White Paper

DESKTOP TO DATA CENTER LINUX

www.novell.com

SUSE® Linux Enterprise Server for High Performance Computing

Novell®

Table of Contents:

What Is High Performance Computing?

The excellent scalability features of Linux, in addition to robust security and performance makes it an excellent choice for server systems, especially in the high performance computing area.

Organizations are increasingly turning to high performance computing (HPC) to deliver massive computing power. This power is required for challenges ranging from high-end, intensive scientific and engineering problems to commercial, data-intensive tasks.

HPC increases the performance of an application. This primarily involves dividing a particular job into as many parallel units as possible and running them simultaneously to achieve faster completion.

The term high performance computing refers to the use of parallel supercomputers and computer clusters, that is, computing systems made of multiple processors linked together in a single system with interconnects. The multiprocessing supercomputer operates on a single operating system (a function known as single system image or SSI) and uses multiple CPUs, where the application-level software is indifferent to the number of processors. The processors share tasks using symmetric multiprocessing (SMP) and Non-Uniform Memory Access (NUMA).

Clusters are a very good example of an HPC setup. A clustered HPC infrastructure has complex design and operation, often involving a large set of interdependent software elements that have to be precisely configured and seamlessly integrated across a growing number of compute nodes. This integration ensures mission-critical standards of performance, stability and reliability.

The hardware running the load must support HPC requirements. This means that the hardware must have the capability to scale; it must have very little latency for its network interconnections and must support resource virtualization. Applications must be parallel enough to be run on a cluster or grid to obtain faster results. They must also scale up well to cope with load increases.

Finally, the operating system is a crucial component of the HPC stack. Not only does it affect the performance, security and reliability of the systems, but it also has a major influence on the type of applications or hardware platform that is chosen.

In an HPC environment, the operating system must have features that enable efficient high performance computing. The operating system must be scalable and cluster friendly, and it must support a wide range of hardware platforms and devices. It must also support virtualization of resources and provide robust security features.

This is where Linux* enters the game. Linux has quickly emerged as a key operating system in the HPC market. The total cost of ownership (TCO) of a Linux system is significantly lower than the TCO of other popular operating systems. The excellent scalability features of Linux, along with robust security and performance, make it an excellent choice for server systems, especially in the clustering and HPC areas.

But first, let's take a closer look at the general evolution of HPC.

Evolution from Science and Engineering to Enterprise Business

Clustering independent, commodity-class machines—and building supercomputers out of them—was a controversial idea as recent as 10–15 years ago. For the last 15 years, HPC technologies have been mainly used

(and are still used) in areas such as academic research, fluid dynamics, oil and gas exploration, computer aided design and testing, and pharmaceutical and military research. The historic cost of HPC or "supercomputers" had limited their use to market segments that could afford these systems.

The evolution of both lower cost hardware and Linux has dramatically reduced the cost of these systems. Compute power has increased on a scale of one thousand times in just a few years, allowing many companies to now use the power of supercomputers in the form of an HPC Linux cluster.

According to IDC, many businesses can benefit today from the power of HPC technology, which is no longer limited to just the scientific market segments.

"In 2008, HPC server market revenue will vault past US$12 billion (2006: $10 billion). Clusters will make up more than 70 percent of this total, and cluster sales will shift even more notably toward blades. Systems in the Workgroup (below $50,000) and Departmental ($50,000 to $250,000) segments will account for more than 60 percent of server revenue—about five of every eight dollars spent. Worldwide, year-over-year server growth will average just under nine percent. Average worldwide growth for the HPC storage market will be higher than for servers, about 11 percent. Total revenue for HPC servers, storage and services will surpass $19 billion (not including costs for software applications, staffing, facilities and power)."[1]

Apart from traditional applications used by technical customers, there has been increased HPC adoption amongst commercial customers for applications such as financial analysis and portfolio management, digital security and surveillance, as well as decision support computing. HPC has been used in businesses with cluster-based supercomputers for data warehouses, line-of-business (LOB) applications and transaction processing.

And, while HPC has been largely limited to enterprises, research and development firms, and academic institutions, there is now a broad swath of midmarket companies adopting HPC because of the availability of economical solutions. The solutions, which were costly and proprietary in the past, have become affordable and open. Further, turnkey cluster building solutions are reducing the complexity of implementing HPC solutions, thus broadening adoptability.

The midmarket, pharmaceutical and biotech firms are hot markets for HPC solutions running on Linux. Recent work on the human genome has triggered growth in genomics and bioinformatics; bioinformatics research requires HPC and open source applications. HPC is also growing in the animation industry for computing, visualization and storage. Now even smaller organizations with small budgets that need HPC can afford it using commodity-class technology.

The Case for High Performance Computing with Linux

The past few years have seen significant changes in the high performance computing landscape (often referred to as high productivity computing). These changes are due, at least in part, to the emergence of open source and new clustering technologies. HPC for Linux is gaining ground at the cost of branded UNIX*.

A few years ago, UNIX variants such as AIX*, HP-UX*, Tru64 UNIX, Solaris* and IRIX ruled. Linux is now displacing many of these systems in the HPC space. Relatively suddenly (by market standards), Intel and AMD replaced RISC processors, Linux unseated UNIX as the dominant OS, and numerous second-tier vendors established positions in the market.

According to IDC many businesses can benefit today from the power of HPC technology, which is no longer limited to just the scientific market segments.

1 *Source: www.hpcwire.com/hpc/1987614.html*

Virtually every industry is adopting Linux clusters to attain the performance improvements needed to deliver on organizational goals.

The high-end SGI Altix 3000 system was the first 64-bit processor system running a fully supported, enterprise-grade Linux operating system, where certified applications for SUSE Linux Enterprise Server and Itanium processors did run unaltered on SGI Altix Servers.

2 "**Beowulf Clusters** *are scalable performance clusters based on commodity hardware, on a private system network, with open source software (Linux) infrastructure. The designer can improve performance proportionally with added machines. The commodity hardware can be any of a number of mass-market, stand-alone compute nodes as simple as two networked computers each running Linux and sharing a file system or as complex as 1024 nodes with a high-speed, low-latency network." (Source: www.beowulf.org/overview/index.html)*

By the late 1990s, computer scientists and engineers had begun to take Linux seriously as a platform for supercomputing. Linux and many of its associated applications and tools are available in the open source community. This fact, combined with the economics of commodity hardware, made Beowulf[2] clusters in inexpensive systems attractive test beds for deploying Linux and open source software in HPC environments.

These first-generation clusters were well suited for highly parallel algorithms with minimal inter-node communication; however, they offered limited capabilities for complex problems requiring large data sets or extensive interprocess communication.

Linux has steadily incorporated HPC features over the years and has become the primary OS for clustering and HPC deployments. Virtually every industry is adopting Linux clusters to attain the performance improvements needed to deliver on organizational goals. Seismic analysis for oil exploration; aerodynamic simulation for motor and aircraft design; molecular modeling for biomedical research; and data mining and financial modeling for business analysis all leverage HPC. Organizations are also adopting clusters based on Linux to ensure constant uptime, while still leveraging the flexibility, reliability and low cost of open source.

Even the large business and research agencies are using Linux for their HPC requirements because Linux on a cluster of x86 servers is more economical. Linux clusters have also become easy to set up and simple to manage. More importantly, there are a lot of resources available for HPC on Linux—many of them free.

The History of SUSE® Linux Enterprise in High Performance Computing

SUSE has always provided Linux code to the HPC market. Since 1993, SUSE engineers have made significant contributions to the advancement and tuning of the Linux kernel and key kernel-related performance technologies.

In the late 1990s and in early 2000, long before Novell acquired the company in January 2004, SUSE Linux AG was working closely with the HPC technology leaders. A historic landmark was SGI's announcement in 2003 to bundle SUSE Linux Enterprise Server 8 on its high-end Altix* 3000 servers and superclusters. The cooperation agreement was made with the intention to extend the Linux OS to new levels of scalability and performance. The Altix system was the first fully supported 64-processor system running a fully supported, enterprise-grade Linux OS, where certified applications for SUSE Linux Enterprise Server and Itanium* processors did run unaltered on SGI* Altix servers.

After the acquisition, that cooperation has continued. Virtually every HPC manufacturer uses SUSE Linux Enterprise or reviews it as a product option. Companies such as SGI, Penguin Computing, Appro International, Cluster Resources, IBM, Dell, and HP have all used the SUSE Linux Enterprise platform from Novell in their HPC solutions. At the International Supercomputing Conference 2004 in Heidelberg, Germany, Novell launched the High Performance Computing Linux Competency Network for Partners, which featured cooperations with companies like SGI, Fujitsu Siemens Computers, HP, IBM and Intel.

With SUSE Linux Enterprise Server 9, Novell released the first enterprise-class

Linux server built on the new Linux 2.6 kernel, which offered unmatched performance and scalability capabilities for large-scale Linux deployments. In consequence, in 2005 SUSE Linux Enterprise Server became SGI's strategic platform of choice. In addition, SGI's proprietary code for HPC enhancements called ProPack, which includes shared memory and scalability, was designed to run with SUSE Linux Enterprise Server 9 Service Pack 1 (SP1).

Also in 2005, Novell announced an integrated, tested and validated solution for HPC. Based on the HP* BladeSystem solution, SUSE Linux Enterprise Server, and additional components from Novell and HP partners, the solution constituted the first HPC offering under the Novell® Validated Configuration Program. The program allowed an application to receive certification when it installed and ran seamlessly with other validation suite solutions in an established HPC architectural setup. This program was designed for customers to take full advantage of HPC on Linux, making it easier to deploy a complex, clustered, Linux-based infrastructure with a flexible, fully tested, fully integrated solution backed by Novell and HP. And important HPC technology partners such as Cluster Resources, which validated Moab Cluster Suite and Moab Grid Suite, joined the SUSE Linux Enterprise and HP hardware validation to become an effective part of this cluster and grid solution.

The Success Story Continues

The success story in HPC continues with SUSE Linux Enterprise Server 10. In July 2006, shortly after the general availability of SUSE Linux Enterprise Server 10, SGI and Novell announced yet another world record in HPC: an SGI Altix 4700 system had achieved a sustained memory bandwidth of 4.35 terabytes (TB) per second in STREAM Triad benchmark tests. The feat was achieved on an SGI Altix 4700 system powered by 1,024 Intel* Itanium 2 processors running under a single copy of SUSE Linux Enterprise Server 10 with SGI ProPack* 5 for Linux. The configuration, which includes 4TB of system memory, was the largest single system image (SSI) attainable on a Linux OS system. This system is installed at the Leibniz Rechenzentrum in Munich, Germany, and is currently ranked 27 on the TOP 500 supercomputer list.[3]

SGI's Altix 4700/450 system incorporates the shared-memory NUMAflex* architecture. A unique feature of NUMAlink is that it is a shared memory, globally shared system interconnect. All physically distributed system memory is mapped into one shared addressable space. NUMAlink based Altix systems offer memory scalability to over 26 TB. SGI Altix platform incorporates the shared-memory NUMAflex architecture and supports up to 512 sockets of Dual-Core Intel Itanium Processors or 1024 cores under one instance of Linux and more then 26 TB of globally shared memory.

There are hundreds of universities and research institutions around the world that use SUSE Linux Enterprise Server in their research labs on HPC systems. The following examples show just a few of the successes in this area:

- *BSC-CNS[4] (Barcelona Supercomputing Center–Centro Nacional de Supercomputación) is the National Supercomputing Facility in Spain and was officially constituted in April 2005 by the Ministry of Education and Science of the Spanish Government, Generalitat de Catalunya (local Catalan Government) and the Technical University of Catalonia (UPC). The mission of BSC-CNS is to investigate, develop and manage information technology in order to facilitate*

"At SGI, our focus is on high-performance computing and robust scalability, and SUSE Linux Enterprise Server is the operating system of choice for many of our Altix and Altix XE customers. As the requirements for high-performance computing continue to grow more complex across industries, the collaboration between SGI and Novell ensures that SUSE Linux Enterprise will continue to be the leading operating system for high-performance clusters that meet those new business needs today."

Irene Qualters
Senior VP, Software
SGI

3 www.top500.org/list/2008/06/100
4 www.bsc.es/plantillaA.php?cat_id=200

"SUSE Linux Enterprise Server has improved our supercomputing capabilities, helping us, for example, to resolve pulsar signals in nano-seconds. We can now run our supercomputing environment on faster and cheaper machines."

Dr. James Murray
Cluster Manager and Research Scientist
Centre for Astrophysics and Supercomputing, Swinburne University of Technology

5 www.nas.nasa.gov/

scientific progress. With this aim, special dedication has been taken to areas such as Computational Sciences, Life Sciences and Earth Sciences. BSC-CNS manages MareNostrum, which is number seven of the most powerful supercomputers in Europe, and number 26 in the world, according to the TOP 500 list from June 2008. In November 2006, its capacity was increased because of the large demand of scientific projects. MareNostrum increased the calculation capacity of the supercomputer in 2006 from 4.812 processors to 10.240 processors with a final calculation capacity of 94.21 trillion floating operations per second (Teraflops), doubling its previous capacity (42.35 Teraflops).

- *Seoul National University (SNU) is South Korea's largest national university, as well as a leading research center. SNU decided to build a new supercomputer, based on SUSE Linux Enterprise Server, to enhance the functionality and scalability lacking in its existing supercomputers. SNU now can provide students, professors, corporations and research institutions across Korea with easy and reliable access to its supercomputing 24x7. SUSE Linux Enterprise Server offers a secure, reliable and flexible OS that meets the university's need for high performance, enterprise-class computing. Researchers and businesses in diverse fields, as well as SNU faculty members and students, currently use the supercomputer to perform extensive calculations that ordinary servers cannot run in subjects such as chemistry, biology, and earth and marine science. The new supercomputer showed actual performance of 5.618 Teraflops in benchmark tests conducted by TOP 500 Supercomputers (www.top500.org) immediately following the installation.*
- *Tokyo Institute of Technology is one of the world's leading science and technology universities, having received acclaim in both Japan and internationally for its outstanding achievements and high educational standards. By introducing SUSE Linux Enterprise Server, the university successfully developed a highly interoperable and scalable grid cluster system, called TSUBAME (Tokyo-tech Supercomputer and Ubiquitously Accessible Mass-storage Environment). In terms of performance and scalability, delivering 85 Teraflops of computing power sees TSUBAME rank number 24 in the world (based on June 2008 TOP 500 ranking) and number four in Asia among supercomputers. The system also features 655 nodes and 10,480 CPUs, the most of any PC cluster-based supercomputer in the world. The university plans to boost the system's performance to exceed 100 Teraflops, making it the most powerful supercomputer outside the United States.*

Universities and research centers are not the only institutions running SUSE Linux Enterprise Server on their HPC systems. Enterprises such as Boeing, NASA, Procter & Gamble, Motorola, Audi and many others use SUSE Linux Enterprise on supercomputers or in high performance environments to both design products and to test for design flaws and safety requirements.

- *NASA Advanced Supercomputing (NAS) Division[5] runs three of its supercomputers on SUSE Linux Enterprise. The* **RTJones Cluster** *(www.nas.nasa.gov/Resources/Systems/rtjones.html) is an SGI Altix-based supercomputer and one of multiple systems selected to evaluate next-generation technology to meet NASA's future supercomputing requirements. The* **Columbia Supercluster** *(www.nas.nasa.gov/Resources/Systems/columbia.html), also based on an SGI Altix system, makes it possible for NASA to achieve breakthroughs in science and engineering for the agency's missions and Vision for Space Exploration, and is also being made available to a broader national science and engineering community. The* **Cray Opteron Cluster** *(www.nas.nasa.gov/Resources/Systems/opteron.html), a 64-*

node cluster, is the latest addition to the set of high-end computing resources at the NAS facility. The system is structured with a single "head" node, which controls the system operation and launches programs, and 63 compute nodes.

- *MTU Aero Engines (MTU), a leading German manufacturer of engines and engine components for civilian and military aircraft, as well as stationary industrial gas turbines, has deployed SUSE Linux Enterprise Server on its high-performance CAE cluster.*
- *Audi AG, one of the world's leading prestige car manufacturers, runs its crash simulation cluster and the aerodynamics and CFD cluster on SUSE Linux Enterprise. Both of the new clusters derive part of their power from the optimization of SUSE Linux Enterprise Server for the Opteron* architecture. This helped Audi to achieve the required performance.*

SUSE Linux Enterprise Server: Leadership in Many Segments

SUSE Linux Enterprise Server is synonymous with high performance Linux running on 64-bit and mainframe systems. SUSE Linux Enterprise Server was the first Linux OS in the market to support 64-bit chip sets. From the x86 products of Intel and AMD, to the Intel Itanium architecture and the System z* mainframes of IBM, SUSE Linux Enterprise Server has been at the forefront of 64-bit processing.

This is one of the reasons that SUSE Linux Enterprise Server has been so successful in the HPC market. This technology segment was an early adopter of the 64-bit chip sets and has always pushed to be a leader in new technologies. Working on some of the world's most challenging problems the industry moved very quickly to the higher-powered chip sets.

SUSE Linux Enterprise Server became the Linux of choice because of its early support of the newer chip sets. When SUSE Enterprise Linux Server 9 shipped in 2004, Novell was clearly in the 64-bit Linux leadership position, and this drove the success of SUSE Linux Enterprise Server on HPC technologies.

Today multi-core processors are setting the stage for new market dynamics. The adoption of multi-core processors is not worth predicting—they're coming, like it or not—but it will have a domino effect that significantly changes the industry. Clusters expanded their share of the HPC market to 65 percent in 2007, according to IDC. In 2007, multi-core parallelism, new acceleration technologies and new low-end products stimulated end user interest and HPC revenue growth as scientists and engineers continued to expand their use of technical servers, especially standards-based clusters, to drive scientific research and product innovation.

That's why enterprises worldwide are turning to SUSE Linux Enterprise Server for high performance computing. SUSE Linux Enterprise Server has been designed to handle mission-critical workloads in the data center. It offers a highly scalable, high performance data center solution that comes with application security, virtualization and integrated systems management across a full range of hardware architectures—and it offers seamless interoperability with existing IT infrastructures.

With advanced memory management and new processor support, Native POSIX Thread Library (NPTL), and advanced multi-pathing and I/O capabilities, SUSE Linux Enterprise Server surpasses RISC/UNIX systems in performance and scalability for large-scale server deployments on commodity blades and servers as well as high-end mainframes.

The following facts illustrate why SUSE Linux Enterprise Server is the preferred OS for HPC:

- *SUSE Linux Enterprise Server supports new processor technologies, including*

"SUSE Linux Enterprise Server was the best operating system for us. In addition to delivering the superior performance necessary for our large-scale system, the platform is flexible enough to support a heterogeneous environment comprised of Windows, open systems, and numerous software packages, and it is easy to use as well. We also knew that a Novell Linux solution supports three of the world's most powerful supercomputers."

Professor Satoshi Matsuoka
Head of the Research and Education Infrastructure Department
Global Scientific Information and Computing Center,
Tokyo Institute of Technology

SUSE Linux Enterprise Server has been designed to handle mission-critical workloads in the data center. It offers a highly scalable, high performance data center solution that comes with application security, virtualization and integrated systems management across a full range of hardware architectures.

"SUSE Linux Enterprise Server gives us the flexibility to choose the most appropriate environment according to the business need. We have been running our HPC cluster on SUSE Linux Enterprise Server for six years now, and the environment has been 100 percent reliable."

Norbert Diehl
Head of IT
MTU Aero Engines

---

6 *As compared to SUSE Linux Enterprise Server 9. Source: In-house tests*

Quad-Core Intel Xeon and Quad-Core AMD* Opteron processors. Through drivers joint-engineered with chip vendors, SUSE Linux Enterprise Server enables multiple virtual machines to run varied data center workloads in native and Xen virtualized environments with outstanding performance, energy efficiency and reliability.*

- *SUSE Linux Enterprise Server performs up to 15 percent faster on systems with multi-core processors.*[6]
- *SUSE Linux Enterprise Server is massively scalable—on specific Itanium systems it scales up to 4,096 processors—and supports over 10 TB of physical memory.*
- *SUSE Linux Enterprise Server supports the latest network-acceleration technologies to gain network performance for I/O intensive applications (Intel I/O AT).*
- *SUSE Linux Enterprise Server runs all major hardware platforms, x86, x64, PowerPC, Itanium 2, and IBM* mainframe servers. This gives you freedom of choice when it comes to hardware acquisition.*
- *Today's enterprises are expected to be 24x7 operations. Maximizing uptime can mean the difference between winning and losing business. SUSE Linux Enterprise Server is built for mission-critical use and minimal downtime, maximizing the probability that systems will remain up and running. With Heartbeat 2, SUSE Linux Enterprise Server ships a fully integrated, multi-node, high-availability solution at no extra cost. The Heartbeat subsystem adds failover functionality to your system, allowing two Linux servers (a primary and a backup) to determine if the other is "alive." If the primary isn't functioning, Heartbeat sends failover resources to the backup. The HPI STONITH module of Heartbeat provides an extensible interface for remotely powering down a cluster node. The idea is quite simple: when the software running on one machine wants to make sure another machine in the cluster is not using a resource, it can pull the plug on the other machine. The integrated cluster software can easily be installed and configured by using the YaST setup tool.*
- *SUSE Linux Enterprise Server ships with file systems perfectly suited for large-scale environments. XFS is a high performance journaling file system, which originated on the SGI IRIX platform. It is completely multi-threaded and can support large files and file systems, extended attributes and variable block sizes, and it improves performance and scalability. Oracle* Cluster File System 2 (OCFS2), the only symmetrical parallel cluster file system to be accepted into the Linux Mainline Kernel, has been designed to host and perform on larger files in a clustered environment, making it a perfect fit for hosting virtual server disk images in a high availability configuration.*
- *Enterprises are continually seeking cost-effective storage methods. With the integration of new iSCSI technology, SUSE Linux Enterprise Server combines SCSI, Ethernet and TCP/IP to create simple, yet powerful high speed, low cost and long distance IP-based storage area networks.*

By looking at the TOP500.org Web site, you can see that SUSE Linux Enterprise is the Linux of choice on the world's largest HPC supercomputers. Of the top 500

supercomputers, more than 85 percent are running on Linux, while 65 of them are clearly identified to run a SUSE Linux Enterprise version—including the following[7]:

- *IBM eServer Blue Gene at Lawrence Livermore National Laboratory (LLNL) is a premier applied science laboratory that is part of the National Nuclear Security Administration (NNSA) within the United States Department of Energy (DOE). Housed in Lawrence Livermore National Laboratory's Terascale Simulation Facility, BlueGene/L (BGL) clocked 478.2 Teraflops on LINPACK, the industry standard measurement for HPC. Built by IBM, BGL is a workhorse supercomputer used to make science simulation of unprecedented detail for NNSA's tri-lab Advanced Simulation and Computing (ASC) Program possible. It leverages the computing expertise and resources of Sandia, Los Alamos and Lawrence Livermore national laboratories. Computer simulations are a cornerstone of NNSA's program to ensure the safety, security and reliability of the nation's nuclear deterrent without underground testing. Recently expanded to accommodate growing demand for high performance systems able to run the most complex nuclear weapons science calculations, BGL now has a peak speed of 596 Teraflops. In partnership with IBM, the machine was scaled up from 65,536 to 106,496 nodes in five rows of racks; the 40,960 new nodes have double the memory of those installed in the original machine.*
- *IBM eServer BlueGene/P, named JUGENE at Jülich Research Centre (Forschungszentrum Juelich or FZJ), was originally focused on nuclear research; however, more recently, it has become involved in a number of grid computing projects, such as UNICORE, DEISA and EUROGRID. It is also part of the European Gauss Centre for Supercomputing.*
- *The biggest supercomputer system installed at an industrial customer—at Total Exploration Production in France—is an SGI Altix ICE 8200 system running on SUSE Linux Enterprise with SGI ProPack. It ranks tenth based on a Linpack performance of 106.1 Tflop/s.*

Today most industries need to react quickly and efficiently to meet the demanding needs of customers and changes in the marketplace. Novell is well positioned to support the historic scientific and engineering community along with the emerging enterprise business needs of organizations which more and more rely on the use of HPC technologies also.

Real-time Linux in the High Performance Computing Market

In the HPC segment, systems are pushing for faster response times—especially in critical-needs environments—and end users are demanding real-time responses for their compute needs. Real-time Linux was driven first by military and scientific needs; today, however, customers in enterprises such as banking and Wall Street trading demand real-time support as well.

Novell, with its partner Silicon Graphics (SGI), began to look into and develop real-time aspects for the Linux kernel. Today SGI has real-time products that run on SUSE Linux Enterprise, and Novell has developed real-time add-ons for SUSE Linux Enterprise that allow for real-time response for the end users that need this level of support.

SUSE Linux Enterprise Real Time from Novell is an industry-standard, fully supported, real-time Linux OS for Intel and AMD-based multiprocessors. This enriched version of the SUSE Linux Enterprise kernel has been specifically engineered to reduce the latency and increase the predictability and reliability

Multi-core parallelism, new acceleration technologies, and new low-end products stimulated end user interest and HPC revenue growth as scientists and engineers continued to expand their use of technical servers, especially standards-based clusters, to drive scientific research and product innovation.

---

7 *Source: www.top500.org/lists/2008/06*
Most of the TOP500 supercomputers are running on "unspecified" versions of Linux. SUSE Linux Enterprise, by being named to run on 65 of them minimum, is clearly the leading operating system for supercomputers.

Real-time Linux in HPC was driven first by military and scientific needs; today, however, customers in enterprises such as banking and Wall Street trading demand real-time support as well.

of time-sensitive, mission-critical applications. It provides guaranteed performance in time-critical environments for hardware-in-the-loop simulations, data acquisition or process control. Key features—from the open source community's latest real-time patch set—include kernel preemption, CPU shielding and assignment, low latency and priority inheritance, and high-resolution timers. It also includes support for the latest open source InfiniBand* stack (OpenFabrics Enterprise Distribution 1.2.5), a switched fabric commodity interconnect that requires little processing overhead that helps you achieve sustainable, real-time performance on Linux.

You can also rely on capabilities that improve quality of service and enable a services-oriented infrastructure. These features include distributed shared memory, Xen* virtualization, and real-time cluster file systems. SUSE Linux Enterprise Real Time significantly enhances all of these features with its low-latency and deterministic processing.

SUSE Linux Enterprise Real Time for deterministic computing needs allows businesses to configure a modular, rich system that helps them consolidate servers, reduce costs and consistently achieve high quality of service.

Turnkey SUSE Linux Enterprise High Performance Computing Clusters

As the HPC segment becomes increasingly mass-market, tools that auto-install and auto-configure HPC clusters out of the box enable new users and new business partners to participate in the benefits of the HPC market. With third-party solutions, such as Moab Cluster Builder for SUSE Linux Enterprise from Cluster Resources Inc., customers and partners are able to deploy a cluster from a single DVD by installing SUSE Linux Enterprise Server and answering as few as four additional questions.

The Cluster Builder technology then prepares the head management node as well as the streamlined compute nodes, installing the software, configuring them to work together and even running validation tests to ensure the cluster is up and running. It completes all these tasks while the administrator does little more than power on the computers as directed.

SUSE Linux Enterprise Server is deployed, and so are resource management technologies, workload management, a job submission portal, cluster management and reporting software, message passing and other standard HPC components. Organizations simply need to add their applications and begin to use the fully deployed HPC cluster. This solution now makes HPC clusters based on SUSE Linux Enterprise the easiest to deploy and the easiest to use in the HPC market.

Interoperability and Hybrid High Performance Computing

Businesses are increasingly dependent on reliable information systems—and to get these reliable systems, they need reliable hardware and OSs. When these elements are combined, organizations can concentrate on providing high quality applications and services quickly. With so many OSs and hardware architectures available, however, they need to know that the choices they make today will not constrict the ability to support their business in the future.

And as Linux and Windows* become the two dominant platforms of the future, there will be an increasing need for these operating systems and the tools that manage them to work well together. Systems that lack well-developed interoperability capabilities can cause inefficiencies throughout the enterprise. For example, limited interoperability between Linux and Windows environments, in both physical and virtual instances, can lead to server sprawl. It can also lead to redundant management tools and inefficient use of IT staff. This translates as well for HPC; it is now clear that the two major platforms

used in the HPC market will also be Linux and Windows.

SUSE Linux Enterprise Server offers a complete open source platform that works seamlessly across hardware architectures and is the only Linux OS that has been optimized to work with Microsoft* Windows. From bidirectional virtualization to standards-based systems management and directory synchronization, interoperability is an intrinsic and differentiated capability that has been built into the fabric of the SUSE Linux Enterprise platform.

Thanks to the business and technical relationship with Microsoft, Novell is the only major Linux company to provide fully approved and licensed interoperability solutions. Many large enterprises see the need in their HPC platforms to use Linux or Windows based on load demand at a certain time. SUSE Linux Enterprise Server is the only Linux OS Microsoft will promote if a business needs to run and migrate mission-critical workloads across Windows and Linux platforms.

HPC clusters can represent a significant investment of company resources. They can be composed of tens, hundreds or even thousands of computers. By using the same cluster for two OSs, companies can gain the flexibility of another cluster without the hardware outlay, thereby maximizing the value of current investments and increasing savings.

There are a number of scenarios in which the ability to easily switch between SUSE Linux Enterprise Server and a Windows OS—which in HPC preferably would be Windows Compute Cluster Server (WCCS) 2003 or Windows HPC Server 2008 (HPCS)—on your HPC cluster is a major advantage. Moreover, if you have Linux-dedicated hardware and you want to develop WCCS skills, a dual-boot cluster can be very helpful. With a dual-boot cluster, you can try WCCS and become familiar with the user interfaces and commands without investing in a new set of hardware. A dual-boot environment also enables you to test usage scenarios without fully disengaging your existing operating system. You can observe and fine-tune your application on WCCS or HPCS while maintaining your existing environment.

As Linux and Windows become the two dominant platforms of the future, there will be an increasing need for these operating systems and the tools that manage them to work well together.

Organizations that build applications to run on both SUSE Linux Enterprise Server and on Windows Compute Cluster Server 2003 or Windows HPC Server 2008 can benefit from a dual-boot environment because it lets them test applications on both operating systems without investing in two sets of hardware. The ability to run applications on Linux and WCCS or HPCS on a single set of hardware can give companies additional flexibility without further investment.

Novell and Microsoft have set a new standard in addressing end user needs. All applications and systems are managed from consoles with which the end user is most familiar. Novell will continue to work with Microsoft to increase interoperability between the two platforms. This ability to load OSs on demand allows end users to meet their engineering, research and enterprise demands without having to worry about what OS their applications run on.

Through Novell partnerships, the hybrid HPC platform businesses can easily manage and dynamically adapt their computing needs. Cluster Resources, for example, offers the Moab Hybrid Cluster solution, which takes the dual-boot environment and dynamically adapts it from Windows to Linux and Linux to Windows based on the submitted workload. As the organization submits more workload

With third-party solutions such as Moab Cluster Builder for SUSE Linux Enterprise, customers or partners are able to deploy a cluster from a single DVD by installing SUSE Linux Enterprise Server and answering as few as four additional questions.

If your business needs to run and migrate mission-critical workloads across both physical and virtual Windows and Linux platforms, SUSE Enterprise Linux is the only Linux operating system Microsoft will promote in these environments.

There are a number of scenarios in which the ability to easily switch between SUSE Linux Enterprise Server and Windows Compute Cluster Server 2003 on your HPC cluster is a major advantage.

of one type the system dynamically adjusts the number of nodes booted into that OS so that the cluster is always running the OS it needs most at that point in time. This solution also can guarantee or reserve portions of the cluster to run a certain OS at certain times to meet known demand surges.

Beyond adaptation, the Moab solution unifies the workload submission experience for end users, unites the workload administration and provides an integrated reporting view of what is happening across both OS environments. This eliminates usability limitations, resolves capacity-planning issues and helps deliver on a true interoperability experience between Windows and SUSE Linux Enterprise.

Supported vs. Non-supported Linux

In the early days of high performance computing, there were very few standard distributions of Linux to work with. Many customers in the HPC segment developed a "build-it-your-self" approach. Linux distributions such as Debian*, Fedora*, and the predecessor of openSUSE®, SUSE Linux, were used to develop these early HPC systems.

Over time, however, many customers and solution providers realized that "free Linux" did not mean "at no cost." And they learned that using a commercial, maintained and supported distribution of Linux saved them money and allowed them to focus on solving the problem or building the product they had set out to accomplish.

One of the biggest values of a maintained Linux OS is the quality of the OS and the support customers can get from the company producing it. An IT organization that is operating a Linux platform without a support and maintenance ecosystem behind it is entirely dependent on the good will of the development community in case of technical difficulties. While this community is large, and while it is possible to subscribe to mailing lists, there are no guaranteed reactions or response times to submitted questions and problems. Furthermore, there is no guarantee that the answers to the questions are correct.

This situation harbors an incalculable risk. If critical questions are not answered quickly or correctly, systems can fail and data can be lost. To minimize this risk, users must safeguard the system by bringing in Linux specialists. This, in turn, always counteracts any cost advantage of a free Linux distribution.

Of course, there are also system specialists who offer maintenance contracts or support services for other Linux systems. However, these are usually small organizations that have only a few employees. How these companies are able to provide a full range of 24x7 support services remains unclear.

Training employees to become Linux specialists is a course of action that also harbors multiple risks. Training and continued education costs should be taken into account, and users must always have a clear plan of what will happen if their in-house Linux specialists leave the company, taking their knowledge with them.

However, the decisive factor is that in-house specialists cannot solve the problem of system compatibility. Level 3 support (where hardware and software vendors—even the original development teams—provide services in critical situations) is normally not available for a Linux system provided by smaller companies or communities, even if an IT

organization has access to a team of specialists. The reasons are obvious: since there is no coordination with hardware and software vendors, neither the development community nor the hardware or software vendor can provide warranties for those systems.

Novell provides broad support services from all relevant hardware vendors. The company has been working closely with leading hardware vendors for years, which affects, for example, the porting of hardware drivers to a Linux platform. As part of its Linux initiative, Novell maintains partnerships and close, cooperative relationships with hardware vendors such as SGI, HP, IBM, Fujitsu Siemens Computers, Dell and Sun.

Novell is strengthening its partnerships to provide even better value for its customers. Examples of these collaborative partnerships include the certifications hardware vendors provide for SUSE Linux Enterprise Server.

Certifications on or for an OS are essential for productive and secure operations. Only the complete certification of a solution stack—that is, integrated certifications for hardware, OSs and software—creates the type of data center security needed to run centralized applications. In addition, future software support and maintenance must not affect certified operating security and performance. In that respect, certified software maintenance is also an integral part of the solution stack.

On the one hand, certifications minimize the risk of system failures. On the other hand, certifications ensure that if a problem or disruption should occur, users can seek support from the source with the highest level of expertise: the original manufacturer.

Because most HPC systems are operating in environments that cannot fail or go down, having a major software company such as Novell to turn to has made a huge difference to many large corporations. SUSE Linux Enterprise Server comes with the full backing of the Novell worldwide ecosystem of technology and solution provider partners and the Novell award-winning support team. Novell brings its 25 years of enterprise-class support and product development to the HPC market and continues to build and provide HPC-ready products to the top HPC solution providers today.

Conclusion

Novell is a key player in the HPC segment. From its 64-bit technical leadership to its enterprise class support teams, Novell is prepared to help customers and HPC partners build the world's best high performance products and applications.

In cooperation with key technology partners such as SGI, HP, Fujitsu, IBM, Appro International, Atipa Technologies, Cluster Resources, ANSYS/Fluent and Penguin Computing, Novell delivers highly scalable and reliable solutions to enterprises and organizations that depend on HPC.

With its leadership in real-time Linux, Novell adds new dimensions to where and how HPC systems can be deployed. Because of the unique partnership with Microsoft, Novell is the best-positioned Linux company in the world to provide interoperability, innovation and support to the HPC market segment.

When you combine 25 years of Novell supporting the enterprise business needs with the expertise and success of SUSE Linux Enterprise Server in the HPC segment, Novell brings an unbeatable mix of knowledge and technical skill to the next generation of HPC use in the enterprise business segments.

Novell is strengthening its partnerships with all leading hardware vendors to provide even better value for its customers. Examples of these collaborative partnerships include the certifications hardware vendors provide for SUSE Linux Enterprise Server.

As most of the HPC systems are operating in environments that cannot fail or go down, having a major software company like Novell to turn to when needed has made a huge difference to many large corporations.

www.novell.com

Contact your local Novell Solutions Provider, or call Novell at:

1 800 714 3400 U.S./Canada
1 801 861 1349 Worldwide
1 801 861 8473 Facsimile

Novell, Inc.
404 Wyman Street
Waltham, MA 02451 USA

Novell®