Software house Liqid is prime contractor in giant supercomputing awards worth $52m

OCTOBER 15 2020

By John Abbott

The US Department of Defense has embraced the concept of disaggregated composable infrastructure to dynamically mix the proportions of CPU, GPU, memory and storage for AI workloads. It recently awarded over $52m in contracts to composable infrastructure startup Liqid, which has just 60 employees.

451 Research®
Now a Part of

S&P Global Market Intelligence

Something a little unusual is going on when a small software firm wins two giant supercomputing awards as a prime contractor. The contracts, which together are worth over $52m, were awarded by the US Department of Defense Army Corps of Engineers to composable infrastructure startup Liqid, which led the two bids with its disaggregated infrastructure software layer. Rival bidders included HPE and Cray (which bid separately before HPE acquired Cray last year).

First to be announced (on September 29) was a $20.6bn award to supply a 17-petaflop composable supercomputer to be deployed at its Engineer Research and Development Lab (ERDC) in Vicksburg, Mississippi. The system will run Liqid's software on top of 86, 784 Intel Xeon Platinum 9200 computer cores, 536 NVIDIA A100 GPUs, 391 terabytes of memory, 4.2 petaflops of all-flash NVME-oF parallel file system storage and Mellanox HDR 200Gb/s InfiniBand adapters. If added to the current Top 500 Supercomputer List, it would rank as the 15th most powerful supercomputer in the world. A second contract was announced on October 1, this time worth $32m and involving two supercomputer systems to be installed at the US Army Supercomputing Research Center, Aberdeen Proving Ground in Maryland. The first will have 1,202 Intel 9200 CPUs and 280 A100 GPUs, the second 1,010 CPUs and 76 GPUs.

Liqid will lean on its partners Intel and NVIDIA to supply the hardware. But it seems to have been the software element that sealed the deal, particularly the ability to dynamically vary the proportion of resources – CPU, GPU, memory and NVMe storage – to better adapt to the performance requirements of graphics-intensive AI applications, which can be uneven and predictable. Liqid only recently moved beyond its previous reliance on PCIe as a fabric to support multiple fabrics, including Ethernet and (crucially for HPC) InfiniBand. Composable infrastructure – making compute, storage and networking infrastructure dynamically configurable on the fly – has been talked about by various vendors for a least five years now, but adoption so far has been limited. Its use as a central software platform for advanced, AI-accelerated HPC installations is the most solid endorsement of the composable concept to emerge to date.