

## FEATURE

## Tech Titans Building Boom

Google, Microsoft, and other Internet giants race to build the mega data centers that will power cloud computing

BY RANDY H. KATZ // FEBRUARY 2009

The **serene** countryside around the Columbia River in the northwestern United States has emerged as a major, and perhaps unexpected, battleground among Internet powerhouses. That's where Google, Microsoft, Amazon, and Yahoo have built some of the world's largest and most advanced computer facilities: colossal warehouses packed with tens of thousands of servers that will propel the next generation of Internet applications. Call it the data center arms race.

The companies flocked to the region because of its affordable land, readily available fiber-optic connectivity, abundant water, and even more important, inexpensive electricity. These factors are critical to today's large-scale data centers, whose sheer size and power needs eclipse those of the previous generation by one or even two orders of magnitude.

These new data centers are the physical manifestation of what Internet companies are calling cloud computing. The idea is that sprawling collections of servers, storage systems, and network equipment will form a seamless infrastructure capable of running applications and storing data remotely, while the computers people own will provide little more than the interface to connect to the increasingly capable Internet.

All of this means that what had been not-so-glamorous bit players of the tech world—the distant data centers—have now come into the limelight. But getting their design, operation, and location right is a daunting task. The engineers behind today's mega data centers—some call them Internet data centers, to differentiate them from earlier ones—are turning to a host of new strategies and technologies to cut construction and operating costs, consume less energy, incorporate greener materials and processes, and in general make the facilities more flexible and easier to expand. Among other things, data centers are adopting advanced power management hardware, water-based cooling systems, and denser server configurations, making these facilities much different from conventional air-conditioned server rooms.

Take Microsoft's data center in Quincy, Wash., with more than 43 600 square meters of space, or nearly the area of 10 American football fields. The company is tight-lipped about the number of servers at the site, but it does say the facility uses 4.8 kilometers of chiller piping, 965 km of electrical wire, 92 900 m<sup>2</sup> of drywall, and 1.5 metric tons of batteries for backup power. And the data center consumes 48 megawatts—enough power for 40 000 homes.

Yahoo, based in Sunnyvale, Calif., also chose the tiny town of Quincy—population 5044, tucked in a valley dotted with potato farms—for its state-of-the-art, 13 000-m<sup>2</sup> facility, its second in the region. The company says it plans to operate both data centers with a zero-carbon footprint by using, among other things, hydropower, water-based chillers, and external cold air to do some of the cooling. As some observers put it, the potato farms have yielded to the server farms.

To the south, Google, headquartered in Mountain View, Calif., opened a vast data center on the banks of the Columbia, in The Dalles, Ore. The site has two active buildings, each 6500 m<sup>2</sup>, and there's been talk of setting up a third. Google, however, won't discuss its expansion plans. Nor will it say how many servers the complex houses or how much energy and water it consumes. When it comes to data centers these days, there's a lot of experimentation going on, and companies are understandably secretive.

Another hush-hush data center project is taking place about 100 km east, in Boardman, Ore. Last year, news emerged that its owner is the Seattle-based online retailer Amazon, a major contender in the cloud-computing market. The facility is said to include three buildings and a 10-MW electrical substation, but the company has declined to confirm details.

Amazon, Google, Microsoft, and Yahoo are trying to keep up with the ever-growing demand for Internet services like searching, image and video sharing, and social networking. But they're also betting on the explosive growth of Web-based applications that will run on their computer clouds. These apps, which display information to users on an Internet browser while processing and storing data on remote servers, range from simple Web-based e-mail like Gmail and Hotmail to more complex services like Google Docs' word processor and spreadsheet. The applications also include a variety of enterprise platforms like Salesforce.com's customer management system. So the race to the clouds is on. It's taking place not only in the Pacific Northwest but also in many other parts of the United States and elsewhere.

With reportedly more than 1 million servers scattered in three dozen data centers around the world, Google is spending billions of dollars to build large-scale facilities in Pryor, Okla.; Council Bluffs, Iowa; Lenoir, N.C.; and Goose Creek, S.C. Last year the company released its Google App Engine, a cloud-based platform that individuals and businesses can use to run applications. Microsoft, whose Quincy facility will house the initial version of its Windows Azure cloud-based operating system, is constructing additional facilities in Chicago, San Antonio, and Dublin, at roughly US \$500 million each. The company is also said to be looking for a site in Siberia.

Designing bigger and better data centers requires innovation not only in power and cooling but also in computer architecture, networking, operating systems, and other areas. Some engineers are indeed quite excited by the prospect of designing computing systems of this magnitude—or as some Googlers call them, "warehouse-size computers."

**The big data centers** that went up during the dot-com boom in the late 1990s and early 2000s housed thousands, even tens of thousands, of servers. Back in those days, facility managers could expand their computing resources almost indefinitely. Servers, storage systems, and network equipment were—and still are—relatively cheap, courtesy of Moore's Law. To get the computing power you needed, all you had to do was set up or upgrade your equipment and turn up the air-conditioning.

But in recent years, this approach has hit a wall. It became just too expensive to manage, power, and cool the huge number of servers that companies employ. Today's largest data centers house many tens of thousands of servers, and some have reportedly passed the 100 000 mark. The data centers of the dot-com era consumed 1 or 2 MW. Now facilities that require 20 MW are common, and already some of them expect to use 10 times as much. Setting up so many servers—mounting them onto racks, attaching cables, loading software—is very time-consuming. Even worse,

with electricity prices going up, it's hugely expensive to power and cool so much equipment. Market research firm IDC estimates that within the next six years, the companies operating data centers will spend more money per year on energy than on equipment.

*The data centers of the dot-com era consumed 1 or 2 megawatts. Now facilities that require 20 megawatts are common, and already some expect to use 10 times as much*

What's more, the environmental impact of data centers is now on the radar of regulators and environmentally conscious stakeholders. The management consulting firm McKinsey & Co. reports that the world's 44 million servers consume 0.5 percent of all electricity and produce 0.2 percent of all carbon dioxide emissions, or 80 megatons a year, approaching the emissions of entire countries like Argentina or the Netherlands.

Data center designers know they must do many things to squeeze more efficiency out of every system in their facilities. But what? Again, companies don't want to give the game away. But in their efforts to demonstrate the greening of their operations, Google, Microsoft, and others have revealed some interesting details. And of course, a number of innovations come from the vendor side, and these equipment suppliers are eager to discuss their offerings.

First, consider the server infrastructure within the data center. Traditionally, servers are grouped on racks. Each server is a single computer with one or more processors, memory, a disk, a network interface, a power supply, and a fan—all packaged in a metal enclosure the size of a pizza box. A typical dual-processor server demands 200 watts at peak performance, with the CPUs accounting for about 60 percent of that. Holding up to 40 pizza-box-size servers, a full rack consumes 8 kilowatts. Packed with blade servers—thin computers that sit vertically in a special chassis, like books on a shelf—a rack can use twice as much power.

The new large-scale data centers generally continue to use this rack configuration. Still, there's room for improvement. For one, Microsoft has said it was able to add more servers to its data centers simply by better managing their power budgets. The idea is that if you have lots of servers that rarely reach their maximum power at the same time, you can size your power supply not for their combined peak power but rather for their average demand. Microsoft says this approach let it add 30 to 50 percent more servers in some of its data centers. The company notes, however, that this strategy requires close monitoring of the servers and the use of power-control schemes to avoid overloading the distribution system in extreme situations.

How the servers are set up at Google's data centers is a mystery. But this much is known: The company relies on cheap computers with conventional multicore processors. To reduce the machines' energy appetite, Google fitted them with high-efficiency power supplies and voltage regulators, variable-speed fans, and system boards stripped of all unnecessary components like graphics chips. Google has also experimented with a CPU power-management feature called dynamic voltage/frequency scaling. It reduces a processor's voltage or frequency during certain periods (for example, when you don't need the results of a computing task right away). The server executes its work more slowly, thus reducing power consumption. Google engineers have reported energy savings of around 20 percent on some of their tests.

**But servers aren't everything** in a data center. Consider the electrical and cooling systems, which have yielded the biggest energy-efficiency improvements in some facilities. In a typical data center, energy flows through a series of transformers and distribution systems, which reduce the high-voltage ac to the standard 120 or 208 volts used in the server racks. The result is that by the time energy reaches the servers, losses can amount to 8 or 9 percent. In addition, the servers' processors, memory, storage, and network interfaces are ultimately powered by dc, so each server needs to make the ac-to-dc conversion, generating additional losses. This scenario is leading some data centers to adopt high-efficiency transformers and distribution systems to improve energy usage, and some experts have proposed distributing high-voltage dc power throughout the data center, an approach they say could cut losses by 5 to 20 percent.

Even more significant, data centers are saving lots of energy by using new cooling technologies. Traditional facilities use raised floors above a cold-air space, called a plenum, and overhead hot-air collection. The major drawback of this scheme is that hot spots—say, a highly active server or switch—may go undetected and the overheated device may fail. To avoid such situations, some data centers cool the entire room to temperatures much lower than what would otherwise be needed, typically to around 13 °C.

Now data centers are turning to new cooling systems that let them keep their rooms at temperatures as high as 27 °C. American Power Conversion Corp., in West Kingston, R.I., offers a special server rack with an air-conditioning unit that blows cold air directly onto the servers. Hewlett-Packard's "smart cooling" techniques use sensors within the machine room to position louvers in the floor that direct cold air at local hot spots, allowing the data center to run at a higher ambient temperature and reducing cooling costs by 25 to 40 percent.

Traditional air-conditioning systems are also making way for higher-efficiency water-based units. Today's large data centers rely on cooling towers, which use evaporation to remove heat from the cooling water, instead of traditional energy-intensive chillers. Several hardware vendors have developed water-based mechanisms to eliminate heat. IBM's Rear Door Heat eXchanger mounts on the back of a standard rack, essentially providing a water jacket that can remove up to 15 kW of heat. It's back to the mainframe days!

Data centers are also trying to use cold air from outside the buildings—a method known as air-side economization. At its Quincy data center, Yahoo says that rather than relying on conventional air-conditioning year-round, it has adopted high-efficiency air conditioners that use water-based chillers and cold external air to cool the server farms during three-quarters of the year.

Finally, the most radical change taking place in some of today's mega data centers is the adoption of containers to house servers. Instead of building raised-floor rooms, installing air-conditioning systems, and mounting rack after rack, wouldn't it be great if you could expand your facility by simply adding identical building blocks that integrate computing, power, and cooling systems all in one module?

That's exactly what vendors like IBM, HP, Sun Microsystems, Rackable Systems, and Verari Systems have come up with. These modules consist of standard shipping containers, which can house some 3000 servers, or more than 10 times as many as a conventional data center could pack in the same space. Their main advantage is that they're fast to deploy. You just roll these modules into the building, lower them to the floor, and power them up. And they also let you refresh your technology more easily—just truck them back to the vendor and wait for the upgraded version to arrive.

Sun was the first company to offer a container module. Its MD S20 includes 280 pizza-box-size quad-core servers plus

Sun's monitoring and control equipment. All in all, it consumes 187.5 kW, which translates to about 12.6 kW/m<sup>2</sup>. Conventional raised-floor data centers have much more modest densities, often as low as 0.5 kW/m<sup>2</sup>. Verari Systems' container houses 1400 blade servers, and it can use either self-contained or chilled water cooling, consuming 400 kW, or 12.9 kW/m<sup>2</sup>. Rackable Systems' ICE Cube uses a self-contained cooling system and a dc power system, allowing the containers to handle a power density of 16 kW/m<sup>2</sup>.

Microsoft's Chicago data center, which will support the company's cloud initiatives, is a hybrid design, combining conventional raised-floor server rooms and a container farm with some 200 units. The space where the containers will sit looks more like a storage warehouse than a typical chilled server room. Pipes hang from the ceiling ready to be connected to the containers to provide them with cooling water and electricity. As one Microsoft engineer describes it, it's the "industrialization of the IT world."

**So what kind of results** have Google and Microsoft achieved? To interpret some of the numbers they have reported, it helps to understand an energy-efficiency metric that's becoming popular in this industry. It's called power usage effectiveness, or PUE, and it's basically the facility's total power consumption divided by the power used only by servers, storage systems, and network gear. A PUE close to 1 means that your data center is using most of the power for the computing infrastructure and that little power is devoted to cooling or lost in the electrical distribution system. Because it doesn't gauge the efficiency of the computing equipment itself, PUE doesn't tell the full story, but it's still a handy yardstick.

A 2007 study by the U.S. Environmental Protection Agency (EPA) reported that typical enterprise data centers had a PUE of 2.0 or more. That means that for every watt used by servers, storage, and networking equipment, an additional watt is consumed for cooling and power distribution. The study suggested that by 2011 most data centers could reach a PUE of 1.7, thanks to some improvements in equipment, and that with additional technology some facilities could reach 1.3 and a few state-of-the-art facilities, using liquid cooling, could reach 1.2.

Curiously, that's where Google claims to be *today*. The company has reported an average PUE of 1.21 for its six "large-scale Google-designed data centers," with one reaching a PUE of 1.15. (The company says its efficiency improvements amount to an annual savings of 500 kilowatt-hours, 300 kilograms of carbon dioxide, and 3785 liters of water per server.) How exactly Google achieves these results remains a closely guarded secret. Some observers say that Google is an extreme case because its massive scale allows it to acquire efficient equipment that would be too costly for other companies. Others are skeptical that you can even achieve a PUE of 1.15 with today's water-based cooling technologies, speculating that the facility uses some unusual configuration—perhaps a container-based design. Google says only that these facilities don't all have the same power and cooling architectures. It's clear, though, that Google engineers are thinking out of the box: The company has even patented a "water-based data center," with containerized servers floating in a barge and with power generated from sea motion.

For its part, Microsoft has recently reported PUE measurements for the container portion of its Chicago data center. The modules have proved to be big energy savers, yielding an annual average PUE of 1.22. The result convinced the company that all-container data centers, although a radical design, make sense. Indeed, Microsoft is pushing the container design even further in its next-generation data centers, with preassembled modules containing not only servers but also cooling and electrical systems converging to create a roofless facility—a data center that could be mistaken for a container storage yard.

But the containerization of the data center also has its skeptics. They say that containers may not be as plug-and-play as claimed, requiring additional power controllers, and that repairing servers could involve hauling entire containers back to vendors—a huge waste of energy. With or without containers, the EPA is now tracking the PUE of dozens of data centers, and new metrics and recommendations should emerge in coming years.

**IT'S CERTAIN** that future data centers will have to take today's improvements far from where they are today. For one thing, we need equipment with much better power management. Google engineers have called for systems designers to develop servers that consume energy in proportion to the amount of computing work they perform. Cellphones and portable devices are designed to save power, with standby modes that consume just a tenth or less of the peak power. Servers, on the other hand, consume as much as 60 percent of their peak power when idle. Google says simulations showed that servers capable of gradually increasing their power usage as their computing activity level increases could cut a data center's energy usage by half.

Designers also need to rethink how data centers obtain electricity. Buying power is becoming more and more expensive, and relying on a single source—a nearby power plant, say—is risky. The companies that operate these complexes clearly need to explore other power-producing technologies—solar power, fuel cells, wind—to reduce their reliance on the grid. The same goes for water and other energy-intensive products like concrete and copper. Ultimately, the design of large-scale computer facilities will require a comprehensive rethinking of performance and efficiency.

And then there's software. Virtualization tools, a kind of operating system for operating systems, are becoming more popular. They allow a single server to behave like multiple independent machines. Some studies have shown that servers at data centers run on average at 15 percent or less of their maximum capacity. By using virtual machines, data centers could increase utilization to 80 percent. That's good, but we need other tools to automate the management of these servers, control their power usage, share distributed data, and handle hardware failures. Software is what will let data centers built out of inexpensive servers continue to grow.

And grow they will. But then what? Will data centers just evolve into ever-larger server-packed buildings? Today's facilities are designed around existing equipment and traditional construction and operational processes. Basically, we build data centers to accommodate our piles of computers—and the technicians who will squeeze around them to set things up or make repairs. Containers offer an interesting twist, but they're still based on existing servers, racks, power supplies, and so on. Some experts are suggesting that we need to design systems *specifically* for data centers. That makes good sense, but just what those facilities will look like in the future is anyone's guess.

#### To Probe Further

For more about Google's data centers and their energy-efficiency strategies, see "The Case for Energy-Proportional Computing," IEEE Computer, December 2007, and also the company's report on this topic at <http://www.google.com/corporate/datacenters>.

Microsoft's Michael Manos and Amazon's James Hamilton offer insights about data centers on their respective blogs: LooseBolts (<http://loosebolls.wordpress.com>) and Perspectives (<http://perspectives.mvdirona.com>).

Data Center Knowledge (<http://www.datacenterknowledge.com>), The Raised Floor (<http://theraisedfloor.typepad.com>), Green Data Center Blog (<http://www.greenm3.com>), and Data Center Links (<http://datacenterlinks.blogspot.com>) offer comprehensive coverage of the main developments in the industry.

#### most recent comments

**BRIAN C 03.17.2010**

As an engineer involved in energy savings I think the actions being taken to reduce centers usage are fine. What I have a problem with is I'm starting to see centers get rebates and energy conservation funds from local governments, utilities, or utility commissions to install lower power servers or server optimization software. The industry is going this way anyway.. I don't think data centers should be using public money to enhance their operations under the banner of going green. There are more effective uses for these funds.I'm scared this could get out of hand.

**ROBERT 07.19.2009**

1.5 tons of batteries? Sounds like enough to power a hummer not 40,000 homes.

**MARK 06.08.2009**

Can you site the factual reference for the amount of energy consumed for this statment?; Take Microsoft's data center in Quincy, Wash., with more than 43 600 square meters of space, or nearly the area of 10 American football fields. The company is tight-lipped about the number of servers at the site, but it does say the facility uses 4.8 kilometers of chiller piping, 965 km of electrical wire, 92 900 m2 of drywall, and 1.5 metric tons of batteries for backup power. And the data center consumes 48 megawatts'enough power for 40 000 homes." And, 48 megawatts per...? day/year, hour? please be specific. Thank you

**MEL ELY 02.17.2009**

What is the impact on business, government and personal communications in the United States, indeed - the world - if one of these centers were destroyed? I assume other nodes in the network would take over? Nevertheless, that would surely result in congestion, slow response and perhaps complete loss of some data. These questions kept coming to mind as I read the article.

**RICHARD REIS 02.06.2009**

As an electrical engineer, I specialize in helping companies use less energy - demand-side management. I do that through lighting efficiency, insulation, limiting infiltration, better HVAC systems, etc. I came to this specialty through as a way to combine my engineering skill and passion for the environment. (Most kWh we use degrade our natural environment through global warming, emission of other pollutants, coal mining, etc.)