


GOOGLE CLOUD GPU: THE BENEFITS OF USING GPUS IN THE CLOUD



It's no secret—there are many advantages for companies to use [cloud computing](#). Physical, on-prem computation requires teams of people to maintain, scale, and orchestrate computer systems. Instead, a company can pay a cloud provider to do all these things.

For companies that do a lot of heavy computing, machine learning, or 3D visualizations, the way you compute your GPUs matters. Here, let's look at the benefits of cloud computing with GPUs, particularly with Google Cloud.

What are GPUs?

We'll start with the basics. A Graphics Processing Unit (GPU) is a specialized electronic circuit. Compared to your standard computer with its Central Processing Unit (CPU), a GPU has a parallel structure that offers faster computing and increased efficiency. GPUs are commonly used for handling 3D graphics or, what the bitcoin craze revealed, [dense computation](#).

Not everyone needs to use GPUs. Cloud GPUs allow a single engineer or designer to have a great amount of computation at their fingertips for a relatively cheap amount. Common use cases for GPUs are in design: from video game renderings, to architectural and engineering renderings, to 3D animation. GPUs can decrease the time spent waiting for a computer to handle these tasks.

The use of GPUs has also surged thanks to [machine learning](#). Engineers can form and test models

much faster on a GPU than on a CPU.

How cloud computation works

Computers are the lifeblood of most companies. Cloud computing helps take care of this necessity, [as a managed service](#). Cloud computing also makes working with GPUs significantly easier and more accessible.

Here are the key benefits of cloud computing for companies:

- **Simplifies business operations.** If a company had to find a way to create its own electricity, it adds just one more responsibility to the company's list of things to do, diverting attention from [service delivery](#) and the customer experience. (Such a task might be a non-starter, stopping many companies from beginning at all.)
- **Encourages specialization.** Cloud computing also allows a company to be more specialized—because you no longer need to maintain a complicated system of computers to do all your work.
- **Reduces cost.** Cloud computation turns computation expenses from [capital expenses to operating expenses](#). This is beneficial to new companies because it lowers the cost of entry into a field. It's accountant friendly because OpEx are easier to account for than CapEx.
- **Promotes experimentation and innovation.** The speed of cloud computation allows your employees and creators to stay in the flow, improving productivity and innovation.

Benefits of Google Cloud GPUs

Now, let's bring together GPUs and cloud computing. Companies should use Google Cloud GPUs if for no other reason than to have a good option for heavy computations. Even if your operation is small and doesn't use many resources, it is still a simple, best practice to start using Google Cloud GPUs. (For ML-optimized hardware in the cloud, Google also [Google Cloud TPUs](#)).

Here are several key benefits to using Google Cloud GPUs:

Saves time

If computers and Kanban product management systems have taught us anything in the past decade, it's been "iteration is key". And designers can rapidly iterate by having faster rendering times. Even a difference of 8 minutes to 10 seconds is a huge difference—it allows a creator to remain in their flow.

Decrease render times from potentially days to just a few hours by using a small [render farm](#). Depending on experience, it is possible to set up a render farm in a few hours to a week. A team's workflow will see render times decrease from a few hours to minutes, and the effect will last for years.

Cloud GPUs can adjust machine learning modelling times from 8-12 hours to 10-15 minutes. I ran a small machine learning model in a Google Colab notebook on a 2MB dataset (tiny!). It took 90 minutes to train on a CPU. When switching to the Google provided GPU, the model trained in only eight minutes.

Frees local resources

Cloud compute doesn't consume local resources. Running a large Machine Learning model or a rendering task can slow or even make a local computer unusable. By outsourcing compute power to the cloud, a person can continue using their computer with ease.

The role of the personal computer has changed. With cloud computing, an individual's computer becomes a command center. It orchestrates computational tasks and rarely handles any workloads itself.

Scales easily

If you're ready to expand your team or workload, GPUs are ready to add, on command,. The process will be familiar, so adding more GPUs in the event of winning larger contracts can easily be executed without learning much more about the technologies.

Minimizes cost

Prices to run these computers are cheap. A new NVIDIA Tesla P100 costs \$2,900. Or, you can use four GPUs for 4 hours on Google Cloud for only [\\$23.36](#). That can be a [huge performance boost](#) over most off-the-shelf, physical solutions.

And at \$.50 to \$1.50 per hour and an ability to pay [by the second](#) with a one-minute minimum, the cost of innovation is cheap!

How to get started with Google Cloud GPUs

Starting with Google Cloud GPUs is a breeze. First, sign up for Google Cloud Platform (GCP). You'll get all the standard benefits that come with GCP: storage options, cloud functions, database management, and integration with frontend applications. You also get [\\$300 in free credit](#).

Then, you can use [Google Colaboratory](#). Colab is Google's version of a [Jupyter Notebook](#) and allows free usage of a single GPU or TPU.

Finally, try it out! Use your own exercise on try [rendering GPUs for game development](#).

Additional resources

To learn more about the cloud for your enterprise, including comparisons how-to's for GCP, AWS, and Azure, browse the [BMC Cloud Blogs](#). For more about machine learning, check out our [Machine Learning Guide](#).