


LY PRODUCTIVE HIGH PERFORMANCE COMPUTING



OVERVIEW

Windows HPC Server 2008 enables computing power to be harnessed, easily and productively, for various high performance computing (HPC) applications.

Because these applications can be very data intensive, it is critical to complement your Windows HPC Server 2008 cluster with an appropriate, highly scalable storage solution.

CLUSTER STORAGE SOLUTIONS THAT FIT YOUR HPC NEEDS

Windows HPC Server 2008 Storage Requirements

Windows® HPC Server 2008 clusters vary in size and throughput requirements; their storage requirements vary as well. Fortunately, there are many storage options available, in addition to several options for managing the storage.

There are many factors to consider when selecting a storage solution for your Windows HPC Server 2008-based cluster. These include:

PERFORMANCE: Traditional supercomputing applications typically consume data upwards of 100 MB/second per core (but with limited scalability) or in the range of 50 MB/second per core (scalable to hundreds of nodes). To gauge your performance needs, review the size of your cluster and the number of copies of applications that may run concurrently. For example, a cluster with a scalable application of 100 quad-core nodes and a data transfer rate of 50 MB/second requires an aggregate throughput of 20 GB/second.

CAPACITY: Once you define the performance requirement, you need to examine the length of time that the data will reside online and the amount of temporary storage you need. For example, if the cluster can generate a throughput of 20 GB/second (more than a terabyte/minute) and if the data needs to be online for five hours, then you need a storage capacity of at least 300 terabytes. Note that most storage vendors will provide you with the capacity characteristics of their solutions.

ACCESSIBILITY: Another consideration is access to storage. Accessibility needs depend on how the application runs in the cluster. In some cases, not all compute nodes need access to all of the storage; if an application can run anywhere in the cluster, however, then all compute nodes need access to all of the storage.

SINGLE NAMESPACE: Some clusters require a single or virtual namespace, where all of the nodes can see all of the files at the same time. Parallel file systems provide a single namespace, while network attached storage (NAS) devices (such as Windows® Storage Server, Windows® Unified Data Storage Server, or devices from third-party vendors) share storage area networks (SANs) behind multiple NAS heads. With Windows Storage Server, a single namespace is presented across multiple Windows Storage Server-based servers; this is called Distributed File System Namespace (DFS-N). DFS-N provides easy management for multiple storage servers, load sharing, and automatic failover capabilities.

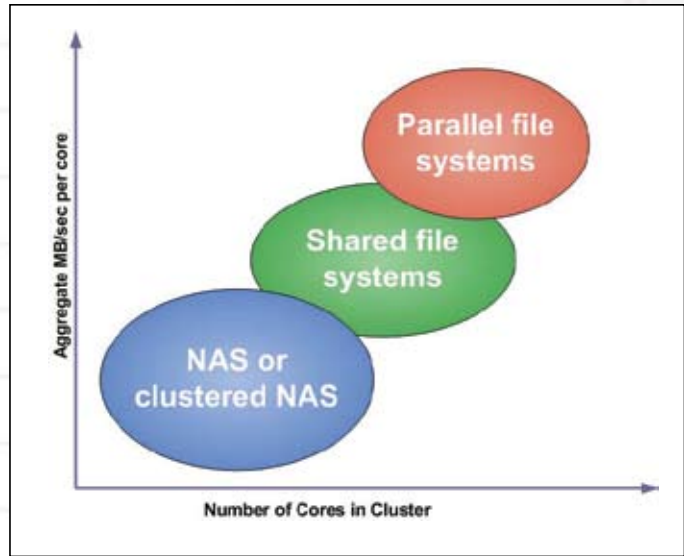
AVAILABILITY: You should also consider the availability of the storage. Availability demands can be met at the server level (high-availability NAS heads), at the storage level (just a bunch of disks [JBOD] or redundant array of independent disks [RAID]), at the storage controller level (failover of logical unit numbers [LUNs]), or at the path level (multipath failover and redundant switches). Many options exist for keeping storage available to the cluster.

Evaluating your needs ensures that the storage architecture you deploy complements your Windows HPC Server 2008 cluster, helps maximize your resources, and provides you with optimal data storage.



WINDOWS HPC SERVER 2008 STORAGE OPTIONS

There are a wealth of storage options available for your Windows HPC Server 2008 cluster from Microsoft and from trusted Microsoft partners. Storage options include NAS and clustered NAS, SAN shared file systems, and, for the highest in aggregate performance, parallel file systems.

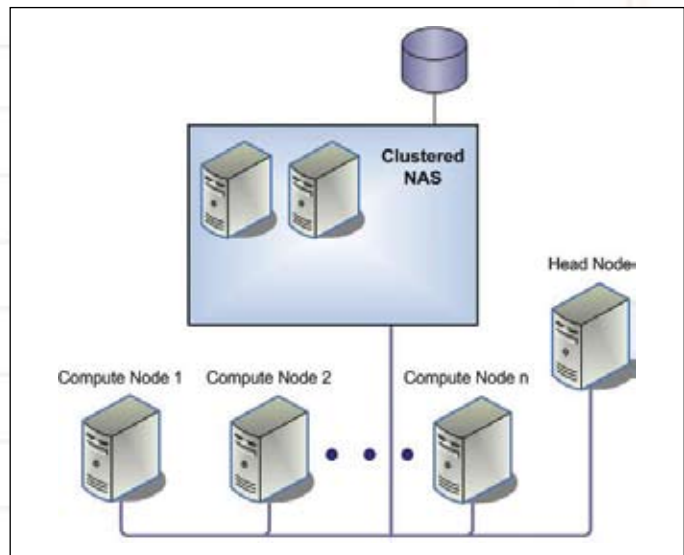


WINDOWS HPC SERVER 2008 STORAGE SOLUTIONS

NAS and Clustered File Systems

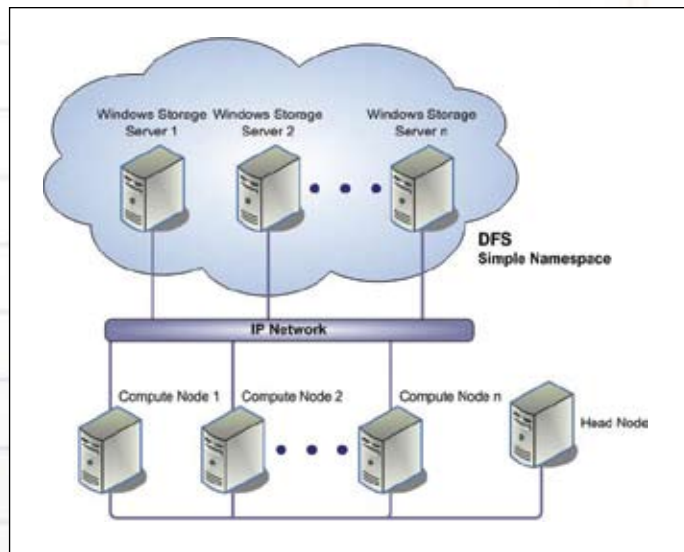
NAS hardware is similar to a traditional file server equipped with direct attached storage (DAS), but the operating system on the NAS server provides only data storage, data access, and the management for these functions. NAS uses file-based protocols such as Network File System (NFS) or Server Message Block (SMB).

There are a number of NAS options available, including single NAS devices, clustered NAS, and virtual filers. A single NAS device is the simplest solution available for your cluster, but it is also the most limited. Many vendors provide out-of-the-box single NAS solutions that can generally scale from one network interface card (NIC) to four or more NICs.



One NAS option is **Windows Storage Server**, which has the advantage of providing a full Windows® ecosystem for your Windows HPC Server 2008 cluster. With an all-Windows HPC stack, interoperability and management consistency are key benefits. Using SMB, rates greater than 100 MB/second per storage server per NIC have been reported on servers with internal RAID. High-end server configurations have demonstrated rates of 500 MB/second per server.

DFSN allows multiple network shares to be aggregated as a virtual file system, presenting the user with a single namespace. This solution can provide the storage capacity you need, while maintaining the interoperability advantages of an all-Windows system.




= HPC++

SAN or Shared or Clustered File Systems

If you need higher single-node performance, a shared file system, accessible to your clustered compute nodes over a SAN, may be appropriate. SAN file systems allow many computers to access the same file system (and even the same file) at the same time. SAN file systems use a faster block storage access method to the compute nodes than NAS networks use.

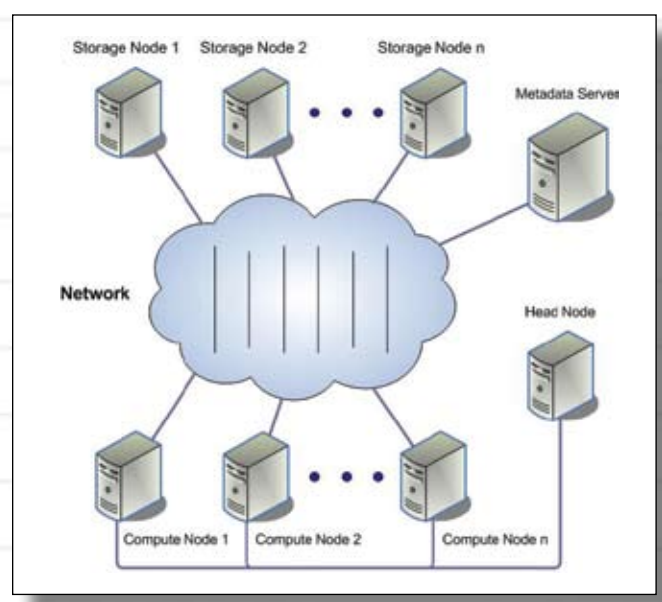
Shared file systems move the control path to the side and let the data path work directly from the compute nodes to the storage. Many use a metadata server (MDS) to arbitrate access to the storage and maintain the consistency of the file system.



Parallel or Clustered File Systems

With parallel file systems, many small compute nodes have their data served by many small storage nodes; the compute nodes can then access the storage nodes at the same time, in parallel. Parallel file systems can support a very large aggregate bandwidth for many thousands of compute nodes.

There are several parallel file systems commercially available that operate by moving some of the file systems intelligence to the storage nodes and accessing files based on object IDs; these parallel file systems are available for Linux now, and some will soon be available for Windows.



WINDOWS HPC SERVER 2008 STORAGE MANAGEMENT SOLUTIONS

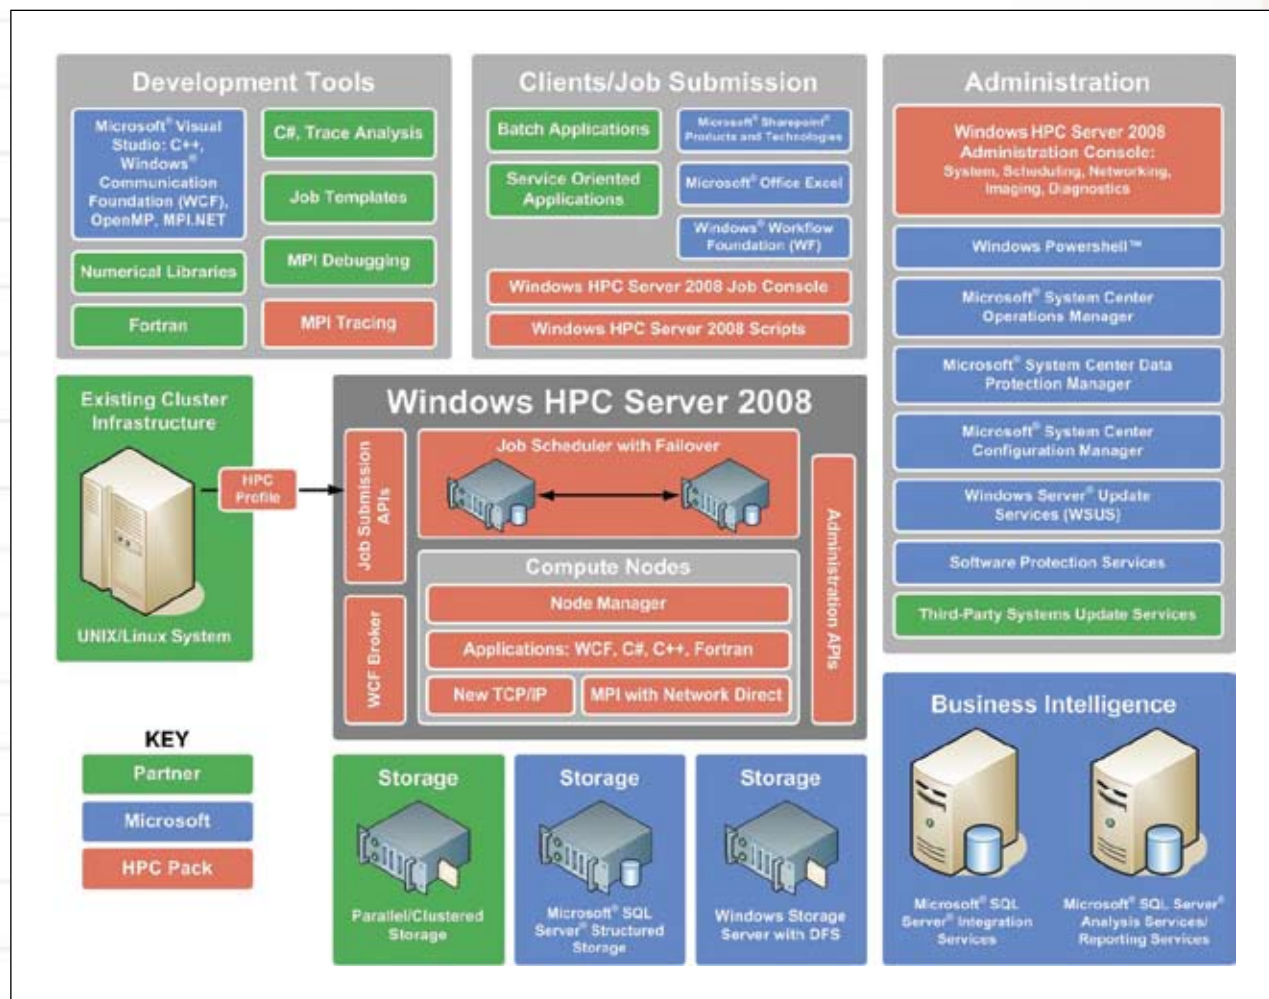
To manage your Windows HPC Server 2008 cluster's storage, you need failure monitoring, change management, performance monitoring, and storage provisioning capabilities. You can use familiar Microsoft tools, including Microsoft System Center Operations Manager 2007, Storage Manager for SANs (SMfS), File Server Resource Manager, Active Directory® Domain Services, and Microsoft® System Center Data Protection Manager 2007 to address these needs.

With the Windows HPC Server 2008 management pack, System Center Operations Manager provides end-to-end service management to more easily identify and resolve issues affecting the performance, health, and availability of Windows-based compute clusters. SMfS offers basic administration of SANs, including device discovery, provisioning, allocation to servers, and management of Microsoft multipath IO. File Server Resource Manager provides quota management, file screening, and storage usage reporting. There are also management tools available from trusted Microsoft partners.



WINDOWS HPC SERVER 2008 STORAGE SOLUTIONS

Windows HPC Server 2008 takes full advantage of the overall Microsoft ecosystem to improve the functionality, usability, and manageability of the cluster.



Windows HPC Server 2008 uses Microsoft® SQL Server® as the repository for the Job Scheduler, greatly improving the scalability and redundancy of the scheduler. Windows HPC Server 2008 also supports integration with the Microsoft® System Center solutions for enterprise management and configuration management. Windows HPC Server 2008 includes a custom management pack that provides advanced management and alerting capabilities through System Center Operations Manager 2007.

FOR MORE INFORMATION

Windows HPC Server 2008 Web site:

- www.microsoft.com/hpc

HPC community Web site:

- <http://windowshpc.net>

Windows Server system storage:

- www.microsoft.com/windowserversystem/storage/default.aspx