

Home (<http://www.hpwire.com/>) | News (<http://www.hpwire.com/news.html>) |
Features (<http://www.hpwire.com/hpwire/topic/featured?limit=25>) |
Blogs (<http://www.hpwire.com/hpwire/topic/hpcblogs?limit=10>) | HPC Markets (<http://markets.hpwire.com/>) |
Whitepapers (<http://www.hpwire.com/whitepapers/topic/hpcwhitepapers?limit=10>) |
Multimedia (<http://www.hpwire.com/multimedia/topic/hpcmedia?limit=10>) |
Events (<http://www.hpwire.com/events/topic/hpcevents/>) | Job Bank (<http://www.hpwire.com/jobs.html>) |
Special Features (<http://www.hpwire.com/specialfeatures/>) | Photo Gallery (<http://www.hpwire.com/specialfeatures/photos>) |

June 19, 2012

TOP500 Gets Dressed Up with New Blue Genes

Michael Feldman

At 16 petaflops, Sequoia recaptures the number one spot for the US.

The 39th TOP500 list was released today at the International Supercomputing Conference in Hamburg, Germany, with a new machine at the top. Sequoia, an IBM Blue Gene/Q machine, delivered a world record 16 petaflops on Linpack, knocking RIKEN's 10-petaflop K Computer into second place. The Japanese K system had held the TOP500 title for a year.



Sequoia, which is housed at Lawrence Livermore National Lab, will provide the NNSA its most advanced simulation platform for maintaining the nuclear weapons stockpile of the US. In its spare time, it will also run unclassified codes for open science research.

The 96-rack Sequoia houses 1.6 million cores, another TOP500 record, and 1.6 petabytes of memory. Peak petaflops is a whopping 20.1 petaflops. The machine is one of six Blue Gene/Q systems of a petaflop or more deployed over the last six months.

Compared to the November 2011 list, [when there was no turnover in the top 10](http://www.hpwire.com/hpwire/2011-11-14/top500_shows_signs_of_stagnation.html) (http://www.hpwire.com/hpwire/2011-11-14/top500_shows_signs_of_stagnation.html),

this time around, there are six brand new machines, plus one, Jaguar, that has benefitted from an upgrade to faster processors. Besides four new Blue Gene/Q's (Sequoia, Mira, Fermi, and JuQUEEN), there is SuperMUC, an IBM iDataPlex cluster at Leibniz Rechenzentrum in Germany, and Curie, a Bull supercomputer installed at the French Atomic Energy Commission (CEA).

The new top 10 looks like this:

1. 16.3 petaflops, Sequoia, United States
2. 10.5 petaflops, K computer, Japan
3. 8.2 petaflops, Mira, United States
4. 2.9 petaflops, SuperMUC, Germany
5. 2.6 petaflops, Tianhe-1A, China
6. 1.9 petaflops, Jaguar, United States
7. 1.7 petaflops, Fermi, Italy
8. 1.4 petaflops, JuQUEEN, Germany
9. 1.4 petaflops, Curie, France
10. 1.3 petaflops, Nebulae, China

Although the US has regained the TOP500 title -- the first time it has been in the top spot since 2009 -- just three of the top 10 are now based in the States, down from five machines, six months ago, continuing a trend that has resulted in more geographical parity. China, Japan, Germany, France and Italy all have supercomputers at the top of the list now.

Taking all 500 supercomputers into account, the US is still the dominant player with 252 systems, but that's down from 263 six months ago. China, is in second place with 68 systems, but it too has lost ground, shedding six since November. Japan (35 systems), the UK (25 systems), France (22 systems) and Germany(20 systems) are the only other nations with more than 10 machines on the list.

With each passing year, the TOP500 becomes a more exclusive club. The least performant machine (the 500th system) is now over 60 teraflops, a Linpack mark that would have earned it the top spot in 2004. Turnover was about average, with the list shedding 170 systems.

Meanwhile, aggregate performance continues its upward climb and is now at 123.4 petaflops, nearly doubling that of the November list, when it totaled 74.2 petaflops. A sizeable chunk of added flops was contributed by new machines that came in at a petaflop or better. Overall, the petaflop club doubled its membership over the last six months, growing from 10 to 20 systems.

From a vendor perspective, IBM cleaned up. The company is responsible for nearly half of the machines on the list, with 213. The next most popular vendor is HP, with 138 systems. Cray (26), Appro (19), Bull (16), SGI (16), and Dell (12), round out top six computer makers. Everyone else is in single digits.

It's even more skewed at the top, where IBM claims five of the top 10. As mentioned before, that's mainly the result of the new Blue Gene/Q installations. No other vendor has more than a single system in this upper tier.

The only area where IBM didn't dominate the field is in processor architecture. Here Intel is king, claiming a 78 percent share overall, split between its various Xeon generations. The latest E5 Xeons, despite being in production only three months, already claim a nine percent share.

GPUs and other accelerators are now installed in 58 systems, up from 39 six months ago. The vast majority of them (53) are using NVIDIA parts. AMD's ATI GPUs and IBM's PowerXCell 8i are installed on two systems, apiece, while Intel's MIC coprocessor made its debut on the TOP500 in an experimental cluster with pre-production Knights Corner chips.

On the interconnect front, InfiniBand now reigns as the most popular technology, with 209 systems, finally beating out Ethernet,

which is installed on 207 machines. The remaining 84 systems use a combination non-standard interconnects (custom, proprietary, Cray, etc.). Although small in number, these specialized networks are installed in systems that represent more than half (55 percent) of the TOP500's aggregate performance.

The next Linpack rankings in November should see many of these trends continue. The top of the list, as always, should be quite interesting, especially since at least a couple of new double-digit petaflop machines, powered by the latest accelerators, are scheduled to make their appearance. The Stampede system at TACC, will be powered by Intel's first Knights Corner coprocessor, while the Titan supercomputer at ORNL, will be featured with the new NVIDIA Kepler parts. If these deployments go as planned, we could, once again, see some major realignments in the top 10.



[\(http://twitter.com/hpcwire/\)](http://twitter.com/hpcwire/)

Share Options

-  [\(/xs/synd/hpcwire/m22121?dest=facebook\)](/xs/synd/hpcwire/m22121?dest=facebook)
-  [\(/xs/synd/hpcwire/m22121?dest=twitter\)](/xs/synd/hpcwire/m22121?dest=twitter)
-  [\(/xs/synd/hpcwire/m22121?dest=digg\)](/xs/synd/hpcwire/m22121?dest=digg)
-  [\(/xs/synd/hpcwire/m22121?dest=yahoo\)](/xs/synd/hpcwire/m22121?dest=yahoo)
-  [\(/xs/synd/hpcwire/m22121?dest=plus\)](/xs/synd/hpcwire/m22121?dest=plus)
-  [\(/xs/synd/hpcwire/m22121?dest=stumbleupon\)](/xs/synd/hpcwire/m22121?dest=stumbleupon)
-  [\(/xs/synd/hpcwire/m22121?dest=live\)](/xs/synd/hpcwire/m22121?dest=live)
-  [\(/xs/synd/hpcwire/m22121?dest=google\)](/xs/synd/hpcwire/m22121?dest=google)
-  [\(/xs/synd/hpcwire/m22121?dest=myspace\)](/xs/synd/hpcwire/m22121?dest=myspace)
-  [\(/xs/synd/hpcwire/m22121?dest=blogger\) \(javascript:;\)](/xs/synd/hpcwire/m22121?dest=blogger)

Subscribe

- [Subscribe to HPCwire \(/xs/register\)](/xs/register)

Discussion

There are 0 discussion items posted.

Join the Discussion

[\(/xs/register\)](/xs/register)



[Become a Registered User Today! \(/xs/register\)](/xs/register)

Registered Users Log in join the Discussion

Email Address:

Remember me

Password (case sensitive)

[Forgot Password? \(/xs/reqpass\)](/xs/reqpass)

Sponsored Links

[Appro and Intel® MIC Architecture, HPC Presentation at ISC 2012!](#)

[\(/ct/e/V7gsREP60tbpaferyJozZLcKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tbpaferyJozZLcKuw5XntEp2A6t1LiSVI=)

Come and see us at Birds of a Feather (BoF) session called "First of a Kind Experimental Cluster for Highly Parallel HPC Workloads" discussing how this collaboration will help develop applications for current and future co-processor computational system.

▶ [Big Data...is it part of your computing challenge?](#)

[\(/ct/e/V7gsREP60tZ599AZHTK6ArcKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tZ599AZHTK6ArcKuw5XntEp2A6t1LiSVI=)

▶ [8 Steps: Optimizing Cache Memory Access & Application Performance](#)

[\(/ct/e/V7gsREP60tbgzOIxKH15qbcKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tbgzOIxKH15qbcKuw5XntEp2A6t1LiSVI=)

▶ [Summit C220I.1F 2U Quad Server](#)

[\(/ct/e/V7gsREP60taoe_OR-](/ct/e/V7gsREP60taoe_OR-)

[DdSsreKuw5XntEp2A6t1LiSVI=\)](DdSsreKuw5XntEp2A6t1LiSVI=)

▶ [Use PGI to Accelerate Your Applications with Less Programming Effort](#)

[\(/ct/e/V7gsREP60tbC9J2IUSx4PLcKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tbC9J2IUSx4PLcKuw5XntEp2A6t1LiSVI=)

▶ [2X in 4 Weeks. Guaranteed.](#)

[\(/ct/e/V7gsREP60tZCpTOerWBTheKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tZCpTOerWBTheKuw5XntEp2A6t1LiSVI=)

▶ [Register to Win Monthly Prizes. Speed Up Your Datacenter](#)

▶ [Most Efficient Transparent way to Optimize Micro Resources](#)

[\(/ct/e/V7gsREP60tYWdallXiZ71LcKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tYWdallXiZ71LcKuw5XntEp2A6t1LiSVI=)

▶ [Acer introduces its new line of servers for technical computing](#)

[\(/ct/e/V7gsREP60tYjCQivskXFB7cKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tYjCQivskXFB7cKuw5XntEp2A6t1LiSVI=)

▶ [Never thought you could own a Cray?](#)

[\(/ct/e/V7gsREP60tala72DgvdL57cKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tala72DgvdL57cKuw5XntEp2A6t1LiSVI=)

▶ [OrangeFS Increases Performance and Broadens Access with New Releases](#)

[\(/ct/e/V7gsREP60tYoKJsMznWvoLcKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60tYoKJsMznWvoLcKuw5XntEp2A6t1LiSVI=)

▶ [Supermicro® FatTwin™ High-Capacity](#)

[\(/ct/e/V7gsREP60taImaNFISx6KrcKuw5XntEp2A6t1LiSVI=\)](/ct/e/V7gsREP60taImaNFISx6KrcKuw5XntEp2A6t1LiSVI=)

RSS Feeds