Client-server demand and legacy systems

Shane Greenstein

University of Illinois

1206 S. Sixth Street

Champaign, IL 61820

sgreenstein@commerce.cba.uiuc.edu

The most important computer users in the US today are business organizations. Virtually every business enterprise buys something from the computing market. Banking and finance are the most computing-intensive sectors, though many sectors of manufacturing, wholesaling and retailing, transportation, and other services are now not far behind in computer intensity. Information systems underlie most efficiency gains, new products and services, and other startling economic improvements in these organizations.

For my money, it is important to understand the demand for all computing platforms, mainframes, minis, and micros. While mainframes and minicomputers have declined over the last decade, they are not dead. To be sure, personal computers are likely to be the most important sector in the future, but the early systems are not disappearing as fast as everyone thought.

This column considers the competition between old and new. Think of "old" as a generic mainframe, like an IBM 4300 or 3900 or any centrally managed large system, and "new" as client-servers. Though no two engineers agree on this definition, think of client-servers as networks of microprocessor-based workstations or personal computers hooked up to more than a common printer. If the server is a centrally managed mainframe routing e-mail, it does not count.

Some years ago, Stanford economist Tim Bresnahan and I began studying the disappearance of users of large systems. (For copies, please contact either CEPR, the Center for Economic Policy Research, Stanford, Calif., (415) 725-1874, or NBER, the National Bureau of Economic Research, Cambridge, Mass., (617) 868-3900.) We obtained surveys of tens of thousands of computing sites from Computer Intelligence Infocorp, an information provider in this industry. We pored over thousands of records of purchases and buyer behavior. We wanted to identify who got rid of their systems. Why some users and not others? What did this say about the nature of competition between old and new?

Our study had some general lessons and some lessons specific to this case. The general lessons are interesting because they will likely arise again in other competitive episodes in the computing market. The specific lessons are interesting because they make good stories (and many people's livelihoods ride on the outcome).

The overriding general observation is this: When an organization makes a major computing-equipment acquisition, it puts many of its routines at risk, potentially revisiting an enterprise's core strategy for structuring operations. Such issues do not resolve themselves easily if management, staff, and the user base do not share a common view of computing priorities.

The details are complex because the technology choices are complex. Effective use of modern computing equipment involves communication equipment and networking technologies, old and new software, and large doses of human intervention. The networks often involve private and public communication lines, private and public switches. On the human side, effective use of computing technology means countless hours of training, learning, and maintenance by staff, and frequent restructuring of important and minor routines by programmers. On the software side, this means producing application-specific and organization-specific programming to refine and retrofit old software.

The management side is also complex. Changing equipment potentially alters an organization's operations, its staffing, and its final product. It has consequences for many poten-

0272-1732/95/$04.00 © 1995 IEEE

tial expenses far beyond the actual purchase of the equipment. Though generalizations are incautious, for most computing, the yearly expenses associated with managing staff, as well as employing software programmers and support help, usually exceed the hardware expenses. Finally, different management patterns, which are associated with different applications and outputs, tend to incur very different magnitudes of adjustment costs. For example, there are large differences in the costs of changing systems for engineering-based applications, batch-oriented back-office functions, or essential real-time applications.

Now for some lessons particular to this case. Here is what we learned in our surveys: In the early 1990s most large computing users stayed with their large systems, surprisingly unwilling to move to a new client-server platform. This seems to make little sense at first blush because the new systems ostensibly had higher technical benefits across a wide set of uses. The key to understanding this puzzle was to have the right view of both the costs and benefits of client-server technology. Many users knew about the potential benefits of client-servers, but abandoning large systems came at larger internal adjustment costs than technologists and engineers anticipated or cared to admit in the trade press.

"Internal" is the important point. No market exists for costs incurred. Resolving problems instead depended on organizational incentives for inducing employees to bear those (often hidden) expenses. We found that the only types of users who initially moved to client-servers were engineers and academics (including most of our friends and colleagues). They tend to have decentralized organizations, and lower, more dispersed internal adjustment costs, but make up only a small fraction of total demand. Many big-system users either did not change or took a more cautious attitude to investing in their computing stock. This helps explain why the aggregate mainframe market demand shifted downward, but much less than predicted by those in the majority of the trade press who predicted (and are still predicting) a revolution in computing technology.

We further investigated whether a user's ties to a manufacturer made a user more resistant to new technology. While stories on this issue abound, there is little systematic statistical investigation of it. One can easily see why this is an important question, since a single firm provides roughly two thirds (by the number of boxes) of all large general-purpose systems in the US. To put one version of this hypothesis bluntly, IBM has few proprietary rights in the most widely used client-server solutions. So, is IBM holding up the diffusion of the technology to customers with whom they have the closest ties?

This issue, even in its blunt form, is difficult to pose in practice because the market structure for complete computer systems is so complicated. Hardware vendors provide software services and maintenance; some even provide customized services. A large third-party software industry for large systems also exists; some of it is available on multiple platforms and some is not. Many users program their own system tools, but buy packaged application software from one vendor or another, and on, and on. The main point here is that the real world was simply too complicated for the most sweeping conspiracy to possibly be true; IBM could not be managing so many different complex situations. To construct the hypothesis at a reasonable level, we had to work hard and understand at a significant level of institutional detail how this market operated.

Anyway, the overriding finding is that no matter how you cut it, IBM took a big hit to sales from traditional customers. Now that is not news, I admit. More to the point, ties to vendors do not matter as much as internal adjustment costs. The one possible exception to this finding occurs among large-system users who buy IBM-proprietary communication technology, often hardware and software products that are complementary to large databases and a large user base. These buyers tended to resist abandoning mainframes for client-servers. However, it was not obvious that the tie to a vendor was as essential as the scale of computing activity, which made for enormous adjustment costs. Also, the new client-server technology had difficulty satisfying these particular types of users' needs.

In general, what did we see here that we are likely to see again? Computer users see constant competition between old and new technology because the appearance of a new technology offering lower costs or superior capability rarely leads to instant replacement of the old technology. After all, some users may be reluctant to retire computers that continue to offer a flow of useful services, even if technical change apparently depreciates the market value of those services. In addition, sellers of the old technology may find their competitive circumstances changed, but react quite naturally with new pricing or new technology strategies. Buyers may also delay purchasing the new technology until anticipated price/performance improvements appear.

The larger point is that the pace of new technology adoption and retirement of the old depends on all the factors that shape the competition between old and new. While it may sound didactic for an economist to emphasize the importance of market processes, my experience has been that these factors arise in many different facets of demand research. Most of these processes are out of a buyer's and user's control, but they influence the costs and benefits of different purchase decisions—even the competition between mainframes and client-servers and, in all likelihood, future events.