BUSINESS 2.0

SEPTEMBER 2002

computing to the Nth degree

Sun Microsystems is betting its future on a radical new technology called N1. If it works, it could revitalize the troubled Silicon Valley pioneer—and change the way the world thinks about computing.

BY ERICK SCHONFELD

THE MEETING HAD RESUMED, AND ROB GINGELL was late to return—not a good move, since it was one of the most important meetings of his career. It was last summer, and Gingell, then Sun Microsystems's chief software engineer, had an excuse: His twin-engine Cessna had broken down, and he'd lost track of time while he gabbed on the phone with his mechanic. That wasn't likely to appease Sun's famously tart-tongued CEO, Scott McNealy, who was getting his introductory briefing on a vital new technology initiative that happened to be Gingell's brainchild. When Gingell huffed back in, 15 minutes late, the boss glared. "You should not be naming anything," McNealy snapped.

Gingell was momentarily taken aback—until he realized that McNealy had started to smile, and that the jibe was a good-natured poke at Gingell for saddling one of the most crucial technology gambles in Sun's storied history with the decidedly clunky name N1. But McNealy would come to tolerate the name, and he loved the concept. In fact, he has essentially staked Sun's future on it.

Until now Sun has spoken about N1 only sparingly and in broad strokes. Many of its details are still under wraps. But in a series of interviews with *Business 2.0,* Sun technologists and executives have laid out the strategy more fully and explained the thinking behind it. It's not an easy thing to grasp: Jonathan Schwartz, senior vice president for software, likens talking about N1 to "a drug-induced experience that becomes very abstract." For the moment, think of N1 as a kind of operating system for computer networks—and Sun's approach to solving one of today's great business quandaries: the mind-boggling complexity of the typical corporate data center. Through software that eventually will be built into everything Sun makes, from servers to workstations to storage devices, N1 acts like an Australian sheepdog corralling the unruly herd of corporate network components so that managing them all for maximum performance becomes a snap.

Beleaguered CIOs everywhere will rejoice.

Or so Sun is betting. "We are reinventing ourselves around N1," says chief technology officer Greg Papadopoulos. And not a moment too soon. Sun is one of the handful of great and durable Silicon Valley companies. It's responsible for some of the most innovative technologies in computing history, from the workstation to the Java programming language. But Sun has been hit hard by the tech slump, executive defections, and a technological landscape that is shifting under its feet. In the fiscal year that ended June 30, Sun's revenues shrank to $13 billion from $18 billion. Its stock price recently sank below $4 a share, down from $64 two years ago. Its flagship servers, which run mostly on Solaris, its proprietary and expensive version of the Unix operating system, are losing ground to machines running cheaper Linux and Windows operating systems. "The water in the room is rising" at Sun, says Ken Smith, who co-manages the Munder Future Technology Fund.

More important, selling computer hardware—the bread-and-butter business that made Sun—increasingly looks like a fool's game. The entire industry is undergoing massive commodification; Sun's gross margins have fallen to 41 percent from 52 percent four years ago. Competitors are already dealing with the problem. IBM, for instance, has shifted its focus more to software and IT services. Hewlett-Packard, newly merged with Compaq, is trying to follow suit. But Sun has neither the inclination nor the wherewithal to build up an army of tens of thousands of IT consultants or to operate hosting centers for corporate clients. Instead, it thinks it can engineer its way to success. "You don't need as big an army if they are competing with bows and arrows and you have machine guns," says Gingell, who now oversees all of Sun's engineers.

N1, Sun believes, is heavy artillery. "It turns out customers don't want software or services," says venture capitalist and Sun board member John Doerr. What they want, he says, is network computing that just works—without all the attendant hassles. Far from trying to find a way out of the hardware business, Sun's strategy is to free itself from the commodity trap by making a new class of computers, based on N1 technology, that take the pain out of networked computing. If the approach works, it could not only breathe new life into Sun but also fundamentally change the way the world thinks about computing. And if it doesn't work ... well, as Gingell puts it, "without N1, Sun would fade away."

One reason it's hard to get a handle on N1 is that Sun hasn't fully translated it into commercial products yet. But N1's distinguishing concept is that it treats the disparate boxes in a customer's data center as one pool of resources. Just as an operating system in an individual computer manages resources, schedules jobs, and deals with certain types of failure, so does N1. The difference is that where a computer's operating system manages components such as the microprocessor, memory, and communications bus, N1 manages the components of the network itself: computers, storage devices, and switches. Today, in a typical corporate data center, a harried CIO often must dispatch squadrons of systems administrators to manually swap out boxes or rewire connections or load new software in response to the regular ebb and flow of business. Sun's N1-powered machines could end all that by pooling the network's resources and allowing CIOs to simply dip in whenever they need to make a change—to add storage or maximize processing power, for instance. In most cases, N1 could automatically make whatever adjustments are needed. N1 "is like the telephone industry moving beyond switchboards and operators to automatic switching," McNealy says. "If we didn't do N1, then we'd need everyone under the age of 30 to work as systems administrators." (For more on how N1 works, see "This Time, the Network Really Is the Computer," right.)

> Sun's strategy is to build a new class of computers, based on N1 technologies, that take the pain out of networked computing.

The closest thing Sun has today that hints at N1's power sits in a glassed-in room at a server ranch in Sunnyvale, Calif. Nearly 2,000 servers with nicknames like Bonanza, Hoss, and Lena purr on a raised floor. The floor is perforated, like the surface of an air hockey game, which allows cold air to pass up and replace the heated jets blasting out the back of each machine. But the truly hot technology—and what gives the scene its N1 flavor—is that the power of every single one of those servers is pooled so that their 7,500 processors can be used as one huge computer. Even the workstations on the chip designers' desks have to contribute one of their two processors to the collective pool. On average, 98 percent of the processors in the network are humming at any given time—three to five times the chip utilization rate of most corporate data centers. On a recent morning, Jim Gateley, who directs the server ranch, had 2,949 computer jobs running and 19,000 more queued for the next available processor. "In a simplified way, we are behaving the way N1 defines computing going forward," Gateley says.

Like many a big idea, N1 was born out of serendipitous office chatter. Gingell and chief hardware engineer Steve McKay were walking down a stairwell after a meeting in the summer of 2000. They had both been thinking about how the Java programming language was giving rise to applications that aren't tied to any particular machine and can run on any operating system. McKay had also been thinking about how advances in networking technologies were making it easier to break up a computer's components, such as storage, and spread them across the network. The result, McKay believed, would be that entire machines would begin to act like components of a larger networked system. He and Gingell agreed that, taken together, these software and hardware trends "would cause us to want to build computers in a different way, like they are a network rather than a box," McKay says.

The engineers' stairwell discussion led to the formation of the N1 team. Gingell recalls that "underlying our thinking was the

question of how we make this a problem that looks like something the company knows how to do." Sun is famous for the primacy of its engineers and the hard-core system designs they've done for decades, so the solution was characteristic: It decided to create an entire computing system out of the network. Gingell dubbed it N1, for Network 1. It was a name that harked back to Sun's roots; the company's first workstation, designed in 1982 by legendary engineers and Sun co-founders Andy Bechtolsheim and Bill Joy, was called the Sun1. "N1 is the blueprint for the next 20 years," says McKay, who now oversees the technology.

He chose Yousef Khalidi to run the N1 team. An intense engineer and a rapid-fire talker, Khalidi had made his mark at Sun by leading the development of its latest cluster technology, which allows a small group of servers to act as one machine (a precursor to N1). The project really came into focus when Khalidi and his team went out to talk to customers about their biggest technology headaches. Throughout 2001 the N1 team visited 35 CIOs and CTOs at Internet service providers, investment banks, and manufacturers. Dave Leonard, CTO of iStructure, the computer hosting subsidiary of Level 3 Communications, was one of them. "Our job is to lower the cost of computing," he told Khalidi. "We can't friggin' do it if we can't increase utilization." Like Leonard, the other customers were desperate to cut the cost of running their networks; as much as 70 percent of that cost is administrative.

Leonard's response to the N1 briefing was typical: "So when can I have it?"

As the team's understanding progressed, Khalidi discovered that there were dozens of groups at Sun that already addressed parts of the problem. For instance, the storage engineers were working on storage "virtualization"—the idea that data contained in many different storage devices appears as one pool of information no matter where those devices may physically reside, and is available to any computer on the network. Another product group was developing low-cost "blade" servers that look like baby N1s. Blades are stripped-down servers, about the size of three pizza boxes stacked on top of each other, that slide into a chassis filled with other blades. Each blade is like a miniserver that runs an application; software in the chassis provisions the blades and manages the interactions between them. The management software in the chassis does with fairly uniform blades the same thing that Sun wants N1 to do with heterogeneous machines in a data center.

Khalidi found a third predecessor to N1 in grid computing. Popular in scientific and technical circles, grid computing is basically a way to create a functional supercomputer by linking hundreds or thousands of servers together to share their processing power for computation-heavy tasks, such as decoding the human genome. But grids are beginning to be deployed in more commercial settings—such as in Sun's server ranch—because they allow for better utilization of computing resources. Even in the most efficient data centers, 70 to 80 percent of the processing power usually sits idle. Through deployment of a grid, applications can be run on any available server and utilization rates can surpass 90 percent. (Grid computing does some of what N1 promises, but N1 also encompasses storage and networking gear and can shift resources on the fly, which grids can't do.)

For its first year of existence, N1 was a closely guarded secret among the technologists working on it. But by August 2001, CTO Papadopoulos felt it was time to tell McNealy about it. McNealy's initial dislike of the N1 name didn't dim the CEO's enthusiasm. Raising computer utilization rates and relieving CIOs of some of the burdens of running a computer network were attractive concepts, whatever Sun called the technology. "My reaction was 'Great, how soon can we start selling it?'" McNealy recalls. The N1 team finished sketching out the technical underpinnings of N1 in December; now groups throughout the company are trying to figure out how to build N1 into everything Sun makes. Sun is already shipping bits and pieces of N1 in some of its latest gear, but won't seriously begin rolling out N1 products until next year. Gingell estimates that by 2005, roughly 85 percent of Sun's revenues could come from products that have N1 properties.

For the most part, N1 will manifest itself as software embedded in other products. Additionally, during the next couple of years, Sun will launch some stand-alone N1 products, such as systems management software for network administrators and a management blade server. To most customers, though, N1 will be invisible. It will just be stitched into the fabric of everything Sun does. And it will be compatible with the systems that customers currently have, whether they run on Sun technology or a rival's. Of course, Sun says, N1 will work best on its machines.

The N1 bet has huge risks. "I don't think you can really point to a period where what Sun is about has been more amorphous," says one analyst.

What could N1 mean for the future of computing? If it works, it opens up entirely new corporate computing options: Servers and storage could be shared by departments, and perhaps one day even by companies. But more broadly, outside IT firms could create their own N1 hosting centers to provide so-called utility computing services. The idea—also being pursued by HP and IBM—is to deliver computing the same way you deliver water or electricity; people would have access to almost limitless computing power whenever they need it, would pay only for what they use, and won't care about the underlying infrastructure. Computing would be cheaper—and almost pain-free.

Or not. For all the potential that Sun sees in N1, the strategy has many skeptics. "I don't think you can really point to a period where what they are about has been more amorphous," says SG Cowen analyst Richard Chu. There is also the real challenge of reorienting the company. People do not give up their religions easily. And at Sun, the religion of the past 20 years has been the supremacy of machines run by its Sparc

chips and its Solaris software. "You worry," Gingell says. "Do we have the courage to transcend our technologies?"

Meanwhile, Sun rivals are attacking the same problems as N1 from different angles. HP already offers data center management software with some N1-like features. And IBM's ambitious Eliza project is trying to create self-healing, self-managing networks. As Sun's chief marketing officer, John Loiacono, puts it: "Somebody is going to solve this problem. If it is not Sun, it will be someone else."

In the long term, of course, success could also have its price. One of the great ironies of the quest to increase the efficiency and utilization rates of existing computer networks is that if someone does figure out how to do it, people may be inclined to buy fewer computers. Even Sun may find it hard to engineer its way out of that one. ❖

Erick Schonfeld (eschonfeld@business2.com) is an editor at large for Business 2.0.