HRG Assessment:

HP High Performance Computing LC Series

High performance computing (HPC), a strategic focus for HP, is a broader concept than the term High Performance Technical Computing (HPTC) implies. The introduction of industry standard platforms into this market by HP, IBM, and others heralds the increasingly commercial use of these systems that were previously used almost exclusively for technical applications. Today, RISC/Unix symmetric multiprocessor (SMP) servers and proprietary cluster platforms dominate HPTC, but there is a rapid movement toward Linux industry standard servers and clusters that is underway.

HRG's investigation of the industry standard, high performance clusters market looks in detail at HP's LC Series clusters. This paper provides an overview of the HPC market, HP's position in that market, and the position of HP's LC Series in the HPC marketplace. Additionally, this paper offers an unbiased examination of HP's competition in this market, provides a summary of market opportunities for the LC Series, and details why HP's LC Series is a good investment for customers looking to buy a HPC solution.

MARKET DYNAMICS

The introduction of the Xeon DP chip set in 2001 provided the necessary floating-point performance typically required by compute intensive applications. Today, Linux clusters are primarily based on Xeon processors. A market is emerging for Itanium 2 processors, in the form of large SMPs and clusters, and AMD Opteron products as stand alone servers and clustered systems. HRG expects this trend to continue with RISC/Unix machines being replaced by flexible, low cost industry standard machines based on Intel or AMD. The vendor with the "best" industry standard system portfolio and cluster strategy should capture significant share of the overall Linux cluster market. The Linux cluster market is expected to grow from $1.1B in 2003 to about $2B in 2005 with about 70% - 80% being in the HPC market segment.

The Linux cluster market in 2003 was more than 1/3 of the overall Linux server market in terms of revenue. HP dominated the worldwide Linux server market with about 29% of revenue market share and Linux servers provided over 25% of HP's HPTC market share of 37%. The worldwide Linux cluster market is expected to grow faster than the worldwide Linux server market over the next 2 – 3 years as transition continues to take place from RISC/Unix technology to industry standard server and operating system technology. This growth will occur because HPC buyers are focused on price/performance, and Linux clusters have a 5x to 20x price/performance advantage over previous generation RISC/Unix platforms.

Harvard Research Group, Inc.
P.O. Box 297
Harvard, MA 01451 - 0297 USA

Tel. (978) 263-3399
Fax. (978) 263-0033
hrg@hrgresearch.com

In market segments that are price sensitive, customers are typically willing to trade a level of software functionality for increases in compute power. The "need for speed" continues with bi-annual chip set performance improvements by Intel and AMD and new additions or replacements in the HPC market are expected to be for the most industry standard Linux clusters.

The price/performance advantages of industry standard clusters is driving the expansion of the HPC marketplace as companies and firms, who previously could not afford this type of computing, can now make an industry standard cluster investment and improve their overall productivity and product quality. The ISV community is well aware of this growth. Most of the top application providers in computer aided engineering (CAE), electronic design automation (EDA), petroleum engineering and seismic analysis, visualization-rendering and gaming, and bio-life sciences have ported their applications from Unix to Linux based industry standard servers in anticipation of this impending shift to industry standard HPC.

Current Events

The transition to Linux clusters is substantiated by a recent survey of HPC scientists who stated that 85 % - 90% of HPC applications could be run on Linux clusters today after simple recompilation from UNIX. In verticals such as life sciences and pharmaceuticals, 65% and 85%, respectively, use clustered platforms and more than half of those are Linux clusters.

Another reason for the high growth rate of Linux clusters compared to the overall Linux server market is that the HPC market segment is not limited by the lack of availability of ISV applications on Linux as in the enterprise market. Many HPC applications are "home grown."

Development tools and utilities familiar to developers in the UNIX world are being migrated to Linux. HP, with its partnership programs has partnered with many tool vendors and offers tool suites from Portland Group, Vampir, Pallas, and SCA Linda as reference sales items. Because researchers and scientists may not have IT skills or parallel development knowledge but need the support, an emerging parallel development utility designed specifically for engineers and research professionals called CxC has been introduced by a start up company named Engineered Intelligence. CxC enables scientific and research professionals to develop parallel applications quickly and apply them to both Linux and Windows clusters for superior execution on entry-level and low-end compute clusters where the investment required for high-speed, low-latency bandwidth is not affordable.

Microsoft Has a Place in HPC

Microsoft is investing in HPC technology and business programs with HP. The two companies are partnering in technical collaboration with industry partners to provide additional choices for emerging customers who may prefer Microsoft Windows 2003 over Linux in an industry standard compute farm configuration. Many HPC applications provide a Windows front end for data calculation and presentation. However in some areas, customers may prefer to stay with an all-Windows solution. This seems to be true in life and materials sciences as well as some isolated CAE sites in which small enterprises do not have the skills or resources to have a two-tier operating system environment. This provides another level of choice for the HP solution.

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

The LC Series

HP's LC Series of high performance, industry standard server cluster solutions consists of three offerings all addressing different segments of the HPC marketplace. The LC 1000 Series, based on the ProLiant DL140 server using the Xeon processor, the LC 2000 Series, based on the ProLiant DL360 server using the Xeon processor, and the LC 3000 Series based on the ProLiant DL145 server using the Opteron 248 processor.

HP's stated design objective is for it or one of its certified partners to provide customers with complete turnkey solutions including software per customers' requirements. LC Series clusters are shipped with servers, network communications, power supplies, etc. They are integrated and fully tested to meet design objectives. Very little setup is required at the customer site except for cabling between multiple racks when a multi-rack cluster is ordered.

Linux -
Replacing UNIX on the Top500 List

A visit to the Top500 list of supercomputers shows that 170 of the top 500 largest, most powerful machines in the world are Intel-based, and 152 of them were built using Xeon processors. We are predicting that most of the 100 or so RISC/Unix platforms at the bottom of the list will be bumped off later this year as a result of the growth of Linux clusters. The number of Intel-based clusters on the list grew by 115 from 2002 to 2003.

LINUX CLUSTER MARKET

The Linux cluster market is experiencing rapid growth as companies increasingly realize the exceptional price/performance advantage that industry standard Linux clusters deliver. Firms looking for a competitive edge and the ability to deliver better services to their customers are the ones that have adopted Linux and Linux cluster technology the fastest. Those market segments where there is fierce competition and where technology can be used to get an edge on the competition include financial services, pharmaceuticals, geophysics (oil/gas), semi-conductors, automotive, and aerospace.

Linux clusters have already taken over the application execution environments that host Web services-based application architectures (or service oriented architectures (SOAs)). These application execution environments are capable of seamlessly running both HPC and commercial application workloads and managing the Linux clusters that underpin the execution environment.

Linux clusters are revolutionizing HPC. The price/performance advantages they hold over other solutions are significant, ranging from 5x to 20x. A complementary technology, around for several years, that had its origin in HPC is distributed workload management, now generally referred to as grid computing

Vendors who sell into the enterprise market in areas such as financial services, healthcare, etc., have said that about 70% of the applications hosted on their systems are compute intensive and most of these are considered mission critical. DataSynapse, a grid middleware vendor that sells into the financial services market and into government and oil/gas verticals, says that almost all applications it encounters that are compute- and data-intensive are considered to be mission critical.

With regard to Linux clusters one key issue is cluster management — whether it is for high performance technical computing or grid computing. Today, there are a number of vendors with hardware platform independent, Linux cluster management solutions for clusters with hundreds and even thousands of nodes. In addition, there are a number of open source products, such as NPACI ROCKS, that can fill the management solution need with varying degrees of functionality. Utilities such as Kickstart, Ganglia, and SystemImager can also be used to provide administrative support for Linux clusters.

Software companies such as Cassatt and Scali provide platform independent, high performance Linux cluster management solutions, while IBM, Linux NetworX, and others provide proprietary solutions targeted for their own hardware products. HP has partnerships with Cassatt, Platform Computing, and Scali as well as a reference

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

relationship with NPACI ROCKS. There is a new class of "cluster" management solutions for grid computing from a number of ISVs (independent software vendors). They include Axceleon, United Devices, and Platform Computing as well as Data Synapse who have partnerships with HP. These cluster management solutions are designed to monitor and manage jobs across a network of heterogeneous computer systems. They can be tightly coupled systems as is the case with the LC Series, or they can be geographically dispersed systems connected with slower network attachments.

Linux Cluster Market Size

In 2000, the HPC market size was approximately $6.5B. Due to the economic downturn of the past two and one half years and a drop of about 13% in server sales overall, the 2003 HPC market size dropped to approximately $5B out of an overall enterprise market size for servers of approximately $45.5B. The overall HPC market size in terms of hardware, software, and services is between $10B - $13B (based on interviews with HP, IBM, and Sun). The overall cluster software market is expected to reach $1.5B by 2006.

The big players in HPC are HP (37%)[1], IBM (28%), Sun (19%), Dell (8%), and SGI (5%). The smaller players include Linux NetworX, Rackable, RackSaver, and RLX. When the small players' market shares are combined, they own only a small percent of the market.

Our conversations with vendors in the HPC market space lead us to believe that revenue growth in HPC will be close to flat or in the single digits overall for the next year or two. But the growth with respect to Linux clusters is expected to be much higher. The overall Linux server market in 2003 grew at a rate of 50% over 2002 with revenue of $2.9B. During this time period, the Linux cluster market was growing at a slightly higher rate. In 2003, the Linux cluster market was about $1.1B and 70% - 80% of that was in HPC. We expect the Linux cluster market to grow beyond $2B by the end of 2005.

HP'S Position In The HPC Market

The HPC server market is just over 10% of the overall server market in terms of revenue, but HPC is the area where new computing technologies are often created to service applications with requirements frequently beyond those of commercial applications. A good example of this is grid computing which was spawned from the HPC area. Financial services organizations are now using grid-computing technologies to solve problems in several minutes that used to take many hours, resulting in better service for customers.

HP has been involved in HPC for a long time with its VMS, Tru64 UNIX, and HP-UX platforms. Many of the servers in the current Top500 list are still Alpha-based Tru64 UNIX and HP Superdome systems with the Linux systems gaining rapidly. Today, HP is the HPC market leader in terms of revenue and is one of only two vendors (SGI is the other one with its Altrix 3000 machines) that have Itanium 2 servers on the Top500 list.

Linux Cluster Solutions

HP views life sciences, manufacturing, EDA, geophysics, visualization, finance, digital content creation, government/defense, basic science, and higher education as vertical markets for Linux clusters. HP has enjoyed a leadership position in Linux cluster solutions in HPC ranging from Alpha servers to ProLiant servers to Itanium servers to Linux workstations. HP has deployed thousands of nodes in Linux clusters for laboratories such as Sandia, Pacific Northwest National Laboratory (PNNL), Forecast System Lab, San Diego Supercomputing Center and Ohio Supercomputing Center, and in commercial enterprises such as British Petroleum, Aramco, General Motors, Toyota, Honda, Audi, Porsche, Boeing, Smith Kline, Johnson & Johnson, Chevron, Exxon, Williams F1 Racing, Norske Hydro, and Eli Lilly.

[1] The percentages indicate the percentage of the hardware revenue for HPC.

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

HP's solution model is to provide choices to customers — stand-alone servers, SMP's, clustered servers supporting HP-UX and Linux, and Microsoft Windows. HP also is providing life cycle solutions for customers with comprehensive services offerings ranging from site preparation to installation and training to facilities management and system exploitation.

HP partners with strategic software vendors such as Scali, Platform Computing, United Devices, Axceleon, and Altair Engineering to provide Linux cluster management for installing software, monitoring the health of the cluster, etc. Job/workload management software such as Platform Computing LSF and PBS-Pro from Altair Engineering can be integrated with HP Linux cluster solutions to allow users to submit and schedule jobs. HP's strategy is to partner with companies that specialize in cluster and job management solutions to provide customers with choice. HP also works with development tool kit vendors such as Engineered Intelligence with its CxC parallel development toolkit, the Portland Group, SCA with its Linda and Paradise utilities, Vampir, Pallus, MATHLab, and Mathmatica.

HP's strategic partner offerings have been involved in deploying clusters ranging from 16 to 1024 nodes. HP is reselling these offerings with service and support provided by partners. Customers can request HP's consulting and integration services for maintenance offerings on an as needed basis. These cluster offerings can differ from site to site. HP has negotiated a special HPC server price for both Red Hat and SUSE enterprise Linux editions giving HP a pricing advantage over the competition. HP provides customer service and support for Linux based servers with special pricing for Linux compute farms. It also offers a preferred software stack named XC, which is emerging as a controlled and fully supported Linux software suite for very large clusters in the 256 to 2048 node size.

Solution Services by HP

Much of HP's expertise in HPC over the past several years has been gained via hands-on work with its customers. In some installations, HP just makes sure that the nodes in a Linux cluster are connected and the cluster is up and running. In other instances, HP provides onsite support and training. HP is able to adjust to customer requirements and does not try to force training and support on customers when they are not needed. Services offerings scale from simple administration to knowledge education and training to systems architecture consulting and systems implementation planning and execution.

PARTNER DRIVEN LINUX CLUSTER OFFERING — HP's LC Series

To test the LC Series against the competition, we have chosen three vendors to review who have Linux cluster offerings. The standard configurations for the vendor's Linux cluster products are given in Table 1.

HPC Application Types

Linux clusters, such as those discussed in this paper, can serve as an application execution environment for serial, messaging, and parallel applications. These clusters generally consist of one control node, responsible for monitoring the execution of applications and the performance of the cluster, and a series of computing nodes that execute application code.

Serial applications are traditionally single machine applications in which a single server is required. Of the three types of applications, serial applications are the most ill suited for Linux clusters, but users are clustering these single server applications into clustered configurations for security, load balancing, and overall administration of the application environment.

Message passing applications are subdivided into tasks or segments. Each task is executed on a compute node with computed results aggregated to create a solution. Message passing applications such as Fluent, a fluid dynamics application, work well on Linux clusters. With message passing applications, the control node initiates the tasks of an application across the compute nodes and loads the appropriate data to each compute node. They

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

generally use a protocol such as MPI (Message Passing Interface) to communicate among compute nodes during execution.

The ideal application type for Linux clusters is parallel applications. HP has partnered with parallel development utility vendors to provide a smooth transition to parallel execution when parallel execution is possible. Parallel applications are found in all industries with potentially minimal command and data message passing between compute nodes. They are multiple executions on many compute nodes of the same application code using different data. Parallel applications include DNA sequence analysis, Genome processing, rendering, weather forecasting, etc. The control node passes data to each compute node and then initiates the applications, constantly fetching new data to direct the parallel servers until the end of the jobs.

HP LC 1000/LC 2000/LC 3000 Series

HP has three Linux cluster choices from the LC Series of clusters — LC 1000, LC 2000, and LC 3000. The LC Series gives users a broad range of choices with each choice designed to satisfy a user's specific requirements. In LC Series clusters, the control node performs job dispatching and monitoring functions. In message passing applications, a compute node executes an application and sends commands and data to other compute nodes. In parallel or serial applications, the control node and compute nodes cooperate by exchanging messages. Regardless of the type of application, a cluster interconnect is needed. Customers can select from Fast Ethernet, Gigabit Ethernet, and Myrinet interconnects for all three LC Series clusters. HP's LC Series Design and Configuration Guides provide customers with a choice of over 180 Linux cluster configurations, ranging from 8 to 128 nodes.

LC 1000 Series clusters are configured with a ProLiant DL380 (two-processor Xeon) server as the control node with multiple ProLiant DL140 compute nodes. The LC1000 is a low cost, high performance IA-32-based cluster with a limited management feature set. It is ideal for cost conscious customers who want to use their own management solution and where low cost/GFLOP is of primary interest.

LC 2000 Series clusters are configured with a ProLiant DL380 control node and multiple ProLiant DL360 G3 compute nodes. LC 2000 Series clusters provide high performance with a full function management feature set and enhanced memory management. The LC 2000 is designed to meet the reliability and management requirements of large enterprise customers needing sophisticated management and scale out capabilities.

Customers who buy LC 3000 Series clusters are generally looking for high performance clusters that start with 32-bit applications and provide seamless scale to 64-bit applications. LC 3000 Series clusters are configured with a ProLiant DL380 control node and multiple ProLiant DL145 compute nodes — AMD Opteron processors with up to 2.4GHz clock speed. Because the DL380 provides transparent monitoring and command processing and 32-bit Linux clusters are the target, the DL145 is not used as a control node. HP feels this is an acceptable solution in the short term.

LC Series clusters are shipped turnkey with a control node and compute nodes, in band and out-of-band management, a cluster interconnect network, power supplies, keyboard, monitor, a mouse used by the control node, and racks (assembled, integrated, and tested with all included components). HP also offers optional software for cluster and file system management, CS Service Carepacks, and Consulting and Integration Services.

The LC Series' Design and Configuration Guides define a suite of complementary storage subsystems based on the ProLiant DL380. The storage subsystems provide entry-level JBOD SCSI solutions through the ProLiant Modular Storage Array (MSA) for NFS, (Network File System), and GFS (Global File System) data management. Seven storage configurations have been defined and configured by HP. They are tested for optimal I/O bandwidth and overall performance on large Terabyte based configurations by HP's strategic ISV partners: PolyServe and Red Hat Sistina.

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

Special HPC editions of Red Hat Enterprise Linux (RHEL) WS 3.0 HPC Edition and SUSE Linux Enterprise Server (SLES) 8 HPC Edition run on LC Series compute nodes. RHEL ES 3.0 and SLES 8 run on control nodes. Linux software is provided with incident support from HP as well as one or three-year subscription services from Novell or Red Hat. Scali Manage and Scali Connect are also optionally available as management solutions for LC Series clusters.

The LC Series also supports strategic, certified ISV applications that have been tested and certified by the ISVs. This gives HP an advantage over Dell, IBM, and RackSaver because it avoids software incompatibilities and functional differences that cause numerous problems for customers implementing HPC solutions.

Table 1: Configuration Summaries for Linux Cluster Offerings

Feature	HP	IBM	Dell	RackSaver
Cluster Brand Name	LC 1000, LC 2000, LC3000	Cluster 1350	HPCC Cluster	BladeRack Series
Primary Form Factors	1U	1U and Blade	1U	1U and Blade
Server Model(s)	Control node: DL380 Compute nodes: DL140 (LC 1000), DL360 (LC 2000), DL145 (LC 3000) Storage is provided via the DL380 control node, or optional storage subsystems are available exclusive of the LC Series. They range from entry NFS Server to HA NFS Server to Global File System Server. They scale to 16 server nodes and 48 TB of storage and are connected via the In Band Management Network for independent I/O from the Cluster Interconnect Network	Control node: xSeries 345 Compute nodes: xSeries 335, BladeCenter HS20, e325 Storage node: xSeries360)	Control node: PowerEdge 2650 (2U) Compute nodes: PowerEdge 1750 (1U), PowerEdge 2650 (2U), PowerEdge 3250 (2U), 1655MC (3U) Storage provided by SCSI-PowerVault 220S	Blade configurations: RS-1100V-44, RS-1100V-44XT, RS1100V-66, RS1100-66XT, RS-1100VM-88, RS-1200V-128
Primary CPU Chip Set/ Processing Speed	Intel Xeon 2.8, 3.0, and 3.2GHz (DL140, DL360, DL380) AMD Opteron 1.6, 1.8, 2.2, and 2.4GHZ (DL145)	Intel Xeon 2.8, 3.0, and 3.2GHz AMD Opteron for e325 (up to 2.2Ghz)	Intel Xeon 2.8, 3.0, and 3.2GHz (PowerEdge 1750/2650) Itanium 2 (PowerEdge 3250) Pentium III (1655MC)	Intel Xeon 2.8, 3.0, and 3.2GHz Itanium 2 AMD Athlon and Opteron
Number of Nodes Supported	Minimum: 8 nodes Maximum: 128, configurable to 1024 nodes through switch additions and up to 96 storage nodes and 142TB of Fibre channel storage	Minimum: 4 nodes Maximum: 512 nodes including up to 32 storage nodes	Maximum: 128 nodes	44, 66, 88, or 132 nodes and up to 176 processors
Configuration Packaging	3 pre-packaged solutions based on a standard 32-node cluster extensible to 64 nodes. Optional cluster configurations up to 180+ configuration designs are accepted for Factory Express	Pre-packaged based on choice of control/compute servers	Can mix and match based on PowerEdge servers – offers pre-bundled order codes for 8, 16, 32, 64 and 128 node and one and two processor configurations	Can order on-line based on pre-selected node options, type of processor, amount of memory and internal storage
Operating System	RHEL ES 3.0, SLES 8, Windows 2003	SLES 8, RHEL 3.0 ES & WS	RHEL 3.0, Windows 2003	RHEL 3.0, SLES 8.0, Windows 2003

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

Table 1 (*continued*): Configuration Summaries for Linux Cluster Offerings

Feature	HP	IBM	Dell	RackSaver
Memory (standard)	1GB (DL140, DL360, DL380) 2GB (DL145)	512MB or 1GB	1GB (PowerEdge 1750, 2650)	512 GB
Memory Expanded	DL380: up to 12GB DL360: up to 8GB DL140: up to 4GB DL145: up to 16GB	Up to 12GB (e325)	Up to 8GB (PowerEdge 1750) Up to 12GB (PowerEdge 2650)	Up to 16GB (Opteron)
Level L2 Cache	1MB to 2MB depending on model	512KB (x335), 1MB (e325)	512KB (PowerEdge 1750)	512KB (Xeon), 1MB (Opteron)
Front Bus Speed	533MHz for Xeon servers	533MHz for Xeon servers	533MHz for Xeon servers	533MHz for Xeon servers
PCI speed	1 - 2 full length 64-bit/133MHz PCI-X slots depending on model	Two 64-bit, 100 MHz PCI-X slots per compute node	Two 64-bit 133MHz PCI-X slots	Two 64-bit 133MHz PCI-X slots
Internal Disk Storage	DL380: 34 – 792GB DL 140: 80 – 160GB DL360: 0 – 298GB DL145: 40 –120GB	x335: 293.6GB SCSI or 240GB IDE	PowerEdge 1750: up to 438GB PowerEdge 2650: up to 365GB	36GB to 146GB
External Disk Storage	Storage options: Entry NFS Storage – 1TB Low End NFS Storage – 2TB HA NFS Storage – 2TB Entry GFS Storage – 4TB Low End GFS Storage – 8 TB Mid Range GFS Storage – 16TB High End GFS Storage – 48 TB	Expanded storage capacity with FAStT200, FAStT600, FAStT700, and FASt900 storage servers and expansion units Integrated RAID-1 with mirroring	PowerVault™ 220S SCSI external storage device on the master node for primary storage	No information available
Interconnect Technology	Fast Ethernet Gigabit Ethernet Myrinet InfiniBand (3rd Party)	Fast Ethernet Gigabit Ethernet Myrinet	Fast Ethernet Gigabit Ethernet Myrinet	Fast Ethernet Gigabit Ethernet Myrinet InfiniBand
Interconnect Models (switches)	HP Procurve 2650 and 2848 series depending on cluster size or Myrinet-2000-Fiber/PCI using LANai 9 RISC @ up to 200MHz	No specific switch included; will interoperate with various Layer 2 switches	Dell PowerConnect™ switches	No specific switch included; will interoperate with various Layer 2 switches
Cluster Software Management Solutions	Scali Manage OSCAR ROCKS	Cluster Systems Management (CSM) IBM Director	Dell OpenManage	RackSwitch OSCAR ROCKS

Source: HRG, Dell, HP, IBM, and RackSaver, June 2004

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

IBM eServer Cluster 1350

IBM is the strongest competitor in HPC to HP. IBM has IA-32 servers (xSeries 335, 1U — the original compute node for the Cluster 1350 — and x345, 2U), Itanium 2 servers (x382, two-way, 2U), AMD Opteron servers (e325, 1U), and BladeCenter HS20 blades (7U chassis with up to 14 two-way Xeon blades) that can be plugged into Linux clusters. The x345 can be used as the control node as well as a storage node.

Any combination of Xeon, HS20, and Opteron nodes can run in a single cluster as compute nodes. The Itanium 2-based x382 server cannot be mixed with the other IBM servers in Cluster 1350. Standard configurations of Cluster 1350 include a management (control) node and up to 511 compute nodes, including up to 32 optional storage nodes that provide shared file storage. IBM provides a special ordering process for customers that require clusters beyond 512 nodes.

Cluster 1350 hardware is manufactured/integrated and tested at an IBM facility before being shipped to a customer. IBM is a single point of contact for warranty and maintenance service for the entire cluster, including OEM components.

Cluster 1350 nodes can run either RHEL or SLES. IBM pre-configures and pre-tests the clusters and ships IBM Cluster Management System (CMS) with them. CMS can distribute configuration files across a Cluster 1350 cluster, define/install nodes, install RPM packages, and display cluster attributes. If a SAN is present, PolyServe's Matrix can optionally be used as a clustered file system.

Dell HPCC Cluster

Dell clusters consist of the PowerEdge 1750 (1U, two-processor Xeon), PowerEdge 2650 (2U, two processor Xeon), PowerEdge 3250 (2U, two processor Itanium 2), and PowerEdge 1655MC building blocks. The PowerEdge 1655MC server can have up to six blades (12 Pentium III processors) consolidated in a 3U chassis.

The control node in HPCC clusters is the PowerEdge 2650. It can also be used as a compute node along with the PowerEdge 1750, PowerEdge 1655MC, and PowerEdge 3250. The majority of Dell's clusters are built around the PowerEdge 1750 and the PowerEdge 2650.

A HPCC Cluster solution is aimed at providing a validated set of components in a single package, which simplifies the design, ordering, and deployment process for customers. HPCC solutions support Fast Ethernet, Gigabit Ethernet, and Myrinet interconnects depending on customer application needs. Configurations start with eight nodes and include 16, 32, 64, and 128 node configurations. Clusters built with the PowerEdge 1655MC are configured with 6, 18, 36, 66, and 132 nodes. Dell has scaled HPCC Cluster solutions to over 1,000 nodes.

Dell teams with Cray to deliver clusters. It provides the hardware components and Cray integrates, tests, and supports the clusters. With assistance from Cray, Dell has delivered some large Linux clusters with more than 1400 nodes.

Dell offers its OpenManage management software with every cluster sold. It also has a partnership with MPI Software Technology, Inc. to support and distribute open source job management components such as OpenPBS (Open Portable Batch System).

RackSaver Clusters

RackSaver has two cluster offerings — the BladeRack and RackSaver 1U clusters. A BladeRack cluster is based on RackSaver's blade offering that delivers compute power in a small footprint. The RackSaver 1U server cluster is a rack-optimized high-density server.

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

The BladeRack series offers three 2U, two processors per node cluster products. One of these is the BladeRack MegaDense cluster. It is able to accommodate up to 88 nodes with 176 processors in a single rack with over 1.4 TFLOPS of compute power. The BladeRack UltraDense cluster can accommodate up to 66 nodes with 132 processors. BladeRack offerings can simultaneously contain the following architectures: Xeon, Itanium 2, Pentium 4, and AMD Athlon (and Athlon64) and Opteron processors.

The HIVE (High-Density Interconnected Vertically-Cooled Enterprise) server can house up to 8 or 11 blades with 16 or 22 processors. HIVE is designed for companies looking to migrate to a blade design. The eight-node HIVE is a rack product and the 11-node HIVE is a pedestal. HIVE clusters can be stacked on top of each other.

RackSaver's cluster server consists of up to 42 nodes (84 processors) of 1U RS-1100 servers in a standard height rack. All RackSaver products can be configured with various interconnect options such as Gigabit Ethernet, Myrinet, etc.

RackSwitch is RackSaver's attempt at cluster management. It is a limited product that monitors temperature, power on/off, etc. RackStor, RackSaver's storage product, is available in a 4U design that allows scaling by simply racking and stacking storage units. The RackStor 4150-200 contains 12 hot swappable SCSI drives for a total of 200GB of storage. Additional storage products are available with up to 1TB storage capacity.

VENDOR COMPARISON

To make a comparison among the four Linux cluster technologies reviewed in this paper, HRG provides a set of features for comparison. Important differentiators are breadth of offerings, cluster management solutions, deployment and installation capabilities, service and support, and pricing. The features include:

1. Fast processors and large cache size per processor
2. Two processors for most HPC applications
3. 1GB or 2GB of memory per processor
4. Dense packaging for compute nodes
5. Interleaved memory, if at the right price per processor
6. Minimal disk size or no disk at all
7. Support for Gigabit Ethernet, Myrinet, and optionally Quadrics
8. At least one PCI slot for high speed interconnects such as Myrinet
9. Linux-based
10. ISV application support in most HPC verticals
11. Ability to handle in-house developed codes
12. Service and support
13. Deployment and installation capability
14. Cluster management capability
15. Pricing

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

The Linux cluster offerings from Dell, HP, IBM, and RackSaver all have the first nine features in the above list. The four vendors generally use Xeon processors and the same interconnect technologies — Gigabit Ethernet and Myrinet for performance and Fast Ethernet for system administration.

Most of the components used to build Linux clusters are commercial off-the-shelf (COTS) products. HP, IBM, and RackSaver offer AMD Opteron servers and Dell does not. These three vendors have an advantage over Dell because the Opteron offering permits users to scale from IA-32 processors to 64-bit (Opteron) with little effort.

ISV application support is an area where HP has leadership over IBM, Dell, and RackSaver. HP is following the path that it has taken with its HP-UX offerings with solution engineering teams by focusing on industry wide support for infrastructure applications such as cluster management, grid management, and industry segmented applications.

Both HP and IBM have considerably more experience working with HPC ISVs than do Dell and RackSaver, but we give the advantage to HP for its aggressive business development and support for strategic ISV partnerships. This is re-enforcement for customer choice — one of HP's stated goals in this marketplace. HP and IBM are equipped to work with users who have in-house developed codes because they were active in HPC when these codes were being developed and many of them were developed on HP and IBM platforms.

HRG gives HP the nod in cluster management solutions. It offers Scali Manage, one of the top Linux cluster management solutions available, for LC Series clusters. HRG ranks the other vendor Linux cluster management solutions in the following order in terms of functionality for managing Linux clusters: CSM, OpenManage, and RackSwitch. HP has also entered into agreements with Axceleon, United Devices, and Meiosys to offer their cluster management solutions for specific industry applications, or where Scali Manage does not fit. HP is currently working on agreements with Platform Computing and DataSynapse to make their distributed job management solutions available for LC Series clusters.

HP's deployment and installation facilities also give it a lead over IBM and a significant lead over Dell and RackSaver. HP's turnkey integration and shipment to the customer with a fully operational cluster (Factory Express Service) leads the competition. The lead over RackSaver is primarily because of the differences in size between HP and RackSaver. Generally, Dell does not do its own cluster deployment and installation; instead, it has a partnership with Cray for installation. It is not uncommon for Dell customers; however, who buy HPCC clusters to do the installation themselves. HRG's perspective is that deployment and installation are two activities in which vendors directly interact with customers and the knowledge gained influences future cluster design and installation procedures.

HP's Factory Express Service provides customers with the ability to customize hardware and software solutions. HP, via its Factory Express Service, assembles hardware, integrates hardware, installs optional software, and tests the entire solution before shipment to the customer for immediate turnkey implementation. HP provides on-site installation and setup if customers desire it. In addition, HP provides documentation that gives customers a complete history of a cluster as it leaves the factory. This includes pre-defined settings, passwords, and cabling instructions in case the customer wishes to extend the cluster over time.

With the latest LC Series announcement in April 2004, clusters can be easily extended from 32 to 64 nodes without having to reconstruct the cluster networks. Customers can order a 32-node cluster and add an additional 32 nodes by adding the optional eight port Gigabit Ethernet or Myrinet 2000 switch cards to the previously installed cluster interconnect chassis.

IBM does not ship Cluster 1350 with Linux installed. LC Series clusters are shipped with Linux installed as an option. Customers who buy Cluster 1350 can pre-order Linux and IBM "drops it in the box," or they can get Linux from Red Hat. In either case, the user installs Linux on Cluster 1350 nodes. HP uses special HPC Linux editions

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

on its LC Series clusters; IBM does not. HP offers Itanium 2-based clusters via its XC Series clusters. Itanium 2 nodes are not standard materials for Cluster 1350. Customers must request Itanium 2 nodes through a separate, non-standard process at IBM.

Customers are concerned with software solutions and any conflicts that exist between infrastructure and industry specific applications. HP's LC Series provides solution testing and porting services for its ISV partners and has an architecture and strategy to provide solutions for various customer needs.

Partner solutions range from HPC cluster management, grid management, parallel development tool kits and compilers, job management, file systems management, checkpoint/restart, and availability management. Partners also provide end-user application support for computational chemistry, petroleum engineering and seismic analysis, and CAE. All of these applications are tested to run on the LC Series platforms. Partner solutions provide customers with a layered approach so they can choose where they want to be and have the confidence that the application(s) will work on their platform. Customers are concerned about certified applications for Linux, and HP works to alleviate this concern by working with its ISV partners. The current LC Series software solution stack for both Linux and Windows compute farms is shown in Table 2.

Table 2: LC Series Software Solution Stack

LC Series Software Portfolio			
Stack	Node	Linux	Windows
Grid	Control Node	Axceleon Enfuzion United Devices Grid Manager	Axceleon Enfuzion
	Compute Node	Axceleon Enfuzion RTE United Devices RTE	Axceleon Enfuzion RTE United Devices RTE
Job Management	Control Node	Altair PBS-PRO Platform Computing LSF	Platform Computing LSF
	Compute Node	Altair PBS-PRO RTE Platform Computing LSF RTE	Platform Computing LSF RTE
Development Tools/Compilers	Control Node	EI CxC Linux Development Pack PGI Linux Compiler Suite	EI CxC Windows Development Pack PGI Windows Compiler Suite
	Compute Node	EI CxC Linux Development RTE	EI CxC Windows Development RTE
Cluster Manager	Control Node	Scali Manage ROCKS (freeware)	Windows 2003 Domain Manager
	Compute Node	Scali Manager RTE ROCKS (freeware)	Windows 2003 Domain Manager
Operating Systems	Control Node	Red Hat Enterprise Linux 3.0 SUSE SLES 8	Windows 2003 SE Windows 2003 EE
	Compute Node	Red Hat HPC WS 3.0 SLES 8 HPC	Windows 2003 Server HPC Edition

Source: HP, June 2004

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

Pricing

Discussing the cost of a single node in a cluster does not take into consideration the discounts that a vendor might offer customers for clusters of 32 nodes, 64 nodes, etc. Vendors who discount their servers to compete with the pricing of another vendor (whose individual server costs are lower) generally make up their "losses" by overcharging for installation, warranties, and service and support, or they provide little in the way of installation assistance and support.

Pricing for individual nodes begins with IA-32-based servers. HP's DL140 configured with a single 3.06GHz Xeon processor, 512MB L2 cache, 1GB memory, and 80GB disk is priced at $1,468. This DL140 configuration is priced 44% lower than a comparably configured PowerEdge 1750 server, and a comparably configured IBM x335 server is slightly more expensive than the PowerEdge 1750.

A ProLiant DL360 G2 with dual 3.06GHz processors, 1MB L2 cache, 4GB memory, and a 36GB disk with two 10/100/1000 Ethernet ports is priced at $5,236. A PowerEdge 1750 with a comparable configuration is priced at $4,905; a RackSaver RS-1100 with the same configuration is priced at $4,499; and an IBM xSeries 335 with a comparable configuration is priced about $500 more than the DL360.

A DL145 with dual 2.2GHz Opteron processors, 1MB L2 cache, 2GB memory, and a 40GB ATA disk is priced at $4,298. A comparably configured IBM e325 Opteron server is priced several hundred dollars higher. RackSaver's RS-1164/OP is priced about the same as the DL145. Dell does not have an Opteron-based server, but it does offer the 64-bit Itanium 2-based PowerEdge 3250 as a compute node in its HPCC cluster offering. A PowerEdge 3250 with two 1.0GHz processors, 2GB memory, and a 36GB hard drive is priced at $7,499, about 74% more than the RackSaver RS-1164OP and a DL 145 with two 2.2GHz Opteron processors and 2GB memory.

Conclusions and Recommendation

The four Linux cluster products compared in this white paper are all good products. The suppliers basically use the same COTS components to build their Linux clusters. Dell and RackSaver compete with HP for cost leadership. But HP has a broader offering than either vendor and can offer customers a Linux cluster that meets their price/performance requirements.
Overall, HP and IBM separate themselves from Dell and RackSaver because they have much more experience in the HPC area as well as comprehensive installation and deployment capabilities. Dell developed a partnership with Cray a couple of years ago to do its cluster installations.

An important aspect of Linux clusters is cluster management. HRG in an independent study selected Scali Manage as one of the best all around Linux cluster management solutions available today. IBM's CSM along with IBM Director are also good cluster management solutions, but they are proprietary.

HP delivers LC Series clusters with Linux installed. According to a spokesperson for IBM, it does not yet install Linux. Customers order or purchase Linux separately and install it themselves. IBM, however, plans to begin installing Linux in the near future. Customers pay HP and IBM for clusters delivered, including OEM components (an exception occurs with respect to IBM if customers do not "pre-order" Linux; otherwise, they have to separately acquire Linux from Red Hat).

HP installs a range of partner driven applications. This provides customers with one-stop shopping and gives them confidence that the hardware and software stacks they purchase are interoperable. It also guarantees customers that they get a complete solution versus a bag of parts that need to be assembled.

Customers who desire a Linux cluster with Itanium 2 nodes can order it directly from HP using a standard ordering process (they would order a XC Series cluster). Customers who desire an Itanium 2-based Linux cluster from IBM must used a non-standard, special ordering process that may delay shipment of the cluster.

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

HP is the unchallenged leader in delivering HPC solutions and provides the most competitive pricing in the industry. And it is the clear leader in breadth of Linux cluster solutions offered. LC Series clusters are available to meet just about any customers' needs. They are configured, tested, and shipped ready to run out of the box. No special ordering process is required, regardless of the node architecture.

For those customers who desire a Linux cluster, HRG recommends that they start with HP and IBM. They have broader offerings and both companies have years of experience delivering clustered solutions for HPC. The real differentiator between these two companies is pricing. HP's server prices are consistently lower than IBM's prices for comparable Intel and AMD servers and either better or very competitive with Dell and RackSaver.

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

Harvard Research Group, Inc.
P.O. Box 297
Harvard, MA 01451-0297 USA

Tel. (978) 263-3399
(978) 263-0033
hrg@hrgresearch.com

Harvard Research Group is an information technology market research and consulting company. The company provides highly focused market research, consulting services, and business modeling tools to vendors and users of computer hardware, software, and services. For more information contact Harvard Research Group as follows:

Harvard Research Group™
P.O. Box 297
Harvard, MA 01451 USA
Tel. (978) 263-3399
Fax (978) 263-0033

E-mail: hrg@hrgresearch.com
http://www.hrgresearch.com