

An innovative solution designed to deliver affordable high-throughput computing



## IBM BladeCenter solution for Bioinformatics



demands for emerging 64-bit computing capabilities. In this constantly changing environment, you need to get the most out of your IT investment, while managing a growing cluster of servers and data and maintaining a high level of service to users.

The IBM BladeCenter® solution for Bioinformatics™ works to deliver affordable, high-throughput and competitive application price and performance—and was designed by IBM Life Sciences to meet the demands of bioinformatics research organizations.

---

### Highlights

---

- *High-throughput computing solution delivers more power for the money*
- *Popular bioinformatics applications with enhanced performance help accelerate research results*
- *Designed for greater ease of deployment*

Until now, maximizing system throughput for an unpredictable, innovative bioinformatics environment may have strained your budget. Selecting and testing a diverse set of public and commercial applications to ensure that they produce reliable results, as well as reprioritizing workloads at a moment's notice, are common activities. Or you may be faced with satisfying user

### An affordable way to manage high-throughput workloads

Most research organizations care about getting the most value for their IT investment. And most applications run “well-enough” in clusters designed for maximum throughput at minimum cost. Cost-optimized computing is a viable option—leveraging the Linux® operating system and commonly used open source applications. The BladeCenter

### Visit

---

**ibm.com** to locate an IBM reseller or for more information.

solution for Bioinformatics can help reduce your overall cost of computing. Its PowerPC 970 processor-based blade servers are cost-effective, less power-hungry and easier to set up than traditional server alternatives. The Linux operating system, as well as open source applications and tools, can also help minimize cost. Services are available for solution implementation and integration within your environment. The resulting solution's aggregate compute power, flexibility and performance can easily compare with higher-priced options.

The IBM BladeCenter technology that enables this solution offers ease of installation, maintenance and upgrade-ability to help reduce operating costs. In addition, its plug-and-play integrated infrastructure makes it easy to add capacity and to consolidate sprawling server farms, storage and networking using just a few "superdense" racks. Overall, the affordable BladeCenter solution for Bioinformatics enables both researchers and IT to do more for less.

### **64-bit capability to solve a wide variety of problems**

The IBM BladeCenter JS21 uses the PowerPC 970 chips which are derived from IBM's 64-bit POWER4™ processor family. A single blade can contain up to 4 cores, each of which has 64-bit data paths and memory addressing, yet it is natively compatible with 32-bit PowerPC software. And it's the first IBM PowerPC chip with AltiVec™ extensions—the single-instruction, multiple-data (SIMD) operations that can significantly accelerate data-intensive processing tasks. The PowerPC 970 is an affordable package which preserves the most impressive traits of the POWER4 processor family. So powerful is the BladeCenter JS21 that it runs up to three times faster than the JS20 server and also is the first BladeCenter to have built-in Advanced POWER Virtualization and the Virtual I/O server.

64-bit processors in the PowerPC 970 enable users to execute applications that utilize memory mapping for data or large table indexing for performance tasks. 16GB of memory per server of



the latest DDR II memory is now affordable, enabling researchers to pursue larger problems.

### **Tailored for bioinformatics**

The IBM BladeCenter solution for Bioinformatics offers a recommended configuration for a typical bioinformatics high-throughput environment. Familiar sequence alignment and analysis applications including BLAST, HMMER, FASTA and ClustalW have been ported and tuned on Linux to leverage the PowerPC 970 processors. Users of HMMER, for example, may find that the PowerPC 970 platform delivers more compute power for the price than most other computing alternatives.



The BladeCenter solution for Bioinformatics includes the GCC C/C++ compiler enabled for POWER4 and SIMD, plus IBM's C/C++ and FORTRAN compilers. Popular open source middleware and development tools tested include:

- *Atlas*
- *BioPERL*
- *BioJAVA*
- *POVRAY*
- *NFS*
- *MASS*
- *MPICH*

For your convenience, the IBM BladeCenter solution for Bioinformatics arrives with instructions for down-loading and installing each open source application or tool.

### **Superior scalability, flexibility and advanced integration**

The BladeCenter solution for Bioinformatics integrates high-performance components in an innovative infrastructure, helping to reduce overall IT costs and the system footprint. The IBM BladeCenter design enables you to combine up to 14 server blades in a single 7U chassis. This allows up to 84 servers in a standard 42U rack. Optionally, BladeCenter JS21s can be installed in a more powerful IBM BladeCenter H giving twice the channels to every blade and an increase of bandwidth by a factor of 10 in a 9U chassis.

The modular design of the BladeCenter solution for Bioinformatics is ideal for high-throughput workloads. It adapts to handle mixed job streams and application workloads. You can initiate jobs on servers when they become available, avoiding the constraints associated with dedicating jobs to specific servers. This same modular design enables you to easily scale up processors, storage and networking capacity.

As a result, you have the flexibility to incrementally expand your high-throughput computing capacity as the modular design enables Intel® Xeon® processor-based (one-core, two-core, four-core) blade servers, AMD Opteron processor-based (one-core, two-core) blade servers and PowerPC 970-based (two-core, four-core) blade servers to coexist in the same chassis. This mixed blade capability helps create workload and configuration flexibility. It even supports multiple operating systems.

Applications can be mapped to the hardware architecture that works best, optimizing the overall system. For example, your 32-bit applications can continue to run on Intel Xeon processor-based blades, while you add PowerPC 970 processor-based blades for applications that perform more cost-effectively on them. This means you can have the flexibility to choose processors and the operating system that workloads demand, while gaining the advantages of a single infrastructure to manage.

The same intelligent tool, IBM Director, manages the entire system, helping simplify and automate IT tasks. With

IBM Director, you deploy, configure, manage and maintain dozens or even hundreds of blade servers. Expanding capacity may be completed in just a few minutes, rather than hours.

What's more, you can pre-load the IBM BladeCenter chassis with extra blades to easily add computing capacity as needed. And you pay for the standby blades only when they are activated. These capabilities make it easier for you to achieve greater density, consolidating your laboratory workloads into one infrastructure.

Whether your interest includes, academic research, Biotechnology, genomics, pharmaceuticals or Clinical care, the IBM BladeCenter solution for

Bioinformatics can demonstrate immediate value in your environment. Participants in this solution include IBM Business Partners, who offer applications, middleware, hardware and services designed specifically for life sciences organizations, and the open source community, to which IBM has a sustained and growing commitment.

#### **For more information**

To learn more about IBM Healthcare Life Science solutions, visit **ibm.com/industries/lifesciences** or contact an IBM Healthcare and Life Sciences specialist at [LS@us.ibm.com](mailto:LS@us.ibm.com).



© Copyright IBM Corporation 2006

IBM Corporation  
Route 100 Somers, NY 10589  
U.S.A.

Printed in the United States of America  
March 2006  
All Rights Reserved.

IBM, the IBM logo, BladeCenter, POWER, POWER4 and PowerPC are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

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

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

AltiVec is a trademark of Motorola, Inc. used under license.

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

IBM reserves the right to alter specifications and other product information without prior notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed.