

By: [Enrico Botezatu](#), Hardware Editor

## **Nvidia Tesla GPGPU Shows Up in Bull's NovaScale Supercomputer**

*The 1068-CPU machine is placed second in the top 500 supercomputers list*

System manufacturer Bull is reportedly working on a top-notch supercomputer, built around 1080 8-core Nehalem processors. The amount of simultaneous execution threads may be overwhelming, but the big news is that the supercomputer will feature 96 GPU cores delivered by Nvidia itself. Currently known as the Bull NovaScale, the supercomputer will be able to deliver 300 Teraflops of computing power. This means that Bull's new creation will be placed second on the top 500 supercomputers list, right after IBM's BlueGene/L machine. The original Bull NovaScale supercomputer was designed to be powered by 1068 nodes, for a whooping amount of 8544 processing cores. However, the Nehalem cores could only deliver a peak floating point performance of 192 Teraflops. According to the French tech website PC Inpact, the computer's design has been updated with eight nodes of Nvidia Tesla S900 GPGPU cards featuring 96 GPUs in the GT-200 family. The Tesla cards are extremely powerful and bring the extra 102 Teraflops of computing power. This way, Bull's supercomputer is one of the most powerful supercomputers, trailing IBM's BlueGene/L that peaks at 596 Teraflops using 212,992 processing cores. Nvidia's Tesla project is a third line of products (apart from tech GeForce and Quadro families), built around solid graphics cores, able to deliver the same computing power as 40 x86 processors. More than that, the graphics cores are also energy-efficient, and a fully-fledged Tesla card is allegedly taking up only 170 watts of power. In contrast, 40 x86 chips would take more than 1600 watts. The company estimates that each Tesla card can deliver 1.1 Teraflops of computing power. Overall, the whole 96 GPUs can achieve about 54 percent of the performance delivered by 1068 8-core Nehalem processors, but the comparison should start here, because computing power is not everything. For instance, supercomputers make heavy use of memory bandwidth, and when it comes to it, the Tesla GPUs don't quite live up to the Nehalems' performance.