Optimizing HPC workloads with Amazon EC2 instances

Increasingly, customers are choosing to run High Performance Computing (HPC) on the AWS cloud for its scale and flexibility.

Engineers, researchers and scientists can maximize productivity and reduce time to results when running their HPC workloads on AWS. HPC system owners can deploy a cloud-based cluster environment in minutes, incorporating a range of AWS services designed specifically for demanding HPC applications such as computational fluid dynamics (CFD), weather and climate modeling, seismic reservoir, and structural simulations. This flexibility, along with the on-demand nature of the cloud, allows for a new approach to meeting the needs of HPC users and applications.

To get the best performance for your workload, AWS recommends choosing the right combination of services to meet the needs of your applications. AWS offers flexible compute, storage, and networking to enable the creation of an HPC environment to match your specific requirements. An HPC architecture can be represented as a stack composed of multiple layers. For each layer, AWS offers services to automate the deployment and simplify the management of the HPC cluster.

An important part of this will be choosing the right Amazon Elastic Compute Cloud (EC2) instances from the wide range of EC2 instance types available. Different families of EC2 instances are optimized for a range of workloads, allowing local data storage and networking along with the required operating system, based on the needs of your applications and the demands of your industry or vertical market. To assist you with identifying the key features and capabilities of your cluster in the cloud, we have described below a number of popular HPC workloads and the most appropriate EC2 instances for each one. This will ensure you have the best match to support your HPC requirements from the broad selection of available EC2 instances.

HPC Solution Layers	AWS Services
Remote Visualization	AWS NICE DCV
Job Scheduler and Orchestration	AWS Batch AWS ParallelCluster
Compute	Amazon EC2
Network	Amazon Elastic Fabric Adapter (EFA)
Storage	Amazon FSx, Amazon S3

Remote visualization

Visualization of results is an important aspect of an HPC workflow. Remote visualization helps accelerate turnaround times for scientists and engineers because users no longer need to download data from AWS to analyze the jobs' results. Similarly, users can benefit from remote visualization techniques also to prepare their inputs directly on AWS, instead of working on fat local client and then upload the file to the cluster to star their simulations.

NICE DCV is a remote visualization technology that enables users to securely connect to graphic-intensive 3D applications hosted on a remote server. In a typical DCV scenario, a graphic-intensive application, such as a 3D modeling or computer-aided design application, is hosted on an EC2 server that provides a high-end GPU, fast I/O capabilities, and large amounts of memory. The DCV client remotely connects to the session and uses the application hosted on the server. Then DCV compresses the visual output of the hosted application and streams it back to the user as an encrypted pixel stream.

HPC on Arm

In the in the past few years, we have seen an increased adoption of ARM architectures for HPC workloads. The milestone has been set by in November 2020 when Fugaku was awarded as the number one supercomputer on the Top500 list, for the second time in a row. This achievement highlighted that the Arm technology is rapidly evolving also for HPC applications due to its ability to combine power efficiency, performance, and scalability.

Amazon EC2 Hpc7g instances are powered by AWS Graviton3E processors designed for HPC workloads, such as computational fluid dynamics (CFD), weather forecasting, and molecular dynamics. AWS Graviton3E processors, with up to 35 percent higher vector instruction processing performance than the Graviton3, are designed to give you the best price/performance for tightly coupled workloads, and deliver 200 Gbps of dedicated network bandwidth that is optimized for traffic between instances in the same VPC.

EC2 offers several high performance instances with AMD and Intel processors that can run HPC workloads, however if you are able to build your application for Arm architecture, you can benefit from the leading price/performance offered by Arm-based AWS Graviton processors.

Amazon EC2 naming convention

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, and is grouped in an instance family based on these capabilities. With hundreds of different instance types, is important to understand how AWS groups them together so that you can easily identify the one that is right for your workload.

Please refer to the EC2 documentation to learn how the instance type naming convention works.

Instance naming explained

Compute workloads and orchestrators

The broad range of potential HPC workloads have individual characteristics that place different demands on the compute element of an HPC environment. By focusing on specific workloads and the industries that run them, we aim to help in selecting the right EC2 instances to meet your needs.

Tightly-coupled HPC workloads

Tightly coupled workloads (e.g. CFD, computer aided engineering (CAE), finite element analysis (FEA), global climate modeling, etc) consist of parallel processes that are dependent on each other to carry out the calculation. Unlike a loosely coupled computation, all processes of a tightly coupled simulation iterate together and require communication with one another.

Amazon EC2 HPC-optimized instances are optimal for running tightly-coupled workloads at scale on AWS.

Tightly-coupled applications typically depend on message passing interface (MPI) for parallel communications, so it's recommended to use EC2 instances that support low latency networking with Elastic Fabric Adapter (EFA), EFA is a network interface that reduces the network bandwidth between EC2 instances. You can find details regarding EFA in the networking section of this document.

There are three different types of HPC-optimized instances:

1. Amazon EC2 Hpc7g instances – powered by Arm-based AWS Graviton3E processors. It's available in three different sizes:

Instance Size	Physical Cores	Memory (GiB)	Memory per Core (GiB)	Network Bandwidth (Gbps[1])	EFA Network Bandwidth (Gbps)
hpc7g.16xlarge	64	128	2	25	200
hpc7g.8xlarge	32	128	4	25	200
hpc7g.4xlarge	16	128	8	25	200

Please note that the quantity of memory is constant across different sizes. Smaller sizes increase the memory per-core and memory bandwidth per-core and can improve your application's performance.

Some of these solvers, such as Siemens Simcenter Star-CCM+ and Ansys Fluent, also supports the ARM architecture, so you can use the Hpc7g instances. Some solvers can benefit from NVIDIA CUDA-enabled GPU acceleration, in which case you would choose EC2 instances powered by NVIDIA GPUs. Find details in the Accelerated Computing section of this document.

2. Amazon EC2 Hpc7a instances – and the previous generation Hpc6a instances are powered by AMD CPUs. These instances are designed to address many common workloads including CFD, weather forecasting and explicit FEA. Hpc6a instances are available in a single size, while Hpc7a is available in different sizes:

Instance Size	CPU	Physical Cores	Memory (GiB)	Memory per Core (GiB)	Network Bandwidth (Gbps[1])	EFA Network Bandwidth (Gbps)
hpc6a.48xlarge	AMD EPYC Milan	96	384	4	25	100
hpc7a.96xlarge	AMD EPYC Genoa	192	768	4	25	300
hpc7a.48xlarge		96		8		
hpc7a.24xlarge		48		16		
hpc7a.12xlarge		24		32		

Similar to Hpc7g instances, the quantity of memory in Hpc7a instances is constant across different instance sizes. Smaller sizes increase the memory per-core and memory bandwidth per-core and will have an impact on your application's performance. It's also important when you're using commercial software that's licensed on a per-core basis.

3. Amazon EC2 Hpc6id instances – are powered by the 3rd generation of Intel Xeon Scalable CPUs (Ice Lake). It's available in a single size:

Instance Size	CPU	Physical Cores	Memory (GiB)	Memory per Core (GiB)	Network Bandwidth (Gbps[1])	EFA Network Bandwidth (Gbps)	Instance Local storage (GiB)
hpc6id.32xlarge	Intel Ice Lake	64	1024	16	25	200	4x 3800 NVMe SSD

Hpc6id instances are designed for memory-bound and data-intensive workloads that need a local scratch disk such as seismic reservoir simulations, implicit FEA, and structural simulations. Hpc6id instances are equipped with large local NMVe disks that can accelerate the performance for I/O intensive applications.

Additionally, with AWS ParallelCluster you can provision resources in a safe, repeatable manner, allowing you to build and rebuild your infrastructure without manual actions or custom scripts. AWS ParallelCluster supports a wide variety of operating systems and batch schedulers so you can migrate your existing HPC workloads with little to no modifications.

Loosely-coupled grid computing

Loosely-coupled workloads (e.g. Monte Carlo calculations, financial risk forecasting and proteomics) may not always require high-performance interconnect, even as the many CPU cores they occupy still test the system's speed limits. These workloads involve many largely independent work items, meaning they do not depend on a high-speed network for optimum performance. For workloads where high single threaded performance is most important, consider Amazon EC2 M5zn instances with a clock speed of 4.5GHz (up to 192 GB of memory or X2iezn with up to 1.536 GB of memory) or Z1d with a clock speed of 4.0GHz (up to 384 GB of memory).

Amazon EC2 R7iz instances are powered by the latest 4th generation Intel Xeon Scalable Processors (Sapphire Rapids) and run at a sustained all-core turbo frequency of 3.9 GHz. With high CPU frequency and DDR5 memory, these instances are a perfect match for your Electronic Design Automation (EDA), financial, actuarial, and data analytics simulation workloads. For other applications that prioritize lower cost per core over single threaded performance select the Amazon EC2 C, Amazon EC2 M or Amazon EC2 R instances (based on memory requirements) or HPC instances.

To orchestrate this type of workload, customers usually prefer AWS Batch. Batch is a container-based managed task execution service. Batch provides job queues with sophisticated scheduling capabilities and compute environments to define the size and shape of the worker nodes. To use an HPC application with Batch you'll first need a containerized version of it. This is a straight-forward procedure when the application is pleasingly parallel and doesn't require cross-node communication using MPI.

Accelerated computing

Some HPC workloads may benefit from the flexibility to activate or increase the mixture of GPU and FPGA accelerators in a compute workload. The additional threads and cores of GPUs can be harnessed to deliver results faster. HPC applications and codes that use CUDA and OpenACC can experience significant performance improvements when using accelerators. Amazon EC2 P5 instances, powered by NVIDIA H100 GPUs, provide higher performance for accelerated HPC applications such as natural language processing, object detection and classification, seismic analysis, and molecular dynamics. P4d instances are powered by eight NVIDIA A100 Tensor Core GPUs and support 400 Gbps instance networking. CUDA and OpenACC applications or rendering workloads can benefit from the A10G Tensor Core GPUs offered on the Amazon EC2 G5 instances or from the NVIDIA T4 GPUs offered on the Amazon EC2 G4dn instances.

The following table summarizes the current and previous generation of Amazon EC2 instances powered by NVIDIA GPUs.

Instance Family	vCPU	Physical Cores	Memory (GiB)	GPU	GPU memory (GiB)	Network performance	Instance Storage
P4d	96	48	1152	8x NVIDIA A100	40	4x 100 Gigabit	8x 1000 GB NVMe SSD
P4de					80		
P5	192	96	2048	Up to 8x NVIDIA H100	80	32x 100 Gigabit	8x3.84 GB NVMe SSD
G4dn	Up to 96	Up to 48	Up to 384	Up to 8x NVIDIA T4	16	Up to 100 Gigabit	Up to 2x 900 NVMe SSD
G5	Up to 192	Up to 96	Up to 768	Up to 8x NVIDIA A10G	24	Up to 100 Gigabit	Up to 2x 3800 NVME SSD

Customers optimize specific elements of their computation using hardware accelerators such as FPGAs and can enable customized hardware acceleration for applications including genomics, analytics and financial risk modeling. FPGAs are featured in the Amazon EC2 F1 instances, and can offer over 100x acceleration when compared to CPUs for a range of use cases. To support the use of the F1 instances, AWS provides a range of tools and resources for developers to rapidly harness the performance of FPGAs.

Choose the right Amazon EC2 instances for your HPC workloads

Do you need accelerators for your jobs?

Specialized hardware accelerators like GPUs enhance the computational capability of workloads. With their parallel processing capabilities, they enable faster execution of complex calculations and data-intensive tasks such as ML, scientific simulations, and graphics rendering, boosting overall performance and efficiency.

No

Do you need high clock speed processors?

Determine how quickly the processors can execute instructions, leading to faster processing and overall improved performance in computationally intensive tasks.

Yes

r7id
m5zn
x2iezn
z1d

Use these if you have significant single threaded bottlenecks, or have licenses that make it costly to add more cores.

No

Can you use ARM-based CPUs for your jobs?

Yes

hpc7g
c7gn

Hpc7g offers the best price performance for HPC workloads.

No

Do you need a local scratch disk?

Yes

hpc6id
r6id

No

hpc7a
hpc6a
c6a
c6i

Yes

Are you using GPU for rendering?

Specialized hardware accelerators like GPUs enhance the computational capability of workloads. With their parallel processing capabilities, they enable faster execution of complex calculations and data-intensive tasks such as ML, scientific simulations, and graphics rendering, boosting overall performance and efficiency.

Yes

g4ad
g4dn
g5

Select one of the above cost-effective and versatile GPU-based instances, suited for ML and graphics-intensive workloads.

No

Are you running ML training?

Yes

p5
p4d
p4de
dl1
trn1

No

inf2

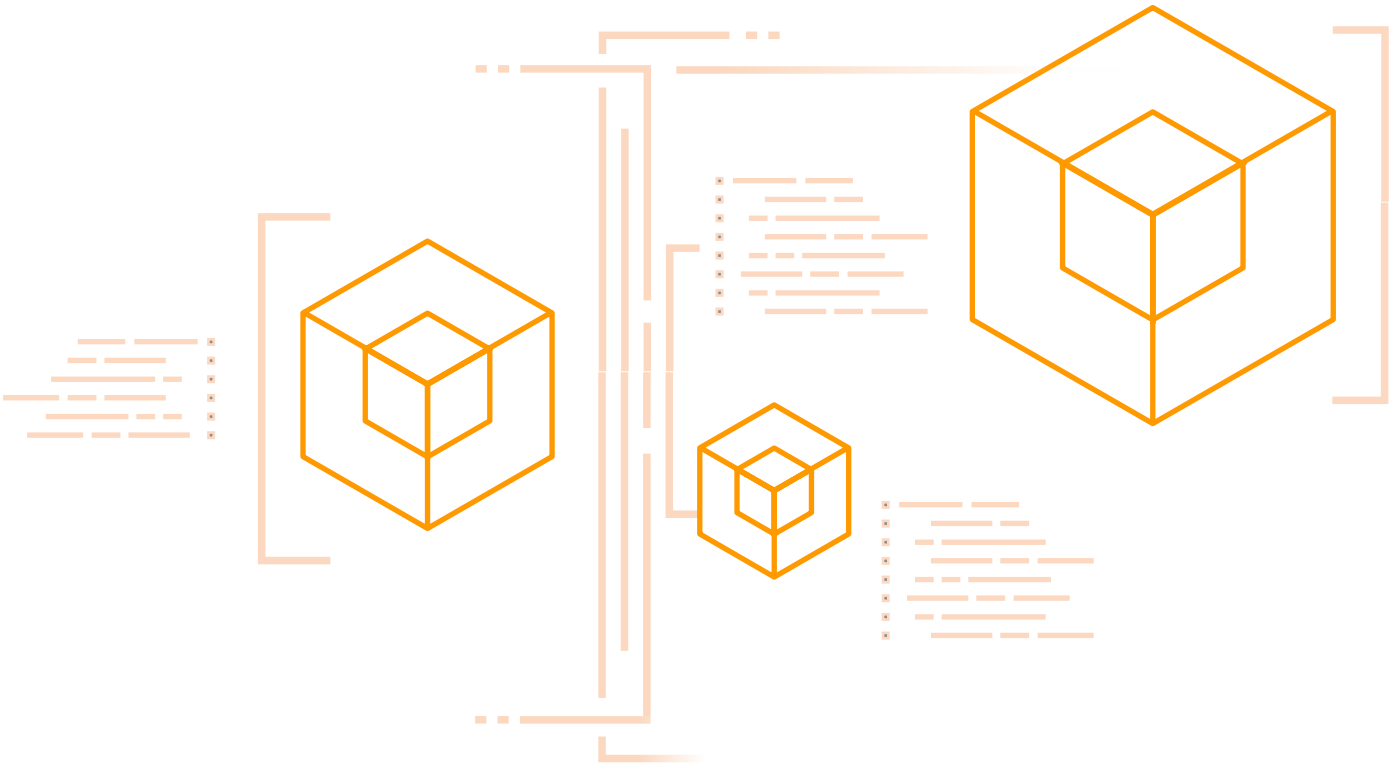

Machine Learning and AI

On-premises HPC administrators struggle to keep pace with the demands of their growing neural network training and inference workloads. AWS machine learning (ML), deep learning and AI-optimized instances leverage GPU power, and high-memory instances for inference—all of which are customizable to fit a user's workload needs.

Together with Accelerated Computing instances previously mentioned, EC2 offers Amazon EC2 Trn1 instances, powered by AWS Trainium accelerators for high-performance deep learning training. AWS Trainium is the second-generation ML accelerator developed by AWS for deep learning training of 100B+ parameter models. Trn1 instances deploy up to 16 AWS Trainium accelerators to deliver a high-performance, low-cost solution for deep learning training.

Instance size	vCPUs	Instance memory	Trainium chips	Accelerator memory	Neuron-Link	Instance networking
Trn1.2xlarge	8	32 GB	1	32 GB	N/A	Up to 12.5 Gbps
Trn1.32xlarge	128	512 GB	16	512 GB	Yes	800 Gbps
Trn1n.32xlarge	128	512 GB	16	512 GB	Yes	1600 Gbps

For fast communication between accelerators, these instances support NeuronLink, a high-speed, nonblocking interconnect. For ultra-large models that do not fit into a single accelerator, data flows directly between accelerators with NeuronLink, bypassing the CPU completely.

Moreover, Trn1 instances support up to 1600 Gbps of EFA network bandwidth to deliver higher performance for network-intensive models. Instances can be deployed in EC2 UltraClusters that enable scaling up to 30,000 Trainium accelerators, which are interconnected with a nonblocking petabit-scale network to provide 6 exaflops of compute performance.

Amazon EC2 Inf2 instances are designed to deliver high performance at the lowest cost in Amazon EC2 for inference and generative AI applications. They are optimized to deploy increasingly complex models, such as large language models (LLMs) and vision transformers, at scale.

EC2 Inf2 instances are powered by the AWS Inferentia2 chips that are designed to deliver high performance at the lowest cost for inference applications. Inferentia2 accelerators deliver up to 4x higher throughput and up to 10x lower latency compared to previous version of the Inferentia accelerator.

Instance size	vCPUs	Instance memory	Inferentia2 chips	Accelerator memory	Neuron-Link	Instance networking
Inf2.xlarge	4	16 GB	1	32 GB	N/A	Up to 15 Gbps
Inf2.8xlarge	32	128 GB	1	32 GB	N/A	Up to 25 Gbps
Inf2.24xlarge	96	384 GB	6	192 GB	Yes	50 Gbps
Inf2.48xlarge	192	768 GB	12	384 GB	Yes	100 Gbps

Network

The optimal network solution for HPC workloads varies based on latency, bandwidth, and throughput requirements. Tightly coupled HPC applications require the lowest latency possible for network connections between compute nodes. For moderately sized tightly coupled workloads, it is possible to select large instance types with a large number of cores so that the application fits entirely within the instance without crossing the network at all.

Alternatively, some applications are network bound and require high network performance. EC2 instances with higher network performance can be selected for these applications. Higher network performance is usually obtained with the largest instance type in a family. Refer to the Instance Type Matrix for more details.

Multiple instances with low latency between the instances are required for large tightly coupled applications. On AWS, this is achieved by launching compute nodes into a Cluster Placement Group (CPG), which is a logical grouping of instances within an Availability Zone. A CPG provides non-blocking and non-oversubscribed connectivity, including full bisection bandwidth between instances. Use CPGs for latency sensitive tightly coupled applications spanning multiple instances.

EFA uses a custom-built operating system bypass technique to enhance the performance of inter-instance communications, which is critical to scaling HPC applications. With EFA, HPC applications using popular HPC technologies like MPI can scale to tens of thousands of CPU cores. EFA supports industry-standard libfabric APIs, so applications that use a supported MPI library can be migrated to AWS with little to no modification. EFA is available as an EC2 networking feature that you can enable on supported EC2 instances at no additional cost. Here you can find how to configure EFA.

Storage

Once you have selected the most appropriate EC2 instances to meet the needs of your workload, your next consideration will be which storage option best meets your requirements. You will want to select both for storage attached directly to your chosen compute instances and storage presented as a file system attached to your cloud based HPC cluster. AWS provides options for storage including Amazon Elastic File System (EFS) or Amazon Elastic Block Store (EBS), which can be added directly to EC2 compute instances, providing local scratch storage or access to a POSIX file system should this be required by your applications.

If your applications require a high performance file system to handle large volumes of data, AWS offers Amazon FSx, a suite of fully managed services from AWS designed to help customers to deploy and manage file systems in the cloud. FSx supports a wide range of workloads with its reliability, security, scalability, and broad set of capabilities. In particular, Amazon FSx for Lustre is a managed service that provides a cost effective and performant solution for HPC architectures requiring a high-performance Lustre parallel file system. Similarly, Amazon FSx for OpenZFS is a fully managed file storage service that provides a ZFS file system. Based on your application needs, you can explore additional Amazon FSx file systems such as Amazon FSx for NetApp ONTAP and Amazon FSx for Windows File Server.

“The Amazon EC2 instance types you specify determines the hardware of the host computer used for your instance. Instance types offer different compute, memory, and storage capabilities, and are grouped in instances categories based on these capabilities.”

Summary

Amazon EC2 instances provide a broad range of options for running high performance workloads in the cloud, enabling you to choose the right compute platform with the right specifications to match codes and application needs precisely. AWS offers nearly limitless configuration of instances and solutions configurations to cater to the demands of individual applications and workloads, and provides virtually unlimited infrastructure, latest generation processors and accelerator technology with optional high speed, low latency networking. Whether migrating your HPC applications to the cloud or wish to burst to the cloud, AWS provides a broad range of scalable, flexible infrastructure and services you can select to match your workloads and tasks.

More and more, organizations are considering the cloud for their HPC workloads, and AWS is committed to helping customers harness the full potential of HPC in the cloud. When you run your HPC workloads on AWS your get the flexibility, scale, and performance to iterate quicker, innovate faster, and reducing time to results.

Accelerate innovation with HPC on AWS

https://aws.amazon.com/hpc ›

1. 500 Mbps network bandwidth outside of the virtual private cloud (VPC) and Amazon Simple Storage Service (Amazon S3), or Amazon Elastic Block Store (EBS).