



# Amazon EC2 Instances Featuring AMD EPYC™ Processors Solution Brief



## Looking for Optimized Performance and Cost

Customers are constantly trying to find the right balance between performance and cost for their applications. The reality is that not all applications require the same level of compute performance. This need to right-size becomes more prevalent as businesses are asked to do more with flat or declining budgets. Here are two challenges customers frequently face:


### LIMITED OPTIONS

Workloads have different characteristics and infrastructure requirements. Limited choices mean that the customer may not be able to right-size their workloads to optimized performance and cost.

### UNDERUTILIZATION

Many workloads only use a fraction of the processor's maximum performance on a given server or cloud instance, so customers may be underutilizing the resources they have.

**AWS improvements** are driven by a never-ending torrent of customer requests that power the Virtuous Cycle:





## A Collaborative Solution

AWS listened to their customers, asked them questions, and noticed that many were having similar challenges of balancing performance and cost for their applications. Knowing these challenges needed to be addressed, AWS and AMD partnered to deliver new Amazon EC2 instances powered by AMD EPYC™ processors that would provide more choices for the customer. These instances allow customers to right-size their workloads while also providing customers 10% lower cost compute and memory compared to comparable instances.

With the AMD partnership, AWS now offers additional EC2 instances that add to the broadest and deepest portfolio of instances and enables customers to optimize both cost and performance for their workloads.

**AMD was founded in 1969 as a Silicon Valley start-up**, focused on leading-edge semiconductor products. Today, AMD develops high-performance computing and visualization products to solve some of the world's toughest and most interesting challenges.



## Benefits of AMD-Based EC2 Instances

### FLEXIBILITY AND CHOICE

EC2 instances featuring AMD EPYC™ processors provide additional choices to help optimize both cost and performance, and may offer a better fit for many workloads that do not fully utilize the compute resources. Such alternatives are available now in the EC2 memory optimized (R5a), general purpose (M5a), and general purpose burstable (T3a) instance families.

### OPTIMIZE THE SAVINGS


EC2 instances featuring AMD EPYC™ processors deliver a 10% lower cost compute and memory compared to comparable instances. Since many workloads utilize only a fraction of the processor's maximum performance, these instances offer a better fit for many workloads.

### SEAMLESS WORKLOAD MIGRATION

Easily migrate applications currently running on existing EC2 instances to the new AMD-based variants with little to no modification. These instances are available in the same sizes and offer application compatibility with the R5, M5, and T3 instances, so customers can start using them just like their other EC2 instances.

**AWS offers choices in a wide range of dimensions** including locations, EC2 instance types, Windows or Linux operating systems, relational and NoSQL database choices, development languages, and pricing/purchase models.

**Since its inception, prices have been reduced 67 times.** AWS studies usage patterns, identifies areas for improvement, and frequently deploys updates across the entire AWS Cloud to make AWS an increasingly better value over time.





## Key Features

### AMD EPYC™ PROCESSORS

Amazon EC2 instances now feature AMD EPYC™ 7000 series processors with an all core turbo clock speed of 2.5 GHz. The AMD-based instances provide additional options for customers and may offer a better fit for many workloads that do not fully utilize the compute resources.


### AWS NITRO SYSTEM

The AWS Nitro System delivers a rich collection of building blocks that offloads many of the traditional virtualization functions to a combination of dedicated hardware and lightweight hypervisor. By doing so, the AWS Nitro System enables high performance, high availability, and high security while also reducing virtualization overhead.

The AWS Nitro System delivers practically all of the compute and memory resources of the host hardware to the user's instances, which frees up additional memory for workloads, boosts performance, and lowers the \$/GiB costs.

### HIGH PERFORMANCE NETWORKING AND STORAGE

Next generation Elastic Network Adapter (ENA) and NVMe Express (NVMe) technology provide AMD-powered EC2 instances with high throughput, low latency bandwidth interfaces for networking, and Amazon Elastic Block Store (Amazon EBS).




## Meet the AMD EPYC™ Powered Instances

### AMAZON EC2 R5A AND R5AD INSTANCES

Amazon EC2 R5a and R5ad instances are memory optimized instances and feature AMD EPYC™ processors with an all core turbo clock speed of up to 2.5 GHz and offer a 10% cost savings over comparable instances. Choose between a selection of instance choices with different storage options (EBS or NVMe SSD), and instance sizes to optimize both cost and performance for specific workload needs.

### USE CASES

Memory-optimized R5a instances are well suited for applications such as high-performance databases, distributed web scale in-memory caches, mid-size in-memory databases, real time big data analytics, data mining, and other enterprise applications. In general, they have fast performance for workloads that process large data sets in memory.



**R5ad instances** have local storage, offering up to 3.6TB of NVMe-based SSDs and are beneficial for applications that need temporary storage of data for scratch space, temporary files, and caches.

## AMAZON EC2 M5A AND M5AD INSTANCES

Amazon EC2 M5a and M5ad instances are general purpose instances that offers a balance of compute, memory, and networking resources and feature AMD EPYC™ processors with an all core turbo clock speed of up to 2.5 GHz and offer a 10% cost savings over comparable instances. Choose between a selection of instance choices with different storage options (EBS or NVMe SSD), and instance sizes to optimize both cost and performance for every type of workload.

### USE CASES

M5a and M5ad instances are ideal for a broad range of workloads, including business critical applications, web and application servers, gaming servers, caching fleets, app development environments, small and mid-size databases, data processing tasks that require additional memory, as well as running backend servers for SAP, Microsoft SharePoint, cluster computing, and other enterprise applications.

**M5ad instances** have local storage, offering up to 3.6TB of NVMe-based SSDs and are beneficial for applications that need temporary storage of data for scratch space, temporary files, and caches.

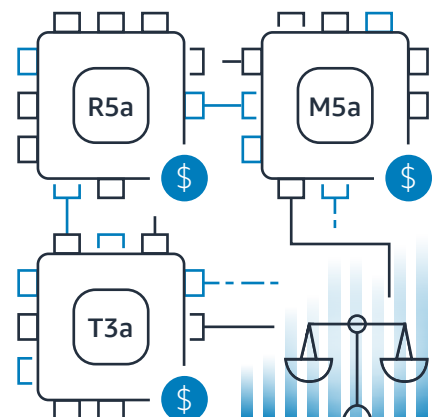
## AMAZON EC2 T3A INSTANCES

Amazon EC2 T3a instances are a low-cost Amazon EC2 instance type designed to provide a baseline level of CPU performance with the ability to burst above the baseline using CPU credits. T3a instances accumulate CPU credits when a workload is operating below baseline threshold. Each earned CPU credit provides the T3a instance the opportunity to burst with the performance of a full CPU core for one minute when needed.

T3a instances start in Unlimited mode by default, giving users the ability to sustain high CPU performance over any desired time frame while keeping costs as low as possible. For most general-purpose workloads, T3a Unlimited instances provide ample performance without any additional charges. If the average CPU utilization of a T3a instance is lower than the baseline over a 24-hour period, the hourly instance price automatically covers all interim spikes in usage. If the instance needs to run at higher CPU utilization for a prolonged period, it can do so for an additional charge of \$0.05 per vCPU-hour. In Standard mode, a T3a instance can burst until it uses up all of its earned credits.

### USE CASES

T3a instances offer a balance of compute, memory, and network resources for a broad spectrum of spiky general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development environments, code repositories, and business-critical applications.



# The Right Infrastructure for the Right Workload

## R5A AND R5AD INSTANCES SPECS

Model	vCPU	Memory (GiB)	Instance Storage (GiB)	Network Bandwidth (Gbps)	EBS Bandwidth (Mbps)
r5a.large	2	16	EBS-Only	Up to 10	Up to 2,120
r5a.xlarge	4	32	EBS-Only	Up to 10	Up to 2,120
r5a.2xlarge	8	64	EBS-Only	Up to 10	Up to 2,120
r5a.4xlarge	16	128	EBS-Only	Up to 10	2,120
r5a.8xlarge	32	256	EBS-Only	Up to 10	3,500
r5a.12xlarge	48	348	EBS-Only	10	5,000
r5a.16xlarge	64	512	EBS-Only	12	7,000
r5a.24xlarge	96	768	EBS-Only	20	10,000
r5ad.large	2	16	1 x 75 NVMe SSD	Up to 10	Up to 2,120
r5ad.xlarge	4	32	1 x 150 NVMe SSD	Up to 10	Up to 2,120
r5ad.2xlarge	8	64	1 x 300 NVMe SSD	Up to 10	Up to 2,120
r5ad.4xlarge	16	128	2 x 300 NVMe SSD	Up to 10	2,120
r5ad.12xlarge	48	348	2 x 900 NVMe SSD	10	5,000
r5ad.24xlarge	96	768	4 x 900 NVMe SSD	20	10,000

## M5A AND M5AD INSTANCES SPECS

Model	vCPU	Memory (GiB)	Instance Storage (GiB)	Network Bandwidth (Gbps)	EBS Bandwidth (Mbps)
m5a.large	2	8	EBS-Only	Up to 10	Up to 2,120
m5a.xlarge	4	16	EBS-Only	Up to 10	Up to 2,120
m5a.2xlarge	8	32	EBS-Only	Up to 10	Up to 2,120
m5a.4xlarge	16	64	EBS-Only	Up to 10	2,120
m5a.8xlarge	32	128	EBS-Only	Up to 10	3,500
m5a.12xlarge	48	192	EBS-Only	10	5,000
m5a.16xlarge	64	256	EBS-Only	12	7,000
m5a.24xlarge	96	384	EBS-Only	20	10,000
m5ad.large	2	8	1 x 75 NVMe SSD	Up to 10	Up to 2,120
m5ad.xlarge	4	16	1 x 150 NVMe SSD	Up to 10	Up to 2,120
m5ad.2xlarge	8	32	1 x 300 NVMe SSD	Up to 10	Up to 2,120
m5ad.4xlarge	16	64	2 x 300 NVMe SSD	Up to 10	2,120
m5ad.12xlarge	48	192	2 x 900 NVMe SSD	10	5,000
m5ad.24xlarge	96	384	4 x 900 NVMe SSD	20	10,000

## T3A INSTANCES SPECS

Model	vCPUs	Memory (GiB)	Baseline Performance/ vCPU	CPU Credits Earned/Hr	Network Burst Bandwidth (Gbps)	EBS Burst Bandwidth (Mbps)
t3a.nano	2	0.5	5%	6	5	1.50
t3a.micro	2	1.0	10%	12	5	1.50
t3a.small	2	2.0	20%	24	5	1.50
t3a.medium	2	4.0	20%	24	5	1.50
t3a.large	2	8.0	30%	36	5	2.05
t3a.xlarge	4	16.0	40%	96	5	2.05
t3a.2xlarge	8	32.0	40%	192	5	2.05



## A Few Lines of Code Can Save You 10%

The partnership between AWS and AMD has enabled new instances that offer cost-effective options that add to the broadest and deepest portfolio of EC2 instances. These instances will help customers optimize cost and performance for their workload. Experience the benefits AMD EPYC™-based EC2 Instances today!

Learn more at <https://aws.amazon.com/ec2/amd/>

