

TECHNOLOGY

Joint effort nabs next wave of US supercomputers

National laboratories collaborate to purchase top-flight machines.

BY ALEXANDRA WITZ

Once locked in an arms race with each other for the fastest supercomputers, US national laboratories are now banding together to buy their next-generation machines.

On 14 November, the Oak Ridge National Laboratory (ORNL) in Tennessee and the Lawrence Livermore National Laboratory in California announced that they will each acquire a next-generation IBM supercomputer that will run at up to 150 petaflops. This means that the machines can perform 150 million billion floating-point operations per second, at least five times as fast as the current leading US supercomputer, the Titan system at the ORNL.

The new supercomputers, which together will cost US\$325 million, should enable new types of science for thousands of researchers who model everything from climate change to materials science to nuclear-weapons performance.

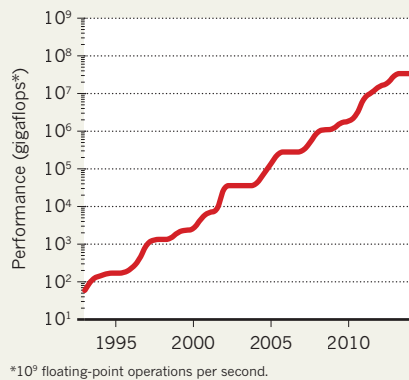
“There is a real importance of having the larger systems, and not just to do the same problems over and over again in greater detail,” says Julia White, manager of a grant programme that awards supercomputing time at the ORNL and Argonne National Laboratory in Illinois. “You can actually take science to the next level.” For instance, climate modellers could use the faster machines to link together ocean and atmospheric-circulation patterns in a regional simulation to get a much more accurate picture of how hurricanes form.

Building the most powerful supercomputers is a never-ending race. Almost as soon as one machine is purchased and installed, lab managers begin soliciting bids for the next one. Vendors such as IBM and Cray use these competitions to develop the next generation of processor chips and architectures, which shapes the field of computing more generally.

In the past, the US national labs pursued separate paths to these acquisitions. Hoping to streamline the process and save money, clusters of labs have now joined together to put out a shared call — even those that perform classified research, such as Livermore. “Our missions differ, but we share a lot of commonalities,” says Arthur Bland, who heads the ORNL computing facility.

NEXT STOP EXAFLOP

The speed of the world's most powerful supercomputer has grown more than five orders of magnitude in the past two decades.



In June, after the first such coordinated bid, Cray agreed to supply one machine to a consortium from the Los Alamos and Sandia national labs in New Mexico, and another to the National Energy Research Scientific Computing (NERSC) Center at the Lawrence Berkeley National Laboratory in Berkeley, California. Similarly, the ORNL and Livermore have banded together with Argonne.

The joint bids have been a learning experience, says Thuc Hoang, programme manager for high-performance supercomputing research and operations with the National Nuclear Security Administration in Washington DC, which manages Los Alamos, Sandia and Livermore. “We thought it was worth a try,” she says. “It requires a lot of meetings about which requirements are coming from which labs and where we can make compromises.”

At the moment, the world's most powerful supercomputer is the 55-petaflop Tianhe-2 machine at the National Super Computer Center in Guangzhou, China. Titan is second, at 27 petaflops. An updated ranking of the top 500 supercomputers was announced on 18 November at the 2014 Supercomputing Conference in New Orleans, Louisiana.

When the new ORNL and Livermore supercomputers come online in 2017, they will almost certainly vault to near the top of the list, says Barbara Helland, facilities-division director of the advanced scientific computing research programme at the

Department of Energy (DOE) office of science in Washington DC.

The new supercomputers, to be called Summit and Sierra, will be structurally similar to the existing Titan supercomputer. They will combine two types of processor chip: central processing units, or CPUs, which handle the bulk of everyday calculations; and graphics processing units, or GPUs, which generally handle three-dimensional computations. Combining the two means that a supercomputer can direct the heavy work to GPUs and operate more efficiently overall. And because the ORNL and Livermore will have similar machines, computer managers should be able to share lessons learned and ways to improve performance, Helland says.

Still, the DOE wants to preserve a little variety. The third lab of the trio, Argonne, will be making its announcement in the coming months, Helland says, but it will use a different architecture from the combined CPU-GPU approach. It will almost certainly be like Argonne's current IBM machine, which uses a lot of small but identical processors networked together. The latter approach has been popular for biological simulations, Helland says, and so “we want to keep the two different paths open”.

Ultimately, the DOE is pushing towards supercomputers that could work at the exascale, or 1,000 times more powerful than the current petascale (see ‘Next stop exaflop’). Those are expected around 2023. But the more power the DOE labs acquire, the more scientists seem to want, says Katie Antypas, head of the NERSC's services department.

“There are entire fields that didn't used to have a computational component to them,” such as genomics and bioimaging, she says. “And now they are coming to us asking for help.” ■

CORRECTION

The News story “‘Forgotten’ NIH smallpox virus languishes on death row” (*Nature* **514**, 544; 2014) wrongly said that the WHO Advisory Committee on Variola Virus Research agreed to commission a report on the bioterror threat from synthesized smallpox — that report was actually commissioned before the committee met.

SOURCE: TOP500.ORG