Enterprise Server 5

Mandriva

Enterprise Server 5
Published 2010-10-01
Copyright © 2009-2010 Mandriva SA
by Loïc Vaillant, Christophe Potigny, Andreas Hasenack, Rafaël Garcia Suarez, Emmanuel Cohen, Vincent Cuirassier, Anne Nicolas, Antoine Ginies, Yann Droneaud, Anthoine Bourgeois, Cédric Delfosse, Nicolas Perrin, and Séverine Wiltgen

Legal Notice

This material may be distributed only subject to the terms and conditions set forth in the Open Publication License, v1.0 or later (the latest version is presently available at opencontent.org (`http://www.opencontent.org/openpub/`).

- Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.
- Distribution of the work or derivative of the work in any standard (paper) book form is prohibited unless prior permission is obtained from the copyright holder.

“Mandriva” and “DrakX” re registered trademarks in the US and/or other countries. The related “Star logo” is also registered. All rights reserved. All other copyrights embodied in this document remain the property of their respective owners.

Table of Contents

List of Tables

Préface

1. About Mandriva Linux

Mandriva Linux is a GNU/Linux distribution supported by Mandriva S.A. which was born on the Internet in 1998. Its main goal was and still is to provide an easy-to-use and friendly GNU/Linux . The values of Mandriva are simplicity, openness and innovation.

On April 7th 2005, the Mandrakesoft company changed its name to Mandriva to reflect its merger with Brazil-based Conectiva. Its core product, Mandrakelinux, became Mandriva Linux.

1.1. Contacting the Mandriva Linux Community

The following are various Web links pointing you to the most important Mandriva Linux-related sources. If you wish to know more about Mandriva as a company, connect to our website (`http://www.mandriva.com/`).

Mandriva Expert (`http://expert.mandriva.com/`) is Mandriva's support platform. It offers a new experience based on trust and the pleasure of rewarding others for their contributions.

We also invite you to subscribe to the various mailing lists (`http://www.mandriva.com/en/mailing_lists`) where the Mandriva Linux community demonstrates its vivacity and keenness.

Please also remember to connect to our security page (`http://www.mandriva.com/security/`). It gathers all security-related material about Mandriva Linux distributions. You will find security and bug advisories, as well as kernel update procedures, the different security-oriented mailing lists which you can join, and Mandriva Online (`https://online.mandriva.com/`). This page is a must for any server administrator or user concerned about security.

1.2. Purchasing Mandriva Products

Mandriva Linux users may purchase products on-line through the Mandriva Store (`http://store.mandriva.com`). You will not only find Mandriva Linux software, operating systems (Free, One, PowerPack) and "live" flash keys, but also special subscription offers, support, third-party software and licenses, documentation, GNU/Linux-related books, and other Mandriva goodies.

For offers destinated to professionnals, you can consult our website (http://www.mandriva.com/enterprise/en) or contact our customer team at **sales@mandriva.com**.

1.3. Contributing to Mandriva Linux

The skills of the many talented folks who use Mandriva Linux can be very useful in the making of the Mandriva Linux system:

- **Packaging.** A GNU/Linux system is mainly made of programs picked up on the Internet. They have to be packaged in order to work together.
- **Programming.** Une There are many, many projects directly supported by Mandriva: find the one which most appeals to you and offer your help to the main developer(s).
- **Internationalisation.** You can help us translate Web pages, programs and their respective documentation.

Consult the development projects (http://www.mandriva.com/en/community/contribute) page to learn more about how you can contribute to the evolution of Mandriva Linux.

2. Conventions Used in this Book

2.1. Typing Conventions

Formatted Example	Meaning
inode	Used to emphasize a technical term.
`ls -lta`	Used for commands and their arguments. (see Section 2.2.1).
`a_file`	Used for file names. It may also be used for RPM package names.
ls(1)	Reference to a `man` page. To read the page, simply type `man 1 ls`, in a command line.
`$ ls *.pid`	Formatting used for text snapshots of what you may see on your screen including computer interactions, program listings, etc.

Formatted Example	**Meaning**
`localhost`	Literal data which does not generally fit in any of the previously defined categories. For example, a key word taken from a configuration file.
OpenOffice.org	Defines application names. Depending on context, the application and command name may be the same but formatted differently. For example, most commands are written in lowercase, while applications names usually begin with an uppercase character.
Files	Indicates menu entries or graphical interface labels. The underlined letter, if present, informs you of a keyboard shortcut, accessible by pressing the **Alt** key plus the letter in question.
Le petit chaperon rouge	Identifies foreign language words.
Warning!	Reserved for special warnings in order to emphasize the importance of words. Read out loud!

Highlights a note. Generally, it gives additional information about a specific topic.

Represents a tip. It could be general advice on how to perform a particular action, or hints about nice features, such as shortcuts to make your life easier.

Be very careful when you see this icon. It always means that very important information about a specific subject will be dealt with.

2.2. General Conventions

2.2.1. Commands Synopsis

The example below shows the symbols you will see when the writer describes the arguments of a command:

```
command <non literal argument> [--option={arg1,arg2,arg3}] [optional arg ...]
```

These conventions are standard and you will find them elsewhere such as in the `man` pages.

The "<" (lesser than) and ">" (greater than) symbols denote a **mandatory** argument not to be copied as is, which should be replaced according to your needs. For example, `<filename>` refers to the actual name of a file. If this name is **`foo.txt`**, you should type `foo.txt`, not `<foo.txt>` or `<filename>`.

The square brackets ("[ ]") denote optional arguments, which you may or may not include in the command.

The ellipsis ("...") means an arbitrary number of arguments may be included.

The curly brackets ("{ }") contain the arguments authorized at this specific place. One of them is to be placed here.

2.2.2. Special Notations

From time to time, you will be asked to press, for example, the keys **Ctrl-R**, which means you need to press and hold the **Ctrl** key and tap the **R** character right after as well. The same applies for the **Alt** and **Shift** keys.

We use capital letters to represent the letter keys. This doesn't mean that you have to type them capitalized. However, there might be programs where typing **R** is not the same than typing **r**. You will be informed when dealing with such programs.

Regarding menus, going to menu item File→Reload user config (**Ctrl-R**) means: click on the File label displayed on the menu (generally located in the upper-left of the window). Then in the pull-down menu, click on the Reload user config item. Furthermore you are informed that you can use the **Ctrl-R** key combination (as described above) to get the same result.

2.2.3. System-Generic Users

Whenever possible, we use two generic users in our examples:

Queen Pingusa	queen	This is our default user, used through most examples in this book.
Peter Pingus	peter	This user can be created afterward by the system administrator and is sometimes used to vary the text.

Starting Your Installation

Mandriva Linux auto-detects a large number of peripherals and this list is way too long for us too quote it entirely in this manual. Nonetheless, some of the strategies described in this chapter will assure the compliance of your hardware and, in some cases, enable you manually configure unrecognized equipment.

Legal disclaimer: the list of Mandriva Enterprise Server 5 supported hardware contains information about peripherals that have been tested or reported to work properly on Mandriva Enterprise Server 5. Because of the large number of possible configurations, Mandriva can not guarantee that a specific peripheral will work with your system.

Chapter 1. Different ways of installing

1.1. Before Installing

Here are a few tips before you install Mandriva Enterprise Server 5. We also give you a few pointers in case of difficulty at the beginning of the installation.

If you're about to install on a new system, make sure that the distribution is recognized. We suggest you use the online database to help you choose your hardware:

- official certification (`http://www.mandriva.com/hardware`): all officially-certified Mandriva Linux hardware;
- community base (`http://hcl.mandriva.com`): community hardware base made up of information collected from users of the distribution.

In order to assure a complete, trouble-free installation, make sure that all your devices are correctly plugged in and powered on. DrakX will detect and automatically configure all devices connected to your server during the installation of Mandriva Linux.

You can choose from many installation types: locally (with CDs or a DVD) or through a network (PXE, FTP, HTTP). No matter what method you use, the Mandriva Linux DrakX, installer will guide you in this installation procedure.

1.2. Installation via CD/DVD

Insert Mandriva Enterprise Server 5 DVD in your CD/DVD player. Start your server. If the start was well done on the removable media, the display below appears. Otherwise, change the boot sequence set in the BIOS to set the CD/DVD priority (see your hardware manual).

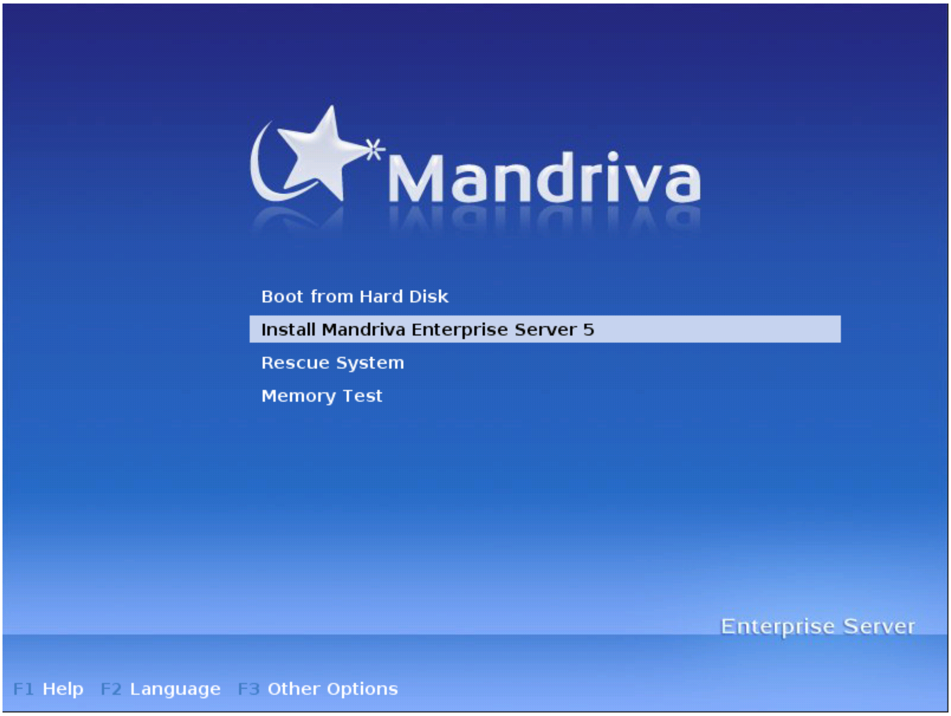

The launch of the installation is by default in graphical mode. However, you can switch to text mode installation. Then type `Esc` before loading the GUI.

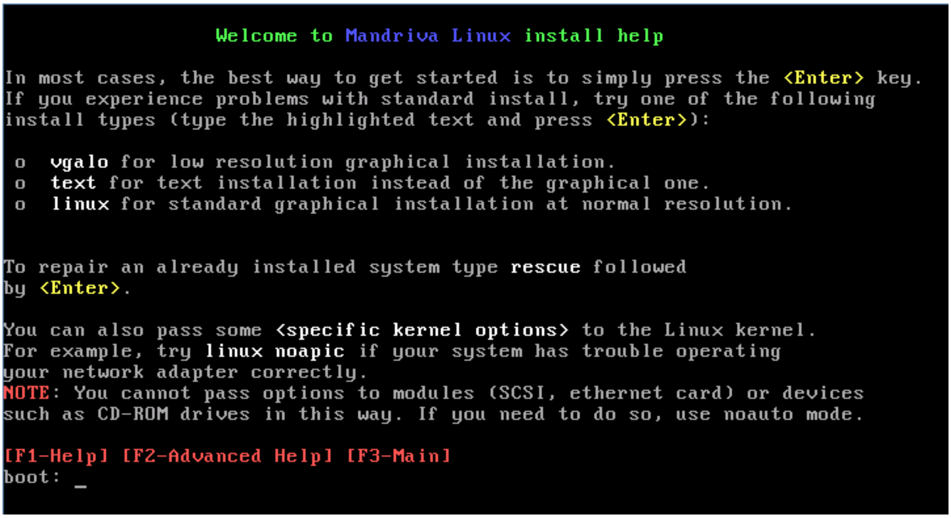

```
                 Welcome to Mandriva Linux install help

In most cases, the best way to get started is to simply press the <Enter> key.
If you experience problems with standard install, try one of the following
install types (type the highlighted text and press <Enter>):

 o  vgalo for low resolution graphical installation.
 o  text for text installation instead of the graphical one.
 o  linux for standard graphical installation at normal resolution.


To repair an already installed system type rescue followed
by <Enter>.

You can also pass some <specific kernel options> to the Linux kernel.
For example, try linux noapic if your system has trouble operating
your network adapter correctly.
NOTE: You cannot pass options to modules (SCSI, ethernet card) or devices
such as CD-ROM drives in this way. If you need to do so, use noauto mode.

[F1-Help] [F2-Advanced Help] [F3-Main]
boot: _
```

A text-mode interface is displayed and you can define start-up options. Here are the most common ones:

- `vgalo`: if you tried a default installation and didn't see the graphical interface, you can try to run the installation in low resolution mode. This happens with certain types of video cards. With Mandriva Linux you are given a number of options to work around problems related to older hardware. To try the installation in low resolution mode, enter **`vgalo`** at the prompt.
- `text` : if your video card is very old and the graphical installation doesn't work at all, you can always choose to install in text mode.
- `noauto` : in some rare cases, your PC may appear to freeze or lock up during the hardware detection phase. If that happens, adding the word `noauto` as a parameter tells the installation program to bypass hardware detection. However, you will need to manually specify hardware parameters later in the installation process. You can add the `noauto` parameter to the previous modes, so depending on your hardware, you may have to specify **`vgalo noauto`** to perform a low-resolution graphical installation without DrakX performing a hardware scan.
- kernel options : most machines don't require specific kernel options. Due to bugs in the design or in the BIOS, there have been a few cases of motherboards incorrectly reporting the amount of memory installed. If you need to manually specify the amount of RAM installed in your PC, use the `mem=xxxM` parameter. For example, to start the installation in normal mode with a computer containing 512 MB of memory, your command line would look like **`linux mem=512M`**. You can also use parameters such as `noapic` and `nolapic` to manage problems linked to interruption or the processor's communication modes.

During the installation, you'll be able to switch to console mode, giving you access to detailed system logs, but also to a shell environment. In that environment, notwithstanding the traditional commands, you can use the `bug` command, which allows you to save on a removable device the system logs, that can be very useful in case of troubleshooting.

1.3. Installing PXE

This chapter describes how to install Mandriva Enterprise Server 5 using the PXE (Pre-boot eXecution Environment) installation. You will learn how to configure the server and the client parts.

1.3.1. What is PXE?

PXE is a protocol designed by Intel that allows computers to boot across a network. PXE is stored in the ROM of recent generation network cards. When the computer boots up, the BIOS loads the PXE ROM into memory and executes it. From the menu displayed, choose an entry for the computer to boot on an operating system loaded through the network.

PXE is implemented by network card vendors following Intel's specification. The following figure represents how it works using a DHCP server and a TFTP server.

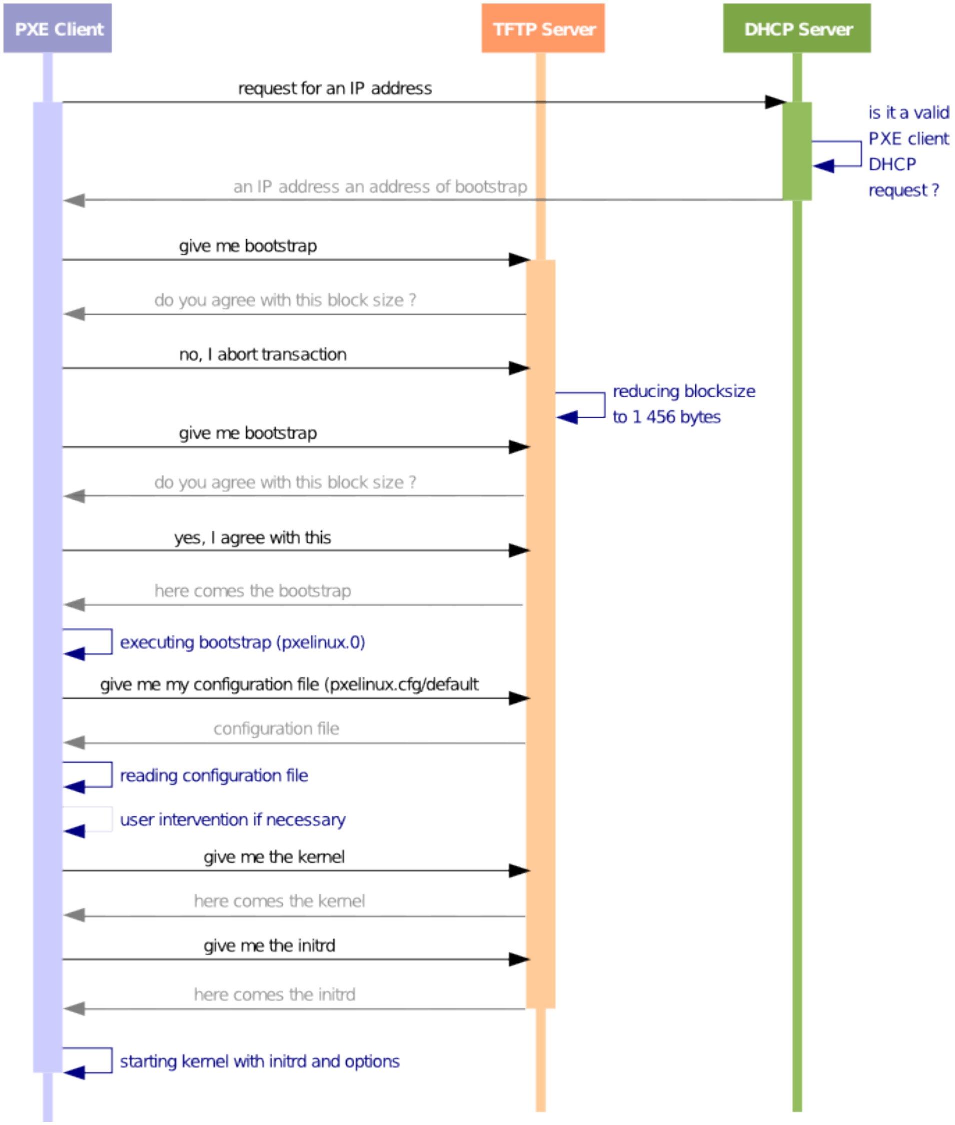
PXE Client
TFTP Server
DHCP Server
request for an IP address
is it a valid PXE client DHCP request ?
an IP address an address of bootstrap
give me bootstrap
do you agree with this block size ?
no, I abort transaction
reducing blocksize to 1 456 bytes
give me bootstrap
do you agree with this block size ?
yes, I agree with this
here comes the bootstrap
executing bootstrap (pxelinux.0)
give me my configuration file (pxelinux.cfg/default
configuration file
reading configuration file
user intervention if necessary
give me the kernel
here comes the kernel
give me the initrd
here comes the initrd
starting kernel with initrd and options

Figure 1-1. How PXE Works?

PXE follows 3 main steps to boot :

- Obtain an IP address from the DHCP server
- Download a bootstrap from the TFTP server.
- Execute the bootstrap

1.3.2. Using PXE to Install Mandriva Enterprise Server 5

Using PXE is quite easy, you need to follow these steps:

Checking your hardware

First check your hardware so that it can use PXE functions. There is no universal way to do it. Your motherboard or network card documentation should specify it. You can also have a look in your BIOS settings, especially for onboard network devices. You will be able to enable or disable a "ROM option" on your network device (the specific method can be rather different depending on the BIOS). Also, in the boot order menu, make sure that the "network boot" is configured before the hard disk boot.

Most recent computers allow you to choose network booting by pressing the **F12** key at the beginning of the boot process.

Booting your system

If your PXE server is well configured, during the booting process, you should see all the 3 steps explained in the previous chapter. Then you should get a prompt: `boot:`. Enter the label of the system you want to install and that's it.

1.3.3. Configuring a PXE Server

Setting up PXE functions:

- DHCP server
- TFTP server
- PXElinux bootstrap
- NFS or HTTP server, we use NFS server in our example

1.3.3.1. Configuration of DHCP Server

This server answers a special DHCP request from the PXE client, depending on the PXE client class. The location of the log files varies according to your server's configuration. You can find them in `/var/log/messages`. The default configuration file is stored in the `/etc` directory.

1.3.3.1.1. dhcpd.conf Options

Here are classical options you need to configure:

Allow booting

The booting flag is used to tell `dhcpd` whether or not to respond to queries from a particular client. This keyword only has meaning when it appears in a host declaration. By default, booting is allowed, but if it's disabled for a particular client, that client won't be able to get an address from the DHCP server.

Not authoritative

If the server is not valid for that segment it will send a `DHCPNACK` message. This is important to set if you have other DHCP servers. If you want to install a PXE server in a network that already has a DHCP server, setting this option will insure that your DHCP answer will always be considered after the authoritative servers. This is usually used with a restricted PXE pool configuration (see below).

Pool

his section of `/etc/dhcpd.conf` defines a pool which contains a range of IP addresses. In our example, the DHCP server allows members of the Class PXE, and denies members of other classes. The sample configuration (see below) restricts the address pool to the PXE clients. This insures that your DHCP will not respond to normal DHCP requests. This allows you to install a PXE server without conflicting with the master DHCP server. If you want to offer IP addresses to normal DHCP and PXE requests, just comment the "Allow member of PXE" line.

The PXE boot server also needs specific options for the DHCP server to work properly:

Class

In our example, we create a PXE Class to determine specific options.

Option vendor-class-identifier

If the `vendor-class-identifier` option of the DHCP request equals PXEClient, this class matches.

Vendor-option-space

Defined to allow some specific options for the class.

Filename

Option `filename` defines the bootstrap client to retrieve. Our TFTP is chrooted so the path is relative to the chrooted directory only. Usually the file represents the PXElinux bootstrap. It's frequently called "linux.0" or "pxelinux.0".

Next-server

It defines the IP of your TFTP server where the bootstrap and its configuration are stored.

Set vendor_class_identifier

Set the vendor-class-identifier field to PXEClient in the DHCP answer. If this field is not set, the PXE client will ignore the answer.

1.3.3.1.2. Sample from a Typical dhcpd.conf File

```
ddns-update-style ad-hoc;
allow booting;
allow bootp;
not authoritative;

 # Definition of PXE-specific options
option space PXE;
option PXE.mtftp-ip code 1 = ip-address;
option PXE.mtftp-cport code 2 = unsigned integer 16;
option PXE.mtftp-sport code 3 = unsigned integer 16;
option PXE.mtftp-tmout code 4 = unsigned integer 8;
option PXE.mtftp-delay code 5 = unsigned integer 8;
option PXE.discovery-control code 6 = unsigned integer 8;
option PXE.discovery-mcast-addr code 7 = ip-address;

class "PXE" {
match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
#filename "/PXEClient/pxegrub";
filename "/PXEClient/pxelinux.0";
option vendor-class-identifier "PXEClient";
vendor-option-space PXE;
option PXE.mtftp-ip 0.0.0.0;
next-server 192.168.200.1;
}

class "known" {
match hardware;
one-lease-per-client on;
ddns-updates on;
ddns-domainname = "mandriva.com";
ddns-hostname = pick-first-value(ddns-hostname, option host-name);
option fqdn.no-client-update on;
set vendor_class_identifier = option vendor-class-identifier;
}
```

```
shared-network "mynetwork" {
subnet 192.168.200.0 netmask 255.255.255.0 {
option subnet-mask 255.255.255.0;
option routers 192.168.200.1;
default-lease-time 28800;
max-lease-time 86400;

pool {
range 192.168.200.1 192.168.200.192;
allow members of "PXE";
}

}
}
```

1.3.3.2. Configuring the TFTP Server

The TFTP server contains the bootstrap and the configuration files. It's launched by the xinetd daemon. The default RPM installation usually works without modifying anything, except that it's not enabled by default. TFTP is managed by the xinetd daemon, so you need to specify that you want xinetd to open a dedicated port. The xinetd configuration for TFTP is in `/etc/xinetd.d/tftp`. Open this file and change `disable=yes` to `disabled=no`. Finally restart xinetd to take it into account: `servicexinetdrestart`.

1.3.3.3. Configuring PXElinux

PXElinux is the PXE bootstrap. This bootstrap is downloaded by the client PXE ROM and executed locally. The aim of this PXE bootstrap is to give a minimalist user interface for choosing the system you want to boot. PXElinux is provided by the SYSLINUX project. The bootstrap is usually named “pxelinux.0” or “linux.0”. Main features are:

- You can define a config file for each PXE client IP address
- The TFTP root directory is the one which contains `pxelinux.0`
- Complies with the vendor-class-identifier requests
- Boots a disk image (e.g. floppy disk image)

1.3.3.3.1. PXElinux tree

```
/var/lib/tftpboot
   |-- X86PC
 |-- linux
  |-- help.txt
  |-- images
  |   |-- 2009.0
  |       |-- all.rdz
  |       `-- vmlinuz
  |-- linux.0
  |-- memdisk
  |-- messages
  |-- pxelinux.cfg
       |-- default
```

All the configuration files are stored in the /X86PC/linux/ directory of the TFTP server: /var/lib/tftpboot. You can create a subdirectory that contains all the files needed by the PXE client. The client is stored in /var/lib/tftpboot/X86PC/linux (PXEPATH). So your linux.0 should be in PXEPATH and set in the dhcp configuration file. Then create a pxelinux.cfg directory in PXEPATH (CFGPATH).

The configuration directory contains all the configuration files (one per PXE client IP address) or a “default” configuration file. The naming mechanism is to convert your IP address in hexadecimal form: 192.168.200.1 gives you the name of the “C0A8C801” CFGFILE. This file contains options for the PXE client so that it will make the request with its 192.168.200.1 IP address. You can use the gethostip script to do this conversion.

In PXEPATH, create a directory called images which will contain subdirectories for each boot system (e.g. linux for Linux boot images or a kernel and its initrd). Each of these contains a boot image or a kernel, its initrd and a help file.

1.3.3.3.2. Configuring PXElinux in pxelinux.cfg

he PXE client downloads the bootstrap (linux.0 in our case), executes it locally and then tries to download its configuration CFGFILE file. This file contains general options and boot images definition. Here are its main options:

DEFAULTkerneloptions

Sets the default command line. If PXELINUX boots automatically, it will act just as if the entries after DEFAULT had been typed in at the boot prompt, except that the option auto is automatically added, indicating an automatic boot.

DISPLAYfilename

Displays the indicated file on the screen at boot time (before the boot prompt is shown).

TIMEOUTtimeout

ndicates how long to wait at start-up. Timeout is canceled as soon as the user types anything on the keyboard, the assumption being that the user will complete the command line already begun. A timeout of zero completely disables the timeout, which is also the default behavior.

F[1-9]fichier

When F[1-9] is pressed, then filename is displayed before going back to the prompt. It can be an easy way to provide help about boot images provided by the PXE server.

After having configured general options, we now need to define images provided by the PXE server. Each defined image starts with the label option. It contains strings that can be typed at the boot prompt. First, let's look at a specific image that can allow booting a local system by forcing the PXE client to exit.

```
label local
     LOCALBOOT 0
```

The general syntax is LOCALBOOTtype. It executes a local disk boot instead of booting a kernel. The argument **0** means to perform a normal boot on the next boot device. The **4** argument performs a local boot with the Universal Network Driver Interface (UNDI) driver still resident in memory. Finally, the **5** argument performs a local boot with the entire PXE stack, including the UNDI driver, still resident in memory. All other values are undefined. If you don't know what the UNDI or PXE stacks are, don't worry, you don't want them: just specify 0.

```
label label
     KERNEL image
     APPEND options...
```

If label is entered as the kernel boots, PXElinux should instead boot image, and the specified APPEND options should be used instead of the ones specified in the global section of the file.

PXE can also allow you to boot the kernel directly. The PXE client will download (using a TFTP server) the kernel in the path given in the KERNEL option, then it downloads the initrd in the path given in the APPENDinitrd= option.

Then the PXE client executes the kernel with the `initrd` and the option given by the `APPEND` option.

Using PXE to test your memory

When you have problems installing or booting a server, you may want to check your main memory. PXE can help you by providing this function.

Get a memtest binary by installing the memtest86+ package:

```
# urpmi memtest86+
```

Then get the binary and copy it to the PXElinux tree:

```
# cp /boot/memtest /var/lib/tftpboot/X86PC/linux/images/memtest
```

Finally add a new entry in the PXElinux configuration file:

```
label memtest kernel images/memtest
```

That's it! You will only need to boot your computer to test on a network device and enter `memtest` on the `linux` prompt.

When you have problems installing or booting a server, you may want to check your main memory. PXE can help you by providing this function.

Get a memtest binary by installing the memtest86+ package:

```
# urpmi memtest86+
```

Then get the binary and copy it to the PXElinux tree:

```
# cp /boot/memtest /var/lib/tftpboot/X86PC/linux/images/memtest
```

Finally add a new entry in the PXElinux configuration file:

```
label memtest kernel images/memtest
```

That's it! You will only need to boot your computer to test on a network device and enter `memtest` on the `linux` prompt.

1.3.3.3.3. Sample Configuration file

This file should be stored in `CFGPATH` and named like a `CFGFILE` (e.g. /var/lib/tftpboot/X86PC/linux/pxelinux.cfg/default)

```
PROMPT 1
```

```
     DEFAULT local
     DISPLAY messages
     TIMEOUT 50

     label local
     LOCALBOOT 0

     label linux
     KERNEL memdisk
     APPEND initrd=images/linux/network.img

     label memtest
     KERNEL images/memtest

label autoinstall
KERNEL images/autoinstall/vmlinuz
APPEND initrd=images/autoinstall/network.rdz ramdisk=32000 vga=788

kickstart=Mandriva/base/auto_inst.cfg.pl useless_thing_accepted
automatic=method:nfs,network:dhcp,interface:eth0,dns:192.168.100.11,
server:192.168.200.1,directory:/install root=/dev/ram3

F1 images/local/help.txt
F2 images/autoinstall/help.txt
```

Chapter 2. Main steps of installation

2.1. Choosing your Language

The first step is to choose the language in which the installation will take place.

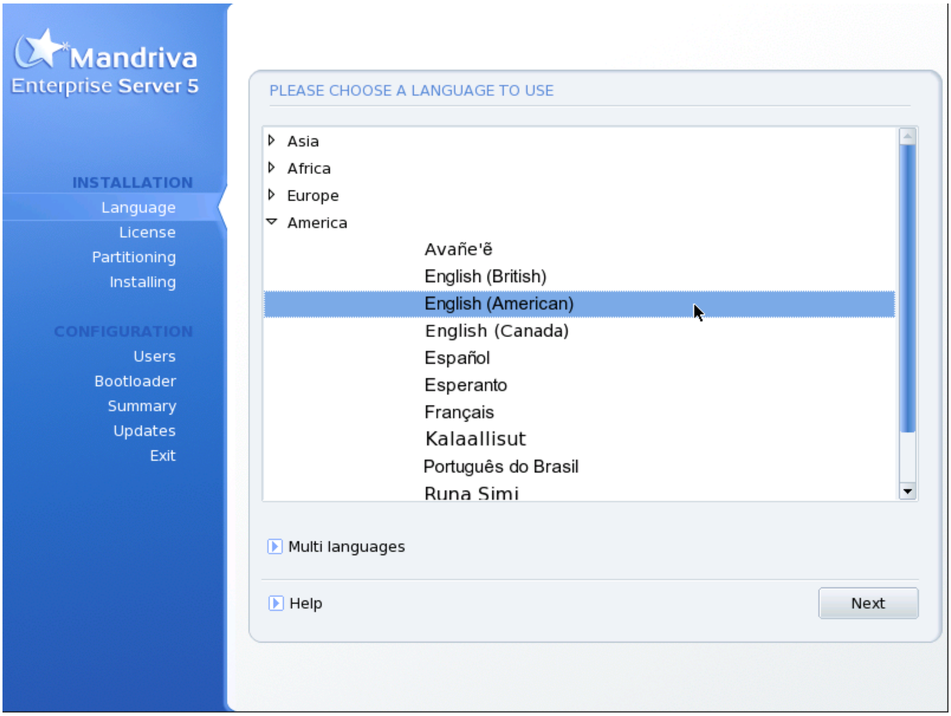

Figure 2-1. Choosing the Default Language

Open the tree relative to the continent you live on, then choose your language. Your language choice will affect the installer, the documentation, and the system in general.

Use the list accessible through the Multi languages button to select other languages to be installed on your workstation, thereby installing the language-specific files for system documentation and applications.

About UTF-8 (unicode) support: Unicode is a character encoding intended to cover all existing languages. However full support for Linux is still under development. For that reason, Mandriva Linux's use of UTF-8 depends on your choice:

1. If you choose a language with a strong legacy encoding (latin1 languages, Russian, Japanese, Chinese, Korean, Thai, Greek, Turkish, and most iso-8859-2 languages), the legacy encoding will be used by default.
2. Other languages use Unicode by default.
3. If you install two or more languages, and those languages don't use the same encoding, then Unicode is used for the whole system.
4. Finally, Unicode can also be forced for use throughout the system at a user's request by selecting the Use Unicode by default option independently of which languages have been chosen.

Note that you're not limited to choosing a single additional language. You may choose several, or even install them all by selecting the All languages option. Selecting support for a language means translations, fonts, spell checkers, etc. are also installed for that language. Make sure you select all languages which are likely to be useful on the machine now, because it may be difficult to configure support for languages you didn't choose at install time later on.

To switch between the various languages installed on your system, you can launch the `localedrake` command as root to change the language used by the entire system. Running the command as a regular user only changes the language settings for that particular user.

2.2. License Terms of the Distribution

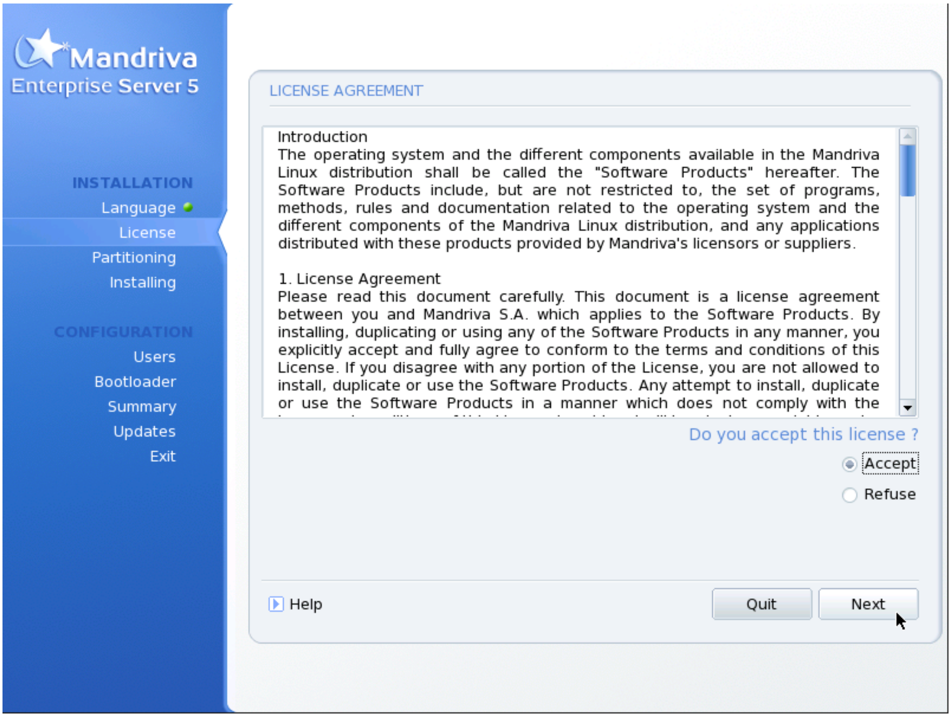

LICENSE AGREEMENT

Introduction
The operating system and the different components available in the Mandriva Linux distribution shall be called the "Software Products" hereafter. The Software Products include, but are not restricted to, the set of programs, methods, rules and documentation related to the operating system and the different components of the Mandriva Linux distribution, and any applications distributed with these products provided by Mandriva's licensors or suppliers.

1. License Agreement
Please read this document carefully. This document is a license agreement between you and Mandriva S.A. which applies to the Software Products. By installing, duplicating or using any of the Software Products in any manner, you explicitly accept and fully agree to conform to the terms and conditions of this License. If you disagree with any portion of the License, you are not allowed to install, duplicate or use the Software Products. Any attempt to install, duplicate or use the Software Products in a manner which does not comply with the

Do you accept this license ?

● Accept

○ Refuse

Before continuing, you should carefully read the terms of the license. It covers the entire Mandriva Linux distribution. If you agree with all the terms it contains, select Accept and click on Next. If not, clicking on Quit reboots your computer.

If you are curious about any technical changes which have occurred in the distribution since the last release, you can click on the Release Notes.

2.3. Installation Class

This step is shown only if an existing Linux partition is found on your machine.

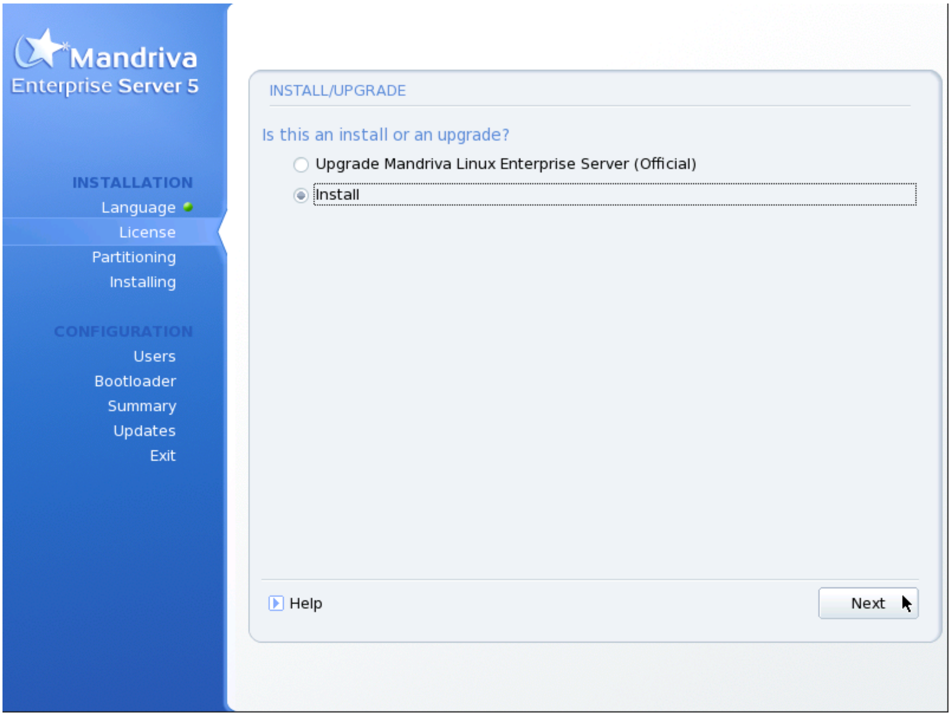

INSTALL/UPGRADE

Is this an install or an upgrade?

○ Upgrade Mandriva Linux Enterprise Server (Official)

● Install

DrakX now needs to know if you want to install from scratch or to upgrade your existing Mandriva Linux system:

Upgrade

> This installation type simply updates the packages currently installed on your Mandriva Linux system. Your current partitioning scheme and user data won't be altered. Most of the other configuration steps remain available and are similar to a standard installation.

Install

> For the most part, this completely wipes out the old system. However, depending on your partitioning scheme, you can prevent some of your existing data (particularly `/home` directories) from being overwritten.

2.4. Configuring Your Keyboard

This step only shows if your language settings don't match any single keyboard. Otherwise, your keyboard layout is automatically selected.

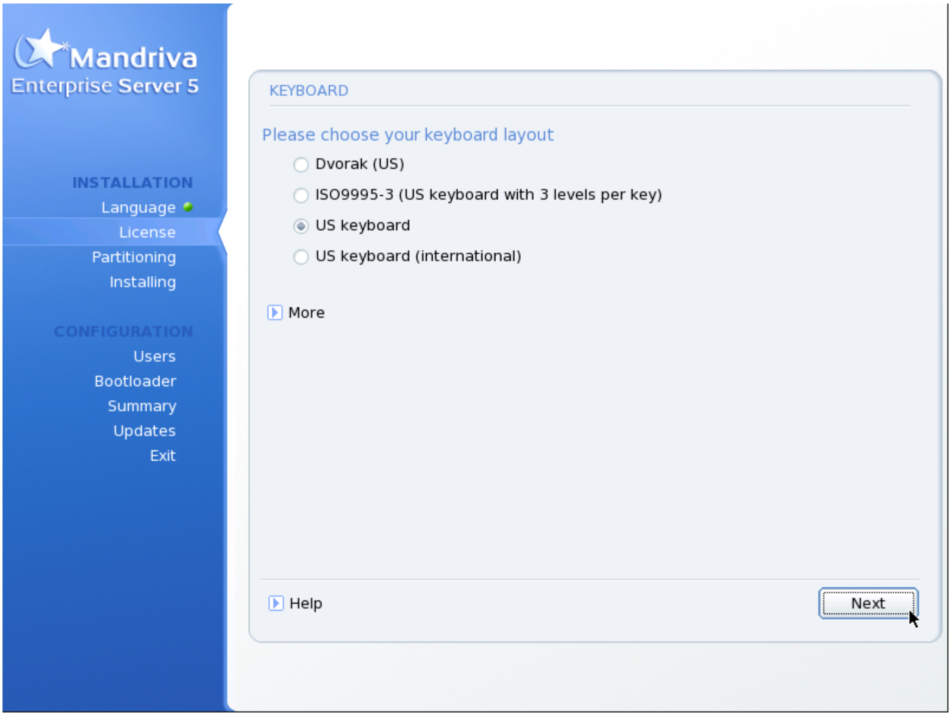

Please choose your keyboard layout

○ Dvorak (US)
○ ISO9995-3 (US keyboard with 3 levels per key)
● US keyboard
○ US keyboard (international)

Depending on the language you choose (see Section 2.1), DrakX automatically selects a particular type of keyboard configuration. Verify that the selection suits you or choose another keyboard layout.

Also, you may not have a keyboard which corresponds exactly to your language: for example, if you are an English-speaking Swiss native, you may have a Swiss keyboard. Or if you speak English and are located in Québec, you may find yourself in the same situation where your native language and country-set keyboard don't match. In either case, this installation step will allow you to select an appropriate keyboard.

Click on the More button for a list of supported keyboards to appear.

If you choose a keyboard layout based on a non-Latin alphabet, the next dialog enables you to choose the key binding which can switch the keyboard between the Latin and non-Latin layouts.

2.5. Security Level

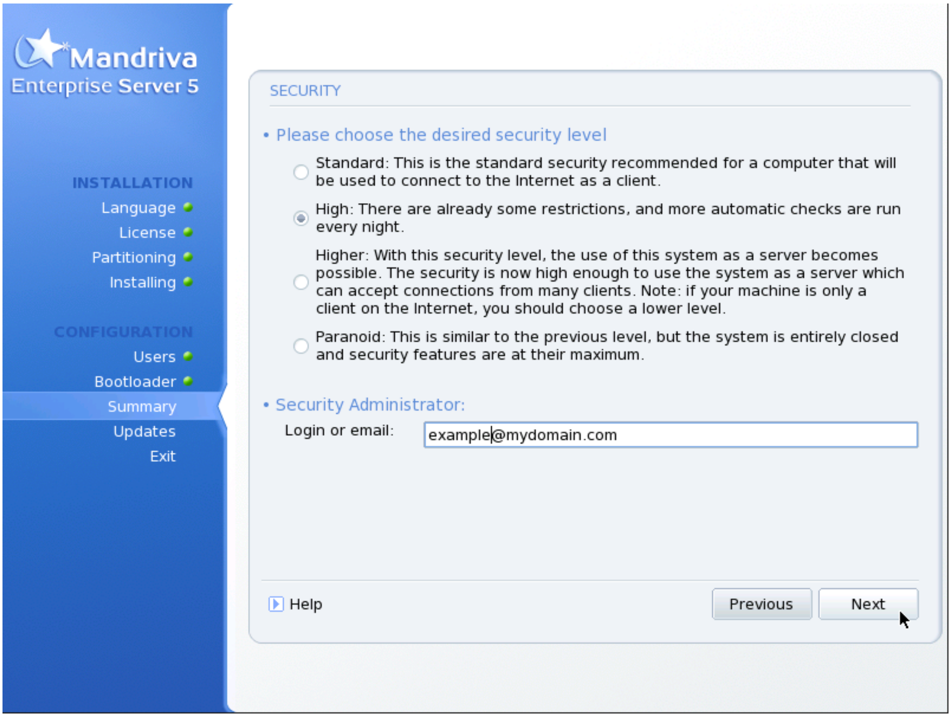

- Please choose the desired security level

○ Standard: This is the standard security recommended for a computer that will be used to connect to the Internet as a client.

● High: There are already some restrictions, and more automatic checks are run every night.

○ Higher: With this security level, the use of this system as a server becomes possible. The security is now high enough to use the system as a server which can accept connections from many clients. Note: if your machine is only a client on the Internet, you should choose a lower level.

○ Paranoid: This is similar to the previous level, but the system is entirely closed and security features are at their maximum.

- Security Administrator:

Login or email: example@mydomain.com

At this point, DrakX allows you to choose your machine's security level. As a rule of thumb, the security level should be set higher if the machine is to contain crucial data, or if it's to be directly exposed to the Internet. The trade-off is that a higher security level is generally obtained at the expense of ease of use.

If you don't know what to choose, keep the default option. You'll be able to change it later with the `draksec` tool in the Mandriva Control Center.

Fill the Security Administrator field with the e-mail address of the person responsible for security. Security-related messages will be sent to that address.

2.6. Partitioning Your Disk

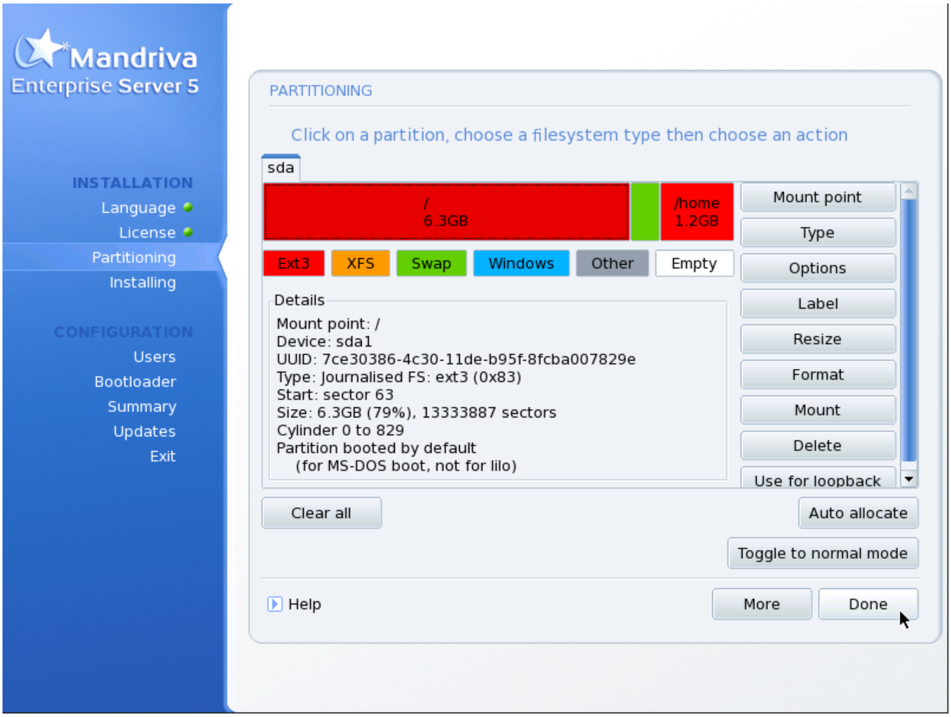

Click on a partition, choose a filesystem type then choose an action

Details

Mount point: /
Device: sda1
UUID: 7ce30386-4c30-11de-b95f-8fcba007829e
Type: Journalised FS: ext3 (0x83)
Start: sector 63
Size: 6.3GB (79%), 13333887 sectors
Cylinder 0 to 829
Partition booted by default
(for MS-DOS boot, not for lilo)

You now have to decide where you want to install Mandriva Linux on your hard drive. It needs to be partitioned which means it must be logically divided in order to create the required space for your new Mandriva Linux system.

Because the process of partitioning a hard drive is usually irreversible and can lead to data loss, it can be intimidating and stressful for the inexperienced user. Fortunately, DrakX includes a wizard which simplifies this process. Before continuing with this step, read through the rest of this section and above all, take your time.

Depending on the configuration of your hard drive, several options are available:

Use free space

This option performs an automatic partitioning of your blank drive(s). If you use this option, there will be no further prompts

Use existing partitions

The wizard detected one or more existing partitions on your hard drive. If you want to use them, choose this option. Then choose the mount points associated with each of the partitions. The legacy mount points are selected by default, and for the most part, it's a good idea to keep them. Then choose the partitions to be formatted or preserved.

Erase entire disk

Choose this option to delete all data and partitions present on your hard drive. You won't be able to undo this operation after you confirm.

If you choose this option, **all** data on your disk will be deleted.

Custom disk partitioning

Choose this option to manually partition your hard drive. Be careful: it's a powerful but dangerous choice and you can very easily lose all your data. This option is only recommended if you have performed custom disk partitioning before, and have enough Linux experience.

By default, Mandriva Enterprise Server 5 adds support for Access Control Lists (ACL) to manage advanced user rights on ext3 partitions.

2.6.1. Advanced use of DiskDrake

DiskDrake allows you to make a partition fit exactly to your needs.

2.6.1.1. The interface

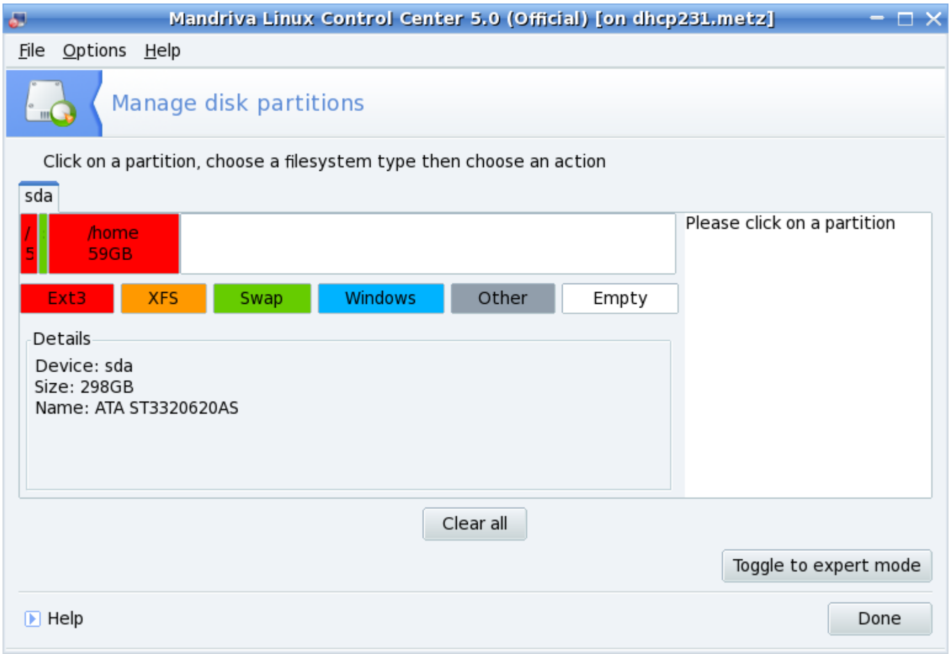

Click on a partition, choose a filesystem type then choose an action

Details

Device: sda
Size: 298GB
Name: ATA ST3320620AS

Please click on a partition

Figure 2-2. DiskDrake's Main Window

DiskDrake enables you to configure each physical hard drive on your machine. If you only have one IDE disk, you will see a single tab called sda below the file-system types. If there is more than one drive, then each drive will have its own tab and will be named according to the kernel name for that drive. `diskdrake` will allow you to manage the partitioning of each drive.

The window (see Figure 2-2) is divided into four zones:

- Top. The structure of your hard drive. When you launch DiskDrake, it will display the current structure of the drive. `diskdrake` will update the display as you make changes.
- Left. A menu relevant to the partition currently selected in the above diagram.
- Right. A description of the selected partition.
- Bottom. Buttons to do general actions. See next section.

2.6.1.2. DiskDrake's Action Buttons

Clear all

Clicking on this button clears all partitions on the current hard drive.

More

Displays a three-button dialog allowing you to:

Save partition table. Allows you to save the current partition table to a file on a disk (a floppy, for example). This may prove useful if a problem arises (such as an error made during drive repartitioning).

Restore partition table. Enables you to restore the partition table as previously saved with Save partition table. Restoring a partition table may recover your data as long as you do not reformat partitions, because the formatting process will overwrite all your data.

Rescue partition table. If you lose your partition table and have no backup, this function scans your hard drive to try and reconstruct the partition table.

Help

Display this documentation in a browser window.

Undo

Cancels the last action. Most modifications done on your partitions are not made permanent until `diskdrake` warns you it will write the partition table. Therefore, this button allows you to undo all modifications done on partitions up to the last time it was saved.

Toggle to expert mode

Allows you to access the expert mode functions (which are even **more** dangerous if you are not sure of what you are doing). Reserved for experts.

Done

Saves your changes and exits DiskDrake.

2.6.2. Using Advanced Manageable Data Partition

Here is a practical example using DiskDrake to obtain a list of the partitions contained on your machine. Your server will mainly be a file server. Therefore it needs a lot of manageable space so that you can increase it easily in case of a full partition. Here is the information we will use:

- Dynamically manageable partitions: we will use `LVM` (Logical Volume Manager). If the file system allows it, it enables you to resize partitions on-the-fly; it also manages disk space without physical limits.
- Advanced ACLs (Access Control Lists) to manage the rights of users: we need extra rights such as on Windows® file systems. The `XFS` file system allows you to increase the size on-the-fly.
- Misc information: total space available will be 350 GB, using 2 disks. Our file system will be mounted on the `/data` partition.

In DiskDrake, choose Custom disk partitioning. We will assume you have already created your system's partitions. At the bottom, click on Toggle to expert mode. Let's follow the steps below to get the desired partition:

1. Getting Space from System Disk.

   As we can see we have 50 GB left from the system disk. Click on it, then on Create. Choose the size so that you can use all the space left. Then in the Filesystem type list, choose Linux Logical Volume Manager. After validating your selection, you should be back to the main screen. Keep your partition selected and click on Add to LVM. In LVM name field, enter a name for this virtual partition. Let's call it `data`. You should now see one more tab named data.

2. Adding More Space from a Second Disk

   Let's now use the second disk to increase the available space. Click on the corresponding tab then on the available space. Click on Create and select all the disk space. In the Filesystem type list, choose Linux Logical Volume Manager. Back to the main screen, click on Add to LVM again. The next screen should now propose 2 items: data which is the name of the first logical group we created, and new. Since we want to increase the partition size of data, select it. You can check that the operation is completed: click on the data tab. The size of the `data` volume should now be increased.

3. Creating a Virtual Partition and File System

   Now that your global volume is ready, you need to create logical partitions inside. In the `data` tab, click on the available space then on Create. We will use 200 GB out of the 250 available. Adjust the size in the field.

Choose `xfs` in the Filesystem type. Fill the Mount point field by typing `/data`. You should also fill the Logical volume name field with `data`.

4. Optimizing File Systems

   Now that your partition and file system is ready, let's optimize some parameters:

   - `noatime` : using this option will not update inode access time on this file system. Data access will be then faster.
   - `grpquota` et `usrquota` : enables you to limit user and group quotas according to the space used per user or group.

In order to simplify administration, you should use significant names for all partitioning steps.

Your file system is now ready. You will be able to modify options and increase the partition's space by using the command line or the `harddrake` tool from the Mandriva Linux Control Center.

2.7. Package Selection

We now enter the software package installation itself. It first consists of selecting the installation media and then the packages to be installed.

2.7.1. Media Handling

If you are doing an installation from a CD, you are first asked to select the CDs you actually have available.

You can also copy all packages on your hard drive. This speeds up installation and eases later package installation as all packages are already available on your hard disk

2.7.2. Choose Packages to Install

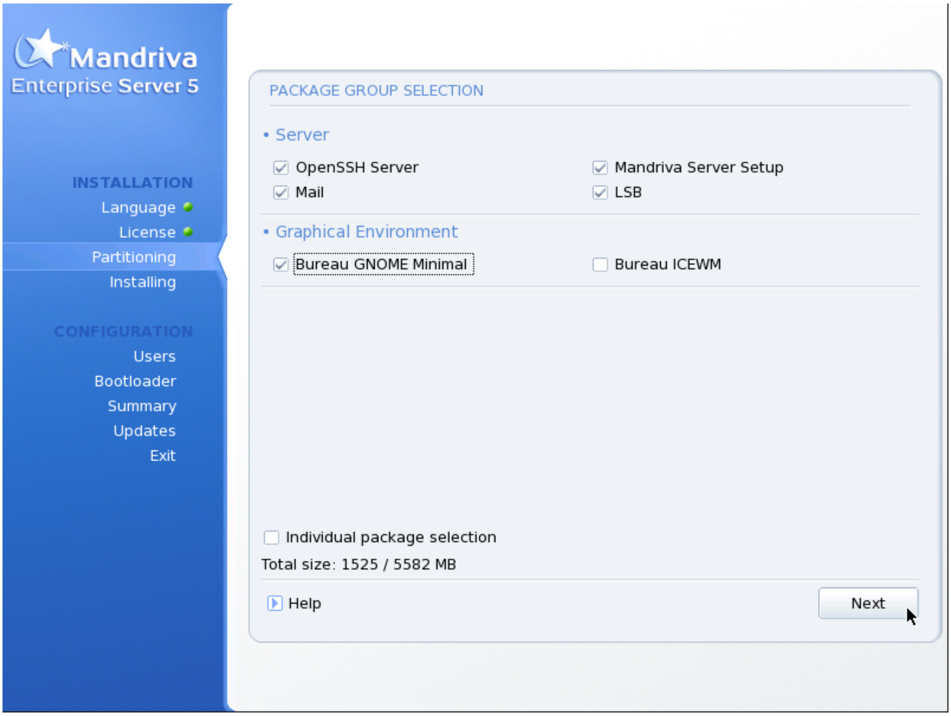

PACKAGE GROUP SELECTION

• Server

☑ OpenSSH Server ☑ Mandriva Server Setup

☑ Mail ☑ LSB

• Graphical Environment

☑ Bureau GNOME Minimal □ Bureau ICEWM

□ Individual package selection

Total size: 1525 / 5582 MB

Mandriva Enterprise Server 5 proposes quite a different way to install packages. This step is executed in two main steps:

Installation step: during the main installation you are able to choose and install packages dealing with the base system. This means the kernel stack, a minimal network stack and a graphical environment.

Graphical environent graphique is not necessary for the proper functioning of a server. However, many administration tools can be accessed only from a graphical environment.

Let's see what you can install during the installation process:

- Server A Remote access server (`OpenSSH`) and a minimal Mail Transport Agent (MTA) (`postfix`) are proposed.

  Mandriva Server Setup (mmc-wizard package) is the installation wizard. It helps you to activate the batteries you need for your server through a web interface.

- Graphical Environment : by default, Mandriva Enterprise Server 5 proposes to install a light Gnome graphical environment. You can also choose a very light one such as IceWM, but this system also operates smoothly without a graphical environment.

Moving the mouse cursor over a group name displays a short explanatory text about that group.

You can check the Individual package selection box, which is useful if you're familiar with the packages being offered, or if you want to have total control over what will be installed.

If you start the installation in Upgrade mode, you can deselect all groups and prevent the installation of any new packages. This is useful to repair or update an existing system.

Minimal Installation

If you deselect all groups when performing a regular installation (as opposed to an upgrade), a new dialog shows up after clicking Next, suggesting different options for a minimal installation:

- With X: installs the minimum number of packages possible to have a working graphical desktop.
- With basic documentation: installs the base system plus basic utilities and their documentation. This installation is suitable to set up a server.
- Truly minimal install installs the absolute minimum number of packages necessary to get a working Linux system. With this installation, you will only have a command-line interface. urpmi will not be installed !

2.7.3. Choosing Individual Packages to Install

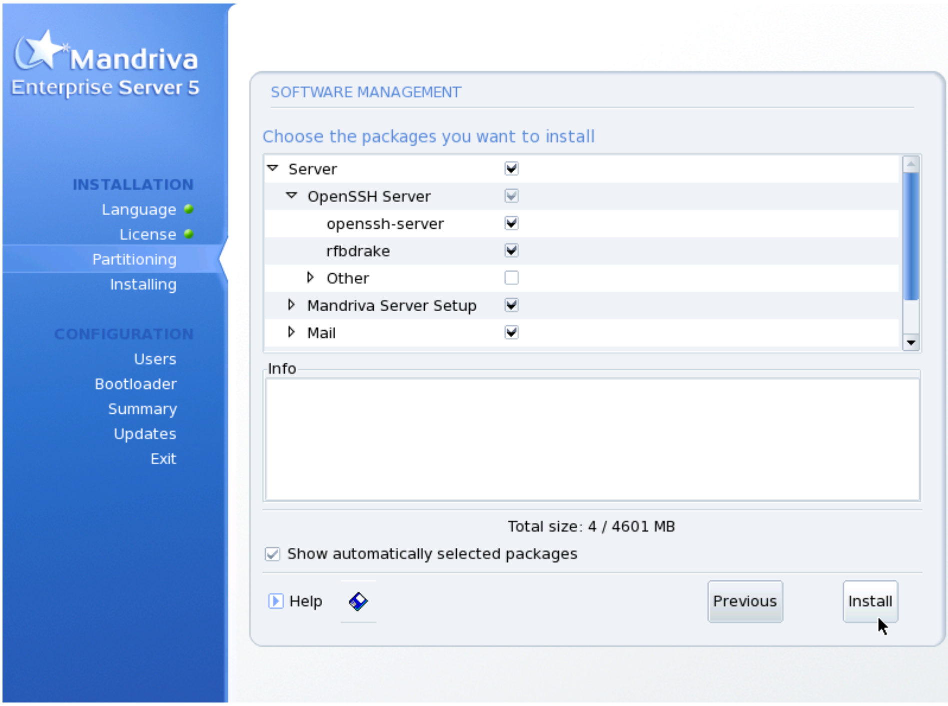

Choose the packages you want to install

- Server ☑
  - OpenSSH Server ☑
    - openssh-server ☑
    - rfbdrake ☑
    - Other □
  - Mandriva Server Setup ☑
  - Mail ☑

Total size: 4 / 4601 MB

☑ Show automatically selected packages

If you choose to install packages individually, the installer presents a tree structure containing all packages classified by groups and subgroups. While browsing the tree, you can select entire groups, subgroups, or individual packages.

Whenever you select a package on the tree, a description appears on the right to tell you the purpose of that package.

If a server package has been selected, either because you specifically chose the individual package or because it was part of a group of packages, you are asked to confirm that you really want those server packages to be installed. By default, Mandriva Linux automatically starts any installed services (servers) at boot time. Even if they are safe and have no known issues at the time the distribution was shipped, it is entirely possible that security holes were discovered after this version of Mandriva Linux was finalized. If you don't know what a particular service is supposed to do or why it's being installed, then click No.

The Show automatically selected packages option is used to disable the warning dialog. These appear whenever the installer automatically selects a package to resolve a dependency issue. Some packages depend on others and the installation of one particular package may require the installation of others. The installer can determine which packages are required to satisfy a dependency and to successfully complete the installation.

The little floppy disk icon at the bottom of the list allows you to load or save the list of packages. This is useful if you have a number of machines that you wish to configure identically. Click on this icon and select whether you wish to Load or Save the list of packages, then select the medium in the following screen and click on OK.

2.8. Adding a User

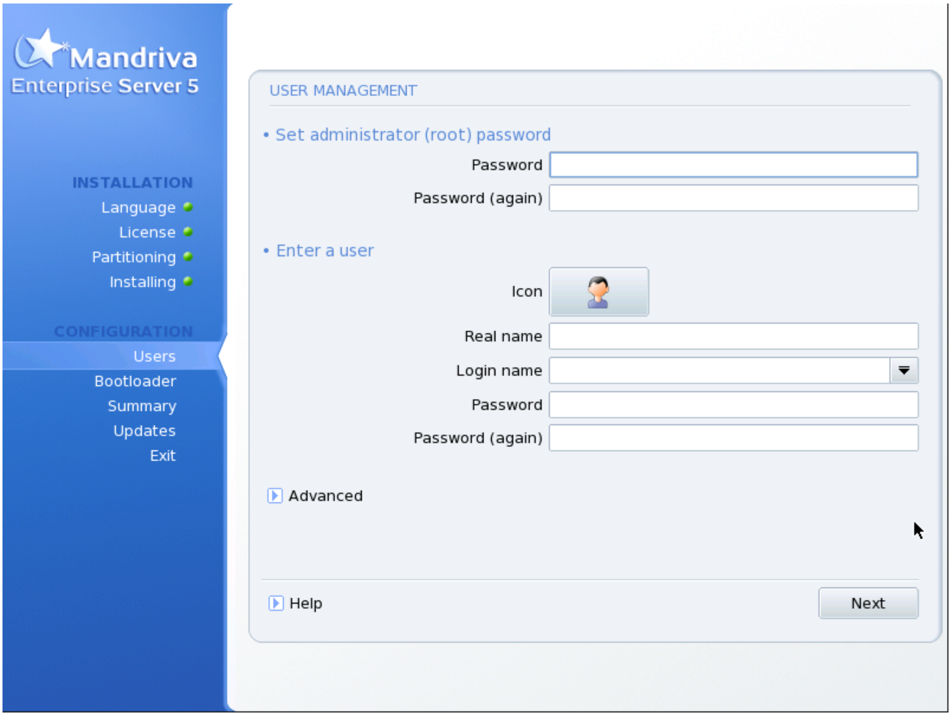

GNU/Linux is a multi-user system which means each user can have his own preferences, files and so on. But unlike the system administrator called `root`, the users you add at this point are not authorized to change anything except their own files and their own configuration, protecting the system from unintentional or malicious changes which could have a serious impact on it.

You must create at least one regular user for yourself — this is the account which you should use for routine, day-to-day usage. Although it’s very easy to log in as `root` to do anything and everything, it may also be very dangerous! A very simple mistake could render your system unusable. If you make a serious mistake as a regular user, the worst that can happen is that you'll lose some information, but you won't affect the entire system.

You are first asked for a real name. DrakX uses the first word you type in this field and copies it, all in lowercase, to the Login name field, which is the name this user must enter to log on to the system. Then enter a password, twice (for confirmation). From a security point of view, a non-privileged (regular) user's password isn't as crucial as the `root` password, but that's no reason to neglect it by making it blank or too simple: after all, **your** files could be the ones at risk.

Once you click on Accept user you can add other users. Add a user for each one of the system's users and click Next when you're finished.

Clicking on Advanced allows you to change the default `shell` for that user (bash by default), and to manually choose the user and group IDs for that user. You will also be able to create non-local users, based on LDAP or NIS directories, Windows domains or Active Directory.

2.9. Installing a Bootloader

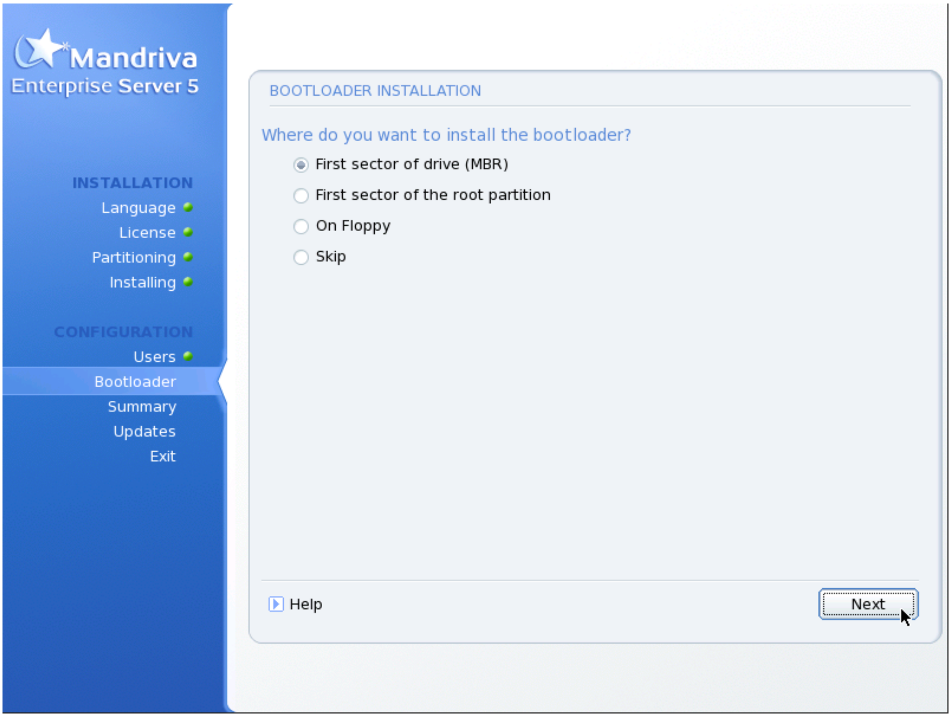

BOOTLOADER INSTALLATION

Where do you want to install the bootloader?

- ● First sector of drive (MBR)
- ○ First sector of the root partition
- ○ On Floppy
- ○ Skip

A bootloader is a small program which is started by the computer at boot time. It's responsible for starting up the whole system. Normally, the bootloader installation is totally automated. DrakX analyzes the disk boot sector and acts according to what it finds. It asks you where it should place the boot loader. Generally, the First sector of drive (MBR) is the safest place.

Choosing Skip won't install a bootloader. Use this option only if you know what you're doing.

2.10. Checking Miscellaneous Parameters

2.10.1. Summary

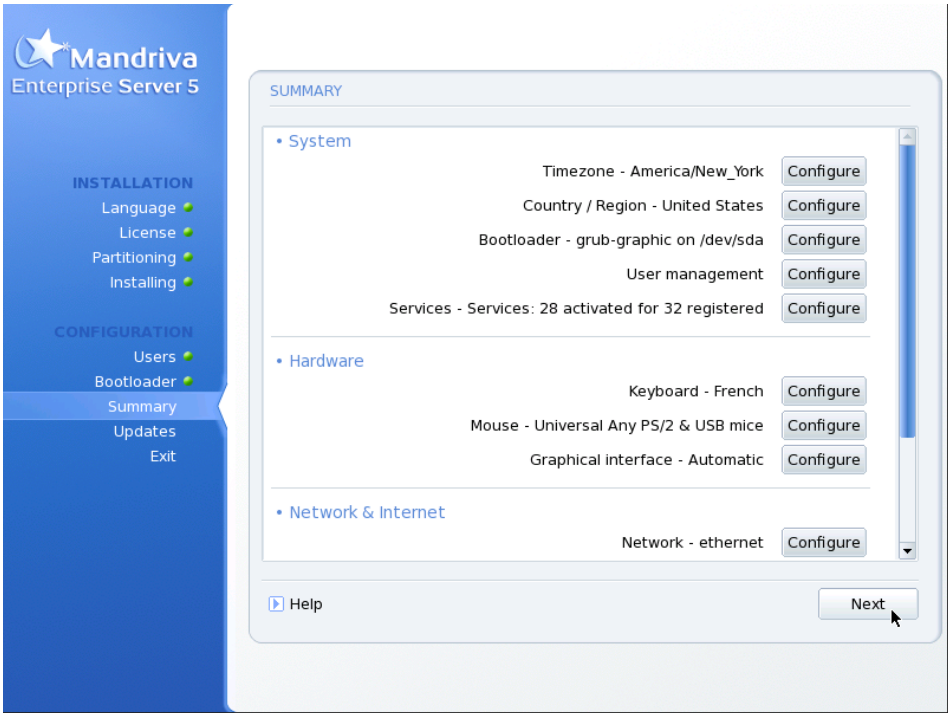

SUMMARY

- System

Timezone - America/New_York

Country / Region - United States

Bootloader - grub-graphic on /dev/sda

User management

Services - Services: 28 activated for 32 registered

- Hardware

Keyboard - French

Mouse - Universal Any PS/2 & USB mice

Graphical interface - Automatic

- Network & Internet

Network - ethernet

As a review, DrakX presents a summary of the information it gathered about your system. Depending on the hardware installed on your machine, you may have some or all of the following entries. Each entry is made up of the hardware item to be configured, followed by a quick summary of the current configuration. Click on the corresponding Configure button to make any changes.

- Keyboard: check the current keyboard map configuration and change it if necessary.
- Country/Region: check the current country selection. If you're not in the country selected by DrakX, click on the Configure button and choose another one. If your country isn't in the list shown, click on Other Countries to get a complete list of countries.
- Timezone: by default, DrakX deduces your time zone based on the country you have chosen. Click on Configure if your time zone isn't correct.
- Mouse: verify the current mouse configuration and change it if necessary.

- Printer: clicking on Configure opens the printer configuration wizard.
- Sound card: if a sound card is detected on your system, it will be displayed here. If you notice the sound card isn't the one actually present on your system, you can click on the button and choose a different driver.
- Graphical Interface: by default, DrakX configures your graphical interface with a resolution that best matches your video card and monitor combination. If that doesn't suit you, or DrakX could not automatically configure it (not configured is displayed), click on Configure to reconfigure your graphical interface. If you can't configure your graphical interface correctly, click on Help to try to solve your problem.
- Network: here's where you can configure your Internet or local network access. Use the Mandriva Control Center after the installation has finished to benefit from full contextual help.
- Proxies: allows you to configure HTTP and FTP proxy addresses if the machine you're installing on is to be located behind a proxy server.
- Security Level: this entry allows you to redefine the security level.
- Firewall: if you plan to connect your machine to the Internet, it's a good idea to protect yourself from intrusions by setting up a firewall.
- Bootloader: to change your bootloader configuration. This should be reserved to advanced users. Refer to contextual help about bootloader configuration in the Mandriva Linux Control Center.
- Services: with this entry you can fine tune which services will run on your machine. It's a good idea to review this setup to check what services you "really" need.

2.10.2. Time-Zone Options

This setup allows to refine the time zone you are currently located in. After you've chosen the location nearest to your time zone, two more options for time management are shown.

Hardware clock set to GMT. GNU/Linux manages time in GMT (Greenwich Mean Time) and translates it to local time according to the time zone you selected. If the clock on your computer is set to local time, you may deactivate this by deselecting Hardware clock set to GMT, which will let GNU/Linux know that the system clock and the hardware clock are in the same time zone. This is useful when the machine also hosts another operating system.

Automatic time synchronization. This option will automatically regulate the system clock by connecting to a remote time server on the Internet. For this feature to work, you must have a working Internet connection. We recommend that you choose a time server located near you or the generic World

Wide entry which will select the best server for you. This option actually installs a time server which can be used by other machines on your local network as well.

2.10.3. Configuring The X Graphical Server

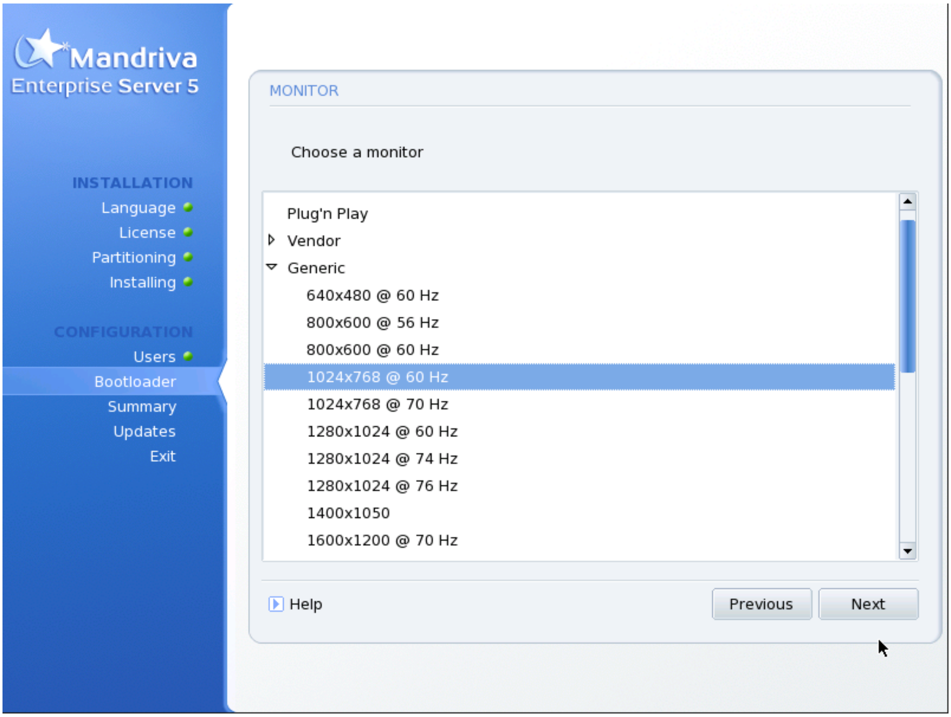

Choose a monitor

Plug'n Play
Vendor
Generic
640x480 @ 60 Hz
800x600 @ 56 Hz
800x600 @ 60 Hz
1024x768 @ 60 Hz
1024x768 @ 70 Hz
1280x1024 @ 60 Hz
1280x1024 @ 74 Hz
1280x1024 @ 76 Hz
1400x1050
1600x1200 @ 70 Hz

X stands for “X Window System” and is the heart of the Linux graphical interface on which all graphical environments are based.

A list of different parameters is available to optimize your graphical display.

Graphic Card

If everything works fine, the installer should detect and configure the video card installed on your machine. If the detection or configuration is incorrect, you can choose the card installed on your system from a list.

Monitor

If the installer fails to detect or configure your monitor properly, you can choose from this list the monitor which is connected to your computer.

Resolution

Here you can choose the resolution and color depth from those available for your graphics hardware. Choose the one which best suits your needs (you will be able to make changes after the installation). A sample of the chosen configuration will be shown in the monitor picture.

Test

Depending on your hardware, this entry might not appear.

The system will try to open a graphical screen at the desired resolution. By answering Yes once you see test message, DrakX proceeds to the next step. If you don't see the test message, it means that some part of the autodetected configuration was incorrect and the test will automatically end after a few seconds and return you to the menu. Change settings until you get a correct graphical display.

Options

This step allows you to choose whether you want your machine to automatically switch to a graphical interface at boot. Obviously, you may want to select No if your machine is to act as a server, or if you were not successful in getting the display configured.

2.10.4. Configure Your Network

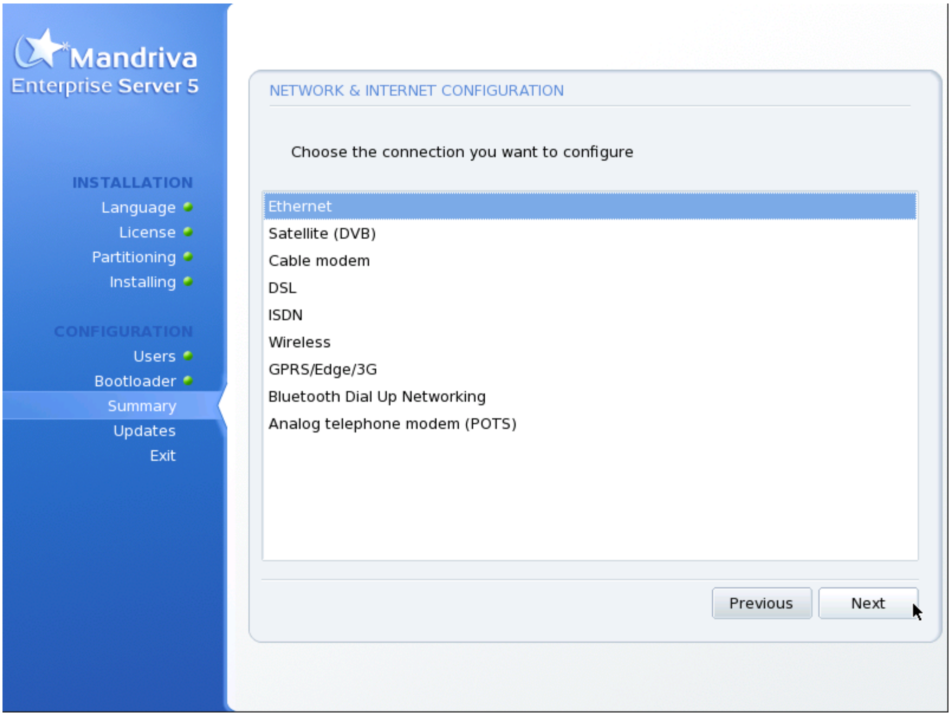

NETWORK & INTERNET CONFIGURATION

Choose the connection you want to configure

Ethernet
Satellite (DVB)
Cable modem
DSL
ISDN
Wireless
GPRS/Edge/3G
Bluetooth Dial Up Networking
Analog telephone modem (POTS)

You will now set up your Internet/network connection. Click Next if you wish to connect your computer to the Internet or to a local network. Mandriva Linux will attempt to auto-detect network devices and modems. If this detection fails, uncheck the Use auto detection box. You may also choose not to configure the network, or to do it later, in which case clicking the Cancel button will take you to the next step.

When configuring your network, the available connection options are: normal modem connection, winmodem connection, ISDN modem, ADSL connection, cable modem, and finally a LAN connection (Ethernet).

We will not detail each configuration option — just make sure that you have all the parameters, such as your IP address, the default gateway, the addresses of your DNS servers, etc., from your Internet Service Provider or system administrator.

2.10.5. Installing a Bootloader

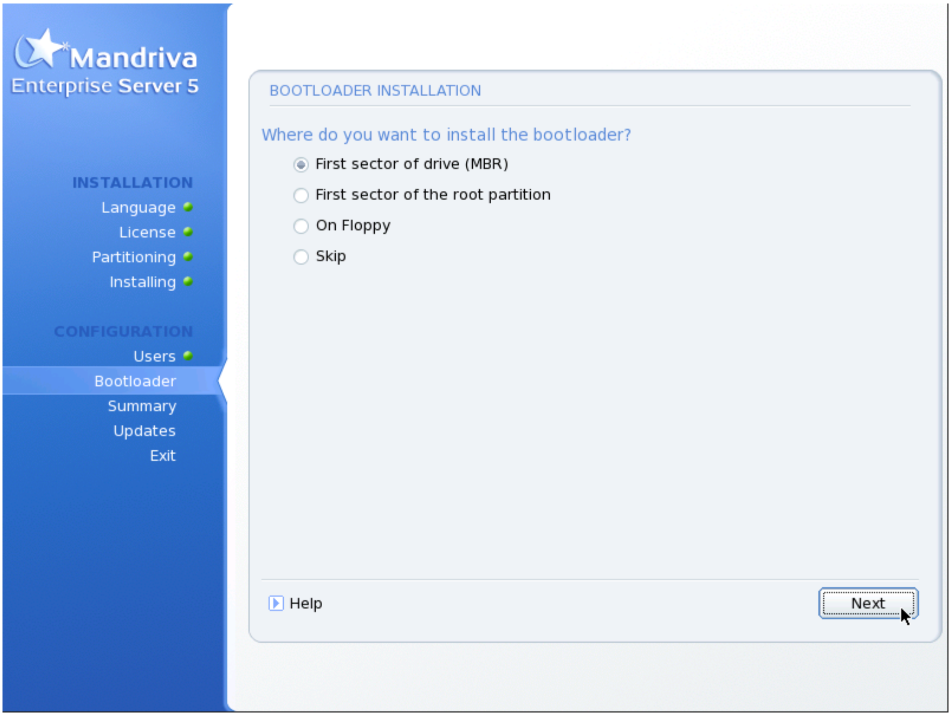

BOOTLOADER INSTALLATION

Where do you want to install the bootloader?

● First sector of drive (MBR)

○ First sector of the root partition

○ On Floppy

○ Skip

This dialog allows you to fine-tune your bootloader:

- Bootloader to use: there are three choices for your bootloader:

  1. GRUB with text menu: if you prefer grub with a text interface.
  2. GRUB with graphical interface: if you prefer grub with a graphical interface.
  3. LILO with text menu: if you prefer LILO with its text menu interface.
  4. LILO with a graphical menu: if you prefer LILO with its graphical interface.

- Boot device: in most cases, you will not change the default (/dev/hda), but if you prefer, the bootloader can be installed on the second hard drive (/dev/hdb), or even on a floppy disk (/dev/fd0).
- Delay before booting the default image: after a boot or a reboot of the computer, this is the delay given to the user at the console to select a boot entry other than the default.
- Enable ACPI: the ACPI for power management standard appeared during year 2002, particularly for laptops. If you know your hardware supports it

and you need it, check this box.

- Force no APIC: If you have noticed hardware problems on your machine (IRQ conflicts, instabilities, machine freeze, etc.) you should try disabling APIC by checking this box.

Be aware that if you choose not to install a bootloader (by selecting Skip), you must ensure that you have a way to boot your Mandriva Linux system! Be sure you know what you are doing before changing any of the options.

Click the Advanced button to view options reserved to expert users.

2.10.6. Selecting Available Services at Boot Time

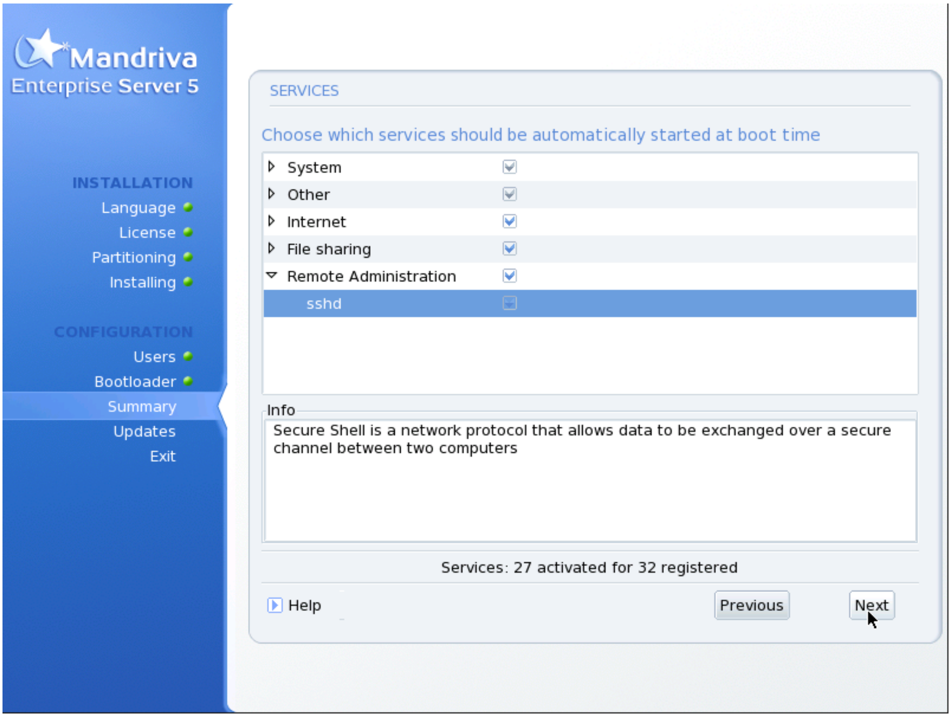

SERVICES

Choose which services should be automatically started at boot time

- System ☑
- Other ☑
- Internet ☑
- File sharing ☑
- Remote Administration ☑
  - sshd □

Info

Secure Shell is a network protocol that allows data to be exchanged over a secure channel between two computers

Services: 27 activated for 32 registered

This dialog is used to select which services you wish to start at boot time.

DrakX lists all services available on the current installation. Review each of them carefully and uncheck those which aren't needed at boot time.

A short explanatory text is displayed about a service when it is selected. However, if you're not sure whether a service is useful or not, it is safer to leave the default setting.

At this stage, be very careful if you intend to use your machine as a server: you probably don't want to start any services which you don't need. Please remember that some services can be dangerous if they're enabled on a server. In general, select only those services you **really** need.

Chapter 3. Mandriva Server Setup (mmc-wizard)

3.1. Use Mandriva Server Setup

Mandriva Server Setup is installation wizard for Mandriva Enterprise Server 5. Its main goal is to make server functionnalities easy to install and give a way to add some more.

If you did not check Mandriva Server Setup box during installation, you can install it at anytime. Just install mmc-wizard package using command line or rpmdrake.

Mandriva Server Setup

When Mandriva Server Setup is installed, it's available through a web browser using following address: https:/IP_server_MES5/mmc-wizard/

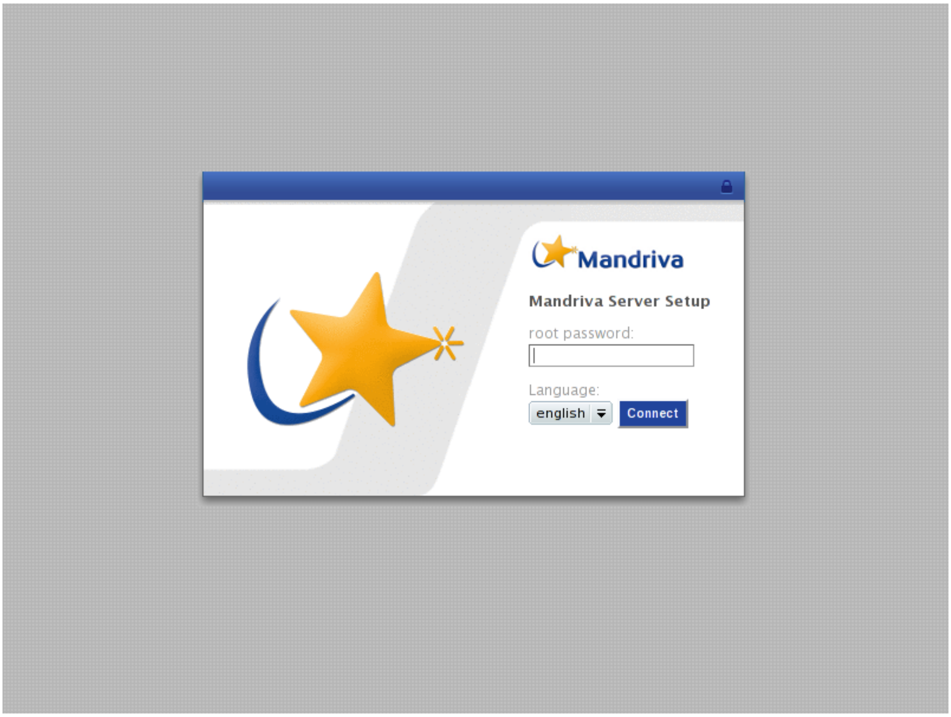

Figure 3-1. Login Mandriva Server Setup

The only way to connect is using root account.

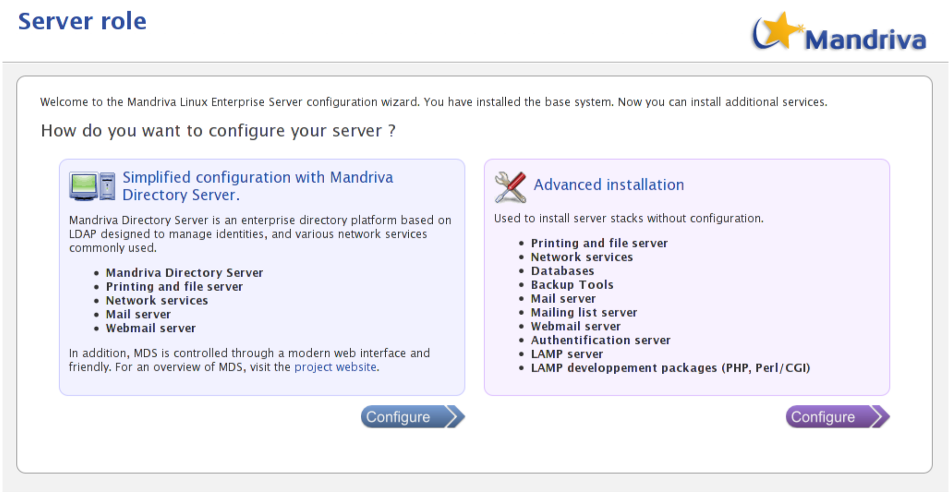

Server role

Welcome to the Mandriva Linux Enterprise Server configuration wizard. You have installed the base system. Now you can install additional services.

How do you want to configure your server ?

Simplified configuration with Mandriva Directory Server.

Mandriva Directory Server is an enterprise directory platform based on LDAP designed to manage identities, and various network services commonly used.

- Mandriva Directory Server
- Printing and file server
- Network services
- Mail server
- Webmail server

In addition, MDS is controlled through a modern web interface and friendly. For an overview of MDS, visit the project website.

Advanced installation

Used to install server stacks without configuration.

- Printing and file server
- Network services
- Databases
- Backup Tools
- Mail server
- Mailing list server
- Webmail server
- Authentification server
- LAMP server
- LAMP developpement packages (PHP, Perl/CGI)

Figure 3-2. Mandriva Server Setup Home

Mandriva Server Setup is divided in 2 main sections:

- Basic configuration with Mandriva Directory Server: you can install and configure very easily Mandriva Directory Server, enterprise directory based on OpenLDAP. Mandriva Directory Server has been developped in a modular way, you can choose émong available modules. -> See "Mandriva Directory Server Stack"
- Advanced Configuration: You will find there all most used middleware stacks for server. These stacks come with a minimal configuration and are not integrated to MDS.

Some of applications stacks can be found in both sections of Mandriva Server Setup. If a stack is already installed (all stack or part of it), second part of stack will not be available anymore (gray check boxes).

3.2. MDS Stacks and simplified configuration

3.2.1. Main Concept

Mandriva Directory Server is an enterprise directory service based on OpenLDAP which allows to manage IT infrastructure user profiles and accounts, so as various widely used network services (Samba, dns, dhcp, mail...).

For a simplified use, Mandriva Directory Server can be managed through a user-friendly and modern web based user interface. The Mandriva Directory Server chapter of this documentation deals about the use of the graphical interface.

For a global view of Mandriva Directory Server, please visit the dedicated project (`http://mds.mandriva.org/`).

3.2.2. Installation and Configuration

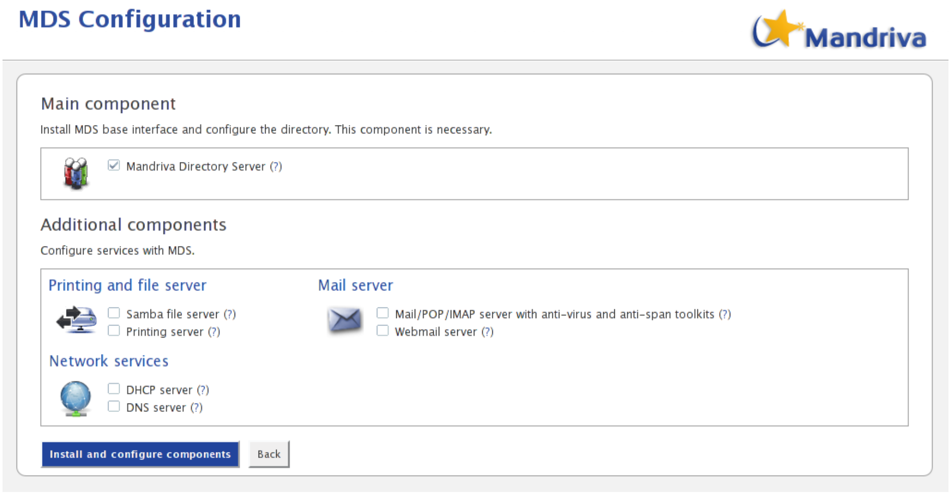

MDS Configuration

Main component

Install MDS base interface and configure the directory. This component is necessary.

☑ Mandriva Directory Server (?)

Additional components

Configure services with MDS.

Printing and file server

□ Samba file server (?)

□ Printing server (?)

Mail server

□ Mail/POP/IMAP server with anti-virus and anti-span toolkits (?)

□ Webmail server (?)

Network services

□ DHCP server (?)

□ DNS server (?)

Figure 3-3. Mandriva Directory Server Installation Page

3.2.2.1. Main Component: Mandriva Directory Server

Mandriva Directory Server Packages and Basic Configuration. Checked by default, this component is necessary to the other modules linked to Mandriva Directory Server.

A window appears during the installation process of this component. Click on the Details button to see the console.

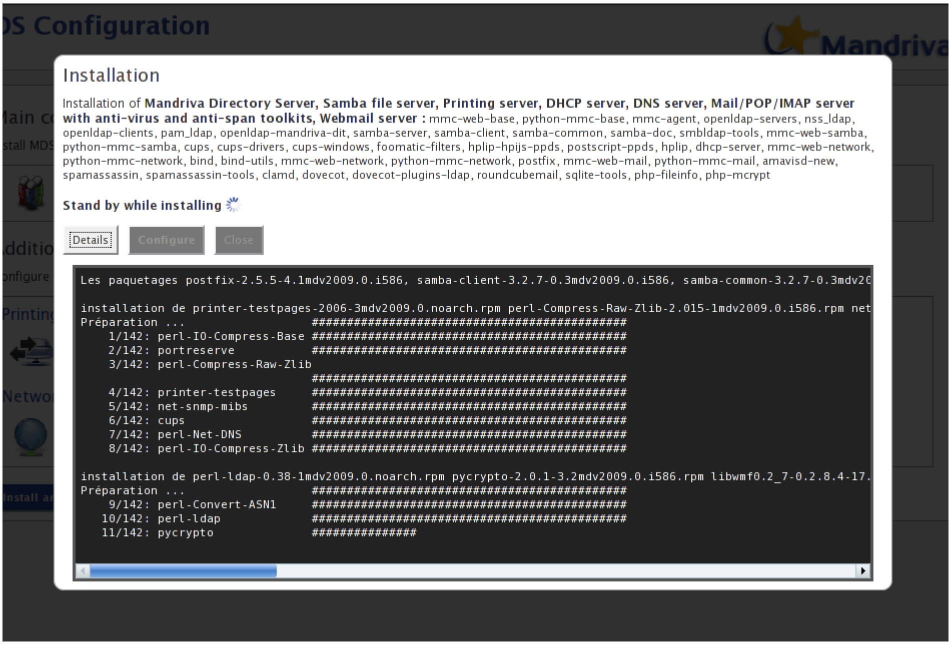

Installation

Installation of **Mandriva Directory Server, Samba file server, Printing server, DHCP server, DNS server, Mail/POP/IMAP server with anti-virus and anti-span toolkits, Webmail server :** mmc-web-base, python-mmc-base, mmc-agent, openldap-servers, nss_ldap, openldap-clients, pam_ldap, openldap-mandriva-dit, samba-server, samba-client, samba-common, samba-doc, smbldap-tools, mmc-web-samba, python-mmc-samba, cups, cups-drivers, cups-windows, foomatic-filters, hplip-hpijs-ppds, postscript-ppds, hplip, dhcp-server, mmc-web-network, python-mmc-network, bind, bind-utils, mmc-web-network, python-mmc-network, postfix, mmc-web-mail, python-mmc-mail, amavisd-new, spamassassin, spamassassin-tools, clamd, dovecot, dovecot-plugins-ldap, roundcubemail, sqlite-tools, php-fileinfo, php-mcrypt

Stand by while installing

```
Les paquetages postfix-2.5.5-4.1mdv2009.0.i586, samba-client-3.2.7-0.3mdv2009.0.i586, samba-common-3.2.7-0.3mdv20

installation de printer-testpages-2006-3mdv2009.0.noarch.rpm perl-Compress-Raw-Zlib-2.015-1mdv2009.0.i586.rpm net
Préparation ...                  #############################################
    1/142: perl-IO-Compress-Base #############################################
    2/142: portreserve           #############################################
    3/142: perl-Compress-Raw-Zlib
                                 #############################################
    4/142: printer-testpages     #############################################
    5/142: net-snmp-mibs         #############################################
    6/142: cups                  #############################################
    7/142: perl-Net-DNS          #############################################
    8/142: perl-IO-Compress-Zlib #############################################

installation de perl-ldap-0.38-1mdv2009.0.noarch.rpm pycrypto-2.0.1-3.2mdv2009.0.i586.rpm libwmf0.2_7-0.2.8.4-17.
Préparation ...                  #############################################
    9/142: perl-Convert-ASN1     #############################################
   10/142: perl-ldap             #############################################
   11/142: pycrypto              ###############
```

Figure 3-4. Installation Details: Console View

At the end of the installation process, click on Configure.

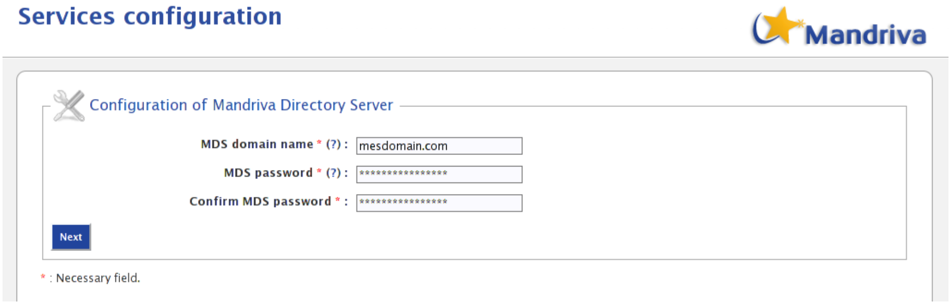

Configuration of Mandriva Directory Server

MDS domain name * (?) : mesdomain.com

MDS password * (?) : ****************

Confirm MDS password * : ****************

* : Necessary field.

Figure 3-5. Mandriva Directory Server Configuration Page

You are asked to fill in the domain name (for example: domain.com) managed by Mandriva Directory Server and a Mandriva Directory Server dedicated administration password. This password will be necessary for every modules Mandriva Directory Server installation.

A summary of the configuration is displayed.

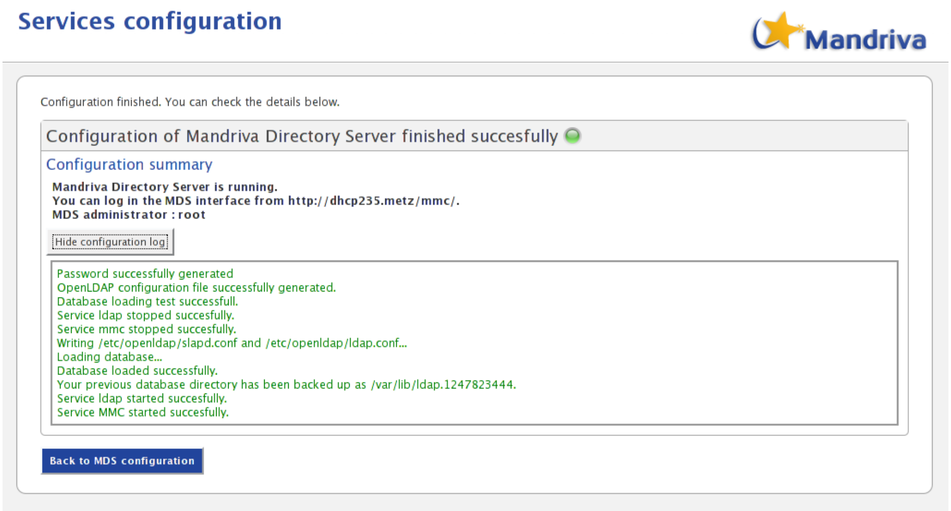

Services configuration

Configuration finished. You can check the details below.

Configuration of Mandriva Directory Server finished succesfully

Configuration summary

Mandriva Directory Server is running.
You can log in the MDS interface from http://dhcp235.metz/mmc/.
MDS administrator : root

```
Password successfully generated
OpenLDAP configuration file successfully generated.
Database loading test successfull.
Service ldap stopped succesfully.
Service mmc stopped succesfully.
Writing /etc/openldap/slapd.conf and /etc/openldap/ldap.conf...
Loading database...
Database loaded successfully.
Your previous database directory has been backed up as /var/lib/ldap.1247823444.
Service ldap started successfully.
Service MMC started succesfully.
```

Figure 3-6. Mandriva Directory Server Configuration Result Page

You can notice that the MDS graphical interface is available through http://IP_serveur_MES5/mmc/. The administrator user of MDS is root.

If you can't connect to the MDS interface from another computer than the server, check if your firewall allows web type requests (80 and 443 ports).

Here is the list of packages installed by this component:

- `mmc-web-base`
- `python-mmc-base`
- `mmc-agent`
- `openldap-servers`
- `nss_ldap`
- `openldap-clients`

3.2.2.2. Printing and Files Server

This component contains the packages and the Mandriva Directory Server configuration modules for the administration of the files and printing shares of Microsoft.

3.2.2.2.1. Samba File Server

During the configuration process, you will be asked to fill in the Samba domain name as well as the Samba server name and the Microsoft domain. Finally, you have to define the Samba administrator password.

The Samba service is configured as the principal domain controler (PDC) for your Microsoft network. Windows computers could join the domain you specified during the configuration step. The "admin" user created during the installation process will allow you to administrate the domain.

For your users to authenticate on the domain, they have to be member of the "Domain Users" group. After the component installation, every newly created user will be member of the "Domain Users" group.

Here is the list of packages installed by the Samba component:

- `samba-server`
- `samba-client`
- `samba-common`
- `samba-doc`
- `smbldap-tools`
- `mmc-web-samba`
- `python-mmc-samba`

3.2.2.2.2. Cups Printing Server

The Cups printing server allows you to share printers installed on your server.

Here is the list of packages installed with the Cups printing server:

- `cups`
- `cups-drivers`
- `cups-windows`
- `foomatic-filters`
- `hplip-hpijs-ppds`
- `postscript-ppds`
- `hplip`

More information about the cups-windows package:

```
The cupsaddsmb command will use the CUPS v6 PostScript printer driver
for Windows available in this package.
To complete the installation, you have to add the Microsoft Postcript
driver files in the /usr/share/cups/drivers directory. These files can
be found on any system running Windows 2000 or higher in the following folder:
%WINDOWS%\SYSTEM32\SPOOL\DRIVERS\W32X86\3

After this step, your /usr/share/cups/drivers directory should contain
these files:
cups6.inf
cups6.ini
cupsps6.dll
cupsui6.dll
ps5ui.dll (from your Windows system)
pscript.hlp (from your Windows system)
pscript.ntf (from your Windows system)
pscript5.dll (from your Windows system)
```

3.2.2.3. Network Services

This component contains the Mandriva Directory Server module packages and configuration to create and manage a LAN (DNS zones and DHCP subnetwork).

During the DNS server module configuration, you can configure the networks which will be able to perform recursive requests on your DNS. A recursive request has an object as domain name of a zone that your DNS server does not know. The DNS server must then contact other DNS servers to resolve the request.

About the resolution of configured zones on your DNS server, there is no restriction on the origin of clients.

You can choose to forward all external DNS requests to another DNS server. Your DNS server will then only resolve the zones you have configured.

The DHCP component does not need any configuration into Mandriva Server Setup.

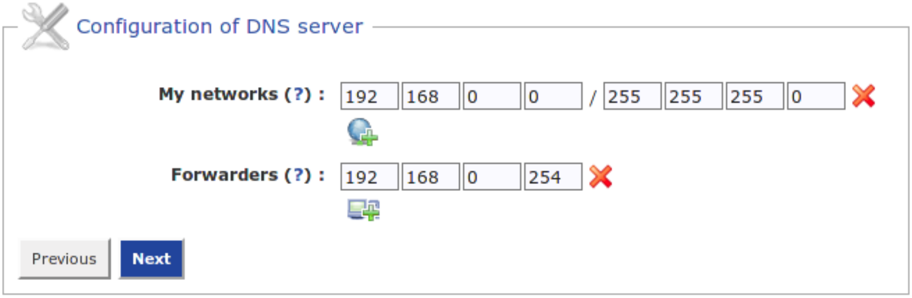

Configuration of DNS server

My networks (?) : 192 168 0 0 / 255 255 255 0

Forwarders (?) : 192 168 0 254

Figure 3-7. DNS service configuration page:

The installed packages for this component are:

- **DHCP Server DHCP.** : `dhcp-server`, `mmc-web-network`, `python-mmc-network`
- **DNS Server.** : `bind`, `bind-utils`, `mmc-web-network`, `python-mmc-network`

3.2.2.4. Mail Server

3.2.2.4.1. Mail/POP/IMAP Server and Antivirus and Antispam Tools

This module installs and configures an SMTP server of MDS (sending and receiving mails), a POP3/IMAP server (reading mails) and spams and viruses detection tools. This configuration allows you to manage as many mail domains you wish.

At the end of the installation process, you will be asked to fill in the SMTP server hostname/FQDN (for example: smtp.domain.com). Specify then which networks are authorized to send mails through Postfix, for example the 192.168.0.0 local network with a 255.255.255.0 mask.

Finally choose the protocols the server Dovecot will provide: imap imaps, pop3 pop3s or imap imaps and pop3 pop3s at the same time.

Don't forget to open the requested ports on the firewall (SMTP: 25, SMTPS: 465, POP3S: 995, IMAPS: 993). Please note that the IMAP and POP3 protocols are not enable from external interfaces.

- The installed packages for this component are:
- `postfix`
- `mmc-web-mail`
- `python-mmc-mail`

- `amavisd-new`
- `spamassassin`
- `spamassassin-tools`
- `clamd`
- `dovecot`
- `dovecot-plugins-ldap`

3.2.2.4.2. Webmail Server

This stack is not directly linked to Mandriva Directory Server. It installs the Roundcube webmail, which will allow your users to have their online messaging. You just have to enable the mail module on your users from the MDS interface if it is not already done. Users will be able to authenticate thanks to their logins and passwords on http://IP_serveur_MES5/roundcubemail.

- The installed packages for this components are:
- `roundcubemail`
- `sqlite-tools`
- `php-fileinfo`
- `php-mcrypt`

3.3. Middleware/server stacks and services (”Advanced configuration”)

3.3.1. Introduction

This module will allow you to install some more server stacks without any configuration.

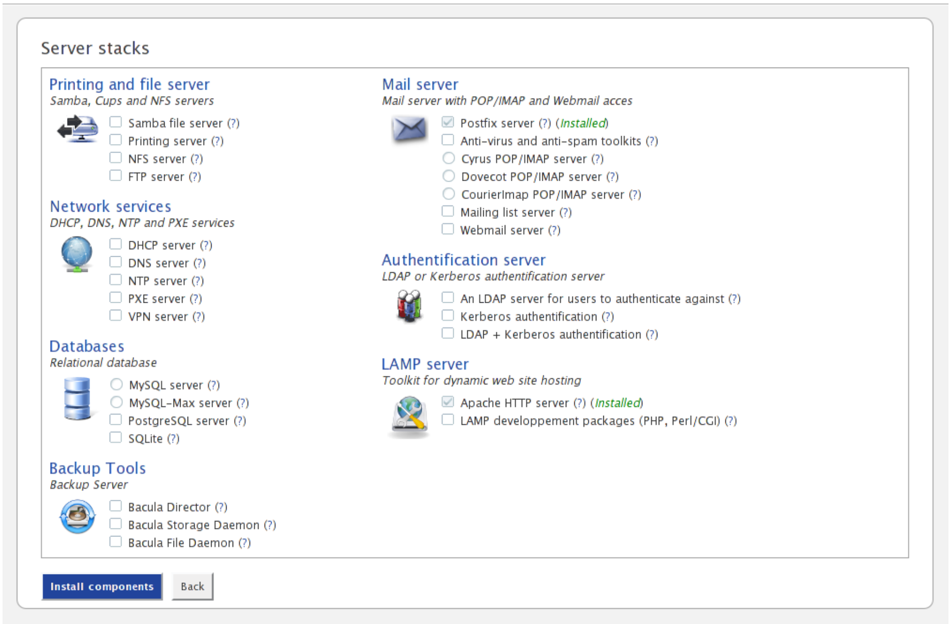

Server stacks

Printing and file server

Samba, Cups and NFS servers

□ Samba file server (?)
□ Printing server (?)
□ NFS server (?)
□ FTP server (?)

Network services

DHCP, DNS, NTP and PXE services

□ DHCP server (?)
□ DNS server (?)
□ NTP server (?)
□ PXE server (?)
□ VPN server (?)

Databases

Relational database

○ MySQL server (?)
○ MySQL-Max server (?)
□ PostgreSQL server (?)
□ SQLite (?)

Backup Tools

Backup Server

□ Bacula Director (?)
□ Bacula Storage Daemon (?)
□ Bacula File Daemon (?)

Mail server

Mail server with POP/IMAP and Webmail acces

☑ Postfix server (?) (*Installed*)
□ Anti-virus and anti-spam toolkits (?)
○ Cyrus POP/IMAP server (?)
○ Dovecot POP/IMAP server (?)
○ CourierImap POP/IMAP server (?)
□ Mailing list server (?)
□ Webmail server (?)

Authentification server

LDAP or Kerberos authentification server

□ An LDAP server for users to authenticate against (?)
□ Kerberos authentification (?)
□ LDAP + Kerberos authentification (?)

LAMP server

Toolkit for dynamic web site hosting

☑ Apache HTTP server (?) (*Installed*)
□ LAMP developpement packages (PHP, Perl/CGI) (?)

Figure 3-8. Server stacks page

Some stacks may be not available to be selected if some of them are already installed or can create conflicts with existing one.

3.3.2. Stacks review

Printing and files server

Samba, CUPS servers and NFS

- Files server Samba

  Files and printers share for Microsoft networks.

  `samba-server`

  `samba-client`

  `samba-common`

  `samba-doc`

  `smbldap-tools`

  `samba-winbind`

- Printing server

  Installation and configuration of network printings using CUPS .

  `cups`

  `cups-drivers`

  `cups-windows`

  `foomatic-filters`

  `hplip-hpijs-ppds`

  `postscript-ppds`

  `hplip`

- NFS server

  Files share with NFS

  `nfs-utils`

  `nfs-utils-clients`

Network services

DHCP, DNS, NTP or PXE services

- DHCP server

  Provides IP parameters to clients.

  `dhcp-server`

- DNS server

  Name resolution on network.

  `bind`

  `bind-utils`

- NTP server

  Time server

  `ntp`

- PXE server

  PXE server (Preboot eXecution Environment).

  `pxelinux`

  `tftp`

  `tftp-server`

`syslinux`

`dhcp-server`

Database

Install some of the most used relationnal databases.

- MySQL server.

  `mysql`

  `mysql-client`

  `phpmyadmin`

- PostgreSQL server.

  `postgresql8.3-server`

  `postgresql8.3-pl`

  `phppgadmin`

- MySQL-Max server

  Alternative MySQL server, binaries were compiled using advanced options.

  `mysql-Max`

  `mysql-client`

  `phpmyadmin`

- SQLite

  CLI tools to manage libsqlite librairy.

  `sqlite-tools`

  `phpsqliteadmin`

Mail server

Mail server, POP/IMAP access and webmail.

- Postfix server

  Mail server

  `postfix`

- Anti-virus and anti-spam toolkits

  AMaVis, ClamAV and SpamAssassin installation.

`amavisd-new`

`spamassassin`

`spamassassin-tools`

`clamd`

- POP/IMAP server

  You can choose among available one:

  POP/IMAP Cyrus server

  `cyrus-imapd`

  `cyrus-imapd-utils`

  `cyrus-sasl`

  `libsasl2-plug-plain` or `lib64sasl2-plug-plain`

  `libsasl2-plug-login` or `lib64sasl2-plug-login`

  POP/IMAP Dovecot server

  `dovecot`

  `dovecot-plugins-ldap`

  POP/IMAP CourierImap server

  `courier-imap`

  `courier-base`

- Mailing-list server

  Sympa installation.

  `sympa`

- Webmail server

  Webmail roundcube installation.

  `roundcubemail`

  `sqlite-tools`

Authentication server

LDAP and/or Kerberos server for users authentication.

- LDAP for users authentication.

  OpenLdap for users authentication.

  `openldap-servers`

`nss_ldap`

`openldap-clients`

`pam_ldap`

`openldap-mandriva-dit`

- Kerberos authentication

  Kerberos authentication server.

  `krb5-server`

  `krb5-workstation`

- LDAP + Kerberos authentication

  LDAP + Kerberos authentication server.

  `openldap-servers`

  `nss_ldap`

  `openldap-clients`

  `pam_ldap`

  `openldap-mandriva-dit`

  `krb5-server`

  `krb5-workstation`

  `libsasl2-plug-gssapi` or `lib64sasl2-plug-gssapi`

LAMP server

A set of free tools to host dynamic web sites.

- Apache HTTP server

  Web server installation.

  `apache-base`

  `apache-mpm-prefork`

  `apache-conf`

  `apache-modules`

  `apache-mod_ssl`

- Development packages for LAMP (PHP, Perl/CGI)

  A set of free tools to host dynamic web sites (script languages, php modules, ...).

  `apache-mod_perl`

`apache-mod_php`

`php-dom`

`php-simplexml`

`php-xml`

`php-xmlrpc`

`php-xsl`

`php-cli`

`php-mysql`

`php-pgsql`

`php-sqlite`

Backup tools

Bacula is a set of programs for backups management but also to restore or check data on a heterogeneous network.

Included version in Mandriva Enterprise Server 5 is Bacula 3. For more information on this project, go to official web site (`http://www.bacula.org`) .

"It comes by night and sucks the vital essence from your computers."

- Bacula Director

  Bacula Director service is a supervisor application for all backup, restore and check operations.

  `bacula-common`

  `bacula-dir-common`

  `bacula-dir-mysql`

- Bacula Storage Daemon

  Bacula Storage Daemon can backup data and metadata and restore it. Storage daemoni can manage read and write operation on storage devices.

  `bacula-sd`

- Bacula File Daemon

  Bacula File Daemon must be installed as a backup client on hosts to be backuped It will provide metadata required by Director.

  `bacula-fd`

3.4. Limit access to Mandriva Server Setup

Once you have configured your server with Mandriva Server Setup we advise you to restrain its access as this can access to your system sensitive parts. For this, two methods can be used:

3.4.1. Disable the mmc-wizard service

Disabling the mmc-wizard service will prevent anyone from doing operations on your server thanks to the configuration interface.

To disable the service you can use the Mandriva Control Center (MCC) available from System->Manage Services, or by launching drakxservices from a command line tool. Stop manually the service and uncheck the "At boot" option.

3.4.2. Forbid the access to the web configuration interface from the network

This method is interesting because the mmc-wizard service remains active, but will only be accessible on the server. Every connexion from a network client will be denied.

For this, edit the `/etc/httpd/conf/webapps.d/mmc-wizard.conf` file and replace the

```
Allow from all
```

line by

```
Allow from 127.0.0.1
  Deny all
```

Chapter 4. Auto-install mode

4.1. Managing auto-installation

Mandriva Enterprise Server 5 provides an auto-installation method. It allows you to reproduce a given installation scenario that you can personnalize according to your needs. It is customizable for a full or partial installation. You will be able to install your server in a partly or fully non interactive way. This section reviews auto-install functionnality.

Comprenhensive documentation is availabe in the `drakx-autoinstall-doc` package. It contains detailed documentation about auto-install configuration and use.

The automated installation feature of DrakX is controlled by the contents of a file named `auto_inst.cfg`. This file is generally located on the boot floppy diskette. PXE can also provide it, so that you can install automatically.

During a manual install, the various declarations are created and the appropriate fields filled in as you make choices from the various screens. Then, when you created the Automated or Replay diskette, selected portions of this structure were simply dumped to a file that will control the actions of DrakX when an Automated or Replay install is done..

An Automated Install requires that all the choices be pre-selected using either the file generated by DrakX, or, manually by you. During each installation, DrakX creates a template based on your choice and called `/root/drakx/auto_inst.cfg.pl`. You can edit this file and modify some of the values.

Auto-installation allows you to automatically manage:

- partitions
- packages choice
- network configuration
- user creation
- authentication
- X configuration
- ...

You can also add some more actions for the post-installation steps, using some `bash` commands for example.

Here is a short extract of an example file:

```
$o = {
 'printer' => {
  'configured' => {}
},
 'default_packages' => [
  'kernel-2.6.27.21-1mnb',
  'vim-enhanced',
  'grub',
  'lilo',
  'vim-minimal',
  ...
  ]
  'net' => {
 'zeroconf' => {
  'hostname' => undef
 },
 'network' => {
  'NETWORKING' => 'yes',
  'DHCP' => 'yes',
  'NET_DEVICE' => 'eth0',
  'NET_INTERFACE' => 'eth0'
  },
  ...
  }
```

In this example, printers are not configured, neither is zeroconf. Default packages choice contains packages like grub and `vim-enhanced`. Network is configured using dhcp on interface `eth0`.

Once your file is ready, you can use it in different ways:

- Using the install CD or DVD and floppy: copy `auto_install.cfg` on a floppy. Then boot with both install medium and floppy. On first screen, press **F1** and on prompt, use this command: `linux kickstart=floppy`.
- Using PXE: you can define a specific image so that it takes into account the auto-installation file. For example: the PXE server and data are stored on an nfs share. You should use this configuration:

  You should then copy `auto_install.cfg` in the `install` directory of your repository. You can rename it using `testauto` for example. Then, if the client uses `mes5auto` image, it will automatically use the auto-installation file without using other media.

You can quite easily check the syntax of `auto_inst.cfg.pl`using a perl check tool. Use the following command:

```
# perl -cw /root/drakx/auto_inst.cfg.pl
auto_inst.cfg.pl syntax OK
```

Introduction to Service Administration

Mandriva Enterprise Server 5 provides graphical tools that can help you to configure your system and services. We will quickly review two of them: Mandriva Control Center and Webmin. Then we will describe each service and provide advanced configuration tips using the command line or the graphical interface.

1. Using Mandriva Control Center

The Mandriva Linux Control Center (MCC) enables system administrators to configure hardware and services used by all users in a friendly way.

Access the Mandriva Linux Control Center through the main menu (System+Configuration→Configure Your Computer.

Some of the Mandriva Linux Control Center components are also available from the command line in text mode by running `drakconf`.

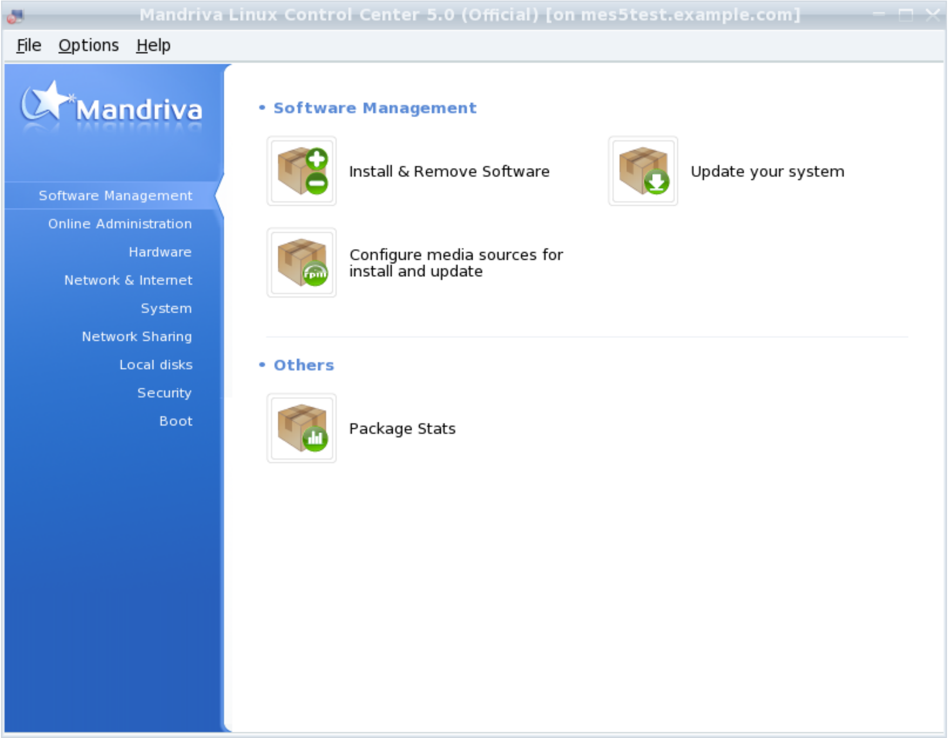

Figure 1. The Control Center's Main Window

Here are some of the available menu entries:

- **Options→Display Logs.** When activated this option displays a Tools Logs window. It shows all system modifications made by the configuration tools launched from within the Centre de contrôle Mandriva Linux.
- **Options→Expert mode.** Gives you access to some of the more advanced tools, which are shown in the table below.
- **Profiles.** This menu gives you access to the configuration profile features.
- **Help→Help.** Opens the help browser which displays documentation about the active configuration tool.
- **Aide→Report Bug.** Allows you to report a bug to the development team.

The tools are sorted into categories. The following table lists them all and refers to the corresponding sections of this manual.

Some more categories appear if the drakwizard package is installed. The documentation for those wizards is available on-disk. Those wizards enable you to do basic configuration of common LAN services such as web, FTP, mail and database servers.

Chapter 5. Base system

Setting up a Basic Mandriva Enterprise Server

This chapter does not aim to present a global overview of GNU/Linux's base operations, but to present Mandriva Enterprise Server 5's new functionalities and key elements. The definition of the base system lies on the following scheme derived from this distribution's architecture.

This chapter tackles features provided by the package management tools (Section 5.1), implementing virtual machines through virtualization (Section 5.3), and load balancing (Section 5.2).

5.1. Software Management Tools

Once your server is installed, you may need to install or uninstall software. Mandriva Enterprise Server 5, offers two options to accomplish this task: the graphical Rpmdrake application (see Section 5.1.4), and the command-line tools. The latter is comprised of urpmi to install and upgrade software, urpme to remove RPM packages, `urpmf` and `urpmq` to search the RPM databases. These are all the underlying programs behind Rpmdrake.

5.1.1. Local repository configuration

In order to facilitate the installation of software packages, it is possible to create a repository of packages directly on the hard drive of your server.

Prerequisites

- You must have a minimum of disk space 3 GB.
- You need the installation CD/DVD.

5.1.1.1. Configuring a local repository in graphical mode

Here's how to configure a local repository with graphical tools.

- Insert the installation CD/DVD.

- An icon representing the media appears on the desktop. Double-click on it to explore the contents of the DVD.
- Copy the i586 directory (or x86_64) on your hard drive to the desired location.
- Once the copy is complete, run the Mandriva Linux Control Center.

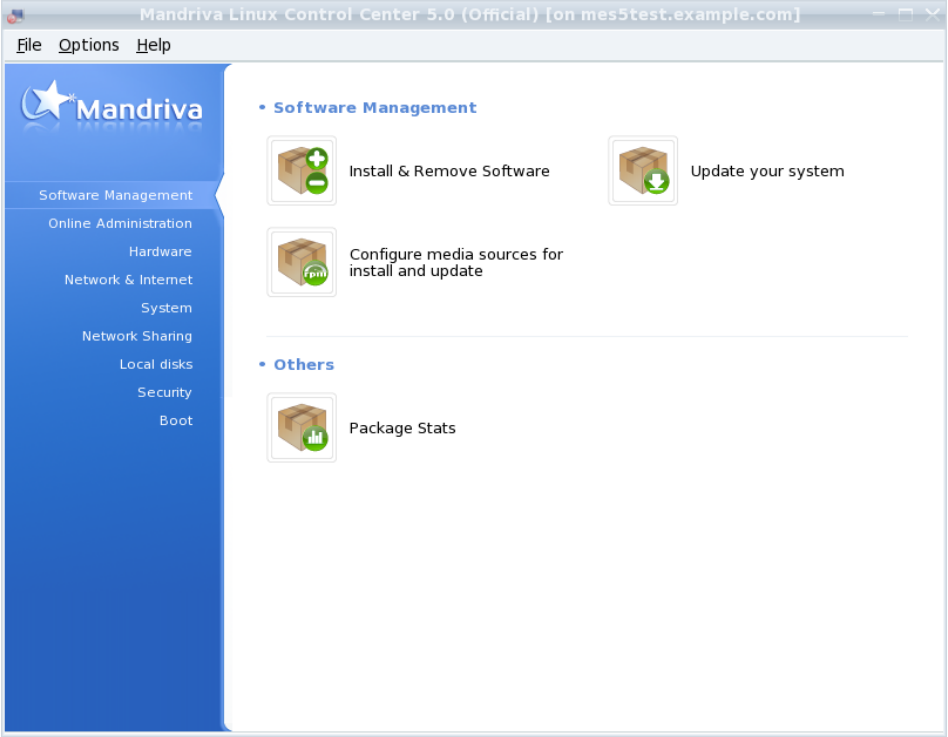

Figure 5-1. The Mandriva Linux Control Center

- Click on Configure media sources for install and update.
- In File, choose Add a custom medium.

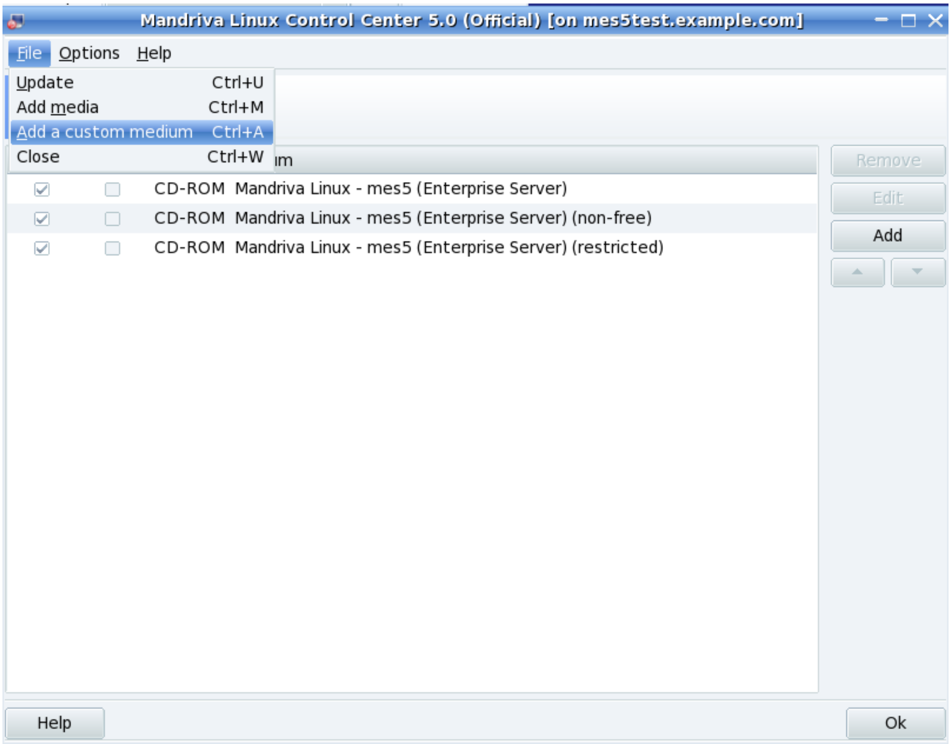
Mandriva Linux Control Center 5.0 (Official) [on mes5test.example.com]
File
Options
Help
Update Ctrl+U
Add media Ctrl+M
Add a custom medium Ctrl+A
Close Ctrl+W
Remove
Edit
Add
Help
Ok

☑ □ CD-ROM Mandriva Linux - mes5 (Enterprise Server)

☑ □ CD-ROM Mandriva Linux - mes5 (Enterprise Server) (non-free)

☑ □ CD-ROM Mandriva Linux - mes5 (Enterprise Server) (restricted)

Figure 5-2. Adding repository

- Fill the fields.

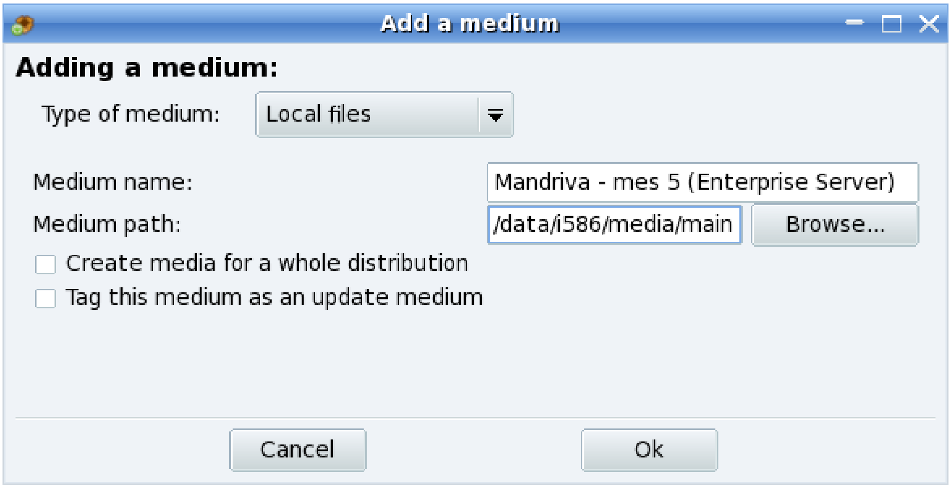

Adding a medium:

Type of medium: Local files

Medium name: Mandriva - mes 5 (Enterprise Server)

Medium path: /data/i586/media/main

☐ Create media for a whole distribution

☐ Tag this medium as an update medium

Figure 5-3. Adding custom repository

There are 3 repository to configure :

- Medium main

Medium name : Mandriva - mes5 (Enterprise Server)

Medium path : `/data/i586/media/main/`

- Medium non-free

Medium name : Mandriva - mes5 (Enterprise Server) non-free

Medium path : `/data/i586/media/non-free/`

- Medium restricted

Medium name : Mandriva - mes5 (Enterprise Server) restricted

Medium path : `/data/i586/media/restricted/`

In our example, the `i586` directory of the DVD has been copied to `/data`

- Disable CDROM media : uncheck Enabled.

5.1.1.2. Configuring a local repository in command line

Here's how to configure a local repository with command lines.

- Insert the installation CD/DVD.
- The CD/DVD is mounted in `/mnt/cdrom`
- Copy the `i586` (or `x86_64`) directory on your hard drive (for example in `/data`):

  ```
  cp -r /media/cdrom/i586 /data/
  ```

- Remove old media:

  ```
  urpmi.removemedia -a
  ```

- Add medias:

  ```
  urpmi.addmedia "Mandriva - mes5 (Entreprise Server) main" file://data/i586/media/
  urpmi.addmedia "Mandriva - mes5 (Entreprise Server) non-free" file://data/i586/me
  urpmi.addmedia "Mandriva - mes5 (Entreprise Server) restricted" file://data/i586/
  ```

- Update with the added media:

  ```
  urpmi.update -a
  ```

5.1.2. Installing and Updating RPMs with urpmi

5.1.2.1. Basic Notions

The main purpose of `urpmi` is to easily download and install RPM packages. RPM software packages often contain dependencies: urpmi recognizes those dependencies, downloads the required missing packages, and removes conflicting packages.

urpmi fetches the list of available RPMs and the RPMs themselves from a source media. Roughly speaking, a source media is described by a name and a location specified by an URL. Currently supported media types include local drives, removable drives (such as CDs), ISO images, and network media via different protocols (http, ftp, ssh and rsync). NFS mounted directories are treated like local drives.

5.1.2.2. Installing RPMs

Here's a description of the basic urpmi options:

```
urpmi <list of package names>
```

That prompts urpmi to fetch and install all packages and their unmet dependencies from the media you've configured. In the process, urpmi might ask questions. For example, if some packages need to be upgraded, or if some new (unspecified) packages have to be installed, that operation will require confirmation. If some packages need to be removed (due to conflicts with the requested packages), urpmi will also ask for confirmation. In some cases, urpmi will also propose a choice between different alternatives, usually proposing the "best" package as a default.

Another very useful urpmi feature is to upgrade all packages to the latest version found on the media. This is done with the following command:

```
urpmi --auto-update
```

urpmi can also install RPM files directly. Instead of using `rpm -i foobar.rpm`, you can pass the path of the RPM file to urpmi and it will then try to resolve the needed dependencies:

```
urpmi /home/user/foobar.rpm
```

Here are some useful options for urpmi :

--auto

Automatic mode: urpmi doesn't ask questions and always selects the default choice.

--test

To test the installation of packages without actually installing anything or modifying the system.

--media *media1,...,mediaN*

This option forces urpmi to only use the specified media, instead of defaulting to all available media. You can also specify a substring of media names for urpmi to select all media containing that substring. For example, `urpmi --auto-update --media updates` will search updates from all media that have "updates" in their name.

See the `urpmi(8)` manpage for more information.

5.1.2.3. Media Management with urpmi

5.1.2.3.1. Adding Media

urpmi is only usable when at least one media is defined. Usually, the OS installation procedure configures a predefined set of media, corresponding to the installation method you've selected: installation CDs, an HTTP or FTP server if you've installed from a networked mirror. System administrators often want to add media. To do so, use the `urpmi.addmedia` program:

```
urpmi.addmedia [options] <name> <url> [with hdlist]
```

In this synopsis, <name> is the name of the new media, <url> the URL where the RPMs are to be found, and the "with" parameter optionally specifies where to find the information file that describes the media's contents..

Supported URLs can be `http://`, `ftp://`, `rsync://`, `ssh://` (this will use `rsync` over `ssh`), `file://`, and `removable://`. `removable://` works like `file://`, but instructs urpmi that the directory is mounted from a removable media, such as a CD or a DVD. If the media requires authentication, you can use the usual URL syntax:

```
<scheme>://<login>:<pass>@host/path
```

Those credentials will not be stored in any world-readable file.

In some cases, if your media points to an external HTTP or FTP server, you might want to use a proxy to access it. Use the `--proxy` and `--proxy-user` options (the second one is needed if your proxy requires authentication.)

5.1.2.3.2. Removing media

This is straight forward. To remove a media named `foo`, simply use the command:

```
urpmi.removemedia foo
```

5.1.2.3.3. Updating Media

Some media never change. This is the case, for example, for CD-ROMs. However, other media, typically updates, grow. New RPMs are added and old ones are removed. Therefore, before using them you should instruct urpmi that their contents might have changed.

To do this, use the `urpmi.update` program. You can either update all media:

```
urpmi.update -a
```

You can also update only specifically named media:

```
urpmi.update updates-one updates-two
```

5.1.2.3.4. Creating Your own Media

The easiest way to create your own media is to let `urpmi.addmedia` do it. However, this will work properly only if you have a small number of RPMs, stored on disk or on a shared NFS mount. To do this, assuming that your RPMs are under a directory like `/var/my-rpms`, enter the command below:

```
urpmi.addmedia my-media /var/my-rpms
```

To create a media containing a large number of RPMs, or to be put on a shared server, you'll need to use the `gendistrib` tool. It comes in the rpmtools package. It generates a mirror tree for one or several media.

A typical media repository, under a root directory `/ROOT/`, has the following structure, (we have two media, named "first" and "second"):

```
ROOT/ - media/
|- first/
|    `- media_info/
|- second/
|    `- media_info/
`- media_info/
```

The RPMs are placed in the `first` and `second` subdirectories. Repository metadata is contained in the top-level `media_info` directory. Per-media metadata is contained in the `first/media_info` and `second/media_info` subdirectories.

Per-media metadata consists of an `hdlist.cz` file that contains the gzipped headers of the media's RPMs, a `synthesis.hdlist.cz` file (much smaller than the hdlist) containing only the information necessary for `urpmi` to resolve dependencies, and optionally, a `pubkey` file if the RPMs are signed (so that urpmi can check that the RPMs it downloads are signed with the key associated to this media.)

Before using `gendistrib`, you must create a `media_info/media.cfg` file to describe this media repository. The syntax of this file is like that of `.ini` files. It contains one section per media:

```
[first] hdlist=hdlist_first.cz name=First
    supplementary media
```

In the previous case, `first` is the directory name, `hdlist_first.cz` is the name of the `hdlist` file that will be created (it must end with `.cz`), and `name=` gives a human-readable description of the media.

Then, you can run gendistrib. You should pass it the /ROOT/ directory as parameter. It will then generate the **hdlist** and **synthesis** files and all other files needed for proper repository operation.

For further information, see the gendistrib(1) manpage.

5.1.2.4. Parallel Installation Command: urpmi-parallel

urpmi-parallel is a useful add-on to urpmi to install packages on many network hosts. It runs a urpmi command in parallel on a specified number of hosts. More precisely, the machine you run the command on (the server) tests the results on each machine in the group (the "clients") one by one, downloads all necessary packages for all machines in the group, distributes the appropriate packages to each machine, then calls urpmi on the machine to do the actual installation.

urpmi must be installed on all client machines, but it's not necessary to have a media defined on client machines.

To use it, follow these steps:

1. Make sure you can ssh from the server to each client machine as root (you can use ssh-add on the server host to avoid entering your passphrase and/or password many times).
2. Install urpmi-parallel-ssh or urpmi-parallel-ka-run on the server machine. The first plugin uses plain ssh to distribute commands to other hosts, while the second one uses ka-run, an efficient parallelization method on top of any remote shell (rsh or ssh), adapted to clusters.
3. Edit **/etc/urpmi/parallel.cfg** to look like this:

   ```
   mynetwork:ssh:host1:host2:host3
   ```

   On this line, mynetwork is the name of the alias you will use to specify the network to urpmi, ssh is the install method (to use ka-run, look up the entry for **/etc/urpmi/parallel.cfg** in urpmi.files(5)), and hostN are the host names of all clients on your network. You can put localhost in this list.
4. Run the urpmi command: for example, to install "package_name":

   ```
   urpmi --parallel mynetwork package_name
   ```

5.1.2.5. Restricted urpmi

urpmi has a more secure counterpart: rurpmi. It is similar to urpmi, but has a stripped-down set of features. It's intended to be used by users without root privileges, but with sudo rights on it, preventing any abuse of this tool to compromise the system.

Its syntax is similar to urpmi, but it prevents installing arbitrary RPMs. They need to come from a registered media. A number of dangerous options, listed in the rurpmi(8) manpage, are also forbidden.

5.1.2.6. Removing RPMs with urpme

The tool used to uninstall RPMs is urpme. For example:

```
urpme <list of package names>
```

This will attempt to remove all listed packages and the packages that depend on them. It will refuse to uninstall "important" packages (the ones that are part of the base system).

See the urpme(8) pmanpage for the reference of all options urpme supports.

Note that urpme isn't able to detect packages that are no longer used. For example, libraries not used by applications. To clean them up, a handy tool is rpm-find-leaves. It will list all RPMs on your system that no other package requires.

5.1.3. Searching for Packages with urpmf and urpmq

5.1.3.1. urpmf

urpmf is a grep-like tool for the urpmi database (the database of all RPMs in the media). By default, it will search through the file names contained in packages, but a variety of options allows more advanced searches.

For example, to find all packages that begin with apache- :

```
urpmf --name '^apache-'
```

The ^ is the beginning-of-line anchor used in standard regular expressions.

To find all packages that contain files whose path name includes /etc/httpd.conf.d :

```
urpmf /etc/httpd.conf.d
```

To find all packages that provide `mail-server` with their version and release number (**-f**) :

```
urpmf --provides -f mail-server
```

See the `urpmf(8)` manpage for more examples and the list of all options.

5.1.3.2. urpmq

The `urpmq` tool allows you to query the urpmi database. It has several modes of operation. Here are a couple of interesting uses:

```
urpmq -i package
```

It lists the information for that package (like `rpm -qi` would do for an installed package). The `--summary` option is similar, but only gives one-line, concise information.

```
urpmq --source package
```

This returns the URL where the package can be retrieved from.

```
urpmq -d package
```

It gives the list of all RPMs that are required by the specified package (recursively).

Inversely, the `urpmq -R package` returns the list of all RPMs that require the specified package.

See the `urpmq(8)` manpage for a list of all options.

5.1.4. Managing Software with Rpmdrake

The urpmi set of tools is command-line based. You can also use a graphical tool called Rpmdrake. It consists of different components which you access by choosing one of the entries of the System+Configuration+Packaging menu in the main menu, or by clicking on Software Management in the Mandriva Linux Control Center.

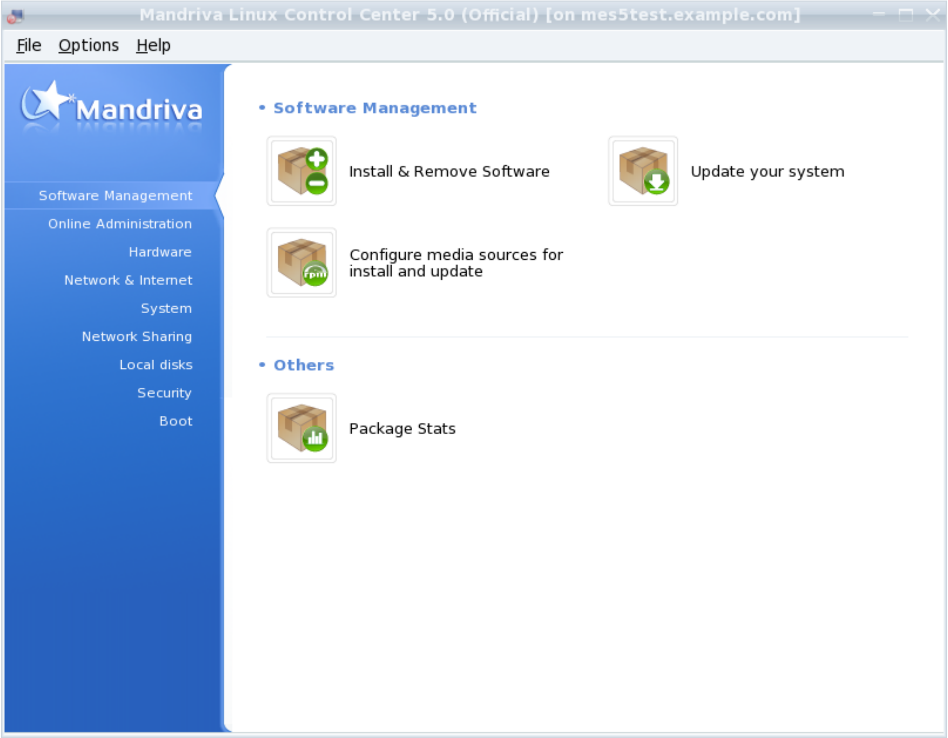

Figure 5-4. Software Management in the Mandriva Linux Control Center

We recommend that you access Rpmdrake via the Mandriva Linux Control Center.

5.1.4.1. Install Software

When launched, Rpmdrake searches the database of available packages. Then the Software Packages Installation interface is displayed.

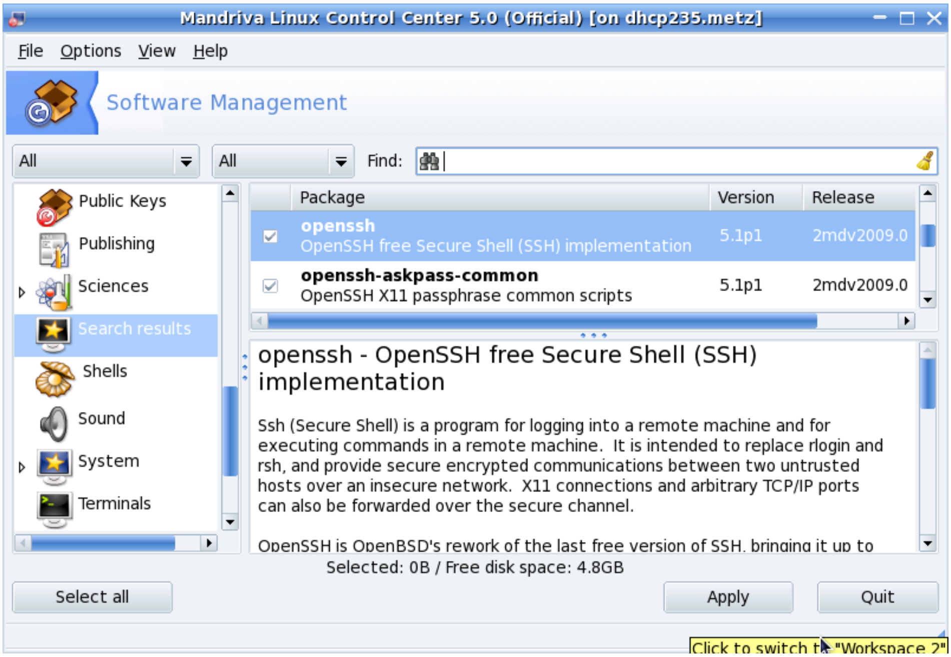

	Package	Version	Release
☑	**openssh** OpenSSH free Secure Shell (SSH) implementation	5.1p1	2mdv2009.0
☑	**openssh-askpass-common** OpenSSH X11 passphrase common scripts	5.1p1	2mdv2009.0

openssh - OpenSSH free Secure Shell (SSH) implementation

Ssh (Secure Shell) is a program for logging into a remote machine and for executing commands in a remote machine. It is intended to replace rlogin and rsh, and provide secure encrypted communications between two untrusted hosts over an insecure network. X11 connections and arbitrary TCP/IP ports can also be forwarded over the secure channel.

OpenSSH is OpenBSD's rework of the last free version of SSH, bringing it up to

Selected: 0B / Free disk space: 4.8GB

Figure 5-5. The Software Package Installation interface

The upper part of the window allows you to manipulate the list of packages you can install. This list is in the middle-left of the window. On its right is an area where the currently selected package is described. At the bottom of the window are two buttons and information about the space required to install the selected packages, as well as how much space is available on your disk.

Additionally, a status bar in the lower part of the window displays messages about actions currently in progress or completed.

5.1.4.1.1. Selecting Packages to Install

In Figure 5-5, the `samba-3.2.7-0.3mdv2009.0` package is selected in the tree-view. In the package-description are shown the required disk space, a short summary (“Samba (SMB) server programs”), as well as a detailed description.

If your software medium repository is configured to use complete package lists (not the summary `synthesis` files, but the full `hdlist` ones, which is the default behaviour after installing Mandriva Linux), you may get more information on the package by choosing the Maximum information radio button in the access area. In addition is a list of the files provided by the package and the change log.

The status bar displays the disk space required by the selected packages, as well as the current free space. Please note that due to dependencies, the disk space required by the selected packages might be greater than the size required by the chosen package itself.

Rpmdrake shows you an alert box if you try to install more software than the available disk space. Delete files you no longer use to proceed with the installation.

Now you can begin the installation by clicking on the Install button. A new window appears, showing you a progress bar of how your installation is proceeding. If you prefer to leave without doing anything, click the Quit button.

While selecting applications to install, it may happen that you choose a package which requires dependencies (additional libraries or another tool) to work correctly. In this case, Rpmdrake displays an information window enabling you to choose whether to accept the selected dependencies, or to Cancel the operation (Figure 5-6).

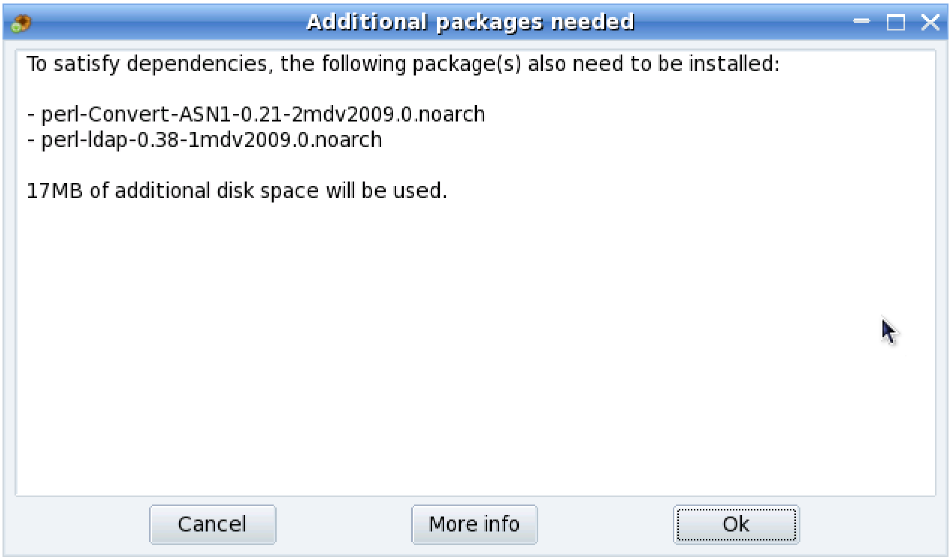

To satisfy dependencies, the following package(s) also need to be installed:

- perl-Convert-ASN1-0.21-2mdv2009.0.noarch
- perl-ldap-0.38-1mdv2009.0.noarch

17MB of additional disk space will be used.

Figure 5-6. Rpmdrake Dependency Alert Box

Let's say you want to install a package which requires dependencies, and various packages are capable of providing that dependency. The list of alternatives is then presented. You may read the additional information presented by clicking the Info button to help you choose the best alternative.

We will now take a closer look at the search and sort functions provided to ease your job as a system administrator:

5.1.4.1.2. Searching for Packages

To search for a specific package, just type its name (or part of the name) in the text area next to the Search button. Then choose from the pull-down list where you want to look for it (either in the package name, in the description provided with the package or in the names of the files stored in the packages). After clicking on the Search button, a new list appears (Search results), showing you the results Rpmdrake found while scanning the package databases.

Let's look at the different sort orders:

Mandriva Linux choices

> Shows the list of packages in the four groups you saw during the installation process. This is the easier sort order because it focuses on a selected part of the available packages: those considered to be the most useful.

All packages, alphabetical

Displays a flat list of all available packages you can install.

packages, by group

Shows the list of packages grouped by function (e.g. games, system, video, etc.).

All packages, by size

Displays a package list sorted by size.

All packages, by selection state

Shows a flat list in which all selected packages are displayed first. The other available packages are listed below.

All packages, by medium repository

Displays packages alphabetically and under the name of the data medium they belong to (see Section 5.1.4.4).

All packages, by update availability

Might show two groups of packages: one for packages you can install, and a second list of packages for which an older version is already installed.

5.1.4.2. Remove Software

This interface works like the "Install Software" one, so we won't repeat its basic functions (Section 5.1.4.1).

5.1.4.3. Mandriva Linux Update

When you launch Mandriva Linux Update, it first asks you to choose an Internet repository to check for updates. You should choose one in a country near you.

A small difference to the "Install Software" interface is the ability to choose the kind of update you want to install by grouping them in certain ways. You may select:

Security updates

These solve security issues and should be installed as soon as possible.

Bugfixes updates

These ones fix application misbehavior.

Normal updates

They just bring slight improvements.

The other difference is a new text section ("Reason for update") inside the package description area. It tells you why this update was made available. This may help you decide if you want to update certain packages or not. When you have a slow Internet connection or you have to pay per MB when you are downloading, it would be wise to read it.

5.1.4.4. The Software Media Manager

This part of Rpmdrake is dedicated to the configuration of the package media repositories. As you can see in Figure 5-7, some media are configured: "Main", "Contrib", etc. With this tool you can add other software media: a CD from a magazine containing RPMs, a Web repository, etc.

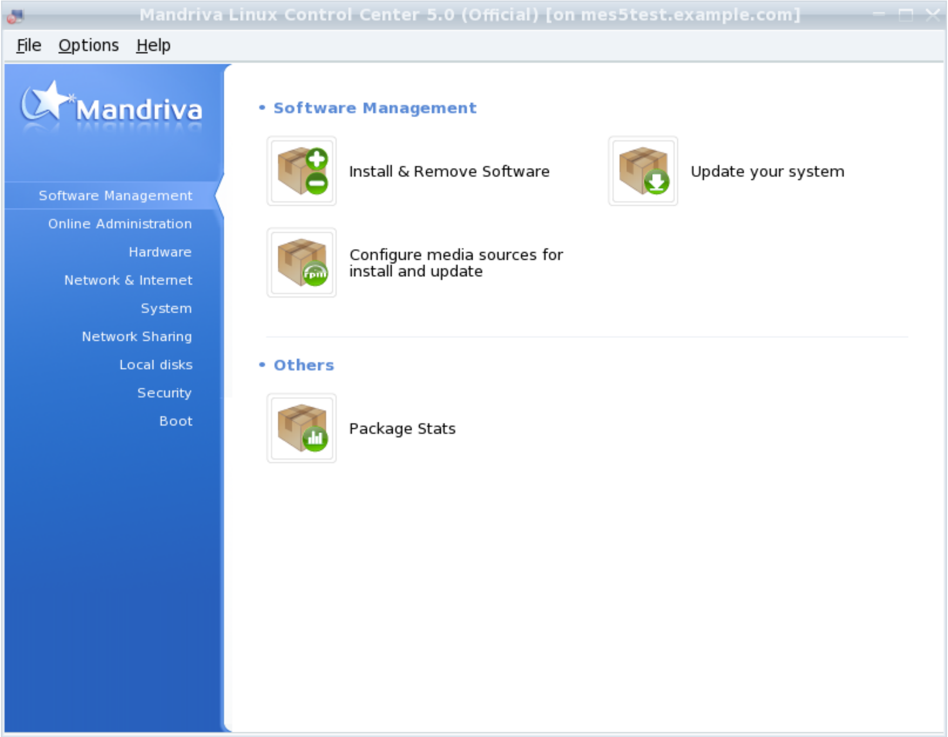

Figure 5-7. The Software Media Manager

The checkboxes in the left-hand column allow you to flag the repositories:

Enabled?

> Uncheck this box to temporarily disable the corresponding medium. The packages contained in this medium won't be available for installation until you enable the medium again.

Updates?

> This box must be checked for update media, that is, media that contain updates of packages that are already installed on your machine, but with an older version number. Therefore only update media are taken into account when looking for updates.

Various buttons allow you to perform actions on the selected media.

Remove

> Allows you to remove a medium which you no longer use. Simply select the medium to be removed in the list and click on Remove.

Edit

Enables you to change the URL or the relative path to the `synthesis/hdlist` file.

Click the Proxy button to configure a proxy server. Note that you can also define a global proxy for all remote media through the main interface's Proxy button.

This option also allows you to change from using `hdlist` files to `synthesis` files.

> Synthesized files only give information about package names, their dependencies and a short summary. You won't be able to search for files inside uninstalled packages, for example, and you won't be able to see the full description for a package if you click on its name.

Add

Permits you to add all publicly available official package sources from Internet repositories. This is useful if you have a fast Internet connection or you only have the first installation CD at hand. Choosing a mirror geographically near to your location is recommended.

After choosing a mirror and clicking OK, package information for the source you chose is downloaded and all included packages will be available for you to install and update your system.

Add custom

Enables you to add a custom software repository source.

Update

Shows a list of already defined media. With it, you can choose the ones from which you want to update the list of available packages. This is useful for remote media to which new packages are being added. Just start the process by clicking on Update.

Manage Keys

It is important that any new package you install is authenticated. To do so, each package can be electronically signed with a "key", and you can allow/disallow keys on a per-medium basis. In Figure 5-8, Mandriva Linux key is allowed for medium Installation CD. Click on Add a key to allow another key for the selected medium, and on Remove key to remove a key from the selected medium.

⚠ As with all security-related questions, make sure you know what you're doing if you decide to remove keys.

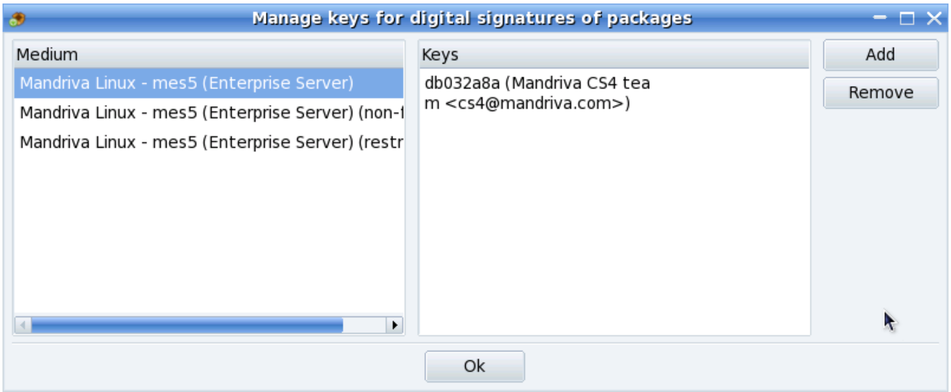

Mandriva Linux - mes5 (Enterprise Server)
Mandriva Linux - mes5 (Enterprise Server) (non-f
Mandriva Linux - mes5 (Enterprise Server) (restr

db032a8a (Mandriva CS4 team <cs4@mandriva.com>)

Figure 5-8. Managing Keys

Proxy

If you are sitting behind a firewall and you still need to access remote media (especially for package updates), you can do so if you have a proxy server which leads to the Internet (at least in an area where you can find a package server). Normally, it should be enough to fill in the Proxy hostname to get it working (Figure 5-9). If you need a user/password combination to get through the proxy, you can also specify it here. Just confirm your changes by clicking OK and you're done.

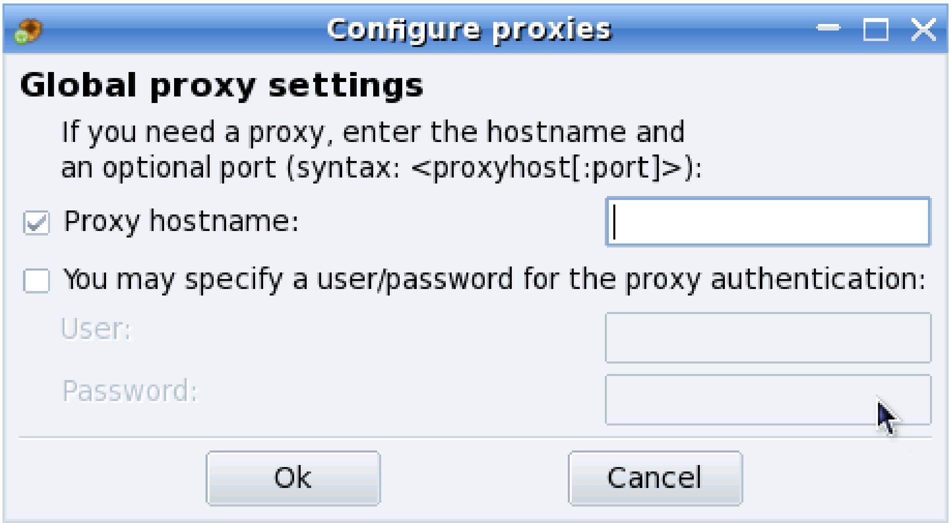

Global proxy settings

If you need a proxy, enter the hostname and an optional port (syntax: <proxyhost[:port]>):

☑ Proxy hostname:

□ You may specify a user/password for the proxy authentication:

User:

Password:

Figure 5-9. Configuring a Proxy server

Parallel

If you run a large network of computers, you may want to install a package on all the computers in parallel. This button opens a dialog window allowing the configuration of the “Parallel” mode. See the `urpmi(8)` man-page for more information.

Global Options

Allows you to configure the program used to download new packages and if the source should be checked against a given key. These choices are used on all installed sources.

Up and Down Arrows

Enable you to change the order in which sources are used when installing packages.

For advanced users

Rpmdrake processes the urpmi configuration file (`/etc/urpmi/urpmi.cfg`) from top to bottom to obtain a list of medium repositories and the packages that each contains.

If a given package appears in more than one medium, and versions differ, then the one with the newest version is used, ignoring the rest.

When a package with the same version appears in more than one medium, only the one appearing first will be used, the rest are ignored.

Either way, you won't miss available packages.

Rpmdrake processes the urpmi configuration file (`/etc/urpmi/urpmi.cfg`) from top to bottom to obtain a list of medium repositories and the packages that each contains.

If a given package appears in more than one medium, and versions differ, then the one with the newest version is used, ignoring the rest.

When a package with the same version appears in more than one medium, only the one appearing first will be used, the rest are ignored.

Either way, you won't miss available packages.

5.1.4.5. Managing a Computer Group

5.1.4.5.1. Defining the Group

By using Rpmdrake's parallel mode, you can define groups of machines to install the same software packages on. This greatly simplifies managing a large number of systems, such as in a LAN. Make sure `park-rpmdrake`, `urpmi-parallel-ssh` and `urpmi-parallel-ka-run` are installed.

This tool is accessible only in expert mode. Choose Options→ Expert mode from the menu and then access the Software Management section of Mandriva Linux Control Center.

Click on New Group, provide a Name for the group, select the Protocol to use (ssh in our example), check the networks to be scanned for hosts in the list or add other networks. By default, only your local network is listed. Then click on Scan.

Hosts on the network must have an `ssh` server running and the corresponding ports (`tcp/22` by default) open if they also host a firewall. The **`rsync`** package must be installed as well. Also, `root` must be allowed to login using `ssh` (`PermitRootLogin yes` in **`/etc/ssh/sshd_config`**, on the hosts).

Wait a few minutes until the network scan finishes and then check the hosts to be made part of the group.

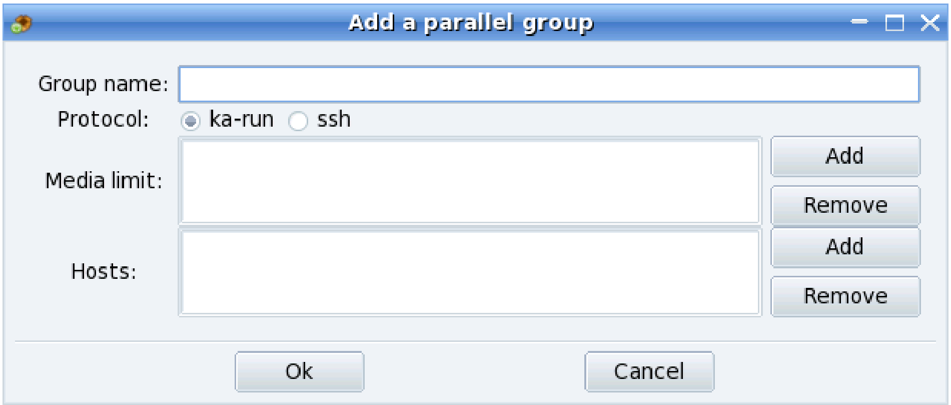

Add a parallel group

Group name:

Protocol: ● ka-run ○ ssh

Media limit:

Hosts:

Figure 5-10. Adding Computers to a Group

You are then asked for `root`'s password on each of the selected hosts to automate the installation of packages on them.

5.1.4.5.2. Managing Packages on a Group of Computers

For package installation or removal in a specific group to succeed, all computers in that particular group must be available.

Select the group and click the Use Group button. You can now install packages on the selected computers group, as you would on a single computer.

You can also install packages on the computers' group from the command line:

```
urpmi --parallel <group_name> <package_name>
```

For example, issuing `urpmi --parallel GraphicDesign gimp` will install GIMP on all computers which are part of the `GraphicDesign` group.

To remove a package from a computer group, use:

```
urpme --parallel <group_name> <package_name>
```

For example, issuing `urpme --parallel GraphicDesign gcc` will remove the C compiler from all computers which are part of the `GraphicDesign` group.

5.2. Load Balancing

This chapter presents solutions to implement load balancing and service distribution to create a high availability service environment. You'll learn about install methods, specific service requirements, as well as how to configure load balancing correctly.

5.2.1. Prerequisites

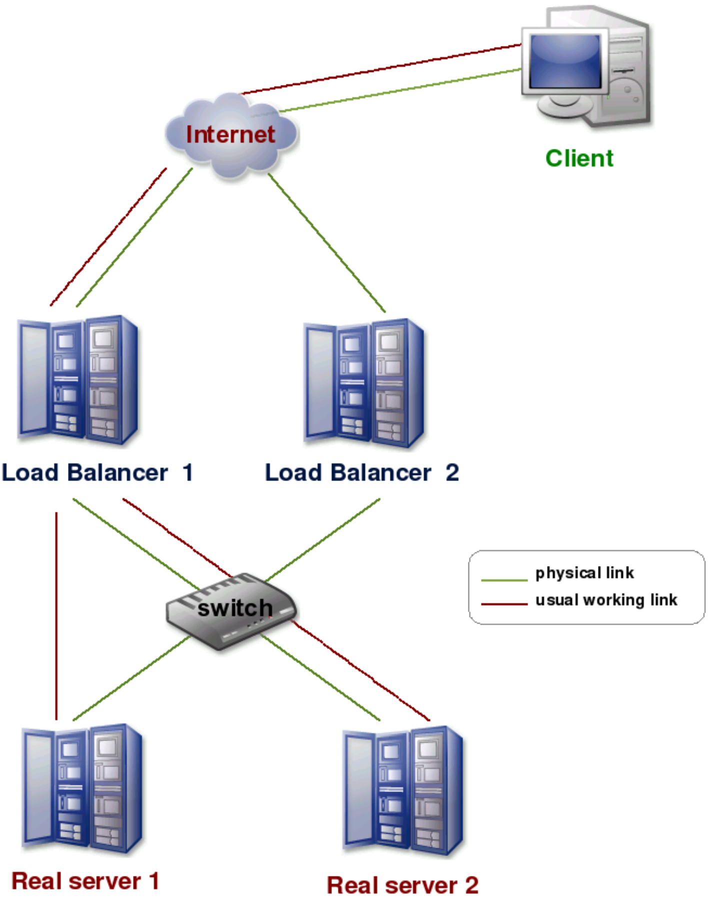

Figure 5-11. General architecture for load balancing

We will look at a simple case: an administrator wants to establish transparent redundancy for his Intranet web server to guarantee service availability. In this case, we have 4 servers allocated as described below:

- Two loadbalancers (identified by `lb1` and `lb2`),

- Two real servers (also called nodes `rs1` and `rs2`) which run the desired service.

Loadbalancers allocate the "load" according to predefined parameters. The real servers are the physical machines running the service. In our case, it's a web server. If the load increases, you can transparently add more hardware.

Here are the IP addresses we will be using:

```
10.0.0.3 - adresse IP de lb1
   10.0.0.4 - adresse IP de lb2
   10.0.0.5 - adresse IP du service Web de l'Intranet

192.168.0.1 - adresse IP local de la passerelle
   192.168.0.12 - adresse IP de rs1
   192.168.0.13 - adresse IP de rs2
```

5.2.2. General Concepts and Web References

Load balancing happens at IP level. It is handled directly by the Linux kernel. The loadbalancer acts as a "router" since it redirects requests (therefore packets) towards a target according to different rules.

Main web references:

- Official site of Linux Virtual Server (LVS) (`http://www.linuxvirtualserver.org/`)
- Most complete HOWTO for LVS. In practice, it replaces the official documentation (`http://www.austintek.com/LVS/LVS-HOWTO/`)
- Keepalived official site (`http://www.keepalived.org/`)
- Linux Advanced Routing HOWTO section on iproute2 (`http://lartc.org/howto/lartc.iproute2.html`)

5.2.3. Installing and Configuring LVS

5.2.3.1. Necessary RPM Packages

On the loadbalancers:

ipvsadm

Contains the executable controlling the ip_vs kernel module that routes the load between servers.

keepalived

Contains keepalived which will be used to controlipvsadm. Keepalived also includes a vrrp daemon to enable service availability.

On the web servers:

iproute2

Enables the advanced ip routing features of the Linux kernel.

5.2.3.2. Installation

1. On the loadbalancers

   As root, install the keepalived RPM which will ask to you to satisfy the required dependencies:

   ```
   [root@lb1 ~]#urpmi keepalived
   To satisfied dependencies, the following packages will be
   installed:
   ipvsadm-1.24-6mdv2009.0
   keepalived-1.1.17-1mdvmes5
   Install these required packages? (0 Mo) (Y/n)
   ```

   You **must** install these packages on both (all) loadbalancers.

2. On the web servers

   Normally, the `iproute2` package should already be installed. If it's not the case, install it:

   ```
   [root@lb1 ~]#urpmi iproute2
   ```

General configuration is only within `/etc/keepalived/keepalived.conf`. All the following configuration elements should be placed in this file (by default, it contains an example configuration).

5.2.3.3. Configuring the Loadbalancers

First, the module `ip_vs` needs to be loaded at boot. Add it to `/etc/modprobe.preload`. Normally, keepalived will load this module automatically.

```
#modprobe ip_vs
```

Keepalived is be used to manage ipvsadm as well as the fail-over between the loadbalancers. We have established redundancy between the loadbalancers for security reasons, to make sure we don't have to worry about this. Keepalived also installs ipvsadm, the ip_vs administration tool.

Let's make sure that it implements ip_forward, because it will not function without it:

```
echo "1" > /proc/sys/net/ipv4/ip_forward
```

Edit /etc/keepalived/keepalived.conf. The global_defs section applies to the general configuration of Keepalived. It allows you to define notification options, and mainly the lvs_id (unique machine identifier).

```
global_defs {
    notification_email {
    mon@mail.com
    }
    notification_email_from keepalived@domaine.tld
    smtp_server 127.0.0.1
    smtp_connect_timeout 30
    lvs_id LVS_MAIN
    }
```

The vrrp_instance section defines use parameters of the VRRP daemon included in Keepalived. This protocol handles dynamic IP addressing between many systems. This allows redundancy of the loadbalancers. If one becomes unavailable, the other picks up the load and grabs the IP address specified in virtual_ipaddress{}. For the second loadbalancer, the state variable contains the SLAVE value. virtual_ipaddress defines all the IP address that the loadbalancers should attribute itself when it is the master. Be careful, virtual_ipaddress only handles one IP per line with a maximum of 32 addresses.

```
vrrp_instance VI_1 {
    state MASTER
    interface eth0
    virtual_router_id 51
    priority 100
    advert_int 1
    authentication {
    auth_type PASS
    auth_pass 1111
    }
    virtual_ipaddress {
    10.0.0.5
    }
    }
```

The following section discusses defining virtual servers for redundancy, as well as the protocols to implement it.

```
# Virtual Servers definitions

    virtual_server 192.168.0.1 80 {
    delay_loop 30
    lb_algo wlc
    lb_kind NAT
    persistence_timeout 50
    protocol TCP
```

If all the real servers are unavailable, keepalived will return the request. It allows you to show a static information page in case of a problem.

```
sorry_server 192.168.0.1 80
```

Add all the real servers:

```
real_server 192.168.0.12 80 {
    weight 20
    }

    real_server 192.168.0.13 80 {
    weight 8
    }

    }
```

5.2.3.4. Configuring the Real Servers

Since the machines are within a local network with the loadbalancers, we must force them to use the loadbalancers as gateways for all traffic. Let's create a static route to do this with `iproute2`.

Let's create the following rules:

```
#!/bin/sh -e
    echo 200 LVS >> /etc/iproute2/rt_tables
    /sbin/ip rule add from 192.168.0.12 table LVS
    /sbin/ip route add default via 192.168.0.1 dev eth0 table LVS
    /sbin/ip route flush cache
```

Let's test the configuration. First, restart keepalived:

```
# service keepalived restart
```

Use your favorite Web browser to access http://10.0.0.5 (`http://10.0.0.5`). If the page loads correctly, your configuration is working. Of course, this is a basic configuration. It may be interesting to add more advanced features such as firewalling, monitoring, etc.

Once the server is installed, you can use `ipvsadm -L` to view the state of your load balancing:

```
# ipvsadm
    IP Virtual Server version 1.2.0 (size=4096)
    Prot LocalAddress:Port Scheduler Flags
    -> RemoteAddress:Port           Forward Weight ActiveConn
    InActConn TCP  10.0.0.5:www wlc persistent 20
    -> 10.0.0.3:www                  Route   25     294        916
    -> 10.0.0.4:www                 Route   8      243        754
```

5.2.3.5. Advanced Configuration

Keepalived includes a monitoring system capable of checking if the servers in your high availability cluster are reachable. You only need to add a section in the Web server's configuration in `/etc/keepalived/keepalived.conf`. The monitoring can be a simple ping:

```
TCP_CHECK { connect_port 80 connect_timeout 3
    }
```

It can go as far as checking that the returned page is the page it should be.

```
HTTP_GET {
    url {
    path /
    digest 69c520033af926d40e7d5d832e3933f6
    }
    connect_timeout 5
    nb_get_retry 3
    delay_before_retry 2
    }
```

To manage the hash, use `genhash` with the following command:

```
[root@lb1 ~]# genhash -s 193.188.255.4 -p 80 -u /
    MD5SUM = 5b5af20a92f1dbe92328005329d1753f
```

You can also launch actions if a specific event happens (typically, fall over to the other loadbalancer if the master becomes unreachable), with the following set up in section `vrrp` :

```
notify_master /path_to_script/script_master.sh
    (or notify_master "/path_to_script/script_master.sh <arg_list>")
    notify_backup /path_to_script/script_backup.sh
    (or notify_backup "/path_to_script/script_backup.sh <arg_list>")
    notify_fault /path_to_script/script_fault.sh
    (or notify_fault "/path_to_script/script_fault.sh <arg_list>")
```

5.3. Virtualization

This chapter discusses the concepts and procedures to implement virtual machines on your Mandriva Enterprise Server 5 using Xen, KVM or graphical tool Virt-manager. This will allow you to run multiple OSes within your Mandriva Linux server.

5.3.1. Xen

5.3.1.1. General Concepts and Main Web References

Xen is a virtual machine monitor for x86 hardware (it runs on i686 and x86-64 CPU classes). It supports running multiple guest operating systems on a single machine. Guest OSes (also called “domains”) require a modified kernel which supports Xen hypercalls to replace physical hardware access. At boot, the Xen kernel is loaded (via GRUB) along with the guest kernel for the first domain (called “domain0”). `domain0` has privileges to access the physical hardware (PCI and ISA devices), to manage other domains, as well as to provide virtual devices (disks and network) for other domains.

Here are the main Web references which discuss the configuration of Xen :

- Official site (`http://www.cl.cam.ac.uk/Research/SRG/netos/xen/`)
- Xen official vendor (`http://citrix.com/`)
- Xen Wiki (`http://wiki.xensource.com/`)

5.3.1.2. Installing and Configuring Domains

5.3.1.2.1. Before Starting

We installed Mandriva Mandriva Enterprise Server 5 on a 12 GB disk. During the installation, we created a partition of 5 GB. Its file system is mounted on / for the `domain0` installation. Additionally, we created a swap partition of 256 MB on `/dev/sda5`.

We now have 7 GB left. We created two partitions of 3.2 GB, that are not yet mounted, for two guest domains os installations, and two swap spaces of 256 MB each.

Here is the partition table :

```
root domain0 -> sda1
```

```
swap domain0 -> sda5

root guest1 -> sda6
swap guest1 -> sda7

root guest2 -> sda8
swap guest2 -> sda9
```

5.3.1.2.2. Installing RPMS

Mandriva Enterprise Server 5 provides the required patches in the kernel so that the Xen functionality is implemented. You will also need both specific kernels and tools to manage virtualization.

```
# urpmi kernel-xen xen
```

5.3.1.2.3. Configuring domain0

Installation of `xen` package provides a bootloader configuration and an initrd creation. What follows is a reminder to configure `domain0`. You have to set up the automatic start of xen services at boot time..

1. Configuring GRUB

   An entry has been added to the GRUB configuration file (`/boot/grub/menu.lst`) so that Xen/Xenolinux can boot. The entry should look like :

   ```
   title XEN 3 / Mandriva Enterprise Server 5
   kernel (hd0,0)/boot/xen.gz dom0_mem=131072
   module (hd0,0)/boot/vmlinuz-2.6.18.8-xen-3.3.0-7mdv root=/dev/sda1 ro
   module (hd0,0)/boot/initrd-2.6.18-8-xen-3.3.0-7mdv.img
   ```

   Review of main parameters :

   kernel

   Tells GRUB where Xen itself is located and parameters to be taken into account by the kernel (in this case, setting domain0 size of allocation memory in kilobytes).

   module

   The first line tells you where to find the Xenolinux kernel that Xen should launch with parameters. These are standard Linux parameters, root device identification and initial read-only mount.

module

The second line contains the path of initrd. It must be the module and not initrd in GRUB configuration, otherwise Xen will not boot.

2. Creating initrd

   We are now going to create `initrd` so that Xen guest domains will be managed at boot :

   ```
   # mkinitrd -v -f /boot/initrd-2.6.18.8-xen-3.3.0-7mdv.img
   2.6.18.8-xen-3.3.0-7mdv
   ```

   When adding the new kernel in menu.lst, it's recommended that you keep your existing entries. You could have to reboot on your old Linux kernel if you have problems.

3. Starting Xen services

   At the end of the installation and configuration process, reboot your system and choose the Xen entry in the GRUB menu.

   During your system's start-up, the first part of the output displays information for Xen itself (low level and hardware). The last part of the output displays information for Xenolinux.

   To create additional domains, start the `xend` daemon control. You can also start the `xendomains` daemon, which launches additional domains at domain0's boot.

   ```
   # chkconfig --add xend
   # chkconfig --add xendomains
   # service xend start
   # service xendomains start
   ```

At this point, you can use the `xm` tool to monitor or maintain domains running on your system.

5.3.1.2.4. Installation of Guest Domains

The first step in creating an additional domain is to **prepare a root filesystem** for it to boot. Typically, this might be stored in a normal partition, an LVM partition, a disk file or on an NFS server. A simple solution to do this is to boot from a system installation disk and install the distribution onto a new partition of your hard drive.

In the following sections, we will give an other option to create an additional domain based on a copy of domain0's root partition, installed with `urpmi` in a physical partition or in a file.

```
# mkdir -p /mnt/xen
```

Partition Copy. The main advantage of this method is that the installation disk is not needed. We copy the root partition of domain0 on another partition of the hard drive to obtain a second filesystem that will be used by the guest domain.

```
# mount /dev/sda6 /mnt/xen
# rsync -avDx / /mnt/xen
# cp -ar /dev/* /mnt/xen/dev/
# umount /mnt/xen
```

The `rsync` command will synchronize only root filesystems, it doesn't follow a mount link to other filesystems.

If your `/usr` and `/var` are separate, redo it for each mount point, except for the network and special entries.

If you have different partition schemes between domain0 and additional domains, don't forget to adapt the `/etc/fstab` file to the guest domain's filesystem.

Install Mandriva Enterprise Server 5 with urpmi. With `urpmi`, you can get a clean installation of a new Mandriva Enterprise Server 5.

1. First, mount your clean partition :

   ```
   # mount /dev/sda6 /mnt/xen
   ```

2. Set up your `urpmi` media :

   ```
   # urpmi.addmedia --distrib --urpmi-root /mnt/xen url
   ```

You must configure `url`to point to the directory containing the installation media.

3. Install the guest domain Mandriva Enterprise Server 5 system base and use `urpmi` to install new packages.

   ```
   # urpmi --urpmi-root /mnt/xen basesystem urpmi
   ```

   At the prompt, you will be asked many questions. Answer them as you would proceed for a classical installation (choose`kernel-xen`).

4. The following steps configure the new features :

   Copy configuration files for `domain0`.

   ```
   # cp /etc/resolv.conf /mnt/xen/etc
   # cp /etc/fstab /mnt/xen/etc
   ```

5. You have to configure guest *domain* `/etc/fstab` file. For example, the partition `/dev/sda6` becoming `/dev/sda1` :

   ```
   /dev/sda1 / ext3 relatime,user_xattr,acl 1 1
   /dev/sda5 swap swap defaults 0 0
   none /proc proc defaults 0 0
   ```

   a. Copy the networking configuration from `domain0`.

   ```
   # cp /etc/sysconfig/network-scripts/ifcfg-eth0 \
   /mnt/xen/etc/sysconfig/network-scripts
   ```

   b. Install dhcp client if necessary.

   ```
   # urpmi --urpmi-root /mnt/xen dhcp-client
   ```

   c. Enter chroot for final configuration.

   ```
   # chroot /mnt/xen
   ```

   d. Create empty necessary configuration files.

   ```
   # touch /etc/mtab /etc/urpmi/urpmi.cfg /var/lib/urpmi/MD5SUM
   ```

   e. Create `shadow` from `passwd`.

   ```
   # pwconv
   ```

   f. Turn on network by creating the `/etc/sysconfig/network` file with the following content :

   ```
   NETWORKING=yes
   ```

   g. Set root password.

   ```
   # passwd root
   ```

   h. Exit chroot's shell.

   ```
   # exit
   ```

   i. Umount guest partition.

   ```
   # umount /mnt/xen
   ```

Installing Mandriva Enterprise Server 5 in a File. This way of working has the advantage of not modifying the partition structure of your hard drive. We will explain how to create and mount a disk image. Then you can choose one of two methods below (urpmi or making a copy of the root partition).

1. First, create an image file for your Mandriva Enterprise Server 5 guest domain. This command creates a disk image of 1GB, filled with zeroes, in a file named `mandriva.img` located in the current directory.

   ```
   # dd if=/dev/zero of=mandriva.img bs=1M count=1 seek=1024
   ```

2. Create a filesystem in the image file. The

   ```
   -j
   ```

   option specifies an `ext3` filesystem.

   ```
   # mke2fs -F -j mandriva.img
   ```

3. Finally, you can mount your image file on a mount point.

   ```
   # mount -o loop mandriva.img /mnt/xen
   ```

You can now use this mount point like a physical partition's mount point and install Mandriva Enterprise Server 5.

In the guest domain's configuration file, the disk options should be modified so that `phy:sdaX` is replaced by `file:path/of/file`. Otherwise, other options are similar.

5.3.1.2.5. Configuration of Additional Domains

Before you can launch an additional domain, you should create a configuration file for guest domains. The following section describes these steps. This file is not essential but will help you to get guest domains to boot without any problems.

Create a Guest Domain Configuration File. The following configuration file (call it "mandriva") should be located in the `/etc/xen` directory in `domain0` if you want to launch it manually, or in the `/etc/xen/auto` directory if you want to launch it automatically by xendomains at boot.

```
kernel = "/boot/vmlinuz-2.6.18.8-xen-3.3.0-7mdv"
ramdisk = "/boot/initrd-xen-2.6.18.8-xen-3.3.0-7mdv.img"
memory = 128
name = "Mandriva"
dhcp = "dhcp"
disk = [ 'phy:sda6,sda1,w', 'phy:sda7,sda5,w' ]
root = "/dev/sda1 ro"
extra = "xencons=tty"
hostname = "mandriva2009"
vif = [ '' ]
```

This file shows the most common options used for a domain definition. Here's a short description of each option :

kernel

Establishes the link to the kernel compiled for Xen.

memory

Set it to the size of the guest domain's memory in megabytes.

name

Name of the additional domain.

dhcp

Uncomment the `dhcp` variable so that the domain will receive its IP address from a DHCP server.

disk

List of block devices exported to the guest domain. In this example, the `sda6` physical partition is named `sda1` in the additional domain and will be the root partition with the following `root` option. The `sda7` partition is a swap. If your disk is an image file, you should replace `phy:sdaX` by `file:path/of/file`. The `w` option explains the read-write rights to access this partition. You could also use the `r` option to set the partition as read-only.

root

Specify the root device parameter on the kernel command line. We must take the partition scheme of the disk option.

extra

Extra string to append to the kernel command line.

hostname

Guest domain hostname.

vif

The configuration of the network interface of the guest.

All the configuration modifications detailed in the next sections must be made in the guest domain, **not** in `domain0` as before. There are two possibilities :

- Launch the additional domain with the `xm create -c /etc/xen/auto/mandriva` command. Then modify the guest domain directly in it.
- Type a chroot command such as :

```
chroot /mnt/xen
```

To exit the chroot shell, type exit and don't forget to unmount the /mnt/xen filesystem (umount /mnt/xen).

Réseautique

Let's use the last method.

1. **Modifying /etc/inittab.** If you don't want to receive annoying messages such as :

   ```
   INIT: Id "2" respawning too fast: disabled for 5 minutes
   INIT: Id "3" respawning too fast: disabled for 5 minutes
   INIT: Id "4" respawning too fast: disabled for 5 minutes
   INIT: Id "5" respawning too fast: disabled for 5 minutes
   INIT: Id "6" respawning too fast: disabled for 5 minutes
   ```

   In your guest console, you should comment unused terminals in the /etc/inittab file like this :

   ```
   1:2345:respawn:/sbin/mingetty tty1
   #2:2345:respawn:/sbin/mingetty tty2
   #3:2345:respawn:/sbin/mingetty tty3
   #4:2345:respawn:/sbin/mingetty tty4
   #5:2345:respawn:/sbin/mingetty tty5
   #6:2345:respawn:/sbin/mingetty tty6
   ```

2. **xend and xendomains Services.** If you installed your system with a domain root filesystem copy, then xend and xendomains services are probably configured to start at boot time. It's useless for an additional domain. We will disable it as follows :

   ```
   # /etc/init.d/xend stop
   # /etc/init.d/xendomains stop
   # chkconfig xend off
   # chkconfig xendomains off
   ```

3. **Keytable Service.** Here is a little hint to avoid failure messages when keymap is loading at guest domain boot. You should disable the keytable service with :

   ```
   # chkconfig keytable off
   # /etc/init.d/keytable stop
   ```

   But it's not enough because this service is called directly in the /etc/rc.d/rc.sysinit file. We will comment out the following three lines (905,906,907) :

   ```
   #if [ -x /etc/init.d/keytable -a -d /usr/lib/kbd/keymaps ]; then
   #    /etc/init.d/keytable start
   #fi
   ```

   Notice that the keytable and numlock (see following section) services are not necessary in additional domains because domain0 has already initialized numlock and all domains use the same keyboard.

4. **Numlock Service.** Finally, in the configuration file of the guest domain, we get this message when the numlock service is started or when we log in :

   ```
   KDGETLED: Argument invalide Error reading current led setting.
   Maybe stdin is not a VT?
   ```

   To avoid this problem, disable the following :

   ```
   # /etc/init.d/numlock stop
   # chkconfig numlock off
   ```

   At this step, we have a Xen system which contains the `domain0` and one or more clean additional domains ready for execution.

5.3.1.3. Managing Guest Domains

5.3.1.3.1. Configuration réseau

This section explains Xen network basics and configuration.

Networking

The network scheme is very simple. `domain0` has real Ethernet interface control. Each guest domain has a virtual network interface. In additional domains, this interface is a normal Ethernet interface but it's named `vifX.Y` with "X" is the number interface (0 for eth0) and Y is the number of its own instance domain. All vif and eth interfaces are connected on the `peth0` bridge for network access.

See XenNetworking (`http://wiki.xensource.com/xenwiki/XenNetworking`) for detailed information on networking with Xen.

You will find a documentation and explanations on les interfaces Ethernet interfaces, MAC addresses, bridging, routing, interface names, VLANs, etc.

5.3.1.3.2. The xm Tool

The xm tool provides many options for managing domains. Here's a quick review :

- Start guest *domains* : Before you start a domain, you should create a config file for it.

  To start an additional domain :

```
# xm create -c /etc/xen/auto/mandriva
```

The `create` command starts a new domain instance. The `-c` option configures the console domain so that it should output immediately and `mandriva` is the name of the additional domain you start. You should now see a list of running domains :

```
# xm list
```

You will find more details about syntax commands and options thorough the `xm help` command.

- Using the Xen console : if you don't specify the `-c` option at start time, you can still use the console.

  ```
  # xm console Mandriva
  ```

  `Mandriva` is the name of the guest domain target. You can replace it with the guest ID displayed in the `xm list`.

  To quit a guest console, just type `Ctrl+]` (`Ctrl+$` dans une console virtuelle locale ou tty).

- Guest domain backup and restore: Administrators of a Xen system can suspend a virtual machine in its current state into a disk file in `domain0`, allowing it to resume at a later time.

  For example, you can suspend a domain called `mandriva` to disk :

  ```
  # xm save Mandriva mandriva.chk
  ```

  This stops the domain named `mandriva` and saves its current state into a file called **`mandriva.chk`**.

  To resume execution of this domain, use the `xm restore` command :

  ```
  # xm restore mandriva.chk
  ```

  This restores the state of the domain and resumes its execution. The domain will restart in the same state and the console may be reconnected using the `xm console Mandriva` command.

5.3.1.4. Troubleshooting

- If you get the following error :

  ```
  Error: Error creating domain (12, 'Cannot allocate memory')
  ```

  add the `dom0_mem` option on the Xen command line in the **`grub.conf`** entry :

- If you get the following error :

  ```
  Error: Error creating domain: Kernel image does not exist :
  ```

```
/boot/vmlinuz-2.6.18.8-xen-3.3.0-7mdv
```

please install the `kernel-xen` package.

5.3.2. KVM

5.3.2.1. Main concepts and web references

KVM (Kernel-based Virtual Machine) is a virtualization solution for x86 architectures based on Intel VT (vmx) and AMD-V (svm) technologies. To determine if such an extension is supported by your hardware, run the following command:

```
# grep '(vmx|svm)' /proc/cpuinfo
```

If you know that your processor includes a virtualisation extension even if it is not displayed, please check if your BIOS is well configured. If the option is not enable, do it. For the modification to be taken into account, a hardware reboot is not enough, you have to completely shut down the server.

With KVM, you can deploy Linux or Windows® virtual machines.

The KVM website is located here: linux-kvm.org (`http://www.linux-kvm.org/page/Main_Page`)

5.3.2.2. Installation

You will need both qemu-img and kvm packages

```
# urpmi qemu-img kvm
```

To use KVM, the kernel module must be loaded. You can verify it thanks to the following command:

```
# lsmod |grep kvm
```

- In case you have the AMD-V technology and if this module is not loaded, load it with:

```
# modprobe kvm-amd
```

- In case you have the Intel VT technology and if this module is not loaded, load it with:

  ```
  # modprobe kvm-intel
  ```

The user which will be used to create your virtual machines must belong to the kvm group (here, the user is named 'test').

```
#usermod -G kvm test
```

Check with the id command that you user belongs to the kvm group:

```
$id
uid=500(test) gid=500(test) groupes=422(kvm),500(test)
```

5.3.2.3. Installing Virtual Machines

Be aware you'll need enough free space on your hard disk to host your virtual machines. You can create a dedicated partition for that use.

First, you have to create an image of the guest virtual machine. The following command create an image of 3GiB maximum size.

```
$ qemu-img create mes5dvd 3G
Formatting 'mes5dvd', fmt=raw, size=3145728 kB
```

Different file formats exist. Please read the qemu-img man page to know the different possible file formats.

We are going to install Mandriva Enterprise Server 5 from the installation DVD media into the mes5dvd image we just have created.

```
kvm -k fr -m 512 -cdrom /dev/cdrom -drive file=mes5dvd
```

the options stands for:

- -k en-us : english keyboard.
- -m 512 : the virtual machine will have 512 MiB of RAM memory.
- -cdrom /dev/cdrom : the virtual machine of the cdrom device will be `/dev/cdrom`, i.e. the physical machine cdrom device.
- -drive file=mes5dvd : the harddisk of the machine will be the mes5dvd file.

The CD or DVD installation media present into you CDROM/DVDROM device is bootable. Once the command is launched, a window appears and the virtual machine installation process begins. Install Mandriva Enterprise Server 5 completely.

To launch your virtual machine:

```
kvm -k en-us -m 512 -cdrom /dev/cdrom -drive file=mes5dvd,boot=on
```

Many options can be passed to KVM (please read man kvm).

With the method described upwards, your virtual machine will have network (same network access than the host machine). KVM deals with the network configuration on its own. If you want to intall several virtual machines, you will have to configure a bridge connection and configure your virtual machines as a NAT. To know how to configure a bridge and how to use it with KVM, you can read the documentation available at linux-kvm.org (`http://www.linux-kvm.org/page/Networking`)

5.3.3. Virt-manager

The "Virtual Machine Manager" application (Virt-manager) is a desktop user interface to configure and manage virtual machines, virtual networks and storage. It is possible to manage Xen, KVM or QEMU guests.

5.3.3.1. How to start Virt-manager

Virt-manager uses libvirtd so it is necessary to start it:

```
#/etc/init.d/libvirtd start
```

To add libvirtd on default start:

```
#chkconfig libvirtd on
```

If you use Xen, start xend daemon. If you use KVM, then load the kvm module.

So, now that you are ready to start Virt-manager, choose Applications > Tools > Emulators > Virtual Machine Manager:

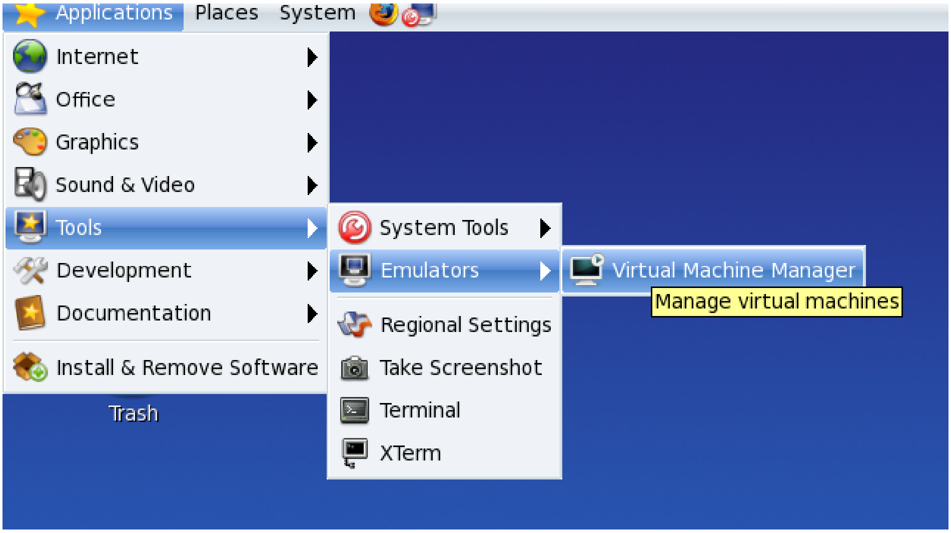

Figure 5-12. Launch Virt-manager

or use the command line :

```
#virt-manager
```

5.3.3.2. Virt-manager use

5.3.3.2.1. Guest creation

When Virt-manager is connected to your virtualisation system (Xen or KVM or QEMU), you can add a new guest by clicking on New.

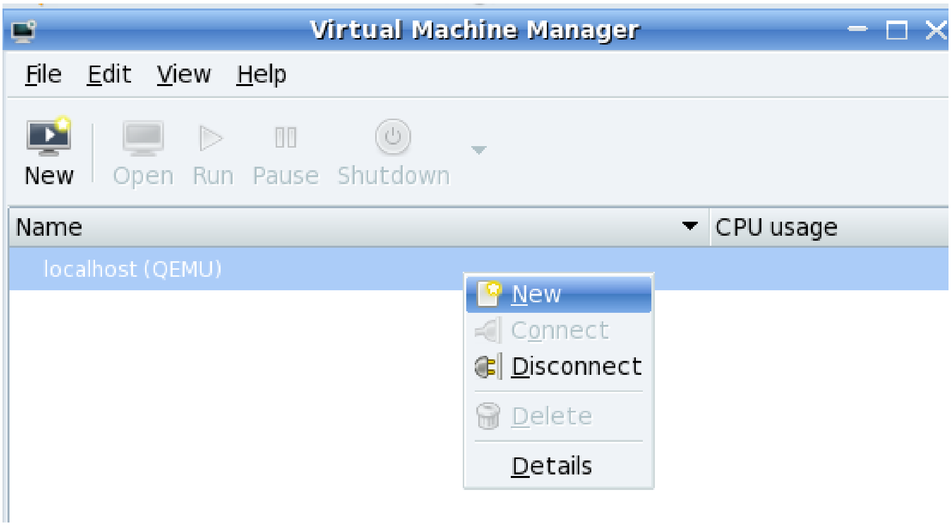

Figure 5-13. Guest's creation

5.3.3.2.2. Virtual networks and storage configuration

To configure your virtual networks and storage, click right on Your domain > Details

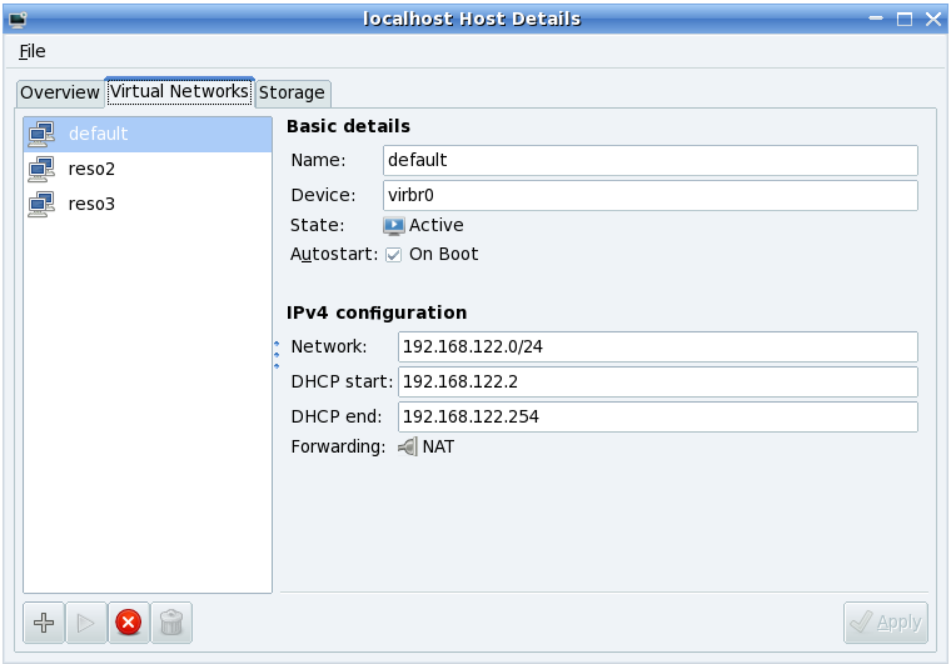

localhost Host Details

Basic details

Name: default

Device: virbr0

State: Active

Autostart: ☑ On Boot

IPv4 configuration

Network: 192.168.122.0/24

DHCP start: 192.168.122.2

DHCP end: 192.168.122.254

Forwarding: NAT

Figure 5-14. Virtual networks configuration

You can find more informations on official Virt-manager website (`http://virt-manager.et.redhat.com/`).

Chapter 6. Mandriva Directory Server

6.1. General review of Mandriva Directory Server

Mandriva Directory Server is an enterprise directory based on OpenLDAP. It can manage users profiles and accounts for a given area.

Mandriva Directory Server has been developped in a ver modular way and comes with additionnal modules so that it can also manage multi protocols files server (NFS, NETBIOS et WEBDAV) and users authentication through directory.

Mandriva Directory Server interface is Mandriva Management Console (MMC).

Mandriva Management Console is Pulse 2 and Mandriva Directory Server common interface. For more information you dan have a look on (`http://pulse2.mandriva.org/`)

Depending on your choice during installation, here is a list of available modules in this interface :

- Base module: users and groups
- Samba module: Windows® (PDC NT4) domain ressources
- Mail module: mail system
- Network module: DNS and DHCP
- Audit module: monitoring and reporting of modifications
- Password policy module: user password rules

In order to connect to Mandriva Management Console, use following adress in web browser: http://server_ip/mmc

6.1.1. Connection to Mandriva Directory Server

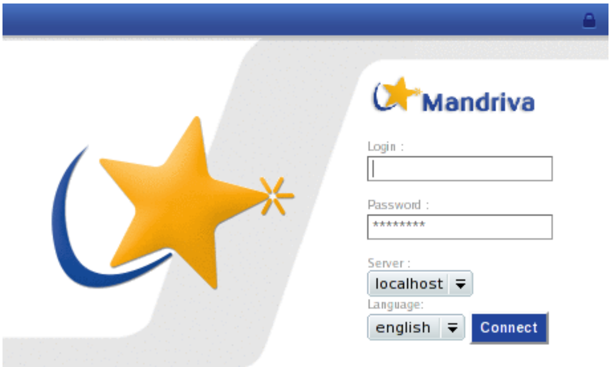

Figure 6-1. Mandriva Directory Server login page

In order to login, you must have Mandriva Directory Server account.

Mandriva Directory Server comes with 2 different kinds of account:

- root account: administration account of Mandriva Management Console. It grants access for all Mandriva Management Console functionnalities. This account is a LDAP account not system account. It has nothing to do with system root account.

  This account will have access for any part of interface at any time.

- All other accounts are users accounts created from Mandriva Management Console. By default they cannot login to Mandriva Management Console. You will have to add specific rights (ACLs).

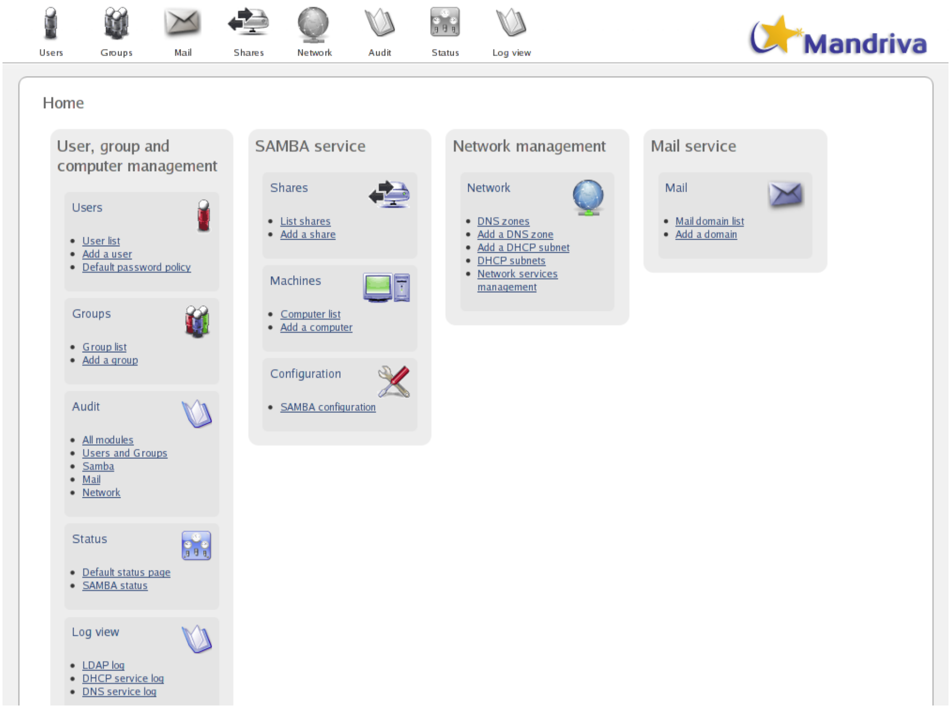

Home

User, group and computer management

Users
- User list
- Add a user
- Default password policy

Groups
- Group list
- Add a group

Audit
- All modules
- Users and Groups
- Samba
- Mail
- Network

Status
- Default status page
- SAMBA status

Log view
- LDAP log
- DHCP service log
- DNS service log

SAMBA service

Shares
- List shares
- Add a share

Machines
- Computer list
- Add a computer

Configuration
- SAMBA configuration

Network management

Network
- DNS zones
- Add a DNS zone
- Add a DHCP subnet
- DHCP subnets
- Network services management

Mail service

Mail
- Mail domain list
- Add a domain

Figure 6-2. Mandriva Directory Server Home Page

A user can choose a use level of interface. By default, this level is set to Normal. Just clic on Normal Mode to switch to Expert. This new level grants acces to some more advanced options.

Normal Mode provides a very simple interface so that use is easy and fast. We will only focus on Normal Mode in this document

6.2. Users accounts management

Users accounts database is managed in a LDAP directory of Mandriva Directory Server. It also stores all users accounts properties.

6.2.1. Users accounts list

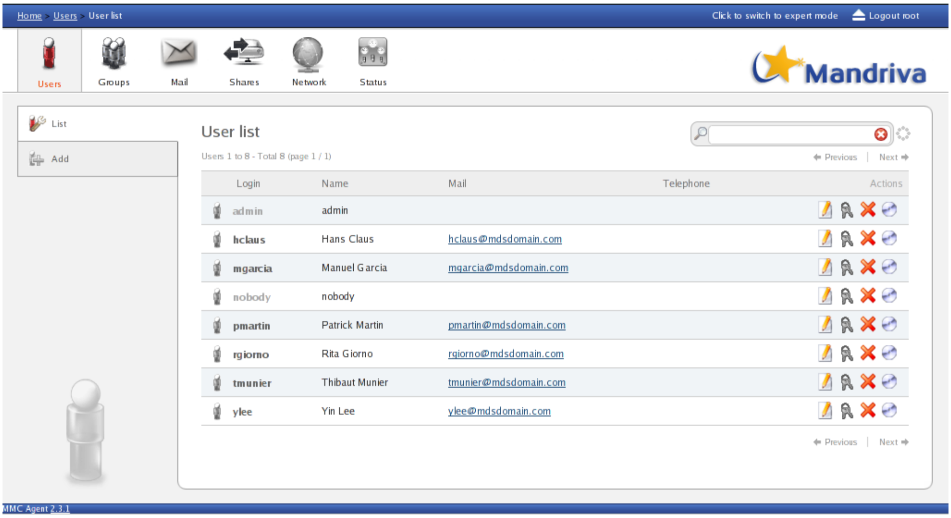

User list

Users 1 to 8 - Total 8 (page 1 / 1)

Login	Name	Mail	Telephone	Actions
admin	admin			
hclaus	Hans Claus	hclaus@mdsdomain.com		
mgarcia	Manuel Garcia	mgarcia@mdsdomain.com		
nobody	nobody			
pmartin	Patrick Martin	pmartin@mdsdomain.com		
rgiorno	Rita Giorno	rgiorno@mdsdomain.com		
tmunier	Thibaut Munier	tmunier@mdsdomain.com		
ylee	Yin Lee	ylee@mdsdomain.com		

Figure 6-3. Users accounts list

This list is divided into 4 columns:

- Login: nickname used to connect
- User name: name and forename of user
- User folder : path to user home directory on server
- Actions : available actions for this account: Modify this account, Modify access rights, Remove account and Backup Home directory.

6.2.2. Add a user account

Clic on Add tab.

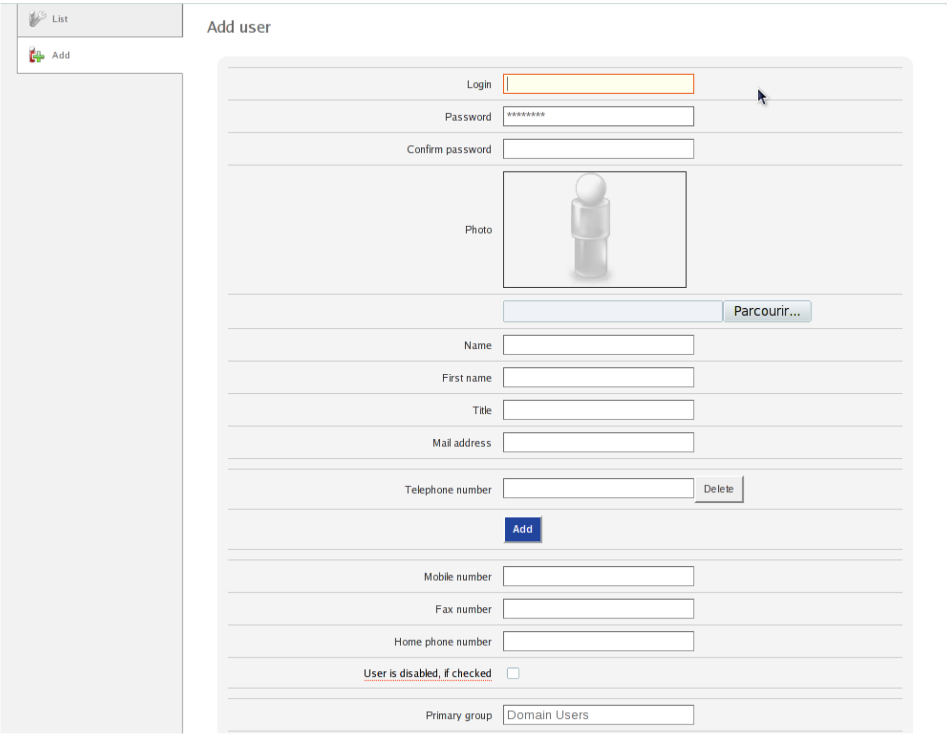

Password ********

Primary group Domain Users

Figure 6-4. Add a user account

Use this page to fill all properties for user account. It may be completed with some more properties depending on Mandriva Management Console modules already installed.

Basic properties are required. Depending on ressources to be provided to user you can add some more properties. All of them can be modified at any time.

Here is a list of basic properties:

- Login: Nickname to login
- Name: user name
- Forename: user forename
- Password, Confirm your password: user password. It must be filled twice. Password will never be printed in a clear way.
- Mail address: user mail address. This will be his main address used to connect on webmail.
- Phone number: user phone number

- Disabled user: If this option is enabled, UNIX account will be disabled. It means user will not be able to open any UNIX session. Defined shell will be set by default to `/bin/false`. Disabling UNIX account will not have any consequences on other properties for this account.

For defining username, it's strongly recommanded to set up a proper policy and follow a given pattern.

Once account is created, you cannot modify username. This only way would be to remove this account and recreate it.

When you switch to expert mode in Mandriva Management Console you will find some more parameters: Home directory, shell or command, UID, GID.

Primary group provides a way to set main group for user. A default primary group is filled. But you can modify it: remove the default one and enter the one you want to add. Completion will help you to fill it.

6.2.3. Remove a user account

Clic on red cross in action zone for a user account. A popup will ask you for confirmation of this. You will be able also to choose to delete personal files of this user or not.

If you enable Remove all user files when you want to remove an account, user home directory will be definitively removed.

6.3. Users groups management

A user can be part of one or several groups. Groups are used to manage users that would have same access rights.

6.3.1. Add a group

In order to add a group, clic on tab Add.

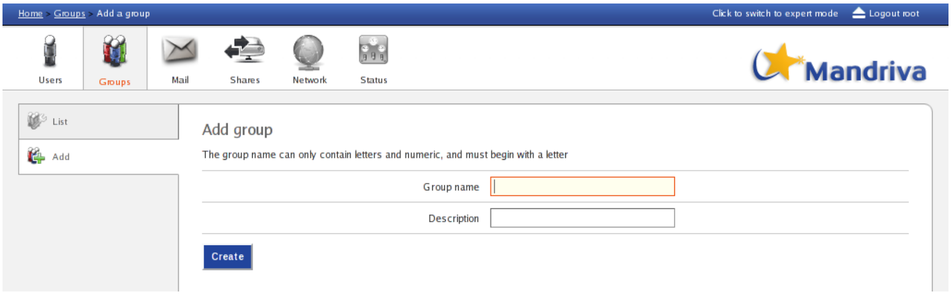

Add group

The group name can only contain letters and numeric, and must begin with a letter

Group name

Description

Figure 6-5. List of users groups

Fill name of group in field then clic on Create.

6.3.2. Edit groups members

In groups list, clic on icon for group edition.

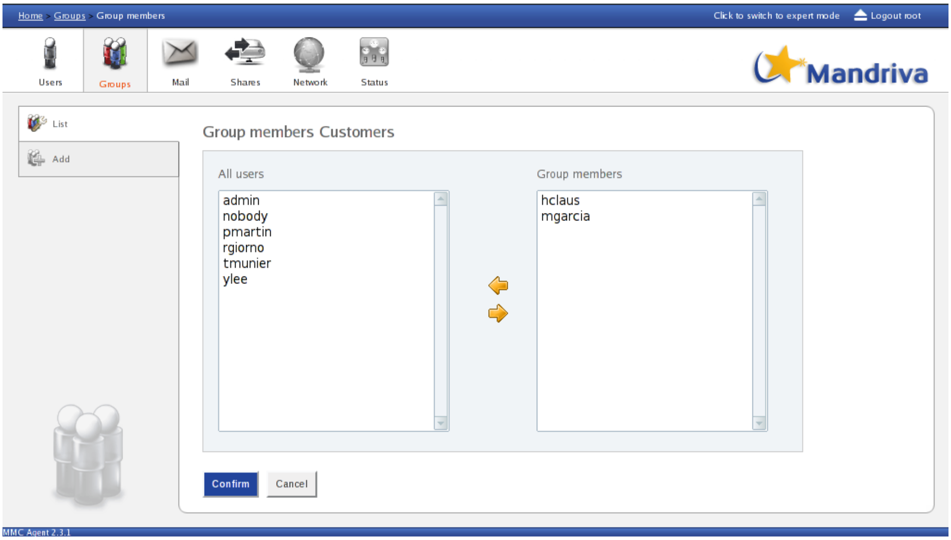

Group members Customers

All users	Group members
admin	hclaus
nobody	mgarcia
pmartin	
rgiorno	
tmunier	
ylee	

Figure 6-6. Edit groups members

Groups members edition page is divided in 2 parts:

- On the left side, you will find list of all registered users of LDAP directory
- On the right side, you will find the current list of users for a given group.

Using the 2 red arrows between these 2 lists, users are added or removed from a group. You can select several users at same time using right clic during selection. You can also select users through the two red arrows between the two lists, users are added or removed from the panel. You can select multiple users by right-clicking with the mouse and moving on users simultaneously. A discontinuous selection is possible by clicking on the users and at the same time pressing the button « Ctrl» (Control) on your keyboard.

It is not possible to remove a user's primary group. To do this, use the edit page of the user and change its primary group.

6.3.3. Delete a group

In groups list, clic on deletio icon (red cross).

Removing a group will not remove users of this group.

6.4. Shares Module (Samba)

Samba provides ressources for Windows® clients: files and printers sharing, authentication service. Samba is availables on Mandriva Directory Server and comes with all functionnalities of a Primary Domain Controler Windows® NT4. Samba module is available whe you clic on Shares tab.

Mandriva Management Console manage main parameters of Samba server of Mandriva Directory Server :

- Domain controler: Mandriva Directory Server can be a domain controler and manage users and hosts accounts for Windows®domain, instead of Windows® NT4 server.
- Files server: Mandriva Management Console can create shared directories on Mandriva Directory Server, available from Windows® clients.

6.4.1. Manage Samba users accounts

Samba users accounts properties can be compared to Windows® users accounts. If Mandriva Directory Server is set up as Primary Domain Controler (PDC) Windows®, you can add on existing accounts some Samba properties. These accounts will then be updated to Windows® accounts. Users will be able to connect on Windows®, access available ressources on domain.

6.4.1.1. Add a new Samba user

In Users module, clic on Add. Fill all required parameters of users as explained in Users management section. Check SAMBA access option.

You can also add Samba properties for an existing user.

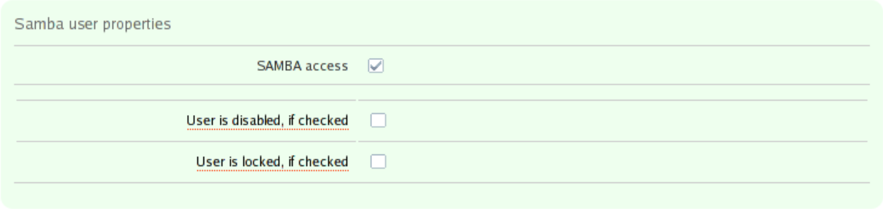

Samba user properties

SAMBA access ☑

User is disabled, if checked □

User is locked, if checked □

Figure 6-7. Samba properties

2 check boxes:

- Disabled user: if this option is enabled user cannot connect anymore to Windows® domain and use ressources managed by Samba.
- Locked user: depending on Samba configuration, you can lock a Samba user. It can happen because of multiple bad login for example. His account is the locked and he cannot connect anymore. Uncheck this option so that he can login again.

In expert mode, more parameters are available: Path for user profil, Start session script, Patch for Home Directory and Connect Home Diretory on network drive.

These parameters are usually fixed once in general options of Samba. If you remain it empy, general values will be used.

6.4.1.2. Remove a Samba user

In user edition page, just uncheck SAMBA access box and validate. User will not be able to access domain ressources anymore.

6.4.2. Samba shares management

6.4.2.1. Add a new share

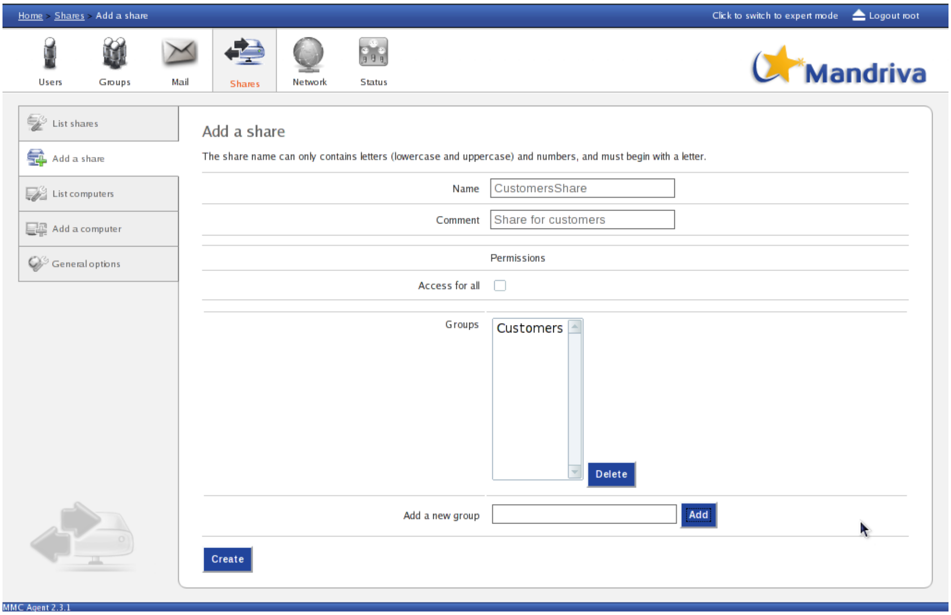

Add a share

The share name can only contains letters (lowercase and uppercase) and numbers, and must begin with a letter.

Name: CustomersShare

Comment: Share for customers

Permissions

Access for all □

Groups: Customers

Add a new group

Figure 6-8. Add a new share

Clic on Add a new share tab.

Fill following information:

- Name: Name of share. This is the one that will appear on network Windows®.
- Comments: ce texte This text will be available from Windows® clients when they list available shares on Mandriva Directory Server server.
- Antivirus on share : activate antivirus scan on opened files on this share. Infected files will be put in quarantine.

- Permissions: access right for share. Define users groups who will have read and write access on share. If Everybody box is checked , everybody will have access to this share. On the contrary if this option is disabled, only selected groups will access to this share.

In expert mode, 2 other options are available: Share is visible on domain and Admin groups for this share.

Everybody option is an easy way to create a new public share for all users..

6.4.2.2. Edit a share

Clic on edit icon in shares list.

You will have all share parameters. Apply your modifications and clic on Validate to save it.

If groups are removed from authorized groups for a share users of these groups will not be able to access share anymore

6.4.3. Hosts management

If Samba is configured as a PDC, you also can manage hosts included in Windows® domain.

Hosts accounts looks like users accounts. they are also stored in Mandriva Directory Server directory.

You will have to check This server is a PDC option in General options tab so that you can get more tabs: Hosts management and Add a new host.

6.4.3.1. Add a new host in domain

When a new host is registered in domain, users can use it to connect to Windows® domain. They will have access to all available shared ressources.

Registration method used by Mandriva Directory Server does not work with all Windows® versions. It's recommanded to register it directly from host interface for domain selection, in Windows® configuraton parameters.

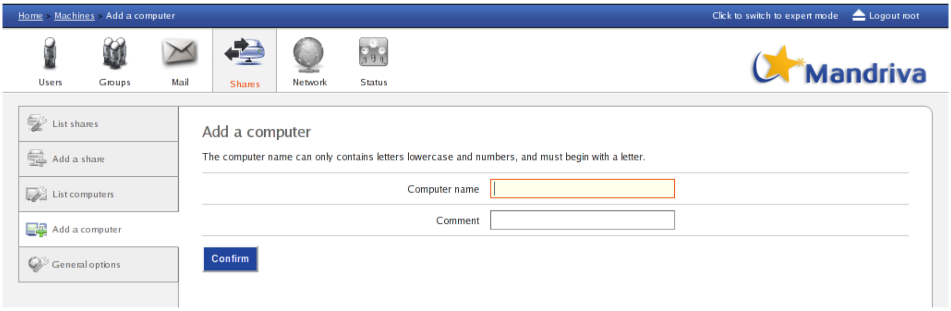

Add a computer

The computer name can only contains letters lowercase and numbers, and must begin with a letter.

Computer name

Comment

Figure 6-9. Add a new host

Clic on Add a new host.

Fill following information:

- Host name: NETBIOS name for this host on Windows® domain.
- Comments : comments for this host.

Finally clic on Add to register it on domain.

6.4.4. General options of Samba

General options tab provide an easy way to modify main options of Samba server.

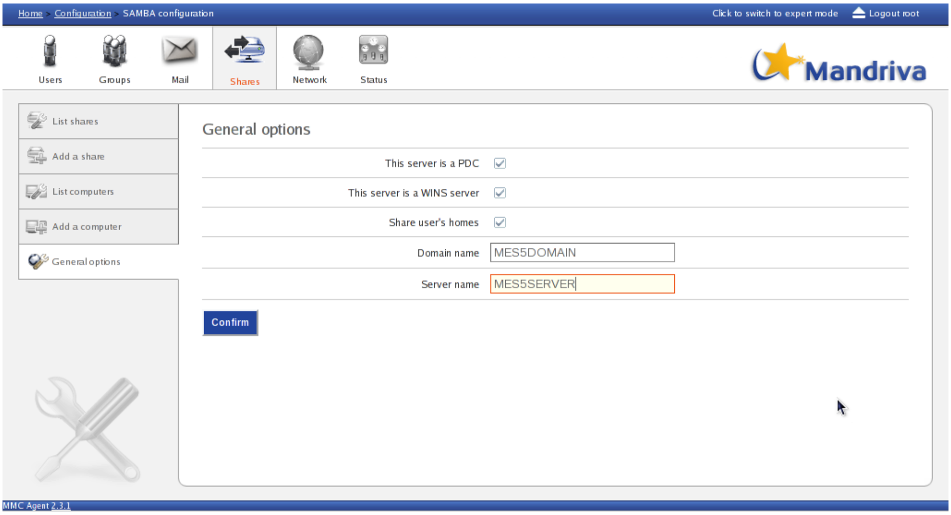

General options

This server is a PDC ☑

This server is a WINS server ☑

Share user's homes ☑

Domain name: MES5DOMAIN

Server name: MES5SERVER

Figure 6-10. Samba configuration

List of available configuration parameters:

- This server is a PDC: If this option is checked, Samba will behave as a Domain Controler Windows® NT4 (PDC). It will be seen on network as PDC and will provide services for all domain members.
- Share users Home Directories: user Home Directory will be seen as an available share, its name will be user name.
- Domain name: If Samba server is a PDC it will be the one controlled by PDC.

  Server name: NETBIOS name ofSamba server.

In expert mode you will find more options available: Path for user profil, Start session script, Home directory path and Connect home directory to network disk.

6.5. Mail Module

Mandriva Directory Server can be used together with mail service.

Mail service will use directory to:

- get information to deliver mails.
- identify users when they wnat to connect to their mail boxes.

Mandriva Management Console can help to configure:

- users mail address and aliases.

- group mail address for users.
- new mail domain if multiple mail domain option is enabled.

6.5.1. Running mail service

Mandriva Directory Server propose 2 ways of working mail service:

- one mail domain: mail service delivers mails for only one mail domain. All users will have same domain.
- Multiple mail domains (virtual domains): mail service can deliver multiple mail domain. These domains will be managed through our interface. If this option is enabled you will get one more tab: Mail in nav bar of Mandriva Management Console.

Depending on the option, you will get different configuration parameters.

6.5.2. Add and modify a mail user

Check Access mail service option is enabled so that a user can user mail service. Access will be complete when mail address is added.

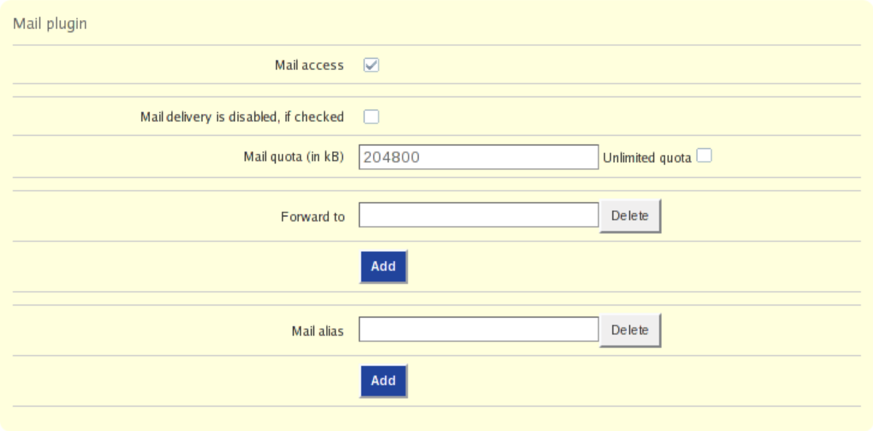

Mail plugin

Mail access ☑

Mail delivery is disabled, if checked □

Mail quota (in kB) 204800 Unlimited quota □

Forward to

Mail alias

Figure 6-11. Add a mail user

Common available parameters:

- Block mail delivery: stop for a while mail delivery only for this user.
- Mail alias: add some more mail address for this user.

In one mail domain configuration, you will find one more parameter:

- Internal recipient user(maildrop)

In multiple mail domain configuration:

- Forward to: all user mails will be forwarded to defined mail(s) address(es)
- Directory for mail delivery: this paramater is available in expert mode only. This directory will contains physically all received mail on server. Without any parameter, it will use default value

6.5.3. Remove a mail user

In user mode edition, uncheck Mail service access. i User will not receive mails anymore.

6.5.4. Add and remove users in a group mail

Use group mail to send easily a mail to several defined users. As an example you can send a mail for all users of `direction` group using `direction@domain.org` mail address. .

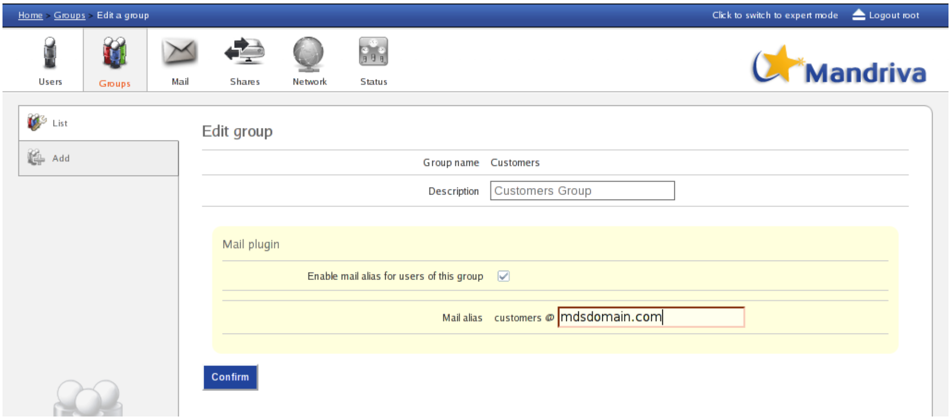

Edit group

Group name Customers

Description Customers Group

Mail plugin

Enable mail alias for users of this group ☑

Mail alias customers @ mdsdomain.com

Figure 6-12. Mail alias for a group of users

Go to Groups module, clic on group edition icon. Check Enable mail alias for members of this group. In multiple mail domain server, add also mail domain.

Clic on validate to save your all your modifications.

Uncheck box to remove a users mail group.

6.5.5. Mail domains management

This option is available for multiple mail domains server only.

In Mail module you fill find a list of all mail domains including a short description for each of them, but also number of users. You can use search bar to look for a specific one.

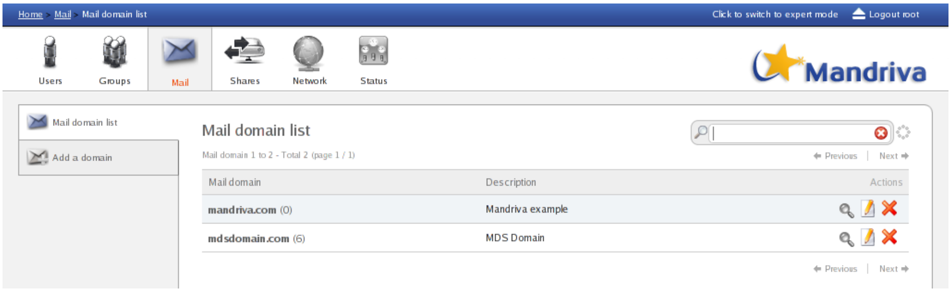

Mail domain list

Mail domain 1 to 2 - Total 2 (page 1 / 1)

Previous | Next

Mail domain	Description	Actions
mandriva.com (0)	Mandriva example	
mdsdomain.com (6)	MDS Domain	

Previous | Next

Figure 6-13. Mail domain list

Available actions for domains in list:

- See domain members: print all users accounts with mail address in this domain
- Domain edition: modify mail domain description.
- Remove domain: all members of this mail domain will not receive mails anymore.

Clic on Add a new domainto add a new mail domain.

Add following information:

- Mail domain: DNS name for mail domain (example : « mandriva.com »).
- Description: add comment to describe this domain.

6.6. Network Module

DNS/DHCP Mandriva Directory Server module can create and manage following components on a LAN:

- DNS zones: NS registers, A and CNAME. Reverse zone are defined automatically.

- DHCP subnets: configure hosts using static IPsbut also dynamic ranges of IPs.

You can also bind DNS zone and DHCP subnet. When you create a static host in a DHCP subnet A register will be automatically created in proper DNS zone.

DNS zones and DHCP configuration are stored in OpenLDAP directory.

6.6.1. Interface

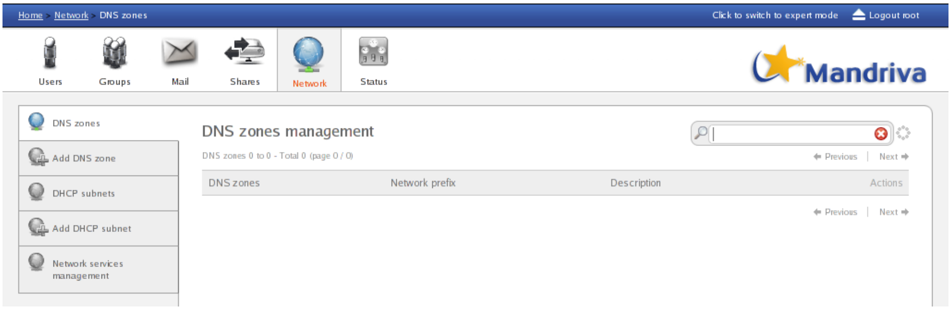

DNS zones management

DNS zones 0 to 0 - Total 0 (page 0 / 0)

DNS zones	Network prefix	Description	Actions

Figure 6-14. DNS/DHCP module

Mandriva Directory Server network module is splitted into 5 tabs in left side of interface:

- the first 2 one are dedicated to DNS module management.
- the 2 following are dedicated to DHCP module management
- the last one is focused on management of services status.

As proposed in all Mandriva Directory Server pages, you can proceed to dynamic search using field on right top part of page.

Again you will find more options in expert mode.

6.6.2. DNS management

6.6.2.1. Add a new zone

Clic on Add a DNS zone tab on left side.

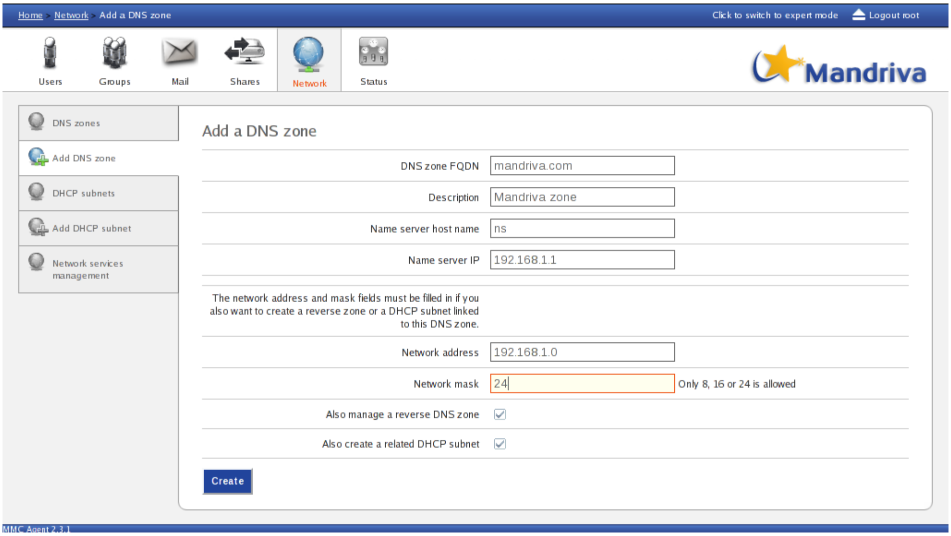

Add a DNS zone

DNS zone FQDN	mandriva.com
Description	Mandriva zone
Name server host name	ns
Name server IP	192.168.1.1

The network address and mask fields must be filled in if you also want to create a reverse zone or a DHCP subnet linked to this DNS zone.

Network address	192.168.1.0
Network mask	24 Only 8, 16 or 24 is allowed
Also manage a reverse DNS zone	☑
Also create a related DHCP subnet	☑

Figure 6-15. Add a DNS zone

You will find here basic information needed to create it:

- FQDN of DNS zone: fully qualified zone name.
- Description: text to describe this new zone
- Host name for name server: host name for name server that will be registered in NS entry of this zone.
- name server IP: This IP address will be associated to previous field.

This page offers also a way to create a DHCP subnet for this DNS zone

If you want to manage both services, you will have to fill following information:

- Network address: network address managed by DHCP service on this specific zone.
- Network mask: subnet mask.
- Manage reverse DNS zone: Revers zone will be automatically created and managed
- Create a DHCP subnet: DNS zone and subnet will be managed together.

As soon as zone is validated, you can see it in Mandriva Directory Server interface. Clic on DNS zones on left side of page.

You will find a list of all registered zones in Mandriva Directory Server :

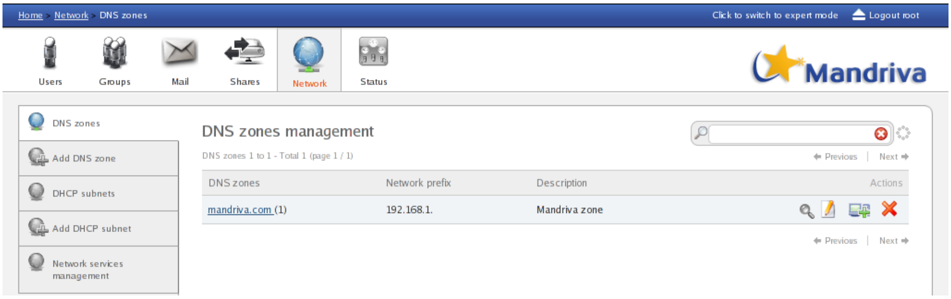

DNS zones management

DNS zones 1 to 1 - Total 1 (page 1 / 1)

DNS zones	Network prefix	Description	Actions
mandriva.com (1)	192.168.1.	Mandriva zone	

Figure 6-16. List of zones managed by Mandriva Directory Server

List is composed by zome name, description and network address. For each record you will find actions button to manage it.

> DHCP subnet created with DNS zone is for now very basic one. You can complete it, see next chapter about DHCP.

6.6.2.2. Add an A record

When zone is created, you can add A records: you can register a DNS name with an IP address

Here is the process:

- Clic on DNS zone on left part.
- Clic on zone name to be edited.
- Clic on Add a new hostbutton.

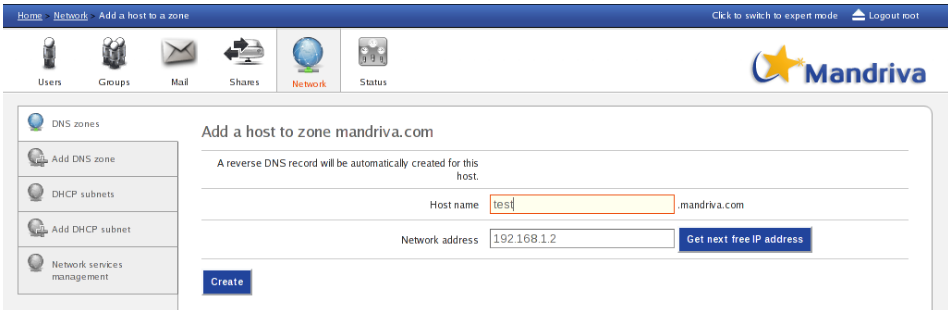

Add a host to zone mandriva.com

A reverse DNS record will be automatically created for this host.

Host name: test .mandriva.com

Network address: 192.168.1.2

Figure 6-17. Add a new A entry

You will find a screen as following:

- Fill host name in Host name field.
- Enter proper address for this new host.

Next IP address button looks for next available IP address.

New host will be added in DNS zone when you clic on Create button.

6.6.2.3. Remove an A record

Just check following steps:

- Click on DNS zones on left side.
- Click on name zone to be edited.
- Click on removal button of the proper record.

Removing a zone will also remove all CNAMEs included in it.

6.6.2.4. Manage CNAME records

CNAME records allow you to associate one or several DNS name to another DNS name. Thus, it is possible to access a computer by using different DNS names in a defined zone.

In Mandriva Directory Server user interface, a CNAME record is obviously related to a A type record.

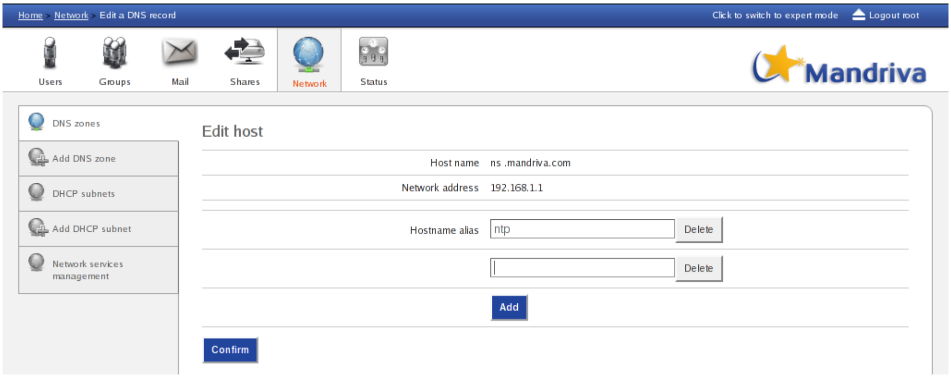

Edit host

Host name ns .mandriva.com

Network address 192.168.1.1

Hostname alias ntp

Figure 6-18. Adding a CNAME type record

To add a CNAME to a host, please read the following:

- Click on the DNS Zones on the left.
- Click on the name zone to edit.
- Click on the edition button related to the record you want to add a CNAME.

It is now possible to add or remove as many CNAMEs as you want thanks to the Add and Remove buttons.

6.6.2.5. Add an MX record

MX (Mail eXchange) records map a domain name to a list of smtp servers for that domain. The smtp server must be an A record of the concerned zone.

Here is the proces to add an MX record:

- Clic on DNS Zones on the left part.
- Clic on Edit zone.

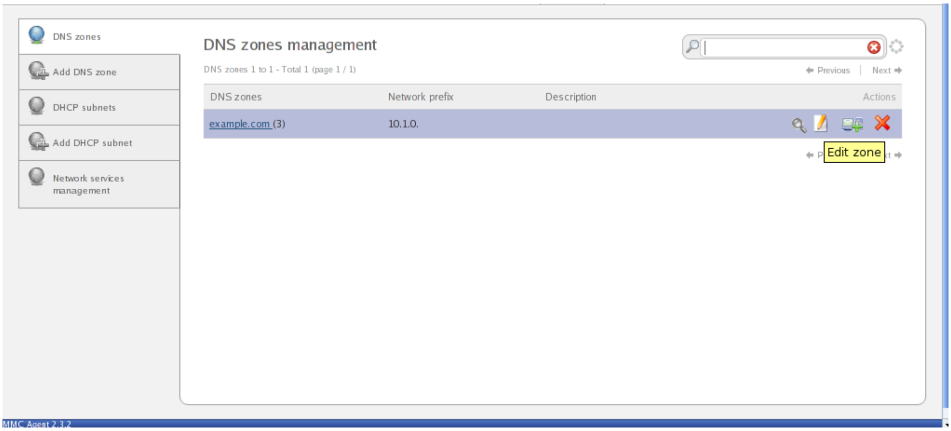

DNS zones management

DNS zones 1 to 1 - Total 1 (page 1 / 1)

DNS zones	Network prefix	Description	Actions
example.com (3)	10.1.0.		

Figure 6-19. Edit a DNS zone

- In the MX records (SMTP servers) area, add your MX record by using the format "distance mxserver". For example: 10 smtp.example.com.

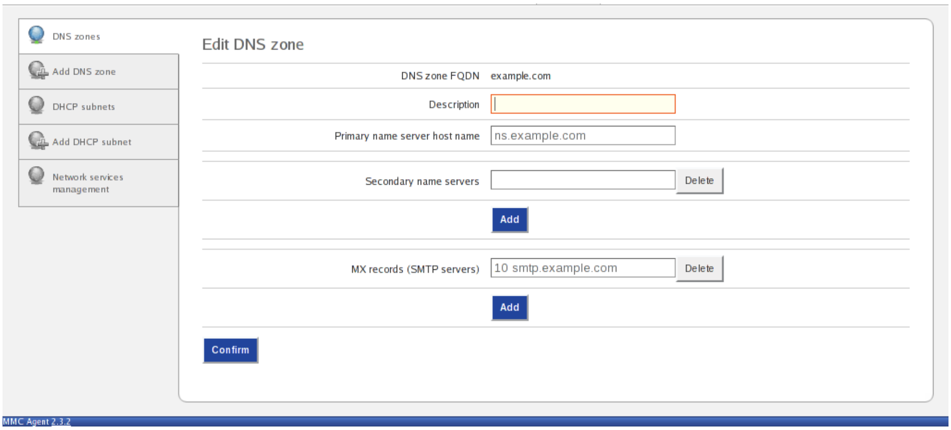

Edit DNS zone

DNS zone FQDN	example.com
Description	
Primary name server host name	ns.example.com
Secondary name servers	
MX records (SMTP servers)	10 smtp.example.com

Figure 6-20. MX record

6.6.2.6. Removing a zone

Removing a zone can be done by clicking on the remove button available on the right, in the zone list.

Removing a zone automatically remove all the former records in this zone, including in LDAP directory.

6.6.3. Managing DHCP

6.6.3.1. Add/Edit a DHCP subnet

To add DHCP subnet management, click on the Add a DHCP subnet tab, on the left.

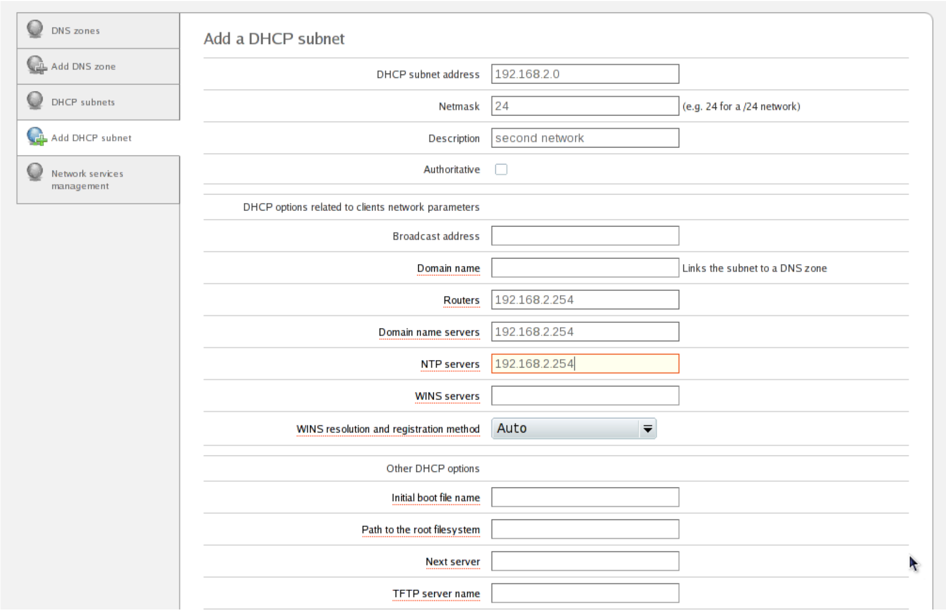

Add a DHCP subnet

DHCP subnet address	192.168.2.0	
Netmask	24	(e.g. 24 for a /24 network)
Description	second network	
Authoritative	□	
DHCP options related to clients network parameters		
Broadcast address		
Domain name		Links the subnet to a DNS zone
Routers	192.168.2.254	
Domain name servers	192.168.2.254	
NTP servers	192.168.2.254	
WINS servers		
WINS resolution and registration method	Auto	
Other DHCP options		
Initial boot file name		
Path to the root filesystem		
Next server		
TFTP server name		

Figure 6-21. Add a DHCP subnet

The presented page helps to fill the basic information required for the subnet creation:

- DHCP subnet Address: It is the network address to create.
- Network Mask: Type the network mask in the full format (8, 16 or 24).
- Description: It is a text field which allows to associate a network description.

The following listed fields allow to define the options which will be given to DHCP clients:

- Broadcast Address: it is the broadcast address that we want to give to clients.
- Domain Name: the indicated domain in this field will be used by clients to suffix their DNS requests. The clients will use the "ntp" name instead of "ntp.mandriva.com". Note that the filled domain is a defined domain in the DNS Zones of Mandriva Directory Server, the DHCP subnet will be automatically linked to this zone.
- Routeur: it is the gateway which will be given to clients by default.
- Domain Name Server: Type in this field a list of IP addresses separated by commas to indicate the available name servers on the network.
- NTP Servers: Type in this field a list of IP addresses separated by commas to indicate the available time servers on the network.

It is also possible to define advanced options for the DHCP server. For example, to allow network machines to boot on a PXE server or to modify the time of a DHCP lease.

Editing a DHCP subnet is made thanks to the same interface. To edit an existing subnet, simply click on the DHCP subnets tab then click on the Edit button corresponding the subnet to modify.

6.6.3.2. Configuring a dynamic IP range

For the machines which are still unregistered in your Mandriva Directory Server, you can define a dynamic address range which will be assigned to the machines whose physical addresses (MAC) are unknown.

To do so, just check the Dynamic Address range for unregistered DHCP Clients and indicate the first and the last addresses of the range.

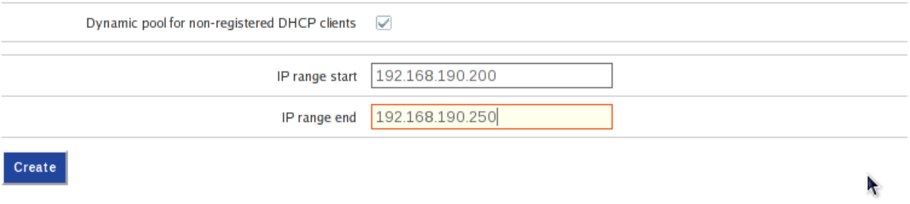

Dynamic pool for non-registered DHCP clients ☑

IP range start 192.168.190.200

IP range end 192.168.190.250

Figure 6-22. Dynamic address range

After the installation process and a restart of the DHCP service, clients will recieve a random address available from this dynamic range.

6.6.3.3. Add/Modify a static host configuration

Mandriva Directory Server allows to book an IP address for a particular machine according to its MAC address. By using this method, the client has the benefits of a given DHCP address (ease of configuration, parameters broadcasting, ...) and always keeps the same IP address.

To record the address reservation, read the following:

- Click on the DHCP Subnets tab on the left.
- Click on the network name to edit.
- Click on the Add a Static Host button.

You can regsiter an address MAC/IP address couple and some specific options to this machine (optional) as indicated in the screen here under:

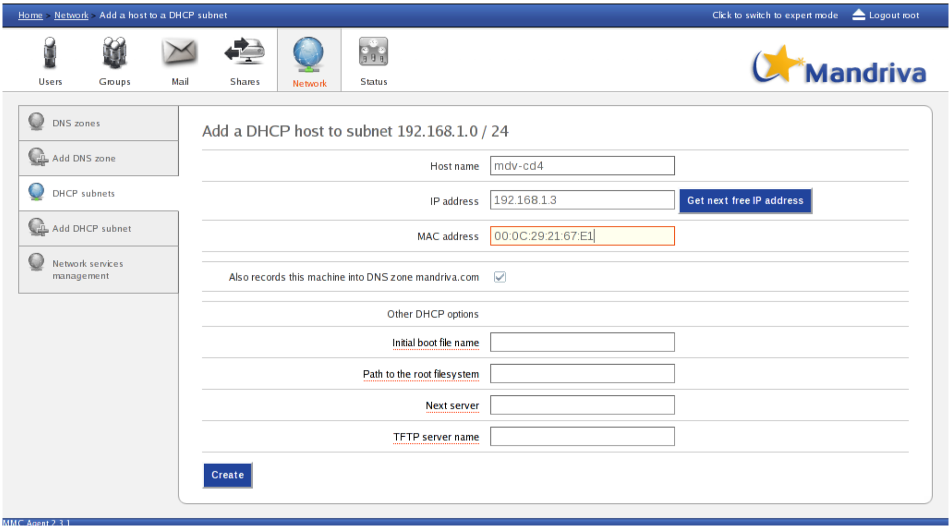

Add a DHCP host to subnet 192.168.1.0 / 24

Host name: mdv-cd4

IP address: 192.168.1.3

MAC address: 00:0C:29:21:67:E1

Also records this machine into DNS zone mandriva.com ☑

Other DHCP options

Initial boot file name:

Path to the root filesystem:

Next server:

TFTP server name:

Figure 6-23. Adding a static host (address reservation)

In this very example, the "mdv-cd4" machine with tbe 00:0C:29:21:67:E1 MAC adress will always receive the 192.168.190.2 IP address by DHCP.

6.6.3.4. Switch from a dynamic to a static configuration

The drawback of the presented method in the preceeding paragraph is that you have to know the physical address (MAC) of the machine. To avoid this problem, Mandriva Directory Server offers the possibility to transform a lease from a dynamic range into a static lease. Thus, you can regsiter a machine in DHCP without knowing its MAC address.

To apply this transformation, read the following:

- Click on the DHCP Subnets tab on the left.
- Click on the network name to edit.

The screen displays the list of static hosts and of machines with IP addresses from the dynamic range.

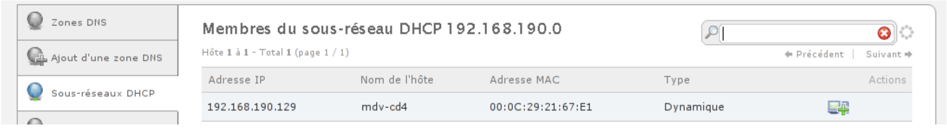

Membres du sous-réseau DHCP 192.168.190.0

Hôte 1 à 1 - Total 1 (page 1 / 1)

Adresse IP	Nom de l'hôte	Adresse MAC	Type	Actions
192.168.190.129	mdv-cd4	00:0C:29:21:67:E1	Dynamique	

Figure 6-24. Transformation of a dynamic host into a static one.

6.6.3.5. Remove a static configuration

To remove a booked address and for the machine to get an IP address from the dynamic range, you just have to click on the Remove button available in front of the machine from the subnet host list.

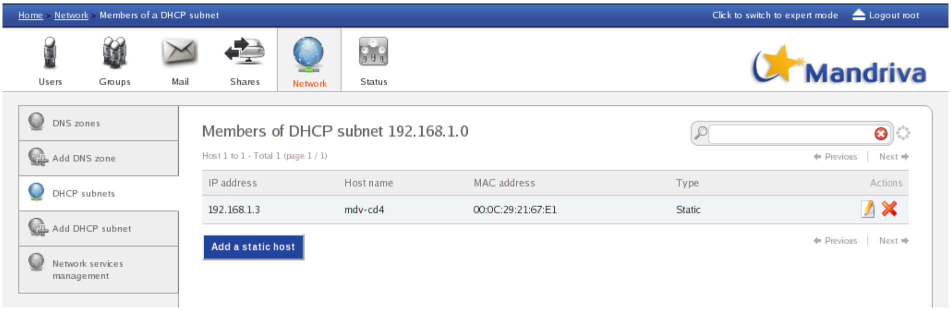

Members of DHCP subnet 192.168.1.0

Host 1 to 1 - Total 1 (page 1 / 1)

IP address	Host name	MAC address	Type	Actions
192.168.1.3	mdv-cd4	00:0C:29:21:67:E1	Static	

Figure 6-25. Removing a static host

6.6.3.6. Removing a DHCP subnet

Removing a DHCP subnet is simply made by clicking on the available button on the right side in the subnet list.

During the removal of a DHCP subnet all the static hosts owned by this very subnet will be removed.

6.6.4. Service Management

6.6.4.1. Stop and Start

The module includes an interface to manage the state of different managed services.

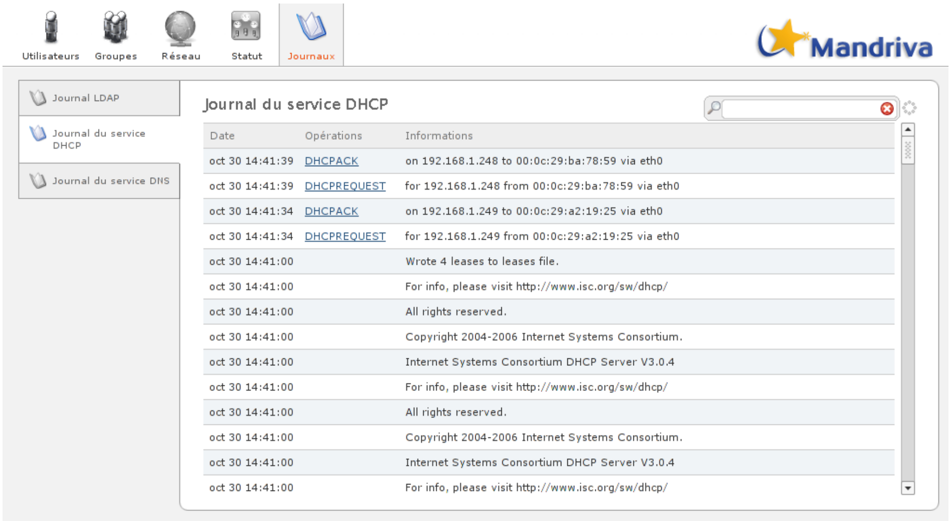

Journal du service DHCP

Date	Opérations	Informations
oct 30 14:41:39	DHCPACK	on 192.168.1.248 to 00:0c:29:ba:78:59 via eth0
oct 30 14:41:39	DHCPREQUEST	for 192.168.1.248 from 00:0c:29:ba:78:59 via eth0
oct 30 14:41:34	DHCPACK	on 192.168.1.249 to 00:0c:29:a2:19:25 via eth0
oct 30 14:41:34	DHCPREQUEST	for 192.168.1.249 from 00:0c:29:a2:19:25 via eth0
oct 30 14:41:00		Wrote 4 leases to leases file.
oct 30 14:41:00		For info, please visit http://www.isc.org/sw/dhcp/
oct 30 14:41:00		All rights reserved.
oct 30 14:41:00		Copyright 2004-2006 Internet Systems Consortium.
oct 30 14:41:00		Internet Systems Consortium DHCP Server V3.0.4
oct 30 14:41:00		For info, please visit http://www.isc.org/sw/dhcp/
oct 30 14:41:00		All rights reserved.
oct 30 14:41:00		Copyright 2004-2006 Internet Systems Consortium.
oct 30 14:41:00		Internet Systems Consortium DHCP Server V3.0.4
oct 30 14:41:00		For info, please visit http://www.isc.org/sw/dhcp/

Figure 6-26. DNS and DHCP Services Management

It is possible to stop and start each service with the Stop and Start buttons. To reload a service, simply click on the reload button

If the DNS servers are configured to be slave on this zone, you must reload the service to send the new parameters to the slave services.

6.6.4.2. Viewing logs

In expert mode, the module offers an interface to view logs for each service. This interface is available through the magnifier button in the service management page or in the Logs tab. Note that log display is automaticaly actualized.

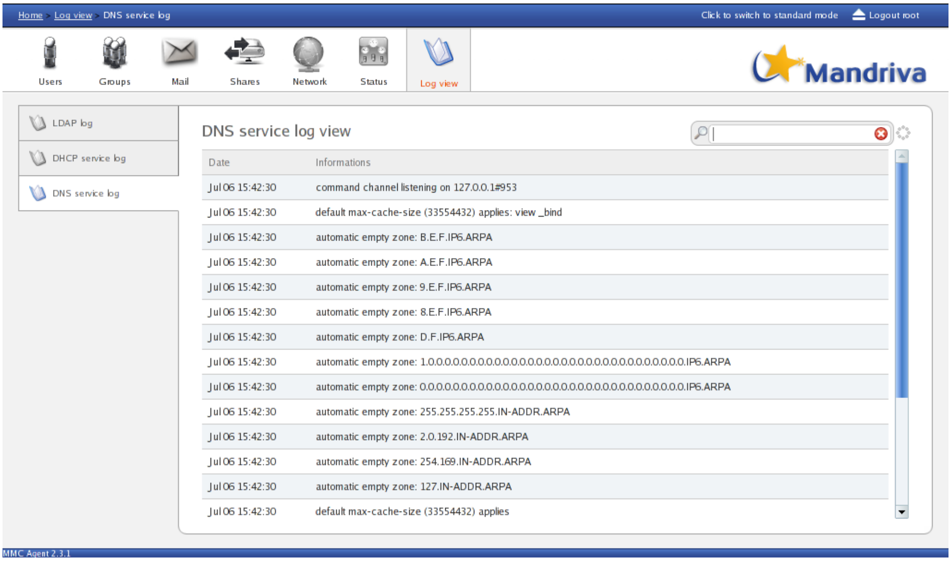

DNS service log view

Date	Informations
Jul 06 15:42:30	command channel listening on 127.0.0.1#953
Jul 06 15:42:30	default max-cache-size (33554432) applies: view _bind
Jul 06 15:42:30	automatic empty zone: B.E.F.IP6.ARPA
Jul 06 15:42:30	automatic empty zone: A.E.F.IP6.ARPA
Jul 06 15:42:30	automatic empty zone: 9.E.F.IP6.ARPA
Jul 06 15:42:30	automatic empty zone: 8.E.F.IP6.ARPA
Jul 06 15:42:30	automatic empty zone: D.F.IP6.ARPA
Jul 06 15:42:30	automatic empty zone: 1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.IP6.ARPA
Jul 06 15:42:30	automatic empty zone: 0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.IP6.ARPA
Jul 06 15:42:30	automatic empty zone: 255.255.255.255.IN-ADDR.ARPA
Jul 06 15:42:30	automatic empty zone: 2.0.192.IN-ADDR.ARPA
Jul 06 15:42:30	automatic empty zone: 254.169.IN-ADDR.ARPA
Jul 06 15:42:30	automatic empty zone: 127.IN-ADDR.ARPA
Jul 06 15:42:30	default max-cache-size (33554432) applies

Figure 6-27. Viewing DNS logs

6.7. Audit module

The audit module allows to trace all transactions made by users in the Mandriva Management Console web interface. Reports web pages allow then to the administrator to know who did what and when.

The following operations are traced :

- LDAP modifications : For example: create, modify, or delete a user.
- filesystem related modifications : For example: create a Samba share.
- service management : For example: stopping DHCP service.

All operations considered as critical are recorded.

6.7.1. Audit tab

To view the actions performed on all modules, click on Audit tab.

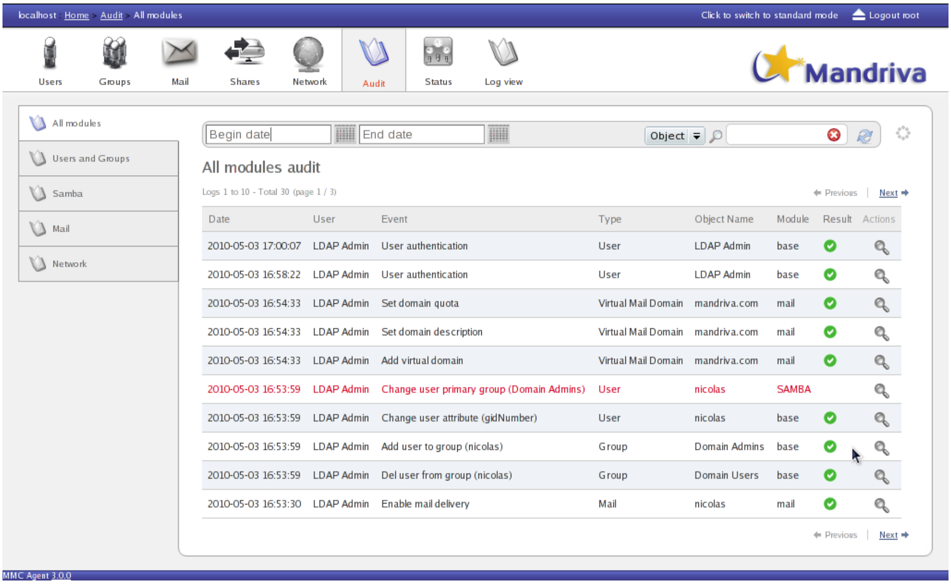

All modules audit

Logs 1 to 10 - Total 30 (page 1 / 3)

Date	User	Event	Type	Object Name	Module	Result	Actions
2010-05-03 17:00:07	LDAP Admin	User authentication	User	LDAP Admin	base	✓	🔍
2010-05-03 16:58:22	LDAP Admin	User authentication	User	LDAP Admin	base	✓	🔍
2010-05-03 16:54:33	LDAP Admin	Set domain quota	Virtual Mail Domain	mandriva.com	mail	✓	🔍
2010-05-03 16:54:33	LDAP Admin	Set domain description	Virtual Mail Domain	mandriva.com	mail	✓	🔍
2010-05-03 16:54:33	LDAP Admin	Add virtual domain	Virtual Mail Domain	mandriva.com	mail	✓	🔍
2010-05-03 16:53:59	LDAP Admin	Change user primary group (Domain Admins)	User	nicolas	SAMBA		🔍
2010-05-03 16:53:59	LDAP Admin	Change user attribute (gidNumber)	User	nicolas	base	✓	🔍
2010-05-03 16:53:59	LDAP Admin	Add user to group (nicolas)	Group	Domain Admins	base	✓	🔍
2010-05-03 16:53:59	LDAP Admin	Del user from group (nicolas)	Group	Domain Users	base	✓	🔍
2010-05-03 16:53:30	LDAP Admin	Enable mail delivery	Mail	nicolas	mail	✓	🔍

Figure 6-28. List of changes in all modules

Events are displayed in reverse chronological order, to directly show the latest events. The report columns are:

- Date : the date of the event ;
- User : login of the user who triggered the event ;
- Event : the event name ;
- Type : the type of the object on which the event happened ;
- Object Name : the name of the object on which the event happened ;
- Result : indicates whether the event has been successfully completed. If yes, a green icon is displayed in the column. Otherwise, the column is empty, and the line is displayed in red ;
- Actions : it is possible to have more details about an event by clicking on the glass icon.

Several filters are available:

- the period of events to display ;

- the name of the object on which you want to display all the associated events. For example, to view all transactions related to the user "jdoe", choose object in the dropdown list, and learn "jdoe" in the search box ;
- the type of objects on which you want to display all events associated. For example, to display all operations related to DNS zone, choose Type in the dropdown list, then DNS Zone ;
- the type of operation. For example, to display all the user delete operations, choose Action in the dropdown list, then Remove user ;
- the user triggers an operation. To view all transactions made by "jdoe", choose User dropdown list, then complete "jdoe" in the search box ;
- tabs on the left can filter the transactions by the modules available in the Mandriva Management Console. For example, to display all operations of the Samba module, click on the left tab Samba.

6.7.2. User audit

On the user edition page, the field Last action shows the timestamp of the last operation that was performed on the user account. When you click the timestamp, the event log of the user account is displayed.

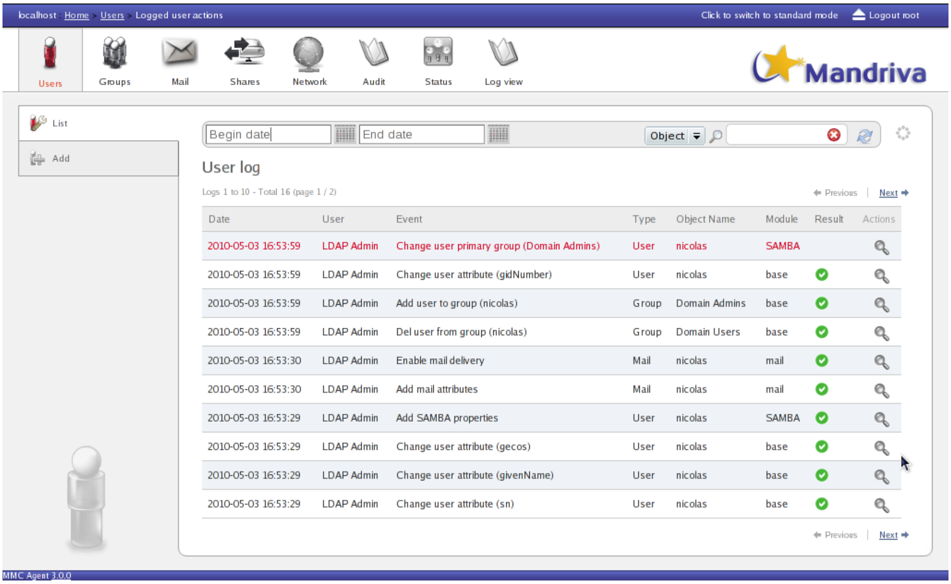

User log

Logs 1 to 10 - Total 16 (page 1 / 2)

← Previous | Next →

Date	User	Event	Type	Object Name	Module	Result	Actions
2010-05-03 16:53:59	LDAP Admin	Change user primary group (Domain Admins)	User	nicolas	SAMBA		🔍
2010-05-03 16:53:59	LDAP Admin	Change user attribute (gidNumber)	User	nicolas	base	✓	🔍
2010-05-03 16:53:59	LDAP Admin	Add user to group (nicolas)	Group	Domain Admins	base	✓	🔍
2010-05-03 16:53:59	LDAP Admin	Del user from group (nicolas)	Group	Domain Users	base	✓	🔍
2010-05-03 16:53:30	LDAP Admin	Enable mail delivery	Mail	nicolas	mail	✓	🔍
2010-05-03 16:53:30	LDAP Admin	Add mail attributes	Mail	nicolas	mail	✓	🔍
2010-05-03 16:53:29	LDAP Admin	Add SAMBA properties	User	nicolas	SAMBA	✓	🔍
2010-05-03 16:53:29	LDAP Admin	Change user attribute (gecos)	User	nicolas	base	✓	🔍
2010-05-03 16:53:29	LDAP Admin	Change user attribute (givenName)	User	nicolas	base	✓	🔍
2010-05-03 16:53:29	LDAP Admin	Change user attribute (sn)	User	nicolas	base	✓	🔍

← Previous | Next →

Figure 6-29. User audit

This report looks the same as the Audit tab report, and also includes some filters.

6.8. LDAP password policy module

This module is made of two parts:

- By default, an LDAP server does not check the passwords quality for user accounts, and does not have rules such as validity period, for example. The first part of this module is therefore a particular configuration of the LDAP server so that it enforces password policies ;
- These policies are LDAP objects, so the password policy module of Mandriva Management Console can adjust the settings.

The Mandriva Management Console allows to manage two types of policies :

- default password policy : it is applied to all users of the LDAP directory, and therefore all users managed with the Mandriva Management Console web interface ;
- password policy per user : if a policy is enabled on a user account, the default policy is ignored for this account.

6.8.1. Available Settings

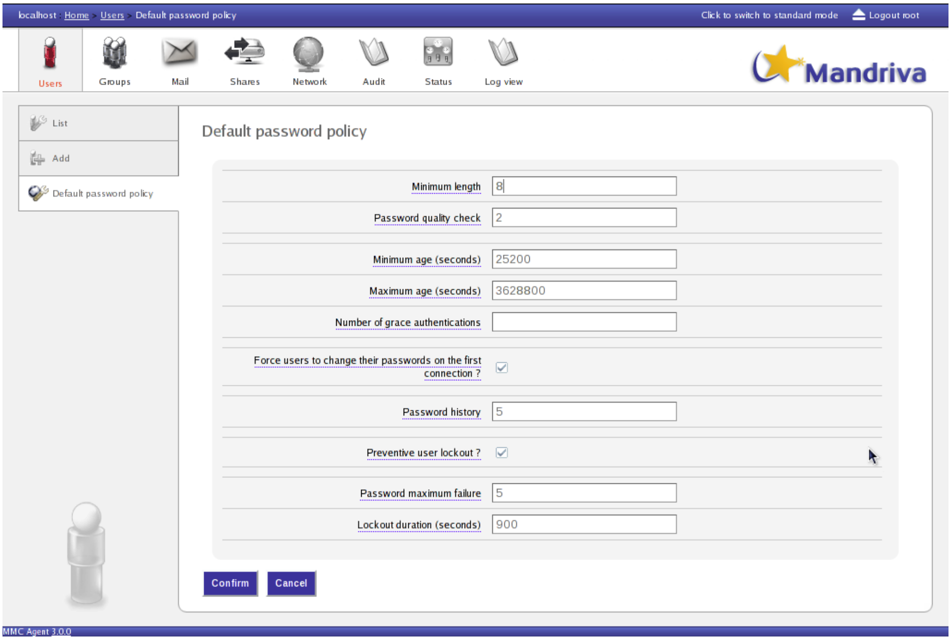

Default password policy

Minimum length	8
Password quality check	2
Minimum age (seconds)	25200
Maximum age (seconds)	3628800
Number of grace authentications	
Force users to change their passwords on the first connection ?	☑
Password history	5
Preventive user lockout ?	☑
Password maximum failure	5
Lockout duration (seconds)	900

Figure 6-30. Default password policy settings

A password policy have the following parameters :

- Minimum length : this attribute contains the minimum number of characters that will be accepted in a password ;
- Password quality check : this attribute indicates how the password quality will be verified while being modified or added. If this attribute is not present, or if the value is 0, quality checking will not be enforced. A value of 1 indicates that the server will check the quality, and if the server is unable to check it (due to a hashed password or other reasons) it will be accepted. A value of 2 indicates that the server will check the quality, and if the server is unable to verify it, it will return an error refusing the password ;
- Minimum age (in seconds) : this attribute holds the number of seconds that must elapse between modifications to the password. If this attribute is not present, 0 seconds is assumed (i.e. the password may be modified whenever and however often is desired) ;
- Maximum age (in seconds) : this attribute holds the number of seconds after which a modified password will expire. If this attribute is not present, or if the value is 0 the password does not expire. ;

- Number of grace authentications : this attribute contains the number of times that an expired password may be used to authenticate a user to the directory. If this attribute is not present or if its value is zero, users with expired password will not be allowed to authenticate ;
- Force users to change their passwords on the first connection ? : This flag specifies whether users must change their passwords when they first bind to the directory after a password is set or reset by the administrator ;
- Password history : This attribute is used to specify the maximum number of used passwords. If the attribute is not present, or if its value is 0, used passwords will not be stored and thus any previously-used password may be reused ;
- Preventive lockout user : This flag indicates, when enabled, that the password may not be used to authenticate after a specified number of consecutive failed bind attempts. The maximum number of consecutive failed bind attempts is specified in the Password maximum failure field below ;
- Password maximum failure : This attribute specifies the number of consecutive failed bind attempts after which the password may not be used to authenticate. If this attribute is not present, or if the value is 0, this policy is not checked, and the value of Preventive user lockout will be ignored ;
- Lockout duration (in seconds) : This attribute holds the number of seconds that the password cannot be used to authenticate due to too many failed authentication attempts. If this attribute is empty, or if the value is 0 the password cannot be used to authenticate until reset by a password administrator.

6.8.2. Password quality test

If the password quality test is enabled, the password must comply to all the following criteria:

- it contains at least one number ;
- it contains at least one lowercase letter ;
- it contains at least one uppercase letter ;
- it contains at least one special character (example: #, $, %, ...) ;
- all characters are different.

6.8.3. Users connection to the Mandriva Management Console

If the password policy module is enabled, a user who connects to the Mandriva Management Console gets a warning message if:

- the user password has been re-initialized by an administrator. The LDAP user account is restricted, and the user should change as soon as possible his/her password ;
- the user account is in grace period. The user must then change his/her password as soon as possible, otherwise the user account will be locked.

6.8.4. Resetting a user password

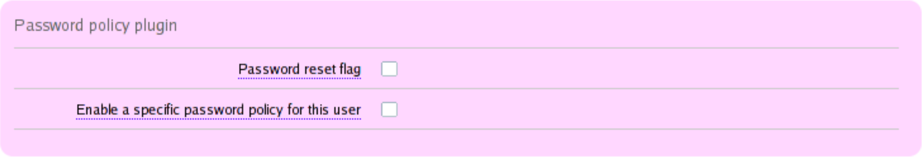

Password policy plugin

Password reset flag □

Enable a specific password policy for this user □

Figure 6-31. User password settings

A password reset flag is available on the edit page of the user account. If set, and if the password policy applied to the user force to change his/her password at first login, a message will warn him/her that he/she must change his/her password (see paragraph above).

6.9. Authentication on workstations

When Samba is installed on Mandriva Enterprise Server 5, it is possible to login on workstations by authenticating on domain using LDAP, either on Windows® or Linux.

For this, it is required that:

- the workstation has joined the Samba domain;
- user accounts have Samba access rights (for Windows® clients only).

6.9.1. Windows® clients

For Windows® 7 (Seven) workstations, it is necessary to perform prior actions. Some Windows® registry settings must indeed be changed - cf. (http://wiki.samba.org/index.php/Windows7):

```
HKLM\System\CCS\Services\LanmanWorkstation\Parameters
DWORD  DomainCompatibilityMode = 1
DWORD  DNSNameResolutionRequired = 0
```

To join a Windows® workstation to a Samba domain, you have to access Systems Properties. For example, right-click on Computer then Properties.

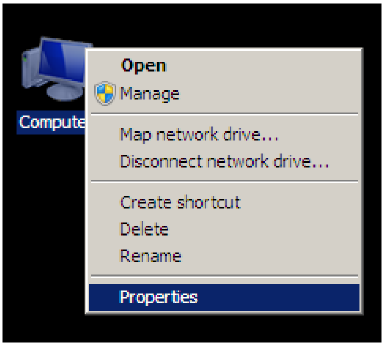

Figure 6-32. Join Windows® to a Samba domain: 1

Go to the Computeur Name tab and click on the Change... button.

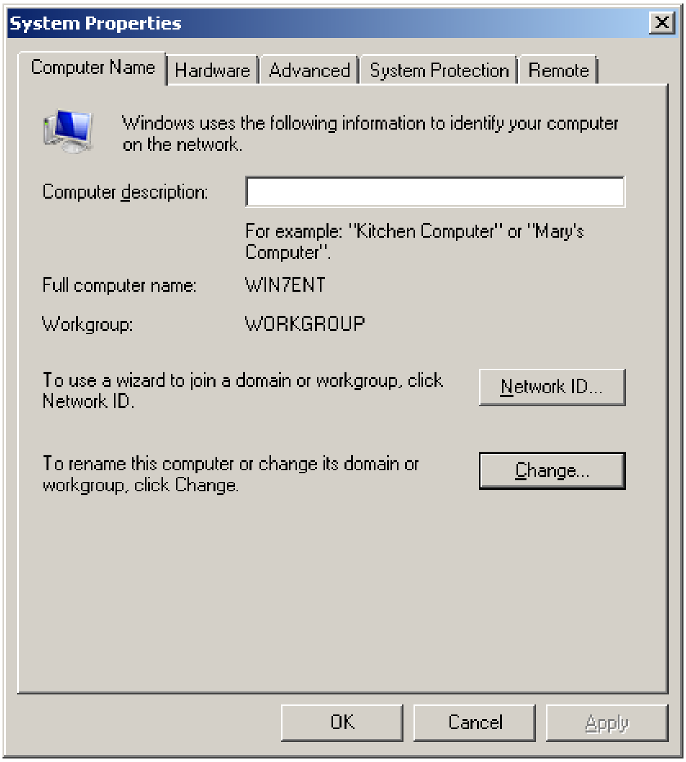

Windows uses the following information to identify your computer on the network.

Computer description:

For example: "Kitchen Computer" or "Mary's Computer".

Full computer name: WIN7ENT

Workgroup: WORKGROUP

To use a wizard to join a domain or workgroup, click Network ID.

To rename this computer or change its domain or workgroup, click Change.

Figure 6-33. Join Windows® to a Samba domain: 2

Fill the domain name and confirm by clicking on the OK button.

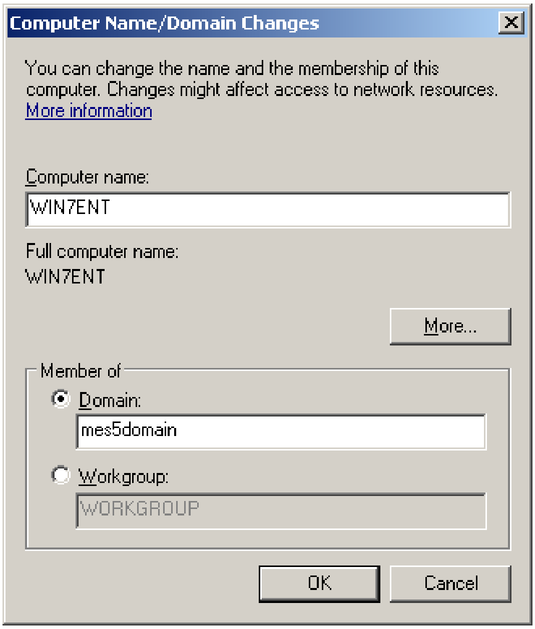

Computer Name/Domain Changes

You can change the name and the membership of this computer. Changes might affect access to network resources. More information

Computer name: WIN7ENT

Full computer name: WIN7ENT

Member of

● Domain: mes5domain

○ Workgroup: WORKGROUP

Figure 6-34. Join Windows® to a Samba domain: 3

An authentication window appears. Fill then the administrator account that is allowed to join the Samba domain and press the OK button.

This account correspond to a domain Administrator account created in the Mandriva Management Console.

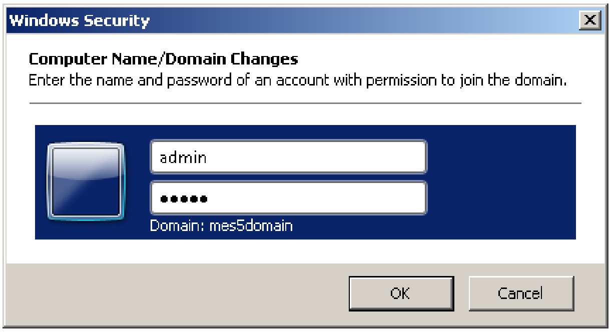

Windows Security

Computer Name/Domain Changes

Enter the name and password of an account with permission to join the domain.

admin

●●●●●

Domain: mes5domain

Figure 6-35. Join Windows® to a Samba domain: 4

This operation may take some time. A new window should appear telling you that you have joined the domain (Mandriva in this example).

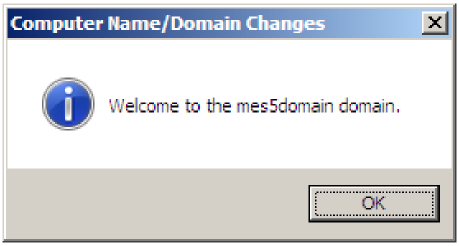

Computer Name/Domain Changes

Welcome to the mes5domain domain.

Figure 6-36. Join Windows® to a Samba domain: 5

You must now reboot the workstation. The Windows® authentication page should offer the domain choice (Mandriva in this example). Log in with a domain User account (user1 in this example).

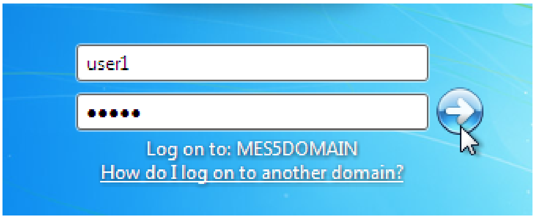

user1

•••••

Log on to: MES5DOMAIN

Figure 6-37. Login in a Samba domain under Windows®

6.9.2. Mandriva workstation

The authentication method used on GNU/Linux will directly be LDAP.

⚠ For this authentication method, it is necessary to open 389 TCP port in the Mandriva Directory Server server.

To configure the Mandriva Linux workstation, use the tool Configure your Computer from the menu. Once in the tool, go to the System tab, then click on Authentication as shown in the screenshot below.

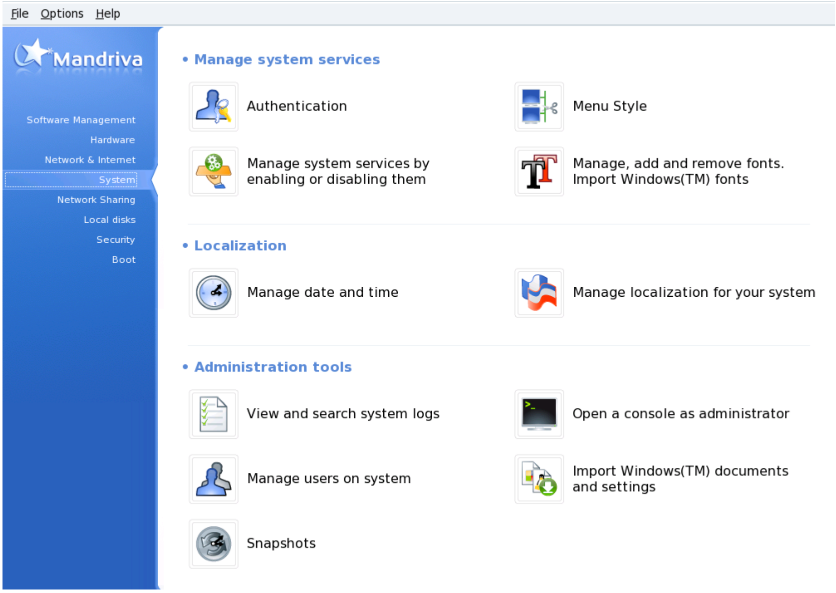

Figure 6-38. LDAP Authentication in Mandriva: 1

Chose the authentication method, ie LDAP.

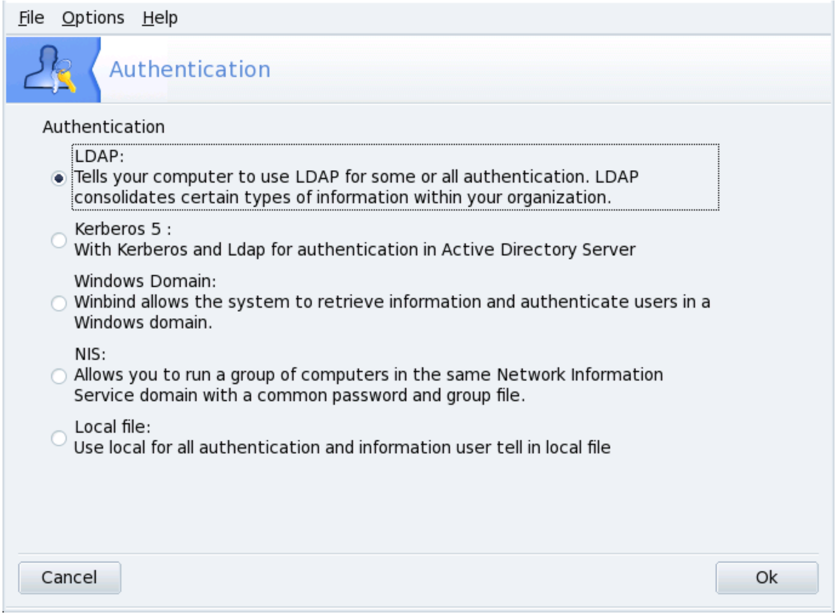

Authentication

● LDAP:
Tells your computer to use LDAP for some or all authentication. LDAP consolidates certain types of information within your organization.

○ Kerberos 5 :
With Kerberos and Ldap for authentication in Active Directory Server

○ Windows Domain:
Winbind allows the system to retrieve information and authenticate users in a Windows domain.

○ NIS:
Allows you to run a group of computers in the same Network Information Service domain with a common password and group file.

○ Local file:
Use local for all authentication and information user tell in local file

Figure 6-39. LDAP Authentication in Mandriva: 2

In the next window, fill the LDAP server in the right field, ie FQDN name (in preference) or the Mandriva Directory Server IP address.

Click then on Fetch base Dn, the Root DN field should automatically fill with the suffix of your LDAP directory. Finally, click on OK button.

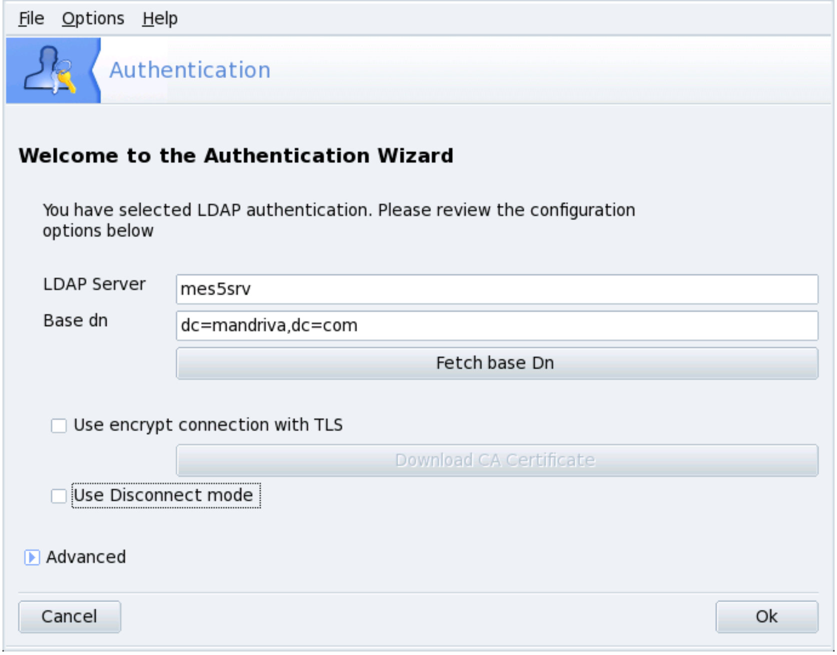

Welcome to the Authentication Wizard

You have selected LDAP authentication. Please review the configuration options below

LDAP Server: mes5srv

Base dn: dc=mandriva,dc=com

□ Use encrypt connection with TLS

□ Use Disconnect mode

Figure 6-40. LDAP Authentication in Mandriva: 3

To view the domain accounts, you can run the following command in a console :

```
# getent passwd
```

Chapter 7. Middleware Stacks

Middleware Management with Mandriva Enterprise Server 5

As described in the stack diagram, Mandriva Enterprise Server 5 offers a number of middleware components, among the most common services for Enterprise servers:

- Identity servers: OpenLDAP directory server and Kerberos network authentication server;
- Database servers: MySQL and PostgreSQL, the most popular open source databases.

7.1. Managing web services: LAMP and Proxy

7.1.1. Managing a LAMP server

This chapter presents the implementation of a LAMP platform (Linux, Apache, MySQL, PHP). It describes the installation method, service specific elements, as well a the initial configuration.

More information is available at:

- Apache official website (`http://httpd.apache.org`) ;
- Apache official documentation (`http://httpd.apache.org/docs/2.2`).

7.1.1.1. Apache installation and tree.

Mandriva Enterprise Server 5 provides a number of packages to install the Apache web server, including respectively the server itself, tools and modules.

- `apache-mpm-prefork` : contains the server daemon, MPM (Multi-Processing Module). If you chose to install Mandriva Server Setup, `apache-mpm-prefork` was automatically installed ;
- `apache-base` : includes Apache tools such as Apache Bench (ab) for load testing, tools and logs;
- `apache-modules` : contains the Apache base modules;

- apache-conf : contains all the Apache configuration files;
- apache-doc : contains the Apache official documentation.

You can complete this list according to your needs and install more modules to handle PHP scripts, SSL secure connections, authentication, etc. These packages are names apache-mod_*.

Here is the Apache tree:

Data: /var/www :

- /var/www/ : data root provided by Apache (DocumentRoot) ;
- /var/www/cgi-bin: CGI script directory;
- /var/www/error : http errors pages;
- /var/www/html : root;
- /var/www/admin : location of the web applications for management
- /var/www/icons : icons in the public domain available for your applications;
- /var/www/perl : perl scripts.

Apache logs : /var/log/httpd/ :

- /var/log/httpd/access_log : page access logs ;
- /var/log/httpd/error_log : error logs.

Apache executables:

- /usr/sbin : all the server's executables and tools;
- /etc/init.d/httpd : initscript for Apache.

Apache configuration files: /etc/httpd :

- /etc/httpd : Apache configuration root (ServerRoot) ;
- /etc/httpd/conf : contains all the basic configuration files of the server ;
- /etc/httpd/conf/fileprotector.conf : rules to protect critical files (eg. : php) ;
- /etc/httpd/conf/httpd.conf : server's main configuration file;
- /etc/httpd/conf/mime.types : MIME type configuration;
- /etc/httpd/conf/vhosts.d : directory containing the virtualhosts' configuration files;
- /etc/httpd/conf/webapps.d : configuration files for web applications;
- /etc/httpd/conf.d : contains links to Apache modules, library, log files, module's configuration files.

7.1.1.2. Configure an Apache server

7.1.1.2.1. Basic Configuration

In the server section, click on the Apache server icon. The welcome screen is divided in two parts, general configuration and virtual server configuration. virtualhost refers to the possibility of hosting many sites on the same Apache server. The global configuration applies to all virtual hosts.

The standard configuration provided often works without any modification. One of the most used features is virtualhosts. To activate them, make sure you have to following 2 lines in /etc/httpd/conf/httpd.conf :

```
NameVirtualHost *:80
     Include conf/vhosts.d/*.conf
```

You only have to declare them in a file indicating the DocumenRoot (where the files are physically) and the ServerName (to identify the virtual host). According to the default configuration provided with Mandriva Linux, you need to provide a file named *.conf placed in /etc/httpd/conf/vhosts.d.

If you are using virtual servers with HTTPS, you must provide only 1 server per IP address.

Let's illustrate our discussion with an example :

```
<VirtualHost *:80>
     DocumentRoot /var/www/html/test1
     ServerName www.test1.com
     </VirtualHost>

     <VirtualHost *:80>
     DocumentRoot /var/www/html/test2
     ServerName www.test2.org
     </VirtualHost>
```

Restart Apache and test.

To facilitate your administration, it is recommended that you plan a configuration file and a log file per virtualhost.

7.1.1.2.2. Apache management toolkit

The initscript `/etc/init.d/httpd` has options to manage the start, stop and information queries:

- Stopping the server:

  ```
  # service httpd stop
         Shutting down httpd:   [  OK  ]
  ```

- Starting the server :

  ```
  # service httpd start
         Starting httpd:     [  OK  ]
  ```

- Restarting the server :

  ```
  # service httpd restart
         Shutting down httpd:   [  OK  ]
         Starting httpd:     [  OK  ]
  ```

- Reload the configuration du serveur :

  ```
  # service httpd reload
         Reloading httpd:     [  OK  ]
  ```

- View Server status :

```
# service httpd status
Apache is running.
httpd: 12137 12136 12135 12134 12133 12132 12131 12130 12122
```

- Extended status :

```
# service httpd extendedstatus
                       Apache Server Status for localhost

   Server Version: Apache/2.2.9 (Mandriva Linux/PREFORK-12mdv2009.0)
          mod_ssl/2.2.9 OpenSSL/0.9.8h PHP/5.2.6 with Suhosin-Patch

   Server Built: Sep 20 2008 03:50:58
   __________________________________________________________________

   Current Time: Wednesday, 06-May-2009 11:42:39 CEST
   Restart Time: Wednesday, 06-May-2009 11:42:37 CEST
   Parent Server Generation: 0
   Server uptime: 1 second
   Total accesses: 0 - Total Traffic: 0 kB
   CPU Usage: u0 s0 cu0 cs0
   0 requests/sec - 0 B/second -
   1 requests currently being processed, 7 idle workers

W_______........................................................
................................................................
................................................................
................................................................

   Scoreboard Key:
   "_" Waiting for Connection, "S" Starting up, "R" Reading Request,
   "W" Sending Reply, "K" Keepalive (read), "D" DNS Lookup,
   "C" Closing connection, "L" Logging, "G" Gracefully finishing,
   "I" Idle cleanup of worker, "." Open slot with no current process

#######################################
#######################################
```

- Check Apache configuration d'Apache (`httpd.conf`) :

```
# service httpd configtest
checking apache configuration integrity :      [  OK  ]
```

- Check virtual host configuration :

```
# service httpd configtest_vhosts
checking apache configuration integrity :      [  OK  ]
```

- Starting the server in debug mode during pre-production tests :

  ```
  # service httpd debug
  Starting httpd (debug mode and in the foreground):
  [Fri Sep 01 05:51:53 2008] [notice] core dump file size limit raised
   to 4294967295 bytes
  [Fri Sep 01 05:51:53 2008] [info] mod_unique_id: using ip addr
   192.168.40.140
  [Fri Sep 01 05:51:54 2008] [notice] Digest: generating secret for
   digest authentication ...
  [Fri Sep 01 05:51:54 2008] [notice] Digest: done
  [Fri Sep 01 05:51:54 2008] [info] mod_unique_id: using ip addr
   192.168.40.140
  ```

  Detailed messages will appear in the console.

Apart from the commands suggested above, you can validate the proper operation of the server with `telnet` on port 80 or 443 :

- A working server:

  ```
  # telnet example.com 80
  Trying 192.168.40.140...
  Connected to example.com (192.168.40.140).
  Escape character is '^]'.
  ```

- A non-working server:

  ```
  # telnet example.com 80
  Trying 192.168.40.140...
  telnet: connect to address 192.168.40.140: Connection refused
  telnet: Unable to connect to remote host: Connection refused
  ```

7.1.1.3. Apache Advance configuration

7.1.1.3.1. https configuration

You must have the `apache-mod_ssl` rpm installed.

By default, a private key and a certificate are issued during the installation. To generate keys, you need `openssl`. Keys and certificates are kept, by default, in `/etc/pki/tls/private/localhost.key` and `/etc/pki/tls/certs/localhost.crt`.

Generating your own key

Run the following operations in the /etc/pki/tls/private/directory :

```
# openssl genrsa -des3  -out server.key 1024
```

Remove the password from the private key: :

```
# openssl rsa -in server.key -out server.pem
```

Generate the new certificate. To do this, enter the command below. You can also enter the default certificate value by editing this file /usr/lib/ssl/openssl.cnf.

```
# openssl req -new -key server.key -out server.csr
```

You can sign it yourself with the following, if you do not have a certificate authority :

```
# openssl x509 -req -days 60 -in server.csr

-signkey server.key -out server.crt
```

configuration files are :

- /etc/cron.daily/certwatch
- /etc/httpd/modules.d/40_mod_ssl.conf
- /etc/httpd/modules.d/41_mod_ssl.default-vhost.conf

You need to edit a section in the file /etc/httpd/modules.d/41_mod_ssl.default-vhost.conf to enable the certificate on part of the server.

```
<IfDefine HAVE_SSL> <IfModule !mod_ssl.c>
     LoadModule ssl_module modules/mod_ssl.so </IfModule>
     </IfDefine>

     <IfModule mod_ssl.c>
     NameVirtualHost 192.168.40.119:443
     <VirtualHost toto:443>
     ServerName toto
     SSLEngine on
     SSLCertificateFile /etc/pki/tls/certs/localhost.crt
     SSLCertificateKeyFile /etc/pki/tls/private/localhost.key
     DocumentRoot /var/www/html/vhost4
     ErrorLog logs/ssl_error_log
     <IfModule mod_log_config.c>
     TransferLog logs/ssl_access_log
     </IfModule>
     </VirtualHost>
```

```
    </IfModule>
```

Reload your server, you should now be able to access the URL `https://192.168.40.119` and accept the certificate for your machine.

7.1.1.4. Securing Apache

7.1.1.4.1. Securing the configuration de base

Securing Apache essentially implies modifying the base configuration provided during installation, mainly located in `/etc/httpd/conf/httpd.conf`.

First, delete the service banner, as it provides valuable information to guide attackers. In `/etc/httpd/conf/httpd.conf` :

```
ServerSignature Off
     ServerTokens Prod
```

To check,use command `HEAD` :

- with the banner :

  ```
  $ HEAD http://dhcp140
  200 OK
  Connection: close
  Date: Wed, 06 May 2009 09:43:52 GMT
  Accept-Ranges: bytes
  ETag: "1bdd3-d7-452ee8885b240"
  Server: Apache/2.2.9 (Mandriva Linux/PREFORK-12mdv2009.0)
  Content-Length: 215
  Content-Type: text/html
  Last-Modified: Sat, 26 Jul 2008 14:59:13 GMT
  Client-Date: Wed, 06 May 2009 09:43:52 GMT
  Client-Peer: 192.168.40.140:80
  Client-Response-Num: 1
  ```

- without the banner :

  ```
  $ HEAD http://dhcp140
  200 OK
  Connection: close
  Date: Wed, 06 May 2009 09:50:04 GMT
  Accept-Ranges: bytes
  ETag: "1bdd3-d7-452ee8885b240"
  Server: Apache
  Content-Length: 215
  Content-Type: text/html
  Last-Modified: Sat, 26 Jul 2008 14:59:13 GMT
  Client-Date: Wed, 06 May 2009 09:50:04 GMT
  ```

```
  Client-Peer: 192.168.40.140:80
  Client-Response-Num: 1
```

We also recommend deleting the name resolution: it creates unnecessary network traffic. In `/etc/httpd/conf/httpd.conf`:

```
HostnameLookups Off
```

Also, let's delete the possibility for users to publish web pages from their `/home` directories (`Apache mod_userdir` module). This module is no longer installed by default, but lets make sure it's not there:

```
# rpm -qa | grep apache-mod_userdir
```

Comment the following line in `/etc/httpd/conf/httpd.conf` :

```
# LoadModule userdir_module modules/mod_userdir.so
```

Forbid page index posting: in the directives definition `<Directory>`, disable `Indexes` as much as possible. Check for the presence of the `<Directory />` directive in `/etc/httpd/conf/httpd.conf`. It must include at least the following elements :

```
<Directory />
    Options -Indexes
    AllowOverride None
     </Directory>
```

7.1.1.4.2. Securing virtualhosts

If you use virtualhosts on this machine, create a "virtualhost catch-all". This way, if a visitor launches a request on the IP address instead of the "virtualhost" name, you will be able to control the page posted.

```
<VirtualHost _default_:*>
     DocumentRoot /var/www/html/defaut
     </VirtualHost
```

Remember to create the matching directory for the specified `DocumentRoot` and place an `index.html` file within this directory.

It is also very useful to have specific log files for each virtualhost to quickly identify future problems. This will give you the following log file tree (to create):

```
/var/log/httpd/
 |-virtualhost1
 | |-access_log
 | `-error_log
 |-virtualhost2
 | |-access_log
 | `-error_log
 ...
```

Below, you have an example of a vitrualhost configuration file using this log file separation:

```
# cat /etc/httpd/conf/vhosts.d/virtualhost1
     <VirtualHost 172.20.30.40>
     DocumentRoot /var/www/html/virtualhost1
     ServerName virtualhost1.domaine.com

     ErrorLog logs/virtualhost1/error_log
     CustomLog logs/virtualhost1/access_log combined
     </VirtualHost>
```

You also need to update the Apache log rotation system to take the new log files into account :

```
# cat /etc/logrotate.d/httpd
/var/log/httpd/*_log /var/log/httpd/virtualhost1/*_log
 /var/log/httpd/apache_runtime_status /var/log/httpd/ssl_mutex {
    rotate 5
    monthly
    missingok
    notifempty
    nocompress
    prerotate
        /etc/rc.d/init.d/httpd closelogs > /dev/null 2><1
    endscript
    postrotate
        /etc/rc.d/init.d/httpd closelogs > /dev/null 2><1
    endscript
}
```

7.1.1.5. Enabling PHP

7.1.1.5.1. Installing and configuring PHP

Only PHP5 is provided by Mandriva Enterprise Server 5. By default, PHP5 is compiled with the hardened-php patch that increases security, such as disabling unsecured code. More info is available at: the hardened PHP project site (`http://www.hardened-php.net/suhosin/`).

The main package to install is apache-mod_php which will modify the Apache configuration to take php into account and restart Apache. You will also have to satisfy some dependencies with some basic php modules required in most configurations.

To setup a "LAMP" environment, you'll also need to install the Apache module enabling SQL support for PHP: php-mysql

PHP configuration is mainly done within the /etc/php.ini file. Each feature is introduced in [bloc] form. Each variable is described as following:

```
variable = value
```

The default file is usually enough to have a working php environment. We will look at improving php security in the next chapter.

Mandriva Enterprise Server 5 contains the following difference: /etc/php.d. To simplify your php.ini file, all dynamic modules have their own dedicated file in /etc/php.d and are included in the global configuration. By default, we have the following files: :

```
# ls /etc/php.d
12_ctype.ini      22_ftp.ini        37_mysqli.ini     57_sysvsem.ini
13_curl.ini       23_gd.ini         42_pgsql.ini      58_sysvshm.ini
18_dom.ini        24_gettext.ini    43_posix.ini      60_tokenizer.ini
21_openssl.ini    26_iconv.ini      47_session.ini    62_xml.ini
21_zlib.ini       28_ldap.ini       54_hash.ini       62_xmlrpc.ini
63_xmlreader.ini  64_xmlwriter.ini  70_pdo.ini        78_sqlite.ini
81_filter.ini     82_json.ini       98_suhosin.ini
```

7.1.1.5.2. PHP: improve the basic configuration

We suggest a number of improvements to the default configuration. All of these should be accomplished within /etc/php.ini.

- Disable global variables if applications support it:

  ```
  register_globals = Off
  ```

- Remove php information from the banner: php availability will not appear in the server banner:

  ```
  expose_php = Off
  ```

- Disable the posting of error messages for php scripts as these may give vulnerability information:

  ```
  display_errors = Off
  ```

- Activate php logs to quickly debug scripts from the logs generated by the `syslogd` deamon:

  ```
  log_errors = Off
         error_log = syslog
  ```

- Disable file upload :

  ```
  file_uploads = Off
  ```

- Disable magic quotes :

  ```
  magic_quotes_gpc = Off
         magic_quotes_runtime = Off
         magic_quotes_sybase = Off
  ```

- Prevent external module loading :

  ```
  enable_dl = Off
  ```

- Forbid treating URLs as files: this parameter would allow you to download from another server :

  ```
  allow_url_fopen = Off
  ```

- Delete **.** from the library path: should be used with caution since some scripts may not work anymore :

  ```
  include_path = "/usr/lib/php/:/usr/share/pear/"
  ```

- Sandbox the execution of php scripts: limits the directories allowed to run php scripts. You can specify many directories, separated with commas :

  ```
  open_basedir = /var/www/html/appli
  ```

  With virtualhosts, this option is more significant if specified in the virtualhosts' configuration file, which refines the behavior of php :

  ```
  <VirtualHost 127.0.0.1>
   DocumentRoot /var/www/html/virtualhost1/html
   ServerName virtualhost1.domaine.com
   php_admin_value open_basedir /var/www/html/virtualhost1
         </VirtualHost>
  ```

- Specify unauthorized functions in scripts which can endanger the system security. You may specify many functions, separated by commas :

```
disable_functions = exec,system
```

To establish the list of functions :

```
$ lynx -dump http://fr.php.net/manual/fr/ref.filesystem.php |
grep 'function\.' | awk -F'.' '{ print $5 }'
```

7.1.2. Proxy setup

This section aims at establishing a caching proxy, mainly for clients using the HTTP protocol (see FTP). Main functions are :

- Bandwidth optimization, when clients request the same resource ;
- Access control and filtering (content, service authentication, schedule).

Squid is an Internet proxy server providing many functions.

7.1.2.1. Concepts

A proxy server acts as an intermediary between the client and the resource to reach. It optimizes the original request and controls its validity. Here, request optimization is done through the use of a local copy of the most requested Internet resources: the cache.

When a client requests a resource, the proxy checks if it has a recent copy of this resource. If this is the case, it returns the local copy, instead of letting the client access the distant resource. This reduces outside traffic accordingly.

SSquid is a service daemon listening on a determined HTTP/HTTPS/FTP port (3128 by default, 8080 is often used for proxy).

When a request is received, the proxy validates the following authorizations :

Is the machine that sends this request included in an IP range authorized for this service? Is the user specified (when authorization is used) allowed to use this service? Does the schedule allow you to access this service at this time?

If one of these conditions is not fulfilled, the proxy returns a significant error. If not, the procedure continues with the verification of the content on cache (in the case of a caching-proxy). One of the following behaviors is then possible :

- The requested resource is not in cache. Squid relays this request to through the web (and keeps a copy in cache) ;
- The requested resource is in cache. If this resource is not too old, this local copy will be returned to the client. (if not, Squid relays the request to through the web and keeps a copy in cache.

7.1.2.2. Installation and file tree

Installation is straight forward. You simply need to install the `squid` package.

Here are the main components of the tree structure :

- `/etc/squid` : contains squid's configuration files and `squid.conf` ;
- `/usr/lib/squid` : outils tools such as `cachemgr.cgi` and `squid_ldap_auth` ;
- `/var/log/squid` : contains squid server logs files ;
- `/var/spool/squid` : contains the server cache.

7.1.2.3. Squid server configuration Squid

Squid configuration is based on the creation of access control list (ACL) for the HTTP resource. This control is based on the filters concerning: source, destination, schedule, protocols (HTTP/HTTPS/FTP) and methods used (GET/POST).

7.1.2.3.1. ACL locations

These ACLs can establish access authorization as well as the resource caching. ACL configuration is located in `/etc/squid/squid.conf`.

Configuring the source

- Password authentication :

  ```
  acl aclname proxy_auth username ...  acl aclname
     proxy_auth_regex [-i] pattern ...
  ```

- IP address of the source :

  ```
  acl aclname src ip-address/netmask
    ... (clients IP address) acl aclname src addr1-addr2/netmask
    ... (range of addresses)
  ```

- Source domain :

  ```
  acl aclname srcdomain .foo.com ...
  # reverse lookup, client IP
  ```

- Number of simultaneous connections for a given client :

  ```
  acl aclname maxconn
    number acl aclname max_user_ip [-s]
    number
  ```

- Browser :

  ```
    acl aclname browser [-i] regexp ...
  # pattern match on
    User-Agent header
  ```

Configuring the destination

- Destination domain :

  ```
  acl aclname dst ip-address/netmask ... (URL
    host's IP address) acl aclname dstdomain .foo.com ...
  # Destination server from URL
  ```

- URL :

  ```
  acl aclname dstdom_regex [-i] xxx ...
  #  regex matching server acl aclname url_regex [-i] ^http://
     ...
  # regex matching on whole URL acl aclname urlpath_regex
     [-i] \.gif$ ...
  # regex matching on URL path
  ```

- Connection schedule :

  ```
  acl aclname time [day-abbrevs]
     [h1:m1-h2:m2]
  ```

- Protocols (HTTP/HTTPS/FTP) and methods (GET/POST/) :

  ```
  acl aclname proto HTTP FTP ...  acl aclname
     method GET POST ...
  ```

This list is not complete, but illustrates the most common ACLs.

7.1.2.3.2. Access configuration

By default, the server listens on port 3128 and no client/network is authorized to pass through the proxy.

```
http_port 3128

#Recommended minimum configuration:
acl all src 0.0.0.0/0.0.0.0
acl localhost src 127.0.0.1/255.255.255.255

http_access allow localhost

http_access deny all
```

Only the proxy (`acl localhost src 127.0.0.1/255.255.255.255`) is authorized to use the (`http_access allow localhost`). All other machines, no matter what their IP is, (`acl all src 0.0.0.0/0.0.0.0`) will have their request denied (`http_access deny all`).

To authorize the local network to use the proxy, it must be declared with an ACL which specifies the appropriate network :

```
acl MyNetworks src 10.0.0.0/24 192.168.0.0/24
```

We must authorize this ACL:

```
http_access allow MyNetworks
```

Which gives the following configuration :

```
acl all src 0.0.0.0/0.0.0.0
acl localhost src 127.0.0.1/255.255.255.255
acl MyNetworks src 10.0.0.0/24 192.168.0.0/24
...
http_access allow localhost
http_access allow MyNetworks
http_access deny all
```

We can consider modifying the “all” ACL instead of adding a new one, but that may cause a problem, since there is no default policy.

Here is a combination of ACLs on the IP and the schedule :

1. You must first define the network(s) concerned :

```
acl MyNetworks src 10.0.0.0/24
       192.168.0.0/24
```

2. Then, the connection schedule (monday to friday, 8 h 30 à 18 h 30) :

   ```
   acl WORKING time MTWHF
    08:30-18:30
   ```

3. Assemble the two ACLs so that machines from the specified network(s) may use the proxy during the defined schedule :

   ```
   http_access allow MyNetwork
    WORKING
   ```

Which gives us the following global configuration :

```
acl all src 0.0.0.0/0.0.0.0
acl localhost src 127.0.0.1/255.255.255.255
acl MyNetworks src 10.0.0.0/24 192.168.0.0/24
acl WORKING time MTWHF 08:30-18:30
...
http_access allow localhost
http_access allow MyNetworks WORKING
http_access deny all
```

7.1.2.3.3. Configuring the cache

The cache is the memory of Squid. Like access, it is possible to use ACLs to authorize or deny caching of certain objects. In general, we will not cache dynamic pages, such as CGI, and general requests with a "?" in the URL.

```
acl QUERY urlpath_regex cgi-bin \?
     no_cache deny QUERY
```

The cache works with a system of disk swap. Most recent objects are kept in memory, while older ones a kept on disk. Hence, you need to adjust the size of the memory cache (directive `cache_mem`, by default 8M) and on disk (`directive cache_dir`, 100M by default) according to your Internet load to reduce swapping.

You also need to establish minimum and maximum size of the objects to put in your cache. (`directives minimum_object_size` – default 0KB, `maximum_object_size` - default 4096KB, `maximum_object_size_in_memory` – default 8KB). To have an efficient cache, make sure you have enough RAM and fast disk drives (SCSI).

The FAQ states the following rule: you need 10MB of RAM for each GB of disk cache. Also, it is recommended that you have twice as much RAM as that which is required by Squid.

7.1.2.4. Access Authentication using LDAP

You can set up authentication to validate that a user, identified by his login and password, has the rights (or not) to surf the Web. This authentication can be based on by many user management structure: LDAP, NCSA (.htpasswd), MSNT/SMB/winbind/NTLM, PAM, getpwam (based on /etc/passwd), sasl, Digest. We will use LDAP.

Authentication is validated by an external programm (`squid_ldap_auth`) which returns true or false for login and password match.

Therefore, you need to create an ACL that calls authentication, (`acl ACLName proxy_auth REQUIRED`), and associate it to the ACL which defines request sources (in this case, external networkds `acl ACLName src 10.0.0.0/24`).

```
auth_param basic children 5
auth_param basic program /usr/lib/squid/squid_ldap_auth -v 3 -b
 ou=Users,dc=example,dc=com localhost
auth_param basic realm Example.com Squid Server
auth_param basic credentialsttl 2 hours
...
acl password proxy_auth REQUIRED
...
acl all src 0.0.0.0/0.0.0.0
acl localhost src 127.0.0.1/255.255.255.255
acl MyNetworks src 10.0.0.0/24 192.168.0.0/24
...
http_access allow localhost
http_access allow MyNetworks password
http_access deny all
```

The `squid` package provides the `squidclient` command. This allows you to quickly test the proper functioning of your server and returns the entire dialog between client and server on the shell.

```
squidclient http://localhost
 HTTP/1.0 200 OK
 Date: Fri, 01 Sep 2008 10:11:24 GMT
 Server: Apache
 Last-Modified: Fri, 28 Jul 2008 10:57:38 GMT
 ETag: "4355-2c-2e3d2c80"
 Accept-Ranges: bytes
```

```
Content-Length: 44
Content-Type: text/html
X-Cache: MISS from unconfigured
Via: 1.0 unconfigured:3128 (squid/2.6.STABLE1)
Proxy-Connection: close
<html><body><h1>It works!</h1></body></html>
```

7.2. Identity service

The Identity service is able to provide authentication and authorization services to users and services in a network. This can be implemented via many different protocols, usually depending on the type of client. Available are:

- Standard local unix authentication
- LDAP
- Kerberos
- Samba

Mandriva Enterprise Server 5 suggests a strong emphasis on LDAP, trying to use it for as many services as possible. To do this we offer a package called `openldap-mandriva-dit`. This package installs a basic LDAP tree which is ready to host the following services:

- Authentication: UNIX®, Samba and Kerberos (using the Heimdal implementation)
- Authorization: UNIX®, Samba, Kerberos and the new Password Policy overlay module from OpenLDAP
- DNS
- Sudo
- DHCP

In this chapter, we will focus on authentication and authorization using OpenLDAP and, later, Kerberos. For details on how to use DNS, DHCP or `sudo` with LDAP, please consult the chapter about each of these services.

After studying this chapter, you should be able to:

- integrate UNIX®, Samba and Kerberos authentication in LDAP
- use password policies
- delegate administrative privileges for services in LDAP to users and groups
- troubleshoot common LDAP issues

7.2.1. General Concepts

Here we will explain some general concepts about LDAP, user and group accounts.

7.2.1.1. Directories

Directory servers can be viewed as databases that store information. LDAP is the protocol used to access these servers:

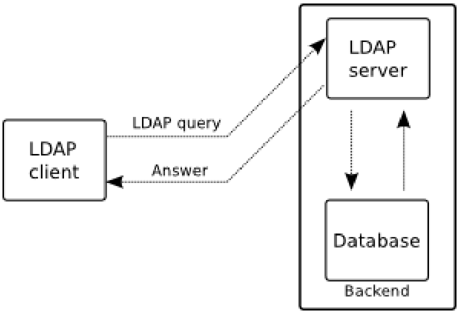

Figure 7-1. LDAP Protocol

The type of database varies a lot between implementations. It could be a standard relational database, a more specialized database such as Berlekey DB, a script that dynamically creates the output, etc.

LDAP tends, however, to be mostly compared to relational databases (SQL). There are some important differences:

- Hierarchy: instead of using tables with records and fields, it uses trees and notes, similar to a filesystem
- Optimization for reading: directories usually have many more read accesses than write attempts
- Distributed: the tree structure allows for branches to be stored elsewhere without compromizing the view of the whole

- Strong standardization: the type of data stored in a directory follows a strong standardization both in the name of the data as well regarding its type. This is usually RFC based.

Here is a small example of a tree that spans more than one server:

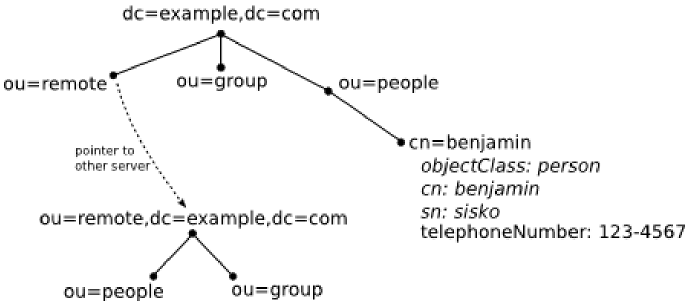

Figure 7-2. LDAP tree

Authentication and authorization services are usually good candidates for using a directory because they need to have a high read performance, they suffer fewer write operations and can be distributed in an organization while still remaining connected.

The set of rules about the data stored in a directory is called schema. The directory schema, among others things, defines:

- object classes: they define what an entry is about and what kind of information it can hold
- attributes: the data itself, similar to a field in standard databases
- indexing and ordering rules: specify how searching a specific attribute should work and how it should be ordered
- type of value: defines what kind of data the attribute should hold (numeric, string, binary blob, etc)

The schema is loved and hated at the same time. It is good for standardization, but it sometimes can make it difficult to add data to a directory. Whereas in a database this is as simple as just adding a new field, in a directory you have to obey the schema. You cannot just add any type of attribute to an entry: only the ones allowed by the object classes being used.

For example, this is the definition of the person object class:

```
objectclass ( 2.5.6.6 NAME 'person'
    DESC 'RFC2256: a person'
    SUP top STRUCTURAL
    MUST ( sn $ cn )
    MAY ( userPassword $ telephoneNumber $ seeAlso $ description ) )
```

This tells us that:

- This object is defined in RFC 2256.
- It is a structural object.
- The sn and cn attributes are mandatory
- The userPassword, telephoneNumber, seeAlso and description are optional attributes to have in an entry using this object class.

Figure 7-2 also illustrates the mandatory and some optional attributes of the `person` object class.

7.2.1.2. Tree Layout, Groups and Privileges

The `openldap-mandriva-dit` package offers a tree layout that can host the described services as well as authentication and authorization. This is what the tree looks like:

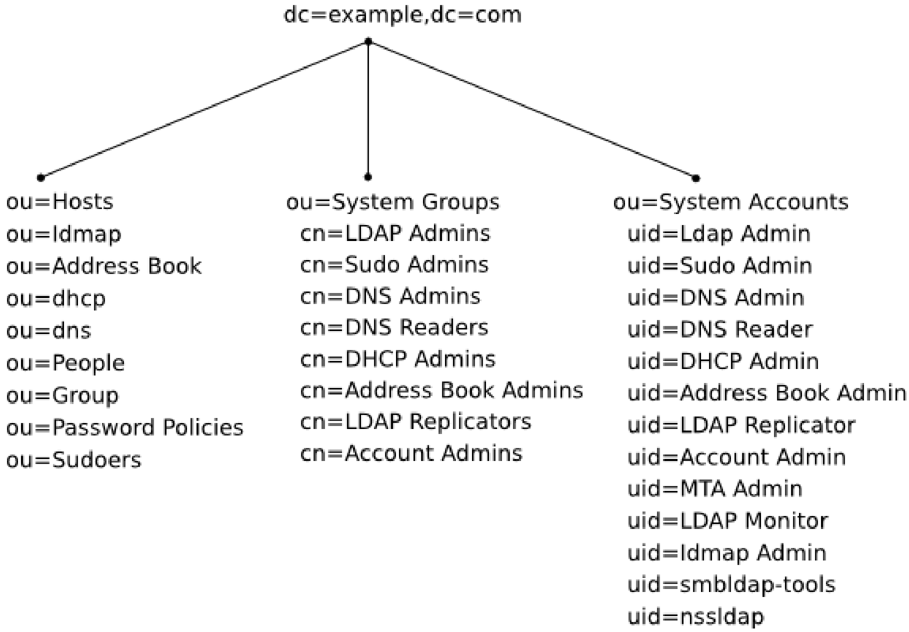

Figure 7-3. Tree Layout

The directory top entry, `dc=example,dc=com` is substituted at installation time with the DNS domain that was detected.

Each group has a default set of privileges assigned to its members. The sole initial member is the respective group administrator in the `SystemAccounts` branch. The privileges are such that each member can administer the respective service branch. So, for example, members of the `SudoAdmins` group can read and write to the `ou=Sudoers` branch.

The other groups not directly associated with a service are described below:

- LDAP Admins: can read and write to any part of the tree. They are also exempt from size and time limits.
- DNS Readers: only members of this group can read from the `ou=DNS` branch.
- LDAP Replicators: members can read any part of the directory, and are also not subject to size and time limits.
- Account Admins: members can create new accounts (unix, samba or kerberos) under the `ou=People`, `ou=Group` and `ou=Hosts` branches.
- MTA Admins: members can write to some specific email related attributes:
  - all attributes used by the `inetLocalMailRecipient` object class
  - the `mail` attribute
- LDAP Monitors: members can read the special `cn=Monitor` tree, which hosts live statistics data about the OpenLDAP server.

There are two extra default accounts supplied with this tree: `smbldap-tools` and `nssldap`. The first one is intended to be used by the `smbldap-tools` suite of programs and is a member of the `AccountAdmins` group. `nssldap` is a generic read account and is not used by default by any service, being provided just for convenience for the administrator who whishes to use it in such a scenario. There is no special ACL in place for this account.

7.2.1.3. Installation

The `openldap-mandriva-dit` package contains an installation script that can be run at the console. The script is installed at `/usr/share/openldap/scripts/mandriva-dit-setup.sh`. This script performs these changes:

- asks the DNS domain (suggesting whatever was auto-detected)
- constructs the top-level directory entry from this domain using `dc` style attributes

- creates and imports an ldif file with the accounts and groups described here
- installs new `slapd.conf` and `mandriva-dit-access.conf` files (making backups of the previous ones) with the default ACLs and other useful configurations (like cache)
- loads the ldif file, backing up the previous database directory

Even though the script performs many tests and backups many files before overwriting them, administrators are advised to backup all data before running it.

This is a sample run using the `dc=example,dc=com` domain:

```
# /usr/share/openldap/scripts/mandriva-dit-setup.sh
Please enter your DNS domain name [mycompany.com]:
example.com


Administrator account

The administrator account for this directory is
uid=LDAP Admin,ou=System Accounts,dc=example,dc=com

Please choose a password for this account:
New password: secretpass
Re-enter new password: secretpass


Summary
=======

Domain:      example.com
LDAP suffix: dc=example,dc=com

Confirm? (Y/n)            Y
config file testing succeeded
Stopping ldap service
Finished, starting ldap service
Running /usr/bin/db_recover on /var/lib/ldap
removing /var/lib/ldap/alock
Starting slapd (ldap + ldaps):               [  OK  ]

Your previous database directory has been backed up as
/var/lib/ldap.1145397294
```

Now the ldap service is up and running. You can use the administrator account with the password we just set up to further populate the tree. This is the only enabled account in the tree so far: all the others are disabled. In order to enable an administrative account it must have a password. For example, let's enable the `DNSAdmin` account:

```
$ ldappasswd -x -D "uid=LDAP Admin,ou=System Accounts,dc=example,dc=com"
```

```
        -W -S "uid=DNS Admin,ou=System Accounts,dc=example,dc=com"
New password: newsecret
Re-enter new password: newsecret
Enter LDAP Password: here is the password for the bind (-D) user:
        LDAP Admin
Result: Success (0)
```

Now the DNSAdmin account has a password and can be used to administer the ou=dns branch of the tree.

7.2.2. Unix Users Authentication

This tree has nothing special regarding posix accounts (traditional Unix accounts). We have the usual:

- ou=People: branch which holds persons accounts (excepting kerberos). So, a user account could be something like uid=john,ou=People,dc=example, dc=com. These entries needs to have at least the posixAccount object class.
- ou=Group: group accounts go into this branch. For example, one would have cn=marketing,ou=Group,dc=example,dc=com. These entries need to have at least the posixGroup object class.

The group that can write to these branches is cn=AccountAdmins,ou=SystemGroups,dc=example,dc=com and the uid=AccountAdmin,ou=SystemAccounts, dc=example,dc=com is its owner and initial member.

By default, initially all system accounts except LDAPAdmin are blocked. To use one, you need to set a password for it. The example below enables the AccountAdmin account by assigning a password:

```
$ ldappasswd -x -D "uid=LDAP Admin,ou=System Accounts,dc=example,dc=com"
        -W -S "uid=Account Admin,ou=System Accounts,dc=example,dc=com"
New password: password for account admin
Re-enter new password: retype it
Enter LDAP Password:
Result: Success (0)
```

This uses the LDAPAdmin account to set a password for AccountAdmin, which is now enabled and can be used to add/remove/modify users and groups.

7.2.2.1. Creating an Account

There are several LDAP frontends available that can create posix accounts. We will show a quick example using the Luma program. It has several plugins and one of those plugins is about user management:

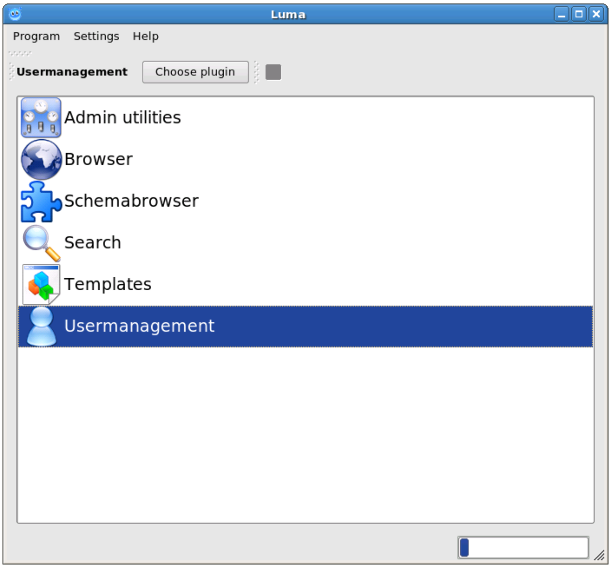

Figure 7-4. Luma plugins

This is the user management plugin:

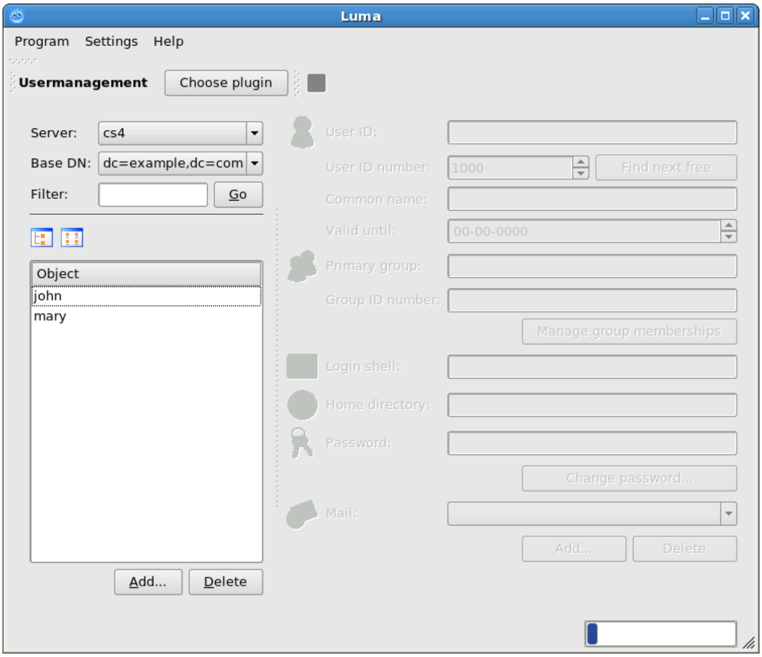

Server: cs4

Base DN: dc=example,dc=com

User ID number: 1000

Valid until: 00-00-0000

Object
john
mary

Figure 7-5. User management plugin

When adding a new user, first we have to select on which server and which branch. Here we go to the `ou=People` branch:

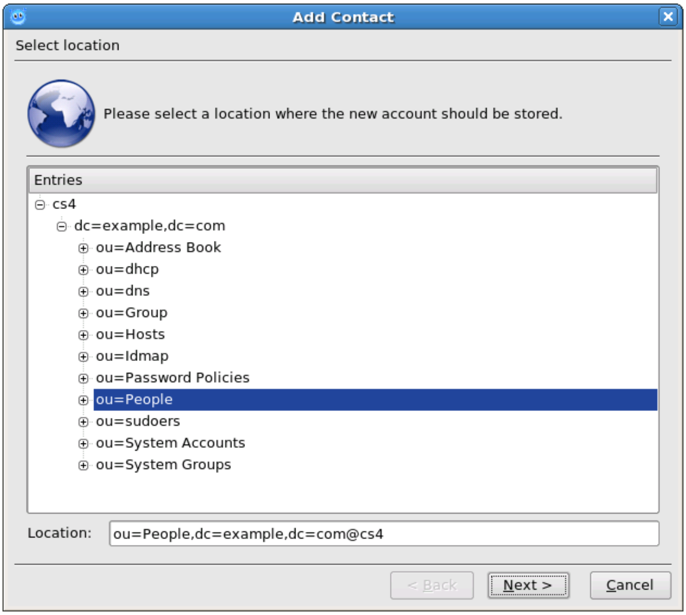

Select location

Please select a location where the new account should be stored.

Entries

```
⊖ cs4
  ⊖ dc=example,dc=com
    ⊕ ou=Address Book
    ⊕ ou=dhcp
    ⊕ ou=dns
    ⊕ ou=Group
    ⊕ ou=Hosts
    ⊕ ou=Idmap
    ⊕ ou=Password Policies
    ⊕ ou=People
    ⊕ ou=sudoers
    ⊕ ou=System Accounts
    ⊕ ou=System Groups
```

Location: ou=People,dc=example,dc=com@cs4

Figure 7-6. Where to add an user

Next we fill in the necessary information for this user:

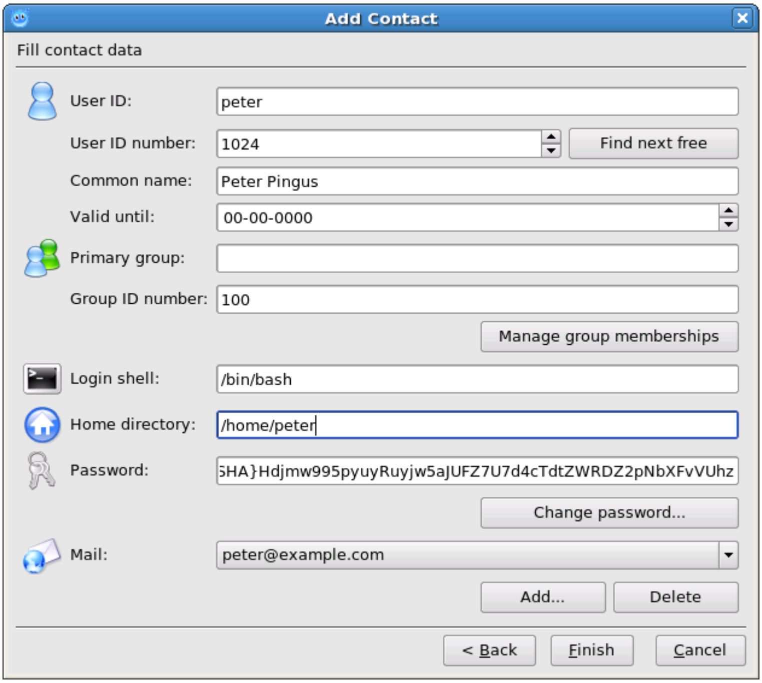

Fill contact data

User ID: peter

User ID number: 1024

Common name: Peter Pingus

Valid until: 00-00-0000

Primary group:

Group ID number: 100

Login shell: /bin/bash

Home directory: /home/peter

Password: SHA}Hdjmw995pyuyRuyjw5aJUFZ7U7d4cTdtZWRDZ2pNbXFvVUhz

Mail: peter@example.com

Figure 7-7. Adding an user

It is necessary to specify at least a primary group for this user. LDAP posix groups will be automatically shown, but if you don't have any then the list will be empty. Here we selected `gidNumber=100`, which belongs to the local `users` group:

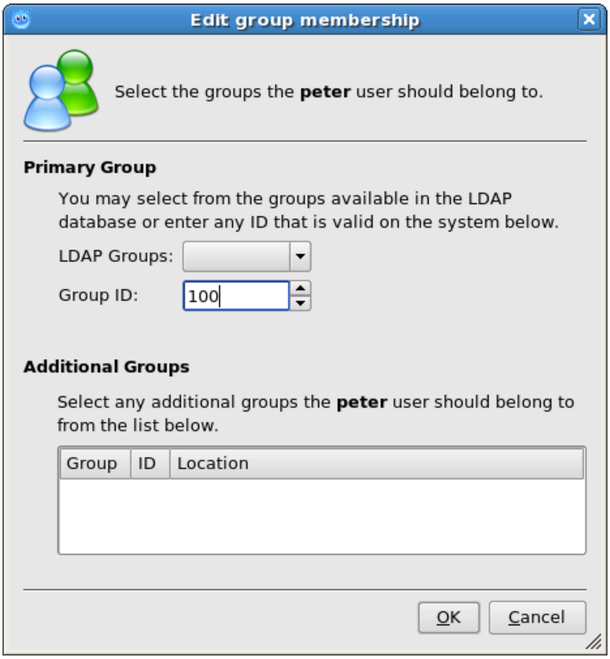

Edit group membership

Select the groups the **peter** user should belong to.

Primary Group

You may select from the groups available in the LDAP database or enter any ID that is valid on the system below.

LDAP Groups:

Group ID: 100

Additional Groups

Select any additional groups the **peter** user should belong to from the list below.

Group	ID	Location

Figure 7-8. Groups

This is how our newly created user looks in LDAP:

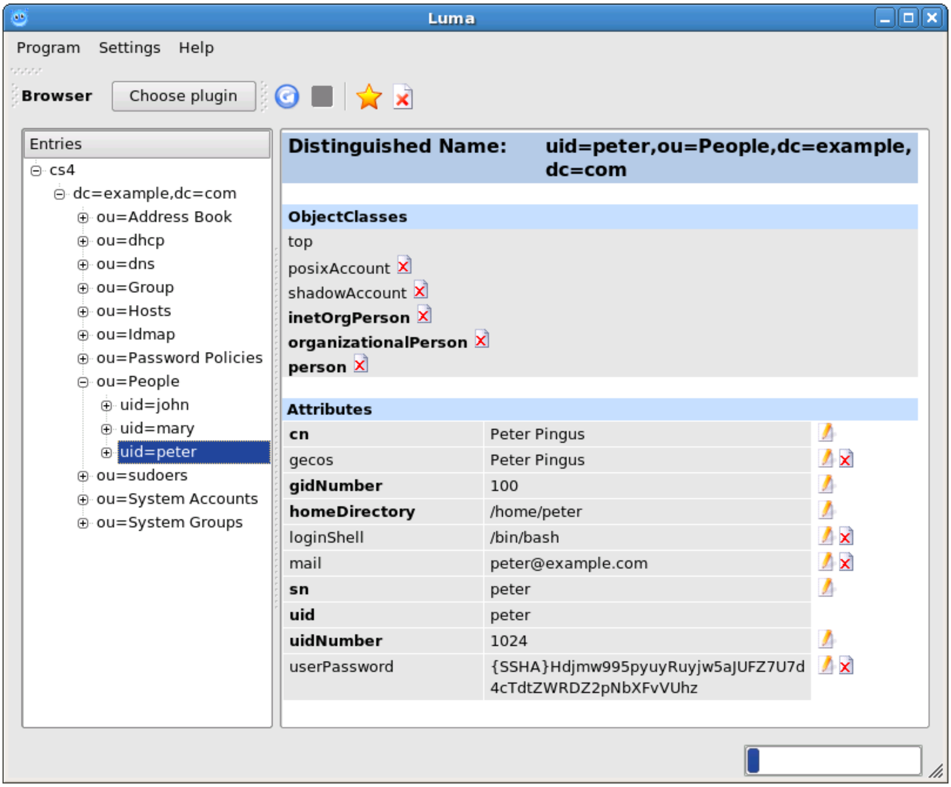

Distinguished Name: **uid=peter,ou=People,dc=example,dc=com**

ObjectClasses

top
posixAccount
shadowAccount
inetOrgPerson
organizationalPerson
person

Attributes

cn	Peter Pingus
gecos	Peter Pingus
gidNumber	100
homeDirectory	/home/peter
loginShell	/bin/bash
mail	peter@example.com
sn	peter
uid	peter
uidNumber	1024
userPassword	{SSHA}Hdjmw995pyuyRuyjw5aJUFZ7U7d4cTdtZWRDZ2pNbXFvVUhz

Figure 7-9. Peter in LDAP

Luma will try to find a free `uidNumber` to allocate for this user. You should always be careful, though, as having two or more users with the same `uidNumber` is wrong.

7.2.2.2. Creating Groups

Unix groups are represented in LDAP using the `posixGroup` class. A typical entry for a group looks like this:

```
dn: cn=ldapusers,ou=Group,dc=example,dc=com
    objectClass: posixGroup
    gidNumber: 1024
    cn: ldapusers
    memberUid: peter
    memberUid: queen
```

This entry defines a group called ldapusers with an identification number of 1024 and, so far, two members called peter and queen.

This structure is simple enough to be created manually. One should only take care to not use the same identification number for more than one group.

System Groups use a different object class: groupOfNames. The main difference is that membership is defined by a full DN instead of just a name as is the case with posixGroup. Unfortunately posixGroup and groupOfUniqueNames cannot be used at the same time because they are both structural classes. It is expected that the revision of RFC2307 will redefine posixGroup as an auxiliary class.

7.2.2.3. Delegating administrative privileges

All members of the cn=AccountAdmins,ou=SystemGroups,dc=example,dc=com group can manage user accounts under ou=People, ou=Group and ou=Hosts (for Samba, which is explained later). Any LDAP Admin, or the Account Admin user himself, can add members to this group.

System Groups have owners, and by default the owner can always edit membership of the group he or she owns. To see who is the owner, check the owner attribute of the System Group in question.

For example, let's add the user PeterPingus, who we just created, to this privileged group so that he can manage accounts:

```
$ ldapmodify -x -D 'uid=Account Admin,ou=System Accounts,
        dc=example,dc=com' -W
Enter LDAP Password: secretpassword
dn: cn=Account Admins,ou=System Groups,dc=example,dc=com
changetype: modify
add: member
member: uid=peter,ou=People,dc=example,dc=com

^D
modifying entry "cn=Account Admins,ou=System Groups,dc=example,dc=com"
```

The same operation can be done with a graphical client, of course. Just add the member attribute pointing to the user dn who you want to add to this group.

7.2.3. Samba Authentication

To use this DIT with Samba, follow these configuration details:

- Layout in LDAP

  This is the layout that has to be configured in `/etc/samba/smb.conf` and `/etc/smbldap-tools/smbldap.conf`:

  - machine accounts: under `ou=Hosts`
  - user accounts: under `ou=People`
  - group accounts: under `ou=Group`
  - idmap branch: under `ou=Idmap`

- `ldapadmindn`

  When it comes to the `ldapadmindn` `/etc/samba/smb.conf` configuration parameter, use a member of the Account Admins group. For example:

  ```
  ldap admin dn = uid=Account Admin,ou=System Accounts,dc=example,dc=com
  ```

- `smbldap-tools`

  In `/etc/smbldap-tools/smbldap_bind.conf`, use the `smbldap-tools` user instead of the directory's rootdn:

  ```
  masterDN="uid=smbldap-tools,ou=System Accounts,dc=example,dc=com"
  ```

  This user is a member of the Account Admins group. If you want to use another account, then make sure it's a member of this same group or else the default OpenLDAP ACLs won't work.

- `smbldap-populate`

  The default smbldap-populate behaviour, at least with version 0.9.2, is to create an administrator account with the following attributes:

  - uidNumber = 0
  - gidNumber = 0
  - name: root
  - member of Domain Admins

  This means that a root user is created in LDAP. We advise against that and suggest you use this command line with smbldap-populate:

  ```
  # smbldap-populate -a Administrator -k 1000 -m 512
  ```

This will create a user with the name Administrator, uidNumber 1000 and gidNumber 512. You can also use uidNumber 500 if you want to match windows' RID for this kind of user, but you may already have a local user with this number.

Later on, the Domain Admins group could be given privileges (see "net rights grant" command), or your shared drives could use the admin users parameter.

- IDMAP

  If using IDMAP's LDAP backend in a member server, set the ldap admin dn configuration parameter in `/etc/samba/smb.conf` to the dn of a member of the Idmap Admins group. For example:

  ```
  ldap admin dn = uid=Idmap Admin,ou=System Accounts,dc=example,dc=com
  ```

  In member servers, there is no need to use the full blown Account Admin user: the Idmap Admins group is the right one as it can only write to the ou=Idmap container.

There is a potential security vulnerability with Idmap in LDAP. Since all domain machines need to have write access to this branch of the directory (and thus need a clear text password stored somewhere), a malicious user with root privileges on such a machine could obtain this password and create any identity mapping in `ou=Idmap`. See this thread (`http://lists.samba.org/archive/samba/2006-March/119196.html`) for more information.

7.2.3.1. Creating samba accounts

The recommended way to use Samba accounts in LDAP is to use the `smbldap-tools` package. This package has several tools for adding, removing and modifying users and groups in an LDAP tree together with the Samba attributes. It can even be used without the Samba attributes, dealing only with the posix ones.

7.2.4. Kerberos Authentication

OpenLDAP can be used as a backend for Heimdal's database, meaning principal accounts can be stored in LDAP. This section will present the steps needed to integrate Heimdal's LDAP backend with OpenLDAP and `openldap-mandriva-dit`. If you don't need to have this funcionality, you can stick with the default Kerberos server which uses the MIT packages and skip this section.

We will start with a new realm which we will call `EXAMPLE.COM`. The rest of this text assumes that `openldap-mandriva-dit` is installed and that the supplied installation script was executed either manually or via Mandriva Server Setup.

When using the LDAP backend, it's advisable to have a script to create users, because Heimdal by default will use the `account` structural object class. Since it's more common to use `inetOrgPerson` (or a derived class), the principal entry would need to be removed and re-added later with `inetOrgPerson`.

Another approach would be to first create the user with whatever means are standard (`smbldap-tools`, manual script, a template in gq or luma, etc.) and then add the kerberos attributes later. We will document both approaches here.

7.2.4.1. Packages

Due to conflicts with MIT's Kerberos packages, Heimdal is packaged as follows in MES5:

- heimdal-libs
- heimdal-server
- heimdal-workstation
- heimdal-devel

Conflicts have been resolved where needed. Only heimdal-libs can be installed concurrently with MIT's libraries.

7.2.4.2. Overview of the changes

Here is a quick overview of the changes needed so that Heimdal can use OpenLDAP as its database backend, as well as use the `openldap-mandriva-dit` DIT:

- configure Heimdal to use LDAP for its backend
- configure OpenLDAP to accept connections from Heimdal via ldapi:// (`ldapi://`)
- test this mapping
- initialize the database
- managing user accounts

7.2.4.3. Heimdal with OpenLDAP

In order to have a database in LDAP, the following `[kdc]` section has to be used in Heimdal's `/etc/krb5.conf`:

```
[kdc]
database = {
    dbname = ldap:ou=People,dc=example,dc=com
    mkey_file = /var/heimdal/mkey
    acl_file = /var/heimdal/kadmind.acl
    }
```

This will instruct Heimdal to use the OpenLDAP server installed on the same host and to use the `ou=People` branch for its principals. The access method Heimdal uses is ldapi:// (`ldapi://`), which is a unix socket on the local filesystem, and authentication is handled by SASL EXTERNAL which we will configure in a moment.

7.2.4.4. Using ldapi:// (`ldapi://`)

OpenLDAP needs to be configured to accept conections via ldapi:// (`ldapi://`), a local unix socket. This is done in the `/etc/sysconfig/ldap` file. Change the SLAPD URL list to the following:

```
    # SLAPD URL list
    SLAPDURLLIST="ldap:/// ldaps:/// ldapi:///"
```

OpenLDAP will need to be restarted, of course.

7.2.4.5. Using SASL EXTERNAL

Heimdal uses SASL EXTERNAL to authenticate itself to the OpenLDAP server when connecting via the ldapi:// (`ldapi://`) socket. When doing this, the binding dn becomes:

```
# ldapwhoami -Y EXTERNAL -H ldapi:///var/run/ldap/ldapi
SASL/EXTERNAL authentication started
SASL username: gidNumber=0+uidNumber=0,cn=peercred,cn=external,cn=auth
SASL SSF: 0
dn:gidNumber=0+uidNumber=0,cn=peercred,cn=external,cn=auth
Result: Success (0)
```

We are going to map this dn to a more meaningful binddn via `authz-regexp`. The `slapd.conf` file provided with `openldap-mandriva-dit` already does this, but here it is for completeness:

```
(...)
ppolicy_default "cn=default,ou=Password Policies,dc=example,dc=com"

authz-regexp "gidNumber=0\\\+uidNumber=0,cn=peercred,cn=external,cn=auth"
"uid=Account Admin,ou=System Accounts,dc=example,dc=com"
authz-regexp ^uid=([^,]+),cn=[^,]+,cn=auth$ uid=$1,ou=People,dc=example,dc=com
```

With this change, and after restarting OpenLDAP, `ldapwhoami` now says we are an Account Admin:

```
# ldapwhoami -Y EXTERNAL -H ldapi:///var/run/ldap/ldapi
SASL/EXTERNAL authentication started
SASL username: gidNumber=0+uidNumber=0,cn=peercred,cn=external,cn=auth
SASL SSF: 0
dn:uid=account admin,ou=system accounts,dc=example,dc=com
Result: Success (0)
```

Note that any process connecting to the ldapi:// (`ldapi://`) socket as root (uid=0, gid=0) will be treated as an Account Administrator after this change!

7.2.4.6. Initializing the realm

We can now initialize the Kerberos realm. After OpenLDAP has been restarted, run the following:

```
# kadmin -l
kadmin> init EXAMPLE.COM
Realm max ticket life [unlimited]:7d
Realm max renewable ticket life [unlimited]:7d
kadmin
```

This will create some default principals under `ou=People`:

```
ou=People
krb5PrincipalName=krbtgt/EXAMPLE.COM@EXAMPLE.COM,ou=People,dc=example,dc=com
krb5PrincipalName=kadmin/changepw@EXAMPLE.COM,ou=People,dc=example,dc=com
krb5PrincipalName=kadmin/admin@EXAMPLE.COM,ou=People,dc=example,dc=com
krb5PrincipalName=changepw/kerberos@EXAMPLE.COM,ou=People,dc=example,dc=co
krb5PrincipalName=kadmin/hprop@EXAMPLE.COM,ou=People,dc=example,dc=com
krb5PrincipalName=default@EXAMPLE.COM,ou=People,dc=example,dc=com
```

7.2.4.7. Managing user and principal accounts

The Heimdal schema allows for principal accounts to be stored in a separate branch to the user accounts. For example, one could have the principal accounts under ou=KerberosPrincipals and user accounts under `ou=People`.

This has the obvious disadvantage of creating a problem with user management: when an user is removed, for example, the corresponding principal account has to be removed also. In other words, we would need a pointer in the user entry for the principal account (the seeAlso attribute is commonly used for things like this). And a script would need to follow this attribute and delete the principal account.

The advantage would be that one user could be associated with many kerberos principals using the seeAlso attribute, like john@REALM and john/admin@REALM.

The biggest disadvantage of this scheme where principals are separated from users is integration with Samba and Ldap simple binds: it's lost. Heimdal will only update the samba password hash if it's stored in the same entry. The same with userPassword: with OpenLDAP using the `smbk5pwdi` module (built with kerberos support), simple binds will only be able to use the kerberos password if everything is in the same entry.

Another option would be to store the principal keys and related attributes right under the user entry. We can do this because the kerberos object classes are auxiliary. So, user John would be, for example, uid=john,ou=people,dc=example,dc=com and the kerberos keys would be stored in this same entry. When this user is removed, so is the principal account. The drawback is that one user can only have one principal, and not several as in the previous case (where john could have john@REALM and john/admin@REALM associated with the same uid=john,ou=people,dc=example,dc=com entry).

But one issue comes up: what do we use to create this user in the first place? If we use kadmin, then it will create an entry of the form `krb5PrincipalName=john@EXAMPLE.COM,ou=People,dc=example,dc=com` with account being the structural object class. Since we tend to use a class derived from person as the structural class (such as inetOrgPerson), there is a conflict. If we use kadmin, we would have to remove the entry and re-add it with inetOrgPerson (and its mandatory attributes).

We can change the structural class that Heimdal will use, but it doesn't add the mandatory attributes so we can't just switch to inetOrgPerson in Heimdal's configuration: it will not work.

A better option would be to first create the user with another tool, such as smbldap or another script, and later add the kerberos attributes. The main advantages are:

- RDN naming will remain consistent with the rest of the entries (no krb5PrincipalName in the RDN if we don't want it)
- structural object class as we want it (for example, inetOrgPerson)
- user and principal accounts together under ou=People

The biggest disadvantage is that the mapping between users and principals would be 1:1, that is, one user could have at most only one kerberos principal associated with its entry.

Both schemes can be used together, however. It's actually more a question about how the accounts will be managed. So, regular users could have their kerberos keys stored in the user's entry, while administration and service keys would be stored under the same branch, but have no user associated with them. It's not very consistent with the tree (after all, ou=People was meant to host actual persons), but it works.

We will now give examples of two possibilities: using kadmin directly and using another script to first create the user account and then add kerberos attributes.

7.2.4.8. Using kadmin directly

We will create a kerberos account for the user "john" using kadmin directly. We don't even have to start Heimdal at this stage because we will be using kadmin in local mode:

```
# kadmin -l
kadmin> add john
Max ticket life [1 day]:10h
Max renewable life [1 week]:1w
Principal expiration time [never]:
Password expiration time [never]:
Attributes []:
john@EXAMPLE.COM's Password: somesecretpassword
Verifying - john@EXAMPLE.COM's Password: somesecretpassword
kadmin>
```

This creates the following entry:

```
dn: krb5PrincipalName=john@EXAMPLE.COM,ou=People,dc=example,dc=com
objectClass: top
objectClass: account
objectClass: krb5Principal
objectClass: krb5KDCEntry
krb5PrincipalName: john@EXAMPLE.COM
uid: john
krb5KeyVersionNumber: 0
krb5MaxLife: 36000
krb5MaxRenew: 604800
```

```
krb5KDCFlags: 126
(...)
```

We can obtain a ticket for this user:

```
# service heimdal start
Starting kdc:                                                  [  OK  ]
# kinit john
john@EXAMPLE.COM's Password: somesecretpassword
# klist
Credentials cache: FILE:/tmp/krb5cc_0
Principal: john@EXAMPLE.COM

Issued           Expires          Principal
Jun 20 15:43:23  Jun 20 22:23:23  krbtgt/EXAMPLE.COM@EXAMPLE.COM
#
```

Notice, however that this `john` user doesn't have the necessary posix attributes to become a system user. We will need something else to create this posix user anyway: Heimdal's role here is over.

7.2.4.9. Adding kerberos attributes to an existing user entry

If the user account already exists in the directory, then all we need to do is add the necessary Heimdal object classes to this account. Being auxiliary, this makes perfect sense.

So, for this example, we will use a pre-configured `smbldap-tools` package to create a sample user and then add the kerberos classes and attributes to it, but any posix user that already exists would work.

Notice we don't add the samba attributes just yet:

```
# smbldap-useradd mary
# getent passwd mary
mary:x:1001:513:System User:/home/mary:/bin/bash
```

The user looks like this in the directory:

```
dn: uid=mary,ou=People,dc=example,dc=com
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: posixAccount
objectClass: shadowAccount
cn: mary
sn: mary
givenName: mary
```

```
    uid: mary
    uidNumber: 1001
    gidNumber: 513
    homeDirectory: /home/mary
    loginShell: /bin/bash
    gecos: System User
    userPassword: {crypt}x
```

We will use the following LDAP modification to add the kerberos attributes and classes to this user:

```
$ ldapmodify -x -D 'uid=Account Admin,ou=System Accounts,dc=example,dc=com'
    -W
Enter LDAP Password: somepassword
     dn: uid=mary,ou=people,dc=example,dc=com
     changetype: modify
     add: objectClass
     objectClass: krb5Principal
     objectClass: krb5KDCEntry
     -
     add: krb5PrincipalName
     krb5PrincipalName: mary@EXAMPLE.COM
     -
     add: krb5KDCFlags
     krb5KDCFlags: 126
     -
     add: krb5KeyVersionNumber
     krb5KeyVersionNumber: 0

    modifying entry "uid=mary,ou=people,dc=example,dc=com"
```

Now `mary` is recognized as a kerberos principal and we can have Heimdal add the keys and other missing attributes by just invoking the password change command as an administrator or in local admin mode:

```
    # kadmin -l
    kadmin> passwd mary
    mary@EXAMPLE.COM's Password: somesecretpass
    Verifying - mary@EXAMPLE.COM's Password:
    somesecretpass  kadmin>
```

This adds the remaining attributes and now "mary" is a full kerberos principal:

```
    $ kinit mary
    mary@EXAMPLE.COM's Password: somesecretpass
    $ klist
    Credentials cache: FILE:/tmp/krb5cc_500
    Principal: mary@EXAMPLE.COM

    Issued            Expires           Principal
```

```
Jun 20 16:14:48  Jun 20 22:54:48  krbtgt/EXAMPLE.COM@EXAMPLE.COM $
```

And with the added bonus of being a posix account as well. So, to remind you: we now have Posix, Samba and Kerberos in a single user entry in LDAP. But there are still three password sources, how to solve this? See next section.

7.2.4.10. Password integration

Probably the most wanted feature of a setup where Heimdal uses OpenLDAP as its database backend is the password integration.

Three very common authentication sources in a network are Samba passwords, posix passwords and kerberos passwords. Just using LDAP doesn't magically integrate these three passwords: LDAP is just a storage and, in fact, each application uses it for itself. For example, these are the attributes used for password or secrets storage by Samba, Heimdal and posix:

- Samba: `sambaNTPassword`, `sambaLMPassword`
- Heimdal: `krb5Key`
- posix: `userPassword`

So, `pam_ldap` can change the userPassword when the user runs the password command from the console, but the Heimdal key and Samba hashes won't be changed. Thus, we have a syncronization problem.

Some administrators run scripts to solve this, or only allow the user to change his/her password via some sort of front-end which will take care of the details of updating all password hashes. Another option that is available is to use the contributed `smbk5pwd` module.

The `smbk5pwd` module is available in the contribs directory of the OpenLDAP tarball and, when built with Samba and Kerberos support, allows for this password integration to work automatically. The module is available by default in the `openldap-servers` package.

This integration happens in three ways:

- EXOP password modifications

  This module intercepts OpenLDAP EXOP password modifications and updates both the Kerberos key and the Samba hashes of the same entry, if they are present. This means that a `ldappasswd` command, for example, will end up changing the Samba and Kerberos passwords. Samba, when using the `ldappasswdsync` option in `smb.conf`, also ends up performing an EXOP password modification and will thus update the Kerberos key without even knowing it.

- kpasswd

  When Heimdal receives a password change request via kadmin or kpasswd, it will check if the target entry contains Samba password hashes. If it does, these hashes will also be updated. The userPassword attribute, used for simple binds, is not touched, but see below.

- simple binds (userPassword)

  Simple binds use the userPassword attribute for password verification. If this attribute contains the special hash specified {K5KEY}, then the password verification will be performed against the kerberos key of the same entry. So, in order to make simple binds use the kerberos password, all we have to do is replace the userPassword attribute with {K5KEY}.

The following configuration changes are necessary in order to use the smbk5pwd module:

```
(...)
modulepath      /usr/lib/openldap
moduleload      back_monitor.la
moduleload      syncprov.la
moduleload      ppolicy.la
moduleload      smbk5pwd.so
password-hash   {K5KEY}
(...)
database bdb
(...)
overlay ppolicy
ppolicy_default "cn=default,ou=Password Policies,dc=example,dc=com"

overlay smbk5pwd
(...)
```

If we don't change the server password hash mechanism to {K5KEY}, then password changes via EXOP will overwrite the userPassword attribute with the new hash instead of leaving it at {K5KEY}.

The smbk5pwd module accepts some configuration directives like smbk5pwd-enable and smbk5pwd-must-change, please see its README file in the openldap-servers documentation directory for details.

If Samba is being used, then the ldappasswdsync option should be set to Only. With this option, Samba will only perform the EXOP password modification and expect the OpenLDAP server to update the Samba hashes, which is exactly what smbk5pwd will do.

To the [global] section of /etc/samba/smb.conf, add:

```
ldap passwd sync = Only
```

Now, test `ldappasswd`, `smbpasswd` and `kpasswd`: a password change performed by any of these should change all three authentication sources.

7.2.5. Tuning

Here we will see some important and basic tuning tips for OpenLDAP. They will help you get the most out of your directory server.

7.2.5.1. `DB_CONFIG`

The default backend for OpenLDAP is called BDB, which stands for Berkeley Data Base. This is a very robust database used by several different projects, but it has many buttons and dials. The `DB_CONFIG` file inside the database directory is used to adjust them.

OpenLDAP (and the Mandriva packages) ship with a default `DB_CONFIG` file, but it better be tuned for each specific environment. An incorrect configuration can lead to severe performance problems and data integrity issues.

Here we will explain some of the more important `DB_CONFIG` options that any OpenLDAP administrator should know.

`set_cachesize<gbytes><bytes><ncache>`

This parameter sets the cache that the BDB library will use. The syntax is:

- `<gbytes>`: size of the cache in gigabytes. So, `1` would mean one gigabyte.
- `<bytes>`: size of the cache in bytes. For example, if the previous parameter were set to `1` and this one to `536870912`, the total size of the cache would be the sum of both values, or `1610612736` bytes. The maximum cache size for each segment (see next option) if 4 gigabytes.
- `<ncache>`: the number of caches. Each cache is allocated in a contiguous memory region. A value of zero or one here indicates just one segment.

`set_lg_bsize<bytes>`

The BDB is a transactional backend. This roughly means that write operations are first noted in a log file and only later commited to the database. The `set_lg_bsize` parameter sets the size of the write buffer (in bytes) for this transaction log. Whenever this amount of bytes was written, the buffer is flushed to disk.

`set_lg_dir<path>`

By default, the transaction log files are written to the same directory where the database resides. This path can be changed so that another disk can be used for these files, improving overall performance.

Most `DB_CONFIG` settings become only effective after the database environment is rebuilt. This can be done with the `db_recover` command. For example, to rebuild the environment in `/var/lib/ldap` the command would be `db_recover -v -h /var/lib/ldap` (`-v` for added verbosity).

Only run `db_recover` when the `slapd` daemon is stopped!

After adjusting `DB_CONFIG`, specially the cache parameter, you should use the `db_stat -m -h /var/lib/ldap` command to check its efficienty. This output will tell you, among other things, the hit percentage the cache is getting. The more the better. Tipical good values are over 90% hit. If the server has enough RAM available, this can be boosted up to 99% or sometimes even 100%.

7.2.5.2. OpenLDAP cache

Besides the BDB cache, OpenLDAP also has its own cache of entries. The `cachesize` configuration parameter takes one value as argument which specifies the number of entries that should be held in the cache. The default value is `1000` (a thousand). This is different to the BDB cache and has a lower impact on performance.

7.2.5.3. Indexes

Indexes are of the utmost importance for any type of database, OpenLDAP included. Without indexes, searches can take seconds to complete and have a high CPU usage.

The index parameter in a database section in slapd.conf is used to specify which attributes should be indexed and with what index type. The most common syntax is:

```
index <attr1[,attr2,...]> <[pres,eq,approx,sub]>
```

An index type varies according to the kind of searches that will be done with that attribute. For example, substring searches (like uid=*john*) will need a sub index type to be fast. These are the most common types:

pres

Called the presence index, is used in tests for the presence of the attribute.

eq

The equality index is used in equality tests. For example, (uid=john).

approx

The approximate index is used in searches that use the approximate test. For example, (uid~=sisko).

sub

This index is used for substring searches, like (uid=*john*). Note that by default OpenLDAP will not apply a substring index to attributes with less than 2 characters (see slapd.conf(5) for details on how to change this).

The next example uses different type of indexes in some attributes:

```
    index   objectClass,uid,uidNumber,gidNumber,memberUid   eq
    index   ou                                              eq
    index   cn,mail,surname,givenname                       eq,sub
    index   entryCSN,contextCSN,entryUUID                   eq
```

Whenever a search is performed on attributes that do not have the appropriate indexes, a warning will be logged by slapd:

```
Jul 31 14:12:36 pandora slapd[15130]: conn=8 op=1 SRCH
    base="dc=example,dc=com" scope=2 deref=0 filter="(ou=remotes)"
Jul 31 14:12:36 pandora slapd[15130]: <= bdb_equality_candidates: (ou)
    index_param failed (18)
```

Whenever this happens, it means the search operation was very slow. Either the attribute should be indexed or the search should be modified exclude this attribute.

When an index is added after the database is already populated, it has to be reindexed. This is done by shutting down the service and running the `slapindex` command. It will also reindex all the other attributes, so it could take some time on large databases.

7.2.5.4. RAM

OpenLDAP can benefit a lot from added RAM. Whenever possible, or whenever noticing that the cache is not being very effective, add RAM to the server. The directory will be happy and so will be your users.

7.2.6. Advanced usage

OpenLDAP has some very interesting overlays that can help run an identity server. Here we will take a look at some of them, namely:

- Password policies: control expiration, password quality, account lockout, etc.
- Unique overlay: guarantee that some attribute has an unique value
- Dynamic groups and lists: create groups and list "on the fly"
- Referential integrity: maintain consistency among entries which reference other entries
- Replication with syncrepl: create copies of your data on other OpenLDAP servers

The Mandriva OpenLDAP server package has most overlays and backends built as separate modules inside the `/usr/lib/openldap` directory. To load an overlay, one has to edit `/etc/openldap/slapd.conf` like this:

```
(...)
modulepath       /usr/lib/openldap
moduleload       overlay-filename
(...)
database bdb
(...)
overlay overlay-name
 overlay-specific-options
(...)
```

So, first load the overlay module with `moduleload` and then later inside the `database` section activate it via the `overlay` directive. Any overlay specific configuration options can now be used.

7.2.6.1. Password Policies

The password policy overlay (ppolicy.la) intercepts LDAP password changes and applies several policies such as:

- password aging
- password history
- password length
- password expiration
- account lockout
- forced password change

Several policies can be defined on the server and users can be assigned to any specific policy or get the default one.

Recent pam_ldap versions (182 or higher) support this OpenLDAP module. To enable it, set pam_lookup_policyyes in /etc/ldap.conf

As an example, we will configure an account with a forced password change, enforce minimum length and use a password history. The full list of policies as well as their usage can be found in the slapo-ppolicy(5) manpage.

1. The default /etc/openldap/slapd.conf created by the openldap-mandriva-dit package already has support for password policies. Just for completeness, the needed changes are emphasized below:

   ```
   (...)
   include /usr/share/openldap/schema/dyngroup.schema
   include /usr/share/openldap/schema/ppolicy.schema
   (...)
   modulepath      /usr/lib/openldap
   moduleload      back_monitor.la
   moduleload      syncprov.la
   moduleload      ppolicy.la # loads the module
   (...)
   database        bdb
   suffix          "dc=example,dc=com"
   (...)
   overlay ppolicy # activates the overlay
   ppolicy_default "cn=default,ou=Password Policies,dc=example,dc=com"
   (...)
   ```

   After restarting the server, it will load the password policy module and start to broadcast it in the list of capabilities this server has.

2. Now, we need to define a policy. Notice that in the configuration file a default policy is already defined. Its contents are very basic, not enforcing anything in particular yet:

   ```
   dn: cn=default,ou=Password Policies,dc=example,dc=com
   cn: default
   objectClass: pwdPolicy
   objectClass: namedObject
   pwdAttribute: userPassword
   ```

   > The `namedObject` object class is defined in the the `kolab.schema` schema file. If you ever decide to not load this schema, password policies will break unless `namedObject` becomes defined elsewhere.

   If we want to enforce the policies we listed earlier, we need to either change this entry or create a new policy. We will do it by changing this entry, which is the default policy for our database. Later it can be cloned to other policies if we wish to do so.

   Below is the command to add the policies we want. Members of the `AccountAdmins` system group can change the policies:

   ```
   $ ldapmodify -x -D 'uid=Account Admin,ou=System Accounts,dc=example,dc=com' -W
   Enter LDAP Password: secretpassword
   dn: cn=default,ou=Password Policies,dc=example,dc=com
   changetype: modify
   add: pwdMustChange
   pwdMustChange: TRUE
   -
   add: pwdCheckQuality
   pwdCheckQuality: 2
   -
   add: pwdInHistory
   pwdInhistory: 2
   -
   add: pwdMinLength
   pwdMinLength: 5

   modifying entry "cn=default,ou=Password Policies,dc=example,dc=com"

   ^D
   ```

   Explaining what we did:

   - `pwdMustChange:TRUE`

     Forces users whose password was reset (more on this later) to change the password. LDAP operations will be severily restricted until this happens.

   - `pwdCheckQuality:2`

Activates password quality checking. This will enforce the `pwdMinLength` setting that we set below. The value 2 means that if for some reason the quality of the password cannot be checked (because the client is giving it to us in a hashed form, for example), then the check will fail.

- `pwdInHistory:2`

  Activates password history checking and makes it remember up to two old passwords.

- `pwdMinLength:5`

  Defines the minimum password length to be 4 characters. If it's less, the new password is rejected.

Now, all entities which do not have a specific instruction to use another policy will be subject to this default policy. Even the powerfull System Accounts!

3. We will begin our testing with the `peter` user by forcing him to change his password. The policy already mandates this, so all we have to do as admins is to mark his password as being reset:

```
$ ldapmodify -x -D 'uid=Account Admin,ou=System Accounts,dc=example,
        dc=com' -W
Enter LDAP Password: secretpassword
dn: uid=peter,ou=People,dc=example,dc=com
changetype: modify
add: pwdReset
pwdReset: TRUE

modifying entry "uid=peter,ou=People,dc=example,dc=com"

^D
```

Now let's see what happens when `peter` tries to, for example, do a search operation:

```
$ ldapsearch -x -LLL -D uid=peter,ou=People,dc=example,dc=com
        -W uid=peter uid
Enter LDAP Password: peter's password
Insufficient access (50)
Additional information: Operations are restricted to
        bind/unbind/abandon/StartTLS/modify password
```

So, `peter` is now severely restricted until he changes his password. If we enable `ppolicy` support in the client, we get a hint about what has happened:

```
$ ldapsearch -x -LLL -D uid=peter,ou=People,dc=example,dc=com
        -W -e ppolicy uid=peter uid
Enter LDAP Password: peter's password
ldap_bind: Success (0); Password must be changed
Insufficient access (50)
```

So, let's go ahead and change the password to, for example, `1234`:

```
$ ldappasswd -x -D uid=peter,ou=People,dc=example,dc=com -W -s
        1234 uid=peter,ou=People,dc=example,dc=com
Enter LDAP Password: peter's password
Result: Constraint violation (19)
Additional info: Password fails quality checking policy
```

We just saw the `pwdMinLength` policy kicking in: we used a password 4 characters long. Let's try one with 5:

```
$ ldappasswd -x -D uid=peter,ou=People,dc=example,dc=com -W
       -s 12345 uid=peter,ou=People,dc=example,dc=com Enter LDAP
Password: Result: Success (0)
```

Ok, this worked. Finally, to demonstrate the password history policy, let's try to change the password back to the previous value:

```
$ ldappasswd -x -D uid=peter,ou=People,dc=example,dc=com
         -W -s peteroldpass uid=peter,ou=People,dc=example,dc=com
Enter LDAP Password: 12345
Result: Constraint violation (19)
Additional info: Password is in history of old passwords
```

To assign a different policy to a specific user, use the `pwdPolicySubentry` attribute and point it to the dn of the policy to be used. For example, to apply the `cn=marketing` policy to the `peter` user this is the change we have to do:

```
$ ldapmodify -x -D 'uid=Account Admin,ou=System Accounts,dc=example,
        dc=com' -W
Enter LDAP Password: secretpassword
dn: uid=peter,ou=People,dc=example,dc=com
changetype: modify
add: pwdPolicySubentry
pwdPolicySubentry: cn=marketing,ou=Password Policies,dc=example,dc=com

modifying entry "uid=peter,ou=People,dc=example,dc=com"

^D
```

In this way, we can have different policies applied to different users. It is currently not possible to directly apply a policy to a group of users: each user of the group would need this change in the `pwdPolicySubentry` attribute.

Be careful when using OpenLDAP's password policies together with applications that have their own policies and/or their own password hash. One example is Samba. It wouldn't work very well to have Samba and OpenLDAP policies at the same time, specially since they control different attributes: Samba uses its own password attributes, while the OpenLDAP policy monitors `userPassword`. It could happen that Samba would allow a password change to happen, but OpenLDAP wouldn't because the policies differ.

7.2.6.2. Unique overlay

The unique overlay is used to prevent a part of the tree from repeating an attribute-value pair. For example, the `ou=People` branch could be monitored for changes to the `uidNumber` attribute. If a change would happen to reuse a number for this attribute that is already being used elsewhere, the change would be denied. So, every write to a monitored attribute will trigger an internal LDAP search for this attribute to check if the value is already used somewhere else in the monitored branch.

It is very important that the monitored attributes have appropriate indexes!

As an example, let's configure the server to prevent duplication of `uidNumber` and `cn` under `ou=People`. All options are described in the `slapo-unique(5)` manpage.

1. Changes to the `slapd.conf` configuration file:

```
(...)
modulepath       /usr/lib/openldap
moduleload       back_monitor.la
moduleload       syncprov.la
moduleload       unique.la # loads the module
(...)
database         bdb
suffix           "dc=example,dc=com"
(...)
overlay unique # activates the overlay
 unique_base  ou=People,dc=example,dc=com
 unique_attributes uidNumber cn

(...)
```

   After restarting the OpenLDAP service, the module will be loaded and we can test.

2. Testing:

   Let's see what happens when we try to change the `uidNumber` of a user to a value that another user already has:

   ```
   $ ldapmodify -x -D 'uid=Account Admin,ou=System Accounts,dc=example,
           dc=com' -W Enter LDAP Password:
   secretpassword dn: uid=queen,ou=People,dc=example,dc=com
   changetype: modify replace: uidNumber uidNumber: 1024

   modifying entry "uid=queen,ou=People,dc=example,dc=com"
   ldap_modify: Constraint violation (19)
   additional info: some attributes not unique
   ```

   The change was denied as expected because the `uidNumber` we choose was already taken by another user.

Take care to not list attributes that could have duplicated values. `gidNumber`, for example, while also used to uniquely identify a posix group, can show up several times with the same value under the `ou=People`: think about users sharing the same primary group.

7.2.6.3. Dynamic groups and lists

Dynamic groups and lists are very similar and can be used to automatically populate an entry with elements.

For example, we could have a dynamic list that would automatically expand to all email addresses we have assigned to people. This would be our `all@example.com` email alias.

In the same way, we could have a group called `allusers` that would always be up-to-date regarding the users in our directory. If one user is removed or added, the group would automatically reflect that change whenever searched.

Whenever an entry with a monitored object class is being returned due to a query, a search is performed to collect attributes that should be returned with this entry. It is this search that makes this entry dynamic. The search parameters are encoded in an attribute of the same entry.

The following example shows how to configure the automatic email alias mentioned earlier which will always expand to all the users in the directory:

1. Changes to the `slapd.conf` configuration file:

   ```
   (...)
   modulepath      /usr/lib/openldap
   moduleload      back_monitor.la
   moduleload      syncprov.la
   ```

```
moduleload      dynlist.la # loads the module
(...)
database        bdb
suffix          "dc=example,dc=com"
(...)
overlay dynlist # activates the overlay
# dynlist-attrset <group-oc> <URL-ad> [<member-ad>]
dynlist-attrset nisMailAlias labeledURI

(...)
```

The `nisMailAlias` parameter is the name of the object class which will trigger the search, and the `labeledURI` is the name of the attribute which has the search specification. It becomes easier to understand when we see the actual entry in the next step.

2. This is the entry of our `allusers` email alias:

```
dn: cn=allusers,ou=Aliases,dc=example,dc=com
cn: allusers
objectClass: nisMailAlias
objectClass: labeledURIObject
labeledURI: ldap:///ou=People,dc=example,dc=com?mail?one?
        (objectClass=inetOrgPerson)
```

The default tree doesn't offer an `ou=Alias` branch. In order to add the `cn=allusers` entry shown, you need to add `ou=Alias` first:

```
dn: ou=Aliases,dc=example,dc=com objectClass: organizationalUnit
ou: Aliases
```

Ok, take a deep breath.

First, remember that `nisMailAlias` is the "trigger". Whenever an entry with that object class is returned, the search specified in `labeledURI` will be performed. In our case, this search means:

- base: `ou=People,dc=example,dc=com`
- scope: `one`
- filter: `(objectClass=inetOrgPerson)`
- returned attribute: `mail`

This search will return the `mail` attribute of all entries under `ou=People` which have the `inetOrgPerson` object class.

It works like this:

```
$ ldapsearch -x -LLL cn=allusers
dn: cn=allusers,ou=Aliases,dc=example,dc=com
cn: allusers
objectClass: nisMailAlias
```

```
objectClass: labeledURIObject
labeledURI: ldap:///ou=People,dc=example,dc=com?mail?one?
        (objectClass=inetOrgPerson)
mail: peter@example.com
mail: queen@example.com
```

Compare this output with the entry we displayed earlier: notice any difference?

There are two new attributes that were not in the original entry: the email address for `peter` and `queen`. These attributes were dynamically added as a result of the search that was performed when this entry was returned. If we remove the mail attribute from `queen`, or the whole entry, next time we search for `allusers` that email entry will also be gone.

Next we will do a similar configuration, but for a dynamic group.

1. Changes to the `slapd.conf` configuration file:

   ```
   (...)
   modulepath      /usr/lib/openldap
   moduleload      back_monitor.la
   moduleload      syncprov.la
   moduleload      dynlist.la # loads the module
   (...)
   database        bdb
   suffix          "dc=example,dc=com"
   (...)
   overlay dynlist # activates the overlay
   # dynlist-attrset <group-oc> <URL-ad> [<member-ad>]
   dynlist-attrset groupOfNames labeledURI member

   (...)
   ```

   We now monitor a different object class: `groupOfNames`. Whenever an entry with it is fetched, the search specified in the labeledURI attribute is performed and all entries that match are listed with the `member` attribute.

2. The dynamic group entry:

   ```
   dn: cn=allusers,ou=Group,dc=example,dc=com
   cn: allusers
   objectClass: groupOfNames
   objectClass: labeledURIObject
   member: uid=LDAP Admin,ou=System Accounts,dc=example,dc=com
   labeledURI: ldap:///ou=People,dc=example,dc=com??one?
           (objectClass=inetOrgPerson)
   ```

The search filter is very similar to the previous case. But now we are no longer interested in the contents of the mail attribute, we just want to list all persons we have in the `ou=People` branch.

3. To test, let's see the contents of our dynamic group:

```
$ ldapsearch -x -LLL -b ou=Group,dc=example,dc=com cn=allusers
dn: cn=allusers,ou=Group,dc=example,dc=com
cn: allusers
objectClass: groupOfNames
objectClass: labeledURIObject
member: uid=LDAP Admin,ou=System Accounts,dc=example,dc=com
member: uid=john,ou=People,dc=example,dc=com
member: uid=mary,ou=People,dc=example,dc=com
member: uid=peter,ou=People,dc=example,dc=com
member: uid=queen,ou=People,dc=example,dc=com
labeledURI: ldap:///ou=People,dc=example,dc=com??one?
        (objectClass=inetOrgPerson)
```

Again, when comparing this with the original entry we added, the result are these new `member` attributes that were not there before. These are created dynamically according to the search results. If we delete one of these users from the database and query the group membership again, it will be updated and not list this deleted user anymore.

Note, however, that a search like below will not work against dynamic groups:

```
$ ldapsearch -x -LLL -b ou=Group,dc=example,dc=com
        "(&(objectClass=groupOfNames)
        (member=uid=queen,ou=People,dc=example,dc=com))"
```

For this to work, the server would have to perform the search specified in `labeledURI` once for each existing group, which is not supported yet. But a compare operation works:

```
$ ldapcompare -x cn=allusers,ou=group,dc=example,dc=com
        member:uid=queen,ou=People,dc=example,dc=com
TRUE
```

7.2.6.4. Referential integrity

Referential integrity is a common feature of relational databases. Basically, it means that when one entry which is required by another one is removed, the operation is denied. One common scenario in the LDAP world would be to remove a user account but not remove the user from supplementary groups. This would lead to phantom users inside groups, i.e., groups with users who no longer exist.

To deal with this scenario, administrators tipically use custom scripts or tools that will update the group membership automatically, or any other entity which depends on the entry being removed. This is not difficult to do: just a few search and update operations.

The `refint` overlay can be used to perform some of this automatically. Continuing with the group example, this overlay can be configured keep the integrity of the `member` attribute, which lists the member of a group. Whenever an entry is removed from the directory, this will trigger a search for any `member` attribute which points to this removed entry. Entries which match this search will have the `member` attribute removed, thus maintaining the integrity.

This shows that the `refint` overlay, instead of denying an operation which would break the integrity, actually removes or renames attributes to cope with a deletion or rename.

The next example removes a user from the database and shows how the group this user belongs to is automatically updated.

1. Changes to the slapd.conf configuration file:

   ```
   (...)
   modulepath       /usr/lib/openldap
   moduleload       back_monitor.la
   moduleload       syncprov.la
   moduleload       refint.la # loads the module
   (...)
   database         bdb
   suffix           "dc=example,dc=com"
   (...)
   overlay refint # activates the overlay
   refint_attributes member
   refint_nothing    "uid=LDAP Admin,ou=System Accounts,dc=example,dc=com"

   (...)
   ```

   The overlay has only two configuration parameters:

   - `refint_attributes`: the integrity attribute. If, for example, an entry called `uid=John,ou=People,dc=example,dc=com` is removed, a search is performed for `member="uid=John,ou=People,dc=example,dc=com"` and this attribute is removed from matching entries.

- `refint_nothing`: it's possible that the last attribute of an entry is removed due to integrity constraints. Some object classes may, however, require at least one attribute. Should this situation present itself, the overlay will populate the last attribute with the dn configured here. So, for example, if the last member of a group is `uid=John,ou=People,dc=example,dc=com` and this user was just removed from the database, in order to not leave the group without members (which is forbidden by the `groupOfNames` object class), the overlay will add `member=uid=LDAPAdmin,ou=SystemAccounts,dc=example,dc=com` to the group entry.

The referential integrity overlay works with full DNs. This means that it's not possible to use it to maintain the integrity of group classes which use the `memberUid` attribute, for example, because this attribute holds just an username and not a DN.

2. Testing:

   Let's assume we have the following group with two members:

   ```
   $ ldapsearch -x -LLL cn=mkt member
   dn: cn=mkt,ou=Group,dc=example,dc=com
   member: uid=peter,ou=People,dc=example,dc=com
   member: uid=queen,ou=People,dc=example,dc=com
   ```

   If we delete the `peter` user from the database, the group gets automatically updated to reflect that this user is gone:

   ```
   $ ldapdelete -x -D 'uid=Account Admin,ou=System Accounts,dc=example,
           dc=com' -W uid=peter,ou=People,dc=example,dc=com
   Enter LDAP Password: secretpassword
   $ ldapsearch -x -LLL cn=mkt member
   dn: cn=mkt,ou=Group,dc=example,dc=com
   member: uid=queen,ou=People,dc=example,dc=com
   ```

If we delete the last remaining user from this group from the database (`queen`), then we will see the `refint_nothing` configuration kicking in:

```
$ ldapdelete -x -D 'uid=Account Admin,ou=System Accounts,dc=example,
        dc=com' -W uid=queen,ou=People,dc=example,dc=com
Enter LDAP Password: secretpassword
$ ldapsearch -x -LLL cn=mkt member
dn: cn=mkt,ou=Group,dc=example,dc=com
member: uid=LDAP Admin,ou=System Accounts,dc=example,dc=com
```

Instead of leaving the group without any members, which is forbidden by the `groupOfNames` object class, the overlay added the `refint_nothing` DN as its sole member.

7.2.7. Replication with syncrepl

It's all very nice and cool to have an LDAP server being used by many services on a network. It's all nicely centralized. But what if the server has a failure? Suddenly, all those services will also stop working, even if they are hosted elsewhere. To cope with situations like this, more than LDAP server has to be installed and offer the same data to the clients. This is done via replication.

There are two replication methods available with OpenLDAP: `slurpd` and `syncrepl`. slurpd is being deprecated and is no longer developed, so syncrepl is the way to go. Here are a few key advantages syncrepl has over slurpd:

- the consumer can start empty, it will get in sync automatically
- push or pull: replication always start with the consumer, but the producer can notify of new changes
- no need to change a producer in any way when adding a new consumer: no restarts, no configuration changes
- replication status can be easily verified

Comparing the syncrepl terminology with the one used in slurpd, a consumer is a slave server and a producer is a master server.

To further illustrate the advantages of syncrepl, here is a comparison of the steps needed to add a new slave/consumer server to an already existing master/producer:

- With slurpd:
  - stop master (or make it read-only)
  - dump master database to ldif
  - tell the master about new slave (configure slapd.conf)
  - scp dump file to slave and import it there
  - start slave
  - start master

- With syncrepl:
  - point consumer to provider and start it (the consumer)

In both implementations, however, there is still only one master/producer possible: there is no support for the so called multi-master replication in OpenLDAP.

The replication is flexible enough to allow us to replicate an entire database or just a subset of it. For example, let's assume a mail server in the DMZ needs access to the company's LDAP server to validate the email addresses before accepting a message. One approach would be to make it query the internal LDAP server:

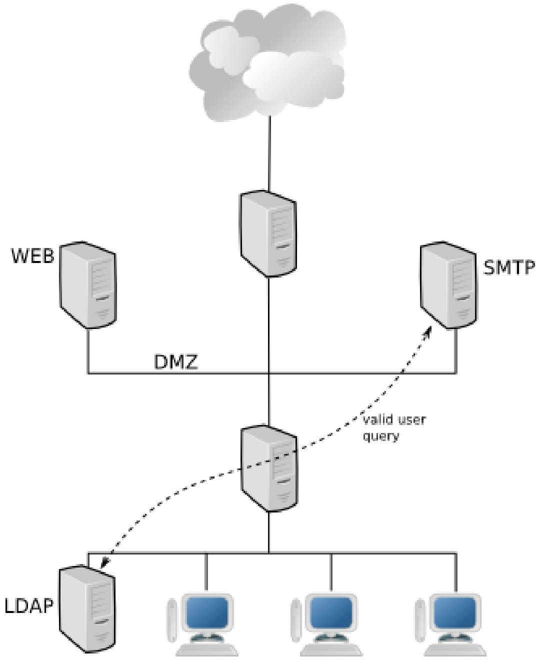

Figure 7-10. Using an internal LDAP server

Another approach, however, would be to place a consumer LDAP server in the DMZ. This server does not need to have the whole database replicated to it: only the information the mail server (and potentially other DMZ servers) needs to do its job. This replication would then filter out many attributes. An additional bonus is that the mail server can keep delivering messages even if the main internal LDAP server is down:

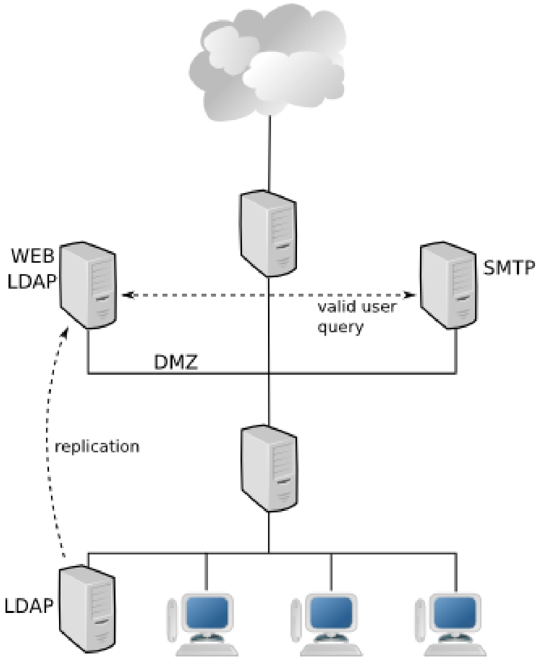

Figure 7-11. Using a replica instead

7.2.7.1. Configuring the Provider

The provider role is done by the `syncprov` overlay. The configuration done by the `openldap-mandriva-dit` package and its setup script already has all the details needed, but we will repeat them here for completeness.

This is what we will need in `slapd.conf`:

```
(...)
modulepath      /usr/lib/openldap
moduleload      back_monitor.la
moduleload      syncprov.la
(...)

database        bdb
suffix          "dc=example,dc=com"
(...)
overlay syncprov
syncprov-checkpoint 100 10
syncprov-sessionlog 100
```

The way syncrepl works (see the OpenLDAP Admin Guide and RFC 4533 for details), write operations are better tracked with the use of a session log.

These are suggested values that vary depending on the number of writes the directory receives in a given amount of time.

- `syncprov-checkpointopsminutes`: update `contextCSN` of the database after `ops` write operations or more than `minutes` minutes have passed since the last checkpoint. Note that these numbers are only checked after a write operation.
- `syncprov-sessionlogops`: size of the log which will record information about write operations.

There is one more change we have to make: configure limits.

`syncrepl` works via an LDAP search operation basically. The consumer will issue this search command on the provider and the response is the data being replicated. So, if the consumer starts empty, the response will contain all the replicated data. Once it is in sync, the responses will only contain the entries that were changed, added or removed. All this means that:

- the consumer needs to be able to read the data: that is, proper ACLs should be in place that grant read access to the data that is supposed to be replicated
- proper limits should be in place for the consumer on the provider: we don't want the consumer to hit a size limit during replication and thus fail to replicate all the entries we need.

In order to address these issues, the common solution is to create an account specifically for replication and tune the ACLs and limits for it. The `openldap-mandriva-dit` does this via the `LDAPReplicators` group and its members:

```
limits group="cn=LDAP Replicators,ou=System Groups,dc=example,dc=com"
    limit size=unlimited
    limit time=unlimited

access to dn.subtree="dc=example,dc=com"
    by group.exact="cn=LDAP Admins,ou=System Groups,dc=example,dc=com" write
    by group.exact="cn=LDAP Replicators,ou=System Groups,dc=example,dc=com" read
    by * break
```

This configuration makes sure that any member of the `LDAPReplicators` group can read all entries and attributes and has no size or time limit applied, which is exactly what we need for proper syncrepl replication. Note that the ACLs can be fine tuned if you prefer to restrict read access to data you know will never be replicated: this is up to the admin.

After making these modifications and restarting the service, this server is ready to serve consumers with replication data. No further changes are necessary here (other than perhaps fine tuning) to add another consumer.

7.2.7.2. Configuring the Consumer

The configuration at the consumer is a bit more complex and resembles the slurpd one. There are two types of replication:

- `refreshOnly`: the consumer initiates the replication by contacting the provider and asking for data. After the data is received, it shuts down the connection and sleeps for a specified amount of time. After this time has elapsed, it wakes up and contacts the provider again to repeat the operation.
- `refreshAndPersist`: again, the consumer initiates the replication by contacting the provider and asking for data. The difference is that after the data is received, the connection stays open. The provider will now use this open connection to notify the consumer of new data as required. So, instead of a cron-like job, consumers using `refreshAndPersist` get new data as soon as it is written to the provider.

For `refresh` replication, this is basically what we need:

```
syncrepl    rid=001
provider=ldap://provider.example.com
starttls=critical
type=refreshOnly
interval=00:01:00:00
searchbase="dc=example,dc=com"
scope=sub
filter="(objectClass=*)"
attrs="*,+"
bindmethod=simple
binddn="uid=LDAP Replicator,ou=System Accounts,ou=global,dc=example,dc=com"
credentials="ldapreplicator"
```

For `refreshAndPersist`, the configuration is like this:

```
syncrepl rid=001 provider=ldap://provider.example.com
starttls=critical type=refreshAndPersist retry="60 +"
searchbase="dc=example,dc=com" scope=sub
filter="(objectClass=*)" attrs="*,+" bindmethod=simple
binddn="uid=LDAP Replicator,ou=System
Accounts,ou=global,dc=example,dc=com"
credentials="ldapreplicator"

updateref   ldap://provider.example.com
```

The more important parameters are described next. Full details can be obtained in the `slapd.conf(5)` manpage.

`rid`

Specifies an unique identifier for this consumer instance

`provider`

URI of the provider

`starttls`

Whether the START TLS operation should be used (`yes`) and if it is critical or not (`critical`). Omit if not desired.

`type`

The type of this replication: either `refreshOnly` or `refreshAndPersist`.

`interval`

In `refreshOnly` mode, this parameter specifies the interval between replication attempts. The format is `dd:hh:mm:ss`.

`retry`

In `refreshAndPersist` mode, this parameter specifies how to deal with the situation where the provider is unreachable. The syntax is a list of `retryinterval` and `numberofretries` pairsi. If the `+` symbol is used in place of the `numberofretries` parameter, then it means indefinetely.

`searchbase,scope,filter,attrs`

These parameters have the same behaviour as regular LDAP search operations and can be used to select with fine detail the data that is to be replicated. The defaults are enough to replicate the whole `searchbase` with all the available attributes, operational and otherwise.

`bindmethod`

Which bind method to use: `simple` or `sasl`.

`binddn`

In the case of `simple` binds, this options specifies the bind dn that is to be used.

`credentials`

Specifies the secret associated with the `simple` or `sasl` bind.

Do not use a hashed password for the `credentials` parameter, it has to be clear text! Remember, in this case the consumer is an LDAP client as any other. When you authenticate against your server, do you type your clear text password or the hashed version?

7.2.7.3. Testing the replication

To test the replication, the simpliest way is to clear the consumer database and start it up empty. The producer logs (at `loglevel256`) should show the consumer's connection (some columns removed for clarity):

```
conn=1 fd=27 ACCEPT from IP=10.0.4.29:4479 (IP=0.0.0.0:389)
conn=1 op=0 BIND
        dn="uid=LDAP Replicator,ou=System Accounts,dc=example,dc=com"
 method=128
conn=1 op=0 BIND
        dn="uid=LDAP Replicator,ou=System Accounts,dc=example,dc=com"
        mech=SIMPLE ssf=0
conn=1 op=0 RESULT tag=97 err=0 text=
conn=1 op=1
        SRCH base="dc=example,dc=com" scope=2 deref=0 filter="(objectClass=*)"
conn=1 op=1 SRCH attr=* +
conn=1 op=2 UNBIND
conn=1 fd=27 closed
```

After a short while, depending on the size of the database, the consumer should be in sync with the provider.

7.2.8. Maintenance tasks

Here we will list some maintenance tasks that should be performed periodically to ensure a smooth running identity service.

7.2.8.1. Cache efficiency

To monitor the BDB cache efficiency, one has to run the db_stat command. Here is an example:

```
# db_stat -m -h /var/lib/ldap/ | head -n 6
40MB 1KB 604B   Total cache size.
1       Number of caches.
40MB 8KB        Pool individual cache size.
0       Requested pages mapped into the process' address space.
975     Requested pages found in the cache (98%).
19      Requested pages not found in the cache.
(...)
```

This shows the current cache size (40MB) and the hit percentage of the main database (98%). A quick way to glance the hits of all database files is to use db_stat -m -h /var/lib/ldap | grep %.

If the cache hit is low (say, under 90%), consider increasing the cache. Don't forget to run db_recover afterwards with the service stopped in order to make the change effective.

Note that you should also always consider the number of requests a database file has had. If there are few requests, the hit percentage may be distorted and is not worth it to act upon.

7.2.8.2. Transaction logs handling

BDB has support for transactions, and OpenLDAP uses this. It means that, with time, the database directory will start filling up with log files:

```
# l /var/lib/ldap/log.*
-rw-------  1 ldap ldap 170K Ago 18 17:40 /var/lib/ldap/log.0000000001
```

By default, each log file will grow up to 10Mbytes in size and then be rotated, which means a new one will be started. Each log file contains all write operations performed on the database, so it would be possible to reconstruct it from scratch if all log files were available. This is called a "catastrophic recovery".

A log file can contain open transactions, that is, changes that were not yet commited to the database. The db_archive command can be used to list the log files that are no longer in use and that could be removed (or backed up elsewhere if you want to have the possibility of running a catastrophic recovery in the future):

```
# db_archive -h /var/lib/ldap
#
```

In this example, no log file was printed by `db_archive`, which means all log files are in use. A different output could be:

```
# db_archive -h /var/lib/ldap
log.0000000001
log.0000000002
log.0000000003
#
```

This means that the displayed log files are no longer in use and could be deleted or backed up. The tool itself can remove these log files automatically if given the `-d` option.

Optionally, the log files can be removed automatically by the library itself whenever they are rotated. For this to happen, one need to specify the `DB_LOG_AUTOREMOVE` flag in the `DB_CONFIG` file:

```
(... other DB_CONFIG options ...)
set_flags DB_LOG_AUTOREMOVE
```

7.2.8.3. Index checking

Indexes are very important for any database, including a directory server. When OpenLDAP gets a search request done on attributes that do not have the proper index, it will log a warning. It is important to periodically, specially in the first days of operation, scan the log files for these warnings and reindex the attributes or change the search parameters:

```
# grep index_param /var/log/ldap/ldap.log
Aug 24 10:04:43 mes5 slapd[27399]: <= bdb_equality_candidates: (ou) index_param fai
Aug 24 10:04:45 mes5 slapd[27399]: <= bdb_equality_candidates: (ou) index_param fai
(...)
Jul 21 15:43:59 mes5 slapd[29666]: <= bdb_equality_candidates: (sambaSIDList) index
Jul 21 15:43:59 mes5 slapd[29666]: <= bdb_equality_candidates: (sambaSIDList) index
(...)
```

This shows that at least two searches were being performed on unindexed attributes. In both cases, the missing index was of the equality type (hence `bdb_equality_candidates`). To remedy the situation, these indexes have to be added and the database has to be reindexed.

7.2.9. Troubleshooting

Here are a few tips for troubleshooting common problems with OpenLDAP.

7.2.9.1. Recognizing users in LDAP

The classical symptom is running `getent passwd john` and the "john" user doesn't exist, but he is in LDAP. There are lots of different systems involved in this simple check, which means many different places things could be wrong.

The path that is passed through when one issues that simple command is more or less this one: glibc, nss, nss_files, nss_ldap, ldap, user entry (`posixAccount` object class). Let's take a look at this:

- Server logs: first of all, check the OpenLDAP server logs to see if a search request is getting there. If it is, than the problem is most likely in the data itself (not present, or incorrect) or in the base search.
- `/etc/nsswitch.conf`: check if `ldap` is correctly added to the `passwd` map and any other map you are querying.
- `nss_ldap`: check that the nss_ldap package is installed.
- `/etc/ldap.conf`: check the server, base search and if ssl is in use or not. If in use (or start tls), check the certificate information on both the server and the client.
- User data: check if the user really exists in the LDAP server and with the correct data (`posixAccount` object class). Verify ACLs. Try to issue manually the search command that nss_ldap is doing.

7.2.9.2. Server logs

OpenLDAP has a comprehensive level of details available for server logs. The most used one, and the recommended one to start debugging, is 256. So, when debugging, make sure you start with `loglevel256` in `slapd.conf`. Other levels are available, see the `slapd.conf(5)` manual page for a full list.

7.2.9.3. Miscellaneous issues

There are several services that are needed for an identity server to operate correctly. Here we will list a few of them:

- Network Time Protocol (NTP): absolutely required by a Kerberos server, it is however always a good idea to run this service on all servers to assure they are all in sync regarding their clocks.

- Domain Name Service (DNS): it is also a requirement to have a proper DNS service available on the network. Many services will probably fail in bizarre ways if DNS is not available or misconfigured.
- Filesystem permissions: the `slapd` daemon runs as an unprivileged user. Sometimes one runs the `db_recover` command as root and forget to change the ownership of the files that were touched back to `ldap`. The initscript takes care of doing this automatically, but sometimes we run things in debug mode without the initscript, and it will fail if the daemon cannot read the database files. The same applies to the main configuration file and SSL certificates: they have to be readable by the `ldap` user.
- If `slapd` suddenly crashed upon startup when the `smbk5pwd.so` overlay was just added, the likely cause is that `/var/heimdal` has strict permissions and the `ldap` user cannot enter this directory. Just give the `ldap` group `r-x` permissions for this directory and it should be fine.

7.3. Database servers

7.3.1. MySQL database server

This chapter provides information for the basic administration of a MySQL database server.

> This document is NOT a tutorial to learn SQL .A complete documentation of available SQL commands in MySQL is available on the MySQL official website (http://mysql.org/doc/#manual (`http://mysql.org/doc/#manual`)).

7.3.1.1. General concepts and references

MySQL is an Open Source relational database management system (RDBMS), developed by the MySQL AB company. This company was created by MySQL developers and offers services around this tool.

This RDBMS relies on many data engines developed by MySQL AB, such as MyISAM (non relational), or by other companies like innoDB (relational) by Innobase Oy.

URLs you should know about :

- MySQL AB company website (http://www.mysql.com/) :product and services;
- MySQL RDBMS (http://dev.mysql.com/) : forums, documentation, downloads... ;
- MySQL download site ; (http://dev.mysql.com/downloads/)
- Official documentation (http://dev.mysql.com/doc/)

Here is the list of available RPMs for Mandriva Enterprise Server 5 :

- mysql : database engine.
- mysql-client : Shell client for the MySQL server (administration tools; backup and restore).
- mysql-common : common files for the previous packages.
- MySQL-Max : this version includes additional features that were not completely tested, or are not required in standard environments. When these features are stable enough, they are included in the standard version. The main advantage of this version is MySQL Cluster and its related services.
- mysql-bench : all the programs and scripts provided by MySQL to run performances tests (benchmarks).

7.3.1.2. Installing and configuring your server

7.3.1.2.1. Basic administration

Although there are existing web or graphic tools to connect to a MySQL server, such as phpMyAdmin (http://www.phpmyadmin.net) ,the basic package provides all the necessary tools to manage the server and the databases.

Just after installation, some basics are needed to run MySQL. The base “mysql” contains the list of databases and MySQL users.

```
# mysql -u root
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1
Server version: 5.0.81 Mandriva Linux - MySQL Standard
Edition (GPL)

Type 'help;' or '\h' for help. Type '\c' to clear the
current input
statement.

mysql> SHOW DATABASES;
+--------------------+
| Database           |
+--------------------+
| information_schema |
| mysql              |
| test               |
+--------------------+
3 rows in set (0.00 sec)

mysql> \r mysql
Connection id:    10
Current database: mysql

mysql> SHOW TABLES;
+---------------------------+
| Tables_in_mysql           |
+---------------------------+
| columns_priv              |
| db                        |
| func                      |
| help_category             |
| help_keyword              |
| help_relation             |
| help_topic                |
| host                      |
| proc                      |
| procs_priv                |
| tables_priv               |
| time_zone                 |
| time_zone_leap_second     |
| time_zone_name            |
| time_zone_transition      |
| time_zone_transition_type |
| user                      |
+---------------------------+
17 rows in set (0.00 sec)

mysql> SELECT user FROM user;
+------+
| user |
+------+
| root |
|      |
| root |
|      |
| root |
+------+
5 rows in set (0.00 sec)
```

No databases are created (except system), but two users are already defined : root and an anonymous user. The first one is the administrator of the database, the second one is a lambda user with access to the `test` database. There are two lines for each user : one allows the user to connect via the host name `localhost` ,the other by the real name of the server or the server IP address.

7.3.1.2.2. Administrator password

The first task to accomplish is to assign a password to these accounts, at least for the administrator. There are many methods to get this done from the shell or straight SQL .Just remember to do it for both users.

```
# mysqladmin -u root
password 'new password'
# mysqladmin -u root -h MES5 password 'new password'
```

The second command can only work if MySQL accepts network connections.

En SQL ,there are two alternatives :

```
# mysql -u root mysql
mysql> SET PASSWORD FOR 'root'@'localhost' =
PASSWORD('new_password');
mysql> SET PASSWORD FOR 'root'@'MES5' =
PASSWORD('new_password');
```

ou

```
# mysql -u root mysql
mysql> UPDATE mysql.user SET Password =
PASSWORD('new_password')
->     WHERE User = 'root';
mysql> FLUSH PRIVILEGES;
```

Here, the command doesn't take the host name into account, which is why you don't need to type it twice. For the anonymous user, you can run the same operations.

Once the password is specified, it's necessary to use it for each connection :

```
# mysql -u root
ERROR 1045 (28000): Access denied for user
'root'@'localhost'
(using password: NO)
# mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1
Server version: 5.0.81 Mandriva Linux - MySQL Standard
Edition (GPL)

Type 'help;' or '\h' for help. Type '\c' to clear the
current input
statement.

mysql>
```

7.3.1.2.3. Deleting anonymous accounts

The anonymous account, although useful to complete the initial tests, is useless for a server in production. Therefore, it should be deleted :

```
# mysql -u root -p mysql
Enter password:
Reading TABLE information FOR completion of TABLE AND
COLUMN names
You can turn off this feature TO get a quicker startup WITH
-A

Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1
Server version: 5.0.81 Mandriva Linux - MySQL Standard
Edition (GPL)

Type ’help;’ or ’\h’ for help. Type ’\c’ to clear the
current input
statement.

mysql> DELETE FROM user WHERE user="";
Query OK, 2 rows affected (0.00 sec)
```

7.3.1.3. Managing databases and users

7.3.1.3.1. Creating/deleting a database

The server administrator must create databases to isolate working data from the applications needing a database :

```
# mysql -u root -p mysql
Enter password:
Reading TABLE information FOR completion of TABLE AND
COLUMN names
You can turn off this feature TO get a quicker startup WITH
-A

Welcome TO the MySQL monitor.  Commands end WITH ; OR \g.
Your MySQL connection id IS 822 TO server version:
5.0.23-log

Type ’help;’ OR ’\h’ FOR help. Type ’\c’ TO clear the
buffer.

mysql> CREATE DATABASE MyBase;
Query OK, 1 row affected (0.00 sec)
mysql> SHOW DATABASES;
+----------+
| DATABASE |
+----------+
| MyBase   |
| mysql    |
| test     |
| tmp      |
+----------+
4 rows IN SET (0.00 sec)
```

This database will be deleted with the following command :

```
mysql> DROP DATABASE
MyBase;
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW DATABASES;
+----------+
| DATABASE |
+----------+
| mysql    |
| test     |
| tmp      |
+----------+
3 rows IN SET (0.00 sec)
```

7.3.1.3.2. Adding and deleting users

Adding and deleting usersSince it is essential to separate data in the databases, it's also important to separate access rights to this data. You just need to create users with the proper rights (administrator of the database, user with read access or other) :

```
mysql> GRANT ALL ON
MyBase.* TO 'MYAdmin'@'localhost' IDENTIFIED BY 'password'
WITH GRANT OPTION;
Query OK, 0 rows affected (0.00 sec)

mysql> GRANT SELECT,INSERT,UPDATE,DELETE ON MyBase.* TO
'MyUser'@'localhost' IDENTIFIED BY 'password';
Query OK, 0 rows affected (0.00 sec)

mysql> FLUSH PRIVILEGES;
Query OK, 0 rows affected (0.00 sec)
mysql> SELECT host,user,password FROM user;
+-----------+----------+------------------+
| host      | user     | password         |
+-----------+----------+------------------+
| localhost | root     | 4b9d71ac52feb410 |
| MES5      | root     | 4b9d71ac52feb410 |
| localhost | MonAdmin | 614e5cb70bb87e92 |
| localhost | MonUser  | 614e5cb70bb87e92 |
+-----------+----------+------------------+
4 rows IN SET (0.00 sec)
mysql> SELECT * FROM db;
```

```
+-----------+---------+----------+-------------+-------------+-------------+-------
| Host      | Db      | User     | Select_priv |
Insert_priv | Update_priv | Delete_priv | Create_priv |
Drop_priv | Grant_priv | References_priv | Index_priv |
Alter_priv | Create_tmp_table_priv | Lock_tables_priv |


+-----------+---------+----------+-------------+-------------+-------------+-------
| %         | test    |          | Y           | Y
| Y           | Y           | Y           | Y         | N
| Y               | Y            | Y            | Y
| Y                |
| %         | test\_% |          | Y           | Y
| Y           | Y           | Y           | Y         | N
| Y               | Y            | Y            | Y
| Y                |
| localhost | MaBase  | MonAdmin | Y           | Y
| Y           | Y           | Y           | Y         | Y
| Y               | Y            | Y            | Y
| Y                |
| localhost | MaBase  | MonUser  | Y           | Y
| Y           | Y           | N           | N         | N
| N               | N            | N            | N
| N                |


+-----------+---------+----------+-------------+-------------+-------------+-------
4 rows IN SET (0.00 sec)
```

We find the users created and their appropriate rights on the appropriate databases.

To delete a user, use the following commands :

```
mysql> REVOKE ALL ON
MyBase.* FROM MyUser@localhost;
Query OK, 0 rows affected (0.00 sec)

mysql> DROP user MyUser@localhost;
Query OK, 0 rows affected (0.00 sec)
```

You need to revoke the user’s rights before you can drop him.

7.3.1.3.3. Insert data in databases

Most software includes a file which contains a data structure and the necessary data to get it going. This file is a succession of SQL commands that the administrator would have had to create the database and insert the data. The administrator will pass a command like this one, once the database and the corresponding user is created :

```
# mysql -u MyAdmin -p
MyBase < myFile.sql
Enter password:
```

7.3.1.4. MySQL server tree

7.3.1.4.1. MySQL data storage

Databases and log files are stored in a directory specified by the `datadir` variable :

```
mysql> SHOW VARIABLES
LIKE 'datadir';
+---------------+-----------------+
| Variable_name | Value           |
+---------------+-----------------+
| datadir       | /var/lib/mysql/ |
+---------------+-----------------+
1 row IN SET (0.00 sec)
```

7.3.1.4.2. Configuring the server and local clients

The server, just like the clients, doesn't need a configuration file and can start with the default values. Nonetheless, during the server's life cycle, the database administrator may want to adapt some variables to improve the performance. Instead of modifying these variables during a future server restart, the configuration file allows you to store these values, and to reload them upon server restart. Also, we could store the necessary information for local clients.

This file is called `my.cnf` .

MySQL searches these files on startup :

- `/etc/my.cnf` : contains global options ;
- `datadir/my.cnf` : contains server specific ;
- `defaults-extra-file` : is the file specified by –defaults-extra-file=# ;

- `~/.my.cnf` : contains user specific options.

7.3.1.4.3. Logs and database

In the `datadir` directory, we find log files, MySQL sockets, as well as directory containing the MySQL databases/MyISAM themselves :

```
# ll /var/lib/mysql
total 21068
-rw-rw----  1 mysql mysql 10485760 aoû 28 10:18 ibdata1
-rw-rw----  1 mysql mysql  5242880 aoû 28 10:18
ib_logfile0-rw-rw----  1 mysql mysql  5242880 aoû 28 10:16
ib_logfile1
drwx--x--x  2 mysql mysql     4096 aoû 28 10:16 mysql/
-rw-rw----  1 mysql mysql    15151 aoû 28 10:16
mysql-bin.000001
-rw-rw----  1 mysql mysql   493595 aoû 28 10:16
mysql-bin.000002
-rw-rw----  1 mysql mysql    11925 aoû 28 10:16
mysql-bin.000003
-rw-rw----  1 mysql mysql      117 aoû 28 10:18
mysql-bin.000004
-rw-rw----  1 mysql mysql      422 aoû 28 10:18
mysql-bin.000005
-rw-rw----  1 mysql mysql       98 aoû 28 10:18
mysql-bin.000006
-rw-rw----  1 mysql mysql      114 aoû 28 10:18
mysql-bin.index
srwxrwxrwx  1 mysql mysql        0 aoû 28 10:18
mysqlmanager.sock=
srwxrwxrwx  1 mysql mysql        0 aoû 28 10:18
mysql.sock=drwx--x--x  2 mysql mysql     4096 aoû  5 12:45
test/
drwx------  2 mysql mysql     4096 aoû 28 10:16 tmp/
```

In fact, we find MySQL databases, `test` and `tmp` ,as well as `ibdata1` containing tables using the InnoDB storage engine.

The `mysql-bin.XXXXXX` contains server binairies, allowing you to replay all the operations accomplished before the a given moment in time (after a restore from backup operation, for example.)

7.3.1.5. Managing the server

7.3.1.5.1. Backup and restore a database

Restoring a database implies, mainly, inserting data in a database. You basically need to setup the backup. One command delivered with MySQL does this: `mysqldump` .

```
# mysqldump
--add-drop-table -u MonAdmin -p MaBase > monFichier.sql
Enter password:
# cat monFichier.sql
-- MySQL dump 10.9
--
-- Host: localhost    Database: MaBase
-- ------------------------------------------------------
-- Server version       4.1.12

/*!40101 SET
@OLD_CHARACTER_SET_CLIENT=@@CHARACTER_SET_CLIENT */;
/*!40101 SET
@OLD_CHARACTER_SET_RESULTS=@@CHARACTER_SET_RESULTS */;
/*!40101 SET
@OLD_COLLATION_CONNECTION=@@COLLATION_CONNECTION */;
/*!40101 SET NAMES utf8 */;
/*!40014 SET @OLD_UNIQUE_CHECKS=@@UNIQUE_CHECKS,
UNIQUE_CHECKS=0 */;
/*!40014 SET @OLD_FOREIGN_KEY_CHECKS=@@FOREIGN_KEY_CHECKS,
FOREIGN_KEY_CHECKS=0 */;
/*!40101 SET @OLD_SQL_MODE=@@SQL_MODE,
SQL_MODE=’NO_AUTO_VALUE_ON_ZERO’ */;
/*!40111 SET @OLD_SQL_NOTES=@@SQL_NOTES, SQL_NOTES=0 */;

--
-- Table structure for table `test`
--

DROP TABLE IF EXISTS `test`;
CREATE TABLE `test` (
`test` char(255) DEFAULT NULL
) ENGINE=MyISAM DEFAULT CHARSET=latin1;

--
-- Dumping data for table `test`
--


/*!40000 ALTER TABLE `test` DISABLE KEYS */;
LOCK TABLES `test` WRITE;
UNLOCK TABLES;
/*!40000 ALTER TABLE `test` ENABLE KEYS */;

/*!40101 SET SQL_MODE=@OLD_SQL_MODE */;
/*!40014 SET FOREIGN_KEY_CHECKS=@OLD_FOREIGN_KEY_CHECKS */;
```

```
/*!40014 SET UNIQUE_CHECKS=@OLD_UNIQUE_CHECKS */;
/*!40101 SET CHARACTER_SET_CLIENT=@OLD_CHARACTER_SET_CLIENT
*/;
/*!40101 SET
CHARACTER_SET_RESULTS=@OLD_CHARACTER_SET_RESULTS */;
/*!40101 SET COLLATION_CONNECTION=@OLD_COLLATION_CONNECTION
*/;
/*!40111 SET SQL_NOTES=@OLD_SQL_NOTES */;
```

In order to avoid having to purge the database to restore, you need to create a backup with the option `-add-drop-table` ,that adds the command to destroy the table before recreating it (if it exists) **before the restore operation** .You need to use the command covered in the Inserting data paragraph. to restore the backup from `myfile.sql` .

```
# mysql -u MyAdmin -p MyBase < myFile.sql
Enter password:
```

MySQL documentation provides more information on the setup of a backup and it's usage. The use of binary logs is also described there.

7.3.1.5.2. Accessing server information

To activate, modify and delete log files (for debug), or optimize server performance according to it's current use, or modify the type of default tables, you have to know the state and the configuration of the server, and modify it without restarting it.

Server configuration

Server configuration is accessible using the following SQL command :

```
mysql> SHOW
VARIABLES;


+---------------------------------+-------------------------------------------
| Variable_name                   | Value
|


+---------------------------------+-------------------------------------------
| back_log                        | 50
|
| basedir                         | /
|
| bdb_cache_size                  | 8388600
|
| bdb_home                        |
/var/lib/mysql/
|
| bdb_log_buffer_size             | 32768
|
| bdb_logdir                      |
|
| bdb_max_lock                    | 10000
|
| bdb_shared_data                 | OFF
|
| bdb_tmpdir                      |
/var/lib/mysql/.tmp/
|
| binlog_cache_size               | 32768
|
...
```

Most variable can be configured in `/etc/my.cnf` and will be taken into account upon server reload. Nonetheless, you may want to modify some of them while the server is live, without having to restart the server.

Moreover, certain values may be modified for the current sessions, or for all sessions (variable list).

You need to have `SUPER` priviliges to modify `GLOBAL` variable. By default, if `SESSION` or `GLOBAL` are not specified, it implies a modification for the current session only.

```
mysql> SET
sort_buffer_size=10000;
mysql> SET SESSION sort_buffer_size=10000;
mysql> SET GLOBAL sort_buffer_size=10000;
```

For a modification to last, you need to make it in the `/etc/my.cnf` file for the value specified on the fly to be also taken into account.

Server state

To adapt server variables while the server is live, you need to see the state of a number of indicators :

```
mysql> SHOW
STATUS;
+--------------------------------+------------+
| Variable_name                  | Value      |
+--------------------------------+------------+
| Aborted_clients                | 436        |
| Aborted_connects               | 3          |
| Binlog_cache_disk_use          | 0          |
| Binlog_cache_use               | 0          |
| Bytes_received                 | 481501964  |
| Bytes_sent                     | 3913113658 |
| Com_admin_commands             | 877        |
...
```

Consulting the state of these variables has an impact on the variables to modify in the previous paragraph.

7.3.1.6. Troubleshooting

7.3.1.6.1. Forgot the Administrator password

If the administrator looses his password, you can access the server after disabling the priviliges to recreate the password.

You need to restart the server with the `–skip-grant-tables` option, and reasign a password to the administrator.

```
# service mysqld stop
Shutting down MySQL: .
[  OK  ]
# mysqld --skip-grant-tables
mysqld: Can’t create/write to file ’/root/tmp/ibi2gyFL’
(Errcode: 13)
090605 12:08:05  InnoDB: Error: unable to create temporary
file; errno: 13
090605 12:08:05 [Warning] Can’t open and lock time zone
table:
Table ’mysql.time_zone_leap_second’ doesn’t exist trying to
live without them
090605 12:08:05 [Note] mysqld: ready for connections.
Version: ’5.0.81’  socket: ’/var/lib/mysql/mysql.sock’
port: 0
Mandriva Linux - MySQL Standard Edition (GPL)
# mysql -u root mysql
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 5.0.24-log

Type ’help;’ or ’\h’ for help. Type ’\c’ to clear the
buffer.

mysql> UPDATE user SET Password=PASSWORD(’testor’) WHERE
User=’root’;
Query OK, 0 rows affected (0.01 sec)
Rows matched: 2  Changed: 0  Warnings: 0

mysql> FLUSH PRIVILEGES;
Query OK, 0 rows affected (0.00 sec)
# killall mysqld
# service mysqld start
Starting MySQL:
[  OK  ]
```

The test checks that it really is necessary to authentify, that some password is not correct and finally, the client connects properly if the right password is used.

```
# mysql -u root mysql
ERROR 1045 (28000): Access denied for user
’root’@’localhost’
(using password: NO)
# mysql -u root -p mysql
```

```
Enter password:
ERROR 1045 (28000): Access denied for user
'root'@'localhost'
(using password: YES)
# mysql -u root -p mysql
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 4 to server version: 5.0.24-log

Type 'help;' or '\h' for help. Type '\c' to clear the
buffer.

mysql>
```

7.3.1.6.2. Others

The official MySQL documentation (`http://dev.mysql.com/doc/refman/5.0/fr/problems.html`) identifies the most common incidents for this service.

7.3.1.7. Updating a user password

The administrator can change the password for a user with the following command :

```
mysql> SET PASSWORD FOR
'myUser'@'localhost' = PASSWORD('password');
```

A user can change his own password by omitting the `FOR` :

```
mysql> SET PASSWORD =
PASSWORD('password');
```

7.3.1.8. Accessing MySQL through the network

By default, MySQL server is installed without allowing access to the server from the network (TCP). It's a security feature for the installation to allow the administrator to create his/her own password. When you need to allow access to MySQL through the network, comment out `skip-networking` in `/etc/my.cnf` and restart the server.

7.3.1.9. Optimizing tables

While a database is live, data is frequently deleted in the tables. Even if the data doesn't show up in the tables, disk space used by this data is not freed.

To free the disk space, you need to reoganize the space occupied by the table. Also, this maintenance operation rebuilds indexes to speed up data access.

From the shell, the `myisamchk` command completes this operation for MyISAM tables while the `isamchk` does it for ISAM tables. This tool is also used to analyse or repair database tables.

```
# myisamchk  --silent -a -r
-S /var/lib/mysql/*/*.MYI
```

You can accomplish the same objective with an SQL command. Note that the table is locked during this operation :

```
mysql> \r test
Connection id:    488
Current DATABASE: test

mysql> CREATE TABLE `test` (   `test` char(255) DEFAULT NULL )
ENGINE=MyISAM DEFAULT CHARSET=latin1;
Query OK, 0 rows affected (0.00 sec)

mysql> OPTIMIZE TABLE test;

+-----------+----------+----------+-----------------------------+
| TABLE     | Op       | Msg_type | Msg_text
|

+-----------+----------+----------+-----------------------------+
| test.test | OPTIMIZE | STATUS   | TABLE IS already up TO date
|

+-----------+----------+----------+-----------------------------+
1 row IN SET (0.00 sec)
```

While the `myisamchk` command is mainly used in cron jobs, the SQL operation offers the advantage of letting the server manage access to tables for optimization. With `myisamchk` ,you have to validate it yourself.

7.3.2. Managing PostgreSQL; server

7.3.2.1. General concept and references

PostgreSQL is relational datase management system (RDBMS) based on POSTGRES, version 4.2, developed at Berkeley University's Computer science department in California. POSTGRES is the origin of many concepts which will be implemented commercially only years later.

PostgreSQL is an open source descendant of the original Berkeley code. It supports a large part of the SQL standard while offering many modern features :

- complexes requests ;
- foreign key ;
- triggers ;
- views ;
- transactionintergrity ;
- concurent access control (MVCC or *MultiVersion Concurrency Control*).

Moreover, PostgreSQL is user extensible in many ways, by adding :

- new data types ;
- new features ;
- new operators ;
- new agregation functions ;
- new indexing methods ;
- new procedural language.

Thanks to its free software licence, PostgreSQL may be used, modified and distributed freely, no matter what the goal is, whether commercial, private or academic.

Main web references :

- PostgreSQL official website (`http://www.postgresql.org/`)
- official documentation (`http://www.postgresql.org/docs/manuals/`)

7.3.2.2. PostgreSQL package installation

Mandriva Enterprise Server 5 provides 4 packages :

- PostgreSQL`8.3` : client as well as the necessary commands to manipulate databases ;
- PostgreSQL`8.3-contrib` : contributions provided in the official PostgreSQL sources ;
- PostgreSQL`8.3-server` : PostgreSQL server ;

For minimal installation, just install the client and server parts.

7.3.2.3. Basic administration

7.3.2.3.1. Getting started

There are existing web and graphic tools to manage a PostgreSQL server, however the packge provides all the necessary tools to manage the server and the databases.

After installation, you only have the basics to start the server and manage it.

```
# psql -U postgres
Password for the postgres user:
Welcome to psql 8.3.7, the interactive interface for
PostgreSQL.

Type:  \copyright  for terms and condition of distribution
\h for a reminder of the most common commands
\? for a quick reminder of psql commands
\g or end with a ; to execute a request
\q to exit

postgres=#postgres=# \l
Database list
Name    | Owner | Encoding
-----------+-------------+----------
postgres  | postgres     | LATIN9
template0 | postgres     | LATIN9
template1 | postgres     | LATIN9
(3 lignes)
postgres=# \c postgres
You are now connected to the «postgres» database.
postgres=# \dt
No relations found.
postgres=# select * from pg_user;
usename   | usesysid | usecreatedb | usesuper | usecatupd |
passwd  | valuntil | useconfig
```

```
----------+----------+-------------+----------+-----------+----------+----------+--
postgres |       10 | t           | t        | t         |
******** |          |
(1 ligne)

postgres=# select * from pg_shadow;
usename  | usesysid | usecreatedb | usesuper | usecatupd |
passwd  | valuntil | useconfig


----------+----------+-------------+----------+-----------+----------+----------+--
postgres |       10 | t           | t        | t         |
|          |
(1 ligne)

postgres=# select * from pg_roles;
rolname  | rolsuper | rolinherit | rolcreaterole |
rolcreatedb | rolcatupdate | rolcanlogin | rolconnlimit |
rolpassword | rolvaliduntil | rolconfig | oid


----------+----------+------------+---------------+-------------+--------------+---
postgres | t        | t          | t             | t
| t           | t            |            -1 | ********
|         |           |  10
(1 ligne)
```

/media/local/svn/doc/trunk/modules/en/CS-middleware-sql.xml:1786:

The `template1` and `template0` commands dont have any specific elements apart from the fact that `template1` is the default database for the `CREATE DATABASE` command. For example, we could delete `template1` and recreate it from `template0` without side effects. This process could be usefull when `template1` is filled up of useless objects.

The `PostgreSQL` database is also created when the group is initialised. This database is meant to become the default for user and application connections. It's a plain copy of `template1` and it can be deleted and recreated if necessary.

7.3.2.3.2. Administrator password

The first task is to create the administrator password.

```
# psql -U postgres
Welcome to psql 8.3.7, the interactive interface for
PostgreSQL.

Type:  \copyright  for terms and condition of distribution
\h for a reminder of the most common commands
\? for a quick reminder of psql commands
\g or end with a ; to execute a request
```

```
\q to exit

postgres=# ALTER ROLE postgres PASSWORD ’password’;
ALTER ROLE
postgres=# select * from pg_shadow;
usename  | usesysid | usecreatedb | usesuper | usecatupd |
passwd                | valuntil | useconfig


----------+----------+-------------+----------+-----------+------------------------
postgres |       10 | t           | t        | t         |
md52bfd7fb3f9637b310a01bae59a2cf2eb |          |
(1 ligne)
```

PostgreSQL official documentation suggests the creation of a user with the CREATEDB and CREATEROLE rights, which is not/media/local/svn/doc/trunk/modules/ middleware-sql.xml:1786: a super-user, and to use this user to manage the database and users. This strategy avoids using the risks of working with the super-user rights for task that do no require such rights.

```
postgres=# CREATE ROLE
admin CREATEDB CREATEROLE LOGIN PASSWORD ’admin’;
CREATE ROLE
postgres=# select * from pg_shadow;
usename  | usesysid | usecreatedb | usesuper | usecatupd |
passwd                | valuntil | useconfig


----------+----------+-------------+----------+-----------+------------------------
postgres |       10 | t           | t        | t         |
md52bfd7fb3f9637b310a01bae59a2cf2eb |          |
admin    |    16384 | t           | f        | f         |
md5f6fdffe48c908deb0f4c3bd36c032e72 |          |
(2 lignes)

postgres=# select * from pg_roles;
rolname  | rolsuper | rolinherit | rolcreaterole |
rolcreatedb | rolcatupdate | rolcanlogin | rolconnlimit |
rolpassword | rolvaliduntil | rolconfig |  oid


----------+----------+------------+---------------+-------------+--------------+---
postgres | t        | t          | t             | t
| t            | t            |           -1 | ********
|               |           |    10
admin    | f        | t          | t             | t
| f            | t            |           -1 | ********
|               |           | 16384
(2 lignes)
```

By default, PostgreSQL does not implement any authentication. You may configure such a process in the PGDATA/pg_hba.conf (where PGDATA means /var/lib/pgsql/data by default on Mandriva Linux) :

```
# psql -U admin
postgresWelcome to psql 8.3.7, the interactive interface
for PostgreSQL.

Type:  \copyright for terms and condition of distribution
\h for a reminder of the most common commands
\? for a quick reminder of psql commands
\g or end with a ; to execute a request
\q to exit

postgres=>

# cat /var/lib/pgsql/data/pg_hba.conf
# PostgreSQL Client Authentication Configuration
File# ==================================================
...
# TYPE  DATABASE    USER        CIDR-ADDRESS
METHOD
# "local" is for Unix domain socket connections only
local   all         all                               md5
# IPv4 local connections:
host    all         all         127.0.0.1/32          md5
# IPv6 local connections:
host    all         all         ::1/128               md5
# service PostgreSQL restart
Stopping PostgreSQL service:
[  OK  ]
Starting PostgreSQL service:
[  OK  ]
# psql -U admin postgres
Password for admin user:
Welcome to psql 8.3.7, the interactive interface for
PostgreSQL.

Type:  \copyright for terms and condition of distribution
\h for a reminder of the most common commands
\? for a quick reminder of psql commands
\g or end with a ; to execute a request
\q to exit/media/local/svn/doc/trunk/modules/en/CS-middleware-sql.xml:1786:

postgres=>
```

7.3.2.4. Database and user management

7.3.2.4.1. Creating and deleting a database

The server administrator must create databases to isolate working data from the data used by applications.

```
]# psql -U admin
postgres
Password for admin user:
Welcome to psql 8.3.7, the interactive interface for
PostgreSQL.

Type:  \copyright for terms and condition of distribution
\h for a reminder of the most common commands
\? for a quick reminder of psql commands
\g or end with a ; to execute a request
\q to exit

postgres=> CREATE DATABASE test;
CREATE DATABASE
postgres=> \l
Database list
Name    | Owner        | Encoding
-----------+--------------+----------
postgres  | postgres     | LATIN9
template0 | postgres     | LATIN9
template1 | postgres     | LATIN9
test      | admin        | LATIN9
(4 lines)
```

This database will be deleted with the following command :

```
postgres=> DROP
DATABASE test;
DROP DATABASE
postgres=> \l
Database list
Name    | Owner        | Encoding
-----------+--------------+----------
postgres  | postgres     | LATIN9
template0 | postgres     | LATIN9
template1 | postgres     | LATIN9
(3 lines)
```

7.3.2.4.2. Adding and deleting users

Since it is useless to separate data in databases, it essential to assign the proper rights to access that data. You simply need to create users with the proper rights (database administrator, user with read access only, etc.) :

/media/local/svn/doc/trunk/modules/en/CS-middleware-sql.xml:1786:

```
# psql -U admin
postgresPassword for admin user:
Welcome to psql 8.3.7, the interactive interface for
PostgreSQL.

Type:  \copyright for terms and condition of distribution
\h for a reminder of the most common commands
\? for a quick reminder of psql commands
\g or end with a ; to execute a request
\q to exit

postgres=> CREATE DATABASE test;
CREATE DATABASE
postgres=> \l
Database list
Name    | Owner        | Encoding
-----------+---------/media/local/svn/doc/trunk/modules/en/CS-middleware-sql.xml:17
postgres  | postgres     | LATIN9
template0 | postgres     | LATIN9
template1 | postgres     | LATIN9
test      | admin        | LATIN9
(4 lines)
postgres=# CREATE USER myadmin PASSWORD 'password';
CREATE ROLE
postgres=# CREATE USER myuser PASSWORD 'password';
CREATE ROLE
postgres=# GRANT ALL ON DATABASE test TO myadmin WITH GRANT
OPTION;
GRANT
postgres=# \c test myadmin
Password for myadmin user:
You are now connected to the "test" database as "myadmin"
user.
test=> GRANT ALL ON DATABASE test TO myuser;
GRANT
test=> \c test myuser
Password for myadmin myuser:
You are now connected to the "test" database as "myuser"
test=> CREATE TABLE matable ( test varchar(255) DEFAULT
NULL );
CREATE TABLE
test=> \c postgres admin
Password for user admin:
You are now connected to the "postgres" database as "admin"
postgres=> SELECT * FROM pg_user;
usename  | usesysid | usecreatedb | usesuper | usecatupd |
passwd  | valuntil | useconfig
```

```
----------+----------+-------------+----------+-----------+----------+----------+--
postgres |       10 | t           | t         | t         |
******** |          |
admin    |    16384 | t           | f         | f         |
******** |          |
monadmin |    16416 | f           | f         | f         |
******** |          |
monuser  |    16417 | f           | f         | f         |
******** |          |
(4 lines)
```

We find the users created with their respective rights for the appropriate database.

Deleting a user is accomplished with the following command :

```
postgres=> REVOKE ALL
ON DATABASE test FROM monuser;
REVOKE
postgres=> DROP ROLE monuser;
DROP ROLE
```

You first have to "revoke" the user's rights and delete all the objects that he owns, then you may "drop" the user.

7.3.2.4.3. Inserting data into databases

Most software includes a file containing the appropriate DB structure and data to launch the application. This file is in fact a succession of SQL operations that the administrator would have to execute to create the content of the database. The administrator run a command close to the following, after creating the database test and it's corresponding user `myuser` :

```
# psql -U myuser test
< MyFile.sql
Password for myuser:
```

7.3.2.5. Day to day administration

7.3.2.5.1. Backup and restore databases

Restoring a database implies, mainly, inserting data in a database. You basically need to setup the backup. One command delivered with PostgreSQL does this: `pg_dump` .

```
# pg_dump --clean -U
postgres test
password :
--
-- PostgreSQL database dump
--

SET client_encoding = 'LATIN9';
SET check_function_bodies = false;
SET client_min_messages = warning;

SET search_path = public, pg_catalog;

DROP TABLE public.products;
DROP SCHEMA public;
--
-- Name: public; Type: SCHEMA; Schema: -; Owner: postgres
--

CREATE SCHEMA public;


ALTER SCHEMA public OWNER TO postgres;

--
-- Name: SCHEMA public; Type: COMMENT; Schema: -; Owner:
postgres
--

COMMENT ON SCHEMA public IS 'Standard public schema';


SET default_tablespace = ";

SET default_with_oids = false;

--
-- Name: products; Type: TABLE; Schema: public; Owner:
postgres; Tablespace:
--

CREATE TABLE products (
product_no integer,
name text,
price numeric
);


ALTER TABLE public.products OWNER TO postgres;

--
-- Data for Name: products; Type: TABLE DATA; Schema:
public; Owner: postgres
--

COPY products (product_no, name, price) FROM stdin;
```

```
\.


--
-- Name: public; Type: ACL; Schema: -; Owner: postgres
--

REVOKE ALL ON SCHEMA public FROM PUBLIC;
REVOKE ALL ON SCHEMA public FROM postgres;
GRANT ALL ON SCHEMA public TO postgres;
GRANT ALL ON SCHEMA public TO PUBLIC;


--
-- PostgreSQL database dump complete
--
```

If you database relies on OID (for foreigh keys for example), you must tell `pg_dump` to save OID as well. For this, use the `-o` option on the command line.

You just need to use the Section 7.3.2.4.3 previous command to restore the content of the backup `myFIle.sql` .

```
# psql -U myuser test
< myFILe.sql
Password for myuser:
```

Further information on backup and restore are available in the PostgreSQL official documentation.

7.3.2.5.2. Accessing server information : server configuration

Held in `/var/lib/pgsql/data/PostgreSQL.conf` ,server configuration is available through the following SQL command :

```
postgres=> SHOW ALL;
name      | setting |                 description



------------------------+---------+-----------------------------------------------
add_missing_from       | off     | Automatically adds
missing TABLE REFERENCES TO FROM clauses.
archive_command        | unset   | WAL archiving command.
australian_timezones   | off     | Interprets ACST, CST,
EST, AND SAT AS Australian time zones.
authentication_timeout | 60      | Sets the maximum time IN
seconds TO complete client authentication.
...

postgres=> SET timezone TO 'Europe/Paris';
```

```
SET
```

7.3.2.6. The PostgreSQL tree

7.3.2.6.1. Data storing

All the databases and their log files are kept in the directory specified by the `datadir` variable :

```
postgres=# SHOW
data_directory;
data_directory
---------------------
/var/lib/pgsql/data
(1 line)
```

7.3.2.6.2. Server configuration

/media/local/svn/doc/trunk/modules/en/CS-middleware-sql.xml:1786:

There are two main configuration files for PostgreSQL :

- `PostgreSQL.conf` : defines all the execution variables for the server ;
- `pg_hba.conf` : manages the rights access for PostgreSQL server.

PostgreSQL applies the these values on startup. Any modifications to these files require a restart.

7.3.2.7. Advanced configuration

7.3.2.7.1. Modifying an account password

The administrator can change a user's password with the following command :

```
postgres=# ALTER
ROLE myuser PASSWORD 'password';
ALTER ROLE
```

A user can modify his password by omitting the `FOR` :

```
postgres=# \c test
```

```
myuser
Password for myuser :
You are now connected to database "test" as "myuser".
test=> ALTER role myuser password 'password';
ALTER ROLE
```

7.3.2.7.2. Optimisation

While a database is live, data is frequently deleted from tables. Even if the data doesn't appear in the tables anymore, disk space is not made available. To do this, you must reorganise the space occupied by the table.

The tool to accomplish this is quite complex and the data optimization being a key element to enhance RDBMS performances, we recommand you read the documentation (`vacuum` (sql)/`vacuumdb`).

The `vacuumdb` command accomplishes this job from the command line. This same tool is also used to analyse and/or repair database tables.

```
# vacuumdb -f -U
postgres test
Password :
VACUUM
```

You can do this with an SQL command as well.

```
# psql -U postgres
test
Welcome to psql 8.3.7, the interactive interface for
PostgreSQL.

Type:  \copyright  for terms and condition of
distribution
\h for a reminder of the most common commands
\? for a quick reminder of psql commands
\g or end with a ; to execute a request
\q to exit

test=# VACUUM FULL products;
VACUUM
```

Chapter 8. Services Stacks

Managing Services Provided by Mandriva Enterprise Server 5

As presented in the stacks schema, Mandriva Enterprise Server 5 offers a number of the most commonly used services for enterprise servers:

- Standard network services: including Bind name server, DHCP dynamic IP address server, and PXE server (PXELinux) to deploy machines.
- File and print sharing services: including a print server (CUPS), a file and domain server (Samba), and other file servers (NFS and ProFTPD).
- Mail services: including SMTP server (Postfix), POP/IMAP server (Cyrus-IMAP), and SPAM management and viruses.

This documentation presents an overview of these services and of their day-to-day management in a Mandriva Enterprise Server 5 environment.

8.1. Managing Main Network Services

8.1.1. Managing Name Servers with BIND

In this chapter, we briefly describe how to configure global options of BIND bind name server, and how to declare new zones (basically, domain names) to be handled by your DNS server. This means other machines on the local network, and possibly on the Internet, will be able to access the machines and services associated to your own domain names.

8.1.1.1. How Does a DNS Server Work?

A DNS server allows you to associate an IP address to a name and vice versa. For example: `www.mandriva.com` (“Name”) is currently associated to `212.85.147.118` (“Address”). To make an analogy, a name server acts somewhat like a telephone directory: you provide it with a name and it gives you the number which allows you to connect to your correspondent. However, this mechanism is generally transparent to the end user: he never needs to remember or type IP addresses thanks to the DNS servers.

BIND (Berkeley Internet Name Domain) is an implementation of a DNS system based on RFC 1034/1035 (Domain Names). BIND provides 3 main parts:

- A hierarchical namespace;
- A name server which contains information about the name tree and the domain configuration. The server program is `named`.
- resolvers are a set of routines in a C library. Their purpose is to translate names to IP addresses. These routines are called by Internet services like FTP. Resolvers get information from BIND directly. But you can also fix an order to ask for resolution and then get the resolution. The order is fixed in the `/etc/nsswitch.conf` file. Usual sources for name resolution are the `/etc/hosts` file and DNS servers:

  ```
  # cat
        /etc/nsswitch.conf ...  hosts: files nisplus nis dns
        ...
  ```

  `files` is used for `/etc/hosts`, and `dns` for BIND.

There are 3 main types of name server:

Primary server

This is an authoritative server regarding data it provides. This kind of server is able to delegate its authority for sub-domains. The primary server provides the most up-to-date data.

Secondary server

These are backup servers for resolvers. They contain the same data as the primary ones, and update by asking the primary server. This process is called (*zone transfer*). The protocol used is DNS NOTIFY. This mechanism allows master servers to notify their slave servers of changes made to a zone's data.

Caching server

This one does not keep data locally. It is not an authoritative server. It asks primary or secondary servers to get answers for resolution requests and keep it in their cache. It is often used to reduce external requests in a LAN and therefore reduce bandwith needs.

If you want to learn more about BIND, we strongly recommend you read the BIND 9 Administrator Reference Manual (`http://www.bind9.net/manuals`), which is available locally on your machine in the `/usr/share/doc` directory after you installed `bind-doc`. Do not hesitate to browse the official BIND web site (`http://www.bind9.net/`).

8.1.1.2. Installing BIND and the BIND Tree

Installing BIND Getting a working configuration is also fairly easy. Advanced configuration can become rather difficult. We provide a few tips about it.

8.1.1.2.1. Packages to Install and the BIND Tree

You will only need one package: BIND. It contains the name server and the tools to help you get the right configuration:

```
# urpmi bind
```

By default, Mandriva Enterprise Server 5 provides the BIND server in `chroot` mode. It means that BIND programs are re-rooted to a directory than the usual one. It cannot call files outside that directory. It provides a convenient way to isolate an untrusted program and which can be dangerous for the server security.

Finally, here's the main tree for a BIND server. `/var/lib/named` is the root directory of the name server:

```
     /etc/init.d/named
     /var/lib/named/
 |-- dev
 |-- etc
 |-- proc
      |-- xxx
 `-- var
      |-- log
      |-- named
      |-- run
      `-- tmp
```

`/var/lib/named/etc` directory contains the main configuration files:

- `named.conf`: this file is read by the `named` daemon at start up. It contains the path to get data files for all managed domains, a kind of name server for each of them, and general configuration for the name server.
- `rndc.conf` and `rndc.key`: these files configure the rndc behavior. rndc controls the operation of a name server. It communicates with the name server over a TCP connection, sending commands authenticated by digital signatures.
- `named`: is the daemon for the BIND name server listening for queries.

You will find logs in the `/var/lib/named/var/log` file. By default, BIND provides a way to log information in different files, depending on the logged request type:

- `default.log`: general information that's not classified in other types of logs, such as start and stop information.
- `notify.log`: the notify protocol deals with changes in data files on master servers.
- `query.log`: all queries for name servers.
- `security.log`: approval and denial of requests.
- `update.log`: dynamic updates.
- `xfer-in.log`: logs the zone transfer requests the name server is receiving when activated.
- `xfer-out.log`: zone transfer requests the name server is sending when activated.

You can add different types of files using other categories, but the main ones are defined above.

8.1.1.2.2. Managing the BIND Service

You can also use the `service` command which provides you with different options:

```
service named {start|stop|status|restart|reload}
```

start

Start `named` daemon reading general configuration and data zone.

stop

Stop `named` daemon.

status

Gives interesting information about the name server: number of managed zones, log level, information about transfer zones, general status of `named` daemon.

```
# service named status
       number of zones: 4
       debug level: 0
       xfers running: 0
       xfers deferred: 0
       soa queries in progress: 0
       query logging is ON
       recursive clients: 0/1000
       tcp clients: 0/100
       server is up and running
```

reload

This parameter does not restart BIND completely. It just rereads the zone files. Use it when you want to take into account modifications done in zone files.

8.1.1.3. Advanced Configuration and Troubleshooting

8.1.1.3.1. How to fix problems?

If the service didn't start, you should look at the `/var/log/messages` file to read the BIND *debug* message. . If you don't find the error, you can use the `named-checkconf` and `named-checkzone` programs to check your configuration:

`named-checkconf`

This command allows you to check the `named.conf` file and find errors you might have in your syntax:

```
# named-checkconf -t /var/lib/named /etc/named.conf
/etc/named.conf:101: zone '0.0.127.in-addr.arpa': type not
present
```

You must use the `-t` option to specify the chroot directory.

`named-checkzone`

This command allows you to check zone files and find errors you might have in your syntax:

```
# named-checkzone
       local /var/lib/named/var/named/reverse/named.local zone
       local/IN: loaded serial 1997022700 OK
```

The first parameter identify the checked zone as mentioned in the `named.conf` file. The second one is the zone file.

With the `bind-utils` package, you can use many utilities and especially the `dig` command to perform advanced queries on DNS servers. For example, to query your local server about `machine2.mydomain.test`, you could run:

```
$ dig machine2.mydomain.test @127.0.0.1
; <<>> DiG 9.2.3 <section<>> machine2.mydomain.test @127.0.0.1
;; global options:  printcmd
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 3287
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 0

;; QUESTION SECTION:
;machine2.mydomain.test.        IN    A

;; ANSWER SECTION:
machine2.mydomain.test. 38400   IN    A    192.168.1.12

;; AUTHORITY SECTION:
mydomain.test.          38400   IN    NS   mycomputer.mydomain.test.

;; Query time: 14 msec
;; SERVER: 127.0.0.1#53(127.0.0.1)
;; WHEN: Fri Jan 23 14:11:10 2004
;; MSG SIZE  rcvd: 81>
```

8.1.1.3.2. SOA Records

SOA (*Start Of Authority*) records define global zone parameters. It can be essential to optimize it, especially when you manage (*slaves*) servers. Here is an example:

```
    $TTL 3360
    $ORIGIN domain.test.
    @       IN      SOA             ns1.domain.test.  admin.domain.test.        (
     2005072801; serial number
     6H;  refresh
     2H;  update retry
     1W;  expiry
     24H;  minimum
    )

 TXT                     "My test domain"


 IN      NS              ns1.linuxeries.org.
 IN      NS              ns6.gandi.net.
```

The general syntax is:

```
name ttl class
   rr name-server email-addr (sn ref ret ex min)
```

Let's detail the main parameters and how you can optimize them:

`$TTL`

The TTL (*Time To Live*) value defines the duration in seconds that the record may be cached. This is used especially when managing slave servers.

`$ORIGIN`

ORIGIN defines a base value from which non qualified name substitutions are made when processing the zone file. If this value is not defined then it will use the value in the `named.conf` file. @ has to be substituted by the last value of `$ORIGIN` in the zone file.

`class`

Defines the record class. It now has the IN value (Internet).

`name-server`

Is a name server that responds authoritatively for the domain. It's specified by its (*Fully Qualified Domain Name or FQDN*) and ends with a **dot**.

`mail-addr`

This has to be the mail address of the email server administrator responsible for this zone who should be contacted in case of a problem. The syntax is:

`mailbox-name@domain.com. .`

`serial number`

This serial number must be incremented each time the zone file is modified. It's a way for the slave server to identify these changes and to address them. There's no definite rule to compose it but one convention is to use a date-based value to simplify task and reading: `yyyymmddss`, yyyy = year, mm = month, dd = day, and ss = second. This way, you can modify it several times a day.

`refresh`

Defines how often the slave will try to refresh the zone from the master. The RFC recommends between 1200 to 43200 seconds.

`update retry`

The time between retries if the slave (secondary) fails to contact the master when REFRESH delays have expired. Typical values are 180 seconds to 900 seconds, or higher.

`expiry`

Slaves will stop responding to zone queries when time has expired and no contact has been made with the master. The RFC recommends between 1209600 and 2419200 seconds (2 to 4 weeks).

`minimum`

A negative caching time, which means the time a `NAME ERROR = NXDOMAIN` record is cached.

By default, values are seconds but you can use s or S (seconds), m or M (minutes), h or H (hours), d or D (days), w or W (weeks).

The dnsreport (`http://member.dnsstuff.com/pages/dnsreport.php`) web site suggests you check your DNS configuration especially when you have a public name server. Enter your zone or domain name, then read the complete resulted report.

8.1.1.4. Managing BIND Data in OpenLDAP Directory

Mandriva Enterprise Server 5 has based all included services on LDAP directories where possible. BIND is one of these services. You need to follow these steps to host BIND data on LDAP:

1. Import the zone into LDAP at the `ou=dns` branch, which is where our ACLs (*Access Control List*) expect the DNS information to be stored.
2. Configure `named.conf` to use LDAP for each imported zone.
3. Configure LDAP authentication parameters in `named.conf` (`ou=dns` can only be read by members of DNS Admins and DNS Readers)

ATo detail these steps, we will use the following zone example based on the `example.com.zone` zone file:

```
$TTL
   86400 $ORIGIN example.com.  @ IN SOA
   aurelio.example.com. hostmaster.example.com. ( 1 ; serial number
   10800 ; refresh 3600 ; retry 604800 ; expires 86400 ) ; TTL @ IN NS
   aurelio.example.com.  @ IN MX 10 mail.example.com.

   gateway         IN      A       10.0.1.1
   dogs            IN      A       10.0.1.7
   mail            IN      A       10.0.1.8
   aurelio         IN      A       10.0.1.9

   dhcp010         IN      A       10.0.1.10
   dhcp011         IN      A       10.0.1.11

   ns1             IN      CNAME   aurelio
   kdc             IN      CNAME   dogs

   localhost       IN      A       127.0.0.1
```

8.1.1.4.1. Configuring named.conf

We have to configure `named.conf` to consult LDAP for these two zones. Below is an example file.

Please remember that named runs in chroot mode and won't use the system `/etc/hosts` file, but its own inside the chroot. Also avoid making loops: for example, don't use a server name that's in the LDAP zone to specify the LDAP server. In general, it's better to use IP addresses instead of hostnames in `named.conf`.

```
options {
        directory "/var/named";
        allow-transfer { none; };
        notify no;
        allow-query { any; };
};

zone "." {
        type hint;
        file "named.ca";
};

zone "0.0.127.in-addr.arpa" {
        type master;
        file "named.local";
};

zone "example.com" {
        type master;
        database "ldap ldap://127.0.0.1/ou=dns,dc=example,dc=com??sub??
          !bindname=uid=DNS%20Reader%2c
ou=System%20Accounts%2cdc=example%2cdc=com,!x-bindpw=dnsreader 86400";
};

zone "1.0.10.in-addr.arpa" {
        type master;
        database "ldap ldap://127.0.0.1/ou=dns,dc=example,dc=com??sub??!
          bindname=uid=DNS%20Reader%2c
ou=System%20Accounts%2cdc=example%2cdc=com,!x-bindpw=dnsreader 86400";
};
...
```

The databasespecifies the database to use for the zone file. In our case, it's LDAP. The URL seems a bit complicated, so let's explain it. The generic format of the URL follows RFC 2255:

```
ldap://server/basedn?attributes?scope?filter?extensions
```

So, if we want to specify a sub-tree search on ou=dns on localhost, with no default filter, attributes or extensions, it would look like this:

```
ldap://localhost?ou=dns,dc=example,dc=com??sub?
```

However, our DIT requires authenticated searches on so we have to add extensions. Extensions are comma-separated lists of names optionally preceded by "!", indicating its usage is critical. We will use bindname and x-bindpw (the latter one, being non-standard, is prefixed by "x-"). The URL now looks like this:

```
ldap://localhost?ou=dns,dc=example,dc=com??sub??!bindname=uid=DNS
    Reader, ou=System
    Accounts,dc=example,dc=com,!x-bindpw=dnsreader
```

Since extensions are comma-separated, we have to separate the commas in the binddn. We also have to escape the spaces. We do this using the standard URL encoding formats. In our case, the URL then finally becomes:

```
ldap://localhost/ou=dns,dc=example,dc=com??sub??!bindname=uid=DNS%20Reader%2c
    ou=System%20Accounts%2cdc=example%2cdc=com,!x-bindpw=dnsreader
```

Be aware that the /etc/named.conf file must be in 0640 mode, owner root:named, because it contains two secrets now: the rndc key and the LDAP credentials used to bind to the directory (just because of the rndc key it should already have these permissions, but make sure).

8.1.1.4.2. Adding Data in LDAP Directory

After preparing your tree with the mandriva-setup.sh script, you can insert these zone files in LDAP at ou=dns (using the DNS Admin user credentials).

You can use the zonetoldap tool provided by the BIND package. It will parse a complete BIND 9 format DNS zone file, and load the contents into an LDAP directory. If the zone already exists, zonetoldap will exit successfully. If the zone does not exist, or only partially exists, zonetoldap will attempt to add all/missing zone data.

The SRV record has to be commented out because zonetoldap can't handle that yet. We will add it manually later.

```
$ zonetoldap -D 'uid=DNS
   Admin,ou=System Accounts,dc=example,dc=com' -W \ -b
   ou=dns,dc=example,dc=com -z example.com -f example.com.zone -h
   localhost -c Enter LDAP Password: secretpass
```

This will produce the following entries under ou=dns,dc=example,dc=com:

```
dc=com
    dc=example
    relativeDomainName=ns1+zoneName=example.com
    relativeDomainName=mail+zoneName=example.com
    relativeDomainName=localhost+zoneName=example.com
    relativeDomainName=kdc+zoneName=example.com
    relativeDomainName=gateway+zoneName=example.com
    relativeDomainName=dogs+zoneName=example.com
    relativeDomainName=dhcp011+zoneName=example.com
    relativeDomainName=dhcp010+zoneName=example.com
    relativeDomainName=aurelio+zoneName=example.com
    relativeDomainName=@+zoneName=example.com
```

>We still have to add the SRV record from the zone file which we commented. The following LDIF can be used:

```
dn: relativeDomainName=_kerberos._udp+zoneName=example.com,dc=example,
     dc=com,ou=dns,dc=example,dc=com
    objectClass: dNSZone
    relativeDomainName: _kerberos._udp
    zoneName: example.com
    SRVRecord: 0 0 88 dogs
```

Let's add it:

```
$ ldapadd -x -D 'uid=DNS Admin,ou=System Accounts,dc=example,dc=com' \
    -W -f srv.ldif
    Enter LDAP Password: secretpass
    adding new entry "relativeDomainName=_kerberos._udp
      +zoneName=example.com,dc=example,dc=com,
    ou=dns,dc=example,dc=com"
```

You should always review the entries produced by the zonetoldap tool, as there may be inconsistencies with other records.

8.1.1.5. Managing a DHCP Server

DHCP (*Dynamic Host Configuration Protocol*) is a network protocol that allows you to provide IP addresses dynamically to clients, as well as to answer global network configurations (DNS, broadcast, etc.).

8.1.1.5.1. How Does DHCP Work?

The diagram below explains the first dialog between a client asking for an IP address and a DHCP server. It's important to understand this as all these types of packets will appear in the log files. Knowing what it contains will help you to identify non working configurations.

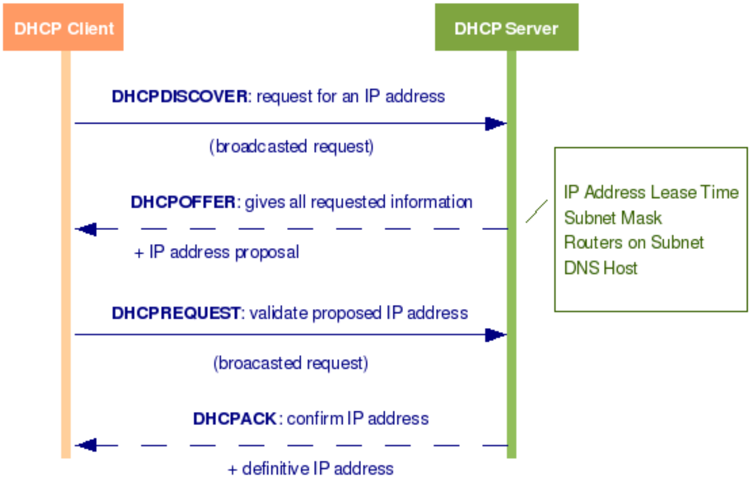

Figure 8-1. How Does DHCP Work?

This is a very basic dialog. It can be much longer (for an IP address denied by a client, for example).

Another reason for a client and a DHCP server to dialog is to manage leases. DHCP servers use IP addresses based on leases: IP addresses are proposed for a limited time to optimize network resources. Then you will find dialogs about leases. As soon as a lease is almost finished, clients can ask the server to renew it. Clients will use DHCPRESQUEST. Also when a lease is nearly finished, the server can launch a DHCPNAK packet which indicates that the lease is expired and will propose to renew it. If there is no answer, then the IP addresses will be available for other requests. Managing leases is one of the most important factors in optimizing DHCP management.

8.1.1.5.2. Installing and Configuring a DHCP Server

8.1.1.5.2.1. Installing Packages and the DHCP Tree

Installing DHCP is rather easy. You will need to install the `dhcp-server` package. As a dependency, `dhcp-common` will also be installed. It contains mainly documentation and the directory for lease files.

```
# urpmi dhcp-server
```

The main DHCP tree files and directories are simple to understand:

- `/etc/dhcpd.conf.sample`: this is a sample of a DHCP server configuration. You can use it as a template. The final file should be named `/etc/dhcp.conf`.
- `/etc/rc.d/init.d/dhcpd`: script to manage the DHCP server daemon.
- `/etc/sysconfig/dhcpd`: contains parameters to be added to the DHCP server daemon when it's started.
- `/usr/sbin/dhcpd`: DHCP server daemon.
- `/usr/sbin/dhcpd-chroot.sh`: script to chroot the DHCP server.
- `/usr/sbin/dhcpd-conf-to-ldap.pl`: script to include all DHCP data in the LDAP directory
- `/usr/sbin/dhcpreport.pl`: script to show information about DHCP server data.
- `/var/lib/dhcp/dhcpd.leases`: file which stores the lease database.

8.1.1.5.2.2. Configuration du serveur DHCP

La configuration d'un serveur DHCP consiste principalement à écrire le fichier `/etc/dhcpd.conf`.

8.1.1.5.2.2.1. General Parameters

Here are the most common parameters:

`ddns-update-style`

This one is mandatory. It deals with dynamic DNS update. The server will not start without it. You should use it at least with the `none` value.

`authoritative`

This parameter is very important. It indicates that the DHCP server should send DHCP-NAK messages to badly configured clients. If this is not done, clients will be unable to get a correct IP address after changing subnets until their old lease has expired, which could take quite a long time. Also, if someone installs a DHCP server in the same network segment, it will not annoy dialogs with clients by sending DHCPNAK messages.

You can also consider as general parameters all network options that will be the same for all hosts and networks. DHCP option statements always start with the `option` keyword, followed by an option name, and followed by option data:

```
option <option_name> <data>
```

8.1.1.5.2.2.2. Host and Network Declarations

You now need to define specific hosts and/or networks for your DHCP policy. Let's take an example:

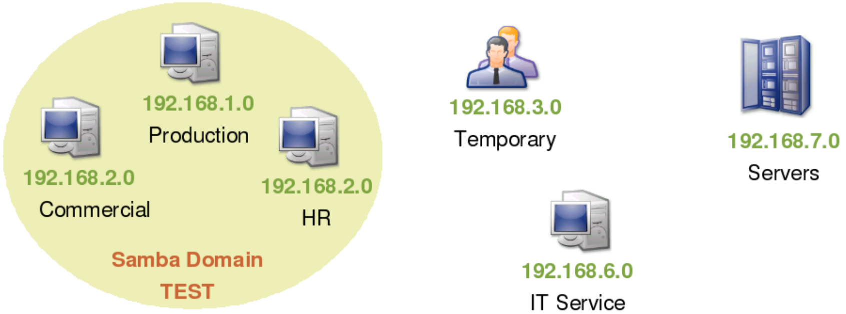

Figure 8-2. Applying Declarations for a Practical DHCP Configuration

We have basically 4 different networks types:

- Samba domain: `commercial`, `HR` and `Production` belong to this domain. Most workstations should have reserved IP so that they can always have the same one. Only `commercial` and `HR` services can access the Internet using the same gateway.
- Temporary users: have minimum rights on network resources. These workstations will have quite short leases. They also use their own gateway to access Internet.
- IT services: don't belong to the Samba domain, but they will have reserved IPs. They use their own gateway.
- Servers: same description as IT services.

The DHCP server allows you to simplify configuration files by using hosts and network declaration. Here are main types of declaration:

`host`

> The `host` declaration allows you to provide a fixed IP address for computers, based on MAC addresses. Then you can provide it not only with an IP address, but also with a hostname.

```
host
  <non_fqdn> { option host-name "<fqdn>"; hardware
  ethernet <MAC_address>; fixed-address <IP_address>;
  }
```

`group`

The `group` declaration allows you to group some host declarations and assign some common parameters, such as a DNS or NetBIOS server.

```
group <name> {
  <option_name> <data>;
  <option_name> <data>;
  ...

  host <non_fqdn> { ... }
  host <non_fqdn> { ... }
  ...
  }
```

`subnet`

`subnet` declarations allow you to provide a specific configuration for dynamic address attribution for a given network. You will have to fix the range of provided IP and network options.

```
subnet <network_address> netmask <netmask_address> {
  <option_name> <data>;
  <option_name> <data>;
  ...
 range <ip_address_min> <ip_address_max>
  }
```

`pool`

`pool` declarations allow you to provide specific parameters for subnetworks in a `subnet` declaration.

```
subnet <network_address>
      netmask <netmask_address> { <option_name> <data>; ...
      pool { <option_name> <data>; range <ip_address_min>
      <ip_address_max> } ...  }
```

Let's use available declarations and parameters to configure the DHCP server file for our example:

```
# cat /etc/dhcpd.conf
ddns-update-style none;
authoritative;
```

```
# common network options
# common domain name server
option domain-name              "domain.tst";
# common domain name server IP address
option domain-name-servers      192.168.7.1;
# common ntp server
option ntp-servers  192.168.7.2;

# Define Samba Domain computers

group sambanet {
# gateway IP address
option routers   192.168.2.254;
# netbios server IP address
option  netbios-name-servers 192.168.2.253;
# do not allow clients that have no host declaration
# to get an IP address
deny unknown-clients;

host com1 {
# fully qualified hostname
option host-name "com1.domain.tst";
# MAC address
hardware ethernet 00:A0:78:8E:9E:AA;
# provided fix IP address
fixed-address 192.168.2.1;
}
...
host hr1 {
option host-name "hr1.domain.tst";
hardware ethernet 00:A0:78:8E:9E:AA;
fixed-address 192.168.2.10;
}
...
}

group sambalone {
option  netbios-name-servers 192.168.2.253;
deny unknown-clients;

host prod1 {
option host-name "prod1.domain.tst";
hardware ethernet 00:A0:78:8E:9E:AA;
fixed-address 192.168.1.1;
}
...
}

subnet 192.168.3.0 netmask 255.255.255.0 {
range 192.168.3.1 192.168.3.200;
allow unknown-clients;
option routers 192.168.3.254;
}

group IT {
```

```
option routers  192.168.6.254;
deny unknown-clients;

host it1 {
option host-name "it1.domain.tst";
hardware ethernet 00:A0:78:8E:9E:AA;
fixed-address 192.168.6.1;
}
...
}

group servers {
option routers  192.168.7.254;
deny unknown-clients;

host dns {
option host-name "dns.domain.tst";
hardware ethernet 00:A0:78:8E:9E:AA;
fixed-address 192.168.7.1;
}
...
}
```

The DHCP server configuration can allow you to use multiple parameters. See dhcpd.conf, dhcp-options and dhcpd.leases man pages for examples.

8.1.1.5.2.3. Managing DHCP Server Service

ou can manage the DHCP service with with the `dhcp initscript` to start, stop, and restart the daemon. Use the following command:

```
service
    {start|stop|restart|status} dhcpd
```

You can also configure how the daemon is to be launched by using the `/etc/sysconfig/dhcpd` file at start-up. Here are its main parameters:

`INTERFACES`

It can limit which network interfaces `dhcpd` listens on. By default, `dhcpd` listens on all interfaces.

OPTIONS

You can add more options to configure the daemon:

- -q: does not print all license text when starting the server (used by default).
- -p: specifies the port on which the server should listen for requests (default is 67). It can be used to configure a DHCP relay server.
- -d: print all dhcpd messages on standard error output for debugging purpose.

8.1.1.5.3. Managing DHCP Data in OpenLDAP Directory

To host DHCP data in OpenLDAP, please follow these steps:

1. Import /etc/dhcpd.conf data into ou=dhcp.
2. Configure /etc/dhcpd.conf to use LDAP (with or without authentication)..

Please also read the README.ldap file in the documentation directory of dhcp-common (/usr/share/doc/dhcp-common/README.ldap).

8.1.1.5.3.1. Configuring dhcpd. conf

You can now remove most of the configuration from /etc/dhcpd.conf, leaving only the LDAP part. This results in the following file:

```
ldap-server "mes5.mandriva";
ldap-port 389;
ldap-username "uid=DHCP Reader,ou=System Accounts,dc=example,dc=com";
ldap-password "dhcpreader";
ldap-base-dn "ou=dhcp,dc=example,dc=com";
ldap-method dynamic;
```

In the example above we chose to use authenticated binds, but anonymous searches can also be used: just leave ldap-username and ldap-password out. After this last change, the DHCP server can be started and it will consult the LDAP tree.

8.1.1.5.3.2. Importing Data in LDAP Directory

The `dhcp-common` package has a contrib script which can be used to import an existing `/etc/dhcpd.conf` file into LDAP: `/usr/sbin/dhcpd-conf-to-ldap.pl`.

For this example, we will import the following configuration file:

```
ddns-update-style none;

     subnet 172.16.10.0 netmask 255.255.255.0 {
     option routers 172.16.10.1;
     option subnet-mask 255.255.255.0;

     option domain-name "example.com";

     option domain-name-servers 10.0.0.5;
     default-lease-time 21600;
     max-lease-time 43200;

     deny unknown-clients;

     host test009.example.com {
     hardware ethernet 00:C0:DF:02:93:71;
     fixed-address 172.16.10.5;
     }
     }
```

The command below creates the ldif file corresponding to our current `dhcpd.conf` configuration. Please note that this script has not yet been tested with all possible DHCP configuration scenarios. Always review the resulting LDIF file.

```
$ perl /usr/sbin/dhcp-common-3.0.3/contrib/dhcpd-conf-to-ldap.pl \
     --basedn "ou=dhcp,dc=example,dc=com" \
     --dhcpdn "cn=DHCP Config,ou=dhcp,dc=example,dc=com" \
     --conf /etc/dhcpd.conf --server mes5.example.com --ldif dhcpd.ldif
     Creating LDAP Configuration with the following options:
     Base DN: ou=dhcp,dc=example,dc=com
     DHCP DN: cn=DHCP Config,ou=dhcp,dc=example,dc=com
     Server DN: cn=mes5.example.com, ou=dhcp,dc=example,dc=com

     Done.
```

Here are the options we used:

- basedn: branch where DHCP information will be stored.
- dhcpdn: entry which will contain the configuration of our server.
- conf: `dhcpd.conf` file which will be migrated to LDAP.

- server: le FQDN of the DHCP server (should match the output of the `hostname` command).
- ldif: le fichier `ldif`.

`dhcpd.ldif` now has the data we will import. Let's verify:

```
dn: cn=mes5.example.com, ou=dhcp,dc=example,dc=com
cn: mes5.example.com
objectClass: top
objectClass: dhcpServer
dhcpServiceDN: cn=DHCP Config,ou=dhcp,dc=example,dc=com

dn: cn=DHCP Config,ou=dhcp,dc=example,dc=com
cn: DHCP Config
objectClass: top
objectClass: dhcpService
dhcpPrimaryDN: cn=mes5.example.com, ou=dhcp,dc=example,dc=com
dhcpStatements: ddns-update-style none
dn: cn=172.16.10.0, cn=DHCP Config,ou=dhcp,dc=example,dc=com
cn: 172.16.10.0
objectClass: top
objectClass: dhcpSubnet
objectClass: dhcpOptions
dhcpNetMask: 24
dhcpStatements: default-lease-time 21600
dhcpStatements: max-lease-time 43200
dhcpStatements: deny unknown-clients
dhcpOption: routers 172.16.10.1
dhcpOption: subnet-mask 255.255.255.0
dhcpOption: domain-name "example.com"
dhcpOption: domain-name-servers 10.0.0.5

dn: cn=test009.example.com, cn=172.16.10.0, cn=DHCP Config,ou=dhcp,
 dc=example,dc=com
cn: test009.example.com
objectClass: top
objectClass: dhcpHost
dhcpHWAddress: ethernet 00:c0:df:02:93:71
dhcpStatements: fixed-address 172.16.10.5
```

This data can now be imported. We will use the DHCP admin account to do so:

```
$ ldapadd -x -D "uid=DHCP Admin,ou=System Accounts,dc=example,dc=com"
 -W -f dhcpd.ldif
Enter LDAP Password: secretpass
adding new entry "cn=mes5.example.com, ou=dhcp,dc=example,dc=com"

adding new entry "cn=DHCP Config,ou=dhcp,dc=example,dc=com"

adding new entry "cn=172.16.10.0, cn=DHCP Config,ou=dhcp,dc=example,
 dc=com"

adding new entry "cn=test009.example.com, cn=172.16.10.0, cn=DHCP
```

```
Config,ou=dhcp,dc=example,dc=com"
```

8.2. File and Print Services

8.2.1. Sharing Printers with CUPS

This chapter is based on CUPS official documentation.

CUPS (Common UNIX Printing System) provides a portable printing layer for UNIX®-based operating systems. It has been developed by Easy Software Products to promote a standard printing solution for all UNIX® vendors and users. CUPS provides the System V and Berkeley command-line interfaces.

8.2.1.1. How Does it Work?

The first time you print to a printer, CUPS creates a queue to keep track of the current status of the printer (everything OK, out of paper, etc.) and any pages you have printed. Most of the time, the queue points to a printer connected directly to your computer via a USB or parallel port. However, it can also point to a printer on your network, a printer on the Internet, or multiple printers depending on the configuration. Regardless of where the queue points, it will look like any other printer to you and your applications.

Every time you print something, CUPS creates a job which contains the queue you are sending the print to, the name of the document you are printing, and the page descriptions. Jobs are numbered (queue-1, queue-2, and so forth) so you can monitor the job as it is printed or cancel it if you see a mistake. When CUPS gets a job for printing, it determines the best programs (filters, printer drivers, port monitors, and backends) to convert the pages into a printable format and then runs them to actually print the job.

When the print job is completely printed, CUPS removes the job from the queue and moves on to any other jobs you have submitted. You can also be notified when the job is finished, or if there are any errors during printing, in several different ways.

CUPS uses the Internet Printing Protocol ("IPP") as the basis for managing print jobs and queues. The Line Printer Daemon ("LPD") Server Message Block ("SMB"), and AppSocket (a.k.a. JetDirect) protocols are also supported with reduced functionality. CUPS adds network printer browsing and PostScript Printer Description ("PPD") based printing options to support real-world printing under UNIX.

CUPS includes an image fileRIP that supports the printing of image files to non-PostScript printers. Sample drivers that use these filters for Dymo, EPSON, HP, and OKIDATA printers are included.

Mandriva Enterprise Server 5 provides CUPS 1.3, with new functions. Here are some of them:

- Network Printer Discovery: CUPS can now find printers on the LAN using SNMP
- LDAP Support: CUPS now supports printer sharing via the Lightweight Directory Access Protocol, version 3
- Export Printers to Samba: the administration page now offers an Export Printers to Samba button which allows administrators to export printer drivers to Microsoft clients via Samba
- Set allowed users: you can now set the list of users and/or groups that are allowed or not allowed to access a printer or class
- Job management: you can cancel all jobs for a printer, move it for another printer
- CUPS API: security and performance have been improved.
- IPP support improvements
- Scheduler improvements

For more information, visit the CUPS official web site: http://cups.org/ (`http://cups.org/`)

8.2.1.2. CUPS installation and file tree

Here are the packages to install to have full CUPS server and the main ppd files to support most printers.

- `cups`: provides all server files
- `cups-drivers`: contains special printer drivers to used with CUPS and their appropriate PPD files
- `hplip-hpijs-ppds`: PPD files for the HPIJS printer driver
- `postscript-ppds`: contains PPD files for older PostScript printers
- `hplip`: an HP driver package to provide Linux support for most Hewlett-Packard DeskJet, LaserJet, PSC, OfficeJet, and PhotoSmart printers and all-in-one peripherals

Now Let's describe the CUPS tree:

- /etc/cups: configuration files such as printers.conf. Overridden by the ServerRoot directive in cupsd.conf.
- /usr/include: CUPS included files.
- /usr/lib: CUPS library files.
- /usr/lib/cups: server programs such as backends and filters. Overridden by the ServerBin directive in cupsd.conf.
- /usr/share/cups: data files such as fonts. Overridden by the DataDir directive in cupsd.conf.
- /usr/share/cups/doc: documentation files. Overridden by the DocumentRoot directive in cupsd.conf.
- /usr/share/locale: localization files.
- /var/cache/cups: cache files such as ppds.dat and remote.cache. Overridden by the CacheDir directive in cupsd.conf.
- /var/log/cups: access_log, error_log, and page_log files. Overridden by the AccessLog, ErrorLog, PageLog, directive in cupsd.conf.
- /var/run/cups: domain socket file and state data such as authentication certificates. Overridden by the StateDir directive in cupsd.conf.
- /var/spool/cups: spooled print jobs. Overridden by the RequestRoot directive in cupsd.conf.

8.2.1.3. Configuring your CUPS server

You can configure CUPS both through configuration file editing and web interface. We recommend you to use web interface for basic configuration and printers management (add, modify, delete) and editing configuration file for advanced purpose, mainly for security.

8.2.1.3.1. Using the CUPS web interface

You can access web interface using https and port 631: https://cups_server:631 (https://cups_server:631).

The Jobs tab allows you to manage print jobs. The Administration tab will give you a web interface to add, modify and delete printers. Simply click on the Add printers or Manage printers buttons.

cupsd is configured by default to show printers shared by other systems and only allow local access to the system and its printers. Administration operations require Basic authentication with membership in the group "sys". Connections are accepted via domain socket (`/var/run/cups/cups.sock`) or "localhost" (127.0.0.1). Users are not allowed to cancel jobs that do not belong to them.

You can change these parameters in the Administration tab in the Server section: click on boxes to choose which option you want to enable or disable.

The server section will also provide access to server logs:

- access and error logs: these contain access and errors regarding CUPS server
- page log: this log contains all access to CUPS' web interface

Let's add a new printer. CUPS 1.2 provides an auto-detect functionality. If printers are up and connected to the network, CUPS should auto detect them, then the new New Printers Found: list will appear. Click on Add this printer and fill the fields with printer information. That's it!

Your printer may not be recognized by CUPS because its ppd file is not in list.

1. In the Administration tab click on Add printer.
2. On the first screen, fill the name field for this printer and location.
3. On the second screen, choose from the list the device to access your printer.
4. On the third screen, fill the exact URI to access your printer.
5. On the last screen, you will be asked to choose a printer in the list. Since you can't find it here, let's provide a ppd file for the printer. You will find it on manufacturer's CD. Click on the Browse button and select this file from tree. That's it!

This procedure may not work since some ppd files are not properly completed. Please have a look at http://linuxprinting.org (`http://linuxprinting.org`) to check for your hardware's compatibility.

Once your printers are added, you can modify their configuration. Click on Printers. Select a printer by clicking on the one you want to modify. You will then see the following:

- Print Test Page: print test a page for configured printer
- Stop Printer: completely disable a given printer

- Reject Jobs: the printer will not be disabled but it will not accept any more jobs.
- Move all Jobs: move all jobs from this printer to a new one you will choose
- Cancel all Jobs: cancel all jobs for this printer
- Unpublish Printer: hide this printer from users
- Modify Printer: modify printer configuration, URI access, driver, etc. ...
- Set Printer Options: modify printing option such as resolution, page size, banner...
- Delete Printer: delete a printer completely
- Set as Default: set this printer as the default printer for users
- Set allowed Users: list all users that can use or not this printer

8.2.1.3.2. Advanced configuration through configuration file edition

You will find configuration files in `/etc/cups`. `/etc/cups/client.conf` enables you to configure client access. `/etc/cups/cupsd.conf` enables you to configure the CUPS server.

The `/etc/cups/cupsd.conf` file contains configuration directives that control how the server functions. Each directive is listed on a line by itself followed by its value. Comments are introduced using the number sign ("#") character at the beginning of a line.

General syntax of `cupsd.conf` looks like the one for Apache. You can specify general parameters using `<directive> <parameter>`.

Let's have a look at the main configuration options:

`AuthType`

The AuthType directive defines the type of authentication to perform:

- `None`: No authentication should be performed (default)
- `Basic`: Basic authentication should be performed using the UNIX password and group files
- `Digest`: Digest authentication should be performed using the `/etc/cups/passwd.md5` file
- `BasicDigest`: Basic authentication should be performed using the `/etc/cups/passwd.md5` file

When using Basic, Digest, or BasicDigest authentication, clients connecting through the localhost interface can also authenticate using certificates. The AuthType directive must appear inside a Location or Limit section.

BrowseAddress

The BrowseAddress directive specifies an address to send browsing information to. Multiple BrowseAddress directives can be specified to send browsing information to different networks or systems.

The @LOCAL name will broadcast printer information to all local interfaces. The @IF(name) name will broadcast to the named interface.

There is no default browse address.

BrowseAllow

The BrowseAllow directive specifies a system or network to accept browse packets from. The default is to accept browse packets from all hosts. Host and domain name matching require that you enable the HostNameLookups directive.

IP address matching supports exact matches, partial addresses that match networks using netmasks of 255.0.0.0, 255.255.0.0, and 255.255.255.0, or network addresses using the specified netmask or bit count. The @LOCAL name will allow browse data from all local interfaces. The @IF(name) name will allow browsing from the named interface.

You can also use the BrowseDeny directive to deny packets. It works the same way as BrowseAllow

Browsing

The Browsing directive controls whether or not network printer browsing is enabled. The default setting is On. This directive does not enable sharing of local printers by itself; you must also use the BrowseAddress or BrowseProtocols directives to advertise local printers to other systems.

DefaultShared

The DefaultShared directive specifies whether printers are shared (published) by default. The default is yes.

JobRetryInterval

The JobRetryInterval directive specifies the number of seconds to wait before retrying a job. This is typically used for fax queues but can also be used with normal print queues whose error policy is retry-job. The default is 30 seconds.

CUPS allows you also to fix access rights to server:

```
<Location />
     Order Deny,Allow
     Deny From All
     Allow From 127.0.0.1
     Allow From @LOCAL
     </Location>
```

In this example, we allow users to access CUPS' home web interface from localhost and from the same LAN as the CUPS server. Using this method, you can configure admin access.

`Location`

The `Location` directive specifies access control and authentication options for the specified HTTP resource or path. The `Allow`, `AuthType`, `Deny`, `Encryption`, `Limit`, `LimitExcept`, `Order`, `Require`, and `Satisfy` directives may all appear inside a location.

Here are common locations on CUPS server:

- `/`: the path for all get operations (get-printers, get-jobs, etc.)
- `/admin`: the path for all administration operations (add-printer, delete-printer, start-printer, etc.)
- `/admin/conf`: path to access to the CUPS configuration files (cupsd.conf, client.conf, etc.)
- `/admin/log`: path for access to the CUPS log files (access_log, error_log, page_log)
- `/classes`: path for all classes
- `/classes/name`: resource for class name
- `/jobs`: path for all jobs (hold-job, release-job, etc.)
- `/printers`: path for all printers
- `/printers/name`: path for printer name

More specific resources override the less specific ones. So the directives inside the `/printers/name` location will override the ones from `/printers`. Directives inside `/printers` will override the ones from `/`. None of the directives are inherited.

Allow

```
   <Location /path>
 ...
 Allow from All
 Allow from None
 Allow from *.domain.com
 Allow from .domain.com
 Allow from host.domain.com
 Allow from nnn.*
 Allow from nnn.nnn.*
 Allow from nnn.nnn.nnn.*
 Allow from nnn.nnn.nnn.nnn
 Allow from nnn.nnn.nnn.nnn/mm
 Allow from nnn.nnn.nnn.nnn/mmm.mmm.mmm.mmm
 Allow from xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx
 Allow from @LOCAL
 Allow from @IF(name)
 </Location>
```

The `Allow` directive specifies a hostname, IP address, or network that is allowed access to the server. `Allow` directives are cumulative, so multiple `Allow` directives can be used to allow access for multiple hosts or networks. The `/mm` notation specifies a CIDR netmask.

The `@LOCAL` name will allow access from all local interfaces. The `@IF(name)` name will allow access from the named interface.

8.2.1.4. Securing your CUPS server

In the default "standalone" configuration, there are few potential security risks - the CUPS server does not accept remote connections, and only accepts shared printer information from the local subnet. When you share printers and/or enable remote administration, you expose your system to potential unauthorized access. This help page provides an analysis of possible CUPS security concerns and describes how to better secure your server.

Authentication Issues

When you enable remote administration, the server will use Basic authentication for administration tasks. The current CUPS server supports Basic, Digest, and local certificate authentication:

- Basic authentication essentially places the clear text of the username and password on the network. Since CUPS uses the system username and password account information, the authentication information could be used to gain access to possibly privileged accounts on the server.

You should enable encryption to hide the username and password information - this is the default on MacOS X and systems with GNU TLS or OpenSSL installed.

- Digest authentication uses an MD5 checksum of the username, password, and domain ("CUPS"), so the original username and password is not sent over the network. The current implementation does not authenticate the entire message and uses the client's IP address for the nonce value, making it possible to launch "man in the middle" and replay attacks from the same client.

  You should enable encryption to hide the username and password information.

- Local certificate authentication passes 128-bit "certificates" that identify an authenticated user. Certificates are created on-the-fly from random data and stored in files under `/var/run/cups/certs`. They have restricted read permissions: root + system-group(s) for the root certificate, and lp + lp for CGI certificates. Because certificates are only available on the local system, the CUPS server does not accept local authentication unless the client is connected to the loopback interface (127.0.0.1 or ::1) or domain socket.

  Ensure that unauthorized users are not added to the system group(s).

Denial of Service Attacks

When printer sharing or remote administration is enabled, the CUPS server, like all Internet services, is vulnerable to a variety of denial of service attacks:

- Establishing multiple connections to the server until the server will accept no more.

  This cannot be protected against by any known software. The `MaxClientsPerHost` directive can be used to configure CUPS to limit the number of connections allowed from a single host, however that does not prevent a distributed attack.

  You should limit access to trusted systems and networks.

- Repeatedly opening and closing connections to the server as fast as possible.

  There is no easy way of protecting against this in the CUPS software. If the attack is coming from outside the local network, it may be possible to filter such an attack. However, once the connection request has been received by the server it must at least accept the connection to find out who is connecting.

- Flooding the network with broadcast packets on port 631.

  It might be possible to disable browsing if this condition is detected by the CUPS software, however if there are large numbers of printers available on the network such an algorithm might think that an attack was occurring when instead a valid update was being received.

  You should block browse packets from foreign or untrusted networks using a router or firewall.

- Sending partial IPP requests; specifically, sending part of an attribute value and then stopping transmission.

  The current code will wait up to 1 second before timing out the partial value and closing the connection. This will slow the server responses to valid requests and may lead to dropped browsing packets, but will otherwise not affect the operation of the server.

  You should block IPP packets from foreign or untrusted networks using a router or firewall.

- Sending large/long print jobs to printers, preventing other users from printing.

  There are limited facilities for protecting against large print jobs (the `MaxRequestSize` attribute), however this will not protect printers from malicious users and print files that generate hundreds or thousands of pages.

  You should restrict printer access to known hosts or networks, and add user-level access controls as needed for expensive printers.

Encryption Issues

CUPS supports 128-bit SSL 3.0 and TLS 1.0 encryption of network connections via the OpenSSL, GNU TLS, and CDSA encryption libraries. In addition to the potential security issues posed by the SSL and TLS protocols, CUPS currently has the following issues.

Certification validation/revocation; currently CUPS does not validate or revoke server or client certificates when establishing a secure connection. This can potentially lead to "man in the middle" and impersonation/spoofing attacks over unsecured networks. Future versions of CUPS will support both validation and revocation of server certificates.

Do not rely on encryption for security when connecting to servers over the Internet or untrusted WAN links.

8.2.1.5. Browsing printers in LDAP directory

CUPS 1.2 allows you to browse printers in a LDAP directory so that you can add it as new printers. As many other services, you will need to specify the way CUPS server can ask LDAP directory, in `/etc/cups/cupsd.conf`:

BrowseLDAPBindDN

The `BrowseLDAPBindDN` directive specifies the LDAP domain name to use when listening for printer registrations. The default is undefined.

BrowseLDAPDN

The `BrowseLDAPDN` directive specifies the LDAP domain name to use when registering local shared printers. The default is undefined

BrowseLDAPPassword

The BrowseLDAPPassword directive specifies the access password to use when connecting to the LDAP server. The default is undefined.

BrowseLDAPServer

The `BrowseLDAPServer` directive specifies the name of the LDAP server to connect to. The default is undefined.

Let's write an example:

```
BrowseLocalProtocols ldap
BrowseRemoteProtocols ldap

BrowseLDAPServer localhost
BrowseLDAPDN ou=printers,dc=example,dc=com
BrowseLDAPBindDN uid=Manager,dc=example,dc=com
BrowseLDAPPassword password
```

8.2.2. Sharing Files with NFS

NFS allows you to export whole directories, even whole filesystems, through the network, therefore enabling file sharing between users. This sharing type is simple to put in place and is essentially used with GNU/Linux and UNIX® systems. NFS is not recommended if you have strong security constraints. Using NFS shouldn't be used often on a local network.

8.2.2.1. Installing a NFS Server

The installation is simple and only requires the `nfs-utils` package.

8.2.2.2. Configuring a NFS Server

The configuration of NFS filesystem exports is done in the `/etc/exports` file. It allows you to establish the access mode to data, as well as rights given to users and machines. It's mandatory to supervise rights assignment to secure data access.

8.2.2.3. Using NFS v4

Mandriva Enterprise Server 5 offers a NFS v4-based server. This new version brings great enhancements security-wise, and the following features:

- Addition of file states concerning locking, read-write, between clients and the server.
- Lease base for file locking allowing clients to recuperate file ownership during lease time. The client must contact the server if he wants to extend the length of the lease.
- Addition of security components such as Kerberos 5 and SPKM3.
- Extension of support for ACL files, notably adding group and user names, allowing this type of direct access.
- Combination of many NFS protocols, allowing better management by firewalls.
- Replication support.
- Client capacity to maintain sessions or to recuperate them even if the server crashes or if there's a network crash.
- Management of a "pseudo" filesystem allowing you to manage all NFS exports from a common root.

Here's a procedure allowing you to put in place a NFS exports tree structure with the help of Kerberos 5. The prerequisites: NFS needs to be installed, and you need an operational Kerberos server.

1. Declaring the Pseudo Filesystem in `/etc/fstab`

   Just add the two following lines:

   ```
   rpc_pipefs /var/lib/nfs/rpc_pipefs rpc_pipefs
         defaults 0 0 nfsd /proc/fs/nfsd nfsd defaults 0
         0
   ```

Then, as for a classic filesystem, mount it:

```
# mount rpc_pipesfs
```

2. Creating the Root of the Exports' Tree Structure

   Since NFS v4 provides a new, pseudo filesystem, all exports are positioned starting at the tree structure's root:

```
/exports/
      |-- test1
      |-- ...
      `-- testn
```

   To create it, specify this parameter: `fsid=0`:

```
# cat /etc/exports
 /export  *(rw,fsid=0,insecure,no_subtree_check,5)
```

3. Creating the Export's Tree Structure

   Then, the tree structure follows the classic scheme:

```
# cat /etc/exports
 /export  *(rw,fsid=0,insecure,no_subtree_check)
 /export/test1  *(rw,nohide,insecure,no_subtree_check)
 ...
```

   Mounting the resource is done in relation to the NFS root:

```
# mount nfs4srv:/test1 /mnt/nfs
```

4. Using Kerberos

   First of all, create a principal for NFS. Also create a key for the server and client using the `kadmin` shell:

```
kadmin.local:  addprinc -randkey nfs/nfs4srv.edge-it.subnet
WARNING: no policy specified for nfs/nfs4srv.example.com@EXAMPLE.COM;
        defaulting to no policy
Principal "nfs/nfs4srv.example.com@EXAMPLE.COM" created.

kadmin.local:  ktadd  -e des-cbc-crc:normal nfs/nfs4srv.example.com
Entry for principal nfs/nfs4srv.example.com with kvno 3, encryption
        type Triple DES cbc mode with HMAC/sha1 added to keytab
 WRFILE:/etc/krb5.keytab.
Entry for principal nfsnfs4srv.example.com with kvno 3, encryption
      type DES cbc mode with CRC-32 added to keytab
      WRFILE:/etc/krb5.keytab.
```

   Copy file `krb5.keytab` on the NFS server and on clients using NFSv4.

   All that's left is to create the exports using `krb5`, with a specific client: `gss/krb5`.

```
# cat /etc/exports
/export gss/krb5(rw,fsid=0,insecure,no_subtree_check,sync,5)
```

   Use the mount command with the `sec=kbr5` option:

```
    # mount -t nfs4 -o sec=krb5 nfs4srv:/ /nnt/nfs
```

8.2.3. FTP File Server

ProFTPD allows you to create and configure a FTP server. You can configure it through a specific Webmin module located in the Servers category.

8.2.3.1. Installation and Tree Structure of ProFTPD

Two packages exist for `proftpd`:

- `proftpd`: package containing the whole ProFTPD server;
- `proftpd-anonymous`: module which activates anonymous connections to the FTP server.

For security reason, we will only install the first package.

```
# urpmi proftpd
```

The tree structure provided by `proftpd` is relatively simple:

- `/etc/proftpd.conf`: ProFTPD's configuration file;
- `/etc/xinetd.d/proftpd-xinetd`: configuration of ProFTPD's daemon launched via `xinetd` (not recommended);
- `/etc/rc.d/init.d/proftpd`: ProFTPD's initscript configured in autonomous mode;
- `/var/log/proftpd`: ProFTPD's log file directory.

8.2.3.2. ProFTPD Toolbox

To verify the configuration file's syntax:

```
# proftpd -t
    Checking syntax of configuration file
    Syntax check complete.
```

To verify that the server works properly:

```
# service proftpd status
    proftpd (pid 32445) is running ...
    # telnet 192.168.40.52 21
    Trying 192.168.40.52...
```

```
Connected to tellure.edge-it.subnet (192.168.40.52).
Escape character is '^]'.
220 ProFTPD 1.2.10 Server (ProFTPD Default Installation)
[192.168.40.52]
```

You should also test a complete session by connecting to the server.

The `proftpd` package installs, not counting the server daemon, a certain number of useful commands to monitor your server's state.

ftpcount

Allows you to count the number of connections on the FTP server at a given time:

```
$ ftpcount
Master proftpd process 32445:
Service class                      -   1 user
```

ftptop

Enables you to view in real time the FTP server's activity, and its connections:

```
ftptop/0.9: Thu Jan  5 11:54:49 2009, up for 1 min
1 Total FTP Sessions: 0 downloading, 0 uploading, 1 idle

PID   S USER     CLIENT               SERVER          TIME COMMAND
32455 I anne     test.domain.subne 0.0.0.0:21      0m47s  idle
```

ftpwho

Allows you to view active connections on the server at a given time, and information concerning running sessions:

```
$ ftpwho -v
standalone FTP daemon [32445], up for  2 hrs 18 min
  542 anne     [ 0m11s]  0m11s idle
        client: workstation.edge-it.subnet [192.168.40.140]
        server: 0.0.0.0:21 (ProFTPD Default Installation)
        location: /
Service class                      -   1 user
```

8.2.3.3. Securing a Proftpd Server

To secure a FTP server, you need to first to set up an efficient firewall. Apart that, securing your server depends on user management and your configuration's customization.

Here's a typical configuration file which we'll detail later on:

```
# cat /etc/proftpd.conf
    ServerName                      "FTP SERVER"
ServerType                      standalone
DeferWelcome                    off
TransferLog                     /var/log/proftpd.xferlog
DefaultRoot                     ~
RequireValidShell               off
ServerIdent                     off
RootLogin   off

ShowSymlinks                    off
DefaultServer                   on
AllowOverwrite                  off

TimeoutNoTransfer               600
TimeoutStalled                  600
TimeoutIdle                     1200

DisplayLogin                    /etc/welcome.msg
DisplayFirstChdir               .message

DenyFilter                      \*.*/
Bind                            192.168.10.55
```

8.2.3.3.1. User Management

A great part of securing your FTP server is linked to how you manage your users. Here are critical steps.

Users without Shells

It's preferable to give FTP access to users with no shell. Here's how to create them:

```
# useradd -s /bin/false user
```

or how to modify them:

```
# usermod -s /bin/false user
```

To authorize access to the FTP server to users with no shell, add in `/etc/proftpd.conf`:

```
RequireValidShell  on
```

User Sessions in chroot Mode

We lock users in a fixed directory, preventing them from going up in the tree structure. To do so, add the following line in `/etc/proftpd.conf`:

```
 DefaultRoot                    ~
```

Prohibit FTP server access to root

Since passwords circulate in clear text, the root password could intercepted by someone. To prevent that, add the following line to `/etc/proftpd.conf`:

```
RootLogin off
```

Prohibit anonymous connections

This is done by default since implementing this functionality demands the installation of an additional package (`proftpd-anonymous`).

Limit connections to a given list of users

The `Limit` instruction allows to specify authorized (or not) users and groups to connect to the server.

```
<Limit LOGIN>
 AllowUser workstation
 AllowGroups devels
 DenyAll
</Limit>
```

8.2.3.3.2. Securing the Server's Configuration

All points mentioned hereafter are relative to modifications in the `/etc/proftpd.conf` file.

ProFTPD uses root privileges only when it's necessary. In the opposite case, it uses the identity defined in the configuration. Steps which need root privileges are:

- accessing ports less than or equal to 1024;
- determining limitations on resources;
- reading configuration information;
- executing code portions associated with the network.

Hiding the banner

This consists in not displaying information about the FTP server type and version:

```
ServerIdent                     off
```

To verify this:

```
# telnet 192.168.40.52 21
 Trying 192.168.40.52...
 Connected to tellure.edge-it.subnet (192.168.40.52).
 Escape character is '^]'.
 220 192.168.40.52 FTP server ready
```

Modifying the default access port

Define a port higher than 1024 (accessible to a non-root user). You can also plan to use an `iptables` command to render this manipulation transparent.

Modifying default messages

The ProFTPD server sends a certain number of messages during the different steps of an FTP session. Some of them can be modified to communicate with the user (security rules reminder, rights, etc.) or to hide messages which could highlight the server type or its version.

Here's the list of instructions corresponding to those messages:

- `DisplayConnect <filename>`: message displayed before the authentication procedure;
- `DisplayFirstChdir <filename>`: message displayed during the first directory change;
- `DisplayLogin <filename>`: message displayed during login;
- `DisplayGoAway <filename>`: message displayed during a refused connection;
- `DisplayQuit <filename>`: message displayed at the end of an FTP session;
- `ServerName <text>`: string displayed during login messages.

To these instructions we can add `DeferWelcome` which, when activated, doesn't display a welcome message when a user authenticates succesfully.

Establishing timeout

These allow you to avoid having to maintain open connections which are not used. There are a certain number of levels on which we can fix timeouts:

- `TimeoutNoTransfer <seconds>`: maximum number of seconds during which an authenticated client can be connected but inactive (by default: 300);
- `TimeoutStalled <seconds>`: maximum number of seconds during which an FTP connection can be in stalled state (by default: 3600);
- `TimeoutIdle <seconds>`: maximum number of seconds during which an FTP connection can be in idle state (by default: 600).

Controlling commands passed by users

The `Limit` instruction allows you to precisely list authorized commands for the FTP server users. These commands can be specified one by one, or by a group of commands.

Individual commands:

- `CWD` (Change Working Directory): change directory;
- `MKD/XMKD` (MaKe Directory): create a directory;
- `RNFR` (ReName FRom), `RNTO` (ReName TO): rename a directory;
- `DELE` (DELEte): destroy a file;
- `RMD` / `XRMD` (ReMove Directory): remove a directory;
- `RETR` (RETRieve): transfer a file from the server to the client;
- `STOR` (STORe): transfer a file from a client to the server;

Groups of commands

- `READ ALL FTP`: commands concerning file reading (doesn't concern the listing of a directory): `RETR`, `SITE`, `SIZE`, `STAT`;
- `WRITE ALL FTP`: commands concerning writing, creating and deleting a directory: `APPE`, `DELE`, `MKD`, `RMD`, `RNTO`, `STOR`, `XMKD`, `XRMD`;
- `DIRS ALL FTP`: commands concerning the display directory files: `CDUP`, `CWD`, `LIST`, `MDTM`, `NLST`, `PWD`, `RNFR`, `XCUP`, `XCWD`, `XPWD`;
- `ALL ALL FTP`: command identical to `READ WRITE DIRS`.

Example:

Limit commands users can do by taking away all file and directory deleting rights in the FTP repository:

```
<Limit RNFR DELE RMD>
 DenyAll
 </Limit>
```

Managing network interfaces

In the case where the machine has many interfaces or may IP addresses, we recommend that you link the server to a unique address.

```
Bind  192.168.10.55
```

Correcting an FTP security breach

The following line allows you to prevent a breach which could be provoked by the `NLST /../*/../*/../*/../*/../*/../*/../*/../*/../*/../` command:

```
DenyFilter  \*.*/
```

8.2.3.3.3. Customizing Access Rules to Files According to Needs

Proftpd allows you to customize the server configuration according to well-defined contexts:

- directories
- virtualhosts (this notion in ProFTPD is very similar to the one found in Apache).

Let's take the case of directories. The FTP repository is located in `/var/lib/ftp`:

- `/var/lib/ftp/datas`: write-accessible application data for developer and in read-only for visitors;
- `/var/lib/ftp/pub`: repository accessible by all users to place files, but no deletions possible;

The corresponding configuration would be:

```
<Directory /var/lib/ftp/datas>
     <Limit WRITE DIRS>
     DenyGroup visitors
     AllowGroup devels
     </Limit>
     </Directory>
     <Directory /var/lib/ftp/pub>
```

```
<Limit ALL>
AllowGroup visitors devels
</Limit>
</Directory>
```

8.2.3.4. Using LDAP Authentication on ProFTPD

By default, Proftpd contains a module allowing you to authentify users of your FTP server from an OpenLDAP directory.

The configuration is relatively simple. We'll start from the following configuration:

```
LDAPServer <ldap_server_ip_address> LDAPDNInfo
    "cn=Manager,dc=example,dc=com" "<password>" LDAPQueryTimeout 5
    LDAPDoAuth on "dc=example,dc=com" LDAPDoUIDLookups on
    "ou=Personnes,dc=example,dc=com" LDAPDoGIDLookups on
    "ou=Groups,dc=example,dc=com" LDAPNegativeCache off
    LDAPHomedirOnDemand off LDAPDefaultAuthScheme MD5
```

The configuration is based on parameters, specifying the essential data to access the LDAP server and interrogate it:

LDAPDNInfo

Holds the DN information for the initial contact to the directory.

LDAPQueryTimeout

Sets the timeout on LDAP requests.

LDAPDoAuth

Authorizes LDAP authentication.

LDAPDoUIDLookups

Sets the default GID to be assigned to users when the `uidNumber` attribute is not found.

LDAPDoGIDLookups

Authorizes LDAP resolution for group rights, and GID to be able to read directories.

LDAPHomedirOnDemand

Forces all LDAP-authentified users to use `HomeDironDemand` by default.

LDAPDefaultAuthScheme

Sets the authentication hash to be used when {hashname} is not specified.

8.2.4. Samba File and Print Server

This section presents the installation of a Samba server (principal domain controller, authentication and authorizations, file and print server) based on an LDAP directory to manage users. It describes the installation method, elements specified for specific services as well as the configuration put in place. You'll also find notes on basic tools used to test the service and to add the first users.

8.2.4.1. General Concepts and Web Reference

Announced in January 1992 by Andrew Tridgell, a student at the computer labs of the Australia National University, nbserver was only supposed to allow the mounting of Windows® shares on a UNIX machine. Since then it has become a complete set of tools to assure network resource management in heterogeneous environments. Nbserver changed its name to SaMBa (1.6.05) in April 1994 to evoke the implemented protocol. From version to version, its behavior got closer to a Windows® NT4 server (domain control, management of authentication to resources, resource sharing, SMB network path, etc.). SaMBa 2.0 was released in January 1999, and SaMBa 3.0 which finally allows you to manage Windows® user groups, was released in September 2003.

SaMBa 4 announces the possibility to put in place an Active Directory-type server.

Samba is a software suite which allows the interconnection of different systems around a common protocol called NetBIOS (*Network Basic Input Output System*) on which are based SMB (*Server Message Block*) and CIFS (*Common Internet File System*) to assure resource sharing (files, printers, serial and parallel ports).

The main goal of he Windows® SMB protocol is to share files. In Samba that goal is enhanced by:

- Determining the presence of other servers by using this protocol on the network (Network Browsing);
- Printing on the network;
- Authentication to access shares, directories and files;
- Lock of files in use;

- Notification of changes made on files and directories;
- Note of the protocol version to use (Dialect);
- Attribute management of extended files;
- Unicode support;
- File-lock management.

Authentication allows notably to put in place Microsoft NT4-type domains.

Main Web references:

- Official Samba Website (`http://www.samba.org/`);
- Samba Documentation (`http://samba.org/samba/docs/`);
- Collection of Samba HOWTOs (`http://samba.org/samba/docs/man/Samba-HOWTO-Col`
- Samba By Example (`http://samba.org/samba/docs/man/Samba-Guide/`);
- smb.conf (`http://samba.org/samba/docs/man/manpages-3/smb.conf.5.html`);
- Samba Wiki (`http://wiki.samba.org/index.php/Main_Page`).

8.2.4.2. Installation and Configuration

At the end of this part, a PDC (Primary Domain Controller) server will be put in place. Users, groups and machines will be stored in an LDAP directory. We take for granted that such a directory is already functional and it's possible to add the necessary data to it. But before entering this part, it's imperative that you install the Samba package and test a simple configuration.

8.2.4.2.1. Necessary Packages

- `samba-server`: holds the `smbd` daemons (authentication and management of share access) and `nmbd` (dialog);
- `network`: name resolution, tole management. It's the basis of a Microsoft NT4 server-type service.
- `samba-client`: contains the clients which allow it to access remote SMB/CIFS resources (needs `mount-cifs` to mount a remote CIFS-type filesystem).
- `samba-winbind`: allows third-party software (PAM, domain member Samba server, Squid, Apache) to authentify themselves on a domain server (Samba or Microsoft Windows®). Such an architecture may need `nss_wins` (for PAM).

- samba-smbldap-tools: this package doesn't contain the whole set of smbldap-tools, but only the /usr/share/samba/scripts/migrate-smbldap script;
- samba-swat: Samba's Web configuration interface;
- samba-doc: all of Samba's documentation;
- samba-vscan-*: every VFS enabling you to scan for viruses on the fly (needs a configured server);
- smbldap-tools: this package contains every smbldap-tools scripts which facilitate data manipulation (users, groups, machines) in the context of an authentified Samba server on an LDAP directory.

Here's Samba's tree structure:

/etc/samba

Holds all of the server's configuration files, essentially samba.conf for general configuration and shares configuration, as well as smb-winbind.conf to configure Winbind.

/usr/lib/samba/vfs

This directory lists every VFS (Virtual File System) module available on your server. In the current version, you have:

- audit: audit tool allowing an exhaustive verification of filesystem access via defined data shares.
- default_quota: enables you to establish by default quotas for users or groups;
- extd_audit: same as audit but with a different log conservation method;
- recycle: allows you to put in place network trash;
- netatalk: permits you to manage Apple compatibility on resource shares;
- vscan: allows you to plug an anti-virus to do on-the-fly scans on one or many shares;

You will find more information in the official documentation: VFS modules in a Samba server (http://www.samba.org/samba/docs/man/Samba-HOWTO-Collection/VFS.html).

`/var/cache/samba`

This directory contains all cache memories in `*.tdb` format files. They contain information such as path to resources, the Netbios names, logins, etc. These files can be the origin of malfunctions when, for example, the cache memory stays unchanged while modifications have been done.

`/var/log/samba`

You will find all of Samba's logs. Be careful, this directory can rapidly reach a gigantic size, especially if you work on a network which contains many machines.

8.2.4.2.2. Configuring a Samba Server in Autonomous Mode

In test phase, it's preferable to restart Samba completely each time you modify `/etc/samba/smb.conf`.

8.2.4.2.2.1. Pre-requisites

Here's the list of elements which must be defined to start:

Workgroup name

In autonomous mode, it's necessary to define the workgroup name to which the server belongs. It's only an arbitrary information but it's mandatory to regroup different machines in the same virtual group.

This information is given by the `workgroup` attribute. This option will also allow you to specify the domain name to control when the server becomes a PDC. For this example, the workgroup's name will be "example".

```
# cat /etc/samba/smb.conf
[global]
...
        workgroup = example
...
```

It's mandatory to specify this attribute.

NetBIOS name of the machine

It's possible to give a NetBIOS name to a machine which will be different than the name affected in the DNS. You can manage this through the `netbios name` attribute:

```
# cat /etc/samba/smb.conf
[global]
...
        netbios name = MES5
...
```

This attribute is optional. If it's omitted, the attribute will take the machine name as defined in the DNS, by default.

Comment about the machine

Similarly, we can associate a comment to it to identify the machine more clearly:

```
# cat /etc/samba/smb.conf
  [global]
  ...
  server string = Samba Server %v
  ...
```

Again, this attribute is optional. If it's omitted, it will stay empty.

8.2.4.2.2.2. Configuring the Service

It's the simplest and fastest way to test the installed packages. No domain notion, only one server with a few users that can connect on a personal share (`[homes]`), and a public share open to all (even anonymous). This simple test allows us to validate that the packages are operational and that the authentication works.

Let's use the following `/etc/samba/smb.conf` file:

```
# cat /etc/samba/smb.conf
[global]
        workgroup = example
        server string = Samba Server %v
        map to guest = Bad User
        log file = /var/log/samba/%m.log
        max log size = 50
        socket options = TCP_NODELAY SO_RCVBUF=8192 SO_SNDBUF=8192

[homes]
        #the [homes] block use corresponds to the default definition of
        # personal user directories
        ; the default path will then be ///home/username//
        # we could add the following attributes to define the shares
        ; comment = Home Directories
        ; browseable = no
        # and users can create/modify/delete files/directories in it
        read only = no

[public]
        # this share defines a public directory
        path = /home/public
        comment = Public Directory
        # we don't have to authentify to access it
        guest ok = yes
        # and users can create/modify/delete files/directories in it
        read only = no
        ; browseable = no
```

As of version 3.4.x, Samba by default stores the account details in the **tdbsam** backend instead of **smbpasswd** backend which was the default for at least Samba 2.0.x versions. In anticipation of the arrival of Samba 3.4.x, it is advisable to use the **tdbsam** backend.

Indeed, the upgrade process does not migrate the accounts. It is therefore advisable to migrate now account details of Samba **smbpasswd** backend to **tdbsam** backend using the command (as root):

```
#pdbedit -i smbpasswd -e tdbsam
```

In the case, inadvisable, in which you want to keep the **smbpasswd** backend, do not forget to add the following option in `smb.conf`:

```
passdb backend = smbpasswd
```

Let's verify that it doesn't have any errors:

```
# testparm
      Load smb config files from /etc/samba/smb.conf
      Processing section "[homes]"
      Processing section "[public]"
      Loaded services file OK.
      WARNING: passdb expand explicit = yes is deprecated
      Server role: ROLE_STANDALONE
      Press enter to see a dump of your service definitions

      [global]
      workgroup = example
      server string = Samba Server %v
      map to guest = Bad User
      log file = /var/log/samba/%m.log
      max log size = 50
      socket options = TCP_NODELAY SO_RCVBUF=8192 SO_SNDBUF=8192

      [homes]
      read only = No

      [public]
      comment = Public Directory
      path = /home/public
      read only = No
      guest ok = Yes
```

No noticeable errors, only a warning concerning the obsolete attribute which is defined by default.

8.2.4.2.2.3. Starting the Service

Now the service can be started:

```
# service smb start
Launching SaMBa service:                                             [  OK  ]
Launching NMB service:                                               [  OK  ]
# ps aux
...
root     27843  0.0  0.7  10724  4028 ?        Ss   12:43   0:00 smbd -D
root     27844  0.0  0.7  10724  4020 ?        S    12:43   0:00 smbd -D
root     27854  0.0  0.4   6568  2068 ?        Ss   12:43   0:00 nmbd -D
...
```

The service is started. Let's see if it answers to requests:

```
# smbclient -L localhost
Password:
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Sharename       Type      Comment
        ---------       ----      -------
        homes           Disk
        public          Disk      Public Directory
        IPC$            IPC       IPC Service (Samba Server 3.2.7)
        ADMIN$          IPC       IPC Service (Samba Server 3.2.7)
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Server               Comment
        ---------            -------

        Workgroup            Master
        ---------            -------
        example
```

The service answers when we ask it to list visible resources. We can play with the browseable attribute in the shares to verify that they appear, or not. For the rest of the tests, the `[homes]` share will have the `browseable = no` attribute.

8.2.4.2.2.4. First Connection in Anonymous Mode

We defined a public share accessible to all, even without authentication. We must now verify that it's really the case:

```
# mkdir -m 777 /home/public
      #  ll /home
total 28
...
drwxrwxrwx  2 root   root   4096 jui 26 13:03 public/
# smbclient //localhost/public
Password:
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]
smb: \> dir
  .                                   D        0  Wed Jul 26 13:03:24 2008
  ..                                  D        0  Wed Jul 26 13:03:24 2008

                48377 blocks of size 16384. 47347 blocks available
smb: \> mkdir test
smb: \> dir
  .                                   D        0  Wed Jul 26 13:03:56 2008
  ..                                  D        0  Wed Jul 26 13:03:24 2008
  test                                D        0  Wed Jul 26 13:03:56 2008

                48377 blocks of size 16384. 47347 blocks available
# ll /home/public/
total 4
drwxr-xr-x  2 nobody nogroup 4096 jui 26 13:03 test/
```

Everything seems OK for the moment.

8.2.4.2.2.5. Creating and Using a User

Let's create a user on the system and on Samba:

```
# useradd -g users -m qatest
# getent passwd qatest
qatest:x:1001:100::/home/qatest:/bin/bash
# ll /home/qatest/ -d
drwxr-xr-x  3 qatest users 4096 jui 26 12:34 /home/qatest//
# passwd qatest
Changing password for user qatest.
      New UNIX password:
BAD PASSWORD: it is based on a dictionary word
Retype new UNIX password:
passwd: all authentication tokens updated successfully.
```

```
# smbpasswd -a qatest
New SMB password:
Retype new SMB password:
# cat /etc/samba/smbpasswd
qatest:1001:8B28C7EF8A97362BAAD3B435B51404EE
        :EB407C0BA4F661A80BCF6B8231A0F6F7:[U          ]:LCT-44C74D6A:
```

Now, let's verify our user:

```
# ssh qatest@localhost
Warning: Permanently added 'localhost' (RSA) to the list of known hosts.
qatest@localhost's password:
[qatest@mes5] $ smbclient -L localhost -U qatest
Password:
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Sharename       Type      Comment
        ---------       ----      -------
        public          Disk      Public Directory
        IPC$            IPC       IPC Service (Samba Server 3.2.7)
        ADMIN$          IPC       IPC Service (Samba Server 3.2.7)
        qatest          Disk      Home directory of qatest
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Server               Comment
        ---------            -------

        Workgroup            Master
        ---------            -------
        example            MES5
```

We see that the `[homes]` share has disappeared from the list (`browseable = no` attribute in place), but that a username share has also appeared (**`qatest`**. In fact, a user has the right to see his own directories if we use the standard definition of `[homes]`.

Let's connect to that share:

```
$ smbclient //localhost/qatest -U qatest
Password:
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]
smb: \> dir
  .                                   D        0  Wed Jul 26 13:14:33 2008
  ..                                  D        0  Wed Jul 26 13:03:24 2008
  tmp                                 D        0  Wed Jul 26 12:34:36 2008
  .screenrc                           H     3793  Wed Jul 26 12:34:37 2008
  .bash_logout                        H       24  Wed Jul 26 12:34:37 2008
  .bash_profile                       H      191  Wed Jul 26 12:34:37 2008
  .bashrc                             H      124  Wed Jul 26 12:34:37 2008
  .bash_completion                    H      145  Wed Jul 26 12:34:37 2008
  .bash_history                       H       33  Wed Jul 26 13:14:33 2008

                48377 blocks of size 16384. 47346 blocks available
smb: \> mkdir test
```

```
smb: \> dir
  .                                   D        0  Wed Jul 26 13:26:04 2008
  ..                                  D        0  Wed Jul 26 13:03:24 2008
  tmp                                 D        0  Wed Jul 26 12:34:36 2008
  .screenrc                           H     3793  Wed Jul 26 12:34:37 2008
  .bash_logout                        H       24  Wed Jul 26 12:34:37 2008
  .bash_profile                       H      191  Wed Jul 26 12:34:37 2008
  .bashrc                             H      124  Wed Jul 26 12:34:37 2008
  .bash_completion                    H      145  Wed Jul 26 12:34:37 2008
  .bash_history                       H       72  Wed Jul 26 13:24:36 2008
  test                                D        0  Wed Jul 26 13:26:04 2008

                48377 blocks of size 16384. 47346 blocks available
$ ll -a
total 40
drwxr-xr-x  4 qatest users 4096 jui 26 13:26 ./
drwxr-xr-x  6 root   root  4096 jui 26 13:03 ../
-rw-r--r--  1 qatest users  145 jui 26 12:34 .bash_completion
-rw-------  1 qatest users   72 jui 26 13:24 .bash_history
-rw-r--r--  1 qatest users   24 jui 26 12:34 .bash_logout
-rw-r--r--  1 qatest users  191 jui 26 12:34 .bash_profile
-rw-r--r--  1 qatest users  124 jui 26 12:34 .bashrc
-rw-r--r--  1 qatest users 3793 jui 26 12:34 .screenrc
drwxr-xr-x  2 qatest users 4096 jui 26 13:26 test/
drwx------  2 qatest users 4096 jui 26 12:34 tmp/
```

Now, let's see the public share:

```
$ smbclient //localhost/public -U qatest
Password:
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]
smb: \> dir
  .                                   D        0  Wed Jul 26 13:03:56 2008
  ..                                  D        0  Wed Jul 26 13:03:24 2008
  test                                D        0  Wed Jul 26 13:03:56 2008

                48377 blocks of size 16384. 47346 blocks available
smb: \> mkdir qatest
smb: \> dir
  .                                   D        0  Wed Jul 26 13:30:35 2008
  ..                                  D        0  Wed Jul 26 13:03:24 2008
  test                                D        0  Wed Jul 26 13:03:56 2008
  qatest                              D        0  Wed Jul 26 13:30:35 2008

                48377 blocks of size 16384. 47346 blocks available
$ ll /home/public/
total 8
drwxr-xr-x  2 qatest users   4096 jui 26 13:30 qatest/
drwxr-xr-x  2 nobody nogroup 4096 jui 26 13:03 test/
```

Our autonomous Samba server is fully operational.

8.2.4.2.3. Samba with LDAP Authentication

The configuration that we've used up until now will be our base for this section. We will only modify the authentication. The latter will now be based on an LDAP directory.

8.2.4.2.3.1. Functional Samba Installation

Simply refer to the preceding section (Section 8.2.4.2.2). We'll keep the same pre-requisites concerning the domain name, the machine's NetBIOS name and the comment associated with it.

```
# vi /etc/samba/smb.conf
[global]
        ...
        workgroup = example
        netbios name = MES5
        server string = Samba Server %v
        ...
```

8.2.4.2.3.2. Domain SID

As long as the service hasn't been started at least once, no SID will have been attributed to the server/domain. It's for that reason that the recuperation of that information is only done after the test in autonomous mode.

The information is kept in `/etc/samba/secrets.tdb`, but it's not directly readable. We'll recuperate it with the `net getlocalsid` command:

```
# net getlocalsid
SID for domain dhcp110 is: S-1-5-21-1518519320-3136826138-1578965553
```

8.2.4.2.3.3. Domain Users and Groups

First of all, let's recall the list of particular SIDs recognized by Windows®. During its installation, Windows® NT4/200x/XP is configured with certain entities (users, group or alias). Each entity has an identified RID. These specific RID must be respected to preserve the integrity of operations. Samba must be fed with those essential domain entities.

If Samba is configured to use tdbsam, the essential entities are automatically created. If LDAP is used, the directory administrator is responsible of their creation: you can use `smbldap-tools` to do so, or the `smbldap-populate` script.

Declared Entity	**RID**	**Type**	**Essential**
Domain Administrator	500	User	No
Domain Guest	501	User	No
Domain KRBTGT	502	User	No
Domain Admins	512	Group	Yes
Domain Users	513	Group	Yes
Domain Guests	514	Group	Yes
Domain Computers	515	Group	No
Domain Controller	516	Group	No
Domain Certificate Admins	517	Group	No
Domain Schema Admins	518	Group	No
Domain Enterprise Admins	519	Group	No
Domain Policy Admins	520	Group	No
Builtin Admins	544	Alias	No
Builtin users	545	Alias	No
Builtin Guests	546	Alias	No
Builtin Power Users	547	Alias	No
Builtin Account Operators	548	Alias	No
Builtin System Operators	549	Alias	No
Builtin Print Operators	550	Alias	No
Builtin Backup Operators	551	Alias	No

Declared Entity	RID	Type	Essential
Builtin Replicator	552	Alias	No
Builtin RAS Servers	553	Alias	No

Table 8-1. Windows Domain Entities

According to those specifications, we can deduce the mandatory entities and optional ones:

- mandatory entries: even though it's not mandatory to have one for the PDC to work correctly, the domain administrator (Domain Administrator/500 - by default, only account to have total control on the system) is very important to manage the domain. A "nobody" user must be created if we want to use the `ldapsam:trusted` instruction (to diminish the dialog between Samba and the POSIX sub-system).

  Concerning groups, the following **must** be created: the Domain Administrator group, the Domain User group, and the Domain Guest group.
- optional entries: according to the table, we should add the corresponding entities, and especially the Domain Computers group in order to associate the machines we'll integrate on that domain to it.

8.2.4.2.3.4. LDAP Directory

This section's goal is not to install a directory. Even if we describe some manipulations, they can in no case be considered as a recipe to put in place a production LDAP directory. On a server used to authentify services such as SSH/PAM, mail, and Apache, a tree structure will already exist.

For this section:

- The basedn considered will be `dc=example,dc=com`
- Users are in the `ou=people,dc=example,dc=com` OU
- Groups are in the `ou=group,dc=example,dc=com` OU
- A `OU` must be created to keep the machine accounts, if we don't want to mix them with the users (`ou=hosts,dc=example,dc=com`).
- We must also have one (or many) account which will have have write-access in those `OU`s, as well as in the basedn to add/modify/delete information relative to Samba, and complementary tools.

You can quickly put in place an LDAP server with a script provided with the distribution. You must install the `openldap-example-dit` package. Then, execute the script to generate the configuration with the data that will be specified. This script creates the tree structure and the necessary accounts for Samba+LDAP, but also necessary to other applications:

```
# /usr/share/openldap/scripts/example-dit-setup.sh
Please enter your DNS domain name [example.com]:
example.com


Administrator account

The administrator account for this directory is
uid=LDAP Admin,ou=System Accounts,dc=example,dc=com

Please choose a password for this account:
New password:
Re-enter new password:


Summary
=======

Domain:        example.com
LDAP suffix:   dc=example,dc=com
Administrator: uid=LDAP Admin,ou=System Accounts,dc=example,dc=com

Confirm? (Y/n)

config file testing succeeded
Stopping ldap service
Finished, starting ldap service
Launching of /usr/bin/db_recover under /var/lib/ldap
removing /var/lib/ldap/alock
Launching of slapd (ldap + ldaps) :                     [  OK  ]

Your previous database directory has been backed up as /var/lib/
ldap.1155740637
All files that were backed up got the suffix "1155740637".

# ldapsearch -x
# extended LDIF
#
# LDAPv3
# base <> with scope subtree
# filter: (objectclass=*)
# requesting: ALL
#

# example.com
dn: dc=example,dc=com
dc: example
objectClass: domain
objectClass: domainRelatedObject
```

```
associatedDomain: example.com

# People, example.com
dn: ou=People,dc=example,dc=com
ou: People
objectClass: organizationalUnit

# Group, example.com
dn: ou=Group,dc=example,dc=com
ou: Group
objectClass: organizationalUnit
description: Container for user accounts

# System Accounts, example.com
dn: ou=System Accounts,dc=example,dc=com
ou: System Accounts
objectClass: organizationalUnit
description: Container for System and Services privileged accounts

# System Goups, example.com
dn: ou=System Groups,dc=example,dc=com
ou: System Groups
objectClass: organizationalUnit
description: Container for System and Services privileged groups

# Hosts, example.com
dn: ou=Hosts,dc=example,dc=com
ou: Hosts
objectClass: organizationalUnit
description: Container for Samba machine accounts
...
# Account Admin, System Accounts, example.com
dn: uid=Account Admin,ou=System Accounts,dc=example,dc=com
uid: Account Admin
objectClass: account
objectClass: simpleSecurityObject
description: Account used to administer all users, groups, machines
        and general accounts

# nssldap, System Accounts, example.com
dn: uid=nssldap,ou=System Accounts,dc=example,dc=com
uid: nssldap
objectClass: account
objectClass: simpleSecurityObject
description: Unprivileged account which can be used by nss_ldap for
when anonymous searches
 are disabled...
```

We find Samba's `OU`s (`People/Group/Hosts`) as well as an OU (`System Accounts`) that contains connection accounts for the different account management tools (account admin, nssldap).

8.2.4.2.3.5. NSS + LDAP

In the part concerning autonomous mode (Section 8.2.4.2.2.5), we saw that we must first create a system user and to make it be recognized by Samba. The same goes for a PDC for which users are registered on an LDAP directory.

To do so, you have to create a Posix user in the LDAP directory and to add the specific information to Samba (using smbldap-tools, for example), and the system must also be able to access this user. We must then configure an NSS (Name Service Switch).

Let's install nssldap.

```
# urpmi nss_ldap
```

For NSS to know that it must consult an LDAP directory to find users and groups, we must modify `/etc/nsswitch.conf`:

```
# /etc/nsswitch.conf
# An example Name Service Switch config file. This file should be
# sorted with the most-used services at the beginning.
#
...
passwd:         files ldap
shadow:         files
group:          files ldap
...
```

> It's important to keep the specified order between `files` and ldap in order for the administration accounts specific to this machine be found without having to interrogate the LDAP directory.

It's not enough to indicate to NSS to interrogate the LDAP directory. We must also show how to do it. To do so, we must inform `/etc/ldap.conf`. In this file, we can specify a user (here, `uid=nssldap,ou=System Accounts,dc=example,dc=com`) which will connect onto the LDAP directory to read information: it's not critical, but it can be useful to track requests or to put in place complex ACL.

```
# vi /etc/ldap.conf
host 127.0.0.1
base dc=example,dc=com
# We uncomment the 2 following lines to track/ACLs
#binddn uid=nssldap,ou=System Accounts,dc=example,dc=com
#bindpw nssldap
nss_base_passwd         ou=People,dc=example,dc=com?one
nss_base_passwd         ou=Hosts,dc=example,dc=com?one
nss_base_shadow         ou=People,dc=example,dc=com?one
nss_base_group          ou=Group,dc=example,dc=com?one
```

In order for Samba to work correctly, it's necessary to allow NSS to read the content of the `OU` hosts: Samba can then identify machines wishing to connect to the domain.

We must now verify that the mechanism works. To do so, you need users and groups in the LDAP directory. If it's not already the case, here's a `ldif` file that you can add to do test:

```
# vi /root/test.ldif
# Test Group
dn: cn=test,ou=Group,dc=example,dc=com
objectClass: top
objectClass: posixGroup
gidNumber: 200
cn: test

# Test User
dn: uid=test,ou=People,dc=example,dc=com
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: posixAccount
objectClass: shadowAccount
cn: test
sn: test
givenName: test
uid: test
uidNumber: 1000
gidNumber: 200
homeDirectory: /home/test
loginShell: /bin/bash
gecos: System User
userPassword: test
# ldapadd -x -D "uid=LDAP Admin,ou=System Accounts,dc=example,dc=com"
        -w <passwd> -f test.ldif
adding new entry "cn=test,ou=Group,dc=example,dc=com"

adding new entry "uid=test,ou=People,dc=example,dc=com"
```

Let's verify that it works adequately with the following commands:

```
# getent passwd test
test:x:1000:200:System User:/home/test:/bin/bash
# getent group test
test:x:200:
# id test
uid=1000(test) gid=200(test) groups=200(test)
```

8.2.4.2.3.6. PAM and LDAP Behavior

It's possible to allow defined users in the LDAP directory to physically connect to the machine via machine via SSH or through a login screen. To do so, we must configure PAM and the associated services through the pam_ldap package.

Let's install pam_ldap.

```
# urpmi pam_ldap
```

We must then modify the /etc/pam.d/system-auth file, since all PAM point to it.

```
# vi /etc/pam.d/system-auth
#%PAM-1.0

auth        required      pam_env.so
auth        sufficient    pam_unix.so likeauth nullok
auth        sufficient    pam_ldap.so use_first_pass
auth        required      pam_deny.so

account     sufficient    pam_unix.so
account     sufficient    pam_ldap.so
account     required      pam_deny.so

password    required      pam_cracklib.so retry=3 minlen=2  dcredit=0
        ucredit=0
password    sufficient    pam_unix.so nullok use_authtok md5 shadow
password    sufficient    pam_ldap.so
password    required      pam_deny.so

session     required      pam_limits.so
session     required      pam_unix.so
```

To verify pam-ldap's behavior, you just need to try to connect with a user you created to make sure that the configuration is OK.

```
# su - test
su: WARNING: can't change directory to /home/test: No file or
        directory of this type
-bash-3.00$ id test
uid=1000(test) gid=200(test) groups=200(test)
```

A warning is displayed because the user's directory (/home/test) doesn't exist.

To do the same test but with SSH, you must modify SSH's configuration and restart it.

```
# vi /etc/ssh/sshd_config
```

```
...
# Set this to 'yes' to enable PAM authentication, account processing,
# and session processing. If this is enabled, PAM authentication will
# be allowed through the ChallengeResponseAuthentication mechanism.
# Depending on your PAM configuration, this may bypass the setting of
# PasswordAuthentication, PermitEmptyPasswords, and
# "PermitRootLogin without-password". If you just want the PAM account
# and session checks to run without PAM authentication, then enable
# this but set ChallengeResponseAuthentication=no
UsePAM yes
...
# service sshd restart
Stopping sshd :                                            [  OK  ]
Launching sshd :                                           [  OK  ]
# ssh test@localhost
Warning: Permanently added 'localhost' (RSA) to the list of known hosts.
Password:
Could not chdir to home directory /home/test: No such file or directory
/usr/X11R6/bin/xauth:  error in locking authority file
       /home/test/.Xauthority
-bash-3.00$ id
uid=1000(test) gid=200(test) groups=200(test)
```

8.2.4.2.3.7. Using smbldap-tools

The `smbldap-tools` package provides a set of tools to manipulate Samba user accounts in the LDAP directory.

```
# ll /usr/sbin/smbldap-*
-rwxr-xr-x  1 root root  5987 mar 21 14:41 /usr/sbin/smbldap-groupadd*
-rwxr-xr-x  1 root root  2473 mar 21 14:41 /usr/sbin/smbldap-groupdel*
-rwxr-xr-x  1 root root  8881 mar 21 14:41 /usr/sbin/smbldap-groupmod*
-rwxr-xr-x  1 root root  2005 mar 21 14:41 /usr/sbin/smbldap-groupshow*
-rwxr-xr-x  1 root root 10294 mar 21 14:41 /usr/sbin/smbldap-passwd*
-rwxr-xr-x  1 root root 14995 mar 21 14:41 /usr/sbin/smbldap-populate*
-rwxr-xr-x  1 root root 20969 mar 21 14:41 /usr/sbin/smbldap-useradd*
-rwxr-xr-x  1 root root  3244 mar 21 14:41 /usr/sbin/smbldap-userdel*
-rwxr-xr-x  1 root root  7633 mar 21 14:41 /usr/sbin/smbldap-userinfo*
-rwxr-xr-x  1 root root 18992 mar 21 14:41 /usr/sbin/smbldap-usermod*
-rwxr-xr-x  1 root root  1958 mar 21 14:41 /usr/sbin/smbldap-usershow*
```

To use these tools, we must configure them so they respect our structure's organization. We must also recuperate the domain's `SID`.

Let's install `smbdap-tools`:

```
# urpmi smbldap-tools
```

To use these tools, we must indicate the `SID`, the domain name, the name of the different OUs of the LDAP directory, the LDAP directory(ies) to use, as well as the POSIX and Samba attributes, by default. We can use the following minimum file (or copy a similar file in the file provided by default):

```
# vi /etc/smbldap-tools/smbldap.conf
      SID="S-1-5-21-2433760973-660784831-1051970529"
      sambaDomain="example

slaveLDAP="127.0.0.1"
slavePort="389"
masterLDAP="127.0.0.1"
masterPort="389"

ldapTLS="0"
verify="require"
cafile="/etc/ssl/cacert.pem"
clientcert=""
clientkey=""

suffix="dc=example,dc=com"
usersdn="ou=People,${suffix}"
computersdn="ou=Hosts,${suffix}"
groupsdn="ou=Group,${suffix}"
idmapdn="ou=Idmap,${suffix}"

sambaUnixIdPooldn="sambaDomainName=${sambaDomain},${suffix}"

scope="sub"

hash_encrypt="SSHA"
crypt_salt_format="%s"

userLoginShell="/bin/false"
userHome="/home/%U"
userHomeDirectoryMode="700"
userGecos="smbldap System User"
defaultUserGid="513"
defaultComputerGid="515"
skeletonDir="/etc/skel"
defaultMaxPasswordAge="45"

userSmbHome="\\MES5\%U"
userProfile="\\MES5\profiles\%U"
userHomeDrive="U:"
userScript=""

mailDomain="example.com"

with_smbpasswd="0"
smbpasswd="/usr/bin/smbpasswd"
with_slappasswd="0"
slappasswd="/usr/sbin/slappasswd"
```

We must also instruct the `/etc/smbldap-tools/smbldap_bind.conf` file: it specifies the accounts which will allow you to connect to the LDAP directory to do different manipulations.

```
# cat /etc/smbldap-tools/smbldap_bind.conf
slaveDN="uid=Account Admin,ou=System Accounts,dc=example,dc=com"
```

```
slavePw="passwd"
masterDN="uid=Account Admin,ou=System Accounts,dc=example,dc=com"
masterPw="passwd"
```

Let's check the behavior of the tools:

```
# smbldap-usershow test
dn: uid=test,ou=People,dc=example,dc=com
objectClass: top,person,organizationalPerson,inetOrgPerson,
        posixAccount,shadowAccount
cn: test
sn: test
givenName: test
uid: test
uidNumber: 1000
gidNumber: 200
homeDirectory: /home/test
loginShell: /bin/bash
gecos: System User
userPassword: test
# smbldap-groupshow test
dn: cn=test,ou=Group,dc=example,dc=com
objectClass: top,posixGroup
gidNumber: 200
cn: test
```

Let's create the domain's users and groups. The smbldap-tools provide a tool (`smbldap-populate`) which creates the LDAP directory's base with the necessary users and groups so it works properly.

Let's add the necessary accounts:

```
# smbldap-populate -a Administrator -k 500 -m 512
Populating LDAP directory for domain example (S-1-5-21-2433760973-
        660784831-1051970529)
(using builtin directory structure)

entry dc=example,dc=com already exist.
entry ou=People,dc=example,dc=com already exist.
entry ou=Group,dc=example,dc=com already exist.
entry ou=Hosts,dc=example,dc=com already exist.
entry ou=Idmap,dc=example,dc=com already exist.
adding new entry: uid=Administrator,ou=People,dc=example,dc=com
adding new entry: uid=nobody,ou=People,dc=example,dc=com
adding new entry: cn=Domain Admins,ou=Group,dc=example,dc=com
adding new entry: cn=Domain Users,ou=Group,dc=example,dc=com
adding new entry: cn=Domain Guests,ou=Group,dc=example,dc=com
adding new entry: cn=Domain Computers,ou=Group,dc=example,dc=com
adding new entry: cn=Administrators,ou=Group,dc=example,dc=com
adding new entry: cn=Account Operators,ou=Group,dc=example,dc=com
adding new entry: cn=Print Operators,ou=Group,dc=example,dc=com
adding new entry: cn=Backup Operators,ou=Group,dc=example,dc=com
adding new entry: cn=Replicators,ou=Group,dc=example,dc=com
adding new entry: sambaDomainName=example,dc=example,dc=com
```

```
Please provide a password for the domain Administrator:
Changing UNIX and samba passwords for Administrator
New password:
Retype new password:
```

8.2.4.2.3.8. Configuring the Samba Server

The [global] section /etc/samba/smb.conf must contain the following information:

```
[global]
...
# we can define many source types
;passdb backend = ldapsam, smbpasswd, guest
;passdb backend = ldapsam:ldap://myfirstldap.edge-it.fr,
        ldapsam:ldap://mysecondldap.edge-it.fr, guest
# here we work with a local LDAP directory
passdb backend = ldapsam:ldap://127.0.0.1

# we connect to the directory with the following LDAP account
ldap admin dn = uid=Account Admin,ou=System Accounts,dc=example.dc=com
# in TLS or SSL
; ldap ssl = start_tls
ldap ssl = off

# the basedn and the OUs where is stored the information
ldap suffix = dc=example.com
ldap machine suffix = ou=hosts
ldap user suffix = ou=people
ldap group suffix = ou=group

# this option allows to not have to configure PAM
# it’s only valid is Samba users don’t have to access directly
#     to an account on the machine
ldapsam:trusted = yes
```

The verification tells us that everything’s OK.

```
# testparm
Load smb config files from /etc/samba/smb.conf
Processing section "[homes]"
Processing section "[public]"
Loaded services file OK.
WARNING: passdb expand explicit = yes is deprecated
Server role: ROLE_STANDALONE
Press enter to see a dump of your service definitions

[global]
        workgroup = EXAMPLE
        server string = Samba Server %v
        map to guest = Bad User
        passdb backend = ldapsam:ldap://127.0.0.1
        log file = /var/log/samba/%m.log
        max log size = 50
        socket options = TCP_NODELAY SO_RCVBUF=8192 SO_SNDBUF=8192
        ldap admin dn = uid=Account Admin,ou=System Accounts,
                        dc=example,dc=com
        ldap group suffix = ou=group
        ldap machine suffix = ou=hosts
        ldap suffix = dc=example,dc=com
        ldap ssl = no
        ldap user suffix = ou=people
        ldapsam:trusted = yes

[homes]
        read only = No
        browseable = No

[public]
        comment = Public Directory
        path = /home/public
        read only = No
        guest ok = Yes
```

Let's test the connection in anonymous mode:

```
# smbclient -L localhost
Password:
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Sharename       Type      Comment
        ---------       ----      -------
        ADMIN$          IPC       IPC Service (Samba Server 3.2.7)
        IPC$            IPC       IPC Service (Samba Server 3.2.7)
        public          Disk      Public Directory
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Server               Comment
        ---------            -------

        Workgroup            Master
        ---------            -------
        example             MES5
```

We must then pass the password of the directory's connection account to Samba:

```
# smbpasswd -W
Setting stored password for "uid=Samba Admin,ou=System Accounts,
        dc=example,dc=com" in secrets.tdb
New SMB password:
Retype new SMB password:
```

We can already verify that the link to the LDAP directory is in place (by typing a bad password, and then the good one).

```
# smbclient -L localhost -U administrator
Password:
session setup failed: NT_STATUS_LOGON_FAILURE
# smbclient -L localhost -U administrator
Password:
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Sharename       Type      Comment
        ---------       ----      -------
        public          Disk      Public Directory
        IPC$            IPC       IPC Service (Samba Server 3.2.7)
        ADMIN$          IPC       IPC Service (Samba Server 3.2.7)
Domain=[MES5] OS=[Unix] Server=[Samba 3.2.7]

        Server               Comment
        ---------            -------

        Workgroup            Master
        ---------            -------
        example
```

Now, let's define our server as a PDC because up until now, we were using it in autonomous mode:

```
# vi /etc/samba/smb.conf
[global]
...
        # PDC
        security = user
        # OS level > 32 to be elected
        os level = 128
        # allows election launch
        # definition of a PDC
        local master = yes
        # domain master browser
        domain master = yes
        # for ce the electios to become PDC
        preferred master = yes
        # server authenticates
        domain logons = yes
...
#  testparm
Load smb config files from /etc/samba/smb.conf
Processing section "[homes]"
Processing section "[public]"
Loaded services file OK.
WARNING: passdb expand explicit = yes is deprecated
Server role: ROLE_DOMAIN_PDC
Press enter to see a dump of your service definitions

[global]
        workgroup = example
        server string = Samba Server %v
        map to guest = Bad User
        passdb backend = ldapsam:ldap://127.0.0.1
        log file = /var/log/samba/%m.log
        max log size = 50
        socket options = TCP_NODELAY SO_RCVBUF=8192 SO_SNDBUF=8192
        domain logons = Yes
        os level = 128
        preferred master = Yes
        domain master = Yes
        ldap admin dn = uid=Samba Admin,ou=System Accounts,dc=edge-it,
                        dc=subnet
        ldap group suffix = ou=group
        ldap machine suffix = ou=hosts
        ldap suffix = dc=edge-it,dc=subnet
        ldap ssl = no
        ldap user suffix = ou=people
        ldapsam:trusted = yes

[homes]
        read only = No
        browseable = No

[public]
        comment = Public Directory
        path = /home/public
        read only = No
        guest ok = Yes
```

8.2.4.2.3.9. Managing Samba File Shares

Until now, we have looked at the general configuration of a Samba server and the set up of its role on the network. Let's not review the types of resource sharing that you can set up.

A share's syntax is relatively simple and must be written after the server's general parameters:

```
[<share_name>]
 comment = "<your comment>"
 path = <path_resource>
 browseable = <yes|no>
 writable = <yes|no>
```

Just as for the general section, a share is introduced by its name being between square brackets. Here are the most common parameters:

- `comment`: enter a significant comment. It appears during the resource path and enables you to identify its content quickly;
- `path`: specifies the local path to the shared resource;
- `browseable`: determines if the resource must appear during the resource path;

- `writable`: indicates the write or read rights on the shared resource.

`logon`-type share

This specific share concerns netlogon scripts or connection scripts of domain users. It allows us to share netlogon scripts generated on the fly through the `root preexec` instruction.

```
[netlogon] comment = Network Logon Service
  path = /data/samba/netlogon guest ok = yes writable = no
  write list = @administrateurs browseable = no root preexec =
  /data/samba/netlogon/logon_script '%m' '%U' '%a' '%g'
  '%L'
```

`profile`-type share

In the situation where you use mobile profiles, these will be part of a specific Samba share which uses the keyword `profiles`.

```
[profiles] path = /data/samba/profiles
       browseable = no guest ok = yes writable = yes
```

`homes`-type share

You can make your user's home directories accessible (*homes*). For this, we use the reserved `homes` keyword. `%u` is a predefined variable containing the user's username.

```
[homes] comment = Home
  Directories browseable = no writable = yes path =
  /data/samba/prives/%u
```

`group`-type share

This share type enables you to define accessible resources for one or many given groups. It avoids having to define a share per group. The principle is simple: it lies on the assignation of rights in the data tree structure. The `hide unreadable` parameter allows to us hide to users directories for which they have no right. For example:

```
# ll /data/samba/groups/ total
       24 drwxrws--- 111 root commercial 4096 avr 28 17:45 clients/
       drwxrws--- 23 root utilisateurs 4096 sep 9 2008 commercial/
       drwxrws--- 6 root support 104 jui 19 14:24 support/
```

The share can then be written this way::

```
[groups] comment =
  Storage Groups path = /data/samba/groups writable = yes
```

```
      browseable = yes hide unreadable = yes
```

8.2.4.2.3.10. Managing Samba Print Shares

We saw file sharing, let's now address printer sharing. Samba allows you to share access and drivers, so you don't have to install those drivers on each client machine.

The `[printers]` specific share allows us to share all available printers through CUPS. The `[print$]` specific printing permits to share drivers.

```
# sharing declared printers
 [printers]
 comment = All Printers
 path = /var/spool/samba
 browseable = no
 guest ok = yes
 writable = no
 printable = yes
 create mode = 0700

 # share printing driver distribution
 [print$]
 path = /var/lib/samba/printers
 browseable = yes
 read only = yes
 write list = @administrateurs
      guest ok = yes
```

As a prerequisite you must have a functional CUPS server. It must also be declared in the [global] section:

```
[global]
 ...
 # we use printers defined in CUPS
 # we load the list
 printcap name = cups
 load printers = yes

 # adding authorized printer for the admin group
 printer admin = @administrateurs
```

If the CUPS server is not on the same machine, you can add the `cups server` instruction followed by the server's address.

Once your shares are defined as well as the CUPS server to use, you need to make Samba take into account the defined printers Samba. Use the `cupsaddsmb` command:

```
cupsaddsmb -a -u
     <administrator_name>
```

In the case where you want to add printers to your CUPS server, it will be necessary to relaunch the Samba server before executing the `cupsaddsmb` command.

8.2.4.2.4. Troubleshooting

Solving a Samba problem is not always an easy task. Here's a checklist to do when crashes or malfunctions occur:

- Check the available disk space on the filesystems hosting the shares.
- Check the available disk space for log files (watch out, if the level is high, also check the space left in `/tmp`). A high level can also affect the server's performance considerably.
- Check the eventual inactivity of the Samba process to better set the timeout variable.
- Check the memory space used by the Samba process.

Samba's official documentation offers a complete guide allowing you to identify exhaustively problem sources and tools to set up to diagnose those problems: see the *Troubleshooting* section (`http://us4.samba.org/samba/docs/man/Samba-HOWTO-Collection/troubleshooting.html`).

8.2.4.2.5. Extended ACLs

Microsoft Windows® uses extended ACLs (Access Control Lists) and attributes in regards with "standard triplé" accessible on a GNU/Linux system (`Read/Write/Execute` for a user, his group, and everyone). It's possible to simulate part of those information on GNU/Linux by using a filesystem which supports such a format, and by adding the packages which facilitate ACL management.

You must then install the following packages:

- acl: this package contains the `getfacl` and `setfacl` commands that allow you to see, add, modify and suppress extended ACLs on a file or directory.
- attr: this package contains the `getfattr` and `setfattr` commands which enable you to see, add, modify and delete extended attributes on a file or directory.

Here are the filesystems which support those ACLs:

- XFS supports those two extended information sets in native mode, like reiserFS and jFS.
- ext2/3 supports these information once the filesystem is patched. It's also necessary to activate the `acl` attribute during the filesystem's mount.

Please read `POSIX Access Control Lists on Linux` for more information.

Make sure you use XFS-512 instead if XFS-256 (the default on many GNU/Linux systems), because this way, extended information are stored in the concerned file's or directory's inode, instead of in the filesystem's metadata. Accessing the information is faster that way.

8.3. Managing Mail Services

In this chapter we discuss a network configuration containing a local SMTP server and POP/IMAP server authentified on a LDAP directory. The protection against viruses and SPAM is assured by `amavisd-new` and the necessary components.and the needed modules. We assume that the LDAP is operational.

8.3.1. POP/IMAP Cyrus-IMAPD Server

Today, Cyrus-IMAP works by default with a SASL (Simple Authentication and Security Layer)authentication security layer. By default, the SASL-secured authentication uses an authentication mode based on PAM.

You will find more information on the project's official website: http://cyrusimap.web.cmu (`http://cyrusimap.web.cmu.edu/`)

8.3.1.1. Installation and Tree Structure of Cyrus-IMAPD

Three packages are necessary:

- `cyrus-sasl`: provides the SASL client and server;
- `cyrus-imapd`: provides the POP/IMAP Cyrus-IMAP server;

- `cyrus-imapd-utils`: provides Cyrus-IMAP utilities, and notably cyradm, which is the server's administration interface.

> Versions don't evolve a lot. However, be rigorous concerning eventual security updates, particularly concerning SASL.

Let's install the required packages:

```
# urpmi cyrus-sasl cyrus-imapd cyrus-imapd-utils
```

Cyrus-IMAP's tree structure is quite simple:

- `/etc/imapd.conf`: configuration file used to access the IMAP server's resources;
- `/etc/cyrus.conf`: Cyrus configuration file;
- `/var/spool/imap/`: directory where mailboxes are stored;
- `/var/log/mail`: log file directory.

8.3.1.2. Configuring Cyrus-IMAP

First, let's check the list of authentication modes available with the SASL version we installed:

```
# saslauthd -v
saslauthd 2.1.22
authentication mechanisms: getpwent kerberos5 pam rimap shadow ldap
```

ldap appears in this list. The mechanism used for authentication is specified in the `/etc/sysconfig/saslauthd` file:

```
# cat /etc/sysconfig/saslauthd
# $Id: CS-service-messaging.xml,v 1.9 2008-09-12 14:19:36 ennael Exp $
# Authentications mechanism (for list see saslauthd -v)
SASL_AUTHMECH=pam
...
```

By default, PAM is the activated authentication mechanism. Let's replace it by LDAP:

```
# cat /etc/sysconfig/saslauthd
# $Id: CS-service-messaging.xml,v 1.9 2008-09-12 14:19:36 ennael Exp $
# Authentications mechanism (for list see saslauthd -v)
SASL_AUTHMECH=ldap
...
```

Now we must define the necessary elements to identify the LDAP directory's contact mode:

```
# cat /etc/saslauthd.conf
ldap_servers: ldap://<ip_ldap_server>
ldap_version: 3
ldap_auth_method: bind
ldap_bind_dn: cn=Manager,dc=example,dc=com
ldap_bind_pw: <password>
ldap_search_base: ou=Users,dc=example,dc=com
ldap_scope: one
ldap_filter: uid=%u
ldap_verbose: on
```

Finally, verify that the service is active on boot-up. If not, it must be configured differently, by activating saslauthd for levels 3 and 5:

```
# chkconfig --level 35 saslauthd on
# chkconfig --list saslauthd
saslauthd  0:Arrêt  1:Arrêt  2:Marche  3:Marche  4:Marche  5:Marche  6:Arrêt
```

Let's configure Cyrus-IMAP. You only need to specify the server's administrator(s) through the admins parameter.

```
# cat /etc/imapd.conf
configdirectory: /var/lib/imap
partition-default: /var/spool/imap
admins: imapadmin cyrus
allowanonymouslogin: no
sieveusehomedir: no
sievedir: /var/lib/imap/sieve
sendmail: /usr/sbin/sendmail
hashimapspool: true
sasl_pwcheck_method: saslauthd
sasl_mech_list: PLAIN
tls_cert_file: /etc/ssl/cyrus-imapd/cyrus-imapd.pem
tls_key_file: /etc/ssl/cyrus-imapd/cyrus-imapd.pem
```

Restart the server:

```
# service cyrus-imapd restart
```

8.3.1.3. Testing the Server’s Behavior

First, we must test saslauthd and the LDAP directory contact with the `testsaslauthd`:

Let’s make sure that Cyrus-IMAP for POP3 (port 110) and IMAP (port 143) works correctly:

```
# telnet 192.168.1.2 143
Trying 192.168.1.2...
Connected to imap.example.com (192.168.1.2).
Escape character is ’^]’.
* OK tellure Cyrus IMAP4 v2.3.12-0.p2.4mdv2009.0 server ready
# telnet 192.168.1.2 110
Trying 192.168.1.2...
Connected to imap.example.com (192.168.1.2).
Escape character is ’^]’.
+OK tellure Cyrus POP3 v2.3.12-0.p2.4mdv2009.0 server ready
<2428069223.1136796638@imap>
```

8.3.1.4. Managing Cyrus-IMAPD

You can manage Cyrus through the cyradm utility. First, connect to the server as a service administrator.

```
# cyradm --user imapadmin localhost
IMAP Password:
tellure> help
authenticate, login, auth         authenticate to server
chdir, cd                         change current directory
createmailbox, create, cm         create mailbox
deleteaclmailbox, deleteacl, dam  remove ACLs from mailbox
deletemailbox, delete, dm         delete mailbox
disconnect, disc                  disconnect from current server
exit, quit                        exit cyradm
help, ?                           show commands
info                              display mailbox/server metadata
listacl, lam, listaclmailbox      list ACLs on mailbox
listmailbox, lm                   list mailboxes
listquota, lq                     list quotas on specified root
listquotaroot, lqr, lqm           show quota roots and quotas for mailbox
mboxcfg, mboxconfig               configure mailbox
reconstruct                       reconstruct mailbox (if supported)
renamemailbox, rename, renm       rename (and optionally relocate) mailbox
server, servername, connect       show current server or connect to server
setaclmailbox, sam, setacl        set ACLs on mailbox
setinfo                           set server metadata
setquota, sq                      set quota on mailbox or resource
version, ver                      display version info of current server
```

Here are the principal user management tasks:

- Adding a user and creating his mailbox: `cm`;

  ```
  localhost> cm user.toto
  localhost> cm user.toto.sent
  localhost> cm user.toto.trash
  ```

- Deleting a mailbox: `dm`;

  ```
  localhost> dm user.toto
  ```

- list the mailboxes created: `lm`

  ```
  localhost> lm
        user.loic (\HasNoChildren)
        user.anne (\HasNoChildren)
        user.benjamin (\HasNoChildren)
  ```

- Set quotas: `sq`

  ```
  localhost> sq 524288000 user.anne
        user.loic (\HasNoChildren)
        user.anne (\HasNoChildren)
        user.benjamin (\HasNoChildren)
  ```

- List quotas: `lq`;

  ```
  localhost> lq user.anne
   STORAGE 0/524288000 (0%)
  ```

- Set ACLs: `setacl`. We use the ACLs listed below:

ACL	**Content**
l	see the list of mailboxes without their content
r	read the content of mailboxes
s	keep the “seen” and “recent” states during IMAP sessions
w	write (modification of the “recent”, “answered” and “draft” messages’ indicators)
i	inserting a message in a mailbox (moving or copying)

ACL	Content
c	create sub-mailboxes in the main box (creating main mailboxes isn't authorized for non-admin users)
d	destroy a message and/or the box itself
a	manage the mailbox (modify ACLs)
none	no rights
read	((=lrs) read the content of the mailbox
append	(=lrsip) read the content of the mailbox and add a message in queue
write	(=lrswipcd) read the content, post in it, add a message in queue, destroy a message or the box itself. In short, all rights except the one to modify ACLs
all	(=lrswipcda) all rights, usually given to the mailboxes' respective owners

Table 8-2. Managing ACLs in Cyrus-IMAP

To set ACLs, use the following command:

```
localhost> setacl user.mailgroup anne read
```

- list ACL: `lam`

```
localhost> lam user.mailgroup
      anne lrswipcda
      loïc lrswipcda
      benjamin lrswipcda
```

> To generate mailboxes on the fly, create a file containing all the mailbox creation commands. Then, send the standard output of the file's display to the `cyradm`.

Example: let's create boxes for users anne, loic and benjamin:

```
# cat liste_boites
```

```
    cm user.anne
    cm user.loic
    cm user.benjamin
# cat liste_boites |cyradm -user admin localhost
```

8.3.1.5. Cyrus Toolbox

Cyrus-IMAP is provided with a number of tools allowing to us test and verify it's good functioning especially on an authentication level:

```
# imtest -a anne localhost
S: * OK tellure Cyrus IMAP4 v2.3.12-0.p2.4mdv2009.0 server ready
C: C01 CAPABILITY
S: * CAPABILITY IMAP4 IMAP4rev1 ACL QUOTA LITERAL+ MAILBOX-REFERRALS
NAMESPACE UIDPLUS ID NO_ATOMIC_RENAME UNSELECT CHILDREN MULTIAPPEND
BINARY SORT THREAD=ORDEREDSUBJECT THREAD=REFERENCES ANNOTATEMORE IDLE
STARTTLS LISTEXT LIST-SUBSCRIBED X-NETSCAPE
S: C01 OK Completed
Please enter your password:
C: L01 LOGIN anne {8}
S: + go ahead
C: <omitted>
S: L01 OK User logged in
Authenticated.
Security strength factor: 0
```

8.3.1.6. Using POP3S and IMAPS

By default, `cyrus-imapd` is installed and includes the activation of `pop3s` and `imaps`. It is strongly recommended you use the secured `IMAP` protocol through `SSL` on port 443. `IMAPS` is a good option to set up a secured mail solution.

Sous Mandriva Enterprise Server 5, lors de l'installation du démon `cyrus-imap`, IMAPS is activated by default. For that to occur, an `SSL` certificate was automatically generated. If you want to change the `SSL`, you can generate a new one with the following command:

```
# openssl req -new -x509 -nodes -out /etc/ssl/cyrus-imapd/cyrus-imapd.pem
-keyout cyrus-imapd.pem -days 365
```

8.3.2. The Postfix SMTP Server

In this section we only address one particular configuration example of Postfix.

8.3.2.1. Basic SMTP Server Concepts

A SMTP (Simple Mail Transfer Protocol) server can be compared to a post office. The post office receives the letter for the zone in which it's located and sorts it. If a letter is destined to someone living in a zone serviced by the post office, it will deliver it in the person's mailbox. In the opposite case, the letter is sent to the post office which services the recipient's zone.

The standard Postfix server's operations are similar. It receives messages from the local network and of other mail servers which identified it as the mail manager for a given domain. The server reads the recipient's address and:

- If the domain name corresponds to the domain locally managed, the message is deposited in the corresponding mailbox.
- In the opposite case, the server looks for the server which manages the concerned zone and sends it the mail.

Postfix is the successor to Sendmail. It's newer and its architecture lies on the notion of modularity.

Postfix's main Web references:

- The official Postfix site; (`http://postfix.org`)
- The official Postfix documentation; (`http://www.postfix.org/documentation.html`)
- Other related documentation. (`http://www.postfix.org/docs.html`)

8.3.2.2. Installation and Tree Structure of Postfix

Installing Postfix is easy: installing the Postfix package will do it. However, Mandriva Enterprise Server 5 provides a number of additional packages:

- `postfix-pcre`: supports PCRE (*Perl Compatible Regular Expression*) for configuration;
- `postfix-ldap` : supports LDAP maps in Postfix to manage authentication on an LDAP directory;

Postfix's tree structure reflects the modularity of its design:

- `/etc/postfix`: directory containing the server's configuration files;
- `/var/log/mail`: directory containing the server's log files, separated into three files (`info`, `warnings`, `errors`) according to the information's importance;
- `/var/spool/postfix`: directory containing every spool directory relative to the server's functioning as described earlier;
- `/etc/sysconfig/postfix`: group of options used to start the server's daemons.

8.3.2.3. Configuring the Postfix Server

The main configuration file is `/etc/postfix/main.cf`. We'll use this file to present the base parameters which assure the server's good functioning in the case described at the beginning of this chapter.

```
# cat /etc/postfix/main.cf
# paramétrage système du serveur
queue_directory = /var/spool/postfix
command_directory = /usr/sbin
daemon_directory = /usr/lib/postfix
mail_owner = postfix

# nom de machine
myhostname = host.domain.com
# nom de domaine de la valeur myhostname
mydomain = domain.com
# domaine apparaissant dans le courrier envoyé de cette machine
myorigin = $mydomain

# interfaces sur lesquelles le service va pouvoir écouter (par défaut
# toutes)
inet_interfaces = all
# domaines pour lequels la machine livrera le courrier localement au
# lieu de les transmettre à une autre machine
mydestination = $myhostname, localhost.$mydomain,
 /etc/postfix/destinations
# map contenant les adresses et/ou utilisateurs locaux
local_recipient_maps = $alias_maps

# code spécifiant une réponse du serveur SMTP lorsque le domaine d'un
# destinataire correspond à $mysdestination ou lorsque l'adresse de
# destination ou l'adresse locale n'existe pas. Par défaut, le code
# est utilisé est 450, code qui propose de renouveler l'envoi (550
# pour ne pas le renouveler)
unknown_local_recipient_reject_code = 450
# réseaux autorisés à utiliser le serveur SMTP
mynetworks = 172.16.51.0/24, 127.0.0.0/8

# spécifie les bases qui seront utilisées par la commande newaliases
# pour générer la table des alias
```

```
alias_database = hash:/etc/postfix/aliases

mail_spool_directory = /var/spool/mail

# spécifie le mode de transport des mails dans le fichier master.cf
# à utiliser après avoir traité les fichiers aliases et .forward
mailbox_transport = cyrus
# bannière affichée lors de l'accès au serveur SMTP
smtpd_banner = $myhostname ESMTP $mail_name ($mail_version) (Mandrake
 Linux)

debug_peer_level = 2
debugger_command =
PATH=/bin:/usr/bin:/usr/local/bin:/usr/X11R6/bin
xxgdb $daemon_directory/$process_name $process_id & sleep 5
# délai en nombre d'heures au bout duquel un avertissement est envoyé
# quand un courrier n'a pas pu être livré
delay_warning_time = 4

# autres paramétrages système du serveur
sendmail_path = /usr/sbin/sendmail.postfix
newaliases_path = /usr/bin/newaliases.postfix
mailq_path = /usr/bin/mailq.postfix
setgid_group = postdrop
manpage_directory = /usr/share/man
sample_directory = /usr/share/doc/postfix-2.0.6/samples
readme_directory = /usr/share/doc/postfix-2.0.6/README_FILES
```

The other base file is `/etc/postfix/master.cf`. It defines the functioning of each server working inside the server:

- `service type`: behaviour type (TCP/IP socket or UNIX);
- `private`: restricted access to the Postfix service;
- `unprivileged`: service executed with or without root privileges;
- `chroot`: service which runs in chroot mode;
- `wakeup time` time after which the process is automatically reactivated;
- `maxproc`maximum number or processes simultaneously executed;
- `command`: executed command.

8.3.2.4. Postfix Toolbox

Postfix offers a certain number of useful tools for day-to-day server administration:

Verify the configuration

List all the parameters used in the `main.cf` file.

```
# postconf
```

Lists only the personalized parameters.

```
# postconf -n
```

Validates Postfix' configuration (`main.cf` file).

```
# postfix check
postfix: fatal: bad string length 0 < 1: manpage_directory =
```

Configure the Postfix daemon

Start | stop | restart | reload the Postfix service's configuration:

```
# service postfix start | stop | restart | reload
```

Force the delivery of queued messages.

```
# service postfix flush
```

Check the Postfix service's state.

```
# service postfix status
       master (pid 6417) est en cours d'exécution...
```

```
# ps -ef | grep postfix
root      6417     1  0 10:18 ? 00:00:00 /usr/lib/postfix/master
postfix   6422  6417  0 10:18 ? 00:00:00 pickup -l -t fifo -u -c -o
 content_filter  -o receive_override_options
postfix   6423  6417  0 10:18 ? 00:00:00 qmgr -l -t fifo -u -c
```

That command allows you to check precisely that the different daemons which compose Postfix are working. You should at least see: `master`, `qmgr`, and `pickup`.

Postfix maps management

Recreate the `aliases.db` map:

```
# newaliases
```

Recreate a specific mapping:

```
# postmap <map>
```

Read a mapping's content:

```
# postmap -q <map>
```

File management

Display every queued message

```
# postqueue -p
-Queue ID- --Size-- ----Arrival Time---- -Sender/Recipient-------
9D3F67D5F*      338 Fri Jan  6 19:31:43  plop@plop.com
                                         anne@tellure.example.subnet

-- 0 Kbytes in 1 Request.
```

or

```
# mailq
-Queue ID- --Size-- ----Arrival Time---- -Sender/Recipient-------
9D3F67D5F*      338 Fri Jan  6 19:31:43  plop@plop.com
                                         anne@tellure.example.subnet

-- 0 Kbytes in 1 Request.
```

Delete queued mail:

```
# postsuper -d <queue_ID>
#  postsuper -d  9D3F67D5F
postsuper: 9D3F67D5F: removed
postsuper: Deleted: 1 message
```

The `postsuper -d ALL` command allows you to suppress every queued mail.

8.3.2.5. Securing a Mail Server

Here are a few tips enabling you to secure at a minimum your Postfix server.

8.3.2.5.1. Postfix in chroot Mode

Securing Postfix can be done by executing the server's daemon's in a "cage". In Postfix's case, this means that the processes have the weakest possible privileges and have access to a limited tree structure, that is `/var/spool/postfix`.

Mandriva Enterprise Server 5's Postfix package is provided with a script which allows you to easily execute Postfix in chroot mode (and to revert back to normal state): `postfix-chroot.sh`. It executes the following:

- Creates the chroot tree structure in `/var/spool/postfix` by default;
- Modifies `/etc/postfix/master.cf` to specify the execution of the daemons in chroot mode.
- Reloading the service.

```
# /usr/sbin/postfix-chroot.sh enable
 setting up chroot at: /var/spool/postfix
copy system files into chroot
```

```
  /etc/localtime -> /var/spool/postfix/etc/localtime
  /etc/host.conf -> /var/spool/postfix/etc/host.conf
  /etc/resolv.conf -> /var/spool/postfix/etc/resolv.conf
  /etc/nsswitch.conf -> /var/spool/postfix/etc/nsswitch.conf
  /etc/hosts -> /var/spool/postfix/etc/hosts
  /etc/services -> /var/spool/postfix/etc/services
copy additional files into chroot
copy nss libraries into chroot
  /lib64/libnss_dns.so.2 -> /var/spool/postfix/lib64/libnss_dns.so.2
  /lib64/libnss_dns-2.8.so -> /var/spool/postfix/lib64/libnss_dns-2.8.so
  /lib64/libnss_nis.so.2 -> /var/spool/postfix/lib64/libnss_nis.so.2
  /lib64/libnss_nis-2.8.so -> /var/spool/postfix/lib64/libnss_nis-2.8.so
  /lib64/libnss_winbind.so.2 -> /var/spool/postfix/lib64/libnss_winbind.so.2
  /lib64/libnss_winbind.so -> /var/spool/postfix/lib64/libnss_winbind.so
  /etc/ldap.conf -> /var/spool/postfix/etc/ldap.conf
  /lib64/libnss_ldap.so.2 -> /var/spool/postfix/lib64/libnss_ldap.so.2
  /lib64/libnss_ldap-2.8.so -> /var/spool/postfix/lib64/libnss_ldap-2.8.so
  /lib64/libnss_compat.so.2 -> /var/spool/postfix/lib64/libnss_compat.so.2
  /lib64/libnss_compat-2.8.so -> /var/spool/postfix/lib64/libnss_compat-2.8.so
  /lib64/libnss_files.so.2 -> /var/spool/postfix/lib64/libnss_files.so.2
  /lib64/libnss_files-2.8.so -> /var/spool/postfix/lib64/libnss_files-2.8.so
  /lib64/libnss_files.so.2 -> /var/spool/postfix/lib64/libnss_files.so.2
  /lib64/libnss_files-2.8.so -> /var/spool/postfix/lib64/libnss_files-2.8.so
Reloading Postfix Service:                   [  OK  ]
```

After this it's important to update the chroot if modifications are done and which could affect the mail environment. To check the presence of such modifications, type the following command:

```
# postfix-chroot.sh check
files /var/spool/postfix/etc/hosts and /etc/hosts differ
Reloading the Postfix service::             [  OK  ]
```

To update:

```
# postfix-chroot.sh check_update
Reloading the Postfix service:             [  OK  ]
```

8.3.2.5.2. Securing Postfix' Configuration

First of all, let's put a few options in `/etc/postfix/main.cf`.

```
smtpd_helo_required = yes
     disable_vrfy_command = yes

smtpd_recipient_restrictions =
    reject_invalid_hostname,
    reject_non_fqdn_hostname,
    reject_non_fqdn_sender,
    reject_non_fqdn_recipient,
    reject_unknown_sender_domain,
    reject_unknown_recipient_domain,
    permit_mynetworks,
    reject_unauth_destination,
    check_recipient_access
    check_client_access dbm:/etc/postfix/client_checks,
    reject_rbl_client sbl.spamhaus.org,
    reject_rbl_client cbl.abuseat.org,
    reject_rbl_client dnsbl.sorbs.net,
    permit

smtpd_data_restrictions =
    reject_unauth_pipelining,
    # http://www.postfix.org/postconf.5.html#reject_unauth_pipelining
    permit
```

8.3.2.5.3. Filtering at the HELO Step

An important step in the dialog with an SMTP server is the `HELO` command. A certain number of checks at this level of the dialog allows you to do a first filtering which is non immaterial. The configuration is done in the `/etc/postfix/main.cf` file:

```
smtpd_recipient_restrictions = check_helo_access
dbm:/etc/postfix/helo_checks
```

This instruction enables us to specify a new file in which the filters to apply are defined. That file is called `/etc/postfix/helo_checks`:

```
# cat /etc/postfix/helo_checks
# We can block the machines that present themselves as being part of
# the domain while they're really outside of it
example.tld REJECT You are not in example.tld

# Ditto for IP addresses (the mail server's IP)
192.168.1.2 REJECT You are not 192.168.1.2

# Ditto for localhost
localhost REJECT You are not me
```

It's useful to know the dialog's steps between an SMTP client and server. This allows you to detect eventual problem levels:

```
$ telnet localhost 25
Trying 127.0.0.1...
Connected to localhost (127.0.0.1).
Escape character is '^]'.
220 dhcp140.example.com ESMTP Postfix (Mandriva MES5)
HELO mandriva.com
250 mes5.example.com
Mail from test@mandriva.com
```

```
250 Ok
RCPT To: a@example.com
250 Ok
DATA
354 End data with <CR><LF>.<CR><LF>
test
.
250 Ok: queued as B4DAE434B
```

The operation occurs in 4 steps:

1. `HELO`: presentation of the host that contacts, the SMTP server answers;
2. `Mail From`: sender's address, the SMTP server validates the address;
3. `RCPT To:`: recipient's address, the SMTP server validates the address;
4. `DATA`data sending, the message ends with a **.** (period).

8.3.2.5.4. Filtering the Sender

In the same manner, let's use a instruction taken from `/etc/postfix/main.cf`

```
smtpd_recipient_restrictions = check_sender_access
dbm:/etc/postfix/sender_checks
```

In /etc/postfix/sender_checks :

```
# cat /etc/postfix/sender_checks
# We ban a domain
# Code SMTP 554 corresponds to: "Transaction failed"
domaine.org 554 Spam not tolerated here

# A domain is put on the blacklist (RBL) but we still want to
# receive mails from that domain: domain.com OK

# For a domain on the blacklist, we want to receive certain addresses
someuser@example3.tld OK
example3.tld REJECT
```

8.3.2.5.5. Filtrer les destinataires

Let's say we want to filter recipients, for example, from an old mailbox that still gets spammed.

In the `/etc/postfix/main.cf` file:

```
smtpd_client_restrictions =
        check_recipient_access regexp:/etc/postfix/rcpt_restrictions
```

In /etc/postfix/rcpt_restrictions

```
/sales@domaine\.info/ REJECT
/bob@domaine\.info/ REJECT
```

If you want to filter on the format of the sender's addresses, in main.cf:

```
smtpd_recipient_restrictions = check_client_access
  pcre:/etc/postfix/client_checks.pcre
```

In the `/etc/postfix/client_checks.pcre` file:

```
/^\@/ 550 Invalid address format.
/[!%\@].*\@/ 550 This server disallows weird address syntax.
```

8.3.2.6. Advanced Postfix Usage

8.3.2.6.1. LDAP Support in Postfix

Postfix offers the possibility to use an LDAP directory to verify the recipient and to deliver his messages. Beforehand, you must install the `postfix-ldap` package.

The technique consists in declaring the LDAP maps, specifying the LDAP server, how to interrogate the server and the necessary information to recuperate.

```
# cat /etc/postfix/main.cf
...
# alias list used for local mail
alias_maps = ldap:ldapuser, ldap:ldapgroup

# maps use for LDAP authentication LDAP
virtual_alias_maps = ldap:ldapuser, ldap:ldapgroup

# definition of the necessary information to recuperate a user's mail address
ldapuser_server_host = 192.168.1.1
ldapuser_server_port = 389
ldapuser_bind = yes
ldapuser_bind_dn = cn=Manager,dc=example,dc=com
ldapuser_bind_pw = secret
```

```
ldapuser_search_base = ou=Personnes,dc=example ,dc=com
ldapuser_timeout = 60
ldapuser_query_filter = (&(objectclass=qmailuser)(mailLocalAddress=%s))
ldapuser_result_attribute = mail
ldapuser_lookup_timeout = 60

# definition of the necessary information to recuperate a group's mail
# address
ldapgroup_server_host = 192.168.1.1
ldapgroup_server_port = 389
ldapgroup_bind = yes
ldapgroup_bind_dn = cn=Manager,dc=example ,dc=com
ldapgroup_bind_pw = secret
ldapgroup_search_base = ou=Groupes,dc=example ,dc=com
ldapgroup_timeout = 60
ldapgroup_query_filter = (&(objectclass=mailalias)(mailAlias=%s))
ldapgroup_result_attribute = rfc822MailMember
ldapgroup_lookup_timeout = 60

# specifies the bases to use by the newalises command to generate the
# alias table
alias_database = hash:/etc/postfix/aliases, ldap:ldapuser,
 ldap:ldapgroup
```

Restart the Postfix server to take the modification into account.

8.3.2.6.2. SMTP Authentication on Postfix

By default the configured Postfix server will not accept messages coming from its network. To authorize another machine to send messages through Postfix, different possibilities exist. Either you authorize the client machine's IP address to use the server to send a message, or you can use SMTP authentication. We'll see how to use the latter..

We must install the following packages:

```
#urpmi libsasl2 libsasl2-devel libsasl2-plug-plain libsasl2-plug-login
```

First of all, we have to generate an SSL certificate for Postfix.

```
# mkdir /etc/postfix/ssl
# cd /etc/postfix/ssl/
# openssl genrsa -des3 -rand /etc/hosts -out smtpd.key 1024
# chmod 600 smtpd.key
# openssl req -new -key smtpd.key -out smtpd.csr
# openssl x509 -req -days 3650 -in smtpd.csr -signkey smtpd.key -out
 smtpd.crt
# openssl rsa -in smtpd.key -out smtpd.key.unencrypted
# mv -f smtpd.key.unencrypted smtpd.key
# openssl req -new -x509 -extensions v3_ca -keyout cakey.pem -out
 cacert.pem -days 3650
```

Then, add the configuration options to Postfix in the /etc/postfix/main.cf file:

```
# cat /etc/postfix/main.cf
smtpd_sasl_auth_enable = yes
broken_sasl_auth_clients = yes
smtpd_tls_auth_only = no
smtp_use_tls = yes
smtpd_use_tls = yes
smtp_tls_note_starttls_offer = yes
smtpd_tls_key_file = /etc/postfix/ssl/smtpd.key
smtpd_tls_cert_file = /etc/postfix/ssl/smtpd.crt
smtpd_tls_CAfile = /etc/postfix/ssl/cacert.pem
smtpd_tls_loglevel = 1
smtpd_tls_received_header = yes
smtpd_tls_session_cache_timeout = 3600s
tls_random_source = dev:/dev/urandom
smtpd_recipient_restrictions = permit_mynetworks
 permit_sasl_authenticated
```

To test the good behavior of the TLS authentication, you can simply connect through telnet:

```
# telnet localhost 25
Trying 127.0.0.1...
Connected to localhost (127.0.0.1).
Escape character is '^]'.
220 localhost ESMTP Postfix (Mandriva MES5)
ehlo localhost
250-localhost
250-PIPELINING
250-SIZE 10240000
250-VRFY
250-ETRN
250-STARTTLS
250-AUTH PLAIN LOGIN
250 8BITMIME
```

If you see line `250-STARTTLS` as well as `S: 250-AUTH` that means that the SMTP authentication is active on your server. From this moment on, only authenticated users will be able to send messages via the SMTP server. Furthermore, the authentication is encrypted through SSL. Remember to configure your mail client accordingly.

8.3.3. Anti-Virus and Anti-Spam Solutions

Avoiding viruses and SPAM means to protect your mail server and your whole network. Here's how Postfix solves this issue. Before handling any message, it transmits incoming mail to the amavisd-new service. The latter, according to its configuration, handles the messages through an anti-virus and/or anti-spam utility. Once that's done, messages are returned in the incoming queue of Postfix which handles them to send them to the recipients.

8.3.3.1. Installing Software Components

1. Install amavisd-new

   Simply install the amavisd-new package. Be careful as the package has numerous dependencies, essentially associated with Perl modules.

2. Install SpamAssassin

   Simply install the spamassassin package.

3. Install the Clamav Anti-virus

   Install the following packages: clamav (common library used to access the anti-virus), clamd (anti-virus server daemon) and clam-db (clamav's anti-virus base).

8.3.3.2. Configuring amavisd-new

The configuration of amavis-d consists of specifying the following in the /etc/amavisd/amavisd.conf file:

- mail domain;
- ncoming/outgoing ports of messages;
- sending address of alerts;
- handling of messages containing viruses;
- specification of the anti-virus.

```
# cat /etc/amavisd/amavisd.conf
 ...
 $mydomain = 'mondomaine.com';
...
$inet_socket_port = 10024;   # listen on this local TCP port(s)
 (see $protocol)
...
# Modifier les adresses mails pour l'envoi des rapports virus / spam
 $virus_admin = "admin\@$mydomain";
 $mailfrom_notify_admin = "admin\@$mydomain";
 $mailfrom_notify_recip = "admin\@$mydomain";
 $mailfrom_notify_spamadmin = "admin\@$mydomain";
...

# Redirection des mails vers Postfix après traitement
$notify_method  = 'smtp:[127.0.0.1]:10025';
$forward_method = 'smtp:[127.0.0.1]:10025'; # set to undef with milter
...

#Traitement des spams et virus détectés
$final_virus_destiny = D_DISCARD;
$final_banned_destiny = D_BOUNCE;
$final_spam_destiny = D_PASS;
$final_bad_header_destiny = D_PASS;
 ...

#Décommenter les lignes correspondant à clamav
['ClamAV-clamd', \&ask_daemon, ["CONTSCAN {}\n",
 "/var/lib/clamav/clamd.socket"],qr/\bOK$/, qr/\bFOUND$/,
 qr/^.*?: (?!Infected Archive)(.*) FOUND$/ ],
...

#Décommenter les lignes correspondant à clamav
@av_scanners_backup = (

  ### http://www.clamav.net/   - backs up clamd or Mail::ClamAV
  ['ClamAV-clamscan', 'clamscan',
    "--stdout --disable-summary -r --tempdir=$TEMPBASE {}", [0], [1],
    qr/^.*?: (?!Infected Archive)(.*) FOUND$/ ],
...
```

VMake sure that the `clamd` daemon is configured to be launched automatically:

```
#chkconfig -list clamd
clamd         0:stop         1:stop         2:stop         3:run        4:run
```

8.3.3.3. Configuring Postfix

To take into account `amavisd-new` in Postfix's configuration, we must modify `main.cf` and `master.cf`:

```
# cat /etc/postfix/main.cf
 ...
# Handing of incoming messages
 content_filter=smtp-amavis:[127.0.0.1]:10024
 smtp-amavis_destination_concurrency_limit=2

 # cat /etc/postfix/master.cf
 ...
```

```
127.0.0.1:10025 inet n - y - - smtpd
-o content_filter=
-o smtpd_restriction_classes=
-o smtpd_client_restrictions=
-o smtpd_helo_restrictions=
-o smtpd_sender_restrictions=
-o smtpd_recipient_restrictions=permit_mynetworks,reject
-o mynetworks=127.0.0.0/8
-o mynetworks_style=host
-o strict_rfc821_envelopes=yes
-o receive_override_options=no_unknown_recipient_checks,
  no_header_body_checks
-o smtpd_client_connection_limit_exceptions=127.0.0.0/8
...
```

Restart Postfix to take modifications into account. To verify its behaviour:

```
# telnet localhost 10024
Trying 127.0.0.1...
Connected to localhost (127.0.0.1).
Escape character is '^]'.
220 [127.0.0.1] ESMTP amavisd-new service ready
 # telnet localhost 10025
Trying 127.0.0.1...
Connected to localhost (127.0.0.1).
Escape character is '^]'.
220 example.com Welcome on example.com mail server
```

8.3.3.4. Configuring SpamAssassin

The main configuration file is `/etc/mail/spamassassin/local.cf`:

```
# cat /etc/mail/spamassassin/local.cf
required_hits 5
rewrite_header Subject [SPAM]
report_safe 0
auto_whitelist_path        /var/spool/spamassassin/auto-whitelist
auto_whitelist_file_mode   0666
dcc_home                   /var/lib/dcc
auto_learn 1
use_razor2 1
use_bayes 1
```

LYou should refine the configuration in regards with the server's desired usage: the stricter SpamAssassin's rules are, the more you're bound to get false positives.

To enhance SpamAssassin's SPAM detection, it's good to understand its behaviour as well as the its configuration's different instructions. During its analysis, SpamAssassin observes the email in its ensemble and gives it a score according to different parameters. For example, if the email contains a lot of keywords usually used in SPAM, its score augments. Furthermore, SpamAssassin, uses a system of complex rules which allows it to model the score according to the type of mail generally handled by the server. Through the server's "progressive learning", this allows you to get precise detection with few false positives.

The `required_hits` instruction defines the score needed for an email to be considered as SPAM. In this case, the `rewrite_header` instruction will be taken into account, and will add the "[SPAM]" tag to the email's subject line. The `use_bayes` instruction indicates to SpamAssassin to use its different internal calculation rules based on its learning. Every mail handled by SpamAssassin will contain in its header a trace of that handling. In your mail client, you can display this information.

```
X-Spam-Status: No, score=-2.6 required=5.0 tests=BAYES_00
   autolearn=disabled version=3.0.4
```

Don't hesitate to look at the score that mails labeled as SPAM (but that weren't) so as to modify the `required_hits` value. Another solution to enhance the score calculation is to define preferred languages for emails. If you receive many emails in English, instruct SpamAssassin with the following instructions:

```
ok_languages en      ok_locales en
```

8.3.3.5. Launching SpamAssassin's Auto-Learning

Here's another way to refine SpamAssassin's behaviour rules by sending it messages you consider to be SPAM.

1. Create a box accessible by all users:

   ```
   # cyradm --user cyrus localhost
   IMAP Password:
   localhost> cm user.SPAM
   localhost> dam user.SPAM SPAM
   localhost> sam user.SPAM anyone all
   localhost> lam user.SPAM
   anyone lrswipcda
   localhost>
   ```

   Users will move all the SPAM they receive on their own accounts to that mailbox.

2. Schedule a script which will launch the auto-learning from those mails, and then destroy them:

   ```
   #Running auto-learn
   ```

```
for i in /var/spool/imap/s/user/SPAM/[0-9]*.; do sa-learn
 --showdots --spam $i; done
#Suppression des messages de la boîte
rm -f /var/spool/imap/s/user/SPAM/[0-9]*.
Reconstruire la base Cyrus
su -l cyrus -c "/usr/lib/cyrus-imapd/reconstruct user.SPAM "
```

The script should be placed in a `crontab` and launched regularly.

Chapter 9. Monitoring

Monitoring softwares allows to monitor services and networks in production.

9.1. Cacti

9.1.1. General review of Cacti

Cacti is a monitoring software based on RRDtool. It is used for capacity planning and metrology (networks, disks, ...) while real-time monitoring role is generally assigned to Nagios.

For more informations, Cacti website is http://www.cacti.net/ and official documentation is available at http://docs.cacti.net/wiki:documentation

9.1.2. Configure the server to make it accessible by Cacti

You have to install net-snmp package:

```
#urpmi net-snmp
```

Activate monitoring of all the discs on the host. This requires editing the `/etc/snmp/snmpd.conf` file, search directive `disk / 10000`, comment on them and add the line:

```
includeAllDisks 5
```

Finish by launching the service SNMP service:

```
#service snmpd start
```

9.1.3. Install and configure a Cacti server

There are pre-requisites for installing Cacti. An HTTP server (Apache is already installed by default in Mandriva Enterprise Server 5) and an DBMS are used. MySQL is used as an example later in this document.

Let's start by installing Cacti package:

```
#urpmi cacti
```

Then, create the MySQL database:

```
#mysql -u root
```

or, if you put a password (highly recommended):

```
#mysql -u root -p
```

and when you are connected to MySQL:

```
create database cacti;
grant all on cacti.* to cactiuser@localhost identified by 'cactiuser';
flush privileges;
quit
```

Finally, you have to create the tables in the database using the SQL script provided:

```
mysql -u root cacti < /usr/share/cacti/sql/cacti.sql
```

Cacti's configuration is done with the Installation Guide directly into a web browser at http://MES5_SERVER_IP/cacti. The default choices are correct for a typical configuration.

Once configured, you then arrive on the Cacti connection interface. By default, the username, the password is "admin". He is asked to change it to the first connection, do.

You are now connected to the Cacti interface and can start preparing machines supervision. Start by monitoring local server.

The interface of Cacti is not localized. In other words, it is only in English.

9.1.4. Monitoring the local server

First, you have to configure SNMP daemon (cf. Configure the server to make it accessible by Cacti).

Log onto Cacti's Web interface administration: http://MES5_SERVER_IP/cacti et follow these basic steps:

1. Add a device

   The first step in creating graphs for your network is adding a device for each network device you want to monitor. A device specifies important details such as the hostname of the network, SNMP parameters, and host type.

   Go to the Devices menu and click Add.

   The first field Description is used to identify the server to the user. On the other side, the Hostname must be either the FQDN of the server or its IP address. Then select the template to use: in the case of a local host, it is possible to use Local Linux Machine, but to check the behavior of the SNMP daemon, consider using ucd/net SNMP Host. The SNMP version used is version 2.

   Summary of configuration choices:

   ```
   Description: Cacti Server
   Hostname: 127.0.0.1
   Host Template: ucd/net SNMP Host
   SNMP Version: Version 2
   Downed Device Detection: Ping and SNMP
   ```

   Once selected, click on Create. The next page displays the results of SNMP and Ping tests.

2. Create graphs for the new device

   Now, you have to add different data types you want to monitor the device. To do this, click on Create Graphs for this Host and select the desired Graph Templates. You should have the availability of the following graphs: network interfaces Traffic (bits/sec), CPU usage, Memory usage, and Available Disk Space. When your choice is made, click on Create.

3. Include the device in a Graph Tree

   To view charts of the new device, you must add it to a Graph Tree. To do this, click on Menu Graph Trees and Default Tree. Change the Name with a more descriptive name: for example Servers, then click Add to add an item. Tree Items page appears, select Host as Tree Item Type and Cacti Server that you created earlier as Host. Finish by clicking on Create.

You can now view the results by clicking on the tab above named Graphs. Under the Servers menu, clicking on host: Cacti Server, you then view the expected graphs.

Wait a few minutes, so that some data are lifts to be visible on the graphs.