

June 14, 2009

Data Center Overload

By TOM VANDERBILT

It began with an Xbox game.

On a recent rainy evening in Brooklyn, I was at a friend's house playing (a bit sheepishly, given my incipient middle age) Call of Duty: World at War. Scrolling through the game's menus, I noticed a screen for Xbox Live, which allows you to play against remote users via broadband. The number of Call of Duty players online at that moment? More than 66,000.

Walking home, I ruminated on the number. Sixty-six thousand is the population of a small city — Muncie, Ind., for one. Who and where was this invisible metropolis? What infrastructure was needed to create this city of ether?

We have an almost inimical incuriosity when it comes to infrastructure. It tends to feature in our thoughts only when it's not working. The [Google](#) search results that are returned in 0.15 seconds were once a stirring novelty but soon became just another assumption in our lives, like the air we breathe. Yet whose day would proceed smoothly without the computing infrastructure that increasingly makes it possible to navigate the world and our relationships within it?

Much of the daily material of our lives is now dematerialized and outsourced to a far-flung, unseen network. The stack of letters becomes the e-mail database on the computer, which gives way to Hotmail or Gmail. The clipping sent to a friend becomes the attached PDF file, which becomes a set of shared bookmarks, hosted offsite. The photos in a box are replaced by JPEGs on a hard drive, then a hosted sharing service like Snapfish. The tilting CD tower gives way to the MP3-laden hard drive which itself yields to a service like Pandora, music that is always "there," waiting to be heard.

But where is "there," and what does it look like?

"There" is nowadays likely to be increasingly large, powerful, energy-intensive, always-on and essentially out-of-sight data centers. These centers run enormously scaled software applications with millions of users. To appreciate the scope of this phenomenon, and its crushing demands on storage capacity, let me sketch just the iceberg's tip of one average individual digital presence: my own. I have photos on Flickr (which is owned by [Yahoo](#), so they reside in a Yahoo data center, probably the one in Wenatchee, Wash.); the [Wikipedia](#) entry about me dwells on a database in Tampa, Fla.; the video on YouTube of a talk I delivered at Google's headquarters might dwell in any one of Google's data centers, from The Dalles in Oregon to Lenoir, N.C.; my LinkedIn profile most likely sits in an [Equinix](#)-run data center in Elk Grove Village, Ill.; and my blog lives at Modwest's headquarters in Missoula, Mont. If one of these sites happened to be down, I might have Twittered a complaint, my tweet paying a virtual visit to (most likely) NTT America's data center in Sterling, Va. And in each of these cases, there would be at least one mirror data center somewhere else — the built-environment equivalent of an external hard drive, backing things up.

Small wonder that this vast, dispersed network of interdependent data systems has lately come to be referred to by an appropriately atmospheric — and vaporous — metaphor: the cloud. Trying to chart the

cloud's geography can be daunting, a task that is further complicated by security concerns. "It's like 'Fight Club,'" says Rich Miller, whose Web site, Data Center Knowledge, tracks the industry. "The first rule of data centers is: Don't talk about data centers."

Yet as data centers increasingly become the nerve centers of business and society — even the storehouses of our fleeting cultural memory (that dancing cockatoo on YouTube!) — the demand for bigger and better ones increases: there is a growing need to produce the most computing power per square foot at the lowest possible cost in energy and resources. All of which is bringing a new level of attention, and challenges, to a once rather hidden phenomenon. Call it the architecture of search: the tens of thousands of square feet of machinery, humming away 24/7, 365 days a year — often built on, say, a former bean field — that lie behind your Internet queries.

INSIDE THE CLOUD

[Microsoft](#)'s data center in Tukwila, Wash., sits amid a nondescript sprawl of beige boxlike buildings. As I pulled up to it in a Prius with Michael Manos, who was then Microsoft's general manager of data-center services, he observed that while "most people wouldn't be able to tell this wasn't just a giant warehouse," an experienced eye could discern revelatory details. "You would notice the plethora of cameras," he said. "You could follow the power lines." He gestured to a series of fluted silver pipes along one wall. "Those are chimney stacks, which probably tells you there's generators behind each of those stacks." The generators, like the huge banks of [U.P.S.](#) (uninterruptible power supply) batteries, ward against surges and power failures to ensure that the data center always runs smoothly.

After submitting to biometric hand scans in the lobby and passing through a sensor-laden multidoor man trap, Manos and I entered a bright, white room filled with librarylike rows of hulking, black racks of servers — the dedicated hardware that drives the Internet. The Tukwila data center happens to be one of the global homes of Microsoft's Xbox Live: within those humming machines exists my imagined city of ether. Like most data centers, Tukwila comprises a sprawling array of servers, load balancers, routers, fire walls, tape-backup libraries and database machines, all resting on a raised floor of removable white tiles, beneath which run neatly arrayed bundles of power cabling. To help keep servers cool, Tukwila, like most data centers, has a system of what are known as hot and cold aisles: cold air that seeps from perforated tiles in front is sucked through the servers by fans, expelled into the space between the backs of the racks and then ventilated from above. The collective din suggests what it must be like to stick your head in a Dyson Airblade hand dryer.

Tukwila is less a building than a machine for computing. "You look at a typical building," Manos explained, "and the mechanical and electrical infrastructure is probably below 10 percent of the upfront costs. Whereas here it's 82 percent of the costs." Little thought is given to exterior appearances; even the word "architecture" in the context of a data center can be confusing: it could refer to the building, the network or the software running on the servers. Chris Crosby, a senior vice president with [Digital Realty Trust](#), the country's largest provider of data-center space, compares his company's product to a car, an assembly-line creation complete with model numbers: "The model number tells you how much power is available inside the facility." He also talks about the "industrialization of the data center," in contrast to the so-called whiteboard model of server design, by which each new building might be drawn up from scratch. The data center, he says, is "our railroad; it doesn't matter what kind of train you put on it."

At Tukwila — as at any big data center — the computing machinery is supported by what Manos calls the "back-of-the-house stuff": the chiller towers, the miles of battery springs, the intricate networks of piping. There's also what Manos calls "the big iron," the 2.5-megawatt, diesel-powered Caterpillar generators

clustered at one end of a cavernous space known as the wind tunnel, through which air rushes to cool the generators. “In reality, the cloud is giant buildings full of computers and diesel generators,” Manos says. “There’s not really anything white or fluffy about it.”

Tukwila is one of Microsoft’s smaller data centers (they number “more than 10 and fewer than 100,” Manos told me with deliberate vagueness). In 2006, the company, lured by cheap hydropower, tax incentives and a good fiber-optic network, built a 500,000-plus-square-foot data center in Quincy, Wash., a small town three hours from Tukwila known for its bean and spearmint fields. This summer, Microsoft will open a 700,000-plus-square-foot data center — one of the world’s largest — in Chicago. “We are about three to four times larger than when I joined the company” — in 2004 — “just in terms of data-center footprint,” Debra Chrapaty, corporate vice president of Global Foundation Services at Microsoft, told me when I met with her at Microsoft’s offices in Redmond, Wash.

Yet when it comes to a large company like Microsoft, it can be difficult to find out what any given data center is used for. The company, for reasons ranging from security to competitive advantage, won’t provide too much in the way of details, apart from noting that Quincy could hold 6.75 trillion photos. “We support over 200 online properties with very large scale,” Chrapaty offered. “And so when you think about Hotmail supporting 375 million users, or search supporting three billion queries a month, or Messenger supporting hundreds of millions of users, you can easily assume that those properties are very large properties for our company.”

Thanks to the efforts of amateur Internet Kremlinologists, there are occasional glimpses behind the silicon curtain. One blogger managed to copy a frame from a 2008 video of a Microsoft executive’s PowerPoint presentation showing that the company had nearly 150,000 servers (a number that presumably would now be much higher, given an estimated monthly server growth of 10,000) and that nearly 80,000 of those were used by its search application, now called Bing. When I discussed the figures with her, Chrapaty would only aver, crisply, that “in an average data center, it’s not uncommon for search to take up a big portion of a facility.”

THE RISE OF THE MEGA-DATA CENTER

Data centers were not always unmarked, unassuming and highly restricted places. In the 1960s, in fact, huge [I.B.M.](#) mainframe computers commanded pride of place in corporate headquarters. “It was called the glasshouse,” says Kenneth Brill, founder of the Uptime Institute, a data-center research and consulting group. “It was located near the executive suite. Here you’d spent \$15 to 30 million on this thing — the executives wanted to show it off.”

Over the past few decades, Brill notes, there has been an oscillation between using centralized I.T. resources and using more dispersed computing power — a battle between mainframes and desktop computers. The latest iteration is what’s called the thin client: the use of centralized servers rather than the software and operating systems of desktops to perform the bulk of computing needs. But thinness in the office has come with increased thickness elsewhere: more servers in ever-larger data centers.

In his book “The Big Switch,” Nicholas Carr draws an analogy between the rise of mega-data centers and the Industrial Revolution. Just as nascent industries, once powered by water wheels, were by the 20th century able to “run their machines with electric current generated in distant power plants,” advances in technology and transmission speeds are permitting computing to function like a utility, a distant but ever-accessible cloud of services.

This has sweeping implications for business and society. Instead of buying software and hiring I.T. employees, companies can outsource things like customer relationship management, or C.R.M., the database software that companies use to track client interactions, to an Internet company like salesforce.com, which sells subscriptions, or seats, to its services. “Customers who have two seats on salesforce.com, like a mom-and-pop flower shop, have access to the same application as a customer that has 65,000 seats, like [Starbucks](http://Starbucks.com) or [Dell](http://Dell.com),” Adam Gross, vice president of platform marketing with salesforce.com, told me at the company’s offices in San Francisco. By contrast, just a few years ago, he went on, “if you were to attack a really large problem, like delivering a C.R.M. application to 50,000 companies, or serving every single song ever, it really sort of felt outside your domain unless you were one of the largest companies in the world. There are these architectures now available for anybody to really attack these massive-scale kinds of problems.”

And while most companies still maintain their own data centers, the promise is that instead of making costly investments in redundant I.T. hardware, more and more companies will tap into the utility-computing grid, piggybacking on the infrastructures of others. Already, [Amazon](http://Amazon.com) Web Services makes available, for a fee, the company’s enormous computing power to outside customers. The division already uses more bandwidth than Amazon’s extensive retailing operations, while its Simple Storage Service holds some 52 billion virtual objects. “We used to think that owning factories was an important piece of a business’s value,” says Bryan Doerr, the chief technology officer of Savvis, which provides I.T. infrastructure and what the company calls “virtualized utility services” for companies like Hallmark. “Then we realized that owning what the factory produces is more important.”

THE ANNIHILATION OF SPACE BY TIME

For companies like Google, Yahoo and, increasingly, Microsoft, the data center is the factory. What these companies produce are services. It was the increasing “viability of a service-based model,” as Ray Ozzie, now the chief software architect at Microsoft, put it in 2005 — portended primarily by Google and its own large-scale network of data centers — that set Microsoft on its huge data-center rollout: if people no longer needed desktop software, they would no longer need Microsoft. This realization brought new prominence to the humble infrastructure layer of the data center, an aspect of the business that at Microsoft, as at most tech companies, typically escaped notice — unless it wasn’t working. Data centers have now become, as Debra Chrapaty of Microsoft puts it, a “true differentiator.”

Indeed, the number of servers in the United States nearly quintupled from 1997 to 2007. (Kenneth Brill of the Uptime Institute notes that the mega-data centers of Google and its ilk account for only an estimated 5 percent of the total market.) The expansion of Internet-driven business models, along with the data retention and compliance requirements of a variety of tighter accounting standards and other financial regulations, has fueled a tremendous appetite for data-center space. For a striking example of how our everyday clicks and uploads help drive and shape this real world real estate, consider [Facebook](http://Facebook.com).

Facebook’s numbers are staggering. More than 200 million users have uploaded more than 15 billion photos, making Facebook the world’s largest photo-sharing service. This expansion has required a corresponding infrastructure push, with an energetic search for financing. “We literally spend all our time figuring how to keep up with the growth,” Jonathan Heiliger, Facebook’s vice president of technical operations, told me in a company conference room in Palo Alto, Calif. “We basically buy space and power.” Facebook, he says, is too large to rent space in a managed “co-location facility,” yet not large enough to build its own data centers. “Five years ago, Facebook was a couple of servers under Mark’s desk in his dorm room,” Heiliger explained, referring to [Mark Zuckerberg](http://MarkZuckerberg.com), Facebook’s founder. “Then it moved to two sorts

of hosting facilities; then it graduated to this next category, taking a data center from an R.E.I.T.” — real estate investment trust — “in the Bay Area and then basically continued to expand that. We now have a fleet of data centers.”

A big challenge for Facebook, or any Internet site with millions of users, is “scalability” — ensuring that the infrastructure will keep working as new applications or users are added (often in incredibly spiky fashion, as when [Oprah Winfrey](#) joined and immediately garnered some 170,000 friends). Another issue is determining where Facebook’s data centers are located, where its users are located and the distance between them — what is called latency. Though the average user might not appreciate it, a visit to Facebook may involve dozens of round trips between a browser and any number of the site’s servers. In 2007, Facebook opened a third data center in Virginia to expand its capacity and serve its increasing number of users in Europe and elsewhere. “If you’re in the middle of the country, the differences are pretty minor whether you go to California or Virginia,” Heiliger said. But extend your communications to, say, India, and delay begins to compound. Bits, limited by the laws of physics, can travel no faster than the speed of light. To hurry things up, Facebook can try to reduce the number of round trips, or to “push the data as close to a user as possible” (by creating new data centers), or to rely on content-data networks that store commonly retrieved data in Internet points of presence (POPs) around the world.

While an anxious Facebook user serially refreshing to see if a friend has replied to an invitation might seem the very picture of the digital age’s hunger for instantaneity, to witness a true imperative for speed, you must visit NJ2, a data center located in Weehawken, N.J., just through the Lincoln Tunnel from Manhattan. There, in an unmarked beige complex with smoked windows, hum the trading engines of several large financial exchanges including, until recently, the Philadelphia Stock Exchange (it was absorbed last year by Nasdaq).

NJ2, owned by Digital Realty Trust, is managed by Savvis, which provides “proximity hosting” — enabling financial companies to be close to the market. At first I took this to mean proximity to Wall Street, but I soon learned that it meant proximity of the financial firms’ machines to the machines of the trading exchanges in NJ2. This is desirable because of the rise of electronic exchanges, in which machine-powered models are, essentially, competing against other machine-powered models. And the temporal window for such trading, which is projected this year by Celent to account for some 46 percent of all U.S. trading volume, is growing increasingly small.

“It used to be that things were done in seconds, then milliseconds,” Varghese Thomas, Savvis’s vice president of financial markets, told me. Intervening steps — going through a consolidated ticker vendor like [Thomson Reuters](#) — added 150 to 500 milliseconds to the time it takes for information to be exchanged. “These firms said, ‘I can eliminate that latency much further by connecting to the exchanges directly,’ ” Thomas explained. Firms initially linked from their own centers, but that added precious fractions of milliseconds. So they moved into the data center itself. “If you’re in the facility, you’re eliminating that wire.” The specter of infinitesimal delay is why, when the Philadelphia Stock Exchange, the nation’s oldest, upgraded its trading platform in 2006, it decided to locate the bulk of its trading engines 80 miles — and three milliseconds — from Philadelphia, and into NJ2, where, as Thomas notes, the time to communicate between servers is down to a millionth of a second. (Latency concerns are not limited to Wall Street; it is estimated that a 100-millisecond delay reduces Amazon’s sales by 1 percent.)

At NJ2, a room hosting one of the exchanges (I agreed not to say which, for security reasons), housed, in typical data-center fashion, rows of loudly humming black boxes, whose activity was literally inscrutable. This seemed strangely appropriate; after all, as Thomas pointed out, the data center hosts a number of

“dark pools,” or trading regimens that allow the anonymous buying and selling of small amounts of securities at a time, so as not, as Thomas puts it, “to create ripples in the market.”

It seemed heretical to think of [Karl Marx](#). But looking at the roomful of computers running automated trading models that themselves scan custom-formatted machine-readable financial news stories to help make decisions, you didn't have to be a Marxist to appreciate his observation that industry will strive to “produce machines by means of machines” — as well as his prediction that the “more developed the capital,” the more it would seek the “annihilation of space by time.”

THE COST OF THE CLOUD

Data centers worldwide now consume more energy annually than Sweden. And the amount of energy required is growing, says Jonathan Koomey, a scientist at Lawrence Berkeley National Laboratory. From 2000 to 2005, the aggregate electricity use by data centers doubled. The cloud, he calculates, consumes 1 to 2 percent of the world's electricity.

Much of this is due simply to growth in the number of servers and the Internet itself. A Google search is not without environmental consequence — 0.2 grams of CO₂ per search, the company claims — but based on E.P.A. assumptions, an average car trip to the library consumes some 4,500 times the energy of a Google search while a page of newsprint uses some 350 times more energy. Data centers, however, are loaded with inefficiencies, including loss of power as it is distributed through the system. It has historically taken nearly as much wattage to cool the servers as it does to run them. Many servers are simply “comatose.” “Ten to 30 percent of servers are just sitting there doing nothing,” Koomey says. “Somebody in some department had a server doing this unique thing for a while and then stopped using it.” Because of the complexity of the network architecture — in which the role of any one server might not be clear or may have simply been forgotten — turning off a server can create more problems (e.g., service outages) than simply leaving it on.

As servers become more powerful, more kilowatts are needed to run and cool them; square footage in data centers is eaten up not by servers but by power. As data centers grow to unprecedented scales — Google recently reported that one of its data centers holds more than 45,000 servers (only a handful of companies have that many total servers) — attention has shifted to making servers less energy intensive. One approach is to improve the flow of air in the data center, through computational fluid-dynamics modeling. “Each of these servers could take input air at about 80 degrees,” John Sontag, director of the technology transfer office at [Hewlett-Packard](#), told me as we walked through the company's research lab in Palo Alto. “The reason why you run it at 57 is you're not actually sure you can deliver cold air” everywhere it is needed. Chandrakant Patel, director of the Sustainable I.T. Ecosystem Lab at H.P., argues there has been “gross overprovisioning” of cooling in data centers. “Why should all the air-conditioners run full time in the data center?” he asks. “They should be turned down based on the need.”

Power looms larger than space in the data center's future — the data-center group Afcom predicts that in the next five years, more than 90 percent of companies' data centers will be interrupted at least once because of power constrictions. As James Hamilton of Amazon Web Services observed recently at a Google-hosted data-center-efficiency summit, there is no Moore's Law for power — while servers become progressively more powerful (and cheaper to deploy) and software boosts server productivity, the cost of energy (as well as water, needed for cooling) stays constant or rises. Uptime's Brill notes that while it once took 30 to 50 years for electricity costs to match the cost of the server itself, the electricity on a low-end server will now exceed the server cost itself in less than four years — which is why the geography of the cloud has migrated to lower-rate areas.

The huge power draws have spurred innovation in the form factor of the data center itself. For its Chicago center, Microsoft is outfitting half the building with shipping containers packed with servers. “Imagine a data center that’s about 30 megawatts, with standard industry average density numbers you can probably fit 25,000 to 30,000 servers in a facility like that,” says Microsoft’s Chrapaty. “You can do 10 times that in a container-based facility, because the containers offer power density numbers that are very hard to realize in a standard rack-mount environment.”

The containers — which are pre-equipped with racks of servers and thus are essentially what is known in the trade as plug-and-play — are shipped by truck direct from the original equipment manufacturer and attached to a central spine. “You can literally walk into that building on the first floor and you’d be hard pressed to tell that building apart from a truck-logistics depot,” says Manos, who has since left Microsoft to join Digital Realty Trust. “Once the containers get on site, we plug in power, water, network connectivity, and the boxes inside wake up, figure out which property group they belong to and start imaging themselves. There’s very little need for people.”

“Our perspective long term is: It’s not a building, it’s a piece of equipment,” says Daniel Costello, Microsoft’s director of data-center research, “and the enclosure is not there to protect human occupancy; it’s there to protect the equipment.”

From here, it is easy to imagine gradually doing away with the building itself, and its cooling requirements, which is, in part, what Microsoft is doing next, with its Gen 4 data center in Dublin. One section of the facility consists of a series of containers, essentially parked and stacked amid other modular equipment — with no roof or walls. It will use outside air for cooling. On our drive to Tukwila, Manos gestured to an electrical substation, a collection of transformers grouped behind a chain-link fence. “We’re at the beginning of the information utility,” he said. “The past is big monolithic buildings. The future looks more like a substation — the data center represents the information substation of tomorrow.”

Tom Vanderbilt is the author of “Traffic: Why We Drive the Way We Do (and What It Says About Us).”

[Copyright 2009 The New York Times Company](#)

[Privacy Policy](#) | [Terms of Service](#) | [Search](#) | [Corrections](#) | [RSS](#) | [First Look](#) | [Help](#) | [Contact Us](#) | [Work for Us](#) | [Site Map](#)
