



MIT



[<< Return to article](#)

Google and Akamai: Cult of Secrecy vs. Kingdom of Openness

The king of search is tapping into what may be the largest grid of computers on the planet. And it remains extraordinarily secretive about its core technologies—perhaps because it senses a potential competitor in dotcom era flameout Akamai.

By Simson Garfinkel
April 21, 2004

“You should never trust this number,” said Martin Farach-Colton, a professor of computer science at Rutgers University, speaking a little more than a year ago. “People make a big deal about it, and it’s not true.”



Farach-Colton was giving a public lecture about his two-year sabbatical working at Google. The number that he was disparaging was in the middle of his PowerPoint slide:

- 150 million queries/day

The next slide had a few more numbers:

- 1,000 queries/sec (peak)
- 10,000+ servers
- More than 4 tera-ops/sec at daily peak
- Index: 3 billion Web pages
- 4 billion total docs
- 4+ petabytes disk storage

A few people in the audience started to giggle: the Google figures didn't add up.

This
generation's
consumer
expects the
next
generation
of wireless



Click Here to Learn How

SPONSORED LINKS

[HP notebooks and desktops.](#)
[Doctor-patient security.](#)

[Learn about the Qualcomm
Launchpad™ Suite of application
Technologies.](#)

I started running the numbers myself. Let's see: "4 tera-ops/sec" means 4,000 billion operations per second; a top-of-the-line server can do perhaps two billion operations per second, so that translates to perhaps 2,000 servers—not 10,000. Four petabytes is 4×10^{15} bytes of storage; spread that over 10,000 servers and you'd have 400 gigabytes per server, which again seems wrong, since Farach-Colton had previously said that Google puts two 80-gigabyte hard drives into each server.

And then there is that issue of 150 million queries per day. If the system is handling a peak load of 1,000 queries per second, that translates to a peak rate of 86.4 million queries per day—or perhaps 40 million queries per day if you assume that the system spends only half its time at peak capacity. No matter how you crank the math, Google's statistics are not self-consistent.

"These numbers are all crazily low," Farach-Colton continued. "Google always reports much, much lower numbers than are true."

Whenever somebody from Google puts together a new presentation, he explained, the PR department vets the talk and hacks down the numbers. Originally, he said, the slide with the numbers said that 1,000 queries/sec was the "minimum" rate, not the peak. "We have 10,000-plus servers. That's plus a lot."

Just as Google's search engine comes back instantly and seemingly effortlessly with a response to any query that you throw it, hiding the true difficulty of the task from users, the company also wants its competitors kept in the dark about the difficulty of the problem. After all, if Google publicized how many pages it has indexed and how many computers it has in its data centers around the world, search competitors like Yahoo!, Teoma, and Mooter would know how much capital they had to raise in order to have a hope of displacing the king at the top of the hill.

Google has at times had a hard time keeping its story straight. When vice president of engineering Urs Hoelzle gave a talk about Google's Linux clusters at the University of Washington in November of 2002, he repeated that figure of 1,000 queries per second—but he said that the measure was made at 2:00 a.m. on December 25, 2001. His point, obvious to everybody in the room, is that even by November 2002, Google was doing a lot more than 1,000 queries per second—just how many more, though, was anybody's guess.

The facts may be seeping out. Last Thanksgiving, the New York Times reported that Google had crossed the 100,000-server mark. If true, that means Google is operating perhaps the largest grid of computers on the planet. "The simple fact that they can build and operate data centers of that size is astounding," says Peter Christy, co-founder of the NetsEdge Research Group, a market research and strategy firm in Silicon Valley. Christy, who has worked in the industry for more than 30 years, is astounded by the scale of Google's systems and the company's competence in operating them. "I don't think that there is anyone close."

It's this ability to build and operate incredibly dense clusters that is as much as anything else the secret of Google's success. And the reason, explains Marissa Mayer, the company's director of consumer Web products, has to do with the way

that Google started at Stanford.

Instead of getting a few fast computers and running them to the max, Mayer explained at a recruiting event at MIT, founders Sergey Brin and Larry Page had to make due with hand-me-downs from Stanford's computer science department. They would go to the loading dock to see who was getting new computers, then ask if they could have the old, obsolete machines that the new ones were replacing. Thus, from the very beginning, Brin and Page were forced to develop distributed algorithms that ran on a network of not-very-reliable machines.

Today this philosophy is built into the company's DNA. Google buys the cheapest computers that it can find and crams them in racks and racks in its six (or more) data centers. "PCs are reasonably reliable, but if you have a thousand of them, one is going to fail every day," said Hoelzle. "So if you can just buy 10 percent extra, it's still cheaper than buying a more reliable machine."

Working at Google, an engineer told me recently, is the nearest you can get to having an unlimited amount of computing power at your disposal.

The Kingdom of Openness

There is another company that has perfected the art of running massive numbers of computers with a comparatively tiny staff. That company is Akamai.

Akamai isn't a household word now, but it did make the front pages when the company went public in November 1999 with what was, at the time, the fourth most successful initial public offering in history. Akamai's stock soared and made billionaires of its founders. In the years that followed, however, Akamai has fallen on hard times. It wasn't just the dot-com crash that caused significant layoffs and the abandonment of the company's California offices: Akamai's cofounder and chief technology officer Danny Lewin was aboard American Airlines Flight 11 on September 11 and was killed when the plane was flown into the World Trade Center. Company morale was devastated.

Akamai's network operates on the same complexity scale as Google's. Although Akamai has only 14,000 machines, those servers are located in 2,500 different locations scattered around the globe. The servers are used by companies like CNN and Microsoft to deliver Web pages. Just as Google's servers are used by practically everyone on the Internet today, so are Akamai's.

Because of their scale, both Akamai and Google have had to develop tools and techniques for managing these machines, debugging performance problems, and handling errors. This isn't software that a company can buy off the shelf—they require laborious in-house development. It is, in fact, software that is one of Akamai's key competitive advantages.

Yes, a few other organizations are also running large clusters of computers. Both NASA's Ames Research Center and Virginia Tech have large clusters devoted to scientific computing. But there are key differences between these systems and the clusters that both Google and Akamai have created. The scientific systems are located in a single place, not spread all over the world. They are generally not directly exposed to the Internet. And perhaps most importantly, the scientific systems are not providing a commodity service to hundreds of millions of Internet

users every day: Google and Akamai must deliver 100 percent uptime. It's easy to go out and buy 10,000 computers—all you need is cash. It's much harder to make those computers all work together as a single service that supports millions of simultaneous users.

To be fair, there are important differences between Google and Akamai—differences that assure that Google won't be breaking into Akamai's business anytime soon, nor Akamai moving into Google's. Both companies have developed infrastructure for running massively parallel systems, but the applications that they are running on top of those systems is different. Google's primary application is a search engine. Akamai, by contrast, has developed a system for delivering Web pages, streaming media, and a variety of other standard Internet protocols.

Another important difference, says Christy, "is that Akamai has had a very hard time creating a clear business model that works, whereas Google has been unbelievably successful." Akamai has thus started looking for new ways that it can sell services that only a massive distributed network can deliver. Struggling for profitability, the company has been aggressively looking for new opportunities for its technology. This might be the reason that Akamai, unlike Google, was willing to be interviewed for this article.

"We started with basic bit delivery—objects, photos, banners, ads," says Tom Leighton, Akamai's chief scientist. "We do it locally. Make it fast. Make it reliable. Make the sites better."

Now Akamai is developing techniques for letting customers run their applications directly on the company's distributed servers. Leighton says that 25 of Akamai's largest customers have done this. The system can handle sudden surges, making it ideal for cases where it is impossible to anticipate demand.

For example, says Leighton, Akamai's network was used to handle a keyboard giveaway contest sponsored by Logitech. Thinking that its contest might be popular, Logitech created an elaborate series of rules, assuring that only so many keyboards would be given away to every state and within any given time period. But Logitech grossly underestimated how many people would click in to the contest. In the past, such underestimates have caused highly publicized Internet events like the Victoria's Secret webcast to crash, frustrating millions of Web surfers and embarrassing the company. But not this time: Logitech's contest ran on the Akamai network without a hitch.

Of course, Logitech could have tried to build the system itself. It could have designed and tested a server capable of handling 100 simultaneous users. That server might cost \$5,000. Then Logitech could have bought 20 of those servers for \$100,000 and put them in a data center. But a single data center could get congested, so it might make more sense to put 10 of them in one data center on the East Coast and 10 in another data center on the West Coast. Still, that system could only handle 2,000 simultaneous users: it might be better to buy 100 servers, for a total cost of \$500,000, and put them at 10 different data centers. But even if they had done this, the engineers at Logitech would have had no way of knowing if the system would actually have worked when it was put to the test—and they would have invested a huge amount of money in engineering that wouldn't have been needed after the event.

And contests aren't the only thing that can run on Akamai's network. Practically any program written in the Java programming language can run on the company's infrastructure. The system can handle mortgage applications, catalogs, and electronic shopping carts. Akamai even runs the backend for Apple's iTunes 99-cent music service.

Perhaps because Akamai is so proud of the system that it has built, the company is very open about the network's technical details. Its network operations center in Cambridge, MA, has a glass wall allowing visitors to see a big screen with statistics. When I visited the company in January, the screen said that Akamai was serving 591,763 hits per second, with 14,372 CPUs online, 14,563 gigahertz of total processing power, and 650 terabytes of total storage. On April 14, the number had jumped to a peak rate of 900,000 hits per second and 43.71 billion requests delivered in a 24-hour period. (Akamai wouldn't disclose the number of CPUs online because that number is part of its quarterly earnings report, to be released on April 28. "But it hasn't changed much," the company's spokesperson told me.)

Mail and Scale

Looking forward, a few business opportunities have obvious appeal to both Google and Akamai. For example, both companies could take their experience in building large-scale distributed clusters to create a massive backup system for small businesses and home PC users. Or they could take over management of home PCs, turning them into smart terminals running applications on remote servers. This would let PC users escape the drudgery of administering their own machines, installing new applications, and keeping anti-virus programs up to date.

And then there is e-mail. Back on April 1, Google announced that it was going to enter the consumer e-mail business with an unorthodox press release: "Search is Number Two Online Activity—Email is Number One: 'Heck, Yeah,' Say Google Founders."

Since then, Google has received considerable publicity for the announced design of its Gmail (Google Mail) offering. The free service promises consumers one gigabyte of mail storage (more than a hundred times the storage offered by other Web mail providers), astounding search through mail archives, and the promise that consumers will never need to delete an e-mail message again. At first many people thought that the announcement was an April Fools joke—a gigabyte per user just seemed like too much storage. But since the vast majority of users won't use that much storage, what Google's promise really says is that Google can buy new hard drives faster than the Internet's users can fill them up. *[Editor's note: Google's proposal to fund Gmail by showing advertisements based on the content of users' e-mail has received significant criticism from a variety of privacy activists. Earlier this month a number of privacy activists circulated a letter asking Google to not launch Gmail until these privacy issues had been resolved. Simson Garfinkel signed that letter as a supporter after this article was written but before its publication.]*

Google's infrastructure seems well-suited to the deployment of a service like Gmail. Last summer Google published a technical paper called The Google File System (GFS), which is apparently the underlying technology developed by Google for allowing high-speed replication and access of data throughout its clusters. With GFS, each user's e-mail could be replicated between several different Google clusters; when users log into Gmail their Web browser could

automatically be directed to the closest cluster that had a copy of their messages.

This is hard technology to get right—and exactly the kind of system that Akamai has been developing for the past six years. In fact, there's no reason, in principle, why Akamai couldn't deploy a similar large-scale e-mail system fairly easily on its own servers. No reason, that is, except for the company's philosophy.

Leighton doesn't think that Akamai would move into any business that required the company to deal directly with end users. More likely, he says, Akamai would provide the infrastructure to some other company that would be in a position to do the billing, customer support, and marketing to end users. "Our focus is selling into the enterprise," he says.

George Hamilton, an analyst at the Yankee Group who covers enterprise computing and networking, agrees. Hamilton calls the idea of Google competing with Akamai "far-fetched." But Google could hire Akamai to supplement Google's technology needs, he says.

Still, such a partnership seems unlikely—at least on the surface. Google might buy Akamai, the way the company bought Pyra Labs in February 2003 to acquire Pyra's Blogger personal Web publishing system. But Akamai, with its culture of openness, doesn't seem like a good match to secretive Google's. Then there is the fact that 20 percent of Akamai's revenue now comes directly from Microsoft, according to Akamai's November 2003 quarterly report. Google's rivalry with Microsoft in Internet search (and now in e-mail) has been widely commented upon in the press; it is unlikely that the company would want to work so closely with such a close Microsoft partner.

Ted Schadler, a vice president at the market research firm Forrester, says that it's possible to envision the two companies competing because they are both going after the same opportunity in massive, distributed computing. "In that sense, they have the same vision. They have to build out a lot of the same technology because it doesn't exist. They are having to learn lots of the same lessons and develop lots of the same technologies and business models."

Schadler says Akamai and Google are both examples of what he calls "programmable Internet business channels." These channels are companies that offer large infrastructure that can offer high quality services on the Internet to hundreds of millions of users at the flick of a switch. Google and Akamai are such companies, but so are Amazon.com, eBay and even Yahoo!. "They are all services that enable business activity—foundation services that [can be] scaled securely," Schadler says.

"If I were a betting man," Schadler adds, "I would say that Google is much more interested in serving the customer and Akamai is more interested in provide the infrastructure—it's retail versus wholesale. There will be lots and lots of these retail-oriented services."

If true, Google might suddenly find itself competing with a company that, like Google itself, seemed to come out of nowhere. Except this time, that company wouldn't have to figure out any of the tricks of running the massive infrastructure

itself.

And that explains why Google is so secretive.

Simson Garfinkel is the author of 12 books on information technology and its impact.

Copyright 2004 Technology Review, Inc. All rights reserved