

## IBM System Cluster 1350



IBM System Cluster 1350

---

### Highlights

---

- ***Helps reduce time to deploy Linux® and Windows® clusters***
- ***Provides industry-leading cluster choice with integrated IBM and OEM components***
- ***Drives power and cooling benefits from IBM Systems innovation***
- ***Offers single point-of-contact for worldwide Cluster 1350 hardware warranty support***

### Reduced time to deployment

IBM high-performance computing (HPC) clustered solutions offer significant price/performance advantages for many high-performance workloads by harnessing the advantages of low-cost servers plus innovative, easily available, open source and commercially available software.

Today, some businesses are using their own resources to build Linux and Windows clusters using commodity hardware, standard interconnects and networking technology; open source software; and in-house or third-party applications. Any savings realized from a potentially lower acquisition cost offered by these systems is offset by the expense and complexity of assembling, integrating, testing and managing these clusters from disparate, piece-part components.

IBM has designed the IBM System Cluster 1350 to help address these challenges. Our clients benefit from IBM's extensive experience with HPC to help minimize complexity and risk. Using advanced Intel® Xeon®, AMD™ Opteron™ and IBM POWER6™-based server nodes, proven cluster management software, and optional high-speed interconnects, the Cluster 1350 offers the best of IBM and third-party technologies. As a result, clients can speed up installation of an HPC cluster, simplify its management and support, and reduce mean time to payback.

The Cluster 1350 is designed to be an ideal solution for a broad range of application environments, including industrial design and manufacturing, financial services, life sciences, government and education. These environments typically require excellent price/performance for handling HPC and business-performance computing workloads. It is also an excellent choice for applications that require horizontal scaling capabilities, such as Web serving and collaboration.

### **Industry-leading cluster choice**

As an HPC clustering solutions provider, IBM offers numerous choices in cluster configurations and platforms. Clusters can be configured by IBM to meet a wide variety of client needs and accommodate a huge array of technical workloads with a broad range of server platforms (rack mounted or blade servers), processor choices, accelerator options, robust storage solutions, networking/communications fabrics and operating systems. All Cluster 1350 components are thoroughly tested in IBM engineering and test laboratories for integrated cluster functionality.

IBM works extremely closely with ISV and hardware partners. It is essential to our clients that IBM HPC clusters employ the industry's finest components so the Cluster 1350 standards are very high and the testing processes are rigorous. Clients who choose IBM are working with a world leader in HPC, as demonstrated by the continued leadership by IBM of the "Top 500 Supercomputers" list published twice per year.<sup>3</sup>

The Cluster 1350 leverages the innovative technology built into IBM System x™ and IBM BladeCenter® servers. Since servers generally make up the majority of any HPC cluster, the IBM technology leadership is crucial to the performance, maintainability, energy efficiency and reliability of the cluster. The System x and BladeCenter servers incorporate an impressive array of energy and thermal management tools and technologies, delivering the benefits of this engineering in power and cooling advantages to the Cluster 1350.

IBM System Cluster 1350 solutions are shipped from the factory assembled, cabled and ready for rapid deployment. IBM offers a variety of implementation services, often in conjunction with IBM Business Partners, to help ensure the smoothest possible system set up at the client site.

HPC experience and leadership; choices in accommodating client needs and workloads; rigorous pre-testing of all components and operating systems; leading-edge server technologies; and

best-of-breed switching technologies along with a single point-of-support for the entire cluster—all mean that IBM is **the** choice for HPC clustering solutions.

### **Expanding possibilities**

The IBM System Cluster 1350 also offers the opportunity to take advantage of IBM General Parallel File System™ (GPFS™) for Linux to expand and enhance cluster data-storage requirements. GPFS is a high-performance, scalable, shared-disk file system that provides fast data access from all nodes in an HPC cluster and NFS export capabilities outside the cluster. Parallel applications running across multiple nodes of the cluster, as well as serial applications running on a single node, can readily access shared files using standard UNIX® file system interfaces. Furthermore, GPFS can be configured for failover from both hard disk drive and server malfunctions.

In addition, the Cluster 1350 incorporates a range of IBM System Storage™ offerings, providing highly reliable data storage for business-critical applications

that require high-speed transfer and large amounts of data. In addition to support for the DS4200, DS4700 Express, and DS4800 disk systems, additional mid-range disk systems include the DS3200, DS3400 or EXP3000. These components can be combined with storage expansion and switch options to provide robust storage solutions for cluster applications.

### **Customized service offerings**

IBM consultants can help configure a complete cluster solution, including all the necessary hardware, designed for the client's specific needs. For example, Installation Support Services are available through the IBM Cluster Enablement Team (CET). CET is a technical team of HPC cluster software engineers, technicians and project managers skilled in the very latest clustering hardware and software technology. Each CET project is professionally managed by a dedicated project manager to provide efficient delivery and deployment of the client's cluster. This team can provide clients with HPC cluster services such as project management; pre-configuration and cluster burn-in; software installation and interoperability (operating system, cluster

managers, GPFS, schedulers and compilers); customized CSM and GPFS training; skills transfer; and code porting and optimization.

In addition, software installation services are available from IBM Global Services or IBM Business Partners. These services may include planning, installation and configuration of the cluster operating system as well as additional cluster software. For even higher levels of support, clients may take advantage of the optional SupportLine for Linux and Windows clusters. The IBM SupportLine services cover all cluster hardware components, the Linux and Windows Compute Cluster Server 2003 operating systems, and GPFS for Linux software.

To further help simplify the deployment effort, IBM offers project management support to coordinate all aspects of delivery and installation, including hardware and software set-up services. Additionally, attractive financing and leasing terms are also available through IBM Global Financing.

Additional services and warranty are offered as follows:

- *Cluster 1350 hardware installation is included at no charge on 42U and 25U racks.*
- *HPC cluster software services and SupportLine for Linux and Windows clusters are available as optional fee-based services.*
- *Cluster Installation Support Services are available through the Cluster Enablement Team (CET) as optional fee-based services.*

Basic limited warranty service on most IBM and non-IBM components: 3-year, next-business-day, on-site support for selected components. Individual nodes retain the warranty and service upgrade offerings for that IBM machine type. Enhanced warranty service plans are available.

#### **HPC cluster management**

IBM offers Cluster Systems Management software (CSM) for Linux, an advanced cluster management software tool that allows single point-of-control management for a cluster of supported Xeon, Opteron, and

POWER6-based servers. This tool, an optional cluster component, helps simplify the management of the cluster and easily scales with the cluster to help improve the efficiency of the system administrator.

CSM allows the administrator to monitor both hardware and software events from a single local or remote location and can generate automatic notification and recovery actions as defined by the administrator. Furthermore, CSM provides node grouping capability, a convenient and essential tool to allow an administrator to apply rules and commands to an individual node, groups of nodes or the entire cluster. These cluster control event monitoring and automated recovery features help enhance the efficiency of the administrator and the reliability of the cluster by enabling highly effective node control and rapid problem detection and resolution.

#### **IBM Systems innovation**

The Cluster 1350 helps enable clients to choose the right advanced server technology to meet the specific application needs of their environment by offering configurations built on Intel Xeon, AMD Opteron, IBM POWER™

and Cell Broadband Engine™ processor-based servers. The latest Cluster 1350 offering provides a broad choice of excellent server technologies with rack-based server options, including IBM System x3550, IBM x3650, IBM x3455, IBM x3655, IBM x3755, and x3450. The x3450 is a purpose-built 1U server that is suited for HPC. Providing both economic and performance benefits, the x3450 delivers high performance with enough DIMM slots for demanding workloads.

Our blade offerings include the Intel processor-based IBM BladeCenter HS21 and HS21-XM, AMD Opteron processor-based LS21/LS41, POWER6-based JS22, and Cell-based QS21 servers. New Intel Xeon quad-core systems enable clients to significantly grow compute power by upgrading or replacing their current systems while maintaining the same hardware footprint or consolidating applications onto fewer, more powerful servers. The increased efficiency helps maximize control over energy spending and space utilization in the data center.

System x and BladeCenter servers are built on the IBM X-Architecture® blueprint, which melds industry-standard components with IBM innovation to produce servers that offer a unique set of capabilities to effectively manage your business. All of these servers offer affordable enterprise power and control for processor-intensive applications. In addition, the BladeCenter HS21 provides additional benefits for clients who value the combination of performance, density, integration and investment protection offered by its revolutionary architecture, which delivers up to double the density available in 1U servers.

The IBM BladeCenter JS22 Express server further extends the IBM Cluster 1350 portfolio by delivering the price/performance capabilities of the IBM POWER6 64-bit processor in a blade architecture. The JS22 Express delivers many new leading-edge technologies—IBM EnergyScale™ technology; IBM Advanced POWER Virtualization (APV) and Altivec™ SIMD acceleration functionality—in a single, high-performance, yet cost-efficient blade server with features designed for high reliability. The IBM Systems Director Active Energy Manager and browser-based Integrated Virtualization

Manager software make it easier than ever to achieve increased utilization and energy efficiency by utilizing the APV and EnergyScale capabilities of the system. This processor technology combined with IBM's innovative BladeCenter architecture offers clients the opportunity to reach new levels of peak computing capacity in a single rack.

The Cell Broadband Engine (Cell/B.E.), supported in the IBM Cluster 1350 bill of materials, is a blade system based on Cell Broadband Engine processors. Cell/B.E. technology is built on proven IBM Power Architecture® technology and often delivers more performance than conventional microprocessors while being more scalable and easier to program than exotic technologies like graphics processing units (GPU), digital signal processors (DSP) or field-programmable gate arrays (FPGA).

#### **Reliable, tested configurations**

Standard configurations of the IBM Cluster 1350 include compute nodes, at least one management node (one redundant management node for failover is optional), and up to 64 storage nodes for a maximum of 1,024 managed nodes. Clients who require configurations larger than a total of 1,026 nodes (1,024 managed nodes

and up to two management nodes) or components not included in standard Cluster 1350 configurations can use a special bid process to request support of these custom configurations. All of these larger configurations utilize standard 42U racks.

Smaller cluster environments may use the IBM Cluster 1350 25U racks, which allow clients to optimize the size and affordability of their cluster to meet specific application needs. For example, clients with database, business intelligence (BI) and general SMB (small and medium business) applications will find that these smaller racks enable extremely cost-effective solutions for the smaller cluster configurations normally required in these environments.

IBM Cluster 1350 clients may choose from a broad variety of cluster interconnect technologies from several of the industry's leading network switch and adapter vendors to meet the specific performance needs of their cluster application environment. These choices span the full range of high-performance networking technologies, including Gigabit/10 Gigabit Ethernet, DDR InfiniBand®, and Myrinet switches and adapters.

### **Innovative power and cooling benefits**

The IBM energy management portfolio tackles the challenge to increase power and thermal efficiency and help reduce costs on many levels. First, inside the system, all IBM System x and IBM BladeCenter servers start with Calibrated Vektored Cooling™ technology. This feature allows dual paths of air to each component, improving uptime and longevity and reducing wasteful air movement and heat generation. Coupled with more energy-efficient power supplies, IBM BladeCenter and System x servers can generate less heat in the critical AC-to-DC power conversion than many alternative systems from the competition.

For clusters within a rack, IBM System x servers are designed to work at full density in a well-planned rack solution. They are also designed to operate at extended temperature ranges to keep the system up and running—even in extreme temperature and failure conditions. IBM rack-based cluster solutions are engineered to optimize air flow and prevent undesirable recirculation within the rack, so servers can run in optimal temperature conditions.

IBM BladeCenter-based clusters help enable you to pack more processors into the same power and cooling envelope as well as better utilize floor space and “right size” data-center design. With IBM BladeCenter, less power per processor means more processing capacity per kilowatt. The BladeCenter runs cooler to deliver greater reliability.

For dense data center environments, IBM provides smart rack-level heat solutions like the super-efficient IBM Rear Door Heat eXchanger. The water-cooled door is designed to dissipate heat generated from the back of the rack to reduce the overall room temperature.

With this combination of benefits at the server and data center level, IBM Systems can provide strong power and cooling benefits to Cluster 1350 clients.

### **Summary**

Creating a computing infrastructure is an exercise in balancing price and performance to deliver the appropriate solution for each client's specific business needs.

The IBM System Cluster 1350 is a comprehensive solution that can help simplify and expedite deployment of a Linux or Windows HPC cluster.

IBM combines all hardware, software, services and support into a single integrated product offering, providing clients the benefit of a single point-of-contact for the entire cluster rather than dealing with multiple vendors for individual components.

IBM's Cluster 1350 is an outstanding choice for any organization that recognizes the economic advantages of a reduced time to deployment of an HPC cluster but has concerns about the time and technical resources required for the end-to-end implementation.

### **For more information**

To learn more about the IBM System Cluster 1350, contact your IBM representative or IBM Business Partner, or visit the following Web sites:

- [ibm.com/systems/clusters](http://ibm.com/systems/clusters)
- [ibm.com/systems/clusters/hardware/factsfeatures.html](http://ibm.com/systems/clusters/hardware/factsfeatures.html)
- [ibm.com/common/ssi](http://ibm.com/common/ssi)
- [ibm.com/systems/x/storage](http://ibm.com/systems/x/storage)
- [www.redhat.com/software/rhel](http://www.redhat.com/software/rhel)
- [www.suse.com/us/business/index.html](http://www.suse.com/us/business/index.html)
- [www.microsoft.com/windowsserver2003/ccs/hpcplus.aspx](http://www.microsoft.com/windowsserver2003/ccs/hpcplus.aspx)

## IBM System Cluster 1350 Summary at a glance

<b>Systems</b>	Blade servers: HS21, HS21-XM, LS21, LS41, JS22, QS21 Rack servers: x3455, x3550, x3650, x3655, x3755, x3450 (July 08)	
<b>Networking</b>	Ethernet	Blade Network Technologies, Cisco, Force10, Nortel, SMC
	InfiniBand	Cisco, Voltaire, QLogic
	Myrinet	Myricom
<b>Accelerator</b>	Clearspeed	PCI-X
<b>External storage</b>	Storage servers System Storage DS4800, DS4700, DS4200, DS3400, DS3200, EXP3000 Storage expansion EXP810 Storage Expansion Unit SAN switch 2005-16B IBM System Storage SAN 16B-2 SAN Switch	
<b>Software</b>	Operating system <ul style="list-style-type: none"><li>• Red Hat Enterprise Linux (RHEL) 5</li><li>• SUSE Linux Enterprise Server (SLES) 10</li><li>• Microsoft® Windows Compute Cluster Server 2003</li><li>• Microsoft Windows HPC Server 2008 (planned October, 2008)</li></ul> Cluster management software <ul style="list-style-type: none"><li>• IBM Cluster Systems Management (CSM) for Linux V1.5.1 (optional)</li><li>• IBM General Parallel File System (GPFS) for Linux V3.2 (optional)</li></ul>	
<b>Rack dimensions</b>	42U primary/expansion rack 79.5" H x 25.2" W x 43.3" D (2020 mm x 640 mm x 1100 mm); 574.2 lbs (261 kg) <sup>1</sup> 25U rack 49.0" H x 23.8" W x 39.4" D (1344 mm x 605 mm x 1001 mm); 221 lbs (100.2 kg) <sup>1</sup> Power and Cooling Calibrated Vectored Cooling, Energy-efficient power supplies, low-voltage processors, IBM Systems Director Active Energy Manager, IBM Power Configurator, Thermal Diagnostics	
<b>Scalability</b>	One management node is required and one redundant management node for failover is optional. A minimum of two and a maximum of 1,024 managed nodes are supported (optional software may limit supported nodes). This total can include up to 64 storage nodes. The maximum configuration is 1,026 nodes including compute, storage and management nodes. Larger configurations are available through a special bid process.	
<b>Services</b>	<ul style="list-style-type: none"><li>• Cluster 1350 hardware installation is included at no charge on 42U and 25U racks.</li><li>• HPC cluster software services and SupportLine for Linux and Windows Clusters are available as optional fee-based services.</li><li>• Cluster Installation Support Services are available through the Cluster Enablement Team (CET) as optional fee-based services.</li></ul>	
<b>Warranty</b>	3-year parts, customer replaceable unit (CRU) or on-site labor, limited warranty, with individual nodes retaining the warranty and service upgrade offerings for that IBM Machine Type; optional warranty service upgrades <sup>2</sup>	



InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

Other company, product and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply. For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven® or ClusterProven®.

Photographs show engineering and design models. Changes may be incorporated in production models.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products, or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, to evaluate the performance of a system they are considering buying.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

MB, GB, and TB = 1,000,000, 1,000,000,000 and 1,000,000,000,000 bytes, respectively, when referring to storage capacity. Accessible capacity is less; up to 3 GB is used in service partition. Actual storage capacity will vary based upon many factors and may be less than stated. Some numbers given for storage capacities give capacity in native mode followed by capacity using data compression technology. Maximum internal hard disk and memory capacities may require the replacement of any standard hard drives and/or memory and the population of all hard disk bays and memory slots with the largest currently supported drives available.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

<sup>1</sup> Weight is dependent on the nodes added and will vary when disks, adapters and peripherals are added.

<sup>2</sup> Standard 1-year warranty applies for x3455 and QS21, individual nodes retain the warranty and service upgrade offerings for that IBM Machine Type. Enhanced warranty service plans are available.

<sup>3</sup> [www.top500.org](http://www.top500.org)

© Copyright IBM Corporation 2008

IBM Corporation  
Integrated Marketing Communications,  
Systems and Technology Group  
Route 100  
Somers, NY 10589

Produced in the United States of America  
June 2008  
All Rights Reserved

This publication was developed for products and/or services offered in the United States. IBM may not offer the products, features or services discussed in this publication in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, BladeCenter, Calibrated Vectored Cooling, ClusterProven, EnergyScale, General Parallel File System, GPFS, POWER, POWER6, Power Architecture, ServerProven, System Storage, System p, System x and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both. A full list of U.S. trademarks owned by IBM may be found at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

UNIX is a registered trademark of The Open Group in the United States, other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries or both.

Intel and Xeon are registered trademarks of Intel Corporation in the United States, other countries or both.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc. in the United States, other countries or both.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries or both.