

waters

OCTOBER 2003

Scale Your Servers

Wall Street is cutting the clutter in its back office with sleek but powerful blade servers and collaborative grid software By Dianne See Morrison

Blade servers have emerged as the weapon of choice for investment banks as they compete for the elusive edge in trading technology. Among the largest firms, JPMorgan Chase, Credit Suisse First Boston, Deutsche Bank and Lehman Brothers are employing blade servers to run front-office applications. Now, fund managers and smaller financial services companies are starting to get interested in what blade servers can do.

What's the big deal? Quite simply: more processing power for less dough—in a smaller package. Blade servers are essentially servers on a blade, or a computer circuit board, that allow users to consolidate all their other servers onto one chassis, sharing power, cooling and storage. "Users can manage 10 to 25 servers from one console as a single platform," explains John Humphreys, a senior analyst of the server group at IDC, a market research firm based in Framingham, Mass.

Once a firm invests in the chassis, it can easily scale those servers by adding more as needed. "It's as simple as buying a new blade and plugging it in," says Humphreys.

Finally, blades are cheaper than standalone servers and they deliver equal, and in some cases, faster processing power. The cost savings of using 10 blades compared to 10 individual servers can range from 15 to 20 percent, according to Humphreys. Those savings are compelling. "I don't know a major Wall Street firm that doesn't at least have a large scale pilot if not its first dozen to hundreds of systems in production," says Evan Bauer, a financial technology consultant and former chief technology officer at Credit Suisse First Boston (CSFB) in New York. Bauer recently installed blade servers at CSFB.

Server, Heal Thyself

It has been called autonomic computing, organic IT, virtual computing and adaptive computing. Whatever the label, the end vision reads much like an online personals ad: “Desperately seeking a computer system. Must be open, intelligent, unafraid of change, have a sense of self, capable of adapting well to new circumstances and able to handle challenges without having a breakdown.”

After years of under-utilization and over-provisioning of IT systems, infrastructure vendors are promising smarter servers that can self-configure, self-heal, self-optimize and self-protect. The goal, they say, is to take the complexity and cost out of the data center; reduce the number of IT staff; and free up banks to concentrate and respond more quickly to their strategic demands.

Autonomic computing is, of course, years away. How many years, vendors like IBM, Sun Microsystems and Hewlett-Packard won't say. Their representatives prefer to call it a “journey” but some recent infrastructure trends have moved up its debut. Grid computing, for example, allows banks to siphon off and redistribute unused processing power to whichever application needs it most at a given time. Blade servers turn out to be cheaper, as much of the physical infrastructure is housed in fewer boxes. Utility data centers, where resources such as servers, networking, storage, and applications are “virtualized,” have also begun to appear from companies like HP. The theory is that investment banks can buy only the amount of computing power they need. The theory is easy but the practice, as always, is trickier. Pricing models are difficult to work out and in some cases, it's still cheaper for a firm to buy its own equipment rather than to outsource.

Analysts report that frustrated IT staffs are already drowning in a sea of acronyms, competing concepts, and marketing hype that makes it difficult to know exactly how to proceed. Banks and their CIOs are understandably suspicious of losing control. Self-optimizing and self-regulating? What if a system administrator disagrees with the diagnosis the system makes? Although a CFO might like the sound of spending less on hardware, CIOs are leery about how systems will actually work with less hardware. Finally, if customers are to hand over some of these processes and outsource to IBM or HP, they must be convinced that giving up control would be worthwhile. “This is an even scarier thought than reducing hardware,” says Amy Wohl, president of Wohl Associates, a Narbeth, Penn.-based technology consulting firm.

The growth in blade servers on Wall Street is echoed in a recent IDC report on worldwide sales of servers. IDC predicts “dramatic growth” in the blade server market. From 2002 to 2003, they expect the blade server market to grow from a \$100 million market, to \$600 million by the end of 2003. By 2007, they expect sales to reach \$6 billion.

Blades Meet The Grid

Of the blade projects implemented on Wall Street, only a small percentage have been made public. JPMorgan Chase is thought to have made the most extensive use of blades, employing them to power its “Compute Backbone” (CBB) project, a grid-computing project used mainly to run risk management systems. Recently, however, they have also put their Portfolio Auto Trading (PAT) application on blades. CSFB, another early adopter, last year upgraded 20 RISC-powered servers with a blade server to run some front-office trading applications.

Although grid computing projects are sold with the idea in mind that banks won't have to purchase new servers—they can use their existing servers to run the daunting risk and portfolio modelling calculations—JPMorgan bought dozens of new Intel-based servers with the Linux operating system in order to move away from the Sun Solaris environment. “We wanted to adopt Linux for the Grid project,” says Ty Panagopoulos, vice president of investment technology for JPMorgan's CBB project.

Like many of the larger banks doing blade projects, CSFB and JPMorgan Chase chose products from Egenera, a three-year old blade server hardware and software provider based in Marlboro, Mass., started by Vern Brownell, CEO and chairman of the company. Brownell drew many of the ideas for his company from his experiences as a former chief technology officer at Goldman Sachs. Egenera's main product is its BladeFrame system, which provides multiple processors on four blades that are housed in a 24 x 30 x 84 chassis. Brownell explains that the blade product is actually both a hardware and software offering and is sold as a unit that is connected to either storage area network (SAN) or network-attached storage devices in a bank's back office. Brownell calls it “the data center in a can.”

Indeed, according to analysts and several banks, Egenera is currently enjoying a spell as Wall Street's darling blade server vendor. Bauer recalls being impressed with the ease with which CSFB installed BladeFrame two years ago, citing the low cost to provision hardware into the blade frame and to administer it. “It is really purposely built for enterprise data centers. A lot of blades use low-speed workstation-style chipsets, but BladeFrame is based on Intel's highest performance server chipsets,” he says. The blades also deliver faster communication rates through a high-speed optical back plane. He adds that because

everything is situated on an internal switch fabric, CIOs don't have to buy for 24 nodes, or buy a Cisco Systems catalyst switch or a data storage switch, for example. "You literally plug the frame in instead of having to plug each blade into those environments," says Bauer.

JPMorgan Chase notes that the applications currently running on blades can be scaled horizontally to cover a vast area of the bank's enterprise. Mark Etherington, Global Head of Infrastructure and Architecture, Equities Division, at JPMorgan Chase says such applications "yell blades." Etherington also notes that while investment banks have been trying to milk more utilization out of their systems for years—such as harnessing PCs overnight by repurposing older Sun or IBM boxes and writing ad hoc software—with the advent of blades and the proper software, they are now in a situation where they "can engineer it properly and get a pretty darn good utilization out of the box."

Etherington also reports that with BladeFrame he was able to reduce 25 percent of the development environment—another cost savings—as he felt he could safely repurpose from a User Acceptance Testing (UAT) environment to a Business Continuity Planning (BCP) environment. "The way you can change the use of CPUs and the way the frame is attached to the SAN and NAS storage behind it, allows you to do that safely," says Etherington.



Switch Blades?

While Egenera remains ahead of the competition, it is charging Wall Street a premium. As Bauer notes, one of the blade manufacturer's few downsides was the plain fact that Egenera was more expensive on a per-chip basis than a standalone server.

"Egenera is not the low-cost solution, but on the other hand, in terms of the time to deploy, cost to deploy and cost to maintain, it was substantially to our advantage," recalls Bauer. In an earlier interview with *Waters*, Steve Yatko, CSFB's global chief technology officer for securities IT, said he expected CSFB's switch to blades to save the company "literally millions of dollars" over time. As for price performance of the RISC hardware products available at the time, CSFB saw a ten-fold improvement. Indeed, the return on investment came in well under six months as the new system also allowed the bank to do more business, Yatko says.

While Egenera is currently the blade server vendor of choice, it could be in for some heavy competition in the upcoming months. The heavy hitters, namely IBM and Sun Microsystems,

are looking to enter the blade and server consolidation space with their own offerings. That wasn't the case just 18 months ago. When CSFB was undertaking its upgrade path to the then-new blade servers, Bauer recalls the lack of competition for Egenera.

"At the time, nothing even came close," says Bauer. He notes that while Big Blue was the second-closest contender, its blade servers were still using a gigabit Ethernet as the back plane. This meant that all connections to the blades themselves ended up going through an Ethernet stack instead of through the more proprietary and speedy switchback plane that BladeFrame uses. IBM lacked the software management suite, and the pan manager across the area network manager that Egenera had.

Etherington notes, however, that Sun's Solaris Resource Manager and Zones on Solaris could provide good solutions to virtual blades or v-blades, while using Linux on IBM's Z series could help alleviate over-provisioning as well. V-blades allow IT managers to run multiple versions of Linux on a single blade. "Those

