



# CST STUDIO SUITE® 2017

Supported GPU Hardware



# 1 Supported Hardware

CST STUDIO SUITE currently supports up to 8 GPU devices in a single host system, meaning each number of GPU devices between 1 and 8 is supported.<sup>1</sup>

The following tables contain some basic information about the GPU hardware currently supported by the GPU Computing feature of CST STUDIO SUITE, as well as the requirements for the host system equipped with the hardware. To ensure compatibility of GPU hardware and host system please check

<https://www.nvidia.com/object/tesla-qualified-servers.html>

Please note that a 64 bit computer architecture is required for GPU Computing. A general hardware recommendation can be found in the FAQ section of the CST support [website](#) (FAQ No.3908).

---

<sup>1</sup>It is strongly recommended to contact CST before purchasing a system with more than four GPU cards to ensure that the hardware is working properly and is configured correctly for CST STUDIO SUITE.

List of supported GPU hardware for CST STUDIO SUITE 2017 <sup>2 3</sup>

Card Name	Series	Platform	Min. CST Version
Tesla P100-SXM2 (Chip)	Pascal	Servers	2017 release
Tesla P100-PCIE-16GB	Pascal	Servers	2017 release
Tesla P100 16GB	Pascal	Servers	2017 release
Tesla P100-PCIE-12GB	Pascal	Servers	2017 SP 2
Quadro P6000 <sup>4</sup>	Pascal	Workstations	2017 SP 2
Quadro GP100	Pascal	Workstations	2017 SP 2
Tesla M60 <sup>4</sup>	Maxwell	Servers/Workst.	2016 SP 4
Tesla M40 <sup>4</sup>	Maxwell	Servers	2016 SP 4
Quadro M6000 24GB <sup>4</sup>	Maxwell	Workstations	2016 SP 4
Quadro M6000 <sup>4</sup>	Maxwell	Workstations	2015 SP 4
Tesla K80	Kepler	Servers	2014 SP 6
Tesla K40 m/c/s/st/d/t	Kepler	Servers/Workst.	2013 SP 5
Quadro K6000	Kepler	Workstations	2013 SP 4
Tesla K20X	Kepler	Servers	2013 release
Tesla K20m/K20c/K20s	Kepler	Servers/Workst.	2013 release
Tesla K10 <sup>4</sup>	Kepler	Servers	2013 release
Quadro 6000 <sup>5</sup>	Fermi	Workstations	2012 SP 6
Tesla Fermi M-Series <sup>5</sup>	Tesla 20/Fermi	Servers	2011 SP 6
Tesla Fermi C-Series <sup>5</sup>	Tesla 20/Fermi	Workstations	2011 SP 6
Tesla M1060 <sup>5</sup>	Tesla 10	Servers	2010 SP 4
Tesla C1060 <sup>5</sup>	Tesla 10	Workstations	2010 release
Quadro FX5800 <sup>5</sup>	Tesla 10	Workstations	2010 release
Quadro Plex 2200D <sup>5</sup>	Tesla 10	Workstations	2010 release



<sup>2</sup>Please note that cards of different series (e.g. "Maxwell" and "Pascal") can't be combined in a single host system for GPU Computing.

<sup>3</sup>Platform = Servers: These GPUs are only available with a passive cooling system which only provides sufficient cooling if it's used in combination with additional fans. These fans are usually available for server chassis only!

Platform = Workstations: These GPUs provide active cooling, so they are suitable for workstation computer chassis as well.

<sup>4</sup> **Important:** The double precision performance of this GPU device is poor, thus, it can't be recommended for PIC-solver and I-solver (double precision) simulations.



<sup>5</sup> **Important:** This hardware is marked as deprecated and won't be supported in upcoming CST STUDIO SUITE versions (2018 and newer).

Hardware Type	 NVIDIA Tesla K20c/K20m/K20s (for Workst./Servers)	 NVIDIA Tesla K20X (for Servers)
Min. CST version required	2013 release	2013 release
Number of GPUs	1	1
Max. Problem Size (Transient Solver)	approx. 50 million mesh cells	approx. 60 million mesh cells
Form Factor	Dual-Slot PCI-Express	Dual Slot PCI-Express
Memory	5 GB GDDR5	6 GB GDDR5
Bandwidth	208 GB/s	250 GB/s
Single Precision Performance	3.52 TFlops	3.95 TFlops
Double Precision Performance	1.17 TFlops	1.32 TFlops
Power Consumption	225 W (max.) requires two auxiliary power connectors	235 W (max.)
PCI Express Requirements	1x PCIe Gen 2 (x16 electrically)	1x PCIe Gen 2 (x16 electrically)
Power Supply of Host System <sup>1</sup>	min. 750 W	min. 750 W
Min. RAM of Host System <sup>2</sup>	24 GB	24 GB

<sup>1</sup>**Important:** The specifications shown assume that only one adapter is plugged into the machine. If you would like to plug in two or more adapters you will need a better power supply (1000W or above) as well as more RAM. Additionally, you need to provide sufficient cooling for the machine. Each Tesla card takes power from the PCI Express host bus as well as the 8-pin and the 6-pin PCI Express power connectors. This is an important consideration while selecting power supplies.

<sup>2</sup>The host system requires approximately 4 times as much memory as is available on the GPU cards. Although it is technically possible to use less memory than this recommendation, the simulation performance of larger models will suffer.

**CST assumes no liability for any problems caused by this information.**

Hardware Type	 NVIDIA Kepler K10 <sup>1</sup> (for Servers)	 NVIDIA Quadro K6000
Min. CST version required	2013 release	2013 SP 4
Number of GPUs	2	1
Max. Problem Size (Transient Solver)	approx. 80 million mesh cells	approx. 120 million cells
Form Factor	Dual-Slot PCI-Express	Dual Slot PCI-Express
Memory	8 GB GDDR5	12 GB GDDR5
Bandwidth	320 GB/s (160 GB/s per GPU)	288 GB/s
Single Precision Performance	4.6 TFlops	5.2 TFlops
Double Precision Performance	0.2 TFlops	1.7 TFlops
Power Consumption	225 W (max.)	225 W (max.)
PCI Express Requirements	1x PCIe Gen 3 (x16 electrically)	1x PCIe Gen 3 (x16 electrically)
Power Supply of Host System <sup>2</sup>	min. 750 W	min. 750 W
Min. RAM of Host System <sup>3</sup>	32 GB	48 GB

<sup>1</sup> The double precision performance of this GPU device is poor, thus, it can't be recommended for PIC-solver, A-Solver and I-solver (double precision) simulations.

<sup>2</sup>**Important:** The specifications shown assume that only one adapter is plugged into the machine. If you would like to plug in two or more adapters you will need a better power supply (1000W or above) as well as more RAM. Additionally, you need to provide a sufficient cooling for the machine. Each Tesla card takes power from the PCI Express host bus as well as the 8-pin and the 6-pin PCI Express power connectors. This is an important consideration while selecting power supplies.

<sup>3</sup>The host system requires approximately 4 times as much memory as is available on the GPU cards. Although it is technically possible to use less memory than this recommendation, the simulation performance of larger models will suffer.

**CST assumes no liability for any problems caused by this information.**

Hardware Type	 NVIDIA Tesla K40m/K40c (for Servers/Workst.)	 NVIDIA Tesla K80 (for Servers)
<b>Min. CST version required</b>	<b>2013 SP 5</b>	<b>2014 SP 6</b>
Number of GPUs	1	2
Max. Problem Size (Transient Solver)	approx. 120 million mesh cells	approx. 240 million mesh cells
Form Factor	Dual-Slot PCI-Express	Dual Slot PCI-Express
Memory	12 GB GDDR5	24 GB GDDR5
Bandwidth	288 GB/s	480 GB/s (240 GB/s per GPU)
Single Precision Performance <sup>1</sup>	5.04 TFlops	8.73 TFlops
Double Precision Performance <sup>1</sup>	1.68 TFlops	2.91 TFlops
Power Consumption	225 W (max.)	300 W (max.)
PCI Express Requirements	1x PCIe Gen 3 (x16 electrically)	1x PCIe Gen 3 (x16 electrically)
Power Supply of Host System	min. 750 W	min. 750 W
Min. RAM of Host System <sup>2</sup>	48 GB	96 GB

<sup>1</sup> Measured with BOOST enabled

<sup>2</sup>The host system requires approximately 4 times as much memory as is available on the GPU cards. Although it is technically possible to use less memory than this recommendation, the simulation performance of larger models will suffer.

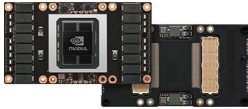

**CST assumes no liability for any problems caused by this information.**

Hardware Type	 NVIDIA Tesla M60 <sup>1</sup> (for Servers/Workst.)	 NVIDIA Tesla M40 <sup>1</sup> (for Servers)
Min. CST version required	2016 SP 4	2016 SP 4
Number of GPUs	2	1
Max. Problem Size (Transient Solver)	approx. 160 million mesh cells	approx. 240 million mesh cells
Form Factor	Dual-Slot PCI-Express	Dual Slot PCI-Express Passive Cooling
Memory	16 GB GDDR5 (8 GB x 2)	24 GB GDDR5
Bandwidth	320 GB/s (160 GB/s per GPU)	288 GB/s
Single Precision Performance	9.64 TFlops	6.84 TFlops
Double Precision Performance	0.301 TFlops	0.213 TFlops
Power Consumption	300 W (max.)	250 W (max.)
PCI Express Requirements	1x PCIe Gen 3 (x16 electrically)	1x PCIe Gen 3 (x16 electrically)
Power Supply of Host System	min. 750 W	min. 750 W
Min. RAM of Host System <sup>2</sup>	64 GB	96 GB

<sup>1</sup> The double precision performance of this GPU device is poor, thus, it can't be recommended for PIC-solver, A-Solver and I-solver (double precision) simulations.

<sup>2</sup>The host system requires approximately 4 times as much memory as is available on the GPU cards. Although it is technically possible to use less memory than this recommendation, the simulation performance of larger models will suffer.

**CST assumes no liability for any problems caused by this information.**

Hardware Type	 <p>NVIDIA Tesla P100 Chip (for Servers)</p>	 <p>NVIDIA Tesla P100 PCIe<sup>1</sup> (for Servers)</p>
<b>Min. CST version required</b>	<b>2017 release</b>	<b>2017 release</b>
Number of GPUs	1	1
Max. Problem Size (Transient Solver)	approx. 160 million mesh cells	approx. 160 / 120 million mesh cells
Form Factor	Chip Passive Cooling	Dual-Slot PCI-Express Passive Cooling
Memory	16 GB CoWoS HBM2	16 / 12 GB CoWoS HBM2
Bandwidth	732 GB/s	732 GB/s / 549 GB/s
Single Precision Performance <sup>2</sup>	10.6 TFlops	9.3 TFlops
Double Precision Performance <sup>2</sup>	5.3 TFlops	4.7 TFlops
Power Consumption	300 W (max.)	250 W (max.)
System interface	NVIDIA NVLink	1x PCIe Gen 3 (x16 electrically)
Power Supply of Host System	min. 750 W	min. 750 W
Min. RAM of Host System <sup>3</sup>	64 GB	64 GB



<sup>1</sup> The 12 GB version has about 25 percent less performance compared to the 16 GB version.

<sup>2</sup> Measured with BOOST enabled

<sup>3</sup>The host system requires approximately 4 times as much memory as is available on the GPU cards. Although it is technically possible to use less memory than this recommendation, the simulation performance of larger models will suffer.

**CST assumes no liability for any problems caused by this information.**



Hardware Type	 NVIDIA Quadro GP 100 (for Workstations)	 NVIDIA Quadro P6000 <sup>1</sup> (for Workstations)
Min. CST version required	2017 SP 2	2017 SP 2
Number of GPUs	1	1
Max. Problem Size (Transient Solver)	approx. 160 million mesh cells	approx. 240 million mesh cells
Form Factor	Dual-Slot PCI-Express	Dual-Slot PCI-Express
Memory	16 GB HBM2	24 GB GDDR5X
Bandwidth	720 GB/s	432 GB/s
Single Precision Performance <sup>2</sup>	10.3 TFlops	12.0 TFlops
Double Precision Performance <sup>2</sup>	5.2 TFlops	0.2 TFlops
Power Consumption	300 W (max.)	300 W (max.)
System interface	1x PCIe Gen 3 (x16 electrically)	1x PCIe Gen 3 (x16 electrically)
Power Supply of Host System	min. 750 W	min. 750 W
Min. RAM of Host System <sup>3</sup>	64 GB	96 GB

<sup>1</sup> The double precision performance of this GPU device is poor, thus, it can't be recommended for PIC-solver, A-Solver and I-solver (double precision) simulations.

<sup>2</sup> Measured with BOOST enabled

<sup>3</sup>The host system requires approximately 4 times as much memory as is available on the GPU cards. Although it is technically possible to use less memory than this recommendation, the simulation performance of larger models will suffer.

**CST assumes no liability for any problems caused by this information.**

Hardware Type	 NVIDIA Quadro M6000 <sup>1</sup> (for Workstations)	 NVIDIA Quadro M6000 24 GB <sup>1</sup> (for Workstations)
Min. CST version required	2015 SP 4	2016 SP 4
Number of GPUs	1	1
Max. Problem Size (Transient Solver)	approx. 120 million mesh cells	approx. 240 million mesh cells
Form Factor	Dual-Slot PCI-Express	Dual Slot PCI-Express
Memory	12 GB GDDR5	24 GB GDDR5
Bandwidth	317 GB/s	317 GB/s
Single Precision Performance	6.8 TFlops	6.8 TFlops
Double Precision Performance	0.2 TFlops	0.2 TFlops
Power Consumption	300 W (max.)	300 W (max.)
PCI Express Requirements	1x PCIe Gen 3 (x16 electrically)	1x PCIe Gen 3 (x16 electrically)
Power Supply of Host System	min. 750 W	min. 750 W
Min. RAM of Host System <sup>2</sup>	48 GB	96 GB

<sup>1</sup> The double precision performance of this GPU device is poor, thus, it can't be recommended for PIC-solver, A-Solver and I-solver (double precision) simulations.

<sup>2</sup>The host system requires approximately 4 times as much memory as is available on the GPU cards. Although it is technically possible to use less memory than this recommendation, the simulation performance of larger models will suffer.

CST assumes no liability for any problems caused by this information.