

DECEMBER 7, 2011 • WASHINGTON, D.C.  
Navigate the Technology Landscape in One Day.

Register Today!

An IDG Enterprise Event



COMPUTERWORLD

Print Article Close Window

## Supercomputer with ARM processors being built in Spain

A prototype supercomputer is gunning to be the world's most power-efficient with Nvidia's Tegra 3 chips

Agam Shah

November 14, 2011 ([IDG News Service](#))

With an eye to bringing power efficiency to high-performance systems, a prototype supercomputer with quad-core ARM processors is being built at the Barcelona Supercomputing Center in Spain.

The prototype server is designed to handle complex workloads and will incorporate Nvidia's quad-core Tegra 3 chips, which were introduced earlier this month for smartphones and tablets. The 1,000 Tegra 3 chips have ARM CPUs and are being paired with discrete graphics processors from Nvidia to speed up scientific and math calculations.

The server could provide inroads for ARM to enter the high-performance computing market, which is dominated by competitors including Intel, Advanced Micro Devices, IBM and Oracle. Today, ARM processors are found in most smartphones and tablets, but have virtually no server market presence.

However, there is a growing interest in using the ARM processors in servers that can deliver high performance while overcoming power constraints, said Steve Scott chief technology officer of the Tesla business unit at Nvidia.

"We're very interested in ARM entering the HPC ecosystem," Scott said.

The Tegra 3 prototype supercomputer won't deliver petaflops of peak performance like some of the fastest computers in the world, Scott said. But the server could make the [Green500](#) list, which measures the most energy efficient supercomputers in the world.

Nvidia could not provide the peak performance numbers or the number of graphics processors in the Tegra 3 server, but matching thousands of ARM cores with its GPUs will help handle scientific workloads while reducing power consumption and computing overhead, Scott said. The BSC already has a prototype ARM server that has 256 Tegra 2 dual-core chips.

By bringing in an ARM CPU, Nvidia is expanding its supercomputing profile, which now largely revolves around its Tesla graphics processors used in supercomputers for complex calculations. A supercomputer being built by the Oak Ridge National Laboratory will pair Tesla GPUs with Advanced Micro Devices' 16-core Opteron CPUs to deliver 20 petaflops of performance. That is faster than Japan's K, which delivers a performance of 8 petaflops and is the world's fastest computer, according to a [list of the world's fastest supercomputers issued by Top500 in June](#).

Nvidia also joins a small contingent of companies experimenting with ARM as an alternative to x86 processors from Intel and AMD. Earlier this month Hewlett-Packard announced server designs with a chip from Calxeda, which includes a quad-core ARM processor and consumes as little as 1.5 watts of power. While ARM may lag x86 server processors on raw performance, analysts have said that a congregation of thousands of ARM CPUs could deliver better performance-per-watt on lightweight and fast-moving workloads such as processing volumes of web transactions.

The prototype Tegra servers are part of an effort largely funded by the European Commission to make systems that can deliver exascale performance while consuming 15 to 30 times less energy than current servers. The Tegra 3 prototype server is yet to undergo final approval, though it has been used for software development. The server specifications for the new system will be defined next year. This project, called Mont Blanc, is being coordinated by the Barcelona Supercomputing Center and has a budget of over 14 million (\$19 million).

But as a newcomer to the server market, ARM faces many challenges. Most server software is designed to run on x86 chips, so code would need to be rewritten to run on ARM processors. Current ARM processors also lack error-correction features and have 32-bit addressing, which limits the memory ceiling to 4GB.

Challenges notwithstanding, Nvidia had to start somewhere, and wants to be a leader in pushing ARM in servers and supercomputers going ahead.

"We understand it's not going to happen overnight," Scott said. "It will take a number of years to develop."

Nvidia is developing a new ARM-based CPU code-named Project Denver, which will go into smartphones, tablets, PCs and supercomputers. It's possible that Nvidia will integrate ARM cores in Tesla products in the long run, Scott said.

"We'll look at it and integrate it when time is right," Scott said.

ARM recently also announced a new 64-bit architecture, which includes many server-specific features. That will get rid of a key roadblock for ARM to make its presence felt in the HPC market.

"To do it right you need to use 64-bit architecture," Scott said.