

Integration Guide for New Dual-Core Intel® Xeon® Processor-Based Servers (or Workstations)

With the Linux* Operating System

Rev 1.0



Revision History

| Date | Revision | Notes | Product Code |
|-------------|-----------------|-----------------|---------------------|
| 2006.05.10 | Rev 1.0 | Initial Version | |
| | | | |
| | | | |

Disclaimers

The information contained in this document is provided for informational purposes only and represents the current view of Intel Corporation ("Intel") and its contributors ("Contributors") as of the date of publication. Intel and the Contributors make no commitment to update the information contained in this document, and Intel reserves the right to make changes at any time, without notice.

DISCLAIMER. THIS DOCUMENT, IS PROVIDED "AS IS." NEITHER INTEL, NOR THE CONTRIBUTORS MAKE ANY REPRESENTATIONS OF ANY KIND WITH RESPECT TO PRODUCTS REFERENCED HEREIN, WHETHER SUCH PRODUCTS ARE THOSE OF INTEL, THE CONTRIBUTORS, OR THIRD PARTIES. INTEL, AND ITS CONTRIBUTORS EXPRESSLY DISCLAIM ANY AND ALL WARRANTIES, IMPLIED OR EXPRESS, INCLUDING WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR ANY PARTICULAR PURPOSE, NON-INFRINGEMENT, AND ANY WARRANTY ARISING OUT OF THE INFORMATION CONTAINED HEREIN, INCLUDING WITHOUT LIMITATION, ANY PRODUCTS, SPECIFICATIONS, OR OTHER MATERIALS REFERENCED HEREIN. INTEL, AND ITS CONTRIBUTORS DO NOT WARRANT THAT THIS DOCUMENT IS FREE FROM ERRORS, OR THAT ANY PRODUCTS OR OTHER TECHNOLOGY DEVELOPED IN CONFORMANCE WITH THIS DOCUMENT WILL PERFORM IN THE INTENDED MANNER, OR WILL BE FREE FROM INFRINGEMENT OF THIRD PARTY PROPRIETARY RIGHTS, AND INTEL, AND ITS CONTRIBUTORS DISCLAIM ALL LIABILITY THEREFOR.

INTEL, AND ITS CONTRIBUTORS DO NOT WARRANT THAT ANY PRODUCT REFERENCED HEREIN OR ANY PRODUCT OR TECHNOLOGY DEVELOPED IN RELIANCE UPON THIS DOCUMENT, IN WHOLE OR IN PART, WILL BE SUFFICIENT, ACCURATE, RELIABLE, COMPLETE, FREE FROM DEFECTS OR SAFE FOR ITS INTENDED PURPOSE, AND HEREBY DISCLAIM ALL LIABILITIES THEREFOR. ANY PERSON MAKING, USING OR SELLING SUCH PRODUCT OR TECHNOLOGY DOES SO AT HIS OR HER OWN RISK.

Licenses may be required. Intel, its contributors and others may have patents or pending patent applications, trademarks, copyrights or other intellectual proprietary rights covering subject matter contained or described in this document. No license, express, implied, by estoppels or otherwise, to any intellectual property rights of Intel or any other party is granted herein. It is your responsibility to seek licenses for such intellectual property rights from Intel and others where appropriate.

Limited License Grant. Intel hereby grants you a limited copyright license to copy this document for your use and internal distribution only. You may not distribute this document externally, in whole or in part, to any other person or entity.

LIMITED LIABILITY. IN NO EVENT SHALL INTEL, OR ITS CONTRIBUTORS HAVE ANY LIABILITY TO YOU OR TO ANY OTHER THIRD PARTY, FOR ANY LOST PROFITS, LOST DATA, LOSS OF USE OR COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF YOUR USE OF THIS DOCUMENT OR RELIANCE UPON THE INFORMATION CONTAINED HEREIN, UNDER ANY CAUSE OF ACTION OR THEORY OF LIABILITY, AND IRRESPECTIVE OF WHETHER INTEL, OR ANY CONTRIBUTOR HAS ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES. THESE LIMITATIONS SHALL APPLY NOTWITHSTANDING THE FAILURE OF THE ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

Intel, Intel NetBurst, Intel Xeon, Intel i960, Intel Inside and the Intel logos are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright © Intel Corporation, 2002. All rights reserved.

*Other brands and names are the property of their respective owners.

Table of Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 4 |
| 1.1 | Document Overview | 4 |
| 1.2 | Target Audience | 4 |
| 1.3 | Document Objective | 4 |
| 2 | Platform Overview | 5 |
| 3 | Platform Hardware Requirements | 7 |
| 3.1 | Motherboard Integration | 7 |
| 3.2 | Fully Buffered DIMM (FBDIMM) Memory | 8 |
| 3.3 | Processor and Thermal Solutions | 8 |
| 3.4 | Server System Infrastructure (SSI) Compliant Chassis | 13 |
| 3.5 | SSI Compliant Power Supply | 14 |
| 3.6 | Other System Components | 16 |
| 4 | Recommended Server Components | 19 |
| 5 | Setup and Configuration of Red Hat® Enterprise Linux® 4 | 20 |
| 6 | Setup and Configuration of VMWare® Server (Linux® Version) | 39 |

1 Introduction



1.1 Document Overview

This document outlines the procedures for deploying a new Dual-Core Intel® Xeon® processor-based server (or workstation) with Red Hat* Enterprise Linux* 4 Update 2.

The process of properly integrating a new server (or workstation) system typically requires access to several source documents, each containing pertinent information. The purpose of this document is to extract the essential information needed from the various sources and place it into one comprehensive document.

See the References section of this document for the complete list of source documents used.

In addition this document will provide any useful “tips and tricks” that may have been discovered during the development of this system.

1.2 Target Audience

- Intel Channel program members who currently integrate Intel® Server Products or are new to server platforms
- System integrators who deploy the new Dual-Core Intel® Xeon® processor-based servers (or workstations)

1.3 Document Objective

Our intent is to better enable system integrators in meeting the competitive challenges they face in the server market and to keep program members up-to-date on emerging server technologies. By following the steps outlined in this document you will be able to deploy a stable new Dual-Core Intel® Xeon® processor-based server (or workstation) using Intel components running Red Hat* Enterprise Linux* 4 Update 2 in a timely and effective manner.

2 Platform Overview

The new platform consists of the Dual-Core Intel® Xeon® processor 5000 sequence, Intel® 5000P or Intel® 5000V or Intel® 5000X Memory Control Hub, and the Intel® 6311ESB or 6312ESB I/O Controller Hub. The new Dual-Core Intel® Xeon® processor-based server enhances platform performance, bandwidth, flexibility, and I/O integration. Intel's innovation around the new platform is a leap ahead of our prior generation servers and offers the best business server platform for IT available in the marketplace.

- Lower Power 64-bit Dual-Core Processors
- Intel® Core™ Micro-Architecture
- Hardware assisted Intel® Virtualization Technology (VT)
- New dual independent point-to-point bus
- Fully Buffered DDR2 DIMM Memory (FBDIMM)
- Intel® I/O Acceleration Technology (optional)
- Embedded RAID technology (optional)
- Quad-Core support
- Intel® EM64T 64 bit computing (standard since 2004)
- PCI Express* (standard since 2004)
- Intel® Execute Disable Bit (XD-bit) (since 2005)
- Intel® Software Optimization Tools (optional)
- Intel® Power Efficiency Tools (optional)

New Dual Independent Point-to-Point Bus

- To balance the higher throughput requirements from dual core CPUs, Intel comes up with the new front side bus architecture – DIB (Dual Independent bus). This new point to point bus enables faster FSB speeds (1066 MHZ and 1333 MHZ), much higher throughput (17-21GB/s transfer rate) and better performance.

Power Efficient Micro Architecture

- Power-efficient Intel® Xeon® processor-based servers feature Intel® Core™ Micro-Architecture, software and management tools to help you maximize performance density while providing tremendous improvements in performance, utilization and reliability
- Improve Total Cost of Ownership (TCO) and server density with 80W Dual-Core Intel® Xeon® processor 5100 series based servers that deliver up to 2x the performance of yesterday's servers with up to 3x the power efficiency

Intel® Virtualization Technology

- Intel® Virtualization Technology is part of a collection of premier Intel designed and manufactured silicon technologies that deliver new and improved computing benefits for home, business users, and IT managers. Virtualization enhanced by Intel® Virtualization Technology will allow a platform to run multiple operating systems and applications in independent partitions. With virtualization, one

computer system can function as multiple 'virtual' systems. With processor and I/O enhancements to Intel's various platforms, Intel® Virtualization Technology can improve the performance and robustness of today's software-only virtual machine solutions.

Intel® Active Server Manager

- Intel® Active Server Manager (ASM) delivers the best combination of integrated management hardware, software, and firmware required to manage today's server environments. Intel® ASM can help IT Track the health and status of connected servers, remotely diagnose and repair systems even when the operating system is not running, and keep software and virus protection up-to-date. Intel ASM features, launching with the Dual-Core Intel® Xeon® processor 5000 sequence, includes:
 - ✓ Support for IPMI 2.0
 - ✓ Remote power control and asset management
 - ✓ Advanced features such as IDE-redirection and remote monitoring (KVM pass-through) capability.

Intel® 6311ESB/6312ESB I/O Controller Hub

- Intel® 6311ESB/6312ESB I/O Controller Hub is a highly integrated I/O chipset. It integrates bridge functionality for PCI Express*, PCI-X*, conventional PCI*, LPC, USB*, SATA*, IDE and SMBus, and Dual-Gigabit Ethernet MAC components as well as numerous board management functions. It provides for all system I/O, allowing for simpler system board architectures and smaller board areas than if discrete components were used.

Intel® I/O Acceleration Technology

- Intel® I/O Acceleration Technology, unlike NIC-centric solutions (such as TCP Offload Engine), is a platform level solution that addresses all packet and payload processing bottlenecks throughout the server platform. It increases CPU efficiency and delivers data to and from applications faster than possible with current server platforms. Most importantly, Intel® I/OAT scales with future platform improvements, providing a path for further reducing infrastructure costs by consolidating hardware and software, and ultimately growing your business.

Fully Buffered DIMM (FBDIMM) Memory Technology

- Fully Buffered DIMM (FBDIMM) memory enables both increased capacity and memory bandwidth requirements needed to keep pace with the processor and I/O performance enhancements on today's dual-core server processors.
- FBDIMM technology offers better RAS (reliability, availability, serviceability) by extending the currently available ECC (error-correcting code, a method of checking the integrity of data in DRAM) to include protection of commands and address data. Additionally, FBDIMM technology automatically retries when an error is detected, allowing for uninterrupted operation in case of transient errors.
- The FBDIMM channel pin count is approximately 69 pins per channel, compared with about 240 pins for today's parallel channel. This results in less routing complexity and less routing area between the memory controller and DIMMs, thereby saving board cost to system manufacturers.

3 Platform Hardware Requirements

This hardware integration section demonstrates the base line of how to integrate a server system using a motherboard based on new Dual-Core Intel® Xeon® processors together with an SSI compliant chassis and power supply, as well as FBDIMM memory and other standard off-the-shelf components. The target audiences of this integration guideline are system designers and system integrators. The validation of server functions and performance optimization will not be discussed here.

As a functional server, new Dual-Core Intel® Xeon® processor-based servers should include the following building blocks:

- Motherboard
- Memory
- Processor and Intel designed thermal solution
- Chassis
- Power supply
- Other system components

3.1 Motherboard Integration

The server boards should be compliant with one of the following Server System Infrastructure (SSI) specifications for building block compatibility and interchangeability between different blocks:

- EEB 3.61
 - EEB (Entry- level Electronics Bay) 3.51 for Entry Pedestal Servers and Workstations.
- TEB 2.11
 - TEB (Thin Electronics Bay) 2.11 for rack mount optimized servers.
- CEB 1.1
 - CEB (Compact Electronics Bay) 1.1 for value form factor servers.

The following system features are defined by the SSI specification:

- Baseboard maximum volumetric and mounting-hole locations
- Power and signal connector pin-outs
- ATX-compliant I/O aperture and dimensions that define its location
- Chassis keep-out volume and board/processor mounting requirements



USEFUL INTEGRATION TIPS:

- SSI specifications can be obtained from SSI website at <http://www.ssiforum.org>
- The motherboard should be SSI-compliant, which defines motherboard dimension, component and chassis attachment interface, board block layout directions, as well as power connectors

3.2 Fully Buffered DIMM (FBDIMM) Memory

Fully Buffered DIMM (FBDIMM) memory provides increased bandwidth and capacity for new Dual-Core Intel® Xeon® processor-based servers. It increases system bandwidth up to 21GB/s (with DDR2-667 FBD memory), and increased memory capacity up to 64GB in a 4 channel 16 DIMM server system.

FBDIMM memory has different versions in terms of the DRAM used in the memory module: DDR2-533 FBD memory and DDR2-667 FBD memory, and has different capacities: 512MB, 1GB, 2GB and 4GB. DIMM pairs must be identical with respect to size, speed and manufacturer.

The picture below illustrates an example of an FBDIMM.



Figure 3.2.2 FBDIMM Example



USEFUL INTEGRATION TIPS

- 1) To take advantage of the Intel® E5000P chipset, Intel recommends using at least 4 FBDIMMs in new Dual-Core Intel® Xeon® processor-based servers to achieve optimal throughput. Example: For 2GB configurations use 4 x 512M FBDIMMs rather than 2 x 1GB FBDIMMs.
- 2) To boot up the system, one FBDIMM should be installed in the first FBDIMM slot.
- 3) Intel has completed FBDIMM validation for different memory configurations. System integrators can find the complete list of tested memory modules at:

<http://developer.intel.com/technology/memory/>

3.3 Processor and Thermal Solutions

The boxed version of the Dual-Core Intel® Xeon® processor 5000 sequence supports a passive 2U+ thermal solution, as well as a combination active/1U passive solution (as figure 3.3.1 and 3.3.2 shows). The 2U passive solution can be used in 2U+ rack chassis and pedestal systems; the active solution is a combined solution that supports pedestal chassis with the fan attached, and 1U rack systems with the fan removed (as figure 3.3.3 shows).




Figure 3.3.1 Active solution




Figure 3.3.2 2U passive solution




Figure 3.3.3 Combined Active Solution

The boxed Dual-Core Intel® Xeon® processor 5000 sequence requires the heat sink to be directly attached to the chassis, in order to securely attach the heat sink. As shown in Figure 3.3.4, verify the Common Enabling Kit (CEK) spring is installed on both processor sockets before motherboard installation. Refer to your motherboard documentation for more information, or contact your motherboard manufacturer to obtain a CEK spring for each processor socket if not included with the motherboard.




Figure 3.3.4 Verify the CEK Spring is Installed for Each Processor Socket

Dual-Core Intel® Xeon® Processor Installation




CAUTION

When unpacking a processor, hold by the edge only to avoid touching the contacts.




A. Open the socket lever

- 1) Push the lever handle down and away from the socket to release it.
- 2) Pull the lever and raise until it stops.



B. Open the load plate

- 1) Push the rear tab with your finger tip to bring the front end of the load plate up slightly.
- 2) Open the load plate as shown.




C. Remove the processor protective cover

- 1) Take the processor out of the box remove the protective shipping cover.




D. Remove socket protective cover

- 1) Grasp the socket protective cover tab and pull away from the load plate as shown.
- 2) Remove the socket protective cover and store for future use.




E. Install the processor

- 1) Orient the processor with the socket so that the processor cut-outs match the socket notches.
- 2) Install the processor as shown.




F. Close load plate and socket lever


- 1) Close the load plate all the way as shown.
- 2) With your finger, push down on the load plate as shown.
- 3) Close the socket lever and ensure that the load plate tab engages under the socket lever when fully closed.




G. Removing the plastic shield for the heat sink




H. Attach the heat sink to the chassis



I. Tighten screws



J. Connect the fan header if you are using the active solution.



3.4 Server System Infrastructure (SSI) Compliant Chassis

Server System Infrastructure (SSI) compliant chassis provides component attachment, adequate airflow, electronic emission protection, as well as data storage and management. The quality of chassis will greatly impact server system performance. Compared to a desktop PC chassis, the requirement for server chassis on thermal, power supplies and data management is much higher.



USEFUL INFORMATION

- Rack mount chassis include 1U rack chassis and 2U+ chassis. The height of 1U chassis is 1.75" and 2U chassis is 3.5".
- SSI specifications can be obtained from SSI website at <http://www.ssiforum.org>

3.4.1 Chassis Form Factor

Pedestal and rack mount are two types of form factors. Pedestal chassis should be compliant to EEB 3.61 or CEB 1.1, and rack mount chassis should be compliant to TEB 2.11.

EEB is mainly used for entry-level servers with larger form factors, and CEB is targeted at value servers with smaller form factors and lower cost.

The EEB chassis can install one motherboard based on the ATX form factor 'stretched' to 12" X 13", a size sometimes known as "full ATX". This represents the maximum size of board in one EEB chassis, though smaller sizes are possible. The maximum motherboard size in CEB Chassis is 12" X 10.5"; there is a 2.5" difference in length. The TEB baseboard is based on the EEB baseboard size of 12" X 13".

There are specific requirements on rack chassis width and height to make sure they are compliant to rack system. The 1U and 2U maximum height dimensions allow for 1.3mm clearance between adjacent systems when they are installed in a typical rack configuration. The 445.0mm maximum width includes chassis width, rails, and tolerance; chassis with typical 3/8" roller-bearing slide rails are limited to 428.0mm. System typical depth is 610mm to 660mm.

Figure 3.4.1

3.5 SSI Compliant Power Supply

The power supply provides adequate power to the processor, motherboard and other key components. As shown by research data, most of the server shut-down issues are caused by a power supply failure, including over-heat protection and over-current protection. So, choosing a reliable power supply is an important consideration for proper server integration.

Based on the SSI form factor, there are three types of SSI power supplies, EPS1U is used in 1U rack mount server, EPS2U is used in 2U rack mount server and EPS12V is used in pedestal server. Additional to that, there are redundant power supplies for 2U and pedestal server, they are ERP2U and ERP12V respectively.

3.5.1 Mechanical form factors

Below are detailed mechanical dimensions for SSI compliant power supplies:

| Form Factor | Application | Height | Width | Length |
|-------------|--------------------|--------|-------|---|
| EPS1U | 1U rack | 40mm | 106mm | 355mm |
| EPS2U | 2U rack | 42.2mm | 106mm | 348.2mm |
| ERP2U | 2U rack Redundant | 83mm | 108mm | 400mm(Regular) 350mm(Hot Swap) |
| EPS12V | Pedestal | 86mm | 150mm | 140mm (<450W) 180mm (450W~750W) 230mm (>800W) |
| ERP12V | Pedestal Redundant | 86mm | 150mm | 260mm |

Table 3.4.3 Power Supply Form Factors

There is no specified requirement on position and length of wire harness of power supply, the base line is wire harness must make integration and wire plug easier and more beneficial for maintaining.

Some time, the power supply will influence cable layout because of card edge dimensions and small gaps between the power supply and the chassis, especially in rack mount server where cable layout is important may impact system thermal performance.

3.5.2 Power level and output rails

Below figure is one typical power budget in one new Dual-Core Intel® Xeon® based server system. There are 8 rails recommended in SSI specification for the new platform compliant power supplies, as below:

Figure 3.4.11 System Power Level in 5000 Series Platform

The following table provides recommendations for the power and current ratings for different form factors:

| Spec | Power (W) | Current Rating (A) | | | | | | | |
|--------|-----------|--------------------|-----|------|------|------|------|------|-------|
| | | +3.3V | +5V | 12V1 | 12V2 | 12V3 | 12V4 | -12V | +5VSB |
| EPS1U | 600 | 10 | 21 | 16 | 16 | 16 | 16 | 0.5 | 3 |
| EPS2U | 700 | 24 | 30 | 16 | 16 | 16 | 16 | 0.5 | 3 |
| ERP2U | 750 | 24 | 30 | 16 | 16 | 16 | 16 | 0.5 | 3 |
| EPS12V | 650 | 24 | 30 | 16 | 16 | 16 | 16 | 0.5 | 3 |
| ERP12V | 700 | 24 | 30 | 16 | 16 | 16 | 16 | 0.5 | 3 |

Table 3.4.4 SSI recommended power levels and current rating

3.5.3 Power supply connectors

Normally, the power supply distribution board shall have the following output connectors and wire harness configuration:

| Connectors | Pins No. |
|---|----------|
| Base board power connector | 24 |
| Processor power connector | 8 |
| +12V4 base board power connector, required for 700W, 750W and 800W power levels | 4 |
| Fan power with fan speed control | 4 |
| Peripheral power connectors | 4 |
| Floppy power connectors | 4 |
| Serial ATA power connectors | 5 |
| Server signal connectors | 5 |

Table 3.4.5 SSI recommended power supply connectors

Integrators must ensure these power connectors provide the proper output and that they are connected properly. Different colors have been used to mark different power output voltages. One power connector may have several colors in one stripe along the colored wire

| 3.3VDC | 3.3VRS | +5VDC | +12V1 | +12V2 | +12V3 | +12V4 | -12V | 5VSB | PS_ON |
|--------|--------|-------|--------|--------|--------|--------|------|--------|-------|
| Orange | Orange | Red | Yellow | Yellow | Yellow | Yellow | Blue | Purple | Green |
| | White | | Black | Yellow | Blue | Green | | | |

Table 3.4.6 SSI recommended power supply triples color

3.5.4 Hot swap and redundant power supplies

It is recommended to use the hot swap and redundant power supplies for new Dual-Core Intel® Xeon® processor-based servers. Hot swapping a power supply is the process of inserting and extracting a power supply from a server while it is operating, without interruption.

3.6 Other System Components

3.6.1 Hard drive

1) SCSI HDD

The SCSI specification was developed to provide a common interface that could be used across all peripheral platforms and system applications. The SCSI interface addresses a wider range of applications, such as Redundant Array of Independent Disks (RAID) storage, and has a broader command set than the parallel ATA interface. The SCSI system contains the SCSI controller (initiator), the SCSI bus (cable or backplane), and one or more target devices. The SCSI controller may be built into the motherboard or housed on a SCSI host bus adapter (HBA) card in a PCI or PCI-X slot. Both configurations are shown in Figure 3.6.1.

SCSI cables can connect up to 16 devices, including the SCSI controller. SCSI cables consist of 34 twisted pairs of multi-stranded flexible copper wires for a total of 68 conductors. SCSI devices inside the server are connected to the SCSI controller by a 68-pin ribbon cable. The ribbon cable has a connector at each end and one or more connectors along its length. External SCSI devices are attached to the SCSI HBA by a round 68-pin cable. Two sets of terminators, one at each end of the SCSI bus, prevent signal reflections within the cables.




Figure 3.6.1 SCSI components

Since 1981, there have been seven generations of the SCSI protocol. Each new generation has doubled the performance of the previous one (Figure 3.6.2). SCSI performance has ranged from an 8-bit, single-ended interface transferring data up to 4 MB/s (SCSI-1) to the latest 16-bit, low-voltage differential interface transferring data at 320 MB/s per channel (Ultra320 SCSI).




Figure 3.6.2 Bandwidths of seven generations of SCSI

SCSI hard drives provide some benefits:

- Fast data transfer (up to 320 MB/s transfer rate per channel)
- High MTBF (Mean Time between Failures)
- Support multiple HDD on one channel, such as a single channel of Ultra 320 can support 15 hard drivers, two channels can support up to 30 hard drivers, and all of hard drivers in the same channel use only one IRG
- SCSI controller can offload some CPU tasks which used to handling the storage commands; CPU utilization rate for SCSI HDD processing is low. This is good for multiple task system.

2) SATA HDD

Serial ATA 1.0 (SATA)

SATA 1.0 specification was developed in 2001, SATA is the first generation of the new disk interface technology replacing Parallel ATA. In desktops, SATA is expected to replace Parallel ATA as the primary internal storage for PCs. SATA1.0 delivers a maximum data transfer rate of 1.5 Gb/sec (150 MB/sec) per port and its future roadmap shows growth to 6.0 Gb/sec (600 MB/sec). Advantages of SATA include a point-to-point interconnect that enables full bandwidth available to each device, lower pin-count, lower voltage, hot-plug capability, thin cabling, longer cable length and register-level compatibility with Parallel ATA. These added features make SATA an option for DAS, NAS and some Storage Area Network (SAN) systems where Parallel ATA may not have been considered.

Serial ATA II (SATA II)

SATA II is the second-generation SATA disk interface technology currently under development by the SATA working group. The SATA II specification picks up where SATA 1.0 left off, and will be deployed in 2 phases. The first phase, called "Extensions to Serial ATA 1.0", focuses primarily on addressing the needs of servers and networked storage. These include queuing, enclosure services, hot plug, cold presence detect, cabling and backplane improvements. The second phase is anticipated to scale performance to 3.0 Gb/sec (300 MB/sec) per port. These combined enhancements will make SATA II a good option for DAS, NAS and SAN storage systems where price/performance and cost are key factors.

3) SAS HDD

SAS (Serial Attached SCSI) is a point-to-point architecture in which all storage devices connect directly to a SAS port rather than sharing a common bus as traditional SCSI devices do. Point-to-point links increase data throughput and improve the ability to locate and fix disk failures. More importantly, the SAS architecture solves the clock skew and signal degradation problems of parallel SCSI at higher signaling rates. SAS inherits its command set from parallel SCSI, frame formats from Fibre Channel, and physical characteristics from Serial ATA.

The SAS and SATA technologies have several common features, including low-voltage differential signaling, 8b/10b encoding, and full duplex communication. The Serial Attached SCSI standards committee designed the SAS infrastructure to be compatible with SATA drives, allowing the coexistence of both storage technologies in the same system and opening the door to SATA scalability. Because the SAS architecture features a proven SCSI command set, advanced command queuing, and advanced verification/error correction, SAS is the ideal solution for mission-critical enterprise storage applications.

- Performance

The speed of the first-generation SAS link is 3.0 gigabits/second (Gb/s). The speed of the second-generation SAS link will be 6.0 Gb/s . SAS links are full duplex; they send and receive information simultaneously, thereby reducing a major source of latency. The SAS interface allows for combining multiple links to create 2x, 3x, or 4x connections for scalable bandwidth.

- SAS/SATA interoperability

The SAS architecture enables system designs that deploy both SAS and SATA devices, a breakthrough for enterprise customers. This capability provides a broad range of storage solutions that give IT managers the flexibility to choose storage devices based on reliability, performance, and cost.

- Greater scalability

Serial Attached SCSI enables highly scalable topologies—internal, external, or a combination of both—to give manufacturers and customers the flexibility to design and deploy a range of solutions. The Serial ATA Tunneling Protocol (STP) enables SAS HBAs to communicate with SATA devices through expanders and, therefore, is key to SATA scalability in the SAS domain.



USEFUL INFORMATION

- More detail information about the disk interface technology can be found at <http://www.intel.com/technology/serialATA/pdf/NP2108.pdf>

4 Recommended Server Components

Before you begin system integration, you may want to become more familiar with some of the server components. The following section contains a brief overview of the products used in this integration guide. Intel recommends that before beginning system integration, the integrator sources all recommended components needed to complete integration. A recommended list of components needed can be found below.

Table 1- Recommended Server Components List

| Server System | | |
|---|--|--|
| Boxed Intel Server Board Qty: 1 | Product code: Intel® Server Board S5000PSL | For more information please refer to: http://intel.com/design/servers/boards/ |
| Processor | | |
| Boxed Dual-Core Intel® Xeon® processor Qty: 2 | Product code: Intel® Xeon® Processor 5000 Series, or Xeon® Processor 5100 Series | For more information please refer to: http://support.intel.com/support/processors/xeon |
| Server Chassis | | |
| Intel Server Chassis Qty: 1 | Product code: Intel® Server Chassis SC5299-E | For more information please refer to: http://intel.com/design/servers/chassis/ |
| Memory Configuration | | |
| 512MB FBDIMM 533MHz Qty: 4 | Product code: Kingston* KVR533D2S8F4/512 | For more information please refer to: http://developer.intel.com/technology/memory/ |

5 Setup and Configuration of Red Hat* Enterprise Linux* 4

Outline of Red Hat* Enterprise Linux* 4 Update 2 Operating System Installation

Perform the following steps to install the Red Hat Linux System:

1. **BIOS Configuration**
2. **Text Based Installation**
3. **Graphics Based Installation**

BIOS Configuration

1. Turn on the server and quickly enter BIOS by pressing **F2** during boot up. Under boot priority section, make sure you first boot device is **CD-ROM**.

Text Based Installation

1. Insert Red Hat* Enterprise Linux* 4 Update 2 disk 1 (of 4).
2. Reset the system, and it will boot up from the bootable CD.
3. As Figure 5-1 shows, the installation will request you to select a method to install the operating system, press **Enter** to install Red Hat Linux in text mode.



Figure 5-1

4. As Figure 5-2 shows, choose **Skip** to skip the media test and start the Linux installation.

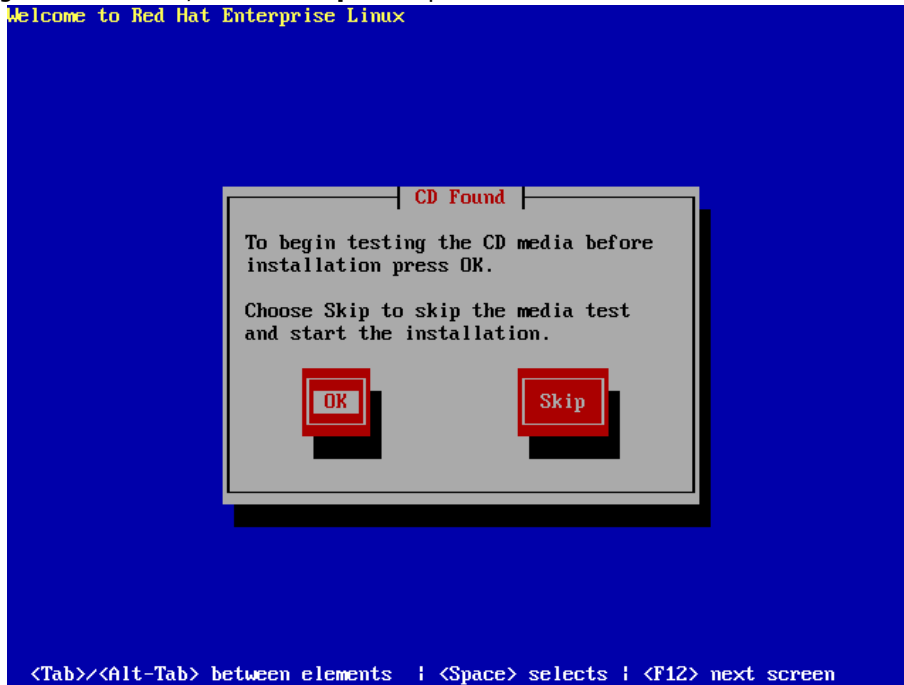


Figure 5-2

Before the installation can begin it must create a new partition and format the hard drive.

Graphics Based Installation

5. The graphic console starts and the “Welcome to Red Hat Enterprise Linux AS installation” is displayed.



Figure 5-3

6. Choose the language you would like to use during the installation

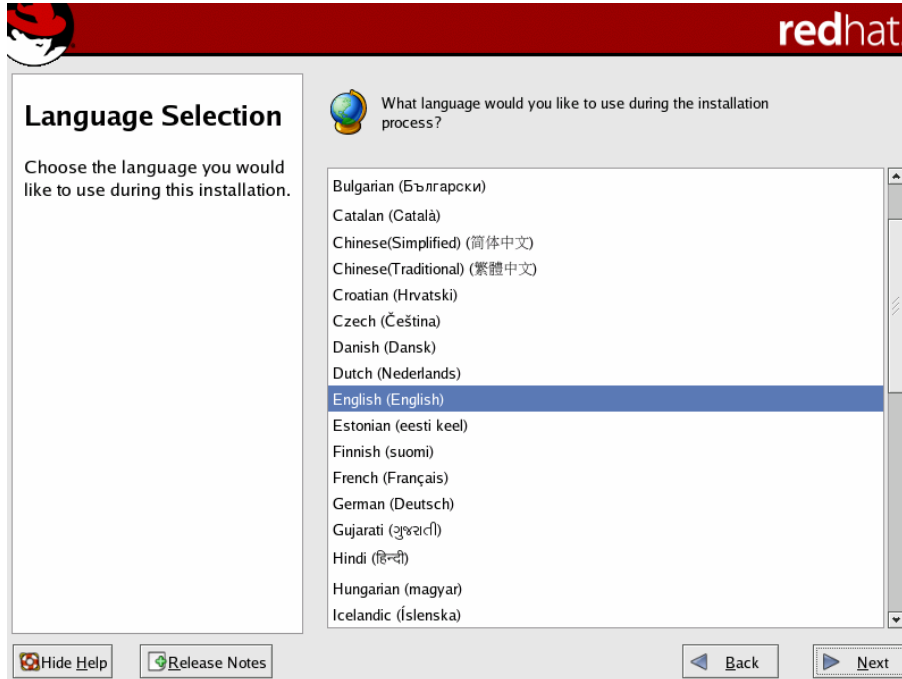


Figure 5-4

7. Choose the layout type for the Keyboard

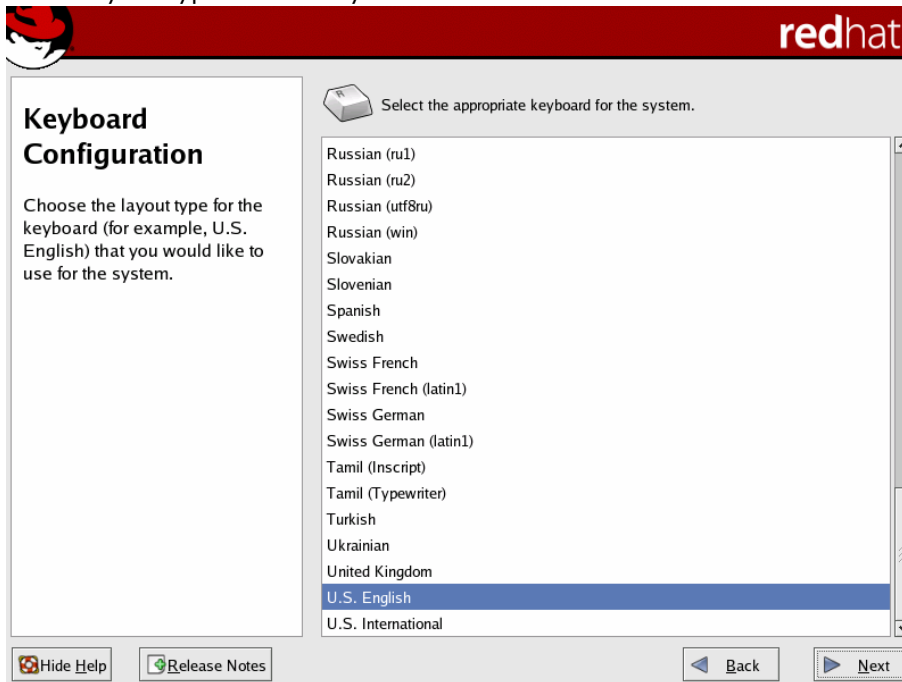


Figure 5-5

8. Disk partitioning is one of the largest obstacles for the new user to install Linux. By selecting automatic partitioning, you do not have to use partitioning tools to create partitions. Select "Manually partition with Disk Druid" to manually partition the hard disk.

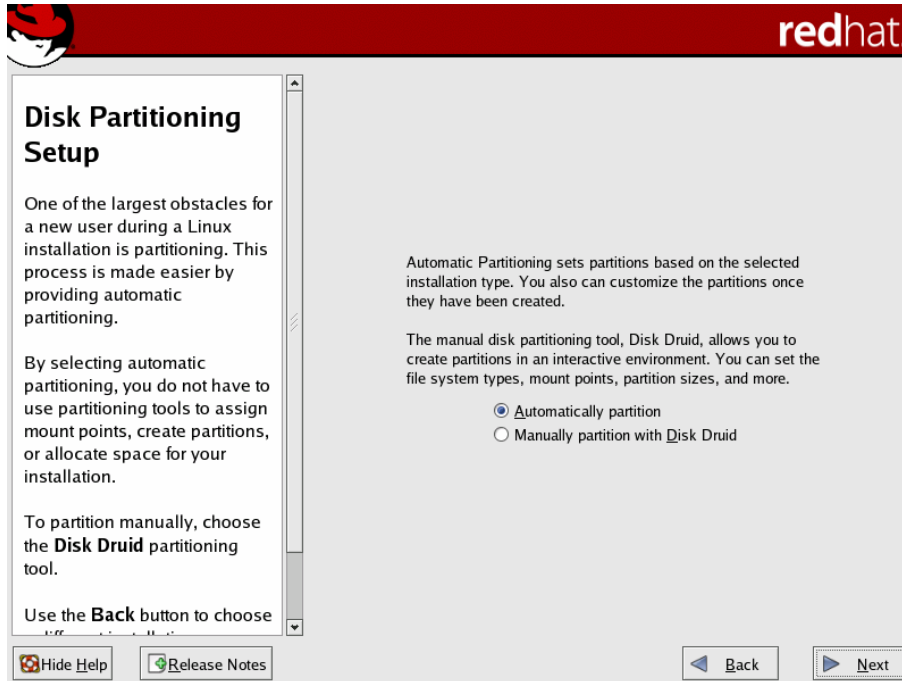


Figure 5-6

9. Note the warning before the Disk Partitioning Setup application begins.

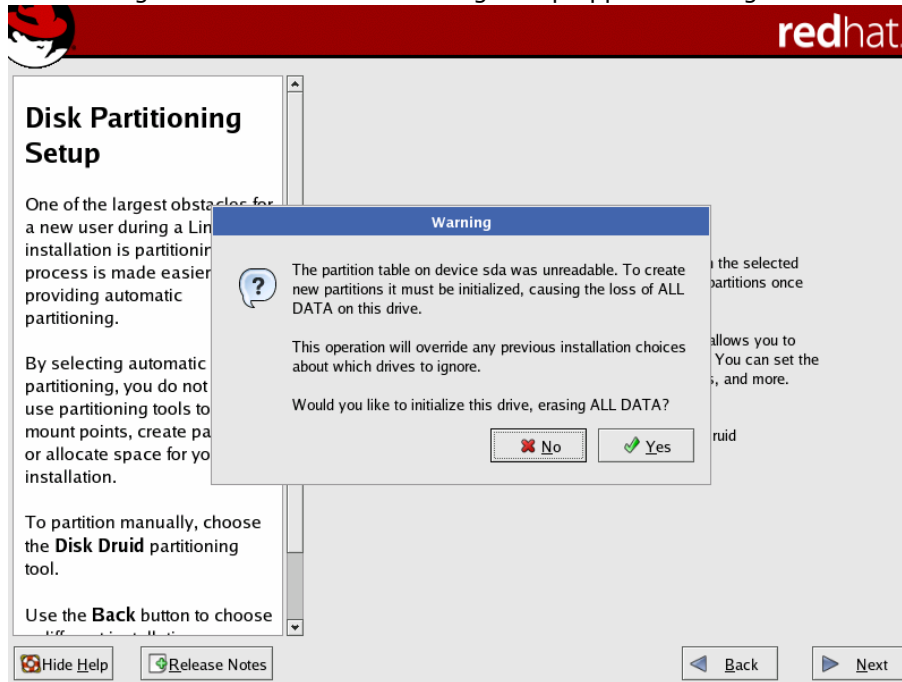


Figure 5-7

This step is necessary to ensure the operating system is installed on a new and error-free partition. Formatting the partition deletes all information that currently exists on the hard drive.

10. Before automatic partitioning can be set up, you must choose how to use the space on the hard drives. Remember: removing all partitions on the system means your data will be **COMPLETELY REMOVED** from the hard drive.

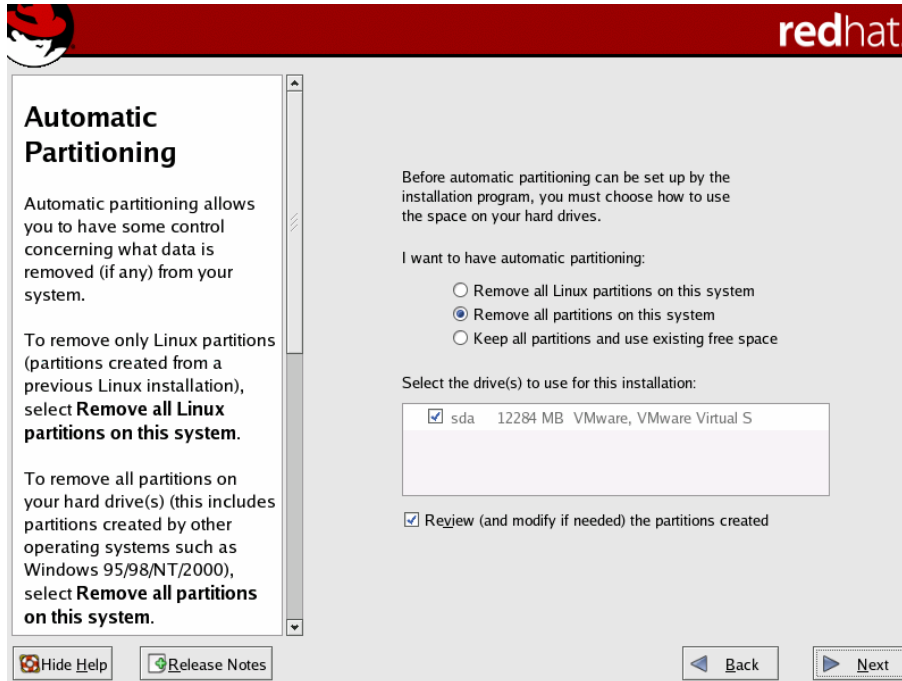


Figure 5-8

11. If you used automatic partitioning, you need confirm the removal of all data and accept the automatically generated partition settings.

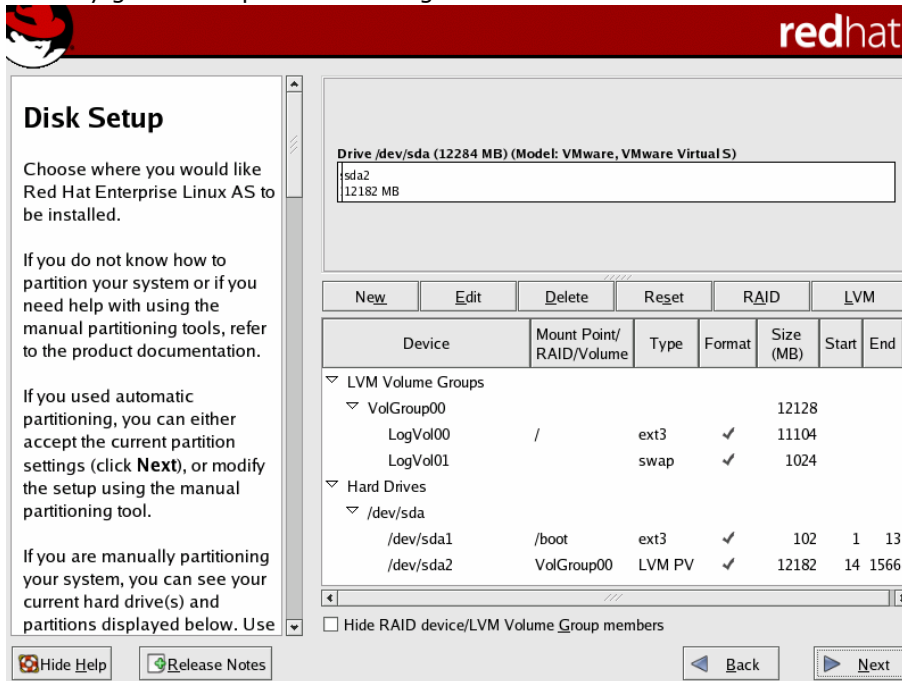


Figure 5-9

12. This step allows you to configure the Linux boot loader to boot other operating systems after the partitioning. You could also choose which operating system (if you have more than one) should boot by default.

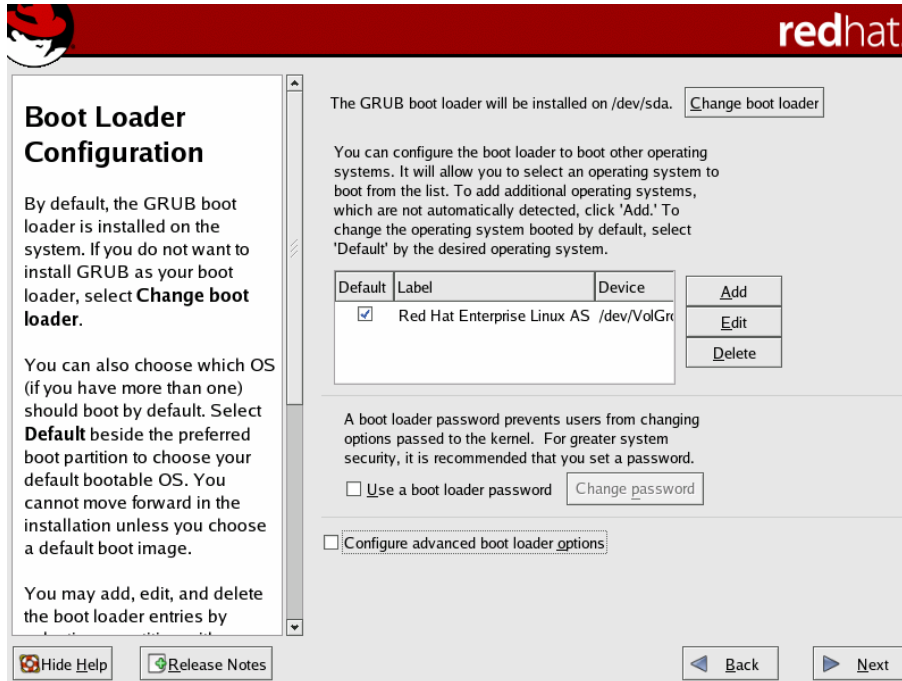


Figure 5-10

13. Refer to Figure 5-11. To configure the network devices, you can either choose to configure by DHCP or manually configure the gateway, DNS and domain information.




Figure 5-11

14. Refer to Figure 5-12. A properly configured firewall can greatly increase the security level of your network. Select "Enable firewall" and the specific services you will allow access to from other computers.

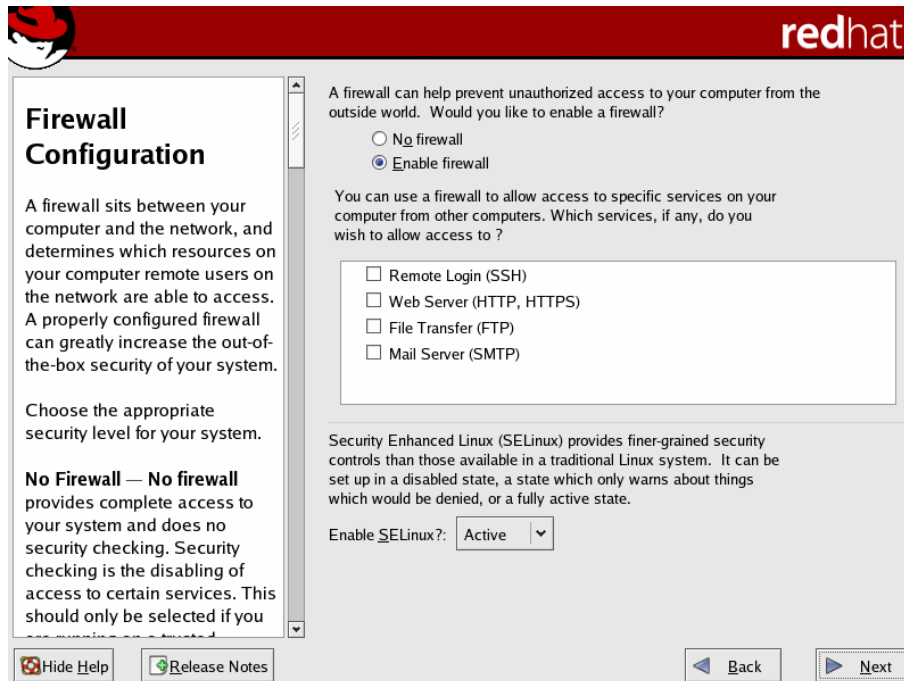


Figure 5-12

15. Refer to Figure 5-13. Select the language to be used as the default language. You can also decide to install other language packs to use with the operating system.




Figure 5-13

16. Select **Time Zone Settings** to set the local time and date settings.

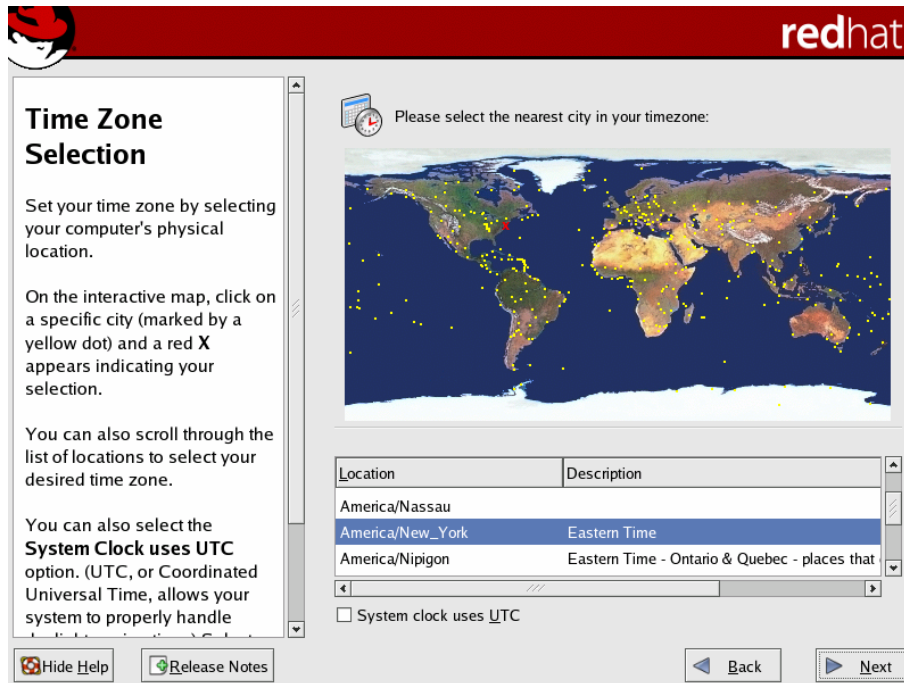


Figure 5-14

Programs that use time in any way access these settings to establish the correct time, therefore these settings are important for successful use of this server.

17. Set the Root Password. Once you have entered the password, please click **Next** to continue.

Use the root account **ONLY** for administration. Once the installation has been completed, create a non-root account for general use and *su-* to gain root access when you need administration rights. This procedure can minimize your chances of a typo or incorrect command doing damage to your system.



Figure 5-15

18. Select the software packages you want to install with the operating system. These packages include extra desktop themes, general and office applications, and development tools.

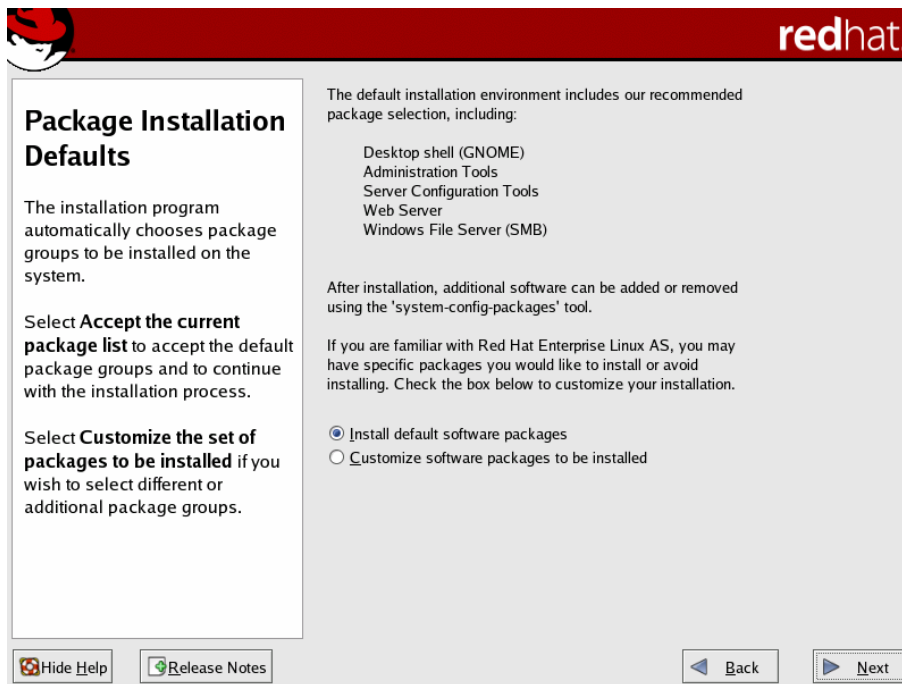


Figure 5-16

19. Choose the package you want to install by selecting the checkbox in each item.

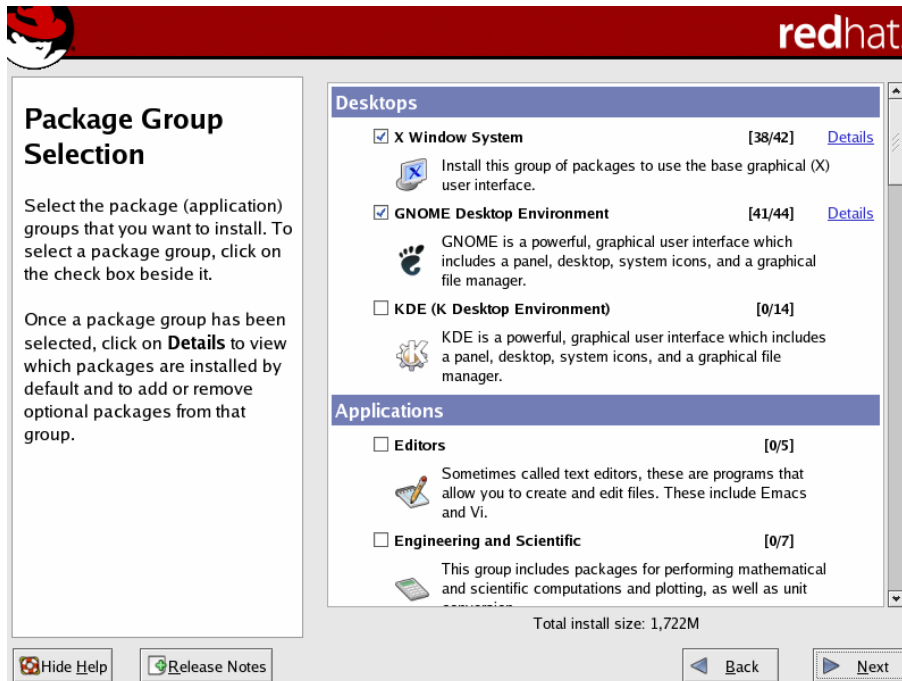


Figure 5-17

20. Congratulations, you are ready to install and copy files.

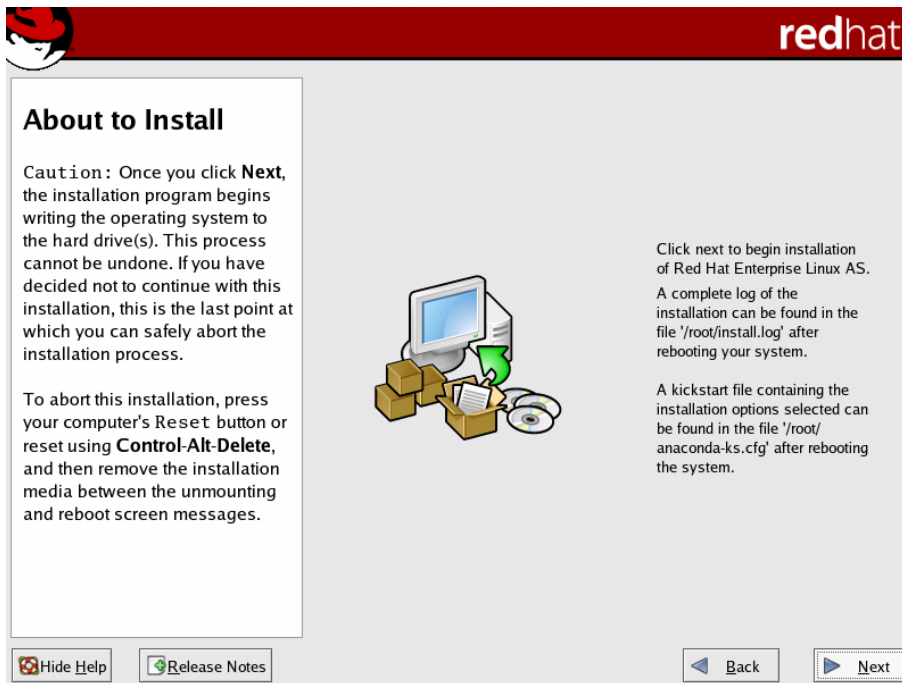


Figure 5-18

21. You are required to have the Red Hat* Linux* installation disks to begin the installation.



Figure 5-19

22. Installation in progress.

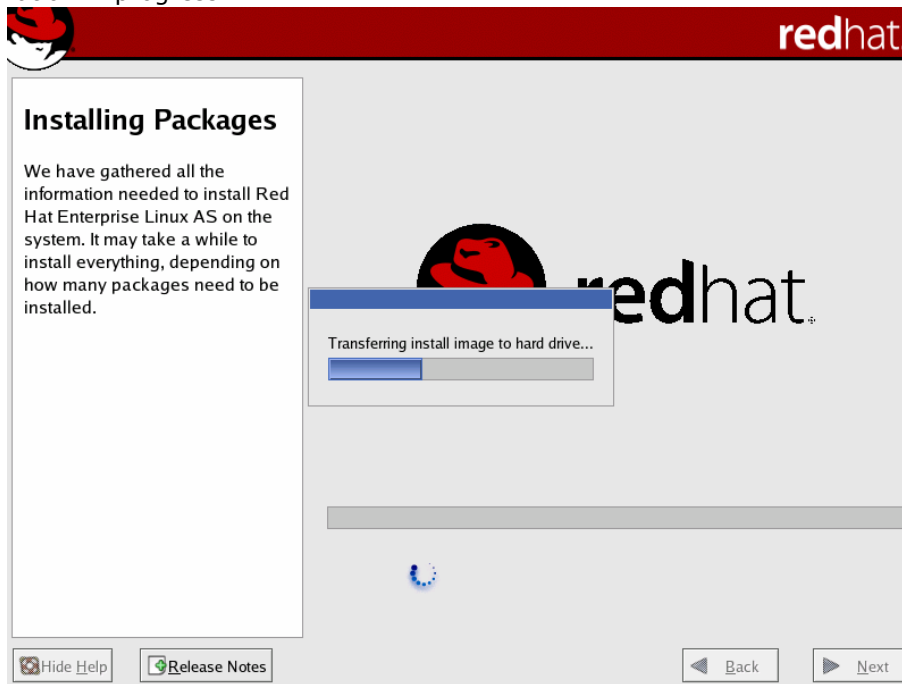


Figure 5-20

23. Installation in progress.

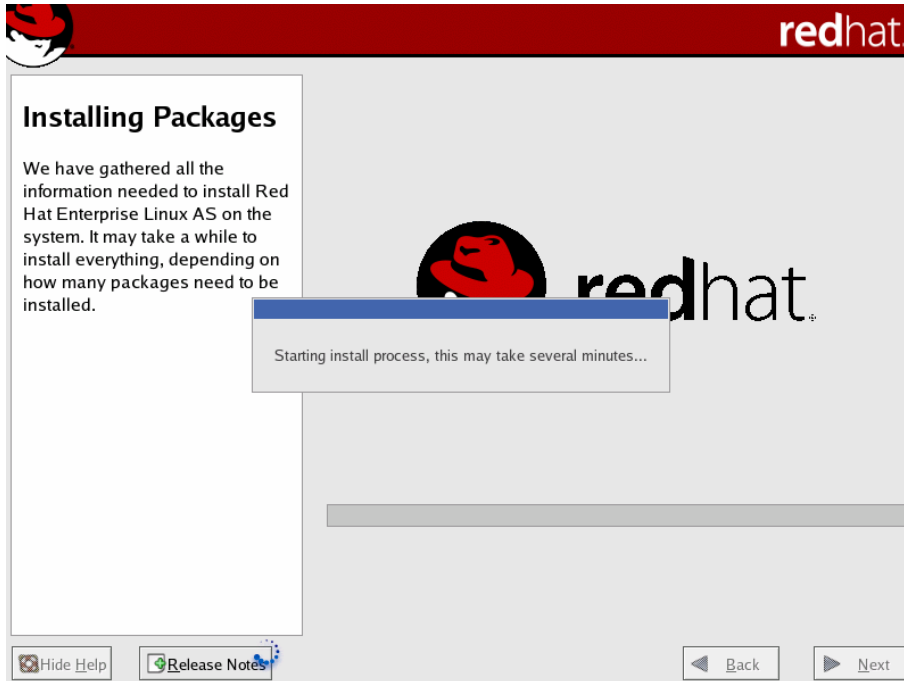


Figure 5-21

24. You will be notified to insert the required disk.

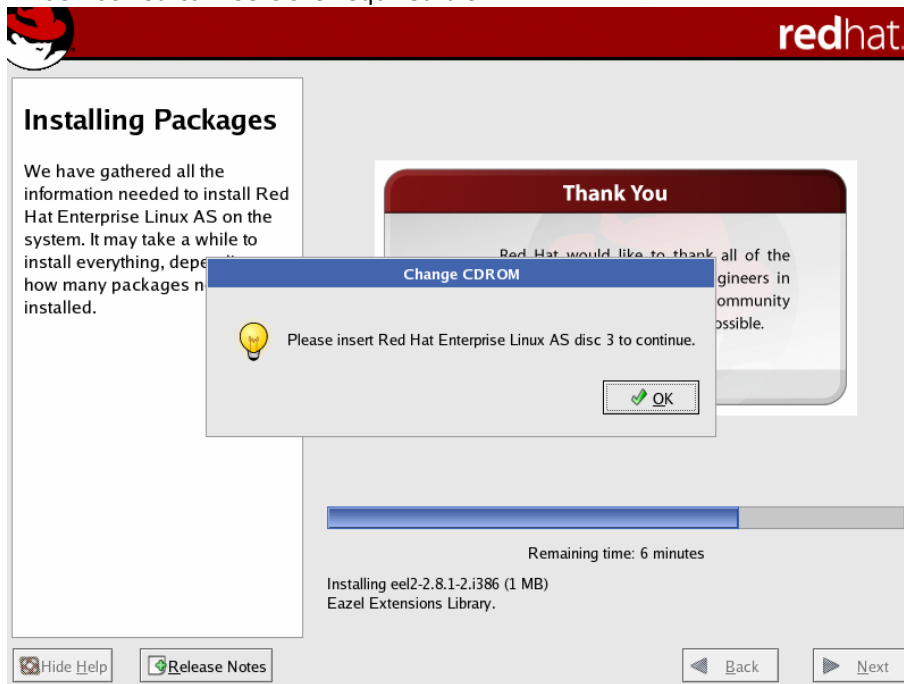


Figure 5-22

25. Post installation process.

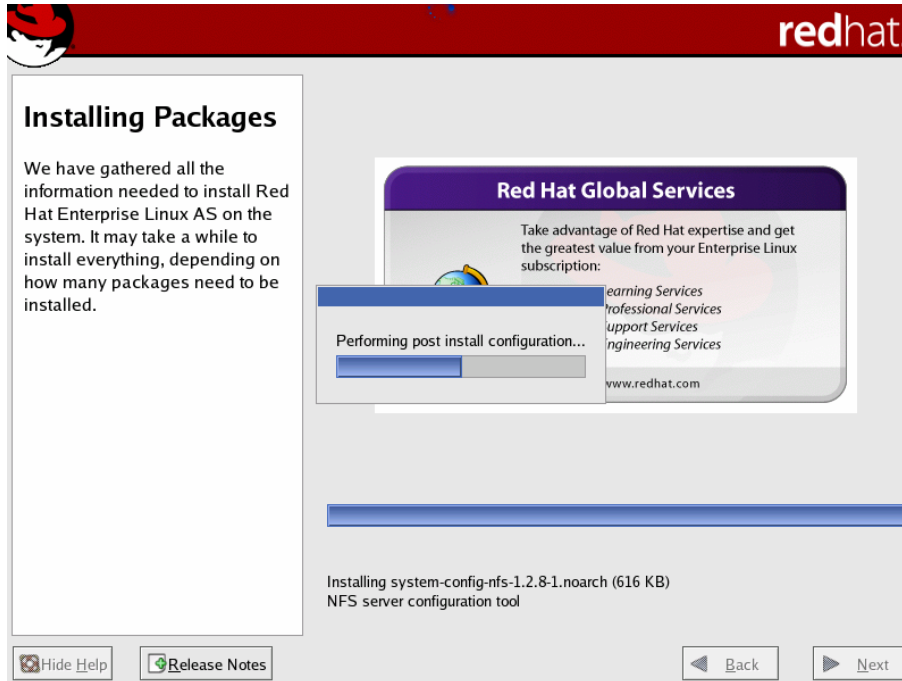


Figure 5-23

26. Congratulations, the installation is complete.

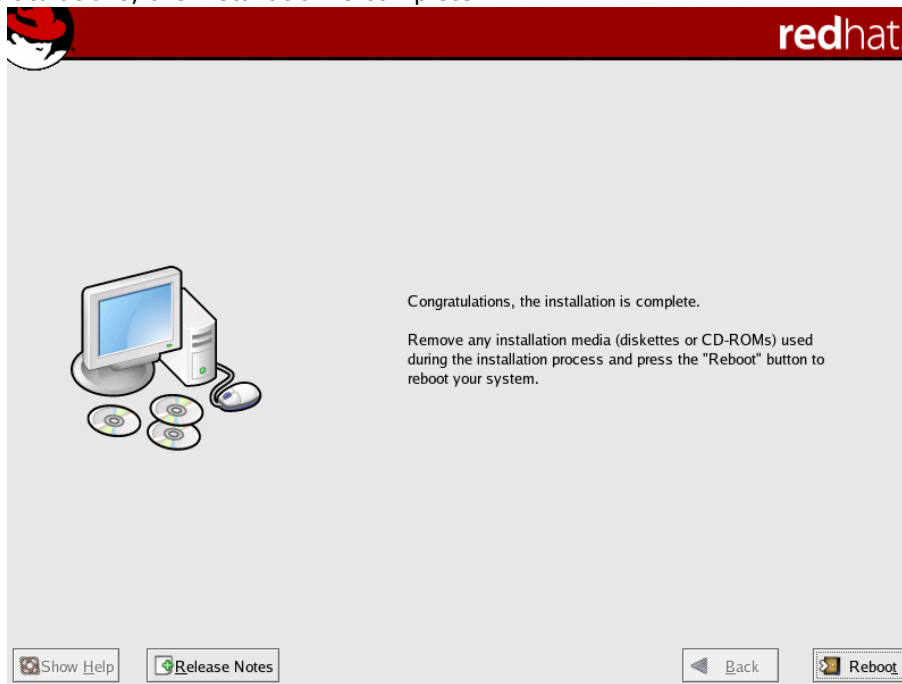


Figure 5-24

27. The installation is finished, but you still need to configure the Linux environment before you can actually use the Linux desktop. Figure 5-25 and Figure 5-26 show the boot up screen.

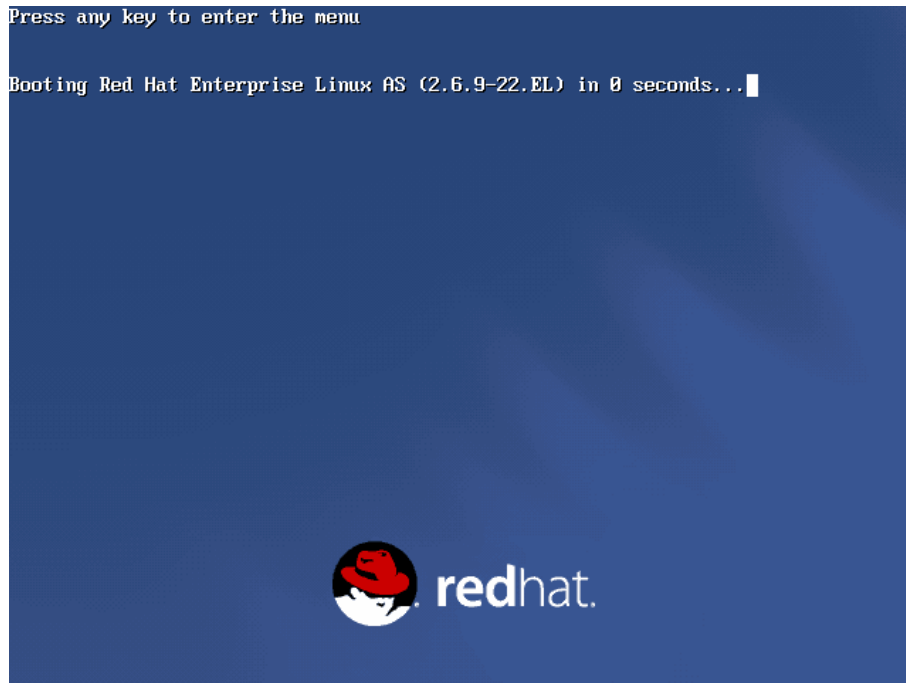


Figure 5-25



Figure 5-26

28. Welcome to the Linux setup agent.



Figure 5-27

29. Click **Yes** to accept the license agreement.

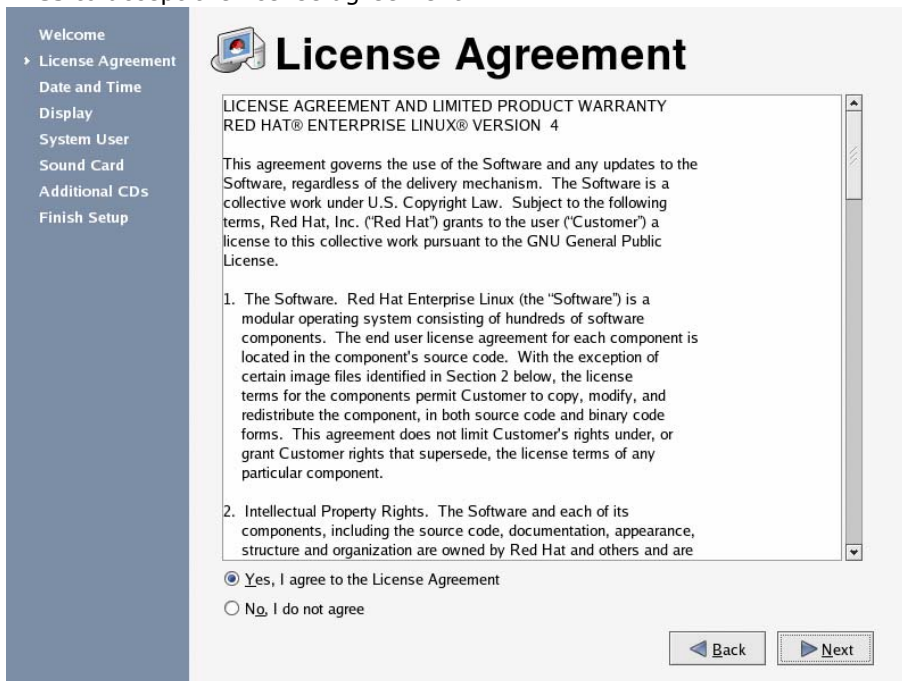


Figure 5-28

30. Set the correct date and time.

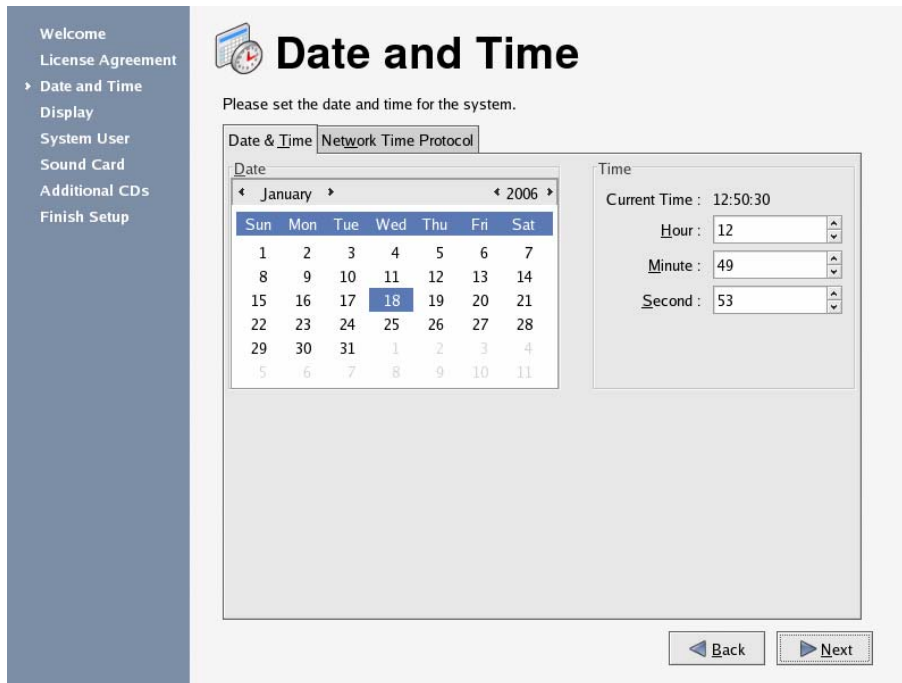


Figure 5-29

31. Configure your display type, including the resolution and color depth.



Figure 5-30

32. It is always recommended that you create a system 'username' for regular (non-administrative) use of the system; this will help secure your system as well.

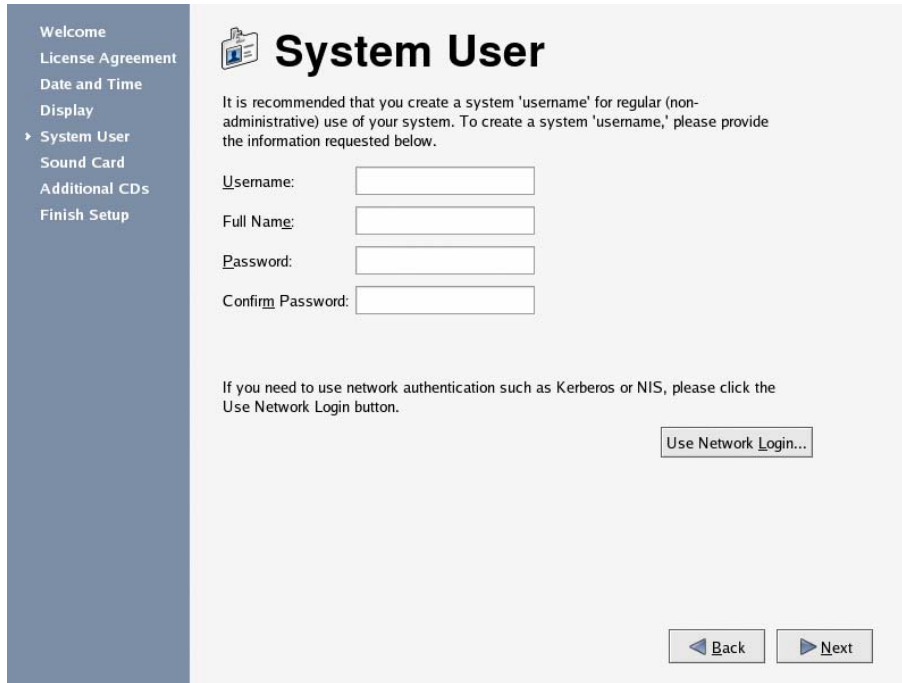


Figure 5-31

33. Configure your sound card.

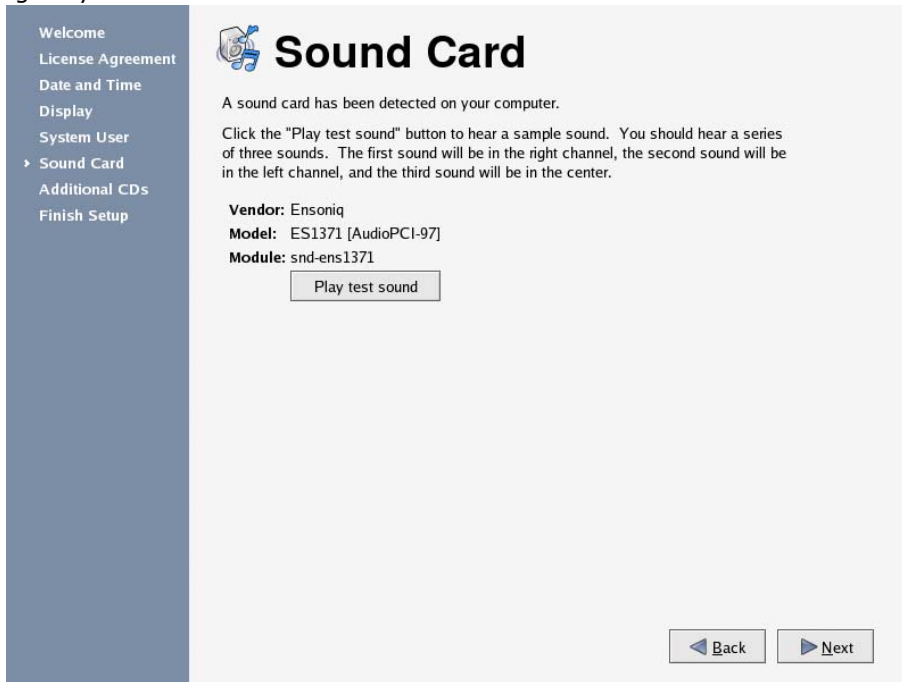


Figure 5-32

34. If you want to install the third party plug-ins applications or the Linux documents, this is the start page.



Figure 5-33

35. You have completed the system setup.



Figure 5-34

36. The Linux desktop after you login in.

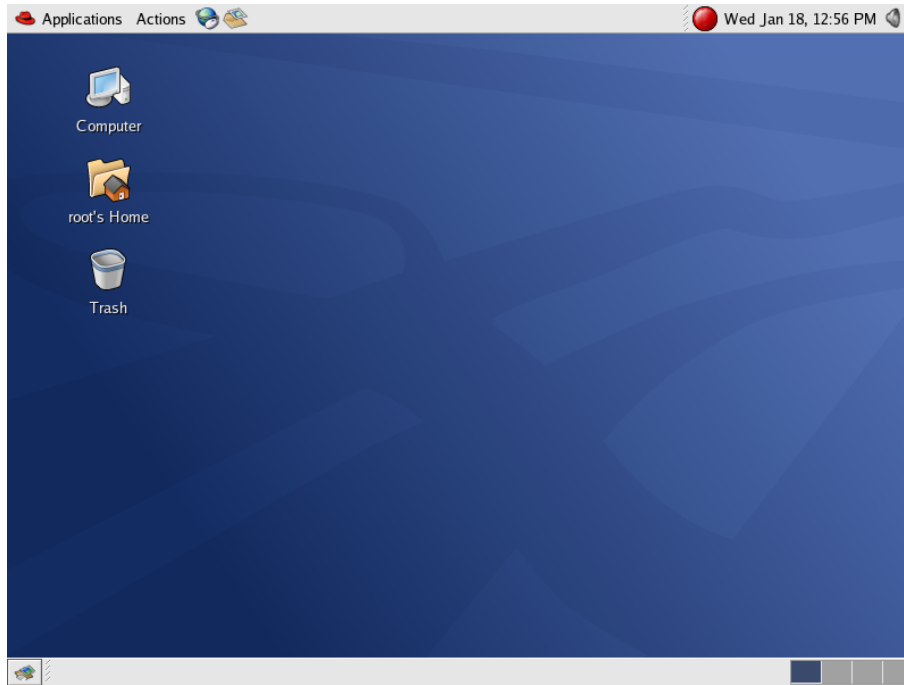


Figure 5-35

6 Setup and Configuration of VMWare* Server (Linux* Version)

Outline of VMWare* Server Installation

Here's the steps you're going to perform to install the VMWare Server.

1. BIOS Configuration
2. Graphic based installation
3. Text based installation and configuration

VMware* Server, is a free virtualization tool for Windows* and Linux* servers. You can learn about more VMware* Server features at <http://www.vmware.com/products/server/>). Before you begin, be sure you have:

- A server system can support Intel® Virtualization technology
- The installation files or CD or disks for the installed operating system

BIOS Configuration

1. Turn on the server and quickly enter BIOS by pressing **F2** during boot up. Under CPU configuration sub-section, make sure you *enabled* the Intel® Virtualization technology.

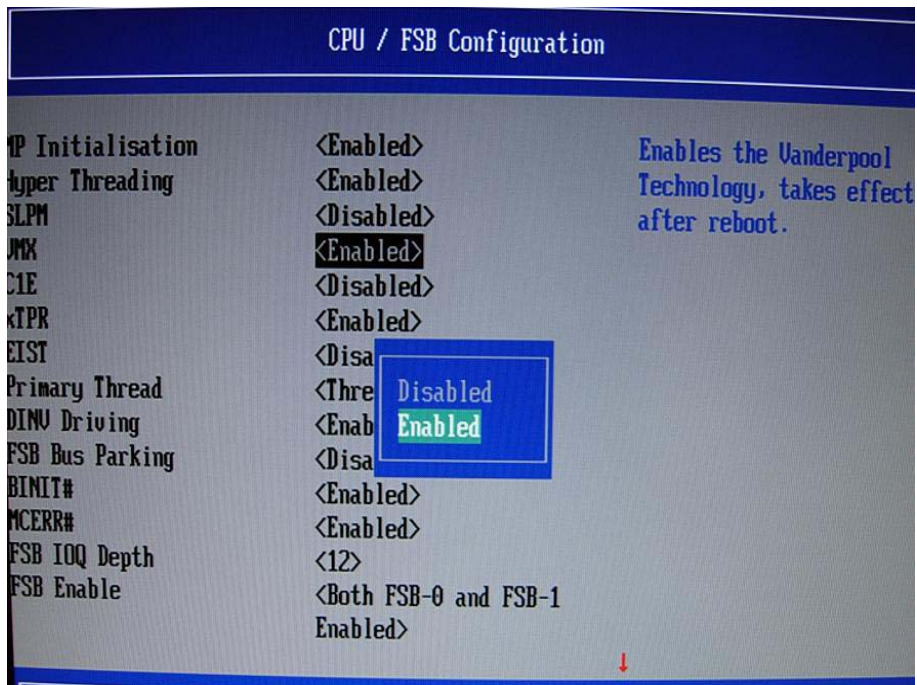


Figure 6-1

The following steps describe the procedure to download the installation file and install from a USB drive. If you are using the CD-ROM instead of downloading the software, the steps are the same except that you start from the CD directory where the installer file is located.

Download your free copy of VMware Server at <http://www.vmware.com/download/server/> before the installation. You will be asked for your VMware registration information before being given access to download files. Please be sure to enter the same email address you provided when registering, because right after the registration, you will get an email contains the serial number(s) to install and run multiple copies of VMware Server.

Graphic Based Installation

2. Save the downloaded file to a USB drive.
3. Log on to your Linux host. In a terminal window, become *root* so you can perform the initial installation steps.

su -

4. Plug the USB drive into one USB port, and the system will automatically mount it.
5. Copy the file into `'/tmp/vmware'`

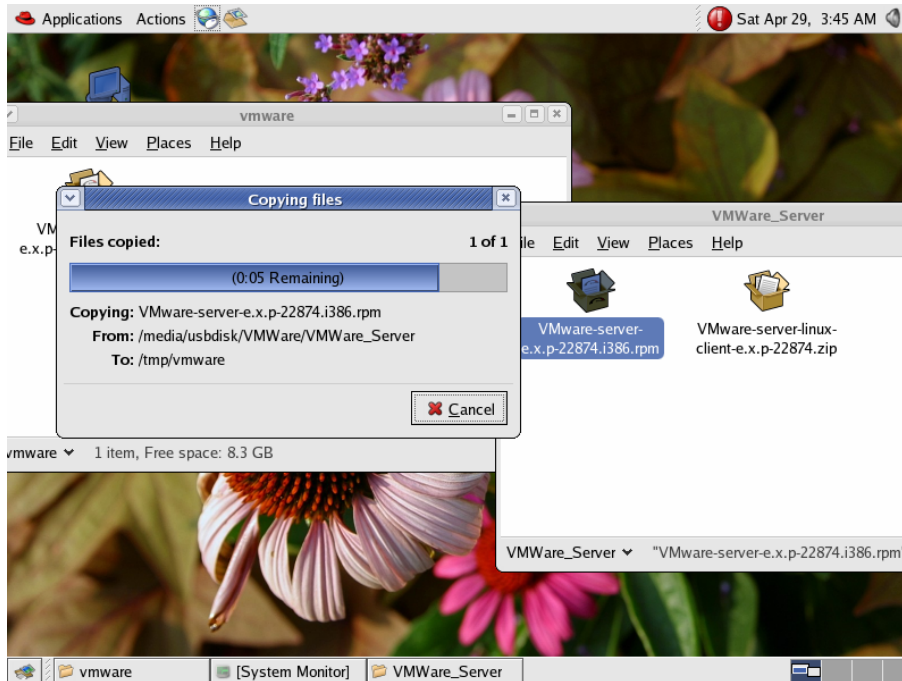


Figure 6-2

6. Change to the `'/tmp/vmware'` directory to verify the copy process is successful.

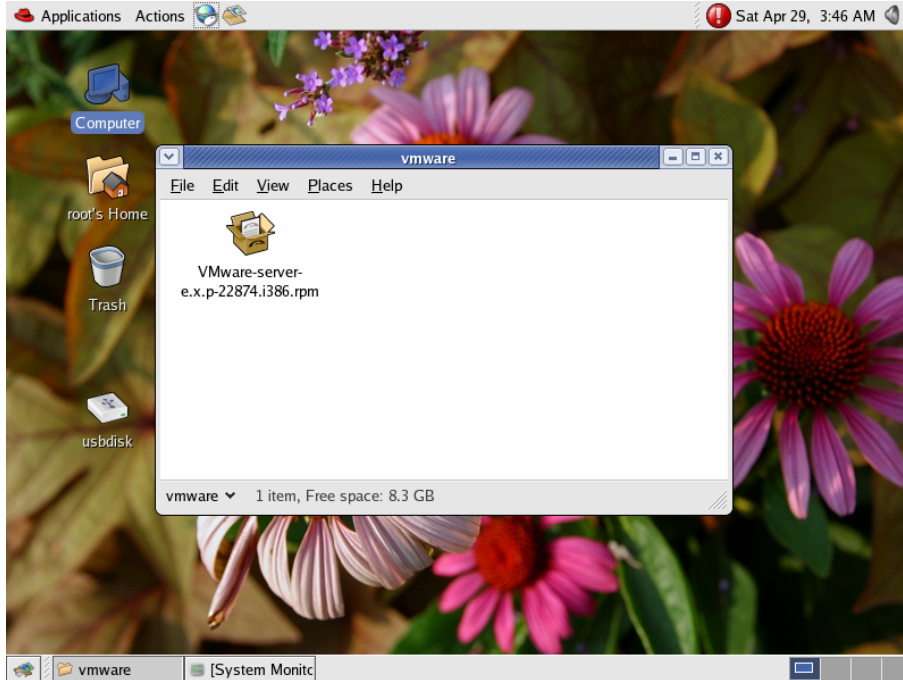


Figure 6-3

Text Based Installation and Configuration

7. Run RPM specifying the installation file.

```
rpm -Uhv VMware-<XXX>.rpm
```

VMware-<XXX>.rpm is the installation file, in place of *<XXX>* the file name contains numbers that correspond to the version and build.

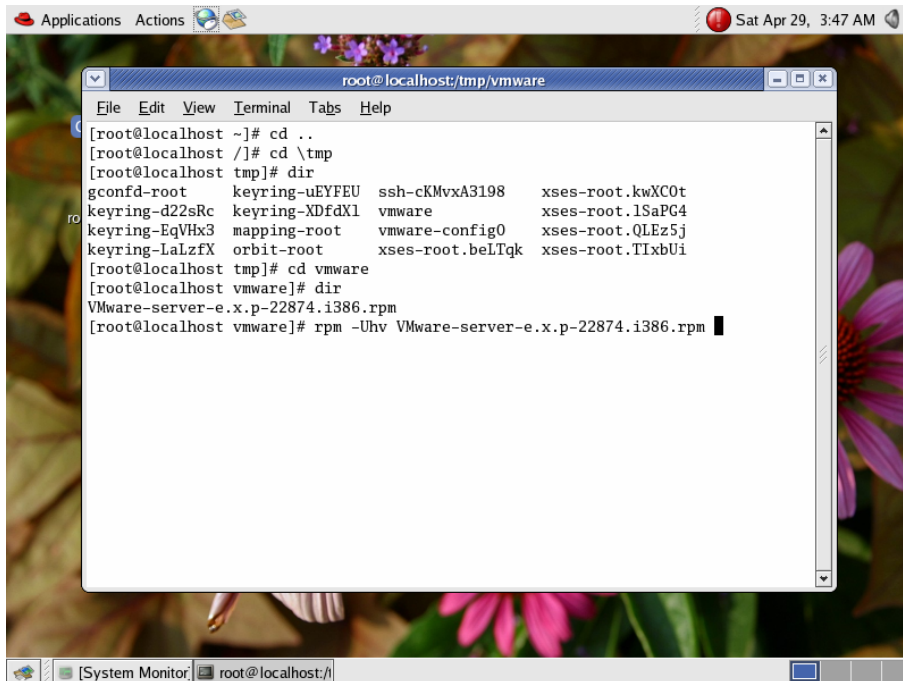


Figure 6-4

8. RPM installation processing...

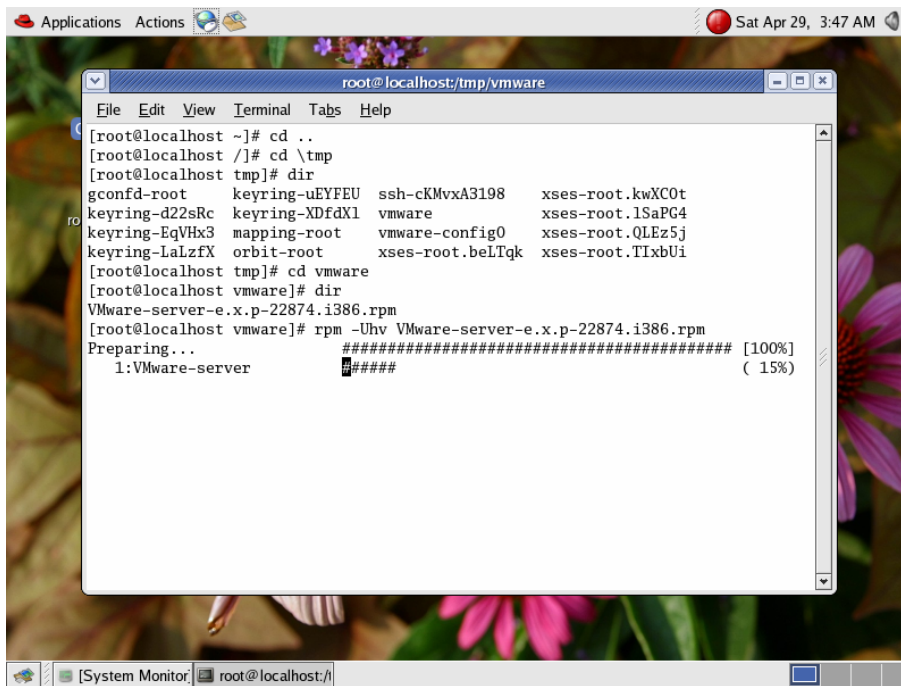


Figure 6-5

9. Run the configuration program from the command line.

vmware-config.pl

Note: Use *vmware-config.pl* to reconfigure VMware Server whenever you upgrade your Linux kernel. It is not necessary to reinstall VMware Server after you upgrade your kernel.

Note: You can also use *vmware-config.pl* to reconfigure the networking options for VMware Server - for example, to add or remove host-only networks.

10. Press **Enter** to read the End User License agreement (EULA). Read through the agreement by pressing the **Spacebar**, and type **Yes** to accept the agreement (as Figure 6-7 shows).

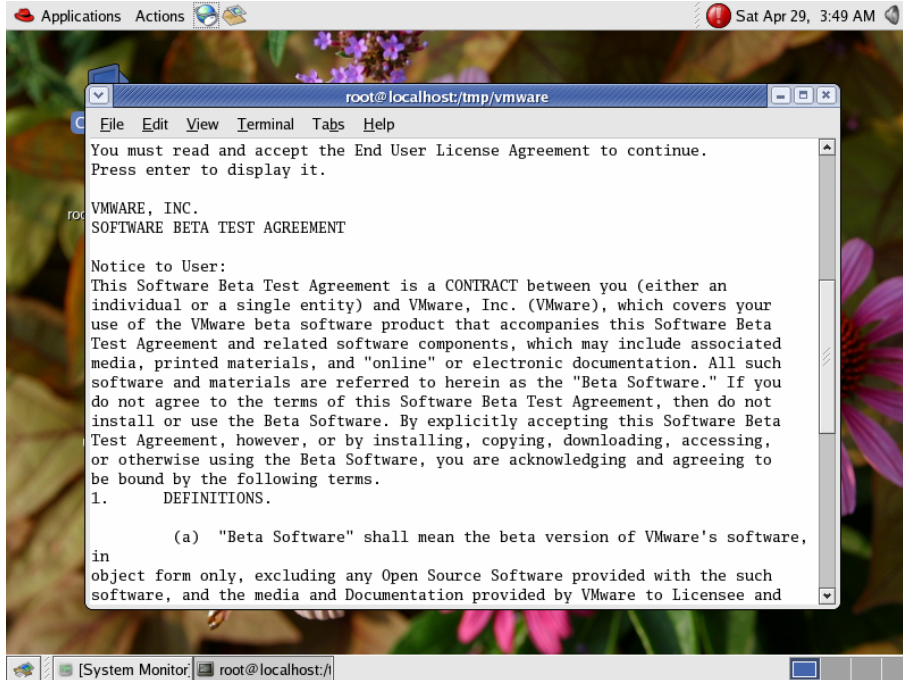


Figure 6-6

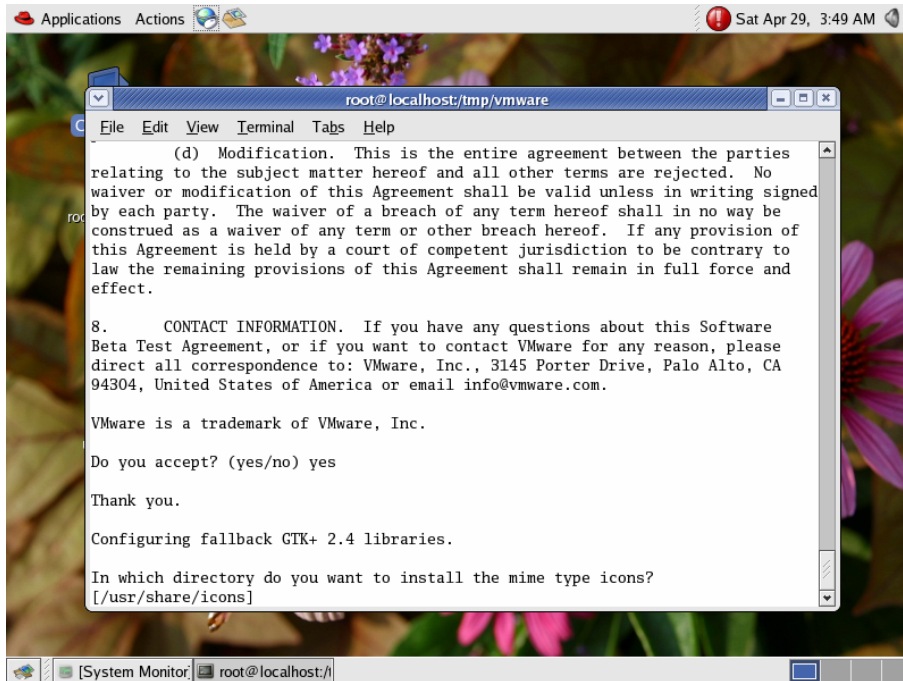


Figure 6-7

11. The remaining prompts are worded in such a way that, in most cases, the default response is appropriate.

Note: if you do not enable host-only networking when you install the software, you can not allow a virtual machine to use both bridged and host-only networking.

12. As shown in Figure 6-8, configure the networking for your virtual machines.

- If you want to use any type of networking with virtual machines, answer **Yes** to this prompt:

Do you want networking for your virtual machines?

Bridged networking is always enabled if you enable networking.

- To enable NAT, answer Yes to the following prompts:

*Do you want to be able to use NAT networking in your virtual machines?
Do you want this script to probe for an unused private subnet?*

This allows you to connect your virtual machines to an external network when you have only one IP network address on the physical network, and that address is used by the host computer.

- To enable host-only networking, answer **Yes** to the following prompts:

*Do you want to be able to use host-only networking in your virtual machines?
Do you want this script to probe for an unused private subnet?*

Host-only networking allows for networking between the virtual machine and the host operating system.

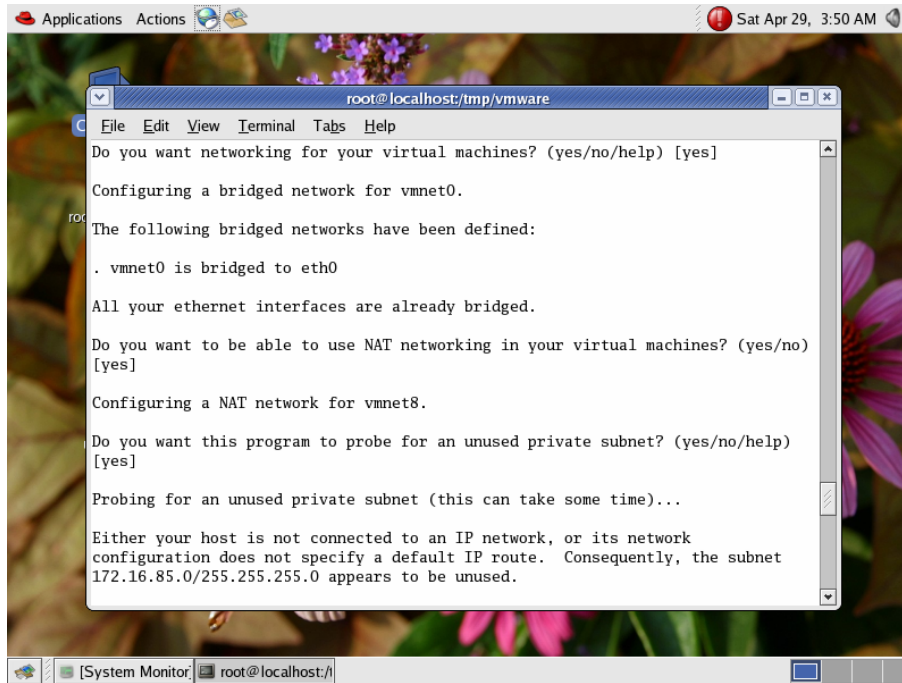


Figure 6-8

13. Enter your VMware Server serial number exactly as it appears (with hyphens) in the email message you received from VMware. When you enter the serial number, it is saved in your license file.

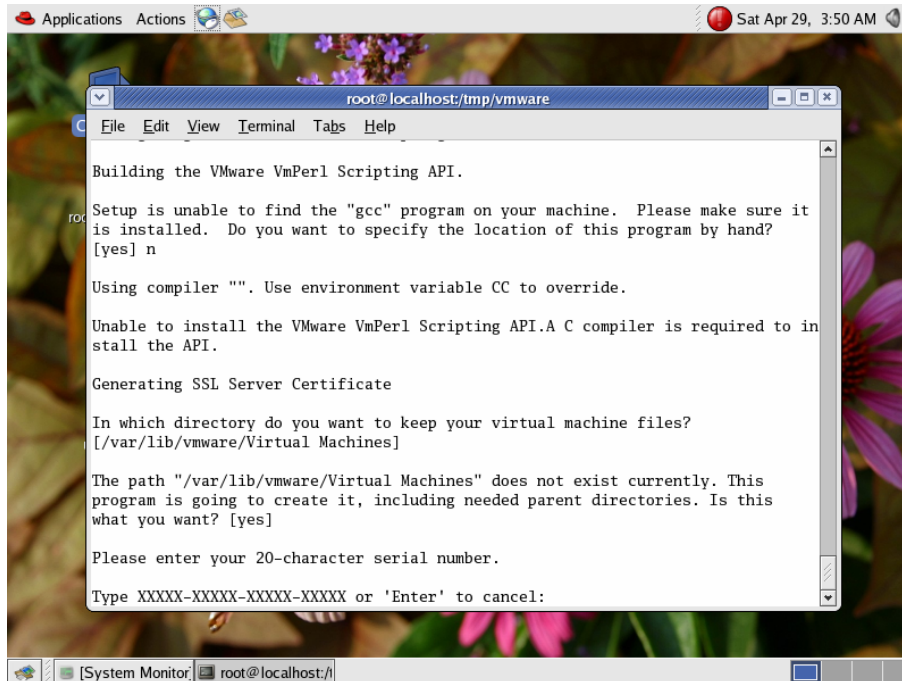


Figure 6-9

14. The configuration program displays a message indicating the configuration completed successfully. If this message is not displayed, run the configuration program again.

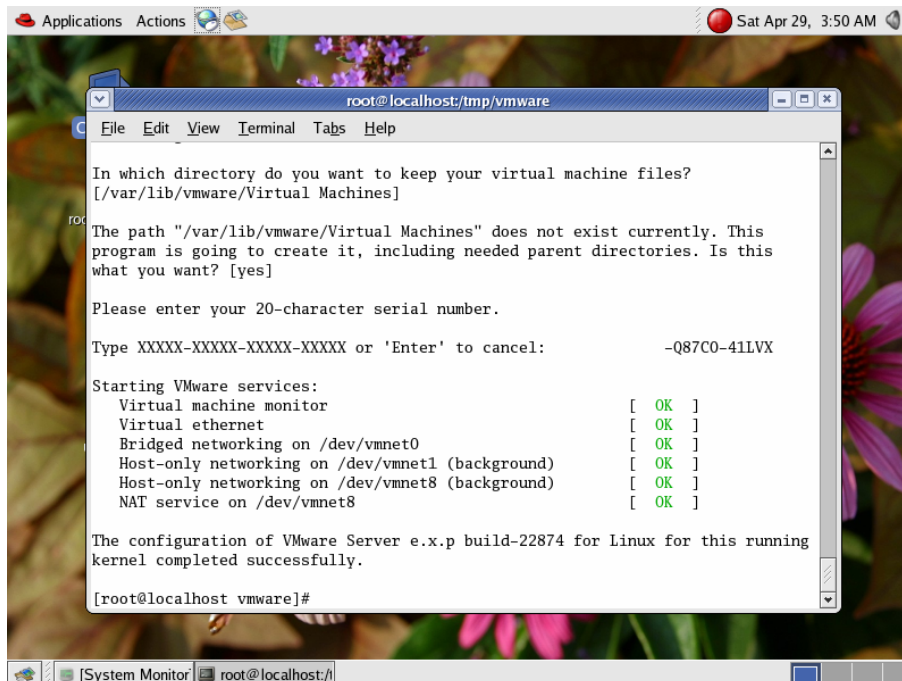


Figure 6-10

15. When done, *exit* from the root account.
16. Run the software from *system tools* directory.

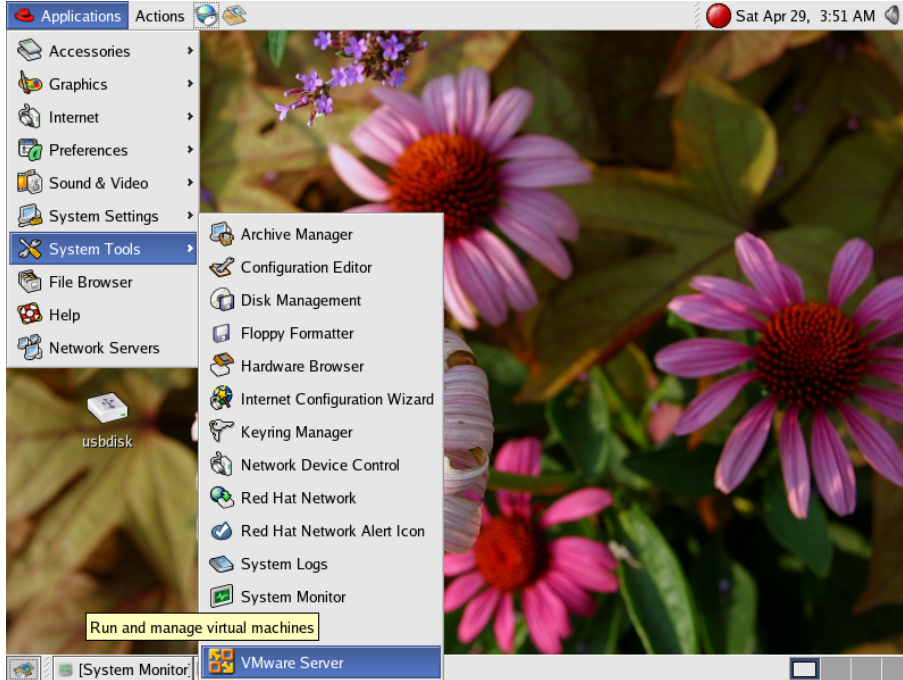


Figure 6-11

17. You will get a prompt asking you to select the VMware host, select *Local Host* if you just want to access virtual machines on the local computer.

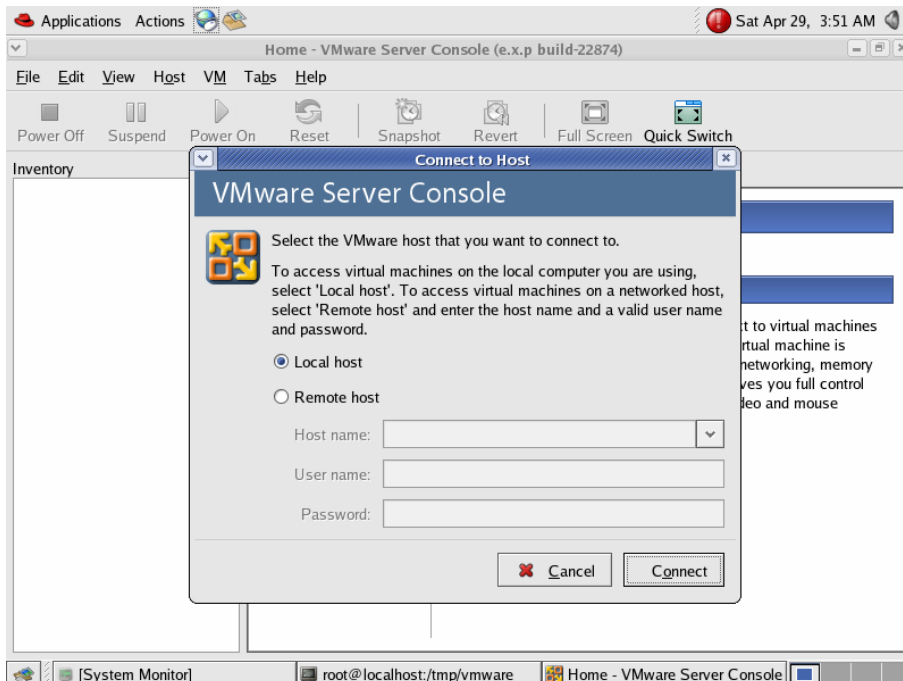


Figure 6-12

18. Initial graphic interface for VMware Server.

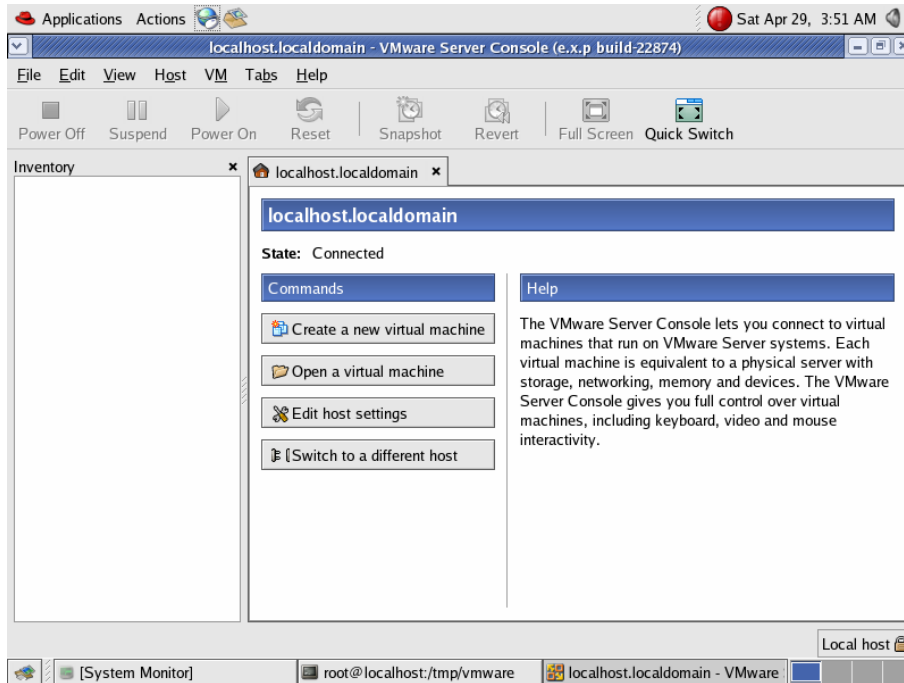


Figure 6-13

19. Product information for VMware Server.

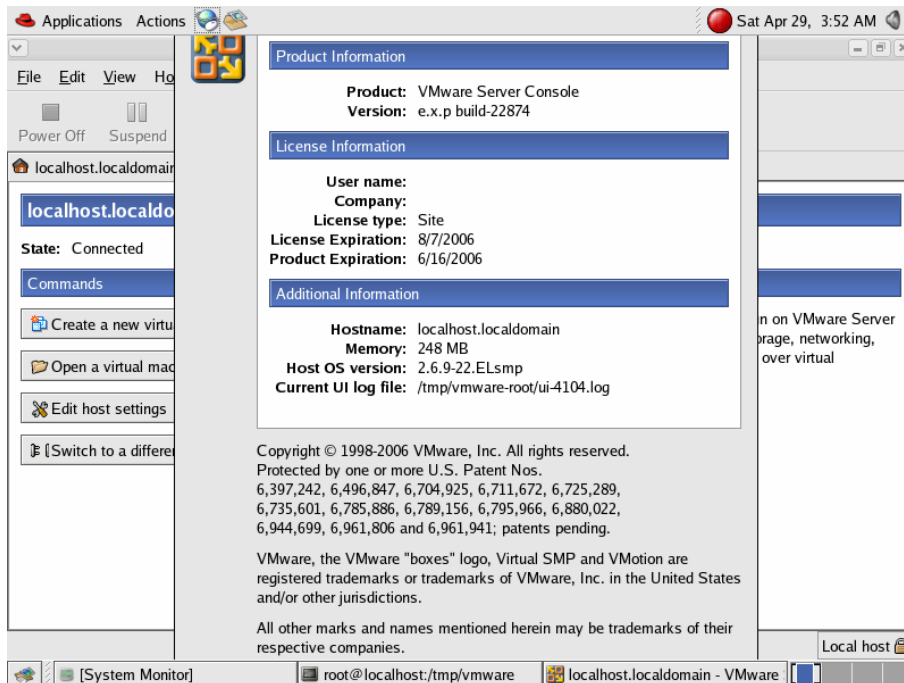


Figure 6-14

20. You can change your serial number using the option in menu "Help->Enter Serial Number".

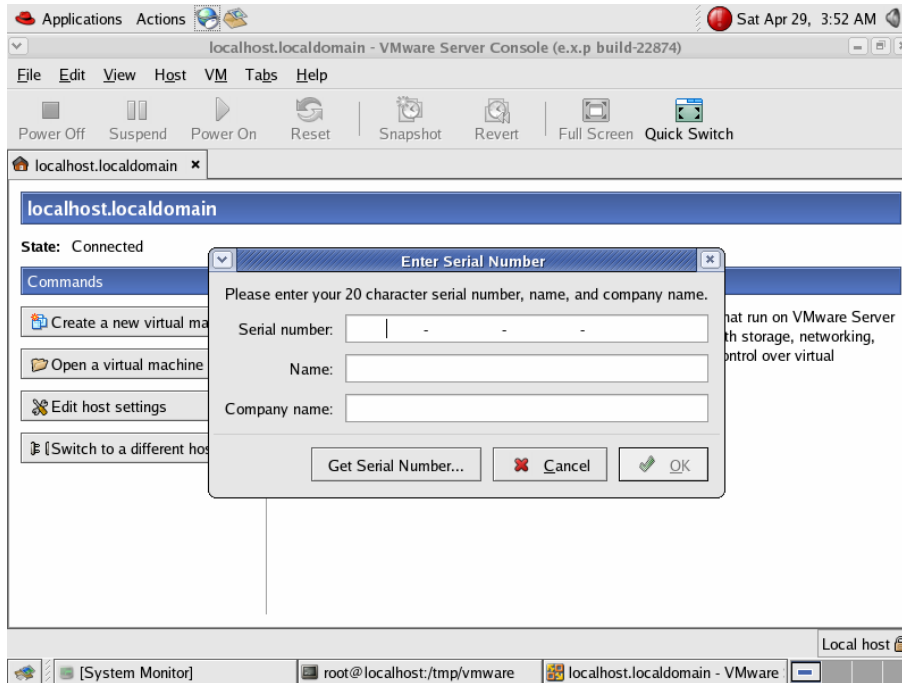


Figure 6-15

21. Installation and configuration completed, you can create new virtual machines with Linux based VMWare Server now.

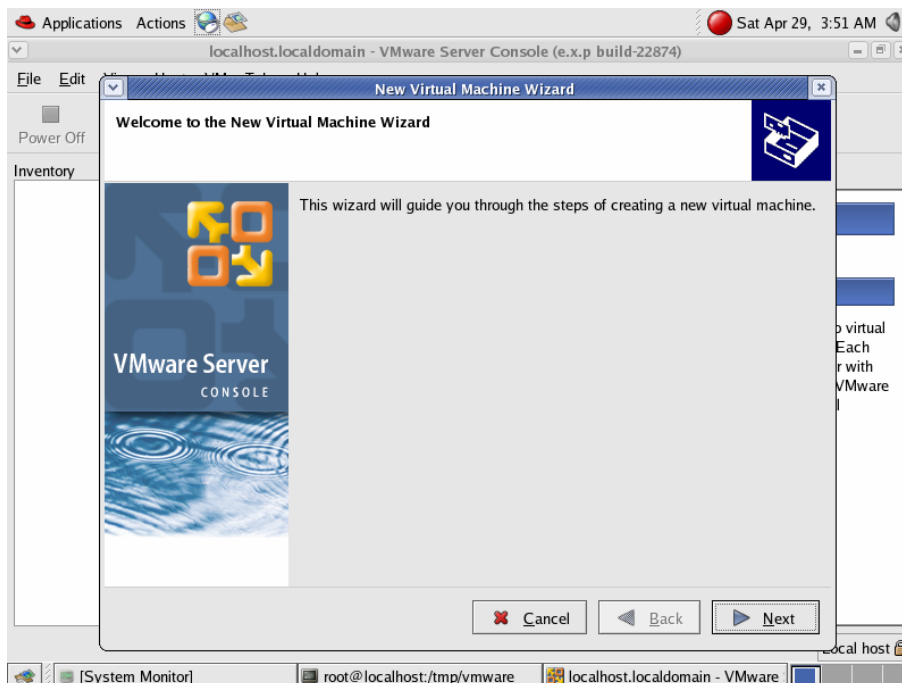


Figure 6-16