Maximizing traffic is priority number one for most sites, but popularity has its price. Here's how some of the big-leaguers coped when the surfing public loved their sites to death.

BY SIMSON L. GARFINKEL

Crash Course

GARRY KASPAROV WAS THE WORLD'S REIGNING CHESS CHAMPION. Deep Blue was the **chess-playing** supercomputer built by IBM Corp.'s Thomas J. Watson Research Center in Yorktown Heights, N.Y. And while chess was ostensibly the mainspring of Deep Blue's existence, IBM hadn't spent **millions of dollars** developing the system just to play games.

Rather, **Deep Blue** was created to explore the limits of massively parallel processing systems. Chess just happened to be the closest thing to a perfect test of those limits, and Garry Kasparov happened to be the closest thing to a perfect test for Deep Blue.

The match was scheduled to begin on Feb. 10, 1996—the launch of a **year-long celebration** leading up to the 50th anniversary of the Association for Computing Machinery's founding in 1947. The ACM had put up a $500,000 purse, but even more attractive to IBM than the money was **the prestige that would accompany a win**. To make the event as public as possible, the company wanted to post the entire match on one of the slickest, quickest and most robust World Wide Web sites in the world.

The project team knew going in that the slightest glitch would be a major embarrassment. They also knew how easy it is to be embarrassed on the Web. Even as the IBM team labored away on its ACM Chess Challenge site, **hordes of football fans** overran and crashed **the Super Bowl site** hosted by Microsoft Corp., one of the few companies that was supposed to know its way around the Web.

ILLUSTRATION BY HENRIK DRESCHER

"Our servers were not crashing. They were performing well considering the load they were running under. It's just that there were not enough of them."

–Bill Koszewski

THE SUPERVISION OF IBM'S CHESS SITE fell to Carol Moore, the company's director of corporate Internet programs. Moore hired K2 Design Inc., a Web design company with which IBM had worked in the past. K2's first problem was time: The site had to be up and running in one month.

IBM brought on Jake Prescott, then creative director at advertising company Ogilvy & Mather Worldwide, as a consultant. A recreational chess player who had worked with Moore on other Web projects, Prescott had been an early promoter of the match—even before Kasparov had signed on—and he wanted the site done right. Prescott became the project's brand steward and head copywriter. He wrote essays on the meaning of this computer chess event that pitted man against machine. He created personalities for the real-time chat rooms to keep the talk on target. And he watched as the part of the site that delivered still video images of Kasparov came together. There was also, of course, a computer-generated chessboard tracking the game's progress, move by move.

As the date closed in, Moore asked her colleagues at IBM how big a server she would need to host the event. "The consensus was that you shouldn't underestimate," she recalls. IBM had hosted the official Web site for the U.S. Open tennis tournament, and that site had attracted more than a million hits a day. Moore figured that the chess tournament would draw 100,000. To be safe, she told K2 to plan for 200,000 hits a day. K2 agreed.

In order to complete the project on time, K2 turned over much of the labor to a small Web hosting outfit in Cambridge, Mass., a reputable company that K2 had recently bought. Planning for the worst, the subcontractor set up a high-end workstation with 512MB of RAM and connected to the Internet via a 10Mbps link.

The site opened for business at noon on Saturday, Feb. 10th. The match started at 3 p.m. in the Lecture Hall of the Pennsylvania Convention Center in Philadelphia. "I was there in the chat lines," recalls Prescott, who had gone so far as to recruit international chess masters for play-by-play commentary. "I could see people rolling in from Norway, Australia—from all over the world—enthusiastic about chess, wanting to know what was going on, talking about the algorithms that were driving Deep Blue."

Shortly after the match began, something went wrong. The server slowed, then started to crawl. "We couldn't get into our own Web site to maintain the chat room," Prescott says. "It was pretty much paralyzed by the time the match got going. Forget about seeing the chessboard and the real-time stuff. Just getting into the Web chat was hard. I kept hammering away at it and got in one time out of 10."

On the stage in Philadelphia, Kasparov was in trouble as well. He would eventually lose the game—at that time his only loss to a machine. But out on the Internet, the reverse was happening. Hundreds of thousands of Internet surfers, curious to see the "ultimate" battle between man and machine, had brought another computer to its knees.

During the peak of the tournament, IBM says, the Web server was nailed with more than 1 million accesses per hour—impressive but, according to the group's estimates, not enough to cause such chaos. For every user who downloaded an HTML page, there were another 50 whose clicks went unanswered. It appeared that IBM needed a site that could handle not 200,000 hits a day but 10 million.

DAVE GROSSMAN'S TELEPHONE started ringing at 7 a.m. on Sunday, Feb. 11. Grossman, a senior staff member at IBM's Internet Division Advanced Technology Group in White Plains, N.Y., was part of an elite team assembled by IBM to build systems capable of hosting huge Web "events." He attacked the problem by bringing in the WOM, or Web Object Manager, a superpowerful server running on IBM's SP-2, a multiprocessor supercomputer that is a close cousin of Deep Blue (see box, Page 44).

The WOM had handled the U.S. Open Web site as well as the site for The Masters golf tournament and had performed flawlessly both times. Transferring the site to the WOM was an obvious solution. But with the second game scheduled for that afternoon, there simply wasn't time.

So Grossman tried the next-best thing. Using a technique called "inlining," he moved the big, bandwidth-hogging files from the server in Cambridge to the SP-2. He left on the original site the Common Gateway Interface (CGI) scripts, the HTML templates, the chat rooms and the control logic that made the whole thing work.

But when the match resumed at 3 p.m., so did the chaos. Once again, there were millions of hits on the original workstation. Once again, the chess pages slowed to a crawl. Once again, the chat rooms

Deep Blue Is Back

Next May, a new, improved Deep Blue will seek revenge against world chess champion Garry Kasparov. IBM scientists say the modified computer is capable of beating last year's model three times out of four.

INSIDe THe WOM

IBM's WOM Web server runs on an IBM SP-2 scalable multiprocessor computer. Each processor is a high-speed RS/6000 computer with 128MB of RAM and an ultrafast connection to a network of 67MHz PowerPC processors. Running on top of that is a special version of IBM's AIX operating system, a modified version of the Apache Web server that runs on the multiple processors, and an operating system that automatically spreads incoming Web requests across the various processors. A 45Mbps T3 connection to the Internet's backbone feeds the whole thing.

Koszewski says the problem was not software but hardware. "Our servers were not crashing," he says. "They were performing well considering the load they were running under. It's just that there were not enough of them."

A different problem had plagued another high-volume Microsoft site in July, two days after MSNBC, the joint Microsoft/National Broadcasting Co. news site, went live (see "The Peacock's New Feathers," *WebMaster*, September 1996). When TWA Flight 800 crashed into the Atlantic, a flood of Internet users called up the site for the latest report, and many were turned away. In that case, says Jeff Reisman, a group program manager at MSNBC, the problem was that servers were running different versions of the Open Database Connectivity (ODBC) driver, which lets the Web server access information stored in relational databases. Reisman says that in an all-out push to get the mammoth site up by the launch date, no one noticed the problem. The result was a slow memory leak that eventually stopped Internet service until the service was restarted.

"We had incredible time pressures," Reisman says apologetically. "We built the hardware in the last two weeks, with many on our team working extremely long hours."

WHILE MANY OBSERVERS blame such incidents on rapidly increasing Internet volume, the problem has been around as long as the Web itself. In August 1994, for example, when *www.playboy.com* went online, a monstrous traffic jam kept out thousands of surfers who were doubtless eager to read the articles. Worried that the site would be swamped, Eileen Kent, Playboy Enterprises Inc.'s vice president of new media, had planned ahead and purchased a T1 line. But it wasn't good enough. As soon as the Playboy site was launched, Kent began receiving reports of rejected connections and pages that would come up only in pieces.

The problem, Kent says, was a bug that caused all TCP connections to try to reconnect after they had been dropped, resulting in a backlog of stuck connections. Playboy increased the default backlog of sockets so the server could handle 80 HTTP connections per second, and the site ran smoothly.

"I do recall sometimes I was really sweating," says Kent. "But it didn't last long."

Today Playboy's site runs on a cluster of Silicon Graphics Inc. Challenge machines, with one Sun Microsystems Inc. Sparcstation thrown in. And Kent has increased the bandwidth: *playboy.com* is now fed by two T3 lines and two OC-3's connected to multiple providers across the country.

Playboy's experience taught it an important lesson: Simply making sure that you have enough bandwidth and RAM isn't enough. If you configure unwisely or forget details, even a site with the fastest network connection running on the most expensive hardware can feel like a beleaguered 486 PC connected to the Internet by a 28.8Kbps modem.

Those were the same lessons learned—and finally applied—by IBM. On Tuesday, Feb. 13, when hundreds of thousands of people clicked into The ACM Chess Challenge site, the WOM smoothly served up the page containing the chessboard. And surfers with the right kind of browser were treated to a Virtual Reality Markup Language version of the board that let them fly around among the chess pieces.

On that day, Kasparov and Deep Blue battled inconclusively. The following morning, *USA Today*'s headline read "Game 3 Chess Match a Draw; IBM Site Up." Deep Blue's fortunes declined over the remainder of the match: In the fifth game of the six-game series, it rejected Kasparov's offer of a draw and went on to lose, four games to two. WM

Simson L. Garfinkel is a freelance technology journalist. His book Practical Unix and Internet Security *was published last year by O'Reilly & Associates. He can be reached at* simsong@vineyard.net.

FINDING IT ONLINE

Association for Computing Machinery
www.acm.org

The ACM Chess Challenge
www.chess.ibm.park.org/

IBM Corp.
www.ibm.com

K2 Design Inc.
www.k2design.com

Microsoft Corp.
www.microsoft.com

Ogilvy & Mather Worldwide
www.ogilvy.com

Playboy Enterprises Inc.
www.playboy.com

Silicon Graphics Inc.
www.sgi.com

Sun Microsystems Inc.
www.sun.com

Super Bowl
www.superbowl.com

WOM: The Boiler Room
www.womplex.ibm.com/

were unusable. And once again the Web site was, in the highly technical language of Grossman's group, "toasted."

The moment Kasparov emerged triumphant, Grossman and his team tore the Web site apart, intending to rebuild it on the WOM. As it turned out, they didn't have to. In the process of dismantling the site, they discovered a strange inefficiency in the software used to build the chessboard.

The program, which had worked well on less trafficked chess matches and had been obtained by K2's subcontractor, was designed to send one graphics interchange format (GIF) file for the blank board—and then up to 32 more pictures, one for each chess piece. Each time someone tried to look at the chessboard, 33 different connections had to be set up and then taken down. That load was breaking the Web site's back.

Dave Grainger and Sean Martin, both from IBM's Hursley Laboratories in England and on assignment to Grossman's team, wrote a program that reduced the 33 GIF files to a single image. In the time left, they wrote another program to process the log files from Deep Blue and display the intermediate results—what Deep Blue was actually "thinking"—on the Web site. Meanwhile, Andy Stanford-Clark, a scientist at Hursley, modified the SP-2's Apache Web server, building into it many programs that were formerly run as CGI scripts.

By Tuesday morning the WOM was ready. And the media were waiting. When Grossman picked up *USA Today*, he saw a front-page headline that read "IBM Site Stalemates Users."

IBM isn't the only high-technology company that has suffered public embarrassment at the hands of the unpredictable Web. At the same time IBM was putting together the chess site, Microsoft was learning from its Super Bowl fiasco and trying to reconfigure its Internet connection to cope with massive demands on its Internet add-on product site. The company installed a T3 connection that could pump data through a fiber-optic cable at 45Mbps. Yet despite that impressive speed, the new connection was overloaded in a matter of weeks.

In early spring, Microsoft hooked up a second T3 connection to relieve the burden. Each connection led directly to one of the nation's two largest Internet providers, Sprint Communications Co. and MCI Communications Corp. That double-barreled solution worked until late May, when Microsoft posted a beta version of Internet Explorer 3.0. Once again, the lines were taxed beyond their capabilities.

The software giant was learning from its mistakes, but it was learning slowly. Preparing for the official release of Explorer 3.0 in August, Microsoft installed eight T3 lines to its Redmond, Wash., site and arranged for mirror sites in San Jose, Calif., and Washington, D.C., to hold copies of Explorer. Finally, it looked like all was well. Microsoft was clocking 20,000 simultaneous connections to its Web servers, according to Internet Explorer Product Manager Bill Koszewski, and the company's data circuits were still at only 50 percent of capacity.

That was the good news. The bad news was that downloading Explorer is only half of the installation process. The second half requires users to register, providing such information as name, address and telephone number. To handle that traffic, Microsoft had set up a pool of 16 Compaq Computer Corp. servers, each equipped with four Pentium processors running in parallel. It was an impressive battery of machines, but it wasn't enough. In the days immediately following the rollout, most people who tried to connect were greeted with the message "Server busy; try again later."

"We owe our users an apology," says Koszewski. "This has happened before."

GAMe PLANS

Those who do not learn from history are doomed to repeat it. Here are some of the lessons of The ACM Chess Challenge site—and tips for avoiding its problems.

Be Prepared. It's impossible to predict how popular a site will be, *but try*. IBM's original chess site was built to handle 200,000 hits a day. It should have been built for 200,000 hits an hour.

Have Backups in Reserve. Always have a second, fully configured machine that you can swap into your system if problems develop with the first. Have a contingency plan in the event that your backup doesn't work.

Know Your System. IBM's "stalemated" Web site ended up on the front page of *USA Today* because a critical piece of software that had been outsourced to a subcontractor didn't work the way it was expected to work. If your system is critical, make sure that somebody on your staff tests every function.

Have Weekend Support. IBM saved itself further embarrassment because it was able to assemble its White Plains team on a Sunday morning and have it work through the day and late into the night.

Test Your Site. If possible, test your Web site with a realistic load. But even if you test it, remember that something can always go wrong.

"Forget about seeing the chessboard and the real-time stuff. Just getting into the Web chat was hard. I kept hammering away at it and got in one time out of 10."

–Jake Prescott