An innovative solution designed to deliver affordable high-throughput computing

IBM @server **BladeCenter Solution for Bioinformatics**

Introducing the IBM @server® BladeCenter Solution for Bioinformatics™. This complete solution works to deliver affordable, high-throughput performance and competitive application price/ 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 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.

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 Solution for Bioinformatics can help reduce your overall cost of computing. Its PowerPC 970-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 @server BladeCenter technology that enables this solution offers ease of installation, maintenance and upgradeability 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 two-way, 1.6 Ghz PowerPC 970 is derived from IBM's 64-bit Power4™

processor. It has 64-bit data paths and memory addressing, yet it is natively compatible with 32-bit PowerPC software. It has tremendous bus bandwidth, pushing effective bus rates to 800 Mhz. And it's the first IBM PowerPC chip with AltiVec™ extensions— the single-instruction, multiple-data (SIMD) operations that can accelerate data-intensive processing tasks. The PowerPC 970 preserves the most impressive traits of the Power4 in an affordable package.

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. Four GB RAM per node is now affordable, enabling researchers to pursue ever larger problems.

Tailored for bioinformatics

The IBM @server 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, BLAT 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*	• *NFS*
• *BioPERL*	• *MASS*
• *BioJAVA*	• *MPICH*
• *POVRAY*	

For your convenience, the IBM @server 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 @server BladeCenter's design enables you to combine up to 14 server blades in a single 7U chassis. This allows up to 84 two-way servers in a standard 42U rack. In fact, over a teraflop of compute power can be packed into a single IBM @server BladeCenter rack.

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 needed — deploying that new capacity faster, maintaining it easier and, most likely, doing it all at a lower total cost.

In addition to the 64-bit PowerPC 970, the IBM @server Blade Center can also host 32-bit Intel™ servers.

The modular design enables Intel Xeon-based (two-way, four-way) and PowerPC 970-based (two-way) 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-based blades, while you add PowerPC 970-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 @server 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.

Empowering life sciences research
The IBM @server BladeCenter Solution for Bioinformatics can demonstrate immediate value in your environment. It was produced by IBM specialists with considerable domain knowledge and technology expertise. 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.

IBM ®

GH09-4717-00