

Integrating LINUX[®] and Windows[®]



- ◆ Your complete Linux/Windows integration guide
- ◆ Detailed coverage of dual-boot issues, data compatibility, and networking
- ◆ Implementing Samba file/print services for Windows workstations
- ◆ Providing cross-platform database access

MIKE MCVNE

OPEN SOURCE TECHNOLOGY SERIES



Integrating Linux and Windows

By [Mike McCune](#)

Publisher : Prentice Hall PTR
Pub Date : December 19, 2000
ISBN : 0-13-030670-3
Pages : 416

The complete solutions guide for every Linux/Windows system administrator!

This complete Linux/Windows integration guide offers detailed coverage of dual-boot issues, data compatibility, and networking. It also handles topics such as implementing Samba file/print services for Windows workstations and providing cross-platform database access. Running Linux and Windows in the same environment? Here's the comprehensive, up-to-the-minute solutions guide you've been searching for!

In *Integrating Linux and Windows*, top consultant Mike McCune brings together hundreds of solutions for the problems that Linux/Windows system administrators encounter most often. McCune focuses on the critical interoperability issues real businesses face: networking, program/data compatibility, dual-boot systems, and more. You'll discover exactly how to:

Use Samba and Linux to deliver high-performance, low-cost file and print services to Windows workstations

Compare and implement the best Linux/Windows connectivity techniques: NFS, FTP, remote commands, secure shell, telnet, and more

Provide reliable data exchange between Microsoft Office and StarOffice for Linux

Provide high-performance cross-platform database access via ODBC

Make the most of platform-independent, browser-based applications

Manage Linux and Windows on the same workstation: boot managers, partitioning, compressed drives, file systems, and more.

For anyone running both Linux and Windows, McCune delivers honest and objective explanations of all your integration options, plus realistic, proven solutions you won't find anywhere else. *Integrating Linux and Windows* will help you keep your users happy, your costs under control, and your sanity intact!



Library of Congress Cataloging-in-Publication Data

McCune, Mike.

Integrating Linux and Windows / Mike McCune.

p. cm. — (Open technology series)

Includes index.

ISBN 0-13-030670-3 (alk. paper)

1. Linux 2. Microsoft Windows (Computer file) 3. Operating systems (computers) I. Title.
II. Series.

QA76.76.O63 M387 2000

005.4'469—dc21

Editorial/Production Supervision:

Wil Mara

Acquisitions Editor:

Miles Williams

Editorial Assistant:

Richard Winkler

Marketing Manager:

Kate Hargett

Manufacturing Manager:

Alexis R. Heydt

Cover Design Director:

Jerry Votta

Cover Designer:

Talar Agasyan

Art Director:

Gail Cocker-Bogusz

Illustrations:

Wil Mara

© 2001 Prentice Hall PTR

Prentice-Hall, Inc.

Upper Saddle River, NJ 07458

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the author and publisher.

The publisher offers discounts on this book when ordered in bulk quantities. For more information, contact: Corporate Sales Department, Prentice Hall PTR, One Lake Street, Upper Saddle River, NJ 07458. Phone: 800-382-3419; FAX: 201-236-7141; E-mail: corpsales@prenhall.com

Names such as company names, trade names, font names, service names, and product names appearing in this book may be registered or unregistered trademarks or service marks, whether or not identified as such. All such names and all registered and unregistered trademarks, service marks, and logos appearing in this book or on its cover are used for identification purposes only and are the property of their respective owners.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Pearson Education Asia Pte. Ltd.

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

Library of Congress Cataloging-in-Publication Data	3
Introduction.....	8
Chapter 1. Having Linux and Windows on the Same PC	10
1.1 Partitions	10
1.2 Filesystems.....	10
1.3 Partition Naming	13
1.4 Linux and Windows 95/98.....	14
1.5 Setting up Linux and Windows 3x/9x on Separate Partitions.....	18
1.6 Partitioning an Existing Hard Drive.....	19
Chapter 2. Accessing ext2 Partitions with Windows	29
2.1 Accessing ext2 Partitions with DOS and Windows 3.1	29
2.2 ltools.....	29
2.3 Accessing ext2 Partitions with Windows 9x	31
2.4 Accessing ext2 Partitions with Windows NT and 2000	33
Chapter 3. Mounting Windows Partitions with Linux	34
3.1 Accessing Compressed DOS/Windows Drives with Linux	34
3.2 Adding a Partition to the <i>fstab</i>	35
Chapter 4. Emulators.....	37
4.1 DOS	37
4.2 Windows.....	37
4.3 VMware	38
4.4 FreeMWare.....	43
4.5 Win4Lin	43
4.6 Conclusion	43
Chapter 5. Internet Applications	45
5.1 Web Server Compatibility.....	45
5.2 FrontPage Extensions	46
5.3 Using Microsoft Office Files on the Web.....	47
5.4 Web Browsers.....	47
5.5 Email.....	49
5.6 Streaming Media.....	51
5.7 Chat	55
5.8 Instant Messaging	56
5.9 Internet Security.....	56
Chapter 6. Business Applications.....	61
6.1 Microsoft Office	61
6.2 Corel WordPerfect Office	61
6.3 Other Commercial Productivity Suites	63
6.4 Open Source Office Suites	64
6.5 Web-Based Suites.....	64
6.6 Reading and Writing Microsoft Office Files	66
6.7 Exporting MS Office Files.....	66
6.8 Importing and Exporting MS Office Files with Linux.....	68
6.9 Using MS Office Documents with Star Office	69
6.10 Checkpoints When Importing and Exporting	71
6.11 Financial Programs	72
6.12 Graphics Programs	73
6.13 The Last Word	76
6.14 Conclusions.....	76

Chapter 7. Databases.....	77
Using Databases.....	77
Choosing a Database.....	78
Connecting Databases.....	78
ODBC.....	79
Chapter 8. Fun and Games.....	80
8.1 Games.....	80
8.2 Game Servers and Extras.....	83
8.3 Classic Games.....	83
Chapter 9. The Linux Desktop.....	84
9.1 Switching Desktops.....	84
9.2 Configuring Desktops.....	85
9.3 Themes.....	87
9.4 Conclusion.....	90
Chapter 10. Running Applications through a Network.....	91
10.1 X-Windows.....	91
10.2 Citrix WinFrame.....	92
10.3 VNC.....	94
10.4 Conclusion.....	109
Chapter 11. Introduction to Windows and Linux Networking.....	110
11.1 Net BIOS.....	110
11.2 TCP/IP and Active Directory.....	110
11.3 Net BIOS over ICP/IP.....	112
Chapter 12. Introduction to Samba.....	113
12.1 How Samba Started.....	113
12.2 How Samba Works.....	113
Chapter 13. Setting Up Samba as a Windows NT Server.....	115
13.1 Setting up Samba as a Stand-Alone Windows NT File Server.....	115
13.2 Adding a Samba Server to an Existing Network.....	124
13.3 Samba as a Primary Domain Controller.....	125
Chapter 14. Connecting Linux to Windows PCs.....	127
14.1 smbclient Command-Line Options.....	128
14.2 smbclient Commands.....	131
14.3 smbtar.....	137
14.4 smbprint.....	138
14.5 smbfs.....	138
14.6 Sharity.....	139
14.7 Conclusion.....	140
Chapter 15. Printing with Samba.....	141
15.1 printtool.....	141
15.2 Testing the Printer.....	145
15.3 Setting up Samba for Printing.....	146
15.4 Automatic Print Driver Installation.....	147
Chapter 16. Using NFS and NIS in Linux and Windows.....	151
16.1 Setting up Linux as an NFS Server.....	151
16.2 Using an NFS Client on Linux.....	154
16.3 Using NFS on Windows.....	156
16.4 Setting up an NIS Server on Linux.....	158
16.5 Setting up an NIS Client on Linux.....	161
16.6 NIS Support for Windows.....	163

Chapter 17. Implementing FTP, Telnet and Other UNIX Protocols in Windows	166
17.1 Setting Up the FTP Server for Windows.....	167
17.2 Setting up FTP for Linux	168
17.3 telnet and Remote Services for Linux.....	168
17.4 Secure Shell (SSH).....	169
Appendix A. Disk Error Codes	178
Appendix B. Samba Documentation	180
The GNU License.....	180
The Samba FAQ	184
Just what is SMB?.....	205
Appendix C. Samba Man Pages.....	215
Lmhosts (5)	215
nmbd.....	216
Samba (7).....	219
smb.conf (5)	221
smbclient (1).....	293
Installation	303
Diagnostics.....	304
Version	304
Author.....	304
smbd (8).....	304
smbpasswd (5)	310
smbpasswd (8)	313
Appendix D. TCP/IP Documentation.....	318
TCP/IP Network Resources List.....	318
Private IP Network Addresses	348

Introduction

In early 1998, I was looking at re-installing Windows 95 for the third time. Granted, I stress computers more than the average user, but this was getting old. I had been playing around with Linux since early 1995 and it looked like a good time to use it as my primary desktop.

I already knew how to install and configure Linux, but I had never used it as a desktop. I found plenty of books on Linux configuration, a few on using it as a server, but nothing on using it as a desktop. Instead I had to scour the Internet for useful information. What I have tried to do for this book is compile what I have learned over the past two years. Hopefully, this will save you the time and frustration of finding it yourselves.

So, how does Linux compare to Windows? As with anything else, each one has its own strengths and weaknesses.

Windows is king of the desktop for good reason. It has a polished interface and more end-user applications than any other operating system. It is also pre-installed on most new PCs, making it an easy, safe choice for most PCs. These factors combine to give Windows about 90% of the desktop market.

Linux is based on UNIX and inherits its security and stability from it. Linux is the most popular choice for public Web servers and it also holds about 25% of the small server market. It is also free (or nearly so) and comes packed with lots of useful tools for programming and server management.

These distinctions aren't permanent, however. Several groups are working on polishing Linux's interface. There is also a rush to develop more end-user applications for Linux. Large PC makers such as Dell, Compaq, and IBM are starting to offer Linux pre-installed on PCs.

While the market for Linux is comparatively small, Linux grew from less than 1% of the desktop market in 1998 to about 4% in 1999. This is amazing considering the Apple Macintosh, which has been around for 15 years, is holding at 5% of the desktop market.

Windows is also working to gain a foothold in the traditional strengths of the UNIX (and Linux) market: security, stability, and scalability. Microsoft put billions of dollars into the recently released Windows 2000 to address these issues. While the jury is still out on whether it succeeded, early reports say that Windows 2000 is much improved over earlier versions of Windows in these areas.

There are also many other reasons for choosing an operating system. They can often draw fanatical devotion (just ask a dedicated Macintosh user). Despite (or maybe because of) its success, Microsoft has some very dedicated enemies. Just search the Internet for "Satan" or "Antichrist" and you will be surprised how many anti-Microsoft sites you hit. The Microsoft Antitrust case was also pushed forward by some dedicated foes. Some users try Linux as an alternative to Windows. It may not be the best way to choose an operating system, but never discount the power of fanatical devotion.

Such devotion is not necessary. Linux and Windows can peacefully coexist on the same computer. It is even possible to run Linux and Windows at the same time! The whole first section of this book is devoted to making coexistence as easy as possible.

The middle section is dedicated to finding useful applications for your Linux systems. Sometimes the same application is available for both Linux and Windows; in other cases,

equivalent applications are available; and in a few cases, the applications are only available for Windows. The good news is that most people can do everything they need to do with either Linux or Windows.

The last section deals with networking. This is a rather advanced topic, but networking is moving from the Fortune 500 into homes and small businesses at a rapid rate. Networking is getting inexpensive enough to offer the same advantages that large businesses have long enjoyed: sharing files, printers, and Internet connections. The increasing use of high-speed Internet connections in the home will continue to drive up demand for home networking.

Fortunately, both Linux and Windows have programs that allow easy connection to each other. Samba allows Linux to act as a Windows file server. Additionally, the NFS and LPD programs allow Windows to use Linux's native protocols.

So which is better, Linux or Windows? That is like asking whether a car or truck is better. They are built for different purposes. Windows plays the traditional role of the car; it is more polished and aimed at the mass market. Linux plays the traditional role of a truck; it is durable and intended to be used as a work vehicle. But like cars and trucks, the roles are starting to overlap. Linux is becoming more polished and easier to use and Windows is concentrating more on security and stability. Windows is still the choice for most users' desktops, but it is no longer the only choice. As you will see in this book, Linux is a solid choice for a server and a viable alternative in the desktop market.

Chapter 1. Having Linux and Windows on the Same PC

[Section 1.1. Partitions](#)

[Section 1.2. Filesystems](#)

[Section 1.3. Partition Naming](#)

[Section 1.4. Linux and Windows 95/98](#)

[Section 1.5. Setting up Linux and Windows 3x/9x on Separate Partitions](#)

[Section 1.6. Partitioning an Existing Hard Drive](#)

1.1 Partitions

There is no need to get rid of Windows to run Linux. In fact, there are many ways to run both of them on the same PC. Each operating system has its own strengths and weaknesses, so often having both on the same PC can be an advantage.

Before the actual installation, we need to go over some basics of Linux and Windows such as partitions and filesystems.

A partition is a way of sectioning off space on a hard drive. Most PCs have their hard drive partitioned into one large drive. It doesn't have to be this way. Drives can be divided into several partitions. This is often done to separate the programs from the data and also for storing multiple operating systems on the same drive.

The first section of a hard drive contains information on the partitions, including where the start and end of each partition is located. It also contains the location of the boot loader, which starts loading the operating system. Each operating system has its own boot loader. Windows 3x, 95, and 98 use `IO.SYS` and `DOS.SYS`, Windows NT uses `NTLDR`, and Linux uses `LILLO`. There are also commercial and shareware boot loaders, such as Norton System Commander, that are designed to make it easier to boot with multiple configurations and multiple operating systems.

1.2 Filesystems

There are also several different filesystems used by Linux and Windows. A filesystem is simply a way of organizing files on a partition. Windows uses FAT, FAT16, FAT32 (VFAT), and NTFS (NT Filesystem). The native filesystem for Linux is ext2, although it supports many other filesystems.

FAT is the original filesystem used by DOS. It is an eight-bit filesystem and will support partitions of up to 32 MB. This was no problem in the early 1980s, when most PCs didn't even have hard drives.

FAT supports the following file attributes:

- **Read-only**— When set, the file can't be deleted or changed.
- **Archive**— Determines whether a file has been changed. This is used by many backup programs.
- **Hidden**— The file doesn't show up in the directory contents.

- **System**— Used for system files. System files are treated differently by the operating system.

Later, as hard drives came into use, the 32 MB limitation of FAT became a burden and an improved FAT16 replaced it. FAT16 increased the available size of the filesystem to 2 GB. Other than the filesystem size, FAT16 is essentially the same as the original FAT filesystem. FAT16 is supported by DOS 4.0 and greater, all versions of Windows, and all current versions of Linux.

With Windows 95 release 2, Microsoft introduced FAT32. This increased the size of the filesystem to 2 terabytes, which is larger than any hard drives currently available for PCs. It is also faster and more robust than FAT16.

NTFS is the native filesystem for Windows NT and 2000. Like FAT32, it also supports 2-terabyte filesystem sizes, but the boot partition is currently limited to 7.8 GB. For some files such as database files, NTFS can support up to 16 exabytes. NTFS offers better reliability and security than any FAT-based filesystem.

The reliability factors are beyond the scope of this book, but NTFS security considerations need to be covered. First of all, everything in the filesystem has an owner. By default, the user who creates an object (anything in the filesystem is an object) is the owner. The owner has full rights to the object unless the rights are taken away.

There are also groups, which contain users. Three special group accounts are: administrator, everyone, and guest.

- The administrator account has all rights to the filesystem. This account can change, create, and delete all objects as well as change the rights of other accounts.
- Everyone is a group that includes all the user accounts on the system. This account is used to change the rights for every user on the system.
- The guest account is a default account with minimal rights. It is often used for accounts such as FTP access accounts, which only need access to a few specific files.

Files in NTFS have the same attributes as files in the FAT filesystem: read-only, hidden, system, and archive. Each user and group can also be assigned rights to objects in the NTFS filesystem. The rights that can be assigned are:

- **List folder contents**— Shows up in a directory listing.
- **Read**— Can read the contents of the object.
- **Read and execute**— Can read and execute the object.
- **Write**— Can change or delete the object.
- **Modify**— Can change the rights on the file.
- **Full control**— Has all of the above rights.

There are three settings for the rights: allow, deny, and inherited.

- **Allow**— Allows rights on the object.
- **Deny**— Takes away rights on the object.

- **Inherited**— If neither allow nor deny is specified, the object will inherit the rights of the directory above it.

To view the rights of an object on NTFS, right-click on the object and choose Properties. Then select the Security tab.

Ext or ext2 is the native filesystem for Linux partitions. Ext is the original filesystem for Linux and ext2 is an improved version of it. Objects (such as files, directories, and devices) in Linux support three properties: read, write, and execute.

- **Read**— If set, allows the object to show up in a directory listing and be read.
- **Write**— If set, allows the object to be written and deleted.
- **Execute**— If set, allows the object to be executed. This must also be set for directories.

An object has three sets of rights: owner, group, and everyone.

- **Owner**— The user who created the file, unless it is changed.
- **Group**— The group that owns the file is the group to which the owner belongs, unless it is changed.
- **Everyone**— The right for all other users on the system.

To view the rights of an object, type `ls -l <object name>`. For example, to find the rights of `index.txt`, type `ls -l index.txt`. The output of this command is as follows:

```
-rwxrwxr--  1 root  root    6230 Dec 21 00:12 index.txt
```

Let's examine what this output means, starting with the first character:

- **First character**— This is a special attribute such as a directory, link, or device driver. A link in Linux is similar to a shortcut in Windows.
- Next three characters (`rwX`)— The owner has read, write, and execute properties.
- Next three characters (`rwX`)— The group has read, write, and execute properties.
- Next three characters (`r--`)— Everyone has the read property.
- `1 root`— The owner's ID number and name.
- `root`— The group name, which is also root.
- `6230`— The size of the file in bytes.
- `Dec 21 00:12`— The last date and time the file was modified.
- `index.txt`— The filename.

You can change the rights of a file with `chmod`. The owner and group can be changed with `chown`.

There is one special account in Linux: `root`. The root account is created automatically when Linux is installed and it has full rights to all objects in the filesystem.

1.3 Partition Naming

Linux and Windows have different ways of naming partitions. Windows simply assigns each partition a letter starting with C. Letters A and B are reserved for floppy drives, since the first PCs came with two floppy drives. The remaining drive letters are assigned as follows:

1. The first primary partition on each drive.
2. The volumes inside the extended partitions on each drive.
3. The remaining primary partitions on each drive.
4. The CD-ROM drive.

For example, if you had two hard drives each with two primary partitions with two volumes in extended partitions on each drive, they would be named as follows:

- C: The first primary partition on the first drive.
- D: The first primary partition on the second drive.
- E: The first extended partition on the first drive.
- F: The second extended partition on the first drive.
- G: The first extended partition on the second drive.
- H: The second extended partition on the second drive.
- I: The second partition on the first drive.
- J: The second partition on the second drive.
- K: The CD-ROM drive.

These drive letters can't be changed in Windows 3x and 9x, but they can easily be changed in Windows NT and 2000.

Figure 1.1. Disk Administration for Windows 2000.



For Windows NT, go to Start -> Programs -> Administrative Tools -> Disk Administrator. For Windows 2000, go to Start -> Programs -> Administrative Tools -> Configure Your Server. From here, choose File Server -> Open Computer Management, then choose Storage -> Disk Management. This allows you to change the partitions and drive letters. You can change the drive letter by simply right-clicking on the drive and then choosing Change Drive Letter and Path. You can then add, edit, or delete the drive. One note, though: You cannot change your boot partition. This is good because the system won't boot if you do!

Linux treats partitions differently. The first two letters denote the type of drive, the next letter is the drive letter, and the last character is the partition number of the drive. There are four main drive types: IDE (*hd*), SCSI (*sd*), ESDI (*ed*), and RAID (*md*, *rd*, or *ida*). For example, *hda1* is an IDE drive (*hd*), it is the first IDE drive (*a*), and the first partition (1). *sd c5* would represent a SCSI drive (*sd*), the third SCSI drive (*c*), and the fifth partition (5).

1.4 Linux and Windows 95/98

First, let's go over what happens when Linux and Windows 9x boot up.

1.4.1 Booting Linux

When Linux boots, it loads the `LILLO` program, which stands for Linux LOader. This then loads the kernel, which is the core of the operating system. Finally, modules are loaded from the `/etc/rc.d` directory. Actually, the Linux boot process is a little more complicated than this, but this description is good enough for our purposes.

The `LILLO` program is configured by using the `/etc/lilo.conf` file. A typical `lilo.conf` would look like this:

```
boot=/dev/hda
```

```
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
image=/boot/vmlinuz-2.2.13-4mdk
    label=linux
    root=/dev/hda5
    read-only
```

Let's go over the lines one at a time:

- `boot`— This is the device that contains the boot files.
- `map`— This is the location of the map file. The map file is a binary file containing disk parameters for the system. The default is `/boot/map`.
- `install`— This is the file that is installed as the boot sector. The default is `/boot/boot.b`.
- `prompt timeout`— This is how long it waits before booting in tenths of a second. This allows time to enter boot parameters manually. In the above example, it is five seconds (50 tenths of a second). If you have a multiple boot system, pressing `<SHIFT>` will bring up the boot choices. You can set up to 16 different boot configurations.
- `image`— This is the kernel. The parameters below are kernel parameters:
 - `label`— The name that shows up on the boot menu.
 - `root`— The location of the filesystem.
 - `read`— The filesystem is mounted read only so that it can be checked for errors with `fsck`. It is then remounted as read/write.

These aren't the only parameters for `lilo.conf`.

1.4.2 Troubleshooting LILO

LILO loads the four letters in "LILO" as it goes through the four stages of loading. This can be helpful in troubleshooting. If LILO stops while loading, the letters displayed tell where it failed:

- **(<nothing>)**— No part of LILO has been loaded. LILO either isn't installed or the partition on which its boot sector is located isn't active.
- `L <error> ...`— The first-stage boot loader has been loaded and started, but it can't load the second-stage boot loader. The two-digit error codes indicate the type of problem. (See the section titled "Disk Error Codes" in [Appendix A](#).) This condition usually indicates a media failure or a geometry mismatch.
- `LI`— The first-stage boot loader was able to load the second-stage boot loader, but could not execute it. This can either be caused by a geometry mismatch or by moving `/boot/boot.b` without running the map installer.

- `LIL`— The second-stage boot loader has been started, but it can't load the descriptor table from the map file. This is typically caused by a media failure or by a geometry mismatch.
- `LIL?`— The second-stage boot loader has been loaded at an incorrect address. This is typically caused by a subtle geometry mismatch or by moving `/boot/boot.b` without running the map installer.
- `LIL-`— The descriptor table is corrupt. This can either be caused by a geometry mismatch or by moving `/boot/map` without running the map installer.
- `LILO`— All parts of `LILO` have been successfully loaded.

If Linux doesn't load, it will give an error code. The meanings of the error codes are listed in [Appendix A](#).

1.4.3 Booting Windows 9x

The way Windows boots is slightly different. The partition points to the boot sector, which loads the two text files `config.sys` and `autoexec.bat`. It then loads the Windows equivalent of the Linux kernel, `win.com`. The configuration file for `win.com` is the Registry, which consists of two binary files: `system.dat` and `user.dat`. While `config.sys` and `autoexec.bat` are text files, Registry files need to be edited with `regedit`.

There are several ways to install Linux and Windows 3x/9x on the same machine. They are listed below, starting with UMSDOS. Before anything is done, back up your hard drive and run `scandisk` on the hard drive to correct any errors. This will save you a lot of trouble down the road.

1.4.4 UMSDOS

The easiest way to install Linux on a Windows 9x machine is to use UMSDOS, which allows Linux to co-exist with Windows on a FAT or FAT32 partition. It allows Linux to boot directly from a command prompt on the FAT partition. On the plus side, you avoid re-partitioning the hard drive or disturbing the existing Windows installation. On the minus side, you do lose the security and robustness of Linux's native ext2 filesystem.

If you plan on using UMSDOS, it is best to choose a distribution with an automated UMSDOS install. Most modern distributions have a UMSDOS install. One Linux distribution designed to easily install as UMSDOS is Phat Linux. It is a hefty download, so unless you have a fast connection and a lot of time, you might want to buy the CD-ROM. As a bonus, you also get technical support with the purchase.

Phat Linux was originally put together by two high school kids who wanted an easy way to install Linux. To install it, unzip the files to the `\phat` directory and run `linux.bat`. This should start booting up into Linux. If it hangs during the bootup, it may be necessary to reboot, press F8 on bootup, and choose the Safe Mode command prompt. After the bootup, there will be a menu that lets you choose the configuration to set up the video, sound, and networking. The `linux.bat` is only three lines, as shown below:

```
loadlin vmlinuz initrd=ramdisk.gz mem=96M
echo Linux failed to load.
pause
```

The first line actually loads Phat Linux. The command is broken down below:

- `loadlin`— This is the loader for UMSDOS.
- `vmlinuz`— This is the Linux kernel.
- `initrd`— Points to a memory image. Phat Linux uses a compressed memory image to start the operating system.
- `ramdisk.gz`— A compressed file that contains the complete Linux filesystem. This can also be a directory.
- `Mem`— The amount of memory in megabytes.

Phat Linux is based on Red Hat with a KDE desktop, so anyone familiar with Red Hat should be able to use Phat Linux.

UMSDOS can easily be installed manually since support for it has been built into the Linux kernel for some time. Most distributions will set up everything for you, but knowing how it works will help with troubleshooting.

The default directory in which to install UMSDOS is `\linux`. This acts like the root directory for Linux. Under this are the standard Linux directories of `bin`, `etc`, `lib`, `root`, `sbin`, `tmp`, `usr`, and `var`.

Then there is the problem of swap space. Normally, Linux creates a separate swap partition. This has the advantage of being fast and robust. With UMSDOS, however, it is usually better to create a swap file. To create a swap file, type the following:

```
dd if=/dev/zero bs=<block size> count=<number of blocks>
of=/<swap
file name> mkswap /<swap file name> <swap file size in
bytes>
sync
swapon /<swap file name>
```

Then add the following to your `/etc/fstab`:

```
/<swap file name> swap swap defaults
```

To determine the block size, run `chkdsk` and it will list the number of bytes in each allocation unit. For example, if our block size is 2048 and we want to create a 16 MB swap file called `swap`, we would enter the following commands:

```
dd if=/dev/zero bs=2048 count=8 of=/swap
mkswap /swap 16384
sync
swapon swap
```

Then we would add the following to the `/etc/fstab`:

```
/swap swap swap defaults
```

1.4.5 Booting with UMSDOS

`loadlin` is the loader for UMSDOS. From a DOS prompt, its command would look like this:

```
loadlin <Linux kernel> root=<root for UMSDOS>
```

For example, if the kernel is located at `c:\linux\boot\vmlinux` and UMSDOS is installed in `c:\linux`, the command for `loadlin` would be:

```
loadlin c:\linux\boot\vmlinux rw root=c:\linux
```

The `rw` option tells `loadlin` to load the Linux filesystem with the read and write options. You can also load it read-only (`ro`) if needed.

You can also copy `loadlin.exe` and `vmlinux` to a bootable DOS floppy if you want to use a boot floppy to load Linux. This will save you the trouble of configuring the boot loader.

1.4.6 Managing UMSDOS Filesystems

UMSDOS puts a file called `--linux-.---` in each directory. This file stores the extended attributes for the Linux files. As we discussed above, DOS only supports read, write, hidden, archive, and system attributes. Linux also has user, group, and executable attributes. These are stored in the `--linux-.---` file.

The `--linux-.---` file is maintained with the `umssync` utility. This utility will create the `--linux-.---` file if it doesn't exist. If the file does exist, `umssync` will update Linux attributes stored in the file. It is a good idea to run this utility often to keep the information up-to-date. The following line can be added to `cron` jobs or `rc.d` (see your user manual for an explanation of how to do this):

```
/sbin/umssync -r99 -c -i+ <root of Linux file system>
```

The `-c` option will only update existing `--linux-.---` files and not create new files. Directories without the `--linux-.---` file in them are ignored by Linux.

1.4.7 Working with DOS and UMSDOS

Files created by DOS are invisible to Linux unless `umssync` is used. If you try to create a file in Linux with the same name as a DOS file, it will say that the file already exists. Other than that, there are no problems with running DOS and UMSDOS on the same partition. You can even use a DOS defragmentation utility on the partition without affecting the UMSDOS filesystem.

1.5 Setting up Linux and Windows 3x/9x on Separate Partitions

Before working with partitions on any drive, make a good backup of all current partitions and have a DOS boot floppy. If done right, the existing files should be unharmed, but there is always a chance of something going wrong.

The easiest way to load Linux on a separate partition is to load it onto a different drive. If I have a primary IDE drive on the first IDE port with Windows 98 on it (`hda1`) and I put a primary drive as the second IDE port for Linux, it would be `hdb`. During the install, I would simply install a Linux partition and swap drive on `hdc`. The only choice I have left then is whether to boot with `LILO` or `loadlin`.

1.5.1 Using LILO to Dual Boot

To use `LILLO`, install `LILLO` onto `hda`. The installation program will usually ask you where to install `LILLO`. You will need to check the documentation for instructions on how to install `LILLO`, since it varies with different installation programs.

Next, you need to edit `/etc/lilo.conf`. It will look something like this:

```
boot=/dev/hdc
map=/boot/map
install=/boot/boot.b
prompt
timeout=50
image=/boot/vmlinuz-2.2.13-4mdk
    label=linux
    root=/dev/hdc5
    read-only
```

Add this to the end of `lilo.conf`:

```
other=/dev/hda1
table=/dev/hda
label=Windows98
```

This will give you the choice of "linux" or "Windows98" on bootup.

1.5.2 Using loadlin to Boot Up

Of course, `loadlin` can also be used to load Linux from the second drive. First put `loadlin` and `vmlinuz` onto your Windows partition. During bootup, press F8 and choose the Safe Mode command prompt. In the above example, we would run `loadlin` from the root of C: with the following parameters:

```
loadlin vmlinuz root=/dev/hdc5 ro
```

This would load the kernel `vmlinuz`, set the Linux root to partition 5 on the primary drive on the second IDE port (`hdc5`), and set the root to read-only (`ro`). Change these settings to match the drive on which that Linux was installed.

1.6 Partitioning an Existing Hard Drive

Although the best way to load Linux on a separate partition is to load it onto its own hard drive, not everyone has more than one hard drive. With hard drives selling for about \$100 these days, it is less of a problem, but if you must, load Linux onto a separate partition on the same hard drive as Windows. There are several programs that allow you to manipulate partitions. If you have 200-1200 MB available on your hard drive, you can create a Linux partition out of the free space.

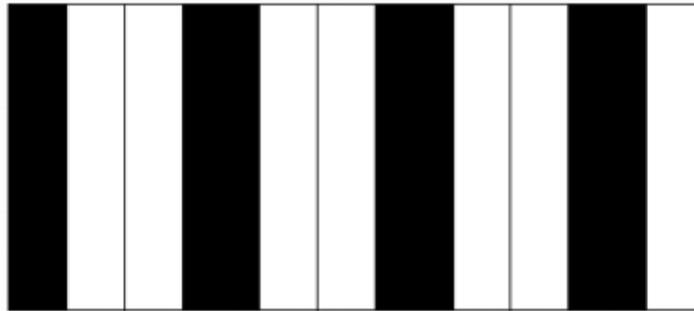
The first thing to do before manipulating the partitions is to back up your hard drive. Next, run `scandisk` on the drive in thorough mode to make sure there are no errors that could cause problems. Then defragment the hard drive with the Windows `defrag` program. This will move all the files to the beginning of the drive and free up space at the end of the drive for a new partition.

Before we repartition the drive, keep in mind that the boot partition of both operating systems must be within the first 1024 cylinders of the hard drive. This is within the first 504 MB on most hard drives. This is a limitation of PC hardware that dates back to when the maximum size of hard drives was 504 MB.

Once we have backed up and checked the drive, we need to decide how to partition the drive: the hard way or the easy way. Let's start with the hard way.

Figure 1.2. Disk partitioning.

WINDOWS FRAGMENTATION



As you use a drive, the files become spread out over the drive (fragmentation). The dark areas represent the files, and the light areas represent the empty space.

EMPTY

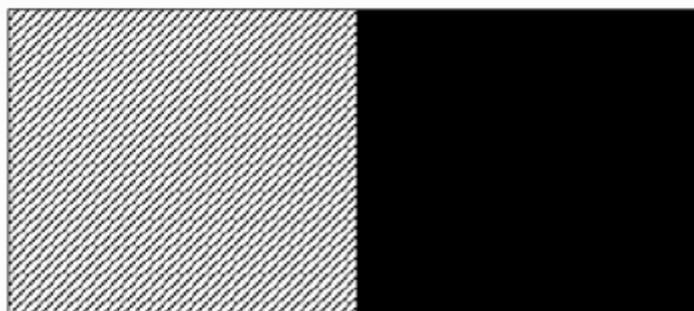
WINDOWS



Defragmenting the drive moves all the files to the beginning of the drive. This leaves end of the drive free.

LINUX

WINDOWS



We can then create a partition at the end of the drive, and install Linux on that partition.

1.6.1 fdisk

If you have all your data backed up (THIS WILL DESTROY ALL DATA ON THE DRIVE), you can use `fdisk` to re-partition the drive. First, boot to DOS in Safe Mode. Then run `\windows\fdisk`. Choose Delete Partition. Then choose Extended. Then delete the primary partition. Once this is done, use `fdisk` to create a DOS partition.

Once you have created the DOS partition, install Windows and all your applications on it. Then restore your data. This is a lot of work and will probably take several hours. There is an easier way.

1.6.2 Resizing Existing Partitions

There are several programs that will resize a partition. There is the commercial program Partition Magic that has a nice graphical interface, and it works with all PC partitions. It is well worth its price. Partition Magic is available at <http://www.powerquest.com>. Many Linux installation programs allow the hard drive to be re-partitioned during installation. Instructions should be included in your distribution's documentation. If your distribution doesn't allow this, you can do it manually with FIPS.

1.6.3 FIPS

If you don't mind a text-based interface, there is FIPS, which is a DOS program that resizes partitions. It is on most Linux distribution CD-ROMs under the subdirectory `/dosutils`. Otherwise, you can get a copy of it at: <http://www.igd.fhg.de/nashaefe/fips/> Compared to many other programs, FIPS has few command-line switches. Those it has are listed below:

- `-t` or `-test`— This doesn't write anything to disk.
- `-d` or `-debug`— This will write errors to the file `FIPSINFO.DBG`. It can help in diagnosing problems.
- `-h` or `-help` or `-?`— The help text.
- `-n<num>`— The drive number to split.

1.6.4 Restrictions of FIPS

There are some restrictions on what you can do with FIPS. The first is that your hard drive must support using INT 13 for low-level disk access. There are a few older Adaptec drive controllers that don't support INT 13, but almost all other drives do, including all newer Adaptec drive controllers.

FIPS also will not work with FAT12. FAT12 is used on partitions that are smaller than 10 MB. This shouldn't cause any problems since it would be useless to split a 10 MB partition anyhow.

You can only split standard FAT and FAT32 partitions. FIPS will not work on extended partitions, NTFS, HPFS (Os2's filesystem), Linux, or any other non-FAT partition.

FIPS will not work with disk managers such as OnTrack. You must uninstall OnTrack, which requires deleting the partition and re-installing (see "fdisk" above).

Also, don't reduce the original partition to less than 4085 clusters. A FAT partition needs at least 4085 clusters. There is FAT12 for smaller partitions.

And lastly, you can't create a new partition on a drive if it already has four partitions. This is because FAT only supports four partitions to a disk.

There are a few special situations that require extra steps to re-partition. If you are using Windows 3.1, any disk compression software (Stacker, DoubleSpace, or SuperStor), or disk mirroring software (Image or Mirror), see the section below entitled "Special Situations with FIPS".

1.6.5 Using FIPS

If your partition doesn't have any of the restrictions listed above, FIPS can be used to resize your partition. There are three things that need to be done before you resize the partition:

1. Back up your data—Something may go wrong.

2. Run `scandisk` or `chkdsk /f`—This will correct errors on the drive.
3. Run `defrag`—This will free up space at the end of the drive.

If there is no space available at the end of the drive, check for hidden or system files by typing:

```
cd \
```

Then type one of the following:

- `dir /a:h /s`— Searches for hidden files.
- `dir /a:s /s`— Searches for system files.

First, check to see what these files do so that you don't delete any important files. Don't, for instance, delete or move the `io.com` or `dos.com` or your system will not boot!

For example, the programs `IMAGE` and `MIRROR` create the files `image.idx` and `mirror.fil`. These files are used to recover a corrupted disk and they are created each time the system boots. To delete these files, first change the attributes with the commands:

```
attrib -r -s -h image.idx
attrib -r -s -h mirror.fil
```

You can then delete the file normally (`del image.idx` or `del mirror.fil`). Next, make a bootable floppy and copy the files `RESTORRB.EXE`, `FIPS.EXE`, and `ERRORS.TXT` to this disk. You can make a bootable floppy with `format a: /s` or `sys a:.` Some PCs are set so that they won't boot off a floppy. Consult your computer's documentation (or a local computer guru) for instructions on how to enable booting from a floppy.

Now that the disk is prepared, boot from the floppy you just created. At the prompt, type `FIPS`. You can exit at any time by pressing `<CTRL> C`. `FIPS` will first try to detect the operating system you are using. Since we booted off a DOS floppy, there should be not problem. It will next try to detect your hard disks. Then it will read the partitions on each drive and display the partition table such as the one shown below (from `FIPS.DOC`):

			Start			End		
Start	Number of							
Part.	bootable	Head	Cyl.	Sector	System	Head	Cyl.	Sector
Sector	Sectors		MB					
1	yes	0	148	1	83h	15	295	63
149184	149184	72						
2	no	1	0	1	06h	15	139	63
63	141057	68						
3	no	0	140	1	06h	15	147	63
141120	8064	3						
4	no	0	0	0	00h	0	0	0
0	0	0						

This is a lot of information. The most important data is the number of megabytes. `FIPS` will next check the root sector for errors. If you have more than one partition, you will be

asked which partition to split. Once you choose a partition, FIPS will show the drive information:

```
Bytes per sector: 512
Sectors per cluster: 8
Reserved sectors: 1
Number of FATs: 2
Number of rootdirectory entries: 512
Number of sectors (short): 0
Media descriptor byte: f8h
Sectors per FAT: 145
Sectors per track: 63
Drive heads: 16
Hidden sectors: 63
Number of sectors (long): 141057
Physical drive number: 80h
Signature: 29h
```

FIPS will then check the drive for errors and free space. If it exits with an error message, make sure you did all of the preparation steps above.

If there are no errors, FIPS will show the size of the original and new partitions. Use the left and right cursor (arrow) keys to change the size of the two partitions. Once you have them at the desired size, press <ENTER>.

FIPS will then recheck the new partition to make sure it is empty. It will show the changes to be made to the partition. You may press `r` to re-edit or `c` to continue. It will then ask if you want to write the changes to disk. Answer `y` to save the changes and FIPS will exit. Reboot the machine to save the changes.

After rebooting, use `scandisk` to check both partitions to make sure they are okay. If there are errors, you can restore the partition by rebooting with the DOS disk and running `restorrb`.

The new partition is a standard DOS partition when first created. When you install Linux, you can use the installation program to delete the new partition and create a Linux partition. Just be sure you delete the correct partition!

Booting Linux to a new partition is no different than booting it on a new hard drive. See the previous sections on dual booting with Linux and Windows 3x/9x.

1.6.6 Special Situations with FIPS

Windows 3.1. If you are using Windows 3.1, you must delete the swap file before splitting the drive. To do this, go the Control Panel (in the Main folder) and choose 386 enhanced. Then choose the Virtual Memory, and change to none. After the drive is split, you can turn the swap file back on.

Stacker, SuperStor, DoubleSpace, and other Disk Compression Programs. These programs create a compressed volume on any disk with a compressed file on it. Then they move all the files to the compressed volume and rename the volumes. The uncompressed volume is typically C: and it contains the boot files and compression program. The compressed volume is D: and it contains the compressed file.

Splitting this drive can be tricky. Be sure to get a good backup because if the compressed file is damaged, the whole drive becomes unreadable. The following steps should allow you to add a new partition to a compressed drive:

1. 1. Make sure you have enough free space on the compressed drive to create the new partition.
2. 2. Use the disk checking software that comes with the drive to check for errors.
3. 3. If you are running Windows 3.1, remove the swap file.
4. 4. Use the disk compression utilities to decrease the size of the compressed volume.
5. 5. Defragment the compressed volume (D:).
6. 6. Use FIPS and split the compressed volume (D:)

If the compressed volume can't be defragmented, try the following steps:

1. 1. Copy your disk defragmentation utility (for example, `defrag.exe` or `diskkorg.exe`) and `attrib.exe` (in `C:\DOS`) to the bootable floppy drive.
2. 2. Boot with the floppy and remove all the hidden and system attributes from the files on the compressed drive (D:). Use `dir /a:h` and `dir/a:s` to find the hidden and system files.
3. 3. Defragment the compressed partition (D:).

After this is done, you should be able to split the drive.

1.6.7 NTFS Partitions

Unfortunately, FIPS and UMSDOS don't work on NTFS partitions. You either have to delete the partition and re-create it or use a commercial partition utility such as Partition Magic. Partition Magic is available in most large computer stores. Information on it is available at <http://www.powerquest.com>.

Before manipulating NTFS partitions, make a rescue disk, which saves important system information to a diskette so it can be restored if something goes wrong. For Windows NT, use Start -> Run, type in `rdisk`, and choose Make Emergency Repair Disk. For Windows 2000, go to Start -> Accessories -> System Tools -> Backup. On the Welcome tab is a button for Emergency Repair Disk. Follow the instructions.

Windows NT and 2000 don't use `fdisk`. They have their own partitioning program called Disk Administrator. During the install, you are asked how big to make the partition. By default, the partition will be the entire drive. NT partitions are limited to 4GB on the boot partition and 7.5GB on all other partitions. Windows 2000 doesn't have this limitation.

Once there is free space on the drive, install Linux on the free space. If Linux and NT are on the same drive, you will want to use NT as the boot loader and install `LILLO` on the Linux partition. Most distributions allow you to specify where to install the boot loader. See your distribution's documentation and help files for information on this.

If you have Linux and NT on separate drives, you can use `LILLO`. See the section below titled "Using LILLO".

Before we start manipulating the NTFS partition, we need to download BootPart by G. Volland. This program can create and manipulate NTFS boot sectors. Go to <http://www.winimage.com/bootpart.htm>.

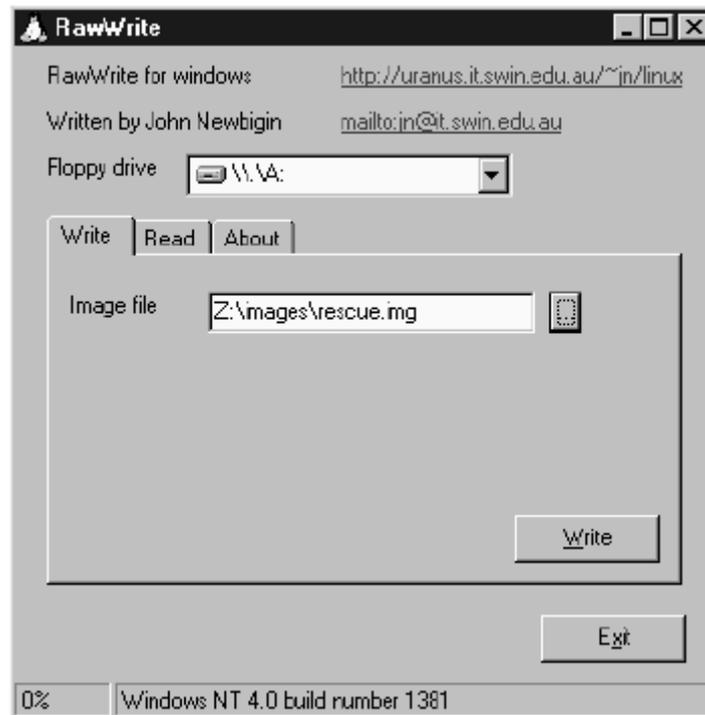
1.6.8 Using NT's Boot Loader

The boot process for Windows NT is different than Windows 3x/9x. The main difference we need to be concerned with is that Windows NT uses a file instead of the partition's boot record to load the operating system. This means that we will have to make a boot file for Linux if we want to use Windows NT's boot loader to load Linux.

To create a Linux boot file, we need to copy the boot sector of the Linux partition to a file. To do this, boot from a Linux rescue diskette and use `dd` to copy the boot sector to a file. If you don't have a rescue diskette, you can create one in Windows by running the

program `rawwrite.exe`, which can be found in the `\dosutils` directory on a Linux installation CD-ROM. Once the program is open, choose the Write option, and then choose the `rescue.img` file in the `\images` directory on the Linux installation CD-ROM ([Figure 1.3](#)).

Figure 1.3. Use RawWrite on your Linux installation CD-ROM to copy the rescue image to the A: drive.



Boot off the rescue diskette and copy the boot partition to the floppy as follows:

```
dd if=<Linux partition> of=<name of file> bs=512 count=1
```

Let's go over the command step by step:

- `if`— The location at which to start copying. In this case, we are starting at the beginning of the Linux partition.
- `of`— The output file's name.
- `bs`— The block size. The boot sector is 512 bytes.
- `count`— The number of blocks copied. We are only copying one block—the boot sector.

Let's use an example of having Linux on `hda2` (the second partition on the first hard drive) and let's name the file `boot.lnx` (the Linux boot sector):

```
dd if=/dev/hda2 of=boot.lnx bs=512 count=1
```

Next, copy the file `boot.lnx` to a DOS formatted floppy:

```
mcopy boot.lnx a:
```

`mcopy` is part of the `mtools` programs. For more information go to <http://www.tux.org/pub/tux/knaff/mtools>. You could also use these commands to copy `boot.lnx` to a floppy:

```
mount -t msdos /dev/fd0 /mnt/floppy
cp boot.lnx /mnt/floppy
umount /mnt/floppy
```

Now, reboot into NT, log in as administrator, and copy `boot.lnx` from A: to the root of C:. Then, edit `boot.ini`. First, take off the system and read-only attributes:

```
attrib -s -r c:\boot.ini
```

Next, edit `boot.ini` with notepad and add a line for booting to Linux:

```
[boot loader]
    timeout=30
    default=multi(0)disk(0)rdisk(0)partition(1)\WINNT
    [operating systems]
        multi(0)disk(0)rdisk(0)partition(1)\WINNT="Windows NT
Workstation
    Version 4.0
        multi(0)disk(0)rdisk(0)partition(1)\WINNT="Windows NT
Workstation
    Version 4.0 [VGA mode] /basevideo /sos
    C:\BOOT.LNX="Linux"
```

The last line will load Linux. When you are done editing and saving `boot.ini`, change the attributes back:

```
attrib +s +r c:\boot.ini
```

When you boot up NT, you will get the following menu, which allows you to select the system to boot:

```
OS Loader V4.00
Please select the operating system to start:
Windows NT Workstation Version 4.0
Windows NT Workstation Version 4.0 [VGA mode]
Linux
```

If you change the boot sector of Linux, you must make a new `boot.lnx`. This usually only happens when you upgrade the Linux kernel.

All these steps can be done automatically with `BootPart`. Simply run the `BootPart` program and choose the location of the Linux partition. It will edit the `boot.ini` and create the Linux boot file automatically.

1.6.9 Using LILO

Windows NT requires its own master boot record on the drive. To boot NT from `LILO`, NT must be loaded on a separate drive. To use `LILO` to boot to NT, Linux must be on the first drive and NT on the second drive.

If we install LILO on the first partition, we must modify `lilo.conf`. Most Linux installation programs allow you to create a LILO menu item for another OS (operating system). Check your manuals to see if it is supported. If not, manually edit the `lilo.conf`, which will look something like this:

```
boot=/dev/hda
map=/boot/map
install=/boot/boot.b
default=linux
keytable=/boot/us.klt
prompt
timeout=50
message=/boot/message image=/boot/vmlinuz
label=linux
root=/dev/hda5
read-only
```

If Windows NT is on the second drive, add the following to the end of `lilo.conf`:

```
other=/dev/hdb1
table=/dev/hda
loader=/boot/any_d.b
label=WindowsNT
```

This is what the added lines mean:

- `other`— This points to the first partition on the second hard drive (`dev/hdb1`).
- `table`— This is where the drive table is. This is required by LILO.
- `loader=/boot/any_d.b`— Required when not booting from the primary drive.
- `label`— The name for the section.

Chapter 2. Accessing ext2 Partitions with Windows

[Section 2.1. Accessing ext2 Partitions with DOS and Windows 3.1](#)

[Section 2.2. ltools](#)

[Section 2.3. Accessing ext2 Partitions with Windows 9x](#)

[Section 2.4. Accessing ext2 Partitions with Windows NT and 2000](#)

2.1 Accessing ext2 Partitions with DOS and Windows 3.1

Two tools are used to access ext2 (Linux native) filesystems from DOS. One is `ext2tools`, which provides read-only access, and is available here: [ftp://sunsite.unc.edu/pub/Linux/system/filesystems/ext2/](http://sunsite.unc.edu/pub/Linux/system/filesystems/ext2/). There is also `ltools`, which gives read and write access to ext2 partitions.

2.2 ltools

`ltools` provides command-line tools, a Web server that allows access with a browser, and a Java tool. The command-line tools mimic the functions of many DOS utilities. They run in DOS and all versions of Windows (3x, 9x, and NT). They should also work with Windows 2000, although they haven't been extensively tested with it.

There is some unusual behavior when writing to the ext2 filesystem with `ltools`. Any write operation will set the filesystem to "not clean." This will cause no problems except that `fsck` (the disk checking tool for Linux) will run the next time the system is booted.

Also, while running `ltools` in a Windows 9x DOS box, the hard drive will lock during use. This will require a floppy disk in your system even if the floppy is not used. The floppy drive will spin, but nothing will be changed on it.

The `ltools` use the Linux-style forward slashes (/) for directories and DOS-style wildcards. For example, all the files in the `/root` directory would be `/root/*.*`.

`ltools` comes with the following programs:

- `ldir [-h | -v | -? | -part] [-s/dev/hd..] [Linux_Directory]`— Performs like the `dir` command in DOS or `ls` in Linux. The switches have the following meanings:
 - **-h -?**— The help screen.
 - **-v**— The version information.
 - **-s**— The Linux drive name e.g., `/dev/hda1`, etc.
 - **-part**— Lists all partitions on the drive.
 - **Linux_Directory**— The directory on the ext2 filesystem. The default is `/`.
- `lread [-h | -v | -?] [-s/dev/hd..] Linux_File [DOS_File]`— Lists or copies a file. Works like `type` or `copy` in DOS. The switches are used to denote:
 - **-h -?**— The help screen.
 - **-v**— The version information.
 - **-s**— The Linux drive name, e.g., `/dev/hda1`, etc.

- **Linux_File**— The Linux file to be viewed or copied.
- **DOS_File**— The DOS file to copy to. This must be a valid DOS filename. Long filenames are supported under Windows 9x.
- **ldrive /dev/hd..**— This sets the default Linux partition. If no partition name is given in the command, it will default to this. For example, to set the default partition to `hda1`, the command would be `ldrive /dev/hda1`.
- **lcd <directory name>**— Sets the default directory for the Linux partition. This will be used if no directory is given otherwise. For example, to set the default directory to `/root`, the command would be `lcd /root/`.
- **ldel [-h | -v | -?] [-s/dev/hd..] Linux_File**— Deletes a file, directory, or link using the following switches:
 - **-h -?**— Help.
 - **-v**— Version.
 - **-s**— The Linux partition.
- **lchange [-h | -v | -?] [-s/dev/hd..] [-fMODE] [-uUID] [-gGID] Linux_File**—
 - **-h -?**— Help.
 - **-v**— Version.
 - **-s**— The Linux partition.
 - **-f**— Changes the file attributes of the Linux file like `chmod`. Uses octal numbers. See your manual for `chmod` for details.
 - **-u**— Changes user ID like `chown`. See your manual for `chown`.
 - **-g**— Changes group ID like `chown`. See your manual for `chown`.
- **lwrite [-fMODE] [-uUID] [-gGID] DOS_PathDOS_File Linux_File**— Copies a file. The switches are the same as `lchange`.
- **ren [-h | -v | -?] [-s/dev/hd..] [-fMODE] [-uUID] [-gGID] Linux_File_old_name Linux_File_new_name**— Renames a Linux file, directory, or symbolic link. The switches are the same as `lchange`.
- **lmkdir [-h | -v | -?] [-s/dev/hd..] [-fMODE] [-uUID] [-gGID] new_Linux_directory**— Makes a Linux directory. The switches are the same as `lchange`.
- **lln [-h | -v | -?] [-s/dev/hd..] [-fMODE] [-uUID] [-gGID] LinuxLinkTarget Linux-LinkSource**— Creates a new symbolic link. A symbolic link is similar to a shortcut in Windows. The switches are the same as `lchange`.

2.2.1 Graphical Interfaces for ltools

`lttools` comes with two graphical interfaces: `LREADsrv` and `LREADgui`. While they do work, if you want a graphical interface for Windows 9x, NT, or 2000, there are better tools available.

LREADsrv. `LREADsrv` is a part of `lttools` that lets you use your Web browser to read and write to the ext2 partition. It also requires some `.GIF` and `.HTM` files to run. When `lttools` are extracted, they put `LREADsrv` and all the required files in the `/bin` directory.

To use `LREADsrv`, first make sure the command-line tools work properly. `LREADsrv` uses these tools to access the ext2 partition. Then run `LREADsrv` and connect your browser to `http://localhost`.

If you already have a Web server on your PC, you don't need to load `LREADsrv`. Simply load the file <http://zzzhlp.htm>. You can also read the drive from a remote machine by putting in the host name or IP address as in <http://10.0.0.1/zzzhlp.htm>.

`LREADsrv` is alpha software. It still has many bugs to iron out, including:

1. It is not a multi-user program. Multiple users can lead to corrupted files or directories!
2. Its error checking is weak. The errors of the underlying applications often won't show up in the browser window.

LREADgui. `LREADgui` is a Java interface for Java Runtime version 1.1 or higher. It gives a graphical interface to `lttools`. For it to work, you must have `lttools` in your path and the `ldrive` set to the proper ext2 partition.

There is also an `LREADjav` that allows up to three different remote hosts to connect to `LREADgui`. Just set the remote hosts and port number (the default is 1605) in the remote menu to `LREADgui`.

`LREADjav`, like all Java programs, requires a relatively fast machine to operate on (a 200 MHz machine is considered slow for a Java program)! The screen also tends to flicker on long directory listings. These problems have more to do with Java itself than the `LREADjav` program.

2.3 Accessing ext2 Partitions with Windows 9x

`FSDEXT2` allows read-only access of an ext2 partition from Windows 9x. It is no longer being actively developed, but it works and is available at <http://www.yip-ton.demon.co.uk/>.

For read and write access, there is `Explore2fs`. It looks and behaves like the Windows Explorer and works with Windows 9x and NT. It is available at <http://uranus.it.swin.edu.au/~jn/linux/explore2fs.htm>.

`Explor2fs` supports a wide range of features, including:

- Drag and drop.
- Support for all block sizes.
- Support for floppy disks 1.44MB and 120MB.
- Supports Zip and Jaz drives.
- Supports Windows 98 extended partition scheme.
- Fast write support. As fast, if not faster, than reading.
- Export file(s).
- Export file as text.
- Export directory.
- View / execute file.
- View symbolic links.
- Import file.
- Import directory.
- umask for new files.

- Delete file.
- Delete directory.
- Make directory.
- Rename.
- Modify file mode (via Properties box).
- Change uid and gid.
- Create device nodes.
- Create symbolic links.
- Large disk support.

Some features to be added to future versions are:

- Format / create filesystem.
- Import `.tgz` file. This is the native compression format for UNIX, like ZIP is for Windows.
- Language support.

When it first starts, Explore2fs is in read-only mode. To enable write support, go to View -> Options, then choose the Debug tab. There is a check box to enable write support. Restart the program to enable write support.

Explore2fs shows only the ext2 partitions. To manipulate a file on the ext2 partition, right-click on the file. This will pull up a menu with the following items on it:

- Properties
- Export File
- Export as Text
- Rename
- Delete
- Import File
- Create

The Properties option pulls up a box that has two tabs: File and Attribute. The File tab shows all the properties of the file and the Attribute tab shows the properties that can be changed.

The Export File and Export File as Text options will bring up a Save As box that will allow you to save the file to a Windows drive. The only difference between the command is the Export File as Text will convert a UNIX text file to a DOS/Windows text file.

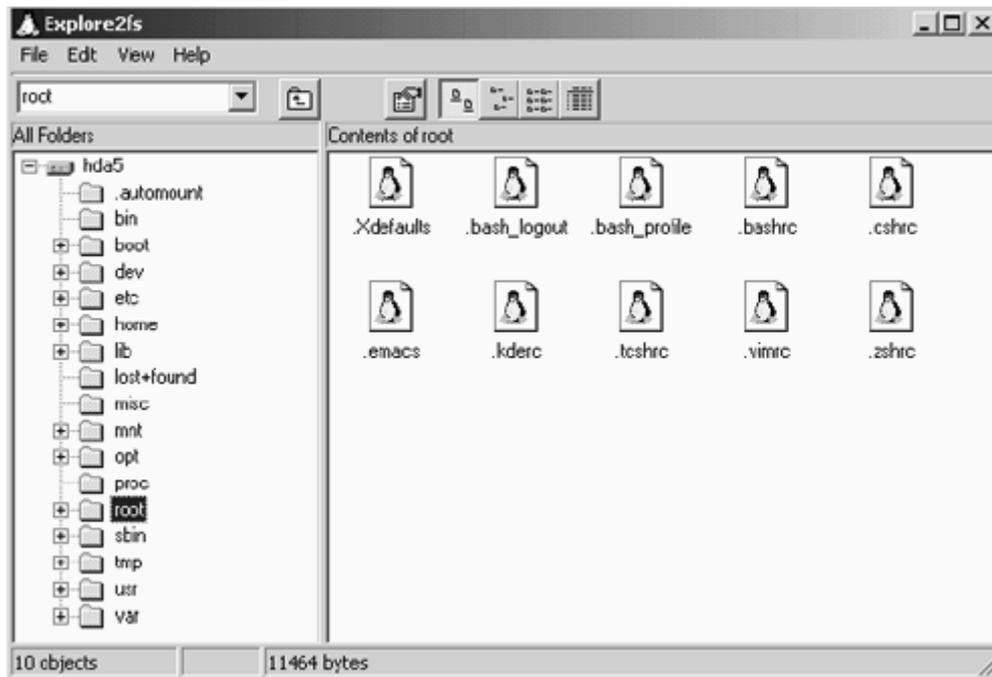
The View option will open the file with the default Windows text editor.

The Rename and Delete options do exactly what they say.

The Import File option will bring up a box that will allow you to choose files to copy from a Windows partition to the current directory on the ext2 drive.

The Create option will allow you to make a new character device, block device, or symbolic link. These are the three basic types of Linux files. See your Linux reference material for an explanation of these files.

Figure 2.1. Explore2fs.



2.4 Accessing ext2 Partitions with Windows NT and 2000

`Ext2fsnt` is a driver that allows Windows NT and 2000 to read from and write to ext2 partitions. It is available at <http://www.chat.ru/~ashedel/ext2fsnt/>. You will also need the RAR archiver to extract the file (<http://www.rarsoft.com>). Just keep in mind that this is beta software, so there is a chance that it can damage your filesystems! It is important to use the included `sync.exe` after manipulating ext2 volumes to make sure the changes are flushed from the cache.

To install `Ext2fsnt`, log in as an administrator, extract the files, and copy the file `ext2.sys` to your `%systemroot%\system32\drivers` directory. The default `%system-root%` is `c:\winnt`.

Next, merge `ext2.reg`. To do this, go to Start-> Run and type in `regedit`. Then choose Registry and Import Registry File. Find the `ext2.reg` file.

Reboot the machine to activate the driver. Edit the `go.cmd` to add your ext2 partitions, then run `go.cmd`.

On a Windows NT system, you can replace the file `%systemroot%\system32\drivers\rs_rec.sys` with the version that comes with `Ext2fsnt` instead of having to run `go.cmd` every time you need to access the ext2 filesystem. Do not replace this driver on a Windows 2000 system! The driver is not fully supported on Windows 2000.

After installation, set up the ext2 partitions as Windows drives. With Windows NT, add them to the Registry. For example, to set up the second partition on the first hard drive as E:, add the following line to to the registry:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session
Manager\DOS_Devices]
"E:"="\\Device\\Harddisk0\\Partition2"
```

For Windows 2000, simply use the Disk Management tool to assign drive letters.

Chapter 3. Mounting Windows Partitions with Linux

Fortunately, most Windows filesystems are supported by the Linux kernel. The `mount` command can be used to mount the partition:

```
mount [options] device directory
```

For example, to mount the first partition on the first drive into the directory `/mnt/drivec`:

```
mount /dev/hda1 /mnt/drivec
```

Now, changing the directory to `/mnt/drivec` (`cd /mnt/drivec`) will show the contents of the first partition, which would be the C: drive in Windows.

The `mount` command will usually detect the filesystem type on the partition. If not, use the `-t` option to specify the filesystem type. The supported Windows filesystems are:

- `iso9600`—Standard CD-ROMs. Newer kernels have support for Rock Ridge Extensions and Joliet Drives (used by Windows 9x CD-ROMs).
- `Msdos`—FAT16 and FAT12.
- `Ntfs`—NTFS.
- `Umsdos`—A UMSDOS partition on a FAT partition.
- `Vfat`—FAT32.

For example, to load an NTFS partition:

```
mount -t ntfs /dev/hda1/ /mnt/drivec
```

A word of caution: The NTFS drivers are read and write, but use the write mode at your own risk!

If your kernel doesn't have support for these filesystems, there are kernel patches for different filesystems, including the UDF filesystem. The UDF filesystem is used on DVD CD-ROMs and some CD-R and CD-RW drives. Unfortunately, the UDF modules for Linux are currently read-only.

Adding new filesystem support to the kernel requires recompiling your kernel. If this isn't covered in the user manual, there are instructions on compiling kernels at <http://www.linuxdoc.org/HOWTO/Kernel-HOWTO.html>.

The kernel patches for the various filesystems are at:

- `FAT32`— <http://www-plateau.cs.berkeley.edu/people/chaf-fee/fat32.html>.
- `Joliet CD-ROMs`— <http://bmrc.berkeley.edu/people/chaffee/joliet.html>.
- `NTFS`— <http://www.informatik.hu-berlin.de/~loewis/ntfs/>.
- `UDF`— <http://trylinux.com/projects/udf/>.

3.1 Accessing Compressed DOS/Windows Drives with Linux

Unlike other drivers, support for compressed drives is not built into most kernels. There are two programs that allow access to compressed DOS/Windows drives: `thsfs` and `DMSDOS`. `thsfs` allows read-only support for Double Space and Drive Space (Microsoft's disk compression programs). `thsfs` is available at `ftp://ftp.ai-lab.fh-furtwan-gen.de/pub/os/linux/local/thsfs.tgz`.

The other program, DMSDOS, uses the loopback device to allow read and write access to compressed drives. It supports Double Space Drive Space, and Stacker (a third-party drive compression program). Like other filesystem support, it is a kernel patch and it is available at: <http://fb9nt.uni-duisburg.de/mitarbeiter/goeckel/software/dmsdos/>.

Once DMSDOS is added to the kernel, the compressed drives can be loaded like ordinary FAT partitions except that a loopback is added. For example, to load a compressed FAT16 drive that is the first partition on the first drive, type:

```
mount /dev/hda1 /DOS mount -t msdos -o loop
/DOS/dblspace.001 /mnt
```

The partition is mounted on `/DOS` and the compressed virtual filesystem (cvf) is mounted on `/mnt`. The cvf is used to mount compressed data and `/DOS` is used to mount decompressed data.

DMSDOS should detect the type of compression used for the drive, but there are options that can be used if there are problems mounting the drive. The DMSDOS options use the following format:

```
DMSDOS -t <partition type> -o cvf_options=option1+option2...
```

There are five option types for DMSDOS. The DMSDOS documentation has a list of all the values for each option:

- `comp:xxx`—This specifies the compression type. Since DMSDOS will automatically detect the compression type, this is seldom used.
- `cf:xxx`—This is a whole number from 1 to 12. 1 is the fastest, but gives the least compression; 12 is the slowest, but gives the most compression.
- `bitfaterrs:xxx`—This is what to do if errors are found on the compressed partition. The default is to set the partition to read-only.
- `loglevel:xxx`—This creates a log file that is useful in troubleshooting. Look at the `dmsdos.h` in the source code for an explanation of log levels.
- `speedup:xxx`—This should not be used unless you know what you are doing. This sets how the drive is accessed. Setting it too low will cause DMSDOS to be very slow and setting it too high could corrupt the drive.

3.1.1 Some Problems with DMSDOS

Of course, accessing a compressed partition is slower than accessing an uncompressed partition because of the compression and decompression involved. This can be kept to a minimum by defragmenting the compressed partition. Double Space, Drive Space, and Stacker come with defragmentation utilities that should be used on a regular basis.

Using DOSEmu with DMSDOS can also cause problems. DOSEmu will often write directly to disk and bypass the DMSDOS file caching. You may need to unmount the compressed partition to use DOSEmu and avoid corrupting the partition or crashing the system.

The DMSDOS documentation contains an extensive list of error codes and kernel messages.

3.2 Adding a Partition to the `fstab`

The file `/etc/fstab` is a text file that stores information on mountable drive systems. Adding an entry to this file will save having to enter long command lines every time you need to mount a partition. A typical `fstab` file would look something like this:

```
/dev/hda1 /boot ext2 defaults 1 2
/dev/hda5 / ext2 defaults 1 1
/dev/hda6 swap_upgrade swap defaults 0 0
/mnt/floppy /mnt/floppy supermount fs=vfat,dev=/dev/fd0 0 0
/mnt/cdrom /mnt/cdrom supermount fs=iso9660,dev=/dev/cdrom 0
0
```

This entry is broken down into six sections, which are separated by a space:

1. The device or remote filesystem to mount.
2. The mount point.
3. The filesystem type.
4. The mount options.
5. Which filesystems to dump. A 0 means don't dump.
6. The `fschk` priority. The root filesystem should be first (1) and the other filesystems should be second (2). 0 means to skip.

Adding a compressed drive to the `fstab` is rather tricky since it requires two entries: one for the compressed volume space (cvs) and another for the drive itself.

```
/DOS/drvspace.001 /DOSF msdos loop 1 0
/dev/hda1 /DOS msdos defaults 1 0
```

The first line loads the cvs as follows:

- `/DOS/drvspace.001`—The compression driver.
- `/DOSF`—The directory used by the compression program.
- `msdos`—The filesystem type.
- `loop`—This option is required for the compression drive.
- `1`—Set the dump to on.
- `0`—Don't `fsck` the drive.

The second line mounts the drive:

- `/dev/hda1`—The actual drive to mount.
- `/DOS`—The mount directory on the Linux partition.
- `msdos`—The filesystem type.
- `defaults`—Use the default mount options.
- `1`—Set dumb to on.
- `0`—Don't `fsck` the drive. (This could corrupt the drive!)

Some of these settings may vary with your setup. Also, depending on your system, you may have to reverse the order of the two lines since the cvs must be loaded first.

Chapter 4. Emulators

Emulators are programs that allow you to run an application written for one OS (operating system) on another OS. In this chapter, we will focus on emulators that allow Windows and DOS programs to run on Linux machines. There are no emulators to run Linux programs on Windows, but on the other hand, most programs are available for Windows.

4.1 DOS

Linux can run most DOS programs using DOSEmu. There are thousands of DOS programs that are still useful. Ten years ago, DOS was the most popular OS for PCs, with both Microsoft (MS-DOS) and Digital Research (DR-DOS) selling their own version of it. Today, you can get copies of DOS for *Caldera* at <http://www.caldera.com> and *FreeDOS* at <http://www.freedos.org>. DOS is still used in many embedded devices such as scanners and cash registers because it is small and has many development tools available for it.

Caldera purchased DR-DOS from Novell in the mid 1990s and now targets it toward embedded applications. You can download it from their Web site, but you must pay for it if you choose to use it. It is the best debugged and most complete version of DOS currently available. It also has a lot of extra features that weren't included in MS-DOS such as multitasking, TCP/IP, and a Web browser.

FreeDOS is a project started in 1994 by Jim Hall. It is under the GPL and thus is free and open source. It is currently used by Linux as part of DOSEmu, which is included in almost every Linux distribution. To use it, simply type `dosemu <dos program>` from the command prompt.

FreeDOS and DOSEmu are good programs to have on a low-end computer. FreeDOS will run on an 8088 (the original IBM PC) with 640KB of memory. Linux will run on a 386 with 4 MB of RAM (although X11 requires 8 MB just to load).

So what can you do with DOS programs? You can turn an old PC into a machine that can surf the Web and use email. If you have a 386/33 with at least 4 MB of RAM (which you can find thrown in dumpsters), you can run the Arachne Web browser found at <http://home.arachne.cz>. This program contains a graphical Web browser, an email client, and a PPP dialer. While it doesn't support many of the fancier features of the Web such as Java, it does support frames and graphics. It can be run inside Linux using DOSEmu and there is a Linux port in the works.

4.2 Windows

Wabi is a commercial program that allows Windows 3x applications to run inside of X-Windows. It is available for many commercial versions of UNIX as well as Linux. The good news is that it runs Windows 3.1 programs very well. The bad news is that it is not currently being developed or supported on the Linux platform. Caldera has extensive documentation on *Wabi* at: <http://www.caldera.com/support/docs/wabil>.

TWIN is a GNU project whose goal is to write UNIX libraries that emulate Windows 32-bit APIs (Application Program Interfaces). Its two goals are to allow Windows to be compiled on UNIX (and Linux) and to allow Intel versions of UNIX to run Windows programs. Its libraries are supported on HPUX, AIX, and Solaris as well as Linux. Although it started as a separate project, the TWIN project has merged with the WINE (WINE Is Not an Emulator) project. For more information on the TWIN project, go to <http://www.willows.com>.

The WINE project's goal is to allow Windows programs to run natively on Intel versions of UNIX. WINE technically is not an emulator since it provides low-level compatibility for Windows programs running on Linux, which gives it a significant speed advantage over using emulation.

WINE currently works under Linux, FreeBSD, and Solaris, and it is included with most current Linux distributions. If you don't have a copy of WINE, you can obtain source code or RPMs at <http://www.winehq.com>.

After WINE is installed, you must configure the `wine.conf` file. The format is the same as a Windows `.ini` file. See [Appendix A](#) for details on configuration.

After WINE is configured, you can install Windows programs. Start X-Windows and open a terminal window. To install a Windows program, type `wine <Windows setup program>`. After the program is set up, you can run it from a terminal window by typing `wine <Windows executable>`.

WINE is currently in development and it has had varied success in running Windows programs. Most Windows 3x programs run without problems, but Win32 programs (Windows 9x and NT programs) have had mixed success. For example, Microsoft PowerPoint 2000 is reported to work perfectly under WINE, but Microsoft Outlook doesn't work at all. For an up-to-date list of Windows programs compatible with WINE, go to <http://www.winehq.com/Apps/query.cgi>.

Windows emulators for Linux are not currently good enough for production use, although they are rapidly improving. Keep an eye on WINE. Corel has been supporting the WINE project and is developing libraries that allow Windows applications to run better with WINE. Corel is currently using WINE to port its WordPerfect Suite to Linux. At the rate of its current development, it will eventually be good enough to run most Windows programs without a problem.

4.3 VMware

VMware is a program that allows a PC to run an OS inside another OS. It currently runs under Windows NT, Windows 2000, and Linux. This is not emulation, but a real session of another OS. The OS (the guest OS) runs as a separate session inside of the main OS (the host OS).

When VMware is started, it opens a window that shows a bootup screen. This looks like a PC booting up, complete with the BIOS setup (see [Figure 4.1](#)). It then boots into a running session of the guest OS (see [Figure 4.2](#)). This is made possible by the VMware virtual platform, which emulates the BIOS of a PC inside of a host OS. This emulation ability allows the installation of virtually any PC OS on top of this virtual platform. If you have problems with the installation, see the section "Known Problems with VMWare."

Both the Linux and Windows versions of VMware install easily. The following is a brief explanation of VMware. Complete installation instructions are available at <http://www.vmware.com/support/linux/doc/> for Linux and <http://www.vmware.com/support/win/doc/> for Windows.

Figure 4.1. VMware creates a virtual machine inside of Windows or Linux.

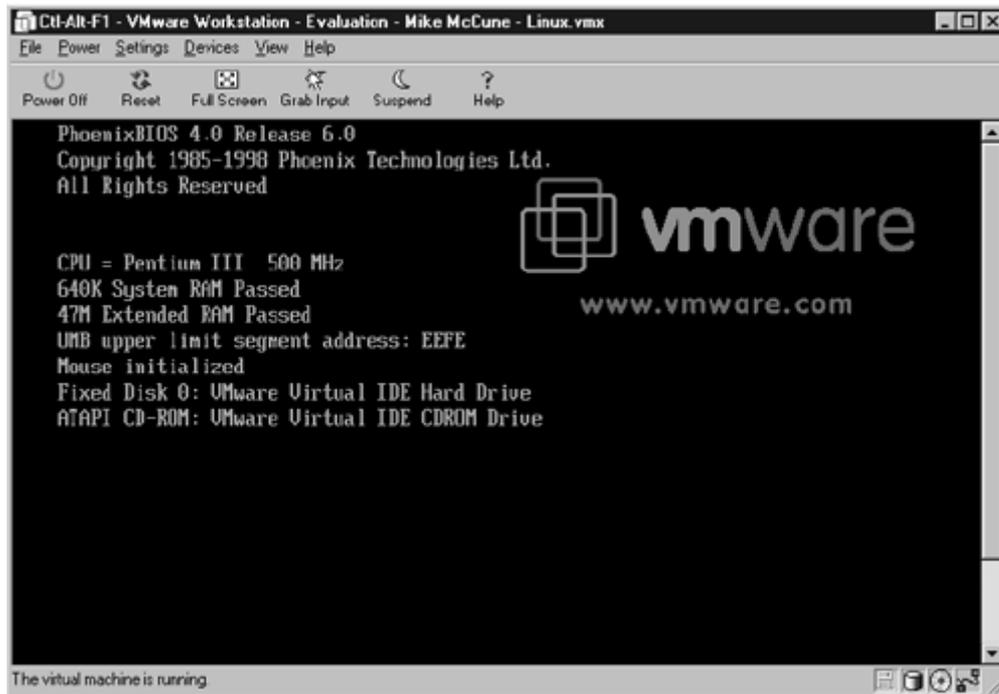
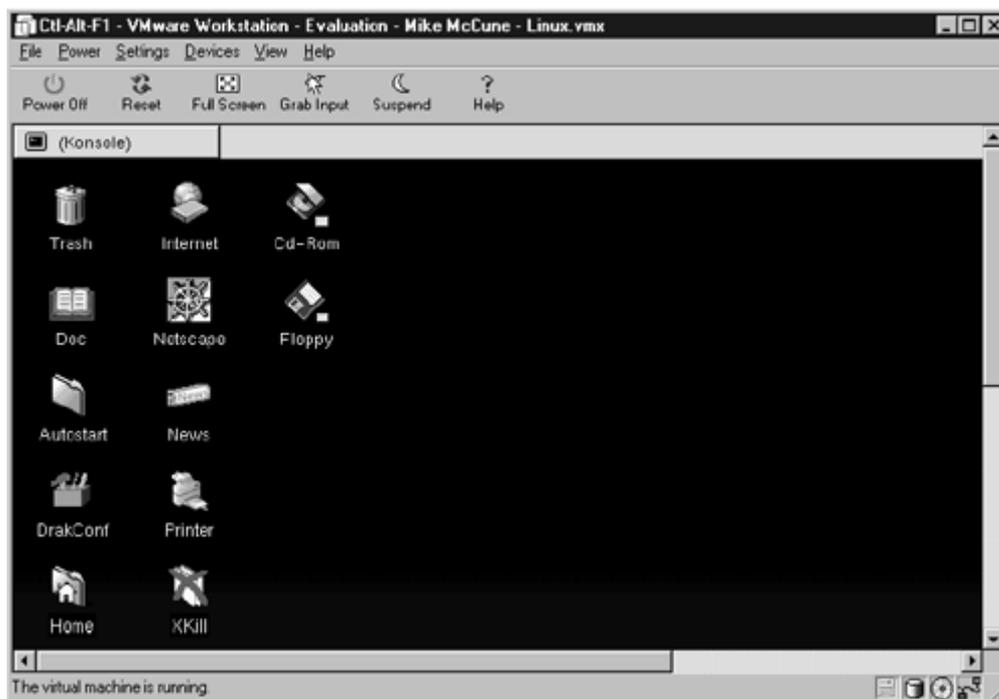


Figure 4.2. VMware is running Linux inside of a Windows NT.



For the Linux version, download the archive for Linux from www.vmware.com. Next, unpack the archive as follows, where `xxx` represents the version number of VMware:

```
tar xzf vmware-forlinux-xxx.tar.gz
```

Once it is extracted, go to the directory that is created by VMware:

```
cd vmware-distrib
```

To install, you must be logged in as administrator or equivalent. Run the installation program by typing:

```
./install.pl
```

When the installation is complete, go to the `vmware` directory and run the configuration utility. This will configure VMware for your system. There will be several questions. In most cases, the defaults should be used:

```
vmware-config.pl
```

When the configuration is complete, you will need to install the license. This is obtained from VMware via email. Once you receive this license, make a directory called `.vmware` in the `vmware` executable directory and copy the file named `license` to that directory.

The installation for Windows NT or 2000 is pretty straightforward. You must be logged in as administrator or equivalent to install the program. First, download and run the installation program. The program will have a name like `vmware-nt-xxx.exe`, where `xxx` is the version number.

To install the license for Windows, double-click on the license file that you get from <http://www.vmware.com>. This will add the following license key to your Registry:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VMware, Inc.\VMware for Windows NT\License
```

After installation, VMware must be configured. To configure VMware for Windows, run the Configuration Wizard in the `vmware` folder (under Programs). For Linux, start X-Windows then type `vmware` from a terminal prompt.

First, you must decide how much RAM and disk space to allocate to the guest OS. Remember that this RAM and disk space will be taken away from the host OS. To allow both operating systems to operate, you need a 266 MHz Pentium II, about 128 MB of total RAM, and at least 500 MB of free disk space on your system, although more is always better. This configuration allows you to allocate up to half of your RAM and 500 MB to 1000 MB of disk space to the host OS. VMware does not re-partition the hard drive; it merely creates a file on your current drive which acts as a virtual drive. Be sure to allow some free space since VMware requires about 50 MB or more of swap space.

Next, if the guest OS is Linux or Windows, install the appropriate VMware tools for that OS. VMware tools for Windows and Linux are included with VMware. These can be accessed by going to Settings -> VMware Tools Install. These and tools for others OSs can also be downloaded from <http://www.vmware.com>. The compressed file is small enough to fit on a floppy, or the tools can be downloaded directly into the guest OS.

To install the tools for Linux, unpack the tools with the following command:

```
tar xzf vmware-toolsxxx.tar.gz (xxx is the version number)
```

Then go to the directory `vmware-tools` and run `./install.pl`.

For Windows, simply run the downloaded file `vmware-toolsxxx.exe`, where `xxx` is the version number.

When installing the tools, be sure to install the toolbox and video drivers. The VMware tools will install their own video driver, which increases the graphics speed of the guest OS. They will also set the video settings of the guest OS to be the same as the host OS. For example, if the host OS is set to 800x600 at 16 million colors, the guest OS will be set to the same. Without the VMware tools, the video settings of the host OS will only support 640x480 at 16 colors.

4.3.1 How Well Does it Work?

VMware works very well as long as your system meets at least the minimum requirements listed above. It allows you to open another OS inside your Windows NT/2000 or Linux system. The host OS runs at almost normal speed and the guest OS runs at about one-half to one-quarter normal speed, depending on what you are doing. Both operating systems share the devices on the PC. The applications in each OS can share data and can even cut and paste between the host and guest operating systems.

4.3.2 Known Problems with VMware

As with any program, not everything works as it is supposed to. These deviations are called bugs by most software companies, although a few companies still insist on calling them features. VMware is a rapidly evolving product with a new version coming out every few months. The VMware support page is at <http://www.vmware.com/support/>. If you don't find the solution to your problem here, fill out an incident report at <http://www.vmware.com/forms/Incident-login.cfm>. There is also a handy search form on this page that they encourage you to use before filing an incident report.

4.3.3 Tuning VMware

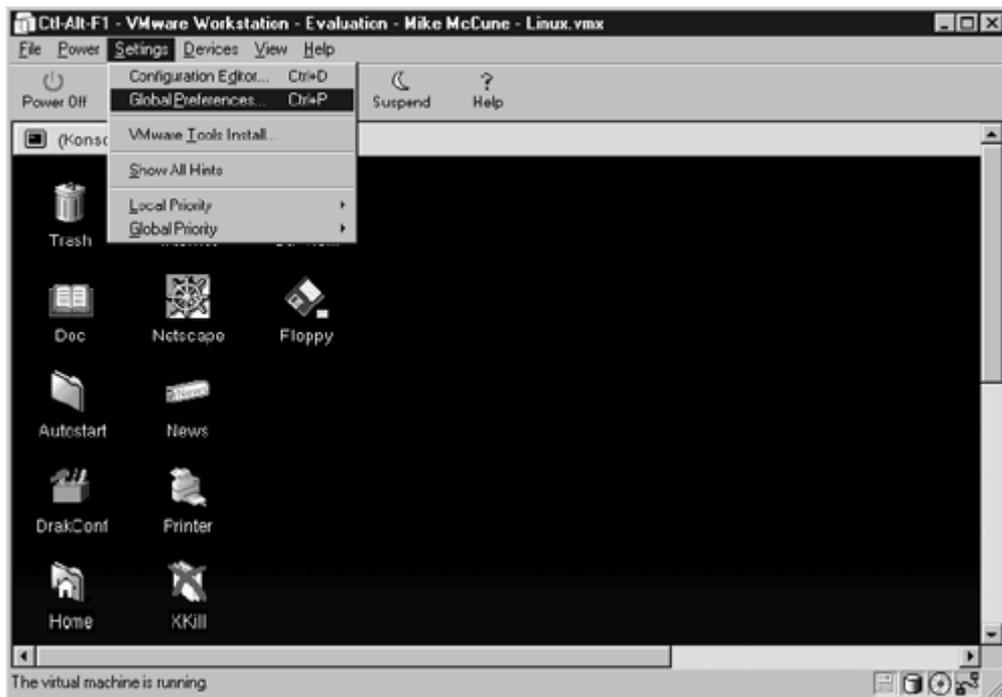
Many of the tuning parameters are mentioned in the installation and configuration discussions above. If the host or guest OS is sluggish, check the configuration parameters first. Make sure that there is enough memory and disk space allocated to each OS. If an OS doesn't have enough memory or disk space, reconfiguring your system may help. If there still isn't enough disk space or memory to run both operating systems properly, consider upgrading your PC.

Other than that, installing the VMware toolkit can have a great effect on the speed of the guest OS. The toolkit has enhanced video drivers designed for the VMware guest session. There is a VMware tools version for Windows 95, NT, and 2000, Linux, and FreeBSD (a version of UNIX).

If Windows is your guest OS, you can get better video performance if you use direct draw mode (not to be confused with Microsoft's Direct Draw). The default for VMware is to use GDI, which is the video system that is used by Windows. Direct draw is faster, but not supported by all systems. If you have video problems after installing direct draw, change back to GDI. To switch to direct draw mode, go to Settings -> Current VM Graphics and change the GDI to Direct Draw.

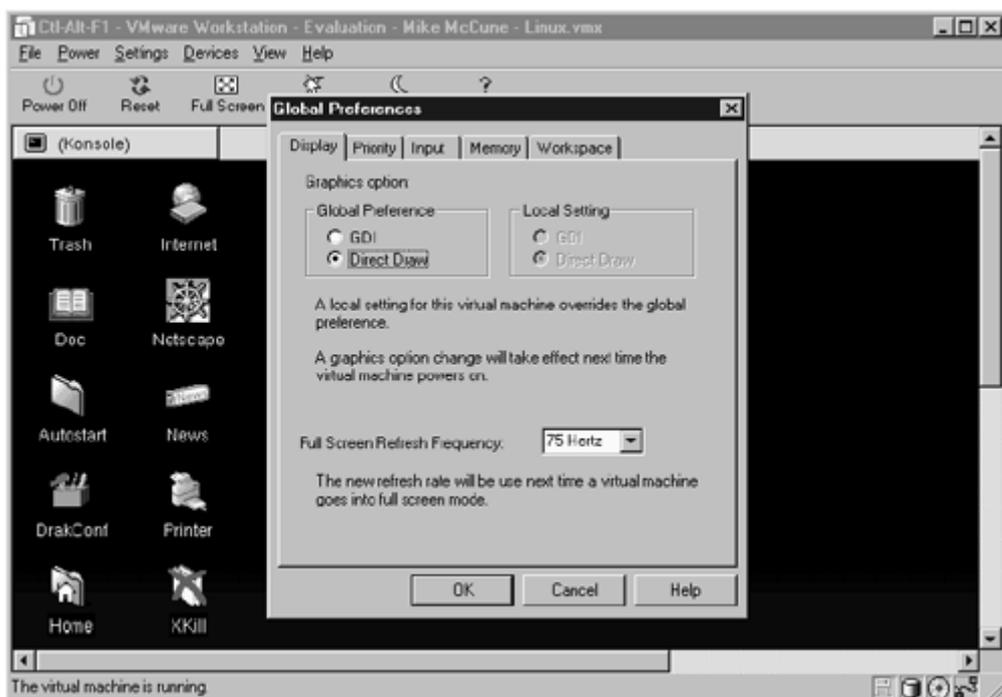
The guest OS is also more responsive in the full-screen mode. To open it in full-screen mode, click on the Full Screen command on the menu. If your screen appears blank or distorted in full-screen mode, check your refresh rate. VMware defaults to a 75 MHz scan rate and some systems won't support it. If this is the case, you can change the scan rate by going to the Settings -> Global Preferences menu as shown in [Figure 4.3](#).

Figure 4.3. Getting to the Global Preferences Screen.



You can toggle back from full-screen to windowed mode by pressing <CTRL>-<ALT>-<ESC>. If you are using direct draw mode, however, your cursor may turn into a black block when you toggle. If this is a problem, you can switch back to GDI mode by using the Display tab on Global Preferences (see [Figure 4.4](#)).

Figure 4.4. Setting the Graphics option to Direct Draw in the Global Preferences Screen.



You can also adjust the priority of the guest OS with the Priority tab (see [Figure 4.5](#)). The priority can be turned down if the guest OS is using too many resources, or it can be turned up to increase the responsiveness of the guest OS.

Figure 4.5. Adjusting the guest operating system's priority in Global Preferences.



Another thing to watch for: If you have a DOS session running in VMware, it will use up CPU cycles even when it is idle. This can be fixed by installing the program `CpuIdle`, which can be downloaded from <http://www.bugcomputer.com/cpuidle/>. Install it according to the instructions that come with it, except don't use the `-cpu` option. When this option is enabled, it can cause problems on some systems.

Also, make sure that CD-ROM autoplay is disabled in both the guest and host OS. Autoplay constantly polls the CD drive for autoplay files, which dramatically slows down performance.

4.4 FreeMWare

FreeMWare is an open source project that allows a Windows session to be opened inside of Linux. While it is early beta right now, it promises to offer an open source alternative to VMware. It is also licensed under the LGPL (LesserGNU Public License). The main difference between LGPL and GPL is that the LGPL is less restrictive on its use or redistribution. For more information on LPGL, go to <http://www.gnu.org/copyleft/lesser.html>. For more information on FreeMWare, go to <http://www.freemware.org>.

4.5 Win4Lin

Win4Lin is the "new kid on the block." It allows Linux users to run a Windows session in Linux, but not the other way around. It is faster than VMware, runs with much less memory (it can easily run with 32 MB of RAM), and sells for about \$40.

I've had the chance to play with the early version of it, and it seems faster, but it isn't as stable as VMware. In its defense, I reviewed a beta version, so these bugs should be worked out over time.

4.6 Conclusion

Currently, the best way to run Windows programs on an Intel Linux platform is with VMware. Most Windows programs will run VMware without a problem. The downside is that it is closed source, commercial software. It also requires you to buy a copy of

Windows to run Windows programs. This may not be a problem since Windows is bundled with most PCs, however.

FreeMWare is a promising open source project that is currently not developed enough for normal use.

The WINE project currently only runs a limited number of Windows programs. It is, however, free and open source (under the GPL) and is progressing rapidly.

Chapter 5. Internet Applications

The Internet has grown from a communications network for government agencies in the late 1960s to the engine that is driving the U.S. economy in the late 1990s. The different strengths of Linux and Windows make it not only possible to mix environments, but desirable.

Due to its UNIX heritage, Linux's native TCP/IP protocol support, stability, and low cost make it ideal as a server. According to a count of servers by leb.net (<http://leb.net/hzo/ioscount/>), Linux is the most popular server on the Internet. Linux also runs Apache, which is by far the most popular Web server on the Internet according to a Netcraft survey (<http://www.netcraft.com/survey/>).

Although Linux is starting to catch up, Windows still has a lot more end-user applications than any other OS. Many applications such as Microsoft Office, Media Player, Out-look, and Exchange are not available for Linux. According to International Data Corporation (<http://www.idc.com/>), Microsoft Windows has about 87% of the desktop OS market as compared to about 4% for Linux. This gives tremendous incentive for companies to release the Windows version of desktop programs first.

This chapter is not going to cover every Internet application or go into details about the configuration of any of them. Those subjects would take much more than one book to cover. Instead, I intend to cover some of the common interoperability issues in dealing with mixed environments on the Internet.

5.1 Web Server Compatibility

Let's start with the Web since it seems to be getting the most press lately. Web servers are relatively simple. They take text files (HTML files) and transfer them using the HTTP protocol. These text files tell the Internet browser how to display a page and where to download the graphics, sounds, programs, and other files needed to render the page. The three most popular Web servers are Apache, Internet Information Server (IIS), and Netscape, which together make up about 85% of the server market.

Apache is a free, open source Web server that, although there is a Windows version available, is run primarily on UNIX platforms such as Linux. IIS is included with Windows NT and 2000 Server and is available only for Windows. Netscape is available in versions for Windows and most flavors of UNIX, including Linux.

Since most Web standards are set by the Internet Engineering Task Force (IETF), Web servers tend to look the same to end-users. The two main problems of compatibility tend to be in connecting to databases and generating dynamic Web pages.

Most large Internet sites today store their data in a database and use this data to generate Web pages on the fly. For information on interoperability between Linux- and Windows- based databases, see [Chapter 7](#).

The traditional way of creating dynamic Web pages is to write CGI scripts. Today, there are several programs that make generating dynamic Web pages easier. Two of the most popular of these tools are PHP and Active Server Pages (ASP). PHP is an open source program designed to run with UNIX and Apache. ASP is a Microsoft product designed to run with Windows and IIS. Both of these programs are available in both UNIX and Windows versions.

5.1.1 PHP

Both the Linux and Windows versions of PHP are available at <http://www.php.net/>. They are available in both source code and binary versions. The Windows binaries are up-to-date, but if you want the latest Linux code, you may have to compile it yourself.

There are some special considerations when connecting PHP to a Microsoft database. On a Windows machine, simply use the ODBC drivers. If you need instructions on setting up ODBC, go to <http://www.php.net/manual/config-odbc.html>.

It is a little more complicated on a Linux machine. To connect to a Microsoft SQL server, use the PHP Sybase modules at <http://www.php.net/extra/ctlib-linux-elf.tar.gz>. SQL Server and Sybase are mostly protocol-compatible.

It's not simple to connect to a Microsoft Access database from a Linux machine either. The easiest and most robust method would be to use Open Database Connectivity (ODBC) drivers. See the section in the databases chapter ([Chapter 7](#)) on ODBC for information on using ODBC with Linux.

5.1.2 Active Server Pages

Active Server Pages (ASP) is included with Microsoft IIS. One use for a Linux machine with ASP is as a back-end database. ASP uses ODBC for database support, and the ODBC drivers from OpenLink can be used to connect to the Linux database for a Microsoft machine running ASP.

IASP by Halcyon Software replicates the ASP functions on Linux and many other platforms. It adds many functions to ASP such as Java Server Pages and full Java support. It can even be used with Microsoft's ASP to add extra functions. There is a free development version and the full version starts at \$495 (<http://www.halcyonsoftware.com>).

Chili!Soft has a beta version of ASP for Linux that is also worth a look. Its price is expected to be \$995 per server (<http://www.chilisoft.com/>).

Another option is to convert ASP pages to PHP and use PHP on both platforms. Considering that PHP is free, this might be a good option if cost is a consideration. There is a free program called ASP2PHP that converts ASP pages to PHP pages. The basic syntax of the program is `asp2php myaspfile.asp`, where `myaspfile.asp` is the ASP file to convert. It also comes with a nice GUI tool, `gtkasp2php`, which makes conversion as easy as point and click. The program is available at <http://home.i1.net/~naken/asp2php/>.

5.2 FrontPage Extensions

FrontPage is a popular Microsoft Web authoring tool. FrontPage Extensions are needed on a Web server to take advantage of FrontPage features such as:

- Letting multiple users simultaneously collaborate on the same Web site and Web server (multi-user authoring).
- Letting users write directly to the Web server with Microsoft FrontPage using a PC or laptop computer from anywhere in the world via the Internet (remote authoring).
- Letting users include forms on their Web site and specify how the results of those forms are handled without the users having to write their own scripts.
- Letting users include discussion webs on their Web site.
- Providing full text search capability on a Web site.
- Letting users include hit counters on their Web site.

Of course, all of these functions can be done without FrontPage Extensions. FrontPage Extensions are only required to take advantage of these functions in Microsoft FrontPage.

Microsoft's IIS automatically installs FrontPage Extensions. Apache doesn't include FrontPage Extensions, but they can be downloaded from <http://www.rtr.com/fpsupport/download.htm>. Instructions for installing FrontPage Extensions on Apache for Linux

are available at <http://www.e-gineer.com/instructions/install-frontpage-extensions-for-apache-on-linux.phtml>.

5.3 Using Microsoft Office Files on the Web

Microsoft Office programs are often used to create Web documents. All the standard Office programs (Word, Excel, and PowerPoint) have a Save As HTML option. This is convenient for quickly creating HTML documents, but Office tends to put non-standard characters in the HTML. These documents then look bad when viewed with a non-Microsoft browser.

HTML documents created with Office have the following non-standard characters:

- The semicolon is missing at the end of numeric character escapes (e.g., `=`).
- Numeric renderings of special characters (e.g., `< > &`).
- Unquoted `<table>` tags containing non-alphanumeric characters.
- PowerPoint mis-nests `` and `` tags.
- PowerPoint uses `` and `` tags to accomplish paragraph breaks.
- Office misses `<tr>` tags in text-only slides.
- Office places extra `</p>` tags.
- Office places unmatched `` tags in headings.
- Office uses "paragraph-long lines," which leads to lines that scroll off the page.

Fortunately, there is a program called demoroniser (<http://www.fourmilab.ch/webtools/demoroniser/>) that will correct these flaws in HTML documents created with Office. It is a freeware Perl script that converts the non-standard characters into standard HTML tags. Perl is installed by default on most Linux distributions. A Win32 version of Perl is available at <http://www.ActiveState.com/ActivePerl/download.htm>). If you have Windows 95, you will also need DCOM (<http://www.microsoft.com/com/resources/downloads.asp>).

The syntax of demoroniser is:

```
demoroniser [ -u ] [ -wcols ] [ infile ] [ outfile ]
```

where:

- `-u`-Help.
- `-w`-Sets the character length of the lines. The default is 72.

5.4 Web Browsers

One advantage of the Web is that the Web page tells the browser what a page should look like and the browser actually renders the Web page. This means that Web pages download fast and there are Web browsers available for every OS. It also means that the browser choice makes a big difference in how Web pages look and not every browser will properly render every Web page.

5.4.1 Browsers Available for Both Linux and Windows

The browser choice for cross-platform compatibility is currently Netscape. It has versions for most major platforms and supports all the major standards, including frames, image maps, plugins, SSL (Secure Sockets Layer), Java, and Java Script. It is included in most Linux distributions and a Windows version can be obtained at <http://www.netscape.com>.

Star Office has a built-in browser that works very well. It renders pages well, but lacks some of the advanced features such as SSL, interlaced GIFs, and Java. See the section on productivity applications for more on Star Office (<http://www.sun.com/staroffice>).

There are a couple of other browsers that deserve a mention: Lynx and Arachne. Lynx is one of the oldest browsers and it is available for almost every OS. Even though it is a text-only browser, it is useful for debugging Web pages. It is very picky about HTML syntax. It is also good for checking to see if Web pages have the proper alt tags that are necessary for text readers used for handicapped access to Web pages (<http://lynx.browser.org/>).

Arachne deserves mention because it is the only graphical DOS browser. It can be run under DOSEmu under Linux and in a DOS window under Windows. It is a small, compact browser that can run on an 8086 (the original IBM PC) with 640KB of memory. There is also a Linux version in the works (<http://home.arachne.cz/>).

5.4.2 Windows Browsers

Since about 90% of all desktops run Windows, there are several browsers that are available for Windows but not Linux. The most popular Web browser for Windows is Microsoft's own Internet Explorer (IE). It supports all the major standards as well as some Microsoft-specific functions such as ActiveX and VBscript, which offer similar functionality to Java and Java Script. The main difference is that ActiveX and VBscript are specific to the Windows platform, whereas Java and Java Script are cross-platform.

Opera is another popular browser for Windows. It supports all the major standards and is still small and fast. It will run on a 386 with 8 MB of RAM. There is also a Linux version in the works. Unlike other browsers, Opera is not free; it costs \$35 (<http://www.opera.com/>).

5.4.3 Up-and-coming Browsers

There are two open source browsers that are poised to give the other browsers a run for their money: KFM and Mozilla. KFM is the file manager for the graphical interface KDE. Just type in the URL in the location window of KFM and it will load a Web page. It doesn't support all of the major standards, but it is being rapidly improved (<http://www.kde.org>).

Mozilla is the open source version of Netscape's browser. It was originally started when Netscape opened up the source code to its browser in early 1998. What started as bug fixes turned into a complete rewrite of the Netscape rendering engine. Mozilla is noticeably faster than the Netscape browser and it is becoming more stable as development continues. It is scheduled to eventually replace the Netscape browser when it becomes stable enough (<http://www.mozilla.org>).

5.4.4 Browser Plugins

Browser plugins are programs that add extra functions such as sound and animation to browsers. Almost all plugins are available for the Windows versions of Internet Explorer and Netscape. The choice of plugins for Linux is much more limited. Currently, Netscape is the only Linux browser that supports plugins. The plugins supported by Netscape for Linux are:

- Real Player—Plays streaming audio and video. See the section on streaming video later in this chapter.
- Cult3D—A 3-D viewer.
- DjVu—A document viewer.
- Flash Player—A viewer for vector graphics and animation.
- Gig—Renders data-driven, interactive graphics.
- Hypercosm3D Player—Renders 3-D computer graphics.

- MpegTV—Plays streaming MPEG video.
- Plugger—Displays inline pictures, sound, and video.
- TANGO Interactive—An interactive, multimedia, collaborative tool.
- Tcl/Tk—A scripting tool.
- Ump—Plays MIDI audio files.
- X11R6.3 Remote Execution—Allows an X11 application to be embedded into a Web page.
- XVIEW—A picture viewer.

Current versions of Netscape plugins can be found at http://www.netscape.com/plugins/search_pi.html.

There is also an Xswallow plugin for the Linux version of Netscape that supports many data types, including:

- vrml1 and vrml2—Popular 3-D formats.
- midi—A music format.
- sun audio—An audio format.
- mpeg—A video and audio format.
- avi—A video format.

Xswallow is available at <http://www.csn.ul.ie/~caolan/docs/Xswallow.html>.

5.5 Email

Email is the original killer application on the Internet. Until the Web took off, email took up most of the Internet's bandwidth. There are several Internet email standards. A brief overview of them would include:

- Mail Transport Protocol:
 - SMTP (Simple Mail Transfer Protocol)-Allows transfer of mail between mail servers.
- Remote Mail Retrieval Protocols:
 - POP3 (Post Office Protocol version 3)-Downloads full messages to the client.
 - IMAP4 (Internet Message Access Protocol version 4)-Allows the actual messages to be stored on a server and only the headers are downloaded.
- Message Formats:
 - RCF822-This is the standard format for ASCII (text) messages.
 - MIME (Multipurpose Internet Mail Extensions)-Allows multi-part messages and non-ASCII (non-text) attachments such as sounds and graphics.

These are very simplified explanations of the major Internet mail standards. A complete list of the standards is available at <http://www.imc.org/rfc.html>.

5.5.1 Mail Clients

There are several popular mail clients that are available for both Linux and Windows. The simplest is PINE (Program for Internet News and Email, or PINE, is not Elm, depending on who you ask).

PINE is a command-line mail program, but it supports all the major Internet mail standards listed above. It also has advanced features such as a message digest, address book, spell checker, and a news reader. Most Linux distributions include a copy of PINE. The Windows version and more information on PINE are at <http://www.washington.edu/pine/>.

Netscape also has a graphical mail program bundled with its Netscape Communicator, which is bundled with most distributions of Linux. More information on it and the Windows version are available at <http://www.netscape.com>.

Star Office also includes a mail program. Both the Linux and Windows versions are available from Sun (<http://www.sun.com/staroffice>). A word of warning: This is a large download.

Since it is the default mail program for Windows, Outlook deserves a mention also. It comes in two flavors: Outlook and Outlook Express. The main difference between the two is that Outlook has some extra features such as the ability to manage multiple mailboxes.

Both versions of Outlook support all the standard Internet mail features as well as the features of Microsoft's Exchange Server. Outlook can be used as a client for any standard mail server or the Microsoft Exchange Server. The Exchange Server uses proprietary protocols to manage mail, calendaring, and global address listings. The Exchange Server is discussed in the next section on mail servers. For more information, refer to <http://www.microsoft.com>.

5.5.2 Mail Servers

There are two types of email in wide use today: Internet mail and collaborative mail. Internet mail supports MSTP, POP3, IMAP, and MIME. Collaborative mail supports all the Internet mail standards plus added features such as shared files, directories, and shared calendars.

Internet Mail Servers. Sendmail is the most used Internet Mail Transport (MTA) on the Internet. It is popular because it is so configurable, but it is also difficult to set up because it is so configurable. It is open source and available for Linux, Windows, and many other operating systems. There is also a commercial version that includes support and configuration tools to make the job easier (see <http://www.sendmail.org>).

While it would take an entire book to cover the configuration of Sendmail, there are some ways to ease the pain of configuration. The main configuration file for Sendmail is `/etc/sendmail.cf`. Writing this file from scratch would be next to impossible. It is best to find a sample `sendmail.cf` that is close to your needs and edit the file. There is also a macro file, `sendmail.mc`, included with Sendmail that can be used to create the `send-mail.cf`. Once the `sendmail.mc` is edited to your needs, create a `sendmail.cf` with the following command:

```
m4 /etc/sendmail.mc > /etc/sendmail.cf
```

Many of the features of Sendmail are not even used by most email users. For example, Sendmail has settings for uucp and BITnet, which are older Internet mail systems that aren't used anymore. Extensive documentation on Sendmail is available on their Web site <http://www.sendmail.org>.

Qmail is another popular MTA. It is much easier to configure than Sendmail, but with that ease comes less maturity. Sendmail has been in use since 1979, while Qmail 1.0 was released in 1997. Like Sendmail, it

is free and open source. For the program and information on Qmail, go to <http://www.qmail.org/>.

Postfix is another popular MTA. Many users swear by its speed and security (<http://www.postfix.cs.uu.nl/>).

These are just a few of the many MTAs available. It is not surprising to see that there are so many MTAs since email is one of the oldest uses for the Internet.

Collaborative Mail. The two most popular collaborative mail programs are Microsoft Exchange and Lotus Notes. Most large companies use these programs to not only send email but to keep a central email directory, share files, and share calendars. The ability to interact with these programs is critical if Linux is to be used in a corporate environment. While there is the open standard Lightweight Directory Access Protocol (LDAP), neither of these two collaborative mail programs use it.

Microsoft Exchange

Microsoft Exchange currently has a slight lead in popularity over Lotus Notes. Exchange consists of two components: the Exchange Server and the Outlook client. Currently, the Exchange Server only runs on Windows NT and 2000. The Outlook client runs on all current versions of Windows.

There are programs that allow Linux to interact with an Exchange environment. On the server side, Hewlett Packard's OpenMail provides a lower cost replacement for Microsoft Exchange Server. It does require that the OpenMail client be installed on all client machines, but once this is done, the OpenMail server looks just like Exchange Server to the Outlook clients. More information on HP OpenMail is available at <http://www.ice.hp.com/cyc/om/00/index.html>.

TradeMail is a program that allows Linux machines to act as Outlook clients. It consists of server-side components that go on the Exchange server and a client on the Linux machine. TradeMail works but is not currently ready for a production environment. For more information on TradeMail, go to <http://cobra.bynari.net/>.

Lotus Notes

Lotus Notes was the first collaborative email program. Like Outlook, it consists of a server and a client. The Notes server is available for both Linux and Windows NT/2000. The Windows server has been out for a while, but the recent addition of the Linux server offers a Notes server without the expense of buying Windows (<http://www.lotus.com/home.nsf/welcome/notes> and <http://www.notes.net/linux>).

There is a Notes client for Windows, but none for Linux. Lotus has said it has no plans to develop a client for Linux. If you need a Notes client for Linux, some people have had luck with running the Windows client with WINE (<http://www.brooklinesw.com/linux/linuxnotes.html>).

5.6 Streaming Media

Streaming media allows live or recorded audio and video to be used on the Internet. While it is relatively new technology, it will become more important as broadband Internet connections become more common.

5.6.1 Streaming Video

Real Networks. Right now, the only complete video streaming package for Linux is Real Networks (<http://www.realnetworks.com/index.html>). Real offers encoding tools (Real Producer and Real Presenter), streaming servers (Real Server), and players (Real Player).

Producing Content with Real

Real Producer Basic is available for Linux, Macintosh, and Windows. It takes live video or audio and converts it to Real Media format. This can then be streamed or downloaded and played with the Real Player. Real Producer Plus has additional features that allow better optimization of Real Media files. There is also Real Producer Pro, which allows the conversion of other audio and video formats into Real format. Unfortunately, Real Producer Pro is only available for Windows. The price of Real Producer Basic is free, and the Pro version sells for up to \$499.

Real Presenter Basic adds embedded audio and video to Microsoft PowerPoint presentations. Obviously, it is available only for Windows. The Basic version is free, but there is also a Plus version that sells for \$69.

Real Servers

Real Server is available for the Macintosh, Windows NT/2000, and a wide variety of UNIX platforms. There are two ways to stream content with Real: HTTP and Real Time Protocol (RTP). For low-volume streaming, simply link a Real file to a Web page and allow HTTP to stream the file. RTP is an IETF standard for data streaming and it is much better suited for streaming media than HTTP. All the details of RTP are available at <http://www.cs.columbia.edu/~hgs/rtp/>.

For higher volume streaming, use Real Server. The free basic player allows up to 25 streams to run simultaneously. Beyond that, the cost of the server is based on the maximum number of streams. For example, a 60-stream server would cost \$1995. The commercial version also includes the ability to password-protect streams and a license to resell the streaming services.

Real Server requires 6 MB of RAM to run plus 40KB for each simultaneous stream in addition to the memory required for the OS. Streaming is not very processor-intensive. A Pentium 120 can easily handle 400 simultaneous 28.8-kps streams.

Bandwidth is the real problem. The preceding example of 400 28.8-kps streams would require over 11.5 mps of bandwidth just for the Real streams. To maintain those 400 streams would require a connection of about 15 mps of bandwidth. Considering that a T1 connection delivers 1.5 mbs at a cost of about \$500 to \$1000 per month, bandwidth can quickly get very expensive.

Real Players

Real Player Basic is available for both Linux and Windows. It allows the end-user to view Real Media files and streams. It also supports many other audio and video formats, including MP3. There is also a Plus version for \$29 that has extra features such as a graphic equalizer and the ability to record streams. The \$49 Real Player Pro also has the ability to create slide shows with video and audio.

Real Player for Windows is better supported than the player for Linux. The current version of Real Player for Windows has been out for months, while the Linux still has a beta version of the previous version of Real Player. This should change now that Real Networks has announced an agreement with Red Hat to bundle its Real Server with Red Hat Linux.

QuickTime by Apple. There are several streaming media servers for Linux that are based on Apple's QuickTime Server. The Apple QuickTime Server is available for the Macintosh and several versions of UNIX (

<http://www.publicsource.apple.com/projects/streaming/>). Since Apple doesn't charge by the stream, it is a much cheaper solution than Real Networks. On the down side, the Linux players for QuickTime files are currently works in progress.

There is also a streaming server called PRISS (Portable RTSP Internet Streaming Server), which is an open source project based on Apple's source code (<http://www.streamingserver.org/priss.html>). This is currently the only open source project for streaming video.

QuickTime players are currently available for MacIntosh and Windows only (<http://www.apple.com/quicktime/download/>). The free QuickTime player allows the playback of QuickTime files and streams. The \$29 QuickTime Pro also adds the ability to create and edit QuickTime files.

There are some efforts to create Linux video players that support QuickTime. The main problem is the CODECs. A CODEC is used to create audio and video files. The problem is that all the CODECs are patented and nobody is willing to pay to use a CODEC for a free Linux player. Nevertheless, there are several Linux players in development.

The best current video player for Linux is xanim, which is included with most Linux distributions. It supports a large number of CODECs, but the problem is that the most common CODECs used by QuickTime are not supported. A list of supported CODECs is available at http://xanim.va.pubnix.com/xa_whatisit.html#format_support.

There is also an open source QuickTime for Linux, Xmovie and Broadcast 2000, which can create and play QuickTime movies using some limited CODECS. More information and downloads are available at <http://heroine.linuxbox.com/quicktime.html>.

Windows Media Player. Windows Media really isn't really a cross-platform solution. The encoding tools, server, and player are all Windows-based and not available for Linux. The player is available for MacIntosh and Windows, though. More information on Windows Media is available at <http://www.microsoft.com/windows/windowsmedia/en/default.asp>, and the tools, server, and player are available at <http://www.microsoft.com/windows/windowsmedia/en/download/default.asp>.

5.6.2 Streaming Audio

There is a wider range of available choices for streaming audio, which works better over a standard dial-up connection. All the streaming video players listed above support streaming audio. There are also several audio-only streaming programs. The most popular audio format on the Web is MP3.

MP3. MP3 started as a way to trade music on the Internet. The main appeal of MP3 was its ability to easily create files from audio CDs that were close to the same quality as the original CD tracks, but were about one-tenth the size of the CD tracks. A typical MP3 song was between one to two megabytes, which could be downloaded over a standard dial-up connection in about 20 to 30 minutes.

MP3 became popular for audio streaming with the introduction of Shoutcast. Shoutcast is a set of tools that allows Windows machines to encode and stream audio using the MP3 format.

There are several components of the Shoutcast package. Shoutcast is a set of plugins for Winamp, which is a popular MP3 player for Windows (<http://www.winamp.com>). For example, the DSP plugin allows Winamp to connect to a Shoutcast server and send compressed audio data to the server (<http://scastweb2-qfe0.spinner.com/download/scast/dsp150b2.exe>).

Shoutcast doesn't come with compression tools. This is because the MP3 CODEC is patented. Microsoft offers free Netshow compression tools that can handle MP3 compression

(<http://mskyus.www.conxion.com/msdownload/netshow/3.01/x86/en/nstools.exe>).

To encode live audio (as opposed to audio files), you will need the Shoutcast Live plugin (http://www.shoutcast.com/download/in_lrec.zip). Another useful plugin is the null output plugin (http://www.shoutcast.com/download/out_null.zip), which allows you to stream audio without a sound card. This is often the case when files are being broadcast instead of being live audio. It is also useful if you don't want to hear the broadcast out of your PC speakers.

The Shoutcast Server is available for all current versions of Windows and many versions of UNIX (<http://www.shoutcast.com/download/files.phtml>). This can be a separate machine or, with the Windows version, the encoding tools and server can be on the same machine.

There is also the open source MP3 server Icecast (<http://www.icecast.org>). It is available for most versions of UNIX and all current versions of Windows. Icecast will stream pre-encoded MP3 files, and there is also a Liveice program for encoding live audio streams.

MP3 Players

One advantage of MP3 is the abundance of players available. Most audio and video players such as Real Player and Windows Media Player support MP3. The Windows Media Player is bundled with all current versions of Windows. Some of the other players for Windows are:

- Winamp (<http://www.winamp.com/>) -Probably the most popular player for Windows.
- Sonique (<http://www.sonique.com/>).
- K-Jofol (<http://www.kjofol.com/>).

For Linux there are:

- xmms (<http://www.xmms.org/>)-This is an X11-based player that is bundled with most Linux distributions.
- Mpg123 (<http://mpg.123.org/>)-This command-line player is included in most Linux distributions. It works well on low-end systems since it doesn't require X11. It also supports proxy servers.
- Freeamp (<http://www.freeamp.org/>)-This player is available for both Linux and Windows.
- Xaudio (<http://www.xaudio.com/>)-Another player available for both Linux and Windows.

MP3 is not free, however. Freunhofer Institute in Germany owns the patents on MP3 and they are starting to charge for its usage. While there are no charges for free players, encoders and commercial players pay a minimum of \$15,000 a year for using the MP3 CODECs. They are also planning to start charging for streaming MP3s as well (<http://www.mp3licensing.com/royalty/index.html>).

Free CODECs are being developed, though. The open source project Vorbis is developing free CODECs for audio and video. They already have working audio encoders and decoders for both Linux and Windows (<http://www.vorbis.com>).

5.7 Chat

Chat allows users to log into a server and post messages to each other. The granddaddy of chat programs is Internet Relay Chat (IRC). Like most other Internet programs, IRC requires a client and a server.

There are many server programs available for both Linux and Windows. Look on <http://www.tucows.com> or <http://www.davecentral.com> for Windows servers. Linux servers can be found on <http://www.linux.tucows.com>, <http://linux.davecentral.com>, or <http://www.freshmeat.net>.

One of the most popular IRC servers for UNIX is IRCd, which runs as a background process (daemon). This program consists of the IRC daemon `ircd` and the authentication daemon `iauth`. These are configured with the files `ircd.conf` and `iauth.conf`, respectively.

IRC servers for Windows aren't nearly as numerous. A good one for Windows is ConferenceRoom. Unlike most UNIX IRC servers, it is shareware. You can download a 30-day evaluation copy from <http://www.webmaster.com/>. After 30 days, it will cost from \$99 for the Personal Edition to about \$4,300 for the Enterprise Edition. In its defense, ConferenceRoom is very robust, full-featured, and easy to configure. It is also available for most versions of UNIX.

Configuring an IRC server is beyond the scope of this book. Most people would not set up an IRC server, but would instead use one of the many IRC servers available on the Internet.

An IRC client is used to attach to an IRC server. There are many IRC clients available for Linux or Windows. They are available at the same Web sites as those listed for IRC servers.

Some popular IRC clients for Windows are mIRC (<http://www.mirc.com/mirc.html>), PIRCH (<http://www.pirchat.com/>), and Visual IRC (<http://www.megalith.co.uk/virc/>).

Linux has both text-only and graphical clients. There are usually several IRC clients bundled with most Linux distributions. Some text-only clients include Blackened (<http://www.blackened.com/blackened/>), IRCII (<http://www.eterna.com.au/ircii/>), and sirc, a small client written in the Perl scripting language (<http://www.iagora.com/~espel/sirc.html>). Some graphical clients include kvirc (<http://www.kvirc.org/>) and x-Chat (<http://xchat.linuxpower.org/>).

There are three settings needed to connect to an IRC server: your username, the server name, and a port number. Your username can be any name that isn't currently being used on the IRC server. Some IRC clients have alternative names in case yours is in use.

There are several IRC server networks. All the servers on a network are connected and share messages. This allows a user anywhere in the world to join an IRC by attaching to a server close to home. Most IRC clients already have a list of servers, so it is a matter of picking a server.

There are three large server networks: DALnet (<http://www.dal.net/>), Efnet (<http://www.efnet.net/>), and Undernet (<http://servers.undernet.org/>). While there are many other networks, these have hundreds of servers throughout the world.

The port number is listed with the server address. This typically looks like this: `us.undernet.org:6777`, where the last number is the port number. Typical ports are 6667, 6668, and 7000.

Once connected, you need to join a channel. To list the available channels, type `/LIST` in the command window. Once a channel is found, join the channel by typing `/JOIN <CHANNEL NAME>`. There are many other commands, so consult your IRC client documentation.

Another useful program is an IRC Bot, which will maintain a channel on an IRC server. Channels on IRC are typically closed when everybody leaves. A bot will maintain the channel even when it is empty. A bot can also be used to manage a channel by banishing users from a channel and by allowing password-protected channels.

Most bots are available for UNIX. Some examples of bots are Eggdrop (<http://www.eggheads.org/>) and Acid Blood (<http://www.darkice.com/site/acidblood/download.htm>) for Linux or Dancer (<http://www.contactor.se/~dast/dancer/>) for Linux or Windows.

5.8 Instant Messaging

Instant messaging is rather new to the Internet. It evolved from the popular AOL Instant Message (AIM) in the early 1990s. It works similar to email, except when you send a message, it shows up instantly on the receiving end, instead of having to be retrieved like email.

The problem with instant messaging is that there is no single standard. Among the most popular instant messaging clients are AIM (<http://www.aol.com>) ICQ (<http://www.icq.com>), Yahoo (<http://www.yahoo.com>), and Microsoft Network, or MSN (<http://www.hotmail.com>). There are Linux clients for AIM, ICQ, and Yahoo. A whole page of them can be found on Fresh Meat, which is a site listing open source software for Linux (<http://www.freshmeat.net/appindex/x11/communication.html>).

Two programs to note are Everybuddy and licq. Everybuddy allows managing several different instant messaging accounts with one program. It currently supports AIM and ICQ with support for Yahoo and MSN to be added soon (<http://www.everybuddy.com>).

licq only works with ICQ, but it supports some of the advanced features of ICQ such as multi-user chat (<http://www.licq.org>).

There is also a project to make an open source instant messaging server and client based on the XML standard (see the chapter on productivity applications for an explanation of XML). Details on the Jabber project, along with servers and clients, are available at <http://www.jabber.com/>.

5.9 Internet Security

Putting a computer on the Internet attaches it to millions of computers worldwide. In most cases, this is good, but it also opens your computer to remote attacks.

The standard load of Linux or Windows is not very secure. Linux and Windows both tend to load every background program available, which opens up avenues of attack for hostile crackers. Turning off unneeded services is a good first step in Internet security.

There are 65,535 ports on a UNIX machine. `/etc/services` is a list of most ports and what services use them. Ports 1 through 1024 are privileged ports, which means only processes running as root or equivalent are allowed to connect to them. These ports are used for common protocols such as FTP, telnet, and HTTP.

There are several ways to tell what service you have running. One way is to list the processes that are running with `ps -ef|more`. This will bring up a list similar to the following:

```
PID  TTY          STAT       TIME  COMMAND
  1  ?            S           0:03  init [5]
  2  ?            SW          0:00  [kflushd]
  3  ?            SW          0:00  [kupdate]
  4  ?            SW          0:00  [kpiod]
  5  ?            SW          0:05  [kswapd]
  6  ?            SW<        0:00  [mdrecoveryd]
341  ?            S           0:00  syslogd -m 0
351  ?            S           0:00  klogd
```

```
366 ?          S          0:00 /usr/sbin/atd
381 ?          S          0:00 crond
396 ?          SW         0:00 [lpd]
435 ?          S          0:01 xfs -port -1 -daemon
449 tty1        SW         0:00 [mingetty]
450 tty2        SW         0:00 [mingetty]
451 tty3        SW         0:00 [mingetty]
452 tty4        SW         0:00 [mingetty]
453 tty5        SW         0:00 [mingetty]
454 tty6        SW         0:00 [mingetty]
455 ?          S          0:00 kdm -nodaemon
--More--
```

The first field is the process ID PID. This is simply a number that is assigned to each process.

The second field, TTY, is the terminal number. A ? is used for an unknown terminal. This usually means it is a system or X11 process.

The third field is the status of the process. These mean the following:

- D—Uninterruptible sleep (usually IO).
- R—Runnable (on run queue).
- S—Sleeping.
- T—Traced or stopped.
- Z—A defunct ("zombie") process.
- W—Has no resident pages.
- >—High-priority process.
- N—Low-priority task.
- L—Has pages locked into memory (for real-time and custom IO).

The fourth field is the amount of time the program has been running. Notice that most of the values here are 0. The reason for this is that most programs only run for a short time then sleep until needed again, which resets the time to 0.

The fifth field is the process name.

To check for open ports, use `netstat -a|more`. This will bring up something like this:

```
Proto Recv-Q Send-Q Local Address           Foreign Address
State
tcp        0      0 localhost:1168          localhost:16001
TIME_WAIT
tcp        0      0 *:1108                  *:*CLOSE
tcp        0      0 mmccune:1100           toc-
d01.blue.aol.co:ftp ESTABLISHED
tcp        0      0 mmccune:1098           cs3.yahoo.com:5050
ESTABLISHED
tcp        0      0 mmccune:1097           cs3.yahoo.com:5050
ESTABLISHED
tcp        0      0 mmccune:1094           cs3.yahoo.com:5050
ESTABLISHED
tcp        0      0 mmccune:1093           cs3.yahoo.com:5050
ESTABLISHED
tcp        0      0 mmccune:1092           cs3.yahoo.com:5050
ESTABLISHED
```

```
tcp      0      0 mmccune:1091      cs3.yahoo.com:5050
ESTABLISHED
tcp      0      0 *:1536            *:*
LISTEN
tcp      0      0 mmccune:1089      msgr-
ns7.hotmail.c:1863 ESTABLISHED
tcp      0      0 *:6000            *:*
LISTEN
tcp      0      0 *:1024            *:*
LISTEN
tcp      0      0 *:printer         *:*
LISTEN
udp      0      0 mmccune:1024      fes-
d015.icq.aol.c:4000 ESTABLISHED
udp      0      0 *:xdmcp           *:*
raw      0      0 *:icmp            *:*
7
raw      0      0 *:tcp             *:*
7
```

The most important field is the last field. If it says `LISTEN`, that means the port is ready to accept connections. The local address field is the host name and process, or port. For example, the line `tcp 0 0 *:printer *:* LISTEN` means that the printer service is listening with the TCP protocol.

The port scanner `nmap` can be downloaded from <http://www.insecure.org/nmap/>. This can be used to scan for open ports on local and network-connected machines. For example, to scan a local machine, type `nmap localhost`. This will give a list of open ports such as:

Port	State	Service
515/tcp	open	printer
1024/tcp	open	unknown
1536/tcp	open	ampr-inter
6000/tcp	open	X11

Don't port-scan other people's machines without permission. They may think you are trying to break into their system!

There is also an online port scanner from Gibson Research Corporation. The site also contains some good general information on computer security (<http://www.grc.com>).

Once we have determined what ports we have open, we need to close the ports that are unneeded. The first place to check is the file `/etc/inetd.conf`. This file contains many of the standard UNIX services such as FTP, telnet, POP, and finger. Unless you need other computers to attach to your computer through any of these services, turn them off by putting a `#` in front of the service as follows:

```
#ftpstreamtcpnowaitroot/usr/sbin/tcpdin.ftpd -l -a
#telnet stream tcp nowait root /usr/sbin/tcpdin.telnetd
```

If you have any questions about what a service does, consult the man page. For example, to find out what FTP does, type `man ftp` at the command prompt. If you really need any of these services, Secure Shell (SSH) provides much better security. See the chapter on SSH.

Most other services are started in the directories under `/etc/rc.d`. This directory has several subdirectories:

```
[root@mmccune rc.d]# ls
init.d/  rc.firewall  rc.sysinit*  rc1.d/  rc3.d/
rc5.d/
rc*      rc.local*    rc0.d/       rc2.d/  rc4.d/
rc6.d/
```

The subdirectory used depends on the run level. To determine the run level, type `runlevel`:

```
[root@mmccune rc.d]# runlevel
N 5
```

In this example, we are using run level 5, so the subdirectory would be `rc5.d`. The contents of `rc5.d` would look something like this:

```
[root@mmccune rc5.d]# ls
K20rstatd@  K42ecdl@      K80nscd@      K99sendmail@
S40atd@
S99local@
K20usersd@  K50snmpd@     K87ypbind@    S05kudzu@
S40cron@
K20rwhod@   K52apmd@     K93portmap@   S10network@
S60lpd@
K23inet@    K55routed@   K96pcmcia@    S20random@
S75keytable@
K27ircd@    K67linuxconf@ K97smb@       S25netfs@
S85sound@
K35usb@     K73hpwebjetd@ K98nfslock@   S30syslog@
S90xfs@
```

The first letter represents Killed or Started. Only files starting with "S" are loaded.

The number represents the order in which the files are loaded. The lower the number, the sooner it is loaded. Some services must be started before others. For example, `network` must be started before `netfs`.

Last is the name of the service. If you need to know what a service does, use the `man` page again. To turn a service off, change the first letter to an "S." For example, to turn off `atd`, type:

```
mv S40atd@ K40atd@
```

Always keep the `syslog`, since it is used by the system. Also, `xfs` is used by X11, so if you use X11, keep `xfs`.

After the services have been pruned, check the system again for open ports. Once this is done, we need to get some tools to monitor and protect the system.

The first thing to do is look for updates. Most known security holes in Linux are quickly fixed, but many users neglect to download fixes. Most Linux distributions post fixes on their Web site. You might also want to subscribe to the Cert (<http://www.cert.org>) and BugNet (<http://www.bugnet.com>) mailing lists.

There are also personal firewalls for both Linux and Windows. A firewall is a device or program that restricts access from the Internet to the local network and computers. A good free firewall for Linux is the Phoenix Adaptive Firewall (<http://www.progressive-systems.com/products/phoenix>). For Windows, there is Zone Alarm

(<http://www.zonelabs.com/>). These programs are free for personal use. Commercial licenses are also available from Phoenix and Zone Labs.

Tripwire is an open source intrusion detection software. The principal behind it is rather simple. Tripwire makes a snapshot of the system settings and stores it on a floppy. Tripwire can then be run later to determine if any system settings have been changed. One thing a cracker does when they break into a system is change the system settings so that they can get back in easily. This is easily detected by Tripwire. Tripwire is available at <http://www.tripwire.org/>.

It is also a good idea to put a port scan detector on your system. One of the best is Psionic's PortSentry (<http://www.psionic.com/abacus/port Sentry/>). If PortSentry detects a port scan, it will drop the connection from the scanning machine and block it from further connections. The program is available in source code only, but it compiles without problems if you follow the installation instructions.

Once the source code is downloaded, extract the files:

```
gunzip portsentry-1.0.tar.gz
tar xvf portsentry-1.0.tar
```

Next, edit the `portsentry_config.h` to add the locations of the files PortSentry uses. The defaults are fine in most cases:

- `CONFIG_FILE`—The path to the PortSentry configuration file. The default location will be `/usr/local/psionic/portsentry/portsentry.conf`.
- `WRAPPER_HOSTS_DENY`—The path and name of the TCP Wrapper `hosts.deny` file. On most systems, the file is in `/etc/hosts.deny`.
- `SYSLOG_FACILITY`—The `syslog` facility for PortSentry to use.
- `SYSLOG_LEVEL`—The `syslog` level at which to send messages.

Then, edit `portsentry.conf`. The `README.install` explains the different settings. Add any host you wish PortSentry to ignore to `portsentry.ignore`.

Finally, compile the program:

```
make linux
make install
```

This will put the program and configuration files in `/usr/local/psionic/portsen-try`. To start PortSentry, go to this directory and type:

```
./portsentry <option>
```

The options are explained in `README.install`. You can only use one option at a time, but you can load PortSentry several times if you want to monitor both TCP and UDP ports. For example, to use advanced scanning (the most secure) on both TCP and UDP ports, type:

```
./portsentry -atcp
./portsentry -autp
```

Chapter 6. Business Applications

For our purposes, business applications are programs running on an individual PC that perform business functions. The top-selling software titles are loaded with various business applications. For instance, the top-selling Microsoft Office is a bundle of common business applications, including a word processor (Word), spreadsheet (Excel), and presentation software (PowerPoint).

Some versions of Microsoft Office contain an email and calendar program (Outlook) and a database (Access). Both of these applications are covered in other chapters.

Financial programs such as Intuit's Quicken and Microsoft Money are also essential business programs. Graphics programs such as Adobe Photoshop, Corel Draw, and Microsoft Visio round out the list of business programs.

6.1 Microsoft Office

With a 95% market share, Microsoft Office is the 800-pound gorilla of productivity suites. If you get a document, spreadsheet, or presentation from someone else, there is a good chance it is a Microsoft Office file.

With that said, for Linux to become popular on the desktop, it must have a program similar to Microsoft Office and it must also be able to read and write Office files.

Linux has several programs that are similar to Microsoft Office. Most of them are also available in Windows versions. None of them is a complete, drop-in replacement for Microsoft Office. Each program has its own strengths and weaknesses that need to be considered.

Working with Office files is a major problem. Since each Office application uses a proprietary format, the file formats have to be reverse-engineered before other programs can open them. The problem is compounded by the fact that the Office applications are tightly integrated with each other and the Windows operating system. Currently, no Linux application is 100% compatible with Office files, but some are better than others.

6.2 Corel WordPerfect Office

Although it doesn't have near the market share of Microsoft Office, Corel WordPerfect Office is still the second most popular productivity suite on the market. It has all the same functions as Microsoft Office, including word processing (WordPerfect), a spreadsheet (Quattro Pro), presentation (Presentations), calendaring (CorelCENTRAL), and Web publishing (Trellix). The deluxe version also comes with a database (Paradox), more fonts, more clip art, the game Railroad Tycoon, and a bean-filled penguin toy (<http://www.corel.com/Office2000/index.htm>).

WordPerfect Office (WP Office) for Linux comes off the same code base as the Windows version and thus both versions look almost identical. Instead of writing separate versions for Linux and Windows, Corel decided to use the WINE libraries for low-level compatibility on Linux. While this has little effect on the speed of the Linux version, it does take it longer to load.

WP Office also comes with its own font server. The standard version comes with over 100 fonts and the deluxe version comes with over 12,000 fonts. The font server also allows the use of Windows fonts and TrueType fonts by simply copying them to the font directory! WP Office has some of the best-looking fonts of any Linux application. Even six-point fonts are easily readable!

Loading a WP Office application takes about 30 seconds to two minutes. Loading a second application only takes a few seconds, since WINE and the font server are already loaded. Once the application is loaded, response is reasonably quick on a decent computer system. Please note that WP Office is NOT for low-end systems. To get decent

performance, you must have a Pentium 200 or better with 64MB of RAM and 540MB of free disk space.

Corel spent a lot of time making sure WP Office could read and write to Microsoft Office documents. Microsoft Word, Excel, and PowerPoint documents are imported into WP Office with little trouble. Later in the chapter, we will discuss importing and exporting MS Office files.

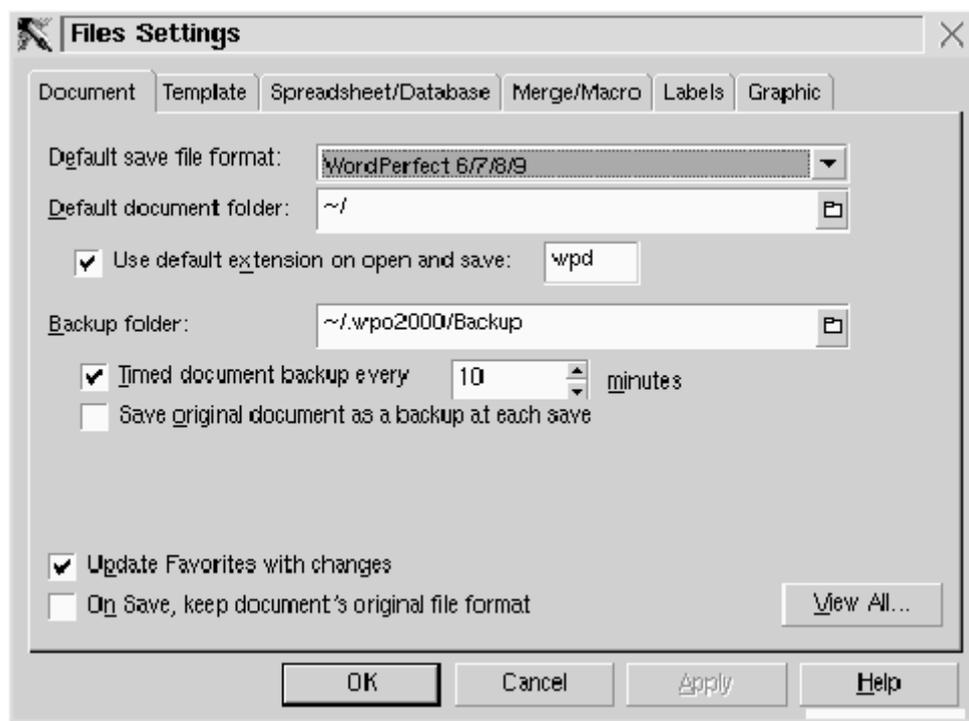
WP Office also has the ability to export files in portable formats such as HTML, XML, PDF, and RTF. HTML is the standard format used by Web pages. WP Office allows documents, spreadsheets, and presentations to be exported directly into HTML to be published on a Web page, eliminating the need for separate HTML editors.

XML is similar to HTML except that users can define their own data formats. While it is still an evolving standard, it promises to allow easier exchange of data. PDF (Portable Document Format) by Adobe allows documents to look exactly the same on different platforms. It is a popular way of distributing government documents and legal briefs. And finally, RTF (Rich Text Format) was first used by Microsoft Word to allow the exporting of Word documents to other word processors. It doesn't support any advanced formatting features of modern word processors, but it is supported by all major word processors.

There are some drawbacks to WP Office. It does have some installation problems on some Linux distributions. It seems to install well on Red Hat, Mandrake, and Corel Linux, but it has problems on other distributions, including the popular SuSE Linux. Many of these problems have work-arounds. See *Corel's support page* at <http://linux.corel.com/support/wpo2000-linux.htm> for more information.

Running WP Office on top of WINE also causes stability problems. WINE is still under development and is not close to being stable yet. Fortunately, a WP Office crash will not take the desktop with it. Just be sure the auto save feature is on by going to Tool -> Settings -> Files -> Document, checking the Time document backup, and setting the number of minutes.

Figure 6.1. Setting the auto save feature. In this case, it is set to save every 10 minutes.



WP Office for Linux is missing two features common to most other office suites: Internet integration and HotSync with Personal Digital Assistants (PDAs). WP Office doesn't have a built-in Web browser or email either. This is not that that big a deal since Netscape is

bundled with WP Office, but sometimes it is easier to check the mail or a Web page without opening another application.

HotSync is included with the Windows version of WP Office and all versions of Star Office. If synchronizing your calendar with your PDA is important, you might want to consider Star Office for Linux.

6.2.1 WordPerfect 8 for Linux

WordPerfect 8 (WP8) is still available for download from the Corel Web site. While it isn't as compatible with Microsoft Office as version 9, it still has most of the features of version 9. It is also much faster, since it uses native Linux code instead of WINE. It is a great choice for low-end systems, since it will even run on an old 486 with 16MB of RAM.

WP8 is available in a free personal edition or a full-blown commercial package that retails for about \$50, although it often sells for much less. The personal edition is a fully functional program. The full version contains extra graphics, clip art, and fonts that are not included with the personal version (<http://linux.corel.com/products/wp8/>).

6.3 Other Commercial Productivity Suites

Lotus SmartSuite has no official plans to make a Linux version, although I wouldn't be surprised if Lotus decides to port SmartSuite to Linux, since they have already ported Notes and Domino to Linux. Also, Lotus' parent company IBM is throwing large amounts of resources into Linux-related development.

6.3.1 Star Office

Sun's Star Office has very good support for Microsoft Office files. It offers equivalent programs for word processing (Writer), spreadsheets (Calc), presentations (Impress), graphics (Draw), email (Mail), calendaring (Schedule), and a database (Base). It also has an integrated Web browser. While it is not as feature-rich as Microsoft Office, it offers all its major functions such as fonts, headers, footers, and template style sheets.

Star Office isn't for everyone, though. Probably the biggest drawbacks are its size and speed. Star Office requires 160MB of disk space and at least 32MB of RAM, although 64MB is required to get decent performance. It also takes a long time to load, even on a fast system. Low-end machines that are often used for Linux boxes are definitely out of the question.

Star Office does have a nifty HotSync feature that is not included with the other office suites for Linux. If synchronizing a PDA with your office suite is important, Star Office has the best support for it under Linux.

Star Office also has its own built-in desktop. On a Windows machine, it will replace the default Windows desktop with its own desktop. While many users are not bothered by this, it can be frustrating to many other users. Star Office can be downloaded for free from Sun's Web site or CDs can be ordered for about \$40 (<http://www.sun.com/staroffice/>).

Sun is releasing Star Office under the GPL open source license. This is good news since it means that Star Office will stay around even if Sun decides not to actively support it. It also means that Star Office can enjoy the benefits of other open source applications (<http://www.openoffice.org/>).

Sun is also beating Microsoft to the punch by offering XML support in Star Office. Microsoft doesn't plan on offering XML support until 2002. The battle for the office suite could get very interesting.

6.3.2 Applixware

While it doesn't support MS Office formats as well as Star Office, many people prefer Applixware's interface and cleaner-looking screen fonts. Instead of taking over the desktop, Applixware loads each application in a separate window.

Applix is also a good choice for low-end systems, since it will run on a 486/66 with 32MB of RAM and about 250MB of free disk space.

Many Linux distributions contain a demo copy of Applix. You can also download it from Applix's Web site. The program is not free, however. A registered copy costs from \$40 for an upgrade to \$190 for the developer's edition (<http://www.applix.com>).

6.4 Open Source Office Suites

Both KDE and Gnome have open source office suites under development. While neither is production-quality yet, they are already feature-rich and complete.

6.4.1 Koffice

Koffice is KDE's office suite. It comes with a word processor (Kword), spreadsheet (Kspread), presentation software (Kpresenter), a charting program (Kchart), a vector-based drawing tool (Killustrator), an imaging program (Kimage), a layer-based image creation program (KimageShop), and a database front-end (katabase).

Koffice is still in development and not ready for everyday use. Having said that, Koffice is surprisingly stable and feature-complete for a product under development. Keep an eye on Koffice. It is an up-and-coming contender!

6.4.2 Gnome Office

Gnome also has an office suite, but unlike Koffice, the Gnome Office is a collection of various open source tools rather than a new suite designed from the ground up. The goals of Gnome Office are to provide the features of Microsoft Office and have good compatibility with Microsoft Office files. Unfortunately, like KDE, the Gnome Office components are still under development.

The heart of Gnome Office is the spreadsheet, Gnumeric. This is a Gnome project designed to create a full-featured GPL spreadsheet. One of its goals is to clone the look and features of Microsoft Excel. AbiWord is Gnome's word processor. It doesn't have all the features of many of the other word processors, but it is very fast and has a clean interface. Gnu Image Manipulation Program (GIMP) was one of the first open source desktop applications available for Linux. GIMP is a photo and image retouching program similar to Adobe Photoshop. While its interface isn't quite as polished as Photoshop, GIMP is every bit as powerful. Dia is a drawing program similar to Visio. Dia is quickly becoming the tool of choice for Gnome developers to do diagrams and flowcharts. Eye of Gnome (EOG) is a quick image viewer. It was designed to allow developers to quickly view image files without having to open them in a larger application like GIMP. Gnome-PIM provides the calendar and address book that is standard in all commercial office suites. Finally, Gnome-DB provides applications and libraries to design easy front-ends for databases. Complete details on the *Gnome Office* are found at <http://www.gnome.org/gnome-office/>.

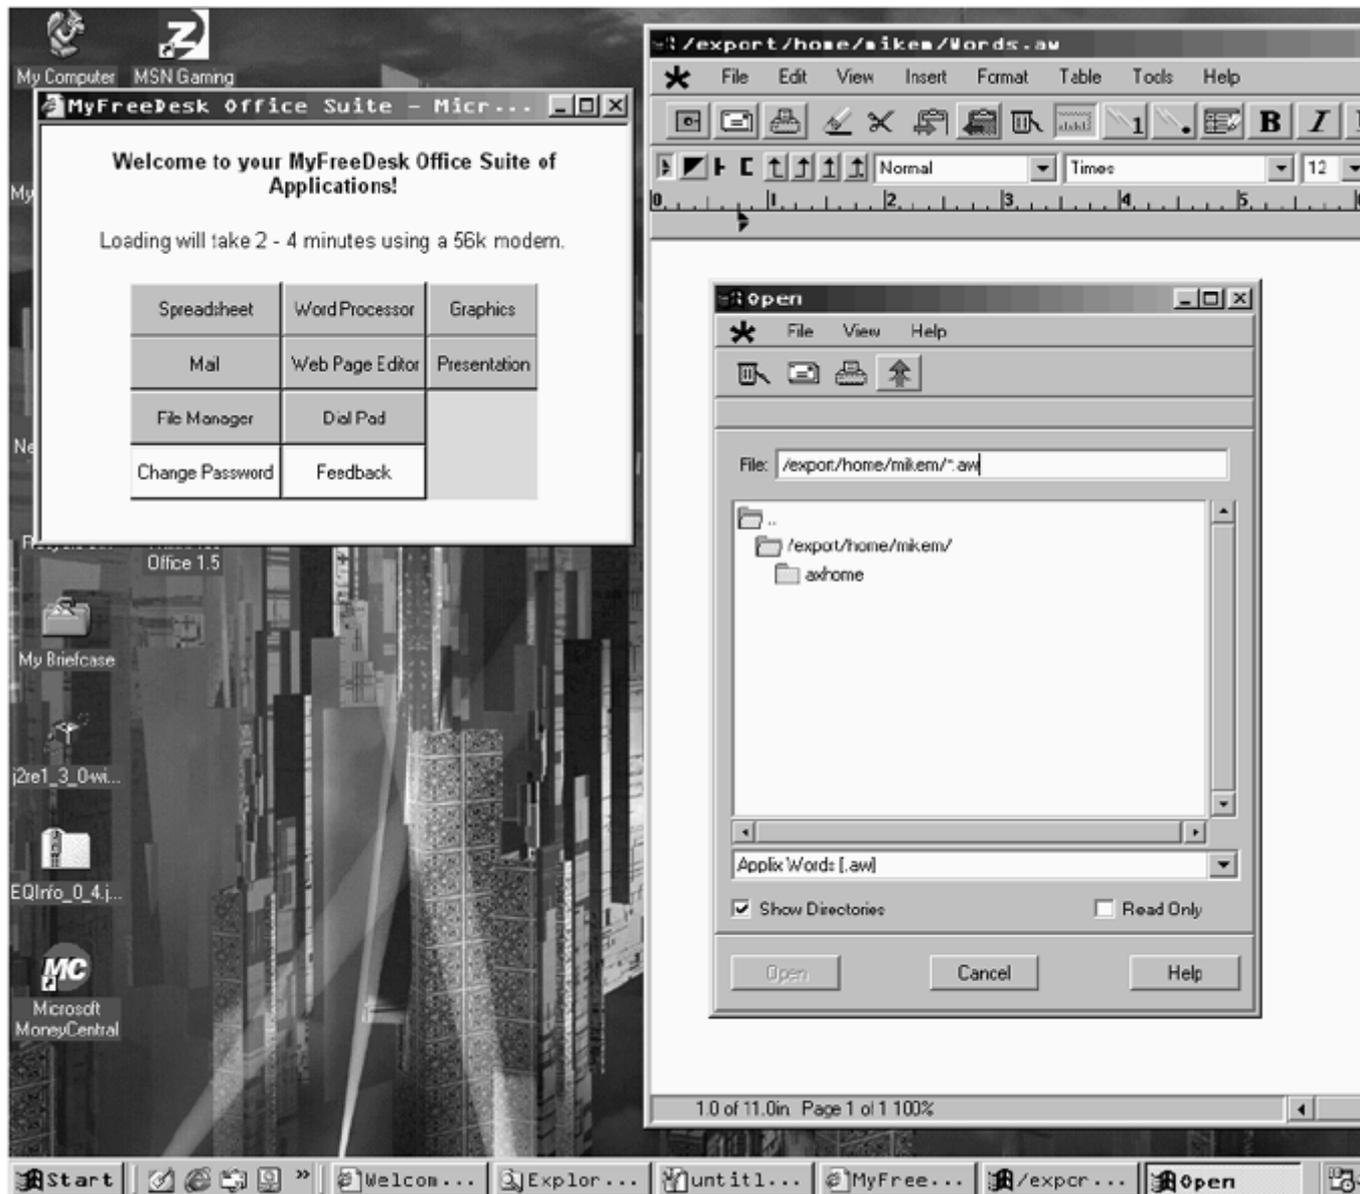
6.5 Web-Based Suites

It's hard to tell whether they will catch on, but Web-based office suites are truly cross-platform. These suites run inside of Netscape or Internet Explorer, which allows them to run on virtually any operating system.

Web-Based Suites (WBS) come in two different flavors: server-based applications and client-based applications. Server-based applications run on the server and send the screen image to the client machine through the Web browser. The file is typically stored

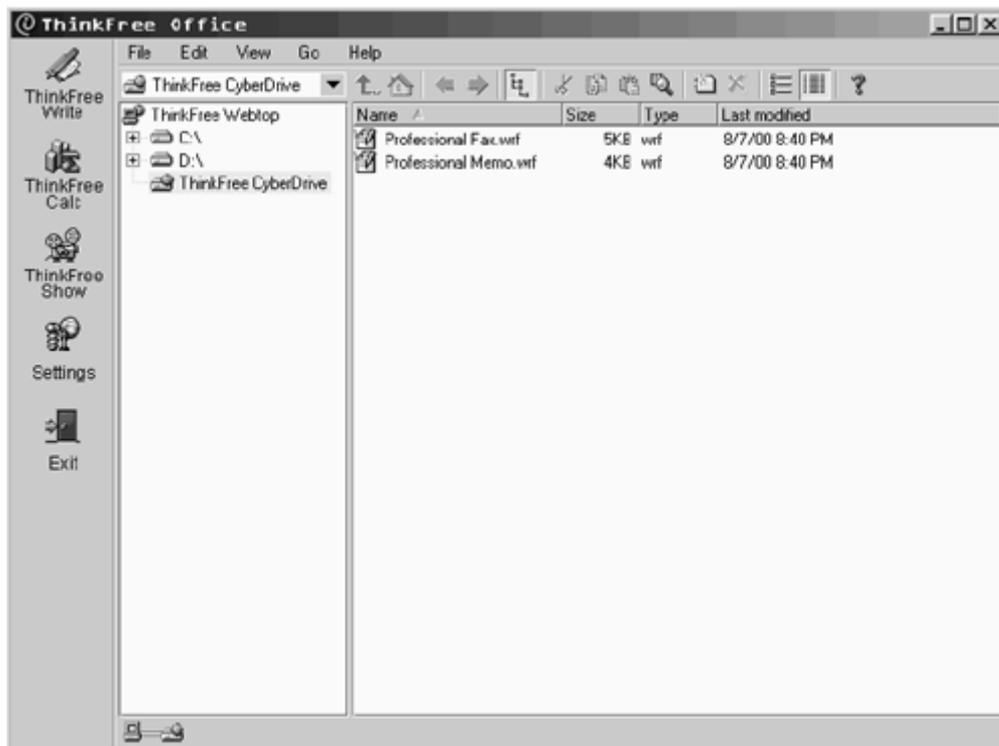
on the server. The compatibility with MS Office documents varies. Some of the server-based applications use MS Office as the back-end and offer near perfect compatibility. Others use Star Office, Applix, or home-built applications. An example of a server-based application is FreeDesk.com, which uses a free, home-grown application (<http://www.freedisk.com>).

Figure 6.2. The Free Desk Office Suite.



Client-based applications require you to download a Java application that runs inside your browser. An example of one of these applications is ThinkFree.com. After signing up, you download a Java program (about 10 MB) that contains a word processor, spreadsheet, and presentation software. After downloading the program, it is started in a browser window. It is more a light-duty suite like MS Works than a full-fledged office suite like MS Office or WP Office. Unlike the server-based applications, client-based applications allow the user to work and store documents offline. These office suites aren't as mature as the standard stand-alone suites. In the case of ThinkFree, the word processor works fairly well, but the spreadsheet and presentation software are still works in progress. ThinkFree is still beta software, so it should improve over time (<http://www.thinkfree.com>).

Figure 6.3. The ThinkFree Office Suite.



Web-based office suites are still relatively new. Will they catch on? They do offer cross-platform compatibility and the ability to store documents on the Web. On the downside, they aren't as mature and full-featured as the stand-alone applications. The server-based applications have the additional disadvantage of not allowing offline use. Only time will tell if WBS will catch on.

6.6 Reading and Writing Microsoft Office Files

Using Office files in other programs has always been tricky. Even using Office files in older versions of Office causes problems. Converting a complex MS Word document from Word 97 format to Word 95 format will often result in a loss of formatting. To further complicate matters, MS Office has changed the data format in every version of Office. This means that Office 95 will not be able to open documents created in Office 97. This often forces people to upgrade their version of MS Office just to be able to read documents created by the newest version of Office. Things may be changing in this regard. MS Office 2000 uses the same format as MS Office 1997.

6.7 Exporting MS Office Files

Fortunately, MS Office allows the user to save files in a wide variety of formats. Besides allowing saves in older formats, MS Office allows saving files in formats of competing programs such as WordPerfect, Lotus 123, Quatro Pro, Freelance, and many others. While the conversion is not perfect, it at least allows MS Office users to exchange files with non-MS Office users as long as the MS Office user knows what programs the other person is using.

If the end format is unknown, you can use a portable format such as Rich Text Format (RTF), Hypertext Markup Language (HTML), Extensible Markup Language (XML), Portable Document Format (PDF), Data Interchange Format (DIF), delimited text, or plain text (TXT or ASCII). All of these formats can potentially lose data, so be sure to save a copy in the native Office format as well.

RTF is a format originally used by MS Word to allow easy exchange of documents with other word processors. It doesn't support some of the advanced features of MS Word, but it is supported by all modern word processors (<http://msdn.microsoft.com/library/specs/richtextformatrtf/specifications/amplertfreaderprog/ramversion15.htm>).

HTML is the format used by the World Wide Web (WWW). One major advantage of HTML is that all operating systems include a Web browser, so HTML files can be viewed without any additional applications (<http://www.w3.org/MarkUp/>).

XML is similar to HTML except that it is more flexible. HTML has a limited set of formatting options that are predefined by the HTML standard. XML allows the creation of user-defined tags, which allows for more formatting options. XML also promises to deliver one of the holy grails of data processing—the separation of data from the formatting. This would allow data to be stored in one place and then be delivered in multiple formats. The problems with XML is that it is still an evolving standard and it isn't supported by all programs (<http://www.w3.org/xml>).

PDF is a standard created by Adobe that allows documents to be displayed on a wide variety of platforms. It is a popular way of creating government and legal documents since it is portable and it supports many advanced formatting features. Adobe has programs that can create and read PDF files on Linux, Windows, and many other platforms. Adobe has a free PDF reader and there is also a free Linux viewer. In addition, many other programs such as word processors can create PDF files.

There are some useful open source programs that help with PDF files. xpdf will open PDF files, although it sometimes has problems with viewing some PDF files. Ghostscript will also convert PostScript printer files to PDF files.

DIF is a format originally used by databases as a way of exchanging files. It is supported by many spreadsheet and database programs.

Delimited text is another way of exchanging data from spreadsheet and database programs. It is much simpler than the DIF format and it is supported by all modern word processing, spreadsheet, and database programs (<http://www.microsoft.com/enable/download/products/office/office97/part06.txt>).

Standard text, also known as ASCII (American Standard Code for Information Interchange), is the standard format for storing text in both DOS and UNIX systems. While it doesn't support any formatting, it can be read by virtually any computer program. DOS and UNIX do use a different end of line character, though. There are programs to convert DOS text into UNIX text, and vice versa. For example, from the Linux prompt, you can type the following to convert a DOS text file to UNIX: `tr -d '\15\32' <dosfile.txt > unixfile.txt`

You can also use `awk` to translate:

- DOS to UNIX: `awk '{ sub("\r$", ""); print }' dosfile.txt > unixfile.txt`
- UNIX to DOS: `awk 'sub("$", "\r")' unixfile.txt > dosfile.txt`

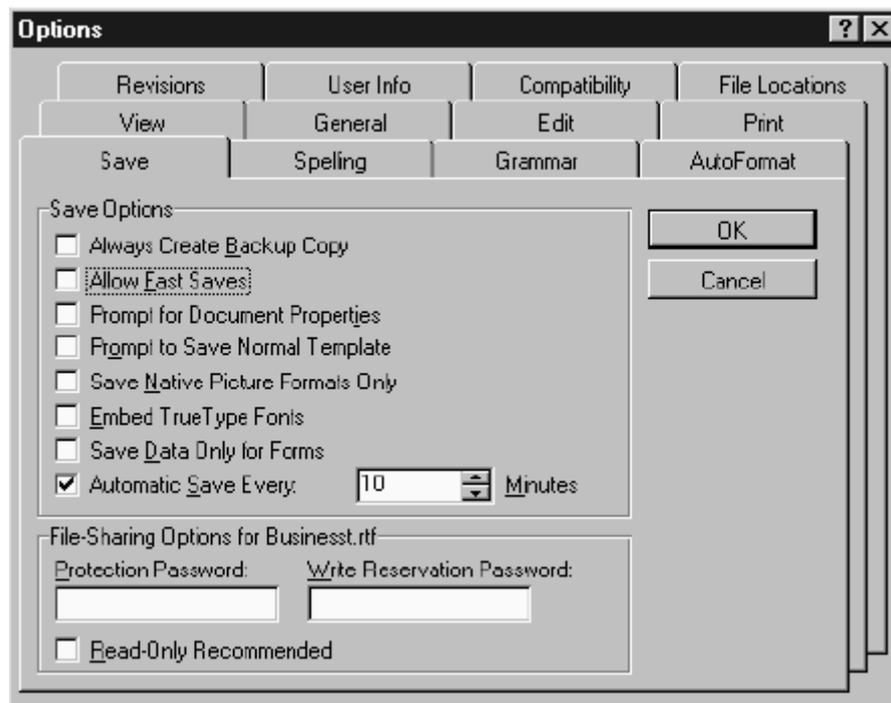
There are some techniques that can make exporting MS Office files easier. Obviously, more complex files have more trouble exporting, so if you plan to export a file, keep it simple. Embedding and linking within documents can also cause problems.

The main difference between embedding and linking is that embedding places a copy of the object in the Office file. Linking places a link to the object in the file. If the original object is changed, the embedded object will not change, but the linked object will be updated to match the original object. Linking is handy because you don't have to update all of the documents when you change the linked object, but it also doesn't export very well. If you are going to export a document, use embedding instead of linking. Or better yet, instead of embedding or linking, simply include the object as a separate file. For information on how to embed and link in an MS Office document, refer to your MS Office user's manual or go to this tutorial.

WordArt, AutoShapes, vertical rules, macros, pivot tables, and equations are also specific to MS Office and may not be exported properly. These options are explained in the MS Office documentation.

Turning off the fast saves feature in MS Word also makes files easier to export. Fast saves has always been a buggy feature and it is better to turn it off anyhow. To turn this off, go to Tools -> Options -> Save and uncheck the Allow Fast Saves option.

Figure 6.4. Turn Off the "Allow Fast Saves" option to allow easier imports of Word Documents.



One last note: MS Office 95 files are easier to import than Office 97 files.

6.8 Importing and Exporting MS Office Files with Linux

Since MS Office has about 90% of the market for Office products, importing and exporting MS Office files from Linux programs is a necessity. I will use Sun's Star Office in this example since it is the most popular office suite for Linux and it supports MS Office files very well. Star Office supports these MS Office features:

- Microsoft Word:
 - AutoShapes.
 - Revision marks.
 - OLE objects.
 - Indexes.
 - Tables, frames, and multi-column formatting.
 - Hyperlinks and bookmarks.
- Microsoft Excel:
 - Pivot tables.
 - Chart types.
 - Conditional formatting.
 - AutoShapes.
 - OLE objects.
 - Most functions/formulas.
- Microsoft PowerPoint:

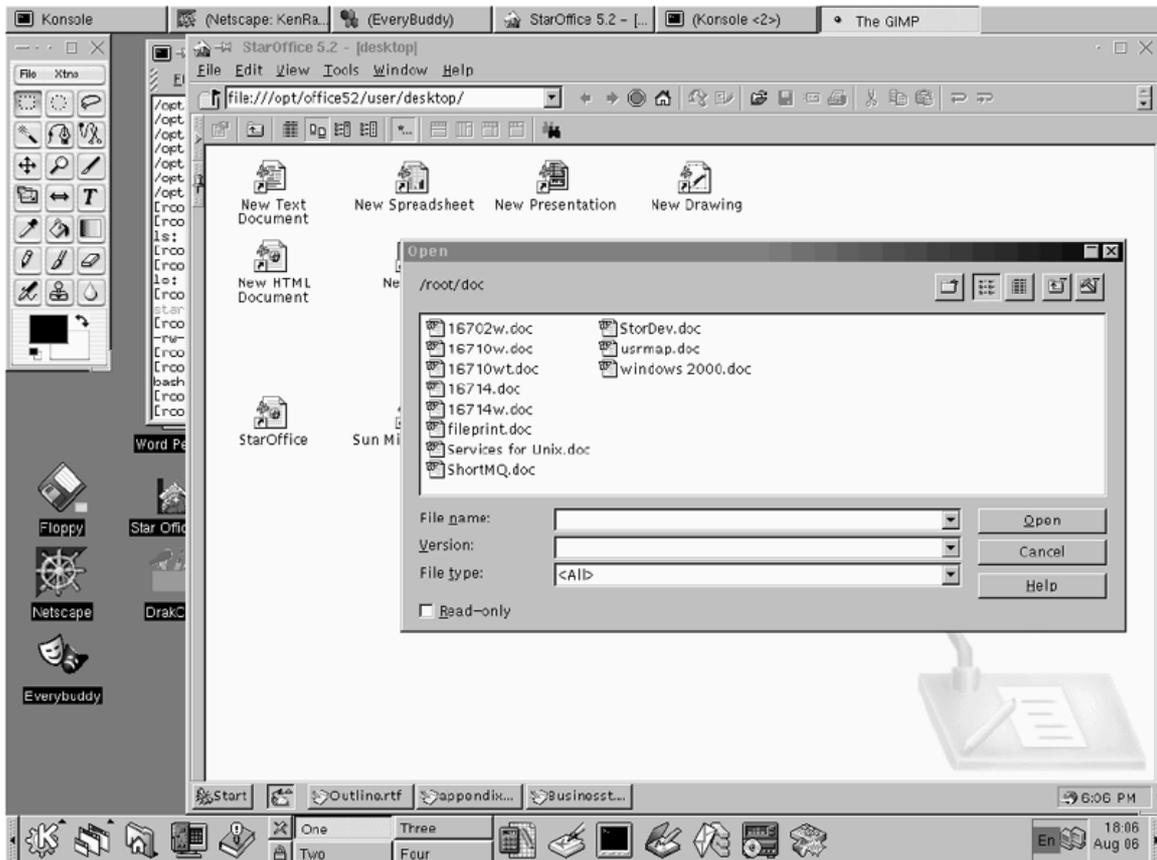
- AutoShapes.
- Tab, line, and paragraph spacing.
- Master background graphics.
- Grouped objects.
- Multimedia effects.

6.9 Using MS Office Documents with Star Office

Star Office has very good support for MS Word and Excel files. Although it may lose some of the formatting, you can almost always get a readable file by opening a Word or Excel document in Star Office. On the other hand, PowerPoint is listed as one of the supported formats, but Star Office is not always able to open PowerPoint files.

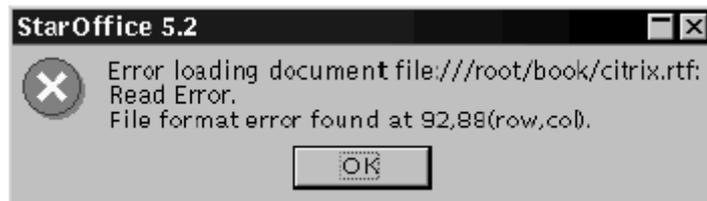
To open an Excel or Word file, simply choose File -> Open from the menu. Word files will appear with the familiar Word icon and Excel files will have the Excel icon in the file listings.

Figure 6.5. Star Office shows the familiar word icons on Microsoft Word documents.



Next, simply click on the file and choose Open. If Star Office recognizes the file, it will simply open it. If there are problems with the file, it will give the following error message:

Figure 6.6. The error message displayed when there is a problem opening a file.



When saving a file, it is a good idea to use the Save As option so as not to overwrite the existing file. Some of the features of Star Office may not translate properly to MS Office format. It is also a good idea to save the document in Star Office format if you are worried about losing the formatting. Star Office uses StarWriter format for text documents and StarCalc for spreadsheets.

Figure 6.7. Saving Star Writer files in their native format.

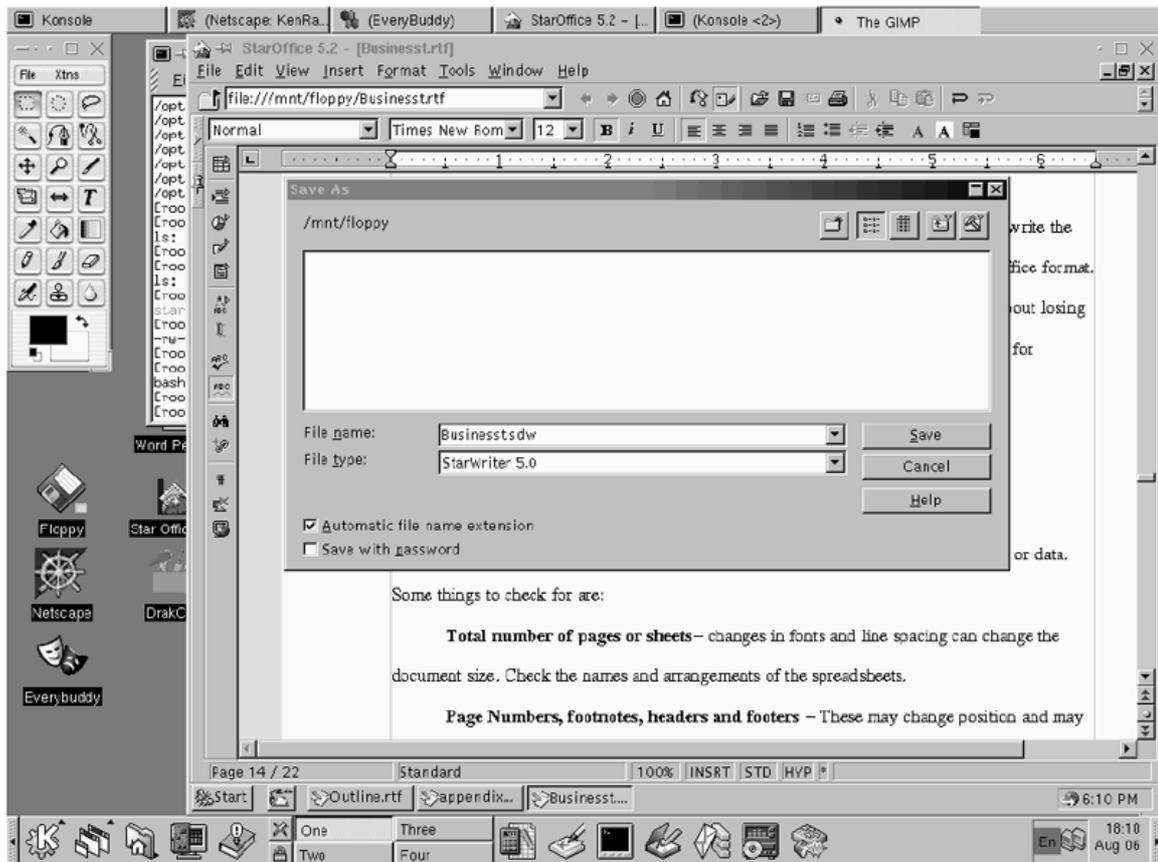
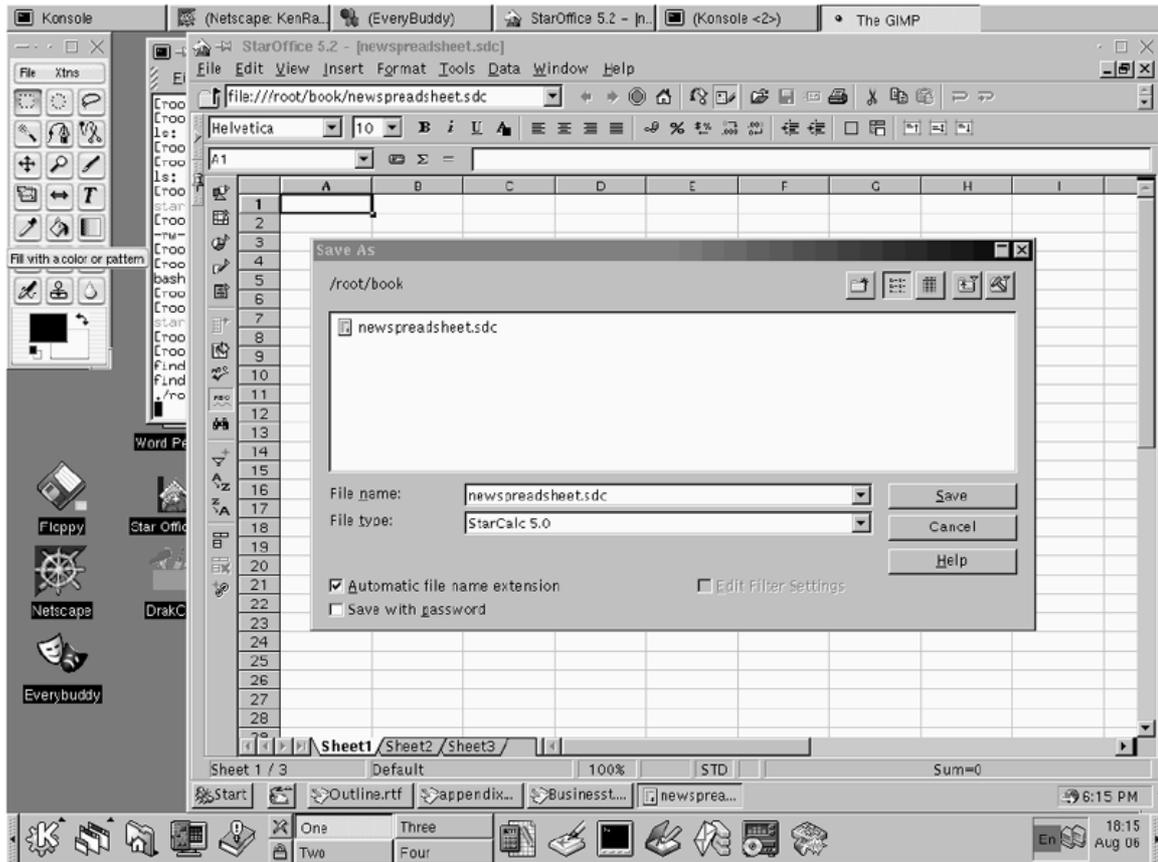


Figure 6.8. Saving StarCalc files in their native format.



6.10 Checkpoints When Importing and Exporting

Anytime a file is imported or exported, there is a chance of losing formatting or data. Some things to check for are:

- **Total number of pages or sheets**— changes in fonts and line spacing can change the document size. Check the names and arrangements of the spreadsheets.
- **Page numbers, footnotes, headers, and footers**— These may change position and may even overlap. Also check the font sizes.
- **Fonts**— The font type and size may have changed.
- **Graphics and charts**— Unsupported graphics will be displayed as gray boxes. Also check the positioning of graphics. Chart styles and data may have changed.
- **Columns**— Check the alignment, gutter, and text flow.
- **Tables**— Check spacing (height and width), shading, and fonts.
- **Document properties**— Check the meta information such as author and date.
- **Special markers**— Check special markers such as index and table of contents entries.
- **Macros and templates**— Macros and templates generally won't import or export. You may have to reformat the document to make it look the same.

- **Formulas**— Long, complex formulas may have problems. Pay particular attention to absolute cell references and operations that rely on the order of the calculation.
- **Special features**— Data validation, help notes, sheet protection, and user-defined functions may not import properly.

While Star Office has very good support for MS Word and Excel files, it has virtually no support for other word processor or spreadsheet formats. It doesn't support WordPerfect or Quatro Pro formats, and it opens Lotus 123 files as word processing documents, which doesn't do much good if you want to edit a Lotus file. If you want to import one of these files in Star Office, you must save it in MS Office, RTF, HTML, or DIF format first.

6.11 Financial Programs

Keeping track of personal finances is one of the most popular uses for home computers, which is why personal financial programs such as Intuit's Quicken and Microsoft Money are some of the most popular programs for home users. While there are no Linux programs as feature-rich as these two programs, Moneydance and GnuCash can fill many of their functions.

Moneydance is currently the most feature-rich Linux financial program. It supports the following features:

- Transaction auto-completion.
- Transaction search.
- Check printing (regular checks or checks with stubs).
- Multiple currencies (you can even define your own!).
- Support for multiple accounts and true double entry.
- An easy-to-use checkbook register-style interface with many shortcuts.
- Custom 3-D and 2-D graphing with the ability to print or export to GIF files.
- Report generator with export to HTML and printing capabilities.
- QIF (Quicken Interchange Format) file import and export.
- Cleared vs. actual account balances updated as you work.
- Multiple languages (Italian, German, French, Brazilian/Portuguese, and UK, and U.S. English) and number/date formats.
- Scheduled transactions and reminders.
- Reconciliation tool to simplify checkbook balancing.
- Automatic notification when upgrades are available.
- Extensive transaction sorting options.
- Support for multiple dates per transaction.
- Support for split transactions.
- Currency fields support math expressions and automatic currency translation.

There are also several features currently in development, including:

- Online banking.
- Budget management and tracking.
- Stock portfolio and investment management.
- Online portfolio updates.

Moneydance is written in Java and requires the Java Development Kit (JDK) to run. It is also not GPL and costs \$50. It runs under Linux, Windows, Macintosh, and most versions of UNIX.

GnuCash, however, is free and GPL. It was created by merging the two open source financial programs GnuCash and X-Accountant. While it isn't as far along as Moneydance, it is developing rapidly. Some of GnuCash's features include:

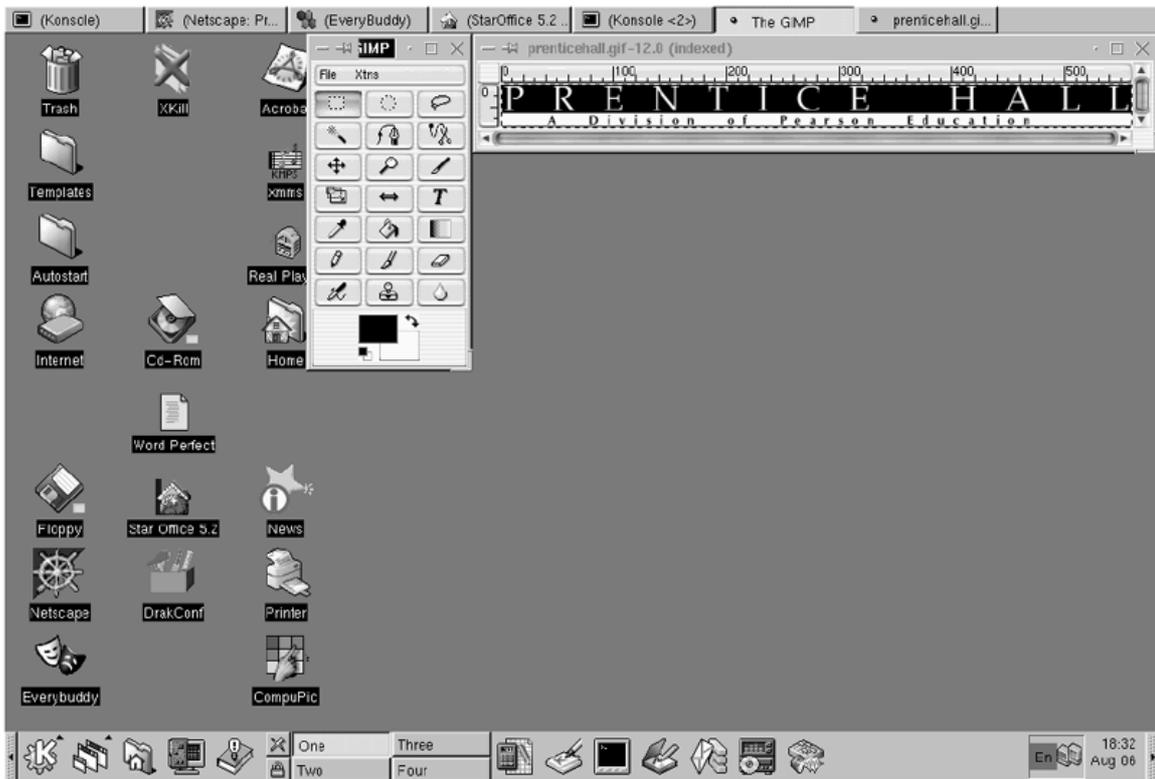
- **An easy-to-use interface**— If you can use the register in the back of your checkbook, you can use GnuCash. Type directly into the register, tab between fields, and use quickfill to automatically complete a transaction.
- Reconcile window with running reconciled and cleared balance makes reconciliation easy.
- **Stock/mutual fund portfolios**— Track stocks individually (one per account) or in a portfolio of accounts (a group of accounts that can be displayed together).
- **Multiple currencies and currency trading**— Multiple currencies are supported and can be bought and sold (traded). Currency movements between accounts are fully balanced when double entry is enabled. (Some aspects of multiple currency support are not fully implemented.)
- Quicken files are automatically merged to eliminate duplicate transactions.
- **Reports**— Displays Balance Sheet, Profit & Loss Statement, portfolio valuation, or prints them as HTML.
- **Chart of accounts**— A master account can have a hierarchy of detail accounts underneath it. This allows similar accounts types (e.g., cash, bank, stock) to be grouped into one master account (e.g., assets).
- **Split transactions**— A single transaction can be split into several pieces to record taxes, fees, and other compound entries.
- **Double entry**— When enabled, every transaction must debit one account and credit another by an equal amount. This ensures that the "books balance", or that the difference between income and outflow exactly equals the sum of all assets, be they bank, cash, stock, or other.
- **Income/expense account types (categories)**— These serve not only to categorize your cash flow, but when used properly with the double-entry feature, these can provide an accurate Profit & Loss Statement.
- **General Ledger**— Multiple accounts can be displayed in one register window at the same time. This can ease the trouble of tracking down typing/entry errors. It also provides a convenient way of viewing a portfolio of many stocks, by showing all transactions in that portfolio.
- Written in C, with Perl, scheme, and tcl support for easy configuration and extensibility.
- File access is locked in a network-safe fashion, preventing accidental damage if several users attempt to access the same file, even if the file is NFS-mounted.
- Provides a byte-stream format, which allows accounts and account groups to be transmitted to other processes via pipes or sockets.
- Gets stock and mutual fund quotes from various Web sites. Updates portfolio automatically (more funds being added regularly).

The GnuCash program and documentation can be found at <http://www.gnucash.org>.

6.12 Graphics Programs

The ability to use graphics is increasing in importance in business. All the office suites have a varying ability to manipulate graphics, but sometimes you need to do more. Stand-alone graphics programs come in two types: photo editors and drawing programs. Photo editors allow you to manipulate photographs and images. These programs imitate many of the tools familiar to artists, including a paintbrush, pen, and airbrush. The most popular of these programs is Adobe Photoshop. The open source program Gnu Image Manipulation Program (GIMP) provides the same functions for Linux. GIMP is every bit as powerful as Photoshop, although its interface is not quite as intuitive. GIMP is included in most Linux distributions and both Linux and Windows versions can also be downloaded from <http://www.gimp.org>.

Figure 6.9. Gimp provides a Photoshop work alike for Linux. A Windows version is also available.



Drawing programs have a different purpose than photo editors. Drawing programs are intended to create charts, graphs, and presentations. Two of the most popular drawing programs are Microsoft Visio and Corel Draw. Fortunately, Corel Draw is available for Linux, Macintosh, and Windows. Depending on how sophisticated the features you need, Corel Draw sells for \$150 to \$700. The top of the line product includes features needed for professional publishing such as color calibration.

Blender and POV-Ray are 3-D rendering tools available for Linux, UNIX, and Windows. Blender is an integrated suite for modeling, rendering, and animation. It is based on tools developed in-house at a digital animation studio. Since Blender is copywrited freeware, it is a good way to try out digital rendering and animation without spending a lot of money (<http://www.neogeo.nl/>).

Figure 6.10. Blender is a free professional quality animation tool available for Windows and Linux.



POV-Ray is one of the best rendering tools available. It uses a text file that describes the objects that it renders. There are many built-in functions like textures, shapes, and colors. It doesn't have a modeling tool, but there are several shareware and commercial modeling tools linked from the POV-Ray Web site. Like Blender, it is copyrighted freeware. Unlike Blender, the source code is also available <http://www.povray.org>.

6.12.1 Graphics Conversion

GIMP, Corel Draw, Photoshop, and most other graphics programs can open and convert many different file formats. Sometimes, though, you run into a graphics file that your program won't support. Other times you have a large number of files that need to be converted and it would take days to open and convert every file. For these situations, you need a graphics conversion program.

There are two types of graphic conversion programs: viewers and bulk converters. A viewer allows you to view the graphic, then save it into another format. This allows you to check and adjust the characteristics of the file before converting it. A bulk converter usually doesn't let you view the graphic before conversion, but it will convert many files at once.

There are dozens of free and shareware graphics conversion programs for Windows. A good list of the various programs is available at <http://www.softseek.com/Graphics and Drawing/Graphics Conversion and Optimization/>.

There are not as many graphics converters available for Linux. One of the best viewers on the market is xv. It is bundled with most Linux distributions. It is a shareware program. There is a charge of \$25 for the program and \$40 will also include printed manuals. The program, source code, and documentation are available at <http://www.trilon.com/xv/xv.html>.

For bulk conversions, use the free utility NetPBM. NetPBM is not a single program, but a set of about 180 command-line programs. There are separate programs for converting between different formats. For example, to convert all PNG (a graphics format popular on

UNIX systems) files to JPEG (a graphics format that creates very small graphics files) files, type the following command:

```
for i in *.png; do pngtopnm $i |ppmtojpeg >`basename $i
.png`.jpg;
done
```

Let's break the command down:

- `for i in *.png;`— This chooses all PNG files.
- `do pngtopnm I`— This is the program that converts PNG to PNM . PNM is an intermediate format used by NetPBM.
- `pnmtojpeg`— This program converts the PNG to JPEG.
- `>'basename $i.png' .jpg;`— This creates a .JPG file with the same name as the .PNG file.
- `done`— Exits the program when done.

The syntax of NetPBM is rather complicated, but it is great for bulk conversions. The program and documentation are available at: <http://netpbm.sourceforge.net/projects/netpbm>.

ImageMagick also has a nifty collection of conversion tools available at <http://www.wizards.dupont.com/cristy/ImageMagick.html>.

6.13 The Last Word

Linux is not nearly as popular on the desktop as Windows, but the rapid development of desktop tools is changing that. Most desktop applications now have equivalent Linux applications. The functions not available for Linux are rapidly being filled by free and commercial applications. Linux's low cost, stability, and low hardware requirements make it not only a viable replacement for Windows, but sometimes a better choice.

6.14 Conclusions

So which office suite for Linux is the best? The answer depends on what you need, of course. If compatibility with MS Office is paramount, nothing currently beats WP Office. WP Office is also the most feature-rich office suite available for Linux.

For a free office suite, it's hard to beat Star Office. It has a good set of features and good compatibility with Word and Excel documents.

Applixware and WordPerfect 8 are the clear choices for low-end systems, since they both run well on a 486 with 32MB of RAM.

The open source office suites Koffice and Gnome Office are not ready for everyday use, yet they should soon challenge the commercial offerings the same way Linux is challenging Windows.

Although they show promise, it is too soon to tell whether WBS can challenge the stand-alone applications.

Chapter 7. Databases

Databases are essential for storing the information that businesses use. Both traditional and online businesses use databases to track orders, merchandise, and customers. Since they are so essential to business, it is not surprising that there are dozens of commercial and free databases and database add-ons.

The most popular commercial databases are available for both Linux and Windows as well as several versions of UNIX. Many are also available for popular mainframes and mini-computers. Some of the more popular commercial databases are:

- DB2—This is IBM's heavy-duty database that originally ran on IBM mainframes and mini-computers. It is also available for Linux, Windows, Os2, and most versions of UNIX (<http://www-4.ibm.com/software/data/db2/udb/>).
- Ingres—This is a database by Computer Associates, which is the second biggest software company behind Microsoft (<http://www.cai.com/products/ingres.htm>)
- Oracle—One of the most popular commercial databases that is available for almost every computer platform (<http://www.oracle.com/>).
- Informix—Offers both a free development version and a commercial version of its database (<http://www.informix.com/>).
- Sybase—Has ported its flagship database, Adaptive Server Enterprise, as an unsupported, free version for Linux (<http://www.sybase.com/adaptiveserver/>).
- Adabas D—This was the first commercial database available for Linux. It is derived from the Adabas Mainframe database. It is bundled with many Linux distributions, including Caldera (<http://www.caldera.com/>) and SuSE (<http://www.suse.de/>). More information can be found at <http://uebb.cs.tu-berlin.de/~krischan/adabasd/LinuxHowTo.html>.

There are also two popular open source databases:

- MySQL—This one is free for Linux users and \$200 for Windows users. Either way, it is one of the least expensive databases on the market. There is also commercial support available for a yearly fee (<http://www.mysql.com>).
- PostgreSQL—This popular UNIX database rivals many of the commercial databases in features. It is available for free, with commercial support available for a yearly fee (<http://www.postgresql.org/>).

There are also two Windows-only databases that are widely used:

- Access—This database is bundled with Microsoft Office Professional. It is intended as a light-duty database and it is popular as a front-end database server.
- SQL Server—This is Microsoft's heavy-duty database. It is popular on Microsoft-based networks.

While this list is not comprehensive, it covers some of the more popular databases.

Using Databases

All the major databases support Structured Query Language (SQL). SQL is a standard from the International Organization for Standardization (ISO) that provides a uniform language for manipulating databases. That being said, all the databases use a slightly different variation of SQL, with some supporting the ISO standard better than others.

SQL has commands for creating databases, adding data, and retrieving data. Create is used to create and change the tables in a database. The major Create commands are:

- `Create Table`—Creates a table.
- `Drop Table`—Deletes a table.
- `Alter Table`—Changes a table.
- `Insert`—Inserts data into a table.
- `Select`—Retrives data from a table.

Of course there are many other commands and subcommands in SQL. To further complicate matters, each database has its own command syntax. If you are serious about using databases, get a book on the specific database you plan to use.

Choosing a Database

Many factors go into choosing a database, which is one of the reasons why so many are on the market. Among the most important factors are fault tolerance, support, and price. Fault tolerance is especially important for databases that are updated often. While there are many factors to consider in database fault tolerance, three of the most important are transactions, hot backups, and locking.

Transactions keep backup copies of any data being updated. This allows old data to be recovered in the event of an error. While this ensures data integrity, it also creates a lot of overhead. Transactions take a lot of processor time and disk space. It is not unusual for a database to run up to ten times slower while using transactions. Also, the files created by enabling transactions are often larger than the database itself.

Hot backups allow data to be backed up without shutting down the database. This feature is a must-have on databases that need to operate 24 hours a day.

Locking prevents two users from updating the same information at the same time. There are two types of locking in common use: record locking and table locking. Record locking locks only the individual record that is being updated. Table locking locks the entire table that the record is in. Table locking only allows one user at a time to update a table. Record locking can allow multiple users to update a table, but it is slower than table locking.

Support and price go hand-in-hand. Large organizations that depend on databases are willing to pay up to several thousand dollars per user for a database from a large company such as IBM or Oracle. The cost of the database and support are much less than the cost of the database going down. The personal and free databases don't have near the same level of support.

All of the commercial databases, including Microsoft SQL Server, support transactions, hot backups, and locking. They also have extensive support available and the prices to match.

Of the free open source databases, PostgreSQL has all the major features of the commercial databases. Support is also available for an extra fee.

MySQL is another popular open source database. It is designed differently than PostgreSQL and the commercial databases. MySQL doesn't support transactions and uses table locking instead of record locking. This makes reading and updating tables much faster on MySQL compared to most other databases. MySQL is a good choice for databases that are read often but rarely updated. It is a popular choice for Web-based databases, since many Web-based databases are used mostly for reading data.

MySQL does have its disadvantages. Multiple users can't update a table at the same time. Also there are no transactions, so if there is an error in updating a table, there is no easy way to recover the data. These limitations can be partially overcome by small tables and frequent backups.

Connecting Databases

Microsoft's Open Database Connectivity (ODBC) software is the most common way to connect to or from a Microsoft database. Many databases and Web development tools

have ODBC connectivity built into them. There are also many stand-alone ODBC programs such as the ones listed in the next section.

ODBC

ODBC consists of server and client software. The server software runs on the database end and the client runs on the machine accessing the database. The server typically runs as an NT service and it is configured in the Control Panel. The client is either an NT service or a stand-alone program.

Separating the server from the client makes ODBC a good cross-platform solution, but it also makes it harder to troubleshoot. ODBC gives rather cryptic errors that don't really tell you what the problem is. It is also hard to tell whether an error was caused by the client or the server.

While Windows includes ODBC drivers, OpenLink Software makes them for most UNIX platforms. A fully operational demo copy can be downloaded from <http://www.openlinksw.com/>. The price for the commercial version starts at \$675.

OpenLink also has a program Virtuoso that uses an intermediary file format to connect to a database. It is available in both Linux and Windows NT versions. Virtuoso prices start at \$599.

There is also a free open source ODBC program available from FXML. It consists of a Windows NT server and clients for Windows, Linux, and UNIX (<http://odbc.linux-ave.com/>). The FXML ODBC server currently works under Windows NT. It is easily configured using the Registry editor. Among the settings that can be configured are the port used, the maximum number of threads and connections, the hosts that are allowed to connect, and the timeouts. If you need further customization, the source is available and can be compiled using Microsoft C++ 6.0.

FXML ODBC ships with clients for COM (used by Microsoft C++), Perl and PHP (for Web pages), and C++ (for Linux and UNIX). If this is not enough, any ODBC interface that supports TCP/IP and XML can be used.

Another inexpensive way to do this would be to use the free database MySQL and its MyODBC drivers. First, download MySQL for Linux and MyODBC for Windows from the MySQL Web site, <http://www.mysql.org>. Install and set up MySQL on the Linux machine. Once the data is imported into MySQL, you can use PHP's built-in support for MySQL to use the data.

To install the MyODBC drivers on the Windows machine. Set up the ODBC driver to connect to the MySQL database on the Linux machine. Open Access and right-click on the database. Select Get External Data -> Link Tables. In the box that comes up, choose ODBC, then select the name of the MyODBC link. You can now edit data and create queries, forms, and reports that will be exported to the MySQL database.

When exporting an Access table to MySQL, you must have a primary key on both the Access table and the MySQL table. If there are any changes made to the table with MySQL, you must re-link the tables. To do this on your Windows machine, go to Start -> Programs -> Tools -> Add-ins -> Linked Table Manager. Go to the ODBC DSN for the table and select the table to re-link.

Chapter 8. Fun and Games

[Section 8.1. Games](#)

[Section 8.2. Game Servers and Extras](#)

[Section 8.3. Classic Games](#)

8.1 Games

Let's face it, for home users, having a lot of cool games is a compelling reason for using a computer. Many of the early adopters of a new technology tend to be active gamers. As Bob Young, the CEO of Redhat Linux put it, "Computer games designed for the Linux operating system are important in spurring the growth of the free-software industry" (<http://www.techweb.com/wire/story/TWB19990520S0014>).

Right now, Windows 9x is the king of the PC gaming platforms. This is due mostly to the fact that Windows 9x has about 80% of the PC market share. In fact, a lot of Linux users keep a Windows 9x partition on their hard drive just for games. Another large part of Windows' success on the game platform is due to Microsoft's DirectX.

DirectX is a set of Application Program Interfaces (APIs) that allows accelerated graphics and sound. Since it is hardware-independent, a program that uses DirectX doesn't need to worry about the type of hardware used on the PC. Currently, only Windows 9x and 2000 support DirectX. Further information on *DirectX* is available at <http://www.microsoft.com/directx/>.

Although Linux is currently behind Windows as a gaming platform, it is rapidly catching up. Interest in Linux gaming is at an all-time high and many popular games are being developed for the Linux platform.

8.1.1 Loki

Loki entertainment software is at the forefront of Linux gaming. Loki was founded in August 1998 with the goal of bringing a wide range of commercial games to Linux. It signed its first contract on December 31, 1998 to port Activision's popular strategy game "Civilization—Call to Power." When the game hit the shelves the next spring, it proved that a Linux game could be equal to its Windows counterpart. It also showed that commercial game companies could make extra money, since porting a game requires a lot less effort and risk than developing a new game (<http://www.lokigames.com>).

This led to several other contracts. Some of the other games Loki has brought to Linux are:

- Myth II - Soulblighter—A fantasy/adventure game from Bungie.
- Railroad Tycoon II and Second Century Expansion Pack—A strategy/simulation game from Pop Top Software and Gathering of Developers.
- Eric's Ultimate Solitaire—A collection of 23 different solitaire games.
- Heretic II—A third-person adventure game from Activision and Raven.
- Heroes of Might and Magic III—A fantasy/strategy game from 3DO.
- Quake III Arena—A First Person Shooter (FPS) from ID Software and Activision.

While Loki did not port Quake, they provided support for the Linux version.

- Heavy Gear II—An FPS by Activision.
- SimCity 3000—The latest version of this classic simulation game from Maxis.

Loki is using open source development tools to create its games. If the needed tools weren't available, Loki developed their own tools and released many of them back to the open source community, including:

- OpenAL (<http://www.openal.org>)—A cross-platform 3-D audio library.
- SDL MPEG Player Library, SMPEG (<http://www.lokigames.com/development/smpeg.php3>)—A general-purpose MPEG video/audio player for Linux.
- Fenris Online Bugtracking System (<http://www.lokigames.com/development/fenris.php3>)—A Web-based, database-supported system for developers to receive input from users regarding bugs in software. It is based on the work of the Bugzilla project (<http://www.mozilla.org/bugs/>).
- SDL Motion JPEG Library, or SMJPEG (<http://www.lokigames.com/development/smjpeg.php3>)—A specialized motion JPEG and audio editing tool for Linux that uses a custom open format.
- Setup 1.0 (<http://www.lokigames.com/development/setup.php3>)—An XML- and GTK-based graphic installer utility. This provides a easy-to-use installer interface for the end-user.

Loki has contributed several enhancements to the gcc compiler, including a fix for the destruction of static objects in shared libraries. The gcc version 3.0 compiler promises to add C++ enhancements such as better debugging tools and an incremental linker. It also promises support for Microsoft C++ extensions.

The latest version of X11 for Linux (XFree86 version 4) offers 3-D acceleration and a more modular design. The modular design will make it easier for vendors to write drivers for Linux (<http://www.xfree86.org>).

There is also the 3-D API OpenGL, which was developed by Silicon Graphics Incorporated (SGI). SGI has long been the preferred platform on which to develop 3-D graphics. It was used for movies such as Jurassic Park and Terminator 2. OpenGL is similar to DirectX in that it gives a standard interface for rendering graphics on different video cards. OpenGL is currently a work in progress for Linux. Information and binaries for Linux are available at <http://oss.sgi.com/projects/>. General information, tools, and programs are available at <http://www.opengl.org/>.

Mesa is another popular 3-D API for Linux. It was originally developed as part of the SSEC Visualization Project at the University of Wisconsin (<http://www.ssec.wisc.edu/~billh/vis.html>). This project's goal was to create scientific modeling programs. Brian Paul's work on this project was later released as Mesa.

Mesa is very similar to OpenGL. It is so similar that most applications written for OpenGL can use Mesa with little or no modification. Many open source programs, including WINE, use the Mesa libraries.

Mesa supports many of the popular 3-D video cards, including 3dfx Voodoo, NVIDIA, Matrox MGA-G200 and MGA-G400, ATI RagePro, Intel i810, NVIDIA Riva, S3 ViRGE, and DRI-based hardware. Information and binaries are available at <http://mesa3d.sourceforge.net/>.

8.1.2 id Software

Loki may be currently pushing games into the Linux arena, but they are certainly not the only ones. id Software was one of the first commercial gaming companies to port their games to Linux. This is not too surprising, since their early games were developed on a Next system, which runs a version of UNIX.

id Software was formed when three game developers, John Carmack, Adrian Carmack, and John Romero, left their jobs at Softdisk publishing after creating the game Invasion of the Vorticons. After over a year of work, they released their first game, Wolfenstein 3D, in 1992. It was a popular FPS where you rescued a woman from Nazis. It laid the

groundwork for many future FPS—3-D graphics, rapid shooting, and the ubiquitous weapon at the bottom of the screen.

In 1993, id released Doom, which is probably the most popular game of all time. It fixed the problem with Wolfenstein that only allowed fixed ceiling height. This gave the game a true 3-D feel by allowing the gamer to look up and down. It also changed the theme to hellish-looking caverns and demon-like monsters. They followed it up with Doom II in 1994.

In 1996, two important events happened. id released Quake, the successor to Doom. It sported better graphics and better multi-player gaming. They also released the source code to Wolfenstein and Doom. In early 1999, id Software released the source code to two other popular games based on Doom: Heretic and Hexen. In late 1999, they re-released Doom under GPL, which would allow others to legally sell games based on the Doom source code.

id continued to improve Quake with the release of Quake II and Quake III. They improved not only the graphics, but the game play as well. In single-player mode, the enemy stood still and only attacked when you approached it. In Quake, the enemy could dodge shots and actively attack.

When Quake III was released in early 1999, it was became the first time a commercial game company released the DOS/Windows, Linux, and MacIntosh versions of their software at the same time (<http://www.idsoftware.com>).

8.1.3 Unreal Tournament

id Software is not the only FPS for Linux, though. Unreal Tournament (UT) puts out an unsupported Linux version. UT is the current state-of-the-art in FPS games. The stated system requirements are a 200MHz processor and 32MB of RAM. Realistically, you need at least a 400MHz processor, 64MB of RAM or better, and a 3dfx-compatible video card. Other 3-D video cards are not supported.

To install UT, you must buy the Windows version and download the Linux binaries from <http://fileplanet.com/index.asp?search=Linux+Unreal+Tournament&file=3205822>. Once the file is downloaded, extract the program as follows:

```
gunzip UT-Linux-400A.tar.gz
tar xvf UT-Linux-400A.tar
```

Be aware that UT requires about 500MB of free disk space. You will also need to have `libxml` and `libglade`. These libraries are included in the download if you don't have them installed already. Also, don't install the game as root, or only root will be able to play the game. Additionally, make sure that the game CD-ROM is mounted as `/mnt/cdrom` or the install program won't find it.

The installation directory is `UT-Linux-400A`. In this directory is the install script `setup.sh`. Running this program inside of X11 is supposed to start the setup program. If the setup script doesn't work, you can start it manually by typing `./setup.data/bin/x86/glibc-2.1/setup.gtk`

After the installation, you can modify the settings with `UnrealTournament.ini` and `User.ini`. Instructions for modifying these files are in the README file located in the installation directory. One glitch in the Linux version is that it doesn't work with the Enlightenment window manager. Since this is the default window manager for Gnome, it is better to use KDE or FVWM when using UT. See the later section on windows managers for details on how to switch.

Game play for UT is highly configurable. It allows you to play both multi-player against other players and in single-player mode against configurable "bots." You can configure both the number of bots and the skill of the bots.

UT includes 60 levels of game play, with many more available for download on the Internet. There are also many UT fan and server sites. A good place to start is at <http://www.unreal.com/index2.html>.

There are also several different variations of the game: Death Match, Capture the Flag (CTF), Domination, Last Man Standing, and Assault. Death Match is the classic play version for multi-player FPS games. Points are given for each kill and subtracted each time you are killed. The person with the most points wins. Capture The Flag is exactly what the name says: The first team to capture the enemy's flag wins. Domination is another variant where teams control locations on a map. When a team controls a location, it gets a point until the other team takes control of the location. Last Man Standing is similar to Death Match, except each player is given a limited number of lives. As the name implies, the last man alive wins. Assault is similar to Capture The Flag, except that each team takes turns attacking or defending. The team that completes the attack in the least amount of time wins.

8.2 Game Servers and Extras

Linux is a popular platform to run game servers. This allows many users on a network to connect to one game and play against each other.

The Quake server is a free, unsupported program that comes in both Linux and Windows versions. Even though it hasn't been updated since 1998, it still works with the newer versions of Quake. Quake servers for several different platforms are available at <http://qwcentral.stomped.com/files/files.htm>. The manual for Quake server is available at <http://qwcentral.stomped.com/qwsvmanual/qwsv.htm>. The instructions are for the Windows version, but the switches are the same.

There are many extra levels and skins for the various versions of Doom and Quake. Links to the latest ones are available at <http://www.idsoftware.com/hot/index.html>.

UT includes a server with a full version of the game. To start it, go to Multiplayer -> Start New Internet Game. The complete instructions on configuring a UT Server are at <http://unreal.epicgames.com/utServers.htm>.

UT also comes with the nifty utility `ngStat`. This allows you to track scoring and statistics for both local and online games. This program collects statistics from the UT log file and displays them inside either Netscape or Internet Explorer. `ngStat` can also be started inside of UT by selecting Stats -> View Global `ngWorldStats`. The latest information and updates for `ngStat` are at <http://www.NetGamesUSA.com/ngStats/UT/>.

If you don't want to run your own server, you can connect to the free UT servers at <http://www.mplayer.com> and <http://www.heat.net>. There is also an online version of `ngStat` at <http://UT.ngWorldStats.com/>.

Of course, UT has extra levels and skins also. Some good places to start looking for them are at <http://www.unrealfiles.com/> and <http://download.cnet.com/downloads/>.

8.3 Classic Games

Many classic games from the 1980s have gained a new lease on life under Linux. There are game emulators available for Amiga, Genesis, Coleco, Commodore, Atari, Gameboy, and most of the original video arcade systems.

Multi Arcade Machine Emulator (MAME) emulates over 2000 games, including most of the popular arcade games. The Linux version of MAME is called XMAME. The program and documentation are available at <http://x.mame.net/> and <http://www.vintagegaming.com/mame.phtml>. Games available for XMAME are in various states of development. For a list of games and the state of their development, go to <http://x.mame.net/download/gamelist.mame>.

Gameboy users are not left out. The games look like the originals, except that the screen is clearer and in color (<http://www.vintagegaming.com/>).

Chapter 9. The Linux Desktop

X11 is only the basic part of the Linux desktop. You also need a window manager for menus, icons, and other desktop features. Unlike Windows, where you have only a single interface, Linux has many choices for the desktop interface (i.e., window managers). The most popular ones are FVWM, KDE, and Gnome.

FVWM (<http://www.fvwm.org/>), which stands for F Virtual Windows Manager, is one of the oldest window managers. It is very stable and uses a relatively small amount of memory. AfterStep (<http://www.afterstep.org/>) is the variation of FVWM that is packaged with most Red Hat base distributions. FVWM is not nearly as friendly or as configurable as the other two popular desktops. It is a good choice for low-end systems, though, since it is much smaller and faster than the other desktops.

If you need a really light window manager, try *twm*. It doesn't look pretty, but it is really easy on the resources (<http://www.visi.com/~hawkeyd/vtwm.html>).

The K Desktop Environment (KDE) project was started in 1996 to provide a complete Integrated Desktop Environment (IDE) for Linux. It includes a window manager, file manager, help system, and utilities. It is the first program designed to give a modern, GUI interface to Linux. KDE is currently the default interface for most Linux distributions (<http://www.kde.org/>).

GNU Network Object Model Environment, or GNOME, (<http://www.gnome.org/>) was started because KDE uses the QT Library from Troll Tech (<http://www.trolltech.com/>). The QT Library is a commercial library and is not under the Gnu Public License (GPL). GNOME was started because many people in the Linux community feared that having commercial products in GPL software would endanger the GPL. While Troll Tech has allowed free use of the QT Library in KDE, GNOME provides an alternative in case things change.

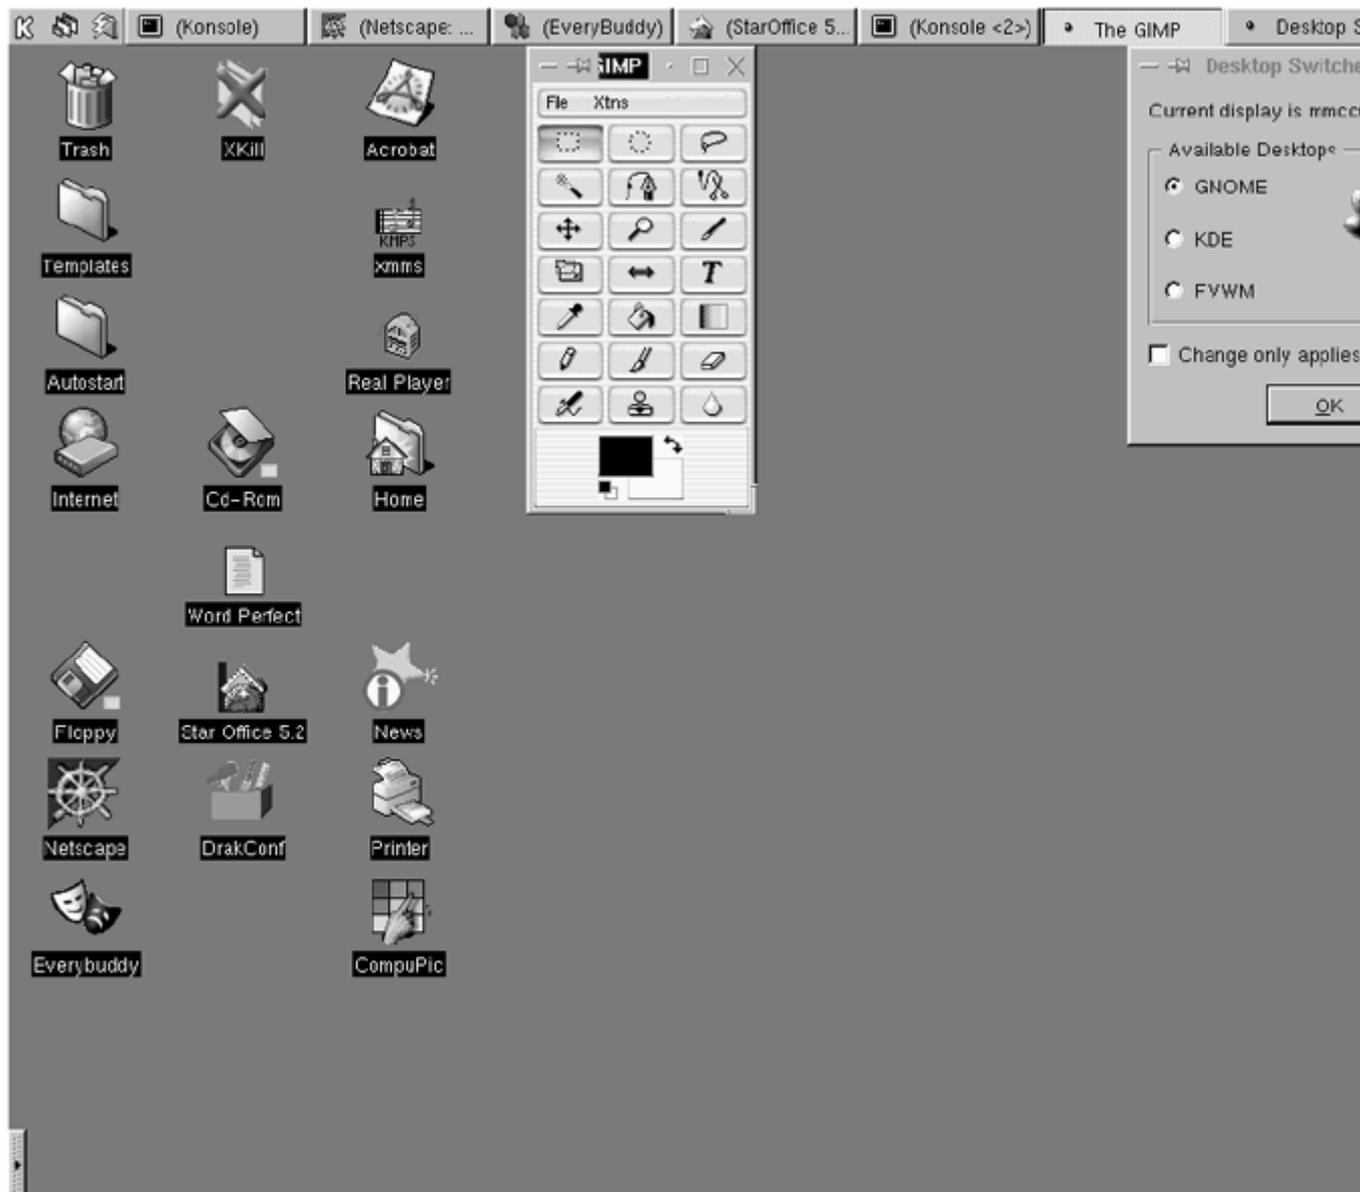
That being said, GNOME offers the richest, most configurable desktop currently available for Linux. GNOME even allows you to change window managers to change the look and feel of the desktop. Although the default window manager for GNOME is Enlightenment (<http://www.enlightenment.org/>), there are many other window managers available (<http://linux-directory.com/links/pages/WindowsManagers/>).

While GNOME is more configurable than FVWM or KDE, it also uses more resources, making it a poor choice for low-end systems. It also has been around less time than FVWM or KDE, so it has had less chance to be debugged.

9.1 Switching Desktops

Giving different desktops a try is relatively easy with a graphical utility called *switchdesk*, which allows the easy changing of desktops.

Figure 9.1. Switchdesk changes the default desktop.



As you can see, switchdesk offers Gnome, KDE, and FVWM. Simply choose one of these to change the desktop for all displays. Click the box Change only applies to current display to only change the desktop on the current display. Since Linux allows you to load more than one display, you could load several desktops at the same time!

Once you have the desktop selected, restart X11 and the new desktop will take effect. If you are using a graphical login, you don't need to use switchdesk—simply choose the window manager when you log on.

9.2 Configuring Desktops

Right-clicking on the desktop of KDE or Gnome will bring up a menu that has an option to change the desktop settings (FVWM doesn't support this). Below are illustrations for changing the setting for Gnome and KDE (see [Figures 9.2](#) and [9.3](#)).

Figure 9.2. Changing the desktop settings in KDE.

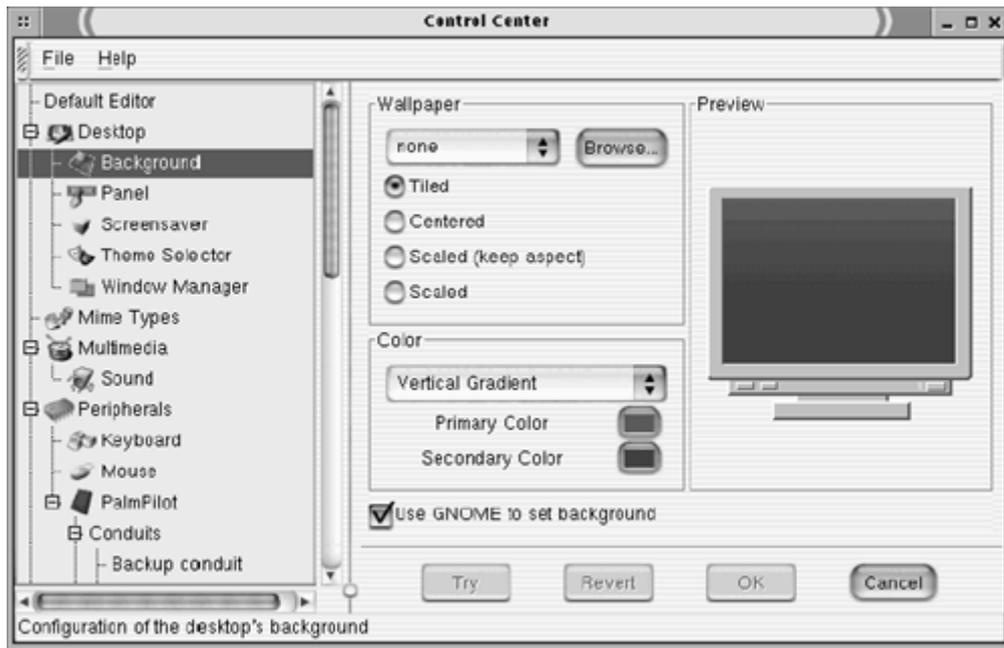
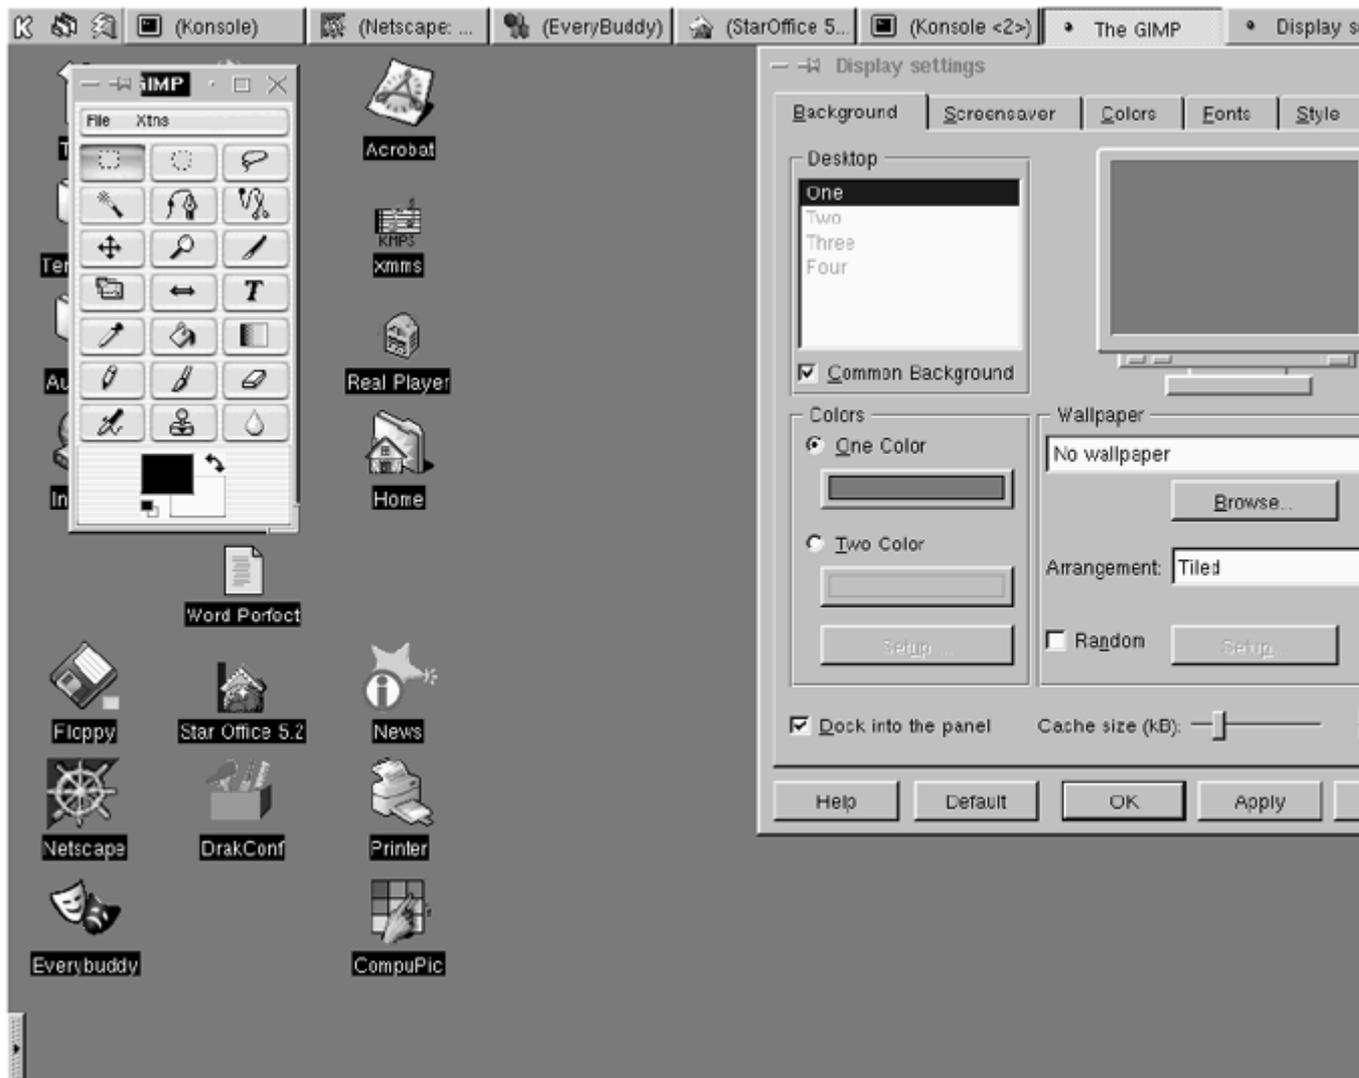


Figure 9.3. Changing the desktop settings in Gnome.



9.3 Themes

All desktop managers allow a great deal of customization of the desktop. For instance, you can download themes, which are preconfigured desktops. A good place to get these is from Themes.org (<http://www.themes.org/>).

To look at themes for your desktop, go to the link for your windows manager. Use Window Maker for FVWM, KDE for KDE, and GTK for Gnome.

From here you can browse the screenshots of different themes and download any that meet your fancy. After the theme is downloaded, it must be installed. Once you download a theme, it will look like `<theme name>.tar.gz`.

For FVWM, copy the theme to the theme file in the WindowMaker directory:

```
cp <theme name>.tar.gz ~/GNUstep/Library/WindowMaker
```

Next, go to the Window Maker directory and extract the files:

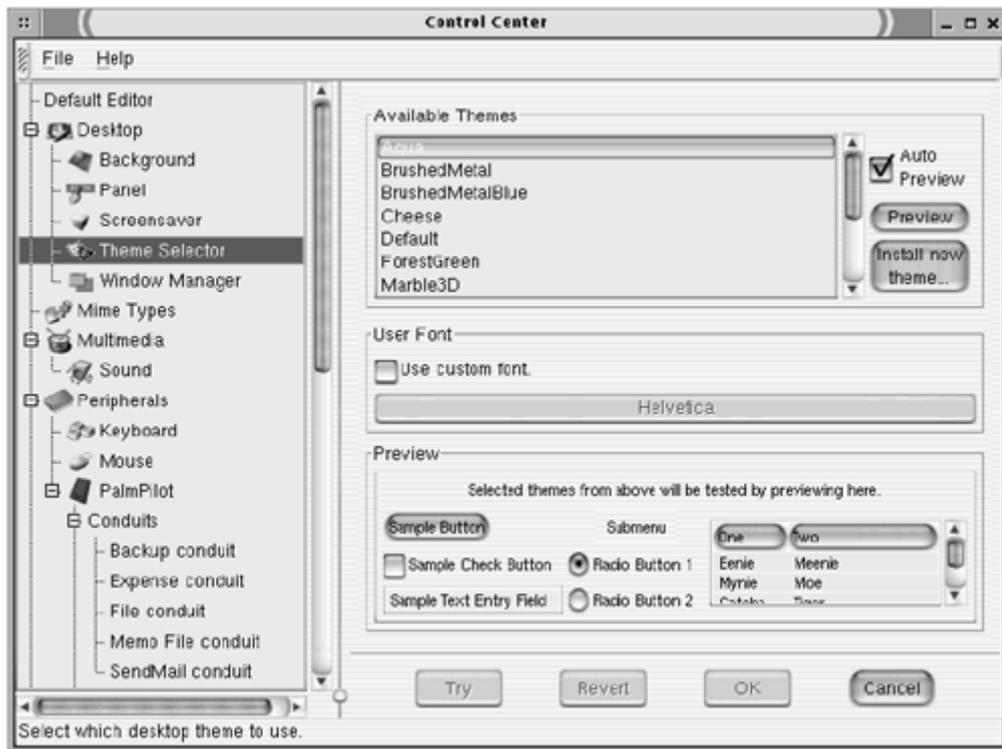
```
cd ~/GNUstep/Library/WindowMaker
gunzip <them name>.tar.gz
tar xvf <theme name>.tar
```

Inside `WindowMaker`, right-click on the desktop, go to Appearance -> Themes, and choose the theme you just installed.

Installing a theme under Gnome is a lot easier. First, go to the menu (the foot logo on the lower left) and choose Settings -> Desktop -> Theme Manager. This will open the Control Center. Click on a theme and it will preview the way the controls and background will look. Since Gnome is the most customizable interface, there are more themes available for Gnome. There are also many other settings in the Control Center that allow you to customize the interface even more.

To install a new theme under Gnome, simply click the button Install New Theme and choose the theme file.

Figure 9.4. Setting up desktop themes in Gnome.



To install a theme under KDE, go to the K menu at the lower left of the screen. Next, choose Settings -> Desktop -> Theme Manager. On the KDE Theme Manager, choose the Installer tab, then choose Add. Choose the theme file from the Add Theme box. Once the theme is selected, choose OK to effect the theme.

Figure 9.5. Opening KDE's Theme Manager.



Figure 9.6. Setting up desktop themes in KDE.



To create your own theme in KDE, click on Create and provide the theme name, author (you), email, Web address, and a description. If you want to share your theme with others, use the Save function to create a file out of it. The default name will be `<theme name>.tar.gz`.

9.4 Conclusion

While Linux isn't on par with Windows in the gaming market, Linux is one of the hottest areas of gaming right now. Linux should catch up or even exceed Windows in the gaming markets in the near future.

Chapter 10. Running Applications through a Network

[Section 10.1. X-Windows](#)

[Section 10.2. Citrix WinFrame](#)

[Section 10.3. VNC](#)

[Section 10.4. Conclusion](#)

10.1 X-Windows

X-Windows version 11 (X11) is the standard graphics interface for Linux and UNIX. It comes in two parts: a client and a server. The client runs the programs and the server runs the display.

The client and the server are often run on the same machine, which allows Linux to run graphical programs locally.

They can also be run on different machines, which allows you to run the programs on one machine and display them on another machine. The client would be running on the machine that has the programs and the server would be running on the machine displaying the screen. This is different from the common perception of the client and server, where the server contains the programs and the client runs them.

X11 splits the load between the client and the server. The client runs the application and the server processes the graphics. This takes some of the load off the server, but it also creates more network traffic. This was done in the early days of computing to allow (relatively) inexpensive workstations to run off a single, expensive server.

10.1.1 X-Windows on Linux

X11 comes with virtually every modern distribution of Linux. Normally, X11 is configured when Linux is installed. To configure X11, you need three things: a supported video card, the model and specifications of your monitor, and the type of your mouse. To find out whether your video card is supported, go to <http://www.xfree86.org/cardlist.html>. Monitor specifications are listed in the documentation that comes with the monitor. If you don't have this, most monitors are listed at <http://ms.ha.md.us/~hawks/hardware/monitor.html>. The Horizontal sync (Hsync), Vertical sync (Vsync), and Maximum Resolution (Max Res) are the most important settings. Your user manual should have a detailed description on how to install and configure X11. An abbreviated version of the configuration instructions are below.

To configure X11, use Xconfigurator. Some Linux distributions use XF86Setup. Most supported cards are found with the autoprobe option. If your card is not found, you can enter the driver and video memory manually. Next, choose the monitor you have. If your monitor is not listed, choose Custom and manually enter the Hsync, Vsync, and Max Res.

Xconfigurator should probe the resolutions and recommend a resolution. You can choose the default or pick your own. Next, X11 will start. If the screen is unreadable, you can always exit X11 by pressing <CTRL>-<ALT>-<BACKSPACE> and run Xconfigurator again.

Once the video card and monitor are configured, the only thing left to configure is the mouse. There are several brands listed, as well as generic serial port and PS2 port mice. You can also choose to emulate a three-button mouse, which emulates the middle button on a three-button mouse by pressing both keys on a two-button mouse. It is a good idea to emulate a three-button mouse, since Linux uses the middle mouse button to cut and paste.

Before you can run X11 on a remote host, you must add it to the server's xhost list. To do this, type `xhost <machine name or IP address>`. You can enable all connections by typing `xhost +`.

Once X11 is configured, it can be started locally with the `startx` script. To start an X11 session on another machine, type `xinit` followed by the machine's name, then the session number (starting with 0). For example, to start session 0 on machine `Server1`, type:

```
xinit server1:0
```

Consult a good manual or your Linux distribution's technical support for any advanced X11 problems. A good newsgroup for X11 problems is `news:comp.windows.x.i386unix`. Just remember the rules for getting help online: Read the manuals first and be specific about the problem.

Most newer distributions use `xdm` or a similar program (like `gdm`) to manage X11. If your system has a graphical login, you are using `xdm`.

`xdm` uses a hidden file in the home directory, `.Xauthority`, to provide security for an X11 session. If you are using X11 remotely, you must either copy `.Xauthority` to the remote machine or use `xhost` to add the machine to the access list (see above).

10.1.2 X-Windows on Microsoft Windows

X11 is included with Linux and UNIX, but it is an add-on program with Microsoft Windows. There are many packages that allow Windows to act as an X11 server, ranging from shareware to commercial programs.

The least expensive program is `MI/X`, which sells for \$25. It provides basic X11 server capability to Microsoft Windows, but it only supports revision 5 of X11, whereas Linux currently supports revision 6. This means that many newer programs will not be displayed properly with `MI/X`.

There are many other X11 servers for Windows, ranging all the way up to about \$600. There is a pretty complete list of them at <http://www.microimages.com/mix/prices.htm>.

10.2 Citrix WinFrame

10.2.1 Overview of the Server

Citrix WinFrame is a product that gives Windows NT the ability to act like X-Windows. It allows multiple users to start a session on an NT server and display the session on a workstation.

Unlike other Windows products, WinFrame allows a separate session for each user. Most other Windows products such as PC Anywhere or VNC will allow multiple users, but not multiple sessions. In other words, multiple users can log on, but they all see the same screen.

Microsoft Terminal Server was licensed from Citrix and provides the same functionality as Citrix WinFrame. Although there are subtle differences in the products, we will discuss Citrix WinFrame, since Citrix also makes a Linux client, which allows Linux machines to run Windows software on an NT server and display it on the Linux machine.

Because the applications are running entirely on the server, we make sure that the server can handle the load. Unlike X-Windows, all the processing is done on the server. This puts less stress on the network, but more stress on the server.

We will begin by figuring out the minimum configuration to handle the clients. Remember, this is the minimum; more is always better.

First, count the number of users and divide the users into normal users and power users. A normal user would only need to use one program at a time, whereas a power user would use several programs at once. Then, allocate 48MB of RAM for NT and WinFrame, plus 10MB for each normal user and 16MB for each power user.

For example, if you have 20 normal users and 10 power users, you would need 200MB for normal users (20 x 10), plus 160MB for power users (10 x 16), plus 48MB for NT and WinFrame, which totals 408MB of RAM.

CPU speed also makes a difference. It is a good idea not to put more than about 50 users per CPU. If you have more than 50 users, get a PC with multiple CPUs (SMP) or get an additional PC. If you are running multiple WinFrame servers, there is load balancing software available from Citrix for \$1,495.

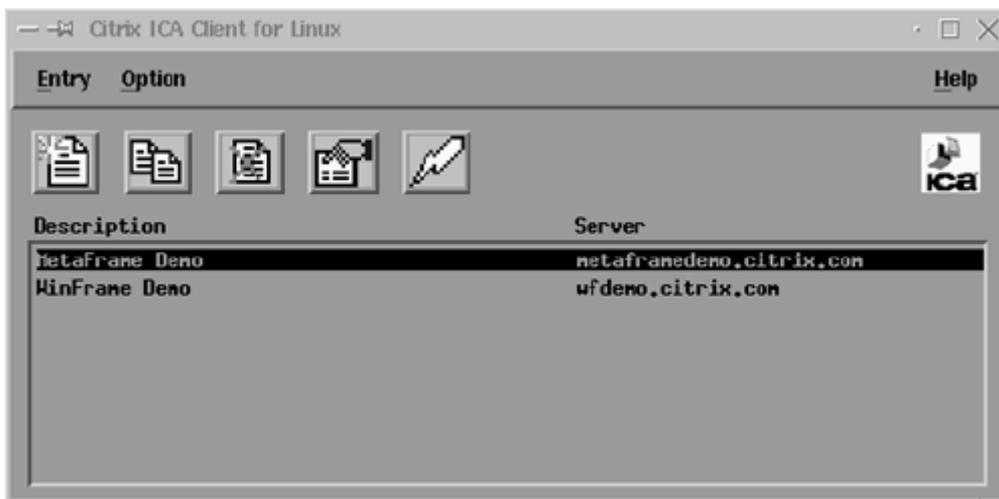
Keep in mind that X11 and WinFrame both require a lot of bandwidth. While a 10MB network will work for a very small network, a 100MB network is really needed. If this isn't enough bandwidth, consider installing multiple NT servers on a switched network.

10.2.2 Client Configuration

For the client side, you have three choices: You can use the Citrix-independent Client Architecture (ICA) client, the Java client, or X-Windows.

The ICA client allows the full functionality of the MetaFrame server. To install it, simply decompress the client file `linuxd.tar`. Next, run the setup program `setupwcf`. By default, it will install the client into the `/usr/lib/ICAClient` directory. To configure it, run the `wcfmgr` and use Entry -> New Option to set up the server to connect to. Supply the optional username, password, and domain. You can set up multiple servers, if needed.

Figure 10.1. The Citrix ICA client for Linux.



To open an ICA session, double-click on the session you wish to open. Enter a username and password, if needed. WinFrame will then open a window with the WinFrame session inside of it.

You can also use the Java client to load a session into your favorite browser. Currently, the only browser that supports the Java client properly is Netscape. When you connect to the server, a Java applet is downloaded and the WinFrame session starts inside your browser. This allows Linux machines without the client to run a WinFrame session, but it is much slower and not as stable. The user must wait for the Java applet to download, and once it is downloaded, it is not as fast as the ICA client. The Java client is also not as stable and it often crashes the browser.

WinFrame also supports X11 with the Citrix UNIX Integration Client. This is a package that loads on the server and acts as an X11 client. It allows the Linux workstation to connect to the WinFrame server using the X11 server that comes with Linux. For example, to start an X11 WinFrame session 0 on Server1, type:

```
xinit Server1:0
```

While this allows connection to WinFrame without installing the ICA client, many of the features are not supported by the UNIX Integration Client. For example, load balancing between multiple servers is not supported. Other features such as client drive mapping, application publishing, persistent cache, compression, COM port mapping, and local printing are not supported either.

Unless there is a good reason to do otherwise, it is best to install the ICA client on all Linux and Windows machines that need to attach to a WinFrame server. This allows the full functionality of the WinFrame services on all the clients.

10.2.3 Application Configuration

Since Windows is designed as a single-user OS, installing applications on a WinFrame server requires extra configuration to use them as multi-user applications. For example, with Microsoft Office, install the program on the server. Then, copy the user configuration files to the user's home directory on the server. These files include: `normal.dot`, `mailbox.pab`, `username.fav`, `username.prf`, `outlprnt`, `frmcache.dat`, `custom.dic`, and `pptools.ppa`. Next, edit the `HKEY_CURRENT_USER` tree of the session's Registry to change the location where it is looking for the configuration files. You will also have to disable the Fastfind feature, since Office will open a new Fastfind session for each user, which will quickly use up all the server's resources.

There are other problems unique to a WinFrame environment. For example, many functions require write access to the system root directory to work. The system root is the directory in which the Windows system files are installed, which is usually the `c:\winnt` or `c:\winframe` directory on the server. One function that requires write access to the system root is the ability to open embedded Web links in an Office document. If there is no write access to the system root, the user will get a message telling them to run `scandisk` or re-install Internet Explorer. As an administrator, you will need to decide whether using these functions is important enough to give the user full access to the system files.

This is not intended to be a full account of all the issues involved in setting up an application on a WinFrame server, but it should give you an idea of the complexity involved. On the positive side, once set up, managing a WinFrame server is a lot easier than managing multiple client machines.

10.2.4 Cost

While WinFrame is a good solution for running Windows applications, it is not cheap. The prices given are suggested retail prices. Your prices may vary.

The first expense is buying NT Server. This costs about \$1,200 for a ten-user license, plus about \$35 for each additional user. Then you must buy a copy of Windows for each PC used. For Linux clients, you must buy a Windows NT license (which costs about \$270) to use it on a Windows NT server.

The WinFrame or MetaFrame server costs \$5,000 for a 15-user server, with each additional license costing about \$200. The cost of applications can vary. The cost of the base version of Microsoft Office is \$280 per user.

So, the cost of a 50-user network would be \$2,600 for the NT Server with a 50-user license. The cost of 50 client licenses would be \$13,500, and a 50-user WinFrame license would be \$12,000. Finally, the cost of 50 copies of Microsoft Office would be \$14,000. This would make our hypothetical 50-user network cost \$42,100, plus the cost of the hardware. Luckily, the WinFrame client will run on even low-end 386 hardware, so a hardware upgrade is usually unnecessary. You would also save in support costs, since the centralized management makes support easier (<http://www.citrix.com>).

10.3 VNC

Virtual Network Computing (VNC) allows a program running on one machine to be displayed on another and it is available for Linux and Windows. It is distributed under the GNU Public License (GPL), so like Linux, it is free and open source. It is similar to X11 and Citrix except for a few differences.

VNC consists of two components: It has a server that runs the programs on one machine and a viewer that displays the programs on another machine. It is available for Linux, Windows, and most other operating systems. There are even viewers for Windows CE and Java. The viewer is very small (the Win32 viewer is 150k) and doesn't require installation, which means it can easily be run from a floppy!

There are some differences in the servers, though. The Linux and UNIX servers are multi-user. This means that more than one machine can open a separate session on the server. The Windows server, however, only allows a single session. It will allow multiple users to log into one session, however. Windows is not a true multi-user OS. If you need to open multiple sessions, use Citrix WinFrame.

One advantage VNC has over X11 is that no state is stored in the viewer. This is a fancy way of saying that if you lose your connection, you can reconnect without losing anything. This is a godsend over modem connections and unreliable networks.

VNC also uses less bandwidth than X11. There are several settings that can be used to optimize VNC for low- and high-bandwidth connections which will be discussed later. A properly optimized VNC session will even work well over a dial-up connection!

10.3.1 Installing VNC Server for Linux

To install VNC, download the packages from <http://www.uk.research.att.com/vnc>. The Linux installation comes with four programs:

- `Vncviewer`— This is the VNC viewer, or client, program for X-Windows.
- `Vncserver`— This is a wrapper script which makes starting an X-Windows VNC server (i.e., desktop) more convenient. It is written in Perl, so to use the script you need that.
- `Vncpasswd`— This program allows you to change the password used to access your X-Windows VNC desktops. The `vncserver` script uses this program when you first start a VNC server.
- `Xvnc`— This is the X-Windows VNC server. It is both an X-Windows server and a VNC server. You normally use the `vncserver` script to start `Xvnc`.

Copy these files to `/usr/local/bin` and you are ready to go. `vncserver` is a script file (written in Perl) that can be edited if needed. This would only be needed if you don't have the packages in the normal places. For example, you may need to edit the first line, `/usr/bin/perl`, if your Perl program is not in the `/usr/bin` directory. You may also need to add a font path and color database path to the section marked, appropriately, Add font path and color database.

Once this is done, start the VNC server by typing `vncserver`. The first time you run it, you will be prompted for a password. This is the password that must be entered at the viewer to connect to the VNC server. You can always change the password later with the `vncpasswd` program.

You can open as many VNC server sessions as needed. It will take the first available X11 session by default. The first display on X11 is 0, so the first VNC server session will be 1. You can also specify which session number to use. For example, to use session 2, type `vncserver :2`. You can kill the server by typing `vncserver -kill :<session number>`.

There are several command-line parameters for VNC server. These are shown by typing `vncserver -help`. Here are what the parameters mean (from the VNC documentation):

- `-name <name>`— Each desktop has a name, which may be displayed by the viewer. It defaults to "x," but you can change it with this option.
- `-geometry <widthxheight>`— Specifies the size of the desktop to be created. The default is 1024x768.
- `-depth <depth>`— Specifies the pixel depth in bits of the desktop to be created. The default is 8.
- `-pixelformat <format>`— Specifies the pixel format for the server to use (BGRnnn or RGBnnn).
- `-kill <number>`— Kills a specified session number.

You can also pass parameters to the `Xvnc` program from the command line. `Xvnc` is the actual VNC server program. `Xvnc` is very similar to X11 and even uses many of the same commands. The parameters that you can pass to `Xvnc` are (from the VNC documentation again):

- `-alwaysshared`— Always treat new clients as shared (i.e., ignore client's shared flag).
- `-nevershared`— Never treat new clients as shared (i.e., ignore client's shared flag).
- `-dontdisconnect`— Don't disconnect existing clients when a new "non-shared" connection comes in. Instead, the new connection is refused. New "shared" connections are still allowed in the normal way.
- `-localhost`— Only allow connections from the same machine. This is useful if you use SSH and want to stop non-SSH connections from any other hosts. See the guide to using VNC with SSH.
- `-cc n`— Sets the color of the visual class used by the server. Some X-Windows applications don't cope too well with the TrueColor visual normally used by an 8-bit-deep `Xvnc`. You can make the server use a PseudoColor visual by specifying `-cc 3`.
- `-economicstranslate`— The server normally uses a lookup table for translating pixel values when the viewer requests a different format from the native one used by the server. This can use up to 256KB per connected viewer, so if you have many viewers, you may wish to specify this option, which will save memory at the expense of a little bit of speed. This is only relevant for 16-bit-deep desktops.

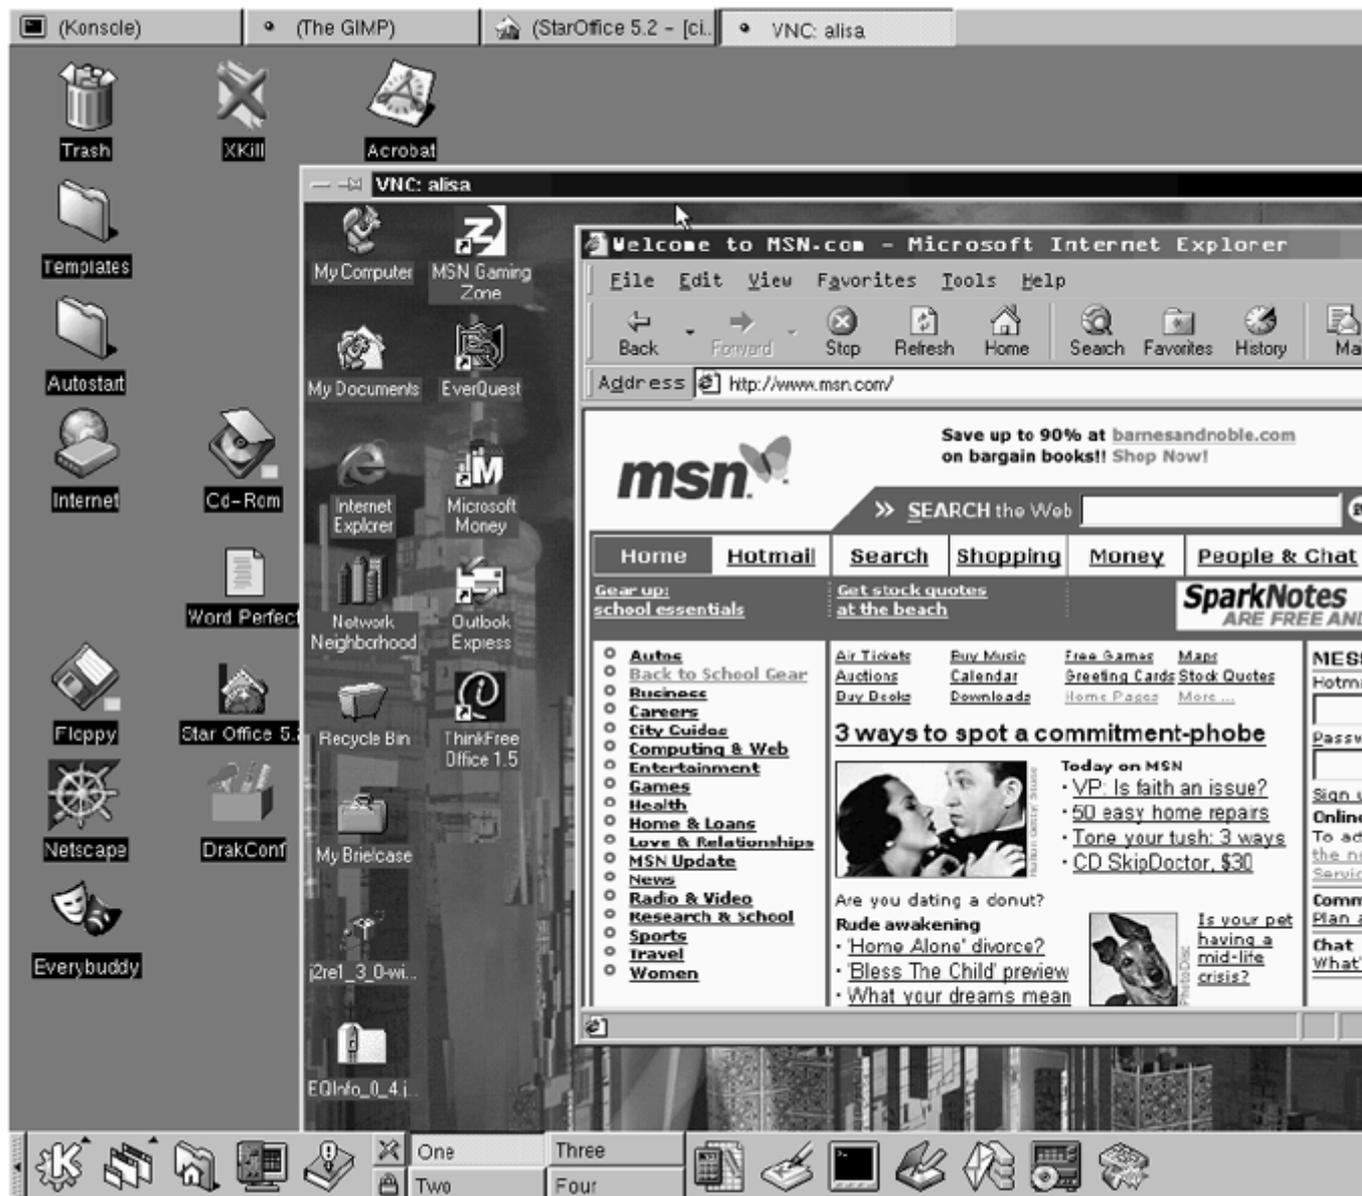
The server runs `~/ .vnc/xstartup` after it is started. This is where you would change your window manager. The default window manager is `twm`, which is small and fast, but very spartan. To change to KDE, for example, change the line `twm &` to `startkde &`.

The server also writes log files in the `~/ .vnc` directory. These files are useful in debugging VNC.

10.3.2 Using the VNC Viewer for Linux

To use the VNC viewer, type the following command: `vncviewer <machine name>: <session number>`. For example, to connect to session 2 on machine `mmccune`, type `vncviewer mmccune:2` You can also use the IP address instead of the machine name as in: `vncviewer 10.0.0.1:2` You will then be prompted for the password. Once it is entered, you should see the remote display as shown in [Figure 10.2](#).

Figure 10.2. Using VNC on a Linux PC to take over a Windows desktop.



Once the viewer is started, you can press `<F8>` to bring up a popup window. When this is up, you can send keystroke commands such as `<CTRL> <ALT> ` to the server. After a command is sent, the popup window will close. You can also use the popup window in and out of full-screen mode, transferring data to and from the Clipboard, and even quitting the viewer.

Full-screen mode allows remote windows to take up the entire screen, without the menus and scroll bars. This is particularly useful if the remote and local screens are set to the same resolution.

One note of caution on full-screen mode: Not all window managers support full-screen mode. If you need full-screen mode, it is better to start the viewer in full-screen mode.

There are many command-line options you can use to customize the VNC viewer. Typing `vncviewer -h` will bring up a list of these options. The options are as follows (from the documentation):

- `-shared`— When you make a connection to a VNC server, all other existing connections are normally closed. This option requests that they be left open, allowing you to share the desktop with someone already using it.
- `-display <Xdisplay>`— Specifies the X-Windows display on which the VNC viewer window should appear.
- `-passwd <password-file>`— If you are on a filesystem which gives you access to the password file used by the server, you can specify it here to avoid typing it in. It will usually be `~/.vnc/passwd`.
- `-viewonly`— Specifies that no keyboard or mouse events should be sent to the server. This is useful if you want to view the desktop without interference. It often needs to be combined with `-shared`.
- `-fullscreen`— Starts in full-screen mode.
- `-geometry <geometry>`— Standard X-Windows position and sizing specification.
- `-bgr233`— Tells the VNC server to send pixels which are only eight bits deep. If your server desktop is deeper than this, then it will translate the pixels before sending them. Less data will generally be sent over the network, which can be a big advantage on slow links, but the server may have to work a bit harder, and you may get some color mismatches. `bgr233` means an eight-bit TrueColorPixel format, with the most significant two bits of each byte representing the blue component, the next three bits representing green, and the least significant three representing red. This format is also used by the Java client.
- `-encodings <encodings>`— This option specifies a list of encodings to use in order of preference, separated by spaces. The default is `copyrect hextile corre rre or raw copyrect hextile corre rre` for a VNC server on the same machine. For example, to use only raw and CopyRect, specify `raw copyrect`.
- `-owncmap`— Tries to use a PseudoColor visual and private color map. This allows the VNC server to control the color map.
- `-truecolour`— Tries to use a TrueColor visual.
- `-depth <d>`— This is only useful on a (real) X-Windows server which supports multiple TrueColor depths. On such a display, the VNC viewer will try to find a visual of the given depth. If successful, the appropriate pixel format will be requested from the VNC server. You cannot use this to force a particular depth from the VNC server. The only option which does this is `-bgr233`.
- `-listen`— Causes the VNC viewer to listen on Port 5500+<display-number> for reverse connections from a VNC server. At present, WinVNC is the only server which supports reverse connections. This is initiated using the **Add**

New Client menu option. It is also used for our internal version of VNC (see <http://www.uk.research.att.com/vnc/internalversion.html>).

You can also use the VNC viewer command-line options to modify X11 resources. You can find the options by typing `vncviewer -xrm`. These options are rarely used, but they are there if needed (from the VNC documentation):

- `shareDesktop` (option `-shared`) — Whether or not to leave other viewers connected. The default is false.
- `viewOnly` (option `-viewonly`) — Blocks mouse and keyboard events. The default is false.
- `fullScreen` (option `-fullscreen`) — Full-screen mode. The default is false.
- `passwordFile` (option `-passwd`) — File from which to get the password (as generated by the `vncpasswd` program). The default is null, i.e., request a password from the user.
- `passwordDialog`— Whether to use a dialog box to get the password (true) or get it from the tty (false). Irrelevant if `passwordFile` is set. The default is false.
- `encodings` (option `-encodings`) — A list of encodings to use in order of preference, separated by spaces. The default is null, which actually means `copyrect hextile corre rre or raw copyrect hextile corre rre` for a VNC server on the same machine.
- `useBGR233` (option `-bgr233`)— Always use the BGR233 (eight-bit) pixel format on the wire, regardless of the visual. The default is false (though BGR233 is used anyway for non-TrueColor visuals with `forceOwnCmap` false).
- `nColours`— When using BGR233, try to allocate this many "exact" colors from the BGR233 color cube. When using a shared color map, setting this resource lower leaves more colors for other X-Windows clients. Irrelevant when using TrueColor. The default is 256 (i.e., all of them).
- `useSharedColours`— If the number of "exact" BGR233 colors successfully allocated is less than 256, then the rest are filled in using the "nearest" colors available. This resource says whether to only use the "exact" BGR233 colors for this purpose, or to use other clients' "shared" colors as well. The default is true (i.e., use other clients' colors).
- `forceOwnCmap` (option `-owncmap`) — Tries to use a PseudoColor visual and a private color map. This allows the VNC server to control the color map. The default is false.
- `forceTrueColour` (option `-truecolour`)— Tries to use a TrueColor visual. The default is false.
- `requestedDepth` (option `-depth`) — If `forceTrueColour` is true, tries to use a visual of this depth. The default is 0 (i.e., any depth).
- `useSharedMemory`— Whether or not to use the MIT shared memory extension if on the same machine as the X-Windows server. The default is true.

- `wmDecorationWidth`, `wmDecorationHeight`— The total width and height taken up by window manager decorations. This is used to calculate the maximum size of the VNC viewer window. The default is width 4, height 24.
- `bumpScrollTime`, `bumpScrollPixels` — When in full-screen mode and the VNC desktop is bigger than the X-Windows display, scrolling happens whenever the mouse hits the edge of the screen. The maximum speed of scrolling is `bumpScrollPixels` pixels every `bumpScrollTime` milliseconds. The actual speed of scrolling will be slower than this, of course, depending on how fast your machine is. The default is 20 pixels every 25 milliseconds.
- `popupButtonCount`— The number of buttons in the popup window. See below for how to customize the buttons.
- `rawDelay`— This is useful for debugging VNC servers by checking exactly which parts of the screen are being updated. For each update rectangle, the VNC viewer puts up a black rectangle for the given time in milliseconds before putting up the pixel data. This only highlights pixel data sent using the raw encoding. The default is 0 (i.e., don't do it).
- `copyRectDelay`— Similar to `rawDelay`, but highlights the areas copied using copyrect encoding.

10.3.3 Installing VNC Server for Windows

VNC server will work with any 32-bit version of Windows such as Windows 95/98, NT, and 2000. With NT, you must have Service Pack 3 or later installed.

First, extract the files to a directory on the local drive and run `setup`. After it is set up, install the default Registry settings using the Start -> Programs -> VNC -> Administrative Tools menu. This will modify some Registry setting to help VNC run better.

There is an icon in Administrative Tools that allows you to start the VNC server. You can run VNC as an application or a service. The application is in the Start -> Programs -> VNC folder.

It is better to run VNC as a service, though. To do this, go to the Administrative Tools menu and choose Install WinVNC service. You can also install it from the command line by typing:

```
c:
cd \Program Files\ORL\VNC
winvnc -install
```

The first time you run the server, it will ask for a password. This will be the password that everyone that attaches to the VNC server will use. There is currently no way to set up VNC server for multiple users or passwords.

Once the VNC server is running, it will add a VNC icon to the system tray. Right-clicking on this icon will display the following menu:

- **Properties**— This will cause the Properties dialog to be displayed, allowing the user to change various WinVNC parameters. See below for the different parameters.
- **Add New Client**— This allows outgoing connections to be made from the server to any "listening" viewer. The name of the target viewer machine can be entered

in the dialog. Connections created this way are treated as shared as of 3.3.3R2. See also the `-connect` option below.

- **Kill All Clients**— This will disconnect all currently connected clients from the server.
- **About WinVNC**— This should be obvious!
- **Close**— Shut down the server.

Moving the mouse over the icon will show the IP address of the local machine. Under Properties, you can find (from the VNC documentation):

- Incoming Connections:
 - **Accept Socket Connections**— The server normally accepts direct, socket-based connections from the `vncviewer` program. Clearing this check box disables direct connection to WinVNC, so that only the CORBA interface used by our internal version may be used to start a connection. For the public version, clearing this will disable any incoming connections.
 - **Display Number**— This allows the user to specify the display number which the server will use. There is normally no need to change this from the default of 0.
 - **Auto**— This check box indicates to WinVNC whether it should use the display number specified in the Display Number box or if it should use the first display number not already in use on the server machine.
 - **Password**— Incoming connections must be authenticated to verify that the person connecting is allowed to connect to this machine. This text box allows your password to be specified for authentication.
 - **Disable Remote Keyboard & Pointer**— Any new incoming connections will be able to view the screen but not send any input.
 - **Disable Local Keyboard & Pointer**— This is experimental and works on NT only. If selected, the local keyboard and mouse will be disabled during a connection. This can be useful if you want to log in to a machine from elsewhere and don't want passersby to be able to use your session.
- **Update Handling**— Use these settings only if your applications are having problems updating the screen. Changing these settings will generally slow down the screen updates. These settings will take effect the next time the VNC server is started (again from the documentation):
 - **Poll Full Screen**— Some applications are incompatible with the methods currently used in WinVNC to trap screen updates. For this reason, it is sometimes useful to be able to poll the entire screen to check for changes, sacrificing performance for accuracy.
 - **Poll Foreground Window**— Polling only the currently selected window for changes is less CPU-intensive than full-screen polling and often gives similar results; for example, when using the Command Prompt, which is not normally compatible with WinVNC.

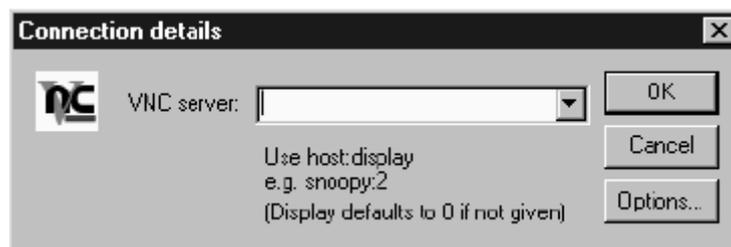
- **Poll Window Under Cursor**— A variation on Poll Foreground Window, this option causes the window under the mouse cursor to be polled for changes. Both options may be enabled simultaneously if required.
- **Poll Console Windows Only**— When this option is set, the only windows which will ever be polled are Command Prompts. This works well in conjunction with Poll Window Under Cursor, to use polling only when the cursor is over a console window.
- **Poll On Event Received Only**— When this option is set, the screen will only be polled for updates when a mouse or keyboard event is received from the remote client. This is provided for low-bandwidth networks, where it may be useful to control how often the screen is polled and changes sent.

The settings above should be enough for almost every situation. There are command-line options and Registry settings that can be set for special situations. The documentation at <http://www.uk.research.att.com/vnc/winvnc.html> explains these advanced settings.

10.3.4 Installing VNC Viewer for Windows

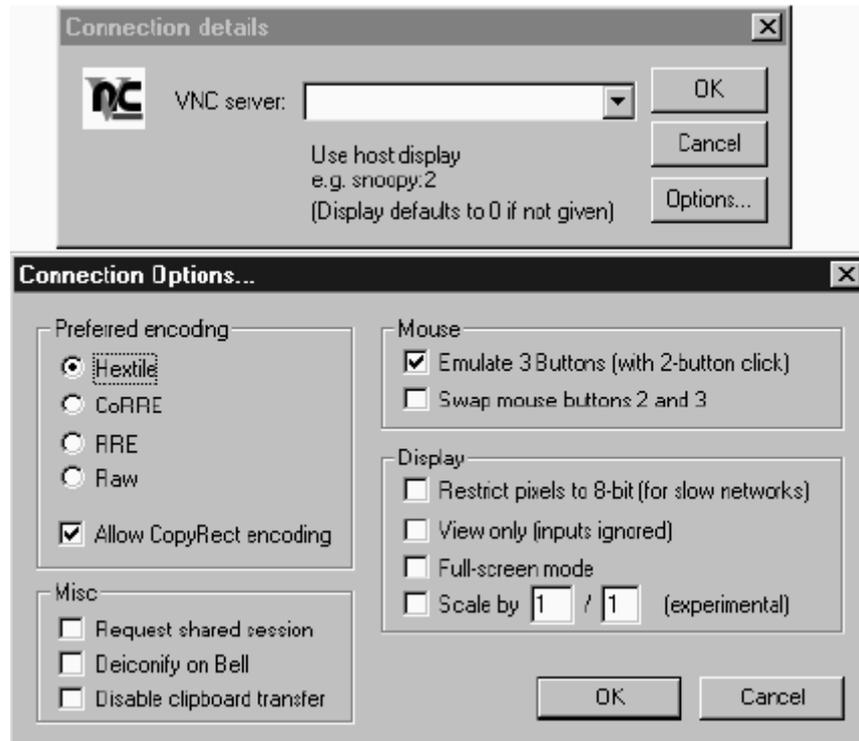
The VNC viewer for Windows can be started at the command line or from the menu as follows: Start -> Programs -> VNC -> Run vncviewer. When it is first started, you will get the following screen (see [Figure 10.3](#)):

Figure 10.3. When VNC for Windows is started, it asks for the server name.



Fill in the VNC server name (or IP address) and the display number. There is also an option box as shown in [Figure 10.4](#).

Figure 10.4. The VNC option box.



The first option is for encoding. The default settings work fine. You can experiment with these if you wish to see if the different encodings work faster.

The mouse options are to emulate a three-button mouse by pressing both keys on a two-button mouse and to swap buttons two and three. Turn on the three-button emulation if you have a two-button mouse. The middle key on a Linux system is the cut and paste key.

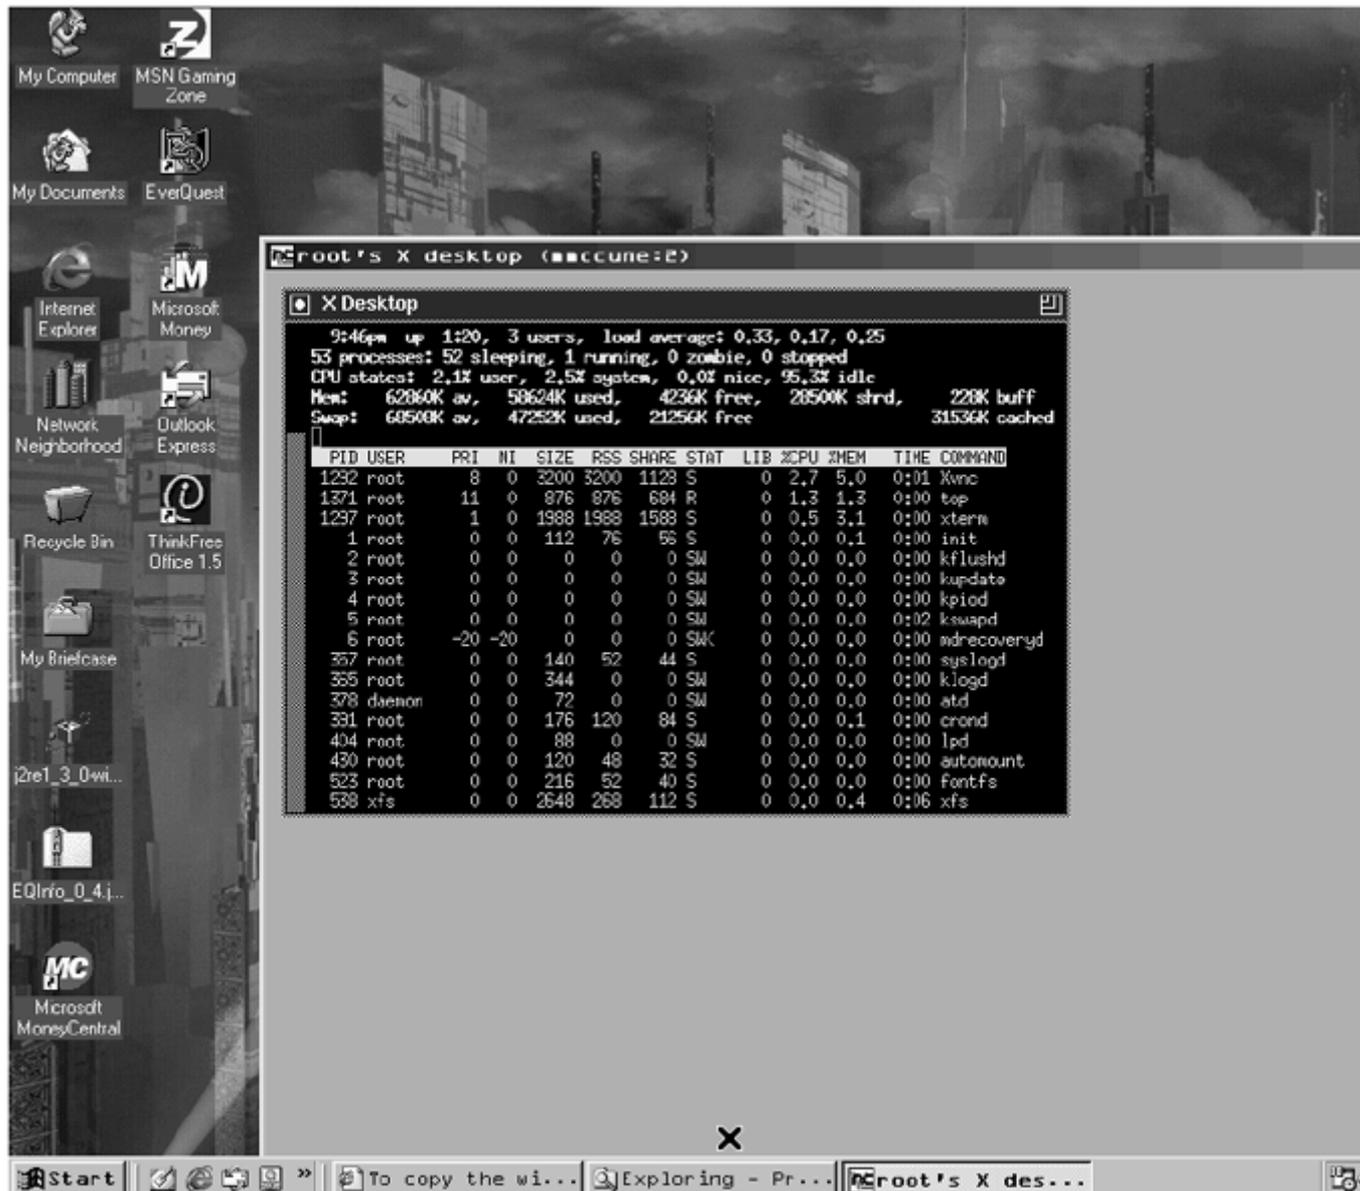
The display options are turned off by default. Turn on the 8-bit display (256 colors) if you have a slow connection (like a modem). The others are pretty much self-explanatory.

The **misc** options are self-explanatory except for the **Deiconify on Bell** option. This setting will allow the beep sounds to be displayed. It also will open the VNC viewer window when the bell sounds. This is a good setting to notify the user of mail or other important events.

To connect, enter the server name (or IP address) and session number, then click OK. You will be prompted for the password.

Once you are connected, you will see a desktop window with an xterm session. Click on the VNC logo at the upper left to bring up the window as shown in [Figure 10.5](#).

Figure 10.5. Using VNC for Windows to open a session on a remote Linux machine.



The menu options are self-explanatory. The connection options are the same as above. One also allows you to send keystrokes to the server. There are also options that allow you to open a new connection or to save the connection information as a file. This will save all the current settings so that they can be opened in another session. Once you get the settings the way you want them, use this option to save them.

Like the Linux viewer, the Windows viewer also has many command-line options. Here they are (from the VNC documentation):

- `-shared`— When you make a connection to a VNC server, all other existing connections are normally closed. This is for security reasons, and because we normally think of VNC as a tool for mobility; your desktop follows you from place to place. This option asks the server to leave any existing connections open, allowing you to share the desktop with someone already using it. Some servers have options to change the default behavior and override this request.
- `-8bit`— The viewer will normally accept whatever pixel format the server offers and do the translation locally. This forces it to request eight-bit True Color (BGR233) from the server, which will reduce network traffic. This is useful over modems.

- `-config file`— You can save all the details of an open connection to a file using a command from the menu. You can then restart that connection at a later date by specifying the name of the file using this switch.
- `-register`— This tells the Windows shell that `.vnc` files are associated with the VNC viewer. You should then be able to double-click on them to start the session. Sometimes Windows seems to need restarting before this takes effect.
- `-scale n/mw`— Specifies a scaling factor for the local display. The values `n` and `m` should be integers. The `/` and `m` can be omitted if `m=1`.
- `-emulate3`— Users with a two-button mouse can emulate a middle button by pressing both buttons at once if this option is enabled on the command line or in the dialog box. Note: On recent versions of the viewer, this is the default, so there's now a `-noemulate3` option to turn it off if wanted.
- `-noemulate3`— Opposite of `-emulate3`.
- `-swapmouse`— This option was more commonly used before three-button emulation was available. Normally, the PC buttons left-middle-right are mapped onto X buttons 1, 2, and 3. This switch causes the buttons to be mapped onto 1, 3, and 2, which may be more useful for two-button users who only have left-right, because they will then get buttons 1 and 2 instead of 1 and 3. If combined with three-button emulation, this also causes the middle button to emulate button 3 instead of button 2. This may be useful if you use button 2 more often.
- `-emulate3timeout`— When using three-button emulation, both mouse buttons must be pressed within a certain period for them to be registered as a single middle-click instead of separate left and right clicks. This option allows that time period to be specified in msec. The default is 100.
- `-emulate3fuzz`— When using three-button emulation, both mouse buttons must be pressed within a certain distance of each other for them to be registered as a single middle-click. This option allows that distance to be specified in pixels. The default is 4.
- `-fullscreen`— This causes connections to start in full-screen mode by default. See below for more details.
- `-listen`— In the internal version of VNC used at AT&T Labs Cambridge, the server can initiate connections to the clients under CORBA control. This switch puts VNC viewer into listening mode where it can accept these connections, but it also has a useful side-effect which may be of interest to those outside AT&T using the public version. A listening VNC viewer does not pop up a connection dialog, but instead installs itself in the system tray. From there, you can easily start up new connections and set default options to be used for them during this instance of the program. Additionally, the latest versions of WinVNC can initiate the connection to a viewer using the Add New Client menu option. For this to work, the viewer must be in listening mode.
- `-disableclipboard`— Clipboard changes caused by cutting or copying at either the viewer or server end are normally transmitted to the other end. This option disables Clipboard transfers.
- `-belldeiconify`— VNC allows for the transmission of a "bell" character, causing a beep at the viewer if it has sound facilities. You can set the sound to be

used for the bell under the VNCviewer section of Sounds in the Control Panel. Often a beep will happen because you are being notified of something such as an email arriving or a compilation finishing. This switch causes a minimized VNC viewer to be unminimized when a bell character is received.

- `-nocursor`, `-dotcursor`, `-normalcursor`— Most VNC servers send their cursor as part of the screen image that is displayed in the viewer. Having a local cursor in addition to this can be distracting. The default is for the viewer to use a small dot to show the position of the local cursor, and this is our recommended mode of use. You can use the `-nocursor` option to turn off this local cursor completely, or `-normalcursor` to leave it at the default Windows "arrow." Some things to note here: When you press a mouse button, it is the local mouse position that is used to send the event. On a slow network, the remote cursor may lag behind the local one a bit. You don't need to wait for it to catch up before you click, but if you have switched off the local cursor display, it can be harder to know exactly where you're clicking! The X-based server has an option which tells it not to show a cursor. This can be useful if combined with `-normalcursor` at the viewer, particularly on slow networks. However, the cursor will then never change shape—it will always be the arrow. We like the default dot the best.
- `-keyboard <kbdname>`— Windows uses an internal and not very helpful name for the keyboard layout currently selected for an application. You can see the one being used by VNC viewer if you select Connection Info from the system menu of the viewer window. If you change the keyboard settings and then make a note of this, you can specify it on the command line to cause VNC viewer to attempt to load this in the future. Note that VNC viewer does not currently support "dead keys", the differences between language and keyboard are confusing, and the way they are handled is different in Windows 95 and NT. But this may help a bit.
- `-logfile <filename>`— VNC viewer (R6 and later) has a logging mechanism which can save some debugging information to a file or display it on a console. This option specifies the name of a file to which a log will be written.
- `-loglevel <n>`— This option controls the amount of logging information sent to the log file. The default is 0, and higher values (up to about 12) will provide more detail.
- `-console`— In addition to, or instead of, logging to a file, this option will cause the debugging information to be sent to a console window.
- `-viewonly`— In view-only mode, no mouse or keyboard events will be sent back to the server. This is useful for teaching sessions or in other situations where you want to observe but don't want to interfere.
- `-restricted`— In restricted mode, most of the items are removed from the menu so that the user cannot, for example, send a `<CTRL>-<ALT>-` to the remote end.

10.3.5 Java VNC Viewer

You don't even need a VNC viewer if you have a Java-capable browser. Netscape, Internet Explorer, and many other browsers support Java. All VNC servers incorporate a small Web server that allows a Java applet to be served to the browser on the client. The

Web server runs on Port 58xx, where xx is the display number. For example, to connect to display 2 on server mmccune, connect to this Web page:

<http://mmccune:5802>

After the applet downloads, the browser will display a VNC session. If you are connecting to a Linux server, you will also have to put in the location of the class files in the `vncserver` script. The Windows server has this built into the server.

On the downside, the Java viewer isn't as fast or as stable as the stand-alone viewer. You can experiment with different encoding schemes to see if they make any difference in speed.

10.3.6 Optimizing VNC

The Linux VNC server is much faster than the Windows version. This is because access to the source code allows the Linux version of VNC to be optimized better. Windows doesn't have readily available source code, so it is less optimized.

There are some optimizations that can be done on VNC server, though. If Windows VNC server is running at near 100% load, check to see if the Poll Full Screen, Poll Foreground, or Poll Windows Under Cursor is on. Turn these off if they are not needed. If you must have any of these turned on, you can turn on Poll Console Windows Only to decrease CPU usage.

There are also some optimizations that can be done on all VNC systems. If VNC is running too slowly, try some of these steps to increase performance:

- **Clean up the desktop**— A busy desktop with lots of color and graphics will take longer to update. Remove any backgrounds or extra graphics that are not needed.
- **Reduce the resolution**— Lower the resolution of the desktop. Don't use a 1280 x 1040, 24-bit desktop when a 16- or 8-bit 800 x 600 will do. This is especially important on a modem connection.
- **Upgrade the card**— An upgrade of an older graphics card on the viewer machine will greatly increase the screen updates.
- **Upgrade the applications**— Some older applications are slow at screen updates. For example, some older versions of Netscape update the screen twice when scrolling.
- **Optimize Java**— Some settings can greatly increase the speed of the screen updates. Try different settings to find the best ones.
- **Turn off outline dragging**— When using the Windows viewer, turn this off in your window manager setup. Most of the time, this is already turned off. If it is enabled, see your Linux documentation for how to do this, since it varies with different setups.
- **Use Secure Shell (SSH)**— SSH compresses the data and will increase the speed on slow connections. Setting up SSH is explained in [Chapter 16](#). Setting up VNC to work with SSH is explained later in this chapter.

10.3.7 VNC Security

VNC encrypts the password, but once the session is started, it is unencrypted. Even though it is unencrypted, it would still be harder to snoop on a VNC session than a telnet or X11 session. If you need more security, it is relatively easy to add since VNC uses a

single TCP/IP socket. There are several packages that can add security to VNC such as Secure Shell (SSH), SSLeay, SOCKS, and TCP Wrappers.

SSH can add encryption to almost any TCP/IP service. The installation of SSH is covered in an earlier chapter. It consists of a client and a server. To use VNC with SSH, run the SSH server (`sshd`) on the machine with the VNC server, start the SSH client on the machine with the VNC viewer, then start up the VNC viewer.

The syntax of the SSH client is rather complex until you break it down. The syntax to connect session 1 on `mmccune` would be:

```
ssh -L 5902:mmccune:5901 mmccune
```

This looks complicated until it is broken down. `ssh` is the SSH client, `-L` is the local machine, `5902` is the port number of the local machine, `mmccune:5901` is the port number on `mmccune` and finally, `mmccune` is the remote SSH host.

So, this would redirect the local Port 5902 to Port 5901 on `mmccune`. To start a secure VNC session, we would open the VNC client on Port 5202 (session 2) on the local machine. This would then be redirected to Port 5901 (session 1) on `mmccune`:

```
vncviewer localhost:2
```

SSH also supports compression if you add the `-c` option to the command line. This is especially important over slow connections such as modem connections. You can also add compression to the configuration options. These and other options are explained in the section on SSH. Another switch we will use later is `-g`, which allows remote hosts to connect to local port forwarding ports.

Now, let's do something more complicated with SSH. Let's say we have two secure networks connected by an insecure Internet connection. On each side of this network, we have Windows workstations. We are using a Linux machine running SSH on each side of the Internet connection. We are running VNC client on a Windows machine on the first network (Windows1) and connecting to a VNC server on a Windows machine on the second network (Windows2).

Windows1 -> Local network 1 -> Linux1 -> Internet Connection -> Linux2 -> Local Network 2 -> Windows2

This is how each machine would be set up:

Windows1— Runs `vncviewer Linux1 :1`.

Linux1— Runs `ssh -g -L 5901:windows2:5900 linux2`.

Linux2— Runs the SSH server, `sshd`.

Windows2— Runs `vncserver display 0`.

The VNC viewer on Windows1 connects to Linux1, which forwards the packets through SSH to Linux2. Linux2 then forwards the packets to Windows2.

This is similar to a Virtual Private Network (VPN), which simply sends encrypted packets across an insecure public network such as the Internet. It won't offer all the features of a VPN; using SSH will often do the job.

Ray Jones' SSLeay (SSL) adds SSL encryption code to VNC. His patches area is available at <http://web.mit.edu/thouis/vnc/>. To use this option, you must recompile the VNC code with the patches and have the SSL libraries installed. This code isn't well-tested, so it is not recommended for a production environment.

SOCKS allows VNC to work through a standard SOCKS port on a firewall. The patches for VNC are available at <http://www.uk.research.att.com/vnc/contrib/socks-patch.txt>. Again, this must be compiled into the code and you must also be sure that the firewall has its SOCKS port enabled. There is a script available at

<http://www.uk.research.att.com/vnc/contrib/rvnc.txt> which helps run VNC through a firewall.

TCP Wrappers allows a machine to restrict which IP addresses can connect to it. The binaries and source code are available at ftp://wik.res.cmu.edu/pub/vncip_bin.zip and ftp://wik.res.cmu.edu/pub/vncip_src.zip. After the files are installed, the `iplist.txt` file must be edited to allow or deny IP addresses. The format looks like the following:

```
allow IP address
deny IP address
deny all
```

The deny all will deny all hosts except the local host.

There are many other add-ons to VNC. For a complete list, go to <http://www.uk.research.att.com/vnc/extras.html>.

10.4 Conclusion

This chapter showed many ways to run a program on one machine and display it on another. This type of setup allows you to run programs on a different platform as well as centralize applications for easy maintenance and upgrades.

Chapter 11. Introduction to Windows and Linux Networking

[Section 11.1. Net BIOS](#)

[Section 11.2. TCP/IP and Active Directory](#)

[Section 11.3. Net BIOS over ICP/IP](#)

11.1 Net BIOS

Windows and Linux developed from different heritages and thus network in different ways. Windows owes its roots to DOS and uses NetBIOS as its native protocol. Windows 2000 is replacing NetBIOS with Active Directory, but Linux still uses the TCP/IP protocol that it inherited from its UNIX background.

Fortunately, both operating systems go to great lengths to assure that they can interoperate with other systems. All versions of Windows have TCP/IP support built into them, and Linux has support for NetBIOS. There are also several programs for both Windows and Linux that allow each system to emulate the network protocols of the other. NetBIOS is an abbreviation for Network Basic Input/Output System. Each device on the network is assigned a 15-byte name consisting of letters, numbers, and special characters (see [Appendix A](#) for rules on naming NetBIOS devices). There is also a 16th byte used for the "resource type." For example, you could name a Windows file server "Fileserver1." If you then wanted to make your "F:" drive attach to "public" on the server, you would use the command `net use F: \\Fileserver1\public`.

The earliest Microsoft implementation of NetBIOS is called NetBEUI. Like anything else, NetBEUI has its advantages and disadvantages. On the upside, it is very fast, since it is a relatively simple protocol. On the downside, it doesn't work well on networks larger than about 20 computers. There are several reasons for this. NetBEUI sends its packets to all the computers on the local network, so the traffic gets really heavy after about a dozen computers. It also is not routable, which means the protocol only works on a local network.

Windows networking has evolved over the years. The oldest method is the workgroup model. A workgroup is just a logical group of computers. In a workgroup, all devices can share resources such as directories and printers with other members of the workgroup. The sharing is controlled by each device or computer. This has the obvious limitation of not having central control of the network's resources. Later on, Microsoft corrected this with the domain model.

Like the workgroup model, the domain model is a logical grouping of network devices. The advantage of a domain is that there is one login to the domain. All network resources have central control. Each domain has a primary and backup domain controller. A domain controller's primary function is to keep track of usernames and passwords. The domain model used by Windows allows users to log into the network only once, but still allows them to use the resources of the entire network. Trust relationships can also be set up between several domains to allow resources to be shared between domains.

11.2 TCP/IP and Active Directory

Windows 2000 will replace NetBEUI and the domain model with Active Directory, which organizes network devices into a hierarchical tree. The best analogy is the organization of a business. At the top is the president, who has authority over everything. Below the president are several vice presidents, and below them are managers, then workers. An Active Directory tree works the same way (see [Figure 11.1](#)). It allows central control of

the network devices, and it also allows the network to be divided into a logical hierarchical structure.

Figure 11.1. A sample Active Directory tree. In this case, it is divided geographically.



Active Directory overcomes some of the disadvantages of NetBEUI. It doesn't rely on broadcast packets for most of its functions, so the network can handle more than 20 hosts. It also doesn't require WINS servers, which are discussed later. Since Active Directory uses TCP/IP, it is routable, so it can be used on large networks covering several locations.

TCP/IP is short for Transmission Control Protocol/Internet Protocol. It was originally developed on ARPANET (Advanced Research Projects Agency Network), the precursor to the Internet. Since most of the machines on ARPANET were running UNIX, TCP/IP was developed largely on UNIX machines. Unlike NetBIOS, TCP/IP is designed for large networks and is routable. Instead of using simple names, TCP/IP uses a 32-bit numeric value called an IP address. IP addresses are written as four sets of numbers between 0-255, separated by periods. An example of an IP address would be 10.23.64.1. The first number in the IP address represents the class of the network.

Networks are broken down into Class A, B, and C networks, depending on size. A Class A network can have up to 16,777,214 hosts, a Class B network can have 16,384 hosts, and a Class C network can have 254 hosts.

Certain numbers in the IP address are reserved for special purposes. The number 0 refers to the current network or host. The number 127 is called a loopback, which is used for diagnostic purposes, and the number 255 is used for broadcasting packets to the entire network.

The other part of the TCP/IP address is the subnet mask. This can be used to subnet or divide a single network into two or more networks. For example, in the above network of 10.23.64.1, a subnet mask of 255.255.0.0 would not divide the network, but a subnet mask of 255.255.240.0 would divide the network. Subnet masks are probably the most complex part of TCP/IP and it would take up too much time to explain them here. See [Appendix A](#) for a more detailed explanation of subnet masks as well as other TCP/IP issues.

IP addresses are friendly for machines, but not too friendly for users. Most IP devices also have a user-friendly name called a host name. Most of us have seen these host names as Web site addresses. For example, the Web site for IBM is <http://www.ibm.com>. The IP address associated with this Web site is 204.146.80.99. Like IP addresses, host names are hierarchical. The "www" represents the Web server, "IBM" represents the domain associated with the IBM Corporation, and "com" represents a commercial Internet site.

The process of turning a human-friendly name into a machine-friendly IP address is done using DNS (the Domain Name Service). DNS is simply a hierarchical system used to find the IP address associated with a host name. When a Linux machine is given a host name, it first looks into the file `/etc/hosts`, which is simply a text file listing host names and addresses. If it doesn't find it here, it queries the primary name server. A name server maintains a database of host names and IP addresses. If the host name is

still not found, the secondary name server and then the tertiary name server will be queried. If the host name is not found, an error message such as "unknown host" will be returned. This is a simplified description of DNS. It will be described in more detail later on.

The IP standard has changed over the years. The current version of IP is 4, which is abbreviated IPv4. The next version of IP is 6 (IPv6), and it has several improvements over IPv4. The most noticeable change is that IPv6 will have 128-byte instead of 32-byte addresses. The Internet is currently running out of IP addresses. The current 32-byte IP addresses will allow about four billion addresses, but 128-byte IP addresses will allow substantially more. IPv6 also has enhanced security and diagnostics built into it. See <http://www.ipv6.org> for the complete specifications of IPv6.

11.3 Net BIOS over ICP/IP

Of course, both Windows and Linux have many features and programs that allow them to interact with each other. Both Windows and Linux machines allow NetBIOS to be implemented over TCP/IP (called NetBT). This allows them to have a common protocol, but it also creates some unique problems. The main problem is that TCP/IP host names can't always be resolved to NetBIOS names. This problem is solved by a WINS (Windows Internet Names Server).

A WINS translates IP host names into NetBIOS names. When a device using NetBT logs into a network, it registers its NetBIOS name with the WINS. If the name is not in use, it is added to the WINS database. Then, if a network device can't find a NetBIOS name, it uses the WINS to look up the address associated with the name.

If the WINS fails to find it, the device can still use broadcasts to try to resolve the name. Most networks have a primary and secondary WINS. This is a simplified explanation of how a WINS works. A more detailed explanation is in Appendix A.

There are also programs that allow Windows computers to use NFS (Network File Services). NFS is the protocol used primarily by UNIX machines to share files over a network. It allows a directory on another machine to be "mounted" so that it appears to be a directory on the local machine.

LPR (Line Printer) allows UNIX to print to a remote machine. Windows has support for LPR built into Windows NT and 2000. There are also programs available for DOS and the other versions of Windows that allow them to use LPR also. NFS and LPR are covered in later chapters.

Linux also has a program that allows it to share directories and printers. Samba is a program that allows non-Windows systems to emulate SMB (Server Message Block), which is used by Windows for file and print sharing. While Samba doesn't support all Windows network functions, it can turn a Linux system into a stand-alone Windows file server, and it also allows Linux to use Windows files and printers. The next two chapters of this book are devoted to Samba, since it is one of the most important tools for connectivity between Windows and Linux.

Chapter 12. Introduction to Samba

[Section 12.1. How Samba Started](#)

[Section 12.2. How Samba Works](#)

12.1 How Samba Started

Samba is a set of programs that allow non-Windows machines to use Server Message Block (SMB) networking, which is the native networking protocol of Microsoft products such as Windows 98 and NT. Samba was started by Australian graduate student Andrew Tridgell. He needed a way to share files between his UNIX machine and a DOS PC.

He could have used UNIX's native networking, Network File Services (NFS, which is discussed later), but he already had NetBIOS running on the DOS PC. Since he couldn't load NetBIOS and NFS at the same time under DOS, he decided to reverse-engineer SMB and implement it on his UNIX machine, making his UNIX machine look like a PC server to the DOS PC.

He released his code on the Internet in 1992. After a few bug fixes, he put the project aside. Then, a few years later, he needed to link his Linux PC with his wife's Windows PC. Having no other choice, he tried his old code and, to his surprise, it worked.

Doing a little research, Andrew discovered that NetBIOS and SMB were at least nominally documented. As he worked to improve his program, he tried to think of a name for it. He entered "SMB" into a spell checker and the first word that came up was "Samba." As its use grew, Samba gathered a whole team of programmers. It is bundled with most distributions of Linux and many commercial versions of UNIX, such as SGI's (formerly Silicon Graphics Inc.) IRIX.

Besides UNIX, there are also versions of Samba for VMS, MVS, OS/2, Stratus-VOS, Amiga, and MPE/iX (these are just other operating systems). Samba has grown into a complete set of tools that implements most of the file and print sharing features of a Microsoft Windows NT Server.

Like Linux, Samba is under the GNU Public License (GPL). This means that the programs and source code are free to use and modify. The only stipulation is that any improvements and fixes to the program must be released back to the public. With millions of users worldwide, both Linux and Samba are improving rapidly due to the feedback and code improvements of many users.

12.2 How Samba Works

Samba uses NetBIOS over TCP/IP (NetBT) to deliver the four major SMB services: file and print services, authentication and authorization, name resolution, and service announcement (also called browsing). These are explained briefly for now and in greater detail later in the book.

File and print services are the most important of these. They allow Windows (and Samba) to share files and printers through the network. These services were the primary reasons for Local Area Networks (LANs) in the first place.

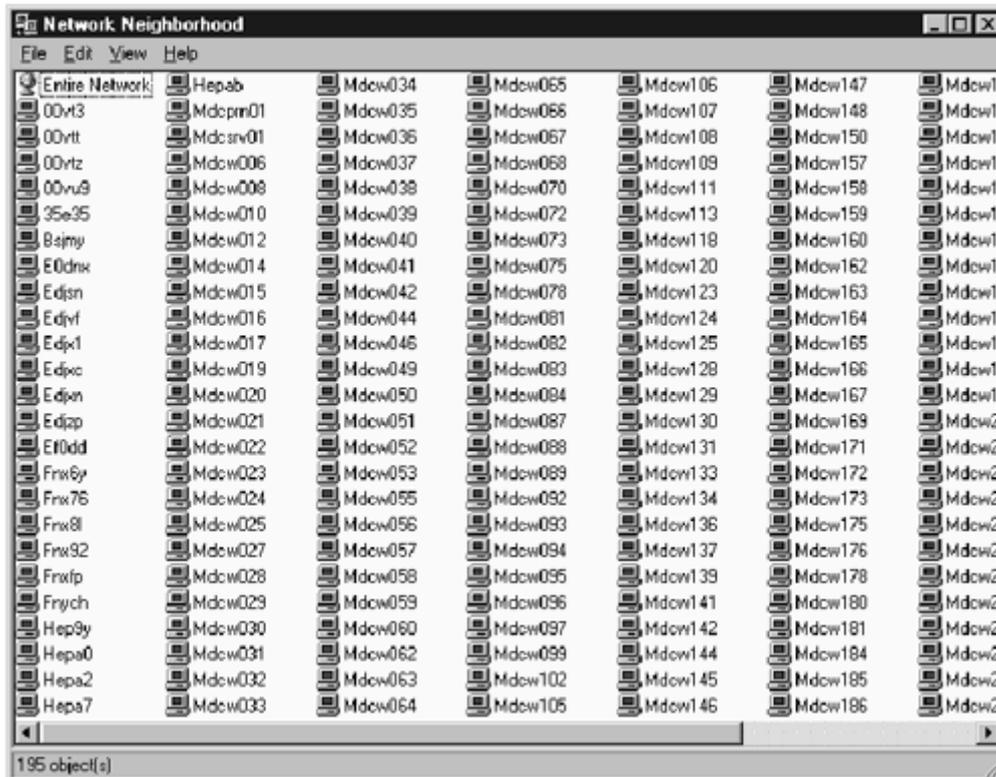
Authentication and authorization use login names and passwords to allow authorized users to access the network. They also restrict what each user can do on the network. Files and printers can be restricted by user. Files can also be made read-only to prevent accidental changes.

Name resolution consists of converting NetBIOS names into IP addresses (NetBT). This is done by the Windows Internet Name Server (WINS) on a Windows network. Samba can perform most of the functions performed by a WINS.

Service announcement, or browsing, allows other devices on the network to know what services are available. The list of services is kept on a machine called the "master

browser." The browsing list can be viewed by opening the "Network Neighborhood" (see [Figure 12.1](#)).

Figure 12.1. "Network Neighborhood" shows the SMB (Windows and Samba) machines on the network.



Samba can perform almost all of the functions of Microsoft Windows NT Server. So, when is Samba a good choice for a Windows network?

While Samba does support most the functions of a Windows Primary Domain Controller (PDC), Backup Domain Controller (BDC) functions are not supported yet. The only way to fully implement BDC services is to use Windows NT Server.

Also, if you already have a Microsoft WINS server, Samba cannot exchange data with it. So if you already have a Microsoft WINS server and need to add another, stick with Windows NT.

Both of these problems may soon be solved by Lightweight Directory Access Protocol (LDAP), which, among other things, allows one login for all a network's resources. Both Microsoft and the Linux community are developing LDAP functionality into their products. For more information on LDAP, see <http://www.openldap.org>.

So why would someone use Samba on a Windows network? Samba combined with Linux is fast, stable, and best of all, free. Windows NT charges per user, with a ten-user version of Windows NT Server costing about \$900. Linux and Samba not only cost less, but they run on lesser hardware. It's not unusual to set up Linux and Samba on an old 486 or low-end Pentium. This is more than enough to run a small, departmental server, while a Windows NT server would require a much better computer to run properly.

Samba is also a good choice for a mixed UNIX and Windows environment. Since Linux has all the standard UNIX networking tools included with it, Samba on a Linux server can communicate transparently with both UNIX and Windows.

With the introduction to Samba out of the way, the next chapter will cover how to set up Samba as a file server.

Chapter 13. Setting Up Samba as a Windows NT Server

Linux and Samba can perform many of the functions of a Windows NT server. The combination can be set up as a stand-alone server on a network or as part of an existing network. We will set up a Linux and Samba server, then discuss the advantages and limitations of using Linux and Samba vs. Windows NT.

13.1 Setting up Samba as a Stand-Alone Windows NT File Server

We will now put our knowledge to practical use. First we will set up Samba as a stand-alone server on a new network. Later on we will add a Samba server to an existing network.

There are IP addresses that are reserved for private networks (networks that are not directly connected to the Internet). See [Table 13.1](#) below for a list of addresses that can be used on private networks.

Class	IP Address	Subnet of Mask
A	10.X.X.X	255.0.0.0
B	172.16.X.X	255.255.0.0
C ^[1]	192.168.X.X	255.255.255.0

^[1] Windows 2000 also uses the class C network address 169.254.X.X for Automatic Private IP addressing.

Next, we need to decide whether to use fixed or dynamic IP addresses on our workstations. Fixed IP addresses are set up on each station individually. Dynamic IP addresses are pulled from the server using Dynamic Host Configuration Protocol (DHCP). Since it is easier to set up the IP configuration on the server rather than each individual workstation, we are going to use DHCP. If we decide to use fixed IP addresses, however, setting up DHCP would be unnecessary.

13.1.1 Setting Up DHCP

So after we install Linux on the server, we need to set up DHCP. The server's IP address is going to be 10.0.0.2 with a subnet mask of 255.255.255.0. We have a router installed at 10.0.0.1. A router simply allows a network to be divided up into sections (subnets). DHCPD, the server program for DHCP, is installed by most Linux distributions. If DHCPD does not come with your distribution, you can download it from Linuxberg <http://www.linuxberg.net/>. Just search for DHCPD. Documentation on DHCP is available at <http://www.linuxdoc.org/docs/ldp/howto/mini/dhcp>.

After DHCPD is installed, the configuration file `/etc/dhcpd.conf` must be created and edited. This file looks a lot like the `smb.conf` file. It has the global, subnet, shared network, and group sections.

The global section has several default values. These can be overridden by the other sections if necessary. We are not using all the values, so see Appendix A for other values. Notice that all lines need to have a semi colon (;) at the end of them.

```
server - identifier mmccune
option subnet - mask 255.255.255.0;
option broadcast - address 10.0.0.255;
```

```
default - lease - time 86400;
max - lease - time 259200;
```

There are two parameters that need to be explained: `default-lease-time` and `max-lease-time`. With DHCP, the IP address given out by the server expires after a given time. This prevents old workstations and other devices from using IP addresses after they are no longer being used. The `default-lease-time` is how long the device keeps the IP address in seconds. In this case, 86400 seconds equals one day (60 seconds x 60 minutes x 24 hours). The `max-lease-time` is how long the device has to renew the IP address. If it doesn't renew by this time, the device's IP address will be disabled and the IP address can be assigned to another device. In this example, the `max-lease-time` is three days (86400 seconds x 3 days).

The subnet section configures the IP subnet the DHCP server is covering. The `subnet` identification and `netmask` are followed by the options enclosed by brackets:

```
subnet 10.0.0.0 netmask 255.255.255.0
{
  range 10.0.0.5 10.0.0.250;
  option netbios - name - servers 10.0.0.2;
  default - lease - time 86400;
}
```

The `range` specifies the addresses that are available for DHCP lease. We have included addresses from 5 to 250. We have left 1-4 and 251-254 for fixed addresses for servers, routers, and other devices that need to have a fixed address. In this case, 1 is the router and 2 is the Samba server (10.0.0.1 and 10.0.0.2).

The `netbios-name-servers` parameter specifies the IP address of the WINS. In this case, we are using the Samba server (10.0.0.2) as the WINS.

The `default-lease-time` overrides the values in the global section. If we leave it blank, the global values will be used.

If we have several subnets, we will need to use the shared network section. For this example, we will use two subnets: 10.0.0.0 and 10.0.1.0.

```
shared - network
{
  subnet 10.0.0.1 netmask 255.255.255.0
  (
    range 10.0.0.5 10.0.0.250;
    option routers 10.0.0.1;
  )
  subnet 10.0.1.0 netmask 255.255.255.0
  {
    range 10.0.1.5 10.0.1.250;
    option routers 10.0.1.1;
  }
}
```

The options for the shared network section are divided into each subnet. The options in the subnet are set up exactly the same as in an individual subnet. In this case, we have a router set up to join the subnets with the two IP addresses 10.0.0.1 and 10.0.1.1.

The group section is for grouping clients together for a specific reason. For example, we could use it to group the Windows 98 workstations separately from the Windows NT workstations. However, since we don't need this for our network, we will not use the group section.

So here is the final `dhcpd.conf` file:

```
server - identifier mmccune
option subnet - mask 255.255.255.0;
```

```
option broadcast - address 10.0.0.255;
default - lease - time 86400;
max - lease - time 259200;
shared - network
{
  subnet 10.0.0.1 netmask 255.255.255.0
  {
    range 10.0.0.510.0.0.250;
    option routers 10.0.0.1;
  }
}
subnet 10.0.1.0 netmask 255.255.255.0
{
  range 10.0.1.510.0.1.250;
  option routers 10.0.1.1;
}
}
```

Next we will set up the `smb.conf` file. We will start with the default file and modify it. The lines starting with `#` are comments, and the lines starting with `;` are options that are not used:

```
# This is the main Samba configuration file. You should read
the
# smb.conf(5) manual page in order to understand the options
listed
# here. Samba has a huge number of configurable options
(perhaps too
# many!), most of which are not shown in this example.
#
# Any line which starts with a ; (semi-colon) or a # (hash)
# is a comment and is ignored. In this example, we will use
a #
# for commentary and a ; for parts of the config file that
you
# may wish to enable.
#
# NOTE: Whenever you modify this file, you should run the
command "testparm"
# to check that you have not made any basic syntactic
errors.
#
#===== Global
Settings=====
[global]
netbios name = mmccune
# workgroup = NT-Domain-Name or Workgroup-Name
  workgroup = WORKGROUP
# server string is the equivalent of the NT Description
field
  server string = Samba Server
# This option is important for security. It allows you to
restrict
# connections to machines which are on your local network.
The
# following example restricts access to two Class C networks
and
```

```
# the "loopback" interface. For more examples of the syntax,
see
# the smb.conf man page.
;   hosts allow = 192.168.1. 192.168.2. 127.
# If you want to automatically load your printer list rather
# than setting them up individually, then you'll need this.
    printcap name = /etc/printcap
    load printers = yes
# It should not be necessary to spell out the print system
type unless
# yours is non-standard. Currently supported print systems
include:
# bsd, sysv, plp, lprng, aix, hpux, and qnx.
;   printing = bsd
# Uncomment this if you want a guest account. You must add
this to /etc/passwd,
# otherwise the user "nobody" is used.
guest account = pcguest
# this tells Samba to use a separate log file for each
machine
# that connects.
    log file = /var/log/samba/log.%m
# Put a capping on the size of the log files (in Kb).
    max log size = 50
# Security mode. Most people will want user level security.
See
# security_level.txt for details.
    security = user
# Use password server option only with security = server
;   password server = <NT-Server-Name>
# Password level allows matching of n characters of the
password for
# all combinations of upper- and lower-case.
password level = 3
username level = 3
# You may wish to use password encryption. Please read
# ENCRYPTION.txt, Win95.txt, and WinNT.txt in the Samba
documentation.
# Do not enable this option unless you have read those
documents.
# By default, Windows 95 doesn't use encrypted passwords,
but Windows 98, NT,
# and 2000 use encrypted passwords.
encrypt passwords = yes
smb passwd file = /etc/smbpasswd
# The following are needed to allow password changing from
Windows to
# update the Linux system password also.
# NOTE: Use these with 'encrypt passwords' and 'smb passwd
file' above.
# NOTE2: You do NOT need these to allow workstations to
change only
```

```
#           the encrypted SMB passwords. They allow the UNIX
password
#           to be kept in sync with the SMB password.
;   unix password sync = Yes
;   passwd program = /usr/bin/passwd %u
;   passwd chat = *New*UNIX*password* %n\n
*ReType*new*UNIX*password* %n\n
*passwd:*all*authentication*tokens*updated*successfully*
# UNIX users can map to different SMB usernames.
username map = /etc/smbusers
# Using the following line enables you to customize your
configuration
# on a per machine basis. The %m gets replaced with the
NetBIOS name
# of the machine that is connecting.
#include = /etc/smb.conf.%m
# Most people will find that this option gives better
performance.
# See speed.txt and the manual pages for details.
    socket options = TCP_NODELAY SO_RCVBUF=8192
SO_SNDBUF=8192
# Configure Samba to use multiple interfaces.
# If you have multiple network interfaces, then you must
list them
# here. See the man page for details.
;   interfaces = 192.168.12.2/24 192.168.13.2/24
# Configure remote browse list synchronization here.
# Request announcement to, or browse list sync from:
#           a specific host or from / to a whole subnet (see
below).
;   remote browse sync = 192.168.3.25 192.168.5.255
# Cause this host to announce itself to local subnets here.
;   remote announce = 192.168.1.255 192.168.2.44
# Browser Control Options:
# Set local master to no if you don't want Samba to become a
master
# browser on your network. Otherwise, the normal election
rules apply.
;   local master = no
# OS level determines the precedence of this server in
master browser
# elections. The default value should be reasonable.
;   os level = 33
# Domain master specifies Samba to be the domain master
browser. This
# allows Samba to collate browse lists between subnets.
Don't use this
# if you already have a Windows NT domain controller doing
this job.
    domain master = yes
# Preferred master causes Samba to force a local browser
election on startup
```

```
# and gives it a slightly higher chance of winning the
election.
    preferred master = yes
# Use only if you have an NT server on your network that has
been
# configured at install time to be a primary domain
controller.
;   domain controller = <NT-Domain-Controller-SMBName>
# Enable this if you want Samba to be a domain logon server
for
# Windows 95 workstations.
;   domain logons = yes
# If you enable domain logons, then you may want a per-
machine or
# per-user logon script.
# Run a specific logon batch file per workstation (machine).
logon script = %m.bat
# run a specific logon batch file per username.
logon script = %U.bat
# Where to store roving profiles (only for Win95 and WinNT).
#   %L substitutes for this server's NetBIOS name; %U
is username.
#   You must uncomment the [Profiles] share below.
logon path = \\%L\Profiles\%U
# All NetBIOS names must be resolved to IP addresses.
# 'name resolve order' allows the named resolution mechanism
to be specified.
# The default order is "host lmhosts wins bcast". "host"
means use the UNIX
# system gethostbyname() function call that will use
/etc/hosts, or
# DNS, or NIS, depending on the settings of
/etc/host.config, /etc/nsswitch.conf,
# and the /etc/resolv.conf file. "host" therefore is system
configuration-
# dependent. This parameter is most often of use to prevent
DNS lookups
# in order to resolve NetBIOS names to IP addresses. Use
with care!
# The example below excludes use of name resolution for
machines that are NOT
# on the local network segment
# - OR - are not deliberately to be known via lmhosts or via
WINS.
; name resolve order = wins lmhosts bcast
# Windows Internet Name Serving Support Section:
# WINS support - Tells the NMBD component of Samba to enable
its WINS Server
wins support = yes
# WINS Server - Tells the NMBD components of Samba to be a
WINS client.
#   Note: Samba can be either a WINS server or a WINS
client, but NOT both.
```

```
; wins server = w.x.y.z
# Only set this if you already have a WINS server on the
network.
# WINS proxy - Tells Samba to answer name resolution queries
on
# behalf of a non-WINS-capable client. For this to work,
there must be
# at least one WINS server on the network. The default is
no.
; wins proxy = yes
# DNS proxy - Tells Samba whether or not to try to resolve
NetBIOS names
# via DNS nslookups. The built-in default for version 1.9.17
is yes;
# this has been changed in version 1.9.18 to no.
    dns proxy = no
# Case preservation can be handy - system default is _no_
# NOTE: These can be set on a per-share basis.
; preserve case = no
; short preserve case = no
# Default case is normally upper-case for all DOS files.
; default case = lower
# Be very careful with case-sensitivity - it can break
things!
; case sensitive = no
#===== Share
Definitions=====
[homes]
    comment = Home Directories
    browseable = no
    writable = yes
# Un-comment the following and create the netlogon directory
for domain logons.
; [netlogon]
;     comment = Network Logon Service
;     path = /home/netlogon
;     guest ok = yes
;     writable = no
;     share modes = no
# Un-comment the following to provide a specific roving
profile share.
# The default is to use the user's home directory.
;[Profiles]
;     path = /home/profiles
;     browseable = no
;     guest ok = yes
# NOTE: If you have a BSD-style print system, there is no
need to
# specifically define each individual printer.
[printers]
    comment = All Printers
    path = /var/spool/samba
    browseable = no
```

```
# Set public = yes to allow user 'guest account' to print.
  guest ok = no
  writable = no
  printable = yes
# This one is useful for people to share files.
[tmp]
  comment = Temporary file space
  path = /tmp
  read only = no
  public = yes
# A publicly accessible directory, but read-only, except for
people in
# the "staff" group.
# I am going to use this directory to put programs such as
Office 2000 here.
[public]
  comment = Public Stuff
  path = /home/samba
  public = yes
  writable = yes
  printable = no
  write list = @staff
# Other examples:
#
# A private printer, usable only by Fred. Spool data will be
placed in Fred's
# home directory. Note that Fred must have write access to
the spool
directory,
# wherever it is.
;[fredsprn]
;  comment = Fred's Printer
;  valid users = fred
;  path = /homes/fred
;  printer = fred's_printer
;  public = no
;  writable = no
;  printable = yes
# A private directory, usable only by Fred. Note that Fred
requires write
# access to the directory.
;[fredsdir]
;  comment = Fred's Service
;  path = /usr/somewhere/private
;  valid users = fred
;  public = no
;  writable = yes
;  printable = no
# A service that has a different directory for each machine
that connects
# allows you to tailor configurations to incoming machines.
You could
# also use the %u option to tailor it by username.
```

```
# The %m gets replaced with the machine name that is
connecting.
;[pchome]
; comment = PC Directories
; path = /usr/pc/%m
; public = no
; writable = yes
# A publicly accessible directory with read/write to all
users. Note that all
# files created in the directory by users will be owned by
the default user, so
# any user with access can delete any other user's files.
Obviously this
# directory must be writable by the default user. Another
user could of course
# be specified, in which case, all files would be owned by
that user instead.
;[public]
; path = /usr/somewhere/else/public
; public = yes
; only guest = yes
; writable = yes
; printable = no
# The following two entries demonstrate how to share a
directory so that two
# users can place files there that will be owned by the
specific users. In this
# setup, the directory should be writable by both users and
should have the
# sticky bit set on it to prevent abuse. Obviously this
could be extended to
# as many users as required.
;[myshare]
; comment = Mary's and Fred's stuff
; path = /usr/somewhere/shared
; valid users = mary fred
; public = no
; writable = yes
; printable = no
; create mask = 0765
```

Notice that we are using `/home/samba` for storing the user programs such as the installation files for Microsoft Office 2000. We have given the group staff the ability to write to this directory. We will also need to allow a non-root user the ability to install programs here.

We made the server a WINS server with the line `wins support = yes`, and set up the printer to allow a guest user to print to it.

After DHCP and Samba are set up, we will need to create the user and group accounts as well as the logon scripts. Since we are using encrypted passwords, we will need to create an encrypted password file:

```
cat /etc/passwd | mksmbpasswd.sh > /etc/smbpasswd
```

Windows 95, OSr2, Windows 98, and Windows NT with Service Pack 3 or greater will support encrypted passwords by default. DOS and older versions of Windows will not support encrypted passwords without an upgrade.

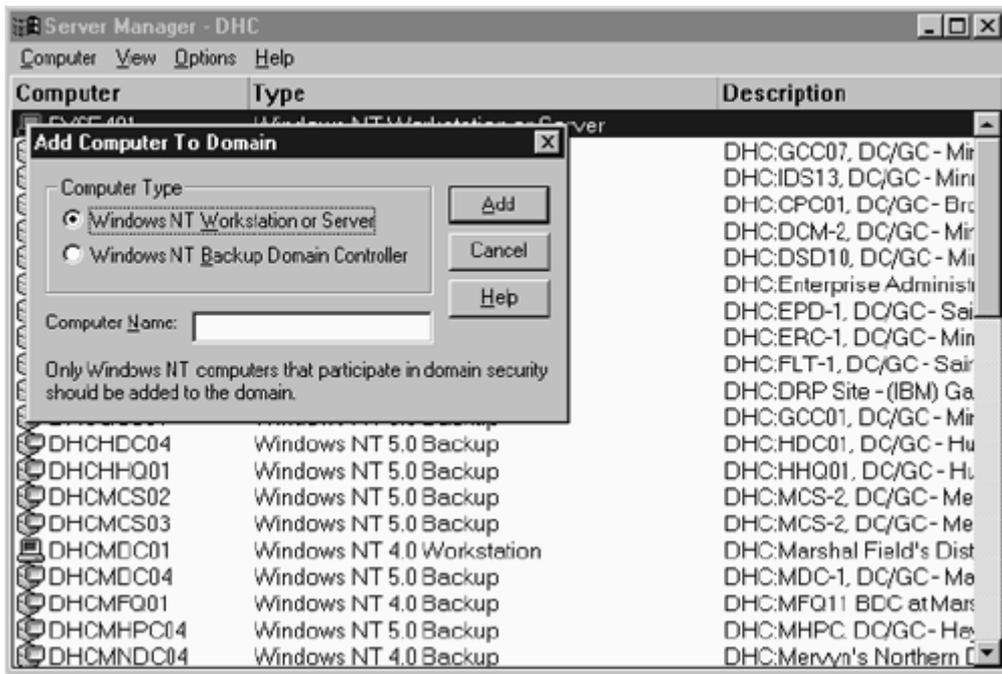
13.2 Adding a Samba Server to an Existing Network

Setting up a Samba server as a stand-alone network is relatively easy. Setting it up on an existing network is more complicated. You first have to get all the settings for the existing network, including the PDC and BDC, WINS servers, routers, and printers.

Samba has no problem joining a Windows domain, but it has some limitations. Samba cannot act as a BDC, and it also cannot exchange data with other WINS servers. Additionally, using Samba as a PDC is still in the experimental stage, so don't use Samba as a PDC in a production environment (you have been warned!).

Entering Samba into a Windows domain is relatively easy. First, create an account for the Samba machine in the Windows domain using the Server Manager for Domains as a stand-alone server or workstation. Just go to Computer -> Add to Domain. Type in the computer name and click Add.

Figure 13.1. Adding a computer to a domain in Windows NT with Server Manager.



Next, we need to configure `smb.conf` on the Samba machine. For this example, we will call the domain `NTDomain`, the PDC will be `NTPDC`, the BDC will be `NTBDC`, and the WINS server will be at IP address `10.0.0.2`. First, we need to stop Samba with this command:

```
samba stop
```

Next, add these settings to `smb.conf`:

```
Security = Domain
Workgroup = NTDomain
Password Server = NTPDC NTBDC ;(This can be set to * if you
want Samba
to automatically locate the password server)
```

```
encrypt passwords = yes
wins server = 10.0.0.2
```

After we restart Samba, (samba start) we need to create a Machine ID for the computer. The Machine ID is how the Domain Controllers identify each device in the Domain. To create this ID, type the following command:

```
smbpasswd -j NTDomain -r NTPDC
```

If successful, you should get a "Joined domain NTDomain" message. If you get another message, first check to see that there are no errors or typos in the previous steps. If this doesn't fix the problem, use the troubleshooting steps from the last chapter. Also, at the end of the appendix are troubleshooting procedures specific to NT Domains.

13.3 Samba as a Primary Domain Controller

As mentioned earlier, Samba can act as a PDC, but the full functions are not yet complete. The functions that work include:

- The ability to act as a PDC for Windows NT 3.51 Service Pack 5 and 4.0 Service Pack 4 clients. This includes adding NT machines to the domain and authenticating users logging into the domain.
- The domain account can be viewed using the User Manager for Domains.
- Viewing resources on the Samba PDC via the Server Manager for Domains of the NT client.
- Windows 95 clients will allow user-level security to be set, but will not currently allow browsing of accounts.
- Machine account password updates.
- Changing of user passwords from an NT client.
- Username <-> RID mapping, which stands for Relative Identification, which identifies a user on the domain.
- Some tools work with this such as the NT Sec tools from pedestal software. Some tools, like `explorer.exe`, do not.
- Partial support for Windows NT group and username mapping.
- Support for an LDAP password database back-end.

Some of the features that are not implemented yet are:

- Trust relationships.
- PDC <=> BDC integration.
- Network printing (although there is a workaround).
- Windows NT Access Control Lists (ACLs) on the Samba shares.

If you still want to use Samba as a PDC, download and compile the latest code first. The easiest way to do this is to use a Revision Control System (RCS), which is included with most Linux distributions. To see if an RCS is installed on your system, type `cvs` at the prompt to pulled up the cvs help message. If you don't have RCS, you can download it from <ftp://ftp.gnu.org/gnu/rcs>.

Once RCS is installed, log into the Samba cvs server. Use your email address as the password:

```
cvs -d :pserver:cvs@cvs.samba.org:/cvsroot login
```

Then download the newest source code into the `/samba` directory:

```
cvs -d :pserver:cvs@cvs.samba.org:/cvsroot co samba
```

The alpha and beta code will have more of the features working than the latest (stable) code, but they will be less stable. In any case, using Samba as a PDC is not recommended in a production environment. For instructions on compiling the code, see [Chapter 12](#).

Setting up the `smb.conf` of a Samba PDC is the same as setting up a normal Samba server except for the following:

```
workgroup = NTDomain
encrypt password = yes
domain logon = yes
prefered master = yes
security = user
```

You might also want to include the following, but don't enable WINS support unless there are no other WINS servers on the local subnet:

```
wins support = yes
logon script = %u.bat (%u is a variable that translates to
the username)
```

If you enable logon scripts, create a netlogon share:

```
[netlogon]
    path = <path of netlogon directory>
    writable = no
    guest = no
```

Then create machine accounts for the PCs. This is equivalent to adding them with the Server Manager for Domains on a Windows NT server.

```
smbpassword -m <netbios name of machine>
```

On the Microsoft workstations, change the domain to match the domain of the Samba server, but make sure the Create Account check box is not selected. When you select OK, you should get the message "Welcome to the <domain name> domain." Then reboot the workstation and log in. You should have a box with three fields: Name, Password, and Domain. This concludes a brief explanation of the workstation configuration.

The Microsoft client's configuration is exactly the same as it would be for connecting to a Microsoft server. Just be sure that TCP/IP is configured and NetBIOS over TCP/IP is enabled (it is usually enabled by default in Windows). This is set in the Network icon in the Control Panel.

Chapter 14. Connecting Linux to Windows PCs

Linux can attach to a networked Windows PC by using Windows' native network capability. First, set up networking in Windows using the Network icon in the control panel.

Several programs are used by Linux to connect to Windows machines. The following programs are included in most Linux distributions and with Samba: `smbclient`, `smbfs`, `smbwrapper`, `smbsh`, `sharity`, `smbtar`, and `smbprint`.

`smbclient` is a command-line utility that allows Linux to browse and copy files between Linux and Windows systems.

`smbfs` and `sharity` allows Windows filesystems to be mounted natively on a Linux system. The Windows directory will look like a normal Linux directory. `sharity` and `smbclient` also support SSL for encrypted transactions.

`smbwrapper` and `smbsh` allow browsing of Windows filesystems from a Linux machine.

`smbtar` allows the backup of a Windows machine to a Linux machine.

`smbprint` allows Linux to print to a shared Windows printer.

There are other packages that use these utilities to access Windows computers. For example, some Linux backup utilities use `smbclient` to back up Windows clients.

`smbclient` is a command-line utility similar to FTP. It allows Linux machines to connect to a Windows share. Once connected, `smbclient` allows Linux to list and copy files as follows:

```
smbclient //alisa/c
Added interface ip=10.0.0.2 bcast=10.0.0.255
nmask=255.255.255.0
Password:
smb: \> ls
  BOOTLOG.TXT          AH      89906   Thu Nov  4 15:24:40 1999
  COMMAND.COM          A       93890   Fri Apr 23 22:22:00 1999
  SUHDLOG.DAT          HR       5166   Fri Aug  6 13:15:00 1999
  ...files deleted here for brevity.....
      64322 blocks of size 131072. 31108 blocks
available
```

Notice that the share names use forward slashes (/) instead of the Windows standard back slashes (\). This is because Linux treats a back slash as an escape character, so in most cases where Windows would use a back slash, Linux would use a forward slash.

We used `ls` to get a list of files, but `dir` would work as well. There are several commands that `smbclient` can accept once it is connected. These can be listed by typing `help`:

```
smb: \> help
ls          dir          du           lcd
cd
pwd         get          mget        put
mput
rename     more        mask        del
open
rm         mkdir       md           rmdir
rd
```

```
prompt          recurse          translate       lowercase
print
printmode       queue          cancel          quit
q
exit            newer          archive         tar
blocksize
tarmode         setmode        help            ?
!
```

To get help on a specific command, type `help <command name>`. For example, if we wanted help on `ls`, we would type `help ls`:

```
smb: \> help ls
HELP ls:
    <mask> list the contents of the current directory
```

As you can see, there is only one parameter for `ls`. It is a mask, which limits the files shown. For example, `ls *.exe` would find all files ending in `exe`.

We will cover the `smbclient` command later. First, we need to cover the command-line options.

14.1 smbclient Command-Line Options

The general form of the `smbclient` command line is:

```
smbclient <servicename> <password> <options>
```

The `servicename` is the name of the Windows service you are attaching to. It is in the form of `//server/share`. The `password` is the password for the share or user logging into the service. This depends on whether Windows is using shared or user-level security. The options are covered below.

- `-s smb.conf`— This points to the location of the `smb.conf`. Like the rest of the Samba suite, `smbclient` gets its settings from the `smb.conf`.
- `-B ip address`— This allows `smbclient` to look up the Windows machine's NetBIOS name by using its IP address.
- `-O Socket Options`— This option allow you to tune the TCP socket used by `smbclient`.
- `-R Name Resolve Order`— There are several ways that Linux uses to resolve names on a network. By default, Linux uses this order to resolve names:
 1. `lmhost`— The file used by Samba for looking up host names. It is stored in the same directory as `smb.conf`.
 2. `host`— This is the default file for looking up names on a Linux system. It is kept in the `/etc` directory.
 3. `wins`— If available, the Linux system will use the WINS servers to resolve names.

4. `bcast`— If the other methods fail, Linux will use a broadcast to attempt to resolve the name. It uses the `interfaces` parameter in `smb.conf` to determine how to broadcast. If the `interfaces` parameter is not set, it will broadcast on all known interfaces. This parameter simply allows the default order of name resolution to be changed; for example, `-R host lmhost bcast` would make the `host` file the first method and it would not use `wins` for resolution.
- `-M Netbios name`— This sends a message to other computers, which is called a WinPopup message on Windows machines. On Windows 9x machines, you must have WinPopup enabled for this feature to work. Also note that this only works with NetBIOS names, not IP addresses or IP host names. After this command is entered, you are asked for the message, followed by <CTRL> D:
 -
 - `smbclient -M alisa`
 - Added interface ip=10.0.0.2 bcast=10.0.0.255 nmask=255.255.255.0
 - Connected. Type your message, ending it with a Control-D
 - Hello!sent 6 bytes
 -
 - `-i scope`— This allows you to set the NetBIOS scope. It is rarely used.
 - `-N`— This suppresses the password prompt. It is especially useful for adding `smbclient` to scripts.
 - `-n netbios name`— This allows you to set the NetBIOS name of the Linux machine. By default, `smbclient` uses the NetBIOS name set in `smb.conf`. If this isn't set, it uses the local host name as the NetBIOS name.
 - `-d debug level`— This controls the number of debugging messages that are printed, which is used to debug problems. The default is 0, which only prints serious errors or warnings. Higher levels (larger numbers) will print more data. A debug level greater than 3 is rarely needed.
 - `-p port`— This allows you to change the TCP port used by `smbclient`. Since Windows machines always use Port 139, this option is needed only in rare cases.
 - `-h`— Shows a help screen that lists the command-line options.
 - `-I destination IP address`— Allows you to specify the IP address of the Windows machine that you are connecting to. This is often used when NetBIOS names can't be resolved to IP addresses. Just keep in mind that most Windows machines also require a NetBIOS name to make a connection.
 - `-E`— This writes the error messages to `stderr` (the error message file) rather than `stdout` (usually the screen). The most common use for this is to redirect error messages to a file.

- `-U username`— Allows you to specify the username used to log on to the Windows PC. The default is the upper-case version of the word "USER". If this doesn't exist, the logname variables are used. You can see these by typing:
 -
 - `env |grep USER`
 - `env |grep LOGNAME`
 -
- `-t terminal code`— This determines how the `smbclient` interprets filenames. It is useful if you are using different filesystems such as the Shift Japanese Industrial Standard (sjis). The default Linux filesystem doesn't need this option to be set.
- `-W workgroup`— This specifies the workgroup that `smbclient` is using. The default `workgroup` is set in the `workgroup` section of the `smb.conf` file.
- `-T tar options`— This allows you back up Windows machines to a Linux machine. It has many options, as listed below. To make things easier, there is a script called `smbtar` (discussed later in the chapter).
- `c tar file`— Specifies which file to back up to. It can also be a device such as a tape drive. For example, you can write to `/dev/rmt0`, which is a tape drive.
- `x tar file`— Restores from a backup file. Like above, it can be a file or device such as a tape drive.
- `I include expression`— Specifies which files or directories to include in the backup or restore. It can use a wildcard such as `*`, or `?` if the `r` option is specified.
- `X exclude expression`— Specifies which files or directories to exclude from the backup or restore. Like above, it can use wildcards if the `r` option is specified.
- `r`— Allows the use of wildcards.
- `b blocksize`— Specifies the number of 512-byte blocks to use in the `tar` file.
- `g`— Specifies incremental mode. It will only back up files with the archive bit set. This is useful for backing up files that have changed since the last backup.
- `q`— Quiet mode. Messages not printed.
- `N file`— Only files newer than the specified file are backed up. Used in conjunction with the `c` option.
- `a`— Resets the archive bit. Used with the `g` and `c` options.

Since the `tar` options are complicated, let's examine several examples of using them. To back up all the the files on the `c` share on the Windows PC `alias` to the file `backup.tar`, type the following:

```
smbclient //alisa/c -Tc backup.tar
```

To restore the entire `c` share on `alisa`, type:

```
smbclient //alisa/c -Tx backup.tar
```

To back up the new files and reset the archive bit, type:

```
smbclient //alisa/c -Tcga backup.tar
```

You can also use tape drives to back up. If you have a tape drive called `rmt0`, the command would be:

```
smbclient //alisa/c -Tcga /dev/rmt0
```

14.2 smbclient Commands

Once `smbclient` is logged into the Windows machine, the first thing you get is a prompt:

```
smb: \>
```

This represents the root of the share on the Windows PC. From here, we can enter `smbclient` commands:

- `? [command]` or `help [command]`— Help for `smbclient` commands. By itself, this command will give help on all the `smbclient` commands, or you can enter a command name for help on the use of the individual command.
- `! [Linux command]`— Allows the running of a Linux command on the local machine without exiting `smbclient`. For example, to view the files on the local hard drive:
 -
 - ```
smb: \>! ls
```
  - ```
MSPRINT.INF                MSPRINT2.INF
```
 - ```
MSPRINT3.INF MSPRINT4.INF
```
  - ```
smb: \>
```
 -
- `archive [level]`— This allows the `mget` command to retrieve files on the Windows machine based on the archive bit. The different levels are:

0 Retrieves all files and leaves the archive bit alone.

1 Retrieves files with the archive bit set and leaves the archive bit alone.

2 Retrives files with the archive bit set and resets the archive bit on those files.

3 Retrives all files and resets the archive bit of all files.

To find out what the archive level is currently set to, type `archive` at the prompt:

```
smb: \> archive
Archive level is 0
smb: \>
```

- **Blocksize**— Specifies the block size for the tar option. The block size is in 512-byte increments. For example, to set the block size to 10 Kilobytes (10K), use the following command:

-
- `smb: /> blocksize 20`
- `blocksize is now 20`
- `smb: />`
-

- **cancel <jobid>**— This cancels the specified jobid in the print queue.

-
- `smb: /> cancel 1`
- `Job 1 cancelled`
- `smb: />`
-

Of course, you must be attached to a printer and know the print job number to do this.

- **cd [directory]**— This allows you to change the directory on the Windows machine. If a directory is specified, it will change to that directory the same as using the Windows `del` command, including the wildcards `*` and `?`.
- **dir [expression] or ls [expression]**— Either of these commands will show a list of files on the Windows machine matching `expression`. It also supports the wildcards `*` and `?`.
- **du**— This command simply shows the amount of space used by all the files in the current directory. If the `recursive` option is turned on, it will show the total for all the directories:

-
- `smb: \> recurse`
- `directory recursion is now off`
- `smb: \> du`
- `64322 blocks of size 131072. 31490 blocks available`
- `Total number of bytes: 20794893`
- `smb: \> recurse`
- `directory recursion is now on`
- `smb: \> du`

- `64322 blocks of size 131072. 31490 blocks available`
- `Total number of bytes: 4240552435`
- `smb: \>`
-

Notice that the total number of bytes with recursion off is 20794893, whereas the total is 4240552435 with recursion on.

- `exit` or `quit` or `q`— Like it says, exits `smbclient`.
- `get <remote file> [local file]`— Copies a remote file from the Windows machine to the Linux machine. You can optionally give the file a different name on the Linux machine with `[local file]`. Otherwise, the Linux filename will be the same as the Windows filename. If `lowercase` is on, the local files will be all lower-case versions of the remote file.

Files are always transmitted in binary mode, unless `translation` is on, which will transmit in ASCII mode, which is used for text files. The ASCII mode translates Windows standard line feeds into Linux standard line feeds. For example:

```
smb: \> lowercase
filename lowercasing is now on
smb: \> lowercase
filename lowercasing is now off
smb: \> translate
CR/LF<->LF and print text translation now on
smb: \> translate
CR/LF<->LF and print text translation now off
smb: \> get config.sys
getting file config.sys of size 59 as config.sys (0.433212
kb/s)
(average 0.433212 kb/s)
smb: \>
```

- `lcd [directory]`— Allows you to change directory on the local (Linux) machine. It uses the same syntax as the `cd` command in Linux.
- `lowercase`— When turned on, this option will convert any files copied from the Windows machine into lower-case filenames on the Linux machine. This is useful since most Linux files are in lower-case.
- `mask <expression>`— This sets the mask for the multiple copy commands `mget` and `mput`. It specifies which files are copied. When `recursive` is on, the directories copied are specified in the `mget` or `mput` commands themselves. `mask` starts out blank and once set, remains the same until it is changed with another `mask` command.
- `md <directory>` or `mkdir <directory>`— These commands allow you to create a directory on the Windows machine. You must have the rights to create the directory, of course.

- `mget <expression>`— This command acts like the Windows `xcopy` command. It will copy all the files matching the expression from the Windows machine to the Linux machine. The expression can contain any valid file characters as well as wildcard characters (`?`, `*`) There are several other commands that affect `mget`.
- `recursion`— If `recursion` is on, `mget` will also copy files in subdirectories and create the directories on the Linux machine. The expression in `mget` will select the directories. If `recursion` is off, the expression in `mget` selects the files instead of the directories.
- `mask`— Specifies the files selected. Once set, it will stay on until changed.
- `prompt`— If on, you will be prompted before each file is copied.
- `newer`— Copies files newer than a specified local file.

In the following example, `mask` and `recursion` are off and `prompt` is on:

```
smb: \staroffice\> ls
.                D          0  Fri Oct  1
18:40:18 1999
..               D          0  Fri Oct  1
18:40:18 1999
so51a_lnx_01_n1.tar  A   7442432  Fri Oct  1
18:47:34 1999
so51a_lnx_01_n2.tar  A  11600384  Fri Oct  1
18:48:52 1999
so51a_lnx_01_n6.tar  A  10398720  Fri Oct  1
18:48:54 1999
so51a_lnx_01_n3.tar  A  10718720  Fri Oct  1
18:49:44 1999
so51a_lnx_01_n5.tar  A  10459136  Fri Oct  1
18:50:00 1999
so51a_lnx_01_n7.tar  A   9664000  Fri Oct  1
18:50:06 1999
so51a_lnx_01_n4.tar  A  10122752  Fri Oct  1
18:50:16 1999
GUILG00         A   9604832  Fri Oct  1
20:18:46 1999
G2player-6.0-0.99092901_i386.rpm  A   739152  Wed Oct 20
18:13:02 1999
MSPRINT2.INF    A   41315   Fri Apr 23
22:22:00 1999
MSPRINT.INF     A   36013   Fri Apr 23
22:22:00 1999
MSPRINT3.INF    A   49761   Fri Apr 23
22:22:00 1999
MSPRINT4.INF    A   55837   Fri Apr 23
22:22:00 1999
PhatLinux32.zip  A 190371364  Thu Nov  4
20:34:48 1999
```

```
kvincviewer-0_0_3_tar.gz          A    159923  Fri Nov  5
21:39:18 1999
smb: \staroffice\> mget *.inf
Get file MSPRINT2.INF? y
getting file MSPRINT2.INF of size 41315 as MSPRINT2.INF
(568.262 kb/s)
  (average 382.173 kb/s)
Get file MSPRINT.INF? y
getting file MSPRINT.INF of size 36013 as MSPRINT.INF
(567.24 kb/s)
  (average 400.243 kb/s)
Get file MSPRINT3.INF? y
getting file MSPRINT3.INF of size 49761 as MSPRINT3.INF
(578.508 kb/s)
  (average 421.07 kb/s)
Get file MSPRINT4.INF? y
getting file MSPRINT4.INF of size 55837 as MSPRINT4.INF
(586.325 kb/s)
  (average 439.997 kb/s)
smb: \staroffice\>
                    64322 blocks of size 131072. 31365 blocks
available
```

- `more`— `more` works like the `more` command in Linux. It allows a text file on the Windows machine to be read a page at a time.
- `mput <expression>`— This command works just like the `mget` command, except that it copies files from the Linux machine to the Windows machine. All the rules for `mget` listed above apply to `mput`.
- `newer <filename>`— Allows `mget` to only copy files that are newer than the Linux file listed in `filename`. For example, to get files newer than the Linux file `core`:
 -
 - `smb: \staroffice\> newer core`
 - `Getting files newer than Sun Nov 7 13:19:06 1999`
 - `smb: \staroffice\>`
 -
- `print <file>`— This command prints a file on the Windows machine using the Linux machine's print services. This can also be accomplished by using the `put` command to copy the file to the print spool directory (see the section on printing).
- `Put <local file> [remote file]`— This command acts like the `get` command except that it copies the file from the Linux machine to the Windows machine. All the rules for `get` apply to `put`.
- `pwd`— This works just like the Linux command `pwd`. It shows the current directory on the Windows machine.

- `queue`— You can view the print queue on the Windows machine with this command.
 - `rd <directory>` or `rmdir <directory>`— Removes a directory from the Windows machine. You must have rights to remove the directory and the directory must be empty. If the directory is not empty, you will get the following error:
 -
 - `smb: \> rd staroffice`
 - `ERRDOS - ERRnoaccess (Access denied.) removing remote directory file`
 - `\staroffice`
 - `smb: \>`
 -
 - `recurse`— This command forces `du`, `mget`, and `mput` to recurse into subdirectories. It is turned off when `smbclient` is opened. To turn it on, type `recurse`:
 -
 - `smb: \> recurse`
 - `directory recursion is now on`
 - `smb: \>`
 -
- If recursion is on, use the `mask` command to specify the files copied by `mget` or `mput`.
- `rm <expression>`— `rm` removes all the files matching `expression` in the current directory of the Windows machine.
 - `setmode <file> [<+ -> <r s h a>]`— This works just like the `attrib` command in Windows. It allows you to change the attributes of a file on the Windows machine as follows:
 - `+`— Adds attribute.
 - `-`— Removes attribute.
 - `r`— Read attribute. When set, makes the files read-only.
 - `s`— System attribute; used for Windows system files.
 - `h`— Hidden attribute; make the file hidden.
 - `a`— Archive bit: this is set on when a file needs to be backed up.
 - `tar <c x> [a b g r q I X N]`— `tar` can back up or restore files from the Windows machine. The commands are exactly the same as the `-T` command-line option for `smbclient`.

- `tarmode <[no] < full inc reset noreset hidden quiet verbose>>` +— This changes the behavior of the `tar` command as follows:
 - `full`— Backs up all files regardless of archive bit.
 - `inc`— Backs up files with the archive bit set.
 - `reset`— Resets the archive bit on all files that are backed up.
 - `system`— Backs up files with system bit set.
 - `hidden`— Backs up files with the hidden bit set.
 - `verbose`— Shows the status during backup.
 - `quiet`— Shows no information during backup.

Placing a `no` in front of a parameter turns it off.

To see what your `tarmode` settings are, type `tarmode` without parameters:

```
smb: \> tarmode
tarmode is now full, system, hidden, noreset, verbose
smb: \>
```

This is the default for `tarmode`. It is set for a `full` backup, including `system` and `hidden` files. It will not reset (`noreset`) the archive bits and will show messages on the screen (`verbose`).

- `translate`— When a file is copied from a Windows machine, `translate` will convert Windows line feed characters into Linux line feed characters. This allows text files to be read normally on the Linux system.

14.3 smbtar

Since the `tar` command for `smbclient` is so complex, the script `smbtar` is included with Samba to simplify backups and restores. The format for `smbtar` is:

```
smbtar options files
```

The options for `smbtar` are as follows:

- `-s server`— This specifies the server (Windows machine) that you are backing up. This option is mandatory.
- `-p password`— The password for the share, if needed.
- `-x service`— The service to connect to. The default is `backup`.

- `-X filenames`— Files to be excluded from the backup.
- `-u user`— The username to connect to the service. The default is the Linux username that you are logged in as.
- `-d directory`— The initial directory to start the backup or restore from.
- `-t tape`— The device or file that is used for the backup or restore. The default is the `tape` variable set in Linux. If there is no `tape` variable, the `tar.out` file is used.
- `-b blocksize`— The tape block size. The default is 20.
- `-N filenames`— Specifies to back up files newer than `filenames`.
- `-i`— Performs an incremental backup. Only files with the archive bit set are backed up.
- `-a`— Resets the archive bit on all files backed up.
- `-r`— Performs a restore instead of a backup.
- `-l log level`— Specifies the debug level. The default is 0. This has the same effect as the `-d` option in `smbclient`.
- `-v`— This sets `smbclient` to verbose mode.
- `filenames`— Without the `X` parameter, you can list the files to back up. If no filenames are listed, it will default to all filenames.

14.4 smbprint

This is another `smbclient` script that allows the printing of a Linux file on a printer attached to a Windows machine. It takes the standard input on the Linux machine and redirects it to the remote printer. The syntax is similar to the following:

```
smbprint <local.file>
```

For `smbprint` to work properly, a `.config` file must be created with the following parameters:

- `server`— The name of the Windows machine.
- `service`— The share name.
- `password`— The password needed for the share, if needed.

14.5 smbfs

Server Message Block Filesystem (`smbfs`) allows you to mount a drive on a Windows machine just like you would mount a drive on a Linux or UNIX machine. To mount a Windows drive, you need two things:

1. `smbfs` support in the Linux kernel.
2. Programs to mount and unmount the `smbfs`.

Most current Linux distributions have `smbfs` support built into the kernel. If your kernel doesn't have `smbfs` support, you must rebuild your kernel with `smbfs` support or use a loadable module. Since kernels and kernel modules are beyond the scope of this book, consult your documentation or a good Linux book on how to do this.

Once kernel support for `smbfs` has been added, you will need to use the utilities `smbmount` and `smbumount` for mounting and unmounting Windows shares on your Linux machine. `smbmount` uses the same command-line parameters as `smbclient`. The general syntax of the command is `smbmount <share> <mount point>`. For example, if we wanted to mount the `c` share on the Windows machine `alisa`, we would use the following command. Note that a directory `/mnt/alisa` was created on the local Linux drive for the mount point.

```
smbmount //alisa/c /mnt/alisa
Added interface ip=10.0.0.2 bcast=10.0.0.255
nmask=255.255.255.0
Password:
[root@mmccune /]#
```

We can now change to the mount directory and run an `ls`.

```
[root@mmccune /]# cd /mnt/alisa
[root@mmccune alisa]# ls -l
total 20330
drwxr-xr-x   1 root   root         512 Aug  6 14:31
Acrobat3
-rwxr-xr-x   1 root   root       76503 Nov  8 12:22
BOOTLOG.PRV
-rwxr-xr-x   1 root   root       76127 Nov  8 12:35
BOOTLOG.TXT
-rwxr-xr-x   1 root   root       93890 Apr 23  1999
COMMAND.COM
-rwxr-xr-x   1 root   root         59 Nov  8 13:14
CONFIG.SYS
(deleted for brevity)
[root@mmccune alisa]#
```

Notice that it looks just like a Linux directory. It lists the standard Linux permissions even though Windows machines don't support the same permissions as Linux. `smbfs` simply fakes the permissions. A file that is listed as writable may be read-only or vice versa. This can lead to problems if you are not careful.

There is another major problem with `smbfs`. If an `smbfs`-mounted drive loses its connection with the Windows machine, the mounted drive will lock up and it won't unmount or mount until the Linux machine is rebooted. This problem may be fixed in future versions of `smbfs`, but for now, it is an annoying problem.

14.6 Sharity

A way around the lock-up problem of `smbfs` is to use `sharity`. `sharity` is a separate program rather than part of the kernel like `smbfs`. It also will allow the SMB share to be re-exported via NFS.

On the downside, `sharity` doesn't support file locking and often has trouble with file handles, which leads to occasional Stale NFS file handle errors.

There are two flavors of `sharity`: `sharity` and `sharity-light`. `sharity` is a commercial program that requires a license. `sharity-light` is a GPL version of the program that lacks many of `sharity`'s advanced features such as:

- Encrypted passwords.
- Mounts multiple shares from the same server.
- Displays correct file modification dates.
- Various caching strategies improve performance.
- Overlapping requests improve performance.
- Automounting facility.
- Better mapping of file operation semantics.
- Multi-user operation.
- Supports International character set.
- Can link to SSL for secure data transport (first implemented in beta version 0.15).

`sharity` and `sharity-light` have a similar syntax to `smbfs`. The basic syntax of `sharity-light` is:

```
shlight //server/service mount-point [options]
```

For a complete list of options, go to <ftp://ftp.obdev.at/pub/Products/Sharity-Light/Sharity-Light.README>.

`sharity`'s home page is at <http://www.obdev.at/Products/Sharity.html>, and `sharity-light`'s home page is at <http://www.obdev.at/Products/shlight.html>.

14.7 Conclusion

These are the tools needed to connect a Linux machine to a Windows machine. In the next chapter, we will discuss printing with Samba.

Chapter 15. Printing with Samba

Setting up a printer can be tricky and it deserves some extra coverage. The printer must be set up for Linux before it can be set up for Samba. Most distributions of Linux will set up a printer during the install. Linux uses the `/etc/printcap` file to configure its printer. Unless you know what you are doing, it is not recommended that you edit this file yourself, as it contains many control characters, which are non-printing characters that are used to configure the printer:

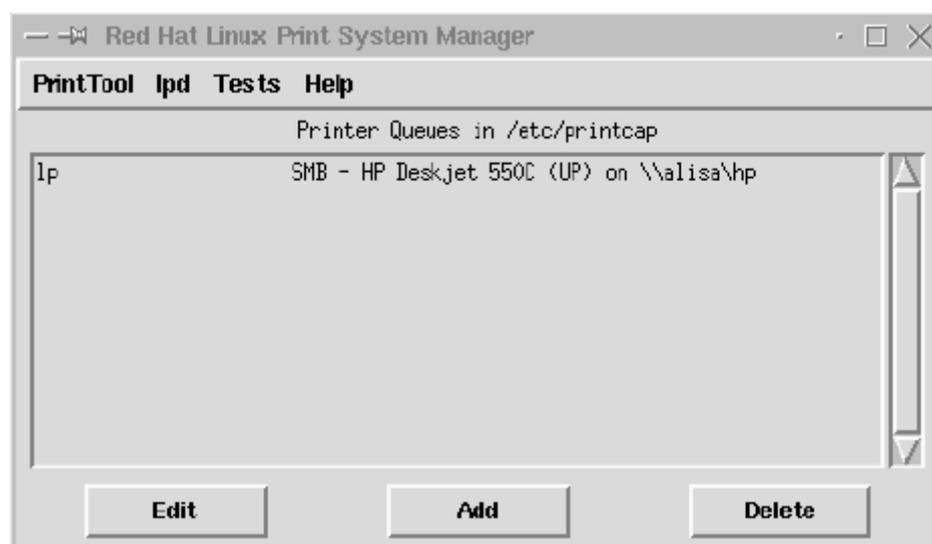
```
lp:\
:sd=/var/spool/lpd/lp:\
:mx#0:\
:sh:\
:lp=/dev/null:\
:af=/var/spool/lpd/lp/acct :\
:if=/var/spool/lpd/lp/filter:
```

If you want to be a masochist you can edit this file with a text editor. For the rest of us, there are configuration utilities for printers. Red Hat and several other distributions come with `printtool`, which runs in X-windows (the graphical interface for Linux and UNIX). Other distributions of Linux may have another tool to configure printers (check your Linux user manual).

15.1 printtool

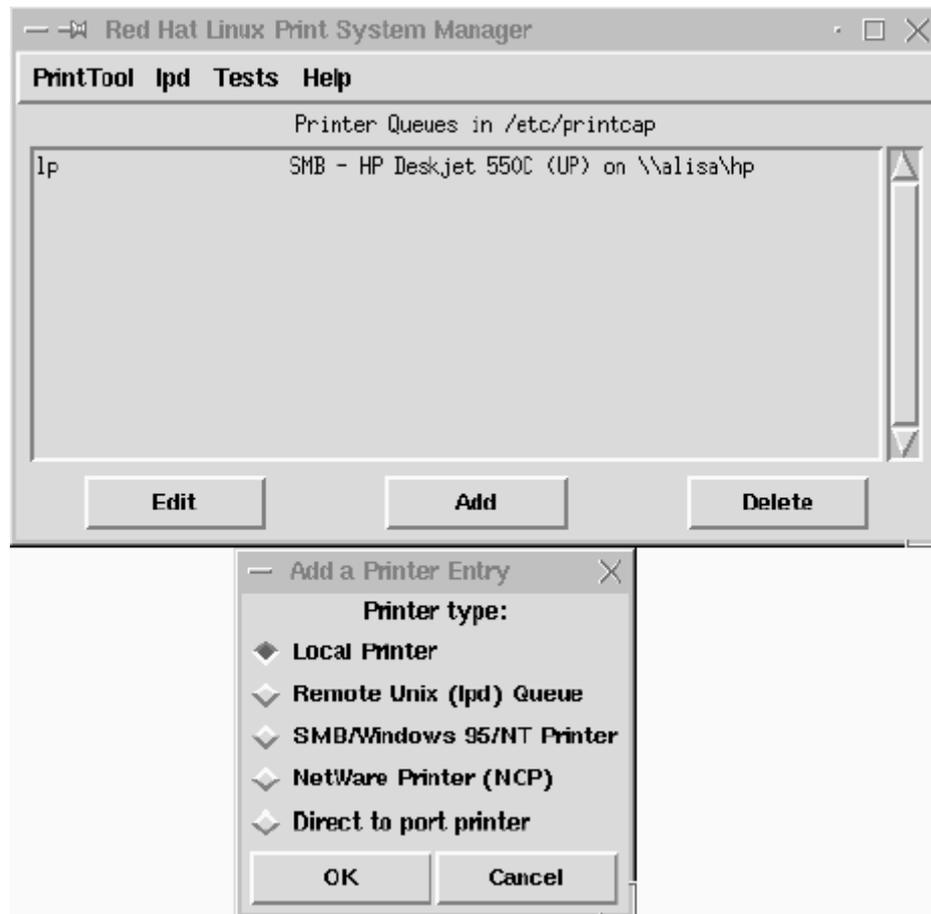
Start `printtool` by typing `printtool&` at the command prompt window inside of X-windows. It will display a Window similar to [Figure 15.1](#).

Figure 15.1. PrintTool's main screen.



Choose the Add button (see [Figure 15.2](#)).

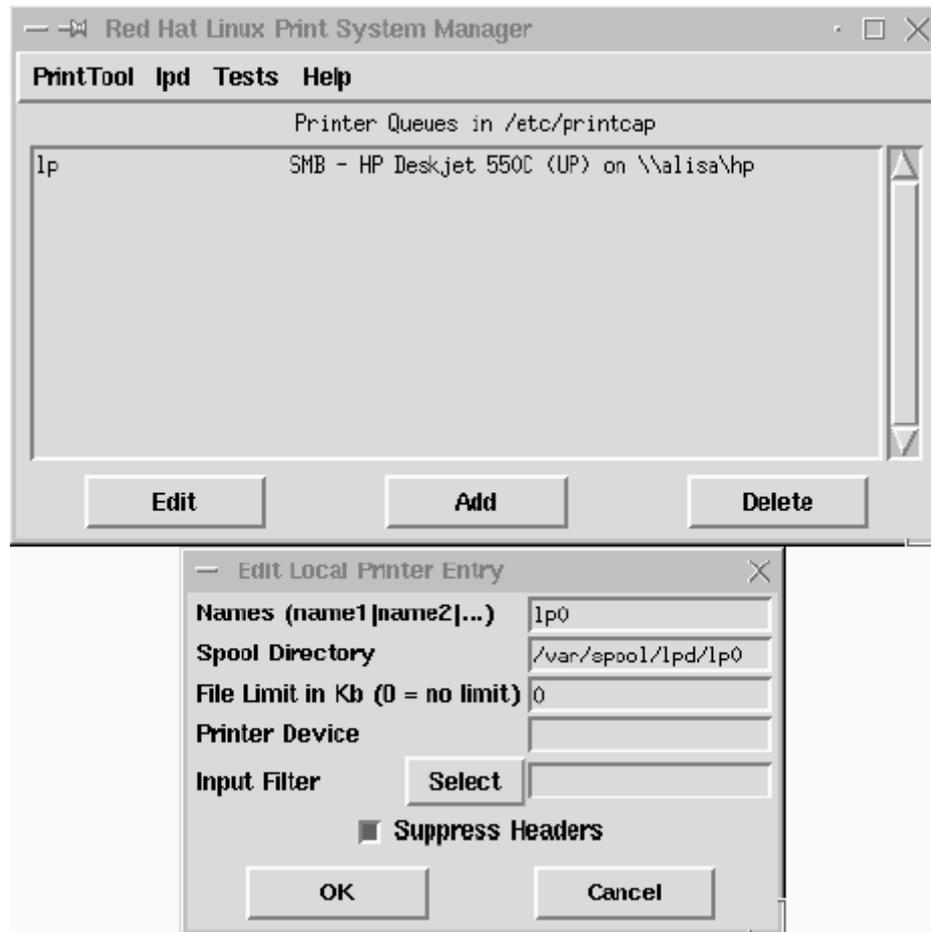
Figure 15.2. The Add Printer menu in PrintTool.



You will be given a choice of Local Printer, Remote Unix (lpd) Queue, SMB/Windows 95/NT Printer and Netware Printer (NCP). Let's go through these one at a time.

- **Local Printer**— A local printer is attached directly to the PC via a parallel or serial port. `printtool` will try to detect your printer. It will then show a window like [Figure 15.3](#).

Figure 15.3. Adding a local printer with PrintTool.



The Name identifies the printer. It can be up to eight characters. Leave the Spool Directory alone, unless you have a good reason to change it. The File Limit is at 0 (unlimited size) by default, but you can limit the size of the spool files.

The Printer Device Field should be filled in. If it isn't, `/dev/lp0` is the first printer port, `/dev/lp1` is the second printer port, etc. `/dev/cua0` is the first serial port (COM1 in Windows), `/dev/cua1` is the second serial port, etc.

At Input Filter, choose the Select button. It should give a list of printers. If your printer isn't listed, perhaps there is a compatible driver for your printer. Check your user manual for help. These Web sites may also help:

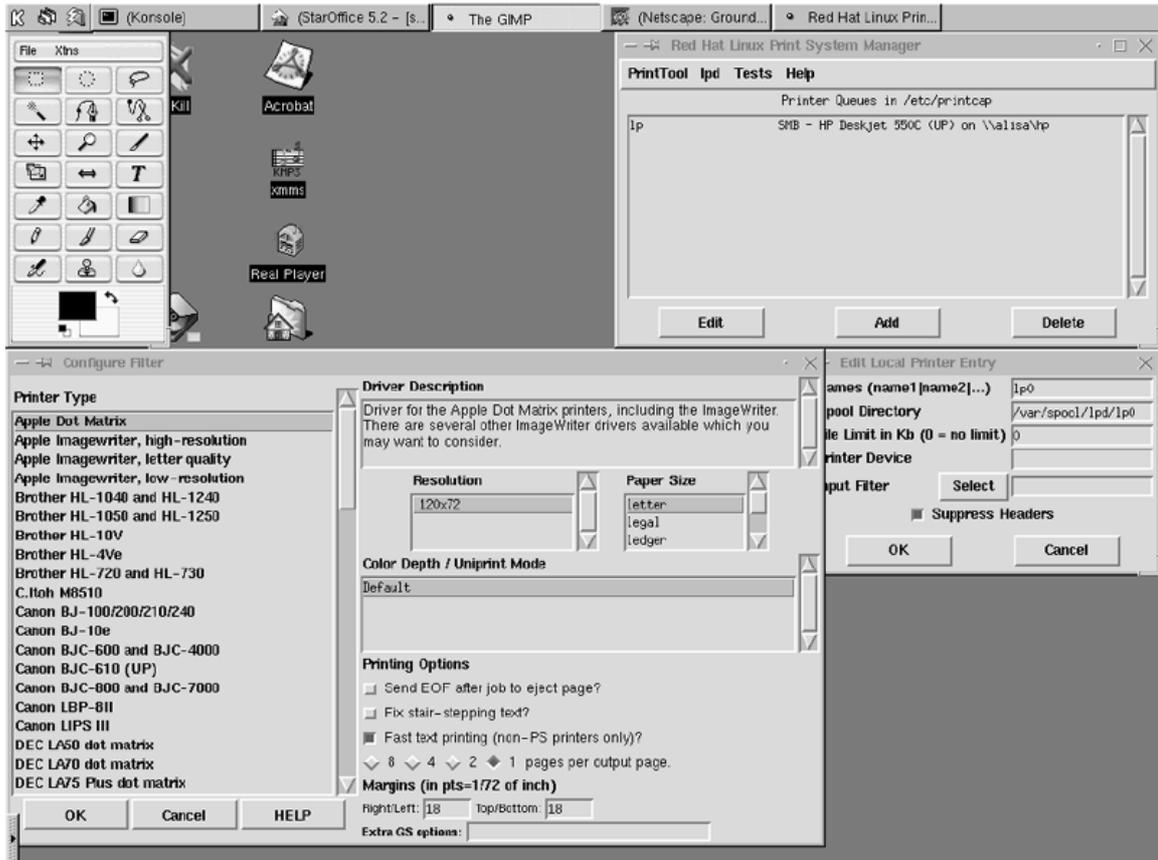
- <http://www.cs.wisc.edu/~ghost/printer.html>
- <http://www.linuxprinting.org/database.html>

Linux inherits native PostScript support from UNIX. PostScript is a standard printing language by Adobe (<http://www.adobe.com>) that is widely used. Hewlett Packard's PCL is the other widely used printing language. One advantage of native PostScript support is that Linux can print PostScript documents on a printer that doesn't support PostScript.

There are some printers that will not work with Linux. WinPrinters are a prime example. These printers use Windows drivers to process the print image and thus will only work with Windows. See the Web pages above for compatible printers.

When you have selected your print driver, choose OK and a listing of printers should appear in the window as shown in [Figure 15.4](#).

Figure 15.4. Selecting a printer driver in PrintTool.



- **Remote Unix Queue**— This is set up the same as the local printer, except there is a Remote Host and Remote Queue to set. The Remote Host is the IP host name or IP address of the remote machine that has UNIX printing enabled. The Remote Queue is the name of the remote printer. This setting is made on the remote machine.
- **SMB/Windows Printer**— The SMB printing option allows a Linux box with Samba installed to print to a shared printer on a Windows machine. This has some extra settings:
 - **Host Name of Printer Server**— This host name must be able to be resolved as an IP address.
 - **IP Number of Server**— Use this if you are having problems resolving the IP addresses of the hosts.
 - **Printer Name**— The share name of the printer.
 - **User**— Logon name of the host machine, if required.
 - **Password**— Password for above logon.
 - **Workgroup**— The workgroup that the host machine is in, if required.
- **Netware (NCP) Printer**— This allow a Linux machine to print to a Novell NetWare server. To use this, you must have NetWare support installed on Linux. It has these extra settings:
 - **Printer Server Name**— The name of the NetWare printer server.

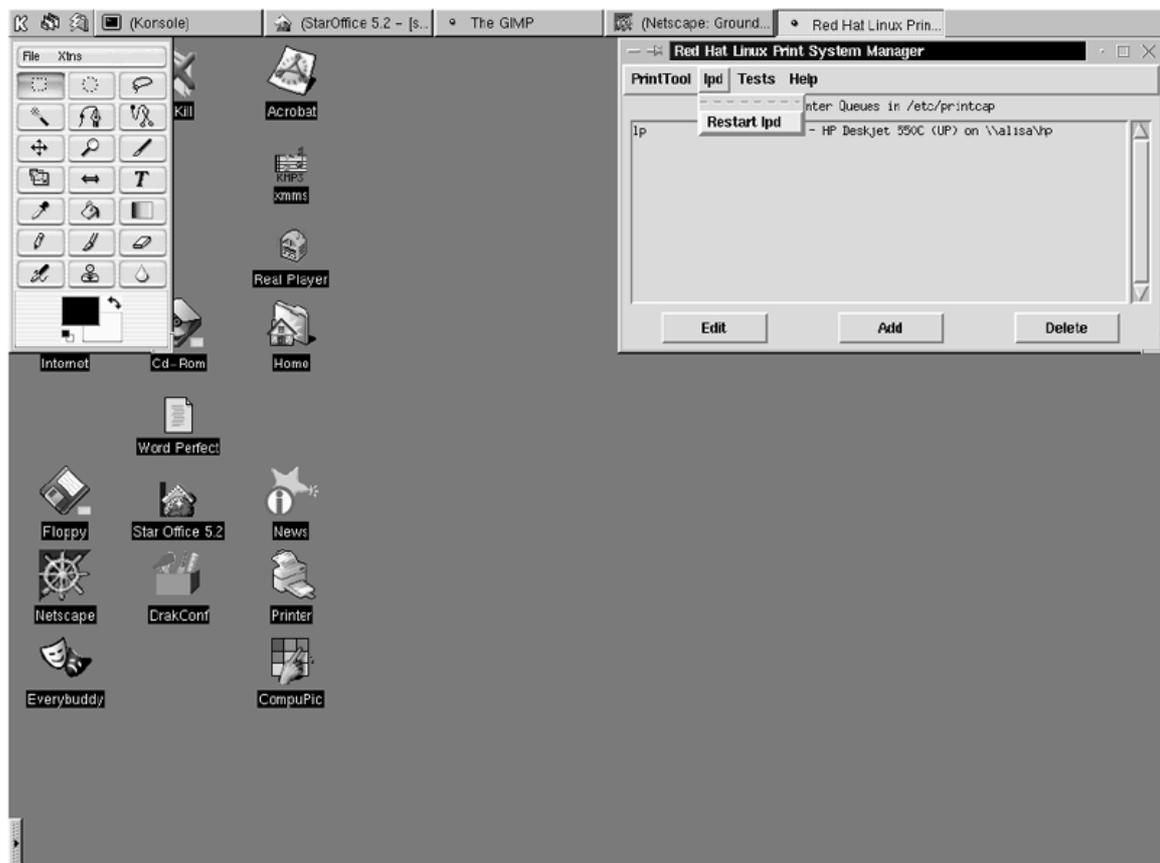
- **Printer Queue Name**— The name of the queue on the NetWare printer server.

15.2 Testing the Printer

There are some options under the Test menu to help diagnose the printer. If you have a local printer, try the Print ASCII directly to port option (this won't work with remote or SMB printers). If this doesn't work, check the cabling and printer configuration. Also, check to see that you selected the correct port (lp0, lp1, etc).

Next, try Print ASCII test page. If this doesn't work, check to see that `lpd` (the printer program for Linux) is loaded and make sure the remote printer is working and is set up properly. There is an option to restart `lpd` on the menu (see [Figure 15.5](#)).

Figure 15.5. Restarting `lpd` with PrintTool.



Next, try Print PostScript test page. If the ASCII test page works and Postscript doesn't, there is a problem with your printer driver. Make sure you have the proper input filter (or driver) selected. Next, check to see if Ghostscript is installed properly. If this isn't the case, try a different filter. Many printers are compatible with other print drivers. See the Web sites listed above for printer compatibility and Ghostscript troubleshooting.

Sometimes, turning off the Suppress Headers option can cause a printer not to print. Try turning it on if your printer still doesn't work.

If you are getting stair-stepped text on a text-only printer, try turning on the Fix Stair Stepping option. If it looks like the printer is accepting data (most printers have a blinking data light) and nothing prints, try turning on the Send EOF option.

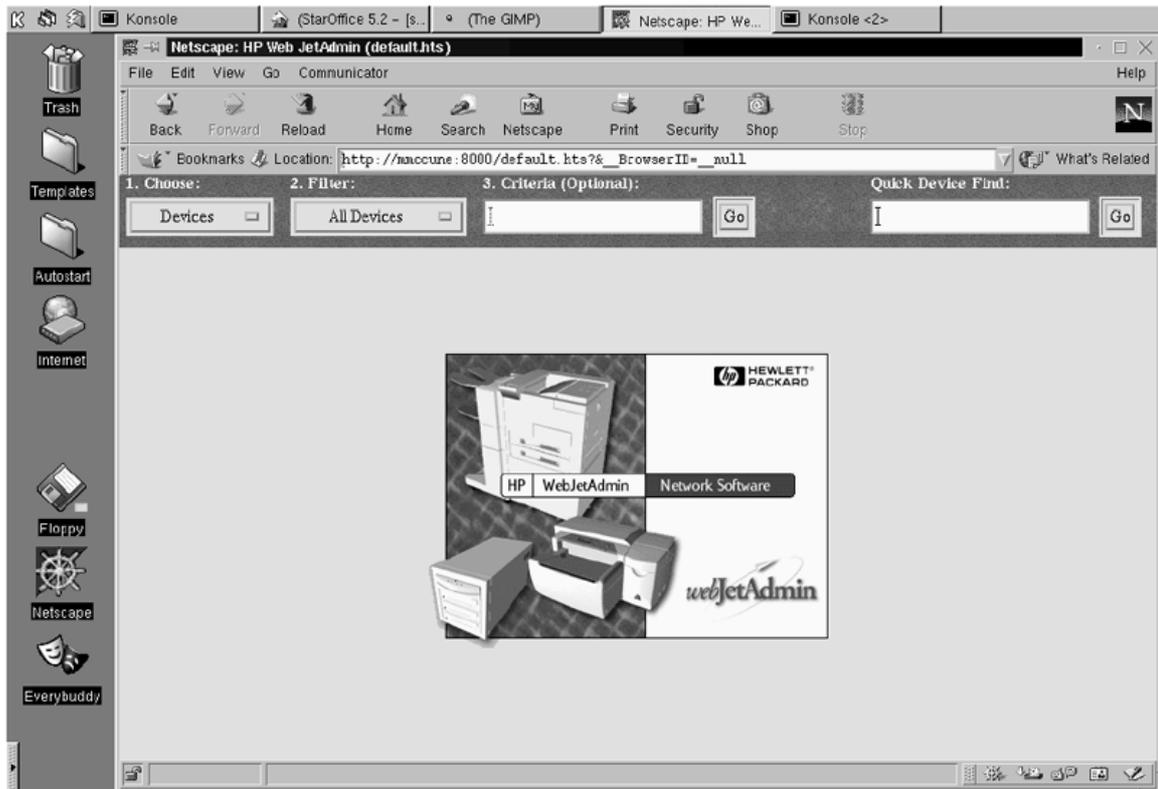
Hewlett Packard (HP) JetDirect printers need to send non-text output as raw on the queue option, which is set on the JetDirect port. A JetDirect printer is an HP printer that has a port that connects it directly to the network.

To set up a JetDirect card, first set up the printer's IP address. This is done through the menu on the front panel. See your printer's instruction manual for instructions.

There are two ways to set up the JetDirect card. One is to simply telnet into the JetDirect card's IP address. Once connected, type `?` for help or `/` for a list of current settings. To change a parameter listed, type: `parameter name: setting`.

For a more advanced administration tool, go to http://www.hp.com/support/net_printing and download the HP Web JetAdmin for Linux (see [Figure 15.6](#)). This will allow you to administer JetDirect cards from your browser window.

Figure 15.6. HP Web JetAdmin for Linux.



Once JetDirect is configured, set it up as a Remote UNIX (lpd) Queue in `printtool`. Be sure to leave the Input Filter field blank.

15.3 Setting up Samba for Printing

Once the `printcap` file is set up, the only thing needed to set up the `printers` for Samba is to add the printers section to the `smb.conf` file.

```
[printers]
  path = <path of spool file>
  printable = yes
```

These are the bare essential parameters for the `printers` section. This will load the `printcap` and allow Samba printing.

A few details can trip up Samba printing if they are not taken care of. Make sure the directory in the `path` parameter is writable by everyone. This allows Samba to write its printer files to the directory. If the `path` variable isn't writable, use `chmod` to change the permissions:

```
chmod 622 <path of spool file>
```

Also, `/dev/null` needs to be writable. Samba needs this file to discard output from external commands. If it isn't writable, use `chmod` again:

```
chmod 622 /dev/null
```

The default setup for Samba will use the Linux print driver. If this is okay, you won't need to do anything else. There are some cases when you would want to use the client's print driver. This is usually the case with Windows clients, since Windows supports a wider selection of printers than Linux.

To use the client's print driver, first you must create a `printcap` entry for it. Use `printtool` to create a printer named `raw`, but don't set up an input filter. Set up the device normally, e.g., `/dev/lp0` for parallel port 1, etc. Then save the setup. It will create an entry similar to this in the `printcap` file:

```
raw:\
  :sd=/var/spool/lpd/raw:\
  :mx#0:\
  :sh:\
  :lp=/dev/lp0:
```

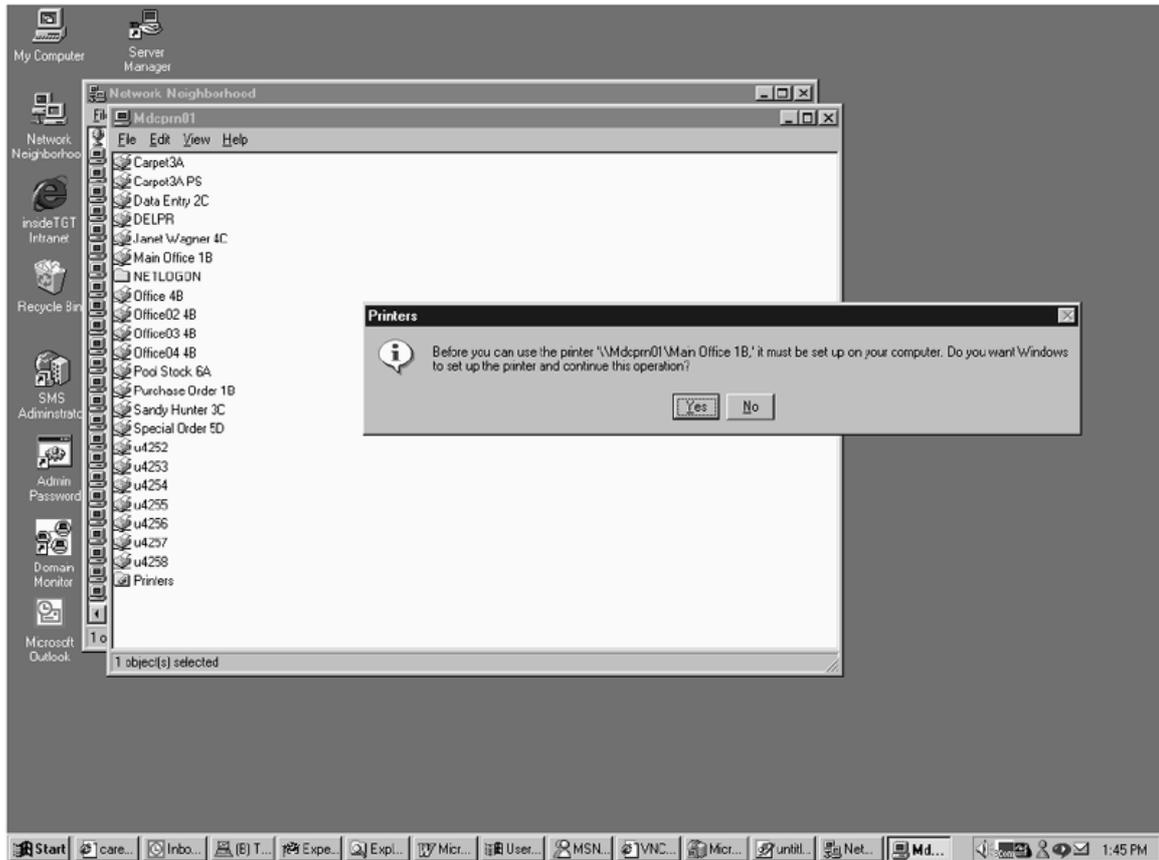
Next, set the `print command` parameter as shown below in the `printers` section. This will take print files sent by the clients and send them to the printer without processing them with the Linux print driver:

```
[printers]
  path = <path of spool file>
  printable = yes
  print command = lpr -P%p %s
```

15.4 Automatic Print Driver Installation

Windows allows the printer driver to be installed automatically from the server. When the client clicks on the Samba printer, it will ask if you want to install the printer driver (see [Figure 15.7](#)).

Figure 15.7. When setting up a network printer, Windows asks if you want to install the printer drivers.



There are three steps to setting up automatic printer driver installation. First, create a directory in which to store the printer drivers. For this example, we will use `/usr/local/samba/printers`. Then, create a share in the `smb.conf` for this directory:

```
[printer$]
  path = /usr/local/samba/printer
  public = yes
  writable = no
  browseable = yes
```

Next, create a list of drivers for the specific printer being used. The drivers that come with Windows are listed in `msprint.inf`, `msprint2.inf`, and `msprint3.inf`. These files are located in the `\windows\inf` directory on a Windows 95/98 PC. They are standard text files. Open one with your favorite text editor (like notepad) and look for your printer. For an updated driver, the `*.inf` file will be on the setup diskette or CD that comes with the printer. This file is usually called `oem*.inf`, although it may vary with different printers.

Next, locate the directory in which Samba is storing the `printers.def` file. This is usually the `/usr` or the `/usr/local/samba/lib` directory, but it may vary with different setups.

To find the file, use the command `locate printers.def`. If this doesn't locate the file, run the command `updatedb`, then try `locate printers.def` again. If the file doesn't exist, you must create one.

Next, copy the file that contains your printer setup (the `msprint.inf` or `oem*.inf`) to the directory on the Samba server that contains `printers.def`. If you copy this file to a diskette and then have trouble reading the diskette in your Samba server see the chapter "Mounting Windows Partitions with Linux." These files can also be copied through the network (see [Chapter 13, "Connecting Linux to Windows PCS"](#)). If

`printers.def` doesn't exist, copy the file to the `lib` directory for Samba, which is in `/usr/local/samba/lib` when compiling Samba and in the `/etc` directory for Red Hat-based distributions. After this is done, add the printer setup to `printers.def`. For this example, we will set up an HP Laserjet 5 printer:

```
make_printerdef MSPRINT3.INF "HP Laserjet 5" >> printers.def
```

Replace `msprint3.inf` with your `*.inf` file and replace "HP Laserjet 5" with the entry in your `*.inf` file for your printer.

After `printers.def` is updated, add the following to your `smb.conf`:

```
[global]
  printer driver file = /usr/local/samba/lib/printers.def
  ;This points to the printer.def file we just modified or
  created.
[lp] ;This is the printer name
  browseable = yes
  printable = yes
  writable = no ; Make it read only
  create mode = 0700 ; Files created here can only be access
  by the
  files' owner
  printer driver = HP Laserjet 5 Printer
  ;The name of the printer exactly as above.
  driver location = \\%h\PRINTER$
  ; %h translates into the computer name. This points to
  the share we
  ; created to store the drivers
```

Finally, we need to copy the drivers to the shared `printers$` directory we created earlier. This shared directory is pointing to the directory `/usr/local/samba/printer`. Type the `make_printerdef` command as above, except do not include the `>> printers.def` command. For example, for the Laserjet 5 listed above, type:

```
make_printerdef MSPRINT3.INF "HP LaserJet 5"
Found:PCL5EMS.DRV.BIDI
End of section found
CopyFiles: @PCL5EMS.DRV,@PJLMON.DLL,UNI,FINSTALL
Datasection: UNI_DATA
Datafile: PCL5EMS.DRV
Driverfile: PCL5EMS.DRV
Helpfile: UNIDRV.HLP
LanguageMonitor: PJL Language Monitor
```

Copy the following files to your `printer$` share location:

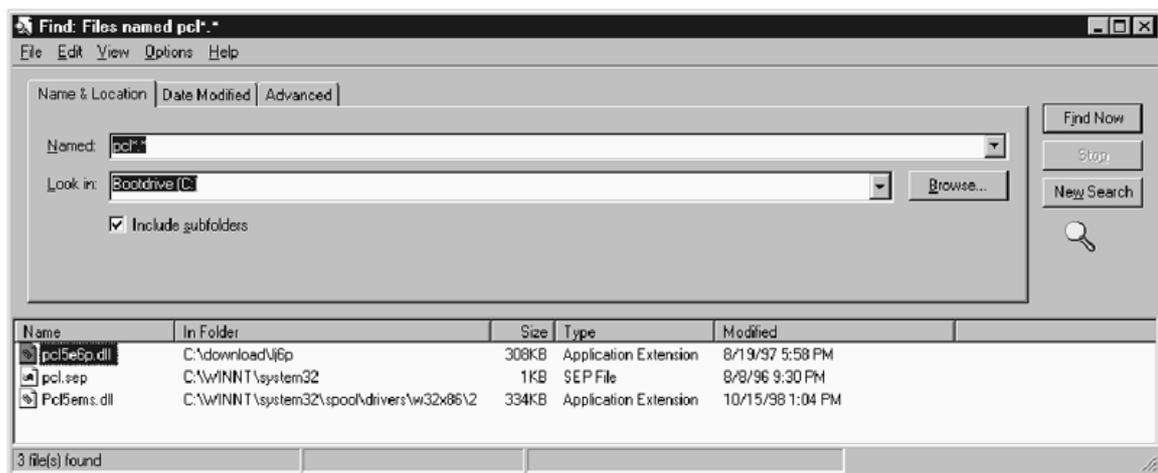
- PCL5EMS.DRV
- PJLMON.DLL
- UNIDRV.DLL
- UNIDRV.HLP
- ICONLIB.DLL
- FINSTALL.DLL

```
• FINSTALL.HLP
  HP LaserJet
5:PCL5EMS.DRV:PCL5EMS.DRV:UNIDRV.HLP:PJL Language

Monitor:EMF:PCL5EMS.DRV,PJLMON.DLL,UNIDRV.DLL,UNIDRV.HLP,ICONLIB.D
  LL,FINSTALL.DLL,FINSTALL.HLP
```

These files will need to be copied from an existing Windows 95/98 machine. First, install the proper driver onto the machine. Then locate the files on the PC. Most of the files are located in `\windows\system` (for Windows 9x) or `\winnt\system32` (for Windows NT and 2000). If they are not there, use the `find` function to locate them as shown in [Figure 15.8](#):

Figure 15.8. The Windows Find Program.



Once everything is configured, restart the `lpd` daemon with the `restart lpd` option in `printtool`. Then restart Samba:

```
samba stop
samba start
```

As usual, test the installation to make sure that everything works properly. When a Windows 95/98 PC connects to the Samba printer share, you should get a screen dialog box asking to install the print driver ([Figure 15.7](#)).

This should install the printer driver. If this doesn't work, check for typos and obvious errors. You should also be able to print to the shared printer. If this doesn't work, use the troubleshooting steps used earlier in the chapter.

Chapter 16. Using NFS and NIS in Linux and Windows

NFS is the native UNIX protocol that allows UNIX machines to share drives through a network. It performs some of the functions as Microsoft's SMB protocol. NFS is a much simpler protocol since it doesn't include authentication and printing. Authentication is handled by the UNIX (or Linux) hosts and printing functions are handled by `lpr` and `lpd`.

NFS is a good protocol for mixed networks. It is not only used by Linux and all other forms of UNIX, it is also available for all other major operating systems, including those for mainframes, midranges, and PCs. It is a much simpler protocol than SMB, and thus uses less system and network resources. It is not a persistent connection, which means that if an NFS client loses contact with the NFS server, it will reconnect when the server comes back online.

Linux has native (built-in) NFS support in the kernel, along with utilities to manage NFS. NFS allows a Linux machine to share a drive or directory with another machine (act as a server). It also allows the Linux machine to mount drives on other machines (act as a client).

Windows doesn't have NFS support, but there are many third-party programs that add NFS support to Windows. There are two ways to connect a Windows machine to an NFS machine. One is to load an NFS program on each Windows machine. The second is to use a machine as a gateway between the NFS machine and the Windows machine. The programs and methods are discussed later on in the chapter.

16.1 Setting up Linux as an NFS Server

Setting up Linux as an NFS server will allow you to share drives and directories on the Linux machine. Although NFS support is included with Linux distributions, it still needs to be set up.

The first thing that needs to be done is to start the `portmap` service. Most distributions start the `portmap` by default. To see if the `portmap` is started, type:

```
rpcinfo -p
  program vers proto  port
  100000    2    tcp    111  portmapper
  100000    2    udp    111  portmapper
```

You should see the `portmapper` attached to the `tcp` and `udp` ports. If the `portmapper` is not started, you need to add it to your `init` scripts. You may need to check your manual or reference book to see how to do this, since it varies with different distributions. The `portmap` program is named `portmap`, `rpc.portmap`, or `rpcbind` and is located in the `/sbin` or `/usr/sbin` directory. The `init` scripts are in the `/etc/rc.d`, `/etc/init.d`, or `/etc/rc.d/init.d` directories.

For example, in the Red Hat distribution of Linux, you would add the script `portmap` to the `/ect/rc.d/init.d` directory. The `portmap` script that comes with Red Hat is listed below:

```
#!/bin/sh
#
# portmap          Start/Stop RPC portmapper
#
# chkconfig: 345 11 89
```

```
# description: The portmapper manages RPC connections, which
are used by \
#             protocols such as NFS and NIS. The portmap
server must be \
#             running on machines which act as servers for
protocols which \
#             make use of the RPC mechanism.
# processname: portmap

# Source function library.
. /etc/rc.d/init.d/functions

# Get config.
. /etc/sysconfig/network

# Check that networking is up.
if [ $ {NETWORKING} = "no" ]
then
    exit 0
fi

[ -f /sbin/portmap ] || exit 0

# See how we were called.
case "$1" in
    start)
        echo -n "Starting portmapper: "
        daemon portmap

        echo
        touch /var/lock/subsys/portmap
        ;;
    stop)
        echo -n "Stopping portmap services: "
        killproc portmap

        echo
        rm -f /var/lock/subsys/portmap
        ;;
    status)
        status portmap
        ;;
    restart|reload)
        $0 stop
        $0 start
        ;;
    *)
        echo "Usage: portmap
{start|stop|status|restart|reload}"
        exit 1
esac

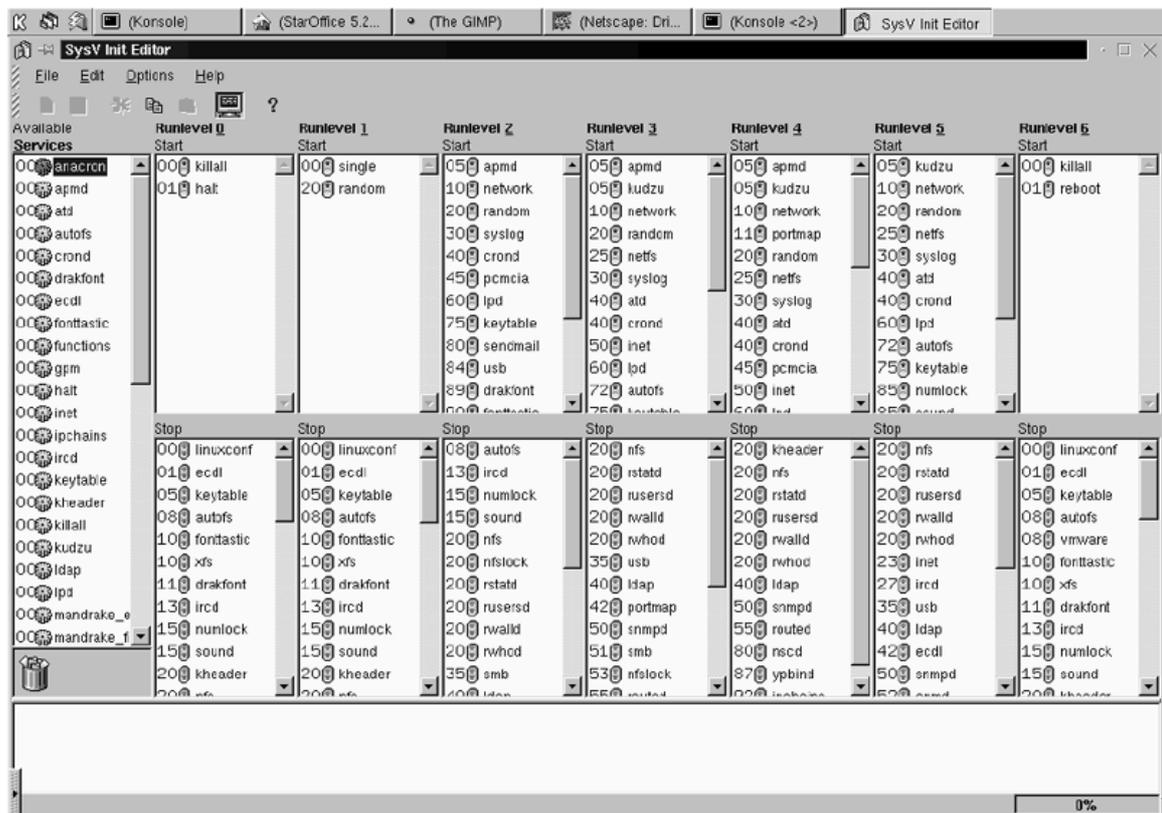
exit 0
```

Adding this script to the `/etc/rc.d/init.d` directory will start the `portmap` services automatically on startup. Please note that most Linux distributions come with the `chkconfig` utility to configure the startup scripts. To turn on `portmap`, type `chkconfig portmap`. To see the services started, type `chkconfig -list portmap`:

```
chkconfig --list portmap
portmap 0:off 1:off 2:off 3:off 4:on 5:off 6:off
```

Red Hat and its variants come with a SysV Init Editor that is located on the Start menu under the System folder. To start the `portmap` service, simply drag the `portmap` to the Runlevel 3 Start window as shown in [Figure 16.1](#).

Figure 16.1. System V(5) Init Editor.



If you prefer the command line, you can use `chkconfig` to turn the `portmap` on:

```
chkconfig --level 3 portmap on
```

Next, we need to define what directories to share. This is done by adding the directories to be shared to the `/etc/exports` file. For example, if we wanted to give read and write access of the directory `/home/alisa` to the computer `alisa`, we would add the following line to export:

```
/home/alisa alisa (rw, root_squash)
```

We also added the `root_squash` options. These prevent a user that is logged on as root on a client machine from gaining root access to the server. The `root_squash`

options should be enabled by default, but it is still a good idea to add them. You can also allow read-only access to the directory by changing `rw` to `ro`.

Finally, start `mountd` and `nfsd`. First, check to see if they are already loaded by running the following:

```
rpcinfo -p
program  vers  proto  port  program
100000   2     tcp    111   portmapper
100000   2     udp    111   portmapper
100024   1     udp    671   status
100024   1     tcp    673   status
100011   1     udp    681   rquotad
100011   2     udp    681   rquotad
100005   1     udp    690   mountd
100005   1     tcp    692   mountd
100005   2     udp    695   mountd
100005   2     tcp    697   mountd
100003   2     udp    2049  nfs
100021   1     udp    1024  nlockmgr
100021   3     udp    1024  nlockmgr
100021   1     tcp    1024  nlockmgr
100021   3     tcp    1024  nlockmgr
```

Notice that `nfs` and `mountd` are already loaded. We now need to reload the export file by running `exportfs -av`. If you don't have the `exportfs` utility, simply restart `nfs` and `mounted` manually (you may have to modify this procedure slightly for your system):

```
killall -HUP /usr/sbin/rpc.mountd
killall -HUP /usr/sbin/rpc.nfsd
```

If `mountd` and `nfs` aren't set to load automatically, set them up in the `init` like we did for `portmap`. For Red Hat, you can use the SysV Init Editor to add the `nfs` services to the Runlevel 3 Start, or you can use `chkconfig`:

```
chkconfig --level 3 nfs on
```

If you get a message like "rpcinfo: can't contact portmapper: RPC: Remote system error - Connection refused" or "No remote programs registered," the `portmap` isn't running or has stopped running. Kill and restart the `portmap`, `mountd`, and `nfsd` services.

This is a very simplified explanation of NFS that covers the basics. There are entire books covering the details of NFS.

16.2 Using an NFS Client on Linux

For NFS to work, NFS support has to be built into the kernel. Most distributions have NFS support built into the kernel. If not, NFS support needs to be added by recompiling the kernel or by adding an NFS kernel module. As stated before, recompiling kernels is beyond the scope of this book. There are many books with details on compiling kernels, or you can go to <http://www.ibiblio.org/mdw/HOWTO/Kernel-HOWTO.html> for details on compiling kernels.

Once NFS is activated in the Kernel, the remote NFS drive can be accessed by simply using the `mount` command:

```
mount -o rsize=4096,wsiz=4096 mmccune:/mnt/home/alisa
/mnt/alisa
```

Of course, you must be root or equivalent to issue the `mount` command. Let's go over what each of these parameters means:

- `-o`— Indicates that there are options to follow.
- `rsize`— The size of the read buffer in bytes.
- `wsiz`— The size of the write buffer in bytes.
- `mmccune`— The NFS server's name.
- `/mnt/home/alisa`— The NFS directory on the server.
- `/mnt/alisa`— The local directory that the NFS directory is mounted onto.

16.2.1 Optimizing NFS

The `rsize` and `wsiz` parameters can be optimized for greater speed. Unfortunately, the only way to do this is by trial and error. Generally, `rsize` and `wsiz` should be as large as possible without sacrificing reliability. Most of the time, it is better to just leave these at the default values.

Optimizing `rsize` and `wsiz`. Start with an `rsize` and `wsiz` of about 1024 and see how long it takes to copy a very large file (64MB or more). Do this several times and average the times. Then add 1024 to the values of `rsize` and `wsiz` and repeat the process until both values are 16,384 bytes.

After you have determined the optimal sizes for these values, we need to make sure they are reliable. This is rather hard to catch. Try running several `ls` commands. Then try reading some large files. If these programs abort or fail without an error message, it is a good sign that `rsize` and `wsiz` are too large and need to be smaller. Most Linux systems perform best when `rsize` and `wsiz` are about 4096, although your results may vary.

16.2.2 Hard and Soft Mounts

The way NFS handles a lost connection is also important. There are hard and soft mounts. A soft mount will send an error to the programs that are accessing the NFS volume. Some programs can handle this, but most can't.

A hard mount is a better option in most cases. It will suspend any programs that are accessing the NFS mount. It is added with the `intr` option, which will stop the program that is using the NFS mount. A `hard intr` is the best option in most cases.

Once the optimum values are determined, they can be added to the `fstab`. The `fstab` is where the drive parameters are defined in Linux. The entry to the `fstab` will have this order, separated by spaces: device, mount point, filesystem type, options separated by commas, `dump`, and `fsckorder`. For example:

```
mmccune:/mnt/home/alisa /mnt/alisa nfs
rsize=4096,wsiz=4096,hard,intr 0 0
```

In this example, the NFS drive is the directory `/mnt/home/alisa` on the server `mmccune`. The `rsize` and `wsize` are 4096 bytes with the `hard` and `intr` options set on. The `dump` and `fscheck` are both set to 0, since this is not a local drive.

16.3 Using NFS on Windows

NFS support is not built into any version of Windows, although there are many add-on programs that allow Windows to function as an NFS server or client. Two well-known commercial packages are Hummingbird's NFS Maestro and NetManage's Chameleon NFS. Microsoft also makes a "Windows for UNIX" add-on pack for Windows NT. There are about a dozen other packages and each has its own strengths and weaknesses. I will not cover all of them here, but I will give an overview of how to implement NFS. Most of the companies that make NFS software for Windows offer evaluation copies, so you can look at a package before making a decision on whether or not to use it.

I will cover the "Windows for UNIX" add-on pack by Microsoft, since this package is often installed on NT machines and it is a necessity in mixed Windows and Linux (or UNIX) environments.

16.3.1 Setting up an NFS server on Windows NT

To create an NFS share on a Windows NT machine, go to the Control Panel and open the Server for NFS applet. From here, choose which drives or directories to share. A drive on NT will appear as `/drive/` on the NFS share. For example, the `D:` drive on NT will appear as `/D/` in the NFS share. You can also use aliases to make share names more manageable. For example, a shared drive of `/D/public/user/share` could be aliased as `share` to make the share easier to remember. The only problem with aliases is that the server needs to be rebooted before the alias takes effect.

You can also set up client groups on NFS shares. This simplifies the administration of NFS shares, since it allows users to be administered as groups. Keep in mind that the security models of Windows and Linux are different. Windows users can belong to multiple groups whereas Linux users can only belong to one group at a time. This normally will not cause any problems, but under certain circumstances, it can cause problems if you are not careful. If this is a problem, you can use `chgrp` to change the group ID.

There is also an option for file locking. When this is enabled, it only allows one user at a time to write to a file. Many programs such as databases implement their own file locking. You would not want to enable this option if a program had its own file locking. Otherwise, it is a good idea to enable this option on files that are shared by multiple users.

The User Name Mapping service is provided with the Windows for UNIX package. It runs as a standard Windows service and it maps UNIX User Ids (UIDs) to Windows Session IDs (SIDs). It maintains the mappings in a table as in the example below:

Windows Username	Windows Domain	UNIX Username	UNIX Domain	UID/GID
JohnDoe	Indwindows	Johnd	Indunix	1090/201
Maryjane	Indwindows	Maryj	Indunix	1223/201

There are also several commercial NFS implementations for Windows. They have the advantage of working under Windows 95/98. Some of these commercial programs are Hummingbird NFS Maestro (<http://www.hummingbird.com/products/nc/nfs/>), NetManage Interdrive (<http://www.netmanage.com/products/>), Omi NFS (<http://www.xlink.com>), and Access NFS (<http://www.ssc-corp.com/nfs>). These commercial programs offer demo programs on their Web sites if you want to try them out.

There is a free implementation of NFS for Windows NT 4 called SSSNT. It doesn't have all the bells and whistles of the commercial programs, but it is relatively easy to install

and set up. It also doesn't run as a system process, so it must be run from a batch file or at the command line. It is free and open source, however, and is available at <http://www.loa.espci.fr/winnt>. This Web site also contains several other free UNIX utilities for Windows.

16.3.2 Setting up an NFS Client

Use the UNIX NFS Client applet in the Control Panel to set up the NFS client for Windows. First, choose the method of authentication. You can use Network Information Services (NIS) or PC NFS Daemon (PCNFSD). NIS is an authentication used on large NFS networks to allow centralized storage of username, passwords, and shares. It is described later in the chapter. If you are using NIS, choose the NIS option and fill in the NIS server name. If you are not using NIS, choose PCNFSD.

Next, choose when the client will be authenticated. You can choose to authenticate at logon or to authenticate when the share is accessed. This is simply a matter of choosing whether to take longer to log on or to take longer to access a share.

Like the Linux client, you can choose the read and write buffer sizes (the default is 64K), the timeout, number of retries, and hard or soft mounts. These can be tuned using the same method as we used to tune the Linux client. As with the Linux client, hard mounts usually work better.

The Windows NFS client also supports NFS version 3. If the NFS server supports version 3, this will provide a significant performance boost on reliable network connections. The downside is that an NFS server crash can lead to file corruption.

By default, the NFS client allows read, write, and execute by the owner of a file and read and execute for all others that have rights to the file. These defaults can be changed in the NFS client setup.

Filename mapping is important in mixed Windows and Linux environments. Linux treats upper- and lower-case letters differently and Windows treats them as the same letter. For example, Linux would treat the filenames `Sample.txt` and `sample.txt` as different filenames, but Windows would treat them as the same filename. There are several ways to set up filename mapping, but the easiest method is to have the Windows NFS client create filenames in lower-case but ignore case when reading files.

The Windows NFS client supports any symbolic links on the server. A symbolic link is a file that points to another file. It is similar to a shortcut on a Windows machine. It can point to a local file or even a file on another machine. By default, the Windows NFS client does not resolve or display unresolved links. It also will not allow you to rename or delete symbolic links. Do not change this unless you have a good reason. Renaming or deleting symbolic links is not a good idea unless you fully understand how symbolic links work. For example, if you delete a symbolic link and re-create a file with the same name, it is no longer a symbolic link. You will have two files instead of a file that points to another file.

16.3.3 Connecting to an NFS Server with Windows

You can find out the NFS shares on a server by typing `show mounts <server name>` on the client. You can then connect to a share by using either the Windows syntax or the UNIX syntax. For example, the next available drive to the share `home` on the server `mmccune`, can be attached from the command prompt in two different ways. This is done to allow compatibility with both UNIX and Window systems.

```
net use * \\mmccune\home (using the Windows syntax)
net use * mmccune:/home (using the UNIX syntax)
```

Once this is done, the share looks and acts like any other network drive on the system.

16.3.4 NFS Gateways

An NFS gateway allows Windows machines to connect to an NFS server through a gateway machine. The advantage of this is that the NFS client software doesn't have to be loaded on all the machines. The disadvantage is that an NFS gateway won't handle high volumes of network traffic. This makes it a good solution for Windows clients that only need occasional access to the NFS server.

An NFS gateway can be set up on a Linux or Windows NT machine. To set up a Linux NFS gateway, first set up the NFS client to mount the NFS server export. Then set up Samba to share the NFS mount. The NFS share will look like a normal share on the server to the Windows clients.

For Windows NT, you will need a program like Hummingbird NFS Maestro Gateway, Access NFS Gateway, or Reflections NFS Gateway. These programs will allow Windows machines to connect to the NFS export through the Windows NT machine running the gateway software.

16.3.5 NFS Security

There are many widely available tools that exploit the security vulnerabilities of NFS. Any connection to the outside should be through a firewall or another security device.

This is not necessarily an NFS security function, but a good general security policy. Edit the `/etc/inetd.conf` file and take out any services you don't need. Any service represents a potential security risk. For example, `rlogind`, `rshd` and `rexecd` allow you to log in and execute programs remotely. While these are handy, if you don't need them, comment them out like the code shown below:

```
#shellstreamtcpnowaitroot/usr/sbin/tcpdin.rshd
#loginstreamtcpnowaitroot/usr/sbin/tcpdin.rlogind
#execstreamtcpnowaitroot/usr/sbin/tcpdin.rexecd
```

Also, make sure you have all the current versions and patches of all programs on your systems. Security holes are fixed and released as security patches. Check the Web sites of Microsoft and your Linux distribution for current security fixes. Also check <http://www.cert.org> or <http://www.securityfocus.com> for current security advisories.

As far as NFS-specific security issues, pay particular attention to the `/etc/exports` file. Here are seven rules for exports taken from CERT advisory CA-94:15:

1. Do *not* self-reference an NFS server in its own `exports` file.
2. Do not allow the exports file to contain a "localhost" entry.
3. Export filesystems only to hosts that require them.
4. Export only to fully-qualified host names.
5. Ensure that export lists do not exceed 256 characters. If you have aliases, the list should not exceed 256 characters *after* the aliases have been expanded. (See CA-94:02.REVISED.SunOS.rpc.mountd.vulnerability.)
6. Use the `showmount (8)` utility to check that exports are correct.
7. Wherever possible, mount filesystems to be exported as read-only and export file-systems as read-only.

16.4 Setting up an NIS Server on Linux

As mentioned above, NIS is a way of centrally managing usernames, passwords, and groups. It is used by all flavors of UNIX, including Linux. Since most Linux distributions don't come with NIS, the first thing that needs to be done is to download it. The current version can be downloaded at <ftp://ftp.kernel.org/pub/linux/utils/net/NIS>. There are also

binary RPM (Red Hat Package Manager) files available at <http://rpmfind.net/linux/RPM>. Just search for `nis-server` and choose the one with the highest version number. Before NIS is installed, you must have `portmap` running (see above). It is also a good idea to have the `time` service running. This can be started by adding the following lines to `/etc/inetd.conf`:

```
time    stream    tcp      nowait  root    internal
time    dgram    udp      wait    root    internal
```

To install the NIS server, first extract the files. The archive will have a name similar to `ypserv-1.3.9.tar.gz`. Extract the archive as follows:

```
gunzip ypserv-1.3.9.tar.gz
tar xvf ypserv-1.3.9.tar
```

This should create a directory called `yp-1.3.9`. There are two ways to compile `ypserv`: with TCP Wrappers or with `securenet`. These programs allow you to limit the computers that have access to the NIS database. The default is `securenet`. You can also use TCP Wrappers, which is more configurable, but harder to configure and more problematic. TCP Wrappers also has problems with memory leaks, which can cause system problems. For demonstration purposes, we will use `securenet`.

Go to the `yp-1.3.9` directory and type `./configure`, then `./BUILD`. If this is going to be a master NIS server, you need to edit `/var/yp/Makefile`. A master NIS server would store the main copies of the NIS database. A slave NIS server would store copies of the master database. If this is the only NIS server, it is obviously going to be a master. A slave is used on large networks to share the load with the master or if there is a slow network connection to the master server.

If you are running a master NIS server, pay particular attention to the `all` section. This is where settings are put into the NIS database. Add and delete any information that needs to be shared. The default settings for the `all` section are:

```
all: passwd group hosts rpc services netid protocols netgrp
mail \
    shadow publickey # networks ethers bootparams printcap \
    # amd.home auto.master auto.home passwd.adjunct
```

Once this is done, go to the object directory. It should be called something like `obj.<os version>`. Type `make install`. This should compile all the programs needed for NIS.

After NIS is compiled, we need to edit the configuration files `/var/yp/securenets` and `/etc/ypserv.conf`. These files may be in the `ypserv-1.3.9` directory and may need to be moved to the proper directories.

`securenet` allows you to limit the hosts that can access the NIS database. The default is to allow everything access, which is usually not what we want. The default `securenet` file looks like this:

```
#
# securenet This file defines the access rights to your NIS
server
# for NIS clients. This file contains netmask/network
# pairs. A client's IP address needs to match with at
least
# one of those.
#
```

```
# One can use the word "host" instead of a netmask of
# 255.255.255.255. Only IP addresses are allowed in
this
# file, not host names.
#
# Always allow access for localhost
255.0.0.0127.0.0.0
# This line gives access to everybody. PLEASE ADJUST!
0.0.0.00.0.0.0
```

Delete the last line and add the hosts that you want to have access to the NIS database. You can specify individual hosts or entire subnets. For example, to give access to all hosts from 10.0.0.1 to 10.0.0.255, add the following line:

```
255.255.255.0          10.0.0.0
```

You can also add individual hosts instead. To add 10.0.0.1 and 10.0.0.2, type the following lines:

```
255.255.255.0    10.0.0.1
255.255.255.0    10.0.0.2
```

Next, edit the `/etc/ypserv.conf` file. This will allow you to increase the security of the NIS database. You can leave the file unedited and NIS will still work normally.

One of the options will allow password shadowing. This will conceal the password from the users but all NIS servers on the network must be running the same version of NIS and have this enabled.

Another option allows the centralized management of NIS keys. If it is not turned on, you will have to manually update the keys on each NIS server on the network.

There are several other options for `ypserv.conf`. See the `man` page for the full options (`man ypserv.conf`).

Once configuration is done, `ypserv` can be started. There is a startup script for Red Hat, `ypserv-1.3.9/etc/ypserv.init`. When it is started, you could get something like this:

```
program 100004 version 1 ready and waiting
program 100004 version 2 ready and waiting
```

You can add this to your `/etc/rc.d/init.d` directory to have it started automatically.

Now to generate the NIS database on the NIS master, go to the `/usr/lib/yp` directory and type:

```
./ypinit -m
```

On an NIS slave, make sure the NIS client is working (see below). Go to `/usr/lib/yp` and type:

```
./ypinit -s masterhost (where the masterhost is the NIS
master's
host name)
```

There are two other programs started in `ypserv.init`. The `rpc.ypxfrd` program copies the NIS master's database to the slaves. The `rpc.yppasswdd` updates the NIS password database when a user changes his or her password.

This is just an introduction to NIS. If you plan on using NIS for a production system, get a good reference on both NFS and NIS.

16.5 Setting up an NIS Client on Linux

Unlike the NIS server, most Linux distributions come with an NIS client. If it is not included, go to <ftp://ftp.kernel.org/pub/linux/utils/net/NIS> and download the following files: `yp-tools-2.2.tar.gz`, `ypbind-mt-1.4.tar.gz`, `ypbind-3.3.tar.gz`, and `ypbind-3.3-glibc5.diff.gz`. You will also need to go to <ftp://uni-paderborn.de/linux/local/yp> and download `yp-clients-2.2.tar.gz`. Then, compile and install the programs.

Of course, you can always go to <http://rpmfind/linux/RPM> and find the latest version of `nis-client` if you don't feel like compiling the program.

Once the NIS client is installed, edit `/etc/yp.conf`. This simply tells the NIS client the name of the NIS server. Below is a sample file showing the different settings:

```
# /etc/yp.conf - ypbind configuration file
# Valid entries are
#
#domain NISDOMAIN server HOSTNAME
#   Use server HOSTNAME for the domain NISDOMAIN.
#
#domain NISDOMAIN broadcast
#   Use broadcast on the local net for domain NISDOMAIN
#
#ypserver HOSTNAME
#   Use server HOSTNAME for the local domain. The
#   IP-address of server must be listed in /etc/hosts.
```

It is a good idea to test the NIS client before adding it to the startup files. First, set the NIS domain. This is set on the NIS server and is different from the TCP/IP domain.

```
cd /bin
./domainname <NIS Domain>
```

Next, make sure the `portmap` is loaded (see above).

Make sure the directory `/var/yp` exists. If it doesn't, create the directory.

Go to the directory `/usr/sbin` and type `ypbind`.

Run `rpcinfo -p` to see if `ypbind` is running. You should see something like this:

```
program    vers    proto    port
100000      2       tcp      111     portmapper
100000      2       udp      111     portmapper
100007      2       udp      792     ypbind
100007      2       tcp      794     ypbind
```

Next, run `rpcinfo -u localhost ypbind`, which should give the following results:

```
program 100007 version 2 ready and waiting
```

Next, run `yycat passwd.byname`, which should show you the entire NIS database. This should confirm that the NIS client is working properly. It can then be added to `/etc/rc.d/init.d`.

There is one optional configuration file that can be added. `/etc/nsswitch.conf` determines the lookup order for NIS entries. This file performs the same function that the `/etc/host.conf` file performs for TCP/IP host names. The sample file below gives a good start:

```
#
# /etc/nsswitch.conf
#
# An example Name Service Switch config file. This file
# should be
# sorted with the most-used services at the beginning.
#
# The entry '[NOTFOUND=return]' means that the search for an
# entry should stop if the search in the previous entry
# turned
# up nothing. Note that if the search failed due to some
# other reason
# (like no NIS server responding) then the search continues
# with the
# next entry.
#
# Legal entries are:
#
#      nisplus          Use NIS+ (NIS version 3)
#      nis              Use NIS (NIS version 2), also
# called YP
#      dns              Use DNS (Domain Name Service)
#      files            Use the local files
#      db               Use the /var/db databases
#      [NOTFOUND=return] Stop searching if not found
#
# so far
#
passwd:      compat
group:      compat
# For libc5, you must use shadow: files nis
shadow:     compat
passwd_compat: nis
group_compat: nis
shadow_compat: nis
hosts:      nis files dns
services:   nis [NOTFOUND=return] files
networks:   nis [NOTFOUND=return] files
protocols:  nis [NOTFOUND=return] files
rpc:        nis [NOTFOUND=return] files
ethers:     nis [NOTFOUND=return] files
netmasks:  nis [NOTFOUND=return] files
netgroup:   nis
bootparams: nis [NOTFOUND=return] files
publickey:  nis [NOTFOUND=return] files
```

```
automount: files
aliases:   nis [NOTFOUND=return] files
```

One final note for NIS: Shadow passwords are not well-supported by NIS. A shadow password is often used by Linux to increase password security. Since they don't give any extra security to NIS, it is best not to use shadow passwords with NIS. See your user's manual for details on shadow passwords and whether your distribution uses them.

16.6 NIS Support for Windows

Windows cannot currently be used as an NIS server, but NISGina is an NIS client that is available for Windows NT. GINA (Graphical Identification and Authentication) is the logon screen for Windows NT. NISGina replaces the `msgina.dll` in Windows NT with its own version of the logon program `nisgina.dll`. Be very careful when setting up NISGina. Setting it up incorrectly will lock you out of the computer and make it unusable! (see <http://www.ldv.ei.tum.de/software/nisgina/>.)

First, download and extract the files. The newest version of NISGina uses Install Shield to automate the installation. Instructions for a manual installation follow.

Assuming that Windows NT is installed in the `/winnt` directory (the default), log in as administrator and copy the `*.dll` and `*.exe` files into the `/winnt/system32` directory.

First, install the `portmap` by typing:

```
installservice portmap
\winnt\system32\portmapservice.exe
```

You should get a message that the service has been installed successfully.

Create the directories `\winnt\var\run` and `\winnt\var\yp`.

Next, start the `portmap` service:

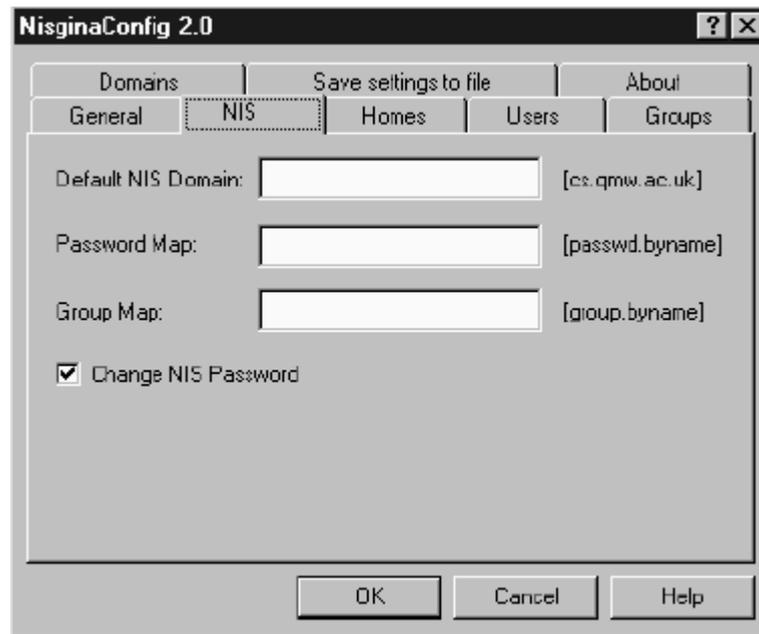
```
net start portmap
```

Install `ypbind` and specify that it will use the `portmap` services:

```
installservice ypbind portmap
\winnt\System32\ypbindservice.exe
```

Now, several Registry entries need to be modified. The easiest way to do this is with `NisginaConfig.exe`, which we copied earlier to the `\winnt\system32` directory. Some versions put an icon in the Control Panel. Once this program is started, go to the NIS tab (see [Figure 16.2](#) below).

Figure 16.2. The NISgina configuration screen.



Set the NIS Domain, Password Map, and Group Map. You can only configure one NIS domain at a time. Once the NIS domains are set, start `yplibind`:

```
net start yplibind
```

Next, we will test the NIS client with the familiar NIS tools: `ypwhich.exe`, `ypcat.exe`, and `ypmatch.exe`. These files were copied to the `/winnt/system32` directory earlier.

`ypwhich` will return the host name or IP address of the NIS server. If this doesn't work, check to be sure that the NIS server is working and that the NIS domain is correct. `ypcat <password map>` should show the NIS password file. If it doesn't, make sure that the password map is correct. For example, if the password map is `passwd.byname`, the command would be:

```
ypcat passwd.byname
xp_port 2f8. Addr 10.0.0.1, port 805
mmccune:x:500:500:Mike McCune:/home/mmccune:/bin/bash
    The rest was deleted for brevity.
```

Lastly, we will use `ypmatch` to find an individual account. If found, it should show the user's encrypted password:

```
ypmatch mmccune
xp_port 2f8. Addr 10.0.0.1, port 805
mmccune:encrypted passwd here:----- (You don't think
I'm going to
    show my password here, do you?)
```

We will now replace Microsoft's logon utility `msgina.dll` with the NIS logon utility `nisgina.dll`. Before we do this, make a recovery disk by running the program `rdisk` and choosing the option to update the recovery disk. We may need to use this if the following step is not done correctly.

Once this is done, open the Registry editor (either `regedit` or `regedt32`). Go to the Registry

hive `[HKLM\SOFTWARE\Microsoft\Windows`

`NT\CurrentVersion\WinLogon`]. If it exists, change the key `GinaDLL` to `nisgina`. If it doesn't exist, create the key.

Once this is done, close the Registry editor and reboot. You should get the familiar login screen. Pressing `<CTRL>-<ALT>-` should get the NISGina logon screen asking for a username and password.

Once you log on using the NIS database, NISGina will check to see if a local NT account exists with the same username. If one exists, NISGina updates the credentials of the local database. If it doesn't exist, NISGina creates a local account with the same username. This allows the username to log on even if the NIS database is not available. Of course, if you can't log on, NISGina is not set up properly and you need to use the rescue disk to recover and start over.

16.6.1 Problems with NISGina

Because NISGina creates a local account, there are problems with roaming profiles. This is because Windows NT treats local accounts on different machines as different accounts, even when the accounts have the same name.

NISGina also keeps a log file, `\winnt\nisgina.log`, which has usernames and encrypted passwords in it. It would be a good idea to make this file readable by the administrator only, since it could pose a security risk.

`ypcat.exe` and `ypmatch.exe` could also be used to view usernames and encrypted passwords and should be similarly restricted. This will not stop anyone from downloading or copying the files from elsewhere, but there is no way to restrict port usage in Windows NT. In Linux (and UNIX), the ports used by `ypcat` and `ypmatch` are restricted for use by the administrator or equivalent users.

16.6.2 NIS Security

NIS was not really designed to be secure. It was designed to allow the centralized management of usernames and passwords on a UNIX network. There are several things you can do to increase security, however.

As with NFS, any connections to the outside should be through a firewall or other security device. NIS should never be used on an unsecure connection.

If security is a real concern, NIS+ is a newer version of NIS that has much better security. The problem is that the NIS+ server for Linux is currently in development. The Linux NIS+ client is well supported, however. If you are already on an NIS+ network, you can download the Linux client at <ftp://ftp.kernel.org/pub/linux/utils/net/NIS+>.

Chapter 17. Implementing FTP, Telnet and Other UNIX Protocols in Windows

FTP (File Transfer Protocol) and telnet are among the oldest Internet protocols that allow two UNIX machines to transfer files (FTP) and log on to another computer (telnet). Each of these programs consists of a client and a server. The client is bundled with all current versions of Linux and Windows. The syntax is `ftp <remote machine>` or `telnet <remote machine>`. The remote machine has to have the FTP or telnet server running. Keep in mind that these protocols transmit information (including usernames and passwords) in the clear, so they would not be suitable for insecure networks like the Internet. If you need to use these over a secure network, use an encryption program such as Secure Shell (SSH), which is explained later in this chapter.

The FTP program has a text-based interface in both Linux and Windows. It is started at the command prompt by typing `ftp`. You can also type in the machine name (or IP address) such as `ftp ftp.cdrom.com`. The remote machine may have a username and password. For anonymous FTP servers such as `cdrom.com`, the username is `anonymous` and the password is your email address.

```
ftp ftp.cdrom.com
Connected to warchive.cdrom.com.
220 warchive.cdrom.com FTP server (Version DG-3.1.37 Sun
Jun 20
21:18:25 PDT 19
99) ready.
Name (ftp://ftp.cdrom.com:root): mmccune@linuxstart.com
331 Guest login ok, send your email address as password.
```

The commands for FTP and those for the `smbclient` that were discussed in [Chapter 14](#) are identical. Once connected, you can type `status` to check the settings. The `type:` setting is important. If it is set to ASCII (text) mode and a binary file is transferred, the file will get mangled. The ASCII setting will preserve the line feeds and character sets in text files.

Once in, the standard `ls` (or `dir`) and `cd` commands can be used to navigate the directories on the remote machine. For a full list of commands, type `?` or `help`.

```
ftp> help
!                debug          mdir            sendport        site $
dir              mget          put             size account
disconnect
mkdir           pwd           status append  exit            mls
quit            struct        ascii          form            mode
quote
system bell    get           modtime        recv
sunique binary
glob            mput          reget          tenex bye       hash
newer           rstatus      tick case     help            nmap
rhel           trace cd     idle          nlist          rename
type cdup      image        ntrans        reset          user
chmod
lcd            open         restart       umask close    ls
prompt         rmdir       verbose cr    macdef
passive
```

`runique` `? delete` `mdelete` `proxy` `send`

Many of these commands are rarely used. Two that you will need to know are `get` and `put`. `get` will transfer a file from the server to a local machine (download). `put` will transfer a file from your local machine to the FTP server (upload). You can also use `mput` and `mget` to transfer multiple files.

17.1 Setting Up the FTP Server for Windows

For Windows NT, you must first install the Internet Information Services. This is done by going to Start -> Programs -> Microsoft Internet Server (Common) -> Internet Information Server Setup. From here, install the FTP services.

After the FTP services are installed, go to Programs -> Microsoft Internet Server (Common) -> Internet Service Manager. When this is opened, you will see a list of services. Double-clicking on FTP Services will bring up the FTP Service Properties configuration utility.

The first tab is the Service tab. Under this tab are:

- **TCP Port**— The default is 21. Leave it at 21 unless you have a good reason to change it.
- **Connection Timeout**— This is the number of seconds until an idle connection is dropped. The default of 900 should work under most circumstances. Increase it if you have a slow connection that is being dropped due to the time-out.
- **Maximum Connections**— The default is 1000. This can be adjusted higher if your system can handle the load, or lower if it can't handle the load.
- **Allow Anonymous Connections**— This is the account that anonymous connections are logged in as. Set the rights on the account accordingly.
- **Allow Only Anonymous Connections**— Don't allow FTP connections under any other user account.
- **Comment**— This doesn't affect anything. It is used to add notes and comments.

The Messages tab allows you to add a welcome message, exit message, and a maximum connections message. These should be self-explanatory.

The Directories tab allows you to set the directories to which the FTP user has access. The alias is how the directory appears to the FTP client. The `<Home>` alias is the root directory of the FTP client. For example, let's say that the `<Home>` alias is set to the `c:\data` directory. When you type `cd /` (or `cd \`) in the FTP client, you will be in the `c:\data` directory on the server.

At the bottom of the Directories tab is a check box to choose between UNIX and MS-DOS directory listings. This just determines how the FTP client sees the directories. The UNIX style will look like the UNIX `ls` command and the MS-DOS style will look like the Windows `dir` command.

The Logging tab has options to enable a log file that keeps track of FTP activity. It logs the addresses of logons and failed logons as well as what files are transferred. Once logging is turned on, there are several parameters to set. Most of them are self-explanatory. One note on the log filename, though. You can use the parameter `d` (for day), `m` (for month), and `y` (for year) to automatically name the file after a day it was created. This is a good way to keep track of FTP log files, since there can be many of them in the log file directory.

The Advanced tab is for access control. This can be used to control which computers can access the FTP service. The Granted Access button will allow all computers to access FTP services except those listed. Conversely, the Denied Access button will block FTP access to all computers except those listed.

The configuration utilities on Windows 2000 are slightly different. To configure FTP, go to Start -> Programs -> Administrative Tools -> Configure Your Server. This will bring up a graphical window with a menu on the left-hand side. Under Web/Media Server, choose the Web Server option. In the middle of the windows, click the link to Start the Windows Components Wizard. Choose the Internet Information Services (IIS) and click Details. Click the check box to turn on the File Transfer Protocol (FTP) Server. Then, choose OK to return to the previous menu. Select Next and the needed files will be copied. Choose Finish and then Exit to exit the install wizard.

To configure FTP services, go to Start -> Programs -> Administrative Tools -> Internet Services Manager. This will pull up a program that looks similar to Explorer. From this, find your server, right-click on the icon, and choose Properties. Under the Master Properties drop-down menu, choose FTP Service and then Edit. This should bring up the same configuration box as the one we used to configure FTP for Windows NT.

17.2 Setting up FTP for Linux

FTP is a standard service with all Linux and UNIX machines. Most of the time, FTP is already running, but if it isn't, it is easy to start. Simply open `/etc/inetd.conf` (the location may vary) in your favorite text editor and add the line (or uncomment it if it is already there):

```
ftp    stream      tcp        nowait   root
/usr/sbin/tcpdin.ftpd -l -a
```

You may have to change the path of the `tcpd` if it is not located in `/usr/sbin`.

Linux uses regular user accounts to handle FTP logins. If you want, you can create a special account and directory for FTP access.

You can also restrict access by only allowing certain hosts to access remote services via the `hosts.allow` and `hosts.deny` files. These are simple text files with a list of host names that can be used to restrict remote services that are started in the `inetd`, like FTP and telnet.

`hosts.allow` is a list of hosts allowed to connect and `host.deny` is a list of hosts denied access.

For example, to deny access to everyone but the addresses 10.0.0.2 through 10.0.0.10, `host.deny` would look like this:

```
ALL: ALL
host.allow would look like this:
ALL: 10.0.0.2/ 10.0.0.10
```

This would allow all services to 10.0.0.2 through 10.0.0.10 (`hosts.allow`) and deny all services to all other hosts (`host.deny`). Of course, you can allow some services and not others by replacing the `ALL` parameter with the service, such as `FTP`.

17.3 telnet and Remote Services for Linux

This service is also started in `/etc/inetd.conf`. To start them, add (or uncomment) the following lines:

```
telnet          stream          tcp          nowait root
/usr/sbin/tcpd  in.telnetd
ftp            stream          tcp          nowait root
/usr/sbin/tcpd  in.ftpd  -l  -a
```

Of course, you can use the `hosts.deny` and `hosts.allow` files to restrict the clients that access these services. Just like FTP, user access is controlled by the user accounts.

17.4 Secure Shell (SSH)

17.4.1 What is SSH?

Secure Shell (SSH) provides authentication and strong encryption to the remote commands in UNIX (and Linux). It consists of a server and a client. The server runs on the machine that you are connecting to and the client runs on the machine that is connecting to the server. There are free and commercial versions of SSH available for most UNIX systems and clients are available for all versions of Windows.

Authentication assures that the remote computer you are connected with is really who it says it is. A common way of breaking into computer systems is to pretend to be somebody else. Encryption assures that the information sent between two computers can't be read by a third party. This assures that information such as usernames and passwords can't be read while they are being broadcast over the network.

SSH provides a secure alternative to remote commands (`rsh`, `rlogin`, and `rexec`). It also provides a secure channel to use for other remote programs such as X11 and `vnc`.

There are two versions of SSH: SSH1 and SSH2. SSH1 is older and runs on almost every operating system. It is free for non-commercial use. SSH2 is newer and has better security. It is only free for educational and personal use. Go to the *Data Fellows Web site* (<http://www.datafellows.com>) or *Van Dyke Technologies Web site* (<http://www.vandyke.com>) for more information on licensing and support for SSH.

Since SSH1 and SSH2 are not fully compatible with each other, you need to decide which one to use. Although there is a free SSH2 client for Windows, it is not well-tested and debugged (meaning use at your own risk). For mixed Linux and Windows environments, either use SSH1 or get a commercial version of SSH2 from Data Fellows or Van Dyke.

17.4.2 What Does It Protect Against?

There are several methods of breaking into UNIX (and Linux) systems. Most of them are based on either pretending to be someone else or capturing authentication data on the network. SSH does not trust anything that comes through the network. An attacker on the network can only cause SSH to disconnect, not take over a session, or capture passwords. Here are some of the attacks that SSH protects against (from the SSH FAQ):

- **IP spoofing**— Where a remote host sends out packets which pretend to come from another, trusted host. SSH even protects against a spoofer on the local network, which can pretend it is your router to the outside.
- **IP source routing**— Where a host can pretend that an IP packet comes from another, trusted host.
- **DNS spoofing**— Where an attacker forges name server records.
- **Packet sniffing**— Interception of clear-text passwords and other data by intermediate hosts.

- **Man in the middle**— Manipulation of data by people in control of intermediate hosts.
- **X11 spoofing**— Attacks based on listening to X-Windows authentication data and spoofed connections to the X11 server.

17.4.3 What SSH Doesn't Protect Against

If an attacker gains root access to your machine, they can subvert SSH as well. Gaining access to your home directory will also compromise SSH security. This is why it is a bad idea to export the home directory with NFS.

17.4.4 SSH Client for Windows

There are several free and commercial SSH clients for Windows. Some free implementations are at:

- Robert O'Callahan's TTSSH, an SSH1 extension to the TeraTerm client—
<http://www.zip.com.au/~roca/ttssh.html>.
- Gordon Chaffee's command-line port of SSH1 and `scp1`—
<http://bmrc.berkeley.edu/people/chaffee/winntutil.html>.
- Sergey Okhapkin's SSH1 and SSH2 servers and clients ported to 32-bit Windows—<http://www.lexa.ru/sos/>.
- PuTTY, Simon Tatham's 32-bit Windows SSH1 client—
<http://www.chiark.greenend.org.uk/~sgtatham/putty.html>.
- FiSSH, Mass Confusion's 32-bit SSH1 client for Windows—
<http://www.massconfusion.com/ssh/>.
- Cynus Win32 port of SSH 1.2.2 by Raju Mathur—
<http://reality.sgi.com/raju/software.html>.

The setup varies with the different packages, so check the documentation that comes with them.

17.4.5 Installing SSH1 for Linux

Like most Linux programs, SSH1 is available in both source code and binaries (compiled programs). The advantages of using source code is that you get the latest updates, but the downside is that it is harder to install since you have to compile it yourself. The binaries are easier to install, but are not always the latest version.

The source code for SSH1 is available at <ftp://ftp.cs.hut.fi/pub/ssh/>. If you have trouble accessing this site, there are several mirrors listed at <http://www.onsight.com/ssh/faq/ssh-faq-2.html>.

Compiled binaries in the RPM format are available at <http://rpmfind.net/linux/RPM/ssh.html>. Just choose the latest version available for your system. For example, the latest RPM version of SSH1 available at this writing is 1.2.27-5, so the link to the latest version would be `ssh-1.2.27-5us.i386`.

17.4.6 Setting Up an SSH1 Server for Linux

The first thing to do is to disable the remote services such as `rexec`, `rsh`, and `rlogin`. This is done by adding a `#` in front of the respective lines in `inetd.conf`. This prevents the clients from logging in insecurely, which would defeat the whole purpose of SSH.

Once this is done, `sshd` (the server program) can be loaded. This can be loaded from the command line. There are several options that can be set in the configuration file, `/etc/sshd_config`. To get details on these settings, go to the man page (type `man sshd` from the prompt). Many of these can be set with command-line options. The command-line options override the settings in the `sshd_config` file.

17.4.7 Using the SSH Clients

SSH is the replacement for `rshell`, `rexec`, and `rlogin`. The general syntax is `ssh<remote computer name>`. You can also send a username (other than your current user-name) by typing `ssh -l <username> <remote computer name>`. To execute a remote command (like `rexec`), type `ssh <remote computer> <command name>`. There are many other options for SSH listed on the man page (type `man ssh`).

`scp` is the replacement for `rcp`. To run it, type `scp <user@host:file to be copied><user@host:where to copy the file>`. You can substitute the `user@host` for files on the local machine. Also, to keep the file attributes, add `-p` to the command line. For example, to copy the file `/root/test.file` from the local machine to the `/root/` directory on the machine `mmccune` using the root account, type `scp -p /root/test.file root@mmccune:/root/Hlt472001731t_Hlt472001731/test.file`. We can use this command to save the attributes of the file also.

The clients for SSH2 are identical to the SSH1 clients except that they have a "2" at the end. The equivalents of `sshd`, `ssh`, and `scp` are `sshd2`, `ssh2`, and `scp2`, respectively. The basic commands are the same, but some of the advanced commands are different. Unless you have both SSH1 and SSH2 installed, SSH2 will create a symbolic link between the SSH commands so that they will run the SSH2 equivalent. See the man pages for the details of the commands.

SSH2 also includes a secure version of FTP called `sftp2`. This allows a secure FTP session between the FTP server and FTP client. The basic syntax of `sftp2` is the same as that of regular FTP. For more advanced commands, see the man page.

17.4.8 Running Other Services over SSH

Backups and Updates. Backups that use `rsh` should work normally. Just change the backup script to use `ssh` instead of `rsh`. There is also a backup utility called `datbkr` that uses SSH at <http://www.psychosis.com/datbkr>.

`rdist` can also be made to run with SSH. `rdist` is a program that allows programs and data on UNIX machines to be updated through the network.

To get `rdist` to work with SSH, you must either recompile `rdist` to use SSH instead of `rsh` or simply use the `-p` option to change the shell used by `rdist`. If you are using password authentication with `rdist` version 6.12 to 6.15 or if you get the message "Warning: Denied agent forwarding because the other end has too old version," you need to upgrade your version of `rdist` (<http://www.magnicomp.com/rdist/>).

Firewalls. Firewalls are devices that control network traffic and isolate one network from another. The most common use of firewalls is to protect internal networks from the Internet.

SSH can send packets through any open port on a firewall by using the `-p` option. For example, most firewalls leave Port 443 open for the Secure Socket Layer (SSL):

```
On the server: sshd -p 443
```

```
On the client: ssh -p 443 <remote host name>
```

Socks 4 and 5 support should work in version 1.2.16 or later. Socks support in version 2.0.11 and later should work as well. Socks is a firewall implementation developed by NEC.

PPP. PPP can run over SSH, but there are problems with TCP packet forwarding. It is also a good idea to enable compression, since this improves the speed of the connection. There is a sample script for running PPP at <http://www.inka.de/~bigred/sw/ssh-ppp-new.txt>. There is also a compiled version of the script at <http://detached.net/vpnstarter/> that is more reliable and easier to set up.

There is a kernel driver that uses UDP that doesn't have the problems that TCP forwarding has at <http://sites.inka.de/sites/bigred/devel/cipe.html>.

UDP Services such as NIS, NFS, and DNS. A general implementation of UDP services can be found at ftp://ftp.tu-chemnitz.de/pub/Local/informatik/sec_rpc/, and even though it has only been implemented for NIS, it should work with NFS and DNS as well.

FTP and POP Services. SSH2 has secure FTP with it. For SSH1, there are two ways to use SSH on FTP. The simplest way is to download a secure version of FTP from <http://www.docs.uu.se/~pem/hacks/>. Another way to do it is to redirect FTP through an SSH connection. To do this, redirect the port FTP uses (Port 21) to an unused local port. For example, to connect to the FTP server mmccune using the local port, 1234:

```
ssh -g -L 1234:mmccune:21 mmccune
```

Then open up an FTP session using the local port, 1234:

```
ftp localhost 1234
```

Please note that the standard FTP servers accept connections through an SSH port, but some FTP servers such as WU-FTP may not. If your FTP server doesn't work, check the documentation to see if it supports passive mode.

POP services are easily redirected through an SSH connection. You can use the script `gwpop` at <ftp://ftp.internatif.org/pub/unix/gwpop/>.

TCP Wrappers. TCP Wrappers simply allow a computer to allow or deny access to certain services based on the connecting machine's IP address. To use this with SSH, SSH must be compiled with TCP Wrapper support. First, configure it with these options:

```
# ./configure --with-lib-wrap=/PATHOF/libwrap.a
```

Then, add the following lines to the SSH Makefile:

```
-I/PATHOF/tcpwrappers  
WRAPLIBS = -L/PATHOF/tcpwrappers -lwrap
```

X11. Like TCP Wrappers, X11 support must be compiled into SSH. The only thing necessary is to configure SSH for X-Windows:

```
./configure --with-x
```

Then compile it normally. After that is done, change the configuration file for the SSH server and client. Add this line to the client, whose default location is `/etc/ssh_config`.

```
ForwardX11 yes
```

Add the following to the server's configuration file located at `/etc/sshd_config`:

```
ForwardX11 yes
```

For SSH2, you only need to change the server's configuration by adding the following line to `/etc/ssh2/sshd2_config`:

```
ForwardX11 yes
```

Keep in mind that if you are using TCP Wrappers, you must add the line `sshd2fdw-X11:` to the `host.allow` file.

SGI GL, Digital Certificates, and PGP Keys. PGP keys are supported in SSH2, but SGI GL and digital certificates are not supported at this time. OpenGL, which runs under X11, is supported.

17.4.9 Administering SSH

The biggest task of administering SSH is managing the keys. These are used to authenticate the clients. There is a script packaged with SSH that is used to collect these keys from the clients. It is called `make-ssh-known-hosts.pl`. This script should be run from `cron` every day to update the keys. `cron` is used to run regularly scheduled tasks in UNIX (and Linux). See your user's manual on how to use `cron`.

This script generates a log file that lists things like clients that change their keys and any other errors or changes. There is a script available at <http://www.uni-karlsruhe.de/~iq25/ssh-faq/comp-host-list> that will process the log file and write it in a more readable form into another file. This is good for finding "man in the middle"-type attacks.

There are also patches available for Kerberos that allow it to work with SSH. Kerberos allows secure authentication across an insecure network. These patches allow SSH to update its encryption keys when Kerberos authentication is used. There are patches available for Kerberos 4 at <http://www.monkey.org/~dugsong/ssh-afs> and Kerberos 5 at http://www.ncsa.uiuc.edu/General/CC/ssh/patch_repository/descriptions/AFS_KRB5.html

17.4.10 Troubleshooting SSH

The best way to prevent problems with SSH is to keep it up-to-date. Patches to SSH are available at <http://www.ssh.org/>. After downloading the patches, go to the SSH source directory and type `patch -p1 -y1 < /path/to/sshpath` (the location of `ssh` and `sshd`). Then run `make distclean`, `configure`, `make`, and `make install` to compile SSH again.

The GMP assembler routines don't work on some systems, so compiling may fail with some error messages from the assembler, so SSH can still be compiled without these routines. Just `configure` SSH with the `-disable-asm` option as follows:

```
make distclean
configure --disable-asm
make
make install
```

When the `configure` process cannot find `xauth`, check to be sure that `xauth` is in your path by typing `whereis xauth`. If this doesn't show the path of `xauth`, add it to the path by typing `export PATH=$PATH: <path of xauth>`. If you don't know where `xauth` is, you can find it by typing `find xauth`. After this is done, run `configure` again.

If compilation aborts with some error message about `libc.so.4`, your system is not configured properly. Run `cd /usr/lib`, then `ln -s libc.sa libg.sa`, then recompile.

If you experience problems with multi-homed hosts (multi-homed simply means they have more than one IP address), which usually occurs on devices that have more than one network card such as bridges and gateways, the client may not have all the IP addresses used by the host. Having a properly configured network will solve most of these problems. In particular, make sure the DNS server has all the IP addresses for the host and that the client is using the DNS server. Also, do an `nslookup` to see if all the host IP addresses appear.

There is also the following Perl script, which loads the SSH client and helps prevent problems with multi-homed hosts. I can't say if it works, but it may be worth a try:

```
Peter Polkinghorne, Computer Centre, Brunel University,
Uxbridge, UB8 3PH, |
Peter.Polkinghorne@brunel.ac.uk +44 1895 274000 x2561
UK |
-----
-----
-----
---
# just run perl on stdin
/usr/local/bin/perl5 /LOCAL/etc/sshrc.pl
-----
---
#
# perl script to xauth for all interfaces - derived from
Paul Pomes script.
#
# check there is a cookie to add:
exit unless $proto_cookie = <>;
# check we need the X11 forwarding
exit unless ($dpy = $ENV{'DISPLAY'});
# extract host/dpynum elements
$dpy =~ /^(.*):(.*)$/ || die "Bad DISPLAY $dpy";
$host = $1;
$dpynum = $2;
# find out the interfaces
open(INTERFACES, "/usr/sbin/ifconfig -a|") || die "Cannot
execute ifconfig";
$xauth = "/usr/openwin/bin/xauth"; open(XAUTH, "|$xauth -q -
")
die "Cannot execute $xauth";
while (<INTERFACES>) {
if (/inet\s+([0-9.]+)/) {
$addr = $1;
next if $addr eq "127.0.0.1";
print XAUTH "add ${addr}:$dpynum $proto_cookie";
```

```
}
}
# add unix domain (NB only once - to match default sshd
behaviour)
print XAUTH "add ${host}/unix:$dpynum $proto_cookie";

close(XAUTH);

close(INTERFACES);
```

If SSH asks you for passwords despite `.rhost`, you may have one of the following problems:

- The remote host's `/etc/ssh_known_hosts` is not world-readable. Fix this with the command `chmod 0+r /etc/ssh_known_hosts`.
- The client host key is not stored in the `known_hosts` file. Note that this has to be the canonical (usually, the fully-qualified) domain name such as `myhost.mydomain.com`
- The client host does not have a reverse mapping in the name server. SSH requires both a reverse mapping and a forward mapping that contains the original IP address. Reverse mapping uses a host name to look up an IP address. Forward mapping uses the host name to look up an IP address.
- A multi-homed client or host does not have all of its IP addresses listed in the DNS entry. Note that versions prior to 1.2.12 have bugs in handling multi-homed hosts. The fix for this is listed in the previous section.
- The user's home directory or `~/.rhosts` is world- or group-writable (see `StrictModes` server configuration option). Setting `StrictModes` to `yes` in the SSH server configuration will check the ownership of the home directories before login. This will stop others from modifying the `.rhost` even if it is left world- or group-writable. On some machines, if the home directory is on an NFS volume, `~/.rhosts` and your home directory may need to be world-readable.

The root account has to use `~/.rhosts`, or `~/.shosts`; `/etc/shosts.equiv` and `/etc/hosts.equiv` are disregarded for root.

- There is confusion between `RhostsRSAAuthentication` and `RSAAuthentication`. Make sure the server and client use the same authentication method. `RSAAuthentication` is the default because it is more secure.

If you have SSH looping with "Secure Connection Refused," SSH will fall back on the insecure `r` command if SSH server is not running on the host machine. This message occurs if SSH can't find the old `r` command. This could be caused by the following:

- You named the SSH client `rsh` without pointing it to the regular `rsh` program—If you did this, you need to compile SSH with the `--with-rsh=PATH` configure option. For example, if you moved the insecure `rsh` to `/etc/old`, compile SSH with this option: `./configure --with-rsh=/etc/old/rsh`.

SSH2 doesn't use `rsh` if it can't open a secure connection. It just refuses to connect.

stone: <http://hp.vector.co.jp/authors/VA000770/stone/>
OpenSSH: <http://www.openssh.com>

Appendix A. Disk Error Codes

0x00	"Internal error"—This code is generated by the sector read routine of the <code>LILLO</code> boot loader whenever an internal inconsistency is detected. This might be caused by corrupt files. Try rebuilding the map file. Another possible cause for this error is an attempt to access cylinders beyond 1024 while using the <code>LINEAR</code> option. See the section titled "BIOS Restrictions" for more details on how to solve the problem.
0x01	"Illegal command"—This shouldn't happen, but if it does, it may indicate an attempt to access a disk which is not supported by the BIOS.
0x02	"Address mark not found"—This usually indicates a media problem. Try again several times.
0x03	"Write-protected disk"—This should only occur on write operations.
0x04	"Sector not found"—This typically indicates a geometry mismatch. If you're booting a raw-written disk image, verify whether it was created for disks with the same geometry as the one you're using. If you're booting from a SCSI disk or a large IDE disk, you should check whether <code>LILLO</code> has obtained correct geometry data from the kernel or if the geometry definition corresponds to the real disk geometry. (See the section titled "Disk Geometry.") Removing <code>COMPACT</code> may help to; so may adding <code>LINEAR</code> .
0x06	"Change line active"—This should be a transient error. Try booting a second time.
0x07	"Invalid initialization"—The BIOS failed to properly initialize the disk controller. You should control the BIOS setup parameters. A warm boot might help too.
0x08	"DMA overrun"—This shouldn't happen. Try booting again.
0x09	"DMA attempt across 64k boundary"—This shouldn't happen. Try omitting the <code>COMPACT</code> option.
0x0C	"Invalid media"—This shouldn't happen and might be caused by a media error. Try booting again.
0x10	"CRC error"—A media error has been detected. Try booting several times, running the map installer a second time (to put the map file at some other physical location or to write "good data" over the bad spot), mapping out the bad sectors/tracks, and

	if all else fails, replacing the media.
0x11	"ECC correction successful"—A read error occurred, but was corrected. LILO does not recognize this condition and aborts the load process anyway. A second load attempt should succeed.
0x20	"Controller error"—This shouldn't happen.
0x40	"Seek failure."—This might be a media problem. Try booting again.
0x80	"Disk timeout"—The disk or drive isn't ready. Either the media is bad or the disk isn't spinning. If you're booting from a floppy, you might not have the drive door closed. Otherwise, trying to boot again might help.
0xBB	"BIOS error"—This shouldn't happen. Try booting again. If the problem persists, remove the <code>COMPACT</code> option or add/remove <code>LINEAR</code> , which might help.

A large amount of documentation is distributed with `LILLO`. The usual location of the documents is `/usr/doc/lilo-x`, where `x` is the version number of `LILLO`.

Appendix B. Samba Documentation

Appendix Objectives

- GNU License
- The Samba FAQ
- Just What is SMB?

This appendix contains some documentation that you might find useful in learning Samba.

The GNU License

Software that is in the public domain can follow several license formats. Some allow you to resell the software, but must pay a royalty. Others allow you only to sell and make money on the costs of the duplication. The GNU license is shown here and is something that you will commonly encounter in the Linux and UNIX World.

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

675 Mass Ave, Cambridge, MA 02139, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any

patent must be licensed for everyone's free use or not licensed at all. The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

1. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
2. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
3. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on The Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

1. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
2. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
3. c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to his License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix: How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of
what it does.>
Copyright (C) 19yy <name of author>
```

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details. You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright c 19yy name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type "show w".
```

```
This is free software, and you are welcome to redistribute it under certain conditions; type
"show c" for details.
```

The hypothetical commands "show w" and "show c" should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than "show w" and "show c"; they could even be mouseclicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program
"Gnomovision" (which makes passes at compilers) written by James Hacker.
<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

The Samba FAQ

Author's Comments: The following FAQ is copied from the online documentation for SAMBA at <http://www.samba.com>. It is copied under the GNU License listed above.

Samba FAQ

This FAQ is automatically generated from the Samba bug tracking system. As such it contains answers that we frequently send to users who report problems to samba-bugs@samba.org. Please report inaccuracies or out of date information so it can be fixed.

Index

CRLF-LF Conversions
Closed Off 1.9.18
Couldn't open status file STATUS..LCK
Domain Controller
Get NTDOM Code
IP Address Change
Linux & mmap()
Logon errors in NT Event Viewer
Macintosh Clients
NT Guest Access
NT SP3 and Encryption
NTDOMAIN code
NetWkstaUserLogon
Not listening for calling name
OpLock Break Errors
PLEASE Start by Reading Docs!
PWL Files
Password Cracking
Pizza Vouchers
SMBFS not Part of Samba
SWAT on Red Hat Linux
Samba 2.x and PAM (especially FreeBSD)
System Error 1240
This is not a helpdesk
Time off by 1 hour
Trapdoor UID
Unix Permissions control Access
User Access Control
Using NT to Browse Samba Shares
Win95 or 98 and Encryption
Win9X in User Level Access mode
Windows98 Passwords
XXX isn't in user level security mode
Y2K
Case sensitive
Comp.protocols.smb
Dont descend & security
File caching
Generic icons displayed
Linux 2.0.x and smbmount
Linux compile problem
Setting times when not owner
Setup.exe and 16 bit programs
Smbclient -N
Smbfs for other Unixes
Smbpasswd: rejected session request
Smbsh and glibc-2.1
Smbtar blocksize
Unsubscribe
Win98 slowdown
Your server software is being unfriendly

CRLF-LF Conversions

We get many requests for CRLF/LF format conversion handling by samba. The problem is that there is no clean way to determine which files should / could be converted and which MUST not be.

Since Unix and DOS/Windows uses alike will use .txt to represent a file containing ASCII text we can not reliably use the file extension. The same applies to the .doc extension. Samba operates around the premise that we should leave all files unchanged. By not implementing CRLF/LF conversions we can not be guilty of damaging anyone's files. When someone comes along with a sound implementation that guarantees file integrity we will jump at the opportunity to implement this feature. Until such time there is no prospect for action on this topic.

Closed Off 1.9.18

Thank you for reporting your difficulties with samba-1.9.18 series code. We regret to advise however that all work on the 1.9.18 code tree has been closed off with the release of samba-1.9.18p10.

All development efforts are now being focussed on stabilizing samba-2.0.0 so it can go into release to stable code as soon as possible.

If you could download the most recent beta release and report back any difficulties you may have we will do our best to close them out before the stable release.

We will still fix major bugs and security holes in 1.9.18p10, but minor bugs will be fixed if the problem still exists in the 2.0 tree.

For information about accessing the latest CVS source code please refer to <http://samba.org/cvs.html>

Couldn't open status file STATUS..LCK

If you run smbstatus before anyone has ever connected to your new Samba installation then you may get the error:

```
Couldn't open status file /var/lock/samba/STATUS..LCK
```

or possibly the error:

```
ERROR smb_shm_open : open failed with code No such file or
directory
ERROR: Failed to initialize share modes!
Can't initialize shared memory - exiting
```

both of these errors are harmless. The appropriate files and memory segments get automatically created the first time you connect to Samba. Try connecting with smbclient and you should find that smbstatus is happy after your first connection.

Domain Controller

```
> Unknown parameter encountered: "domain controller"
> Ignoring unknown parameter "domain controller"
```

As of 1.9.18 the "domain controller" parameter has changed. You should not need it, but in it's place may need "networkstation user logon = yes". Please check the smb.conf man page BEFORE using this option so you understand it's significance.

Get NTDOM Code

The domain controller code is now integrated into the main Samba source code.

Please see <http://samba.org/cvs.html> for information on how to download the latest version. You may also wish to join the samba-ntdom mailing list. See <http://lists.samba.org/> for details.

There is also a separate Samba NTDOM FAQ available in the documentation section of the samba web site at <http://samba.org/samba/>

IP Address Change

The following is an example of a problem we see from time to time:

"Samba was working fine. We had to change the IP address of the Samba server. Following the change Samba does not work. We have tried EVERYTHING—it still does not work!"

What to check:

1. Follow all instructions in DIAGNOSIS.txt from the samba docs directory.
2. Locate you browse.dat and wins.dat files. They may be found in the following typical locations:
- 3.
4. `/usr/local/samba/var/locks`
5. `/var/locks/samba`
6. `/opt/samba/var/locks`
- 7.

If you can not locate where samba stores these files you can always run:

```
testparm | grep lock
```

8. Shut down samba.
9. Delete the browse.dat and wins.dat files
10. Restart samba.
11. Check that any files you deleted have been recreated.
12. Now follow DIAGNOSIS.txt again.

Cause:

Samba will place into these files entries for itself with your old IP address. When you restart Samba it preloads it's name cache with this information and expects to be able to resolve it's own address to the same address as it has just read from these files.

Deleting the files means samba takes a little longer to stabilize on startup but otherwise will now operate correctly.

In Samba 2.0 this problem has been fixed properly by storing signature information in the relevant files.

Linux & mmap()

Early versions of Linux did not support shared writeable mmap(). I believe it was introduced in kernel version 2.0.

There are 3 possible fixes:

1. upgrade to a newer version of the Linux kernel
2. don't compile Samba with FAST_SHARE_MODES defined
3. use Samba 1.9.18 which provides FAST_SHARE_MODES via SysV IPC shared memory. I believe Linux 1.2 supported this.

In Samba 2.0 configure script will auto-detect the OS capabilities and will enable shared memory only if available.

Logon errors in NT Event Viewer

The logon errors in the NT event viewer are caused by Samba trying to detect broken NT password servers.

Some NT servers will accept any username/password for session setup requests and always validate it, returning a positive session setup response without the guest bit set. Samba checks for this by deliberately sending an incorrect password when calling the password server in server level security. If the incorrect password succeeds then Samba logs an error and refuses to use the password server.

You can remove this check from the code if you want, but as we have not yet worked out what causes a NT server to show this behavior there is a risk that your NT server will start behaving incorrectly and thus make your Samba server insecure.

Future versions Samba will have a new security option "security = domain" which will use the same protocols that NT uses for domain authentication (currently Samba uses the method that MS documents, rather than that which Microsoft actually use). Once that in place this problem should be solved.

Macintosh Clients

> Are there any Macintosh clients for Samba?

Yes. Thursby now have a CIFS Client/Server called DAVE—see <http://www.thursby.com/> They test it against Windows 95, Windows NT and samba for compatibility issues. At the time of writing, DAVE was at version 1.0.1. The 1.0.0 to 1.0.1 update is available as a free download from the Thursby web site (the speed of finder copies has been greatly enhanced, and there are bug-fixes included).

Alternatives—There are two free implementations of AppleTalk for several kinds of UNIX machines, and several more commercial ones. These products allow you to run file services and print services natively to Macintosh users, with no additional support required on the Macintosh. The two free implementations are Netatalk, <http://www.umich.edu/~rsug/netatalk/>, and CAP, <http://www.cs.mu.oz.au/appletalk/atalk.html>. What Samba offers MS Windows users, these packages offer to Macs. For more info on these packages, Samba, and Linux (and other UNIX-based systems) see http://www.eats.com/linux_mac_win.html

NT Guest Access

What you are seeing is normal and deliberate.

MS Windows NT can be configured with the guest account enabled. When this is the case no logon attempt will ever fail. Instead NT will allow the user access as the guest account IF the username and/or password are incorrect. In a situation where Samba is using and NT system to validate user passwords, if the NT server guest account is enabled then a user logging on as "root" will always be validated even if the password was incorrect. There is NO way that samba can tell from the reply packet from NT whether the password was correct and normal user privilege has been granted, or whether the password was incorrect and the user has been given only "guest" privileges. In short, if we were NOT to do what we do, then there would be no way of telling whether or not the password server allows guest only logons. Were we to just accept the validation response from such a server the a user could easily gain "root" level access to a Samba server.

Now you would not really want us to change the current behavior, would you?

```
> Hi folks ... I don't know if you have seen this, have
corrected
this yet
> or it is my configuration.
> I am using our company PDC for passwd
authentication and it works OK
> except for one snag.
> The authentication process between the our Samba server &
the PDC always
> includes one unsuccessful pass thru attempt.
> This initial pass thru validation has an
incorrect user password
> (1F1F1F1F.....). A SMB reject from the PDC forces the
Samba Svr to
> immediately send a second validation with the correct
> encrypted Bell Master Domain user password.
> It would be nice to get rid of the first bad validation
attempt.
```

NT SP3 and Encryption

Microsoft changed WinNT in service pack 3 to refuse to connect to servers that do not support SMB password encryption.

There are two main solutions:

1. enable SMB password encryption in Samba. See ENCRYPTION.txt in the Samba docs (this is best done with samba versions more recent than 1.9.18)
2. disable this new behavior in NT. See WinNT.txt in the Samba docs

Note that Samba-1.9.18 and later support encrypted passwords without need to recompile and link with the libdes (DES) library. Refer to the man page for smb.conf and to ENCRYPTION.txt for information about use of encrypted passwords.

Please refer to the following URL for more information on this subject:

<http://support.microsoft.com/support/kb/articles/q166/7/30.asp>

NTDOMAIN code

If you are trying the NTDOMAIN version of Samba then please join the samba-technical list and listen there for a while. I also suggest you read the list archives. See <http://lists.samba.org/> for more info.

Samba support for NT domain control is still very experimental. Only try it if you are a programmer willing to experiment.

All bug reports regarding this experimental code should be directed ONLY at the samba-technical or samba-ntdom mailing lists.

In response to concerns over profile handling: Profile handling will be looked at seriously once we get the domain code to stabilize. Until then, what we document in Samba and what works can be in conflict. We stress again the highly experimental nature of all the NTDOMAIN code.

In response to concerns over compilation problems: Code updates to the head branch code tree are to samba-2.1.0 NOT samba-2.0.0. The Samba-2.1.0 code may not compile and may not work at any time. Please use this at your own risk.

Samba-2.0.0 is being readied for release within a few days. It is the release candidate stable code but does not have fully functional PDC support. This is precursor code to the samba-2.1.0 branch.

NetWkstaUserLogon

The password server behavior changed because we discovered that bugs in some NT servers allowed anyone to login with no password if they chose an account name that did not exist on the password server. The NT password server was saying "yes, it's OK to login" even when the account didn't exist at all! Adding the NetWkstaUserLogon call fixed the problem, and follows the "recommended" method that MS have recently documented for pass through authentication.

The problem now is that some NT servers (in particular NT workstation?) don't support the NetWkstaUserLogon call. The call also doesn't work for accounts in trust relationships.

The eventual solution for this will be to replace the password server code in Samba with NT domain code as that is developed. For now you have the choice of compiling Samba either with or without the NetWkstaUserLogon call in the password server code.

In 1.9.18p3 and later you can disable the NetWkstaUserLogon call with an option in your smb.conf using the "networkstation user login" option.

Not listening for calling name

```
> Session request failed (131,129) with myname=HOBBS dest-
name=CALVIN
> Not listening for calling name
```

If you get this when talking to a Samba box then it means that your global "hosts allow" or "hosts deny" settings are causing the Samba server to refuse the connection. Look carefully at your "hosts allow" and "hosts deny" lines in the global section of smb.conf.

It can also be a problem with reverse DNS lookups not functioning correctly, leading to the remote host identity not being able to be confirmed, but that is less likely.

OpLock Break Errors

```
> I'm receiving the same error from the following versions
of samba:
>
> smbd version 1.9.18p10 started
> smbd version 2.0.0beta2 started.
>
> The following error message appears multiple times in the
log files with
> either version, and eventually locks up the entire samba
server:
>
> [1998/12/11 22:29:07, 0]
smbd/oplock.c:request_oplock_break(909)
> request_oplock_break: no response received to oplock
break request to pid
> 19883 on port 2328 for dev = 810000a, inode = 332533
>
```

```
> I'm running this on Digital UNIX V4.0B on an AlphaStation
255. We've also
> seen this same error message on DU V3.2C running various
flavors of samba.
>
```

The advisory message means that you have either a defective network card on one of your clients, or else an MS Windows application is refusing to respond to an oplock break request from another MS Windows client that wishes to access an already locked file.

PLEASE Start by Reading Docs!

We are glad you are interested in Samba, but please read the documentation!

English: <http://samba.anu.edu.au/samba>

German: <http://samba.sernet.de>

Japanese: <http://samba.bento.ad.jp>

French: <http://www.bde.espci.fr/homepage/Patrick.Mevzek/samba>

From the README file:

DOCUMENTATION

There is quite a bit of documentation included with the package, including man pages, and lots of .txt files with hints and useful info. This is also available from the web pages. There is a growing collection of information under docs/faq; by the next release expect this to be the default starting point.

FTP SITE

Please use a mirror site! The list of mirrors is in docs/MIRRORS.txt. The master ftp site is samba.anu.edu.au in the directory pub/samba.

MAILING LIST

There is a mailing list for discussion of Samba. To subscribe send mail to listproc@samba.anu.edu.au with a body of "subscribe samba Your Name" Please do NOT send this request to the list alias instead.

To send mail to everyone on the list mail to samba@listproc.anu.edu.au

There is also an announcement mailing list where new versions are announced. To subscribe send mail to listproc@samba.anu.edu.au with a body of "subscribe samba-announce Your Name". All announcements also go to the samba list.

NEWS GROUP

You might also like to look at the usenet news group comp.protocols.smb as it often contains lots of useful info and is frequented by lots of Samba users. The newsgroup was initially setup by people on the Samba mailing list. It is not, however, exclusive to Samba, it is a forum for discussing the SMB protocol (which Samba implements). The samba list is gatewayed to this newsgroup.

PWL Files

If you are worried about the security of PWL files then I suggest you look at <http://samba.org/pub/samba/docs/security.html>

Password Cracking

```
> Could you please send me Frank Stevensons program for
cracking .pwl
> files.
>
> If you have any other programs for cracking windows 95 pwl
files, could
> you send them to me or tell me where I can find them.
>
```

No. These are not part of samba—we will provide only samba components. For obvious reasons we can not offer any other software like password cracking tools. These are available from sites like BugTraq.

Pizza Vouchers

FAQ Answer about pizza vouchers:

The note about pizza vouchers in the Samba documentation started out as a bit of a joke. Since then I've been amazed at the number of people who have managed to send a pizza voucher in one way or another! I've had many happy evenings eating pizza thanks to these generous folks!

Some people also write to us wondering how to get a pizza voucher to Australia. Here are the techniques that have worked for others:

1. a few people have successfully talked their local Pizza Hut shop into sending vouchers to Australia. Apparently the staff look at you a bit strangely at first but at least 2 people have succeeded!
2. Others have rung up a local pizza outlet in Canberra (where I live) and have offered their credit card numbers over the phone. They then emailed me telling me which shop to ring up to order pizza. This has worked with two different shops here in Canberra—Pizza Hut and Dominos Pizza.
3. I've received several pizza vouchers that are valid in various places around the world and have started a small collection. Some day I hope to eat my way through them when I visit some of these places.
4. I've received several .gif files of pizzas and even some nice pictures and cardboard pizzas!

Please remember that the "pizza voucher for Samba" thing started as a joke. It's a lot of fun but you certainly are not required to send anything. I won't starve without them and maybe I'll put on a bit less weight :-)

Also remember all the other people who help with Samba. If someone helps you particularly then maybe just send them a thank you email? This sort of thing really is appreciated!

Anyway, if you still want to send a pizza my address is:

3 Ballow Crescent

Macgregor A.C.T

2615 Australia

Phone: +61 2 6254 8209

and the nearest pizza outlet is "Kippax Pizza Hut"
Thanks!

SMBFS not Part of Samba

We regret to advise that including `smbmount` and `smbumount` in the Samba tarball has been a mistake. These programs are part of the `smbfs` package and are NOT maintained by the Samba-Team. We are currently contemplating the removal of these programs from the samba tarball since they are a constant cause of complaint and we do NOT have spare capacity to handle additional load.

SMBFS is available only for Linux. In Samba 2.0 we are introducing a tool called `smbsh` that offers similar functionality but is portable to a large number of Unixes.

SWAT on Red Hat Linux

```
> I hope this is not something stupid I overlooked somehow.  
I  
browsed the  
> mailingarchive and I noticed I wasn't the only one with  
this problem.  
> This is what happend:
```

Now please excuse my revised version of what you did: You followed the Sinatra method of software installation—the "I did it my way!" And you found it does not work because something was over looked. Maybe?

Here is how my Sinatra method works:

1. Unpack the samba tarball
2. rename the directory "samba-2.0.0beta5" to "samba-2.0.0".
3. `cd samba-2.0.0/packaging/redHat`
4. `sh makerpms.sh`
5. `cd /usr/src/redhat/RPMS/i386`
6. `rpm -Uvh samba-2.0.0-1.i386.rpm`

NOTE: This installation does a "kill-HUP" on the PID for `inetd`, so SWAT is available.

1. launch a web browser from a client (any client)
2. Go to the URL http://samba_host:901
3. User = root, enter root's password.

Now you are ready to configure samba.

```
>  
> First I used the Redhat section of the install procedure  
which uses an rpm  
> statement to buils a rpm package. I installed the package  
and tried to use  
> SWAT, the grafical web interface.
```

If you built the Red Hat RPMs, then all you had to do was install the binary RPM and all should work. If it does not work then you most likely have a TCP/IP configuration problem, not a samba problem.

```
> This failed because it assumes the
```

```
> samba directory tree in /usr/local/samba and the redhat
rpm
file installs stuff
> in /usr/share or similar.
```

Please explain this statement. Nowhere in the samba2.spec file do we show any dependency on /usr/local/samba, and if you built the packages using the shell script provided, then all paths will have been changed to reflect the standard Linux path layouts.

```
> Eventually, after making a some
> simlinks I thought it would be better to compile and
install
the
> distrubution-independed way (just make and make install
that is)
> Fine, now SWAT seems to work. However I have to put it in
demo mode because
> When it asks me to login I can't use my root login with
root-password
> anymore. It's not valid anymore to SWAT.
```

But the Samba installer is meant for people who REALLY know what they are doing and therefore does not attempt to install SWAT. Oh, yes, it installs the binary files but it does not modify /etc/inetd.conf and /etc/services.

```
>
> The question is: what is being validaded here: my Linux
passwd word for root
> or is somthing involved here with the samba passwords? Or
is this authentication
> handled by SWAT itself?
```

Red Hat Linux uses PAM (Pluggable Authentication Modules) and the authentication method is specified in the /etc/pam.d files. SWAT uses the /etc/pam.d/{samba,login} control files.

```
>
>
> I have RedHat 5.2 which I believe uses the egcs compiler
(not sure though..) I
> run it on an i386 pentium 166 MMX, so this is pretty
standard
I think. I used
> beta 4 of the Samba release
>
> Hope this helps...
```

Samba 2.x and PAM (especially FreeBSD)

Samba 2.x detects whether your OS has PAM (Pluggable Authentication Module) support at compile time and uses it if it is available. This leads to a problem on systems

that have PAM support in the libraries but where PAM is not configured. These "sleeping" PAM implementations cause all unix password authentication attempts to fail. We have fixed this for the next release of Samba (version 2.0.4) by adding a—`withpam` configure option. If you don't use that option then PAM won't be used.

System Error 1240

System error 1240 means that the client is refusing to talk to a non-encrypting server. Microsoft changed WinNT in service pack 3 to refuse to connect to servers that do not support SMB password encryption.

There are two main solutions:

1. enable SMB password encryption in Samba. See `ENCRYPTION.txt` in the Samba docs
2. disable this new behaviour in NT. See `WinNT.txt` in the Samba docs

This is not a helpdesk

The address samba-bugs@samba.org is meant for reporting bugs or for obtaining help where you suspect that a Samba bug is involved. It is not meant as a general helpdesk service. It's not that we don't want to help (we do!) but we get thousands of emails and have only a few people to deal with them. Trying to help with each Samba install personally is way beyond our abilities.

Instead I suggest that you try one of the following resources:

1. The Samba web site at <http://samba.org/samba/>
2. The mailing list samba@samba.org.

Read <http://lists.samba.org/> for more info on that.

1. The newsgroup `comp.protocols.smb`
2. a newsgroup specific to your OS, such as `comp.os.linux.*`

If you can afford it you could also contact one of the many companies that provide commercial Samba support. See <http://samba.org/samba/> and follow the links.

Time off by 1 hour

It sounds like either your PCs or your server don't have their timezones set up correctly for daylight savings time or just disagree on the time zone you are in.

Without knowing what sort of system you are running Samba on it is difficult to know what to change. If you want to use a "quick fix" then maybe the "time offset" command in `smb.conf` will be useful.

Trapdoor UID

```
> Log message "you appear to have a trapdoor uid system"
```

This can have several causes. It might be because you are using a uid or gid of 65535 or -1. This is a VERY bad idea, and is a big security hole. Check carefully in your `/etc/passwd` file and make sure that no user has uid 65535 or -1. Especially check the "nobody" user, as many broken systems are shipped with nobody setup with a uid of 65535.

It might also mean that your OS has a trapdoor uid/gid system

This means that once a process changes effective uid from root to another user it can't go back to root. Unfortunately Samba relies on being able to change effective uid from root to non-root and back again to implement its security policy. If your OS has a trapdoor uid system this won't work, and several things in Samba may break. Less things will break if you use user or server level security instead of the default share level security, but you may still strike problems.

The problems don't give rise to any security holes, so don't panic, but it does mean some of Samba's capabilities will be unavailable. In particular you will not be able to connect to the Samba server as two different uids at once. This may happen if you try to print as a guest" while accessing a share as a normal user. It may also affect your ability to list the available shares as this is normally done as the guest user.

Complain to your OS vendor and ask them to fix their system.

Note: the reason why 65535 is a VERY bad choice of uid and gid is that it casts to -1 as a uid, and the setreuid() system call ignores (with no error) uid changes to -1. This means any daemon attempting to run as uid 65535 will actually run as root. This is not good!

Unix Permissions control Access

Typical question:

```
> So here's the problem (and I apologise if this is just a
configuration
> issue). On any of our NT Workstation
clients (with the reg
edit to allow plain
> passwords), once a drive is mapped using userid and
password
with \\host\anyuser,
> that user can subsequently issue \\host\root and mount the
unix box at the root
> directory without any authentication required .. they have
access to the whole
> machine.
```

Samba does NOT second guess the security policy you wish to impose upon your site. You, as a system administrator for your Unix system, have the responsibility to determine by means of setting correct user, group and other permissions who can access what on your Unix system.

If you really want to restrict users so they can not connect to any other users' home directory then you will need to set the Unix permissions on your users home directory to drwx——(Octal 0x700).

User Access Control

```
> In windows when i set up a share in "user mode" i get the
message:
> "You cannot view the list of users at this time. Please
try again later."
>
> I know you have lists of users for access and aliasing
purposes,
but i
> have read nothing to support the idea that these lists
control
```

```
the Domain
> Users List...
```

Samba does NOT at this time support user mode access control for Window 9x although we hope to support it in an upcoming release.

Using NT to Browse Samba Shares

```
> WIN-NT workstations (nt4.0, service pack 3)
> samba with
> security = user
> encrypt passwords = yes
> guest account = guest
>
> start the explorer on a win-nt workstation and select
network.
I find
> my unix server running samba, but I can not see the list
of
shares
> unless I am a user, who is known in the smbpasswd of the
unix machine.
> The guest account "guest" exists on my
unix machine. For
testing I even
> made him a regular user with a password.
>
> With my network monitor I can see, that the win-nt
workstation
uses the
> current login, to connect to IPC$ on the samba server
> (for example "administrator"), not the guest account.
```

This is exactly how Windows NT works. You MUST have a valid account on the Windows NT box you are trying to see the resource list on. If your currently logged in account details do NOT match an account on the NT machine you are trying to access then you will be presented with a logon box for that machine. When you enter the name of an account on that machine / domain, together with a valid password then the resource list is made available. If the account details are not correct then no resource list is shown.

Samba follows the behaviour of Windows NT exactly.

Samba can be compiled with the GUEST_SESSION_SETUP option at 0, 1 or 2. The default is 0. If this is set to 1 or 2 then Windows NT machines that DO NOT have an account on the Samba server will see the resource list. Unfortunately Windows client bugs mean that using this option will probably cause more problems than it will solve. We do not suggest that you use it.

Win95 or 98 and Encryption

FAQ answer about Win95 (with TCP/IP update or OSR2 version) and Win98 and Samba: Microsoft changed Win95 upon release of their OSR2 version to refuse to connect to servers that do not support SMB password encryption. This change was also released in a TCP/IP update for Win95 and is included in all versions of Win98 that we've seen so far.

There are two main solutions:

1. enable SMB password encryption in Samba. See ENCRYPTION.txt in the Samba docs
2. disable this new behavior in Win95/98. See Win95.txt in the Samba docs

The Samba docs directory is included with any recent Samba distribution or available at <ftp://samba.anu.edu.au/pub/samba/docs/>

NOTE: You must reboot your machine for these registry changes to take effect.

Win9X in User Level Access mode

```
> I've Samba running as a NT-Server, but there is a problem:
>
> When i want to share something on the win95-client, this
client wants a
> userlist from the NT-Server (Samba).
> How can I make Samba providing a userlist?
```

Sorry. This is not yet supported. Some time after we release samba-2.0.0 we will commence the long task of implementing this functionality. For now you can should not put your Win95 or Win98 system into User Level Access mode.

Windows98 Passwords

Please refer to the following URL:

<http://support.microsoft.com/support/kb/articles/q187/2/28.asp>

*** NOTE ***

After making the registry changes referred to in this document you MUST reboot your Windows 98 PC for the changes to take effect.

AA

Note—the best way to solve this problem is to enable encrypted passwords on your Samba server. Windows 98 works well with Samba when Samba is running in encrypted password mode. Samba versions 1.9.18 and later support encrypted passwords providing it is correctly configured in smb.conf and an smbpasswd file has been created. To enable encrypted passwords on Samba, read the file docs/ENCRYPTION.txt in the Samba distribution.

If you wish to change Windows 98 to send plaintext passwords again, look on the Win98 CD in the directory "/tools/mtsutil/" you should find the files: "ptxt_on.inf" and "ptxt_off.inf" here is a clip from the mtsutil.txt file in that directory:

```
=====
PTXT_ON.INF - SENDS PLAIN-TEXT PASSWORDS TO YOUR NETWORK SERVER
=====
```

For security reasons, Windows 98 will not allow you to send plain-text passwords. The password is encrypted by default. However, Samba servers can be configured to require plain-text passwords, so you will not be able to connect to Samba servers configured in this mode unless you change a Registry entry to enable plain-text passwords.

Caution: Enabling plain-text passwords could compromise security.

To enable plain-text passwords, add the Registry entry for EnablePlainTextPassword (as a Dword) and set the value to 1 in the following Registry location:

```
HKEY_LOCAL_MACHINE\System
\CurrentControlSet\Services\VxD\Vnetsup
```

To set the value for EnablePlainTextPassword to 1:

1. Select PTXT_ON.INF found in the \Tools\MTSutil folder on the Windows 98 CD.
2. Right-Click PTXT_ON.INF.

-or-

Hold down the SHIFT key and press the function key, F10.

1. Choose INSTALL to add the EnablePlainTextPassword entry and set its value to 1.

```
=====
PTXT_OFF.INF - SENDS ENCRYPTED PASSWORDS TO YOUR NETWORK SERVER
=====
```

To re-enable the sending of encrypted passwords to your network server, add the Registry entry EnablePlainTextPassword (as a Dword) and set the value to 0 in the following Registry location:

```
HKEY_LOCAL_MACHINE\System
\CurrentControlSet\Services\VxD\Vnetsup
```

To set the value for EnablePlainTextPassword to 0:

1. Select PTXT_OFF.INF found in the \Tools\MTSutil folder on the Windows 98 CD.
2. Right-Click PTXT_OFF.INF.

-or-

Hold down the SHIFT key and press the function key, F10.

1. Choose INSTALL to add the EnablePlainTextPassword entry and set its value to 0.

XXX isn't in user level security mode

This error message means you are using server level security with a password server that isn't in user level security mode. The password server code relies on being able to send username/password pairs and getting back a yes/no response. This isn't possible unless the server is in user level security mode.

The most common reasons for this problem are:

1. you are trying to use a Win95 box as the password server. That won't work.
2. you are using a Samba server as the password server and that server is configured in share level security.

Y2K

Samba is year 2000 compliant so long as the underlying operating system that Samba is running upon are year 2000 compliant (Linux is, as are most modern UNIX systems, along with VMS, MVS and many others that Samba runs on.)

For a much more detailed discussion and the latest information see <http://samba.anu.edu.au/samba/y2k.html>

case sensitive

Many Microsoft clients and applications cannot handle case sensitive servers. They often change the case of a filename before sending it over the wire.

In Samba, Just use "short preserve case = yes" and "preserve case = yes". Never use "case sensitive = yes"

comp.protocols.smb

FAQ about comp.protocols.smb:

The newsgroup comp.protocols.smb is quite separate from the samba mailing list. Someone from outside the Samba team ways gatewaying messages from the Samba mailing list to comp.protocols.smb for a while but hopefully this has stopped now.

comp.protocols.smb does contain a lot of discussion about Samba so it is often a useful resource for Samba users.

For info on accessing usenet newsgroups like comp.protocols.smb please ask a local internet guru. The Samba Team has no way of knowing how your local news system is setup so we can't help you.

dont descend & security

Dont descend is not meant to be a security feature, it's an administrative convenience. It can be easily bypassed.

It is meant for things like /proc to prevent utilities like file manager from recursing themselves to death in a filesystem that has links back into itself.

Please use the underlying unix security of file permissions to give you real security.

file caching

Some people report problems with "caching" of data. Generally the bug report goes like this:

- create a file on a Unix box
- view the file on a PC via Samba
- change the file on the Unix box
- look at the file again on the PC via Samba and the changes are not visible

The first thing to realize is that this is the expected behavior! The SMB protocol uses a thing called "opportunistic locking". This allows the client to "safely" do client side caching of file data. The problem is that this caching is only safe if all programs access the files via SMB. As soon as you access the data via a non-SMB client then you will get data inconsistencies.

The solution is simple! Disable oplocks in smb.conf for those shares that need to be accessed simultaneously from Unix and windows. See the "oplocks" and "veto oplock files" options in smb.conf(5)

Samba-1.9.18 and the samba-2.x series support oplocks.

Samba-1.9.17 series does NOT.

In addition, you may care to explore the effect of making the following registry entries under MS Windows NT4:

```
=====
===
=====
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
LanmanWorkstation\Parameters]
"BufFilesDenyWrite"=dword:00000000
```

```
"BufNamedPipes"=dword:00000000
"UseOpportunisticLocking"=dword:00000000
"DormantFileLimit"=dword:00000000
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
LanmanWorkstation\Parameters\Linkage]
"UtilizeNtCaching"=dword:00000000
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\
Filesystem]
"Win95TruncateExtensions"=dword:00000000
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
LanManServer\Parameters]
"EnableOpLockForceClose"=dword:00000001
"EnableOpLocks"=dword:00000000
=====
===
=====
```

The following registry entry may help under Windows 9X also:

```
[HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\VxD\VR
EDIR]
"DiscardCacheOnOpen"=string:00000001
```

generic icons displayed

Some Samba users have reported a problem where the icons for programs stored on the server are not displayed, with generic "exe" icons being displayed instead. This problem seems to be caused by the client detecting the connection to the server as "slow". When this happens the client tries to make things a bit faster by not reading icon information from the remote .exe files.

The problem happens particularly when running Samba with "security=server". In server level security Samba delays the negprot and session setup replies while talking to the password server (this is necessary). To make matters worse Samba initially sends a deliberately bad password to the password server in order to detect a known bug in some NT servers where they say "yes" to all passwords. The NT server replies to this bad password very slowly (probably in an attempt to stop password cracking over the network). All this introduces a approximately 3 second delay in making the connection which is sufficient to trigger the Win95 "don't display icons" code.

We don't have a solution to this problem, but we can say that (apart from cosmetics) it is harmless.

linux 2.0.x and smbmount

If you had read the following (cut and pasted from the 1.9.18p1 Makefile) you will have noticed the first sentence specifically says that pre-2.1.70 Linux kernels cannot use the version of smbmount included with Samba. As it states you should use the smbfs utilities available via anon ftp from the site below.

```
# If you are using Linux kernel version 2.1.70 and later,
you
should
# uncomment the following line to compile the smbmount
utilities
```

```
# together with Samba. If you are using Linux kernel version
2.0.x
# you must use the smbfs utilities from
# ftp://ftp.gwdg.de/pub/linux/misc/smbfs
```

linux compile problem

```
[snip]
> Machine OS: Linux: Slackware 3.4
[snip]
> What Happened:
[snip]
> Compiling smbpass.c
> gcc: Internal compiler error: program ccl got fatal signal
6
> make: *** [smbpass.o] Error 1
```

There appears to be a problem with the GCC binaries supplied with Slackware 3.4 (and possibly a few other systems). For some reason they fall over when trying to compile the encrypted passwords code in Samba. There are 2 solutions, the first is a short term fix which doesn't always work, and as far as I can work out the second should always work:

1. Remove the -O from FLAGS1 in your Makefile.
2. Download a new set of GCC binaries (there is one somewhere under <ftp://sunsite.unc.edu/pub/Linux>) or rebuild it from the source.

setting times when not owner

```
> If I open a Microsoft Office's file like Word for instance
from my workstation,
> and I'm not the owner of that file but I have the right to
write on it,
> Microsoft Office will change the file's date to the
current
date even if I did
not make any modification to the file.
This doesn't appear if
I'm the owner of
> the file. The result of this matter is that the date of
the
files on the network
> doesn't mean nothing because the date change of a file as
soon as somebody else
> than the owner open it .
```

This is the expected behaviour and is a result of POSIX semantics. You can ask Samba to override this, however.

Please see the "dos filetimes" option in the smb.conf man page.

setup.exe and 16 bit programs

Running 16 bit programs from Windows NT on a Samba mapped drive

The Windows NT redirector has a bug when running against a Samba or Windows 95 mapped drive and attempting to run a 16 bit executable.

The problem occurs when the pathname to a 16 bit executable contains a non 8.3 filename compliant directory component, Windows NT will fail to load the program and complain it cannot find the path to the program.

It can be verified that this is a bug in Windows NT and not Samba as the same problem can be reproduced exactly when attempting to run the same program with the same pathname from a Windows 95 server (ie. the problem still exists even with no Samba server involved).

Microsoft have been made aware of this problem, it is unknown if they regard it as serious enough to provide a fix for this.

One of the reasons this problem is reported frequently is that InstallShield setup.exe executables are frequently written as 16 bit programs, and so hit this problem.

As a workaround, you may create (on a Samba server at least) a symbolic link with an 8.3 compliant name to the non 8.3 compliant directory name, and then the 16 bit program will run. Alternatively, use the 8.3 compliant mangled name to specify the path to run the binary.

This will be fixed when Samba adds the NT-specific SMB calls in Samba 2.0 as once the NT SMB calls are used this problem no longer occurs (which is why the problem doesn't occur when running against a drive mapped to a Windows NT server).

smbclient -N

```
> When getting the list of shares available on a host using
the command
> smbclient -N -L <server>
> the program always prompts for the password if the server
is a Samba server.
> It also ignores the "-N" argument when querying some (but
not all) of our
> NT servers.
```

No, it does not ignore -N, it is just that your server rejected the null password in the connection, so smbclient prompts for a password to try again.

To get the behaviour that you probably want use

```
smbclient -L host -U%
```

this will set both the username and password to null, which is an anonymous login for SMB. Using -N would only set the password to null, and this is not accepted as an anonymous login for most SMB servers.

smbfs for other unixes

```
> mount -ufs \\NTSERVER\MOUNT_DIR /sun_mount_point
```

smbfs is only available for Linux.

In Samba 2.0 we are introducing a utility called smbsh that will provide similar functionality but is portable to a wide range of Unixes.

smbpasswd: rejected session request

```
> I have installed samba everything seems to be working fine
except the
> smbpasswd executable file i can change a users password as
root with
> no problem but if a users tries to change his own
password..
it would
> give this error:
>
> smbpasswd: machine 127.0.0.1 rejected the session request.
> Error was : code 131
>
> im not sure what i have to add for it to allow this users
to change
> their own password.. it seems to be defaulting to the
local ip
> address.. ???
>
> please help .. thanks for your time.
>
```

Firstly, make sure that you do NOT have "bind interfaces only = yes" in your smb. conf file.

Secondly, if you have specified "hosts allow = xxx.xxx.xxx.xxx/yy" please add to it "localhost". ie: hosts allow = 123.45.67.0/24 127.

smbpasswd needs to be able to connect to smbd on the local machine, hence it is trying to connect to the 127.0.0.1 address.

smbsh and glibc-2.1

smbsh doesn't work with glibc-2.1 on Linux systems. That includes RedHat 6.0 and many other recent Linux distributions. It is very hard to fix this as the glibc maintainers have deliberately removed the necessary hooks for smbsh to work. They don't like the idea of user space filesystems.

The only thing we can suggest right now is to use smbfs instead.

smbtar blocksize

There was a slight error in early versions of smbtar which prevents the blocksize parameter from working correctly. This was fixed in Samba version 1.9.18.

unsubscribe

For information on unsubscribing or changing your subscribed address please see the instructions at <http://lists.samba.org/>

win98 slowdown

```
> Further to my previous post, I have made an interesting
discovery. This
> particular slowdown only occurs from clients that are
running
```

> Windows 98.

The Windows98 explorer (and possibly other programs) incorrectly set the "sync" bit in write requests to network shares. This causes an enormous slowdown as Samba (quite correctly) does a fsync() on the file after each write. Combine this with the fact that Windows98 explorer uses very small write sizes (around 1.5k) and you get really terrible results.

In Samba 1.9.18p10 and later we modified Samba to by default ignore these incorrect sync requests. This results in an enormous performance increase when using Windows98 explorer.

You can get the old (slow) behavior back using the "strict sync" option.

your server software is being unfriendly

If you get "your server software is being unfriendly" when you try to connect to a server using smbclient then it means that smbclient established a TCP connection to the server but got garbage (or nothing) back when it tried to do a NBT "session request" on the open socket. The "unfriendly" bit comes from the fact that the client is expecting one of a number of possible error codes as defined in the spec (see RFC1001/1002) but instead it got something totally different.

This usually means that you aren't successfully talking to a SMB server at all, and that the socket is connected to something else. A common cause if Samba is the server is that smbd failed to startup correctly and exited before it got to the point of answering the session request. Faulty/missing config files can do this or if you are launching via inetd then maybe your inetd.conf or /etc/services is setup incorrectly.

Just what is SMB?

Author's Note: I have added web addresses for items that were just links. Otherwise the document is the same as the one that is available from <http://www.samba.org>

Just what is SMB?

V1.1

Richard Sharpe

24-Apr-1999

Copyright c1996,1997,1998,1999 Richard Sharpe

Copying

Please see the section on Copying this document for details of my policy on use of this document.

Disclaimer

This document attempts to provide a service to people involved with the SMB (soon to be CIFS) protocol in some way. Every attempt has been made to ensure that the information is correct, but no warranties are implied. Richard Sharpe can not be held liable for any loss or consequences resulting from your use or misuse of this information.

If you have any comments, please send me mail at sharpe@ns.aus.com.

Acknowledgments

I would like to thank Andrew Tridgell for getting me started in this area by suggesting that I might like to start on smblib, Dan Shearer for much encouragement and information,

Paul Blackman for helping with this page, and a number of other people who have not given me approval to name them.

I would also like to thank the many people who have sent me positive comments and constructive feedback.

Trademarks

Microsoft, MS, Windows, Windows 95, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation. Microsoft Corporation in no way endorses this document, nor is the author in any way affiliated with Microsoft Corporation.

All other trademarks are the sole property of their respective owners.

Table of Contents

Introduction

What's New?

What is SMB?

SMB Clients and Servers Currently Available

SMB Servers

SMB Clients

Further resources on the web

Copying this document

Introduction

This document explains what the SMB protocol is and discusses the many client and server implementations of SMB that are available. The document grew out of my interest in implementing SMBlib, a portable library of SMB client routines.

SMB is an important protocol because of the large number of PCs out there that already have client and server implementations running on them. All Windows for Workgroups, Windows 95 and Windows NT systems are (or are capable of) running SMB as either a client, a server, or both.

What's New

While there are many things out there that are new, perhaps the thing of greatest interest as far as the SMB protocol is concerned is CIFS, the Common Internet File System.

What is SMB?

SMB, which stands for Server Message Block, is a protocol for sharing files, printers, serial ports, and communications abstractions such as named pipes and mail slots between computers.

The earliest document I have on the SMB protocol is an IBM document from 1985. It is a copy of an IBM Personal Computer Seminar Proceedings from May 1985. It contains the **IBM PC Network SMB Protocol**. The next document I have access to is a Microsoft/Intel document called **Microsoft Networks/OpenNET-FILE SHARING PROTOCOL** from

1987. The protocol was subsequently developed further by Microsoft and others. Many of the documents that define the SMB protocol(s) are available at ftp.microsoft.com in the SMB documentation area.

SMB is a client server, request-response protocol. The diagram to the left illustrates the way in which SMB works. The only exception to the request-response nature of SMB (that is, where the client makes requests and the server sends back responses) is when the client has requested opportunistic locks (oplocks) and the server subsequently has to break an already granted oplock because another client has requested a file open with a mode that is incompatible with the granted oplock. In this case, the server sends an unsolicited message to the client signalling the oplock break.

Servers make file systems and other resources (printers, mailslots, named pipes, APIs) available to clients on the network. Client computers may have their own hard disks, but they also want access to the shared file systems and printers on the servers.

Clients connect to servers using TCP/IP (actually NetBIOS over TCP/IP as specified in RFC1001 and RFC1002), NetBEUI or IPX/SPX. Once they have established a connection, clients can then send commands (SMBs) to the server that allow them to access shares, open files, read and write files, and generally do all the sort of things that you want to do with a file system. However, in the case of SMB, these things are done over the network.

As mentioned, SMB can run over multiple protocols. The following diagram shows this:

SMB can be used over TCP/IP, NetBEUI and IPX/SPX. If TCP/IP or NetBEUI are in use, the NetBIOS API is being used.

NetBIOS over TCP/IP seems to be referred to by many names. Microsoft refers to it as NBT in some places and NetBT in others (specifically in their Windows NT documentation and in the Windows NT registry). Others refer to it as RFCNB. NetBEUI is sometimes referred to as NBF (NetBIOS Frame Format?) by Microsoft.

NetBIOS Names. If SMB is used over TCP/IP or NetBEUI, then NetBIOS names must be used in a number of cases. NetBIOS names are up to 15 characters long, and are usually the name of the computer that is running NetBIOS. Microsoft, and some other implementors, insist that NetBIOS names be in upper case, especially when presented to servers as the CALLED NAME.

SMB Protocol Variants. Since the inception of SMB, many protocol variants have been developed to handle the increasing complexity of the environments that it has been employed in.

The actual protocol variant client and server will use is negotiated using the negprot SMB which must be the first SMB sent on a connection.

The first protocol variant was the Core Protocol, known to SMB implementations as PC NETWORK PROGRAM 1.0. It could handle a fairly basic set of operations that included:

- connecting to and disconnecting from file and print shares
- opening and closing files
- opening and closing print files
- reading and writing files
- creating and deleting files and directories
- searching directories
- getting and setting file attributes
- locking and unlocking byte ranges in files

Subsequent variants were introduced as more functionality was needed. Some of these variants and the related version of LAN Manager are:

Table .		
SMB Protocol Variant	Protocol Name	Comments
PC NETWORK PROGRAM 1.0	Core Protocol	The original version of SMB as defined in IBM's PC Network Program. Some versions were called PCLAN1.0
MICROSOFT NETWORKS 1.03	Core Plus Protocol	Included Lock&Read and Write&Unlock SMBs with different versions of raw read and raw write SMBs
MICROSOFT NETWORKS 3.0	DOS LAN Manager 1.0	The same as LANMAN1.0, but OS/2 errors must be translated to DOS errors.
LANMAN1.0	LAN Manager 1.0	The full LANMAN1.0 protocol.
DOS LM1.2X002	LAN Manager 2.0	The same as LM1.2X002, but errors must be translated to DOS errors.
LM1.2X002	LAN Manager 2.0	The full LANMAN2.0 protocol.
DOS LANMAN2.1	LAN Manager 2.1	The same as LANMAN2.1, but errors must be translated to DOS errors.
LANMAN2.1	LAN Manager 2.1	The full LANMAN2.1 protocol.
Windows for Workgroups 3.1a	LAN Manager 2.1?	Windows for Workgroups 1.0?
NT LM 0.12	NT LAN Manager 1.0?	Contains special SMBs for NT
Samba	NT LAN Manager 1.0?	Samba's version of NT LM 0.12?
CIFS 1.0	NT LAN Manager 1.0	Really NT LM 0.12 plus a bit?

Some variants introduced new SMBs, some simply changed the format of existing SMBs or responses, and some variants did both.

Security. The SMB model defines two levels of security:

Share level. Protection is applied at the share level on a server. Each share can have a password, and a client only needs that password to access all files under that share. This was the first security model that SMB had and is the only security model available in the Core and CorePlus protocols. Windows for Workgroups' vserver.exe implements share level security by default, as does Windows 95.

User Level. Protection is applied to individual files in each share and is based on user access rights. Each user (client) must log in to the server and be authenticated by the server. When it is authenticated, the client is given a UID which it must present on all subsequent accesses to the server. This model has been available since LAN Manager 1.0.

Browsing the network. Having lots of servers out in the network is not much good if users cannot find them. Of course, clients can simply be configured to know about the servers in their environment, but this does not help when new servers are to be introduced or old ones removed.

To solve this problem, browsing has been introduced. Each server broadcasts information about its presence. Clients listen for these broadcasts and build up browse lists. In a NetBEUI environment, this is satisfactory, but in a TCP/IP environment, problems arise. The problems exist because TCP/IP broadcasts are not usually sent outside the subnet in which they originate (although some routers can selectively transport broadcasts to other subnets).

Microsoft have introduced browse servers and the Windows Internet Name Service (WINS) to help overcome these problems.

CIFS: The latest incarnation? Microsoft and a group of other vendors (Digital Equipment, Data General, SCO, Network Appliance Corp, etc) are engaged in developing a public version of the SMB protocol. It is expected that CIFS 1.0 will be essentially NT LM 0.12 with some modifications for easier use over the Internet.

An Example SMB Exchange. The protocol elements (requests and responses) that clients and servers exchange are called SMBs. They have a specific format that is very similar for both requests and responses. Each consists of a fixed size header portion, followed by a variable sized parameter and data portion.

After connecting at the NetBIOS level, either via NBF, NetBT, etc, the client is ready to request services from the server. However, the client and server must first identify which protocol variant they each understand.

The client sends a *negprot* SMB to the server, listing the protocol dialects that it understands. The server responds with the index of the dialect that it wants to use, or 0xFFFF if none of the dialects was acceptable.

Dialects more recent than the Core and CorePlus protocols supply information in the *negprot* response to indicate their capabilities (max buffer size, canonical file names, etc).

Once a protocol has been established. The client can proceed to logon to the server, if required. They do this with a *sesssetupX* SMB. The response indicates whether or not they have supplied a valid username password pair and if so, can provide additional information. One of the most important aspects of the response is the UID of the logged on user. This UID must be submitted with all subsequent SMBs on that connection to the server.

Once the client has logged on (and in older protocols-Core and CorePlus-you cannot logon), the client can proceed to connect to a tree.

The client sends a *tcon* or *tconX* SMB specifying the network name of the share that they wish to connect to, and if all is kosher, the server responds with a TID that the client will use in all future SMBs relating to that share.

Having connected to a tree, the client can now open a file with an open SMB, followed by reading it with read SMBs, writing it with write SMBs, and closing it with close SMBs.

SMB Clients and Servers Currently Available

There are a few SMB clients available today and a relatively large number of servers available from a range of vendors.

The main clients are from Microsoft, and are included in Windows for WorkGroups 3.x, Windows 95, and Windows NT. They are most evident when you use the File Manager or the Windows 95 Explorer, as these allow you to connect to servers across the network. However they are also used when you open files using a UNC (universal naming convention).

Some other clients that I am aware of are:

- smbclient from Samba

- smbfs for Linux

- SMBlib (an SMB client library that is in development)

Server implementations are available from many sources. Some that I am aware of are:

- Samba

- Microsoft Windows for Workgroups 3.x

- Microsoft Windows 95

- Microsoft Windows NT

- The PATHWORKS family of servers from Digital

- LAN Manager for OS/2, SCO, etc

- VisionFS from SCO

- TotalNET Advanced Server from Syntax

- Advanced Server for UNIX from AT&T (NCR?)

- LAN Server for OS/2 from IBM

The next two sections will discuss each of the above in turn.

SMB Servers

Before discussing SMB servers, it is useful to discuss the difference between Workgroups and Domains.

Workgroups. A workgroup is a collection of computers that each maintain their own security information. With Windows for Workgroups, each server is pretty much in share level security. Windows 95 can pass user authentication off to an NT or LAN Manager server.

However, the point of a workgroup is that security is distributed, not centralized.

Domain. A domain is a collection of computers where security is handled centrally. Each domain has one or more domain controllers. There is

usually a primary domain controller and several backup domain controllers. The domain controllers maintain account style information related to users (clients), like account names, encrypted passwords, authorized hours of use, groups the user belongs to, etc.

Samba. Samba is a freely available SMB server for UNIX, OpenVMS (recently ported and maybe not very stable) developed by Andrew Tridgell and maintained by a loosely knit group of people all over the world. Samba runs on a great many UNIX variants (Linux, Solaris, SunOS, HP-UX, ULTRIX, DEC OSF/1, Digital UNIX, Dynix (Sequent), IRIX (SGI), SCO Open Server, DG-UX, UNIXWARE, AIX, BSDI, NetBSD, NEXTSTEP, A/UX, etc).

Samba implements the NT LM 0.12 protocol dialect. Samba can now participate in a domain (both as a PDC and a Member of a domain), and it can participate in browsing and can be a browse master. Samba can also process logon requests for Windows 95 systems

Samba implements user level security, but shares can be public where access is mapped to the owner etc of the share.

Microsoft Windows Servers. Microsoft has a number of SMB server implementations for the Windows range of operating systems. These are not separate products, rather, they are integral to the appropriate version of the Windows operating system. However, they can be switched off either through the Control Panel or at the command line (**net stop server** at DOS prompt).

It is clear from the fact that the Windows 95 and Windows NT SMB servers react differently to certain sequences of SMBs, that Microsoft do not use the same code for each of these servers (although the Windows for Workgroups and Windows 95 implementations may be derived from the same code).

Windows for Workgroups 3.11 implements the Windows for Workgroups 3.0a protocol variant, and implements share level security.

Windows 95 implements the NT LM 0.12 protocol level and implements both share and user level security.

Windows NT implements the NT LM 0.12 protocol level and implements both share and user level security.

LAN Manager and LAN Manager for UNIX (LM/X). Microsoft and AT&T GIS ported various LAN Manager versions to the UNIX operating system. This code formed the basis of many SMB servers available for UNIX operating systems from many vendors.

Some examples are: LM/X for SCO, LM Server for HP-UX (Advanced Server/9000), etc.

The most recent version of this software seems to be LAN Manager for UNIX Version 2.2, which implements the LANMAN2.1 protocol variant.

VisionFS. VisionFS is a written-from-scratch SMB server from SCO. It is available for Solaris 2.x, HP-UX and SCO (both SCO OpenServer and UNIXware).

TotalNET Advanced Server. This product is from Syntax. It is a completely independently written SMB server, that was perhaps the first SMB server for UNIX. These days, it comes with additional modules providing AppleShare and NetWare serving all in the one product.

Advanced Server for UNIX. After LM/X, NCR (which used to be ATT GIS) (perhaps with help from Microsoft) ported the Windows NT SMB server code to UNIX to provide the same level of functionality as Windows NT.

PATHWORKS. PATHWORKS is the name of a product family from Digital equipment corporation. It included both servers and clients, with the servers running on:

VAX and Alpha VMS

VAX and MIPS ULTRIX

DEC OSF/1 for AXP and Digital UNIX (DEC OSF/1 renamed)

OS/2

The clients ran on DOS, Windows, Windows for Workgroups, Windows NT and Windows 95 and are explained below.

Digital's clients and server implement SMB over DECnet as well as TCP/IP and more recently, NetBEUI. The SMB over DECnet specification has never been released.

Digital's original PATHWORKS servers were for VAX/VMS and implemented the CorePlus protocol (MICROSOFT NETWORKS 1.03 dialect). This product went through several versions and culminated in version 4.2. After a time, a version was done for ULTRIX and called PATHWORKS for ULTRIX V1, the highest version of which was 1.3. Both of these product streams were internally developed.

Subsequently, Digital used the AT&T and Microsoft LAN Manager for UNIX (LM/X) code. This was released as PATHWORKS V5.0 for OpenVMS (LAN Manager) and PATHWORKS V5.0 for Digital UNIX (LAN Manager). This product implements LAN Manager for UNIX V2.2 and the highest SMB dialect that it recognizes is LANMAN2.1 (and DOS LANMAN2.1). The reason for the LAN Manager in brackets at the end of each product name is that the products also support NetWare functionality.

PATHWORKS V5 is able to participate in a Windows NT based domain, albeit only as a Backup Domain Controller or a member server.

Recently, Digital has announced PATHWORKS V6.0 for UNIX (Advanced Server), which is based on AT&T's ASU (Advanced Server for UNIX) product.

LAN Server for OS/2. This is an IBM product that seems to be derived in some way from Microsoft's LAN Manager code.

SMB Clients

There are several SMB clients out there:

Microsoft Clients

Windows NT

Windows 95

Windows for Workgroups 3.11

Digital's PATHWORKS clients

Samba's smbclient

Linux's smbfs

SMBlib

Further Resources On The Web

The following are some other web pages that you can visit that are relevant to the SMB protocol:

Samba <http://samba.anu.edu.au/samba>

SMBlib <http://samba.anu.edu.au/samba/smbliib>

SCO's VisionFS <http://www.sco.com/products/visionfs>

Syntax's TotalNET Advanced Server <http://www.syntax.com/totalnet/tasbody2.htm>

Digital's PATHWORKS products <http://www.digital.com/info/pathworks>

Microsoft's Windows NT products <http://www.microsoft.com/NTServer>

IBM's LAN Server products <http://www.austin.ibm.com/pspinfo/linfo.html>

IBM's PC Integration with AIX [HTTP://www.austin.ibm.com/resource/technology/aixpcint.html](http://www.austin.ibm.com/resource/technology/aixpcint.html)

Data General's Support of Advanced Server for UNIX http://www.dg.com/products/html/dg_ux.html

smbfs <http://www.boutell.com/lsm/lsmbyid.cgi/000948> LSM entry (and smbfs ftp <ftp://ftp.gwdg.de/pub/linux/misc/smbfs> location)

CIFS Home page <http://www.microsoft.com/intdev/cifs>

*Network Appliance's Support for
CIFS*http://www.netapp.com/news/level3b/news_rel_960613.html

*HP Ships NT Server Network Operating System on Enterprise-Class HP-
UX Platform*<http://www.hp.com/pressrel/apr96/08apr96d.htm>

*AT&T GIS announces Advanced Server for UNIX
Systems*<http://www.att.com/press/0894/940822.nca.html>

*Thursby's Dave, Macintosh Client Software for Microsoft
Networking*<http://www.thursby.com>

Solstice LM Server<http://www.sun.com/sunsoft/solstice/Networking-products/lmserver.html>

*Triteal's TEDfs, an SMB server for CDE (Unix)
machines.*<http://www.triteal.com/WinTEX/evalguide/index.html>

Copying this document

I have had a number of requests for permission to use this document in other material. In one case, I was asked if someone could include this document as an appendix in a book. In another case, I was asked if the document could be handed to customers and potential customers. In both cases I felt that the request was reasonable.

My view on these matters is that this document was written to be read.

However, I would ask that you send me email stating your intended use and requesting my permission.

FeedBack

This document will be updated from time to time. If you have any comments, please feel free to send me email at sharpe@ns.aus.com
Visit me at NS Computer Software and Services P/L for more info on where I currently work.

**Copyright c1996, 1997, 1998, 1999, Richard Sharpe
Last updated 24-Apr-1999.**

Appendix C. Samba Man Pages

Appendix Objectives

- ✓ Lmhosts(5)
- ✓ Nmbd
- ✓ Samba(7)
- ✓ Samba.conf
- ✓ Smbclient(1)
- ✓ Smbd(8)
- ✓ Smbpasswd(5)
- ✓ Smbpasswd(8)
- ✓ Smbstatus(1)

The following are manual pages that should be useful to you when working with Samba.

Lmhosts (5)

Samba

23 Oct 1998

Name

lmhosts—The Samba NetBIOS hosts file

Synopsis

lmhosts is the **Samba** NetBIOS name to IP address mapping file.

Description

This file is part of the **Samba** suite.

lmhosts is the **Samba** NetBIOS name to IP address mapping file. It is very similar to the **/etc/hosts** file format, except that the hostname component must correspond to the NetBIOS naming format.

File Format

It is an ASCII file containing one line for NetBIOS name. The two fields on each line are separated from each other by white space. Any entry beginning with # is ignored. Each line in the `lmhosts` file contains the following information:

- **IP Address**— in dotted decimal format.
- **NetBIOS Name**— This name format is a maximum fifteen character host name, with an optional trailing '#' character followed by the NetBIOS name type as two hexadecimal digits.

If the trailing '#' is omitted then the given IP address will be returned for all names that match the given name, whatever the NetBIOS name type in the lookup. An example follows:

```
#
# Sample Samba lmhosts file.
#
192.9.200.1 TESTPC
192.9.200.20 NTSERVER#20
192.9.200.21 SAMBASERVER
```

Contains three IP to NetBIOS name mappings. The first and third will be returned for any queries for the names "TESTPC" and "SAMBASERVER" respectively, whatever the type component of the NetBIOS name requested.

The second mapping will be returned only when the "0x20" name type for a name "NTSERVER" is queried. Any other name type will not be resolved.

The default location of the `lmhosts` file is in the same directory as the `smb.conf` file.

Version

This man page is correct for version 2.0 of the Samba suite.

See Also

`smb.conf` (5), `smbclient` (1), `smbpasswd` (8), `samba` (7).

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of *Open Source software*, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@samba.org.

See **samba** (7) to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

nmbd

Samba

23 Oct 1998

Name

nmbd—NetBIOS name server to provide NetBIOS over IP naming services to clients

Synopsis

nmbd [-D] [-o] [-a] [-H *lmhosts* file] [-d *debuglevel*] [-l *log file basename*] [-n *primary NetBIOS name*] [-p *port number*] [-s *configuration file*] [-i *NetBIOS scope*] [-h]

Description

This program is part of the **Samba** suite.

nmbd is a server that understands and can reply to NetBIOS over IP name service requests, like those produced by SMBD/CIFS clients such as Windows 95/98, Windows NT and LanManager clients. It also participates in the browsing protocols which make up the Windows "Network Neighborhood" view.

SMB/CIFS clients, when they start up, may wish to locate an SMB/CIFS server. That is, they wish to know what IP number a specified host is using.

Amongst other services, **nmbd** will listen for such requests, and if its own NetBIOS name is specified it will respond with the IP number of the host it is running on. Its "own NetBIOS name" is by default the primary DNS name of the host it is running on, but this can be overridden with the **-n** option (see OPTIONS below). Thus **nmbd** will reply to broadcast queries for its own name(s). Additional names for **nmbd** to respond on can be set via parameters in the **smb.conf(5)** configuration file.

nmbd can also be used as a WINS (Windows Internet Name Server) server. What this basically means is that it will act as a WINS database server, creating a database from name registration requests that it receives and replying to queries from clients for these names.

In addition, **nmbd** can act as a WINS proxy, relaying broadcast queries from clients that do not understand how to talk the WINS protocol to a WIN server.

Options

- **D** If specified, this parameter causes **nmbd** to operate as a daemon. That is, it detaches itself and runs in the background, fielding requests on the appropriate port. By default, **nmbd** will NOT operate as a daemon. **nmbd** can also be operated from the **inetd** meta-daemon, although this is not recommended.
- **-a** If this parameter is specified, each new connection will append log messages to the log file. This is the default.
- **-o** If this parameter is specified, the log files will be overwritten when opened. By default, the log files will be appended to.
- **-H filename** NetBIOS *lmhosts* file.

The *lmhosts* file is a list of NetBIOS names to IP addresses that is loaded by the **nmbd** server and used via the name resolution mechanism **name resolve order** described in **smb.conf (5)** to resolve any NetBIOS name queries needed by the server. Note that the contents of this file are *NOT* used by **nmbd** to answer any name queries. Adding a line to this file affects name NetBIOS resolution from this host *ONLY*.

The default path to this file is compiled into Samba as part of the build process. Common defaults are */usr/local/samba/lib/lmhosts*, */usr/samba/lib/lmhosts* or */etc/lmhosts*. See the **lm-hosts (5)** man page for details on the contents of this file.

- **-d debuglevel** debuglevel is an integer from 0 to 10.

The default value if this parameter is not specified is zero.

The higher this value, the more detail will be logged to the log files about the activities of the server. At level 0, only critical errors and serious warnings will be logged. Level 1 is a reasonable level for day to day running—it generates a small amount of information about operations carried out.

Levels above 1 will generate considerable amounts of log data, and should only be used when investigating a problem. Levels above 3 are designed for use only by developers and generate HUGE amounts of log data, most of which is extremely cryptic.

Note that specifying this parameter here will override the **log level** parameter in the **smb.conf (5)** file.

- **-l logfile** The **-l** parameter specifies a path and base filename into which operational data from the running **nmbd** server will be logged. The actual log file name is generated by appending the extension ".nmb" to the specified base name. For example, if the name specified was "log" then the file log.nmb would contain the debugging data. The default log file path is compiled into Samba as part of the build process. Common defaults are */usr/local/samba/var/log.nmb*, */usr/samba/var/log.nmb* or */var/log/log.nmb*.
- **-n primary NetBIOS name** This option allows you to override the NetBIOS name that Samba uses for itself. This is identical to setting the **NetBIOS name** parameter in the **smb.conf** file but will override the setting in the **smb.conf** file.
- **-p UDP port number** UDP port number is a positive integer value.

This option changes the default UDP port number (normally 137) that **nmbd** responds to name queries on. Don't use this option unless you are an expert, in which case you won't need help!

- **-s configuration file** The default configuration file name is set at build time, typically as */usr/local/samba/lib/smb.conf*, but this may be changed when Samba is autoconfigured.

The file specified contains the configuration details required by the server. See **smb.conf (5)** for more information.

- **-i scope** This specifies a NetBIOS scope that **nmbd** will use to communicate with when generating NetBIOS names. For details on the use of NetBIOS scopes, see *rfc1001.txt* and *rfc1002.txt*. NetBIOS scopes are very rarely used, only set this parameter if you are the system administrator in charge of all the NetBIOS systems you communicate with.
- **-h** Prints the help information (usage) for **nmbd**.

Files

/etc/inetd.conf

If the server is to be run by the **inetd** meta-daemon, this file must contain suitable startup information for the meta-daemon.

/etc/rc

(or whatever initialization script your system uses).

If running the server as a daemon at startup, this file will need to contain an appropriate startup sequence for the server.

/usr/local/samba/lib/smb.conf

This is the default location of the **smb.conf** server configuration file. Other common places that systems install this file are `/usr/samba/lib/smb.conf` and `/etc/smb.conf`.

When run as a **WINS** server (see the **wins support** parameter in the **smb.conf (5)** man page), **nmbd** will store the WINS database in the file `wins.dat` in the `var/locks` directory configured under wherever Samba was configured to install itself.

If **nmbd** is acting as a **browse master** (see the local master parameter in the **smb.conf (5)** man page), **nmbd** will store the browsing database in the file `browse.dat` in the `var/locks` directory configured under wherever Samba was configured to install itself.

Signals

To shut down an **nmbd** process it is recommended that SIGKILL (-9) *NOT* be used, except as a last resort, as this may leave the name database in an inconsistent state. The correct way to terminate **nmbd** is to send it a SIGTERM (-15) signal and wait for it to die on its own.

nmbd will accept SIGHUP, which will cause it to dump out it's namelists into the file `namelist.debug` in the `/usr/local/samba/var/locks` directory (or the `var/locks` directory configured under wherever Samba was configured to install itself). This will also cause **nmbd** to dump out it's server database in the `log.nmb` file. In addition, the debug log level of **nmbd** may be raised by sending it a SIGUSR1 (`kill -USR1 <nmbd-pid>`) and lowered by sending it a SIGUSR2 (`kill -USR2 <nmbd-pid>`). This is to allow transient problems to be diagnosed, whilst still running at a normally low log level.

Version

This man page is correct for version 2.0 of the Samba suite.

See Also

inetd (8), **smbd (8)**, **smb.conf (5)**, **smbclient (1)**, **testparm (1)**, **testprns (1)**, and the Internet RFC's **rfc1001.txt**, **rfc1002.txt**. In addition the CIFS (formerly SMB) specification is available as a link from the Web page: <http://samba.org/cifs/>.

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@samba.org.

See **samba (7)** to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

Samba (7)

Samba

23 Oct 1998

Name

Samba—A Windows SMB/CIFS fileserver for UNIX

Synopsis

Samba

Description

The Samba software suite is a collection of programs that implements the Server Message Block (commonly abbreviated as SMB) protocol for UNIX systems. This protocol is sometimes also referred to as the Common Internet File System (CIFS), LanManager or NetBIOS protocol.

Components

The Samba suite is made up of several components. Each component is described in a separate manual page. It is strongly recommended that you read the documentation that comes with Samba and the manual pages of those components that you use. If the manual pages aren't clear enough then please send a patch or bug report to samba-bugs@samba.org.

- **smbd**
- The **smbd** (8) daemon provides the file and print services to SMB clients, such as Windows 95/98, Windows NT, Windows for Workgroups or LanManager. The configuration file for this daemon is described in **smb.conf** (5).
- **nmbd**
- The **nmbd** (8) daemon provides NetBIOS nameserving and browsing support. The configuration file for this daemon is described in **smb.conf** (5).
- **smbclient**
- The **smbclient** (1) program implements a simple ftp-like client. This is useful for accessing SMB shares on other compatible servers (such as Windows NT), and can also be used to allow a UNIX box to print to a printer attached to any SMB server (such as a PC running Windows NT).
- **testparm**
- The **testparm** (1) utility allows you to test your **smb.conf** (5) configuration file.
- **testprns**
- the **testprns** (1) utility allows you to test the printers defined in your printcap file.
- **smbstatus**
- The **smbstatus** (1) utility allows you list current connections to the **smbd** (8) server.
- **nmblookup**
- the **nmblookup** (1) utility allows NetBIOS name queries to be made from the UNIX machine.
- **make_smbcodepage**
- The **make_smbcodepage** (1) utility allows you to create SMB code page definition files for your **smbd** (8) server.
- **smbpasswd**
- The **smbpasswd** (8) utility allows you to change SMB encrypted passwords on Samba and Windows NT(tm) servers.

Availability

The Samba software suite is licensed under the GNU Public License (GPL). A copy of that license should have come with the package in the file COPYING. You are encouraged to distribute copies of the Samba suite, but please obey the terms of this license.

The latest version of the Samba suite can be obtained via anonymous ftp from samba.org in the directory `pub/samba/`. It is also available on several mirror sites worldwide.

You may also find useful information about Samba on the newsgroup `comp.protocols.smb` and the Samba mailing list. Details on how to join the mailing list are given in the README file that comes with Samba.

If you have access to a WWW viewer (such as Netscape or Mosaic) then you will also find lots of useful information, including back issues of the *Samba* mailing list, at <http://samba.org/samba/>.

Version

This man page is correct for version 2.0 of the Samba suite.

Contributions

If you wish to contribute to the Samba project, then I suggest you join the Samba mailing list at samba@samba.org. See the Web page at <http://samba.org/listproc> for details on how to do this.

If you have patches to submit or bugs to report then you may mail them directly to samba-bugs@samba.org. Note, however, that due to the enormous popularity of this package the Samba Team may take some time to respond to mail. We prefer patches in `diff -u` format.

Credits

Contributors to the project are now too numerous to mention here but all deserve the thanks of all Samba users. To see a full list, look at <ftp://samba.org/pub/samba/alpha/change-log> for the pre-CVS changes and at <ftp://samba.org/pub/samba/alpha/cvs.log> for the contributors to Samba post-CVS. CVS is the Open Source source code control system used by the Samba Team to develop Samba. The project would have been unmanageable without it.

In addition, several commercial organizations now help fund the Samba Team with money and equipment. For details see the *Samba Web* pages at <http://samba.org/samba/samba-thanks.html>.

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@samba.org.

smb.conf (5)

Samba

23 Oct 1998

Name

smb.conf—The configuration file for the Samba suite

Synopsis

smb.conf The **smb.conf** file is a configuration file for the Samba suite. **smb.conf** contains runtime configuration information for the Samba programs. The **smb.conf** file is designed to be configured and administered by the **swat (8)** program. The complete description of the file format and possible parameters held within are here for reference purposes.

File Format

The file consists of sections and parameters. A section begins with the name of the section in square brackets and continues until the next section begins. Sections contain parameters of the form

```
'name = value'
```

The file is line-based—that is, each newline-terminated line represents either a comment, a section name or a parameter.

Section and parameter names are not case sensitive.

Only the first equals sign in a parameter is significant. Whitespace before or after the first equals sign is discarded. Leading, trailing and internal whitespace in section and parameter names is irrelevant. Leading and trailing whitespace in a parameter value is discarded. Internal whitespace within a parameter value is retained verbatim.

Any line beginning with a semicolon (";") or a hash ("#") character is ignored, as are lines containing only whitespace.

Any line ending in a '\ ' is "continued" on the next line in the customary UNIX fashion.

The values following the equals sign in parameters are all either a string (no quotes needed) or a boolean, which may be given as yes/no, 0/1 or true/false. Case is not significant in boolean values, but is preserved in string values. Some items such as create modes are numeric.

Section Descriptions

Each section in the configuration file (except for the **[global]** section) describes a shared resource (known as a "*share*"). The section name is the name of the shared resource and the parameters within the section define the shares attributes.

There are three special sections, **[global]**, **[homes]** and **[printers]**, which are described under "**special sections**". The following notes apply to ordinary section descriptions.

A share consists of a directory to which access is being given plus a description of the access rights which are granted to the user of the service. Some housekeeping options are also specifiable.

Sections are either filesystem services (used by the client as an extension of their native file systems) or printable services (used by the client to access print services on the host running the server).

Sections may be designated **guest** services, in which case no password is required to access them. A specified UNIX **guest account** is used to define access privileges in this case.

Sections other than guest services will require a password to access them. The client provides the username. As older clients only provide passwords and not usernames, you may specify a list of usernames to check against the password using the **"user="** option in the share definition. For modern clients such as Windows 95/98 and Windows NT, this should not be necessary.

Note that the access rights granted by the server are masked by the access rights granted to the specified or guest UNIX user by the host system. The server does not grant more access than the host system grants.

The following sample section defines a file space share. The user has write access to the path /home/bar. The share is accessed via the share name "foo":

```
[foo]
path = /home/bar
writeable = true
```

The following sample section defines a printable share. The share is readonly, but printable. That is, the only write access permitted is via calls to open, write to and close a spool file. The 'guest ok' parameter means access will be permitted as the default guest user (specified elsewhere):

```
[aprinter]
path = /usr/spool/public
read only = true
printable = true
guest ok = true
```

Special Sections

- **The [global] section**
- Parameters in this section apply to the server as a whole, or are defaults for sections which do not specifically define certain items. See the notes under **"PARAMETERS"** for more information.
- **The [homes] section**
- If a section called 'homes' is included in the configuration file, services connecting clients to their home directories can be created on the fly by the server.

When the connection request is made, the existing sections are scanned. If a match is found, it is used. If no match is found, the requested section name is treated as a user name and looked up in the local password file. If the name exists and the correct password has been given, a share is created by cloning the [homes] section.

Some modifications are then made to the newly created share:

- The share name is changed from 'homes' to the located username
- If no path was given, the path is set to the user's home directory.

If you decide to use a **path=** line in your [homes] section then you may find it useful to use the **%S** macro. For example:

```
path=/data/pchome/%S
```

would be useful if you have different home directories for your PCs than for UNIX access. This is a fast and simple way to give a large number of clients access to their home directories with a minimum of fuss.

A similar process occurs if the requested section name is "homes", except that the share name is not changed to that of the requesting user. This method of using the [homes] section works well if different users share a client PC.

The [homes] section can specify all the parameters a normal service section can specify, though some make more sense than others. The following is a typical and suitable [homes] section:

```
[homes]
writeable = yes
```

An important point is that if guest access is specified in the [homes] section, all home directories will be visible to all clients **without a password**. In the very unlikely event that this is actually desirable, it would be wise to also specify **read only access**.

Note that the **browseable** flag for auto home directories will be inherited from the global browseable flag, not the [homes] browseable flag. This is useful as it means setting browseable=no in the [homes] section will hide the [homes] share but make any auto home directories visible.

- **The [printers] section**

This section works like [homes], but for printers.

If a [printers] section occurs in the configuration file, users are able to connect to any printer specified in the local host's printcap file.

When a connection request is made, the existing sections are scanned. If a match is found, it is used. If no match is found, but a **[homes]** section exists, it is used as described above. Otherwise, the requested section name is treated as a printer name and the appropriate printcap file is scanned to see if the requested section name is a valid printer share name. If a match is found, a new printer share is created by cloning the [printers] section.

A few modifications are then made to the newly created share:

- The share name is set to the located printer name
- If no printer name was given, the printer name is set to the located printer name
- If the share does not permit guest access and no username was given, the username is set to the located printer name.

Note that the [printers] service **MUST** be printable—if you specify otherwise, the server will refuse to load the configuration file.

Typically the path specified would be that of a world-writeable spool directory with the sticky bit set on it. A typical [printers] entry would look like this:

```
[printers]
path = /usr/spool/public
writeable = no
guest ok = yes
printable = yes
```

All aliases given for a printer in the printcap file are legitimate printer names as far as the server is concerned. If your printing subsystem doesn't work like that, you will have to set up a pseudo-printcap. This is a file consisting of one or more lines like this:

```
alias|alias|alias|alias...
```

Each alias should be an acceptable printer name for your printing subsystem. In the **[global]** section, specify the new file as your printcap. The server will then only recognize names found in your pseudo-printcap, which of course can contain whatever aliases you like. The same technique could be used simply to limit access to a subset of your local printers.

An alias, by the way, is defined as any component of the first entry of a printcap record. Records are separated by newlines, components (if there are more than one) are separated by vertical bar symbols ("|").

Note

On SYSV systems which use `lpstat` to determine what printers are defined on the system you may be able to use **"printcap name = lpstat"** to automatically obtain a list of printers. See the **"printcap name"** option for more details.

Parameters

Parameters define the specific attributes of sections.

Some parameters are specific to the **[global]** section (e.g., **security**). Some parameters are usable in all sections (e.g., **create mode**). All others are permissible only in normal sections. For the purposes of the following descriptions the **[homes]** and **[printers]** sections will be considered normal. The letter 'G' in parentheses indicates that a parameter is specific to the **[global]** section. The letter 'S' indicates that a parameter can be specified in a service specific section. Note that all 'S' parameters can also be specified in the **[global]** section—in which case they will define the default behavior for all services.

Parameters are arranged here in alphabetical order—this may not create best bedfellows, but at least you can find them! Where there are synonyms, the preferred synonym is described, others refer to the preferred synonym.

Variable Substitutions

Many of the strings that are settable in the config file can take substitutions. For example the option `path = /tmp/%u` would be interpreted as `path = /tmp/john` if the user connected with the username john.

These substitutions are mostly noted in the descriptions below, but there are some general substitutions which apply whenever they might be relevant. These are:

- **%S** = the name of the current service, if any.
- **%P** = the root directory of the current service, if any.
- **%u** = user name of the current service, if any.
- **%g** = primary group name of **%u**.
- **%U** = session user name (the user name that the client wanted, not necessarily the same as the one they got).
- **%G** = primary group name of **%U**.
- **%H** = the home directory of the user given by **%u**.
- **%v** = the Samba version.
- **%h** = the internet hostname that Samba is running on.
- **%m** = the NetBIOS name of the client machine (very useful).

- **%L** = the NetBIOS name of the server. This allows you to change your config based on what the client calls you. Your server can have a "dual personality".
- **%M** = the internet name of the client machine.
- **%N** = the name of your NIS home directory server. This is obtained from your NIS auto.map entry. If you have not compiled Samba with the **— with-automount** option then this value will be the same as **%L**.
- **%p** = the path of the service's home directory, obtained from your NIS auto.map entry. The NIS auto.map entry is split up as "%N:%p".
- **%R** = the selected protocol level after protocol negotiation. It can be one of CORE, COREPLUS, LANMAN1, LANMAN2 or NT1.
- **%d** = The process id of the current server process.
- **%a** = the architecture of the remote machine. Only some are recognized, and those may not be 100% reliable. It currently recognizes Samba, WfWg, WinNT and Win95. Anything else will be known as "UNKNOWN". If it gets it wrong then sending a level 3 log to samba-bugs@samba.org should allow it to be fixed.
- **%l** = The IP address of the client machine.
- **%T** = the current date and time.

There are some quite creative things that can be done with these substitutions and other smb.conf options.

Name Mangling

Samba supports "*name mangling*" so that DOS and Windows clients can use files that don't conform to the 8.3 format. It can also be set to adjust the case of 8.3 format filenames.

There are several options that control the way mangling is performed, and they are grouped here rather than listed separately. For the defaults look at the output of the testparm program.

All of these options can be set separately for each service (or globally, of course).

The options are:

"mangle case = yes/no" controls if names that have characters that aren't of the "default" case are mangled. For example, if this is yes then a name like "Mail" would be mangled. Default *no*.

"case sensitive = yes/no" controls whether filenames are case sensitive. If they aren't then Samba must do a filename search and match on passed names. Default *no*.

"default case = upper/lower" controls what the default case is for new filenames. Default *lower*.

"preserve case = yes/no" controls if new files are created with the case that the client passes, or if they are forced to be the "default" case. Default *Yes*.

"short preserve case = yes/no" controls if new files which conform to 8.3 syntax, that is all in upper case and of suitable length, are created upper case, or if they are forced to be the "default" case. This option can be use with **"preserve case = yes"** to permit long filenames to retain their case, while short names are lowered. Default *Yes*.

By default, Samba 2.0 has the same semantics as a Windows NT server, in that it is case insensitive but case preserving.

Note About Username/Password Validation

There are a number of ways in which a user can connect to a service. The server follows the following steps in determining if it will allow a connection to a specified service. If all the steps fail then the connection request is rejected. If one of the steps pass then the following steps are not checked.

If the service is marked **"guest only = yes"** then steps 1 to 5 are skipped.

- 1: If the client has passed a username/password pair and that username/password pair is validated by the UNIX system's password programs then the connection is made as that username. Note that this includes the `\\server\service%username` method of passing a username.
- 2: If the client has previously registered a username with the system and now supplies a correct password for that username then the connection is allowed.
- 3: The client's netbios name and any previously used user names are checked against the supplied password, if they match then the connection is allowed as the corresponding user.
- 4: If the client has previously validated a username/password pair with the server and the client has passed the validation token then that username is used. This step is skipped if `"revalidate = yes"` for this service.
- 5: If a `"user="` field is given in the `smb.conf` file for the service and the client has supplied a password, and that password matches (according to the UNIX system's password checking) with one of the usernames from the `user=` field then the connection is made as the username in the `"user="` line. If one of the username in the `user=` list begins with a ``@'` then that name expands to a list of names in the group of the same name.
- 6: If the service is a guest service then a connection is made as the username given in the `"guest account ="` for the service, irrespective of the supplied password.

Complete List of Global Parameters

Here is a list of all global parameters. See the section of each parameter for details. Note that some are synonyms.

```
announce as                default service            getwd
cache                      local group
map
announce version          dfree command              homedir
map                        local master
auto services             dns proxy                  hosts
equiv                    lock dir
bind interfaces only      domain admin group
interfaces                lock
directory
browse list               domain admin users
keepalive                 log file
change notify timeout    domain controller         kernel
oplocks                  log level
character set             domain group map          ldap
bind as                  logon drive
client code page          domain groups              ldap
passwd file              logon home
coding system             domain guest group        ldap
port                     logon path
config file               domain guest users        ldap
server                   logon script
deadtime                  domain logons              ldap
suffix                   lpq cache
time
debug timestamp          domain master              lm
announce                  machine
```

```

password
debuglevel          domain user map      lm
interval           timeout
default            encrypt passwords    load
printers           mangled
stack
max disk size      panic action         root
directory         ssl version
max log size       passwd chat          security
stat cache
max mux            passwd chat debug    server
string            stat cache
size
max open files     passwd program       shared
mem size          strip dot
max packet         password level       smb
passwd file        syslog
max ttl           password server      smbrun
syslog only
max wins ttl       preferred master     socket
address           time offset
max xmit           preferred master     socket
options           time server
message command    preload              ssl
timestamp
logs
min wins ttl       printcap             ssl CA
certDir           unix
password sync
name resolve order printcap name        ssl CA
certFile          unix
realname
netbios aliases   printer driver file  ssl
ciphers           update
encrypted
netbios name       protocol             ssl
client cert       use rhosts
nis homedir        read bmpx            ssl
client key        username
level
nt pipe support    read prediction      ssl
compatibility     username map
nt smb support     read raw             ssl
hosts             valid chars
null passwords     read size            ssl
hosts resign      wins proxy
ole locking compatibility remote announce      ssl
require clientcert wins server
os level           remote browse sync   ssl
require servercert wins support
packet size        root                 ssl
server cert        workgroup

```

```

server key          root dir          ssl
                    write raw

```

Complete List of Service Parameters

Here is a list of all service parameters. See the section of each parameter for details. Note that some are synonyms.

```

admin users          default case      fake
oplocks             hosts allow
allow hosts         delete readonly  follow
symlinks           hosts deny
alternate permissions delete veto files force
create mode        include
available          deny hosts       force
directory mode     invalid users
blocking locks     directory        force
group              locking
browsable          directory mask   force
user               lppause command
browseable         directory mode   fstype
lpq command
case sensitive     dont descend     group
lpresume
command
casesignames       dos filetime resolution guest
account            lprm command
comment            dos filetimes    guest ok
magic output
copy               exec             guest
only              magic script
create mask        fake directory create hide dot
files             mangle case
create mode        times           hide
files             mangled map
mangled names      postscript
queueresume command user
mangling char      preexec         read
list              username
map archive        preserve case    read
only              users
map hidden         print command
revalidate         valid users
map system         print ok         root
postexec          veto files
map to guest       printable        root
preexec           veto oplock
files
max connections    printer          set
directory          volume
min print space    printer driver   share
modes              wide links

```

only guest		printer driver location	short
preserve case	writable		
only user		printer name	status
write list			
oplocks		printing	strict
locking	write ok		
path		public	strict
sync	writeable		
postexec		queuepause command	sync
always			

Explanation of Each Parameter

- **admin users (S)**
- This is a list of users who will be granted administrative privileges on the share. This means that they will do all file operations as the super-user (root).
- You should use this option very carefully, as any user in this list will be able to do anything they like on the share, irrespective of file permissions.
- **Default:**
 -
 - `no admin users`
 -
- **Example:**
 -
 - `admin users = jason`
 -
- **allow hosts (S)**
- A synonym for this parameter is "**hosts allow**"

This parameter is a comma, space, or tab delimited set of hosts which are permitted to access a service.

If specified in the **[global]** section then it will apply to all services, regardless of whether the individual service has a different setting.

You can specify the hosts by name or IP number. For example, you could restrict access to only the hosts on a Class C subnet with something like "`allow hosts = 150.203.5.`". The full syntax of the list is described in the man page **hosts_access (5)**. Note that this man page may not be present on your system, so a brief description will be given here also.

Note

IF you wish to allow the **smbpasswd (8)** program to be run by local users to change their Samba passwords using the local **smbd (8)** daemon, then you *MUST* ensure that the localhost is listed in your **allow hosts** list, as **smbpasswd (8)** runs in client-server mode and is seen by the local **smbd** process as just another client.

You can also specify hosts by network/netmask pairs and by netgroup names if your system supports netgroups. The *EXCEPT* keyword can also be used to limit a wildcard list. The following examples may provide some help:

Example 1: allow localhost and all IPs in 150.203.*.* except one

```
hosts allow = localhost, 150.203. EXCEPT 150.203.6.66
```

Example 2: allow localhost and hosts that match the given network/netmask

```
hosts allow = localhost, 150.203.15.0/255.255.255.0
```

Example 3: allow a localhost plus a couple of hosts

```
hosts allow = localhost, lapland, arvidsjaur
```

Example 4: allow only hosts in NIS netgroup "foonet" or localhost, but deny access from one particular host

```
hosts allow = @foonet, localhost hosts deny = pirate
```

Note that access still requires suitable user-level passwords.

See **testparm (1)** for a way of testing your host access to see if it does what you expect.

Default: none (i.e., all hosts permitted access)

Example: allow hosts = 150.203.5. localhost myhost.mynet.edu.au

- **alternate permissions (S)**

This is a deprecated parameter. It no longer has any effect in Samba2.0. In previous versions of Samba it affected the way the DOS "read only" attribute was mapped for a file. In Samba2.0 a file is marked "read only" if the UNIX file does not have the "w" bit set for the owner of the file, regardless if the owner of the file is the currently logged on user or not.

- **announce as (G)**

This specifies what type of server **nmbd** will announce itself as, to a network neighborhood browse list. By default this is set to Windows NT. The valid options are : "NT", "Win95" or "WfW" meaning Windows NT, Windows 95 and Windows for Workgroups respectively. Do not change this parameter unless you have a specific need to stop Samba appearing as an NT server as this may prevent Samba servers from participating as browser servers correctly.

Default: announce as = NT **Example** announce as = Win95

- **announce version (G)**

This specifies the major and minor version numbers that nmbd will use when announcing itself as a server. The default is 4.2. Do not change this parameter unless you have a specific need to set a Samba server to be a downlevel server.

Default: announce version = 4.2

Example: announce version = 2.0

- **auto services (G)**

This is a list of services that you want to be automatically added to the browse lists. This is most useful for homes and printers services that would otherwise not be visible. Note that if you just want all printers in your printcap file loaded then the **"load printers"** option is easier.

Default: `no auto services`

Example: `auto services = fred lp colorlp`

- **available (S)**

This parameter lets you *"turn off"* a service. If `available = no`, then *ALL* attempts to connect to the service will fail. Such failures are logged.

Default: `available = yes`

Example: `available = no`

- **bind interfaces only (G)**

This global parameter allows the Samba admin to limit what interfaces on a machine will serve smb requests. It affects file service **smbd** and name service **nmbd** in slightly different ways.

For name service it causes **nmbd** to bind to ports 137 and 138 on the interfaces listed in the **"interfaces"** parameter. **nmbd** also binds to the "all addresses" interface (0.0.0.0) on ports 137 and 138 for the purposes of reading broadcast messages. If this option is not set then **nmbd** will service name requests on all of these sockets. If **"bind interfaces only"** is set then **nmbd** will check the source address of any packets coming in on the broadcast sockets and discard any that don't match the broadcast addresses of the interfaces in the **"interfaces"** parameter list. As unicast packets are received on the other sockets it allows **nmbd** to refuse to serve names to machines that send packets that arrive through any interfaces not listed in the **"interfaces"** list. IP Source address spoofing does defeat this simple check, however so it must not be used seriously as a security feature for **nmbd**.

For file service it causes **smbd** to bind only to the interface list given in the **"interfaces"** parameter. This restricts the networks that **smbd** will serve to packets coming in those interfaces. Note that you should not use this parameter for machines that are serving PPP or other intermittent or non-broadcast network interfaces as it will not cope with non-permanent interfaces.

In addition, to change a user's SMB password, the **smbpasswd** by default connects to the **"localhost"**—`127.0.0.1` address as an SMB client to issue the password change request. If **"bind interfaces only"** is set then unless the network address `127.0.0.1` is added to the **"interfaces"** parameter list then **smbpasswd** will fail to connect in its default mode. **smbpasswd** can be forced to use the primary IP interface of the local host by using its **"-r remote machine"** parameter, with **"remote machine"** set to the IP name of the primary interface of the local host.

Default: `bind interfaces only = False`

Example: `bind interfaces only = True`

- **blocking locks (S)**

This parameter controls the behavior of **smbd** when given a request by a client to obtain a byte range lock on a region of an open file, and the request has a time limit associated with it.

If this parameter is set and the lock range requested cannot be immediately satisfied, Samba 2.0 will internally queue the lock request, and periodically attempt to obtain the lock until the timeout period expires.

If this parameter is set to "False", then Samba 2.0 will behave as previous versions of Samba would and will fail the lock request immediately if the lock range cannot be obtained.

This parameter can be set per share.

Default: `blocking locks = True`

Example: `blocking locks = False`

- **browsable (S)**

Synonym for **browseable**.

- **browse list(G)**

This controls whether **smbd** will serve a browse list to a client doing a NetServerEnum call. Normally set to true. You should never need to change this.

Default: `browse list = Yes`

- **browseable**

This controls whether this share is seen in the list of available shares in a net view and in the browse list.

Default: `browseable = Yes`

Example: `browseable = No`

- **case sensitive (G)**

See the discussion in the section **NAME MANGLING**.

- **casesignames (G)**

Synonym for "**case sensitive**".

- **change notify timeout (G)**

One of the new NT SMB requests that Samba 2.0 supports is the "ChangeNotify" requests. This SMB allows a client to tell a server to "*watch*" a particular directory for any changes and only reply to the SMB request when a change has occurred. Such constant scanning of a directory is expensive under UNIX, hence an **smbd** daemon only performs such a scan on each requested directory once every **change notify timeout** seconds.

change notify timeout is specified in units of seconds.

Default: `change notify timeout = 60`

Example: `change notify timeout = 300`

Would change the scan time to every 5 minutes.

- **character set (G)**

This allows a **smbd** to map incoming filenames from a DOS Code page (see the **client code page** parameter) to several built in UNIX character sets. The built in code page translations are:

- **ISO8859-1** Western European UNIX character set. The parameter **client code page** *MUST* be set to code page 850 if the character set parameter is set to iso8859-1 in order for the conversion to the UNIX character set to be done correctly.

- **ISO8859-2** Eastern European UNIX character set. The parameter **client code page** *MUST* be set to code page 852 if the **character set** parameter is set to ISO8859-2 in order for the conversion to the UNIX character set to be done correctly.
- **ISO8859-5** Russian Cyrillic UNIX character set. The parameter **client code page** *MUST* be set to code page 866 if the **character set** parameter is set to ISO8859-2 in order for the conversion to the UNIX character set to be done correctly.
- **KOI8-R** Alternate mapping for Russian Cyrillic UNIX character set. The parameter **client code page** *MUST* be set to code page 866 if the **character set** parameter is set to KOI8-R in order for the conversion to the UNIX character set to be done correctly.

BUG. These MSDOS code page to UNIX character set mappings should be dynamic, like the loading of MS DOS code pages, not static.

See also **client code page**. Normally this parameter is not set, meaning no filename translation is done.

Default: `character set = <empty string>`

Example: `character set = ISO8859-1`

- **client code page (G)**

This parameter specifies the DOS code page that the clients accessing Samba are using. To determine what code page a Windows or DOS client is using, open a DOS command prompt and type the command "chcp". This will output the code page. The default for USA MS-DOS, Windows 95, and Windows NT releases is code page 437. The default for western european releases of the above operating systems is code page 850.

This parameter tells **smbd** which of the `codepage. XXX` files to dynamically load on startup. These files, described more fully in the manual page **make_smbcodepage (1)**, tell **smbd** how to map lower to upper case characters to provide the case insensitivity of filenames that Windows clients expect.

Samba currently ships with the following code page files:

- **Code Page 437-MS-DOS Latin US**
- **Code Page 737-Windows '95 Greek**
- **Code Page 850-MS-DOS Latin 1**
- **Code Page 852-MS-DOS Latin 2**
- **Code Page 861-MS-DOS Icelandic**
- **Code Page 866-MS-DOS Cyrillic**
- **Code Page 932-MS-DOS Japanese SJIS**
- **Code Page 936-MS-DOS Simplified Chinese**
- **Code Page 949-MS-DOS Korean Hangul**
- **Code Page 950-MS-DOS Traditional Chinese**

Thus this parameter may have any of the values 437, 737, 850, 852, 861, 932, 936, 949, or 950. If you don't find the codepage you need, read the comments in one of the other codepage files and the **make_smbcodepage (1)** man page and write one. Please remember to donate it back to the Samba user community.

This parameter co-operates with the **"valid chars"** parameter in determining what characters are valid in filenames and how capitalization is done. If you set both this parameter and the **"valid chars"** parameter the **"client code page"** parameter *MUST* be set before the **"valid chars"** parameter in the **smb.conf** file. The **"valid chars"** string will then augment the character settings in the "client code page" parameter.

If not set, **"client code page"** defaults to 850.

See also : "**valid chars**"

Default: `client code page = 850`

Example: `client code page = 936`

- **codingsystem (G)**

This parameter is used to determine how incoming Shift-JIS Japanese characters are mapped from the incoming "**client code page**" used by the client, into file names in the UNIX filesystem. Only useful if "**client code page**" is set to 932 (Japanese Shift-JIS).

The options are:

- **SJIS** Shift-JIS. Does no conversion of the incoming filename.
 - **JIS8, J8BB, J8BH, J8@B, J8@J, J8@H** Convert from incoming Shift-JIS to eight bit JIS code with different shift-in, shift out codes.
 - **JIS7, J7BB, J7BH, J7@B, J7@J, J7@H** Convert from incoming Shift-JIS to seven bit JIS code with different shift-in, shift out codes.
 - **JUNET, JUBB, JUBH, JU@B, JU@J, JU@H** Convert from incoming Shift-JIS to JUNET code with different shift-in, shift out codes.
 - **EUC** Convert an incoming Shift-JIS character to EUC code.
 - **HEX** Convert an incoming Shift-JIS character to a 3 byte hex representation, i.e. `:AB`.
 - **CAP** Convert an incoming Shift-JIS character to the 3 byte hex representation used by the Columbia AppleTalk Program (CAP), i.e. `:AB`. This is used for compatibility between Samba and CAP.
- **comment (S)**

This is a text field that is seen next to a share when a client does a queries the server, either via the network neighborhood or via "net view" to list what shares are available.

If you want to set the string that is displayed next to the machine name then see the server string command.

Default: `No comment string`

Example: `comment = Fred's Files`

- **config file (G)**

This allows you to override the config file to use, instead of the default (usually **smb.conf**). There is a chicken and egg problem here as this option is set in the config file!

For this reason, if the name of the config file has changed when the parameters are loaded then it will reload them from the new config file.

This option takes the usual substitutions, which can be very useful.

If the config file doesn't exist then it won't be loaded (allowing you to special case the config files of just a few clients).

Example: `config file = /usr/local/samba/lib/smb.conf.%m`

- **copy (S)**

This parameter allows you to *clone* service entries. The specified service is simply duplicated under the current service's name. Any parameters specified in the current section will override those in the section being copied.

This feature lets you set up a "template" service and create similar services easily. Note that the service being copied must occur earlier in the configuration file than the service doing the copying.

Default: `none`

Example: `copy = otherservice`

- **create mask (S)**

A synonym for this parameter is "**create mode**".

When a file is created, the necessary permissions are calculated according to the mapping from DOS modes to UNIX permissions, and the resulting UNIX mode is then bit-wise "AND"ed with this parameter. This parameter may be thought of as a bit-wise MASK for the UNIX modes of a file. Any bit *not* set here will be removed from the modes set on a file when it is created.

The default value of this parameter removes the "group" and "other" write and execute bits from the UNIX modes.

Following this Samba will bit-wise "OR" the UNIX mode created from this parameter with the value of the "force create mode" parameter which is set to 000 by default.

This parameter does not affect directory modes. See the parameter "**directory mode**" for details.

See also the "**force create mode**" parameter for forcing particular mode bits to be set on created files. See also the "**directory mode**" parameter for masking mode bits on created directories.

Default: `create mask = 0744`

Example: `create mask = 0775`

- **create mode (S)**

This is a synonym for **create mask**.

- **deadtime (G)**

The value of the parameter (a decimal integer) represents the number of minutes of inactivity before a connection is considered dead, and it is disconnected. The deadtime only takes effect if the number of open files is zero.

This is useful to stop a server's resources being exhausted by a large number of inactive connections.

Most clients have an auto-reconnect feature when a connection is broken so in most cases this parameter should be transparent to users.

Using this parameter with a timeout of a few minutes is recommended for most systems.

A deadtime of zero indicates that no auto-disconnection should be performed.

Default: `deadtime = 0`

Example: `deadtime = 15`

- **debug timestamp (G)**

Samba2.0 debug log messages are timestamped by default. If you are running at a high "**debug level**" these timestamps can be distracting. This boolean parameter allows them to be turned off.

Default: `debug timestamp = Yes`

Example: `debug timestamp = No`

- **debug level (G)**

The value of the parameter (an integer) allows the debug level (logging level) to be specified in the **smb.conf** file. This is to give greater flexibility in the configuration of the system.

The default will be the debug level specified on the command line or level zero if none was specified.

Example: `debug level = 3`

- **default (G)**

A synonym for default service.

- **default case (S)**

See the section on "**NAME MANGLING**". Also note the "**short preserve case**" parameter.

- **default service (G)**

This parameter specifies the name of a service which will be connected to if the service actually requested cannot be found. Note that the square brackets are *NOT* given in the parameter value (see example below).

There is no default value for this parameter. If this parameter is not given, attempting to connect to a nonexistent service results in an error.

Typically the default service would be a **guest ok, read-only** service.

Also note that the apparent service name will be changed to equal that of the requested service, this is very useful as it allows you to use macros like **%S** to make a wildcard service.

Note also that any ' _ ' characters in the name of the service used in the default service will get mapped to a ' / '. This allows for interesting things.

Example:

```
default service = pub
[pub]
    path = /%S
```

- **delete readonly (S)**

This parameter allows readonly files to be deleted. This is not normal DOS semantics, but is allowed by UNIX.

This option may be useful for running applications such as rcs, where UNIX file ownership prevents changing file permissions, and DOS semantics prevent deletion of a read only file.

Default: `delete readonly = No`

Example: `delete readonly = Yes`

- **delete veto files (S)**

This option is used when Samba is attempting to delete a directory that contains one or more vetoed directories (see the "**veto files**" option). If this option is set to False (the default) then if a vetoed directory contains any non-vetoed files or directories then the directory delete will fail. This is usually what you want.

If this option is set to True, then Samba will attempt to recursively delete any files and directories within the vetoed directory. This can be useful for integration with file serving systems such as **NetAtalk**, which create meta-files within directories you might normally veto DOS/Windows users from seeing (e.g. `.AppleDouble`)

Setting `"delete veto files = True"` allows these directories to be transparently deleted when the parent directory is deleted (so long as the user has permissions to do so).

See also the **veto files** parameter.

Default: `delete veto files = False`

Example: `delete veto files = True`

- **deny hosts (S)**

The opposite of **"allow hosts"**—hosts listed here are *NOT* permitted access to services unless the specific services have their own lists to override this one. Where the lists conflict, the **"allow"** list takes precedence.

Default: `none` (i.e., no hosts specifically excluded)

Example: `deny hosts = 150.203.4. badhost.mynet. edu.au`

- **dfree command (G)**

The `dfree` command setting should only be used on systems where a problem occurs with the internal disk space calculations. This has been known to happen with Ultrix, but may occur with other operating systems. The symptom that was seen was an error of "Abort Retry Ignore" at the end of each directory listing.

This setting allows the replacement of the internal routines to calculate the total disk space and amount available with an external routine. The example below gives a possible script that might fulfill this function.

The external program will be passed a single parameter indicating a directory in the filesystem being queried. This will typically consist of the string `"/"`. The script should return two integers in ascii. The first should be the total disk space in blocks, and the second should be the number of available blocks. An optional third return value can give the block size in bytes. The default blocksize is 1024 bytes.

Note

Your script should *NOT* be `setuid` or `setgid` and should be owned by (and writeable only by) root!

Default: By default internal routines for determining the disk capacity and remaining space will be used.

Example: `dfree command = /usr/local/samba/bin/dfree`

Where the script `dfree` (which must be made executable) could be:

```
#!/bin/sh
```

```
df $1 | tail -1 | awk '{print $2" "$4}'
```

or perhaps (on Sys V based systems):

```
#!/bin/sh
/usr/bin/df -k $1 | tail -1 | awk '{print $3" "$5}'
```

Note that you may have to replace the command names with full path names on some systems.

- **directory (S)**

Synonym for **path**.

- **directory mask (S)**

This parameter is the octal modes which are used when converting DOS modes to UNIX modes when creating UNIX directories.

When a directory is created, the necessary permissions are calculated according to the mapping from DOS modes to UNIX permissions, and the resulting UNIX mode is then bit-wise "AND"ed with this parameter. This parameter may be thought of as a bit-wise MASK for the UNIX modes of a directory. Any bit *not* set here will be removed from the modes set on a directory when it is created.

The default value of this parameter removes the "group" and "other" write bits from the UNIX mode, allowing only the user who owns the directory to modify it.

Following this Samba will bit-wise "OR" the UNIX mode created from this parameter with the value of the "force directory mode" parameter. This parameter is set to 000 by default (i.e. no extra mode bits are added).

See the "**force directory mode**" parameter to cause particular mode bits to always be set on created directories.

See also the "**create mode**" parameter for masking mode bits on created files.

Default: `directory mask = 0755`

Example: `directory mask = 0775`

- **directory mode (S)**

Synonym for **directory mask**.

- **dns proxy (G)**

Specifies that **nmbd** when acting as a WINS server and finding that a NetBIOS name has not been registered, should treat the NetBIOS name word-for-word as a DNS name and do a lookup with the DNS server for that name on behalf of the name-querying client. Note that the maximum length for a NetBIOS name is 15 characters, so the DNS name (or DNS alias) can likewise only be 15 characters, maximum.

nmbd spawns a second copy of itself to do the DNS name lookup requests, as doing a name lookup is a blocking action.

See also the parameter **wins support**.

Default: `dns proxy = yes`

- **domain admin group (G)**

This is an **EXPERIMENTAL** parameter that is part of the unfinished Samba NT Domain Controller Code. It has been removed as of November 98. To work with the latest code builds that may have more support for Samba NT Domain Controller functionality please

subscribe to the mailing list **Samba-ntdom** available by sending email to list-proc@samba.org

- **domain admin users (G)**

This is an **EXPERIMENTAL** parameter that is part of the unfinished Samba NT Domain Controller Code. It has been removed as of November 98. To work with the latest code builds that may have more support for Samba NT Domain Controller functionality please subscribe to the mailing list **Samba-ntdom** available by sending email to list-proc@samba.org

- **domain controller (G)**

This is a **DEPRECATED** parameter. It is currently not used within the Samba source and should be removed from all current smb.conf files. It is left behind for compatibility reasons.

- **domain group map (G)**

This option allows you to specify a file containing unique mappings of individual NT Domain Group names (in any domain) to UNIX group names. This allows NT domain groups to be presented correctly to NT users, despite the lack of native support for the NT Security model (based on VAX/VMS) in UNIX. The reader is advised to become familiar with the NT Domain system and its administration.

This option is used in conjunction with "**local group map**" and "**domain user map**". The use of these three options is trivial and often unnecessary in the case where Samba is not expected to interact with any other SAM databases (whether local workstations or Domain Controllers).

The map file is parsed line by line. If any line begins with a '#' or a ';' then it is ignored. Each line should contain a single UNIX group name on the left then a single NT Domain Group name on the right, separated by a tabstop or '='. If either name contains spaces then it should be enclosed in quotes. The line can be either of the form:

```
UNIXgroupname \\DOMAIN_NAME\\DomainGroupName
```

OR:

```
UNIXgroupname DomainGroupName
```

In the case where Samba is either an **EXPERIMENTAL** Domain Controller or it is a member of a domain using "**security = domain**", the latter format can be used: the default Domain name is the Samba Server's Domain name, specified by "**workgroup = MYGROUP**".

Any UNIX groups that are *NOT* specified in this map file are assumed to be either Local or Domain Groups, depending on the role of the Samba Server.

In the case when Samba is an **EXPERIMENTAL** Domain Controller, Samba will present *ALL* such unspecified UNIX groups as its own NT Domain Groups, with the same name.

In the case where Samba is member of a domain using "**security = domain**", Samba will check the UNIX name with its Domain Controller (see "**password server**") as if it was an NT Domain Group. If the Domain Controller says that it is not, such unspecified (unmapped) UNIX groups which also are not NT Domain Groups are treated as Local Groups in the Samba Server's local SAM database. NT Administrators will recognise these as Workstation Local Groups, which are managed by running **USRMGR.EXE** and selecting a remote Domain named "\\WORKSTATION_NAME", or by running **MUSR-MGR.EXE** on a local Workstation.

This may sound complicated, but it means that a Samba Server as either a member of a domain or as an **EXPERIMENTAL** Domain Controller will act like an NT Workstation (with a local SAM database) or an NT PDC (with a Domain SAM database) respectively,

without the need for any of the map files at all. If you **want** to get fancy, however, you can.

Note that adding an entry to map an arbitrary NT group in an arbitrary Domain to an arbitrary UNIX group *REQUIRES* the following:

- that the UNIX group exists on the UNIX server.
- that the NT Domain Group exists in the specified NT Domain
- that the UNIX Server knows about the specified Domain;
- that all the UNIX users (who are expecting to access the Samba Server as the correct NT user and with the correct NT group permissions) in the UNIX group be mapped to the correct NT Domain users in the specified NT Domain using "**domain user map**".

Failure to meet any of these requirements may result in either (or both) errors reported in the log files or (and) incorrect or missing access rights granted to users.

- **domain groups (G)**

This is an **EXPERIMENTAL** parameter that is part of the unfinished Samba NT Domain Controller Code. It has been removed as of November 98. To work with the latest code builds that may have more support for Samba NT Domain Controller functionality please subscribe to the mailing list **Samba-ntdom** available by sending email to listproc@samba.org

- **domain guest group (G)**

This is an **EXPERIMENTAL** parameter that is part of the unfinished Samba NT Domain Controller Code. It has been removed as of November 98. To work with the latest code builds that may have more support for Samba NT Domain Controller functionality please subscribe to the mailing list **Samba-ntdom** available by sending email to listproc@samba.org

- **domain guest users (G)**

This is an **EXPERIMENTAL** parameter that is part of the unfinished Samba NT Domain Controller Code. It has been removed as of November 98. To work with the latest code builds that may have more support for Samba NT Domain Controller functionality please subscribe to the mailing list **Samba-ntdom** available by sending email to listproc@samba.org

- **domain logons (G)**

If set to true, the Samba server will serve Windows 95/98 Domain logons for the **workgroup** it is in. For more details on setting up this feature see the file DOMAINS.txt in the Samba documentation directory `docs/` shipped with the source code.

Note that Win95/98 Domain logons are *NOT* the same as Windows NT Domain logons. NT Domain logons require a Primary Domain Controller (PDC) for the Domain. It is intended that in a future release Samba will be able to provide this functionality for Windows NT clients also.

Default: `domain logons = no`

- **domain master (G)**

Tell **nmbd** to enable WAN-wide browse list collation. Setting this option causes **nmbd** to claim a special domain specific NetBIOS name that identifies it as a domain master

browser for its given **workgroup**. Local master browsers in the same **workgroup** on broadcast-isolated subnets will give this **nmbd** their local browse lists, and then ask **smbd** for a complete copy of the browse list for the whole wide area network. Browser clients will then contact their local master browser, and will receive the domain-wide browse list, instead of just the list for their broadcast-isolated subnet.

Note that Windows NT Primary Domain Controllers expect to be able to claim this **workgroup** specific special NetBIOS name that identifies them as domain master browsers for that **workgroup** by default (i.e. there is no way to prevent a Windows NT PDC from attempting to do this). This means that if this parameter is set and **nmbd** claims the special name for a **workgroup** before a Windows NT PDC is able to do so then cross subnet browsing will behave strangely and may fail.

By default ("auto") Samba will attempt to become the domain master browser only if it is the Primary Domain Controller.

Default: `domain master = auto`

Example: `domain master = no`

- **domain user map (G)**

This option allows you to specify a file containing unique mappings of individual NT Domain User names (in any domain) to UNIX user names. This allows NT domain users to be presented correctly to NT systems, despite the lack of native support for the NT Security model (based on VAX/VMS) in UNIX. The reader is advised to become familiar with the NT Domain system and its administration.

This option is used in conjunction with "**local group map**" and "**domain group map**". The use of these three options is trivial and often unnecessary in the case where Samba is not expected to interact with any other SAM databases (whether local workstations or Domain Controllers).

This option, which provides (and maintains) a one-to-one link between UNIX and NT users, is *DIFFERENT* from "**username map**", which does *NOT* maintain a distinction between the name(s) it can map to and the name it maps.

The map file is parsed line by line. If any line begins with a '#' or a ';' then the line is ignored. Each line should contain a single UNIX user name on the left then a single NT Domain User name on the right, separated by a tabstop or '='. If either name contains spaces then it should be enclosed in quotes. The line can be either of the form:

```
UNIXusername \\DOMAIN_NAME\\DomainUserName
```

or:

```
UNIXusername DomainUserName
```

In the case where Samba is either an **EXPERIMENTAL** Domain Controller or it is a member of a domain using "**security = domain**", the latter format can be used: the default Domain name is the Samba Server's Domain name, specified by "**workgroup = MYGROUP**".

Any UNIX users that are *NOT* specified in this map file are assumed to be either Domain or Workstation Users, depending on the role of the Samba Server.

In the case when Samba is an **EXPERIMENTAL** Domain Controller, Samba will present *ALL* such unspecified UNIX users as its own NT Domain Users, with the same name.

In the case where Samba is a member of a domain using "**security = domain**", Samba will check the UNIX name with its Domain Controller (see "**password server**") as if it was an NT Domain User. If the Domain Controller says that it is not, such unspecified (unmapped) UNIX users which also are not NT Domain Users are treated as Local Users in the Samba Server's local SAM database. NT Administrators will recognise these as Workstation Users, which are managed by running **USRMGR.EXE** and selecting a remote Domain named "\\WORKSTATION_NAME", or by running **MUSRMRG.EXE** on a local Workstation.

This may sound complicated, but it means that a Samba Server as either a member of a domain or as an **EXPERIMENTAL** Domain Controller will act like an NT Workstation (with a local SAM database) or an NT PDC (with a Domain SAM database) respectively, without the need for any of the map files at all. If you **want** to get fancy, however, you can.

Note that adding an entry to map an arbitrary NT User in an arbitrary Domain to an arbitrary UNIX user *REQUIRES* the following:

- that the UNIX user exists on the UNIX server.
- that the NT Domain User exists in the specified NT Domain.
- that the UNIX Server knows about the specified Domain.

Failure to meet any of these requirements may result in either (or both) errors reported in the log files or (and) incorrect or missing access rights granted to users.

- **dont descend (S)**

There are certain directories on some systems (e.g., the `/proc` tree under Linux) that are either not of interest to clients or are infinitely deep (recursive). This parameter allows you to specify a comma-delimited list of directories that the server should always show as empty.

Note that Samba can be very fussy about the exact format of the "dont descend" entries. For example you may need `"/proc"` instead of just `"/proc"`. Experimentation is the best policy:-)

Default: `none` (i.e., all directories are OK to descend)

Example: `dont descend = /proc,/dev`

- **dos filetime resolution (S)**

Under the DOS and Windows FAT filesystem, the finest granularity on time resolution is two seconds. Setting this parameter for a share causes Samba to round the reported time down to the nearest two second boundary when a query call that requires one second resolution is made to **smbd**.

This option is mainly used as a compatibility option for Visual C++ when used against Samba shares. If oplocks are enabled on a share, Visual C++ uses two different time reading calls to check if a file has changed since it was last read. One of these calls uses a one-second granularity, the other uses a two second granularity. As the two second call rounds any odd second down, then if the file has a timestamp of an odd number of seconds then the two timestamps will not match and Visual C++ will keep reporting the file has changed. Setting this option causes the two timestamps to match, and Visual C++ is happy.

Default: `dos filetime resolution = False`

Example: `dos filetime resolution = True`

- **dos filetimes (S)**

Under DOS and Windows, if a user can write to a file they can change the timestamp on it. Under POSIX semantics, only the owner of the file or root may change the timestamp. By default, Samba runs with POSIX semantics and refuses to change the timestamp on a file if the user `smbd` is acting on behalf of is not the file owner. Setting this option to True allows DOS semantics and `smbd` will change the file timestamp as DOS requires.

Default: `dos filetimes = False`

Example: `dos filetimes = True`

- **encrypt passwords (G)**

This boolean controls whether encrypted passwords will be negotiated with the client. Note that Windows NT 4.0 SP3 and above and also Windows 98 will by default expect encrypted passwords unless a registry entry is changed. To use encrypted passwords in Samba see the file ENCRYPTION.txt in the Samba documentation directory `docs/` shipped with the source code.

In order for encrypted passwords to work correctly **smbd** must either have access to a local **smbpasswd (5)** file (see the **smbpasswd (8)** program for information on how to set up and maintain this file), or set the **security=** parameter to either **"server"** or **domain** which causes **smbd** to authenticate against another server.

- **exec (S)**

This is a synonym for **preexec**.

- **fake directory create times (S)**

NTFS and Windows VFAT file systems keep a create time for all files and directories. This is not the same as the `ctime`—status change time—that Unix keeps, so Samba by default reports the earliest of the various times Unix does keep. Setting this parameter for a share causes Samba to always report midnight 1-1-1980 as the create time for directories.

This option is mainly used as a compatibility option for Visual C++ when used against Samba shares. Visual C++ generated makefiles have the object directory as a dependency for each object file, and a make rule to create the directory. Also, when NMAKE compares timestamps it uses the creation time when examining a directory. Thus the object directory will be created if it does not exist, but once it does exist it will always have an earlier timestamp than the object files it contains.

However, Unix time semantics mean that the create time reported by Samba will be updated whenever a file is created or deleted in the directory. NMAKE therefore finds all object files in the object directory but the last one built are out of date compared to the directory and rebuilds them. Enabling this option ensures directories always predate their contents and an NMAKE build will proceed as expected.

Default: `fake directory create times = False`

Example: `fake directory create times = True`

- **fake oplocks (S)**

Oplocks are the way that SMB clients get permission from a server to locally cache file operations. If a server grants an oplock (opportunistic lock) then the client is free to assume that it is the only one accessing the file and it will aggressively cache file data. With some oplock types the client may even cache file open/close operations. This can give enormous performance benefits.

When you set `"fake oplocks = yes"` **smbd** will always grant oplock requests no matter how many clients are using the file.

It is generally much better to use the real **oplocks** support rather than this parameter.

If you enable this option on all read-only shares or shares that you know will only be accessed from one client at a time such as physically read-only media like CDROMs, you will see a big performance improvement on many operations. If you enable this option on shares where multiple clients may be accessing the files read-write at the same time you can get data corruption. Use this option carefully!

This option is disabled by default.

- **follow symlinks (S)**

This parameter allows the Samba administrator to stop **smbd** from following symbolic links in a particular share. Setting this parameter to `"No"` prevents any file or directory that

is a symbolic link from being followed (the user will get an error). This option is very useful to stop users from adding a symbolic link to `/etc/passwd` in their home directory for instance. However it will slow filename lookups down slightly. This option is enabled (i.e. **smbd** will follow symbolic links) by default.

- **force create mode (S)**

This parameter specifies a set of UNIX mode bit permissions that will **always** be set on a file created by Samba. This is done by bitwise "OR"ing these bits onto the mode bits of a file that is being created. The default for this parameter is (in octal) 000. The modes in this parameter are bitwise "OR"ed onto the file mode after the mask set in the **"create mask"** parameter is applied.

See also the parameter **"create mask"** for details on masking mode bits on created files.

Default: `force create mode = 000`

Example: `force create mode = 0755`

would force all created files to have read and execute permissions set for "group" and "other" as well as the read/write/execute bits set for the "user".

- **force directory mode (S)**

This parameter specifies a set of UNIX mode bit permissions that will **always** be set on a directory created by Samba. This is done by bitwise "OR"ing these bits onto the mode bits of a directory that is being created. The default for this parameter is (in octal) 0000 which will not add any extra permission bits to a created directory. This operation is done after the mode mask in the parameter **"directory mask"** is applied.

See also the parameter **"directory mask"** for details on masking mode bits on created directories.

Default: `force directory mode = 000`

Example: `force directory mode = 0755`

would force all created directories to have read and execute permissions set for "group" and "other" as well as the read/write/execute bits set for the "user".

- **force group (S)**

This specifies a UNIX group name that will be assigned as the default primary group for all users connecting to this service. This is useful for sharing files by ensuring that all access to files on service will use the named group for their permissions checking. Thus, by assigning permissions for this group to the files and directories within this service the Samba administrator can restrict or allow sharing of these files.

Default: `no forced group`

Example: `force group = agroup`

- **force user (S)**

This specifies a UNIX user name that will be assigned as the default user for all users connecting to this service. This is useful for sharing files. You should also use it carefully as using it incorrectly can cause security problems.

This user name only gets used once a connection is established. Thus clients still need to connect as a valid user and supply a valid password. Once connected, all file operations will be performed as the `"forced user"`, no matter what username the client connected as.

This can be very useful.

Default: `no forced user`

Example: `force user = auser`

- **fstype (S)**

This parameter allows the administrator to configure the string that specifies the type of filesystem a share is using that is reported by **smbd** when a client queries the filesystem type for a share. The default type is **"NTFS"** for compatibility with Windows NT but this can be changed to other strings such as "Samba" or "FAT" if required.

Default: `fstype = NTFS`

Example: `fstype = Samba`

- **getwd cache (G)**

This is a tuning option. When this is enabled a caching algorithm will be used to reduce the time taken for `getwd()` calls. This can have a significant impact on performance, especially when the **widelinks** parameter is set to `False`.

Default: `getwd cache = No`

Example: `getwd cache = Yes`

- **group (S)**

Synonym for **"force group"**.

- **guest account (S)**

This is a username which will be used for access to services which are specified as **"guest ok"** (see below). Whatever privileges this user has will be available to any client connecting to the guest service. Typically this user will exist in the password file, but will not have a valid login. The user account **"ftp"** is often a good choice for this parameter. If a username is specified in a given service, the specified username overrides this one. On some systems the default guest account "nobody" may not be able to print. Use another account in this case. You should test this by trying to log in as your guest user (perhaps by using the `"su -"` command) and trying to print using the system print command such as **lpr (1)** or **lp (1)**.

Default: `specified at compile time, usually "nobody"`

Example: `guest account = ftp`

- **guest ok (S)**

If this parameter is `"yes"` for a service, then no password is required to connect to the service. Privileges will be those of the **guest account**.

See the section below on **security** for more information about this option.

Default: `guest ok = no`

Example: `guest ok = yes`

- **guest only (S)**

If this parameter is `"yes"` for a service, then only guest connections to the service are permitted. This parameter will have no affect if **"guest ok"** or **"public"** is not set for the service.

See the section below on **security** for more information about this option.

Default: `guest only = no`

Example: `guest only = yes`

- **hide dot files (S)**

This is a boolean parameter that controls whether files starting with a dot appear as hidden files.

Default: `hide dot files = yes`

Example: `hide dot files = no`

- **hide files(S)**

This is a list of files or directories that are not visible but are accessible. The DOS "hidden" attribute is applied to any files or directories that match.

Each entry in the list must be separated by a ' / ', which allows spaces to be included in the entry. '*' and '?' can be used to specify multiple files or directories as in DOS wildcards.

Each entry must be a Unix path, not a DOS path and must not include the Unix directory separator '/ '.

Note that the case sensitivity option is applicable in hiding files.

Setting this parameter will affect the performance of Samba, as it will be forced to check all files and directories for a match as they are scanned.

See also "**hide dot files**", "**veto files**" and "**case sensitive**".

Default

`No files or directories are hidden by this option (dot files are hidden by default because of the "hide dot files" option).`

Example `hide files =`

`./*/DesktopFolderDB/TrashFor%m/resource. frk/`

The above example is based on files that the Macintosh SMB client (DAVE) available from **Thursby** creates for internal use, and also still hides all files beginning with a dot.

- **homedir map (G)**

If "**nis homedir**" is true, and **smbd** is also acting as a Win95/98 **logon server** then this parameter specifies the NIS (or YP) map from which the server for the user's home directory should be extracted. At present, only the Sun auto.home map format is understood. The form of the map is:

`username server:/some/file/system`

and the program will extract the servername from before the first ':'. There should probably be a better parsing system that copes with different map formats and also Amd (another automounter) maps.

NB: A working NIS is required on the system for this option to work.

See also "**nis homedir**", **domain logons**.

Default: `homedir map = auto.home`

Example: `homedir map = amd.homedir`

- **hosts allow (S)**

Synonym for **allow hosts**.

- **hosts deny (S)**

Synonym for **denyhosts**.

- **hosts equiv (G)**

If this global parameter is a non-null string, it specifies the name of a file to read for the names of hosts and users who will be allowed access without specifying a password. This is not to be confused with **allow hosts** which is about hosts access to services and is more useful for guest services. **hosts equiv** may be useful for NT clients which will not supply passwords to samba.

Note

The use of **hosts equiv** can be a major security hole. This is because you are trusting the PC to supply the correct username. It is very easy to get a PC to supply a false username. I recommend that the **hosts equiv** option be only used if you really know what you are doing, or perhaps on a home network where you trust your spouse and kids. And only if you *really* trust them :-).

Default `No host equivalences`

Example `hosts equiv = /etc/hosts.equiv`

- **include (G)**

This allows you to include one config file inside another. The file is included literally, as though typed in place.

It takes the standard substitutions, except **%u**, **%P** and **%S**.

- **interfaces (G)**

This option allows you to setup multiple network interfaces, so that Samba can properly handle browsing on all interfaces.

The option takes a list of ip/netmask pairs. The netmask may either be a bitmask, or a bitlength.

For example, the following line:

```
interfaces = 192.168.2.10/24 192.168.3.10/24
```

would configure two network interfaces with IP addresses 192.168.2.10 and 192.168.3.10. The netmasks of both interfaces would be set to 255.255.255.0.

You could produce an equivalent result by using:

```
interfaces = 192.168.2.10/255.255.255.0 192.168.3.10/255.255.255.0
```

if you prefer that format.

If this option is not set then Samba will attempt to find a primary interface, but won't attempt to configure more than one interface.

See also "**bind interfaces only**".

- **invalid users (S)**

This is a list of users that should not be allowed to login to this service. This is really a *"paranoid"* check to absolutely ensure an improper setting does not breach your security. A name starting with a '@' is interpreted as an NIS netgroup first (if your system supports NIS), and then as a UNIX group if the name was not found in the NIS netgroup database.

A name starting with '+' is interpreted only by looking in the UNIX group database. A name starting with '&' is interpreted only by looking in the NIS netgroup database (this requires NIS to be working on your system). The characters '+' and '&' may be used at the start of the name in either order so the value "+&group" means check the UNIX group database, followed by the NIS netgroup database, and the value "&group" means check the NIS netgroup database, followed by the UNIX group database (the same as the '@' prefix).

The current servicename is substituted for %S. This is useful in the **[homes]** section.

See also **"valid users"**.

Default: No invalid users

Example: invalid users = root fred admin @wheel

- **keepalive (G)**

The value of the parameter (an integer) represents the number of seconds between **"keepalive"** packets. If this parameter is zero, no keepalive packets will be sent. Keepalive packets, if sent, allow the server to tell whether a client is still present and responding.

Keepalives should, in general, not be needed if the socket being used has the SO_KEEPALIVE attribute set on it (see **"socket options"**). Basically you should only use this option if you strike difficulties.

Default: keep alive = 0

Example: keep alive = 60

- **kernel oplocks (G)**

For UNIXs that support kernel based **oplocks** (currently only IRIX but hopefully also Linux and FreeBSD soon) this parameter allows the use of them to be turned on or off.

Kernel oplocks support allows Samba **oplocks** to be broken whenever a local UNIX process or NFS operation accesses a file that **smbd** has oplocked. This allows complete data consistency between SMB/CIFS, NFS and local file access (and is a *very cool* feature:-).

This parameter defaults to *"On"* on systems that have the support, and *"off"* on systems that don't. You should never need to touch this parameter.

- **ldap bind as (G)**

This parameter is part of the *EXPERIMENTAL* Samba support for a password database stored on an LDAP server. These options are only available if your version of Samba was configured with the **—with-ldap** option.

This parameter specifies the entity to bind to an LDAP directory as. Usually it should be safe to use the LDAP root account; for larger installations it may be preferable to restrict Samba's access. See also **ldap passwd file**.

Default: none (bind anonymously)

Example: ldap bind as = "uid=root, dc=mydomain, dc=org"

- **ldap passwd file (G)**

This parameter is part of the *EXPERIMENTAL* Samba support for a password database stored on an LDAP server. These options are only available if your version of Samba was configured with the **—with-ldap** option.

This parameter specifies a file containing the password with which Samba should bind to an LDAP server. For obvious security reasons this file must be set to mode 700 or less.

Default: none (bind anonymously)

Example: ldap passwd file =
/usr/local/samba/private/ldappasswd

- **ldap port (G)**

This parameter is part of the *EXPERIMENTAL* Samba support for a password database stored on an LDAP server. These options are only available if your version of Samba was configured with the **—with-ldap** option.

This parameter specifies the TCP port number of the LDAP server.

Default: ldap port = 389.

- **ldap server (G)**

This parameter is part of the *EXPERIMENTAL* Samba support for a password database stored on an LDAP server back-end. These options are only available if your version of Samba was configured with the **—with-ldap** option.

This parameter specifies the DNS name of the LDAP server to use when storing and retrieving information about Samba users and groups.

Default: ldap server = localhost

- **ldap suffix (G)**

This parameter is part of the *EXPERIMENTAL* Samba support for a password database stored on an LDAP server back-end. These options are only available if your version of Samba was configured with the **—with-ldap** option.

This parameter specifies the node of the LDAP tree beneath which Samba should store its information. This parameter **MUST** be provided when using LDAP with Samba.

Default: none

Example: ldap suffix = "dc=mydomain, dc=org"

- **lm announce (G)**

This parameter determines if **nmbd** will produce Lanman announce broadcasts that are needed by **OS/2** clients in order for them to see the Samba server in their browse list. This parameter can have three values, "true", "false", or "auto". The default is "auto". If set to "false" Samba will never produce these broadcasts. If set to "true" Samba will produce Lanman announce broadcasts at a frequency set by the parameter "**lm interval**". If set to "auto" Samba will not send Lanman announce broadcasts by default but will listen for them. If it hears such a broadcast on the wire it will then start sending them at a frequency set by the parameter "**lm interval**".

See also "**lm interval**".

Default: lm announce = auto

Example: lm announce = true

- **lm interval (G)**

If Samba is set to produce Lanman announce broadcasts needed by **OS/2** clients (see the "**lm announce**" parameter) then this parameter defines the frequency in seconds

with which they will be made. If this is set to zero then no Lanman announcements will be made despite the setting of the "**lm announce**" parameter.

See also "**lm announce**".

Default: `lm interval = 60`

Example: `lm interval = 120`

- **load printers (G)**

A boolean variable that controls whether all printers in the printcap will be loaded for browsing by default. See the "**printers**" section for more details.

Default: `load printers = yes`

Example: `load printers = no`

- **local group map (G)**

This option allows you to specify a file containing unique mappings of individual NT Local Group names (in any domain) to UNIX group names. This allows NT Local groups (aliases) to be presented correctly to NT users, despite the lack of native support for the NT Security model (based on VAX/VMS) in UNIX. The reader is advised to become familiar with the NT Domain system and its administration.

This option is used in conjunction with "**domain group map**" and "**domain name map**". The use of these three options is trivial and often unnecessary in the case where Samba is not expected to interact with any other SAM databases (whether local workstations or Domain Controllers).

The map file is parsed line by line. If any line begins with a '#' or a ';' then it is ignored. Each line should contain a single UNIX group name on the left then a single NT Local Group name on the right, separated by a tabstop or '='. If either name contains spaces then it should be enclosed in quotes. The line can be either of the form:

```
UNIXgroupname \\DOMAIN_NAME\\LocalGroupName
```

or:

```
UNIXgroupname LocalGroupName
```

In the case where Samba is either an **EXPERIMENTAL** Domain Controller or it is a member of a domain using "**security = domain**", the latter format can be used: the default Domain name is the Samba Server's Domain name, specified by "**workgroup = MYGROUP**".

Any UNIX groups that are *NOT* specified in this map file are treated as either Local or Domain Groups depending on the role of the Samba Server.

In the case when Samba is an **EXPERIMENTAL** Domain Controller, Samba will present *ALL* unspecified UNIX groups as its own NT Domain Groups, with the same name, and *NOT* as Local Groups.

In the case where Samba is member of a domain using "**security = domain**", Samba will check the UNIX name with its Domain Controller (see "**password server**") as if it was an NT Domain Group. If the Domain Controller says that it is not, such unspecified (unmapped) UNIX groups which also are not NT Domain Groups are treated as Local Groups in the Samba Server's local SAM database. NT Administrators will recognise these as Workstation Local Groups, which are managed by running **USRMGR.EXE** and selecting a remote Domain named "\\WORKSTATION_NAME", or by running **MUSR-MGR.EXE** on a local Workstation.

This may sound complicated, but it means that a Samba Server as either a member of a domain or as an **EXPERIMENTAL** Domain Controller will act like an NT Workstation (with a local SAM database) or an NT PDC (with a Domain SAM database) respectively, without the need for any of the map files at all. If you **want** to get fancy, however, you can.

Note that adding an entry to map an arbitrary NT group in an arbitrary Domain to an arbitrary UNIX group *REQUIRES* the following:

- that the UNIX group exists on the UNIX server.
- that the NT Domain Group exists in the specified NT Domain
- that the UNIX Server knows about the specified Domain;
- that all the UNIX users (who are expecting to access the Samba Server as the correct NT user and with the correct NT group permissions) in the UNIX group be mapped to the correct NT Domain users in the specified NT Domain using "**domain user map**".

Failure to meet any of these requirements may result in either (or both) errors reported in the log files or (and) incorrect or missing access rights granted to users.

- **local master (G)**

This option allows **nmbd** to try and become a local master browser on a subnet. If set to False then **nmbd** will not attempt to become a local master browser on a subnet and will also lose in all browsing elections. By default this value is set to true. Setting this value to true doesn't mean that Samba will *become* the local master browser on a subnet, just that **nmbd** will *participate* in elections for local master browser.

Setting this value to False will cause **nmbd** *never* to become a local master browser.

Default: `local master = yes`

- **lock dir (G)**

Synonym for "**lock directory**".

- **lock directory (G)**

This option specifies the directory where lock files will be placed. The lock files are used to implement the "**max connections**" option.

Default: `lock directory = /tmp/samba`

Example: `lock directory = /usr/local/samba/var/locks`

- **locking (S)**

This controls whether or not locking will be performed by the server in response to lock requests from the client.

If "`locking = no`", all lock and unlock requests will appear to succeed and all lock queries will indicate that the queried lock is clear.

If "`locking = yes`", real locking will be performed by the server.

This option *may* be useful for read-only filesystems which *may* not need locking (such as cdrom drives), although setting this parameter of "`no`" is not really recommended even in this case.

Be careful about disabling locking either globally or in a specific service, as lack of locking may result in data corruption. You should never need to set this parameter.

Default: `locking = yes`

Example: `locking = no`

- **log file (G)**

This options allows you to override the name of the Samba log file (also known as the debug file).

This option takes the standard substitutions, allowing you to have separate log files for each user or machine.

Example: `log file = /usr/local/samba/var/log.%m`

- **log level (G)**

Synonym for "debug level".

- **logon drive (G)**

This parameter specifies the local path to which the home directory will be connected (see "logon home") and is only used by NT Workstations.

Note that this option is only useful if Samba is set up as a **logon server**.

Example: `logon drive = h:`

- **logon home (G)**

This parameter specifies the home directory location when a Win95/98 or NT Workstation logs into a Samba PDC. It allows you to do

```
"NET USE H: /HOME"
```

from a command prompt, for example.

This option takes the standard substitutions, allowing you to have separate logon scripts for each user or machine.

Note that this option is only useful if Samba is set up as a **logon server**.

Example: `logon home = "\\remote_smb_server\%U"`

Default: `logon home = "\\%N\%U"`

- **logon path (G)**

This parameter specifies the home directory where roaming profiles (USER.DAT / USER.MAN files for Windows 95/98) are stored.

This option takes the standard substitutions, allowing you to have separate logon scripts for each user or machine. It also specifies the directory from which the "desk-top", "start menu", "network neighborhood" and "programs" folders, and their contents, are loaded and displayed on your Windows 95/98 client.

The share and the path must be readable by the user for the preferences and directories to be loaded onto the Windows 95/98 client. The share must be writeable when the logs in for the first time, in order that the Windows 95/98 client can create the user.dat and other directories.

Thereafter, the directories and any of the contents can, if required, be made read-only. It is not advisable that the USER.DAT file be made read-only—rename it to USER.MAN to achieve the desired effect (a *MAND*atory profile).

Windows clients can sometimes maintain a connection to the [homes] share, even though there is no user logged in. Therefore, it is vital that the logon path does not include a reference to the homes share (i.e. setting this parameter to `\\%N\HOMES\profile_path` will cause problems).

This option takes the standard substitutions, allowing you to have separate logon scripts for each user or machine.

Note that this option is only useful if Samba is set up as a **logon server**.

Default: `logon path = \\%N\%U\profile`

Example: `logon path = \\PROFILESERVER\HOME_DIR\%U\PROFILE`

- **logon script (G)**

This parameter specifies the batch file (.bat) or NT command file (.cmd) to be downloaded and run on a machine when a user successfully logs in. The file must contain the DOS style cr/lf line endings. Using a DOS-style editor to create the file is recommended. The script must be a relative path to the [netlogon] service. If the [netlogon] service specifies a **path** of /usr/local/samba/netlogon, and logon script = STARTUP.BAT, then the file that will be downloaded is:

```
/usr/local/samba/netlogon/STARTUP.BAT
```

The contents of the batch file is entirely your choice. A suggested command would be to add `NET TIME \\SERVER /SET /YES`, to force every machine to synchronize clocks with the same time server. Another use would be to add `NET USE U: \\SERVER\UTILS` for commonly used utilities, or `NET USE Q: \\SERVER\ISO9001_QA` for example.

Note that it is particularly important not to allow write access to the [netlogon] share, or to grant users write permission on the batch files in a secure environment, as this would allow the batch files to be arbitrarily modified and security to be breached.

This option takes the standard substitutions, allowing you to have separate logon scripts for each user or machine.

Note that this option is only useful if Samba is set up as a **logon server**.

Example: `logon script = scripts\%U.bat`

- **lppause command (S)**

This parameter specifies the command to be executed on the server host in order to stop printing or spooling a specific print job.

This command should be a program or script which takes a printer name and job number to pause the print job. One way of implementing this is by using job priorities, where jobs having a too low priority won't be sent to the printer.

If a "%p" is given then the printname is put in its place. A "%j" is replaced with the job number (an integer). On HPUX (see **printing=hpux**), if the "-p%p" option is added to the lpq command, the job will show up with the correct status, i.e. if the job priority is lower than the set fence priority it will have the PAUSED status, whereas if the priority is equal or higher it will have the SPOOLED or PRINTING status.

Note that it is good practice to include the absolute path in the lppause command as the PATH may not be available to the server.

See also the "**printing**" parameter.

Default: Currently no default value is given to this string, unless the value of the "**printing**" parameter is `SYSV`, in which case the default is:

```
lp -i %p-%j -H hold
```

or if the value of the "**printing**" parameter is `softq`, then the default is:

```
qstat -s -j%j -h
```

Example for HPUX: `lppause command = /usr/bin/lpalt %p-%j -p0`

- **lpq cache time (G)**

This controls how long lpq info will be cached for to prevent the **lpq** command being called too often. A separate cache is kept for each variation of the **lpq** command used by the system, so if you use different **lpq** commands for different users then they won't share cache information.

The cache files are stored in `/tmp/lpq.xxxx` where `xxxx` is a hash of the **lpq** command in use.

The default is 10 seconds, meaning that the cached results of a previous identical **lpq** command will be used if the cached data is less than 10 seconds old. A large value may be advisable if your **lpq** command is very slow.

A value of 0 will disable caching completely.

See also the "**printing**" parameter.

Default: `lpq cache time = 10`

Example: `lpq cache time = 30`

- **lpq command (S)**

This parameter specifies the command to be executed on the server host in order to obtain "**lpq**"-style printer status information.

This command should be a program or script which takes a printer name as its only parameter and outputs printer status information.

Currently eight styles of printer status information are supported; BSD, AIX, LPRNG, PLP, SYSV, HPUX, QNX and SOFTQ. This covers most UNIX systems. You control which type is expected using the "**printing =**" option.

Some clients (notably Windows for Workgroups) may not correctly send the connection number for the printer they are requesting status information about. To get around this, the server reports on the first printer service connected to by the client. This only happens if the connection number sent is invalid.

If a `%p` is given then the printername is put in its place. Otherwise it is placed at the end of the command.

Note that it is good practice to include the absolute path in the **lpq command** as the PATH may not be available to the server.

See also the "**printing**" parameter.

Default: depends on the setting of `printing =`

Example: `lpq command = /usr/bin/lpq %p`

- **lpresume command (S)**

This parameter specifies the command to be executed on the server host in order to restart or continue printing or spooling a specific print job.

This command should be a program or script which takes a printer name and job number to resume the print job. See also the "**lppause command**" parameter.

If a `%p` is given then the printername is put in its place. A `%j` is replaced with the job number (an integer).

Note that it is good practice to include the absolute path in the **lpresume command** as the PATH may not be available to the server.

See also the "**printing**" parameter.

Default:

Currently no default value is given to this string, unless the value of the "**printing**" parameter is `SYSV`, in which case the default is:

```
lp -i %p-%j -H resume
```

or if the value of the "**printing**" parameter is `softq`, then the default is:

```
qstat -s -j%j -r
```

Example for HPUX: `lpresume command = /usr/bin/lpalt %p-%j -p2`

- **lprm command (S)**

This parameter specifies the command to be executed on the server host in order to delete a print job.

This command should be a program or script which takes a printer name and job number, and deletes the print job.

If a `%p` is given then the printername is put in its place. A `%j` is replaced with the job number (an integer).

Note that it is good practice to include the absolute path in the **lprm command** as the PATH may not be available to the server.

See also the **"printing"** parameter.

Default: depends on the setting of "printing ="

Example 1: `lprm command = /usr/bin/lprm -P%p %j`

Example 2: `lprm command = /usr/bin/cancel %p-%j`

- **machine password timeout (G)**

If a Samba server is a member of an Windows NT Domain (see the **"security=domain"** parameter) then periodically a running **smbd** process will try and change the **MACHINE ACCOUNT PASSWORD** stored in the file called `<Domain>.<Machine>.mac` where `<Domain>` is the name of the Domain we are a member of and `<Machine>` is the primary **"NetBIOS name"** of the machine **smbd** is running on. This parameter specifies how often this password will be changed, in seconds. The default is one week (expressed in seconds), the same as a Windows NT Domain member server.

See also **smbpasswd (8)**, and the **"security=domain"** parameter.

Default: `machine password timeout = 604800`

- **magic output (S)**

This parameter specifies the name of a file which will contain output created by a magic script (see the **"magic script"** parameter below).

Warning: If two clients use the same **"magic script"** in the same directory the output file content is undefined.

Default: `magic output = <magic script name>.out`

Example: `magic output = myfile.txt`

- **magic script (S)**

This parameter specifies the name of a file which, if opened, will be executed by the server when the file is closed. This allows a UNIX script to be sent to the Samba host and executed on behalf of the connected user.

Scripts executed in this way will be deleted upon completion, permissions permitting.

If the script generates output, output will be sent to the file specified by the **"magic output"** parameter (see above).

Note that some shells are unable to interpret scripts containing carriage-return-linefeed instead of linefeed as the end-of-line marker. Magic scripts must be executable *"as is"* on the host, which for some hosts and some shells will require filtering at the DOS end.

Magic scripts are *EXPERIMENTAL* and should *NOT* be relied upon.

Default: `None. Magic scripts disabled.`

Example: `magic script = user.csh`

- **mangle case (S)**

See the section on **"NAME MANGLING"**.

- **mangled map (S)**

This is for those who want to directly map UNIX file names which can not be represented on Windows/DOS. The mangling of names is not always what is needed. In particular you may have documents with file extensions that differ between DOS and UNIX. For example, under UNIX it is common to use ".html" for HTML files, whereas under Windows/DOS ".htm" is more commonly used.

So to map "html" to "htm" you would use:

```
mangled map = (*.html *.htm)
```

One very useful case is to remove the annoying ";1" off the ends of filenames on some CDROMS (only visible under some UNIXs). To do this use a map of (*;1 *).

Default: no mangled map

Example: mangled map = (*;1 *)

- **mangled names (S)**

This controls whether non-DOS names under UNIX should be mapped to DOS-compatible names ("mangled") and made visible, or whether non-DOS names should simply be ignored.

See the section on "**NAME MANGLING**" for details on how to control the mangling process.

If mangling is used then the mangling algorithm is as follows:

- The first (up to) five alphanumeric characters before the rightmost dot of the filename are preserved, forced to upper case, and appear as the first (up to) five characters of the mangled name.
- A tilde "~" is appended to the first part of the mangled name, followed by a two-character unique sequence, based on the original root name (i.e., the original filename minus its final extension). The final extension is included in the hash calculation only if it contains any upper case characters or is longer than three characters.

Note that the character to use may be specified using the "**mangling char**" option, if you don't like '~'.

- The first three alphanumeric characters of the final extension are preserved, forced to upper case and appear as the extension of the mangled name. The final extension is defined as that part of the original filename after the rightmost dot. If there are no dots in the filename, the mangled name will have no extension (except in the case of "**hidden files**"—see below).
- Files whose UNIX name begins with a dot will be presented as DOS hidden files. The mangled name will be created as for other filenames, but with the leading dot removed and "_____" as its extension regardless of actual original extension (that's three underscores).

The two-digit hash value consists of upper case alphanumeric characters.

This algorithm can cause name collisions only if files in a directory share the same first five alphanumeric characters. The probability of such a clash is 1/1300.

The name mangling (if enabled) allows a file to be copied between UNIX directories from Windows/DOS while retaining the long UNIX filename. UNIX files can be renamed to a new extension from Windows/DOS and will retain the same basename. Mangled names do not change between sessions.

Default: mangled names = yes

Example: mangled names = no

- **mangling char (S)**

This controls what character is used as the *"magic"* character in **name mangling**. The default is a '~' but this may interfere with some software. Use this option to set it to whatever you prefer.

Default: `mangling char = ~`

Example: `mangling char = ^`

- **mangled stack (G)**

This parameter controls the number of mangled names that should be cached in the Samba server **smbd**.

This stack is a list of recently mangled base names (extensions are only maintained if they are longer than 3 characters or contains upper case characters).

The larger this value, the more likely it is that mangled names can be successfully converted to correct long UNIX names. However, large stack sizes will slow most directory access. Smaller stacks save memory in the server (each stack element costs 256 bytes).

It is not possible to absolutely guarantee correct long file names, so be prepared for some surprises!

Default: `mangled stack = 50`

Example: `mangled stack = 100`

- **map archive (S)**

This controls whether the DOS archive attribute should be mapped to the UNIX owner execute bit. The DOS archive bit is set when a file has been modified since its last backup. One motivation for this option is to keep Samba/your PC from making any file it touches from becoming executable under UNIX. This can be quite annoying for shared source code, documents, etc...

Note that this requires the **"create mask"** parameter to be set such that owner execute bit is not masked out (i.e. it must include 100). See the parameter **"create mask"** for details.

Default: `map archive = yes`

Example: `map archive = no`

- **map hidden (S)**

This controls whether DOS style hidden files should be mapped to the UNIX world execute bit.

Note that this requires the **"create mask"** to be set such that the world execute bit is not masked out (i.e. it must include 001). See the parameter **"create mask"** for details.

Default: `map hidden = no`

Example: `map hidden = yes`

- **map system (S)**

This controls whether DOS style system files should be mapped to the UNIX group execute bit.

Note that this requires the **"create mask"** to be set such that the group execute bit is not masked out (i.e. it must include 010). See the parameter **"create mask"** for details.

Default: `map system = no`

Example: `map system = yes`

- **map to guest (G)**

This parameter is only useful in **security** modes other than "**security=share**"—i.e. user, server, and domain.

This parameter can take three different values, which tell **smbd** what to do with user login requests that don't match a valid UNIX user in some way.

The three settings are:

- **"Never"**— Means user login requests with an invalid password are rejected. This is the default.
- **"Bad User"**— Means user logins with an invalid password are rejected, unless the username does not exist, in which case it is treated as a guest login and mapped into the **guest account**.
- **"Bad Password"**— Means user logins with an invalid password are treated as a guest login and mapped into the **"guest account"**. Note that this can cause problems as it means that any user incorrectly typing their password will be silently logged on a **"guest"**—and will not kn—there will have been no message given to them that they got their password wrong. Helpdesk services will **hate** you if you set the **"map to guest"** parameter this way :-).

Note that this parameter is needed to set up **"Guest"** share services when using **security** modes other than share. This is because in these modes the name of the resource being requested is **not** sent to the server until after the server has successfully authenticated the client so the server cannot make authentication decisions at the correct time (connection to the share) for **"Guest"** shares.

For people familiar with the older Samba releases, this parameter maps to the old compile-time setting of the GUEST_SESSSETUP value in local.h.

Default: `map to guest = Never` **Example:** `map to guest = Bad User`

- **max connections (S)**

This option allows the number of simultaneous connections to a service to be limited. If **"max connections"** is greater than 0 then connections will be refused if this number of connections to the service are already open. A value of zero mean an unlimited number of connections may be made.

Record lock files are used to implement this feature. The lock files will be stored in the directory specified by the **"lock directory"** option.

Default: `max connections = 0`

Example: `max connections = 10`

- **max disk size (G)**

This option allows you to put an upper limit on the apparent size of disks. If you set this option to 100 then all shares will appear to be not larger than 100 MB in size.

Note that this option does not limit the amount of data you can put on the disk. In the above case you could still store much more than 100 MB on the disk, but if a client ever asks for the amount of free disk space or the total disk size then the result will be bounded by the amount specified in **"max disk size"**.

This option is primarily useful to work around bugs in some pieces of software that can't handle very large disks, particularly disks over 1GB in size.

A **"max disk size"** of 0 means no limit.

Default: `max disk size = 0`

Example: `max disk size = 1000`

- **max log size (G)**

This option (an integer in kilobytes) specifies the max size the log file should grow to. Samba periodically checks the size and if it is exceeded it will rename the file, adding a ".old" extension.

A size of 0 means no limit.

Default: `max log size = 5000`

Example: `max log size = 1000`

- **max mux (G)**

This option controls the maximum number of outstanding simultaneous SMB operations that samba tells the client it will allow. You should never need to set this parameter.

Default: `max mux = 50`

- **maxopenfiles (G)**

This parameter limits the maximum number of open files that one **smbd** file serving process may have open for a client at any one time. The default for this parameter is set very high (10,000) as Samba uses only one bit per unopened file.

The limit of the number of open files is usually set by the UNIX per-process file descriptor limit rather than this parameter so you should never need to touch this parameter.

Default: `max open files = 10000`

- **max packet (G)**

Synonym for "packetsize"(packetsize).

- **max ttl (G)**

This option tells **nmbd** what the default "time to live" of NetBIOS names should be (in seconds) when **nmbd** is requesting a name using either a broadcast packet or from a WINS server. You should never need to change this parameter. The default is 3 days.

Default: `max ttl = 259200`

- **max wins ttl (G)**

This option tells **nmbd** when acting as a WINS server (**wins support =true**) what the maximum "time to live" of NetBIOS names that **nmbd** will grant will be (in seconds). You should never need to change this parameter. The default is 6 days (518400 seconds).

See also the "**min wins ttl**" parameter.

Default: `max wins ttl = 518400`

- **max xmit (G)**

This option controls the maximum packet size that will be negotiated by Samba. The default is 65535, which is the maximum. In some cases you may find you get better performance with a smaller value. A value below 2048 is likely to cause problems.

Default: `max xmit = 65535`

Example: `max xmit = 8192`

- **message command (G)**

This specifies what command to run when the server receives a WinPopup style message.

This would normally be a command that would deliver the message somehow. How this is to be done is up to your imagination.

An example is:

```
message command = csh -c 'xedit %s;rm %s' &
```

This delivers the message using **xedit**, then removes it afterwards. *NOTE THAT IT IS VERY IMPORTANT THAT THIS COMMAND RETURN IMMEDIATELY.* That's why I have the '&' on the end. If it doesn't return immediately then your PCs may freeze when sending messages (they should recover after 30secs, hopefully).

All messages are delivered as the global guest user. The command takes the standard substitutions, although **%u** won't work (**%U** may be better in this case).

Apart from the standard substitutions, some additional ones apply. In particular:

- "%s" = the filename containing the message.
- "%t" = the destination that the message was sent to (probably the server name).
- "%f" = who the message is from.

You could make this command send mail, or whatever else takes your fancy. Please let us know of any really interesting ideas you have.

Here's a way of sending the messages as mail to root:

```
message command = /bin/mail -s 'message from %f on %m' root  
<  
%s; rm %s
```

If you don't have a message command then the message won't be delivered and Samba will tell the sender there was an error. Unfortunately WfWg totally ignores the error code and carries on regardless, saying that the message was delivered.

If you want to silently delete it then try:

```
"message command = rm %s".
```

Default: no message command

Example: message command = csh -c 'xedit %s;rm %s' &

- **min print space (S)**

This sets the minimum amount of free disk space that must be available before a user will be able to spool a print job. It is specified in kilobytes. The default is 0, which means a user can always spool a print job.

See also the **printing** parameter.

Default: min print space = 0

Example: min print space = 2000

- **min wins ttl (G)**

This option tells **nmbd** when acting as a WINS server (**wins support = true**) what the minimum "time to live" of NetBIOS names that **nmbd** will grant will be (in seconds). You should never need to change this parameter. The default is 6 hours (21600 seconds).

Default: min wins ttl = 21600

- **name resolve order (G)**

This option is used by the programs in the Samba suite to determine what naming services and in what order to resolve host names to IP addresses. The option takes a space separated string of different name resolution options.

The options are :`lmhosts`, `host`, `wins` and `bcast`. They cause names to be resolved as follows:

- **lmhosts** : Lookup an IP address in the Samba `lmhosts` file.
- **host** : Do a standard host name to IP address resolution, using the system `/etc/hosts`, NIS, or DNS lookups. This method of name resolution is operating system depended for instance on IRIX or Solaris this may be controlled by the `/etc/nsswitch.conf` file).
- **wins** : Query a name with the IP address listed in the **wins server** parameter. If no WINS server has been specified this method will be ignored.
- **bcast** : Do a broadcast on each of the known local interfaces listed in the **interfaces** parameter. This is the least reliable of the name resolution methods as it depends on the target host being on a locally connected subnet.

Default: `name resolve order = lmhosts host wins bcast`

Example: `name resolve order = lmhosts bcast host`

This will cause the local `lmhosts` file to be examined first, followed by a broadcast attempt, followed by a normal system hostname lookup.

- **netbios aliases (G)**

This is a list of NetBIOS names that **nmbd** will advertise as additional names by which the Samba server is known. This allows one machine to appear in browse lists under multiple names. If a machine is acting as a **browse server** or **logon server** none of these names will be advertised as either browse server or logon servers, only the primary name of the machine will be advertised with these capabilities.

See also "**netbios name**".

Default: `empty string (no additional names)`

Example: `netbios aliases = TEST TEST1 TEST2`

- **netbios name (G)**

This sets the NetBIOS name by which a Samba server is known. By default it is the same as the first component of the host's DNS name. If a machine is a **browse server** or **logon server** this name (or the first component of the hosts DNS name) will be the name that these services are advertised under.

See also "**netbios aliases**".

Default: `Machine DNS name.`

Example: `netbios name = MYNAME`

- **nis homedir (G)**

Get the home share server from a NIS map. For UNIX systems that use an automounter, the user's home directory will often be mounted on a workstation on demand from a remote server.

When the Samba logon server is not the actual home directory server, but is mounting the home directories via NFS then two network hops would be required to access the users home directory if the logon server told the client to use itself as the SMB server for home directories (one over SMB and one over NFS). This can be very slow.

This option allows Samba to return the home share as being on a different server to the logon server and as long as a Samba daemon is running on the home directory server, it will be mounted on the Samba client directly from the directory server. When Samba is

returning the home share to the client, it will consult the NIS map specified in "**homedir map**" and return the server listed there.

Note that for this option to work there must be a working NIS system and the Samba server with this option must also be a **logon server**.

Default: `nis homedir = false`

Example: `nis homedir = true`

- **nt pipe support (G)**

This boolean parameter controls whether **smbd** will allow Windows NT clients to connect to the NT SMB specific `IPC$` pipes. This is a developer debugging option and can be left alone.

Default: `nt pipe support = yes`

- **nt smb support (G)**

This boolean parameter controls whether **smbd** will negotiate NT specific SMB support with Windows NT clients. Although this is a developer debugging option and should be left alone, benchmarking has discovered that Windows NT clients give faster performance with this option set to "`no`". This is still being investigated. If this option is set to "`no`" then Samba offers exactly the same SMB calls that versions prior to Samba2.0 offered. This information may be of use if any users are having problems with NT SMB support.

Default: `nt support = yes`

- **null passwords (G)**

Allow or disallow client access to accounts that have null passwords.

See also **smbpasswd (5)**.

Default: `null passwords = no`

Example: `null passwords = yes`

- **ole locking compatibility (G)**

This parameter allows an administrator to turn off the byte range lock manipulation that is done within Samba to give compatibility for OLE applications. Windows OLE applications use byte range locking as a form of inter-process communication, by locking ranges of bytes around the 2^{32} region of a file range. This can cause certain UNIX lock managers to crash or otherwise cause problems. Setting this parameter to "`no`" means you trust your UNIX lock manager to handle such cases correctly.

Default: `ole locking compatibility = yes`

Example: `ole locking compatibility = no`

- **only guest (S)**

A synonym for "**guest only**".

- **only user (S)**

This is a boolean option that controls whether connections with usernames not in the **user=** list will be allowed. By default this option is disabled so a client can supply a username to be used by the server.

Note that this also means Samba won't try to deduce usernames from the service name. This can be annoying for the **[homes]** section. To get around this you could use **"user =**

%S" which means your **"user"** list will be just the service name, which for home directories is the name of the user.

See also the **user** parameter.

Default: `only user = False`

Example: `only user = True`

- **oplocks (S)**

This boolean option tells `smbd` whether to issue oplocks (opportunistic locks) to file open requests on this share. The oplock code can dramatically (approx. 30% or more) improve the speed of access to files on Samba servers. It allows the clients to aggressively cache files locally and you may want to disable this option for unreliable network environments (it is turned on by default in Windows NT Servers). For more information see the file `Speed.txt` in the Samba docs/ directory.

Oplocks may be selectively turned off on certain files on a per share basis. See the "veto oplock files" parameter. On some systems oplocks are recognized by the underlying operating system. This allows data synchronization between all access to oplocked files, whether it be via Samba or NFS or a local UNIX process. See the **kernel oplocks** parameter for details.

Default: `oplocks = True`

Example: `oplocks = False`

- **os level (G)**

This integer value controls what level Samba advertises itself as for browse elections. The value of this parameter determines whether `nmbd` has a chance of becoming a local master browser for the **WORKGROUP** in the local broadcast area. Setting this to zero will cause `nmbd` to always lose elections to Windows machines. See `BROWSING.txt` in the Samba docs/ directory for details.

Default: `os level = 32`

Example: `os level = 65 ; This will win against any NT Server`

- **packet size (G)**

This is a deprecated parameter that has no effect on the current Samba code. It is left in the parameter list to prevent breaking old `smb.conf` files.

- **panic action (G)**

This is a Samba developer option that allows a system command to be called when either `smbd` or `nmbd` crashes. This is usually used to draw attention to the fact that a problem occurred.

Default: `panic action = <empty string>`

- **passwd chat (G)**

This string controls the "chat" conversation that takes place between `smbd` and the local password changing program to change the users password. The string describes a sequence of response-receive pairs that `smbd` uses to determine what to send to the `passwd` program and what to expect back. If the expected output is not received then the password is not changed.

This chat sequence is often quite site specific, depending on what local methods are used for password control (such as NIS etc).

The string can contain the macros "%o" and "%n" which are substituted for the old and new passwords respectively. It can also contain the standard macros "\n", "\r", "\t" and "\s" to give line-feed, carriage-return, tab and space.

The string can also contain a '*' which matches any sequence of characters.

Double quotes can be used to collect strings with spaces in them into a single string. If the send string in any part of the chat sequence is a fullstop "." then no string is sent.

Similarly, if the expect string is a fullstop then no string is expected.

Note that if the "unix password sync" parameter is set to true, then this sequence is called *AS ROOT* when the SMB password in the smbpasswd file is being changed, without access to the old password cleartext. In this case the old password cleartext is set to "" (the empty string).

See also "unix password sync", "passwd program" and "passwd chat debug".

Example:

```
passwd chat = "*Enter OLD password*" %o\n "*Enter NEW  
password*" %n\n  
"*Reenter NEW password*" %n\n "*Password  
changed*"
```

Default:

```
passwd chat = *old*password* %o\n *new*password* %n\n  
*new*password* %n\n *changed*
```

- **passwd chat debug (G)**

This boolean specifies if the passwd chat script parameter is run in "debug" mode. In this mode the strings passed to and received from the passwd chat are printed in the **smbd** log with a "debug level" of 100. This is a dangerous option as it will allow plaintext passwords to be seen in the **smbd** log. It is available to help Samba admins debug their "passwd chat" scripts when calling the "passwd program" and should be turned off after this has been done. This parameter is off by default.

See also "passwd chat", "passwd program".

Example: `passwd chat debug = True`

Default: `passwd chat debug = False`

- **passwd program (G)**

The name of a program that can be used to set UNIX user passwords. Any occurrences of %u will be replaced with the user name. The user name is checked for existence before calling the password changing program.

Also note that many passwd programs insist in "reasonable" passwords, such as a minimum length, or the inclusion of mixed case chars and digits. This can pose a problem as some clients (such as Windows for Workgroups) uppercase the password before sending it.

Note that if the "unix password sync" parameter is set to "True" then this program is called *AS ROOT* before the SMB password in the **smbpasswd** file is changed. If this UNIX password change fails, then **smbd** will fail to change the SMB password also (this is by design).

If the "unix password sync" parameter is set this parameter *MUST USE ABSOLUTE PATHS* for ALL programs called, and must be examined for security implications. Note that by default "unix password sync" is set to "False".

See also "unix password sync".

Default: `passwd program = /bin/passwd`

Example: `passwd program = /sbin/passwd %u`

- **password level (G)**

Some client/server combinations have difficulty with mixed-case passwords. One offending client is Windows for Workgroups, which for some reason forces passwords to upper case when using the LANMAN1 protocol, but leaves them alone when using COREPLUS!

This parameter defines the maximum number of characters that may be upper case in passwords.

For example, say the password given was "FRED". If **password level** is set to 1, the following combinations would be tried if "FRED" failed:

"Fred", "fred", "fRed", "frEd", "freD"

If **password level** was set to 2, the following combinations would also be tried:

"FRed", "FrEd", "FreD", "fRED", "fReD", "frED", ..

And so on.

The higher value this parameter is set to the more likely it is that a mixed case password will be matched against a single case password. However, you should be aware that use of this parameter reduces security and increases the time taken to process a new connection.

A value of zero will cause only two attempts to be made—the password as is and the password in all-lower case.

Default: `password level = 0`

Example: `password level = 4`

- **password server (G)**

By specifying the name of another SMB server (such as a WinNT box) with this option, and using "**security = domain**" or "**security = server**" you can get Samba to do all its username/password validation via a remote server.

This options sets the name of the password server to use. It must be a NetBIOS name, so if the machine's NetBIOS name is different from its internet name then you may have to add its NetBIOS name to the lmhosts file which is stored in the same directory as the **smb.conf** file.

The name of the password server is looked up using the parameter "**name resolve order=**" and so may resolved by any method and order described in that parameter.

The password server much be a machine capable of using the "LM1.2X002" or the "LM NT 0.12" protocol, and it must be in user level security mode.

NOTE: Using a password server means your UNIX box (running Samba) is only as secure as your password server. *DO NOT CHOOSE A PASSWORD SERVER THAT YOU DON'T COMPLETELY TRUST.*

Never point a Samba server at itself for password serving. This will cause a loop and could lock up your Samba server!

The name of the password server takes the standard substitutions, but probably the only useful one is **%m**, which means the Samba server will use the incoming client as the password server. If you use this then you better trust your clients, and you better restrict them with hosts allow!

If the "**security**" parameter is set to "**domain**", then the list of machines in this option must be a list of Primary or Backup Domain controllers for the **Domain**, as the Samba server is cryptographically in that domain, and will use cryptographically authenticated RPC calls to authenticate the user logging on. The advantage of using "**security=domain**" is

that if you list several hosts in the "**password server**" option then **smbd** will try each in turn till it finds one that responds. This is useful in case your primary server goes down. If the "**security**" parameter is set to "**server**", then there are different restrictions that "**security=domain**" doesn't suffer from:

- You may list several password servers in the "**password server**" parameter, however if an **smbd** makes a connection to a password server, and then the password server fails, no more users will be able to be authenticated from this **smbd**. This is a restriction of the SMB/CIFS protocol when in "**security=server**" mode and cannot be fixed in Samba.
- If you are using a Windows NT server as your password server then you will have to ensure that your users are able to login from the Samba server, as when in "**security= server**" mode the network logon will appear to come from there rather than from the users workstation.

See also the "**security**" parameter.

Default: `password server = <empty string>`

Example: `password server = NT-PDC, NT-BDC1, NT-BDC2`

- **path (S)**

This parameter specifies a directory to which the user of the service is to be given access. In the case of printable services, this is where print data will spool prior to being submitted to the host for printing.

For a printable service offering guest access, the service should be readonly and the path should be world-writable and have the sticky bit set. This is not mandatory of course, but you probably won't get the results you expect if you do otherwise.

Any occurrences of **%u** in the path will be replaced with the UNIX username that the client is using on this connection. Any occurrences of **%m** will be replaced by the NetBIOS name of the machine they are connecting from. These replacements are very useful for setting up pseudo home directories for users.

Note that this path will be based on "**root dir**" if one was specified.

Default: `none`

Example: `path = /home/fred`

- **postexec (S)**

This option specifies a command to be run whenever the service is disconnected. It takes the usual substitutions. The command may be run as the root on some systems.

An interesting example may be do unmount server resources:

```
postexec = /etc/umount /cdrom
```

See also **preexec**.

Default: `none` (no command executed)

Example: `postexec = echo "%u disconnected from %S from %m (%I)" >> /tmp/log`

- **postscript (S)**

This parameter forces a printer to interpret the print files as postscript. This is done by adding a **%!** to the start of print output.

This is most useful when you have lots of PCs that persist in putting a control-D at the start of print jobs, which then confuses your printer.

Default: `postscript = False`

Example: `postscript = True`

- **preexec (S)**

This option specifies a command to be run whenever the service is connected to. It takes the usual substitutions.

An interesting example is to send the users a welcome message every time they log in. Maybe a message of the day? Here is an example:

```
preexec = csh -c 'echo \"Welcome to %S!\" |  
/usr/local/samba/bin/smbclient -M %m -I %I' &
```

Of course, this could get annoying after a while :-)

See also **postexec**.

Default: `none` (no command executed)

Example: `preexec = echo \"%u connected to %S from %m (%I)\" >>
/tmp/log`

- **preferred master (G)**

This boolean parameter controls if **nmbd** is a preferred master browser for its workgroup. If this is set to true, on startup, **nmbd** will force an election, and it will have a slight advantage in winning the election. It is recommended that this parameter is used in conjunction with "**domain master = yes**", so that **nmbd** can guarantee becoming a domain master. Indeed the default ("auto") enables "preferred master" if Samba is configured as the domain master browser.

Use this option with caution, because if there are several hosts (whether Samba servers, Windows 95 or NT) that are preferred master browsers on the same subnet, they will each periodically and continuously attempt to become the local master browser. This will result in unnecessary broadcast traffic and reduced browsing capabilities.

See also **os level**.

Default: `preferred master = auto`

Example: `preferred master = yes`

- **prefered master (G)**

Synonym for "**preferred master**" for people who cannot spell :-).

- **preload**

Synonym for "**auto services**".

- **preserve case (S)**

This controls if new filenames are created with the case that the client passes, or if they are forced to be the "`default`" case.

Default: `preserve case = yes`

See the section on "**NAME MANGLING**" for a fuller discussion.

- **print command (S)**

After a print job has finished spooling to a service, this command will be used via a `system()` call to process the spool file. Typically the command specified will submit the spool file to the host's printing subsystem, but there is no requirement that this be the

case. The server will not remove the spool file, so whatever command you specify should remove the spool file when it has been processed, otherwise you will need to manually remove old spool files.

The print command is simply a text string. It will be used verbatim, with two exceptions: All occurrences of "%s" will be replaced by the appropriate spool file name, and all occurrences of "%p" will be replaced by the appropriate printer name. The spool file name is generated automatically by the server, the printer name is discussed below.

The full path name will be used for the filename if "%s" is not preceded by a '/'. If you don't like this (it can stuff up some lpq output) then use "%f" instead. Any occurrences of "%f" get replaced by the spool filename without the full path at the front.

The print command *MUST* contain at least one occurrence of "%s" or "%f"—the "%p" is optional. At the time a job is submitted, if no printer name is supplied the "%p" will be silently removed from the printer command.

If specified in the "[global]" section, the print command given will be used for any printable service that does not have its own print command specified.

If there is neither a specified print command for a printable service nor a global print command, spool files will be created but not processed and (most importantly) not removed.

Note that printing may fail on some UNIXs from the "nobody" account. If this happens then create an alternative guest account that can print and set the "guest account" in the "[global]" section.

You can form quite complex print commands by realizing that they are just passed to a shell. For example the following will log a print job, print the file, then remove it. Note that ';' is the usual separator for command in shell scripts.

```
print command = echo Printing %s >> /tmp/print.log; lpr -P
%p
%s; rm %s
```

You may have to vary this command considerably depending on how you normally print files on your system. The default for the parameter varies depending on the setting of the "printing=" parameter.

Default: For "printing=" BSD, AIX, QNX, LPRNG or PLP : `print command = lpr -r -P%p %s`

For "printing=" SYS or HPUX : `print command = lp -c -d%p %s; rm %s`

For "printing=" SOFTQ : `print command = lp -d%p -s %s; rm %s`

Example: `print command = /usr/local/samba/bin/myprint-script %p %s`

- **print ok (S)**

Synonym for **printable**.

- **printable (S)**

If this parameter is "yes", then clients may open, write to and submit spool files on the directory specified for the service.

Note that a printable service will ALWAYS allow writing to the service path (user privileges permitting) via the spooling of print data. The "read only" parameter controls only non-printing access to the resource.

Default: `printable = no`

Example: `printable = yes`

- **printcap (G)**

Synonym for **printcapname**.

- **printcap name (G)**

This parameter may be used to override the compiled-in default printcap name used by the server (usually `/etc/printcap`). See the discussion of the **[printers]** section above for reasons why you might want to do this.

On System V systems that use **lpstat** to list available printers you can use "`printcap name = lpstat`" to automatically obtain lists of available printers. This is the default for systems that define SYSV at configure time in Samba (this includes most System V based systems). If "**printcap name**" is set to **lpstat** on these systems then Samba will launch "`lpstat -v`" and attempt to parse the output to obtain a printer list.

A minimal printcap file would look something like this:

```
print1|My Printer 1
print2|My Printer 2
print3|My Printer 3
print4|My Printer 4
print5|My Printer 5
```

where the `|` separates aliases of a printer. The fact that the second alias has a space in it gives a hint to Samba that it's a comment.

Note

Under AIX the default printcap name is `/etc/qconfig`. Samba will assume the file is in AIX `"qconfig"` format if the string `"qconfig"` appears in the printcap filename.

Default: `printcap name = /etc/printcap`

Example: `printcap name = /etc/myprintcap`

- **printer (S)**

This parameter specifies the name of the printer to which print jobs spooled through a printable service will be sent.

If specified in the **[global]** section, the printer name given will be used for any printable service that does not have its own printer name specified.

Default: none (but may be `"lp"` on many systems)

Example: `printer name = laserwriter`

- `printer driver (S)`

This option allows you to control the string that clients receive when they ask the server for the printer driver associated with a printer. If you are using Windows95 or WindowsNT then you can use this to automate the setup of printers on your system.

You need to set this parameter to the exact string (case sensitive) that describes the appropriate printer driver for your system. If you don't know the exact string to use then

you should first try with no **"printer driver"** option set and the client will give you a list of printer drivers. The appropriate strings are shown in a scrollbox after you have chosen the printer manufacturer.

See also **"printer driver file"**.

Example: printer driver = HP LaserJet 4L

- **printer driver file (G)**

This parameter tells Samba where the printer driver definition file, used when serving drivers to Windows 95 clients, is to be found. If this is not set, the default is :

```
SAMBA_INSTALL_DIRECTORY/lib/printers.def
```

This file is created from Windows 95 `"msprint.def"` files found on the Windows 95 client system. For more details on setting up serving of printer drivers to Windows 95 clients, see the documentation file in the docs/ directory, `PRINTER_DRIVER.txt`.

Default: None (set in compile).

Example:

```
printer driver file =  
/usr/local/samba/printers/drivers.def
```

See also **"printer driver location"**.

- **printer driver location (S)**

This parameter tells clients of a particular printer share where to find the printer driver files for the automatic installation of drivers for Windows 95 machines. If Samba is set up to serve printer drivers to Windows 95 machines, this should be set to

```
\\MACHINE\aPRINTER$
```

Where MACHINE is the NetBIOS name of your Samba server, and PRINTER\$ is a share you set up for serving printer driver files. For more details on setting this up see the documentation file in the docs/ directory, `PRINTER_DRIVER.txt`.

Default: None

Example: printer driver location = \\MACHINE\PRINTER\$

See also **"printer driver file"**.

- **printer name (S)**

Synonym for **printer**.

- **printing (S)**

This parameters controls how printer status information is interpreted on your system, and also affects the default values for the **"print command"**, **"lpq command"** **"lppause command"**, **"lpresume command"**, and **"lprm command"**.

Currently eight printing styles are supported. They are **"print-ing=BSD"**, **"printing=AIX"**, **"printing=LPRNG"**, **"printing=PLP"**, **"printing=SYSV"**, **"printing=HPUX"**, **"printing=QNX"** and **"printing=SOFTQ"**.

To see what the defaults are for the other print commands when using these three options use the **"testparm"** program.

This option can be set on a per printer basis

See also the discussion in the **[printers]** section.

- **protocol (G)**

The value of the parameter (a string) is the highest protocol level that will be supported by the server.

Possible values are :

- CORE: Earliest version. No concept of user names.
- COREPLUS: Slight improvements on CORE for efficiency.
- LANMAN1: First "modern" version of the protocol. Long filename support.
- LANMAN2: Updates to Lanman1 protocol.
- NT1: Current up to date version of the protocol. Used by Windows NT. Known as CIFS.

Normally this option should not be set as the automatic negotiation phase in the SMB protocol takes care of choosing the appropriate protocol.

Default: `protocol = NT1`

Example: `protocol = LANMAN1`

- **public (S)**

Synonym for "guest ok".

- **queuepause command (S)**

This parameter specifies the command to be executed on the server host in order to pause the printerqueue.

This command should be a program or script which takes a printer name as its only parameter and stops the printerqueue, such that no longer jobs are submitted to the printer.

This command is not supported by Windows for Workgroups, but can be issued from the Printer's window under Windows 95 & NT.

If a "%p" is given then the printername is put in its place. Otherwise it is placed at the end of the command.

Note that it is good practice to include the absolute path in the command as the PATH may not be available to the server.

Default: depends on the setting of "printing ="

Example: `queuepause command = disable %p`

- **queueresume command (S)**

This parameter specifies the command to be executed on the server host in order to resume the printerqueue. It is the command to undo the behavior that is caused by the previous parameter ("queuepause command").

This command should be a program or script which takes a printer name as its only parameter and resumes the printerqueue, such that queued jobs are resubmitted to the printer.

This command is not supported by Windows for Workgroups, but can be issued from the Printer's window under Windows 95 & NT.

If a "%p" is given then the printername is put in its place. Otherwise it is placed at the end of the command.

Note that it is good practice to include the absolute path in the command as the PATH may not be available to the server.

Default: depends on the setting of "printing ="

Example: `queuepause command = enable %p`

- **read bmpx (G)**

This boolean parameter controls whether **smbd** will support the "Read Block Multiplex" SMB. This is now rarely used and defaults to off. You should never need to set this parameter.

Default: `read bmpx = No`

- **read list (S)**

This is a list of users that are given read-only access to a service. If the connecting user is in this list then they will not be given write access, no matter what the "**read only**" option is set to. The list can include group names using the syntax described in the "**invalid users**" parameter.

See also the "**write list**" parameter and the "**invalid users**" parameter.

Default: `read list = <empty string>`

Example: `read list = mary, @students`

- **read only (S)**

Note that this is an inverted synonym for "**writeable**" and "**write ok**".

See also "**writeable**" and "**write ok**".

- **read prediction (G)**

Note

This code is currently disabled in Samba2.0 and may be removed at a later date. Hence this parameter has no effect.

This options enables or disables the read prediction code used to speed up reads from the server. When enabled the server will try to pre-read data from the last accessed file that was opened read-only while waiting for packets.

Default: `read prediction = False`

- **read raw (G)**

This parameter controls whether or not the server will support the raw read SMB requests when transferring data to clients.

If enabled, raw reads allow reads of 65535 bytes in one packet. This typically provides a major performance benefit.

However, some clients either negotiate the allowable block size incorrectly or are incapable of supporting larger block sizes, and for these clients you may need to disable raw reads.

In general this parameter should be viewed as a system tuning tool and left severely alone. See also "**write raw**".

Default: `read raw = yes`

- **read size (G)**

The option **"read size"** affects the overlap of disk reads/writes with network reads/writes. If the amount of data being transferred in several of the SMB commands (currently SMBwrite, SMBwriteX and SMBreadbrow) is larger than this value then the server begins writing the data before it has received the whole packet from the network, or in the case of SMBreadbrow, it begins writing to the network before all the data has been read from disk.

This overlapping works best when the speeds of disk and network access are similar, having very little effect when the speed of one is much greater than the other.

The default value is 2048, but very little experimentation has been done yet to determine the optimal value, and it is likely that the best value will vary greatly between systems anyway. A value over 65536 is pointless and will cause you to allocate memory unnecessarily.

Default: `read size = 2048`

Example: `read size = 8192`

- **remote announce (G)**

This option allows you to setup **nmbd** to periodically announce itself to arbitrary IP addresses with an arbitrary workgroup name.

This is useful if you want your Samba server to appear in a remote workgroup for which the normal browse propagation rules don't work. The remote workgroup can be anywhere that you can send IP packets to.

For example:

```
remote announce = 192.168.2.255/SERVERS 192.168.4.255/STAFF
```

the above line would cause **nmbd** to announce itself to the two given IP addresses using the given workgroup names. If you leave out the workgroup name then the one given in the **"workgroup"** parameter is used instead.

The IP addresses you choose would normally be the broadcast addresses of the remote networks, but can also be the IP addresses of known browse masters if your network config is that stable.

See the documentation file BROWSING.txt in the docs/ directory.

Default: `remote announce = <empty string>`

Example: `remote announce = 192.168.2.255/SERVERS
192.168.4.255/ STAFF`

- **remote browse sync (G)**

This option allows you to setup **nmbd** to periodically request synchronization of browse lists with the master browser of a samba server that is on a remote segment. This option will allow you to gain browse lists for multiple workgroups across routed networks. This is done in a manner that does not work with any non-samba servers.

This is useful if you want your Samba server and all local clients to appear in a remote workgroup for which the normal browse propagation rules don't work. The remote workgroup can be anywhere that you can send IP packets to.

For example:

```
remote browse sync = 192.168.2.255 192.168.4.255
```

the above line would cause **nmbd** to request the master browser on the specified subnets or addresses to synchronize their browse lists with the local server.

The IP addresses you choose would normally be the broadcast addresses of the remote networks, but can also be the IP addresses of known browse masters if your network config is that stable. If a machine IP address is given Samba makes NO attempt to

validate that the remote machine is available, is listening, nor that it is in fact the browse master on it's segment.

Default: `remote browse sync = <empty string>`

Example: `remote browse sync = 192.168.2.255 192. 168.4.255`

- **revalidate (S)**

Note that this option only works with "**security=share**" and will be ignored if this is not the case.

This option controls whether Samba will allow a previously validated username/password pair to be used to attach to a share. Thus if you connect to `\\server\share1` then to `\\server\share2` it won't automatically allow the client to request connection to the second share as the same username as the first without a password.

If "**revalidate**" is "**True**" then the client will be denied automatic access as the same username.

Default: `revalidate = False`

Example: `revalidate = True`

- **root (G)**

Synonym for "**root directory**".

- **root dir (G)**

Synonym for "**root directory**".

- **root directory (G)**

The server will "`chroot ()`" (i.e. Change it's root directory) to this directory on startup. This is not strictly necessary for secure operation. Even without it the server will deny access to files not in one of the service entries. It may also check for, and deny access to, soft links to other parts of the filesystem, or attempts to use "`..`" in file names to access other directories (depending on the setting of the "**wide links**" parameter).

Adding a "**root directory**" entry other than "/" adds an extra level of security, but at a price. It absolutely ensures that no access is given to files not in the sub-tree specified in the "**root directory**" option, *including* some files needed for complete operation of the server. To maintain full operability of the server you will need to mirror some system files into the "**root directory**" tree. In particular you will need to mirror `/etc/passwd` (or a subset of it), and any binaries or configuration files needed for printing (if required). The set of files that must be mirrored is operating system dependent.

Default: `root directory = /`

Example: `root directory = /homes/smb`

- **root postexec (S)**

This is the same as the "**postexec**" parameter except that the command is run as root. This is useful for unmounting filesystems (such as cdroms) after a connection is closed. See also "**postexec**".

- **root preexec (S)**

This is the same as the "**preexec**" parameter except that the command is run as root. This is useful for mounting filesystems (such as cdroms) before a connection is finalized. See also "**preexec**".

- **security (G)**

This option affects how clients respond to Samba and is one of the most important settings in the **smb.conf** file.

The option sets the `"security mode bit"` in replies to protocol negotiations with **smbd** to turn share level security on or off. Clients decide based on this bit whether (and how) to transfer user and password information to the server.

The default is `"security=user"`, as this is the most common setting needed when talking to Windows 98 and Windows NT.

The alternatives are `"security = share"`, `"security = server"` or `"security=domain"`.

*****NOTE THAT THIS DEFAULT IS DIFFERENT IN SAMBA2.0 THAN FOR PREVIOUS VERSIONS OF SAMBA *****.

In previous versions of Samba the default was `"security=share"` mainly because that was the only option at one stage.

There is a bug in WfWg that has relevance to this setting. When in user or server level security a WfWg client will totally ignore the password you type in the "connect drive" dialog box. This makes it very difficult (if not impossible) to connect to a Samba service as anyone except the user that you are logged into WfWg as.

If your PCs use usernames that are the same as their usernames on the UNIX machine then you will want to use `"security = user"`. If you mostly use usernames that don't exist on the UNIX box then use `"security = share"`.

You should also use `security=share` if you want to mainly setup shares without a password (guest shares). This is commonly used for a shared printer server. It is more difficult to setup guest shares with `security=user`, see the `"map to guest"` parameter for details.

It is possible to use **smbd** in a *"hybrid mode"* where it offers both user and share level security under different **NetBIOS aliases**. See the **NetBIOS aliases** and the **include** parameters for more information.

The different settings will now be explained.

- **"security=share"** When clients connect to a share level security server then need not log onto the server with a valid username and password before attempting to connect to a shared resource (although modern clients such as Windows 95/98 and Windows NT will send a logon request with a username but no password when talking to a `security=share` server). Instead, the clients send authentication information (passwords) on a per-share basis, at the time they attempt to connect to that share.

Note that **smbd** **ALWAYS** uses a valid UNIX user to act on behalf of the client, even in `"security=share"` level security.

As clients are not required to send a username to the server in share level security, **smbd** uses several techniques to determine the correct UNIX user to use on behalf of the client.

A list of possible UNIX usernames to match with the given client password is constructed using the following methods:

- If the **"guest only"** parameter is set, then all the other stages are missed and only the **"guest account"** username is checked.
- If a username is sent with the share connection request, then this username (after mapping—see **"username map"**), is added as a potential username.
- If the client did a previous *"logon"* request (the SessionSetup SMB call) then the username sent in this SMB will be added as a potential username.
- The name of the service the client requested is added as a potential username.
- The NetBIOS name of the client is added to the list as a potential username.
- Any users on the **"user"** list are added as potential usernames.

If the **"guest only"** parameter is not set, then this list is then tried with the supplied password. The first user for whom the password matches will be used as the UNIX user. If the **"guest only"** parameter is set, or no username can be determined then if the share is marked as available to the **"guest account"**, then this guest user will be used, otherwise access is denied.

Note that it can be *very* confusing in share-level security as to which UNIX username will eventually be used in granting access.

See also the section **"NOTE ABOUT USERNAME/PASSWORD VALIDATION"**.

- **"security=user"**

This is the default security setting in Samba2.0. With user-level security a client must first `log-on` with a valid username and password (which can be mapped using the **"username map"** parameter). Encrypted passwords (see the **"encrypted passwords"** parameter) can also be used in this security mode. Parameters such as **"user"** and **"guest only"**, if set are then applied and may change the UNIX user to use on this connection, but only after the user has been successfully authenticated.

Note that the name of the resource being requested is *not* sent to the server until after the server has successfully authenticated the client. This is why guest shares don't work in user level security without allowing the server to automatically map unknown users into the **"guest account"**. See the **"map to guest"** parameter for details on doing this.

See also the section **"NOTE ABOUT USERNAME/PASSWORD VALIDATION"**.

- **"security=server"**

In this mode Samba will try to validate the username/password by passing it to another SMB server, such as an NT box. If this fails it will revert to **"security = user"**, but note that if encrypted passwords have been negotiated then Samba cannot revert back to checking the UNIX password file, it must have a valid `smbpasswd` file to check users against. See the documentation file in the docs/ directory ENCRYPTION.txt for details on how to set this up.

Note that from the clients point of view **"security=server"** is the same as **"security=user"**. It only affects how the server deals with the authentication, it does not in any way affect what the client sees.

Note that the name of the resource being requested is *not* sent to the server until after the server has successfully authenticated the client. This is why guest shares don't work in server level security without allowing the server to automatically map unknown users into the **"guest account"**. See the **"map to guest"** parameter for details on doing this.

See also the section **"NOTE ABOUT USERNAME/PASSWORD VALIDATION"**.

See also the **"password server"** parameter. and the **"encrypted passwords"** parameter.

- **"security=domain"**

This mode will only work correctly if `smbpasswd` has been used to add this machine into a Windows NT Domain. It expects the **"encrypted passwords"** parameter to be set to `true`. In this mode Samba will try to validate the username/password by passing it to a Windows NT Primary or Backup Domain Controller, in exactly the same way that a Windows NT Server would do.

Note that a valid UNIX user must still exist as well as the account on the Domain Controller to allow Samba to have a valid UNIX account to map file access to.

Note that from the clients point of view **"security=domain"** is the same as **"security=user"**. It only affects how the server deals with the authentication, it does not in any way affect what the client sees.

Note that the name of the resource being requested is *not* sent to the server until after the server has successfully authenticated the client. This is why guest shares don't work

in domain level security without allowing the server to automatically map unknown users into the **"guest account"**. See the **"map to guest"** parameter for details on doing this.

e,(BUG:) There is currently a bug in the implementation of **"security=domain"** with respect to multi-byte character set usernames. The communication with a Domain Controller must be done in UNICODE and Samba currently does not widen multi-byte user names to UNICODE correctly, thus a multi-byte username will not be recognized correctly at the Domain Controller. This issue will be addressed in a future release.

See also the section **"NOTE ABOUT USERNAME/PASSWORD VALIDATION"**.

See also the **"password server"** parameter. and the **"encrypted passwords"** parameter.

Default: `security = USER`

Example: `security = DOMAIN`

- **server string (G)**

This controls what string will show up in the printer comment box in print manager and next to the IPC connection in `"net view"`. It can be any string that you wish to show to your users.

It also sets what will appear in browse lists next to the machine name.

A `"%v"` will be replaced with the Samba version number.

A `"%h"` will be replaced with the hostname.

Default: `server string = Samba %v`

Example: `server string = University of GNUs Samba Server`

- **set directory (S)**

If `"set directory = no"`, then users of the service may not use the `setdir` command to change directory.

The `setdir` command is only implemented in the Digital Pathworks client. See the Pathworks documentation for details.

Default: `set directory = no`

Example: `set directory = yes`

- **share modes (S)**

This enables or disables the honoring of the `"share modes"` during a file open. These modes are used by clients to gain exclusive read or write access to a file.

These open modes are not directly supported by UNIX, so they are simulated using shared memory, or lock files if your UNIX doesn't support shared memory (almost all do).

The share modes that are enabled by this option are `DENY_DOS`, `DENY_ALL`, `DENY_READ`, `DENY_WRITE`, `DENY_NONE` and `DENY_FCB`.

This option gives full share compatibility and enabled by default.

You should **NEVER** turn this parameter off as many Windows applications will break if you do so.

Default: `share modes = yes`

- **shared mem size (G)**

It specifies the size of the shared memory (in bytes) to use between **smbd** processes. This parameter defaults to one megabyte of shared memory. It is possible that if you have a large server with many files open simultaneously that you may need to increase this parameter. Signs that this parameter is set too low are users reporting strange problems trying to save files (locking errors) and error messages in the `smbd` log looking like `"ERROR smb_shm_alloc : alloc of XX bytes failed"`.

Default: `shared mem size = 1048576`

Example: `shared mem size = 5242880 ; Set to 5mb for a large number of files.`

- **short preserve case (G)**

This boolean parameter controls if new files which conform to 8.3 syntax, that is all in upper case and of suitable length, are created upper case, or if they are forced to be the "default" case. This option can be use with "**preserve case =yes**" to permit long filenames to retain their case, while short names are lowered. Default Yes.

See the section on **NAME MANGLING**.

Default: `short preserve case = yes`

- **smb passwd file (G)**

This option sets the path to the encrypted smbpasswd file. By default the path to the smbpasswd file is compiled into Samba.

Default: `smb passwd file= <compiled default>`

Example: `smb passwd file = /usr/samba/private/smbpasswd`

- **smbrun (G)**

This sets the full path to the **smbrun** binary. This defaults to the value in the Makefile.

You must get this path right for many services to work correctly.

You should not need to change this parameter so long as Samba is installed correctly.

Default: `smbrun=<compiled default>`

Example: `smbrun = /usr/local/samba/bin/smbrun`

- **socket address (G)**

This option allows you to control what address Samba will listen for connections on. This is used to support multiple virtual interfaces on the one server, each with a different configuration.

By default samba will accept connections on any address.

Example: `socket address = 192.168.2.20`

- **socket options (G)**

This option allows you to set socket options to be used when talking with the client.

Socket options are controls on the networking layer of the operating systems which allow the connection to be tuned.

This option will typically be used to tune your Samba server for optimal performance for your local network. There is no way that Samba can know what the optimal parameters are for your net, so you must experiment and choose them yourself. We strongly suggest you read the appropriate documentation for your operating system first (perhaps "**man setsockopt**" will help).

You may find that on some systems Samba will say "Unknown socket option" when you supply an option. This means you either incorrectly typed it or you need to add an include file to `includes.h` for your OS. If the latter is the case please send the patch to samba-bugs@samba.org.

Any of the supported socket options may be combined in any way you like, as long as your OS allows it.

This is the list of socket options currently settable using this option:

- `SO_KEEPALIVE`
- `SO_REUSEADDR`

- SO_BROADCAST
- TCP_NODELAY
- IPTOS_LOWDELAY
- IPTOS_THROUGHPUT
- SO_SNDBUF *
- SO_RCVBUF *
- SO_SNDLOWAT *
- SO_RCVLOWAT *

Those marked with a * take an integer argument. The others can optionally take a 1 or 0 argument to enable or disable the option, by default they will be enabled if you don't specify 1 or 0.

To specify an argument use the syntax `SOME_OPTION=VALUE` for example `SO_SNDBUF=8192`. Note that you must not have any spaces before or after the = sign.

If you are on a local network then a sensible option might be

```
socket options = IPTOS_LOWDELAY
```

If you have a local network then you could try:

```
socket options = IPTOS_LOWDELAY TCP_NODELAY
```

If you are on a wide area network then perhaps try setting `IPTOS_THROUGHPUT`. Note that several of the options may cause your Samba server to fail completely. Use these options with caution!

Default: `socket options = TCP_NODELAY`

Example: `socket options = IPTOS_LOWDELAY`

- **ssl (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is **NOT** enabled by default in any current binary version of Samba.

This variable enables or disables the entire SSL mode. If it is set to "no", the SSL enabled samba behaves exactly like the non-SSL samba. If set to "yes", it depends on the variables "`ssl hosts`" and "`ssl hosts resign`" whether an SSL connection will be required.

Default: `ssl=no` **Example:** `ssl=yes`

- **ssl CA certDir (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is **NOT** enabled by default in any current binary version of Samba.

This variable defines where to look up the Certification Authorities. The given directory should contain one file for each CA that samba will trust. The file name must be the hash value over the "Distinguished Name" of the CA. How this directory is set up is explained later in this document. All files within the directory that don't fit into this naming scheme are ignored. You don't need this variable if you don't verify client certificates.

Default: `ssl CA certDir = /usr/local/ssl/certs`

- **ssl CA certFile (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

This variable is a second way to define the trusted CAs. The certificates of the trusted CAs are collected in one big file and this variable points to the file. You will probably only use one of the two ways to define your CAs. The first choice is preferable if you have many CAs or want to be flexible, the second is preferable if you only have one CA and want to keep things simple (you won't need to create the hashed file names). You don't need this variable if you don't verify client certificates.

Default: `ssl CA certFile = /usr/local/ssl/certs/trustedCAs.pem`

- **ssl ciphers (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

This variable defines the ciphers that should be offered during SSL negotiation. You should not set this variable unless you know what you are doing.

- **ssl client cert (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

The certificate in this file is used by **smbclient** if it exists. It's needed if the server requires a client certificate.

Default: `ssl client cert = /usr/local/ssl/certs/smbclient.pem`

- **ssl client key (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

This is the private key for **smbclient**. It's only needed if the client should have a certificate.

Default: `ssl client key = /usr/local/ssl/private/smbclient.pem`

- **ssl compatibility (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

This variable defines whether SSLeay should be configured for bug compatibility with other SSL implementations. This is probably not desirable because currently no clients with SSL implementations other than SSLeay exist.

Default: `ssl compatibility = no`

- **ssl hosts (G)**

See "**ssl hosts resign**".

- **ssl hosts resign (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is **NOT** enabled by default in any current binary version of Samba.

These two variables define whether samba will go into SSL mode or not. If none of them is defined, samba will allow only SSL connections. If the "**ssl hosts**" variable lists hosts (by IP-address, IP-address range, net group or name), only these hosts will be forced into SSL mode. If the "**ssl hosts resign**" variable lists hosts, only these hosts will NOT be forced into SSL mode. The syntax for these two variables is the same as for the "**hosts allow**" and "**hosts deny**" pair of variables, only that the subject of the decision is different: It's not the access right but whether SSL is used or not. See the "**allow hosts**" parameter for details. The example below requires SSL connections from all hosts outside the local net (which is 192.168.*.*).

Default: `ssl hosts = <empty string> ssl hosts resign = <empty string>`

Example: `ssl hosts resign = 192.168.`

- **ssl require clientcert (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is **NOT** enabled by default in any current binary version of Samba.

If this variable is set to "`yes`", the server will not tolerate connections from clients that don't have a valid certificate. The directory/file given in "**ssl CA certDir**" and "**ssl CA certFile**" will be used to look up the CAs that issued the client's certificate. If the certificate can't be verified positively, the connection will be terminated. If this variable is set to "`no`", clients don't need certificates. Contrary to web applications you really *should* require client certificates. In the web environment the client's data is sensitive (credit card numbers) and the server must prove to be trustworthy. In a file server environment the server's data will be sensitive and the clients must prove to be trustworthy.

Default: `ssl require clientcert = no`

- **ssl require servercert (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is **NOT** enabled by default in any current binary version of Samba.

If this variable is set to "yes", the **smbclient** will request a certificate from the server. Same as "**ssl require clientcert**" for the server.

Default: `ssl require servercert = no`

- **ssl server cert (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

This is the file containing the server's certificate. The server **_must_** have a certificate. The file may also contain the server's private key. See later for how certificates and private keys are created.

Default: `ssl server cert = <empty string>`

- **ssl server key (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

This file contains the private key of the server. If this variable is not defined, the key is looked up in the certificate file (it may be appended to the certificate). The server ***must*** have a private key and the certificate ***must*** match this private key.

Default: `ssl server key = <empty string>`

- **ssl version (G)**

This variable is part of SSL-enabled Samba. This is only available if the SSL libraries have been compiled on your system and the configure option "`-with-ssl`" was given at configure time.

Note that for export control reasons this code is ****NOT**** enabled by default in any current binary version of Samba.

This enumeration variable defines the versions of the SSL protocol that will be used. "`ssl2or3`" allows dynamic negotiation of SSL v2 or v3, "`ssl2`" results in SSL v2, "`ssl3`" results in SSL v3 and "`tls1`" results in TLS v1. TLS (Transport Layer Security) is the (proposed?) new standard for SSL.

Default: `ssl version = "ssl2or3"`

- **stat cache (G)**

This parameter determines if **smbd** will use a cache in order to speed up case insensitive name mappings. You should never need to change this parameter.

Default: `stat cache = yes`

- **stat cache size (G)**

This parameter determines the number of entries in the **stat cache**. You should never need to change this parameter.

Default: `stat cache size = 50`

- **status (G)**

This enables or disables logging of connections to a status file that **smbstatus** can read. With this disabled **smbstatus** won't be able to tell you what connections are active. You should never need to change this parameter.

Default: status = yes

- **strict locking (S)**

This is a boolean that controls the handling of file locking in the server. When this is set to "yes" the server will check every read and write access for file locks, and deny access if locks exist. This can be slow on some systems.

When strict locking is "no" the server does file lock checks only when the client explicitly asks for them.

Well behaved clients always ask for lock checks when it is important, so in the vast majority of cases "**strict locking = no**" is preferable.

Default: strict locking = no

Example: strict locking = yes

- **strict sync (S)**

Many Windows applications (including the Windows 98 explorer shell) seem to confuse flushing buffer contents to disk with doing a sync to disk. Under UNIX, a sync call forces the process to be suspended until the kernel has ensured that all outstanding data in kernel disk buffers has been safely stored onto stable storage. This is very slow and should only be done rarely. Setting this parameter to "no" (the default) means that **smbd** ignores the Windows applications requests for a sync call. There is only a possibility of losing data if the operating system itself that Samba is running on crashes, so there is little danger in this default setting. In addition, this fixes many performance problems that people have reported with the new Windows98 explorer shell file copies.

See also the "**sync always**" parameter.

Default: strict sync = no

Example: strict sync = yes

- **strip dot (G)**

This is a boolean that controls whether to strip trailing dots off UNIX filenames. This helps with some CDRoms that have filenames ending in a single dot.

Default: strip dot = no

Example: strip dot = yes

- **sync always (S)**

This is a boolean parameter that controls whether writes will always be written to stable storage before the write call returns. If this is false then the server will be guided by the client's request in each write call (clients can set a bit indicating that a particular write should be synchronous). If this is true then every write will be followed by a **fsync()** call to ensure the data is written to disk. Note that the "**strict sync**" parameter must be set to "yes" in order for this parameter to have any affect.

See also the "**strict sync**" parameter.

Default: sync always = no

Example: sync always = yes

- **syslog (G)**

This parameter maps how Samba debug messages are logged onto the system syslog logging levels. Samba debug level zero maps onto syslog LOG_ERR, debug level one

maps onto LOG_WARNING, debug level two maps to LOG_NOTICE, debug level three maps onto LOG_INFO. The parameter sets the threshold for doing the mapping, all Samba debug messages above this threshold are mapped to syslog LOG_DEBUG messages.

Default: `syslog = 1`

- **syslog only (G)**

If this parameter is set then Samba debug messages are logged into the system syslog only, and not to the debug log files.

Default: `syslog only = no`

- **time offset (G)**

This parameter is a setting in minutes to add to the normal GMT to local time conversion. This is useful if you are serving a lot of PCs that have incorrect daylight saving time handling.

Default: `time offset = 0`

Example: `time offset = 60`

- **time server (G)**

This parameter determines if **nmbd** advertises itself as a time server to Windows clients. The default is False.

Default: `time server = False`

Example: `time server = True`

- **timestamp logs (G)**

Samba2.0 will add timestamps to all log entries by default. This can be distracting if you are attempting to debug a problem. This parameter allows the timestamping to be turned off.

Default: `timestamp logs = True`

Example: `timestamp logs = False`

- **unix password sync (G)**

This boolean parameter controls whether Samba attempts to synchronize the UNIX password with the SMB password when the encrypted SMB password in the **smbpasswd** file is changed. If this is set to true the program specified in the **"passwd program"** parameter is called ***AS ROOT***—to allow the new UNIX password to be set without access to the old UNIX password (as the SMB password change code has no access to the old password cleartext, only the new). By default this is set to **"false"**.

See also **"passwd program"**, **"passwd chat"**.

Default: `unix password sync = False`

Example: `unix password sync = True`

- **unix realname (G)**

This boolean parameter when set causes samba to supply the real name field from the unix password file to the client. This is useful for setting up mail clients and WWW browsers on systems used by more than one person.

Default: `unix realname = no`

Example: `unix realname = yes`

- **update encrypted (G)**

This boolean parameter allows a user logging on with a plaintext password to have their encrypted (hashed) password in the smbpasswd file to be updated automatically as they log on. This option allows a site to migrate from plaintext password authentication (users authenticate with plaintext password over the wire, and are checked against a UNIX account database) to encrypted password authentication (the SMB challenge/response authentication mechanism) without forcing all users to re-enter their passwords via smbpasswd at the time the change is made. This is a convenience option to allow the change over to encrypted passwords to be made over a longer period. Once all users have encrypted representations of their passwords in the smbpasswd file this parameter should be set to "off".

In order for this parameter to work correctly the "**encrypt passwords**" parameter must be set to "no" when this parameter is set to "yes".

Note that even when this parameter is set a user authenticating to smbd must still enter a valid password in order to connect correctly, and to update their hashed (smbpasswd) passwords.

Default: `update encrypted = no`

Example: `update encrypted = yes`

- **use rhosts (G)**

If this global parameter is a true, it specifies that the UNIX users ".rhosts" file in their home directory will be read to find the names of hosts and users who will be allowed access without specifying a password.

Note

The use of **use rhosts** can be a major security hole. This is because you are trusting the PC to supply the correct username. It is very easy to get a PC to supply a false username. I recommend that the **use rhosts** option be only used if you really know what you are doing.

Default: `use rhosts = no`

Example: `use rhosts = yes`

- **user (S)**

Synonym for "**username**".

- **users (S)**

Synonym for "**username**".

- **username (S)**

Multiple users may be specified in a comma-delimited list, in which case the supplied password will be tested against each username in turn (left to right).

The **username=** line is needed only when the PC is unable to supply its own username. This is the case for the COREPLUS protocol or where your users have different WfWg usernames to UNIX usernames. In both these cases you may also be better using the `\\server\share%user` syntax instead.

The **username=** line is not a great solution in many cases as it means Samba will try to validate the supplied password against each of the usernames in the `username=` line in turn. This is slow and a bad idea for lots of users in case of duplicate passwords. You may get timeouts or security breaches using this parameter unwisely.

Samba relies on the underlying UNIX security. This parameter does not restrict who can login, it just offers hints to the Samba server as to what usernames might correspond to the supplied password. Users can login as whoever they please and they will be able to do no more damage than if they started a telnet session. The daemon runs as the user that they log in as, so they cannot do anything that user cannot do.

To restrict a service to a particular set of users you can use the **"valid users="** parameter.

If any of the usernames begin with a '@' then the name will be looked up first in the `yp netgroups` list (if Samba is compiled with `netgroup` support), followed by a lookup in the UNIX groups database and will expand to a list of all users in the group of that name.

If any of the usernames begin with a '+' then the name will be looked up only in the UNIX groups database and will expand to a list of all users in the group of that name.

If any of the usernames begin with a '&' then the name will be looked up only in the `yp netgroups` database (if Samba is compiled with `netgroup` support) and will expand to a list of all users in the `netgroup` group of that name.

Note that searching through a groups database can take quite some time, and some clients may time out during the search.

See the section **"NOTE ABOUT USERNAME/PASSWORD VALIDATION"** for more information on how this parameter determines access to the services.

Default: The guest account if a guest service, else the name of the service.

Examples:

```
username = fred
username = fred, mary, jack, jane, @users, @pcgroup
```

- **username level (G)**

This option helps Samba to try and "guess" at the real UNIX username, as many DOS clients send an all-uppercase username. By default Samba tries all lowercase, followed by the username with the first letter capitalized, and fails if the username is not found on the UNIX machine.

If this parameter is set to non-zero the behavior changes. This parameter is a number that specifies the number of uppercase combinations to try whilst trying to determine the UNIX user name. The higher the number the more combinations will be tried, but the slower the discovery of usernames will be. Use this parameter when you have strange usernames on your UNIX machine, such as "AstrangeUser".

Default: `username level = 0`

Example: `username level = 5`

- **username map (G)**

This option allows you to specify a file containing a mapping of usernames from the clients to the server. This can be used for several purposes. The most common is to map usernames that users use on DOS or Windows machines to those that the UNIX box

uses. The other is to map multiple users to a single username so that they can more easily share files.

The use of this option, therefore, relates to UNIX usernames and not Windows (specifically NT Domain) usernames. In other words, once a name has been mapped using this option, the Samba server uses the mapped name for internal *AND* external purposes.

This option is *DIFFERENT* from the "**domain user map**" parameter, which maintains a one-to-one mapping between UNIX usernames and NT Domain Usernames: more specifically, the Samba server maintains a link between *BOTH* usernames, presenting the NT username to the external NT world, and using the UNIX username internally.

The map file is parsed line by line. Each line should contain a single UNIX username on the left then a '=' followed by a list of usernames on the right. The list of usernames on the right may contain names of the form @group in which case they will match any UNIX username in that group. The special client name '*' is a wildcard and matches any name. Each line of the map file may be up to 1023 characters long.

The file is processed on each line by taking the supplied username and comparing it with each username on the right hand side of the '=' signs. If the supplied name matches any of the names on the right hand side then it is replaced with the name on the left. Processing then continues with the next line.

If any line begins with a '#' or a ';' then it is ignored

If any line begins with an '!' then the processing will stop after that line if a mapping was done by the line. Otherwise mapping continues with every line being processed. Using '!' is most useful when you have a wildcard mapping line later in the file.

For example to map from the name "admin" or "administrator" to the UNIX name "root" you would use:

```
root = admin administrator
```

Or to map anyone in the UNIX group "system" to the UNIX name "sys" you would use:

```
sys = @system
```

You can have as many mappings as you like in a username map file.

If your system supports the NIS NETGROUP option then the netgroup database is checked before the /etc/group database for matching groups.

You can map Windows usernames that have spaces in them by using double quotes around the name. For example:

```
tridge = "Andrew Tridgell"
```

would map the windows username "Andrew Tridgell" to the unix username tridge. The following example would map mary and fred to the unix user sys, and map the rest to guest. Note the use of the '!' to tell Samba to stop processing if it gets a match on that line.

```
!sys = mary fred
guest = *
```

Note that the remapping is applied to all occurrences of usernames. Thus if you connect to "\\server\fred" and "fred" is remapped to "mary" then you will actually be connecting to "\\server\mary" and will need to supply a password suitable for "mary" not "fred". The only exception to this is the username passed to the "**password server**" (if you have one). The password server will receive whatever username the client supplies without modification.

Also note that no reverse mapping is done. The main effect this has is with printing. Users who have been mapped may have trouble deleting print jobs as PrintManager under WfWg will think they don't own the print job.

Default: `no username map`

Example: `username map = /usr/local/samba/lib/users.map`

- **valid chars (S)**

The option allows you to specify additional characters that should be considered valid by the server in filenames. This is particularly useful for national character sets, such as adding u-umlaut or a-ring.

The option takes a list of characters in either integer or character form with spaces between them. If you give two characters with a colon between them then it will be taken as an lowercase:uppercase pair.

If you have an editor capable of entering the characters into the config file then it is probably easiest to use this method. Otherwise you can specify the characters in octal, decimal or hexadecimal form using the usual C notation.

For example to add the single character 'Z' to the charset (which is a pointless thing to do as it's already there) you could do one of the following

```
valid chars = Z
valid chars = z:Z
valid chars = 0132:0172
```

The last two examples above actually add two characters, and alter the uppercase and lowercase mappings appropriately.

Note that you **MUST** specify this parameter after the "**client code page**" parameter if you have both set. If "**client code page**" is set after the "**valid chars**" parameter the "**valid chars**" settings will be overwritten.

See also the "**client code page**" parameter.

Default:

Samba defaults to using a reasonable set of valid characters for English systems

Example: `valid chars = 0345:0305 0366:0326 0344: 0304`

The above example allows filenames to have the Swedish characters in them.

Note

It is actually quite difficult to correctly produce a "**valid chars**" line for a particular system. To automate the process tino@augzburg.net has written a package called "**validchars**" which will automatically produce a complete "**valid chars**" line for a given client system. Look in the examples/validchars/ subdirectory of your Samba source code distribution for this package.

- **valid users (S)**

This is a list of users that should be allowed to login to this service. Names starting with '@', '+' and '&' are interpreted using the same rules as described in the "**invalid users**" parameter.

If this is empty (the default) then any user can login. If a username is in both this list and the **"invalid users"** list then access is denied for that user.

The current servicename is substituted for **"%S"**. This is useful in the **[homes]** section.

See also **"invalid users"**.

Default: No valid users list. (anyone can login)

Example: valid users = greg, @pcusers

- **veto files(S)**

This is a list of files and directories that are neither visible nor accessible. Each entry in the list must be separated by a **'/'**, which allows spaces to be included in the entry. **'*'** and **'?'** can be used to specify multiple files or directories as in DOS wildcards.

Each entry must be a unix path, not a DOS path and must **not** include the unix directory separator **'/'**.

Note that the **"case sensitive"** option is applicable in vetoing files.

One feature of the veto files parameter that it is important to be aware of, is that if a directory contains nothing but files that match the veto files parameter (which means that Windows/DOS clients cannot ever see them) is deleted, the veto files within that directory **are automatically deleted** along with it, if the user has UNIX permissions to do so.

Setting this parameter will affect the performance of Samba, as it will be forced to check all files and directories for a match as they are scanned.

See also **"hide files"** and **"case sensitive"**.

Default: No files or directories are vetoed.

Examples:

Example 1.

```
Veto any files containing the word Security, any ending in
.tmp, and any directory containing the word root.
veto files = /*Security*/*.tmp/*root*/
```

Example 2.

```
Veto the Apple specific files that a NetAtalk server
creates.
veto files = /.AppleDouble/.bin/.AppleDesktop/Network Trash
Folder/
```

- **veto oplock files (S)**

This parameter is only valid when the **"oplocks"** parameter is turned on for a share. It allows the Samba administrator to selectively turn off the granting of oplocks on selected files that match a wildcarded list, similar to the wildcarded list used in the **"veto files"** parameter.

Default: No files are vetoed for oplock grants.

Examples:

You might want to do this on files that you know will be heavily contended for by clients. A good example of this is in the NetBench SMB benchmark program, which causes heavy client contention for files ending in **".SEM"**. To cause Samba not to grant oplocks on these files you would use the line (either in the **[global]** section or in the section for the particular NetBench share:

```
veto oplock files = /*.SEM/
```

- **volume (S)**

This allows you to override the volume label returned for a share. Useful for CDRoms with installation programs that insist on a particular volume label.

The default is the name of the share.

- **wide links (S)**

This parameter controls whether or not links in the UNIX file system may be followed by the server. Links that point to areas within the directory tree exported by the server are always allowed; this parameter controls access only to areas that are outside the directory tree being exported.

Default: `wide links = yes`

Example: `wide links = no`

- **wins proxy (G)**

This is a boolean that controls if **nmbd** will respond to broadcast name queries on behalf of other hosts. You may need to set this to "yes" for some older clients.

Default: `wins proxy = no`

- **wins server (G)**

This specifies the IP address (or DNS name: IP address for preference) of the WINS server that **nmbd** should register with. If you have a WINS server on your network then you should set this to the WINS server's IP.

You should point this at your WINS server if you have a multi-subnetted network.

Note

You need to set up Samba to point to a WINS server if you have multiple subnets and wish cross-subnet browsing to work correctly.

See the documentation file BROWSING.txt in the docs/ directory of your Samba source distribution.

Default: `wins server =`

Example: `wins server = 192.9.200.1`

- **wins support (G)**

This boolean controls if the **nmbd** process in Samba will act as a WINS server. You should not set this to true unless you have a multi-subnetted network and you wish a particular **nmbd** to be your WINS server. Note that you should **NEVER** set this to true on more than one machine in your network.

Default: `wins support = no`

- **workgroup (G)**

This controls what workgroup your server will appear to be in when queried by clients. Note that this parameter also controls the Domain name used with the "**security=domain**" setting.

Default: `set at compile time to WORKGROUP`

Example: `workgroup = MYGROUP`

- **writable (S)**

Synonym for "**writeable**" for people who can't spell :-). Pronounced "ritter-bull".

- **write list (S)**

This is a list of users that are given read-write access to a service. If the connecting user is in this list then they will be given write access, no matter what the "**read only**" option is set to. The list can include group names using the `@group` syntax.

Note that if a user is in both the read list and the write list then they will be given write access.

See also the "**read list**" option.

Default: `write list = <empty string>`

Example: `write list = admin, root, @staff`

- **write ok (S)**

Synonym for **writeable**.

- **write raw (G)**

This parameter controls whether or not the server will support raw writes SMB's when transferring data from clients. You should never need to change this parameter.

Default: `write raw = yes`

- **writeable**

An inverted synonym is "**read only**".

If this parameter is "**no**", then users of a service may not create or modify files in the service's directory.

Note that a printable service ("**printable = yes**") will **ALWAYS** allow writing to the directory (user privileges permitting), but only via spooling operations.

Default: `writeable = no`

Examples:

```
read only = no
writeable = yes
write ok = yes
```

Warnings

Although the configuration file permits service names to contain spaces, your client software may not. Spaces will be ignored in comparisons anyway, so it shouldn't be a problem—but be aware of the possibility.

On a similar note, many clients—especially DOS clients—limit service names to eight characters. **Smbd** has no such limitation, but attempts to connect from such clients will fail if they truncate the service names. For this reason you should probably keep your service names down to eight characters in length.

Use of the **[homes]** and **[printers]** special sections make life for an administrator easy, but the various combinations of default attributes can be tricky. Take extreme care when designing these sections. In particular, ensure that the permissions on spool directories are correct.

Version

This man page is correct for version 2.0 of the Samba suite.

See Also

smbd (8), **smbclient (1)**, **nmbd (8)**, **testparm (1)**, **testprns (1)**, **Samba**, **nmblookup (1)**, **smbpasswd (5)**, **smbpasswd (8)**.

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@amba.org.

See **samba (7)** to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

smbclient (1)

Samba

23 Oct 1998

Name

smbclient—ftp-like client to access SMB/CIFS resources on servers

Synopsis

```
smbclient servicename [password] [-s smb.conf] [-B IP addr] [-O socket options][  
-R name resolve order] [-M NetBIOS name] [-i scope] [-N] [-n NetBIOS name] [-d  
debuglevel] [-P] [-p port] [-l log basename] [-h] [-l dest IP] [-E] [-U username] [-L NetBIOS  
name] [-t terminal code] [-m max protocol] [-W workgroup] [-T<cx>IXFqgbNan] [-D  
directory] [-c command string]
```

Description

This program is part of the **Samba** suite.

smbclient is a client that can "talk" to an SMB/CIFS server. It offers an interface similar to that of the ftp program (see **ftp (1)**). Operations include things like getting files from the server to the local machine, putting files from the local machine to the server, retrieving directory information from the server and so on.

Options

- **servicename** servicename is the name of the service you want to use on the server. A service name takes the form `//server/service` where *server* is the NetBIOS name of the SMB/CIFS server offering the desired service and *service* is the name of the service offered. Thus to connect to the service *printer* on the SMB/CIFS server *smbserver*, you would use the servicename

-
- `//smbserver/printer`
-

Note that the server name required is NOT necessarily the IP (DNS) host name of the server! The name required is a NetBIOS server name, which may or may not be the same as the IP hostname of the machine running the server.

The server name is looked up according to either the **-R** parameter to **smbclient** or using the **name resolve order** parameter in the `smb.conf` file, allowing an administrator to change the order and methods by which server names are looked up.

- **password** password is the password required to access the specified service on the specified server. If this parameter is supplied, the **-N** option (suppress password prompt) is assumed.

There is no default password. If no password is supplied on the command line (either by using this parameter or adding a password to the **-U** option (see below)) and the **-N** option is not specified, the client will prompt for a password, even if the desired service does not require one. (If no password is required, simply press ENTER to provide a null password.)

Note

Some servers (including OS/2 and Windows for Workgroups) insist on an uppercase password. Lowercase or mixed case passwords may be rejected by these servers.

Be cautious about including passwords in scripts.

- **-s smb.conf** This parameter specifies the pathname to the Samba configuration file, `smb.conf`. This file controls all aspects of the Samba setup on the machine and `smbclient` also needs to read this file.

- **-B IP addr** The IP address to use when sending a broadcast packet.
- **-O socket options** TCP socket options to set on the client socket. See the socket options parameter in the **smb.conf (5)** manpage for the list of valid options.
- **-R name resolve order** This option allows the user of smbclient to determine what name resolution services to use when looking up the NetBIOS name of the host being connected to.

The options are :*lmhosts*, *host*, *wins* and *bcast*". They cause names to be resolved as follows:

- **lmhosts**: Lookup an IP address in the Samba *lmhosts* file. The *lmhosts* file is stored in the same directory as the **smb.conf** file.
- **host**: Do a standard host name to IP address resolution, using the system */etc/hosts*, NIS, or DNS lookups. This method of name resolution is operating system depended for instance on IRIX or Solaris this may be controlled by the */etc/nsswitch.conf* file).
- **wins**: Query a name with the IP address listed in the **wins server** parameter in the *smb.conf* file. If no WINS server has been specified this method will be ignored.
- **bcast**: Do a broadcast on each of the known local interfaces listed in the *interfaces* parameter in the *smb.conf* file. This is the least reliable of the name resolution methods as it depends on the target host being on a locally connected subnet. To specify a particular broadcast address the **-B** option may be used.

If this parameter is not set then the name resolve order defined in the **smb.conf** file parameter (**name resolve order**) will be used.

The default order is *lmhosts*, *host*, *wins*, *bcast* and without this parameter or any entry in the **"name resolve order"** parameter of the **smb.conf** file the name resolution methods will be attempted in this order.

- **-M NetBIOS name** This options allows you to send messages, using the "WinPopup" protocol, to another computer. Once a connection is established you then type your message, pressing ^D (control-D) to end.

If the receiving computer is running WinPopup the user will receive the message and probably a beep. If they are not running WinPopup the message will be lost, and no error message will occur.

The message is also automatically truncated if the message is over 1600 bytes, as this is the limit of the protocol.

One useful trick is to cat the message through **smbclient**. For example:

```
cat mymessage.txt | smbclient -M FRED
```

will send the message in the file *mymessage.txt* to the machine FRED.

You may also find the **-U** and **-I** options useful, as they allow you to control the FROM and TO parts of the message.

See the **message command** parameter in the **smb.conf (5)** for a description of how to handle incoming WinPopup messages in Samba.

Note

Copy WinPopup into the startup group on your WfWg PCs if you want them to always be able to receive messages.

- **-i scope** This specifies a NetBIOS scope that smbclient will use to communicate with when generating NetBIOS names. For details on the use of NetBIOS scopes, see rfc1001.txt and rfc1002.txt. NetBIOS scopes are very rarely used, only set this parameter if you are the system administrator in charge of all the NetBIOS systems you communicate with.
- **-N** If specified, this parameter suppresses the normal password prompt from the client to the user. This is useful when accessing a service that does not require a password.

Unless a password is specified on the command line or this parameter is specified, the client will request a password.

- **-n NetBIOS name** By default, the client will use the local machine's hostname (in uppercase) as its NetBIOS name. This parameter allows you to override the host name and use whatever NetBIOS name you wish.
- **-d debuglevel** debuglevel is an integer from 0 to 10, or the letter "A".

The default value if this parameter is not specified is zero.

The higher this value, the more detail will be logged to the log files about the activities of the client. At level 0, only critical errors and serious warnings will be logged. Level 1 is a reasonable level for day to day running—it generates a small amount of information about operations carried out.

Levels above 1 will generate considerable amounts of log data, and should only be used when investigating a problem. Levels above 3 are designed for use only by developers and generate HUGE amounts of log data, most of which is extremely cryptic. If debuglevel is set to the letter "A", then *all* debug messages will be printed. This setting is for developers only (and people who *really* want to know how the code works internally).

Note that specifying this parameter here will override the **log level** parameter in the **smb.conf (5)** file.

- **-P** This option is no longer used. The code in Samba2.0 now lets the server decide the device type, so no printer specific flag is needed.
- **-p port** This number is the TCP port number that will be used when making connections to the server. The standard (well-known) TCP port number for an SMB/CIFS server is 139, which is the default.
- **-l logfile** If specified, logfile specifies a base filename into which operational data from the running client will be logged.

The default base name is specified at compile time.

The base name is used to generate actual log file names. For example, if the name specified was "log", the debug file would be `log.client`.

The log file generated is never removed by the client.

- **-h** Print the usage message for the client.
- **-I IP address** IP address is the address of the server to connect to. It should be specified in standard "a.b.c.d" notation.

Normally the client would attempt to locate a named SMB/CIFS server by looking it up via the NetBIOS name resolution mechanism described above in the **name resolve order** parameter above. Using this parameter will force the client to assume that the server is on the machine with the specified IP address and the NetBIOS name component of the resource being connected to will be ignored.

There is no default for this parameter. If not supplied, it will be determined automatically by the client as described above.

- **-E** This parameter causes the client to write messages to the standard error stream (stderr) rather than to the standard output stream.

By default, the client writes messages to standard output—typically the user's tty.

- **-U username** This specifies the user name that will be used by the client to make a connection, assuming your server is not a downlevel server that is running a protocol level that uses passwords on shares, not on usernames.

Some servers are fussy about the case of this name, and some insist that it must be a valid NetBIOS name.

If no username is supplied, it will default to an uppercase version of the environment variable `USER` or `LOGNAME` in that order. If no username is supplied and neither environment variable exists the username "GUEST" will be used.

If the `USER` environment variable contains a ``%'` character, everything after that will be treated as a password. This allows you to set the environment variable to be `USER=username%password` so that a password is not passed on the command line (where it may be seen by the `ps` command).

If the service you are connecting to requires a password, it can be supplied using the **-U** option, by appending a percent symbol ("`%`") then the password to username. For example, to attach to a service as user "`fred`" with password "`secret`", you would specify.

```
-U fred%secret
```

on the command line. Note that there are no spaces around the percent symbol.

If you specify the password as part of username then the **-N** option (suppress password prompt) is assumed.

If you specify the password as a parameter *AND* as part of username then the password as part of username will take precedence. Putting nothing before or nothing after the percent symbol will cause an empty username or an empty password to be used, respectively.

The password may also be specified by setting up an environment variable called `PASSWORD` that contains the users password. Note that this may be very insecure on some systems but on others allows users to script `smbclient` commands without having a password appear in the command line of a process listing.

Note

Some servers (including OS/2 and Windows for Workgroups) insist on an uppercase password. Lowercase or mixed case passwords may be rejected by these servers.

Be cautious about including passwords in scripts or in the `PASSWORD` environment variable. Also, on many systems the command line of a running process may be seen via the `ps` command to be safe always allow `smbclient` to prompt for a password and type it in directly.

- **-L** This option allows you to look at what services are available on a server. You use it as `"smbclient -L host"` and a list should appear. The **-I** option may be useful if your NetBIOS names don't match your tcp/ip dns host names or if you are trying to reach a host on another network.
- **-t terminal code** This option tells `smbclient` how to interpret filenames coming from the remote server. Usually Asian language multibyte UNIX implementations use different character sets than SMB/CIFS servers (*EUC* instead of *SJIS* for example). Setting this parameter will let `smbclient` convert between the UNIX filenames and the SMB filenames correctly. This option has not been seriously tested and may have some problems.

The terminal codes include `sjis`, `euc`, `jis7`, `jis8`, `junet`, `hex`, `cap`. This is not a complete list, check the Samba source code for the complete list.

- **-m max protocol level** With the new code in Samba2.0, `smbclient` always attempts to connect at the maximum protocols level the server supports. This parameter is preserved for backwards compatibility, but any string following the **-m** will be ignored.
- **-W WORKGROUP** Override the default workgroup specified in the `workgroup` parameter of the `smb.conf` file for this connection. This may be needed to connect to some servers.
- **-T tar options** `smbclient` may be used to create **tar (1)** compatible backups of all the files on an SMB/CIFS share. The secondary tar flags that can be given to this option are:
 - **c** Create a tar file on UNIX. Must be followed by the name of a tar file, tape device or `"-"` for standard output. If using standard output you must turn the log level to its lowest value `-d0` to avoid corrupting your tar file. This flag is mutually exclusive with the **x** flag.
 - **x** Extract (restore) a local tar file back to a share. Unless the **-D** option is given, the tar files will be restored from the top level of the share. Must be followed by the name of the tar file, device or `"-"` for standard input. Mutually exclusive with the **c** flag. Restored files have their creation times (mtime) set to the date saved in the tar file. Directories currently do not get their creation dates restored properly.
- **I** Include files and directories. Is the default behavior when filenames are specified above. Causes tar files to be included in an extract or create (and

therefore everything else to be excluded). See example below. Filename globbing works in one of two ways. See **r** below.

- **X** Exclude files and directories. Causes tar files to be excluded from an extract or create. See example below. Filename globbing works in one of two ways now. See **r** below.
- **b** Blocksize. Must be followed by a valid (greater than zero) blocksize. Causes tar file to be written out in blocksize*TBLOCK (usually 512 byte) blocks.
- **g** Incremental. Only back up files that have the archive bit set. Useful only with the **c** flag.
- **q** Quiet. Keeps tar from printing diagnostics as it works. This is the same as tarmode quiet.
- **r** Regular expression include or exclude. Uses regular regular expression matching for excluding or excluding files if compiled with HAVE_REGEX_H. However this mode can be very slow. If not compiled with HAVE_REGEX_H, does a limited wildcard match on * and ?.
- **N** Newer than. Must be followed by the name of a file whose date is compared against files found on the share during a create. Only files newer than the file specified are backed up to the tar file. Useful only with the **c** flag.
- **a** Set archive bit. Causes the archive bit to be reset when a file is backed up. Useful with the **g** and **c** flags.

Tar Long File Names

smbclient's tar option now supports long file names both on backup and restore. However, the full path name of the file must be less than 1024 bytes. Also, when a tar archive is created, smbclient's tar option places all files in the archive with relative names, not absolute names.

Tar Filenames

All file names can be given as DOS path names (with \ as the component separator) or as UNIX path names (with / as the component separator).

Examples

- Restore from tar file backup.tar into myshare on mypc (no password on share).

```
smbclient //mypc/myshare "" -N -Tx backup.tar
```

- Restore everything except users/docs

```
smbclient //mypc/myshare "" -N -TXx backup.tar users/docs
```

- Create a tar file of the files beneath users/docs.

```
smbclient //mypc/myshare "" -N -Tc backup.tar users/docs
```

- Create the same tar file as above, but now use a DOS path name.

```
smbclient //mypc/myshare "" -N -tc backup.tar users\edocs
```

- Create a tar file of all the files and directories in the share.

```
smbclient //mypc/myshare "" -N -Tc backup.tar *
```

- **-D initial directory** Change to initial directory before starting. Probably only of any use with the tar **-T** option.
- **-c command string** command string is a semicolon separated list of commands to be executed instead of prompting from stdin. **-N** is implied by **-c**.

This is particularly useful in scripts and for printing stdin to the server, e.g. `-c 'print -'`.

Operations

Once the client is running, the user is presented with a prompt:

```
smb: \>
```

The backslash ("`\`") indicates the current working directory on the server, and will change if the current working directory is changed.

The prompt indicates that the client is ready and waiting to carry out a user command. Each command is a single word, optionally followed by parameters specific to that command. Command and parameters are space-delimited unless these notes specifically state otherwise. All commands are case-insensitive. Parameters to commands may or may not be case sensitive, depending on the command.

You can specify file names which have spaces in them by quoting the name with double quotes, for example "a long file name".

Parameters shown in square brackets (e.g., "[parameter]") are optional. If not given, the command will use suitable defaults. Parameters shown in angle brackets (e.g., "<parameter>") are required.

Note that all commands operating on the server are actually performed by issuing a request to the server. Thus the behavior may vary from server to server, depending on how the server was implemented.

The commands available are given here in alphabetical order.

- **? [command]** If "command" is specified, the `?` command will display a brief informative message about the specified command. If no command is specified, a list of available commands will be displayed.

- **!** **[shell command]** If "shell command" is specified, the ! command will execute a shell locally and run the specified shell command. If no command is specified, a local shell will be run.
- **cd [directory name]** If "directory name" is specified, the current working directory on the server will be changed to the directory specified. This operation will fail if for any reason the specified directory is inaccessible.

If no directory name is specified, the current working directory on the server will be reported.

- **del <mask>** The client will request that the server attempt to delete all files matching "mask" from the current working directory on the server.
- **dir <mask>** A list of the files matching "mask" in the current working directory on the server will be retrieved from the server and displayed.
- **exit** Terminate the connection with the server and exit from the program.
- **get <remote file name> [local file name]** Copy the file called "remote file name" from the server to the machine running the client. If specified, name the local copy "local file name". Note that all transfers in smbclient are binary. See also the **lowercase** command.
- **help [command]** See the ? command above.
- **lcd [directory name]** If "directory name" is specified, the current working directory on the local machine will be changed to the directory specified. This operation will fail if for any reason the specified directory is inaccessible.

If no directory name is specified, the name of the current working directory on the local machine will be reported.

- **lowercase** Toggle lowercasing of filenames for the **get** and **mget** commands.

When lowercasing is toggled ON, local filenames are converted to lowercase when using the **get** and **mget** commands. This is often useful when copying (say) MSDOS files from a server, because lowercase filenames are the norm on UNIX systems.

- **ls <mask>** See the **dir** command above.
- **mask <mask>** This command allows the user to set up a mask which will be used during recursive operation of the **mget** and **mput** commands.

The masks specified to the **mget** and **mput** commands act as filters for directories rather than files when recursion is toggled ON.

The mask specified with the **.B** mask command is necessary to filter files within those directories. For example, if the mask specified in an **mget** command is "source*" and the mask specified with the mask command is "*.c" and recursion is toggled ON, the **mget** command will retrieve all files matching "*.c" in all directories below and including all directories matching "source*" in the current working directory.

Note that the value for mask defaults to blank (equivalent to "") and remains so until the mask command is used to change it. It retains the most recently specified value indefinitely. To avoid unexpected results it would be wise to change the value of **.l** mask back to "" after using the **mget** or **mput** commands.

- **md <directory name>** See the **mkdir** command.
- **mget <mask>** Copy all files matching mask from the server to the machine running the client.

Note that mask is interpreted differently during recursive operation and nonrecursive operation—refer to the **recurse** and **mask** commands for more information. Note that all transfers in .B smbclient are binary. See also the lowercase command.

- **mkdir <directory name>** Create a new directory on the server (user access privileges permitting) with the specified name.
- **mput <mask>** Copy all files matching mask in the current working directory on the local machine to the current working directory on the server.

Note that mask is interpreted differently during recursive operation and non-recursive operation—refer to the **recurse** and **mask** commands for more information. Note that all transfers in .B smbclient are binary.

- **print <file name>** Print the specified file from the local machine through a printable service on the server.

See also the **printmode** command.

- **printmode <graphics or text>** Set the print mode to suit either binary data (such as graphical information) or text. Subsequent print commands will use the currently set print mode.
- **prompt** Toggle prompting for filenames during operation of the **mget** and **mput** commands.

When toggled ON, the user will be prompted to confirm the transfer of each file during these commands. When toggled OFF, all specified files will be transferred without prompting.

- **put <local file name> [remote file name]** Copy the file called "local file name" from the machine running the client to the server. If specified, name the remote copy "**remote file name**". Note that all transfers in smbclient are binary. See also the **lowercase** command.
- **queue** Displays the print queue, showing the job id, name, size and current status.
- **quit** See the **exit** command.
- **rd <directory name>** See the **rmdir** command.
- **recurse** Toggle directory recursion for the commands **mget** and **mput**.

When toggled ON, these commands will process all directories in the source directory (i.e., the directory they are copying .IR from) and will recurse into any that match the mask specified to the command. Only files that match the mask specified using the **mask** command will be retrieved. See also the **mask** command.

When recursion is toggled OFF, only files from the current working directory on the source machine that match the mask specified to the **mget** or **mput** commands will be copied, and any mask specified using the **mask** command will be ignored.

- **rm <mask>** Remove all files matching mask from the current working directory on the server.
- **rmdir <directory name>** Remove the specified directory (user access privileges permitting) from the server.
- **tar <c|x>[IXbgNa]** Performs a tar operation—see the **-T** command line option above. Behavior may be affected by the **tarmode** command (see below). Using **g** (incremental) and **N** (newer) will affect tarmode settings. Note that using the **"-"** option with **tar x** may not work—use the command line option instead.
- **blocksize <blocksize> Blocksize.** Must be followed by a valid (greater than zero) blocksize. Causes tar file to be written out in **blocksize*TBLOCK** (usually 512 byte) blocks.
- **tarmode <full|inc|reset |noreset>** Changes tar's behavior with regard to archive bits. In full mode, tar will back up everything regardless of the archive bit setting (this is the default mode). In incremental mode, tar will only back up files with the archive bit set. In reset mode, tar will reset the archive bit on all files it backs up (implies read/write share).
- **setmode <filename> <perm=[+ |-]rsha>** A version of the DOS attrib command to set file permissions. For example:

```
setmode myfile +r
```

would make myfile read only.

Notes

Some servers are fussy about the case of supplied usernames, passwords, share names (AKA service names) and machine names. If you fail to connect try giving all parameters in uppercase.

It is often necessary to use the **-n** option when connecting to some types of servers. For example OS/2 LanManager insists on a valid NetBIOS name being used, so you need to supply a valid name that would be known to the server.

smbclient supports long file names where the server supports the LANMAN2 protocol or above.

Environment Variables

The variable **USER** may contain the username of the person using the client. This information is used only if the protocol level is high enough to support session-level passwords.

The variable **PASSWORD** may contain the password of the person using the client. This information is used only if the protocol level is high enough to support session-level passwords.

Installation

The location of the client program is a matter for individual system administrators. The following are thus suggestions only.

It is recommended that the smbclient software be installed in the /usr/local/samba/bin or /usr/samba/bin directory, this directory readable by all, writeable only by root. The client program itself should be executable by all. The client should *NOT* be setuid or setgid!

The client log files should be put in a directory readable and writeable only by the user.

To test the client, you will need to know the name of a running SMB/CIFS server. It is possible to run **smbd (8)** as an ordinary user—running that server as a daemon on a user-accessible port (typically any port number over 1024) would provide a suitable test server.

Diagnostics

Most diagnostics issued by the client are logged in a specified log file. The log file name is specified at compile time, but may be overridden on the command line.

The number and nature of diagnostics available depends on the debug level used by the client. If you have problems, set the debug level to 3 and peruse the log files.

Version

This man page is correct for version 2.0 of the Samba suite.

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@samba.org.

See **samba (7)** to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

smbd (8)

Samba

23 Oct 1998

Name

smbd—server to provide SMB/CIFS services to clients

Synopsis

smbd [-D] [-a] [-o] [-d debuglevel] [-l log file] [-p port number] [-O socket options] [-s configuration file] [-i scope] [-P] [-h]

Description

This program is part of the **Samba** suite.

smbd is the server daemon that provides filesharing and printing services to Windows clients. The server provides filespace and printer services to clients using the SMB (or

CIFS) protocol. This is compatible with the LanManager protocol, and can service LanManager clients. These include MSCLIENT 3.0 for DOS, Windows for Workgroups, Windows 95, Windows NT, OS/2, DAVE for Macintosh, and smbfs for Linux.

An extensive description of the services that the server can provide is given in the man page for the configuration file controlling the attributes of those services (see **smb.conf (5)**). This man page will not describe the services, but will concentrate on the administrative aspects of running the server.

Please note that there are significant security implications to running this server, and the **smb.conf (5)** manpage should be regarded as mandatory reading before proceeding with installation.

A session is created whenever a client requests one. Each client gets a copy of the server for each session. This copy then services all connections made by the client during that session. When all connections from its client are closed, the copy of the server for that client terminates.

The configuration file, and any files that it includes, are automatically reloaded every minute, if they change. You can force a reload by sending a SIGHUP to the server. Reloading the configuration file will not affect connections to any service that is already established. Either the user will have to disconnect from the service, or `smbd` killed and restarted.

Options

- **-D** If specified, this parameter causes the server to operate as a daemon. That is, it detaches itself and runs in the background, fielding requests on the appropriate port. Operating the server as a daemon is the recommended way of running `smbd` for servers that provide more than casual use file and print services.

By default, the server will NOT operate as a daemon.

- **-a** If this parameter is specified, each new connection will append log messages to the log file. This is the default.
- **-o** If this parameter is specified, the log files will be overwritten when opened. By default, the log files will be appended to.
- **-d debuglevel** `debuglevel` is an integer from 0 to 10.

The default value if this parameter is not specified is zero.

The higher this value, the more detail will be logged to the log files about the activities of the server. At level 0, only critical errors and serious warnings will be logged. Level 1 is a reasonable level for day to day running—it generates a small amount of information about operations carried out.

Levels above 1 will generate considerable amounts of log data, and should only be used when investigating a problem. Levels above 3 are designed for use only by developers and generate HUGE amounts of log data, most of which is extremely cryptic.

Note that specifying this parameter here will override the **log level** parameter in the **smb.conf (5)** file.

- **-l log file** If specified, *log file* specifies a log filename into which informational and debug messages from the running server will be logged. The log file generated is never removed by the server although its size may be controlled by the **max log size** option in the **smb.conf (5)** file. The default log file name is specified at compile time.
- **-O socket options** See the **socket options** parameter in the **smb.conf (5)** file for details.

- **-p port number** port number is a positive integer value. The default value if this parameter is not specified is 139.

This number is the port number that will be used when making connections to the server from client software. The standard (well-known) port number for the SMB over TCP is 139, hence the default. If you wish to run the server as an ordinary user rather than as root, most systems will require you to use a port number greater than 1024—ask your system administrator for help if you are in this situation.

In order for the server to be useful by most clients, should you configure it on a port other than 139, you will require port redirection services on port 139, details of which are outlined in `rfc1002.txt` section 4.3.5.

This parameter is not normally specified except in the above situation.

- **-s configuration file** The file specified contains the configuration details required by the server. The information in this file includes server-specific information such as what printcap file to use, as well as descriptions of all the services that the server is to provide. See **smb.conf (5)** for more information. The default configuration file name is determined at compile time.
- **-i scope** This specifies a NetBIOS scope that the server will use to communicate with when generating NetBIOS names. For details on the use of NetBIOS scopes, see `rfc1001.txt` and `rfc1002.txt`. NetBIOS scopes are *very* rarely used, only set this parameter if you are the system administrator in charge of all the NetBIOS systems you communicate with.
- **-h** Prints the help information (usage) for `smbd`.
- **-P** Passive option. Causes `smbd` not to send any network traffic out. Used for debugging by the developers only.

Files

`/etc/inetd.conf`

If the server is to be run by the `inetd` meta-daemon, this file must contain suitable startup information for the meta-daemon. See the section `INSTALLATION` below.

`/etc/rc`

(or whatever initialization script your system uses).

If running the server as a daemon at startup, this file will need to contain an appropriate startup sequence for the server. See the section `INSTALLATION` below.

`/etc/services`

If running the server via the meta-daemon `inetd`, this file must contain a mapping of service name (e.g., `netbios-ssn`) to service port (e.g., 139) and protocol type (e.g., `tcp`). See the section `INSTALLATION` below.

`/usr/local/samba/lib/smb.conf`

This is the default location of the `smb.conf` server configuration file. Other common places that systems install this file are `/usr/samba/lib/smb.conf` and `/etc/smb.conf`.

This file describes all the services the server is to make available to clients. See **smb.conf (5)** for more information.

Limitations

On some systems `smbd` cannot change `uid` back to root after a `setuid()` call. Such systems are called "trapdoor" `uid` systems. If you have such a system, you will be unable to connect from a client (such as a PC) as two different users at once. Attempts to connect the second user will result in "access denied" or similar.

Environment Variables

Printer

If no printer name is specified to printable services, most systems will use the value of this variable (or "lp" if this variable is not defined) as the name of the printer to use. This is not specific to the server, however.

Installation

The location of the server and its support files is a matter for individual system administrators. The following are thus suggestions only.

It is recommended that the server software be installed under the `/usr/local/samba` hierarchy, in a directory readable by all, writeable only by root. The server program itself should be executable by all, as users may wish to run the server themselves (in which case it will of course run with their privileges). The server should NOT be `setuid`. On some systems it may be worthwhile to make `smbd` `setgid` to an empty group. This is because some systems may have a security hole where daemon processes that become a user can be attached to with a debugger. Making the `smbd` file `setgid` to an empty group may prevent this hole from being exploited. This security hole and the suggested fix has only been confirmed on old versions (pre-kernel 2.0) of Linux at the time this was written. It is possible that this hole only exists in Linux, as testing on other systems has thus far shown them to be immune.

The server log files should be put in a directory readable and writeable only by root, as the log files may contain sensitive information.

The configuration file should be placed in a directory readable and writeable only by root, as the configuration file controls security for the services offered by the server. The configuration file can be made readable by all if desired, but this is not necessary for correct operation of the server and is not recommended. A sample configuration file "smb.conf.sample" is supplied with the source to the server—this may be renamed to "smb.conf" and modified to suit your needs.

The remaining notes will assume the following:

- **smbd** (the server program) installed in `/usr/local/samba/bin`
- **smb.conf** (the configuration file) installed in `/usr/local/samba/lib`
- log files stored in `/var/adm/smblogs`

The server may be run either as a daemon by users or at startup, or it may be run from a meta-daemon such as `inetd` upon request. If run as a daemon, the server will always be ready, so starting sessions will be faster. If run from a meta-daemon some memory will be saved and utilities such as the `tcpd` TCP-wrapper may be used for extra security. For serious use as file server it is recommended that **smbd** be run as a daemon.

When you've decided, continue with either **RUNNING THE SERVER AS A DAEMON** or **RUNNING THE SERVER ON REQUEST**.

Running the Server as a Daemon

To run the server as a daemon from the command line, simply put the **-D** option on the command line. There is no need to place an ampersand at the end of the command line—the **-D** option causes the server to detach itself from the tty anyway.

Any user can run the server as a daemon (execute permissions permitting, of course). This is useful for testing purposes, and may even be useful as a temporary substitute for something like ftp. When run this way, however, the server will only have the privileges of the user who ran it.

To ensure that the server is run as a daemon whenever the machine is started, and to ensure that it runs as root so that it can serve multiple clients, you will need to modify the system startup files. Wherever appropriate (for example, in `/etc/rc`), insert the following line, substituting port number, log file location, configuration file location and debug level as desired:

```
/usr/local/samba/bin/smbd -D -l /var/adm/smblogs/log -s  
/usr/local/samba/lib/smb.conf
```

(The above should appear in your initialization script as a single line. Depending on your terminal characteristics, it may not appear that way in this man page. If the above appears as more than one line, please treat any newlines or indentation as a single space or TAB character.)

If the options used at compile time are appropriate for your system, all parameters except **-D** may be omitted. See the section **OPTIONS** above.

Running the Server on Request

If your system uses a meta-daemon such as `inetd`, you can arrange to have the `smbd` server started whenever a process attempts to connect to it. This requires several changes to the startup files on the host machine. If you are experimenting as an ordinary user rather than as root, you will need the assistance of your system administrator to modify the system files.

You will probably want to set up the NetBIOS name server **nmbd** at the same time as **smbd**. To do this refer to the man page for **nmbd (8)**.

First, ensure that a port is configured in the file `/etc/services`. The well-known port 139 should be used if possible, though any port may be used.

Ensure that a line similar to the following is in `/etc/services`:

```
netbios-ssn 139/tcp
```

Note for NIS/YP users—you may need to rebuild the NIS service maps rather than alter your local `/etc/services` file.

Next, put a suitable line in the file `/etc/inetd.conf` (in the unlikely event that you are using a meta-daemon other than `inetd`, you are on your own). Note that the first item in this line matches the service name in `/etc/services`. Substitute appropriate values for your system in this line (see **inetd (8)**):

```
netbios-ssn stream tcp nowait root /usr/local/samba/bin/smbd  
-dl -l/var/adm/smblogs/log -s/usr/local/samba/lib/smb.conf
```

(The above should appear in `/etc/inetd.conf` as a single line. Depending on your terminal characteristics, it may not appear that way in this man page. If the above appears as more than one line, please treat any newlines or indentation as a single space or TAB character.)

Note that there is no need to specify a port number here, even if you are using a nonstandard port number.

Lastly, edit the configuration file to provide suitable services. To start with, the following two services should be all you need:

```
[homes]
writeable = yes
[printers]
writeable = no
printable = yes
path = /tmp
public = yes
```

This will allow you to connect to your home directory and print to any printer supported by the host (user privileges permitting).

Testing the Installation

If running the server as a daemon, execute it before proceeding. If using a meta-daemon, either restart the system or kill and restart the meta-daemon. Some versions of inetd will reread their configuration tables if they receive a HUP signal.

If your machine's name is "fred" and your name is "mary", you should now be able to connect to the service `\\fred\mary`.

To properly test and experiment with the server, we recommend using the smbclient program (see **smbclient (1)**) and also going through the steps outlined in the file *DIAGNOSIS.txt* in the *docs/* directory of your Samba installation.

Version

This man page is correct for version 2.0 of the Samba suite.

Diagnostics

Most diagnostics issued by the server are logged in a specified log file. The log file name is specified at compile time, but may be overridden on the command line.

The number and nature of diagnostics available depends on the debug level used by the server. If you have problems, set the debug level to 3 and peruse the log files.

Most messages are reasonably self-explanatory. Unfortunately, at the time this man page was created, there are too many diagnostics available in the source code to warrant describing each and every diagnostic. At this stage your best bet is still to grep the source code and inspect the conditions that gave rise to the diagnostics you are seeing.

Signals

Sending the smbd a SIGHUP will cause it to re-load its smb.conf configuration file within a short period of time.

To shut down a users smbd process it is recommended that SIGKILL (-9) *NOT* be used, except as a last resort, as this may leave the shared memory area in an inconsistent state. The safe way to terminate an smbd is to send it a SIGTERM (-15) signal and wait for it to die on its own.

The debug log level of smbd may be raised by sending it a SIGUSR1 (`kill -USR1 <smbd-pid>`) and lowered by sending it a SIGUSR2 (`kill -USR2 <smbd-pid>`). This is to allow transient problems to be diagnosed, whilst still running at a normally low log level.

Note that as the signal handlers send a debug write, they are not re-entrant in smbd. This you should wait until smbd is in a state of waiting for an incoming smb before issuing

them. It is possible to make the signal handlers safe by un-blocking the signals before the select call and re-blocking them after, however this would affect performance.

See Also

hosts_access (5), **inetd (8)**, **nmbd (8)**, **smb.conf (5)**, **smbclient (1)**, **testparm (1)**, **test-prns (1)**, and the Internet RFC's **rfc1001.txt**, **rfc1002.txt**. In addition the CIFS (formerly SMB) specification is available as a link from the Web page : <http://samba.org/cifs/>.

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@samba.org.

See **samba (7)** to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

smbpasswd (5)

Samba

23 Oct 1998

Name

smbpasswd-The Samba encrypted password file

Synopsis

smbpasswd is the Samba encrypted password file.

Description

This file is part of the **Samba** suite.

smbpasswd is the **Samba** encrypted password file. It contains the username, Unix user id and the SMB hashed passwords of the user, as well as account flag information and the time the password was last changed. This file format has been evolving with Samba and has had several different formats in the past.

File Format

The format of the smbpasswd file used by Samba 2.0 is very similar to the familiar Unix **passwd (5)** file. It is an ASCII file containing one line for each user. Each field within each line is separated from the next by a colon. Any entry beginning with # is ignored. The smbpasswd file contains the following information for each user:

- **name** This is the user name. It must be a name that already exists in the standard UNIX passwd file.

- **uid** This is the UNIX uid. It must match the uid field for the same user entry in the standard UNIX passwd file. If this does not match then Samba will refuse to recognize this **smbpasswd** file entry as being valid for a user.
- **Lanman Password Hash** This is the *LANMAN* hash of the users password, encoded as 32 hex digits. The *LANMAN* hash is created by DES encrypting a well known string with the users password as the DES key. This is the same password used by Windows 95/98 machines. Note that this password hash is regarded as weak as it is vulnerable to dictionary attacks and if two users choose the same password this entry will be identical (i.e. the password is not "salted" as the UNIX password is). If the user has a null password this field will contain the characters "NO PASSWORD" as the start of the hex string. If the hex string is equal to 32 "X" characters then the users account is marked as *disabled* and the user will not be able to log onto the Samba server.

WARNING !!. Note that, due to the challenge-response nature of the SMB/CIFS authentication protocol, anyone with a knowledge of this password hash will be able to impersonate the user on the network. For this reason these hashes are known as "*plain text equivalent*" and must *NOT* be made available to anyone but the root user. To protect these passwords the **smbpasswd** file is placed in a directory with read and traverse access only to the root user and the **smbpasswd** file itself must be set to be read/write only by root, with no other access.

- **NT Password Hash** This is the *Windows NT* hash of the users password, encoded as 32 hex digits. The *Windows NT* hash is created by taking the users password as represented in 16-bit, little-endian UNICODE and then applying the *MD4* (internet rfc1321) hashing algorithm to it.

This password hash is considered more secure than the **Lanman Password Hash** as it preserves the case of the password and uses a much higher quality hashing algorithm. However, it is still the case that if two users choose the same password this entry will be identical (i.e. the password is not "salted" as the UNIX password is).

WARNING !!. Note that, due to the challenge-response nature of the SMB/CIFS authentication protocol, anyone with a knowledge of this password hash will be able to impersonate the user on the network. For this reason these hashes are known as "*plain text equivalent*" and must *NOT* be made available to anyone but the root user. To protect these passwords the **smbpasswd** file is placed in a directory with read and traverse access only to the root user and the **smbpasswd** file itself must be set to be read/write only by root, with no other access.

- **Account Flags** This section contains flags that describe the attributes of the users account. In the **Samba2.0** release this field is bracketed by '[' and ']' characters and is always 13 characters in length (including the '[' and ']' characters). The contents of this field may be any of the characters.
- **"U"** This means this is a "*User*" account, i.e. an ordinary user. Only **User** and **Workstation Trust** accounts are currently supported in the **smbpasswd** file.
- **"N"** This means the account has no password (the passwords in the fields **Lanman Password Hash** and **NT Password Hash** are ignored). Note that this will only allow users to log on with no password if the **null passwords** parameter is set in the **smb.conf (5)** config file.
- **"D"** This means the account is disabled and no SMB/CIFS logins will be allowed for this user.

- **"W"** This means this account is a *"Workstation Trust"* account. This kind of account is used in the Samba PDC code stream to allow Windows NT Workstations and Servers to join a Domain hosted by a Samba PDC.

Other flags may be added as the code is extended in future. The rest of this field space is filled in with spaces.

- **Last Change Time** This field consists of the time the account was last modified. It consists of the characters LCT- (standing for "Last Change Time") followed by a numeric encoding of the UNIX time in seconds since the epoch (1970) that the last change was made.
- **Following fields** All other colon separated fields are ignored at this time.

Notes

In previous versions of Samba (notably the 1.9.18 series) this file did not contain the **Account Flags** or **Last Change Time** fields. The Samba 2.0 code will read and write these older password files but will not be able to modify the old entries to add the new fields. New entries added with **smbpasswd (8)** will contain the new fields in the added accounts however. Thus an older **smbpasswd** file used with Samba 2.0 may end up with some accounts containing the new fields and some not.

In order to convert from an old-style **smbpasswd** file to a new style, run the script **convert_smbpasswd**, installed in the Samba bin/ directory (the same place that the **smbd** and **nmbd** binaries are installed) as follows:

```
cat old_smbpasswd_file | convert_smbpasswd > new_smbpasswd_file
```

The `convert_smbpasswd` script reads from stdin and writes to stdout so as not to over-write any files by accident.

Once this script has been run, check the contents of the new **smbpasswd** file to ensure that it has not been damaged by the conversion script (which uses `awk`), and then replace the <old **smbpasswd** file> with the <new **smbpasswd** file>.

Version

This man page is correct for version 2.0 of the Samba suite.

See Also

smbpasswd (8), **samba (7)**, and the Internet RFC1321 for details on the MD4 algorithm.

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison, samba-bugs@samba.org.

See **samba (7)** to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

smbpasswd (8)

Samba

23 Oct 1998

Name

smbpasswd-change a users SMB password

Synopsis

smbpasswd [-a] [-d] [-e] [-D debug level] [-n] [-r remote_machine] [-R name resolve order] [-m] [-j DOMAIN] [-U username] [-h] [-s] username

Description

This program is part of the **Samba** suite.

The **smbpasswd** program has several different functions, depending on whether it is run by the *root* user or not. When run as a normal user it allows the user to change the password used for their SMB sessions on any machines that store SMB passwords.

By default (when run with no arguments) it will attempt to change the current users SMB password on the local machine. This is similar to the way the **passwd (1)** program works.

smbpasswd differs from how the **passwd** program works however in that it is not *setuid root* but works in a client-server mode and communicates with a locally running **smbd**.

As a consequence in order for this to succeed the **smbd** daemon must be running on the local machine. On a UNIX machine the encrypted SMB passwords are usually stored in the **smbpasswd (5)** file.

When run by an ordinary user with no options. **smbpasswd** will prompt them for their old smb password and then ask them for their new password twice, to ensure that the new password was typed correctly. No passwords will be echoed on the screen whilst being typed. If you have a blank smb password (specified by the string "NO PASSWORD" in the **smbpasswd** file) then just press the <Enter> key when asked for your old password.

smbpasswd can also be used by a normal user to change their SMB password on remote machines, such as Windows NT Primary Domain Controllers. See the **(-r)** and **-U** options below.

When run by root, **smbpasswd** allows new users to be added and deleted in the **smbpasswd** file, as well as allows changes to the attributes of the user in this file to be made. When run by root, **smbpasswd** accesses the local **smbpasswd** file directly, thus enabling changes to be made even if **smbd** is not running.

Options

- **-a** This option specifies that the username following should be added to the local **smbpasswd** file, with the new password typed (type <Enter > for the old password). This option is ignored if the username following already exists in the **smbpasswd** file and it is treated like a regular change password command. Note that the user to be added **must** already exist in the system password file (usually */etc/passwd*) else the request to add the user will fail.

This option is only available when running **smbpasswd** as root.

- **-d** This option specifies that the username following should be *disabled* in the local **smbpasswd** file. This is done by writing a "D" flag into the account control space in the **smbpasswd** file. Once this is done all attempts to authenticate via SMB using this username will fail.

If the **smbpasswd** file is in the "old" format (pre-Samba 2.0 format) there is no space in the users password entry to write this information and so the user is disabled by writing "X" characters into the password space in the **smbpasswd** file. See **smbpasswd (5)** for details on the "old" and new password file formats.

This option is only available when running **smbpasswd** as root.

- **-e** This option specifies that the username following should be *enabled* in the local **smbpasswd** file, if the account was previously disabled. If the account was not dis-abled this option has no effect. Once the account is enabled then the user will be able to authenticate via SMB once again.

If the **smbpasswd** file is in the "old" format then **smbpasswd** will prompt for a new password for this user, otherwise the account will be enabled by removing the "D" flag from account control space in the **smbpasswd** file. See **smbpasswd (5)** for details on the "old" and new password file formats.

This option is only available when running **smbpasswd** as root.

- **-D debuglevel** debuglevel is an integer from 0 to 10. The default value if this para-meter is not specified is zero.

The higher this value, the more detail will be logged to the log files about the activities of **smbpasswd**. At level 0, only critical errors and serious warnings will be logged.

Levels above 1 will generate considerable amounts of log data, and should only be used when investigating a problem. Levels above 3 are designed for use only by developers and generate HUGE amounts of log data, most of which is extremely cryptic.

- **-n** This option specifies that the username following should have their password set to null (i.e. a blank password) in the local **smbpasswd** file. This is done by writing the string "NO PASSWORD" as the first part of the first password stored in the **smbpasswd** file.

Note that to allow users to logon to a Samba server once the password has been set to "NO PASSWORD" in the **smbpasswd** file the administrator must set the following parameter in the [global] section of the **smb.conf** file:

```
null passwords = true
```

This option is only available when running **smbpasswd** as root.

- **-r remote machine name** This option allows a user to specify what machine they wish to change their password on. Without this parameter **smbpasswd** defaults to the local host. The "*remote machine name*" is the NetBIOS name of the SMB/CIFS server to contact to attempt the password change. This name is resolved into an IP address using the standard name resolution mechanism in all programs of the **Samba** suite. See the **-R name resolve order** parameter for details on changing this resolving mechanism.

The username whose password is changed is that of the current UNIX logged on user. See the **-U username** parameter for details on changing the password for a different username.

Note that if changing a Windows NT Domain password the remote machine specified must be the Primary Domain Controller for the domain (Backup Domain Controllers only have a read-only copy of the user account database and will not allow the password change).

Note that Windows 95/98 do not have a real password database so it is not possible to change passwords specifying a Win95/98 machine as remote machine target.

- **-R name resolve order** This option allows the user of `smbclient` to determine what name resolution services to use when looking up the NetBIOS name of the host being connected to.

The options are `:"lmhosts"`, `"host"`, `"wins"` and `"bcast"`. They cause names to be resolved as follows:

- **lmhosts:** Lookup an IP address in the Samba `lmhosts` file.
- **host:** Do a standard host name to IP address resolution, using the system `/etc/hosts`, NIS, or DNS lookups. This method of name resolution is operating system dependent. For instance on IRIX or Solaris, this may be controlled by the `/etc/nsswitch.conf` file).
- **wins:** Query a name with the IP address listed in the **wins server** parameter in the **smb.conf** file. If no WINS server has been specified this method will be ignored.
- **bcast:** Do a broadcast on each of the known local interfaces listed in the **interfaces** parameter in the `smb.conf` file. This is the least reliable of the name resolution methods as it depends on the target host being on a locally connected subnet.

If this parameter is not set then the name resolve order defined in the **smb.conf** file parameter **name resolve order** will be used.

The default order is `lmhosts`, `host`, `wins`, `bcast` and without this parameter or any entry in the **smb.conf** file the name resolution methods will be attempted in this order.

- **-m** This option tells **smbpasswd** that the account being changed is a *MACHINE* account. Currently this is used when Samba is being used as an NT Primary Domain Controller. PDC support is not a supported feature in Samba2.0 but will become supported in a later release. If you wish to know more about using Samba as an NT PDC then please subscribe to the mailing list samba-ntdom@samba.org.

This option is only available when running **smbpasswd** as root.

- **-j DOMAIN** This option is used to add a Samba server into a Windows NT Domain, as a Domain member capable of authenticating user accounts to any Domain Controller in the same way as a Windows NT Server. See the **security=domain** option in the **smb.conf (5)** man page.

In order to be used in this way, the Administrator for the Windows NT Domain must have used the program *"Server Manager for Domains"* to add the primary NetBIOS name of the Samba server as a member of the Domain.

After this has been done, to join the Domain invoke **smbpasswd** with this parameter. **smbpasswd** will then look up the Primary Domain Controller for the Domain (found in the **smb.conf** file in the parameter **password server** and change the machine account password used to create the secure Domain communication. This password is then stored by **smbpasswd** in a file, read only by root, called `<Domain>.<Machine>.mac` where `<Domain>` is the name of the Domain we are joining and `<Machine>` is the primary NetBIOS name of the machine we are running on.

Once this operation has been performed the `smb.conf` file may be updated to set the **security=domain** option and all future logins to the Samba server will be authenticated to the Windows NT PDC.

Note that even though the authentication is being done to the PDC all users accessing the Samba server must still have a valid UNIX account on that machine.

This option is only available when running **smbpasswd** as root.

- **-U username** This option may only be used in conjunction with the **-r** option. When changing a password on a remote machine it allows the user to specify the user name on that machine whose password will be changed. It is present to allow users who have different user names on different systems to change these passwords.
- **-h** This option prints the help string for **smbpasswd**, selecting the correct one for running as root or as an ordinary user.
- **-s** This option causes **smbpasswd** to be silent (i.e. not issue prompts) and to read it's old and new passwords from standard input, rather than from `/dev/tty` (like the **passwd (1)** program does). This option is to aid people writing scripts to drive **smbpasswd**
- **username** This specifies the username for all of the *root only* options to operate on. Only root can specify this parameter as only root has the permission needed to modify attributes directly in the local **smbpasswd** file.

Notes

Since **smbpasswd** works in client-server mode communicating with a local **smbd** for a non-root user then the **smbd** daemon must be running for this to work. A common problem is to add a restriction to the hosts that may access the **smbd** running on the local machine by specifying a **"allow hosts"** or **"deny hosts"** entry in the **smb.conf** file and neglecting to allow *"localhost"* access to the **smbd**.

In addition, the **smbpasswd** command is only useful if **Samba** has been set up to use encrypted passwords. See the file **ENCRYPTION.txt** in the docs directory for details on how to do this.

Version

This man page is correct for version 2.0 of the Samba suite.

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@samba.org.

See **samba (7)** to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

smbstatus (1)

Samba

23 Oct 1998

Name

smbstatus—report on current Samba connections

Synopsis

smbstatus [-b] [-d] [-L] [-p] [-S] [-s configuration file] [-u username]

Description

This program is part of the **Samba** suite.

smbstatus is a very simple program to list the current Samba connections.

Options

- **-b** gives brief output.
- **-d** gives verbose output.
- **-L** causes smbstatus to only list locks.
- **-p** print a list of **smbd** processes and exit. Useful for scripting.
- **-S** causes smbstatus to only list shares.
- **-s configuration file** The default configuration file name is determined at compile time. The file specified contains the configuration details required by the server. See **smb.conf (5)** for more information.
- **-u username** selects information relevant to *username* only.

Version

This man page is correct for version 2.0 of the Samba suite.

See Also

smb.conf (5), **smbd (8)**

Author

The original Samba software and related utilities were created by Andrew Tridgell samba-bugs@samba.org. Samba is now developed by the Samba Team as an Open Source project similar to the way the Linux kernel is developed.

The original Samba man pages were written by Karl Auer. The man page sources were converted to YODL format (another excellent piece of Open Source software, available at <ftp://ftp.icce.rug.nl/pub/unix/>) and updated for the Samba2.0 release by Jeremy Allison. samba-bugs@samba.org.

See **samba (7)** to find out how to get a full list of contributors and details on how to submit bug reports, comments etc.

Appendix D. TCP/IP Documentation

Appendix Objectives

- TCP/IP Network Resources List
- Private IP Network addresses

AUTHOR'S NOTE

In this Appendix I have chosen those manual pages that I feel will be most useful to you in working with Samba. There are other commands that you will find useful as you progress in learning Samba.

TCP/IP Network Resources List

The following list is an excellent reference for Internet TCP/IP Resources. It is used with the Uri's permission. For an uptodate list you will want to access the main URL listed below. If you are not very familiar with what kind of resources there are, you will want to look at the breadth of material listed below.

Uri's TCP/IP Resources List FAQs, TUTORIALS, GUIDES, WEB PAGES & SITES, AND BOOKS ABOUT TCP/IP BY URI RAZ

This posting contains a list of various resources (books, web sites, FAQs, newsgroups, and useful net techniques) intended to help a newbie to learn about the TCP/IP suite of protocols.

I have written this document over the last few years. Yet, I could not have made this document without the assistance of other people. I would, therefore, like to thank to Andrew Gierth, Trevor Jenkins, Mark Daugherty, Michael Hunter, David Peter, Erick Engelke, Jose Carrilho, Jose Carrilho, Al Vonkeman, Zia R. Siddiqui, Jarle Aase, Daryl Banttari, Daniel K. Kim, Brian Schwarz, James Marshall, Diane Boling and Gisle Vanem who helped me in many ways, and to all the people who worked to produce all the materials listed in this document.

This article is available as a web page at:

http://www.private.org.il/tcpip_rl.html

http://t2.technion.ac.il/~s2845543/tcpip_rl.html

http://www.best.com/~mphunter/tcpip_resources.html

This article is available via FTP at:

<ftp://rtfm.mit.edu/pub/usenet-by-group/news.answers/internet/tcp-ip/re-source-list>

ftp://rtfm.mit.edu/pub/usenet-by-hierarchy/comp/protocols/tcp-ip/TCP_IP_Re-sources_List

Note

If you have any comments, addition suggestions, corrections, etc, to the article itself, please send them to me at the *technion*. My email address is <mailto:s2845543@t2.technion.ac.il>

There are plenty of copies of this article on the web. Please do not create another one, as when the copies go out of date all the requests to remove dead links, add new links, fix typos, etc which I already did in the latest version go to me.

If you have any questions about TCP/IP in general, which are not directly related to this article, please post them to an appropriate newsgroup, as my time is limited, and as it will serve you better.

WARNING: job offers from outside Israel will be treated as spam.

1 Books About TCP/IP

Richard Stevens' TCP/IP illustrated.
Published by Addison-Wesley.

Volume 1 - describes the TCP/IP protocols. ISBN 0201633469

Volume 2 - describes the TCP/IP stack as implemented in 4.4BSD-Lite, at the source code level. ISBN 020163354X

Volume 3 - describes HTTP, NNTP, and more. ISBN 0201634953

Richard Steven's UNIX Network Programming.
Published by Prentice Hall.

Described here is the 2nd edition of the book.

The 1st edition (ISBN 0139498761) will be sold until the third volume of of the 2nd edition will be out.

Volume 1 - "Networking APIs: Sockets and XTI". Describes UNIX network programming in & out, including a lot of code examples, covering IPv4 & IPv6, sockets and XTI, TCP & UDP, raw sockets, programming techniques, multicasting & broadcasting, and what not. The best TCP/IP programming book around, IMHO. ISBN 013490012X

Volume 2 - "Interprocess Communications". ISBN 0130810819

Volume 3 - "Applications" Name is probable, to be published.

Douglas Comer's Internetworking with TCP/IP.
Published by Prentice-Hall.

Volume 1 - describes the TCP/IP protocols, architecture and principles.
ISBN 0132169878

Volume 2 - describes a TCP/IP implementation (with C code), implemented on the XINU operating system. ISBN 0131255274

Volume 3 - describes network programming, and has a sockets version (ISBN 013260969X), a TLI version (ISBN 0132609770), and a winsock version (ISBN 0138487146)

TCP/IP Explained

By Philip Miller

Published by Digital Press

ISBN 1555581668

A fine book about TCP/IP, covering all the layers, starting with an overview of the lowest 2 OSI layers, through IP(+ICMP), UDP, TCP, routing (RIP + OSPF + EGP + BGP), broadcasting and multicasting, DNS, SNMP, several apps (FTP, Telnet, SMTP, ...), with chapters about IPv6 and Internet Security. The book is readable, with lots of diagrams and packet trace decodes. Some points missing, such as TCP congestion avoidance.

Troubleshooting TCP/IP - Analyzing the Protocols of the Internet

By Mark A. Miller

Published by M & T Books

ISBN 1558514503

A good troubleshooting guide, with good explanations of most protocols, starting from network layer, through ARP, DNS, routing, and up to the applications, including SMTP, FTP, and TELNET. Coverage includes SNMP, ATM, IPv6. Case studies, included for every subject, include sniffer output and explanations.

High-Speed Networks: TCP/IP and ATM Design Principles

By William Stallings

Published by Prentice-Hall

ISBN 0135259657

This book explains how to design high-speed networks (ATM, 100 Mbps & Gbps ethernet) intended to carry high volume data (WWW, still images, video on demand, etc). Coverage includes explanation of ATM and Fast & Gigabit Ethernet, the mathematical background needed for performance analysis, traffic management (IP & ATM), routing, and compression.

TCP/IP: Architecture, Protocols, and Implementation with IPv6 and IP Security

By Sidnie Feit

Published by McGraw-Hill

ISBN 0070213895

This book covers TCP/IP in one volume, starting from the physical layer, through IP, UDP & TCP, the various applications (WWW, mail, etc) to network management.

SNMP, SNMPv2, SNMPv3, and RMON1 and RMON2

By William Stallings

Published by Addison-Wesley

ISBN 0201485346

An encyclopedic book about SNMP & RMON. Covers the material in depth and clarity, giving good background of the subject.

Networking with Microsoft TCP/IP

By Drew Heywood

Published by New Riders

ISBN 1562057138

An excellent book about management of Microsoft Windows TCP/IP networks, starting from the basics of explaining networking technologies, through installation of TCP/IP on DOS and all MS Windows versions, routing, managing (DHCP, WINS, DNS), troubleshooting, IIS & FrontPage.

TCP/IP Network Administration

By Craig Hunt.

Published by O'Reilly

ISBN 093717582X

An excellent book about management of TCP/IP networks, covering every subject that needed, including DNS, routing, sendmail, configuring, and trouble-shooting. This book is UNIX oriented.

Networking Personal Computers with TCP/IP - Building TCP/IP Networks

By Craig Hunt

Published by O'Reilly

ISBN 1565921232

A good book about management of TCP/IP networks, which is PC oriented, covering DOS, Windows, Windows-95, and Windows-NT.

Teach Yourself TCP/IP in 14 days.

By Timothy Parker

Published by SAM'S Publishing.

ISBN 0672305496

This book is intended for network managers, and gives an overview of TCP/IP from ground up, in a short schedule.

PPP Design and Debugging

By James Carlson

Published by Addison-Wesley

ISBN 0201185393

An excellent book about PPP. This compact book is packed with info about PPP, covering it in both depth and width, covering LCP, negotiation & authentication, network layer protocols, bandwidth management, etc, including trace interpretation, C code & pseudo code, and lots of resources and references.

NOSintro - TCP/IP over Packet Radio (An Introduction to the KA9Q Network Operating System)

By Ian Wade

Published by Dowermain

ISBN 1897649002

NOSintro describes in detail how to use Phil Karn's KA9Q Network Operating System, and is a classic reference work in this area. It includes full information on how to install & configure KA9Q, and how to make it work in a packet radio environment.

The book is very well illustrated, with many diagrams & hands-on examples of keyboard commands.

Extracts from the book are available at <http://www.netro.co.uk/nosintro.html>

IPv6: The New Internet Protocol

By Christian Huitema

Published by Prentice-Hall.

ISBN 0138505055

This book, written by Christian Huitema - a member of the Internet Architecture Board, gives an excellent description of IPv6, how it differs from IPv4, and the hows and whys of it's development.

Unix Network Programming

By W. Richard Stevens

Published by Prentice-Hall

ISBN 0139498761

Obsoleted by the second edition, to be covered soon.

Unix System V. Network Programming

By Steven A. Rago

Published by Addison-Wesley

ISBN 0201563185

This books gives a good coverage of UNIX network programming. Though it is centered around SVR4, it covers many subjects, including STREAMS, TLI, sockets, RPC, and kernel level communications, including ethernet & SLIP drivers.

The Design and Implementation of the 4.4 BSD Operating System.

By Marshall Kirk McKusick, Keith Bostic, Michael J. Karels and John S. Quarterman.

Published by Addison-Wesley.

ISBN 0201549794

This book describes the internals of the 4.4 BSD operating system, including the Net/2 TCP/IP stack implementation. A good explanation of the most commonly used implementation of TCP/IP.

Linux Kernel Internals

By M. Beck, H. Bohme, M. Dziadzka, U. Kunitz, R. Magnus, and D. Verworner.

Published by Addison-Wesley

ISBN 0201331438

This book describes the internals of the Linux operating system, version 2.0, with a chapter devoted to the TCP/IP stack.

Windows Sockets Network Programming

By Bob Quinn and Dave Shute

Published by Addison-Wesley

ISBN 0201633728

An excellent book about winsock programming, with chapters about porting apps from BSD Unix & sockets, DLLs, debugging, and nice appendice.

The two following books are not directly related to TCP/IP, but are recommended as good books for windows programmer who write TCP/IP clients & servers, and are complementary to the above book:

1. Win32 Network Programming

By Ralph Davis

Published by Addison-Wesley

ISBN 0201489309

This book shows programmers how to build networked apps using the 32-bit features of Win95 and NT, and includes a floppy with all the examples' code.

2. Multithreading Applications in Win32

By Jim Beveridge and Robert Wiener

Published by Addison-Wesley

ISBN 0201442345

This book shows developers how, when and where to use multi-threading in Win32 applications, and includes a CD-ROM.

Interconnections

By Radia Perlman

Published by Addison-Wesley

ISBN 0201563320

This is a good book about bridging and routing, which has both a wide coverage and a technical depth. The book covers TCP/IP routing in only one chapter, which is extensive, but gives a much wider perspective on bridging, brouting, and routing in general.

Routing in the Internet

By Christian Huitema

Published by Prentice Hall

ISBN 0131321927

A clear and thorough, though a bit dated, book about routing. Covers all major routing protocols (RIP, OSPF, IGRP & EIGRP, IS-IS, EGP, BGP3, BGP4 & CIDR), and covers multicast, mobility, and resource reservation.

Internet Routing Architectures

By Bassam Halabi
Published by Cisco Press
ISBN 1562056522

A clear and through book about interdomain routing network design, with many clear examples with diagrams. Focuses on BGP4 and is, naturally, oriented toward Cisco's way of doing it (which is not much of a limit, considering Cisco's dominance of the routers market).

OSPF, Anatomy of an Internet Routing Protocol

By John T. Moy
Published by Addison-Wesley
ISBN 0201634724

A great book about OSPF, including it's history, multicast routing, management, debugging, comparisons to other routing protocols, and the companion book (OSPF Complete Implementation) goes through a complete implementation of OSPF (included on a CD), with a port to FreeBSD 2.1 and a Windows-95 simulator.

BGP4

By John W. Stewart III
Published by Addison-Wesley
ISBN 0201379511

A small (<150 pages) book, covering BGP4 in full using clear language and drawings. The four chapters include an introduction, the protocol, operations, and extensions (scaling, route flap dampening, authentication, negotiation, etc).

Data and Computer Communications

By William Stallings
Published by Prentice-Hall.
ISBN 0024154253

A very good book about computer communications basics. Includes information about TCP/IP and IPv6.

Computer Networks

By Andrew S. Tanenbaum
Published by Prentice-Hall.
ISBN 0133499456

A very good book about computer communications basics. Describes communications according to the OSI seven layers model, but includes information about TCP/IP and IPv6.

Information Warfare and Security

By Dorothy E. Denning
Published by Addison-Wesley
ISBN 0201433036

A book covering all aspects of information warfare with clear explanations and many references. Gives an excellent framework to Internet security.

2 Major On-Line Resources

1. The *IETF's* home page is <http://www.ietf.org/>

This is the authoritative source for RFCs (which include all the standards for TCP/IP), FYIs, drafts, and other infos about the internet and TCP/IP.

A good place to look for RFCs is the *Kashpureff Family's* site, at <http://www.kashpureff.org/nic/>, which has a copy of all RFCs and drafts, as well as a search engine to search for keywords through either RFCs or drafts.

Cabletron has a repository of RFCs and drafts. Drafts are indexed by subject, while STDs & RFCs by title & number. A search engine is supplied to search through titles or bodies.

Another source for RFCs is rfc-info@isi.edu - to get further info, send a message with any subject, and with the body having one line, containing either "help", or "help: ways_to_get_rfc".

Note

the RFCs are the documents giving the official documentation to the various internet protocols. For specs / description / details / info about any internet protocol, first look at the Kashpureff Family site or get the RFCs index via email.

An excellent index of RFCs is available in an appendix in Comer's first volume, but it is current as of the publishing date only.

Comment: as many people seem to look for RFCs on CD-ROMs, I list here two titles I know of:

1. Infomagic has a 2 CDs set titled "STANDARDS" which contains, among other things, all the RFCs & IENs.
2. Walnut-Creek has a CD-ROM titled "Internet Info" which contains some of the RFCs & IENs, among other stuff.

Network Research Group home page - <http://www-nrg.ee.lbl.gov/nrg.html>

Internet Assigned Numbers Authority home page - <http://www.iana.org/>

Internet Engineering Task Force home page - <http://www.ietf.org/>

Internet Research Task Force home page - <http://www.irtf.org/>

Internet Societal Task Force home page - <http://www.iskf.isoc.org/>

Internet SOCIety home page - <http://www.isoc.org/>

Internet Architecture Board home page - <http://www.iab.org/>

Internet Engineering Steering Group - <http://www.ietf.org/iesg.html>

Internet Engineering & Planning Group - <http://www.iepg.org/>

Internet Mail Consortium - <http://www.imc.org/>

The Generic Top Level Domain Memorandum of Understanding - <http://www.gtld-mou.org/>

Internet Ad-Hoc Committee home page - <http://www.iahc.org/>

ICANN - The Internet Corporation for Assigned Names and Numbers - <http://www.icann.org/>

ICANN Watch - <http://www.icannwatch.org/>

Open Root Server Confederation - <http://www.open-rsc.org/>

RFC editor's web page - <http://www.rfc-editor.org/>

Overview of the DNS Controversy - <http://www.flywheel.com/ircw/overview.html>

Another article by Robert Shaw - <http://www.itu.int/intreg/dns.html>

The *National Telecommunications and Information Administration's* Proposals for Management of Internet Names and Addresses page.

<http://www.ntia.doc.gov/ntiahome/domainname/domainhome.htm>

The *AlterNIC's* home page is <http://www.alternic.com/>

This site carries RFCs, internet drafts, and materials relating to freedom of speech, encryption, and more.

The *FAQs.org* sites carries RFCs, STDs, and FYIs. Subject indices are available, and RFCs can be viewed by active/all basis, and with several levels of details. Page's URL is <http://www.faqs.org/rfcs/>

2. The comp.answers & news.answers newsgroups contain (or at least should) all FAQ postings for the newsgroups dealing with computers.

The following newsgroups contain discussion related to TCP/IP:

- Newsgroups FAQs are posted periodically to their top-hierarchy answers newsgroup (e.g. comp.os.vms => comp.answers). Those groups, along with news.newusers.questions, are great places to look for FAQs & tips in.
- the comp.protocols hierarchy, which covers various networking protocols, such as tcp/ip, kermi, and iso.

notice that some TCP/IP related protocols have discussion groups of their own (e.g. NFS, SNMP, NTP, PPP).

- the comp.dcom hierarchy, including groups that discuss lans, modems, and ethernet.
 - the comp.mail hierarchy, which covers various electronic mail programs (pine, elm, sendmail, etc).
 - The news hierarchy, which covers the various subjects related to usenet, including the NNTP protocol.
3. All the newsgroups' FAQs, as well as other introductory documents are stored at <ftp://rtfm.mit.edu/pub/>. A good introductory to TCP/IP from the site is the file

<ftp://rtfm.mit.edu/pub/net/internet.text>. The FAQs can be accessed on the web at <http://www.faqs.org/> as well.

As the rtfm.mit.edu & faqs.org sites might be heavily loaded, and as many sites mirror the FAQs archive, it is advisable to search for FAQs at geographically nearer sites. A list of many mirror sites (allowing access via FTP, WWW, Gopher, mail, etc) is available at: <ftp://rtfm.mit.edu/pub/faqs/news-answers/introduction>

The comp.protocols.tcp-ip group has a FAQ, previously maintained by George V. Neville-Neil, now by Mike Oliver, is located at:

<ftp://rtfm.mit.edu/pub/faqs/internet/tcp-ip/tcp-ip-faq/>

<http://www.itprc.com/tcpipfaq/default.htm>

<http://t2.technion.ac.il/~s2845543/tcpip-faq/default.htm>

The comp.protocols.tcp-ip.ibmpc newsgroup has a FAQ, written by Bernard D. Aboba, which can be found at at:

<ftp://ftp.netcom.com/pub/ma/mailcom/IBMTCP/ibmtcp.zip>

<http://www.inetassist.com/faqs/tcpibmpc.htm>

The comp.protocols.tcp-ip.domains newsgroup has a FAQ, maintained by Chris Peckham, which can be found at:

<http://www.users.pfmc.net/~cdp/cptd-faq/>

<ftp://rtfm.mit.edu/pub/usenet/news.answers/internet/tcp-ip/domains-faq/>

The sockets programming FAQ, by Vic Metcalfe, is located at:

<ftp://rtfm.mit.edu/pub/usenet/news.answers/unix-faq/socket>

<http://www.faqs.org/faqs/unix-faq/socket/index.html>

The alt.winsock newsgroup has a FAQ, by Nancy Cedeno Alegria, located at:

<http://www.well.com/user/nac/alt-winsock-faq.html>

<http://www.faqs.org/faqs/windows/winsock-faq/index.html>

<ftp://rtfm.mit.edu/pub/usenet/news.answers/windows/winsock-faq>

The Winsock Programmer's FAQ, by Warren Young, is located at:

<http://www.cyberport.com/~tangent/programming/winsock/>

<http://www.faqs.org/faqs/windows/winsock/programmer-faq/index.html>

<ftp://rtfm.mit.edu/pub/usenet/news.answers/windows/winsock/programmer-faq>

The windows-sockets page, by Bob Quinn, is located at: <http://www.sockets.com/>

The *sockaddr.com* - Programming Resources for WinSock site, is located at: <http://www.sockaddr.com/>

The Raw IP Networking FAQ, by Thamer Al-Herbish, is available at: <http://www.whitefang.com/rin/>

Stardust has winsock pages, located at:

<http://www.stardust.com/wsresource/wsresrce.html>

<http://www.winsock.com/>

Wandel & Goltermann have brought up the *decodes.com*

The size lists is intended to be a "Resource for Network Protocol Analysis". <http://www.decodes.com/>

The Secure Sockets Layer Discussion List FAQ is located at:

<http://www.consensus.com/security/ssl-talk-faq.html>

<ftp://ftp.consensus.com/pub/security/ssl-talk-faq.txt>

Info about *Ssh (Secure Shell)* may be found at:

<http://www.ssh.org/>

<http://www.cs.hut.fi/ssh/>

<http://www.faqs.org/faqs/computer-security/ssh-faq/index.html>

<ftp://rtfm.mit.edu/pub/usenet-by-group/comp.security.ssh/>

Info about *SOCKS (secure sockets using proxies / firewalls)* -

http://www.socks.nec.com/ftp://coast.cs.purdue.edu/pub/doc/faq/faq_socks

The *DNS Resources Directory*, an excellent resource, may be found at -

<http://www.dns.net/dnsrd/>

Jarle Aase's FTP Protocol Resource Center site may be found at -

<http://war.igaa.com:8080/ftp/>

Info about various *TCP/IP protocols* originating from UNIX utilities, such as *r** services, *lpd*, and *talk*, can be found in a page I've written up for the purpose of concentrating the info at a single point. <http://t2.technion.ac.il/~s2845543/mini-tcpip.faq.html>

4. The *comp.protocols.tcp-ip.ibmcp* is gated to a mailing list as well, and it is served by listserv@list.nih.gov, under the name PCIP.

The alt.winsock newsgroup is gated to a mailing list as well.

The mailing list is named <http://www.winsock@microdyne.com>. The [un]subscribe address is, of course, <http://www.winsock-request@microdyne.com>

There's an IPv6 mailing list. It's named ipng, and it is served served by <http://www.Majordomo@sunroof.eng.sun.com>

5. RFC #1180 (RFC1180), titled "A TCP/IP Tutorial", is a good tutorial, with a focus on how an IP packet travels from source to destination. RFC #2151 (FYI30), titled "A Primer On Internet and TCP/IP Tools" is a good introductory to TCP/IP tools, such as ping, finger, and traceroute.

3 Misc Web Pages

1. The *Unix Guru Universe's* home page is <http://www.ugu.com/> You could find in this site references to all kinds of info relating to UNIX, including TCP/IP.

There are three great sites for all of *MS-Windows's* versions, which cover a lot of info relating to connecting MS-Windows to TCP/IP networks.

The sites are:

<http://www.windows.com/>

<http://www.windows-95.com/>

<http://www.windows98.org/>

<http://www.support.microsoft.com/> [requires registration]

The *Network Professionals Resource Center* contains links to many FAQs, computers & networking magazines' home pages, etc. <http://www.inetassist.com/>

The *Network Management Server* carries FAQs, white papers, free software, etc related to network management. <http://www.netman.cit.buffalo.edu/>

The *Direct Cable Connection, Null-modem, Serial Ports* site explains how to connect two windows machines to each other using serial or parallel ports to create a two nodes network. <http://www.php.indiana.edu/~jrricha/dcc1.html>

2. You can find many books on the web:
 1. Macmillan's Personal Bookshelf <http://www.mcp.com/personal/>
 2. McGraw-Hill's BetaBooks <http://www.pbg.mcgraw-hill.com/betabooks/>
 3. National Academy Press's Reading Room <http://www.nap.edu/info/browse.htm>
 4. The *Network Administrators' Guide* By Olaf Kirch <http://www.sunsite.unc.edu/mdw/LDP/nag/nag.html>
 5. *Computer Networks and Internets* By Douglas E. Comer <http://www.netbook.cs.purdue.edu/>

Books related pages:

6. The *Xinu BUG* Page at the University of Canberra, Australia. <http://www.willow.canberra.edu.au/~chrisc/bugs.html>

7. List of enhancements to *Comer's TCP code* by Simon Ilyushchenko
<http://www.internasoft.com/simon/tcp/>

You can find on-line networking magazines:

8. *Network Magazine*<http://www.networkmagazine.com/>
9. *Data Communications*<http://www.data.com/>
10. *Network Computing*<http://www.networkcomputing.com/>

A copy of "Netizens: An Anthology" is available at

11. <http://www.columbia.edu/~rh120/> - HTML
12. <ftp://wuarchive.wustl.edu/doc/misc/acn/netbook/> - ASCII
3. The following links would supply intro info on TCP/IP:
 1. gopher://gopher-chem.ucdavis.edu/11/Index/Internet_aw/
 2. *Optimized Engineering Technical Compendium* (LANs & IP)
<http://www.optimized.com/COMPENDI/>
 3. Introduction to
TCP/IP<http://www.pclt.cis.yale.edu/pclt/COMM/TCPIP.HTM>
 4. *Introduction to the Internet
Protocols*<http://www.oac3.hsc.uth.tmc.edu/staff/snewton/tcp-tutorial/>
 5. Under the hood of the 'net: An overview of the TCP/IP Protocol Suite, By
Jason Yanowitz. <http://www.info.acm.org/crossroads/xrds1-1/tcpjmy.html>
 6. IP overview, by Cisco.
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ip.htm
 7. *Tech-NIC's technical page* <http://www.tech-nic.dk/html/technical.html>
 8. *Thomas's Technical Links*<http://www.psp.demon.co.uk/tfl/techlinks.htm>
 9. Several nice tutorials from *Scan Technologies*<http://www.scan-technologies.com/tutorials.htm>
 10. *The IP Address and Classes*<http://www.sangoma.com/fguide.htm> (linked
from <http://www.sangoma.com/tutorial.htm>)
 11. Subnetting:

What's A Netmask?

<http://www.digitalmx.com/wires/subnet.html>

IP Address Subnetting

Tutorial<http://www.ziplink.net/~ralphb/IPSubnet/index.html>

IP Subnet Calculations

<http://www.swcp.com/~jgentry/topo/unit3.htm>

Daryl's TCP/IP Primer

Addressing and Subnetting on the Near Side of the 'Net

<http://ipprimer.windsorcs.com/>

Breeze Through Subnet Masking, by John Lambert,
MCSE

<http://support.wrq.com/tutorials/tcpip/tcpipfundamentals.html>

Al Vokeman's netmask calculator

The calculator is implemented via JavaScript (not CGI), making it quick, but requires JavaScript supported and enabled.

<http://www.telusplanet.net/public/sparkman/netcalc.htm>

Another CIDR subnet mask calculator can be found at

http://minnie.cs.adfa.edu.au/Gateways/range_check.html

12. A *CIDR FAQ*<http://www.rain.net/faqs/cidr.faq.html>
13. *Cisco's Internetworking Terms and Acronyms*<http://www.cio-sys.cisco.com/univercd/data/doc/cintrnet/ita.htm>
14. "TCP/IP Tutorial and Technical Overview" from IBM
<http://www.publib.boulder.ibm.com/cgi-bin/bookmgr/BOOKS/GG243376/CCONTENTS>
15. An Overview of TCP/IP Protocols and the Internet
<http://www.hill.com/library/staffpubs/tcpip.html>
16. *IP Addressing Fundamentals*<http://www.support.wrq.com/tutorials/tcpip/tcpipfundamentals.html>
17. Understanding IP Addressing: Everything You Ever Wanted To Know
<http://www.3com.com/nsc/501302.html>
18. Understanding IP Addressing
<http://www.noc.gate.net/doclib/faqs/help/net.html>
19. *hedrick-intro to the Internet Protocols*<http://www.duth.gr/InfoBase/intro.ip.toc.html>
20. A short page about TCP/IP security by Chris Chambers, Justin Dolske, and Jayaraman Iyer. <http://www.cis.ohio-state.edu/~dolske/gradwork/cis694q/>
21. *Von Welch* has a network performance page at
http://www.ncsa.uiuc.edu/People/vwelch/net_perf/

One of the subpages explains *TCP*

windowshttp://www.ncsa.uiuc.edu/People/vwelch/net_perf/tcp_windows.html

22. *Marc Slemko' Path MTU Discovery and Filtering ICMP*
<http://www.worldgate.com/~marcs/mtu/>
23. *Cliff Green's Introduction to Internet Protocols for Newbies*
http://www.halcyon.com/cliffg/uwteach/shared_info/internet_protocols.html
24. *Catalyst's Introduction to TCP/IP Programming*<http://www.catalyst.com/tcpintro.html>
25. *Frank Dekervel's DNS Tutorial*http://www.geocities.com/SiliconValley/Network/4504/dns_tut.txt

Gary Kessler's Setting Up Your Own

DNS<http://www.hill.com/library/staffpubs/dns.html>

The DNS Security Extensions, by Cricket Liu.

<http://www.acmebw.com/papers/dnssec.pdf>

26. *RPC*

http://www.pandonia.canberra.edu.au/OS/I14_1.html

<http://www.ja.net/documents/NetworkNews/Issue44/RPC.html>

<http://www.mmt.bme.hu/~kiss/docs/dce/rpc.html>

<http://glacier.unl.edu/~samal/class/DIST/lectures/rpc.html>

27. DHCP

Ralph Droms' DHCP Resources site <http://www.dhcp.org/>

Alan Dobkin's DHCP Resources<http://nws.cc.emory.edu/webstaff/alan/net-man/computing/dhcp/>

28. BSD socket programming tutorials

Quick - <http://ftp.std.com/homepages/jimf/sockets.html>

Intro - http://ccnqa.uwaterloo.ca/~mvlioy/stuff/ipc_intro_tut.txt

Advanced - http://ccnqa.uwaterloo.ca/~mvlioy/stuff/ipc_adv_tut.txt

Windows socket programming tutorials

<http://users.neca.com/vmis/wsockexp.htm>

<http://users.neca.com/vmis/wsockprg.htm>

An Introduction to Socket Programming<http://www.uwo.ca/its/doc/courses/notes/socket/index.html>

Beej's Guide to Network Programming<http://www.ecst.csuchico.edu/~beej/guide/net/>

Vijay Mukhi's Winsock Programming page<http://users.neca.com/vmis/wsockprg.htm>

Network Programmer's Guide<http://www.ibr.cs.tu-bs.de/~harbaum/docs/netprog/contents.html>

Unix Network Programminghttp://gaia.cs.umass.edu/ntu_socket/

Spencer's Socket Site<http://www.lowtek.com/sockets/>

29. *Xiaomu's WinSock page*<http://omni.cc.purdue.edu/~xniu/winsock.htm>

30. Routing protocols

In general -

http://www.bind.com/http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/routing.htm

IGRP -

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/igrp.htm

IGRP & Enhanced IGRP -

<http://www.cisco.com/warp/public/103/index.shtml>

IGRP & Enhanced IGRP -

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/en_igrp.htm

RIP -

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/rip.htm

OSPF, BGP, IPv6, GateD -

<http://www.roedu.net/~cmatei/network/>

EGP -

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/5143.htm

BGP4 - <http://www.cisco.com/warp/public/459/18.html>

BGP -

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

OSPF - <http://www.cisco.com/warp/public/104/1.html>

OSPF -

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ospf.htm

OSPF - <http://www.3com.com/nsc/501304.html>

NLSP (Novell) - <http://www.3com.com/nsc/501309.html>

Multi Layer Routing - <http://infonet.aist-nara.ac.jp/member/nori-d/mlr/>

31. Merit GateD Consortium

This site contains wealth of information about *GateD*, including source distributions, documentation, etc. <http://www.gated.org/>

GNU Zebra site

The *GNU Zebra* project is a router software implementing OSPFv2, BGP4, RIPv1, and RIPv2. It has a special architecture that differs from

GateD in that it allows to offloads the computation from the CPU to special ASICs and in it's modularity. <http://www.zebra.org/>

32. Switching:

"Layer 3 and 4 Switching", article from *Performance Computing*.

<http://www.performancecomputing.com/columns/packets/9812.shtml>

"IP Switching: Issues and Alternatives," by R. Jain.

<http://www.cis.ohio-state.edu/~jain/talks/ipsw.htm>

"IP Switching", course given by Shishir Agrawal.

http://www.cis.ohio-state.edu/~jain/cis788-97/ip_switching/index.htm

"L5: A Self Learning Layer 5 Switch", a report from IBM

<suddenly unavailable online>

33. *Internet Performance Measurement and Analysis Project* home page.

<http://www.merit.edu/ipma/>

34. Host Name to Latitude/Longitude [http://cello.cs.uiuc.edu/cgi-](http://cello.cs.uiuc.edu/cgi-bin/slamm/ip2ll/)

[bin/slamm/ip2ll/](http://cello.cs.uiuc.edu/cgi-bin/slamm/ip2ll/)

35. *Internet Weather Report*

<http://www.internetweather.com/>

<http://www3.mids.org/weather/>

<http://www.internettrafficreport.com/>

36. *Connected: An Internet*

Encyclopedia<http://www.freesoft.org/CIE/index.htm>

37. The Network Engineer's Toolkit Site <http://www.wanresources.com/>

38. TCP/IP For Internet Administrators <http://techref.ezine.com/tc/>

39. *Roll Your Own Intranet* page <http://users.neca.com/vmis/roll.htm>

40. *Materials on TCP/IP*

Networking<http://spectral.mscs.mu.edu/NetworksClass/Materials/>

41. *Windows and TCP/IP for Internet*

Access<http://learning.lib.vt.edu/wintcpip/wintcpip.html>

42. Al's WinSock Tuning FAQ [http://www.cerberus-](http://www.cerberus-sys.com/~belleisl/mtu_mss_rwin.html)

[sys.com/~belleisl/mtu_mss_rwin.html](http://www.cerberus-sys.com/~belleisl/mtu_mss_rwin.html)

43. *Henning Schulzrinne's RTP (Real Time Protocol)* site

<http://www.cs.columbia.edu/~hgs/rtp/>

Queen's University Real - Time Transport Protocol

(QRTP)<http://htm4.ee.queensu.ca:8000/ling/QRTP.html>

44. *My own IP -> Geographical Location Detective's* page

<http://t2.technion.ac.il/~s2845543/IP2geo.html>

45. Computer Networking and Internet Protocols

By Keith W. Ross and James F. Kurose
<http://www.seas.upenn.edu/~ross/book/Contents.htm>

46. Shawn J. Rappaport's Internetworking page
<http://www.futureone.com/~opeth/internetwork.html>
47. *Slow start* & delayed ack explained <http://www.sun.com/sun-on-net/performance/tcp.slowstart.html>
48. Information about NetBIOS and NetBEUI can be found at

<http://www.s390.ibm.com/bookmgr-cgi/bookmgr.cmd/BOOKS/bk8p7001/CCONTENTS>

<http://ourworld.compuserve.com/homepages/timothydevans/nbf.htm>
49. *RGB's TCP/IP Whitepapers & Guides*
<http://www.rgb.co.uk/support/guides/tcpip.htm>
50. ADTRAN PPP Internetworking Primer
<http://www.alliancedatacom.com/dial-up-point-to-point-technology.htm>
51. *IP Masquerade for Linux*<http://ipmasq.cjb.net/>

TCP/IP courses from universities:

52. *The Cooperative Association for Internet Data Analysis* maintains a list of pointers to Internet Engineering related university courses.
<http://www.caida.org/iec.evi/courses/index.html>
53. Dr. Reuven Cohen

Internet Networking

Technion - Israel Institute of Science
<http://www.cs.technion.ac.il/Courses/cs236341/>
54. Dr. Shlomi Dolev

Computer Communications and Distributed Algorithms

Ben-Gurion University <http://www.cs.bgu.ac.il/~ccda982/> (slides are in hebrew)
55. Dr. Ofer Hadar

Introduction To Computer Networks

Technion - Israel Institute of Science
<http://www.cs.technion.ac.il/~cs236334/> (slides are in hebrew)
56. Prof. Deborah Estrin

Computer Communications

University of South California <http://catarina.usc.edu/cs551/cs551.html>
57. Dave Hollinger

Network Programming<http://www.cs.rpi.edu/courses/netprog/index.html>

58. Prof. Jim Kurose

Computer Networks<http://gaia.cs.umass.edu:80/cs653/>

59. David C. Blight

Telecommunication Networks<http://www.ee.umanitoba.ca/~blight/c24759-97/>

60. Phil Scott

Data Communications, Computer Networks<http://ironbark.bendigo.latrobe.edu.au/staff/pscott/pscott.home.html>

61. David Cyganski

Telecommunications Transmission Technologies<http://bugs.wpi.edu:8080/EE535/>

62. S. Keshav

Engineering Computer Networks<http://www.cs.cornell.edu/cs519/>

63. Prof. Ralph Droms

Purdue University

Computer Networks<http://www.netbook.cs.purdue.edu/cs363/index.html>

64. Simon Cleary

RMIT university

Computer Networks and Protocols<http://www.cse.rmit.edu.au/~rdssc/courses/ds454/>

65. Phil Scott

La Trobe university

Computer Networks<http://ironbark.bendigo.latrobe.edu.au/subjects/bitcne/>

The following links would supply info about IPv6:

66. *IP Next Generation*

This is the first site to visit to get any information about IPv6, from overviews, through RFCs & drafts, to implementations (including availability of stacks on various platforms & source code for IPv6 stacks)
<http://playground.sun.com/pub/ipng/html/ipng-main.html>

- 67. UK IPv6 Resource Centre <http://www.cs-ipv6.lancs.ac.uk/>
- 68. 6bone Home Page <http://www.6bone.net/>
- 69. IP Next Generation Overview
<http://info.isoc.org/HMP/PAPER/PT1/html/pt1.html.hinden>
- 70. IPv6: The New Internet Protocol

By William Stallings

<http://www.comsoc.org/pubs/surveys/stallings/stallings-orig.html>

- 71. The New and Improved Internet Protocol

By William Stallings <http://www.byte.com/art/9609/sec5/art2.htm>

- 72. The IPng Group's home page <http://ganges.cs.tcd.ie/4ba2/ipng/>
- 73. IPv6 RFCs & links collection <http://www.aloni.com/IPv6.htm>
- 74. IPv6: The New Version of the Internet Protocol

By Steve Deering <http://trail.isi.edu/deering/>

- 75. The Future of the Internet: *IPng and the TCP/IP Protocol* <http://ccnga.uwaterloo.ca/~dkidston/presentations/IPng/index.html>
- 76. IPv6 specifications - Latest RFCs and Internet Drafts Collection
<http://seusa.sumitomo.com/htmls/randd/ipv6/doc.html>
- 77. Process' IPv6 Resource Center. <http://www.process.com/ipv6.htm>
- 78. IPv6: The Next Generation Internet Protocol

By Gary C. Kessler http://www.hill.com/library/staffpubs/ipv6_exp.html

- 79. *IPv6: Next Generation Internet Protocol* <http://www.3com.com/nsc/ipv6.html>
- 80. *Literature Research IPv6 (IPng)*, by Mike Crawford.
<http://www.mediaport.org/~iamano/lr.zip>
- 81. Existing Routing Protocols and IPv6 <http://www.pub.ro/~cmatei/network/ipv6/>
- 82. *The IPv6 organization site*. <http://www.ipv6.org/>

For information about the Internet's future:

- 83. *Internet II* site. <http://www.internet2.org/>
- 84. *Next Generation Internet Initiative* <http://www.ngi.gov/>
- 85. *The Quality of Service Forum site*. http://www.qosforum.com/tech_resources.htm

The following links would supply info about IP multicasting:

- 86. *The IP Multicast Initiative home page* <http://www.ipmulticast.com/>
- 87. *The Mbone (multicast bone) FAQ* <http://www.cs.columbia.edu/~hgs/internet/mbone-faq.html>
- 88. *The multicast backbone home page* <http://www.mbone.com/>
- 89. *An Introduction to IP Multicast* <http://ganges.cs.tcd.ie/4ba2/multicast/index.html>
- 90. *Introduction to IP Multicast Routing*
<http://www.3com.com/nsc/501303.html>

91. A collection of documents explaining multicast routing.
ftp://ftpeng.cisco.com/ipmulticast/multicast_training.html

The following links would supply info about IP security:

92. *Internet Security Survey* - <http://www.trouble.org/survey/>
93. *Phrack Magazine's* site - <http://www.phrack.com/>
94. The *SKIP* site - <http://www.skip.org/>

Simple Key management for Internet ProtocolsSKIP

encrypts info at the IP layer, enabling all applications which communicate via IP (using either TCP or UDP) to benefit from security.

95. *Peter Gutmann's "Security and Encryption-related Resources and Links"* contains a huge collection of links to security sites.
<http://www.cs.auckland.ac.nz/~pgut001/links.html>
96. COAST's Hotlist: *Computer Security, Law & Privacy* is another huge collection of links to security & privacy issues.
<http://www.cs.purdue.edu/homes/spaf/hotlists/csec-plain.html>
97. *Telstra* has a Security Papers & Documents page, most of them relating to network security. <http://www.telstra.com.au/pub/docs/security/>
98. *SunWorld* has a good article about IPsec
<http://www.sun.com/sunworldonline/swol-06-1998/swol-06-ipsec.html>

Pages about research into networking can be found at:

99. *Networking Research at the PSC*<http://www.psc.edu/networking/>
100. List of Publications by Raj Jain's Group <http://www.cis.ohio-state.edu/~jain/papers.html>
101. *Luigi Rizzo* - Research work
<http://www.iet.unipi.it/~luigi/research.html>
102. *UCLA Internet Research Lab*<http://irl.cs.ucla.edu/>
103. *TCP Over Satellite* work group <http://tcpsat.lerc.nasa.gov/tcpsat/>
104. *Mobile Computing Paper Collection at NTHU*<http://piggy.cs.nthu.edu.tw/paper/Mobile/index.html>
105. *Rutgers university DataMan mobile computing laboratory*<http://www.cs.rutgers.edu/dataman/>
106. *Network Bibliography*<http://www.cs.columbia.edu/~hgs/netbib/>
107. *ValueRocket Consulting*<http://www.valuerocket.com/papers/>

Mark Daugherty's TCP/IP page contains IPv4 Datagram Reference Chart in AutoCad format (.dxf) and as a 9 pages Word document, as well as lots of other links to such stuff as well known port numbers, FAQs, ethernet resources, etc, in his home-page.

<http://mdaugherty.home.mindspring.com/index.html>

<http://mdaugherty.home.mindspring.com/tcpip.html> [TCP/IP page]

The *protocols.com* site has posters of many protocols in both HTML and PDF formats, though the later requires (free) registration.

<http://www.protocols.com/pbook/tcpip.htm> [HTML posters]

<http://www.protocols.com/pbook/pdf/index.html> [PDF posters]

The *Information Technology Professional's Resource Center* contains plenty of links to networking subjects, including IP, Cisco, guides, magazines' home pages, networking security, and more. <http://www.itprc.com/>

Randy Baker's Introduction to Data Communications page
<http://www.georcoll.on.ca/staff/rbaker/idccbt.htm>

TechFest's Networking page. <http://www.techfest.com/networking/>

First Monday is a journal about the Internet which is published on the internet, with all it's articles peer-reviewed.

It's archives contain articles about TCP/IP, indexed at
<http://www.firstmonday.dk/subjects/technical.html>

The *Institute for Global Communications (IGC)* has an excellent page of TCP/IP resources, starting from some general background, through pointers to platform specific links and comm-hardware links. <http://www.igc.org/igc/help/tcpip.html>

Cisco's site contains a couple of internetworking guides:

108. IP Protocols page
<http://cio.cisco.com/warp/public/732/IP/index.html>
109. IP Technical Tips page <http://www.cisco.com/warp/public/105/>
110. Internetworking Technology Overview
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/index.htm
111. Internetwork Design Guide
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/index.htm>

IBM's Austin site contains a couple of TCP/IP guides:

112. TCP/IP Tutorial and Technical Overview
http://www.austin.ibm.com/resource/aix_resource/Pubs/redbooks/html-books/gg243376.04/3376fm.html
 113. Using the Information Superhighway
http://www.austin.ibm.com/resource/aix_resource/Pubs/redbooks/html-books/gg242499.00/2499fm.html
 114. Accessing the Internet
http://www.austin.ibm.com/resource/aix_resource/Pubs/redbooks/html-books/sq242597.00/2597fm.html
4. *Richard Stevens'* home page <http://www.kohala.com/~rstevens/>

Douglas Comer's home page <http://www.cs.purdue.edu/people/comer>

Andrew Tannenbaum's home page <http://www.cs.vu.nl/~ast/>

William Stallings's home page <http://www.shore.net/~ws>

James Carlson's home page <http://people.ne.mediaone.net/carlson/ppp>

Raj Jain's home page <http://www.cis.ohio-state.edu/~jain/>

5. *O'Reilly* <http://www.ora.com/>

Prentice Hall<http://www.prenhall.com/>

Addison Wesley<http://www.aw.com/>

MacMillan<http://www.mcp.com/>

McGraw-Hill<http://www.mcgraw-hill.com/>>

MIS:Press<http://www.mispress.com/> (M & T Books)

New Riders<http://www.newriders.com/>

InfoMagic

home page <http://www.infomagic.com/>

ftp site <ftp://ftp.infomagic.com/>

Walnut Creek's

home page <http://www.cdrom.com/>

ftp site <ftp://ftp.cdrom.com/>

6. *GNU* project <http://www.gnu.org/>

OpenBSD's home page <http://www.openbsd.org/>

FreeBSD's home page <http://www.freebsd.org/>

NetBSD's home page <http://www.netbsd.org/>

Linux's home page <http://www.linux.org/>

Trinux's home page <http://www.trinux.org/>

Linux Kernel Archive<http://www.kernel.org/>

The *Internet Software Consortium*, a non-profit organization, carries and supports BIND, DHCP, and INN. The software is supplied for free, as well as limited support via mailing list. A support contract comes, naturally, with a fee. <http://www.isc.org/>

Erick Engelke has a web page titled "WATTCP Locator", supplying lots of info about WATTCP, a TCP/IP package for DOS. The latest version of WATTCP is pointed to from this page. <http://www.supro.com/wattcp/wattcp.html>

Gisle Vanem has upgraded the WATTCP tcp/ip stack to include DHCP, RARP, file-based lookup, BSD-compatible API. Supports several compilers and DOS-extenders. WATT-32 is found at <http://www.bgnett.no/~giva/index.html>

Phil Karn's KA9Q (DOS TCP/IP stack) is under Karn's home page. <http://people.qualcomm.com/karn/code/ka9qnos/>

Michael Bernardi's MS-DOS Applications for Internet Use FAQ, which contains a list of TCP/IP stacks & applications for DOS.

<ftp://ftp.demon.co.uk/pub/doc/ibmpc/dos-apps.txt>

<http://www.dendarii.demon.co.uk/FAQs/dos-apps.html>

Dan Kegel has a page titled "MS-DOS TCP/IP Programming", which is crammed with links & info about TCP/IP for DOS.

<http://www.alumni.caltech.edu/~dank/trumpet/>

The *Public Netperf* Homepage is available, courtesy of HP, at

<http://www.netperf.org/>

The *Linux Router Project*, making a floppy sized distribution of Linux used to build and maintain routers, terminal servers, etc. <http://www.linuxrouter.org/>

7. A good search engine could supply further info.

The *Yahoo* engine, at <http://www.yahoo.com/>, has a good index, including a page about TCP-IP.

Some other good search engines are

AltaVista at <http://www.altavista.digital.com/>

InfoSeek at <http://www.infoseek.com/>

Hotbot at <http://www.hotbot.com/>

The *Networked Computer Science Technical Reference Library* site is an archive of computer science articles, which can be searched through using an impressive search engine. <http://www.ncstrl.org/>

The *DejaNews* site archives all the posts to usenet. The site, at <http://www.dejanews.com/>, enables users to search through posts sent over the past few years using different methods, which may be combined, such as words from articles, authors, and newsgroups. The ability to find past posts discussing unfamiliar subjects is an endless source of information, and may supply immediate answers to questions asked on usenet in the past.

If you wish to have a post of yours not archived in dejanews add the header "X-No-Archive: Yes" to your posting's header, or write it as your article's first line. Notice that this wouldnt prevent other people from quoting your article, thus causing the quoted material to be archived.

Other useful features of DejaNews:

- Get poster profiles.

This gives a count of how many posts did a poster send to each newsgroup, with a poster identified by it's email address.

- Search for newsgroups discussing given subjects.

As the search is done by frequency of words in posts, the results should be taken with a grain of salt, e.g.

NEWSGROUPS WHERE PEOPLE TALK ABOUT: christianity

All the newsgroups in the following list contain christianity in some article. The confidence rating indicates how sure we are that people talk about your query in the newsgroup. Clicking on the newsgroup name will show you all of the articles within the group which match your query.

Confidence	Newsgroup
99%	alt.atheism
63%	rec.games.frp.misc
54%	rec.music.christian
39%	alt.religion.christian
38%	soc.religion.christian
38%	soc.penpals
33%	austin.general

The *Norwegian University of Science and Technology*, located at Trondheim, has an FTP search engine on the web, located at <http://ftpsearch.ntnu.no/ftpsearch>, that can find files on anonymous FTP servers world wide.

The search is similar to the one done by archie, and can be very useful for finding source code for utilities, FAQs, etc.

A quick search for the word ping produced the following output:

```
ftp.cc.uec.ac.jp (Japan)
 1 ftp.cc.uec.ac.jp /.0/4.4BSD-
Lite/usr/src/sbin/ping
 2 ftp.cc.uec.ac.jp /.0/4.4BSD-
Lite/usr/src/sys/i386/floppy/ping
 3 ftp.cc.uec.ac.jp /.0/Linux/redhat-
4.1/i386/RedHat/instimage/usr/bin/ping
 4 ftp.cc.uec.ac.jp /.0/Linux/redhat-
devel/i386/RedHat/instimage/usr/bin/ping
ftp.dwc.edu (Educational)
 5 ftp.dwc.edu
/.03/redhat/i386/RedHat/instimage/usr/bin/ping
 6 ftp.dwc.edu
/.03/redhat/sparc/RedHat/instimage/usr/bin/ping
```

```
7 ftp.dwc.edu
/.03/redhat/sparc/misc/src/trees/rescue/bin/ping
ftp.fujixerox.co.jp (Japan)
8 ftp.fujixerox.co.jp /.1/NetBSD-
current/src/sbin/ping
[more links snipped]
```

Other files search engine are located at <http://www.filez.com/> and <http://castor.acs.oakland.edu/cgi-bin/vsl-front/> which can find files for specific platforms (e.g. unix, windows, mac) or specific formats (e.g. wav, midi, fonts, source code).

4 Newsgroups Discussing Networking & TCP/IP

news:alt.comp.dcom.sys.xyplex Discussions relating to Whittaker/Xyplex

news:alt.dcom.slip-emulators Pseudo-SLIP/PPP with shell accounts. TIA, SLAP, etc.

news:alt.dcom.telecom Unmoderated discussion of telecommunications technology.

news:alt.dcom.telecom.radius Remote Authentication Dial-In User Service

news:alt.mbone Global InterNet multicast network discussions

news:alt.winsock Windows Sockets.

news:alt.winsock.programming Programming Windows Sockets.

news:alt.winsock.trumpet The Trumpet newsreader.

news:comp.dcom.cabling Cabling selection, installation and use.

news:comp.dcom.cell-relay Forum for discussion of Cell Relay-based products.

news:comp.dcom.frame-relay Technology and issues regarding frame relay networks.

news:comp.dcom.isdn The Integrated Services Digital Network (ISDN).

news:comp.dcom.lans.ethernet Discussions of the Ethernet/IEEE 802.3 protocols.

news:comp.dcom.lans.fddi Discussions of the FDDI protocol suite.

news:comp.dcom.lans.hyperchannel Hyperchannel networks within an IP network.

news:comp.dcom.lans.misc Local area network hardware and software.

news:comp.dcom.lans.token-ring Installing and using token ring networks.

news:comp.dcom.modems Data communications hardware and software.

news:comp.dcom.modems.cable Cable modems and internet access via cable tv.

news:comp.dcom.net-analysis Network Testing and Analysis Procedures and Results.

news:comp.dcom.net-management Network management methods and applications.

news:comp.dcom.sys.bay-networks Bay Networks hardware, software, other products.

news:comp.dcom.sys.cisco Info on Cisco routers and bridges.

news:comp.dcom.sys.nortel Nortel telecommunications products and systems.

news:comp.dcom.telecom Telecommunications digest. (Moderated)

news:comp.dcom.telecom.tech Discussion of technical aspects of telephony.

news:comp.dcom.wan All topics concerned with wide area networking.

news:comp.dcom.xdsl Discussion area for different DSL technologies.

news:comp.os.linux.networking Networking and communications under Linux.

news:comp.os.ms-windows.apps.winsock.mail Winsock email applications.

news:comp.os.ms-windows.apps.winsock.misc Other Winsock applications.

news:comp.os.ms-windows.apps.winsock.news Winsock news applications.

news:comp.os.ms-windows.networking.misc Windows and other networks.

news:comp.os.ms-windows.networking.ras Windows RAS networking.

news:comp.os.ms-windows.networking.tcp-ip Windows and TCP/IP networking.

news:comp.os.ms-windows.networking.win95 Win95 to Novell, TCP/IP, other nets.

news:comp.os.ms-windows.networking.windows Windows' built-in networking.

news:comp.os.ms-windows.nt.admin.networking Windows NT network administration.

news:comp.os.ms-windows.programmer.networks Network programming.

news:comp.os.ms-windows.programmer.tools.winsock Winsock programming.

news:comp.os.os2.networking.tcp-ip TCP/IP under OS/2.

news:comp.protocols.dns.bind Berkeley Internet Name Domain (BIND). (Moderated)

news:comp.protocols.dns.ops DNS operations (where not BIND specific). (Moderated)

news:comp.protocols.dns.std DNS standards activities, including IETF. (Moderated)

news:comp.protocols.iso The ISO protocol stack.

news:comp.protocols.kerberos The Kerberos authentication server.

news:comp.protocols.misc Various forms and types of protocol.

news:comp.protocols.nfs Discussion about the Network File System protocol.

news:comp.protocols.ppp Discussion of the Internet Point to Point Protocol.

news:comp.protocols.smb SMB file sharing protocol and Samba SMB server/client.

news:comp.protocols.snmp The Simple Network Management Protocol.

news:comp.protocols.tcp-ip TCP and IP network protocols.

news:comp.protocols.tcp-ip.domains Topics related to Domain Style names.

news:comp.protocols.tcp-ip.ibmpc TCP/IP for IBM(-like) personal computers.

news:comp.protocols.time.ntp The network time protocol.

news:comp.security.firewalls Anything pertaining to network firewall security.

news:comp.security.misc Security issues of computers and networks.

news:comp.std.wireless Examining standards for wireless network technology. (Moderated)

news:comp.unix.large UNIX on mainframes and in large networks.

news:trumpet.questions Questions & general discussion of Trumpet Winsock.

news:trumpet.bugs Reporting & discussion of bugs or "features" in Trumpet Winsock.

news:microsoft.public.win16.programmer.networks

news:microsoft.public.win32.programmer.networks

news:microsoft.public.win95.dialupnetwork

news:microsoft.public.win95.networking

news:microsoft.public.win98.comm.dun

news:microsoft.public.win98.comm.modem

news:microsoft.public.win98.networking

news:microsoft.public.windowsnt.protocol.routing

news:microsoft.public.windowsnt.protocol.tcpip

5 Misc Networking Pages

1. A networking terms dictionary is available <http://www.ktek.com/ktek/Lans-Wans.html>
2. The *comp.protocols.ppp* FAQ is available at

<http://www.faqs.org/faqs/ppp-faq/part1/index.html>

<ftp://rtfm.mit.edu/pub/usenet-by-group/comp.protocols.ppp/>

The *comp.protocols.snmp* FAQ FAQ is available at

<http://www.pantherdig.com/snmpfaq/index.html>

<ftp://ftp.cs.utwente.nl/pub/src/snmp/>

There is a DHCP FAQ, written by *John Wobus*, available at <http://web.syr.edu/~jmwobus/comfaqs/dhcp.faq.html>

The *Amiga TCP/IP* FAQ, written by Mike Meyer, is available at

<http://www.phone.net/ATCPFAQ/amtcp.txt>

<http://www.phone.net/ATCPFAQ/amtcp.html>

There's a site for the Kermit project at <http://www.kermit-project.org/>

3. The *comp.dcom.lans.ethernet* FAQ is available at

<http://www.faqs.org/faqs/LANs/ethernet-faq/index.html>

<ftp://rtfm.mit.edu/pub/usenet-by-hierarchy/news/answers/LANs/ethernet-faq>

Charles Spurgeon's Ethernet Page is at <http://wwwhost.ots.utexas.edu/ethernet/ethernet-home.html>

The *comp.dcom.lans.token-ring* FAQ is available at <http://home.sprynet.com/sprynet/jtmesser/faq/contents.html>

The *comp.dcom.cabling* FAQ is available at

<http://www.faqs.org/faqs/LANs/cabling-faq/index.html>

<ftp://rtfm.mit.edu/pub/usenet-by-group/comp.dcom.cabling/>

The *comp.dcom.cell-relay* FAQ is available at <http://cell-relay.indiana.edu/cell-relay/FAQ/ATM-FAQ/FAQ.html>

4. The *Daedalus* project at Berkeley deals with wireless networking and mobile computing, and its web page contains links to some articles. <http://daedalus.cs.berkeley.edu/>

Two pages describing T1 with technical details are

<http://www.laruscorp.com/t1tut.htm>

<http://www.gsnetworks.com/ezvu/ezvu500/TUTORIAL/PROTOCOL/genT1.txt>

5. The *Big-LAN* FAQ, created for the big-lan@listserv.syr.edu mailing list, which discusses "[the] issues in designing and operating Campus-Size Local Area Networks,..." is available at <ftp://rtfm.mit.edu/pub/usenet-by-hierarchy/news/answers/LANs/big-lan-faq>

The *comp.security.firewalls* newsgroup has a FAQ, available at

<http://www.clark.net/pub/mjr>

<ftp://ftp.greatcircle.com/pub/firewalls/FAQ>

There's also a firewalls mailing list, served by <mailto:Majordomo@GreatCircle.com> archived at <ftp://ftp.greatcircle.com/pub/firewalls/archive/>

Daniel K. Kim has built a Searchable Check Point FireWall-1 discussion archive site (other mailing lists archived as well). <http://msgs.securepoint.com/>

6. A large collection of communication tutorials may be found at IOL's training page, which has links to materials on TCP/IP, LAN technologies, programming & administrations manuals, and more. <http://www.iol.unh.edu/training/index.html>

Data Communications magazine has a collection of technical tutorials available at its site, covering such subjects as ATM, IP, high speed networking, etc. <http://www.data.com/Tutorials/>

The *University of Leeds ATM MultiMedia group* has a collection of articles, links, etc about ATM. <http://www.scs.leeds.ac.uk/atm-mm/links.html>

3COM has a page containing links to a collection of networking articles. http://www.3com.com/technology/tech_net/white_papers/index.html

7. The *comp.unix.programmer* FAQ can be found at:

<http://www.erlenstar.demon.co.uk/unix/>

<http://www.whitefang.com/unix/>

<ftp://rtfm.mit.edu/pub/usenet/comp.unix.programmer/faq>

8. The windows 95 FAQ, which covers, among other subjects, subjects relating to TCP/IP, networking, and modems, can be found at:

<http://www.orca.bc.ca/win95/>

<ftp://rtfm.mit.edu/pub/usenet/news.answers/windows/win95/faq/>

9. *Committee T1's World Wide Web Site* <http://www.t1.org/>

The *ATM Forum's* home page can be found at <http://www.atmforum.com/>

The *Frame Relay Forum's* home page can be found at <http://www.frforum.com/>

The *Frame Relay Resource Center* <http://www.alliancedatacom.com/>

The *cable modems* home page <http://www.cablemodems.com/>

The *GigaBit Ethernet Alliance* home page <http://www.gigabit-ethernet.org/>

10. *Protocols* for WAN, LAN, ATM data communications and telecommunications. <http://www.protocols.com/>

Oceanwave Technical Resources. <http://www.oceanwave.com/technical-resources/>

Rohit's Srivastava's High Speed Networking & Programming page. <http://members.tripod.com/~srohit/compu.html>

Network Design Tutorials and Other Resources. <http://www.alaska.net/~research/Net/tutorial.htm>

Networking Technologies - Software Toolkits and Documentationhttp://www.nsrc.org/lowcost_tools/net-tech.html

Network Troubleshooting site. <http://www.networktroubleshooting.com/>

Tomi Engdahl's Telecommunication Electronics Page.
<http://junitec.ist.utl.pt/einfo/telecom.html>

Edwin Kremer's Security References. <http://www.cs.ruu.nl/~edwin/hot-1st.html>

Standards (and Cross References) <http://www.cmpcmm.com/cc/standards.html>

Lynn Larrow's Modems, Networking and Communications Links page.
<http://www.webcom.com/~llarrow/comfaq.html>

Randy's Home Page. <http://ic.net/~nunez/>

Hill Associates IT Technology Training networking articles.
<http://www.hill.com/library/staffpubs/index.html>

Private IP Network Addresses

AUTHOR'S NOTE

The following document is an example of an RFC (Request for Comment) that was used at the main way of setting standards on the Internet. The RFC is produced by a small working group that may be affiliated with any of a number of companies or organizations. If an RFC passed the public review stage it would become a standard. Because the Internet and technology has become so complex the standards are usually set by standards groups that use a number of different numbering schemes such as IEEE, ISO, ASCII, etc.

The following RFC was used to define private address network ranges that can be used by anyone without having to get prior approval. These addresses are not routed across the Internet. You can use these addresses on an isolated network, an Intranet, or local LAN/WAN side of a firewall that is also connected to the Internet. Your network will also be inaccessible from the Internet when the local addresses are accessed from outside the firewall because they are not being routed. However, your router can be used to allow you to access the Internet from your local LAN/WAN.

Network Working Group

Request for Comments: 1918

Obsoletes: 1627, 1597

BCP: 5

Category: Best Current Practice

Y. Rekhter

Cisco Systems

B. Moskowitz

Chrysler Corp.

D. Karrenberg

RIPE NCC

G. J. de Groot

RIPE NCC

E. Lear

Silicon Graphics, Inc.

February 1996

Address Allocation for Private Internets

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

1. Introduction

For the purposes of this document, an enterprise is an entity autonomously operating a network using TCP/IP and in particular determining the addressing plan and address assignments within that network.

This document describes address allocation for private internets. The allocation permits full network layer connectivity among all hosts inside an enterprise as well as among all public hosts of different enterprises. The cost of using private internet address space is the potentially costly effort to renumber hosts and networks between public and private.

2. Motivation

With the proliferation of TCP/IP technology worldwide, including outside the Internet itself, an increasing number of non-connected enterprises use this technology and its addressing capabilities for sole intra-enterprise communications, without any intention to ever directly connect to other enterprises or the Internet itself.

The Internet has grown beyond anyone's expectations. Sustained exponential growth continues to introduce new challenges. One challenge is a concern within the community that globally unique address space will be exhausted. A separate and far more pressing concern is that the amount of routing overhead will grow beyond the capabilities of Internet Service Providers. Efforts are in progress within the community to find long term solutions to both of these problems. Meanwhile it is necessary to revisit address allocation procedures, and their impact on the Internet routing system.

To contain growth of routing overhead, an Internet Provider obtains a block of address space from an address registry, and then assigns to its customers addresses from within that block based on each customer requirement. The result of this process is that routes to many customers will be aggregated together, and will appear to other providers as a single route [RFC1518], [RFC1519]. In order for route aggregation to be effective, Internet providers encourage customers joining their network to use the provider's block, and thus renumber their computers. Such encouragement may become a requirement in the future.

With the current size of the Internet and its growth rate it is no longer realistic to assume that by virtue of acquiring globally unique IP addresses out of an Internet registry an organization that acquires such addresses would have Internet-wide IP connectivity once the organization gets connected to the Internet. To the contrary, it is quite likely that when the organization would connect to the Internet to achieve Internet-wide IP connectivity the organization would need to change IP addresses (renumber) all of its public hosts (hosts that require Internet-wide IP connectivity), regardless of whether the addresses used by the organization initially were globally unique or not.

It has been typical to assign globally unique addresses to all hosts that use TCP/IP. In order to extend the life of the IPv4 address space, address registries are requiring more justification than ever before, making it harder for organizations to acquire additional address space [RFC1466].

Hosts within enterprises that use IP can be partitioned into three categories:

Category 1: hosts that do not require access to hosts in other enterprises or the Internet at large; hosts within this category may use IP addresses that are unambiguous within an enterprise, but may be ambiguous between enterprises.

Category 2: hosts that need access to a limited set of outside services (e.g., E-mail, FTP, netnews, remote login) which can be handled by mediating gateways (e.g., application layer gateways). For many hosts in this category an unrestricted external access (provided via IP connectivity) may be unnecessary and even undesirable for privacy/security reasons. Just like hosts within the first category, such hosts may use IP addresses that are

unambiguous within an enterprise, but may be ambiguous between enterprises.

Category 3: hosts that need network layer access outside the enterprise (provided via IP connectivity); hosts in the last category require IP addresses that are globally unambiguous.

We will refer to the hosts in the first and second categories as "private". We will refer to the hosts in the third category as "public".

Many applications require connectivity only within one enterprise and do not need external (outside the enterprise) connectivity for the majority of internal hosts. In larger enterprises it is often easy to identify a substantial number of hosts using TCP/IP that do not need network layer connectivity outside the enterprise.

Some examples, where external connectivity might not be required, are:

- A large airport which has its arrival/departure displays individually addressable via TCP/IP. It is very unlikely that these displays need to be directly accessible from other networks.
- Large organizations like banks and retail chains are switching to TCP/IP for their internal communication. Large numbers of local workstations like cash registers, money machines, and equipment at clerical positions rarely need to have such connectivity.
- For security reasons, many enterprises use application layer gateways to connect their internal network to the Internet. The internal network usually does not have direct access to the Internet, thus only one or more gateways are visible from the Internet. In this case, the internal network can use non-unique IP network numbers.
- Interfaces of routers on an internal network usually do not need to be directly accessible from outside the enterprise.

3. Private Address Space

The Internet Assigned Numbers Authority (IANA) has reserved the following three blocks of the IP address space for private internets:

```
10.0.0.0           -           10.255.255.255   (10/8
prefix)
172.16.0.0        -           172.31.255.255   (172.16/12
prefix)
192.168.0.0       -           192.168.255.255  (192.168/16
prefix)
```

We will refer to the first block as "24-bit block", the second as "20-bit block", and to the third as "16-bit" block. Note that (in pre-CIDR notation) the first block is nothing but a single class A network number, while the second block is a set of 16 contiguous class B network numbers, and third block is a set of 256 contiguous class C network numbers.

An enterprise that decides to use IP addresses out of the address space defined in this document can do so without any coordination with IANA or an Internet registry. The address space can thus be used by many enterprises. Addresses within this private address space will only be unique within the enterprise, or the set of enterprises which choose to cooperate over this space so they may communicate with each other in their own private internet.

As before, any enterprise that needs globally unique address space is required to obtain such addresses from an Internet registry. An enterprise that requests IP addresses for its external connectivity will never be assigned addresses from the blocks defined above.

In order to use private address space, an enterprise needs to determine which hosts do not need to have network layer connectivity outside the enterprise in the foreseeable future and thus could be classified as private. Such hosts will use the private address space defined above. Private hosts can communicate with all other hosts inside the enterprise, both public and private. However, they cannot have IP connectivity to any host outside of the enterprise. While not having external (outside of the enterprise) IP connectivity private hosts can still have access to external services via mediating gateways (e.g., application layer gateways).

All other hosts will be public and will use globally unique address space assigned by an Internet Registry. Public hosts can communicate with other hosts inside the enterprise both public and private and can have IP connectivity to public hosts outside the enterprise. Public hosts do not have connectivity to private hosts of other enterprises.

Moving a host from private to public or vice versa involves a change of IP address, changes to the appropriate DNS entries, and changes to configuration files on other hosts that reference the host by IP address.

Because private addresses have no global meaning, routing information about private networks shall not be propagated on inter-enterprise links, and packets with private source or destination addresses should not be forwarded across such links. Routers in networks not using private address space, especially those of Internet service providers, are expected to be configured to reject (filter out) routing information about private networks. If such a router receives such information the rejection shall not be treated as a routing protocol error.

Indirect references to such addresses should be contained within the enterprise. Prominent examples of such references are DNS Resource Records and other information referring to internal private addresses. In particular, Internet service providers should take measures to prevent such leakage.

4. Advantages and Disadvantages of Using Private Address Space

The obvious advantage of using private address space for the Internet at large is to conserve the globally unique address space by not using it where global uniqueness is not required.

Enterprises themselves also enjoy a number of benefits from their usage of private address space: They gain a lot of flexibility in network design by having more address space at their disposal than they could obtain from the globally unique pool. This enables operationally and administratively convenient addressing schemes as well as easier growth paths.

For a variety of reasons the Internet has already encountered situations where an enterprise that has not been connected to the Internet had used IP address space for its hosts without getting this space assigned from the IANA. In some cases this address space had been already assigned to other enterprises. If such an enterprise would later connects to the Internet, this could potentially create very serious problems, as IP routing cannot provide correct operations in presence of ambiguous addressing. Although in principle Internet Service Providers should guard against such mistakes through the use of route filters, this does not always happen in practice. Using private address space provides a safe choice for such enterprises, avoiding clashes once outside connectivity is needed. A major drawback to the use of private address space is that it may actually reduce an enterprise's flexibility to access the Internet. Once one commits to using a private address, one is committing to renumber part or all of an enterprise, should one decide to provide IP connectivity between that part (or all of the enterprise) and the Internet. Usually the cost of renumbering can be measured by counting the number of hosts that have to transition from private to public. As was discussed earlier, however, even if a network uses globally unique addresses, it may still have to renumber in order to acquire Internet-wide IP connectivity.

Another drawback to the use of private address space is that it may require renumbering when merging several private internets into a single private internet. If we review the examples we list in Section 2, we note that companies tend to merge. If such companies prior to the merge maintained their uncoordinated internets using private address space, then if after the merge these private internets would be combined into a single private internet, some addresses within the combined private internet may not be unique. As a result, hosts with these addresses would need to be renumbered.

The cost of renumbering may well be mitigated by development and deployment of tools that facilitate renumbering (e.g. Dynamic Host Configuration Protocol (DHCP)). When deciding whether to use private addresses, we recommend to inquire computer and software vendors about availability of such tools. A separate IETF effort (PIER Working Group) is pursuing full documentation of the requirements and procedures for renumbering.

5. Operational Considerations

One possible strategy is to design the private part of the network first and use private address space for all internal links. Then plan public subnets at the locations needed and design the external connectivity.

This design does not need to be fixed permanently. If a group of one or more hosts requires to change their status (from private to public or vice versa) later, this can be accomplished by renumbering only the hosts involved, and changing physical connectivity, if needed. In locations where such changes can be foreseen (machine rooms, etc.), it is advisable to configure separate physical media for public and private subnets to facilitate such changes. In order to avoid major network disruptions, it is advisable to group hosts with similar connectivity needs on their own subnets.

If a suitable subnetting scheme can be designed and is supported by the equipment concerned, it is advisable to use the 24-bit block (class A network) of private address space and make an addressing plan with a good growth path. If subnetting is a problem, the 16-bit block (class C networks), or the 20-bit block (class B networks) of private address space can be used.

One might be tempted to have both public and private addresses on the same physical medium. While this is possible, there are pitfalls to such a design (note that the pitfalls have nothing to do with the use of private addresses, but are due to the presence of multiple IP subnets on a common Data Link subnetwork). We advise caution when proceeding in this area.

It is strongly recommended that routers which connect enterprises to external networks are set up with appropriate packet and routing filters at both ends of the link in order to prevent packet and routing information leakage. An enterprise should also filter any private networks from inbound routing information in order to protect itself from ambiguous routing situations which can occur if routes to the private address space point outside the enterprise.

It is possible for two sites, who both coordinate their private address space, to communicate with each other over a public network. To do so they must use some method of encapsulation at their borders to a public network, thus keeping their private addresses private.

If two (or more) organizations follow the address allocation specified in this document and then later wish to establish IP connectivity with each other, then there is a risk that address uniqueness would be violated. To minimize the risk it is strongly recommended that an organization using private IP addresses choose randomly from the reserved pool of private addresses, when allocating sub-blocks for its internal allocation.

If an enterprise uses the private address space, or a mix of private and public address spaces, then DNS clients outside of the enterprise should not see addresses in the private address space used by the enterprise, since these addresses would be ambiguous. One way to ensure this is to run two authority servers for each DNS zone containing both publicly and privately addressed hosts. One server would be visible from the public address space and would contain only the subset of the enterprise's addresses which were reachable using public addresses. The other server would be reachable only from the private network and would contain the full set of data, including the private addresses and whatever public addresses are reachable the private network. In order to ensure consistency, both servers should be configured from the same data of which the publicly visible zone only contains a filtered version. There is certain degree of additional complexity associated with providing these capabilities.

6. Security Considerations

Security issues are not addressed in this memo.

7. Conclusion

With the described scheme many large enterprises will need only a relatively small block of addresses from the globally unique IP address space. The Internet at large benefits through conservation of globally unique address space which will effectively lengthen the lifetime of the IP address space. The enterprises benefit from the increased flexibility provided by a relatively large private address space. However, use of private addressing requires that an organization renumber part or all of its enterprise network, as its connectivity requirements change over time.

8. Acknowledgments

We would like to thank Tony Bates (MCI), Jordan Becker (ANS), Hans-Werner Braun (SDSC), Ross Callon (BayNetworks), John Curran (BBN Planet), Vince Fuller (BBN Planet), Tony Li (cisco Systems), Anne Lord (RIPE NCC), Milo Medin (NSI), Marten Terpstra (BayNetworks), Geza Turchanyi (RIPE NCC), Christophe Wolfhugel (Pasteur Institute), Andy Linton (connect.com.au), Brian Carpenter (CERN), Randy Bush (PSG), Erik Fair (Apple Computer), Dave Crocker (Brandenburg Consulting), Tom Kessler (SGI), Dave Piscitello (Core Competence), Matt Crawford (FNAL), Michael Patton (BBN), and Paul Vixie (Internet Software Consortium) for their review and constructive comments.

9. References

[RFC1466] Gerich, E., "Guidelines for Management of IP Address Space", RFC 1466, Merit Network, Inc., May 1993.

[RFC1518] Rekhter, Y., and T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, September 1993.

[RFC1519] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

10. Authors' Addresses

Yakov Rekhter
Cisco systems
170 West Tasman Drive
San Jose, CA, USA
Phone: +1 914 528 0090
Fax: +1 408 526-4952
EMail: yakov@cisco.com

Robert G Moskowitz
Chrysler Corporation
CIMS: 424-73-00
25999 Lawrence Ave
Center Line, MI 48015
Phone: +1 810 758 8212
Fax: +1 810 758 8173
EMail: rgm3@is.chrysler.com

Daniel Karrenberg
RIPE Network Coordination Centre
Kruislaan 409
1098 SJ Amsterdam, the Netherlands
Phone: +31 20 592 5065
Fax: +31 20 592 5090
EMail: Daniel.Karrenberg@ripe.net

Geert Jan de Groot
RIPE Network Coordination Centre
Kruislaan 409
1098 SJ Amsterdam, the Netherlands
Phone: +31 20 592 5065
Fax: +31 20 592 5090
EMail: GeertJan.deGroot@ripe.net

Eliot Lear
Mail Stop 15-730
Silicon Graphics, Inc.
2011 N. Shoreline Blvd.
Mountain View, CA 94043-1389
Phone: +1 415 960 1980
Fax: +1 415 961 9584
EMail: lear@sgi.com

Integrating Linux[®] and Windows[®]

MIKE McCUNE

The complete solutions guide
for every Linux/Windows system
administrator!

Running Linux and Windows in the same environment? Here's the comprehensive, up-to-the-minute solutions guide you've been searching for!

In **Integrating Linux and Windows**, top consultant Mike McCune brings together hundreds of solutions for the problems that Linux/Windows system administrators encounter most often. McCune focuses on the critical interoperability issues real businesses face: networking, program/data compatibility, dual-boot systems, and more. You'll discover exactly how to:

- Use Samba and Linux to deliver high-performance, low-cost file and print services to Windows workstations
- Compare and implement the best Linux/Windows connectivity techniques: NFS, FTP, remote commands, secure shell, telnet, and more
- Provide reliable data exchange between Microsoft Office[®] and StarOffice[™] for Linux
- Provide high-performance cross-platform database access via ODBC
- Make the most of platform-independent, browser-based applications
- Manage Linux and Windows on the same workstation: boot managers, partitioning, compressed drives, file systems, and more

If you're running both Linux and Windows, McCune delivers honest and objective explanations of all your integration options, plus realistic, proven solutions you won't find anywhere else. This book will help you keep your users happy, your costs under control, and your sanity intact!

ABOUT THE AUTHOR

MIKE McCUNE is a Chicago-based consultant specializing in PC-based networks and workstations running Windows, Linux, and NetWare. He has been consulting since 1990 for clients including IBM, Hewlett-Packard, ITT, and GE.

\$39.99 U.S. | \$60.00 Canada
Prentice Hall
Upper Saddle River, NJ 07458
www.phptr.com



ISBN 0-13-030670-3



0 76092 00966 5

9 780130 306708

9 0000