

Big guns for hire: Supercomputing-in-the-cloud

Todd R. Weiss

13.07.2010 kl 18:15 | Computerworld (US)

For much of its existence, **Woodward Governor Co.** has been building fuel nozzles and other parts for jet aircraft engines the old-fashioned way -- carefully designing them, building physical prototypes out of raw materials, conducting tests and finally creating production parts. Typically, it took the Fort Collins, Colo., company weeks to produce workable plans for specialized new nozzle designs. Now, that time frame has been trimmed down to several hours by cutting out the steps needed to make physical prototypes.

Instead, Woodward is using powerful cloud-based supercomputers that can run complex software to quickly create and evaluate designs that can go right into production. In other words, the nozzle designs can go from creation to production 80% faster because the company no longer needs to first build models to prove the designs, a change that's projected to save the company about a half-million dollars each year.

"Woodward had some fairly robust computer workstations internally" to produce designs, but the company didn't have anything powerful enough to run large simulations, says Robert Graybill, an HPC consultant who worked with Woodward to help add on-demand supercomputer power to the company's production capabilities. Woodward's pilot project lasted 10 weeks, and additional testing will continue through October, followed by reviews of the project to determine if and how Woodward might use on-demand supercomputing in the future.

Graybill is the CEO and president of McLean, Va.-based **Nimbis Services Inc.**, an intermediary that helps match businesses that need more computing horsepower with HPC providers that can deliver it on an as-needed, pay-as-you-go basis. (Woodward IT representatives were not available for comment.)

Dan Olds, an analyst at Gabriel Consulting Group Inc. in Beaverton, Ore., says on-demand supercomputing is an idea that's blossoming. It's particularly useful for businesses that sometimes need processing help but can't afford to buy their own full-time, on-site HPC equipment.



Consultant Robert Graybill of Nimbis Services says on-demand supercomputing is "making a full circle" from the days of renting time on mainframes. "Cloud computing is really driven by users that need additional computing capacity."

"And it really opens the door for other users where this isn't even on their horizon yet," Olds says. "They could do these things on their own with the modest equipment they have, but it literally could take years." Alternately, users could simplify the problems they're trying to solve so their projects could run on systems they can afford, "but that might not give them the results they need," Olds says. On-demand supercomputing helps companies avoid both those compromises without breaking their budgets.

The market for such services has been fairly limited in the past, said Charles King, principal analyst at Pund-IT Inc. in Hayward, Calif. "But a weak, rocky economy makes on-demand services a lot more affordable and attractive than purchasing and maintaining a dedicated supercomputing cluster," he said in an e-mail.

Check on software compatibility

The biggest challenges are more likely related to software than hardware, King says. Many companies that use HPC systems have developed their own applications or have tweaked packaged applications for their own use. So how providers of on-demand supercomputing services work with customers to run their existing software is a key issue.

For its part, Woodward got started with on-demand supercomputing when it was invited to participate in a pilot project through the University of Southern California's Information Sciences Institute (ISI), where Graybill is the director of innovation. The goal was to explore how manufacturers could use HPC **cloud computing** services to improve their industrial design and modeling processes. The work was done under a \$3.67 million contract from the Defense Advanced Research Projects Agency (DARPA).

Woodward began accessing HPC services through an **IBM Computing On Demand** facility in Poughkeepsie, N.Y. It's a cluster of 256 System x3550 servers and 128 System x3450 servers with a total of 16GB of RAM, providing more than 19 trillion floating point operations per second (TFLOPS) of sustained performance, according to IBM. Woodward is able to use the extra computing power as needed, Graybill says.

The Woodward pilot project was one of four similar projects that DARPA funded in the past couple of years. The three other projects involved virtual metal forming for ACE Clearwater Enterprises in Torrance, Calif., chassis weight reduction for AlphaSTAR Corp. of Long Beach, Calif., and electromagnetic interference signature analysis for AltaSim Technologies Inc. of Columbus, Ohio.

Reducing waste

To date, the extra computing horsepower is helping Woodward produce new jet engine nozzle designs much more efficiently, while also reducing waste materials by 50%, according to figures from the pilot project. Costs were reduced by \$275,000 per engineer annually by using supercomputers to eliminate the need for physical modeling, saving more than \$500,000 per year, the study concluded.

"The access to cloud-based supercomputing really does change how you design your products," because supercomputing expands your available technologies, tools and resources, Graybill says. "It impacts your time to market and your quality in good ways."

And because it's able to buy the HPC time only when needed, Woodward doesn't have to deploy expensive on-site HPC clusters, which would bring additional energy costs and require maintenance by additional IT administrators, he explains. ☐

"It's sort of making a full circle from the days of buying time on a mainframe in the old days," he says. "This cloud computing is really driven by users that need additional computing capacity."

Similar kinds of public-private HPC projects have been undertaken through the **Ohio Supercomputer Center** (OSC), where businesses based in that state can get both computing capacity and expert advice to help improve their bottom lines.

"We have been helping with building economic development in Ohio, giving businesses from large to small access to supercomputing," says Ashok Krishnamurthy, the interim co-director of the center, which is based in Columbus. More than 25 businesses have directly participated in OSC's Blue Collar Computing program, which provides supercomputing power to companies that typically hadn't considered using it previously. More than 250 smaller companies have also used the services by accessing them through related organizations such as the Edison Welding Institute (EWI), he says.



Ashok Krishnamurthy says the Ohio Supercomputer Center's "Blue Collar Computing" initiative has given hundreds of small Ohio companies access to supercomputing power.

Welding gains from HPC analysis, too

EWI, a member organization for companies that do welding, works with the OSC to

allow member companies to log into a Web portal, the E-Weld Predictor, to simulate complex welds using supercomputers before doing the actual welding work, Krishnamurthy says. "This simulates a whole bunch of prototypes to cut down the time it takes to create a welding process from six months to two weeks."

"This is huge, because it reduces the time it takes to get new welding procedures in place," says Chris Conrardy, chief technology officer and vice president of technology and innovation at EWI. "It lets you optimize the weld and reduce the risks of bad welds."

What's really novel about this, Conrardy explains, is that the welding process is not linear. "It has lots of variables and historically has been difficult to model on computers. All sorts of things can happen to the properties of the materials depending on how you apply the heat," he says. And for particularly critical applications -- a pipeline, say, or a pressurized vessel -- "folks can spend an awful lot of time and money figuring out how to weld those things."

Typical workstations couldn't handle all of the calculations that are involved in this process, he says.

The E-Weld Predictor allows a user to enter the specifics of the materials to be joined, the dimensions and other needed details. After a few minutes, the system generates a PDF document with details about how the joining can be accomplished successfully, he says. "It's basically a report with predictions and procedures and results for the weld. It's something that we think is going to be increasingly a big deal."

Around 100 EWI member companies used the E-Weld Predictor in its first month, Conrardy says.

The E-Weld Predictor harnesses the OSC's 1,650-node IBM Glenn 1350 supercomputer. The Glenn 1350 runs AMD Opteron multicore processors and IBM cell processors and offers a peak performance of more than 75 TFLOPS, according to OSC.

The assistance provided by OSC can give a smaller company a huge boost a market filled with competitors, Krishnamurthy says. In particular, the time saved by using the supercomputing systems can be key to companies that are competing for contracts or need to manufacture a product as quickly as possible. "Shorter time to market makes a huge difference," he says. "Six months down to two weeks is a big deal."

The OSC is also working with DARPA, Krishnamurthy says, to help smaller companies that are providing parts and equipment through the U.S. Department of Defense supply chain. One example is a defense components manufacturer, which cannot be named because of security concerns, that designs power control units for the U.S. Navy, says Krishnamurthy. Often the vendor couldn't meet strict bid specifications because it didn't have the resources to do component testing, so it would request waivers for such testing, he says. The OSC and one of its partners, AltaSim Technologies, helped by showing the defense contractor how it could simulate the performance and construction of very large power control units using offsite supercomputing power.

The OSC's assistance helped the company save up to \$1 million in the development of prototype power units, says Jeffrey Crompton, principal at AltaSim. The supercomputer simulation work will potentially save even more if the parts eventually go into production -- as much as \$100 million or more in savings over traditional testing, prototyping and manufacturing, he explains.

A nod to Amazon

Graybill says he gives Amazon.com Inc. the bulk of the credit for creating the on-demand computing marketplace. In 2006, **Amazon began offering cloud-based computing services** to businesses that needed additional IT resources. The online retailer originally set up its Elastic Compute Cloud 2 (EC2) on-demand computing service as a way to utilize and monetize its internal IT infrastructure outside of peak holiday shopping periods, Graybill says. "They had [that huge infrastructure] to manage the loads on their busiest days, so they had spare cycles to sell at other times. That started this revolution, in my opinion," he says.

But while Amazon tapped into a need among businesses that never before had such easy access to externally managed IT resources, the retailer's offering is not necessarily the same thing as on-demand high-performance computing. Specialized HPC systems like those from SGI and IBM offer better and faster interconnects between processor nodes and memory, which means they are capable of dramatically superior performance than the more general computing architectures offered by Amazon and others, Graybill explains.

"Not to pick on Amazon, but these are different systems for different needs," he says. That said, Amazon recently added **more sophisticated monitoring tools** and **improved integration with traditional data centers** for services like connecting an internal cloud with Amazon's offerings.

Regardless of a customer's preferences, Graybill maintains that on-demand HPC brings the power of supercomputing to those who otherwise wouldn't be able to afford it -- "to the disenfranchised, so to speak."

Todd R. Weiss, an award-winning former Computerworld reporter, is a freelance writer covering technology news, cool tech gear, CRM, ERP and more. You can follow him on Twitter at <http://twitter.com/TechManTalking> or write to him at toddrweiss@gmail.com.

http://www.computerworld.com/s/article/9178932/Big_guns_for_hire_Supercomputing_in_the_cloud_

<http://news.idg.no/cw/art.cfm?id=CD1235BA-1A64-67EA-E4FABF08413A78DF>