

Server blades: Where are we now and where are we heading?

By Lance A. Leventhal

Blades became a billion-dollar market in 2004 (according to IDC), with the potential to reach \$9 billion by 2008. As Senator Everett Dirksen once put it, "A billion here, a billion there, and pretty soon you're talking about real money." Blades are also a major topic in most surveys of IT, telecom, and data center managers. We will soon reach the time when all major installations will have blades. Of course, there are many problems yet to be solved. We need better tools to manage, boot, configure, assign, monitor, and track blades. Automation is absolutely essential, since no one will want to handle them individually or examine their status one-by-one. Most installations are small today (typically under 100 blades), but numbers will grow. We can expect that thousands and even millions of blades will become commonplace, judging from the rapid growth in demand and numbers of other servers. Scalability will quickly become a major issue, since solutions that work well for a few blades may not handle large numbers.

Everything must be scalable, including:

- Management tools
- Interfaces
- Storage
- Network connections
- Repair and upgrading methods
- Power
- Cooling

We must be ready to handle the numbers that will come faster than we can imagine. Research that has dealt with large numbers of processing elements or network nodes may well apply here as well.

Software will eventually be the major obstacle. Obviously, users want it to be inexpensive, easy to use, and simple to integrate with other software. Oh, and expandable to handle increased numbers

and types of elements, capable of managing storage and peripherals, and accessible over a network. Did I mention that the vendor should provide installation and training – and be large enough and competitive enough to be around for a while? The major OEMs must form alliances that ensure software will run on their blades, and the smaller vendors must join forces with them. System managers need a unified approach rather than hundreds of separate screens with different features and functionality.

Of course, users also want interchangeable standards. They want to buy whatever is best (or cheapest), combine old equipment with new, and not worry about a vendor dropping a product line, being acquired or merged, or deciding to focus on other areas. So far, the major OEMs (Dell, HP, IBM, Sun) have been loath to comply, fearing competition from copy-cat manufacturers with low labor costs.

Data centers already face problems in handling blades. Crowding more units into smaller spaces leads to serious cooling and power distribution issues. Older centers may have to be retrofitted with extra equipment, and new centers must have much greater capabilities. Designers will have to be aware of standards such as Network Equipment Building Standards (NEBS), Telecommunications Industry Association (TIA), and American Society of Heating, Refrigerating, and Air Conditioning Engineers (ASHRAE).

There is no question that blades are here to stay. Lower costs, smaller sizes, greater modularity, and simpler maintenance and expansion all make the concept an obvious choice. But we will obviously need to solve challenges in software, interconnect, management, and facility planning to get the most advantage from them.

Lance A. Leventhal is the Technical Editor of the Blade Systems Alliance and Program Chairperson of the Server Blade Summit.

For further information, contact Lance at:

Blade Systems Alliance
 P.O. Box 1258
 Rancho Santa Fe, CA 92067
 Tel: 858-756-3327
 Fax: 858-756-2656
 E-mail: Lance@bladesystems.org
 Website: www.bladesystems.org