IDC
Analyze the Future

WHITE PAPER

The Third Wave of Evolution in Server Blades

Sponsored by: Sun Microsystems

John Humphreys Jed Scaramella

August 2006

EXECUTIVE SUMMARY

This IDC white paper examines an emerging segment of the server market aimed at bringing together the benefits of the blade architecture without compromising the higher-end features traditionally found in rackmount servers. Historically, customer adoption in the blade market has been relegated to the Web and application tiers of the datacenter as early products required customers to compromise performance in exchange for density. This situation is changing. With the Sun Blade 8000 Modular System, Sun Microsystems set out to break this relationship and provide users with a product that combines the performance of a rackmount system with the serviceability, scalability, modularity, and ease of deployment associated with blade systems.

With the Sun Blade 8000, Sun is committed to providing blade products that focus on modularity and serviceability in its x64 line for the most demanding applications in customer datacenters. Sun is applying this design approach across server form factors and product lines to build the best possible compute, I/O, power, storage, cooling, and system management modules. The approach is woven into the design of Sun's x64 enterprise servers and, in particular, the Sun blade server product line.

By focusing at the high end of the blade space, Sun can break down barriers that have thus far inhibited blade adoption and leverage its long expertise in designing and manufacturing high-end systems that run the most mission-critical parts of an organization's IT infrastructure. The Sun Blade 8000 helps to take blades and x64 systems into areas of the IT environment historically reserved for rackmount systems.

This paper details customer trends in the blade server market, the Sun Blade 8000 Modular System, as well as the principles behind Sun's Modular System design program.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

INTRODUCTION

Server vendors have long provided a variety of choices for IT organizations, and segments of the market are increasingly juxtaposed. For example, since 2002, vendors of blade servers and rack-optimized servers have offered alternative solutions for flexible, low-cost IT infrastructure, particularly in the application and Web access tiers. For more demanding mission-critical workloads, IT consumers have been limited by a lack of suitable blade options and have turned to larger SMP servers with greater capacities and capabilities. These circumstances have led to today's datacenter design points:

- To meet demand for increased capacity for smaller application and Web services solutions, IT organizations prefer to scale out an infrastructure of blade servers or rack-optimized servers in small increments. Applications typically run on a single monolithic server.

- For larger and more critical applications, IT organizations favor scaling up to midrange and larger symmetrical multiprocessing (SMP) servers. Larger and more critical applications typically run on a single server engineered to a higher degree of reliability, availability, and serviceability (RAS) than a blade or rack-optimized server.

This IDC white paper examines an emerging segment of the server market aimed at bringing together the benefits of the blade architecture and a systems engineering approach to meet the scale-up preference for enterprise-class computing. The strategy, as named and implemented by Sun, is *modular computing*.

Modular Computing Design at Sun

Sun defines modular computing as a three-step design process. The first step is decoupling a server system into independent modular components. Sun's six modules are *compute* (i.e., processor technology), *input/output* (I/O), *power* and *cooling* (i.e., datacenter facilities), *storage*, and *system management*. The second step in the modular computing approach is developing best-of-breed solutions for each module. The criteria for judging best of breed are enterprise-class requirements for capacity, I/O and connectivity, and RAS. The third step is combining modules to create a server well suited to enterprise customer needs and workloads.

Although a server system primarily comprises hardware modules, the role of software in a modular design remains critical. In the end, enterprise IT organizations are solution oriented, service driven, and operating system agnostic. Best-of-breed enterprise servers must offer a range of choices of operating environments, software tools, database support, and line-of-business software solutions. While the enterprise IT organization is deploying technology infrastructure, from the enterprise business perspective, the IT organization is supporting a business function. Only through software can a server really be declared a system.

Modular Computing and Server Blades

IDC believes that Sun's modular approach provides a strong framework for interpreting and forecasting evolution in the growing server market. Evolving customer requirements are affecting all six modules listed earlier. Perhaps the best place to start is with processor technology (i.e., the compute module) and then follow with the other modules:

- With the advent of multicore processors, a single server blade is approaching the capabilities previously associated with midrange SMP servers (those priced from $25,000 to $499,999, as defined by IDC). Especially for processors running the x86 instruction set, blade servers will support all major operating environments — Unix, Linux, and Windows. These operating environments, in turn, provide platforms for solution software. From an enterprise requirements perspective, the capacity of single, multisocket blade servers to support enterprise-class workloads is increasing. Moreover, the flexibility to choose among operating environments and software tools, a feature closely associated with enterprise servers, is emerging as a value-add feature of server blades.

- Improved processor throughput has an immediate effect on I/O and network connectivity requirements. Backplane bandwidth must be increased to avoid creating a choke point for total server system performance. In-chassis I/O aggregation is needed to interface with different interconnect fabrics (e.g., Gigabit Ethernet [GbE], Fibre Channel, and InfiniBand).

- Increasing processor capability and compute density also have an immediate effect on electrical power and cooling requirements. For blade servers, the chassis must be engineered to supply both power and cooling in sufficient quantity. Power and cooling components will need to be redundant and replaceable without causing server disruption to meet enterprise RAS requirements.

- On-board, compact, high-performance disk storage is commonplace on blade servers. System images and state information are replicated elsewhere to support rapid swapping of server blades, but onboard storage is crucial to performance. In a blade server architecture, large-scale and longer-term storage is typically provisioned across a network with storage arrays or storage area networks (SANs), which further underscores the importance of robust I/O infrastructure.

- Buyers of enterprise IT systems are keenly aware of the long-term costs of system maintenance. Hence, system management software will play a strategic role for customers who are consolidating workloads onto fewer servers to gain control over server sprawl and reduce costs. When managing the IT budget, enterprise IT professionals are fully aware of the importance of downstream operating expenses in addition to up-front capital expenses.

The modular computing approach is a reminder that system components must be balanced to provide optimal throughput. Additional processing capability benefits can easily be nullified by weaknesses in any other module. For example, inadequate cooling will lead to thermal decay and shorter, less reliable server life. Additionally, multiple server blades that demand individual system administration will raise the total cost of ownership.

Trends in the Blade Server Marketplace

Blade servers have completed two cycles of evolution and are poised to begin a third cycle. In the first cycle, IT organizations deployed blade servers in the edge tier to support workloads such as Web serving, data streaming, and file-and-print services. In the second cycle, triggered by more robust server blade platforms with greater processing capacity and improved RAS, IT organizations began to host application-tier workloads. Blade servers were candidate platforms for mail and other collaborative applications, for example.

IDC believes that the third cycle of evolution will be marked by the deployment of mission-critical applications to blade servers. When blade servers demonstrate enterprise-class RAS and midsize server capacity, then they will be deployed to support critical business processes. Database, decision support, and business analytic applications will be deployed with server blades. Enterprise resource planning, customer relationship management, and online transaction processing workloads will be hosted as well.

Modular Computing and Blade Server Benefits

Blade servers entered the market with a variety of promised benefits, particularly when compared with either rack-optimized servers or larger SMP servers. IDC believes, and IDC demand-side, customer-based research confirms, that blades and modular systems offer clear advantages over some architectures and are equal to others. With special attention to advantages and parity, we offer the following brief review of blade server benefits:

- **Cost of deployment.** Large numbers of servers can be deployed less expensively as blade servers than as rack-optimized servers. The costs of power and network cabling are significantly reduced because a blade chassis provides these shared services, whereas each rack-optimized server must be wired individually. The cost of deployment for both blade servers and monolithic servers is similar.

- **Scalability.** Blade servers are more easily scaled to match workload growth, especially for growth in the quantity of workloads (versus expansion of a single workload), than fixed-capacity monolithic servers. Blades are added to the system as needed. This approach results in a low-cost deployment because the blade server chassis is already in place, thus reducing overprovisioning. Monolithic servers are often overprovisioned to a greater degree, and outgrowing a midsize server leads to a larger incremental cost. Blade servers and rack-optimized servers offer similar levels of scalability.

- **Hardware serviceability.** Unlike both rack-optimized and monolithic servers, the blade architecture includes a field-replaceable server blade, which requires less skill and time to service or replace. The stateless blade does not require lengthy reconfiguration. While elements of a monolithic server can be swapped (e.g., a failed processor board, memory segment, or networking component), the process ordinarily is more involved and requires greater skill, resulting in a longer period of downtime. In terms of serviceability, blades have an advantage when compared with rack-optimized servers.

- **Power and cooling efficiency.** Blade servers require less power and cooling than rack-optimized servers. Consolidated power supplies are key. Whereas each rackmounted server requires two power supplies, blade servers benefit from their use of efficient shared power infrastructures. The blade chassis also provides an opportunity to improve cooling efficiency over heat dissipation strategies that must occur for each server in a rack. Blade servers and monolithic servers offer similar levels of power and cooling efficiency.

In summary, IDC views the blade server as an important new architecture that continues to mature and migrate toward enterprise-class workloads. Blade servers offer benefits and advantages that depend on an IT organization's goals and objectives. As technologies continue to improve across server architectures, sharp contrasts in benefits and capabilities will attenuate.

Modular Computing at Sun: The Galaxy Architecture

Sun's modular computing architecture is founded on the idea that it is not the form factor that determines the system architecture but rather the features and benefits that help determine the best form factor for the job. With this modular architecture, Sun has worked to combine the benefits of rack-optimized systems with those of blade servers and to enable customers to draw upon standards for networking their blade platforms. Sun's goal is to provide customers with choices for the most suitable architecture for their broad enterprise computing needs.

Sun's clean-sheet approach to modular computing allows customers to leverage the company's holistic system design — which includes computing, storage, I/O, cooling, power, and system management — to roll out a platform that is both expansive and manageable. Sun's attention to the importance of industry-standard, modular components is a cornerstone of the solution. By taking this higher-level view, Sun designed its products to work at the datacenter level, rather than at the rack or blade chassis level, to ensure that industry-standard systems and full management capabilities — hardware, software, power, and cooling — are at the forefront of its modular computing design.

Benefits of design at the datacenter level on industry-standard components include:

- Ease of integration into installed systems and existing infrastructure
- More exacting levels of manageability and RAS
- Consistent I/O throughout the datacenter
- Investment protection

Sun's portfolio of products for its modular computing solution draws from the company's server product lines. The company's modular computing initiative utilizes best-of-breed products and industry-standard components from across its enterprise and volume server product lines. The combination of high-end technologies and volume economics has resulted in a product line that delivers flexibility and enables customers to choose from a broad portfolio of products to best match their needs with system capabilities.

Sun's blade products have focused on offering flexibility, manageability, and RAS in an industry-standard design. This holistic approach to modular computing extends beyond just the hardware. Sun has extended the concept of modularity into the software space, specifically focusing on modularity in the operating system with Solaris 10 Containers, in the management layer with N1 System Manager and N1 Service Provisioning System, and in the application server space with Java Enterprise Systems and Java System Suites.

Sun Solutions with Blades

The Sun Blade 8000 offers enterprise-class RAS as well as hot-pluggable and redundant I/O aggregation, fans, power supplies, disks, and I/O modules. Additionally, the I/O throughput capabilities with the Sun Blade 8000 scale up as the number of cores deployed increases. Each server module can hold four Opteron dual-core processors, 64GB of memory, and two hot-swappable disk drives for up to 160GB of storage per server module.

Modular I/O

Modular I/O is a big differentiator for Sun. The Sun Blade 8000 can reach a maximum theoretical I/O bandwidth of 1.92Tbps across 10 server modules with up to 80 cores — approximately 192Gbps to each server module. This bandwidth is achieved using standard PCI Express links from the module, across the passive backplane, to I/O modules in the back of the chassis. Each server module contains six separate PCI Express links. Two x8 links connect to PCI Express ExpressModules that provide unique dedicated I/O to each server module. The use of ExpressModules in this manner allows each server module to have its own "I/O personality." The remaining four links can be used to aggregate I/O across multiple server modules in a PCI Express Network Express Module with up to four modules per chassis.

This combination of I/O allows customers to dedicate I/O bandwidth and to share different fabrics, including GbE, Fibre Channel, and InfiniBand, within a chassis. It also ensures that customers will have I/O headroom as communication needs expand. Both the ExpressModules and Network Express Modules are hot pluggable and field serviceable without having to open the chassis or server modules.

Additionally, each server module is equipped with Sun's Integrated Lights Out Management (ILOM) service processor. This processor allows the server modules to be managed in the same way a standalone rack-optimized Sun Galaxy server would be managed, ensuring that best practices and administrator training do not need to change simply because the form factor has changed.

Manageability

The common management tools and framework across the x64 line of systems form a key value proposition for Sun. This line includes not just the Sun Fire X4600 and Sun Blade 8000 products but also the Sun Fire X2100, X4100, X4200, and X4500 systems.

The foundation of this framework is the ILOM service processor. This standard feature provides users with remote and secure KVM capabilities via Ethernet or serial port access. A standard Web browser GUI and a command line interface leverage standards such as SNMP and IPMI 2.0 for the collection of important system health and performance information. The incorporation of the service processor into each server module allows users to treat their server modules like all other servers and directly manage the server module without having to manage the chassis.

In addition, customers can use the N1 System Manager for infrastructure management across the entire Sun server portfolio. This software package allows users to manage groups of servers, provision operating system patches and firmware updates remotely, monitor hardware and operating system events, as well as have lights-out and role-based access to the systems. N1 System Manager supports and works with third-party solutions such as certified operating systems including Solaris, Red Hat, and SUSE Linux and with virtualization products from VMware and Xen.

The Sun Blade 8000 for Enterprise Workloads

From a target market perspective, the Sun Blade 8000 is designed for the needs of enterprise workloads, including database, business processing, and data warehousing applications. In addition, as virtualized servers enable another round of server consolidations, the RAS of the platform becomes even more critical as customers place more applications on a single device. In these consolidations, robust and reliable platforms are a must. Moreover, because of the high I/O both to the server module and exiting the chassis, the platform is positioned to deliver balanced performance in I/O-intensive business, high-performance computing (HPC), and technical computing applications that typically run in clustered configurations.

IDC Analysis

In their current state, blade servers are at the beginning of their technological evolution. The significant majority of deployments use dual-processor blades to support infrastructure-tier workloads. IT organizations are not yet convinced of their viability for mission-critical applications.

The global market has already begun to show interest in taking blades deeper into the datacenter by envisioning the ways in which blade systems are architected overall. As a result, blade vendors are responding by designing platforms with the needs of the most demanding enterprise in mind. This shift in design is being driven both by advancements in technology — such as dual-core/multicore improvements in heating and cooling requirements — and by customer demand for deploying blades in support of higher-end applications. For example, the demand for high-end Unix on a blade platform is already emerging in the financial services market — typically a leading vertical and early adopter of technology.

Blade and Rack Servers

Blades are a hot topic of discussion in the current market, but rack-optimized servers are equally important to the future of modular computing. In the datacenter, modularity can be applied to storage systems, rack-optimized servers, and networking, as well as operating systems, middleware, and applications. When modular systems and applications are combined, a powerful and highly leveraged solution emerges for addressing increasingly demanding enterprise requirements.

A key to the growth of modular computing is the I/O element. Blade servers offer unified I/O within a single chassis in support of multiple blade servers. This capability not only results in improved ease of management for these elements, but it also helps encourage the market adoption of key future I/O technologies, such as 10GbE. The modular design of the blade platform allows customers to adopt new technologies when they are ready. It also enables customers to hedge against future uncertainty inherent with rapidly evolving technologies. This concept of choice and investment protection is key to those users who have already found blades such a powerful solution.

While IDC is anticipating a 42.7% CAGR for blade shipments from 2006 to 2010, rack-optimized servers will also grow at a good clip, from an estimated 4.2 million units in 2006 to 5.5 million units in 2010. By 2010, the two form factors will account for nearly three-quarters of the total worldwide market of units shipped.

Ultimately, the vision of modular computing ensures that the right tools can be brought together to create a highly customized solution for a broad set of enterprise IT needs.

Opportunities

Power and cooling, more cost-effective solutions for mission-critical workloads, and investment protection are top-of-mind issues, and Sun's positioning of the Sun Blade 8000 Modular System addresses these issues directly. At the same time, IT organizations are increasingly concerned about rack-optimized server alternatives, which are less efficient to cool, less serviceable, and slower to deploy. By combining these two architectures, Sun can assist customers in determining which set of features is needed in specific sections of their datacenters.

Sun's support for integrated systems that mix Galaxy and Sun Fire servers is designed to provide IT organizations with value, flexibility, and choice. Experimenting with a Sun Blade 8000 will not create a silo but rather can provide a testbed for understanding how this new hardware option fits best in the tier architecture.

Additionally, Sun's modular computing strategy addresses one clear challenge that continues to face blade server vendors: investment protection. By designing the platform with a five-year life cycle in mind, Sun has enabled customers to upgrade modules independently. Thus, the power and cooling modules can be upgraded over time to support next-generation blades that require additional resources. Additionally, by using industry standards, Sun is providing the market with an element of future-proofing in both its server modules and its modular computing solution as a whole. Sun has clearly identified the importance of utilizing industry-standard components in its products and is working to ensure that the message is understood by the wider market.

Challenges

Sun and other vendors must continue to engage customers in industrywide educational discussions that articulate how to map IT requirements to the different server choices available. Monolithic, rack-optimized, and blade servers all have their place, but IT organizations are not always certain which choice is the best fit. Further, the ability to integrate these different server categories may be viewed skeptically by some IT organizations. Explaining and demonstrating the modular computing approach will be essential.

An additional challenge that the market as a whole must address is that of I/O and bandwidth. While much chatter exists in the market today about 10GbE, it is not yet a reality for the majority of customers. Hardware vendors need to be more active in cultivating the ecosystem surrounding 10GbE to heighten awareness of the technology and subsequent demand while simultaneously working to drive down the prices of 10GbE components. Blade server solutions can leverage 10GbE more efficiently than rack-optimized servers by virtue of the chassis I/O design and the streamlined use of I/O within the chassis between the blades. Sun will need to communicate the ability of the Sun Blade 8000 Modular System — as well as other products in its server lines — to scale I/O when increasing the number of processors. This is a critical message to take to the market today.

CONCLUSION

IDC believes that the Sun Blade 8000 Modular System meets the needs of customers in the third cycle of blade server adoption. The modular system reflects Sun's focus on modular computing design for the enterprise application tier of IT. Sun has aimed at developing best-of-breed modules that come together to form a high-performance, high-RAS server. When blade servers move onward from edge workloads to enterprise workloads, design features included in the modular line will be required.