

Blades Have The Edge

Superslim machines are fomenting a quiet revolution in the server room

By Jane Wright

THE VAST, LABYRINTHINE COMPUTER ROOM of Veritas DGC Inc., in Houston, used to be a crowded, bustling place. Small groups of technicians would often hunker down there for hours in wired warrens formed by refrigerator-size racks housing more than 10 000 server—"nodes," in networking jargon. The technicians, known as "nodelers," worked day and night in that sprawling hive to keep all those machines up and running, processing the advanced geophysical and seismic models that are at the heart of Veritas's business.

These days, the oil industry consulting firm's computer room is much quieter and less crammed. The nodelers are still nodeling, but they spend considerably less time in the chilly computer room. Here's why: the number of servers that are down at any given time has dropped from 2 percent to 1 percent. It seems like a tiny improvement, but when you have 10 000 servers, that single percentage point means a hundred more units dutifully crunching data rather than idly waiting for a nodeler.



Similar transformations are happening in computer rooms all over the world, and they are all the result of the same decision: to switch from conventional servers to blade servers. Blades, for most information technology departments, offer huge improvements over conventional, rack-mounted units. This compact, slim computer (hence the name "blade") is typically based on the same Intel or AMD processors and Windows or Linux operating systems as most other servers, but it consumes a lot less power and takes up a lot less room.

Blades fit into enclosures that hold several units, usually vertically and side by side, like books on a shelf. To replace a bad blade, a technician need only pop it out of its enclosure and put in a new one. And blades are equipped with management programs that let staff easily set them up for specific applications or arrange them in special configurations.

The greatest benefit of blades, therefore, is their flexibility. They are much easier to install, manage, and repair than conventional servers—especially when you have hundreds or thousands of units. That was the case, for example, for Weta Digital Ltd., a special-effects company in Wellington, New Zealand, founded by filmmaker Peter Jackson. It used more than 500 IBM blade servers, each with two Intel processors, to bring to life the digital characters and scenes of the final installment of The Lord of the Rings trilogy, The Return of the King (2003), which won an Academy Award for best visual effects.

The Barcelona Supercomputer Center in Spain also turned to blades when it decided to build a massive computer cluster late last year. The center mounted 2282 IBM blade servers in 27 racks, which took up less than 50 square meters [see photo, "[A Mass of Blades](#)"]. A similar system built with regular servers would likely require more than double that area.

The machine, dubbed MareNostrum, took fourth place last November in the biannual Top500 ranking of supercomputers compiled by the University of Mannheim, in Germany, and the University of Tennessee. MareNostrum is now the most powerful computer in Europe, according to that list.

BLADES ARE THE FASTEST-GROWING SEGMENT in the overall server market, which last year had worldwide sales of more than US \$40 billion, according to Gartner Inc., in Stamford, Conn. Like other new technologies, blades still face a number of technical challenges, but more and more companies are already switching to them. The organizations that are buying blades today are mostly large ones, such as telecommunications and pharmaceutical companies, Internet service providers, and financial institutions.

Buyers are adopting blades for various reasons. Some, like Veritas, were seeking to concentrate their sprawling server farms and had ended up with computer rooms that were overcrowded and power hungry; blades in this case allowed them to free some rack space without losing computing power.

Others are going with blades because they are easily replaceable when failures occur. Some brokerage firms, for example, are using blades to run their equity trading systems, which need to be highly reliable so brokers can see up-to-the-second stock prices. These firms have so many servers that failures are routine events, and units need to be "hot-swappable": you pop the failed unit out of the enclosure and when you put in a new one, it goes on line automatically.

Yet other buyers have found that blades, despite costing an average of 10 percent more than comparably configured conventional units, bring technical benefits that outweigh their higher prices. Because of the comparative ease with which they can be upgraded or reconfigured, blades have lower administrative and service costs, which can make up for the higher initial outlay in as few as two or three years. Blade makers claim that as few as two significant changes to a server—upgrades or reconfigurations, say—over the typical three- to five-year life of a machine can justify the purchase of blades. In fact, the majority of blades' customers have hundreds, if not thousands, of servers to which they make changes about once a year.

Blades, which first hit the market in 2001, still account for less than 3 percent of the total server market, or about \$1 billion. But their adoption is accelerating [see chart, "[Blades Take Off](#)"]. Gartner's estimate of 340 000 units sold in 2004 is about double the total for the year before. Gartner expects that number to triple over the next four years, at which point one in every eight servers in operation in the world is likely to be a blade.

Today, heavyweights such as Dell, Fujitsu, Hewlett-Packard, IBM, and Sun are the biggest suppliers of blades; together these five companies supply more than 90 percent of the blade market [see chart, "[Slicing the Pie](#)"]. But the small group of newer and smaller vendors that share the rest of the pie—notably Egenera Inc., in Marlboro, Mass., and Verari Systems Inc. (formerly RackSaver), in San Diego—are finding that clever innovation can make up for what they lack in size and clout.

CLEARLY, A REVOLUTION IS AFOOT in the market for servers, the workhorse machines around which corporate and other networks are built. Servers anchor networks as repositories for programs and data, providing users access to databases, files, e-mail, Web pages, and countless other resources. They make up the second largest segment—a solid 15 percent—of the \$260-billion-plus worldwide computing hardware market.

Despite having been around for more than 30 years, conventional servers still have many drawbacks: they take up huge amounts of space and they use a lot of electricity. Even worse—perhaps their biggest shortcoming—is their relative inflexibility. Almost all servers are one-application machines; if you find yourself with more financial servers than you need, for example, you have to totally reprogram them before you can use any of the machines as a customer-service server.

First, consider the physical and electrical drawbacks. Ordinary servers are often rack-mounted units that are about the size of pizza boxes, but some are mainframes the size of refrigerators. These servers contain essentially the same components as your everyday desktop PC: a motherboard with processor and memory; a hard-disk drive; a network card; a power supply; input/output interfaces for keyboard, mouse, video, and USB connections; and often CD-ROM and possibly floppy disk drives.

A blade server, on the other hand, carries only the absolute essentials: a small motherboard with processor, memory, and networking circuitry, and maybe a hard drive. The result is a device that is less than half the size of the pizza-box-type server.

The special enclosures to which blades attach provide them with all the other resources they need to work. The enclosures contain the blades' power supply, a few fans for cooling, the disk drives, and interfaces for mouse, video, keyboard, USB, and network connections. By sharing these components instead of having their own, those blades require much less space, cabling, and power in comparison with conventional servers.

The upshot is that a standard 48.3-centimeter-wide, 1.88-meter-tall rack that would normally hold 42 pizza-box-size servers can instead house two to seven times that many blades [see illustration, "[The Blade Effect](#)"]. Think of a stuffed bookcase instead of a stack of pizza boxes. Ensclosed in those enclosures, the blades use no cables; the enclosure's built-in connectors feed the blades with power and data. In this way, blade systems reduce the cabling by up to 75 percent compared with traditional servers, whose cabling spaghetti has always been a headache to system administrators.

The advantages in power consumption can also be significant. Because they share power supplies and have fewer functional chips and other components, blades consume less electricity and generate less heat. How much electricity blades can save varies on a case-by-case basis, depending on the types of processors and applications being used, along with other parameters. But according to Gartner's analysis of equipment specifications and other vendor data, blades can need up to 50 percent less electricity than comparably configured traditional rack-mounted servers.

THE PROMISE OF BLADES, however, lies not so much in their hardware as in the software used to manage them. Blades' management programs let system administrators automate and simplify what are otherwise tedious, time-consuming tasks.

Consider the conversion of a financial server to a customer-service server. With most conventional servers, you would have to go to the computer room and disconnect the server's network and storage cables from the financial system and reconnect them to the customer-service system. Next, you would have to erase the server's drive and copy an operating system and the new applications onto it. Finally, you would need to configure the server's networking information to give it access to the necessary resources in the customer-service system, such as specific databases and storage-area networks.

Blades get around these chores by letting administrators configure an enclosure's slots to specific

tasks. For instance, an administrator can set things up so that the hard drive of any blade inserted into an enclosure's No. 3 slot (enclosures normally have 10 to 20 slots) will automatically get a copy of a Linux-based operating system along with a Web-server program. If a blade is inserted in slot No. 4, on the other hand, its hard drive would get a copy of Windows NT along with a database-server program.

This multitude of operations involved in installing and configuring a server is known as provisioning. For conventional servers, many operations have to be done explicitly, by physically detaching and reattaching cables and then manually reinstalling software and reconfiguring networking information. True, a number of programs are available to facilitate provisioning with conventional servers, but those are not always standard features; blades, on the other hand, encapsulate these management tools by design. In some cases, installing a new blade can be as easy as inserting a VCR tape into a player.

The real breakthrough of blades, therefore, is the ease with which they can be installed, configured, and, especially, reconfigured. The networks, storage systems, and other resources a conventional server can access in a computer room depend mostly on its physical connections. What a blade server can access also depends on its connections, of course, but those in turn depend on the enclosure's connections—and they can be changed with mere mouse clicks, rather than by physically rearranging cables.

Suppose, for example, that an enclosure is connected to 10 different networks and 10 different storage devices. Would a blade in this enclosure also be connected to all those networks and devices? If you wish so, yes. The management software would let you set this blade's connections without ever touching a cable. You could associate the blade with, say, the human resources network and a neighboring blade with a backup storage device. The enclosure's networking and switching components would act as traffic cops that appropriately direct data from the blades to the outer world and vice versa.

Another useful feature manufacturers are adding to blades is the ability to split a single physical blade server into many smaller, virtual ones. It lets an administrator create logical shells that each run a different operating system and its respective applications simultaneously on a single server. It's an old trick—called virtualization—that can be done with conventional servers, too. But with blades, the flexibility inherent in reconfiguring their software and connections makes virtualization tools much more powerful, letting companies more easily "break up" their computing resources and assign the pieces on demand to different users.

The upshot is that the more blades proliferate, the closer administrators come to eliminating a longstanding problem in computer rooms everywhere: that machines are working far below their full computing capacity. Industry data for Intel-based servers show an average utilization rate of only 20 percent. The cause is a lack of flexibility to change servers' configurations more dynamically, giving more computing power to whichever application needs it most at any given time; the main culprit is the one-application/one-server model. With more easily configurable servers like blades and with better provisioning and virtualization tools, administrators can tweak their machines and—finally—approach the long-sought 100 percent utilization rate.

SO IF BLADES ARE TECHNICALLY SUPERIOR in many ways, why have they captured barely 3 percent of the market after almost five years? The short answer is that, on average, they still cost more than conventional servers. But that's not to say that blades have to hit a certain price to compete successfully with conventional servers. The blade format costs more because it provides more value than the conventional format. So buyers need to take into account such things as how many machines they have, how much they spend on maintenance, and what levels of reliability

they need, among other factors, to find out how they can benefit from blades.

In fact, few companies can afford a wholesale migration to blades, which generally cost a few hundred dollars more per server than the similarly configured rack-mounted units. Instead, many are opting for a gradual changeover as their older servers become obsolete, restricting blades' adoption to certain applications, such as critical Web-based programs that need to be always available to users.

That slow-deployment scenario has been aggravated by incompatibility problems. Because blades are a relatively new technology, standards both for hardware and software have yet to be widely adopted. For example, enclosures from IBM can't be used to house blades from Sun. And the blade management software from HP can do little to assist an administrator in managing blades from Fujitsu. In fact, many customers fear that they may be tied to a single vendor or, even worse, that they may eventually get stuck with products from different manufacturers that don't talk to each other.

Then there are a couple of unexpected—and rather unwelcome—surprises that have frustrated many customers. Some of them eagerly filled their racks with blades only to discover that their air-conditioning systems weren't powerful enough to cool such dense blocks of computing machinery, which overheated and shut down. As it turns out, even though the heat dissipation of an individual blade is smaller than that of a normal server, filling a whole rack with blades causes the total number of servers to increase so dramatically that it may end up creating serious hot spots in the computer room.

Worse still, other customers reported they found out they didn't have enough electric power to run whole racks of blades (in fact, the whole idea of being constrained by power, rather than space, is something new to many administrators). Most U.S. office buildings can provide 10 000 watts of power for a single rack of computer equipment. With normal servers containing one Intel processor each, there is no problem: a rack will need just about that much—10 000 W—of power. Now fill that same rack with blades, and you may need 50 percent more power, or 15 000 W. That's an amount of power the typical computer room may not have for a single rack. In fact, some customers, after turning their blades on, got memos from their buildings' administrators saying they were causing brownouts on the premises!

Careful planning can help avert these problems, but sometimes a computer room simply doesn't have the cooling and electrical capacities needed to handle all the blades that a company would like to install. The bottom line is that, although vendors continue to claim that a rack can house hundreds of blades, the reality is that power issues are limiting many buyers to filling their racks just partially. This situation has slowed the adoption of blades, to the dismay of vendors and customers alike.

THE DEVELOPMENT OF BLADE TECHNOLOGY, like the early development of the PC, demands a high degree of experimentation and innovative approaches. To date, some of the most promising ideas have been brought to market by a group of smaller vendors. It was one of them, in fact, that pioneered the very idea of a super-slim server. RLX Technologies, in Woodlands, Texas, started developing its first blade in late 1999, and in 2001, it announced its first products. RLX didn't remain alone for long, as a number of companies soon got on the blades bandwagon. By the end of 2002, HP, IBM, and others released their blade offerings.

Since then, things haven't been easy for the small vendors. Two years ago, there were about 10 of them, but as competition got stiffer, that number has dropped to less than a handful. Even RLX dropped out of the blades market, announcing early this year it would focus solely on server management software. The surviving ones—such as Egenera and Verari—have been able to

keep up with their much bigger rivals because blade hardware is based mostly on commodity components such as Intel chips. So the competition is based not so much on which components they use as on how cleverly they can make them work together.

Egenera, for example, offers provisioning features flexible enough to let administrators take a blade's processor and put it to work as if it were part of another blade that needs more processing power. The same can be done with other components, such as a storage connection or a network card. Such changes can be programmed to occur automatically, according to policies established in advance that reallocate resources as needed. Egenera had sales of \$41.2 million in its last fiscal year and says it has secured more than \$100 million in venture funding over the past five years.

At the same time, all vendors, big and small, are competing to be the first to solve blades' electrical and heating issues. Some are investigating new cooling strategies, looking beyond the traditional bottom-to-top, hot-aisle/cool-aisle model commonly used in computer rooms today. Verari stands out for its unique cooling system, which directs cold air at the blades from different directions. IBM, Sun, and others have even contemplated a return to the old days of cooling computers with circulating liquids. So far no ideal solution has emerged, and therefore the race is on. The vendor that provides the most agile blades while resolving the electrical and cooling problems stands to stake out an enviable position in this emergent and promising market.

Meanwhile, at Veritas's computer room in Houston, things remain calm for the nodelers. Peter Harris, manager of the advanced systems group, says that the company now has more than 5000 blades in operation and that installing and configuring one of these blades takes minutes—sometimes just seconds—rather than hours. Needless to say, life is better for the nodelers. They have more time to work on other projects rather than servicing the same nodes over and over again. In some sense, they aren't nodelers anymore. Shall we call them "bladers"?

About the Author

Jane Wright is a research director at Gartner Inc., in Stamford, Conn., where she focuses on servers and related hardware and software. Before joining Gartner, she worked as an analyst at Digital Equipment Corp. and as a systems engineer at IBM Corp.

To Probe Further

For more on blade-based computer clusters, see <http://public.lanl.gov/feng/Bladed-Beowulf.pdf>.

Information about blade standards is available at <http://www.bladesystems.org>.

For more on the hardware market, see "Hardware Hangover," *IEEE Spectrum*, January 2003.

IMAGE: BRYAN CHRISTIE