

Preface

Named after the easily replaceable working part of a razor blade, or a readily accessible sword in its sheath, “server blades” are relatively new entrants in the high-volume mainstream server markets. And it is the IBM eServer™ BladeCenter® system that has distinguished server blade architecture for the enterprise, defining a new way for clients to build their data centers and computing infrastructures. The record growth and adoption of BladeCenter systems and the general industry appeal of blade architectures can be attributed to its many distinct benefits. These benefits can be portioned into seven categories, as follows.

1. *Density.* The physical orientation of blades and modules within the enclosure (chassis) results in very efficient use of space, generally yielding server densities greater than traditional-form-factor servers. Blade servers can provide rack densities several times greater than those obtained using traditional-form-factor servers.
2. *Hardware cost.* The contents and structure of the server blade chassis result in lower effective cost of additional server nodes as node density increases. This is due to “node amortization,” which means that the costs of chassis frames, backplane, power supplies, fans, CD, floppy disk drive, and other I/O are shared and amortized across server CPU nodes. Once a threshold number of server CPU nodes is reached, the effective cost of adding more nodes becomes lower than for traditional server form factors (e.g., 1U).
3. *Cable management.* As the number of traditional-form-factor servers increases, the numbers of cables required for power, management, networking, keyboard, video, mouse, etc. increase linearly. A typical full industry-standard rack of standard 1U servers can require as many as several hundred cables. A full rack of server blades can require as few as a couple of dozen cables, resulting in increased reliability, reduced complexity, and reduced cost of ownership.
4. *Setup and configuration cost.* Fewer chassis and less mounting hardware reduce several undesirable dimensions of server management. These include the time required to build and configure a rack of servers, the opportunity for error, and the skill level required to perform setup and configuration. This translates to savings for the supplier on preconfigured racks and for custom configuring within data centers.
5. *Infrastructure integration.* Traditional data center infrastructures have been characterized by physically separate boxes for servers and for networking and

storage area network (SAN) gear. Server blade architectures can embed networking and SAN switches into the chassis. Once “physical integration” is achieved, a platform for operational and administrative integration is established for seamless quality of service throughout the network and server, and for a single systems management view of both the server and network infrastructure.

6. *Reliability, availability, and serviceability.* In contrast to traditional rack-mounted servers, a server blade node can be replaced without the use of tools to manipulate mounting hardware and with no need to remove cables. A server blade can be removed by simply opening the chassis bezel and releasing a locking mechanism. In addition, server blades have optimized accessibility to “hot-swap” components, which are not easily accessible for many rack-mounted 1U-form-factor servers. The server blade design also allows for much simpler network and disk I/O component access as well as easier servicing and upgrades of switches and routers within the chassis.
7. *Flexible function and technologies.* Replaceable processor blades and other modules, such as network routers, SAN switches, and power supplies, provide customers with the opportunity to easily benefit from new and alternative technologies without replacing the entire server system. This flexibility keeps the chassis useful and provides a level of protection against obsolescence.

This issue of the *IBM Journal of Research and Development* contains eleven papers that discuss how these benefits are manifested and describe the many dimensions of the BladeCenter system. The opening paper, by Desai et al., provides an overview of the BladeCenter system and sets the stage for the remaining ten papers. Although this overview and the following papers were comprehensive at the time of their writing, the reader should keep in mind that the BladeCenter system continues to evolve. IBM and the many other BladeCenter Ecosystem companies continually introduce new capabilities for this platform.

The paper by Hughes, Patel, et al. describing the midplane addresses the heart of the BladeCenter system. The complex confluence of many data, address, power, and control signals required expert engineering to architect and package this removable and redundant component. The midplane embodies agnostic interconnections among the processor blades, switches, and modules, allowing versatile network fabric options, from Ethernet, Fibre Channel, and InfiniBand®** I/O to

**Registered trademark of InfiniBand Trade Association in the United States, other countries, or both.

proprietary high-speed protocols. The open midplane hardware specification has facilitated third-party development of many switches and modules and is a chief catalyst in the rapid development of the BladeCenter industry ecosystem.

The versatility of the processor blade is seen in the multiple processor architectures it supports. The processor blade paper, by Hughes et al., illustrates the use of Intel, IBM, and AMD processors, along with two-socket and four-socket CPU system architectures. Such architectural variety in the processor blade contributes greatly to integrated data center solutions, as described in the BladeCenter solutions paper by Fore et al. and the BladeCenter T paper by Vanderlinden et al. Integrated “solutions in a box” have proven to be valuable for clients in the retail, financial, communications, and health sciences market sectors. The BladeCenter T system is a prime example of alternative implementations of the BladeCenter architecture which can be customized for a particular market application such as the telecommunications industry.

In their paper on power, packaging, and cooling, Crippen et al. reveal the significant engineering achievement of cooling and packaging BladeCenter systems. At its introduction, the BladeCenter system delivered the highest-performing and thermally hottest industry-standard microprocessors, in the densest package in the industry. This was accomplished without compromising enterprise computing capability such as memory and network capacity.

The benefits of infrastructure integration are seen in the network paper, by Hunter et al., and the storage paper, by Holland et al. The variety and depth of the BladeCenter switch portfolio has set it apart in the industry. The leading companies in the network and storage industries, such as Brocade, Cisco, McData, Myrinet, Nortel, Qlogic, and Topspin, have built their technology into BladeCenter systems. The benefits of physically integrating the elements of a data center infrastructure are complemented by the integration of the data center management functions within a single view for the systems administrator. The modular replaceable switches, along with the management module and CPU blades, add to the overall steady-state reliability of BladeCenter systems by reducing the mean time to replace (MTTR). In addition, this modularity greatly eases the building of scale-out infrastructures for growing data centers and high-performance computing facilities.

The dominant benefits of blade architectures are realized only if the system components are effectively managed. Complete management ranges from the lowest-level hardware element to the firmware and software and up to the system and rack levels. Brey et al., in their paper on chassis management, and Pruett et al., in their paper

on system management, describe the broad management capabilities of BladeCenter systems. Leading these capabilities are the software tools in IBM Director, which complement the hardware and firmware tools within the management module. Both of these tools are provided as standard with the BladeCenter system.

The paper by Piazza et al. on thermal diagnostics is an example of how the BladeCenter system is conducive to the development of advanced thermal management technologies. As components become more highly integrated and run at higher frequencies and data rates, thermal problems become aggravated. This is further complicated by the demand for sustained or increased density and system availability. The thermal management research described in this paper directly addresses this growing issue, which is among the top concerns of IT professionals today.

The seven benefits of BladeCenter systems discussed at the beginning of this preface are realized through a balanced system design which comprises the hardware, firmware, software, electrical, thermal, and mechanical domains. The eleven papers in this issue offer the reader an illustrated description of how this complex balanced design was achieved. The reader will also gain insight into how and why BladeCenter systems have been the impetus for the emergence of an entire industry, complete with standards bodies, conferences, ecosystems, clones, and the rethinking of IT.

In closing this preface, and on behalf of the entire IBM BladeCenter team, I wish to express our appreciation and gratitude to our chief development and marketing partner, the Intel Corporation.

With great admiration I extend sincere thanks to my colleagues on the original team who conceived the BladeCenter design, to those who developed it for its debut in the market, and to the many who have now joined the BladeCenter community, further shaping it into one of the most profound dimensions in computing today.

Thomas M. Bradicich
Chief Technology Officer
xSeries/BladeCenter Servers
IBM Systems and Technology Group
Guest Editor