

**Get Ready for Open Source:
SUSE Linux Enterprise Desktop
Book 2**

Novell Training Services

2007

Proprietary Statement

Copyright © 2007 Novell, Inc. All rights reserved. The reproduction, photocopying, storing on a retrieval system, or transmitting of this manual is protected under a Creative Commons Attribution-NonCommercial-ShareAlike 2.5 license.

You are free to share (copy, distribute, display, and perform the work) and to remix (make derivative works) under the following conditions: you must attribute the work in the manner specified by the author or licensor, you may not use this work for commercial purposes. In addition, if you alter, transform, or build upon this work, you may distribute the result work only under a license identical to this one.

For any reuse or distribution, you must make clear to others the license terms of this work. For clarification or to apply for a waiver to any of these conditions, contact Novell, Inc.

1800 South Novell Place
Provo, UT 84606-2399

Disclaimer

Novell, Inc. makes no representations or warranties with respect to the contents or use of this manual, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose.

Further, Novell, Inc. reserves the right to revise this publication and to make changes in its content at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any NetWare software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose.

Further, Novell, Inc. reserves the right to make changes to any and all parts of NetWare software at any time, without obligation to notify any person or entity of such changes.

This Novell Training Manual is published solely to instruct students in the use of Novell networking software. Although third-party application software packages are used in Novell training courses, this is for demonstration purposes only and shall not constitute an endorsement of any of these software applications.

Further, Novell, Inc. does not represent itself as having any particular expertise in these application software packages and any use by students of the same shall be done at the students' own risk.

Software Piracy

Throughout the world, unauthorized duplication of software is subject to both criminal and civil penalties.

If you know of illegal copying of software, contact your local Software Antipiracy Hotline.

For the Hotline number for your area, access Novell's World Wide Web page at <http://www.novell.com> and look for the piracy page under "Programs".

Or, contact Novell's anti-piracy headquarters in the U.S. at 800-PIRATES (747-2837) or 801-861-7101.

Acknowledgements

Michael Eicks (Novell Inc.)

Trademarks

Novell, Inc. has attempted to supply trademark information about company names, products, and services mentioned in this manual. The following list of trademarks was derived from various sources.

Novell, Inc. Trademarks

Novell, the Novell logo, NetWare, BorderManager, ConsoleOne, DirXML, GroupWise, iChain, ManageWise, NDPS, NDS, NetMail, Novell Directory Services, Novell iFolder, Novell SecretStore, Ximian, Ximian Evolution and ZENworks are registered trademarks; CDE, Certified Directory Engineer and CNE are registered service marks; eDirectory, Evolution, exteNd, exteNd Composer, exteNd Directory, exteNd Workbench, Mono, NIMS, NLM, NMAS, Novell Certificate Server, Novell Client, Novell Cluster Services, Novell Distributed Print Services, Novell Internet Messaging System, Novell Storage Services, Nsure Resources, Nterprise, Nterprise Branch Office, Red Carpet and Red Carpet Enterprise are trademarks; and Certified Novell Administrator, CNA, Certified Novell Engineer, Certified Novell Instructor, CNI, Master CNE, Master CNI, MCNE, MCNI, Novell Education Academic Partner, NEAP, Ngage, Novell Online Training Provider, NOTP and Novell Technical Services are service marks of Novell, Inc. in the United States and other countries. SUSE is a registered trademark of SUSE LINUX AG, a Novell company. For more information on Novell trademarks, please visit <http://www.novell.com/company/legal/trademarks/tmlist.html>.

Other Trademarks

Adaptec is a registered trademark of Adaptec, Inc. AMD is a trademark of Advanced Micro Devices. AppleShare and AppleTalk are registered trademarks of Apple Computer, Inc. ARCServ is a registered trademark of Cheyenne Software, Inc. Btrieve is a registered trademark of Pervasive Software, Inc. EtherTalk is a registered trademark of Apple Computer, Inc. Java is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries. Linux is a registered trademark of Linus Torvalds. LocalTalk is a registered trademark of Apple Computer, Inc. Lotus Notes is a registered trademark of Lotus Development Corporation. Macintosh is a registered trademark of Apple Computer, Inc. Netscape Communicator is a trademark of Netscape Communications Corporation. Netscape Navigator is a registered trademark of Netscape Communications Corporation. Pentium is a registered trademark of Intel Corporation. Solaris is a registered trademark of Sun Microsystems, Inc. The Norton AntiVirus is a trademark of Symantec Corporation. TokenTalk is a registered trademark of Apple Computer, Inc. Tru64 is a trademark of Digital Equipment Corp. UnitedLinux is a registered trademark of UnitedLinux. UNIX is a registered trademark of the Open Group. WebSphere is a trademark of International Business Machines Corporation. Windows and Windows NT are registered trademarks of Microsoft Corporation.

All other third-party trademarks are the property of their respective owners.

Contents

- 1 Administer Linux with YaST 3**
 - 1.1 Manage the Network 4
 - 1.1.1 Change Hostname Domain and DNS Information 8
 - 1.1.2 Change Routes 10
 - 1.1.3 View your Network Settings 12
 - 1.2 Manage Users and Groups 14
 - 1.2.1 User Administration 16
 - 1.2.2 Administer Groups 20
 - 1.3 Install and Remove Software 23
 - 1.4 Configure a Printer 28
 - 1.4.1 New Queue for Existing Printer 29
 - 1.4.2 Directly Connected Printers 31
 - 1.4.3 Network Printers 33

- 2 The Linux Shell and Tools 35**
 - 2.1 Organize Files at the Command Line 37
 - 2.1.1 Move in the File System 38
 - 2.1.2 List Directory Content 41
 - 2.1.3 Move a File 45
 - 2.1.4 Copy a File 47
 - 2.1.5 Create a Directory 50
 - 2.1.6 Delete Directories and Files 51
 - 2.1.7 Link Files 53
 - 2.1.8 Find Files 56
 - 2.1.9 Archive Files 63
 - 2.1.10 Manage File Permission and Ownership 69
 - 2.2 Understand bash Commands and Features 75
 - 2.2.1 Use Completion of Commands and Filenames 75
 - 2.2.2 Use the bash History 76
 - 2.2.3 Use bash Variables 78

2.2.4	Use bash Aliases	80
2.2.5	Use Search Patterns to Expand Names	82
2.3	Understand Advanced Command Execution	85
2.3.1	Use Piping and Redirection	85
2.3.2	Understand Regular Expressions	89
2.3.3	Manage Shell Processes	98
2.4	The vi Text Editor	102
2.4.1	Start vi	103
2.4.2	Move the Cursor	104
2.4.3	Learn the Working Modes	105
2.4.4	Learn Advanced vi Commands	110
2.4.5	Create Macros	112
2.4.6	Create Abbreviations	114
2.5	The Basics of bash Programming	115
2.5.1	Variables in bash Scripts	116
2.5.2	Compute Numbers	118
2.5.3	Comparing Files, Strings and Numbers	122
2.5.4	Read from Input	132
2.5.5	Examine Cases	134
2.5.6	Use Loops	138
2.5.7	Use Program Options and Parameters	143
2.5.8	Create Shell Functions	146
2.6	Understand the sed Stream Editor	151
3	Linux Services for Developers	161
3.1	The Version Control System Subversion	161
3.1.1	Install Subversion	163
3.1.2	Set up a Subversion Repository	165
3.1.3	Start a Project	167
3.1.4	Prepare for Working with Subversion	168
3.1.5	Use Subversion	170
3.1.6	Manage Subversion	178
3.1.7	Use Subversion in a Network	181
3.2	The Web Server Apache	183
3.2.1	Install Apache	184
3.2.2	Start and Test the Web Server	186
3.2.3	Locate the Document Root of the Web Server	188
3.2.4	Locate the Apache Configuration Files	188

3.2.5	Understand the Basic Rules of the Configuration Files	189
3.2.6	Understand the Default Apache Configuration . . .	191
3.2.7	Limit Access to the Web Server	192

Appendix **198**

A The rpm Command **199**

Index **202**

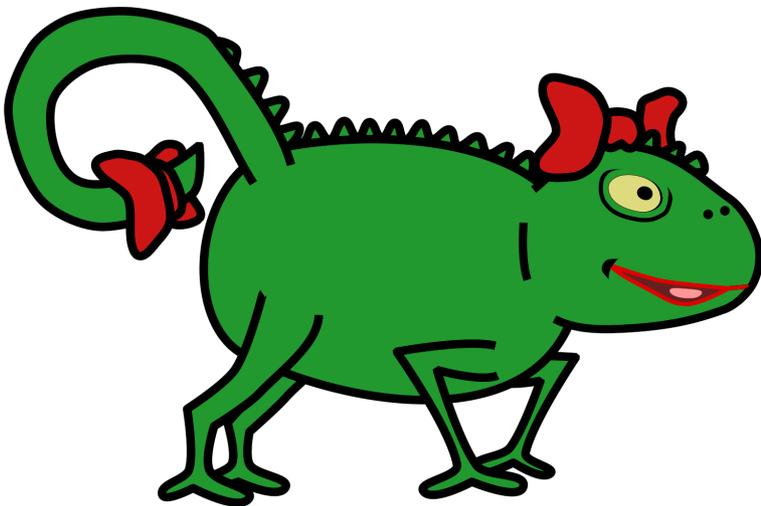
Hello, my name is *Suzie*...

...and I will help Geeko guide you through this course. The course is called “Get Ready for Open Source: SUSE Linux Enterprise Desktop – Book 2.”

In the last course, you learned a lot about open source software and Linux. You learned about important applications on Linux that can help you at school.

In this course, you will get experience as a software developer. But first you have to learn some important administration tasks.

This course is based on Novell SUSE Linux Enterprise Desktop 10. You will find all the software you need on the course CD.



1 Administer Linux with YaST

You can use the tool YaST to complete many configuration tasks on your SUSE Linux Enterprise Desktop. YaST stands for “Yet another Setup Tool.”

To start the graphical user interface, select the YaST icon in the **System** group of the Application Browser.



Then, enter the root password.



The YaST Control Center dialog appears.



In the left frame of the dialog there is a list of module categories. If you select an item in this list, you can see the modules in the right frame.

There are a lot of modules available. In this manual, you will learn the modules to perform the following tasks:

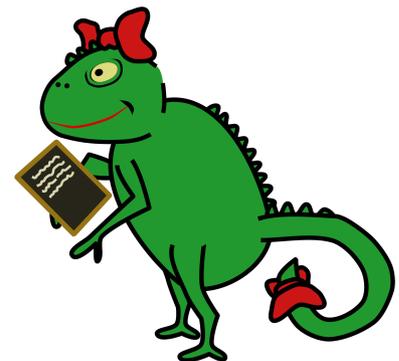
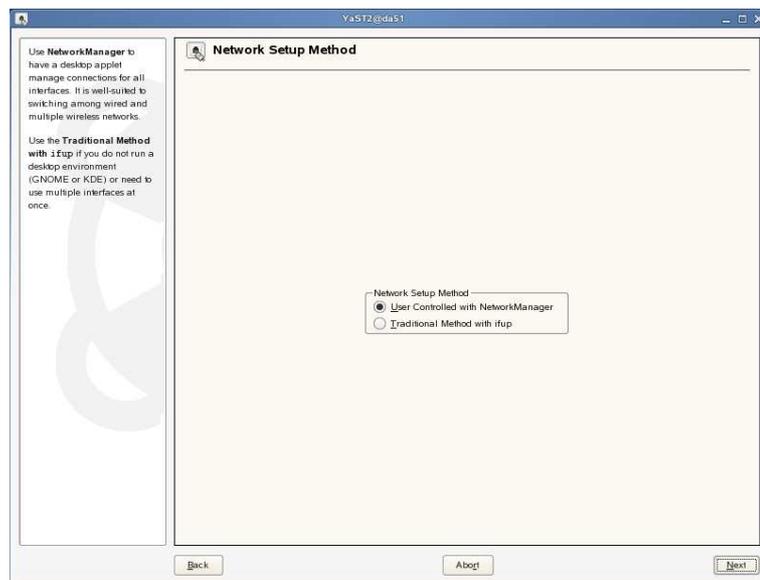
- ☞ Manage the network
- ☞ Manage users and groups
- ☞ Install and remove software
- ☞ Configure a printer

1.1 Manage the Network

To start the network configuration module of YaST, select **Network Devices** → **Network Card**.



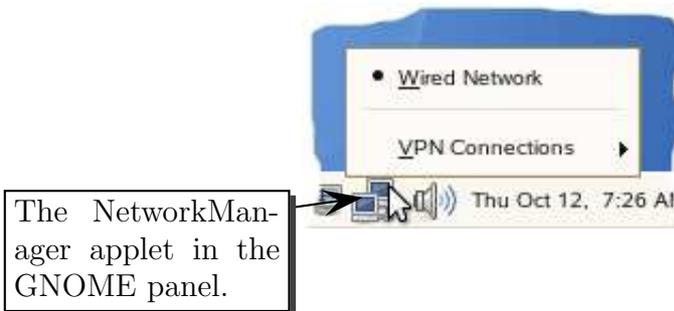
In the first dialog, select the setup method.



You can choose between the following network setup methods:

User Controlled with NetworkManager. By default, only the system administrator root is allowed to change the network configuration. If you are running SUSE Linux Enterprise Desktop on a laptop computer and often changing networks (e. g., home network, school network, cafe network), you often have to switch to user root and back to normal user. To avoid this, use the NetworkManager.

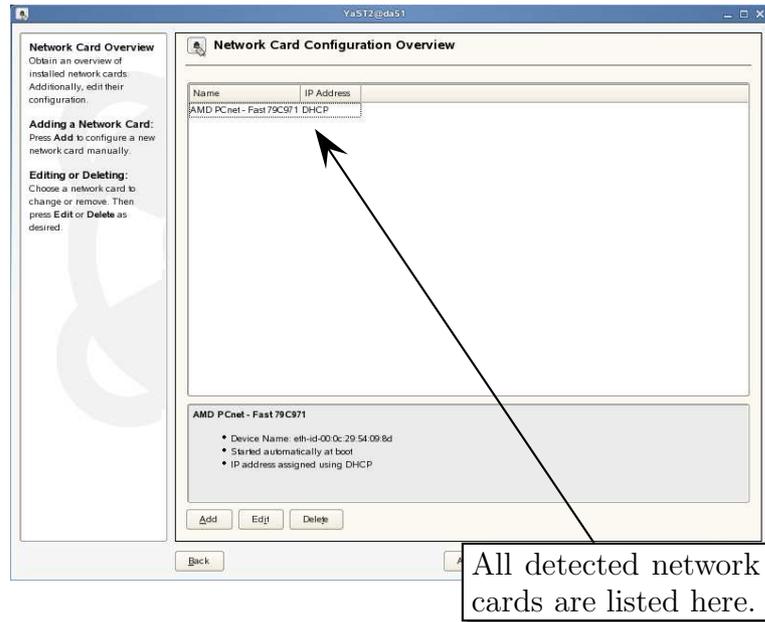
With the NetworkManager method a normal user is allowed to change the network configuration. To do this, you have to use a desktop applet that manages the connections for all network interfaces.



Traditional Method with ifup. The traditional method uses the command `ifup` to set up a network interface. To change the network configuration, you have to have root permission.

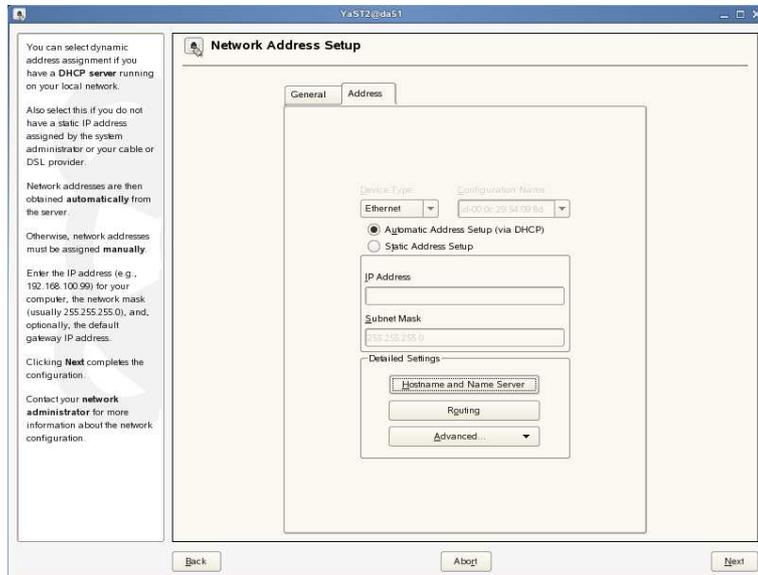
The traditional method is recommended if you run your SUSE Linux Enterprise Desktop on a desktop machine, because you do not have to change the network configuration very often on such a machine.

Select **Next** to go to the Network Card Configuration Overview.



If you want to use the default settings, select **Next**. Otherwise, select the card you want to configure; then select **Edit**.

The Network Address Setup dialog appears.



In this dialog, enter the following information to integrate the network device into an existing network:

Automatic Address Setup (via DHCP). Select this option if the network card should receive an IP address from a DHCP server. This server must be available in the network.

Static Address Setup. If you choose this option, you need to enter the IP address of the network interface or of the computer in the network under **IP Address**.

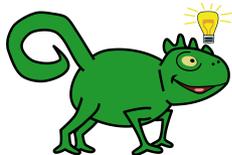
Each computer in the network has at least one address for each network interface, which must be unique in the entire network. According to the currently valid standard (IPv4), this address consists of a sequence of four bytes, separated by dots (such as “10.10.0.69”).

When choosing the IP address, you need to know if the computer will be directly connected to the Internet. In this case, use an assigned official IP address. Otherwise, use an address from a private address space.

Subnet Mask. The network mask (referred to as subnet mask in YaST), determines in which network an IP address is located. The mask divides the IP address into a network part and a host part, thus defining the number of possible computers in a network. All computers within the same network can reach each other directly (without a router in between).

Hostname and Name Server. Computers in the network can be addressed directly by using their IP addresses or a unique name. A name server (DNS) has to be there to convert names into IP addresses and vice versa.

Routing. If the computer should only reach computers in the same subnet, you need not enter any routes.



Exercise: Understand IP Addresses

Search the internet for “network class” and “network mask”

to understand how an IP address is composed.

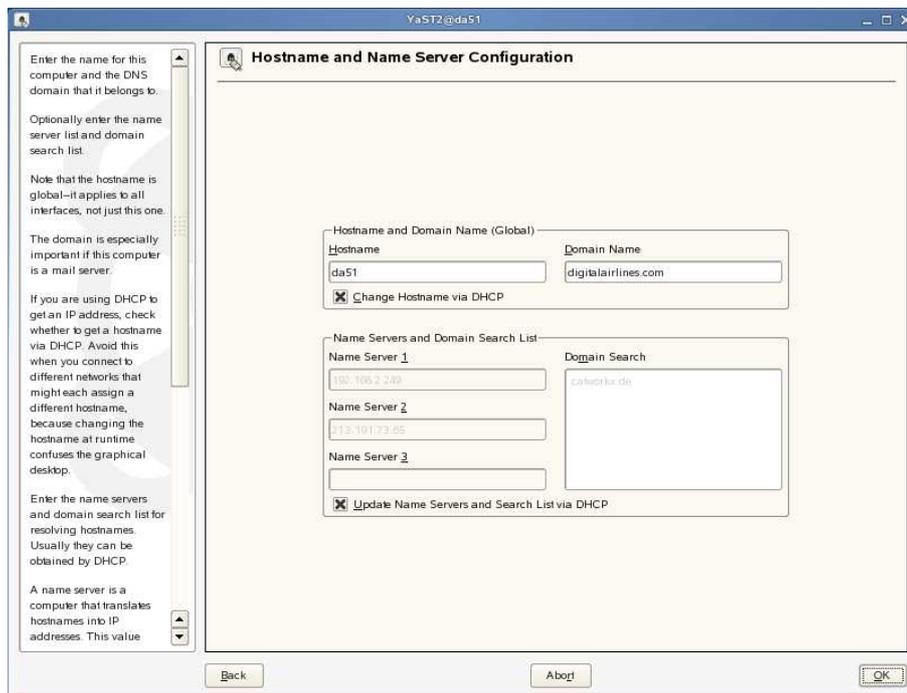
1.1.1 Change Hostname Domain and DNS Information

If you are using DHCP and you select **Hostname and Name Server**, the following appears:



If you want to change data delivered by DHCP (e. g., IP number), select **Modify**. But if you want to modify other information, select **Accept**.

The Hostname and Name Server Configuration dialog appears.



In this dialog, you can enter the following:

Hostname. Enter a name for the computer. This name should be unique within the network.

Domain Name. This is the name of the DNS domain the computer belongs to. Domains help to divide networks. All computers in a defined organizational area normally belong to the same domain.

A computer can be addressed uniquely by giving its FQDN (Fully Qualified Domain Name). This consists of the host name and the domain name, such as da51.digitalairlines.com. In this example, the domain is digitalairlines.com.

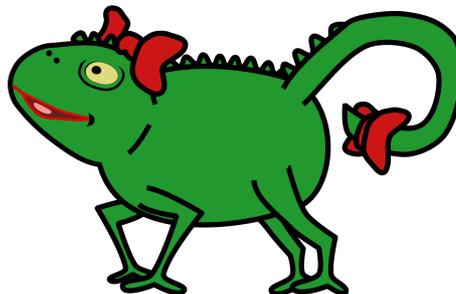
List of name servers. To address other computers in the network with their host names, identify the name server, which converts computer names into IP addresses and vice versa.

You can specify a maximum of three name servers.

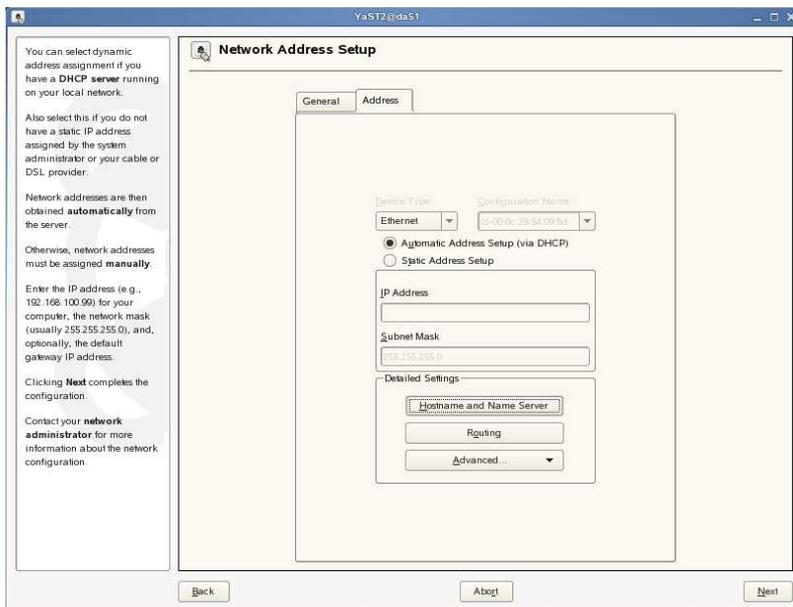
Domain search list. In the local network, you can address other hosts with their host names. You do not have to use the FQDN, because the domain search list specifies the domains used by the system to expand the host name to the FQDN.

This FQDN is then passed to the name server to be converted. For example, da51 is completed by using digitalairlines.com to the FQDN da51.digitalairlines.com. Then this name is passed to the name server to be converted.

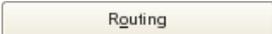
If the search list contains more than one domain, the completion takes place one after the other. The FQDN is passed to the name server until an entry returns an associated IP address. Separate the domains with a comma or white space.

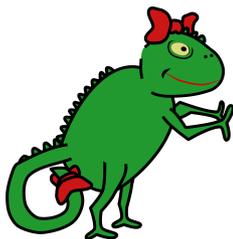
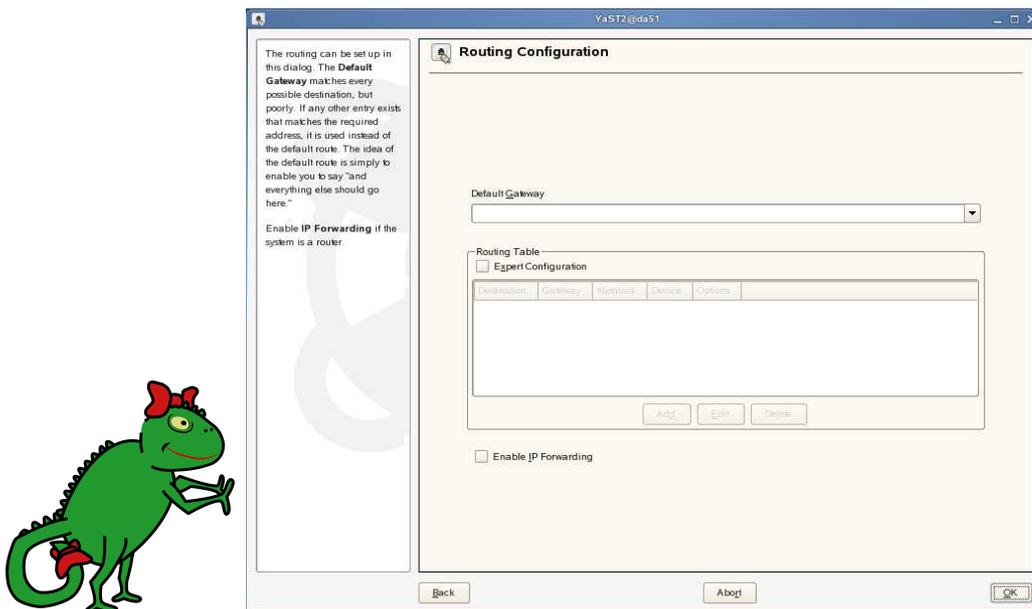


Select to go back to the Network Address Setup dialog.



1.1.2 Change Routes

If you need to enter a default gateway (a computer that forwards information from one network to other networks) or create a routing table, select  in the Network Address Setup dialog. The following appears:



You can define the following:

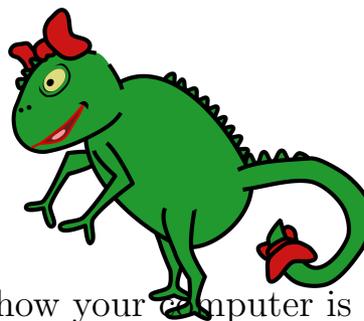
Default Gateway. If the network has a gateway, its address can be specified in the network configuration.

All data not addressed to the local network is then forwarded directly to the gateway.

Routing Table. You can create entries in the routing table of the system after selecting **Expert Configuration**.

Enable IP Forwarding. If you select this option, IP packages that are not dedicated for your computer are routed.

Select to go back to the Network Address Setup dialog.



Exercise: Manage the Network

Use the YaST network module to view how your computer is connected to the network.

- ❶ Does it use DHCP or does it have a static IP address?

- ❷ What is your current IP address?

- ❸ What is your DNS server?

- ❹ What is your gateway?

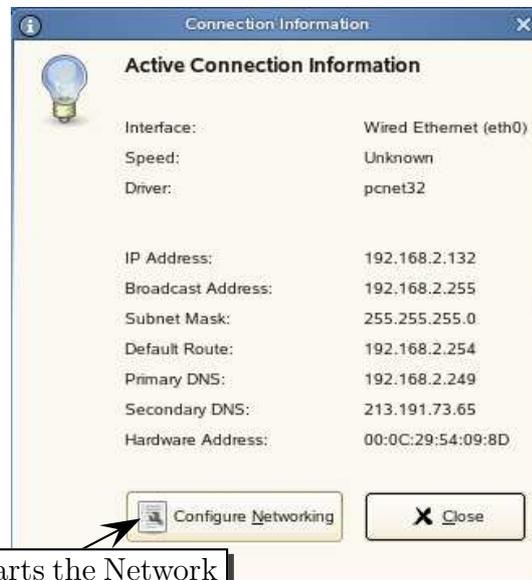
1.1.3 View your Network Settings

After you have saved the configuration with YaST, you can use the ethernet card in your computer. If you are using the NetworkManager method, the NetworkManager applet in GNOME panel informs you when the network connection is established successfully.

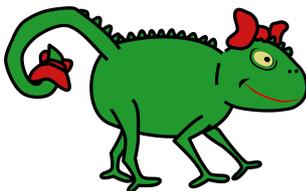


To see the data of your network connection, right-click the NetworkManager icon in GNOME panel and select **Connection Information** from the pop-up menu.

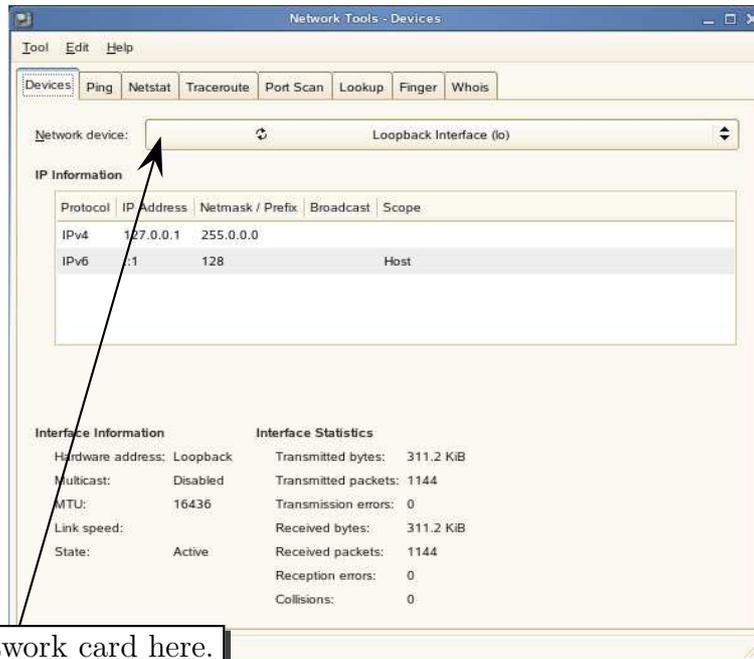
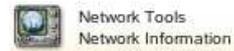
A dialog appears:



This button starts the Network Card module of YaST directly.



If you do not use NetworkManager, you can get the data of your network connection by selecting **Network Tools** in the **System** group of the **Application Browser**.

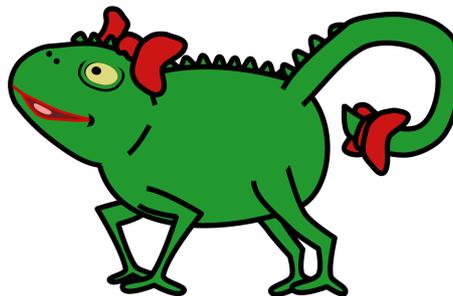


The information on your network connection is written in the **Devices** tab. If you have problems with your connection, you can use the other tabs to test your network.



Exercise: View Your Network Settings

Verify the information from the last exercise by using the **Network Tools**.



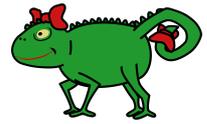
1.2 Manage Users and Groups

A Linux operating system handles several users at the same time (multiuser) and allows these users to perform several tasks on the same computer simultaneously (multitasking).

For this reason, every user has to log in before using the system. To log in, you need the following information:

- ☞ A username
- ☞ A password

As the operating system can operate numbers much better than strings, users are handled internally as numbers. This number is called UID (User ID).



Every Linux system has a privileged user, called *root*, the administrator of the system. This user always has the UID 0. Users can be grouped according to characteristics or activities. For example:

- ☞ Normal users are usually in the group *users*.
- ☞ All users who create web pages can be grouped in the group *webedit*.

The file permissions for the directory containing the web pages has to be set in the way the group *webedit* is able to write (save files).

Each group is also handled by the system with a number, called the GID (Group ID). The following group types are available:

- ☞ Normal groups (GID \geq 100)
- ☞ Groups used by the system (GID $<$ 100)
- ☞ The root group (GID = 0)

YaST has a module to manage users and groups. You can access YaST user and group account administration in the following ways:

User administration. From the YaST Control Center, select **Security and Users** → **User Management**.



User Management

Group administration. From the YaST Control Center, select **Security and Users** → **Group Management**.



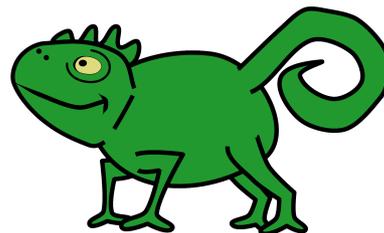
Group Management



Exercise: Manage Users and Groups

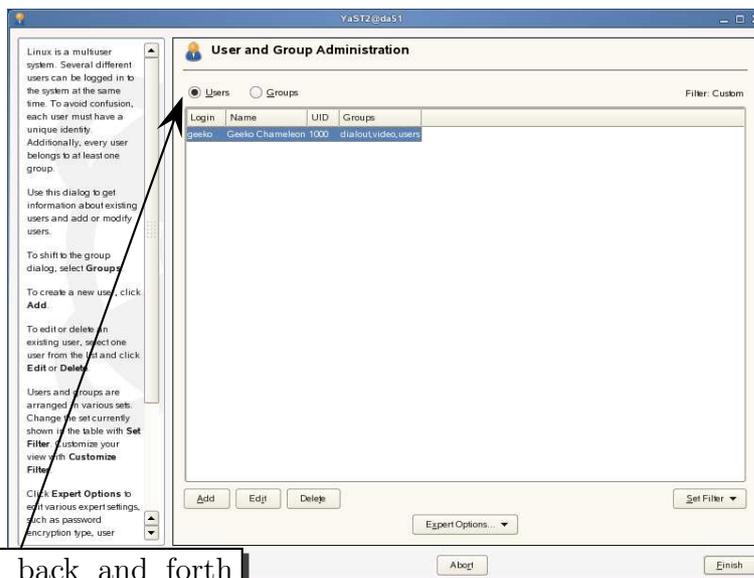
Consider six groups at your school that need different permission.

- ① _____
- ② _____
- ③ _____
- ④ _____
- ⑤ _____
- ⑥ _____



1.2.1 User Administration

The user account management window lists the existing user accounts (as in the following):



You can switch back and forth between administering users and administering groups by selecting these radio buttons.

A list of users (accounts on your server) appears with information such as login name, full name, UID, and groups each user belongs to.

If you want to change the users listed, select **Set Filter**; then select one of the following:

Local Users. User accounts you have created on your local computer for logging into the computer

System Users. User accounts created by the system to run services

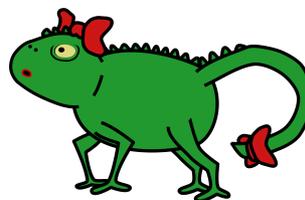
Custom. A customized view of users based on the settings configured with **Customize Filter**

Customize Filter. This option lets you combine user sets (such as **Local Users** and **System Users**) to display a customized view (with **Custom**) of the user list

Additional sets of users are added to the **Set Filter** drop-down list as you configure and start services on your server.

You can create a new user account or edit an existing account by selecting

or . The following appears:



Enter or edit information in the following fields:

User's Full Name. Enter a user name (such as Geeko Chameleon).

Username. Enter a user name that is used to log in to the system (such as geeko).

Password and Confirm Password. Enter (twice) a password for the user account. When entering a password, distinguish between uppercase and lowercase letters.

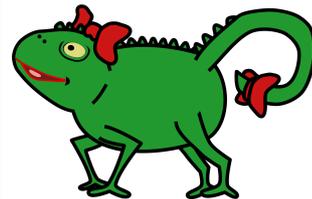
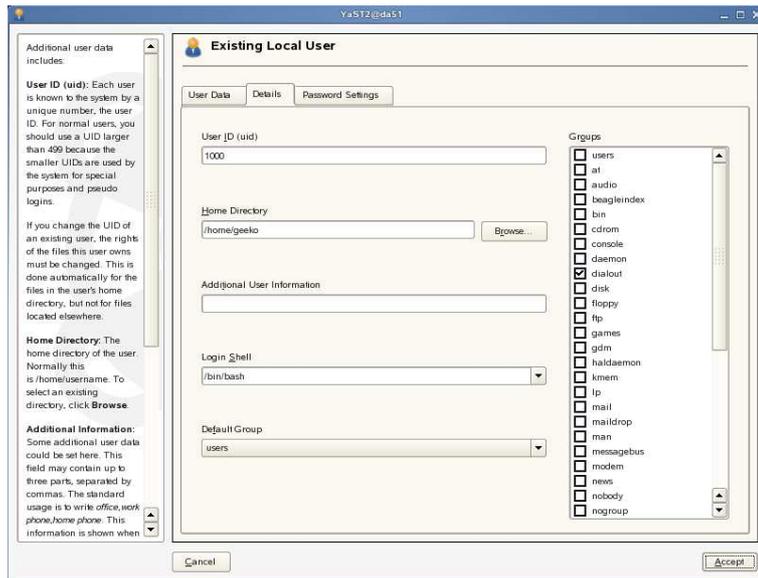
Valid password characters include letters, digits, blanks, and `#* , . ; : . _ - + ! $ % & / | ? { [()] } =`.

The password should not contain any special characters (such as accented characters), as you might find it difficult to type these characters on a different keyboard layout when logging in from another country.

With the current password encryption (Blowfish), the password length should be between 5 and 72 characters.

Disable User Login. If you select this option, the user is not able to login anymore. But the user account is not deleted. This option is useful if you want to disable a user account temporarily.

To set the properties of the user (such as the UID, the home directory, the login shell, group affiliation, and additional user account comments), select the **Details** tab. The following appears:



Enter or edit information in the following fields:

User ID (uid). For normal users, you should use a UID greater than 999, because the lower UIDs are used by the system for special purposes and pseudo logins.

If you change the UID of an existing user, the permission of the files of this user must be changed, too. This is done automatically for the files in the user's home directory, but not for files located elsewhere.

Home Directory. The home directory of the user. Normally this is `/home/username`.

You can select an existing directory by selecting .

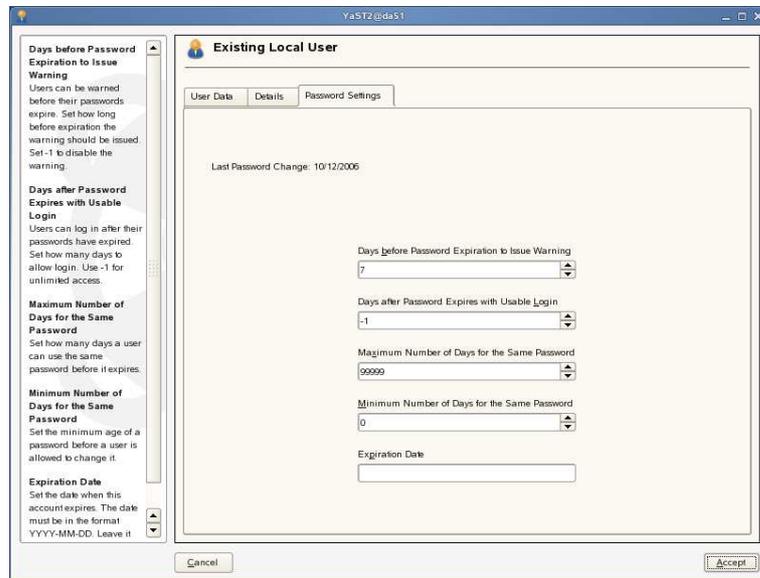
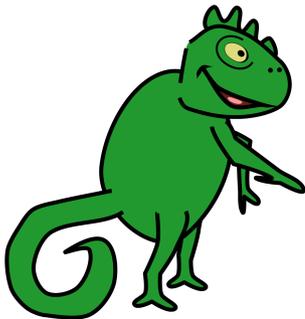
Additional User Information. This field contains up to three parts separated by commas. It is often used to enter such information as office, work phone, or home phone.

Login Shell. From the drop-down list, select the default login shell for this user.

Default Group. This is the group the user belongs to. Select a group from the list of all groups configured on your system.

Groups. Select all additional memberships you want to assign to the user.

To set various password parameters (such as duration of a password), select the **Password Settings** tab. The following appears:



Enter or edit information in the following fields:

Days before Password Expiration to Issue Warning.

Enter the number of days before the password expires that the user is informed about it.

Enter -1 to disable the warning.

Days after Password Expires with Usable Login. Enter the number of days after the password expires that users can continue to log in.

Enter -1 for unlimited access.

Maximum number of days for the same password. Enter the number of days a user can use the same password before it expires.

Minimum number of days for the same password. Enter the minimum age of a password before a user can change it.

Expiration date. Enter the date when the account expires. The format is `YYYY-MM-DD`.

Leave the field empty if the account never expires.

Save the settings for the new or edited user by selecting .

A new user appears in the list.

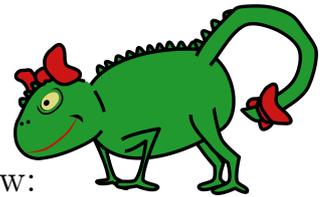
Configure your server with the new settings by selecting .



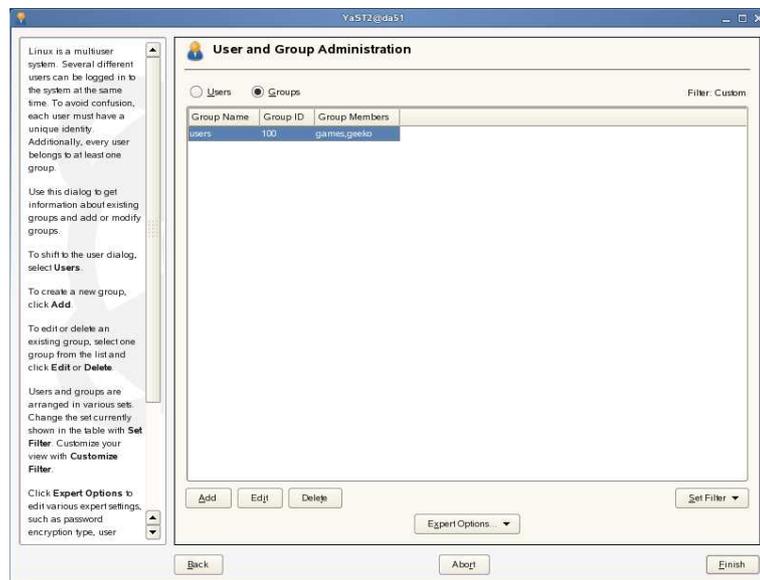
Exercise: Create a New User

Create a new user account on your local computer. This account is needed for exercises later in this book.

1.2.2 Administer Groups



You can administer groups from the following window:



A list of groups appears with information such as Group Name, Group ID (GID), and Group Members.

To change the groups listed, select **Set Filter**; then select one of the following:

Local Groups. Groups created on your local server to provide permissions for members assigned to the group.

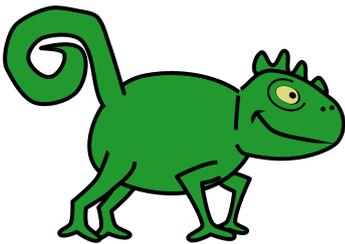
System Groups. Groups created by the system to run services.

Custom. A customized view of groups based on the settings configured with **Customize Filter**.

Customize Filter. This option lets you combine group sets (such as **Local Groups** and **System Groups**) to display a customized view (with **Custom**) of the groups list.

Additional sets of groups are added to the **Set Filter** drop-down list when you configure and start services on your server.

You can create a new group or edit an existing group by selecting **Add** or **Edit**. If you select **Edit**, the following dialog appears.



The fields allow you to enter or edit information:

Group Name. The name of the group. Avoid long names. Default name lengths are between two and eight characters.

Group ID (gid). The number assigned to the group. The number must be between 0 and 60000. GIDs up to 99 represent system groups. GIDs beyond 99 can be used for normal groups. YaST warns you if you try to create a GID that already exists.

Password. (Optional) If you want the members of the group to identify themselves while switching to this group (see `man newgrp`), assign a password.

For security reasons, the password is represented by dots.

Confirm Password. Enter the password once again to avoid typing errors.

Group Members. Select which user should be a group member.

A second list may appear (if you select). It shows for which users this group is the primary group. This list cannot be edited by YaST.

After you have finished entering or editing the group information, select . You are returned to the Group Administration dialog. Save the configuration settings by selecting .

The information you enter when creating or editing users and groups with YaST is saved in the following user administration files:

 /etc/passwd

 /etc/shadow

 /etc/group



Exercise: Administer Groups

How many local groups and system groups are available on your computer?

1.3 Install and Remove Software

After performing a standard SUSE Linux Enterprise Desktop 10 installation, you will often have to install additional software.

RPM (*RPM Package Management*) is the default file type for software to install. It contains the software and a lot of information (e. g., dependencies to other software packages).

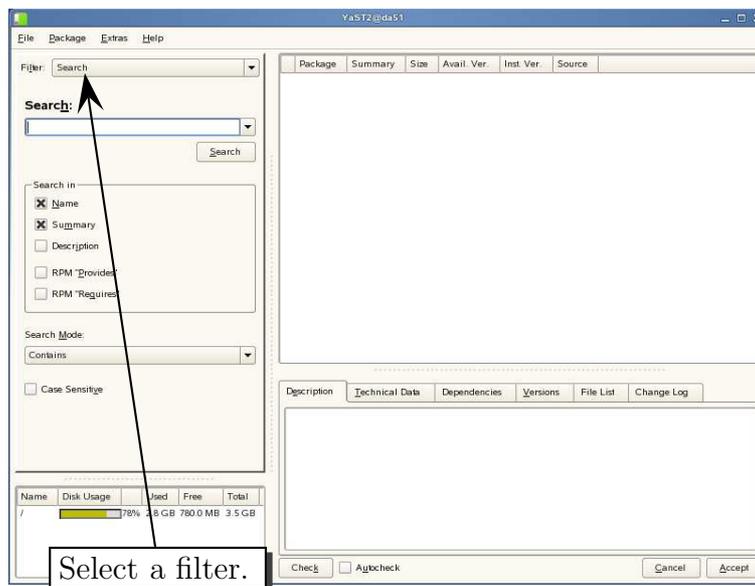


Start the YaST module **Software** → **Software Management**.



The installed packages and the packages available on the installation media are analyzed, and dependencies between packages are checked.

After this, the following dialog for searching packages appears:



To find the software you want to install, you can select a filter listed in the drop-down list labeled **Filter** in the top left corner of the window. The following filters are available:

Patterns. Displays all software that is available on the known installation media. It is grouped in pre-defined sets of packages that logically belong together.

Package Groups. Displays all software that is available on the known installation media. It is grouped by category.

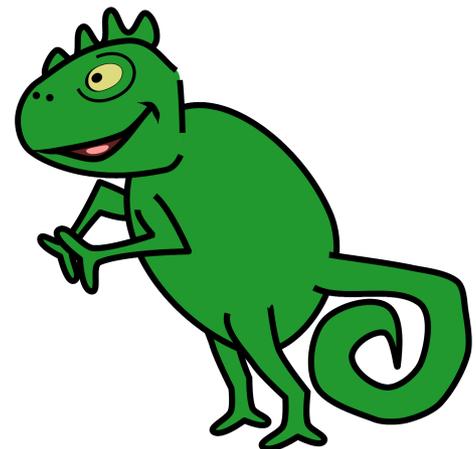
Languages. Displays all language related files (e.g., spell checker, help)

Installation Sources. Lists all registered installation sources and displays the available packages of these sources.

Search. Enter a search term and define where you want YaST to search for the software package.

Installation Summary. Displays all the packages with a marked status.

To find a package, select the **Search** filter and enter the package name, parts of the package name, or some keywords in the **Search** field; then select .



The matched packages are listed in the area on the right side. The installation state is shown by a small symbol in front of the package name. The most commonly used symbols include the following:

<input type="checkbox"/>	Do not install	This package is not installed and it will not be installed.
<input checked="" type="checkbox"/>	Install	This package will be installed. It is not installed yet.
<input checked="" type="checkbox"/>	Keep	This package is already installed. Leave it untouched.
	Update	This package is already installed. Update it or reinstall it (if the versions are the same).
	Delete	This package is already installed. Delete it.
	Taboo	This package is not installed and should not be installed under any circumstances, especially not because of unresolved dependencies that other packages might have or get. Packages set to "taboo" are treated as if they did not exist on any installation media.
	Protected	This package is installed and should not be modified, especially not because of unresolved dependencies that other packages might have or get. Use this status for third-party packages that should not be overwritten by newer versions that may come with the distribution.
<input checked="" type="checkbox"/>	Autoinstall	This package will be installed automatically because some other package needs it. Hint: You may have to use "taboo" to get rid of such a package.

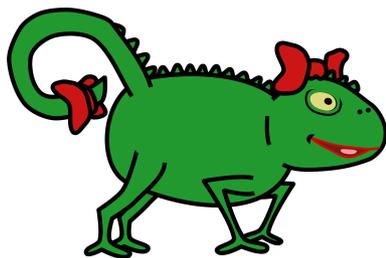
To view a list of all possible symbols, select **Help** → **Symbols**.

Select the symbol of the package you want to install several times until the install symbol appears; then select .

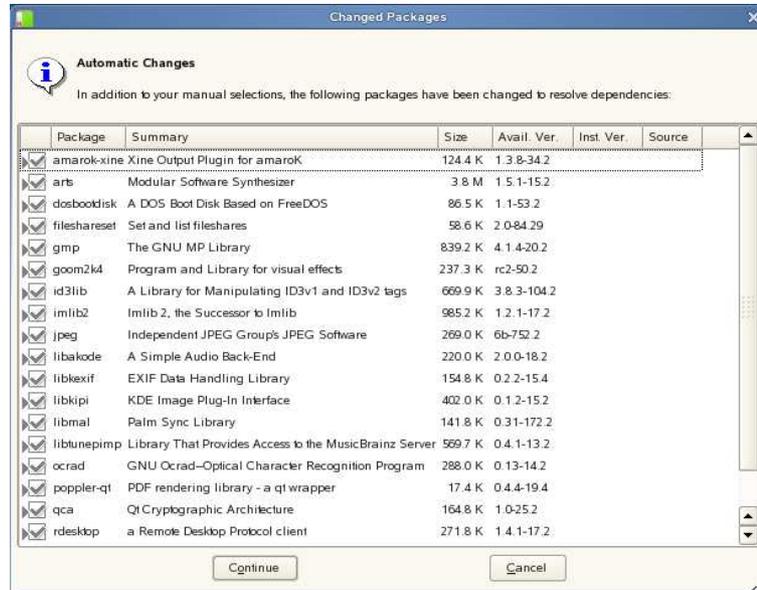


Exercise: Install Software – Part I

Install the package `findutils-locate`. This tool is introduced later in this book.

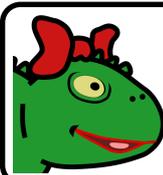


You might see a dialog indicating that dependencies between the packages cannot be solved and some other packages need to be installed, too. In most cases, you can confirm this dialog.



Exercise: Install Software – Part II

Install the package `ksh`. This is a Linux shell like the `bash` and introduced later in this book.

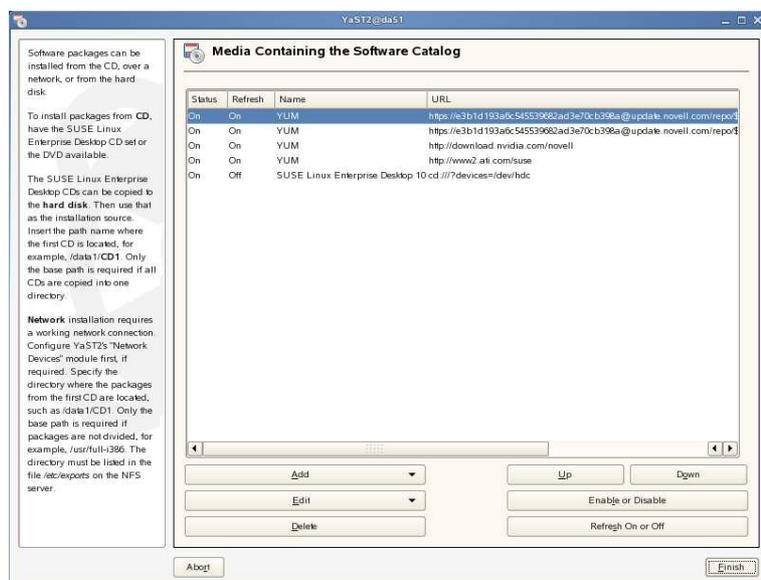


If the wrong CD or DVD is in your drive, a warning appears.

The software installation dialog lists only the packages that are on the current installation media. If you want to add more installation sources, select **Software** → **Installation Source** from the YaST Control Center.

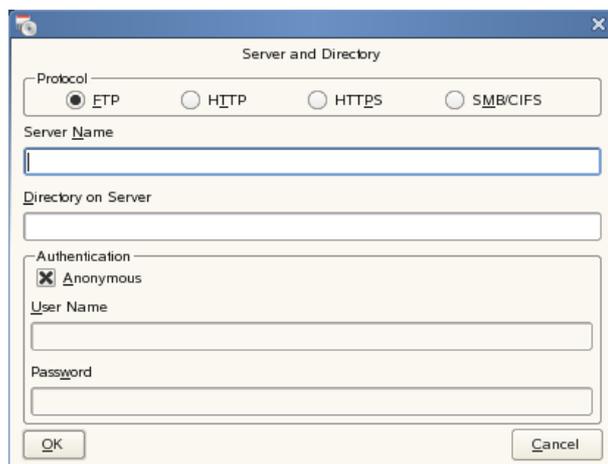


The following appears:



To add a new source, select the **Add** drop-down list and select the type of installation source.

Depending on the type of source, you might have to provide additional information (such as the IP address of an installation server). For example:



To edit the configuration of an existing installation source, select the source in the list; then select **Edit**.

If you want to disable an installation source temporarily, select the source in the list; then select **Enable** or **Disable**.

To remove an installation source permanently, select the source from the list; then select **Delete**.

YaST uses the first installation source in the list that has the software package you want to install. To change the order of a source in the list, select the source and **Up** or **Down**.

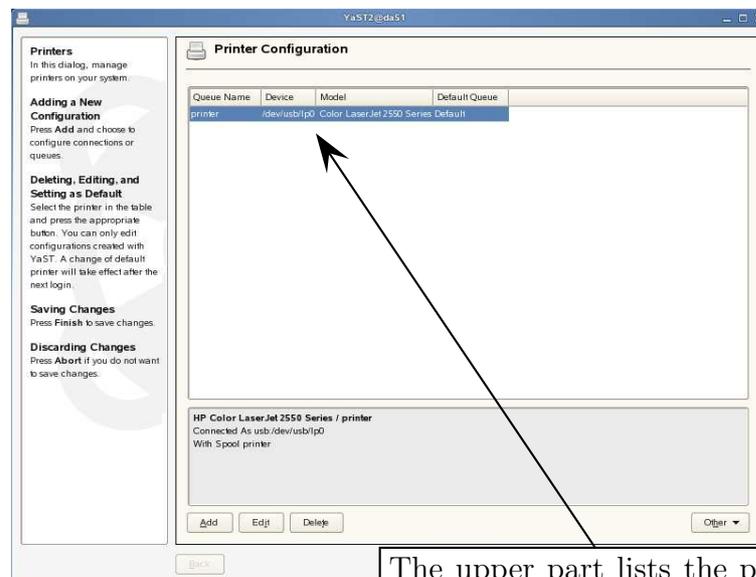
1.4 Configure a Printer

The software that manages your printers and print jobs is called CUPS (*Common Unix Printing System*). It is installed by default.

YaST provides a module for printer installation and configuration functionality. The connected printers are detected and configured automatically during the installation. You can start the module manually by selecting **Hardware** → **Printer**.



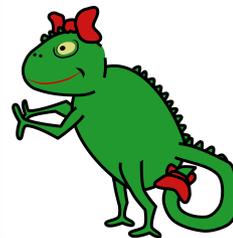
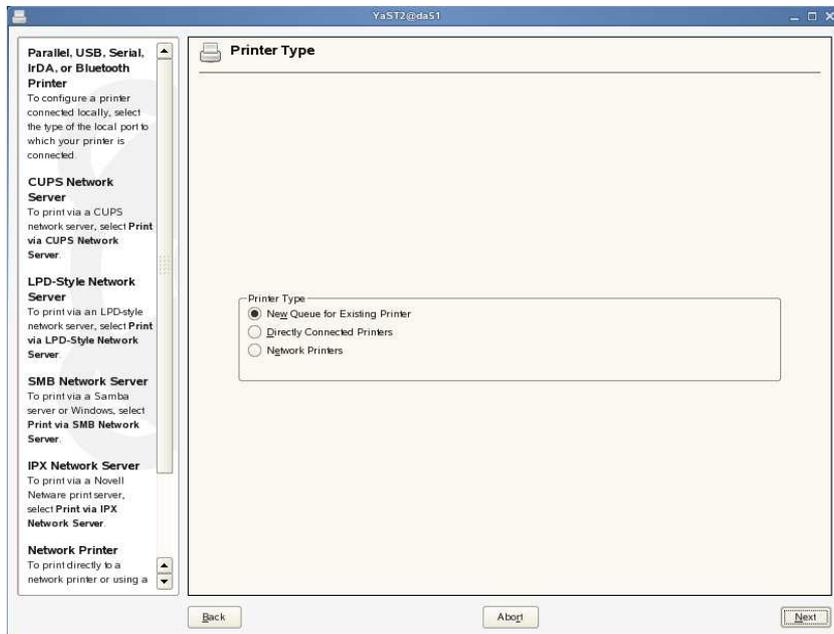
The Printer Configuration dialog appears:



The upper part lists the printers that have already been configured and all the automatically detected printers.

If there are printers listed in the upper part, the lower part of the dialog shows details of the selected printer.

To add a printer that is not listed in the upper part of the dialog (e. g., a network printer), select **Add**. The following appears:



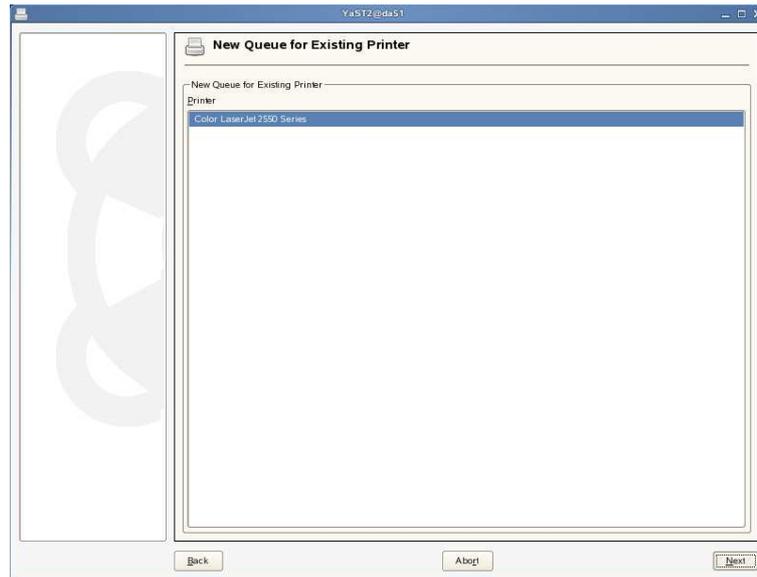
Depending on your selection here, the next dialog offers more specific options. The possible selections are:

- ☞ New Queue for Existing Printer
- ☞ Directly Connected Printers
- ☞ Network Printers

1.4.1 New Queue for Existing Printer

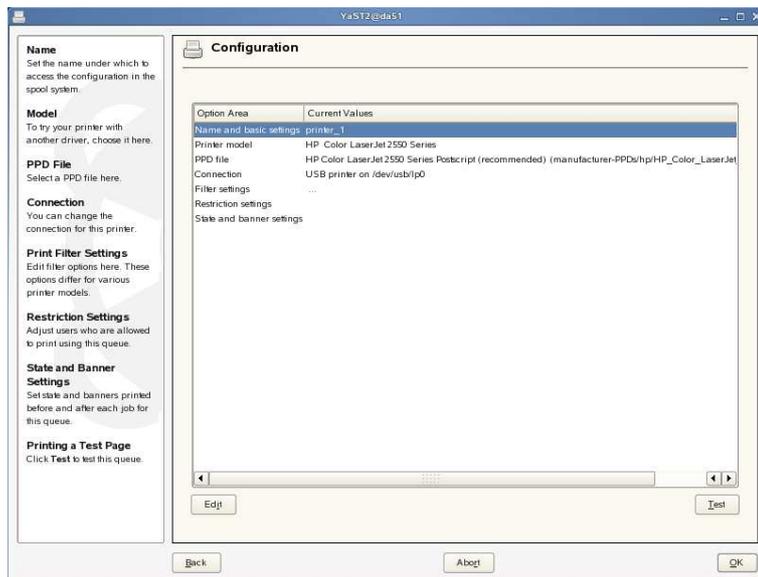
This option appears if there is already a printer configured. Adding a queue to an existing printer is useful if you want to use, for instance, a different resolution or print quality by printing to a different queue.

Selecting **New Queue for Existing Printer** and pressing **Next** opens a dialog where you select the printer you want to add a queue.



Select a printer; then select **Next**.

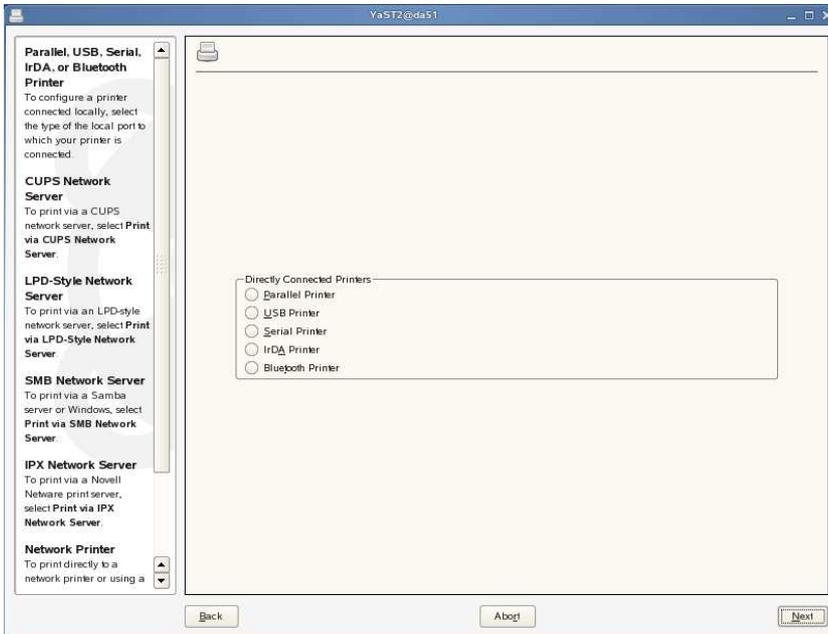
In the following Configuration dialog, select an entry; then **Edit**.



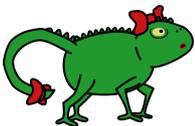
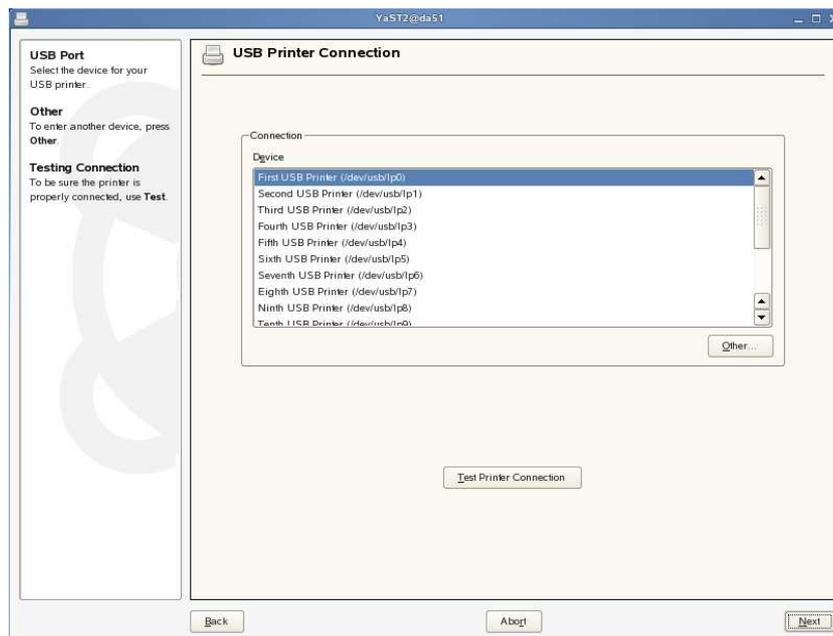
Define the queue according to your needs. You can test your printer configuration by selecting **Test**. **OK** creates the queue and returns you to the Printer Configuration dialog.

1.4.2 Directly Connected Printers

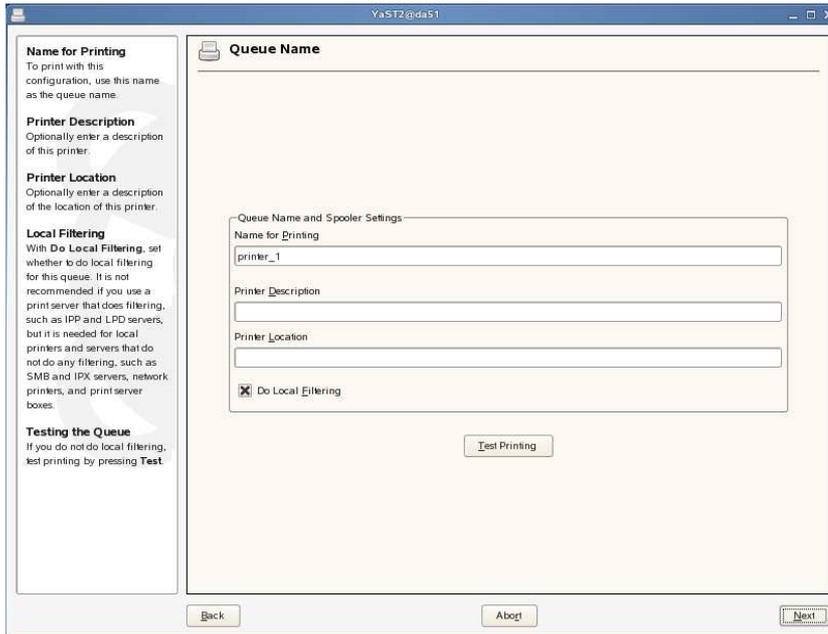
Most directly connected printers are detected automatically. If not, select the connection type:



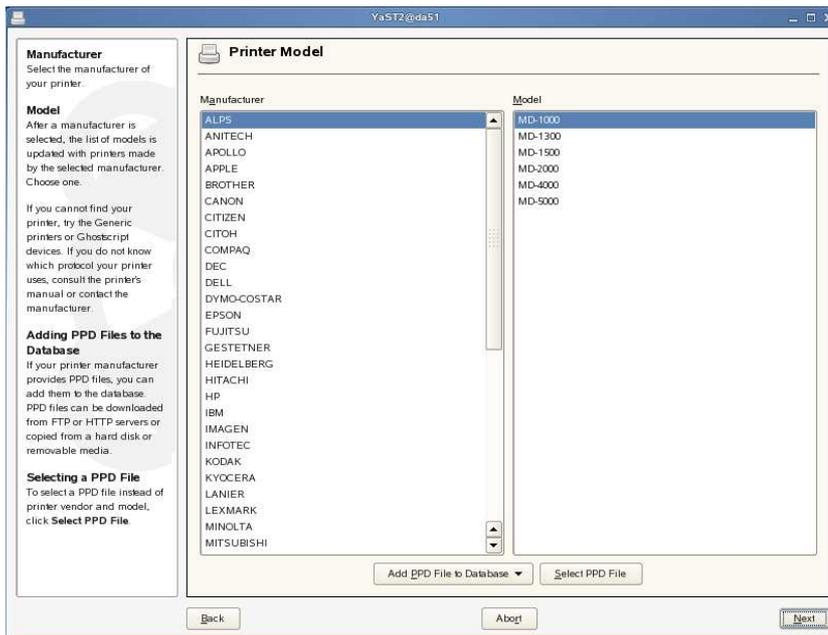
In the next dialog select the interface your printer is connected with.



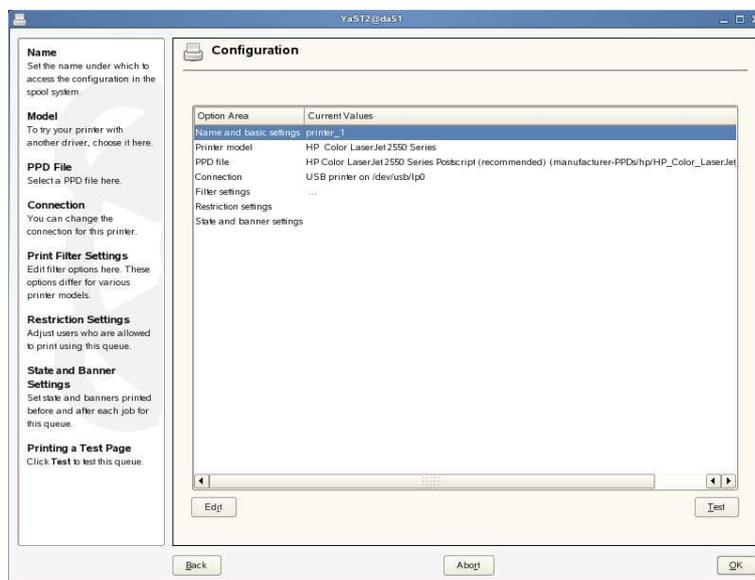
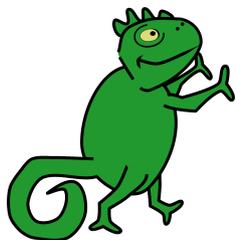
After selecting **Next**, enter a name for the new print queue. You can also enter a description of the printer and the printer's location.



In the next dialog, select the printer model. First, in the left frame, select a manufacturer; then, in the right frame, select a model.

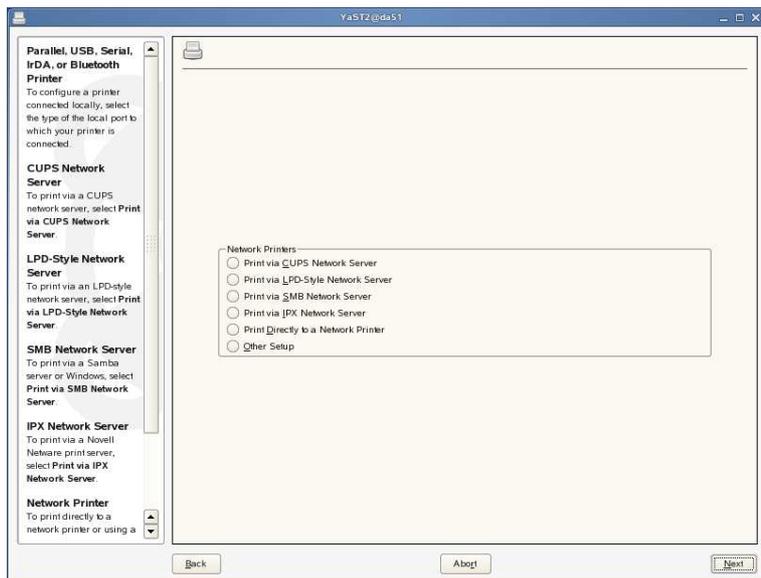


Finally, the Configuration dialog—you already know it—is shown.



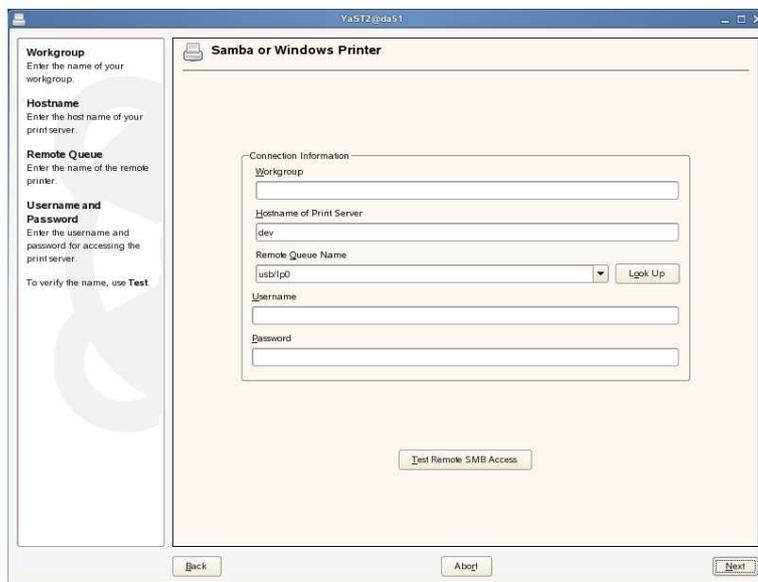
1.4.3 Network Printers

CUPS supports a couple of network protocols to communicate with a print server:



Ask the administrator of your network what protocol is used by the print server. Enter more detailed information (e. g., server name, queue name) in the following dialog. Depending on the selected protocol, this dialog varies.

For example, the dialog for SMB printers look like this:

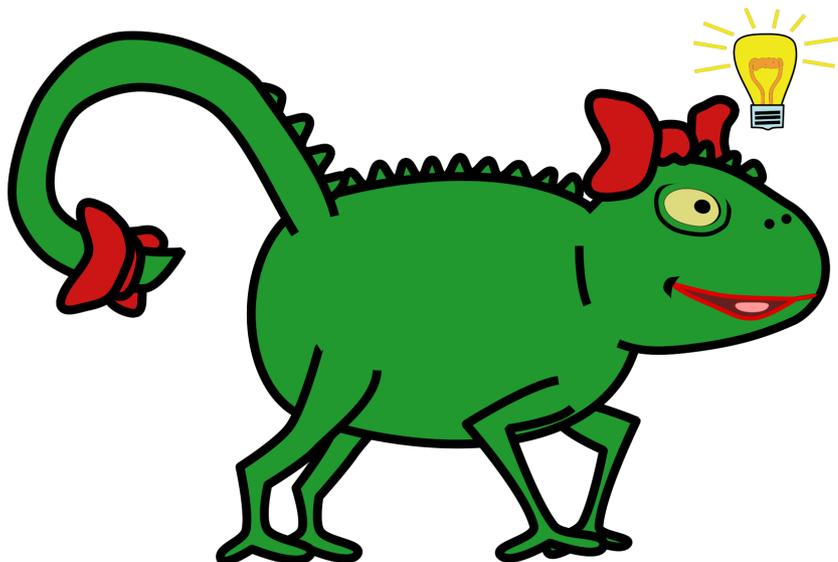


It is out of the scope of this course to describe how to connect to a network printer.



Exercise: Configure a Printer

Use YaST to view your printer configuration. Create a second queue for your current printer. This new queue should use another (lower) resolution.



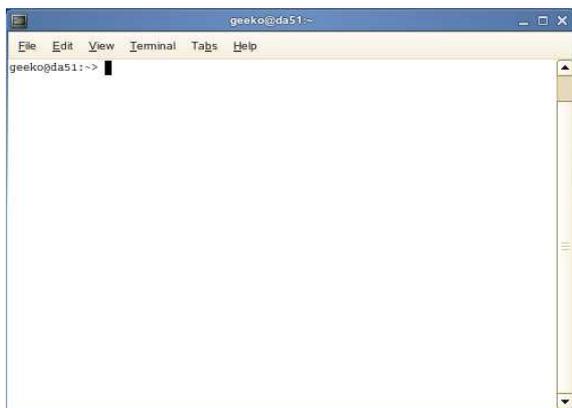
2 The Linux Shell and Tools

You cannot communicate directly with the operating system kernel. This is why you need to use a program that acts as an interface between the user and the operating system.

In the operating systems of the UNIX family, this program is called *shell*.

To have a look at the shell, select **Gnome Terminal** or **X Terminal** from the main menu.

Gnome Terminal



X Terminal



You also have up to six virtual terminals (F1 – F6) running on your computer. By pressing **Ctrl** + **Alt** + **Fx**, you can switch between individual terminals. By pressing **Ctrl** + **Alt** + **F7**, you can switch back to your graphical user interface.



The shell takes entries that users make, interprets them, converts them to system calls, and delivers system messages back to the user. This is why it is called a “command interpreter.” UNIX has a couple of shells, most of which Linux provides.

The following are examples of some popular shells:

- ☞ The Bourne shell (linked to Bourne Again shell)
- ☞ The Bourne Again shell
- ☞ The Korn shell
- ☞ The C shell (linked to Tenex C shell)
- ☞ The Tenex C shell

The various shells provide different functions and have different syntaxes.

The Bourne shell (sh) and the C shell (tsh) are not used anymore. The Bourne Again shell (bash) and the Tenex C shell (tcsh) are more user-friendly.

To show syntax differences between the Bourne Again shell (on the left) and the Tenex C shell (on the right), we use simple shell programs that display the numbers 1 to 10 on the screen.

```

1 #!/bin/bash
2
3 declare -i maximum=10
4 declare -i current=0
5
6 while ;; do
7     if [ $current -eq $maximum ]; then
8         break
9     fi
10
11     current=$((current+1))
12     echo $current
13 done

```

```

1 #!/bin/tcsh
2
3 @ maximum = 10
4 @ current = 0
5
6 while (1)
7     if ( $current == $maximum ) then
8         break
9     endif
10
11     @ current = $current + 1
12     echo $current
13 end

```

The most important tasks of the shell are:

- ❶ To interpret the command line. The shell takes the user's input and interprets it.
- ❷ To configure the user environment. When a shell is started, configuration files are read. Variables or aliases can be set.
- ❸ Shell programming. You can simplify or automate complex tasks.

The standard Linux shell is bash. Every shell can be started like a program and you can switch to a different shell at any time. For example, you can switch to the C shell by entering `tcsh` and you can switch to the Korn shell by entering `ksh`. A new bash is started with the command `bash`.

A shell is started at a text console right after a user logs in. This is called the *login shell*. The shell that is started is determined by the user database.

Unlike most other programs, the shell does not shut itself down. You need to enter the command `exit` to return to the previous shell.



Exercise: Understand the Linux Shells

The prompt is slightly different in each shell. Which prompt belongs to which shell?

A – Korn Shell

B – Bourne Again Shell

C – Tenex C Shell

①

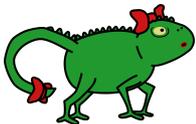
```
1 geeko@da51 :~>
```

②

```
1 /home/geeko>
```

③

```
1 geeko@da51 :/home/geeko>
```



2.1 Organize Files at the Command Line

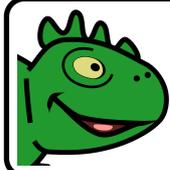
When you start a shell, a prompt indicates that the shell is ready to receive your commands.



The prompt may look similar to this one.

The prompt in this example contains the following information:

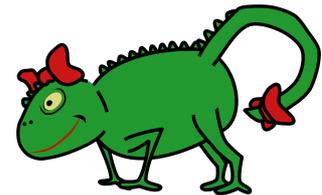
geeko – Current user’s login
 da51 – Computer’s hostname
 ~ – Current position in the file system



The path for the user’s home directory can be abbreviated to the tilde (“~”).

If you are logged in as user root, the login is not shown. The prompt looks like this:

```
1 da51:~ #
```



2.1.1 Move in the File System

You can use the command `cd` (change directory) to change between directories.

Some examples include the following:

 `cd Documents`

Change to the subdirectory Documents

```
1 geeko@da51:~> cd Documents
2 geeko@da51:~/Documents>
```

 `cd /tmp`

Change directly to the directory `/tmp` (absolute path)

```
1 geeko@da51:~> cd /tmp
2 geeko@da51:/tmp>
```

 `cd`

Change from any directory to the home directory

```
1 geeko@da51:/tmp> cd
2 geeko@da51:~>
```

☞ `cd ..`

Move one directory level higher

```
1 geeko@da51:~> cd ..
2 geeko@da51:~/home>
```

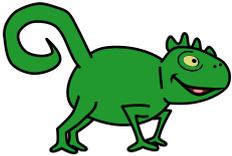
☞ `cd ../../`

Move two directory levels higher

☞ `cd -`

Move to the last valid directory

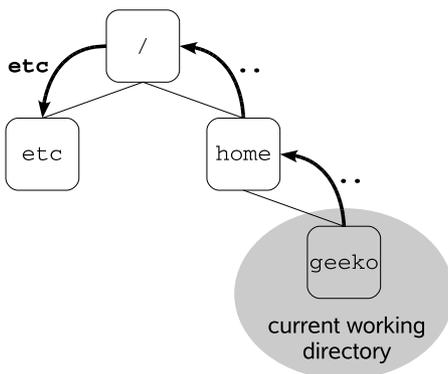
```
1 geeko@da51:~/home> cd /tmp
2 geeko@da51:/tmp> cd -
/home
4 geeko@da51:~/home>
```



The directory names in a path are separated by a slash (“/”). The path can be specified in two ways:

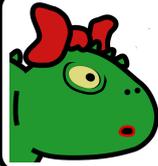
- ☞ As a relative path starting from the current directory
- ☞ As an absolute path starting from the root of the entire file system tree

The absolute path always begins with a slash (“/”), the symbol for the root directory.



In this figure, the current position in the file system is geeko's home directory. To change into the `/etc` directory, you can use the following commands:

- ☞ absolute: `cd /etc`
- ☞ relative: `cd ../../etc`

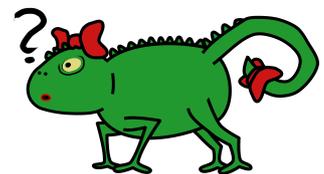


The length of the path cannot be more than 4096 characters, including the slashes.

☞ Exercise: Move in the File System

Write down to which directory you switch after entering the following commands:

Current Position	Command	Final Position
<code>/var/log/</code>	<code>cd cups</code>	
<code>/usr/X11R6/</code>	<code>cd bin</code>	
<code>/usr/X11R6/</code>	<code>cd /bin</code>	
<code>/usr/share/</code>	<code>cd ..</code>	
<code>/etc/postfix/</code>	<code>cd /</code>	
<code>~</code>	<code>cd ..</code>	
<code>/home</code>	<code>cd ..</code>	
<code>/media/cdrom/</code>	<code>cd ../../lib</code>	
<code>/sbin</code>	<code>cd</code>	



2.1.2 List Directory Content

The command `ls` (list) lists directory contents. You can specify the directory without entering an option the contents of the current directory are listed.

```
1 geeko@da51:~> ls
2 bin Desktop Documents public_html
geeko@da51:~>
```

The following are the most important options you can use with `ls`:

☞ Without any option

Displays the contents of the current directory in several columns (file and directory names only)

```
1 geeko@da51:~> ls /var
2 adm cache games lib lock log mail opt run spool tmp X11R6 yp
geeko@da51:~>
```

☞ `-a`

Also displays hidden files (such as `.bashrc`)

```
1 geeko@da51:~> ls -a /var
2 . adm games lock mail run tmp yp
3 .. cache lib log opt spool X11R6
4 geeko@da51:~>
```

☞ `-F`

After each name, a character indicates the file type. (“/” for directories, “*” for executable files, “|” for FIFO files, “@” for symbolic links).

```
1 geeko@da51:~> ls -F /var
2 adm/ games/ lock/ mail@ run/ tmp/ yp/
3 cache/ lib/ log/ opt/ spool/ X11R6/
4 geeko@da51:~>
```

☞ `-d`

Lists directory entries instead of contents.

```
1 geeko@da51:~> ls -d /var
2 /var
geeko@da51:~>
```

☞ `-l` (“long list”)

Gives a detailed list of all the files. For each file name, information on permission, modification time, and size is included.

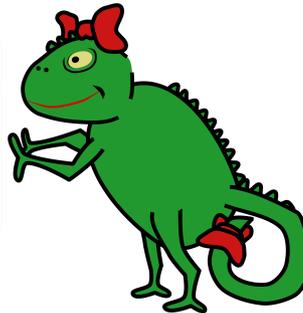
```

1 geeko@da51:~> ls -l /var
2 insgesamt 4
3
4 drwxr-xr-x  8 root  root   224 2006-09-04 10:21 adm
5
6 drwxr-xr-x 12 root  root   304 2006-09-04 10:26 cache
7
8 drwxrwxr-x  2 games games 2816 2006-09-04 10:37 games
9
10 drwxr-xr-x 39 root  root   968 2006-09-04 10:42 lib
11
12 drwxrwxr-t  4 root  uucp   128 2006-09-04 11:23 lock
13
14 drwxr-xr-x 11 root  root   936 2006-09-04 11:30 log
15
16 lrwxrwxrwx  1 root  root    10 2006-09-04 09:59 mail -> spool/mail
17
18 drwxr-xr-x  3 root  root    72 2006-09-04 10:06 opt
19
20 drwxr-xr-x 19 root  root   952 2006-09-04 11:30 run
21
22 drwxr-xr-x 11 root  root   296 2006-09-04 10:15 spool
23
24 drwxrwxrwt  3 root  root    80 2006-09-04 11:38 tmp
25
26 drwxr-xr-x  5 root  root   152 2006-09-04 10:13 X11R6
27
28 drwxr-xr-x  3 root  root   104 2006-09-04 10:15 yp
29
30 geeko@da51:~>

```

Have a look at the first ten characters of the output (“-rw-r--r--”). The first character indicates the type of file:

Character	File Type
-	Normal file
d	Directory
l	Link



The remaining nine characters show the file permission.

You can assign the following three kinds of permission to a file or directory:

Read (r) This permission allows the file to be read or the contents of a directory to be listed.

Write (w) This permission allows a file to be modified. It allows files to be created or deleted within a directory.

Execute (x) This permission allows a file to be executed. It allows you to make changes to this directory.

If permission is set, the character is shown. Otherwise a “-” appears.

The permission characters are grouped (“rwx rwx rwx”):

Character 1 to 3 Represent the permission of the file owner.

Character 4 to 6 Represent the permission of the owning group.

Character 7 to 9 Represent the permission of all other users.

The special permission SUID, SGID, and Sticky are shown by the characters “s” and “t”.

SUID Set user ID

```
1 geeko@da51:~ > ls -l /usr/bin/passwd
2 -rwsr-xr-x 1 root shadow 79765 2006-03-24 12:19
  /usr/bin/passwd
4 geeko@da51:~ >
```

SGID Set group ID

```
1 geeko@da51:~ > ls -l /usr/bin/wall
2 -rwxr-sr-x 1 root tty 10192 2006-03-22 05:24 /usr/bin/wall
geeko@da51:~ >
```

Sticky Sticky bit

```
1 geeko@da51:~ > ls -ld /tmp
2 drwxrwxrwt 15 root root 608 2006-04-06 12:45 /tmp
geeko@da51:~ >
```

The number next to the permission shows the number of hard links to that file. Hard links are explained on page 55 in this chapter.

The name of the file owner is shown next to the number of hard links. The name of the owning group is shown next to the file owner.

Then, the date and time of the changes in the file and the filename follow. For symbolic links, the file name of the target is also shown.

 -t

Files are sorted by date of alteration.

```
1 geeko@da51:~> ls -t /var
2 tmp run log lock lib games cache adm spool yp X11R6 opt mail
geeko@da51:~>
```

Combined with the option `-r`, the files are listed the other way round (the latest file is displayed at the end).

 -R

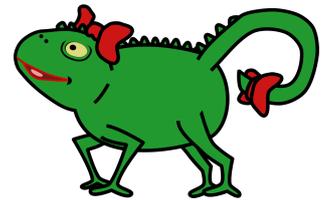
Output is recursive, including all subdirectories.

```
1 geeko@da51:~> ls -R /var
2 /var:
3 adm
4 cache
5 games
6 lib
7 lock
8 log
```

```

mail
10 opt
   run
12 spool
   tmp
14 X11R6
   yp
16
   /var/adm:
18 autoinstall
   backup
20 fillup-templates
   mount
22 perl-modules
   SuSEconfig
24
   /var/adm/fillup-templates:
26 group.aaa_base
   passwd.aaa_base
28 ...

```



 -u

The files are sorted by date of last access.

```

1 geeko@da51:~> ls -u /var
2 adm  cache  games  lib  lock  log  opt  run  spool  tmp  X11R6  yp  mail
geeko@da51:~>

```

You also can combine the `ls` options.

```

1 geeko@da51:~> ls -la /var
2 insgesamt 5
4 drwxr-xr-x 14 root  root    360 2006-09-04 10:04 .
   drwxr-xr-x 22 root  root    520 2006-09-04 11:19 ..
   drwxr-xr-x  8 root  root    224 2006-09-04 10:21 adm
6   drwxr-xr-x 12 root  root    304 2006-09-04 10:26 cache
   drwxrwxr-x  2 games games 2816 2006-09-04 10:37 games
8   drwxr-xr-x 39 root  root    968 2006-09-04 10:42 lib
   drwxrwxr-t  4 root  uucp   128 2006-09-04 11:23 lock
10  drwxr-xr-x 11 root  root    936 2006-09-04 11:30 log
   lrwxrwxrwx  1 root  root     10 2006-09-04 09:59 mail -> spool/mail
12  drwxr-xr-x  3 root  root     72 2006-09-04 10:06 opt
   drwxr-xr-x 19 root  root    952 2006-09-04 11:30 run
14  drwxr-xr-x 11 root  root    296 2006-09-04 10:15 spool
   drwxrwxrwt  3 root  root     80 2006-09-04 11:38 tmp
16  drwxr-xr-x  5 root  root    152 2006-09-04 10:13 X11R6
   drwxr-xr-x  3 root  root    104 2006-09-04 10:15 yp
18 geeko@da51:~>

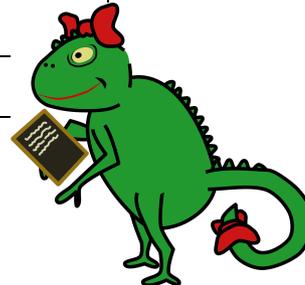
```



Exercise: List Directory Content – Part I

How many subdirectories are inside the following directories?

Directory	Number of Subdirectories
/	
/sbin	
/var/log/	
/etc/init.d/	
/usr/share/doc/	
/bin	



Exercise: List Directory Content – Part II

Write down the owner and the owning group of the following files and directory. What are the permission of the following files and directories?

File/Directory	User	Group	Permission
~/ .nautilus/			
/tmp			
/var/spool/mail/			
/etc/init.d/autofs			
/usr/bin/passwd			
/etc/passwd			
/etc/shadow			
/usr/bin/wall			

2.1.3 Move a File

You can use the command `mv` (move) to move or rename files. The syntax is

`mv source destination`

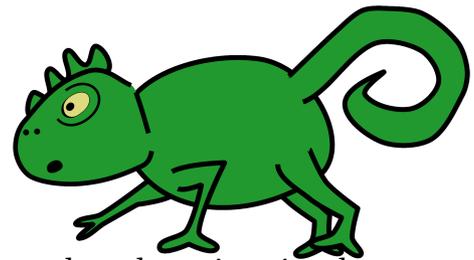
For example, to move one file (`index.html`) to another directory (`/tmp`), the command looks like this:

```
1 geeko@da51:~> mv index.html /tmp
2 geeko@da51:~>
```



Attention! The move has not been confirmed as successful. Normally, the bash does not confirm that a task has been completed successfully. Only error messages are shown.

```
1 geeko@da51:~> mv index.html /tmp
2 mv: cannot stat 'index.html': No such file or directory
geeko@da51:~>
```



Attention! If a file with the same name already exists in the target directory, it is replaced without consulting you.

You can also use the command `mv` to rename a file, for example:

```
1 geeko@da51:~> mv index.html index_en.html
2 geeko@da51:~>
```

The file `index.html` is renamed to `index_en.html`.

The difference between these two examples is: in the first example, the second parameter is a directory name and in the second example, the second parameter is a filename.

If you want to move more than one file, you can use the wildcards “?” (for any character) and “*” (for none, one, or several characters) in the first parameter.

```
1 geeko@da51:~> mv *.html /tmp
2 geeko@da51:~>
```

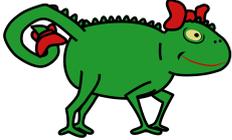
All files with the suffix “`.html`” in the current directory are moved into the `/tmp` directory.

You can use `mv` with some options. The most two important are:

 `-i`

Consults you before moving or renaming a file. This prevents existing files with the same name from being replaced.

```
1 geeko@da51:~> mv -i index.html /tmp
2 mv: overwrite '/tmp/index.html'? n
3 geeko@da51:~> mv index.html /tmp
4 geeko@da51:~>
```



 `-u`

Only moves files that are newer than the target files of the same name.

```
1 geeko@da51:~> ls -l index.html
2 -rw-r--r-- 1 geeko users 26432 2005-12-08 15:48 index.html
3 geeko@da51:~> ls -l /tmp/index.html
4 -rw-r--r-- 1 geeko users 864 2006-09-04 11:26 /tmp/index.html
5 geeko@da51:~> mv -u index.html /tmp/
6 geeko@da51:~> ls -l index.html
7 -rw-r--r-- 1 geeko users 26432 2005-12-08 15:48 index.html
8 geeko@da51:~> ls -l /tmp/index.html
9 -rw-r--r-- 1 geeko users 864 2006-09-04 11:26 /tmp/index.html
10 geeko@da51:~> mv index.html /tmp/
11 geeko@da51:~> ls -l /tmp/index.html
12 -rw-r--r-- 1 geeko users 26432 2005-12-08 15:48 /tmp/index.html
13 geeko@da51:~> ls -l index.html
14 /bin/ls: index.html: No such file or directory
15 geeko@da51:~>
```

Explanation: The file `~/index.html` was modified in December 2005 and is 26432 bytes (lines 1–2). The file `/tmp/index.html` was modified in September 2006 and is 864 bytes (lines 3–4). After moving the file `~/index.html` to the directory `/tmp` (line 5) using the option `-u`, the file `~/index.html` still exists (lines 6–7) and the destination file still has the same data (864 bytes, lines 8–9). That means, the file was not moved. Without `-u` (line 10), the command `mv` overwrites the destination file although it is newer than `~/index.html` (lines 11–12). The file `~/index.html` does not exist anymore (lines 13–14).

2.1.4 Copy a File

You can copy files and directories with the command `cp` (copy). The syntax is

`cp source destination`

When using the command `cp`, please do not forget:

- ☞ The command `cp` replaces existing files without consulting you.
- ☞ If you only want to copy the contents of a directory (without the directory itself), the target directory must already exist, for example, to back up a directory using a different name.

Without any option, `cp` copies files. If you want to copy a directory, you have to use the option `-r`.

```
1 geeko@da51:~> cp proposals/ /tmp
2 cp: omitting directory 'proposals/'
3 geeko@da51:~> cp -r proposals/ /tmp
4 geeko@da51:~>
```

In this example, the directory `~/proposals/` with all its subdirectories is copied to the directory `/tmp`. The result is a directory `/tmp/proposals/`.

To only copy the contents of the directory `proposals/` (including hidden files and subdirectories) to the directory `/tmp`, enter the following:

```
1 geeko@da51:~ > cp -r proposals/. /tmp
2 geeko@da51:~ >
```

If you do not want to copy the hidden files, enter the following:

```
1 geeko@da51:~ > cp -r proposals/* /tmp
2 geeko@da51:~ >
```

Besides `-r`, the most important options of `cp` are:

☞ `-a, --archive`

Copies a directory and subdirectories (like `-r`); symbolic links, file permission, owners, and time stamps are not changed.

```
1 geeko@da51:~> ls -a index.html
2 -rw-r--r-- 1 geeko users 864 2006-09-04 11:26 index.html
3 geeko@da51:~> cp -a index.html /tmp/
4 geeko@da51:~> ls -l /tmp/index.html
5 -rw-r--r-- 1 geeko users 864 2006-09-04 11:26 /tmp/index.html
6 geeko@da51:~> cp index.html /tmp/
7 geeko@da51:~> ls -l /tmp/index.html
8 -rw-r--r-- 1 geeko users 864 2006-09-04 13:35 /tmp/index.html
9 geeko@da51:~>
```

Explanation: In this example, the last change of the file `~/index.html` was at 11:26 (lines 1–2). After copying the file with the `-a` option (line 3), the modification time of the copy is 11:26, too (lines 4–5). When not using `-a` (line 6), the modification time is set to the current time (here 13:35, lines 7–8).

`-i, --interactive`

Asks before overwriting.

```
1 geeko@da51:~> cp -i index.html /tmp
2 cp: overwrite '/tmp/index.html'? n
3 geeko@da51:~> cp index.html /tmp
4 geeko@da51:~>
```

`-s, --symbolic-link`

Creates symbolic links instead of copying.

```
1 geeko@da51:~> cp -s /tmp/index.html .
2 geeko@da51:~> ls -l index.html
3 lrwxrwxrwx 1 geeko users 15 2006-09-04 13:37 index.html -> /tmp/index.html
4 geeko@da51:~>
```

`-u, --update`

Copies a file only when the source file is newer than the destination file or when the destination file is missing.

```
1 geeko@da51:~> ls -l index.html
2 -rw-r--r-- 1 geeko users 26432 2005-12-08 15:48 index.html
3 geeko@da51:~> ls -l /tmp/index.html
4 -rw-r--r-- 1 geeko users 864 2006-09-04 13:40 /tmp/index.html
5 geeko@da51:~> cp -u index.html /tmp/
6 geeko@da51:~> ls -l /tmp/index.html
7 -rw-r--r-- 1 geeko users 864 2006-09-04 13:40 /tmp/index.html
8 geeko@da51:~> cp index.html /tmp/
9 geeko@da51:~> ls -l /tmp/index.html
10 -rw-r--r-- 1 geeko users 26432 2006-09-04 13:42 /tmp/index.html
11 geeko@da51:~>
```

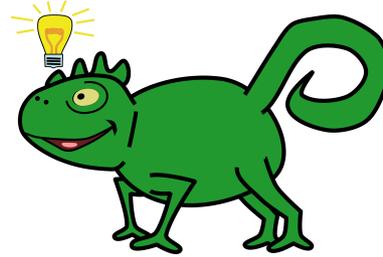
Explanation: The file `~/index.html` was modified in December 2005 and is 26432 bytes (lines 1–2). The file `/tmp/index.html` was modified in September 2006 and is 864 bytes (lines 3–4). After copying the file `~/index.html` to the directory `/tmp` (line 5) using the option `-u`, the destination file is still the same (864 bytes, lines 6–7). Without `-u` (line 8), the command `cp` overwrites the destination file, although it is newer than `~/index.html` (lines 9–10).

Exercise: Move and Copy Files

Do the following:

- ❶ Copy the file `/etc/init.d/syslog` into your home directory. How does the file permission change?
- ❷ Rename the file `~/syslog` into `testfile`.

- ③ Copy the file `grep-demo` from the `examples` directory of the book CD into your home directory.



2.1.5 Create a Directory

You can use the command `mkdir` (make directory) to create new directories.

```
1 geeko@da51:~> mkdir Proposal
2 geeko@da51:~>
```

With the option `-p`, you can create a complete path, as in the following:

```
1 geeko@da51:~> ls
2 bin Desktop Documents public_html
3 geeko@da51:~> mkdir Invoice/Customer
4 mkdir: cannot create directory 'Invoice/Customer': No such file or directory
5 geeko@da51:~> mkdir -p Invoice/Customer
6 geeko@da51:~> ls
7 bin Desktop Documents Invoice public_html
8 geeko@da51:~> ls Invoice/
9 Customer
10 geeko@da51:~>
```

Explanation: It is possible to create a new directory `Invoice` including a directory `Customer` (lines 3-4). Use the option `-p`: `mkdir -p` (line 5).

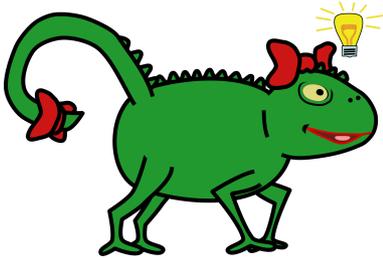


Exercise: Create a Directory

Do the following:

- ① Create a new directory `Marketing` in your home directory.
- ② Move the file `testfile` into the `Marketing` directory.
- ③ Create a new directory `School` that includes a subdirectory `Mathematics` in the `/tmp` directory.
- ④ Copy the `Marketing` directory from your home directory into the directory `/tmp/School/Mathematics`.

- 5 Create a new directory `German` in the directory `/tmp/School/`.



2.1.6 Delete Directories and Files

You can use the `rmdir` (remove directory) command to delete the directory or directories (e. g., `rmdir Proposal`).

```
1 geeko@da51:~> rmdir Proposal
2 geeko@da51:~>
```

The directory or directories must be empty before you can delete them.

```
1 geeko@da51:~> rmdir Invoice
2 rmdir: Invoice: Directory not empty
3 geeko@da51:~> ls Invoice/
4 Customer
5 geeko@da51:~>
```

You can use the command `rm` (remove) to delete files, as in the following:

```
1 geeko@da51:~> rm file
2 geeko@da51:~>
```

If you use `rm` with wildcards, be careful that only the files are deleted you want to delete.

```
1 geeko@da51:~> ls
2 bin      Documents  file2    file4    file6    Invoice
3 Desktop  file1     file3    file5    file7    public_html
4 geeko@da51:~> rm file*
5 geeko@da51:~>
```

This example deletes all files, that begin with `file`, in the current directory without asking for confirmation.



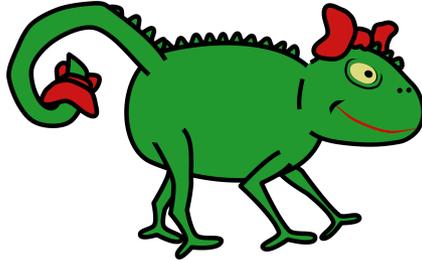
Attention! Files deleted with the command `rm` cannot be restored.

If the user does not have the permission to delete a file, this file is ignored and an error message is printed.

```

1 geeko@da51:~> rm /etc/vimrc
2 rm: remove write-protected regular file '/etc/vimrc'? y
3 rm: cannot remove '/etc/vimrc': Permission denied
4 geeko@da51:~>

```



The following are some important options you can use with `rm`:

☞ `-i`

Consults you before deleting.

```

1 geeko@da51:~> rm -i file
2 rm: remove regular empty file 'file'? y
3
4 geeko@da51:~>

```

☞ `-r` (recursively)

Allows full directories to be deleted.

```

1 geeko@da51:~> rmdir Invoice
2 rmdir: Invoice: Directory not empty
3
4 geeko@da51:~> rm -r Invoice
5
6 geeko@da51:~>

```

☞ `-f` (force)

By default, `rm` consults you if the file that should be deleted is read-only. Using this option, the files are deleted without consulting you.

```

1 geeko@da51:~> ls -l file
2 -r--r--r-- 1 geeko users 0 2006-09-04 11:44 file
3
4 geeko@da51:~> rm file
5 rm: remove write-protected regular empty file 'file'? n
6
7 geeko@da51:~> rm -f file
8
9 geeko@da51:~>

```

Explanation: The file `file` is read-only (lines 1–2). When you use `rm` without any option the command prompts for confirmation (lines 3–4). (This can be annoying if you delete a couple of files using wildcards.) Using `-f`, the files are deleted without consulting you (line 5).



Exercise: Delete Directories and Files

Do the following:

- ❶ Delete the file `/tmp/School/Mathematics/Marketing/testfile`.
- ❷ Delete the empty directory `/tmp/School/German/`.
- ❸ Delete the directory `School` with all its content from the `/tmp` directory.

2.1.7 Link Files

File system formats in Linux keep data and administration information separately. The data are organized according to the respective file system format.

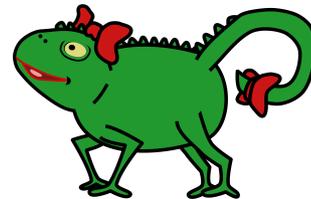
Each file is described by an inode (index node or information node). To see the inode number, enter `ls -li`.

```

1 geeko@da51:~> ls -li
2 insgesamt 0
3 84800 drwxr-xr-x 2 geeko users 48 2006-09-22 11:20 bin
4 84811 drwxr-xr-x 2 geeko users 80 2006-09-22 11:20 Desktop
5 84801 drwx----- 2 geeko users 80 2006-09-22 11:20 Documents
6 84808 drwxr-xr-x 2 geeko users 80 2006-09-22 11:20 public_html
geeko@da51:~>

```

Each of these inodes is 128 bytes and contains all the information on this file except the file name. This includes details of the owner, access permission, size, various time details (time of modification, of access, of modification of the inode), and the links to the data blocks of this file.



The command `ln` creates a link. A link is a reference to a file. With a link, you can access a file from anywhere in the file system using different names for it. This means that the file itself exists only once on the system, but it can be accessed by using different names.

In Linux, two kinds of links are available:

- ☞ Hard links
- ☞ Symbolic links

Use the command `ln` to create a hard link. This link points to the inode of a file. Thereafter, the file can be accessed by using both names—by the name of the file itself and by the name of the link. That means you are no longer able to tell the difference between the original file and the link.

The following is an example of using the command `ln`:

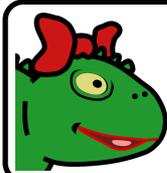
```

1 geeko@da51:~/Proposal > ls -li
2 total 4
3 88658 -rw-r--r-- 1 geeko users 82 2004-04-06 14:21 old
4 geeko@da51:~/Proposal > ln old new
5 geeko@da51:~/Proposal > ls -li
6 total 8
7 88658 -rw-r--r-- 2 geeko users 82 2004-04-06 14:21 old
8 88658 -rw-r--r-- 2 geeko users 82 2004-04-06 14:21 new
9 geeko@da51:~/Proposal >

```

Explanation: `ls -li` displays the file inode (lines 1–3). In this example, the inode of the file `old` is 88658. With the `ln` command, you can create a hard link (line 4). The new link has the same inode as the original file (lines 5–8).

Hard links can only be used when both the file and the link are in the same file system (on the same partition), because inode numbers are only unique within the same file system.



You can see the number of hard links pointing to an inode in the `ls` output after the permission.

A symbolic link is assigned its own inode—the link refers to a file, so a distinction can always be made between the link and the actual file. You can create a symbolic link with the command `ln` and the option `-s`.

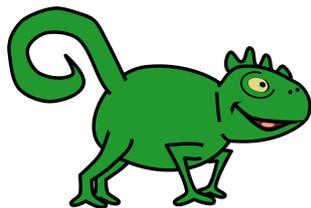
The following is an example of creating a symbolic link:

```

1 geeko@da51:~/Proposal > ls -li
2 total 4
3 88658 -rw-r--r-- 1 geeko users 82 2004-04-06 14:21 old
4 geeko@da51:~/Proposal > ln -s old new
5 geeko@da51:~/Proposal > ls -li
6 total 4
7 88658 -rw-r--r-- 1 geeko users 82 2004-04-06 14:21 old
8 88657 lrwxrwxrwx 1 geeko users 3 2004-04-06 14:27 new -> old
9 geeko@da51:~/Proposal >

```

Explanation: In this example, the inode of the file `old` is 88658 (lines 1–3). With the `ln -s` command you can create a soft link (line 4). In this example, the link has the inode 88657 (lines 5–8).



The disadvantage of a symbolic link is that it can point to a non-existing object if the object no longer exist.

If you delete the file `old` in the example above, `new` will point to a non-existing file. The common `ls` does not display that the link is “broken”:

```
1 geeko@da51:~/Proposal > rm old
2 geeko@da51:~/Proposal > ls -li
total 0
4 88657 lrwxrwxrwx 1 geeko users 3 2004-04-06 14:27 new -> old
geeko@da51:~/Proposal >
```

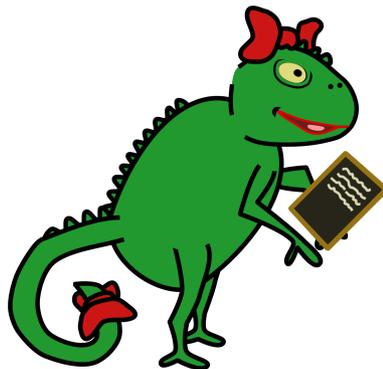
With symbolic links, you can also create links to directories.



Exercise: Link Files

Do the following:

- ❶ How many links does the following directory include?
`/boot/boot/boot/boot/boot/boot/boot/boot/boot/`
- ❷ Create a symbolic link named `symlink` in your home directory that points to the file `Marketing/testfile`.
- ❸ Create a hard link named `hardlink` in your home directory that points to the file `Marketing/testfile`.
- ❹ Remove the file `Marketing/testfile`.
- ❺ Test links with the command `cat filename`. The command `cat` displays the content of a text file.



2.1.8 Find Files

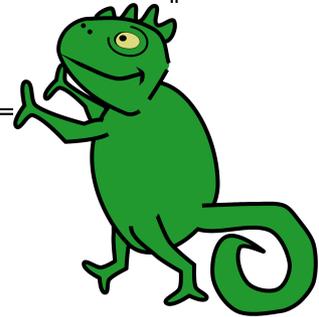
The shell provides commands to find files. The most important commands are:

☞ find

☞ locate

Example (Suppose the following files exist):

☞ File	☞ File1a	☞ File2a
☞ file	☞ File1b	☞ FILE
☞ File1	☞ File2	☞ MyFile



Find Files with file

To search for files on the command line, use the command `find`. The syntax is:

```
find path criterion action
```

For example:

```
1 geeko@da51:~> find . -name "File?"
2 ./File1
3 ./File2
4 geeko@da51:~>
```

The command has a lot of options. Some of them are explained in the following:

path The part of the file system to search (the specified directory and all its subdirectories). If nothing is specified, the file system below the current directory is used.

criterion The properties the file should have. They refer to the following:

☞ `-ctime [+/-]days`

Searches for files whose last change took place no later than (or no earlier than) a specified number of days ago.

☞ **-group *name***

Searches for files owned by the group *name*. Instead of a name, you can use the numeric GID.

☞ **-name *pattern***

Searches for files whose names contain the given pattern. If the pattern contains wild cards, it must be enclosed by quotation marks. Otherwise, it will be interpreted by the shell and not by the `find` command.

☞ **-size [+/-]*size***

Matches files that are bigger or smaller than a certain size. As an argument, the size (in blocks of 512 bytes) is given. The suffix “c” switches to byte (e.g., `-size -260c`) and “k” to blocks of 1024 bytes (e.g., `-size +600k`). A preceding “+” stands for all larger files and a “-” for all smaller files.

☞ **-type *file_type***

Searches for a file type. A file type can be one of the following: “d” for a directory, “f” for a file, or “l” for a symbolic link.

☞ **-user *name***

Searches for files owned by user *name*. Instead of a name, you can use the numeric UID.



action Options that influence the following conditions or control the whole search, such as the following:

☞ **-print (default)**

☞ **-exec *command***

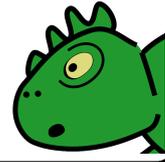
With the option `-exec`, you can call up another command. This option is frequently used to link `find` and `grep` (see section 2.3.2 on page 90) as in the following:

```
1 geeko@da51:~ > find ~ -name "letter*" -type f -exec grep appointment {} \;
2 appointment for next meeting: 23.08.
  /home/geeko/letters/letter_Smith
4 geeko@da51:~ >
```

Explanation: In this example, the command `find` searches for files whose names start with `letter`, and then passes the names of the files found with `-exec` to the following command (in this case `grep appointment {}`).

The two brackets `{}` are placeholders for the file names that are found and passed to the command `grep`. The semicolon closes the `-exec` instruction. As the semicolon is a special character, it is masked by a backslash.

When you use `grep` alone, it searches for a specific expression in a given file. If you use `grep` with `find`, you can search for a specific expression in a file whose location is not known.



Exercise: Find Files

Please answer the following questions:

- How many files in the `/bin` directory are larger than 1024kByte?

- How many files in the directory `/etc` belong to the group `lp`?

- Where in the filesystem is the command `grep` located?

Locate Files with `locate`



Attention! The command `locate` is not available in a standard installation. Before you can use this command, you have to install the package `findutils-locate`.

The command `locate` is equivalent to `find -name`.

As the command `find` searches through the selected part of the file system, this process might be quite slow. `locate` searches through a database

that was created especially for this purpose (`/var/lib/locatedb`). That quicken the search.

On SUSE Linux Enterprise Desktop, the database is automatically created and updated every day. If you change something after this update, you have to update manually using the command `updatedb`.



Attention! After installing `findutils-locate`, create the database by using the command `updatedb` (as user `root`). Otherwise, you will get an error message:

```
1 geeko@da51:~> locate File?
2 locate: /var/lib/locatedb: No such file or directory
```

You can execute a command at the command line as user `root` by using the `su` command. If you want to execute just one command, use the option `-c command`. To authenticate, enter the root password.

```
1 geeko@da51:~> su -c updatedb
2 Passwort:
geeko@da51:~>
```

The following example shows the output of `locate`:

```
1 geeko@da51:~> locate File?
2 /home/geeko/File1
3 /home/geeko/File1a
4 /home/geeko/File1b
5 /home/geeko/File2
6 geeko@da51:~>
```

`locate` displays all files whose names contain the search string.



To learn more about `locate`, enter `man locate`.



Exercise: Locate Files

Make sure that the package `findutils-locate` is installed and update the database that is used by `locate`.

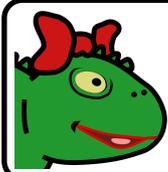
Use `locate` to find the following files:

- ☞ `updatedb`: _____
- ☞ `XF86Config`: _____
- ☞ `smb.conf`: _____
- ☞ `hardlink`: _____

- ☞ boot.log: _____
- ☞ installkernel: _____

Find Executable Files with `which`

The command `which` searches all paths listed in the variable `PATH` for the specified command and returns the full path of the command.



A variable is a labeled memory space, where you can store information.

In the variable `PATH`, the most important directories where the shell looks for executable files are listed.

To see the content of a variable, use the command `echo` and add a “\$” in front of the variable’s name.



To see the content of the variable `PATH`, enter `echo $PATH`.

```
1 geeko@da51:~> echo $PATH
2 /home/geeko/bin:/usr/local/bin:/usr/bin:/usr/X11R6/bin:/bin:/usr/games:/opt/gnom
e/bin:/opt/kde3/bin:/usr/lib/jvm/jre/bin:/usr/lib/mit/bin:/usr/lib/mit/sbin
4 geeko@da51:~>
```

If several versions of a command are in different directories and you want to know which version is executed when you enter it without specifying a path, use the command `which`. The following is an example of using the command `which`:

```
1 geeko@da51:~ > which find
2 /usr/bin/find
3 geeko@da51:~ > which cp
4 /bin/cp
5 geeko@da51:~> which locate
6 /usr/bin/locate
7 geeko@da51:~ >
```



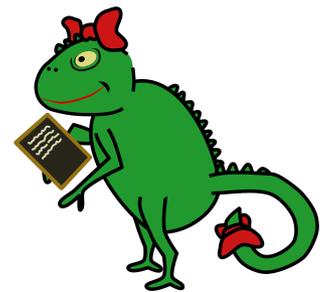
For more information on `which`, enter `man which`.



Exercise: Find Executable Files

Use which to locate the following commands:

- ☞ grep: _____
- ☞ which: _____
- ☞ mkdir: _____
- ☞ nautilus: _____
- ☞ OOo-writer: _____
- ☞ ktab: _____
- ☞ fvwm2: _____



Specify the Type of a Command

With the command `type`, you can find out what kind of command is executed when the command is entered:

- ☞ a shell built-in command (an essential command that is hardcoded in the shell)
- ☞ an external command (called by the shell)
- ☞ an alias
- ☞ a function

The option `-a` displays all instances of a command with this name in the file system. The following is an example of using the command `type`:

```

1 geeko@da51:~ > type type
2 type is a shell built in
3 geeko@da51:~ > type grep
4 grep is /usr/bin/grep
5 geeko@da51:~ > type -a grep
6 grep is /usr/bin/grep
7 grep is /bin/grep
8 geeko@da51:~ >

```



Exercise: Specify the Type of a Command

What kind of files are the following:

echo

- A shell built in
- An external command: _____
- An alias to: _____
- A function

ls

- A shell built in
- An external command: _____
- An alias to: _____
- A function

which

- A shell built in
- An external command: _____
- An alias to: _____
- A function

_which

- A shell built in
- An external command: _____
- An alias to: _____
- A function

test

- A shell built in
- An external command: _____
- An alias to: _____
- A function

2.1.9 Archive Files

tar (tape archiver) is the most commonly used tool for data backup. It archives files in a special format, either directly on a corresponding medium (such as magnetic tape or formatted floppy disk) or to a so-called archive file. By default, the data is *not* compressed.



Names of archive files have the suffix `.tar`. If archive files are compressed (usually with the command `gzip`; see page 66), the suffix is either `.tar.gz` or `.tgz`.

The syntax is: `tar option files`

The command first expects an option (which is why it can also be used without a minus sign) and the name of the directory to be backed up. All directories and files beneath this are also saved. One very important option is `-f filename` which is used to specify the name of the archive to be written (or the device file).

Directories are typically backed up with:

```
1 geeko@da51:~ > tar -cvf /tmp/backup.tar /home/geeko
```

Explanation: This command backs up the complete contents of the user's home directory and stores it in the file `/tmp/backup.tar`.

The option `-c` (*create*) creates the archive. The option `-v` (*verbose*) provides a detailed list containing the name of each file being stored. After the option `-f` (*file*), specify the name of the archive to be created.

When an archive is created, absolute paths turns to relative and that is the first `/` in the path is removed.

```
1 tar: Removing leading '/' from member names
```

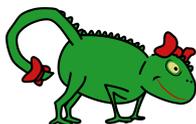
To unpack files of an archive, enter:

```
1 geeko@da51:~ > tar -xvf /tmp/backup.tar
```

Explanation: This command unpacks all files stored in the archive into the current directory. The directory structure is created according to the relative path specifications in the `tar` archive.

If you only want to unpack one file, specify the filename:

```
1 geeko@da51:~ > tar -xvf /tmp/backup.tar home/geeko/.bashrc
```



More useful options are:

-C

Unpacks to a specific directory using the option `-C` followed by the directory name.

```
1 geeko@da51:~ > tar -xvf /tmp/backup.tar -C /data
```

-d

Compares files in the archive with those in the file system.

```
1 geeko@da51:~ > tar -dvf /tmp/backup.tar -C /home/geeko
```

-j

Compresses or decompresses the tar archive using `bzip2` (see page 67).

```
1 geeko@da51:~ > tar -cjvf /tmp/backup.tar.bz2 /home/geeko
```

-r

Adds files to an archive.

```
1 geeko@da51:~ > tar -vf /tmp/backup.tar -r new_file
```

-t

Displays the contents of an archive:

```
1 geeko@da51:~ > tar -tvf /tmp/backup.tar
```

-u

Only adds files to an archive that are newer than the file version in the archive (*update*).

```
1 geeko@da51:~ > tar -uvf /tmp/backup.tar -C /home/geeko
```

--exclude

The given files are not included when a directory is archived.

```
1 geeko@da51:~ > tar -cvf /tmp/backup.tar /home/geeko --exclude="*.iso"
```

Explanation: The home directory of user `geeko` is archived without all ISO files in his home directory.

-z

Compresses or decompresses the tar archive using `gzip` (see page 66).

```
1 geeko@da51:~ > tar -czvf /tmp/backup.tar.gz /home/geeko
```



Exercise: Create a tar Archive

Create the `tar` archive `/tmp/bin.tar` containing your entire `/bin` directory. Then go to the directory `/tmp/` and unpack only one file, e.g., the file `chown`. Compare owner and time stamp of the unpacked and the original file.

To compress `tar` archives, use the options `-j` or `-z`. In Linux, you can also use some tools to compress and decompress data.

You can use the command `gzip` to compress and decompress data. The compressed data has the suffix `.gz`. This command is only useful for compressing individual files.

The syntax of `gzip` is:

```
gzip [options] file
```

For example:

```

1 geeko@da51:~> ls
2 backup.tar bin Desktop Documents public_html
3 geeko@da51:~> gzip backup.tar
4 geeko@da51:~> ls
5 backup.tar.gz bin Desktop Documents public_html
6 geeko@da51:~>

```

To save several files or entire directories in a compressed file, use the command `tar`.



The following are some useful options of `gzip`:

☞ `-c`

Compresses the file without modifying the original file. The result is written to the standard output (usually the screen).

☞ `-d`

Decompresses the specified file instead of compressing it.

☞ `-1` to `-9`

Controls the compression speed. `-1` compresses faster but produces larger files. `-9` requires more computing time but produces smaller files. The default setting is `-6`.

 `-r`

Also compresses and decompresses files in all subdirectories.

Instead of `gzip -d`, you can also use the command `gunzip`. This command decompresses a file compressed by `gzip` and removes the suffix `.gz`.

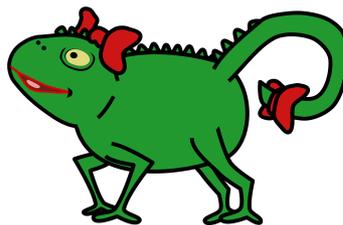
```
1 geeko@da51:~> ls
2 backup.tar.gz  bin  Desktop  Documents  public_html
3
4 geeko@da51:~> gunzip backup.tar.gz
5
6 geeko@da51:~> ls
7 backup.tar  bin  Desktop  Documents  public_html
8
9 geeko@da51:~>
```



Exercise: Compress a File – Part I

Compress the archive `/tmp/bin.tar` using `gzip` with three different compression speeds. Write down the file size of the compressed file here:

-1 _____
 -6 _____
 -9 _____



`bzip2` also compresses files. It creates files that are usually about twenty to thirty percent smaller than the files compressed by `gzip`. But `bzip2` needs more computing time to compress. The compressed data has the suffix `.bz2`.

The syntax of `bzip2` is:

`bzip2 [options] file`

For example:

```
1 geeko@da51:~> ls
2 backup.tar  bin  Desktop  Documents  public_html
3
4 geeko@da51:~> bzip2 backup.tar
5
6 geeko@da51:~> ls
7 backup.tar.bz2  bin  Desktop  Documents  public_html
8
9 geeko@da51:~>
```

Some important options are:

☞ `-c`

Compresses the file without modifying the original file. The result is written to the standard output (usually the screen).

☞ `-d`

Decompresses the specified file instead of compressing it.

☞ `-1 to -9`

Controls the compression speed. `-1` compresses faster but produces larger files. `-9` requires more computing time but produces smaller files. The default setting is `-9`.

`bunzip2` decompresses a file compressed by `bzip2` and removes the suffix `.bz2`. `bunzip2` is equivalent to `bzip2 -d`.

☞ **Exercise: Compress a File – Part II**

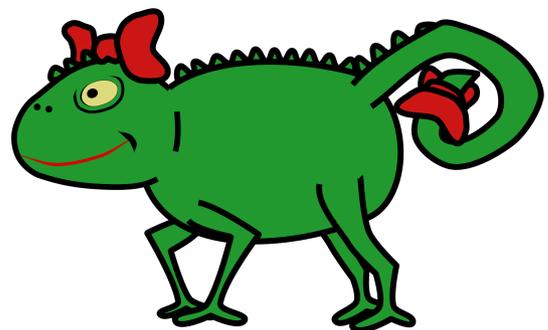
Compress the archive `/tmp/bin.tar` using `bzip2` with three different compression speeds. Write down the file size of the compressed file here:

`-1` _____

`-6` _____

`-9` _____

Copy the `tar` archive into your home directory and unpack the files into a new directory `Training`.



2.1.10 Manage File Permission and Ownership

You can use the command `ls -l` to display the contents of the current directory with the assigned permission for each file or subdirectory.

For example, entering `ls -l` displays the following permission for `MyFile`:

```
1 geeko@da51:~> ls -l MyFile
2 -rw-r--r-- 1 geeko users 0 2006-02-07 19:10 MyFile
geeko@da51:~>
```

You can use the command `chmod` to change the permission of a file. Both the owner of a file and root can use this command.

There are options to change the permission for the owner (“u”), group (“g”), other (“o”), or all (“a”). To add permission use the option “+”. To remove permission, use the option (“-”).

The following are examples of using `chmod`:

Example	Result
<code>chmod u+x</code>	The owner can execute the file.
<code>chmod g=rw</code>	All group members can read and write.
<code>chmod u=rwx</code>	The owner has all permissions.
<code>chmod u=rwx,g=rw,o=r</code>	All permissions for the owner, read and write for the group, read for all other users.
<code>chmod +x</code>	All users (owner, group, others) are allowed to execute the file.
<code>chmod a+x</code>	All users (owner, group, others) are allowed to execute the file (“a” for all).

In the following example, the user `geeko` allows the other members of the group `users` to write to the file `hello.txt` by using `chmod`:

```
1 geeko@da51:~ > ls -l hello.txt
2 -rw-r--r-- 1 geeko users 0 2006-04-06 12:40 hello.txt
geeko@da51:~ > chmod g+w hello.txt
4 geeko@da51:~ > ls -l hello.txt
-rw-rw-r-- 1 geeko users 0 2006-04-06 12:40 hello.txt
```

With the option `-R` (recursive) and a specified directory, you can change the access permission of all files and subdirectories under the specified directory.

**Exercise: Manage File Permissions – Part I**

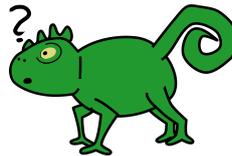
From the directory `Training` and its entire contents, remove the execute permission for “group” and “others”.

Besides using letters (`rwX`), you can also use the octal way of representing the permission letters with groups of numbers. Every file and directory in a Linux system has a numerical permission value assigned to it. This value has three digits.

The first digit represents the permission assigned to the file or directory owner. The second digit represents the permission assigned to the group associated with the file or directory. The third digit represents the permission assigned to others.

Each digit is the sum of the following three values assigned to it:

- ☞ Read: 4
- ☞ Write: 2
- ☞ Execute: 1



For example, a file named `MyFile.txt` has the permission of `754`.

This means the owner of the file can read, write, and execute the file ($4+2+1$), the group associated with the file can read and execute the file ($4+1$), and others can read it (4). If you use numbers, you can add them, as in the following:

Owner	Group	Others
<code>rwX</code>	<code>r-X</code>	<code>r--</code>
<code>421</code> ($4+2+1=7$)	<code>4-1</code> ($4+1=5$)	<code>4--</code> (4)

The following are examples of using numbers instead of letters:

Example	Result
<code>chmod 754 hello.txt</code>	All permissions for the owner, read and execute for the group, read for all other users (<code>rwX r-x r--</code>).
<code>chmod 777 hello.txt</code>	All users (user, group, others) can read, write and execute the file (<code>rwX rwX rwX</code>).

You also can set the special permission using `chmod`. The following table describes the special file permission for files and directories:

Char.	No.	Name	Files	Directories
<code>t</code>	1	Sticky bit	not applicable	Users can only delete files when they are the owner, or when they are root or the owner of the directory. This is usually applied to the directory <code>/tmp</code> .
<code>s</code>	2	SGID (set GroupID)	When a program is started, this sets the group ID of the process to that of the file's group.	Files created in this directory belong to the directory's group and not to the user's primary group. New directories created in this directory inherit the SGID bit.
<code>s</code>	4	SUID (set UserID)	Sets the user ID of the process to that of file's owner when the program is started.	Not applicable

You set the sticky bit with `chmod`, either by using letters (such as `chmod o+t /tmp`) or the numbers (e.g., `chmod 1777 /tmp`).

The sticky bit is listed in the permission for others (“t”), as in the following:

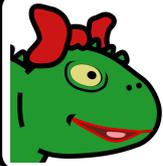
```
1 geeko@da51:~ > ls -ld /tmp
2 drwxrwxrwt 15 root root 608 2006-04-06 12:45 /tmp
geeko@da51:~ >
```

The following is an example for SUID:

```
1 geeko@da51:~ > ls -l /usr/bin/passwd
2 -rwsr-xr-x 1 root shadow 79765 2006-03-24 12:19 /usr/bin/passwd
geeko@da51:~ >
```

Each user is allowed to change his password, but you need root permission to write it into the file `/etc/shadow`. The following is an example for SGID:

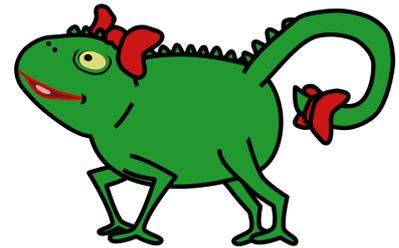
```
1 geeko@da51:~ > ls -l /usr/bin/wall
2 -rwxr-sr-x 1 root tty 10192 2006-03-22 05:24 /usr/bin/wall
geeko@da51:~ >
```



You can use the command `wall` to send messages to all virtual terminals. It is executed with the permission of the group `tty`.



Attention! These kinds of permissions should be used very carefully. They are only explained here briefly.



Exercise: Manage File Permissions – Part II

What does the following permission values for files mean?

☞ Permission 777

Owner _____

Group _____

Others _____

☞ Permission 755

Owner _____

Group _____

Others _____

☞ Permission 600

Owner _____

Group _____

Others _____

☞ Permission 644

Owner _____

Group _____

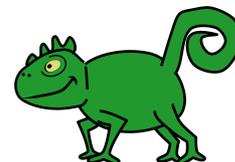
Others _____

☞ Permission 2755

Owner _____

Group _____

Others _____



The user root can use the command `chown` to change the user and group affiliation of a file by using the following syntax:

```
chown new_user.new_group file
```

For example:

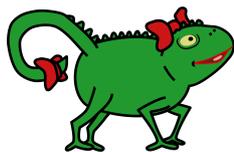
```
1 da51:~ # chown geeko.users MyFile
2 da51:~ #
```

To change the owner, but not the group, use the following command syntax:

```
chown new_user file
```

For example:

```
1 da51:~ # chown geeko MyFile
2 da51:~ #
```



To change the group, but not the user, use the following command syntax:

```
chown .new_group file
```

For example:

```
1 da51:~ # chown .users MyFile
2 da51:~ #
```

As root, you can also change the file's group affiliation with the command `chgrp`. Use the following syntax:

```
chgrp new_group file
```

For example:

```
1 da51:~ # chgrp users MyFile
2 da51:~ #
```

A normal user can use the command `chown` to allocate a file of your own to a new group. To do this, use the following syntax:

```
chown .new_group file
```

For example:

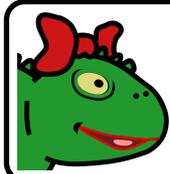
```
1 geeko@da51:~> chown .dialout MyFile
2 geeko@da51:~>
```

You can also use `chgrp`. Use the following syntax:

```
chgrp new_group file
```

For example:

```
1 geeko@da51:~> chgrp dialout MyFile
2 geeko@da51:~>
```



You can only change your file's group affiliation if you are a member of the new group.

In the following example, root uses `chown` to change the ownership of the file `hello.txt` from `geeko` to `tux`:

```
1 da51:~ # ls -l hello.txt
2 -rw-r--r-- 1 geeko users 0 2006-04-06 12:43 hello.txt
da51:~ # chown tux.users hello.txt
4 da51:~ # ls -l hello.txt
5 -rw-r--r-- 1 tux users 0 2006-04-06 12:43 hello.txt
6 da51:~ #
```

In the following example, root uses `chown` to limit access to the file `list.txt` for members of the group `training`:

```
1 da51:~ # ls -l list.txt
2 -rw-r----- 1 geeko users 0 2006-04-06 12:43 list.txt
da51:~ # chown .training list.txt
4 da51:~ # ls -l list.txt
5 -rw-r----- 1 geeko training 0 2006-04-06 12:43 list.txt
6 da51:~ #
```

Of course, root and the file owner can still access the file. Although the group has changed, the owner permission remains the same.

The most important option for `chown` and `chgrp` is `-R` to change file owner and owning group recursively. That means that all files in a subdirectory are also changed.

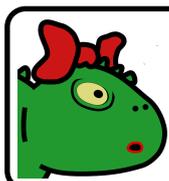


Exercise: Manage File Ownership

Change the owning group of the directory `Training` and its content to `dialout`.

2.2 Understand bash Commands and Features

A modern program has to be graphical with colorful dialogs, buttons, drag & drop functions, etc. The bash shell seems to be unglamorous compared with a Microsoft operating system. Typing on a keyboard seems to be much harder than clicking with a mouse.



Did you ever tried to change the permission of 120 files using Windows?

Nevertheless, bash provides some functions that simplifys your work. The coolest will be explained in this chapter.

2.2.1 Use Completion of Commands and Filenames

The bash shell supports a function that completes commands and file-names. Just enter the first characters of a command (or a filename) and press `[Tab]`. The bash shell completes the name of the command or file. For example, in your home directory press ...

 `[M]`, `[K]`, `[D]`, `[Tab]` → `mkdir`

 `[C]`, `[D]`, `[]`, `[/]`, `[H]`, `[Tab]` → `cd /home`

☞ `V`, `I`, `[]`, `.`, `B`, `A`, `Tab` → `vi .bash` → `[]`,
`Tab` → `vi .bash_history`

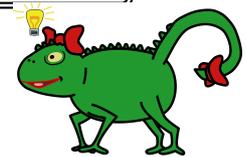
If there is more than one possible match, the bash shell shows all possibilities when you press the `Tab` key a second time.

For example:

`M`, `K`, `Tab`, `Tab` →

1	<code>mkallcomposecaches</code>	<code>mkfontdir</code>	<code>mkmanifest</code>
2	<code>mkcfm</code>	<code>mkfontscale</code>	<code>mknod</code>
	<code>mk_cmds</code>	<code>mkhtmlindex</code>	<code>mkpasswd</code>
4	<code>mkcomposecache</code>	<code>mkhybrid</code>	<code>mktemp</code>
	<code>mkdir</code>	<code>mkinfodir</code>	<code>mkzftree</code>
6	<code>mkdirhier</code>	<code>mkinodedb2</code>	<code>mkzimage_cmdline</code>
	<code>mkfifo</code>	<code>mkisofs</code>	

This feature facilitates entering long filenames.



2.2.2 Use the bash History

bash stores the commands you enter so you have easy access to them. By default, the commands are written into the file `.bash_history` in the user's home directory. The size of this file is set to a maximum of 1,000 entries.

To display the content of the file, use the command `history`.

```
1 geeko@da51:~> history
2   1  ls -i
3   2  ls -li
4   3  history
geeko@da51:~>
```

To display the commands stored in the history cache (one at a time), use the arrow keys. `Up-arrow` shows the previous command; `Down-arrow` shows the next command. After finding the command you are looking for, edit it as you like; then execute it by pressing `Enter`.

When browsing the entries of the history, you can also select specific commands. Typing one or several letters, or pressing `Page Up` or `Page Down`, displays the preceding or next command beginning with this letter in the history cache.

If you enter a part of the command (not necessarily the beginning of the command), press `Ctrl` + `R` to look for the commands containing the search string in the history list. The commands that match are displayed. The searching process starts with the last command executed.



Exercise: Use the bash History

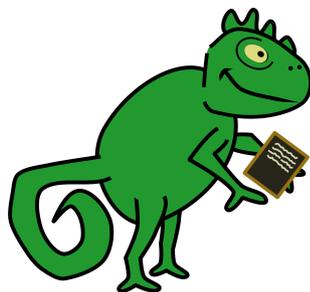
What are the last commands in your history including the following words:

☞ test. _____

☞ mkdir. _____

☞ mv. _____

☞ cp. _____



2.2.3 Use bash Variables

With shell and environment variables, you are able to configure the behavior of the shell and to customize its environment.

Environment variables control the behavior of a program that is started from a shell. Shell variables control the behavior of shell itself.

The convention is to write environment variables (e. g. PATH) in uppercase letters. If you set your own variables, they should also be written in capitals for the sake of clarity. In shell scripts, use lowercase letters for variable names.

Some important environment variables include the following:

PATH – If you start a program, the system looks for it in the directories specified here (the directories are separated by “:”). The order the directories are arranged in is important, since they are searched in turn.

HOME – The user’s home directory.

USER – The user’s login name.

To display the value of a shell or environment variable, enter `echo $variable`, as in the following:

```
1 geeko@da51:~ > echo $HOME
2 /home/geeko
geeko@da51:~ >
```

For clarity, we will write the “\$” in front of the variable’s name when we refer to a variable in the following text.



To set the value of a variable or to create a new variable, use the syntax `variable=value`, as in the following:

```
1 geeko@da51:~ > MYVAR=myvalue
2 geeko@da51:~ > echo $MYVAR
myvalue
4 geeko@da51:~ >
```

The value can be a number, a character, or a string. If the string includes a space, to write the value in full quotes, as in the following:

```
1 geeko@da51:~ > MYVAR="my value"
2 geeko@da51:~ > echo $MYVAR
my value
4 geeko@da51:~ >
```

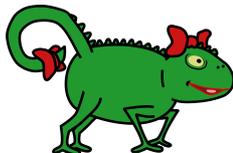
You can also use `echo` to display text on the screen.

```
1 geeko@da51:~> echo Hello Geeko
2 Hallo Geeko
geeko@da51:~>
```



Exercise: Use bash Variables – Part I

What is the content of the `$PATH` variable on your computer?



If you want to print something on the screen, such as Variable `$HOME` is set to `content_of_$HOME` the method above does not work.

```
1 geeko@da51:~> echo Variable $HOME is set to $HOME
2 Variable /home/geeko is set to /home/geeko
geeko@da51:~>
```

Setting the text in double quotes does not work, either.

```
1 echo "Variable $HOME is set to $HOME"
2 Variable /home/geeko is set to /home/geeko
geeko@da51:~>
```

If you use single quotes the result looks like the following.

```
1 geeko@da51:~> echo 'Variable $HOME is set to $HOME'
2 Variable $HOME is set to $HOME
geeko@da51:~>
```

There are several solutions, for example:

❶ A combination of quotes

```
1 geeko@da51:~> echo 'Variable $HOME is set to' "$HOME"
2 Variable $HOME is set to /home/geeko
geeko@da51:~>
```

❷ A combination of quotes and nonquotes

```
1 geeko@da51:~> echo 'Variable $HOME is set to' "$HOME"
2 Variable $HOME is set to /home/geeko
geeko@da51:~>
```

❸ Masking “\$” with a backslash

```
1 geeko@da51:~> echo "Variable \$HOME is set to $HOME"
2 Variable $HOME is set to /home/geeko
geeko@da51:~>
```



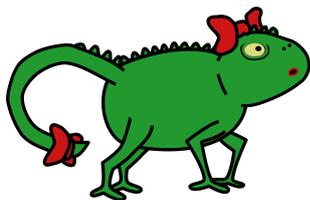
Exercise: Use bash Variables – Part II

Redo all examples in this section by using the variable `$PATH` instead of `$HOME`.

2.2.4 Use bash Aliases

Defining aliases allows you to create shortcuts for commands and their options or to create commands with entirely different names. For example, on a SUSE Linux Enterprise Desktop 10, whenever you enter the commands `dir`, `md`, or `ls`, you are using aliases.

The command `alias` shows you aliases defined on your system. For example, you will see that `dir` is an alias for `ls -l` and `md` is an alias for `mkdir -p`.



The following are examples of aliases that define new commands:

```

1 geeko@da51:~> alias md
2 alias md='mkdir -p'
3 geeko@da51:~> alias dir
4 alias dir='ls -l'

```

To find out if a command is an alias, use the `type` command. `type` tells you whether the command specified is a built-in shell command, a regular command, a function, or an alias.

For regular commands, the `type` lists the path to the corresponding executable. For aliases, it lists the elements aliased:

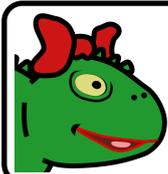
```

1 geeko@da51:~> type -a ls
2 ls is aliased to '/bin/ls $LS_OPTIONS'
3 ls is /bin/ls

```

Explanation: The above example shows that `ls` is an alias used to add some options to the command. The `-a` option shows both the contents of the alias and the path to the original `ls` command. The output shows that `ls` uses the options stored in the variable `LS_OPTIONS`. These options are responsible for `ls` listing different file types in different colors.

Most of the system aliases are defined in the file `/etc/bash.bashrc`. To define aliases, use the `alias` command. To delete them, use the `unalias` command.



To define an alias, use the following syntax:
`alias aliasname="command options"`

To delete an alias, use the following syntax:
`unalias aliasname`



An alias defined in this way is only valid for the current shell, as in the following:

```

1 geeko@da51:~> alias ls="echo Hello"
2 geeko@da51:~> ls
3 Hello
4 geeko@da51:~> bash
5 geeko@da51:~> ls
6 bin Desktop Documents public_html
7 geeko@da51:~>

```

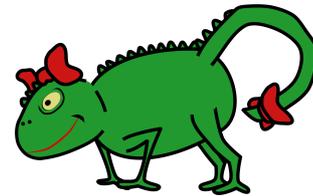
But if you terminate this shell, the alias is gone. To avoid this, store the alias in one of the shell's configuration files. On the SUSE Linux Enterprise Desktop, the file `~/.alias` is created for aliases defined by each user.

This file is read in by `~/.bashrc`, where a command is included to that effect. Aliases are not relevant to shell scripts at all, but can be a real time saver when using the shell interactively.



Exercise: Use bash Aliases

Create a new alias `spy` that displays the content of a `tar` archive compressed with `gzip`.



2.2.5 Use Search Patterns to Expand Names

Occasionally, you might want to perform operations on a series of files without having to name all the files. In this case, you can use of the following search patterns:



Replaces a single character (except `/`).

```
1 geeko@da51:~> ls /bin/s?
2 /bin/sh /bin/su
geeko@da51:~>
```



Replaces any string length, including zero characters (except `.` at the beginning of a file name and `/`).

```
1 geeko@da51:~> ls /bin/s*
2 /bin/sash /bin/setkeycodes /bin/sh /bin/sort
3 /bin/scsidev /bin/setleds /bin/showconsolefont /bin/stty
4 /bin/sed /bin/setmetamode /bin/showkey /bin/su
5 /bin/setfont /bin/setserial /bin/sleep /bin/sync
6 geeko@da51:~>
```

☞ [0-9]

Replaces any of the characters enclosed (here: numbers from 0 to 9).

```
1 geeko@da51:~> ls /bin/s[e-h]*
2 /bin/sed /bin/setleds /bin/sh
3 /bin/setfont /bin/setmetamode /bin/showconsolefont
4 /bin/setkeycodes /bin/setserial /bin/showkey
geeko@da51:~>
```

☞ [abcd]

Replaces one of the characters a, b, c and d.

```
1 geeko@da51:~> ls /bin/s[abcd]*
2 /bin/sash /bin/scsidev
```

☞ [a-ew-z]

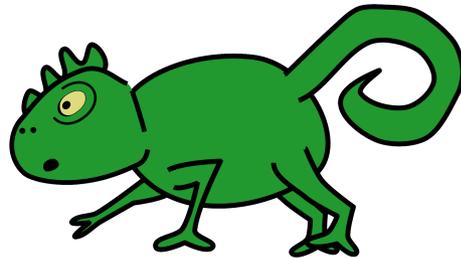
Replaces any character from the ranges a-e and w-z.

```
1 geeko@da51:~> ls /bin/s[a-cm-t]*
2 /bin/sash /bin/scsidev /bin/sort /bin/stty
geeko@da51:~>
```

☞ [!abc]

Replaces any character, except a, b, and c.

```
1 geeko@da51:~> ls /bin/s[!eh]*
2 /bin/sash /bin/scsidev /bin/sleep /bin/sort /bin/stty /bin/su /bin/sync
geeko@da51:~>
```



! **Attention!** The meaning of some search patterns differs from their meaning as regular expressions (see 2.3.2 on page 89).

If you use search patterns (wild cards), the shell compares them with all filenames and, if they match, replaces the expression with all filenames found.

! **Attention!** The order of characters is different for normal users and for the user root, because both use a different character encoding by default.

For normal users (UTF-8 encoding) the lowercase character follows the uppercase character immediately. For root (POSIX encoding) all the lowercase characters are behind the uppercase characters. This means that between “A” and “C” in POSIX, the only character is “B”. But, in UTF-8, there are “a,” “B,” and “b” in between.

The variable `$LANG` specifies the language. In the following example, the language is set to US English using UTF-8 encoding.

```
1 geeko@da51:~> echo $LANG
2 en_US.UTF-8
geeko@da51:~>
```

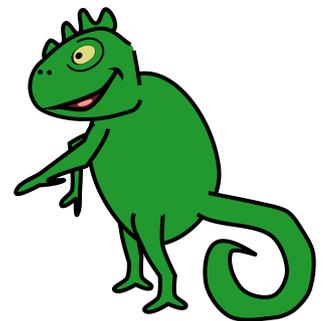
UTF-8 (Unicode) handles all kind of letters and not only Latin letters, while POSIX (ASCII) only does Latin letters.



Exercise: Use Search Patterns to Expand Names

How many programs in the `/bin` directory have a name...

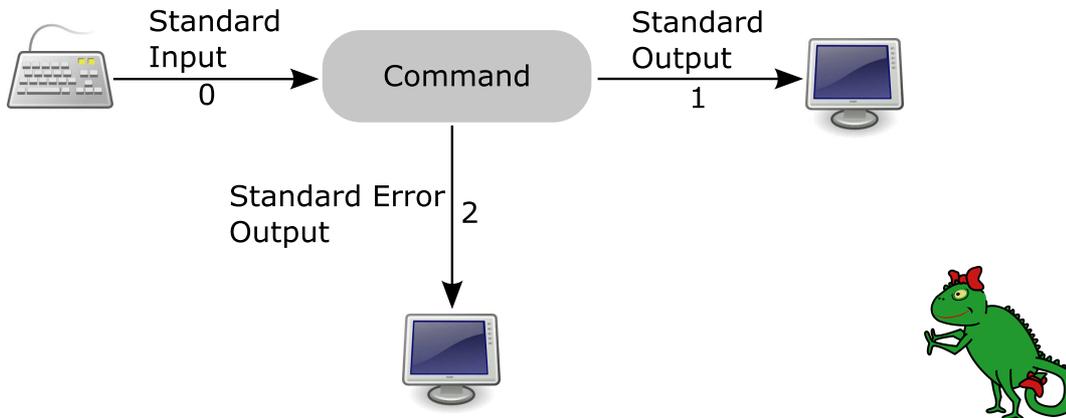
- ... that begins with “sa”? _____
- ... that begins with “x”, “y”, or “z”? _____
- ... that ends with “s”? _____
- ... with three letters? _____
- ... that include at least one number? _____



2.3 Understand Advanced Command Execution

2.3.1 Use Piping and Redirection

Linux has three standard data channels:



The following describes these channels:

Channel	Number	Description
Standard input (stdin)	0	The currently running program reads the input from this channel (usually the keyboard).
Standard output (stdout)	1	The program sends its output to this channel (usually the monitor).
Standard error output (stderr)	2	Errors are issued through this channel (usually the monitor).

Each channel can be redirected by the shell. For example, standard input can come from a file or standard output and standard error output can be directed to a file. The following are the redirection characters:

- < – Redirects standard input.
- > – Redirects standard output (“>” without a preceding number is just an abbreviation for “1>”).
- 2> – Redirects standard error output.



“>” overwrites an existing file. If the output should be appended to an existing file, you have to use “>>” or “2>>”.

The following is an example of a standard input, standard output, and standard error output. The default behavior of the shell is shown here:

```
1 geeko@da51:~ > ls /opt /recipe
2 /bin/ls: /recipe: No such file or directory
  /opt:
4 gnome kde3
```

Explanation: The content of the directory `/opt` is shown in line 3 and 4. The directory `/recipe` does not exist. An error message is shown in line 2.

If the standard error output is redirected to `/dev/null`, only the standard output is displayed on the screen:

```
1 geeko@da51:~ > ls /opt /recipe 2> /dev/null
2 /opt:
  gnome kde3
```

To redirect standard output and standard error output to a file (such as “output”), enter the following:

```
1 geeko@da51:~ > ls /opt /recipe > output 2>&1
2 geeko@da51:~ >
```

Explanation: First, the standard output is redirected to the file `output` (`> output`); then the standard error output is directed to the standard output (`2>&1`). The “&” refers to the file descriptor that follows (“1” for the standard output).

To display the contents of the file list, use the command `cat`, as in the following:

```
1 geeko@da51:~> cat list
2 /bin/ls: /recipe: No such file or directory
  /opt:
4 gnome
  kde3
```

This option of process communication is available not only in the shell, but can also be used in programs directly. All known files in the system can be used as input or output.



Occasionally, you might want to use a file as input for a program that expects input from the keyboard. To do this, the standard input is redirected, as in the following:

```
1 geeko@da51:~ # echo "Hello Tux,"
2 >
3 > how are you?
4 > Is everything okay?" > greetings
geeko@da51:~ # mail tux < greetings
```

Explanation: First, the text is redirected to the file `greetings` through the command “>” (lines 1–3). The mail program `mail` receives its input from the file `greetings` (not from the keyboard), and then the email program sends the email to the user `geeko` (line 4).



Exercise: Use Piping and Redirection – Part I

Create a file labeled `Redirection` in your home directory that includes...

- ❶ ... the output of the command `ls /`; followed by ...
- ❷ ... the content of the variable `$PATH`; followed by ...
- ❸ ... the error message of the command `cd /root`.

The output of one command can be used as the input for another command by using the pipe (“|”):

```
command1 | command2
```

In a pipe, a maximum of 4 KB of unprocessed data can exist. If the process creating the output tries to write to a full pipe, it is stopped and only allowed to continue if the writing process can be completed. On the other side, the reading process is stopped if it tries to read an empty pipe.

```
1 geeko@da51:~ > ls -l /etc | less
```

Occasionally, you might want output from a command displayed on the screen and written to a file. To do this, use the command `tee`:

```
1 geeko@da51:~ > ls -l | tee output
```

In this example, the output of the command is displayed on the screen as well as written to the file `output`.



To redirect the output of several consecutive commands on the command line, the commands must be separated with semi-colons and enclosed in parentheses (`command1 ; command2 ; ...`):

```

1 geeko@da51:~> (id ; ls ~) > output
2 geeko@da51:~> cat output
uid=1000(geeko) gid=100(users)
4 groups=14(uucp),16(dialout),33(video),100(users)
bin
6 Desktop
Documents
8 output
public_html
10 geeko@da51:~>

```



Exercise: Use Piping and Redirection – Part II

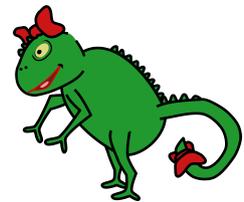
The command `wc` (word count) counts words in a text file. With the option `-l` it counts lines.

Pipe the output of the command `ls -l` into the command `wc -l`.

What does the result of this construction mean?

The shell starts a separate subshell for processing the individual commands. To redirect the linked commands, the shell must be forced to execute the command chain in the same subshell by enclosing the expression in parentheses.

Upon completion, every program returns a value that states the success of the execution. If this return value is 0, the command completed successfully. If an error occurred, the return value is greater than 0. (Depending on the program, different return values indicate different errors.)



You can use the command `echo $?` to display a return value.

```

1 geeko@da51:~> ls
2 bin Desktop Documents public_html
geeko@da51:~> echo $?
4 0
geeko@da51:~>

```

The return value can be used to trigger the execution of another command:

Link	Result
<code>command1 && command2</code>	<code>command2</code> is only executed if <code>command1</code> is completed without any errors.
<code>command1 command2</code>	<code>command2</code> is only executed if <code>command1</code> is completed with an error.

The following illustrates using both “||” and “&&”:

```

1 geeko@da51:~> ls recipe || ls ~
2 /bin/ls: recipe: No such file or directory
  bin Desktop Documents output public_html test
4 geeko@da51:~> ls recipe && ls ~
  /bin/ls: recipe: No such file or directory
6 geeko@da51:~>

```

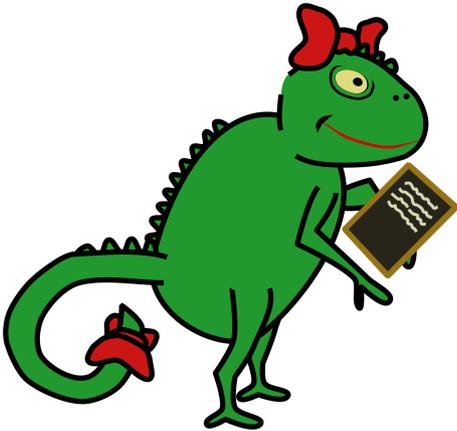
Explanation: The file `recipe` does not exist and the command `ls recipe` leads to an error (lines 1–2). Because of this, the command `ls ~` in line 1 is executed and in the line 4 is not.



Exercise: Use Piping and Redirection – Part III

What is the return value of the following commands?

- ☞ `ls -l | wc -l` _____
- ☞ `cd /root` _____
- ☞ `less wall` _____
- ☞ `echo $LOREM` _____
- ☞ `echo $?` _____



2.3.2 Understand Regular Expressions

Regular expressions are strings consisting of metacharacters and literals (regular characters and numerals). In the context of regular expressions,

metacharacters are those characters that do not represent themselves but have special meanings.

They can act as placeholders for other characters or can be used to indicate a position in a string.



Attention! It is important to remember that some metacharacters used by the shell for filename expansion have a meaning different from the one discussed here.

Many commands rely on regular expressions for pattern matching, for example, the command `grep` and its variant `egrep`. These commands are used to search files for certain patterns using the syntax `grep search_pattern filename`

The command searches `filename` for all text that matches `search_pattern` and prints the lines that contains the pattern. You can also specify several files, in which case the output will not only print the matching line, but also the corresponding filenames.

Several options are available to specify that only the line number should be printed, for instance, or that the matching line should be printed together with leading and trailing context lines.

You can specify search patterns in the form of regular expressions, although the basic `grep` command is limited in this regard. To search for more complex patterns, use the `egrep` command (or `grep -E`) instead, which accepts extended regular expressions.

As a simple way to deal with the difference between the two commands, make sure you use `egrep` in all of your shell scripts. The regular expressions used with `egrep` need to comply with the standard syntax of regular expressions.



You can read details on this topic in the manual page of `grep`.

To avoid having special characters in search patterns interpreted by the shell, enclose the pattern in quotation marks.

The following is an example of using `egrep` and `grep`:

```

1 geeko@da51:~> egrep (b|B)lurb file*
2 bash: syntax error near unexpected token '|'
3 geeko@da51:~> grep "(b|B)lurb" file*
4 geeko@da51:~> egrep "(b|B)lurb" file*
5 file1:blurb
6 filei2:Blurb
7 geeko@da51:~>

```



On the book CD, you can find an archive file “grep-demo.tar.gz” with text files with some meaningless content. We will use these text files for the following explanations.

Uncompress the archive into your home directory.

The following are options you can use with the command `grep`:

 `-i`

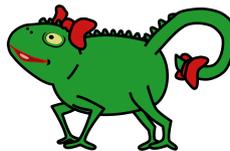
Enables case insensitivity.

```

1 geeko@da51:~> grep "duis" *.txt
2 loremipsum.txt:delenit au gue duis dolore te feugate nulla facilisi.
3 loremipsum.txt:au gue duis dolore te feugate nulla facilisi.
4 geeko@da51:~> grep -i "duis" *.txt
5 loremipsum.txt:Duis te feugifacilisi.
6 loremipsum.txt:Duis autem dolor in hendrerit
7 loremipsum.txt:delenit au gue duis dolore te feugate nulla facilisi.
8 loremipsum.txt:Duis te feugifacilisi per suscipit
9 loremipsum.txt:Duis te feugifacilisi.
10 loremipsum.txt:Duis autem dolor in hendrerit in
11 loremipsum.txt:au gue duis dolore te feugate nulla facilisi.
12 geeko@da51:~>

```

Explanation: The first `grep` command (line 1) finds two lines in the text file (lines 2–3). With the option `-i` (line 4), `grep` ignores upper and lower case and lists the lines with a capital “D” in the word “duis” (lines 5–11).



 `-l`

Shows only the names of files that contain the search string.

```

1 geeko@da51:~> grep "duis" *.txt
2 loremipsum.txt:delenit au gue duis dolore te feugate nulla facilisi.
3 loremipsum.txt:au gue duis dolore te feugate nulla facilisi.
4 geeko@da51:~> grep -l "duis" *.txt
5 loremipsum.txt
6 geeko@da51:~>

```

Explanation: If `grep` finds more than one file, it writes the filename at the beginning of each line (lines 2–3). Using the option `-l` (line 4), only the filename is written on the screen (line 5).

 `-r`

Searches entire directory trees recursively. This option only works if you specify a directory and not files in the the `grep` command.

```
1 geeko@da51:~> grep "duis" *
2 loremipsum.txt:delenit au gue duis dolore te feugat nulla facilisi.
3 loremipsum.txt:au gue duis dolore te feugat nulla facilisi.
4 geeko@da51:~> grep -r "duis" *
5 dirgrep/loremipsum2.txt:delenit au gue duis dolore te feugat nulla facilisi.
6 dirgrep/loremipsum2.txt:au gue duis dolore te feugat nulla facilisi.
7 loremipsum.txt:delenit au gue duis dolore te feugat nulla facilisi.
8 loremipsum.txt:au gue duis dolore te feugat nulla facilisi.
geeko@da51:~>
```

Explanation: The current directory is scanned in line 1. Only the hits of the file `loremipsum.txt` are shown in the output (lines 2–3). With the option `-r` (line 4), the sub-directory `dirgrep` is also scanned and the hits in the file `loremipsum2.txt` are listed (lines 5–8).

 `-v`

Shows all lines that do not contain the search string.

```
1 geeko@da51:~> grep "duis" *.txt
2 loremipsum.txt:delenit au gue duis dolore te feugat nulla facilisi.
3 loremipsum.txt:au gue duis dolore te feugat nulla facilisi.
4 geeko@da51:~> grep -v "duis" *.txt
5 loremipsum.txt:LOREM IPSUM DOLOR SIT AMET
6 loremipsum.txt:
7 loremipsum.txt>Lorem ipsum dolor sit amet,
8 loremipsum.txt:consectetuer adipiscing elit,
9 ...
10 numbers.txt:0
11 numbers.txt:98798
12 numbers.txt:765765
geeko@da51:~>
```



 `-n`

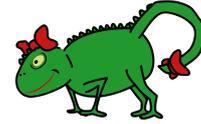
Shows the line numbers.

```
1 geeko@da51:~> grep -n "duis" *.txt
2 loremipsum.txt:18:delenit au gue duis dolore te feugat nulla facilisi.
3 loremipsum.txt:39:au gue duis dolore te feugat nulla facilisi.
4 geeko@da51:~>
```

 `-h`

Shows no filenames.

```
1 geeko@da51:~> grep -h "duis" *.txt
2 delenit au gue duis dolore te feugat nulla facilisi.
  au gue duis dolore te feugat nulla facilisi.
4 geeko@da51:~>
```



The following list shows the most important metacharacters and their meanings:

☞ ^

The beginning of the line.

```
1 geeko@da51:~> grep -n "^et" *.txt
2 loremipsum.txt:16:et accusan et iusto odio dignissim
  loremipsum.txt:37:et iusto odio dignissim qui blandit
4 geeko@da51:~>
```

Explanation: With `^et`, “et” is matched if it appears at the beginning of the line.

☞ \$

The end of the line.

```
1 geeko@da51:~> grep -n "at$" *.txt
2 loremipsum.txt:14:vel illum dolore eu feugiat
4 geeko@da51:~>
```

Explanation: With `$at`, “at” is matched if it appears at the end of line.

☞ \<

The beginning of the word

```
1 geeko@da51:~> grep -n "\<se" *.txt
2 loremipsum.txt:5:sed diem nonummy nibh euismodtincidunt
  loremipsum.txt:27:sed diem nonummy nibh euismodtincidunt
4 geeko@da51:~>
```

Explanation: With `\<se`, all words are matched that start with “se”.

☞ \>

The end of the word.

```
1 geeko@da51:~> grep -n "se\>" *.txt
2 loremipsum.txt:13:in vulputate velit esse molestie consequat,
  loremipsum.txt:34:vulputate velit esse molestie consequat,
4 geeko@da51:~>
```

Explanation: With `se\>`, all words are matched that end with “se”.

☞ [characters]

One character from the set.

```

1 geeko@da51:~> grep -n "[ex]e" *.txt
2 loremipsum.txt:6:ut lacreet dolore magna aliquam erat volutpat.
  loremipsum.txt:8:quis nostrud exerci tution ullamcorper
4 loremipsum.txt:20:quis nostrud exerci taion
  loremipsum.txt:28:ut lacreet dolore magna aliquam erat volutpat.
6 loremipsum.txt:30:quis nostrud exerci tution ullamcorper
geeko@da51:~>

```

Explanation: `[ex]e` matches all lines containing “ee” or “xe”.

Using a hyphen you can specify a range of characters.

```

1 geeko@da51:~> grep -n "^[a-c]" *.txt
2 loremipsum.txt:4:consectetuer adipiscing elit,
  loremipsum.txt:26:adipiscing elit,
4 loremipsum.txt:39:au gue duis dolore te feugat nulla facilisi.
geeko@da51:~>

```

Explanation: `^[a-c]` matches all lines beginning with the characters “a”, “b”, or “c”.

☞ `[^characters]`

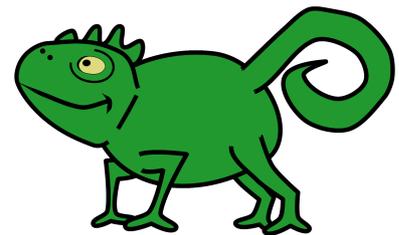
None of the specified characters.

```

1 geeko@da51:~> grep -n "^[^a-c]" *.txt
2 loremipsum.txt:1:LOREM IPSUM DOLOR SIT AMET
  loremipsum.txt:3:Lorem ipsum dolor sit amet,
4 loremipsum.txt:5:sed diem nonummy nibh euismodtincidunt
  ...
6 numbers.txt:11:0
  numbers.txt:98798
8 numbers.txt:765765
geeko@da51:~>

```

Explanation: The `grep` command matches all lines that do not begin with the characters “a”, “b”, or “c”.



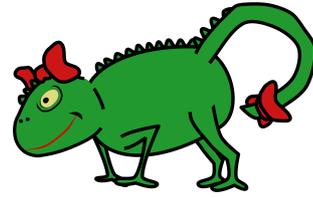
Any single character.

```

1 geeko@da51:~> grep -n "^.$" *.txt
2 numbers.txt:11:0
geeko@da51:~>

```

Explanation: `^.$` matches all lines with only one character (between the beginning and ending of the line).

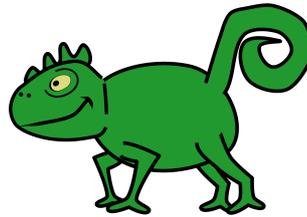


☞ + (Works only with `egrep`)

One or more of the preceding expression.

```
1 geeko@da51:~> egrep -n "^0+$" *.txt
2 numbers.txt:1:0000
  numbers.txt:11:0
4 geeko@da51:~>
```

Explanation: `^0+$` matches all lines with one or more “0” (between the beginning and ending of the line).



☞ *

Any number (including none) of the preceding single character.

```
1 geeko@da51:~> grep -n "^0*$" *.txt
2 loremipsum.txt:2:
  numbers.txt:1:0000
4 numbers.txt:11:0
geeko@da51:~>
```

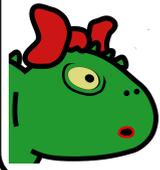
Explanation: `^0*$` matches all lines with zero, one or more “0” (between the beginning and ending of the line).

☞ `{min,max}` (Works only with `egrep`)

The preceding expression appears *min* times at minimum and *max* times at maximum.

```
1 geeko@da51:~> egrep -n "[0-9]{4,5}" *.txt
2 numbers.txt:1:0000
  numbers.txt:12:98798
4 numbers.txt:13:765765
geeko@da51:~>
```

Explanation: `[0-9]{4,5}` matches all lines with four (at minimum) and five (at maximum) numbers.



Are you surprised why there is a six-digit number matched as well in line 4?

The mistake with the previous expression is, that we did not specify what is before or after the fourth or fifth numbers. Have a look the following listing.

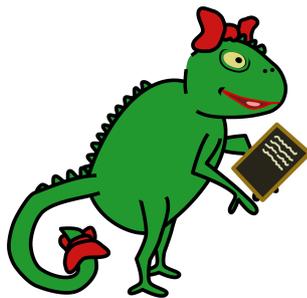
```
1 geeko@da51:~> egrep -n "^ [0-9]{4,5}$" *.txt
2 numbers.txt:1:0000
  numbers.txt:12:98798
4 geeko@da51:~>
```

☞ `|` (Works only with `egrep`)

The expression before or after the pipe.

```
1 geeko@da51:~> egrep -n "ea|gue" *.txt
2 loremipsum.txt:10:ut aliquip ex ea commodo consequat.
  loremipsum.txt:18:delenit au gue duis dolore te feugiat nulla facilisi.
4 loremipsum.txt:31:suscipit lobortis nisl ut aliquip ex ea commodo consequat.
  loremipsum.txt:39:au gue duis dolore te feugiat nulla facilisi.
6 geeko@da51:~>
```

Explanation: `ea|gue` matches all lines with the expression “ea” or “gue”.



(...) (Works only with egrep)

Enclose alternatives for grouping with others.

```
1 geeko@da51:~> egrep -n "(d|D)uis" *.txt
2 loremipsum.txt:11:Duis te feugifacilisi.
  loremipsum.txt:12:Duis autem dolor in hendrerit
4 loremipsum.txt:18:delenit au gue duis dolore te feugat nulla facilisi.
  loremipsum.txt:23:Duis te feugifacilisi per suscipit
6 loremipsum.txt:32:Duis te feugifacilisi.
  loremipsum.txt:33:Duis autem dolor in hendrerit in
8 loremipsum.txt:39:au gue duis dolore te feugat nulla facilisi.
geeko@da51:~>
```

Explanation: (d|D)uis matches “duis” and “Duis”.

? (Works only with egrep)

Zero or one of the preceding characters.

```
1 geeko@da51:~> egrep -n "it,$" *.txt
2 loremipsum.txt:4:consectetuer adipiscing elit,
  loremipsum.txt:26:adipiscing elit,
4 geeko@da51:~> egrep -n "it,?$" *.txt
5 loremipsum.txt:4:consectetuer adipiscing elit,
6 loremipsum.txt:12:Duis autem dolor in hendrerit
  loremipsum.txt:23:Duis te feugifacilisi per suscipit
8 loremipsum.txt:26:adipiscing elit,
  loremipsum.txt:37:et iusto odio dignissim qui blandit
10 loremipsum.txt:38:praesent luptatum zzril delenit
geeko@da51:~>
```

Explanation: it,\$ (line 1) matches all lines that end with “it,”. it,?\$ (line 1) matches all lines that end with “it,” or “it”.

\

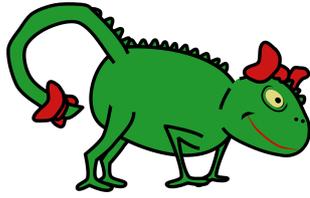
Masks the following character to remove its special meaning.

```
1 geeko@da51:~> grep -n "\.$" *.txt
2 loremipsum.txt:6:ut lacreet dolore magna aliquam erat volutpat.
  loremipsum.txt:10:ut aliquip ex ea commodo consequat.
4 loremipsum.txt:11:Duis te feugifacilisi.
  loremipsum.txt:18:delenit au gue duis dolore te feugat nulla facilisi.
6 loremipsum.txt:22:nisl ut aliquip ex en commodo consequat.
  loremipsum.txt:24:lobortis nisl ut aliquip ex en commodo consequat.
8 loremipsum.txt:28:ut lacreet dolore magna aliquam erat volutpat.
  loremipsum.txt:31:suscipit lobortis nisl ut aliquip ex ea commodo consequat.
10 loremipsum.txt:32:Duis te feugifacilisi.
  loremipsum.txt:39:au gue duis dolore te feugat nulla facilisi.
12 geeko@da51:~>
```

Explanation: \.\$ matches all lines that end with a full stop. Without masking the full stop represents any single character (see page 94).



To learn more about the structure of regular expressions, read the corresponding manual page with `man 7 regex`.



Exercise: Understand Regular Expressions – Part I

What do the following regular expressions mean? How many lines match from the file `loremipsum.txt`?

☞ `^.t` _____

☞ `it\>` _____

☞ `m{2}` _____

☞ `[^.,]$` _____



Exercise: Understand Regular Expressions – Part II

Create regular expressions that match the following demands. How many hits do you get in the file `loremipsum.txt`?

☞ All lines with at least one upper-case character.

☞ All lines with the word “in” or “ad”.

☞ All lines with a three letter word.

2.3.3 Manage Shell Processes

The Linux shell environment allows processes to run either in the foreground or in the background.

Processes executed in the foreground are started in a terminal window and run until the process completes; the terminal window does not return to a prompt until the program's execution is complete.

Background process execution occurs when a process is started and the terminal window returns to a prompt before the process completes.

Existing processes can be switched from foreground to background execution under the following circumstances:

- ☞ The process must be started in a terminal window or console shell.
- ☞ The process does not require input from the terminal window.

If the process meets this criteria, it can be moved to the background.



Attention! Processes that require input within the terminal can be moved to the background as well, but when input is requested, the process will be suspended until it is brought to the foreground and the requested input is provided.

Commands in a shell can be started in the foreground or in the background. Processes in the foreground can directly receive transmitted signals.

For example, if you enter `xeyes` to start the `xeyes` program, it is running in the foreground. If you press `Ctrl+Z`, the process stops:

```
1 [1]+ Stopped xeyes
2 geeko@da51 :~>
```

You can continue running a stopped process in the background by entering `bg`, as in the following:

```
1 geeko@da51 :~> bg
2 [1]+ xeyes &
geeko@da51 :~>
```

The ampersand (line 2) displayed in the output means that the process is now running in the background. Appending an ampersand to a command starts the process in the background (instead of the foreground), as in the following:

```
1 geeko@da51:~> xeyes &
2 [2] 4351
geeko@da51:~>
```

With this, the shell from which you started the program is available again for user input immediately.

In the above example, both the job ID ([2]) and the process ID of the program (4351) are returned.

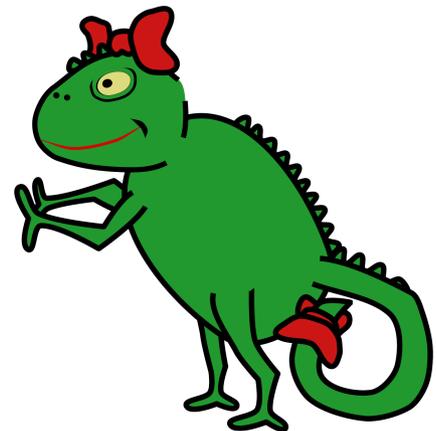
Each process started from the shell is assigned a job ID by the job control of the shell. The command `jobs` lists the contents of job control, as in the following:

```
1 geeko@da51:~> jobs
2 [1]+ Stopped xeyes
   [2] Running xeyes &
4 [4]- Running sleep 99 &
geeko@da51:~>
```

In this example, the process with job ID 3 is already terminated. The processes 2 and 4 are running in the background (notice the ampersand), and process 1 is stopped.

The “+” sign indicates the process that will respond to `fg` without options, and the “-” sign indicates the process that inherits the “+” sign once the process with the “+” sign ends.

The next background process will be assigned the job ID of 5 (highest number+1).



Not only can you continue running a stopped process in the background by using the command `bg`, you can also switch a process to the foreground by entering `fg job_ID`, as in the following:

```
1 geeko@da51:~> fg 1
2 xeyes
```

The shell also informs you about the termination of a process running in the background:

```
1 [4]- Done sleep 99
```

The job ID is displayed in square brackets. Done means the process terminated properly. If you see **Terminated** instead, it means that the process was requested to terminate. **Killed** indicates a forceful termination of the process.



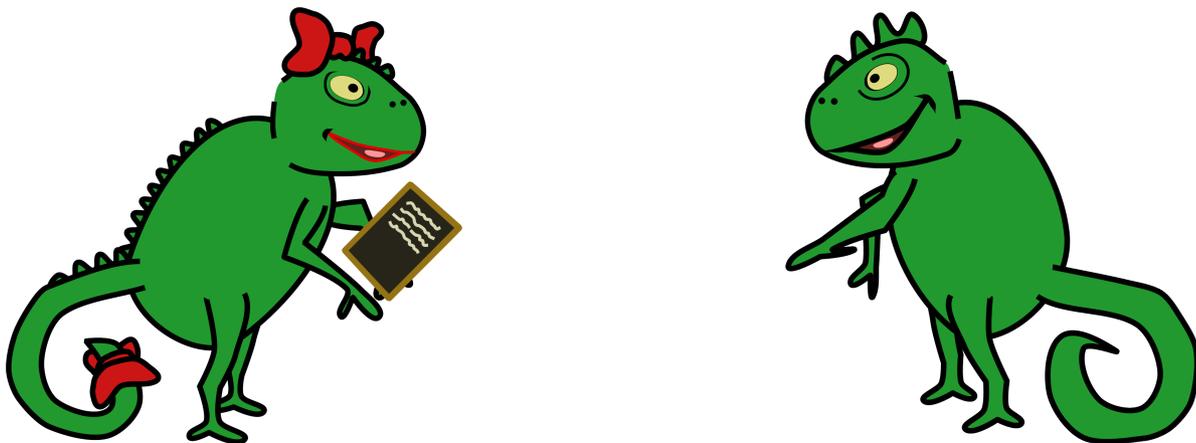
Exercise: Manage Shell Processes

Start the Acrobat Reader showing a PDF file from the command line, by entering:

```
acroread /usr/share/doc/manual/sled-gnome-user_en/  
sled-gnome-user_en.pdf
```

Stop the process. What happens when you try to use the Acrobat Reader?

Move the `acroread` process in the background. What happens when you try to use the Acrobat Reader?



2.4 The vi Text Editor

Because most of the services of a Linux computer are configured by editing an ASCII file, you need a text editor. There are a lot of text editors available in Linux, for example:

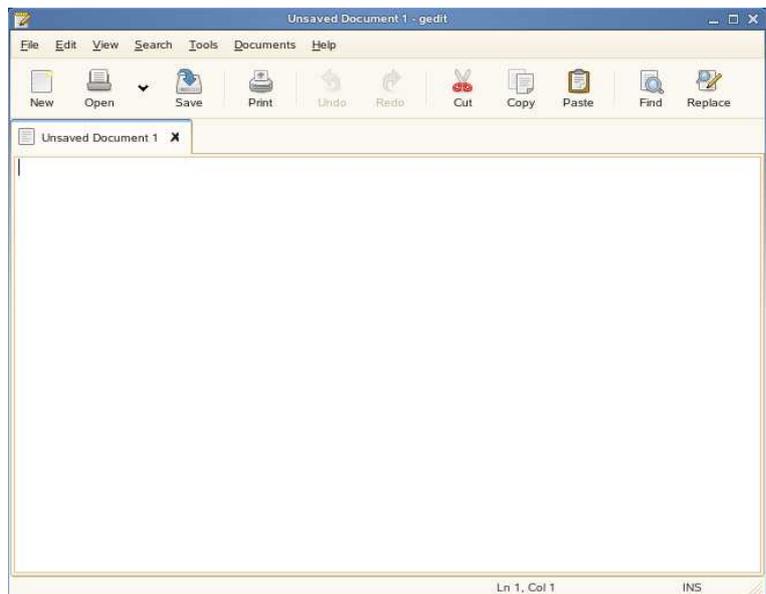
- ☞ vi
- ☞ emacs
- ☞ xemacs
- ☞ xedit
- ☞ gedit
- ☞ kwrite

Every text editor has advantages and disadvantages. There are two kinds of editors:

- ☞ Command line editors
- ☞ Graphical editors

The main advantage of command line editors is that you do not need a graphical user interface to use them. They run in a shell.

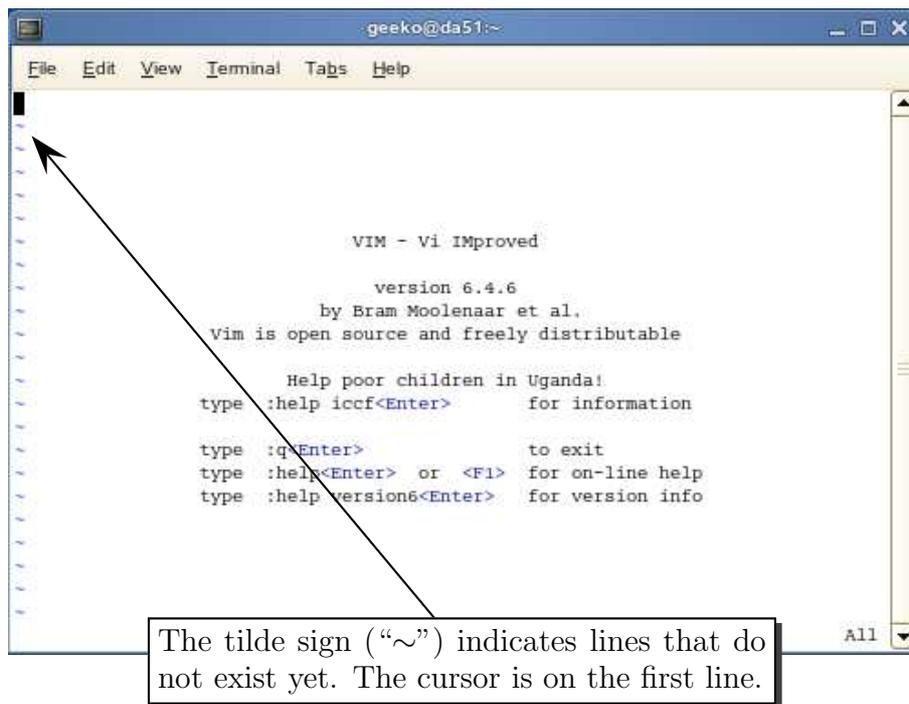
Graphical editors are (normally) easy to use and do not need big explanations. An example is the editor gedit that can be started from the main menu (application group Tools).



Although many factors can be involved when selecting an editor for everyday use, the reason vi is used by most administrators is that it is available on every Linux and UNIX system. Because of this, you should be able to use vi. In SUSE Linux Enterprise Desktop, vim (vi improved) by Bram Moolenaar is the standard vi editor. When you enter vi, vim is started via a link to it.

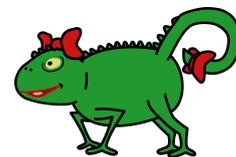
2.4.1 Start vi

You can start vi by entering vi or vim.



You can also enter vi followed by various options. To edit a file, enter vi and the name of a file, as in the following example: `vi exercise`

If a file does not yet exist, it is created. The text of the file appears in an editor at the command line.



2.4.2 Move the Cursor

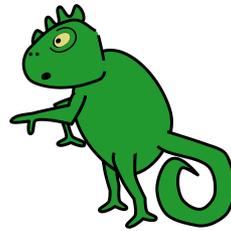
You can move the cursor with the **K**, **J**, **H**, and **L** keys.



Exercise: Move Cursor in vi – Part I

Open a text file with vi. Insert the direction the cursor moves by pressing the following keys:

Key	Direction
K	
J	
H	
L	



If available, you also can use the arrow keys (**Up-arrow**, **Down-arrow**, **Left-arrow**, and **Right-arrow**) to move the cursor.

With **Ctrl** + **V** you can move one screen page down, with **Alt** + **V** one screen page up.

Some more interesting keys for moving the cursor:

- E** – Moves the cursor to the end of the current word.
- W** – Moves the cursor to the beginning of the next word.
- B** – Moves the cursor to the beginning of the previous word.
- O** – Moves the cursor to the beginning of the current line.
- \$** – Moves the cursor to the end of the current line.
- G** + **G** – Moves the cursor to the beginning of the first line.
- Shift** + **G** – Moves the cursor to the beginning of the last line.



Exercise: Move Cursor in vi – Part II

Open a text file with vi and practice the movements mentioned above.

2.4.3 Learn the Working Modes

In contrast to many other editors, vi is mode-oriented. When vi is first started, it is in normal mode. Anything you enter in this mode is considered a command. You must switch to input mode before you can type any text. This can be frustrating to users who are unfamiliar with vi.

In addition to switching modes, you must learn which keys perform which actions, because you cannot use the mouse. However, the number of commands needed for everyday work is fairly small, and you can get used to them quickly.

To enter text, you must first switch the editor to input mode by typing **I** (insert) or pressing the **Ins** key. At the bottom of the screen, you see the message `-- INSERT --`.



Press **Esc** once to take you back to the normal mode. From normal mode you can switch to command-line mode by entering ":". The cursor jumps to the last line after ":" and waits for a command entry.

A command will only be carried out in command-line mode after you press **Enter**. Then you are automatically back in normal mode.

The following is a summary of the available modes (the first three are the most important ones):

- Normal mode
- When vi starts, it is automatically in this mode. In normal mode, vi can be given commands. E. g. the command "i" (entered by pressing **I**) puts it into insert mode and the command ":" switches it to command-line mode.
- Sometimes the normal mode is also called "Command mode".

Insert mode

- In this mode, vi accepts all input as text. You can enter the insert mode for example by pressing **I** or **A** in the normal mode. Return to normal mode with **Esc**.

A special variant of the insert mode is the “replace mode”, where the entered text replaces the previous text. You can switch to replace mode with **Shift**+**R**.

```
quis nostrud exerci tation
ullamcorper suscipit lobortis
nisl ut aliquip ex en commodo
Duis te feugifacilisi per
```

```
-- REPLACE --
```

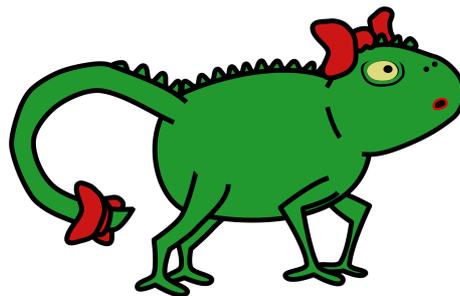
vi indicates the replace mode.

- Command-line mode – In this mode, vi accepts commands from the command line. You can enter the command-line mode for example by pressing **:**, **/** or **?**. Pressing **Enter** causes the command to be executed and automatically returns to the normal mode. **Esc** aborts the command-line mode and returns to normal mode.

```
ullamcorper suscipit lobortis
nisl ut aliquip ex en commodo c
Duis te feugifacilisi per susci
```

```
:
```

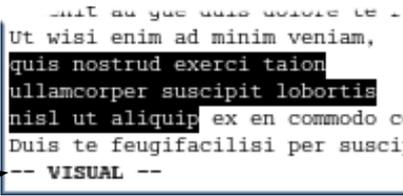
vi indicates the command-line mode.



Visual mode

- In this mode you can select text for editing, copying and so on. You can enter the visual mode for

example by pressing **V** in the normal mode. You can return to normal mode by pressing **Esc**.



A screenshot of the vi editor in visual mode. The text on the screen is:

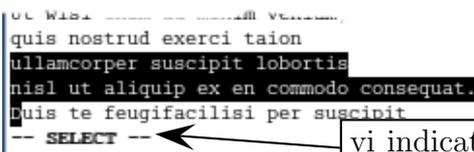

```

  ...nit au que quis doiore te -
  Ut wisi enim ad minim veniam,
  quis nostrud exerci taion
  ullamcorper suscipit lobortis
  nisl ut aliquip ex en commodo c
  Duis te feugifacilisi per susci
  -- VISUAL --
  
```

 A black rectangular selection box highlights the text from the second line to the fourth line. Below the text, the prompt "-- VISUAL --" is displayed. A callout box with an arrow points to this prompt, containing the text "vi indicates the visual mode."

Select mode

- The select mode is similar to the visual mode, but there are different commands available. You can switch to the select mode by entering **gh** in the normal mode. You can return to normal mode by pressing **Esc**.



A screenshot of the vi editor in select mode. The text on the screen is:

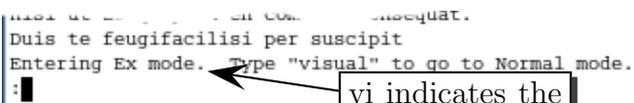

```

  ...ut wisi enim ad minim veniam,
  quis nostrud exerci taion
  ullamcorper suscipit lobortis
  nisl ut aliquip ex en commodo consequat.
  Duis te feugifacilisi per suscipit
  -- SELECT --
  
```

 A black rectangular selection box highlights the text from the second line to the fourth line. Below the text, the prompt "-- SELECT --" is displayed. A callout box with an arrow points to this prompt, containing the text "vi indicates the select mode."

Ex mode

- In the ex mode, vi does not work screen oriented. That means you do not have any cursor for positioning. You have to use the commands of the command-line mode to edit the text. From the normal mode, you can enter the ex mode by pressing **Shift+Q**. To exit the ex mode, you have to enter **vi** or **visual**.



A screenshot of the vi editor in ex mode. The text on the screen is:


```

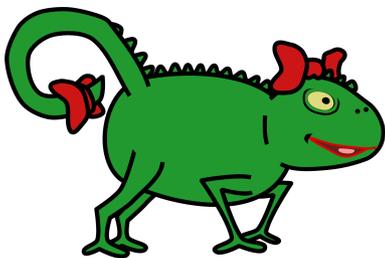
  ...ut wisi enim ad minim veniam,
  Duis te feugifacilisi per suscipit
  Entering Ex mode. Type "visual" to go to Normal mode.
  :
  
```

 A cursor is positioned at the end of the prompt ":". A callout box with an arrow points to this prompt, containing the text "vi indicates the ex mode mode."

You can use the following commands in normal mode:

Command	Result
i	Switches vi to insert mode.
a	Switches vi to insert mode (similar to i, but the cursor moves one step to the right).
A	Switches vi to insert mode (similar to i, but the cursor moves to the end of the current line).
x	Deletes the character where the cursor is.
dd	Deletes the line in which the cursor is located and copies it to the buffer.
D	Deletes the rest of the current line from the cursor position.
yy	Copies the line in which the cursor is located to the buffer.
p, P	Inserts the contents of the buffer after/before current cursor position.
o, O	Inserts a new blank line below/above the current line.
rcharacter	Replaces the current character by <i>character</i> .
J	Removes the line break at the end of the current line.
~	Changes between upper- and lower-case.
ZZ	Saves the current file and ends vi.
u	Undoes the last operation.
. (Full stop)	Repeats the last operation.
n	Repeats the search in the same direction.
N	Repeats the search in the opposite direction.

If you want to use a command for several units, place the corresponding number in front of the command. For example, `3x` deletes three characters, `5dd` deletes five lines, and `7yy` copies seven lines to the buffer.



You can use the following commands in command-line mode:

Command	Result
:q	Ends vi (if no changes were made).
:q!	Ends vi without saving changes in the file.
:wq or :x	Saves the current file and ends vi.
:w	Saves the current file.
:w <i>file</i>	Saves the current file under the name <i>file</i> . (Note: You continue editing the original file, not the new file.)
:e <i>file</i>	Opens the file <i>file</i> (if no changes were made at the current file).
:e <i>file</i>	Closes the current file without saving and opens the file <i>file</i> .
:ha	Prints the current file.
/ <i>pattern</i>	Searches forward from the cursor position for <i>pattern</i> .
? <i>pattern</i>	Searches backward from the cursor position for <i>pattern</i> .

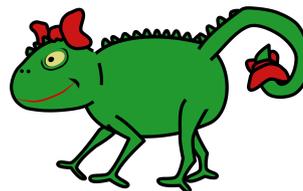


Exercise: Learn the Working Modes

Use vi to enter the following poetry of William Shakespeare. Save the file as Poetry.txt in your home directory.

```

1 Some glory in their birth, some in their skill,
2 Some in their wealth, some in their bodies' force,
  Some in their garments, though new-fangled ill,
4 Some in their hawks and hounds, some in their horse;
  And every humour hath his adjunct pleasure,
6 Wherein it finds a joy above the rest:
  But these particulars are not my measure;
8 All these I better in one general best.
  Thy love is better than high birth to me,
10 Richer than wealth, prouder than garments' cost,
  Of more delight than hawks or horses be;
12 And having thee, of all men's pride I boast:
  Wretched in this alone, that thou mayst take
14 All this away and me most wretched make.
```



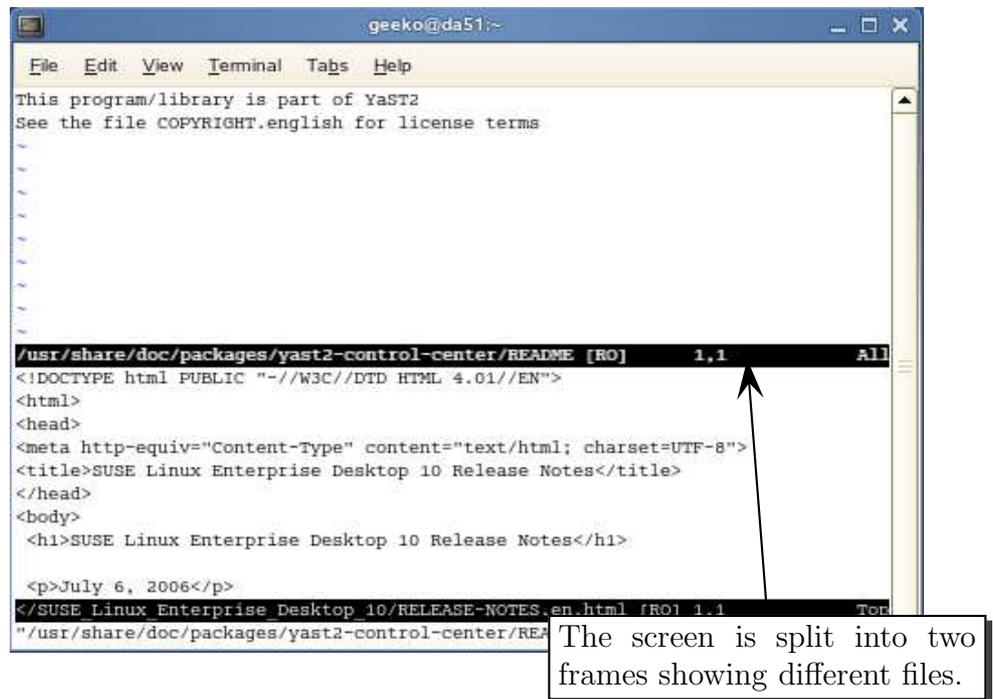
2.4.4 Learn Advanced vi Commands

In the following some more advanced vi features are introduced. They are interesting especially for developers.

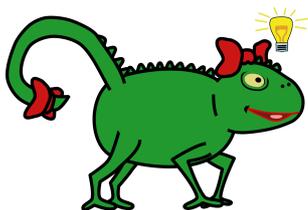
Split the Screen

Using the command `:sp filename`, the vi screen is split in two frames, and the specified file is shown in upper frame.

You can switch between the two frames by typing `Ctrl+W`, `W` in the command-line mode.



With `Ctrl+W`, `_` the current frame is maximized. `Ctrl+W`, `=` resizes all frames to the same size.



Auto Completion

In insert mode you can complete a word by pressing `Ctrl`+`P`. vi searches for words starting in the same way in the text and offers them to complete the current word. By pressing the key combination several times, you can switch between all alternatives.

Auto completion also works for lines by pressing `Ctrl`+`X`, `Ctrl`+`L`.

Use Marks

You can set a mark in a text by moving the cursor to the position you want to mark and entering `mcharacter`. For `character` you can use a character of a – z or A – Z.

If you want to jump to the mark later, enter `'character`.

Enable Syntax Highlighting

When entering `:syntax on` in command-line mode, syntax highlighting is enabled. vi knows a lot of programming languages and writes a computer program in multiple colors. This simplifies reading the program code.

```

geeko@da51:~
File Edit View Terminal Tabs Help
!DOCTYPE html PUBLIC "-//W3C//DTD HTML 4.01//EN">
<html>
<head>
<meta http-equiv="Content-Type" content="text/html; charset=UTF-8">
<title>SUSE Linux Enterprise Desktop 10 Release Notes</title>
</head>
<body>
<h1>SUSE Linux Enterprise Desktop 10 Release Notes</h1>

<p>July 6, 2006</p>

<p>For the latest version of the SUSE Linux Enterprise Desktop 10 Release Notes
, see http://www.novell.com/documentation/sled10/readme/RELEASE-NOTES.en.html.</
p>
<div>
<dl>
<dt>1.0 Xgl</dt>
<dd><ol>
<li><a href="#01">1.1 Registering SLED 10 to Enable Xgl</a></li>
<li><a href="#02">1.2 New Xgl Configuration Option</a></li>
<li><a href="#03">1.3 Semi-transparent Application Windows with Xgl</a></li>
<li><a href="#04">1.4 Using Full Screen Mode on VMware Workstations with Xgl</a>
</li>
</ol>
</dd>
</dl>
</div>
:syntax on                               1,1                               Top

```

If you want to enable syntax highlighting permanent, add the following to the file `~/.vimrc`.

```
1 syntax on
```

By default, this file does not exist.

To disable syntax highlighting, enter `:syntax off` in command-line mode.

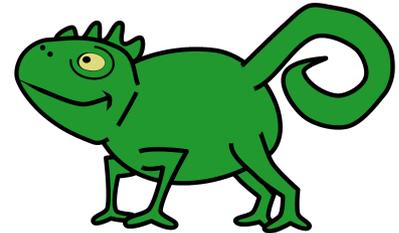
i There is a tutorial for vi available on your Linux computer. Enter `vimtutor` in the shell command line to start it.

i In the directory `/usr/share/doc/packages/vim/` there are two example configuration files `vimrc_example1` and `vimrc_example2`. Both include a lot of explanations and comments. These lines begin with double quotes.



Exercise: Enable Syntax Highlighting

Enable vi's syntax highlighting permanent on your computer.



2.4.5 Create Macros

Some tasks require a lot of redundant working steps. Suppose there is a file with last names and firstnames:

```
1 George
2 William
   Paul
4
   Smith
6 Carter
   Bailey
```

Now you want to append the first name behind the last name. If you have file with a thousands of names (and not just three), a macro helps you to fulfill this task.

Do the following steps:

- ❶ Move the cursor to the very beginning of the file (upper left corner).
- ❷ Enter `[Q]`, followed by any other characters to start recording a macro. The second character would be the name of the macro.

```

44-
ullamcorper suscipit lob
nisl ut aliquip ex en con
Duis te feugifacilisi pe
recording

```

vi is recording
a macro.

- ❸ Enter the commands that are necessary for your task. Referring to the upper example, the steps could be like the following:
 - ❶ `[D]`, `[D]` (Deletes the first line and copies the name to the clipboard.)
 - ❷ `[3]`, `[J]` (Moves the cursor one line down.)
 - ❸ `[P]` (Inserts the copied name under the current line.)
 - ❹ `[K]` (Moves the cursor one line up.)
 - ❺ `[Shift]+[A]` (Moves the cursor to the end of the line and switches into the insert mode.)
 - ❻ `[,]` (Enter a comma.)
 - ❼ `[Esc]` (Leave the insert mode.)
 - ❽ `[Shift]+[J]` (Appends the following line to the current one.)
 - ❾ `[G]`, `[G]` (Moves the cursor to the beginning of the file.)
- ❹ Finish macro recording by entering `[Q]`.

Now you can execute the macro by entering `[@]` followed by the “name” of the macro.

i For more detailed information about macros, enter `:help recording`.



Exercise: Create Macros

Create a macro, that inserts three spaces at the beginning of a line.

2.4.6 Create Abbreviations

An abbreviation works only in the insert or command-line mode. They have to be defined in command-line mode.

The general syntax to define an abbreviation is:

abbreviation_command short_form long_form

The *abbreviation_command* depends on the mode the abbreviation should be valid for:

iabbrev – Abbreviation is valid only for insert mode.

cabbrev – Abbreviation is valid only for command-line mode.

abbrev – Abbreviation is valid for insert and command-line mode.

The *short_form* is the text you want to enter. *long_form* is the text that should be inserted instead.

For example, after defining the following abbreviation,

```
1 iabbrev asap as soon as possible
```

any time, you enter *asap* in insert mode, *asap* is converted to *as soon as possible*.

It is also possible to use control sequences in the *short_form* and *long_form*:

Meaning	Enter
Space	<Space>
Enter	<CR>
Tab	<Tab>
Del	
Pos1	<Home>
End	<End>
Up-arrow	<Up>
Down-arrow	<Down>
Left-arrow	<Left>
Right-arrow	<Right>

For example, after defining the following abbreviation

```
1 iabbrev siny Sincerly yours<CR>Geeko
```

any time, you enter `siny` in insert mode, it is converted to

```
1 Sincerly yours
2 Geeko
```

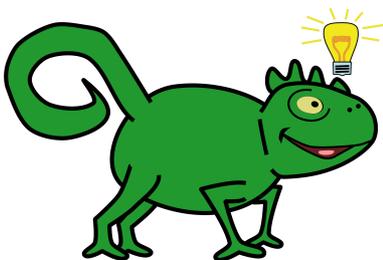


For more detailed information about macros, enter `:help abbreviations`.



Exercise: Create Abbreviations

Create an abbreviation that adds `#!/bin/bash` followed by two line breaks.



2.5 The Basics of bash Programming

As the Linux shell is programmable, you can use the `vi` text editor to write a program. Normally a bash script is written in a text file and

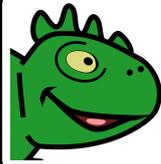
begins with the following line (called a *shee-bang line*):

```
#!/bin/bash
```

Before using a new bash script, you have to make it executable with the command `chmod u+x bash_script`).

Inside the bash script, you can use all the commands you learned before, such as `echo` to display text on the screen or request the content of a variable (see section 2.2.3 on page 78).

```
1 #!/bin/bash
2 echo 'The content of variable $PATH is:' $PATH
```



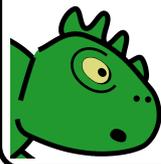
Normally the hash sign (“#”) is used to mark the following as comment. Comments are not executed. The only exception from this rule is the shee-bang line.

If you execute the script, the output may look like this:

```
1 geeko@da51:~> ./bash_script
2 The content of variable $PATH is: /home/geeko/bin:/usr/local/bin:/usr/bin:/usr/X
3 11R6/bin:/bin:/usr/games:/opt/gnome/bin:/opt/kde3/bin:/usr/lib/jvm/jre/bin:/usr/
4 lib/mit/bin:/usr/lib/mit/sbin
geeko@da51:~>
```

2.5.1 Variables in bash Scripts

In the example above, we use the environment variable `$PATH`. It is possible to define your own variables that are only valid inside the bash script. The syntax is *variable_name=value*.



It is a convention to use lowercase letters for variable names in a program code. Only constants (which do not change their value during the program’s run time) are written in uppercase letters.

```

1 #!/bin/bash
2
4 myvar=13
4 echo 'The content of variable $myvar is:' $myvar

```

To get the content of a variable (e. g., with the `echo` command), you have to add a “\$” in front of the variable’s name (see section 2.1.8 on page 61).

A variable can store numbers, characters, and text strings.

```

1 #!/bin/bash
2
4 a=2
4 b=x
4 c=Geeko
6 d="Geeko and Suzie"
8
8 echo '$a is a number:' $a
8 echo '$b is a character:' $b
10 echo '$c is a string:' $c
10 echo '$d is a string with whitespaces:' $d

```

The output looks like the following:

```

1 $a is a number: 2
2 $b is a character: x
4 $c is a string: Geeko
4 $d is a string with whitespaces: Geeko and Suzie

```



Exercise: Types of Variables

Numbers can be added. Strings can’t. Make aware that there are differences between numbers and strings in the kind which operation is meaningful.

Write down four operations that are typical for numbers and four that are typical for strings.

Operations for numbers:

- ① _____
- ② _____
- ③ _____
- ④ _____

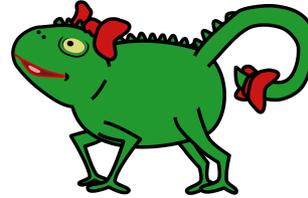
Operations for strings:

- ① _____
- ② _____
- ③ _____
- ④ _____

2.5.2 Compute Numbers

Variables are often combined with mathematical operators. That means that you can compute them. Variables with numbers can be:

- ☞ Added (use operator “+”)
- ☞ Subtracted (use operator “-”)
- ☞ Multiplied (use operator “*”)
- ☞ Divided (use operator “/”)
- ☞ Modulated (use operator “%”)



Compute Integers

In the following example, we write the sum of the variables \$a and \$b into a new variable.

```
1 #!/bin/bash
2
3 a=3
4 b=2
5 c=$a+$b
6
7 echo '$a+$b=$c'
```

The output looks like the following.

```
1 $a+$b=3+2
```

As you can see, the variable C contains the string “3+2,” but not the sum “5”. To get the sum, you have the following possibilities:

- ① Allow the variable \$c only to store numbers.

To do this, use the command `declare` with the option “-i” to declare the variable \$c.

```
1 #!/bin/bash
2
3 a=3
4 b=2
5 declare -i c=$a+$b
6
7 echo '$a+$b=$c'
```

- ② Mark the formula as expression that has to be evaluated.

To calculate at the command line, use the command `expr`.

```
1 geeko@da51:~ > expr 3 + 2
2 5
geeko@da51:~ >
```



Attention! In contrast to the method using `declare`, with `expr` there has to be a space before and after the operator(s).

If you use `expr` in a shell script, you have to quote the expression with backticks (“`”) or `$(...)`. Characters that have a special meaning in the bash (such as “*” and brackets) have to be masked by a backslash.

```
1 #!/bin/bash
2
3 a=3
4 b=2
5 c=$(expr $a + $b)
6
7 echo '$a+$b=$c'
```

The expression inside `$(...)` is replaced by the result of the expression.

The output looks like this:

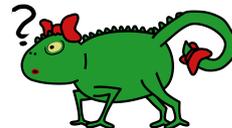
```
1 $a+$b=5
```

Now let us do the same thing with the division operator.

```
1 #!/bin/bash
2
3 a=3
4 b=2
5 c=$(expr $a / $b)
6
7 echo '$a/$b=$c'
```

The result is not what you might expect:

```
1 $a/$b=1
```



The results of a division is an integer, and an integer number does not have any decimal places. You can get the rest of a division with the modulo operator:

```
1 #!/bin/bash
2
3 a=3
4 b=2
5 c=$(expr $a / $b)
6 d=$(expr $a % $b)
7
8 echo '$a/$b=' '$c Rest:$d'
```

Instead of `expr` you also can use the following syntax `$((...))`. For example:

```
1 #!/bin/bash
2
3 a=3
4 b=2
5 c=$(( $a + $b ))
6
7 echo '$a+$b=' '$c'
```



Exercise: Compute Integers

Write a bash program that computes the squares of 6, 9, and 12.

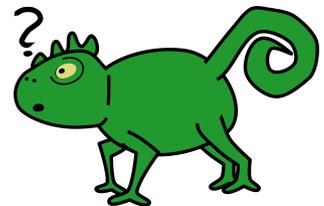
Compute Float Numbers

By default, the bash is not able to calculate with float numbers.

```
1 #!/bin/bash
2
3 a=3
4 b=2.7
5 c=$(expr $a + $b)
6
7 echo '$a+$b=' '$c'
```

This program leads to the result:

```
1 expr: non-numeric argument
2 $a+$b=
```



To solve this problem, use the shell tool called `bc`. `bc` is an arbitrary precision calculator. At the command line, you can use `bc` for calculations.

```

1 geeko@da51:~> bc
2 bc 1.06
  Copyright 1991-1994, 1997, 1998, 2000 Free Software Foundation, Inc.
4 This is free software with ABSOLUTELY NO WARRANTY.
  For details type 'warranty'.
6 3+4
  7
8 (10+4)/2
  7
10 2+0.44
  2.44
12 7/3
  2
14 quit

```

Explanation: When entering `bc`, the calculator starts, and you can enter an expression to calculate (lines 6, 8, 10 and 12). To see the result of the calculation, press `Enter` (lines 7, 9, 11, and 13). You can group expressions by using brackets (line 8). You can use `bc` to calculate with float numbers (line 10). By default, `bc` uses the same kind of processing for division (without decimal places) as the shell does (lines 12+13). To exit `bc`, enter `quit` (line 14).

`bc` provides its own programming language, but that would not matter now. To use `bc` within a bash script, use a simple “trick”: piping the formular into `bc`. Use the `echo` command and the “|” pipe.

```

1 geeko@da51:~> echo "2.4+7" | bc
2 9.4

```

To store the result of an calculation using `bc`, put it into backticks (“”).

```

1 geeko@da51:~> a=`echo "2.4+7" | bc`
2 geeko@da51:~> echo $a
  9.4

```

Instead of backticks, you can also use the following syntax with `$(...)`.

```

1 geeko@da51:~> a=$(echo "2.4+7" | bc)
2 geeko@da51:~> echo $a
  9.4

```

It’s up to you, which syntax you prefer. In this manual, the second one is used.

To see decimal places, enter `scale=number` in front of the mathematical operation, separated by semicolon.

For example:

```

1 geeko@da51:~> echo "scale=2; 9/2" | bc
2 4.50

```



Exercise: Compute Float Numbers

Write a bash program that computes the squares of 1.5, 2.75 and 7.2

2.5.3 Comparing Files, Strings and Numbers

To find out if three is greater than two, you do not need to write a bash program. But in interactive programs a lot of questions appear and the program execution depends on the results. Common questions are, for example:

- ☞ Does the file *filename* exist?
- ☞ How old is the file?
- ☞ How does the user answer the program's question?
- ☞ Is the value of a variable still inside the valid range?
- ☞ How often has a loop already been completed?

Use the `test` command for testing. The result of a test is either “0” (true) or “1” (false). This result is stored in the variable “\$?”.

The `test` command knows a lot of parameters. First we want to introduce the parameters concerning files:

- ☞ `-d target`

Tests if *target* is an existing directory.

```
1 #!/bin/bash
2
3 a=/var
4 b=/etc/profile
5
6 test -d $a
7 echo $a: $?
8
9 test -d $b
10 echo $b: $?
```

The output of this program looks like the following:

```
1 /var: 0
2 /etc/profile: 1
```

☞ `-e target`

Tests if *target* is an existing file.

```
1 #!/bin/bash
2
3 a=/var
4 b=2
5
6 test -e $a
7 echo $a: $?
8
9 test -e $B
10 echo $b: $?
```

The output looks like the following:

```
1 /var: 0
2 2: 1
```

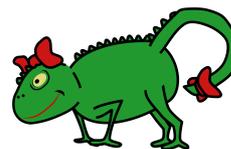
☞ `-f target`

Tests if *target* exists and if it is a regular file.

```
1 #!/bin/bash
2
3 a=/var
4 b=/etc/profile
5
6 test -f $a
7 echo $a: $?
8
9 test -f $b
10 echo $b: $?
```

The output of this program looks like the following. Compare it with the output of parameter `-d`:

```
1 /var: 1
2 /etc/profile: 0
```



☞ `-r target, -w target, -x target`

Tests if the file permission for the current user is set:

- ☞ reading (`-r`)
- ☞ writing (`-w`)
- ☞ executing (`-x`)

```

1 #!/bin/bash
2
3 a=/etc/profile
4
5 test -r $a
6 echo "Read: $?"
7
8 test -w $a
9 echo "Write: $?"
10
11 test -x $a
12 echo "Execute: $?"

```

The file `/etc/profile` has the following permission: `rw-r--r--`. User `Geeko` is only allowed to read the file. The output of the above script looks like the following:

```

1 Read: 0
2 Write: 1
3 Execute: 1

```

👉 `-s target`

Tests if a file is empty.

```

1 #!/bin/bash
2
3 a=/etc/profile
4
5 test -s $a
6 echo $?

```

The output of the script should be “0”.

👉 `file1 -nt file2`

Tests if `file1` is newer than `file2`. The date of the last modification is used, not the date of the creation.

```

1 #!/bin/bash
2
3 a=/etc/profile
4 b=/var/log/Xorg.0.log
5
6 test $a -nt $b
7 echo "$a -nt $b:" $?

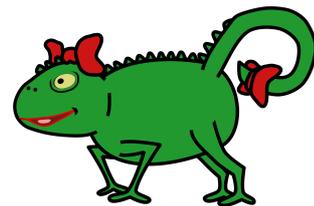
```

The result should be “1” (false), and the output looks like the following:

```

1 /etc/profile -nt /var/log/Xorg.0.log: 1

```



 `file1 -ot file2`

Is the counterpart of `-nt` and tests if `file1` is older than `file2`.

```
1 #!/bin/bash
2
4 a=/etc/profile
4 b=/var/log/Xorg.0.log
6 test $a -ot $b
6 echo "$a -ot $b:" $?
```

The result should be “1” (false), and the output looks like the following:

```
1 /etc/profile -ot /var/log/Xorg.0.log: 0
```



Exercise: Comparing Files, Strings and Numbers – Part I

Write a bash program that...

- ... tests if the file `/var/log/messages` exists,
- ... if it is empty and
- ... that displays the file permission.

`test` is also able to test and compare strings.

 `-z string`

The result is “0” (true), if the string is “0”.

```
1 #!/bin/bash
2
4 a=geeko
6 test -z $a
6 echo "a: $?"
8 test -z $b
8 echo "b: $?"
```

The output looks like the following:

```
1 a: 1
2 b: 0
```

Explanation: As the variable `$b` is not defined, it has a length of “0” and the test result is “true”.

 `-n string`

This parameter is the counterpart of `-z`. The test result is “true” if the string contains one or more characters.

```

1 #!/bin/bash
2
3 a=geeko
4
5 test -n $a
6 echo "a: $?"
7
8 test -n $b
9 echo "b: $?"

```

The output looks like the following:

```

1 a: 0
2 b: 1

```

☞ `string1 == string2`

The test result is “true” if *string1* and *string2* are the same.

```

1 #!/bin/bash
2
3 a=geeko
4 b=suzie
5 c=geeko
6
7 test $a == $b
8 echo "$a == $b: $?"
9
10 test $a == $c
11 echo "$a == $c: $?"

```

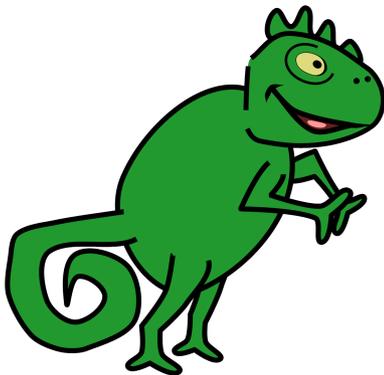
Explanation: Three variables are defined (lines 3–5). The variables `$a` and `$c` are the same. In line 7, the variables `$a` and `$b` are compared. In line 10, the variables `$a` and `$c` are compared.

The output looks like the following:

```

1 geeko == suzie: 1
2 geeko == geeko: 0

```



☞ `string1 != string2`

This operator is the counterpart of the “`==`” operator. The test result is “true” if both strings are *not* the same.

```

1 #!/bin/bash
2
3 a=geeko
4 b=suzie
5 c=geeko
6
7 test $a != $b
8 echo "$a != $b: $?"
9
10 test $a != $c
11 echo "$a != $c: $?"

```

The output looks like the following:

```

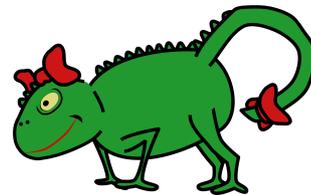
1 geeko != suzie: 0
2 geeko != geeko: 1

```



Exercise: Comparing Files, Strings and Numbers – Part II

Write a bash program that tests whether the string “Linux” and the string “School” are equal.



There are also some operators which compare numbers:

 `integer1 -eq integer2`

Tests if *integer1* is equal to *integer2*.

```

1 #!/bin/bash
2
3 a=3
4 b=2
5
6 test $a -eq $b
7 echo "$a -eq $b: $?"

```

The output looks like the following:

```

1 3 -eq 2: 1

```

 `integer1 -ne integer2`

Is the counterpart of `-eq`. It testes if *integer1* is *not* the same as *integer2*.

```

1 #!/bin/bash
2
3 a=3
4 b=2
5
6 test $a -ne $b
7 echo "$a -ne $b: $?"

```

The output looks like the following:

```

1 3 -ne 2: 0

```

☞ `integer1 -gt integer2`

Tests if *integer1* is strictly greater than *integer2*.

```

1 #!/bin/bash
2
3 a=3
4 b=2
5
6 test $a -gt $b
7 echo "$a -gt $b: $?"
8
9 test $b -gt $a
10 echo "$b -gt $a: $?"

```

The output looks like the following:

```

1 3 -gt 2: 0
2 2 -gt 3: 1

```

☞ `integer1 -ge integer2`

Tests if *integer1* is greater than or the same as to *integer2*.

```

1 #!/bin/bash
2
3 a=3
4 b=2
5 c=3
6
7 test $a -ge $b
8 echo "$a -ge $b: $?"
9
10 test $a -ge $c
11 echo "$a -ge $c: $?"

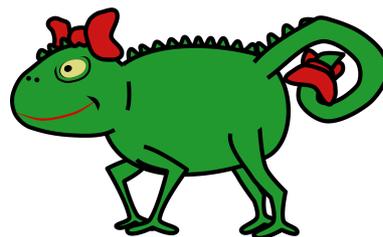
```

The output looks like the following:

```

1 3 -ge 2: 0
2 3 -ge 3: 0

```



 `integer1 -lt integer2`

Tests if *integer1* is strictly less than *integer2*.

```

1 #!/bin/bash
2
3 a=3
4 b=2
5
6 test $a -lt $b
7 echo "$a -lt $b: $?"
8
9 test $b -lt $a
10 echo "$b -lt $a: $?"

```

The output looks like the following:

```

1 3 -lt 2: 1
2 2 -lt 3: 0

```

 `integer1 -le integer2`

Tests if *integer1* is less than or equal to *integer2*.

```

1 #!/bin/bash
2
3 a=3
4 b=2
5 c=3
6
7 test $a -le $b
8 echo "$a -le $b: $?"
9
10 test $a -le $c
11 echo "$a -le $c: $?"

```

The output looks like the following:

```

1 3 -le 2: 1
2 3 -le 3: 0

```

Exercise: Comparing Files, Strings and Numbers – Part III

Write a bash program that tests whether the number “13” is greater than the number “11”.

It is also possible to compare float numbers. To do this, you have to use the `bc` command. The syntax differs from the shell syntax:

 `expression1 == expression2`

Is *expression1* the same as *expression2*?

 `expression1 != expression2`

Is *expression1* not the same as *expression2*?

☞ `expression1 > expression2`

Is `expression1` strictly greater than `expression2`?

☞ `expression1 >= expression2`

Is `expression1` greater than or equal to `expression2`?

☞ `expression1 < expression2`

Is `expression1` strictly less than `expression2`?

☞ `expression1 <= expression2`

Is `expression1` less than or equal to `expression2`?



All these operators are shown in the following program:

```

1 #!/bin/bash
2
3 a=3.7
4
5 b=$(echo "$a == 3.7" | bc)
6 echo "$a == 3.7: $b"
7
8 b=$(echo "$a != 3.7" | bc)
9 echo "$a != 3.7: $b"
10
11 b=$(echo "$a > 3.7" | bc)
12 echo "$a > 3.7: $b"
13
14 b=$(echo "$a >= 3.7" | bc)
15 echo "$a >= 3.7: $b"
16
17 b=$(echo "$a < 3.7" | bc)
18 echo "$a < 3.7: $b"
19
20 b=$(echo "$a <= 3.7" | bc)
21 echo "$a <= 3.7: $b"

```

The output looks like the following:

```

1 3.7 == 3.7: 1
2 3.7 != 3.7: 0
3 3.7 > 3.7: 0
4 3.7 >= 3.7: 1
5 3.7 < 3.7: 0
6 3.7 <= 3.7: 1

```



Attention! The return values of `bc` are different to the bash: “0” means the relation is false and “1” means the relation is true.



Exercise: Comparing Files, Strings and Numbers – Part IV

Write a bash program that tests whether the number “13.5” is greater than the number “22”.

Sometimes it is necessary to combine two or more tests, e. g., to test if the value of a variable is greater than a number *but* less than another number. There are two combinations possible:

Logical OR. Test one returns “0” (true) *or* test two returns “0” (true).

		Test1	
		true	false
Test2	true	true	true
	false	true	false

Use the parameter `-o` to create an OR connection between two tests.

```

1 #!/bin/bash
2
3 a=/var
4 b=/etc/profile
5
6 test -f $a -o -f $b
7 echo "Result: $?"

```

Explanation: This script tests whether the file `/var`, *or* `/etc/profile` is a regular file. For `/var` the result is false, because it is a directory. But `/etc/profile` is a regular file. This is why the whole test returns “0” (true).

Logical AND. Test one returns “0” (true) *and* test two returns “0” (true).

		Test1	
		true	false
Test2	true	true	false
	false	false	false

Use the parameter `-a` to create an AND connection between tests.

```

1 #!/bin/bash
2
3 a=7
4 b=2
5
6 test $a -gt $b -a $(expr $a % 2) -eq 0
  echo "Result: $?"

```

Explanation: This script tests if the value of `$a` is greater than the value of `$b` ($7 > 2$; true) *and* if the value of `$a` is even (false). Because of the last comparison is false, the whole test is false.

The next logical operator we want to introduce is NOT. This operator is called “negation” operator. It is written by an exclamation mark.

```

1 #!/bin/bash
2
3 a=7
4 b=2
5
6 test ! $a -gt $b
  echo "Result: $?"

```

Explanation: `$a -gt $b` tests if the value of `$a` is greater than the value of `$b`. In the example above, that is true. But, because of the negation, the “true” is changed to “false”.



In `bc` the logical AND is written by “`&&`” and logical OR by “`||`”. For NOT, use the exclamation mark, too.



Exercise: Comparing Files, Strings and Numbers – Part V

Write a bash program that tests whether the number “13.5” is greater than the number “22” or smaller than “15”.

2.5.4 Read from Input

It does not make sense to write a bash program with eight lines just to compare two numbers. And if you need a comparison of two other numbers, you have to write a new program?

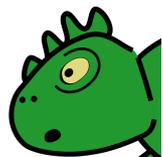
Of course, it is not true, and you can ask the program’s user to enter some data. Therefore the command `read` is used.

`read` waits until the user finishes his input by pressing `Enter`. You can specify a variable where the input is stored.

```
1 #!/bin/bash
2
3 read a
4 echo "You entered: $a"
```

Explanation: After this small program starts, it waits for input. The input will be stored in the variable `$a` (line 3). After the user finishes his input by pressing `Enter`, the input is written on the screen (line 4).

If you do not specify a variable after the `read` command, the input is stored in the variable `$REPLY`.



You can also use `read` to specify more than one variable. The input is separated by white spaces into tokens. The first token is stored in the first variable, the second token in the second variable, and so forth. All remaining text tokens are stored in the last variable.

```
1 #!/bin/bash
2
3 read a b c
4 echo "First token: $a"
5 echo "Second token: $b"
6 echo "Third token: $c"
```

If you enter `This is a short text.`, the output looks like the following:

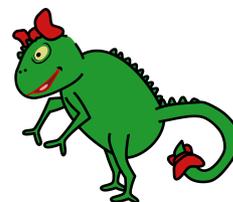
```
1 First token: This
2 Second token: is
   Third token: a short text.
```

If there are more variables than text tokens, the remaining variables stay empty.



Exercise: Read from Input

Write a bash program that reads two integer values from input and tests if the first number is larger than the second one.



2.5.5 Examine Cases

The bash supports two possibilities to control the program flow:

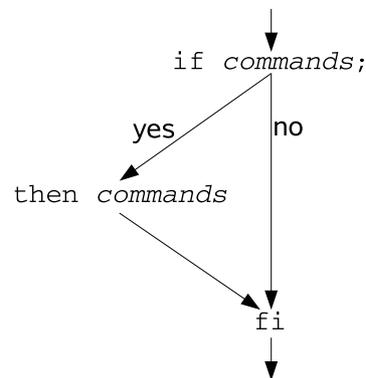
☞ `if ...fi`

☞ `case ...esac`

If you want to differentiate between a few cases, use the first possibility.

The most simple way to use the `if` statement is the following syntax:

```
if commands; then
    commands
fi
```



First, the command after `if` is evaluated. If the result is “0” (true), the commands after the corresponding `then` are executed. If the result of the `if` command is “1” (false), nothing happens and the program continues after `fi`.

In most cases, the commands after `if` is a test with the `test` command.

```

1 #!/bin/bash
2
3 echo "Enter something:"
4 read a
5
6 if test -e $a; then
7     echo "$a is an existing file."
8 fi
```

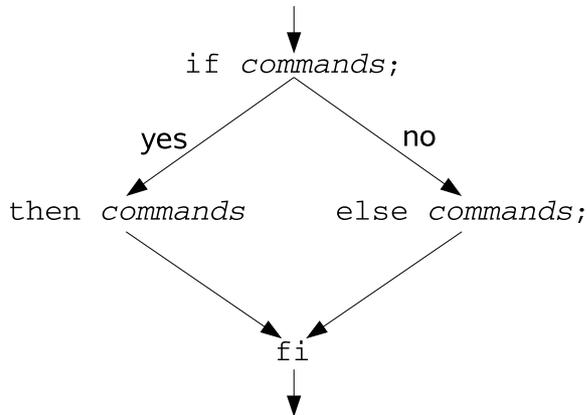
Explanation: This program asks the user to enter something (lines 3 – 4). Then it tests if the entered string is an existing file (line 6). If so, a message is written on the screen (line 7).

You can add `else` to the construct above:

```

if commands; then
    commands
else
    commands
fi

```



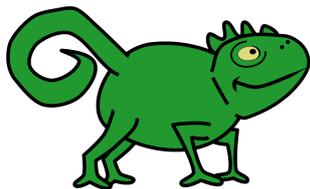
The commands after `else` are executed, if the result of the `if` command is false.

```

1 #!/bin/bash
2
3 echo "Enter something:"
4 read a
5
6 if test -e $a; then
7     echo "$a is an existing file."
8 else
9     echo "$a is no file."
10 fi

```

Explanation: The program above is extended by lines 8 and 9. Now, if the input is no filename, another message is written.

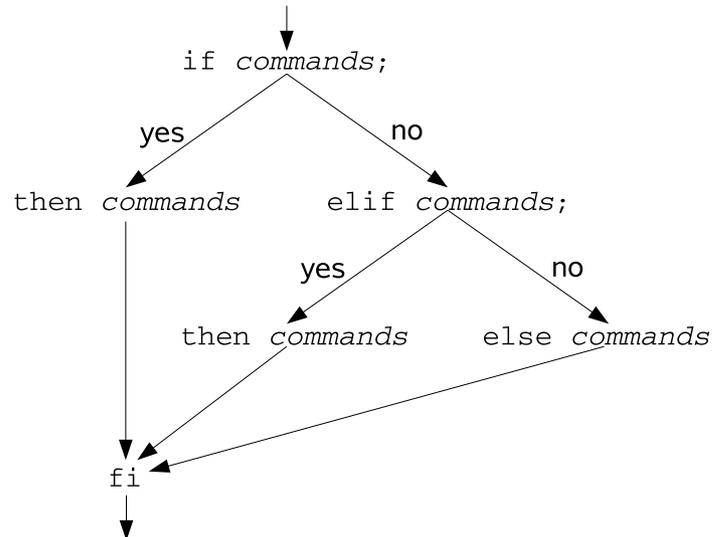


If you want to check for more than one condition, put `elif` between the `then` command and the `else` command.

```

if commands; then
    commands
elif commands; then
    commands
else commands
    commands
fi

```



First, the command after `if` is evaluated. If the result is “0” (true), the commands after the corresponding `then` are executed. If the result of the `if` command is false, then the commands after `elif` are evaluated, whether they are true or false. If neither the `if` command nor the `elif` command is true, the `else` commands are executed.

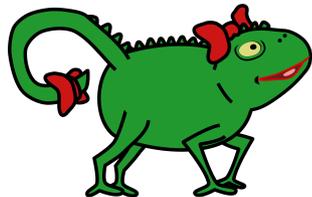
You can use more than one `elif` in a condition to build long chains of conditions.

```

1  #!/bin/bash
2
3  echo "Enter something:"
4  read a
5
6  if test -e $a; then
7      echo "$a is an existing file."
8  elif test $a == geeko; then
9      echo "Entered geeko."
10 else
11     echo "$a is no file and not geeko."
12 fi

```

Explanation: The above program is extended by lines 8 and 9. Now, the entered string is tested whether it is the same as “geeko”. It is tested, if the entered text is not an existing file.



Exercise: Examine Cases – Part I

Write a bash program that reads two integer values from input and tests ...

- ☞ ... whether the first number is greater than the second one,
- ☞ ... whether the first number is equal to the second one, or
- ☞ ... whether the first number is smaller than the second one.

It should display different messages on the screen for each case.

If you have a lot of `if` conditions, it is better to use `case` instead.

The syntax is:

```
case condition in
    pattern | pattern ...)
    commands
    ;;
    pattern | pattern ...)
    commands
    ;;
esac
```

The patterns are tested if they match the condition. If yes, the commands in the pattern are executed. If not, the next pattern is tested.

```
1 #!/bin/bash
2
3 echo "Enter something:"
4 read a
5
6 case $a in
7     [Yy]es|[Tt]rue)
8         echo "You entered Yes/True."
9         ;;
10    [Nn]o|[Ff]alse)
11        echo "You entered No/False."
12        ;;
13 esac
```

Explanation: The entered string is tested whether it is “Yes”, “yes”, “True”, or “true” (line 7). If it does not match, it is tested for “No”, “no”, “False”, or “false” (line 9).

If a pattern matches, and its commands are executed, the program finishes the `case-esac` region and continues after `esac`.

If you want to execute commands in case none of the patterns match, add an asterisk pattern:

```

1 #!/bin/bash
2
3 echo "Enter something:"
4 read a
5
6 case $a in
7     [Yy]es|[Tt]rue)
8         echo "You entered Yes/True."
9         ;;
10    [Nn]o|[Ff]alse)
11        echo "You entered No/False."
12        ;;
13    *)
14        echo "Wrong input."
15 esac

```



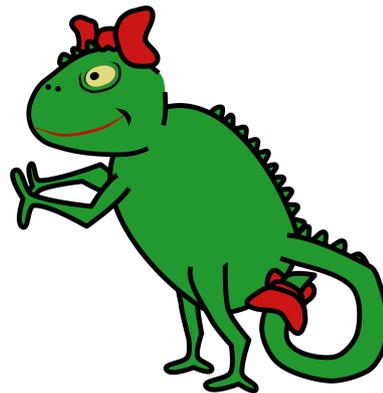
Exercise: Examine Cases – Part II

Write a bash program that reads one one-digit integer number from input and displays the name of the number on the screen (e. g., One for “1”).

2.5.6 Use Loops

By using a loop, you can make the interpreter execute a part of the program several times. The bash knows three loops:

- ❶ while
- ❷ until
- ❸ for



The while Loop

The syntax of the while loop is:

```
while condition; do
    commands
done
```

Normally, the condition is a test and the commands inside the loop are executed as long as the condition is true.

```
1 #!/bin/bash
2
3 a=1
4
5 while test $a -le 8; do
6     echo "\$a=$a"
7     a=$((expr $a + 1))
8 done
```

Explanation: In line 3, the variable `$a` is set to “1”. In line 5, it is tested if `$a` is less than (or equal to) “8”. While the value of `$a` is less than “8”, the commands inside the loop are executed. Line 6 creates some output. In line 7 the value of `$a` is increased by “1”.

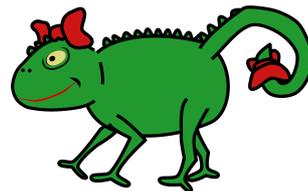
The output of the program above looks like the following:

```
1 $a=1
2 $a=2
3 $a=3
4 $a=4
5 $a=5
6 $a=6
7 $a=7
8 $a=8
```



Exercise: Use Loops – Part I

Write a bash program that reads one one-digit integer number from input and counts from one to the entered number. It should display the name of the number on the screen (e.g., One for “1”) and not the number. Use a `while` loop.



The until Loop

The `until` loop is similar to the `while` loop. The syntax is similar, too:

```
until condition; do
    commands
done
```

The difference between `while` and `until` is, that if you use `until`, the *commands* are executed *until* the condition is true.

To get the same result as in the program above, the code has to be:

```
1 #!/bin/bash
2
3 a=1
4
5 until test $a -gt 8; do
6     echo "\$a=$a"
7     a=$(expr $a + 1)
8 done
```

Explanation: `while` is replaced by `until`. “`-le`” is replaced by “`-gt`”.



Exercise: Use Loops – Part II

Write a bash program that reads one one-digit integer number from input and counts from one to the entered number. It should display the name of the number on the screen (e.g., One for “1”) and not the number. Use an `until` loop.



The for Loop

The `for` loop differs from `while` and `until`. It is used to repeat code several times. The number of repetition is fixed and defined before.

`for` can be used in two different ways. The simplest syntax is:

```
for variable in list; do
    commands
done
```

The variable is set to the first item of the list, and then the commands are executed. The next time, the variable is set to the second item of the list. The commands are executed again, and so on.

```

1 #!/bin/bash
2
3 for a in Geeko Suzie 100 0 Novell; do
4     echo "$a"
5 done

```

The output of this program looks like the following:

```

1 Geeko
2 Suzie
3 100
4 0
5 Novell

```

The list can be created while the program is running.

```

1 #!/bin/bash
2
3 for a in $(ls /bin/b*); do
4     echo "$a"
5 done

```

Explanation: All files that are in the /bin directory and starts with “b” are used as a list.

The output looks like the following:

```

1 /bin/basename
2 /bin/bash
3 /bin/bluepinat

```

Another syntax of for is:

```

for ((variable; condition; change)); do
    commands
done

```

The parts are the following:

variable – The start value of a variable is set.

condition – The loop is executed as long as the condition is true.

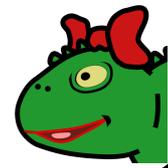
change – At the end of the loop, the value of the variable is changed in the defined way.

The following is an example:

```
1 #!/bin/bash
2
3 for ((a=10; a<=15; a++)); do
4     echo "$a"
5 done
```

Explanation: The first time line 3 is executed, variable `$a` is set to “10”. It tests, whether the condition (`$a` is less than or the same as “15”) is true. The command(s) of the loop (line 4) are executed once. After `done` (line 5), the value of variable `$a` is increased by “1”. It tests, whether the condition is true again.

If you append a double plus (“++”) at a variable in a `for` command, the value is increased by 1. If you append a double minus (“--”) at a variable, the value is decreased by 1.



The output of the program looks like the following:

```
1 10
2 11
3 12
4 13
5 14
6 15
```



It is not mandatory that you use the same variable in the three parts of the `for` command.



Exercise: Use Loops – Part III

Write a bash program that reads one one-digit integer number from input and counts from one to the entered number. It should display the name of the number on the screen (e.g., One for “1”) and not the number. Use a `for` loop.

Break a Loop

You can exit a loop before it is executed in the specified times. To do this you have to use the command **break**.

```

1 #!/bin/bash
2
3 a=1
4
5 while true; do
6     echo "$a"
7     a=$(expr $a + 1)
8     if test $a == 8; then
9         break
10    fi
11 done

```

Explanation: The variable `$a` is set to “1” in line 3. The `while` command (line 5) returns always “true” and the loop would be executed endless, if there was no `break` (line 9). If the value of `$a` is equal to “8”, the loop is interrupted.

The output looks like the following:

```

1 1
2 2
3 3
4 4
5 5
6 6
7 7

```



2.5.7 Use Program Options and Parameters

As described in the previous chapters, a command usually knows a lot of options and parameters. You can use them for your program, too.

The number of the passed options is stored in the variable “`$#`”.

```

1 #!/bin/bash
2
3 echo "$# options passed."

```

Now, let us start this short programm with some parameters:

```

1 geeko@da51:~> ./script.sh geeko suzie 100 Novell
2 4 options passed.

```

The command used to start the program is stored in the variable “\$0”, the first parameter in variable “\$1”, the second in “\$2”, and so on.

```
1 #!/bin/bash
2
3 echo "\$0=$0"
4 echo "\$1=$1"
5 echo "\$2=$2"
6 echo "\$3=$3"
```

When starting the program with some parameters, the output, for example, looks like the following:

```
1 geeko@da51:~> ./script.sh geeko suzie 100 Novell
2 $0=./script.sh
3 $1=geeko
4 $2=suzie
5 $3=100
```

This method is not very flexible, because you need to know the number of parameters. In the last example, the parameter `Novell` is not computed, because the program expects only three parameters. Additionally, you can only access the command plus nine options.

The following variation is used more often: Test the first option. Then, remove the first item of the list and the second becomes the first item. This process is done by the command `shift`.

Normally, you put this test into a loop. A simple example looks like the following:

```
1 #!/bin/bash
2
3 while test $# -gt 0; do
4     echo "Option: $1"
5     shift
6 done
```

Explanation: The `test` command tests whether the list of options still has some items (line 3). If so, the first item is printed on the screen (line 4) and removed from the list (line 5). The loop is repeated until all items are removed.

The output of the program above looks like the following:

```
1 geeko@da51:~> ./script.sh geeko suzie 100 Novell
2 Option: geeko
3 Option: suzie
4 Option: 100
5 Option: Novell
```

To remove more than one item from the list, use `shift number`.



Normally a program used to examine the options and parameters looks like the following:

```

1 #!/bin/bash
2
3 while test $# -gt 0; do
4     case $1 in
5         -a) echo "Option a"
6             shift
7             ;;
8         -b) echo "Option b"
9             shift
10            ;;
11        -f) echo "Option f with argument $2"
12            shift 2
13            ;;
14        *)  echo "Wrong option: $1"
15            break
16            ;;
17    esac
18 done

```

Let us do some tests:

①

```

1 geeko@da51:~> ./script.sh -a -f /home/geeko
2 Option a
3 Option f with argument /home/geeko

```

②

```

1 geeko@da51:~> ./script.sh -f /home/geeko -a
2 Option f with argument /home/geeko
3 Option a

```

③

```

1 geeko@da51:~> ./script.sh -b /home/geeko
2 Option b
3 Wrong option: /home/geeko

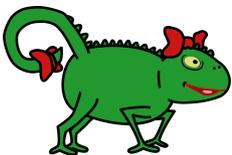
```

④

```

1 geeko@da51:~> ./script.sh /home/geeko -f
2 Wrong option: /home/geeko

```



Exercise: Use Program Options and Parameters

Describe what happens in the four examples above. Refer to the line numbers of the previous program code.

- ① _____

- ② _____

- ③ _____

- ④ _____

2.5.8 Create Shell Functions

Functions are parts of a program that are repeated often in a program at different locations. The program code is extracted from the main program into a function. The function is called from the *main program* when needed.

A function has the following has the following syntax:

```
name () { commands }
```



The function is called from the main program by its name. It is possible to add some parameters.

For example:

```
1 #!/bin/bash
2
3 sum() {
4     a=$(expr $1 + $2)
5     echo "$1+$2="$a
6 }
7
8 sum 4 5
9 sum 2 9
10 sum 12 1
11 sum 204 2387
12 sum 5 3
13 sum 21 39
```

Explanation: The main program begins in line 8. It only calls the function “sum” six times with different parameters.

The function (lines 3–6) adds the given parameters and writes the sum on the screen. To refer to the parameters, the variables “\$1” and “\$2” are used.

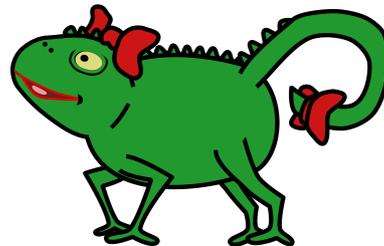
Referring to function parameters works the same way as referring to program parameters (see section 2.5.7 on page 143).



The output of the program above is:

```
1 4+5=9
2 2+9=11
3 12+1=13
4 204+2387=2591
5 5+3=8
6 21+39=60
```

With parameters, you can transfer information into the function. To transfer information back into the main program, you can use the command `return` followed by a number.



For example:

```

1  #!/bin/bash
2
3  validate() {
4      if [ $1 = yes ]; then
5          return 1
6      else
7          return 0
8      fi
9  }
10
11 a=1
12 while true; do
13     echo $a" -> To continue, enter \"yes\""
14     read b
15     if validate $b; then
16         break
17     fi
18     a=$(expr $a + 1)
19 done
20 echo "Stop!"

```

Explanation: The main program starts in line 11. It contains an endless loop (lines 12–19) and asks for some user input (line 14). The loop is only terminated, when the return value of the function `validate()` is “true” (equal to “0”).

The function gets the input of the user as command parameter. If the user entered `yes`, the function returns the value “1”. In all other cases, it returns “0”.

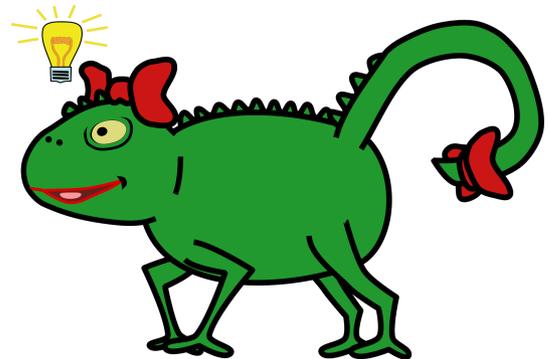
If you enter `yes` four times, the output looks like this:

```

1  -> To continue, enter "yes"
2  yes
3  -> To continue, enter "yes"
4  yes
5  -> To continue, enter "yes"
6  yes
7  -> To continue, enter "yes"
8  no
9  Stop!

```

By default, a variable is valid for the whole program. It is called, *global variable*. But it is possible to create variables in a function that are valid only for the function. This type of variable is called a *local variable* and it is created with the command `local`.



```

1 #!/bin/bash
2
3 localvar() {
4     local a=1
5     echo "function:" $a
6 }
7
8 a=0
9 echo "main before:" $a
10 localvar
11 echo "main after:" $a

```

Explanation: The main program starts in line 8. A variable `$a` is created and set to “0”. In line 10 a function is called. In this function a local variable is defined, that also is labeled “a” (line 4).

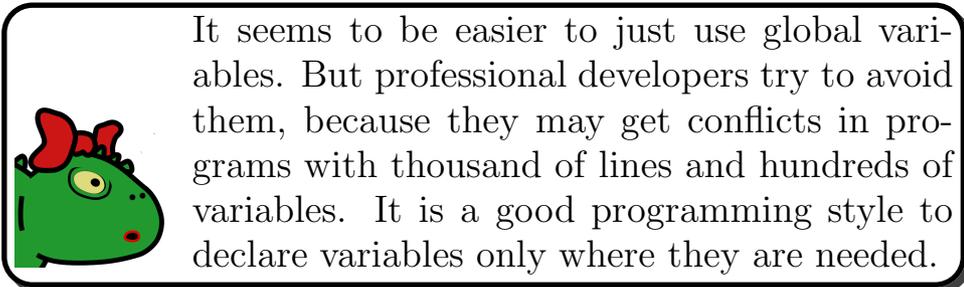
The output of the program is:

```

1 main before: 0
2 function: 1
3 main after: 0

```

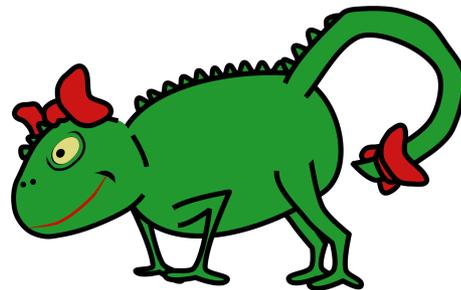
If a local variable has the same name as a global variable, the global variable is overlaid by the local variable inside the function. After executing the function, the global variable gets its original value back.



Exercise: Create Shell Functions

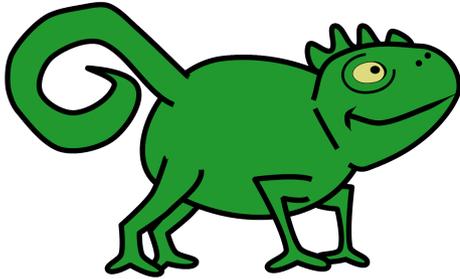
Write a bash program that reads numbers from the input and calculates the mean average of these numbers. If you enter “finish”, the program should end.

Use at least one function to check whether the user entered a number or “finish”.



General Programming Conventions

- ❶ Write comments in your program.
- ❷ Use white-spaces to format your program.
- ❸ Use self-explaining names for functions and variables.
- ❹ Use lowercase letters for variable names and function names.
- ❺ For constants, (they do not change their value during the program's run time) use uppercase letters.
- ❻ If you need variables for counting loops (e. g., in a for statement), use `$i`, `$j`, or `$k`.
- ❼ If you need a variable for only a short time (e. g., to buffer a value), use `$a`, `$b`, and `$c` for integer numbers and `$x`, `$y`, and `$z` for float numbers.



2.6 Understand the sed Stream Editor

sed is a stream editor that manipulates text files. That means it is different to many other text editors, like vi or gedit.

sed does not work interactively. It is controlled by command-line options or a script, like the ex mode of vi.

The stream editor loads the current line of the file into an internal buffer for manipulation. The results are sent to standard out. The original file is never changed. If you want to save the output in a file, you have to redirect the standard output (see section 2.3.1 on page 85).

You can start sed in general by

```
sed options 'command' filename
```

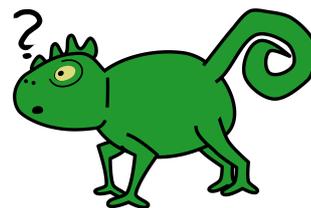
The simplest sed command is “p” (print), which prints out a line.

```
1 geeko@da51:~ > sed 'p' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
3 LOREM IPSUM DOLOR SIT AMET
4
5
6 Lorem ipsum dolor sit amet,
7 Lorem ipsum dolor sit amet,
8 ...
```

As you can see in the previous example, each line is printed out twice. sed sends each line once to standard out by default. To suppress this, you have to add the option `-n`.

```
1 geeko@da51:~ > sed -n 'p' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
3
4 Lorem ipsum dolor sit amet,
5 consectetur adipiscing elit,
6 ...
```

Let's look at the *command* part of the sed call.



First of all, we can specify an address declaration with the line numbers `sed` should examine.

Different kinds are possible:

No declaration – All lines

number – Only the line *number*; e. g.

```
1 geeko@da51:~ > sed -n '1p' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
```

start,end – The lines from *start* to *end*; e. g.

```
1 geeko@da51:~ > sed -n '2,4p' loremipsum.txt
2
3 Lorem ipsum dolor sit amet,
4 consectetur adipiscing elit,
```

\$ – The last line; e. g.

```
1 geeko@da51:~ > sed -n '36,$p' loremipsum.txt
2 facilisis at vero eros et accumsan
3 et iusto odio dignissim qui blandit
4 praesent luptatum zzril delenit
5 au gue duis dolore te feugiat nulla facilisi.
```

/regexp/ – The lines that match the regular expression *regexp*; e. g.

```
1 geeko@da51:~ > sed -n '/^D/p' loremipsum.txt
2 Duis te feugifacilisi.
3 Duis autem dolor in hendrerit
4 Duis te feugifacilisi per suscipit
5 Duis te feugifacilisi.
6 Duis autem dolor in hendrerit in
```

Explanation: All lines are printed that begin with an uppercase “D”.

number,/regexp/ – The lines from *number* to the first line that matches *regexp*; e. g.

```
1 geeko@da51:~ > sed -n '9,/D/p' loremipsum.txt
2 suscipit lobortis nisl
3 ut aliquip ex ea commodo consequat.
4 Duis te feugifacilisi.
```

Explanation: All lines are printed from line 9 to the first line that begins with an uppercase “D”.

You can also use a regular expression as starting point in the same way.

The regular expressions of sed are similar to those you learned in section 2.3.2 on page 93. However some more characters need to be escaped with an backslash:

☞ `\(...\)`

☞ `\{...\}`

You can combine two or more sed commands using the option “-e”:

```
1 geeko@da51:~ > sed -n -e '3p' -e '6p' -e '9p' loremipsum.txt
2 Lorem ipsum dolor sit amet,
  ut lacreet dolore magna aliquam erat volutpat.
4 suscipit lobortis nisl
```

Explanation: The lines 3, 6 and 9 are printed on the screen.

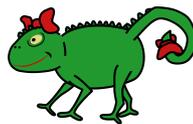
Alternatively, you can separate the single commands by a semicolon.

```
1 geeko@da51:~ > sed -n '3p;6p;9p' loremipsum.txt
2 Lorem ipsum dolor sit amet,
  ut lacreet dolore magna aliquam erat volutpat.
4 suscipit lobortis nisl
```



Exercise: Understand sed – Part I

Use sed to view the lines 10 to 15 from the `loremipsum.txt` file.



Some more of the most important sed commands are listed in the following:

☞ “a” (append)

Adds one or more lines *after* the specified line(s).

```
1 geeko@da51:~ > sed '1aHello Geeko' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
  Hello Geeko
4
  Lorem ipsum dolor sit amet,
6 consectetuer adipiscing elit,
  ...
```

If you want to add more than one line, you have to mask the line break with a backslash.

```
1 geeko@da51:~ > sed '1aHello Geeko\  
2 > Hello Suzie' loremipsum.txt  
LOREM IPSUM DOLOR SIT AMET  
4 Hello Geeko  
Hello Suzie  
6  
Lorem ipsum dolor sit amet,  
8 consectetur adipiscing elit,  
...
```

☞ “i” (insert)

Adds one or more lines *before* the specified line(s).

```
1 geeko@da51:~ > sed '1iHello Geeko' loremipsum.txt  
2 Hello Geeko  
LOREM IPSUM DOLOR SIT AMET  
4  
Lorem ipsum dolor sit amet,  
6 consectetur adipiscing elit,  
...
```

☞ “r filename”

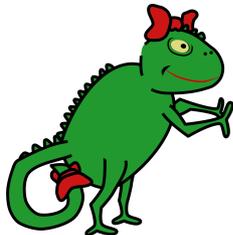
Inserts the content of the specified file behind the specified line(s).

```
1 geeko@da51:~ > sed '1r numbers.txt' loremipsum.txt  
2 LOREM IPSUM DOLOR SIT AMET  
0000  
4 123  
...  
6 98798  
765765  
8  
Lorem ipsum dolor sit amet,  
10 consectetur adipiscing elit,  
sed diem nonummy nibh euismodtincidunt  
12 ...
```

☞ “d” (delete)

Deletes the specified line(s).

```
1 geeko@da51:~ > sed '1,3d' loremipsum.txt  
2 consectetur adipiscing elit,  
sed diem nonummy nibh euismodtincidunt  
4 ut lacreet dolore magna aliquam erat volutpat.
```



☞ “c”

Replaces the specified lines.

```
1 geeko@da51:~ > sed '1,3cHello Geeko' loremipsum.txt
2 Hello Geeko
   consectetuer adipiscing elit,
4 sed diem nonummy nibh euismodtincidunt
   ut lacreet dolore magna aliquam erat volutpat.
```

Explanation: In this example lines 1 to 3 are replaced by the line “Hello Geeko”.

☞ “w”

Writes the output of sed into the specified file. You can use this command instead of redirecting the output using “>”.

```
1 geeko@da51:~ > sed -n '1w output.txt' loremipsum.txt
```

Explanation: A new file `output.txt` is created. It includes the first line of the file `loremipsum.txt`.

If you do not use the option “-n” like in the example, the content of `loremipsum.txt` is written to the screen (and not to `output.txt`). This option does not affect the file `output.txt`.



☞ “q” (quit)

Quits the stream editor.

```
1 geeko@da51:~ > sed '3q' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
4 Lorem ipsum dolor sit amet,
```



Exercise: Understand sed – Part II

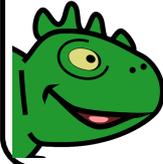
Use sed to remove the first ten lines of `loremipsum.txt`. The output should be saved in the file `shortlorem.txt`.

☞ `y/character(s)/character(s)/`

Changes single characters. The first character from the first set is repaced by the first character of the second set. The second character from the first set is replaced by the second character of the second set, and so on.

```
1 geeko@da51:~ > sed 'y/oim/XYZ/' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
4 LXreZ YpsuZ dXlXr sYt aZet,
   cXnsectetuer adYpYscYng eLYt,
6 sed dYeZ nXnuZZy nYbh euYsZXdtYncYdunt
```

Explanation: The lowercase “o” is replaced by “X”. The lowercase “i” is replaced by “Y” and the lowercase “m” is replaced by “Z”.



Instead of using the slash you can use any other character. The last command could also look like `sed 'y?oim?XYZ?' loremipsum.txt` or like `sed 'yBoimBXYZB' loremipsum.txt`. The character next to command “y” defines which character is used to divide the two character sets and finishes the command.

👉 `s/pattern/string/flag(s)` (substitute)

Replaces the first *pattern* (search pattern) with the second *string*. Normally, the search pattern includes regular expressions.

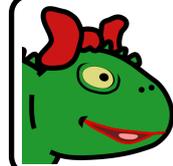
```
1 geeko@da51:~ > sed 's/dol/XXXXX/' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
3
4 Lorem ipsum XXXXXor sit amet,
  consectetuer adipiscing elit,
5 sed diem nonummy nibh euismodtincidunt
  ut lacreet XXXXXore magna aliquam erat volutpat.
6
7 ...
8
```

Explanation: The character combination “dol” is replaced by “XXXXX”.



Exercise: Understand sed – Part III

Use sed to replace “quis” by “quod” in the file `loremipsum.txt`. The output should be saved in file `quodlorem.txt`



Instead of using the slash you can use any other character.

Regular expressions are greedy, which means they try to match a string as large as possible. For example,

```
1 geeko@da51:~ > echo AABBBCCAABB | sed s/A.*A/XX/
2 XXBB
geeko@da51:~ >
```

Explanation: The regular expression matches a string that begins with an “A”, followed by zero or more arbitrary characters and ends with an “A”. Because regular expressions are greedy, the first and last “A” matche and not (for example) the first and the second “A”.

By default the “s” replaces only the first appearance of the search pattern. You can control the behavior of “s” by appending another substitution flag.

- g** – Any appearance of the search pattern is replaced in a line.
- i** – The case of the search pattern is ignored.
- p** – The line where the search pattern matches is put out. In combination with the sed option “-n” only the matching lines are put out.

```
1 geeko@da51:~ > sed -n 's/dol/XXXXX/p' loremipsum.txt
2 Lorem ipsum XXXXor sit amet,
  ut lacreet XXXXore magna aliquam erat volutpat.
4 Duis autem XXXXor in hendrerit
  vel illum XXXXore eu feugiat
```

- w filename** – The lines that match the search pattern are written to a file such as
`sed 's/dol/XXXXX/w output.txt'`
`loremipsum.txt`

- s** – The string is treated as single line. This makes possible to match a newline character (described by “\n”)

Of course, you can combine the substitution flags. E. g.
`sed 's/dol/XXXXX/gi' loremipsum.txt`

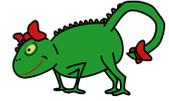


Sometimes you need the pattern that matches the regular expressing in the *string* part of the sed command again. For this there are nine register where you can store a match. You have to mark an expression you want to save with the grouping brackets `\(...\)` (see section 2.3.2 on page 97). You can refer to nine matches per line by `\1` to `\9`.

```
1 geeko@da51:~ > sed 's/\([a-c]\)/X\1X/g' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET

4 Lorem ipsum dolor sit XaXmet,
  XcXonseXcXtetuer XaXdipisXcXing elit,
6 sed diem nonummy niXbXh euismodtinXcXidunt
```

Explanation: The regular expression matches all “a”, “b” and “c” and encloses them with an “X” before and after.



☞ “n” (next)

The line that follows the search pattern is examined.

```
1 geeko@da51:~ >
```

☞ “h” (hold), “g”, “G” and “x”

The current line is loaded into a buffer (*pattern space*) from sed. Then, it is manipulated and sent to the output channel. Using the command “h”, you can copy the content of the pattern space into the holding buffer. With “G” you can insert the content of the holding buffer to the current line.

```
1 geeko@da51:~ > sed '1h;3G' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
4 Lorem ipsum dolor sit amet,
  LOREM IPSUM DOLOR SIT AMET
6 consectetur adipiscing elit,
```

Explanation: The first line is copied into the holding buffer and is inserted after the third line.

Command “g” replaces the current line by the content of the holding buffer.

```
1 geeko@da51:~ > sed '1h;3g' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
4 LOREM IPSUM DOLOR SIT AMET
  consectetur adipiscing elit,
6 sed diem nonummy nibh euismodtincidunt
```

The command “x” exchanges the content of pattern space and holding buffer.

```
1 geeko@da51:~ > sed '1h;2x;4G' loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET
  LOREM IPSUM DOLOR SIT AMET
4 Lorem ipsum dolor sit amet,
  consectetur adipiscing elit,
6
  sed diem nonummy nibh euismodtincidunt
8 ut lacreet dolore magna aliquam erat volutpat.
```

Explanation: The first command (“1h”) copies the line of `loremipsum.txt` into the holding buffer. “2x” exchanges the content of the second line (an empty line) with the content of the

holding buffer. After the fourth line the new content of the holding buffer (the empty line) is inserted, because of command “4G”.

You can write all or part of the sed commands in an external file and call it with sed option “-f”:

```
sed -f commands.txt loremipsum.txt
```

The following rules have to be valid for the external command file:

- ❶ A line that starts with a hash (“#”) is a comment line and not interpreted.
- ❷ In front of and after a command, no whitespaces are allowed.
- ❸ If you write more than one command in a line, you have to divide the commands by semicolons.

The whole command script is executed for each line of the text file.



Exercise: Understand sed – Part IV

Write a sed script, that replaces “dolor” with “XXXXX” and each “o” with “O” in the file `loremipsum.txt`. All lines beginning with “Duis” should be removed. The output should be saved in the file `scriptlorem.txt`.

Although the examples here appear a little bit simple, you might use sed later in your bash programs to manipulate files. The introduced functions of adding and removing items of a text file is important, e. g., if you want to evaluate automatically generated log files of physics experiments.

3 Linux Services for Developers

A common misconception of software developers is that they are sitting alone in their dark garages writing programs. But, the reality is different. Professional developers often belong to a team and are responsible for a part of a program. Therefore, they have to work together, mostly in an office.

Writing programs in a team requires tools that help manage the development. In this chapter, we introduce two important services for a developer:

☞ A version control system

☞ A web server



3.1 The Version Control System Subversion

When two or more people are working at the same document at the same time, you have got a problem: only the version of the person that saves the file last stored. The version of the first person is replaced. It is not possible to restore the first version, even if it was better than the second one.

To solve parts of this problem, you can add a timestamp in your filename. Then, for example, at the end of each day someone collects all new versions of the file, compares them and creates a “final version of the day” with all changes. This might work for small files and small groups, but is not very user friendly, even if you use the command `diff` to find the changes in a file. The syntax is:

```
diff file1 file2
```

The output looks like this:

```
1 geeko@da51:~> diff loremipsum.txt loremipsum2.txt
2 11d10
3 < Duis te feugifacilisi.
4 18a18
5 > Duis te feugifacilisi.
6 22c22
7 < nisl ut aliquip ex en commodo consequat.
8 ---
9 > ut aliquip ex commodo consequat.
10 34,36d33
11 < vulputate velit esse molestie consequat,
12 < vel illum dolore eu feugiat nulla
13 < facilisis at vero eros et accumsan
```

Explanation: Each difference starts with a number–character–number combination. The first number is the line number of the first file. The second number is the line number of the second file. The character between describes the kind of change. The most important characters are:

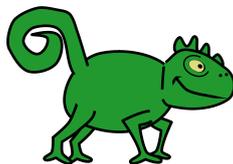
- a – A line is added (line 4–5).
- d – A line is deleted (line 2–3).
- c – The content of the line changed (line 6–9).

If there are more differences in several consecutive lines, the starting line and the ending line are separated by a colon (line 10). The lines that changed are also listed. A preceding “<” indicates that the line is not available in the second file anymore. A preceding “>” indicates that the line new in the second file.

A better choice is to use a versioncontrol software. This has two advantages:

- ❶ Several people can work on one file simultaneously.
- ❷ Previous versions of a file are stored and can be restalled.

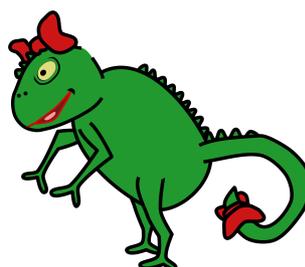
The dinosaur of versioncontrol systems is the open-source software Concurrent Versions System (CVS). But, because of some missing functionalities (e. g., renaming files), the system Subversion was designed to be the successor to CVS.



i There is a free book about Subversion available at <http://svnbook.red-bean.com>.

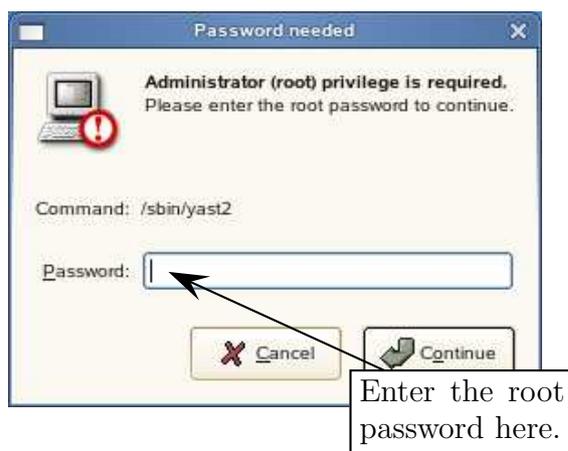


You can use Subversion for all kinds of files, and not only for program code (e. g., OpenOffice documents). We recommend using it when you prepare a presentation with others or you want to be sure that you can access previous document versions.

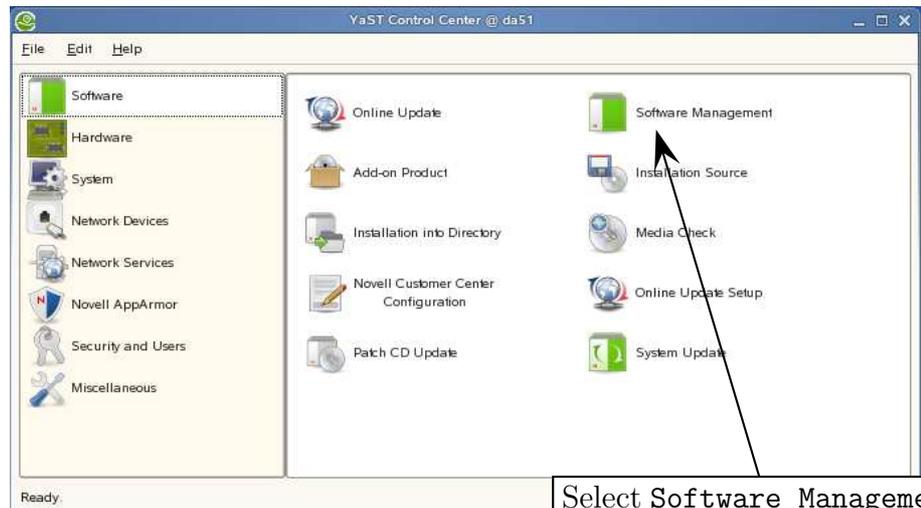


3.1.1 Install Subversion

Subversion is available on the SUSE Linux Enterprise DVD, but it is not installed by default. To install it, select the YaST icon in the **System** group of the Application Browser.

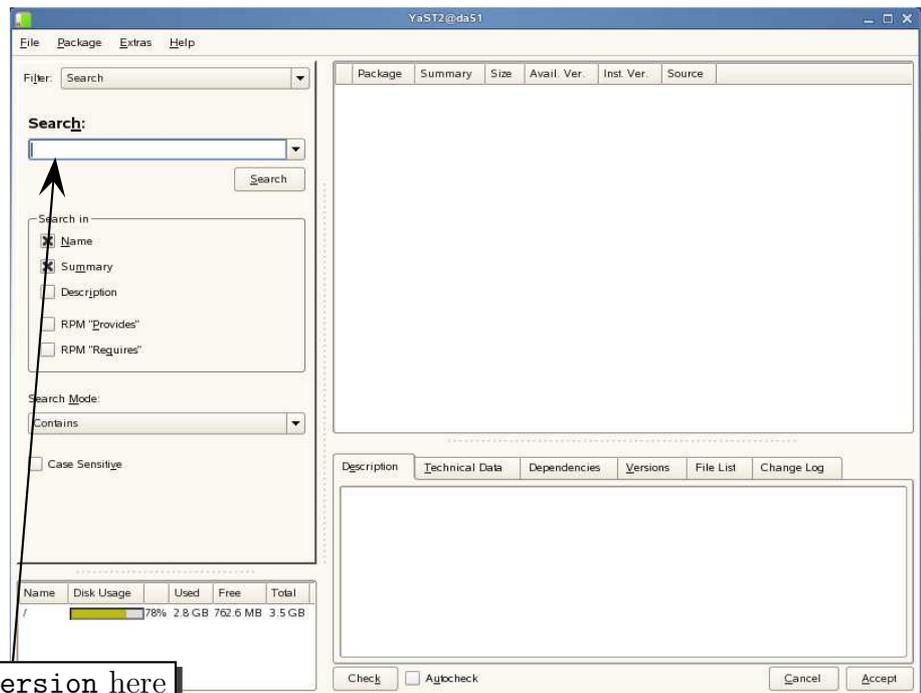


The YaST Control Center appears.

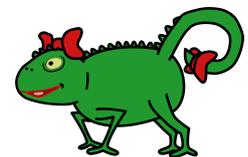


Select Software Management.

The Software Management dialog appears.



Enter subversion here and select Search.

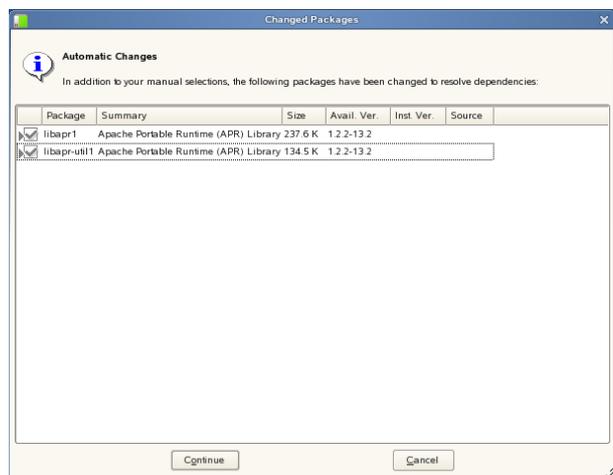


Two packages should be found:

- ☞ subversion
- ☞ subversion-devel

Select only `subversion` and then .

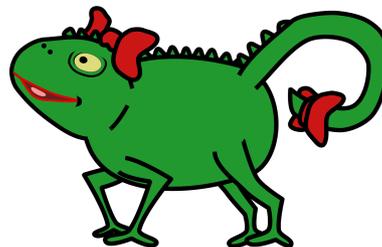
There should appear a message that two more packages are required to resolve the dependencies.



Select to go on.

☞ Exercise: Install Subversion

Install Subversion as described above.



3.1.2 Set up a Subversion Repository

In Subversion, you have a central repository for all projects. You do not work in this repository. You always use a local copy, and this is synchronized with the repository.

The repository can be stored on your local computer but also on a computer in the network. The last case is more usual. Subversion supports a lot of network protocols and it does not make any differences where the repository is located. In this course, we will use a local repository.

The first step in working with Subversion is to create a repository. Therefore, you have to use the command `svnadmin`.

`svnadmin` provides some options. The option we want to introduce here is `create`. With this option, a new Subversion repository is created. The syntax is:

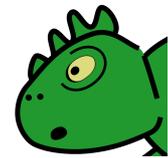
```
svnadmin create directory
```

The directory you specify must already exist. The repository is created inside. You can create a repository only in a directory you have write permission for. So, normally you will use the command `svnadmin` as user `root`.

To switch to user `root`, you have the following possibilities:

- ☞ Log out as normal user and log in as user `root`. → Bad solution
- ☞ Switch to one of the six virtual terminals by pressing `Ctrl` + `Alt` + `Fx` and log in as `root`. → Only a little bit better solution
- ☞ Start a terminal emulation from the main menu and enter the command `su - (switch user)`. You are then asked to enter the root password. → Best solution

The prompt changes when you are logged in as `root`. The username is not shown in the prompt anymore and at the end of the prompt there is a hash sign (“#”).



Before you create a repository think of where you want to create it. We recommend to use the `/srv` directory.

```
1 da51:~ # mkdir /srv/repository
2 da51:~ #
```

Then you can create the repository.

```
1 da51:~ # svnadmin create /srv/repository/
2 da51:~ #
```

Some files are created.

```

1 da51:~ # ls -l /srv/repository/
2 insgesamt 8
3 drwxr-xr-x 2 root root 128 2006-10-12 08:20 conf
4 drwxr-xr-x 2 root root 48 2006-10-12 08:20 dav
5 drwxr-sr-x 5 root root 256 2006-10-12 08:20 db
6 -r--r--r-- 1 root root 2 2006-10-12 08:20 format
7 drwxr-xr-x 2 root root 360 2006-10-12 08:20 hooks
8 drwxr-xr-x 2 root root 104 2006-10-12 08:20 locks
9 -rw-r--r-- 1 root root 229 2006-10-12 08:20 README.txt
10 da51:~ #

```

You also have to make sure that normal users (or better the users of a special group) have the write permission for this directory.

```

1 da51:~ # chgrp -R users /srv/repository/
2 da51:~ # chmod -R 775 /srv/repository/
3 da51:~ #

```

When you have finished this, you can log out of the terminal emulation as user root, by using the `exit` command.

```

1 da51:~ # exit
2 exit
3 geeko@da51:~>

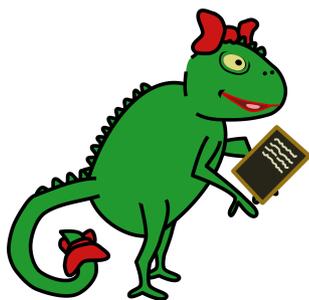
```

You are now logged in as a normal user.



Exercise: Create a Subversion Repository

Create a Subversion repository labeled `repository` in the `/srv` directory. Ensure that all users are allowed to change the content of the repository.



3.1.3 Start a Project

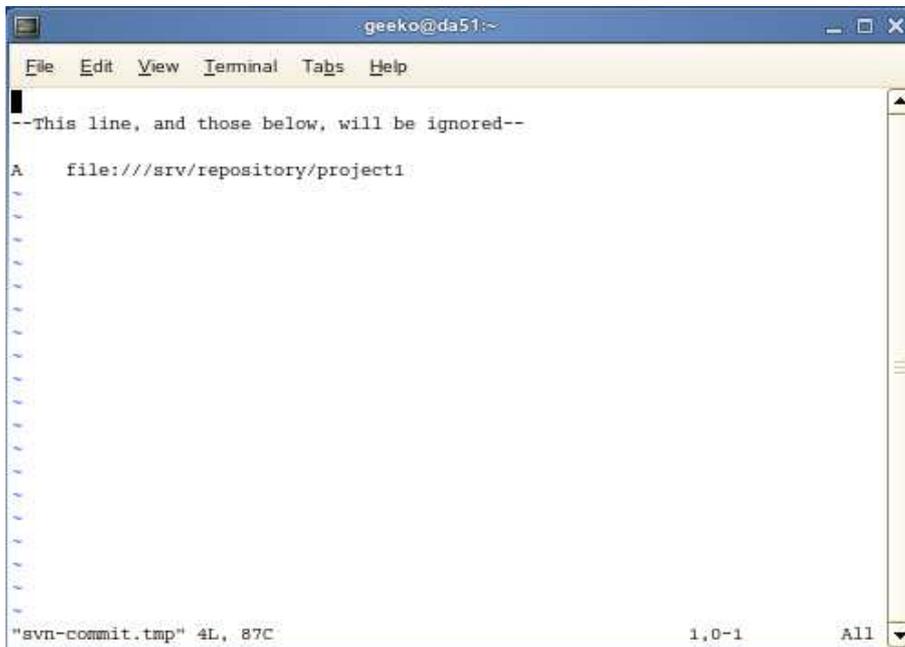
At the moment, you have an empty repository. Now you have to create a directory in Subversion where you can store the files for a project.

This can be done by the command `svn` with the option `mkdir`.

```
1 geeko@da51:~> svn mkdir file:///srv/repository/project1
```

Explanation: The string `file:///srv/repository/project1` specifies where (`/srv/repository/`) the new project directory should be created and what its name is (`project1`). The prefix `file://` says that it is in the local file system. This notation is similar to the address of web pages (e.g., `http://www.novell.com`).

The vi editor starts, and you can enter a description of your project.



```
geeko@da51:~
File Edit View Terminal Tabs Help
--This line, and those below, will be ignored--
A file:///srv/repository/project1
"svn-commit.tmp" 4L, 87C 1,0-1 All
```

After leaving vi, the directory is created, and you should get a message like this:

```
1 Committed revision 1.
```



Exercise: Create a Subversion Directory

Create a directory `project1` in your Subversion repository.

3.1.4 Prepare for Working with Subversion

In most cases, you already have some project files when you create your Subversion directory. If these files are stored in the same directory in the file system, you can import all files in one step.

The command to do this is `svn import`. You also have to specify the directory you want to import and the location of the subversion directory. For example:

```
1 geeko@da51:~> svn import my_project file:///srv/repository/project1/
2 Adding      my_project/numbers.txt
3 Adding      my_project/loremipsum.txt
4 Adding      my_project/dirgrep
5 Adding      my_project/dirgrep/loremipsum2.txt
6
7 Committed revision 2.
8 geeko@da51:~>
```

Explanation: `my_project` is the name of the local directory that should be imported. `file:///srv/repository/project1/` is the location of the Subversion directory where the files are imported.

After entering the command, vi starts again, and you can enter a note. After exiting vi, the files are added.

The notes you have to enter in vi can also be specified after the `svn` option `-m "notes"`.



Exercise: Import Files into Subversion

On the book CD there is an archive file `grep-demo.tar.gz`, which you already know from section 2.3.2 on page 91. Extract the archive into an empty directory `my_project`. Add this directory to `project1`.

Remove the `my_project` directory from your home directory after.



If you want to change the files in the repository, you have to create a local copy in your home directory first. This is done by the command `svn checkout`. You have to specify which Subversion directory you want to check out and where the local copy should be located. For example:

```
1 geeko@da51:~> svn checkout file:///srv/repository/project1 project1
2 A   project1/numbers.txt
3 A   project1/loremipsum.txt
4 A   project1/dirgrep
5 A   project1/dirgrep/loremipsum2.txt
6 Checked out revision 2.
geeko@da51:~>
```

Explanation: `file:///srv/repository/project1` is the location of the Subversion directory that is copied. `project1` is the name of the directory including the local copy. Because `project1` does not include any path description, the directory is created at the current position in the file system.

All files from the repository directory are copied. There is also a hidden directory `.svn` created, which includes meta information of Subversion.

```
1 geeko@da51:~> ls -la project1/
2 total 9
3 drwxr-xr-x  4 geeko users 160 2006-10-12 09:31 .
4 drwxr-xr-x 25 geeko users 1448 2006-10-12 09:31 ..
5 drwxr-xr-x  3 geeko users 104 2006-10-12 09:31 dirgrep
6 -rw-r--r--  1 geeko users 1344 2006-10-12 09:31 loremipsum.txt
7 -rw-r--r--  1 geeko users  55 2006-10-12 09:31 numbers.txt
8 drwxr-xr-x  7 geeko users 296 2006-10-12 09:31 .svn
geeko@da51:~>
```



Exercise: Create a Local Copy

Create a local copy of the `project1` Subversion directory. The name of the local copy should also be `project1`.

3.1.5 Use Subversion

During your daily work with Subversion you will do the following tasks frequently:

- ➡ **Transfer your work into the repository.** When you change the files of your local copy, you have to put your changes frequently (e. g., once a day) into the repository, so that the members of your team can get the latest version.

This is called *commit*.

- ➡ **Get the changes from others.** To be sure that you use the latest version before you start working with your local copy, you have to get all changes from your team members.

This is called *update*.

- ☞ **Add new files to the repository.** Add a new file to the Subversion directory to ensure that your team members also get it.
- ☞ **Remove files from the repository.** Remove files from the Subversion directory to ensure that your team members do not use it anymore.
- ☞ **Move or rename files inside the repository.** Move or rename a file inside the Subversion directory.
- ☞ **Solve conflicts.** In a large team, you cannot avoid that file is changed by two people in the same time period (that means, between update and commit). Subversion provides functionalities that help solve these conflicts.



Transfer your Work into the Repository

Because you are working with a local copy and not in the Subversion repository directly, you have to synchronize these two directories periodically. On the one hand your changes have to be transferred into the repository and on the other hand the changes from others in the repository have to be transferred to you.

To transfer your changes to the repository, use the command `svn commit`. You can add the name of a file or subdirectory to transfer changes in the file respectively in the subdirectory.

```
1 geeko@da51:~/project1> svn commit loremipsum.txt
2 Sending          loremipsum.txt
  Transmitting file data .
4 Committed revision 3.
geeko@da51:~/project1>
```

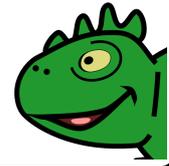
Explanation: Because the file `loremipsum.txt` is specified in the `svn commit` command, only this file is transferred, although there are also changes in other files in the project.

If you do not add any options the files of the current directory and all subdirectories are transferred.

```
1 geeko@da51:~/project1> svn commit
2 Sending          dirgrep/loremipsum2.txt
3 Sending          numbers.txt
4 Transmitting file data ...
  Committed revision 4.
6 geeko@da51:~/project1>
```

Explanation: Without any file specification all changes in files of the current directory and all subdirectories are transferred.

If there are no changes, no file will be transferred.



Before the files are transferred, vi starts, and you can enter a note about your changes.

Get the Changes from Others

Similar to `svn commit`, works `svn update`. This command transfers all files that changed into the repository in your working copy.

Without any file specification, all files and subdirectories are checked for updates. But you also can update a single file or subdirectory by appending its filename to the `svn update` command.

```
1 geeko@da51 :~/project1> svn update
2 U    loremipsum.txt
   Updated to revision 6.
4 geeko@da51 :~/project1>
```

Explanation: There is a newer version of the file `loremipsum.txt` available in the repository. It replaces the version of the file in your working copy.

The “U” at the beginning of line 2 indicates that a file is updated.



Exercise: Update Your Working Copy

In section 1.2.1 on page 20, you created a new user account. Open a new terminal emulation and log in as this new user by entering `su - username`.

This second user should checkout the `project1` Subversion directory into his home directory. Change something in the project files, and commit this changes into the repository.

Now, get these changes into the Subversion working copy of your normal user.

Add New Files to the Repository

To add a new file to the repository, use the command `svn add filename`.

```
1 geeko@da51:~/project1> svn add help.txt
2 A      help.txt
geeko@da51:~/project1>
```

The file is marked to be added. It will be transferred into the repository by the command `svn commit`.



Exercise: Add a File to the Repository

Create a new file as normal user and add it to the repository. Switch to your second user account and get the new file by updating your files.

Remove Files from the Repository

To remove a file from a project in the Subversion repository, use the command `svn delete filename`.

```
1 geeko@da51:~/project1> svn delete help.txt
2 D      help.txt
geeko@da51:~/project1>
```

The file is removed from your working copy immediately. It is removed from the repository after a `svn commit`.



Exercise: Remove a File from the Repository

Remove the file you created in the last exercise as normal user. Switch to your second user account and update your working copy.

Move or Rename Files Inside the Repository

If you want to rename a file or move it in another subdirectory of the project, use the command `svn rename old new`.

Internally a copy of the old file is added with the new filename. The original file with the old file is deleted.

```
1 geeko@da51:~/project1> svn move dirgrep/loremipsum2.txt dirgrep/README
2 A      dirgrep/README
3 D      dirgrep/loremipsum2.txt
4 geeko@da51:~/project1>
```

Explanation: The “A” at the beginning of line 2 indicates, that a file is added. The “D” at the beginning of line 3 indicates, that a file is deleted.

The changes are made directly in your working copy. In the repository they are made after a `svn commit`.



Exercise: Rename a File

Rename the file `dirgrep/loremipsum2.txt` to `dirgrep/README` as normal user. Switch to your second user account and update your working copy.

Copy Files Inside the Repository

If you want to copy a file or move it into another subdirectory of the project, use the command `svn copy source destination`.

```
1 geeko@da51:~/project1> svn copy loremipsum.txt dirgrep/loremipsum3.txt
2 A      dirgrep/loremipsum3.txt
3 geeko@da51:~/project1>
```

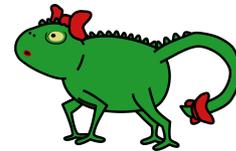
Explanation: The “A” at the beginning of line 2 indicates, that a file is added.

The changes are made directly in your working copy. In the repository, they are made after a `svn commit`.

Solve Conflicts

You cannot avoid that file is changed by two people in the same time period (that means, between update and commit).

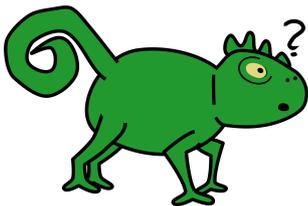
Example: User Tux updates his working copy of the repository project1 at 8:20 a. m. User Geeko updates his working copy of the same repository at 8:30 a. m. Tux finishes his work at the files at 10:00 a. m. and commits his changes into the repository. User Geeko works until 5:00 p. m. on the files and now he wants to commit his changes.



There appears to be a conflict in the upper example, if both do not touch the same files of the project. But, if so, user Geeko gets a conflict message when he tries to commit his files. For example:

```
1 geeko@da51:~/project1> svn commit
2 Sending      loremipsum.txt
3 svn: Commit failed (details follow):
4 svn: Out of date: '/project1/loremipsum.txt' in transaction '9-1'
5 svn: Your commit message was left in a temporary file:
6 svn:      '/home/geeko/project1/svn-commit.tmp'
geeko@da51:~/project1>
```

Explanation: In line 4 you can see that the conflict appears in the file `/project1/loremipsum.txt`.



To see what changed in the file in the meantime, use the command `svn diff filename`.

The output looks like this:

```

1 geeko@da51:~/project1> svn diff loremipsum.txt
2 Index: loremipsum.txt
3 =====
4 --- loremipsum.txt      (revision 6)
5 +++ loremipsum.txt      (working copy)
6 @@ -36,4 +36,4 @@
7     facilisis at vero eros et accumsan
8     et iusto odio dignissim qui blandit
9     praesent luptatum zzril delenit
10    -au gue duis dolore te feugat nulla facilisi.
11    +au duis dolore te feugat nulla facilisi.
12 geeko@da51:~/project1>

```

Explanation: In line 4 and 5 is noted, which symbols are used for the two files. A “-” represents the file `loremipsum.txt` of the repository (line 4), a “+” represents the file `loremipsum.txt` of the working copy (line 5).

In line 6 it is written which lines changed. The first two numbers (separated by comma) are marked with a “-”. They represent the line number and column number of the file in the repository. The second two numbers are marked with a “+”. They represent the line number and column number of the file in the working copy.

Then some lines of the file are listed (lines 7 – 11). Lines with a prefixed “-” (line 10) are only included in the file version of the repository. Lines with a prefixed “+” (line 11) are only included in the file version of the working copy.

If you see the differences, you can decide now whether you want to change these lines in your working directory.

If you want to throw away all your changes, use the command `svn revert filename`.

```

1 geeko@da51:~/project1> svn revert loremipsum.txt
2 Reverted 'loremipsum.txt'
3 geeko@da51:~/project1>

```

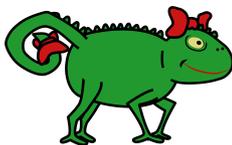
If you want to resolve the conflict, try `svn update filename`. If the changes in the repository version and the working copy do not overlap, both files are merged by Subversion.

```

1 geeko@da51:~/project1> svn update loremipsum.txt
2 G   loremipsum.txt
3 Updated to revision 12.
4 geeko@da51:~/project1>

```

Explanation: The “G” at the beginning of line 2 stands for merged. After merging the file, the changes can be committed.



If the changes in the repository version and the working copy overlap, you get a conflict when trying `svn update`.

```
1 geeko@da51:~/project1> svn update loremipsum.txt
2 C   loremipsum.txt
   Updated to revision 15.
4 geeko@da51:~/project1>
```

Explanation: The “C” at the beginning of line 2 indicates that there is a conflict.

If you open the file with a text editor now, you can see marks in the lines, where the conflict appeared. For example:

```
1 ...
2 facilisis at vero eros et accumsan
   et iusto odio dignissim qui blandit
4 praesent luptatum zzril delenit
   <<<<<<< .mine
6 duis dolore de feugiat nulla facilisi.
   =====
8 au gue duis dolore magna aliquam erat volutpat.
   >>>>>>> .r15
```

Explanation: The text between <<<<<<< .mine and ===== is the version in your working copy (line 6). The text between ===== and >>>>>>> .r15 is the version in the repository (line 8). The number `r15` specifies the version number.

Now you can decide which version you prefer and solve the conflict. After saving the corrected version of the file, tell Subversion that the conflict is solved. This is done by the command `svn resolved filename`.

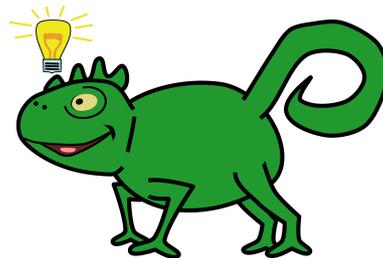
```
1 geeko@da51:~/project1> svn resolved loremipsum.txt
2 Resolved conflicted state of 'loremipsum.txt'
   geeko@da51:~/project1>
```

Now you can commit your changes.



Exercise: Solve Conflicts

Use your two user accounts to create a conflict in the file `loremipsum.txt` that can not be solved by merging files. Resolve the conflict manually.



3.1.6 Manage Subversion

You may ask yourself why you should enter notes in vi all the time? Subversion logs this information with every commit.

```

1 geeko@da51:~/project1> svn log
2 -----
3 r5 | tux | 2006-10-12 07:54:34 -0400 (Thu, 12 Oct 2006) | 2 lines
4
5 Renamed loremipsum2.txt to help.txt
6
7 -----
8 r4 | tux | 2006-10-12 07:51:24 -0400 (Thu, 12 Oct 2006) | 2 lines
9
10 Typo
11
12 -----
13 r3 | geeko | 2006-10-12 07:51:04 -0400 (Thu, 12 Oct 2006) | 2 lines
14
15 Added empty lines after each paragraph.
16
17 -----
18 r2 | geeko | 2006-10-12 07:49:59 -0400 (Thu, 12 Oct 2006) | 2 lines
19
20 Initial import
21
22 -----
23 r1 | geeko | 2006-10-12 07:49:55 -0400 (Thu, 12 Oct 2006) | 2 lines
24
25 Project created
26
27 -----
28 geeko@da51:~/project1>

```

Explanation: A horizontal line separates the revisions of the project. In the first line of each revision, the following information are listed:

Revision number | User name | Date and time of the commit | Number of comment lines

Then the notes are written.

If you want to see only the logs from revision *X* to *Y* add the option `-r X:Y`.

```

1 geeko@da51:~/project1> svn log -r 4:5
2 -----
3 r4 | tux | 2006-10-12 07:51:24 -0400 (Thu, 12 Oct 2006) | 2 lines
4
5 Typo
6
7 -----
8 r5 | tux | 2006-10-12 07:54:34 -0400 (Thu, 12 Oct 2006) | 2 lines
9
10 Renamed loremipsum2.txt to help.txt
11
12 -----
13 geeko@da51:~/project1>

```

Explanation: The order of log entries differs when you use `-r X:Y` or `-r Y:X`.

If you want to see, which files were changed, use the `--verbose` option.

```

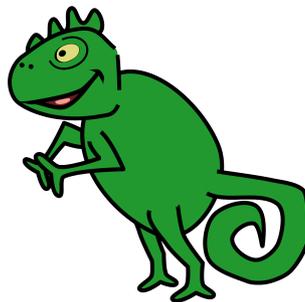
1 geeko@da51:~/project1> svn log --verbose
2 -----
3 r5 | tux | 2006-10-12 07:54:34 -0400 (Thu, 12 Oct 2006) | 2 lines
4 Changed paths:
5   A /project1/dirgrep/help.txt (from /project1/dirgrep/loremipsum2.txt:3)
6   D /project1/dirgrep/loremipsum2.txt
7
8 Renamed loremipsum2.txt to help.txt
9
10 -----
11 r4 | tux | 2006-10-12 07:51:24 -0400 (Thu, 12 Oct 2006) | 2 lines
12 Changed paths:
13   M /project1/loremipsum.txt
14
15 Typo
16
17 -----
18 r3 | geeko | 2006-10-12 07:51:04 -0400 (Thu, 12 Oct 2006) | 2 lines
19 Changed paths:
20   M /project1/loremipsum.txt
21
22 Added empty lines after each paragraph.
23
24 -----
25 r2 | geeko | 2006-10-12 07:49:59 -0400 (Thu, 12 Oct 2006) | 2 lines
26 Changed paths:
27   A /project1/dirgrep
28   A /project1/dirgrep/loremipsum2.txt
29   A /project1/loremipsum.txt
30   A /project1/numbers.txt
31
32 Initial import
33
34 -----
35 r1 | geeko | 2006-10-12 07:49:55 -0400 (Thu, 12 Oct 2006) | 2 lines
36 Changed paths:
37   A /project1
38
39 Project created
40
41 -----
42 geeko@da51:~/project1>

```



Exercise: View Subversion Logs

View your Subversion logs and check whether your notes are self explanatory.



If you want to see a previous version of a text file, you can use the command `svn cat` to view the content. Use the option `-r` to specify the revision number. For example:

```
1 geeko@da51:~/project1> svn cat -r 2 loremipsum.txt
2 LOREM IPSUM DOLOR SIT AMET

4 Lorem ipsum dolor sit amet,
   consectetur adipiscing elit,
6 sed diem nonummy nibh euismodtincidunt
   ...
8 et iusto odio dignissim qui blandit
   praesent luptatum zzril delenit
10 au gue duis dolore te feugiat nulla facilisi.
   geeko@da51:~/project1>
```

Once you know what the revision number of a file is, you can rebuild it, or parts of it. Use the command `svn merge` to merge a old version with the current one. Therefore you have to know the current revision number (see the log). With the option `-r current:previous`, you specify the revisions you want to merge. For example:

```
1 geeko@da51:~/project1> svn merge -r 5:2 file:///srv/repository/project1
2 U   loremipsum.txt
   D   dirgrep/help.txt
4 A   dirgrep/loremipsum2.txt
   geeko@da51:~/project1>
```

Explanation: The current revision is 5. The project should be rolled back to revision 2. The file `loremipsum.txt` is replaced by the old revision (line 2). The file `dirgrep/help.txt` was not available in revision 2. It is deleted (line 3). The file `dirgrep/loremipsum2.txt` is added (line 4), because it was deleted in the meantime.

It is also possible to get conflicts here. They are indicated by a “C” at the beginning of a line. Use the methods you learned before (see section 3.1.5 on page 174) to solve them.

If you want to rebuild a file you (or someone else) removed from the repository, use the command `svn copy`. With the option `-r revision` you can specify the revision where the still file exists. You also have to specify where the copy of the deleted file should be located in your working copy and what its name is. For example:

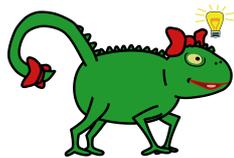
```
1 geeko@da51:~/project1> svn copy -r 4 file:///srv/repository/project1/dirgrep/lor
2 emipsum2.txt ./dirgrep/loremipsum2.txt
   A   dirgrep/loremipsum2.txt
4   geeko@da51:~/project1>
```

Explanation: The file `/srv/repository/project1/dirgrep/loremipsum2.txt` was deleted by renaming in revision 5 (see line 5 in the log on page 178). Revision 4 is the last revision that exists. It is copied from this revision into the working copy at `./dirgrep/loremipsum2.txt`.

If the project you manage with Subversion is terminated and you do not need the files in the repository anymore, use `svn rm` to remove the project.

```
1 geeko@da51:~> svn rm file:///srv/repository/project1
2
3 Committed revision 3.
4 geeko@da51:~>
```

The files of your working copy you can remove manually.



3.1.7 Use Subversion in a Network

If you want to access a Subversion repository that is not located on your local computer but on a server in your network, There are two possibilities:

- 1 You can access the server via HTTP. Therefore a web server using WebDAV must run on the server, too.
- 2 You can access a Subversion server that runs on the machine.

We want to introduce the second possibility here and what you have to do to set up a Subversion server.



Set Up a Subversion Server

After installing the Subversion software package there is also a program called `svnserv` installed. This program is the Subversion server. You can start it by entering `svnserv -d` at the command line. There is no further output.

```
1 geeko@da51:~> svnserv -d
2 geeko@da51:~>
```

But, now you can use the Subversion commands on another computer to work with the repository. Instead of `file://path` you now have to use `svn://IP-address/path` or `svn://hostname/path`. For example:

☞ Use the IP address of the server.

```
1 geeko@da53:~> svn checkout svn://10.10.0.51/srv/repository/project1/
2 A    project1/numbers.txt
  A    project1/loremipsum.txt
4 A    project1/dirgrep
  A    project1/dirgrep/help.txt
6 Checked out revision 5.
geeko@da53:~>
```

☞ Use the hostname of the server.

```
1 geeko@da54:~> svn checkout svn://da51.digitalairlines.com/srv/repository/projec
2 t1/
  A    project1/numbers.txt
4 A    project1/loremipsum.txt
  A    project1/dirgrep
6 A    project1/dirgrep/help.txt
  Checked out revision 5.
8 geeko@da54:~>
```

This kind of server connection is not secure, because...

- ❶ the information that is transferred via the network is not encrypted
- ❷ every user in the network can access the repository

To solve the first problem, you can connect with the server using the `svn+ssh` protocol. When you try to connect using the `svn+ssh` protocol, you have to authenticate.

```
1 geeko@da54:~> svn checkout svn+ssh://192.168.2.132/srv/repository/project1/
2 Password:
```

You need a local account on the Subversion server to use this method. But, it is still possible to use the simple `svn` protocol for other users.

To stop the Subversion server, enter
`killall svnserve.`



To enforce user authentication, you can start the Subversion server with the option `-t` (tunneled).

```
1 da51:~ # svnserve -t
2 ( success ( 1 2 ( ANONYMOUS EXTERNAL ) ( edit-pipeline ) ) ) )
```

It is not possible to connect by using the simple svn protocol now.

```
1 geeko@da54:~> svn checkout svn://192.168.2.132/srv/repository/project1/
2 svn: Can't connect to host '192.168.2.132': Connection refused
geeko@da54:~>
```

But svn+ssh works fine.

```
1 geeko@da54:~> svn checkout svn+ssh://192.168.2.132/srv/repository/project1/
2 Password:
```



Exercise: Set Up a Subversion Server

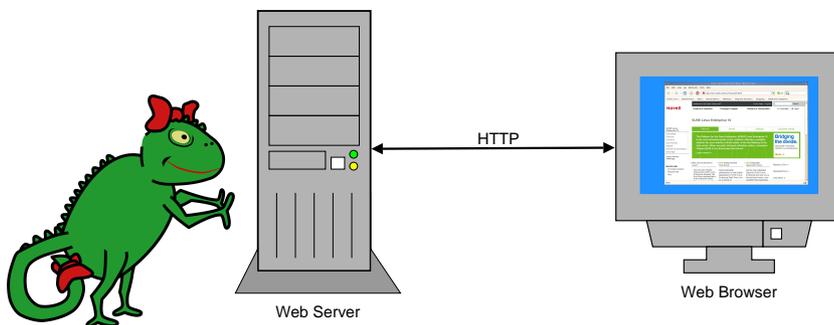
Start a Subversion server on your computer. Check out the files of project1 of your neighbour.

3.2 The Web Server Apache

Once your program get a stable version, you want to share it with other developers. To do this, you need a web server to provide some documentation and your software packages for download. In this chapter, you will learn how to set up your own web server.

A web server delivers data requested by a web browser. The data have different formats such as HTML files, image files, Flash animations, or sound files.

Web browsers and web servers communicate with each other by using HTTP (Hyper Text Transfer Protocol). The following diagram shows the relationship between browser, server, and HTTP:



Besides delivering data to the web browser, a web server performs tasks such as limiting access to specific web pages, logging access to a file, and encrypting the connection between server and browser.

The commonly used web server is Apache.

3.2.1 Install Apache

The Apache web server is not included in SUSE Linux Enterprise Desktop 10. But you can use the software packages from SUSE Linux Enterprise Server 10, because both are based on the same code. In the directory `packages/` on the book CD, you will find the following files:

- 👉 `apache2-2.2.0-21.2.i586.rpm`
- 👉 `apache2-example-pages-2.2.0-21.2.i586.rpm`
- 👉 `apache2-prefork-2.2.0-21.2.i586.rpm`

There are several ways to install RPM packages. In section 1.3 on page 23 you learned how to install RPM packages from the SUSE Linux Enterprise Desktop DVD using YaST.

In this chapter, you learn how to install software from the command line using the command `rpm`.

Start a terminal emulation. Then switch to user root by entering the command `su -`. You are asked to enter the root password.

To install a software package, use the command `rpm`. In the following, the syntax is described briefly.

i All options of `rpm` can be found in the manual page (`man rpm`).

To install all RPM packages of the `/media/cdrecorder/packages/` directory, enter the following:

```
1 da51:~ # rpm -ihv /media/cdrecorder/packages/*.rpm
```

The options are:

- i – Install package.
- h – Show the installation process using hash signs.
- v – Enable verbose output.

After the options, the path is entered where the RPMs can be found. The wildcard (“*”) means no or any characters can match the wildcard.

When the installation is finished, the output looks like the following:

```

1 da51:~ # rpm -ihv /media/cdrecorder/packages/*rpm
2 Preparing... ##### [100%]
   1:apache2-example-pages ##### [ 33%]
   2:apache2-prefork ##### [ 67%]
   3:apache2 ##### [100%]
6 Starting SuSEconfig, the SuSE Configuration Tool...
Running module permissions only
8 Reading /etc/sysconfig and updating the system...
Executing /sbin/conf.d/SuSEconfig.permissions...
10 Checking permissions and ownerships - using the permissions files
   /etc/permissions.d/postfix
12   /etc/permissions.d/susehelp
   /etc/permissions
14   /etc/permissions.easy
   /etc/permissions.local
16 setting /usr/sbin/suexec2 to root:root 4755. (wrong permissions 0755)
Finished.
18 Updating etc/sysconfig/apache2...
looking for old 2.0 modules to be renamed...
20 Done.
da51:~ #

```

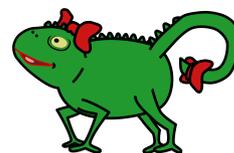


You will find more information on the rpm command in Appendix A on page 199.



Exercise: Install Apache

Install the Apache web server as described above.



3.2.2 Start and Test the Web Server

After installing the required software, start the web server. Switch to user root and enter `rcapache2 start`.

```
1 da51:~ # rcapache2 start
2 Starting httpd2 (prefork) done
da51:~ #
```

If you want to stop the Apache web server, enter `rcapache2 stop`.

```
1 da51:~ # rcapache2 stop
2 Shutting down httpd2 (waiting for all children to terminate) done
da51:~ #
```

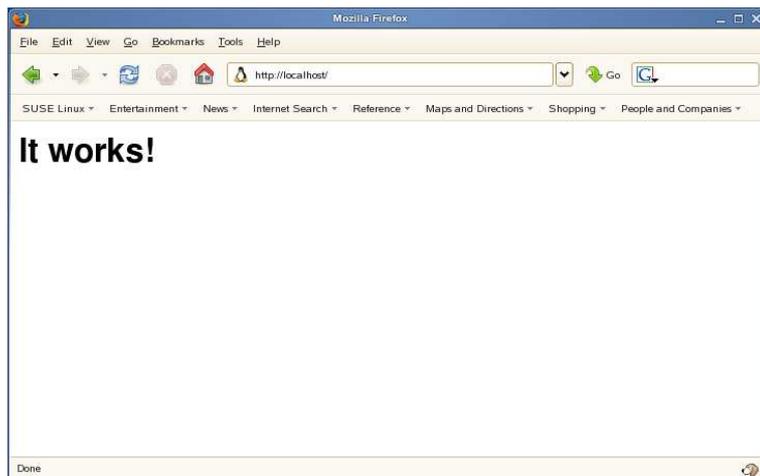
Apache doesn't start automatically when you reboot your computer. If you want the web server to start at boot time, enter `insserv apache2`.

```
1 da51:~ # insserv apache2
2 da51:~ #
```

To test the web server, open a web browser and enter the following address:

`http://localhost`

The browser displays the following page:



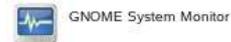
If your SUSE Linux Enterprise Desktop 10 is connected to a network, you (and other hosts on the network) can remotely access the web server by entering the following:

`http://IP-address`

If your network provides a DNS server, you can use the hostname instead of the IP address.

When Apache is running, the system starts some Apache processes, because each web server request needs its own Apache process. Handling more than one requests at the same time is called *pre-forking*.

If you start the GNOME System Monitor from the **System** group in the Application Browser, you can see some more interesting details.



To get the following window, where you can see all processes (and not only yours), select **View** → **All Processes**. Then, enable viewing threads by selecting **View** → **Threads**. In **Edit** → **Preferences**, activate **User** in the **Process Fields** list in the **Processes** tab.



Process Name	User	Status	Resident Memory	% CPU	Nice	ID	Arguments
gnome-volume-manager	geeko	Sleeping	1.2 MiB	0	0	6677	gnome-volu
hald	root	Sleeping	1.4 MiB	0	0	2839	/usr/sbin/ha
hald-addon-acpi	root	Sleeping	80.0 KiB	0	0	2846	hald-addon
hald-addon-storage	root	Sleeping	88.0 KiB	0	0	2902	hald-addon
httpd2-prefork	root	Sleeping	1.1 MiB	0	0	19411	/usr/sbin/htt
httpd2-prefork (thread)	wwwrun	Sleeping	1.1 MiB	0	0	19422	/usr/sbin/htt
httpd2-prefork (thread)	wwwrun	Sleeping	1.1 MiB	0	0	19416	/usr/sbin/htt
httpd2-prefork (thread)	wwwrun	Sleeping	1.1 MiB	0	0	19415	/usr/sbin/htt
httpd2-prefork (thread)	wwwrun	Sleeping	1.1 MiB	0	0	19414	/usr/sbin/htt
httpd2-prefork (thread)	wwwrun	Sleeping	1.2 MiB	0	0	19413	/usr/sbin/htt
httpd2-prefork (thread)	wwwrun	Sleeping	1.2 MiB	0	0	19412	/usr/sbin/htt
init [5]	root	Sleeping	44.0 KiB	0	0	1	init [5]
kacpid	root	Sleeping	0 bytes	0	-5	8	
kauditd	root	Sleeping	0 bytes	0	-5	5868	

The Apache processes are called `httpd2-prefork`.

In the **User** column, you can see that only one process runs with root permission. This process started the other Apache processes running with the permission of user `wwwrun`. These processes will answer to HTTP requests.



Exercise: Start Apache

Start your web server and test whether you can access the web server of your neighbor.

3.2.3 Locate the Document Root of the Web Server

The default directory of the data provided by Apache is `/srv/www/htdocs/`. This directory is also called *DocumentRoot* of the web server.

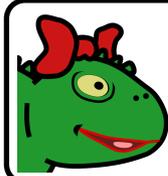
After the installation of `apache2-example-pages-2.2.0-21.2.i586.rpm`, the directory `/srv/www/htdocs/` contains some Apache example pages.

```
1 da51:~ # ls /srv/www/htdocs/
2 apache_pb.gif  apache_pb2.gif  apache_pb2_ani.gif  gif  robots.txt
3 apache_pb.png  apache_pb2.png  favicon.ico          index.html
4 da51:~ #
```

You can replace the data in the DocumentRoot directory to display your own web server content. As the web server runs with the permission of user `wwwrun`, ensure that this user can read the files in the DocumentRoot directory.

If you create subdirectories in DocumentRoot, you can access these subdirectories with the following web address scheme:

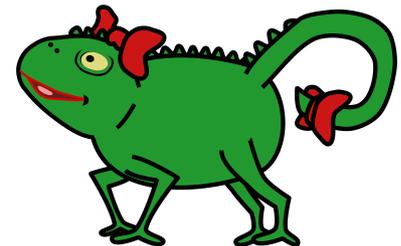
`http://IP-address/subdirectory`



If no specific file is requested in the address, Apache looks for a file named `index.html`. You can change this name in the Apache configuration files.

3.2.4 Locate the Apache Configuration Files

The Apache web server has several configuration files located in the directory `/etc/apache2/`.



The following is a list of the most important Apache configuration files:

httpd.conf. This is the main Apache configuration file.

default-server.conf. This file contains the basic web server setup. However, all options set in this file can be overwritten by other configuration files.

uid.conf. This configuration file sets the user and group ID for Apache. By default, Apache uses the user `wwwrun` and the group `www`.

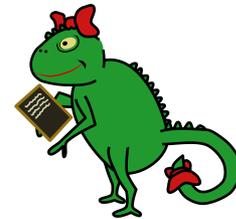
listen.conf. In this configuration file, you can specify the IP addresses and TCP/IP ports Apache is listening to. By default, Apache listens to all assigned interfaces on port 80.

server-tuning.conf. You can use this configuration file to fine tune the performance of Apache. The default values should be fine, unless you need a web server that has to handle a lot of requests at the same time.

error.conf. In this file, you configure what to do if Apache cannot handle a request correctly.

ssl-global.conf. Configure the connection encryption with SSL in this configuration file.

In this manual, we only describe the file `default-server.conf`.



3.2.5 Understand the Basic Rules of the Configuration Files

The options of the Apache configuration files are called *directives*. Directives are case sensitive, e. g., “include” is not the same as “Include.”

Directives can be grouped so that they do not apply to the global server configuration. In the following, the directives only apply to the directory `/srv/www/htdocs/`:

```
1 <Directory "/srv/www/htdocs">
2     Options None
3     AllowOverride None
4     Order allow,deny
5     Allow from all
6 </Directory>
```

The directives are grouped by `<Directory "/srv/www/htdocs">` and `</Directory>` which limits their validity to the directory `/srv/www/htdocs/` only.



You can use the “#” character to write comments in the configuration file. All lines starting with a “#” are ignored by the Apache program.

Every time you edit the Apache configuration files, you have to reload the web server by entering `rcapache2 reload`.

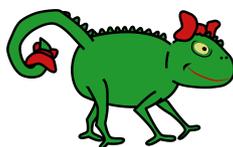
```
1 da51:~ # rcapache2 reload
2 Reload httpd2 (graceful restart) done
da51:~ #
```

Sometimes, you have to stop and restart the web server by entering `rcapache2 restart`.

```
1 da51:~ # rcapache2 restart
2 Syntax OK
3 Shutting down httpd2 (waiting for all children to terminate) done
4 Starting httpd2 (prefork) done
da51:~ #
```

If you changed something in the configuration files and you want to verify whether the syntax is correct, enter `apache2ctl configtest`. If the syntax is correct, the command displays the following message:

```
1 da51:~ # apache2ctl configtest
2 Syntax OK
da51:~ #
```



3.2.6 Understand the Default Apache Configuration

The file `/etc/apache2/default-server.conf` is the main configuration file of the Apache web server. The following is an overview of the most important directives used in that file:

- ☞ **DocumentRoot**. Specifies the DocumentRoot of the web server.

For example:

```
1 DocumentRoot "/srv/www/htdocs"
```

- ☞ **<Directory "directory"> .../<Directory>**. All directives within this block apply only to the specified directory.

For example:

```
1 <Directory "/srv/www/htdocs">
2     Options None
3     AllowOverride None
4     Order allow,deny
5     Allow from all
6 </Directory>
```

- ☞ **Options**. With this directive, additional options can be applied to logical blocks, e. g., directories.

For example, in line 2 of the listing above there are no more options.

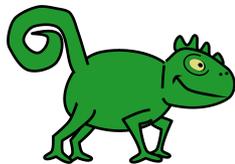
- ☞ **AllowOverride**. Determines whether directives are allowed to be overwritten by a configuration found in a `.htaccess` file in a directory.

For example, in line 3 of the listing above, it is not allowed to overwrite any directives.

- ☞ **Alias *fakename* "realname"**. Allows you to create an alias of a directory.

For example:

```
1 Alias /icons/ "/usr/share/apache2/icons/"
```



- ☞ **ScriptAlias**. Allows you to create an alias of a directory containing scripts for dynamic content generation.

For example:

```
ScriptAlias /cgi-bin/ "/srv/www/cgi-bin/"
```

In most cases, the default settings don't need to be changed.

 An overview of all Apache directives can be found at <http://httpd.apache.org/docs-2.0/mod/directives.html>.

3.2.7 Limit Access to the Web Server

Normally Apache delivers data to all hosts in the network the web server is connected to. Sometimes it can be useful to limit access to the content delivered by Apache.

The following are the commonly used methods:

- ☞ Limit access using IP addresses
- ☞ Limit access using user authentication

Limit Access Using IP Addresses

If you want to use IP addresses to limit access to the web server, the following directives are available:

Order. This directive defines in which order in which the allow and deny directives are evaluated.

You have the following options:

- ☞ **deny,allow.** First, the deny directives are evaluated, then the allow directives. Access is allowed by default. Each client that does not match a deny directive or does match an allow directive, has access to the server.
- ☞ **allow,deny.** First, the allow directives are evaluated, then the deny directives. Access is denied by default. Each client that does not match an allow directive or does match a deny directive, does not have access to the server.

- ☞ **Mutual-failure.** Only those hosts listed in the allow list and not in the deny list have access. This has the same effect as `Order allow,deny`.

For example:

```

1 <Directory "/srv/www/htdocs">
2     Options None
3     AllowOverride None
4     Order allow,deny
5     Allow from all
6 </Directory>
```

Allow from. IP addresses or networks listed here are allowed to access the web server.

Deny from. IP addresses or networks listed here are not allowed to access the web server.

You can use the following options with the `Deny from` and the `Allow from` directives:

- ☞ **all.** This option applies to all hosts.
- ☞ **A (partial) domain-name.** This option applies to hosts that names match or end in the given expression (such as `novell.com`). Only complete domain components are matched. “`novell.com`” matches `www.novell.com`, but not `foonovell.com`.
- ☞ **A full IP address.** This option applies to a specific IP address (such as `10.0.0.23`).
- ☞ **A partial IP address.** This option applies to IP addresses starting with the given IP address fragment (such as `10.0.0`).
- ☞ **A network/netmask pair.** This option applies to IP addresses matching to the given network/netmask pair (such as `10.0.0.0/255.255.255.0` or `10.0.0.0/24`).

All these directives have to be defined within a `<Directory>` block. They control the access to all data in that directory and all subdirectories.

For example:

```

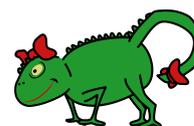
1 <Directory "/srv/www/htdocs">
2 Order deny,allow
3 Deny from all
4 Allow from 10.0.0.0/24
5 </Directory>
```

Explanation: In this example, only hosts from the network 10.0.0.0/24 have access to the data in the directory `/srv/www/htdocs/`.



Exercise: Limit Access Using IP Addresses

Ask for your neighbor's IP address. Then configure your Apache web server, that your neighbour's computer is denied access to your web server.



Limit Access Using User Authentication

Using IP addresses, you can control the hosts that accesses the web server, but not the users using these computers.

Apache offers another option of access control called *basic authentication*. If you protect content on your web server with this method, users are required to log in before they can access the data.

To use basic authentication, you first have to create user accounts for the web server. To do this, use the command `htpasswd2`.

The following command creates a password file and an account for the user `geeko`.

```
1 da51:~ # htpasswd2 -c /etc/apache2/htpasswd geeko
2 New password:
  Re-type new password:
4 Adding password for user geeko
da51:~ #
```

Explanation: After entering this command, `htpasswd2` prompts you for a password for the new user. The passwords are stored in the file `/etc/apache2/htpasswd`.

You can specify a different location for the password file, but you have to make sure that it is readable for the user `wwwrun` and not located within the `DocumentRoot` of your server.

To use a password file for the first time, use `htpasswd2` with the `-c` option to create the file. If you want to add more users later, use the following command:

```
htpasswd2 /etc/apache2/htpasswd username
```

To delete a user from the password file, use the following command:

```
htpasswd2 -D /etc/apache2/htpasswd username
```

For example:

```
1 da51:~ # htpasswd2 -D /etc/apache2/htpasswd geeko
2 Deleting password for user geeko
da51:~ #
```

After you created user accounts, configure Apache to prompt for a password when accessing restricted data. Add the following lines to the restricted directory's block:

```
1 AuthType Basic
2 AuthName "Restricted Files"
  AuthUserFile /etc/apache2/htpasswd
4 Require user geeko
```

The following describes each line:

AuthType Basic. This directive defines the authentication method. For the type described in this section, the value is **Basic**.

AuthName "Restricted Files". With this directive, you have to choose a label for the restricted directory of your web server. This name is used for the authentication process between browser and web server.

AuthUserFile /etc/apache2/htpasswd. This directive defines the password file used for the restricted directory.

Require user geeko. This directive lists the user included in the password file who is allowed to access the directory. You can add more than one user by separating the user names with spaces. You can also use the directive **Require valid-user**. This defines that all valid combinations of user names and passwords are granted access.

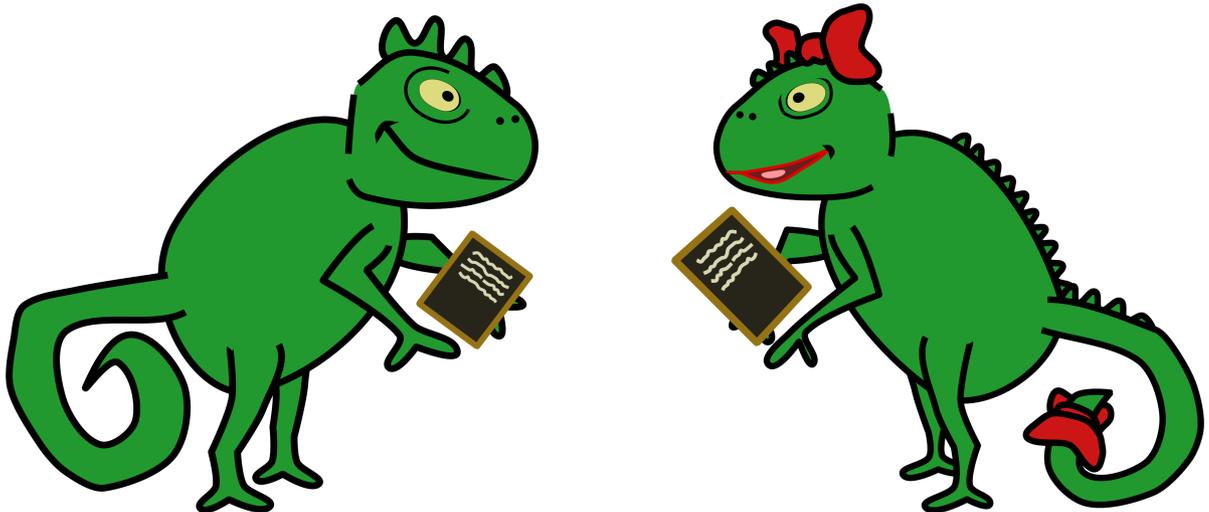
If the user wants to access web pages in the protected (sub-)directory, he is prompted to enter his user name and password.





Exercise: Limit Access Using User Authentication

In section 1.2.1 on page 20, you created a new user account. Configure your Apache web server so that this user gains access to your web server by using a user name and password.



Appendix

A The rpm Command



You can install software in the RPM format by using YaST or by using the command `rpm`. YaST ensures the automatic resolution of dependencies, while `rpm` only controls them. Resolution must be performed by hand.

Using the `rpm` command, you can install software (`rpm -i`), uninstall software (`rpm -e`), and query information from the RPM database (`rpm -q`).

The following are examples of using `rpm`:

- 👉 `rpm -qa`. Displays a list of all installed packages.
- 👉 `rpm -qi package`. Displays summary information about the specified package, for example:

```
1 da51:~ # rpm -qi bash
2 Name      : bash                      Relocations: (not relocatable)
3 Version   : 3.1                       Vendor: SUSE LINUX Products GmbH
4 , Nuernberg, Germany
5 Release   : 24.4                       Build Date: Fri Jun 16 09:14:10 2006
6 Install Date: Thu Oct 12 05:48:17 2006 Build Host: barbella.suse.de
7 Group     : System/Shells              Source RPM: bash-3.1-24.4.src.rpm
8 Size      : 2627789                    License: GPL
9 Signature : DSA/SHA1, Fri Jun 16 09:19:35 2006, Key ID a84edae89c800aca
10 Packager  : http://bugs.opensuse.org
11 URL       : http://www.gnu.org/software/bash/bash.html
12 Summary   : The GNU Bourne-Again Shell
13 Description :
14 Bash is an sh-compatible command interpreter that executes commands
15 read from standard input or from a file. Bash incorporates useful
16 features from the Korn and C shells (ksh and csh). Bash is intended to
17 be a conformant implementation of the IEEE Posix Shell and Tools
18 specification (IEEE Working Group 1003.2).
19
20
21
22 Authors :
23 -----
24 Brian Fox <bfox@gnu.org>
25 Chet Ramey <chet@ins.cwru.edu>
26 Distribution: SUSE Linux Enterprise 10 (i586)
da51:~ #
```

For information about a package that has not been installed, add the option `-p` and specify the complete path to the package in question. Entering `rpm -qp1/software/package.rpm` lists all the files in `package.rpm`, as in the following example:

```

1 da51:~ # rpm -qp1 apache2-example-pages-2.2.0-21.2.i586.rpm
2 /srv/www/htdocs/apache_pb.gif
3 /srv/www/htdocs/apache_pb.png
4 /srv/www/htdocs/apache_pb2.gif
5 /srv/www/htdocs/apache_pb2.png
6 /srv/www/htdocs/apache_pb2_ani.gif
7 /srv/www/htdocs/favicon.ico
8 /srv/www/htdocs/index.html
9 /srv/www/htdocs/robots.txt
10 da51:~ #

```

You can also use the following options with the `rpm` command:

Option	Description
<code>--checksig</code>	Makes sure a package is complete and correct (PGP signature check).
<code>-e</code> or <code>--erase</code>	Uninstalls a package.
<code>-F</code> or <code>--freshen</code>	Updates a package only if the package is already installed in an older version.
<code>--help</code>	Lists options for the <code>rpm</code> command.
<code>-i</code> or <code>--install</code>	Installs a package.
<code>-i --force</code>	Installs a package even if <code>rpm</code> says that the package is already installed.
<code>-i --nodeps</code>	Installs a package regardless of unresolved dependencies.
<code>-qa</code>	Lists all installed packages.
<code>-qf</code>	Displays which package a file belongs to.
<code>-qi</code>	Gives brief information about the given package.
<code>-ql</code>	Lists all files that belong to a package.
<code>-qp1</code>	Lists all files of a package that are not installed.
<code>-U</code> or <code>--upgrade</code>	Updates a package. (In contrast to <code>-f</code> , the package is installed, even if there is no older version of the package installed.)
<code>-U --oldpackage</code>	Installs an older version, even if the newer version is already installed (downgrade).

Option	Description
<code>-V</code> or <code>--verify</code>	Compares information about the installed files in the package with information about the files taken from the original package and stored in the rpm database. Among other things, verifying compares the size, permissions, type, owner, and group of each file. Any discrepancies are displayed.

i A complete summary of all options can be found in the manual page (`man rpm`).

Index

- ++, 142
- , 142
- .alias, 82
- .bash_history, 76
- .bashrc, 82
- .htaccess, 191
- /etc/group, 22
- /etc/passwd, 22
- /etc/shadow, 22
- <, 85–88
- >, 85–88
- #
 - #, *see* shee-bang
- ##, 143, 144
- \$0, 144
- \$1, 144, 147
- \$HOME, 78
- \$PATH, 78
- \$PATH, 61
- \$REPLY, 133
- \$USER, 78
- &&, 89
- ‘, 119, 121
- |, 87, 121
- ||, 89
- 2>, 85–88
- 2>>, 86

- absolute path, 39

- Acrobat Reader, 101
- acroread, 101
- administration, *see* YaST
- alias, 80–82
- alias, 80
- AND, 131
- Apache, 183–196
 - DocumentRoot, 188
 - access, 192–196
 - authentication, 194–196
 - configuration, 188–196
 - encryption, 189
 - error, 189
 - files, 188
 - installation, 184–185
 - start, 186–187
 - test, 186–187
- apache2ctl, 190
- archive, 63–68

- background, 98–101
- backticks, 119, 121
- bash, 75–84
 - alias, 80–82
 - command completion, 75–76
 - function, 146–149
 - history, 76–77
 - name expansion, 82–84
 - process, 98–101
 - programming, 115–149

- regular expression, 89–91, 93–98
 - variable, 78–80, 116–132, 148–149
- bash, 37
- bc, 120–121, 129–131
- break, 143
- bunzip2, 68
- bzip2, 67–68

- case, 134, 137–138
- cat, 56
- cd, 38–40
- channel, 85
- character encoding, 83
- chgrp, 73–75
- chmod, 69–73
- chown, 73–75
- command completion, 75–76
- compare, *see* test
- constants, 116
- conventions, 150
- cp, 47–50
- CUPS, 28
- CVS, 162

- declare, 118
- DHCP, 7
- diff, 161
- directory
 - create, 50–51
 - delete, 51–53
- DNS, 7
- DocumentRoot, *see* Apache
- domain, 9

- echo, 78, 121
- editor, 102–115
- egrep, 90–93, 95–97

- elif, 135–136
- else, 134–136
- encoding, 83
- error output, *see* stderr
- esac, 134
- escape, 80
- exit, 37, 167
- expr, 119

- fi, 134–136
- file, 37–75
 - archive, 63–68
 - copy, 47–50
 - delete, 51–53
 - find, 56–63
 - group, 43
 - link, *see* link
 - move, 45–47
 - owner, 43
 - ownership, 73–75
 - permission, 69–73
 - search, 56–63
- find, 57–59
- float numbers, 120–121
- for, 138, 140–142
- foreground, 98–101
- FQDN, 9
- function, 146–149

- gateway, 10
- gedit, 102
- GID, 14, 20, 22
- global variable, 148
- greedy, 156
- grep, 90–93
- group, 14–22
 - filter, 20
 - ID, *see* GID
 - local, 20

- member, 22
 - password, 22
 - root, 14
 - system, 21
 - user, 22
- group ID, *see* GID
- gunzip, 67
- gzip, 66–67
- hard link, *see* link, hard
- history, 76–77
- HTML, 183
- htpasswd2, 194–195
- HTTP, 183
- httpd, 187
- if, 134–136
- ifup, 5
- index.html, 188
- inode, 53
- input, *see* stdin, 132–133
- installation media, 26–28
- IP address, 7
- IPv4, 7
- job ID, 100–101
- killall, 182
- ksh, 26, 37
- link, 53–56
 - hard, 43, 53–55
 - symbolic, 43, 53, 55–56
- ln, 53–56
- local, 148
- local variable, 148
- locate, 59–61
- loop, 138–143
 - break, 143
- ls, 41–45
- main program, 146
- mkdir, 50–51
- multitasking, 14
- multiuser, 14
- mv, 45–47
- name expansion, 82–84
- name server, 7, 9
- negation, 132
- network, 4–13
 - manager, *see* NetworkManager
 - mask, 7
 - printer, 33
 - tools, 13
- NetworkManager, 5, 12
- newline, 157
- operator, 118
 - logical, 131–132
- option, 143–146
- OR, 131
- output, *see* stdout
- parameter, 143–147
- PATH, 78
- pipe, 85–89, 121
- port, 189
- POSIX, 84
- printer, 28–34
 - manufacturer, 32
 - network, 33
 - queue, 29
- process ID, 100–101
- prompt, 38
- quotes, 79
- rcapach2, 186, 190
- read, 132–133

- redirection, 85–89
- regular expression, 89–91, 93–98, 152, 156
- relative path, 39
- repository, *see* Subversion
- return, 147
- return value, 88–89
- rm, 51–53
- rmdir, 51
- root, 14
- route, 7, 10–11
- router, 7
- rpm, 184–185

- search pattern, 82–84
- sed, 151–159
 - append, 153
 - comment, 159
 - delete, 154
 - holding buffer, 158–159
 - insert, 153–154
 - newline, 157
 - pattern space, 158–159
 - quit, 155
 - regular expression, 152, 156
 - replace, 155–158
 - substitute, 156
 - substitution flag, 157
- SGID, 43, 71, 72
- shee-bang, 116
- shell, 35–159
 - bash, 75–84
 - function, 146–149
 - programming, 115–149
- shift, 144
- software, 23–28
 - dependencies, 23, 26
 - filter, 24
 - icons, 25
 - installation media, 26–28
 - special permission, 71
- SSL, 189
- stderr, 85–88
- stdin, 85–88
- stdout, 85–88
- Sticky, 43
- sticky bit, 71
- stream editor, *see* sed
- su, 60, 166, 172, 184
- subnet mask, 7
- Subversion, 161–183
 - add files, 173
 - check out, 169
 - conflict, 174–177
 - copy files, 174
 - delete files, 173
 - import, 169
 - install, 163–165
 - log, 178–179
 - manage, 178–181
 - merge files, 176
 - move files, 173
 - network, 181–183
 - rename files, 173
 - repository, 165–167
 - server, 181–183
 - set up, 165–167
 - update, 172
- SUID, 43, 71
- svn, 167–181
 - add, 173
 - cat, 180
 - checkout, 169
 - commit, 171
 - copy, 174, 180
 - delete, 173

- diff, 175
- import, 169
- log, 178–179
- merge, 180
- mkdir, 167
- move, 173
- resolved, 177
- revert, 176
- rm, 181
- update, 172, 176
- svnadmin, 166
- svnserve, 181–183
- switch user, 166
- symbolic link, *see* link, symbolic
- tar, 63–66
- TCP/IP, 189
- tcsh, 37
- test, 122–132, 134
 - combination, 131–132
 - file, 122–125
 - float, 129–131
 - integer, 127–129
 - number, 127–131
 - string, 125–127
- type, 62–63, 81
- UID, 14, 18
- unalias, 81
- until, 138–140
- user, 14–22
 - disable, 18
 - filter, 16
 - group, 19
 - ID, *see* UID
 - local, 16
 - name, 14, 17
 - password, 14, 17, 19–20
 - root, 14
 - shell, 19
 - switch, 166
 - system, 16
 - user ID, *see* UID
- UTF-8, 84
- variable, 61, 78–80
 - declare, 118
- version control system, *see* Subversion
- vi, 102–115
 - abbreviation, 114–115
 - auto completion, 111
 - command mode, 105
 - command-line mode, 106, 108
 - ex mode, 107
 - frames, 110
 - macro, 112–113
 - marks, 111
 - modes, 105–109
 - normal mode, 105, 107–108
 - replace mode, 106
 - screen, 110
 - select mode, 107
 - split, 110
 - syntax highlighting, 111
 - tutorial, 112
 - visual mode, 107
- wc, 88
- web browser, 183–184
- web server, *see* Apache
- WebDAV, 181
- which, 61–62
- while, 138–139
- wwwrun, 187, 189
- YaST, 3–35
 - group, 15, 20–22

network, 4–13
printer, 28–34
software, 23–28
user, 15–20

