

Demand Computing (and Receive)

Ever wish you could adjust your compute capacity like changing the volume on a radio? You're not alone. New services are ready to take on systems-level computing.

By Jennifer Crebs

Last year, Applied Biosystems was in the midst of figuring out how to rapidly design genotyping assays based on the millions of validated SNPs released by the HapMap project. ABI's internal compute farm would have required three months to get the job done, and it wouldn't have been cost-effective to add servers on a temporary basis. Francisco De La Vega, the company's senior director of computational genetics, decided to try a different avenue: on-demand computing. He signed on as a beta tester of Sun Microsystems' grid service, tested out the algorithms on the Sun architecture, and ported his Linux-based code to the Solaris system. The job ran and finished in about a week.

On-demand, or utility, computing is meant to address situations just like this. Instead of extending power by physically setting up a gaggle of CPUs *in situ*, utility service providers make compute resources available to users, who are then charged in accordance with what they actually use. Much as public utilities like electricity or gas charge based on usage, resources under the utility computing rubric follow a metered price plan. For labs with neither the time nor capital to set up a home-grown cluster, this can be a real boon. Recently, vendors have increased their offerings for the life science market, making it easier than ever to access off-site compute resources at fairly low prices.

HP was one of the first companies to

get in on the utility computing craze several years ago, says Nick van der Zweep, director of utility computing at the company. One of its first utility products, called Instant Capacity, was developed in the late '90s and allowed users to instantly ramp up CPU power by activating nodes on an as-needed basis. These days, the company has added a range of on-demand services, which are all designed to adapt to variable computing needs. For those who don't want to buy equipment and pay for increased power on an as-needed basis, the company also offers a pay-for-use lease on an entire system.

Late last year, HP began to offer its Flexible Computing Services, by which customers can access the company's data centers as needed with cost-per-unit pricing. Compute power is based on platforms ranging from 32- and 64-bit Intel and Opteron processor systems to HP's Intel-based Integrity servers that clock in at 128 to 256 CPUs per box.

According to van der Zweep, flexible computing is targeted more to the high-performance technical market, and offers a couple of variations. Customers can either choose a basic service to access the machines with their own applications, or they can use the machines outfitted with HP's own batch and provisioning software. Upon enrolling to use the services, users receive start-up consultations, training, and a 48-hour pilot project to assess the

types of applications needed, van der Zweep says.

Sun also offers a "continuum of solutions" for researchers looking for more power, says Joerg Schwarz, the company's director of life sciences. For some, building and operating a cluster may be the way to go. But "once you have expanded your capacity, you're stuck with it," Schwarz says. It was with this in mind that Sun recently launched a service that provides a grid-based option for researchers who experience spikes in demand for computing muscle.

The Sun Grid Compute Facility provides users with online access to computational services for \$1 per CPU hour. The grid utility runs the Solaris x64 operating system on an Opteron platform, and users receive a gigabyte of disk space in which to install applications. Access to the system requires neither negotiating with a sales rep nor waiting for service to kick in. Users can set up an account online via PayPal, then load resources, run the job, and simply download results, says Aisling MacRunnels, senior director of utility computing at Sun.

Sun is currently working with a number of partners to furnish the grid with standard pre-installed applications, such as Blast and annotated protein databases. In this way, researchers will require even less set-up time when running sequence alignments with the utility. Schwarz expects that the application will be in place later on this fall. **GT**