



The Register On-demand webcast
Solutions for the Virtual Era [View Now](#)  
Find out what makes for a successful mainstream virtualization deployment.

Original URL: http://www.theregister.co.uk/2010/08/09/darpa_extremescale_awards/

DARPA awards petaflops super challenge squeeze

Small scale ExtremeScale

By [Timothy Prickett Morgan](#)

Posted in HPC, 9th August 2010 20:09 GMT

[Free whitepaper – The Register Guide to Enterprise Virtualization](#)

If you were thinking about entering the US Defense Advanced Research Project Agency's ExtremeScale supercomputing challenge [issued in March](#) [1], you missed your chance. DARPA's awarded grants to four design teams, plus on other that'll run benchmarks one the HPC prototypes.

The heavy hitters in the HPC community as well as in academia in the US were awarded grants to proceed on designing prototype machines for the Ubiquitous High Performance Computing ExtremeScale challenge.

The Register On-Demand webcast
The Virtual Desktop Unleashed
Solving the real-world problems
Join the experts as they spell out the common pitfalls of virtualisation.
In association with 

There are a lot of different goals with the ExtremeScale challenge, as we detailed in March, but the upshot is that DARPA wants a petaflops supercomputer, including its networking, storage, and compute elements as well as its cooling, to be crammed in a space a little larger than a standard server rack - four inches wide by 78 inches high and 40 inches deep - and consume only 57 kilowatts to power and cool the device.

The machine has to deliver a peak petaflops of performance and 50 gigaflops per watt sustained power efficiency while running the Linpack Fortran number-crunching test. The system has to be able to do

single-precision and double-precision math, 16-bit, 32-bit, and 64-bit integer math, and chew on a streaming sensor array like you might have scanning the skies of incoming missiles.

DARPA has also asked for parallel programming to be made easier on these machines than it currently is on massively parallel CPU or CPU-GPU hybrids, for multiple ExtremeScale units to be linked together, and for the machines to run five specific workloads.

These include a massive streaming sensor data problem, a large dynamic graph-based informatics problem, a decision problem that includes search with hypothesis testing and planning, and two as-yet-unnamed applications culled from the IT bowels of the Department of Defense.

The IT community players that have been granted ExtremeScale UHPC contracts today DARPA include one team led Intel and another from Nvidia. Intel has not yet detailed who is on its team or what its award is, but Nvidia said in a statement it has tapped supercomputer maker Cray, the US Department of Energy's Oak Ridge National Laboratory, and six universities to be on its team.

Nvidia said it was awarded \$25m to pursue the ExtremeScale impossible dream. Nvidia is still working out the final details on what those six universities will be doing.

Oak Ridge and Nvidia [have already been working on hybrid CPU-GPU compute clusters](#) [2], and last October announced the DOE had kicked in funds to study the use of Nvidia Tesla GPUs in x64 clusters.

Oak Ridge is also, of course, where the 1.76 petaflops Jaguar Cray Opteron-Linux cluster is housed as well. So, in essence, DARPA is paying for work that the Nvidia team was already pursuing piecemeal, but shaping it for a particular physical environment and code set. Anything worth selling once is worth selling two, or maybe three times.

Intel had aspirations for its ill-fated Larrabee GPU, but [something went terribly wrong](#) [3] with that project and Larrabee was killed off as a discrete graphics processor in May and then the core guts of the project [was reinvented](#) [4] as the Knights GPU co-processor in June.

First knights

The first of the Knights GPU co-processors, Knights Ferry, uses a 32-core x64 chip with attached vector units for each core and a shared, coherent cache running across all cores. The chip inside the Knights Ferry GPU co-processor is code-named Aubrey Isle.

The next generation of chips, called Knights Corner, will have 50 of these hybrid x64-vector cores. The interesting thing about this is that Intel may not need to put Xeon chips inside its ExtremeScale boxes at all, except for maybe a head node and a storage node.

Ditto for Nvidia is it gets around to embedding a proper x64 processor inside of its GPUs.

Whatever Intel and Nvidia are planning for their prototype systems, they are not saying yet, and that stands to reason since the prototypes are not due until 2018. Whatever Intel does, it surely will not be a bunch of Xeon processors circa 2018. The compute density of standard x64 processors, while no doubt impressive, will not be enough to meet DARPA's goals eight years from now.

DARPA also awarded ExtremeScale contracts to MIT's Computer Science and Artificial Intelligence Laboratory and to Sandia National Laboratory, another of the big nuke labs run by the DOE that consumes a huge amount of flops and budget dollars.

Sandia is noteworthy in that it was also where Intel plunked its ground-breaking one teraflops massively parallel ASCI Red super, and where Cray got the development dough to create the massively parallel Red Storm Opteron-Linux super.

Red Storm is, of course, the prototype that eventually became the XT3, XT4, XT5, and now XT6 and XE6 supercomputer lines from Cray. Egos are always involved when it comes to supercomputers, and you can bet that Sandia wants to have the first ExtremeScale box, too.

The mystery about the DARPA ExtremeScale awards is why SGI was not among those chosen - perhaps SGI is on the Intel team? In early June, SGI [announced](#) [5] Project Mojo, its own plan to pack a petaflops into a server rack - and using various kinds of accelerators.

The SGI project is not trying to cram the whole data center into a rack as the ExtremeScale project does, but SGI seemed pretty confident that it could get one petaflops of computing into a rack - and in a number of different ways - by the middle of 2011 or so. Condensing this down to get storage and cooling into the rack should be something SGI is interested in.

And what about IBM and Hewlett-Packard? Don't they accept ExtremeScale challenges?

DARPA said it has chosen a team of boffins at the Georgia Institute of Technology to create the application, benchmark, and metrics team to analyze how well or poorly the four ExtremeScale prototypes that Intel, Nvidia, MIT, and Sandia put together.

DARPA had not returned phone calls for more information on the ExtremeScale systems and the monetary awards each team received by press time. ®

Links

1. http://www.theregister.co.uk/2010/03/03/darpa_uhpc_extremescale_challenge/
2. http://www.theregister.co.uk/2009/10/01/oak_ridge_fermi/
3. http://www.theregister.co.uk/2010/05/25/intel_isc_announcements/
4. http://www.theregister.co.uk/2010/06/01/intel_knights_co_processor/
5. http://www.theregister.co.uk/2010/06/03/sgi_petaflops_cabinet/

Related stories

[Shipping container fuels Australia's bid to be supercomputing power](#) (9 August 2010)

http://www.theregister.co.uk/2010/08/09/ivec_hp_supercomputer_pod_acquisition/

[Tilera to stuff 200 cores onto single chip](#) (22 June 2010)

http://www.theregister.co.uk/2010/06/22/tilera_quanta_server_stratton/

[Project Mojo – the server rack that packs petaflops](#) (3 June 2010)

http://www.theregister.co.uk/2010/06/03/sgi_petaflops_cabinet/

[DARPA asks you to cram petaflops super into single rack](#) (3 March 2010)

http://www.theregister.co.uk/2010/03/03/darpa_uhpc_extremescale_challenge/

[DARPA: Can we have a one-cabinet petaflop supercomputer?](#) (26 June 2009)

http://www.theregister.co.uk/2009/06/26/darpa_supercomputer_slim_size/

© Copyright 1998–2010