

## Part 2 Installing Linux

Jarrold Spiga explains the Fedora installation, including printer and Internet connection configuration, and system updates.



### Skill level

Beginner

### Requirements

PC with DVD drive  
• Internet connection  
• printer (optional).

### Time to complete

4 hours

### Bonus DVD software

Fedora Core 3

## Making Fedora Core 3 CDs

This month's cover DVD contains the Fedora Core 3 CD ISO images. To install Fedora Core 3, you'll need to use a CD writing application (such as Nero Burning ROM or Adaptec's Easy CD Creator) to write the images to four CDs. The following steps should work for most burning apps:

- Copy the first ISO image to your local hard disk.
- Open the CD/DVD burning application and cancel any wizard screens that pop up.
- Select File > Open, and select the first ISO image. A write media window should appear.
- Insert blank media into your drive and hit the button that starts the burn.
- Repeat the process for the remaining ISO images. Alternatively, you can download a Fedora Core 3 DVD ISO images from the Fedora Web site (<http://fedora.redhat.com>).

Last month, we explained in detail how to prepare yourself and your system for installing Linux. This month, we've included Fedora Core 3 on the cover DVD (as a set of ISO images) and this guide will lead you through installing it. We'll also show you how to configure peripherals and set your system up to download updates automatically.

### FINAL PREPARATION

While this was covered in last month's guide, we'll reiterate it now. Ensure all your data is backed up before you begin. This guide is comprehensive, but it can't cover every possible scenario you could encounter when you're installing Fedora Core 3.

Once your system is backed up and you've completed all the preparation steps detailed in the last issue, insert the first Fedora installation CD into your about-to-be-Linux machine's CD/DVD drive and reboot your computer.

Your system should boot to the CD and a text-based welcome screen will appear.

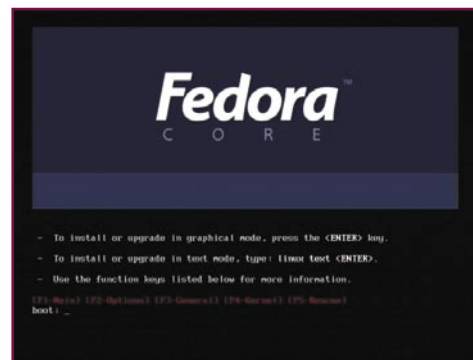
If your system doesn't boot to your optical drive, restart it and enter the BIOS. Check the boot device priorities, ensuring the system will boot to your CD/DVD drive first. If you have done this and find that you still can't boot to the CD, try creating a boot floppy disk and boot to that (refer to the box "Boot images gone floppy" for instructions on how to do this).

### STARTING THE INSTALLER

1 Once the text-based welcome screen has loaded, hit Enter to start the graphical installation procedure. Your system will then load a compact



2 While testing your media is an optional step, you may want to do it if you're unsure of its quality.



1 Booting to the Fedora Core 3 installer results in this text screen. Hit Enter to start using a graphical installer.


version of the Linux kernel into memory and ask you whether you want to test your installation media.

2 In the first part of this guide, we told you to check the media as a preparatory step. In some cases, however (such as if the media is scratched or if you've freshly downloaded and burned ISO images), it pays to run another scan — there's nothing worse than getting 90% of the way through an installation only to find that you can't read from your CD.

But if you're short on time or you're confident that your media is good, it's safe to skip this step.

3 Next up, Anaconda (the Fedora installer) will load into memory. The first thing that it does is some initial "probing" — Linux-speak for hardware detection. This will generate an initial X Windows configuration for your system (X Windows is the GUI for Linux). In a nutshell, it attempts to detect details about your mouse, keyboard, monitor and graphics card in order to display the GUI and accept input during the installation process.

4 After Anaconda has finished probing, X



3 The Linux installer probes your system for the settings that it needs to display the graphical installer.

## “How should I know if it works? That’s what beta testers are for. I only coded it.” – Linus Torvalds

► Windows should load and the Fedora Core 3 Install Welcome screen should display. If your system doesn’t show this screen, chances are that Anaconda encountered a problem when detecting your hardware. Try rebooting and repeat the installation process from scratch, or use the text-based installer.

To load the text-based installer, reboot and when the text-based welcome screen appears (see Image 1), type:

```
linux text
```

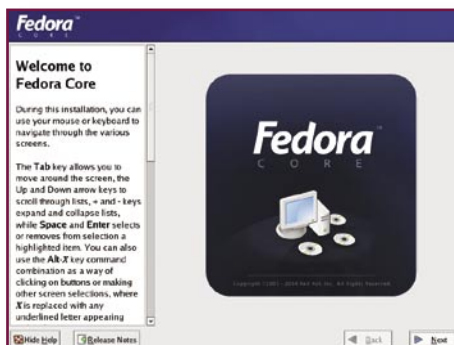
and hit Enter. There is very little difference between the graphical and text-based installers. Both achieve identical installations, but the graphical installer is a lot easier to use.

### IT’S GONE ALL GUI!

From the graphical welcome screen, you can see how the installer’s interface is laid out. The frame on the left contains help and useful information about each step in the installation. The main frame contains the information you’ll be working with at each step. The bottom frame contains Next and Back buttons, similar to those you’ll find in any Windows installer. Lastly, a button allowing you to read the release notes is always visible.

It’s a good idea to have a quick skim through the release notes (located in the root directory of the first installation CD/DVD in HTML format). These notes may give you advance warning of any problems you may have when running Fedora Core. For instance, Fedora Core 2 shipped with a known problem that prevented the kernel from loading into memory on systems featuring an ASUS P4P800 motherboard. There was no workaround for this at the time of release.

5 After you hit the Next button, you’ll get to



4 If the probe was successful, this welcome screen will appear. If not, reboot and use the text-mode installer.

the installation language selection screen. Most readers will probably want to select English (English) here. It’s important to note that this is NOT the language used when Fedora is installed — this setting will only apply for the duration of the install.

6 Next, select the layout of the keyboard that is currently plugged into your system. Keyboards sold in Australia almost always use the US English layout, although the option to use Dvorak or Asian keyboards is available. Unlike the previous language selection setting, this one applies to your system once Fedora Core 3 is installed.

### PROBING PROBLEMS

Anaconda generally does a good job of probing for hardware, but it’s not perfect. One of the areas where it sometimes fails is in detecting specifics about your mouse. It can usually figure out which port it’s connected to, but it often has trouble with the specifics. Does your mouse have a scroll wheel? How many buttons does it have?

In some cases, you’ll be asked to select an appropriate mouse. If the mouse cursor can be moved and clicked successfully, don’t change the setting at this stage — selecting an incompatible mouse type will prevent your mouse working, making it trickier to complete the installation.

Another area where Anaconda can occasionally trip up is in detecting your monitor. As with the mouse selection screen, a monitor configuration screen may or may not appear depending on your configuration.

If you see this screen, Anaconda was obviously able to figure out that your monitor was capable of displaying 800 x 600 at 60Hz. However, you may not be able to use higher resolutions/refresh rates unless Linux is aware of what your monitor can handle. ►



5 The first language selection screen applies only to the Fedora installation process.



## Boot images gone floppy

The Fedora Core 3 boot floppy images can be transferred to a single floppy disk using the rawrite utility, included on the cover DVD. First, grab a high-density (1.44MB) 3.5in floppy disk and back up all of the data on it (assuming there’s data on there that you want to keep). Bring up a command prompt by going to Start > Run > cmd (or command under Windows 95/98/Me). At the command prompt, enter:

```
format a:
```

where a: is the drive containing your floppy disk. Confirm that you want to format the disk and leave the volume name blank. With the Fedora Core 3 DVD in your drive, execute the following commands (assuming that D: is your DVD drive):

```
c:\> d:
```

```
d:\> cd /dosutils
```

```
d:\dosutils> rawrite
```

```
Enter disk image source file name: ..\images\bootdisk.img
```

```
Enter target diskette drive: a:
```

```
Please insert a formatted diskette into drive A: and
```

```
press --ENTER-- : [Enter]
```

When the rawrite command has finished copying the image to the floppy disk, you should then be able to boot to it and commence your Fedora Core 3 installation.

## Manage disks logically

The Logical Volume Manager (LVM) is another advanced disk management scheme supported by Linux. In a nutshell, it allows you to pool the storage space located on many drives and treats them as one large volume (similar to a striped RAID array, but with more flexibility). You can then break up this one large logical drive into the smaller partitions that Linux uses.

Using LVM can be risky since it's much more susceptible to disk failure than standalone drive configurations. However, when used in conjunction with growing partitions and hardware RAID controllers, it's perfect for high availability and high-performance disk subsystems. As such, LVM is generally only used in very high-end systems and servers.

There's one restriction with LVM: the partition that contains your boot information (/boot, or the root partition if you have not created a /boot partition) can't be created within a logical volume, it must reside on a physical disk.

The most important settings are the horizontal and vertical sync figures. The vertical sync figure is also known as the refresh rate. Every monitor has a range of frequencies that it supports at each resolution — exceed these frequencies and your monitor won't display the image as intended.

If your monitor is listed, select it and continue with the install. If you can't find your monitor in the list, leave the detected setting as your monitor type. When Linux is up and running, you can then try to identify the horizontal and vertical frequencies for your monitor, allowing you to use higher refresh rates and resolutions.

### CHOOSE YOUR OWN ADVENTURE

The next screen is the "choose your own adventure" option. Fedora's developers have grouped together four unique sets of installation packages for different purposes:

**Personal Desktop.** If you're looking for a quick and easy installation for a general purpose desktop machine, the personal desktop option will install packages that suit tasks related to personal computing.

Included are common word processor, spreadsheet, email and Web browser applications, a sprinkling of games and X Windows. However, some Linux power tools are missing.

### Option summary

- ✓ Great for general purpose use; smallest installation footprint.
- ✗ Isn't ideal for everyone, especially power users.

**Workstation.** Emulating a traditional Unix system, Workstation installs a more advanced set of packages tailored to application development, programming and scripting.

Included are development environments, support for numerous Unix-oriented programming languages, compilers, debuggers and toolkits, along with the mandatory office apps.

### Option summary

- ✓ A solid, yet flexible Linux install.
- ✗ A waste of disk space for beginners; strictly professional.

**Server.** If you're looking to Linux for running Internet-based servers (such as a Web server, FTP server, email servers, and so on), the Server option may be the one for you.

When running a server it's best to minimise the installation footprint to keep as many resources as possible available to the server apps, so hardly anything else but the servers will be installed. However, you will have the opportunity to install X Windows later in the installation.

### Option summary

- ✓ System optimised for server needs; small installation footprint.
- ✗ Not for general purpose use.

**Custom.** If none of those quite fit your needs, or if you want to be able to pick and choose your own packages, the Custom install is for you.

If you select this option, you'll get to browse through the large package library a little later in the installation.

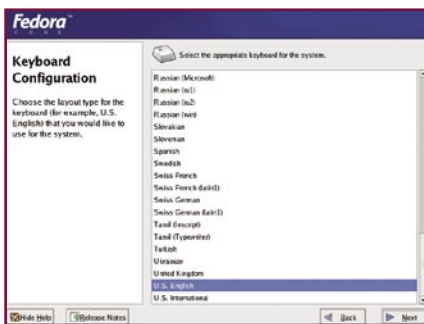
### Option summary

- ✓ Completely customisable.
- ✗ Package selection can be confusing and difficult.

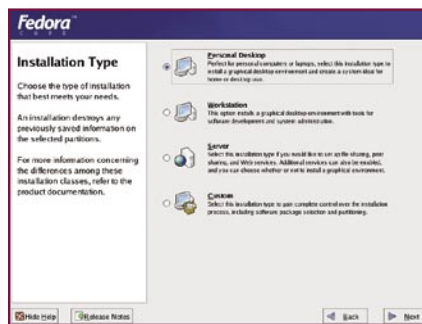
After selecting which package you'd like to install, hit Next to proceed.

### THE DISKS IN MY COMPUTER

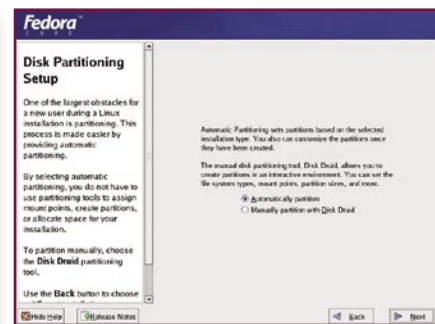
The next step of the installation is arguably the most difficult — partitioning your hard drive for Linux gives you a crash course in the Linux file system. The part that confuses most people boils down to the different file system structures



**6** Select the appropriate keyboard layout for your system — this is usually US English.



**7** Choose the type of installation you're performing. If you're unsure, pick Personal Desktop.



**8** Automatic partitioning is an easy option if you only have Linux on your system. However, most users will want to partition manually.

► under both Windows and Linux. Although they're designed differently, the process of implementing them is quite similar.

The Windows (and MS-DOS-based) file system is designed so that the operating system treats each physical disk (and each partition) as a separate entity. A drive letter is assigned to each disk or partition. Network users may even have drive letters assigned to UNC paths, but these are still treated as separate entities.

To an extent, this changed with the release of Windows 95. We were introduced to a new parent of all drives named "My Computer". My Computer under Windows is conceptually equivalent to a Linux root (/) file system.

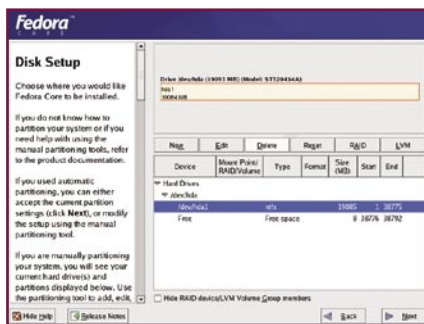
Under My Computer, an icon for every drive on your system appears. In Linux terms, you could say that every drive is mounted under My Computer, along with a couple of additional items, including Control Panel.

In Linux, everything (except for the swap partition) is accessible from the root file system. In last month's guide, we recommended that you set up at least three Linux partitions — one for the root file system, another for your swap data and the third as a small boot partition. Linux will mount the boot partition under the root file system at /boot to make it accessible.

We also mentioned that advanced users may reserve space for additional partitions, such as a /home partition — in other words, a partition mounted under the root file system at /home. If you're new to Linux, this is an additional layer of complexity you needn't consider now — keep things simple and use one massive root partition as well as a small /boot partition.

This doesn't just apply to hard disk partitions — floppy disks, CDs and DVDs are all mounted in a similar fashion (usually under /mnt/floppy or /mnt/cdrom).

Under Windows XP and 2003, you can also mount other drives under another path on your system (provided you're using an NTFS file system on your PC — this feature is incompatible with FAT32 or VFAT volumes).



**9** Our real-world installation began with a Windows partition that consumed the whole hard disk.

## DON'T CUT CORNERS

**8** The next choice that you have to make in the installer is whether you want to automatically partition your system, or if you want to do it manually using a tool called Disk Druid.

Before you decide, consider this: automatic partitioning sounds much easier and simpler, but there's a catch. Automatic partitioning can be unpredictable. For instance, if you've reserved 100MB at the beginning of your drive for a boot partition, there's no guarantee that the automatic partitioning process will use the space for this purpose.

The automatic partitioning process won't allow you to configure your hard disk in a software RAID configuration, either.

If Linux is going to be the only operating system on this PC, using automatic partitioning will get the job done quickly and without fuss.

## AUTOMATIC FOR THE PEOPLE

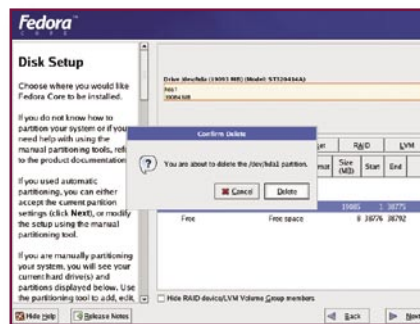
If you do let the installer partition your system for you, you'll still have to tell it whether it should fiddle with existing partitions on your hard disk, and which hard disk you want to install Linux on.

You have three disk partitioning behaviours to choose from:

- Remove all Linux partitions on this system. You should select this option if you have previously attempted to install Linux (Fedora or otherwise). If you select this option, the installer will completely remove all Linux partitions from the hard disk and recreate and reformat them before installing the new Fedora Core. The installer won't touch any other partitions on your hard disk, including FAT32 (VFAT) and NTFS partitions.

- Remove all partitions on this system. Don't choose this option if you're planning to dual-boot to an existing Windows installation. This behaviour will remove all of the partitions on your hard drive and replace them with the Linux partition scheme that matches your installation type best.

- Keep all partitions and use existing free space. The best option for dual-booters who have



**10** Windows is removed to make way for Linux. Be aware that all of the partition's data will be purged.

## Growing partitions

Most home users won't have to worry about growing partitions, but the concept is a good one to understand for when you become more proficient with Linux.

Consider the case where you have a file server with limited disk capacity. Create a partition for all of your users' data (/home) that is separate from your root file system. This way, if your users fill up their home directories, Linux won't crash because there's still space available for Linux to write other data to disk. You know that you need to reserve a certain amount of space for this partition now, but you're not sure of how much space you'll need in six months' time.

In this case, create the partition so that it takes up the amount of space you think will be required in the long term, but select the option to "Fill all space up to (MB)" and enter the current capacity requirements in the associated text box. The partition that will be created will meet your size requirements now, and space will be reserved on the hard disk for it to "grow" at a later date.

Let's say that in six months' time, the user data partition doesn't need to be increased, but the root file system is running a little low on space. You can then reduce the amount of space reserved for the growth of your user data, and then grow your root directory a little.

This gives administrators extra flexibility. Once you assign disk space to a partition, you can't deallocate it. By using this method, you can reserve some space on your hard disk so you can grow partitions in the future.

However, most Linux home users won't need to worry about this.



## Custom boots

**LILO (the Linux Loader) is another boot loader and is installed as the default boot loader in place of GRUB by a number of other distros. Accomplished users of Linux may be more familiar with LILO and may want to use this boot loader instead.**

**Other users may already have a familiar and flexible boot loader installed on their systems. The Fedora installer also gives you the option to install no additional boot loader. If you choose this option, you'll need to configure your existing boot loader to load Linux properly or you won't be able to access your new OS.**

**The Fedora installation routing allows such users to change the boot loader — upon pressing the relevant button, a window will pop up asking you which boot loader you'd like to use: GRUB, LILO or none.**

**With any luck, the installer will correctly detect all bootable operating systems on your PC. If not, add other entries to the list by using the Add button. Likewise, if you want to make sure a particular operating system is not accessible in the boot loader, select that OS and hit the Delete button (this won't delete the whole operating system — just its entry in GRUB).**

**You can set GRUB to load a default operating system after a period of time (10 seconds, by default). The operating system with a tick under the Default column will be the one that GRUB loads.**

**You can also set a boot loader password if you're worried about anyone gaining unauthorised access to your system.**

► made appropriate preparations, this option will instruct the installer to create new partitions for Linux in the existing unallocated space on your hard disk.

Before the installer can proceed, you'll also have to tell it which hard disk you want to install Fedora Core on to. The information box should provide you with the disk capacity, manufacturer and model, which should allow you to be sure of which hard disk is about to be changed.

While you're at this spot, look at the device name — it should be called "hda" or something similar. Linux identifies hard disks by both the interface that the hard disk is plugged in to, as well as the channel the drive is connected on. The hard disk attached to the primary master IDE controller is always assigned the name "hda". The secondary master is "hdb"; primary slave is "hdc", while the secondary slave is "hdd". A similar convention is used for SCSI and native SATA controllers — "sca" is assigned to the first drive, "scb" the second and so on.

Towards the bottom of this screen, you have the option to "review (and modify if needed) the partitions created". Keep this option selected.

### MANUAL DELIVERS THE POWER BETTER

If you decide to manually partition your system, the Disk Druid graphical interface will appear. The screen layout is similar to the Disk Management MMC snap-in under Computer Management in Windows — a disk partitions diagram shows (logically) where each partition resides on the hard disk, while a pane below this gives you the details of each partition.

In between the two, a row of buttons allows you to create, destroy and modify partitions. There are also RAID and LVM (Logical Volume Manager) buttons, but we'll come to these options shortly.

► To demonstrate manual partitioning as clearly as possible, we'll step through a real world example. A Pentium III 1GHz machine was sitting in a corner of the APC Labs. After cleaning the dust off the case and plugging it in, we proceeded

with the installation up until this point. As you can see in the image, this system had a single NTFS partition (containing a Windows 2000 installation) occupying the whole hard disk.

10 The intention was to make this a dedicated Linux PC. As such, the first step was to eliminate the existing NTFS partition and free up space on the drive, which we can then use to create Linux partitions. If you plan to use a dual-boot configuration, you'll probably want to keep the NTFS or FAT32 (VFAT) partitions that Windows is installed on to preserve the data on them — Linux partitions will be created in the gaps that were made during the preparation stages (see last month's guide).

Once you have enough free space on the disk, it's time to create some partitions. The New button will allow you to do this. If you're thinking about creating RAID volumes, take a look at the RAID information box.

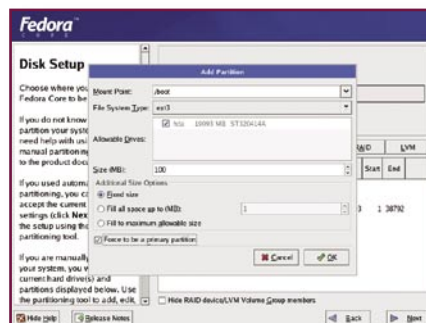
The /boot partition should be created as close to the beginning of the hard disk as possible to work around the 1,024 cylinder restriction that some BIOSes have. To ensure that this occurs, create the /boot partition first.

### A BIT ABOUT EXT3

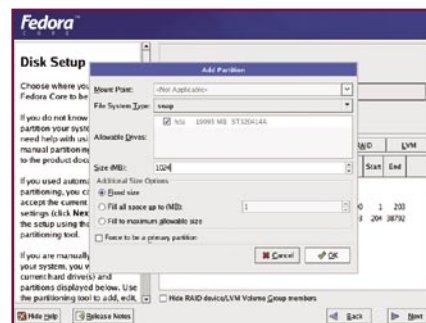
11 Change the mount point to /boot, and use the ext3 file system type. Until recently, Linux used the ext2 file system with indexes and blocks in a similar (but different) manner to NTFS. While this file system has performed admirably for Linux over the years, one of its limitations was that as volumes got larger, they became more difficult to manage if the system crashed.

After a system crash, a series of tests would be run on the file system to try to fix any problems and ensure it wasn't corrupt. On occasion, an irreparable corruption would occur, preventing you from being able to mount that file system. As volumes got larger, the probability of these problems occurring increased.

Essentially, the ext3 file system is identical to ext2, but it also keeps a journal of all of the changes made to a file system — that way if



11 Create the /boot partition at the beginning of the drive. This will avoid any BIOS limitations that your system may have.



12 Next, create your swap partition. Note that this partition doesn't get mounted on the file system.

## A RAID of light

In last month's guide, we mentioned that it is possible to set up software RAID configurations through the Linux installer. While this can offer a flexible and fault-tolerant solution to your Linux installation, the performance of software RAID configurations is not as good as hardware-based solutions. Your RAID and hard disk configuration can also have a significant impact on performance and fault tolerance.

One of the main advantages to software RAID configurations is that you can get access to the upper RAID levels, including RAID 5. You can also set up RAID configuration similar to that described in the Matrix RAID Workshop article (APC December 2004, page 100).

The first step that you need to do on your RAID journey is to reserve partitions on your hard disk(s) for RAID devices (and yes, it is possible to set up RAID arrays using a single disk). Clicking on the New button within Disk Druid to create a new partition can do this. This partition will be reserved for use in a RAID device.

Once the "Create new partition" window pops up, select "Software RAID" as the partition type, indicate which hard disk to create the partition on and indicate how much hard disk space you want to reserve for this RAID partition. You won't be able to select a mount point at this stage — it gets selected when the RAID device is created. There's no need for RAID partitions to be set as primary partitions. Also, ensure that your partition is a fixed size. When you hit OK, a new partition will be displayed in a similar fashion to any other partition.

Repeat the process for any other RAID partitions you want to set up — you'll need at least two, preferably on different disks for better performance and fault tolerance.

Once you've created all of your RAID partitions, it's time to combine them into RAID devices. Click on the RAID button from Disk Druid's main screen. The next screen that pops up gives you a few options:

■ **Create a software partition.** This takes you through the process of creating

another partition on your hard disk to reserve for use in a RAID device. You should have already done this in the previous steps.

■ **Create a RAID device.** This option allows you to create a RAID device using two or more RAID partitions. Choose this method.

■ **Clone a drive to create a RAID device.** This option will allow you to select an existing drive on your system and automatically create new partitions on another drive which identically match that of the original drive. All partitions on both drives will then be mirrored.

From here, you can then select which RAID partitions you want in your new array, what type of RAID array you're creating, the mount point for the array and what file system type you want for the array (in most cases, ext3).

When you hit OK, you should be able to see the details of the new RAID array in Disk Druid's bottom frame. If you're creating more than one RAID array, simply repeat the process.

▶ the system crashed and the file system was in an inconsistent state, Linux would be able to recover by using the information in the journal to fix the file system. You may see ext3 referred to as a journalled file system for this reason. Journalling helps the file system recover from problems reliably.

### BACK TO THE TASK AT HAND

When creating any partition, you'll also need to select which drive you want the partition to reside on.

Just below this option, you're asked to specify the size of the partition. Red Hat recommends the /boot partition is at least 100MB. The additional size options can give you additional flexibility on your hard disk.

Fixed size partitions are the norm — where the operating system generally won't allow you to grow or shrink the partition. The other additional size options will be covered in a moment.

Force the /boot partition to be a primary partition. Without the aid of a boot floppy or CD/DVD, Linux can't boot if the partition containing the boot information is on a logical partition. Other Linux partitions needn't be set up as primary partitions.

Once you've entered these settings, hit OK to create the partition.

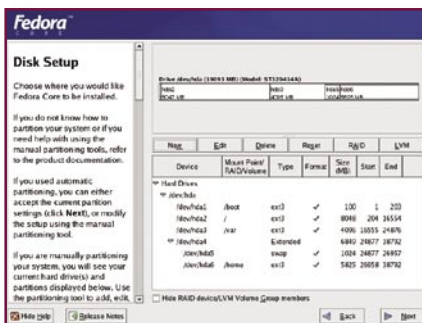
### SWAP SPACE

A Linux swap partition has to be created for the Linux virtual memory system to be able

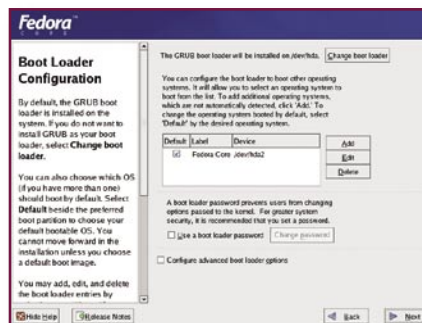
to manage applications when the system is running low on free physical memory. Start creating your swap partition by hitting the New button again.

12 Since swap space is not actually mounted on the root file system, you can disregard the Mount Point option here. However, the swap partition uses a different file system type than partitions that hold data. Select swap from the list of file system types. You'll also need to select which drive should hold the swap space if you have more than one hard disk.

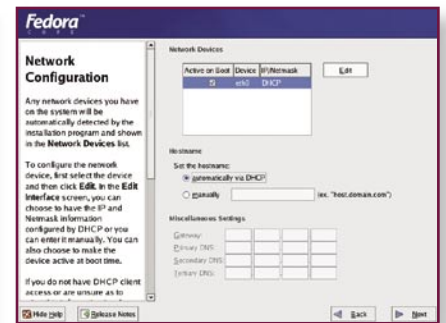
Setting the amount of hard disk to use for swap space is crucial to the performance of your Linux PC. If the swap space is too small, the virtual memory subsystem



13 Create your root partition along with any other partitions you want. Then double-check everything.



14 Verify the default boot loader configuration. The default options will suit most users.



15 While Fedora hopes the network configuration is supplied via DHCP, you can make changes.

▶ will have a hard time managing how much data to page in and out of memory. If the swap space is too large, a whole chunk of your hard disk is going to waste.

A good rule of thumb is to set the swap space so that it is twice the size of the amount of RAM you have installed. For instance, if you have 512MB of RAM, assign 1,024MB to your swap partition. If you have limited hard disk space, consider assigning less space.

However, the reverse doesn't apply — you won't necessarily get better performance by making your swap partition larger than double the amount of RAM that you have. Twice the amount is almost always perfect.

There's no need to force your swap partition to be a primary partition — as mentioned previously, only the partition containing your boot data needs to be on a primary partition.

## THE ROOT OF THE MATTER

All Linux installations must have a root (/) partition. You should create the root partition using a procedure similar to the one you used to create your /boot partition. Click New, select ext3 as the file system type, leave the mount point as "/" and select the hard disk that you want the partition to reside on.

If you don't plan to create any other partitions, select the "Fill to maximum allowable size" option in the Additional Size Options frame. This option will make the partition that you're creating take up all of the remaining contiguous space on the disk.

Again, there's no need to set the root partition as a primary partition - unless you haven't created a dedicated /boot partition.

## LOOKING THROUGH WINDOWS

If you have Windows partitions formatted in

FAT or FAT32 (VFAT), these can be accessed from within Linux.

In order to mount an existing VFAT partition under Linux, select the partition from within Disk Druid and click Edit. Select a mount point from the window that pops up (use a location that is easily identified, such as /mnt/win).

Be sure to leave the partition intact — formatting it will erase the data on the partition.

Leave any NTFS partitions on your system alone. If you don't assign a mount path to them, Linux won't be able to access them. NTFS support under Linux is still experimental and mounting these partitions is likely to corrupt them irreparably.

## CHECK ONE, TWO

13 At this point (even if you've let the installer set up your partitions automatically), it's worth double-checking that your system is partitioned properly. In particular, make sure that:

- you have a /boot partition and that it is at least 100MB;
- the /boot partition starts before the 1,024th cylinder on your hard disk (look at the start column);
- a / (root) partition has been created;
- a Linux swap partition has been created which is at least equal to and no more than double the amount of memory in your system;
- you haven't mounted any NTFS-formatted partitions in the Linux file system;
- you're only formatting the new Linux partitions that you're creating (including the swap partition).

Once you're happy with the way your partitions are set up, hit Next to proceed with the installation.

## BOOT LOADERS

14 The boot loader is a crucial component that instructs your PC how to load any operating system that is installed. The boot loader is most commonly loaded into the Master Boot Record (MBR) of the hard disk that is first booted to. During system startup, your PC's BIOS will arrange for this code to be executed.

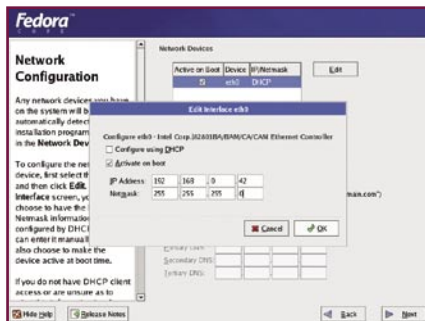
If you only have Windows installed on your system, you'll rarely see the Windows boot loader. By default, it's transparent to the boot process. With Linux (or if you're dual-booting), the boot loader will greet you every time you start up.

By default, the Fedora Core installer will install the "Grand Unified Bootloader" (GRUB) on the MBR and automatically configure it so that it can boot to all bootable partitions. This screen allows you to customise the way the boot loader runs on your system.

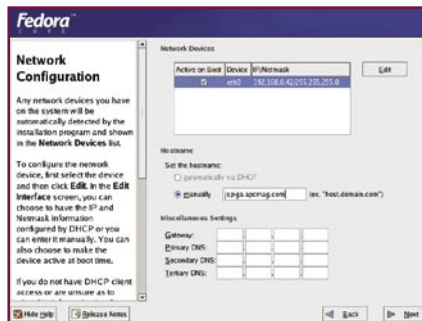
## INTERFACE CONFIGURATION

15 The next network configuration screen will really only apply to you if your system is connected to a network, or if you'll be hooking up a cable or DSL modem directly to your PC. By default, Fedora assumes that your network interfaces are connected to a broadcast domain that features a DHCP server and that this DHCP server will automatically configure your network settings. However, this is often not the case — your ISP may require a specific configuration, or you may want to use static settings if you plan on using your Fedora box as a server.

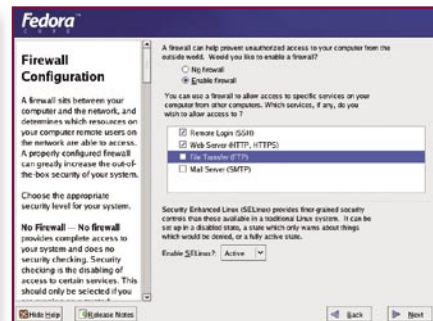
This screen is not used to configure dialup modem connections and most ADSL connections (in particular, ADSL connections that use point-to-point protocol (PPP) connections such as PPPoE or PPPoA to connect to the Internet). ▶



16 If you don't rely on DHCP, edit your Network Device settings and specify your IP address and subnet mask.



17 Most DHCP servers don't assign a hostname, so it's a good idea to manually set this yourself.



18 The installer can configure a simple firewall configuration in a few clicks — and it's more robust than the Windows Firewall.

► If you want your system to pull networking information directly from another DHCP server on your network, you can leave the network devices configuration alone and jump to the next heading.

**16** If you'd like to configure a static IP address to a network interface, select the interface in question and hit the Edit button. A small window pops up that will allow you to specify an IP address and subnet mask for the interface (once you disable the "Configure using DHCP option") and whether you want the interface to start up during the system boot process. You can also toggle this last feature from the main network configuration screen.

### BY ANY OTHER NAME

The system's hostname can be set automatically via DHCP, or you can specify one yourself. Most DHCP servers will generally assign IP information only, so you may want to set your hostname manually regardless of whether you use a DHCP server or not. If no DHCP server assigns a hostname to your system, Fedora will assume the localhost.localdomain name.

Some cable providers require that the system a cable modem is connected to has a specific hostname — if this is the case for your ISP, this is where you should define the hostname.

**17** To change the host name, select the Manual button and type in the hostname you wish to use. Fedora generally expects a fully qualified domain name (FQDN) for your hostname (as shown in the diagram). If you enter a single word, the installer will append ".localdomain" to your hostname.

### OTHER NETWORK SETTINGS

Depending on the options you've chosen, you might be able to customise the miscellaneous network settings. If

you're relying on DHCP to assign you a network configuration, they'll be greyed out — otherwise you should be able to set your gateway and DNS settings. What these details should be will depend on your network configuration. If you're not sure what they should be, it's usually safe to leave them blank. This is especially true if you're using a PPP connection to connect to the Internet — these details will usually be provided to your system when you go online.

Once you've completed the basic network configuration, hit Next to continue.

### THIS AIN'T NO WINDOWS FIREWALL!

Linux has always had a very strong and robust firewall solution. While it hasn't always been easy to configure, it's stronger and more feature-packed than any desktop firewall software package available. The IPTables firewall is also a part of the Linux kernel, as opposed to "bolting on" to the side of the networking stack of the operating system. A firewall that has been implemented in the kernel can protect the PC more adequately.

**18** The Fedora installer can create a simple default firewall configuration in a snap. Of course, if you'd prefer to disable the firewall, the option is there. But for most desktop systems, a simple firewall will improve security without reducing too much functionality — especially if you're connecting this system to the Internet.

If you plan to run servers on your Linux system, set up the firewall so that other network clients can communicate with these services. A small Web server and FTP server will be connected to the example PC. We'd also like to be able to bring up secure remote shells (an SSH session).

The supplied list of services is brief, to say the least. There are hundreds of

other services that can run on your Linux box, some of which you may want to make accessible to other machines. To add a definition for these ports to the firewall, add them into the Other Ports field using the <port number>:<protocol> convention, and separate each entry with a comma. For instance, if you installed a DNS and POP3 server and wanted to enable these services, enter the following in the Other Ports field:

```
53:tcp, 53:udp, 110:tcp, 110:udp
```

A DNS server will usually listen for TCP and UDP traffic on port 53, while a POP3 server usually listen on port 110 for TCP and UDP connections. A thorough list of standard ports can be found at [www.iana.org/assignments/port-numbers](http://www.iana.org/assignments/port-numbers). Alternatively, refer to the box "Well-known ports" for a list of the standard ports for more common applications.

### SPEAK MY LANGUAGE

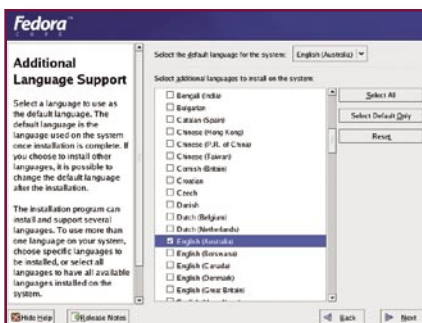
**19** You may have more than one language installed on your system, but the language selected as the default at the top of this screen will be displayed for all users until the language preference is changed within the Linux desktop.

### DO THE TIME WARP

In particular, a Linux system will perform maintenance tasks on an hourly, daily, weekly and monthly basis at close to 4:00 every morning. Because some of these maintenance tasks can be quite intensive, you wouldn't want them to run at times where you're likely to be using the system. Setting the appropriate time zone information will ensure that these tasks are run when most of us are fast asleep.

**20** Across the map of the world, you'll see a scattering of yellow dots. Each one corresponds to a city that Fedora has a time zone entry for. To set your time zone, select the yellow dot that is closest to your location.

The selected city should appear highlighted in the box below the map. Double-check which city has been selected, and ensure that location has the same time zone details as your current location. For instance, if you live in Grafton, New South Wales, select Australia/Sydney as your location because it shares the same time zone. Australia/Brisbane might be closer geographically, but Queensland does not adopt daylight savings time. ►



**19** After configuring your firewall, specify the languages you want to use within Fedora after the installation is complete.



**20** Setting your system time is important — it affects the time displayed on your system and the running of routinely performed scheduled tasks.



▶ The check box at the bottom of the screen allows you to specify whether the system clock (the time that is set from within the BIOS) is set to Universal Time Clock (UTC) — also referred to as Greenwich Mean Time (GMT). Generally, this option does not need to be set.

When you're done, click Next.

## ACCOUNTABILITY

Similar to the way root (/) is the core of the Linux file system, the root user is the account that has complete and total access to anything on the Linux system. The root user may run applications, change settings and modify files — even if they haven't been given the express rights to do so. With root access, a user can do pretty much anything they want on a Linux system.

**21** It's imperative to set a very secure password for this account on your system. It's just as important that you remember this password — forget it, and you won't be able to use your Linux system to anywhere near its full potential.

You must enter a root password — the installer won't allow you to use no password. In fact, an error message will pop up for passwords containing fewer than six characters, since shorter passwords are easier to guess than longer ones.

Once the root password is successfully set, Anaconda will read the package information from the installation CD/DVD. This process may take a couple of minutes, depending on your system.

## STANDARD PACKAGES

Earlier, you will have selected your Linux installation type — Personal Desktop, Workstation, Server or Custom. If you chose either of the latter two options, the Package Group Selection screen should be displayed next.

Otherwise, you'll see a brief summary of the packages for the installation type you selected.

This screen should also offer you two simple choices: Accept the Current Package List, and allow the installer to copy over all of the default packages for your installation type, or Customise the Set of Packages to be Installed.

If you're new to Linux, there's a good chance you won't know what many of the package names are, or whether you really want to install a certain package or not. In this case, it's probably best to accept the current package list — you can always install additional packages later.

## PACKAGE PICKING

**22** If you chose a Server or Custom installation, or decided to customise the Personal Desktop or Workstation collections, you'll be greeted with the Package Group Selection screen. From this screen, you can select which broad package groups should be installed. Simply check the box adjacent to each package group you'd like to install and the most popular packages belonging to that group will be installed.

If you'd prefer more control over exactly which packages are installed, click on the details link next to a package group. A window will appear, listing which optional packages may be installed. For some package groups, you won't be able to remove certain packages that are required for the installation of the package group as a whole.

As you add and remove packages, the total install size towards the bottom of the window will change. Be warned: if this figure is higher than the amount of space you have available under your root partition (assuming you haven't set up

several different partitions for various parts of your file system), you won't have enough hard disk space to install all the packages you want.

A detailed explanation of every package available is beyond the scope of this article. However, most packages have a brief description associated with them to give you an idea of what that package does. If you're not sure whether you need a certain package or not, don't install it at this stage — you can install these packages from the Linux desktop after the installation has completed.

To make the most of this Linux guide (and so you can follow future guides), ensure that you install the following package groups:

- The X Windows system. This is the Linux Graphical User Interface (GUI) and without it, you'll get stuck at a daunting command line interface.
- The GNOME Desktop Environment. You may install KDE as an alternative, but this guide will focus on the use of GNOME.
- The Graphical Internet Package. This installs a Web browser and email client for the X Windows system.

If you scroll down to the end of the list of package groups, two special groups appear — Minimal and Complete. As the names suggest, selecting the Minimal group will install only the required packages, while Complete will install everything.

Once you're happy with the set of packages that will be installed on your system, hit Next.

## THE ACTUAL INSTALLATION

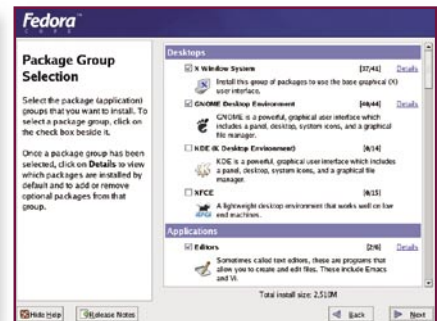
**23** The installer should now be ready to import Linux to your system. This is the point of no return. If you have any doubts about whether you want to install Linux, this is the point to restart



**21** Your root account can control everything on your system. Ensure that you set a secure password.



**22** Depending on your installation type, you may be given the option to customise the packages to be installed on your system.



**23** Select which package groups that you'd like installed on your system. You can have control over specific packages using the details links.

▶ your computer. If you're sure you want to go ahead with the installation, click Next at the confirmation screen.

**24** Next, the Fedora installer will commit the partition changes made in Disk Druid to your hard disk and reformat each partition with the appropriate file system. The actual time it takes to format each partition depends on how big it is as well as the speed of your drives. While the process typically takes a few minutes, a progress bar will show you how far into the operation the installer is.

After all partitions have been formatted, the installer will perform a few management tasks before copying the packages across to the hard disk.

**25** Once file copying has begun, a different progress meter will appear to indicate how much of the installation has completed. Depending on the speed of your PC and the packages you're installing, the copying process can take anywhere from a couple of minutes to an hour.

Once all the packages are installed, the installer will perform post-installation configuration to ensure each application's default configuration is applied properly. Again, a progress bar will indicate progress.

With any luck, all installation tasks are complete and you'll be asked to remove the installation media from your drive(s) and restart your system.

What happens when you reboot?

After you hit Reboot, you'll be returned to a text interface, and you'll probably see a whole lot of text scroll up rapidly. The installer is now going through the process of shutting down your system by stopping services, closing applications, unmounting the file systems, and sending the halt signal to your PC, causing it to reboot.



**24** Once you proceed past this screen, there's no turning back. Only move forward if you're sure you want to install Linux.

## Well-known ports

Many applications rely on connecting to a server on a standard, well-known port in order to make communication easier. However, if you're running a firewall between the client and the server, you'll need to allow the communication through the well-known port on the server. The following is a list of the most commonly used well-known ports (and protocols), and the application usually associated with each port.

**20:tcp, 20:udp**

**FTP** — File Transfer Protocol: Data Transfer (on active connections)

**21:tcp, 21:udp**

**FTP** — File Transfer Protocol: Control

**22:tcp, 22:udp**

**SSH** — Secure Shell

**23:tcp, 23:udp**

**Telnet**

**25:tcp, 25:udp**

**SMTP** — Simple Mail Transfer Protocol

**53:tcp, 53:udp**

**DNS** — Domain Name Server

**67:tcp, 67:udp, 68:tcp, 68:udp**

**DHCP/BOOTP** — Dynamic Host Control Protocol or Bootstrap Protocol

**80:tcp, 80:udp**

**HTTP** — Hypertext Transfer Protocol

**110:tcp, 110:udp**

**POP3** — Post Office Protocol, version 3

**119:tcp, 119:udp**

**NNTP** — Network News Transfer Protocol

**123:tcp, 123:udp**

**NTP** — Network Time Protocol

**137:tcp, 137:udp**

**NETBIOS** Name Service

**138:tcp, 138:udp**

**NETBIOS** Datagram Service

**139:tcp, 139:udp**

**NETBIOS** Session Service

**143:tcp, 143:udp**

**IMAP** — Internet Message Access Protocol

**161:tcp, 161:udp**

**SNMP** — Simple Network Management Protocol

**443:tcp, 443:udp**

**HTTPS** — Hypertext Transfer Protocol over TLS/SSL

The last option on the firewall configuration screen relates to trusted network devices. If there's a specific network interface that you want to disable the firewall on, ensure the trusted interfaces are ticked. All traffic will be allowed to traverse these interfaces without being passed through the firewall rule set.

In a few seconds, the PC will load its BIOS and go through the steps of booting your system. After everything is initialised, the BIOS will hand things over to the boot loader.

**26** If you've opted for GRUB, the installer will throw up a simple visual display, asking you to select which operating system you'd like to install. If you're dual-booting, options for both Windows (sometimes DOS will appear in its place, but selecting this will load Windows) and Fedora Core will appear, with the default OS already

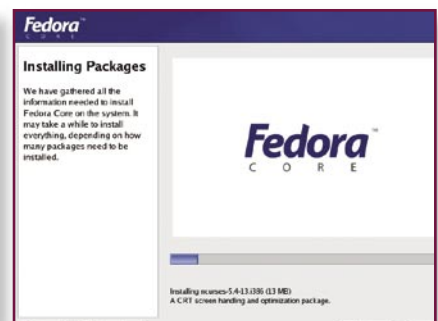
highlighted. If you don't want to load the default operating system, you have 10 seconds (by default) to change your mind. Simply key to the OS you want to boot, then hit Enter.

Pressing any key at the GRUB interface will cause the default operating timer to stop — if you want to boot to the default operating system, select it and then hit Enter.

When Linux loads, a lot of text quickly scrolls up the screen — similar to the beginning of the Linux installation. As



**25** After committing your partition settings, the installer formats your partitions. . .



**26** . . . then the installation files are copied to your system. Reboot once the installer tells you to.

▶ before, the text generated during this loading process lists exactly what is happening. Novice users can disregard these messages for the time being — you'll understand them more as you progress.

**27** After the kernel has loaded, the X Windows GUI should start. Fedora will then display a bar in the centre of the screen that shows you how it's progressing. This process should take considerably longer on the first boot than it will subsequently.

**28** Once everything is loaded, you'll be asked to complete the configuration. The first step in this first-run configuration process is to set the date and time for your system.

Network Time Protocol (NTP) can be enabled if desired. NTP is a time synchronisation protocol that contacts a remote server and synchronises your computer's with a remote server's. The local time zone setting will still be applied to your clock — if you connect to an NTP server in a different time zone, the correct time will still appear on your system (assuming you've configured your time zone correctly).

Once the clock is set, hit Next to proceed.

## USER ACCOUNTS

As mentioned earlier, the root account has access to do practically anything on a Linux machine. For that reason, avoid logging on to the system as the root user wherever possible. Executing the wrong command by accident as the root user can leave your system in an inoperable state. And the chances of a new Linux user entering in such a command is much higher than with an experienced user. Therefore, you should set up at least one additional user account and use this account for running Linux generally.

**29** Creating a user account is easy — set a username (it must be completely in lower case), enter your full name (for identification) and the password for this account. Make sure this password is different to the one for the root user.

On this same screen, you should also see a Use Network Login button. To manage a medium to large number of computers and users, many business networks have network authentication services running on them, such as the Network Information Service (NIS), Lightweight Directory Access Protocol (LDAP) or Microsoft's Active Directory Service (ADS). In general, beginner users shouldn't worry about the options that this button reveals since they're beyond the scope of this guide. But if such an authentication service is running on your network and you're familiar with it, try getting this to work so you can use the same login information on your Linux box that you use on the rest of the network.

## LINUX UPDATE

You're probably familiar with the Windows Update Web site — you load it up in Internet Explorer, it scans your system for the presence of any updates that haven't been installed, then downloads and installs them. Fedora users have a couple of similar services available to them:

■ **Red Hat Network (RHN):** Primarily aimed at users who administer a number of Red Hat Linux or Fedora Linux PCs, this Web-based systems management tool allows you to update, manage and monitor a number of machines via a simple interface. Subscribing to Red Hat Network usually incurs a fee — it's not a free service. Full details about RHN can be found at <http://rhn.redhat.com>.

■ **up2date:** up2date can connect to Red Hat,

update sources and manually download and install any updates that are a part of your current distribution branch. This free service is not as straightforward as the RHN subscription, but is fine for managing a couple of Linux PCs.

**30** The next screen you encounter will ask whether you want to register this system on the Red Hat Network. In most cases, select the option that says "No, I do not want to register my system". This guide explains how to update your system using the up2date tool.

## ADDITIONAL CDS

**31** After specifying your system update preference, the post-installation configurator asks whether you want to install any additional software from the Fedora installation media. For the time being, just hit Next — we'll explain how to install more software on your Fedora box in an upcoming part of this guide.

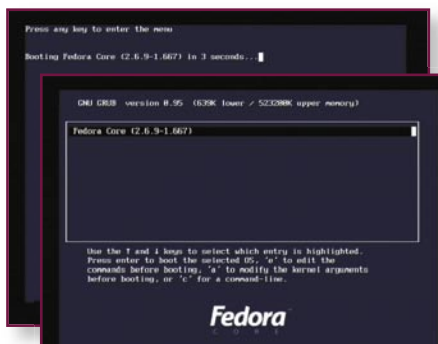
You should now be informed that the post-installation configuration has been completed. Click Next to proceed to the login screen.

## LOGGING ON

**32** Fedora's login screen is pretty self-explanatory — all you need to do to log in to the system is enter your username followed by Enter, then enter your password and Enter. Don't log in as the root user at this stage — use the most recently created account.

Some advanced options are available at the bottom of this screen:

■ **Language.** The left-most option at the bottom of the screen, if it's written in a language you can't understand, allows the user to change the language that is used on the system for their session. By default, this option is only available if language(s) ▶



**27** The GRUB boot screen will appear upon reboot. Make sure that you boot to Linux this first time around.



**28** Once Linux is up, you'll have to go through some last-minute configuration tasks.



**29** Configure the default display resolution and colour depth. Your end users can customise this setting on their desktop if they wish.

► were installed towards the end of the main installation process.

■ **Session.** This option allows you to select which desktop environment you'd like to use for your session. It's available only if you performed a Desktop or Workstation installation of Fedora, or if you selected the GNOME or KDE environments from the package group selection screen. After a fresh installation of Fedora, the default environment is GNOME, but this setting can be changed from within X Windows.

The session screen also allows you to load a "Failsafe Terminal". This will load a third environment called the Tabbed Windows Manager (TWM) — a very basic GUI environment that will allow you work in X Windows if neither GNOME or KDE run on your system.

■ **Reboot.** Shuts down Linux and causes the system to reboot.

■ **Shutdown.** Like the reboot option, this option enables your system to shut down gracefully.

Linux is similar to Windows in terms of shutting down. Both operating systems should be shut down properly whenever possible in order to allow all applications to terminate and the file systems to be unmounted. Simply pulling the plug on your Linux PC could corrupt data on your hard drives and may well lead to an unbootable system.

33 After logging in to Linux, you should be greeted with the GNOME desktop. What you see will be similar to Windows in many ways. A menu accessible from the GNOME button (a red hat under Fedora Linux, or a paw print under other Linux distributions) in the bottom left corner of the screen contains shortcuts to applications installed on your system. There's a quick launch bar, complete with icons for frequently used apps, and a notification bar (bottom right)



**30** Create alternative user accounts to maintain your system's security and integrity.

## Network printers

**In a network of computers, it makes little sense to connect a printer to every PC. Nor does it make sense to move and reconnect a printer to a different PC each time someone wants to print something.**

**If you'd like to connect to a networked printer queue from your Linux PC, change the queue type from "Locally connected" to the networked printer queue type the shared printer is using. For instance, if a Windows machine is sharing the printer, select the "Networked Windows (SMB)" option.**

**You'll then need to supply the wizard with the necessary network and authentication information for the Linux PC to communicate with the printer. If you don't know this information, contact your network administrator.**

**Once you've selected the device your printer is plugged in to, click Forward. The next step is to select the model of the printer in order for the Linux PC to**

**use the correct printer driver.**

**If your printer is not listed, refer to your printer documentation — your printer may be able to "emulate" another printer; that is, you may be able to use another manufacturer's or model of printer driver. Alternatively, your printer manufacturer may have some additional instructions on how to set up your printer under Linux. However, keep in mind that not all printers are Linux-compatible.**

**Once you've selected your printer model, hit Forward again. A confirmation dialog will now be displayed. Clicking on the Finish button will save your printer configuration and ask whether you'd like to print a test page. Select Yes to check that everything is configured and working correctly.**

**If you need to add more printer queues, repeat the process. Otherwise, close the Printer Configuration tool.**

shows the current time.

At this stage, you'll probably want to familiarise yourself with the desktop. Go through some of the system settings and tinker with the appearance of different items. If you come across a setting that you need root access in order to change, you'll be prompted for the root password. Such prompts make it easier for you to keep logged in to your user account for day-to-day tasks, only using root access when it's required.

### PRINTER CONFIGURATION

Under Windows, you set up a new printer by going to the Printers applet within Control Panel and running the Add New Printer

wizard. In Linux, the Printer Configuration applet can be found under GNOME > System Settings > Printing.

34 If you're not already logged in as the root user, you're prompted for the root password to access this applet. Changes to printer settings occur system-wide and has the potential to affect other users of the system.

To add a new printer, click the New toolbar icon. An Add New Print Queue wizard appears. Click Forward to move through the wizard.

35 Enter a name (containing no spaces) for the local printer queue. Multiple queues can be set up for a single printer, each with different print priority, default paper settings and so on. This helps with the management of printer resources if you have a large number of users sharing a multifunction printer. However, if you don't need to manage your printer so stringently, you'll only need to set up the one queue.

At this screen, you can also jot down a brief description of the printer. Doing so is optional, but can help users identify the print queue. Click Forward to proceed.

Next, you'll be asked to select which queue type you're using: that is, how the printer is connected to your computer. If the printer is connected directly to your Linux PC, it will most likely connect via a Parallel or USB interface. If so, ensure that



**31** When asked if you want to install additional software, proceed for the time being. We'll cover that topic in a future guide.



► “Locally connected” is the queue type and select the interface (device) that your printer is connected to from the list below the queue type drop-down menu.

A parallel port under Linux is given the device name `/dev/lp#`, where # is a number assigned to each individual port. If your printer is connected to LPT1 under Windows, it will be connected to `/dev/lp0` under Linux. Similarly, LPT2 corresponds to `/dev/lp1`. Regardless of whether you have a printer plugged into your parallel port, there will be a `/dev/lp#` entry for each parallel port.

USB printers are detected in a different way — a device will only be present once a printer has been plugged in. So, the first discovered USB printer will be assigned the device name `/dev/usb/lp0`; the second discovered printer will be assigned `/dev/usb/lp1` and so on.

If a USB printer is connected and switched on but no USB printer device is detected, try switching your printer off and then on again, and click the “Rescan devices” button.

## SETTING UP INTERNET ACCESS

Unless your system connects to the Internet via LAN, you’ll have to go through the Internet Connection Wizard to get your dialup, ISDN or xDSL modem working properly. Go to GNOME > System Tools > Internet Configuration Wizard. Again, this is a privileged mode applet so if you aren’t logged in as the root user, you’ll be asked to enter the root user password.

Once the application loads, select the type of device you want to use to connect to the Internet and click Forward after ensuring the modem is switched on. This guide shows you how to set up dialup and ADSL connections.

## DIALUP

**36** After hitting the Forward button, your system will be scanned for any attached dialup modems. With any luck, your modem is detected as being attached to a serial port on your system, such as `/dev/ttyS0` (remember, software modems and WinModems won’t work under Linux, as explained in last month’s guide).

The device enumeration convention used for communications devices is similar to how printers are enumerated: if your modem is connected to COM1 under Windows, the equivalent Linux device is `/dev/ttyS0`. COM2 is equivalent to `/dev/ttyS1`, and so on.

Confirm that the detected settings are correct. If Linux was unable to detect your modem, you’ll need to enter the settings manually. Refer to your modem documentation to ensure that you’re using the correct settings. Most modern modems should be able to use a baud rate of 57,600 or higher, using hardware flow control.

**37** The next screen allows you to configure ISP information. Enter the dialup access area code and telephone number for your ISP. Also, complete the login name and password fields — these authentication details will be sent to your ISP on connection to give you access to the Internet.

In most cases, the ISP will supply you with IP address settings, so the default “Automatically obtain IP address settings” option featured on the next page should remain selected. However, if you’ve been allocated a dedicated IP address configuration, select the “Statically set IP address” option and enter the details supplied by your ISP.

Once you’ve finished filling out these settings, your dialup connection should be ready to go.

## ADSL

**38** In most cases, setting up an ADSL service is a little simpler than using a dialup modem service. All you do is select which Ethernet device your modem is connected to (if you use a USB ADSL modem, this should appear as one of your Ethernet devices) and provide your authentication details.

## FINAL NETWORK CONFIGURATION

Once you’ve completed the Internet Connection Wizard, the Network Configuration window should display (this is accessible from GNOME > System Settings > Network). From this window, you can activate (connect) or deactivate (disconnect) any network interface on your PC.

Highlight your newly created Internet connection, and click the Edit toolbar button. If you want the connection to be established whenever your computer starts, select the appropriate checkbox.

**39** Note that if you check the “Allow all users to enable and disable the device” option, any user can connect the system to the Internet; otherwise, the root user will have to do so. Enabling this option might be handy in a home environment where a number of different users are logging on to and using the Linux PC.

Lastly, click on the Advanced tab and check the “Restart if connection dies” box if you want your PC to automatically redial on disconnection. When you’re finished tailoring your connection, click OK. Close the Network Configuration window and be sure to save your settings when prompted.

## BEING UP2DATE

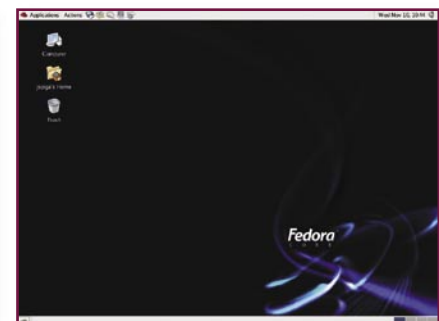
**40** Look at the notification area in the bottom right corner of the screen. ►



**32** Once your post-installation configuration is complete, hit “Next” to be taken to the login screen.



**33** A progress screen shows you how the system is going while starting up required applications once you log on.



**34** The GNOME desktop has similarities to Windows’ and Mac OS’, with its GNOME button, task bars and Notification bars.

▶ You should see an icon with a tick, cross or exclamation mark in it. If not, go to GNOME > System Tools > Red Hat Network Alert Icon.

If configured correctly, the Red Hat Network Alert Notification tool should automatically notify you if there are any updates available for your system. It is also used for the Red Hat Network subscriptions feature (which you should not have signed up to use during the installation process).

While it's possible to configure both the Red Hat Network Alert Notification tool and `up2date` to download updates from the GUI interface, we'll be using the YUM command line tool to tell us when there are updates available. The advantage to using a command line tool is that if something goes wrong and you can't get into X Windows, you'll still be able to securely update your system.

In order to configure the Red Hat Network Notification tool to inform you of when there are updates available, ensure that you are logged on as the root user (this is one of the few occasions where it is advised that you log in as root) and that you have established a connection to the Internet. Right-click the Red Hat Network Alert icon (in the bottom-right corner of the screen) and select Configuration. Clicking Forward will take you through to the Terms of Service agreement — make sure that you agree to these terms before proceeding.

The next screen allows you to specify HTTP Proxy information. If you're using a proxy server to access the Internet, fill this screen out; otherwise, leave it blank, indicating that you have a direct connection to the Internet. Click Forward to apply the information you've entered, and the Alert tool will then connect to the

update servers and then download a list of the current version numbers for all packages that you have installed on your system. By now, the icon in the notification area should be green, indicating that it's transferring data.

The process of searching for updates can take anywhere from a few minutes to over an hour. Once the list of updates has been retrieved, the alert icon will have changed to red with a pulsating exclamation mark if there are updates waiting, or green with a white tick if you're up to date.

### MIRROR, MIRROR...

However before downloading any of the updates, you should ensure that your Linux system is configured to download updates from a mirror near you. Fedora is pre-configured to download updates from a default server but as you can imagine, this server is swamped with traffic due to the high number of people who don't update the following settings. Using a local mirror increases the rate at which updates can transfer to your system, making the update process significantly quicker.

First, load up your web browser (the icon next to the GNOME icon) and go to <http://fedora.redhat.com/download/mirrors.html> and find a mirror site that is close to you. In this example, we're going to be using the <http://mirror.pacific.net.au/linux/redhat/fedora/> mirror site in order to download updates, since we know that this server is not situated too far away from our ISP on the Internet. Click the link to the mirror of your choice — you'll need to investigate the directory structure of your chosen mirror when updating your settings because each mirror uses a slightly different directory naming convention.

Once you're connected to your mirror, minimise your web browser — to change

the relevant settings, there are two files that we have to edit: `/etc/yum.conf` and `/etc/sysconfig/rhn/sources`.

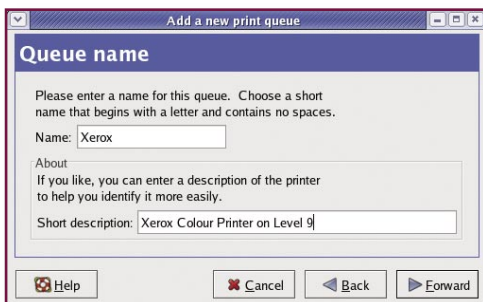
To load `yum.conf` in a text editor, double click on the Computer icon on your desktop, then the Filesystem icon and then find the `etc` icon. Now search for the `yum.conf` file (it should be right at the bottom), right-click on it and select Edit With > `gedit`.

Once the file has opened in the editor, scroll down to the `[base]` section of the file. You should see two entries under this heading: a name and a baseurl. The key directive that needs changing here is the baseurl — change it so that it reflects the directory structure of the mirror that you are using.

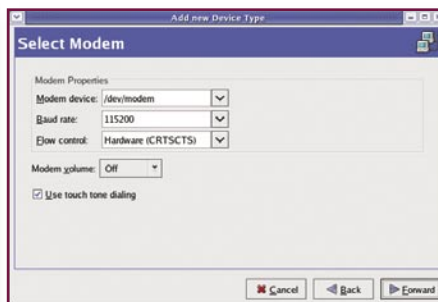
Two variables are often used in this file to ensure that these settings remain valid regardless of which version of Fedora Core is installed on your system. The `$releasever` variable is replaced with the version of Fedora core that you're using — 3 in this case, since you're using Fedora Core 3. The `$basearch` variable corresponds to the base architecture that you're using — in most cases, this will be `i386` (but Athlon 64 users will be using an `x86_64` base architecture). Both of these variables should be interchanged with the appropriate directory names on the mirror server.

In our example, we are changing the baseurl to [http://mirror.pacific.net.au/linux/redhat/fedora/\\$releasever/\\$basearch/os/](http://mirror.pacific.net.au/linux/redhat/fedora/$releasever/$basearch/os/). Make sure that there is a slash on the end of your URL. Also, browse to this location on the mirror server and make sure that there is a directory named headers contained within — if this directory is missing, the updater will not be able to download the descriptions and version numbers of the updates packages and will fail to update your system.

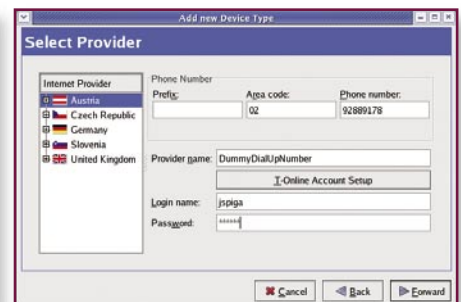
Below the `[base]` section lies the



**35** Enter in a name for your printer queue. This should help identify both the properties of the queue and the printer the queue uses.



**36** Linux will attempt to detect your modem. If it is detected successfully, verify that the correct modem has been found.



**37** Enter in your ISP information. These details will be saved on your Linux PC so you don't have to enter these credentials every time you connect.

► [updates-released] section. The base section is there just in case anything needs to be rolled back to the version that was installed with your version of Fedora core, while the updates-released section contains all of your updates.

Again, update the baseurl directive in this section to reflect the location of the updates on your mirror server. In our case, our baseurl is changed to **http://mirror.pacific.net.au/linux/redhat/fedora/updates/\$releasever/\$basearch/**. Again, make sure that there is a slash at the end of the URL and ensure that the location in question has a headers directory contained within it.

### USE THE SOURCE

After you've made these changes, save the file and close the editor. Then follow a similar procedure to the one shown above to load up the /etc/sysconfig/rhn/sources file.

This file contains the locations that the system updater (yum) will download updates from. Scroll down, and you'll eventually come across two lines that will start with yum fedora-core-3 and yum updates-released-fc3. Again, both of these entries should be updated to reflect paths on the mirror of your choice. In our case, our entries are:

```
yum fedora-core-3 http://
  pacific.net.au/linux/redhat/
  fedora/3/$ARCH/os/

yum updates-released-fc3 http://
  mirror.pacific.net.au/linux/
  redhat/fedora/updates/3/$ARCH
```

Notice the \$ARCH variable in this file? It specifies the architecture of the system that you're using in much the same way as the



**38** Setting up an ADSL modem is easy — just select the Ethernet interface that your modem is plugged in to.

\$basearch variable was used above.

Below these entries, you may want to define a number of other mirrors that can be used to download updates from — just in case the mirror you'd prefer to use is not available. All you need to do is copy the two lines containing your download locations, replace yum with yum-mirror, and modify the URL so that it points to a valid update location. Once again, be sure that these locations contain headers directories.

### DOWNLOADING THE UPDATES

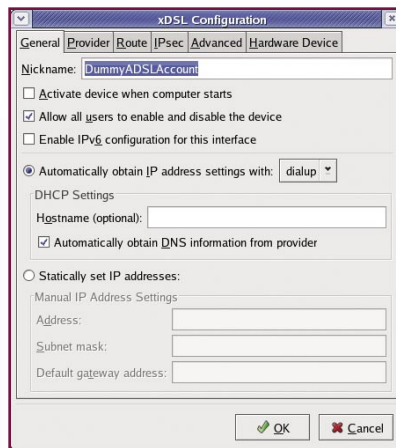
Once you're done adding mirrors, save your file and close the editor — it's now time to update your system. Right-click on your desktop and select Open Terminal. Once the terminal window pops up, enter the following command:

```
yum check-sources
```

The system update application will then connect to the server you are using as your local mirror and retrieves the header information for all of the packages installed on your server. This can transfer up to 50MB of package headers to your system, depending on what you have installed. However, this is a one-off task — it doesn't need to be done on each subsequent update of your system.

When the headers have been downloaded, you'll be dumped back to a command prompt. Now, the downloading of updates can begin. Enter:

```
yum update
```



**39** "Allow all users to enable or disable the device" means that Internet connections can be established by non-root users.

Yum will then compare all of the package headers on your system to the headers that have just been downloaded. A list of newer package versions available for download will shortly be displayed and you'll be asked whether you want to install these updates. Review the list carefully — when you become more proficient at Linux, in some cases, there may be packages that you don't want to update.

Also, you should notice that the kernel packages are never updated. Newer kernels are always installed first, allowing you to switch to them later. We will explain in detail why this is a good idea in another guide.

If you're happy with the updates yum is planning to download, press Y. Yum will then proceed to download all available updates. And within a few minutes (or hours, depending on how many updates there are and your connection speed), you'll have a completely up-to-date system.


### LOGGING OUT

Always shut down the system properly to prevent data loss. To do this, click on GNOME > Log Out. Your desktop will appear less bright and the Log out dialog will appear. Several actions are available: Log out, which will take you back to the login screen; Shut down; and Restart the computer.

One other option is available at this screen — a check box that says "Save current setup". If this option is enabled, the position of desktop elements will be saved in your current profile. The next time you log in, your desktop will appear exactly the same as you left it. **ETC**

### Next month . . .

In next month's Mastering Linux guide, you'll get your first look at the daunting command line. We'll also explain how the Linux file system works.



**40** The Red Hat Network Alert icon can check whether updates are available and will inform you if your system is up-to-date.

