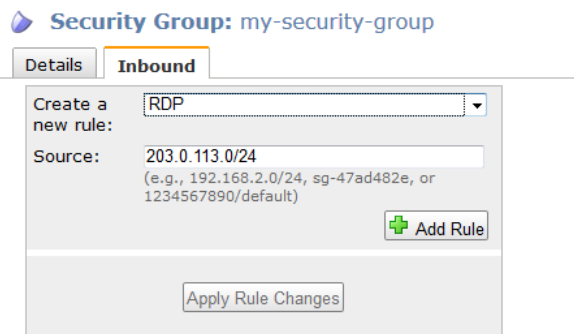
```
PROMPT> ec2-authorize group_name -p 22 -s your_ip_address/24
GROUP default
PERMISSION default ALLOWS tcp 22 22 FROM CIDR your_ip_address/24
```

## Adding a Rule for Inbound RDP Traffic

### To add a rule to a security group for inbound RDP traffic

1. In the navigation pane of the Amazon EC2 console, click **Instances**. Select your instance and look at the **Description** tab; **Security groups** lists the security groups that are associated with the instance. Click the **view rules** link to display a list of the rules that are in effect for the instance.
2. In the navigation pane, click **Security Groups**. Select one of the security groups associated with your instance.
3. In the details pane, on the **Inbound** tab, select **RDP** from the **Create a new rule** list.





4. In the **Source** box, specify the public IP address of your computer, in CIDR notation. For example, if your IP address is 203.0.113.25, specify 203.0.113.25/32 to list this single IP address in CIDR notation. If your company allocates addresses from a range, specify the entire range, such as 203.0.113.0/24.
5. Click **Add Rule**. An asterisk appears on the **Inbound** tab, indicating that the rule hasn't been applied.
6. When you're finished adding rules, click **Apply Rule Changes**. The new rules are applied to all instances that are associated with the security group.

If you prefer, you can use the [ec2-authorize](#) command as follows. Be sure to run this command on your local system, not on the instance itself.

```
PROMPT> ec2-authorize default -p 3389 -s your_ip_address/24
GROUP default
PERMISSION default ALLOWS tcp 3389 3389 FROM CIDR your_ip_address/24
```

## Assigning a Security Group to an Instance

You can assign a security group to an instance when you launch the instance. When you add or remove rules, those changes are automatically applied to all instances to which you've assigned the security group.

After you launch an instance in EC2-Classic, you can't change its security groups. After you launch an instance in a VPC, you can change its security groups. For more information, see [Changing an Instance's Security Groups](#) in the *Amazon Virtual Private Cloud User Guide*.

# Amazon EC2 and Amazon Virtual Private Cloud (VPC)

Amazon Virtual Private Cloud (Amazon VPC) enables you to define a virtual network in your own logically isolated area within the Amazon Web Services (AWS) cloud, known as a *virtual private cloud (VPC)*. You can launch your AWS resources, such as instances, into your VPC. Your VPC closely resembles a traditional network that you might operate in your own datacenter, with the benefits of using AWS's scalable infrastructure. You can configure your VPC; you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings. You can connect instances in your VPC to the Internet. You can connect your VPC to your own corporate datacenter, making the AWS cloud an extension of your datacenter. To protect the resources in each subnet, you can use multiple layers of security, including security groups and network access control lists. For more information, see [Amazon Virtual Private Cloud User Guide](#).

## Benefits of Using a VPC

By launching your instances into a VPC instead of EC2-Classic, you gain the ability to:

- Assign static private IP addresses to your instances that persist across starts and stops
- Assign multiple IP addresses to your instances
- Define network interfaces, and attach one or more network interfaces to your instances
- Change security group membership for your instances while they're running
- Control the outbound traffic from your instances (egress filtering) in addition to controlling the inbound traffic to them (ingress filtering)
- Add an additional layer of access control to your instances in the form of network access control lists (ACL)
- Run your instances on single-tenant hardware

## Differences Between EC2-Classic and EC2-VPC

Instances run in one of two supported platforms: EC2-Classic and EC2-VPC. Your AWS account is capable of launching instances either into both platforms or only into EC2-VPC, on a region by region basis. If you can launch instances only into EC2-VPC, we create a default VPC for you. A default VPC combines the benefits of the advanced features provided by EC2-VPC with the ease of use of EC2-Classic. For more information, see [Supported Platforms](#).

The following table summarizes the differences between instances launched in EC2-Classic, instances launched in a default VPC, and instances launched in a nondefault VPC.

Characteristic	EC2-Classic	Default VPC	Nondefault VPC
Public IP address (from Amazon's public IP address pool)	Your instance receives a public IP address.	Your instance launched in a default subnet receives a public IP address by default, unless you specify otherwise during launch.	Your instance doesn't receive a public IP address by default, unless you specify otherwise during launch.
Private IP address	Your instance receives a private IP address from the EC2-Classic range each time it's started.	Your instance receives a static private IP address from the address range of your default VPC.	Your instance receives a static private IP address from the address range of your VPC.
Multiple private IP addresses	We select a single private IP address for your instance; multiple IP addresses are not supported.	You can assign multiple private IP addresses to your instance.	You can assign multiple private IP addresses to your instance.
Elastic IP address	An EIP is disassociated from your instance when you stop it.	An EIP remains associated with your instance when you stop it.	An EIP remains associated with your instance when you stop it.
DNS hostnames	DNS hostnames are enabled by default.	DNS hostnames are enabled by default.	DNS hostnames are disabled by default.

Characteristic	EC2-Classical	Default VPC	Nondefault VPC
Security group	<p>A security group can reference security groups that belong to other AWS accounts.</p> <p>You can create up to 500 security groups in each region.</p>	<p>A security group can reference security groups for your VPC only.</p> <p>You can create up to 100 security groups per VPC.</p>	<p>A security group can reference security groups for your VPC only.</p> <p>You can create up to 100 security groups per VPC.</p>
Security group association	<p>You can assign an unlimited number of security groups to an instance when you launch it.</p> <p>You can't change the security groups of your running instance. You can either modify the rules of the assigned security groups, or replace the instance with a new one (create an AMI from the instance, launch a new instance from this AMI with the security groups that you need, disassociate any Elastic IP address from the original instance and associate it with the new instance, and then terminate the original instance).</p>	<p>You can assign up to 5 security groups to an instance.</p> <p>You can assign security groups to your instance when you launch it and while it's running.</p>	<p>You can assign up to 5 security groups to an instance.</p> <p>You can assign security groups to your instance when you launch it and while it's running.</p>
Security group rules	<p>You can add rules for inbound traffic only.</p> <p>You can add up to 100 rules to a security group.</p>	<p>You can add rules for inbound and outbound traffic.</p> <p>You can add up to 50 rules to a security group.</p>	<p>You can add rules for inbound and outbound traffic.</p> <p>You can add up to 50 rules to a security group.</p>
Tenancy	<p>Your instance runs on shared hardware.</p>	<p>You can run your instance on shared hardware or single-tenant hardware.</p>	<p>You can run your instance on shared hardware or single-tenant hardware.</p>

## Amazon VPC Documentation

For more information about Amazon VPC, see the Amazon VPC documentation.

Guide	Description
<a href="#">Amazon Virtual Private Cloud Getting Started Guide</a>	Provides a hands-on introduction to Amazon VPC.
<a href="#">Amazon Virtual Private Cloud User Guide</a>	Provides detailed information about how to use Amazon VPC.

Guide	Description
<a href="#">Amazon Virtual Private Cloud Network Administrator Guide</a>	Helps network administrators configure your customer gateway.

## Supported Platforms

Amazon EC2 supports the following platforms. Your AWS account is capable of launching instances either into both platforms or only into EC2-VPC, on a region by region basis.

Platform	Introduced In	Description
EC2-Classic	The original release of Amazon EC2	Your instances run in a single, flat network that you share with other customers.
EC2-VPC	The original release of Amazon VPC	Your instances run in a virtual private cloud (VPC) that's logically isolated to your AWS account.

## Supported Platforms in the Amazon EC2 Console

The Amazon EC2 console indicates which platforms you can launch instances into for the selected region, and whether you have a default VPC in that region.

Verify that the region you'll use is selected in the navigation bar. On the Amazon EC2 console dashboard, look for **Supported Platforms** under **Account Attributes**. If there are two values, `EC2-Classic` and `EC2-VPC`, you can launch instances into either platform. If there is one value, `EC2-VPC`, you can launch instances only into EC2-VPC.

If you can launch instances only into EC2-VPC, we create a default VPC for you. Then, when you launch an instance, we launch it into your default VPC, unless you create a nondefault VPC and specify it when you launch the instance.

### EC2-VPC

The dashboard displays the following under **Account Attributes** to indicate that the account supports only the EC2-VPC platform, and has a default VPC with the identifier `vpc-1a2b3c4d`.



Supported Platforms

EC2-VPC

Default VPC

vpc-1a2b3c4d

If your account supports only EC2-VPC, you can select a VPC from the **Network** list, and a subnet from the **Subnet** list when you launch an instance using the launch wizard.

Network ⓘ	vpc-1a2b3c4d (172.31.0.0/16) (default) 	Create new VPC
Subnet ⓘ	No preference (default subnet in any Availability Zor) 	Create new subnet

### EC2-Classic, EC2-VPC

The dashboard displays the following under **Account Attributes** to indicate that the account supports both the EC2-Classic and EC2-VPC platforms.

#### Supported Platforms

EC2-Classic  
EC2-VPC

If your account supports EC2-Classic and EC2-VPC, you can launch into EC2-Classic using the launch wizard by selecting **Launch into EC2-Classic** from the **Network** list. To launch into a VPC, you can select a VPC from the **Network** list, and a subnet from the **Subnet** list.

### Related Topic

For more information about how you can tell which platforms you can launch instances into, see [Detecting Your Supported Platforms](#) in the *Amazon Virtual Private Cloud User Guide*.


## Differences Between Instances in EC2-Classic and EC2-VPC

With EC2-Classic, we assign each instance a private IP address from a shared private IP address range. We also assign each instance a public IP address from Amazon's pool of public IP addresses. Instances access the Internet directly through the AWS network edge.

With EC2-VPC, we assign each instance a private IP address from the private IP address range of your VPC. You can control the IP address range, subnets, routing, network gateways, network ACLs, and security groups for your VPC. You can specify whether your instance receives a public IP address during launch. Instances with public IP addresses or Elastic IP addresses can access the Internet through a logical Internet gateway attached to the AWS network edge. For more information about EC2-VPC, see [What is Amazon VPC?](#) in the *Amazon Virtual Private Cloud User Guide*.

The following diagram shows instances in each platform. Note the following:

- Instances C1, C2, C3, and C4 are in the EC2-Classic platform. C1 and C2 were launched by one account, and C3 and C4 were launched by a different account. These instances can communicate with each other, can access the Internet directly, and can access other Amazon Web Services such as Amazon Simple Storage Service (Amazon S3).
- Instances V1 and V2 are in different subnets in the same VPC in the EC2-VPC platform. They were launched by the account that owns the VPC; no other account can launch instances in this VPC. These instances can communicate with each other and can access the following through the Internet gateway: instances in EC2-Classic, other Amazon Web Services (such as Amazon S3), and the Internet.



For more information about the differences between EC2-Classic and EC2-VPC, see [Amazon EC2 and Amazon Virtual Private Cloud \(VPC\)](#) (p. 414).

## Amazon EC2 Instance IP Addressing

We provide your instances with IP addresses and DNS hostnames. These can vary depending on whether you launched the instance in the EC2-Classic platform or in a virtual private cloud (VPC).

For information about the EC2-Classic and EC2-VPC platforms, see [Supported Platforms](#) (p. 417). For information about Amazon VPC, see [What is Amazon VPC?](#) in the *Amazon Virtual Private Cloud User Guide*.

### Topics

- [Private Addresses and Internal DNS Hostnames \(p. 420\)](#)
- [Public IP Addresses and External DNS Hostnames \(p. 420\)](#)
- [Differences Between EC2-Classic and EC2-VPC \(p. 421\)](#)
- [Determining Your Public, Private, and Elastic IP Addresses \(p. 422\)](#)
- [Assigning a Public IP Address \(p. 423\)](#)
- [Multiple IP Addresses \(p. 423\)](#)

## Private Addresses and Internal DNS Hostnames

You can use private IP addresses and internal DNS hostnames for communication between instances in the same network (EC2-Classic or a VPC). Private IP addresses are not reachable from the Internet. For more information about private IP addresses, see [RFC 1918](#).

When you launch an instance, we allocate a private IP address for the instance using DHCP.

Each instance that you launch into a VPC has a default network interface. The network interface specifies the primary private IP address for the instance. If you don't select a primary private IP address, we select an available IP address in the subnet's range. You can specify additional private IP addresses, known as secondary private IP addresses. Unlike primary private IP addresses, secondary private IP addresses can be reassigned from one instance to another. For more information, see [Multiple IP Addresses \(p. 423\)](#).

Each instance is provided an internal DNS hostname that resolves to the private IP address of the instance in EC2-Classic or your VPC. We can't resolve the DNS hostname outside the network that the instance is in.

If you create a custom firewall configuration in EC2-Classic, you must allow inbound traffic from port 53 (with a destination port from the ephemeral range) from the address of the Amazon DNS server; otherwise, internal DNS resolution from your instances fails. If your firewall doesn't automatically allow DNS query responses, then you'll need to allow traffic from the IP address of the Amazon DNS server. To get the IP address of the Amazon DNS server on Linux, use the following command: **grep nameserver /etc/resolv.conf**. To get the IP address of the Amazon DNS Server on Windows, use the following command: **ipconfig /all | findstr /c:"DNS Servers"**.

For instances launched in EC2-Classic, a private IP address is associated with the instance until it is stopped or terminated.

For instances launched in a VPC, a private IP address remains associated with the network interface when the instance is stopped and restarted, and is released when the instance is terminated.

## Public IP Addresses and External DNS Hostnames

You can use public IP addresses and external DNS hostnames for communication between your instances and the Internet or other AWS products, such as Amazon Simple Storage Service (Amazon S3). Public IP addresses are reachable from the Internet.

When you launch an instance in EC2-Classic, we automatically allocate a public IP address for the instance. When you launch an instance into EC2-VPC, you can control whether your instance receives a public IP address. The public IP address is associated with the eth0 network interface (the primary network interface).

A public IP address is assigned to your instance from Amazon's pool of public IP addresses, and is not associated with your AWS account. When a public IP address is disassociated from your instance, it is released back into the public IP address pool, and you cannot reuse it.

You cannot manually associate or disassociate a public IP address from your instance. Instead, in certain cases, we release the public IP address from your instance, or assign it a new one:

- We release the public IP address for your instance when it's stopped or terminated. Your stopped instance receives a new public IP address when it's restarted.
- We release the public IP address for your instance when you associate an Elastic IP address (EIP) with your instance, or when you associate an EIP with the primary network interface (eth0) of your instance in a VPC. When you disassociate the EIP from your instance, it receives a new public IP address.
- If the public IP address of your instance in a VPC has been released, it will not receive a new one if there is more than one network interface attached to your instance.

If you require a persistent public IP address that can be associated to and from instances as you require, use an Elastic IP address (EIP) instead. You can allocate your own EIP, and associate it to your instance. For more information, see [Elastic IP Addresses \(EIP\)](#) (p. 428).

We don't automatically assign a public IP address to an instance that you launch in a nondefault subnet. Therefore, if you want an instance in a nondefault subnet to communicate with the Internet, you must either enable the public IP addressing feature during launch, or associate an Elastic IP address with the primary or any secondary private IP address assigned to the network interface for the instance.

We provide each instance that has a public IP address with an external DNS hostname. We resolve an external DNS hostname to the public IP address of the instance outside the network of the instance, and to the private IP address of the instance from within the network of the instance. If your instance is in a VPC and you assign it an Elastic IP address, it receives a DNS hostname if DNS hostnames are enabled. For more information, see [Using DNS with Your VPC](#) in the *Amazon Virtual Private Cloud User Guide*.

The private IP address and public IP address for an instance are directly mapped to each other through network address translation (NAT). For more information about NAT, see [RFC 1631: The IP Network Address Translator \(NAT\)](#).

**Note**

Instances that access other instances through their public NAT IP address are charged for regional or Internet data transfer, depending on whether the instances are in the same region.

## Differences Between EC2-Classical and EC2-VPC

The following table summarizes the differences between IP addresses for instances launched in EC2-Classical, instances launched in a default subnet, and instances launched in a nondefault subnet.

Characteristic	EC2-Classical	Default Subnet	Nondefault Subnet
Public IP address (from Amazon's public IP address pool)	Your instance receives a public IP address.	Your instance launched in a default subnet receives a public IP address by default, unless you specify otherwise during launch.	Your instance doesn't receive a public IP address by default, unless you specify otherwise during launch.
Private IP address	Your instance receives a private IP address from the EC2-Classical range each time it's started.	Your instance receives a static private IP address from the address range of your default VPC.	Your instance receives a static private IP address from the address range of your VPC.



**Amazon Elastic Compute Cloud User Guide**  
**Determining Your Public, Private, and Elastic IP**  
**Addresses**

Characteristic	EC2-Classical	Default Subnet	Nondefault Subnet
Multiple IP addresses	We select a single private IP address for your instance; multiple IP addresses are not supported.	You can assign multiple private IP addresses to your instance.	You can assign multiple private IP addresses to your instance.
Network interfaces	IP addresses are associated with the instance; network interfaces aren't supported.	IP addresses are associated with a network interface. Each instance has one or more network interfaces.	IP addresses are associated with a network interface. Each instance has one or more network interfaces.
Elastic IP address	An EIP is disassociated from your instance when you stop it.	An EIP remains associated with your instance when you stop it.	An EIP remains associated with your instance when you stop it.
DNS hostnames	DNS hostnames are enabled by default.	DNS hostnames are enabled by default.	DNS hostnames are disabled by default.

## Determining Your Public, Private, and Elastic IP Addresses

You can use the EC2 console to determine the private IP addresses, public IP addresses, and EIPs of your instances.

### To determine your instance's IP addresses using the Amazon EC2 console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Instances** in the navigation pane.
3. Select an instance. The console displays information about the instance in the lower pane.
4. If an EIP has been associated with the instance, get its address from the **Elastic IP** field. If no EIP has been associated with the instance, you can get the public IP address from the **Public DNS** field.
5. Get the private IP address from the **Private IP** field.

You can also determine the public and private IP addresses of your instances using instance metadata. For more information about instance metadata, see [Instance Metadata and User Data \(p. 290\)](#).

### To determine your instance's IP addresses using instance metadata

1. Connect to the instance.
2. Use the following command to access the private IP address:

```
GET http://169.254.169.254/latest/meta-data/local-ipv4
```

3. Use the following command to access the public IP address:

```
GET http://169.254.169.254/latest/meta-data/public-ipv4
```

Note that if an EIP is associated with the instance, the value returned is that of the EIP.

## Assigning a Public IP Address

If you launch an instance in EC2-Classic, it is assigned a public IP address by default. You can't modify this behavior.

If you launch an instance into a VPC, a public IP addressing feature is available for you to control whether your instance is assigned a public IP address. The public IP address is assigned to the network interface with the device index of eth0. This feature depends on certain conditions at the time you launch your instance.

### Important

You can't manually disassociate the public IP address from your instance after launch. Instead, it's automatically released in certain cases, after which you cannot reuse it. For more information, see [Public IP Addresses and External DNS Hostnames](#) (p. 420). If you require a persistent public IP address that you can associate or disassociate at will, assign an Elastic IP address to the instance after launch instead. For more information, see [Elastic IP Addresses \(EIP\)](#) (p. 428).

### To access the public IP addressing feature when launching an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Launch Instance**.
3. Choose an AMI and click its **Select** button, then choose an instance type and click **Next: Configure Instance Details**.
4. On the **Configure Instance Details** page, if a VPC is selected in the **Network** list, a **Public IP** check box is displayed. This check box, if selected, will assign a public IP address to your instance. If you selected a default subnet, the **Public IP** check box is selected by default.

The following rules apply:

- A public IP address can only be assigned to a single network interface with the device index of eth0. The **Public IP** check box is not available if you're launching with multiple network interfaces, and is not available for the eth1 network interface.
- You can only assign a public IP address to a new network interface, not an existing one.

This feature is only available during launch. However, whether you assign a public IP address to your instance during launch or not, you can associate an Elastic IP address with your instance after it's launched. For more information, see [Elastic IP Addresses \(EIP\)](#) (p. 428).

## API and Command Line Tools for Public IP Addressing

To enable or disable the public IP addressing feature, use the `NetworkInterface.n.AssociatePublicIpAddress` parameter with the [RunInstances](#) request, or the `--associate-public-ip-address` option with the [ec2-run-instances](#) command.

## Multiple IP Addresses

In EC2-VPC, you can specify multiple private IP addresses for your instances. The number of network interfaces and private IP addresses that you can specify for an instance depends on the instance type. For more information, see [Private IP Addresses Per ENI Per Instance Type](#) (p. 433).

It can be useful to assign multiple private IP addresses to an instance in your VPC to do the following:

- Host multiple websites on a single server by using multiple SSL certificates on a single server and associating each certificate with a specific IP address.

- Operate network appliances, such as firewalls or load balancers, that have multiple private IP addresses for each network interface.
- Redirect internal traffic to a standby instance in case your instance fails, by reassigning the secondary private IP address to the standby instance.

#### Topics

- [How Multiple IP Addresses Work \(p. 424\)](#)
- [Assigning a Secondary Private IP Address \(p. 424\)](#)
- [Configuring the Operating System on Your Instance to Recognize the Secondary Private IP Address \(p. 426\)](#)
- [Assigning an Elastic IP Address to the Secondary Private IP Address \(p. 426\)](#)
- [Viewing Your Secondary Private IP Addresses \(p. 426\)](#)
- [Unassigning a Secondary Private IP Address \(p. 427\)](#)

## How Multiple IP Addresses Work

The following list explains how multiple IP addresses work with network interfaces:

- You can assign a secondary private IP address to any network interface. The network interface can be attached to or detached from the instance.
- You must choose a secondary private IP address that's in the CIDR block range of the subnet for the network interface.
- Security groups apply to network interfaces, not to IP addresses. Therefore, IP addresses are subject to the security group of the network interface in which they're specified.
- Secondary private IP addresses can be assigned and unassigned to elastic network interfaces attached to running or stopped instances.
- Secondary private IP addresses that are assigned to a network interface can be reassigned to another one if you explicitly allow it.
- When assigning multiple secondary private IP addresses to a network interface using the command line tools or API, the entire operation fails if one of the secondary private IP addresses can't be assigned.
- Primary private IP addresses, secondary private IP addresses, and any associated Elastic IP addresses remain with the network interface when it is detached from an instance or attached to another instance.
- Although you can't move the primary network interface from an instance, you can reassign the secondary private IP address of the primary network interface to another network interface.
- You can move any additional network interface from one instance to another.

The following list explains how multiple IP addresses work with Elastic IP addresses:

- Each private IP address can be associated with a single Elastic IP address, and vice versa.
- When a secondary private IP address is reassigned to another interface, the secondary private IP address retains its association with an Elastic IP address.
- When a secondary private IP address is unassigned from an interface, an associated Elastic IP address is automatically disassociated from the secondary private IP address.

## Assigning a Secondary Private IP Address

You can assign the secondary private IP address to the network interface for an instance as you launch the instance, or after the instance is running.

### To assign a secondary private IP address when launching an instance in EC2-VPC

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click the **Launch Instance** button.
3. Choose an AMI and click its **Select** button, then choose an instance type and click **Next: Configure Instance Details**.
4. On the **Configure Instance Details** page, choose a VPC from the **Network** list, and a subnet from the **Subnet** list.
5. In the **Network Interfaces** section, do the following, and then click **Next: Add Storage**:
  - a. Click **Add Device** to add another network interface. The console enables you specify up to 2 network interfaces when you launch an instance. After you launch the instance, click **Network Interfaces** in the navigation pane to add additional network interfaces. The total number of network interfaces that you can attach varies by instance type. For more information, see [Private IP Addresses Per ENI Per Instance Type](#) (p. 433).
  - b. For each network interface, you can specify a primary private IP address, and one or more secondary private IP addresses. For this example, however, accept the IP address that we automatically assign.
  - c. Under **Secondary IP addresses**, click **Add IP**, and then enter a private IP address in the subnet range, or accept the default, `Auto-assign`, to let us select an address.

#### Important

After you have added a secondary private IP address to a network interface, you must connect to the instance and configure the secondary private IP address on the instance itself. For more information, see [Configuring the Operating System on Your Instance to Recognize the Secondary Private IP Address](#) (p. 426).

6. On the next **Add Storage** page, you can specify volumes to attach to the instance besides the volumes specified by the AMI (such as the root device volume), and then click **Next: Tag Instance**.
7. On the **Tag Instance** page, specify tags for the instance, such as a user-friendly name, and then click **Next: Configure Security Group**.
8. On the **Configure Security Group** page, select an existing security group or create a new one. Click **Review and Launch**.
9. On the **Review Instance Launch** page, review your settings, and then click **Launch** to choose a key pair and launch your instance. If you're new to Amazon EC2 and haven't created any key pairs, the wizard prompts you to create one.

### To assign a secondary private IP to an existing instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, click **Network Interfaces**, and then right-click the network interface attached to the instance.
3. Select **Manage Private IP Addresses**.
4. In the **Manage Private IP Addresses** dialog box, do the following:
  - a. Click **Assign a secondary private address**.
  - b. In the **Address** field, enter a specific IP address that's within the subnet range for the instance, or leave the field blank and we'll select an IP address for you.
  - c. (Optional) Select **Allow reassignment** to allow the secondary private IP address to be reassigned if it is already assigned to another network interface.
  - d. Click **Yes, Update**, and then click **Close**.

**Note**

You can also assign a secondary private IP address to an instance by clicking **Instances** in the navigation pane, right-clicking your instance, and selecting **Manage Private IP Addresses**. You can configure the same information in the dialog as you did in the steps above.

## Configuring the Operating System on Your Instance to Recognize the Secondary Private IP Address

After you assign a secondary private IP address to your instance, you need to configure the operating system on your instance to recognize the secondary private IP address.

If you are using Amazon Linux, the `ec2-net-utils` package can take care of this step for you. It configures additional network interfaces that you attach while the instance is running, refreshes secondary IP addresses during DHCP lease renewal, and updates the related routing rules. If you require manual control over your network configuration, you can remove the `ec2-net-utils` package. For more information, see [Configuring Your Network Interface Using `ec2-net-utils` \(p. 435\)](#).

If you are using another Linux distribution, see the documentation for your Linux distribution. Search for information about configuring additional network interfaces and secondary IP addresses. If the instance has two or more interfaces on the same subnet, search for information about using routing rules to work around asymmetric routing.

For information about configuring a Windows instance, see [Configuring a Secondary Private IP Address for Your Windows Instance](#) in the *Amazon Elastic Compute Cloud Microsoft Windows Guide*.

## Assigning an Elastic IP Address to the Secondary Private IP Address

### To assign an EIP to a secondary private IP address in EC2-VPC

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Elastic IPs** in the navigation pane.
3. Right-click the IP address, and then click **Associate**.
4. In the **Associate Address** dialog box, select the network interface from the **Network Interface** drop-down list, and then select the secondary IP address from the **Private IP address** drop-down list.
5. Click **Associate**.

## Viewing Your Secondary Private IP Addresses

### To view the private IP addresses assigned to a network interface in EC2-VPC

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Select the network interface whose private IP addresses you want to view.
4. On the **Details** tab in the details pane, check the **Primary IP** and **Secondary Private IPs** fields for the primary private IP address and any secondary private IP addresses assigned to the network interface.

### To view the private IP addresses assigned to an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Instances** in the navigation pane.
3. Select the instance whose private IP addresses you want to view.
4. On the **Description** tab in the details pane, check the **Private IPs** and **Secondary private IPs** fields for the primary private IP address and any secondary private IP addresses assigned to the instance through its network interface.

## Unassigning a Secondary Private IP Address

If you no longer require a secondary private IP address, you can unassign it from the instance or the network interface. When a secondary private IP address is unassigned from an elastic network interface, the Elastic IP address (if it exists) is also disassociated.

### To unassign a secondary private IP address from an instance


1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Instances** in the navigation pane.
3. Right-click an instance, and then click **Manage Private IP Addresses**.
4. In the **Manage Private IP Addresses** dialog box, beside the secondary private IP address to unassign, click **Unassign**.
5. Click **Yes, Update**, and then close the dialog box.

### To unassign a secondary private IP address from a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interface** in the navigation pane.
3. Right-click an instance, and then click **Manage Private IP Addresses**.
4. In the **Manage Private IP Addresses** dialog box, beside the secondary private IP address to unassign, click **Unassign**.
5. Click **Yes, Update**, and then click **Close**.

**Manage Private IP Addresses** Cancel X

You can assign and unassign secondary private IP addresses on this network interface. Leave the address field blank and an available address will be assigned or enter an IP address that you want to assign.

▼  eth0: eni-ea67dc83 - Primary network interface - 10.0.0.0/24

10.0.0.174	46.51.219.123	Primary IP
<del>10.0.0.34</del>		Undo

[Assign a secondary private address](#)

Allow reassignment

Are you sure you want to perform the following changes?

- 1 private IP address will be unassigned from eni-ea67dc83

Close Yes, Update

## Elastic IP Addresses (EIP)

An *Elastic IP address* (EIP) is a static IP address designed for dynamic cloud computing. With an EIP, you can mask the failure of an instance by rapidly remapping the address to another instance. Your EIP is associated with your AWS account, not a particular instance, and it remains associated with your account until you choose to explicitly release it.

There's one pool of EIPs for use with the EC2-Classic platform and another for use with your VPC. You can't associate an EIP that you allocated for use with a VPC with an instance in EC2-Classic, and vice-versa. For more information about EC2-Classic and EC2-VPC, see [Supported Platforms \(p. 417\)](#).

### Topics

- [Elastic IP Addresses in EC2-Classic \(p. 428\)](#)
- [Elastic IP Addresses in a VPC \(p. 428\)](#)
- [Differences Between EC2-Classic and EC2-VPC \(p. 429\)](#)
- [Allocating an Elastic IP Address \(p. 429\)](#)
- [Describing Your Elastic IP Addresses \(p. 430\)](#)
- [Associating an Elastic IP Address with a Running Instance \(p. 430\)](#)
- [Associating an Elastic IP Address with a Different Running Instance \(p. 430\)](#)
- [Releasing an Elastic IP Address \(p. 430\)](#)
- [Using Reverse DNS for Email Applications \(p. 431\)](#)
- [API and CLI Overview \(p. 431\)](#)
- [Elastic IP Address Limit \(p. 431\)](#)

## Elastic IP Addresses in EC2-Classic

By default, we assign each instance in EC2-Classic two IP addresses at launch: a private IP address and a public IP address that is mapped to the private IP address through network address translation (NAT). The public IP address is allocated from the EC2-Classic public IP address pool, and is associated with your instance, not with your AWS account. You cannot reuse a public IP address after it's been disassociated from your instance.

If you use dynamic DNS to map an existing DNS name to a new instance's public IP address, it might take up to 24 hours for the IP address to propagate through the Internet. As a result, new instances might not receive traffic while terminated instances continue to receive requests. To solve this problem, use an EIP.

When you associate an EIP with an instance, the instance's current public IP address is released to the EC2-Classic public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. In addition, stopping the instance also disassociates the EIP from it.

To ensure efficient use of EIPs, we impose a small hourly charge when they are not associated with a running instance, or when they are associated with a stopped instance.

## Elastic IP Addresses in a VPC

We assign each instance in a default VPC two IP addresses at launch: a private IP address and a public IP address that is mapped to the private IP address through network address translation (NAT). The public IP address is allocated from the EC2-VPC public IP address pool, and is associated with your instance, not with your AWS account. You cannot reuse a public IP address after it's been disassociated from your instance.



We assign each instance in a nondefault VPC only a private IP address, unless you specifically request a public IP address during launch. To ensure that an instance in a nondefault VPC that has not been assigned a public IP address can communicate with the Internet, you must allocate an Elastic IP address for use with a VPC, and then associate that EIP with the elastic network interface (ENI) attached to the instance.

When you associate an EIP with an instance in a default VPC, or an instance in which you assigned a public IP to the eth0 network interface during launch, its current public IP address is released to the EC2-VPC public IP address pool. If you disassociate an EIP from the instance, the instance is automatically assigned a new public IP address within a few minutes. However, if you have attached a second network interface to the instance, the instance is not automatically assigned a new public IP address; you'll have to associate an EIP with it manually. The EIP remains associated with the instance when you stop it.

To ensure efficient use of EIPs, we impose a small hourly charge when they are not associated with a running instance, or when they are associated with a stopped instance or an unattached network interface. While your instance is running, you are not charged for one EIP associated with the instance, but you are charged for any additional EIPs associated with the instance.

For information about using an EIP with an instance in a VPC, see [Elastic IP Addresses](#) in the *Amazon Virtual Private Cloud User Guide*.

## Differences Between EC2-Classic and EC2-VPC

The following table lists the differences between EIPs on EC2-Classic and EC2-VPC.

Characteristic	EC2-Classic	EC2-VPC
Allocation	When you allocate an EIP, it's for use only in EC2-Classic.	When you allocate an EIP, it's for use only in a VPC.
Association	You associate an EIP with an instance.	An EIP is a property of an elastic network interface (ENI). You can associate an EIP with an instance by updating the ENI attached to the instance. For more information, see <a href="#">Elastic Network Interfaces (ENI) (p. 431)</a> .
Reassociation	If you try to associate an EIP that's already associated with another instance, the address is automatically associated with the new instance.	If you try to associate an EIP that's already associated with another instance, it succeeds only if you allowed reassociation.
Instance stop	If you stop an instance, its EIP is disassociated, and you must re-associate the EIP when you restart the instance.	If you stop an instance, its EIP remains associated.
Multiple IP	Instances support only a single private IP address and a corresponding EIP.	Instances support multiple IP addresses, and each one can have a corresponding EIP. For more information, see <a href="#">Multiple IP Addresses (p. 423)</a> .

## Allocating an Elastic IP Address

### To allocate an EIP for use with EC2-Classic

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.



2. Click **Elastic IPs** in the navigation pane.
3. Click **Allocate New Address**.
4. Select **EC2** from the **EIP** list, and then click **Yes, Allocate**.

## Describing Your Elastic IP Addresses

### To view your Elastic IP addresses

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Elastic IPs** in the navigation pane.
3. To filter the displayed list, start typing part of the EIP or the ID of the instance to which it is assigned in the search box.

## Associating an Elastic IP Address with a Running Instance

### To associate an Elastic IP address with a running instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Elastic IPs** in the navigation pane.
3. Select an EIP and click **Associate Address**.
4. In the **Associate Address** dialog box, select the instance from the **Instance** list box and click **Associate**.

## Associating an Elastic IP Address with a Different Running Instance

### To reassociate an Elastic IP address

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Elastic IPs** in the navigation pane.
3. Select the EIP, and then click the **Disassociate** button.
4. Click **Yes, Disassociate** when prompted.
5. Select the EIP, and then click **Associate**.
6. In the **Associate Address** dialog box, select the new instance from the **Instance** list, and then click **Associate**.

## Releasing an Elastic IP Address

If you no longer need an EIP, we recommend that you release it (the address must not be associated with an instance). You incur charges for any EIP that's allocated for use with EC2-Classic but not associated with an instance.

### To release an Elastic IP address

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Elastic IPs** in the navigation pane.

3. Select the Elastic IP address, click the **Release Address** button, and then click **Yes, Release** when prompted.

## Using Reverse DNS for Email Applications

If you intend to send email to third parties from an instance, we suggest you provision one or more Elastic IP addresses and provide them to us in the [Request to Remove Email Sending Limitations form](#). AWS works with ISPs and Internet anti-spam organizations (such as Spamhaus) to reduce the chance that your email sent from these addresses will be flagged as spam.

In addition, assigning a static reverse DNS record to your Elastic IP address used to send email can help avoid having email flagged as spam by some anti-spam organizations. You can provide us with a reverse DNS record to associate with your addresses through the aforementioned form. Note that a corresponding forward DNS record (A Record) pointing to your Elastic IP address must exist before we can create your reverse DNS record.

## API and CLI Overview

The following table summarizes the available Elastic IP address command line interface (CLI) and corresponding API actions.

Description	CLI Command	API Action
Acquires an Elastic IP address for your account.	<a href="#">ec2-allocate-address</a>	<a href="#">AllocateAddress</a>
Associates an Elastic IP address with an instance or a network interface.	<a href="#">ec2-associate-address</a>	<a href="#">AssociateAddress</a>
Describes one or more of your Elastic IP addresses.	<a href="#">ec2-describe-addresses</a>	<a href="#">DescribeAddresses</a>
Disassociates an Elastic IP address from the instance or network interface it's associated with.	<a href="#">ec2-disassociate-address</a>	<a href="#">DisassociateAddress</a>
Releases an Elastic IP address allocated to your account.	<a href="#">ec2-release-address</a>	<a href="#">ReleaseAddress</a>

## Elastic IP Address Limit

By default, all AWS accounts are limited to 5 EIPs, because public (IPv4) Internet addresses are a scarce public resource. We strongly encourage you to use an EIP primarily for load balancing use cases, and use DNS hostnames for all other inter-node communication.

If you feel your architecture warrants additional EIPs, please complete the [Amazon EC2 Elastic IP Address Request Form](#). We will ask you to describe your use case so that we can understand your need for additional addresses.

## Elastic Network Interfaces (ENI)

An elastic network interface (ENI) is a virtual network interface that you can attach to an instance in a VPC. An ENI can include the following attributes:

- a primary private IP address
- one or more secondary private IP addresses
- one Elastic IP address per private IP address
- a MAC address
- one or more security groups
- a source/destination check flag
- a description

You can create a network interface, attach it to an instance, detach it from an instance, and attach it to another instance. The attributes of a network interface follow the network interface as it is attached or detached from an instance and reattached to another instance. When you move a network interface from one instance to another, network traffic is redirected to the new instance.

Each instance in a VPC has a default network interface. The default network interface has a primary private IP address in the IP address range of its VPC. You can create and attach additional network interfaces. The maximum number of network interfaces that you can use varies by instance type. For more information, see [Private IP Addresses Per ENI Per Instance Type \(p. 433\)](#).

Attaching multiple network interfaces to an instance is useful when you want to:

- Create a management network.
- Use network and security appliances in your VPC.
- Create dual-homed instances with workloads/roles on distinct subnets.
- Create a low-budget, high-availability solution.

### Topics

- [Private IP Addresses Per ENI Per Instance Type \(p. 433\)](#)
- [Creating a Management Network \(p. 433\)](#)
- [Use Network and Security Appliances in Your VPC \(p. 434\)](#)
- [Creating Dual-homed Instances with Workloads/Roles on Distinct Subnets \(p. 434\)](#)
- [Create a Low Budget High Availability Solution \(p. 434\)](#)
- [Best Practices for Configuring Network Interfaces \(p. 435\)](#)
- [Configuring Your Network Interface Using ec2-net-utils \(p. 435\)](#)
- [Creating a Network Interface \(p. 436\)](#)
- [Deleting a Network Interface \(p. 437\)](#)
- [Viewing Details about a Network Interface \(p. 437\)](#)
- [Attaching a Network Interface When Launching an Instance \(p. 437\)](#)
- [Attaching a Network Interface to a Stopped or Running Instance \(p. 438\)](#)
- [Detaching a Network Interface from an Instance \(p. 439\)](#)
- [Changing the Security Group of a Network Interface \(p. 439\)](#)
- [Changing the Source/Destination Checking of a Network Interface \(p. 440\)](#)
- [Associating an Elastic IP Address with a Network Interface \(p. 440\)](#)
- [Disassociating an Elastic IP Address from a Network Interface \(p. 441\)](#)
- [Changing Termination Behavior for a Network Interface \(p. 442\)](#)
- [Adding or Editing a Description for a Network Interface \(p. 442\)](#)
- [Adding or Editing Tags for a Network Interface \(p. 443\)](#)
- [API and Command Overview \(p. 443\)](#)

## Private IP Addresses Per ENI Per Instance Type

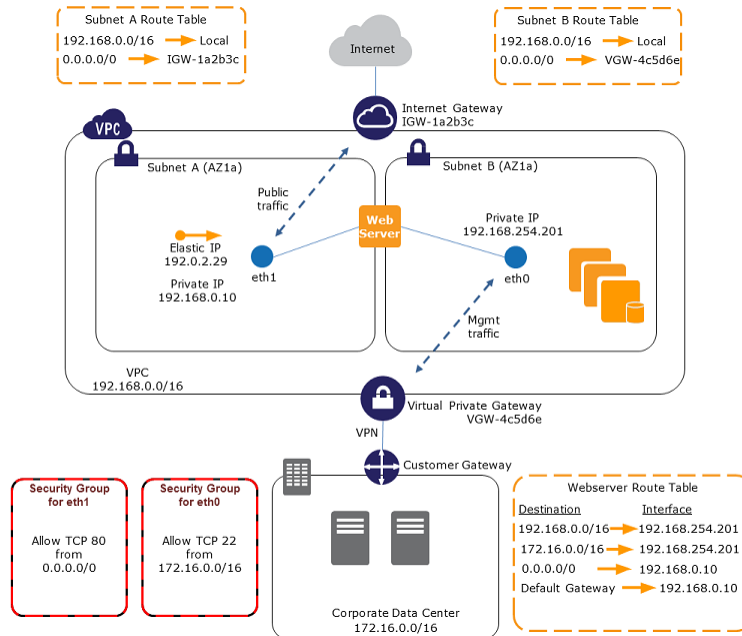
The following table lists the maximum number of elastic network interfaces (ENI) per instance type, and the maximum number of private IP addresses per ENI. ENIs and multiple private IP addresses are only available for instances running in a VPC. For more information, see [Multiple IP Addresses](#) (p. 423).

Instance Type	Maximum ENIs	IP Addresses per ENI
cc2.8xlarge	8	30
cg1.4xlarge	8	30
c1.xlarge	4	15
c1.medium	2	6
hi1.4xlarge	8	30
m2.2xlarge	4	30
m2.xlarge	4	15
m2.4xlarge	8	30
cr1.8xlarge	8	30
hs1.8xlarge	8	30
m1.xlarge	4	15
m1.large	3	10
m1.medium	2	6
m1.small	2	4
m3.2xlarge	4	30
m3.xlarge	4	15
t1.micro	2	2

## Creating a Management Network

You can create a management network using network interfaces. In this scenario, the secondary network interface on the instance handles public-facing traffic and the primary network interface handles back-end management traffic and is connected to a separate subnet in your VPC that has more restrictive access controls. The public facing interface, which may or may not be behind a load balancer, has an associated security group that allows access to the server from the Internet (for example, allow TCP port 80 and 443 from 0.0.0.0/0, or from the load balancer) while the private facing interface has an associated security group allowing SSH access only from an allowed range of IP addresses either within the VPC or from the Internet, a private subnet within the VPC or a virtual private gateway.

To ensure failover capabilities, consider using a secondary private IP for incoming traffic on a network interface. In the event of an instance failure, you can move the interface and/or secondary private IP address to a standby instance.



## Use Network and Security Appliances in Your VPC

Some network and security appliances, such as load balancers, network address translation (NAT) servers, and proxy servers prefer to be configured with multiple network interfaces. You can create and attach secondary network interfaces to instances in a VPC that are running these types of applications and configure the additional interfaces with their own public and private IP addresses, security groups, and source/destination checking.

## Creating Dual-homed Instances with Workloads/Roles on Distinct Subnets

You can place a network interface on each of your web servers that connects to a mid-tier network where an application server resides. The application server can also be dual-homed to a back-end network (subnet) where the database server resides. Instead of routing network packets through the dual-homed instances, each dual-homed instance receives and processes requests on the front end, initiates a connection to the back end, and then sends requests to the servers on the back-end network.

## Create a Low Budget High Availability Solution

If one of your instances serving a particular function fails, its network interface can be attached to a replacement or hot standby instance pre-configured for the same role in order to rapidly recover the service. For example, you can use an ENI as your primary or secondary network interface to a critical service such as a database instance or a NAT instance. If the instance fails, you (or more likely, the code running on your behalf) can attach the ENI to a hot standby instance. Because the interface maintains its private IP addresses, Elastic IP addresses, and MAC address, network traffic will begin flowing to the standby instance as soon as you attach the ENI to the replacement instance. Users will experience a brief loss of connectivity between the time the instance fails and the time that the ENI is attached to the standby instance, but no changes to the VPC route table or your DNS server are required.

## Best Practices for Configuring Network Interfaces

- You can attach a network interface to an instance when it's running (hot attach), when it's stopped (warm attach), or when the instance is being launched (cold attach).
- You can detach secondary (eth*N*) network interfaces when the instance is running or stopped. However, you can't detach the primary (eth0) interface.
- You can attach a network interface in one subnet to an instance in another subnet in the same VPC, however, both the network interface and the instance must reside in the same Availability Zone.
- When launching an instance from the CLI or API, you can specify the network interfaces to attach to the instance for both the primary (eth0) and additional network interfaces.
- Launching an instance with multiple network interfaces automatically configures interfaces, private IP addresses, and route tables on the operating system of the instance. A warm or hot attach of an additional network interface may require you to manually bring up the second interface, configure the private IP address, and modify the route table accordingly. (Instances running Microsoft Windows Server or Amazon Linux automatically recognize the warm or hot attach and configure themselves.)
- Attaching another network interface to an instance is not a method to increase or double the network bandwidth to or from the dual-homed instance.

## Configuring Your Network Interface Using ec2-net-utils

Amazon Linux AMIs may contain additional scripts installed by Amazon Web Services, known as ec2-net-utils. These scripts optionally automate the configuration of your elastic network interfaces (ENIs).

Use the following command to install the package on Amazon Linux if it's not already installed, or update it if it's installed and additional updates are available:

```
yum install ec2-net-utils
```

The following components are part of ec2-net-utils:

### udev rules (/etc/udev/rules.d)

Identifies network interfaces when they are attached, detached, or reattached to a running instance, and ensures that the hotplug script runs (53-ec2-network-interfaces.rules). Maps the MAC address to a device name (75-persistent-net-generator.rules, which generates 70-persistent-net.rules).

### hotplug script

Generates an interface configuration file suitable for use with DHCP (/etc/sysconfig/network-scripts/ifcfg-eth*N*). Also generates a route configuration file (/etc/sysconfig/network-scripts/route-eth*N*).

### DHCP script

Whenever the network interface receives a new DHCP lease, this script queries the instance metadata for Elastic IP addresses. For each Elastic IP address, it adds a rule to the routing policy database to ensure that outbound traffic from that address uses the correct network interface. It also adds each private IP address to the network interface as a secondary address.

### ec2ifup eth*N*

Extends the functionality of the standard **ifup**. After this script rewrites the configuration files `ifcfg-ethN` and `route-ethN`, it runs **ifup**.

### ec2ifdown eth*N*

Extends the functionality of the standard **ifdown**. After this script removes any rules for the network interface from the routing policy database, it runs **ifdown**.

### ec2ifscan

Checks for network interfaces that have not been configured and configures them.

Note that this script isn't available in the initial release of ec2-net-utils.

To list any configuration files that were generated by ec2-net-utils, use the following command:

```
ls -l /etc/sysconfig/network-scripts/*-eth?
```

To disable the automation on a per-instance basis, you can add `EC2SYNC=no` to the corresponding `ifcfg-ethN` file. For example, use the following command to disable the automation for the `eth1` interface:

```
sed -i -e 's/^EC2SYNC=yes/EC2SYNC=no/' /etc/sysconfig/network-scripts/ifcfg-eth1
```

If you want to disable the automation completely, you can remove the package using the following command:

```
yum remove ec2-net-utils
```

## Creating a Network Interface

### To create a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Click **Create Network Interface**.
4. In the **Create Network Interface** dialog box, provide the following information for the network interface, and then click **Yes, Create**.
  - a. In **Description**, enter a descriptive name.
  - b. In **Subnet**, select the subnet. Note that you can't move the network interface to another subnet after it's created, and you can only attach the network interface to instances in the same subnet.
  - c. In **Private IP**, enter the primary private IP address. If you don't specify an IP address, we'll select an available private IP address from within the selected subnet.
  - d. In **Security Groups**, select one or more security groups.

The screenshot shows the 'Create Network Interface' dialog box. It has a title bar with 'Create Network Interface' and a 'Cancel' button. The dialog contains the following fields:

- Description:** A text input field.
- Subnet:** A dropdown menu with the text '--- Select a Subnet ---'.
- Private IP:** A text input field with 'auto-assign' selected.
- Security Groups:** A list box with the text 'Select a subnet first'.


At the bottom of the dialog are two buttons: 'Cancel' and 'Yes, Create'.

## Deleting a Network Interface

You must first detach a network interface from an instance before you can delete it. Deleting a network interface releases all attributes associated with the network interface and releases any private IP addresses or Elastic IP addresses to be used by another instance.

### To delete a network interface


1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Select a network interface, and then click the **Delete** button.
4. In the **Delete Network Interface** dialog box, click **Yes, Delete**.



## Viewing Details about a Network Interface

### To view details about a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Select the network interface.
4. View the details on the **Details** tab.

 **Network Interface:** eni-a66ed5cf

**Details** **Tags**

<b>Network Interface ID:</b>	 eni-a66ed5cf	<b>Subnet:</b>	subnet-cd8a35a4
<b>VPC:</b>	vpc-f28a359b	<b>Zone:</b>	ap-southeast-1b
<b>MAC Address:</b>	02:78:d7:00:8a:1e	<b>Description:</b>	Primary network interface
<b>Security Groups:</b>	quick-start-1	<b>Owner:</b>	053230519467
<b>Status:</b>	in-use	<b>Private IP:</b>	10.0.1.233
<b>Private DNS:</b>	-	<b>Secondary Private IPs:</b>	10.0.1.20
<b>Source/Dest. Check:</b>	enabled	<b>Attachment ID:</b>	eni-attach-a99c57c0
<b>Instance:</b>	i-886401dc	<b>Attachment Owner:</b>	053230519467
<b>Device Index:</b>	0	<b>Attachment Status:</b>	attached

## Attaching a Network Interface When Launching an Instance

You can attach an additional network interface to an instance when you launch it into a VPC.

### Note

If an error occurs when attaching a network interface to your instance, this causes the instance launch to fail.



### To attach a network interface when launching an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Launch Instance**.
3. Choose an AMI and click its **Select** button, then choose an instance type and click **Next: Configure Instance Details**.
4. On the **Configure Instance Details** page, choose a VPC from the **Network** list, and a subnet from the **Subnet** list.
5. In the **Network Interfaces** section, the console enables you specify up to 2 network interfaces when you launch an instance. After you launch the instance, click **Network Interfaces** in the navigation pane to add additional network interfaces. The total number of network interfaces that you can attach varies by instance type. For more information, see [Private IP Addresses Per ENI Per Instance Type](#) (p. 433). You can also enter an IP address for the primary network interface (eth0). When you've finished, click **Next: Add Storage**.
6. On the next **Add Storage** page, you can specify volumes to attach to the instance besides the volumes specified by the AMI (such as the root device volume), and then click **Next: Tag Instance**.
7. On the **Tag Instance** page, specify tags for the instance, such as a user-friendly name, and then click **Next: Configure Security Group**.
8. On the **Configure Security Group** page, select an existing security group or create a new one. Click **Review and Launch**.
9. On the **Review Instance Launch** page, details about the primary and additional network interface are displayed. Review the settings, and then click **Launch** to choose a key pair and launch your instance. If you're new to Amazon EC2 and haven't created any key pairs, the wizard prompts you to create one.

## Attaching a Network Interface to a Stopped or Running Instance


You can attach a network interface to any of your stopped or running instances in your VPC from either the **Instances** page or the **Network Interfaces** page of the Amazon EC2 console.

### To attach a network interface to a stopped or running instance using Instances

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Instances** in the navigation pane.
3. Right-click the instance, and then select **Attach Network Interface**.
4. In the **Attach Network Interface** dialog box, select the network interface, and then click **Attach**.

### To attach a network interface to a stopped or running instance using Network Interfaces

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Select the network interface.
4. Click the **Attach** button.
5. In the **Attach Network Interface** dialog box, select the instance, and then click **Yes, Attach**.



## Detaching a Network Interface from an Instance

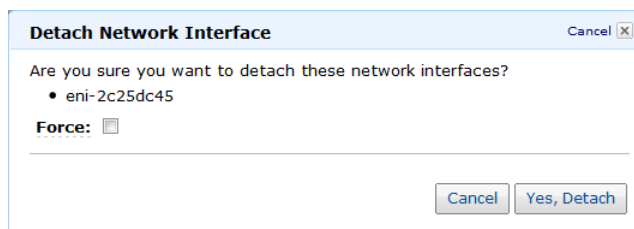
You can detach a secondary network interface at any time, using either the **Instances** or **Network Interfaces** pane of the Amazon EC2 console.

### To detach a network interface from an instance using Instances

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Instances** in the navigation pane.
3. Right-click the instance, and then select **Detach Network Interface**.
4. In the **Detach Network Interface** dialog box, select the network interface, and then click **Detach**.

### To detach a network interface from an instance using Network Interfaces

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Select the network interface, and then click the **Detach** button.
4. In the **Detach Network Interface** dialog box, click **Yes, Detach**. If the network interface fails to detach from the instance, select **Force**, and then try again.



## Changing the Security Group of a Network Interface


You can change the security groups that are associated with a network interface. When you create the security group, be sure to specify the same VPC as the subnet for the network interface.

### Note

You can't change security group membership for interfaces owned by other Amazon Web Services, such as Elastic Load Balancing, using the Amazon EC2 console, command line interface, or API actions. To modify a security group owned by another service, use the console, command line interface, or API for that service.

### To change the security group of a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Select the network interface.
4. Right-click the network interface, and then select **Change Security Groups**.
5. In the **Change Security Groups** dialog box, select the security groups to use, and then click **Save**.




## Changing the Source/Destination Checking of a Network Interface

The Source/Destination Check attribute controls whether source/destination checking is enabled on the instance. Disabling this attribute enables an instance to handle network traffic that isn't specifically destined for the instance. For example, instances running services such as network address translation, routing, or a firewall should set this value to `disabled`. The default value is `enabled`.

### To change source/destination checking for a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Right-click the network interface, and then select **Change Source/Dest Check**.
4. In the **Change Source/Dest Checking** dialog box, select **Enabled** (if enabling), or **Disabled** (if disabling), and then click **Save**.



## Associating an Elastic IP Address with a Network Interface

If you have an Elastic IP address, you can associate it with one of the private IP addresses for the network interface. You can associate one Elastic IP address with each private IP address.

### To associate an Elastic IP address with a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Right-click the network interface, and then select **Associate Address**.
4. In the **Associate Address** dialog box, select the Elastic IP address from the **Address** list.
5. In **Associate to private address**, select the private IP address to associate with the Elastic IP address.
6. Click **Allow Reassociation** to allow the Elastic IP address to be associated with the specified network interface if it's currently associated with another instance or network interface, and then click **Yes, Associate**.

The screenshot shows the 'Associate Address' dialog box. At the top, it says 'Associate Address' with a 'Cancel' button. Below that, it asks 'Select the address which you wish to associate to eni-69ce7500.' There are two dropdown menus: 'Address:' with 'Select an Address' and 'Associate to' with '10.0.1.152\*'. Below the second dropdown, it says 'private IP address: \* denotes the primary private IP address'. At the bottom, there is an 'Allow Reassociation:' checkbox which is unchecked. At the very bottom, there are two buttons: 'Cancel' and 'Yes, Associate'.

## Disassociating an Elastic IP Address from a Network Interface

If the network interface has an Elastic IP address associated with it, you can disassociate the address, and then either associate it with another network interface or release it back to the address pool. Note that this is the only way to associate an Elastic IP address with an instance in a different subnet or VPC using a network interface, as network interfaces are specific to a particular subnet.

### To disassociate an Elastic IP address from a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Right-click the network interface, and then select **Disassociate Address**.
4. In the **Disassociate Address** dialog box, click **Yes, Disassociate**.

The screenshot shows the 'Disassociate Address' dialog box. At the top, it says 'Disassociate Address' with a 'Cancel' button. Below that, it asks 'Are you sure that you wish to disassociate this IP address?'. There are two lines of text: 'Network Interface: eni-69ce7500' and 'Public IP: 44.01.222.154 - 10.0.1.152\*'. At the bottom, there are two buttons: 'Cancel' and 'Yes, Disassociate'.

## Changing Termination Behavior for a Network Interface

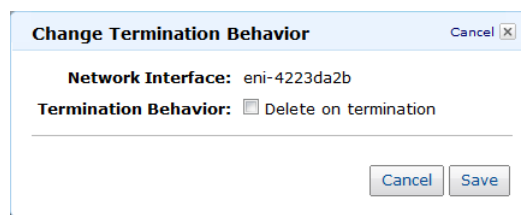
You can set the termination behavior for a network interface attached to an instance so that it is automatically deleted when you delete the instance it's attached to.

### Note

By default, network interfaces that are automatically created and attached to instances using the EC2 console are set to terminate when the instance terminates. However, network interfaces created using the `ec2-create-network-interface` command aren't set to terminate when the instance terminates.

### To change termination behavior for network interfaces

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Right-click the network interface, and then select **Change Termination Behavior**.
4. In the **Change Termination Behavior** dialog box, select the **Delete on termination** check box if you want the network interface to be deleted when you terminate an instance.



**Change Termination Behavior** Cancel X

**Network Interface:** eni-4223da2b

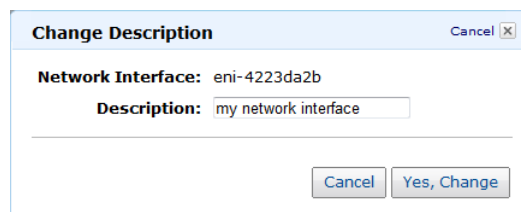
**Termination Behavior:**  Delete on termination

Cancel Save

## Adding or Editing a Description for a Network Interface

### To add or edit a description for a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Right-click the network interface, and then select **Change Description**.
4. In the **Change Description** dialog box, enter a description for the network interface, and then click **Yes, Change**.



**Change Description** Cancel X

**Network Interface:** eni-4223da2b

**Description:**

Cancel Yes, Change

## Adding or Editing Tags for a Network Interface

Tags are metadata that you can add to a network interface. Tags are private and are only visible to your account. Each tag consists of a key and an optional value. For more information about tags, see [Tagging Your Amazon EC2 Resources \(p. 532\)](#).

### To add or edit tags for a network interface

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Network Interfaces** in the navigation pane.
3. Select the network interface.
4. In the details pane, click the **Tags** tab, and then click **Add/Edit Tags**.
5. In the **Tag Network Interfaces** dialog box, enter a key and an optional value for each tag that you want to add, and then click **Save Tags**.

## API and Command Overview

The following table summarizes the available network interface commands and corresponding API actions.

Description	Command	API Action
Attaches a network interface to an instance.	<code>ec2-attach-network-interface</code>	<code>AttachNetworkInterface</code>
Creates a network interface in the specified subnet.	<code>ec2-create-network-interface</code>	<code>CreateNetworkInterface</code>
Deletes a network interface.	<code>ec2-delete-network-interface</code>	<code>DeleteNetworkInterface</code>
Describes a network interface attribute.	<code>ec2-describe-network-interface-attribute</code>	<code>DescribeNetworkInterfaceAttribute</code>
Describes one or more of your network interfaces.	<code>ec2-describe-network-interfaces</code>	<code>DescribeNetworkInterfaces</code>
Detaches a network interface from an instance.	<code>ec2-detach-network-interface</code>	<code>DetachNetworkInterface</code>
Modifies a network interface attribute.	<code>ec2-modify-network-interface-attribute</code>	<code>ModifyNetworkInterfaceAttribute</code>
Resets a network interface attribute.	<code>ec2-reset-network-interface-attribute</code>	<code>ResetNetworkInterfaceAttribute</code>

# Storage

---

## Topics


- [Amazon Elastic Block Store \(Amazon EBS\)](#) (p. 446)
- [Amazon EC2 Instance Store](#) (p. 508)
- [Amazon Simple Storage Service \(Amazon S3\)](#) (p. 517)
- [Block Device Mapping](#) (p. 517)

Amazon EC2 provides you with flexible, cost effective, and easy-to-use data storage options for your instances. Each option has a unique combination of performance and durability. These storage options can be used independently or in combination to suit your requirements.

After reading this section, you should have a good understanding about how you can use the data storage options supported by Amazon Elastic Compute Cloud to meet your specific requirements. These storage options include the following:

- Amazon Elastic Block Store (Amazon EBS)
- Amazon EC2 instance store
- Amazon Simple Storage Service (Amazon S3)

The following figure shows the relationship between these types of storage.



### Amazon EBS

Amazon EBS provides durable, block-level storage volumes that you can attach to a running Amazon EC2 instance. You can use Amazon EBS as a primary storage device for data that requires frequent and granular updates. For example, Amazon EBS is the recommended storage option when you run a database on an instance.

An Amazon EBS volume behaves like a raw, unformatted, external block device that you can attach to a single instance. They persist independently from the running life of an Amazon EC2 instance. After an EBS volume is attached to an instance, you can use it like any other physical hard drive. As illustrated in the previous figure, you can attach multiple volumes to an instance. You can also detach an EBS volume from one instance and attach it to another instance.

To keep a back-up copy of your data, you can create a *snapshot* of an EBS volume, which is stored in Amazon S3. You can create a new Amazon EBS volume from a snapshot, and attach it to another instance.

### Amazon EC2 Instance Store

Each Amazon EC2 instance, unless it's a micro or M3 instance, can access storage from disks that are physically attached to the host computer. This disk storage is referred to as *instance store*. Instance store provides temporary block-level storage for Amazon EC2 instances. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance; if you stop or terminate an instance, any data on instance store volumes is lost.

### Amazon S3

Amazon S3 is a repository for Internet data. Amazon S3 provides access to reliable and inexpensive data storage infrastructure. It is designed to make web-scale computing easier by enabling you to store and retrieve any amount of data, at any time, from within Amazon EC2 or anywhere on the web. For example, you can use Amazon S3 to store backup copies of your data and applications.

### Adding Storage

Every time you launch an instance from an AMI, a root storage device is created for that instance. The root storage device contains all the information necessary to boot the instance. You can specify storage volumes in addition to the root device volume when you create an AMI or launch an instance using *block device mapping*.

You can also attach EBS volumes to a running instance.



# Amazon Elastic Block Store (Amazon EBS)

Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes are highly available and reliable storage volumes that can be attached to any running instance that is in the same Availability Zone. Amazon EBS volumes that are attached to an Amazon EC2 instance are exposed as storage volumes that persist independently from the life of the instance. With Amazon EBS, you only pay for what you use. For more information about Amazon EBS pricing, see the Projecting Costs section of the [Amazon Elastic Block Store page](#).

Amazon EBS is recommended when data changes frequently and requires long-term persistence. Amazon EBS volumes are particularly well-suited for use as the primary storage for file systems, databases, or for any applications that require fine granular updates and access to raw, unformatted, block-level storage. Amazon EBS is particularly helpful for database-style applications that frequently encounter many random reads and writes across the data set.

You can attach multiple volumes to the same instance within the limits specified by your AWS account. Your account has a limit on the number of Amazon EBS volumes that you can use, and the total storage available to you. For more information about these limits, and how to request an increase in your limits, see [Request to Increase the Amazon EBS Volume Limit](#).

## Topics

- [Features of Amazon EBS \(p. 446\)](#)
- [Amazon EBS Volumes \(p. 447\)](#)
- [Amazon EBS Snapshots \(p. 484\)](#)
- [Amazon EBS Volume Performance \(p. 494\)](#)
- [Amazon EBS API and Command Overview \(p. 506\)](#)

## Features of Amazon EBS

- You can create Amazon EBS storage volumes from 1 GB to 1 TB in size and mount them as devices on your Amazon EC2 instances. You can mount multiple volumes on the same instance, but each volume can only be attached to one instance at a time. For more information, see [Creating or Restoring an Amazon EBS Volume \(p. 449\)](#).
- With Amazon EBS provisioned IOPS (input/output operations per second) volumes, you can provision a specific level of I/O performance, up to 4000 IOPS per volume. This allows you to predictably scale to thousands of IOPS per Amazon EC2 instance. For more information, see [Provisioned IOPS Volumes \(p. 449\)](#).
- Amazon EBS volumes behave like raw, unformatted block devices. You can create a file system on top of these volumes, or use them in any other way you would use a block device (like a hard drive). For more information on creating file systems and mounting volumes, see [Making an Amazon EBS Volume Available for Use \(p. 459\)](#).
- You can create point-in-time snapshots of Amazon EBS volumes, which are persisted to Amazon S3. Snapshots protect data for long-term durability and they can be used as the starting point for new Amazon EBS volumes. The same snapshot can be used to instantiate as many volumes as you wish. These snapshots can be copied across AWS regions. For more information, see [Amazon EBS Snapshots \(p. 484\)](#).
- Amazon EBS volumes are created in a specific Availability Zone, and can then be attached to any instances in that same Availability Zone. To make a volume available outside of the Availability Zone, you can create a snapshot and restore that snapshot to a new volume anywhere in that region. You can copy snapshots to other regions and then restore them to new volumes there, making it easier to leverage multiple AWS regions for geographical expansion, data center migration, and disaster recovery. For more information, see [Creating an Amazon EBS Snapshot \(p. 485\)](#) and [Copying an Amazon EBS Snapshot \(p. 488\)](#).

- A large repository of public data set snapshots can be restored to Amazon EBS volumes and seamlessly integrated into AWS cloud-based applications. For more information, see [Using Public Data Sets \(p. 453\)](#).
- Performance metrics, such as bandwidth, throughput, latency, and average queue length, are available through the AWS Management Console. These metrics, provided by Amazon CloudWatch, allow you to monitor the performance of your volumes to make sure that you are providing enough performance for your applications without paying for resources you don't need. For more information, see [Amazon EBS Volume Performance \(p. 494\)](#).

## Amazon EBS Volumes

An Amazon EBS volume is a durable, block-level storage device that you can attach to a single EC2 instance. You can use Amazon EBS volumes as primary storage for data that requires frequent updates, such as the system drive for an instance or storage for a database application. Amazon EBS volumes persist independently from the running life of an EC2 instance. After a volume is attached to an instance, you can use it like any other physical hard drive. Amazon EBS provides two volume types: Standard and Provisioned IOPS. They differ in performance characteristics and price, allowing you to tailor your storage performance and cost to the needs of your applications. For more information, see [Amazon EBS Volume Types \(p. 448\)](#).

### Topics

- [Benefits of Using Amazon EBS Volumes \(p. 447\)](#)
- [Amazon EBS Volume Types \(p. 448\)](#)
- [Creating or Restoring an Amazon EBS Volume \(p. 449\)](#)
- [Using Public Data Sets \(p. 453\)](#)
- [Attaching an Amazon EBS Volume to an Instance \(p. 455\)](#)
- [Making an Amazon EBS Volume Available for Use \(p. 459\)](#)
- [Describing Volumes \(p. 462\)](#)
- [Monitoring the Status of Your Volumes \(p. 464\)](#)
- [Detaching an Amazon EBS Volume from an Instance \(p. 475\)](#)
- [Deleting an Amazon EBS Volume \(p. 477\)](#)
- [Expanding the Storage Space of a Volume \(p. 479\)](#)

## Benefits of Using Amazon EBS Volumes

### Data Availability

When you create an Amazon EBS volume in an Availability Zone, it is automatically replicated within that zone to prevent data loss due to failure of any single hardware component. After you create a volume, you can attach it to any Amazon EC2 instance in the same Availability Zone. After you attach a volume, it appears as a native block device similar to a hard drive or other physical device. At that point, the instance can interact with the volume just as it would with a local drive; the instance can format the Amazon EBS volume with a file system such as ext3 (Linux) or NTFS (Windows) and install applications.

An Amazon EBS volume can be attached to only one instance at a time within the same Availability Zone. However, multiple volumes can be attached to a single instance. If you attach multiple volumes to a device that you have named, you can stripe data across the volumes for increased I/O and throughput performance.

You can get monitoring data for your Amazon EBS volumes at no additional charge (this includes data for the root device volumes for Amazon EBS-backed instances). For more information, see [Monitoring Volumes with CloudWatch \(p. 464\)](#).

### Data Persistence

An Amazon EBS volume is off-instance storage that can persist independently from the life of an instance. You continue to pay for the volume usage as long as the data persists.

By default, Amazon EBS volumes that are attached to a running instance automatically detach from the instance with their data intact when that instance is terminated. The volume can then be reattached to a new instance, enabling quick recovery. If you are using an Amazon EBS-backed instance, you can stop and restart that instance without affecting the data stored in the attached volume. The volume remains attached throughout the stop-start cycle. This enables you to process and store the data set indefinitely, only using the processing and storage resources when required. The data set persists on the volume until the volume is deleted explicitly. After a volume is deleted, it can't be attached to any instance.

By default, Amazon EBS volumes that are created and attached to an instance at launch are deleted when that instance is terminated. You can modify this behavior by changing the value of the flag `DeleteOnTermination` to `false` when you launch the instance. This modified value causes the volume to persist even after the instance is terminated, and enables you to attach the volume to another instance.

### Snapshots

Amazon EBS provides the ability to create snapshots (backups) of any Amazon EC2 volume and write a copy of the data in the volume to Amazon S3, where it is stored redundantly in multiple Availability Zones. The volume does not need to be attached to a running instance in order to take a snapshot. As you continue to write data to a volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes. These snapshots can be used to create multiple new Amazon EBS volumes, expand the size of a volume, or move volumes across Availability Zones.

When you create a new volume from a snapshot, it's an exact copy of the original volume at the time the snapshot was taken. By optionally specifying a different volume size or a different Availability Zone, you can use this functionality to increase the size of an existing volume or to create duplicate volumes in new Availability Zones. The snapshots can be shared with specific AWS accounts or made public. When you create snapshots, you incur charges in Amazon S3 based on the volume's total size. For a successive snapshot of the volume, you are only charged for any additional data beyond the volume's original size.

Amazon EBS snapshots are incremental backups, meaning that only the blocks on the volume that have changed after your most recent snapshot are saved. If you have a volume with 100 GiB of data, but only 5 GiB of data have changed since your last snapshot, only the 5 GiB of modified data is written to Amazon S3. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.

To help categorize and manage your volumes and snapshots, you can tag them with metadata of your choice. For more information, see [Tagging Your Amazon EC2 Resources \(p. 532\)](#).

## Amazon EBS Volume Types

Amazon EBS provides two volume types: Standard and Provisioned IOPS (input/output operations per second). They differ in performance characteristics and price, allowing you to tailor your storage performance and cost to the needs of your applications.

### Standard Volumes

Standard volumes offer cost effective storage that is ideal for applications with light or bursty I/O requirements. These volumes deliver approximately 100 IOPS on average, with burst capability of up to hundreds of IOPS, and they can range in size from 1 GB to 1 TB. Standard volumes can be striped together in a RAID configuration for larger size and greater performance. The following example usage scenarios work well with Amazon EBS standard volumes:

- File server
- Log processing
- Low-traffic websites

- Analytics
- System boot volume

### Provisioned IOPS Volumes

Provisioned IOPS volumes are designed to meet the needs of I/O-intensive workloads, particularly database workloads, that are sensitive to storage performance and consistency in random access I/O throughput. You specify an IOPS rate when you create the volume, and Amazon EBS delivers within 10 percent of the provisioned IOPS performance 99.9 percent of the time over a given year.

A Provisioned IOPS volume can range in size from 10 GB to 1 TB and you can provision up to 4000 IOPS per volume. The ratio of IOPS provisioned to the volume size requested can be a maximum of 30; for example, a volume with 3000 IOPS must be at least 100 GB. You can stripe multiple volumes together in a RAID configuration for larger size and greater performance. The following example usage scenarios work well with Amazon EBS Provisioned IOPS volumes:

- Business applications
- Database workloads, such as:
  - MongoDB
  - Microsoft SQL Server
  - MySQL
  - PostgreSQL
  - Oracle

There are several factors that can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, and workload demand. For more information about getting the most out of your Provisioned IOPS volumes, see [Amazon EBS Volume Performance \(p. 494\)](#).

## Creating or Restoring an Amazon EBS Volume

### Topics

- [AWS Management Console \(p. 450\)](#)
- [Command Line Interface \(p. 450\)](#)
- [API \(p. 451\)](#)

To use Amazon EBS, you first create a volume that can be attached to any Amazon EC2 instance within the same Availability Zone. You can create an Amazon EBS volume with data from a snapshot stored in Amazon S3 or as an empty volume. You can also create and attach Amazon EBS volumes when you launch instances. For more information about volumes and snapshots, see [Amazon Elastic Block Store \(Amazon EBS\) \(p. 446\)](#).

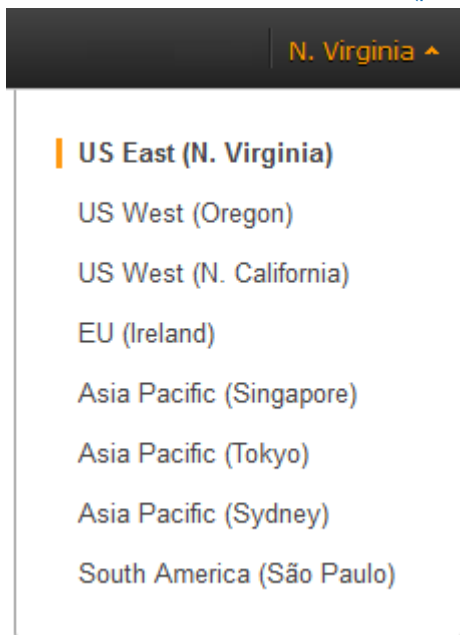
New volumes created from existing Amazon S3 snapshots load lazily in the background. This means that after a volume is created from a snapshot, there is no need to wait for all of the data to transfer from Amazon S3 to your Amazon EBS volume before your attached instance can start accessing the volume and all its data. If your instance accesses data that hasn't yet been loaded, the volume immediately downloads the requested data from Amazon S3, and continues loading the rest of the data in the background.

Accessing data for the first time from Amazon S3 might cause latency during the loading period. To avoid the possibility of an increased latency during the background loading of the data, you first access/read all performance-critical data from the volume (or read the entire volume) to ensure it has been loaded to the Amazon EBS volume from Amazon S3 before running the application.

## AWS Management Console

### To create or restore an Amazon EBS volume

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, select the region that meets your needs. This choice is important because some Amazon EC2 resources can be shared between regions, while others can't. For more information, see [Resource Locations](#) (p. 528).



3. Click **Volumes** in the navigation pane.
4. Above the upper pane, click **Create Volume**.
5. In the **Create Volume** dialog box, in the **Volume Type** list, select **Standard** or **Provisioned IOPS**. For more information, see [Amazon EBS Volume Types](#) (p. 448).
6. In the **Size** box and **GiB** list, select the size of the volume (in GiB or TiB).
7. For Provisioned IOPS volumes, in the **IOPS** box, enter the maximum number of input/output operations per second (IOPS) that the volume can support.
8. In the **Availability Zone** list, select the Availability Zone in which to launch the instance. For more information, see [Regions and Availability Zones](#) (p. 7).
9. To restore an Amazon EBS volume, in the **Snapshot** list, select the ID of the snapshot from which you are launching the volume. Skip this step if you aren't restoring an Amazon EBS volume.
10. Click **Yes, Create**.

#### Note

If you specify both a volume size and a snapshot ID, the size must be equal to or greater than the snapshot size. Any AWS Marketplace product codes from the snapshot are propagated to the volume.

## Command Line Interface

### To create or restore an Amazon EBS volume

1. To create a volume, use the `ec2-create-volume` command:

```
PROMPT> ec2-create-volume --size 80 --availability-zone us-east-1a
```

To restore a volume, use the [ec2-create-volume](#) command and specify the snapshot using the `--snapshot` option:

```
PROMPT> ec2-create-volume --snapshot snap-1f727354 --size 80 --availability-zone us-east-1a
```

Amazon EC2 returns information about the volume that is similar to the following example.

```
VOLUME vol-c7f95aae 80 us-east-1a creating 2010-03-30T13:54:37+0000 standard
```

2. To check whether the volume is ready, use the [ec2-describe-volumes](#) command.

```
PROMPT> ec2-describe-volumes volume_id
```

Amazon EC2 returns information about the volume that is similar to the following example.

```
VOLUME vol-c7f95aae 80 us-east-1a available 2010-03-30T13:54:37+0000 standard
```

### Note

Any AWS Marketplace product codes from the snapshot are propagated to the volume.

## API

To create an Amazon EBS volume, use the [CreateVolume](#) action. Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=CreateVolume  
&Size=size  
&AvailabilityZone=zone  
&AUTHPARAMS
```

The following is an example response.

```
<CreateVolumeResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <requestId>06eabcdf-95b8-424f-87dc-da8e976f7858</requestId>  
  <volumeId>vol-d11fbbb8</volumeId>  
  <size>80</size>  
  <snapshotId/>  
  <availabilityZone>us-east-1a</availabilityZone>  
  <status>creating</status>  
  <createTime>2010-03-23T09:16:24.000Z</createTime>  
  <volumeType>standard</volumeType>  
</CreateVolumeResponse>
```

To restore an Amazon EBS volume, use the [CreateVolume](#) action. Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=CreateVolume  
&Size=size  
&SnapshotId=snapshot  
&AvailabilityZone=zone  
&AUTHPARAMS
```

The following is an example response.

```
<CreateVolumeResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <requestId>06eabcdf-95b8-424f-87dc-da8e976f7858</requestId>  
  <volumeId>vol-d11fbbb8</volumeId>  
  <size>80</size>  
  <snapshotId>snap-1f727354</snapshotId>  
  <availabilityZone>us-east-1a</availabilityZone>  
  <status>creating</status>  
  <createTime>2010-03-23T09:16:24.000Z</createTime>  
  <volumeType>standard</volumeType>  
</CreateVolumeResponse>
```

## Using Public Data Sets

### Topics

- [Public Data Set Concepts](#) (p. 453)
- [Finding Public Data Sets](#) (p. 453)
- [Launching an Instance](#) (p. 453)
- [Creating a Public Data Set Volume](#) (p. 454)
- [Mounting the Public Data Set Volume](#) (p. 455)

This section describes how to use Amazon EC2 public data sets.

### Public Data Set Concepts

Amazon EC2 provides a repository of public data sets that can be seamlessly integrated into AWS cloud-based applications. Amazon stores the data sets at no charge to the community and, like with all AWS services, you pay only for the compute and storage you use for your own applications.

Previously, large data sets such as the mapping of the Human Genome and the US Census data required hours or days to locate, download, customize, and analyze. Now, anyone can access these data sets from an Amazon EC2 instance and start computing on the data within minutes. You can also leverage the entire AWS ecosystem and easily collaborate with other AWS users. For example, you can produce or use prebuilt server images with tools and applications to analyze the data sets. By hosting this important and useful data with cost-efficient services such as Amazon EC2, AWS hopes to provide researchers across a variety of disciplines and industries with tools to enable more innovation, more quickly.

For more information, go to the [Public Data Sets Page](#).

### Available Public Data Sets

Public data sets are currently available in the following categories:

- **Biology**—Includes Human Genome Project, GenBank, and other content.
- **Chemistry**—Includes multiple versions of PubChem and other content.
- **Economics**—Includes census data, labor statistics, transportation statistics, and other content.
- **Encyclopedic**—Includes Wikipedia content from multiple sources and other content.

### Finding Public Data Sets

Before you launch a public data set, you must locate it.

#### To find a public data set

1. Go to the [Public Data Sets Page](#).
2. Locate a public data set and write down its snapshot ID for your operating platform (Windows or Linux/UNIX).

### Launching an Instance

You'll attach a volume based on the public data set to an instance. Launch the instance as you typically launch an instance. For more information, see [Launch Your Instance](#) (p. 266).



## Creating a Public Data Set Volume

To use a public data set, you create an Amazon EBS volume, specifying the snapshot ID of the public data set.

### AWS Management Console

#### To create an Amazon EBS volume

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, click **Volumes**.
3. Above the upper pane, click **Create Volume**.
4. In the **Create Volume** dialog box, in the **Volume Type** drop-down list, select **Standard** or **Provisioned IOPS**. For more information, see [Amazon EBS Volume Types \(p. 448\)](#).
5. In the **Size** box and **GiB** drop-down list, select the size of the volume (in GiB or TiB).
6. For Provisioned IOPS volumes, in the **IOPS** box, enter the maximum number of input/output operations per second (IOPS) that the volume can support.
7. In the **Availability Zone** drop-down list, select the availability zone in which to launch the instance. For more information, see [Regions and Availability Zones \(p. 7\)](#)
8. In the **Snapshot** drop-down list, select the ID of the snapshot from which you are launching the volume (optional)
9. Click **Yes, Create**.

### Command Line Tools

#### To create an Amazon EBS volume

1. Enter the following command.

```
PROMPT> ec2-create-volume --snapshot public-data-set-snapshot-id --zone availability-zone
```

Amazon EBS returns information about the volume similar to the following example.

```
VOLUME vol-4d826724 85 us-east-1a creating 2008-02-14T00:00:00+0000
```

2. To check whether the volume is ready, use the following command.

```
PROMPT> ec2-describe-volumes vol-4d826724
```

Amazon EBS returns information about the volume similar to the following example.

```
VOLUME vol-4d826724 85 us-east-1a available 2008-07-29T08:49:25+0000
```

## API

#### To create an Amazon EBS volume

- Construct the following Query request.

```
https://ec2.amazonaws.com/  
?Action=CreateVolume  
&AvailabilityZone=zone  
&SnapshotId=public-data-set-snapshot-id  
&AUTHPARAMS
```

The following is an example response.

```
<CreateVolumeResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <volumeId>vol-4d826724</volumeId>  
  <size>85</size>  
  <status>creating</status>  
  <createTime>2008-05-07T11:51:50.000Z</createTime>  
  <availabilityZone>us-east-1a</availabilityZone>  
  <snapshotId>snap-59d33330</snapshotId>  
</CreateVolumeResponse>
```

## Mounting the Public Data Set Volume

Mount the public data set volume as you typically mount an EBS volume. For more information, see [Making an Amazon EBS Volume Available for Use \(p. 459\)](#).

## Attaching an Amazon EBS Volume to an Instance

### Topics

- [AWS Management Console \(p. 456\)](#)
- [Command Line Interface \(p. 457\)](#)
- [API \(p. 458\)](#)

This section describes how to attach an Amazon EBS volume that you created to an instance.

The following table lists the available device names on Amazon EC2. You can specify these names when attaching a volume to a running instance or when attaching a volume when launching an instance using a block device mapping. The block device driver for the instance assigns the actual volume names when mounting the volumes, and these names can be different than the names that Amazon EC2 recommends. For more information about the instance store, see [Amazon EC2 Instance Store \(p. 508\)](#). For information about the root device storage, see [Amazon EC2 Root Device Volume \(p. 12\)](#).

Instance Type	Possible for Connection	Reserved for Root	Instance Store Volumes	Recommended for EBS Connection
Linux / UNIX	/dev/sd[a-z] /dev/sd[a-z][1-15] /dev/hd[a-z] /dev/hd[a-z][1-15]	/dev/sda1	/dev/sd[b-e]	/dev/sd[f-p] /dev/sd[f-p][1-6]

Instance Type	Possible for Connection	Reserved for Root	Instance Store Volumes	Recommended for EBS Connection
Windows	xvd[a-z] xvdb[a-z] /dev/sda[1-2] /dev/sd[b-e]	/dev/sda1	xvd[a-e]	xvd[f-p] (with Red Hat PV drivers)  xvd[f-z] (with Citrix PV drivers)

### Important

For Linux/UNIX instance types, we've received reports that some custom kernels might have restrictions that limit use to `/dev/sd[f-p]` or `/dev/sd[f-p][1-6]`. If you're having trouble using `/dev/sd[q-z]` or `/dev/sd[q-z][1-6]`, try switching to `/dev/sd[f-p]` or `/dev/sd[f-p][1-6]`.

### Important

You cannot attach volumes that share the same device letters both with and without trailing digits. For example, if you attach a volume as `/dev/sdc` and another volume as `/dev/sdc1`, only `/dev/sdc` will be visible to the instance. If you wish use trailing digits in device names, you must use trailing digits on all device names that share the same base letters.

Depending on the block device driver of your instance's kernel, the device may be attached with a different name than what you specify. For example, if you specify a device name of `/dev/sdh`, your device may be renamed `/dev/xvdh` or `/dev/hdh` by the kernel; in some cases, even the trailing letter may also change (where `/dev/sda` could become `/dev/xvde`). Amazon Linux AMIs create a symbolic link from the renamed device path to the name you specify, but other AMIs may behave differently.

Hardware virtual machine (HVM) AMIs (such as the base Windows and Cluster Compute images) do not support the use of trailing numbers on device names (`xvd[a-p][1-15]`). You can view the virtualization type of your instance on the **Instances** page of the Amazon EC2 dashboard.

Depending on the size of your instance, Amazon EC2 provides instance store volumes on `sd[b-e]` (on Linux/UNIX) or `xvd[a-e]`. Although you can connect your Amazon EBS volumes using these device names, we highly recommend that you don't because the behavior can be unpredictable.

An Amazon EC2 Windows AMI comes with an additional service installed, the **Ec2Config Service**. The Ec2Config service runs as a local system and performs various functions to prepare an instance when it first boots up. After the devices have been mapped with the drives, the Ec2Config service then initializes and mounts the drives. The root drive is initialized and mounted as `c:\`. The instance stores that comes attached to the instance are initialized and mounted as `d:\`, `e:\`, and so on. By default, when an Amazon EBS volume is attached to a Windows instance, it may show up as any drive letter on the instance. You can change the settings of the Ec2Config service to set the drive letters of the Amazon EBS volumes per your specifications. For more information, see [Using Ec2Config](#) in the *Amazon Elastic Compute Cloud Microsoft Windows Guide*.

## AWS Management Console

### To attach an Amazon EBS volume to an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Volumes** in the navigation pane.

The console displays a list of current volumes.

3. Select a volume and click **Attach Volume**.

The **Attach Volume** dialog box appears.

4. Select the instance to attach the volume to from the **Instance** list (only instances in the same Availability Zone as the volume are displayed).
5. Select how the device is exposed to the instance from the **Device** list.
6. Click **Attach** to attach the volume to the instance. The volume and instance must be in the same Availability Zone.

#### Note

- Windows instances use either Red Hat or Citrix paravirtual (PV) drivers. If your Windows instance is using Citrix PV drivers, you can attach up to a total of 25 Amazon EBS volumes using the Amazon EC2 CLI; Windows instances with Red Hat PV drivers are limited to 16 volumes. Regardless of which drivers you are using, you can attach up to 16 Amazon EBS volumes using the Amazon EC2 console. To upgrade your Windows instance from Red Hat to Citrix paravirtual drivers, see [Upgrading Your PV Drivers on Your Windows AMI](#).
- Although it is technically possible to attach more than 25 volumes to a Windows instance with Citrix PV drivers, this is likely to cause performance issues and is not recommended.
- If you want to push an Amazon S3 object to an Amazon EBS volume, follow the procedure above to attach the volume, then use a data transfer application (such as FTP/SFTP or SCP) on your running instance to transfer the data.

If a volume has an AWS Marketplace product code:

- The volume can only be attached to the root device of a stopped instance.
- You must be subscribed to the AWS Marketplace code that is on the volume.
- The configuration (instance type, operating system) of the instance must support that specific AWS Marketplace code. For example, you cannot take a volume from a Windows instance and attach it to a Linux instance.
- AWS Marketplace product codes are copied from the volume to the instance.

## Command Line Interface

To attach an Amazon EBS volume to an instance, use [ec2-attach-volume](#).

```
PROMPT> ec2-attach-volume volume_id -i instance_id -d device
```

Amazon EC2 returns information using this syntax.

```
ATTACHMENT volume_id instance_id device attaching date_time
```

This example attaches volume `vol-4d826724` to instance `i-6058a509` in Linux and UNIX and exposes it as device `/dev/sdh`.

```
PROMPT> ec2-attach-volume vol-4d826724 -i i-6058a509 -d /dev/sdh
```

```
ATTACHMENT vol-4d826724 i-6058a509 /dev/sdh attaching 2010-03-30T13:58:58+0000
```

This example attaches volume `vol-4d826724` to instance `i-6058a509` in Windows and exposes it as device `xvdf`.

```
PROMPT> ec2-attach-volume vol-4d826724 -i i-6058a509 -d xvdf
```

```
ATTACHMENT vol-4d826724 i-6058a509 xvdf attaching 2010-03-30T13:58:58+0000
```

**Note**

The volume and instance must be in the same Availability Zone.

**Note**

- Windows instances use either Red Hat or Citrix paravirtual (PV) drivers. If your Windows instance is using Citrix PV drivers, you can attach up to a total of 25 Amazon EBS volumes using the Amazon EC2 CLI; Windows instances with Red Hat PV drivers are limited to 16 volumes. Regardless of which drivers you are using, you can attach up to 16 Amazon EBS volumes using the Amazon EC2 console. To upgrade your Windows instance from Red Hat to Citrix paravirtual drivers, see [Upgrading Your PV Drivers on Your Windows AMI](#).
- Although it is technically possible to attach more than 25 volumes to a Windows instance with Citrix PV drivers, this is likely to cause performance issues and is not recommended.

**Note**

If a volume has an AWS Marketplace product code:

- The volume can only be attached to the root device of a stopped instance.
- You must be subscribed to the AWS Marketplace code that is on the volume.
- The configuration (instance type, operating system) of the instance must support that specific AWS Marketplace code. For example, you cannot take a volume from a Windows instance and attach it to a Linux instance.
- AWS Marketplace product codes are copied from the volume to the instance.

For an overview of the AWS Marketplace, see <https://aws.amazon.com/marketplace/help/200900000>. For details on how to use the AWS Marketplace, see [AWS Marketplace](#).

## API

To attach an Amazon EBS volume to an instance, use [AttachVolume](#). Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=AttachVolume  
&VolumeId=volume-id  
&InstanceId=instance-id  
&Device=device  
&AUTHPARAMS
```

The following is an example response.

```
<AttachVolumeResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <requestId>59dbff89-35bd-4eac-99ed-be587EXAMPLE</requestId>  
  <volumeId>vol-4d826724</volumeId>  
  <instanceId>i-6058a509</instanceId>  
  <device>/dev/sdh</device>  
  <status>attaching</status>  
  <attachTime>2008-05-07T11:51:50.000Z</attachTime>  
</AttachVolumeResponse>
```

**Note**

The volume and instance must be in the same Availability Zone.

**Note**

- Windows instances use either Red Hat or Citrix paravirtual (PV) drivers. If your Windows instance is using Citrix PV drivers, you can attach up to a total of 25 Amazon EBS volumes using the Amazon EC2 CLI; Windows instances with Red Hat PV drivers are limited to 16 volumes. Regardless of which drivers you are using, you can attach up to 16 Amazon EBS volumes using the Amazon EC2 console. To upgrade your Windows instance from Red Hat to Citrix paravirtual drivers, see [Upgrading Your PV Drivers on Your Windows AMI](#).
- Although it is technically possible to attach more than 25 volumes to a Windows instance with Citrix PV drivers, this is likely to cause performance issues and is not recommended.

**Note**

If a volume has an AWS Marketplace product code:

- The volume can only be attached to the root device of a stopped instance.
- You must be subscribed to the AWS Marketplace code that is on the volume.
- The configuration (instance type, operating system) of the instance must support that specific AWS Marketplace code. For example, you cannot take a volume from a Windows instance and attach it to a Linux instance.
- AWS Marketplace product codes are copied from the volume to the instance.

For an overview of the AWS Marketplace, see <https://aws.amazon.com/marketplace/help/200900000>. For details on how to use the AWS Marketplace, see [AWS Marketplace](#).

## Making an Amazon EBS Volume Available for Use

After you attach an Amazon EBS volume to your instance, it is exposed as a block device. You can format the volume with any file system and then mount it. After you make the Amazon EBS volume available for use, you can access it in the same ways that you access any other volume. You can also take snapshots of your Amazon EBS volume for backup purposes or to use as a baseline when you create another volume.

This topic describes how to make your Amazon EBS volume available for use.

### Making the Volume Available on Linux

#### To make an Amazon EBS volume available for use on Linux

1. Connect to your instance using SSH.
2. Depending on the block device driver of your instance's kernel, the device may be attached with a different name than what you specify. For example, if you specify a device name of `/dev/sdh`, your device may be renamed `/dev/xvdh` or `/dev/hdh` by the kernel; in some cases, even the trailing letter may also change (where `/dev/sda` could become `/dev/xvde`). Amazon Linux AMIs create a symbolic link from the renamed device path to the name you specify, but other AMIs may behave differently.

Use the `lsblk` command to view your available disk devices and their mount points (if applicable) to help you determine the correct device name to use.

```
$ lsblk
NAME MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
```

```
xvda1 202:1 0 8G 0 disk /  
xvdh 202:112 0 100G 0 disk
```

In this example, `/dev/xvda1` is mounted as the root device, and `/dev/xvdh` is attached, but it has not been mounted yet. The `df -h` and `sudo fdisk -l` commands also provide valuable information about disk space and partitions.

- (Optional) Use the following command to create an ext3 file system on the volume.

**Caution**

This step assumes that you're mounting an empty volume. If you're mounting a volume that already has data on it (for example, a volume that was restored from a snapshot), don't use `mkfs` before mounting the volume (skip to the next step instead). Otherwise, you'll format the volume and delete the existing data.

```
$ sudo mkfs -t ext3 device_name
```

- Use the following command to create the directory.

```
$ sudo mkdir volume_name
```

- Use the following command to mount the volume.

```
$ sudo mount device_name volume_name
```

- (Optional) To enable the instance to reconnect to an Amazon EBS volume on reboot, add the device to the `fstab` or create a script that automatically mounts the volume on startup.

## Make the Volume Available on Windows


### To use an Amazon EBS volume

- Log in to your instance using Remote Desktop.
- Windows Server 2012: Go to the Start screen.
  - Windows Server 2008: On the taskbar, click **Start**, and then click **Run**.
- Type `diskmgmt.msc` and press **Enter**. The **Disk Management** utility opens.


**Caution**

If you're mounting a volume that already has data on it (for example, a public data set), make sure you don't reformat the volume and delete the existing data.

- Select the disk that represents the new Amazon EBS volume.



5. On the **Disk Management** menu, select **Action - All Tasks - Online**.



6. A new disk needs to be initialized before it can be used. To initialize the disk:
  - a. In the Disk Management utility, select the new Amazon EBS volume disk.
  - b. On the **Disk Management** menu, select **Action - All Tasks - Initialize Disk**.
  - c. In the **Initialize Disk** dialog, select the disk to initialize, select the desired partition style, and press **OK**.

Your new Amazon EBS volume is now available for use. Any data written to this file system is written to the Amazon EBS volume and is transparent to applications using the device.



## Describing Volumes

### Topics

- [AWS Management Console](#) (p. 462)
- [Command Line Interface](#) (p. 462)
- [API](#) (p. 463)

You can list information about a volume, including the specific instance to which the volume is attached.

### AWS Management Console

#### To view information about an Amazon EBS volume

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Volumes** in the navigation pane.

The console displays a list of current volumes and the instances they are attached to.

3. To view more information about a volume, select it.

Information about the volume appears in the lower pane.

### Command Line Interface

To describe volumes and list information about the volumes that you own, use the [ec2-describe-volumes](#) command.

```
PROMPT> ec2-describe-volumes
```

Amazon EC2 returns information similar to the following.

```
VOLUME vol-4d826724 80          us-east-1a in-use      2010-03-30T13:58:58+0000
  standard
ATTACHMENT vol-4d826724 i-6058a509 /dev/sdh attached 2010-03-30T13:54:55+0000
VOLUME vol-50957039 13          us-east-1a available 2010-03-24T08:01:44+0000
  standard
VOLUME vol-6682670f 1          us-east-1a in-use      2010-03-30T08:11:01+0000
  standard
ATTACHMENT vol-6682670f i-69a54000 /dev/sdh attached 2010-03-30T09:21:14+0000
```

This information includes the volume ID, capacity, status (in-use or available), and creation time of each volume. If the volume is attached, an attachment line shows the volume ID, the instance ID to which the volume is attached, the device name exposed to the instance, its status (attaching, attached, detaching, detached), and when it attached.

#### Tip

You can use **ec2-describe-volumes** to filter the results to only the volumes that match the criteria you specify.

To describe instances and list volumes that are attached to running instances, use the [ec2-describe-instances](#) command.

```
PROMPT> ec2-describe-instances
```

Amazon EC2 returns information similar to the following.

```
RESERVATION    r-f25e6f9a      111122223333    default
INSTANCE       i-84b435de      ami-b232d0db    ec2-184-73-201-68.compute-
1.amazonaws.comdomU-12-31-39-00-86-35.compute-1.internal    running gsg-keypair
0    ml.small 2010-03-30T08:43:48+0000    us-east-1a    aki-94c527fd
ari-96c527ff  monitoring-disabled  184.73.201.68    10.254.137.191    ebs
BLOCKDEVICE    /dev/sda1        vol-cf13b3a6    2010-03-30T08:01:44.000Z
true
BLOCKDEVICE    /dev/sdh         vol-c7f95aae    2010-03-30T13:58:58.000Z
true
```

### Tip

You can use `ec2-describe-instances` to filter the results to only the instances that match the criteria you specify.

For more information about block device mapping, see [Block Device Mapping \(p. 517\)](#).

## API

To describe volumes and list information about all volumes that you own, use the [DescribeVolumes](#) action. Construct the following request.

```
https://ec2.amazonaws.com/
?Action=DescribeVolumes
&AUTHPARAMS
```

The following is an example response.

```
<DescribeVolumesResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <volumeSet>
    <item>
      <volumeId>vol-4282672b</volumeId>
      <size>80</size>
      <status>in-use</status>
      <createTime>2008-05-07T11:51:50.000Z</createTime>
      <attachmentSet>
        <item>
          <volumeId>vol-4282672b</volumeId>
          <instanceId>i-6058a509</instanceId>
          <size>80</size>
          <snapshotId>snap-12345678</snapshotId>
          <availabilityZone>us-east-1a</availabilityZone>
          <status>attached</status>
          <attachTime>2008-05-07T12:51:50.000Z</attachTime>
        </item>
      </attachmentSet>
    </item>
    ...
  </volumeSet>
```

### Tip

You can use `DescribeVolumes` to filter the results to only the volumes that match the criteria you specify.

## Monitoring the Status of Your Volumes

Amazon Web Services (AWS) automatically provides data, such as Amazon CloudWatch metrics and volume status checks, that you can use to monitor your Amazon Elastic Block Store (Amazon EBS) volumes.

### Topics

- [Monitoring Volumes with CloudWatch](#) (p. 464)
- [Monitoring Volumes with Status Checks](#) (p. 466)
- [Monitoring Volume Events](#) (p. 469)
- [Working with an Impaired Volume](#) (p. 471)
- [Working with the AutoEnableIO Volume Attribute](#) (p. 473)

## Monitoring Volumes with CloudWatch

CloudWatch metrics are statistical data that you can use to view, analyze, and set alarms on the operational behavior of your volumes.

The following table describes the types of monitoring data available for your Amazon EBS volumes.

Type	Description
Basic	Data is available automatically in 5-minute periods at no charge. This includes data for the root device volumes for Amazon EBS-backed instances.
Detailed	Provisioned IOPS volumes automatically send one-minute metrics to CloudWatch.

When you get data from CloudWatch, you can include a `Period` request parameter to specify the granularity of the returned data. This is different than the period that we use when we collect the data (5-minute periods). We recommend that you specify a period in your request that is equal to or larger than the collection period to ensure that the returned data is valid.

You can get the data using either the Amazon CloudWatch API or the AWS Management Console. The console takes the raw data from the Amazon CloudWatch API and displays a series of graphs based on the data. Depending on your needs, you might prefer to use either the data from the API or the graphs in the console.

### Amazon EBS Metrics

You can use the Amazon CloudWatch `GetMetricStatistics` API to get any of the Amazon EBS volume metrics listed in the following table. Similar metrics are grouped together in the table, and the metrics in the first two rows are also available for the local stores on Amazon EC2 instances.

Metric	Description
<code>VolumeReadBytes</code>	The total number of bytes transferred in a specified period of time. Data is only reported to Amazon CloudWatch when the volume is active. If the volume is idle, no data is reported to Amazon CloudWatch.
<code>VolumeWriteBytes</code>	
	Units: Bytes

Metric	Description
VolumeReadOps VolumeWriteOps	<p>The total number of I/O operations in a specified period of time.</p> <p><b>Note</b> To calculate the average I/O operations per second (IOPS) for the period, divide the total operations in the period by the number of seconds in that period.</p> <p>Units: Count</p>
VolumeTotalReadTime VolumeTotalWriteTime	<p>The total number of seconds spent by all operations that completed in a specified period of time. If multiple requests are submitted at the same time, this total could be greater than the length of the period. For example, for a period of 5 minutes (300 seconds): if 700 operations completed during that period, and each operation took 1 second, the value would be 700 seconds.</p> <p>Units: Seconds</p>
VolumeIdleTime	<p>The total number of seconds in a specified period of time when no read or write operations were submitted.</p> <p>Units: Seconds</p>
VolumeQueueLength	<p>The number of read and write operation requests waiting to be completed in a specified period of time.</p> <p>Units: Count</p>
VolumeThroughputPercentage	<p>Used with Provisioned IOPS volumes only. The percentage of I/O operations per second (IOPS) delivered of the total IOPS provisioned for an EBS volume. Provisioned IOPS volumes deliver within 10 percent of the provisioned IOPS performance 99.9 percent of the time over a given year.</p> <p><b>Note</b> During a write, if there are no other pending I/O requests in a minute, the metric value will be 100 percent. Also, a volume's I/O performance may become degraded temporarily due to an action you have taken (e.g., creating a snapshot of a volume during peak usage, running the volume on a non-EBS-optimized instance, accessing data on the volume for the first time).</p> <p>Units: Percent</p>
VolumeConsumedReadWriteOps	<p>Used with Provisioned IOPS volumes only. The total amount of read and write operations consumed in a specified period of time.</p> <p>Units: Count</p>

### Graphs in the AWS Management Console

After you create a volume, you can go to the Amazon EC2 console and view the volume's monitoring graphs. They're displayed when you select the volume on the **Volumes** page in the EC2 console. A **Monitoring** tab is displayed next to the volume's **Description** tab. The following table lists the graphs that are displayed. The column on the right describes how the raw data metrics from the Amazon CloudWatch API are used to produce each graph. The period for all the graphs is 5 minutes.

<b>Graph Name</b>	<b>Description Using Raw Metrics</b>
Read Bandwidth (KiB/s)	Sum(VolumeReadBytes) / Period / 1024
Write Bandwidth (KiB/s)	Sum(VolumeWriteBytes) / Period / 1024
Read Throughput (Ops/s)	Sum(VolumeReadOps) / Period
Write Throughput (Ops/s)	Sum(VolumeWriteOps) / Period
Avg Queue Length (ops)	Avg(VolumeQueueLength)
% Time Spent Idle	Sum(VolumeIdleTime) / Period * 100
Avg Read Size (KiB/op)	Avg(VolumeReadBytes) / 1024
Avg Write Size (KiB/op)	Avg(VolumeWriteBytes) / 1024
Avg Read Latency (ms/op)	Avg(VolumeTotalReadTime) * 1000
Avg Write Latency (ms/op)	Avg(VolumeTotalWriteTime) * 1000

For the average latency graphs and average size graphs, the average is calculated over the total number of operations (read or write, whichever is applicable to the graph) that completed during the period.

The AWS Management Console contains a console for Amazon CloudWatch. In the Amazon CloudWatch console you can search and browse all your AWS resource metrics, view graphs to troubleshoot issues and discover trends, create and edit alarms to be notified of problems, and see at-a-glance overviews of your alarms and AWS resources. For more information, see [AWS Management Console](#) in the *Amazon CloudWatch Developer Guide*.

## Monitoring Volumes with Status Checks

Volume status checks enable you to better understand, track, and manage potential inconsistencies in the data on an Amazon EBS volume. They are designed to provide you with the information that you need to determine whether your EBS volumes are impaired, and to help you control how a potentially inconsistent volume is handled.

Volume status checks are automated tests that return a pass or fail status. If all checks pass, the status of the volume is `ok`. If a check fails, the status of the volume is `impaired`. If the status is `insufficient-data`, the checks may still be in progress on the volume. You can view the results of volume status checks to identify any impaired volumes and take any necessary actions.

When EBS determines that a volume's data is potentially inconsistent, the default is that it disables I/O to the volume from any attached EC2 instances, which helps to prevent data corruption. After I/O is disabled, the next volume status check fails, and the volume status is `impaired`. In addition, you'll see an event that lets you know that I/O is disabled, and that you can resolve the impaired status of the volume by enabling I/O to the volume. We wait until you enable I/O to give you the opportunity to decide whether to continue to let your instances use the volume, or to run a consistency check using a command such as `fsck` (Linux/UNIX) or `chkdsk` (Windows) before doing so.

### Note

Volume status is based on the volume status checks, and does not reflect the volume state. Therefore, volume status does not indicate volumes in the `error` state (for example, when a volume is incapable of accepting I/O.)

If the consistency of a particular volume is not a concern for you, and you'd prefer that the volume be made available immediately if it's impaired, you can override the default behavior by configuring the

volume to automatically enable I/O. If you enable the `AutoEnableIO` volume attribute, the volume status check continues to pass. In addition, you'll see an event that lets you know that the volume was determined to be potentially inconsistent, but that its I/O was automatically enabled. This enables you to check the volume's consistency or replace it at a later time.

The following table lists statuses for standard and provisioned IOPS volumes.


Overall Volume Status	I/O Enabled Status	I/O Performance Status (Provisioned IOPS volumes only)
<code>ok</code>	Enabled (I/O Enabled or I/O Auto-Enabled)	Normal (Volume performance is as expected)
<code>warning</code>	Enabled (I/O Enabled or I/O Auto-Enabled)	Degraded (Volume performance is below expectations)  Severely Degraded (Volume performance is well below expectations)
<code>impaired</code>	Enabled (I/O Enabled or I/O Auto-Enabled)  Disabled (Volume is offline and pending recovery, or is waiting for the user to enable I/O)	Stalled (Volume performance is severely impacted)  Not Available (Unable to determine I/O performance because I/O is disabled)
<code>insufficient-data</code>	Enabled (I/O Enabled or I/O Auto-Enabled)  Insufficient Data	Insufficient Data

To view and work with status checks, you can use the AWS Management Console, the API, or the command line interface.

### AWS Management Console

#### To view status checks

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, under **ELASTIC BLOCK STORE**, click **Volumes**.
3. On the **EBS Volumes** page, the **Status Checks** column lists the operational status of each volume.
4. To view an individual volume's status, select the volume, and then click the **Status Checks** tab.



- If you have a volume with a failed status check (status is impaired), see [Working with an Impaired Volume \(p. 471\)](#).

Alternatively, you can use the **Events** pane to view all events for your instances and volumes in a single pane. For more information, see [Monitoring Volume Events \(p. 469\)](#).

### Command Line Tools

The following examples show how to use the `ec2-describe-volume-status` command to describe the status of your volumes.

To do this	Run this command
Get the status of all volumes	<code>ec2-describe-volume-status</code>
Get the status of all volumes with a volume status of impaired	<code>ec2-describe-volume-status --filter "volume-status.status=impaired"</code>

The following is an example of the output for two volumes, including one that is impaired.

```
Type          VolumeId      AvailabilityZone  VolumeStatus
VOLUME        vol-0452396b us-east-1a       ok
VOLUME        vol-2670f74a us-east-1b       impaired
Type          Name          Status
VOLUMESTATUS io-enabled    failed
Type  EventType                NotBefore                NotAfter  EventId
      EventDescription
EVENT potential-data-inconsistency 2011-12-01T14:00:00.000Z          evol-61a54008 This is an example
Type  ActionCode                EventId                EventType
      EventDescription
ACTION enable-volume-io          evol-61a54008          potential-data-inconsistency
      This is an example
```

If you have a volume with a failed status check (status is `impaired`), see [Working with an Impaired Volume](#) (p. 471).

For more information, see [ec2-describe-volume-status](#) in the *Amazon Elastic Compute Cloud Command Reference Guide*.

## API

You can use the `DescribeVolumeStatus` action to retrieve the status of your volumes. For usage information, see [DescribeVolumeStatus](#) in the *Amazon Elastic Compute Cloud API Reference*.

## Monitoring Volume Events

When EBS determines that a volume's data is potentially inconsistent, it disables I/O to the volume from any attached EC2 instances by default. This causes the volume status check to fail, and creates a volume status event that indicates the cause of the failure.

To automatically enable I/O on a volume with potential data inconsistencies, change the setting of the `AutoEnableIO` volume attribute. For more information about changing this attribute, see [Working with an Impaired Volume](#) (p. 471).

Each event includes a start time that indicates the time at which the event occurred, and a duration that indicates how long I/O for the volume was disabled. The end time is added to the event when I/O for the volume is enabled.

Volume status events include one of the following descriptions:

### Awaiting Action: Enable IO

Volume data is potentially inconsistent. I/O is disabled for the volume until you explicitly enable it. The event description changes to **IO Enabled** after you explicitly enable I/O.

### IO Enabled

I/O operations were explicitly enabled for this volume.

### IO Auto-Enabled

I/O operations were automatically enabled on this volume after an event occurred. We recommend that you check for data inconsistencies before continuing to use the data.

### Normal

For provisioned IOPS volumes only. Volume performance is as expected.

### Degraded

For provisioned IOPS volumes only. Volume performance is below expectations.

### Severely Degraded

For provisioned IOPS volumes only. Volume performance is well below expectations.

### Stalled

For provisioned IOPS volumes only. Volume performance is severely impacted.

You can view events for your volumes using the AWS Management Console, the API, or the command line interface.

## AWS Management Console

### To view events for your volumes

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, click **Events**. You see a list of all instances and volumes and any associated events.
3. On the **Events** page, in **Viewing**, select **Volumes** to view only volume status. You can also filter on specific status types.



4. To view the event details of a specific volume, select the volume in the **Events** page.

The screenshot shows the Amazon EC2 console's Events page. On the left is a navigation pane with categories like INSTANCES, IMAGES & SECURITY, ELASTIC BLOCK STORE, and NETWORK & SECURITY. The main area displays a table of events for all resources and statuses. The table has columns for Resource Id, Event Status, Resour, Zone, Event Type, Description, Start Time, Duration, and Ev. One event is selected, and its details are shown below the table.

N	Resource Id	Event Status	Resour	Zone	Event Type	Description	Start Time	Duration	Ev
<input type="checkbox"/>	vol-f138498c	IO Enabled	Volume	us-east-1c	potential-data-inconsistency	IO Enabled	2012-11-07 13:49 PST	9 minutes, 18 seconds	IO
<input type="checkbox"/>	vol-28215055	IO Enabled	Volume	us-east-1c	potential-data-inconsistency	IO Enabled	2012-11-01 15:24 PDT	6 minutes, 16 seconds	IO
<input type="checkbox"/>	vol-387e9b44	Awaiting Action: Enable IO	Volume	us-east-1c	potential-data-inconsistency	Awaiting Action: Enable IO	2012-11-24 23:19 PST	2 days, 11 hours	IO
<input type="checkbox"/>	vol-087e9b74	IO Enabled	Volume	us-east-1c	potential-data-inconsistency	IO Enabled	2012-11-14 22:35 PST	2 minutes, 7 seconds	IO
<input checked="" type="checkbox"/>	vol-087e9b74	I/O throughput (0 to 50%) is well below expected	Volume	us-east-1c	io-performance:severely-degraded	I/O throughput (0 to 50%) is well below expected	2012-11-27 11:10 PST	5 minutes, 42 seconds	IO
<input type="checkbox"/>	vol-c9987cb5	IO Auto Enabled	Volume	us-east-1c	potential-data-inconsistency	IO Auto Enabled	2012-11-07 13:11 PST	11 minutes, 53 seconds	IO
<input type="checkbox"/>	vol-c9987cb5	I/O throughput (> 90%) is as expected	Volume	us-east-1c	io-performance:degraded	I/O throughput (> 90%) is as expected	2012-11-14 17:08 PST	1 minute	IO

**1 Event selected**

**Event:** vol-087e9b74

**Event Status:** I/O throughput (0 to 50%) is well below expectations

**IO Status:** Enabled

**Zone:** us-east-1c

**Event Type:** io-performance:severely-degraded

**Attached To:** Loading...

**Start Time:** 2012-11-27 11:10 PST

**End Time:**

For more information about Events, see Working with Volume Status and Events in the Amazon EC2 User Guide. If you need technical assistance with your volume, post your issue to the Developer Forums or visit our Support Center.

If you have a volume where I/O is disabled, see [Working with an Impaired Volume \(p. 471\)](#). If you have a volume where I/O performance is below normal, this might be a temporary condition due to an action you have taken (e.g., creating a snapshot of a volume during peak usage, running the volume on a non-EBS-optimized instance, accessing data on the volume for the first time, etc.).

### Command Line Tools

To view events for your volumes, use the `ec2-describe-volume-status` command as follows.

```
ec2-describe-volume-status
```

Amazon EC2 returns output similar to the following:

```
ec2-describe-volume-status vol-11111111, vol-22222222
Type      VolumeId      AvailabilityZone VolumeStatus
VOLUME    vol-11111111 us-east-1a      ok
VOLUME    vol-22222222 us-east-1b      impaired
Type      Name          Status
VOLUMESTATUS io-enabled    failed
Type      EventType     NotBefore          NotAfter EventId      EventDescription
EVENT     potential-data-inconsistency 2011-12-01T14:00:00.000Z      848721011
This is an example
Type      ActionCode     EventId      EventType
EVENT     enable-volume-io 848721011    potential-data-inconsistency
This is an example
```

For more information, see `ec2-describe-volume-status` in the *Amazon Elastic Compute Cloud Command Line Reference*.

If you have a volume with a failed status check, see [Working with an Impaired Volume \(p. 471\)](#).

## API

You can use the `DescribeVolumeStatus` action to retrieve the status of your volumes and any associated events. For usage information, see [DescribeVolumeStatus](#) in the *Amazon Elastic Compute Cloud API Reference*.

## Working with an Impaired Volume

This section discusses your options if a volume is impaired because the volume's data is potentially inconsistent.

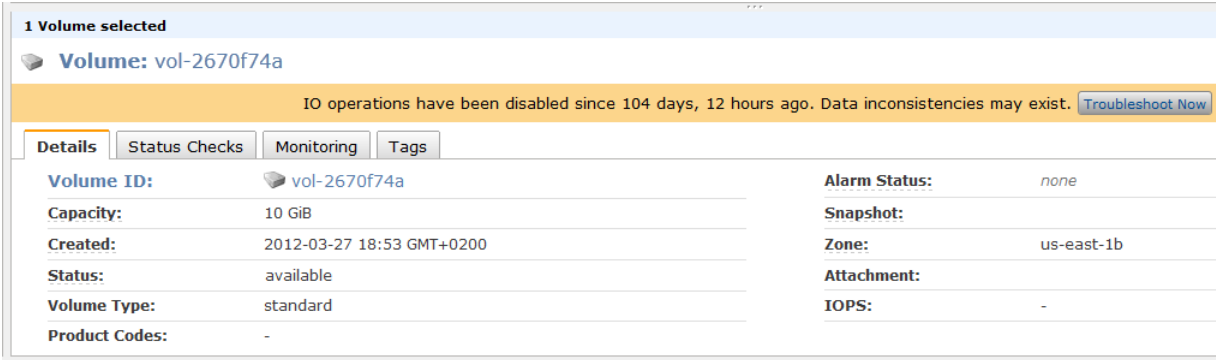
### Topics

- [Option 1: Perform a Consistency Check on the Volume Attached to its Instance \(p. 471\)](#)
- [Option 2: Perform a Consistency Check on the Volume Using Another Instance \(p. 472\)](#)
- [Option 3: Delete the Volume If You No Longer Need It \(p. 473\)](#)

### Option 1: Perform a Consistency Check on the Volume Attached to its Instance

The simplest option is to enable I/O and then perform a data consistency check on the volume while the volume is still attached to its Amazon EC2 instance.

#### To perform a consistency check on an attached volume

1. Stop any applications from using the volume.
2. Enable I/O on the volume.
  - a. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
  - b. In the navigation pane, under **ELASTIC BLOCK STORE**, click **Volumes**.
  - c. Select the volume on which you want to enable I/O operations.
  - d. On the **Details** tab, click **Enable Volume IO**.

1 Volume selected	
Volume: vol-2670f74a	
I/O operations have been disabled since 104 days, 12 hours ago. Data inconsistencies may exist. <a href="#">Troubleshoot Now</a>	
Details   Status Checks   Monitoring   Tags	
Volume ID:	vol-2670f74a
Capacity:	10 GiB
Created:	2012-03-27 18:53 GMT+0200
Status:	available
Volume Type:	standard
Product Codes:	-
Alarm Status:	none
Snapshot:	
Zone:	us-east-1b
Attachment:	
IOPS:	-
  - e. In **Enable Volume IO**, click **Yes, Enable**.
3. Check the data on the volume.
  - a. (Optional) Run the **fsck** (Linux) or **chkdsk** (Windows) command.
  - b. Review the application logs.
  - c. (Optional) If the volume has been impaired for more than 20 minutes you can contact support. Click **Troubleshoot Now**, and then on the **Troubleshoot Status Checks** dialog box, click **Contact Support** to submit a support case.

For information about using the command line interface to enable I/O for a volume, see [ec2-enable-volume-io](#) in the *Amazon Elastic Compute Cloud Command Line Reference*. For information about using the API to enable I/O for a volume, see [EnableVolumeIO](#) in the *Amazon Elastic Compute Cloud API Reference*.

## Option 2: Perform a Consistency Check on the Volume Using Another Instance

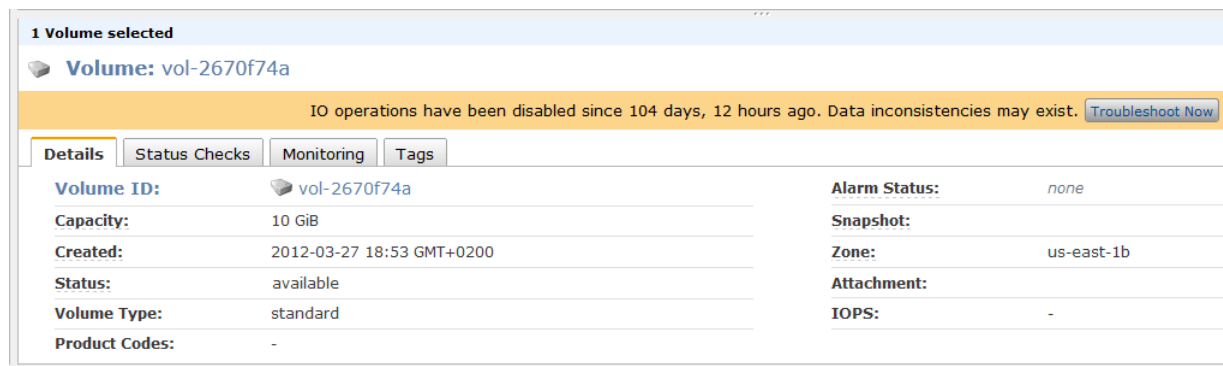
Use the following procedure to check the volume outside your production environment.

### Important

This procedure may cause the loss of write I/Os that were suspended when volume I/O was disabled.

### To perform a consistency check on a volume in isolation

1. Stop any applications from using the volume.
2. Detach the volume from the instance.
  - a. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
  - b. In the navigation pane, under **ELASTIC BLOCK STORE**, click **Volumes**.
  - c. Select the volume that you want to detach.
  - d. Click **More**, and then click **Force Detach**.
  - e. In the **Force Detach Volume** dialog box, click **Yes, Force**.
3. Enable I/O on the volume.
  - a. In the navigation pane, under **ELASTIC BLOCK STORE**, click **Volumes**.
  - b. Select the volume that you detached in the previous step.
  - c. On the **Details** tab, click **Enable Volume IO**.



- d. In the **Enable Volume IO** dialog box, click **Yes, Enable**.
4. Attach the volume to another instance. For information, see [Launch Your Instance \(p. 266\)](#) and [Attaching an Amazon EBS Volume to an Instance \(p. 455\)](#).
  5. Check the data on the volume.
    - a. (Optional) Run the **fsck** (Linux) or **chkdsk** (Windows) command.
    - b. Review the application logs.
    - c. (Optional) If the volume has been impaired for more than 20 minutes, you can contact support. Click **Troubleshoot Now**, and then on the **Troubleshoot Status Checks** dialog box, click **Contact Support** to submit a support case.

For information about using the command line interface to enable I/O for a volume, see [ec2-enable-volume-io](#) in the *Amazon Elastic Compute Cloud Command Line Reference*. For information about using the API to enable I/O for a volume, see [EnableVolumeIO](#) in the *Amazon Elastic Compute Cloud API Reference*.

### Option 3: Delete the Volume If You No Longer Need It

If you want to remove the volume from your environment, simply delete it. For information about deleting a volume, see [Deleting an Amazon EBS Volume](#) (p. 477).

If you have a recent snapshot that backs up the data on the volume, you can create a new volume from the snapshot. For information about creating a volume from a snapshot, see [Creating or Restoring an Amazon EBS Volume](#) (p. 449).

## Working with the AutoEnableIO Volume Attribute


When EBS determines that a volume's data is potentially inconsistent, it disables I/O to the volume from any attached EC2 instances by default. This causes the volume status check to fail, and creates a volume status event that indicates the cause of the failure. If the consistency of a particular volume is not a concern, and you prefer that the volume be made available immediately if it's `impaired`, you can override the default behavior by configuring the volume to automatically enable I/O. If you enable the `AutoEnableIO` volume attribute, I/O between the volume and the instance is automatically reenabled and the volume's status check will pass. In addition, you'll see an event that lets you know that the volume was in a potentially inconsistent state, but that its I/O was automatically enabled. When this event occurs, you should check the volume's consistency and replace it if necessary. For more information, see [Monitoring Volume Events](#) (p. 469).

This section explains how to view and modify the `AutoEnableIO` attribute of a volume using the AWS Management Console, the command line interface, or the API.

### AWS Management Console


#### To view the AutoEnableIO attribute of a volume

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, under **ELASTIC BLOCK STORE**, click **Volumes**.
3. Select the volume.
4. In the lower pane, click the **Status Checks** tab.
5. In the **Status Checks** tab, **Auto-Enable Volume IO** displays the current setting for your volume, either `Enabled` or `Disabled`.




### To modify the AutoEnableIO attribute of a volume

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, under **ELASTIC BLOCK STORE**, click **Volumes**.
3. Select the volume.
4. At the top of the **Volumes** page, click **More**.
5. Click **Change Auto-Enable IO Setting**.



6. In the **Modify Auto-Enable Volume IO** dialog box, set the **Auto-Enable Volume IO** option to automatically enable I/O for an impaired volume. To disable the feature, clear the option.



7. Click **Save**.

Alternatively, instead of completing steps 4-6 in the previous procedure, go to the **Status Checks** tab and click **Auto-Enable IO Setting**.

## Command Line Tools

Use the `ec2-describe-volume-attribute` command as follows to view the `AutoEnableIO` attribute of a volume.

```
ec2-describe-volume-attribute vol-12345678 --auto-enable-io
```

For more information, see [ec2-describe-volume-attribute](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

Use the `ec2-modify-volume-attribute` command as follows to modify the `AutoEnableIO` attribute of a volume:

```
ec2-modify-volume-attribute vol-12345678 --auto-enable-io true
```

For more information, see [ec2-modify-volume-attribute](#) in the *Amazon Elastic Compute Cloud Command Line Reference*.

## API

To view the `AutoEnableIO` attribute of a volume, use the `DescribeVolumeAttribute` action. For usage information, see [DescribeVolumeAttribute](#) in the *Amazon Elastic Compute Cloud API Reference*.

To modify the `AutoEnableIO` attribute of a volume, use the `ModifyVolumeAttribute` action. For usage information, see [ModifyVolumeAttribute](#) in the *Amazon Elastic Compute Cloud API Reference*.

## Detaching an Amazon EBS Volume from an Instance

### Topics

- [AWS Management Console \(p. 475\)](#)
- [Command Line Interface \(p. 476\)](#)
- [API \(p. 477\)](#)

You can detach an Amazon EBS volume from an instance explicitly or by terminating the instance. However, a volume must be unmounted inside the instance before being detached. Failure to do so results in the volume being stuck in the busy state while it is trying to detach, which could possibly damage the file system or the data it contains. You can reattach a volume that you detached (without unmounting it), but it may not get the same mount point and the data on the volume might be out of sync if there were writes to the volume in progress when it was detached.

### Note

If an Amazon EBS volume is the root device of an instance, it cannot be detached unless the instance is in the stopped state.

If the root volume is detached from an instance with an AWS Marketplace product code, then the AWS Marketplace product codes from that volume are no longer associated with the instance.

This example unmounts the volume and then explicitly detaches it from the instance. This is useful when you want to terminate an instance or attach a volume to a different instance. To verify that the volume is no longer attached to the instance, see [Describing Volumes \(p. 462\)](#).

## AWS Management Console

### To detach an Amazon EBS volume

1. First, unmount the volume.

For Linux/UNIX, use the following command to unmount the `/dev/sdh` device.

```
# umount -d /dev/sdh
```

For Windows, open **Disk Management**, right-click the volume to unmount, and select **Change Drive Letter and Path**. Then, select the mount point to remove and click **Remove**.

2. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
3. Click **Volumes** in the navigation pane.

The console displays a list of current volumes.

4. Select a volume and click **Detach Volume**.

A confirmation dialog box appears.

5. Click **Yes, Detach**.

The volume is detached from the instance.

### Caution

If your volume stays in the *detaching* state, you can force the detachment by clicking **Force Detach**. Forcing the detachment can lead to data loss or a corrupted file system. Use this option only as a last resort to detach a volume from a failed instance, or if you are detaching a volume with the intention of deleting it. The instance doesn't get an opportunity to flush file system caches or file system metadata. If you use this option, you must perform file system check and repair procedures.

If you've tried to force the volume to detach multiple times over several minutes and it stays in the *detaching* state, you can post a request for help to the [Amazon EC2 forum](#). To help expedite a resolution, include the volume ID and describe the steps that you've already taken.

## Command Line Interface

### To detach an Amazon EBS volume explicitly

1. For Linux/UNIX, use the following command to unmount the `/dev/sdh` device.

```
# umount -d /dev/sdh
```

For Windows, open **Disk Management**, right-click the volume to unmount, and select **Change Drive Letter and Path**. Then, select the mount point to remove and click **Remove**.

2. To detach the volume, use the `ec2-detach-volume` command.

```
PROMPT> ec2-detach-volume vol-4d826724
```

Amazon EC2 returns information similar to the following example.

```
ATTACHMENT vol-4d826724 i-6058a509 /dev/sdh detaching 2010-03-30T13:58:58+0000
```

### Caution

If your volume stays in the *detaching* state, you can force the detachment by adding the `--force` option to your command. Forcing the detachment can lead to data loss or a corrupted file system. Use this option only as a last resort to detach a volume from a failed instance, or if you are detaching a volume with the intention of deleting it. The instance

doesn't get an opportunity to flush file system caches or file system metadata. If you use this option, you must perform file system check and repair procedures. If you've tried to force the volume to detach multiple times over several minutes and it stays in the *detaching* state, you can post a request for help to the [Amazon EC2 forum](#). To help expedite a resolution, include the volume ID and describe the steps that you've already taken.

### To detach an Amazon EBS volume by terminating the instance

- Use the `ec2-terminate-instances` command.

```
PROMPT> ec2-terminate-instances i-6058a509
```

Amazon EC2 returns information similar to the following example.

```
INSTANCE    i-6058a509    running shutting-down
```

## API

To detach an Amazon EBS volume, use the [DetachVolume](#) action. Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=DetachVolume  
&VolumeId=volume-id  
&InstanceId=instance-id  
&AUTHPARAMS
```

The following is an example response.

```
<DetachVolumeResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <requestId>59dbff89-35bd-4eac-99ed-be587EXAMPLE</requestId>  
  <volumeId>vol-4d826724</volumeId>  
  <instanceId>i-6058a509</instanceId>  
  <device>/dev/sdh</device>  
  <status>detaching</status>  
  <attachTime>2008-05-08T11:51:50.000Z</attachTime>  
</DetachVolumeResponse>
```

### Caution

If your volume stays in the *detaching* state, you can force the detachment by adding the `Force` parameter to your API request. Forcing the detachment can lead to data loss or a corrupted file system. Use this option only as a last resort to detach a volume from a failed instance, or if you are detaching a volume with the intention of deleting it. The instance doesn't get an opportunity to flush file system caches or file system metadata. If you use this option, you must perform file system check and repair procedures.

If you've tried to force the volume to detach multiple times over several minutes and it stays in the *detaching* state, you can post a request for help to the [Amazon EC2 forum](#). To help expedite a resolution, include the volume ID and describe the steps that you've already taken.

## Deleting an Amazon EBS Volume

### Topics



- [AWS Management Console](#) (p. 478)
- [Command Line Interface](#) (p. 478)
- [API](#) (p. 478)

After you no longer need a volume, you can delete it. After deletion, its data is gone and the volume can't be attached to any instance. However, before deletion, you can store a snapshot of the volume, which you can use to recreate the volume later.

This section describes how to delete a volume.

## AWS Management Console

### To delete a volume

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Volumes** in the navigation pane.

The console displays a list of current volumes.

3. Select a volume and click **Delete Volume**.

A confirmation dialog box appears.

4. Click **Yes, Delete**.

The volume is deleted.

## Command Line Interface

To delete a volume, use the `ec2-delete-volume` command.

```
PROMPT> ec2-delete-volume vol-4282672b
```

Amazon EC2 returns information similar to the following example.

```
VOLUME vol-4282672b
```

## API

To delete a volume, use the `DeleteVolume` action. Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=DeleteVolume  
&VolumeId=volume-id  
&AUTHPARAMS
```

The following is an example response.

```
<DeleteVolumeResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <requestId>59dbff89-35bd-4eac-99ed-be587EXAMPLE</requestId>  
  <return>true</return>  
</DeleteVolumeResponse>
```

## Expanding the Storage Space of a Volume

Sometimes, it is necessary for you to increase the storage space of an existing volume without losing the data that is on the volume. This topic explains how to expand the storage space of an Amazon EBS volume by migrating your data to a larger volume, and then extending the file system on the volume to recognize the newly-available space. After you verify that your new volume is working properly, you may delete the old volume.

### Topics

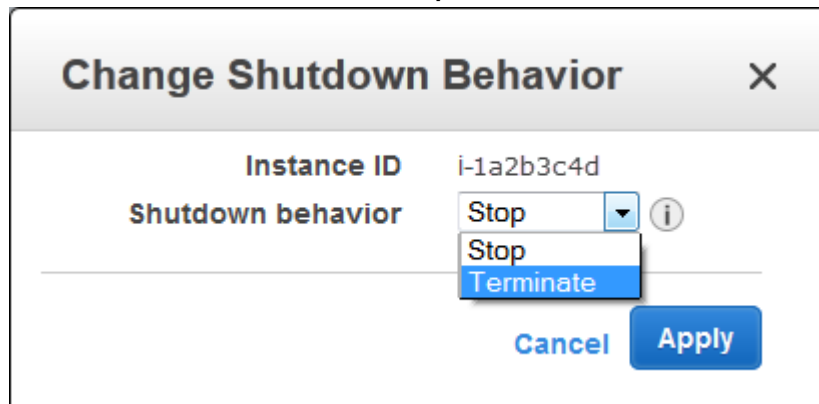
- [Migrating your Data to a Larger Volume](#) (p. 479)
- [Extending a Linux File System](#) (p. 481)
- [Extending a Windows File System](#) (p. 482)
- [Deleting the Old Volume](#) (p. 484)

## Migrating your Data to a Larger Volume

### To migrate your data to a larger Amazon EBS volume

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. If the volume to be expanded is the root volume of an instance, ensure that the instance's **Shutdown Behavior** value is set to **Stop** and not **Terminate**. If it is already set to **Stop**, go on to step 3.
  - a. In the navigation pane, click **Instances**, right-click on the instance to check, and select **Change Shutdown Behavior**.
  - b. If the **Shutdown behavior** is set to **Terminate**, select **Stop** from the list and click **Apply**.

If the **Shutdown behavior** is set to **Stop**, click **Cancel**.



3. Stop the instance that is attached to the volume to be expanded. For more information about how to stop an instance, see [Stopping and Starting Your Instances](#) (p. 285).
4. Create a snapshot of the volume.
  - a. In the navigation pane, click **Volumes**, right-click on the volume to be expanded, and select **Create Snapshot**.
  - b. Enter a **Name** and **Description** for the snapshot, and click **Yes, Create**.
5. Create a new volume from the snapshot.

- a. In the navigation pane, click **Snapshots**. It can take several minutes to create a snapshot.
  - b. When the status of the snapshot just created is set to **completed**, right-click the snapshot and select **Create Volume from Snapshot**.
  - c. In the **Create Volume** dialog box, select the desired **Volume Type**, enter the new **Size** for the volume, set the **Availability Zone** to the same Availability Zone as the instance, and click **Yes, Create**.
6. Detach the old volume.
- a. In the navigation pane, click **Volumes**, select the old root volume from the list of volumes, and make note of the value of *device name* in **Attachment Information**. The **Attachment Information** value takes the following form:

```
instance information:device name
```

- b. Right-click the old root volume and select **Detach Volume**.
  - c. In the **Detach Volume** dialog box, click **Yes, Detach**. It may take several minutes for the volume to be detached.
7. Attach the expanded volume
- a. In the navigation pane, click **Volumes**, select the new root volume from the list of volumes, right-click the new volume, and select **Attach Volume**.
  - b. Select the instance from the **Instances** list, enter the same device name retrieved in step 6, and click **Yes, Attach**.
8. If the instance was stopped in step 3, restart the instance.
- a. In the navigation pane, click **Instances**, right-click the instance, and select **Start**.
  - b. In the **Start Instances** dialog box, select **Yes, Start**. If the instance fails to start, and the volume being expanded is a root volume, verify that you attached the expanded volume using the same device name as the original volume.

### Important

Only EC2-VPC instances with Elastic IP addresses retain their public IP address when they are stopped. If your instance is running in EC2-Classic, the EIP address is disassociated when the instance is stopped, and you must re-associate the EIP after restarting the instance. For more information, see [Elastic IP Addresses \(EIP\) \(p. 428\)](#). If your instance is not using an EIP, then you need to retrieve your new public DNS name for your instance from the Instances page of the Amazon EC2 console to connect to it.

After the instance has started, you should check the file system size to see if your instance recognizes the larger volume space. Use the **df -h** command to check the file system size.

```
# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.9G  943M  6.9G  12% /
tmpfs           1.9G   0    1.9G   0% /dev/shm
```

If the value in the `Size` column does not indicate your newly-expanded volume size, you need to extend the file system on your device so that your instance can use the new space. This procedure varies depending on the operating system running on the instance. For information about how to do this on different operating systems, see the following topics:

- [Extending a Linux File System \(p. 481\)](#)
- [Extending a Windows File System \(p. 482\)](#)

## Extending a Linux File System

In Linux, you use the **resize2fs** command to resize the file system to the larger size of the new volume. This command works even if the volume you wish to extend is the root volume.

### Note

If the volume you are extending has been partitioned, you need to increase the size of the partition before you can resize the file system. This can be accomplished with the `fdisk` or `parted` commands. Because repartitioning can inadvertently result in data loss on a volume, we highly recommend that you create a snapshot of the volume being repartitioned as a backup in case of any data loss.

### To extend a Linux file system

1. Log in to your Linux instance using an SSH client. For more information about connecting to a Linux instance, see [Connecting to Your Linux/UNIX Instances Using SSH \(p. 279\)](#).
2. Use the **df** command to report the existing file system disk space usage. In this example, the `/dev/xvda1` device has already been expanded to 70 GB, but the file system only sees the original 8 GB size.

```
# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      7.9G  943M  6.9G  12% /
tmpfs           1.9G   0    1.9G   0% /dev/shm
```

3. Use the **resize2fs** command to resize the file system to the new size of the volume. This command must be issued from root; otherwise, a permissions error is returned.

```
# resize2fs /dev/xvda1
resize2fs 1.42.3 (14-May-2012)
Filesystem at /dev/xvda1 is mounted on /; on-line resizing required
old_desc_blocks = 1, new_desc_blocks = 5
Performing an on-line resize of /dev/xvda1 to 18350080 (4k) blocks.
The filesystem on /dev/xvda1 is now 18350080 blocks long.
```

4. Use the **df** command to report the existing file system disk space usage, which should now show the full 70 GB.

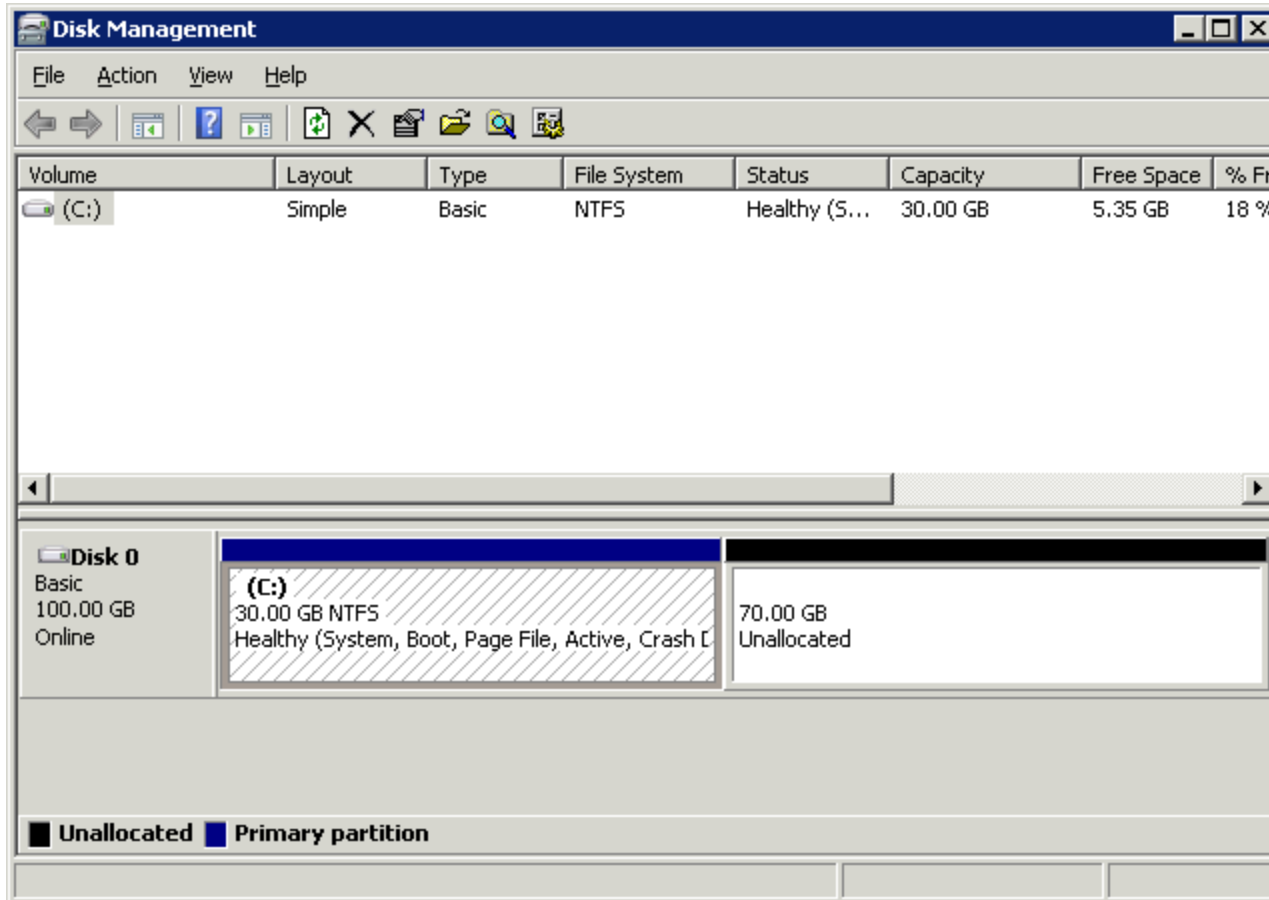
```
# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      69G  951M   68G   2% /
tmpfs           1.9G   0    1.9G   0% /dev/shm
```

## Extending a Windows File System


In Windows, you use the Disk Management utility to extend the disk size to the new size of the volume.

### To extend a Windows file system


1. Log in to your Windows instance using Remote Desktop.
2.
  - Windows Server 2012: Go to the Start screen.
  - Windows Server 2008: On the taskbar, click **Start**, and then click **Run**.
3. Type **diskmgmt.msc** and press **Enter**. The **Disk Management** utility opens.



4. Right-click the expanded drive and select **Extend Volume**.



5. In the Extend Volume Wizard, click **Next**, then set the **Select the amount of space in MB** box to the number of megabytes by which to extend the volume. Normally, you set this to the maximum available space. Complete the wizard.



## Deleting the Old Volume

After the new volume has been attached and extended in the instance, you can delete the old volume if it is no longer needed.

### To delete the old volume

1. In the Amazon EC2 console, click **Volumes** in the navigation pane.
2. Right-click the old volume and select **Delete Volume**.
3. In the **Delete Volume** dialog box, click **Yes, Delete**.

## Amazon EBS Snapshots

An Amazon EBS snapshot is a point-in-time backup copy of an Amazon EBS volume that is stored in Amazon S3. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. When you delete a snapshot, only the data exclusive to that snapshot is removed. Active snapshots contain all of the information needed to restore your data (from the time the snapshot was taken) to a new Amazon EBS volume.

### Topics

- [Creating an Amazon EBS Snapshot \(p. 485\)](#)
- [Deleting an Amazon EBS Snapshot \(p. 487\)](#)
- [Copying an Amazon EBS Snapshot \(p. 488\)](#)
- [Describing Snapshots \(p. 491\)](#)
- [Sharing Snapshots \(p. 492\)](#)

When you create a new Amazon EBS volume, you may create it based on an existing snapshot; the new volume begins as an exact replica of the original volume that was used to create the snapshot. New volumes created from existing Amazon S3 snapshots load lazily in the background, so you can begin using them right away. If your instance accesses a piece of data which hasn't yet been loaded, the volume immediately downloads the requested data from Amazon S3, and then continues loading the rest of the volume's data in the background. For more information about creating snapshots, see [Creating an Amazon EBS Snapshot \(p. 485\)](#).

You can share your snapshots with specific individuals, or make them public to share them with the entire AWS community. Users with access to your snapshots can create their own Amazon EBS volumes from your snapshot, but your snapshots remain completely intact. For more information about how to share snapshots, see [Sharing Snapshots \(p. 492\)](#).

Amazon EBS snapshots are constrained to the region in which they are created. However, you may copy snapshots across AWS regions, making it easier to leverage multiple AWS regions for geographical expansion, data center migration and disaster recovery. You may copy any accessible snapshots that are in the "available" status. For more information, see [Copying an Amazon EBS Snapshot \(p. 488\)](#).

## Creating an Amazon EBS Snapshot

### Topics

- [AWS Management Console \(p. 486\)](#)
- [Command Line Interface \(p. 486\)](#)
- [API \(p. 486\)](#)

After writing data to an Amazon EBS volume, you can periodically create a snapshot of the volume to use as a baseline for new volumes or for data backup. If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed after your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.

Snapshots occur asynchronously and the status of the snapshot is `pending` until the snapshot is complete.

By default, only you can launch volumes from snapshots that you own. However, you can choose to share your snapshots with specific AWS accounts or make them public. For more information, see [Sharing Snapshots \(p. 492\)](#).

When a snapshot is created, any AWS Marketplace product codes from the volume are propagated to the snapshot.

You can take a snapshot of an attached volume that is in use. However, snapshots only capture data that has been written to your Amazon EBS volume at the time the snapshot command is issued. This may exclude any data that has been cached by any applications or the operating system. If you can pause any file writes to the volume long enough to take a snapshot, your snapshot should be complete. However, if you can't pause all file writes to the volume, you should unmount the volume from within the instance, issue the snapshot command, and then remount the volume to ensure a consistent and complete snapshot. You may remount and use your volume while the snapshot status is `pending`.

To create a snapshot for Amazon EBS volumes that serve as root devices, you should stop the instance before taking the snapshot.

To unmount the volume in Linux/UNIX, use the following command:

```
umount -d device_name
```



Where *device\_name* is the device name (for example, `/dev/sdh`).

To unmount the volume in Windows, open Disk Management, right-click the volume to unmount, and select **Change Drive Letter and Path**. Select the mount point to remove, and then click **Remove**.

## AWS Management Console

### To create a snapshot

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Snapshots** in the navigation pane.

The console displays a list of current snapshots.

3. Click **Create Snapshot**.

The Create Snapshot dialog box appears.

4. Select the volume to create a snapshot for and click **Create**.

Amazon EC2 begins creating the snapshot.

## Command Line Interface

To create a snapshot, use the `ec2-create-snapshot` command.

```
PROMPT> ec2-create-snapshot volume_id
```

Amazon EC2 returns information similar to the following example.

```
SNAPSHOT          snap-88fe11e0   vol-c7f95aae   pending 2010-03-30T14:10:38+0000
111122223333      80
```

When the snapshot is complete, its status changes to `completed`.

## API

To create a snapshot, use the `CreateSnapshot` action. Construct the following request.

```
https://ec2.amazonaws.com/
?Action=CreateSnapshot
&VolumeId=volume-id
&AUTHPARAMS
```

The following is an example response.

```
<CreateSnapshotResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <requestId>fe0d60a3-b804-484e-b558-92518bebe7af</requestId>
  <snapshotId>snap-05b4aa6c</snapshotId>
  <volumeId>vol-d11fbbb8</volumeId>
  <status>pending</status>
  <startTime>2010-03-23T09:43:51.000Z</startTime>
  <progress/>
  <ownerId>999988887777</ownerId>
  <volumeSize>20</volumeSize>
```

```
<description/>  
</CreateSnapshotResponse>
```

## Deleting an Amazon EBS Snapshot

### Topics

- [AWS Management Console](#) (p. 487)
- [Command Line Interface](#) (p. 487)
- [API](#) (p. 487)

This section describes how to delete a snapshot.

### Note

- If you make periodic snapshots of a volume, the snapshots are incremental so that only the blocks on the device that have changed since your last snapshot are saved in the new snapshot. Even though snapshots are saved incrementally, the snapshot deletion process is designed so that you need to retain only the most recent snapshot in order to restore the volume.
- You cannot delete a snapshot of the root device of an Amazon EBS volume used by a registered AMI. You must first de-register the AMI before you can delete the snapshot.

## AWS Management Console

### To delete a snapshot

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Snapshots** in the navigation pane.

The console displays a list of current snapshots.

3. Select a snapshot and click **Delete Snapshot**.

A confirmation dialog box appears.

4. Click **Yes, Delete**.

The snapshot is deleted.

## Command Line Interface

To delete a snapshot, use the `ec2-delete-snapshot` command.

```
PROMPT> ec2-delete-snapshot snapshot_id
```

Amazon EC2 returns information similar to the following example.

```
SNAPSHOT snap-78a54011
```

## API

To delete a snapshot, use the `DeleteSnapshot` action. Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=DeleteSnapshot  
&SnapshotId=snapshot-id  
&AUTHPARAMS
```

The following is an example response.

```
<DeleteSnapshotResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">  
  <requestId>59dbff89-35bd-4eac-99ed-be587EXAMPLE</requestId>  
  <return>true</return>  
</DeleteSnapshotResponse>
```

## Copying an Amazon EBS Snapshot

### Topics

- [AWS Management Console \(p. 489\)](#)
- [Command Line Interface \(p. 490\)](#)
- [API \(p. 490\)](#)

Amazon Elastic Block Store (Amazon EBS) provides block level storage volumes for use with Amazon Elastic Compute Cloud (Amazon EC2) instances. With Amazon EBS, you can create point-in-time snapshots of volumes and store them on Amazon Simple Storage Service (Amazon S3). After you've stored a snapshot in Amazon S3, you can copy it from one AWS region to another, or within the same region, using the Amazon EC2 console, Amazon EC2 CLI, or the API. You can use snapshot copies to create backups of data, to create new Amazon EBS volumes, or to create Amazon Machine Images (AMIs).

You can have up to five snapshot copy requests in progress to a single destination per account. You can copy any accessible Amazon EBS snapshots that have "completed" status, including shared snapshots and snapshots that you've created. You can also copy AWS Marketplace, VM Import/Export, and AWS Storage Gateway snapshots, but you must verify that the snapshot is supported in the destination region.

The first snapshot copy of a volume is always a full copy. Each subsequent snapshot copy is incremental, meaning that only the blocks on the volume that have changed since your last snapshot copy to the same destination are transferred. Incremental snapshots make the copy process faster. Support for incremental snapshots is specific to a region pair. For example, if you copy a snapshot from the US East (Northern Virginia) Region to the US West (Oregon) Region, the first snapshot copy of the volume is a full copy. However, subsequent snapshot copies of the same volume transferred between the same regions are incremental. A snapshot copy can only be done incrementally as long as there is one full copy of the same volume available in the destination region.

### Note

You cannot copy an Amazon Relational Database Service (Amazon RDS) snapshot.

When you copy a snapshot, you are only charged for the data transfer and storage used to copy the snapshot data across regions and to store the copied snapshot in the destination region. You are not charged if the snapshot copy fails. However, if you cancel a snapshot copy that is not yet complete, or delete the source snapshot while the copy is in progress, you are charged for the bandwidth of the data transferred. Unlike the create snapshot operation, which is incremental, the copy snapshot operation copies all of the bytes in the snapshot every time. The snapshot is copied across regions using the secure Amazon S3 Copy and the snapshot copy receives a snapshot ID that's different from the original snapshot's ID.

You can use a copy of an Amazon EBS snapshot in the following ways:

- **Geographic Expansion:** You can launch your applications in a new region.
- **Migration:** You can migrate an application to a new region, to enable better availability and minimizing cost.
- **Disaster Recovery:** You can back up your data and logs across different geographical locations at regular intervals. In case of disaster, you can restore your applications using point-in-time backups stored in the secondary region. This minimizes data loss and recovery time.

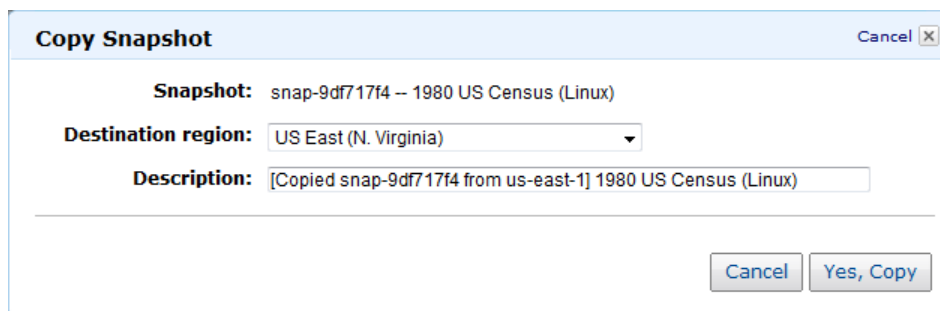
The Amazon EC2 console, Amazon EC2 CLI, and the API are designed to provide an intuitive customer experience. We use the push model in the console design to minimize user clicks for the Amazon EBS snapshot use cases discussed earlier. You can easily initiate a copy from the console by starting with the source region. We use a pull model in the Amazon EC2 CLI and the API, because these experiences factor in how customers use automation. You only need to know the source snapshot ID and source region to initiate the copy using the Amazon EC2 CLI or API.

## AWS Management Console

### To copy a snapshot using the Amazon EC2 console

You can create a copy of an Amazon EBS snapshot using the Amazon EC2 console.

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, under **ELASTIC BLOCK STORE**, click **Snapshots**.
3. In the **EBS Snapshots** pane, right-click the snapshot to copy, and then click **Copy Snapshot**.
4. In the **Copy Snapshot** dialog box, update the following as necessary:
  - **Snapshot:** Select another snapshot, if appropriate.
  - **Destination region:** Select the region where you want to write the copy of the snapshot.
  - **Description:** By default, the description includes information about the source snapshot so that you can identify a copy from the original. You can change this description as necessary.
5. Click **Yes, Copy**.



**Copy Snapshot** Cancel

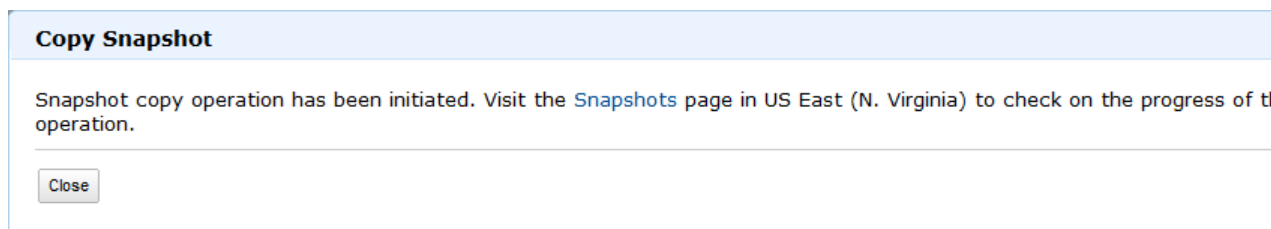
**Snapshot:** snap-9df717f4 -- 1980 US Census (Linux)

**Destination region:** US East (N. Virginia)

**Description:** [Copied snap-9df717f4 from us-east-1] 1980 US Census (Linux)

Cancel Yes, Copy

6. In the **Copy Snapshot** confirmation dialog box, you can click **Snapshots** to go to the **EBS Snapshots** pane in the region specified, or click **Close**.



**Copy Snapshot**

Snapshot copy operation has been initiated. Visit the [Snapshots](#) page in US East (N. Virginia) to check on the progress of the operation.

Close

To view the progress of the copy process later, switch the Amazon EC2 console to the destination region, and then refresh the **EBS Snapshots** pane. Copies in progress are listed at the top of the **EBS Snapshots** pane.

## Command Line Interface

You can create a copy of an Amazon EBS snapshot using the `ec2-copy-snapshot` command. You can view the progress of the copy process using the `ec2-describe-snapshots` command. For more information, see [ec2-copy-snapshot](#) and [ec2-describe-snapshots](#).

The Amazon EC2 CLI tools share a `-region` parameter that defines the region to which the API call will be directed. When copying a snapshot using the `ec2-snapshot-copy` command, you can use the `-region` parameter to specify the destination region.

### To copy a snapshot using the CLI

1. At a command prompt, switch to the destination region, and then type the following:

```
ec2-copy-snapshot -r source-region -s source-snapshot-id [-d description]  
[-region destination-region]
```

2. To track the progress of the copy, at a command prompt, switch to the destination region, and then type the following:

```
ec2-describe-snapshots [snapshot_id ...] [-a] [-o owner...] [-r user_id]  
[["--filter"name=value"] ...]
```

## API

### To copy a snapshot using the API

- Use the [CopySnapshot](#) action to construct the following request:

```
https://ec2.amazonaws.com/?Action=CopySnapshot  
&AWSAccessKeyId=AKIAIOSFODNN7EXAMPLE  
&Description=My%20snapshot  
&Signature=VjpsFePIKxDc1IUy92W3SBApdLiap7nno4pEc9iEXAMPLE  
&SignatureMethod=HmacSHA256  
&SignatureVersion=2  
  
&SourceRegion=us-west-1  
  
&SourceSnapshotId=snap-2EXAMPLE  
&Timestamp=2012-12-11T02%3A03%3A35.713Z  
&Version=2012-12-01
```

The following is an example response.

```
<CopySnapshotResponse xmlns="http://ec2.amazonaws.com/doc/2012-12-01/">  
<requestId>60bc441d-fa2c-494d-b155-5d6a3EXAMPLE</requestId>  
<snapshotId>snap-8EXAMPLE</snapshotId>  
</CopySnapshotResponse>
```

## Describing Snapshots

### Topics

- [AWS Management Console](#) (p. 491)
- [Command Line Interface](#) (p. 491)
- [API](#) (p. 491)

This section describes how to view snapshots that you have created.

### AWS Management Console

#### To describe snapshots

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Snapshots** in the navigation pane.

The console displays a list of all snapshots to which you have access, and their status.

3. To reduce the list, select an option from the **Viewing** list. For example, to view only your snapshots, select **Owned by Me**.

The console displays a new list of snapshots.

4. To view more information about a snapshot, select it.

Information about the snapshot appears in the lower pane.

### Command Line Interface

To describe snapshots, use the `ec2-describe-snapshots` command.

```
PROMPT> ec2-describe-snapshots snap-78a54011
```

Amazon EC2 returns information about the snapshot.

```
SNAPSHOT snap-78a54011 vol-4d826724 pending 2008-02-15T09:03:58+0000 60%
```

If the snapshot is in the process of being created, the status is `pending` and a percentage complete is displayed (e.g., 60%). After the snapshot is complete, its status changes to `completed`.

#### Tip

If you have a large number of snapshots, you can use `ec2-describe-snapshots` to filter the results to only the snapshots that match the criteria you specify.

### API

To describe snapshots, use the `DescribeSnapshots` action. Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=DescribeSnapshots  
&SnapshotId=snapshot-id  
&AUTHPARAMS
```

The following is an example response.

```
<DescribeSnapshotsResponse xmlns="http://ec2.amazonaws.com/doc/2013-10-01/">
  <requestId>26245bae-2c70-4846-82d3-65784e298820</requestId>
  <snapshotSet>
    <item>
      <snapshotId>snap-05b4aa6c</snapshotId>
      <volumeId>vol-d11fbbb8</volumeId>
      <status>completed</status>
      <startTime>2010-03-23T09:43:52.000Z</startTime>
      <progress>100%</progress>
      <ownerId>999988887777</ownerId>
      <volumeSize>20</volumeSize>
      <description/>
    </item>
  </snapshotSet>
</DescribeSnapshotsResponse>
```

### Tip

If you have a large number of snapshots, you can use `DescribeSnapshots` to filter the list to only the snapshots that match criteria you specify.

## Sharing Snapshots

### Topics

- [AWS Management Console \(p. 492\)](#)
- [Command Line Interface \(p. 493\)](#)
- [API \(p. 493\)](#)

This section describes how to share your Amazon EBS snapshots with your co-workers or others in the AWS community by modifying the permissions of the snapshot. Users that you have authorized can quickly use your Amazon EBS shared snapshots as the basis for creating their own Amazon EBS volumes. If you choose, you can also make your data available publicly to all AWS users. Users to whom you have granted access can create their own Amazon EBS volumes based on your snapshot and your original snapshot remains intact.

### Note

Snapshots are constrained to the region in which they are created. If you would like to share a snapshot with another region, you need to copy the snapshot to that region. For more information about copying snapshots, see [Copying an Amazon EBS Snapshot \(p. 488\)](#).

### Important

When you share a snapshot (whether by sharing it with another AWS account or making it public to all), you are giving others access to all the data on your snapshot. Share snapshots only with people with whom you want to share *all* your snapshot data.

## AWS Management Console

### To modify snapshot permissions

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Click **Snapshots** in the navigation pane.

The console displays a list of current snapshots and their status.

3. Select a snapshot and click **Permissions**.

The **Modify Snapshot Permissions** dialog box appears.

4. Choose whether to make the snapshot public or to share it with select AWS accounts:

### Important

Making your snapshot public shares all snapshot data with everyone. Snapshots with AWS Marketplace product codes cannot be made public.

- To make the snapshot public, select **Public** and click **Save**.
- To expose the snapshot only to specific AWS accounts, select **Private**, enter the IDs of those AWS accounts, and click **Save**.

The console modifies permissions for the snapshot.

## Command Line Interface

### To modify snapshot permissions

1. Use the `ec2-describe-snapshot-attribute` command to first describe the snapshot's permissions.

```
PROMPT> ec2-describe-snapshot-attribute snap_id --create-volume-permission
```

If there are no permissions set on the snapshot, the output is empty.

2. Choose whether to make the snapshot public, or to share it with a specific AWS account.

### Important

Making your snapshot public shares all snapshot data with everyone. Snapshots with AWS Marketplace product codes cannot be made public.

- To make the snapshot public, use the `ec2-modify-snapshot-attribute` command as follows.

```
PROMPT> ec2-modify-snapshot-attribute snap_id -c --add all
```

Amazon EC2 returns permission information for the snapshot.

```
createVolumePermission snap_id ADD group all
```

- To share the snapshot with a particular AWS account, use the `ec2-modify-snapshot-attribute` command as follows

```
PROMPT> ec2-modify-snapshot-attribute snap_id -c --add account_id
```

Amazon EC2 returns permission information for the snapshot.

```
createVolumePermission snap_id ADD account_id
```

## API

### To modify snapshot permissions

1. Use the `DescribeSnapshotAttribute` action to describe the snapshot's permissions. Construct the following request.



```
https://ec2.amazonaws.com/  
?Action=DescribeSnapshotAttribute  
&SnapshotId=snapshot-id  
&AUTHPARAMS
```

The following is an example response.

```
<DescribeSnapshotAttributeResponse xmlns="http://ec2.amazonaws.com/doc/2013-  
10-01/">  
  <requestId>d0d21738-e3da-4077-947d-c9e48472d831</requestId>  
  <snapshotId>snap-05b4aa6c</snapshotId>  
  <createVolumePermission/>  
</DescribeSnapshotAttributeResponse>
```

2. Choose whether to make the snapshot public, or to share it with a specific AWS account.

### Important

Making your snapshot public shares all snapshot data with everyone. Snapshots with AWS Marketplace product codes cannot be made public.

- To make the snapshot public, use the [ModifySnapshotAttribute](#) action. Construct the following request.

```
https://ec2.amazonaws.com/  
?Action=ModifySnapshotAttribute  
&SnapshotId=snapshot-id  
&CreateVolumePermission.Add.1.Group=all  
&AUTHPARAMS
```

- To share the snapshot with a particular AWS account, use the [ModifySnapshotAttribute](#) action as follows.

```
https://ec2.amazonaws.com/  
?Action=ModifySnapshotAttribute  
&SnapshotId=snapshot-id  
&CreateVolumePermission.Add.1.UserId=111122223333  
&AUTHPARAMS
```

## Amazon EBS Volume Performance

Several factors can affect the performance of Amazon EBS volumes, such as instance configuration, I/O characteristics, workload demand, and storage configuration. After you learn the basics of working with EBS volumes, it's a good idea to look at the I/O performance you require and at your options for increasing EBS performance to meet those requirements.

### Topics

- [Amazon EBS Performance Tips](#) (p. 495)
- [Amazon EC2 Instance Configuration](#) (p. 495)
- [I/O Characteristics](#) (p. 497)
- [Workload Demand](#) (p. 497)

- [Pre-Warming Amazon EBS Volumes \(p. 498\)](#)
- [RAID Configuration \(p. 501\)](#)
- [Benchmark Volumes \(p. 503\)](#)

## Amazon EBS Performance Tips

- When you consider the performance requirements for your EBS storage application, it is important to start with an EC2 configuration that is optimized for EBS and that can handle the bandwidth that your application storage system requires. For more information, see [Amazon EC2 Instance Configuration \(p. 495\)](#).
- When you measure the performance of your EBS volumes, especially with provisioned IOPS volumes, it is important to understand the units of measure involved and how performance is calculated. For more information, see [I/O Characteristics \(p. 497\)](#).
- There is a relationship between the maximum performance of your EBS volumes, the amount of I/O you are driving to them, and the amount of time it takes for each transaction to complete. Each of these factors (performance, I/O, and time) affects the others, and different applications are more sensitive to one factor or another. For more information, see [Workload Demand \(p. 497\)](#).
- There is a 5 to 50 percent reduction in IOPS when you first access each block of data on a newly created or restored EBS volume. You can avoid this performance hit by accessing each block in advance. For more information, see [Pre-Warming Amazon EBS Volumes \(p. 498\)](#).
- Some instance types can handle more IOPS than you can provision for a single volume. You can join multiple provisioned IOPS volumes together in a RAID 0 configuration to use the available bandwidth for these instances. You can also provide redundancy for your volumes with a RAID 1 (mirrored) configuration. For more information, see [RAID Configuration \(p. 501\)](#).
- You can benchmark your storage and compute configuration to make sure you achieve the level of performance you expect to see before taking your application live. For more information, see [Benchmark Volumes \(p. 503\)](#).
- Amazon Web Services provides performance metrics for EBS that you can analyze and view with Amazon CloudWatch and status checks that you can use to monitor the health of your volumes. For more information, see [Monitoring the Status of Your Volumes \(p. 464\)](#).
- Frequent snapshots provide a higher level of data durability, but they also degrade the performance of your application while the snapshot is in progress. This trade off becomes critical when you have data that changes rapidly. Whenever possible, plan for snapshots to occur during off-peak times in order to minimize workload impact. For more information, see [Amazon EBS Snapshots \(p. 484\)](#).

## Amazon EC2 Instance Configuration

When you plan and configure EBS volumes for your application, it is important to consider the configuration of the instances that you will attach the volumes to. In order to get the most performance out of your EBS volumes, you should attach them to an EBS-optimized instance that can support the bandwidth available to the volumes. This is especially important when you use provisioned IOPS volumes, or when you stripe multiple volumes together in a RAID configuration.

### Use EBS-Optimized Instances

Any performance-sensitive workloads that require minimal variability and dedicated EC2 to EBS traffic, such as production databases or business applications, should use provisioned IOPS volumes that are attached to an EBS-optimized instance. EC2 instances that are not launched as EBS-optimized instances offer no guarantee of network resources. The only way to ensure sustained reliable network bandwidth between your EC2 instance and your EBS volumes is to launch the EC2 instance as EBS-optimized. Provisioned IOPS volumes should always be attached to EBS-optimized instances because they are designed to deliver the expected performance only when attached to EBS-optimized instances.

### Choose an EC2 Instance with Enough Bandwidth

Launching an instance that is EBS-optimized provides you with a dedicated connection between your EC2 instance and your EBS volume. However, it is still possible to provision EBS volumes that exceed the available bandwidth for certain instance types, especially when multiple volumes are striped in a RAID configuration. The following table shows which instance types are available to be launched as EBS-optimized and the maximum amount of IOPS the instance can support. Be sure to choose an EBS-optimized instance that provides dedicated EBS throughput that exceeds the amount of EBS bandwidth needed for your application; otherwise, the EBS to EC2 connection will become a performance bottleneck.

Instance Type	EBS-optimized Available	Maximum Throughput (MB/s)*	Max 16K IOPS
t1.micro	No	32 MB/s	2,000
m1.small	No	64 MB/s	4,000
m1.medium	No	64 MB/s	4,000
m1.large	Yes	64 MB/s	4,000
m1.xlarge	Yes	128 MB/s	8,000
m3.xlarge	Yes	64 MB/s	4,000
m3.2xlarge	Yes	128 MB/s	8,000
c1.medium	No	32 MB/s	2,000
c1.xlarge	Yes	128 MB/s	8,000
cc2.8xlarge	N/A**	800 MB/s	48,000
m2.xlarge	No	64 MB/s	4,000
m2.2xlarge	Yes	64 MB/s	4,000
m2.4xlarge	Yes	128 MB/s	8,000
cr1.8xlarge	N/A**	800 MB/s	48,000
hi1.4xlarge	N/A**	800 MB/s	48,000
hs1.8xlarge	N/A**	800 MB/s	48,000
g2.2xlarge	Yes	128 MB/s	8,000
cg1.4xlarge	N/A**	800 MB/s	48,000

The `t1.micro` instance has a maximum 16 KB IOPS value of 2,000, but since this instance type is not available to be launched as EBS-optimized, that value is an absolute best-case scenario and is not guaranteed; to guarantee 2,000 16 KB IOPS, we must use an instance that supports EBS optimization. The `m1.large` instance type does support EBS optimization, but the maximum 16 KB IOPS available for this instance type is 4,000; if two 4,000 provisioned IOPS EBS volumes are attached to this instance in a RAID configuration, the EC2 to EBS connection bandwidth limit will keep these volumes from providing the maximum IOPS that were provisioned for them. In this case, we must use an EBS-optimized EC2 instance that supports 8,000 16 KB IOPS, such as the `m1.xlarge` instance type.

Several instance types in the table above show "N/A" in the EBS-optimized column; these instance types are EBS-optimized by default, so there are no special options or flags required at launch time to enable the dedicated EC2 to EBS bandwidth. Since the maximum provisioned IOPS value for EBS volumes is

4,000, you would have to use multiple EBS volumes simultaneously to reach the level of I/O performance available to these instance types.

You must use EBS-optimized instances to get the full performance benefits of Amazon EBS Provisioned IOPS volumes. For more information, see [EBS-Optimized Instances \(p. 107\)](#).

## I/O Characteristics

On a given volume configuration, certain I/O characteristics drive the performance behavior on the back end. Provisioned IOPS volumes deliver consistent performance whether an I/O operation is random or sequential, and also whether an I/O operation is to read or write data. I/O size, however, does make an impact on IOPS because of the way they are measured. In order to fully understand provisioned IOPS volumes and how they will perform in your application, it important to know what IOPS are and how they are measured.

### What are IOPS?

IOPS are input/output operations per second. Amazon EBS measures these I/O operations in 16 KB chunks. When you provision a 4,000 IOPS volume and attach it to an EBS-optimized instance that can provide the necessary bandwidth, you can transfer 4,000 16 KB chunks of data per second (for a bandwidth of approximately 64 MBs or 500 Mbps).

This configuration could transfer 2,000 32 KB chunks, or 1,000 64 KB chunks of data per second as well, before it reached a bandwidth limit that equals 4000 16 KB I/O operations per second. If your I/O chunks are greater than 16 KB, you may experience a smaller number of IOPS than you provisioned, but the bandwidth you experience should equal the amount that would be used by your provisioned amount of 16 KB IOPS.

For 16 KB or smaller I/O operations, you should see the amount of IOPS that you have provisioned, provided that you are driving enough I/O to keep the drives busy. For smaller I/O operations, you may even see an IOPS value that is higher than what you have provisioned (when measured on the client side), and this is because the client may be coalescing multiple smaller I/O operations into a single 16 KB chunk.

If you are not experiencing the expected IOPS or throughput you have provisioned, ensure that your EC2 bandwidth is not the limiting factor; your instance should be EBS-optimized and your instance type EBS dedicated bandwidth should exceed the IOPS you have provisioned. For more information, see [Amazon EC2 Instance Configuration \(p. 495\)](#). Another possible cause for not experiencing the expected IOPS is that you are not driving enough I/O to the EBS volumes. For more information, see [Workload Demand \(p. 497\)](#).

## Workload Demand

Workload demand plays an important role in getting the most out of your provisioned IOPS volumes. In order for your volumes to deliver the amount of IOPS that you have provisioned, they need to have enough I/O requests sent to them. There is a relationship between the demand on the volumes, the amount of IOPS that are provisioned for them, and the latency of the request (the amount of time it takes for the I/O operation to complete).

### Average Queue Length

The queue length is the number of pending I/O requests for a device. Optimal average queue length will vary for every customer workload, and this value depends on your particular application's sensitivity to IOPS and latency. If your workload is not delivering enough I/O requests to maintain your optimal average queue length, then your volume might not consistently deliver the IOPS that you have provisioned. However, if your workload maintains an average queue length that is higher than your optimal value, then your per-request I/O latency will increase; in this case, you should provision more IOPS for your volume. We recommend that you target an optimal average queue length of 1 for every 200 Provisioned IOPS

and tune that value based on your application requirements. For example, a volume with 1000 Provisioned IOPS should target an average queue length of 5.

**Note**

Per-request I/O latency may increase with higher average queue lengths.

**Latency**

Latency is the true end-to-end client time of an I/O operation; in other words, when the client sends a IO, how long does it take to get an acknowledgement from the storage subsystem that the IO read or write is complete. If your I/O latency is higher than you require, check your average queue length to make sure that your application is not trying to drive more IOPS than you have provisioned. You can maintain high IOPS while keeping latency down by maintaining a low average queue length (which is achieved by provisioning more IOPS for your volume).

## Pre-Warming Amazon EBS Volumes

When you create a new EBS volume or restore a volume from a snapshot, the back-end storage blocks are allocated to you immediately. However, the first time you access a block of storage, it must be either wiped clean (for new volumes) or instantiated from its snapshot (for restored volumes) before you can access the block. This precursory action takes time and can cause a 5 to 50 percent loss of IOPS for your volume the first time each block is accessed. Performance is restored after the data is accessed once.

You can avoid this performance hit in a production environment by writing to or reading from all of the blocks on your volume before you use it; this process is called *pre-warming*. Writing to all of the blocks on a volume is preferred, but that is not an option for volumes that were restored from a snapshot, because that would overwrite the restored data. For a completely new volume that was created from scratch, you should write to all blocks before using the volume. For a new volume created from a snapshot, you should read all the blocks that have data before using the volume.

**Topics**

- [Pre-warming Amazon EBS Volumes on Linux \(p. 498\)](#)
- [Pre-Warming Amazon EBS Volumes on Windows \(p. 500\)](#)

## Pre-warming Amazon EBS Volumes on Linux

### To Pre-Warm a New Volume on Linux

For a new volume, use the **dd** command to write to all blocks on a volume. This procedure will write zeroes all of the blocks on the device and it will prepare the volume for use. Any existing data on the volume will be lost.

1. Attach the newly created volume to your Linux instance.
2. Use the **lsblk** command to list the block devices on your instance.

```
[ec2-user ~]$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvdf  202:80  0  30G  0 disk
xvda1 202:1   0   8G  0 disk /
```

Here you can see that the new volume, `/dev/xvdf`, is attached, but not mounted (because there is no path listed under the `MOUNTPOINT` column).

**Important**

If your new volume is mounted, unmount it. The next step should not be performed on a mounted device.

3. Use the **dd** command to write to all of the blocks on the device. The **if** (input file) parameter should be set to one of the Linux virtual devices, such as `/dev/zero`. The **of** (output file) parameter should be set to the drive you wish to warm. The **bs** parameter sets the block size of the write operation; for optimal performance, this should be set to 1 MB.

```
[ec2-user ~]$ sudo dd if=/dev/zero of=/dev/xvdf bs=1M
```

**Note**

This step may take several minutes up to several hours, depending on your EC2 instance bandwidth, the IOPS provisioned for the volume, and the size of the volume.

When the operation is finished, you will see a report of the write operation.

```
dd: writing `'/dev/xvdf': No space left on device
15361+0 records in
15360+0 records out
32212254720 bytes (32 GB) copied, 2449.12 s, 13.2 MB/s
```

Your volume is now ready for use. For more information, see [Making an Amazon EBS Volume Available for Use \(p. 459\)](#).

### To Pre-Warm a Volume Restored from a Snapshot on Linux

For volumes that have been restored from snapshots, use the **dd** command to read from all blocks on a volume. This procedure will read all of the blocks on the device and it will prepare the volume for use. All existing data on the volume will be preserved.

1. Attach the newly restored volume to your Linux instance.
2. Use the **lsblk** command to list the block devices on your instance.

```
[ec2-user ~]$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvdf  202:80  0 30G  0 disk
xvda1 202:1   0  8G  0 disk /
```

Here you can see that the new volume, `/dev/xvdf`, is attached, but not mounted (because there is no path listed under the `MOUNTPOINT` column).

**Note**

The following step can be performed on a mounted or unmounted device.

3. Use the **dd** command to read all of the blocks on the device. The **if** (input file) parameter should be set to the drive you wish to warm. The **of** (output file) parameter should be set to the Linux null virtual device, `/dev/null`. The **bs** parameter sets the block size of the read operation; for optimal performance, this should be set to 1 MB.

```
[ec2-user ~]$ sudo dd if=/dev/xvdf of=/dev/null bs=1M
```

**Note**

This step may take several minutes up to several hours, depending on your EC2 instance bandwidth, the IOPS provisioned for the volume, and the size of the volume.

When the operation is finished, you will see a report of the read operation.

```
15360+0 records in
15360+0 records out
32212254720 bytes (32 GB) copied, 2480.02 s, 13.0 MB/s
```

Your volume is now ready for use. For more information, see [Making an Amazon EBS Volume Available for Use](#) (p. 459).

## Pre-Warming Amazon EBS Volumes on Windows

There are multiple ways to pre-warm EBS volumes on Windows. The most simple solution is to provide a full format of the volume. Use the following command to perform a full format of a new volume:

### Warning

The following command will destroy any existing data on the volume.

```
C:\>format drive_letter: /p:1
```

You can also perform a full format by right-clicking on the drive in a Windows Explorer window and clicking **Format**. Because this operation destroys all data on the volume, it is only appropriate for *new* volumes. For a more powerful pre-warming tool, that allows you to pre-warm volumes that have been restored from a snapshot or that contain existing data, you should consider **dd** for Windows.

### Install dd for Windows

The **dd** for the Windows program provides a similar experience to the **dd** program that is commonly available for Linux and UNIX systems, and it allows you to pre-warm both new EBS volumes and volumes that have been restored from snapshots. At the time of this writing, the most recent beta version contains the `/dev/null` virtual device that is required to pre-warm volumes restored from snapshots. Full documentation for the program is available at <http://www.chrysocome.net/dd>.

1. Download the most recent binary version of **dd** for Windows from <http://www.chrysocome.net/dd>. You must use version 0.6 beta 3 or newer to pre-warm restored volumes.
2. (Optional) Create a folder for command line utilities that is easy to locate and remember, such as `C:\bin`. If you already have a designated folder for command line utilities, you can use that folder instead in the following step.
3. Unzip the binary package and copy the `dd.exe` file to your command line utilities folder (for example, `C:\bin`).
4. Add the command line utilities folder to your `Path` environment variable so you can execute the programs in that folder from anywhere.

### Important

The following steps don't update the environment variables in your current command prompt windows. The command prompt windows that you open after you complete these steps will contain the updates. This is why it's necessary for you to open a new command prompt window to verify that your environment is set up properly.

- a. Click **Start**, right-click **Computer**, and then click **Properties**.
- b. Click **Advanced system settings**.
- c. Click **Environment Variables**.
- d. Under **System Variables**, click the variable called **Path** and then click **Edit**.
- e. In **Variable value**, append a semicolon and the location of your command line utility folder (`;%C:\bin\`) to the end of the existing value.
- f. Click **OK** to close the **Edit System Variable** window.



## To Pre-Warm a Volume Using dd for Windows

1. Use the `dd --list` command to list the available devices on your system.

```
C:\>dd --list
rawwrite dd for windows version 0.6beta3.
Written by John Newbigin <jn@it.swin.edu.au>
This program is covered by terms of the GPL Version 2.

Win32 Available Volume Information
\\.\Volume{6c195fc1-0511-11e3-9220-806e6f6e6963}\
  link to \\?\Device\HarddiskVolume1
  fixed media
  Mounted on \\.\c:

\\.\Volume{c608dec6-4280-11e3-89dd-123141015292}\
  link to \\?\Device\HarddiskVolume2
  fixed media
  Mounted on \\.\z:
```

You can use the Volume name, the Device name, or the drive letter (for example, `\\.\z:`) to identify the device you want to pre-warm in the following steps. If you want to completely wipe the drive, proceed to [Step 2 \(p. 501\)](#); if you want to preserve the data on a volume, proceed to [Step 3 \(p. 501\)](#).

2. (Optional: destructive pre-warm command) Execute the following command to write zeros (from the virtual device, `/dev/zero`) to the specified device that you identified in [Step 1 \(p. 501\)](#). The `bs=1M` flag uses a one megabyte block size (which speeds up the write process) and the `--progress` flag provides a status update as the write operation is conducted.

### Warning

The following command will destroy any existing data on the volume. For volumes with existing data (such as volumes restored from snapshots), use the command in [Step 3 \(p. 501\)](#).

```
C:\>dd if=/dev/zero of=\\.\drive_letter: bs=1M --progress
```

3. (Optional: read-only pre-warm command) Execute the following command to read all blocks on the specified device (and send the output to the `/dev/null` virtual device). This command will preserve any data on the volume.

```
C:\>dd if=\\.\drive_letter: of=/dev/null bs=1M --progress
```

4. When the operation completes, you are ready to use your new volume. For more information, see [Making an Amazon EBS Volume Available for Use \(p. 459\)](#).

## RAID Configuration

### Striping Options for EBS Volumes

With Amazon EBS, you can use any of the standard RAID configurations that you could use with a traditional bare metal server, as long as that particular RAID configuration is supported by the operating system for your instance. This is because all RAID is accomplished at the software level. For greater I/O performance than you can achieve with a single volume, RAID 0 can stripe multiple volumes together; for on-instance redundancy, RAID 1 can mirror two volumes together.

The following table compares these two common striping options.



Configuration	Use	Advantages	Disadvantages
RAID 0	When I/O performance is more important than fault tolerance as in, for example, a heavily used database (where data replication is already set up separately).	I/O is distributed across the volumes in a stripe. If you add a volume, you get the straight addition of throughput.	Performance of the stripe is limited to the worst performing volume in the set.
RAID 1	When fault tolerance is more important than I/O performance as in, for example, a critical application.	Safer from the standpoint of data durability.	Does not provide a write performance improvement; requires more bandwidth than non-RAID configurations because the data is written to multiple volumes simultaneously.

**Important**

RAID 5 and RAID 6 are not recommended for Amazon EBS because the configurations do not provide optimum performance.

**To Create a RAID Array on Linux**

Creating a RAID 0 array allows you to achieve a higher level of performance for a file system than you can provision on a single EBS volume. A RAID 1 array offers a "mirror" of your data for extra redundancy. Before you perform this procedure, you need to decide how large your RAID array needs to be and how many IOPS you want to provision.

The resulting size of a RAID 0 array will be the *sum* of the sizes of the volumes within it, and the bandwidth will be the *sum* of the available bandwidth of the volumes within it. The resulting size and bandwidth of a RAID 1 array is *equal* to the size and bandwidth of the volumes in the array. For example, two 500 GiB EBS volumes with 4,000 provisioned IOPS each will create a 1 TiB RAID 0 array with an available bandwidth of 8,000 16K IOPS or a 500 GiB RAID 1 array with an available bandwidth of 4,000 16K IOPS.

1. Create the EBS volumes for your array. For more information, see [Creating or Restoring an Amazon EBS Volume \(p. 449\)](#).

**Important**

Create volumes with identical size and IOPS performance values for your array. Make sure you do not create an array that exceeds the available bandwidth of your EC2 instance. For more information, see [Amazon EC2 Instance Configuration \(p. 495\)](#).

2. Attach the EBS volumes to the instance that you want to host the array. For more information, see [Attaching an Amazon EBS Volume to an Instance \(p. 455\)](#).
3. (Optional) Pre-warm your volumes. EBS volumes can experience between a 5 to 50 percent performance hit the first time you access a block on the volume. Pre-warming allows you to access each block on the volumes before you need to use them so you can achieve full performance in a production workload. For more information, see [Pre-Warming Amazon EBS Volumes \(p. 498\)](#).
4. Use the **mdadm** command to create a logical RAID device from the newly attached EBS volumes. Substitute the number of volumes in your array for *number\_of\_volumes* and the device names for each volume in the array (such as `/dev/xvdf`) for *device\_name*.

**Note**

You can list the devices on your instance with the **lsblk** command to find the device names. (RAID 0 only) To create a RAID 0 array, execute the following command (note the `--level=stripe` option to stripe the array):

```
[ec2-user ~]$ sudo mdadm --create --verbose /dev/md0 --level=stripe --raid-  
devices=number_of_volumes device_name1 device_name2
```

(RAID 1 only) To create a RAID 1 array, execute the following command (note the `--level=mirror` option to mirror the array):

```
[ec2-user ~]$ sudo mdadm --create --verbose /dev/md0 --level=mirror --raid-  
devices=number_of_volumes device_name1 device_name2
```

5. Create a file system on your RAID device. For example, to create an `ext4` file system, execute the following command:

```
[ec2-user ~]$ sudo mkfs.ext4 /dev/md0
```

6. Create a mount point for your RAID array.

```
[ec2-user ~]$ sudo mkdir /mnt/md0
```

7. Finally, mount the RAID device on the mount point you created:

```
[ec2-user ~]$ sudo mount -t ext4 /dev/md0 /mnt/md0
```

Your RAID device is now ready for use.

## Benchmark Volumes

This section demonstrates how you can test the performance of Provisioned IOPS volumes by simulating workloads similar to those of a database application. The process is as follows:

1. Launch an EBS-optimized instance
2. Create new Provisioned IOPS volumes
3. Attach the volumes to your EBS-optimized instance
4. Create a RAID array from the volumes, then format and mount it
5. Install a tool to benchmark I/O performance
6. Benchmark the I/O performance of your volumes
7. Delete your volumes and terminate your instance so that you don't continue to incur charges

## Set Up Your Instance

To get optimal performance from a Provisioned IOPS volume, we recommend that you use an EBS-optimized instance. EBS-optimized instances deliver dedicated throughput between Amazon EC2 and Amazon EBS, with options between 500 and 1,000 Mbps, depending on the instance type.

To create an EBS-optimized instance, select **Launch as an EBS-Optimized instance** when launching the instance using the EC2 console, or specify `--ebs-optimized` when using the command line. Be sure that you launch one of the instance types that supports this option. For the example tests in this topic, we recommend that you launch an `m1.xlarge` instance. For more information, see [EBS-Optimized Instances \(p. 107\)](#).

To create a Provisioned IOPS volume, select **Provisioned IOPS (io1)** when creating the volume using the EC2 console, or specify `--type io1 --iops/iops` when using the command line. For information about attaching these volumes to your instance, see [Attaching an Amazon EBS Volume to an Instance \(p. 455\)](#).

For the example tests, we recommend that you create a RAID array with 6 volumes, which offers a high level of performance. Because you are charged by the gigabytes used and the number of provisioned IOPS, not the number of volumes, there is no additional cost for creating multiple, smaller volumes and using them to create a stripe set. If you're using Oracle ORION to benchmark your volumes, it can simulate striping the same way that Oracle ASM does, so we recommend that you let ORION do the striping. If you are using a different benchmarking tool, you'll need to stripe the volumes yourself.

For information about creating a striped volume on Windows, see [Create a Striped Volume in Windows](#).

To create a six-volume stripe set on Linux, use a command like the following.

```
$ sudo mdadm --create /dev/md0 --level=0 --chunk=64 --raid-devices=6 /dev/sdf /dev/sdg /dev/sdh /dev/sdi /dev/sdj /dev/sdk
```

For this example, the file system we use is XFS. You should use the file system that meets your requirements. On Amazon Linux, use the following command to install the XFS file system.

```
$ sudo yum install -y xfsprogs
```

Then, use the following commands to create and mount the XFS file system.

```
$ sudo mkdir -p /media/p_iops_vol0 && sudo mkfs.xfs /dev/md0 && sudo mount -t xfs /dev/md0 /media/p_iops_vol0 && sudo chown ec2-user:ec2-user /media/p_iops_vol0/
```

Finally, to get the most accurate results while running these tests, you must pre-warm the volume. For a completely new volume that was created from scratch, you should write to all blocks before using the volume. For a new volume created from a snapshot, you should read all the blocks that have data before using the volume. For example, on Linux you can read each block on the volume using the following command.

```
$ dd if=/dev/md0 of=/dev/null
```

On Windows, a full format of the volume pre-warms it. Use the `format <drive letter> /p:1` command to write zeros to the entire disk.

### Important

Unless you pre-warm the volume, you might see between a 5 to 50 percent reduction in IOPS when you first access it.

## Install Benchmark Tools

The following are among the possible tools you can install and use to benchmark the performance of a Provisioned IOPS volume.

Tool	Platform	Description
<a href="#">fio</a>	Linux, Windows	For benchmarking I/O performance. (Note that fio has a dependency on libaio-devel.)

Tool	Platform	Description
Oracle Orion Calibration Tool	Linux, Windows	For calibrating the I/O performance of storage systems to be used with Oracle databases.
<a href="#">SQLIO</a>	Windows	For calibrating the I/O performance of storage systems to be used with Microsoft SQL Server.  For information about how to improve the performance of your Microsoft SQL Server databases, see <a href="#">Optimizing Databases</a> at the <i>MSDN</i> website.

## Example Benchmarking Commands

These benchmarking tools support a wide variety of test parameters. You should use commands that approximate the workloads your volumes will support. These commands are intended as examples to help you get started.

Run the following commands on an EBS-optimized instance with attached Provisioned IOPS volumes that have been pre-warmed.

When you are finished testing your volumes, see these topics for help cleaning up: [Deleting an Amazon EBS Volume \(p. 477\)](#) and [Terminate Your Instance \(p. 287\)](#).

### fiio Commands

Run **fiio** on the stripe set that you created.

By default, **fiio** is installed in `/usr/local/bin`, which isn't in root's path. You should update your path, or use **chmod** to enable **fiio** to be run as `ec2-user`.

The following command performs 16 KB random write operations.

```
fiio --directory=/media/p_iops_vol0
--name fio_test_file --direct=1 --rw=randwrite --bs=16k --size=1G
--numjobs=16 --time_based --runtime=180 --group_reporting
```

The following command performs 16 KB random read operations.

```
fiio --directory=/media/p_iops_vol0
--name fio_test_file --direct=1 --rw=randread --bs=16k --size=1G
--numjobs=16 --time_based --runtime=180 --group_reporting
```

For more information about interpreting the results, go to this Linux tutorial: [Inspecting disk IO performance with fiio](#).

### Oracle ORION Commands

Run ORION on the Provisioned IOPS volumes, having it simulate Oracle ASM striping instead of providing it with a stripe set that uses Windows striping.

In the directory where you installed ORION, create a file, `piops_test.lun`, to specify the volumes for your stripe set. The following example file specifies six Provisioned IOPS volumes to be striped.

```
\\.\D:  
\\.\E:  
\\.\F:  
\\.\G:  
\\.\H:  
\\.\I:
```

The following command performs 16 KB random I/O operations (80 percent reads and 20 percent writes), simulating 64 KB RAID-0 stripes.

```
orion -run advanced -testname piops_test -size_small 16 -size_large 16  
-type rand -simulate raid0 -stripe 64 -write 80 -matrix detailed -num_disks 6
```

After the command is finished, ORION generates output files with the results in the same directory. For more information about ORION, see its [Documentation](#).

### SQLIO Commands

Run SQLIO on the stripe set that you created.

Create a file, `param.txt`, to specify your striped set. The contents of this file should look something like this (here, `d:\` corresponds to the striped set, and the test uses 6 threads and a 10 GB file).

```
d:\bigtestfile.dat 6 0x0 10240
```

The following command performs 16 KB random data writes.

```
sqlio -kW -s600 -frandom -t8 -o8 -bl6 -LS -BH -Fparam.txt
```

The following command performs 16 KB random data reads.

```
sqlio -kR -s600 -frandom -t8 -o8 -bl6 -LS -BH -Fparam.txt
```

The results are displayed in the Command Prompt window. For more information about SQLIO, see the `readme.txt` file in your SQLIO installation directory.

## Amazon EBS API and Command Overview

The following table summarizes the available Amazon EBS commands and corresponding API actions for creating and using Amazon EBS volumes.

Command and API Action	Description
<a href="#">ec2-attach-volume</a> <a href="#">AttachVolume</a>	Attaches the specified volume to a specified instance, exposing the volume using the specified device name.  A volume can be attached to only a single instance at any time. The volume and instance must be in the same Availability Zone. The instance must be in the <code>running</code> or <code>stopped</code> state.

Command and API Action	Description
<a href="#">ec2-copy-snapshot</a> <a href="#">CopySnapshot</a>	<p>Copies a point-in-time snapshot of an Amazon Elastic Block Store (Amazon EBS) volume and stores it in Amazon Simple Storage Service (Amazon S3). You can copy the snapshot within the same region or from one region to another. You can use the snapshot to create new Amazon EBS volumes or Amazon Machine Images (AMIs).</p>
<a href="#">ec2-create-snapshot</a> <a href="#">CreateSnapshot</a>	<p>Creates a snapshot of the volume you specify.</p> <p>After the snapshot is created, you can use it to create volumes that contain exactly the same data as the original volume.</p>
<a href="#">ec2-create-volume</a> <a href="#">CreateVolume</a>	<p>Creates a new Amazon EBS volume using the specified size and type, or based on a previously created snapshot.</p>
<a href="#">ec2-delete-disk-image</a>	<p>Deletes a partially or fully uploaded disk image for conversion from Amazon S3.</p>
<a href="#">ec2-delete-snapshot</a> <a href="#">DeleteSnapshot</a>	<p>Deletes the specified snapshot.</p> <p>This command does not affect currently running Amazon EBS volumes, regardless of whether they were used to create the snapshot or were derived from the snapshot.</p>
<a href="#">ec2-delete-volume</a> <a href="#">DeleteVolume</a>	<p>Deletes the specified volume. The command does not delete any snapshots that were created from the volume.</p>
<a href="#">ec2-describe-snapshot-attribute</a> <a href="#">DescribeSnapshotAttribute</a>	<p>Describes attributes for a snapshot.</p>
<a href="#">ec2-describe-snapshots</a> <a href="#">DescribeSnapshots</a>	<p>Describes the specified snapshot.</p> <p>Describes all snapshots, including their source volume, snapshot initiation time, progress (percentage complete), and status (<i>pending</i>, <i>completed</i>, and so on.).</p>
<a href="#">ec2-describe-volumes</a> <a href="#">DescribeVolumes</a>	<p>Describes your volumes, including size, volume type, source snapshot, Availability Zone, creation time, status (<i>available</i> or <i>in-use</i>). If the volume is <i>in-use</i>, an attachment line shows the volume ID, the instance ID to which the volume is attached, the device name exposed to the instance, its status (<i>attaching</i>, <i>attached</i>, <i>detaching</i>, <i>detached</i>), and when it attached.</p>
<a href="#">ec2-describe-volume-attribute</a> <a href="#">DescribeVolumeAttribute</a>	<p>Describes an attribute of a volume.</p>
<a href="#">ec2-describe-volume-status</a> <a href="#">DescribeVolumeStatus</a>	<p>Describes the status of one or more volumes. Volume status provides the result of the checks performed on your volumes to determine events that can impair the performance of your volumes.</p>
<a href="#">ec2-detach-volume</a> <a href="#">DetachVolume</a>	<p>Detaches the specified volume from the instance it's attached to.</p> <p>This command does not delete the volume. The volume can be attached to another instance and will have the same data as when it was detached.</p>

Command and API Action	Description
<a href="#">ec2-enable-volume-io</a> EnableVolumeIO	Enables I/O operations for a volume that had I/O operations disabled because the data on the volume was potentially inconsistent.
<a href="#">ec2-import-volume</a> ImportVolume	Creates a new import volume task using metadata from the specified disk image, and imports the image to Amazon EC2.
<a href="#">ec2-modify-snapshot-attribute</a> ModifySnapshotAttribute	Modifies permissions for a snapshot (i.e., who can create volumes from the snapshot). You can specify one or more AWS accounts, or specify <code>all</code> to make the snapshot public.
<a href="#">ec2-modify-volume-attribute</a> ModifyVolumeAttribute	Modifies a volume's attributes to determine whether a volume should be automatically enabled for I/O operations.
<a href="#">ec2-reset-snapshot-attribute</a> ResetSnapshotAttribute	Resets permission settings for the specified snapshot.

## Amazon EC2 Instance Store

Each Amazon EC2 instance, unless it's a micro or M3 instance, can access storage from disks that are physically attached to the host computer. This disk storage is referred to as *instance store*.

### Topics

- [Instance Storage Concepts](#) (p. 508)
- [Instance Stores Available on Instance Types](#) (p. 509)
- [Instance Store Swap Volumes](#) (p. 510)
- [Instance Store Device Names](#) (p. 511)
- [Instance Store Usage Scenarios](#) (p. 512)
- [Adding Instance Store Volumes to an AMI](#) (p. 515)
- [Optimizing Disk Performance](#) (p. 516)

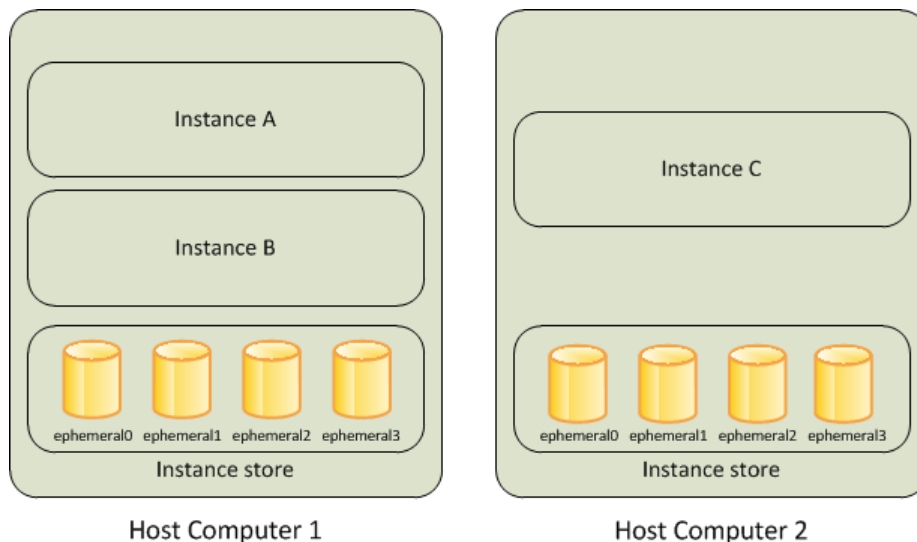
## Instance Storage Concepts

An Amazon EC2 instance store provides temporary block-level storage for use with an Amazon EC2 instance. The size of an instance store ranges from 150 GiB to up to 48 TiB, and varies by instance type. Larger instance types have larger instance stores. For more information, see [Instance Stores Available on Instance Types](#) (p. 509).

An instance store consists of one or more instance store volumes. Instance store volumes must be configured via block device mapping at launch time and mounted on the running instance before they can be used. By default, instances launched from an Amazon EBS-backed instance have no mounted instance store volumes. Instances launched from an instance store-backed AMI have a mounted instance store volume for the virtual machine's root device volume, and can have other mounted instance store volumes, depending on the instance type. For more information about instance store-backed AMIs and Amazon EBS-backed AMIs, see [Storage for the Root Device](#) (p. 45).

Instance store volumes are usable only from a single instance during its lifetime; they can't be detached and then attached to another instance. If you create an AMI from an instance, the data on its instance store volumes isn't preserved and isn't present on the instance store volumes for the instances that you

launch from this AMI. While an instance store is dedicated to a particular instance, the disk subsystem is shared among instances on a host computer, as shown in the following figure.



The data in an instance store persists only during the lifetime of its associated instance. If an instance reboots (intentionally or unintentionally), data in the instance store persists. However, data on instance store volumes is lost under the following circumstances:

- Failure of an underlying drive
- Stopping an Amazon EBS-backed instance
- Terminating an instance

Therefore, do not rely on instance store volumes for valuable, long-term data. Instead, keep your data safe by using a replication strategy across multiple instances, storing data in Amazon S3, or using Amazon EBS volumes. For more information, see [Amazon Elastic Block Store \(Amazon EBS\)](#) (p. 446).

When you launch an instance, whether it's launched from an Amazon EBS-backed AMI or an instance store-backed AMI, you can attach instance store volumes to the instance using block device mapping. For more information, see [Adding Instance Store Volumes to an AMI](#) (p. 515).

## Instance Stores Available on Instance Types

Amazon EC2 instances are divided into different instance types, which determine the size of the instance store available on the instance by default. When you launch an instance, you can specify an instance type or use the default instance type, which is an `m1.small` instance.

The instance type also determines the type of hardware for your instance store volumes. A high I/O instance (`hi1.4xlarge`) uses solid state drives (SSD) to deliver very high random I/O performance. This is a good option when you need storage with very low latency, but you don't need it to persist when the instance terminates, or you can take advantage of fault tolerant architectures. For more information see [HI1 Instances](#) (p. 102).

The following table shows the instance types along with the size and quantity of the instance store volumes available to each instance type; these instance store volumes are included as part of the instance's hourly cost.



Instance Type	Instance Store Volumes
t1.micro	None (use Amazon EBS volumes)
m1.small	1 x 160 GB†
m1.medium	1 x 410 GB
m1.large	2 x 420 GB (840 GB)
m1.xlarge	4 x 420 GB (1680 GB)
m3.xlarge	None (use Amazon EBS volumes)
m3.2xlarge	None (use Amazon EBS volumes)
c1.medium	1 x 350 GB†
c1.xlarge	4 x 420 GB (1680 GB)
m2.xlarge	1 x 420 GB
m2.2xlarge	1 x 850 GB
m2.4xlarge	2 x 840 GB (1680 GB)
hi1.4xlarge	2 x 1024 GB SSD (2048 GB)
hs1.8xlarge	24 x 2048 GB (49 TB)
cr1.8xlarge	2 x 120 GB SSD (240 GB)
cc2.8xlarge	4 x 840 GB (3360 GB)
cg1.4xlarge	2 x 840 GB (1680 GB)

† The `c1.medium` and `m1.small` instance types also include a 900 MB instance store swap volume, which may not be automatically enabled at boot time. For more information, see [Instance Store Swap Volumes \(p. 510\)](#).

## Instance Store Swap Volumes

Swap space in Linux can be used when a system requires more memory than it has been physically allocated. When swap space is enabled, Linux systems can swap infrequently used memory pages from physical memory to swap space (either a dedicated partition or a swap file in an existing file system) and free up that space for memory pages that require high speed access.

The `c1.medium` and `m1.small` instance types have a limited amount of physical memory to work with, and they are given a 900 MB swap volume at launch time to act as virtual memory for Linux AMIs. Although the Linux kernel sees this swap space as a partition on the root device, it is actually a separate instance store volume, regardless of your root device type.

Amazon Linux AMIs automatically enable and use this swap space, but your AMI may require some additional steps to recognize and use this swap space. To see if your instance is using swap space, you can use the **swapon -s** command.

```
[ec2-user@ip-12-34-56-78 ~]$ swapon -s
Filename                                Type              Size    Used    Priority
/dev/xvda3                             partition        917500  0       -1
```

The above instance has a 900 MB swap volume attached and enabled. If you don't see a swap volume listed with this command, you may need to enable swap space for the device. Check your available disks using the **lsblk** command.

```
[ec2-user@ip-12-34-56-78 ~]$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda1 202:1 0 8G 0 disk /
xvda3 202:3 0 896M 0 disk
```

Here, the swap volume `xvda3` is available to the instance, but it is not enabled (notice that the `MOUNTPOINT` field is empty). You can enable the swap volume with the **swapon** command.

**Note**

You need to prepend `/dev/` to the device name listed by **lsblk**. Your device may be named differently, such as `sda3`, `sde3`, or `xvde3`. Use the device name for your system in the command below.

```
[ec2-user@ip-12-34-56-78 ~]$ sudo swapon /dev/xvda3
```

Now the swap space should show up in **lsblk** and **swapon -s** output.

```
[ec2-user@ip-12-34-56-78 ~]$ lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
xvda1 202:1 0 8G 0 disk /
xvda3 202:3 0 896M 0 disk [SWAP]
[ec2-user@ip-12-34-56-78 ~]$ swapon -s
Filename                                Type              Size      Used      Priority
/dev/xvda3                              partition         917500    0         -1
```

You will also need to edit your `/etc/fstab` file so that this swap space is automatically enabled at every system boot.

```
[ec2-user@ip-12-34-56-78 ~]$ sudo vim /etc/fstab
```

Append the following line to your `/etc/fstab` file (using the swap device name for your system):

```
/dev/xvda3 none swap sw 0 0
```

## Instance Store Device Names

Within an instance store, instance store volumes are exposed as block devices. The virtual devices for instance store volumes are `ephemeral[0-3]`. Instance types that support one instance store volume have `ephemeral0`. Instance types that support two instance store volumes have `ephemeral0` and `ephemeral1`. Instance types that support four instance store volumes have `ephemeral0`, `ephemeral1`, `ephemeral2`, and `ephemeral3`. Each instance store volume is pre-formatted with the `ext3` file system. However, you can reformat volumes with the file system of your choice after you launch your instance. A Windows instance uses a built-in tool, EC2Config Service, to reformat the instance store volumes available on an instance with the NTFS file system.

Every instance store-backed AMI and instance has a mapping of the instance store volumes attached to the instance. Each entry in the mapping consists of a device name and the volume that it's mapped to. The instance store volumes are available to the instance, but you can't access them until they are mounted. A Windows instance uses the EC2Config Service to mount the instance store volumes for an instance. On Linux, the instance type determines which instance store volumes are mounted for you and which

are available for you to mount yourself. The block device driver for the instance assigns the actual volume name when mounting the volume, and the name assigned can be different than the name that Amazon EC2 recommends.

The following table lists the devices names reserved for instance store volumes by instance type, and the default state of the storage device (formatted, mounted, available) on an instance store-backed instance. For example, an `m1.small` instance has `ephemeral0` (ext3, 15 GiB) and `swap` (Linux swap, 896 MB).

Device Name	Linux/UNIX instance store-backed instance	Windows instance store-backed instance
<code>/dev/sda1</code> , <code>/dev/xvda1</code> , <code>/dev/xvde1</code>	Formatted and mounted as root ( <code>/</code> ).	Formatted and mounted as <code>C:\</code> .
<code>/dev/sda2</code> , <code>xvdb</code>	Formatted and mounted as <code>/mnt</code> or <code>/media/ephemeral0</code> on <code>m1.small</code> and <code>c1.medium</code> instances.	Formatted and mounted on small instance types.
<code>/dev/sda3</code> , <code>/dev/xvda3</code> , <code>/dev/xvde3</code>	Available as swap space on <code>m1.small</code> and <code>c1.medium</code> instances. For more information, see <a href="#">Instance Store Swap Volumes (p. 510)</a>	Not available.
<code>/dev/sdb</code> , <code>xvdb</code>	Formatted and mounted as <code>/mnt</code> or <code>/media/ephemeral0</code> on <code>m1.medium</code> , <code>m1.large</code> , <code>m1.xlarge</code> , <code>c1.xlarge</code> , <code>cc2.8xlarge</code> , <code>cr1.8xlarge</code> , <code>m2.xlarge</code> , <code>m2.2xlarge</code> , <code>m2.4xlarge</code> , and <code>hi1.4xlarge</code> .	Formatted and mounted on <code>m1.medium</code> , <code>m1.large</code> , <code>m1.xlarge</code> , <code>c1.xlarge</code> , <code>m2.xlarge</code> , and <code>m2.2xlarge</code> .
<code>/dev/sdc</code> , <code>xvdc</code>	Available on <code>m1.large</code> , <code>m1.xlarge</code> , <code>cc2.8xlarge</code> , <code>c1.xlarge</code> , <code>cr1.8xlarge</code> , and <code>hi1.4xlarge</code> .	Formatted and mounted on <code>m1.large</code> , <code>m1.xlarge</code> , <code>c1.xlarge</code> , and <code>m2.4xlarge</code> .
<code>/dev/sdd</code> , <code>xvdd</code>	Available on <code>m1.xlarge</code> , <code>c1.xlarge</code> , and <code>cc2.8xlarge</code> .	Formatted and mounted on <code>m1.xlarge</code> and <code>c1.xlarge</code> .
<code>/dev/sde</code> , <code>xvde</code>	Available on <code>m1.xlarge</code> and <code>c1.xlarge</code> .	Formatted and mounted on <code>m1.xlarge</code> and <code>c1.xlarge</code> .

An instance can have multiple instance store volumes mapped to a device. However, the number and size of these volumes must not exceed the instance store available for the instance type. For more information, see [Instance Stores Available on Instance Types \(p. 509\)](#).

## Instance Store Usage Scenarios

Instance store volumes are ideal for temporary storage of information that changes frequently, such as buffers, caches, scratch data, and other temporary content, or for data that is replicated across a fleet of instances, such as a load-balanced pool of web servers.

### Making Instance Stores Available on Your Instances

Instances that use Amazon EBS for the root device do not, by default, have instance store available at boot time. Also, you can't attach instance store volumes after you've launched an instance. Therefore, if you want your Amazon EBS-backed instance to use instance store volumes, you must specify them using

a block device mapping when you create your AMI or launch your instance. Examples of block device mapping entries are: `/dev/sdb=ephemeral0` and `/dev/sdc=ephemeral1`. For more information about block device mapping, see [Block Device Mapping \(p. 517\)](#)

The following procedure describes how to launch an Amazon EBS-backed `m1.large` Windows instance with instance store volumes.

### Launch Amazon EBS-backed Windows Instances with Instance Store Volumes

1. Locate an Amazon EBS-backed Windows AMI.
2. Add block device mapping entries.

Console: For more information, see [To add volumes to an instance \(p. 523\)](#).

CLI: Use the `ec2-run-instances` command to launch an `m1.large` instance; specify volumes for the block device mapping using the following options: `-b "/dev/xvdb=ephemeral0"` and `-b "/dev/xvdc=ephemeral1"`.

3. Connect to the instance.
4. On the **Start** menu, choose **Computer**.
5. Devices listed:
  - Local Disk C:/ 9.98GiB
  - Local Disk D:/ 419GiB
  - Local Disk E:/ 419GiB
6. Double-click **Local Disk C:/**. You can see the list of installed applications. This is your root drive.
7. Double-click **Local Disk D:/** and then double-click **Local Disk E:/**. These drives are empty. They are the instance stores that come with your `m1.large` instance, and they are available for you to use with your applications.

The following procedure describes how to launch an Amazon EBS-backed `m1.large` Linux instance with instance store volumes.

### Accessing Instance Stores on Amazon EBS-backed Linux Instances

1. Locate an Amazon EBS-backed Linux/UNIX AMI.
2. Add block device mapping entries.

Console: For more information, see [To add volumes to an instance \(p. 523\)](#).

CLI: Use the `ec2-run-instances` command to launch an `m1.large` instance; specify volumes for the block device mapping using the following options: `-b "/dev/xvdb=ephemeral0"` and `-b "/dev/xvdc=ephemeral1"`.

3. Connect to the instance.
4. Verify the instance stores currently mounted on the disk.
5. Notice a 10GiB root partition mounted at `/` and a 420 GiB mounted on an `/media/ephemeral0` volume. Your `m1.large` instance comes with 2 420 GiB instance store volumes; the second volume is available but must be mounted before it can be used.
6. To format and mount the other 420 GiB instance store volume:
  - a. Create a file system of your choice on the device `/dev/sdc` (requires root privileges).
  - b. Create a directory on which to mount the device.
  - c. Mount the device on the newly created directory.

7. Verify that the device has been mounted.
8. Optionally, list the files on the root device.

You can also map instance store volumes to block devices when you create an AMI. The instances launched from such an AMI have instance store volumes at boot time. For information about adding a block device mapping while creating an AMI, see [Creating Your Own AMIs \(p. 62\)](#).

The following procedure describes how to access the instance store volumes from within an Amazon EC2 instance store-backed `m1.large` Windows instance.

### Tasks for Accessing Instance Stores on Amazon EC2 instance store-backed Windows Instances

1	Locate an Amazon EC2 instance store-backed Windows AMI.
2	Launch an <code>m1.large</code> instance.
3	Connect to the instance.
4	On the <b>Start</b> menu, choose <b>Computer</b> .
5	Devices listed: <ul style="list-style-type: none"><li>• Local Disk C:/ 9.98GiB</li><li>• Local Disk D:/ 419GiB</li><li>• Local Disk E:/ 419GiB</li></ul>
6	Double-click Local Disk C:/. You see the list of all installed applications. This is your root drive.
7	Double-click Local Disk D:/ and then double-click Local Disk E:/. These are empty. They are the instance store volumes that come with your <code>m1.large</code> instance, and they are available to use with your applications just like any physical drive.

Depending on the instance type, some instance store volumes on Amazon EC2 instance store-backed Linux and UNIX instances are not mounted when the instance is launched. For example, on an `m1.large` Linux and UNIX instance, the device `/dev/sdc`, although formatted and available, must be mounted before it can be used.

The following procedure describes how to access the instance store from within Amazon EC2 instance store-backed `m1.large` Linux instance.

### Accessing Instance Stores on Amazon EC2 instance store-backed Linux Instances

1	Locate an Amazon EC2 instance store-backed Linux/Unix AMI.
2	Launch an <code>m1.large</code> instance.
3	Connect to the instance.
4	Check out the file systems currently mounted on the disk.
5	Notice 10 GiB root partition mounted on the root and 420 GiB mounted on an <code>ephemeral0</code> device. Your <code>m1.large</code> instance comes with 2 420 GiB instance store volumes; the second volume is available but must be mounted before it can be used.

6	To mount the other 420GiB: a. Create a directory on which to mount the device. b. Mount the device on the newly created directory.
7	Check to see whether the device has been mounted.
8	Optionally, list the files on the root device.

## Suppressing Instance Stores at Launch Time

You can prevent a particular instance storage volume from attaching to the instance. You can do this for both Amazon EC2 instance store-backed instances and Amazon EBS-backed instances. For example, specifying the mapping `/dev/sdc=none` when launching an instance prevents `/dev/sdc` from attaching to the instance. For more information about block device mapping, see [Block Device Mapping \(p. 517\)](#).

## Adding Instance Store Volumes to an AMI

Amazon EBS-backed AMIs don't include an instance store by default. However, you might want instances launched from your Amazon EBS-backed AMIs to include instance store volumes. This section describes how to create an AMI that includes instance store volumes.

### AWS Management Console

#### To add instance store volumes to an AMI

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, click **Instances**.
3. Select an instance and select **Create Image** from the **Actions** list.
4. In the **Create Image** dialog, add a meaningful name and description to your image.
5. Click **Add New Volume**.
6. For each instance store volume, select a volume from the **Type** list and a device name from **Device**.
7. Click **Create Image**.

### Command Line Interface

Use the `ec2-register` command to specify a block device mapping that includes the instance store volumes for the image. For more information about block device mapping, see [Block Device Mapping \(p. 517\)](#).

#### To add instance store volumes to an AMI

1. Use the `ec2-register` command with the desired block device mapping information.

The following example registers an AMI with an 80GiB root device volume at `/dev/sda1` created from the `snap-12345678` snapshot. The root volume's `DeleteOnTermination` flag is set to `false`. The second block device mapping in the request maps `/dev/sdc` to `ephemeral0`.

You can omit the size in the block device mapping if you want to create a volume that is the size of the snapshot. If you do specify a size, it must be equal to or larger than the size of the snapshot. You can also resize the partition or create a new partition later.

### Tip

If you're using the command line interface on a Windows computer, you must put quotation marks around any part of the command that includes an equal sign. For example: `... -b "/dev/sdc=ephemeral0" ...`

```
PROMPT> ec2-register -n My_Image_Name -d My_image_description --root-device-  
name /dev/sda1 -b /dev/sda1=snap-12345678:80:false -b /dev/sdc=ephemeral0
```

In response, you get the ID for your new AMI.

```
IMAGE      ami-61a54008
```

2. Any instance you launch from this AMI includes the instance store volumes that you specified when you created the AMI. You can confirm that the instance store devices are available from within the instance itself using instance metadata. For more information, see [Viewing the Instance Block Device Mapping for Instance Store Volumes \(p. 526\)](#).

## Optimizing Disk Performance

Because of the way that Amazon EC2 virtualizes disks, the first write to any location on a standard instance store volume performs more slowly than subsequent writes. (The performance of the solid state drives (SSD) used by high I/O instances is not affected this way.) For most applications, amortizing this cost over the lifetime of the instance is acceptable. However, if you require high disk performance, we recommend that you initialize your drives by writing once to every drive location before production use. If you require further improvements in latency or throughput, we recommend using Amazon EBS.

To initialize the stores, use the following Linux/UNIX `dd` commands, depending on which store you want to initialize (`/dev/sdb`, etc.).

### Note

Make sure to unmount the drive before performing this command.  
Initialization can take a long time (about 8 hours for an extra large instance).

To initialize the instance store volumes, use the following commands on the `m1.large`, `m1.xlarge`, `c1.xlarge`, `m2.xlarge`, `m2.2xlarge`, and `m2.4xlarge` instance types:

```
dd if=/dev/zero of=/dev/sdb bs=1M  
dd if=/dev/zero of=/dev/sdc bs=1M  
dd if=/dev/zero of=/dev/sdd bs=1M  
dd if=/dev/zero of=/dev/sde bs=1M
```

For information about the instance storage that is available for each instance type, see [Instance Stores Available on Instance Types \(p. 509\)](#).

To perform initialization on all instance store volumes at the same time, use the following command:

```
dd if=/dev/zero bs=1M|tee /dev/sdb|tee /dev/sdc|tee /dev/sde > /dev/sdd
```

Configuring drives for RAID initializes them by writing to every drive location. When configuring software-based RAID, make sure to change the minimum reconstruction speed:

```
echo $((30*1024)) > /proc/sys/dev/raid/speed_limit_min
```

# Amazon Simple Storage Service (Amazon S3)

Amazon S3 is a repository for Internet data. Amazon S3 provides access to reliable, fast, and inexpensive data storage infrastructure. It is designed to make web-scale computing easy by enabling you to store and retrieve any amount of data, at any time, from within Amazon EC2 or anywhere on the web. Amazon S3 stores data objects redundantly on multiple devices across multiple facilities and allows concurrent read or write access to these data objects by many separate clients or application threads. You can use the redundant data stored in Amazon S3 to recover quickly and reliably from instance or application failures.

Objects are the fundamental entities stored in Amazon S3. Objects consist of object data and metadata. The metadata is a set of name-value pairs that describes the object. The data portion is opaque to Amazon S3.

Every object stored in Amazon S3 is contained in a bucket. Buckets organize the Amazon S3 namespace at the highest level and identify the account responsible for that storage. Amazon S3 buckets are similar to Internet domain names. Objects stored in the buckets have a unique key value and are retrieved using a HTTP URL address. For example, if an object with a key value `/photos/mygarden.jpg` is stored in the `myawsbucket` bucket, then it is addressable using the URL `http://myawsbucket.s3.amazonaws.com/photos/mygarden.jpg`.

For more information about the features of Amazon S3, see the [Amazon S3 product page](#).

## Amazon S3 and Amazon EC2

The combination of Amazon S3 and Amazon EC2 provides resizable computing capacity with highly scalable and fast data storage infrastructure.

Amazon EC2 uses Amazon S3 for storing Amazon Machine Images (AMIs). You use AMIs for launching EC2 instances. In case of instance failure, you can use the stored AMI to immediately launch another instance, thereby allowing for fast recovery and business continuity.

Amazon EC2 also uses Amazon S3 to store snapshots (backup copies) of the data volumes. You can use snapshots for recovering data quickly and reliably in case of application or system failures. You can also use snapshots as a baseline to create multiple new data volumes, expand the size of an existing data volume, or move data volumes across multiple Availability Zones, thereby making your data usage highly scalable. For more information about using data volumes and snapshots, see [Amazon Elastic Block Store \(p. 446\)](#).

Amazon S3 stores AMIs and snapshots redundantly on multiple devices across multiple facilities, thus providing an instance with highly durable storage infrastructure.

Amazon EC2 doesn't have a built-in capability to work with Amazon S3. If you have permission, you can access a bucket using its URL. If you are a developer, you can use the API to access data in Amazon S3. You can also use the [AWS Tools for Windows PowerShell](#) to access Amazon S3 objects. Otherwise, there are a variety of tools that other people have written that you can use to access your data in Amazon S3; some of the common ones are discussed in the AWS forums. For more information about using Amazon S3, see the [Amazon Simple Storage Service Developer Guide](#).

## Block Device Mapping

Each Amazon EC2 instance that you launch has an associated root device volume, either an Amazon Elastic Block Store (EBS) volume or an instance store volume. You can use block device mapping to specify additional EBS volumes or instance store volumes to attach to an instance when it's launched. You can also attach additional EBS volumes to a running instance; see [Attaching an Amazon EBS Volume](#)



to an Instance (p. 455). However, the only way to attach instance store volumes to an instance is to use block device mapping to attach them as the instance is launched.

For more information about root device volumes, see [Changing the Root Device Volume to Persist \(p. 14\)](#).

#### Topics

- [Block Device Mapping Concepts \(p. 518\)](#)
- [AMI Block Device Mapping \(p. 521\)](#)
- [Instance Block Device Mapping \(p. 523\)](#)

## Block Device Mapping Concepts

A *block device* is a storage device that moves data in sequences of bytes or bits (blocks). These devices support random access and generally use buffered I/O. Examples include hard disks, CD-ROM drives, and flash drives. A block device can be physically attached to a computer or accessed remotely as if it were physically attached to the computer. Amazon EC2 supports two types of block devices:

- instance store volumes (virtual devices whose underlying hardware is physically attached to the host computer for the instance)
- Amazon EBS volumes (remote storage devices)

A *block device mapping* defines the block devices to be attached to an Amazon EC2 instance. You can specify a block device mapping as part of creating an AMI so that the mapping is used by all instances launched from the AMI. Alternatively, you can specify a block device mapping when you launch an instance, so this mapping overrides the one specified in the AMI from which you launched the instance.

## Specifying a Block Device Mapping

Use a block device mapping to attach instance store volumes and EBS volumes to an Amazon EC2 instance.

When you create a block device mapping, you specify this information for each block device that you need to attach to the instance:

- [Linux/UNIX] The device name within Amazon EC2, as shown in this table. The block device driver for the instance assigns the actual volume name when mounting the volume, and the name assigned can be different than the name that Amazon EC2 recommends.

<b>Reserved for the root device</b>	/dev/sda1
<b>Recommended for instance store volumes</b>	/dev/sd[b-e]
<b>Recommended for EBS volumes</b>	/dev/sd[f-p] /dev/sd[f-p][1-6]
<b>Possible for EBS volumes</b>	/dev/sd[a-z] /dev/sd[a-z][1-15] (see Tip) /dev/hd[a-z] /dev/hd[a-z][1-15]

**Tip**

On some devices, the names that you specify in your block device mapping can conflict with the default block device names. To avoid this issue, do not use names of the form `/dev/sda[2-15]`.


- [Windows] The device name within Amazon EC2, as shown in this table. The block device driver for the instance assigns the actual volume name when mounting the volume, and the name assigned can be different than the name that Amazon EC2 recommends.

<b>Reserved for the root device</b>	<code>/dev/sda1</code>
<b>Recommended for instance store volumes</b>	<code>xvd[a-e]</code>
<b>Recommended for EBS volumes</b>	<code>xvd[f-p]</code>
<b>Possible for EBS volumes</b>	<code>xvd[a-p]</code> <code>/dev/sda[1-2]</code> <code>/dev/sd[b-e]</code>

- [Instance store volumes only] The virtual device: `ephemeral[0-3]`.
- [EBS volumes only] The ID of the snapshot to use to create the block device (`snap-xxxxxxx`). This value is optional as long as you specify a volume size.
- [EBS volumes only] The size of the volume, in GiB. The specified size must be greater than or equal to the size of the specified snapshot.
- [EBS volumes only] Whether to delete the volume on instance termination (`true` or `false`). The default value is `true`.
- [EBS volumes only] The volume type (`standard` or `io1`). The default value is `standard`.
- [EBS volumes only] The number of input/output operations per second (IOPS) that the volume supports. (Not used with `standard` volumes.)

## Example Block Device Mapping

This figure shows an example block device mapping for an Amazon EBS-backed instance. It maps `/dev/sdb` to `ephemeral0` and maps two EBS volumes, one to `/dev/sdh` and the other to `/dev/sdj`. It also shows the EBS volume that is the root device volume, `/dev/sda1`.



Note that this example block device mapping is used in the example commands and APIs in this topic. You can find example commands and APIs that create block device mappings here:

- [Specifying a Block Device Mapping for an AMI \(p. 521\)](#)
- [Updating the Block Device Mapping when Launching an Instance \(p. 523\)](#)

## How Devices are Made Available in the Operating System

Device names like `/dev/sdh` and `xvdh` are used by Amazon EC2 to describe block devices. The block device mapping is used by Amazon EC2 to specify the block devices to attach to an Amazon EC2 instance. After a block device is attached to an instance, it must be mounted by the operating system before you can access the storage device. When a block device is detached from an instance, it is unmounted by the operating system and you can no longer access the storage device.

With a Linux instance, the device names specified in the block device mapping are mapped to their corresponding block devices when the instance first boots. The instance type determines which instance store volumes are formatted and mounted by default. You can mount additional instance store volumes at launch, as long as you don't exceed the number of instance store volumes available for your instance type. For more information, see [Amazon EC2 Instance Store \(p. 508\)](#). The block device driver for the instance determines which devices are used when the volumes are formatted and mounted. For more information, see [Attaching an Amazon EBS Volume to an Instance \(p. 455\)](#).

With a Windows instance, the device names specified in the block device mapping are mapped to their corresponding block devices when the instance first boots, and then the Ec2Config service initializes and mounts the drives. The root device volume is mounted as `C:\`. The instance store volumes are mounted as `D:\`, `E:\`, and so on. When an EBS volume is mounted, it can be mounted using any available drive letter. However, you can configure how the Ec2Config Service assigns drive letters to EBS volumes; for more information, see [Using EC2Config](#).

## Viewing Block Device Mappings

You can view information about each block device in a block device mapping. For details, see:

- [Viewing the EBS Volumes in an AMI Block Device Mapping \(p. 522\)](#)
- [Viewing the EBS Volumes in an Instance Block Device Mapping \(p. 525\)](#)
- [Viewing the Instance Block Device Mapping for Instance Store Volumes \(p. 526\)](#)

## AMI Block Device Mapping

Each AMI has a block device mapping that specifies the block devices to attach to an instance when it is launched from the AMI. An AMI that Amazon provides includes a root device only. To add additional block devices to an AMI, you must create your own AMI.

### Specifying a Block Device Mapping for an AMI

There are two ways to specify volumes in addition to the root volume when you create an AMI. If you've already attached volumes to a running instance before you create an AMI from the instance, the block device mapping for the AMI includes those same volumes. For EBS volumes, the existing data is saved to a new snapshot and it's this new snapshot that's specified in the block device mapping. For instance store volumes, the data is not preserved.

For an EBS-backed AMI, you can add EBS volumes and instance store volumes using a block device mapping. For an instance store-backed AMI, you can add only instance store volumes using a block device mapping.

### AWS Management Console

#### To add volumes to an AMI

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, click **Instances**.
3. Select an instance and select **Create Image** from the **Actions** list.
4. In the **Create Image** dialog box, click **Add New Volume**.
5. [Optional] Select a volume from the **Type** list and a device name from **Device**.
6. Click **Create Image**.

### Command Line Interface

Use **ec2-register -b "devicename=blockdevice"** to create an AMI with a block device mapping.

#### **devicename**

The device name within Amazon EC2

#### **blockdevice**

To omit a mapping for the device from the AMI, specify `none`.

To add an instance store device, specify `ephemeral[0-3]`.

[EBS-backed instance only] To add an EBS volume, specify `snapshot-id:size:[true|false]`. To add an empty EBS volume, omit the snapshot ID. To indicate whether the EBS volume should be deleted on termination, specify `true` or `false`; the default value is `true`.

For example, run this command at a command prompt to register an Amazon EBS-backed Linux AMI with an EBS root volume, an instance store volume, an EBS volume based on a snapshot, and an empty 100 GiB EBS volume. (Be sure to set the `EC2_PRIVATE_KEY` and `EC2_CERT` environment variables first.)

```
ec2-register -n ImageName --root-device-name /dev/sda1 -s snap-e1eb279f  
-b "/dev/sdb=ephemeral0" -b "/dev/sdh=snap-d5eb27ab" -b "/dev/sdj=:100"
```

This command maps `/dev/sda1` to an EBS root volume based on a snapshot, `/dev/sdb` to `ephemeral0`, `/dev/sdh` to an EBS volume based on a snapshot, and `/dev/sdj` to an empty EBS volume that is 100 GiB in size. The output is the ID for your new AMI.

```
IMAGE ami-72aa081b
```

For more information, see [ec2-register](#).

To verify that the AMI was created successfully, look for it in the Amazon EC2 Console (click **AMIs** in the navigation pane, then select **Owned by me** from the **Filter** drop-down list) or the output of the following command:

```
ec2-describe-images -o self
```

Alternatively, if you've already attached volumes to an instance, you can use one of these methods as appropriate.

- Amazon EBS-backed AMI
  - [ec2-create-image](#)
- Instance store-backed AMI
  - [Creating an Instance Store-Backed AMI From an Existing AMI \(p. 68\)](#)
  - [Creating an Instance Store-Backed Windows AMI](#)

## Viewing the EBS Volumes in an AMI Block Device Mapping

You can easily enumerate the EBS volumes in the block device mapping for an AMI.

### AWS Management Console

Use the AWS Management Console as follows to enumerate the EBS volumes in the block device mapping for an AMI:

#### To view the EBS volumes for an AMI

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>
2. In the navigation pane, click **AMIs**.
3. Select **EBS images** from the **Filter** drop-down list to get a list of EBS-backed AMIs.
4. Select the desired AMI, and look at the **Details** tab. At a minimum, the following information is available for the root device:
  - **Root Device Type** (ebs)
  - **Root Device Name** (for example, `/dev/sda1`)
  - **Block Devices** (for example, `/dev/sda1=snap-e1eb279f:8:true`)

If the AMI was created with additional EBS volumes using a block device mapping, the **Block Devices** field displays the mapping for those additional EBS volumes as well. (Recall that this screen doesn't display instance store volumes.)

## Command Line Interface

Use `ec2-describe-images` *ami\_id* to enumerate the EBS volumes in the block device mapping for an AMI.

For example, run this command at a command prompt to get information about an AMI, including its block device mapping. (Be sure to set the `EC2_PRIVATE_KEY` and `EC2_CERT` environment variables first.)

```
ec2-describe-images ami-72aa081b
```

The output includes the block device mapping for your AMI. The following information is available for EBS volumes:

- The device name within Amazon EC2
- The ID of the snapshot used when creating the block device (optional)
- The size of the volume, in GiB
- The volume type.

Here's example output for a Linux AMI.

```
BLOCKDEVICEMAPPING    /dev/sda1    snap-e1eb279f    8    standard
BLOCKDEVICEMAPPING    /dev/sdh     snap-d5eb27ab    200   standard
BLOCKDEVICEMAPPING    /dev/sdj     snap-d5eb27ab    100   standard
```

For more information, see [ec2-describe-images](#).

## Instance Block Device Mapping

By default, an instance that you launch includes any storage devices specified in the block device mapping of the AMI from which you launched the instance. You can specify changes to the block device mapping for an instance when you launch it, and these updates overwrite or merge with the block device mapping of the AMI. However, you can't modify the block device mapping entry for the root device volume.

## Updating the Block Device Mapping when Launching an Instance

You can add EBS volumes and instance store volumes to an instance when you launch it. Note that updating the block device mapping for an instance doesn't make a permanent change to the block device mapping of the AMI from which it was launched.

### AWS Management Console

#### To add volumes to an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the Amazon EC2 console dashboard, click **Launch Instance**.
3. On the **Choose an Amazon Machine Image (AMI)** page, choose the AMI to use and click **Select**.

4. Follow the wizard to complete the **Choose an Instance Type** and **Configure Instance Details** pages.
5. On the **Add Storage** page, you can modify the root volume, EBS volumes, and instance store volumes as follows:
  - To change the size of the root volume, locate the **Root** volume under the **Type** column, and change its **Size** field.
  - To suppress an EBS volume specified by the block device mapping of the AMI used to launch the instance, locate the volume and click its **Delete** icon.
  - To add an EBS volume, click **Add New Volume**, ensure you select **EBS** from the **Type** list, and fill in the fields (**Device**, **Snapshot**, and so on).
  - To suppress an instance store volume specified by the block device mapping of the AMI used to launch the instance, locate the volume, and click its **Delete** icon.
  - To add an instance store volume, click **Add New Volume**, select **Instance Store** from the **Type** list, and select a device name from **Device**.
6. Complete the remaining wizard pages, and then click **Launch**.

## Command Line Interface

Use **ec2-run-instances -b "devicename=blockdevice"** to define an entry for the block device mapping for an instance.

### **devicename**

The device name within Amazon EC2

### **blockdevice**

To omit a mapping for the device from the AMI, specify `none`.

To add an instance store device, specify `ephemeral[0-3]`.

[EBS-backed instance only] To add an EBS volume, specify `snapshot-id:size:[true|false]`. To add an empty EBS volume, omit the snapshot ID. To indicate whether the EBS volume should be deleted on termination, specify `true` or `false`; the default value is `true`.

For example, suppose that an EBS-backed Linux AMI specifies the following block device mapping:

- `/dev/sdb=ephemeral0`
- `/dev/sdh=snap-92d333fb`
- `/dev/sdj=:100`

To prevent `/dev/sdj` from attaching to an instance launched from this AMI, use this option.

```
-b "/dev/sdj=none"
```

To increase the size of `/dev/sdh` to 300 GiB, use this option.

```
-b "/dev/sdh=:300"
```

Notice that we didn't need to specify the snapshot ID for `/dev/sdh`, because specifying the device name is enough to identify the volume.

To attach an additional instance store volume, `/dev/sdc`, use this option. If the instance type doesn't support more than one instance store volume, this option has no effect.

```
-b "/dev/sdc=ephemeral1"
```

For more information, see [ec2-run-instances](#).

## Viewing the EBS Volumes in an Instance Block Device Mapping

You can easily enumerate the EBS volumes mapped to an instance.

### Note

For instances launched before the release of the 2009-10-31 API, AWS can't display the block device mapping. You must detach and reattach the volumes so that AWS can display the block device mapping.

## AWS Management Console

Use the AWS Management Console to enumerate the EBS volumes in the block device mapping for an instance.

### To view the EBS volumes for an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>
2. In the navigation pane, click **Instances**.
3. Select **EBS root device** from the **Filter** list. This displays a list of Amazon EBS-backed instances.
4. Locate and click the desired instance and look at the details displayed in the **Description** tab. At a minimum, the following information is available for the root device:
  - **Root device type** (ebs)
  - **Root device** (for example, sda1)
  - **Block devices** (for example, sda1, sdh, and sdj)

If the instance was launched with additional EBS volumes using a block device mapping, the **Block devices** box displays those additional EBS volumes as well as the root device. (Recall that this dialog box doesn't display instance store volumes.)

<b>Root device type</b>	ebs
<b>Root device</b>	/dev/sda1
<b>Block devices</b>	/dev/sda1
	/dev/sdf

5. To display additional information about a block device, click its entry next to **Block devices**. This displays the following information for the block device:
  - **EBS ID** (vol-xxxxxxx)
  - **Root device type** (ebs)
  - **Attachment time** (yyyy-mmT hh:mm:ss.TZD)
  - **Block device status** (attaching, attached, detaching, detached)
  - **Delete on termination** (Yes, No)



## Command Line Interface

Use **ec2-describe-instances** *instance\_id* to enumerate the EBS volumes in the block device mapping for an instance.

For example, run this command at a command prompt to get information about an instance, including its block devices. (Be sure to set the EC2\_PRIVATE\_KEY and EC2\_CERT environment variables first.)

```
ec2-describe-instances i-xxxxxxxx
```

The output includes the block devices for your instance. The following information is available for EBS volumes:

- The device name within Amazon EC2
- The volume ID
- The time when the attachment was initiated
- Whether to delete the volume on instance termination
- The volume type (*standard* or *io1*)
- The number of input/output operations per second (IOPS) that the volume supports. (Not used with *standard* volumes.)

The output for a Linux instance looks something like this.

```
BLOCKDEVICE /dev/sda1 vol-xxxxxxxx yyyy-mmThh:mm:ss.STZD true
BLOCKDEVICE /dev/sdh vol-xxxxxxxx yyyy-mmThh:mm:ss.STZD true
BLOCKDEVICE /dev/sdj vol-xxxxxxxx yyyy-mmThh:mm:ss.STZD true
```

To view this information plus the status (attaching, attached, detaching, detached), run the **ec2-describe-instances** command with the **-v** option to see the full response.

For more info, see [ec2-describe-instances](#).

## Viewing the Instance Block Device Mapping for Instance Store Volumes

When you view the block device mapping for your instance, you can see only the EBS volumes, not the instance store volumes. You can use instance metadata to query the complete block device mapping. The base URI for all requests for instance metadata is <http://169.254.169.254/latest/>.

First, connect to your running instance. If the instance is running Linux, the GET command is already available. If the instance is running Windows, install wget on the instance, and replace GET with wget in the examples below.

Use this query on a running instance to get its block device mapping.

```
GET http://169.254.169.254/latest/meta-data/block-device-mapping/
```

The response includes the names of the block devices for the instance. For example, the output for an instance store-backed `m1.small` instance looks like this.

```
ami
ephemeral0
```

```
root
swap
```

The ami device is the root device as seen by the instance. The instance store volumes are named `ephemeral[0-3]`. The swap device is for the page file. If you've also mapped EBS volumes, they appear as `ebs1`, `ebs2`, and so on.

To get details about an individual block device in the block device mapping, append its name to the previous query, as shown here.

```
GET http://169.254.169.254/latest/meta-data/block-device-mapping/ephemeral0
```

For more info, see [Instance Metadata and User Data \(p. 290\)](#).

# Resources and Tags

---

The following topics describe features that enable you to manage Amazon EC2 resources, such as AMIs, instances, Amazon EBS volumes, and snapshots.

## Topics

- [Resource Locations \(p. 528\)](#)
- [Listing and Filtering Your Resources \(p. 529\)](#)
- [Tagging Your Amazon EC2 Resources \(p. 532\)](#)

## Resource Locations

The following table describes which Amazon EC2 resources are global, regional, or based on Availability Zone.

Resource	Type	Description
AWS Account	Global	You can use the same AWS account in all regions.
DevPay Product Codes	Global	You can use the same DevPay product codes in all regions.
SSH Key Pairs	Global or Regional	You can use the SSH key pairs that you create using Amazon EC2 only in the region where you created them. You can create and upload an RSA key pair that you can use in all regions. For more information, see <a href="#">Amazon EC2 Key Pairs (p. 385)</a> .
Amazon EC2 Resource Identifiers	Regional	Each resource identifier, such as an AMI ID, instance ID, EBS volume ID, or EBS snapshot ID, is tied to its region and can be used only in the region where you created the resource.

Resource	Type	Description
User-Supplied Resource Names	Regional	Each resource name, such as a security group name or an SSH key pair name, is tied to its region and can be used only in the region where you created the resource. Although you can create resources with the same name in multiple regions, they aren't related to each other.
AMIs	Regional	An AMI is tied to the region where its files are located within Amazon S3. You can copy an AMI from one region to another. For more information, see <a href="#">Copying AMIs (p. 84)</a> .
Elastic IP Addresses	Regional	An Elastic IP address is tied to a region and can be associated only with an instance in the same region.
Security Groups	Regional	A security group is tied to a region and can be assigned only to instances in the same region. You can't enable an instance to communicate with an instance outside its region using security group rules. Traffic from an instance in another region is seen as WAN bandwidth.
EBS Snapshots	Regional	An EBS snapshot is tied to its region and can only be used to create volumes in the same region. You can copy a snapshot from one region to another. For more information, see <a href="#">Copying an Amazon EBS Snapshot (p. 488)</a> .
EBS Volumes	Availability Zone	An Amazon EBS volume is tied to its Availability Zone and can be attached only to instances in the same Availability Zone.
Instances	Availability Zone	An instance is tied to the Availability Zones in which you launched it. However, note that its instance ID is tied to the region.

#### Related Topic

- [Regions and Availability Zones \(p. 7\)](#)

## Listing and Filtering Your Resources

As part of using Amazon EC2, you might use different *resources*. These include images, instances, Amazon EBS volumes, and snapshots. You can get a list of some types of resource using the Amazon EC2 console. You can get a list of each type of resource using its corresponding command or API action. For example, you can list Amazon Machine Images (AMI) using `ec2-describe-images` or `DescribeImages`, and you can list instances using `ec2-describe-instances` or `DescribeInstances`.

#### Topics

- [Filters \(p. 530\)](#)
- [Listing Resources \(p. 530\)](#)
- [Filtering Resources \(p. 531\)](#)

## Filters

The resulting lists of resources can be long, so you might want to filter the results to include only the resources that match certain criteria. For example, you can list only the public 64-bit Windows AMIs that use an Amazon EBS volume as the root device volumes. As another example, you can list only the Amazon EBS snapshots that have been tagged with a certain value.

You can specify multiple filter values. For example, you can list all the instances whose type is either `m1.small` or `m1.large`. The resource must match at least one of the values to be included in the resulting resource list. You can also specify multiple filters. For example, you can list all the instances whose type is either `m1.small` or `m1.large`, and that have an attached EBS volume that is set to delete when the instance terminates. The instance must match all your filters to be included in the results.

You can also use wildcards with the filter values. An asterisk (\*) matches zero or more characters, and a question mark (?) matches exactly one character. For example, you can use `*database*` as a filter value to get all EBS snapshots that include `database` in the description. If you were to specify `database` as the filter value, then only snapshots whose description equals `database` would be returned. Filter values are case sensitive. We support only exact string matching, or substring matching (with wildcards).

### Tip

Your search can include the literal values of the wildcard characters; you just need to escape them with a backslash before the character. For example, a value of `\*amazon?\` searches for the literal string `*amazon?\`.

## Listing Resources

You can list your resources using the console, command-line interface, or API actions.

### Listing Resources Using the AWS Management Console

You can view the most common EC2 resource types using the AWS Management Console. To view additional resources, use the command line interface or the API actions.

#### To list EC2 resources using the console

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, click the option that corresponds to the resource, such as **AMIs** or **Instances**.

- EC2 Dashboard
  - Events
  - Tags
  
- 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

## Listing Resources Using the CLI and API

Each resource type has a corresponding command or API request that you use to list resources of that type. For example, you can use the [ec2-describe-instances](#) command or the [DescribeInstances](#) action to list your instances.

The response contains information for all your instances.

## Filtering Resources

You can filter a list of your resources using the console, command-line interface, or API actions.

### Filtering Resources Using the AWS Management Console

You can perform some basic filtering and sorting of the most common resource types using the AWS Management Console. For additional filtering capabilities, use the command line interface or the API actions.

**To list volumes in the `us-east-1b` Availability Zone with a status of `available`**

1. In the navigation pane, click **Volumes**.
2. In the **Viewing** pane, select **Detached Volumes** from the drop-down list. (A detached volume is available to be attached to an instance in the same Availability Zone.)
3. In the **Viewing** pane, enter `us-east-1b` in the search field.
4. Any volumes that meet this criteria are displayed.

### To list public 64-bit Windows AMIs backed by Amazon EBS

1. In the navigation pane, click **AMIs**.
2. In the **Filter** pane, select **Public images**, **EBS images** and **Windows** from the drop-down lists.
3. Enter `x86_64` in the search field.
4. Any AMIs that meet this criteria are displayed.

## Filtering Resources Using the CLI and API

You can add one or more filters to an `ec2-describe-XXX` command or a `DescribeXXX` action to scope the results to resources that match certain criteria. For a list of the available filters for a given EC2 resource, go to the `ec2-describe-XXX` topic for that resource in the [Amazon Elastic Compute Cloud Command Line Reference](#), or go to the `DescribeXXX` topic for that resource in the [Amazon Elastic Compute Cloud API Reference](#).

### Tip

If you're using the command line tools on Windows, you might need to use quotation marks as shown in the examples in the [Amazon Elastic Compute Cloud Command Line Reference](#).

# Tagging Your Amazon EC2 Resources

To help you manage your instances, images, and other Amazon Elastic Compute Cloud (Amazon EC2) resources, you can assign your own metadata to each resource in the form of *tags*. This topic describes tags and shows you how to create them.


### Topics

- [Tag Basics \(p. 532\)](#)
- [Tag Restrictions \(p. 533\)](#)
- [Tagging Your Resources for Billing \(p. 534\)](#)
- [Working with Tags in the Console \(p. 535\)](#)
- [API and CLI Overview \(p. 539\)](#)

## Tag Basics

Tags enable you to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. Each tag consists of a key and a value, both of which you define. For example, you could define a set of tags for your account's Amazon EC2 instances that helps you track each instance's owner and stack level. We recommend that you devise a set of tag keys that meets your needs for each resource type. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.

The following diagram illustrates how tagging works. In this example, you've assigned two tags to each of your instances, one called `Owner` and another called `Stack`. Each of the tags also has an associated value.



Tags don't have any semantic meaning to Amazon EC2 and are interpreted strictly as a string of characters. Also, tags are not automatically assigned to your resources.

You can work with tags using the AWS Management Console, the Amazon EC2 command line interface (CLI), and the Amazon EC2 API.

You can assign tags only to resources that already exist. When you use the Amazon EC2 console, you can access a list of tags to add to an instance, which will be applied immediately after the instance is created. If you add a tag that has the same key as an existing tag on that resource, the new value overwrites the old value. You can edit tag keys and values, and you can remove tags from a resource at any time. You can set a tag's value to the empty string, but you can't set a tag's value to null.

If you're using AWS Identity and Access Management (IAM), you can control which users in your AWS account have permission to create, edit, or delete tags. For more information about IAM, see [Controlling Access to Amazon EC2 Resources](#) (p. 399).

## Tag Restrictions

The following basic restrictions apply to tags:

- Maximum number of tags per resource—10
- Maximum key length—127 Unicode characters
- Maximum value length—255 Unicode characters
- Tag keys and values are case sensitive.
- Do not use the `aws:` prefix in your tag names or values because it is reserved for AWS use. You can't edit or delete tag names or values with this prefix.

You can't terminate, stop, or delete a resource based solely on its tags; you must specify the resource identifier. For example, to delete snapshots that you tagged with a tag key called `DeleteMe`, you must first get a list of those snapshots using `DescribeSnapshots` with a filter that specifies the tag. Then you use `DeleteSnapshots` with the resource identifiers of the snapshots (for example, `snap-1a2b3c4d`). You can't call `DeleteSnapshots` with a filter that specified the tag. For more information about using filters when listing your resources, see [Listing and Filtering Your Resources](#) (p. 529).

You can tag public or shared resources, but the tags you assign are available only to your AWS account and not to the other accounts sharing the resource.



You can't tag all resources, and some you can only tag using API actions or the command line. The following table lists all Amazon EC2 resources and the tagging restrictions that apply to them, if any. Resources with tagging restrictions of None can be tagged with API actions, the CLI, and the console.

Resource	Tagging support	Tagging restrictions
AMI	Yes	None
Bundle Task	No	
Customer Gateway	Yes	None
DHCP Option	Yes	None
EBS Volume	Yes	None
Elastic IP	No	
Instance	Yes	None
Internet Gateway	Yes	None
Key Pair	No	
Load Balancer	No	
Network ACL	Yes	None
Network Interface	Yes	None
Placement Group	No	
Reserved Instance	Yes	None
Reserved Instance Listing	No	
Route Table	Yes	None
Spot Instance Request	Yes	Tag with CLI and API only
Security Group - EC2 Classic	Yes	None
Security Group - VPC	Yes	None
Snapshot	Yes	None
Subnet	Yes	None
Virtual Private Gateway	Yes	None
VPC	Yes	None
VPN Connection	Yes	None

For more information about tagging using the AWS console, see [Working with Tags in the Console \(p. 535\)](#). For more information about tagging using the API or command line, see [API and CLI Overview \(p. 539\)](#).

## Tagging Your Resources for Billing

You can use tags to organize your AWS bill to reflect your own cost structure. To do this, sign up to get your AWS account bill with tag key values included. Then, to see the cost of your combined resources, you can organize your billing information based on resources that have the same tag key values. For

example, you can tag several resources with a specific application name, and then organize your billing information to see the total cost of that application across several services. For more information, see [Cost Allocation and Tagging](#) in *About AWS Account Billing*.

## Working with Tags in the Console

Using the Amazon EC2 console, you can see which tags are in use across all of your Amazon EC2 resources in the same region. You can view tags by resource and by resource type, and you can also view how many items of each resource type are associated with a specified tag. You can also use the Amazon EC2 console to apply or remove tags from one or more resources at a time.

### Topics

- [Displaying Tags](#) (p. 535)
- [Adding and Deleting Tags on an Individual Resource](#) (p. 536)
- [Adding and Deleting Tags to a Group of Resources](#) (p. 537)
- [Adding a Tag When You Launch an Instance](#) (p. 538)
- [Filtering a List of Resources by Tag](#) (p. 539)

## Displaying Tags

You can display tags in two different ways in the Amazon EC2 console. You can display the tags for an individual resource or for all resources.

### To display tags for individual resources

When you select a resource-specific page in the Amazon EC2 console, it displays a list of those resources. For example, if you select **Instances** from the navigation pane, the console displays a list of Amazon EC2 instances. When you select a resource from one of these lists (e.g., an instance), if the resource supports tags, you can view and manage its tags. On most resource pages, you can view the tags in the **Tags** tab on the details pane. The following image shows the **Tags** tab for an instance with two tags: Name = DNS Server and Purpose = Network Management.

You can add a column to the resource list that displays all values for tags with the same key. This column enables you to sort and filter the resource list by the tag. There are two ways to add a new column to the resource list to display your tags.

- On the **Tags** tab, click **Show Column** for the tag.
- Click the **Show/Hide Columns** gear-shaped icon, and in the **Show/Hide Columns** dialog box, select the tag key under **Your Tag Keys**.

### To display tags for all resources

You can display tags across all resources by selecting **Tags** from the navigation pane in the Amazon EC2 console. The following image shows the **Tags** pane, which lists all tags in use by resource type.

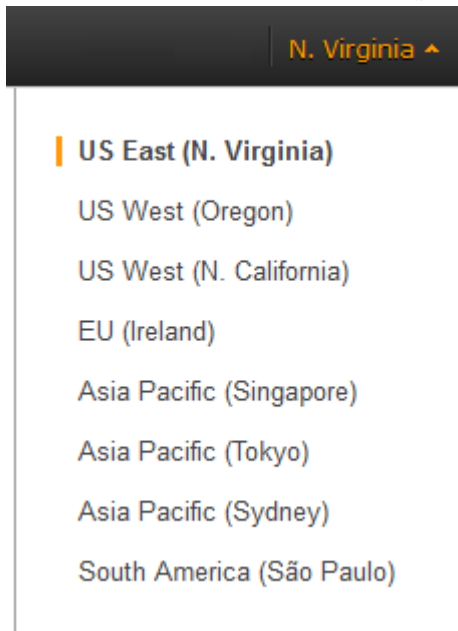
	Tag Key	Tag Value	Total	Instances	AMIs	Volumes
Manage Tag	Name	DNS Server	1	1	0	0
Manage Tag	Owner	TeamB	2	0	0	2
Manage Tag	Owner	TeamA	2	0	0	2
Manage Tag	Purpose	Project2	1	0	0	1
Manage Tag	Purpose	Logs	1	0	0	1
Manage Tag	Purpose	Network Management	1	1	0	0
Manage Tag	Purpose	Project1	2	0	0	2

## Adding and Deleting Tags on an Individual Resource

You can manage tags for an individual resource directly from the resource's page. If you are managing an AMI's tags, the procedures are different from that of other resources. All procedures are explained below.

### To add a tag to an individual resource

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, select the region that meets your needs. This choice is important because some Amazon EC2 resources can be shared between regions, while others can't. For more information, see [Resource Locations](#) (p. 528).



3. In the navigation pane, click a resource type (for example, **Instances**).
4. Select the resource from the resource list.
5. Select the **Tags** tab in the details pane.
6. Click the **Add/Edit Tags** button.

7. In the **Add/Edit Tags** dialog box, specify the key and value for each tag, and then click **Save**.

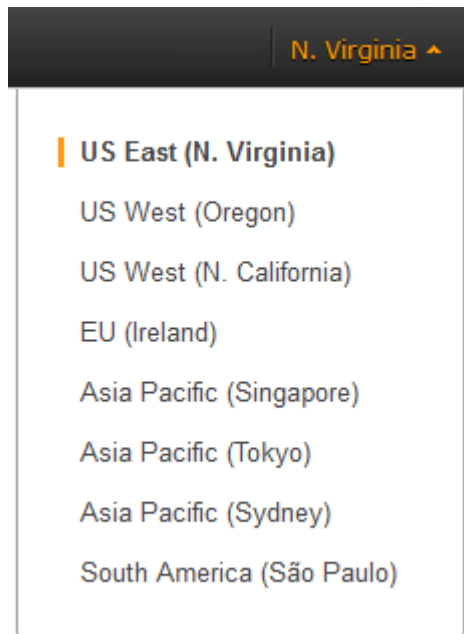
### To delete a tag from an individual resource

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, select the region that meets your needs. This choice is important because some Amazon EC2 resources can be shared between regions, while others can't. For more information, see [Resource Locations](#) (p. 528).
3. In the navigation pane, click a resource type (for example, **Instances**).
4. Select the resource from the resource list.
5. Select the **Tags** tab in the details pane.
6. Click **Add/Edit Tags**, click the **Delete** icon for the tag, and click **Save**.

## Adding and Deleting Tags to a Group of Resources

### To add a tag to a group of resources

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, select the region that meets your needs. This choice is important because some Amazon EC2 resources can be shared between regions, while others can't. For more information, see [Resource Locations](#) (p. 528).



3. In the navigation pane, click **Tags**.
4. At the top of the content pane, click **Manage Tags**.
5. From the **Filter** drop-down list, select the type of resource (for example, instances) that you want to add tags to.
6. In the resources list, select the check box next to each resource that you want to add tags to.
7. In the **Key** and **Value** boxes under **Add Tag**, type the tag key and values you want, and then click **Add Tag**.

#### Note

If you add a new tag with the same tag key as an existing tag, the new tag overwrites the existing tag.

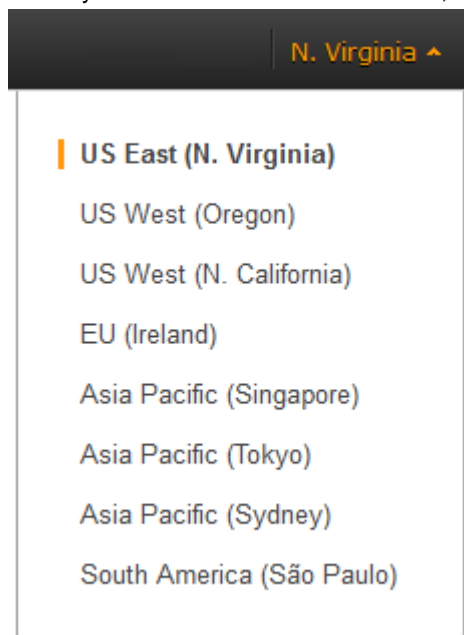
### To remove a tag from a group of resources

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, select the region that meets your needs. This choice is important because some Amazon EC2 resources can be shared between regions, while others can't. For more information, see [Resource Locations \(p. 528\)](#).
3. In the navigation pane, click **Tags**.
4. At the top of the content pane, click **Manage Tags**.
5. To view the tags in use, click the **Show/Hide Columns** gear-shaped icon, and in the **Show/Hide Columns** dialog box, select the tag keys you want to view, and then click **Close**.
6. From the **Filter** drop-down list, select the type of resource (for example, instances) that you want to remove tags from.
7. In the resource list, select the check box next to each resource that you want to remove tags from.
8. Under **Remove Tag**, click in the **Key** box to select a key, or type its name, and then click **Remove Tag**.

## Adding a Tag When You Launch an Instance

### To add a tag using the Launch Wizard

1. From the navigation bar, select the region for the instance. This choice is important because some Amazon EC2 resources can be shared between regions, while others can't. Select the region that meets your needs. For more information, see [Resource Locations \(p. 528\)](#).



2. Click the **Launch Instance** button on the EC2 dashboard.
3. The **Choose an Amazon Machine Image (AMI)** page displays a list of basic configurations called Amazon Machine Images (AMIs). Choose the AMI that you want to use and click its **Select** button. For more information about selecting an AMI, see [Finding a Suitable AMI \(p. 47\)](#).
4. On the **Configure Instance Details** page, configure the instance settings as necessary, then click **Next: Add Storage**.
5. On the **Add Storage** page, you can specify additional storage volumes for your instance. Click **Next: Tag Instance** when done.

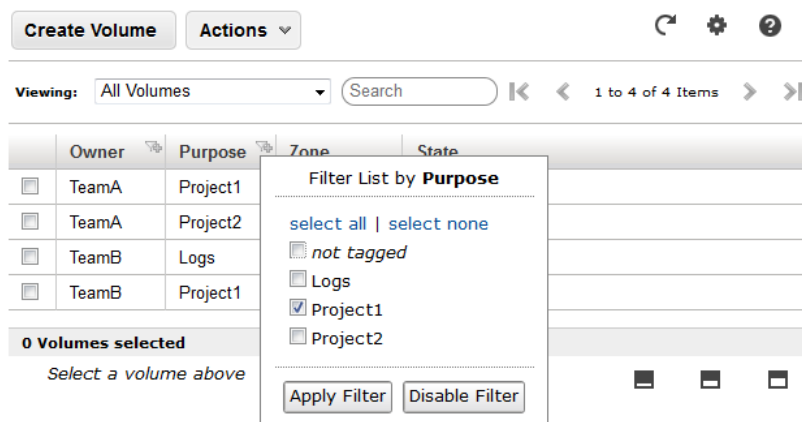
6. On the **Tag Instance** page, specify tags for the instance by providing key and value combinations. Click **Create Tag** to add more than one tag to your resource. Click **Next: Configure Security Group** when you are done.
7. On the **Configure Security Group** page, you can choose from an existing security group that you own, or let the wizard create a new security group for you. Click **Review and Launch** when you are done.
8. Review your settings. When you're satisfied with your selections, click **Launch**. Select an existing key pair or create a new one, select the acknowledgment check box, and then click **Launch Instances**.

## Filtering a List of Resources by Tag

You can filter your list of resources based on one or more tag keys and tag values.

### To filter a list of resources by tag

1. Display a column for the tag as follows:
  - a. Select one of the resources.
  - b. Select the **Tags** tab in the details pane.
  - c. Locate the tag in the list and click **Show Column**.
2. Click the filter icon in the top right corner of the column for the tag to display the filter list.
3. Select the tag values, and then click **Apply Filter** to filter the results list.



## API and CLI Overview

Use the following API and CLI commands to add, update, list, and delete the tags for your resources. The documentation for each command provides examples.

Description	Command	API Action
Adds or overwrites one or more tags for the specified resource or resources.	<code>ec2-create-tags</code>	<a href="#">CreateTags</a>

Description	Command	API Action
Deletes the specified tags from the specified resource or resources.	<a href="#">ec2-delete-tags</a>	<a href="#">DeleteTags</a>
Describes one or more tags for your resources.	<a href="#">ec2-describe-tags</a>	<a href="#">DescribeTags</a>

You can also filter a list of resources according to their tags. For example syntax, see [Filtering Resources \(p. 531\)](#). For lists of supported filters, see the relevant `ec2-describe` command in the [Amazon Elastic Compute Cloud Command Line Reference](#), or the relevant `Describe` API action in the [Amazon Elastic Compute Cloud API Reference](#).

# Setting Up the Amazon EC2 Command Line Interface Tools on Linux/UNIX

---

The Amazon EC2 command line interface tools (also called the *CLI tools*) wrap the Amazon EC2 API actions. These tools are written in Java and include shell scripts for both Windows and Linux/UNIX/Mac OSX. For a detailed reference guide that describes the commands, see the [Amazon Elastic Compute Cloud Command Line Reference](#).

## Note

Alternatively, you can use the AWS Command Line Interface (CLI), which provides commands for a broad set of AWS products, including Amazon EC2. To get started with the AWS CLI, see [AWS Command Line Interface User Guide](#). For more information about the AWS CLI commands for Amazon EC2, see [ec2](#).

Before you can use the Amazon EC2 CLI tools on your computer or your instance, you must install the tools and set the environment variables used by the tools. Use the set of directions for your operating system:

- [Set Up the Amazon EC2 CLI Tools on Amazon Linux \(p. 541\)](#)
- [Set Up the Amazon EC2 CLI Tools on RHEL, Ubuntu, or Mac OS \(p. 542\)](#)
- [Installing the Amazon EC2 CLI Tools on Windows](#) (in the *Amazon Elastic Compute Cloud Microsoft Windows Guide*)

## Set Up the Amazon EC2 CLI Tools on Amazon Linux

Instances that you launch using an Amazon Linux AMI already include the Amazon EC2 CLI tools.

Each time you use the Amazon EC2 CLI tools on your instance, you must provide your identity. Your access keys identify you to the Amazon EC2 CLI tools. There are two types of access keys: access key IDs and secret access keys. You should have stored your access keys in a safe place when you created them. Although you can retrieve your access key ID from the [Your Security Credentials](#) page, you can't



retrieve your secret access key. Therefore, if you can't find your secret access key, you'll need to create new access keys before you can use the CLI tools.

The easiest way to provide your access keys to the Amazon EC2 CLI is to set the `AWS_ACCESS_KEY` and `AWS_SECRET_KEY` environment variables. First, add the following lines to `~/.bashrc` and save the file.

```
export AWS_ACCESS_KEY=your-aws-access-key-id
export AWS_SECRET_KEY=your-aws-secret-key
```

After you've updated `~/.bashrc`, run the following command:

```
$ source ~/.bashrc
```

To verify that your CLI tools are set up correctly, run the following command:

```
$ ec2-describe-regions
```

If you get an error that required option `-O` is missing, check the setting of `AWS_ACCESS_KEY`, fix any errors, and try the command again.

If you get an error that required option `-W` is missing, check the setting of `AWS_SECRET_KEY`, fix any errors, and try the command again.

The default region for the Amazon EC2 CLI tools is `us-east-1`. For information about configuring the Amazon EC2 CLI tools to use a different region, see [\(Optional\) Set the Region \(p. 547\)](#).

## Set Up the Amazon EC2 CLI Tools on RHEL, Ubuntu, or Mac OS

You must complete the following setup tasks before you can use the Amazon EC2 CLI tools on your own computer.

### Topics

- [Download the CLI Tools \(p. 542\)](#)
- [Tell the Tools Where Java Lives \(p. 543\)](#)
- [Tell the CLI Tools Where They Live \(p. 545\)](#)
- [Tell the CLI Tools Who You Are \(p. 545\)](#)
- [\(Optional\) Tell the CLI Tools To Use a Proxy Server \(p. 546\)](#)
- [Verify the Tools Setup \(p. 546\)](#)
- [\(Optional\) Set the Region \(p. 547\)](#)

## Download the CLI Tools

The CLI tools are available as a ZIP file on this site: [Amazon EC2 CLI Tools](#). The ZIP file is self-contained; no installation is required. You can simply download the file and unzip it.

**Optional step before you unzip:** For detailed information about authenticating the download before unzipping the file, see [Verify the Signature of the Tools Download \(p. 547\)](#).

## Tell the Tools Where Java Lives

The Amazon EC2 CLI tools require Java. If you don't have Java 1.6 or later installed, download and install Java. Either a JRE or JDK installation is acceptable. To view and download JREs for a range of platforms, see [Java Downloads](#).

### Important

Instances that you launch using the Amazon Linux AMI already include Java.

The Amazon EC2 CLI read the `JAVA_HOME` environment variable to locate the Java runtime. This environment variable should specify the full path of the directory that contains a subdirectory named `bin` that contains the Java executable you installed (`java.exe`).

### To set the `JAVA_HOME` environment variable on Linux/UNIX and Mac OS

1. You can verify whether you have Java installed and where it is located using the following command:

```
$ which java
```

The following is example output.

```
/usr/bin/java
```

If the previous command does not return a location for the Java binary, you need to install Java. For help installing Java on your platform, see [Java Downloads](#).

2. Find the Java home directory on your system. The **which java** command executed earlier returns Java's location in the `$PATH` environment variable, but in most cases this is a symbolic link to the actual program; symbolic links do not work for the `JAVA_HOME` environment variable, so you need to locate the actual binary.
  - a. (Linux only) For Linux systems, you can recursively run the **file** command on the **which java** output until you find the binary.

```
$ file `which java`  
/usr/bin/java: symbolic link to `/etc/alternatives/java'
```

The `/usr/bin/java` location is actually a link to `/etc/alternatives/java`, so you need to run the **file** command on that location to see whether that is the real binary.

```
$ file /etc/alternatives/java  
/etc/alternatives/java: symbolic link to `/usr/lib/jvm/java-6-openjdk-  
amd64/jre/bin/java'
```

This returns a new location, which is the actual binary. Verify this by running the **file** command on this location.

```
$ file /usr/lib/jvm/java-6-openjdk-amd64/jre/bin/java  
/usr/lib/jvm/java-6-openjdk-amd64/jre/bin/java: ELF 64-bit LSB execut  
able...
```

This location is the actual binary (notice that it is listed as an executable). The Java home directory is where `bin/java` lives; in this example, the Java home directory is `/usr/lib/jvm/java-6-openjdk-amd64/jre`.

- b. (Mac OSX only) For Mac OSX systems, the `/usr/libexec/java_home` command returns a path suitable for setting the `JAVA_HOME` variable.

```
$ /usr/libexec/java_home
/System/Library/Java/JavaVirtualMachines/1.6.0.jdk/Contents/Home
```

3. Set `JAVA_HOME` to the full path of the Java home directory.
  - a. (Linux only) For the Linux example above, set the `JAVA_HOME` variable to the directory where `bin/java` was located in [Step 2.a \(p. 543\)](#).

```
$ export JAVA_HOME="/usr/lib/jvm/java-6-openjdk-amd64/jre"
```

**Note**

If you are using Cygwin, `JAVA_HOME` should contain a Windows path.

- b. (Mac OSX only) For the Mac OSX example above, set the `JAVA_HOME` variable to `$(/usr/libexec/java_home)`. The following command sets this variable to the output of the `java_home` command; the benefit of setting the variable this way is that it updates to the correct value if you change the location of your Java installation later.

```
$ export JAVA_HOME=$(/usr/libexec/java_home)
```

4. You can verify your `JAVA_HOME` setting using this command.

```
$ $JAVA_HOME/bin/java -version
```

If you've set the environment variable correctly, the output looks something like this.

```
java version "1.6.0_27"
OpenJDK Runtime Environment (IcedTea6 1.12.6) (6b27-1.12.6-1ubuntu0.12.04.2)
OpenJDK 64-Bit Server VM (build 20.0-b12, mixed mode)
```

5. Add this environment variable definition to your shell start up scripts so that it is set every time you log in or spawn a new shell. The name of this startup file differs across platforms (in Mac OSX, this file is commonly called `~/.bash_profile` and in Linux, it is commonly called `~/.profile`), but you can find it with the following command:

```
$ ls -al ~ | grep profile
```

If the file does not exist, you can create it. Use your favorite text editor to open the file that is listed by the previous command, or to create a new file with that name. Then edit it to add the variable definition you set in [Step 3 \(p. 544\)](#).

6. Verify that the variable is set properly for new shells by opening a new terminal window and testing that the variable is set with the following command.

**Note**

If the following command does not correctly display the Java version, try logging out, logging back in again, and then retrying the command.

```
$ $JAVA_HOME/bin/java -version
```

## Tell the CLI Tools Where They Live

The Amazon EC2 CLI tools read the `EC2_HOME` environment variable to locate supporting libraries. Before using these tools, set `EC2_HOME` to the directory path where you unzipped them. This directory is named `ec2-api-tools-w.x.y.z` (where `w`, `x`, `y`, and `z` are components of the version number). It contains sub-directories named `bin` and `lib`.

In addition, to make things a little easier, you can add the `bin` directory for the CLI tools to your system path. The examples in the *Amazon Elastic Compute Cloud User Guide* assume that you have done so.

You can set the `EC2_HOME` and `PATH` environment variables as follows. Add them to your shell start up scripts so that they're set every time you log in or spawn a new shell.

### To set the `EC2_HOME` and `PATH` environment variables on Linux/UNIX

1. Use this command to set the `EC2_HOME` environment variable.

```
export EC2_HOME=<path-to-tools>
```

#### Note

If you are using Cygwin, `EC2_HOME` must use Linux/UNIX paths (for example, `/usr/bin` instead of `C:\usr\bin`). Additionally, the value of `EC2_HOME` cannot contain any spaces, even if the value is quoted or the spaces are escaped.

2. You can update your `PATH` as follows.

```
export PATH=$PATH:$EC2_HOME/bin
```

## Tell the CLI Tools Who You Are

Your access keys identify you to the Amazon EC2 CLI tools. There are two types of access keys: access key IDs and secret access keys. You should have stored your access keys in a safe place when you created them. Although you can retrieve your access key ID from the [Your Security Credentials](#) page, you can't retrieve your secret access key. Therefore, if you can't find your secret access key, you'll need to create new access keys before you can use the CLI tools.

Every time you issue a command, you must specify your access keys using the `--aws-access-key` and `--aws-secret-key` (or `-O` and `-W`) options. Alternatively, you might find it easier to store your access keys using the following environment variables:

- `AWS_ACCESS_KEY`—Your access key ID
- `AWS_SECRET_KEY`—Your secret access key

If these environment variables are set properly, their values serve as the default values for these required options, so you can omit them from the commands. You can add them to your shell start up scripts so that they're set every time you log in or spawn a new shell.

You can set these environment variables as follows.

```
export AWS_ACCESS_KEY=your-aws-access-key-id
export AWS_SECRET_KEY=your-aws-secret-key
```

## (Optional) Tell the CLI Tools To Use a Proxy Server

If the computer you have installed the Amazon EC2 CLI tools on requires the use of a proxy server, you must tell the CLI tools to use the proxy server with the `EC2_JVM_ARGS` environment variable.

The following table contains the proxy configuration properties that can be set for the `EC2_JVM_ARGS` variable. The properties that are required will depend on the type of proxy server being used. For example, the `http.proxyDomain` and `http.proxyWorkstation` properties are only used with a Windows NTLM proxy.

Property	Description
<code>https.proxyHost</code>	HTTPS proxy host. Use when <code>EC2_URL</code> specifies an HTTPS host.
<code>https.proxyPort</code>	HTTPS proxy port. Use when <code>EC2_URL</code> specifies an HTTPS host.
<code>http.proxyHost</code>	HTTP proxy host. Use when <code>EC2_URL</code> specifies an HTTP host.
<code>http.proxyPort</code>	HTTP proxy port. Use when <code>EC2_URL</code> specifies an HTTP host.
<code>http.proxyDomain</code>	Proxy domain (HTTPS and HTTP)
<code>http.proxyWorkstation</code>	Proxy workstation (HTTPS and HTTP)
<code>http.proxyUser</code>	Proxy user name (HTTPS and HTTP)
<code>http.proxyPass</code>	Proxy password (HTTPS and HTTP)
<code>http.nonProxyHosts</code>	A list of hosts that should be reached directly, bypassing the proxy. Each item in the list is separated by ' '.

You set the `EC2_JVM_ARGS` variable with the **export** command:

```
export EC2_JVM_ARGS="-Dhttps.proxyHost=my.proxy.com -Dhttps.proxyPort=8080"
```

## Verify the Tools Setup

Let's quickly verify that your Amazon EC2 CLI tools are set up correctly. Run the following command to view your available regions.

```
$ ec2-describe-regions
```

If your environment variables are set correctly, the output lists regions and their corresponding service endpoints.

If you get an error that required option **-O** is missing, check the setting of `AWS_ACCESS_KEY`, fix any errors, and try the command again.

If you get an error that required option **-W** is missing, check the setting of `AWS_SECRET_KEY`, fix any errors, and try the command again.

## (Optional) Set the Region

By default, the Amazon EC2 CLI tools use the US East (Northern Virginia) region (`us-east-1`) with the `ec2.us-east-1.amazonaws.com` service endpoint URL. To access a different region with the CLI tools, you must set the `EC2_URL` environment variable to the proper service endpoint URL.

### To set the service endpoint URL

1. To list your available service endpoint URLs, call the **ec2-describe-regions** command, as shown in the previous section.
2. Set the `EC2_URL` environment variable using the service endpoint URL returned from the **ec2-describe-regions** command as follows.

```
export EC2_URL=https://<service_endpoint>
```

If you've already launched an instance using the console and wish to work with the instance using the CLI, you must specify the endpoint URL for the instance's region. You can verify the region for the instance by checking the region selector in the console navigation bar.

For more information about the regions and endpoints for Amazon EC2, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

## Verify the Signature of the Tools Download

Whenever you download an application from the Internet, you should authenticate the identity of the software publisher and check that the application has not been altered or corrupted since it was published. This protects you from installing a version of the application that contains a virus or other malicious code.

If you determine that the software for the CLI tools has been altered, do not unzip or install the file that you downloaded. Instead, contact Amazon Web Services.

### Topics

- [Overview](#) (p. 547)
- [Install the GPG Tools](#) (p. 548)
- [Authenticate the Public Key](#) (p. 549)
- [Verify the Signature of the Package](#) (p. 550)

## Overview

The first step is to establish trust with the software publisher: download the public key of the software publisher, check that the owner of the public key is who they claim to be, and then add the public key to your keyring. Your keyring is a collection of known public keys. You can then explicitly trust the public key, or trust it implicitly because the public key is trusted by someone with whom you have a pre-existing trust relationship.

After you've established the authenticity of the public key, you can use it to verify the signature of the application. Using security tools, you'll calculate a signature from the publisher's public key and your downloaded copy of the application. If the calculated signature matches the signature the software

developer has published for the application, you can have confidence that the application has not been altered.

Amazon EC2 CLI tools are signed using GnuPG, an open implementation of the Pretty Good Privacy (OpenPGP) standard for secure digital signatures. GnuPG provides authentication and integrity checking through a 128-bit digital signature. Amazon EC2 publishes a public key and signatures you can use to verify the downloaded Amazon EC2 CLI tools. For more information about PGP and GnuPG (GPG), see <http://www.gnupg.org>.

## Install the GPG Tools

If your operating system is Linux or UNIX, the GPG tools are likely already installed. To test whether the tools are installed on your system, type `gpg` at a command-line prompt. If the GPG tools are installed, you get a GPG command prompt. If the GPG tools are not installed, you get an error stating that the command cannot be found. You can install the GnuPG package from a repository.

### To install GPG Tools on Debian-based Linux

- From a terminal, run the following command.

```
apt-get install gnupg
```

### To install GPG Tools on Red Hat-based Linux

- From a terminal, run the following command.

```
yum install gnupg
```

### To install GPG Tools on Windows

- Download and install [Gpg4win](#), an implementation of GnuPG that runs on Windows.

During installation of Gpg4win, you can use the default values suggested by Gpg4win. As part of the installation process, you are asked whether you want to define a root certificate. Defining a root certificate is a way to establish trust with many software publishers; setting a root certificate establishes trust with all publishers trusted by that certificate. If you want to define a root certificate, follow the instructions in the text box. If not, click the check box to skip this step. Either option is fine for the purpose of verifying the signature of the Amazon EC2 CLI tools package.

### To install GPG Tools on Mac OS

- Download and install [GPGTools](#), an implementation of GnuPG that runs on Mac OS.

In addition to installing pre-compiled implementations of GnuPG, you can also download and compile the source code from <http://gnupg.org/download/index.en.html>.

After you have installed a set of GPG tools, use them to create a public-private key set. You'll need this later to sign changes to your trust status. If you have previously installed the GPG tools and already have a public-private key set, you can skip this step.

### To create a private key for the GPG tools

1. From the command line, run the following command.

```
gpg --gen-key
```

2. Answer the questions that follow. You can use the suggested default values. Make a note of the passphrase you use to create the private key. You'll need this value later.

## Authenticate the Public Key

The next step in the process is to authenticate the EC2 Packages public key and add it as a trusted key in your GPG keyring.

### To authenticate the EC2 Packages public key

1. Create a text file named `ec2-packages-public.key` and copy the public key from [EC2 Packages Public Key](#) into the text file. This includes everything from `-----BEGIN PGP PUBLIC KEY BLOCK-----` to `-----END PGP PUBLIC KEY BLOCK-----`. Save the text file.

#### Note

This text file must use ASCII encoding.

2. Import the EC2 Packages public key into your keyring using the following command in the directory where you saved the file `ec2-packages-public.key`.

```
gpg --import ec2-packages-public.key
```

The command returns results similar to the following. Make a note of the key value; you'll need it in the next step. In the example below, the key value is `0349E66A`.

```
gpg: key 0349E66A: public key "AWS EC2 Packages <ec2-packages@amazon.com>"
imported
gpg: Total number processed: 1
gpg:          imported: 1 (RSA: 1)
gpg: 3 marginal(s) needed, 1 complete(s) needed, PGP trust model
gpg: depth: 0  valid:  2  signed:  1  trust: 0-, 0q, 0n, 0m, 0f, 2u
gpg: depth: 1  valid:  1  signed:  0  trust: 0-, 0q, 0n, 0m, 1f, 0u
gpg: next trustdb check due at 2014-07-20
```

3. Verify the fingerprint by running the following command, where `key-value` is replaced by the value from the previous step.

```
gpg --fingerprint key-value
```

This command returns results similar to the following. Compare the key fingerprint to that published on [EC2 Packages Public Key](#). They should match. If they don't, do not continue to install the CLI tools, and contact Amazon Web Services.

```
pub 4096R/0349E66A 2011-04-04
Key fingerprint = A262 37CF 2294 C30E 9844 96C9 116B 3651 0349 E66A
uid AWS EC2 Packages <ec2-packages@amazon.com>
```



4. If the fingerprint matches the one published on the [aws.amazon.com](http://aws.amazon.com) website, you may choose to trust EC2 Packages public key. To do so, run the following command, where *key-value* is replaced by the key value from Step 2.

```
gpg --edit-key key-value
```

Type the following command at the GPG tools prompt.

```
trust
```

The GPG tools ask you to establish a level of trust for EC2 Packages, with a prompt such as the following. To learn more about these trust options, go to the [The GNU Privacy Handbook](#).

```
Please decide how far you trust this user to correctly verify other users'
keys
(by looking at passports, checking fingerprints from different sources,
etc.)

 1 = I don't know or won't say
 2 = I do NOT trust
 3 = I trust marginally
 4 = I trust fully
 5 = I trust ultimately
 m = back to the main menu

Your decision?
```

Type 4 and press the Enter key.

5. Sign the key with your private key (created when you installed the GPG tools) to set the new trust level. Do this using the following command.

```
sign
```

You are asked to confirm, type *y*, and press the Enter key. You are asked for the passcode that you used when you created your private key. Type your passcode and press the Enter key.

6. Save your changes using the following command.

```
save
```

This saves your changes to the keyring. It should also exit the PGP session. If it doesn't, press CTRL+Z to exit the PGP session and return to the main terminal.

## Verify the Signature of the Package

With the GPG tools installed and the EC2 Packages public key authenticated and the EC2 Packages public key trusted, you're ready to check the signature of the Amazon EC2 CLI tools package.

### To verify the signature of Amazon EC2 CLI tools package

1. Download the Amazon EC2 CLI tools package, `ec2-api-tools.zip`, from [Amazon EC2 CLI Tools](#).

2. Create a new text file name `ec2-api-tools.zip.asc` and copy the contents of the Amazon EC2 [command line tools signature](#) into this file. Copy everything including the `-----BEGIN PGP SIGNATURE-----` to `-----END PGP SIGNATURE-----` lines. Save the file.
3. Verify the signature of the CLI tools by typing the following at a command line prompt in the directory where you saved the file `ec2-api-tools.zip.asc` and the CLI package `ec2-api-tools.zip`. Both files must be present.

```
gpg --verify ec2-api-tools.zip.asc ec2-api-tools.zip
```

The output should be something like the following.

```
gpg: Signature made Mon Mar 12 14:51:33 2012 PDT using RSA key ID 0349E66A  
gpg: Good signature from "AWS EC2 Packages <ec2-packages@amazon.com>"
```

If the output contains the phrase 'Good signature from "AWS EC2 Packages <ec2-packages@amazon.com>"' the signature has successfully been verified, and you can proceed to unzip and install the Amazon EC2 CLI tools. If the output includes the phrase "BAD signature", check that you performed the procedure correctly. If you continue to get this response, contact Amazon Web Services and do not unzip or install the file that you downloaded.

# Making API Requests

---

We provide the Query API for Amazon EC2, as well as software development kits (SDK) for Amazon Web Services (AWS) that enable you to access Amazon EC2 from your preferred programming language.

## Topics

- [Required Knowledge](#) (p. 552)
- [Available Libraries](#) (p. 552)
- [Query Requests](#) (p. 553)
- [Troubleshooting API Request Errors](#) (p. 555)
- [Ensuring Idempotency](#) (p. 557)
- [SOAP Requests](#) (p. 559)

## Required Knowledge

If you plan to access Amazon EC2 through an API, you should be familiar with the following:

- XML
- Web services
- HTTP requests
- One or more programming languages, such as Java, PHP, Perl, Python, Ruby, C#, or C++.

## Available Libraries

AWS provides libraries, sample code, tutorials, and other resources for software developers who prefer to build applications using language-specific APIs instead of submitting a request over HTTP or HTTPS. These libraries provide basic functions that automatically take care of tasks such as cryptographically signing your requests, retrying requests, and handling error responses, so that it is easier for you to get started.

For more information, see [AWS SDKs and Tools](#).

# Query Requests

Query requests are HTTP or HTTPS requests that use the HTTP verb GET or POST and a Query parameter named `Action`. For a list of Amazon EC2 API actions, see [Actions](#).

## Topics

- [Structure of a GET Request \(p. 553\)](#)
- [Endpoints \(p. 554\)](#)
- [Query Parameters \(p. 554\)](#)
- [Query API Authentication \(p. 554\)](#)
- [Query Response Structures \(p. 555\)](#)

## Structure of a GET Request

The Amazon EC2 documentation presents the GET requests as URLs, which can be used directly in a browser.

### Tip

Because the GET requests are URLs, you must URL encode the parameter values. In the Amazon EC2 documentation, we leave the example GET requests unencoded to make them easier to read.

The request consists of the following:

- **Endpoint:** The URL that serves as the entry point for the web service.
- **Action:** The action that you want to perform (for example, use `RunInstance` to run an instance).
- **Parameters:** Any parameters for the action; each parameter is separated by an ampersand (&). You must also include authorization parameters that AWS uses to ensure the validity and authenticity of the request. For more information, see [Query API Authentication \(p. 554\)](#).

The following is an example request that launches instances:

```
https://ec2.amazonaws.com/?Action=RunInstances&ImageId=ami-2bb65342&MaxCount=3&MinCount=1&Placement.AvailabilityZone=us-east-1b&Monitoring.Enabled=true&AWSAccessKeyId=0GS7553JW74RRM612K02EXAMPLE&Version=2013-10-01&Expires=2010-10-10T12:00:00Z&Signature=lBP67vCvG1DMBQ1dofZxg8E8SUEXAMPLE&SignatureVersion=2&SignatureMethod=HmacSHA256
```

To make these example requests even easier to read, the Amazon EC2 documentation presents them in the following format:

```
https://ec2.amazonaws.com/?Action=RunInstances
&ImageId=ami-2bb65342
&MaxCount=3
&MinCount=1
&Placement.AvailabilityZone=us-east-1b
&Monitoring.Enabled=true
&AWSAccessKeyId=0GS7553JW74RRM612K02EXAMPLE
&Version=2013-10-01
&Expires=2010-10-10T12:00:00Z
&Signature=lBP67vCvG1DMBQ1dofZxg8E8SUEXAMPLE
```

```
&SignatureVersion=2  
&SignatureMethod=HmacSHA256
```

The first line specifies the endpoint of the request. After the endpoint is a question mark (?), which separates the endpoint from the parameters.

The `Action` parameter indicates the action to perform. For a complete list of actions, see [Actions](#) in the *Amazon Elastic Compute Cloud API Reference*.

The remaining lines specify additional parameters for the request.

## Endpoints

An endpoint is a URL that serves as an entry point for a web service. You can select a regional endpoint for Amazon EC2 when you make your requests to reduce latency. For more information about regions, see [Region and Availability Zone Concepts](#) (p. 7). For information about the endpoints for Amazon EC2, see [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.

If you specify the general endpoint, `ec2.amazonaws.com`, we use the endpoint for `us-east-1`. To use a different region, specify its associated endpoint. For example, if you specify `ec2.us-west-2.amazonaws.com` as the endpoint, we direct your request to the `us-west-2` endpoint.

## Query Parameters

Each Query request must include required common parameters to handle authentication and selection of an action. For more information, see [Common Query Parameters](#) in the *Amazon Elastic Compute Cloud API Reference*.

Some operations take lists of parameters. These lists are specified using the `param.n` notation, where `n` is an integer starting from 1.

The following example adds multiple devices to a block device mapping using a list of `BlockDeviceMapping` parameters.

```
http://ec2.amazonaws.com/?Action=RunInstances  
&ImageId.1=ami-72aa081b  
...  
&BlockDeviceMapping.1.DeviceName=/dev/sdj  
&BlockDeviceMapping.1.Ebs.NoDevice=true  
&BlockDeviceMapping.2.DeviceName=/dev/sdh  
&BlockDeviceMapping.2.Ebs.VolumeSize=300  
&BlockDeviceMapping.3.DeviceName=/dev/sdc  
&BlockDeviceMapping.3.VirtualName=ephemeral1  
&AUTHPARAMS
```

## Query API Authentication

You can send Query requests over either the HTTP or HTTPS protocol.

Regardless of which protocol you use, you must include a signature in every Query request. Amazon EC2 uses Signature Version 2. For more information, see [Signature Version 2 Signing Process](#) in the *Amazon Web Services General Reference*.

In the example Query requests we present in the Amazon EC2 documentation, we omit the parameters related to authentication to make it easier for you to focus on the parameters for the action. We replace



- Modify actions, such as `RunInstances` and `CreateVolumes`. These requests create or modify resources, so they have a lower request limit than describe calls.
- The `CreateKeyPair`, `GetConsoleOutput`, `AuthorizeSecurityGroupIngress`, and `RevokeSecurityGroupIngress` actions. These requests take the most time and resource to complete, so they have the lowest request limit.

If an API request exceeds the API request rate for its category, the request returns the `RequestLimitExceeded` error code. To prevent this error, ensure that your application doesn't retry API requests at a high rate. You can do this by using care when polling and by using exponential back-off retries.

## Polling

Your application might need to call an API repeatedly to check for an update in status. Before you start polling, give the request time to potentially complete. When you begin polling, use an appropriate sleep interval between successive requests. For best results, use an increasing sleep interval.

## Retries or batch processing

Your application might need to retry an API request after it fails, or to process multiple resources (for example, all your volumes). To lower the rate of API requests, use an appropriate sleep interval between successive requests. For best results, use an increasing or variable sleep interval.

## Calculating the sleep interval

When you have to poll or retry an API request, we recommend using an exponential backoff algorithm to calculate the sleep interval between API calls. The idea behind exponential backoff is to use progressively longer waits between retries for consecutive error responses. For more information, and implementation examples of this algorithm, see [Error Retries and Exponential Backoff in AWS](#).

## Eventual Consistency

The Amazon EC2 API follows an eventual consistency model, due to the distributed nature of the system supporting the API. This means that the result of an API command you run that affects your Amazon EC2 resources might not be immediately visible to all subsequent commands you run. You should keep this in mind when you carry out an API command that immediately follows a previous API command.

Eventual consistency can affect the way you manage your resources. For example, if you run a command to create a resource, it will eventually be visible to other commands. This means that if you run a command to modify or describe the resource that you just created, its ID might not have propagated throughout the system, and you will get an error responding that the resource does not exist.

To manage eventual consistency, you can do the following:

- Confirm the state of the resource before you run a command to modify it. Run the appropriate `Describe` command using an exponential backoff algorithm to ensure that you allow enough time for the previous command to propagate through the system. To do this, run the `Describe` command repeatedly, starting with a couple of seconds of wait time, and increasing gradually up to five minutes of wait time.
- Add wait time between subsequent commands, even if a `Describe` command returns an accurate response. Apply an exponential backoff algorithm starting with a couple of seconds of wait time, and increase gradually up to about five minutes of wait time.

### Eventual Consistency Error Examples

The following are examples of error codes you may encounter as a result of eventual consistency.

- `InvalidInstanceID.NotFound`

If you successfully run the `RunInstances` command, and then immediately run another command using the instance ID that was provided in the response of `RunInstances`, it may return an `InvalidInstanceID.NotFound` error. This does not mean the instance does not exist.

Some specific commands that may be affected are:

- `DescribeInstances`: To confirm the actual state of the instance, run this command using an exponential back-off algorithm.
- `TerminateInstances`: To confirm the state of the instance, first run the `DescribeInstances` command using an exponential back-off algorithm.

**Important**

If you get an `InvalidInstanceID.NotFound` error after running `TerminateInstances`, this does not mean that the instance is or will be terminated. Your instance could still be running. This is why it is important to first confirm the instance's state using `DescribeInstances`.

- `InvalidGroup.NotFound`

If you successfully run the `CreateSecurityGroup` command, and then immediately run another command using the instance ID that was provided in the response of `CreateSecurityGroup`, it may return an `InvalidGroup.NotFound` error. To confirm the state of the security group, run the `DescribeSecurityGroups` command using an exponential back-off algorithm.

## Ensuring Idempotency

An *idempotent* operation completes no more than one time.

When you launch an instance, the request typically returns before the operation has completed. You determine whether the operation was successful by monitoring the state of the instance (it goes from `pending` to `running`). If the operation times out or there are connection issues, you might need to retry the request. However, if the original request and a retry both successful, you'll end up with more instances than you intended to launch.

If you launch you instance using the `ec2-run-instances` command or the `RunInstances` API action, you can optionally provide a client token to ensure that the request is idempotent. If you repeat a request, the same response is returned for each repeated request. The only information that might vary in the response is the state of the instance.

The client token is a unique, case-sensitive string of up to 64 ASCII characters. It is included in the response when you describe the instance. The client token is valid for at least 24 hours after the termination of the instance. You should not reuse a client token in another call later on.

If you repeat a request with the same client token, but change another request parameter, Amazon EC2 returns an `IdempotentParameterMismatch` error.

You can use the same client token for the same request across different regions. For example, if you send an idempotent request to launch an instance in the `us-east-1` region, and then use the same client token in a request in other regions, we'll launch instances in each of those regions.

The following table shows common response codes and the recommended course of action.

Code	Retry	Comments
200 (OK)	No effect	The request has succeeded and any further retries have no effect.



Code	Retry	Comments
400 (Client Error)	Not recommended	The request will never succeed (for example, a specified parameter value is not valid). If the request involves a resource that is in the process of changing states, repeating the request could possibly succeed (for example, launching an instance using an Amazon EBS volume that is about to become <code>available</code> ).
500 (Server Internal Error)	Recommended	The error is generally transient. Repeat the request with an appropriate back-off strategy.
503 (Server Unavailable)	Recommended	The error can occur when there is extreme load. Repeat the request with an appropriate back-off strategy.

## Idempotency Support

The following commands and actions support idempotent operations using a client token:

- CopyImage
- CreateReservedInstancesListing
- `ec2-copy-image`
- `ec2-create-reserved-instances-listing`
- `ec2-modify-reserved-instances`
- `ec2-run-instances`
- ModifyReservedInstances
- RunInstances

The following commands and actions are idempotent.

- AssociateAddress
- DisassociateAddress
- `ec2-associate-address`
- `ec2-disassociate-address`
- `ec2-terminate-instances`
- TerminateInstances

## Example Idempotent Command

Use the `ec2-run-instances` command as follows to make an idempotent request:

```
PROMPT> ec2-run-instances ami-b232d0db -k my-key-pair --client-token 550e8400-  
e29b-41d4-a716-446655440000
```

The `--client-token` option requires a unique, case-sensitive string of up to 64 ASCII characters.

## Example Idempotent Query

Use the [RunInstances](#) action as follows to make an idempotent request:

```
https://ec2.amazonaws.com/?Action=RunInstances
&ImageId=ami-3ac33653
&MaxCount=1
&MinCount=1
&KeyName=my-key-pair
&ClientToken=550e8400-e29b-41d4-a716-446655440000
&AUTHPARAMS
```

The *ClientToken* parameter requires a unique, case-sensitive string of up to 64 ASCII characters.

## SOAP Requests

We have deprecated the SOAP API for Amazon EC2, and will end support for it in 2013.

# Document History

The following table describes important additions to the Amazon EC2 documentation set. We also update the documentation frequently to address the feedback that you send us.

**Current API version: 2013-10-01.**

Feature	API Version	Description	Release Date
Launching an instance from the AWS Marketplace	2013-10-01	Added information about launching an instance from the AWS Marketplace using the Amazon EC2 launch wizard. For more information, see <a href="#">Launching an AWS Marketplace Instance (p. 271)</a> .	31 October 2013
New launch wizard	2013-10-01	Added information about the redesigned EC2 launch wizard. For more information, see <a href="#">Launching an Instance (p. 266)</a> .	10 October 2013
Modifying Instance Types of Amazon EC2 Reserved Instances	2013-10-01	Added information about a new Reserved Instances feature for modifying the instance type of Linux/UNIX and Windows (without SQL Server) Reserved Instances within the same family (e.g., m1, m2, m3, c1). For more information, see <a href="#">Modifying Your Reserved Instances (p. 227)</a> .	09 October 2013
New Amazon Linux AMI release	2013-09-30	Added information about the release of Amazon Linux AMI 2013.09. For more information, see <a href="#">Amazon Linux (p. 87)</a> .	30 September 2013
Modifying Amazon EC2 Reserved Instances	2013-08-15	Added information about a new EC2 Reserved Instances feature for modifying Reserved Instances in a region. For more information, see <a href="#">Modifying Your Reserved Instances (p. 227)</a> .	11 September 2013
Assigning a public IP address	2013-07-15	Added information about a new public IP addressing feature for launching instances in a VPC. For more information, see <a href="#">Assigning a Public IP Address (p. 423)</a> .	20 August 2013

Feature	API Version	Description	Release Date
Granting resource-level permissions	2013-06-15	Added information about new Amazon Resource Names (ARNs) and condition keys for Amazon EC2. For more information, see <a href="#">IAM Policies for Amazon EC2 (p. 401)</a> .	8 July 2013
Incremental Snapshot Copies	2013-02-01	Added information about incremental snapshot copies. For more information, see <a href="#">Copying an Amazon EBS Snapshot (p. 488)</a> .	11 June 2013
New Tags pane	2013-02-01	Added information about the new Tags pane in the Amazon EC2 console. For more information, see <a href="#">Tagging Your Amazon EC2 Resources (p. 532)</a> .	04 April 2013
New Amazon Linux AMI release	2013-02-01	Added information about the release of Amazon Linux AMI 2013.03. For more information, see <a href="#">Amazon Linux (p. 87)</a> .	27 March 2013
Support for additional EBS-optimized instance types	2013-02-01	The following instance types can now be launched as EBS-optimized instances: High-CPU Extra Large (c1.xlarge), High-Memory Double Extra Large (m2.2xlarge), M3 Extra Large (m3.xlarge), and M3 Double Extra Large (m3.2xlarge).  For more information, see <a href="#">EBS-Optimized Instances (p. 107)</a> .	19 March 2013
Copy an AMI from one region to another	2013-02-01	You can copy an AMI from one region to another, enabling you to launch consistent instances in more than one AWS region quickly and easily.  For more information, see <a href="#">Copying AMIs (p. 84)</a> .	11 March 2013
Launch instances into a default VPC	2013-02-01	Your AWS account is capable of launching instances into either the EC2-Classic or EC2-VPC platform, or only into the EC2-VPC platform, on a region-by-region basis. If you can launch instances only into EC2-VPC, we create a default VPC for you. When you launch an instance, we launch it into your default VPC, unless you create a nondefault VPC and specify it when you launch the instance.  For more information, see <a href="#">Supported Platforms (p. 417)</a> .	11 March 2013
High-memory cluster (cr1.8xlarge) instance type	2012-12-01	Have large amounts of memory coupled with high CPU and network performance. These instances are well suited for in-memory analytics, graph analysis, and scientific computing applications. For more information, see <a href="#">Instance Types (p. 94)</a> .	21 January 2013
High storage (hs1.8xlarge) instance type	2012-12-01	High storage instances provide very high storage density and high sequential read and write performance per instance. They are well-suited for data warehousing, Hadoop/MapReduce, and parallel file systems. For more information, see <a href="#">HS1 Instances (p. 104)</a> .	20 December 2012

Feature	API Version	Description	Release Date
EBS snapshot copy	2012-12-01	You can use snapshot copies to create backups of data, to create new Amazon EBS volumes, or to create Amazon Machine Images (AMIs). For more information, see <a href="#">Copying an Amazon EBS Snapshot (p. 488)</a> .	17 December 2012
Updated EBS metrics and status checks for Provisioned IOPS volumes	2012-10-01	Updated the EBS metrics to include two new metrics for Provisioned IOPS volumes. For more information, see <a href="#">Monitoring Volumes with CloudWatch (p. 464)</a> . Also added new status checks for Provisioned IOPS volumes. For more information, see <a href="#">Monitoring Volumes with Status Checks (p. 466)</a> .	20 November 2012
Support for Microsoft Windows Server 2012	2012-10-01	Amazon EC2 now provides you with several pre-configured Windows Server 2012 AMIs. These AMIs are immediately available for use in every region and for every 64-bit instance type. The AMIs support the following languages: <ul style="list-style-type: none"> <li>• English</li> <li>• Chinese Simplified</li> <li>• Chinese Traditional</li> <li>• Chinese Traditional Hong Kong</li> <li>• Japanese</li> <li>• Korean</li> <li>• Portuguese</li> <li>• Portuguese Brazil</li> <li>• Czech</li> <li>• Dutch</li> <li>• French</li> <li>• German</li> <li>• Hungarian</li> <li>• Italian</li> <li>• Polish</li> <li>• Russian</li> <li>• Spanish</li> <li>• Swedish</li> <li>• Turkish</li> </ul>	19 November 2012
Linux Kernels	2012-10-01	Updated AKI IDs; reorganized distribution kernels; updated PVOps section.	13 November 2012
Amazon Linux AMIs	2012-10-01	Added information to Adding Packages and the location of the latest Linux AMIs. For more information, see <a href="#">Adding Packages (p. 91)</a> .	02 November 2012

Feature	API Version	Description	Release Date
M3 Instances	2012-10-01	Added information about the new Amazon Elastic Compute Cloud (Amazon EC2) M3 extra-large and M3 double-extra-large instance types. For more information, see <a href="#">Instance Types (p. 94)</a> .	31 October 2012
Amazon EC2 Spot Instance Request Status	2012-10-01	Added information about Amazon EC2 Spot Instance request status, which makes it easy to determine the state of your Amazon EC2 Spot requests.	14 October 2012
New Amazon Linux AMI release	2012-08-15	Added information about the release of Amazon Linux AMI 2012.09. For more information, see <a href="#">Amazon Linux (p. 87)</a> .	11 October 2012
Amazon EC2 Reserved Instance Marketplace	2012-08-15	The Reserved Instance Marketplace matches sellers who have Amazon EC2 Reserved Instances that they no longer need with buyers who are looking to purchase additional capacity. Reserved Instances bought and sold through the Reserved Instance Marketplace work like any other Reserved Instances, except that they can have less than a full standard term remaining and can be sold at different prices.	11 September 2012
Provisioned IOPS (input/output operations per second) for Amazon EBS	2012-07-20	Provisioned IOPS volumes deliver predictable, high performance for I/O intensive workloads, such as database applications, that rely on consistent and fast response times. For more information, see <a href="#">Amazon EBS Volume Types (p. 448)</a> .	31 July 2012
High I/O instances for Amazon EC2	2012-06-15	High I/O instances provides very high, low latency, disk I/O performance using SSD-based local instance storage. For more information, see <a href="#">H1 Instances (p. 102)</a> .	18 July 2012
IAM roles on Amazon EC2 instances	2012-06-01	IAM roles for Amazon EC2 provide: <ul style="list-style-type: none"> <li>• AWS access keys for applications running on Amazon EC2 instances.</li> <li>• Automatic rotation of the AWS access keys on the Amazon EC2 instance.</li> <li>• Granular permissions for applications running on Amazon EC2 instances that make requests to your AWS services.</li> </ul>	11 June 2012

Feature	API Version	Description	Release Date
Spot Instance features that make it easier to get started and handle the potential of interruption.		Using Auto Scaling, you can now manage your Spot Instances: <ul style="list-style-type: none"> <li>Place bids for Amazon EC2 Spot Instances using Auto Scaling launch configurations, and set up a schedule for placing bids for Spot Instances. For more information, see <a href="#">Managing Spot Instances with Auto Scaling (p. 143)</a>.</li> <li>Get notifications when instances are launched or terminated. For more information, see <a href="#">Using Auto Scaling to Get Notifications for Spot Instances (p. 156)</a>.</li> <li>Use AWS CloudFormation templates to launch Spot Instances in a stack with AWS resources. For more information, see <a href="#">Using CloudFormation Templates to Launch Spot Instances (p. 158)</a>.</li> </ul>	7 June 2012
EC2 instance export and timestamps for status checks for Amazon EC2	2012-05-01	Added support for exporting Windows Server instances that you originally imported into EC2. Added support for timestamps on instance status and system status to indicate the date and time that a status check failed.	25 May 2012
EC2 instance export, and timestamps in instance and system status checks for Amazon VPC	2012-05-01	Added support for EC2 instance export to Citrix Xen, Microsoft Hyper-V, and VMware vSphere. Added support for timestamps in instance and system status checks.	25 May 2012
Cluster Compute Eight Extra Large instances	2012-04-01	Added support for cc2.8xlarge instances in Amazon Virtual Private Cloud (Amazon VPC).	26 April 2012
AWS Marketplace AMIs	2012-04-01	Added support for AWS Marketplace AMIs.	19 April 2012
New Linux AMI release		Added information about the release of Amazon Linux AMI 2012.03. For more information, see <a href="#">Amazon Linux (p. 87)</a> .	28 March 2012
New AKI version		Added information about the release of new AKI version 1.03 and the release of AKIs for the AWS GovCloud (US) region. For more information, see <a href="#">Using Your Own Linux Kernels (p. 80)</a> .	28 March 2012
Medium instances, support for 64-bit on all AMIs, and a Java-based SSH Client		Added support for a new instance type and 64-bit information. Added procedures for using the Java-based SSH client to connect to Linux/UNIX instances.	7 March 2012

Feature	API Version	Description	Release Date
Reserved Instance pricing tiers		Added a new section discussing how to take advantage of the discount pricing that is built into the Reserved Instance pricing tiers. For more information, see <a href="#">Understanding Reserved Instance Pricing Tiers (p. 200)</a> .	5 March 2012
Elastic Network Interfaces (ENIs) for EC2 instances in Amazon Virtual Private Cloud		Added new section about elastic network interfaces (ENIs) for EC2 instances in a VPC. For more information, see <a href="#">Elastic Network Interfaces (ENI) (p. 431)</a> .	21 December 2011
New GRU Region and AKIs		Added information about the release of new AKIs for the SA-East-1 Region. This release deprecates the AKI version 1.01. AKI version 1.02 will continue to be backward compatible.	14 December 2011
New offering types for Amazon EC2 Reserved Instances	2011-11-01	You can choose from a variety of Reserved Instance offerings that address your projected use of the instance: <i>Heavy Utilization</i> , <i>Medium Utilization</i> , and <i>Light Utilization</i> . See <a href="#">Reserved Instances (p. 193)</a> .	01 December 2011
Amazon EC2 instance status		You can view additional details about the status of your instances, including scheduled events planned by AWS that might have an impact on your instances. These operational activities include instance reboots required to apply software updates or security patches, or instance retirements required where there are hardware issues. See <a href="#">Monitoring the Status of Your Instances (p. 346)</a> .	16 November 2011
Amazon EC2 Cluster Compute Instance Type		Added support for Cluster Compute Eight Extra Large (cc2.8xlarge) to Amazon EC2.	14 November 2011
New PDX Region and AKIs		Added information about the release of new AKIs for the new US-West 2 Region.	8 November 2011
Amazon EC2 Spot Instances in Amazon VPC		Added information about the support for Amazon EC2 Spot Instances in Amazon VPC. With this update, users will be able to launch Spot Instances in the Amazon Virtual Private Cloud (Amazon VPC). By launching Spot Instances in Amazon VPC, users of Spot Instances can enjoy all of the controls and advanced security options of Amazon VPC. For more information, see <a href="#">Launching Spot Instances in Amazon Virtual Private Cloud (p. 160)</a> .	11 October 2011



Feature	API Version	Description	Release Date
New Linux AMI release		Added information about the release of Amazon Linux AMI 2011.09. This update removes the beta tag from the Amazon Linux AMI, supports the ability to lock the repositories to a specific version, and provides for notification when updates are available to installed packages including security updates. For more information, see <a href="#">Amazon Linux (p. 87)</a> .	26 September 2011
Simplified VM import process for users of the CLI tools		The VM Import process for CLI users is simplified with the enhanced functionality of <code>ec2-import-instance</code> and <code>ec2-import-volume</code> , which now will perform the upload of the images into Amazon EC2 after creating the import task. In addition, with the introduction of the <code>ec2-resume-import</code> command, users can restart an incomplete upload at the point the task stopped. For more information, see <a href="#">Importing Your Virtual Machine into Amazon EC2 (p. 324)</a> .	15 September 2011
Support for importing in VHD file format		VM Import can now import virtual machine image files in VHD format. The VHD file format is compatible with the Citrix Xen and Microsoft Hyper-V virtualization platforms. With this release, VM Import now supports RAW, VHD and VMDK (VMware ESX-compatible) image formats. For more information, see <a href="#">Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2 (p. 318)</a> .	24 August 2011
Support for Microsoft Windows Server 2003 R2		VM Import now supports Windows Server 2003 (R2). With this release, VM Import supports all versions of Microsoft Windows Server supported by Amazon EC2.	24 August 2011
Update to the Amazon EC2 VM Import Connector for VMware vCenter		Added information about the 1.1 version of the Amazon EC2 VM Import Connector for VMware vCenter virtual appliance (Connector). This update includes proxy support for Internet access, better error handling, improved task progress bar accuracy, and several bug fixes. For more information, see <a href="#">Importing EC2 Instances (p. 301)</a> .	27 June 2011
Enabling Linux AMI to run user-provided kernels		Added information about the AKI version change from 1.01 to 1.02. This version updates the PVGRUB to address launch failures associated with t1.micro Linux instances. For more information, go to <a href="#">Using Your Own Linux Kernels (p. 80)</a> .	20 June 2011

Feature	API Version	Description	Release Date
Spot Instances Availability Zone pricing changes		Added information about the Spot Instances Availability Zone pricing feature. In this release, we've added new Availability Zone pricing options as part of the information returned when you query for Spot Instance requests and Spot Price history. These additions make it easier to determine the price required to launch a Spot Instance into a particular Availability Zone. For more information, see <a href="#">Spot Instances (p. 115)</a> .	26 May 2011
AWS Identity and Access Management		Added information about AWS Identity and Access Management (IAM), which enables users to specify which Amazon EC2 actions a user can use with Amazon EC2 resources in general. For more information, see <a href="#">Controlling Access to Amazon EC2 Resources (p. 399)</a> .	26 April 2011
Enabling Linux AMI to run user-provided kernels		Added information about enabling a Linux AMI to use PVGRUB Amazon Kernel Image (AKI) to run a user-provided kernel. For more information, go to <a href="#">Using Your Own Linux Kernels (p. 80)</a> .	26 April 2011
Dedicated instances		Launched within your Amazon Virtual Private Cloud (Amazon VPC), Dedicated Instances are instances that are physically isolated at the host hardware level. Dedicated Instances let you take advantage of Amazon VPC and the AWS cloud, with benefits including on-demand elastic provisioning and pay only for what you use, while isolating your Amazon EC2 compute instances at the hardware level. For more information, see <a href="#">Using EC2 Dedicated Instances</a> in the <i>Amazon Virtual Private Cloud User Guide</i> .	27 March 2011
Reserved Instances updates to the AWS Management Console		Updates to the AWS Management Console make it easier for users to view their Reserved Instances and purchase additional Reserved Instances, including Dedicated Reserved Instances. For more information, see <a href="#">Reserved Instances (p. 193)</a> .	27 March 2011
Support for Windows Server 2008 R2		Amazon EC2 now provides you with several pre-configured Windows Server 2008 R2 AMIs. These AMIs are immediately available for use in every region and in most 64-bit instance types, excluding t1.micro and HPC families. The AMIs will support multiple languages.	15 March 2011
New Amazon Linux reference AMI		Added information about the new Amazon Linux reference AMI, which replaces the CentOS reference AMI. Removed information about the CentOS reference AMI, including the section named Correcting Clock Drift for Cluster Instances on CentOS 5.4 AMI. For more information, see <a href="#">AMIs for GPU Instances (p. 106)</a> .	15 March 2011

Feature	API Version	Description	Release Date
Metadata information		Added information about metadata to reflect changes in the 2011-01-01 release. For more information, see <a href="#">Instance Metadata and User Data (p. 290)</a> and <a href="#">Instance Metadata Categories (p. 297)</a> .	11 March 2011
Amazon EC2 VM Import Connector for VMware vCenter		Added information about the Amazon EC2 VM Import Connector for VMware vCenter virtual appliance (Connector). The Connector is a plug-in for VMware vCenter that integrates with VMware vSphere Client and provides a graphical user interface that you can use to import your VMware virtual machines to Amazon EC2. For more information, see <a href="#">Importing EC2 Instances (p. 301)</a> .	3 March 2011
Force volume detachment		You can now use the AWS Management Console to force the detachment of an Amazon EBS volume from an instance. For more information, see <a href="#">Detaching an Amazon EBS Volume from an Instance (p. 475)</a> .	23 February 2011
Instance termination protection		You can now use the AWS Management Console to prevent an instance from being terminated. For more information, see <a href="#">Enabling Termination Protection for an Instance (p. 288)</a> .	23 February 2011
Correcting Clock Drift for Cluster Instances on CentOS 5.4 AMI		Added information about how to correct clock drift for cluster instances running on Amazon's CentOS 5.4 AMI.	25 January 2011
Restructured Documentation		Implemented the following updates: <ul style="list-style-type: none"> <li>• Consolidated all developer guide information into this user guide.</li> <li>• Restructured and updated the information about creating AMIs. For more information, see <a href="#">Creating Your Own AMIs (p. 62)</a>.</li> <li>• Updated the introduction to the user guide. For more information, see <a href="#">What is Amazon EC2? (p. 1)</a></li> </ul>	14 January 2011
VM Import		Added information about VM Import, which allows you to import a virtual machine or volume into Amazon EC2. For more information, see <a href="#">Using the Command Line Tools to Import Your Virtual Machine to Amazon EC2 (p. 318)</a> .	15 December 2010
Basic monitoring for instances		Added information about basic monitoring for EC2 instances. For more information, see <a href="#">Monitoring Instances (p. 336)</a> .	12 December 2010

Feature	API Version	Description	Release Date
Cluster GPU instances		Amazon EC2 offers cluster GPU instances (cg1.4xlarge) for high-performance computing (HPC) applications. For more information about the specifications for each Amazon EC2 instance type, see <a href="#">Instance Type Details</a> .	14 November 2010
Filters and Tags		Added information about listing, filtering, and tagging resources. For more information, see <a href="#">Listing and Filtering Your Resources (p. 529)</a> and <a href="#">Tagging Your Amazon EC2 Resources (p. 532)</a> .	19 September 2010
Idempotent Instance Launch		Added information about ensuring idempotency when running instances. For more information, see <a href="#">Ensuring Idempotency (p. 557)</a> .	19 September 2010
Micro instances		Amazon EC2 offers the t1.micro instance type for certain types of applications. For more information, see <a href="#">Micro Instances (p. 95)</a> .	8 September 2010
AWS Identity and Access Management for Amazon EC2		Amazon EC2 now integrates with AWS Identity and Access Management (IAM). For more information, see <a href="#">Controlling Access to Amazon EC2 Resources (p. 399)</a> .	2 September 2010
Cluster instances		Amazon EC2 offers cluster compute instances for high-performance computing (HPC) applications. For more information about the specifications for each Amazon EC2 instance type, see <a href="#">Instance Type Details</a> .	12 July 2010
Amazon VPC IP Address Designation		Amazon VPC users can now specify the IP address to assign an instance launched in a VPC.	12 July 2010
Amazon CloudWatch Monitoring for Amazon EBS Volumes		Amazon CloudWatch monitoring is now automatically available for Amazon EBS volumes. For more information, see <a href="#">Monitoring Volumes with CloudWatch (p. 464)</a> .	14 June 2010
High-memory extra large instances		Amazon EC2 now supports a High-Memory Extra Large (m2.xlarge) instance type. For more information about the specifications for each Amazon EC2 instance type, see <a href="#">Instance Type Details</a> .	22 February 2010
Reserved Instances with Windows		Amazon EC2 now supports Reserved Instances with Windows. For more information about Reserved Instances, see <a href="#">Reserved Instances (p. 193)</a> .	22 February 2010