



M-STAR

**GPU-Native, Next-Gen
CFD Software**

Graphic Processing Unit (GPU) FAQs

M-Star CFD is designed to run on GPUs – pairing modern algorithms with modern GPU architectures to deliver maximum science. There are a lot of GPU options, so figuring out which one you need can be overwhelming. Before you invest, refer to the FAQs below for system requirements, guidelines and recommendations.

1. Why do I need a GPU?

For workstations or desktop setups, you will need a GPU to not only run the **solver** but also complete **pre- and post-processing**. Finding the right GPU for your needs that can support all three functions is important to ensure maximum performance.

It is not recommended to run on CPU due to the fact that GPUs – even consumer-grade level ones – will always provide a better user experience. Additionally, many advanced features in M-Star that use custom UDFs are not supported on CPU.

2. What brand GPU does M-Star support?

We only support NVIDIA GPUs. We do not support AMD GPUs at this time.

3. Which GPU cards can I use with M-Star?

M-Star's code will run on any NVIDIA GPU since 2014 (Compute Capability 3.5 or newer), including **Tesla A100, Tesla V100 and GeForce RTX 3090**. GPUs can be classified in three different categories: gaming, workstation and server.

Gaming

Minimum required performance at a lower cost for individual workstations

- RTX 2080
- RTX 2080 Ti
- Titan RTX
- RTX 3080
- RTX 3090

Workstation

Higher performance at a higher cost for individual workstations

- Titan V
- Quadro GP100
- Quadro GV100
- RTX A6000
- Quadro RTX 8000

Server

Higher performance and excellent scalability at a higher cost for server enclosures (not individual)

- Tesla P100
- Tesla V100
- Tesla A100
- Tesla A30

4. What should I consider when selecting a GPU that's right for me?

Deciding on a GPU card is a balancing act between cost and performance. Specific factors to weigh include:

- **Memory Size:** The capacity of the GPU. Memory size ultimately **limits the size of the problem you can solve**. Simulations are made up of cells and particles. The more memory you have, the more cells and particles you have.
- **Memory Bandwidth:** The speed at which a GPU moves memory from RAM to compute units, an important indicator of performance.
- **Single Precision:** Another indicator of speed. Single precision performance measures how fast a GPU can perform calculations. *Double precision performance is not required for M-Star.
- **NVLINK Support:** How you connect multiple GPUs together. NVLINK should always be installed between GPUs to ensure top performance. [Learn more.](#)

4 million cells per GB
Most users will want a capacity of tens of millions of cells.

Let's compare common gaming, workstation and server cards by these factors:

	Card Name	Memory Size	Memory Bandwidth	Single Precision	Cost	NVLINK Support
Gaming	GeForce RTX 2080	8 GB	448 GB/s	8.9 TFlops	USD 799	Yes
	RTX 2080 Ti	11 GB	616 GB/s	11.7 TFlops	USD 1000	Yes
	Titan RTX	24 GB	672 GB/s	12.4 TFlops	USD 2500	Yes
	RTX 3080	10 GB	760 GB/s	25.1 TFlops	USD 700	No
	RTX 3090	24 GB	936 GB/s	29.3 TFlops	USD 1500	Yes
Workstation	Titan V	12 GB	653 GB/s	12.2 TFlops	USD 3000	No
	Quadro GP100	16 GB	732 GB/s	10.3 TFlops	USD 7000	Yes
	Quadro GV100	32 GB	868 GB/s	12.8 TFlops	USD 9000	Yes
	RTX A6000	48 GB	768 GB/s	38.7 TFlops	USD 5000	Yes
	Quadro RTX 8000	48 GB	672 GB/s	16.3 TFlops	USD 5500	Yes
Server	Tesla P100	16 GB	730 GB/s	9.3 TFlops	USD 3000	Yes
	Tesla V100	32 GB	900 GB/s	14 TFlops	USD 8000	Yes
	Tesla A100	40 GB	1550 GB/s	19.5 TFlops	USD 9500	Yes
	Tesla A30	24 GB	933 GB/s	10.3 TFlops	USD 5000	Yes

5. How many GPUs do I need?

This depends on the size of the problem you want to solve. One of the biggest benefits of multiple GPUs is that you can essentially double your memory, allowing you to fit a bigger problem on it. **Most M-Star users deploy 1-2 GPUs.**

Note that not all GPUs support NVLINK connection. If you choose to connect more than one GPU, they should be the same card.

6. What GPU hardware do you recommend?

Whenever possible, we suggest using NVIDIA Tesla GPUs – enterprise-grade equipment intended to be used for heavy computational loads. Features include:

- TCC Mode on Windows: Important for multi-GPU jobs.
- ECC memory: More reliable than non-ECC RAM; important for high-value data.
- Higher memory capacity: Important for handling larger problems.

7. What about GeForce- or Quadro-based hardware?

GeForce-based hardware is more consumer-grade GPUs intended primarily for rendering 3D scenes. These are compatible with the M-Star solver and can be highly cost-effective in regard to memory and CUDA cores. However, multi-GPU use of GeForce hardware is generally limited to Linux. For workstations being used for pre- and post-processing in addition to running simulations, we recommend having a GeForce GPU dedicated to pre- and post-processing, especially when running Windows.

Quadro-based GPUs work fine and can be a cost-effective means to run M-Star.

For more details on minimum requirements, recommended workstations and recommended servers, please review our documentation, [here](#).