

White Paper



Linux and High Performance Computing

HariPrasad Nellitheertha

Introduction

Operating systems (OS) are a crucial component of the enterprise software stack. Not only do they have implications on performance, security and reliability of the systems, they also have a major influence on the type of applications or platform that is chosen. Technically, operating systems act as an interface between the hardware and the applications handling the communication and transfer of data between the two. They also manage the hardware resources and help in ensuring that all users of the system get a fair share of the resources to operate.

There are and have been several popular OS with a widespread user base. Some of the popular OS over the years include DOS, Windows (all the flavors – XP, NT, 2000 etc), Linux, UNIX (and all its flavors), OS/2, MacOS and others. For special applications, such as embedded devices and devices with real time requirements, there are other OS that are being used. Some common features that most general purpose OS provide include the ability to multitask (where more than one task or process can be actively running on the machine), ability to handle multiple processors, file systems for systematic data access and storage and networking facilities to communicate with other machines. Naturally, each OS has its own set of features that differentiates it from others.

One OS that is gaining a lot of support is Linux. Started as a part time project by a Finnish student in 1991, it has a developer base of hundreds of programmers today. Industry giants such as IBM, HP and Intel have committed their resources behind it. Linux is a feature rich, general purpose operating system that can also be customized to suit special purposes such as real time systems or embedded devices. Released under the GNU General Public License (GPL) scheme, this open source product has seen a tremendous growth in the market and is challenging the leading OS in the desktop and server markets. This paper provides a brief introduction to Linux and discusses its various features. It also highlights the features that make Linux an excellent choice for high performance computing.

Linux

The term Linux actually refers to only the operating system kernel that is being developed by an open source community led by Linus Torvalds. An operating system is packaged around the Linux kernel by adding on applications, user interfaces and various other functions. The correct technical term for an operating system built around the Linux kernel would be GNU/Linux. Examples of operating systems built on the Linux kernel are SUSE Linux 9, Red Hat Enterprise Linux, Debian and others. Throughout this paper, the term “Linux” is used interchangeably to mean both a kernel and an operating system (the difference is highlighted when required).

The Linux kernel was started as a college project by Linus Torvalds in 1991, as an experiment in place of the Minix operating system. The release of Linux 1.0 happened in 1994 and it supported only Intel 386 processor based machines. Linux 1.2, released in 1995, supported a wider variety of platforms such as Alpha, SPARC and MIPS. Linux 2.0 came out in 1996 and supported many more architectures along with support for symmetric multiprocessors (SMP). Major releases of the Linux kernel come out roughly once every two years and Linux 2.4 which came out in 2001 helped the growth of Linux with its support for ISA Plug and Play, USB and other options. Linux 2.6 which came out in 2003 has numerous improvements over Linux 2.4 with much better scalability, better scheduling algorithms, many more file systems and other new features. Till the release of the 2.6 kernel, the development model of the kernel consisted of an odd numbered series of releases (2.1, 2.3, 2.5 and so on) which constituted the development tree containing the latest (but unstable) set of features, and an even numbered series (2.2, 2.4, 2.6) which made available a stable release for production use. Post the release of the 2.6 kernel both the stable and development trees have been merged.

While the kernel itself is developed by an active open source community with members from various corporations, universities and freelance developers working together, many companies package the kernel along with other necessary software and bring out operating systems based on the Linux kernel. This includes organizations like Red Hat, Novell, Debian, Mandrake etc. These companies too follow the open source model and make available a free version of their package and source code for download and use. Support and service, of course, comes with a fee and forms the basis for the revenue model for these organizations.

Features of Linux

The following technical features are available on most Linux distributions featuring the Linux 2.6 kernel. Some of them, however, may not be available on particular distributions. No attempt is being made here to list down distribution-wise features.

- Excellent multitasking capabilities: several programs can run at the same time.
- Multi-user capabilities: more than one user can be logged on to the same system at the same time
- Multi-platform: Linux runs on a plethora of platforms including PowerPC, Alpha, MIPS, AMD 64, Intel Itanium, SPARC, zSeries apart from, of course, the i386 platform
- Linux supports multi-processor systems. Systems with many tens of processors can run Linux smoothly and can extract performance benefits.

- Linux has built-in support within the kernel to run multiple independent threads of the same process.
- Demand load executables: read from disk only those parts of the program that are required for execution.
- Shared copy-on-write pages among executables. This means that multiple processes can use the same memory for running their code. Only upon an attempt to write to one of those pages is a separate copy made.
- Virtual memory using paging to disk.
- A swap subsystem which can use either a separate partition or a file on partition as the swapping destination
- Core dumps upon system failure, for use by a program debugger or a dump analyzer for servicing.
- Licensed under the open source GPL license model. This means that the entire source code of the Linux kernel is available for use and improvement to the public.
- Multiple virtual consoles: several independent login sessions.
- Support for numerous file systems including ext2, ext3, Reiserfs, JFS, XFS, HFS, Vfat, HPFS, UFS, CDFS and others.
- TCP/IP networking including FTP, Telnet, NFS etc
- Support for numerous networking protocols including TCP, IPV4, IPV6, IPX and DDP

Linux and High Performance Computing

High Performance Computing (HPC) deals with methods to increase the performance of an application. Primarily this involves dividing a particular job into as many parallel units as possible and running them simultaneously to achieve faster completion. Clusters are a very good example of an HPC setup. To have a good HPC setup, support is required from various layers. The hardware that is running the load needs to support HPC requirements. This means that the hardware must have the capability to scale; it must have very little latencies 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 very well with increase in load. Finally, the operating system must also have features that enable efficient HPC. The OS must be scalable and must be cluster friendly. The OS must support a wide variety of platforms and devices. It must also support virtualization of resources and robust security features.

Linux has steadily incorporated HPC features over the years and is a favorite for clustering and HPC deployments. It provides a suite of features that enables HPC applications to run very efficiently. We discuss below the various features available in Linux that make it especially suited to a high performance computation scenario.

Scalability features

- Supports the new features by Intel on 32 bit processors such as Physical Address Extension (PAE) which allows operating systems to access up to 64GB of RAM
- Number of unique users and groups that can be supported is up to 4 billion
- Number of process IDs (PIDs) is up to 1 billion
- File systems of up to 16TB in size are supported
- Supports up to 4095 major devices (device classes such as SCSI disks, Ethernet cards, IDE disks, keyboards, etc) and up to 1 million sub devices under each category.
- Advanced process scheduler which employs the O(1) algorithm that provides guaranteed performance even when the system is loaded with thousands of processes.
- Features such as e-poll which enables a Linux server to handle hundreds of nodes running any network intensive application Platform and Device Support

Platform and Device Support

- Supports Non Uniform Memory Access (NUMA) machines and high end Symmetric Multi Processor (SMP) machines
- Supports multiple architectures including Intel Itanium, AMD 64bit processors, IBM PowerPC, SPARC, ARM etc.
- Supports Intel Hyper Threading technology
- Support for dual and multi core processors
- Support for all the latest (and fast) network interconnects such as Infiniband, Gigabit Ethernet and Fiber Channel.

These interconnects enable scale-out of clusters and Grids with very little latencies

Clustering Support

- Linux is a very powerful platform when it comes to clustering. There are hundreds of Linux clusters around the world running successfully.
- Linux has very good technical features that support clustering. Linux supports clusters of up to many hundreds of nodes. The Linux kernel provides for high performance and very low latencies, making it very suitable for high performance computing applications.
- Features like the Read Copy Update (RCU), per-CPU counters, fine grained locking and other Linux Scalability Effort (LSE) features have reduced latencies to a great extent.
- Linux offers a wide variety of clustering file systems such as Oracle Cluster File System (OCFS), Lustre, IBM's General Parallel File System (GPFS), GFS and others. These work on distributed and shared storages and make cluster implementations much easier and efficient.
- As mentioned earlier, support for multiple low latency network interconnects and support for SAN/NAS make Linux very suitable for clustering solutions, especially database clustering solutions.

Virtualization Support

- Support for multiple operating system virtualization tools such as Xen, Linux VServer and Virtuozzo.
- Linux based virtualization tools also support non-Linux guest operating systems such as FreeBSD, NetBSD and Windows thus suiting migration scenarios very well.
- Linux also excellent support for Storage Area Networks (SAN) that will enable shared storage across a cluster or Grid.

Security

- Linux has very robust security features. It has far fewer virus attacks and spyware intrusions than other operating systems.
- The open source development model of Linux has ensured that the quality of the Linux code is very high. This has led to fewer bugs and fewer security loopholes.
- Linux has a very active developer community, along with support from vendors. It also has a powerful set of debugging and serviceability tools including debuggers, crash dump tools, tracing and logging mechanisms and probing tools. This means that the turnaround time for bugs and security issues is very quick. Thus the window for exploitation by malicious users and attacks is small.
- Linux has UNIX style userid and password protection and users have limited privileges. Non-root users cannot compromise the security of the system.
- Security Enhanced Linux (SE-Linux) system which is being developed by the National Security Agency is a big step towards further enhancing Linux security. It implements features such as access control to kernel subsystems, enforcement of separation of information based on confidentiality, control of threats of tampering and by-passing of application security and confinement of damage caused by malicious programs. These features are getting into regular distributions as well and are making Linux robust from the security point of view.

Cost

- Linux, being an open source product, is a very cost effective solution.
- The total cost of ownership of a Linux box is significantly lower than that of other popular operating systems. This makes Linux a very good choice for scale-out architectures like clusters or Grids where hundreds of nodes need to be deployed.
- There have been comparisons¹ that show that for a new setup, Linux is up to 26% cheaper than a similar Windows offering. Similarly, for pre-existing hardware, Linux can be up to 36% cheaper.

¹“Linux vs. Windows, Total Cost of Ownership Comparison”, Cybersource Pty Ltd

Observations

Linux offers a wide variety of options in almost every category of software. Linux is highly customizable and also has excellent technical features. These reasons are contributing to the growth of Linux in server and other markets.

- Linux promises excellent scalability features
 - Supports 512 (and more) processors in single systems
 - Supports up to 16TB of hard disk storage
 - Supports tens of hardware platforms
 - Multiple platforms supported on 64 bit systems (Itanium, PowerPC, AMD etc) thus suitable for high performance computing applications.
- Linux also supports a plethora of devices including newer high speed network interconnects such as Infiniband and Gigabit Ethernet.
- Linux also supports a plethora of devices including newer high speed network interconnects such as Infiniband and Gigabit Ethernet.
- Linux offers many virtualization solutions at the OS level and at the storage level.
- The above features make Linux an excellent platform for scale-out architecture for High Performance Computing and Grid applications.
- Linux offers excellent performance for clustering applications. Several technical features aid the clustering ability of Linux.
- Linux is superior in terms of security. It is not possible to configure systems to log in without user IDs and passwords. Root privileges are required in Linux to access and modify system data. There are far fewer viruses and spyware programs for Linux when compared with other operating systems.
- Far fewer bugs have been seen on Linux over the years. Also, due to the open source nature of support, bug fixing turnaround time is much lesser in case of Linux.
- The total cost of ownership over a period of few years is lesser with Linux than Windows. This includes cost of acquisition (hardware and software), installation, configuration, maintenance, administration, training and others.

Conclusion

Linux has emerged as a key operating system competing for market share in the desktop and server market. The total cost of ownership of a Linux system is significantly lower than other popular operating systems. The excellent scalability features of Linux along with robust security and performance makes it an excellent choice for server systems, especially in the clustering and high performance computing areas.

References

1. "Linux vs. Windows, Total Cost of Ownership Comparison", Cybersource Pty Ltd
2. "Linux vs. Windows, A comparison of Linux and Windows", <http://www.michaelhorowitz.com/Linux.vs.Windows.html>
3. "Linux vs. Windows: Why Linux will Win", <http://librenix.com/?inode=6565>
4. "Linux Features", <http://www.tldp.org/HOWTO/INFOSHEET2.html>
5. "An Overview of Red Hat Enterprise Linux 2.1 Reliability, Availability, Scalability and Manageability (RASM) Features", Red Hat White Paper.
6. "The Wonderful World of Linux 2.6", <http://www.kniggit.net/wwol26.html>
7. "SUSE® LINUX Enterprise Server 9 Technical Feature List", A Novell SUSE document.
8. "Security Enhanced Linux", <http://www.nsa.gov/selinux/T>



For more information, contact askus@infosys.com

About Infosys

Many of the world's most successful organizations rely on Infosys to deliver measurable business value. Infosys provides business consulting, technology, engineering and outsourcing services to help clients in over 30 countries build tomorrow's enterprise.

For more information about Infosys (NASDAQ:INFY), visit www.infosys.com.