

BladeCenter solutions

The architecture and design of the IBM eServer™ BladeCenter® system make it an ideal platform for many information technology (IT) solutions. Its high availability, modularity, flexible infrastructure integration, and on demand configuration options have prompted developers within IBM and in end-user IT organizations to create solutions tailored to a particular industry or function. This paper describes the BladeCenter system in the application of several of these solutions—their target enterprises, end-users, and key benefits. High-performance and scientific clusters, branch-office-in-a-box, financial services, and hosted-client solutions are reviewed.

L. S. Fore
D. W. Cosby
G. Pruett
D. B. Rhoades
C. O. Schulz
B. Smith
J. L. Wooldridge

Introduction

Many features went into the design of the IBM eServer* BladeCenter* system in an effort to meet a wide range of customer requirements in different enterprises and industries. Customers required high availability, denser form factors, rapid deployment capabilities, flexible architectures, simplified infrastructure, and enhanced management capabilities. All of the above requirements had to be achieved in a cost-effective manner using open standards.

- **High availability:** The redundant BladeCenter design offers high availability with no single point of failure for interconnects, power, cooling, networking, storage, or management [1–3].
- **Density:** The BladeCenter system leads the industry in volumetric density and power density; its density is double that of previous 1U rack-optimized servers, resulting in the highest power density (watts per unit volume) used in IBM products to date [4].
- **Deployment:** The shared BladeCenter infrastructure created a need for more sophisticated network-based deployment and configuration tools. Its management software combines multiple management tools and technologies to create an integrated solution for rapidly deploying, managing, and configuring the chassis components and attached storage area networks [5].
- **Flexibility:** The BladeCenter architecture is flexible enough to allow the use of multiple processors (IBM, Intel, AMD), multiple processor architectures (IBM

Power Architecture*, Intel Xeon**, and AMD Opteron** processors), and multiple operating systems [1]. To ensure that the architecture will be flexible enough to support multiple input/output (I/O) fabric protocols, serialized/deserialized (SerDes) is used as the internal high-speed communication electrical interface [2].

- **Simplification:** BladeCenter architecture combines server, network, storage, and management technologies to consolidate and simplify infrastructures built around the scale-out model [6, 7].
- **Management:** The BladeCenter platform itself was designed for remote management [3], with a built-in management module and remote console support [5]. Compute, network, and storage components operate under a common chassis management scheme. BladeCenter chassis management allows administrators to quickly install, configure, inventory, and diagnose their equipment from anywhere on the network [3].
- **Costs:** The BladeCenter architecture reduces acquisition costs by providing shared chassis infrastructure—shared power supplies, compact disk–read-only memory (CD-ROM), floppy disk drive, and Universal Serial Bus (USB) ports for all processor blades in the chassis [5]. Also, by incorporating the server connectivity into the chassis, the cost for even the most elaborate storage solutions can be significantly less than for a comparable solution with individual server boxes [7].

©Copyright 2005 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/05/\$5.00 © 2005 IBM

- *Open standards:* IBM blade products have shifted the server packaging paradigm in the industry and resulted in an open server blade and switch architecture [4]. Industry standards are used throughout the BladeCenter design, from the I/O interconnects to the networking [6].

This paper describes four solutions based on the BladeCenter platform. The cost-effectiveness of the design lends itself to *high-performance computing solutions* that offer one of the industry's lowest cost/performance ratios for clusters. A *branch office and retail store solution* is designed to meet evolving customer demands and run advanced business applications specific to the enterprise. The architecture provides low-latency operations—a requirement for many *financial industry solutions*. And features support the low cost and highly secure *hosted-client solutions* that are efficiently deployed and managed on the BladeCenter infrastructure in the data center.

High-performance and scientific computing clusters

The historical notion of high-performance computing (HPC) is associated with large-scale computations needed to solve complex scientific problems. This idea is most notable in the areas of advanced science and research and is recognized in such areas as academia, life science, material science, simulation, and other disciplines steeped in the traditional pursuit of large-scale mathematical problem solving. Adding to the canonical HPC definition today's emerging nontraditional areas such as finance, digital media, and decision support changes the concept of HPC, and the change manifests a new paradigm, a call for new solutions that require definition in scope and characterization.

Clustering, as a solutions offering in HPC, has existed for quite a few years [8]. It can be defined as a collection of compute, network, storage components, and special software that interact closely in the pursuit of HPC problem solving. Clustering for HPC was once reserved for expensive and proprietary offerings. HPC users always desire the most capability possible for the funds available. To that extent, the academic and national laboratories have continually stimulated experimental ideas which provide economical breakthroughs. Some of those efforts have concentrated on the concept of clustering inexpensive discrete systems to take advantage of additive compute capability. These efforts in economical clustering have yielded some successful implementations, such as the System X supercomputer at Virginia Tech and the Stone Soupercomputer built at Oak Ridge National Laboratory in 1997. The best results to date center on the use of the Linux** operating system (OS) and the broad cooperation of multi-industry

disciplines to build these systems. The results are best described as *Linux clusters*.

Over the last several years, Linux clusters have become a focus for the advancement of economic clustering. Several contributing factors to this development have been the maturation of the Linux OS, improved tools (compilers, libraries, debuggers) for development, compute node performance improvements, improvement of interconnect networks, and developments in storage technology.

With new developments in BladeCenter technology, it has become far better suited to high-performance computing solutions based on Linux clustering implementations. Some of the key advances that make this possible are shown in **Table 1**.

The TOP500** list of supercomputer sites [9] is the product of the TOP500 project, started in 1993 and generally recognized in the HPC community as a means of ranking machines on the basis of their relative performance using the Linpack benchmark [10]. Linpack measures the performance of a machine by solving a system of linear equations, a method often used in HPC programs. It is considered a processor-intensive and interconnect-intensive benchmark. As such, the use of Linpack tends to give the broadest indication of the overall compute strength of a system regardless of its architecture. Hence, BladeCenter capability in the environment of high-performance computing is illustrated by its position on this list. As can be seen in **Table 2**, in the space of a year it climbed from ranking 44th to being the fourth fastest computer in the world. Altogether, it appeared 42 times on the TOP500 list and had an overall share of 8.4% of the TOP500 listings.

The high density, deployment, and management issues associated with HPC clusters make the BladeCenter architecture an ideal platform. With the high density of IBM BladeCenter servers, diskless nodes, and an open system environment, an HPC clustering built on BladeCenter systems offers an excellent price/performance ratio; very high reliability, availability, and serviceability (RAS); and significant cost efficiencies.

Branch office in a box

The value of information technology (IT) in corporate headquarters was established decades ago. Information technology has historically penetrated the branch office of every successful organization distributing goods and services. The continuous drive for lower cost and efficiency gives retail and service organizations the ability to deliver low customer prices and high product availability. The innovative application of technology provides the foundation on which suppliers collect and assemble the data necessary to recognize buying trends, predict future trends, and use that data to gain

Table 1 Notable additions to the BladeCenter design that promote clustering.

<i>Category</i>	<i>Development</i>	<i>Benefit</i>
Compute nodes	HS40 (four-way Intel Xeon blades) JS20 (IBM PowerPC* blades) HS20 (two-way Intel EM64T blades)	Larger cache 64-bit processor 64-bit capability
Networks	Optical passthrough module (OPM) Myrinet host bus adapters (HBAs) InfiniBand** switch module and HBAs	Allows fiber connections High-speed interconnect High-speed interconnect
Storage	Small Computer Systems Interface (SCSI) on two-way Xeon blades	Performance

Table 2 Ranking of BladeCenter systems in the TOP500 list of supercomputer sites.

<i>Period</i>	<i>Rank</i>	<i>No. of processors</i>	<i>Category</i>
Nov 04	4	3,564	Supercomputer center
Jun 04	34	1,064	Finance and banking
Nov 03	44	1,176	Digital media

competitive advantages. It is common to recognize trends from aggregate data assembled at the regional or corporate level.

Customers expect to receive more information about products and services on demand at their point of contact with the organization. In retail organizations, this might take the form of personalized marketing, automated shopping and checkout, and improved inventory management using radio frequency identification (RFID) [11] technologies.

Technologies and requirements for branch offices and retailers

A set of new applications driven by the aforementioned trends are creating new requirements for the branch office and retail store infrastructure. Because of their rich media content, these new applications will be more data-, network-, and compute-intensive than current solutions. Service providers and retailers will present their customers with more information, will collect more information about their customers, and will use information to make the experience more efficient and more personal. Those deploying solutions will also carefully consider the expense and value of new applications, with an expectation of a high return on investment. Some of the new applications currently being deployed include the following:

- Customer identity recognition and data collection.

- In-store information-assisted shopping based on customer connectivity (e.g., smart carts, scan as you shop, Web access).
- Multimedia information displays.
- Interactive kiosks, e.g., design your own kitchen, create customized compact disks (CDs) and digital video disks (DVDs).
- Employee communications.
- Automated inventory.

In addition to the requirements discussed above, branch office operations are becoming more highly dependent on IT, making fault tolerance and high availability a key issue. Data security is a concern, as both regulations and customer wariness can inhibit the collection and use of personal information. Finally, although there is a need to place considerable technology within the store, usually little or no local IT operations expertise will be available.

These new applications and environmental factors call for a new set of technical capabilities that will be vital to new branch office or store IT:

- High-availability platforms (servers, OS, and middleware) for application deployment.
- Remote management and highly automated local operations.
- Data security.
- Voice over Internet Protocol (VoIP) for a low-cost integrated communication infrastructure.
- Radio frequency identification (RFID).
- Flexible server and storage capacity.

Rollout and configuration

The BladeCenter physical package enables easy transportation and installation, including assembly and wiring at the user location, without requiring on-site skilled personnel. An entire server farm requiring racks of servers, KVM (keyboard, video, mouse), monitors,

Table 3 Infrastructure savings with a BladeCenter system for 42 servers: integrated drive electronics (IDE), dual storage area networks (SANs), dual Ethernet, KVM, and redundant power.

	42 1U servers	42 blades	Reduction (%)
Rack space	42U	21U	–50
Ethernet cables	84	6/24	–71
Fibre Channel cables	84	12	–86
KVM cables	42	0/3	–93
Systems management cables	42	6	–86
Power cords	84	12	–86
Power distribution units	8	4	–50
KVM switches	6	0/1	–83

switches, and other appliances can be consolidated into one 7U-high rack [4].

The consolidation of infrastructure (network switching, server management, and KVM and servers [12]) minimizes the cabling that must be done, either centrally or at distributed back-office locations. This integration provides significant cost savings by reducing or eliminating redundant external networking switches, KVM, and management connections [5], as shown in **Table 3**. Many parts, plugs, and wires that can be misconnected, wrongly configured, misplaced, or mishandled are eliminated. Additionally, when the user receives remote technical assistance, these items do not have to be verified as they would if present.

Remote administration

Distributed back-office locations typically do not have the benefit of the most expert technical experts available at the corporate IT headquarters. To that end, the BladeCenter remote management capabilities [3] are provided. The system was designed to be managed over the network [5]. Its integrated management function provides all of the benefits of the remote management adapters typically deployed in each standalone server, but without the cabling complexity of connecting 14 server management cards. Remote administrators can update device firmware, cycle device power, and perform all management functions that can be performed by a local system administrator. The BladeCenter integrated remote KVM provides the look and feel of sitting at the workstation, though the server may be network-attached hundreds of miles away. Extensive Simple Network Management Protocol (SNMP) implementations can notify remotely located expert administrators of nearly every conceivable event, including chassis environmental elements, predictive failure analysis, and cable disconnects. Thus branch offices can benefit from

dedicated—but remotely located—expert support personnel who can function as though they are on site [5].

Flexible growth, expansion, and remote deployment

The BladeCenter architecture was designed to embrace shared storage [4]. Whether attached by Ethernet-based iSCSI (Internet Small Computer System Interface) [13] or Fibre Channel, dedicated storage subsystems are typically more highly available and more easily maintained than most dedicated server file systems; the economy of shared storage systems makes these features more affordable. Rather than being distributed in individual servers, excess storage that must accommodate future expansion and growth is available in a common pool that can be partitioned and assigned to servers as needed. By leveraging the existing Ethernet infrastructure for connectivity, iSCSI implements shared storage at a much more attractive price point than typical Fibre Channel pricing. While Fibre Channel solutions (2 Gb/s) typically provide higher performance, iSCSI performance (1-Gb/s Ethernet) is more than sufficient for most back-office implementations.

As applications and customer needs evolve, the need for additional servers often arises. Whether the need is for an Intel- or an IBM PowerPC*-based server, the BladeCenter architecture can accommodate the expansion, and it supports common operating systems—Microsoft Windows**, Linux, IBM AIX*, etc. [6].

Because the architecture provides common, consistent interfaces and a consolidated infrastructure in a fixed environment, rich software deployment support is feasible and cost-effective. Many flexible deployment options are possible. Spare blades can be deployed, then imaged as needed to provide redundancy and availability. In the event of catastrophic software failure, bare metal restores can be initiated by automatic notification of blade inserts with the option of administrator approval.

Solutions for financial services

In the financial services industry (FSI), the growth in the number of electronic transactions has accelerated as companies work to differentiate their offerings, lower their operational costs, and achieve greater efficiency. To accomplish this, they must reduce the response time between systems and between systems and clients during online transactions. They must also increase the number of systems and clients that can concurrently participate in transactions, and they must be able to quickly introduce new line-of-business solutions to the marketplace via both applications and new hardware.

In FSI, the issue of latency involves any flow of data or information to customers, the implementation or validation of new ideas, or the steps needed to pursue new business opportunities. For example, people who do electronic equities trading or manage investment risk need to perform complex Monte Carlo simulations as quickly as possible to capitalize on market fluctuations, or they may miss an opportunity exploited by a competitor. This is the essential need FSI customers have for higher speeds and reduced latency.

To reduce latency, FSI firms seek the ultimate commodity hardware and operating system that gives them the ability to implement complete line-of-business opportunities faster than the prior technology allowed. A BladeCenter system is a choice low-cost solution that can give FSI clients best-of-breed technology for their application environments. It conforms to open standards, which provides users with many choices in processors (Intel Xeon, IBM PowerPC, two-way, four-way, etc.) [1], switching (Ethernet with vendors such as Cisco, Nortel, and IBM, and Fibre Channel [6] with vendors Broadcom and QLogic), or new I/O standards (e.g., InfiniBand I/O [14]). This allows building blocks to be defined as standards within IT organizations that can be configured to achieve maximum reduction of latency in FSI solutions.

Many new lines of business opportunities have traditionally been built upon smaller quantities of large servers (mainframe or UNIX**). The specific IT business process and operating procedures associated with these larger servers has become a prohibitive factor in delivering new business solutions to the market quickly. Reduction latency and cost performance has made possible a blade-based scale-out solution for many of these new applications today. In addition to the hardware performance and cost benefits, the distributed server environment associated with scaled-out BladeCenter servers does not have the legacy of regimented infrastructure and process controls associated with the larger mainframe servers in FSI IT organizations. This allows IT departments to run a more efficient engineering organization with BladeCenter systems.

In FSI firms, Linux server farms are generally managed by the UNIX engineering organizations, which traditionally have responsibility for the various versions of UNIX in data centers. These administrators have requirements for many of the same remote systems management functionality that they have on their higher-end UNIX platforms. The BladeCenter design consolidates the configuration, control, and management in one central control point, called the management module, which is standard in each chassis. The management module introduced SOL capability [3, 15] that allows Linux system administrators to perform terminal server operations similar to those on their UNIX servers without the need for additional expensive external hardware and associated serial cables. The SOL interface allows Linux administrators to develop and execute custom and policy-based scripts that control power, cycle the OS, or load and reload the OS and applications. With the development of custom scripts, SOL allows IT administrators to manage a chassis of blades with the same effort as an individual server.

The majority of the early FSI installations involved new lines of business applications that executed on Linux rather than the traditional AIX*, Sun Solaris**, and HP-UX versions of UNIX. The current economics of BladeCenter and Linux technology enable these IT organizations to launch new business applications in a timely, competitive, and cost-sensitive business environment.

BladeCenter infrastructure consolidation and choice of technologies allow its integration into large data centers without requiring customers to change their operational standards. This consolidation provides flexibility and infrastructure cost savings. Ecosystem options are available from the industry leaders in infrastructure connectivity (Cisco, Nortel, Broadcom, QLogic, TopSpin, Myranet, etc.) in addition to passthrough devices for Ethernet (copper and optical) and Fibre Channel (optical) that allow customers to leverage their current infrastructure.

The flexibility of BladeCenter I/O is vital to FSI worldwide data center installations, since the various geographies may have different operational standards within their IT infrastructure. This usefulness of this flexibility is also intensified by the fact that many clients in the financial services sector have merged with other financial companies over time (industry consolidation), and these organizations may have different standards within their current IT infrastructure.

For example, a data center in New York may have a 1-GB copper Ethernet network interconnecting its servers, and a data center in London may have a 1-GB optical Ethernet. The BladeCenter modular building block solution allows for maximum integration and

abstracts the change from the rest of the organizations. This modular design approach also allows worldwide engineering teams to focus on core components in their area of expertise. It is common to have engineering teams certifying BladeCenter components in different countries, thus enabling the FSI firm to leverage key skills in the company to further maximize time to market of new business applications while also reducing IT costs.

The BladeCenter environment, with its space planning, power, and cooling, is another example of leveraging IT skills in a worldwide perspective. For security reasons after 9–11, FSI firms have geographically dispersed their operational data centers. The ability to have a worldwide standard on blade hardware enables facilities engineering teams to standardize data center designs and accommodate different regulatory agency power requirements. This is another way in which the modularity of the BladeCenter system and its ability to uniquely customize installations reduces organizational support costs.

Although the majority of early installations in the financial services sector were Linux environments, this critical mass of IBM blades in worldwide data centers is now promoting installations of Windows-based environments. Since the BladeCenter system is already certified and implemented in these infrastructures, all of the necessary work associated with the certification process for the ecosystem, environment, and management does not have to be repeated in order to have the Windows operating environment certified; and because it is already certified and implemented in FSI infrastructures, the necessary work associated with the interconnect infrastructure, environmental areas such as power and cooling, and management does not have to be repeated.

Hosted-client solutions

IBM hosted-client solutions provide the infrastructure required for hosting desktop sessions on remote server hardware. In a hosted-client environment, the desktop OS and applications execute on servers in a remote data center.

By using the BladeCenter architecture for hosted clients, the data center infrastructure is simple to deploy [5], administer [3, 5, 6], and troubleshoot [3] at a significantly lower cost of ownership [2, 4] and higher density [5] than other form-factor servers. End-user storage, securely hosted in the data center, is remote from the user desktop. BladeCenter offers a wide range of local storage options and remote storage connectivity options [7].

End-users interact with a client device that can be supplied with a variety of software and hardware component options. The hosted-client work described in this paper is a result of the combined efforts of IBM

eServer development teams, Tivoli* software developers, and researchers at the IBM Thomas J. Watson Research Center in Hawthorne, New York.

The subject hosted-client solutions refer specifically to hosting Windows and Linux OS and application sessions and providing a network-based connection from the client to those sessions. The hosted-client architecture allows existing Windows and Linux applications to be hosted. The hosted client supports the transition from existing environments with thousands of desktop personal computers (PCs) to a centrally managed desktop infrastructure that is more secure, manageable, and highly available. Hosted-client architecture allows end-users to restructure around thin clients and data-center-hosted servers while preserving the current application investment. The BladeCenter platform is ideal for hosting the scale-out compute infrastructure [4], storage [7], and networking [6] required to host the desktop OS.

Requirements in desktop computing

Key customer requirements in desktop computing are lower total cost of ownership, improved security, higher availability, and better manageability, as shown in **Table 4**.

Hosted-client solution benefits

The hosted-client environment provides server-class reliability to desktop users and allows the desktop compute and storage infrastructure to be distributed across geographically dispersed sites. The hosted-client solution promotes a new paradigm for desktop computing, with a number of advantages:

- Minimizes desktop clutter.
- Minimizes desktop noise, power, and HVAC (heating, ventilation, air conditioning) requirements.
- Improves data security; eliminates data stored on the desktop and moves it to the data center.
- Improves disaster recovery; end-user state and data moved from the desktop to high-availability clustered data center servers.
- Reduces network latency between desktop applications and server applications.
- Significantly reduces the cost of managing desktop environments.
- Provides the ability to customize, test, and refine a solution-delivery infrastructure in a test environment and captures the configuration and deployment process to minimize the cost in deployment to a production environment.
- Offers repeatable, consistent, and integrated deployment/undeployment of resources from infrastructure.

Table 4 Customer needs in desktop computing.

<i>Customer needs</i>	<i>IBM focus</i>	<i>BladeCenter feature implementation</i>
Lower total cost of ownership	• Reduce costly moves, adds, and changes • Reduce end-user support and distributed maintenance expenses	• Changes are centralized • Support can be done in the data center
Improved security	• Reduce user accessibility to software, data, and certain hardware • Limit user access to hosted applications via screen data only • Provide centralized backup and data protection	• Prevention of unwarranted changes to hardware or software configurations
Higher availability	• Provide enterprise server levels of reliability and availability not affordable in a desktop	• Spare blades for failover
Better manageability	• Control software installations and access to hardware • Centralize hardware maintenance	• Ability to centrally manage software licenses • Lowered risk of theft and unauthorized modifications

- Provides on demand capability to share resources across various application environments to minimize total cost of ownership.
- Provides an integrated, validated, and tested environment which includes components that address major aspects of the hosted-client solution life cycle.

Hosted-client architecture

As depicted in **Figure 1**, the hosted-client architecture is composed of a three-tiered infrastructure. The tiers include the client tier (end-user interaction), the compute engine tier (hosts the desktop OS and applications), and the state management tier (manages all state associated with an end-user's desktop experience).

Hosted-client solutions enable access to desktop services from a wide range of client devices, including thin clients with embedded OSs, tablet and hand-held devices, and traditional repurposed fat clients running locked-down OSs. A *thin client* may take the form of a terminal, such as a Neoware** [16] product, or it may be a reprovisioned desktop with a locked-down OS software load. A *fat client* is a PC desktop running a local OS with local applications that connects into remote-hosted sessions, such as a Citrix-hosted application. The client device provides KVM for the remote session. Depending on the environment, the client device may also provide connections to local printers, I/O devices (USB memory key, floppy disk, CD/DVD-ROM), and audio.

The hosted-client compute engine can be constructed in multiple ways. The three primary solution approaches are the physical 1:1, terminal sessions, and virtualized. In the physical 1:1 solution, each end-user is running a separate OS instance on a dedicated PC blade. In the terminal session solution, each user is connected to a terminal session, and the applications execute within this

session. Multiple concurrent sessions are supported on a single OS instance. In the virtualized solution, each end-user is running a separate OS instance. The end-user's desktop OS is running in a virtual machine. A hypervisor, such as VMware ESX Server** or Microsoft Virtual Server, allows a single physical server to host multiple "guest" operating systems in individual virtual machines. All three approaches leverage a remote desktop network protocol. The three most common protocols in use today are Microsoft Remote Deployment Protocol [17], Citrix ICA** [18], and XDM (Linux/UNIX X Windows) [19].

The goal of hosted-client solutions is to provide an integrated, tested reference architecture that simplifies deployment, management, and maintenance of the hosting infrastructure while retaining the capability to customize the offering to specific customer needs through related service offerings. This solution is expected to provide automated deployment and provisioning of servers, OSs, applications, and management tools. The choice of the BladeCenter system as the hosting infrastructure provides the scale-out features, network deployment, configuration tools, and management capabilities required.

The management software includes the following:

- IBM Director for BladeCenter hardware management [20].
- IBM Remote Deployment Manager (RDM) [21] for bare-metal OS installation.
- Connection broker, which provides an administrative interface into the solution. It controls session allocation and provides integration with IBM Director for troubleshooting and problem determination.




Figure 1

Hosted-client architecture. The goal for the eServer BladeCenter system is stateless blades and no DASD (direct access storage device).

- Tivoli Provisioning Manager [22] (TPM) for build-out and maintenance of the server and application environments.
- Customizable templates to define and configure the TPM data-center model, including the application environment, networking requirements, resource pools, and dynamic provisioning attributes.
- TPM workflows [23] and scripts that utilize IBM Director and RDM to perform bare-metal OS and software stack installation on blades [24].
- TPM workflows to install and configure the hosting software (e.g., Citrix, VMware ESX Server, or Microsoft Virtual Server).
- Infrastructure to propagate BladeCenter events to TPM using IBM Director as the management proxy.

The connection broker and IBM Director provide the administrative operational interface into the hosted-client system; this interface is used primarily by the help desk for troubleshooting and correcting end-user problems, such as restarting sessions or forcibly logging off users.

The TPM provides a more sophisticated user interface for developing workflows to automate server resource allocations; this interface is used primarily by data center architects and programmers to customize or troubleshoot workflows.

Solution components

The hosted-client infrastructure is managed by three primary software components: an active directory server, a connection broker, and a virtualization engine management server (TPM, IBM Director, RDM). The role of each component is summarized below:

- *Active directory:* Provides Kerberos¹ and Lightweight Directory Access Protocol (LDAP) services to the clients. A ServiceConnectionPoint class object of the LDAP schema is used to publish connection broker mappings to network segments—it is a primary,

¹ Kerberos is an Internet Engineering Task Force standard for providing authentication. It works by having a central server grant a "ticket" that is honored by all networked nodes running Kerberos.

secondary, ternary connection broker for clients on particular network segments or Domain Name Server suffixes. The LDAP schema may also define user profiles, including sessions, user class, and user capabilities (e.g., whether the user is allowed to attach USB storage devices to client). The hosted-client user profile information may be stored either in an unused attribute in an existing schema or in an extended LDAP schema that defines a new multivalued attribute.

- *Connection broker*: Tracks session state, allocates users to appropriate session, and provides failover to a different OS instance. The connection broker includes a simple operator user interface, including command-line interface and graphical user interface.
- *Virtualization engine* [25]: Includes the IBM Director and TPM tools used for the hosted-client solutions. IBM Director provides hardware management. TPM provides provisioning of additional host platforms, reprovisioning of host platforms, cloning of virtual machines (VMs), power control of VMs, movement of VMs to another blade, and suspension of idle hosted-client sessions.

The blades in the BladeCenter chassis provide the compute engine hardware infrastructure. The hosting software stacks execute on the compute engine (i.e., Citrix Metaframe** for terminal sessions or on the hypervisor and Windows XP guest VMs for the virtualized solution). The SAN provides Fibre Channel-attached storage for the databases and for the VMs.

The management server is provided on a separate standalone server, which is typically outside the blade chassis. The software stack for the management server is shown in **Figure 2**.

Conclusion

Four solutions leveraging the BladeCenter architecture and design were reviewed. The high-performance computing community has embraced the open ecosystem, the scaling capabilities, and the performance of BladeCenter clusters. The branch-office-in-a-box solution benefits from the balanced, flexible integration of high-performance servers and the management infrastructure. This creative integration consolidates the distributed back-office and retail branch, enabling unique, significant efficiencies. Solutions in the financial services sector have taken advantage of the modularity, the flexibility of its design, the integration of management and configuration of chassis hardware components, and, most significantly, the low latency of the intrablade communications. The hosted-client solution encompasses the software and hardware infrastructure required for hosting desktop




Figure 2

Software stack for hosted-client solutions. (TIO = Tivoli Intelligent Orchestrator.)

sessions on remote IBM BladeCenter servers. This provides server-class reliability with storage securely hosted in the data center, and simplifies and reduces the cost of desktop management.

Acknowledgments

The authors gratefully acknowledge constructive comments from Dr. Tom Bradicich of the IBM Systems and Technology Group.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Intel Corporation, Advanced Micro Devices, Inc., Linus Torvalds, TOP500.Org, InfiniBand Trade Association, Microsoft Corporation, The Open Group, Sun Microsystems, Inc., VMware, Inc., Neoware Systems Inc., and Citrix Systems, Inc. in the United States, other countries, or both.

References

1. J. E. Hughes, M. L. Scollard, R. Land, J. Parsonese, C. C. West, V. A. Stankevich, C. L. Purrington, D. Q. Hoang, G. R. Shippy, M. L. Loeb, M. W. Williams, B. A. Smith, and D. M. Desai, "BladeCenter Processor Blades, I/O Expansion Adapters, and Units," *IBM J. Res. & Dev.* **49**, No. 6, 837–859 (2005, this issue).
2. J. E. Hughes, P. S. Patel, I. R. Zapata, T. D. Pahel, Jr., J. P. Wong, D. M. Desai, and B. D. Herrman, "BladeCenter Midplane and Media Interface Card," *IBM J. Res. & Dev.* **49**, No. 6, 823–836 (2005, this issue).
3. T. Brey, B. E. Bigelow, J. E. Bolan, H. Cheselka, Z. Dayar, J. M. Franke, D. E. Johnson, R. N. Kantesaria, E. J. Klodnicki, S. Kochar, S. M. Lardinois, C. M. Morrell, M. S. Rollins, R. R. Wolford, and D. R. Woodham, "BladeCenter Chassis Management," *IBM J. Res. & Dev.* **49**, No. 6, 941–961 (2005, this issue).
4. M. J. Crippen, R. K. Alo, D. Champion, R. M. Clemo, C. M. Grosser, N. J. Gruendler, M. S. Mansuria, J. A. Matteson, M. S. Miller, and B. A. Trumbo, "BladeCenter Packaging, Power, and Cooling," *IBM J. Res. & Dev.* **49**, No. 6, 887–904 (2005, this issue).
5. G. Pruett, A. Abbondanzio, J. Bielski, T. D. Fadale, A. E. Merkin, Z. Rafalovich, L. A. Riedle, and J. W. Simpson, "BladeCenter Systems Management Software," *IBM J. Res. & Dev.* **49**, No. 6, 963–975 (2005, this issue).
6. S. W. Hunter, N. C. Strole, D. W. Cosby, and D. M. Green, "BladeCenter Networking," *IBM J. Res. & Dev.* **49**, No. 6, 905–919 (2005, this issue).
7. W. G. Holland, P. L. Caporale, D. S. Keener, A. B. McNeill, and T. B. Vojnovich, "BladeCenter Storage," *IBM J. Res. & Dev.* **49**, No. 6, 921–939 (2005, this issue).
8. G. Pfister, *In Search of Clusters: The Ongoing Battle of Lowly Parallel Computers*, Prentice-Hall, Upper Saddle River, NJ, 1998.
9. TOP500 project; see <http://www.top500.org/>.
10. J. J. Dongarra, P. Luszczek, and A. Petitet, "The LINPACK Benchmark: Past, Present, and Future," December 2001; see <http://www.netlib.org/utk/people/JackDongarra/PAPERS/hpl.pdf>.
11. Association for Automatic Identification and Mobility; see <http://www.aimglobal.org/technologies/rfid/> and <http://www.aimglobal.org/standards/rfidstds/RFIDStandard.asp>.
12. D. M. Desai, T. M. Bradicich, D. Champion, W. G. Holland, and B. M. Kreuz, "BladeCenter System Overview," *IBM J. Res. & Dev.* **49**, No. 6, 809–821 (2005, this issue).
13. SNIA IP Storage Forum; see <http://www.snia.org/ipstorage/about/iscsi/>.
14. InfiniBand Trade Association; see <http://www.infinibandta.org/home>.
15. IBM Corporation, Serial over LAN Setup Guide—IBM eServer BladeCenter; see <http://www-307.ibm.com/pc/support/site.wss/document.do?lnocid=MIGR-54666>.
16. IBM Corporation, Neoware thin clients; see <http://www.pc.ibm.com/us/neoware>.
17. Microsoft Corporation, Remote Desktop Protocol (RDP) Features and Performance, white paper; see <http://download.microsoft.com/download/8/4/f/84fc80a7-661f-4c96-b5d6-cf73903b09f2/rdpfandp.doc>.
18. D. Miles, "Implementing Windows Terminal Server and Citrix MetaFrame on IBM eServer xSeries Servers," *IBM Redpaper* (April 2003); see <http://www.redbooks.ibm.com/redpapers/pdfs/redp3629.pdf>.
19. K. Taylor, "XDM and X Terminal mini-HOWTO," Revision v1.00, Linux Online Inc., May 15, 2002; see <http://www.linux.org/docs/ldp/howto/XDM-Xterm/index.html>.
20. R. Marchini, "IBM BladeCenter Systems Management with IBM Director V4.1 and Remote Deployment Manager V4.1," *IBM Redpaper* (February 4, 2004); see <http://www.redbooks.ibm.com/abstracts/redp3776.html>.
21. IBM Corporation, "Using RDM 4.11 to Deploy BladeCenter 8839," June 23, 2004; see ftp://ftp.software.ibm.com/pc/pcbbbs/pc_servers/usingrdm4_11todeploybladecenter8839.pdf.
22. E. Manoel, S. C. Brumfield, K. Converse, M. DuMont, L. Hand, G. Lilly, M. Moeller, A. Nemati, and A. Waisanen, "Provisioning On Demand Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator," *IBM Redbooks* (December 21, 2003); see <http://www.redbooks.ibm.com/abstracts/sg248888.html>.
23. E. Manoel, M. Hicks, M. Moeller, I. Naick, M. Poulson, and J. Surmannet, "Developing Workflows and Automation Packages for IBM Tivoli Intelligent ThinkDynamic Orchestrator," *IBM Redbooks* (May 17, 2004); see <http://www.redbooks.ibm.com/abstracts/SG246057.html?Open>.
24. R. Credle, A. Chechi, B. Fenoglio, J. Haskins, K. Kekessie, J. Kenefick, G. Pruett, B. Rajaraman, D. Ulrich, and T. Verhoeven, "IBM Web Infrastructure Orchestration," *IBM Redbooks* (November 19, 2003); see <http://www.redbooks.ibm.com/redbooks/pdfs/sg247003.pdf>.
25. IBM Corporation, IBM eServer Software Information Center, Virtualization Engine; see http://publib.boulder.ibm.com/infocenter/eserver/v1r1/en_US/index.htm?info/icmain.htm.

Received January 2, 2005; accepted for publication February 23, 2005; Internet publication October 7, 2005

Lynn S. Fore *IBM Systems and Technology Group, 3039 Cornwallis Road, Research Triangle Park, North Carolina 27709 (sls@us.ibm.com)*. Mrs. Fore is a Senior Engineer in the xSeries*/BladeCenter Architecture and Technology Department. She received a B.S. degree in computer engineering from Clemson University in 1987. She joined IBM in 1988 and worked in system firmware development in the Personal Computing Division for five years, followed by system management development for six years. Mrs. Fore has been part of the xSeries/BladeCenter Architecture and Technology team since 1998.

David W. Cosby *IBM Systems and Technology Group, 3039 Cornwallis Road, Research Triangle Park, North Carolina 27709 (cosbydw@us.ibm.com)*. Mr. Cosby works in BladeCenter ecosystem development. He received dual B.S. degrees in electrical and computer engineering from North Carolina State University in 1991, followed by an M.S. degree in computer engineering in 1993. He joined the IBM Networking Hardware Division in 1992, focusing on high-performance device drivers. Before joining the BladeCenter team, Mr. Cosby worked in the IBM Microelectronics Division as the technical leader for network processor system integration.

Gregory Pruett *IBM Systems and Technology Group, 3039 Cornwallis Road, Research Triangle Park, North Carolina 27709 (pruett@us.ibm.com)*. Mr. Pruett is a Senior Technical Staff Member in xSeries software solutions. He works as a software architect and strategist, specializing in provisioning and remote deployment solutions. He holds B.S. degrees in mathematics and computer science from Furman University and an M.S. degree in computer science from the University of North Carolina. He has eight years of experience in systems management and deployment software, working as lead developer and architect on the IBM Director and IBM Remote Deployment Manager products. Mr. Pruett has 30 issued or pending patents in areas including RFID, BladeCenter design, bare-metal OS installation, and on demand provisioning.

David B. Rhoades *IBM Systems and Technology Group, 3039 Cornwallis Road, Research Triangle Park, North Carolina 27709 (drhoades@us.ibm.com)*. Mr. Rhoades is a Senior Technical Staff Member in the IBM Systems and Technology Group. He currently works on the BladeCenter Product Strategy/Xpert team as the development interface with Wall Street account teams and other large account opportunities related to BladeCenter and systems-management products. He received a B.S. degree in computer and electrical engineering from Purdue University in 1983. He joined the IBM Entry Systems Division in Boca Raton in 1984, working on the storage subsystems for the PC/XT*, PC AT*, and PS/2 products. Mr. Rhoades has numerous issued or pending patents and has been recognized as a Master Inventor.

Charles O. Schulz *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (cschulz@watson.ibm.com)*. Mr. Schulz is a Senior Engineer at the IBM Thomas J. Watson Research Center, supporting the IBM xSeries products. He received his B.S. and M.S. degrees in electrical and electronic engineering from North Dakota State University in 1971 and 1972, respectively. Before joining IBM in 1990, he held engineering and management positions at various aerospace and computer companies. He has extensive experience in high-reliability and fault-tolerant computer design as well as computer design for real-time control of aircraft and critical

aircraft systems. His current research interests include computer architecture for high-performance scalable and partitioned servers. Mr. Schulz has one issued patent and fifteen pending, as well as various technical publications on computer architecture and design.

Ben Smith *IBM Systems and Technology Group, 3039 Cornwallis Road, Research Triangle Park, North Carolina 27709 (smithbe1@us.ibm.com)*. Mr. Smith joined IBM in 1999. He received a B.S. degree in engineering technology from the University of South Florida in 1981. He has spent his career in high-performance computing (HPC) working on the design and use of HPC hardware and software solutions. Mr. Smith has recently been working as chief architect for the Cluster 1350 product offering.

James L. Wooldridge *IBM Systems and Technology Group, 3600 Carillon Point, Kirkland, Washington 98033 (wooldrij@us.ibm.com)*. Mr. Wooldridge is an xSeries software architect. He received a B.S. degree in computer science from Washington State University in 1987. Since joining IBM in 1989, he has worked on a number of IBM products based on Intel processors and Microsoft operating systems. He worked on the IBM OS/2* network fileserver and the Windows NT** port to PowerPC, and he led the enablement and optimization of Windows on nonuniform memory access (NUMA)-based Intel servers. Recently he has worked on the Windows Enterprise Workload Manager (EWLM), resource management, x86 virtualization, and Intel-based hosted-client solutions. Mr. Wooldridge has one issued patent and three pending.