



About this research note:

Technology Insight notes describe emerging technologies, tools, or processes as well as analyze the tactical and strategic impact they will have on the enterprise.

Windows HPC Server 2008 Elevates High Performance Computing

Publish Date: October 2, 2008

High performance computing (HPC) has been primarily used by academic and research facilities. With Microsoft's introduction of its Windows HPC Server 2008, some of the limitations of prior generations of HPC solutions are being removed. HPC is becoming accessible to businesses of all sizes.

INFO~TECH
research group

Passionate About Research.
Driven By Results.

www.infotech.com

© 1998-2008 Info-Tech Research Group



Executive Summary

Microsoft's introduction of HPC Server 2008 offers additional capabilities that will allow more enterprises to create high performance computing (HPC) environments. This introduction focuses much more on integration into existing business data centers and much less on traditional HPC environments.

The enhancements for this version of Microsoft's HPC product include:

- » Integration with existing infrastructure.
- » Simplified cluster management.
- » Additional support for multi-threaded applications.
- » Improved job scheduling and application submission.

Overall, the biggest changes are focused on management and ease of use by end users of all types. Those changes, combined with additional application and ISV support, make Microsoft HPC Server 2008 worth considering for Microsoft-centric data centers.



Technology Point

Microsoft has replaced its previous high performance computing (HPC) product, Microsoft Compute Cluster (MCC) Server 2003, with [Microsoft HPC Server 2008](#). Like the previous version (refer to the Info-Tech Advisor research note "[Windows for High Performance Computing Offers Price Advantages](#)") the package is built on the current Windows Server platform.

HPC Server 2008, like its predecessor, allows groups of computers to be linked into a shared computing cluster. One server, the head node, acts as a control point for distributing jobs between available servers and coordinating resources and results.

HPC environments have been available for government, education, research, weather forecasting, oil exploration, and so on for some time. These groups created software specifically designed to meet the needs of their industry and compute requirements. What has changed is the need for this level of processing power to be available to enterprises not focused on these traditional HPC environments.

The growth of social networks, gaming, data mining, and Web searches has begun to bring the need for massive compute power to more and more enterprises. Along with that need has come more emphasis on multi-threaded applications, designed to operate on multi-core processors and multi-processor servers. The programming skills required to create this new generation of business applications can be leveraged to build applications that can now utilize HPC clusters.

While the majority of HPC environments are composed of Linux based servers utilizing Beowulf clustering, (as discussed in the Info-Tech Advisor research note, "[Beowulf: An Epic Adventure in Clustering](#)") Microsoft offers a viable (if not free) option. By allowing HPC servers to share existing infrastructure within the data center, enterprises with heavy Microsoft Server investments can now leverage existing servers and connectivity to build HPC environments from a few, to hundreds, of servers.

What It Is & How It Works

HPC offers the ability for multiple servers to process information in parallel. In a true HPC environment the applications need to be designed to allow for multi-threading and parallel operations. This technology is often referred to as a cluster because of the use of multiple processors all working together. HPC does not necessarily provide high availability but it does provide the processing capabilities needed for compute intensive applications.

In the Microsoft (and other) HPC environment one server is designated as the head node. This unit acts as a coordination point for distributing work to the other servers in the HPC cluster. In most cases, any server that is to be used as part of an HPC environment will not be running any applications other than those assigned to it by the head node.



The head node is the primary contact point for both job scheduling from users, as well as network and storage connectivity to/from the processing nodes in the HPC cluster. Depending on the specific data access requirements servers in the HPC cluster can be networked with Ethernet (1Gb or 10Gb) or Infiniband. For integration into standard data center operations, most enterprises will find the Ethernet solution both adequate and inexpensive. Only the most high end requirements would dictate the use of Infiniband.

Applications requiring the capabilities of HPC do need to be designed to be able to utilize the multi-server environment. This means that the software must support multiple independent threads of code. If an application requires sequential operations (calculation A feeds data to calculation B which feeds data to calculation C and so on) it may not be a good candidate for parallel and multi-threaded operations. If the same set of calculations will be run against different sets of data with little or no requirement for interaction between those calculations, then a HPC environment may provide the high performance mechanism needed.

Key Considerations

Microsoft has added some significant improvements to this latest release of its HPC environment. When taken as a whole it shows the migration of standard data center operations toward a more distributed and coordinated structure. Some of the key changes and improvements found in HPC Server 2008 are focused on allowing existing data centers to create HPC environments without significant changes to the operations.

Integration with Existing Infrastructure

Previous implementations of the Microsoft HPC environment highly recommended using a separate physical infrastructure for connecting the processing nodes. This eliminated the potential for overwhelming the enterprise network with HPC data flows and simplified the overall management of the cluster.

With this release both existing servers and existing networking infrastructure can be used to create a HPC cluster. For many business applications built around multi-threading and parallel processing mechanisms, the amount of data that is required by individual servers or passed between nodes should not be enough to degrade operations for other servers.

For those situations where data flow does begin to impact normal network operations the creation of a new physical network can be eased by using existing Ethernet technology for the interconnections. This physically separate infrastructure can continue to be managed using the same tools and skill sets as is required by the rest of the enterprise infrastructure.



Simplified Cluster Management

Microsoft Server 2008 and HPC Server 2008 have added management tools that are specifically aimed at making the management of clusters (HPC or high availability) easier. These include the ability to define which servers in the network belong to a cluster and what application can be allocated to those nodes. In addition, all servers in all cluster configurations can now be monitored with Microsoft System Center Operations Manager. This allows one view of the complete network, including those servers that are allocated to the HPC environment.

Being able to assign policies to both applications and server nodes ensures that appropriate resources will be available for the required processing. Nodes can be added and removed from the cluster and the current status of individual servers and whole clusters can be monitored. For Microsoft centric data centers this single-pane-of-glass environment will make implementation and management of HPC clusters significantly easier since existing administrator skill sets can now be extended to these additional servers.

Additional Support for Multi-Threaded Applications

Emphasis on multi-threaded applications is not directly related to the HPC environments. The rapidly growing use of multi-core processors has made the use of parallel and multi-threaded techniques for software development more main-stream. Segmenting applications into separately executable threads for multi-core and multi-processor hardware translates almost directly into the mechanisms required for HPC applications.

Microsoft provides tools for creating [parallel processes in Microsoft .Net](#) as well as general forums for parallel computing development. The latest versions of Microsoft Excel have built-in capabilities to evaluate the processing requirements of the internal calculations and to automatically allocate those to different cores, processors, or servers.

In addition, Microsoft's Visual Studio and C Sharp have added functionality to allow parallel and multi-threaded operations. These, along with support for Message Passing Interface (MPI), allow the creation and integration of HPC applications that can communicate and work with other operating system environments. Combined with Microsoft's push for ISV created applications, enterprises should soon find HPC capable business software more readily available and, when necessary, easier to create. Financial and business intelligence applications will have much of the initial development focus.

Improved Job Scheduling and Application Submission

The standard management package included with the HPC Server 2008 provides tools for defining policies that can be applied to specific application types. This can include how many simultaneous threads can be run or even how many (and which) cores to use on specific servers in the HPC cluster.



Management tools that monitor and report on the operations by server and by job allow the automation of scheduling and distribution of applications based on processing needs, priority and availability of infrastructure. Additionally, applications can be deployed so that the user has no knowledge of the HPC back-end.

Simply executing the application, such as an Excel based calculation, will cause the head end node to:

- » Evaluate the user (is this user/system authenticated for HPC operations?)
- » Determine submitted application requirements
- » Monitor available resources
- » Distribute application threads to available nodes
- » Monitor application performance (has it reached/exceeded time constraints or resource limits)
- » Report completion to the user application

While the end user may not need to be aware of the HPC operations, the application itself must be modified to allow for the multi-threaded and parallel operations. This is a step in the right direction but is still not the end-game of HPC development.

Key Takeaways

As enterprises move to utilize more computing resources in more powerful environments HPC will become a more common method for all enterprises. However, there are things to consider before making the move to HPC.

1. **Applications are still focused on scientific and research areas.** Even with the growing development of multi-threaded applications for standard business environments the majority of HPC development is still focused on traditional high end processing needs. While there is a push by Microsoft for ISV development of more common business applications, it will take a while to produce the software needed by most enterprises.
2. **In-house development may need new expertise and training.** Any enterprise that decides to begin developing or modifying applications for HPC (or even multi-core) environments can, and should, consider the additional training that will be required. This includes not just programming methodologies, but also the ability to define which components of an application can be migrated to semi-autonomous and parallel operations. This skill set is becoming more common with the rapid introduction of multi-threaded development for multi-core processors, but for many enterprises in-house developers will not yet have had to deal with this requirement.



3. **Not all applications will benefit from HPC.** HPC is not for all applications. The reason that government, education, research, weather forecasting, oil exploration and heavy duty simulation laboratories have led in this area is that they have needed the massive processing power HPC delivers. Many common applications, such as database operations, use symmetrical multiprocessing or load-balancing techniques to provide similar high performance results. Before expending any resources on testing HPC infrastructure and operating environments the real processing needs of business applications should be evaluated. Maybe the addition of load-balancing or more powerful servers will meet the need without changing the code. Quick and cheap can sometimes provide better solutions than the more complex development needs of HPC.

Bottom Line

HPC has been primarily used by academic and research facilities. With Microsoft's introduction of its Windows HPC Server 2008, some of the limitations of prior generations of HPC solutions are being removed. HPC is becoming accessible to businesses of all sizes.

Info-Tech's products and services combine actionable insight and relevant advice with ready-to-use tools and templates that cover the full spectrum of IT concerns. Our practical approach is designed to have a clear and measurable positive impact on your organization's bottom line.

We serve over 21,000 clients at 8,000 organizations around the world. Since 1998, we have focused on making the work of IT professionals easier - and on helping them achieve greater personal and corporate success.

[More About Info-Tech](#)