



- [About AWS](#) ▾
- [Products](#) ▾
- [Solutions](#) ▾
- [Resources](#) ▾
- [Support](#) ▾
- [Your Account](#) ▾

Infrastructure Services

- [Amazon Elastic Compute Cloud \(Amazon EC2\)](#)
- [Amazon SimpleDB](#)
- [Amazon Simple Storage Service \(Amazon S3\)](#)
- [Amazon Simple Queue Service \(Amazon SQS\)](#)
- [AWS Premium Support](#)

Payments & Billing

On-Demand Workforce

Web Search & Information

Amazon Fulfillment & Associates

[Home](#) > [Products](#) > [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) > [Amazon EC2 Instance Types](#)

Amazon EC2 Instance Types

[Amazon Elastic Compute Cloud \(Amazon EC2\)](#) provides the flexibility to choose from a number of different instance types to meet your computing needs. Each instance provides a predictable amount of dedicated compute capacity and is charged per instance-hour consumed.

Available Instance Types

Standard Instances

Instances of this family are well suited for most applications.

Small Instance (default)*

1.7 GB memory
1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit)
160 GB instance storage (150 GB plus 10 GB root partition)
32-bit platform
I/O Performance: Moderate
Price: \$0.10 per instance hour

Large Instance

7.5 GB memory
4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each)
850 GB instance storage (2×420 GB plus 10 GB root partition)
64-bit platform
I/O Performance: High
Price: \$0.40 per instance hour

Extra Large Instance

15 GB memory 8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each)
1,690 GB instance storage (4×420 GB plus 10 GB root partition)
64-bit platform
I/O Performance: High
Price: \$0.80 per instance hour

High-CPU Instances

Instances of this family have proportionally more CPU resources than memory (RAM) and are well suited for compute-intensive applications.

High-CPU Medium Instance

1.7 GB of memory
5 EC2 Compute Units (2 virtual cores with 2.5 EC2 Compute Units each)
350 GB of instance storage
32-bit platform
I/O Performance: Moderate
Price: \$0.20 per instance hour

High-CPU Extra Large Instance

7 GB of memory
20 EC2 Compute Units (8 virtual cores with 2.5 EC2 Compute Units each)

1690 GB of instance storage
64-bit platform
I/O Performance: High
Price: \$0.80 per instance hour

Selecting Instance Types

Amazon EC2 instances are grouped into two families: Standard and High-CPU. Standard Instances have memory to CPU ratios suitable for most general purpose applications; High-CPU instances have proportionally more CPU resources than memory (RAM) and are well suited for compute-intensive applications. When choosing instance types, you should consider the characteristics of your application with regards to resource utilization and select the optimal instance family and size.

One of the advantages of EC2 is that you pay by the instance hour, which makes it convenient and inexpensive to test the performance of your application on different instance families and types. One good way to determine the most appropriate instance family and instance type is to launch test instances and benchmark your application.

Measuring Compute Resources

Transitioning to a utility computing model fundamentally changes how developers have been trained to think about CPU resources. Instead of purchasing or leasing a particular processor to use for several months or years, you are renting capacity by the hour. Because Amazon EC2 is built on commodity hardware, over time there may be several different types of physical hardware underlying EC2 instances. Our goal is to provide a consistent amount of CPU capacity no matter what the actual underlying hardware.

Amazon EC2 uses a variety of measures to provide each instance with a consistent and predictable amount of CPU capacity. In order to make it easy for developers to compare CPU capacity between different instance types, we have defined an Amazon EC2 Compute Unit. The amount of CPU that is allocated to a particular instance is expressed in terms of these EC2 Compute Units. We use several benchmarks and tests to manage the consistency and predictability of the performance of an EC2 Compute Unit. One EC2 Compute Unit provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor. This is also the equivalent to an early-2006 1.7 GHz Xeon processor referenced in our original documentation. Over time, we may add or substitute measures that go into the definition of an EC2 Compute Unit, if we find metrics that will give you a clearer picture of compute capacity.

To find out which instance will work best for your application, the best thing to do is to launch an instance and benchmark your own application. One of the advantages of EC2 is that you pay by the hour which makes it convenient and inexpensive to test the performance of your application on different instance types.

I/O Performance

Amazon EC2 provides virtualized server instances. While some resources like CPU, memory and instance storage are dedicated to a particular instance, other resources like the network and the disk subsystem are shared among instances. If each instance on a physical host tries to use as much of one of these shared resources as possible, each will receive an equal share of that resource. However, when a resource is under-utilized you will often be able to consume a higher share of that resource while it is available.

The different instance types will provide higher or lower minimum performance from the shared resources depending on their size. Each of the instance types has an I/O performance indicator (*moderate* or *high*). Instance types with high I/O performance have a larger allocation of shared resources. Allocating larger share of shared resources also reduces the variance of I/O performance. For many applications, *moderate* I/O performance is more than enough. However, for those applications requiring greater or more consistent I/O performance, you may want to consider instances with *high* I/O performance.

*The Small Instance type is equivalent to the original Amazon EC2 instance type that has been available since Amazon EC2 launched. This instance type is currently the default for all customers. To use other instance types, customers must specifically request them via the *RunInstances* API.