


# Amazon Elastic Cloud Compute (EC2): Amazon EC2 Overview



**Andru Estes**

Principal Author

 andru-estes





# **Amazon EC2 and AMIs**





## Amazon EC2

On-demand, scalable computing service offered  
by AWS for running your very own VMs.  
Get what you need, when you need it.  
*Infrastructure as a Service (IaaS)*



**Amazon EC2 allows you to  
wait minutes, not months!**




## **Why Use EC2?**

**You only pay for exactly what you use while its running**

**You avoid wasted capacity by using the ability to grow and shrink as needed**




# EC2 Instance Basics




## Compute

The amount of virtual CPUs




## Memory

The amount of RAM



## Network

Bandwidth capabilities



## Storage

Either virtual hard drive, temporary, or network storage



# Amazon EC2 Key Pairs



A set of Regional security credentials used to prove your identity when connecting to an Amazon EC2 instance

Comprised of a public key and private key

The public is stored on the EC2 instance, and you download and use the private key for connecting

Use these key pairs to:

- SSH into Linux instances
- Decrypt the administrator password for Windows instances (used to RDP)



# Amazon Machine Image (AMI)

**“An Amazon Machine Image (AMI) is an image provided by AWS that provides the information required to launch an instance.”**

**You must select an AMI before launching any EC2 instance!**

---

Citation: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html>



# Types of AMIs

## Public

Anyone with an AWS Account can explore and launch these as needed

## Private

Create your very own AMI and keep private, or choose to share it with other AWS Accounts



# What to Know About AMIs



**Prebaking:** Embedding a significant portion of your application artifacts within an AMI so it is loaded at instance launch and ready to go, faster



**Use Case #1:** Create a “golden” image with all necessary security patches updates to deploy monthly



**Use Case #2:** Standardize deployments to ensure each EC2 instance has the exact same installations and custom software configurations





# **Amazon EC2 Sizes and Instance Types**



# Choosing the Correct EC2 Instance

## Class and Generation

The type of instance and the generation of the instance type

## Instance Size

How big or small you need your resources to be



# Breaking Down Instance Type Nomenclature

**t4g.large**



# Breaking Down Instance Type Nomenclature

**t4g.large**  
↑  
T-family



# Breaking Down Instance Type Nomenclature

t4g.large



Generation of the instance



# Breaking Down Instance Type Nomenclature

t4g.large



Size of the instance




**There are several different categories of instance types...**




# Types of EC2 Instances - Part 1



**General Purpose:** Balance of compute, memory, and networking. T-family instances are burstable. Primarily M & T-type instances.



**Compute Optimized:** Perfect for heavily compute-bound workloads. High-performing processors. Modeling and batch processing. C-type.



**Memory Optimized:** Primarily used for workloads of processing large data sets in memory, like caching. High memory per CPU ratio. R, X, & Z-type.



# Types of EC2 Instances - Part 2



**Accelerated Computing:** Hardware acceleration. Useful for video transcoding and graphics rendering. Incorporate GPUs for performance boost. Primarily P and G instances.



**Storage Optimized:** Used with workloads that need high, sequential read and write access to local storage. Low-latency, and best IOPs for apps. Look for the I (eye) instance family.



**HPC Optimized:** Best price performance for running HPC workloads at scale. Genomic projects, simulations, or ML. HPC-instance family.



# Summary

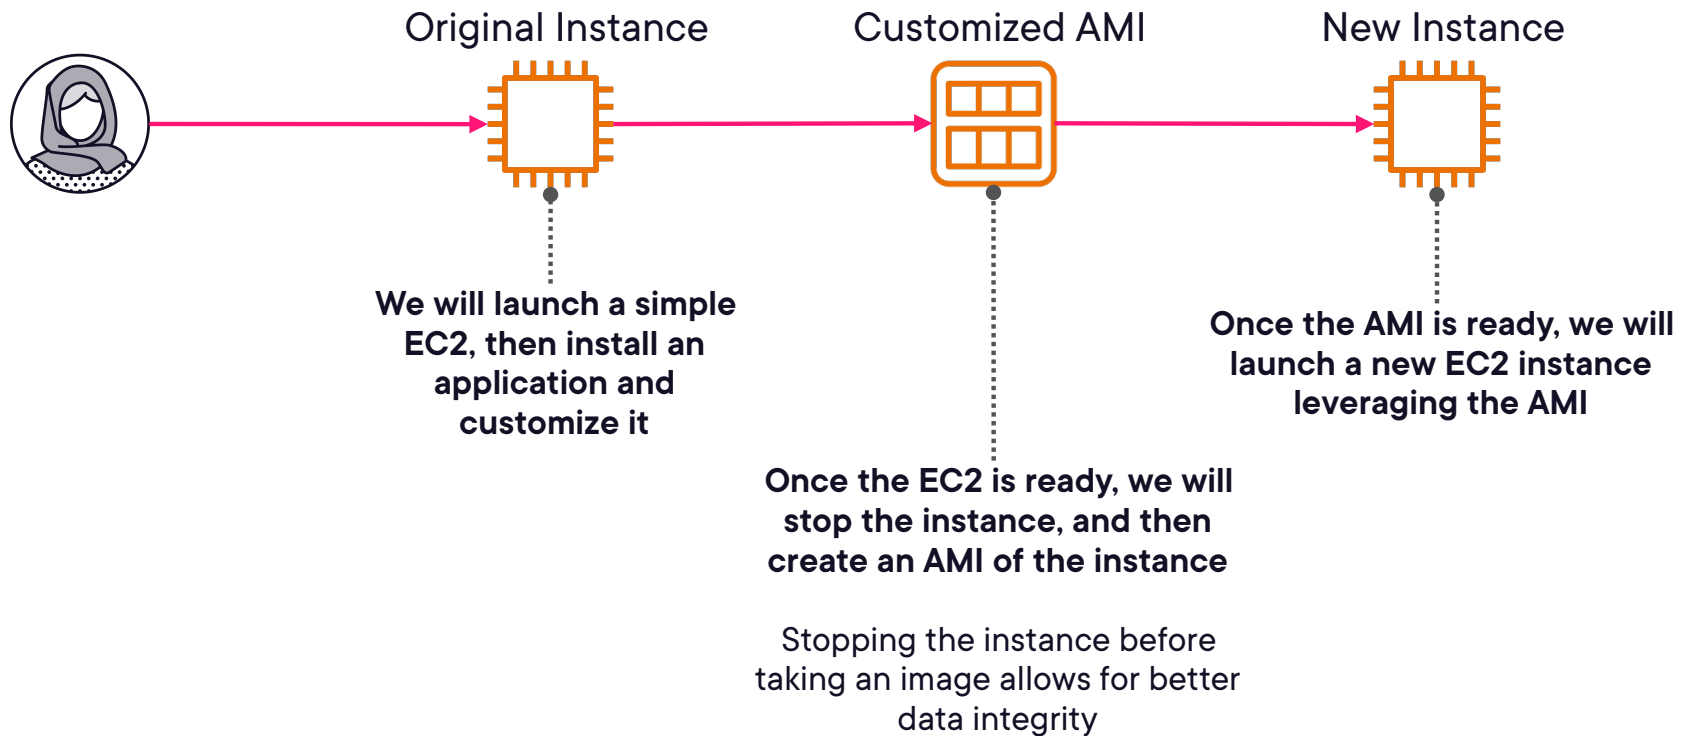
**Choose the correct instance type and family for your use case**

**Scale power up or down based on the size of you instances**

**The bigger an instance, the more expensive the instance will be**



# Demo: Launching an EC2 Instance and Creating an AMI





# Amazon EC2 User Data

<https://t.me/learningnets>



**User data is passed to an instance to perform automated configuration tasks and run scripts right after the instance starts.**



**This process is commonly referred to as bootstrapping.**



# User Data Formats

**Simple plain text**

**Base64-encoded text**



# User Data Format Examples - Plain Text

```
#!/bin/bash
```

```
yum update -y  
yum install -y httpd
```

```
systemctl start httpd.service  
systemctl enable httpd.service
```

```
systemctl restart httpd.service
```





**Both of those examples do  
the exact same thing!**



# User Data Use Cases



**Bootstrapping Applications:** You want to deploy a web application on an EC2 instance and you need to install and configure the application automatically during boot.



**Configuring Instance Settings:** You want to configure instance settings, such as the time zone, locales, and keyboard layout on a fleet of EC2 instances.

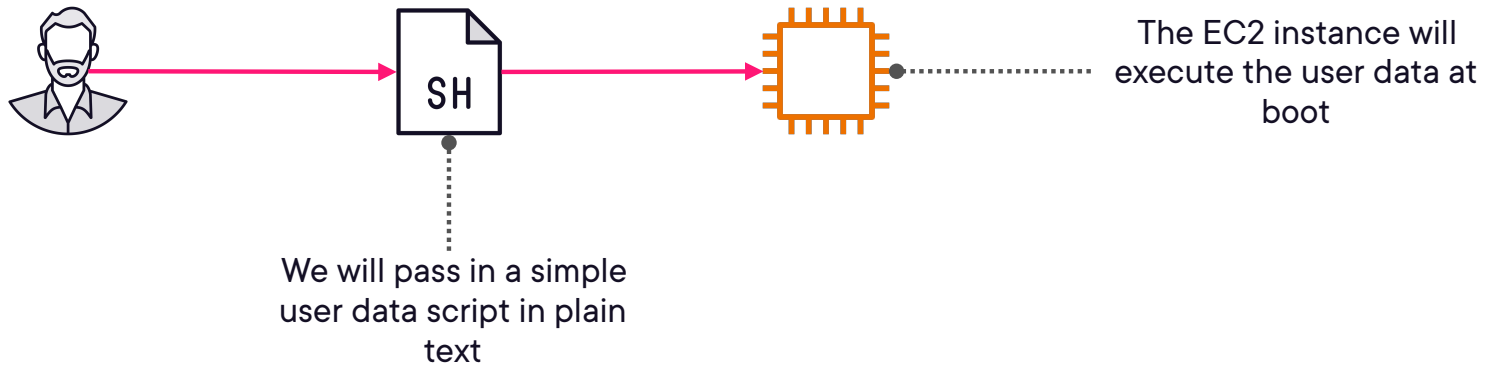


**Joining to a Domain:** You can use EC2 user data to pass a script that joins the instance to the domain and configures the necessary settings.

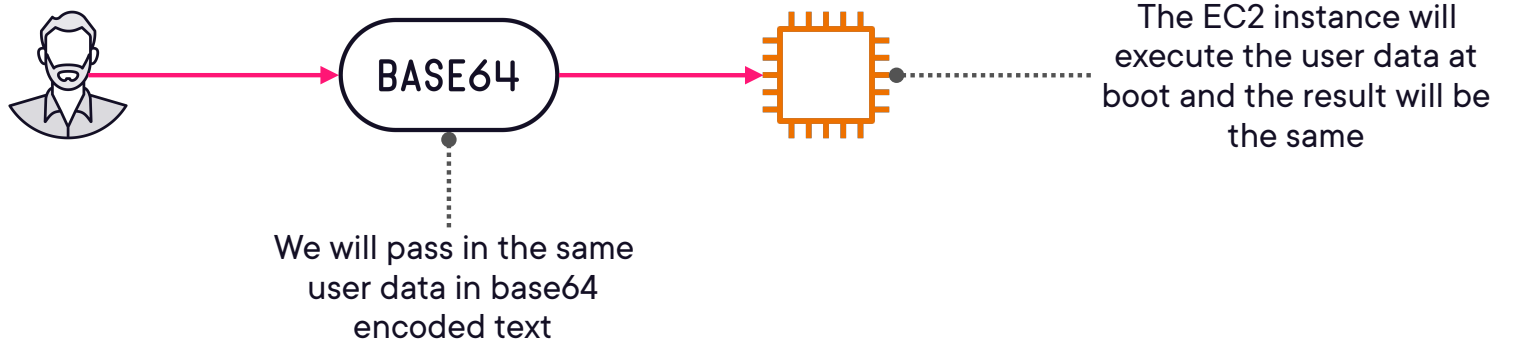


# Demo: Passing in EC2 User Data

1



2





# **EC2 Hibernate**





## EC2 Hibernate

Feature of EC2 that signals the operating system to perform hibernation, aka **suspend-to-disk**

Works by saving the content in the RAM to an attached long-term volume (EBS)

Useful for quick recovery failures where an application takes a long time to boot up, you can use this to pre-warm the instance

Source: <https://unsplash.com/>



# Starting After Hibernation

1<sup>ST</sup>

EBS root volume is restored to its previous state

2<sup>ND</sup>

RAM contents are reloaded

3<sup>RD</sup>

Processes that were previously running on the instance are resumed

4<sup>TH</sup>

Previously attached data volumes are reattached and the instance retains its instance ID





# **Module Summary and Exam Tips**



**EC2 instances are a VM  
hosted in AWS instead of your  
own data center.**



**Select the capacity you need  
right now, grow and shrink on  
demand, and only pay for  
what you use!**



# Amazon EC2 Key Pairs

**Private and public key pair used to connect to your EC2 instances**

**SSH uses the private key directly to connect**

**RDP uses the private key to decrypt the admin password**



# EC2 User Data Review

User data allows you to pass in a bootstrap script to your instance

A bootstrap script is a script that runs when the instance first runs

Useful for installing applications, as well as updates and more




**Remember that there are numerous EC2 instance family types to choose from for your use cases!**




# EC2 Hibernate Review



**EC2 feature for instances that performs a suspend-to-disk task**



**The suspend-to-disk process saves data living in RAM to an attached long-term EBS volume**



**This is very useful for quick recovery failures of an application that can take a long time to boot up**

