

The rise of the Mega-Data Center

Stacey Higginbotham

The likes of Intel, Dell, and Hewlett-Packard are offering their chips and computers in new ways thanks to the emergence of cloud computing.

Behind popular Web services such as Facebook, Google (GOOG), and Amazon's (AMZN) AWS are racks and racks of computers serving up millions of pages or providing raw computing power. The use of thousands of servers to deliver one application or act as a pool of computing resources has changed the way that chipmakers and computer vendors are building their products. It has also led to the rise of the mega-data center.

Intel (INTC) estimates that by 2012, up to a quarter of the server chips it sells will go into such mega-data centers. Dell (DELL), which nearly two years ago created its Data Center Solutions Group to address the needs of customers buying more than 2,000 servers at a time, now says the division is the fourth- or fifth-largest server vendor in the world. In the meantime, suppliers are creating product lines and spending money on R&D to adjust to the needs of these mega-data-center operators, as those operators are fulfilling an increasing demand for applications and services delivered via the cloud.

Commoditized hardware saves money

The mega-data centers running computing clouds are becoming more distinct from both their corporate cousins, which have to run multiple applications, and the high-performance computing (HPC) systems that combine multiple CPUs with expensive networking equipment. In a Webinar held Feb. 18, Russ Daniels, chief technology officer of Cloud Strategy Services at Hewlett-Packard (HPO), explained some of the differences to one of the company's customers.

"In HPC and grid computing...we tend to focus on workloads that would be important enough to deserve specialized hardware," Daniels said. "Cloud computing is the same technological approach of doing work in parallel but done in the context of a commoditized network architecture and hardware."

In a nod to the shift in computing, HP last year reorganized its high-performance computing and commodity servers designed for mega-data centers into its Scalable Computing Initiative. But so far, it's Dell that's created a business around building customized servers for each customer using off-the-shelf hardware. Indeed, Dell understands that tiny savings in hardware spread out over thousands of servers mean huge price cuts for customers.

Like a car rental firm

For a data-center customer that doesn't need a swappable fan in place, the savings of \$10 offered by placing a permanent fan inside the server, multiplied by thousands of servers, adds up to real dollars. Instead of discounting its normal servers for large-volume buyers, Dell offers them exactly what they want and still makes money on the sales.

Jason Waxman, general manager of high-density computing in Intel's server systems, says his employer is learning the same lessons, especially when applied to the cost of power to run those data centers. In a Feb. 18 conference call to discuss Intel's ties to cloud computing, he compared mega-data-center owners to a car rental firm, noting that when a consumer buys an automobile he or she looks for the best individual features, but when Hertz buys a fleet of cars, it wants the set of features that costs the least to operate.

For Intel, that means power savings. Waxman said that since 25% of the costs of running one of these mega-data centers can be traced to power consumption, Intel is designing

motherboards that can be cooled more efficiently, offering software that keeps servers from running too hot, and participating in a variety of projects to bring power costs down.

On the chip side, many of these gains have trickled down to all server products and will continue to do so, but if the operators of these mega-data centers become too successful at delivering computing and services through the cloud, the pool of customers for HP, Dell, Rackable and IBM (IBM) may get a lot smaller.

Business Week, New York, 20 fev. 2009, Technology, online. Disponível em <www.businessweek.com>. Acesso em: 26 fev. 2009.

A utilização deste artigo é exclusiva para fins educacionais