

Server Blades:

An Emerging Server Architecture

By Mike Roberts

The emerging server blade architecture promises to alleviate many of the data center issues brought about by the shift to scale out architectures featuring standards-based, rack-dense servers. Server blades have the potential to increase server density, improve manageability, lower power consumption, and enhance deployment and serviceability, all resulting in lower total cost of ownership. Although the initial incarnation of blades is well suited for expansion into the front end of the IT infrastructure, the architecture is likely to evolve and play a significant role throughout the data center.

Data center administrators and IT personnel have struggled with the infusion of rack-dense servers into their environments as they continue to shift from monolithic, proprietary architectures to more cost-effective, standards-based servers. Although the growth in rack-dense servers has been strong, vendors have not focused on issues of deployment, serviceability, cost, and management as much as improving density. Server blades, a new architecture for the next generation of servers, will alleviate many of the existing issues and will begin to change how vendors design standards-based servers.

What is a server blade?

Although vendors will likely create numerous incarnations of the server blade, it is basically a server on a card—a single motherboard that contains a full computer system, including processor(s), memory, network connections, and associated electronics (see Figure 1). Most server blades initially will include some level of local disk storage with provisions to access external storage, as discussed later in this article.

Each blade runs its own instance of the operating system; consequently, administrators can designate individual blades for different applications or end users and can insert or remove (hot plug)

them without affecting the operation of other blades. Multiple blades plug into the midplane of a chassis, or enclosure, that provides shared, redundant infrastructure components such as power, cooling, networking, and cabling.

Why server blades?

Essentially, blades allow system vendors to provide customers with solutions that improve efficiency in the data center, which ultimately reduces the total cost of ownership. Improving data center efficiency hinges on the following points.

Reduce hardware costs

Server blades are relatively inexpensive because each blade does not have a separate chassis and infrastructure like a traditional server. By leveraging power, cooling, management hardware, and cables over multiple systems, the per-server cost can be dramatically reduced, as shown in Figure 2.

Vendors can now integrate components that customers would purchase separately—such as remote management cards, Ethernet switches, and even KVM (keyboard, video, and mouse) switches—into the chassis at a very low cost. Leveraging high-quality components over multiple servers also makes sense from




Figure 1. A shift in server design

a reliability standpoint, since the system has fewer components that could fail. Because of the initial investment cost associated with the chassis, blades are better suited for concentrated data center environments rather than distributed sites with small numbers of servers.

Simplify deployment and service efforts

Deploying multiple servers can be a time-consuming and resource-intensive process. An administrator installs each server in a rack, cables it to power and to the network, and provisions it with software. The cables, in particular, create one of the primary deployment headaches in high-density environments.

Using blades, an administrator installs and cables the chassis only; the individual server blades need no cabling. Adding new compute resources is as easy as plugging in a new blade, just like adding a hard drive today. The chassis minimizes the amount of cabling because redundant power supplies are leveraged across multiple blades. In addition, built-in switches consolidate Ethernet data and management networks, and even KVM connections (see Figure 3).




Figure 2. Blade economics

Automated software provisioning tools allow administrators to quickly and easily deploy software to one or many blades in a single operation over the network. After provisioning the blades, an administrator can manage them completely using remote-management tools. Customers should expect vendors to supply fully featured tools for remote software provisioning, hardware configuration, and systems management. These tools will allow customers to manage blades and traditional integrated servers from a single framework, lowering resource costs.

Servicing a failed blade involves removing the failed blade and plugging in a replacement while the failed blade waits for repair. An administrator also can easily implement a global hot-spare server, which can be remotely provisioned to take the attributes and tasks from a blade that requires service.

Optimize use of space in data centers

Because data center real estate is expensive, the market has witnessed a rapid movement toward dense rack-optimized servers during the last 24 months. Server blades can increase server density by 100 to 800 percent, depending on the design and vendor, when compared to today's 1U rack-optimized system (see Figure 4). In addition, since each blade is an individual server, the chassis can be used to consolidate multiple applications in a single box. This consolidation further assists customers in centralizing their compute resources within the data center.

Increase resource utilization

Data center managers must size the infrastructure of each client (whether an internal organization or external customer in the case of xSPs) for peak demand. Often, a significant portion of the hardware is underutilized while still generating heat and consuming capital, administrative resources, power, and space.

By migrating to a server blade architecture and using related improvements in workload management software, managers

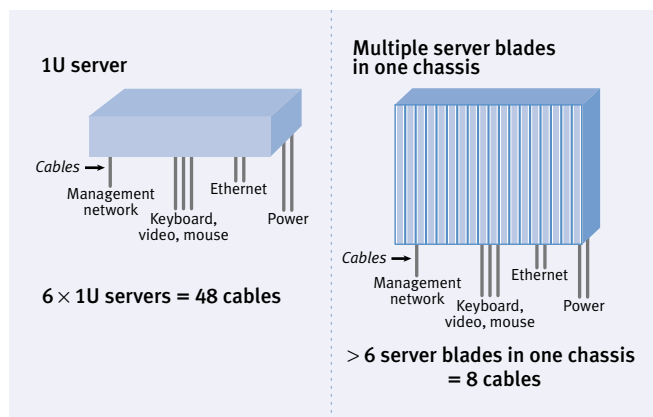


Figure 3. Cable consolidation by blade chassis

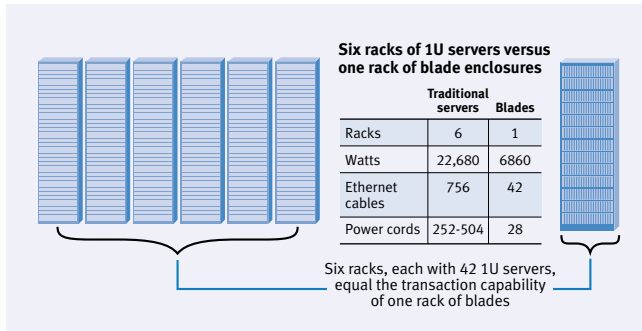


Figure 4. Density of traditional servers versus blades

can size the data center infrastructure not on the peak demand of an individual client or group, but on the demand of the data center as a whole. Sizing is largely a function of management software, but blades can enhance software sizing capability with their low cost and ability to simplify the process of adding compute resources.

Reduce power consumption

To reduce the power consumption of blades, most vendors will use low-power processors in some of their blades. Even blades that do not use low-power processors will use less power than “integrated” servers, because blades have fewer power-consuming components. In addition, vendors can afford to use higher efficiency power supplies because they are leveraged over multiple servers. Blades that consume less power also produce less heat, which reduces the amount of electricity required by cooling systems.

What are the challenges?

The miniaturization and increased density of servers present challenges for storage, server management, standardization, and server-networking collision points.

Limitations for storage expansion

One concern about server blades is their ability to handle storage expansion. Although the first server blades on the market will have onboard disks, they will generally be limited in number. This limitation is fine for applications in the first and second tiers of the IT infrastructure, which are not generally storage intensive. These applications rely on back-end servers to connect to the storage infrastructure.

Storage cannot easily be internally expanded because most blade architectures prohibit traditional peripheral component interconnect (PCI) expansion slots. In addition, blades may not easily accommodate multiple external interface electronics and connectors, which forces vendors to choose between Ethernet and Fibre Channel for connectivity to external storage.

Ethernet will likely be the storage connection of choice for the first generation of high-density blades, because blades must have Ethernet connections to access a LAN (local area network) or WAN (wide area network). Because of its cost-effectiveness and pervasiveness in the data center, Ethernet is an optimal storage connection for blades. Using Ethernet, server blades can connect to network-attached storage (NAS) and Internet Protocol storage (IPS). Emerging IPS technologies like iSCSI (Internet SCSI) enable access to block-level storage resources over Ethernet. Fibre Channel connectivity will more likely be relegated to less dense, higher performance, modular products.

Need for robust management tools

Data centers are already struggling with the proliferation of equipment and its impact on resources; server blades can only exacerbate this issue. For blades to be successful, vendors will need to deliver robust tools that support not only standard management interfaces but also remote access, automated software provisioning, or hardware configuration tools. Data center managers should consider these new tools when planning their infrastructure so they may take full advantage of this new architecture’s capabilities.

Multiple hardware standards

CompactPCI® (cPCI) is currently the only standard that defines blade form factor and connectors. Blades and enclosures built to this standard can operate in a standard cPCI enclosure even if they come from a variety of vendors. This compatibility provides some investment protection for the customer, although servicing and supporting blades from multiple vendors may incur additional costs.

However, cPCI is not currently suited to high-performance, high-density, low-cost environments because of its architecture. The cPCI blades on the market are targeted primarily toward telecommunications and industrial applications instead of traditional data center computing; thus, access to an “ecosystem” of third-party blades provides little benefit for most data center customers.

Server blade vendors that do not use cPCI will build their solutions on standards-based chipsets, controllers, and processors, as well as standard protocols like Ethernet and Simple Network Management Protocol/Common Information Model (SNMP/CIM). However, the connectors and signal routing will differ depending on each vendor’s implementation. Although customers will not be able to mix blades from multiple vendors in an enclosure, they will benefit from this implementation because it allows vendors to design products that offer better features at a lower cost, while occupying less rack space than products built on standards such as cPCI. As a result, in systems that do not use cPCI blades, server blades still will interact with each other and with “integrated”

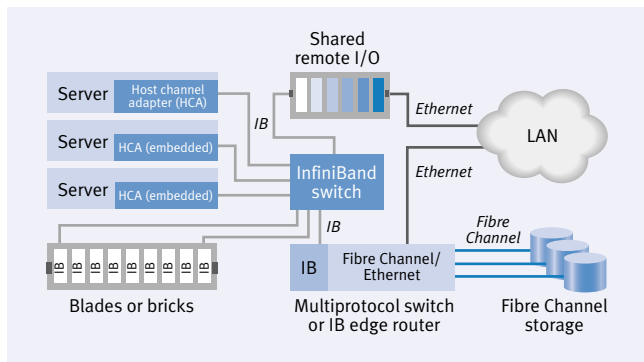


Figure 5. Sample InfiniBand fabric topology

servers in a standard manner, run standard operating systems, and be managed using standard management tools.

Network-server collision point

Because Ethernet cables emerging from each blade would be unmanageable in a high-density environment, vendors are likely to integrate Ethernet switches into the enclosure. If an Ethernet switch is used as a cable consolidation point, the only Ethernet cables leaving the enclosure are the uplinks from the switches. Although this consolidation is elegant from a cabling perspective, it will cause administrators to change the way they think about server deployment. Network administrators and server administrators will need to work closely as they deploy and manage their server-network combination. Vendors also need to provide tools for managing these Ethernet switches.

What lies in the future for blades?

A new I/O architecture, InfiniBand™ (IB), may drive the modular server landscape because it provides a universal interconnect fabric that is ideal in situations where space is limited and multiple fabrics such as Ethernet, Fibre Channel/SCSI, IPC (inter-process communication) are not possible (see Figure 5). IB can offer the following benefits for server blades:

- ▶▶ IB provides a high-performance, low-latency, blade-to-blade interconnect for high-performance computing and database clustering.
- ▶▶ IB provides connectivity to external storage and Ethernet networks through IB-to-Fibre Channel or IB-to-Ethernet gateways or “edge routers.” This architecture lets vendors use the capabilities of IB but does not force customers to re-cable the data center for IB.
- ▶▶ IB allows multiple blades or multiple enclosures to share legacy PCI peripherals by providing connectivity to external PCI expansion modules. Not only does this approach save

space, but it also allows customers to leverage devices over multiple servers.

- ▶▶ Native IB storage, not yet available, will allow IB-enabled servers to connect directly into IB storage without changing protocols.

IB interconnects and new management tools can provide the foundation for true “utility computing” models, where data center administrators have a pool of compute and storage resources that they can dynamically allocate as required. This technology fosters a dynamic and modular data center, whether it belongs to a service provider or a brick-and-click enterprise.

A new IT infrastructure

Early versions of server blades will be primarily high-density, low-power devices with relatively low performance. This type of blade is suited for first-tier applications such as static Web servers, security, network services, and streaming media because the applications can be easily and inexpensively load balanced. The performance of an application depends on the aggregate performance of the servers rather than the performance of an individual server.

Higher performance, less dense blade designs will help drive blade usage into more mainstream applications in the corporate data center. These designs can offer the individual performance characteristics and features available in today’s rack-dense servers along with the cost, deployment, serviceability, and density benefits of server blades. The blades will be well suited to high-performance Web servers, dedicated application servers, server-based or thin-client computing, and high-performance computing (HPC) clusters.

The introduction of server blades and associated technology like IB will usher in a new IT infrastructure. IT managers should start planning now for server blade installations by evaluating IP-based storage solutions, remote software provisioning and management solutions, scale-out architectures, and load-balancing technologies. 🌐

Mike Roberts (mike_j_roberts@dell.com) is a product manager at Dell. He is responsible for the Dell server blade product line. Mike has a B.A. from Marquette University and an MBA from the University of Texas at Austin.

FOR MORE INFORMATION

For more information, please visit <http://www.dell.com>