White Paper

Enterprise Linux

SUSE® Linux Enterprise Server for High Performance Computing

Table of Contents

SUSE
We adapt. You succeed.

What Is High Performance Computing?

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 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.

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 become the 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 15–20 years ago. For the last 20 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 commercial companies to 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.

IDC expects the HPC technical server market revenues (not including costs for software applications, staffing, facilities and power) to grow at an 8.5 percent compound annual growth rate (CAGR) over the five-year forecast to reach revenues of roughly US$14 billion by 2017. What's more, IDC forecasts that HPC systems software will grow to US$1.5 billion in 2017.

The brightest spot in HPC revenues remains the high-end Supercomputer segment (HPC systems sold for US$500,000 and up). However, the lower segments with systems selling for under US$250,000 (Departmental segment in the US$100,000 to US$249,000 price band) are seeing good growth. One of the fastest-growing HPC market segments is storage, with a revenue of US$4.1 billion in 2012, and expected US$6.0 billion in 2017.[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.

1 *Source: IDC's Top 10 HPC Market Predictions for 2014, IDC Web Conference 11 February 2014*

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.

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.

A few years ago, UNIX variants such as AIX, HP-UX, Tru64 UNIX, Solaris and IRIX ruled. However, Linux has displaced 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.

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.

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:* **ftp://ftp.sjtu.edu.cn/sites/www.beowulf.org/overview/index.html***)*

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

Linux has steadily incorporated HPC features over the years and has become the primary OS for clustering and HPC deployments. Virtually every industry has adopted 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.

Already back in the late 1990s, SUSE started its close cooperation 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.

Today, virtually every HPC manufacturer uses SUSE Linux Enterprise or reviews it as a product option. Companies such as SGI, Cray, NEC, IBM, Dell and HP have all used the SUSE Linux Enterprise platform in their HPC solutions.

With SUSE Linux Enterprise Server 9, SUSE released the first enterprise-class Linux server built on the 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 included shared memory and scalability, was designed to run with SUSE Linux Enterprise Server 9 Service Pack 1.

Also in 2005, SUSE announced an integrated, tested and validated solution for HPC. Based on the HP BladeSystem solution, SUSE Linux Enterprise Server, and additional components from SUSE and HP partners, the solution constituted the first HPC offering under the 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 SUSE and HP.

To help customers reduce the complexity and risk associated with buying an HPC cluster solution, SUSE joined the Intel Cluster Ready program a few years back. Developed in conjunction with hardware and software vendors, the Intel Cluster Ready program is designed to simplify purchasing, deployment, and management of HPC clusters. SUSE Linux Enterprise Server is powering many of the certified Intel Cluster Ready systems.

The Success Story Continues

The success story continued with SUSE Linux Enterprise Server versions 10 and 11. As an example, working with SUSE and IBM, Leibniz Rechenzentrum (LRZ), the central computing center for Munich's universities and other academic institutions, implemented SuperMUC, which in 2012 became Europe's fastest supercomputer (ranked 4th in the Top 500 list of June 2012)[3]. As part of the Gauss Center for Supercomputing (GCS), which operates the most powerful high performance computing infrastructure in Europe, LRZ offers supercomputer resources to a large scientific community. The institution wanted to provide researchers across Europe with a reliable and powerful high performance computing platform, which would enable users to make faster progress in their complex research projects. SuperMUC consists of approximately 9,400 general-purpose computing nodes with a peak performance of three Petaflop/s. The system comprises more than 155,000 Intel Xeon processor cores and more than 300 TB main memory, and is connected to disk storage systems with a total capacity of 11.5 PB. LRZ chose to run its high performance computing system SuperMUC with SUSE Linux Enterprise Server, leveraging SUSE's proven expertise in high performance computing and leading automation tools such as AutoYaST, which allows systems to be installed without manual intervention.

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.

3 **http://top500.org/lists/2012/06/**

"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 customers. As the requirements for high-performance computing continue to grow more complex across industries, the collaboration between SGI and SUSE 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

To date, 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 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 one of the most powerful supercomputers in Europe.*
- *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 more than 5.6 Petaflops of peak computing power sees the current HPC system TSUBAME 2.5 rank number 11 in the world (based on November 2013 Top 500 ranking). The system features over 1,400 nodes and more than 74,000 CPUs, and is among the three most powerful supercomputers in Asia.*
- *The Pittsburgh Supercomputing Center (PSC) competes to support scientific research from fluid dynamics to climate modeling and genomics. PSC won a National Science Foundation grant with a SUSE Linux Enterprise Server-based SGI UV 1000 cache-coherent shared-memory system. The system now hosts 1,316 users and 373 research projects at universities across the United States with unparalleled ease of use for rapidly testing new ideas.*

Universities and research centers are not the only institutions running SUSE Linux Enterprise Server on their HPC systems. Organizations and enterprises such as NASA, Boing, BMW, 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] operates NASA's High-End Computing Capability Project (HECC), which provides world-class computing resources to users across the nation to accelerate the development of innovative technologies, ensure new scientific discoveries, develop complex engineering systems, and reduce risks in support of NASA missions. Pleiades Supercomputer (***www.nas.nasa.gov/hecc/resources/pleiades.html****), represents NASA's state-of-the-art technology for meeting the agency's supercomputing requirements, enabling NASA scientists and engineers to conduct modeling and simulation for NASA missions. This distributed-memory SGI ICE cluster is connected with InfiniBand in a dual-plane hypercube technology, and runs on SUSE Linux Enterprise Server 11.*
- *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.*

4 **www.bsc.es/**
5 **www.nas.nasa.gov/**

"SUSE Linux Enterprise Server provides the stable foundation that we need to ensure optimal performance from our IT infrastructure, and one that will support us as we continue to expand our high-performance computing capabilities."

YONG PAN
Deputy Manager of Operations and Maintenance
Beijing Computing Center

- *With the Pangea system at Total Exploration, France, SGI developed the fastest supercomputer used in an industrial environment. SGI's close partnership with SUSE enabled Total to use SUSE Linux Enterprise Server as the operating system for Pangea that uniquely combines high scalability, flexible architecture and the ability to handle large volumes of CPUs.*

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, SUSE 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 have set the stage for new market dynamics. The adoption of multi-core processors has had a domino effect that significantly changed the industry. Clusters expanded their share of the HPC market to 65 percent, and 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 11 has been the first Linux operating system to include the Linux 3.x kernel, which enhances the overall performance by 20 percent. The performance of memory and compute intensive workloads is increased through support for transparent huge pages. Support for transparent per-CPU load balancing on multi-queue devices and faster packet filtering improve network*

"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."

PROFESSOR SATOSHI MATSUOKA

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

performance. Enhanced control groups (I/O throttling and memory cgroup controller optimization) help you get the most out of your systems.

- *SUSE Linux Enterprise Server supports new processor technologies. Through drivers joint-engineered with chip vendors, SUSE Linux Enterprise Server enables multiple virtual machines to run varied data center workloads in native and virtualized environments with outstanding performance, energy efficiency and reliability.*
- *SUSE Linux Enterprise Server is massively scalable—on several systems it scales up to 4,096 logical and supports 16 TB of physical memory (certified).*
- *SUSE Linux Enterprise Server supports 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, x86-64, POWER, Itanium, and IBM System z. 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.*
- *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. Btrfs (B-tree file system) provides for improved services availability and data integrity with copy on write functionalities, integrated volume management and checksums. New snapshot and rollback capabilities offer improved resilience.*
- *Enterprises are continually seeking cost-effective storage methods. With the integration of 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.*
- *SUSE Linux Enterprise Server includes Open Fabrics Enterprise Distribution (OFED) which enables affordable, high-speed remote direct memory access (RDMA) capable fabrics and unified interconnects based on InfiniBand and 10 Gigabit Ethernet. Support for Fiber Channel over Ethernet (FCoE) allows existing fiber channel storage to be accessed using an Ethernet adapter. Data Center Bridging (DCB) enhancements allow network traffic with differing requirements to operate and co-exist on Ethernet networks.*

6 *Source:* **http://top500.org/lists/2013/11/**

"We have relied on SUSE Linux Enterprise Server in high performance computing for approximately 15 years, and have always been very satisfied with the operating system. Specific features such as its early support of the Andrew File System distinguished SUSE Linux Enterprise Server from the competition."

HERBERT HUBER
Division Head of Super Computing
Leibniz Supercomputing Centre

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 96 percent are running on Linux, and many of them run a SUSE Linux Enterprise version—including SuperMUC at Leibniz Rechenzentrum, Titan at Oak Ridge National Laboratory, Pleiades at NASA/Ames Research Center, TSUBAME 2.5 at Tokyo Institute of Technology and Pangea at Total Exploration in France.[6]

Today most industries need to react quickly and efficiently to meet the demanding needs of customers and changes in the marketplace. SUSE 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.

SUSE, 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 SUSE 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 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 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 OpenFabrics Enterprise Distribution, a switched fabric commodity interconnect that requires little processing overhead that helps you achieve sustainable, real-time performance on Linux.

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.

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.

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, SUSE 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 HPC Server—on your HPC cluster is a major advantage. Moreover, if you have Linux-dedicated hardware and you want to develop Windows HPC Server skills, a dual-boot cluster can be very helpful. With a dual-boot cluster, you can try Windows HPC Server 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 Windows HPC Server while maintaining your existing environment.

Organizations that build applications to run on both SUSE Linux Enterprise Server and on Windows HPC Server 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 Windows HPC Server on a single set of hardware can give companies additional flexibility without further investment.

SUSE 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. SUSE will continue to work with Microsoft to increase interoperability between thetwo 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.

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-yourself" 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.

SUSE 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, SUSE maintains partnerships and close, cooperative relationships with hardware vendors such as SGI, HP, IBM, Fujitsu and Dell.

SUSE 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 SUSE to turn to has made a huge difference to many large corporations. SUSE Linux Enterprise Server comes with the full backing of our worldwide ecosystem of technology and solution provider partners and award-winning support team. SUSE brings its 20+ 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

SUSE is a key player in the HPC segment. From its 64-bit technical leadership to its enterprise class support teams, SUSE 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, Dell, HP, Fujitsu, IBM, Cray and NEC, SUSE delivers highly scalable and reliable solutions to enterprises and organizations that depend on HPC.

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

When you combine 20+ years of SUSE supporting the enterprise business needs with the expertise and success of SUSE Linux Enterprise Server in the HPC segment, SUSE brings an unbeatable mix of knowledge and technical skill to the next generation of HPC use in the enterprise business segments.

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

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

1 800 796 3700 U.S./Canada
1 801 861 4500 Worldwide

SUSE
Maxfeldstrasse 5
90409 Nuremberg
Germany

www.suse.com