



SOLUTION BRIEF

Everything You Need for Artificial Intelligence and Machine Learning

Get top-of-the line resources at your fingertips with Vultr.

[VULTR.COM](https://vultr.com)

Artificial intelligence (AI) and machine learning (ML) have had astounding growth over their brief history. They work based on neural networks, a series of algorithms that mimic how the human brain creates connections between elements. This concept has been around for a while but can now be applied to make highly granular decisions and solve complex problems.

The feats of AI computing are only possible through graphical processing units (GPUs). GPUs are like central processing units (CPUs) built with more specific tasks in mind. Initially, GPUs helped render complex graphics like realistic light and shadow in 3D animation. Since then, the use cases for GPUs have expanded as developers discovered their ability to handle AI and ML-related projects.

The GPU focuses on the concurrency of aspects in a single problem in parallel. It doesn't deal with the switching contexts, unlike the more generalized CPU. GPUs are especially well-suited to AI and ML problems as they can support complex, multi-step processes attacked by parallel computing. Parallel computing and AI rely on the subject area's interconnected parts to produce a solution.

Due to their complexity and degree of specialization, GPUs come with a heavy price tag. On-premise GPU options need constant updates and maintenance, requiring additional training or hiring of personnel with GPU-specialized skills. Large cloud providers offering GPU services cost more and can be impractical for the number of resources and experimentation involved with AI and ML.


Vultr offers flexible and customizable solutions for AI and ML

AI and ML projects are experimental and investigative. While this is great for learning, it doesn't make planning for the future easy. AI and ML projects require a lot of resources for testing, training, and updating models. Unfortunately, these resource needs aren't always consistent and they vary from project to project, making them difficult to predict.

Vultr customizability

The hyperscalers provide a variety of GPU offerings and support for AI and ML, but their support is tailored to their platforms, creating vendor lock-in. These providers offer pricing plans that encourage long-term commitment with various costs for switching hardware during the commitment period.

For example, AWS has discounts for long-term reserved instances but is challenging to set up and has few GPU options. Similarly, Google Cloud requires you to attach GPUs to standard virtual machines (VMs), which can make pricing confusing. Meanwhile, Azure is for full-production GPU instances, making it difficult to scale up.



Vultr offers a wide variety of hardware packages by default, set up with everything you need. Your environment suits your specifications, hardware dependencies are accounted for, and you're ready to go in a plug-and-play, maintained environment. You don't have to choose a predetermined platform. If you change your approach, Vultr allows you to change your configuration.

Processing power

No matter how much processing power your project needs, Vultr provides a wide variety of GPU offerings. As your knowledge of the problem and your workload expands, you may migrate to the higher processing power: [the A100](#).

While the A100 is marginally faster on clock speed, it provides better memory clock speed and bandwidth to perform faster on more significant problems.

The fractional advantage

Vultr's virtualization model also employs fractional GPU usage. The fractional model allows the GPU to split into smaller parts, each operating independently of the others, allowing your engineers to spin up GPUs in the exact size needed without the expense of an entire piece of hardware. Multiple teams can work on different projects using the same GPU, each taking a slice, without needing any client setup or configuration.

Vultr has cloud data centers across six continents. Teams worldwide can work in parallel on the same hardware to promote consistency and collaboration. The Vultr control panel allows you to monitor fractional GPU use so you can scale it according to your teams' needs and budget.

Additional tools

Besides GPUs, Vultr offers the tools needed to perform cloud operations and support geographically dispersed teams. Vultr provides a flexible environment for various [operating systems](#) (OS) or deploying a custom ISO so that you can deploy apps such as Docker, WordPress, or GitLab with a single click. Vultr provides Kubernetes Management, SSD Storage, Backups, Flexible Networking, and DDOS protection options.

Vultr provides all these tools within a single infrastructure, allowing AI and ML teams to take advantage of the automated management of GPU containers. Developers can also use Vultr's SSD storage options to improve the [performance](#) of neural network applications. You can create environments your teams are already familiar with, reducing the learning curve.

Vultr offers the advantage of flexibility. From fractional GPUs to a suite of choices, your development teams can choose the most effective environments for their project while controlling costs.

The Vultr approach

For an AI or ML team to be as effective as possible, they need the freedom to experiment, prototype, and adjust. There's no one-size-fits-all solution, and the initial direction teams choose may not always be correct, which is why Vultr strives to be as approach-agnostic as possible.

The full array of services and hardware empowers teams to use what they need when they need it. Vultr's role is to support and enable, not force or interfere through unnecessary lock-ins or extra costs.

Changing approaches shouldn't mean changing cloud vendors or learning new infrastructure. From its array of hardware, software, geographic locations, and pricing options, Vultr provides the flexibility to support your success.

Power AI/ML initiatives with Vultr

The cycle of bringing AI or ML projects from investigation to production often requires the development team to explore different approaches. The path to a worthwhile AI and ML project may have several branches, each requiring variations in developers' tools.

Organizations must invest in flexible solutions to ensure their AI and ML projects are successful. These experimental forms often need you to adjust resource use, stay conscious of GPU capabilities, and track ML model size. Vultr's flexibility and implementation speed enable you to meet these needs.

▶ Check out the full range of Vultr's alternative [GPU offerings](#) to support your AI and ML projects.