

Q: What is the biggest security challenge your IT organization will face when providing distributed access to enterprise applications?

ROLLOVER TO ANSWER 

cit

COMPUTERWORLD

 Print Article  Close Window

The supercomputer on your desktop

While there's still a need for ginormous hardware, some traditional high-performance applications are already on the desktop and others are heading that way

John Brandon

February 16, 2010 ([Computerworld](#))

High-performance computing (HPC) has almost always required a supercomputer -- one of those room-size monoliths you find at government research labs and universities. And while those systems aren't going away, some of the applications traditionally handled by the biggest of Big Iron are heading to the desktop.

One reason is that processing that took an hour on a standard PC about eight years ago now takes six seconds, according to Ed Martin, a manager in the automotive unit at computer-aided design software maker Autodesk Inc. Monumental improvements in desktop processing power, graphics processing unit (GPU) performance, network bandwidth and solid-state drive speed combined with 64-bit throughput have made the desktop increasingly viable for large-scale computing projects.

Thanks to those developments, a transition to "a supercomputer on your desk" is in full force.

Earthquake simulations, nuclear-stockpile simulations and DNA research are staying put on traditional supercomputers for now. But as processor technology advances to multiple cores in the next 10 years, even those activities, or portions of them, could conceivably make their way to the desktop.

In the meantime, here are some examples of high-performance applications that are already running on smaller computers.



BROCADE 

SAVE THE DATE


Attend Brocade Conference 2010 September 20-22
and discover how to transform your network.

Brocade DCX • Brocade FCX • BigIron • NetIron
ServerIron • SAN Extension • SAN Switches

Building better drugs for anesthesia

Today, doctors know how to administer anesthesia-inducing drugs, and they know the effects, but they do not actually know what the drugs' molecules are doing when the patient drifts off to sleep. This analysis requires intense computational power to see not only when the anesthetic enters the respiratory system, but also how it starts making changes.

At [Temple University](#), researchers have developed models that measure the effects of applying anesthesia on molecules within nerve cells. The models currently run on a supercomputer, but plans are underway to perform the calculations on [an Nvidia GPU cluster](#) with four nodes. This will both save money and give researchers more flexibility to conduct tests when they're ready to do so (instead of having to wait for their scheduled time to use a supercomputer).



Researchers are modeling the effects of applying anesthetics on molecules within nerve cells.

In that scenario, each GPU has the computational power of a small HPC cluster. GPU calculations involve mathematical calculations on the scale of those normally used to, say, render pixels in a video game.

Dr. Axel Kohlmeyer, a researcher on the project, says the best way to understand the simulation is to imagine a box filled with rubber balls, where each ball is a slightly different size and moves at a slightly different rate, interconnected with springs. Some springs are stronger or weaker than others, and some of the balls move faster or react differently. In the simulation, Kohlmeyer can follow the movements of all molecules to see the effects of anesthetics in the human body.

"Groups of particles will form and go where they like to be as determined by the magnitude of their interactions," says Kohlmeyer, explaining how the simulation evolves to the point where the interactions become balanced. Temperature variants produce vibrations and introduce new molecular activity. "The computational model is actually simple, but the challenge is you need so many millions of interactions. We do not want to just know the interactions at one point, but rather how they change over time."

Having to repeat the calculations very often is another part of the challenge, he adds.

For Kohlmeyer, the goal is to discover when the condition of not feeling anything actually occurs in the human body. This could lead to the creation of new kinds of anesthetics or help doctors determine why problems such as memory loss can occur after surgery.


Bone surgical simulation

Researchers at the Ohio Supercomputer Center (OSC) in Columbus, Ohio, have found that not every simulation requires a traditional supercomputer. Don Stredney, the director and interface lab research scientist for [biomedical applications at OSC](#), found a limitation that's common with supercomputers: Batch processes are static and run on a scheduled time frame. They cannot provide real-time interactions, so they can't mimic a real surgical procedure. Desktop workstations

that cost \$6,000 to \$10,000 allow his team to run simulations that show, in real-time, how a surgery [changes a patient's anatomy](#), he says.

Stredney says his industry benefited from innovations in computer gaming because the standard consumer GPU became much more powerful, resulting in better realism at a much lower cost. Stredney says his researchers use commodity PCs running standard GPUs such as those from AMD's ATI unit and Nvidia Corp., but not high-end GPU clusters. However, they find that when the data sets grow too large with some simulations, they need to return to the supercomputer.

What drives Stredney's group back to the supercomputer, he says, is the "exponentially increasing size of data sets, images in the gigabyte-per-slice range and multiscale data sets that are now routinely being acquired at half-terabyte levels." Ever-larger data sets and the complex interaction required for real-time visual and auditory simulations "require more sophisticated systems," he says.



Simulations can show, in real-time, how a surgery changes a patient's anatomy.

Injection molding in automotive design

Injection-molding simulations are invaluable to car makers, Autodesk's Martin says. Injection molding is a process for producing parts from plastic materials. Simulations show whether an injection mold -- such as a bumper, for instance -- will cause denting and how the mold will fit with other parts of the car. They also reveal any defects. Designers consider many variables: the temperature of the mold, its geometric shape and how the injection-mold process will work with certain materials.



Automotive simulations such as this one can save car makers a lot of money.

A single physical prototype of a fender can cost more than \$1 million, Martin explains, so the better the simulation, the fewer prototypes that have to be built and the lower the production costs.

Simulations "used to require a significant cluster-computing installation, but we are achieving the same level of power with current desktop computers," says Martin, who says the desktop advances that played the biggest roles in making that

possible were the move to multicore processing, the use of multiple GPUs and 64-bit throughput. Martin uses standard desktop computers that can be purchased at Wal-Mart, with the latest 3D-capable GPUs and [Intel](#) dual-core CPUs.

The automotive industry is also using desktop software to model car designs. With such tools, the manufacturers can create complex renderings that rival -- or even outshine -- the graphics in the

best video games. Interestingly, when these models are created on desktop PCs, they are often used in marketing literature and in TV commercials. Martin says it was possible to take the models created on supercomputers and transport them to PCs, but today's PC-to-PC file transfers are simpler because the file formats are compatible across applications and the transfers take place over a standard network, making it possible for the images to reach a greater number of users.

Here's an example of modeling on the desktop: Designers often create car-paint models, examining the metallic flakes and how they look on certain plastics, or whether the paint looks dull in certain lighting conditions. The more accurate the models, the more process-intensive they are. In the past, an HPC environment was required because of the high number of polymers used in car paint -- around 8,000. Martin says modern desktops can now handle this kind of high-speed processing.

Web-based computational search


One of the most interesting ways HPC is coming to the desktop is through the Web. The best example of this right now is on WolframAlpha, a "computational knowledge engine" on the Web that's designed to, [in the site's own words](#), "collect and curate all objective data... and make it possible to compute whatever can be computed about anything."

The sea change here is that, while the searches you launch at WolframAlpha.com are still conducted on a supercomputer, the results return almost instantly, right in your browser. For example, when you type the simple word *moon*, you trigger a complex calculation for the moon's orbit relative to the Earth, along with average distances over a historical period.

[Schoeller Porter](#), an architect at Wolfram Research Inc., says the calculations can occur in near-real-time because of the decreased cost of HPC components. In the past, this kind of complex calculation would require setting up a batch request for a supercomputer that might take a few minutes or hours, or even a day. That wasn't necessarily because the calculations themselves take a long time, but because the processing resources were so expensive.

To test the yourself, type these calculation requests -- asking about the tides in Honolulu in six months -- into the WolframAlpha search engine:

- International space station jun 24
- skychart Timbuktu yesterday at 8:00pm
- NACA 2014 15 degrees
- tides in honolulu in 6 months
- $y'' + \sin y = x$
- GATTAACCC



Wednesday, May 5, 2010		
high tide	10:11 pm HST	+1.4 feet
Thursday, May 6, 2010		
low tide	5:42 am HST	0 feet
high tide	12:08 pm HST	+0.7 feet
low tide	4:23 pm HST	+0.5 feet

(computed using historical data, not taking into account weather, etc. heights relative to all-time average lowest daily tide)

WolframAlpha's Web-based searches bring complex computational data into a browser. A portion of the results for a recent query, about tides in Honolulu, is shown above.

When you run searches on WolframAlpha, you can change any variable at will from your desktop. You'll notice that some of the queries take longer than others, depending on how long it takes the supercomputer to determine the result.

While [Google](#) searches also rely on a supercomputer or cluster on the back end, they do not involve the same computational algorithms as WolframAlpha. Google searches for information and provides relevant links; WolframAlpha feeds the user actual information, not links, and that information may be pulled from a knowledge base or generated by calculations it performs.

Weather forecasting

Bill Magro, director of HPC software solutions at Intel Corp., says weather forecasting went through a major evolution in the past two decades, crossing a line from providing just a general idea of weather patterns to showing fine-grained details, such as the details of individual storm cells. This happened mostly because the immense data sets required for weather forecasting can now sit within the desktop's memory, and the processing also can finally keep pace.

Today, portions of weather simulations for meteorological use can run on desktop computers at local television stations, not just at the government agencies that monitor weather.

Such modeling is extremely complicated because weather occurs in open 3D space and there are millions of variables

Special effects on the desktop

Even though the movie industry still relies on traditional high-performance computing to render special effects, Hollywood was an early adopter of 64-bit workstations to help offload some of that work onto the desktop.

The special effects in movies like *2012* and [Avatar](#) usually require incredible processing power. There's an obvious

for temperature, location, wind and other factors. Weather models take the spatial extent -- say, a 5-kilometer grid with one data point for each 5km in each direction -- and then "shrink it down to 1km and then down to 500 meters," Magro explains. "Because you are shrinking in three dimensions, the computational requirement goes up by a factor of eight. It's not that there is more data, but you are modeling with more detail. It really does come down to having enough RAM and enough computer power."

The desktop is already being used for weather simulations at some local TV stations. In the future, as processor technology advances even further, we could all run detailed personal weather forecasts on our PCs. Magro says desktop computers could be used by individuals to see models for their immediate area, or to forecast weather for vacations six months into the future, for example, based on the conditions they are experiencing and have entered into the model.

More apps likely to follow

David Bader, the executive director of high-performance computing at the Georgia Institute of Technology, envisions several additional uses for HPC on the desktop in the next 10 years.

There are possibilities for discovering new trends among social networks or identifying their key influencers. Recommendations could be made for improving the energy efficiency of behaviors such as traffic routing to avoid congestion, scheduling computer applications to minimize energy usage, and monitoring a smart power grid, Bader says.



Special-effects tools like Autodesk's Smoke Text Preset, which created the 3D movie title above, run on the desktop.

difference between an animated waterfall that looks computer-generated (and made with consumer software) and one that looks realistic. Companies such as Industrial Light and Magic, Uncharted Territory and Dreamworks rely on vast computing clusters and server farms for rendering. A single frame of video can take hours to render.

CGI effects are still compute-intensive: In just one scene in which techniques are used to make light flow evenly over an object and then move in unison with a camera, there are millions of pixels on the screen at the same time. Bruno Sargeant, the film and TV industry manager at Autodesk, says shader models -- which act like a catalyst for pixels and use mathematical algorithms to make pixels move realistically -- are also helping.

But high-end workstations are starting to make their mark. Both Adobe Systems Inc., with its 64-bit Mercury rendering engine, and Autodesk, with an application called Smoke that's coming out in mid-November for Mac Pro computers, are leading the charge to make CGI effects work well on desktop computers, without the need for a massive server farm or cluster. Sargeant says the planets have aligned just recently with much faster bus speeds, multicore processing, solid-state drives and GPU acceleration.

"This area is extremely competitive," says Bjorn Andersson, an analyst at Artair Group, a business development consultancy in Sweden. "The success of some movies relies on what level of groundbreaking effects they are the first ones to use. That in turn depends on two factors -- access to raw compute power and developing new software functionality."

Other possibilities include the construction of 3D scenes from massive collections of public photographs and live transcriptions of teleconferences with speaker IDs.

Although experts agree that the largest supercomputers will be around for some time to come, it will be interesting to see what other traditional HPC applications will eventually come to reside on the desktop.

John Brandon is a veteran of the computing industry, having worked as an IT manager for 10 years and as a tech journalist for another 10. He has written more than 2,500 feature articles and is a regular contributor to Computerworld.