The HP Unified Cluster Portfolio: A Flexible Approach to Applying Cluster Computing to HPC Applications

Executive Overview

High performance computing (HPC) has a never-ending demand for increased problem-solving capabilities. For many application mixes, clustering multiple industry-standard platforms together, and adding software to facilitate the use and maintenance of these systems, has become a viable lower-cost alternative to large-scale "supercomputing" systems.

HP has implemented a standards- and packaged-based clustering strategy for delivering products into the HPC market segments. This strategy allows HP to concentrate on key value-add components in the system infrastructure (e.g., platform chipsets), while maintaining a high degree of compatibility for application availability and lower manufacturing costs.

As an example of this, Figure 1 illustrates the wide variety of platforms that are available for clustering and cluster packages from HP. Users are able to select the platform that is the most suitable for their combination of price/performance, application and support requirements.

Figure 1. The wide variety of standards-based platforms available for clustering.

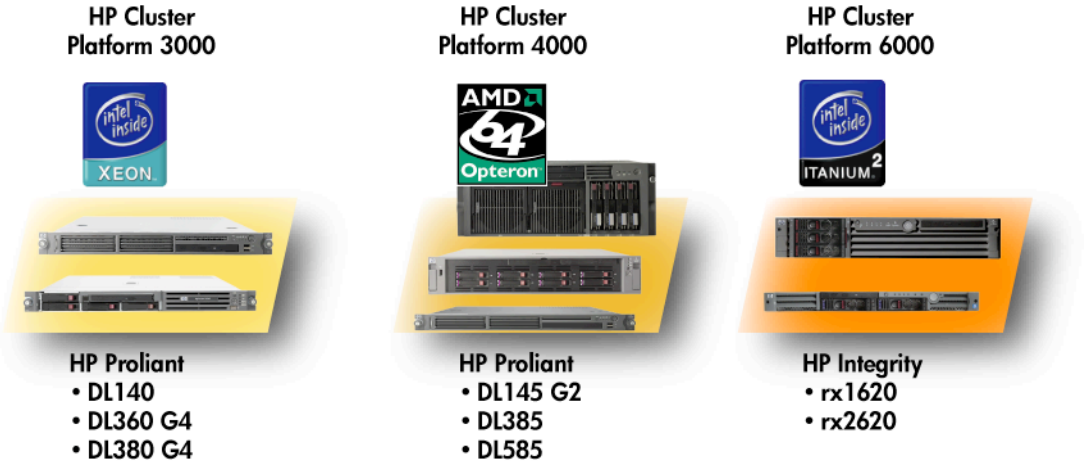

The HP Unified Cluster Portfolio provides a broad range of integrated and customer-ready solutions that enable rapid and confident deployment of HPC clusters. The portfolio effectively solves current clustering challenges by combining the flexibility of a custom solution with the simplicity, reliability and value of a pre-configured, factory-built product. A wide range of qualified options ensures flexible choices, simple implementation, and successful results. The HP Unified Cluster Portfolio features flexible, supported platforms; a wide range of open source and commercial middleware; and the latest in industry-standard technology.

Introduction

Clustering is a technology that leverages low-cost "commodity" platforms and interconnect products to provide customers with higher levels of performance and system availability. Trends in networking capabilities, hardware price/performance, and clustering software have moved much of the functions of IT organizations toward a more centralized model—a model that naturally lends itself to clustering. Vendors have responded to this movement to a centralized model by packaging and providing comprehensive integrated cluster solutions.

The Case for Clustering

Clustering technology has become the leading choice for compute-intensive applications, such as science and research, computer aided engineering, life and materials science, financial modeling, weather research, and geosciences. IDC, a leading industry analyst firm, believes that clusters will account for more than half of all technical computer sales by 2008[1]. Driving this growth are multiple factors, including:

- Increased demand for lower acquisition costs
- Demand for greater problem-solving capabilities
- Proliferation of low-cost, standards-based processors and system software
- Advances in usability and management software
- Advances in applications software

Clusters and HPC Applications Requirements

HPC applications demand a wide range of architectural solutions—from high computational requirements to large amounts of storage. Figure 2 illustrates various architectural elements required by different market segments.

Of critical importance to many HPC applications is the availability of storage and visualization elements that scale commensurate with the computing capabilities. Many scientific applications rely on a balance of all three capabilities —computation, storage, and visualization.

[1]"HP's Strategy for Delivering Cluster Technology to Technical Computing Environments", IDC, Christopher Willard, Ph.D., November 2004.

Figure 2. Architectural requirements in HPC market segments.

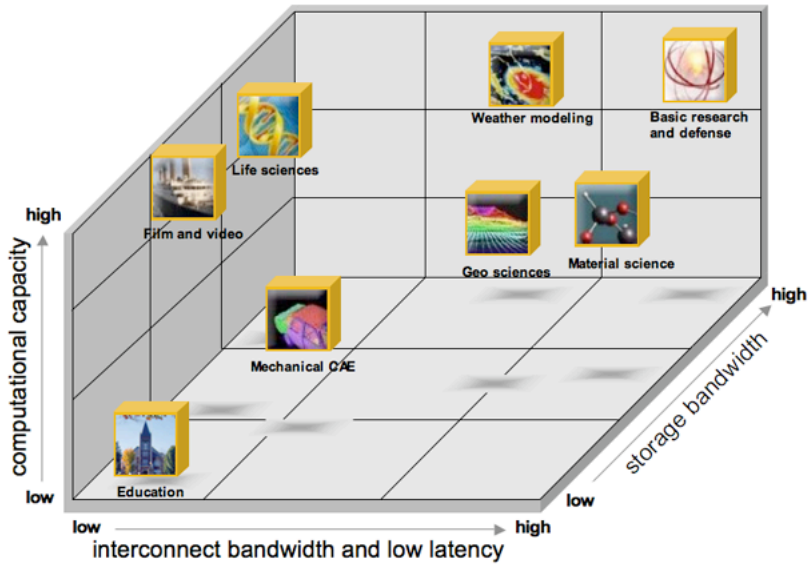

An important conclusion that we can draw from Figure 2 is that clustering solutions not only need to provide extreme levels of performance, but the solutions need to be flexible enough to span the wide variety of architectural requirements of HPC applications.

Challenges

While there are many benefits of clustering lower-cost systems into a practical problem-solving capability, the process is not without challenges. Some of these challenges include:

Complexity—A cluster is, by definition, a collection of independent computer systems, each with its own processor(s), memory, operating system, and related management requirements. The integration and ongoing management of tens, hundreds, or thousands of independent systems presents a considerable technological challenge.

Application availability—In order to take advantage of a cluster, applications must be structured to employ the multiple computing elements, typically with some form of parallelism. Because there is overhead involved in communicating between nodes in a cluster, application algorithms must comprehend the overhead, and break problems into parallelizable chunks that are large enough to be practical to distribute across the cluster. Consequently, applications must be designed (or redesigned) to manage this kind of parallel processing.

Scalability—It is one thing to connect hundreds of small systems together; it is quite another to reap the benefits of the clustered configuration. As previously mentioned, every connection between nodes introduces overhead to the configuration; it is quite easy to construct a cluster where the interconnect presents more overhead than is provided by the platforms within the cluster. Scalability is the ability to provide real world increases in performance commensurate with the number of computing elements in the cluster.

The HP Unified Cluster Portfolio addresses many of these challenges by providing a modular approach to selecting and integrating platforms, interconnects, and management software.

The HP Unified Cluster Portfolio

Introduction

The HP Unified Cluster Portfolio is a modular package of hardware, software, and services for scalable computation, data management, and visualization. It features a range of platforms, an extensive library of open source and commercial middleware, and the latest in industry-standard interconnect technology in a factory-built, tested, and supported solution.

The Unified Cluster Portfolio's modular framework (Figure 3) removes the complexity normally associated with deploying a custom cluster solution by providing a pre-tested, pre-configured, and fully integrated solution. The customer's cluster solution is built at HP's manufacturing integration centers and final integration occurs at the customer site. All software offerings are tested and verified by HP partners to run on a unified platform, enabling rapid deployment and a comprehensive environment for high performance computing. In addition, integration and optimization of applications reduces the risk and complexity of customized porting and provides a unified application environment for diverse workloads and users.

Providing customers with a broad choice of tested options is a key component of the Unified Cluster Portfolio. The Unified Cluster Portfolio provides a choice of processors, interconnects, operating systems, and proven open source and commercial middleware, and application software. These choices are possible because HP leverages the expertise of professionals worldwide in the open source community and commercial companies.

Figure 3. Overview of the modular design of the Unified Cluster Portfolio

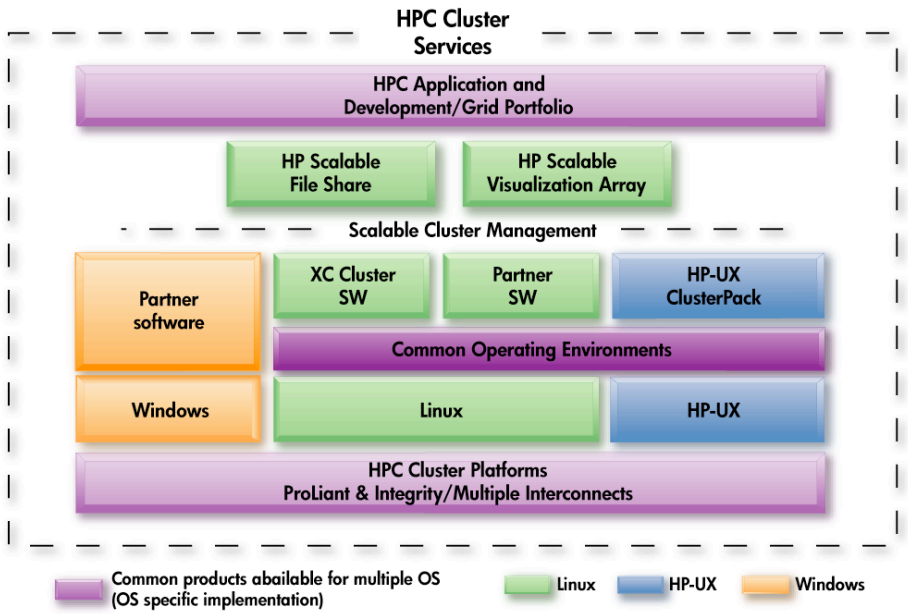

Additional HP expertise and technologies, such as the HP StorageWorks Scalable File Share, maximize performance, ensuring reliable and scalable performance. In addition, customers benefit from the single, unified design methodology and production of scale because it lowers cost, improves quality, and drives breakthrough price/performance.

As a framework for our description of the HP Unified Cluster Portfolio, we will examine four main elements that make up a cluster solution: platforms, interconnects, usability and management software, and integration services.

Cluster Nodes

The foundation of exemplary cluster performance and price/performance is the platforms used as nodes within the cluster. HP relies on industry-standard processors[2] as the core technology in cluster nodes. These processors are surrounded by an infrastructure that satisfies the intense compute and I/O needs of HPC applications. Typically these platforms are one to four processor systems, such as the HP Integrity rx1620 and rx4640 models, but other systems may be allowed as cluster nodes.

Figure 4. Generic cluster architecture showing different node types.

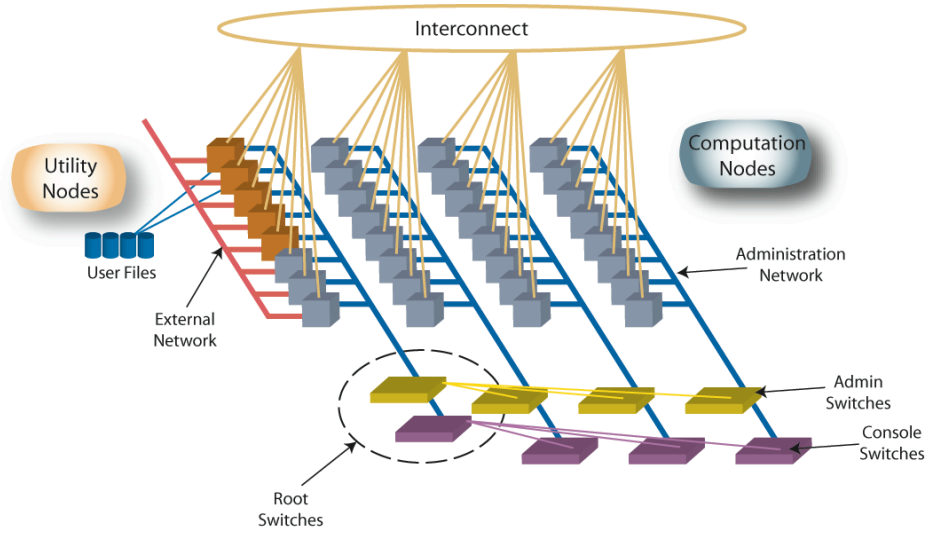

Platforms may be used as compute (application) nodes, service (control) nodes, and/or file services nodes as shown in Figure 4.

HP Cluster Platforms at-a-glance

As shown in Table 1, HP offers three choices of platforms, all based on industry-standard processor technology. Please refer to appendix A for detailed specification of each of these platforms.

[2] Excluded from this document are the HP PA-RISC and Alpha processor-based systems. Customers using these systems are being migrated to the standard processor technologies mentioned.

Table 1. The foundation of the Unified Cluster Portfolio is the choice of available platforms.

	Compute node and processor	Control node	Interconnects	Operating Systems
HP Cluster Platform 3000	HP ProLiant DL140 and DL360 servers with Intel EM64T Xeon™ processors	HP ProLiant DL380 servers	Myrinet, InfiniBand, or Ethernet	Linux or Microsoft Windows
HP Cluster Platform 4000	HP ProLiant DL145 and DL585 servers with AMD Opteron™ processors	HP ProLiant DL145 or DL385 servers	Quadrics, InfiniBand, Myrinet, or Ethernet	Linux or Microsoft Windows
HP Cluster Platform 6000	HP Integrity rx1620 and rx2620 servers with Intel® Itanium® 2 processors	HP Integrity rx1620, rx2620 or rx4640 servers	Quadrics, InfiniBand, or Ethernet	Linux or HP-UX

All of the HP Unified Cluster Platforms are based on industry-standard, "off-the-shelf" processors. Using industry-standard processors provides the highest level of price/performance—because of the large volumes of these processors that are manufactured and consumed worldwide—and the largest library of available applications.

HP Cluster Platform 3000

The HP Cluster Platform 3000 features the Intel® Xeon™ processor-based ProLiant DL140 or DL360 servers as compute nodes. Compute nodes are combined with one or more control/utility nodes, which enable a flexible cluster design to meet wide range of requirements. ProLiant DL380 servers are offered as the Control (head) node, and as optional Utility nodes, enabling expansion of cluster capabilities for administrative functions such as file services or additional log-in servers. All nodes feature Xeon processors with Extended memory 64 Technology (EM64T), with 1 or 2 MB Cache and an 800 MHz FSB. New EM64T technology allows users to run either 32-bit or 64-bit applications[3]. Both servers are optimized for rack mounting (1U and/or 2U form factors) and provide extensive features important for clustering, such as embedded remote management and optional redundant power.

HP Cluster Platform 4000

The HP Cluster Platform 4000 features ProLiant servers with AMD Opteron™ processors, a choice of Quadrics, Myrinet, InfiniBand or Ethernet interconnects, support for either Linux or Microsoft Windows environments, and other middleware options. ProLiant DL145 or DL585 servers are available as compute nodes, supporting a mix of dual and four way processor nodes within a single cluster. Compute nodes are combined with one or more control/utility nodes, which enable a flexible cluster design to meet wide range of requirements. ProLiant DL385 servers, along with the DL145 servers, are offered as the Control (head) node, and as optional Utility nodes, enabling expansion of cluster capabilities for administrative functions such as file services or additional log-in servers. The HP Cluster Platform 4000 is an excellent performer in environments where the workload is a mix of 32-bit and 64-bit applications[4], such as electronic design automation, life sciences, and geosciences. Applications that are compute intensive and run well on clusters of IA-32-based systems, such as many computational fluid dynamics applications, will also see advantages with AMD Opteron processor-based clusters.

HP Cluster Platform 6000

The HP Cluster Platform 6000 features HP Integrity servers with Itanium 2 processors and a choice of Quadrics, InfiniBand, or Ethernet interconnects. Co-developed by HP and Intel, the Intel Itanium architecture enables higher processor efficiency, performance, and scalability.

[3] Depending upon the operating system selected.
[4] Depending upon the operating system selected.

Support for either Linux or HP-UX operating environments and numerous middleware options are offered. The Cluster Platform 6000 offers dual processor Integrity rx1620 and rx2620s servers, as well as the four-way rx4640 servers, which can be combined in the same cluster. The HP Cluster Platform 6000 delivers superior 64-bit performance and can easily scale up or scale out for complicated workloads, such as seismic analysis, molecular modeling, and structural mechanics.

All of the Itanium-based platforms provide a high performance infrastructure (chipset) with specific HPC enhancements, a high degree of upgradeability to newer technology, and a high degree of expandability.

Interconnects

There are a number of industry standard interconnects available—some directly resold by HP and some that may be acquired through a third-party vendor (although HP is able to perform the integration, please see " Integration, Services and Support", below). It should be noted that there is no "best" cluster interconnect, as each one has certain advantages and disadvantages. The proper solution depends upon the kinds of problems being solved and how the cluster interconnect performs in that particular situation.

HP sells and services—as standard products—the combinations of interconnect/platform/operating environments shown in Table 2. In cases where an interconnect may not be under standard support, HP Services organizations may provide custom integration.

Table 2. Interconnect support on platform/operating environment combinations.

	Linux/ CP3000	Linux/ CP4000	Linux/ CP6000	HP-UX/ CP6000
Gigabit Ethernet	yes	yes	yes	yes
Quadrics	yes	yes	yes	no
Myrinet	yes	yes	no	no
InfiniBand	yes	yes	yes	yes

For additional information on each interconnect, please see Appendix B, "Interconnect Specifications."

Gigabit Ethernet

Gigabit Ethernet is best suited as a cluster interconnect for capability clusters and throughput clusters where latency is not critical—batch processing of lots of small jobs, for example.

For Cluster Platforms with Gigabit Ethernet as the system interconnect, HP's ProCurve family of switches are deployed either as a single cluster switch or in tiered hierarchy utilizing federations of larger ProCurve switches, configured with smaller switches directly connected to nodes. Gigabit Ethernet is supported on both HP-UX and Linux.

Myrinet

Myrinet is a cost-effective, high-performance, packet-communication and switching technology that is widely used to interconnect clusters, provided by Myricom. Introduced in August 1994, there are now many thousands of Myrinet installations, ranging in size to more than 1,000 hosts and more than 2,000 processors. These sites include many of the world's premier cluster-computing systems. Over 38 % of the November 2004 TOP500 supercomputers use

Myrinet technology. Myrinet switch networks can scale to tens of thousands of nodes. Interfaces execute firmware to offload protocol processing from the node processors.

Quadrics

Quadrics Supercomputers World Ltd (QSW) provides QsNet, a high bandwidth, ultra low latency interconnect for high-end scalable systems. The technology has been developed from the outset to support the requirements of supercomputer class systems, with the emphasis on performance, resilience, security, and data integrity. QsNet can provide a peak bandwidth, after protocol, of 340Mbytes/second in each direction. The design of the data network guarantees that all nodes obtain this level of performance, whatever the size of system.

Quadrics also provides a set of optimized libraries for common, distributed memory programming models to exploit the capabilities of the base hardware. The kernel communication layer allows system services to take advantage of the performance of QsNet. Quadrics' RMS software completes the picture, providing monitoring, partitioning, and scheduling tools to manage large scale systems.

QsNet is supported on all platforms under the Linux operating environment.

InfiniBand

The InfiniBand architecture is a new I/O specification defined by the InfiniBand Trade Association. InfiniBand utilizes a point-to-point linking technology as the basis for an I/O "fabric" to increase the aggregate data rate between servers and storage devices. The InfiniBand product is available in two switch sizes and can be configured to support up to 512 nodes. The switch enclosure can be connected to the clustered nodes using either PCI-X or PCI-E adapters.

Interconnect Summary

When deciding which interconnect to use, there are three main variables to consider: price, bandwidth, and latency. Generally, interconnects with higher performance (high bandwidth and lower latency) are priced higher. To help assess which interconnect to use, determine the type of applications that will be run on the cluster and analyze the expected interconnect traffic. Some variables to look at:

- How much communication is needed by the application(s)? For example, "embarrassingly parallel" applications—those that don't have much need for communication across the interconnect—might be able to get away with a lower-cost and lower performing interconnect.
- How much filesystem traffic will traverse the interconnect? For example, some CAE applications require fairly small amounts of input data, but generate huge intermediate (local to the node) scratch files. In this case, little node-to-node communication is required, and the interconnect will fairly lightly loaded.
- How high of a reliability factor is required? In some situations, certain nodes (and databases that reside there) may have a high availability requirement; it might be wise to provide a redundant path to these nodes via additional physical interconnect paths.

Operating Systems

Linux

HP's team of experts has extensive experience with Linux clusters, having delivered and supported some of the most powerful Linux clusters to date in both public and private sectors. All three HP Cluster Platforms are qualified and tested for the Linux OS with your choice of Red Hat EL 3.0 or 4.0, SUSE SLES 8 or SLES 9—each available with HP support services. A Red Hat compatible Linux OS is integrated with HP XC System Software, HP StorageWorks Scalable File Share products, and the HP Scalable Visualization Cluster, based on SEPIA technology (see "Complementing Technologies," below).

HP-UX

The HP-UX operating system is offered on the Cluster Platform 6000 (Integrity Servers). A proven operating system, HP-UX is one of the premier UNIX implementations in the market. The stability and completeness of HP-UX is especially important for enterprise customers who demand a robust operating environment for mission critical applications. An extensive portfolio of applications and tools are available for the HP-UX OS. The core software component of the HP-UX clustering capabilities is ClusterPack.

Windows

For customers running in a Microsoft Windows environment, HP offers cluster solutions with Windows 2003 as a featured option with the HP Cluster Platform 3000 and 4000. By maintaining a Windows environment, customers are able to smoothly expand their technical computing capabilities with minimal training and disruption.

Common Operating Environments

HP-MPI

A key feature of the Unified Cluster Portfolio is the availability of the HP-MPI libraries and drivers, which enable software developers, including major ISVs, to develop a single set of executables across multiple interconnect technologies. By providing a common reference design (HP-MPI with Red Hat 3.0 based) and defining a full environment, a single cluster is able to support a mix of ISV applications. This approach has been enthusiastically endorsed by the ISV community, resulting in rapid availability of an extensive portfolio of commercial applications.

HP-MPI complies with the MPI-1.2 and MPI-2 standards and is a high performance native implementation. HP-MPI enhancements provide optimized point-to-point and collective communication routines, including shared memory for intra-node communication and remote direct memory access (RDMA) technology for inter-node communication whenever possible. HP-MPI supports 32- and 64-bit applications, single- and multi-threaded, and provides tools to debug and instrument MPI execution. Applications built with HP-MPI are interconnect-independent; it is not necessary to recompile or relink an application in order for it to operate on a different HP XC System Interconnect.

With HP XC System Software Version 2, HP-MPI V2.1 is fully integrated with SLURM resource management. In addition, HP Cluster Platforms support popular MPI implementations, such as MPICH.

Management and Usability Software

One of the obvious characteristics of clusters is that they are composed of a fairly large number of individual (and largely independent) platforms. While this provides excellent price/performance, it also introduces management and usability issues, because each node has its own copy of an operating system, middleware, and applications and user support databases. To make the promise of clustering a reasonable proposition, these systems must be easy to use and to manage.

HP has always been at the forefront of management middleware technology (for example, HP OpenView, introduced in 1996, has more than 135,000 total installations); cluster management is no exception. HP provides a choice of management and usability options for both Linux and HP-UX operating environments, including HP's XC System Software, the HP-UX Clusterpack, and third party solutions such as Scali Manage and Scali MPI Connect. Of course, customers have the flexibility of also being able to install solutions provided by the open source community such as OSCAR and ROCKs.

Linux—XC System Software V2.1

The XC System Software is based on a standard Linux distribution combined with several open source packages. This open source base combined with technology from Hewlett-Packard and its partners achieves a complete solution for development, execution, and management of simultaneous parallel and serial applications.

The XC System Software accommodates the diverse system configurations and user workloads common in HPC production environments. Specific features of this system design include:

- Optimized application performance—Resource management, scheduling, and job launch technology is provided by products from Platform Computing and the open source-based Simple Linux Utility for Resource Management (SLURM). Application availability is ensured by preserving the underlying Linux API/ABI and integration of HP-MPI to provide an interconnect-neutral (but MPICH compatible) application MPI library. Local storage access is supported for standard NFS and StorageWorks Scalable File Share (SFS)—see "Complementing Technologies," below.
- Scalable cluster management—The HP XC System Software provides single-point of control for installation, node power management, monitoring, and general systems administration.
- Automated system initialization—The HP XC System Software provides the maximum amount of automation possible while allowing for high levels of flexibility. Software installation is done first on the cluster's head node, followed by a topology discovery stage. Using this information, the administrator assigns roles for various nodes (such as log-in or I/O) as required. Each node is booted over the network with an auto install kernel, and then a Base Image is distributed, along with node-specific configuration information.
- Advanced resource management—LSF HPC from Platform Computing Corporation is deployed within the HP XC Cluster for managing batch and interactive jobs. In addition to launching jobs, LSF provides job management and information capabilities. As noted above, LSF is integrated with an underlying resource manager, SLURM. With adherence to industry standard interfaces, the XC software is fully capable of supporting other popular resource/job and grid managers that may already be in use at customers' sites.
- Utilization and co-development of open source technologies—HP XC System Software is an HP-supported and defined implementation of an open source architecture for high performance clusters. Open source technologies develop at a pace rarely matched by proprietary solutions. By choosing to build upon and integrate with, rather than compete with, the established cluster technologies that comprise the HP XC Cluster solution, HP and its collaborative partners extend the functionality of clusters, instead of merely replicating it.
- Support for global file system with high performance I/O—The HP XC System Software supports NFS 3 and is also enabled for Lustre client services for high-performance and high-availability file I/O.

Table 3 summarizes the components of the HP XC System Software.

Table 3. Components of the HP XC System Software stack.

Function	Implementation
Distribution / Kernel	Red Hat 3.0/2.4.21 kernel compatible
Cluster Alias	LVS—Linux Virtual Server
Outbound Net	NAT—Linux Network Address Translation
Batch	Platform LSF V6.0
Launch	SLURM - Simple Linux Utility for Resource Management
System Files Management	SystemImager +limited configuration tools +database
Console	Console Management Facility (CMF)
Operations	Pdsh, syslogging

Monitoring	◦ Nagios [with plug ins]—Open source service and networking monitoring program ◦ SuperMon - High performance cluster monitoring
MPI	HP-MPI v2.1

HP-UX—ClusterPack

HP's industry-standard and award-winning operating environment, HP-UX, represents the highest proportion of installed and sold systems by HP (about three times that of Linux over the past several years). HP-UX represents a solid, proven operating environment for HPC applications. The core component of the HP-UX clustering capabilities is ClusterPack.

The HP ClusterPack is an integrated cluster management solution specifically designed for HP-UX and HPC applications. The hpc/ClusterPack™ integrates HP's enterprise system management solution, ServiceControl Manager (SCM), and Platform Computing's latest offering, Clusterware Pro™ V5.0. ClusterPack includes features for installation and configuration, single-point system administration, and distributed workload management.

The ClusterPack is available as a standard option for the ClusterPlatform 6000, with Gigabit Ethernet and InfiniBand, in configurations of up to 128 nodes. In addition, ClusterPack is supported on clusters assembled with multiple larger Integrity servers, including the HP Integrity Superdome.

Scali

Scali and HP have partnered to deliver Scali Manage™ and Scali MPI Connect™, advanced Linux cluster solutions. Scali Manage provides comprehensive tools for system installation, configuration, management, and monitoring. Scali MPI Connect is a fully integrated message passing interface solution. It enables companies to take advantage of diverse high performance interconnects and utilize demanding application software capabilities.

Software Development Environment

HP and its software partners offer a powerful and comprehensive HPC software development environment for producing and optimizing high performance parallel applications. HP and/or partners provide compilers, math libraries, and other development tools for shared and distributed memory systems in HP-UX and Linux cluster environments.

For additional information, and lists of software development tools that are available, please see http://www.hp.com/techservers/clusters/sw_development.html.

Complementary Scalable Technologies

Advances in computational performance have outpaced capabilities in scalable storage, I/O, and visualization tools. In order to provide customers that have these requirements with a balanced approach to scaling their cluster capabilities, HP provides two complementing technologies:

- StorageWorks Scalable File Share, based on Lustre technology.
- HP Scalable Visualization, based on Sepia technology.

Both of these technologies are available in the Unified Cluster Portfolio, and are described in greater detail below.

Scalable File Share (SFS)

HP SFS is a powerful file server that gives users of Linux clusters a storage option that is easy to use and easy to administer. HP's industry-leading, open source-based file server increases aggregate bandwidth by distributing files in parallel across clusters of industry-standard server and storage components.

HP SFS can span dozens to thousands of Linux clients, simplifying the ability to run clustered applications. HP SFS allows applications to see a single high bandwidth file system image regardless of the number of Linux clients. And because the interface for these Lustre file systems is POSIX compliant, with all of the standard interfaces in place, programs can run without modification. Administration and support is simplified by automatically managing the file server, processes, and databases that track all cluster components.

HP SFS includes three disk array storage options, based on the level of resiliency required: HP StorageWorks 3000 Enterprise Storage Array (EVA3000) RAID 5, the HP StorageWorks SFS20 RAID 5+1, and SFS20 RAID 5. The SFS20 RAID 5 option, based on the SFS20 SATA-based enclosures, offers a low cost option that also provides high performance and reliability. Optimized specifically for most HPC science and engineering applications running on Linux clusters, this option is the leading choice of universities, government laboratories, and HPC industries.

Scalable Visualization Array (SVA)

The HP Scalable Visualization Cluster (SVC) provides a high performance and scalable visualization capability for clusters. The HP SVA is a Linux cluster visualization system that distributes image data among a number of workstations, each working in parallel to render a portion of the image. Consequently, the data on each node is obviously smaller, and the rendering process takes substantially less time.

The HP SVC lowers the cost for higher performance visualization because it allows true interactive processing of large datasets. The SVA architecture provides a large display surface and high resolution, resulting in easier collaboration, more accurate data interrogation, and better use of the available visualization display resources.

Integration, Services, and Support

Expert end-to-end HPC services provide a single point of accountability, ensuring successful implementation—starting with planning, continuing through migration and transition, and ending with ongoing maintenance and optimization. HP provides expertise in systems integration and administration, applications migration and optimization, systems architecture, customer education, as well as implementation program management.

HP provides multiple organizations for assisting customers in integration, research, and evaluation of HP clustering products.

HP Services

HP Services provides a complete portfolio of services tailored to customers' business and IT needs. The professionals at HP Services help cut the cost and complexity of supporting hardware, software, and network environments. HP Services provides a single-point-of-accountability lifecycle support; rapid-response coverage around the clock and around the world; and flexible service plans.

HP Services offerings include: deployment services for risk-reducing installation and startup, implementation, and integration; availability services to proactively reduce downtime and meet service-level commitments; performance services providing objective technical assistance; and innovative support management services to complement internal capabilities with multivendor expertise.

HPC Consulting and Integration Group

HP 's HPC consulting and integration (C&I) provides a single point of accountability ensuring successful implementation, starting with planning through migration and transition to ongoing maintenance and optimization. The C&I group works closely with HP Engineering and HP partners, including Linux NetworX, Platform Computing, Scyld Computing, Sistina, and Scali, to offer cluster solutions to meet complex customer needs.

The HPC Enterprise Center

The Worldwide HPC Enterprise Center (HPCEC) develops expertise, hosts technical corporate visits, performs large benchmarks, delivers training, collaborates with key customers, and provides feedback to HP Engineering customer requirements on supercomputing, scalable applications, and solutions.

The HPCEC staff has in-depth expertise in scientific and technical disciplines as well as distributed parallel scalable applications and technologies. They also manage the HP Scalable Cluster Center (see below). The HPC Expertise Centers are equipped with HPC systems, technologies, and solutions from HP.

The HP Scalable Cluster Center

The HP Scalable Cluster Center, located in Littleton, MA, serves both as a development and proving facility for HP 's clustering products, and as a production facility to support benchmarking and R&D.

The HP Scalable Cluster Center is used by HP scientists and engineers for the development, testing and qualification of HPC cluster technologies, ongoing benchmarks and scalability demonstrations, and ISV application certification, optimization and scalability characterization. The Center's supercomputing power is also used by HP Labs to accelerate the pace of innovation in quantum computing, life sciences, and advanced computing technologies.

Appendix A. Platform Specifications

HP Cluster Platform 3000[5]

Platform	HP ProLiant DL140 G2	HP ProLiant DL360 G4 and G4P	HP ProLiant DL380 G4
Node Type	Compute or Control	Compute	Control or Utility
Processor	Intel Xeon 3.6 GHz with 2 MB L2 Cache/800 MHz FSB; Xeon 2.8 or 3.4 GHz with 1 MB Cache/800 MHz	Intel Xeon 3.0, 3.2 GHz (G4P only), 3.4 or 3.6 GHz G4:1 MB Cache/800 MHz G4P: 2 MB Cache/800 MHz	Intel Xeon 3.2, 3.4 or 3.6 GHz with 1 or 2 MB Cache/800 MHz; Xeon 3.0 GHz with 2 MB Cache
Processors per node	2 if compute, 1 or 2 if control node	1 or 2	1 or 2
Memory per node	1-16 GB	1-12 GB (8 GB for G4)	1-12 GB
Disk/media bays	2 Non Hot Plug SATA OR SCSI Drives	Up to 2 drives (SCSI or SATA)	Up to 6 drives (SCSI)
I/O Slots	2 PCI-X 64-bit/133MHz slot, with support for one PCI-Express with optional riser board	1 full length and 1 half length PCI-X slots (optional PCI Express)	3 PCI-X slots, with support for up to 2 PCI-Express with risers
Number of nodes	Minimum: 1 control node and 4 compute nodes Maximum: Total of 512 nodes (utility + compute nodes)		
Cluster interconnect	Myrinet XP, InfiniBand, Gigabit Ethernet		
Management network	10/100 network or Gigabit Ethernet		
SAN Storage in cabinet	Optional MSA1000 Storage subsystem mounted in cabinet as part of cluster		
Operating system options	Linux: Red Hat Enterprise Linux 3.0/4.0, SUSE SLES8/SLES9, Microsoft Windows 2003 HPC (up to 64 nodes, GigE only)		
Operating environment tools	Linux: HP-MPI, Scali MPI Connect		
Cluster management software options	HP XC System Software v2.1 Scali Manage		
Packaging: (available in two styles)	• High density packaging designed for optimal footprint and lowest entry price, available up to 128 compute nodes (delivered in 4 cabinets) • Modular packaging designed for expandability		
Services	HP Enhanced Services available for HP Cluster Platforms. HP offers on-site field installation. HP software support available for operating systems and HP software options.		
Warranty	Standard hardware warranty on platforms, based on underlying components		

[5] Please check the HP web site (http://www.hp.com/techservers/clusters/) for the latest specifications.

HP Cluster Platform 4000[6]

Platform	HP ProLiant DL145 G2	HP ProLiant DL385	HP ProLiant DL585
Node Type	Compute or Control	Control or Utility	Compute
Processor	AMD Opteron 2.2 or 2.4 GHz (Plus new dual core as released)	AMD Opteron 2.4 or 2.6 GHz (Plus new dual core as released)	AMD Opteron 2.2, 2.4 or 2.6 GHz single core AMD Opteron 1.8, 2.2 GHz dual core
Processors per node	2 if compute, 1 or 2 if control node	1 or 2	2 or 4
Memory per node	2-16 GB	2-16 GB	2-64 GB (32 GB at 400 MHz)
Disk/Media bays	2 drives	6 drives	4 drives
I/O Slots	2 x 133 MHz PCI-X (one full length, one low profile) Optional 1 x PCI-Express@x16 in place of full length PCI0X slot	2 - 64-bit/100 MHz PCI-X 1 - 64-bit/133 MHz PCI-X	6 - 64-bit/100 MHz PCI-X 2 - 64-bit/133 MHz PCI-X
Number of Nodes	Minimum: 1 control node and 4 compute nodes Maximum: Total of 512 nodes (utility + compute nodes)		
Cluster interconnect	Myrinet XP, Quadrics QSNet II, InfiniBand, Gigabit Ethernet		
Management network	10/100 network or Gigabit Ethernet		
SAN storage in cabinet	Optional MSA1000 storage subsystem mounted in cabinet as part of cluster		
Operating system options	Linux: Red Hat Enterprise Linux 3.0/4.0, SUSE SLES8/SLES9, Microsoft Windows 2003 HPC (up to 64 nodes, GigE only)		
Operating environment tools	Linux: HP-MPI, Scali MPI Connect		
Cluster management software options	HP XC System Software v2.1, Scali Manage		
Packaging (available in two styles)	• High-density packaging designed for smaller footprint and lowest entry price, available up to 128 compute nodes (delivered in 4 cabinets) • Modular packaging designed for expandability		
Services	HP Enhanced Services are available for HP Cluster Platforms. HP offers on-site field installation. HP software support is available for operating systems and HP software options.		
Warranty	Standard hardware warranty on platforms, based on underlying components		

[6] Please check the HP web site (http://www.hp.com/techservers/clusters/) for the latest specifications.

HP Cluster Platform 6000[7]

Platform	HP Integrity rx1620	HP Integrity rx2620	HP Integrity rx4640
Node type	Compute or Control	Compute, Control or Utility	Control or Utility
Processor	Intel Itanium 2 1.3 GHz (3 MB Cache, 400 MHz FSB); 1.6 GHz (3 MB cache, 533 MHz FSB)	Intel Itanium 2 1.3 GHz (3 MB Cache); 1.6 GHz (6 MB cache	Intel Itanium 2 1.5 GHz (4 MB Cache) 1.6 GHz (6 or 9 MB Cache)
Processors per node	2 if compute, 1 or 2 if control node	2 if compute, 1 or 2 if control node	2 or 4
Memory per node	1-16 GB	1-24 GB	1-128 GB
Disk/Media bays	2 drives	3 drives	2 drives
I/O Slots	2-64 bit/133 MHz PCI-X	4-64-bit/100 MHz PCI-X	4-64-bit/66 MHz PCI-X 2-64-bit/133 MHz PCI-X
Form factor	1 U	2 U	4 U
Number of Nodes	Minimum: 1 control node and 4 compute nodes Maximum: Total of 512 nodes (utility + compute nodes)		
Cluster interconnect	Quadrics QSNet II, InfiniBand, Gigabit Ethernet		
Management network	10/100 network or Gigabit Ethernet		
SAN Storage in cab	Optional MSA1000 Storage subsystem mounted in cabinet as part of cluster		
Operating System options	Linux: Red Hat Enterprise Linux 3.0/4.0 and SUSE SLES 8/9 (GigE only)[2], HP-UX 11i v2 (up to 128 nodes, GigE and InfiniBand only)		
Operating environment tools	Linux: HP-MPI and Scali MPI Connect HP-UX: HP TCOE including HP-MPI, HP MLIB		
Cluster management software options	HP XC System Software v2.1 Scali Manage ClusterPack for HP-UX		
Packaging: (available in two styles)	• High density packaging: designed for optimal footprint and lowest entry price, available up to 128 compute nodes (delivered in 4 cabinets) • Modular packaging: designed for expandability		
Services	HP Enhanced Services available for HP Cluster Platforms. HP offers on-site field installation. HP software support available for operating systems and HP software options.		
Warranty	Standard hardware warranty on platforms, based on underlying components		

[7] Please check the HP web site (http://www.hp.com/techservers/clusters/) for the latest specifications.

For more information

http://www.hp.com/techservers/clusters/
HP Clusters for technical computing home page.

http://www.hp.com/techservers/products/sfs.html
HPC storage solutions for technical computing home page.

http://www.microsoft.com/windowsserver2003/hpc/default.mspx
Microsoft's High Performance Computing for Windows Server 2003 home page.

http://www.hp.com/techservers/hpccn/sci_vis/
SEPIA web pages in the HP Collaboration and Competency Network (HP CCN)

http://www.hp.com/techservers/products/consulting_and_integration
Consulting and Integration home site.

http://www.hp.com/techservers/clusters/hpc_clusterpack.html
HP hpc/ClusterPack information.

http://www.hp.com/techservers/clusters/sw_development.html
Cluster software development information.

4AA0-0828ENW, 06/2005