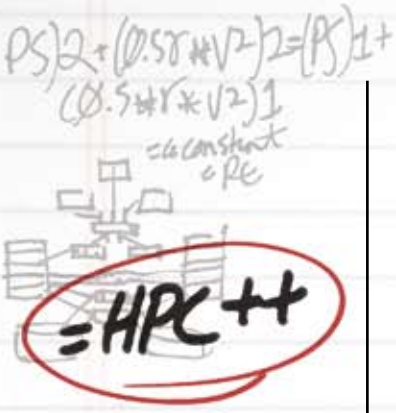

OVERVIEW

Innovation is the key to success in today's competitive environment. Many of these problems are inherently complex and demand vast amounts of computing power. HPC has become a fundamental enabler of innovation by providing dedicated compute resources to solve complex simulations and long-running calculations.

SITUATION

Historically, deploying, managing and troubleshooting HPC systems has been cumbersome, resulting in reduced cluster efficiency. Customers require a comprehensive platform that is easy to deploy, manage, and integrate with existing infrastructure. Meeting these requirements allows customers to accomplish more in less time, and thereby improve the productivity of their system administrators, application developers, and end users.

INTRODUCING WINDOWS® HPC SERVER 2008

Windows HPC Server 2008 provides a productive, cost-effective and high-performance computing (HPC) solution that runs on x64-bit hardware. Windows HPC Server 2008, can be deployed, managed and extended using familiar tools and technologies.

SOLUTION

Windows HPC Server 2008 combines the power of a Windows 64-bit Server platform with rich, out-of-the-box functionality to improve the productivity, and reduce the complexity, of your HPC environment. Windows HPC Server 2008, provides a comprehensive set of deployment, administration, and monitoring tools that are easy to deploy, manage, and integrate with your existing infrastructure.

Windows HPC Server 2008 enables broader adoption of HPC by providing a rich and integrated end-user experience scaling from the desktop application to the clusters. A wide range of software vendors, in various verticals, have designed their applications to work seamlessly with Windows HPC Server 2008 so that users can submit and monitor jobs from within familiar applications without having to learn new or complex user interfaces.

Developing parallel programs requires integrated development environments along with support for distributed computing standards. Visual Studio® 2008 provides a comprehensive parallel programming environment for Windows HPC Server 2008. In addition to supporting OpenMP, MPI, and Web Services, Windows HPC Server 2008 also supports third-party numerical library providers, performance optimizers, compilers, and a native parallel debugger for developing and troubleshooting parallel programs.

ENHANCED PRODUCTIVITY

Windows HPC Server 2008 allows you to accomplish more, in less time, with reduced effort, by leveraging users' existing skills and integrating with the tools already in use in your enterprise. Administrators and developers can save time with a common productivity and development platform across technical workstation and cluster. Administrators can use existing Microsoft management tools to centrally manage their entire Windows Server infrastructure. Command line interfaces are fully supported for administrators, and end users are not required to master command-line interfaces (CLIs) to harness the power of HPC. Application developers can work with familiar development tools, such as Visual Studio's native parallel debugger, to develop and troubleshoot parallel programs.

Teams can manage their projects and collaborate using workflows by using Microsoft Office SharePoint® Server 2007.

SCALABLE PERFORMANCE

Windows HPC Server 2008 is built on proven Windows Server® 2008 x64-bit technology. Windows HPC Server 2008 can efficiently scale to thousands of processing cores and includes management tools that help systems administrators proactively monitor system health and maintain system stability. Integration with Windows Server 2008 Enterprise and Microsoft® Windows SQL Server® provides failover capabilities in the event of system failure.

Microsoft

Windows HPC Server 2008

QUICKLY DEPLOY A MANAGEABLE INFRASTRUCTURE

Windows HPC Server 2008 includes wizards, tools, built-in management consoles, and a To Do List (Figure 1) to simplify the provisioning of compute nodes while enhancing Windows Server 2008 Deployment Services to allow you to create, modify, and deploy systems images across a cluster. New management tools and enhancements include the following:

NEW! The time and effort for setup and configuration are reduced through the use of system templates. Templates allow an IT professional to create standardized system images or apply patches that can be consistently deployed across the cluster.

ENHANCED! Operating system images can quickly be deployed using Windows Deployment Services (WDS). Progress can be monitored through the new HPC Pack Administration console.

ENHANCED! Manage clusters programmatically though command-line interfaces, PowerShell, or other scripting languages such as PERL.

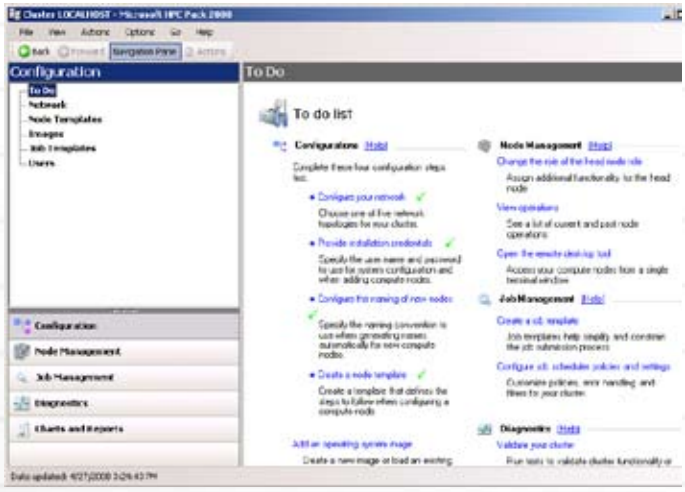

Figure 1: The To Do List simplifies cluster configuration tasks.

MONITORING, SYSTEM HEALTH, AND REPORTING

Built on Windows Server 2008 64-bit technology, Windows HPC Server 2008 includes scalable management tools built using the new System Center interface. The heat map (Figure 2) provides an overview of system utilization. New tools and enhancements for administrators include:

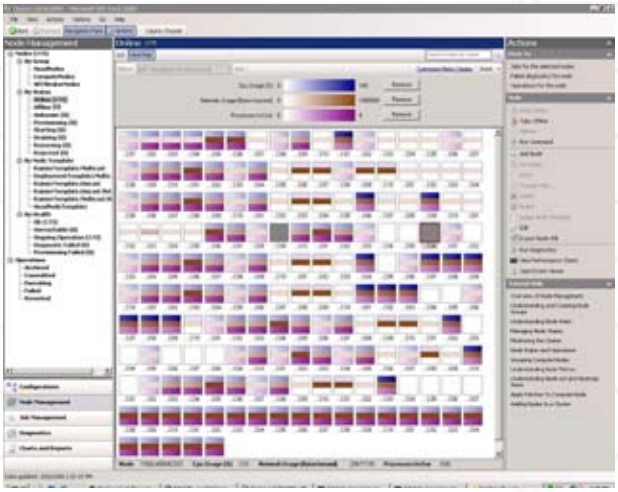

Figure 2: The heat map provides an at-a-glance view of system utilization.

NEW! Windows Server 2008 Enterprise includes Failover Services. The combination of Windows Failover Services and SQL Server database clustering provide head node redundancy in the event of a hardware failure.

NEW! Administrators can create groups based on hardware, software, or networking characteristics. Groups allow for simplified administration and allow optimal utilization by mapping jobs requirements to an appropriate set of servers.

NEW! Built-in diagnostic tools allow administrators to quickly identify and diagnose hardware, software, or network problems across the cluster. Additional capabilities such as: automating routine, redundant tasks, and intelligent reporting and monitoring are possible through the use of System Center Operations Manager 2007.

Microsoft

A FOUNDATION FOR HIGH-PERFORMANCE SERVICE ORIENTED APPLICATIONS (SOA)

The Microsoft HPC Pack 2008 includes a flexible job scheduler that includes a command-line, graphical user interface (GUI), and application programming interface for submitting jobs to the cluster. The 2008 job scheduler addresses both batch and newer service oriented applications. It can also be deployed together with partner job schedulers for additional advanced policies or mixed environments. New job scheduling tools and enhancements include:

NEW! Job scheduling supports new SOA workloads. Compute nodes can communicate with the submitting systems through Windows Communication Foundation (WCF) brokers. WCF brokers are dedicated nodes that scale out cluster performance, and act as proxies facilitating communication between public network clients and compute nodes on the private networks. Additional WCF brokers can be added to scale as required.

ENHANCED! The HPC Pack job scheduler is multi-core aware, allowing sophisticated scheduling by core, by socket, or by server. The job scheduler supports existing policies such as: priority-based first come, first served (FCFS), backfilling, non-exclusive and license-aware scheduling, and also includes new policies for job profiling, preemption, and growing and shrinking of jobs.

NEW! The HPC Pack reduces the complexity of integrating with existing clusters by supporting industry standards such as the Open Grid Forum's (OGF) HPC Profile for job scheduler interoperability. Additionally, the Subsystem for UNIX-based Applications (SUA) integrates applications to maximize existing investments while extending UNIX- and LINUX-based applications to Windows systems.

NETWORKING AND HIGH-SPEED INTERCONNECTS

Windows HPC Server 2008 significantly reduces the complexity of deploying multiple networks across a cluster by providing an updated networking wizard, new network diagnostics tools, and a simplified way of performing network driver management. New functionality and enhancements include:

ENHANCED! The Microsoft Message Passing Interface (MS MPI) is based on the Argonne National Labs implementation (MPICH2) of the MPI2 standard. MS MPI can utilize any interconnect that is supported on Windows Server 2008.

NEW! NetDirect is Microsoft's new Remote Direct Memory Access (RDMA) interface for high-speed, low-latency networks such as those running on Infiniband, 10 Gigabit Ethernet and Myrinet. With an architecture that directly bypasses OS and TCP/IP overhead, NetDirect achieves better performance for massively parallel programs that can utilize very low-latency, high-bandwidth, and enables efficient CPU utilization.

WINDOWS HPC SERVER 2008 ARCHITECTURE

The Windows HPC Server 2008 architecture is shown in Figure 3.

The Windows HPC Server 2008 head node:

- Controls and mediates all access to the cluster resources.
- Is the single point of management, deployment, and job scheduling for the cluster.
- Can failover to a backup head node in the case of failure.

Windows HPC Server 2008 uses the existing corporate infrastructure and Microsoft Active Directory® for:

- Security
- Account management
- Operations management using tools such as Systems Center Operations Manager 2007.

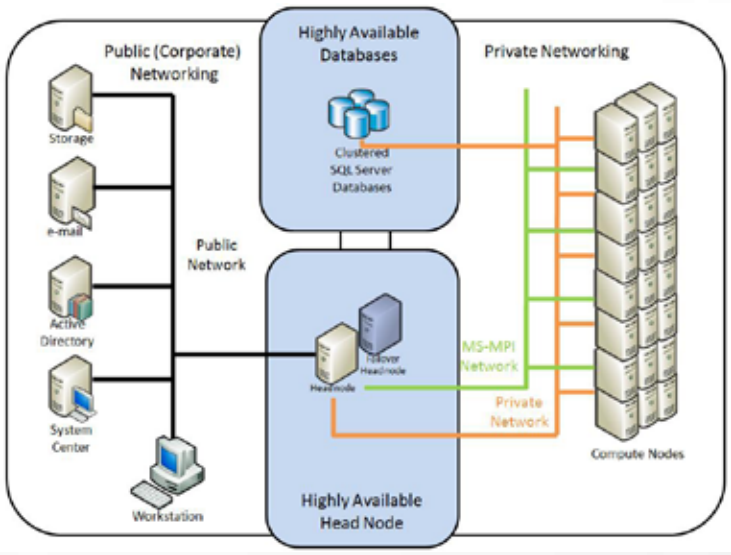

Figure 3: Windows HPC Server 2008 architecture.

Microsoft

INTERACTIVE SESSIONS THROUGH THE WCF

Figure 4 shows an interactive session through the WCF, which includes the following steps:

1. Clients create a session by specifying the WCF service used to perform the calculation.
2. The job scheduler assigns a broker and launches a service instance pool on multiple nodes.
3. Client connects to the broker and the job scheduler provides the End Point Reference (EPR) of the broker to the client.
4. Client connects to the EPR.
5. Standard WCF request/response messages occur.

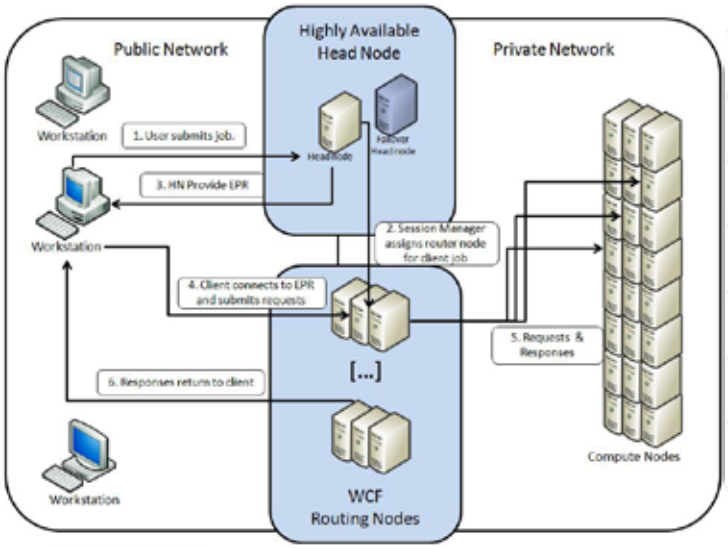

WINDOWS HPC SERVER 2008 SYSTEM REQUIREMENTS

The minimum system hardware requirements are similar to the hardware requirements for Windows Server 2008, Windows HPC Server 2008 supports up to 128 gigabytes (GB) of RAM. Supported processors include AMD Opteron, AMD Athlon 64, Intel Xeon with Intel EM64T, and Intel Pentium with Intel EM64T.

CPU Requirement	x64 architecture computer with Intel Pentium or Xeon family processors with Intel Extended Memory 64 Technology (EM64T) processor architecture; AMD Opteron family processors; AMD Athlon family processors; compatible processor(s).
Multiprocessor support	Windows Server 2008 HPC Edition and Windows Server 2008 Standard Edition support up to four processors per server. Windows Server 2008 Enterprise Edition supports up to eight processors per server.
Minimum RAM	512 MB
Maximum RAM	128 GB
Minimum Disk Space for Setup	50 GB
Disk Volumes	A single system volume is required for the head and compute nodes. Redundant array of independent disks (RAID) is supported but not required. The system volume must be MBR. Additional volumes can be MBR or GPT.
Network Interface Cards	At least one network interface card (NIC) is required. If a private network is used, the head node requires at least two NICs, and compute nodes at least one. Each node may also require a high-speed NIC for a MPI network.

FURTHER INFORMATION

For more information about Windows HPC Server 2008 and HPC please visit http://www.microsoft.com/hpc

Microsoft