

Guide to Building Your Own Tesla Personal Supercomputer System

This guide is to help you build a Tesla Personal Supercomputer. If you have experience in building systems/workstations, you may want to build your own system. Otherwise, the easiest thing is to buy an off-the-shelf Tesla Personal Supercomputer from one of [these resellers](#).

As with building any system, it is at your own risk and responsibility. There are many possible components to choose from when you build such a system. NVIDIA provides general guidance, but cannot test every configuration and combination of components.

Minimum Specifications of Main Components

- 3x Tesla C1060
- Quad-core CPU: 2.33 GHz (Intel or AMD)
- 12 GB of system memory (4GB per Tesla GPU)
- Linux 64-bit or Windows XP 64-bit
- System acoustics < 45 dBA
- 1200 W power supply

Example of Complete 4 Tesla C1060 System Configuration

	4 Tesla C1060 Configuration
Motherboard	Foxconn Destroyer nForce 780a
Tesla GPUs	4x Tesla C1060
CPU	AMD Phenom 9850 2.5 GHz quad-core
Memory	16 GB (4x 4GB) DDR2 DIMMs from G.Skill
Power Supply	Coolmax CUQ-1350B 1350W
Case	Lian Li PC-P80
Hard drive	640 GB
DVD drive	DVD burner
CPU fan, heat sink	For the AMD Phenom
Total Retail Price	~\$ 8500*

* Prices obtained from online retailers in the US

3 Tesla C1060 Configurations

	3 Tesla Configuration	3 Tesla Configuration
Motherboards	Foxconn Destroyer 780a MSI K9A2 Platinum 790FX	Intel SkullTrail D5400XS
Tesla GPUs	3x Tesla C1060	3x Tesla C1060s
Display GPU	1x Quadro FX or NVS	1x Quadro NVS or FX (single slot)
Memory	12 GB DDR2 (3x 4GB DDR2 DIMMs)	12 GB DDR2 (3x 4GB DDR2 DIMMs) for SkullTrail
Power Supply	1200W	1200W
Note	--	Slot 0: Tesla C1060 Slot 1: Tesla C1060 Slot 2: Quadro NVS or FX (single slot) Slot 3: Tesla C1060

Motherboards

The Tesla C1060 computing processor is a dual-wide PCI-e x16 Gen2 board. It can be used in a Gen 1 PCI-e x16 slot as well, but this leads to lower system bandwidth between the CPU and GPU, which may impact application performance (depends on the application).

So, you need to use motherboards that have 3 or 4 PCI-e x16 slots that are dual slot apart from each other.

Motherboards that have 4x dual-wide PCI-e x16 slots include:

- Foxconn Destroyer 780a
 - This motherboard has a on-board NVIDIA GPU
 - With 4x Tesla C1060s installed, the PCI-e slots operate as PCI-e x8 Gen2 bandwidth
- MSI K9A2 Platinum 790FX
 - Works with 3x Tesla C1060s + 1 Quadro FX5800 (for display and compute)
 - With 4x GPU boards installed, the PCI-e slots operate as PCI-e x8 Gen2 bandwidth

There are several motherboards that have 3x dual-wide PCI-e x16 slots (including the two above). An example is:

- Intel D5400XS (Skulltrail)
 - 3x PCI-e x16 Gen1 dual-wide slots + 1 PCI-e x16 Gen1 single-wide slot
 - Put one Tesla C1060 board each in PCI-e slot 0, slot 1, and slot 3
 - Put Quadro FX1700 or FX3700 (single slot graphics cards) in PCI-e slot 2

CPUs

Choice of CPU is determined by the motherboard you use. We recommend that you use at least a 2.33 GHz quad-core CPU such as:

- Intel Core 2 or Xeon or Core i7 (Bloomfield) quad-core
- AMD Phenom or Opteron quad-core

System Memory

Since each Tesla C1060 has 4GB of GPU memory, we recommend 4GB of system memory per Tesla C1060. So, for a 3x Tesla C1060 system, include at least 12 GB of system memory and for 4x Tesla C1060 system, configure with 16GB of system memory.

Most of the motherboards have only 4 DIMM slots, therefore you should use 4GB DDR2 DIMMs to get 12GB or 16GB of system memory (4GB DDR3 DIMMs are currently not broadly available).

Power Supplies

There are several choices for power supplies. Choose a good power supply that is rated at least 1200 Watts or better still 1350 Watts. An example is the Coolmax CUQ-1350B 1350W power supply. Note, each Tesla C1060 requires either two 6-pin power connectors or one 8-pin power connector. This is important consideration while selecting power supplies.

Hard Drive, DVD, other components

Hard and DVD drive choices are left to you. In general, it is best to have at least 160GB of hard drive.

Computer Cases / Chassis

The case / chassis choice is important, since when you plug in four Tesla C1060 boards in, it needs a case that has 8 slots (which is larger than the usual ATX case).

Commercially available off-the-shelf cases with 8+ slots are

- Lian-Li PC-P80
- ThermalTake ArmorPlus
- ABS Canyon 695

It may be possible to take a 7-slot chassis and cut an extra 8th slot to accommodate the four dual-wide GPU boards. Try at your own risk!

System Cooling

Some system cases (like the Lian-Li) come with system fans of their own. In general, we recommend at least one case fan placed so that it is blowing air on the side of the Tesla boards (i.e., blowing air directly at the motherboard). The key is to keep the ambient temperature inside the case below 45C.

Operating Systems

We recommend either Linux 64-bit or Window XP 64-bit for best operation of this high-memory, high-performance system.

Verifying your System

Once you assemble your system and install an operating system on it, please download the CUDA driver, CUDA toolkit and optionally the CUDA SDK examples from [CUDA Zone](#).

After you download the CUDA driver, toolkit and SDK examples, run the following commands from the CUDA SDK:

- `deviceQuery`
 - This will report the number of Tesla GPUs in the system
- `bandwidthTest --memory=pinned --device=N`
 - Run for each C1060, where N=0, 1, 2, 3 for 4 C1060s in the system
 - This will report the PCI-E bandwidth to and from the CPU and each of the GPUs
 - Peak PCI-E x16 Gen2 bandwidth is between 5 and 6 GBytes/sec. PCI-E x16 Gen1 and PCI-E x8 Gen2 are about half that.
- `nbody --benchmark --n=131072 --device=N`
 - Run simultaneously for as many instances as Tesla GPUs in your system.
 - That is, for 4 Tesla C1060s, run 4 instances with N=0,1,2,3 simultaneously
 - This will run the nbody program on all the Tesla GPUs

We will roll out a test kit soon that will automate this and also serve as a good stability test for your system.

Reporting problems/questions

NVIDIA offers no direct support to individuals building their own Tesla Personal Supercomputers. However, we recommend using the [CUDA forums](#) to ask other CUDA developers about their experiences in building these systems.