

Installing SUSE 9.1

The most important part of getting Linux up and running is installing the system. Unfortunately, this is also where most users encounter problems because of differences between the types of information that you need to know when installing Linux versus Windows. This chapter will demystify the process by helping you through the installation, pointing out any stumbling blocks that you may hit upon, and offering suggestions for resolving them.

The program used to install SUSE Linux is known as YaST, which stands for Yet another System Tool. (The “Yet another . . .” naming convention is a standard Unix/Linux naming convention, intended to humorously reflect on the number of similar tools that different people and companies have developed to do specific tasks in their favorite, customized fashion. YaST provides a framework that supports independent modules that perform a variety of administrative tasks, including modules for installation, all system administration and configuration tasks, and subsequent system updates. The YaST interface that you use for installation is therefore very similar to the interfaces that you will use for system configuration and administrative tasks when you have completed your SUSE Linux installation. Powerful and well-designed, YaST will quickly become your friend.

Selecting Your Installation Method

You can install SUSE in numerous ways. Different installation methods are useful in different circumstances. The most common and recommended installation method is to use the installation media provided with the boxed SUSE Linux product. This book focuses on installing SUSE Linux Professional 9.1 via the CDs provided with the SUSE Linux product. Installing SUSE Linux using the DVD that is also provided in the boxed SUSE product follows essentially the same process, but with the added bonus of not having to switch CDs.

Note

The DVD included with this book provides the SUSE Linux Professional 9.1 distribution.

You can install SUSE Linux in the following ways:

- ◆ **Compact disc** — The easiest and most common form of installation, because almost every modern computer system includes a CD drive. This is the standard way to perform a fresh installation of SUSE Linux on a computer system.
- ◆ **DVD** — A popular form of installation that saves you from having to swap out multiple CDs, but the computer system on which you are installing SUSE must contain a DVD drive.



In This Chapter

Partitioning your disks

Package selection

Configuring your network

Creating a user

Setting up X



Different Installation Sources

This chapter focuses on installing SUSE Linux from the DVD that was packaged with this book or from the installation discs you have purchased. However, your installation discs and the installation DVD that is packaged with this book (like all SUSE installation media) also support a number of other installation sources. If you want to make sure that you get the latest SUSE installation on your system, you may want to select Manual Installation from the initial menu of the boot DVD, select the Start Installation option, and then select the Network source medium. This enables you to select from a variety of different installation sources, including FTP installation, which enables you to install SUSE from a network source, such as one of SUSE's up-to-date repositories. (Other network installation mechanisms include HTTP, NFS, SMB, and TFTP, although FTP is the most common.) To install SUSE from a network source, you must first have used the installer's Network Modules screen to install the drivers for the network card in your computer, and the computer on which you are installing SUSE must also be connected to the Internet. Although this requires some knowledge about your computer system's hardware, it is a great way to get the latest and greatest version of SUSE Linux. As noted earlier, the DVD packaged with this book provides the most recent version of SUSE Linux Professional Edition available at the time that this book was written. To get the latest and greatest version of SUSE Linux and all of its patches, you can always install this version and then update it using the YaST Online Update module that is discussed in Chapter 9.

Because of the amount of storage available on a DVD, the SUSE Linux DVD also includes some packages that are not available on the CD installation set.

- ♦ **Manual installation** — Manual installation requires that you boot from a SUSE CD but provides more control over the source of the packages used when installing SUSE Linux. For example, this installation method enables you to install SUSE from a centralized network repository where the SUSE Linux packages are located, using network protocols such as FTP (File Transfer Protocol), HTTP (Hypertext Transfer Protocol), NFS (Network File System), SMB (Server Message Block, the Windows file sharing protocol), and even TFTP (Trivial File Transfer Protocol). This is a common installation method when you want to install SUSE on a large number of networked computer systems. Manual installation also enables you to install SUSE from an existing hard drive partition where the SUSE packages are installed. You can use Manual installation to install SUSE from a portable, external hard drive.
- ♦ **AutoYaST** — AutoYaST is an advanced installation method that enables a system administrator to create a profile file that can be used to automate installing SUSE Linux on any number of identically configured systems.

As you can see, each installation method has its own advantages and disadvantages, and some are specifically targeted toward technically sophisticated users or system administrators who are installing SUSE into existing networked environments. The remainder of this chapter focuses on installing from CD or DVD, but also provides an overview of using SUSE's network-based installation.

Starting Your Installation

Inside your SUSE box you should find the SUSE manuals (which are considered amongst the best Linux manuals available) and the media case.

The media case contains five CDs and two double-sided DVDs. One of the installation DVDs is installable, while the other contains the SUSE Linux source code. The installable DVD has two sides, one used to install SUSE on standard Pentium-class PCs, and the other containing an installable version of SUSE Linux for 64-bit systems. Each side of the DVDs is labeled in extremely fine print around the center ring of the DVD. Depending on the hardware in your computer system, installing from DVD is the least time-consuming installation method.

Insert the first CD or the bootable DVD in your system's optical drive. If you are booting from DVD, make sure that the side that you want to boot from is facing up in your DVD drive.

Next, enable booting from the optical media drive on your computer to start the installation routine. During boot up of the machine, you need to enter the BIOS and set the order in which your system will probe attached devices looking for bootable media. You can enter your system's BIOS setup routines by pressing a special key when booting the machine. Typically, this is the F2, Delete, or F1 key—check your system's boot screen for BIOS Setup instructions, which are usually displayed at the bottom of the screen. When you've entered the BIOS setup screens, different BIOS have different ways of configuring your system's boot sequence. You may find the options you are looking for under Startup Items, Boot Options, or under your Advanced settings. Make sure that your CD or DVD drive is probed before your floppy disk, hard drives, or network. Once set, save the new settings, and your machine will reboot.

At this point, your system should boot from the first SUSE CD or the DVD, and you will see the welcome screen (see Figure 1-1).



Tip

If your system does not display a screen like the one in Figure 1-1, reboot and hold down the Shift key while your computer system boots. This will reboot your system into a text-mode installer that follows the same general sequence as the graphical boot process described in this chapter, but has fewer dependencies on the capabilities of the graphics card in your machine.



Figure 1-1: The SUSE Welcome screen

As you can see, SUSE caters to many languages, reflecting its European origins and widespread audience (which should only continue to increase now that SUSE is owned by Novell).

Selecting boot options

When the boot splash screen has finished, you will be asked to select how you wish to install SUSE, as well as some other helpful options for booting your system (see Figure 1-2).



Figure 1-2: Boot options

The boot menu offers more than just installation options, although the most common selection is the standard Installation item. We discuss the other six options in detail because at some point in the life of a SUSE user you will likely need to use the others.

- ♦ **Boot from Hard Disk**— This is the default setting if you do not interact with the boot sequence. It's the default because your system automatically reboots as part of the installation process to load the kernel that is installed on your hard drive during the initial phases of the installation process. If you forget to remove the installation media, the system will still boot off the hard disk and the install routine can continue.
- ♦ **Installation**— This is the standard option that most users should select. It will boot from the CD and start the install routine (YaST). We discuss the rest of the process in the remainder of this chapter.
- ♦ **Installation - ACPI Disabled**— Advanced Configuration and Power Interface (ACPI) is a feature of most new processors that controls power management and the way interrupts are handled by the system hardware. You should select this option if you encounter problems during the installation process, such as if your computer system “goes to sleep” (blanks the screen and powers down the drives) and if pressing the appropriate keystroke does not wake it up again.
- ♦ **Installation - Safe Settings**— As with the ACPI Disabled installation method, this turns off some of the features of the kernel that can cause problems with buggy or old system hardware. You should select this option if you encounter problems during installation, and they do not seem related to power management.

- ♦ **Manual Installation**— This will boot the kernel normally, but will drop you into a text-based version of YaST that enables you to select advanced installation mechanisms such as a network or hard-drive install. You must select this option if you want to use these installation methods.
- ♦ **Rescue System**— The Rescue System enables you to correct system problems, such as disk corruption or lost passwords, by booting from the installation media and subsequently correcting system problems. The Rescue System is quite a feature-rich system that you can use to load and edit filesystems, as well as change the settings of an installed system.
- ♦ **Memory Test**— SUSE has been very kind and integrated a memory test suite in the system boot menu. The memory test will run long and exhaustive tests on your system's memory and warn you of any anomalies that it encounters on the way. We have used this a few times with systems that don't quite seem to be running as we expect, and it has been able to tell us that a DIMM (Dual In-Line Memory Module) has indeed failed.

In this chapter, we will select the standard Installation option in the boot menu.

Note

SUSE has changed the original boot splash screen to be something more akin to the Windows boot up (see Figure 1-3). While this is fine for first-time users, it is something that will infuriate hard-core Linux users. SUSE is aware this may be a problem for some users, and pressing F2 while the system boots up will allow you to see the kernel and init messages.



Figure 1-3: Booting SUSE installation

So far, the system has booted a minimal Linux kernel that is sufficient to run the installation process and execute the SUSE installer and the various utilities that it uses to probe and configure your system. SUSE's YaST installer now begins to collect information that it will use to configure your system to match your personal and hardware requirements.

 Tip

The installer uses a very different boot process than that used by a standard SUSE Linux system. The standard Linux boot up sequence will be discussed in more detail in Chapter 4.

Configuring language settings

When the system has booted, you will be asked to configure your language settings (see Figure 1-4). SUSE has put a lot of effort into supporting as many languages as possible to accommodate a large audience. All language options are shown in their respective dialects and associated fonts. When your language has been selected, the installer will instantly change the system language and allow you to continue the installation process in that language.

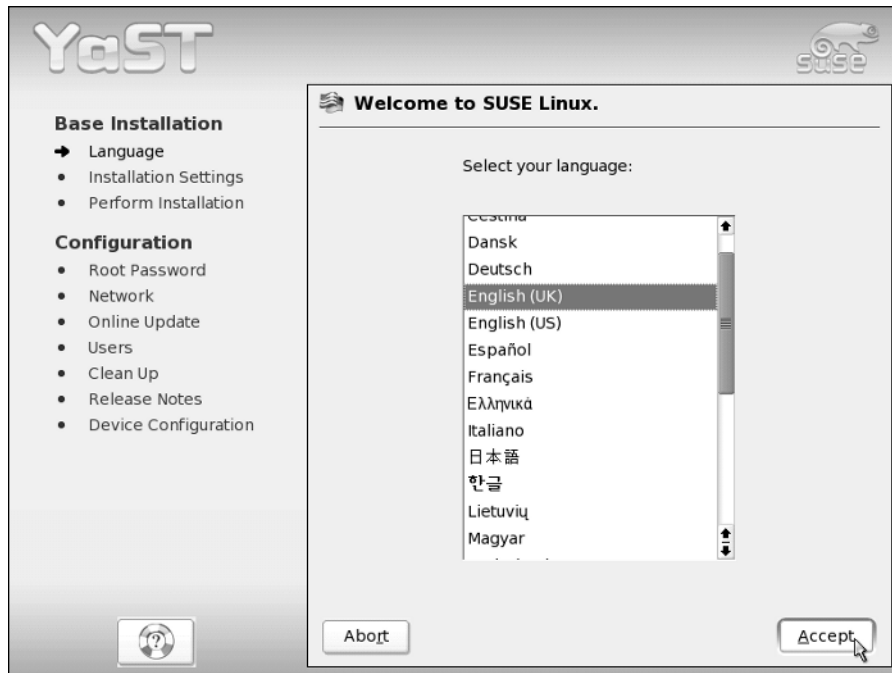


Figure 1-4: Selecting the system language

When you've selected your language, click the Accept button or use the keyboard shortcut Ctrl+A.

During the installation routine, you can control the screen with your keyboard using *accelerators*. Any option on the screen can be selected by pressing the Control key (Ctrl) and the accelerator code, signified by an underlined character in a button or a GUI element. For example, in Figure 1-4, pressing Ctrl+R will abort the installation while Ctrl+A will accept the setting you selected and proceed to the next screen.

Selecting and Customizing Installation

Next, you may be asked for the type of installation that you wish to perform (see Figure 1-5). For example, if you are installing SUSE Linux Professional on a system on which an operating system is already installed, YaST allows you to upgrade an existing SUSE installation, which is usually a very smooth process.

Note

If you are installing SUSE Linux Personal edition, you will not see this screen because the Personal edition does not support upgrades. In that case, skip ahead to the discussion of Figure 1-6.

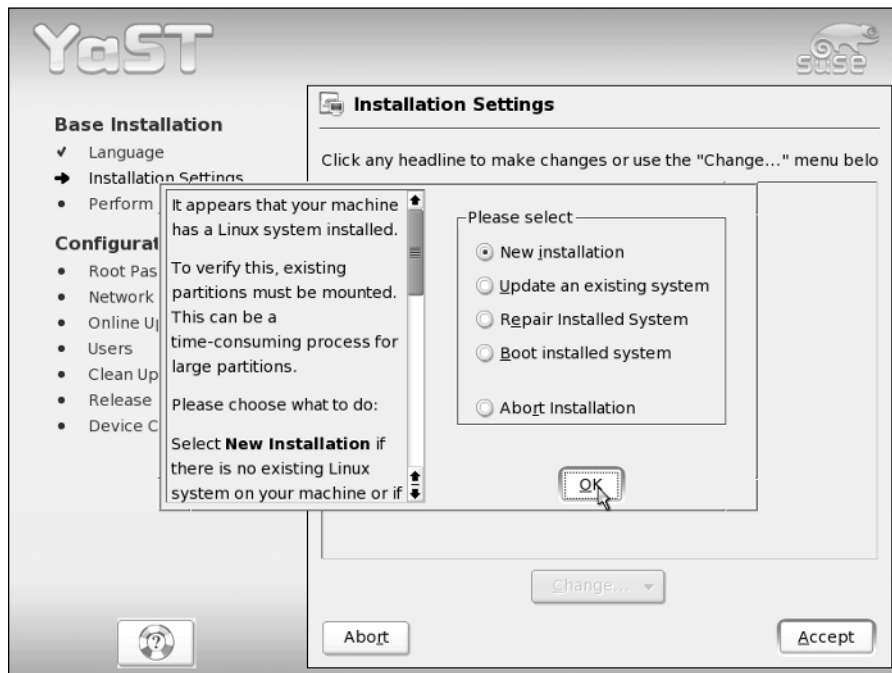


Figure 1-5: Selecting installation type

The installation options shown in Figure 1-5 are the following:

- ♦ **New installation** — Indicates that you want to install SUSE Linux for the first time, regardless of whether that disk already contains an operating system. Any existing SUSE or other Linux installation on the hard disk will be ignored. If you are installing SUSE for the first time on a disk that already contains an operating system other than Linux, such as Microsoft Windows, this is still the right option to select. Information about installing SUSE on a disk where another operating system is already installed (and you want to preserve that operating system) is provided in the disk partitioning part of the installation, described in the next section of this chapter, “Partitioning your disks.” The final stage of the installation process asks you to confirm your selections before it makes any changes to your disk — don’t worry!

- ♦ **Update an existing system** — Upgrades an existing SUSE Linux installation to the version on the media that you are installing from. When upgrading an existing SUSE installation, YaST checks the existing RPM database (the Red Hat Package Manager database where information about installed software is located) and upgrades the components that are already installed, resolving any dependencies it needs. Any errors in RPM dependencies can then be changed manually if the need arises.
- ♦ **Repair installed system** — If you had hoped to upgrade an existing system but cannot because you receive error messages about disk or filesystem problems, selecting this option causes the SUSE installer to perform a number of maintenance procedures, such as checking the consistency of the filesystem(s) in your existing installation. If these complete successfully, you can then proceed with the upgrade. If not, you may need to reboot your system and select Rescue Mode from the first installer screen in order to correct these problems before you can proceed.
- ♦ **Boot installed system** — Enables you to check your existing system before you decide whether to do a new installation or upgrade it. Booting an installed system will use the install kernel to mount an existing SUSE installation and traverse the `init` tree to allow the user to repair or use the system. You may encounter problems or see error messages because the modules of the installed system will very likely not work with the SUSE install kernel. This will enable your system to boot, but you may not be able to access specific devices that are associated with modules that could not load.
- ♦ **Abort installation** — Aborts the installation process and reboots your system. If you really do not want to proceed with installing SUSE Linux at this time, make sure that you remove the SUSE boot media from your optical drive quickly.

In most cases, you will want to select New installation. The remainder of this section concentrates on a new SUSE installation.

Selecting New installation takes you to the Installation Settings screen (see Figure 1-6).

Clicking any heading in this section enables you to modify that aspect of your installation. Similarly, selecting the Change button displays a pop-up from which you can select any of the headings on this screen to change or examine the relevant aspects of the installation to guarantee that they meet your requirements.

- ♦ **System** — Displays a dialog showing the hardware that the installer detected in your system. You cannot change these values.
- ♦ **Mode** — The installation mode. This was selected in the previous screen, but here you can change this during customization if needed.
- ♦ **Keyboard layout** — Select the keyboard set used for the system and the installation process.
- ♦ **Mouse** — Select the mouse type for installation. This also helps with the configuration of X later in the process.
- ♦ **Partitioning** — One of the most important aspects of installing a Linux system. Partitioning configures the target hard drive for the installation of an operating system.
- ♦ **Software** — Selection of predefined software profiles, as well as individual software packages.
- ♦ **Booting** — Configuration of the Linux boot loader. The boot loader bootstraps a loader at bootup that allows the user to boot not only Linux, but also any other operating systems in the system.

- ♦ **Time zone**—Set the time zone of the system based on either your location or specifically setting the GMT offset.
- ♦ **Language**—Set the language of the system and also the installation process. This was already set earlier in the installation process.
- ♦ **Default Runlevel**—Set the initial boot runlevel for the system. Runlevels are discussed in Chapter 4. For now, the default value (runlevel 5) is acceptable.
- ♦ **Reset to defaults**—Remove all changes you have made and start from scratch. This is useful for testing installation mixtures and seeing how these affect your system.

Throughout the remainder of the installation, we talk in more detail about what these settings do to your system and we also discuss the ways in which you can change these settings.

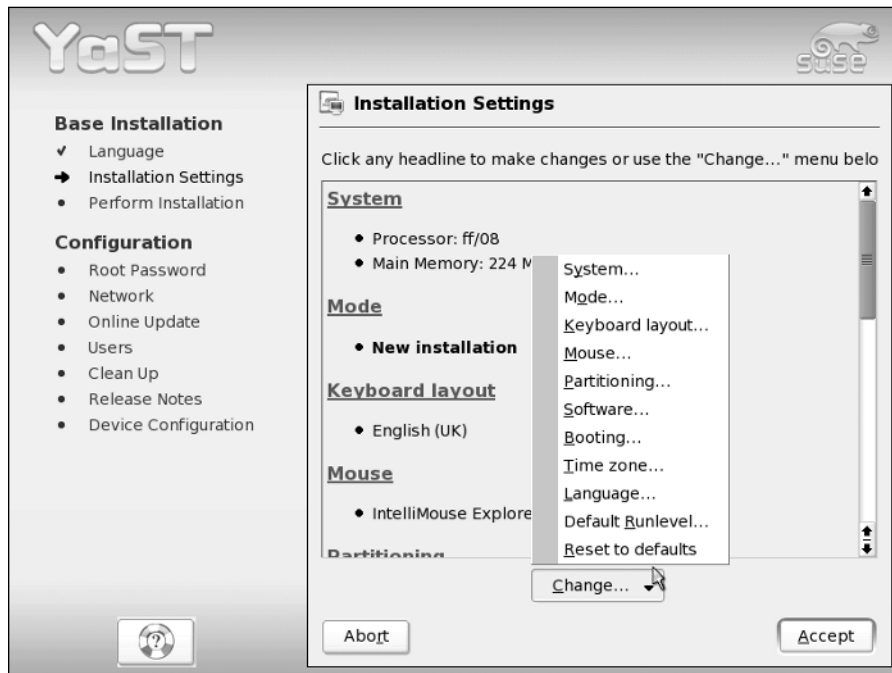


Figure 1-6: Installation customization

Partitioning your disks

YaST initially chooses a partitioning scheme based on your disk layout. It is very likely that the installation default will be fine (see Figure 1-7) for a first-time user. For other users, YaST allows you to control the layout of partitions on the disk, the type of filesystems that are used on those partitions, and any options that will be used when mounting them.

One key thing to know when defining and experimenting with disk partitioning is that none of the changes that you are defining are actually performed until you explicitly tell YaST to proceed with the installation. You can make as many changes or experiment with different partitioning schemes as much as you want without actually committing those changes. Aborting the SUSE Linux installation at any time before this point will leave your system's disk exactly as it was when you started the installation process.

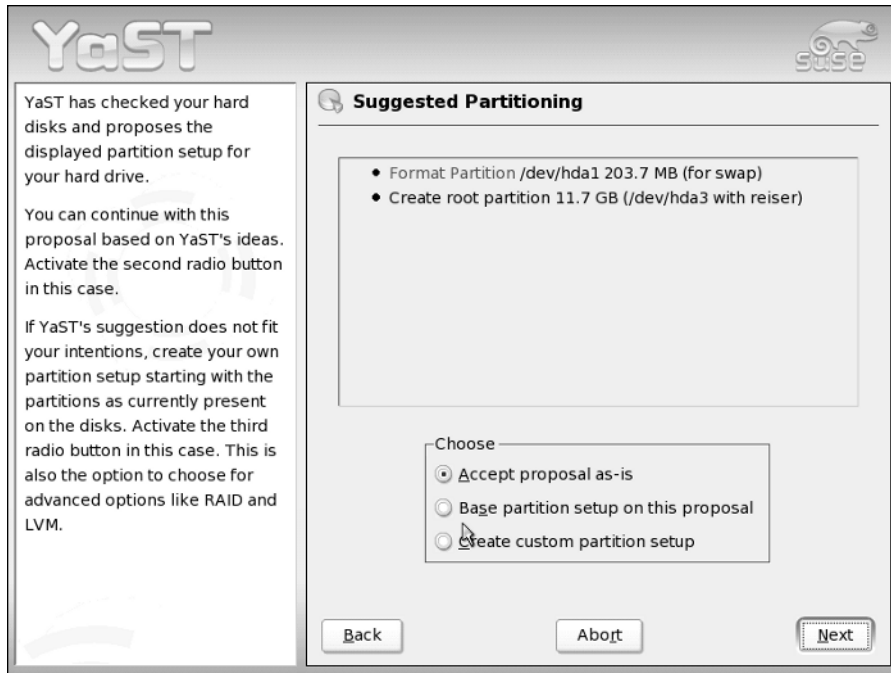


Figure 1-7: Partitioning

What you do next depends on your requirements:

- ♦ If you want to accept the default partition layout selected by YaST, skip ahead to the section of this chapter entitled “Selecting software for installation.”
- ♦ If you are installing SUSE on a computer where another operating system, such as Microsoft Windows, is already installed and you want to preserve that operating system or any of the data that it contains, read the next section for details on the changes that YaST will make to your disk.
- ♦ If you are an experienced Linux user and want to specify your own customized partitioning scheme, select Partitioning from the menu shown back in Figure 1-6 to have full control over the partition layout of the system. Select the Create custom partition setup radio button shown in Figure 1-7, and click Next. Then, select the Custom partitioning - for experts radio button, and click Next (see Figure 1-8). This will present you with the option to create and delete partitions, as well as other advanced options such as software RAID and cryptographic filesystems.

If you are creating your own partitioning scheme and do not already have an operating system on your computer that you want to preserve, skip to the section “Primary and extended partitions.”



Figure 1-8: Selecting custom partitioning

Resizing existing operating system partitions

Nowadays, it is quite common to have systems that can boot multiple operating systems. Such computer systems enable users to take advantage of the power and applications available in each operating system by selecting between available operating systems when you boot the system. These are typically referred to as *dual-boot* systems because most people install at most two operating systems on a single machine. However, because more than two operating systems can be installed on a single disk, the proper name is *multi-boot*, which is the term used in this section. The number of operating systems that you can boot and run on a single computer is really limited only by the amount of disk space available on your computer system.

With SUSE Linux, the most common type of multi-boot system is a system that can boot either SUSE Linux or some version of Microsoft Windows. Windows will be used as an example throughout the rest of this section, although the same general concepts are true when setting up multi-boot systems that will run SUSE Linux and any other operating system.

Explaining how to install Windows on an existing SUSE Linux system is certainly not relevant to a discussion of installing SUSE Linux. However, the reverse is not true. Installing SUSE Linux on a system that already runs Windows, and on which you want to be able to continue to run Windows, is a common wish. This is quite easy to do and involves only resizing your existing Windows partition(s) so that sufficient contiguous space is available for installing SUSE.

If you have selected a new installation on a system that already contains an operating system such as Windows that you want to preserve, and if the disk or Windows partition in that system has sufficient free space to install SUSE Linux, YaST will propose a solution based on resizing your existing Windows partition and automatically creating appropriate swap and root partitions. If at all possible, you should accept this default selection.

If you do not have sufficient free space to install SUSE Linux and YaST cannot automatically resize your existing operating system partitions, your only alternative (besides adding another disk to your system) is to abort the SUSE install process, remove the installation media, and reboot into your other operating system. You must then free up sufficient disk space and clean up the organization of your operating system's partition(s) using a utility such as Windows' Disk Defragmenter. If there is sufficient unused space on your Windows partition, you should then be able to restart the SUSE installation process and let YaST select appropriate partitioning and resizing values for you.



For more on setting up dual-boot (multi-boot) systems, see Chapter 4.

Primary and extended partitions

In this section, we start with a clean disk to create the partitions needed to install SUSE. If you wish to remove the partitions on an existing installation of an operating system, select the partition and press the Delete button. You will be asked to confirm this, and the partition will be removed.

If you select Create, you are prompted for the type of partition you wish to create (see Figure 1-9). In the PC world, the BIOS can access only four primary partitions. These can be thought of as four physical boundaries on the disk, with separate data and filesystems on each. With Linux, you need at least two partitions, and if you have Windows on another partition, and a data or home disk on the other, you may quickly run out of ways to expand the way your disk is laid out. To combat this, logical and extended partitions were designed. An *extended partition* is a placeholder for further logical partitions, and it is a good idea to create one extended partition (which takes up one of your primary partitions) and create logical partitions to accommodate further partitioning schemes in the future.

The most common way to partition disks for home Linux use is to have one primary partition for the Linux root partition, a second primary partition for the swap partition, and then an extended partition for any other (logical) partitions that may be needed. Using extended and logical partitions grows the amount of total partitions you can have in a system to 16, which is usually more than enough.

Select the Primary partition radio button and click OK to proceed.

Defining filesystems

Once a primary partition has been created, you need to define the format in which a filesystem should be created on that partition, its size, and the mount point for that filesystem. Linux and Unix use the definition of mount points in the same way that Windows uses drive letters. The advantage with Linux is that the whole system is hierarchical in nature, and therefore access to data on disks, network drives, and partitions can be kept under one manageable tree structure.

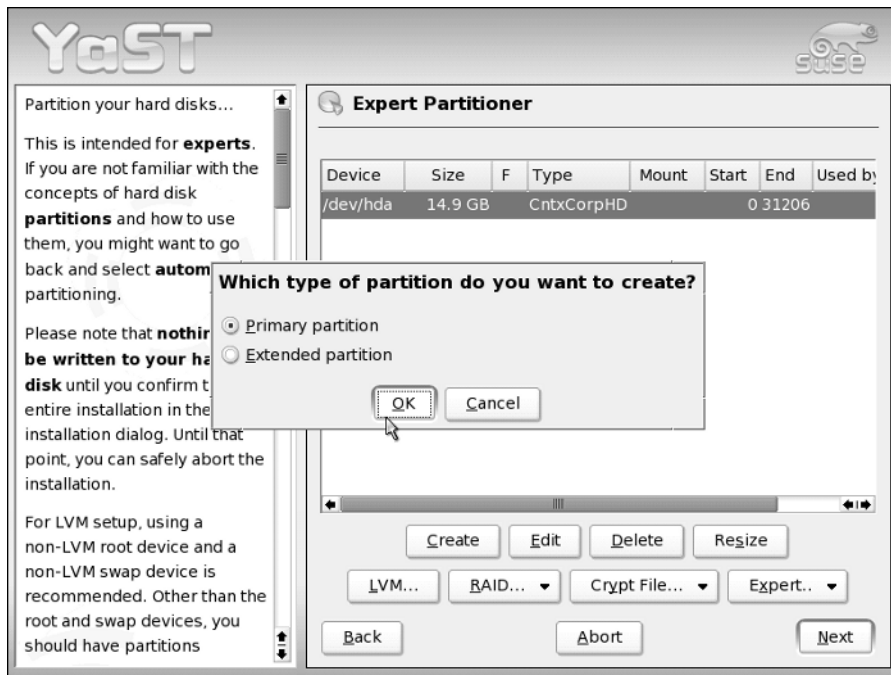


Figure 1-9: Creating a partition

Swap partitions

The first partition you need to create is the *swap partition*. Most modern operating systems use swap partitions, also referred to as *swap space*, to support virtual memory. Virtual memory is a technique for enabling a system to use more memory than is physically available to the operating system. Processes on the system that are inactive or are waiting for input are copied from physical memory into swap space, known as *swapping out a process*. At this point, the physical memory associated with those processes can be reused by the operating system. When the process can run again, such as when input is available, it is copied from the swap space back into memory and continues execution. This is known as *swapping in a process*. The way in which processes are swapped in and out of memory is simple in theory, but is triggered by a number of internal metrics that are maintained and constantly updated by the kernel.

You should always create a swap partition on a Linux or Unix machine as the workload on any system can never be fully quantified beforehand and running out of physical memory without swap space causes processes to crash or be unable to execute in the first place.

The window to create a filesystem/partition can be quite daunting for new users (see Figure 1-10). SUSE and the other distributions try to make the process as simple and usable as possible. Selecting the format of the filesystem is primarily a concern when creating data partitions or for advanced users, as discussed in the next section. When creating a swap partition, you must select Swap as its format. You will notice that the mount point will also change to be swap because the swap partition is not mounted like a data partition but is used internally by the Linux system.

Cross-Reference

Filesystems are discussed in more detail in Chapter 3.

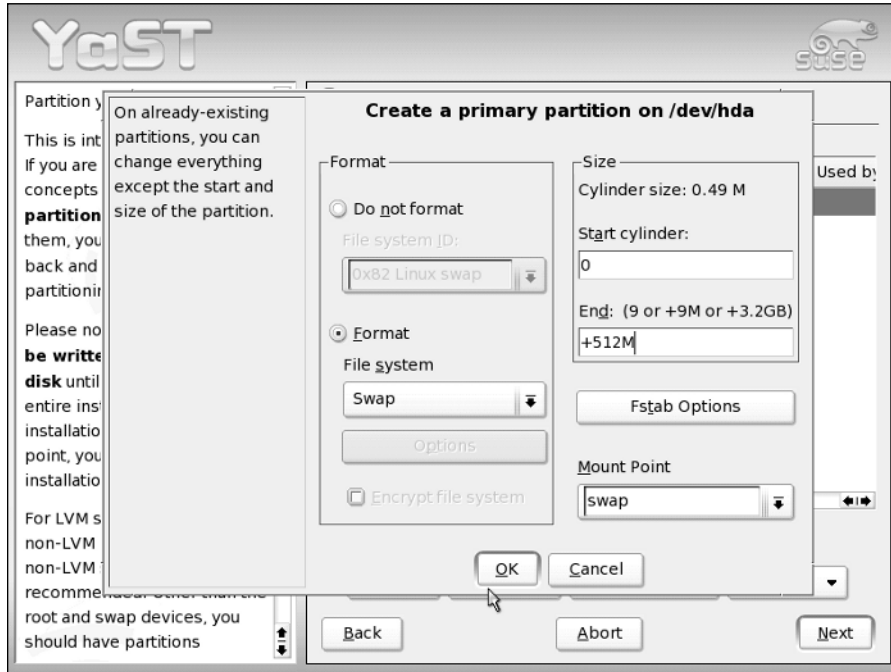


Figure 1-10: Creating filesystems

Start and end cylinders are often new concepts to new Linux users who are used to data sizes being defined in mega- and gigabytes. YaST allows you to enter the size of a partition in human readable form, such as MB and GB. The start cylinder, as this is the first partition on the disk, is 0 (the start of the usable space on the disk), and the end cylinder is what we need to change. It is usually customary to select a swap size that is 1.5 times the amount of physical RAM in the system, but this is subject to much conjecture. A reasonable swap size should be considered based on the workload of the machine you will be using, and as most modern PC systems have at least 512MB, it is safe to use the standard 1.5 times physical memory. To specify that you want the swap partition to be 750MB, enter **+750M** in the End cylinder entry box. The **+** signifies that you want to add space, the number is the unit of space needed, and the **M** specifies that the amount of data is expressed in megabytes. You can also specify **G** for gigabytes, which you will be using in the following example of creating a root partition.

After entering the size of your new swap partition, click OK to proceed.

Tip

At a bare minimum, the filesystems that need to be created are the swap space and a root (**/**) filesystem. However, for ease of use and manageability, the creation of a **/home** partition can help keep your personal data separate from the system partition and also allow you to keep your data in the unlikely event that you want to do a total reinstall of Linux. See the section on “Data partitions” later in this chapter for more information.

In this example you are creating the bare minimum, the swap and root partitions.

The root partition

Once the swap space has been created, you need to configure the root (/) partition (see Figure 1-11). The root (/) partition is the most important data partition on any Linux or Unix system, and is the only non-swap filesystem partition that is required in order to boot a Unix or Linux system. The *root partition* takes its name from the fact that it is the partition mounted at the root of the Unix/Linux filesystem, which is the directory known as “/”. A filesystem must be mounted on this directory to successfully boot a Linux system. The root filesystem contains core directories required to boot Linux, such as the directory through which devices are accessed (/dev); the directory containing system administration, configuration, and initialization files (/etc); the directory in which critical system libraries, kernel modules, security, and internationalization information are located (/lib); and directories containing critical system binaries (/sbin, /bin, and so on).

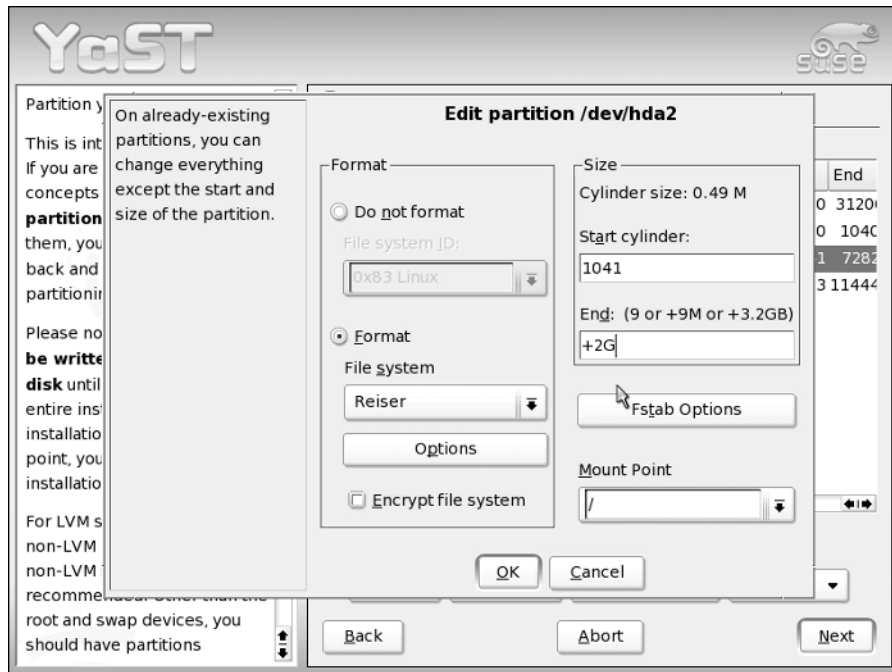


Figure 1-11: Creating the root partition

By default, creating this partition will automatically use the remaining unallocated space on the hard drive, which is fine for our example. However, if you need to create another partition, /home, for example, you specify the size of the partition explicitly as you did with the swap space. See the next section, “Data partitions,” for an overview of why you may want to create additional partitions.

The default type of filesystem used in SUSE is the Reiser filesystem, often referred to as the ReiserFS. It was one of the first available journaling filesystems for Linux, and a lot of the work was funded by both SUSE and mp3.com. A *journaling filesystem* dedicates a specific part of the filesystem for use as a cache of pending writes to the filesystem; this ensures that filesystem updates occur in a clean, atomic fashion; and allows a fast recovery if the system is not cleanly shut down. Ordinarily, when a Linux system is shut down, it ensures that all

pending writes to each filesystem have completed, and then detaches the filesystems (known as *unmounting* them) to guarantee that all system data is consistent before the system is turned off. Using a journaling filesystem does not mean it is safer to just power off the machine, as data loss can still occur when data is not completely written to the disk.

Once the root partition has been created, you can review your changes (see Figure 1-12) and proceed with the software installation by clicking Next. If you want to create additional filesystems during the installation process, read the next section before clicking Next.

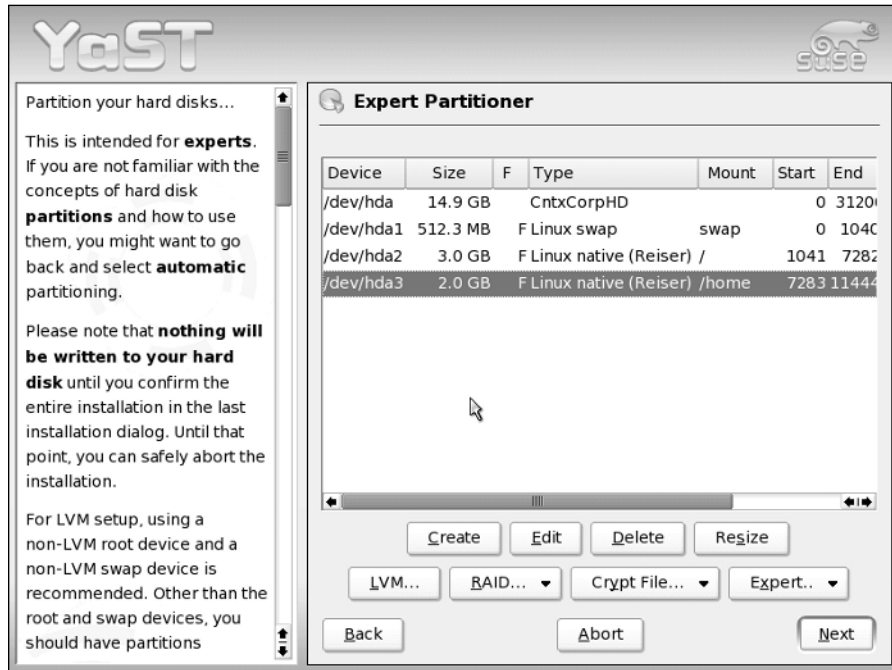


Figure 1-12: Reviewing changes to the partition scheme

Data partitions

Data partitions is a generic term for partitions that are formatted as a filesystem and in which both the system and its users can store data. The partition designated as the root filesystem is a special case of a data partition because it is required in order to boot a Linux system.

The preceding sections explained how to create the swap and root partitions that must be present to successfully boot a Linux system. However, you can also create other data partitions, format them as filesystems, and specify their mount points during the installation process. On Linux systems, a *mount point* is simply a Linux directory through which a filesystem is made available to the system, known as *mounting* that filesystem. Using regular directories as mount points is a clever part of the design of Unix and Linux. If you run out of disk space

on a given partition, you can add another disk to your system, create data partitions there, copy the data from existing directories to those partitions, and then mount the new partitions on the directory where the data was originally located, effectively increasing the amount of storage available to an existing system.

Today's larger disks make it attractive to create other data partitions. You have several reasons to consider creating multiple data partitions on today's disks:

- ♦ When you boot a Linux system, the system checks the consistency of each of its filesystems (as defined in the file `/etc/fstab`—more about this in Chapter 3). Checking the consistency of a single, huge, nonjournaled filesystem can take quite a bit of time.
- ♦ Filesystem corruption can occur as a result of a number of problems, such as a system crash, sudden power loss, or hardware problems. Whenever a filesystem is corrupted, repairing it (which is mandatory) can cause you to lose data. Creating multiple partitions reduces the extent to which filesystem corruption can affect a single data partition.
- ♦ Keeping data on multiple partitions limits the chance that you can lose data during a subsequent system upgrade. Some upgrades reformat the root partition or recreate its directory structure. If your user data is stored on other data partitions, they will not be affected by changes to the root filesystem.
- ♦ Some Linux backup software backs up data on a per-partition basis. Backing up a single huge partition can take quite a bit of time. Also, if your backups fail (such as when a tape is corrupted), you may not be able to use the backups to restore your system. Creating multiple partitions limits problems related to a backup failure to a single partition.

Chapter 3 provides more detail about creating multiple partitions and the types of filesystems supported by Linux and provides additional reasons why you may want to create multiple partitions on your Linux system. Most types of Linux filesystems can be resized once they have been created, enabling you to customize your system's partitioning, even after the system has been installed and is running.

If you want to create multiple partitions during the installation process, you can do this by making sure that the root partition does not completely fill your disk and then creating additional partitions in the remaining space on your disk. Common parts of a Linux system that you may want to put on separate data partitions are `/boot`, `/home`, `/opt`, `/tmp`, `/var`, `/usr`, and `/usr/local`. For more information on these partitions and the types of information stored there, see Chapter 3.

Selecting software for installation

The software that is automatically selected as part of a default SUSE installation provides you with nearly every type of software required for day-to-day work. This section offers additional details about the other types of installations provided by the SUSE installer to provide a full and thorough SUSE learning experience.

To customize the software that is included as part of your SUSE installation, you must click the Software heading in YaST's Installation Settings panel, or click Change and select Software from the popup menu. Doing either of these displays the pane shown in Figure 1-13.

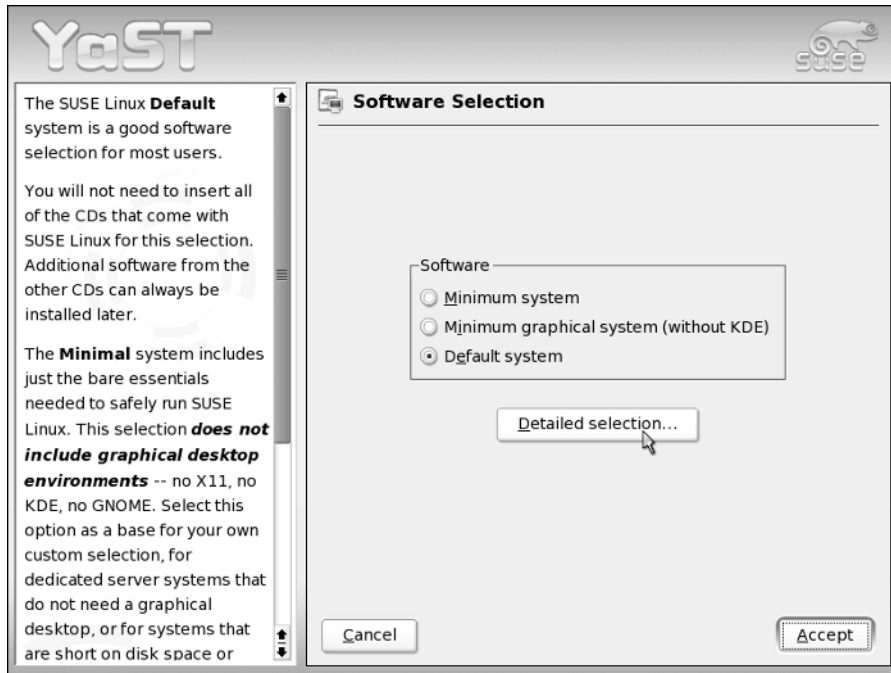


Figure 1-13: Defining the software installation

The predefined choices available to users are as follows:

- ♦ **Minimum system** — The bare minimum usable SUSE system with no graphical user interface. This is something that we use a lot during server installation to dramatically lower the amount of software to install and to be extremely granular in the selection of software packages.
- ♦ **Minimum graphical system (without KDE)** — This will install a more usable system but without the optimized KDE (K Desktop Environment). We deal with KDE and desktop environments in Chapter 8 of this book, but for now, just know that the option is available. Some application software configuration tools (such as those used by Oracle) require the X Window system but do not require all of the bells and whistles provided by desktop environments such as KDE.
- ♦ **Default system** — This is the standard SUSE system, which includes KDE, OpenOffice, web browsers, a music player, and so on. This is what we will use to base our customized installation on and is what most users will be happy installing.

Selecting the Detailed selection button displays the YaST package manager (see Figure 1-14). This is a powerful installation tool that “other” distributions lack and saves a lot of time by calculating package dependencies automatically, providing a search function for not only package names, but also package descriptions.

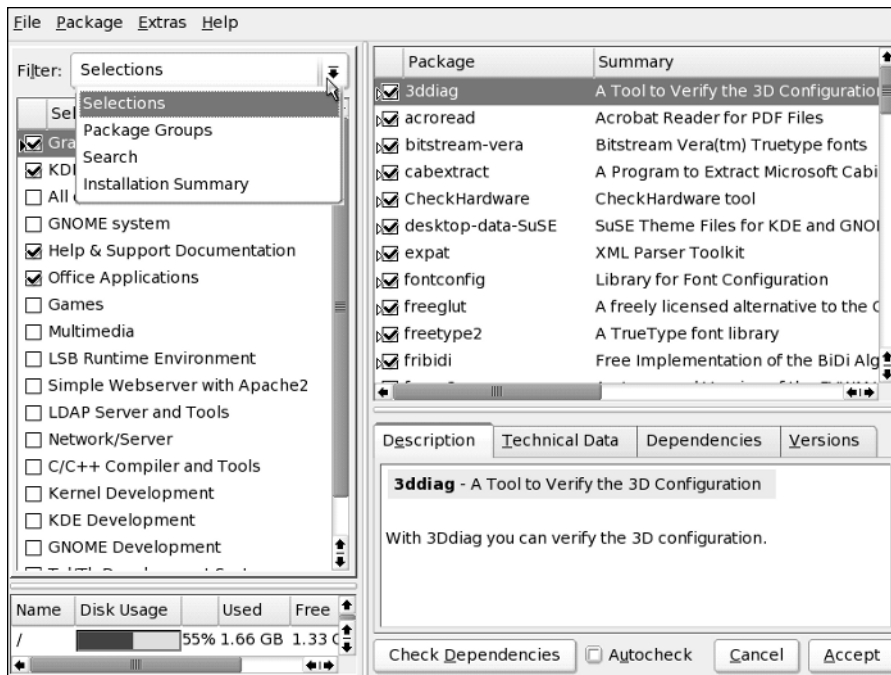


Figure 1-14: Using the YaST package manager

The left panel of the package selection screen gives a broad overview of packages and disk space used, and also indicates how much space will be required when installing all selected packages.

The Filter drop-down box provides a powerful way to limit what packages you can select. As we stated in the Introduction, we ourselves differ on our views of a few Linux idiosyncrasies, including text editors (Roger and Bill: emacs; Justin: vim), desktop environments (Justin: GNOME; Roger: KDE; Bill: a window manager), and also Linux on the desktop in general. SUSE is aware of differing views throughout the whole Linux community, and therefore using package selections allows users to specify things such as which desktop environment and editor they want to install — why waste disk space if you're not going to use something. The same is true for games, multimedia, and specific server software. The amount of disk space required to install your system can be reduced or enhanced by selecting specific packages.

We will keep the default package selection as chosen by SUSE and add a new package that is not installed by default.

Selecting Search from the drop-down box allows you to enter search criteria for a package and returns all results based on the "Search in" criteria selected. Figure 1-15 shows a search for the Blackbox window manager. As you can see, YaST returned not only the package Blackbox but also other packages that contain the word *blackbox* in their summary definitions, which can be seen in the Description window.

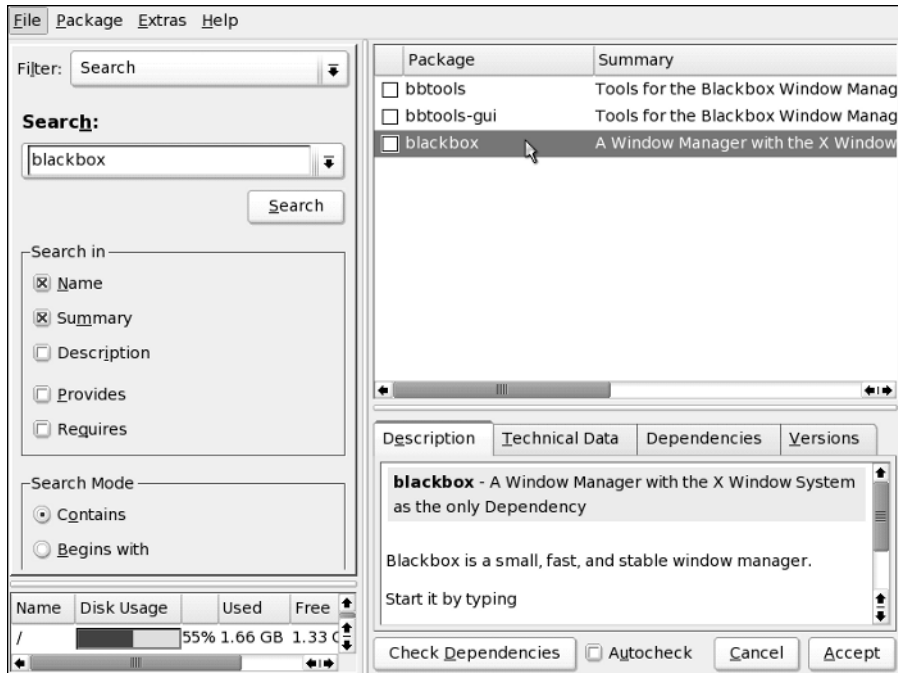


Figure 1-15: Searching for individual packages

When a package is selected, your disk usage will be increased to reflect the size of the install domain.

Select Accept to add those packages you select to the install list and take you back to the package installation summary screen.

Selecting a boot loader

The next item you can change is the configuration of the boot loader. A boot loader is central to the deployment of Linux as it controls the booting of operating systems on the PC. To customize the boot loader that is used by your SUSE installation, you must click the Booting heading in YaST's Installation Settings panel, or click the Change button and select Booting from the popup menu. Doing either of these displays the pane shown in Figure 1-16, YaST's Boot Loader Setup screen.

Linux systems typically use one of two boot loaders, LILO (Linux Loader) or GRUB (Grand Unified Boot Loader). Both are very powerful and flexible, and are controlled by easily edited configuration files (`/etc/lilo.conf` and `/etc/grub.conf`, respectively). The key difference between the two boot loaders is how they interact with these configuration files. If you use

LILO and update its configuration file, you must rerun the `lilo` command to update the system boot information that is stored on your disk. GRUB automatically rereads its configuration file whenever you boot your system and therefore does not require that you update any other system boot information.

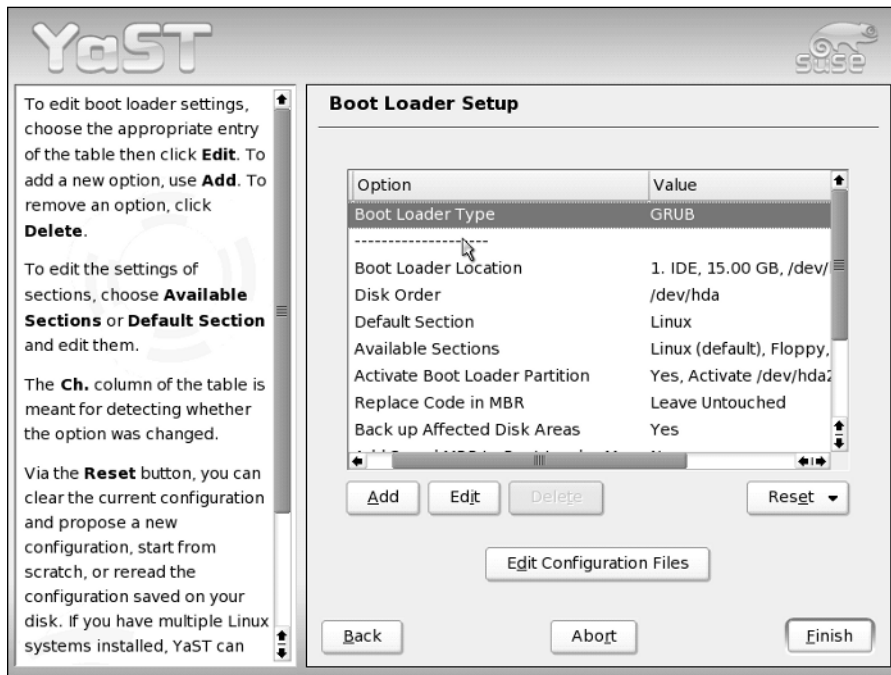


Figure 1-16: Boot loader configuration

A few years ago, the general consensus was to move away from the LILO boot loader to the GRUB boot loader. GRUB provides a more robust boot loader, and the default configuration is fine for most users. If YaST detects a Windows installation, it adds this as a boot option, providing a means to dual-boot Windows and Linux on the same system.

Selecting an item in the Option list and pressing the Edit button provides you with a selection window for that specific option (see Figure 1-17). If you want to use the traditional LILO boot loader, you can make that change here. You can also decide not to use a boot loader at all and boot your Linux system manually with the SUSE install CD. (There are other options, but we won't discuss them here because they are intended for use only in special cases and are not something that most users will want to do.)

When you make any changes that you want on the Boot Loader setup screen, click the Finish button to return to the standard YaST installer screen.

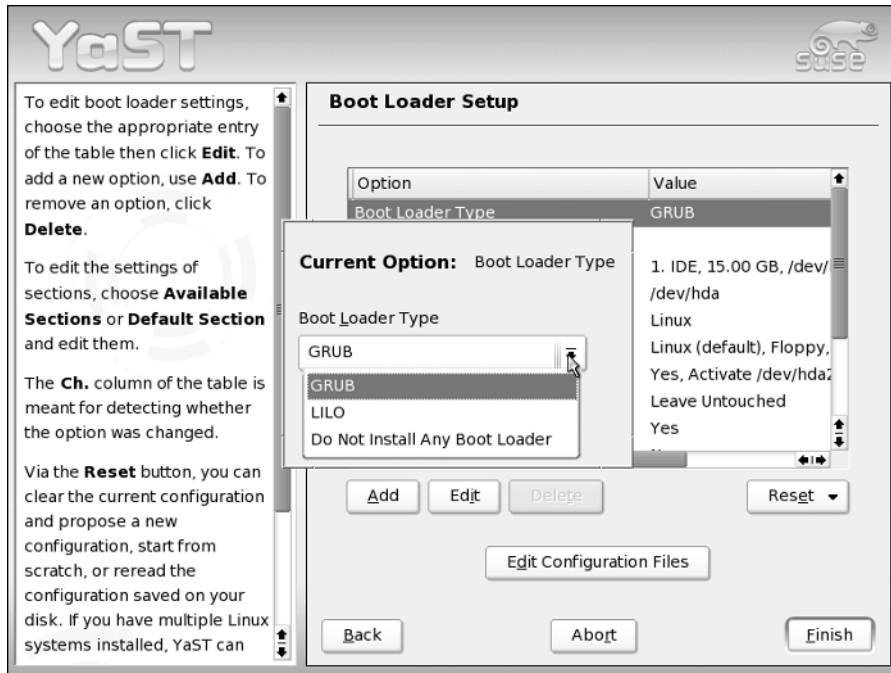


Figure 1-17: Selecting the boot loader

Changing the default runlevel

Runlevels are discussed in detail in Chapter 4. As a quick summary, a system's *runlevel* determines the services that are automatically started when your system boots. The YaST Installation Settings screen can be used to change the default runlevel of the system by clicking the Default Runlevel heading in YaST's Installation Settings panel or by clicking the Change button and selecting Default Runlevel from the popup menu.

As you can see in Figure 1-18, you can choose to boot your SUSE system in a variety of different ways: without networking functionality (runlevel 2), multiuser with network (runlevel 3), or multiuser with X Windows (runlevel 5). The default runlevel in a standard installation is runlevel 5, multiuser with X Windows. You should keep this as your default runlevel unless you have a specific reason to change it.

When you make any changes that you want to your system's default runlevel, click OK to set the selected runlevel as your system default. The Set Default Runlevel popup closes, and YaST's Installation Setting panel displays.

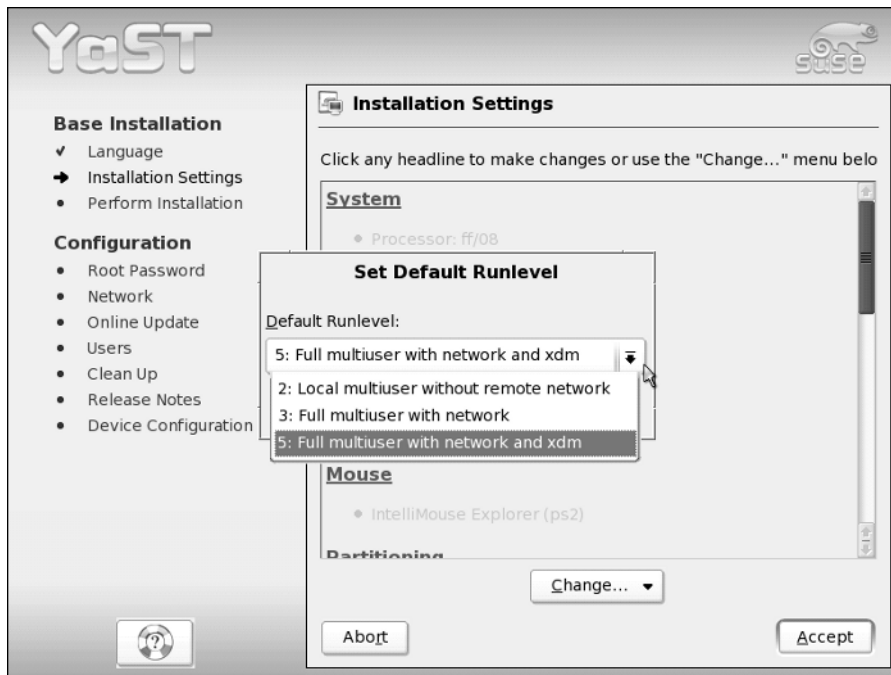


Figure 1-18: Changing the default runlevel

Running the Installation

When you've made any changes to the installation, select the Accept button in the Installation Settings window, and the installation process begins. You will be asked if you definitely want to create the partitions you defined and the filesystems that sit above them (see Figure 1-19). Partitioning the disk is a destructive process and *will* remove any data that those partitions replace.



This is your last chance to abort your installation without making any changes to your disk. You should continue only if you are sure that the selected settings are correct. If you are installing SUSE for the first time on a new computer system, you have nothing to worry about. If you are installing SUSE on an existing computer system on which you need to preserve existing data, double-check your settings before proceeding. You can double-check that your partitioning scheme is, in fact, correct for your environment and make changes as necessary by selecting Partitioning from the Installation Settings screen. Then triple-check your selections before proceeding.

Selecting “Yes, install” will destructively create the partitions and filesystems and install the packages you selected.



Figure 1-19: Confirmation of installation

During package installation, you can view the progress, the CDs needed, and also an overview of the package use by selecting the Details button. To switch back to the slideshow, select the Slideshow button.

This may be a good time to get a coffee, as nothing interesting happens while packages are installed. If you are installing from CDs, after the packages from CD 1 are installed, the system will automatically reboot itself to use the specific kernel for your architecture, as well as to commit packages installed on the system.

If you think back to the discussion of the install boot options at the beginning of the chapter, you will remember that the default is to boot off the hard drive first. This helps a lot if you leave the install media in the drive and are drinking your coffee in another room.

When the system has rebooted, YaST asks you for the remaining media to install the rest of the packages. In the case of a minimal installation, or all packages being installed, YaST proceeds automatically to the system configuration.

Configuring your root password

The first item that needs configuration is the root password (see Figure 1-20). We will talk about the root user in Chapter 2, but for now it is enough to know this is the user who has the privileges to change anything on the system, has access to all files on the system, and is known as a *superuser*.

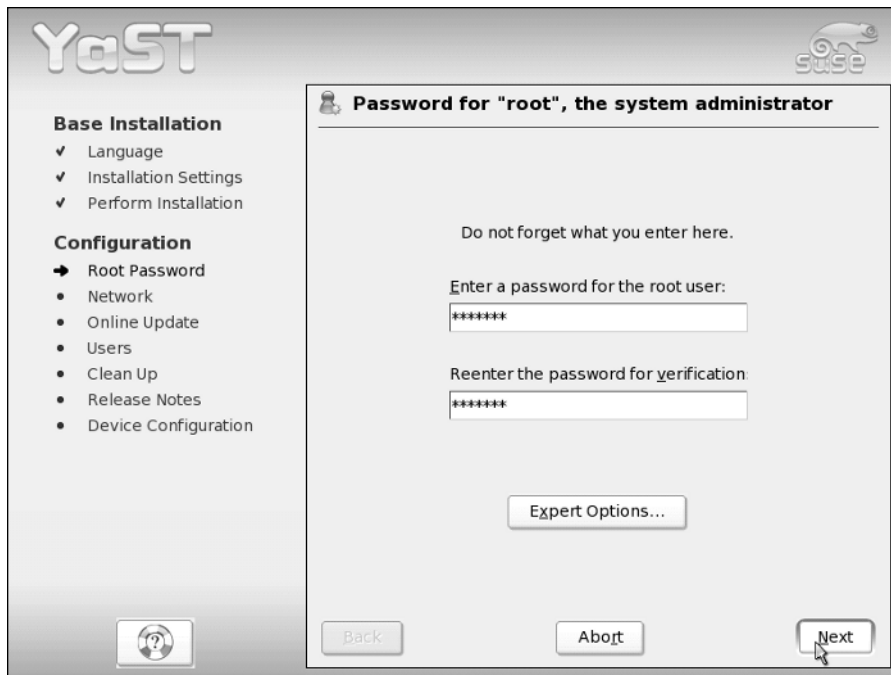


Figure 1-20: Setting the root password

The password should be something that you can remember, but also difficult to guess. A combination of letters and numbers is always a good way of making a strong password. Using your name, family member names, and so on should be avoided as these can be easy targets for passwords. Click Next once you've entered your root password and re-entered it for verification.

Configuring your network access

If any network interface cards have been detected in the system, you will be asked to configure them for network access (see Figure 1-21). By default, YaST sets the first Ethernet card it finds as your system's primary Ethernet interface and assigns it an address that is configured via the Dynamic Host Configuration Protocol (DHCP).



Cross-Reference

You can find discussions about DHCP servers in Chapter 19.

For most people using SUSE in a business environment, a DHCP server may already be running, and an address, domain name system (DNS) server list, and router configuration will already be available. Home users and users setting up a server system will find it necessary to configure these details manually. Home users with simple broadband or dialup connections will often automatically receive this information from their Internet service providers (ISPs) and therefore may not need to change these settings.

To change the configuration of the network card, click the Change button. You are then asked to select the network card that you wish to work with (see Figure 1-22).

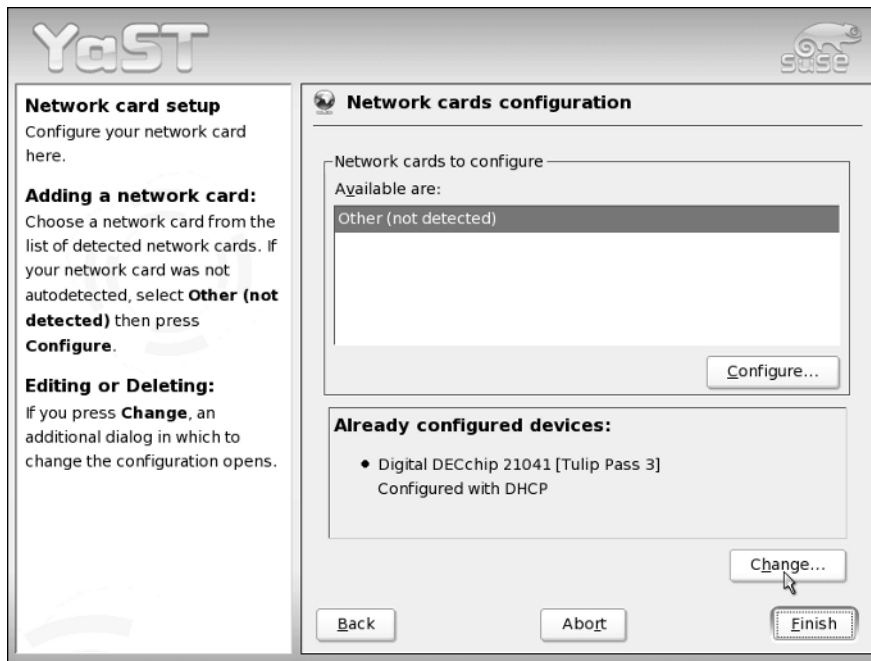


Figure 1-21: Configuring network cards

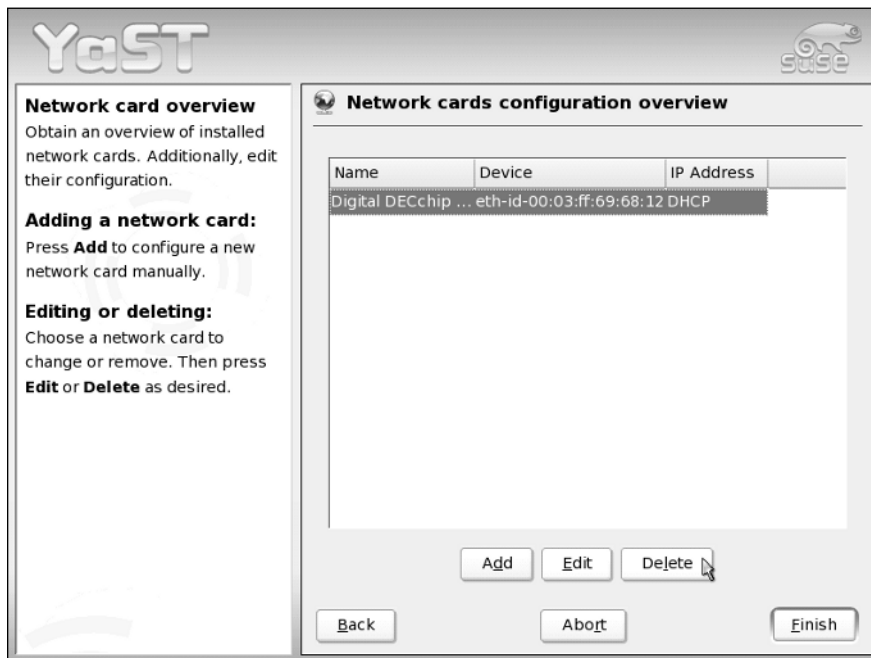


Figure 1-22: Selecting your network card

In this example case, you have only one network card that has been configured by the system. If you had more cards configured (for example, one network card with DHCP and one with manual configuration), you would see the network card type, the configuration options, and also the device ID as seen by Linux for each.

During this configuration, you can remove and add any cards that have been defined, as well as change the configuration of an existing card. In this example, you will change the default SUSE settings to a manual configuration, so select the Edit button. A screen like the one shown in Figure 1-23 appears.

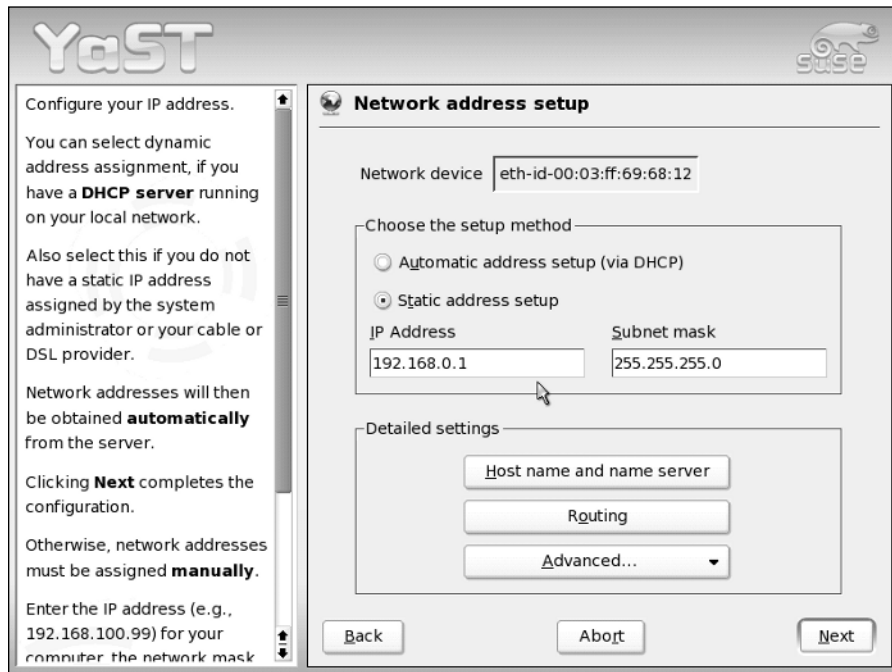


Figure 1-23: Configuration of network cards

In this example configuration, we will set the IP address of the network card to 192.168.0.1/255.255.255.0, with a router/gateway of 192.168.0.8 and a DNS server of 192.168.0.254. If you are unfamiliar with these terms at this stage, see Chapter 6 for additional information.

To change the configuration of the network card from automatic to manual, select Static address setup. This enables you to edit the IP and subnet mask fields. As you can see in Figure 1-23, we have set the IP address/netmask to that of the configuration we talked about in the preceding paragraph.

Setting up your host name and DNS addresses

To set up the host name of the Linux machine and the addresses of your Domain Name System servers, select the Host name and name server button. A screen like that shown in Figure 1-24 appears.

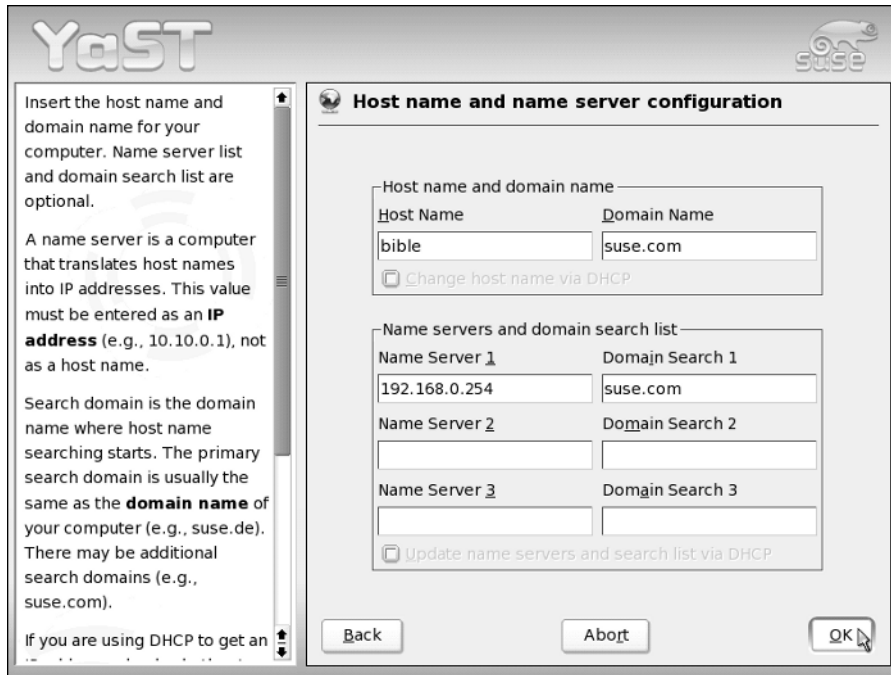


Figure 1-24: Configuring DNS and host name

The host name of your Linux machine can be anything you like, such as a person's name, a descriptive name, or something random. The only thing that you have to bear in mind is that the host name and domain name can contain only letters and numbers as well as a hyphen or an underscore. The host name can be only one string of characters and cannot contain a space or a period. As the name suggests, the domain name dictates the network "domain" that this machine falls into. This domain may well be something in line with your company's policy or could be something you have set up yourself.

Tip

When integrating a new system into an existing networked environment, you should always follow the same naming conventions that are already being used, especially for the domain name. If you do not, other systems on the network may not be able to locate your system correctly, and certain services on your system may not be able to interoperate with existing network services.

Enter the name server address into the Name Server 1 field. You can also enter up to two other separate DNS server entries. Your administrator or ISP should be able to give you this information.

The Domain Search entry is used to control how your machine looks up the address of other machines connected via TCP/IP. For example, if you use a Domain Search entry such as `suse.com`, you can communicate with any machine in the SUSE domain by just its host name. For example, with `suse.com` as the Domain Search entry, you can communicate with the machine you are setting up in this example by just using the host name of `bible`. If you do have `suse.com` as a Domain Search field, however, you have to specify the fully qualified domain name of the machine you wish to communicate with (in the case of this example, that is `bible.suse.com`).

When you have set the DNS configuration for your system, press OK to save your changes.

Configuring the default gateway

Next, you will probably need to configure the router/gateway for your system. To do this, click the Routing button. You should see a screen like the one shown in Figure 1-25.

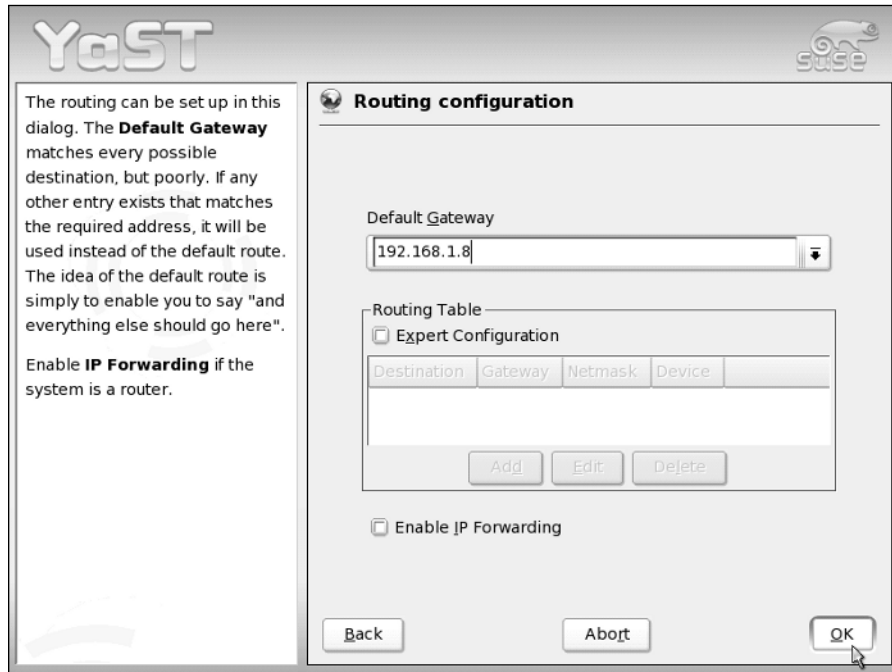


Figure 1-25: Configuring a default gateway

Your default gateway address is the IP address of the host to which TCP/IP packets that are not destined for your local network are sent for further processing. For example, your gateway address will be that of your asymmetric digital subscriber line (ADSL) router if that is how you connect to the Internet. In other cases, your network or system administrator will be able to provide you with this information.

When you have set the gateway address, click OK to proceed. If you have finished configuring all of the network cards that you need to configure, you can select Finish in the Network cards configuration overview window. This tells YaST to save the changes to your network configuration and restart the system networking.

Testing your connection and online updates

Taking a page from the "other" operating systems, SUSE now allows you to run the update service as soon as the system has been installed (see Figure 1-26). Online updates are discussed in more detail in Chapter 9 and are not discussed here as part of our sample installation. If you are feeling adventurous, then testing your Internet connection and running the online update is a good idea, but it is not necessary as part of the installation process. Click Next after you have made your choice.

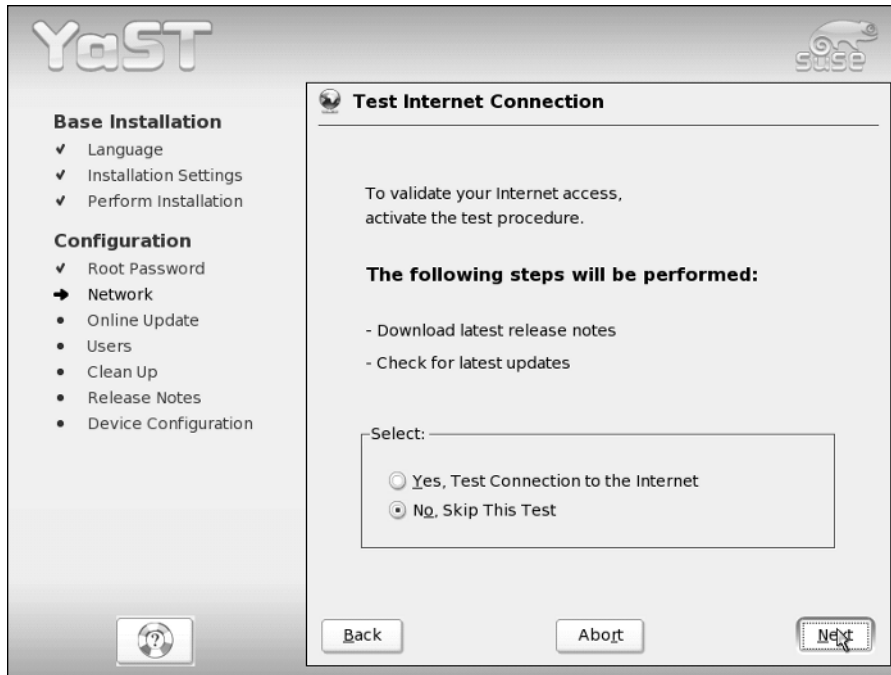


Figure 1-26: Online update

Configuring your modem

Modems are notorious for being something of a problem area for Linux because many of the internal PCI modems that are on sale are not true modems, but are what is known as *winmodems* or *soft modems*. The idea behind winmodems is that some of the functionality of the modem can be offloaded from hardware into software—the software in question being the Windows operating system. As these devices are designed to work only with Windows and in theory require a Microsoft operating system to work, it is not surprising that there are difficulties getting them to work on Linux. So there are three possibilities:

- ♦ **You have an old-fashioned external serial modem.** These will always work under Linux.
- ♦ **You have a winmodem.** This may or may not work with Linux.
- ♦ **You have an internal true hardware modem.** In almost all cases this will work with Linux.

Winmodem support has improved considerably, and some previously unsupported modems now work with Linux. SUSE 9.1 supports at least the SmartLink and Lucent Itmodem types. During the installation, if YaST detects that you have a supported winmodem, it will install the necessary package to provide driver support.

Assuming that your modem is supported, YaST asks you for the necessary information to set up the modem and your dial-up connection. In the first screen of the setup you are asked for a dial prefix if necessary (a prefix you might need to get an outside line, for example) and to choose between tone dialing and pulse dialing (tone dialing will be the correct choice unless your telephone exchange is very antiquated). The other two choices here (Speaker on and Detect Dial tone) you will almost certainly want to leave as they are (selected by default).

The next screen asks you for your country and offers a small selection of preconfigured providers (ISPs). This means that the access number is already known to the system for these providers. One or two have gone further and allow you to use a preconfigured username and password to sign up with them, or even to get full anonymous access with payment being collected through your phone charges.

If you already have an account with an ISP that is not listed, you need to press the New button and add the details of the provider's name, the access phone number, and your username and password.

When this is done, press Finish and the modem configuration should be complete. You will then be able to connect using the `kintern` program, which you can access through the KDE menu (Internet ⇄ Dialup). You can set `kintern` to run whenever you log in to KDE; if you do this, you can log in by clicking its tiny icon, which will be resident in the "system tray" area of the KDE panel (by default in the bottom-right corner of the screen).

**Note**

More information about using `winmodems` under Linux can be obtained from www.linmodems.org/. You may be able to obtain a driver for your modem from this site, even if it is not supported by SUSE. You can also download a tool called `scanModem`, which detects the exact type of modem that you have. When you know this, you can search the site for information about whether it is possible to get it working.

ISDN and ADSL connections

SUSE has very good support for internal Integrated Services Digital Network (ISDN) cards, and these can also be set up at this point in the installation. In most cases, the card will be automatically set up, and you just have to provide the specific information given to you by your ISP.

Asymmetric digital subscriber line (ADSL) providers sometimes offer a particular hardware device to connect with. These are sometimes USB devices. Unfortunately, there are a large number of different types and not all of them work with Linux. There are also different standards in different countries, and as a result, getting these devices to work on Linux has always been something of a problem. If YaST detects such a device during the installation, it attempts to set it up, but there are still many cases in which USB ADSL devices fail to work with SUSE Linux.

If at all possible, rather than using a USB device for ADSL, you should choose one of the ADSL routers with Ethernet output. This type of ADSL connection will always work, and if your provider offers the choice, you should definitely go for this type of connection. If your provider offers a "wires-only" service, you can buy such a router and use it to connect; again, there should be no problems at all. All you need to do is follow the instructions provided by the manufacturer or ISP for setting up your network connection to talk to the router and make the necessary settings in YaST's networking module.

Adding a new user

Just as Windows provides the infrastructure to authenticate users via a central database, the Unix world can use the Network Information System (NIS) or Lightweight Directory Access Protocol (LDAP) to store user account details. Most home users need to configure only a standalone machine and should select that option (see Figure 1-27).



See Chapter 24 for more information on the configuration of LDAP.



Figure 1-27: User configuration

Click Next to create a new local user, and the Add a New Local User screen appears (see Figure 1-28). Most of the information needed for creating a new user is self-explanatory.



One nice feature of the SUSE user creation process is that you can set yourself as the user that receives any mail destined for root by selecting the Receive System Mail option. Whether you set this up or not, it is always a good idea to read the root user mail (if you are the owner of the root user account!) to see any automated mails that the system sends as well as mails from the mail subsystem. This includes bounced emails, system errors, and package updates that have been installed.

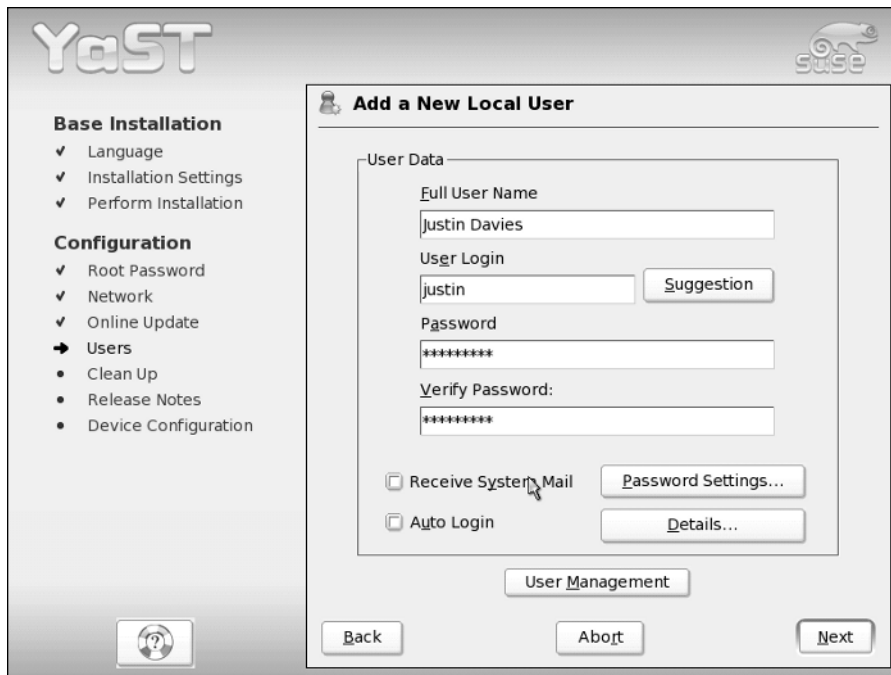


Figure 1-28: Creating a local user

Taking another page from Mac OS X and Windows XP, SUSE enables you to set up an account to automatically log into the system for you when the machine boots up. For home users, this provides a simpler way to use their system, but it is impractical and insecure in business environments. For example, if you are the user who is automatically logged in on boot up, someone else can gain access to your files simply by turning the machine on.

When you have finished adding your user information, click Next to continue, and SuSEconfig will run.

SuSEconfig

The SUSE system configuration is controlled by SUSE-specific control files that the system application SuSEconfig uses to write application-specific configuration files. This allows the user to configure services and server processes without having to understand application-specific configuration files. When all packages have been installed, SuSEconfig picks up the default configuration files installed by SUSE and writes out specific application configurations.

SuSEconfig is a core element of the SUSE system and allows YaST to maintain configuration files for services it can control. Any time you make changes to a service via YaST, SuSEconfig will be called to commit those changes.

Reviewing the release notes

When SuSEconfig has finished its initial installation, you will be shown the SUSE release notes (see Figure 1-29). These notes contain general information about changes from previous versions of SUSE Linux, as well as a technical overview of the previous version. This file also provides errata from the SUSE manual and is worth a read to get a general idea as to what has happened since the last release. Ninety percent of users at this point have not touched the included manuals with SUSE, except perhaps to move them out of the way when locating the installation media, so this at least gives you an introduction to the features and functionality of the product.

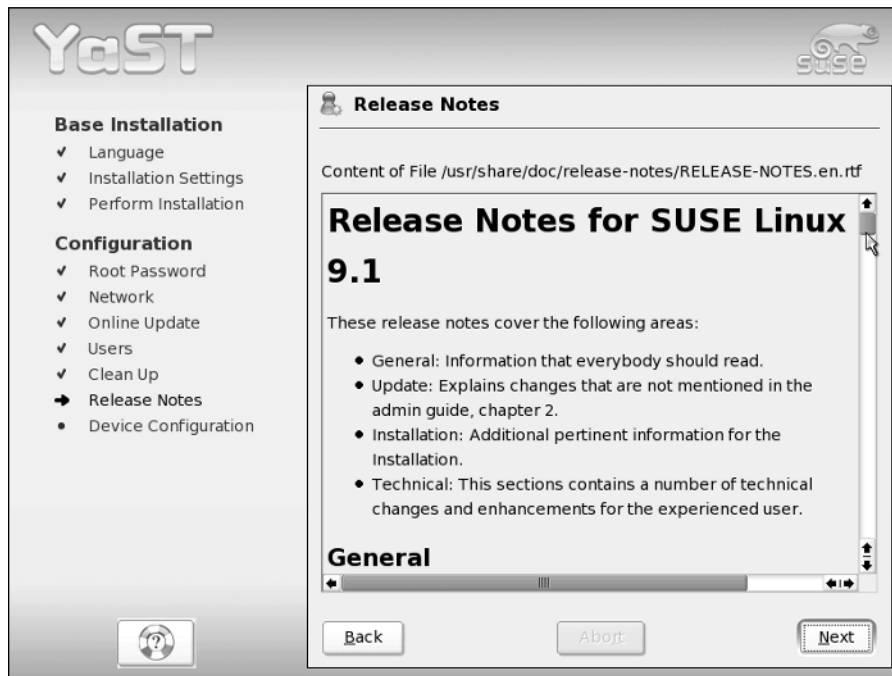


Figure 1-29: The release notes

Configuring your hardware

When you have read the release notes, click Next and you will be asked to configure your hardware (see Figure 1-30). The YaST installer and the YaST system configuration manager will run the same modules to configure hardware. For now you will configure the video card so that you can use X/KDE/GNOME.

Configuring your monitor

Select Change ⇄ Graphics Cards, and the Sax2 tool will be loaded to configure all aspects of X Windows, including monitor, graphics card, mouse, and keyboard. The Sax2 tool is a utility written by SUSE to simplify configuring the X Window system, which is otherwise a complex task. Sax2 makes it easy to configure the X Window system and related hardware such as your mouse and graphics card.

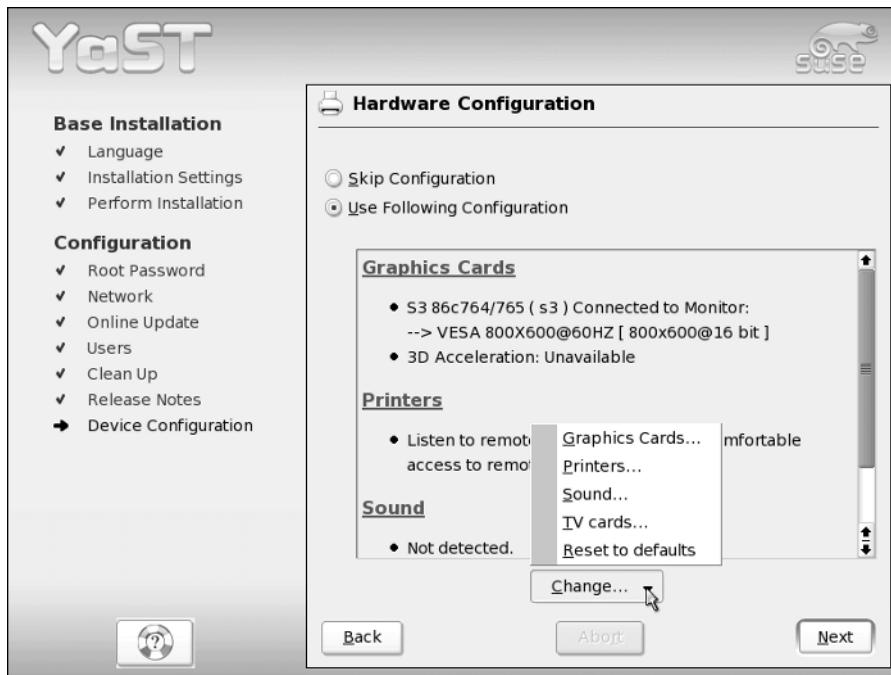


Figure 1-30: Hardware configuration

Sax2 is sometimes cautious in its estimation of the resolution you can use for X Windows and some manual tweaking may be needed to get the most out of your graphics card. The first thing to change or double-check is the monitor resolution, as the highest physical resolution the graphics card can display is limited by what resolution your monitor can handle. Click the Monitor selection on the left and either edit the configuration by clicking Change configuration or feel safe in the knowledge that the resolution is high enough for you (see Figure 1-31).

Every Linux book and piece of documentation on X Windows configuration has a disclaimer about configuring your graphics system. This book is no different because there are real dangers if you set up your monitor incorrectly. Because the graphics card drives the monitor, it is imperative that you either configure the graphics system with standard lower settings, or take a look in the documentation that came with *both* your monitor and your graphics card and figure out the correct settings for them. Sax2 comes with well-defined Video Electronics Standards Association (VESA) settings for general setup and also specific configurations from the major manufacturers of graphics systems. The remainder of this section discusses a low-specification graphics setup that should be safe for most people. However, you really should know how hard you can push your system so as not to damage your monitor by overdriving what your graphics card gives to it. Most of today's monitors have built-in settings to protect against hardware damage, but you should be especially careful when configuring the X Window system on an older monitor.

When you choose to change the configuration of your monitor, you will be presented with its current configuration. Click Properties to change the monitor configuration (see Figure 1-32).

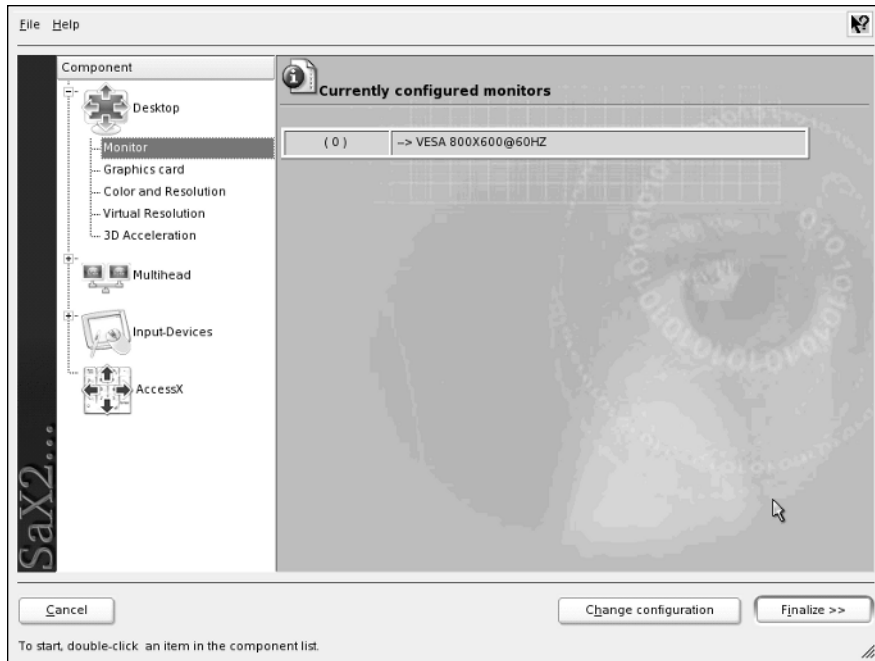


Figure 1-31: Current monitor configuration

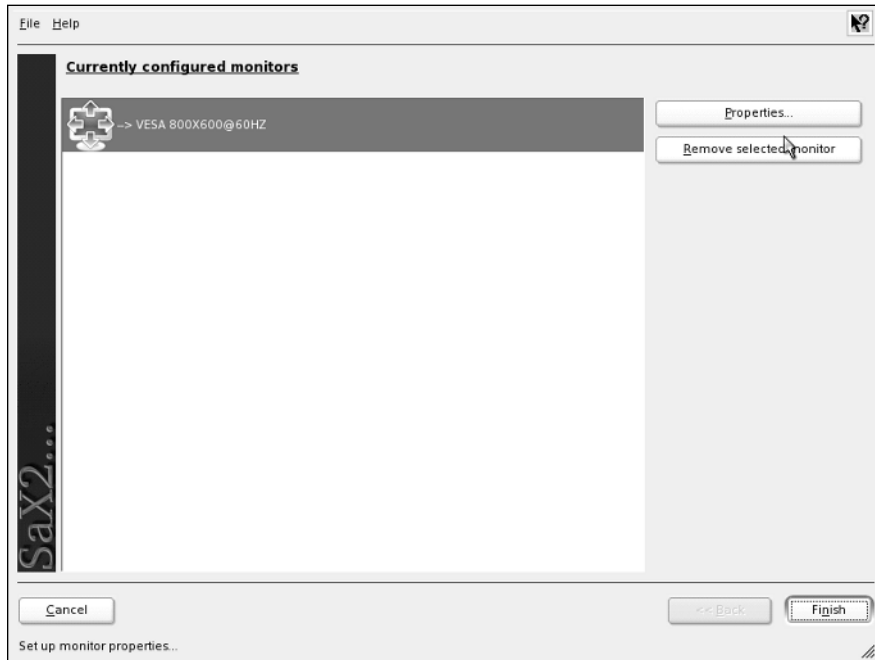


Figure 1-32: Choosing to change current monitor properties

Figure 1-32 shows a general configuration setting made by selecting a VESA standard resolution (see Figure 1-33). Most modern monitors adhere in one way or another to VESA standards. These are standard frequency and graphics resolutions that will work on most graphics cards and monitors. As noted previously, you should use the settings specific to your monitor, or select the specific manufacturer from the list on the left and the specific card you use from the list on the right to guarantee that your graphics will work in your hardware configuration.



Figure 1-33: Setting monitor options



Tip

If the settings detected for your monitor do not work, you may want to return to this screen and select Framebuffer as the specification for your monitor. A framebuffer provides an abstraction of your system's graphics hardware that the X Window system may be able to use if it cannot directly communicate with your graphics hardware.

Once set, click OK and then Finish to save your monitor settings.

Graphics card configuration

As you have changed the resolution of your monitor, Sax2 automatically changes the display resolution from your graphics card (see Figure 1-34). A piece of advice on color depth is that the higher the depth, the more processing needs to be done by the graphics cards. The higher end of the graphics card market can deal with this quite easily, but if you are using a graphics fallback such as framebuffer that uses the system CPU for its graphics routines, you will impact the speed of your system by increasing the bit depth and resolution. For 90 percent of people using Linux or any other operating system, a bit depth of 16-bit is enough. Eight-bit is generally too low for day-to-day use, and you will notice that the display of colors at this depth is noticeably worse than 16- or 24-bit. You can change these settings by selecting the appropriate items from the pull-down menus.

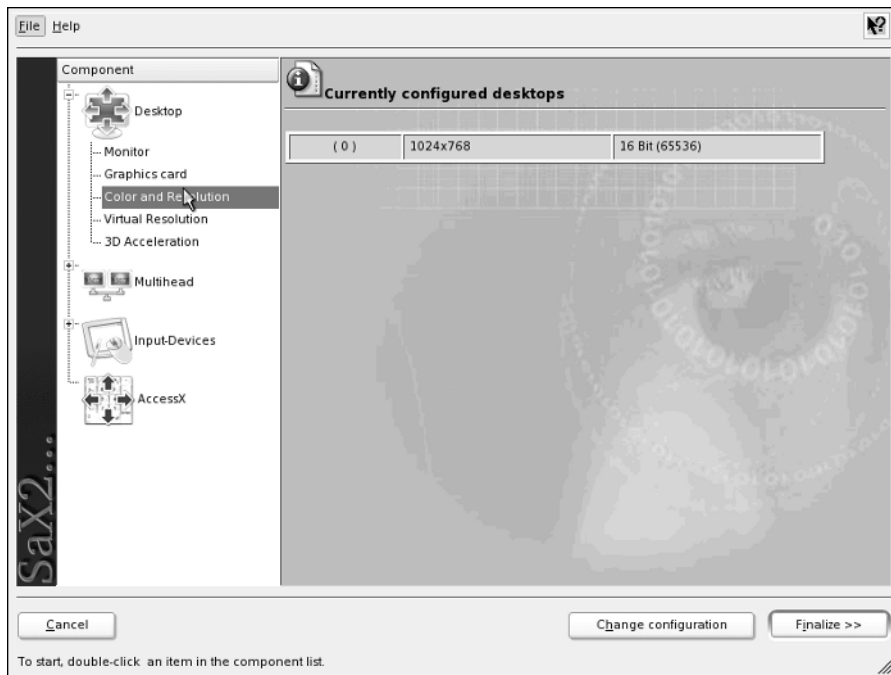


Figure 1-34: Graphics card configuration

Testing your graphics configuration

When you are happy with your graphics configuration, click Finalize to save your changes. You will be prompted to test your configuration before it is saved to the system configuration (see Figure 1-35). It is highly advisable that you select Test at this point because if you are booting directly into X Windows and your graphics do not work as you thought, it can be annoying and worrying to new users to see the system repeatedly try to load the X server when there is a misconfiguration.

Sax2 will try to load an X server with the configuration you gave it previously, and if things go well you should see a screen similar to that in Figure 1-36. The arrows provide a way to move the screen around the monitor (the same as you can do with your monitor using the buttons on the front). When you are happy with the configuration, click Save.

When the configuration has been saved, you will be told that it will not take effect until you restart your graphics system. As this is the installation process, your graphics system has not been started yet, and this warning is there for people who run Sax2 after installation where X Windows may be running at the same time as the Sax2 program.

Configuring your sound card

YaST will detect the sound card and will typically set it up completely automatically. During a standard installation, you are not required to intervene in this process; it just happens. In almost all cases, that is all you need to know. The rest of this section concerns what you can do after installation if it turns out that sound was not configured correctly during the installation.

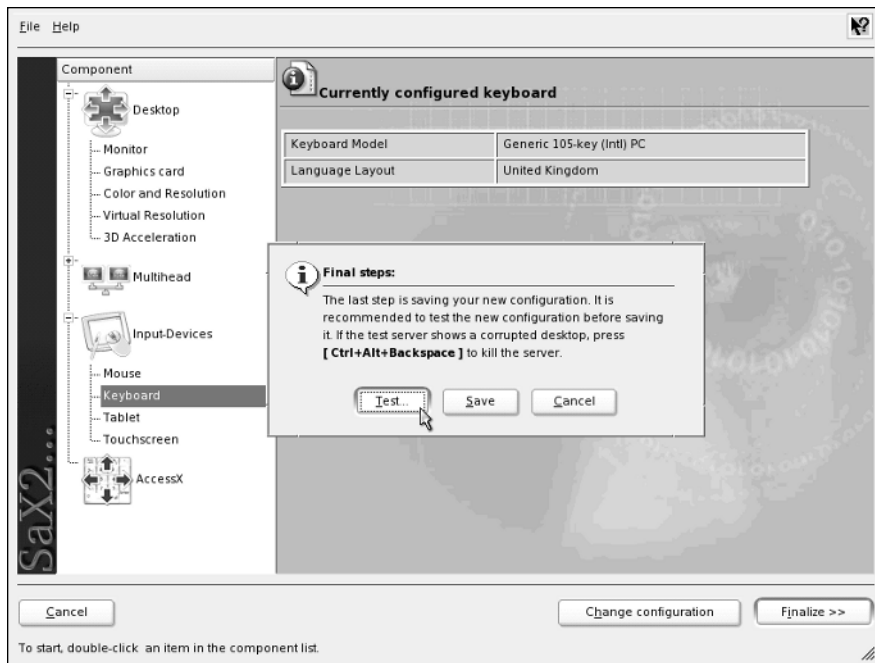


Figure 1-35: Choosing to test your graphics configuration

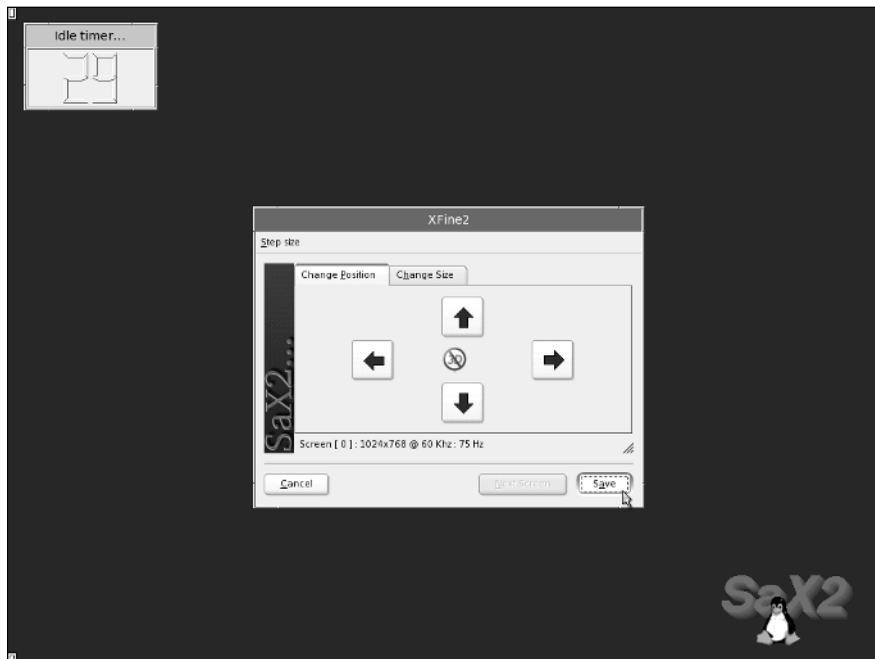


Figure 1-36: Testing your graphics configuration

After the installation is complete, confirm that sound is working by attempting to play a music CD or music file (for example, an MP3 file using the `xmms` player program). If you don't hear sound at this stage, first check the physical connection to the speakers. Then (if you are using KDE) check that the KDE volume control is at a sensible setting and not disabled or turned down to zero. In the unlikely event that sound still fails to work, you can rerun the YaST sound module in expert mode. The sound module is found in YaST's Hardware section. You will find three possible setup modes:

- ♦ **Quick automatic setup**— This is the default and is the one that is used during installation.
- ♦ **Normal setup**— This allows you to test the sound card. There is a volume control and a test button. When you have set the volume, a test sound is played when you press test.
- ♦ **More detailed installation of sound cards**— If you choose this option, you will be taken to a screen where any configurable options for the particular sound card that has been detected can be set. Depending on the particular card, these may include settings to enable a joystick and MPU (midi processing unit) port settings.

If even experimentation with the detailed installation options fails, you can try the low-level `alsaconf` program. As root, type `alsaconf` to start the program. It is a text-based program that, in fact, provides the back end for YaST's sound configuration module, but running it standalone gives you the opportunity to use its command-line options, including `alsaconf -l`, which writes a log file to `/tmp/alsaconf.log` that may give a clue as to the problem.



Tip

The ALSA (Advanced Linux Sound Architecture) home page can be found at www.alsa-project.org/. This is the best place to start if you have any difficulties with configuration of sound on Linux.

Completing Installation

It has been a long road, but you have successfully installed SUSE at this point (see Figure 1-37). Pat yourself on the back if you are a brand-new user to the world of Linux. You have done a lot of new things by installing SUSE— not the least of which is that you have begun a journey on which you'll learn lots of new ideas and philosophies, as well as giving you a stable operating system to use.

If you are an experienced Linux user, you should be quietly happy about how much SUSE has come along from other distributions and how easy it has been to install it. Don't worry, however; as with everything Linux-related, you can make it as easy or as hard as you like, and you will see how in later chapters.

Just to whet your appetite, Figure 1-38 shows an image of what you will see once your system has booted up to the system proper.

As this is a new installation with default settings, typing in your username and password and pressing the login button automatically loads the K Desktop Environment. The version of KDE that SUSE ships with has been optimized to integrate with the SUSE system, and you will see how well a job the developers have done in Chapter 8. Enjoy, play around with the system, and as the developers in Germany say: Have a lot of fun!





Figure 1-37: Installation completed



Figure 1-38: The final installed system

