


Cloud GPU

High-performance bare metal servers and virtual machines for AI/ML, HPC, rendering, and remote workstations. Full range of sizes, from multi-GPU configurations, to fractions of GPUs.



NVIDIA GPUs for every use case

Vultr offers machines powered by a variety of NVIDIA GPUs, so you can accelerate a wide range of workloads.




NVIDIA A100

Delivering unprecedented acceleration and powering the world's highest-performing AI, data analytics, and HPC workloads.



NVIDIA A40

Combining professional graphics with powerful compute and AI, to meet today's design, creative, and scientific challenges.



NVIDIA A16

Enabling virtual desktops and workstations with the power and performance to tackle any project from anywhere.

Breakthrough GPU Virtualization

Talon's unique approach partitions physical NVIDIA GPUs into discrete virtual GPUs, each with their own memory and compute. Perfect for AI inference and model building workloads such as natural language processing, voice recognition and computer vision.



Powerful add-ons

Mount high performance scalable storage to your instance, making space management significantly more flexible. Available add-ons include backups, snapshots, firewalls, flexible networking, and DDoS protection.



Backups



Snapshots



Firewall



Flexible
Networking



DDoS
Protection

Powerful platform worldwide



100% SLA guaranteed

Enterprise grade stability and performance.



Custom ISOs

Deploy any operating system of your choosing.



Best-in-class AMD & Intel CPUs

The latest generation AMD and Intel CPUs guarantee fast performance.



NVMe SSD

Blazing fast read/write speeds on NVMe solid-state disk hardware.



Virtual private clouds

Establish isolated regional VPCs for your applications.



Public IPv6 network

Develop for the next generation internet protocol.



Root administrator access

Full "root" access and a dedicated IP address included with all VMs.



API & Terraform

Use our API or Terraform provider to quickly create and control your instances.

More cloud, less money

The Vultr Talon platform makes accelerated computing affordable and easy to use.
In just minutes, provision VMs with fractional NVIDIA GPUs.



NVIDIA A100

Delivering unprecedented acceleration and powering the world's highest-performing AI, data analytics, and HPC workloads.

GPU	GPU RAM	vCPUs	RAM	Storage	Bandwidth	Monthly Price	Hourly Price
1/20	4 GB	1 vCPU	6 GB	70 GB	1 TB	\$90 /mo	\$0.143 /hr
1/10	8 GB	1 vCPU	12 GB	140 GB	1 TB	\$180 /mo	\$0.268 /hr
1/7	10 GB	2 vCPU	15 GB	170 GB	2 TB	\$250 /mo	\$0.372 /hr
2/7	20 GB	3 vCPU	30 GB	350 GB	3 TB	\$450 /mo	\$0.670 /hr
3/7	40 GB	6 vCPU	60 GB	700 GB	6 TB	\$875 /mo	\$1.302 /hr
1	80 GB	12 vCPU	120 GB	1400 GB	10 TB	\$1,750 /mo	\$2.604 /hr
2	160 GB	24 vCPU	240 GB	2800 GB	10 TB	\$3,500 /mo	\$5.208 /hr
4	320 GB	48 vCPU	480 GB	5600 GB	15 TB	\$7,000 /mo	\$10.417 /hr
8	640 GB	96 vCPU	960 GB	10200 GB	25 TB	\$14,000 /mo	\$20.833 /hr



NVIDIA A40

Combining professional graphics with powerful compute and AI, to meet today's design, creative, and scientific challenges.

GPU	GPU RAM	vCPUs	RAM	Storage	Bandwidth	Monthly Price	Hourly Price
1/6	8 GB	4 vCPU	20 GB	360 GB	5 TB	\$210 /mo	\$0.313 /hr
1/2	24 GB	12 vCPU	60 GB	1110 GB	10 TB	\$625 /mo	\$0.930 /hr
1	48 GB	24 vCPU	120 GB	1440 GB	15 TB	\$1,250 /mo	\$1.860 /hr
4	192 GB	96 vCPU	480 GB	8900 GB	60 TB	\$5,000 /mo	\$7.440 /hr
8	384 GB	192 vCPU	960 GB	17800 GB	120 TB	\$10,000 /mo	\$14.880 /hr