

When More CPUs Are Better



Moving to parallel processing could be the right step for your IS department, but perform the necessary evaluation before you leap ahead.

As a member of a design or development team, when you sense that a new technology may be more appropriate than the mainframe or a single server, it can be intimidating to face an IS department and its advisors who are traditionalists. I found myself in such a predicament a few years ago in recommending parallel processing. Today, that would be easier. Buying multiprocessor systems is no longer for the early adopters alone. Gartner Group recently estimated that this \$1 billion market will grow 50 percent annually.

For the record, let's clarify usage regarding parallel processing and its more specific incarnations. *Parallel processing* is the umbrella term for a computing technique in which many operations are performed simultaneously, using many microprocessors. A "coarse-grain" system, often identified as *symmetric multiprocessing* (SMP), contains relatively few but powerful processors. A "fine-grain," or *massively parallel processing* (MPP), machine contains up to thousands of smaller ones.

Increasing name recognition and accelerating market share may open the door to considering an idea in your company, but they are not a sufficient reason to recommend the purchase of parallel processors. Applying more generic business justifications to your project is a better way to help determine whether you should take a

closer look at parallel processing. Here are some useful considerations. (For a full discussion of this technology applied to data access, see page 26.)

No Pat Answers

In price/performance ratio, parallelism far outstrips mainframes and single RISC or Intel servers for many applications. Actual numbers will vary depending upon the application.

It can be reassuring that early adopters have gone before you. Multiprocessor systems have been used successfully for online transaction processing, data warehousing, decision support, simulation and financial modeling. Vendors can provide sample success stories to help with your cost/benefit analysis.

Database vendors finally have arrived on these platforms. Informix took three years to rewrite its core database architecture to provide clustered SMP and another year to accommodate MPP functionality. Sybase, Oracle and IBM now also have integrated parallel processing capability into their database architectures.

Scalability is a major benefit. SMP is available from the low end, with two- and four-way machines, up to the high end, where systems with eight or more processors are outperforming mainframes for many applications. When anticipated correctly, upgrades are relatively straightforward.

SMP *cluster systems* (multiple processors on multiple linked machines) can overcome some of the operating system limitations in handling SMP. Hewlett-Packard, for example, is touting the advantage of parallel server technology over more traditional MPP architectures.

There are multiple sources. All major hardware vendors have a multiprocessor product line, so you can shop around and count on continued competition, variety and innovation among major systems.

Choosing SMP, clustering or MPP should depend upon your specific applications and goals. Solutions change almost as fast as definitions. It used to be that data warehousing, which must store large volumes of data and provide reasonable response times on queries for a large number of users, was an application for MPP architectures. When detailed queries on massive amounts of operational data were required in realtime—for instance, if you were Mattel and needed to know how many Busy Gal Barbies sold in Boston in the past hour—MPP was the answer. Now, SMP solutions and clusters should be considered, too.

SMP is sure to become more popular as more tools and applications are developed for these environments. And clusters do not necessarily require new hardware, while MPP systems usually do.

Stop, Look and Listen

It is best to have an experienced, vendor-neutral guide when going through this the first time. A few years ago, I was helping an internal IS team responsible for recommending a server for a huge financial modeling, or number-crunching, application. The IS department had initially recommended adopting Windows NT on an Intel processor.

Once we modeled the application conceptually and simulated it on the users' mainframe, we found single-processor server architectures to be inappropriate.

By Sally Atkins

The application took 24 hours to run! There was not enough mainframe capacity in this \$40 billion corporation to handle such a new application. We ended up testing with an eight-way Sun Sparcserver. The simulations, not possible before the new financial modeling had been done, ran in a quarter of the time they would have taken on the mainframe. Seeing the proof of the simulations left no business choice but to adopt the new parallel technology.

We took several lessons from this experience. One was to slow down the users and their eager programmers long

enough to do simulations before entering the recommendation stage. The full model allowed us to understand the nature of the processes and data we were dealing with, and we gained the will to go ahead and build applications that were not feasible with single-processor architectures.

We also learned to use the knowledge bases of the hardware, operating system and database vendors in the presales process, just as we'd consult with any other member of the team. Some vendor alliances are stronger than others. We found that engineering partnerships are a lot deeper and more valuable to users

than mere marketing alliances. These show up in presales modeling efforts and can help you choose vendor partners that are capable of working together.

Moving to parallel processing, especially in a distributed client/server environment, is a big step. It may not be appropriate for every need, but today there is no good reason not to consider it as an alternative. ■

Sally Atkins is president of *IST Consulting*, an affiliate of *NetSource, Inc.*, based in Boston. She can be reached at Sally@kins.com.

